Cisco Nexus 4000

Design Guide

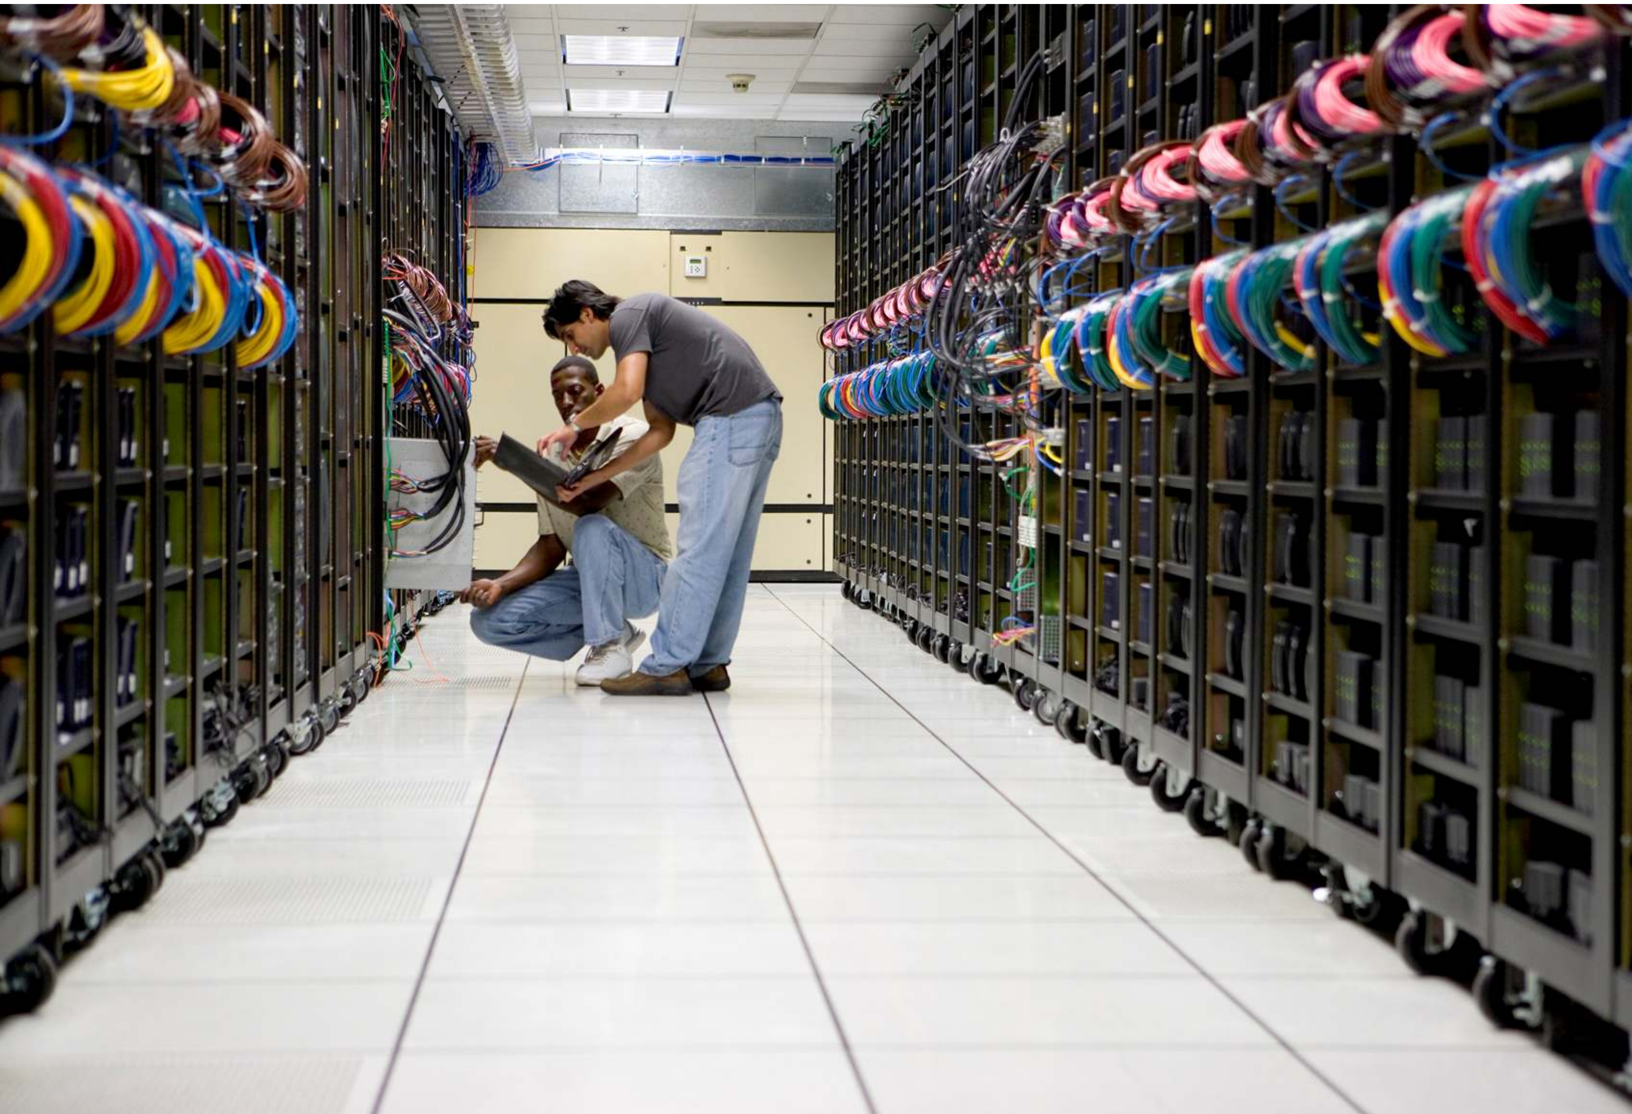

Contents

Introduction

This white paper is targeted at server, storage, and network administrators who plan to deploy IBM BladeCenter servers with the unified fabric solution of the Cisco Nexus® 4000 Series Switches. The Cisco® Nexus 4000 Series is Cisco's next-generation blade switch made available for IBM's blade chassis. It provides 20 full line rate interfaces at 1G or 10G, of which 14 are for the blade servers and 6 are for uplink connectivity. Building a strong Layer 2 Ethernet infrastructure, the Cisco Nexus 4000 also provides the capability to allow Fibre Channel over Ethernet (FCoE) traffic to traverse this network.

The Cisco Nexus 4000 Series Switches provide the Fibre Channel Forwarder (FCF) function that allows FCoE devices to log in to the fabric and forwards Fibre Channel (FC) frames to appropriate devices over an Ethernet infrastructure. The Cisco Nexus 4000 Series is the first to unify Ethernet and Fibre Channel into a single integrated blade switch. With the addition of the FCF, it is the next step in the evolution to allow FCoE traffic to be forwarded to a FCF in a multihop topology. This is known as a FCoE Initialization Protocol (FIP) snooping bridge.

This design guide provides guidelines and best practices in implementing the different types of deployments that are possible with the Cisco Nexus 4000 Series. Various topologies can be implemented with the Cisco Nexus 4000, and this design guide describes the factors to consider with respect to these various topologies and how best to implement them.

Cisco Nexus 4000 Series in an IBM Blade Chassis

The Cisco Nexus 4000 Series blade switch is supported only by models BCH and BCH-T of IBM's blade chassis. The blade switch can be inserted only into the high-speed slots (HSS) in the back of the chassis, which are slots 7, 8, 9, and 10. Figure 1 depicts what the back of the IBM BCH blade chassis looks like and where the Cisco Nexus 4000 can be inserted.

Figure 1. IBM BCH Chassis with Cisco Nexus 4000

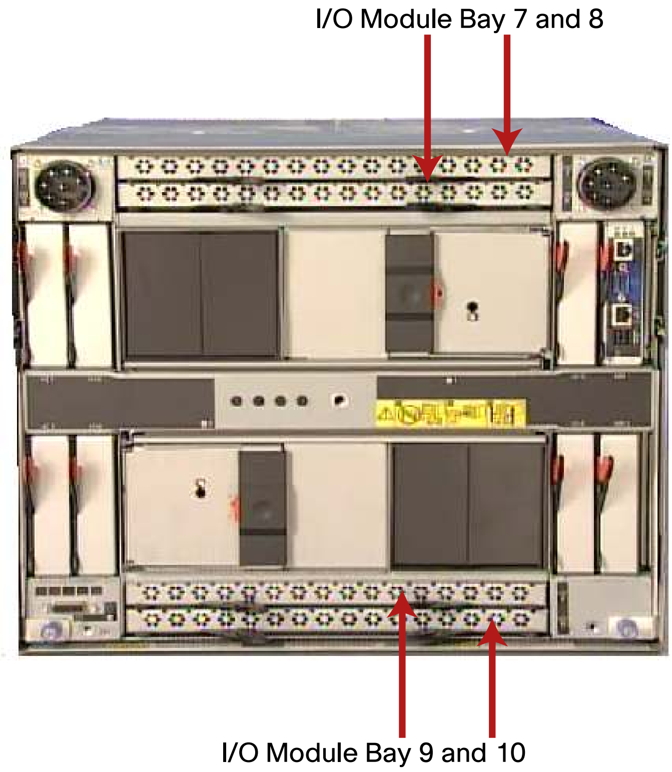

The blade server internal LAN on Motherboard (LoM) and additional mezzanine cards are statically pinned to the I/O module bays in the back of the chassis. The LoM ports are statically pinned to I/O bays 1 and 2. There are two available mezzanine slots in the blade servers, and the mezzanine cards that are supported for the Cisco Nexus 4000 have to be of type “CFFH.” The cards can be inserted only into the HSS slots inside the blade server. The supported list of mezzanine cards includes the following vendors:

- Broadcom
    - Either dual-port or quad-port
    - Supports only Ethernet (no FCoE support)
- QLogic
    - Supports only dual ports
    - Supports both Ethernet and FCoE traffic

Note: As of January 2010, Emulex is currently only qualified as an Enternet NIC. IBM is in the process of qualifying Emulex Converged Network Adapters (CNAs) for FCoE.

Blade Server to Cisco Nexus 4000 Mapping

With the blade server connectivity ports statically mapped to the I/O modules, it is critical to know how they are mapped. The mezzanine cards fit only in the HSS slots of the blade servers that are then mapped to the HSS I/O bays in the back of the chassis. So depending on which mezzanine card is inserted, the mapping of those ports is as follows:

- Dual-port cards (either Broadcom or QLogic)
    - Port 1 is mapped to I/O bay 7
    - Port 2 is mapped to I/O bay 9

Note: The Cisco Nexus 4000 blade switch will be inserted in bays 7 and 9.

- Quad-port cards (currently only Broadcom)
    - Port 1 is mapped to I/O bay 7
    - Port 2 is mapped to I/O bay 9
    - Port 3 is mapped to I/O bay 8
    - Port 4 is mapped to I/O bay 10

Figure 2 depicts the mapping for the blade server’s ports to I/O bay modules.

Figure 2. Blade Server Mappings

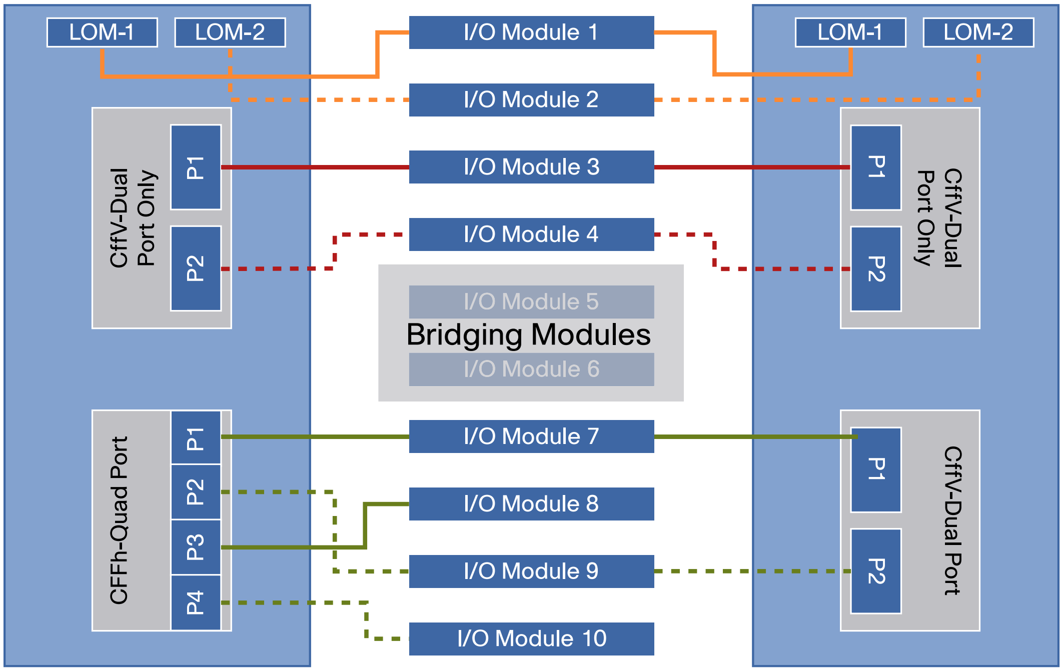

Note: I/O modules 5 and 6 are bridging modules, and blade server I/O ports are not mapped to those I/O bays.

Correlation of Blade Server to Cisco Nexus 4000 Interfaces

After understanding the internal blade server ports and how they are mapped to the blade switches in the chassis, the next correlation needed is the port mapping within the Cisco Nexus 4000 to the blade servers. Since there are 14 server-facing ports and 6 uplinks on the Nexus 4000, the server-facing ports are mapped as shown in Figure 3.

Figure 3. Cisco Nexus 4000 Port Mapping to IBM Blade Servers

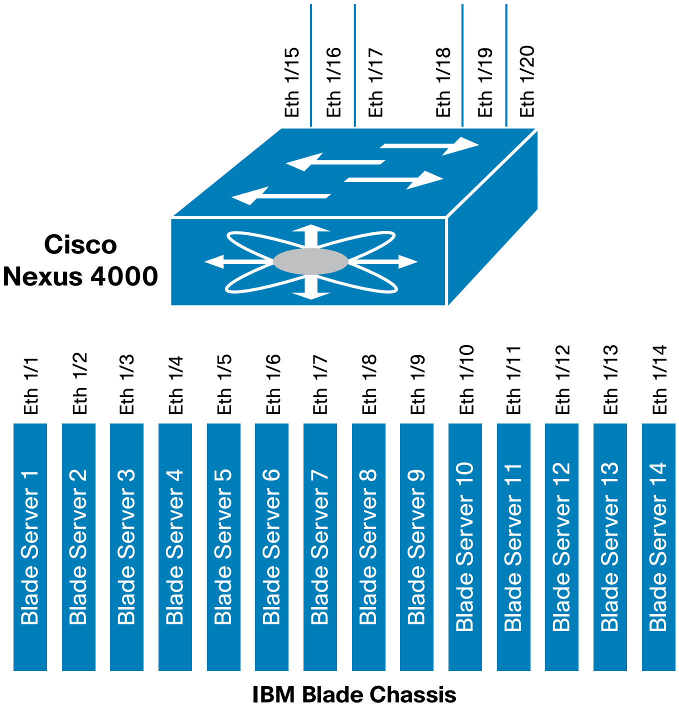

Note: The numbering of the uplink interfaces goes from left to right in the blade switch and starts at interface number Ethernet 1/15.

Cisco Nexus 4000 in an Ethernet Only Solution

The Cisco Nexus 4000 Series can be deployed as a typical Ethernet blade switch as the access layer. Running FCoE on the Cisco Nexus 4000 is not required but is an added advantage in unifying the I/O module within the blade chassis, which reduces the number of additional I/O modules needed if FC traffic is required. With the configuration of the Cisco Nexus 4000 in the access layer, the deployment will currently fall into one of these two categories:

1. Classic spanning-tree deployment with the well-known V-shaped or U-shaped or inverted U topology
2. Deployment with port channels dual-connected to the upstream Cisco Nexus 7000 (running virtual PortChannel [vPC]), Nexus 5000 (running vPC), or Catalyst® 6500 (running Virtual Switching System [VSS])

Note: The Cisco Nexus 4000 release with version code 4.1(2)E1(1) does not natively support vPC, meaning that configuration of vPC to the downlinks (server-facing links) is not supported.

We will focus on deployment number 2 in this design guide, which will be the best-practice deployment with the Cisco Nexus 4000. This does not preclude building designs based on deployment option 1, which is also supported. Figure 4 shows the supported topologies for an Ethernet-only environment.

Figure 4. Other Supported Ethernet-Only Topologies

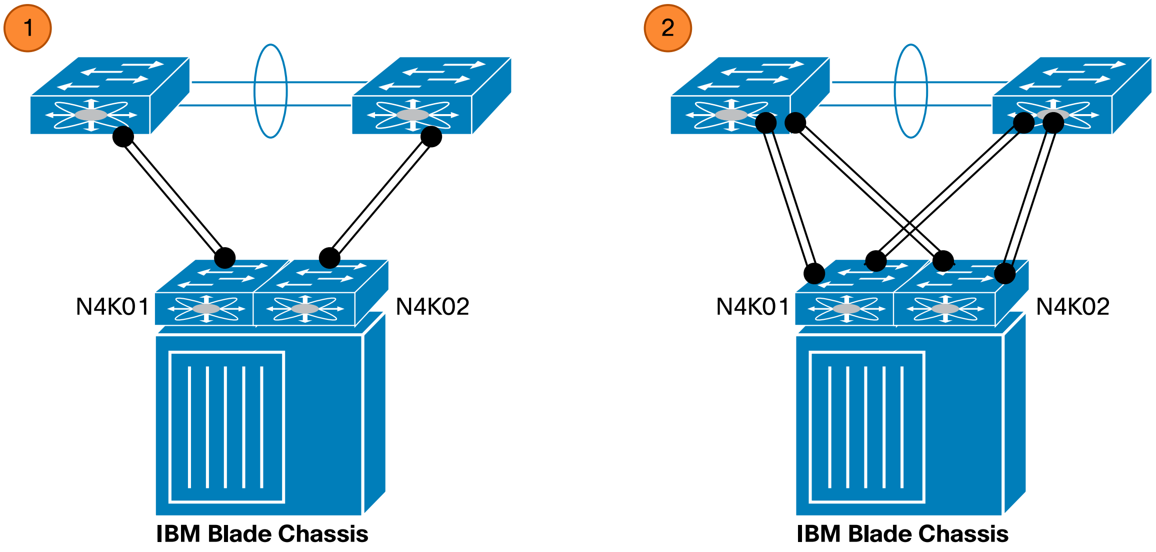

Note: In topology 1, there are no spanning-tree issues, since they are connected to only one upstream switch. As a best practice, it is still recommended that you not make the Cisco Nexus 4000 the root bridge. In topology 2, the use of spanning tree will be blocked on the secondary path. The diagram shows multiple links, but these are not required.

Preferred Connectivity Between the Aggregation Layer and the Cisco Nexus 4000

With both the Cisco Nexus 5000 and Nexus 7000 supporting vPC, the Nexus 4000 may uplink to either of these aggregation switches and gains the benefits of a vPC environment. Figure 5 depicts what the physical topology would look like.

Figure 5. Cisco Nexus 4000 Connectivity to the vPC Aggregation Layer

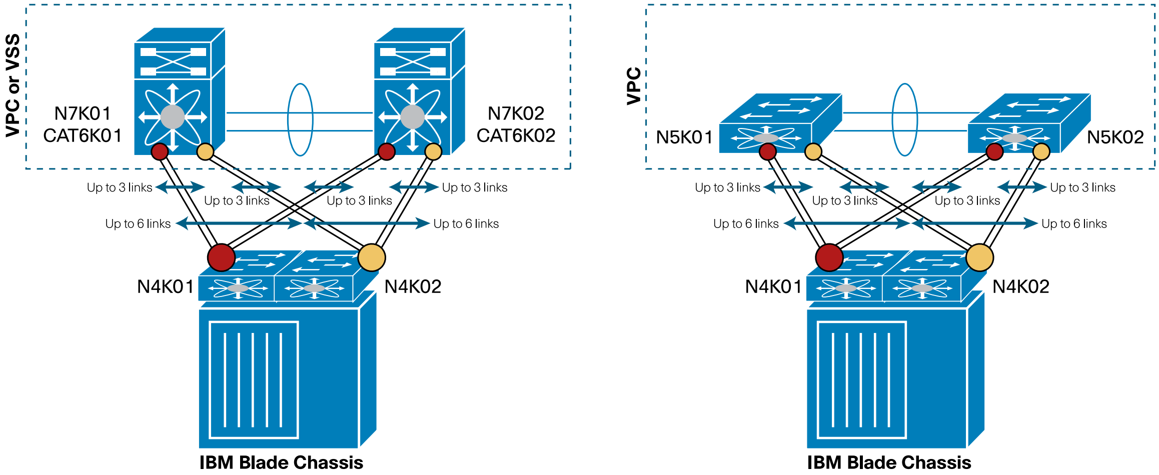

Note: The benefit of connecting the Nexus 4000 to a vPC-capable aggregation layer is that doing so allows for all the links to be forwarding and eliminates the need for a spanning tree.

Note: Connecting the Cisco Nexus 4000 to the currently available Nexus 2000 (model 2148) is not supported or recommended, as the Cisco Nexus 2148 has only copper Gigabit Ethernet connectivity and Bridge Protocol Data Unit (BDPU) Guard is disabled.

Spanning-Tree Best Practices

With an Ethernet-only topology in which the Cisco Nexus 4000 is uplinked to a vPC environment, the use of a spanning tree is not a concern. For design considerations for vPC on the Cisco Nexus 5000 or Nexus 7000 or for VSS on the Catalyst 6000, please check Cisco's website for the appropriate design guides. Since the Cisco Nexus 4000 currently does not natively run vPC, the server-facing ports cannot be vPC. This would require operating system specific bonding drivers to be able to create a network interface card (NIC) team from the interfaces connected to the Nexus 4000.

Also, as a best practice, it is recommended that you configure the server-facing ports to be edge ports (with PortFast enabled). If the port is configured as a trunk, that port can be configured as an edge trunk.

Cisco Nexus 4000 Link Aggregation Control Protocol for vPC Configurations

The configuration of the upstream switch will vary depending which upstream switch the Cisco Nexus 4000 is connected to. Please refer to the relevant documentation for the Nexus 5000, Nexus 7000, or Catalyst 6000 for vPC or VSS configuration guidelines.

As a best practice on the Cisco Nexus 4000, enabling Link Aggregation Control Protocol (LACP) is recommended. Figure 6 shows an example configuration/topology for a vPC environment with the Nexus 4000.

Figure 6. vPC Topology Configuration with the Cisco Nexus 4000

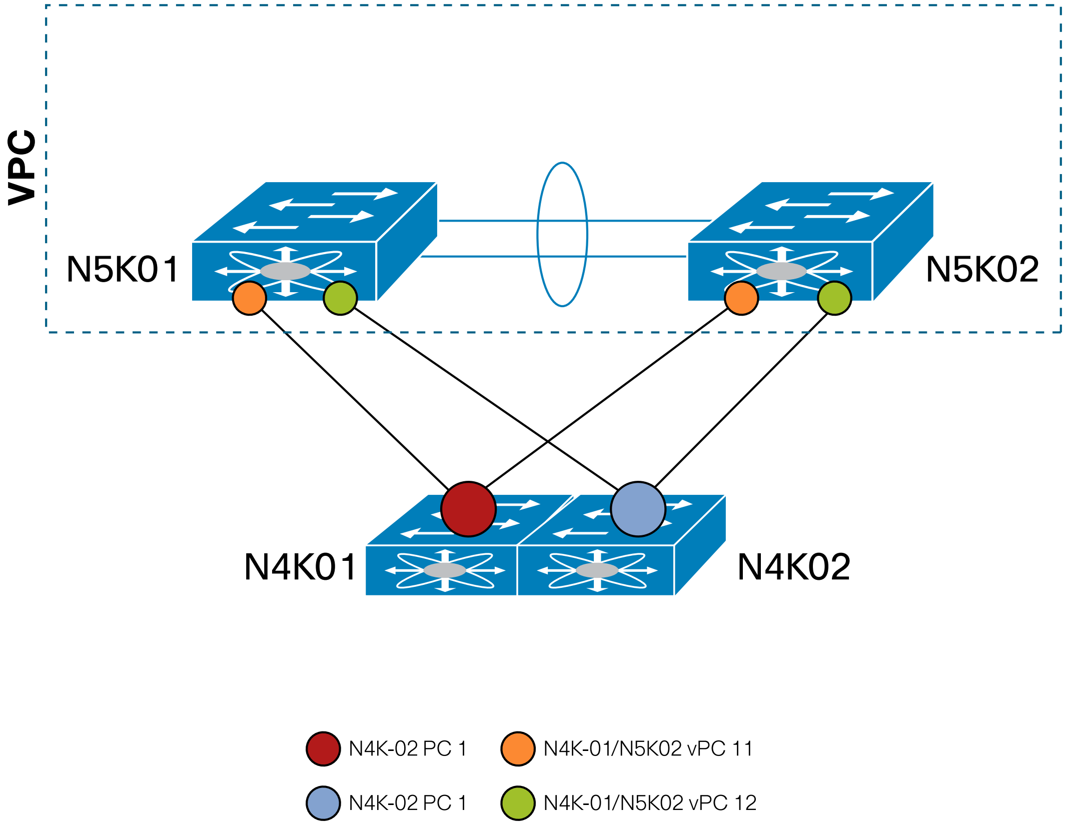

If the port channel is configured as active and the upstream switch is not configured for port channeling, the port-channel ports will have the “Individual” (I) state and run regular spanning tree.

You need to configure the Cisco Nexus 4000 and the upstream switch for LACP for the negotiation to complete and the port channel to form:

```
bch1-n4k-b7(config)# int eth1/15-16
bch1-n4k-b7(config-if-range)# channel-group 1 mode passive
```

The port channel on the Cisco Nexus 4000 then comes online. This indicates that the LACP negotiation is functioning between the upstream vPC system and the Nexus 4000:

```
bch1-n4k-b7# show port-channel summary
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
----------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
----------------------------------------------------------------------
1    Po1(SU)    Eth      LACP      Eth1/15(P)    Eth1/16(P)

The port channel on the Cisco Nexus 5000 also goes up, thanks to the LACP
negotiation:
n5k-1# show vpc brief
[…]
```

```
vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po10   up     10-14,21-24,50,60

vPC status
----------------------------------------------------------------------
id   Port   Status Consistency Reason                     Active vlans
--   ----   ------ ----------- -------------------------- ------------
11   Po11   up     success     success                    10-14,21-24
                                                          ,50,60
```

If the port-channel ports are suspended, there is a good chance that a mismatch exists between the port-channel ports and the switches that are supposed to be brought up in the port channel.

Cisco Nexus 4000 in an FCoE Solution

The Cisco Nexus 4000 is a lossless Ethernet switch that is capable of sending FCoE traffic. Since the Nexus 4000 is not an FCF, the CNAs in the blade chassis will not do a fabric login on the Nexus 4000. This will require forwarding the login process for those CNAs to an actual FCF. The only Cisco switches that function as FCFs are the Cisco Nexus 5000 Series Switches.

For FCoE to maintain the dual storage area network (SAN) fabrics in the Cisco Nexus 4000 topology, the Nexus 4000 uplinks will need to be linked to a single Nexus 5000. This type of topology will allow both FCoE and Ethernet traffic to flow through the same physical links. Figure 7 depicts the supported FCoE topology.

Figure 7. Cisco Nexus 4000 FCoE Supported Topology

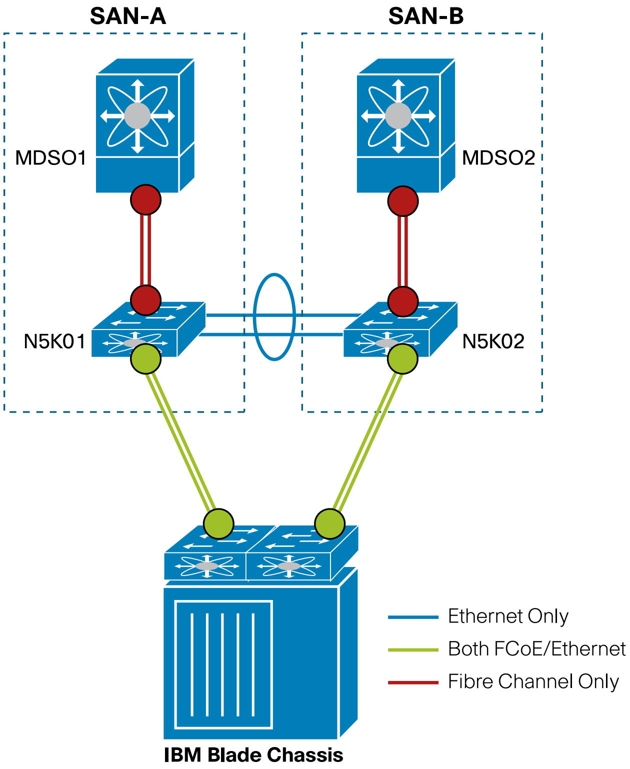

Note: The Ethernet links between the Nexus 5000 switches do not allow FCoE traffic. This ensures the separation of the SAN-A and SAN-B storage fabrics.

FCoE Initialization Protocol Login Process

FCoE devices, such as CNAs and FCoE targets, require a fabric login process. The negotiation protocol that they use is called FCoE Initialization Protocol (FIP), which has its own specific Ethertype, also called FIP. This fabric login process involves three major steps (Figure 8):

- VLAN discovery: Checks for a VLAN that is FCoE capable
- FCF discovery: Solicits for an FCF MAC address in the fabric
- Fabric login (FLOGI): Does a fabric login to the FCF, and the fabric assigns a MAC address to the FCoE device

Figure 8. FIP Login Process

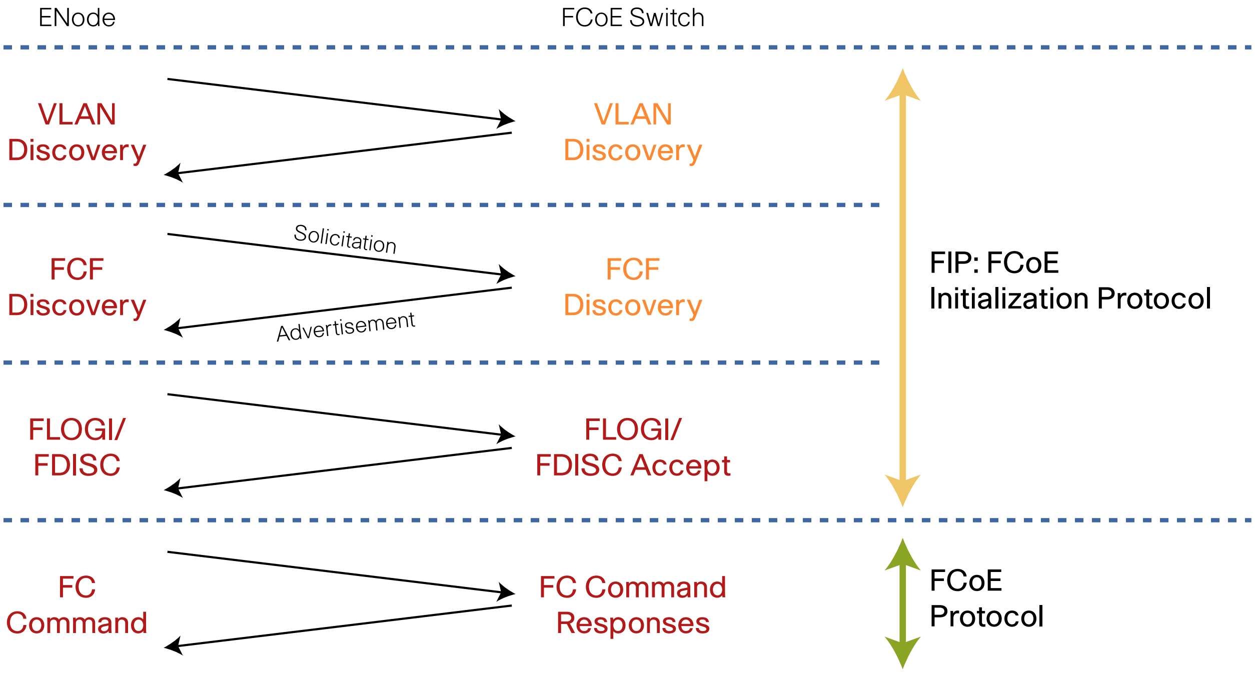

When FCoE devices are directly connected to an FCF, the devices go through a secure fabric login process similar to the login process for FC devices in a FC network. Since FCoE devices connected to the Cisco Nexus 4000 are not directly connected to an FCF and do not log in to the Nexus 4000, the FIP packets will need to be forwarded to an actual FCF (such as a Cisco Nexus 5000).

There is a potential security risk with this login process in that it is possible to spoof the MAC address of the FCoE device. The Cisco Nexus 4000 snoops the FIP packets and provides access control lists (ACLs) during the entire process, thus making the login process secure.

In the configuration of the Cisco Nexus 4000, similar to the Nexus 5000 VLAN configuration for FCoE, the VLAN needs to be enabled for FIP. Also, the VLAN number that is enabled for FIP on the Nexus 4000 has to match the VLAN number on the Nexus 5000 that is enabled for FCoE. It is a best practice to create a VLAN for FCoE traffic that is separate from the standard VLANs for Ethernet traffic.

Spanning-Tree Best Practices

Because this topology is not using vPC upstream, spanning-tree considerations are necessary. As a best practice, the Cisco Nexus 4000 should not be configured to be the root bridge. Check the Nexus 5000 spanning-tree priority for the particular VLANs that will be used for the network. The Nexus 5000 typically has a default priority for the VLANs starting at 32,000. To change the priority on the Nexus 4000, set the spanning-tree priority to anything higher than the priority set on the Nexus 5000. The example below shows how to configure this on the Nexus 4000:

```
bch1-n4k-b7# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
bch1-n4k-b7(config)# spanning-tree vlan 1-3967 priority 36864
bch1-n4k-b7(config)# show spanning-tree root
                                        Root  Hello Max Fwd
Vlan                   Root ID          Cost  Time  Age Dly  Root Port
---------------- -------------------- ------- ----- --- ---  ----------------
VLAN0001         32769 000d.ecb1.1ffc       2    2   20  15      Ethernet1/15
VLAN0030         32798 000d.ecb1.1ffc       2    2   20  15      Ethernet1/15
VLAN0050         32818 000d.ecb1.1ffc       2    2   20  15      Ethernet1/15
VLAN0100         32868 000d.ecb1.1ffc       2    2   20  15      Ethernet1/15
VLAN0182          4278 000d.ecb1.1ffc       2    2   20  15      Ethernet1/15
VLAN0250         33018 000d.ecb1.1ffc       2    2   20  15      Ethernet1/15
VLAN0251         33019 000d.ecb1.1ffc       2    2   20  15      Ethernet1/15
bch1-n4k-b7(config)#
```

Note: The Nexus 4000 is not the root bridge for all VLANs.

Cisco Nexus 4000 FCoE with vPC Upstream Configuration

On the Cisco Nexus 4000, it is not possible to allow FCoE traffic through a vPC uplink. To configure FCoE and also be able to use the benefits of a vPC upstream on the Nexus 4000, the following deployment is recommended. Figure 9 depicts this supported topology.

Figure 9. FCoE with vPC Topology

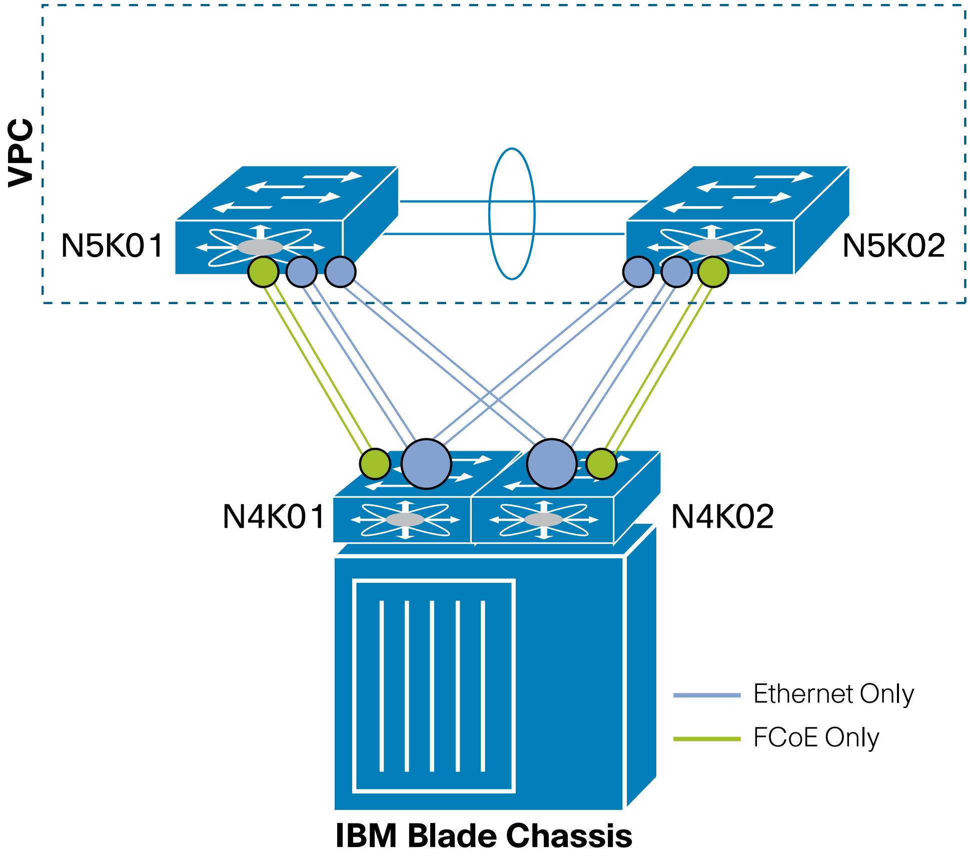

With up to six uplinks available for upstream connectivity, it is possible to use two links for FCoE and two links for vPC connectivity upstream. That leaves two uplinks available to be used for either FCoE or vPC for additional bandwidth if needed. Even though Figure 9 depicts the vPC connectivity to the Cisco Nexus 5000, it can also be connected to the Nexus 7000 or Catalyst 6000 for VSS. The FCoE links must be connected to the Nexus 5000, since it is the only FCF.

With standard Ethernet traffic flowing through the vPC links, those VLANs will not be affected by the spanning tree. For FCoE VLANs, it is still a best practice to make sure that the Nexus 5000 is the root bridge.

Configuring FCoE on the Cisco Nexus 4000

The creation of the virtual fibre channel (vfc) interface on the Cisco Nexus 5000 for FCoE-enabled devices connected to the Nexus 4000 requires these devices to bind via the MAC addresses of the CNAs. Configuring FCoE for the Nexus 4000 and Nexus 5000 devices involves the following tasks:

- Cisco Nexus 4000 Configuration
    - Enable FIP snooping
    - Configure VLAN for FIP snooping
    - Configure FCoE uplinks
    - Configure FCoE server-facing links
- Cisco Nexus 5000 Configuration
    - Enable FCoE
    - Configure VLAN-to-VSAN mapping

- Configure downlinks to the Nexus 4000
- Configure vfc for Nexus 4000 CNAs

Cisco Nexus 4000 FCoE Configuration

By default, the Cisco Nexus 4000 does not have the FIP snooping feature enabled. Before enabling the feature, verify that the FIP snooping license is installed. Otherwise, a 120-day temporary license is provided once the feature is enabled. To show the license status, run the following command:

```
bch1-n4k-b7# show license usage
Feature  Ins  Lic   Status Expiry Date Comments  Count
-----------------------------------------------------------------------------------
--------------
BASIC_STORAGE_SERVICES_PKG           Yes    -   In use Never
-----------------------------------------------------------------------------------
--------------
```

FIP Snooping VLAN

When creating the VLAN to do FIP snooping, the VLAN numbers must match up with the VLAN that is enabled for FCoE from the Cisco Nexus 5000. Once that VLAN is created, the Nexus 4000 needs to specify that VLAN as being enabled for FIP snooping. The following command shows how to do this:

```
bch1-n4k-b7# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
bch1-n4k-b7(config)# vlan 30
bch1-n4k-b7(config-vlan)# fip-snooping enable
bch1-n4k-b7(config-vlan)# show running-config vlan 30
version 4.1(2)E1(1)
vlan 30
  fip-snooping enable
```

Configuring FCoE Uplinks

As a best practice, it is recommended that you configure a minimum of two physical uplinks and create a port channel out of them. Since the uplinks will allow multiple VLANs, this port channel should be set as trunk mode. For FCoE traffic to work properly on the Cisco Nexus 4000, an additional configuration step is needed on this FCoE uplink. The command for the interface, be it a port channel (recommended) or a physical port, is “fip-snooping port-mode fcf.” This will designate the interface to forward FIP packets to the location of the upstream FCF. The following example shows how this is done for a port channel:

```
bch1-n4k-b7#configure terminal
bch1-n4k-b7(config)# feature lacp
bch1-n4k-b7(config)# interface port-channel 2
bch1-n4k-b7(config-if)# interface eth1/15-16
bch1-n4k-b7(config-if)# channel-group 2 mode active
bch1-n4k-b7(config)# interface port-channel 2
bch1-n4k-b7(config-if)# switchport mode trunk
bch1-n4k-b7(config-if)# fip-snooping port-mode fcf
bch1-n4k-b7(config-if)# no shut
```

Configuring FCoE Server-Facing Links

If all of the blade servers in the blade chassis will be used for FCoE, all of the server-facing ports can be configured at the same time. To configure these server-facing ports, have them set in trunk mode to allow multiple VLANs to traverse the links. If you want to specify which VLANs are allowed to traverse the links, you must at a minimum allow the FCoE VLAN and the native VLAN, in order for FCoE to work. Another recommendation is to set these ports to “spanning-tree port type edge trunk,” which is the equivalent of PortFast. That way these ports will not join the spanning tree, and convergence time will be quicker. The following example shows how to configure these ports, where the native VLAN for Ethernet traffic is set to something other than VLAN 1.

```
bch1-n4k-b7#configure terminal
bch1-n4k-b7(config)#interface ethernet 1/1-14
bch1-n4k-b7(config-if)# switchport mode trunk
bch1-n4k-b7(config-if)# switchport trunk allowed vlan 30, 100
```

Note: The above command is not needed, but if you would like to specify the allowed VLANs, make sure the FCoE VLAN is on the allowed list, as shown above.

```
bch1-n4k-b7(config-if)# switchport trunk native vlan 100
bch1-n4k-b7(config-if)# spanning-tree port type edge trunk
```

Note: ***** Warning:*** The edge port type (PortFast) should be enabled only on ports connected to a single host. Connecting hubs, concentrators, switches, bridges, and so on to this interface when the edge port type (PortFast) is enabled can cause temporary bridging loops. Use with caution.

Cisco Nexus 5000 FCoE Configuration

The Cisco Nexus 5000 can be either in switching mode (the default) or in N port virtualization (NPV) mode when configuring FCoE. To support FCoE configuration with the Nexus 4000 devices, the Nexus 5000 must have NX-OS 4.1(3)N1(1) or higher. That code, also known as Cronulla, is when the Nexus 5000 introduced support for FIP-enabled devices. The first thing to do on the Nexus 5000 is to enable FCoE, following the configuration commands below:

```
n5k-1# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
n5k-1(config)# feature fcoe
FC license checked out successfully
fc_plugin extracted successfully
FC plugin loaded successfully
FCoE manager enabled successfully
FC enabled on all modules successfully
```

Configuring VLAN-to-VSAN Mapping

When enabling FCoE on the Cisco Nexus 5000, a single VLAN must be mapped to a single VSAN. By default, when FCoE is enabled on the Nexus 5000, VSAN 1 is created automatically. In most Cisco MDS SAN environments, VSAN 1 is not used in production, and it is recommended that you use a different VSAN number for production FC traffic. The same holds true for VLAN mapping for FCoE. It is always best to use a different VLAN number for FCoE specifically and not allow other Ethernet traffic to be used in that VLAN. In the example below, VSAN 10 is used for FC traffic and VLAN 30 is configured as the FCoE-capable VLAN.

```
n5k-1# configure terminal
n5k-1(config)# vsan database
n5k-1(config-vsan-db)# vsan 2
n5k-1(config-vsan-db)# vlan 30
n5k-1(config-vlan)# fcoe vsan 2
n5k-1(config-vlan)# show vlan fcoe
VLAN      VSAN      Status
--------  --------  --------
30        2         Operational
```

Configuring Downlinks to the Cisco Nexus 4000

It is recommended that you configure a port channel using LACP. The only requirement for configuring the downlinks from the Cisco Nexus 5000 to the Nexus 4000 is that they be trunk ports and allow the native VLAN and the FCoE VLAN to traverse the port channel. The example below shows how to do this:

```
n5k-1#configure terminal
n5k-1(config)# feature lacp
n5k-1(config)# interface port-channel 2 mode active
n5k-1(config-if)# interface eth1/9-10
n5k-1(config-if)# channel-group 2
n5k-1(config)# interface port-channel 2
n5k-1(config-if)# switchport mode trunk
n5k-1(config-if)# no shut
```

Configure Virtual Fibre Channel Interfaces for Cisco Nexus 4000 FCoE CNAs

For FCoE devices connected directly to the Cisco Nexus 5000, we recommend that you configure the vfc interface by binding it to the physical interface on the Nexus 5000. When configuring the vfc interfaces for FCoE-capable devices connected on the Nexus 4000, the binding will need to be bound by MAC address. There are a few ways to find this MAC address on the blade server so that the vfc can be created. You can go to the actual blade server and use the CNA vendor’s tool to figure out what the MAC address is. The other option is to go to the Nexus 4000 switch and run the command “show fip-snooping vlan-discovery.” This command shows what the FCoE MAC address is and what port it is connected to. Here is the output of that command:

```
bch1-n4k-b7# show fip-snooping vlan-discovery
  Legend:
-------------------------------------------------------------------------------
      Interface      VLAN          FIP MAC
-------------------------------------------------------------------------------
      Eth1/1         100         00:c0:dd:04:0c:f5
```

With this information, you have two options for creating the vfc. You can use the command-line interface (CLI) or the Device Manager GUI. The following shows how to create the vfc via the CLI.

```
n5k-1# configure terminal
n5k-2(config)# interface vfc 101
n5k-2(config-if)# bind mac-address 00:c0:dd:04:0c:f5
n5k-2(config-if)# no shutdown
n5k-2(config-if)# show vsan membership
vsan 1 interfaces:
fc2/1     fc2/2  fc2/3  fc2/4  san-port-channel 1  vfc101
vsan 2 interfaces:
vsan 4079(evfp_isolated_vsan) interfaces:
vsan 4094(isolated_vsan) interfaces:

n5k-1(config-if)# vsan database•this will get to the VSAN database
 n5k-1(config-vsan-db)# vsan 2 interface vfc101
n5k-2(config-vsan-db)# show vsan membership
vsan 1 interfaces: fc2/1      fc2/2  fc2/3  fc2/4  san-port-channel 1
vsan 2 interfaces:
vfc101
vsan 4079(evfp_isolated_vsan) interfaces:
n5k-1# show interface vfc101
```

```
vfc101 is up
Bound MAC is 00:c0:dd:04:0c:f5
FCF priority is 128
Hardware is Virtual Fibre Channel
Port WWN is 20:67:00:0d:ec:b2:b9:bf
Admin port mode is F, trunk mode is on
snmp link state traps are enabled
    Port mode is F, FCID is 0xcd0000
    Port vsan is 2 [snip]
```

FCoE Load Balancing

The default load balancing for the Cisco Nexus 4000 is Layer 3, which from a Fibre Channel perspective is equivalent to source-destination load balancing. The equivalent load balancing for FC that is exchange based requires that the Nexus 4000 load balancing be set to Layer 4. In the current configuration, load balancing on the Nexus 4000 is set on a system wide basis, and the same is true for the Nexus 5000.

With the Ethernet port-channeling configuration for the vfcs flowing through, if a link in the port channel fails, the FCoE session will not log out of the FCF (Nexus 5000) but will redirect the I/O to the available link in the port channel. The packets that were on the failed link will be dropped and the I/O will be retried.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Printed in USA

C07-574724-00 01/10