

Dell-Cisco Spanning Tree Interoperability Testing and Recommendations

Victor Lama
Fabric Specialist – G500
April 2011

DOCUMENT REVISIONS:

Initial Release v1.0 – April 9, 2011

Revised Release v2.0 – Added Switch Configuration Details and revised Overview section – April 12, 2011

OVERVIEW:

The following lab test was conducted in response to client questions and requests for clarification regarding Spanning Tree Protocol (STP) interoperability between Dell and Cisco switches. More specifically, the focus of the testing was on the expected failover and reconvergence times in such a mixed environment when a link failure occurs. This testing is **not** formally sanctioned by Dell Engineering. This is a personal initiative.

Since the most commonly deployed version of Spanning Tree is the Rapid Spanning Tree Protocol, the Cisco switches in this lab were configured for rapid-PVST+ (Rapid per-VLAN Spanning-Tree Plus) and the Dell switch was configured for the open standard 802.1w RSTP, which was incorporated into the 2004 version of the 802.1D standard.

The objective of this testing was not to recreate every possible permutation of STP in a switched environment, but instead to understand the basic underlying semantics of communication between Cisco's proprietary version of STP and the open standard. As you will see, that communication has gaps that lead to limitations in functionality. It is imperative that network designers understand these limitations and make qualitative decisions as to whether the requirements of their network will be met notwithstanding – or whether an alternate design must be deployed. In short, the results of this test show that the Dell Power Connect 6200 series switch behaved as any IEEE 802.1D-compliant bridge would when interacting with Cisco's proprietary PVST+ protocol.

Part II of this document covers my **personal** recommendations to overcome the sub-optimal results of deploying a mixed STP environment. One possible alternative is to deploy Multiple Spanning-Tree (MST) or, when possible, STP can be altogether avoided by leveraging Dell's innovative Simple Switch Mode of operation.

TEST SCENARIO:

Hardware:

- 2 Cisco 3550 switches at the distribution layer (D1 and D2)
- 1 Dell Power Connect 6224 switch as the access switch (A1)
- 3 Category 6 UTP Cables

Software

- Cisco 3550 switches running IOS Version 12.2(35) SE5
- Dell 6224 Switch running OS Version 2.2.0.3




Figure 1 - Lab Design. Looped Topology

The next few examples show the detailed configurations of the Cisco D1 and D2 switches, as well as the Dell 6224 A1 access switch. The configurations are followed by a series of 'show' commands to illustrate the state of Spanning-Tree convergence prior to testing. Pertinent notes are in *blue italics*.

Example 1 - Applicable configuration of Cisco D1 switch

```
switch-D1#show running-config
Building configuration...
Current configuration : 2599 bytes
!
version 12.2
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname switch-D1
!
ip subnet-zero
ip routing
!
(Spanning Tree mode set to rapid-PVST+)
spanning-tree mode rapid-pvst
spanning-tree extend system-id
(Bridge set as Primary Root for VLAN 1 and 10)
spanning-tree vlan 1,10 priority 4096
!
vlan internal allocation policy ascending
!
interface FastEthernet0/1
description dot1q_trunk_to_Dell_6224_1/g1
switchport trunk encapsulation dot1q
switchport mode trunk
!
interface FastEthernet0/24
description dot1q_crosslink_to_Cisco_D2_Fa0/24
switchport trunk encapsulation dot1q
switchport mode trunk
!
interface Vlan1
no ip address
!
interface Vlan10
description Management_VLAN
ip address 10.10.10.2 255.255.255.0
!
ip classless
!
end
!
switch-D1#
```

Example 2 - Applicable configuration of Cisco D2 switch

```
switch-D2#show running-config
Building configuration...
Current configuration : 2330 bytes
!
version 12.2
service timestamps debug uptime
service timestamps log uptime
!
hostname switch-D2
!
ip subnet-zero
!
(Spanning-Tree mode set to rapid-PVST+)
spanning-tree mode rapid-pvst
spanning-tree extend system-id
!
vlan internal allocation policy ascending
!
interface FastEthernet0/1
description dot1q_trunk_to_Dell_6224_1/g2
switchport trunk encapsulation dot1q
switchport mode trunk
!
interface FastEthernet0/24
description dot1q_crosslink_to_Cisco_D1_Fa0/24
switchport trunk encapsulation dot1q
switchport mode trunk
!
interface Vlan1
no ip address
!
interface Vlan10
description Management_VLAN
ip address 10.10.10.3 255.255.255.0
!
ip classless
!
end
!
switch-D2#
```

Example 3 - Applicable Dell 6224 Switch Configuration

```
Dell_6224#show running-config
!Current Configuration:
!System Description "PowerConnect 6224P, 2.2.0.3, VxWorks5.5.1"
!System Software Version 2.2.0.3
!
configure
vlan database
vlan 10
exit
hostname "6224#1"
exit
ip routing
interface vlan 10
routing
ip address 10.10.1.1 255.255.255.0
exit
exit
!
interface ethernet 1/g1
description dot1q_trunk_to_Cisco_D1-Fa0/1
switchport mode general
no switchport general acceptable-frame-type tagged-only
switchport general allowed vlan add 10 tagged
exit
!
interface ethernet 1/g2
description dot1q_trunk_to_Cisco_D2-Fa0/1
switchport mode general
no switchport general acceptable-frame-type tagged-only
switchport general allowed vlan add 10 tagged
exit
exit

Dell_6224#
```

The next few examples display the output of the 'show spanning-tree' command for all three switches. The output shows the converged status of the Spanning-Tree Protocol on each switch.

Example 4 – ‘Show Spanning-Tree’ Output for Cisco Switch D1

switch-D1#sh spanning-tree

(Note that the Cisco switches show Spanning-Tree statistics for each VLAN configured – per-VLAN Spanning-Tree +)

VLAN0001

Spanning tree enabled protocol rstp

Root ID Priority 4097

Address 000e.8364.6d80

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 000e.8364.6d80

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

Aging Time 300

Interface	Role	Sts	Cost	Prio.	Nbr	Type
-----------	------	-----	------	-------	-----	------

(Both trunk ports are Designated and Forwarding)

Fa0/1	Desg	FWD	19	128.1		P2p
-------	------	-----	----	-------	--	-----

Fa0/24	Desg	FWD	19	128.24		P2p
--------	------	-----	----	--------	--	-----

VLAN0010

Spanning tree enabled protocol rstp

Root ID Priority 4106

Address 000e.8364.6d80

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 000e.8364.6d80

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

Aging Time 300

Interface	Role	Sts	Cost	Prio.	Nbr	Type
-----------	------	-----	------	-------	-----	------

(Both trunk ports are Designated and Forwarding)

Fa0/1	Desg	FWD	19	128.1		P2p
-------	------	-----	----	-------	--	-----

Fa0/24	Desg	FWD	19	128.24		P2p
--------	------	-----	----	--------	--	-----

switch-D1#

Example 5 – ‘Show Spanning-Tree’ Output for Cisco Switch D2

switch-D2#sh spanning-tree

(Note that the Cisco switches show Spanning-Tree statistics for each VLAN configured – per-VLAN Spanning-Tree +)

VLAN0001

Spanning tree enabled protocol rstp

(Recognizes Cisco Switch D1 as root for both VLANs)

Root ID Priority 4097
Address 000e.8364.6d80
Cost 19
Port 24 (FastEthernet0/24)
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority 32769 (priority 32768 sys-id-ext 1)
Address 000d.bc6e.9300
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
Aging Time 300

Interface	Role	Sts	Cost	Prio.	Nbr	Type
-----------	------	-----	------	-------	-----	------

(Fa0/24 is the root port and Fa0/1 is Designated. Both Forwarding)

Fa0/1	Desg	FWD	19	128.1	P2p	
Fa0/24	Root	FWD	19	128.24	P2p	

VLAN0010

Spanning tree enabled protocol rstp

Root ID Priority 4106
Address 000e.8364.6d80
Cost 19
Port 24 (FastEthernet0/24)
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority 32778 (priority 32768 sys-id-ext 10)
Address 000d.bc6e.9300
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
Aging Time 300

Interface	Role	Sts	Cost	Prio.	Nbr	Type
-----------	------	-----	------	-------	-----	------

(Fa0/24 is the root port and Fa0/1 is Designated. Both Forwarding)

Fa0/1	Desg	FWD	19	128.1	P2p	
Fa0/24	Root	FWD	19	128.24	P2p	

switch-D2#

Example 6 – ‘Show Spanning-Tree’ Output for Dell Switch A1

Dell_6224#show spanning-tree

(Note the singular instance of RSTP (the CST) displayed by the Dell 6224 802.1w-compliant bridge)

Spanning tree Enabled mode rstp

CST Regional Root: 80:00:5C:26:0A:91:1F:C2

Regional Root Path Cost: 0

ROOT ID *(Recognizes Cisco D1 as Root Bridge)*

Address 10:01:00:0E:83:64:6D:80

Path Cost 200000

Root Port 1/g1

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

Bridge ID

Priority 32768

Address 80:00:5C:26:0A:91:1F:C2

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

Interfaces

Name	State	Prio.Nbr	Cost	Sts	Role	PortFast	RootPort

1/g1	Enabled	128.1	200000	FWD	Root	No	No
1/g2	Enabled	128.2	200000	DSC	Altn	No	No

(Port 1/g1 is the root port and it's in Forwarding State; port 1/xg2 is Alternate in Discarding state)

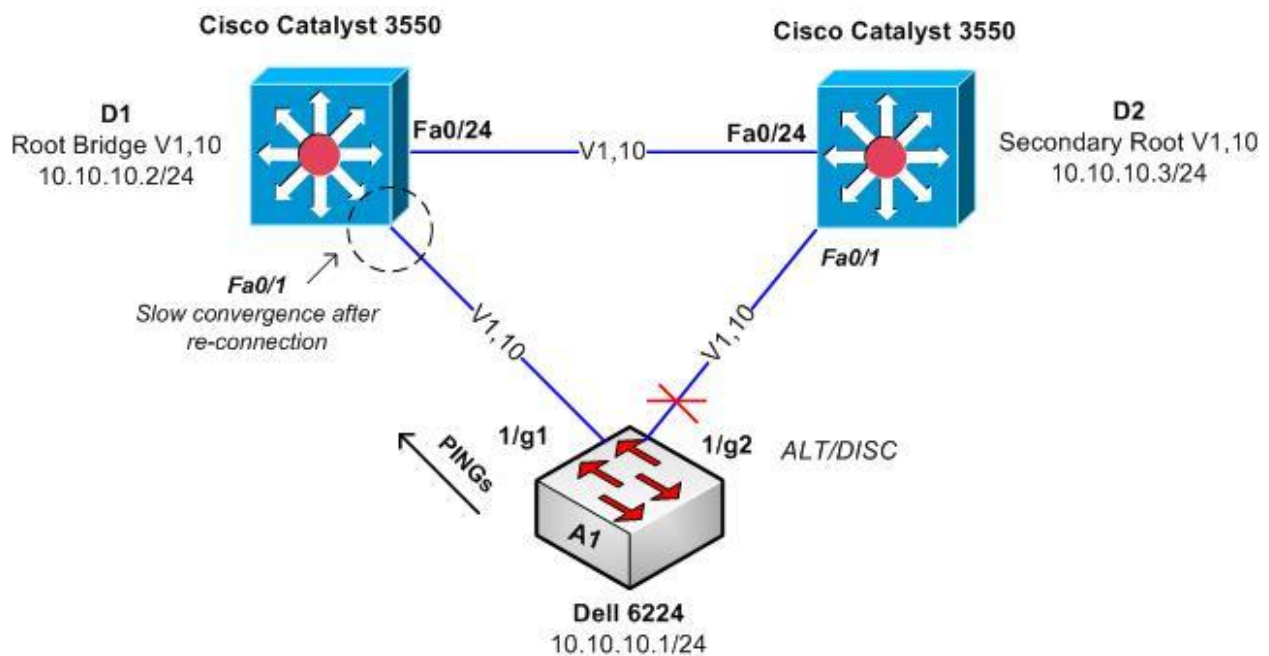


Figure 2 – Lab Test Details.

Test Method:

1. Initiated continuous PING from A1 to D1's SVI for VLAN 10
2. Disconnected A1 to D1 link; failover to D2 link fast (dropped one PING)
3. Reconnected A1 to D1 link; network took over 30 seconds to reconverge (30 PINGs dropped)

Example 7 – Results of PING Test

(Step 1 – Initiated Continuous PING of D1's VLAN 10 interface; timeout of 1000 ms (1s) for each PING)

```
Dell_6224#ping 10.10.10.2 count 60 timeout 1000
```

```
Pinging 10.10.10.2 with 64 bytes of data:
```

```
64 bytes from 10.10.10.2: icmp_seq = 0. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 1. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 2. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 3. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 4. time < 10ms
```

(Step 2 – Disconnect A1-D1 link at D1. PING seq=5 dropped as A1's ALT port transitions to Forwarding state)

```
64 bytes from 10.10.10.2: icmp_seq = 6. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 7. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 8. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 9. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 10. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 11. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 12. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 13. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 14. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 15. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 16. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 17. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 18. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 19. time < 10ms
```

(Step 3 – A1-D1 link restored. Port Fa0/1 on D1 goes through STP Listening and Learning states – 30 seconds – exactly 30 PINGs dropped)

```
64 bytes from 10.10.10.2: icmp_seq = 49. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 50. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 51. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 52. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 53. time < 10ms
```

```
64 bytes from 10.10.10.2: icmp_seq = 54. time < 10ms
```

```
----10.10.10.2 PING Statistics----
```

```
55 packets transmitted, 29 packets received, 53% packet loss
```

```
round-trip (ms) min/avg/max = 0/0/0
```

```
Dell_6224#
```

Test Results Analysis:

The results obtained during failover testing were consistent with the expected results when a third party switch, running 802.1w RSTP, interacts with a Cisco switch running rapid-PVST+.

What exactly happened and why was reconvergence so slow?

In order to interoperate with a standards-compliant bridge, a Cisco switch running its proprietary rapid-PVST+ protocol must interact with the mono Spanning Tree instance, otherwise known as the Common Spanning Tree (CST), of the 802.1w-compliant bridge. Whereas PVST+ executes an STP instance for each VLAN, including VLAN 1, CST assumes only one spanning-tree instance for the entire bridged network, regardless of the number of VLANs configured. The CST instance is represented by the logical topology that results from STP convergence of VLAN 1.

When connected to a third party switch, the Cisco switch will exchange STP information with the CST by sending its Bridge Protocol Data Units (BPDUs) for VLAN 1 to the Spanning-Tree Protocol's **well-known** multicast address of **0180.C200.0000** (SAP encapsulation - DSAP 42, SSAP 42) on an 802.1Q trunk. Note that the Cisco switch will also send BPDUs for VLAN 1 on its reserved multicast address of **0100.0cCC.CCCD** (SNAP encapsulation).

The **0180.C200.0000** address is the multicast address that 802.1w/802.1D-2004-compliant bridges listen to for STP configuration and topology change notifications, as well as for rapid-STP proposal and agreement BPDUs. **Proposal and agreement BPDUs are used to negotiate fast-convergence parameters in the event of a failure.** Therefore, for a Cisco bridge running rapid-PVST+ to be able to converge with a third party bridge running 802.1w RSTP, VLAN 1 must be allowed on all 802.1Q trunks that interconnect them.

For all 802.1Q tagged VLANs (more specifically, all VLANs besides VLAN 1), Cisco switches send their BPDUs **only** to the reserved Cisco multicast address of **0100.0cCC.CCCD**. Therefore, unless the third party switch is also listening to this multicast address, its STP process will only have visibility to, and an understanding of, the logical topology of the CST. That means the rapid failover and convergence offered by RSTP will only be evident on VLAN 1. Regardless of the existence of other VLANs, from an STP perspective, the standards-compliant bridge “sees” only VLAN 1, and therefore only one logical STP topology exists.

Example 8 - Output from Cisco D1 rapid-PVST+ Events Debug

(Notice that BPDUs for VLAN 1 are being sent to both the Cisco multicast address as well as the standard's multicast address. Note also that BPDUs will be sent out both active trunk ports.)

01:02:59: RSTP(1): sending BPDU out Fa0/1

01:02:59: STP SW: TX: sstp copy: **0100.0ccc.cccd**<-000e.8364.6d81 type/len 0032

01:02:59: **encap SNAP** linktype sstp vlan 1 len 64 on v1 Fa0/1

01:02:59: AA AA 03 00000C 010B SSTP

01:02:59: CFG P:0000 V:02 T:02 F:3C R:1001 000e.8364.6d80 00000000

01:02:59: B:1001 000e.8364.6d80 80.01 A:0000 M:1400 H:0200 F:0F00

(Con't)

```
01:02:59: T:0000 L:0002 D:0001
01:02:59: STP SW: TX: 0180.c200.0000<-000e.8364.6d81 type/len 0027
01:02:59: encap SAP linktype ieee-st vlan 10 len 60 on v10 Fa0/1 01:02:59: 42 42 03 SPAN
01:02:59: CFG P:0000 V:02 T:02 F:3C R:1001 000e.8364.6d80 00000000
01:02:59: B:1001 000e.8364.6d80 80.01 A:0000 M:1400 H:0200 F:0F00
01:02:59: RSTP(1): sending BPDU out Fa0/24
01:02:59: STP SW: TX: sstp copy: 0100.0ccc.cccd<-000e.8364.6d98 type/len 0032
01:02:59: encap SNAP linktype sstp vlan 1 len 64 on v1 Fa0/24
01:02:59: AA AA 03 00000C 010B SSTP
01:02:59: CFG P:0000 V:02 T:02 F:3C R:1001 000e.8364.6d80 00000000
01:02:59: B:1001 000e.8364.6d80 80.18 A:0000 M:1400 H:0200 F:0F00
01:02:59: T:0000 L:0002 D:0001
01:02:59: STP SW: TX: 0180.c200.0000<-000e.8364.6d98 type/len 0027
01:02:59: encap SAP linktype ieee-st vlan 10 len 60 on v10 Fa0/24
01:02:59: 42 42 03 SPAN
01:02:59: CFG P:0000 V:02 T:02 F:3C R:1001 000e.8364.6d80 00000000
01:02:59: B:1001 000e.8364.6d80 80.18 A:0000 M:1400 H:0200 F:0F00
```

*(Notice below that the BPDUs for VLAN 10 are being sent **only** to the Cisco reserved multicast address, which the Dell switch forwards without acting on them. Those BPDUs will also be sent out both trunk ports.)*

```
01:02:59: RSTP(10): sending BPDU out Fa0/1
01:02:59: STP SW: TX: 0100.0ccc.cccd<-000e.8364.6d81 type/len 0032
01:02:59: encap SNAP linktype sstp vlan 10 len 64 on v10 Fa0/1
01:02:59: AA AA 03 00000C 010B SSTP
01:02:59: CFG P:0000 V:02 T:02 F:3C R:100A 000e.8364.6d80 00000000
01:02:59: B:100A 000e.8364.6d80 80.01 A:0000 M:1400 H:0200 F:0F00
01:02:59: T:0000 L:0002 D:000A
01:02:59: RSTP(10): sending BPDU out Fa0/24
01:02:59: STP SW: TX: 0100.0ccc.cccd<-000e.8364.6d98 type/len 0032
01:02:59: encap SNAP linktype sstp vlan 10 len 64 on v10 Fa0/24
01:02:59: AA AA 03 00000C 010B SSTP
01:02:59: CFG P:0000 V:02 T:02 F:3C R:100A 000e.8364.6d80 00000000
01:02:59: B:100A 000e.8364.6d80 80.18 A:0000 M:1400 H:0200 F:0F00
01:02:59: T:0000 L:0002 D:000A
```

Example 9 – Output from Cisco D1 rapid-PVST+ Events Debug During Failover Test

(A1-D1 link failed by disconnecting cable from D1. Not shown here, but Dell 6224 uplink port facing D2 immediately transitions into the Forwarding state.)

```
00:03:51: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to down
00:03:52: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to down
```

(Con't)

*(A1-D1 link reconnected at D1. Take note that the proposal sent by D1 is responded to by A1 only for VLAN 1. The proposal for VLAN 10 goes **unanswered** for a duration exceeded by two times the forward-delay interval – 30 seconds by default. During this time, PING packets are being dropped.)*

```
00:04:00: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to up
00:04:01: RSTP(1): initializing port Fa0/1
00:04:01: RSTP(1): Fa0/1 is now designated
00:04:01: RSTP(10): initializing port Fa0/1
00:04:01: RSTP(10): Fa0/1 is now designated
00:04:01: RSTP(1): transmitting a proposal on Fa0/1
00:04:01: RSTP(10): transmitting a proposal on Fa0/1
00:04:01: RSTP(1): received an agreement on Fa0/1
00:04:01: RSTP(10): transmitting a proposal on Fa0/1
00:04:02: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed state to up
00:04:03: RSTP(10): transmitting a proposal on Fa0/1
00:04:05: RSTP(10): transmitting a proposal on Fa0/1
00:04:07: RSTP(10): transmitting a proposal on Fa0/1
00:04:09: RSTP(10): transmitting a proposal on Fa0/1
00:04:11: RSTP(10): transmitting a proposal on Fa0/1
00:04:13: RSTP(10): transmitting a proposal on Fa0/1
00:04:15: RSTP(10): transmitting a proposal on Fa0/1
00:04:16: RSTP(10): Fa0/1 fdwhile Expired
00:04:17: RSTP(10): transmitting a proposal on Fa0/1
00:04:19: RSTP(10): transmitting a proposal on Fa0/1
00:04:21: RSTP(10): transmitting a proposal on Fa0/1
00:04:23: RSTP(10): transmitting a proposal on Fa0/1
00:04:25: RSTP(10): transmitting a proposal on Fa0/1
00:04:27: RSTP(10): transmitting a proposal on Fa0/1
00:04:29: RSTP(10): transmitting a proposal on Fa0/1
00:04:31: RSTP(10): Fa0/1 fdwhile Expired
switch-D1#
```

The output above clearly explains why, upon restoring the downed link between A1 and D1, the D1 switch port on the A1-D1 link (designated port) had to go through all the STP port states before shifting to the forwarding state, which resulted in the loss of many PINGs. The RSTP proposal and agreement BPDUs for VLAN 10 that would normally have been exchanged by the access and distribution layer switches operating PVST+ – or by non-Cisco switches on VLAN 1 – were ignored by the Dell switches for all but VLAN 1. With no agreement in place for VLAN 10, the Cisco switch port had no choice but to go through all the STP port states, which took anywhere between 30 and 50 seconds (two times the Forward Delay timer plus any time remaining on the Max-age timer).

The initial failover to the A1-D2 link is a different matter, though. This occurred rapidly, as expected. The reason is that the speed with which switch A1's port state machine moved into the Forwarding state from the ALT/DISC state did not depend on any interaction with an external switch during the failure

itself – no BPDU exchanges were needed. When A1 first converged, it recognized its uplink to D2 as an alternate path to the root bridge, so it rightly placed its local port in the ALT/DISC state – at the ready in the event that the root port fails.

For all VLANs, except VLAN 1, the Dell 6224 switch simply received the multicast packets destined for the reserved Cisco multicast address and flooded them out all corresponding VLAN ports, much the same way broadcasts are handled. In other words, the BPDUs for the tagged 802.1Q VLANs were simply tunneled through the non-Cisco region.

There are non-Cisco switches that do listen to the Cisco multicast address of 0100.-0cCC.CCCD for the sake of interoperability, such as Extreme Networks, Force 10, some Foundry (now Brocade) switches and Juniper's VSTP.

PART II

OVERCOMING THE LIMITATIONS OF RSTP/rapid-PVST+ INTEROPERABILITY

Stand Alone Switch Recommendation

Because of the inherent limitation to fast convergence when interacting with PVST+, it is recommended that the Dell and Cisco switches become part of a Multiple Spanning Tree Protocol region (MSTP). MSTP, which is codified in IEEE 802.1s, uses many of the same rapid convergence primitives defined in RSTP. The major difference is the fact that, with MSTP, multiple VLANs are mapped to a particular MSTP instance, thereby dramatically reducing the number of spanning tree instances that would normally be run by Cisco's PVST+ protocol. Both the Cisco and Dell switches must be configured with the same VLAN-to-MST instance mapping. More pertinent to rapid convergence is the fact that both switches will be fully communicating using the well-known multicast address of **0180.C200.0000**. This allows for agreement on the characteristics of logical topologies and expedited convergence parameters.

The same test was conducted using MSTP. Reconvergence was nearly sub-second upon failover and reconnection of the downed link.

Configuring MST is out of the scope of this document.

Chassis-based Blade Switches

The Dell Power Connect M6220, M6348 and M8024 are fully-featured Ethernet blade switches for the Dell m1000e blade enclosure. By deploying them in Simple Switch Mode (SSM), they offer the capability to be easily deployed as port aggregators for streamlined interoperability with Top-of-Rack (ToR) platforms from multiple vendors, including Cisco. With SSM, STP does not have to be configured at all.

Port aggregation is enabled when the blade switch is placed in Simple Switch Mode (SSM). SSM allows a simplified way of deploying a loop-free layer-2 switched environment without having to contend with

the challenges of integrating STP with third party switches or having to settle for suboptimal reconvergence times.

SSM works by “pinning” a blade switch’s downlink ports to one or more uplink ports to form an Aggregator Group (AG). A downlink port is one that is directly connected to the server blade’s LOMs and mezzanine cards through the traces on the midplane of the blade chassis. The uplink ports are external and connected to the ToR switch via copper or fiber optic cables, depending on the requirements of the deployment scenario. Multiple AGs can be configured within a stand-alone switch or across stacked switches for added flexibility and redundancy.

To keep traffic from different AGs separated from each other, and thus prevent bridging loops, a switch operating in SSM mode will automatically assign a VLAN identifier to each AG. This precludes the need to configure STP. And for more optimal oversubscription ratios, multiple uplink ports can be associated with the same AG. As ports are added, an SSM-enabled switch will dynamically configure a LACP bundle without any administrator intervention. Traffic will then be distributed across the LAG using a configurable hashing algorithm. The result is a simplified, loop-free switched environment with predictable traffic patterns and deterministic and immediate failover.

To summarize, the advantages of deploying SSM are:

- Port Aggregation is easy to configure. Simply map internal ports to external ports, assign a VLAN to the group (if required), and it’s ready to go.
- Automatically configures multiple external ports into an LACP trunk group.
- By using Aggregator Groups, the feature provides loop-free operation without using STP.
- Works across a stack of switches (M6220 and M6348) so that you can now manage switches as one via the easy-to-use interface.
- Seamless interoperability - Uplink looks like NIC ports to the network.
- Port Aggregation is completely interoperable. Dynamic (via LACP) and static link aggregation is supported on the external ports.

Note, that a simplified switching deployment using SSM means that certain functionality is necessarily removed:

- No SNMPv3
- Restricted management interface
- Limited configuration
- Certain commands, such as routing-related features and QoS, are not supported.

Configuring SSM is out of the scope of this document. Please refer to the Dell White Paper entitled, “Simple Switch Mode: Port Aggregator Feature White Paper” published in February of 2010 for details and configuration guidelines.