CCNA Exam Prep: Back to Networking Basics with Hank Preston and Patrick Gargano

Spanning-Tree Protocol (STP) Unleashed!

Exploring how STP works to break Layer 2 loops in a redundant switched topology.

- Investigate the starting default STP results
- Review the STP timers
- Enable RSTP
- Enable PortFast and BPDU Guard
- Enable Root Guard
- Enable Loop Guard

Setup and Scenario

In this set of lab-based demonstrations, you are the network engineer for a growing organization tasked with updating the network to support new network needs. The network traditionally only had two switches, but to provide better performance and redundancy, two more switches were added following the recommended enterprise network model of a distribution layer and an access layer. Unfortunately, users are now complaining of connectivity and access issues.

The NOC team is asking you to investigate what the issue is and correct it.

The following lab-based demonstrations will look at how STP operates by default and how to improve its performance and security.

*Be sure to **START** the lab before continuing to the demo labs.*

Part 1 - Investigate the default STP results

Since the network has grown from two switches to four, and extra links were added to provide redundancy, users are complaining of poor performance and connectivity issues. Start by investigating the current state of the STP for VLANs 10 and 20.

Step 1

Issue the `show spanning-tree vlan 10` command on all four switches (DLS1, DLS2, ALS1 and ALS2) and identify which switch is the root bridge for VLAN 10. Repeat the same command but for VLAN 20.

> Note: To ensure the correct placement of the root bridge at the start of this lab, and to ensure that the lab steps can be successfully replicated, the following STP bridge priority values have been statically assigned:
>
> - DLS1 Bridge Priority: 24576
> - DLS2 Bridge Priority: 16384
> - ALS1 Bridge Priority: 20480
> - ALS2 Bridge Priority: 12288
>
> In a "real network", the default priority of 32768 would be in place on all switches.

ALS2 is the root bridge for both VLAN 10 and VLAN 20 because it has the lowest bridge priority. When the bridge priority is tied, STP uses the MAC address as the tie-breaker. Since it is difficult to predict and control the MAC address of a switch in virtual lab environments, the bridge priority was used in this lab to simulate the suboptimal root bridge election.

```
ALS2# show spanning-tree vlan 10

    VLAN0010
      Spanning tree enabled protocol ieee
      Root ID    Priority    12298
                 Address     5254.0005.130e
                 This bridge is the root
                 Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

      Bridge ID  Priority    12298  (priority 12288 sys-id-ext 10)
                 Address     5254.0005.130e
                 Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
                 Aging Time  300 sec

    Interface           Role Sts Cost      Prio.Nbr Type
    ------------------- ---- --- --------- -------- --------------------------------
    Gi0/1               Desg FWD 4         128.2    P2p
    Gi0/2               Desg FWD 4         128.3    P2p
    Gi1/0               Desg FWD 4         128.5    P2p
```

Step 2

From the `show spanning-tree vlan 10` output, investigate the port roles and port states on all four switches.

Since ALS2 is the root bridge, it has all of its connected ports in a designated (forwarding) state for VLAN 10 and VLAN 20

Because ALS1, DLS1, and DLS2 are not root bridges, only one port must be elected root on each of these three switches. The root port is the port with the lowest cost to the root bridge. As DLS2 has a lower BID than DLS1 and ALS1, all other ports on DLS2 are set to designated. Other ports on ALS1 and DLS1 are nondesignated (Altn), except for ports connected to hosts, and for DLS1 G0/1 which is designated for the link between DLS1 and ALS1.

> Note: The Cisco proprietary protocol PVST+ uses the term "alternate" for nondesignated ports. The blocked port is also known as the alternate port with RSTP. It is an alternate path to the root, which is less desirable. The port is blocking, where incoming frames are dropped.

Part 2 - Investigate STP timer values

Before optimizing the current STP setup, start by investigating the current STP timers and how they impact convergence time.

Step 1

Turn on the STP topology events debugging on DLS1

```
DLS1# debug spanning-tree events
```

You have turned on this debug to observe STP convergence, which is how much time STP needs to establish a new path after a link failure in real time.

Step 2

Shut down the forwarding root port on DLS1, G0/2, and observe how long it takes STP to notice the failure and make a redundant link forwarding. Notice in the debug output that the G1/0 port transitions to listening for VLANs 1, 10, 20. The port will then transition to learning for VLANs 1, 10, 20 and then to forwarding for all three VLANs. The port changes for VLAN 10 are highlighted in the output below.

```
DLS1(config)# interface G0/2
DLS1(config-if)# shutdown

*Sep 13 14:56:08.059: STP: VLAN0001 new root port Gi1/0, cost 8
*Sep 13 14:56:08.059: STP: VLAN0001 Gi1/0 -> listening
*Sep 13 14:56:08.059: STP[1]: Generating TC trap for port GigabitEthernet0/2
*Sep 13 14:56:08.059: STP: VLAN0010 new root port Gi1/0, cost 8
*Sep 13 14:56:08.060: STP: VLAN0010 Gi1/0 -> listening
*Sep 13 14:56:08.060: STP[10]: Generating TC trap for port GigabitEthernet0/2
*Sep 13 14:56:08.060: STP: VLAN0020 new root port Gi1/0, cost 8
*Sep 13 14:56:08.060: STP: VLAN0020 Gi1/0 -> listening
*Sep 13 14:56:08.060: STP[20]: Generating TC trap for port GigabitEthernet0/2
*Sep 13 14:56:10.060: STP: VLAN0001 sent Topology Change Notice on Gi1/0
*Sep 13 14:56:10.063: STP: VLAN0010 sent Topology Change Notice on Gi1/0
*Sep 13 14:56:10.065: STP: VLAN0020 sent Topology Change Notice on Gi1/0
*Sep 13 14:56:23.060: STP: VLAN0001 Gi1/0 -> learning
*Sep 13 14:56:23.061: STP: VLAN0010 Gi1/0 -> learning
*Sep 13 14:56:23.061: STP: VLAN0020 Gi1/0 -> learning
*Sep 13 14:56:26.560: STP: VLAN0001 sent Topology Change Notice on Gi1/0
*Sep 13 14:56:26.561: STP[1]: Generating TC trap for port GigabitEthernet0/1
*Sep 13 14:56:26.561: STP: VLAN0001 Gi0/1 -> blocking
*Sep 13 14:56:26.564: STP: VLAN0010 sent Topology Change Notice on Gi1/0
*Sep 13 14:56:26.564: STP[10]: Generating TC trap for port GigabitEthernet0/1
*Sep 13 14:56:26.565: STP: VLAN0010 Gi0/1 -> blocking
*Sep 13 14:56:26.568: STP: VLAN0020 sent Topology Change Notice on Gi1/0
*Sep 13 14:56:26.568: STP[20]: Generating TC trap for port GigabitEthernet0/1
*Sep 13 14:56:26.568: STP: VLAN0020 Gi0/1 -> blocking
*Sep 13 14:56:38.059: STP[1]: Generating TC trap for port GigabitEthernet1/0
*Sep 13 14:56:38.063: STP: VLAN0001 sent Topology Change Notice on Gi1/0
*Sep 13 14:56:38.064: STP: VLAN0001 Gi1/0 -> forwarding
*Sep 13 14:56:38.065: STP[10]: Generating TC trap for port GigabitEthernet1/0
*Sep 13 14:56:38.066: STP: VLAN0010 sent Topology Change Notice on Gi1/0
*Sep 13 14:56:38.066: STP: VLAN0010 Gi1/0 -> forwarding
*Sep 13 14:56:38.066: STP[20]: Generating TC trap for port GigabitEthernet1/0
*Sep 13 14:56:38.068: STP: VLAN0020 sent Topology Change Notice on Gi1/0
*Sep 13 14:56:38.068: STP: VLAN0020 Gi1/0 -> forwarding
```

You should see a total of 30 seconds elapse from when the port starts listening to when it finally transitions to forwarding. #### STP Timers STP uses three different timers to ensure proper loop-free convergence: * __hello time:__ The time between each BPDU that is sent on a port. Equals 2 seconds, by default * __forward delay:__ The time that is spent in the listening and learning state. Equals 15 seconds, by default * __maximum age:__ Controls the maximum length of time that a bridge port stores its BPDU information; equals 20 seconds, by default The transition between port states takes from 30 to 50 seconds, depending on the topology change. You can adjust STP timers. You can tune the hello time from 1 up to 10 seconds, the forward delay from 4 up to 30 seconds, and the maximum age from 6 up to 40 seconds. However, the timer values should never be changed without consideration. When changing the timers, you should

apply changes only on the root bridge. The root bridge will then propagate the timer values to the other switches. Return to ALS2 and run the ```show spanning-tree vlan 10``` command again to see the timer values on the current root bridge.

```
ALS2# show spanning-tree vlan 10

    VLAN0010
    .
                This bridge is the root
                Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
```

Step 3

Turn the interface G0/2 on DLS1 back on and disable STP debugging:

```
    DLS1(config)# interface G0/2
    DLS1(config-if)# no shutdown
    DLS1(config-if)# do undebug all
```

You have turned the interface back on to observe how the topology changes after a failed interface comes back up. On DLS1 and ALS1 use the ```show spanning-tree vlan 10``` command to observe how STP port roles get redefined after a failed interface comes back up. The port roles are again the same as they were before shutting down the interface. The port roles on ALS2 have not changed since it is the current root bridge for VLAN 10, and all ports are designated. > After you bring G0/2 on the DLS1 back up, it will take about 30 seconds for STP to make ports either forwarding or blocking. ## Part 3 - Enable IEEE 802.1w Rapid Spanning-Tree Now that you have investigated the default STP values and current topology, you will start by enabling RSTP. ### Step 1 Enable RSTP on all switches (DLS1, DLS2, ALS1, ALS2):

```
spanning-tree mode rapid-pvst
```

If all switches in the network, except one, are running RSTP, the interfaces that lead to legacy STP switches will automatically fall back to legacy STP. If you are using Cisco switches, they will fall back to PVST+. You can check whether all the switches have RSTP configured by observing the convergence time.

Step 2

Verify that RSTP is enabled on all four switches (DLS1, DLS2, ALS1, ALS2):

```
show spanning-tree vlan 10

    VLAN0010
    Spanning tree enabled protocol rstp
      <... output omitted ...>
```

If you look through the rest of the output from the `show spanning-tree vlan 10` command you will see that the port roles and states have remained the same as they were for PVST+.

Step 3

Re-enable the `debug spanning-tree events` command and then shut down the interface G0/2 on DLS1 and observe the convergence time of RSTP:

```
    DLS1# debug spanning-tree events
    Spanning Tree event debugging is on
    DLS1(config)# interface G0/2
    DLS1(config-if)# shutdown

    *Sep 13 15:59:14.084: RSTP(1): updt roles, root port Gi0/2 going down
    *Sep 13 15:59:14.084: RSTP(1): Gi1/0 is now root port
    *Sep 13 15:59:14.084: RSTP(1): syncing port Gi0/1
    *Sep 13 15:59:14.084: RSTP(1): syncing port Gi0/3
    *Sep 13 15:59:14.085: RSTP(10): updt roles, root port Gi0/2 going down
    *Sep 13 15:59:14.085: STP(10): Gi1/0 is now root port
    *Sep 13 15:59:14.085: RSTP(10): syncing port Gi0/1
    *Sep 13 15:59:14.085: RSTP(10): syncing port Gi0/3
    *Sep 13 15:59:14.086: RSTP(20): updt roles, root port Gi0/2 going down
    *Sep 13 15:59:14.086: RSTP(20): Gi1/0 is now root port
    *Sep 13 15:59:14.086: RSTP(20): syncing port Gi0/1
    *Sep 13 15:59:14.086: RSTP(20): syncing port Gi0/3
    *Sep 13 15:59:14.093: STP[1]: Generating TC trap for port GigabitEthernet1/0
    *Sep 13 15:59:14.094: STP[10]: Generating TC trap for port GigabitEthernet1/0
    *Sep 13 15:59:14.094: STP[20]: Generating TC trap for port GigabitEthernet1/0
    *Sep 13 15:59:14.097: RSTP(1): transmitting a proposal on Gi0/1
    *Sep 13 15:59:14.099: RSTP(1): transmitting a proposal on Gi0/3
    *Sep 13 15:59:14.100: RSTP(10): transmitting a proposal on Gi0/1
    *Sep 13 15:59:14.105: RSTP(10): transmitting a proposal on Gi0/3
    *Sep 13 15:59:14.108: RSTP(20): transmitting a proposal on Gi0/1
    *Sep 13 15:59:14.110: RSTP(20): transmitting a proposal on Gi0/3
    *Sep 13 15:59:14.124: RSTP(1): updt roles, received superior bpdu on Gi0/1
    *Sep 13 15:59:14.124: RSTP(1): Gi0/1 is now alternate
    *Sep 13 15:59:14.127: RSTP(1): synced Gi1/0
```

```
*Sep 13 15:59:14.127: RSTP(10): updt roles, received superior bpdu on Gi0/1
*Sep 13 15:59:14.127: RSTP(10): Gi0/1 is now alternate
*Sep 13 15:59:14.128: RSTP(10): synced Gi1/0
*Sep 13 15:59:14.129: RSTP(20): updt roles, received superior bpdu on Gi0/1
*Sep 13 15:59:14.129: RSTP(20): Gi0/1 is now alternate
*Sep 13 15:59:14.130: RSTP(20): synced Gi1/0
```

If you want to observe the port state recalculation, you must trigger a topology change. One of the options is to shut down the interface. How much time will it take spanning tree to converge now that you have enabled the rapid version? The convergence time of RSTP is much shorter than the convergence time of STP. The entire convergence happens at the speed of BPDU transmission. That convergence can be less than 1 second. Notice in the output above that as soon as the G0/2 interface is shutdown, G1/0 becomes the new root port, G1/1 remains an alternate port in blocking/discarding state, and G0/1 now transitions to alternate role (BLK) since the bridge ID is lower on ALS1. You will also notice the transmission of proposals on G0/3 for a total of 30 seconds as DLS1 tries to negotiate RSTP with the CORE_ROUTER. The G0/3 port role remains designated, but the port transitions quickly from blocking/discarding through listening and learning states to ensure that no STP loops exists on that port. After 30 seconds the port transitions to forwarding. ### Step 4 Turn the interface G0/2 on DLS1 back on and disable STP debugging:

```
DLS1(config)# interface G0/2
DLS1(config-if)# no shutdown
DLS1(config-if)# do undebug all
```

Part 4 - Tune RSTP for Root and Backup Root Bridge Election

You do not want the network to choose the root bridge by itself. If all switches have default STP priorities, the switch with the lowest MAC address will become the root bridge. The oldest switch will have the lowest MAC address since the lower MAC addresses were factory-assigned first. To manually set the root bridge, you can change the priority of the switch.

In the example topology, you do not want the access layer switch ALS2 to become the root bridge. If ALS2 was the root bridge, the links between the distribution layer switches would get blocked. The traffic between DLS1 and DLS2 would then need to go through ALS2, which is not optimal.

Consider this example traffic diagram where traffic between hots connected to access layer switches all must traverse ALS2.

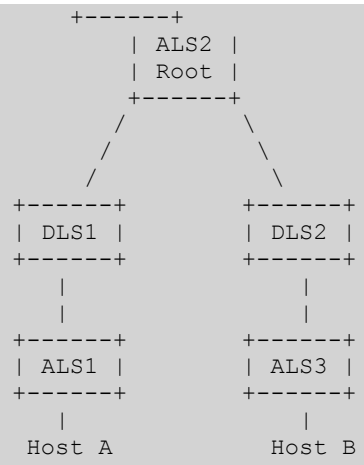

You want distribution or core switches to become the root bridge.

There are two ways of changing switch priority:

- Setting the exact value. The value must be between 0 and 61,440 and in increments of 4096.

  ```
  spanning-tree vlan vlan-id priority bridge-priority
  ```

  - The default value is 32,768.
  - 4 of the 16 priority bits represent the VLAN ID

- Setting the primary root bridge with a macro. Use the spanning-tree vlan vlan-id root primary command.

  ```
  spanning-tree vlan vlan-id root {primary | secondary}
  ```

  - If you issue the `show running-configuration` command, you will see the switch priority as a number—not the `primary` or `secondary` keyword.

Step 1

Configure DLS1 as the root bridge for VLANs 1 and 10 and secondary root bridge for VLAN 20. Use the `spanning-tree vlan vlan-id priority value` command to override the values that were pre-configured at the start of the lab. Use the value 4096 for the

primary root and 8192 for the secondary root. These values are lower than the current priority values across all four switches.

```
DLS1(config)# spanning-tree vlan 1 priority 4096
DLS1(config)# spanning-tree vlan 10 priority 4096
DLS1(config)# spanning-tree vlan 20 priority 8192
```

Step 2

Configure DLS2 as the root bridge for VLAN 20 and secondary root bridge for VLANs 1 and 10. Use the `spanning-tree vlan vlan-id priority value` command.

```
DLS2(config)# spanning-tree vlan 20 priority 4096
DLS2(config)# spanning-tree vlan 1 priority 8192
DLS2(config)# spanning-tree vlan 10 priority 8192
```

The reason for sharing the root bridge role across DLS1 and DLS2 for VLANs 10 and 20 is to provide load-balancing across redundant links. DLS1 is configured as the root bridge for VLAN 10 and DLS2 is configured as the root bridge for VLAN 20. ALS1 and ALS2 will forward traffic for VLAN 10 through DLS1, while they will also forward traffic for VLAN 20 through DLS2. > VLAN 1 is included in this step to avoid suboptimal forwarding of VLAN 1 traffic. ### Step 3 Verify that DLS1 is now the root bridge for VLAN 10 and the secondary root bridge for VLAN 20.

```
DLS1# show spanning-tree vlan 10

VLAN0010
  Spanning tree enabled protocol rstp
  Root ID    Priority    4106
             Address     5254.000e.8c2b
             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     5254.000e.8c2b
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec

DLS1# show spanning-tree vlan 20

VLAN0020
  Spanning tree enabled protocol rstp
  Root ID    Priority    4116
             Address     5254.0009.c377
             Cost        4
             Port        5 (GigabitEthernet1/0)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    8212  (priority 8192 sys-id-ext 20)
             Address     5254.0014.4977
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec
```

Step 4

Verify that DLS2 is now the root bridge for VLAN 20 and the secondary root bridge for VLAN 10.

```
DLS2# show spanning-tree vlan 10

VLAN0010
  Spanning tree enabled protocol rstp
  Root ID    Priority    4106
             Address     5254.000e.8c2b
             Cost        4
             Port        5 (GigabitEthernet1/0)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    8202  (priority 8192 sys-id-ext 10)
             Address     5254.001f.3178
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec

DLS2# show spanning-tree vlan 20

VLAN0020
  Spanning tree enabled protocol rstp
  Root ID    Priority    4116
             Address     5254.001f.3178
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    4116   (priority 4096 sys-id-ext 20)
             Address     5254.001f.3178
```

```
               Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
               Aging Time  300 sec
```

Part 5 - Enable PortFast and BPDU Guard on Access Ports

If a switch port connects to another switch, the STP initialization cycle must transition from state to state to ensure a loop-free topology.

However, for access devices such as PCs, laptops, servers, and printers, the delays incurred with STP initialization can cause problems such as DHCP timeouts. Cisco designed PortFast, which, together with BPDU, can be used as enhancements to STP to reduce the time that is required for an access device to enter the forwarding state.

When the PortFast feature is enabled on a switch port that is configured as an access port, that port bypasses the typical STP listening and learning states. This feature allows the port to transition from the blocking to the forwarding state immediately. You can use PortFast on access ports that are connected to a single workstation or to a server to allow those devices to connect to the network immediately rather than waiting for the spanning tree to converge.

The STP PortFast BPDU guard enhancement allows network designers to enforce the STP domain borders and keep the active topology predictable. The devices behind the ports that have STP PortFast enabled are not able to influence the STP topology. At the reception of BPDUs, the BPDU guard operation disables the port that has PortFast configured. The BPDU guard transitions the port into "errdisable" state, and a message appears on the console.

Step 1

Enable PortFast and BPDU Guard on ALS1 and ALS2 G1/0, G1/1, and G1/2. Use the `interface range` command to speed up this configuration task.

```
ALS1(config)# interface range g1/0-2
ALS1(config-if-range)# 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

ALS1(config-if-range)# spanning-tree bpduguard enable

ALS2(config)# interface range g1/0-2
ALS2(config-if-range)# 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

ALS2(config-if-range)# spanning-tree bpduguard enable
```

> It is also possible to enable PortFast and BPDU Guard by default on the switch with the `` spanning-tree portfast default `` and the `` spanning-tree portfast bpduguard default `` commands. ### Step 2 Use the ``show spanning-tree interface G1/0 detail`` and ``show spanning-tree interface G1/1 detail`` command on ALS1 and ALS2 to verify that PortFast and BPDU Guard are both enabled.

```
ALS1# show spanning-tree interface g1/0 detail
 Port 5 (GigabitEthernet1/0) of VLAN0010 is designated forwarding
   Port path cost 4, Port priority 128, Port Identifier 128.5.
   Designated root has priority 4106, address 5254.000e.8c2b
   Designated bridge has priority 20490, address 5254.0014.cbfc
   Designated port id is 128.5, designated path cost 4
   Timers: message age 0, forward delay 0, hold 0
   Number of transitions to forwarding state: 2
   The port is in the portfast edge mode
   Link type is point-to-point by default
   Bpdu guard is enabled
   BPDU: sent 1531, received 0
```

Step 3

To illustrate the importance of BPDU guard to prevent connecting a switch/bridge to a port that should only be connected to hosts, we will connect ALS1 and ALS2 interfaces G1/2 - a port we just configured for portfast and bpduguard.

Right-click ALS1 and choose "Add link". Connect to ALS2 and be sure to select interfaces **GigE1/2** on both switches.

The interfaces were pre-configured in `shutdown` state. Go ahead and enable each interface.

```
ALS1(config)#interface GigabitEthernet 1/2
ALS1(config-if)#no shut
ALS1(config-if)#end
```

```
ALS2(config)#interface GigabitEthernet 1/2
ALS2(config-if)#no shut
ALS2(config-if)#end
```

As soon as you enable the second interface, you should see messages indicating the misconfiguration.

```
*Sep 16 16:48:23.755: %SPANTREE-2-BLOCK_BPDUGUARD: Received BPDU on port Gi1/2 with BPDU Guard enabled. Disabling port.
*Sep 16 16:48:23.756: %PM-4-ERR_DISABLE: bpduguard error detected on Gi1/2, putting Gi1/2 in err-disable state
*Sep 16 16:48:24.760: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet1/2, changed state to down
*Sep 16 16:48:25.762: %LINK-3-UPDOWN: Interface GigabitEthernet1/2, changed state to down
```

And check the status of the interfaces to see the state `err-disabled`

```
ALS1#show interfaces g1/2
GigabitEthernet1/2 is down, line protocol is down (err-disabled)
```

> An `err-disabled` port will not automatically recover. It requires administrator intervention. This is done by fixing whatever network problem caused the disabled state, and then issuing a `shutdown` followed by `no shutdown` command on the interface.

Remove the link from the topology by right clicking the added link in the topology and choosing "Delete".

Step 4

Recall earlier that DLS1 transmitted RSTP proposals on G0/3 to CORE_ROUTER for 30 seconds when a topology change was detected. Since DLS1 G0/3 is connected to a router and not a L2 switch, you can enable PortFast even though G0/3 is in trunking mode.

```
DLS1(config)# interface G0/3
DLS1(config-if)# spanning-tree portfast trunk
%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
```

Step 5

Use the `show spanning-tree vlan 10 interface g0/3 detail` command to verify the PortFast status of the trunk.

```
DLS1(config-if)# do show spanning-tree vlan 10 interface g0/3 detail
 Port 4 (GigabitEthernet0/3) of VLAN0010 is designated forwarding
.
   The port is in the portfast edge trunk mode
.
```

Part 6 - Enable Root Guard

The root guard feature of Cisco switches prevents a switch or rogue device from becoming a root bridge in a spanning tree domain. The root guard feature is designed to provide a way to enforce the placement of root bridges in the network. Root guard limits the switch ports from which the root bridge can be negotiated. If a port where root guard is enabled receives BPDUs that are superior to BPDUs that the current root bridge is sending, then the port transitions to a root-inconsistent state, which is effectively equal to an STP listening state, and no data traffic is forwarded across that port.

Step 1

Enable Root Guard on DLS1 and DLS2 ports that connect to ALS1 and ALS2.

```
DLS2(config)# interface range g0/1-2
DLS1(config-if-range)# spanning-tree guard root
DLS1(config-if-range)#
*Sep 13 17:55:22.294: %SPANTREE-2-ROOTGUARD_CONFIG_CHANGE: Root guard enabled on port GigabitEthernet0/1.
*Sep 13 17:55:22.298: %SPANTREE-2-ROOTGUARD_CONFIG_CHANGE: Root guard enabled on port GigabitEthernet0/2.

DLS2(config)#interface range g0/1-2
DLS2(config-if-range)# spanning-tree guard root
DLS2(config-if-range)#
*Sep 13 18:04:23.877: %SPANTREE-2-ROOTGUARD_CONFIG_CHANGE: Root guard enabled on port GigabitEthernet0/1.
*Sep 13 18:04:23.881: %SPANTREE-2-ROOTGUARD_CONFIG_CHANGE: Root guard enabled on port GigabitEthernet0/2.
```

The root guard feature prevents a switch from becoming a root bridge on configured ports. Root guard is best deployed toward ports that connect to switches that should not be the root bridge. ### Step 2 Use the ``show spanning-tree vlan 10 interface g0/1`` and ``show spanning-tree vlan 10 interface g0/2`` commands on DLS1 and DLS2 to verify that Root Guard is enabled.

```
DLS1# show spanning-tree vlan 10 interface g0/1 detail
 Port 2 (GigabitEthernet0/1) of VLAN0010 is designated forwarding
```

```
   .
   .
      Root guard is enabled on the port
   .
```

Step 3

Now we'll verify that root guard protects root status by re-configuring ALS1 with a better (ie lower) priority than either distribution layer switch.

The lowest possible priority value that can be configured is `0`. Set this on ALS1 for all possible VLANs.

```
ALS1(config)#spanning-tree vlan 1-4094 priority 0
```

As soon as you change the priority on ALS1, DLS1 and DLS2 will report errors.

```
   DLS1# *Sep 16 17:00:22.663: %SPANTREE-2-ROOTGUARD_BLOCK: Root guard blocking port GigabitEthernet0/1 on VLAN0001.

   DLS2# *Sep 16 17:00:21.974: %SPANTREE-2-ROOTGUARD_BLOCK: Root guard blocking port GigabitEthernet0/2 on VLAN0001.
```

Look at the spanning-tree details for interface g0/1 on DLS1. Verify that it shows the expected error for all three vlans.

```
   DLS1#show spanning-tree interface g0/1 detail
    Port 2 (GigabitEthernet0/1) of VLAN0001 is broken  (Root Inconsistent)
      .

    Port 2 (GigabitEthernet0/1) of VLAN0010 is broken  (Root Inconsistent)
      .

    Port 2 (GigabitEthernet0/1) of VLAN0020 is broken  (Root Inconsistent)
      .
```

The status is also displayed in the output of `show spanning-tree`.

```
   DLS1# show span vlan 10

   VLAN0010
   .

   Interface           Role Sts Cost      Prio.Nbr Type
   ------------------- ---- --- --------- -------- --------------------------------
   Gi0/1               Desg BKN*4         128.2    P2p *ROOT_Inc
   Gi0/2               Desg FWD 4         128.3    P2p
   Gi0/3               Desg FWD 4         128.4    P2p
   Gi1/0               Desg FWD 4         128.5    P2p
   Gi1/1               Desg FWD 4         128.6    P2p
```

Fix the network by removing the incorrect bridge priority value on ALS1.

```
ALS1(config)#no spanning-tree vlan 1-4094 priority
```

Unlike ports that are `err-disabled` due to BPDU guard, root inconsistent ports are automatically re-enabled when the configuration problem is fixed.

```
*Sep 16 17:11:11.875: %SPANTREE-2-ROOTGUARD_UNBLOCK: Root guard unblocking port GigabitEthernet0/1 on V
    *Sep 16 17:11:13.173: %SPANTREE-2-ROOTGUARD_UNBLOCK: Root guard unblocking port GigabitEthernet0/1
```

Part 7 - Enable Loop Guard

A Layer 2 loop is created when an STP blocking port in a redundant topology erroneously transitions to the forwarding state. This situation usually happens because one of the ports of a physically redundant topology (not necessarily the STP blocking port) no longer receives STP BPDUs

When one of the ports in a physically redundant topology no longer receives BPDUs, STP conceives that the topology is loop-free. Eventually, the blocking port from the alternate or backup port becomes designated and moves to a forwarding state. This situation creates a loop.

The loop guard feature performs additional checks. If BPDUs are not received on a nondesignated port, and loop guard is enabled, that port is moved into the STP loop-inconsistent blocking state, instead of the listening, learning, and eventually forwarding state. Without the loop guard feature, the port takes on the designated port role. The port moves to the STP forwarding state and creates a loop.

Root guard is mutually exclusive with loop guard. Root guard is used on designated ports, and it does not allow the port to become nondesignated. Loop guard works on nondesignated ports and does not allow the port to become designated through the expiration of the Max Age timer. Root guard cannot be enabled on the same port as loop guard. When loop guard is configured on a port, it disables any

root guard that is already configured on the same port. Loop guard must be enabled on the nondesignated ports (more precisely, on root and alternate ports) for all possible combinations of active topologies.

Step 1

Enable Loop Guard on ALS1 G0/1 and G0/2, as well as on ALS2 G0/1 and G0/2.

```
ALS1(config)# interface range g0/1-2
ALS1(config-if-range)# spanning-tree guard loop

ALS2(config)# interface range g0/1-2
ALS2(config-if-range)# spanning-tree guard loop
```

> It is also possible to configure Loop Guard by default on a switch with the `` spanning-tree loopguard default `` command. ### Step 2 Verify that Loop Guard is enabled on ALS1 and ALS2 G0/1 and G0/2.

```
ALS1# show spanning-tree vlan 10 interface g0/1 detail
 Port 2 (GigabitEthernet0/1) of VLAN0010 is root forwarding
.
   Loop guard is enabled on the port
.
```

Note: Network problems that trigger loop guard typically occur due to Layer 1 (ie cabling; transceiver) issues, or strange configurations on hosts that act as bridges. These are difficult to simulate, but configuring loop guard is a best practice to protect against these unlikely but disruptive problems.

Great job! You finished the lab! You've learned how to:

- investigate the spanning-tree defaults
- enable RSTP
- enable features such as PortFast, BPDU Guard, Root Guard, & Loop Guard