


Lab 3.10.7 Port Level Tuning to Control STP Behavior



Objective

The purpose of this lab is to use PortFast, UplinkFast, BPDU guard, root guard, and UDLD to control STP behavior on a port.

Note: This lab uses fiber connections between the ALSwitch1 and DLSwitch1 and DLSwitch2. If the available equipment does not have fiber connections, use CAT 5 crossover cables between the Gigabit Ethernet interfaces. However, instructions and tasks for Step 8 (uplinkfast) and Step 11 (UDLD) cannot be followed exactly and certain results will not be as indicated or expected.

Scenario

A new redundant switched network has just been implemented. The default behavior of Spanning-Tree Protocol (STP) has created some undesirable results. The ports take up to 50 seconds to reach forwarding state. This prevents DHCP clients from receiving an IP address during normal boot-up. PortFast will be used to prevent this problem in the future.

Enabling PortFast can create a security risk in a switched network. A port configured with PortFast will go into blocking state if it receives a Bridge Protocol Data Unit (BPDU). An unauthorized device can send BPDUs into the PortFast interface and set a port to blocking. When the port is in blocking state it will accept all BPDUs. This could lead to false STP information that enters the switched network and causes unexpected STP behavior. Bridge Guard Data Unit (BGDU) will be used to prevent unauthorized BPDUs from entering the switched network through PortFast enabled ports.

When the active uplink between the two switches is broken, it takes the redundant link 30 seconds to complete the spanning-tree process before bringing up the backup, or blocked, link. This results in a temporary network outage for users. UplinkFast will be used to reduce STP convergence time.

ALSwitch2 is connected with a slower and more unreliable connection. The network administrator wants to prevent the ALSwitch2 from becoming the root bridge or from being in the path to the root bridge. ALSwitch2 should be avoided as much as possible. Root guard will be used to prevent ALSwitch2 from becoming the root bridge.

ALSwitch1 is connected to the distribution layer with Gigabit Ethernet links. If the transmit or receive link in a fiber cable is disconnected or cut, then it could lead to a unidirectional link. Unidirectional links can transmit or receive data, but not both. Unidirectional links have an adverse effect on the network. Use UniDirectional Link Detection (UDLD) protocol to prevent unidirectional links from occurring.

The network design is as follows.

| Catalyst Type | Switch | VTP Domain | VTP Mode |
|---------------|-----------|------------|----------|
| 3550 | DLSwitch1 | CORP | Server |
| 3550 | DLSwitch2 | CORP | Client |
| 2950 | ALSwitch1 | CORP | Client |
| 2950 | ALSwitch2 | CORP | Client |

The VLAN configuration information is as follows.

| VLAN ID | VLAN Name | VLAN Subnet | DLSwitch1 and DLSwitch2 | ALSwitch1 and ALSwitch2 |
|---------|------------|----------------|-------------------------|---------------------------------|
| 1 | Native | 172.16.1.0/24 | All Ports | Gi0/1-2 Fa0/1-4 Fa0/12-24 |
| 10 | Accounting | 172.16.10.0/24 | | Fa0/5-8 |
| 20 | Marketing | 172.16.20.0/24 | | FA0/9-12 |
| Trunk | | 802.1Q | 802.1Q | 802.1Q |

Step 1

Do not cable the lab until all switch configurations and `vlan.dat` files have been erased.

Delete the vlan database if it exists on any switches and clear the configuration.

```
switch#delete flash:vlan.dat
Delete filename [vlan.dat]?
```

```

Delete flash:vlan.dat? [confirm]
switch#
switch#erase startup-config
Erasing the nvram filesystem will remove all files! Continue? [confirm]
DLSwitchA#reload

System configuration has been modified. Save? [yes/no]:n
Proceed with reload? [confirm]

```

Cable the lab according to the diagram. Crossover Cat 5 cables must be used since the devices are similar.

Configure the hostname, passwords, and Telnet access to all the switches. Configure the interface VLAN 1 IP address on each switch.

Step 2

Observe the default behavior of Spanning-Tree (STP) using the **show spanning-tree** command on all switches.

1. Which switch became the root bridge?
-

2. What command was used to find the root bridge?
-

Step 3

Configure the trunking interfaces to create a trunk link between the switches. Set the port to trunking with 802.1Q encapsulation on DLSwitch1 and DLSwitch2.

Note: An error may appear because the port is set to auto encapsulation. If this occurs, enter the **switchport mode trunk** command after the **switchport trunk encapsulation dot1q** command.

```

DLSwitch1(config)#interface range gigabitethernet 0/2 , fastethernet 0/3
DLSwitch1(config-if-range)#switchport trunk encapsulation dot1q
DLSwitch1(config-if-range)#switchport mode trunk
DLSwitch1(config-if-range)#^Z

```

```

DLSwitch2(config)#interface range gigabitethernet 0/1 , fastethernet 0/3
DLSwitch2(config-if-range)#switchport trunk encapsulation dot1q
DLSwitch2(config-if-range)#switchport mode trunk
DLSwitch2(config-if-range)#^Z

```

The 2950 switches do not need the encapsulation configured. These switches default to 802.1Q. Some IOS versions do not offer any other options. Console into each access layer switch and configure trunking.

```

ALSwitch1(config)#interface range gigabitethernet 0/1 , gigabitethernet 0/2
ALSwitch1(config-if-range)#switchport mode trunk
ALSwitch1(config-if-range)#^Z

ALSwitch2(config)#interface range fastethernet 0/1 , fastethernet 0/3
ALSwitch2(config-if-range)#switchport mode trunk
ALSwitch2(config-if-range)#^Z

```

Verify the trunk configuration on each switch with the **show interfaces trunk** command.

Step 4

Console into DLSwitch1 and configure the vtp domain CORP, server mode, and the appropriate VLANs and names as shown below.

```
DLSwitch1#vlan database
DLSwitch1(vlan)#vtp domain CORP
DLSwitch1(vlan)#vtp server
DLSwitch1(vlan)#vlan 10 name Accounting
DLSwitch1(vlan)#vlan 20 name Marketing
DLSwitch1(vlan)#exit
```

Configure DLSwitch2 as a VTP client as shown below.

```
DLSwitch2#vlan database
DLSwitch2(vlan)#vtp client
DLSwitch2(vlan)#exit
```

Step 5

Configure ALSwitch1 and ALSwitch2 as VTP clients and assign ports to the respective VLANs in each switch as shown below. The **interface range** command can be used to configure several interfaces at the same time.

```
ALSwitch1#vlan database
ALSwitch1(vlan)#vtp client
ALSwitch1(vlan)#exit
ALSwitch1#config terminal
ALSwitch1(config)#interface range fastethernet 0/5 - 8
ALSwitch1(config-if-range)#switchport access vlan 10
ALSwitch1(config-if-range)#interface range fastethernet 0/9 - 12
ALSwitch1(config-if-range)#switchport access vlan 20
ALSwitch1(config-if-range)#^Z

ALSwitch2#vlan database
ALSwitch2(vlan)#vtp client
ALSwitch2(vlan)#exit
ALSwitch2#config terminal
ALSwitch2(config)#interface range fastethernet 0/5 - 8
ALSwitch2(config-if-range)#switchport access vlan 10
ALSwitch2(config-if-range)#interface range fastethernet 0/9 - 12
ALSwitch2(config-if-range)#switchport access vlan 20
ALSwitch2(config-if-range)#^Z
```

Console into each switch and verify the VTP and VLAN configurations with the **show vtp status** and **show vlan** commands.

Step 6

Configure DLSwitch1 as the root bridge.

Change the root bridge priority for each VLAN on DLSwitch1 to 4096.

```
DLSwitch1(config)#spanning-tree vlan 1 priority 4096
DLSwitch1(config)#spanning-tree vlan 10 priority 4096
DLSwitch1(config)#spanning-tree vlan 20 priority 4096
```

Verify that DLSwitch1 is the root bridge for each VLAN with the **show spanning-tree** command.

```
DLSwitch1#show spanning-tree
```

```
VLAN0001
```

```
Spanning tree enabled protocol ieee
```

```
Root ID    Priority    4097  
Address    000a.b701.f700
```

```
This bridge is the root
```

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

```
Bridge ID  Priority    4097 (priority 4096 sys-id-ext 1)
```

```
Address    000a.b701.f700
```

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

```
Aging Time 300
```

| Interface Name | Port ID Prio.Nbr | Cost | Sts | Designated Cost Bridge ID | Port ID Prio.Nbr |
|-------------------|---------------------|------|-----|------------------------------|---------------------|
| Fa0/3 | 128.3 | 19 | FWD | 0 4097 000a.b701.f700 | 128.3 |
| Gi0/2 | 128.26 | 4 | FWD | 0 4097 000a.b701.f700 | 128.26 |

```
VLAN0010
```

```
Spanning tree enabled protocol ieee
```

```
Root ID    Priority    4106  
Address    000a.b701.f700
```

```
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    000a.b701.f700
```

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

```
Aging Time 300
```

| Interface Name | Port ID Prio.Nbr | Cost | Sts | Designated Cost Bridge ID | Port ID Prio.Nbr |
|-------------------|---------------------|------|-----|------------------------------|---------------------|
| Fa0/3 | 128.3 | 19 | FWD | 0 4106 000a.b701.f700 | 128.3 |
| Gi0/2 | 128.26 | 4 | FWD | 0 4106 000a.b701.f700 | 128.26 |

```
VLAN0020
```

```
Spanning tree enabled protocol ieee
```

```
Root ID    Priority    4116  
Address    000a.b701.f700
```

```
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    000a.b701.f700
```

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

```
Aging Time 300
```

| Interface Name | Port ID Prio.Nbr | Cost | Sts | Designated Cost Bridge ID | Port ID Prio.Nbr |
|-------------------|---------------------|------|-----|------------------------------|---------------------|
| Fa0/3 | 128.3 | 19 | FWD | 0 4116 000a.b701.f700 | 128.3 |
| Gi0/2 | 128.26 | 4 | FWD | 0 4116 000a.b701.f700 | 128.26 |

Step 7

Observe the default behavior of spanning tree. Connect a workstation to any of the switch ports on either access layer switch and turn on the workstation. After the NIC is initialized by the operating system, the port will turn yellow. The port is now active and starting the spanning-tree process. Watch the workstation boot up and watch the color of the link light. The workstation should make it through most of the startup before the link turns green and active. This is where DHCP has the opportunity to get an IP address while spanning tree is in listening and learning state.

It should take about 30 seconds for a new device to become active in a port.

Configure PortFast on the switch ports.

Configure FastEthernet ports 0/5 through 12 for PortFast on the access layer switches.

```
ALSwitch1(config)#interface range fastethernet 0/5 - 12
ALSwitch1(config-if-range)#spanning-tree portfast
```

Warning: PortFast should only be enabled on ports that are connected to a single host. If hubs, concentrators, switches, and bridges are connected to the interface when PortFast is enabled, temporary bridging loops can occur. Use with caution.

PortFast will be configured in eight interfaces with the **range** command. However, it will only be effective when the interfaces are in a non-trunking mode.

```
ALSwitch2(config)#interface range fastethernet 0/5 - 12
ALSwitch2(config-if-range)#spanning-tree portfast
```

Verify that PortFast is operating on the access layer switches.

Remove the workstation from the switch and plug it into any port configured with PortFast. The port should become active immediately. The access layer switch indicator light will become green without the yellow learning and listening period. Use the **show spanning-tree** command to check the state of each link.

3. How could PortFast create bridging loops?

Step 8

Observe what happens when the status of an uplink changes.

Remove the uplink cable between ALSwitch1 and DLSwitch1 while monitoring the backup link port. Observe if the light on the switch is indicating a yellow blocked port or use the **show spanning-tree** command.

It should take about 30 seconds for the backup uplink ports to become active. Reconnect the cable between ALSwitch1 and DLSwitch1.

UplinkFast will now be enabled on ALSwitch2.

```
ALSwitch2(config)#spanning-tree uplinkfast
```

Use the following command to verify the UplinkFast configuration.

```
ALSwitch2#show spanning-tree summary total
```

```
Root bridge for: none.
Extended system ID is enabled.
PortFast BPDU Guard is disabled
EtherChannel misconfiguration guard is enabled
UplinkFast is enabled
BackboneFast is disabled
Default pathcost method used is short
```

| Name | Blocking | Listening | Learning | Forwarding | STP Active |
|---------|----------|-----------|----------|------------|------------|
| 3 vlans | 0 | 0 | 0 | 3 | 3 |

```
Station update rate set to 150 packets/sec.
```

```
UplinkFast statistics
-----
```

```
Number of transitions via uplinkFast (all VLANs) : 0
Number of proxy multicast addresses transmitted (all VLANs) : 0
```

Disconnect the cable between ALSwitch1 and DLSwitch2 while monitoring the backup uplink port. The backup port should come up in less than ten seconds.

Step 9

Use the global configuration mode to enable the BPDU guard feature on ALSwitch1.

```
ALSwitch1(config)#spanning-tree portfast bpduguard
```

When the BPDU guard feature is enabled on the switch, STP shuts down PortFast enabled interfaces that receive BPDUs instead of putting them into a blocking state. PortFast-enabled interfaces do not receive BPDUs in a valid configuration. The receipt of a BPDU by a PortFast-enabled interface indicates an invalid configuration such as the connection of an unauthorized device. The BPDU guard feature blocks BPDUs by placing the interface in the ErrDisable state. The BPDU guard feature provides a secure response to invalid configurations because the interface must be manually placed back in service.

Configure port FastEthernet0/1 on ALSwitch1 to access mode with PortFast enabled.

```
ALSwitch1(config)#interface fastethernet 0/1
ALSwitch1(config-if)#switchport mode access
ALSwitch1(config-if)#spanning-tree portfast
ALSwitch1(config-if)#exit
```

Connect a cable between FastEthernet 0/1 on ALSwitch1 to FastEthernet 0/1 on DLSwitch1.

The following error should appear.

```
05:31:56: %SPANTREE-2-RX_PORTFAST: Received BPDU on PortFast enabled port.
Disabling FastEthernet0/1.
05:31:56: %PM-4-ERR_DISABLE: bpduguard error detected on Fa0/1, putting
Fa0/1 in err-disable state
05:31:57: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1,
changed state to down
```

The switch receives the error and shuts down the port. This protects the switch from accepting false BPDUs.

Step 10

Prevent ALSwitch2 from becoming the root or from being in the path to the root.

The Layer 2 network of a service provider (SP) can include many connections to switches that are not owned by the SP. STP can reconfigure itself in this type of topology and select a customer switch as the STP root switch. The root-guard feature can be configured on interfaces that connect to switches outside of the customer network. STP calculations can be used to identify an interface in the customer network as the root port. Root guard will place this interface in the root-inconsistent or blocked state to prevent the customer switch from becoming the root switch or from being in the path to the root.

UplinkFast must be disabled because it cannot be used with root guard.

```
ALSwitch2(config)#no spanning-tree uplinkfast
```

Configure all the DLSwitch1 and DLSwitch2 ports that connect to ALSwitch2 with root guard.

```
DLSwitch1(config)#interface fastethernet 0/3
DLSwitch1(config-if)#spanning-tree guard root
DLSwitch1(config-if)#exit

DLSwitch2(config)#interface fastethernet 0/3
DLSwitch2(config-if)#spanning-tree guard root
DLSwitch2(config-if)#exit
```

Configure ALSwitch2 with a lower STP priority than DLSwitch1 for VLAN 1. ALSwitch2 would become the root for VLAN1 without root guard.

```
ALSwitch2(config)#spanning-tree vlan 1 priority 0
```

Issue the **show spanning-tree** command on DLSwitch1.

DLSwitch1 will still be the root bridge for VLAN 1 on ALSwitch1 and DLSwitch2. Root guard prevented ALSwitch2 from becoming the root bridge.

Interface FastEthernet 0/3 on both the DLSwitch1 and DLSwitch2 are in the blocking state for VLAN 1 which essentially prevents any VLAN 1 traffic from traversing the ALSwitch2 links.

Step 11

Disconnect one of the connectors between ALSwitch1 and DLSwitch1. Observe the line status on the switches. A unidirectional link has just been created.

A unidirectional link occurs when traffic sent by the local device is received by the neighbor but traffic from the neighbor is not received by the local device. This indicates that the transmit or receive part of the connection is broken. This can be caused by a cut or disconnected cable.

UDLD is a Layer 2 protocol that enables devices connected through fiber-optic or twisted-pair Ethernet cables to monitor the physical configuration of the cables and detect a unidirectional link.

All connected devices must support UDLD for the protocol to identify and disable unidirectional links. When UDLD detects a unidirectional link, it shuts down the affected port and sends out an alert. Unidirectional links can cause a variety of problems such as spanning-tree topology loops.

Now reconnect the transmit or receive cable to the switch.

Enable UDLD with the global configuration command **udld enable** on the DLSwitch1, DLSwitch2, and ALSwitch1.

Note: This command only affects fiber-optic interfaces. Use the **udld** interface configuration command to enable UDLD on other interface types.

```
ALSwitch1(config)#udld enable
DLSwitch1(config)#udld enable
DLSwitch2(config)#udld enable
```

Disconnect one of the fiber connectors between ALSwitch1 and DLSwitch1. Observe what happens to the line status on the two switches.

UDLD will administratively shut down the port.