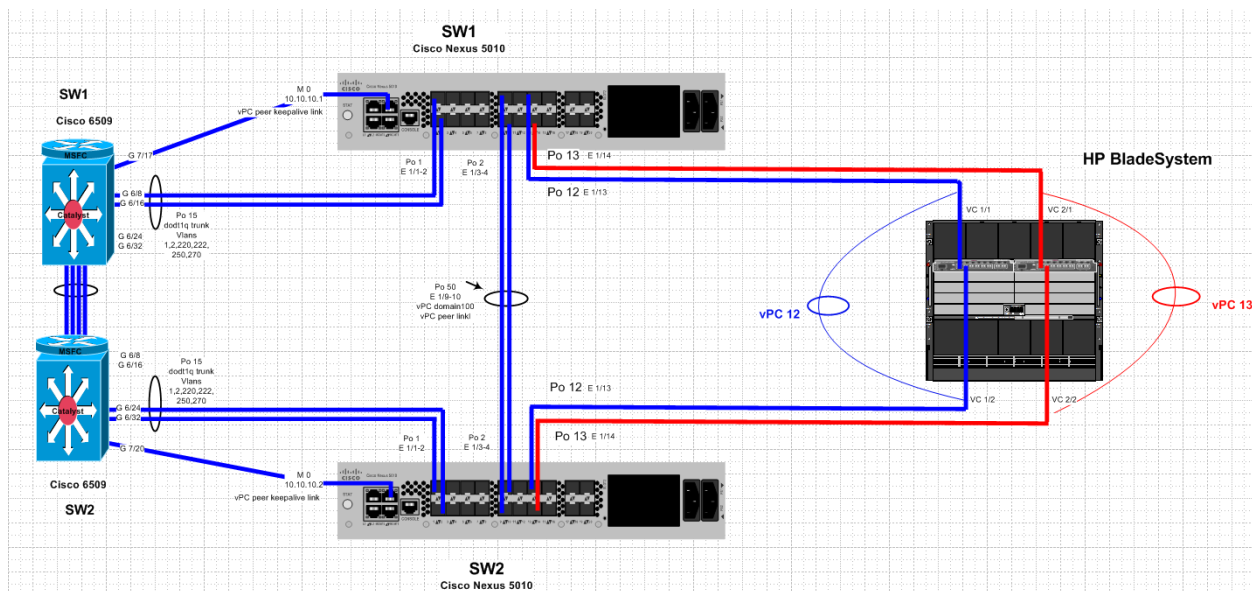


HP Virtual Connect Flex-10 and Nexus vPC (virtual PortChannel) Configuration

This paper will outline the steps to configure the Virtual Connect Flex-10 modules and Cisco Nexus 5000/7000 series switches as a virtual PortChannel. HP Virtual Connect Ethernet modules will work seamlessly with Cisco Nexus infrastructures and this new network design. A virtual PortChannel (vPC) provides some basic benefits.

- Allows a single device to use a PortChannel across two upstream devices
- Eliminates Spanning Tree Protocol blocked ports
- Provides a loop-free topology
- Uses all available uplink bandwidth

This diagram below shows the environment configuration. There are two Nexus 5010 switches, One HP Blade Enclosure with two Virtual Connect Flex-10 modules. The uplink port (blue) in VC 1/1 is connected to Nexus SW1, Port 1/13: and uplink VC 1/2 (blue) is connected to Nexus SW2 Port 1/13. These ports are configured as PortChannel Po12. The uplink port (red) VC 2/1 is connected to Nexus SW1 Port, 1/14 and uplink VC 2/2 (red), is connected to Nexus SW2 Port 1/14. These ports are configured as PortChannel Po13.



Setting up the Nexus switches

The following configuration outline details how to setup the virtual port channels for this configuration. Your case may vary but the steps are the same.

Here are the basic steps:

- Enabling lacp and vpc feature, on both switch (*notice: pagp is not supported on NX-OS*)
- Create a vPC domain and enter vpc-domain mode
- Configure vPC peer keepalive link
- Create vPC peer link
- Move port-channel to vPC
- Create Po10 and Po11
- Move it to vPC

Log into the nexus switch to configure Global Settings: LACP and VPC

```
switch# Config-t
```

```
switch (config)# feature lacp *This is now set at the global level and not at the individual port
```

```
switch (config)# feature vpc enable/start vpc (virtual port-channel) feature/process
```

```
switch (config)# show feature Verify the above features are enabled
```

Feature Name	Instance	State
-----	-----	-----
tacacs	1	disabled
lacp	1	enabled
interface-vlan	1	disabled
private-vlan	1	disabled
udld	1	disabled
vpc	1	enabled
fcoe	1	disabled
fex	1	enabled

Configuring the management interface (must be done on both switches)

```
switch# int mgmt0 goes to management interface
```

`(config-if)# ip address 10.10.10.1/24` *set ip address and subnet mask (it might be done already)*

`switch(config-if)# vrf context management` *goes to vrf management context for configuration*

`switch#(config-vrf)# ip route 0.0.0.0/0 10.10.10.1` *configure the virtual routing*

`switch (config-vrf)# exit` *back to previous context*

Creating a vPC domain and configuring the peer keepalive link

Establish VPC trunk between Nexus SW 1 and SW 2:

The Nexus OS version 4.1(3)N1(1) is installed on both N5K switches. The vpc feature was enabled on the two N5K switches with the keepalive address pointing to the peer Vlan 1 address: A vpc peer link was created as required, this case bonding to 2 x 10GE ports in a portchannel. Initial vpc link operation was verified with the 'show vpc' command at the N5K CLI (output shown later).

`switch(config)# vpc domain 100` *create a vpc domain*

`switch(config-vpc-domain)#peer-keepalive destination 10.10.10.2` *establish link to the 2nd switch mgmt interface (assuming it has already been configured with ip address)*

→Note:

-----: Management VRF will be used as the default VRF ::-----

Configuring the vPC peer link

The process as described below is:

1. Add interface eth 1/9 and 1/10 of both switch to port-channel Po50
2. Make Po50 a peer link

[on each N5K]

`switch (config-vpc-domain)# int eth 1/3-4` *goes to interfaces where the port-channel 50 will be set*

`switch (config-if)# channel-group 50 mode active`

1. Etherchannel naming is now replaced with Portchannel (Nexus OS)
2. Creating a channel-group X will automatically create a port-channel (PoX)

`switch (config-if)#int po50` *goes to port-channel interface po1*

`switch (config-if)#vpc peer-link` *and make it peer*

Please note that spanning tree port type is changed to "network" port type on vPC peer-link.

This will enable spanning tree Bridge Assurance on vPC peer-link provided the STP Bridge Assurance

(which is enabled by default) is not disabled.

(optional) switch (config-if)#switchport mode trunk like any ISL, the port should be turned to trunk mode (according Cisco)

(optional) switch (config-if)#switchport trunk allowed vlan 1-2,220, 222, 250, 270 list of allowed vlan

Show Interface: Output should look like this

```
interface port-channel50
  description ** Nexus interswitch trunk **
  switchport mode trunk
  switchport trunk allowed vlan 1-2,220, 222, 250, 270
  vpc peer-link
```

```
interface Ethernet1/9
  description ** interswitch trunk **
  switchport mode trunk
  switchport trunk allowed vlan 1-2,220, 222, 250, 270
  channel-group 50 mode active
```

```
interface Ethernet1/10
  description ** interswitch trunk **
  switchport mode trunk
  switchport trunk allowed vlan 1-2,220, 222, 250, 270
  channel-group 50 mode active
```

Each N5K switch has a single 10GE port connected to Virtual Connect Interconnect bay 1. An LACP portchannel was created on each switch (arbitrarily numbered 12) and a vpc identifier was associated with that portchannel using the same number 12. Note that an identical number for the portchannel and the vpc identifier is **not** required. What is required is that the vpc identifier be identical for the peered portchannels on the respective switches. The Virtual Connect-connected port was then added to the channel on each switch: [A second identical vpc was created with the 2nd Virtual Connect Module link;

that vpc is numbered 13]. Portchannel creation and channel member additions were configured as follows on Nexus SW 1:

Configuring Port-channel linked to Virtual-Connect

Configuration needs to be done on both switches

```
interface port-channel12
  description ** VC to HPBS Bay 1 **
  switchport mode trunk
  switchport trunk allowed vlan 1-2,220, 222, 250, 270
  vpc 12
spanning-tree port type edge trunk
```

```
interface port-channel13
  description ** VC to HPBS Bay 2 **
  switchport mode trunk
  switchport trunk allowed vlan 1-2,220, 222, 250, 270
  vpc 13
spanning-tree port type edge trunk
```

```
interface Ethernet1/13
  description ** Po12 member to HPBS **
  switchport mode trunk
  switchport trunk allowed vlan 1-2,220, 222, 250, 270
  channel-group 12 mode active
```

```
interface Ethernet1/14
  description ** Po13 member to HPBS **
  switchport mode trunk
  switchport trunk allowed vlan 1-2,220, 222, 250, 270
```

channel-group 13 mode active

The last configuration command on the portchannel interface (**bold**) is critical. It is used to enable the PortFast feature which removes the portchannel (and all of the channel members) from the STP topology. It must be attached to the portchannel, not the member interfaces. Otherwise spanning-tree will block the channel and thus all member ports and no traffic will be forwarded from the channel to the Blade System.

Switch Setup Verification Steps:

Operational verification from the N5K side is first seen with the 'show vpc' command"

ANC-DCNX-SW1# sho vpc

Legend:

(*) - local vPC is down, forwarding via vPC peer-link

vPC domain id : 100

Peer status : peer adjacency formed ok

vPC keep-alive status : peer is alive

Configuration consistency status: success

vPC role : primary

vPC Peer-link status

id Port Status Active vlans

-- --- -----

1 Po50 up 1-2,220, 222, 250, 270

vPC status

id	Port	Status	Consistency	Reason	Active vlans
----	------	--------	-------------	--------	--------------

12	Po12	up	success	success	1-2,220, 222, 250, 270
----	------	----	---------	---------	------------------------

13	Po13	up	success	success	1-2,220, 222, 250, 270
----	------	----	---------	---------	------------------------

The 'show vpc' output is *NOT* enough to verify transport. vPC's may be up and pass a successful consistency check, and link-layer criteria for an up-up interface state may be satisfied. However, spanning-tree may still put the portchannel in a blocked state. This was encountered when the mistake was made placing the PortFast enabling command (spanning-tree port type edge trunk) on the channel port member interfaces.

To verify forwarding state, run the 'show spanning-tree' command. As you can see below, Port Channels 12 and 13 are in a forwarding state with additional information indicating vPC configuration in the type field.

```
ANC-DCNX-SW1# show spanning-tree
```

```
VLAN0001
```

```
Spanning tree enabled protocol rstp
```

```
Root ID Priority 100
```

```
Address 0002.0000.bc01
```

```
Cost 3
```

```
Port 4096 (port-channel1)
```

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

```
Bridge ID Priority 32769 (priority 32768 sys-id-ext 1)
```

```
Address 000d.0000.c53c
```

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

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

```
-----
```

```
Po1 Root FWD 3 128.4096 Network P2p
```

```
Po12 Desg FWD 1 128.4107 (vPC) Edge P2p
```

```
Po13 Desg FWD 1 128.4108 (vPC) Edge P2p
```

```
Po50 Desg FWD 1 128.4145 (vPC peer-link) Network P2p
```

```
VLAN0002
```

```
Spanning tree enabled protocol rstp
```

```
Root ID Priority 100
```

```
Address 0002.0000.bc02
```

Cost 3

Port 4096 (port-channel1)

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

Bridge ID Priority 32770 (priority 32768 sys-id-ext 2)

Address 000d.0000.c53c

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

Interface Role Sts Cost Prio.Nbr Type

Po1 Root FWD 3 128.4096 Network P2p

Po12 Desg FWD 1 128.4107 (vPC) Edge P2p

Po13 Desg FWD 1 128.4108 (vPC) Edge P2p

Po50 Desg FWD 1 128.4145 (vPC peer-link) Network P2p

VLAN0220

Spanning tree enabled protocol rstp

Root ID Priority 100

Address 0002.0000.bc02

Cost 3

Port 4096 (port-channel1)

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

Bridge ID Priority 32770 (priority 32768 sys-id-ext 2)

Address 000d.0000.c53c

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

Interface Role Sts Cost Prio.Nbr Type

Po1 Root FWD 3 128.4096 Network P2p

Po12 Desg FWD 1 128.4107 (vPC) Edge P2p

Po13 Desg FWD 1 128.4108 (vPC) Edge P2p

Po50 Desg FWD 1 128.4145 (vPC peer-link) Network P2p

VLAN0222

Spanning tree enabled protocol rstp

Root ID Priority 100

 Address 0002.0000.bcde

 Cost 3

 Port 4096 (port-channel1)

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

Bridge ID Priority 32990 (priority 32768 sys-id-ext 222)

 Address 000d.0000.c53c

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

Interface Role Sts Cost Prio.Nbr Type

Po1 Root FWD 3 128.4096 Network P2p

Po12 Desg FWD 1 128.4107 (vPC) Edge P2p

Po13 Desg FWD 1 128.4108 (vPC) Edge P2p

Po50 Desg FWD 1 128.4145 (vPC peer-link) Network P2p

VLAN0250

Spanning tree enabled protocol rstp

Root ID Priority 100

 Address 0002.0000.bc02

 Cost 3

 Port 4096 (port-channel1)

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

Bridge ID Priority 32770 (priority 32768 sys-id-ext 2)

 Address 000d.0000.c53c

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

Interface Role Sts Cost Prio.Nbr Type

```

-----
Po1      Root FWD 3      128.4096 Network P2p
Po12    Desg FWD 1      128.4107 (vPC) Edge P2p
Po13    Desg FWD 1      128.4108 (vPC) Edge P2p
Po50     Desg FWD 1      128.4145 (vPC peer-link) Network P2p

```

VLAN0270

Spanning tree enabled protocol rstp

Root ID Priority 100

Address 0002.0000.bc02

Cost 3

Port 4096 (port-channel1)

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

Bridge ID Priority 32770 (priority 32768 sys-id-ext 2)

Address 000d.0000.c53c

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

```

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

```

```

Po1      Root FWD 3      128.4096 Network P2p
Po12    Desg FWD 1      128.4107 (vPC) Edge P2p
Po13    Desg FWD 1      128.4108 (vPC) Edge P2p
Po50     Desg FWD 1      128.4145 (vPC peer-link) Network P2p

```

Lastly, once the configuration is in port forwarding state, MAC addresses of physical and virtual hosts will be resident in the L2 cache (note the HP IEEE OUI hex prefixes). This test can also be indirectly verified from the host servers (properly configured at layer 2 and 3) successfully receiving an ICMP echo reply from their gateways or other valid network address.

The 'show mac-address-table' gives us what we are looking for at layer 2.

ANC-DCNX-SW1# show mac-address-table

VLAN	MAC Address	Type	Age	Port
1	0000.0c07.ac01	dynamic	0	Po1
1	0002.fc49.bc0a	dynamic	0	Po1
1	0005.7446.2443	dynamic	10	Po1
1	0005.7446.244b	dynamic	20	Po1
1	0023.7d43.3c2b	dynamic	30	Po12
1	0023.7d43.3c3b	dynamic	20	Po13
2	0000.0c07.ac02	dynamic	0	Po1
2	0002.fc49.bc0a	dynamic	0	Po1
2	0005.7446.2443	dynamic	10	Po1
2	0017.a477.009a	dynamic	0	Po12
2	0017.a477.009e	dynamic	0	Po13
222	0000.0c07.ac01	dynamic	0	Po1
222	0002.fc49.bc0a	dynamic	0	Po1
222	0005.7446.2443	dynamic	10	Po1
222	0017.a477.009c	dynamic	0	Po12
222	0017.a477.00a0	dynamic	0	Po13

Total MAC Addresses: 16

Virtual Connect Network Setup (VCM):

The virtual connect environment was setup with an active/active configuration so that all uplink ports from the VC interconnect bays would be active links to the Nexus switches.

VCM Configuration Outline:

Create two Virtual Connect Shared Uplink Sets

Server_VLANs_A Select Uplink ports from VC Bay 1: Port X1, Port X2

Server_VLANs_B Select Uplink ports from VC Bay 2: Port X1, Port X2

Create VC Networks mapping to each VLAN ID

Select Shared Uplink Set – Server_VLANs_A, at the bottom input the VC Network names and VLAN id

VLAN2_A vlan ID 2

VLAN220_A vlan ID 220

VLAN222_A vlan ID 222

VLAN250_A vlan ID 250

VLAN270_A vlan ID 270

Repeat steps for SUS Server_VLANS_B, making sure to name the VC network with a suffix _B with the same vlan ID.

The picture below shows the VC network with all physical links once the VPC is setup.

Each virtual connect link for the configured SUS was placed in an Active state after the Nexus switches were properly configured

Shared Uplink Set	Port Status	Connector Type	Uplink Port
OK Server_VLANS_A	OK Linked/Active	10 Gb SFP-DAC	OK Blade_Srv2: Bay 1: Port X1
	OK Linked/Active	10 Gb SFP-DAC	OK Blade_Srv2: Bay 1: Port X2
OK Server_VLANS_B	OK Linked/Active	10 Gb SFP-DAC	OK Blade_Srv2: Bay 2: Port X1
	OK Linked/Active	10 Gb SFP-DAC	OK Blade_Srv2: Bay 2: Port X2

Ethernet Networks	Shared Uplink Set (VLAN ID)	Port Status	Connector Type	Connected To	Uplink Port
Failed DMZ_FTP		Unavailable Not Linked 0 Mb	SFP-RJ45		Unavailable Blade_Srv2: Bay 1: Port X6
		Unavailable Not Linked 0 Mb	SFP-RJ45		Unavailable Blade_Srv2: Bay 2: Port X6
Failed DMZ_Web		Unavailable Not Linked 0 Mb	SFP-RJ45		Unavailable Blade_Srv2: Bay 1: Port X5
		Unavailable Not Linked 0 Mb	SFP-RJ45		Unavailable Blade_Srv2: Bay 2: Port X5
OK VLAN2_A	Server_VLANS_A(2)	OK Linked/Active 10 Gb	SFP-DAC	00:0d:ec:d4:c5:00 (00:0d:ec:d4:c5:14)	OK Blade_Srv2: Bay 1: Port X1
		OK Linked/Active 10 Gb	SFP-DAC	00:0d:ec:d4:e8:40 (00:0d:ec:d4:e8:54)	OK Blade_Srv2: Bay 1: Port X2
OK VLAN2_B	Server_VLANS_B(2)	OK Linked/Active 10 Gb	SFP-DAC	00:0d:ec:d4:c5:00 (00:0d:ec:d4:c5:15)	OK Blade_Srv2: Bay 2: Port X1
		OK Linked/Active 10 Gb	SFP-DAC	00:0d:ec:d4:e8:40 (00:0d:ec:d4:e8:55)	OK Blade_Srv2: Bay 2: Port X2
OK VLAN222_A	Server_VLANS_A (222)	OK Linked/Active 10 Gb	SFP-DAC	00:0d:ec:d4:c5:00 (00:0d:ec:d4:c5:14)	OK Blade_Srv2: Bay 1: Port X1
		OK Linked/Active 10 Gb	SFP-DAC	00:0d:ec:d4:e8:40 (00:0d:ec:d4:e8:54)	OK Blade_Srv2: Bay 1: Port X2
OK VLAN222_B	Server_VLANS_B (222)	OK Linked/Active 10 Gb	SFP-DAC	00:0d:ec:d4:c5:00 (00:0d:ec:d4:c5:15)	OK Blade_Srv2: Bay 2: Port X1
		OK Linked/Active 10 Gb	SFP-DAC	00:0d:ec:d4:e8:40 (00:0d:ec:d4:e8:55)	OK Blade_Srv2: Bay 2: Port X2
OK VMotion					

Example configuration of VCM link state if all VC uplink ports are in a single Virtual Connect Network.

Link state changes to Active/Passive between physical VC modules. Bay 1, ports X1, X2 have been placed in a standby state, while Bay 2 ports are active.

This behavior is by design; VCs loop prevention mechanism has placed VC Bay 1 into standby to prevent any loops in the topology. Each VC module has by default 2 cross connect ports X7 and X8, when present the adjacent VC modules are stacked together. Without this loop prevention behavior a network loop would be formed.

Port	Port Role	Port Status	Connector Type	Connected To	PID	Speed Duplex	Delete
AZ_1: Bay 1: Port X1	NA	OK Linked/Standby	10 Gb SFP-DAC	00:0d:ec:d5:aa:80 (00:0d:ec:d5:aa:88)	Auto	X	
AZ_1: Bay 1: Port X2	NA	OK Linked/Standby	10 Gb SFP-DAC	00:0d:ec:d5:aa:00 (00:0d:ec:d5:aa:08)	Auto	X	
AZ_1: Bay 2: Port X1	NA	OK Linked/Active	10 Gb SFP-DAC	00:0d:ec:d5:aa:80 (00:0d:ec:d5:aa:89)	Auto	X	
AZ_1: Bay 2: Port X2	NA	OK Linked/Active	10 Gb SFP-DAC	00:0d:ec:d5:aa:00 (00:0d:ec:d5:aa:09)	Auto	X	

Summary:

This paper has outlined the basic setup steps to create a virtual port channel using the new Nexus switches and Virtual Connect Flex-10. With this design you can create a highly available network infrastructure resulting in switch and path redundancy.

Source for configuration guide:

http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/configuration_guide_c07-543563.html

Sources for IOS vs NX-OS comparisons:

http://docwiki.cisco.com/wiki/Cisco_NX-OS/IOS_Configuration_Fundamentals_Comparison