



## Rapid Spanning Tree (RSTP) Traffic Engineering

« Port Channels | Multiple Spanning Tree (MST) Traffic Engineering »

Last updated: January 18, 2017

### Rapid Spanning Tree (RSTP) Traffic Engineering

#### Note:

Configure the tasks "VLANs & Trunking" and "Port Channels" prior to beginning this lab.

#### Objective

- Influence layer 2 traffic flows by modifying the STP Root Bridge and Root Port elections.

#### Task

- Per the previous task the network should be configured as follows:
  - The first server's link to the first 5K is in VLAN 10 with address 10.0.0.1X/24, where X is the server number.
  - The first server's link to the second 5K is in VLAN 20 with address 20.0.0.1X/24, where X is the server number.
  - The second server's link to the first 5K is in VLAN 20 with address 20.0.0.1X/24, where X is the server number.
  - The second server's link to the second 5K is in VLAN 10 with address 10.0.0.1X/24, where X is the server number.
  - 5Ks and 7Ks have SVIs for VLANs 10 & 20 with IP addressing 10.0.0.X/24 and 20.0.0.X/24 respectively, where X is the device number.
  - 5Ks have access ports to the servers and trunk ports to the 7Ks, and the 7Ks have trunks to the 5Ks.
  - Port-Channel5 is configured between the 5Ks, Port-Channel7 is configured between the 7Ks, and both are trunks.
- Modify your previous configuration as follows:
  - Re-enable the diagonal links between the 5Ks and 7Ks, and configure them as trunks.
  - Allow all VLANs on all trunks on both the 5Ks and 7Ks.
- Modify Spanning-Tree Protocol between the switches as follows:
  - All switches should use 32 bits for spanning-tree port path costs.
  - Set your first 7K to be the STP Root Bridge for VLAN 10, with the other 7K being the backup Root Bridge.
  - Set your first 5K to be the STP Root Bridge for VLAN 20, with the other 5K being the backup Root Bridge.
  - Your first server's VLAN 10 traffic to the second server should follow the path:
    - First server -> First 5K -> First 7K -> Second 7K -> Second 5K -> Second server
    - E.g. for Rack 1, Server 1 -> N5K1 -> N7K1 -> N7K2 -> N5K2 -> Server 2.
  - Your first server's VLAN 20 traffic to the second server should follow the path:
    - First server -> Second 5K -> First 7K -> First 5K -> Second server
    - E.g. for rack 1, Server 1 -> N5K2 -> N7K1 -> N5K1 -> Server 2.
- Once complete you should have IP reachability between the servers in both VLANs 10 & 20.
- Verify the new Root Bridge elections for both VLANs 10 & 20 and that layer 2 traffic is flowing between the servers as specified.

#### Configuration [Click to collapse](#)

```
N5K1:
vlan 10,20
spanning-tree pathcost method long
spanning-tree vlan 20 priority 24576
!
interface Ethernet1/1
switchport
switchport mode access
switchport access vlan 10
no shutdown
!
interface Ethernet1/2
switchport
switchport mode access
switchport access vlan 20
no shutdown
!
interface Ethernet1/3 - 4
switchport
switchport mode trunk
channel-group 5
no shutdown
!
interface port-channel5
switchport mode trunk
speed 10000
!
interface Ethernet1/5 - 6
switchport
switchport mode trunk
no shutdown
```

```
N5K2:
vlan 10,20
spanning-tree pathcost method long
spanning-tree vlan 20 priority 28672
!
interface Ethernet1/1
switchport
switchport mode access
switchport access vlan 20
no shutdown
!
interface Ethernet1/2
switchport
switchport mode access
switchport access vlan 10
no shutdown
```

```

!
interface Ethernet1/3 - 4
  switchport
  switchport mode trunk
  channel-group 5
  no shutdown
!
interface port-channel5
  switchport mode trunk
  speed 10000
!
interface Ethernet1/5
  switchport
  switchport mode trunk
  spanning-tree vlan 20 cost 1
  no shutdown
!
interface Ethernet1/6
  switchport
  switchport mode trunk
  spanning-tree vlan 10 cost 1
  no shutdown

N7K1:
feature lacp
!
vlan 10,20
spanning-tree pathcost method long
spanning-tree vlan 10 priority 24576
!
interface Ethernet1/1 - 2
  lacp rate fast
  switchport
  switchport mode trunk
  channel-group 7 mode active
  no shutdown
!
interface port-channel7
  switchport
  switchport mode trunk
  lacp min-links 2
!
interface Ethernet1/3
  switchport
  switchport mode trunk
  spanning-tree vlan 20 cost 1
  no shutdown
!
interface Ethernet1/4
  switchport
  switchport mode trunk
  no shutdown

N7K2:
feature lacp
!
vlan 10,20
spanning-tree pathcost method long
spanning-tree vlan 10 priority 28672
!
interface Ethernet1/1 - 2
  lacp rate fast
  switchport
  switchport mode trunk
  channel-group 7 mode active
  no shutdown
!
interface port-channel7
  switchport
  switchport mode trunk
  lacp min-links 2
!
interface Ethernet1/3 - 4
  switchport
  switchport mode trunk
  no shutdown

```

## Verification

NX-OS runs Rapid Per-VLAN Spanning-Tree Protocol by default. This means that for each VLAN that is created, a separate instance of STP is created, with each of these running the 802.1w RSTP algorithm. Beyond this, the default behavior of STP on NX-OS is essentially identical to that of Catalyst IOS.

For the purposes of traffic engineering, the Root Bridge election can be modified on a per-VLAN basis by changing the STP Bridge-ID Priority value (lower is preferred), and the Root Port election can be modified on a per-port per-VLAN basis by changing either the STP Port Cost or STP Port Priority. Changing the Port Cost is the more common modification, as the Port Priority only comes into play when you are choosing between multiple links with the same Root Path Cost, and the links connect to the same upstream switch.

Note that in this example the range of STP cost values was increased from a 16-bit field to a 32-bit field (the long pathcost method), as newer higher-speed links such as 10GigE, 40GigE, 100GigE, and beyond start to look the same from a cost point of view when the default short pathcost method is used. Also note that this command is applicable only when running Rapid PVST on NX-OS, as Multiple Spanning-Tree (MST) always uses the longer pathcost length.

Verification of this task can be performed by viewing the Root Bridge and Root Port election on a per-switch basis, or by viewing the MAC address table, as the STP topology ultimately controls which interfaces can participate in MAC address learning. Below we see that for VLAN 10, N7K1 is elected the Root Bridge. This implies that all of its VLAN 10 links will be Designated ports in the Forwarding state.

```

N7K1# show spanning-tree vlan 10

VLAN0010
  Spanning tree enabled protocol rstp
  Root ID    Priority    24586
            Address    68bd.abd7.6042
            This bridge is the root

```

<https://t.me/learningnets>

```

Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
  Bridge ID Priority 24586 (priority 24576 sys-id-ext 10)
  Address 68bd.abd7.6042
  Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Po7	Desg FWD		1000	128.4102	P2p
Eth1/3	Desg FWD		2000	128.131	P2p
Eth1/4	Desg FWD		2000	128.132	P2p

The servers' MAC addresses for VLAN 10 are being learned in ports Eth1/3 and Po7, which implies that N5K1 and N7K2 on the other end of these links, respectively, have chosen those ports as their Root Ports.

```

N7K1# show mac address-table vlan 10
Note: MAC table entries displayed are getting read from software.
Use the 'hardware-age' keyword to get information related to 'Age'

Legend:
  * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
  age - seconds since last seen,+ - primary entry using vPC Peer-Link, E - EVPN entry
  (T) - True, (F) - False , ~~~ - use 'hardware-age' keyword to retrieve age info
VLAN/BD  MAC Address  Type  age  Secure NTFY Ports/SWID.SSID.LID
-----
G 10    68bd.abd7.6042  static  -    F  F  sup-eth1(R)
* 10    0000.0000.0010  dynamic  ~~~  F  F  Eth1/3
* 10    0000.0000.0021  dynamic  ~~~  F  F  Po7
* 10    00de.fb12.1a01  dynamic  ~~~  F  F  Po7
* 10    00de.fb12.1a7c  dynamic  ~~~  F  F  Eth1/3

```

Per the output below, N5K2 chose E1/6 as the Root Port to reach N7K1. This modification has been achieved by lowering the port cost on the link from N5K2 to N7K2, which results in a lower end-to-end path cost to reach N7K1 via N7K2.

```

N5K2# show spanning-tree vlan 10

VLAN0010
  Spanning tree enabled protocol rstp
  Root ID Priority 24586
  Address 68bd.abd7.6042
  Cost 1001
  Port 134 (Ethernet1/6)
  Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

  Bridge ID Priority 32778 (priority 32768 sys-id-ext 10)
  Address 00de.fb12.1a01
  Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Interface Role Sts Cost Prio.Nbr Type
-----
Po5 Desg FWD 1000 128.4100 P2p
Eth1/2 Desg FWD 2000 128.130 P2p
Eth1/5 Altn BLK 2000 128.133 P2p
Eth1/6 Root FWD 1 128.134 P2p

```

Note that in the above output we cannot see all the details of the path. Instead we can only see the local port costs as well as the total path cost. To see the received path cost from other neighbors, the `detail` keyword needs to be used for the `show spanning-tree` output.

```

N5K2# show spanning-tree vlan 10 detail

VLAN0010 is executing the rstp compatible Spanning Tree protocol
Bridge Identifier has priority 32768, sysid 10, address 00de.fb12.1a01
Configured hello time 2, fex hello time 12, max age 20, forward delay 15
Current root has priority 24586, address 68bd.abd7.6042
Root port is 134 (Ethernet1/6), cost of root path is 1001
Topology change flag not set, detected flag not set
Number of topology changes 20 last change occurred 0:20:42 ago
  from Ethernet1/6
Times: hold 1, topology change 35, notification 2
  hello 2, max age 20, forward delay 15
Timers: hello 0, topology change 0, notification 0

Port 4100 (port-channel5) of VLAN0010 is designated forwarding
  Port path cost 1000, Port priority 128, Port Identifier 128.4100
  Designated root has priority 24586, address 68bd.abd7.6042
  Designated bridge has priority 32778, address 00de.fb12.1a01
  Designated port id is 128.4100, designated path cost 1001
  Timers: message age 0, forward delay 0, hold 0
  Number of transitions to forwarding state: 4
  Link type is point-to-point by default
  BPDU: sent 3827, received 14

Port 130 (Ethernet1/2) of VLAN0010 is designated forwarding
  Port path cost 2000, Port priority 128, Port Identifier 128.130
  Designated root has priority 24586, address 68bd.abd7.6042
  Designated bridge has priority 32778, address 00de.fb12.1a01
  Designated port id is 128.130, designated path cost 1001
  Timers: message age 0, forward delay 0, hold 0
  Number of transitions to forwarding state: 5
  Link type is point-to-point by default
  BPDU: sent 6468, received 0

Port 133 (Ethernet1/5) of VLAN0010 is alternate blocking
  Port path cost 2000, Port priority 128, Port Identifier 128.133
  Designated root has priority 24586, address 68bd.abd7.6042
  Designated bridge has priority 24586, address 68bd.abd7.6042

```

```

Designated port id is 128.132, designated path cost 0, Topology change is set
Timers: message age 15, forward delay 0, hold 0
Number of transitions to forwarding state: 2
Link type is point-to-point by default
BPDU: sent 7, received 1037

Port 134 (Ethernet1/6) of VLAN0010 is root forwarding
Port path cost 1, Port priority 128, Port Identifier 128.134
Designated root has priority 24586, address 68bd.abd7.6042
Designated bridge has priority 28682, address 0026.980c.2142
Designated port id is 128.132, designated path cost 1000
Timers: message age 15, forward delay 0, hold 0
Number of transitions to forwarding state: 2
Link type is point-to-point by default
BPDU: sent 5, received 1036

```

In the above detailed output we can see the comparison of Eth1/5 with a total path cost of 2000 (port path cost 2000 + designated path cost 0) and Eth1/6 with a total path cost of 1001 (port path cost 1 + designated path cost 1000). Since  $1001 < 2000$ , Eth1/6 is chosen as the root port.

In the below output we see the detailed view from N5K2 of the VLAN 20 topology.

```

N5K2# show spanning-tree vlan 20 detail

VLAN0020 is executing the rstp compatible Spanning Tree protocol
Bridge Identifier has priority 28672, sysid 20, address 00de.fb12.1a01
Configured hello time 2, fex hello time 12, max age 20, forward delay 15
Current root has priority 24596, address 00de.fb12.1a7c
Root port is 133 (Ethernet1/5), cost of root path is 2
Topology change flag not set, detected flag not set
Number of topology changes 22 last change occurred 0:19:36 ago
    from Ethernet1/5
Times: hold 1, topology change 35, notification 2
    hello 2, max age 20, forward delay 15
Timers: hello 0, topology change 0, notification 0

Port 4100 (port-channel5) of VLAN0020 is alternate blocking
Port path cost 1000, Port priority 128, Port Identifier 128.4100
Designated root has priority 24596, address 00de.fb12.1a7c
Designated bridge has priority 24596, address 00de.fb12.1a7c
Designated port id is 128.4100, designated path cost 0, Topology change is set
Timers: message age 15, forward delay 0, hold 0
Number of transitions to forwarding state: 3
Link type is point-to-point by default
BPDU: sent 78, received 1094

Port 129 (Ethernet1/1) of VLAN0020 is designated forwarding
Port path cost 2000, Port priority 128, Port Identifier 128.129
Designated root has priority 24596, address 00de.fb12.1a7c
Designated bridge has priority 28692, address 00de.fb12.1a01
Designated port id is 128.129, designated path cost 2
Timers: message age 0, forward delay 0, hold 0
Number of transitions to forwarding state: 7
Link type is point-to-point by default
BPDU: sent 6559, received 0

Port 133 (Ethernet1/5) of VLAN0020 is root forwarding
Port path cost 1, Port priority 128, Port Identifier 128.133
Designated root has priority 24596, address 00de.fb12.1a7c
Designated bridge has priority 32788, address 68bd.abd7.6042
Designated port id is 128.132, designated path cost 1, Topology change is set
Timers: message age 16, forward delay 0, hold 0
Number of transitions to forwarding state: 2
Link type is point-to-point by default
BPDU: sent 504, received 632

Port 134 (Ethernet1/6) of VLAN0020 is designated forwarding
Port path cost 2000, Port priority 128, Port Identifier 128.134
Designated root has priority 24596, address 00de.fb12.1a7c
Designated bridge has priority 28692, address 00de.fb12.1a01
Designated port id is 128.134, designated path cost 2
Timers: message age 0, forward delay 0, hold 0
Number of transitions to forwarding state: 3
Link type is point-to-point by default
BPDU: sent 1090, received 40

```

For VLAN 20 traffic, N5K2 is comparing Port-Channel5 with a total path cost of 1000 (port path cost 1000 + designated path cost 0) and Eth1/5 with a total path cost of 2 (port path cost 1 + designated path cost 1). Note that this path modification also required a change on the upstream switch N7K2, which lowered the total path cost of Eth1/5 to less than Port-Channel5's.

[« Port Channels | Multiple Spanning Tree \(MST\) Traffic Engineering »](#)

[^ back to top](#)