# Network Configuration Example

## Configuring Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric

Release
NCE 74

Modified: 2016-08-01

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net



*Network Configuration Example Configuring Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric*
NCE 74


The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

## END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at http://www.juniper.net/support/eula.html. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

# Table of Contents

CHAPTER 1

# Configuring Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric



## About This Network Configuration Example

This network configuration example describes the Fibre Channel over Ethernet to Fibre Channel (FCoE-FC) gateway capability of the Juniper Networks® QFX3500 Switch and the benefits of using it. It also includes a step-by-step procedure for configuring an FCoE-FC gateway.

## Understanding FCoE-FC Gateway Benefits

Configuring a QFX3500 switch or a Node device on a QFabric™ family of products as a Fibre Channel over Ethernet to Fibre Channel (FCoE-FC) gateway enables you to converge your Ethernet network and your FC storage area network (SAN) traffic. The FCoE-FC gateway handles both FC traffic encapsulated in Ethernet (FCoE traffic) and native FC traffic from the FC SAN. The FCoE-FC gateway encapsulates native FC traffic from the FC SAN in Ethernet before forwarding it to the Ethernet network as FCoE traffic. The FCoE-FC gateway also decapsulates FCoE traffic from the Ethernet network before forwarding it to the FC SAN as native FC traffic.

The ability to converge Ethernet traffic (as FCoE traffic) and native FC traffic is cost-effective because it eliminates the need to encapsulate native FC traffic and decapsulate FCoE traffic on the FC SAN edge switch. Thus, the FC SAN switch does not need extra adapters to handle FCoE traffic. Instead, the FC SAN edge switch only needs to handle native FC traffic.

Using a QFX3500 Switch as an FCoE-FC gateway involves configuring native FC interfaces on ports connected to the FC SAN, and configuring an FCoE VLAN interface that includes Ethernet ports connected to the FCoE (Ethernet) network. You configure the native FC interfaces and the FCoE VLAN interface as part of a local FC fabric on the FCoE-FC

gateway. In addition to creating an FCoE-FC gateway fabric, you must also configure proper class-of-service treatment to ensure lossless transport of the storage traffic across the Ethernet network.

**Related Documentation**

- Understanding Interfaces on an FCoE-FC Gateway on page 6
- Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 17

## Understanding Interfaces on an FCoE-FC Gateway

When the QFX Series functions as an FCoE-FC gateway to connect FCoE devices on an Ethernet network to a Fibre Channel (FC) switch in a storage area network (SAN), it handles FCoE traffic from hosts and native FC traffic from the FC switch. To support this architecture, each local FC fabric configured on the gateway (in the **fc-fabrics** configuration hierarchy) must have:

- An Ethernet-network-facing F_Port interface for the FCoE VLAN to connect to FCoE device VN_Ports in the form of an FCoE VLAN interface. Multiple VF_Ports are initiated on the F_Port interface, one VF_Port for each ENode that logs in to the FC network.
- One or two blocks of six proxy N_Port (NP_Port) interfaces to connect to FC switch fabric ports (F_Ports).

Each FC fabric is local to the gateway on which you configure it. This means that both the FC switch and the FCoE devices must be connected to the same gateway (QFX3500 switch or QFabric system Node device), and that all of the interfaces configured for the local fabric also must be on that gateway. FC fabric traffic does not flow between different Node devices in a QFabric system.

This topic describes:

- Native FC Interfaces to the FC Switch on page 6
- FIP Login Session Limits on page 8
- Trusted and Untrusted Interfaces on page 12
- Buffer-to-Buffer Credit Recovery on page 12
- FCoE VLAN Interface to FCoE Devices on page 13
- Assigning Interfaces to a Fibre Channel Fabric on page 16
- Deleting a Fibre Channel Interface on page 16

### Native FC Interfaces to the FC Switch

You must configure either 6 or 12 of the physical interfaces on the gateway as native FC NP_Port interfaces to connect to FC switch F_Port interfaces. By default, all of the gateway interfaces are Ethernet interfaces, so you must explicitly configure the interfaces that you want to use as FC interfaces.

You can configure the FC-capable ports xe-0/0/0 through xe-0/0/5 as fc-0/0/0 through fc-0/0/5, and ports xe-0/0/42 through xe-0/0/47 as fc-0/0/42 through fc-0/0/47 to

create blocks of native FC interfaces. You cannot individually configure a single port as a native FC interface. Within these port blocks, you cannot mix FC interfaces with Ethernet interfaces. All of the ports in a block must be either native FC interfaces or Ethernet interfaces.

You cannot configure ports xe-0/0/6 through xe-0/0/41 and ports xe-0/1/1 through xe-0/1/15 as native FC ports; they can only be Ethernet ports. Native FC ports do not handle Ethernet traffic (including FCoE traffic); they handle only native FC traffic and must connect to native FC ports.

You can configure:

- Six native FC interfaces by configuring either ports xe-0/0/0 through xe-0/0/5 as fc-0/0/0 through fc-0/0/5 *or* ports xe-0/0/42 through xe-0/0/47 as fc-0/0/42 through fc-0/0/47.
- Twelve native FC interfaces by configuring ports xe-0/0/0 through xe-0/0/5 as fc-0/0/0 through fc-0/0/5 *and* ports xe-0/0/42 through xe-0/0/47 as fc-0/0/42 through fc-0/0/47.
- No native FC interfaces by leaving ports xe-0/0/0 through xe-0/0/5 and ports xe-0/0/42 through xe-0/0/47 in their default state as Ethernet interfaces.

Each native FC interface can belong to only one local FC fabric configured on the gateway. You can configure up to 12 FC fabrics on a gateway, but each FC fabric must use different native FC interfaces to connect to an FCF. (Although the native FC ports are configured in blocks, each individual port can belong to a different FC fabric.) Native FC interfaces can be configured as loopback interfaces.


- Port Mode on page 7
- NPIV on page 7
- Port Speed on page 8


### Port Mode

The gateway presents a proxy N_Port (NP_Port) interface to the FC switch. An NP_Port connects to a single FC switch F_Port using a point-to-point link (in other architectures an N_Port can also connect in a point-to-point link to another N_Port, but that is not a valid configuration on the gateway).

You must explicitly configure each native FC interface connected to an FC switch as an NP_Port. The gateway NP_Ports act as a proxy for the FCoE device virtual N_Ports (VN_Ports) when the VN_Ports attempt to connect to the FC switch.

When the FC switch is a trusted switch, configure the fabric as **fcoe-trusted** to reduce overhead caused by the VN_Port to VF_Port (VN2VF_Port) FIP snooping filters that are automatically installed on untrusted ports.

### NPIV

FC requires a unique point-to-point link between the FC switch and each host N_Port. The gateway creates an independent virtual link for each FCoE device session by mapping

each FCoE device to a virtualized N_Port through the gateway's proxy function. This process is called N_Port ID virtualization (NPIV).

NPIV makes each virtual link look like a dedicated point-to-point link to the FC switch. In this way, multiple FCoE devices, multiple applications, and multiple virtual machines on an FCoE device can connect to an FC switch using one physical port instead of using a physical port for each host connection. The virtual link creates a secure boundary between traffic from different sources that are on a single physical port.

FCoE-FC gateway mode implements NPIV as follows:

1. An NP_Port on the gateway comes up and logs in to the attached F_Port on the FC switch. The FC switch sees the gateway port as a physical FC device N_Port and assigns it a unique FCID. This establishes the physical point-to-point link between the gateway and the FC switch.
2. The gateway receives a FIP discovery message from an FCoE device that seeks to log in to the FC network. To the FCoE device, the gateway presents a virtual F_Port (VF_Port) interface and appears to be an FCF.
3. The gateway converts the FCoE device's message into an FC fabric discovery (FDISC) message and sends it through the least-loaded physical NP_Port to the FC switch. The FDISC message requests an FCID for the new virtual link.
4. The FC switch processes the request, accepts it, assigns a unique FCID for the connection, and sends the response.
5. The gateway maps the FC switch response to the host FCoE device's VN_Port and sends a FIP acceptance advertisement to the FCoE device.
6. The FCoE device accepts the FCID.

If the FC switch rejects the FDISC, the gateway relays the rejection to the FCoE device VN_Port.

### Port Speed

The gateway supports configuring FC port speeds of 2 Gbps, 4 Gbps, or 8 Gbps. FC ports can also autonegotiate the port speed to 2, 4, or 8 Gbps.

## FIP Login Session Limits

A FIP login session is a fabric login (FLOGI) or fabric discovery (FDISC) login to the FC SAN fabric. (A session here does not refer to an end-to-end server-to-storage session; there is no limit to the number of end-to-end server-to-storage sessions.) You can limit the maximum number of FIP login sessions on each gateway Node device (QFX3500 switch or QFabric system Node device configured in FCoE-FC gateway mode), on each local gateway FC fabric, and on each individual NP_Port interface in a local FC fabric:

- Gateway Node devices and Node groups—The total number of FIP login sessions on the gateway Node or Node group (the sum of the sessions on all of the NP_Port interfaces in all of the local FC fabrics on the gateway Node or Nodes) cannot exceed the limit. When a gateway reaches the maximum session limit, the gateway sends subsequent multicast discovery advertisements (MDAs) with the availability bit set

to 0 (zero) to prevent additional ENode login attempts. If the maximum number of sessions is running on the gateway, ENodes cannot use the gateway to log in new sessions to the FC switch. When the number of sessions falls below the maximum, the gateway sets the availability bit in MDAs to 1 so that ENodes can again log in new sessions. When a session slot becomes available, the system accepts the first session request to fill the slot.

- FC fabric—The total number of FIP login sessions on an FC fabric (the sum of the sessions on all of the NP_Port interfaces that belong to the fabric) cannot exceed the limit. When a fabric reaches the maximum session limit, the gateway sends MDAs associated with that fabric with the availability bit set to 0 to prevent additional ENode login attempts.

NOTE: **Other FC fabrics on the same gateway can still accept ENode logins as long as the maximum session limit for those fabrics and the maximum session limit for the gateway (the Node device) have not been met.**

- NP_Port interfaces—The total number of FIP login sessions cannot exceed the interface's limit. When an interface reaches the maximum session limit, the gateway removes it from the load-balancing list for that FC fabric to prevent the gateway from attempting to assign new sessions to the interface. Other interfaces in the FC fabric can still accept logins until the FC fabric or gateway reaches its maximum session limit. However, the interface that reached the maximum session limit cannot be assigned new sessions until the number of sessions on the interface falls below the limit.

BEST PRACTICE: **Configure a maximum session limit for each NP_Port interface that is less than or equal to the number of FIP sessions the directly connected FC switch port is configured to support. This prevents the gateway from attempting to assign new login sessions to an interface when the connected FC switch port reaches its maximum number of sessions.**


- FCoE Trusted and Untrusted Interface Session Limits on page 9
- Configuring Consistent Session Limits on page 10
- Decreasing Session Limits on page 10
- Increasing Session Limits on page 11
- Effect of Deactivating and Then Reactivating the Configuration on Session Limits on page 11


### FCoE Trusted and Untrusted Interface Session Limits

The maximum number of VN2VF_Port FCoE login sessions that each gateway can support is 2500 sessions, regardless of whether interfaces are trusted or untrusted. (In software releases earlier than Junos OS Release 12.3, the session limit on untrusted interfaces and untrusted fabrics was 376 sessions.)

## Configuring Consistent Session Limits

The system does not perform commit checks to enforce consistent session limit configuration. For example, the system does not prevent you from configuring a higher limit for ENode sessions than the total session limit for the gateway Node device, or from configuring a higher limit on an interface than on the fabric to which the interface belongs.

To prevent unexpected FIP login rejections, you should configure consistent Node device, fabric, and interface session limits. For example:

- The session limit of an interface should not exceed the session limit of the fabric to which it belongs.
- For interfaces that belong to the same fabric, the sum of the interface session limits should not exceed the fabric session limit.
- The fabric session limit should not exceed the session limit of the gateway Node device.
- For fabrics that belong to the same gateway Node device, the sum of the fabric session limits should not exceed the Node device session limit.

Session limit configuration considerations include:

- The fabric session limit restricts how many sessions can run on the NP_Port interfaces that belong to that fabric. If the combined session limits of the interfaces exceed the fabric session limit, the total number of sessions on the interfaces is the fabric limit.

  For example, if a fabric has three NP_Port interfaces, and each NP_Port interface has a limit of 500 sessions (total of 1500 sessions for the three interfaces), but the fabric has a limit of 1000 sessions, the combined number of sessions on the three interfaces is limited to 1000 sessions.
- The gateway Node device session limit restricts how many sessions can run on the fabrics that belong to that gateway. If the combined session limits of the fabrics exceed the gateway Node device session limit, the total number of sessions on the fabrics is the gateway Node device limit.

  For example, if a gateway has two fabrics, and each fabric has a limit of 1000 sessions (total of 2000 sessions for the two fabrics), but the gateway has a limit of 1500 sessions, the combined number of sessions on the two fabrics is limited to 1500 sessions.

Hierarchically, the gateway Node device session limit is the maximum limit for all sessions on the gateway, regardless of fabric and interface session limits. In the same way, the fabric session limit supersedes the interface session limit.

When session limits are exceeded, no new logins are accepted until a session slot becomes free.

## Decreasing Session Limits

If you decrease the session limit, the currently logged in sessions are terminated as follows:

- Gateway Node devices and Node groups—Decreasing the session limit terminates all of the sessions on the Node device (all sessions on all interfaces on all fabrics). If the gateway Node device is part of a Node group, all sessions on all members of the Node group are terminated.
- Fabric—Decreasing the session limit terminates all of the sessions on all of the interfaces that belong to the fabric.
- NP_Port interfaces—Decreasing the session limit terminates all of the sessions on the interface and also terminates all of the sessions on any other interfaces that belong to the same fabric.

After you decrease a session limit, the sessions are terminated even if the new session limit is greater than the number of currently active sessions. For example:

- An interface has 300 active sessions.
- The current session limit is 1000 sessions.
- You decrease the session limit to 500 sessions and commit the new configuration.
- All 300 sessions are logged out, even though the new session limit is greater than the number of sessions running.

After the session limit change takes effect, the ENodes log in again and establish new sessions, up to the new session limits.

### Increasing Session Limits

Increasing the session limits does not disrupt logged in sessions.

### Effect of Deactivating and Then Reactivating the Configuration on Session Limits

If you decrease session limits, all ENodes are logged out. Deactivating and then reactivating the configuration can have the same effect as decreasing the session limit, which results in the ENodes being logged out.

The ENode logouts occur because when you deactivate the configuration, the system reverts to the default session limit of 2500 sessions (the maximum number of sessions). When you reactivate the configuration, the system uses the configured session limit. Unless the configured session limit is equal to the maximum session limit, reactivating the configuration decreases the session limit, which causes the ENodes to be logged out.

For example, if you:

1. Configure and commit a limit of 400 sessions.
2. Allow ENodes to log in and start sessions.
3. Deactivate the configuration.
4. Reactivate the configuration.
5. The ENode sessions are logged out because deactivating the session increased the session limit from 400 to 2500.

Because an increase in the session limit does not affect existing sessions, the running ENode sessions are not affected. However, reactivating the configuration decreased the session limit from 2500 back to 400. The session limit decrease causes the ENode sessions to be logged out.

## Trusted and Untrusted Interfaces

By default, gateway fabric interfaces are untrusted interfaces. If you do not configure a gateway fabric as an FCoE trusted fabric to set all of the gateway fabric interfaces as trusted interfaces, the gateway installs VN2VF_Port FIP snooping filters on the fabric ports.

If you configure a gateway fabric as an FCoE trusted fabric, the gateway does not install VN2VF_Port FIP snooping filters on the fabric interfaces. This is usually done when the gateway is connected to an FCoE transit switch that has VN2VF_Port FIP snooping enabled.

Regardless of whether an interface is trusted or untrusted, the maximum session limit is 2500 sessions.

NOTE: **The session limit for a Node group is the same as the session limit for an individual Node device, 2500 sessions. Even if more than one Node device in a Node group is acting as an FCoE-FC gateway, the total maximum number of sessions on all Node devices in the Node group is 2500 sessions.**

The default maximum login session value for Node devices (on QFabric systems, the maximum applies to each Node device), trusted fabrics, and interfaces in trusted fabrics is 2500 sessions.

## Buffer-to-Buffer Credit Recovery

Buffer-to-buffer credits represent the number of receive buffers an interface can use to store FC frames. Buffer-to-buffer credit determines buffer-to-buffer flow control. When an interface transmits a frame, it decrements its buffer-to-buffer credit count by one. When the destination interface forwards the frame and frees a buffer, it sends a receiver ready (R_RDY) primitive to the transmitting interface. Each R_RDY primitive the transmitting interface receives increments its buffer-to-buffer credit count by one.

Both interfaces on an FC link track buffer-to-buffer credits. As long as buffer-to-buffer credits are available, the transmitter continues to send frames. If the number of buffer-to-buffer credits reaches zero (0), transmission stops until buffer-to-buffer credits are available, as indicated by the reception of an R_RDY primitive. Buffer-to-buffer credits can compensate for long cable distances to limit throughput and prevent buffer overflow.

However, if frame corruption or errors transmitting R_RDY primitives occur, the buffer-to-buffer credit counters on the sending and receiving interfaces do not have the same values. This causes the permanent loss of buffer-to-buffer credits. When credits are lost, the buffer credit count can decrement to zero and indicate that there is no available buffer space even if buffer space is actually available. This can result in unnecessary link idle time.

To recover lost buffer-to-buffer credits, you can configure a buffer-to-buffer credit state change number (BB_SC_N). BB_SC_N must be configured on both ends of the connection. If only one end of the connection is configured for BB_SC_N, the feature is disabled. The two directly connected FC interfaces communicate the BB_SC_N value during fabric login (FLOGI).

When you enable BB_SC_N on the interfaces on both ends of an FC link, the interfaces exchange buffer-to-buffer state change send (BB_SCs) and buffer-to-buffer state change receive (BB_SCr) primitives to track the number of frames sent and the number of R_RDY primitives received. The state change number determines the number of frames and R_RDY primitives the interfaces exchange between consecutive BB_SCn primitives and between consecutive BB_SCr primitives. The state change primitives inform each interface of the other interface's frame count and R_RDY count states.

The state counters should match so that each interface knows and agrees with the other interface's state. If the interface at either end of the link detects a discrepancy, it knows that a frame or an R_RDY primitive was corrupted or dropped.

For example, if a receiving interface has sent two R_RDY primitives but the BB_SCr that the interface receives from the sending interface only counts one R_RDY primitive received, it reveals that one R_RDY primitive was not delivered successfully and that one buffer-to-buffer credit was lost. When one of the interfaces on the link detects a discrepancy, the interfaces can take corrective action and recover the lost buffer-to-buffer credits.

Enabling the buffer-to-buffer credit recovery feature does not impact buffer resources and has an insignificant impact on processing resources.

If buffer-to-buffer credit recovery is not used, then when there is no buffer credit on a port, a timeout and recovery mechanism prevents buffer overflow.

## FCoE VLAN Interface to FCoE Devices

Each FC fabric configured on the gateway includes at least one FCoE VLAN interface to connect the FCoE devices on the FCoE VLAN to the FC switch. (Including the FCoE VLAN interface and the native FC interfaces in the FC fabric configuration connects them.) FCoE VLANs can include any Ethernet interface on the switch that is in tagged-access or trunk mode. The best practice is to configure Ethernet interfaces that belong to FCoE VLANs in **tagged-access** port mode.

NOTE: **The Ethernet interfaces that connect to FCoE devices must include a native VLAN to transport FIP traffic, because FIP VLAN discovery and notification frames are exchanged as untagged packets.**

FCoE VLANs should carry only FCoE traffic. You should not mix FCoE traffic and standard Ethernet traffic on the same VLAN.

NOTE: FCoE VLANs (any VLAN that carries FCoE traffic) support only Spanning Tree Protocol (STP) and link aggregation group (LAG) Layer 2 features.

Each FCoE VLAN interface can belong to only one FC fabric configured on the gateway. A gateway FC fabric can have more than one FCoE VLAN, but each FCoE VLAN in the FC fabric must belong only to that FC fabric. You can configure more than one FC fabric on a gateway; each FC fabric must use different FCoE VLAN interfaces to connect to FCoE devices.

NOTE: Storm control must be disabled on all Ethernet interfaces that belong to the FCoE VLAN to prevent FCoE traffic from being dropped.



## Port Mode

You must explicitly configure the FCoE VLAN interface in F_Port mode. All members of the FCoE VLAN use the FCoE VLAN interface as the connection to the gateway NP_Port interfaces and ultimately to the FC switch.

All of the 10-Gigabit Ethernet interfaces that are members of an FCoE VLAN should be configured as **tagged-access** port mode interfaces. However, the system also supports configuring these interfaces in **trunk** port mode.

BEST PRACTICE: Use tagged-access port mode for Ethernet interfaces that are connected to converged network adapters (CNAs) in FCoE access devices.

Use trunk port mode when an Ethernet interface is an interswitch link (ISL)—that is, when the port is connected to another switch. For example, if a port is connected to a transit switch that is performing VN2VF_Port FIP snooping, configure the port in trunk mode and as an FCoE trusted port.

The tagged-access port mode was not available in Junos OS Release 11.3 and earlier releases. In Release 11.3 and earlier, only trunk port mode was used for Ethernet interfaces that belong to an FCoE VLAN. Because tagged-access mode is now available, using trunk mode for interfaces connected to FCoE CNAs is not recommended.

If an existing configuration uses trunk mode for ports connected to FCoE CNAs, you can change the port mode to tagged-access without disrupting traffic. Although we recommend changing the port mode of these ports from trunk mode to tagged-access mode as a best practice, it is not mandatory.

New configurations should use tagged-access mode for interfaces that connect to FCoE devices.

There are several advantages of configuring Ethernet ports connected to FCoE devices in **tagged-access** mode instead of in **trunk** mode:

- It is standard practice to configure ISL ports as trunk ports.
- It is standard practice not to configure ports that connect to servers as trunk ports.
- When an interface goes down, if that interface is in **trunk** mode, then the FCoE sessions on that interface are terminated only after the gateway stops receiving FIP keepalive messages from the ENode and exceeds 2.5 times the FIP keepalive timeout advertisement value. If the interface is in **tagged-access** mode and the interface goes down, the gateway sends a FIP message to terminate the sessions on the interface.
- Similarly, if an ENode session moves from one interface to another interface, if the original interface is in **trunk** mode, the session is not removed from the interface until the gateway stops receiving FIP keepalive messages and exceeds 2.5 times the FIP keepalive advertisement timeout value. But if the interface is in **tagged-access** mode, the gateway detects that the session is no longer on the interface, does not refresh the FIP keepalive timer, and thus ages out the session.

NOTE: FIP is enabled on the FCoE VLAN, which is a Layer 3 interface. As with other Layer 3 interfaces under Junos OS, when the last member (10-Gigabit Ethernet interface) of the FCoE VLAN is deleted, the FCoE VLAN interface is internally marked as "down." When the Layer 3 FCoE VLAN interface is marked as "down", FIP stops running on it. When the last member interface is deleted from an FCoE VLAN and FIP stops running, the result could be an immediate timeout for the VN_Ports that were connected on that interface, regardless of whether the port mode is tagged-access or trunk.

## Disabling Storm Control on FCoE Interfaces

Storm control is enabled by default on all interfaces. When a QFX3500 switch or a QFX3500 Node device is acting as an FCoE-FC gateway, disable storm control on the QFX3500 switch or QFX3500 Node device, and if desired, enabled it on ports that are not part of an FCoE-FC gateway VLAN. Storm control is not supported on the FCoE interfaces of an FCoE-FC gateway VLAN. Configuring storm control on an Ethernet interface and including that interface in an FCoE-FC gateway may have undesirable effects, including FCoE packet loss.

You can disable storm control in either of two ways:

- Disable storm control on all interfaces, then enable storm control on the interfaces you want to use storm control. (From the default configuration, you cannot disable storm control on individual interfaces because the default configuration enables storm control on **all** interfaces, not on individual interfaces.)

For example, if you want interfaces xe-0/0/20, xe-0/0/21, and xe-0/0/22 to use storm control, disable storm control on all interfaces, then enable storm control on those three interfaces:

1. Disable storm control on all interfaces:

   ```
   user@switch# delete ethernet-switching-options storm-control interface all
   ```

2. Enable storm control on interfaces xe-0/0/20, xe-0/0/21, and xe-0/0/22:

   ```
   user@switch# set ethernet-switching-options storm-control interface xe-0/0/20
   user@switch# set ethernet-switching-options storm-control interface xe-0/0/21
   user@switch# set ethernet-switching-options storm-control interface xe-0/0/22
   ```

- Disable storm control for all unknown unicast traffic on all interfaces by including the following statement in your configuration:

  ```
  user@switch# set ethernet-switching-options storm-control interface all no-unknown-unicast
  ```

### NPIV Support

The gateway supports FCoE device NPIV. For example, a single physical FCoE device can have multiple virtual machines running on it. Each virtual machine can instantiate a separate virtual connection to the gateway, which results in its own virtual link to the FC switch. In this way, an FCoE device can have multiple separate connections to the FC SAN on a single physical port.

This is similar to the NPIV function the gateway performs with the FC switch to support multiple virtual FCoE device connections on one physical NP_Port.

The gateway presents multiple VF_Port interfaces on each FCoE VLAN interface to support the requirement for unique, secure virtual links.

### VN2VF_Port FIP Snooping

The FCoE-facing ports that belong to an FCoE VLAN on a gateway are enabled for VN2VF_Port FIP snooping automatically. You can disable VN2VF_Port FIP snooping on any individual interface by configuring it as a trusted interface.

## Assigning Interfaces to a Fibre Channel Fabric

You assign at least one FCoE VLAN interface and at least one native FC interface to each FC fabric you configure on the gateway. All of the interfaces that belong to an FC fabric must reside on the same gateway device. Interfaces on different gateways cannot belong to the same FC fabric, because an FC fabric is local to a single gateway device.

## Deleting a Fibre Channel Interface

To delete an FC interface or an FCoE VLAN interface, you must delete the interface from the fabric first and then delete the interface from the switch.

**Related Documentation**

- Understanding FCoE-FC Gateway Benefits on page 5
- Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric on page 17

# Example: Setting Up Fibre Channel and FCoE VLAN Interfaces in an FCoE-FC Gateway Fabric

To transmit Fibre Channel (FC) traffic between FCoE devices and a storage area network (SAN) FC switch, you configure a local FC fabric on the gateway. The gateway FC fabric includes FCoE and native FC interfaces, and a VLAN to carry FCoE traffic from FCoE-capable devices. The gateway FC fabric creates the path between the FCoE devices and the SAN.

This example describes how to configure the interfaces, VLAN, and FC fabric to connect FCoE devices to the FC switch and route traffic between the VLAN and FC interfaces:



## Requirements

This example uses the following hardware and software components:

- A configured and provisioned Juniper Networks QFX3500 Switch to act as an FCoE-FC gateway
- FCoE-capable devices in an Ethernet network equipped with converged network adapters (CNAs)
- An FC switch to transmit and receive native FC traffic
- FC storage devices in the SAN
- Junos OS Release 11.1 or later for the QFX Series

NOTE: **This configuration example has been tested using the software release listed and is assumed to work on all later releases.**

## Overview

No interfaces are configured for FC network connectivity by default. You need to configure the FC fabric and its interfaces explicitly. Each FC fabric consists of a combination of at least one FCoE VLAN interface between the FCoE-FC gateway and the FCoE devices, and one or more native FC interfaces between the FCoE-FC gateway and the FC switch.

An FCoE VLAN interface connects the FCoE-FC gateway to FCoE devices. FCoE traffic between the devices and the FCoE-FC gateway requires a dedicated VLAN used only for FCoE traffic. You cannot mix standard Ethernet traffic and FCoE traffic on the FCoE VLAN.

NOTE: IGMP snooping is not supported on FCoE VLANs. Disable IGMP snooping on FCoE VLANs.

Disable storm control on all Ethernet interfaces that belong to the FCoE VLAN to prevent FCoE traffic from being dropped. Configuring storm control on an Ethernet interface and including that interface in an FCoE-FC gateway may have undesirable effects, including FCoE packet loss. After disabling storm control on all interfaces, enable storm control on any interfaces that are not part of an FCoE-FC gateway on which you want to use storm control.

When FCoE frames enter the FCoE-FC gateway, the gateway:

1. Strips the Ethernet encapsulation from the FCoE frames.
2. Sends the resulting native FC frames to the FC switch through the gateway's native FC interfaces.

Each FC interface and FCoE VLAN interface can belong to only one FC fabric. Different FC fabrics must use different native FC interfaces and different FCoE VLAN interfaces. Multiple FC fabrics on the FCoE-FC gateway can connect to the same FC switch, but they must use different FC interfaces and different FCoE VLAN interfaces.

The Ethernet interfaces that belong to the FCoE VLAN should be configured in tagged-access port mode and must include the native VLAN because FIP VLAN discovery and notification frames are exchanged as untagged packets. These Ethernet interfaces require a maximum transmission unit (MTU) size of at least 2180 bytes to accommodate the FC payload and FCoE encapsulation. (Sometimes the MTU is rounded up to 2500 bytes. If larger frames are expected on the interface, set the MTU size accordingly.)

This example shows a simple configuration to illustrate the basic steps for creating:

- The FCoE-device-facing VLAN and its 10-Gigabit Ethernet interfaces
- The VLAN interface
- The FC-switch-facing native FC interfaces
- One FC fabric on the FCoE-FC gateway

Configuring these elements results in traffic being routed between the VLAN and FC interfaces, thus connecting the FCoE devices to the FC switch through the FCoE-FC gateway.

A VLAN called **blue** transports FCoE traffic between FCoE devices and the FCoE-FC gateway using an FCoE VLAN interface called **vlan.100**. The FCoE-FC gateway's **vlan.100** interface presents an F_Port interface to the FCoE devices on the VLAN. For each FCoE device ENode that logs in to the FCoE-FC gateway, the gateway instantiates a virtual F_Port (VF_Port) interface. This creates a virtual link between the ENode VN_Port and the FCoE-FC gateway. The FCoE-FC gateway's native FC interfaces transport FC traffic between the gateway and the FC switch.

Configuring both the FCoE VLAN interface and the native FC interfaces as part of a gateway fabric associates them in the switch and makes the connection between the FCoE servers and the FC switch.

## Topology

The topology for this example consists of one QFX3500 switch with FC-capable ports to connect to the FC switch and with Ethernet ports in tagged-access mode to connect to the FCoE devices. Table 1 on page 19 and Figure 1 on page 20 show the configuration components of this example.

**Table 1: Components of the Fibre Channel Interface Configuration Topology**

| Property | Settings |
|---|---|
| Switch hardware | QFX3500 switch in gateway mode |
| FCoE VLAN name and tag ID | **blue**, tag **100**<br><br>IGMP snooping disabled on the FCoE VLAN. |
| Interfaces in VLAN **blue** | Interfaces: **xe-0/0/6**, **xe-0/0/7**, **xe-0/0/8**, **xe-0/0/9**, **xe-0/0/10**, **xe-0/0/11**<br>Port mode: **tagged-access**<br>MTU: **2180**<br>Native VLAN: **1**<br><br>NOTE: You can bundle two or more of the VLAN interfaces in a link aggregation group (LAG) if you wish.<br><br>NOTE: FCoE VLANs (any VLAN that carries FCoE traffic) support only Spanning Tree Protocol (STP) and link aggregation group (LAG) Layer 2 features.<br><br>NOTE: This example disables storm control on all interfaces. In your configuration, explicitly enable storm control on interfaces on which you want to use storm control. Configuring storm control on an Ethernet interface and including that interface in an FCoE-FC gateway may have undesirable effects, including FCoE packet loss. |
| FCoE VLAN interface | **vlan.100**<br>Port mode: **f-port** |
| Native Fibre Channel interfaces | Interfaces: **fc-0/0/0**, **fc-0/0/1**, **fc-0/0/2**, **fc-0/0/3**, **fc-0/0/4**, **fc-0/0/5**<br>Port mode: **np-port**<br>Speed: **4 Gbps** |
| Fibre Channel fabric **fcproxy1** | Fabric type: **proxy**<br>Fabric ID: **1**<br>FC interfaces: **fc-0/0/0**, **fc-0/0/1**, **fc-0/0/2**, **fc-0/0/3**, **fc-0/0/4**, **fc-0/0/5** |

Figure 1: Fibre Channel Interface Configuration Topology

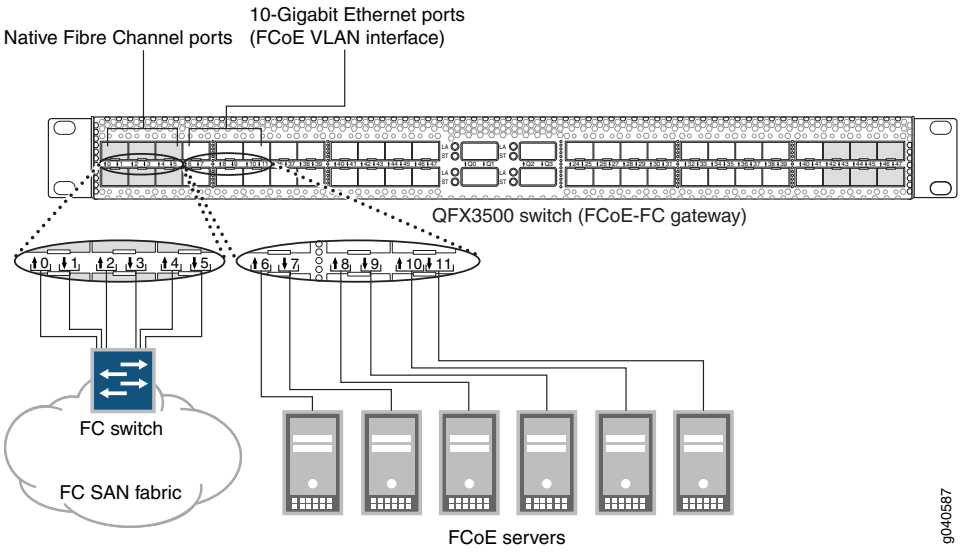


This configuration example creates a VLAN for FCoE traffic and routes its traffic to an FCoE VLAN interface that is part of the FC fabric. It also creates the FC interfaces needed to connect to the FC switch.

To set up FC interfaces and FCoE VLAN interfaces:

- Configure a VLAN to use as a dedicated FCoE VLAN:
  - Configure the interfaces the FCoE VLAN uses as Ethernet switching interfaces in tagged-access port mode.
  - Disable storm control on the interfaces.
  - Configure the interfaces the FCoE VLAN uses with the native VLAN.
  - Configure the FCoE VLAN to use the desired Ethernet interfaces.
  - Disable IGMP snooping on the FCoE VLAN. (IGMP snooping is enabled by default on all VLANs, but is not supported on FCoE VLANs).
- Configure the FCoE VLAN interface.
- Define the interface for the FCoE VLAN (associate the VLAN with the FCoE VLAN interface).
- Configure the physical FC interfaces (either one or two 6-port blocks) that connect to the FC switch.
- Configure the logical FC interfaces that connect to the FC switch.
- Configure the FCoE-FC gateway fabric:
  - Configure the fabric ID.
  - Configure the fabric as a proxy fabric.
  - Add the FCoE VLAN interface and the native FC interfaces to the fabric.

To keep the example simple, the configuration steps show six Ethernet interfaces in the FCoE VLAN and six native FC interfaces in the FC fabric. Use the same configuration procedure to add more interfaces to the FCoE VLAN or to the FC fabric.

## Configuration

**CLI Quick Configuration**

To quickly configure FCoE and native FC interfaces on an FCoE-FC gateway and route traffic between the FCoE VLAN and FC interfaces, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans blue vlan-id 100
set vlans native vlan-id 1
set interfaces xe-0/0/6 unit 0 family ethernet-switching port-mode tagged-access vlan
  members blue
set interfaces xe-0/0/7 unit 0 family ethernet-switching port-mode tagged-access vlan
  members blue
set interfaces xe-0/0/8 unit 0 family ethernet-switching port-mode tagged-access vlan
  members blue
set interfaces xe-0/0/9 unit 0 family ethernet-switching port-mode tagged-access vlan
  members blue
set interfaces xe-0/0/10 unit 0 family ethernet-switching port-mode tagged-access vlan
  members blue
set interfaces xe-0/0/11 unit 0 family ethernet-switching port-mode tagged-access vlan
  members blue
set interfaces xe-0/0/6 unit 0 family ethernet-switching native-vlan-id 1
set interfaces xe-0/0/7 unit 0 family ethernet-switching native-vlan-id 1
set interfaces xe-0/0/8 unit 0 family ethernet-switching native-vlan-id 1
set interfaces xe-0/0/9 unit 0 family ethernet-switching native-vlan-id 1
set interfaces xe-0/0/10 unit 0 family ethernet-switching native-vlan-id 1
set interfaces xe-0/0/11 unit 0 family ethernet-switching native-vlan-id 1
set interfaces xe-0/0/6 mtu 2180
set interfaces xe-0/0/7 mtu 2180
set interfaces xe-0/0/8 mtu 2180
set interfaces xe-0/0/9 mtu 2180
set interfaces xe-0/0/10 mtu 2180
set interfaces xe-0/0/11 mtu 2180 delete ethernet-switching-options storm-control
  interface all
set vlans blue interface xe-0/0/6.0
set vlans blue interface xe-0/0/7.0
set vlans blue interface xe-0/0/8.0
set vlans blue interface xe-0/0/9.0
set vlans blue interface xe-0/0/10.0
set vlans blue interface xe-0/0/11.0
set protocols igmp-snooping vlan blue disable
set interfaces vlan unit 100 family fibre-channel port-mode f-port
set vlans blue l3-interface vlan.100
set chassis fpc 0 pic 0 fibre-channel port-range 0 5
set interfaces fc-0/0/0 unit 0 family fibre-channel port-mode np-port
set interfaces fc-0/0/1 unit 0 family fibre-channel port-mode np-port
set interfaces fc-0/0/2 unit 0 family fibre-channel port-mode np-port
set interfaces fc-0/0/3 unit 0 family fibre-channel port-mode np-port
set interfaces fc-0/0/4 unit 0 family fibre-channel port-mode np-port
set interfaces fc-0/0/5 unit 0 family fibre-channel port-mode np-port
set interfaces fc-0/0/0 fibrechannel-options speed 4g
```

```
set interfaces fc-0/0/1 fibrechannel-options speed 4g
set interfaces fc-0/0/2 fibrechannel-options speed 4g
set interfaces fc-0/0/3 fibrechannel-options speed 4g
set interfaces fc-0/0/4 fibrechannel-options speed 4g
set interfaces fc-0/0/5 fibrechannel-options speed 4g
set fc-fabrics fcproxy1 fabric-id 1
set fc-fabrics fcproxy1 fabric-type proxy
set fc-fabrics fcproxy1 interface vlan.100
set fc-fabrics fcproxy1 interface fc-0/0/0.0
set fc-fabrics fcproxy1 interface fc-0/0/1.0
set fc-fabrics fcproxy1 interface fc-0/0/2.0
set fc-fabrics fcproxy1 interface fc-0/0/3.0
set fc-fabrics fcproxy1 interface fc-0/0/4.0
set fc-fabrics fcproxy1 interface fc-0/0/5.0
```

**Step-by-Step Procedure**

Configure FCoE and FC interfaces in an FCoE-FC gateway FC fabric and set up traffic routing between the FCoE VLAN and FC interfaces:

1. Configure the VLAN for FCoE traffic:

   ```
   [edit vlans]
   user@switch# set blue vlan-id 100
   ```

2. Configure the native VLAN:

   ```
   [edit vlans]
   user@switch# set native vlan-id 1
   ```

3. Configure the Ethernet interfaces for the FCoE VLAN in tagged-access mode and as members of the FCoE VLAN (VLAN **blue**):

   ```
   [edit interfaces]
   user@switch# set xe-0/0/6 unit 0 family ethernet-switching port-mode
     tagged-access vlan members blue
   user@switch# set xe-0/0/7 unit 0 family ethernet-switching port-mode
     tagged-access vlan members blue
   user@switch# set xe-0/0/8 unit 0 family ethernet-switching port-mode
     tagged-access vlan members blue
   user@switch# set xe-0/0/9 unit 0 family ethernet-switching port-mode
     tagged-access vlan members blue
   user@switch# set xe-0/0/10 unit 0 family ethernet-switching port-mode
     tagged-access vlan members blue
   user@switch# set xe-0/0/11 unit 0 family ethernet-switching port-mode
     tagged-access vlan members blue
   ```

4. Configure the native VLAN on the Ethernet interfaces in the FCoE VLAN:

   ```
   [edit interfaces]
   user@switch# set xe-0/0/6 unit 0 family ethernet-switching native-vlan-id 1
   user@switch# set xe-0/0/7 unit 0 family ethernet-switching native-vlan-id 1
   user@switch# set xe-0/0/8 unit 0 family ethernet-switching native-vlan-id 1
   user@switch# set xe-0/0/9 unit 0 family ethernet-switching native-vlan-id 1
   user@switch# set xe-0/0/10 unit 0 family ethernet-switching native-vlan-id 1
   user@switch# set xe-0/0/11 unit 0 family ethernet-switching native-vlan-id 1
   ```

5. Set the MTU to 2180 for each Ethernet interface:

   ```
   [edit interfaces]
   user@switch# set xe-0/0/6 mtu 2180
   ```

```
user@switch# set xe-0/0/7 mtu 2180
user@switch# set xe-0/0/8 mtu 2180
user@switch# set xe-0/0/9 mtu 2180
user@switch# set xe-0/0/10 mtu 2180
user@switch# set xe-0/0/11 mtu 2180
```

6. Disable storm control on all interfaces (afterward, be sure to enable storm control on any interfaces on which you want to use storm control):

```
user@switch# delete ethernet-switching-options storm-control interface all
```

7. Assign the Ethernet interfaces to the FCoE VLAN:

```
[edit vlans blue interface]
user@switch# set xe-0/0/6.0
user@switch# set xe-0/0/7.0
user@switch# set xe-0/0/8.0
user@switch# set xe-0/0/9.0
user@switch# set xe-0/0/10.0
user@switch# set xe-0/0/11.0
```

8. Disable IGMP snooping on the FCoE VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan blue disable
```

9. Configure the FCoE VLAN interface and port mode for the FCoE traffic:

```
[edit interfaces]
user@switch# set vlan unit 100 family fibre-channel port-mode f-port
```

10. Define the FCoE VLAN interface as the interface for the FCoE VLAN:

```
[edit vlans]
user@switch# set blue l3-interface vlan.100
```

11. Configure the physical FC interfaces the fabric uses to connect to the FC switch:

```
[edit chassis fpc 0 pic 0]
user@switch# set fibre-channel port-range 0 5
```

NOTE: When you configure ports as FC ports, the port designation changes from xe-n/n/n.n format to fc-n/n/n.n format to indicate that the interface is an FC interface. FC interfaces do not support 10-Gbps interface speed but instead conform to FC interface speeds of 2 Gbps, 4 Gbps, or 8 Gbps.

12. Configure the native FC interfaces and port mode:

```
[edit interfaces]
user@switch# set fc-0/0/0 unit 0 family fibre-channel port-mode np-port
user@switch# set fc-0/0/1 unit 0 family fibre-channel port-mode np-port
user@switch# set fc-0/0/2 unit 0 family fibre-channel port-mode np-port
user@switch# set fc-0/0/3 unit 0 family fibre-channel port-mode np-port
user@switch# set fc-0/0/4 unit 0 family fibre-channel port-mode np-port
user@switch# set fc-0/0/5 unit 0 family fibre-channel port-mode np-port
```

13. Configure the native FC interface port speed:

```
[edit interfaces]
```

```
user@switch# set fc-0/0/0 fibrechannel-options speed 4g
user@switch# set fc-0/0/1 fibrechannel-options speed 4g
user@switch# set fc-0/0/2 fibrechannel-options speed 4g
user@switch# set fc-0/0/3 fibrechannel-options speed 4g
user@switch# set fc-0/0/4 fibrechannel-options speed 4g
user@switch# set fc-0/0/5 fibrechannel-options speed 4g
```

14. Configure the FC fabric name and unique ID:

```
[edit fc-fabrics]
user@switch# set fcproxy1 fabric-id 1
```

15. Define the FC fabric as an FCoE-FC gateway:

```
[edit fc-fabrics]
user@switch# set fcproxy1 fabric-type proxy
```

16. Assign the FCoE VLAN interface to the fabric:

```
[edit fc-fabrics]
user@switch# set fcproxy1 interface vlan.100
```

17. Assign the native FC interfaces to the fabric:

```
[edit fc-fabrics]
user@switch# set fcproxy1 interface fc-0/0/0.0
user@switch# set fcproxy1 interface fc-0/0/1.0
user@switch# set fcproxy1 interface fc-0/0/2.0
user@switch# set fcproxy1 interface fc-0/0/3.0
user@switch# set fcproxy1 interface fc-0/0/4.0
user@switch# set fcproxy1 interface fc-0/0/5.0
```

## Verification

To verify that the native FC interfaces and FCoE VLAN interface have been created, added to the FC fabric, and are operating properly, perform these tasks:


- Verifying That the Native FC Interfaces and the FCoE VLAN Interface Have Been Created on page 24
- Verifying That the FCoE VLAN Includes the Correct Ethernet Interfaces on page 25
- Verifying That the FC Fabric Includes the Correct Interfaces on page 26
- Verifying Native FC Interface Operation on page 26
- Verifying That IGMP Snooping Has Been Disabled on the FCoE VLAN on page 27


### *Verifying That the Native FC Interfaces and the FCoE VLAN Interface Have Been Created*

**Purpose** Verify that the six native FC interfaces and the FCoE VLAN interface have been created on the switch and are configured in the correct mode.

**Action** List all of the FC interfaces configured on the switch using the **show fibre-channel interfaces** command:

```
user@switch> show fibre-channel interfaces
                          Native              Config  Oper
Interface     Idx   Type  Fabric-id     NPIV  Mode    Mode    State
fc-0/0/0.0    70    FC    1             YES   NP      NP      up
fc-0/0/1.0    71    FC    1             YES   NP      NP      up
```

```
fc-0/0/2.0    72    FC     1            YES   NP      NP      up
fc-0/0/3.0    73    FC     1            YES   NP      NP      up
fc-0/0/4.0    74    FC     1            YES   NP      NP      up
fc-0/0/5.0    75    FC     1            YES   NP      NP      up
vlan.100      67    FCOE   1            YES   F       F       up
```

**Meaning** The **show fibre-channel interfaces** command lists all native FC interfaces and FCoE VLAN interfaces configured on the switch. The command output shows that the FC interfaces **fc-0/0/0.0**, **fc-0/0/1.0**, **fc-0/0/2.0**, **fc-0/0/3.0**, **fc-0/0/4.0**, and **fc-0/0/5.0** have been created and that those six interfaces:

- Are native Fibre Channel interfaces (type **FC**).
- Belong to the FC fabric with a configured fabric ID of **1**.
- Are capable of N_Port ID virtualization (NPIV).
- Have a configured mode and an operational mode of proxy N_Port (**NP**), which means that they should be connected to an FCF or an FC switch, not to an FCoE device, and that they carry native FC traffic.
- Show an operational state of **up**.

The command output also shows that the FCoE VLAN interface **vlan.100** has been created and that interface:

- Is an FCoE VLAN interface (type **FCOE)**.
- Belongs to the FC fabric with a configured fabric ID of **1**.
- Is capable of N_Port ID virtualization (NPIV).
- Has a configured mode and an operational mode of F_Port (**F**), which means that its interfaces connect to FCoE devices and carry FCoE traffic.
- Shows an operational state of **up**.

### *Verifying That the FCoE VLAN Includes the Correct Ethernet Interfaces*

**Purpose** Verify that the FCoE VLAN **blue** has been created with the correct VLAN tag (**100**) and with the correct Ethernet interfaces.

**Action** List all of the interfaces configured on the switch in VLAN **blue** using the **show vlans** command:

```
user@switch> show vlans blue
Name            Tag     Interfaces
blue            100
                      xe-0/0/6.0, xe-0/0/7.0, xe-0/0/8.0, xe-0/0/9.0, xe-0/0/10.0

                        xe-0/0/11.0
```

**Meaning** The **show vlans blue** command lists the interfaces that are members of the FCoE VLAN **blue**. The command output shows that the **blue** VLAN has a tag ID of 100 and includes the interfaces **xe-0/0/6.0**, **xe-0/0/7.0**, **xe-0/0/8.0**, **xe-0/0/9.0**, **xe-0/0/10.0**, and **xe-0/0/11.0**.

*Verifying That the FC Fabric Includes the Correct Interfaces*

**Purpose** Verify that the FC fabric configuration is configured on the switch with the correct native FC and FCoE VLAN interfaces.

**Action** List all of the interfaces configured on FC fabrics on the switch using the **show fibre-channel fabric** command:

```
user@switch> show fibre-channel fabric
Name            Fabric-id      Type      Interfaces
fcproxy1        1              PROXY
                                         fc-0/0/0.0
                                         fc-0/0/1.0
                                         fc-0/0/2.0
                                         fc-0/0/3.0
                                         fc-0/0/4.0
                                         fc-0/0/5.0
                                         vlan.100
```

**Meaning** The **show fibre-channel fabric** command lists the interfaces that are members of each FC fabric. The command output shows that the only fabric configured on the switch is named **fcproxy1**, has a fabric-id of **1**, and is a **proxy** fabric in an FCoE-FC gateway. The command output also shows that the native FC interfaces **fc-0/0/0.0**, **fc-0/0/1.0**, **fc-0/0/2.0**, **fc-0/0/3.0**, **fc-0/0/4.0**, and **fc-0/0/5.0**, and the FCoE VLAN interface **vlan.100** belong to **fcproxy1**.

*Verifying Native FC Interface Operation*

**Purpose** Verify that the native FC interfaces are online and display the number of FC sessions on each interface.

**Action** List all of the native FC NP_Port interface states and sessions by FC fabric using the **show fibre-channel proxy np-port** command:

```
user@switch> show fibre-channel proxy np-port
Fabric: fcproxy1, Fabric-id: 1
NP-Port     State              Sessions       LB state   LB weight
fc-0/0/0.0  online             3              ON         4
fc-0/0/1.0  online             3              ON         4
fc-0/0/2.0  online             2              ON         4
fc-0/0/3.0  online             2              ON         4
fc-0/0/4.0  online             2              ON         4
fc-0/0/5.0  online             2              ON         4
```

**Meaning** The **show fibre-channel proxy np-port** command lists the interfaces that are configured as native FC proxy N_Port interfaces. The command output shows:

- The fabric name is **fcproxy1** and its fabric ID is **1**.
- The interfaces are **online**.
- The number of FC sessions (virtual links) running on each interface.
- The load-balancing (LB) state is **ON** for all of the interfaces.
- The LB weight reflects the port speed of each interface, which is **4** Gbps.

*Verifying That IGMP Snooping Has Been Disabled on the FCoE VLAN*

**Purpose** Verify that IGMP snooping is disabled on the FCoE VLAN.

**Action** List the IGMP snooping protocol information for the FCoE VLAN using the **show configuration protocols igmp-snooping vlan blue** command:

```
user@switch> show configuration protocols igmp-snooping vlan blue
    disable;
```

**Meaning** The **show configuration protocols igmp-snooping vlan blue** command lists the IGMP snooping configuration for the FCoE VLAN. The command output shows that IGMP snooping is disabled on the FCoE VLAN.

### Results

Display the results of the configuration:

```
user@switch> show configuration
fc-0/0/0 {
  fibrechannel-options {
    speed 4g;
  }
  unit 0 {
    family fibre-channel {
      port-mode np-port;
    }
  }
}
fc-0/0/1 {
  fibrechannel-options {
    speed 4g;
  }
  unit 0 {
    family fibre-channel {
      port-mode np-port;
    }
  }
}
fc-0/0/2 {
  fibrechannel-options {
    speed 4g;
  }
  unit 0 {
    family fibre-channel {
      port-mode np-port;
    }
  }
}
fc-0/0/3 {
  fibrechannel-options {
    speed 4g;
  }
  unit 0 {
```

```
      family fibre-channel {
        port-mode np-port;
      }
    }
  }
  fc-0/0/4 {
    fibrechannel-options {
      speed 4g;
    }
    unit 0 {
      family fibre-channel {
        port-mode np-port;
      }
    }
  }
  fc-0/0/5 {
    fibrechannel-options {
      speed 4g;
    }
    unit 0 {
      family fibre-channel {
        port-mode np-port;
      }
    }
  }
  xe-0/0/6 {
    mtu 2180;
    unit 0 {
      family ethernet-switching {
        port-mode tagged-access;
        vlan {
          members blue;
        }
        native-vlan-id 1;
      }
    }
  }
  xe-0/0/7 {
    mtu 2180;
    unit 0 {
      family ethernet-switching {
        port-mode tagged-access;
        vlan {
          members blue;
        }
        native-vlan-id 1;
      }
    }
  }
  xe-0/0/8 {
    mtu 2180;
    unit 0 {
      family ethernet-switching {
        port-mode tagged-access;
        vlan {
          members blue;
```

```
        }
        native-vlan-id 1;
      }
    }
  }
  xe-0/0/9 {
    mtu 2180;
    unit 0 {
      family ethernet-switching {
        port-mode tagged-access;
        vlan {
          members blue;
        }
        native-vlan-id 1;
      }
    }
  }
  xe-0/0/10 {
    mtu 2180;
    unit 0 {
      family ethernet-switching {
        port-mode tagged-access;
        vlan {
          members blue;
        }
        native-vlan-id 1;
      }
    }
  }
  xe-0/0/11 {
    mtu 2180;
    unit 0 {
      family ethernet-switching {
        port-mode tagged-access;
        vlan {
          members blue;
        }
        native-vlan-id 1;
      }
    }
  }
  vlan {
    unit 100 {
      family fibre-channel {
        port-mode f-port;
      }
    }
  }
  fc-fabrics {
    fcproxy1 {
      fabric-id 1
      fabric-type proxy
      interface {
        vlan.100
        fc-0/0/0.0;
        fc-0/0/1.0;
```

```
        fc-0/0/2.0;
        fc-0/0/3.0;
        fc-0/0/4.0;
        fc-0/0/5.0;
      }
    }
}
protocols {
  igmp-snooping {
    vlan blue {
      disable;
    }
  }
}
vlans {
  blue {
    vlan-id 100
    interface {
      xe-0/0/6.0;
      xe-0/0/7.0;
      xe-0/0/8.0;
      xe-0/0/9.0;
      xe-0/0/10.0;
      xe-0/0/11.0;
    }
    l3-interface vlan.100
  }
  native {
    vlan-id 1;
  }
}
```

TIP: To quickly configure the interfaces, issue the load merge terminal command and then copy the hierarchy and paste it into the switch terminal window.

**Related Documentation**

- Understanding FCoE-FC Gateway Benefits on page 5
- Understanding Interfaces on an FCoE-FC Gateway on page 6