



Port Configuration Examples

This appendix describes port configuration examples relating to FCoE topologies and it includes the following sections:

- [VE Port Configuration Example, page B-1](#)
- [FCoE VE Port Topology Example, page B-1](#)
- [Enabling FCoE and Verifying QoS Configuration, page B-2](#)
- [Configuring VE Ports, page B-5](#)

VE Port Configuration Example

This section provides a sample configuration of the Cisco Nexus 5000 Series switch FCoE VE Port implementation. The configuration covers the switches in switch mode. FCoE initiators are used in this lab. You can attach either FC F Port storage directly to a Nexus 5000 Series switch FC GEMs, or use an FCoE target.



Note

This example can be used for configuring VE ports between two Cisco Nexus 5000 Series switches in both fabrics. It does not include server configurations.

FCoE VE Port Topology Example

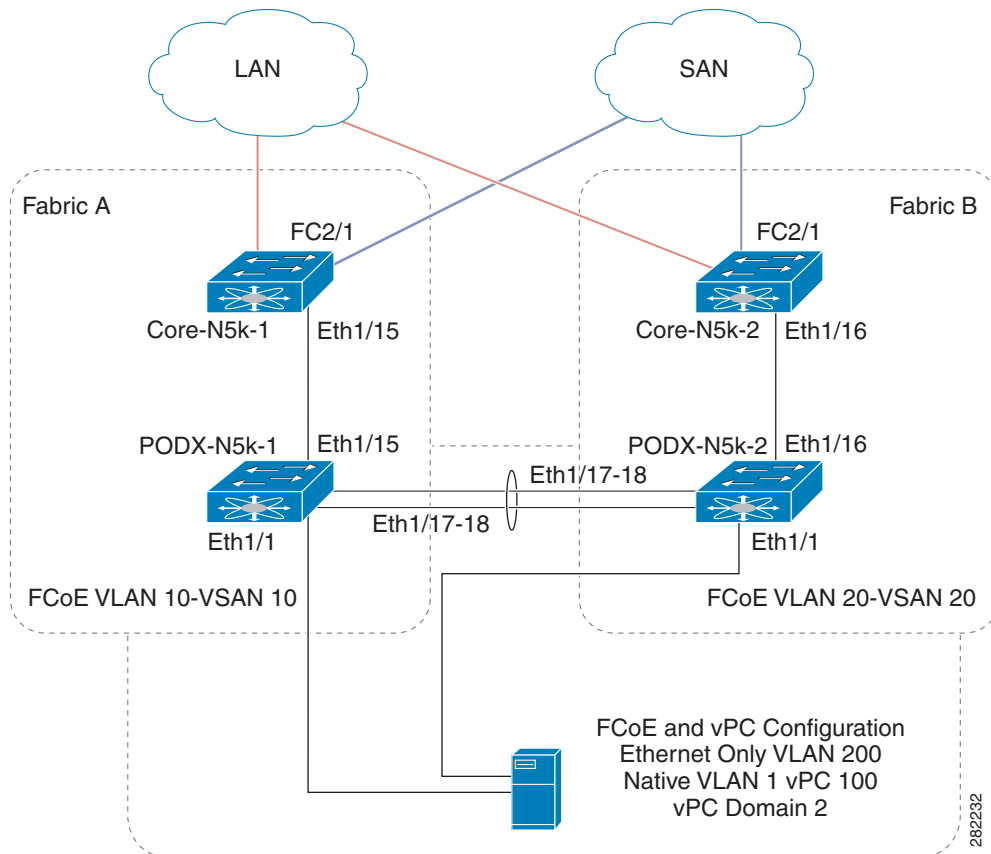
[Figure B-1](#) shows the topology that was used for the configuration example. The following configuration parameters are used in this topology:

- FCoE VLAN for Fabric A: 10
- FCoE VSAN for Fabric A: 10
- FCoE VLAN for Fabric B: 20
- FCoE VSAN for Fabric B: 20
- Ethernet Only VLAN across both fabrics: 200

You should choose these values before the time of configuration.

Send documentation comments to n5kdocfeedback@cisco.com

Figure B-1 FCoE VE Port Topology



Note

The FCoE VLAN/VSAN numbering does not have to be the same within the fabric. As a best practice, use different FCoE VLANs and VSAN numbers between the two fabrics to avoid confusion. Configurations have often been set up to assign ODD VLAN/VSANs for one fabric and EVEN VLANs/VSANs for the other fabric. This is just one example of keeping the numbers separate between the two fabrics

Enabling FCoE and Verifying QoS Configuration

Step 1 Enable FCoE.

```
switch# configure terminal
switch(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
```

Step 2 (Optional) If you do not want to use the default Quality of Service (QoS) settings, specify your own policies:

Send documentation comments to n5kdocfeedback@cisco.com



Note Note: if you use custom policies, **class-fcoe** must be included in your QoS policies.

```
switch(config) system qos
switch(config-sys-qos)# service-policy type qos input fcoe-customized-in-policy-name
switch(config-sys-qos)# service-policy type queuing input
fcoe-customized-in-policy-name
switch(config-sys-qos)# service-policy type queuing output
fcoe-customized-out-policy-name
switch(config-sys-qos)# service-policy type network-qos fcoe-customized-nq-policy-name
```

Step 3 Verify that the FCoE policy maps can be found in the running configuration:



Note Note: If you specified customized QoS policy map names in [Step 2](#), make sure you replace the default map names with your customized map names.

```
switch(config-sys-qos)# show policy-map system

Type network-qos policy-maps
=====

policy-map type network-qos system
  class type network-qos class-fcoe
    match qos-group 1

    pause no-drop
    mtu 2158
  class type network-qos class-default
    match qos-group 0

    mtu 1500

Service-policy (qos) input:    system
policy statistics status:    disabled

Class-map (qos):    class-fcoe (match-any)
  Match: cos 3
  set qos-group 1

Class-map (qos):    class-default (match-any)
  Match: any
  set qos-group 0

Service-policy (queuing) input:  default-in-policy
policy statistics status:    disabled

Class-map (queuing):    class-fcoe (match-any)
  Match: qos-group 1
  bandwidth percent 50

Class-map (queuing):    class-default (match-any)
  Match: qos-group 0
  bandwidth percent 50

Service-policy (queuing) output:  default-out-policy
policy statistics status:    disabled

Class-map (queuing):    class-fcoe (match-any)
  Match: qos-group 1
  bandwidth percent 50
```

Send documentation comments to n5kdocfeedback@cisco.com

```
Class-map (queuing): class-default (match-any)
  Match: qos-group 0
  bandwidth percent 50
```

Quality of Service configuration on the Nexus 5000 series consists of three main constructs:

- Class-map and policy-map type qos: for classification purposes
- Class-map and policy-map type network: for network properties such as drop and no drop, queue size
- Class-map and policy-map type queueing: for bandwidth allocation

This exercise consists of changing the bandwidth allocation and the COS settings for FCoE.

Without proper configuration of class-fcoe in QoS, the following problems may occur:

- vFC interfaces do not come up (CNAs require advertisement of DCB parameters for FCoE)
- Drops noticed for I/Os



Note

QoS has the following guidelines:

- A classification policy-map only applies in input
- A network policy-map applies globally (system)
- A queueing policy-map normally is meaningful in output, but since the exercise uses it to control the bandwidth allocation from CNA to the Cisco Nexus 5000 Series switch, in this case it is applied in input

Beginning in Cisco NX-OS Release 5.0(2)N1(1), you can modify the buffer allocation for no-drop classes:

```
switch(config-pmap-nq)# policy-map type network-qos nqos_policy
switch(config-pmap-nq)# class type network-qos nqos_class
switch(config-pmap-nq-c)# pause no-drop buffer-size <size> pause-threshold <threshold>
resume-threshold <threshold>
```

Step 4 Verify the FCoE system class is active:

```
switch(config-sys-qos)# show queuing interface ethernet 1/1
Ethernet1/1 queuing information:
TX Queuing
  qos-group sched-type oper-bandwidth
  0 WRR 50
  1 WRR 50
RX Queuing
  qos-group 0
  q-size: 370240, HW MTU: 1500 (1500 configured)
  drop-type: drop, xon: 0, xoff: 2314
  Statistics:
  Pkts received over the port : 0
  Ucastpkts sent to the cross-bar : 0
  Mcastpkts sent to the cross-bar : 0
  Ucastpkts received from the cross-bar : 0
  Pkts sent to the port : 0
  Pkts discarded on ingress : 0
  Per-priority-pause status : Rx (Inactive), Tx (Inactive)
  qos-group 1
  q-size: 79360, HW MTU: 2158 (2158 configured)
  drop-type: no-drop, xon: 128, xoff: 252
```

Send documentation comments to n5kdocfeedback@cisco.com

```

Statistics:
Pkts received over the port : 0
Ucastpkts sent to the cross-bar : 0
Mcastpkts sent to the cross-bar : 0
Ucastpkts received from the cross-bar : 0
Pkts sent to the port : 0
Pkts discarded on ingress : 0
Per-priority-pause status : Rx (Inactive), Tx (Inactive)
Total Multicast crossbar statistics:
Mcastpkts received from the cross-bar : 0

```

- Step 5** Repeat [Step 1](#) through [Step 4](#) on both upstream Cisco Nexus 5000 Series switches (CORE_N5k-1 and CORE_N5k-2 in this example).

Configuring VE Ports

FCoE VLAN and VSAN numbering in this example is as follows:

- Fabric A uses FCoE VLAN 10 and VSAN 10
- Fabric B uses FCoE VLAN 20 and VSAN 20



Note

There are two switches in Fabric A and two switches in Fabric B. The FCoE VLAN/VSANs must match between the switches in the same fabric in order to bring up the VE port between them.

- Step 1** Configure the VSAN on the Nexus 5000 Series switch for Fabric A:

```

switch(config)#
switch(config)# vsan database
switch(config-vsan-db)# vsan 10

```

- Step 2** Configure the FCoE VLAN to VSAN mapping and verify that it is up and operational for Fabric A:

```

switch(config)# vlan 10
switch(config-vlan)# fcoe vsan 10
switch(config-vlan)#
switch(config-vlan)# show vlan fcoe
Original VLAN ID      Translated VSAN ID      Association State
-----
10                    10 Operational
switch(config-vlan)#

```

- Step 3** Repeat [Step 1](#) and [Step 2](#) on the upstream Nexus 5000 Series switch in Fabric A.

- Step 4** Configure the VSAN on the Nexus 5000 for Fabric B:

```

switch(config)#
switch(config)# vsan database
switch(config-vsan-db)# vsan 20

```

- Step 5** Configure the FCoE VLAN to VSAN mapping and verify that it is up and operational for Fabric B

```

switch(config)# vlan 20
switch(config-vlan)# fcoe vsan 20
switch(config-vlan)#
switch(config-vlan)# show vlan fcoe
Original VLAN ID      Translated VSAN ID      Association State
-----
20                    20 Operational
switch(config-vlan)#

```

Send documentation comments to n5kdocfeedback@cisco.com

- Step 6** Repeated [Step 1](#) and [Step 2](#) on the upstream Nexus 5000 Series switch in Fabric B.
- Step 7** Configure the underlying 10-Gigabit Ethernet port that the vFC interface will be bound to. The VE port will use this interface as the physical transport for FCoE traffic between the two switches. This interface needs to be configured to trunk the appropriate FCoE VLAN as well as the Ethernet VLAN (in this example, we are using VLAN 200 to carry Ethernet traffic).

The 10-Gigabit Ethernet interfaces connecting the switches in this lab are shown in the topology above:

- Fabric A uses FCoE VLAN 10 and VSAN 10
- Fabric B uses FCoE VLAN 20 and VSAN 20
- PODX-N5K-1 (Fabric A) uses Ethernet 1/15 to connect to CORE N5K1
- PODX-N5K-2 (Fabric B) uses Ethernet 1/16 to connect to CORE N5K2

Configuration for both switches in Fabric A:

```
switch(config)# vlan 200
switch(config)# interface ethernet 1/15
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk allowed vlan 10, 200
switch(config-if)#
```

Configuration for both switches in Fabric A:

```
switch(config)# vlan 200
switch(config)# interface ethernet 1/16
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk allowed vlan 20, 200
switch(config-if)#
```

- Step 8** Configure the vFC interface on the switch that will be bound to the VE port and add this vFC interface to VSAN 44 in the VSAN database:

The vFC numbers for the VE ports are as follows:

- Fabric A uses FCoE VLAN 10 and VSAN 10
- Fabric B uses FCoE VLAN 20 and VSAN 20
- POD1-N5K-1 (Fabric A) uses Ethernet 1/15 to connect to CORE N5K1
- POD1-N5K-2 (Fabric B) uses Ethernet 1/16 to connect to CORE N5K2
- POD1-N5K-1 (Fabric A) uses vfc 15 and binds it to Ethernet 1/15
- POD1-N5K-2 (Fabric B) uses vfc 16 and binds it to Ethernet 1/16

Configuration for both switches in Fabric A:

```
switch(config)# int vfc 15
switch(config-if)# switchport mode e
switch(config-if)# switchport trunk allowed vsan 10
switch(config-if)# bind interface eth 1/15
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)# vsan database
switch(config-vsan-db)# vsan 10 interface vfc 15
switch(config-vsan-db)# show vsan membership
vsan 1 interfaces:
fc2/1          fc2/2          fc2/3          fc2/4
fc2/5          fc2/6          fc2/7          fc2/8
vsan 10 interfaces:
vfc15
vsan 4079(evfp_isolated_vsan) interfaces:
```

Send documentation comments to n5kdocfeedback@cisco.com

```
vsan 4094(isolated_vsan) interfaces:
switch(config-vsan-db) # exit
```

Configuration for both switches in Fabric B:

```
switch(config)# int vfc 16
switch(config-if)# switchport mode e
switch(config-if)# switchport trunk allowed vsan 20
switch(config-if)# bind interface eth 1/16
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)# vsan database
switch(config-vsan-db)# vsan 20 interface vfc 16
switch(config-vsan-db)# show vsan membership
vsan 1 interfaces:
fc2/1          fc2/2          fc2/3          fc2/4
fc2/5          fc2/6          fc2/7          fc2/8
vsan 20 interfaces:
vfc16
vsan 4079(evfp_isolated_vsan) interfaces:
vsan 4094(isolated_vsan) interfaces:
switch(config-vsan-db) # exit
```



Note

Don't forget that these interface configurations must be configured on both sides of the ISL connecting the two switches in the same fabric.

Step 9

Verify that the vFC is up and operational. By default, the vFC will show as trunking. Make sure that it is bound to the correct physical interface and that VSAN 44 shows as allowed and active as well as up on the vFC interface.

Verify both switches in Fabric A:

```
switch(config)# show int vfc 15
vfc15 is trunking
Bound interface is Ethernet1/15
Hardware is Virtual Fibre Channel
Port WWN is 20:0e:00:0d:ec:b4:43:7f
Peer port WWN is 00:00:00:00:00:00:00:00
Admin port mode is E, trunk mode is on
snmp link state traps are enabled
Port mode is TE
Port vsan is 10
Trunk vsans (admin allowed and active) (10)
Trunk vsans (up) (10)
Trunk vsans (isolated) ()
Trunk vsans (initializing) ()
1 minute input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
1 minute output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
13 frames input, 1028 bytes
0 discards, 0 errors
13 frames output, 1180 bytes
0 discards, 0 errors
last clearing of "show interface" counters never
Interface last changed at Sat Nov 6 17:58:39 2010
```

Verify both switches in Fabric A:

```
switch(config)# show int vfc 16
vfc16 is trunking
Bound interface is Ethernet1/16
Hardware is Virtual Fibre Channel
Port WWN is 20:0e:00:0d:ec:b4:43:7d
Peer port WWN is 00:00:00:00:00:00:00:00
```

Send documentation comments to n5kdocfeedback@cisco.com

```
Admin port mode is E, trunk mode is on
snmp link state traps are enabled
Port mode is TE
Port vsan is 20
Trunk vsans (admin allowed and active) (20)
Trunk vsans (up) (20)
Trunk vsans (isolated) ()
Trunk vsans (initializing) ()
1 minute input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
1 minute output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
 13 frames input, 1028 bytes
 0 discards, 0 errors
 13 frames output, 1180 bytes
 0 discards, 0 errors
last clearing of "show interface" counters never
Interface last changed at Sat Nov 6 17:58:39 2010
```