Nexus 7000 Peer−Switch Configuration (Hybrid Setup)

Document ID: 116140

Contributed by Andy Gossett and Rajesh Gatti, Cisco TAC Engineers.
Aug 09, 2013

Contents

Introduction

This document describes how to configure peer−switch on the Cisco Nexus 7000 Series switches in order to allow non−virtual port channel (non−vPC) connections to load balance between VLANs.

When peer−switch is enabled, each Nexus 7000 switch shares a virtual bridge ID, which allows both switches to act as root for the VLAN. For devices with a connection to each Nexus 7000 switch in the vPC domain that are not capable of port channeling, the Layer 2 (L2) topology relies on Spanning Tree Protocol (STP) in order to block the redundant links. The peer−switch feature allows for pseudo−STP configurations to allow non−vPC connections to load balance STP states between the two Nexus 7000 switches. This document discusses in detail the reason for the pseudo−STP configurations and how they affect non−vPC and vPC links.

A mix of vPC and non−vPC links is called a hybrid setup.

The MAC addresses for each switch used in the configuration example in this document are:

- Nexus 7000 vPC Switch 1 (N7K−1): 00:24:98:6f:3b:41
- Nexus 7000 vPC Switch 2 (N7K−2): 00:24:98:6f:3b:42
- Non−vPC Switch 1 (SW−1): 00:24:98:6f:3b:44
- Non−vPC Switch 2 (SW−2): 00:24:98:6f:3b:43

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Spanning Tree Protocol (STP)
- Virtual port channel (vPC)

Components Used

The information in this document is based on the Cisco Nexus 7000 Series Switches with Supervisor 1 Module.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Configure

Note: Use the Command Lookup Tool (registered customers only) in order to obtain more information on the commands used in this section.

Note: The Output Interpreter Tool (registered customers only) supports certain ***show*** commands. Use the Output Interpreter Tool in order to view an analysis of ***show*** command output.

Normal vPC Behavior for Hybrid Setup

This is a network diagram of a hybrid setup without peer−switch enabled. Both Nexus 7000 switches are configured with a priority of 8192 for all VLANs. N7K−1 wins the bridge election because it has the lower bridge ID. Therefore, you expect SW−1 to block on the link from N7K−2. SW−2 is connected to the Nexus 7000 switches via a vPC and will be in a forwarding status. SW−2 receives Bridge Protocol Data Units (BPDUs) only from the primary switch in the vPC, which is N7K−1 in this example.

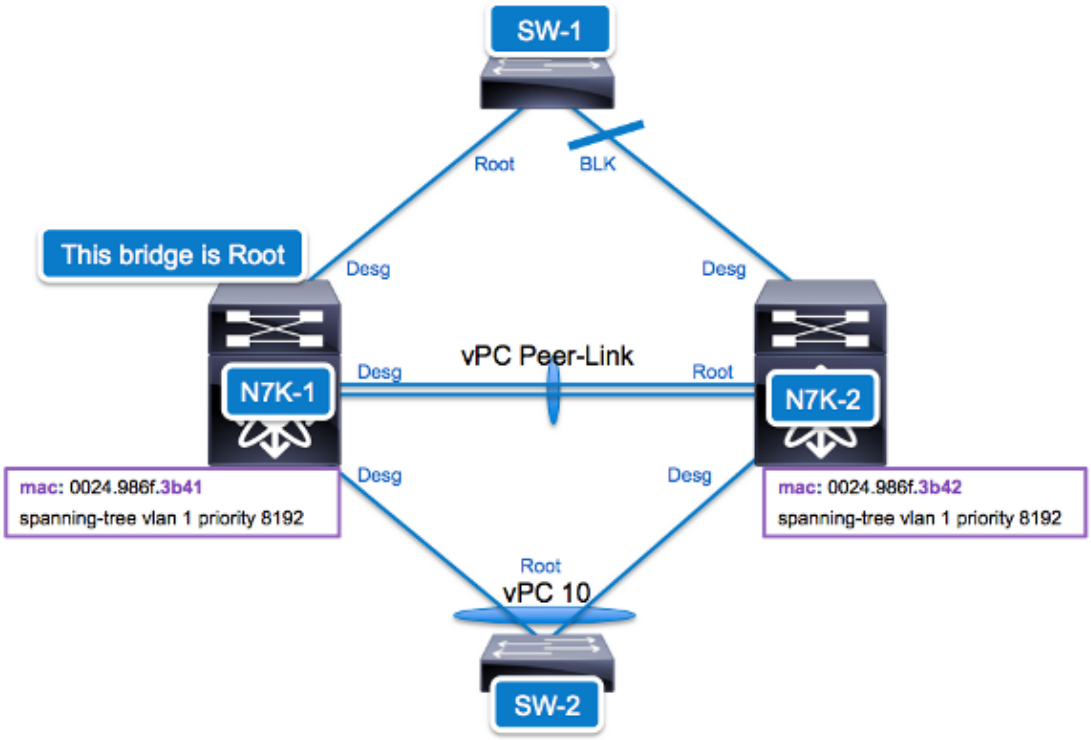


```
SW-1# show span vlan 1VLAN0001
  Spanning tree enabled protocol rstp
  Root ID     Priority    8193
              Address     0024.986f.3b41
              Cost        4
              Port        295 (Ethernet2/39)
              Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    32769  (priority 32768 sys-id-ext 1)
              Address     0024.986f.3b44
              Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Eth2/39          Root FWD 4         128.295  P2p
Eth2/40          Altn BLK 4         128.296  P2p

SW-1# show span vlan 1 detail

 VLAN0001 is executing the rstp compatible Spanning Tree protocol
  Bridge Identifier has priority 32768, sysid 1, address 0024.986f.3b44
  Configured hello time 2, max age 20, forward delay 15
  Current root has priority 8193, address 0024.986f.3b41
  Root port is 295 (Ethernet2/39), cost of root path is 4
  Topology change flag not set, detected flag not set
  Number of topology changes 4 last change occurred 0:29:13 ago
          from Ethernet2/39
  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 295 (Ethernet2/39) of VLAN0001 is root forwarding
   Port path cost 4, Port priority 128, Port Identifier 128.295
   Designated root has priority 8193, address 0024.986f.3b41
   Designated bridge has priority 8193, address 0024.986f.3b41
   Designated port id is 128.260, designated path cost 0, Topology change is set
   Timers: message age 16, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   Link type is point-to-point by default
```

```
   BPDU: sent 4, received 898

 Port 296 (Ethernet2/40) of VLAN0001 is alternate blocking
   Port path cost 4, Port priority 128, Port Identifier 128.296
   Designated root has priority 8193, address 0024.986f.3b41
   Designated bridge has priority 8193, address 0024.986f.3b42 <-- Although same priority,
   Designated port id is 128.272, designated path cost 2          advertising Bridge ID is higher
   Timers: message age 16, forward delay 0, hold 0                 and therefore this link is BLK
   Number of transitions to forwarding state: 2
   Link type is point-to-point by default
   BPDU: sent 6, received 895
```

Enable Peer−Switch on Both Nexus Switches

This is a network diagram of a a hybrid setup with peer−switch enabled. When peer−switch is enabled, each Nexus 7000 switch shares a virtual bridge ID which allows both switches to act as root for the VLAN. The vPC peer−link is always in a forwarding status and runs L2 Gateway Interconnection Protocol (L2GIP) in order to prevent bridging loops.

Each Nexus 7000 switch sends BPDUs with a root bridge identified by the virtual bridge ID. On vPC links, the designated bridge ID also uses the virtual bridge ID. For non−vPC links, the designated bridge ID is the physical bridge ID of the corresponding Nexus 7000 switch. This allows the non−vPC switch (SW−1) to make a root decision based upon BPDU advertisements instead of port priority.

Note: For proper behavior, VLAN priorities on both Nexus 7000 switches should be configured the same.

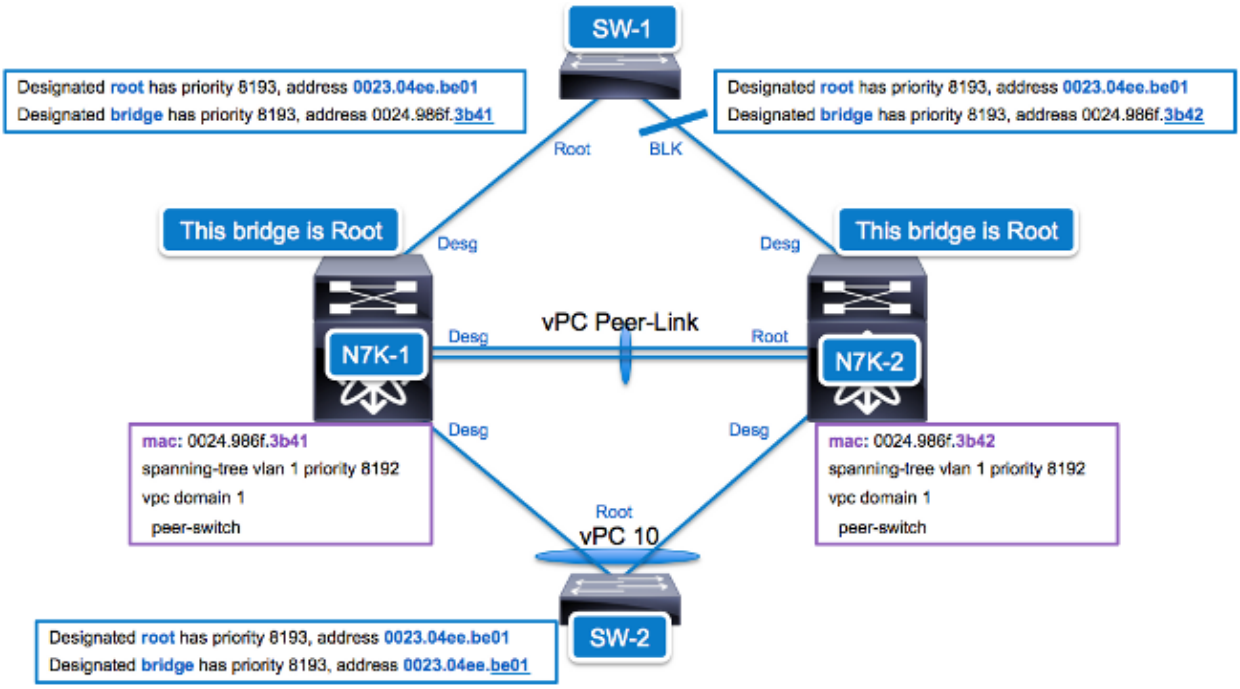

Non−vPC Connection

With peer−switch enabled, each Nexus 7000 switch generates BPDUs with the root bridge set to the virtual bridge ID and the designated bridge set to the physical bridge ID. Since the priorities are the same, all non−vPC connections always forward on the link connected to the Nexus 7000 switch with the lower bridge ID (N7K−1 in this example) and block on the links connected to the Nexus 7000 switch with the higher bridge ID (N7K−2 in this example).

```
SW-1# show span vlan 1
```

```
VLAN0001
  Spanning tree enabled protocol rstp
  Root ID    Priority    8193
             Address     0023.04ee.be01
             Cost        4
             Port        295 (Ethernet2/39)
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32769  (priority 32768 sys-id-ext 1)
             Address     0024.986f.3b44
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Eth2/39          Root FWD 4         128.295  P2p
Eth2/40          Altn BLK 4         128.296  P2p
SW-1# show span vlan 1 detail
 VLAN0001 is executing the rstp compatible Spanning Tree protocol
  Bridge Identifier has priority 32768, sysid 1, address 0024.986f.3b44
  Configured hello time 2, max age 20, forward delay 15
  Current root has priority 8193, address 0023.04ee.be01
  Root port is 295 (Ethernet2/39), cost of root path is 4
  Topology change flag not set, detected flag not set
  Number of topology changes 6 last change occurred 0:25:38 ago
          from Ethernet2/39
  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 295 (Ethernet2/39) of VLAN0001 is root forwarding
   Port path cost 4, Port priority 128, Port Identifier 128.295
   Designated root has priority 8193, address 0023.04ee.be01    <---Root Bridge = virtual ID
   Designated bridge has priority 8193, address 0024.986f.3b41  <---Designated Bridge ID = N7K-1
   Designated port id is 128.260, designated path cost 0, Topology change is set
   Timers: message age 16, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   Link type is point-to-point by default
   BPDU: sent 4, received 2280
 Port 296 (Ethernet2/40) of VLAN0001 is alternate blocking
   Port path cost 4, Port priority 128, Port Identifier 128.296
   Designated root has priority 8193, address 0023.04ee.be01    <---Root Bridge = virtual ID
   Designated bridge has priority 8193, address 0024.986f.3b42  <---Designated Bridge ID = N7K-2
   Designated port id is 128.272, designated path cost 0
   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 2278
```

vPC Connection

With peer-switch enabled, vPC connections receive BPDUs with both the root bridge and designated bridge set to the virtual bridge ID.

```
SW-2# show span vlan 1

VLAN0001
  Spanning tree enabled protocol rstp
  Root ID    Priority    8193
             Address     0023.04ee.be01
              Cost        3
             Port        4105 (port-channel10)
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32769  (priority 32768 sys-id-ext 1)
             Address     0024.986f.3b43
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec
```

```
Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Po10             Root FWD 3         128.4105 P2p

SW-2# show span vlan 1 detail

 VLAN0001 is executing the rstp compatible Spanning Tree protocol
  Bridge Identifier has priority 32768, sysid 1, address 0024.986f.3b43
  Configured hello time 2, max age 20, forward delay 15
  Current root has priority 8193, address 0023.04ee.be01
  Root port is 4105 (port-channel10), cost of root path is 3
  Topology change flag not set, detected flag not set
  Number of topology changes 5 last change occurred 0:21:40 ago
          from port-channel10
  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 4105 (port-channel10) of VLAN0001 is root forwarding
   Port path cost 3, Port priority 128, Port Identifier 128.4105
   Designated root has priority 8193, address 0023.04ee.be01          <--- Virtual Bridge ID
   Designated bridge has priority 8193, address 0023.04ee.be01       <--- Virtual Bridge ID
   Designated port id is 128.4105, 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 96, received 2804
```

Enable Load Balancing between VLANs on Non−vPC Links

Under default peer−switch configuration, all VLANs on the non−vPC switch are forwarding on a single link. In order to load balance between the VLANs, the designated and root priorities advertised can be manually set by use of spanning tree pseduo−information configurations. Cisco recommends that the root priority under the pseduo−information be lower than the best spanning tree priority in order to prevent topology change notifications (TCNs) under failover conditions. The designated priorities can be load balanced between the two Nexus 7000 switches in the vPC domain.

In this example, the global spanning tree priorities on both Nexus 7000 switches have been set to 8192. Under the pseudo−information, the root priority has been configured as 4096, which is lower than the best priority of 8192. Therefore, the switch that is participating with peer−switch enabled becomes the root for the VLAN. In order to load balance between the two switches, the designated priorities are alternated for VLAN 9 and VLAN 10. For the non−vPC connections to SW−1, VLAN 9 is forwarded on the link to N7K−1, and VLAN 10 is forwarded on the link to N7K−2.

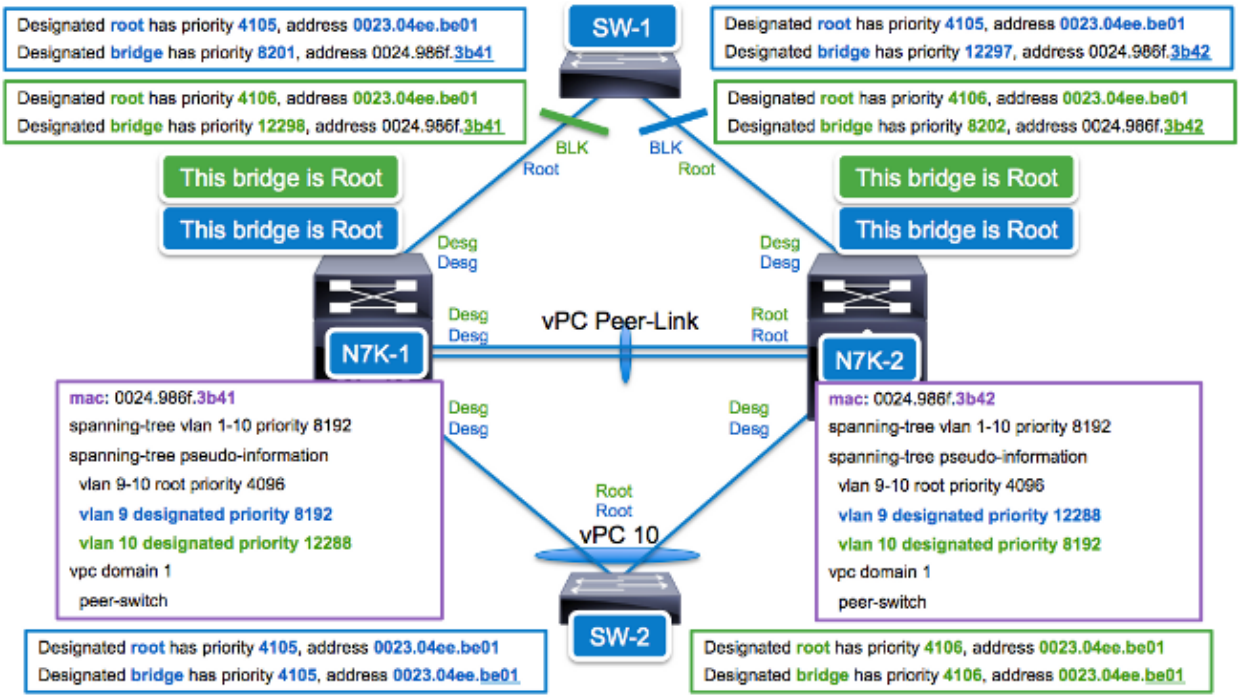

Non−vPC Connection

For VLAN 9, SW−1 sees the pseudo root bridge priority and bridge ID as the same value from both N7K−1 and N7K−2. However, both N7K−1 and N7K−2 send their configured pseudo designated priorities. Therefore, SW−1 sees the designated bridge priority of 8201 (8192 + 9) from N7K−1 and the designated bridge priority of 12297 (12288 + 9) from N7K−2; SW−1 chooses the link toward N7K−1 as the forwarding link on VLAN 9.

```
SW-1# show span vlan 9

VLAN0009
  Spanning tree enabled protocol rstp
  Root ID    Priority    4105
             Address     0023.04ee.be01
             Cost        4
             Port        295 (Ethernet2/39)
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32777  (priority 32768 sys-id-ext 9)
             Address     0024.986f.3b44
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Eth2/39          Root FWD 4         128.295  P2p
Eth2/40          Altn BLK 4         128.296  P2p

SW-1# show span vlan 9 detail

 VLAN0009 is executing the rstp compatible Spanning Tree protocol
  Bridge Identifier has priority 32768, sysid 9, address 0024.986f.3b44
  Configured hello time 2, max age 20, forward delay 15
  Current root has priority 4105, address 0023.04ee.be01
  Root port is 295 (Ethernet2/39), cost of root path is 4
  Topology change flag not set, detected flag not set
  Number of topology changes 16 last change occurred 0:06:56 ago
          from Ethernet2/39
  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 295 (Ethernet2/39) of VLAN0009 is root forwarding
  Port path cost 4, Port priority 128, Port Identifier 128.295
  Designated root has priority 4105, address 0023.04ee.be01     <--- Root Virtual Bridge ID
  Designated bridge has priority 8201, address 0024.986f.3b41   <--- Designated N7K-1, 8201
  Designated port id is 128.260, designated path cost 0
  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 31, received 3486

Port 296 (Ethernet2/40) of VLAN0009 is alternate blocking
  Port path cost 4, Port priority 128, Port Identifier 128.296
  Designated root has priority 4105, address 0023.04ee.be01    <--- Root Virtual Bridge ID
  Designated bridge has priority 12297, address 0024.986f.3b42 <--- Designated is N7K-2, 12297
  Designated port id is 128.272, designated path cost 0
  Timers: message age 15, forward delay 0, hold 0
  Number of transitions to forwarding state: 4
  Link type is point-to-point by default
  BPDU: sent 31, received 3496
```

Similarly for VLAN 10, SW−1 sees the pseudo root bridge priority and bridge ID as the same value from both N7K−1 and N7K−2. Again, both N7K−1 and N7K−2 send their configured pseudo designated priorities. For VLAN 10, SW−1 sees the designated bridge priority of 12298 (12288 + 10) from N7K−1 and the designated bridge priority of 8202 (8192 + 10) from N7K−2; SW−1 chooses the link toward N7K−2 as the forwarding link for VLAN 10. In this way, the non−vPC connected switches can load balance VLAN STP state between N7K−1 and N7K−2.

```
SW-1# show span vlan 10 detail

 VLAN0010 is executing the rstp compatible Spanning Tree protocol
  Bridge Identifier has priority 32768, sysid 10, address 0024.986f.3b44
  Configured hello time 2, max age 20, forward delay 15
  Current root has priority 4106, address 0023.04ee.be01
  Root port is 296 (Ethernet2/40), cost of root path is 4
  Topology change flag not set, detected flag not set
  Number of topology changes 7 last change occurred 0:07:13 ago
          from Ethernet2/40
  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 295 (Ethernet2/39) of VLAN0010 is alternate blocking
   Port path cost 4, Port priority 128, Port Identifier 128.295
   Designated root has priority 4106, address 0023.04ee.be01     <--- Root Virtual Bridge ID
   Designated bridge has priority 12298, address 0024.986f.3b41  <--- Designated N7K-1, 12298
   Designated port id is 128.260, designated path cost 0, Topology change is set
   Timers: message age 16, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   Link type is point-to-point by default
   BPDU: sent 4, received 3497

 Port 296 (Ethernet2/40) of VLAN0010 is root forwarding
    Port path cost 4, Port priority 128, Port Identifier 128.296
   Designated root has priority 4106, address 0023.04ee.be01    <--- Root Virtual Bridge ID
   Designated bridge has priority 8202, address 0024.986f.3b42  <--- Designated N7K-2, 8202
   Designated port id is 128.272, designated path cost 0
   Timers: message age 16, forward delay 0, hold 0
   Number of transitions to forwarding state: 3
   Link type is point-to-point by default
   BPDU: sent 10, received 3492
```

vPC Connection

For vPC links, the root and designated fields use the pseudo root priority and virtual bridge ID, respectively.

```
SW-2# show span vlan 9

VLAN0009
  Spanning tree enabled protocol rstp
  Root ID    Priority    4105
             Address     0023.04ee.be01
             Cost        3
             Port        4105 (port-channel10)
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32777  (priority 32768 sys-id-ext 9)
             Address     0024.986f.3b43
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Po10             Root FWD 3         128.4105 P2p

SW-2# show span vlan 10

VLAN0010
  Spanning tree enabled protocol rstp
  Root ID    Priority    4106
             Address     0023.04ee.be01
             Cost        3
             Port        4105 (port-channel10)
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32778  (priority 32768 sys-id-ext 10)
             Address     0024.986f.3b43
             Hello Time  2  sec  Max Age 20 sec  Forward Delay 15 sec

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Po10             Root FWD 3         128.4105 P2p

SW-2#show span vlan 9 detail

 VLAN0009 is executing the rstp compatible Spanning Tree protocol
  Bridge Identifier has priority 32768, sysid 9, address 0024.986f.3b43
  Configured hello time 2, max age 20, forward delay 15
  Current root has priority 4105, address 0023.04ee.be01
  Root port is 4105 (port-channel10), cost of root path is 3
  Topology change flag not set, detected flag not set
  Number of topology changes 12 last change occurred 0:04:29 ago
          from port-channel10
  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 4105 (port-channel10) of VLAN0009 is root forwarding
   Port path cost 3, Port priority 128, Port Identifier 128.4105
   Designated root has priority 4105, address 0023.04ee.be01    <--- Root Virtual Bridge ID
   Designated bridge has priority 4105, address 0023.04ee.be01  <--- Root Virtual Bridge ID
   Designated port id is 128.4105, 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 119, received 4867

SW-2# show span vlan 10 detail
```

```
VLAN0010 is executing the rstp compatible Spanning Tree protocol
 Bridge Identifier has priority 32768, sysid 10, address 0024.986f.3b43
 Configured hello time 2, max age 20, forward delay 15
 Current root has priority 4106, address 0023.04ee.be01
 Root port is 4105 (port-channel10), cost of root path is 3
 Topology change flag not set, detected flag not set
 Number of topology changes 6 last change occurred 0:04:36 ago
         from port-channel10
 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 4105 (port-channel10) of VLAN0010 is root forwarding
  Port path cost 3, Port priority 128, Port Identifier 128.4105
  Designated root has priority 4106, address 0023.04ee.be01    <--- Root Virtual Bridge ID
  Designated bridge has priority 4106, address 0023.04ee.be01  <--- Root Virtual Bridge ID
  Designated port id is 128.4105, designated path cost 0, Topology change is set
  Timers: message age 17, forward delay 0, hold 0
  Number of transitions to forwarding state: 2
  Link type is point-to-point by default
  BPDU: sent 96, received 5179
```

Caveats

See Cisco bug ID CSCub74914: Pseudo STP priorities incorrectly set on vPC links in peer−switch setup

Verify

There is currently no verification procedure available for this configuration.

Troubleshoot

There is currently no specific troubleshooting information available for this configuration.

Related Information

- ***Cisco Nexus 7000 Series NX−OS Interface Configuration Guide, Release 5.x: Configuring vPCs: vPC Peer Switch***
- ***Design and Configuration Guide: Best Practices for Virtual Port Channels (vPC) on Cisco Nexus 7000 Series Switches***
- ***Technical Support & Documentation − Cisco Systems***

Updated: Aug 09, 2013 Document ID: 116140