

Nexus 7000 STP Priority Change in Peer Switch Setup Impact Assessment and Configuration Example



Document ID: 118858

Contributed by Rajesh Gatti, Cisco TAC Engineer.
Apr 09, 2015

Contents

Introduction

Prerequisites

- Requirements
- Components Used

Configure

- What is the impact when you follow the Cisco best practice?
- Conclusion
- Caveat
- Known Bugs Related to Peer Switch

Verify

Troubleshoot

Related Information

Introduction

This document describes to customers the impact of standardization of the Virtual Port Channel (vPC) peer switch configuration in scenarios where it does not follow recommendations, such as mismatched Spanning Tree Protocol (STP) priorities.

The peer switch feature allows a pair of Cisco Nexus 7000 Series devices to appear as a single STP root in the Layer 2 topology. This feature eliminates the need to pin the STP root to the vPC primary switch and improves vPC convergence if the vPC primary switch fails. The values that you apply for the spanning tree priority must be identical on both vPC peers.

There have been deployments where the vPC peer switch configuration had mismatched STP priorities in the production environment.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Nexus 7000 architecture
- vPC feature

Components Used

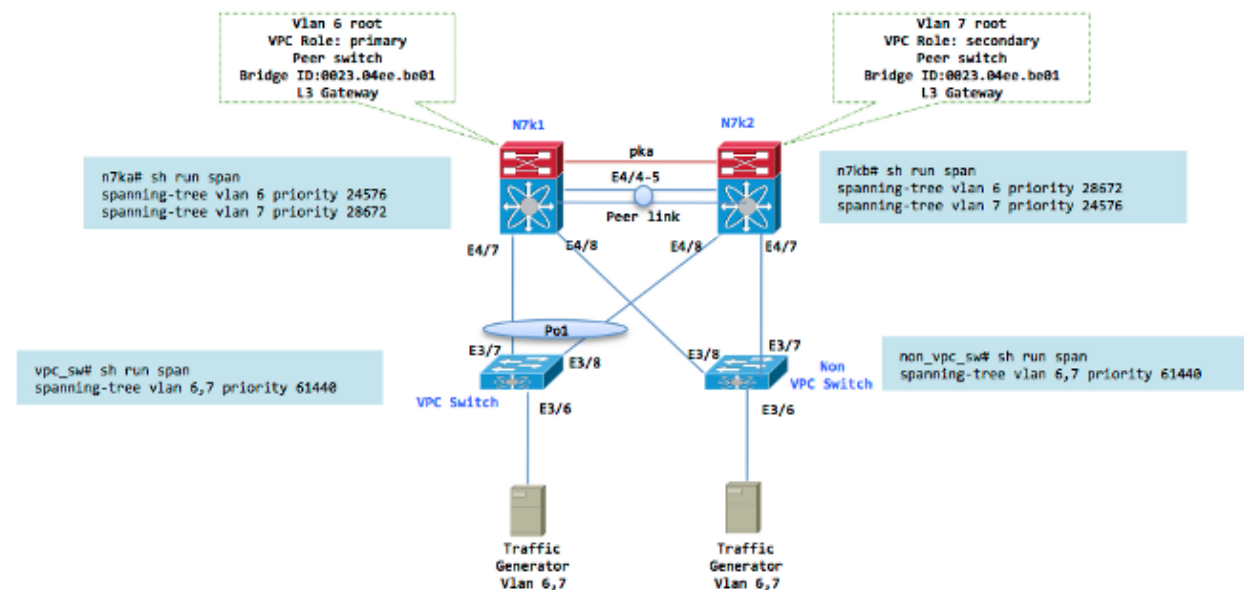
The information in this document is based on these software and hardware versions:

- Nexus 7000 with Release 6.2.10 and later
- M1/F2 series line card
- Rapid Spanning Tree Protocol (RSTP) is deployed on all switches

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

In this network diagram there is a simple vPC setup defined by a pair of Nexus 7000s. The downstream access switches are configured to be part of a vPC and nonvPC respectively. The traffic generator pushes intraVLAN and interVLAN traffic through the network.



Notice that the STP priorities are different for the vPC VLANs defined even though the switches operate in peer switch mode.

The outputs shown here are for VLAN 6.

```
n7ka# show span vlan 6
```

```
VLAN0006
Spanning tree enabled protocol rstp
Root ID    Priority    24582
           Address    0023.04ee.be01
           This bridge is the root
           Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
Bridge ID  Priority    24582 (priority 24576 sys-id-ext 6)
           Address    0023.04ee.be01
           Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
Interface   Role Sts Cost      Prio.Nbr Type
-----
Po1         Desg FWD 1         128.4096 (vPC) P2p
Po3         Desg FWD 1         128.4098 (vPC peer-link) Network P2p
```

Eth4/8 Desg FWD 2 128.520 P2p

n7kb# **show span vlan 6**

```
VLAN0006
Spanning tree enabled protocol rstp
Root ID    Priority    24582
           Address    0023.04ee.be01
           Cost        1
           Port        4098 (port-channel3)
           Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
Bridge ID Priority    28678 (priority 28672 sys-id-ext 6)
           Address    0023.04ee.be01
           Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
Interface    Role Sts Cost        Prio.Nbr Type
-----
Po1            Desg FWD 1            128.4096 (vPC) P2p
Po3            Root FWD 1            128.4098 (vPC peer-link) Network P2p
Eth4/7         Desg FWD 2            128.519 P2p
```

vpc_sw# **show span vlan 6**

```
VLAN0006
Spanning tree enabled protocol rstp
Root ID    Priority    24582
           Address    0023.04ee.be01
           Cost        1
           Port        4096 (port-channel1)
           Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
Bridge ID Priority    61446 (priority 61440 sys-id-ext 6)
           Address    6c9c.ed4e.6f43
           Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
Interface    Role Sts Cost        Prio.Nbr Type
-----
Po1            Root FWD 1            128.4096 P2p
```

non_vpc_sw# **show span vlan 6**


```
VLAN0006
Spanning tree enabled protocol rstp
Root ID    Priority    24582
           Address    0023.04ee.be01
           Cost        2
           Port        392 (Ethernet3/8)
           Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
Bridge ID Priority    61446 (priority 61440 sys-id-ext 6)
           Address    0022.557a.4343
           Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
Interface    Role Sts Cost        Prio.Nbr Type
-----
Eth3/7         Altn BLK 2            128.391 P2p
Eth3/8         Root FWD 2            128.392 P2p
```

Even though the configuration does not match the recommended Cisco best practice, there are no issues with traffic forwarding between the VLANs (intraVLAN or interVLAN).

What is the impact when you follow the Cisco best practice?

In a production environment, if the need arises to make the STP priority the same on both vPC peers, the first question that is asked is what is the impact.

In the topology shown here, changes were made to the STP priority for VLAN 6 and 7 on both Nexus 7000s. Since in a peer switch setup, both Nexus 7000s independently generate the Bridge Protocol Data Unit (BPDU), the change causes one of the Nexus 7000s to advertise the BPDU with the same attribute as its peer.



A nonVPC switch could now receive a superior BPDU from a different path which could cause a change in the forwarding state of a specific link. The state change is shown in the next sample output where e3/7 has gone into the forwarding state after the peer switch change. Since RSTP is in use, this state change is subsecond. However, it results in generation of a Topology Change Notification (TCN).

```
non_vpc_sw# show span vlan 6
VLAN0006
Spanning tree enabled protocol rstp
Root ID Priority 24582
      Address 0023.04ee.be01
      Cost 2
      Port 391 (Ethernet3/7)
      Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
Bridge ID Priority 61446 (priority 61440 sys-id-ext 6)
      Address 0022.557a.4343
      Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
Interface Role Sts Cost Prio.Nbr Type
-----
Eth3/7 Root FWD 2 128.391 P2p => E3/7 was in Altn BLK state before
Eth3/8 Altn BLK 2 128.392 P2p
```

```
non_vpc_sw# show span vlan 6 det
VLAN0006 is executing the rstp compatible Spanning Tree protocol
Bridge Identifier has priority 61440, sysid 6, address 0022.557a.4343
Configured hello time 2, max age 20, forward delay 15
Current root has priority 24582, address 0023.04ee.be01
Root port is 391 (Ethernet3/7), cost of root path is 2
Topology change flag not set, detected flag not set
Number of topology changes 14 last change occurred 0:01:37 ago
      from Ethernet3/7
Times: hold 1, topology change 35, notification 2
      hello 2, max age 20, forward delay 15
Timers: hello 0, topology change 0, notification 0
```

```
n7ka# show span vlan 6 det
VLAN0006 is executing the rstp compatible Spanning Tree protocol
Bridge Identifier has priority 24576, sysid 6, address 0023.04ee.be01
Configured hello time 2, max age 20, forward delay 15
We are the root of the spanning tree
Topology change flag set, detected flag not set
Number of topology changes 28 last change occurred 0:01:37 ago
```

```
    from port-channel3  
Times: hold 1, topology change 35, notification 2  
      hello 2, max age 20, forward delay 15  
Timers: hello 0, topology change 18, notification 0
```

```
n7kb# show span vlan 6 det  
VLAN0006 is executing the rstp compatible Spanning Tree protocol  
Bridge Identifier has priority 24576, sysid 6, address 0023.04ee.be01  
Configured hello time 2, max age 20, forward delay 15  
We are the root of the spanning tree  
Topology change flag set, detected flag not set  
Number of topology changes 20 last change occurred 0:01:37 ago
```

```
    from Ethernet4/7  
Times: hold 1, topology change 35, notification 2  
      hello 2, max age 20, forward delay 15  
Timers: hello 0, topology change 28, notification 0
```

```
vpc_sw# show span vlan 6 det  
VLAN0006 is executing the rstp compatible Spanning Tree protocol  
Bridge Identifier has priority 61440, sysid 6, address 6c9c.ed4e.6f43  
Configured hello time 2, max age 20, forward delay 15  
Current root has priority 24582, address 0023.04ee.be01  
Root port is 4096 (port-channel1), cost of root path is 1  
Topology change flag not set, detected flag not set  
Number of topology changes 23 last change occurred 0:01:37 ago
```

```
    from port-channel1  
Times: hold 1, topology change 35, notification 2  
      hello 2, max age 20, forward delay 15  
Timers: hello 0, topology change 0, notification 0
```

The single TCN causes a flush of the MAC table, but as long as traffic (unicast) is bidirectional there is no interruption in the traffic flow. This has been tested in the lab with both intraVLAN and interVLAN traffic.

Note: If your setup consists purely of vPC switches, then there are no TCNs generated as there is no change in the STP states for downstream switches. There will not be any impact to the traffic flow.

Conclusion

Standardization of the STP priority in order to follow Cisco best practice has no impact to the traffic flow based on the lab test.

Caveat

The lab test only introduces a limited set of variables which might not match a real-world environment from a complexity point of view. Cisco advises that you ensure that these changes are implemented in change windows in order to avoid any surprises.

Known Bugs Related to Peer Switch

- CSCut31625 – Enhancement: Peer-switch exclude vlan for non-root vlans
- CSCuq57422 – vPC: Peer-Switch is not Supported on Non-Root Peers
- CSCub74914 – Pseudo STP priorities incorrectly set on vPC links in peer-switch setup
- CSCuf35758 – N7K:Peer switch feature conflict for non vpc vlans

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

- *Best Practices for Virtual Port Channels (vPC) on Cisco Nexus 7000 Series Switches*
- *Technical Support & Documentation – Cisco Systems*

Updated: Apr 09, 2015

Document ID: 118858
