SPANNING TREE PROTOCOL IN LAYER2/LAYER 3 ENVIRONMENTS

Configuring STP with Juniper Networks EX Series Ethernet Switches

Table of Contents

Table of Figures

Introduction

This document is targeted at SEs and other technical audiences as an implementation guide for configuring the Spanning Tree Protocol (STP) with Juniper Networks® EX Series Ethernet Switches.

Introduction

This implementation guide will provide a basic understanding of all three versions of the Spanning Tree Protocol, namely 802.1d-1998, RSTP (802.1w) and MSTP (802.1s), as well as background on common enterprise networks and how to deploy EX Series switches in a mixed environment with Juniper Networks MX Series 3D Universal Edge Routers and Cisco switches.

Spanning Tree Protocol Overview

The plug-and-play Spanning Tree Protocol (STP), a Layer 2 protocol that ensures a loop-free topology in bridged LANs, has been around for decades. Spanning Tree-enabled switches go through a root election process based on Bridge Protocol Data Unit (BPDU). All other switches then build the shortest, lowest-cost path to the root switch and block ports that are not along these shortest paths, resulting in a loop-free, tree-like topology. Without a protocol such as Spanning Tree, Layer 2 bridged networks are susceptible to broadcast/multicast and/or unknown unicast storms.

In the next few bullets, we will discuss three different versions of standards-based STP, and then talk briefly about Cisco's Spanning Tree implementation—Per-VLAN Spanning Tree Plus (PVST+):

- **Spanning Tree Protocol (STP–802.1d 1998)**: The original version of STP (legacy STP) only supports a single instance of Spanning Tree in a bridged network, typically referred to as Common Spanning Tree (CST). On 802.1Q trunks that carry multiple VLANs, one VLAN—usually the default or VLAN 1—will dictate the forwarding topology for all other VLANs.
  In STP, when a port is enabled or there is any change in the STP topology, it can take up to 50 seconds (MAX_Age + 2 x FWD_Delay, with default timers) for the bridged network to reconverge.

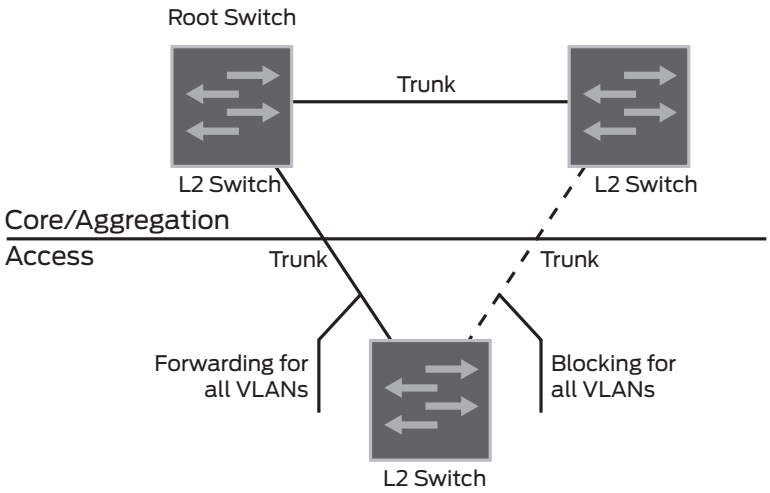

Figure 1: Common Spanning-Tree instance

- **Rapid Spanning Tree Protocol (RSTP–802.1w/802.1d 2004):** Similar to legacy STP, RSTP implements a single instance of Spanning Tree but dramatically improves the convergence time. For example, in legacy STP, a host port coming online can take up to 30 seconds (2 x FWD_Delay) before it begins forwarding. Direct or indirect node/link failure convergence can even take up to 50 seconds (MAX_Age + 2 x FWD_Delay) with legacy STP. Such convergence times are not acceptable in today's high-performance enterprise networks.
  RSTP overcomes slow convergence times by introducing new port roles (edge port, alternate port and backup port) and new convergence behavior under specific connectivity conditions (for example, point-to-point or shared medium). These enhancements enable RSTP to achieve sub-second forwarding and network convergence. RSTP is backward-compatible with legacy STP; if legacy STP BPDU is detected on a link, then an RSTP-capable bridge will revert to legacy STP on that given port.

- **Multiple Spanning Tree Protocol (MSTP–802.1s/802.1Q 2003):** MSTP allows multiple instances of Spanning Tree to run concurrently over a bridged network. A Multiple Spanning Tree Instance (MSTI) is associated with a group of VLANs. There can be up to 64 MSTIs, and each instance can have its own independent forwarding topology by configuring different MSTI root switches. This will allow all links on all switches to be utilized. For example in Figure 2, both uplinks of the access layer switch will be forwarding traffic instead of just one as seen in legacy STP or RSTP. MSTP is an extension of RSTP features (edge ports, alternate ports and so on) and can interoperate with STP/RSTP switches.
- **Per-VLAN Spanning Tree Plus/Rapid-PVST+ (PVST+/RPVST+):** PVST+ is Cisco's version of Spanning Tree. In PVST+, each VLAN will have its own instance of Spanning Tree. On an 802.1Q trunk, VLAN 1 follows the legacy STP (802.1d) standard and advertises on the IEEE STP multicast, whereas the other VLANs on a trunk link advertise on Cisco's reserved multicast address. PVST+ has the same convergence time as legacy STP.
  Since PVST+ convergence time is sub-optimal, Cisco incorporated RSTP features into PVST+, which they call RPVST. However, RPVST+ still advertises on the reserved multicast address for all VLANs except VLAN 1.
- **VLAN Spanning Tree Protocol (VSTP):** VSTP is compatible with PVST+/RPVST+. Similar to PVST+/RPVST+, VSTP creates a per-VLAN spanning tree instance for every VLAN configured for VSTP. VSTP BPDUs are tagged with the corresponding VLAN identifier and are sent to multicast address 01-00-0C-CC-CC-CD to interact with Cisco PVST+/RPVST+ bridges. EX Series switches can support up to 128 VSTP instances as of Junos OS release 9.4R1.

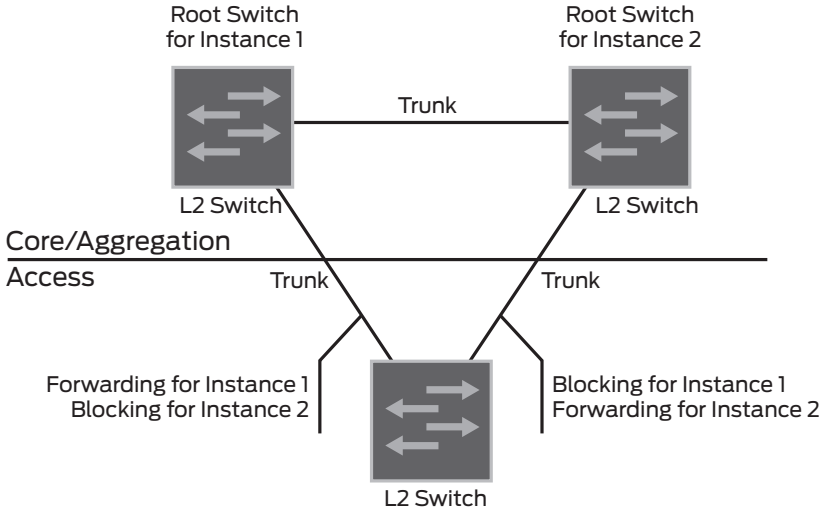

Figure 2: MSTP allows both links of the access layer switch to be used at the same time while maintaining a loop-free network

The Juniper Networks Position on Layer 2 versus Layer 3

The EX Series can be deployed in either Layer 2 or Layer 3 environments, or a combination of both. This section describes the advantages and disadvantages of deploying Layer 3 versus Layer 2 solutions.

STP adds another layer of complexity to network planning and troubleshooting, forcing network administrators to carefully architect their networks. When administrators are tasked with planning for a Layer 2 network, they know STP is necessary to prevent Layer 2 loops. However, they also face the challenge of determining which Spanning Tree Protocol—STP, RSTP or MSTP—to implement, realizing such decisions require tradeoffs between efficiency/functionality and simplicity/complexity. For example, consider the decision of how to handle traffic forwarding over redundant links: should administrators implement MSTP, which is highly efficient but complex, or should they use STP/RSTP, which is less efficient but easier to configure?

When STP does misbehave due to user configuration error or equipment/software failure, the consequences can be devastating. Bridge loops might form within the bridged network, resulting in broadcast/multicast/unknown-unicast storms that can take down the entire network. Troubleshooting such STP-related issues is no easy task and requires STP experts to spend a great deal of time narrowing down and finding the root cause of the failure.

Administrators also need to plan for additional features such as BPDU Guard, Root Guard or Loop Guard to prevent inadvertent network changes by rogue switches. In addition, to achieve comprehensive integrated high availability (HA) in both Layer 2 and Layer 3, administrators must be cognizant of the relationship between routing protocols, Virtual Router Redundancy Protocol (VRRP) and STP.

The EX Series includes full OSPF and IP multicast as part of the base software license, making Layer 3 to the access layer a viable solution and eliminating STP from the aggregation and core portions of the network. Network administrators must carefully plan for IP allocation in their network; in a routed network, administrators shouldn't worry about Layer 2 loops because all uplinks should be IP segmented. Nor should administrators be concerned about unutilized links because OSPF implements equal-cost multipath routing. Network administrators can expect sub-second convergence with features such as Bi-directional Forwarding Detection (BFD) routing protocol.

Eliminating STP from a majority of the network means that there is one less protocol to plan, test, implement, troubleshoot and manage. Network administrators can expect lower capital and operational expenses when deploying a Layer 3 solution with the EX Series.

Configuration Guide

This section covers the three-tiered core-, aggregation- and access-layer topologies commonly used in today's enterprise network (see Figure 3). The Juniper Networks product portfolio can be deployed in a three-tiered network. At the core layer, for instance, network architects desire a low-latency switch with non-blocking fabrics and wire rate linecards—solutions such as Juniper Networks MX Series 3D Universal Edge Routers. At the aggregation layer, depending on the connectivity choice (10 GbE or GbE), either Juniper Networks MX480 Ethernet Services Router or EX4200-24F (24-port 100/1000BASE-X fiber) switches with Virtual Chassis technology can be used. At the access layer, Juniper Networks EX4200 line of Ethernet switches are a perfect fit.

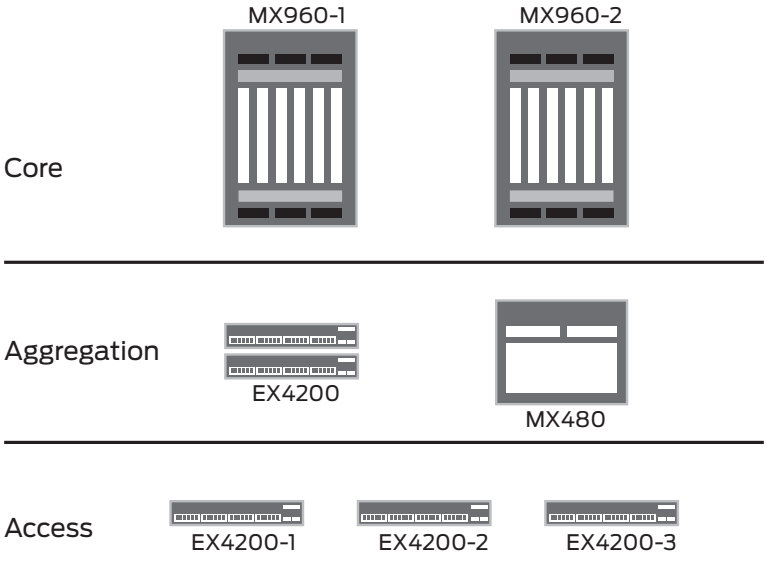

Figure 3: Typical three-tier enterprise switched network

Configuration Options

This section covers multiple configuration options for the topology shown in Figure 3.

Option 1: Layer 2 Everywhere (Not Recommended)

As described in the previous section, deploying Spanning Tree on all three layers is not recommended due to the added complexity. However, we will still cover the configuration as it is deployed in some customer networks.

This section will show how to configure STP/RSTP and MSTP on each of the switches. First, the commands for configuring STP/RSTP on the switches are covered.

STP/RSTP Configuration

Core Switches:

Since STP and RSTP configurations are similar, this section will focus on RSTP. To configure STP, simply change the keyword from "RSTP" to "STP."

In Spanning Tree, all enabled Spanning Tree bridges first go through a root election process. The switch with the lowest bridge ID is declared the root. A bridge ID consists of a bridge priority and MAC address. The root switch is the base for Spanning Tree where all of the other non-root switches build their shortest path to the root switch and block any redundant links. The placement of the root switch is important since the majority of traffic will pass through the root switch.

Rather than rely on chance, it is important to configure the bridge priority on the switch to make sure that it becomes the root switch. The bridge priority defaults to 32768 but is configurable from 0 to 65536. If switches have the same bridge priority, then the bridge with the lowest MAC address wins. In Figure 3, the MX960-1 should be the root switch, since the majority of client traffic will be to and from the local and/or remote servers (typically server switches and/or WAN routers, not shown in the diagram, are connected to the core switches). Therefore, MX960-1 is configured with a bridge-priority of 4K (4096), which makes it the lowest bridge ID and, consequently, the root switch.

It is equally important to configure a backup root switch. If a root switch fails, all remaining switches repeat the election process to select a new root switch; configuring a backup allows you to preselect which device will take over in the case of a primary failure—in this case, the core MX960-2 switch. Therefore, MX960-2 is configured with a bridge-priority of 8K (8192), making its bridge ID higher than MX960-1 but lower than the other switches. This ensures that it becomes the new root switch if the primary root switch fails.

```
juniper@MX960-1# set protocols rstp bridge-priority 4k
juniper@MX960-2# set protocols rstp bridge-priority 8k
```

Aggregation Switches:

Because RSTP root is already designated on the core switch, no extra RSTP configuration is required on the switches in the aggregation layer; RSTP is enabled by default on the EX Series.

On the MX Series router, the same command as above (core layer) is used to enable RSTP, but without the bridge-priority command.

Access Switches:

An Edge Port feature allows a port to transition to a forwarding state without a 30-second delay (2 x FWD_Delay). This feature is ideal for ports that are connected to PCs/servers, routers, firewalls and so on.

RSTP automatically incorporates Edge Port based on certain characteristics of a port such as full-duplex mode and the fact that no BPDU has been detected, but this feature has the option to be manually configured for an interface as well. In the case of STP, the Edge Port feature must be manually configured for specific interfaces using the following command:

```
juniper@EX4200# set protocols stp interface ge-0/0/0.0 edge
```

If a BPDU is detected on the Edge Port, the protocol is intelligent enough to reclassify the port as a non-edge port. Thus if a loop is detected, the port can transition to a blocking state.

MSTP Configuration

MSTP is the desired Spanning Tree Protocol. MSTP allows network architects to load-balance VLANs, as well as to pass traffic through those VLANs over multiple uplinks and still maintain a loop-free network. When configuring MSTP, the region name, revision mapping and VLAN mappings must all be the same on all of the switches; otherwise these switches will not be part of the same region. In the example shown in Figure 3, all switches in the enterprise network are in the same region.

The majority of enterprise networks will have no more than two Spanning Tree instances; each MSTI will be mapped to a group of VLANs – for example, odd VLANs for Instance 1 and even VLANs for Instance 2. To achieve MSTI load balancing, one of the core switches will be the MSTI regional root for Instance 1 and the backup MSTI regional root for Instance 2, and vice versa for the other core switch.

Core Switches:

```
juniper@MX960-1# set protocols mstp msti 1 bridge-priority 4k vlan [v1 v3]
juniper@MX960-1# set protocols mstp msti 2 bridge-priority 8k vlan [v2 v4]
juniper@MX960-2# set protocols mstp msti 1 bridge-priority 8k vlan [v1 v3]
juniper@MX960-2# set protocols mstp msti 2 bridge-priority 4k vlan [v2 v4]
```

Aggregation and Access Switches:

```
juniper@MX480# set protocols mstp msti 1 vlan [v1 v3]
juniper@MX480# set protocols mstp msti 2 vlan [v2 v4]
juniper@EX4200-1# set protocols mstp msti 1 vlan [v1 v3]
juniper@EX4200-1# set protocols mstp msti 2 vlan [v2 v4]
```

Depending on port connectivity and density requirements (10 GbE or GbE), either the MX480 router or the EX4200 switches can be used at the aggregation layer.

Option 2: Layer 3 at the Core and Between Core and Aggregation Layers; Layer 2 Between Aggregation and Access Layers and at the Access Layer

For this option, OSPF or other routing protocols must be configured at the core and aggregation switches. Since the focus of this document is Spanning Tree, Layer 3 configurations will not be covered.

For Layer 2, MSTP would be ideal between the aggregation and access layers. The root switch should be one of the two switches in the aggregation layer. Since Spanning Tree commands were covered in the previous section, no examples will be provided here.

Option 3: Layer 3 Everywhere

OSPF is part of the base license package for the EX Series; therefore, configuring Layer 3 on all routers/switches in Figure 3 is another option. OSPF configurations will not be covered, since the focus of this document is Spanning Tree.

Even though Layer 3 is implemented to the access layer, it is recommended that Spanning Tree remain enabled on the EX Series to prevent inadvertent loops that can be caused by end users.

Spanning Tree Interoperability with Cisco Switches

This section will cover STP interoperability between the Juniper Networks EX Series switches and Cisco switches, including the interaction between STP/RSTP and PVST+. Alternative solutions to PVST+ by using MSTP (802.1s) will also be discussed.

PVST+

On a given 802.1Q trunk, VLAN 1 BPDUs are always advertised as untagged to the well-known IEEE Spanning Tree multicast address, 01:80:C2:00:00:00. The other VLAN(s) on the trunk advertise tagged PVST+ BPDUs to Cisco's reserved multicast address, 01:00:0C:CC:CC:CD. This allows Cisco to have a per-VLAN instance that can utilize multiple redundant trunk links by load-balancing the VLANs across them. As shown in Figure 4, odd VLANs are forwarded on one of the uplink trunks and blocked on the other uplink trunk, and vice versa for the even VLANs.

One drawback to PVST+ is that it is not embraced by any standard. PVST+ also places a large burden on switch resources such as CPU and memory, as each switch must maintain a separate STP database instance and process BPDUs on a per-port per-VLAN basis. While that might be OK with just a few VLANs, it obviously cannot scale when 4,000 VLANs are supported.

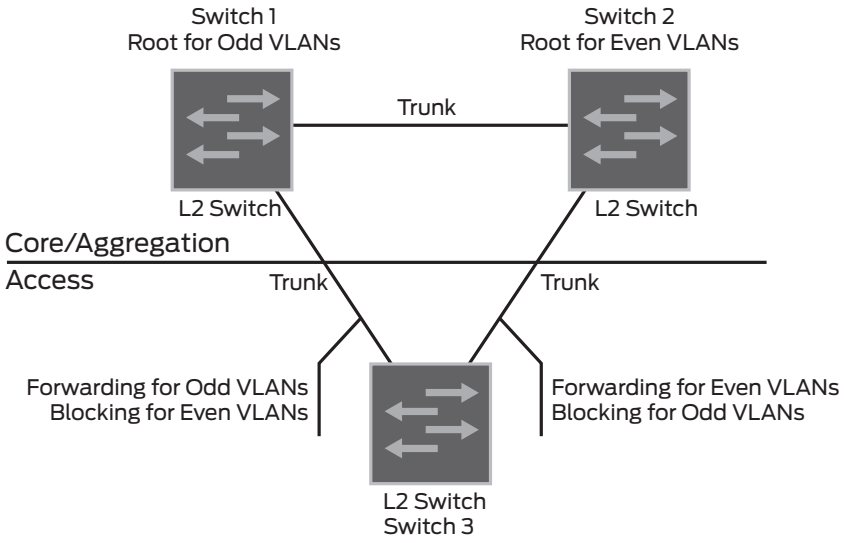

Figure 4: PVST+ implementation

Interoperating STP/RSTP with PVST+

By default, RSTP is enabled on all Ethernet switching-enabled ports on the EX Series. RSTP behaves the same as STP, as both only implement a single Spanning Tree instance for all VLANs over a trunk.

Cisco switches, however, default to Per-VLAN Spanning Tree Plus (PVST+) protocol.

Before continuing, it's important to understand how the EX Series interacts with the various types of BPDUs the Cisco switches send out. Since IEEE STP BPDU is standards-based and advertises to the well-known IEEE reserved multicast address, the EX Series will process the IEEE STP BPDU. However, since PVST+ BPDU is a non standards-based protocol and advertises to Cisco's reserved multicast address, the EX Series will treat it as any other multicast packet and flood it on all forwarding ports within the VLAN in which the respective BPDUs are received (see Figure 5).

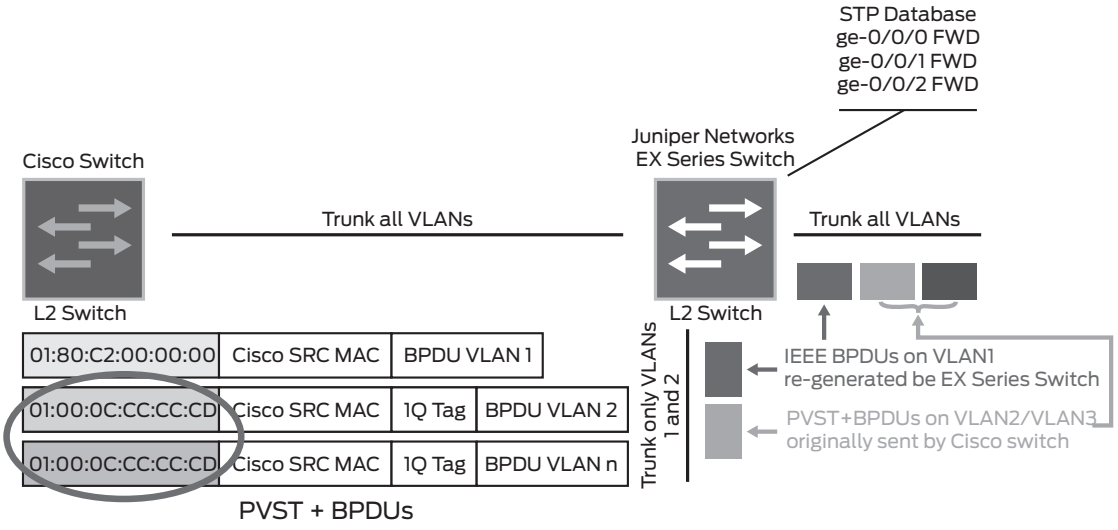

Figure 5: When EX Series switches receive BPDUs from a Cisco switch, only the IEEE STP BPDU is processed; the PVST+ BPDUs are flooded as regular unknown L2 multicast packets

Now that it's understood how EX Series switches handle the various types of BPDUs that Cisco switches send out, the next step will be to see how a typical Layer 2 topology forwarding path will behave with a mixture of Juniper Networks switches running STP/RSTP and Cisco switches running PVST+.

In Figure 6, VLANs 1 through 4 are traversing all 802.1Q trunk links. Switches 1 and 2 are Cisco switches running PVST+. Switch 1 is the STP root for VLANs 1 and 3 with a bridge priority of 4096. It is also the STP backup root switch for VLANs 2 and 4 with a bridge priority of 8192. Switch 2 is the STP root switch for VLANs 2 and 4 with a bridge priority 4096 and is the STP backup root switch for VLANs 1 and 3 with a bridge priority of 8192.

There is one additional caveat to PVST+. By default, the path cost assigned by PVST+ is based on Cisco's STP cost convention, not the IEEE STP cost. Therefore, the path cost needs to be changed so that it complies with IEEE STP path cost and all similar links (like FastE, GbE) have the same link-cost value. On the Cisco devices, the command "spanning-tree pathcost method long" must be entered to make it compliant with IEEE STP path cost. This command will be applied to all Cisco switch ports.

Switch 3 in Figure 6 is a Juniper Networks EX Series switch running RSTP. Since RSTP only has a single Spanning Tree instance for all VLANs across a trunk and there is no VLAN awareness, all traffic will be forwarded or blocked based on the VLAN 1 topology. Therefore, the link between Switch 1 and Switch 3 will be forwarding for **all** VLANs (VLAN 1 through 4) and the other link will be blocking for **all** VLANs (VLAN 1 through 4).

The net result is that Switch 3 is blocking one of the uplinks; there is no logical L2 loop in the topology presented in Figure 6, regardless of which Cisco switch is the root for any of the VLANs (except VLAN 1) and regardless of the fact that these Cisco switches are running PVST+.

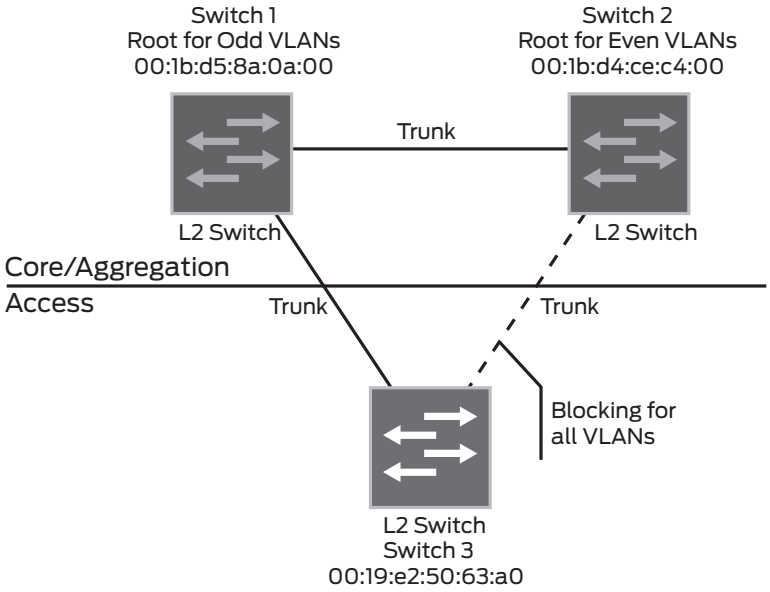

Figure 6: Cisco at the core/aggregation layers running PVST+ and Juniper Networks EX Series Ethernet Switches at the access layer running RSTP; only one link is forwarding traffic

The output below is taken from Switch 3. From the "show spanning-tree bridge" command, the EX Series switch properly acknowledged the root bridge is Switch 1. The "show spanning-tree interface" command shows the port states and, based on the output, only interface ge-0/0/0.0 is forwarding.

```
juniper@EX4200> show spanning-tree bridge
STP bridge parameters
Context ID                          : 0
Enabled protocol                    : RSTP
  Root ID                           : 4097.00:1b:d5:8a:0a:00
  Root cost                         : 20000
  Root port                         : ge-0/0/0.0
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Message age                       : 1
  Number of topology changes        : 6
  Time since last topology change   : 315 seconds
  Local parameters
    Bridge ID                       : 32768.00:19:e2:50:63:a0
    Extended system ID              : 0
    Internal instance ID            : 0

juniper@EX4200> show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/0.0     128:514      128:1     4097.001bd58a0a00     20000  FWD    ROOT
ge-0/0/1.0     128:516      128:55    8192.001bd4cec400     20000  BLK    ALT
```

In Figure 7, the topology has been reversed, with the EX Series at the core/aggregation layers and the Cisco switch at the access layer, supporting all four VLANs.

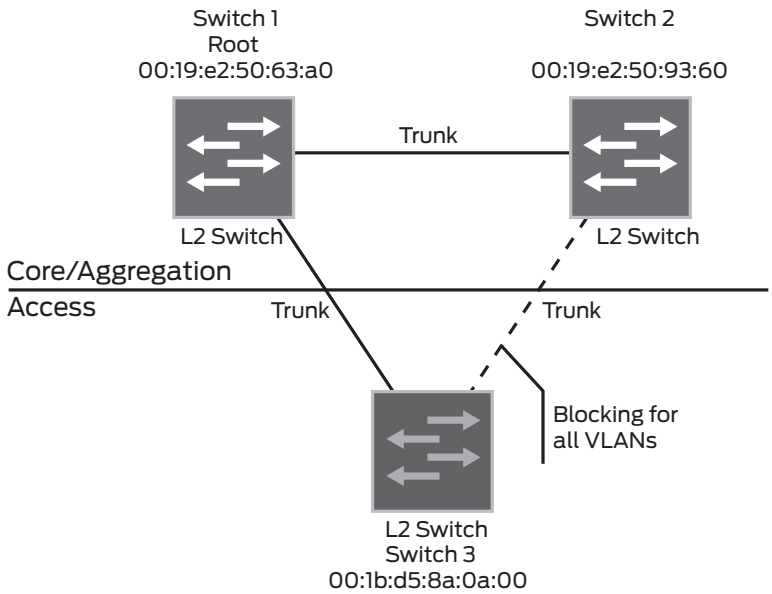

Figure 7: Juniper Networks EX Series Ethernet Switches at the core/aggregation layers running RSTP and Cisco at the access layer running PVST+; only one uplink on the Cisco switch is forwarding traffic

Again, RSTP is enabled on the EX Series (Switches 1 and 2) and PVST+ is enabled on the Cisco switch (Switch 3). Switch 1 is the root of the CST with a priority of 4096; Switch 2 is the backup root switch of the CST with a priority of 8192.

Switch 3 (the Cisco switch running PVST+) recognizes Switch 1 as the STP root for VLAN 1 and recognizes Switch 2 as having a lower bridge-priority than itself on VLAN 1, therefore blocking its uplink trunk port to Switch 2 on VLAN 1.

For VLANs 2-4, Switch 3 sends out Cisco PVST+ BPDUs on the trunks. Because Switch 1 and Switch 2 do not understand PVST+ BDPUs, the switches will treat them as regular L2 multicast packets and flood them. Switch 3 is the PVST+ root for VLANs 2-4 and receives its own PVST+ BPDUs on VLANs 2-4 on both of its uplinks; it ends up blocking one of these two uplink trunk ports.

Therefore, the result is a forwarding topology similar to that in Figure 6, again with only one trunk-port uplink forwarding on all VLANs on Switch 3 and the other uplink trunk port blocking on all VLANs. Following is the output from the Cisco switch (Switch 3) showing the topology from Figure 7. Notice the way that all of the VLANs are blocking on Gi1/0/2. Also, note that the port role for VLAN 1 is different from the other VLANs; on VLAN 1, it is classified as an alternate port while on the other VLANs, it is classified as a backup, indicating that Switch 3 receives its own BPDUs on these VLANs on port Gi1/0/2.

```
3750E#show spanning-tree
VLAN0001
  Spanning tree enabled protocol ieee
  Root ID    Priority    4096
             Address     0019.e250.63a0
             Cost        20000
             Port        14 (GigabitEthernet1/0/14)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32769  (priority 32768 sys-id-ext 1)
             Address     001b.d58a.0a00
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Gi1/0/1             Root FWD 20000     128.13   P2p
Gi1/0/2             Altn BLK 20000     128.14   P2p


VLAN0002
Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Gi1/0/1             Desg FWD 20000     128.13   P2p
Gi1/0/2             Back BLK 20000     128.14   P2p
VLAN0003
Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Gi1/0/1             Desg FWD 20000     128.13   P2p
Gi1/0/2             Back BLK 20000     128.14   P2p

VLAN0004
Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Gi1/0/1             Desg FWD 20000     128.13   P2p
Gi1/0/2             Back BLK 20000     128.14   P2p
```

Also note that on the Cisco switch, when changing the native VLAN ID of a trunk to something other than VLAN 1, the CST topology is still based on VLAN 1; this is because the Cisco device will still send the IEEE STP BPDUs untagged with VLAN 1 parameters over an 802.1Q trunk. And for the VLAN ID that is declared as native-vlan, the Cisco switch will send an untagged PVST+ BPDU to the Cisco reserve multicast MAC address 01:00:0C:CC:CC:CD.

In Figure 8 below, VLAN 2 is the native VLAN between Cisco and Juniper Networks. The IEEE STP BPDU sent by the Cisco switch still contains VLAN 1 parameters and is untagged. Also note that the BPDU for VLAN 2 is still sent on the PVST+ multicast address by the Cisco switch and is untagged.

Please note: It is imperative to match the native VLAN ID on both sides of an 802.1Q trunk.

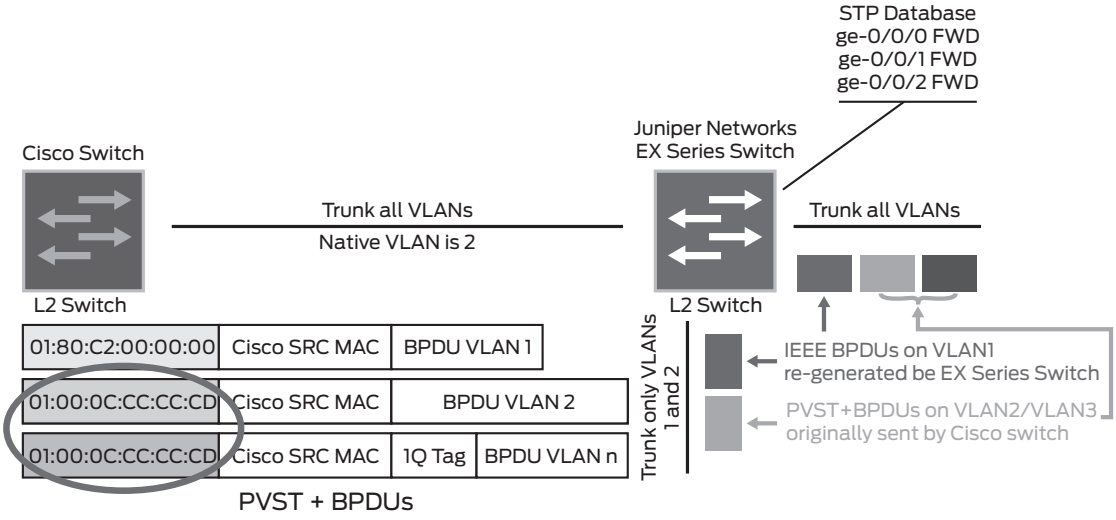

Figure 8: CST is still based on VLAN 1, even though the native VLAN is 2

Multiple Spanning Tree

MSTP helps overcome these STP/RSTP shortcomings, as well as mitigate the demanding resource requirements and lack of standards compliance that are associated with PVST+. The benefit of MSTP is its ability to map a group of VLANs to a Multiple Spanning Tree Instance rather than just one VLAN per instance as seen with PVST+. Different root switches and different STP parameters can be configured for each MSTI, resulting in a different Layer 2 forwarding topology. For example, as shown in Figure 4, though there could be many VLANs configured between the access and core/aggregation layers of an enterprise network, the majority of these networks have no more than two Spanning Tree forwarding topologies. Hence, instead of maintaining X instances of STP for X number of VLANs under PVST+, it only takes two MSTIs to achieve the same result with MSTP—half the VLANs mapped to MSTI 1 and half the VLANs mapped to MSTI 2, as shown in Figure 9. If the network administrator desires additional Spanning Tree topologies, more MSTIs can be added (up to a maximum of 64).

The required MSTP configurations for EX Series and Cisco switches shown in Figure 9 are provided below. Note that all switches are in the same MSTP region, and the root switches for each MSTI are on the core switches.

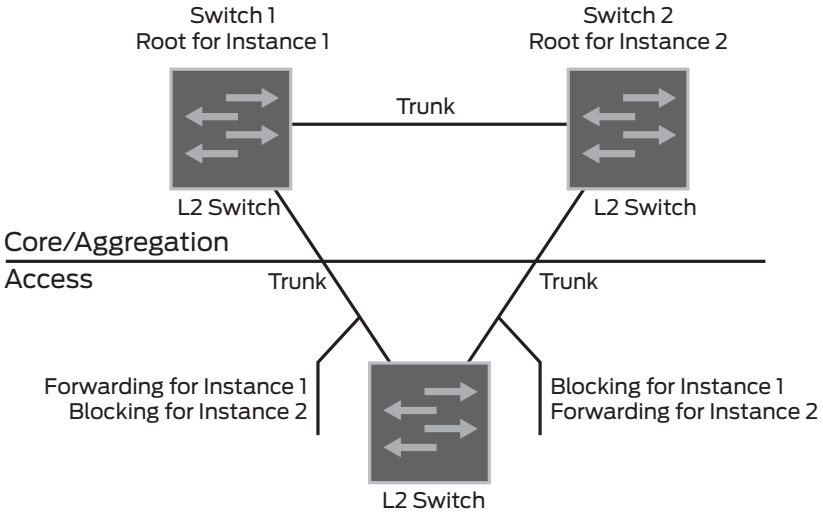

Figure 9: All switches running MSTP

Switch 1, a Cisco switch, has the following configuration:

```
3750E-1(config)#spanning-tree mode mst
3750E-1(config)#spanning-tree mst 0,1 priority 4096
3750E-1(config)#spanning-tree mst configuration
3750E-1(config-mst)#instance 1 vlan 1,3
3750E-1(config-mst)#instance 2 vlan 2,4
```

The output below validates the configurations shown above. Switch 1 is root for MST0 and MST1 because the path cost is 0 and there is no root port, whereas the MST2 root is Switch 2.

```
3750E-1#show spanning-tree root

                                        Root    Hello Max Fwd
MST Instance           Root ID          Cost    Time  Age Dly  Root Port
---------------- -------------------- --------- ----- --- ---  ------------
MST0              4096 001b.d58a.0a00         0    2   20  15
MST1              4097 001b.d58a.0a00         0    2   20  15
MST2              4098 001b.d4ce.c400     20000    2   20  15  Gi1/0/2
```

Switch 2, also a Cisco switch, has the following configuration:

```
3750E-2(config)#spanning-tree mode mst
3750E-2(config)#spanning-tree mst 2 priority 4096
3750E-2(config)#spanning-tree mst configuration
3750E-2(config-mst)#instance 1 vlan 1,3
3750E-2(config-mst)#instance 2 vlan 2,4
```

The output below validates the configurations shown above. Switch 2 is root for MST2 because the path cost is 0 and there is no root port, whereas Switch 1 is root for MST1 and MST0.

```
3750E-2#show spanning-tree root

                                        Root    Hello Max Fwd
MST Instance           Root ID          Cost    Time  Age Dly  Root Port
---------------- -------------------- --------- ----- --- ---  ------------
MST0              4096 001b.d58a.0a00         0    2   20  15  Gi2/0/2
MST1              4097 001b.d58a.0a00     20000    2   20  15  Gi2/0/2
MST2              4098 001b.d4ce.c400         0    2   20  15
```

Switch 3, Juniper Networks EX Series switch, has the following configuration:

```
juniper@EX4200# set protocols mstp msti 1 vlan [v1 v3]
juniper@EX4200# set protocols mstp msti 2 vlan [v2 v4]
```

The output below shows that both links on the EX Series switches are being utilized. The link between Switch 1 and Switch 3 is forwarding for MSTI 1 and blocking for MSTI 2; the link between Switch 2 and Switch 3 is forwarding for MSTI 2 and blocking for MSTI 1.

```
juniper@EX4200> show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/2.0     128:515        128:1   4096.001bd58a0a00  200000000 FWD    ROOT
ge-0/0/3.0     128:516       128:55  32768.001bd4cec400  200000000 BLK    ALT

Spanning tree interface parameters for instance 1

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/2.0     128:515        128:1   4097.001bd58a0a00  200000000 FWD    ROOT
ge-0/0/3.0     128:516       128:55  32769.001bd4cec400  200000000 BLK    ALT

Spanning tree interface parameters for instance 2

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/2.0     128:515        128:1  32770.001bd58a0a00  200000000 BLK    ALT
ge-0/0/3.0     128:516       128:55   4098.001bd4cec400  200000000 FWD    ROOT
```

VLAN Spanning Tree Protocol (VSTP)

VSTP, introduced in Junos OS 9.4R1 for the EX3200 and EX4200 switches, allows users to create separate spanning-tree instances for every VLAN configured (see Figure 10). VSTP BPDUs are transmitted as tagged packets to a multicast destination MAC address of 01-00-0C-CC-CC-CD on 802.1Q trunked ports. On the access port, VSTP BPDUs are sent to the IEEE well-known multicast MAC address (01-80-C2-00-00-00) with the corresponding spanning-tree information for the VLAN to which the access port belongs.

VSTP is equivalent to Cisco's RPVST+ but is backward-compatible with Cisco's PVST+. By default, VSTP BPDUs are equivalent to RPVST+ (for a detailed description of PVST+/RPVST+, please review the Spanning Tree Overview section). However, if the adjacent switch is sending PVST+ BPDUs, then the EX Series switch will send the VSTP BPDU equivalent to PVST+ BPDUs for backward compatibility.

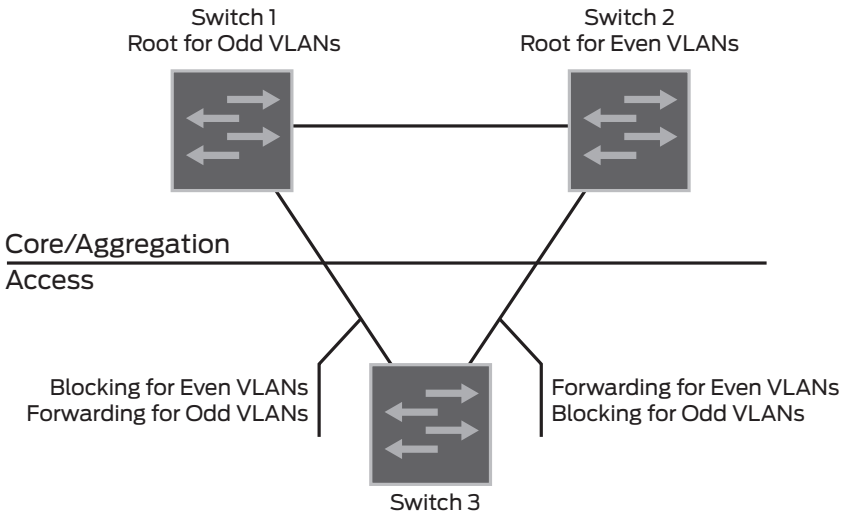

Figure 10: VLAN load-balancing between PVST+ (core/aggregation switches) and VSTP (access switch)

When the LAN infrastructure includes a mix of Juniper and Cisco switches and is configured for per-VLAN spanning-tree, it is recommended that the spanning-tree mode on the Cisco switches be changed from PVST+ to RPVST+. PVST+ does not have the same fast convergence features (portfast, uplinkfast and backbonefast) that are natively built into RPVST+.

Only one spanning-tree protocol can be active on an EX Series switch at any given time. To enable VSTP on the EX Series switches, RSTP-related configurations must first either be deleted or deactivated. Then, the following syntax needs to be implemented on every VLAN that needs to be configured for VSTP:

"set protocols vstp vlan <vlan-id or name>"

Bridge priority, port priority and port cost are configured under the VLAN of the VSTP protocol. VSTP/PVST+ or IEEE BPDUs will not be processed, generated or flooded for any VLAN that is not configured under VSTP.

Below is a sample configuration for enabling VSTP on EX Series switches. In this example there are five VLANs, although only VLANs 2-5 are configured for VSTP. This document will discuss VLAN 1 in the next couple of paragraphs.

```
juniper@EX4200> show configuration protocols
vstp {
    vlan 2;
    vlan 3;
    vlan 4;
    vlan 5;
}
```

Base on the topology in Figure 10 and the above configuration, the following output can be expected on the EX Series switches:

```
juniper@EX4200> show spanning-tree interface

Spanning tree interface parameters for VLAN 2

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/2.0     128:515        128:1   4098.001bd58a0a00  200000000 FWD    ROOT
ge-0/0/3.0     128:516       128:55  32770.001bd4cec400  200000000 BLK    ALT

Spanning tree interface parameters for VLAN 3

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/2.0     128:515        128:1  32771.001bd58a0a00  200000000 BLK    ALT
ge-0/0/3.0     128:516       128:55  4099.001bd4cec400  200000000 FWD    ROOT

Spanning tree interface parameters for VLAN 4

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/2.0     128:515        128:1   4100.001bd58a0a00  200000000 FWD    ROOT
ge-0/0/3.0     128:516       128:55  32772.001bd4cec400  200000000 BLK    ALT

Spanning tree interface parameters for VLAN 5

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/2.0     128:515        128:1  32773.001bd58a0a00  200000000 BLK    ALT
ge-0/0/3.0     128:516       128:55   4101.001bd4cec400  200000000 FWD    ROOT
```

From the above output, note there is a spanning-tree instance for each VLAN configured for VSTP. The forwarding and blocking links match Figure 10; odd-numbered VLANs are blocking on ge-0/0/2.0 and forwarding ge-0/0/3.0 and vice versa for the even-numbered VLANs.

There are a few things about Junos OS 9.4 worth mentioning:

- VSTP is supported on the EX3200 and EX4200 lines of switches only. The EX8200 line will support VSTP in a future release.
- The number of supported VSTP instances in 9.4R1 is 128. This number will be increased to 253 in 9.4R2.
- EX Series switches will not process IEEE BPDUs with Cisco switches on VLAN 1 (please read below for details). This will be supported in future release.

With Cisco switches, VLAN 1 complies with the IEEE spanning-tree specification. Currently the EX Series switches cannot process BPDUs on VLAN 1 with Cisco switches. If there is a trunk port configured between Cisco switches and Juniper EX Series switches, then the VLAN that is configured as the native-vlan than that VLAN (native-vlan) will treat the BPDUs as a regular multicast and flood the BPDU (see Figure 11). Therefore, the EX Series switches will rely on other switches to block the redundant links (see Figure 12).

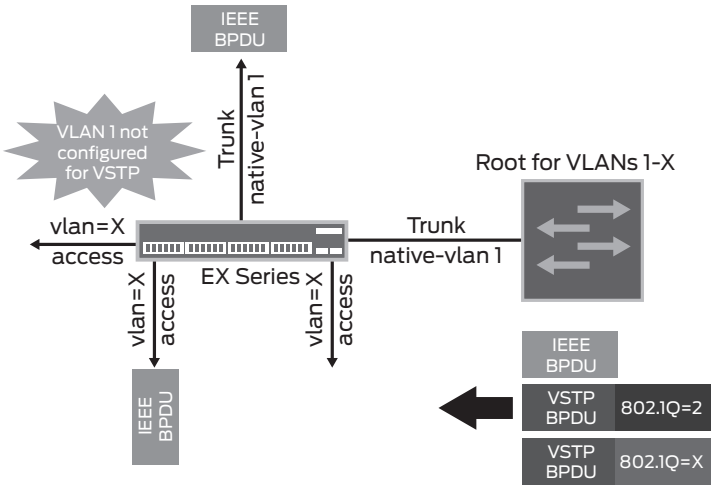

Figure 11: BPDUs received on the native-vlan will be treated as regular multicast packet.

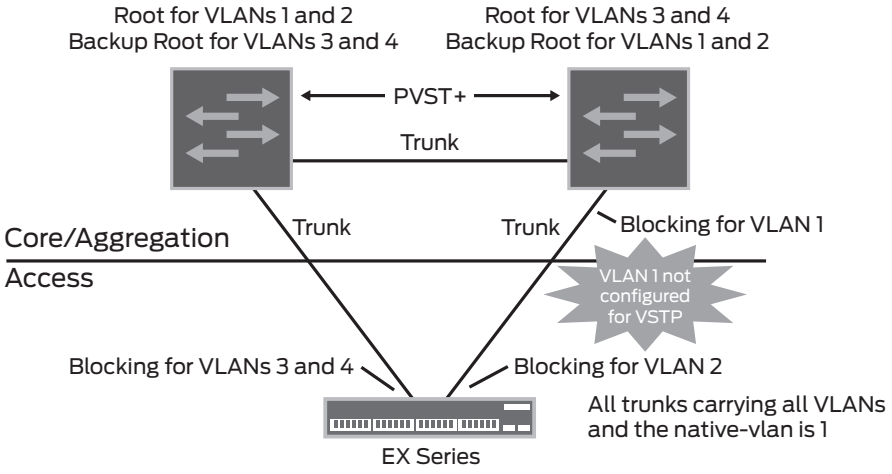

Figure 12: The upper-left switch is blocking the redundant link for VLAN 1.

Conclusion

Juniper Networks EX Series Ethernet Switches can be deployed in either a Layer 2 or Layer 3 environment. However, if a Layer 2 topology is a must, then the Spanning Tree Protocol is necessary to prevent any unwanted L2 bridge loops. This implementation guide has covered the fundamentals and configurations of all three standards-based Spanning Tree Protocols. By following these guidelines, Juniper customers can confidently integrate EX Series switches into their current network infrastructure without compromising desired STP stability, functionality and flexibility.

About Juniper Network

Juniper Networks is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at www.**juniper.net**.

Corporate and Sales Headquarters
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC Headquarters
Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters
Juniper Networks Ireland
Airside Business Park
Swords, County Dublin, Ireland
Phone: 35.31.8903.600
EMEA Sales: 00800.4586.4737
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

8010002-004-EN Jan 2012