PVST Simulation on MST Switches

Document ID: 116464

Contributed by Aninda Chatterjee, Cisco TAC Engineer.

Nov 21, 2013

Contents

Introduction

This document describes the purpose and functionality of Per VLAN Spanning Tree (PVST) simulation on Multiple Spanning Tree (MST) switches. It also addresses the basic rules that must be followed in order to avoid PVST simuation inconsistencies and the reason for these inconsistencies.

Prerequisites

Requirements

Cisco recommends that you have basic knowledge of MST concepts, such as Common and Internal Spanning Tree (CIST) and boundary ports.

Components Used

This document is not restricted to specific software and hardware versions.

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.

Background Information

Often, MST regions are connected to other domains − Per VLAN Spanning Tree Plus (PVST+) or rapid−PVST+ regions. These switches that run PVST+ (or rapid) cannot process MST−type Bridge Protocol Data Units (BPDUs). For this reason, there must be a backward−compatibility mechanism that runs so that these two domains can interact with each other seamlessly. This is what PVST simulation addresses and

achieves.

This simulation must be run only on boundary ports − these are ports that are directly connected to the PVST+ domain switches. The receipt of a Shared Spanning Tree Protocol (SSTP) BPDU on the port of a switch that runs MST causes the PVST simulation mechanism to trigger.

Topology

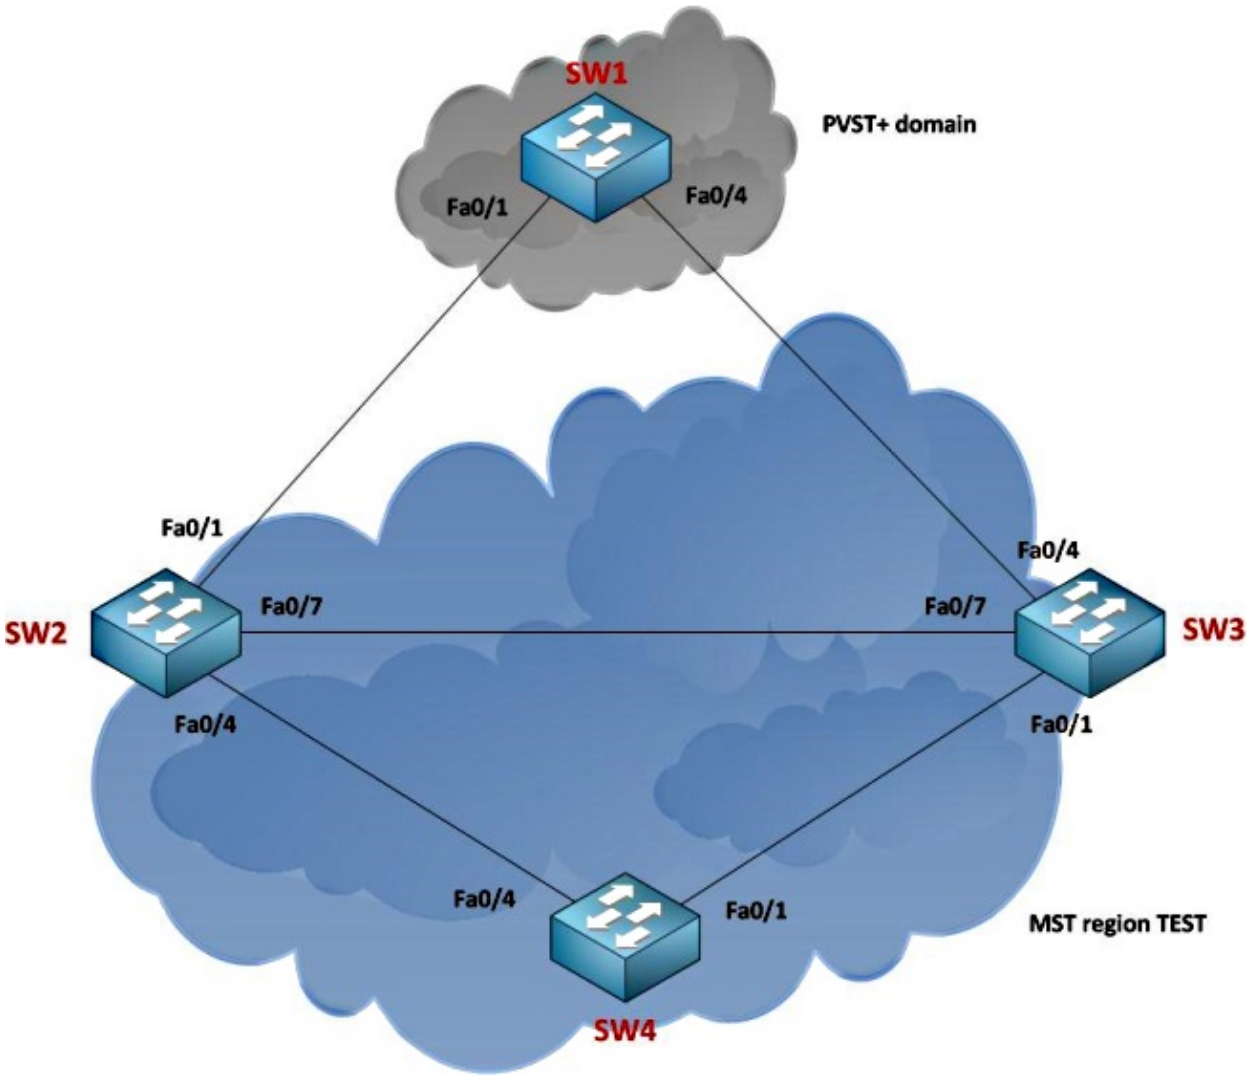

Basic Configuration on MST Switches

In this topology, Switch 1 (SW1) runs PVST+, while switches SW2, SW3, and SW4 run MST and are all in the same region.

MST Configurations on SW2, SW3, and SW4

```
SW2#show spanning-tree mst configuration
Name      [TEST]
Revision  1     Instances configured 2
Instance  Vlans mapped
--------  ----------------------------------------------
0         1
1         2-4094
---------------------------------------------------------
```

```
SW3#show spanning-tree mst configuration
Name      [TEST]
Revision  1     Instances configured 2
Instance  Vlans mapped
--------  ---------------------------------------------------
0         1
1         2-4094
-------------------------------------------------------------

SW4#show spanning-tree mst configuration
Name      [TEST]
Revision  1     Instances configured 2
Instance  Vlans mapped
--------  ---------------------------------------------------
0         1
1         2-4094
-------------------------------------------------------------
```

PVST Simulation

With such a topology (a mixture of MST and non−MST regions), the root bridge of CIST is in one of two places:

- Within a MST region
- Within a non−MST region.

PVST simulation runs seamlessly with two critical rules:

- If the root bridge for CIST is within a non−MST region, the spanning−tree priority of VLANs 2 and above within that domain must be better (lesser) than that of VLAN 1.

- If the root bridge for CIST is within a MST region, VLANs 2 and above defined in the non−MST domains must have their spanning−tree priorities worse (greater) than that of the CIST root.

If you do not adhere to these two rules, you encounter the ***PVST simulation failure***. These two rules, in a way, are identical to the root−guard feature and are actually derived from it.

The next sections examine the rules (scenarios) individually in order explain how PVST simulation works.

Scenario 1: The Root Bridge for CIST is in the PVST+ Domain

In this scenario, SW1 is the root. Here is its configuration:

```
spanning-tree vlan 1 priority 8192
spanning-tree vlan 2-4094 priority 4096
```

SW2 has this configuration:

```
spanning-tree mst 0 priority 12288
spanning-tree mst 1 priority 0
```

SW3 has this configuration:

```
spanning-tree mst 0 priority 16384
```

SW4 has this configuration:

```
spanning-tree mst 0 priority 16384
```

SW1 does not hear any BPDUs that it can understand, so it elects itself as the root for all VLANs, and starts to send BPDUs toward the MST−region switches. When SW2 receives a SSTP BPDU on Fa0/1, it understands that the interface is connected to a PVST+ domain. It subsequently sets the flag in order to enable PVST simulation on this interface.

A critical concept to understand is that ***only the Institute of Electrical and Electronics Engineers (IEEE) BPDU for VLAN 1 is processed for the root bridge election***. This is compared to ***only the instance 0 information from the MST region***. No other instance information is used in order to elect the root bridge for CIST. No other VLAN information from the PVST+ domain other than VLAN 1 is used in order to elect the CIST root bridge.

A question arises here of what happens with the other BPDUs. SW1 allows these VLANs across its trunk link to SW2:

```
SW1#show interfaces fa0/1 trunk
Port        Mode             Encapsulation  Status        Native vlan
Fa0/1       on               802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/1       1-4094
Port        Vlans allowed and active in management domain
Fa0/1       1-2,10,17,29,34,38,45,56,67,89,100,200,300,333,500,666,999
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/1       1-2,10,17,29,34,38,45,56,67,89,100,200,300,333,500,666,999
```

SW1 generates one BPDU for every VLAN, and sends them to SW2. These BPDUs are simply used for consistency checks as part of the PVST simulation. However, their information is not copied anywhere.

```
SW1#show spanning-tree vlan 1
VLAN0001
  Spanning tree enabled protocol ieee
  Root ID    Priority    8193
             Address     0022.0dba.9d00
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
  Bridge ID  Priority    8193   (priority 8192 sys-id-ext 1)
             Address     0022.0dba.9d00
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300
Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Fa0/1               Desg FWD 19        128.3    P2p
Fa0/4               Desg FWD 19        128.6    P2p


SW2#show spanning-tree mst 0
##### MST0    vlans mapped:   1
Bridge        address 0022.916d.5380  priority      12288 (12288 sysid 0)
Root          address 0022.0dba.9d00  priority      8193  (8192 sysid 1)
              port    Fa0/1           path cost     200000
Regional Root this switch
Operational   hello time 2 , forward delay 15, max age 20, txholdcount 6
Configured    hello time 2 , forward delay 15, max age 20, max hops    20
Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/1            Root FWD 200000    128.3    P2p Bound(PVST)
Fa0/4            Desg FWD 200000    128.6    P2p
Fa0/7            Desg FWD 200000    128.9    P2p
```

These outputs show that Fa0/1 of SW2 is elected as the root port. As discussed earlier, SW1 sends one BPDU

per VLAN for every VLAN allowed across its trunk link. This is confirmed from a debug on SW1:

```
STP: VLAN0001 Fa0/1 tx BPDU: config protocol=ieee
Data &colon; 0000 00 00 00 200100220DBA9D00 00000000 200100220DBA9D00 8003
0000 1400
STP: VLAN0010 Fa0/1 tx BPDU: config protocol=ieee
Data &colon; 0000 00 00 00 100A00220DBA9D00 00000000 100A00220DBA9D00 8003
0000 1400 0200 0F00
STP: VLAN0017 Fa0/1 tx BPDU: config protocol=ieee
Data &colon; 0000 00 00 00 101100220DBA9D00 00000000 101100220DBA9D00 8003
0000 1400 0200 0F00

*snip*
```

When these BPDUs arrive on SW2, the VLAN 1 BPDU is processed, which is reflected in the outputs. The other BPDUs then go through the PVST simulations root−guard−based consistency check.

In this setup, the consistency check passes and there is no PVST simulation failure. In order to generate a failure, increase the priority of VLAN 2 to greater than 8192 on SW1.

```
SW1#conf t
SW1(config)#spanning−tree vlan 2 priority 12288
```

This message displays on SW2:

```
%SPANTREE−2−PVSTSIM_FAIL: Blocking root port Fa0/1: Inconsitent inferior PVST
BPDU received on VLAN 2, claiming root 12290:0022.0dba.9d00
```

Here is what was stored on Fa0/1 of SW2 as root bridge information:

```
SW2#show spanning−tree interface fa0/1 detail
 Port 3 (FastEthernet0/1) of MST0 is broken  (PVST Sim. Inconsistent)
   Port path cost 200000, Port priority 128, Port Identifier 128.3.
   Designated root has priority 8193, address 0022.0dba.9d00
   Designated bridge has priority 8193, address 0022.0dba.9d00
   Designated port id is 128.3, designated path cost 0
   Timers: message age 4, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   Link type is point−to−point by default, Boundary PVST
   BPDU: sent 100, received 4189
```

The information that comes from SW1 is ***12290:0022.0dba.9d00***, and this is compared to ***8193.0022.0dba.9d00***. Since the port is a root port, and it has received an inferior BPDU, it enters into a PVST simulation failure state and displays the error message seen previously. This is because the boundary port cannot be in two different states at once − the receipt of the inferior BPDU dictates that the port should move to designated, whereas via VLAN 1 information dictates that the port should remain a root port. This confusion is prevented with PVST simulation. The port is also moved to a PVST simulation inconsistent state.

```
SW2#show spanning−tree
MST0
  Spanning tree enabled protocol mstp
  Root ID    Priority    8193
             Address     0022.0dba.9d00
             Cost        200000
             Port        3 (FastEthernet0/1)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
  Bridge ID  Priority    12288  (priority 12288 sys−id−ext 0)
             Address     0022.916d.5380
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

Interface           Role Sts Cost      Prio.Nbr Type
```

```
------------------ ---- --- --------- -------- -------------------------------
Fa0/1              Root BKN*200000    128.3    P2p Bound(PVST) *PVST_Inc
Fa0/4              Desg FWD 200000    128.6    P2p
Fa0/7              Desg FWD 200000    128.9    P2p
```

Scenario 2: The Root Bridge for CIST is in the MST Region

This situation reverses the roles from the previous scenario. The root bridge for the CIST is now in the MST region. SW2 is the root bridge.

```
SW2#show spanning-tree mst 0
##### MST0    vlans mapped:   1
Bridge        address 0022.916d.5380  priority      12288 (12288 sysid 0)
Root          this switch for the CIST
Operational   hello time 2 , forward delay 15, max age 20, txholdcount 6
Configured    hello time 2 , forward delay 15, max age 20, max hops    20<

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/1            Desg FWD 200000    128.3    P2p Bound(PVST)
Fa0/4            Desg FWD 200000    128.6    P2p
a0/7            Desg FWD 200000    128.9    P2p
```

Fa0/1 is still the boundary port, and the PVST simulation runs on this interface. This now plays a very important role again. ***The PVST+ domain expects one BPDU per VLAN, but MST does not do that.*** PVST simulation takes the bridge information of instance 0 (priority + MAC address), and it creates one BPDU for every VLAN that is allowed across its interface with this information. It simply tags each of these BPDUs with the appropriate VLAN IDs.

This can be verified with a debug on SW1:

```
STP: VLAN0001 rx BPDU: config protocol = ieee, packet from FastEthernet0/1  ,
linktype IEEE_SPANNING , enctype 2, encsize 17
STP: enc 01 80 C2 00 00 00 00 22 91 6D 53 83 00 26 42 42 03
STP: Data     000000000030000022916D5380000000003000022916D538080030000140002
000F00
STP: VLAN0001 Fa0/1:0000 00 00 00 30000022916D5380 00000000 30000022916D5380
8003 0
STP: VLAN0002 rx BPDU: config protocol = ieee, packet from FastEthernet0/1  ,
linktype SSTP , enctype 3, encsize 22STP: enc 01 00 0C CC CC CD 00 22 91 6D 53
83 00 32 AA AA 03 00 00 0C 01 0B
STP: Data     000000000030000022916D5380000000003000022916D538080030000140002
000F00
STP: VLAN0002 Fa0/1:0000 00 00 00 30000022916D5380 00000000 30000022 916D5380
8003 0000 1400 0200 0F00

STP: VLAN0010 rx BPDU: config protocol = ieee, packet from FastEthernet0/1  ,
linktype SSTP , enctype 3, encsize 22
STP: enc 01 00 0C CC CC CD 00 22 91 6D 53 83 00 32 AA AA 03 00 00 0C 01 0B
STP: Data     000000000030000022916D5380000000003000022916D538080030000140002
000F00
STP: VLAN0010 Fa0/1:0000 00 00 00 30000022916D5380 00 000000 30000022916D5380
8003 0000 1400 0200 0F00
```

In order to generate a failure condition for this, change the priority for VLAN 2 on SW1 to a value lower than 12,288.

```
SW1#conf t
SW1(config)#spanning-tree vlan 2 priority 8192
```

Here is the output on SW2:

```
%SPANTREE-2-PVSTSIM_FAIL: Blocking designated port Fa0/1: Inconsitent superior PVST
BPDU received on VLAN 2, claiming root 8194:0022.0dba.9d00
```

The information that comes from SW1 is ***8192:0022.0dba.9d00,*** and this is compared to ***12288:0022.916d.5380***. Since the port is a designated port, and it received a superior BPDU, it enters into a PVST simulation failure state and displays the previous error message. The port is also moved into a PVST simulation inconsistent state.

```
SW2#show spanning-tree mst 0
##### MST0    vlans mapped:   1
Bridge        address 0022.916d.5380  priority      12288 (12288 sysid 0)
Root          this switch for the CIST
Operational   hello time 2 , forward delay 15, max age 20, txholdcount 6
Configured    hello time 2 , forward delay 15, max age 20, max hops    20

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Fa0/1            Desg BKN*200000    128.3    P2p Bound(PVST) *PVST_Inc
Fa0/4            Desg FWD 200000    128.6    P2p
Fa0/7            Desg FWD 200000    128.9    P2p
```

Summary

PVST simulation is run on boundary ports and works in two ways:

- If the MST region has the root bridge for CIST, PVST simulation is required in order to replicate instance 0 information, and create one BPDU for every VLAN that is allowed across the trunk and tag it with the appropriate VLAN information.
- If the root bridge for CIST is outside of the MST region, then PVST simulation is required to process VLAN 1 information only. The other BPDUs (VLANs 2 and above) are used for consistency checks and information from these VLANs is never copied as root bridge information.

For PVST simulation to work without failures, these two conditions must be met:

- If the root bridge for CIST is within a non−MST region, the spanning−tree priority of VLANs 2 and above within that domain must be better (lesser) than that of VLAN 1.
- If the root bridge for CIST is within a MST region, VLANs 2 and above defined in the non−MST domains must have their spanning−tree priorities worse (greater) than that of the CIST root.

If these conditions are not met, the boundary port is put into a PVST simulation inconsistent state until the problem is corrected.

Updated: Nov 21, 2013 Document ID: 116464