CIS 3250
Spanning Tree Protocols

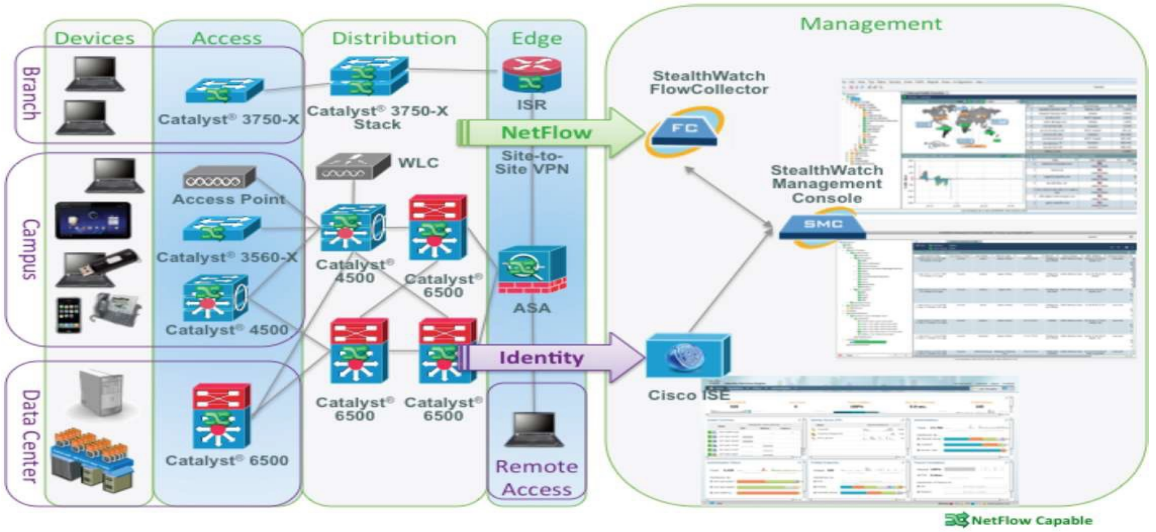

The Role of Redundancy in a Converged Switched Network

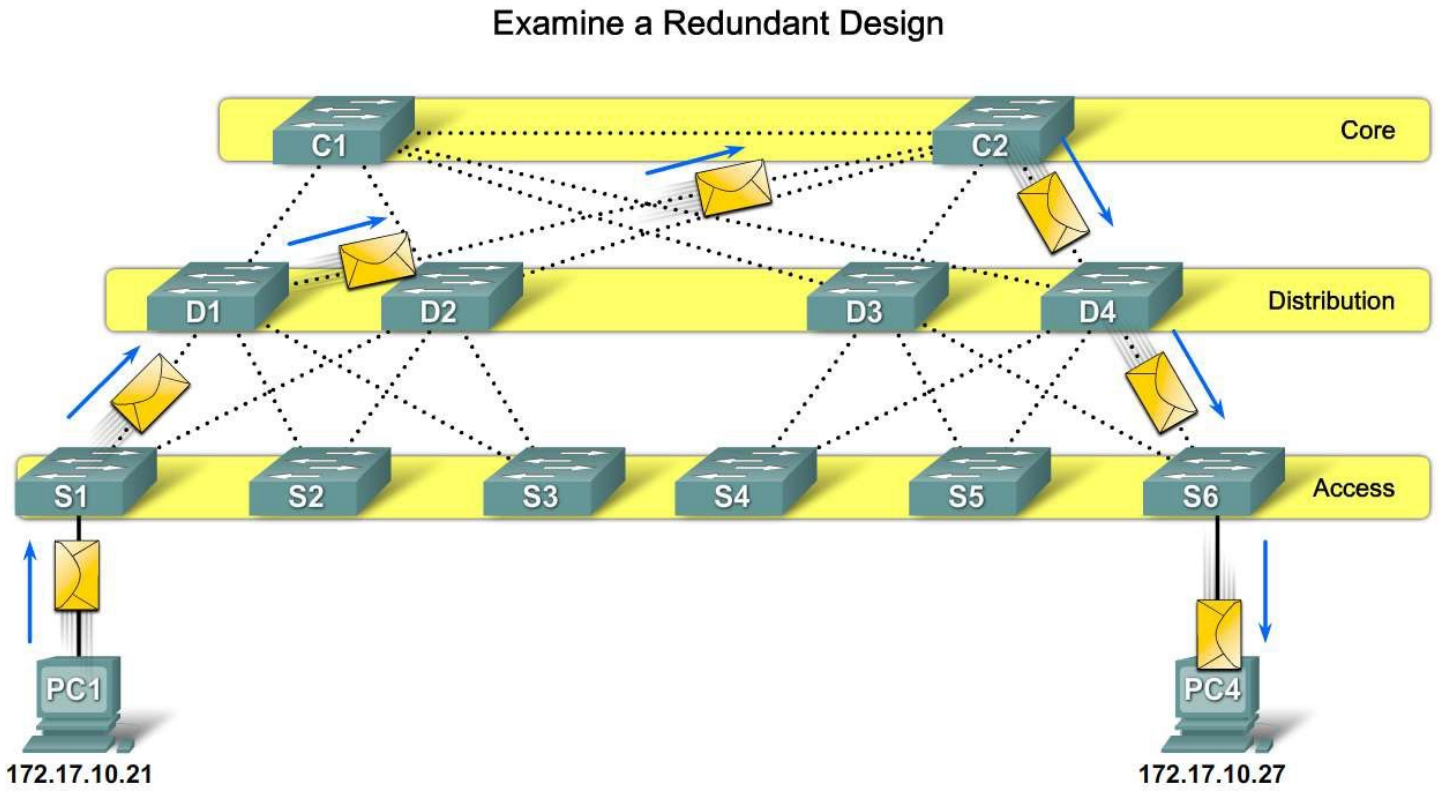

The Role of Redundancy in a Converged Switched Network

redundancy can disable a hierarchical network

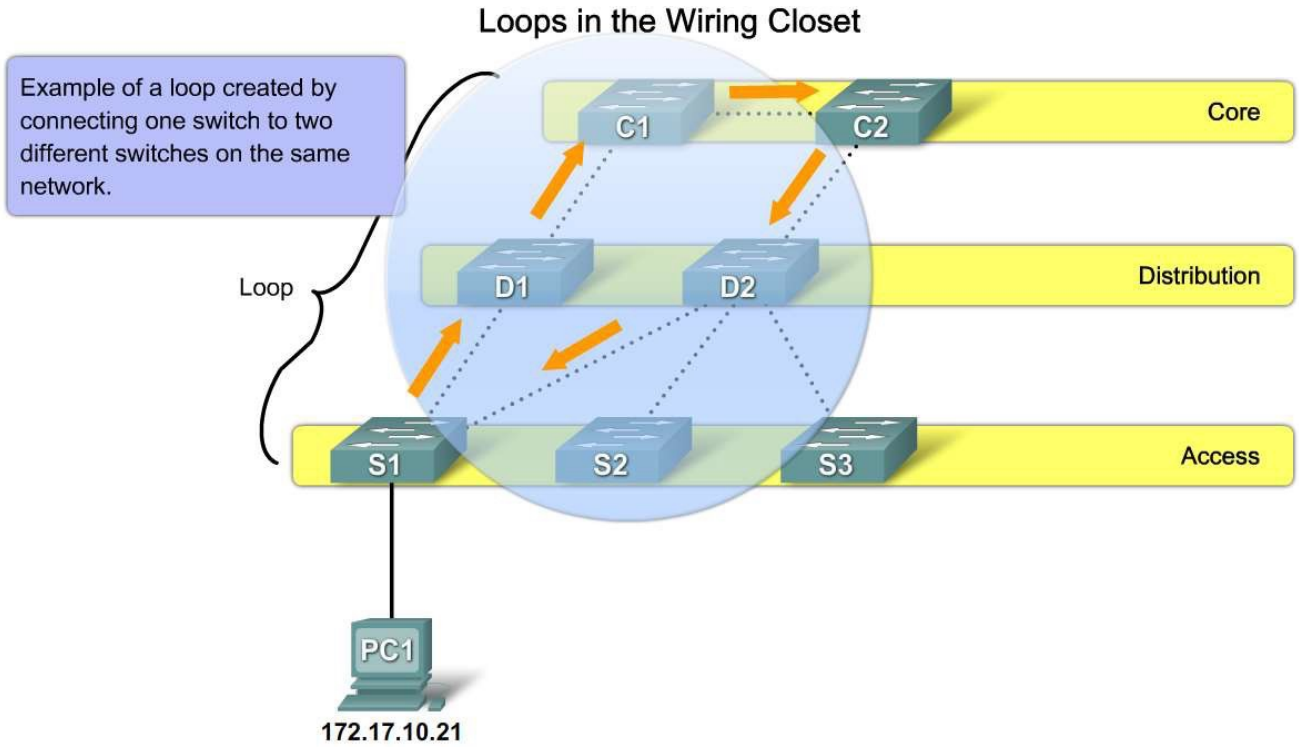

Loops in the Wiring Closet

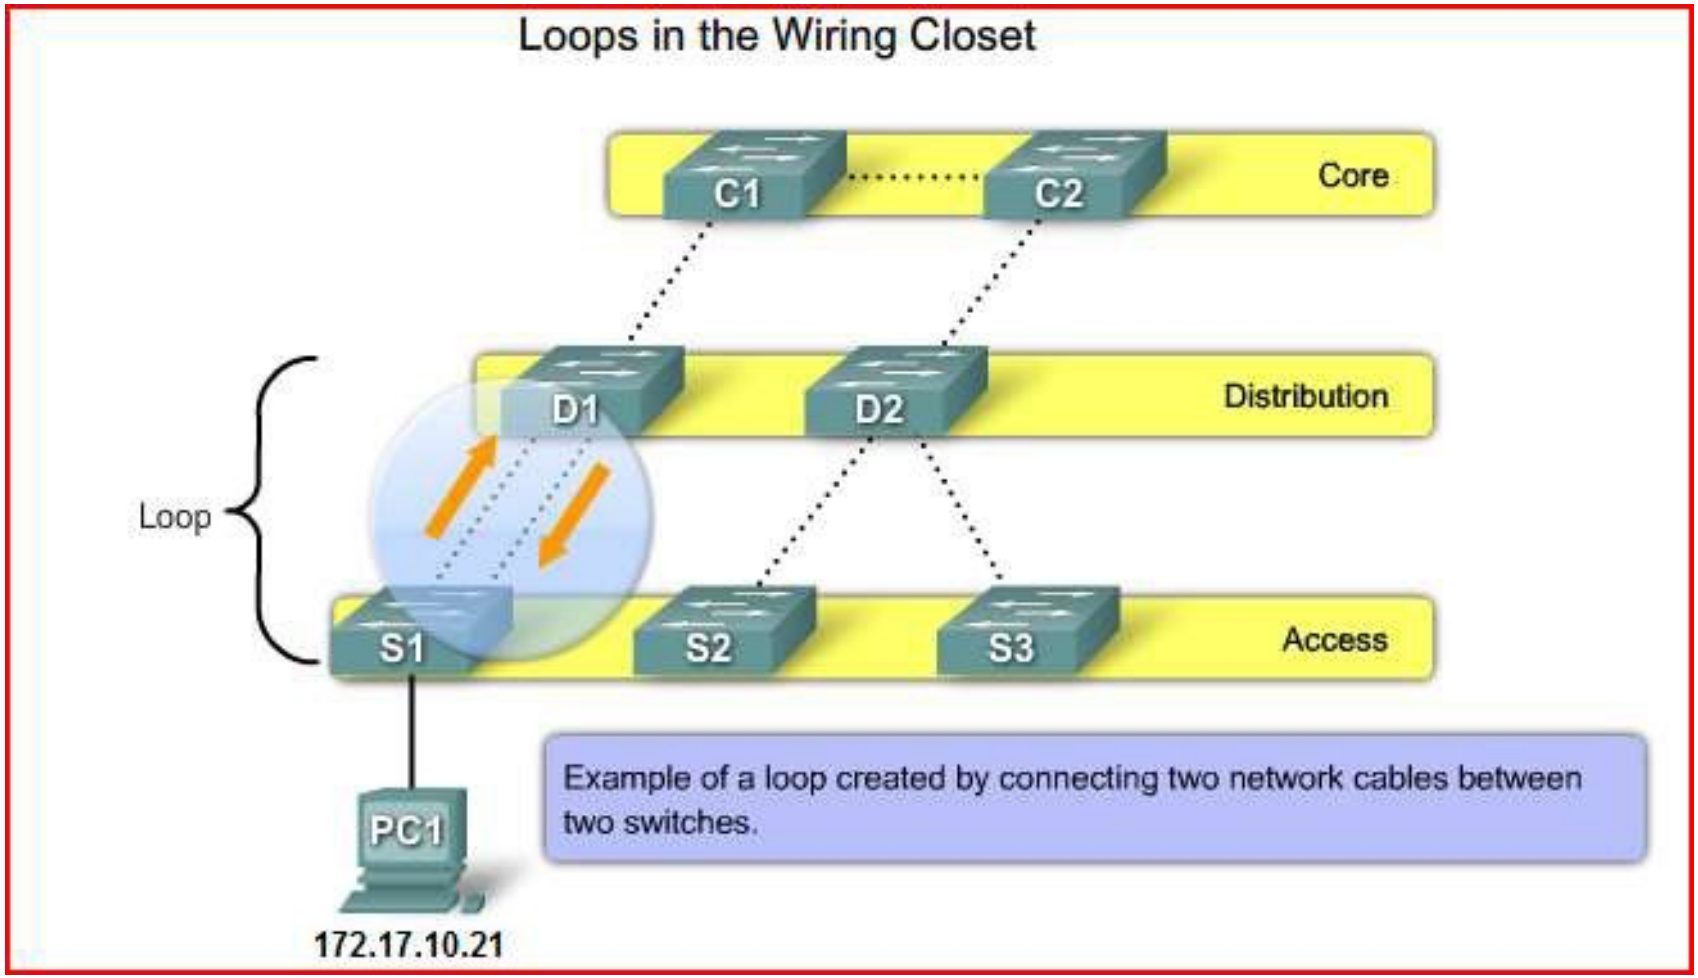

Loops in the Cubicles

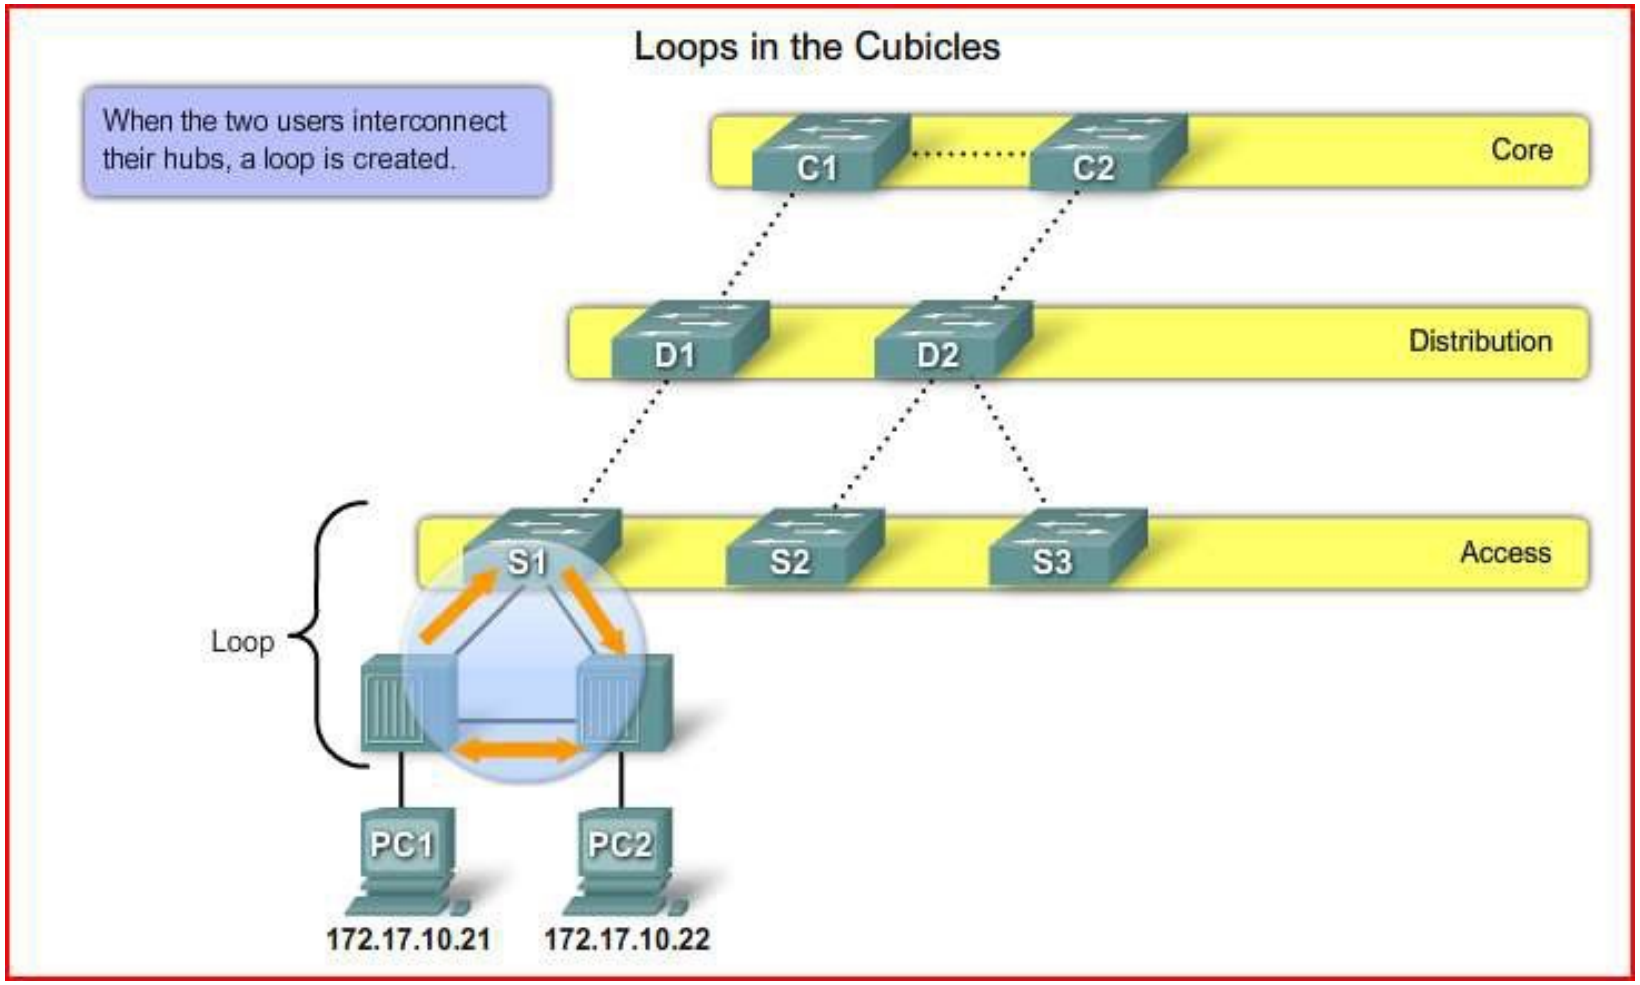

Redundancy in a Converged Switched Network

Layer 2 loops can occur in well-managed networks

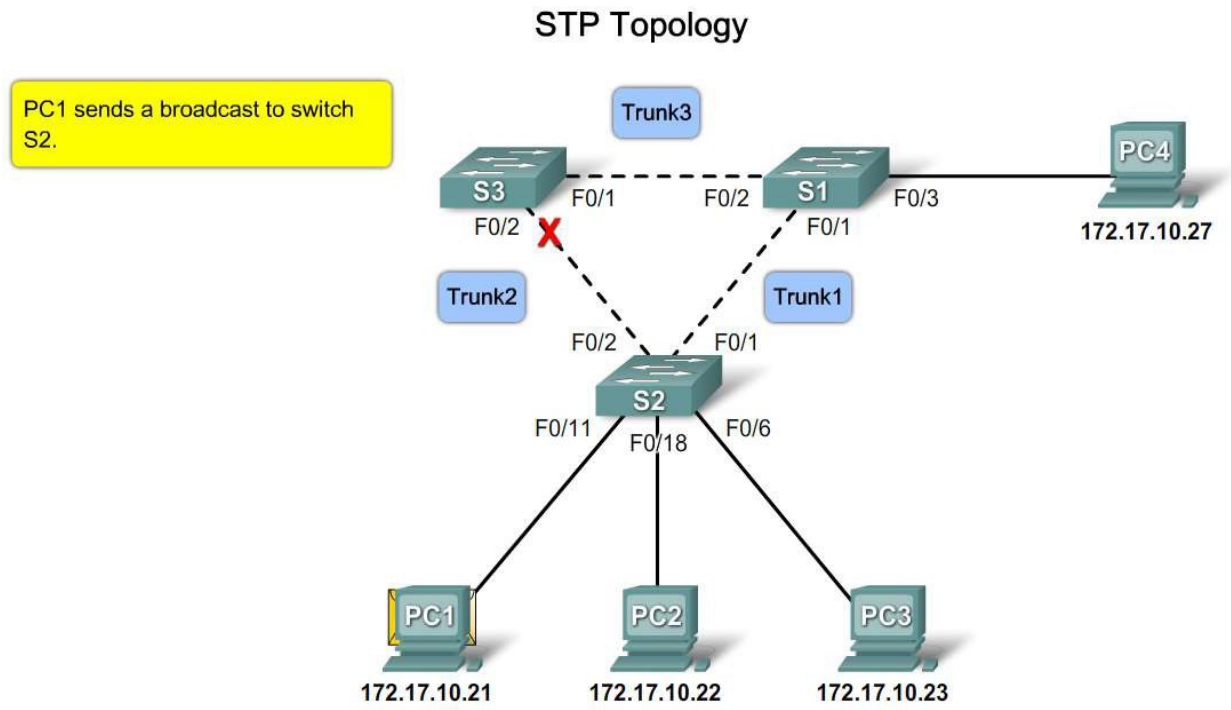

Spanning Tree Protocol - STP

- STP ensures that there is only one logical path between all destinations on the network by intentionally blocking redundant paths that could cause a loop.
- A port is considered blocked when network traffic is prevented from entering or leaving that port. This does not include bridge protocol data unit (BPDU) frames that are used by STP to prevent loops.

Spanning Tree Protocol (STP)

- “STP often accounts for more than 50% of the configuration, troubleshooting, and maintenance headaches in real-world campus networks (especially if they are poorly designed).”
- “Complex protocol that is generally poorly understood.”
- Radia Perlman – Developer of STP

STP works to Eliminate Layer 2 Loops in a Converged Network

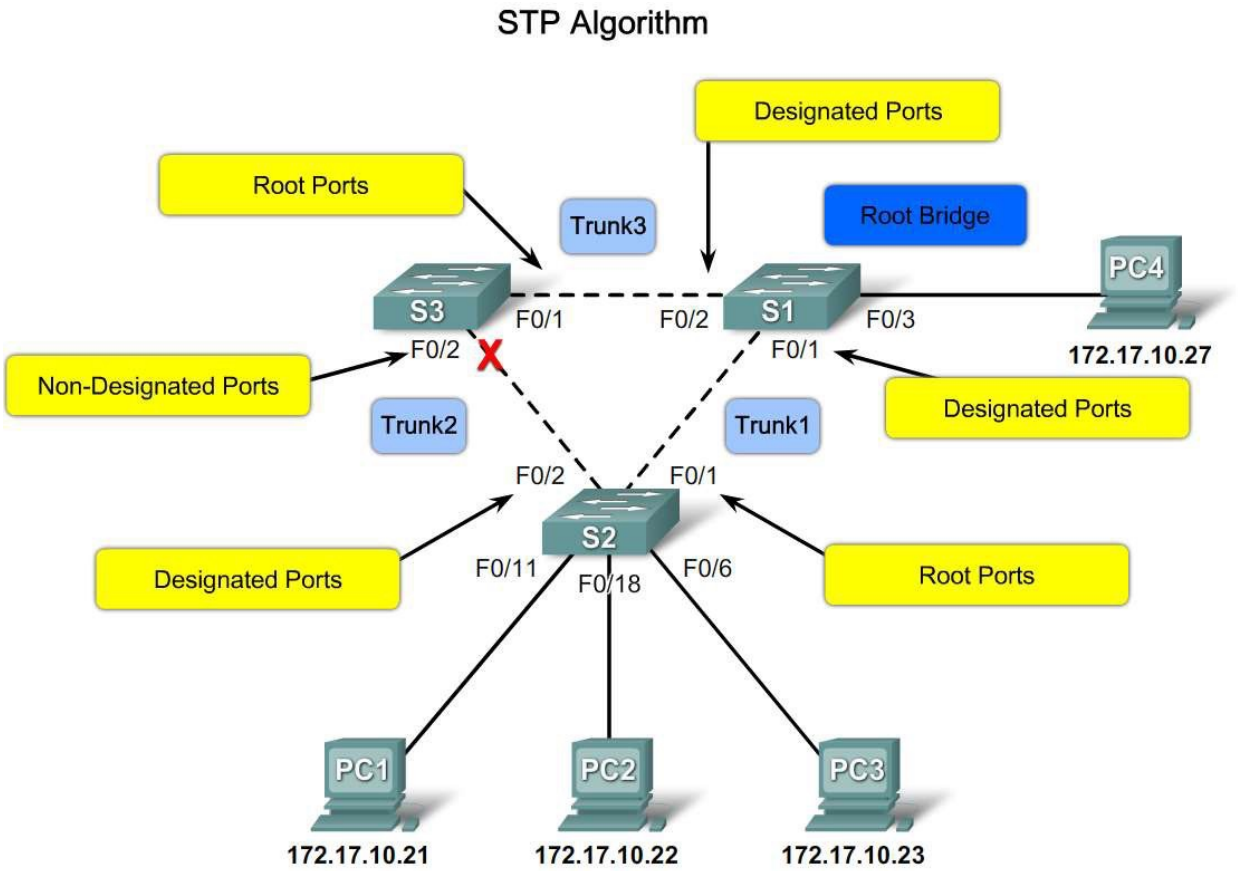

Spanning Tree Algorithm - STA

- The STA designates a single switch as the root bridge and uses it as the reference point for all path calculations.
- All switches participating in STP exchange Bridge Protocol Data Units (BPDU) frames to determine which switch has the lowest bridge ID (BID) on the network.
- The switch with the lowest BID automatically becomes the root bridge for the STA calculations. The root bridge election process will be discussed in detail later.

Layer 2 Loops

- Broadcasts and Layer 2 loops can be a dangerous combination.
- Ethernet frames have no TTL field.
- After an Ethernet frame starts to loop, it will probably continue until someone shuts off one of the switches or breaks a link.
- IP has a mechanism to prevent loops.

Ethernet Frame Format

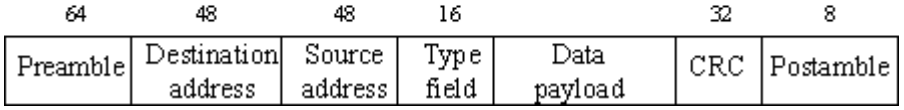

IP Packet

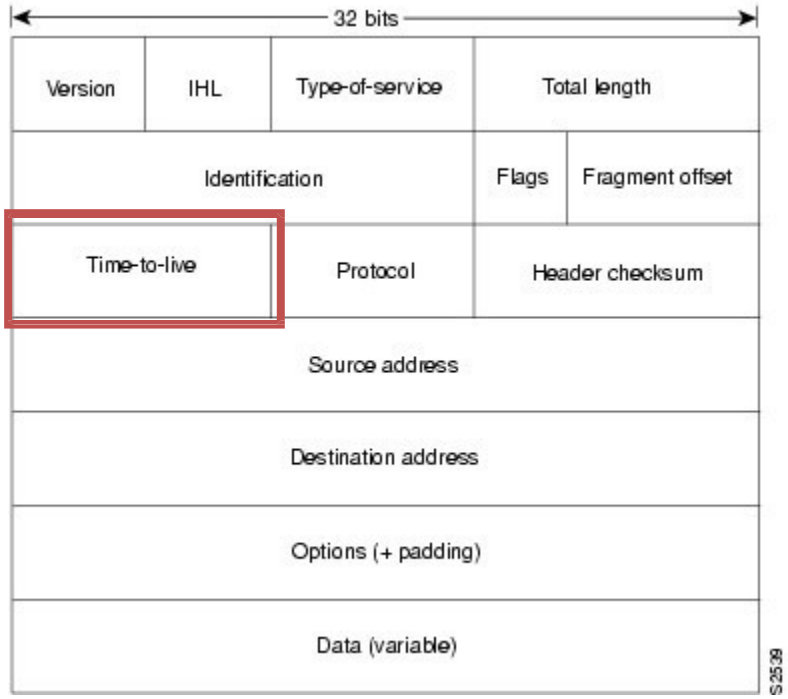

Layer 2 Loops

- Bridge loops can occur any time there is a redundant path or loop in the bridge network.
- The switches will flip-flop the MAC address table entries (creating extremely high CPU utilization).
- Unicasts, unknown unicasts, and broadcasts are all problems.

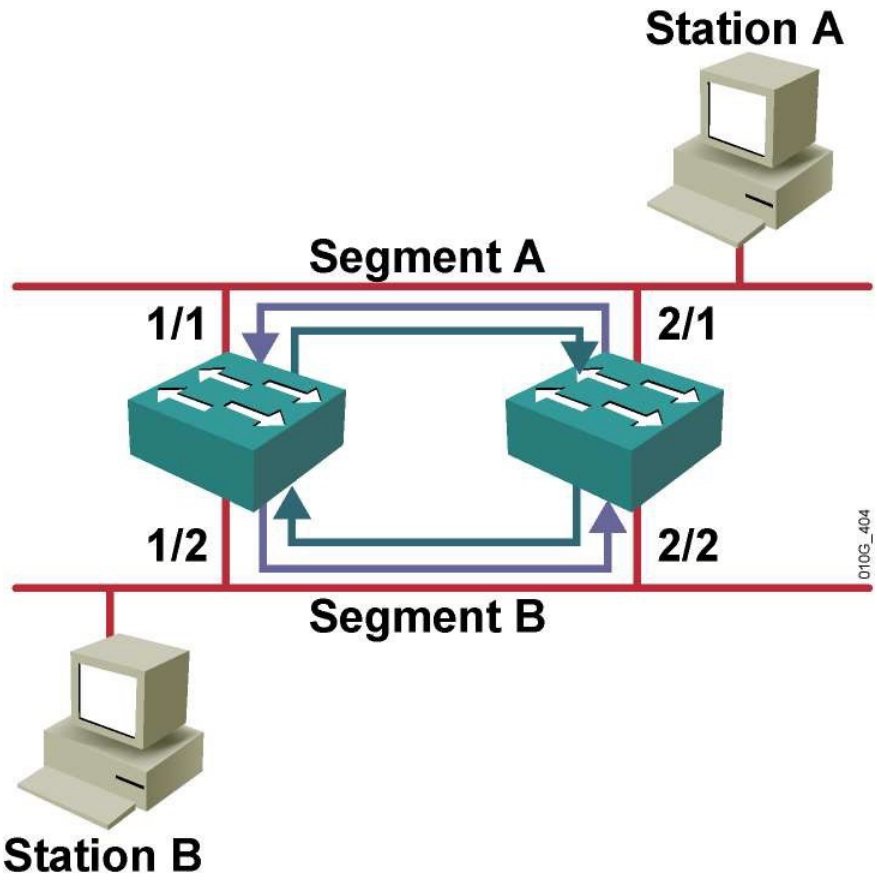

Spanning Tree Algorithm

- STP executes an algorithm called the Spanning Tree Algorithm (STA).
  - STA chooses a reference point called a **root bridge.**
  - Then determines the available paths to that reference point.
  - If more than two paths exist, STA picks the best path and blocks the rest.

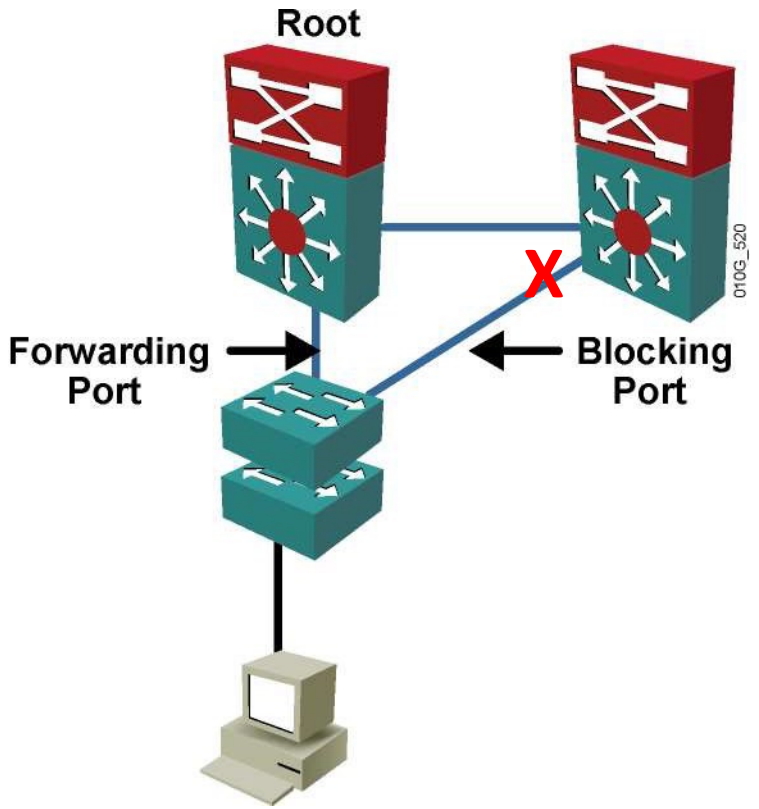

Bridge Protocol Data Units - BPDU

- The BPDU is the message frame exchanged by switches for STP.
- Each BPDU contains a BID that identifies the switch that sent the BPDU.
- The BID contains a priority value, the sending switch's MAC address, and an optional extended system ID. The combination of these three fields determines the lowest BID value.
- You will learn more about the root bridge, BPDU, and BID later.

Bridge ID (BID)

- Bridge ID (BID) is used to identify each bridge/switch.
- The BID is used in determining the center of the network, in respect to STP, known as the root bridge.

Bridge ID Without the Extended System ID

Bridge ID with the Extended System ID

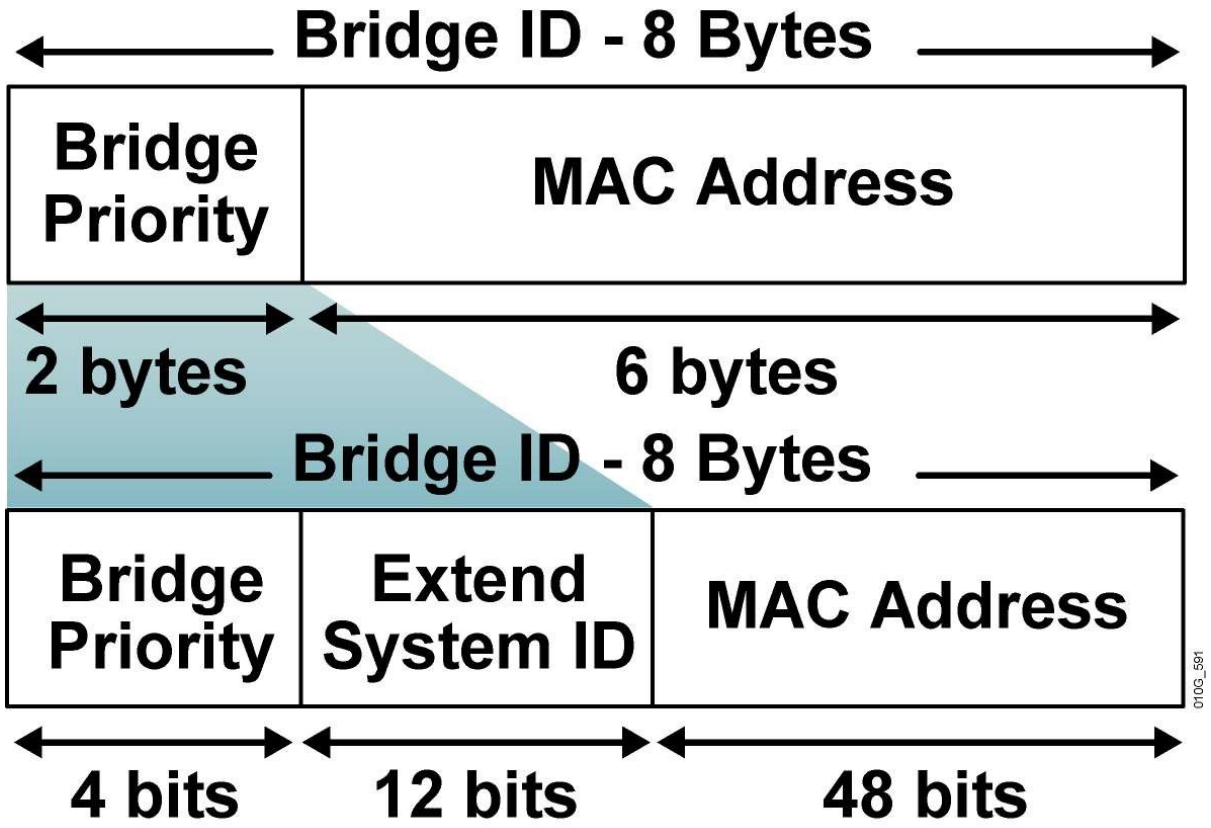

Bridge ID (BID)

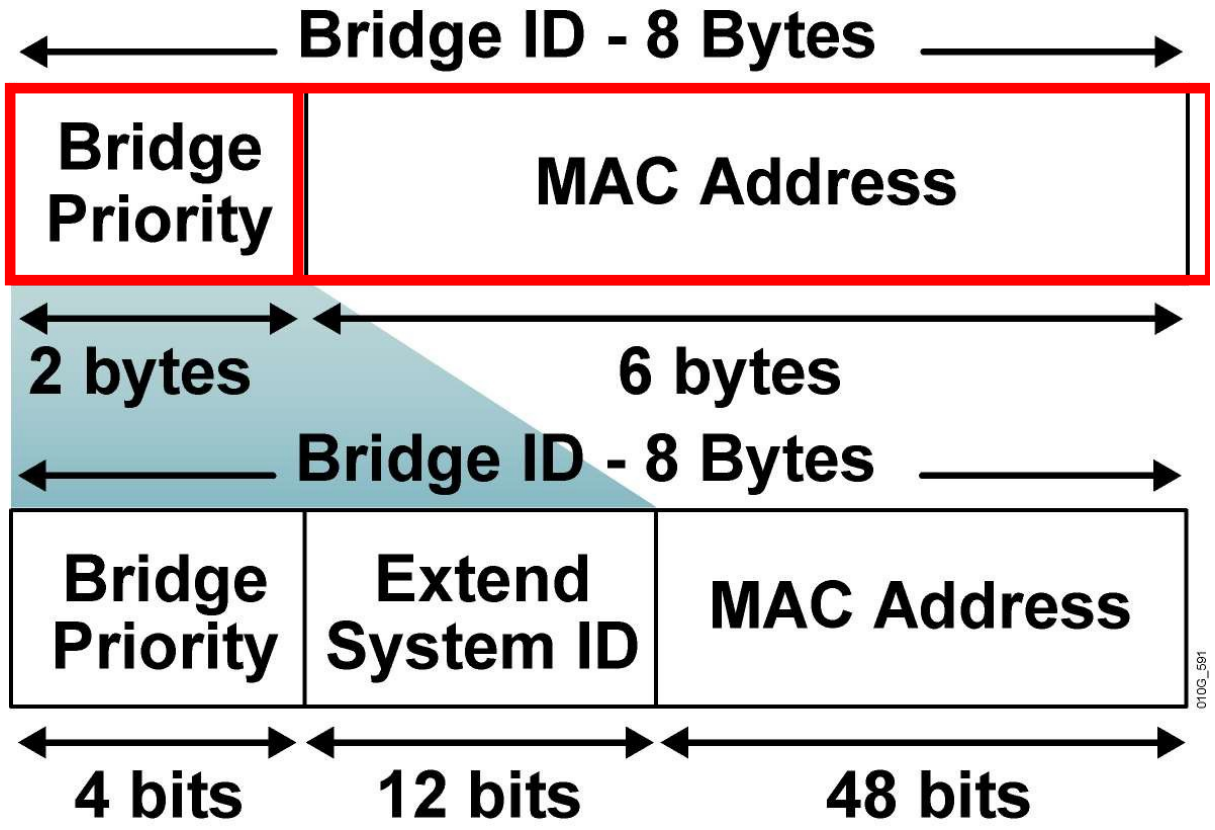

- Consists of two components:
  - **A 2-byte Bridge Priority**: Cisco switch defaults to 32,768 or 0x8000.
    - Usually expressed in **decimal format**
  - **A 6-byte MAC address**
    - Usually expressed in **hexadecimal format**.

Bridge ID (BID)

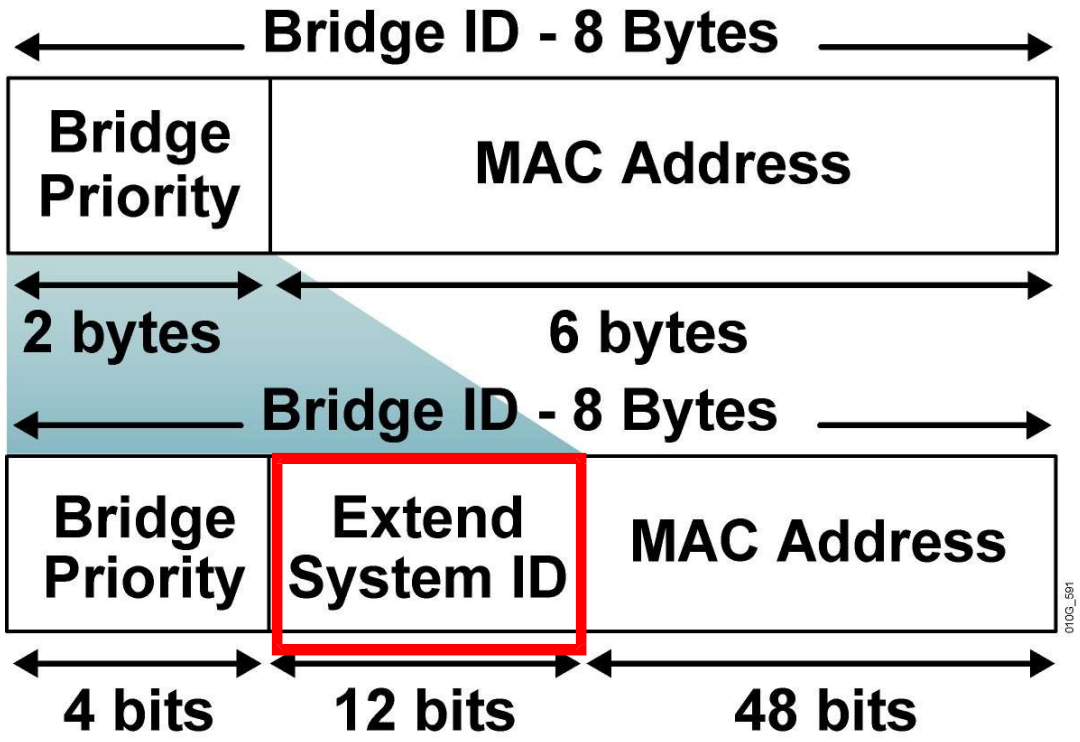

- Each switch has a **unique BID**.
- **Original 802.1D** standard, the BID = **Priority** Field +**MAC address** of the switch.
  - All VLANs were represented by a Common Spanning Tree (CST) – one spanning tree for all VLANs (later).
- **Per-VLAN Spanning Tree (PVST)** requires that a separate instance of the spanning tree protocol be run for each VLAN
  - **BID field is required to carry VLAN ID (VID)**.
  - **Extended system ID** to carry a VID.

Bridge ID (BID)

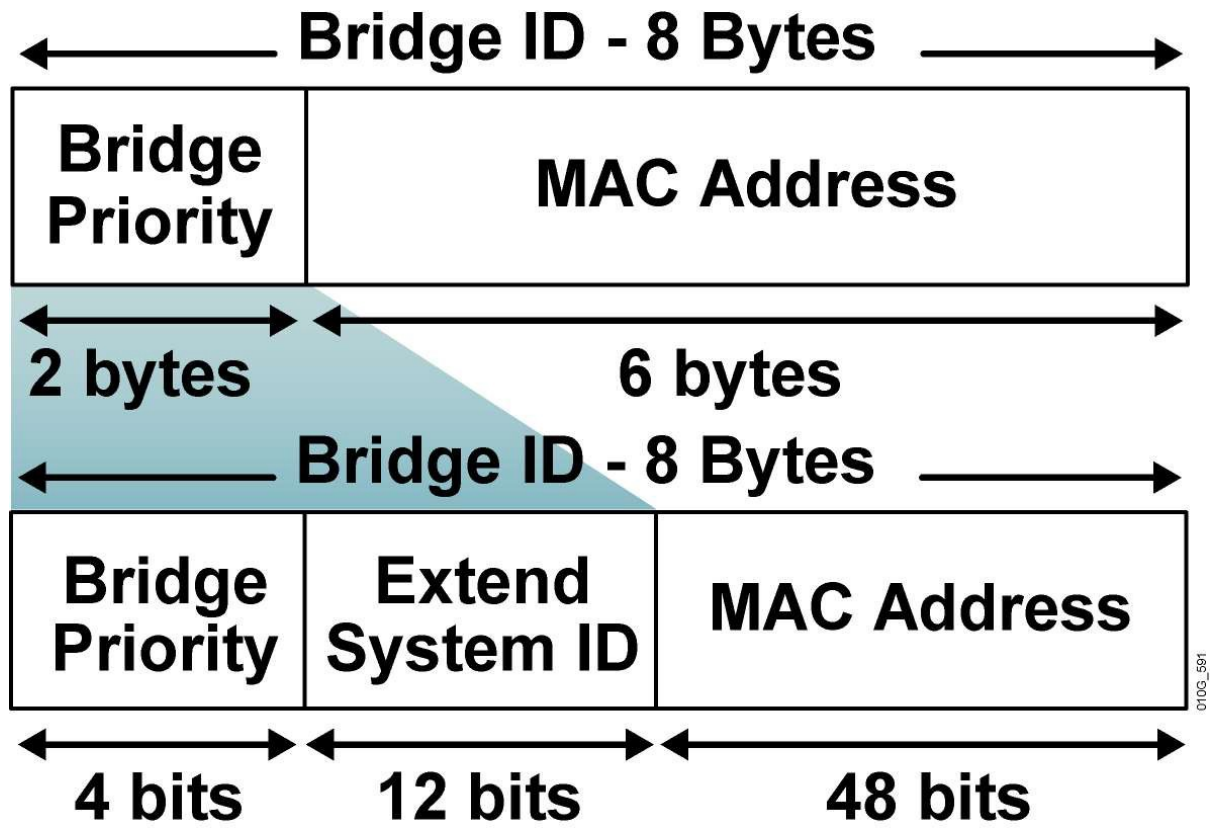

- Used to elect a root bridge (coming)
- **Lowest** Bridge ID is the root.
- If all devices have the same priority, the bridge with the lowest MAC address becomes the root bridge. (Yikes)
- **Note**: For simplicity, in our topologies we will use Bridge Priorities without the Extended System ID. (Same process, just done per VLAN.)

Path Cost – Original Spec (Linear)

Link Speed	Cost (Revised IEEE Spec)	Cost (Previous IEEE Spec)
10 Gbps	2	1
1 Gbps	4	1
100 Mbps	19	10
10 Mbps	100	100

- Bridges use the concept of **cost** to evaluate how close they are to other bridges.
- Used to create the loop-free topology .
- Originally, 802.1D defined cost as **1 billion/bandwidth** of the link in Mbps.
  - Cost of 10 Mbps link = 100
  - Cost of 100 Mbps link = 10
  - Cost of 1 Gbps link = 1
- Running out of room for faster switches including 10 Gbps Ethernet

Path Cost – Revised Spec (Non-Linear)

Link Speed	Cost (Revised IEEE Spec)	Cost (Previous IEEE Spec)
10 Gbps	2	1
1 Gbps	4	1
100 Mbps	19	10
10 Mbps	100	100

- IEEE modified the most to use a ***non-linear scale*** with the new values of:
  - 4 Mbps 250 (cost)
  - 10 Mbps 100 (cost)
  - 16 Mbps 62 (cost)
  - 45 Mbps 39 (cost)
  - 100 Mbps 19 (cost)
  - 155 Mbps 14 (cost)
  - 622 Mbps 6 (cost)
  - 1 Gbps 4 (cost)
  - 10 Gbps 2 (cost)

- You can change the path cost by modifying the cost of a port.
- Exercise caution when you do this!
- BID and Path Cost are used to develop a loop-free topology .
- Coming very soon!

Five-Step STP Decision Sequence

- When creating a loop-free topology, STP always uses the same five-step decision sequence:

Five-Step decision Sequence

Step 1 - Lowest BID

Step 2 - Lowest Path Cost to Root Bridge

Step 3 - Lowest Sender BID

Step 4 – Lowest Port Priority

Step 5 - Lowest Port ID

- Bridges use Configuration BPDUs during this five-step process.
- We will assume all BPDUs are configuration BPDUs until otherwise noted.

FYI: BPDU key concepts

BPDU key concepts:

- Bridges save a copy of only the best BPDU seen on every port.
- When making this evaluation, it considers all of the BPDUs received on the port, as well as the BPDU that would be sent on that port.
- As every BPDU arrives, it is checked against this five-step sequence to see if it is more attractive (lower in value) than the existing BPDU saved for that port.
- Only the lowest value BPDU is saved.
- Bridges send configuration BPDUs until a more attractive BPDU is received.
- Okay, lets see how this is used...

Elect one Root Bridge

The STP algorithm uses three simple steps to ***converge*** on a loop-free topology:

STP Convergence

Step 1 Elect one Root Bridge
Step 2 Elect Root Ports
Step 3 Elect Designated Ports

- When the network first starts, all bridges are announcing a chaotic mix of BPDUs.
- All bridges immediately begin applying the five-step sequence decision process.
- Switches need to elect a single Root Bridge.
- Switch with the **lowest BID** wins!
- Note: Many texts refer to the term “highest priority” which is the “lowest” BID value.
- This is known as the **“Root War.”**

Spanning Tree Algorithm - Shortest Path

- The STA calculates the shortest path to the root bridge.
- Each switch uses the STA to determine which ports to block.
- While the STA determines the best paths to the root bridge for all destinations in the broadcast domain, no data traffic is forwarded.
- The STA considers both path and port costs.
- The path costs are calculated using port cost values associated with port speeds for each switch port along a given path.
- The sum of the port cost values determines the overall path cost to the root bridge. If there is more than one path to choose from, STA chooses the path with the lowest path cost.
- More about path and port costs in later topics.

Summarize How STP works to Eliminate Layer 2 Loops in a Converged Network

How do port roles support the operation of STP ?

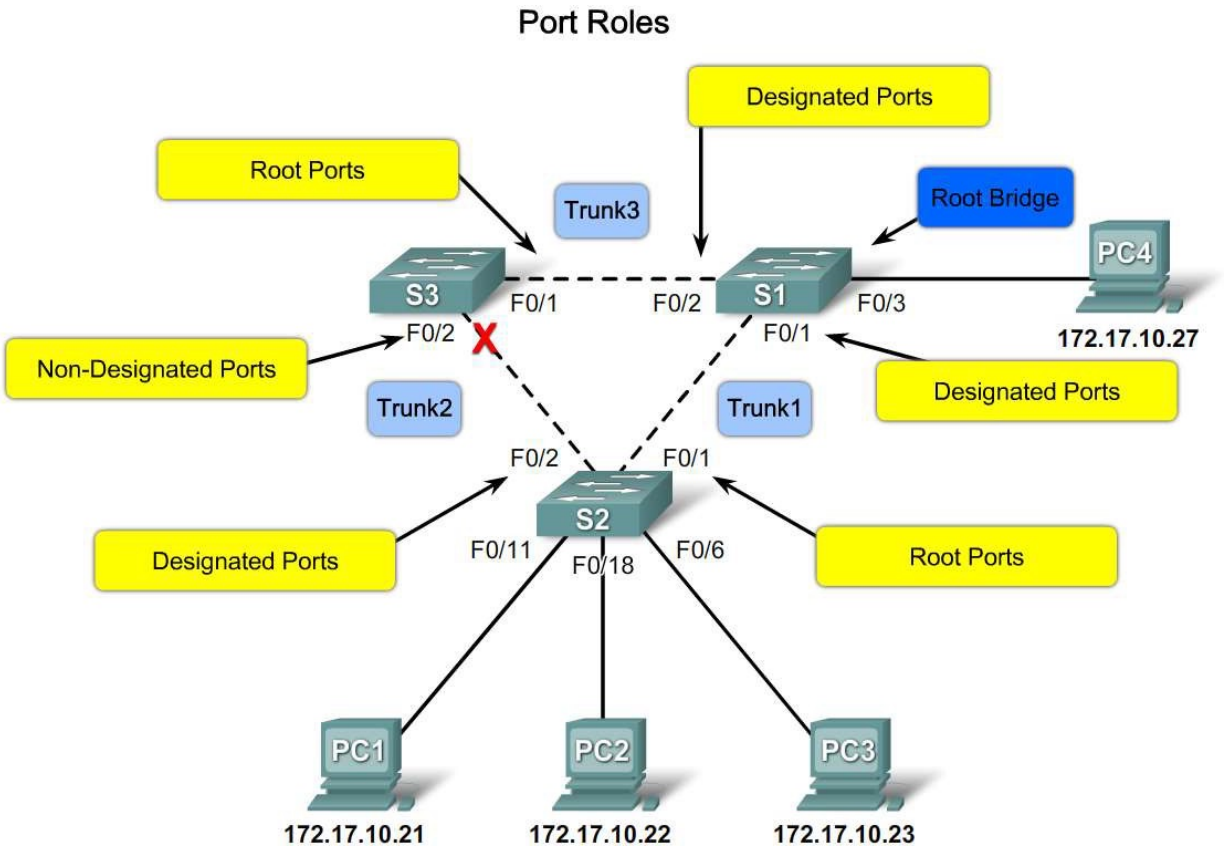

Port Roles

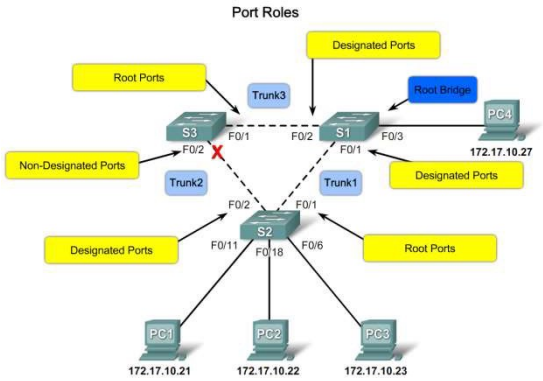

- **Root ports** - Switch ports closest to the root bridge. In the example, the root port on switch S2 is F0/1 configured for the trunk link between switch S2 and switch S1. The root port on switch S3 is F0/1, configured for the trunk link between switch S3 and switch S1.
- **Designated ports** - All non-root ports that are still permitted to forward traffic on the network. In the example, switch ports F0/1 and F0/2 on switch S1 are designated ports. Switch S2 also has its port F0/2 configured as a designated port.
- **Non-designated ports** - All ports configured to be in a blocking state to prevent loops.

Port Roles

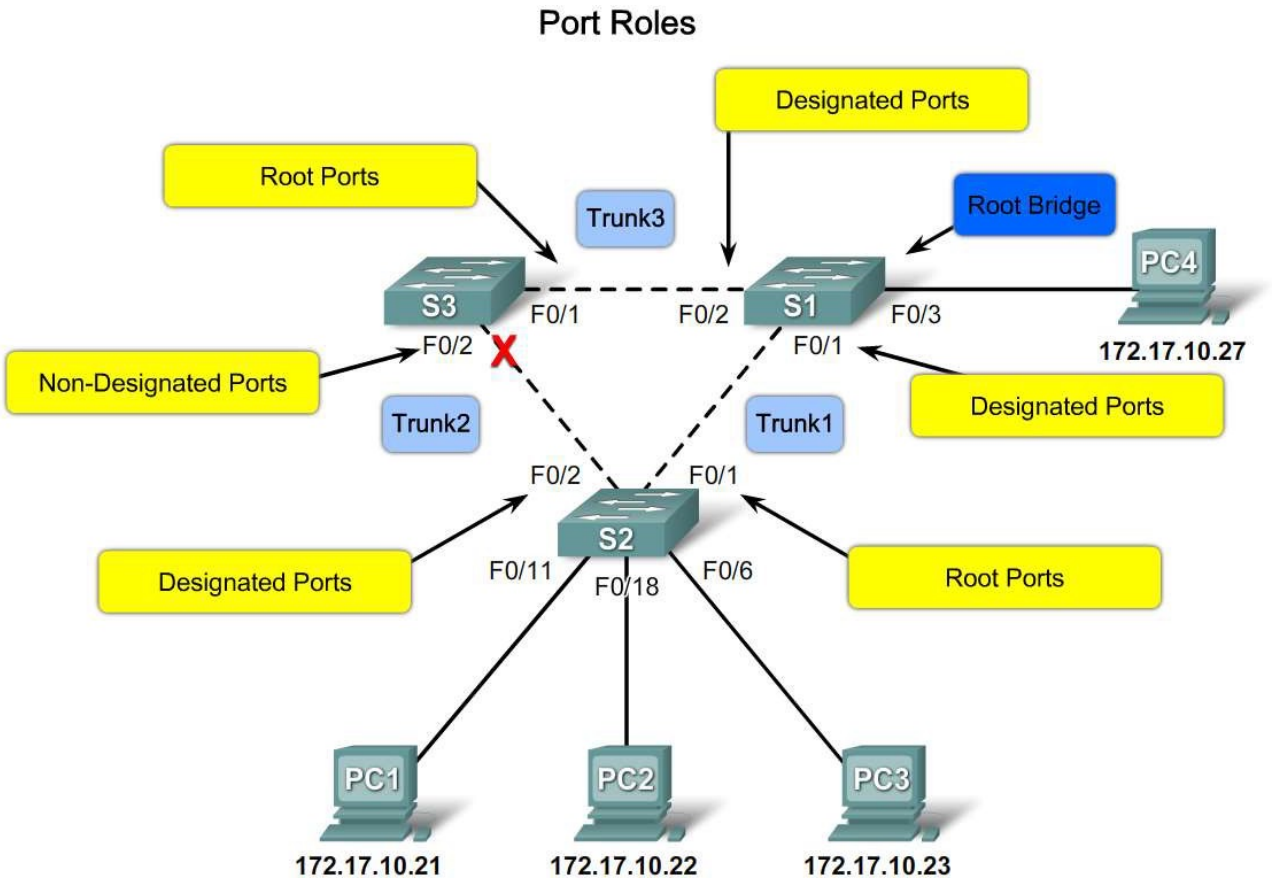
Port Roles
Designated Ports
Root Ports
Trunk3
Root Bridge
PC4
S3
F0/1
F0/2
S1
F0/3
F0/2
F0/1
172.17.10.27
Non-Designated Ports
Designated Ports
Trunk2
Trunk1
F0/2
F0/1
S2
Designated Ports
F0/11
F0/18
F0/6
Root Ports
PC1
PC2
PC3
172.17.10.21
172.17.10.22
172.17.10.23

How the STP Algorithm Uses Three Steps to Converge on a Loop-Free Topology

- What is convergence for a switched network?
- The 3-step process STP uses to create a loop free topology

STP Convergence Steps:

Step 1: elect a root bridge
Step 2: elect the root ports
Step 3: elect the designated and non-designated ports

STP Convergence

- Convergence is an important aspect of the spanning-tree process.
- Convergence is the time it takes for the network to determine which switch is going to assume the role of the root bridge, go through all the different port states, and set all switch ports to their final spanning-tree port roles where all potential loops are eliminated.
- The convergence process takes time to complete because of the different timers used to coordinate the process.

Root Bridge

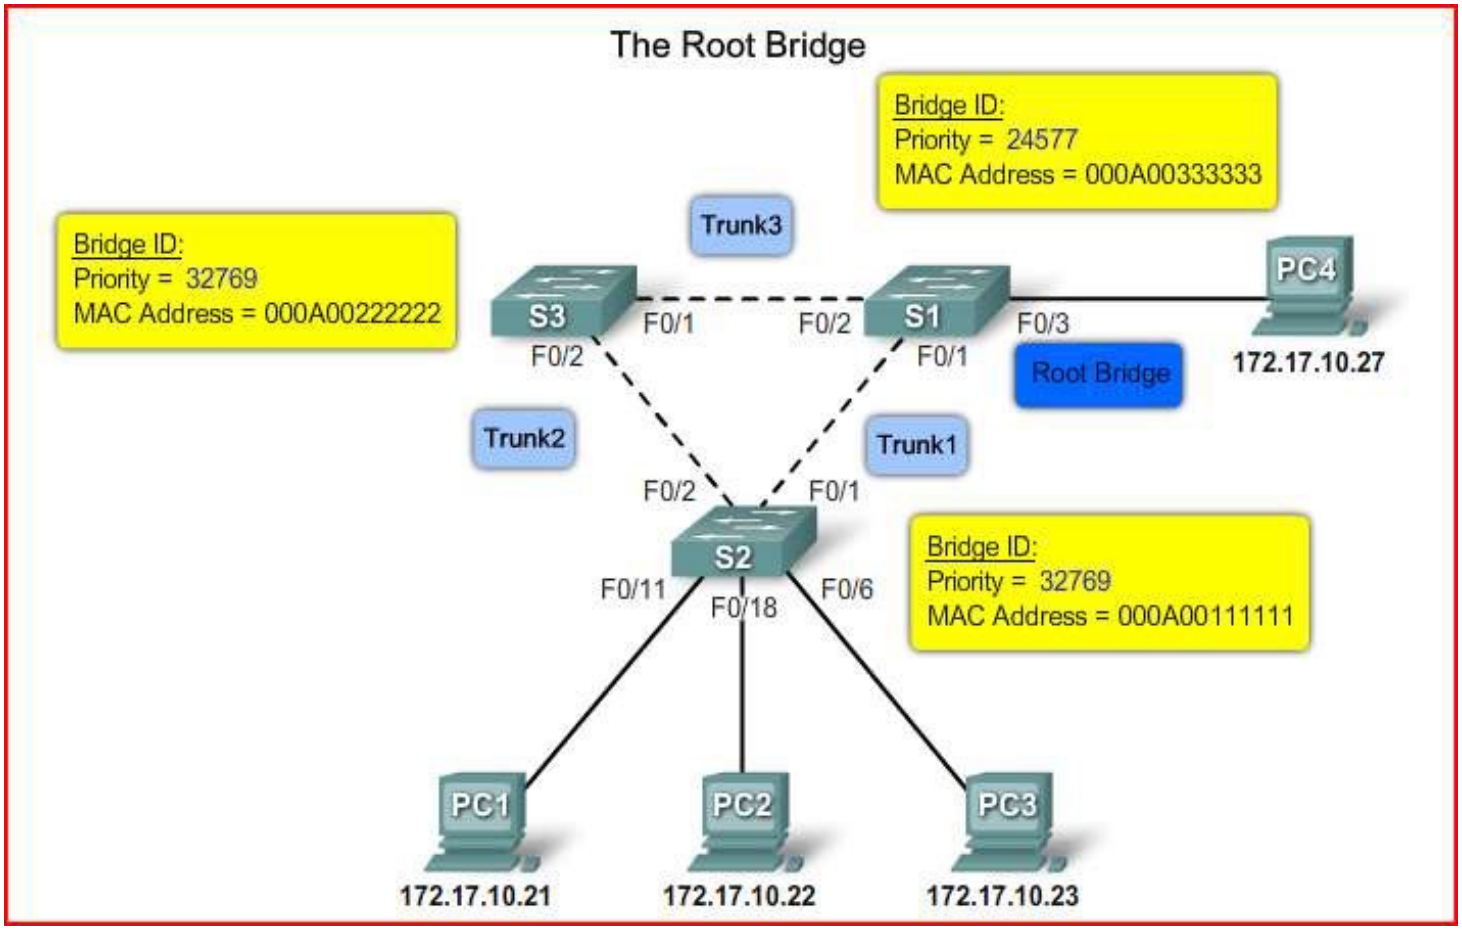
The Root Bridge
Bridge ID:
Priority = 24577
MAC Address = 000A00333333
Bridge ID:
Priority = 32769
MAC Address = 000A00222222
Trunk3
S3
F0/1
F0/2
S1
F0/3
PC4
F0/2
F0/1
Root Bridge
172.17.10.27
Trunk2
Trunk1
F0/2
F0/1
S2
Bridge ID:
Priority = 32769
MAC Address = 000A00111111
F0/11
F0/18
F0/6
PC1
PC2
PC3
172.17.10.21
172.17.10.22
172.17.10.23

Root Bridge Election

- The BID is made up of a priority value, an extended system ID, and the MAC address of the switch.
- All switches in the broadcast domain participate in the election process.
- After a switch boots, it sends out BPDU frames containing the switch BID and the root ID every 2 seconds.
- By default, the root ID matches the local BID for all switches on the network. The root ID identifies the root bridge on the network.
- Initially, each switch identifies itself as the root bridge after bootup.

Root Bridge Election 2

- As the switches forward their BPDU frames, adjacent switches in the broadcast domain read the root ID information from the BPDU frame.
- If the root ID from the BPDU received is lower than the root ID on the receiving switch, the receiving switch updates its root ID identifying the adjacent switch as the root bridge.
- Note: It may not be an adjacent switch, but any other switch in the broadcast domain.
- The switch then forwards new BPDU frames with the lower root ID to the other adjacent switches.
- Eventually, the switch with the lowest BID ends up being identified as the root bridge for the spanning-tree instance.

BID Fields

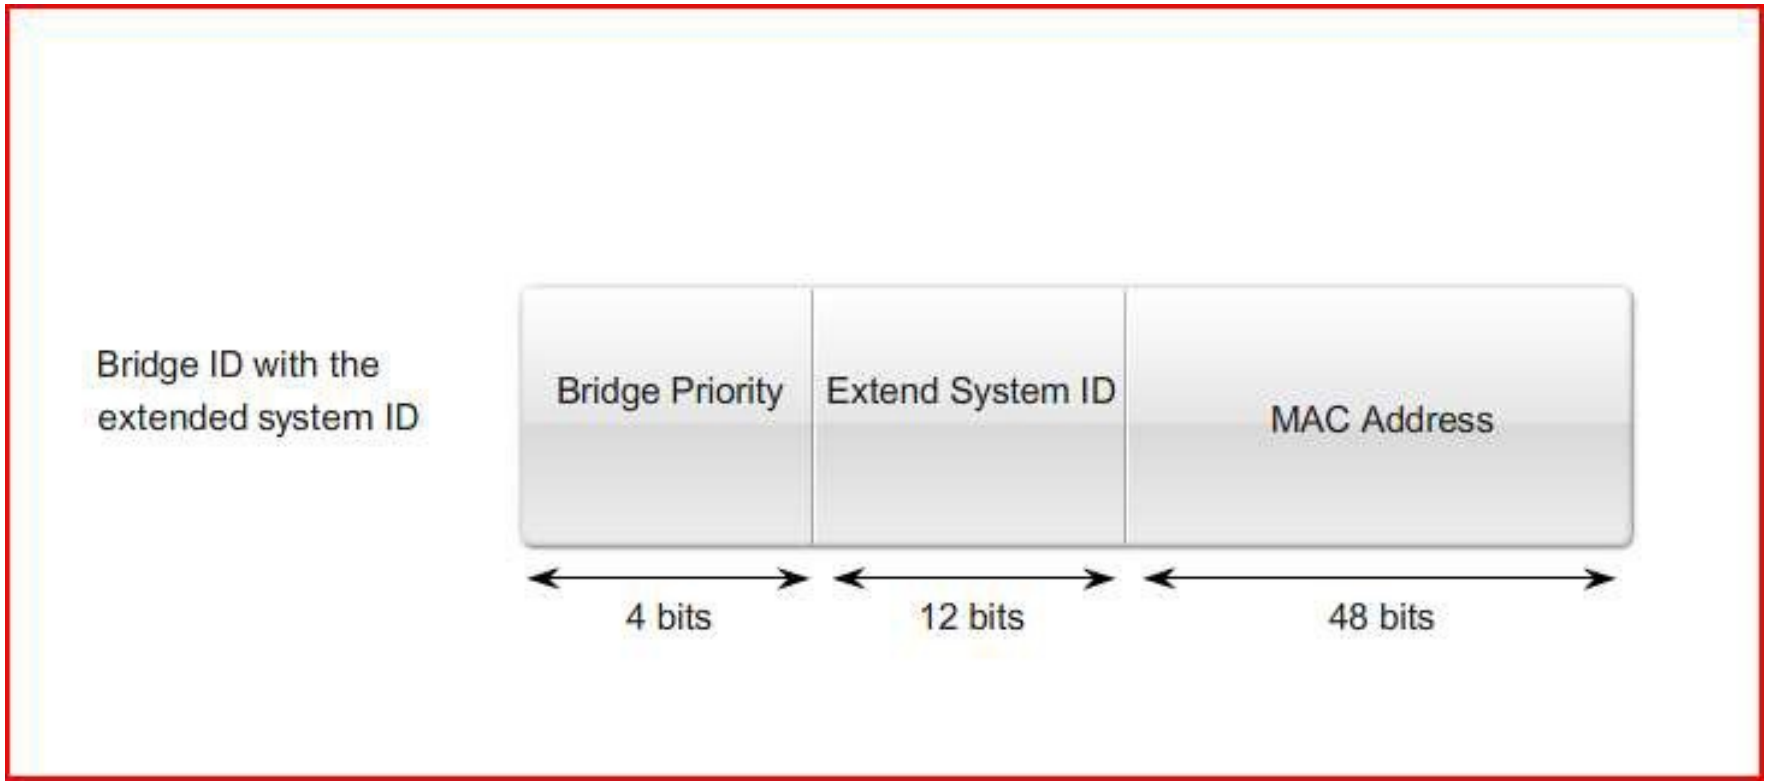
Bridge ID with the extended system ID
Bridge Priority
Extend System ID
MAC Address
4 bits
12 bits
48 bits

Best Path to Root

- When the root bridge has been designated for the spanning-tree instance, the STA starts the process of determining the best paths to the root bridge from all destinations in the broadcast domain.
- The path information is determined by summing up the individual port costs along the path from the destination to the root bridge.
- The default port costs are defined by the speed at which the port operates.
- In the table (next slide), you can see that 10-Gb/s Ethernet ports have a port cost of 2, 1-Gb/s Ethernet ports have a port cost of 4, 100-Mb/s Fast Ethernet ports have a port cost of 19, and 10-Mb/s Ethernet ports have a port cost of 100.

Two-key STP Concepts

- STP calculations make extensive use of two key concepts in creating a loop-free topology:
  - Bridge ID
  - Path Cost

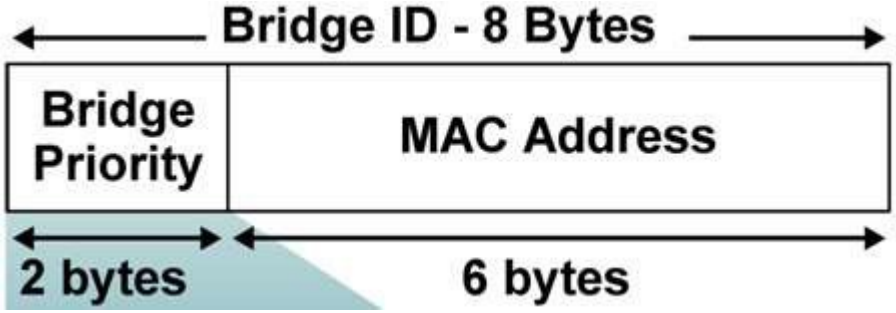

Link Speed	Cost (Revised IEEE Spec)	Cost (Previous IEEE Spec)
10 Gbps	2	1
1 Gbps	4	1
100 Mbps	19	10
10 Mbps	100	100

Configuring Port Costs

Configure Port Cost

```
S2#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
S2(config)#interface f0/1
S2(config-if)#spanning-tree cost 25
S2(config-if)#end
S2#
```

Reset Port Cost

```
S2#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
S2(config)#interface f0/1
S2(config-if)#no spanning-tree cost
S2(config-if)#end
S2#
```

Path Cost

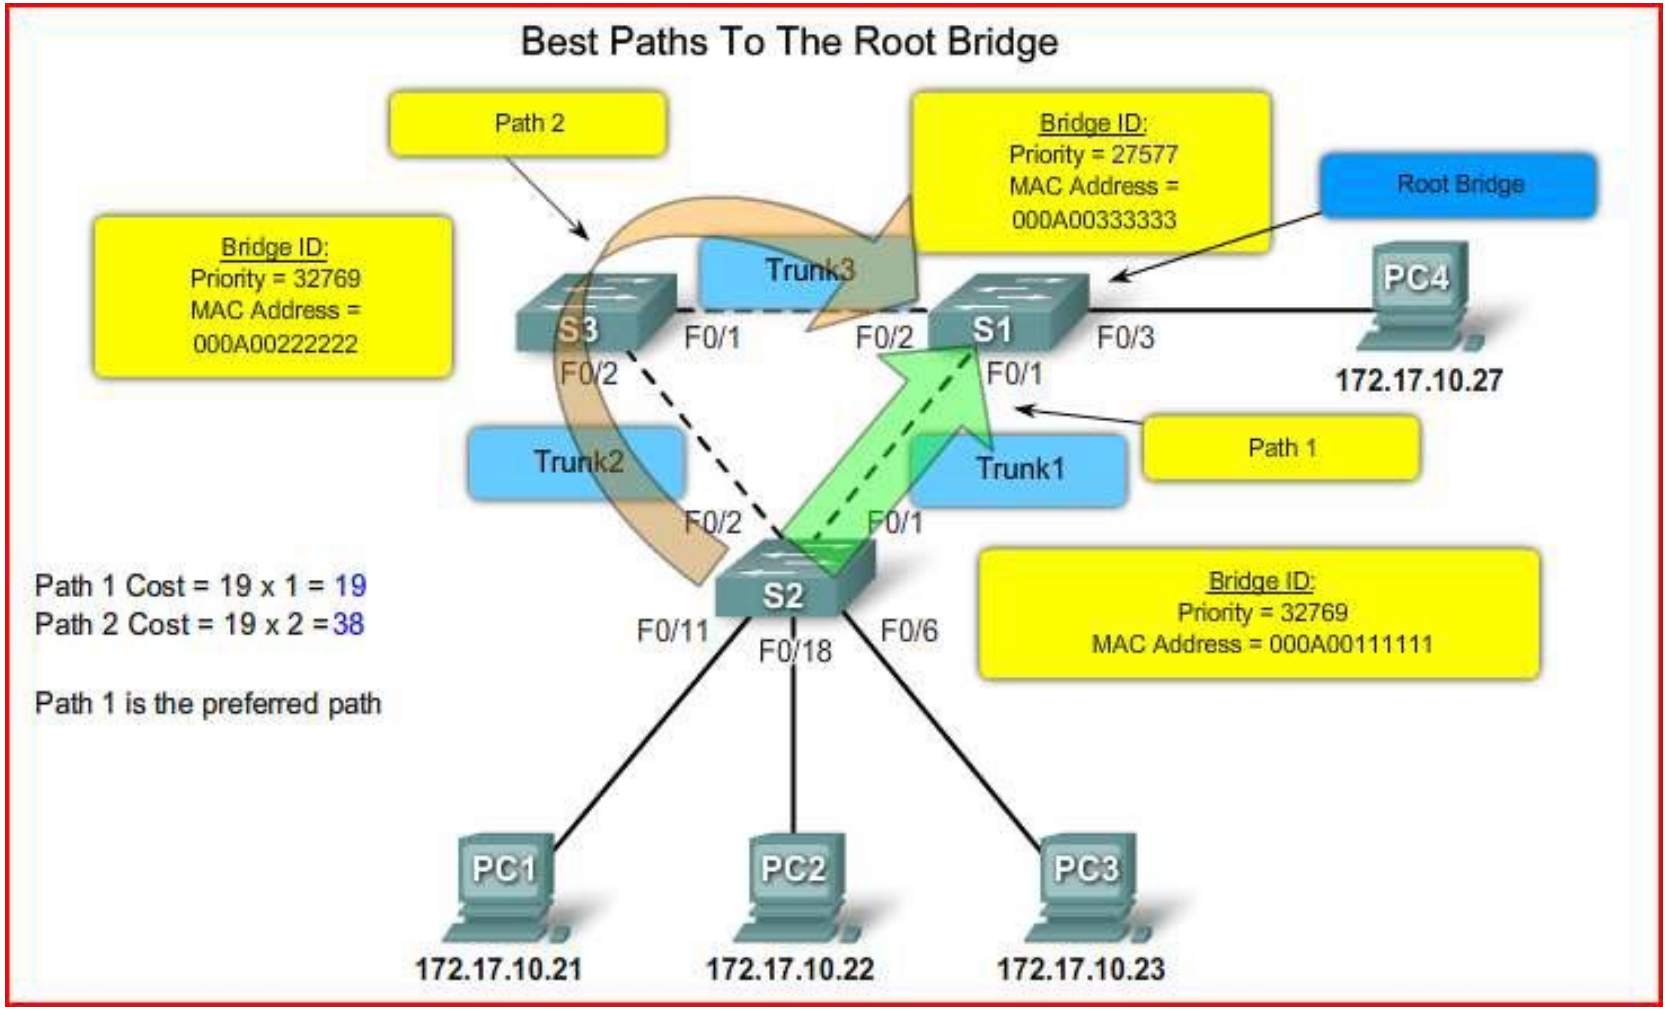
Best Paths To The Root Bridge
Path 2
Bridge ID:
Priority = 27577
MAC Address =
000A00333333
Root Bridge
Bridge ID:
Priority = 32769
MAC Address =
000A00222222
Trunk3
S3
F0/1
F0/2
S1
F0/3
PC4
F0/2
F0/1
172.17.10.27
Path 1
Trunk2
Trunk1
F0/2
F0/1
S2
Path 1 Cost = 19 x 1 = 19
Path 2 Cost = 19 x 2 = 38
Path 1 is the preferred path
Bridge ID:
Priority = 32769
MAC Address = 000A00111111
F0/11
F0/18
F0/6
PC1
PC2
PC3
172.17.10.21
172.17.10.22
172.17.10.23

Verifying Port and Path Cost

```
S2#show spanning-tree

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority 27577
              Address     000A.0033.3333
              Cost        19
              Port        1
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    32769  (priority 32768 sys-id-ext 1)
              Address     000A.0011.1111
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
F0/1             Root FWD 19         128.1    Edge P2p
F0/2             Desg FWD 19         128.2    Edge P2p
```

BPDU Fields

Field #	Bytes	Field
4	2	Protocol ID
	1	Version
	1	Message type
	1	Flags
8	8	Root ID
	4	Cost of path
	8	Bridge ID
	2	Port ID
12	2	Message age
	2	Max age
	2	Hello time
	2	Forward delay

BPDU Example

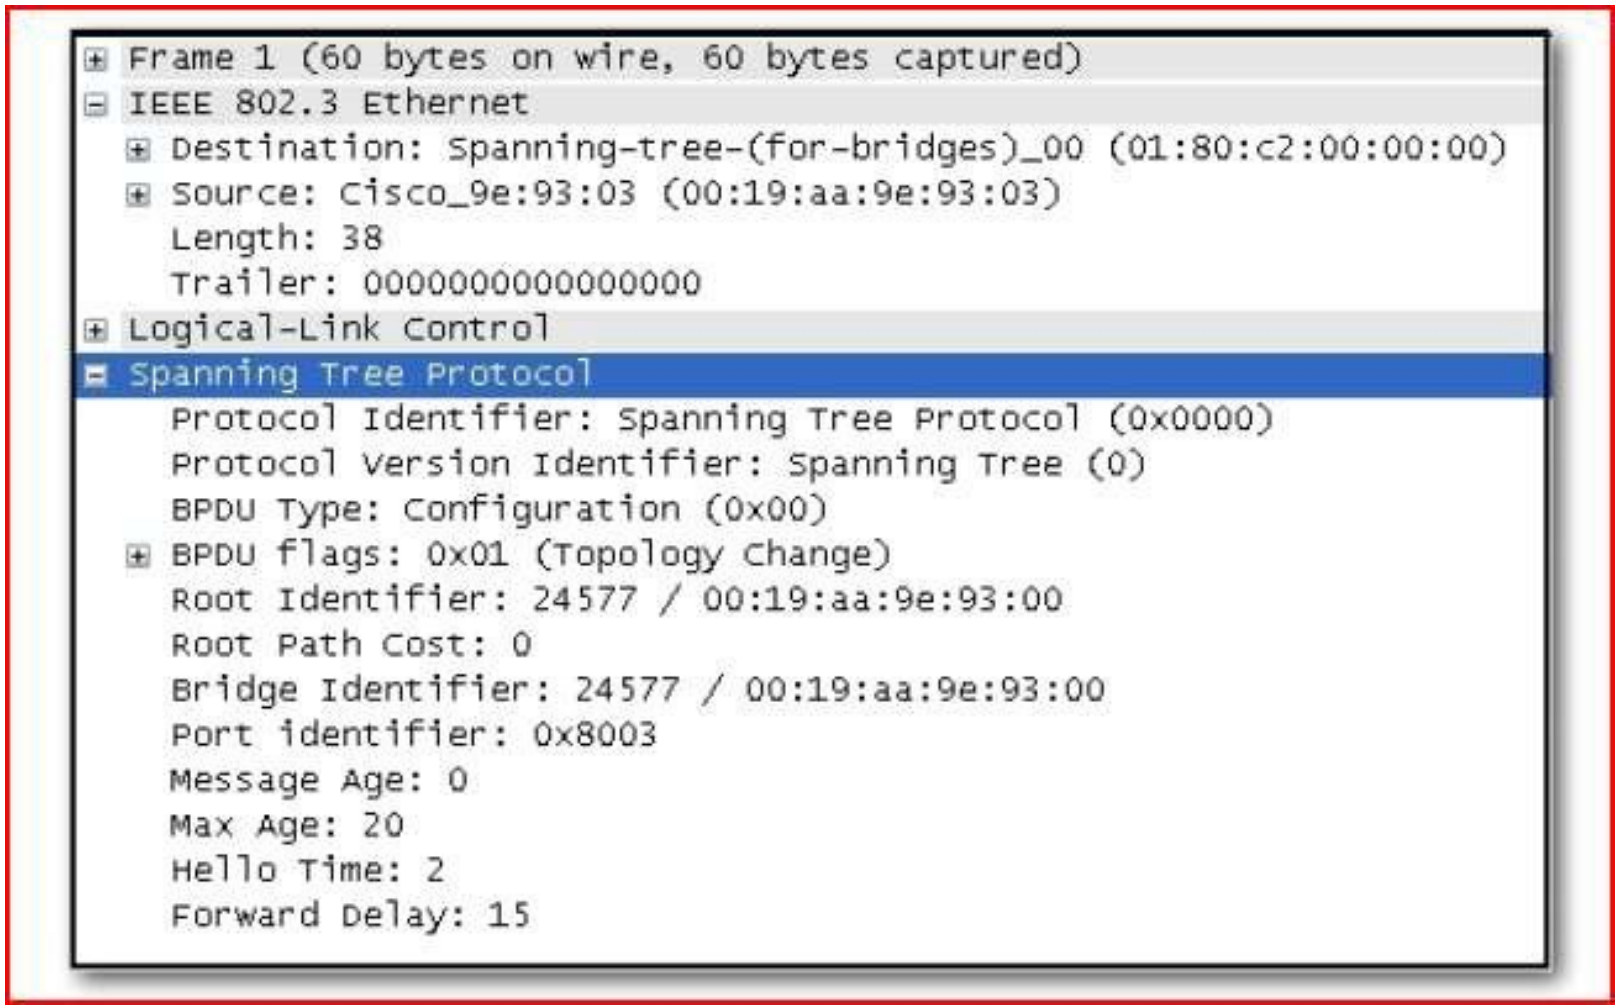

BPDU Process - 1

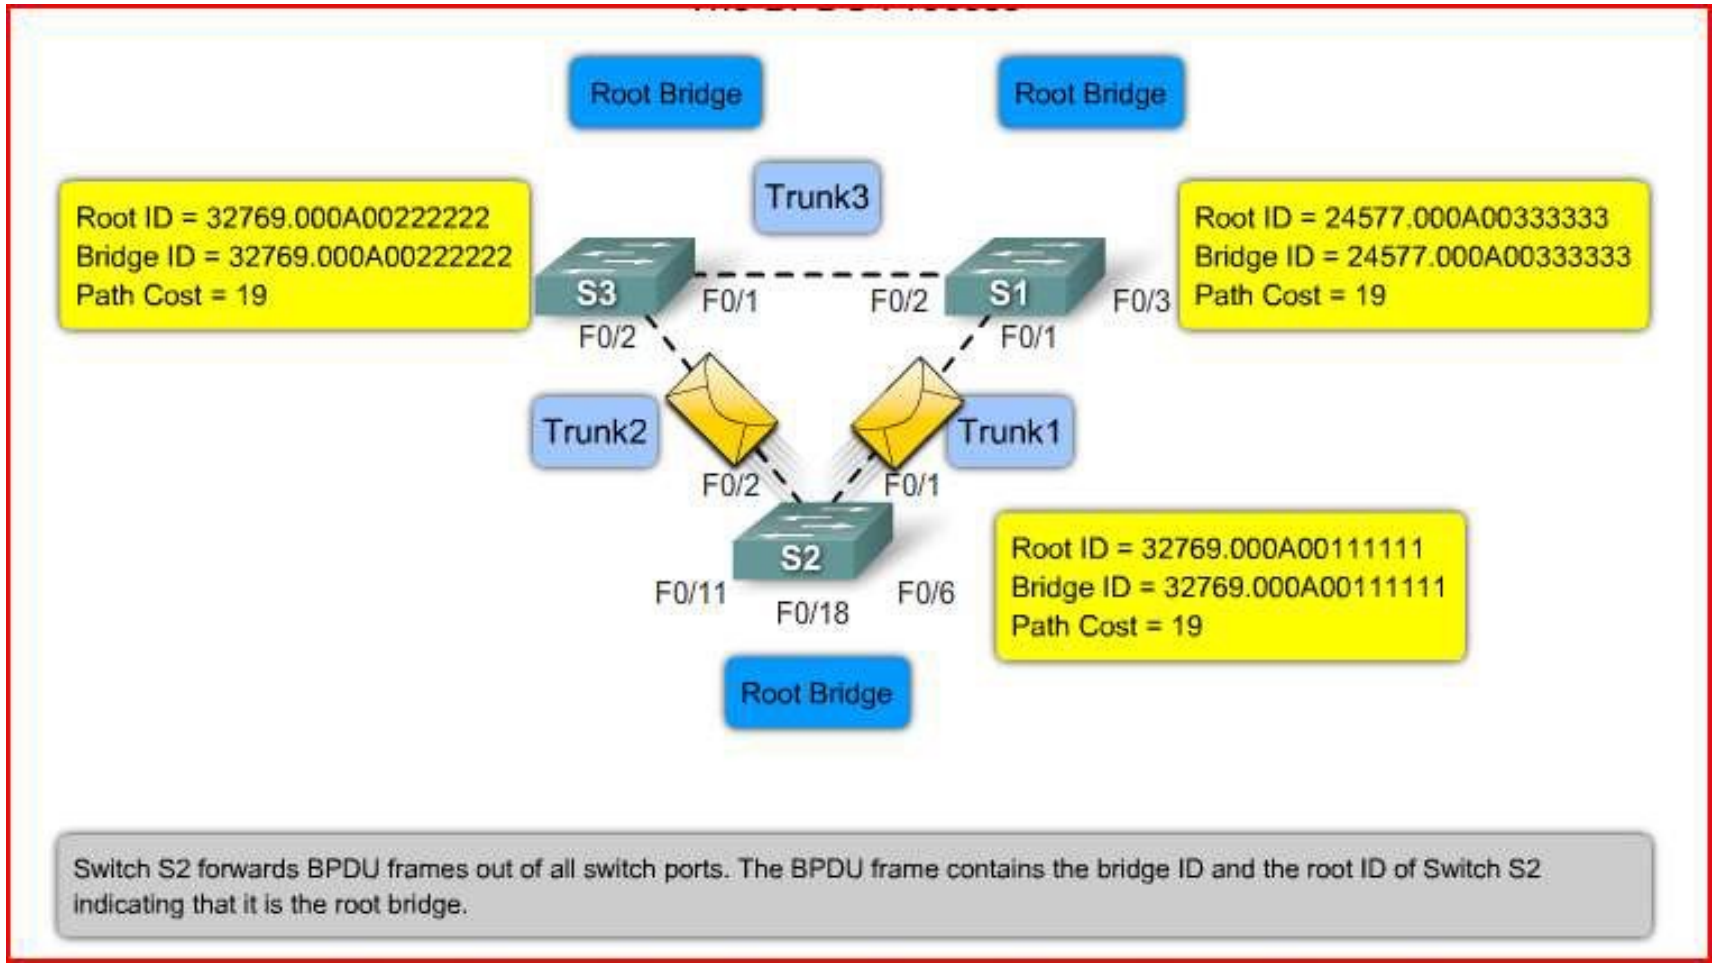

BPDU Process -2

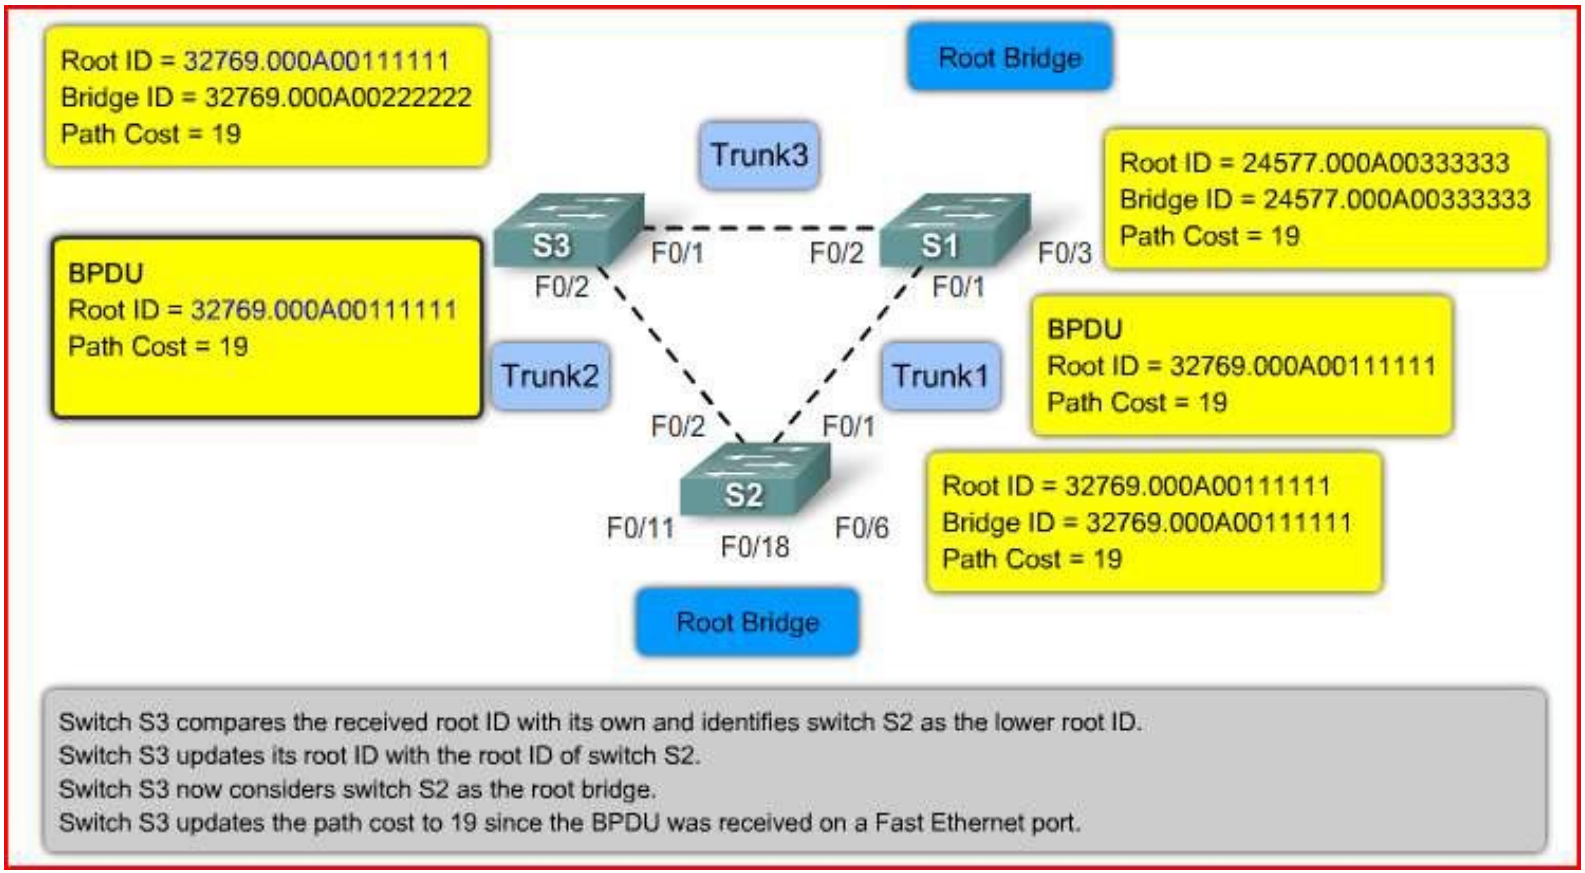

BPDU Process - 3

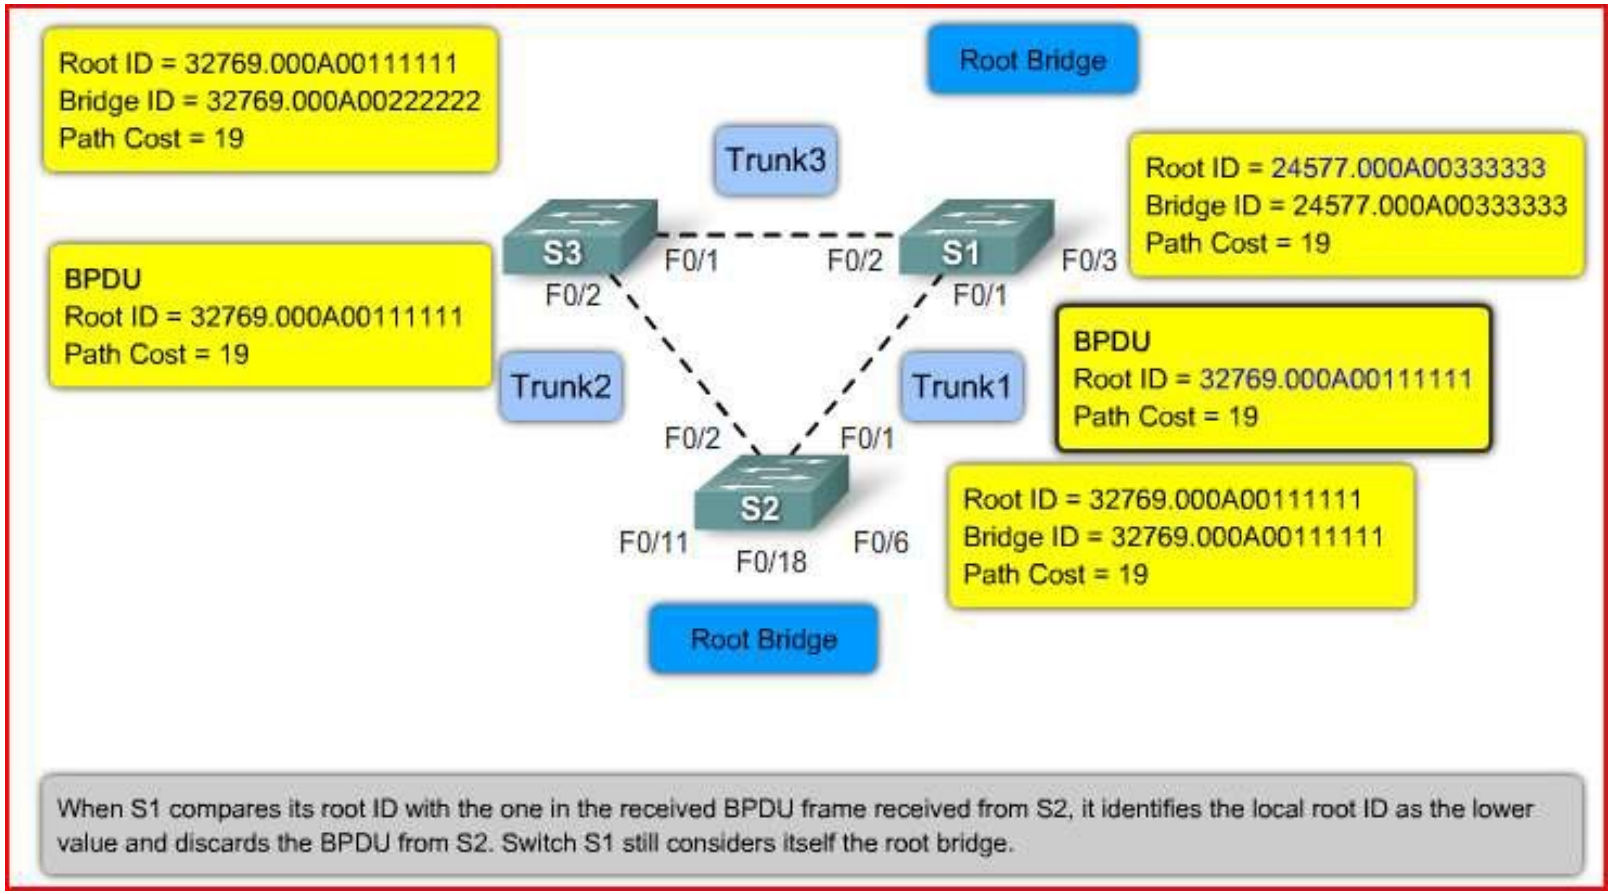

BPDU Process - 4

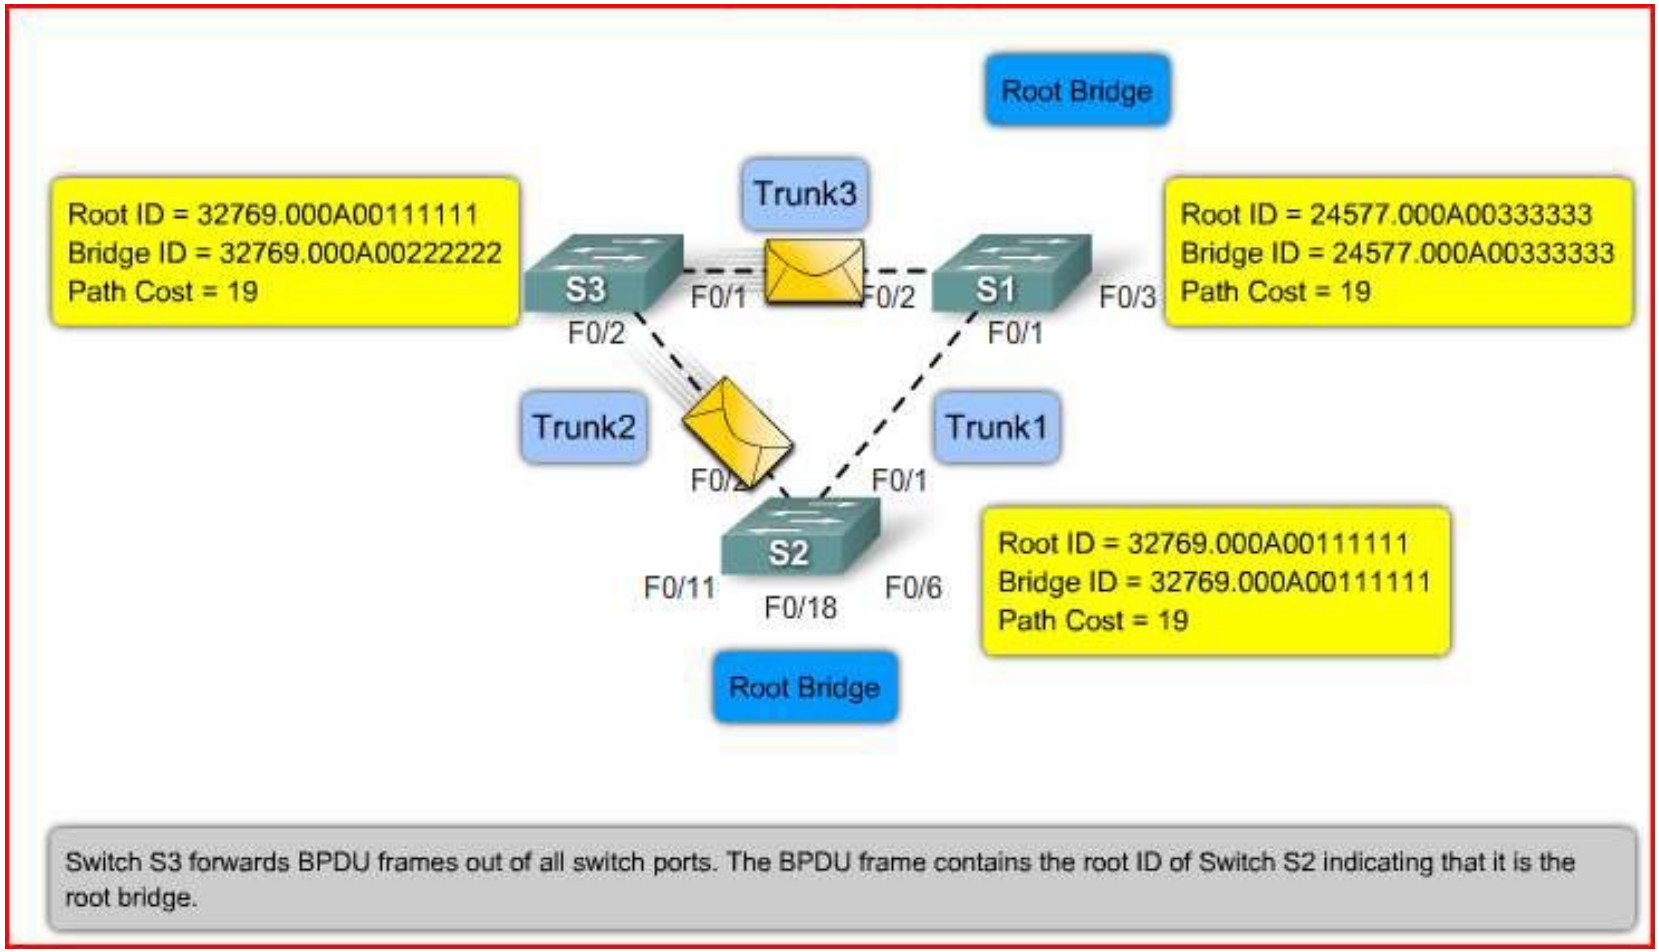

BPDU Process - 5

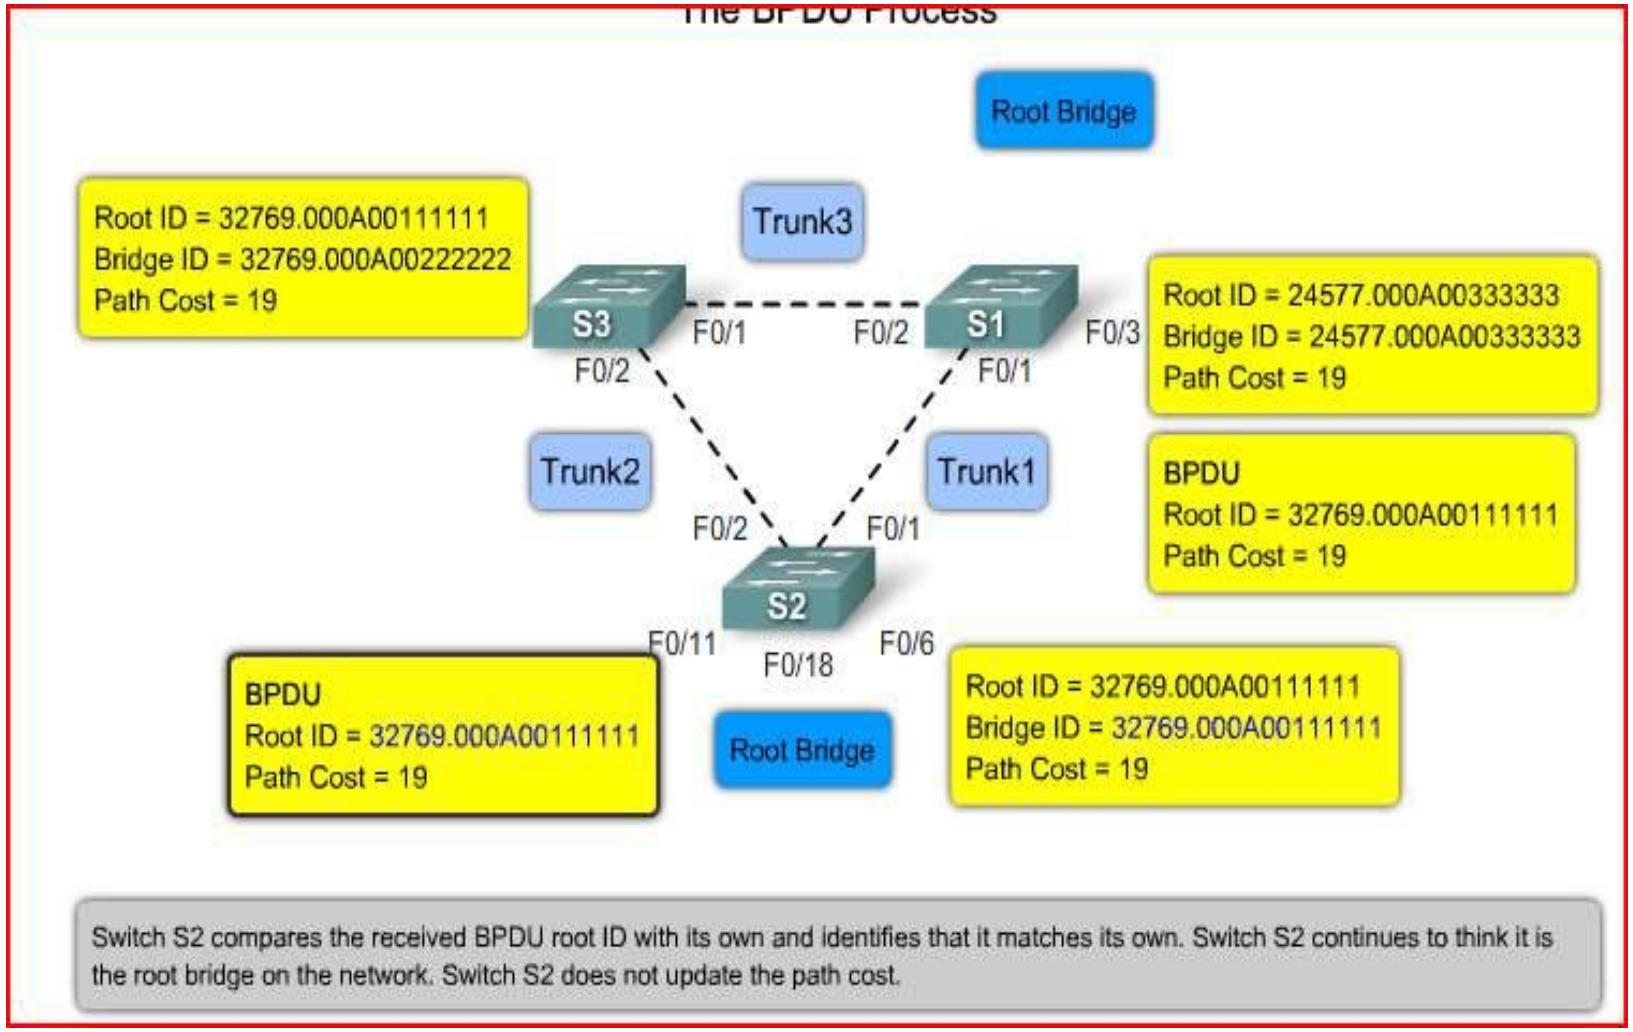

BPDU Process - 6

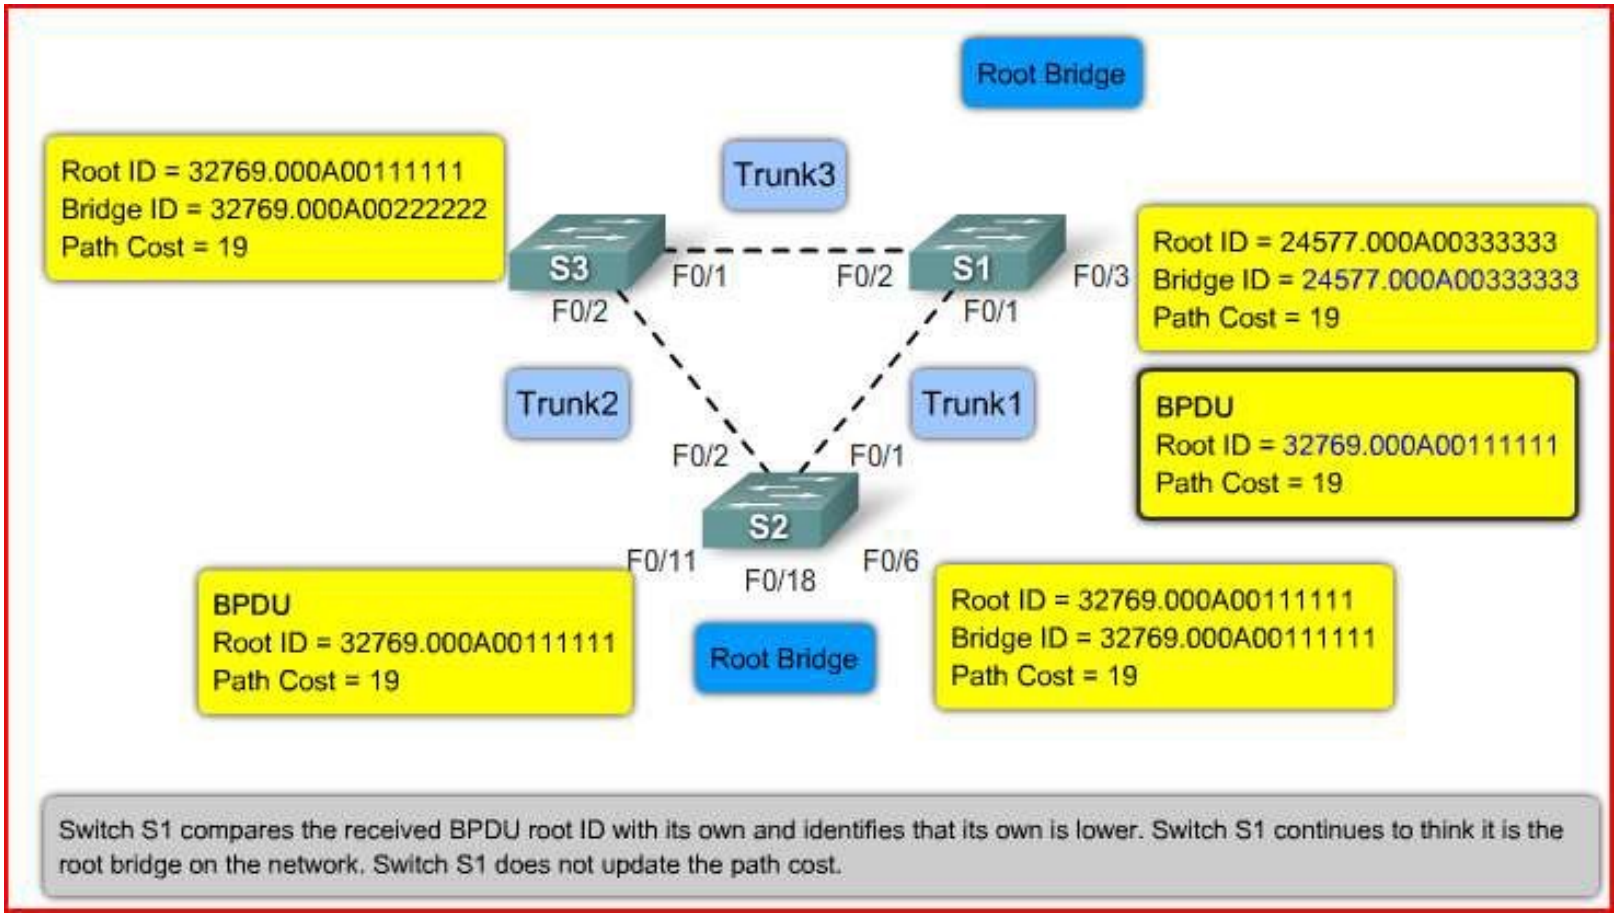

BPDU Process - 7

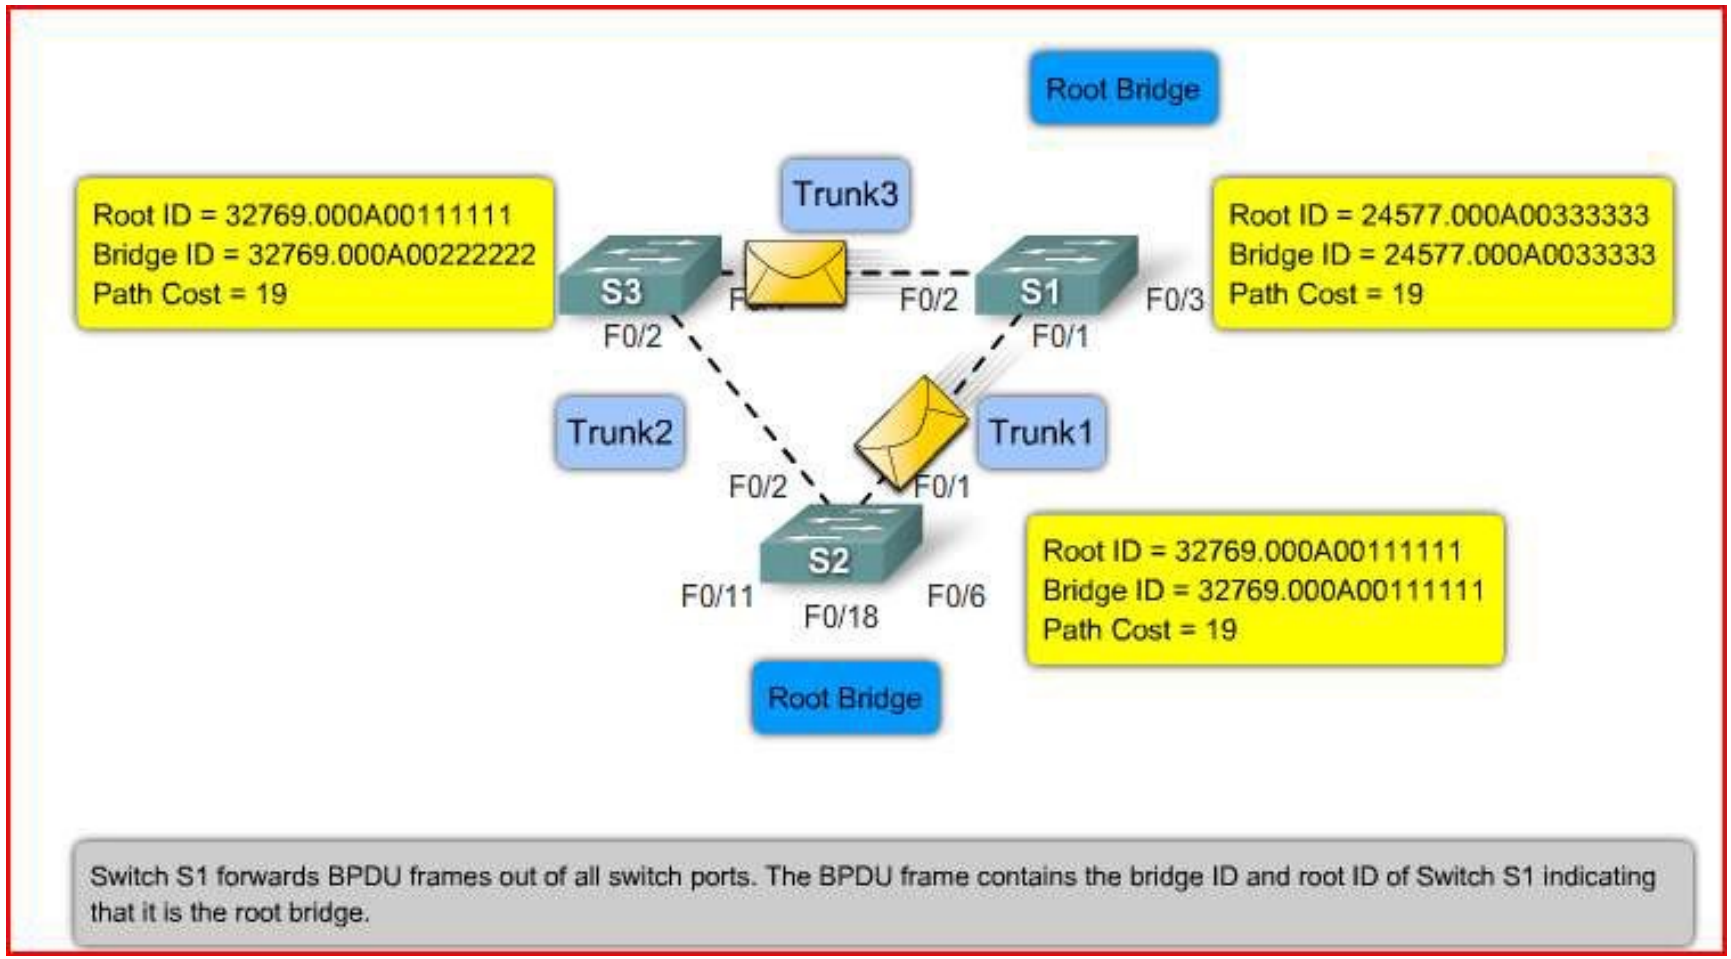

BPDU Process - 8

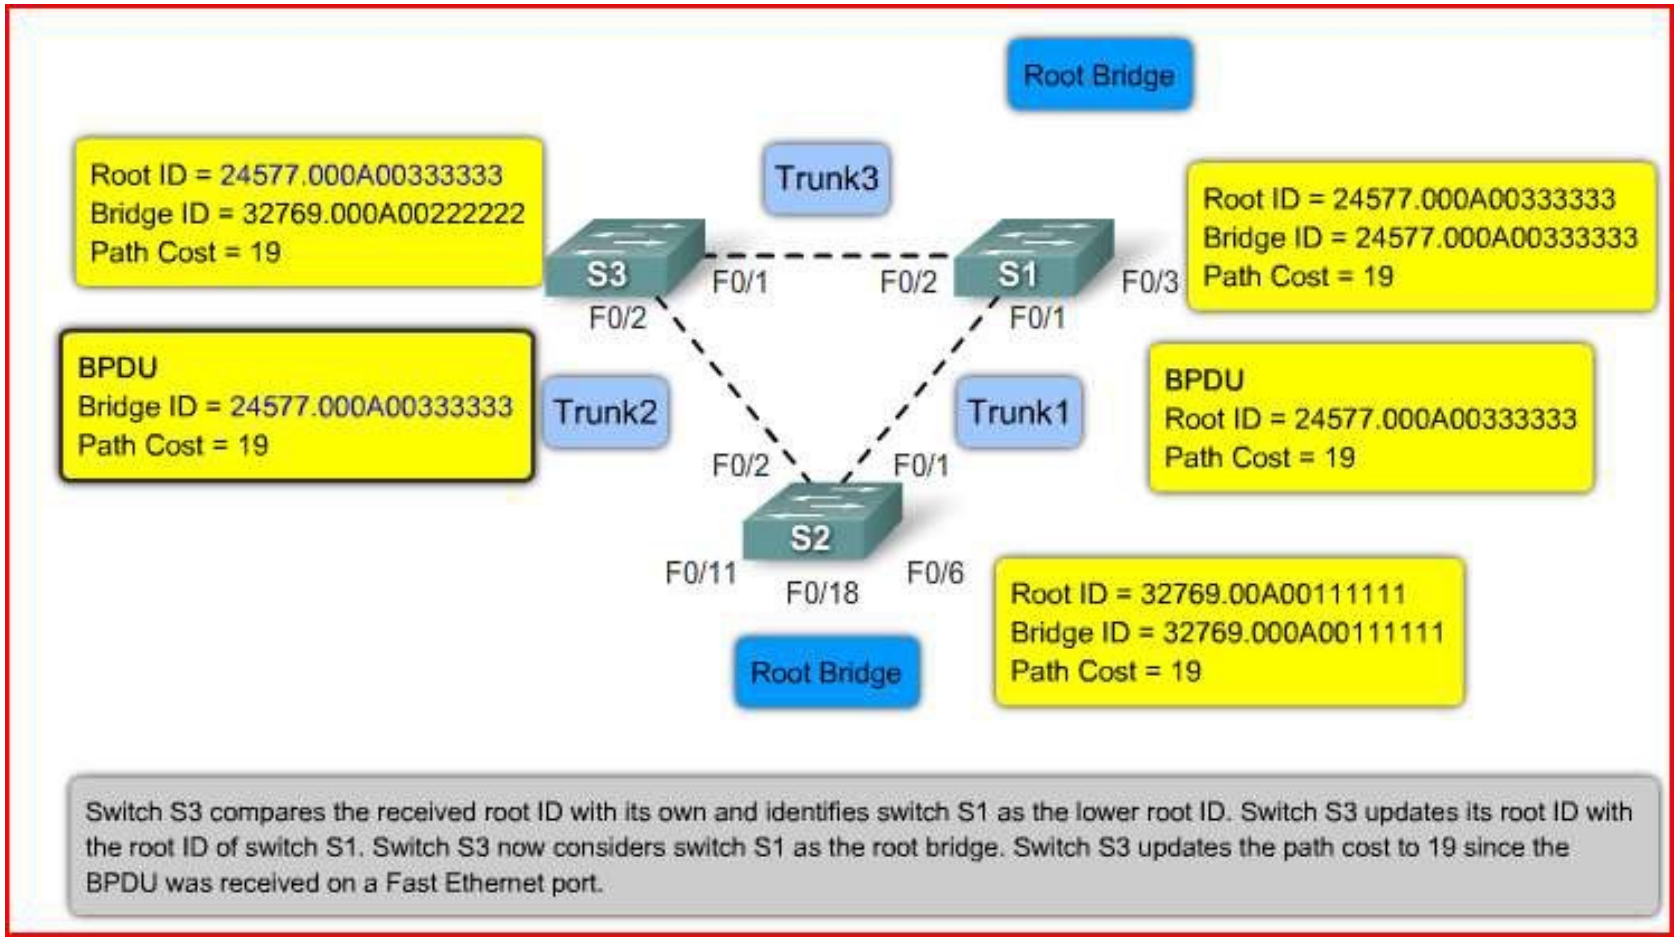

BPDU Process - 9

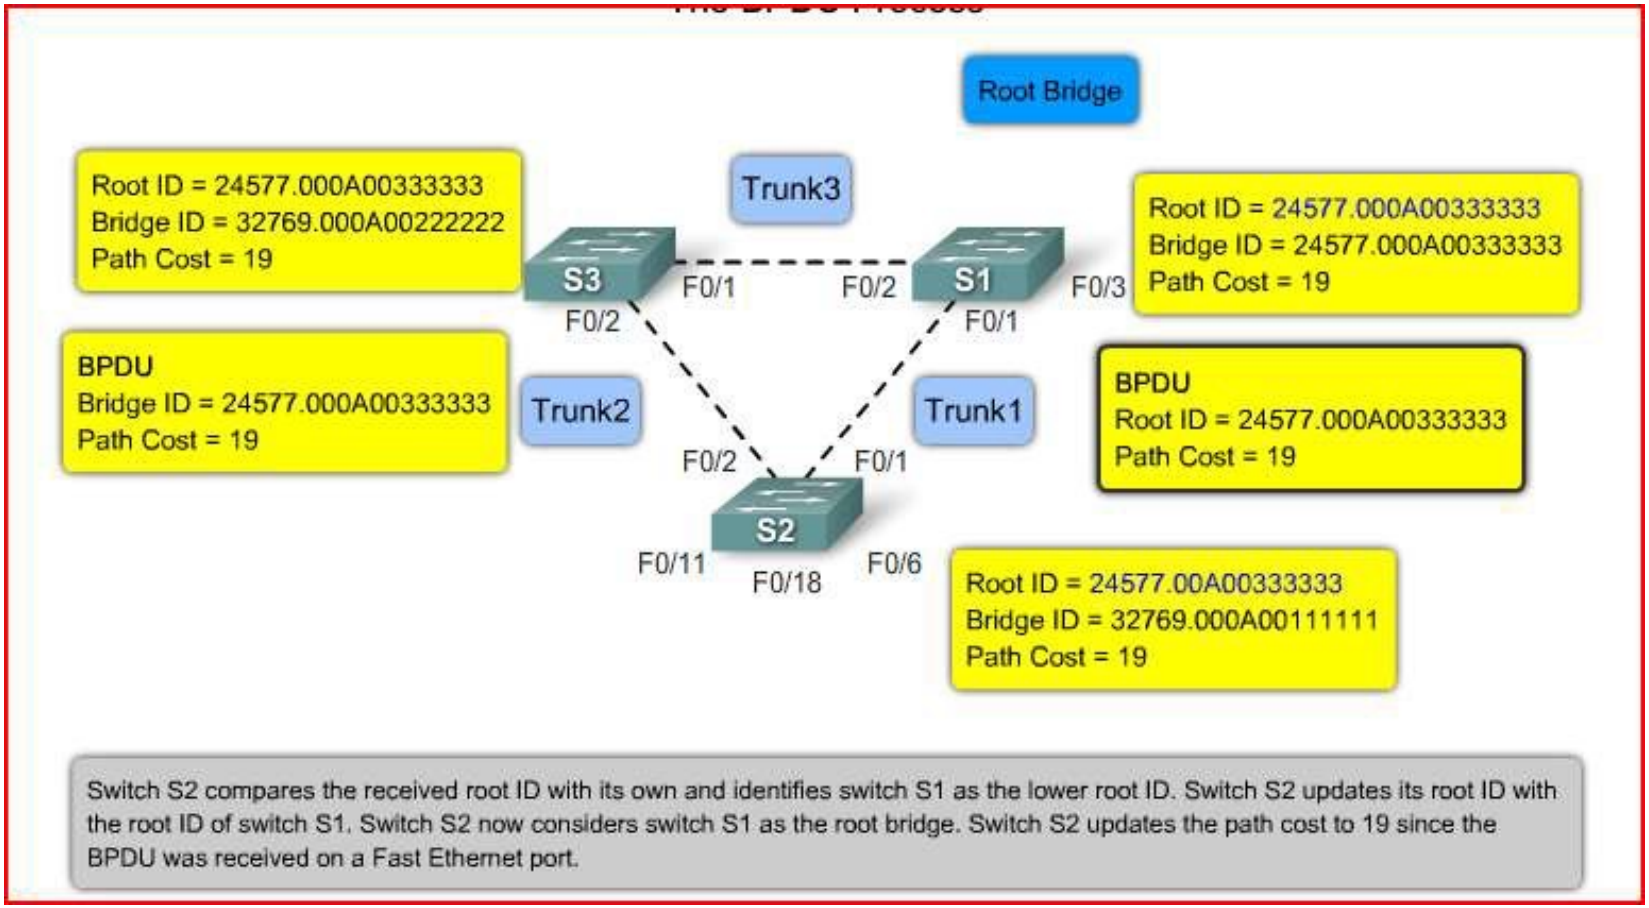

BID Fields

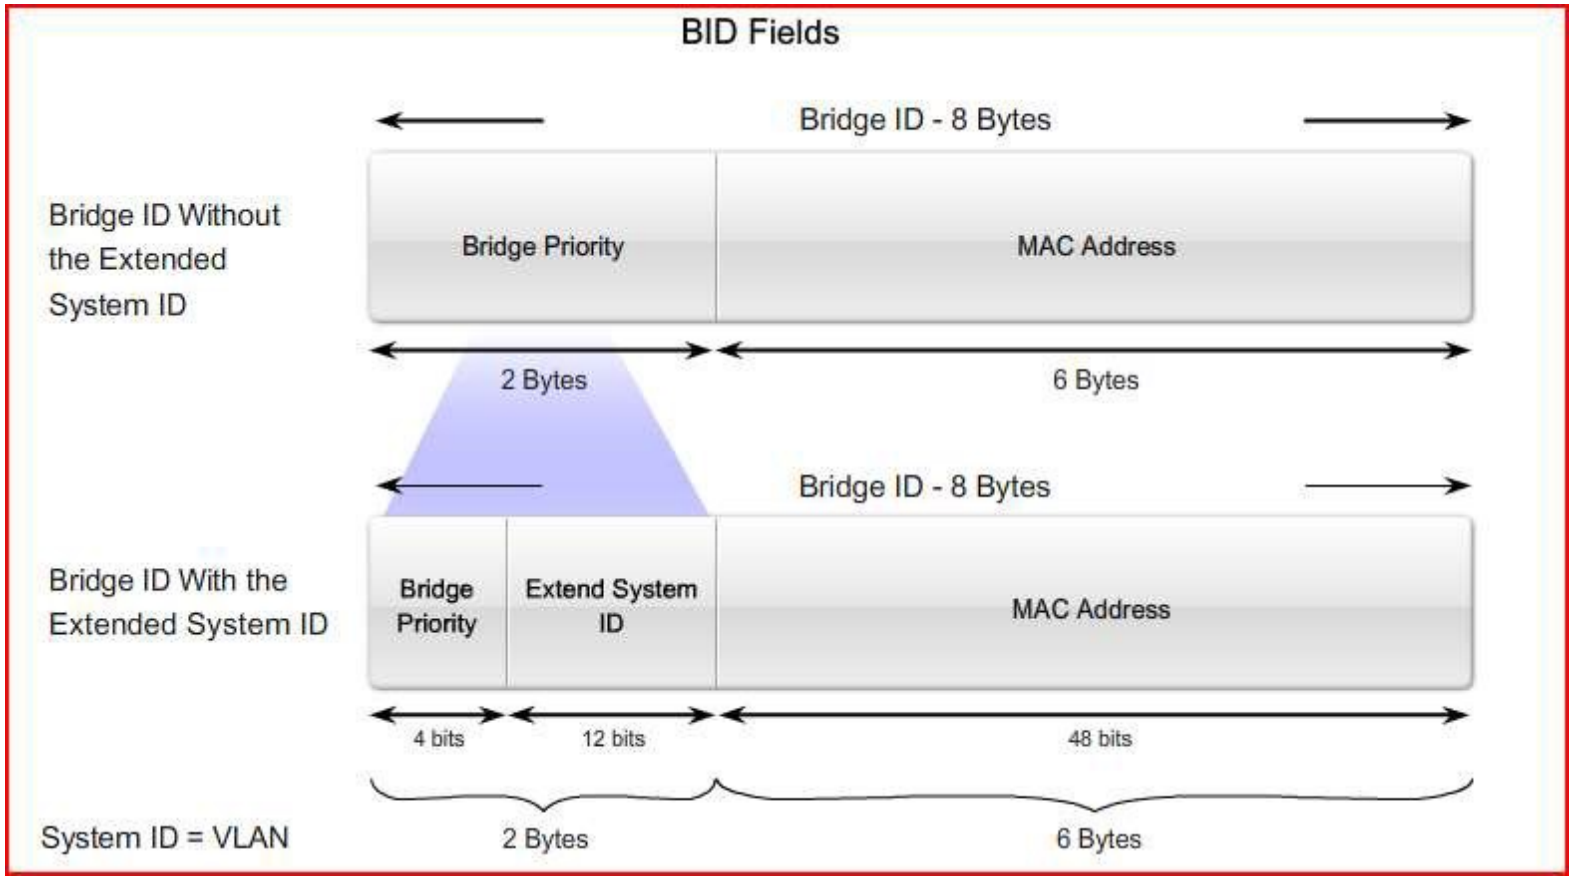
BID Fields
Bridge ID - 8 Bytes
Bridge ID Without the Extended System ID
Bridge Priority
MAC Address
2 Bytes
6 Bytes
Bridge ID - 8 Bytes
Bridge ID With the Extended System ID
Bridge Priority
Extend System ID
MAC Address
4 bits
12 bits
48 bits
System ID = VLAN
2 Bytes
6 Bytes

Bridge Priority

- The bridge priority is a customizable value that you can use to influence which switch becomes the root bridge.
- The switch with the lowest priority, which means lowest BID, becomes the root bridge (the lower the bridge priority value, the higher the STP priority).
- For example, to ensure that a specific switch is always the root bridge, you set the priority to a lower value than the rest of the switches on the network.
- The default value for the priority of all Cisco switches is 32768. The priority range is between 1 and 65536; therefore, 1 is the highest priority.

Priority-Based Decision

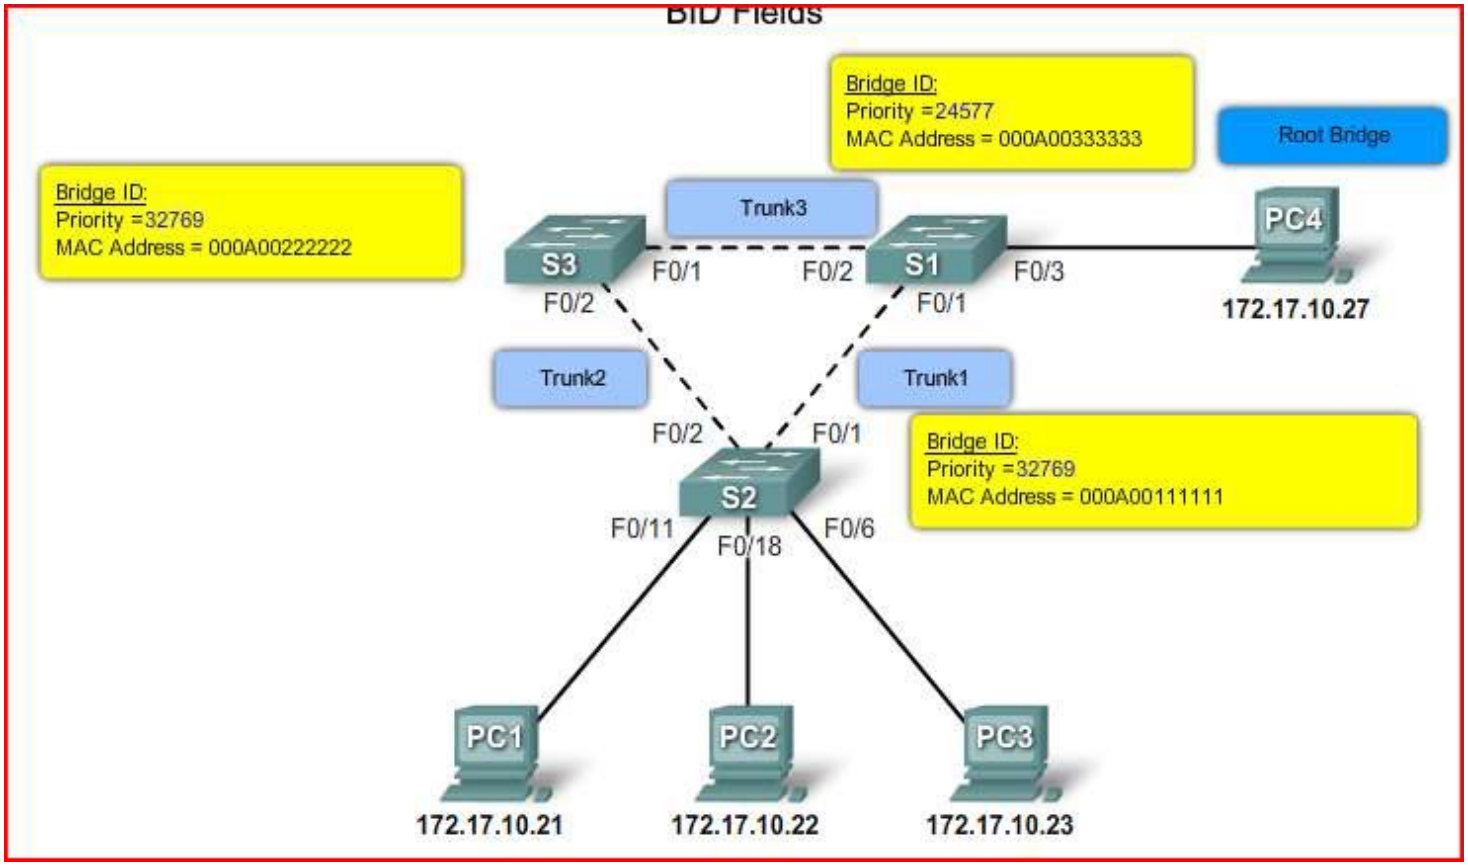

Extended System ID

- The early implementation of STP was designed for networks that did not use VLANs. There was a single common spanning tree across all switches. When VLANs started to become common for network infrastructure segmentation, STP was enhanced to include support for VLANs. As a result, the extended system ID field contains the ID of the VLAN with which the BPDU is associated.
- When the extended system ID is used, it changes the number of bits available for the bridge priority value, so the increment for the bridge priority value changes from 1 to 4096. Therefore, bridge priority values can only be multiples of 4096.
- The extended system ID value is added to the bridge priority value in the BID to identify the priority and VLAN of the BPDU frame.
- More about per VLAN spanning tree (PVST) in a later section of this chapter.

MAC Address

- When two switches are configured with the same priority and have the same extended system ID, the switch with the MAC address with the lowest hexadecimal value has the lower BID.
- Initially, all switches are configured with the same default priority value. The MAC address is then the deciding factor.
- It is recommended to configure the desired root bridge switch with a lower priority to ensure that it is elected root bridge. This also ensures that the addition of new switches to the network does not trigger a new spanning-tree election, which could disrupt network communication while a new root bridge is being selected.

MAC-Address Based

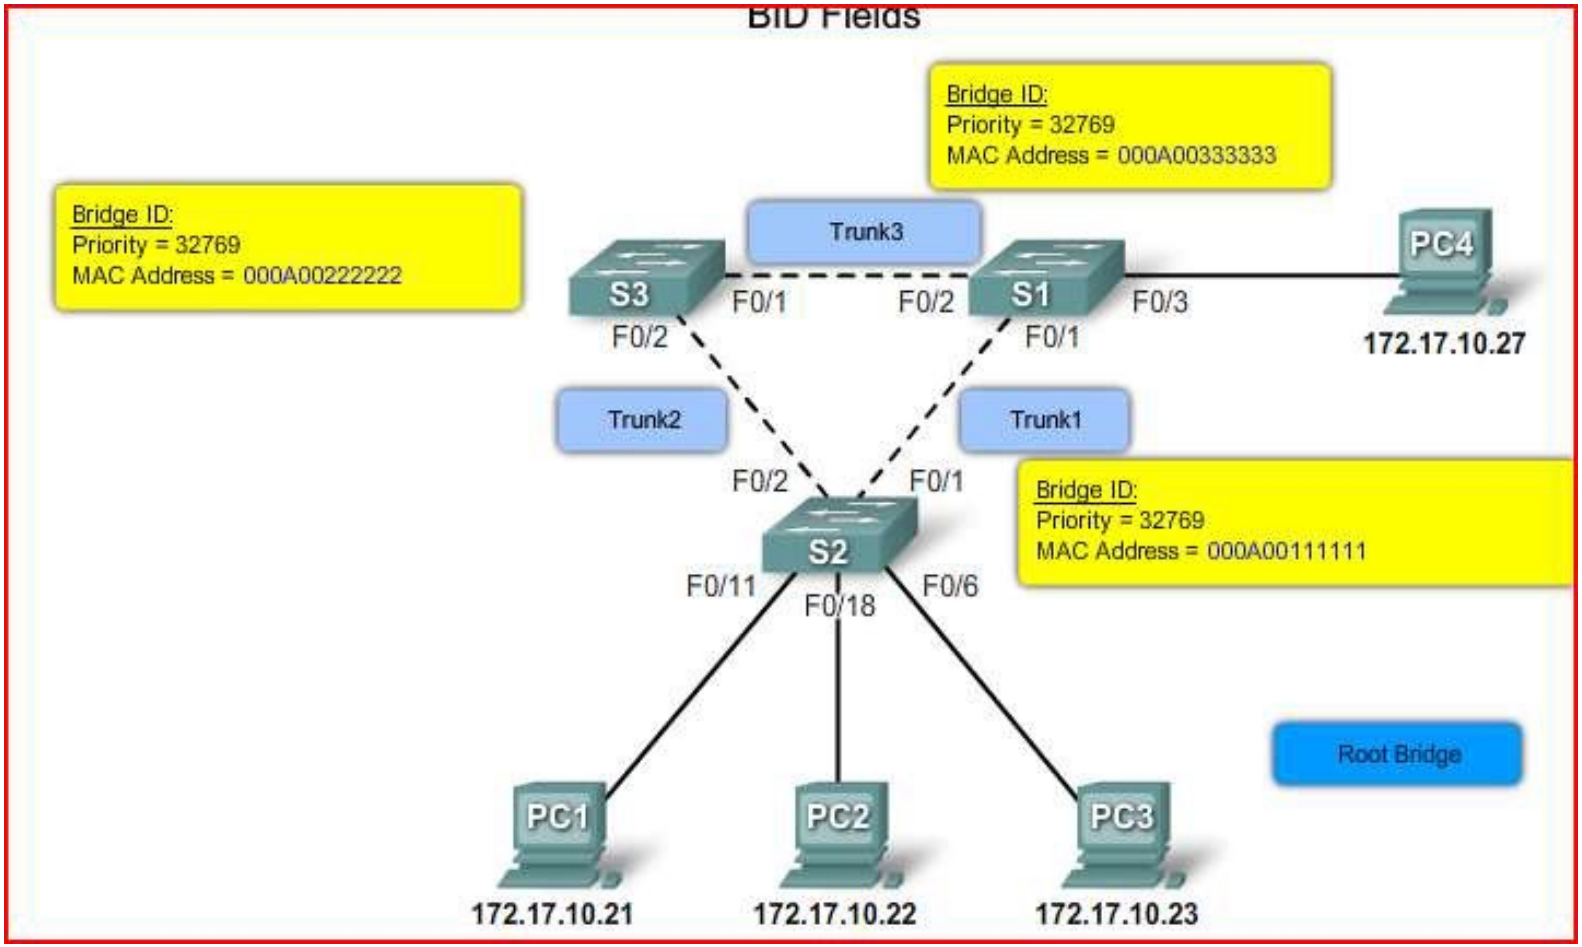

Configure BID

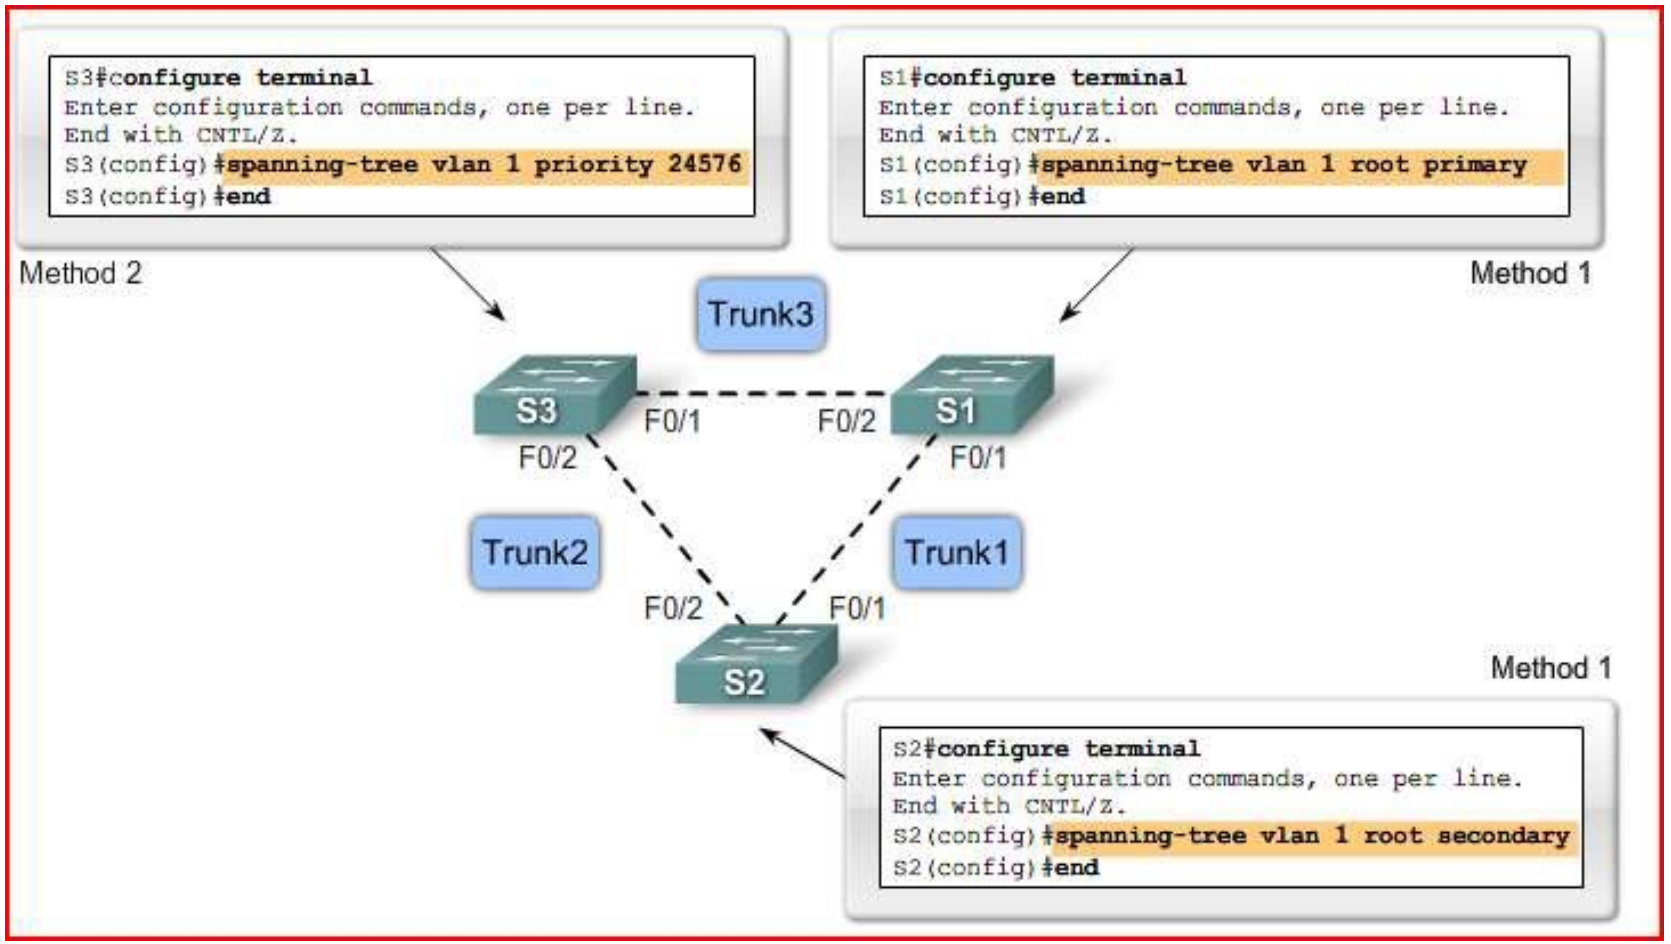

Verify BID

```
S1#show spanning-tree

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority    24577
              Address     000A.0033.3333
              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     000A.0033.3333
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/1            Desg FWD 4         128.1    Shr
Fa0/2            Desg FWD 4         128.2    Shr
S1#
```

How does STP works to Eliminate Layer 2 Loops in a Converged Network

The role of STP port states and BPDU timers in the operation of STP

Port States

Processes	Blocking	Listening	Learning	Forwarding	Disable
Receives and process BPDUs	✓	✓[1]	✓	✓	✗
Forward data frames received on interface	✗	✗	✗	✓	✗
Forward data frames switched from another interface	✗	✗	✗	✓	✗
Learn MAC addresses	✗	✗	✓	✓	✗

[1]Return to blocking if not lowest cost path to root bridge

BPDU Timers

Hello time	The hello time is the time between each BPDU frame that is sent on a port. This is equal to 2 seconds by default, but can be tuned to be between 1 and 10 seconds.
Forward delay	The forward delay is the time spent in the listening and learning state. This is by default equal to 15 seconds for each state, but can be tuned to be between 4 and 30 seconds.
Maximum age	The max age timer controls the maximum length of time a switch port saves configuration BPDU information. This is 20 seconds by default, but can be tuned to be between 6 and 40 seconds.

Port States

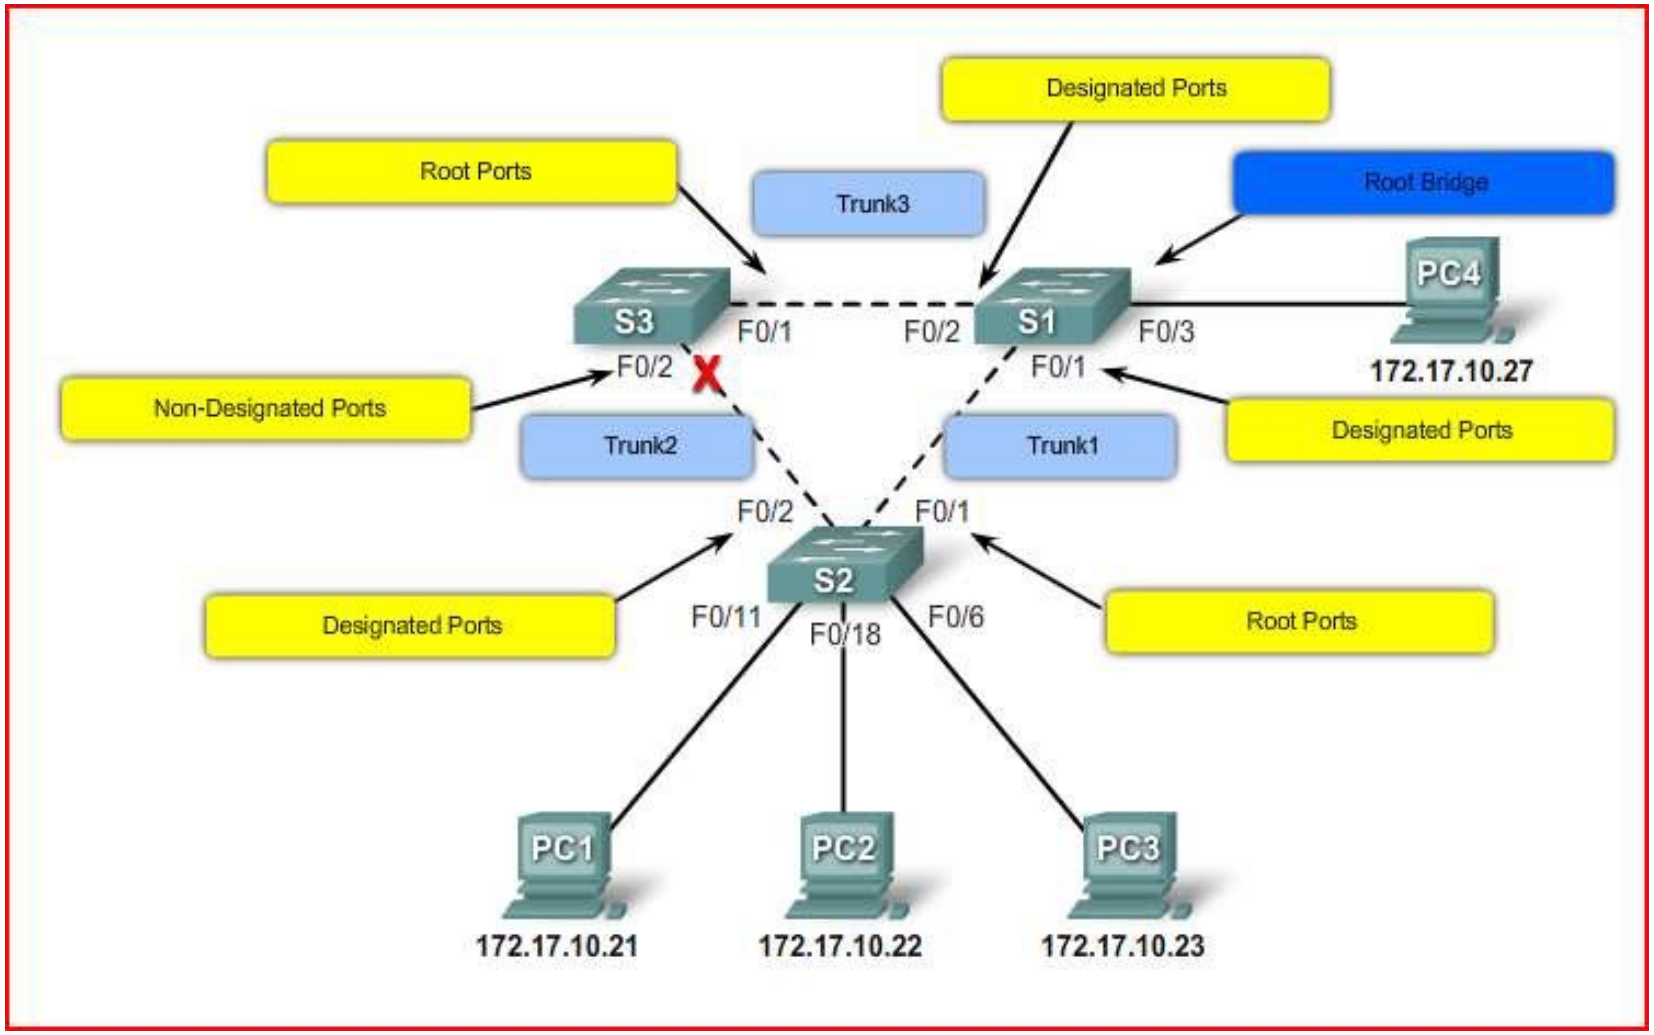
Designated Ports
Root Ports
Trunk3
Root Bridge
PC4
S3
F0/1
F0/2
S1
F0/3
F0/2
F0/1
172.17.10.27
Non-Designated Ports
Trunk2
Trunk1
Designated Ports
F0/2
F0/1
S2
Designated Ports
F0/11
F0/18
F0/6
Root Ports
PC1
PC2
PC3
172.17.10.21
172.17.10.22
172.17.10.23

Root Port

- The root port exists on non-root bridges and is the switch port with the best path to the root bridge.
- Root ports forward traffic toward the root bridge. The source MAC address of frames received on the root port are capable of populating the MAC table. Only one root port is allowed per bridge.
- In the example, switch S1 is the root bridge and switches S2 and S3 have root ports defined on the trunk links connecting back to S1.

Designated Port

- The designated port exists on root and non-root bridges.
- For root bridges, all switch ports are designated ports.
- For non-root bridges, a designated port is the switch port that receives and forwards frames toward the root bridge as needed.
- Only one designated port is allowed per segment. If multiple switches exist on the same segment, an election process determines the designated switch, and the corresponding switch port begins forwarding frames for the segment. Designated ports are capable of populating the MAC table.
- In the example, switch S1 has both sets of ports for its two trunk links configured as designated ports. Switch S2 also has a designated port configured on the trunk link going toward switch S3.

Non-designated Ports

- The non-designated port is a switch port that is blocked, so it is not forwarding data frames and not populating the MAC address table with source addresses.
- A non-designated port is not a root port or a designated port. For some variants of STP, the non-designated port is called an alternate port.
- In the example, switch S3 has the only non-designated ports in the topology. The non-designated ports prevent the loop from occurring.

Port Roles

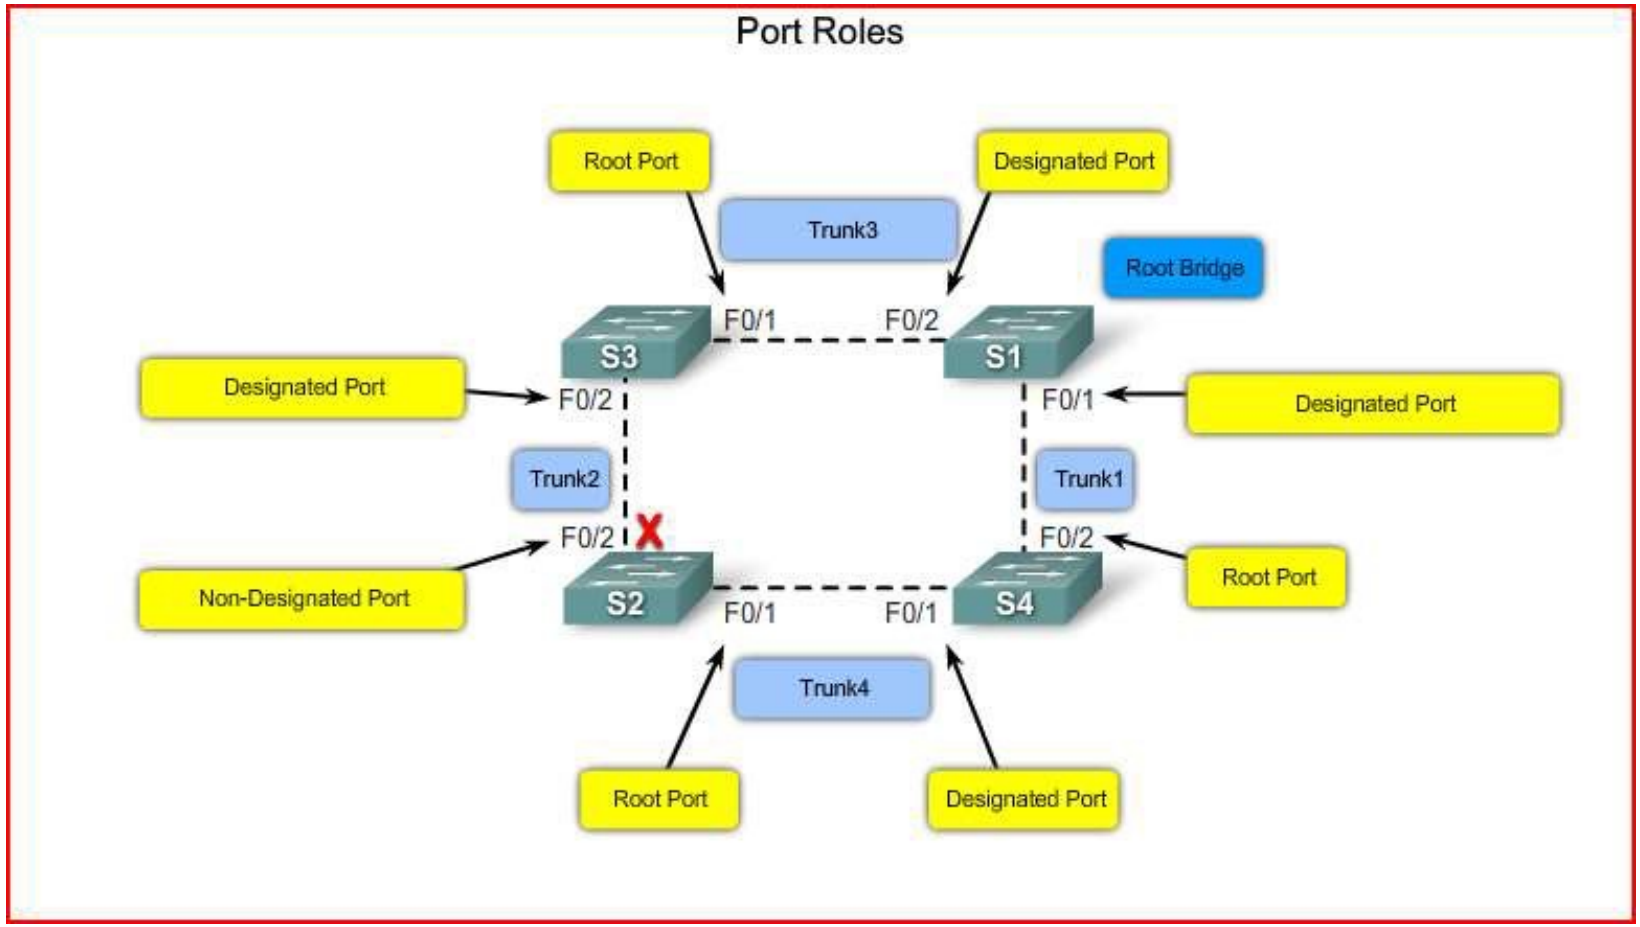
Port Roles
Root Port
Designated Port
Trunk3
Root Bridge
F0/1
F0/2
S3
S1
Designated Port
F0/2
F0/1
Designated Port
Trunk2
Trunk1
F0/2
F0/2
Non-Designated Port
S2
F0/1
F0/1
S4
Root Port
Trunk4
Root Port
Designated Port

STA Assigns Port Roles

- The switch compares the path costs on all switch ports participating in the spanning tree. The switch port with the lowest overall path cost to the root is automatically assigned the root port role because it is closest to the root bridge. In a network topology, all switches that are using spanning tree, except for the root bridge, have a single root port defined.
- When there are two switch ports that have the same path cost to the root bridge and both are the lowest path costs on the switch, the switch needs to determine which switch port is the root port. The switch uses the customizable port priority value, or the lowest port ID if both port priority values are the same.

Port ID

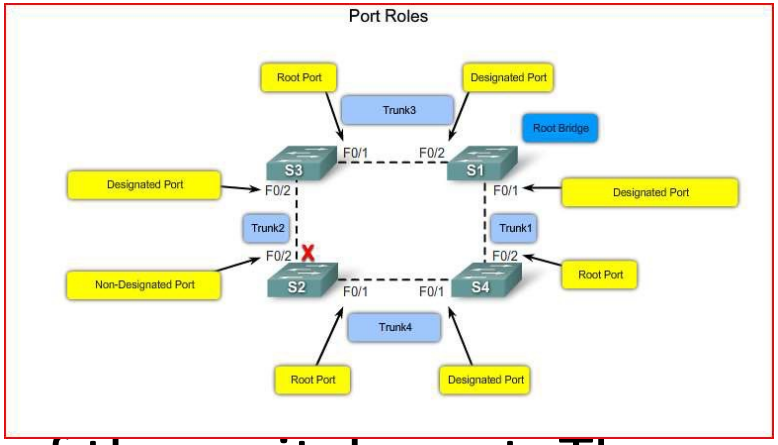

- The port ID is the interface ID of the switch port. The figure shows four switches. Port F0/1 and F0/2 on switch S2 have the same path cost value back to the root bridge. However, port F0/1 on switch S2 is the preferred port because it has a lower port ID value.
- The port ID is appended to the port priority. For example, switch port F0/1 has a default port priority value of 128.1, where 128 is the configurable port priority value, and .1 is the port ID. Switch port F0/2 has a port priority value of 128.2, by default.

Port Priority Configuration

```
S2#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
S2(config)#interface f0/1
S2(config-if)#spanning-tree port-priority 112
S2(config-if)#end
S2#
```

NOTE: Default priority is 128

Verifying Port Roles and Priority

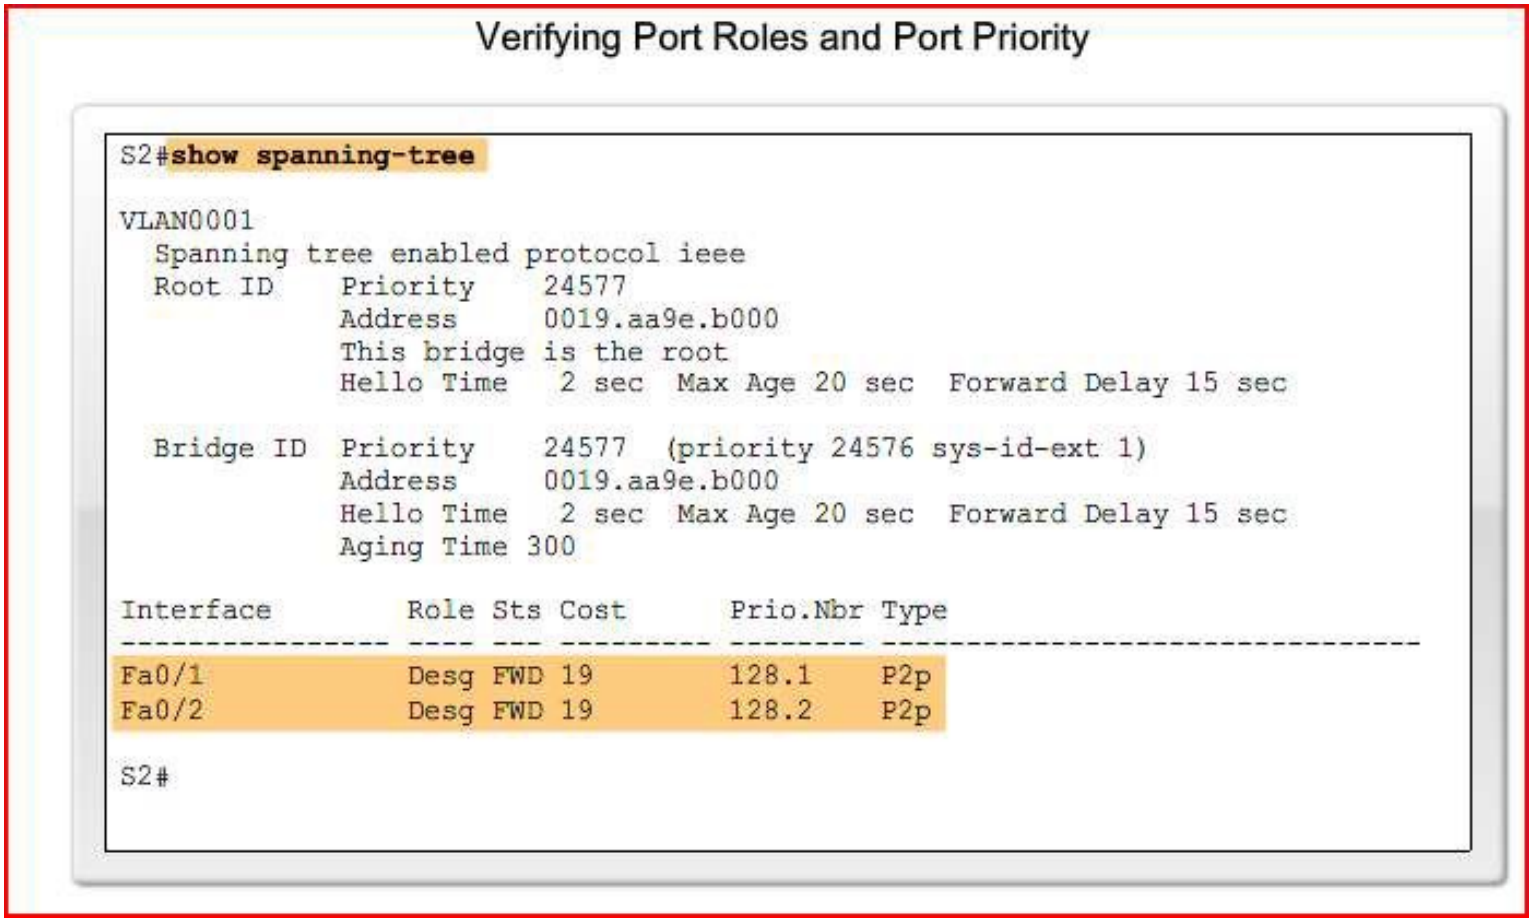

Port States

Processes	Blocking	Listening	Learning	Forwarding	Disable
Receives and process BPDUs	[1]YES	YES	YES	YES	NO
Forward data frames received on interface	NO	NO	NO	YES	NO
Forward data frames switched from another interface	NO	NO	NO	YES	NO
Learn MAC addresses	NO	NO	YES	YES	NO

Port States

- STP determines the logical loop-free path throughout the broadcast domain. The spanning tree is determined through the information learned by the exchange of the BPDU frames between the interconnected switches. To facilitate the learning of the logical spanning tree, each switch port transitions through five possible port states and three BPDU timers.
- The spanning tree is determined immediately after a switch is finished booting up. If a switch port were to transition directly from the blocking to the forwarding state, the port could temporarily create a data loop if the switch was not aware of all topology information at the time. For this reason, STP introduces five port states.
- Blocking, Listening, Learning, Forwarding, Disabled

Blocking

- The port is a non-designated port and does not participate in frame forwarding.
- The port receives BPDU frames to determine the location and root ID of the root bridge switch and what port roles each switch port should assume in the final active STP topology.

Listening

- STP has determined that the port can participate in frame forwarding according to the BPDU frames that the switch has received thus far.
- At this point, the switch port is not only receiving BPDU frames, it is also transmitting its own BPDU frames and informing adjacent switches that the switch port is preparing to participate in the active topology.

Learning, Forwarding, Disabled

- Learning - The port prepares to participate in frame forwarding and begins to populate the MAC address table.
- Forwarding - The port is considered part of the active topology and forwards frames and also sends and receives BPDU frames.
- Disabled - The Layer 2 port does not participate in spanning tree and does not forward frames. The disabled state is set when the switch port is administratively disabled.

BPDU Timers

Hello time	The hello time is the time between each BPDU frame that is sent on a port. This is equal to 2 seconds by default, but can be tuned to be between 1 and 10 seconds.
Forward delay	The forward delay is the time spent in the listening and learning state. This is by default equal to 15 seconds for each state, but can be tuned to be between 4 and 30 seconds.
Maximum age	The max age timer controls the maximum length of time a switch port saves configuration BPDU information. This is 20 seconds by default, but can be tuned to be between 6 and 40 seconds.

BPDU Timers 2

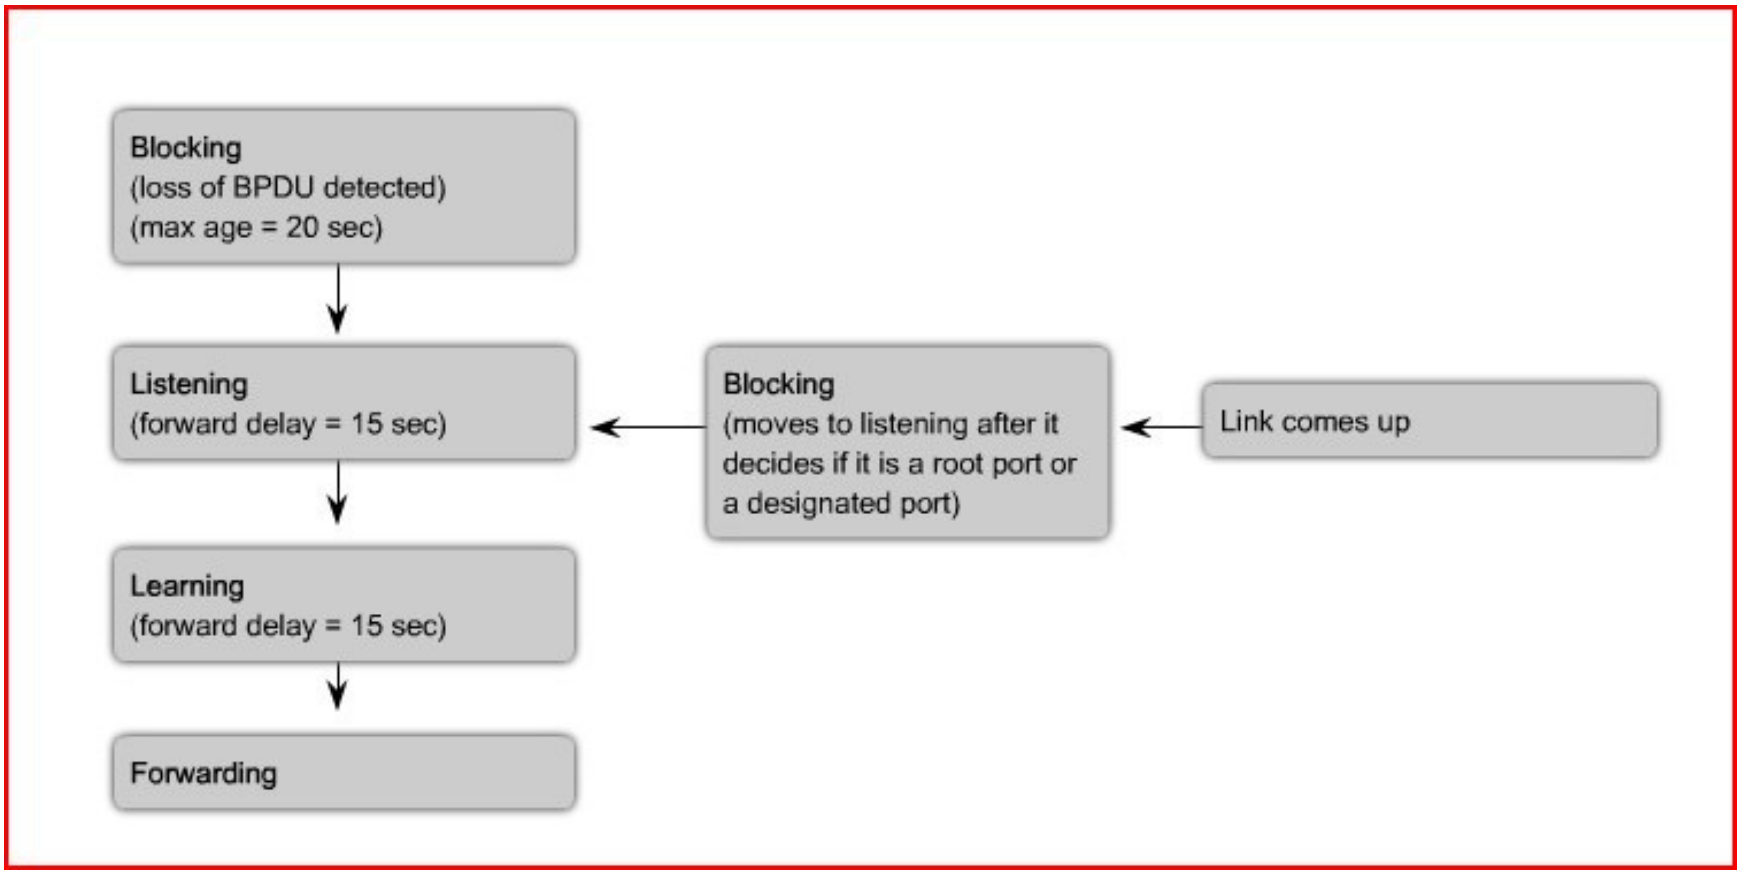

Network Diameter

- It is recommended that the BPDU timers not be adjusted directly because the values have been optimized for the seven-switch diameter.
- Adjusting the spanning-tree diameter value on the root bridge to a lower value automatically adjusts the forward delay and maximum age timers proportionally for the new diameter.
- Typically, you do not adjust the BPDU timers nor reconfigure the network diameter. However, if after research, a network administrator determined that the convergence time of the network could be optimized, the administrator would do so by reconfiguring the network diameter, not the BPDU timers.

Configure Network Diameter

```
S1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
S1(config)#spanning-tree vlan 1 root primary diameter 5
S1(config)#end
S1#
```

Cisco Portfast

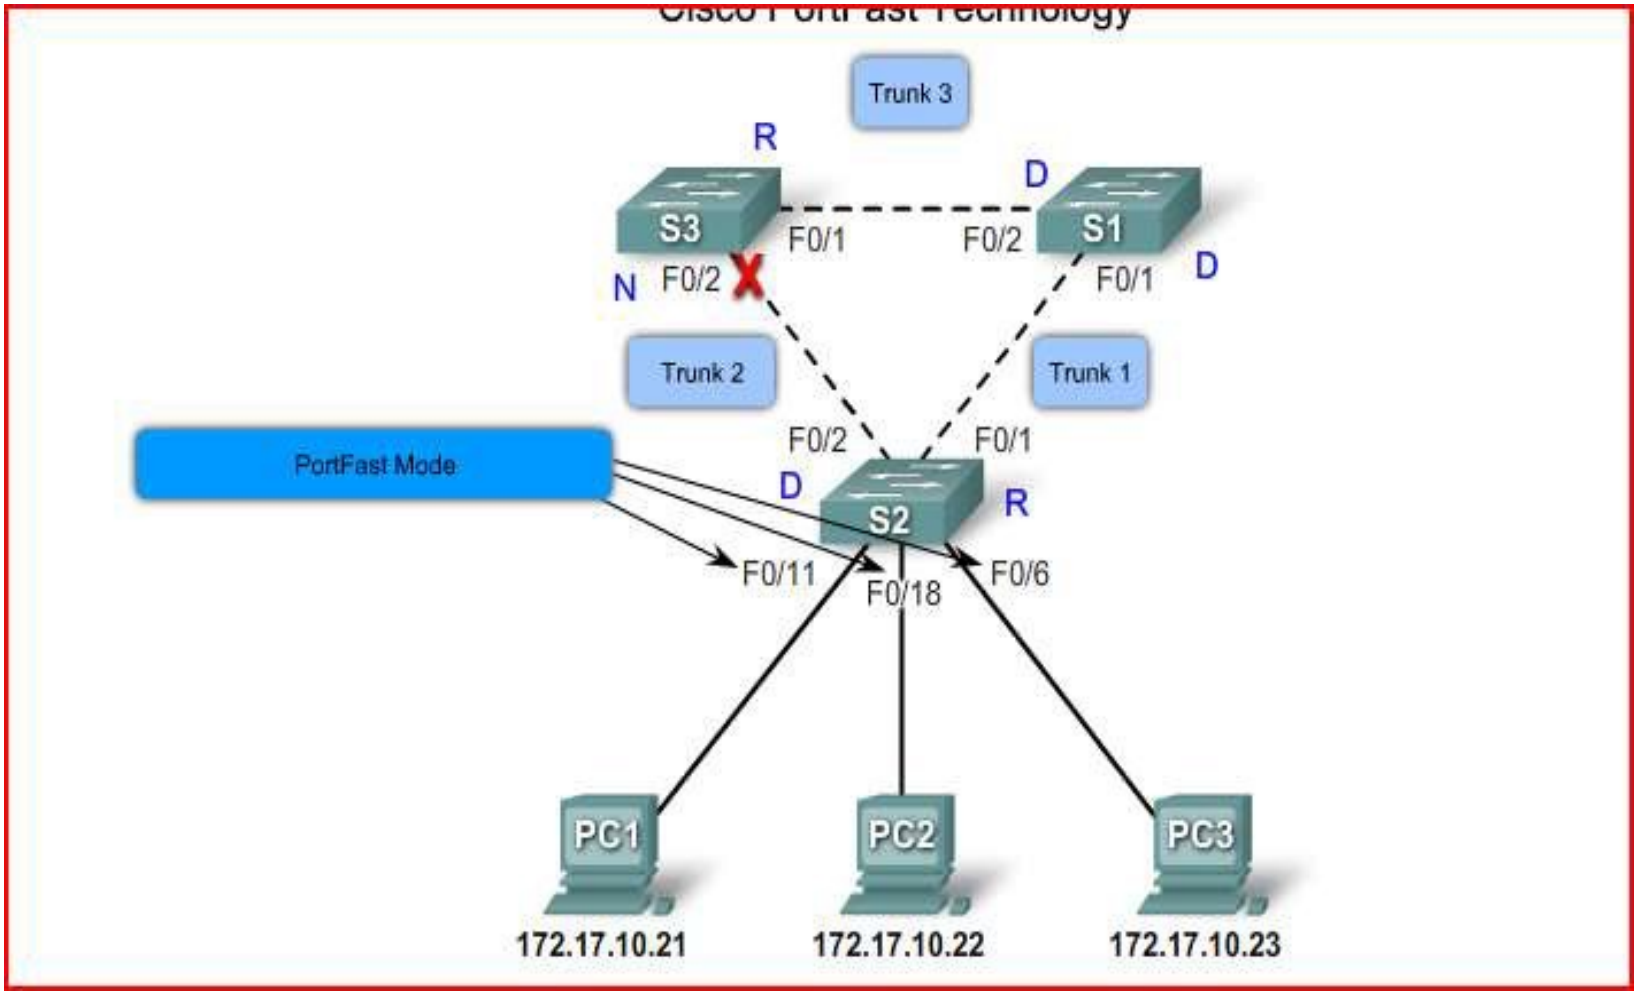

Portfast Configuration

Enable PortFast

```
S2(config)# interface FastEthernet 0/11
S2(config-if)# spanning-tree portfast
 Warning: portfast should only be enabled on ports connected to a single
 host. Connecting hubs, concentrators, switches, bridges, etc... to this
 interface when portfast is enabled, can cause temporary bridging loops.
 Use with CAUTION

 Portfast has been configured on FastEthernet0/11 but will only
 have effect when the interface is in a non-trunking mode.
S2(config-if)# end
```

Disable PortFast

```
S2(config)# interface FastEthernet 0/11
S2(config-if)# no spanning-tree portfast
S2(config-if)# end
```

Verify Portfast

```
S2#show running-config
<output omitted>
!
interface FastEthernet0/11
 switchport mode access
 spanning-tree portfast
!
<output omitted>
end
S2#
```

How does the STP Algorithm Uses Three Steps to Converge on a Loop-Free Topology

The STP decision sequence is used to elect a root bridge for a network

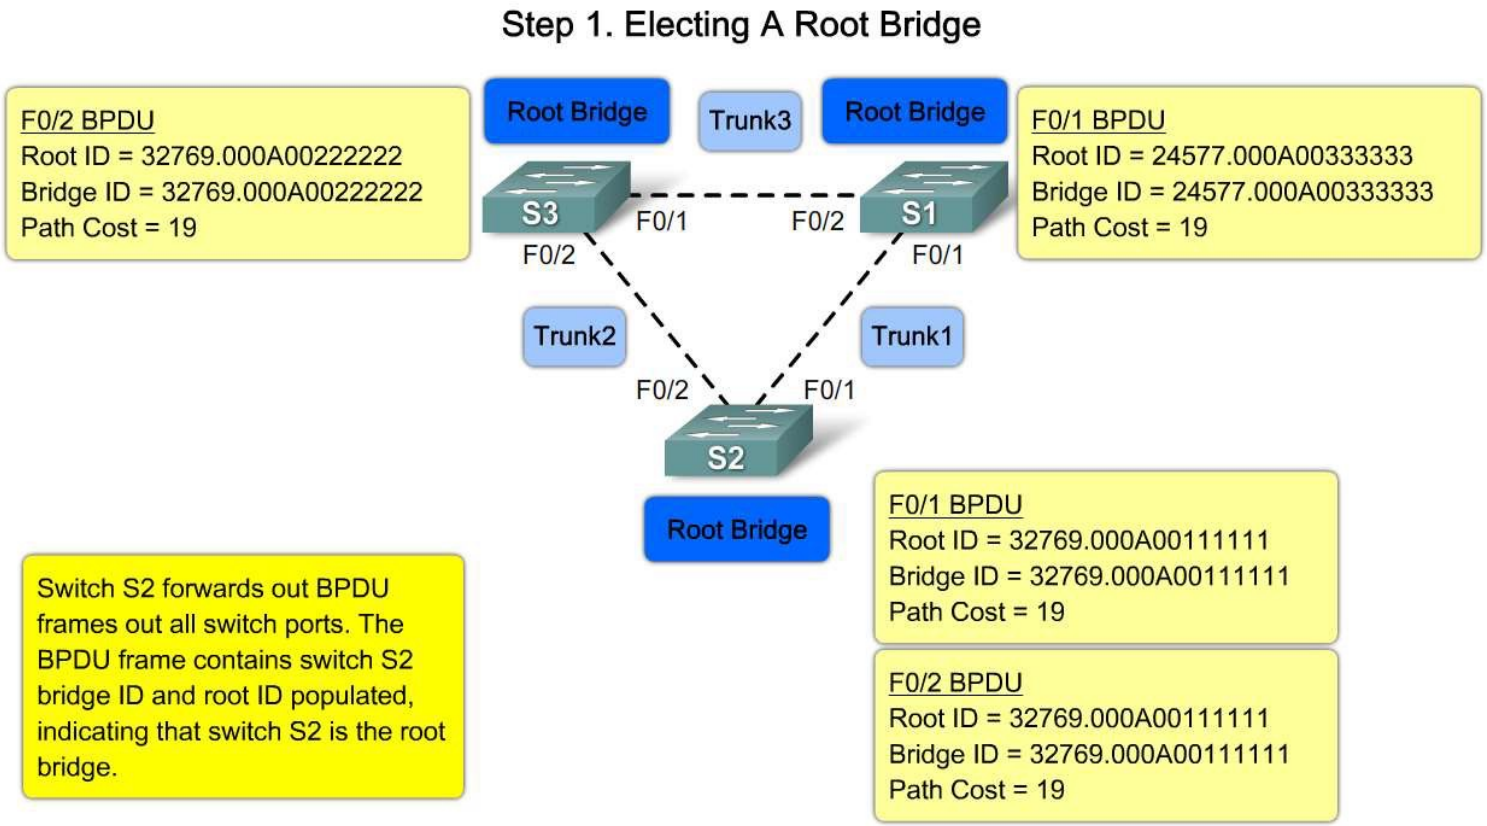

Step 1: Electing a Root Bridge

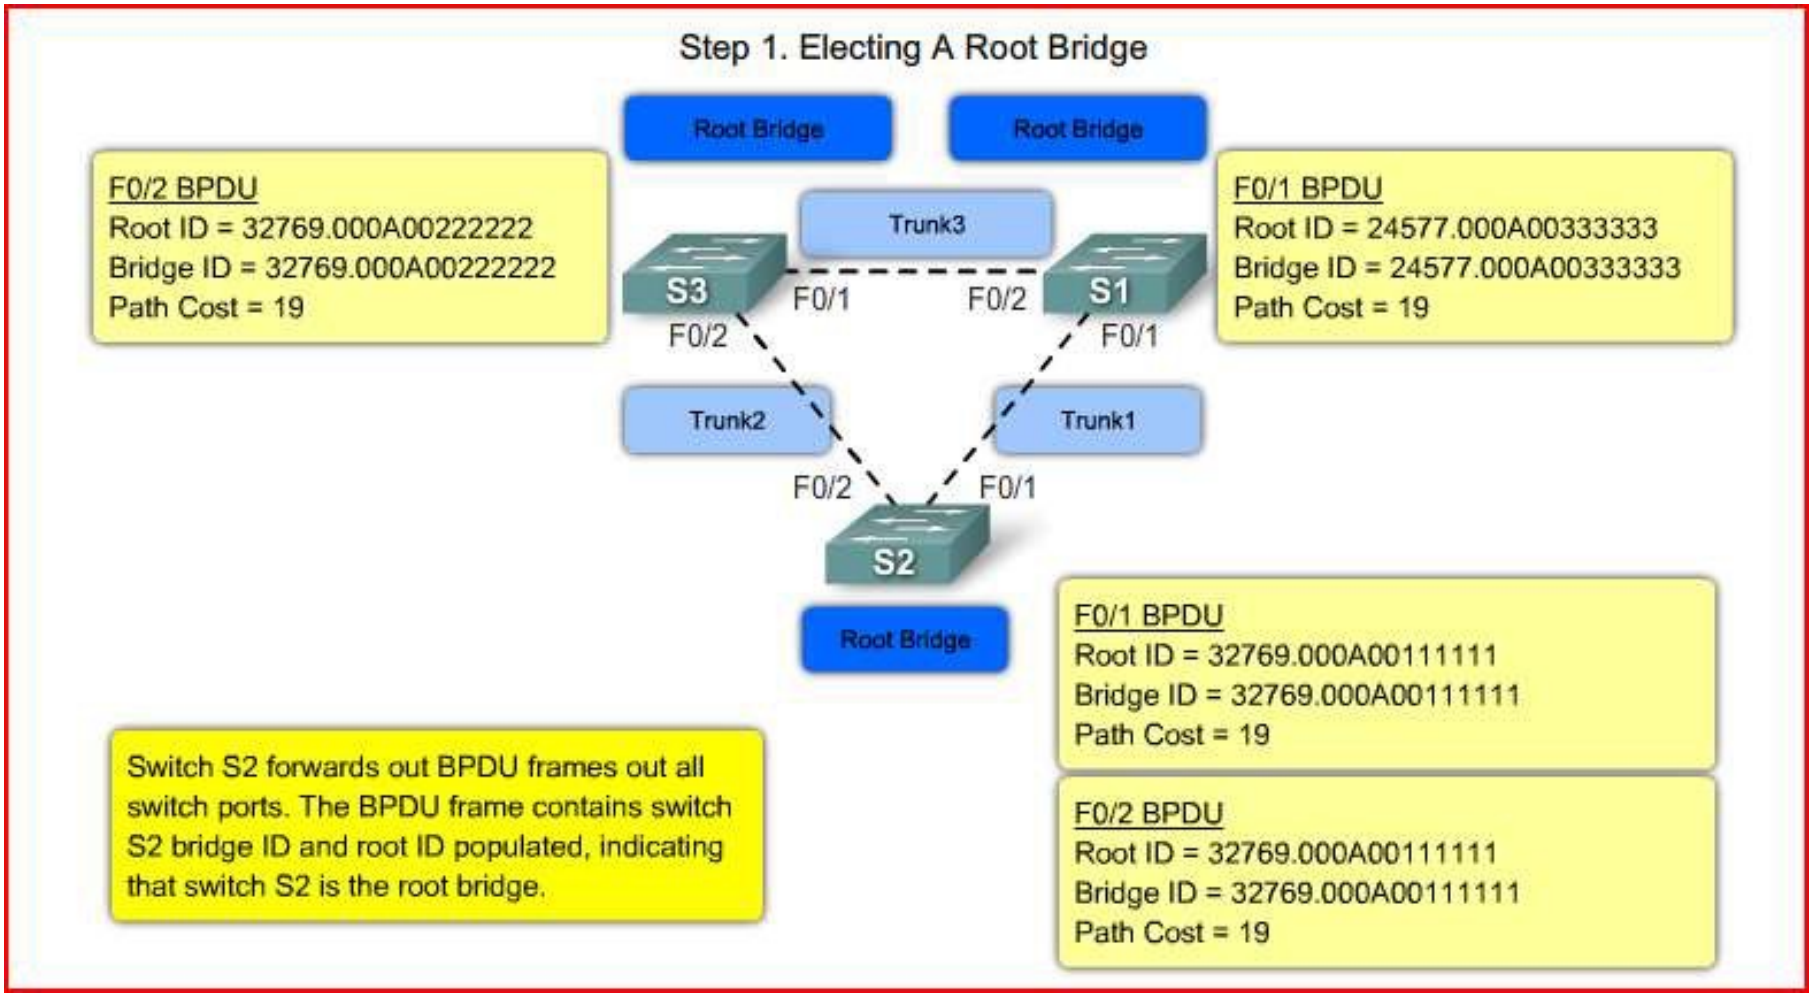

The Root Bridge

- The root bridge is the basis for all spanning-tree path cost calculations and ultimately leads to the assignment of the different port roles used to prevent layer 2 loops.
- A root bridge election is triggered after a switch has finished booting up, or when a path failure has been detected on a network.
- Initially, all switch ports are configured for the blocking state, which by default lasts 20 seconds.
- Maximum STP network diameter is seven switch hops. This allows the entire root bridge election process to occur within 14 seconds, which is less than the time the switch ports spend in the blocking state.

Bridge Protocol Data Units

- Immediately after the switches have finished booting up, they start sending BPDU frames advertising their BID in an attempt to become the root bridge.
- Initially, all switches in the network assume that they are the root bridge for the broadcast domain.
- The flood of BPDU frames on the network have the root ID field matching the BID field, indicating that each switch considers itself the root bridge. These BPDU frames are sent every 2 seconds based on the default hello timer value.
- Switches compare the root ID from the received BPDU frame with their own root ID. If the root ID from the received BPDU frame is lower than its own root ID, the root ID field is updated indicating the new best candidate for the root bridge role.

Lowest ID is the Winner

- As each switch receives the BPDU frames from its neighboring switches, they compare the root ID from the received BPDU frame with the root ID configured locally.
- If the root ID from the received BPDU frame is lower than the root ID it currently has, the root ID field is updated indicating the new best candidate for the root bridge role.

After the Election

- The switches continue to forward their BPDU frames advertising the root ID of the root bridge every 2 seconds.
- Each switch is configured with a max age timer that determines how long a switch retains the current BPDU configuration in the event it stops receiving updates from its neighboring switches.
- By default, the max age timer is set to 20 seconds.
- If a switch fails to receive 10 consecutive BPDU frames from one of its neighbors, it assumes that a logical path in the spanning tree has failed and that the BPDU information is no longer valid.
- This triggers another spanning-tree root bridge election.

Verify Root Bridge Election

```
S1#show spanning-tree

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority    24577
              Address     000A.0033.3333
              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      000A.0033.3333
              Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------
Fa0/1            Desg FWD 19        128.1    Shr
Fa0/2            Desg FWD 19        128.2    Shr

S1#
```

Explain How the STP Algorithm Uses Three Steps to Converge on a Loop-Free Topology

The process of electing a root port on a switch

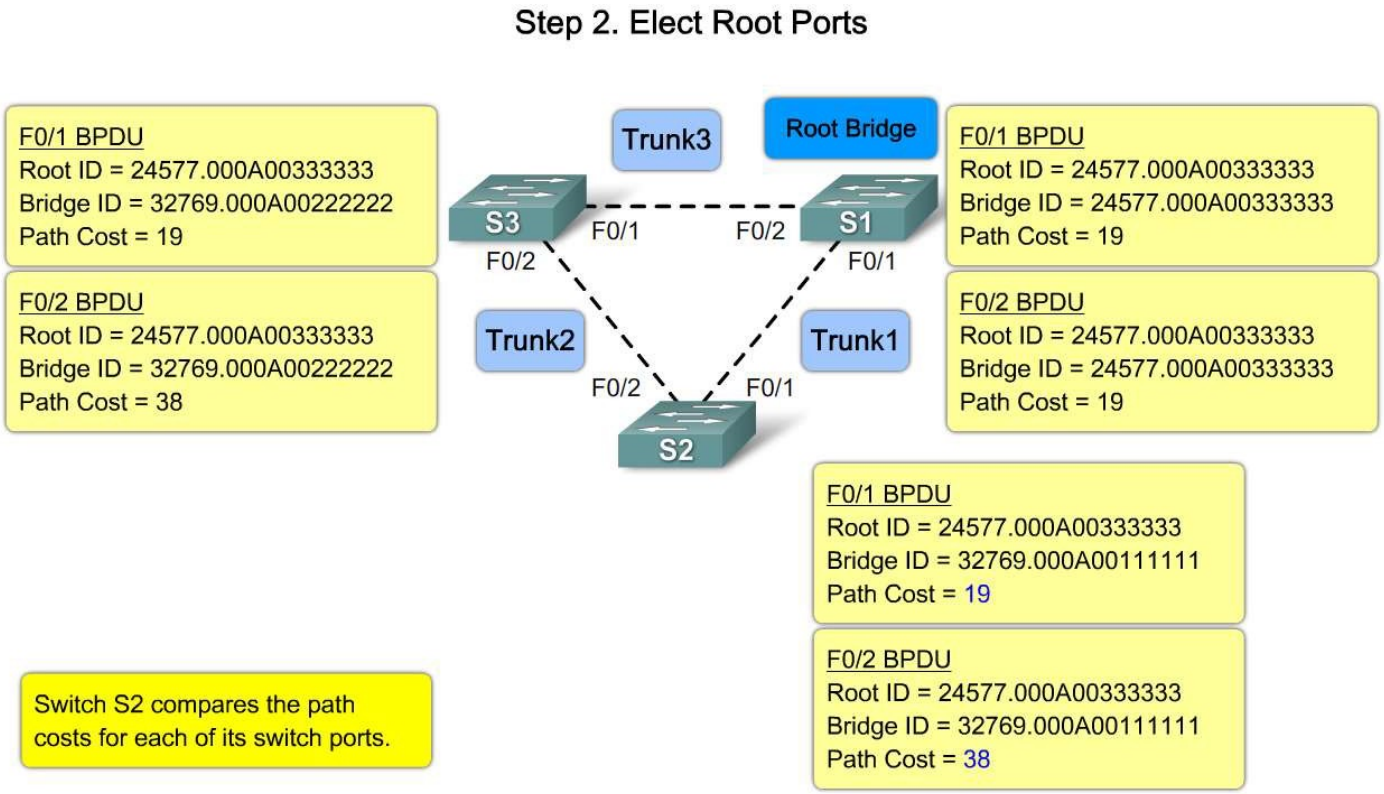

Step 2: Elect Root Ports

- Every switch in a spanning-tree topology, except for the root bridge, has a single root port defined.
- The root port is the switch port with the lowest path cost to the root bridge.
- Path cost determines which switch port becomes the root port.
- Switch ports with equivalent path costs to the root use the configurable port priority value.
- They use the port ID to break a tie. When a switch chooses one equal path cost port as a root port over another, the losing port is configured as the non-designated to avoid a loop.

Root Ports

- The process of determining which port becomes a root port happens during the root bridge election BPDU exchange.
- Path costs are updated immediately when BPDU frames arrive indicating a new root ID or redundant path.
- At the time the path cost is updated, the switch enters decision mode to determine if port configurations need to be updated.
- The port role decisions do not wait until all switches settle on which switch is going to be the final root bridge.
- As a result, the port role for a given switch port may change multiple times during convergence, until it finally settles on its final port role after the root ID changes for the last time.

Verify Root Ports

```
S2#show spanning-tree

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority    24577
              Address     000A.0033.3333
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
  Bridge ID   Priority    32769  (priority 32768 sys-id-ext 1)
              Address     000A.0011.1111
              Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------
Fa0/1            Root FWD 19        128.1    P2p
Fa0/2            Desg FWD 19        128.2    P2p

S2#
```

How does the STP Algorithm Uses Three Steps to Converge on a Loop-Free Topology

Describe the process of electing designated ports and non-designated ports on a switch

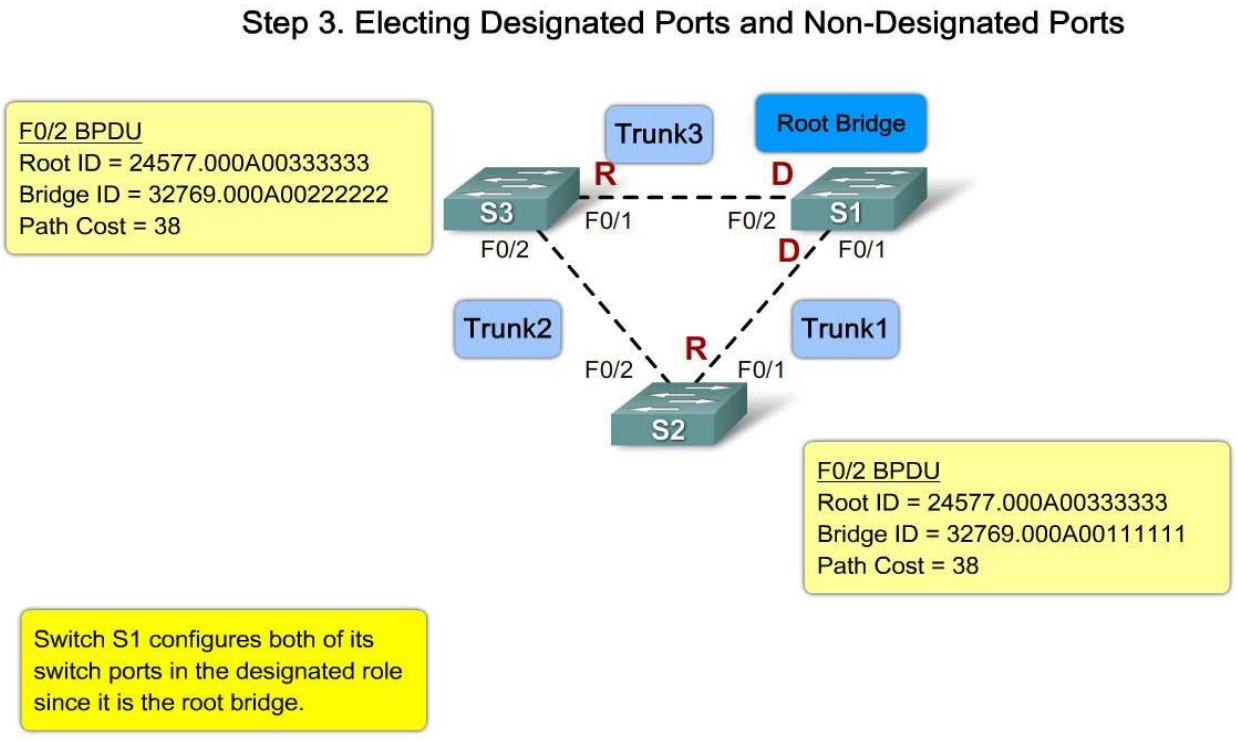

Step 3: Electing Designated and Undesignated Ports

- There can only be one designated port on each segment in a switched network.
- When two non-root port switch ports are connected on the same LAN segment, a competition for port roles occurs. The two switches exchange BPDU frames to sort out which switch port is designated and which one is non-designated.
- Generally, when a switch port is configured as a designated port, it is based on the BID.
- The first priority is the lowest path cost to the root bridge and only if the port costs are equal, is the BID considered.

Blocking State

- When two switches exchange their BPDU frames, they examine the sending BID of the received BPDU frame to see if it is lower than its own.
- The switch with the lower BID wins the competition and its port is configured in the designated role. The losing switch configures its switch port to be non-designated and, therefore, in the blocking state to prevent the loop from occurring.

Verify DP and Non-DP

```
S3#show spanning-tree

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID     Priority    24577
              Address     000A.0033.3333
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
  Bridge ID   Priority    32769  (priority 32768 sys-id-ext 1)
              Address     000A.0022.2222
              Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/1            Root FWD 19        128.1    P2p
Fa0/2            Altn BLK 19        128.2    P2p

S3#
```

STP Topology Change Notification

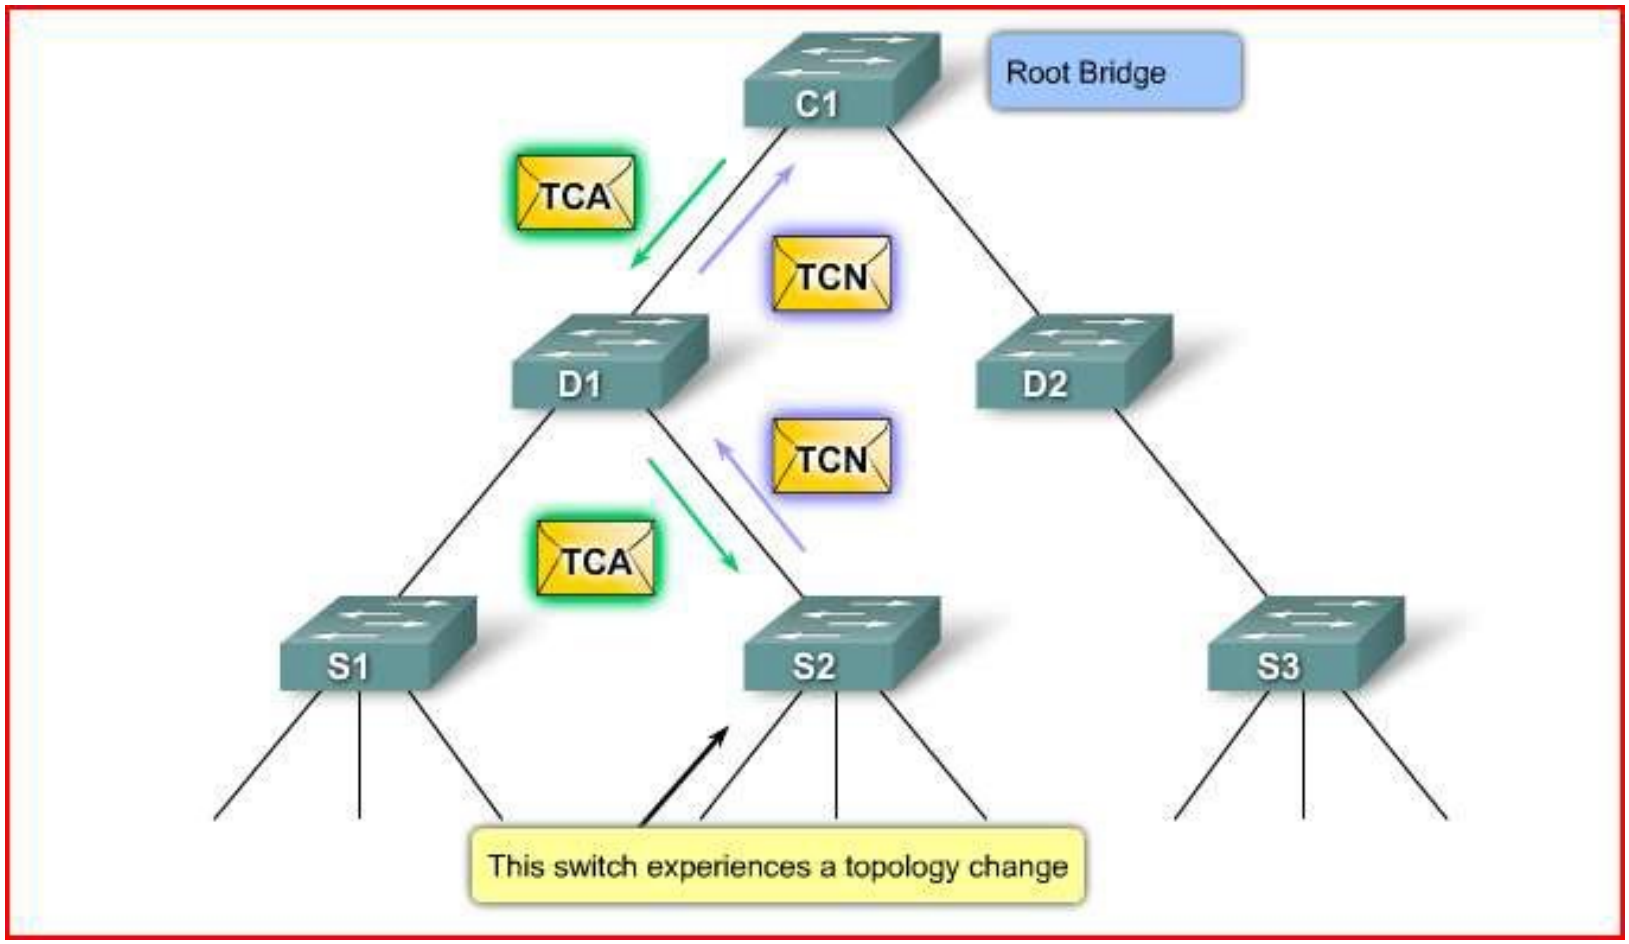

Topology Change Notification

- A switch considers it has detected a topology change either when a port that was forwarding is going down (blocking for instance) or when a port transitions to forwarding and the switch has a designated port.
- When a change is detected, the switch notifies the root bridge of the spanning tree.
- The root bridge then broadcasts the information into the whole network.

Broadcast Notification

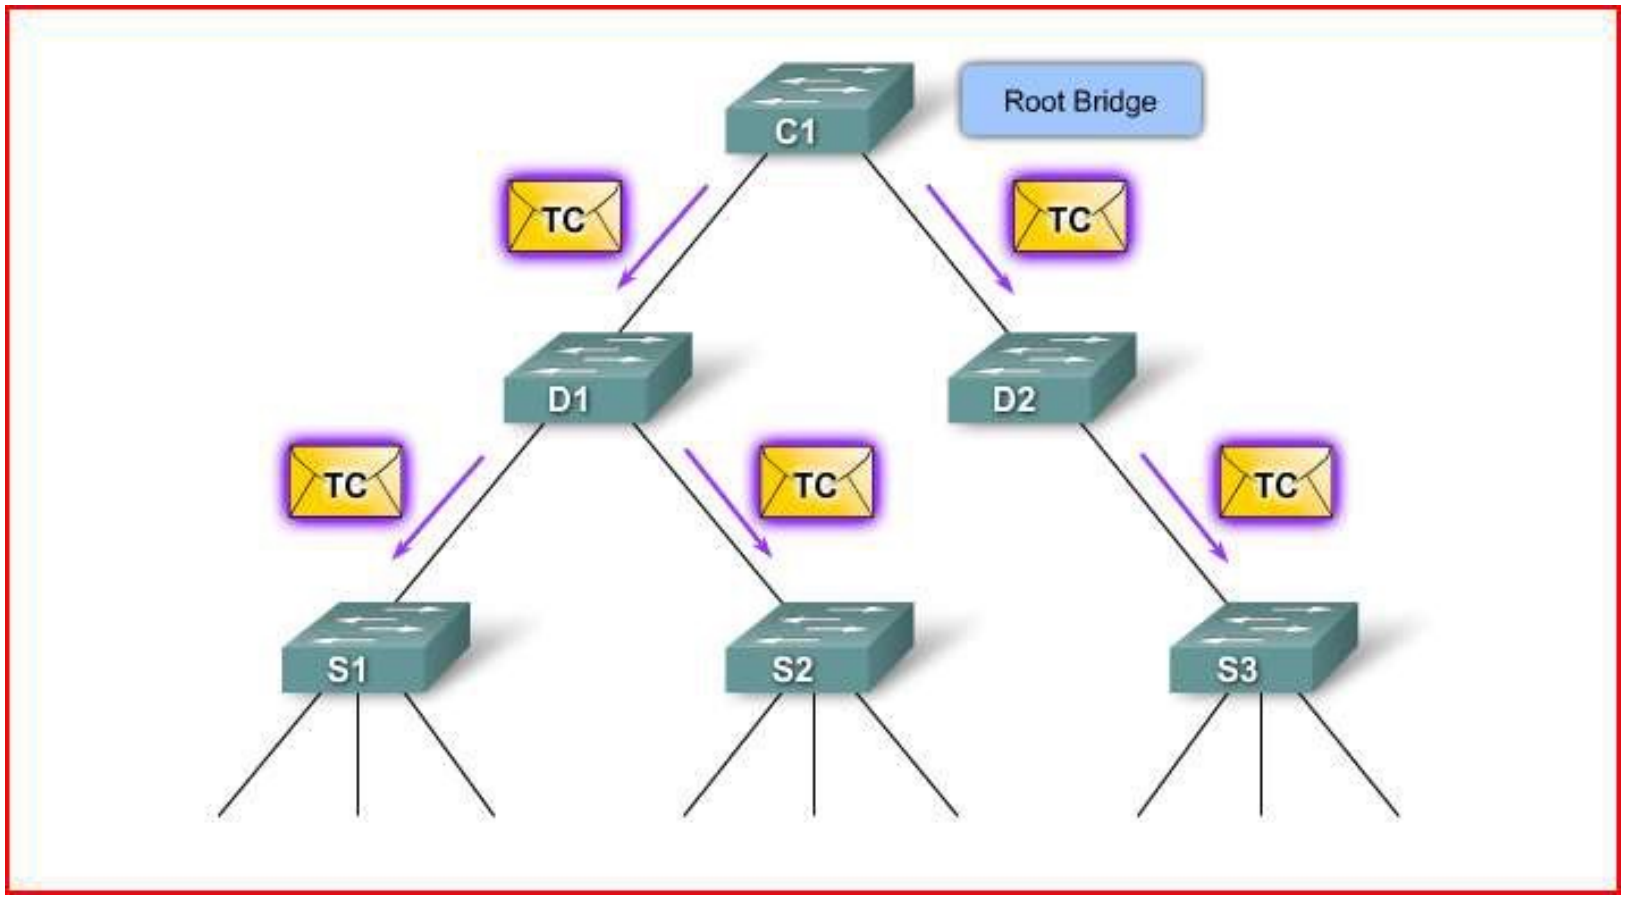

Broadcast Notification

- Once the root bridge is aware that there has been a topology change event in the network, it starts to send out its configuration BPDUs with the topology change (TC) bit set. These BPDUs are relayed by every switch in the network with this bit set. As a result, all switches become aware of the topology change and can reduce their aging time to forward delay. Switches receive topology change BPDUs on both forwarding and blocking ports.

Cisco and STP Variants

Cisco Proprietary	PVST • Uses the Cisco proprietary ISL trunking protocol • Each VLAN has an instance of spanning tree • Ability to load balance traffic at layer-2 • Includes extensions BackboneFast, UplinkFast, and PortFast
	PVST+ • Supports ISL and IEEE 802.1Q trunking • Supports Cisco proprietary STP extensions • Adds BPDU guard and Root guard enhancements
	rapid-PVST+ • Based on IEEE802.1w standard • Has faster convergence than 802.1D
IEEE Standard	RSTP • Introduced in 1982 provides faster convergence than 802.1D • Implements generic versions of the Cisco proprietary STP extensions • IEEE has incorporated RSTP into 802.1D, identifying the specification as IEEE 802.1D-2004
	MSTP • Multiple VLANs can be mapped to the same spanning-tree instance • Inspired by the Cisco Multiple Instances Spanning Tree Protocol (MISTP) • IEEE 802.1Q-2003 now includes MSTP

PVST

- Per-VLAN spanning tree protocol (PVST) - Maintains a spanning-tree instance for each VLAN configured in the network.
- It uses the Cisco proprietary ISL trunking protocol that allows a VLAN trunk to be forwarding for some VLANs while blocking for other VLANs. Because PVST treats each VLAN as a separate network, it can load balance traffic at Layer 2 by forwarding some VLANs on one trunk and other VLANs on another trunk without causing a loop.
- To learn more about these extensions, visit: http://www.cisco.com/en/US/docs/switches/lan/catalyst4000/7.4/configuration/guide/stp_enha.html.

PVST+

- Per-VLAN spanning tree protocol plus (PVST+) - Cisco developed PVST+ to provide support for IEEE 802.1Q trunking.
- PVST+ provides the same functionality as PVST, including the Cisco proprietary STP extensions. PVST+ is not supported on non-Cisco devices.
- PVST+ includes the PortFast enhancement called BPDU guard, and root guard. To learn more about BPDU guard, visit: http://www.cisco.com/en/US/tech/tk389/tk621/technologies_tech_note09186a008009482f.shtml

Implement Rapid per VLAN Spanning Tree (rapid PVST+) in a LAN

Features of the PVST+, RSTP and rapid PVST+ variants of STP

Cisco and STP Variants

Cisco Proprietary	PVST • Uses the Cisco proprietary ISL trunking protocol • Each VLAN has an instance of spanning tree • Ability to load balance traffic at layer-2 • Includes extensions BackboneFast, UplinkFast, and PortFast
	PVST+ • Supports ISL and IEEE 802.1Q trunking • Supports Cisco proprietary STP extensions • Adds BPDU guard and Root guard enhancements
	rapid-PVST+ • Based on IEEE802.1w standard • Has faster convergence than 802.1D
IEEE Standard	RSTP • Introduced in 1982 provides faster convergence than 802.1D • Implements generic versions of the Cisco proprietary STP extensions • IEEE has incorporated RSTP into 802.1D, identifying the specification as IEEE 802.1D-2004
	MSTP • Multiple VLANs can be mapped to the same spanning-tree instance • Inspired by the Cisco Multiple Instances Spanning Tree Protocol (MISTP), • IEEE 802.1Q-2003 now includes MSTP

PVST+

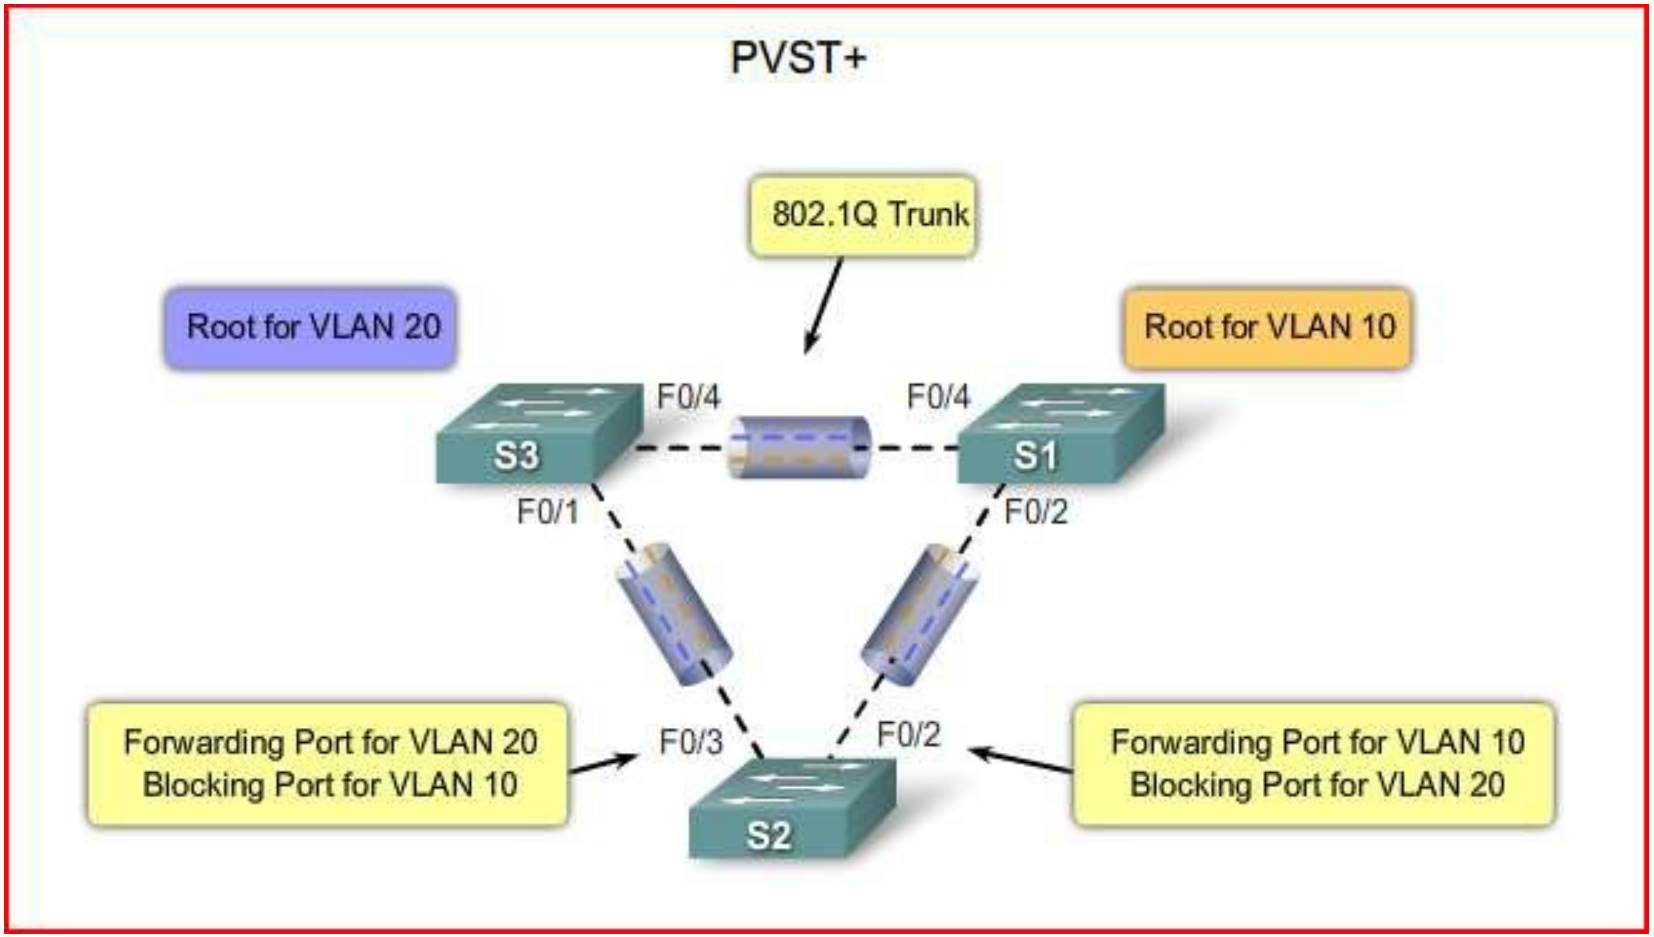
PVST+
802.1Q Trunk
Root for VLAN 20
Root for VLAN 10
F0/4
F0/4
S3
S1
F0/1
F0/2
Forwarding Port for VLAN 20
Blocking Port for VLAN 10
F0/3
F0/2
Forwarding Port for VLAN 10
Blocking Port for VLAN 20
S2

PVST+ and Bridge ID

- In the original 802.1D standard, an 8-byte BID is composed of a 2-byte bridge priority and a 6-byte MAC address of the switch.
- There was no need to identify a VLAN because there was only one spanning tree in a network.
- PVST+ requires that a separate instance of spanning tree run for each VLAN.
- To support PVST+, the 8-byte BID field is modified to carry a VLAN ID (VID).
- The bridge priority field is reduced to 4 bits and a new 12-bit field, the extended system ID field, contains the VID. The 6-byte MAC address remains unchanged.

PVST+ Bridge ID

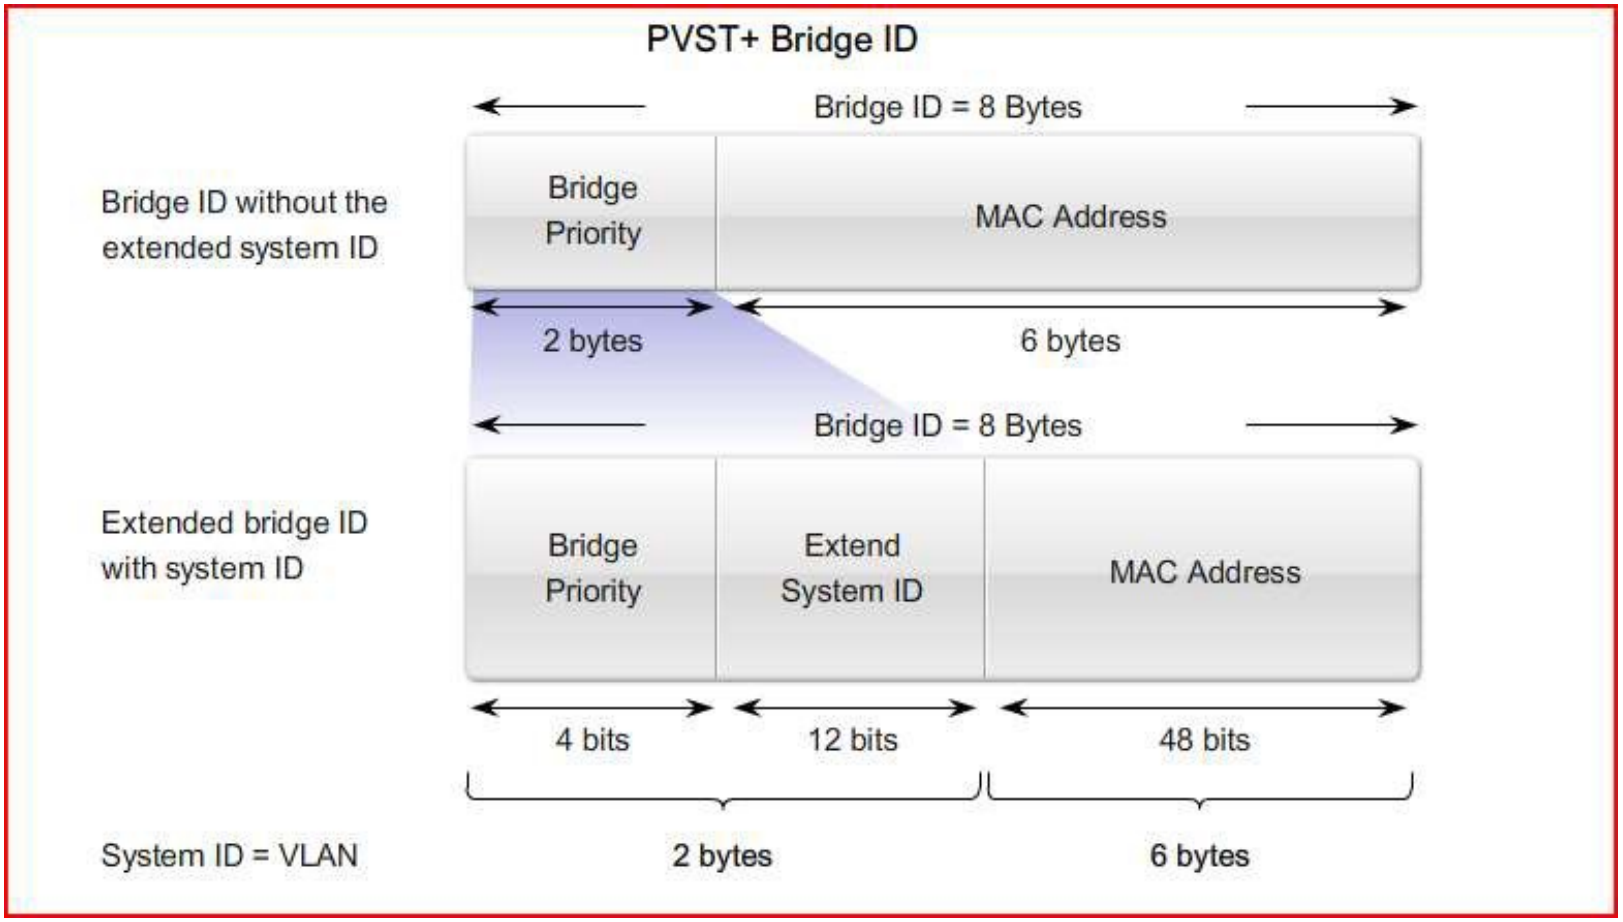

NOTE: If no priority is configured, every switch has the same default priority, and the election of the root bridge for each VLAN is based on the MAC address.

Default Settings

Default Switch Configuration

Feature	Default Setting
Enable state	Enabled on VLAN 1
Spanning-tree mode	PVST+ (Rapid PVST+ and MSTP are 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 Mb/s: 4, 100 Mb/s: 19, 10 Mb/s: 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 Mb/s: 4, 100 Mb/s: 19, 10 Mb/s: 100
Spanning-tree timers	Hello time: 2 seconds Forward-delay time: 15 seconds Maximum-aging time: 20 seconds Transmit hold count: 6 BPDUs

PVST

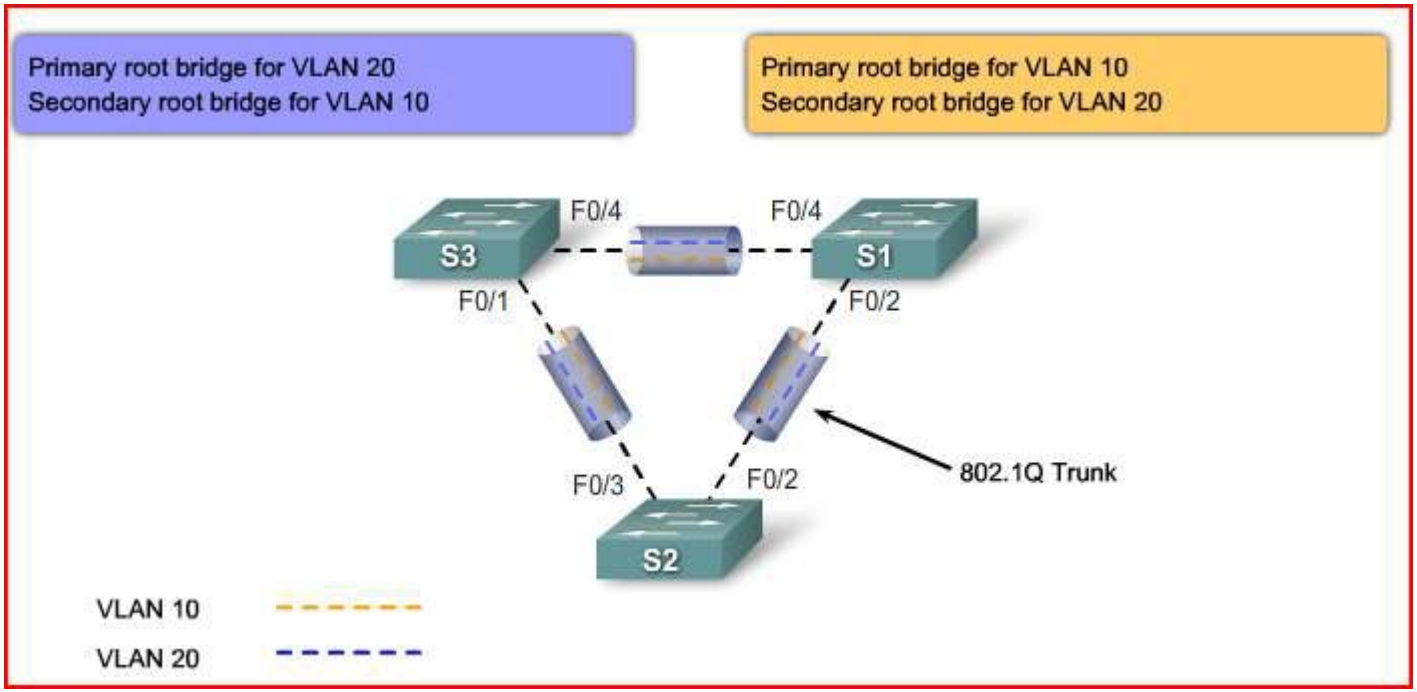
Primary root bridge for VLAN 20
Secondary root bridge for VLAN 10
Primary root bridge for VLAN 10
Secondary root bridge for VLAN 20
S3
S1
S2
F0/4
F0/4
F0/1
F0/2
F0/3
F0/2
802.1Q Trunk
VLAN 10
VLAN 20

Configuring PVST

```
S3(config)#spanning-tree vlan 20 root primary
```

This command forces switch S3 to be the primary root for VLAN 20.

```
S3(config)#spanning-tree vlan 10 root secondary
```

This command forces switch S3 to be the secondary root for VLAN 10.

```
S1(config)#spanning-tree vlan 10 root primary
```

This command forces switch S1 to be the primary root for VLAN 10.

```
S1(config)#spanning-tree vlan 20 root secondary
```

Configuring PVST Continued

```
S3(config)#spanning-tree vlan 20 priority 4096
```

This command sets the priority for switch S3 to be the lowest possible making it most likely that switch S3 will be the primary root for VLAN 20.

```
S1(config)#spanning-tree vlan 10 priority 4096
```

This command sets the priority for switch S1 to be the lowest possible making it most likely that switch S1 will be the primary root for VLAN 10.

Verifying STP

```
S1#show spanning-tree active
<output omitted>
VLAN0010
  Spanning tree enabled protocol ieee
  Root ID     Priority    4106
              Address     0019.aa9e.b000
              This bridge is the root
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
  Bridge ID   Priority    4106   (priority 4096 sys-id-ext 10)
              Address     0019.aa9e.b000
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300
Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/2            Desg FWD 19        128.2    P2p
Fa0/4            Desg FWD 19        128.4    P2p
<output omitted>
```

Verifying PVST

```
S1#show run
Building configuration...

Current configuration : 1595 bytes
!
version 12.2
<output omitted>
!
spanning-tree mode pvst
spanning-tree extend system-id
spanning-tree vlan 1 priority 24576
spanning-tree vlan 10 priority 4096
spanning-tree vlan 20 priority 28672
!
<output omitted>
```

Implement Rapid per VLAN Spanning Tree (rapid PVST+) in a LAN

Features of PVST+

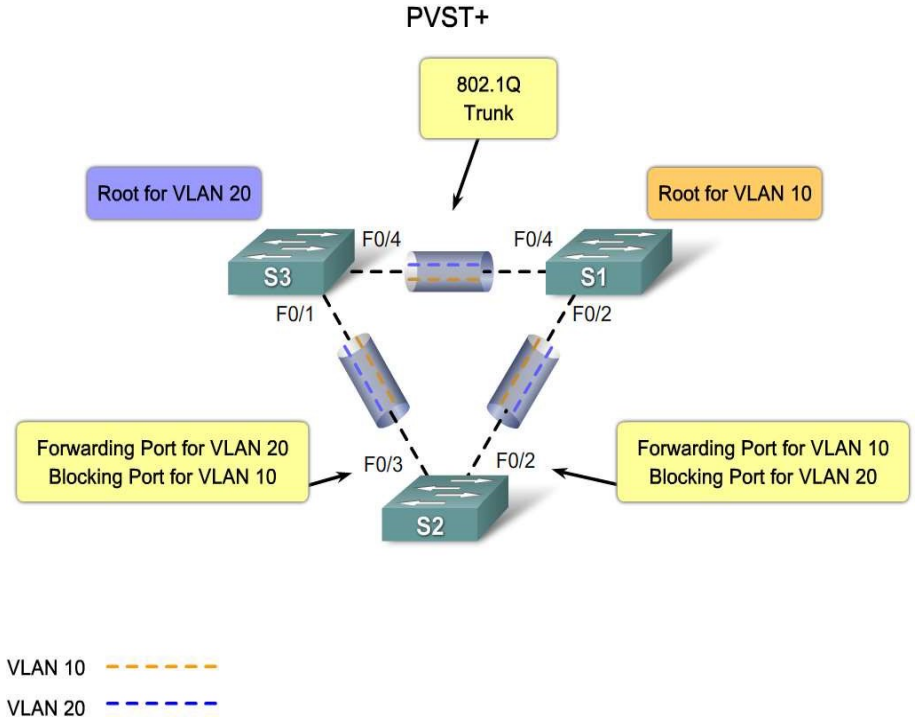

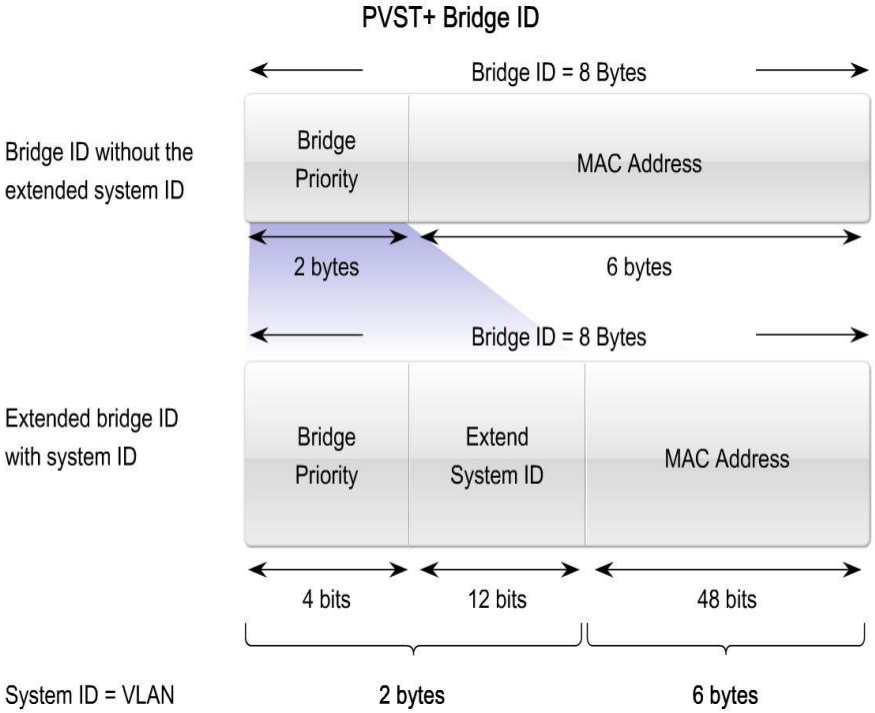

Implement Rapid per VLAN Spanning Tree (rapid PVST+) in a LAN

Features of RSTP

What is RSTP?

Characteristics of RSTP:

- Is the preferred protocol for preventing Layer 2 loops in a switched network
- Transparently integrates Cisco-proprietary enhancements
- Performs better than the Cisco-proprietary enhancements
- Not compatible with Cisco-proprietary enhancements
- Defines different port states and port roles
- Is backward compatible with 802.1D
- Has kept most configuration parameters unchanged
- Has the same BPDU format as the IEEE 802.1D BPDU
- Does not need 802.1D timers

RSTP

- RSTP speeds the recalculation of the spanning tree when the Layer 2 network topology changes.
- RSTP can achieve much faster convergence in a properly configured network, sometimes in as little as a few hundred milliseconds.
- RSTP redefines the type of ports and their state.
- There is a difference between the role of the port and the state of the port.
- If a port is configured to be an alternate or a backup port it can immediately change to a forwarding state without waiting for the network to converge.

RSTP BPDU

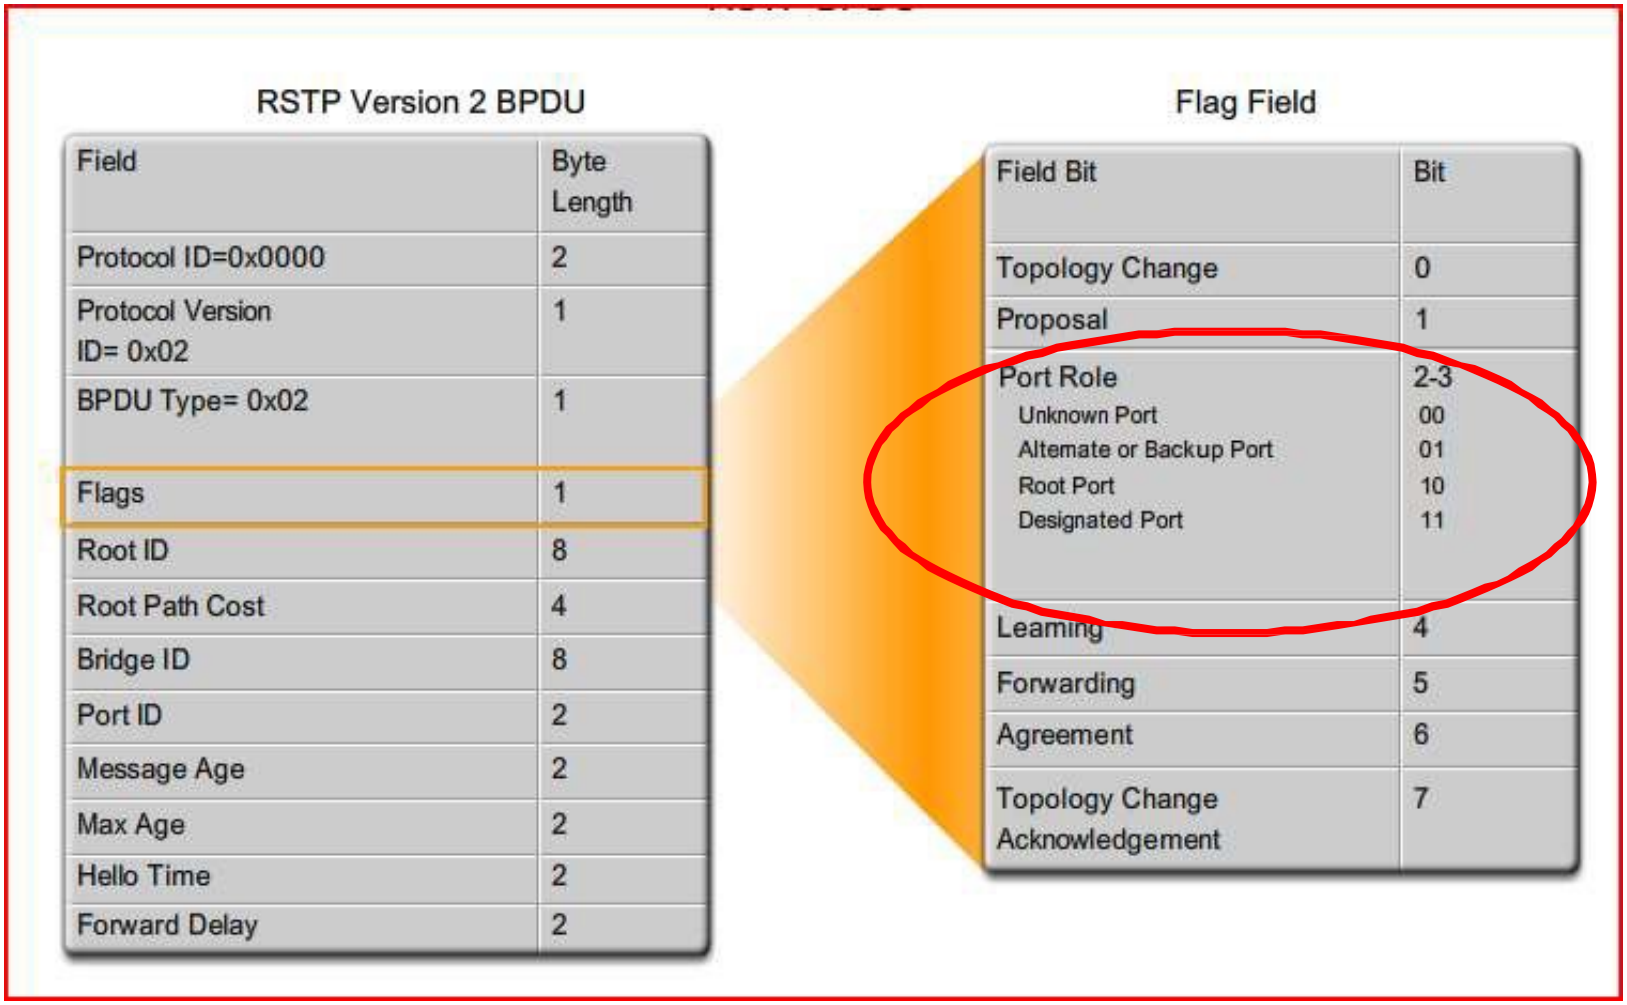

RSTP Version 2 BPDU

Field	Byte Length
Protocol ID=0x0000	2
Protocol Version ID= 0x02	1
BPDU Type= 0x02	1
Flags	1
Root ID	8
Root Path Cost	4
Bridge ID	8
Port ID	2
Message Age	2
Max Age	2
Hello Time	2
Forward Delay	2

Flag Field

Field Bit	Bit
Topology Change	0
Proposal	1
Port Role Unknown Port Alternate or Backup Port Root Port Designated Port	2-3 00 01 10 11
Learning	4
Forwarding	5
Agreement	6
Topology Change Acknowledgement	7

RSTP and BPDUs

- Protocol information can be immediately aged on a port if hellos are not received for three consecutive hello times, 6 seconds by default, or if the max age timer expires.
- Because BPDUs are used as a keepalive mechanism, three consecutively missed BPDUs indicate lost connectivity between a bridge and its neighboring root or designated bridge.
- The fast aging of the information allows failures to be detected quickly.

Implement Rapid per VLAN Spanning Tree (rapid PVST+) in a LAN

- Describe RSTP edge ports

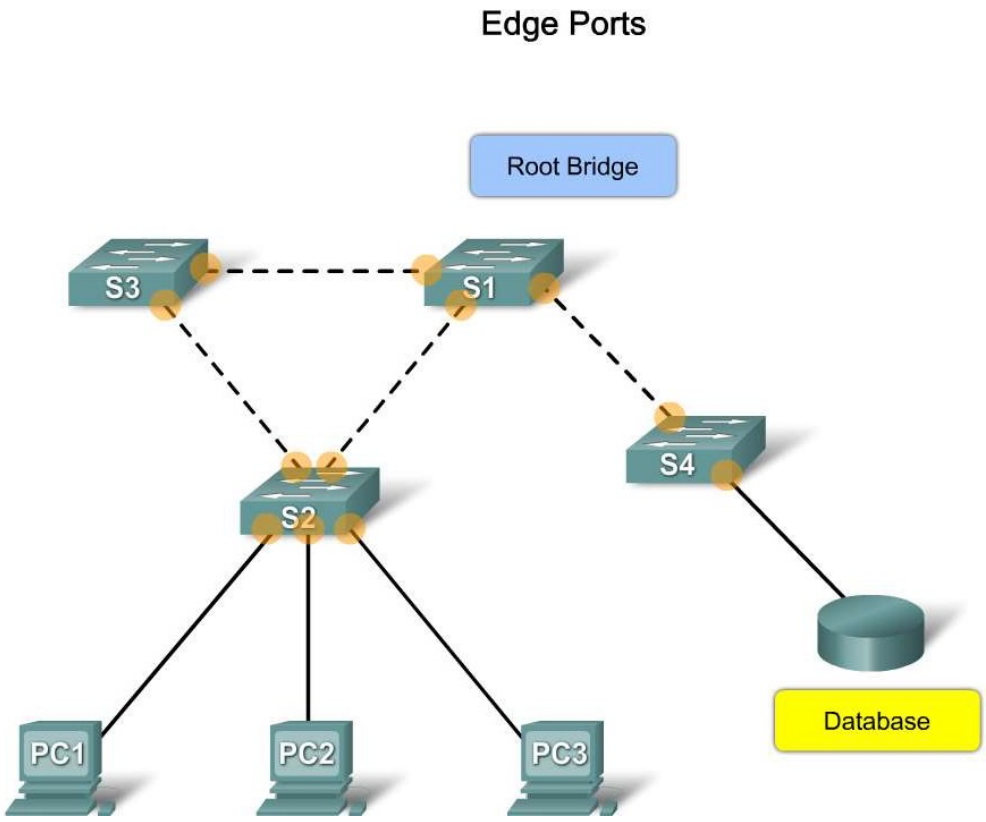

RSTP Edge Ports

- An RSTP edge port is a switch port that is never intended to be connected to another switch device. It immediately transitions to the forwarding state when enabled.
- The edge port concept is well known to Cisco spanning-tree users, because it corresponds to the PortFast feature in which all ports directly connected to end stations anticipate that no switch device is connected to them.
- The PortFast ports immediately transition to the STP forwarding state, thereby skipping the time-consuming listening and learning stages.
- Neither edge ports nor PortFast-enabled ports generate topology changes when the port transitions to a disabled or enabled status.

Edge Ports 2

- Unlike PortFast, an RSTP edge port that receives a BPDU loses its edge port status immediately and becomes a normal spanning-tree port.
- The Cisco RSTP implementation maintains the PortFast keyword using the spanning-tree portfast command for edge port configuration.
- This makes an overall network transition to RSTP more seamless.
- Configuring an edge port to be attached to another switch can have negative implications for RSTP when it is in sync state because a temporary loop can result, possibly delaying the convergence of RSTP due to BPDU contention with loop traffic.

RSTP Port Types

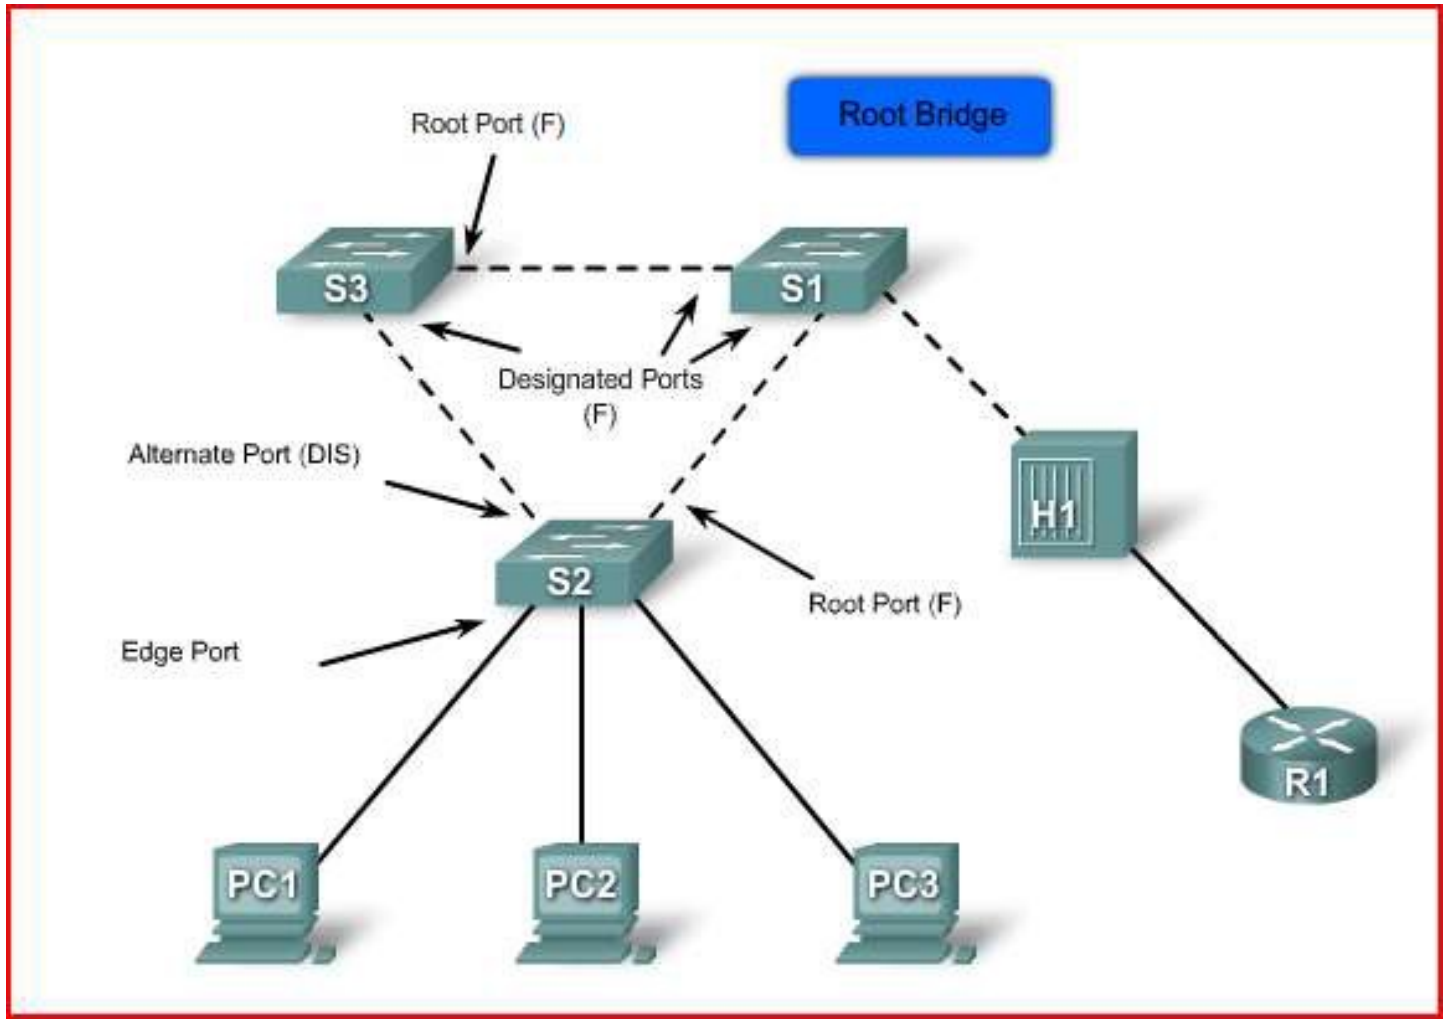

RSTP Link Types

- The link type provides a categorization for each port participating in RSTP.
- The link type can predetermine the active role that the port plays as it stands by for immediate transition to forwarding state if certain conditions are met.
- These conditions are different for edge ports and non-edge ports.
- Non-edge ports are categorized into two link types, point-to-point and shared. The link type is automatically determined, but can be overwritten with an explicit port configuration.

Link Type Parameter

- Edge ports, the equivalent of PortFast-enabled ports, and point-to-point links are candidates for rapid transition to a forwarding state.
- Before the link type parameter is considered, RSTP must determine the port role.
- Root ports do not use the link type parameter. Root ports are able to make a rapid transition to the forwarding state as soon as the port is in sync.
- Alternate and backup ports do not use the link type parameter in most cases.
- Designated ports make the most use of the link type parameter. Rapid transition to the forwarding state for the designated port occurs only if the link type parameter indicates a point-to-point link.

Implement Rapid per VLAN Spanning Tree (rapid PVST+) in a LAN

- Describe the RSTP link types

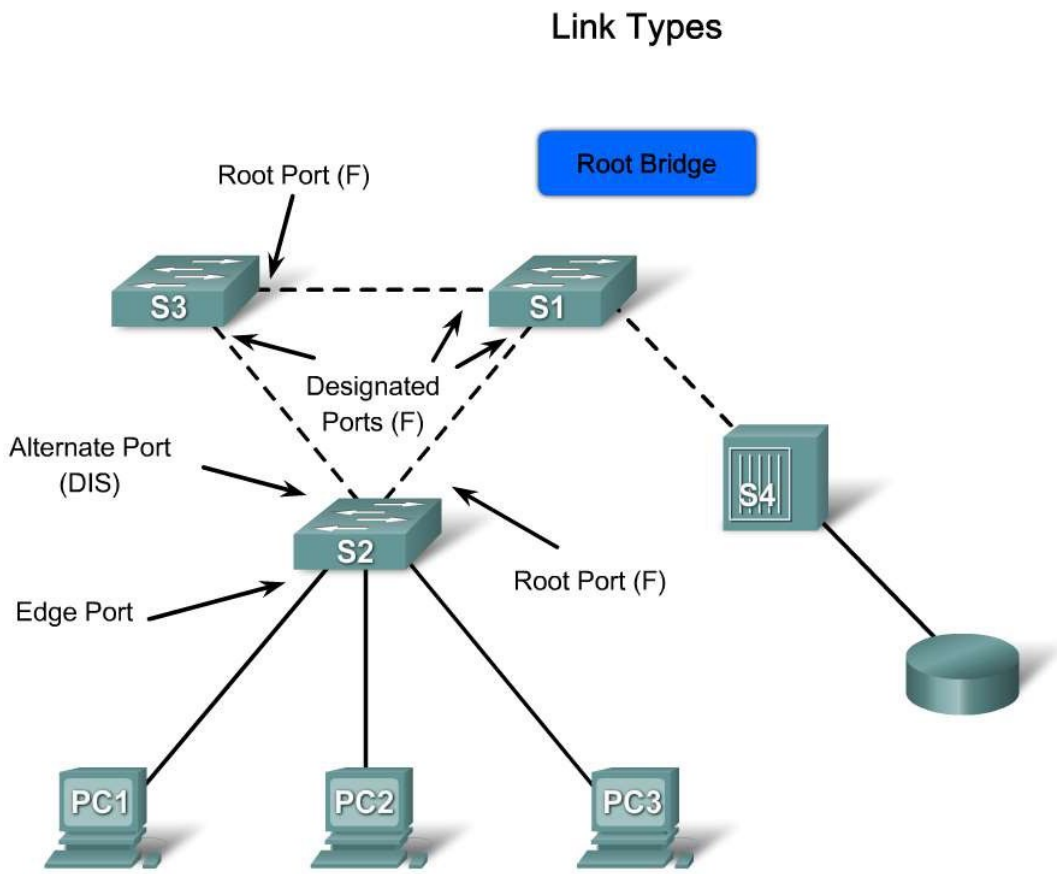

Implement Rapid per VLAN Spanning Tree (rapid PVST+) in a LAN

- Describe the RSTP port states and port roles

RSTP Port States

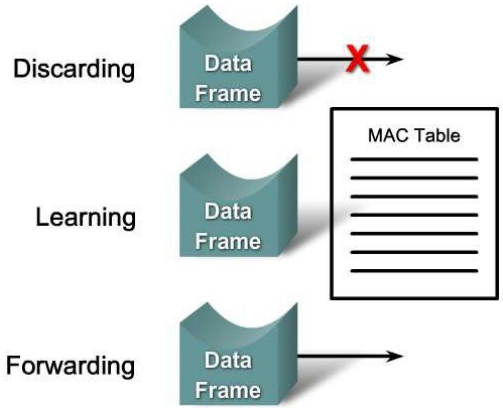

RSTP Port States

Port State	Action
Discarding	This state is seen in both a stable active topology and during topology synchronization and changes. The discarding state prevents the forwarding of data frames, thus "breaking" the continuity of a layer 2 loop.
Learning	This state is seen in both a stable active topology and during topology synchronization and changes. The learning state accepts data frames to populate the MAC table in an effort to limit flooding of unknown unicast frames.
Forwarding	This state is seen only in stable active topologies. The forwarding switch ports determine the topology. Following a topology change, or during synchronization, the forwarding of data frames occurs only after a proposal and agreement process.

RSTP Port States

- RSTP provides rapid convergence following a failure or during re-establishment of a switch, switch port, or link.
- An RSTP topology change causes a transition in the appropriate switch ports to the forwarding state through either explicit handshakes or a proposal and agreement process and synchronization.
- You will learn more about the proposal and agreement process later.
- With RSTP, the role of a port is separated from the state of a port. For example, a designated port could be in the discarding state temporarily, even though its final state is to be forwarding.
- The three possible RSTP port states: discarding, learning, and forwarding.

Port States

Operational Port State	STP Port State	RSTP Port State
Enabled	Blocking	Discarding
Enabled	Listening	Discarding
Enabled	Learning	Learning
Enabled	Forwarding	Forwarding
Disabled	Disabled	Discarding

RSTP Port States

- Alternate and Blocking
  - These two port roles correspond to the blocking state of 802.1D.
- A blocked port is defined as not being the designated or root port.
- A blocked port receives a more useful BPDU than the one it sends out on its segment.
- Remember that a port absolutely needs to receive BPDUs in order to stay blocked. RSTP introduces these two roles for this purpose.

Alternate Port

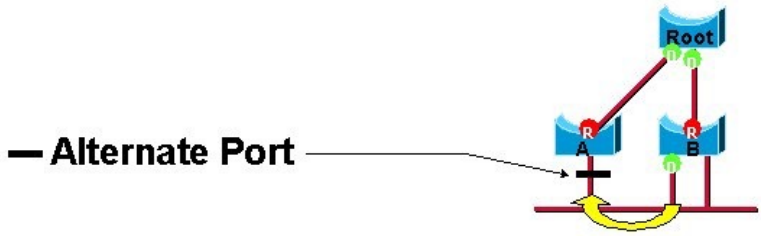

- An alternate port receives more useful BPDUs from another bridge and is a port blocked. This is shown in this diagram:

Backup Port

- An alternate port provides an alternate path to the root bridge and therefore can replace the root port if it fails.
- A backup port provides redundant connectivity to the same segment and cannot guarantee an alternate connectivity to the root bridge.

Implement Rapid per VLAN Spanning Tree (rapid PVST+) in a LAN

Configure rapid-PVST+

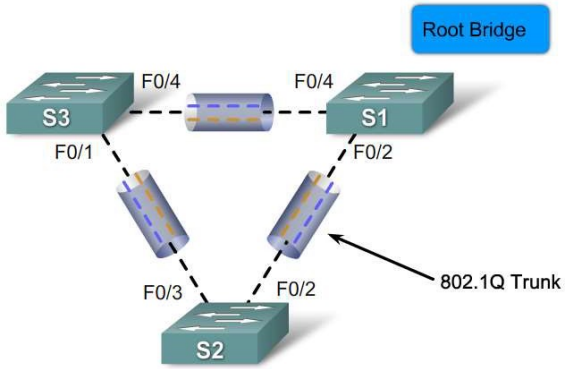

Cisco IOS Command Syntax	
Enter global configuration mode.	`configure terminal`
Configure rapid PVST+ spanning-tree mode.	`spanning-tree mode rapid-pvst`
Specify an interface to configure, and enter interface configuration mode. The VLAN ID range is 1 to 4094. The port-channel range is 1 to 6.	`interface`
Specify that the link type for this port is point-to-point.	`spanning-tree link-type point-to-point`
Return to privileged EXEC mode.	`end`
Clear all detected STP.	`clear spanning-tree detected-protocols`

Configuring RSTP

```
S1#configure terminal
S1(config)#spanning-tree mode rapid-pvst
S1(config)#interface f0/2
S1(config-if)#spanning-tree link-type point-to-point
S1(config-if)#end
S1#clear spanning-tree detected-protocols
```

Verifying RSTP

```
S1# show spanning-tree vlan 10

VLAN0010
  Spanning tree enabled protocol rstp
  Root ID     Priority    4106
              Address     0019.aa9e.b000
              This bridge is the root
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
  Bridge ID   Priority    4106   (priority 4096 sys-id-ext 10)
              Address     0019.aa9e.b000
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time 300
Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/2            Desg LRN 19        128.2    P2p
Fa0/4            Desg LRN 19        128.4    P2p
<output truncated>

S1#
```

Verifying RSTP

```
S1# show run

<Output omitted>

spanning-tree mode rapid-pvst
spanning-tree extend system-id
spanning-tree vlan 1 priority 24576
spanning-tree vlan 10 priority 4096
spanning-tree vlan 20 priority 28672
!
<output omitted>

S1#
```

Layer 3 Switching

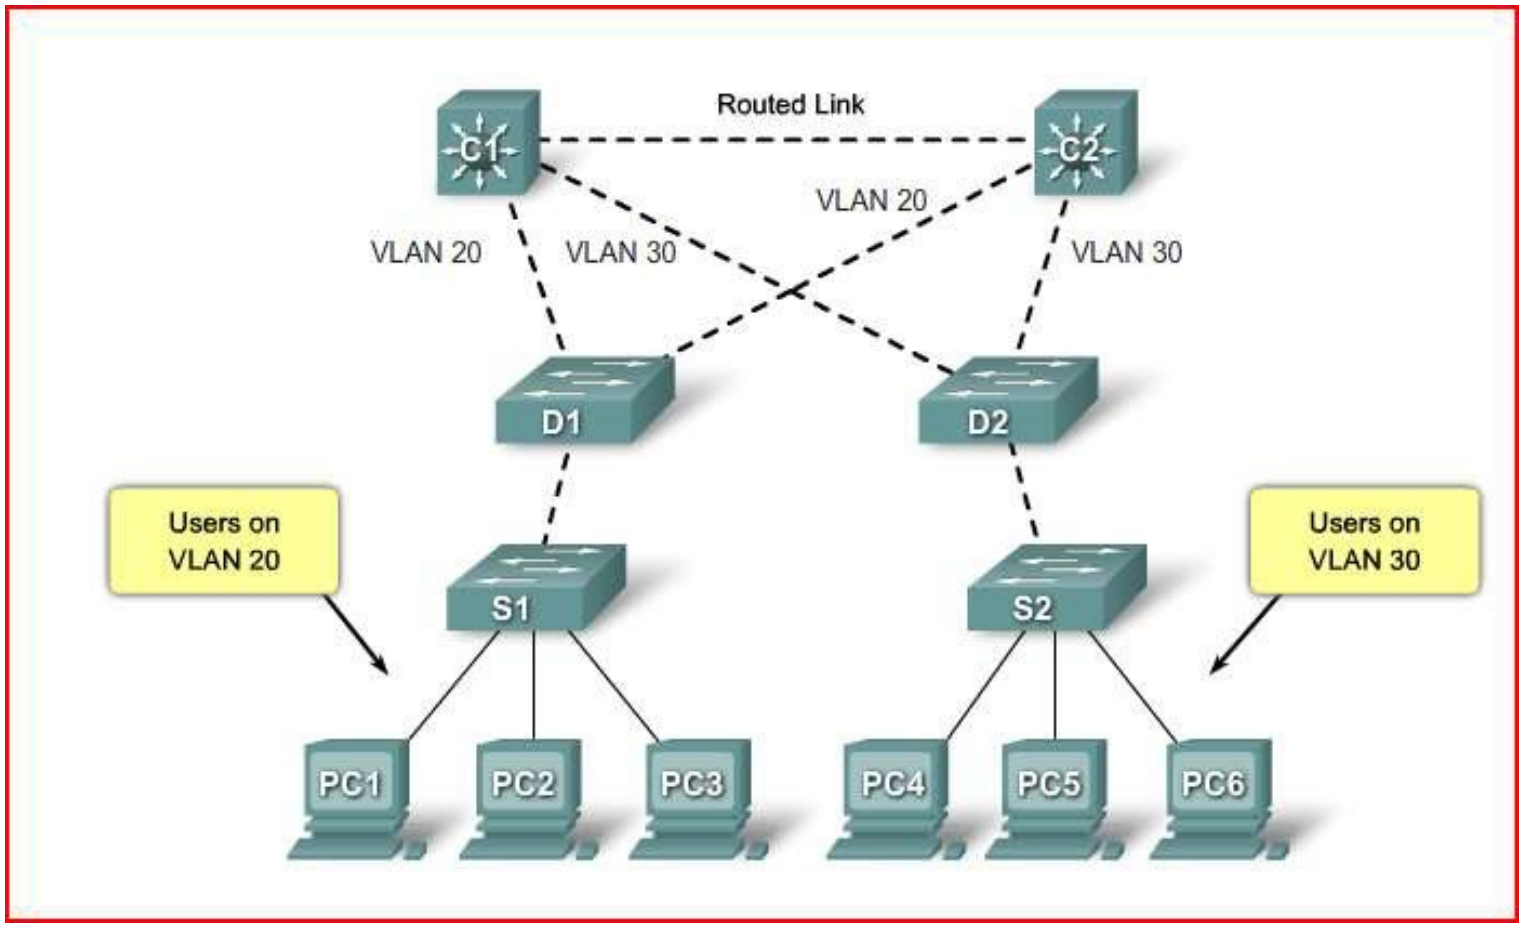

Layer 3 Switching

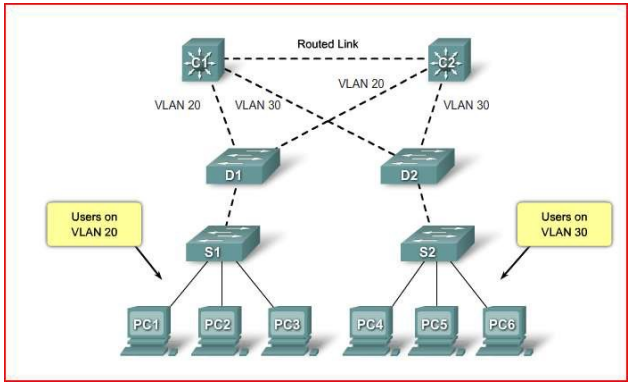

- Layer 3 switching means routing approximately at the speed of switching. A router performs two main functions:
- It builds a forwarding table. The router generally exchanges information with peers by way of routing protocols.
- It receives packets and forwards them to the correct interface based on the destination address.
- High-end Cisco Layer 3 switches are now able to perform this second function, at the same speed as the Layer 2 switching function. In the figure:
- There is no speed penalty with the routing hop and an additional segment between C1 and C2.
- Core switch C1 and core switch C2 are Layer 3 switches. VLAN 20 and VLAN 30 are no longer bridged between C1 and C2, so there is no possibility for a loop.

Implement Rapid per VLAN Spanning Tree (rapid PVST+) in a LAN

- Need to design STP to avoid problems

Final Points

Keep STP Even If It Is Unnecessary

- Do not disable STP.
- STP is not very processor-intensive
- the few BPDUs sent on each link do not reduce bandwidth.
- But a bridge network without STP can go down in a fraction of a second

Keep Traffic off the Administrative VLAN

- A high rate of broadcast or multicast traffic on the administrative VLAN adversely effects the CPU's ability to process vital BPDUs.
- Keep user traffic off the administrative VLAN.

Do Not Have a Single VLAN Span the Entire Network

- VLAN 1 serves as an administrative VLAN, where all switches are accessible in the same IP subnet.
- A bridging loop on VLAN 1 affects all trunks and can bring down the network.
- Segment the bridging domains using high-speed Layer 3 switches.

Implement Rapid per VLAN Spanning Tree (rapid PVST+) in a LAN

- Describe how to identify and solve the key STP configuration issues

Troubleshoot a Failure

To troubleshoot a bridging loop, you need to know:

- The topology of the bridge network
- The location of the root bridge
- The location of the blocked ports and the redundant links

Summary

- Spanning Tree Protocol (STP) is used to prevent loops from being formed on redundant networks
- STP uses different port states & timers to logically prevent loops
- There is at least one switch in a network that serves as the root bridge
  - –Root bridge is elected using information found in BPDU frames
- Root ports are determined by the spanning tree algorithm and are closest to the root bridge

Summary

- STP lengthy convergence time (50 seconds) facilitated the development of:
  - RSTP
    - convergence time is slightly over 6 seconds
  - Rapid PVST+
    - adds VLAN support to RSTP is the preferred spanning-tree protocol on a Cisco switch network