HP ProCurve and Cisco Catalyst Interoperability

Link Aggregation and VLANs

Paul Crandell
Senior Systems Engineer
ProCurve Networking
Hewlett-Packard Company
January 2002

Overview

The purpose of this paper is to outline the interoperability issues when connecting HP ProCurve and Cisco Catalyst switches. This part of the paper is written for a user who is familiar with Cisco products and wishes to interoperate with HP products.

The vast majority of interoperability issues stem from the use of these features:

- Link aggregation
- VLANs

In order to discuss the interoperability issues, we need to lay some groundwork by resolving some differences in the way HP (and most of the industry) and Cisco use the term “trunking”.

HP Definition

Trunking: To HP users, a trunk is a single logical, higher bandwidth link formed by bundling several links together. This technology is also known as port aggregation or link aggregation. Cisco users typically call this an EtherChannel, channel or bundle. For the sake of clarity in this paper, we will use the term ***port trunk*** for this feature.

Cisco Definition

Trunking: To Cisco users, a trunk is a link between two switches carrying traffic belonging to multiple VLANs. HP has no particular name for this feature. For the sake of clarity in this paper, we will use the term ***VLAN trunk*** for this feature.

Link Aggregation

Switched networks often require more bandwidth between two switches than a single link can provide. To address this issue, we can apply link aggregation, port aggregation, or port trunking. All three of these terms are synonymous; all three refer to the ability to combine multiple links into a single logical link.

Most link aggregation implementations allow up to four ports of like speed and media (copper, fiber optic, etc.) to be connected so that they function as a single, higher-speed port between two devices. Link aggregation trunks can be used as switch-to-switch links or switch-to-server links. Generally, the links in a trunk must be of the same speed and media.

As stated above, link aggregation is a technology used to bundle multiple links together into a single logical link. This technology is primarily used to provide increased bandwidth between two switches. In a link aggregation trunk, up to four links can be bundled together providing up to four times the bandwidth.

Most switch vendors provide some technique for combining links in this fashion. HP switches support two specific implementations—namely HP Port Trunking and Cisco FEC (Fast EtherChannel).

Cisco (actually Kalpana, then purchased by Cisco) developed a proprietary technology for bundling links together called EtherChannel, Fast EtherChannel, or Gigabit EtherChannel depending on the link speed. EtherChannel is the only link aggregation technology available on Cisco switches.

Manually Configured Port Trunks

HP Switches support two types of port trunks: HP Trunking (HPT) and Fast EtherChannel (FEC). An HPT type port trunk interoperates with manually configured FEC port trunks on Cisco switches. In this example, a 4-port manual port trunk is created between an HP switch and the Cisco Catalyst 6500 and 3500XL.

The following is the HP Switch CLI configuration file showing a manually configured HP Trunk type port trunk. The critical commands in creating the port trunk are in red.

HP Switch Configuration

```
; J4865A Configuration Editor

hostname "HP ProCurve Switch 4108GL"

time daylight-time-rule None

cdp run

interface A1
 no lacp
 exit
interface A2
 no lacp
 exit
interface A3
 no lacp
 exit
interface A4
 no lacp
 exit

trunk A1-A4 Trk1 Trunk

ip default-gateway 10.1.1.1

snmp-server community "public" Unrestricted

vlan 1
 name "DEFAULT_VLAN"
 untagged A5-A12,A17-A24,C1-C24,D1-D24,E1-E24,G1-G24,Trk1
 no untagged A13-A16
 ip address 10.1.1.12 255.255.255.0
 exit
```

Port trunk type FEC can also be used in this configuration. This would change the keyword “Trunk” to “FEC”. FEC mode supports dynamic configuration of the partner Cisco switch as described in the next section.

Cisco Catalyst 4000/5000/6000 Configuration

The following is the Cisco Catalyst 4000/5000/6000 CatOS configuration file. This file manually configures a FEC port trunk. The critical commands in creating the port trunk are in red.

```
begin
!
# ***** NON-DEFAULT CONFIGURATION *****
!
!
#time: Thu Oct 18 2001, 22:33:42
!
#version 6.1(3)
!
#!
#vtp
set vtp domain pctests
set vlan 1 name default type ethernet mtu 1500 said 100001 state active
set vlan 10,20,30,40
!
#ip
set interface sc0 1 140.1.1.11/255.255.255.0 140.1.1.255
!
# default port status is enable
!
#module 1 : 2-port 1000BaseX Supervisor
!
#module 2 : 4-port Multilayer Switch
!
#module 3 : 8-port 1000BaseX Ethernet
!
#module 4 : 48-port 10/100BaseTX Ethernet
set port channel 4/1-4 mode on silent
end
```

Cisco Catalyst 2900XL/3500XL Configuration

The following is a Cisco Catalyst 2900XL/3500XL configuration file. This file manually configures a FEC port trunk using destination address based load balancing (default). The critical commands in creating the port trunk are in red.

```
version 12.0
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname Switch
!
ip subnet-zero
!
interface FastEthernet0/1
 port group 1 distribution destination
!
interface FastEthernet0/2
 port group 1 distribution destination
!
interface FastEthernet0/3
 port group 1 distribution destination
!
interface FastEthernet0/4
 port group 1 distribution destination
!
interface FastEthernet0/5
!
interface FastEthernet0/6 (same through 0/24)
!
interface FastEthernet0/24
!
interface GigabitEthernet0/1
!
interface GigabitEthernet0/2
!
interface VLAN1
 ip address 10.1.1.15 255.255.255.0
 no ip directed-broadcast
 no ip route-cache
!
!
ip default-gateway 10.1.1.1
!
line con 0
 transport input none
 stopbits 1
line vty 5 15
!
end
```

Cisco IOS Routers

Some Cisco IOS Routers, such as the 7000, support FastEtherChannel link aggregation. The following commands create a 4 port link aggregation trunk. Critical commands are in red.

```
interface port-channel 1
  ip address 10.1.1.15  255.255.255.0
  exit
!
interface fastethernet 0/0
  no ip address
  channel-group 1
  exit
!
interface fastethernet 0/1
  no ip address
  channel-group 1
  exit
!
interface fastethernet 0/3
  no ip address
  channel-group 1
  exit
!
interface fastethernet 0/4
  no ip address
  channel-group 1
  exit
```

Dynamic Port Trunks

Some Cisco switches use a proprietary protocol called PAgP (Port Aggregation Protocol) to dynamically configure channels (port trunks). Cisco Catalyst 4000/5000/6000 switches support PAgP. Catalyst 2900XL/3500XL switches do not. In PAgP a port can be in one of 4 modes:

- Off: prevents the port from channeling.
- On: forces the port to channel without PAgP (manually configured). With the **on** mode, a usable EtherChannel exists only when a port group in **on** mode is connected to another port group in **on** mode.
- Desirable: places a port into an active negotiating state, in which the port initiates negotiations with other ports by sending PAgP packets.
- Auto: places a port into a passive negotiating state, in which the port responds to PAgP packets it receives but does not initiate PAgP packet negotiation. (Default)

Only auto and desirable are PAgP modes. When a port is in desirable or auto mode, PAgP transmits special frames to negotiate the EtherChannel. After PAgP identifies correctly matched EtherChannel links, it groups the ports into an EtherChannel. Ports configured in on or off mode do not transmit or listen for PAgP frames.

You can modify the **auto** and **desirable** modes with the **silent** and **non-silent** keywords. Silent is the default. These keywords deal with configuring Spanning Tree on the Cisco. **Silent** is the correct keyword for use with HP Switches. Auto silent is the default configuration for all ports.

When ports on an HP switch are configured as FEC channel (port trunk), the HP switch transmits the special PAgP negotiation frames. A PAgP-capable Cisco switch at the other end will dynamically configure the channel if the ports are in auto or desirable mode. This is advantage to configuring the HP for FEC instead of HP Trunking.

The following table describes how HP and Cisco switches must be configured when interoperating. Note that the HP switch must be manually configured in all cases because of an issue discussed in the next section on LACP.

	PAgP capable switch: (Catalyst 4000/5000/6000)	Non-PAgP capable switch: (Catalyst 2900XL/3500XL)
HP Switch / HP Trunking	HP: manual configuration Cisco: manual configuration	HP: manual configuration Cisco: manual configuration
HP Switch / FEC	HP: manual configuration Cisco: automatic configuration if in the Desirable or Auto state	HP: manual configuration Cisco: manual configuration

LACP on HP Switches

HP Switches support the dynamic creation of port trunks using IEEE 802.1ad standard LACP (Link Aggregation Control Protocol). A complete discussion of this protocol is

beyond the scope of this paper, however, LACP functions similarly to PAgP. Currently, Cisco does not support LACP on any of its switches.

Unfortunately, LACP and PAgP do not interoperate. HP ProCurve and Cisco Catalyst switches, therefore, cannot ***both*** dynamically create a port trunk when interoperating.

On the HP switch, LACP must be disabled on the ports you will be using for the port trunk to the Cisco. This is shown in the example configuration files above and below. The Menu interface on HP switches does this automatically for you when you configure an HP Trunk or an FEC port trunk.

HP Switch Configuration

In this example, a 4-port semi-dynamic port trunk is created between the HP and Cisco Catalyst 4000/5000/6000 switches. Semi-dynamic means that the HP switch is manually configured while the Cisco switch is dynamically configured. Note that some extraneous commands have been edited from the example configuration files for the sake of brevity.

The following is the HP Switch CLI configuration file showing a manually configured FEC port trunk. The critical commands in creating the port trunk are in red.

```
 interface A1
 no lacp
 exit
interface A2
 no lacp
 exit
interface A3
 no lacp
 exit
interface A4
 no lacp
 exit

trunk A1-A4 Trk1 FEC

ip default-gateway 10.1.1.1

vlan 1
 name "DEFAULT_VLAN"
 untagged A5-A12,A17-A24,C1-C24,D1-D24,E1-E24,G1-G24,Trk1
 no untagged A13-A16
 ip address 10.1.1.12 255.255.255.0
 exit
```

Cisco Catalyst 4000/5000/6000 Configuration

The following is the Cisco Catalyst 4000/5000/6000 CatOS configuration file. The critical commands in creating the port trunk are in red.

```
#vtp
set vtp domain pctests
set vlan 1 name default type ethernet mtu 1500 said 100001 state active
set vlan 10,20,30,40
!
#ip
set interface sc0 1 140.1.1.11/255.255.255.0 140.1.1.255
!
# default port status is enable
!
#module 1 : 2-port 1000BaseX Supervisor
!
#module 2 : 4-port Multilayer Switch
!
#module 3 : 8-port 1000BaseX Ethernet
!
#module 4 : 48-port 10/100BaseTX Ethernet
set port channel 4/1-4 mode auto silent (this is the default state)
end
```

The default state for all ports on the Catalyst 4000/5000/6000 is auto silent mode; the command listed above is entered to be explicit.

Cisco Catalyst 2900XL/3500XL Configuration

The Catalyst 2900XL/3500XL does not support PAgP.

VLAN Trunking

A “trunk” in Cisco jargon is a single link (or port trunk) carrying traffic belonging to multiple VLANs. Cisco originally developed a proprietary method of VLAN trunking called ISL. IEEE 802.1Q, the industry standard method, was developed later. HP supports IEEE 802.1Q VLANs on all HP switches. Cisco supports both ISL and 802.1Q VLAN trunks, although lately, Cisco has been moving away from ISL.

Configuring VLAN trunks on Cisco switches is significantly different than the comparable task on HP switches. This is mainly due to Cisco hiding much of the complexity of configuring VLANs in their user interface. A Cisco network administrator never hears the terms tagged ports or untagged ports when configuring VLANs. This difference is what causes most problems when Cisco and HP switches interoperate using VLANs.

Configuring VLAN trunks on Cisco Switches

When a port is defined to be a VLAN trunk, regardless of whether it is defined as an ISL VLAN trunk or an 802.1Q VLAN trunk, the Cisco OS automatically configures the port to be a member of every VLAN defined on the switch. That is, traffic from every VLAN will flow across the link.

When traffic from multiple VLANs flow across the link, the switch marks the VLAN membership of each frame before transmitting it, so the neighboring switch knows to which VLAN the frame belongs. The method of marking the frame is one of the primary differences between ISL and 802.1Q.

802.1Q Tagging

We will not cover ISL encapsulation in this paper, but in 802.1Q VLAN trunks, a 4-byte tag is inserted in each frame just after the Layer 2 header. 12 bits of this 4-byte tag are used to identify the VLAN to which a frame belongs. When a frame does or does not carry the 802.1Q tag is confusing.

To illustrate this concept, let’s trace a frame from its source host (Host X) to its destination host (Host Y). Assume for a moment that both the source and destination hosts are on the same VLAN (VLAN 100), but are connected to different switches.

1. The “normal” Ethernet frame is transmitted from Host X and is received by Switch A on Port 1. Port 1 on Switch A is configured to belong to VLAN 100 and only traffic from VLAN 100 hosts is expected on this port. The port is configured to be a member of VLAN 100 in the untagged state. This means that the switch expects to receive untagged frames on this port and these frames belong to VLAN 100.
2. Switch A looks in its forwarding table and determines, based on the Layer 2 destination address, that the frame must be forwarded across the link to Switch B. Switch B is connected to Port 2 on Switch A. Port 2 is a VLAN trunk and is configured to be a member of all VLANs in the tagged state. In order for a port to carry traffic from multiple VLANs it must be in the tagged state. Before Switch A transmits the frame to Switch B, it inserts the 802.1Q tag showing that the frame belongs to VLAN 100. The tagged frame is then transmitted to Switch B.

3. Switch B receives the frame and determines it belongs to VLAN 100. Switch B does a lookup of the Layer 2 destination address and determines that Host Y is on Port 12. Port 12 is a member only of VLAN 100 and is untagged. Switch B then removes the 802.1Q tag from the frame restoring it to its original state and transmits the frame to Host Y.

In general, a frame carries the 802.1Q tag when it is required to distinguish it from frames belonging to other VLANs on a common link. According to the IEEE 802.1Q specification, a specific link can carry the traffic from one untagged VLAN and many tagged VLANs. As long as only one VLAN is untagged, the switches can clearly identify the traffic for all VLANs.

Native VLAN

The concept of a native VLAN for a VLAN trunk is an important one when working with Cisco switches. Before being defined to be a trunk port, a port belongs to a VLAN (even if it is the Default VLAN). When the port is defined to be a VLAN trunk, that VLAN to which the port belongs is the native VLAN for that VLAN trunk. When a Cisco switch is first initialized, one VLAN exists (VLAN 1, the default VLAN) and every port is a member of it. This means that unless the port is assigned to another VLAN, the native VLAN for a VLAN trunk will be the Default VLAN (VLAN 1).

When a VLAN trunk is created on a Cisco switch, the native VLAN for the trunk is *untagged*. All other VLANs are tagged. The Cisco network administrator has no flexibility on this issue.

It is critical that the configuration of the HP switch exactly match the Cisco's. Before configuring the HP switch, the network administrator should determine the native VLAN for the trunk (typically VLAN 1) on the Cisco switch. When configuring the VLAN trunk port on the HP switch, the native VLAN for that trunk should be configured as untagged and all other VLANs should be configured as tagged.

VLAN Trunking Examples

HP Switch Configuration

In this example, a VLAN trunk is created between an HP and a Cisco switch. The VLAN trunk is on port A1 and contains VLAN 1, VLAN 10, VLAN 20, and VLAN 30. Notice that VLAN 1 is untagged and the other VLANs are tagged in the HP switch configuration file. Note that some extraneous commands have been edited from the example configuration files for the sake of brevity. The critical commands are highlighted in red.

```
; J4865A Configuration Editor
; Created on release #G.04.XX

hostname "HP ProCurve Switch 4108GL"

time daylight-time-rule none

cdp run

ip default-gateway 10.1.1.1

snmp-server community "public" Unrestricted

vlan 1
 name "DEFAULT_VLAN"
 untagged A1-A3,B1-B24,C1-C24,D1-D24
 ip address 10.1.1.11 255.255.255.0
 exit

vlan 10
 name "VLAN10"
 untagged B1-B24
 tagged A1
 exit

vlan 20
 name "VLAN20"
 untagged C1-C24
 tagged A1
 exit

vlan 30
 name "VLAN30"
 untagged D1-D24
 tagged A1
 exit
```

Cisco Catalyst 4000/5000/6000 Configuration

The following is the Cisco Catalyst 4000/5000/6000 CatOS configuration file. The commands critical to creating the VLAN trunk are highlighted in red.

```
#vtp
set vtp domain pctests
set vlan 1 name default type ethernet mtu 1500 said 100001 state active
set vlan 10,20,30
stemaxhop 7 backupcrf off
!
#ip
set interface sc0 1 10.1.1.11/255.255.255.0 10.1.1.255
!
# default port status is enable
!
#module 1 : 2-port 1000BaseX Supervisor
!
#module 2 : 4-port Multilayer Switch
!
#module 3 : 8-port 1000BaseX Ethernet
!
#module 4 : 48-port 10/100BaseTX Ethernet
set trunk 4/1  nonegotiate dot1q 1-1005,1025-4094
!
```

Cisco Catalyst 2900XL/3500XL Configuration

The following is a Cisco Catalyst 2900XL/3500XL configuration file. The commands critical to creating the VLAN trunk are highlighted in red.

```
!
ip subnet-zero
!
interface FastEthernet0/1
 switchport trunk encapsulation dot1q
 switchport mode trunk
 switchport trunk native vlan 1
!
interface FastEthernet0/2
!
interface FastEthernet0/3   (same to 24)
!
interface FastEthernet0/24
!
interface GigabitEthernet0/1
!
interface GigabitEthernet0/2
!
interface VLAN1
 ip address 10.1.1.15 255.255.255.0
 no ip directed-broadcast
 no ip route-cache
!
interface VLAN10
 no ip directed-broadcast
 no ip route-cache
!
interface VLAN20
 no ip directed-broadcast
 no ip route-cache
!
interface VLAN30
 no ip directed-broadcast
 no ip route-cache
!
ip default-gateway 10.1.1.1
!
line con 0
 transport input none
 stopbits 1
line vty 5 15
```

Cisco IOS Routers

IEEE 802.1Q VLAN trunking is supported Cisco IOS Routers only if they have 10/100BaseT interfaces (for example, the 2600, 3000, 4000, 7000, and the MSM). The following commands are extracted from a 2621 configuration file. Critical commands are

```
interface FastEthernet0/1
 no ip address
 duplex auto
 speed auto
!
interface FastEthernet0/1.1
 encapsulation dot1Q 1 native
 ip address 10.1.1.15 255.255.255.0
!
interface FastEthernet0/1.2
 encapsulation dot1Q 10
 ip address 10.1.10.1 255.255.255.0
!
interface FastEthernet0/1.3
 encapsulation dot1Q 20
ip address 10.1.20.1 255.255.255.0
!
interface FastEthernet0/1.4
 encapsulation dot1Q 30
 ip address 10.1.30.1 255.255.255.0
```

in red.

Link Aggregation and VLAN Trunking

HP Switch Configuration

The last example combines the first two examples. In this example, we create a port trunk across 4 ports and then use the port trunk as a VLAN trunk.

The following is the HP switch configuration file. The critical commands are highlighted in red.

```
interface A1
 no lacp
 exit
interface A2
 no lacp
 exit
interface A3
 no lacp
 exit
interface A4
 no lacp
 exit

trunk A1-A4 Trk1 Trunk

ip default-gateway 10.1.1.1

snmp-server community "public" Unrestricted

vlan 1
 name "DEFAULT_VLAN"
 untagged A5-A12,A17-A24,C1-C24,D1-D24,E1-E24,G1-G24,Trk1
 no untagged A13-A16
 ip address 10.1.1.12 255.255.255.0
 exit

vlan 10
 tagged Trk1
 exit

vlan 20
 tagged Trk1
 exit

vlan 30
tagged Trk1
 exit
```

Cisco Catalyst 4000/5000/6000 Configuration

The following is the Cisco Catalyst 4000/5000/6000 CatOS configuration file. The commands critical to creating the VLAN trunk are highlighted in red.

```
#vtp
set vtp domain pctests
set vlan 1 name default type ethernet mtu 1500 said 100001 state active
set vlan 10,20,30
!
#ip
set interface sc0 1 10.1.1.11/255.255.255.0 10.1.1.255

!
#port channel
set port channel 4/1-4 21
!
# default port status is enable
!
!
#module 1 : 2-port 1000BaseX Supervisor
!
#module 2 : 4-port Multilayer Switch
!
#module 3 : 8-port 1000BaseX Ethernet
!
#module 4 : 48-port 10/100BaseTX Ethernet
set trunk 4/1  nonegotiate dot1q 1-1005,1025-4094
set trunk 4/2  nonegotiate dot1q 1-1005,1025-4094
set trunk 4/3  nonegotiate dot1q 1-1005,1025-4094
set trunk 4/4  nonegotiate dot1q 1-1005,1025-4094
set port channel 4/1-4 mode on silent
!
#module 5 empty
!
#module 6 empty
!
#module 15 empty
!
#module 16 empty
end
```

Cisco Catalyst 2900XL/3500XL Configuration

The following is a Cisco Catalyst 2900XL/3500XL configuration file. The commands critical to creating the VLAN trunk are highlighted in red.

```
interface FastEthernet0/1
 port group 1 distribution destination
 switchport trunk encapsulation dot1q
 switchport mode trunk
!
interface FastEthernet0/2
 port group 1 distribution destination
 switchport trunk encapsulation dot1q
 switchport mode trunk
!
interface FastEthernet0/3
 port group 1 distribution destination
 switchport trunk encapsulation dot1q
 switchport mode trunk
!
interface FastEthernet0/4
 port group 1 distribution destination
 switchport trunk encapsulation dot1q
 switchport mode trunk
!
interface FastEthernet0/5  (same to 24)
!
interface FastEthernet0/24
!
interface GigabitEthernet0/1
!
interface GigabitEthernet0/2
!
interface VLAN1
 ip address 10.1.1.15 255.255.255.0
 no ip directed-broadcast
 no ip route-cache
!
interface VLAN10
 no ip directed-broadcast
 no ip route-cache
!
interface VLAN20
 no ip directed-broadcast
 no ip route-cache
!
interface VLAN30
 no ip directed-broadcast
 no ip route-cache!
!
ip default-gateway 10.1.1.1
!
```

Cisco IOS Router Configuration

Cisco IOS routers such as the 7000 or MSM modules for the Catalyst support the combination of link aggregation and VLAN trunking. The following is an excerpt from the configuration file of the MSM module in a Catalyst 6500. The critical commands are in red.

```
interface Port-channel1
 no ip address
 no ip directed-broadcast
!
interface Port-channel1.1
 encapsulation dot1q 1
 ip address 10.1.1.15  255.255.255.0
!
interface Port-channel1.10
 encapsulation dot1q 20
 ip address 10.1.10.1  255.255.255.0
no ip directed-broadcast
!
interface Port-channel1.20
 encapsulation dot1q 10
 ip address 10.1.20.1  255.255.255.0
no ip directed-broadcast
!
interface Port-channel1.30
 encapsulation dot1q 30
 ip address 10.1.30.1 255.255.255.0
no ip directed-broadcast
!
!
interface GigabitEthernet0/0/0
 no ip address
 no ip directed-broadcast
 no negotiation auto
 channel-group 1
!
interface GigabitEthernet1/0/0
 no ip address
 no ip directed-broadcast
 no negotiation auto
 channel-group 1
!
interface GigabitEthernet3/0/0
 no ip address
 no ip directed-broadcast
 no negotiation auto
 channel-group 1
!
interface GigabitEthernet4/0/0
 no ip address
 no ip directed-broadcast
 no negotiation auto
 channel-group 1
```