

Migrating from Cisco to HP ProCurve Networks

- Introduction 2
- Replacing Cisco hardware with HP ProCurve..... 3
 - Replacing access layer switches 4
 - Replacing distribution layer components..... 5
 - Adding HP ProCurve at the Layer 3 core 6
 - A further option: Using Layer 3 throughout the network 6
- Dealing with nomenclature issues 7
- Configuring VLAN interoperability 8
- Details of Spanning Tree Protocols..... 9
 - Single Instance Spanning Tree Protocol 9
 - PVST: Spanning tree on a per-VLAN basis 10
 - MSTP: Separate STP for each VLAN, with redundant links blocked 11
- Combining Cisco and HP ProCurve spanning tree networks..... 11
 - Software version 11
 - Running MSTP on Cisco 11
 - Achieving PVST-MSTP interoperability 12
 - Spanning tree frames on HP ProCurve and Cisco.. 12
 - Configuring PVST-MSTP interoperability..... 13
 - Hardening spanning tree 13
- Gateway Redundancy Protocols: HSRP/VRRP..... 14
- Interoperability and IP Routing 14
- Summary 15

Introduction

This white paper is a high-level guide showing the basic design and migration considerations when moving from a Cisco network to a standards-based HP ProCurve network, or when interoperating with both Cisco and HP ProCurve networks.

Although there are several areas of concern when migrating from Cisco to HP ProCurve networks, the two most important considerations are virtual LAN (VLAN) setup and Spanning Tree Protocol (STP) interoperability. Because HP ProCurve works to provide standards-based solutions, the routing protocols should conform to those standards.

For gateway or router redundancy, HP ProCurve supports the standards-based Virtual Router Redundancy Protocol (VRRP), not Cisco's proprietary Hot Standby Router Protocol (HSRP). This does not mean that either protocol is inferior. However, interoperability dictates a plan that implements standard protocols. (Note that HSRP and VRRP are not routing protocols, as they do not advertise IP routes or affect the routing table in any way.)

Consider a typical scenario where an installation with Cisco switches is being migrated to a ProCurve network. Here, there are three primary interoperability considerations: 1) VLAN setup, 2) configuration of standards-based spanning tree, and, finally, 3) spanning tree interoperability.

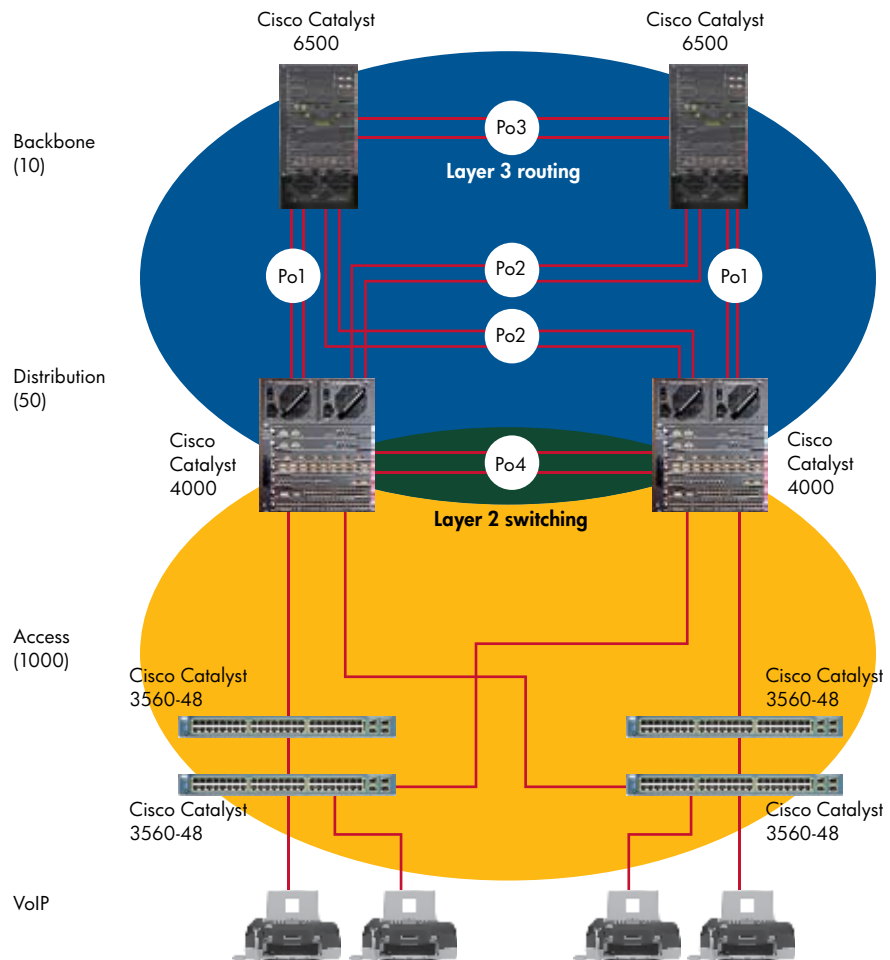
This white paper covers the following topics:

- Migrating from a Cisco infrastructure to an HP ProCurve infrastructure
- VLAN interoperability
- Spanning tree interoperability
- Hardening spanning tree
- IP routing

Replacing Cisco hardware with HP ProCurve

This basic scenario seen in figure 1 shows Cisco hardware throughout the network. In the figure, Cisco Catalyst 3500 XL Series Switches are at the access layer, which, in ProCurve terminology, is the “edge.” Cisco Catalyst 4000 Series Switches are at the distribution layer; these are both Layer 2 and Layer 3 switches. At the core are Cisco Catalyst 6500 Series Switches using pure Layer 3 switches.

Figure 1. Typical enterprise starting point, before migration

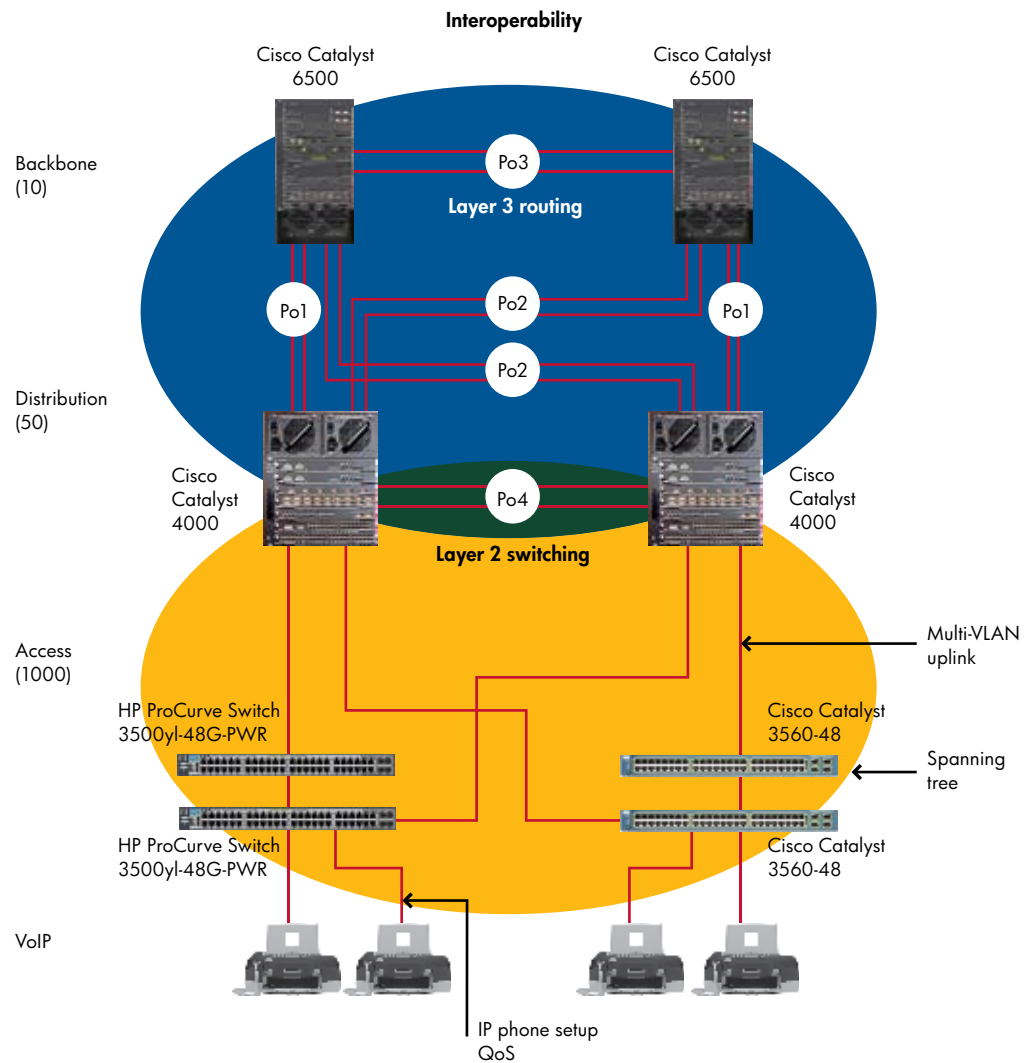


Replacing access layer switches

The first step in migration is to replace the switches at the access layer with HP ProCurve Edge switches (see figure 2). To replace the access layer switches, the HP ProCurve Edge switches must be connected to Cisco Catalyst distribution switches. This leads to some interoperability points, the first of which is the uplink and the VLAN in the uplink. Cisco terminology identifies these as “trunks,” while in ProCurve networks they are known as a “tagged port.”

The other very important consideration is spanning tree. By default, Cisco supports Per VLAN Spanning Tree (PVST), while ProCurve supports only standards-based protocols—in this case, IEEE 802.1s Multiple Spanning Tree Protocol (MSTP) or IEEE 802.1w Rapid Spanning Tree Protocol (RSTP). Connecting Cisco IP phones may require further tuning of Quality of Service (QoS) parameters.

Figure 2. Replacing access layer switches with HP ProCurve switches

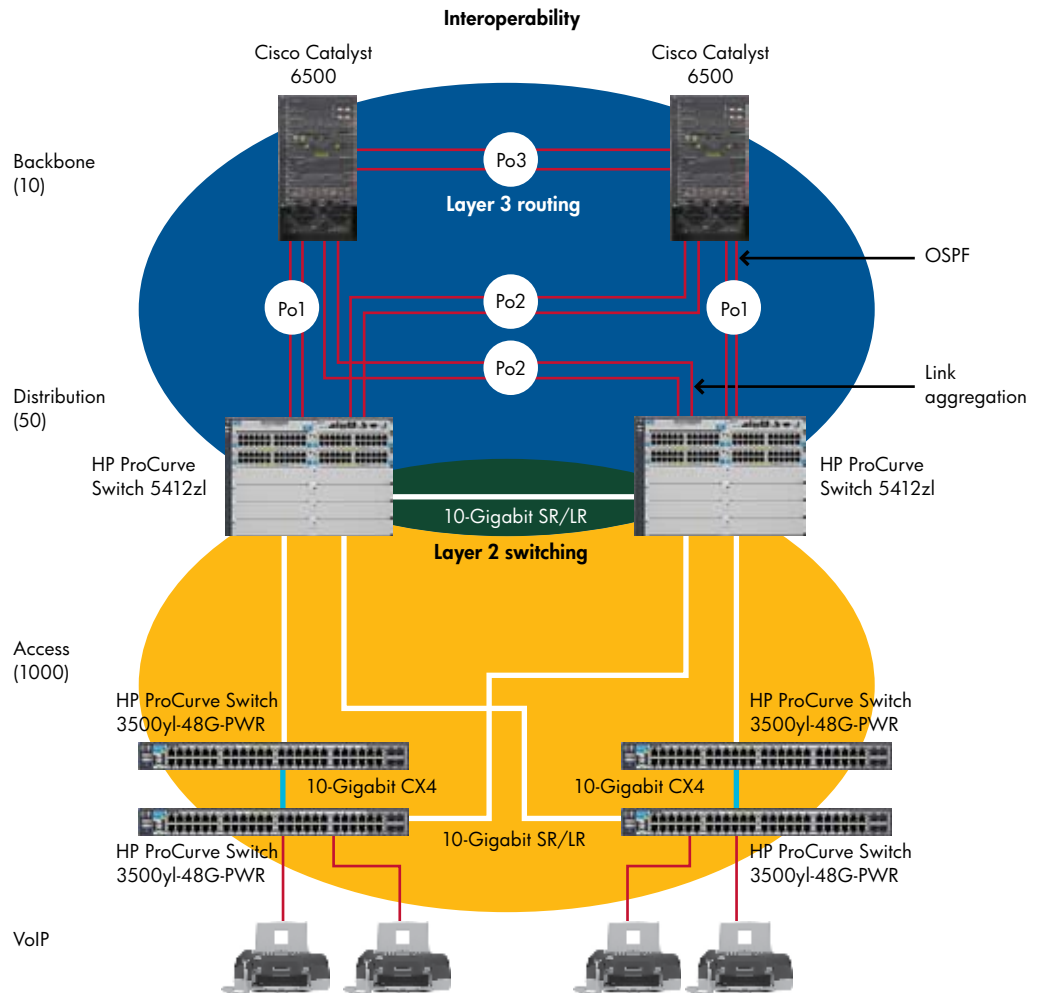


Replacing distribution layer components

The second step of integration is the use of HP ProCurve Switch 5400zl or 8200zl Series at the distribution layer (see figure 3). Now the interoperability concerns shift to the Cisco Catalyst 6500 remaining on the network.

At the distribution layer, interoperability issues are related more to connectivity and routing. Connectivity will typically require “port channel interfaces,” a generic term referring to link aggregation. (In HP ProCurve terminology, link aggregation is known as “trunking.”) Another consideration is the routing protocol, which will usually be Open Shortest Path First (OSPF).

Figure 3. Replacing distribution layer components with HP ProCurve switches



Adding HP ProCurve at the Layer 3 core

The next step in migration is to replace the Catalyst 6500 switches with HP ProCurve 8212zl or 8206zl switches at the Layer 3 core, as illustrated in figure 4. (Depending upon the size of the network, an 8206zl or 8212zl switch could also be used in the distribution layer.)

A further option: Using Layer 3 throughout the network

At this point, an option is to use Layer 3 switching throughout the entire network (see figure 5). Because the network now consists of all HP ProCurve switches and standards-based protocols, interoperability is no longer an issue. The only consideration now is whether this plan makes sense for the customer.

Figure 4. Replacing the Layer 3 core with HP ProCurve switches

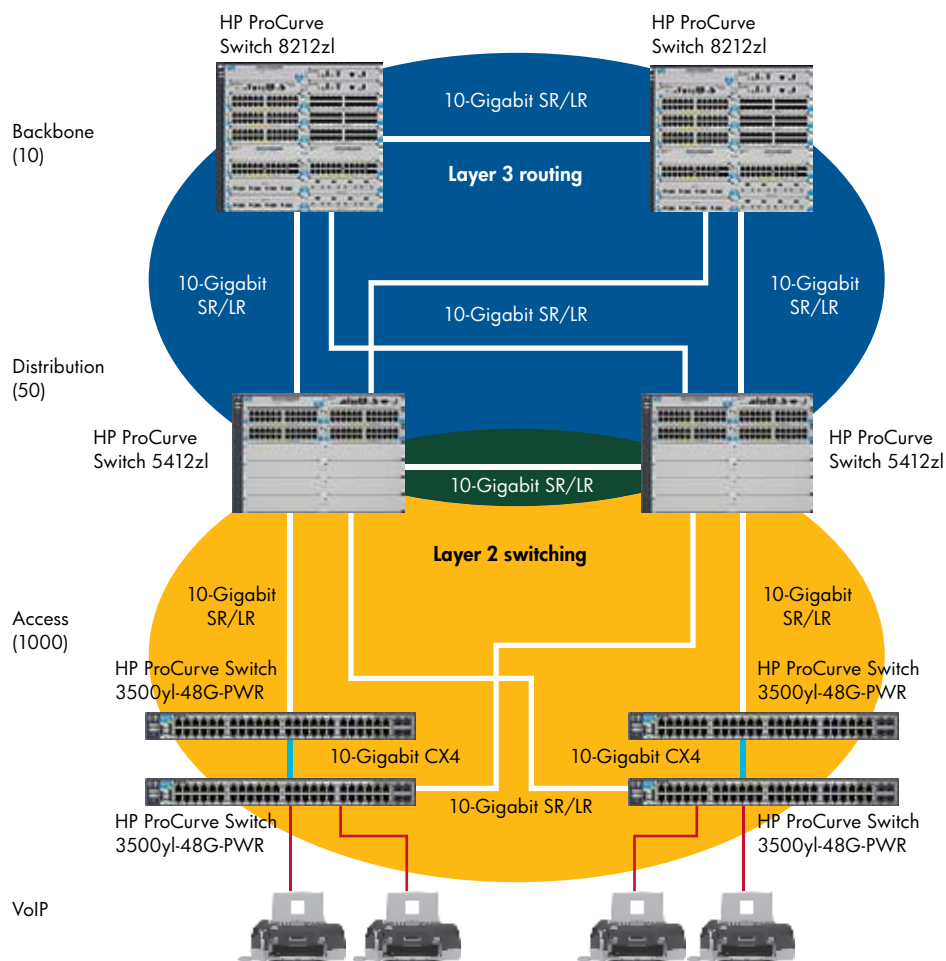
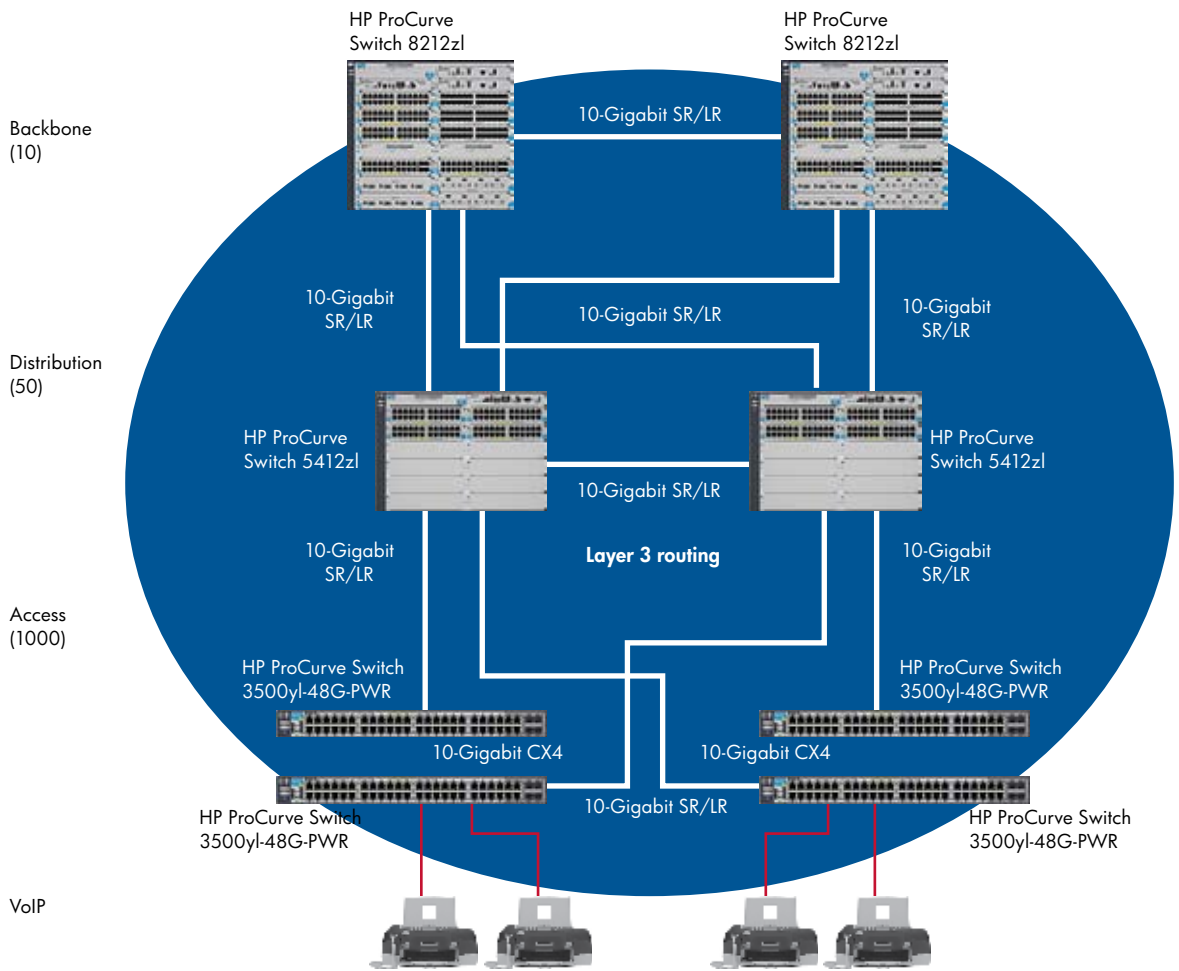


Figure 5. Using Layer 3 routing throughout the network



Dealing with nomenclature issues

Table 1 details the basic terminology differences to be aware of during a migration. Engineers performing the migration will need to keep these differences in mind.

Table 1. Important nomenclature differences between HP ProCurve and Cisco networks

Switch port role	Cisco	HP ProCurve
End-user ports (PCs, printer,...)	Access port, VLAN	Port in one VLAN only (untagged port)
IP phone ports	Use auxiliary VLAN (voice)	Use voice VLAN (tagged or untagged)
End-user + IP phone ports	Access port with auxiliary VLAN (voice) Multi-VLAN access port (MVAP)	Use untagged VLAN for end user and tagged voice VLAN for IP phone
Server ports for multiple VLANs	Trunk port	Port tagged in multiple VLANs
Switch-to-switch ports for multiple VLANs	Trunk port	Port tagged in multiple VLANs
Aggregated ports	Port channel interface (Po)	

Remember these general guidelines:

- An “access port” on Cisco is an “untagged port” on ProCurve.
- A “trunk port” on Cisco is a “tagged port” on ProCurve.
- The “port channel” on Cisco is “link aggregation” and is called a “trunk” on ProCurve.
- When a single end-user device such as a PC is connected, the connection from a VLAN point of view is untagged. Cisco refers to this kind of port as an “access port.”
- When a phone and cascaded PC are connected, normally one VLAN is untagged for the data traffic of the PC, and the other is tagged for the voice traffic of the phone. On Cisco devices, this is called an “access-port with auxiliary VLAN.” Note that Cisco recently changed this nomenclature and now calls it a “multi-VLAN access port.”
- Trunking from the ProCurve side is meant to aggregate multiple ports together, while on Cisco it is meant to transport multiple VLANs over one port.
- Link aggregation on the Cisco side is called “channeling.”

The differences are often primarily a variation in terminology. With regard to standards-based switching and routing, ProCurve and Cisco hardware often perform the same functions, but different language is used to describe them.


Configuring VLAN interoperability

Figure 6 illustrates the many types of connections to bear in mind when configuring VLAN interoperability. For example, consider an end-user port, one that may have a PC or printer attached to it. On Cisco, this is called an “access port,” while on HP ProCurve switches it is referred to as “untagged.” In the end, any frame sent by the PC is just an Ethernet frame.

Another type of device is an IP phone. Usually the IP phone will use tagged traffic, but it could be untagged also. When an end-user device and IP phone are combined, data traffic is untagged, while the phone’s voice traffic is tagged.

Figure 6. Types of connections to consider for VLAN interoperability

- End-user ports (PCs, printer)
- IP phone ports
- End-user + IP phone ports
- Server ports for one VLAN
- Server ports for multiple VLANs
- Switch-to-switch ports with one VLAN
- Switch-to-switch ports with multiple VLANs
- Aggregated ports



Details of Spanning Tree Protocols

The spanning tree operation is the point in the process that requires extra care when planning a migration. Understanding the issues of interoperability will help to streamline the process.

The original Spanning Tree Protocol is defined in the IEEE 802.1D Standard. This protocol creates a spanning tree within a mesh network of switches, disabling links that are not part of that tree. The result is a single active path between any two network nodes.

Note that the ratified IEEE 802.1 standards use an uppercase letter denoting a major release that incorporates up to three revisions. Revisions are officially denoted by lowercase letters. Confusion sometimes results because of typographical errors in documentation.

The primary standard is IEEE 802.1D-1998, and is called Spanning Tree or IEEE Spanning Tree. The other, an update to 1998, is IEEE 802.1w, known as Rapid Spanning Tree. A related standard is IEEE 802.1Q, a trunking standard that allows multiple bridged networks to transparently share the same physical network link without leakage of information between networks. Three of the standards referenced in the following discussion of PVST and MSTP interoperability are either parts of the IEEE 802.1D-1998 or 802.1Q standards.

Single Instance Spanning Tree Protocol

The standard Spanning Tree Protocol (IEEE 802.1D-1998) is single instance. This version of STP is not aware of VLANs, and this is also true for MSTP (IEEE 802.1Q-2002 or 802.1s). As shown in figure 7, a port that is blocked is blocked for all VLANs. That is, when one instance is blocked, the entire port is blocked, blocking all VLANs that may be carried on that port.

A server can be on an untagged port—an access port in HP ProCurve—or it can use multiple VLANs, in which case its traffic is tagged. Similarly, the uplink, or switch-to-switch port, can be in one VLAN or in multiple VLANs. In fact, multiple VLANs are very common; this is known as a “tagged port” on HP ProCurve switches, or a “trunk port” on Cisco. In either case, it means that the port carries multiple VLAN traffic. The aggregated port, or the link aggregation, is a trunk on HP ProCurve switches, while on Cisco it is a “port channel,” or simply a “channel.”

When planning interoperability between Cisco and HP ProCurve switches, it is important to consider the multi-VLAN ports. Remember that the Cisco “trunk” is typically the native VLAN, and that the native VLAN on a Cisco trunk must map the untagged VLAN on the HP ProCurve tagged port.

Another best practice is to make sure that the same VLAN is allowed and configured on both sides. That is, the allowed VLANs on the Cisco trunk should match the tagged and the untagged VLANs on HP ProCurve switches.

One very common question that arises when configuring Cisco trunk and HP ProCurve tagged is, “Do I need VLAN 1?” On Cisco switches, VLAN 1 is very important; for example, it carries IEEE STP Bridge Protocol Data Units (BPDUs) and frames such as VLAN Trunk Protocol (VTP).

On most HP ProCurve switches, especially those designated “zl” and “yl,” frames such as STP BPDUs are sent between the switches untagged and separate from any VLAN assignment. Although they are found between the switches and use specific MAC addresses known by the switches, BPDUs are part of the “control plane” of information. They are not part of the data and are not attached to any specific VLAN.

However, when BPDU data is received by a Cisco switch, it will be attached to the native VLAN, so it will be sent by Cisco the same way, into the native VLAN. This is important for understanding interoperability in spanning tree operations.

Hints and tips:

- Pay attention to multi-VLAN ports.
- Make sure that the native VLAN on the Cisco trunk is the untagged VLAN on the ProCurve tagged port.
- Ensure that the same VLANs are allowed and configured on both sides.
- Remember that, unlike with Cisco, BPDUs (spanning tree, LLDP, and LACP) are not attached to the untagged port or any VLAN on HP ProCurve.

PVST: Spanning tree on a per-VLAN basis

In the switching world, a loop can cause a condition often referred to as a “broadcast storm,” and these loops typically occur within a VLAN. With Cisco Per VLAN Spanning Tree (PVST), ports are blocked on a per-VLAN basis, with one instance of spanning tree per VLAN. As shown in figure 8, this allows a different root spanning tree to be defined for each VLAN, with, for example, one root for VLAN 1, one root for VLAN 2, and one root for VLAN 3.

Note: Cisco devices actually use Per VLAN Spanning Tree Plus (PVST+), which provides the same functionality as PVST but uses IEEE 802.1Q trunking technology rather than the now-obsolete

Inter-Switch Link (ISL). According to Cisco, PVST+ is an enhancement to the IEEE 802.1Q specification and is not supported on non-Cisco devices.

Cisco can also use Rapid-Per-VLAN-Spanning Tree (Rapid-PVST), a proprietary protocol that allows a switch port to transition more rapidly than PVST+.

PVST solves the problem of a loop in the VLAN. However, PVST has the drawback of requiring one configuration per VLAN. Furthermore, PVST is very CPU intensive, because it uses a one-to-one mapping of VLANs and STP. For instance, 1,000 VLANs will require 1,000 instances of STP.

Figure 7. The original Spanning Tree Protocol, in which a port is blocked for all VLANs

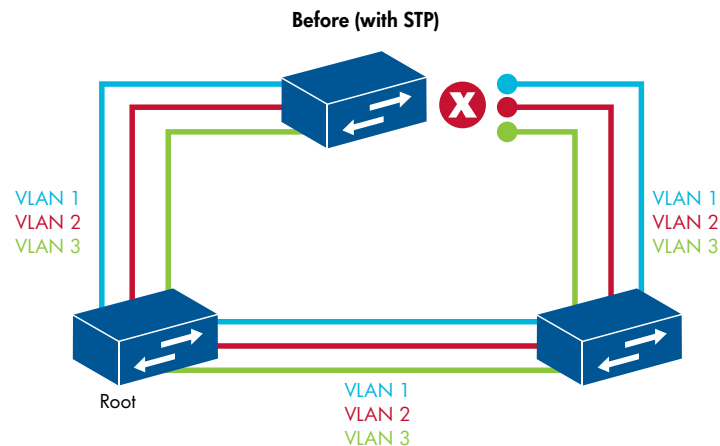



Figure 8. PVST+ allows separate instances for each VLAN

- PVST tags BPDUs so that it can run separate instances for each VLAN.
- This yields a 1 to 1 mapping of VLAN and STP instance, which means many separate processes.
- This is CPU intensive on what are already typically overburdened CPUs.



MSTP: Separate STP for each VLAN, with redundant links blocked

Multiple Spanning Tree Protocol (MSTP) allows better use of VLANs. Originally defined in IEEE 802.1s and later merged into IEEE 802.1Q-2003, MSTP is a “per-VLAN” protocol that configures a separate spanning tree for each VLAN group and blocks the links that are redundant within each spanning tree.

As shown in figure 9, the left-hand switch is the root for spanning tree instance 1, and the switch on the right is the switch for instance 2. The ports are blocked on a per-instance basis. If VLANs are set on this port, then the VLANs that are mapped into this instance are blocked. Thus, blocking is not on a per-VLAN basis, but on a per-instance basis.

Instances are chosen to match the number of possible paths through the Layer 2 network. This is usually only two or three instances. Setup is necessary only for these instances, making it comparatively simple compared to the hundreds of instances needed for PVST.

Combining Cisco and HP ProCurve spanning tree networks

To combine Cisco and HP ProCurve spanning tree networks, MSTP can be run on the Cisco devices, or PVST Cisco networks can be combined with MSTP HP ProCurve networks.

Software version

To support the compliant IEEE 802.1s-2002 standard, ensure that the Cisco devices have the latest software installed. Check the latest Cisco documentation for details.

Running MSTP on Cisco


When combining Cisco and HP ProCurve networks by running MSTP on Cisco devices, make sure that all parameters correspond.

This includes the same set of parameters, the same names, the same revision, and the same VLAN-to-instance mapping on all switches.

Hints: Use these hints and tips to help interoperate with MSTP on Cisco and HP ProCurve switches:

- Set the MSTP configuration parameters on both networks to be identical.
- Ensure that the same name, revision number, and mapping are used between VLANs and instances.
- To obtain standard MSTP BPDU, use trunk ports on Cisco uplinks.
- If an untagged uplink is required, do not use an access port. Instead, define the Cisco port as a trunk and allow only the native VLAN.
- Note that Cisco supports a pre-version of MSTP, which appears the same as MSTP. The difference is not apparent in the command set. However, this pre-version does not interoperate with standard MSTP. Ensure that the Cisco device has the right version of IOS, and the real MSTP, not the pre-MSTP version.

Figure 9. MSTP allows Spanning Tree Protocol with per-instance configuration



Achieving PVST-MSTP interoperability

In general, HP ProCurve attains interoperability with Cisco PVST+ configurations as follows:

- IEEE 802.1D with PVST+ in native VLAN on Cisco with VLAN 1 allowed on the Cisco trunk
- IEEE 802.1w with Rapid PVST+ in native VLAN on Cisco with VLAN 1 allowed on the Cisco trunk

There are two primary factors that govern interoperability. First, the chosen Spanning Tree Protocols must use a standard IEEE untagged BPDU format. Second, the untagged BPDU in the correct format must have the IEEE standard destination MAC address.

Spanning tree frames on HP ProCurve and Cisco

Understanding the frames employed by the different protocols will provide a better understanding of the details of STP interoperability. Figure 10 shows the frames for several different spanning tree BPDUs.

The top frame in the figure represents an IEEE 802.1D-1998 spanning tree frame. It is untagged and uses a standard MAC address. The next frame, labeled IEEE 802.1w, is a frame for rapid spanning tree; just as is the case with the top frame, this is also untagged and also uses a standard MAC address.

MSTP is also untagged and uses a standard MAC address. However, the MSTP BPDU includes information regarding instances (shown on the right of figure 10), as well as information regarding Rapid Spanning Tree Protocol (RSTP) and Common Spanning Tree (CST). This information is how MSTP interacts with other switches, whether they are running CST or RSTP. By use of this common information, the BPDU can be understood by all switches. HP ProCurve switches send the standard BPDU, and the switches are able to use and recognize the Cisco switch when the Cisco switch has a standard implementation of spanning tree.

Frames on Cisco running PVST+ are different than the IEEE 802.1D standard. The bottom of figure 10 shows a representation of the essential information in these frames that is used in PVST+ operation. (Note that PVST+ BPDUs are not the same between tagged and untagged VLAN versions.)

Figure 10. Frames for various spanning tree BPDUs

IEEE 802.1D 1998	Untagged IEEE destination MAC: 01:80:c2:00:00:00			
IEEE 802.1w	Untagged IEEE destination MAC: 01:80:c2:00:00:00			
MSTP (IEEE 802.1s)	Untagged IEEE destination MAC: 01:80:c2:00:00:00	RSTP and MSTP common CST information	MSTP-specific parameters IST info. MSTI info. ... additional MSTI info.	
PVST+ on Cisco trunk ports	Untagged for native VLAN Cisco destination MAC: 01:00:0c:cc:cc:cd	Untagged for native VLAN (only if VLAN 1 is allowed on the trunk) IEEE destination MAC: 01:80:c2:00:00:00	Tagged Cisco destination MAC: 01:00:0c:cc:cc:cd	

If PVST+ BPDUs are sent out on a trunk port as tagged, they use a Cisco destination MAC address, as shown on the right side of the figure. The Cisco MAC address is not known by ProCurve switches. If these BPDUs are received into an HP ProCurve tagged port containing the same VLAN IDs, the HP ProCurve switch will forward these BPDUs.

The center of the frame information for PVST+ in figure 10 shows that if the VLAN is untagged (as with the native VLAN on Cisco), and if VLAN 1 is allowed (not necessarily the native VLAN, but with just VLAN 1 allowed), the native VLAN will allow untagged BPDUs to be sent. This untagged BPDU will also contain the IEEE destination MAC address recognized by non-Cisco devices such as HP ProCurve switches. This is the most common customer use model, because VLAN 1 carries other information, such as CDP and VTP, that is useful to customer implementations. As a result, most customers allow VLAN 1 on their trunk (multi-VLAN link).

With VLAN 1 allowed on the trunk and a native VLAN, the BPDU will be in the IEEE standard format. Thus, the PVST+ BPDU will interact with HP ProCurve switches, because, in this case, both will be using the IEEE standard MAC address and use untagged BPDUs.

If VLAN 1 is not allowed, then the BPDU will be untagged. However, it will be the Cisco MAC address, not the standard MAC address. In this case, the HP ProCurve switch will not understand this BPDU; the HP ProCurve switch will not see itself as a destination and will therefore forward the Cisco proprietary BPDU.

What this means: For interoperability between PVST+ and the IEEE standard spanning tree, the Cisco device needs a configured native VLAN and, a very important point, *VLAN 1 must be allowed on the trunk port.*

Configuring PVST-MSTP interoperability

For PVST-MSTP interoperability, keep these points in mind:

- On Cisco trunk inter-switch links, make sure that VLAN 1 is allowed. (Otherwise, non-standard BPDU will be sent in the native VLAN.)
- Take special care of the root and secondary root setup into the native VLAN, as Cisco and HP ProCurve switches will interoperate through the standard BPDUs.
- To obtain faster convergence, set Rapid-PVST instead of PVST+ on Cisco.
- On Cisco hardware, be sure to use the “path cost long” method.

Hardening spanning tree

Spanning tree instability can be caused by a number of factors, including:

- Uni-directional links
- Rogue devices attempting to communicate with STP
- Permanent STP topology changes due to flapping ports, or end-user ports not set to edge mode (portfast)
- Loops not detected by STP

Spanning tree can be secured or hardened. The techniques are virtually the same on HP ProCurve and Cisco devices; however, the names, although similar, are often different. For example, what is called “loop-protect” on HP ProCurve is called “keepalive” on Cisco. And BPDU-protection on HP ProCurve is called BPDU guard on Cisco. Table 2 shows the major differences in terminology.

Table 2. Nomenclature differences for securing STP

HP ProCurve	Cisco
Remote Fault Notification (RFN) using autonegotiation	Remote Fault Notification (RFN) using autonegotiation
Uni-Directional Link Detection (UDLD)	Uni-Directional Link Detection (UDLD)
BPDU-protection	BPDU guard
Loop-protect	Keepalive
Root-guard	Root guard
—	Loop guard

Gateway Redundancy Protocols: HSRP/VRRP

Cisco supports Hot Standby Router Protocol (HSRP), and HP ProCurve supports Virtual Router Redundancy Protocol (VRRP). These protocols do not interoperate. In the process of migration, consider replacing the two Cisco core devices with two HP ProCurve devices at the core during the same scheduled outage.

Although Cisco ostensibly supports VRRP, its version of the protocol does not appear to interoperate with HP ProCurve switches, and may not interoperate with devices from other vendors. Thus, setting VRRP on the Cisco switch is not an option for interoperability with HP ProCurve switches.

To sum up, ensure that both cores are from the same vendor, whether HP ProCurve or Cisco. If you replace one core, replace the other at the same time.

Interoperability and IP Routing

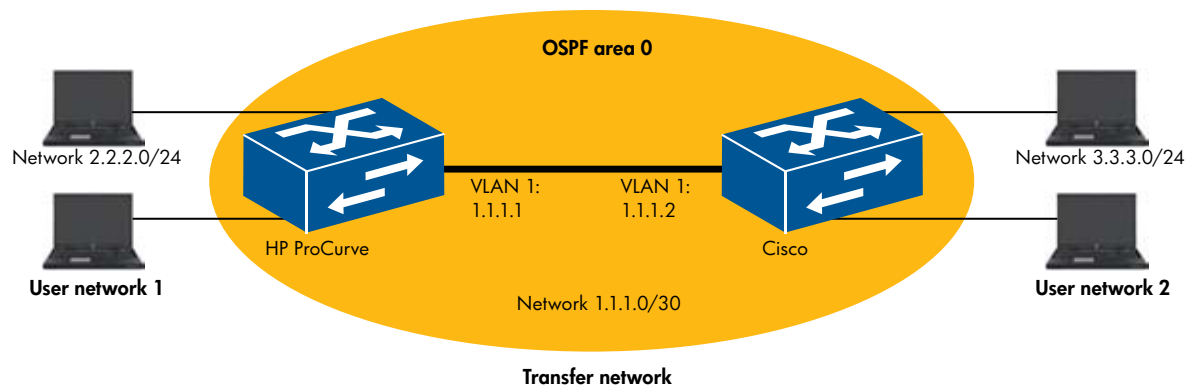
HP ProCurve does not support Cisco proprietary protocols. For the two systems to interoperate, the Cisco device must run standard dynamic routing protocols that HP ProCurve does support.

Information in the routing table may be statically supplied, by defining IP addresses on VLANs or by defining static routes to remote networks. It can be dynamically learned via route exchange protocols, such as Routing Information Protocol (RIP), Open Shortest Path First (OSPF), or Broader Gateway Protocol Version 4 (BGP4). It may also be redistributed between all sources of routing information.

The routing table is filled with information about reaching other networks. This information comes from the IP interface configuration, from static routes, and from RIP and OSPF. (With HP ProCurve switches, BGP4 is currently supported only in the HP ProCurve Routing Switch 9300m Series.) HP ProCurve switches can also redistribute the traffic into RIP or the reverse.

Figure 11 shows a simple example of an interoperating configuration, with HP ProCurve switches on the left side and Cisco on the right side.

Figure 11. Example of HP ProCurve–Cisco interoperability using OSPF



There are some minor differences to consider between the two different OSPF implementations:

- Cisco OSPF is enabled with network statements globally.
- HP ProCurve OSPF is enabled within the VLAN context.
- There are redistribution differences.
- HP ProCurve is always non-broadcast multiple access (NBMA).
- Cisco uses the highest loopback IP address for router ID, while HP ProCurve devices typically use the lowest.
- With HP ProCurve, the loopback is always /32 mask.
- With HP ProCurve, the OSPF link cost is "1" by default.

These differences can be dealt with during implementation planning.

Summary

Whether migrating from Cisco to HP ProCurve networks or interoperating Cisco- and HP ProCurve-based networks, the important areas of concern should be VLAN setup and Spanning Tree Protocol. Migration entails a step-by-step process, beginning with replacement of the access layer switches, followed by replacement of distribution layer components, and, finally, the addition of HP ProCurve at the Layer 3 core.

Nomenclature is likely to be an issue, because commands on HP ProCurve and other standards-based devices sometimes use different terminology than Cisco uses. Furthermore, the choice of Spanning Tree Protocol will be important, with Multiple Spanning Tree Protocol (MSTP) providing the best results on the HP ProCurve platform. This choice, in turn, calls for some additional choices on the Cisco side: whether to run MSTP on Cisco, which may require an updated version of IOS, or to interoperate PVST on Cisco with MSTP on HP ProCurve devices. As for routing, the HP ProCurve standards-based solution should dictate that the routing protocols conform to those standards.

Technology for better business outcomes

To learn more, visit www.hp.com/go/procurve

© Copyright 2008, 2009 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

4AA1-8198ENW Rev. 1, July 2009

