



Configuring FC NPV

- [Supported Hardware, on page 1](#)
- [FC NPV Overview, on page 1](#)
- [FC NPV Mode, on page 2](#)
- [Server Interfaces, on page 2](#)
- [NP Uplinks, on page 3](#)
- [SAN Port Channels, on page 6](#)
- [FLOGI Operation, on page 11](#)
- [NPV Traffic Management, on page 11](#)
- [FC NPV Traffic Management Guidelines, on page 12](#)
- [FC NPV Guidelines and Limitations, on page 13](#)
- [Licensing Requirements for FC NPV, on page 15](#)
- [Configuring NPV, on page 16](#)
- [Verifying FC NPV, on page 21](#)
- [FC NPV Core Switch and FC NPV Edge Switch Configuration Example, on page 24](#)

Supported Hardware

FC NPV is supported on N9K-C93180YC-FX, N9K-C9336C-FX2-E, and N9K-C93360YC-FX2 switches.

N9K-C93180YC-FX and N9K-C93360YC-FX2 supports only the following FC SFPs:

- DS-SFP-FC8G-SW
- DS-SFP-FC16G-SW
- DS-SFP-FC32G-SW

N9K-C9336C-FX2-E supports only the following FC SFP:

- DS-SFP-4x32G-SW


FC NPV Overview

A switch is in NPV mode after enabling NPV. NPV mode applies to an entire switch. All end devices connected to a switch that are in NPV mode must log in as an N port to use this feature (loop-attached devices are not

supported). All links from the edge switches (in NPV mode) to the NPV core switches are established as NP ports (not E ports), which are used for typical inter-switch links. NPIV is used by the switches in NPV mode to log in to multiple end devices that share a link to the NPV core switch.

The following figure shows an interface-level view of an FC NPV configuration.

Figure 1: FC NPV Interface Configuration



FC NPV Benefits

FC NPV provides the following:

- Increased number of hosts that connect to the fabric without adding domain IDs in the fabric
- Connection of FC and FCoE hosts and targets to SAN fabrics using FC interfaces
- Automatic traffic mapping
- Static traffic mapping

FC NPV Mode

In FC NPV mode, the edge switch relays all traffic to the core switch and shares the domain ID of the core switch.

FC NPV is enabled by installing and enabling **feature-set fcoe-npv**. You cannot configure FC NPV mode on a per-interface basis. FC NPV mode applies to the entire switch.

Server Interfaces

- In Cisco Nexus 9000 Series switches, server interfaces can be FC or vFC interfaces.
- Server interfaces are F ports on the edge switch that connect to the servers. A server interface may support multiple end devices by enabling the N port identifier virtualization (NPIV) feature. NPIV provides a means to assign multiple FC IDs to a single N port, which allows the server to assign unique FC IDs to different applications.



Note To use NPIV, enable the NPIV feature and reinitialize the server interfaces that will support multiple devices.

- FC server interfaces should be in trunk mode off. Trunk mode on is not supported.
- vFC server interfaces should be in trunk mode on.
- Server interfaces are automatically distributed among the NP uplinks to the core switch. All of the end devices connected to a server interface are mapped to the same NP uplink.
- When you connect a 16G host adapter to a 32G SFP port on Cisco Nexus 93360YC-FX and/or 93360YC-FX2 switches, the link may not come up when the speed is configured as auto speed. Or sometimes, it defaults to 8G speed. Then, to use 16G speed, you must manually configure the port using the command **switchport speed 16000**.
- 8G speed is not supported for server and target interfaces.

NP Uplinks

- In Cisco Nexus 9000 Series switches, NP uplink interfaces can be native Fibre Channel interfaces, virtual fiber channel interfaces, SAN port channel interfaces, or virtual ethernet port-channel interfaces.
- All interfaces from the edge switch to the core switch are configured as proxy N ports (NP ports).
- An NP uplink is a connection from an NP port on the edge switch to an F port on the core switch. When an NP uplink is established, the edge switch sends a fabric login message (FLOGI) to the core switch then (if the FLOGI is successful) registers itself with the name server on the core switch. Subsequent FLOGIs from end devices connected to this NP uplink are forwarded as-is to the core switch. Subsequent FLOGIs from the same VSAN are forwarded as fdisc.



Note In the switch CLI configuration commands and output displays, NP uplinks are called External Interfaces.

- The default speed of NP links is set to auto.
- The features below must be enabled on the core switch:
 - **feature npiv**
 - **feature fport-channel-trunk**
- If the FC uplink speed is 8G, the fill pattern should be configured as IDLE on the core switch.



Note Following is an example of configuring IDLE fill pattern on a Cisco MDS switch:

```
Switch(config)# int fc2/3
Switch(config)# switchport fill-pattern IDLE speed 8000
Switch(config)# sh run int fc2/3
```

```
interface fc2/3
switchport speed 8000
switchport mode NP
switchport fill-pattern IDLE speed 8000
no shutdown
```

**Note**

- To enable trunking and for the FLOGI from NP uplink of Cisco Nexus 9000 Series switches to be successful on the core switch, both the core and the Cisco Nexus 9000 Series switch should be configured with the OUI of each other.

Configure the OUI on the core and Nexus 9000 switch *only* if the OUI value is not registered by default on either of them.

The OUI is found and configured as follows:

```
N9K(config-if)# show wwn switch
Switch WWN is 20:00:2c:d0:2d:50:ea:64
N9K(config-if)#
```

On the core, we see the output below if the OUI (0x2cd02d) is already registered.

```
MDS9710(config-if)# sh wwn oui | i 2cd02d
0x2cd02d Cisco Default
MDS9710(config-if) #
```

If the OUI is not registered with the core, configure it manually.

```
MDS9710(config-if)# wwn oui 0x2cd02d
```

- Beginning with Cisco NX-OS Release 7.3(0)D1(1), the OUI is configurable on a Cisco MDS 9700 Series core switches.
- If the uplinks to core switch are FCoE enabled, then the FKA advertisement period is taken from configured value on the core switch. If the uplinks to core switch are FC enabled, FKA a period is taken from configured value on local NPV switch.

**Note**

The following example shows the FCoE uplink. Because the switch has an FCoE link, the value is taken from FCF:

```
switch(config)# sh run fcoe_mgr | i i fka
fcoe fka-adv-period 12

switch(config)# sh fcoe

FCF details for interface vfc-pol142
FCF-MAC is 54:7f:ee:ec:71:84
FC-MAP is 0e:fc:00
FCF Priority is 128
FKA Advertisement period for FCF is 8 seconds <<<<<
```

The following example shows the FC uplink:

```
switch(config)# sh run | i i fka
fcoe fka-adv-period 10

switch(config)# sh fcoe
FCF details for interface san-port-channel29
FCF-MAC is 2c:d0:2d:50:e4:29
FC-MAP is 0e:fc:00
FCF Priority is 129
FKA Advertisement period for FCF is 10 seconds
```

SAN Port Channels

About SAN Port Channels

- A SAN port channel is a logical interface that combines a set of FC interfaces connected to the same fibre channel node and operates as one link.
- SAN port channels support bandwidth utilization and availability.
- SAN port channels on Cisco Nexus 9000 Series switches are mainly used to connect to MDS cores and to provide optimal bandwidth utilization and transparent failover between the uplinks of a VSAN.

Configuring SAN Port Channels

When you configure a SAN port channel, it gets created with default values. You can modify all the default values, except the channel mode. You must connect each switch to same number of interfaces on either side of a SAN port channel. Otherwise, you see a SAN port channel error.

SAN Port Channel Guidelines and Limitations

- The number of SAN port channels and vFC port channels, together, can be only 8 on the Cisco Nexus 9000 Series switch.
- The maximum number of FC interfaces that can be combined into a SAN port channel is limited to 16.
- The default channel mode on Cisco Nexus 9000 Series switches for SAN port channels is **active**; this cannot be changed.
- Operating speed and member addition to san-po limitation on N9K-C9336C-FX2-E are available in [FC NPV Guidelines and Limitations, on page 13](#).

Creating a SAN Port Channel

This section explains how to create a SAN port channel.

Step 1 switch# **configure terminal**

Enters the global configuration mode.

Step 2 switch(config)# **interface san-port-channel** *channel-number*

Creates the specified SAN port channel using the default mode (on). The SAN port channel number is in the range of 1 to 256

The following example shows the SAN port channel creation:

```
switch(config)# interface san-port-channel 1
switch(config-if)#
```

About SAN Port Channel Modes

A SAN port channel is configured with channel mode active by default. When active, the member ports initiate port-channel-protocol negotiation with the peer port(s) regardless of the channel-group mode of the peer port. If the peer port, while configured in a channel group, does not support the port-channel protocol, or responds with a nonnegotiable status, the port channel will be disabled. The active port-channel mode allows automatic recovery without explicitly enabling and disabling the port-channel-member ports at either end.

About Deleting SAN Port Channels

When you delete the SAN port channel, the corresponding channel membership is also deleted.

If you delete the SAN port channel for one port, then the individual ports within the deleted SAN port channel retain the compatibility parameter settings (speed, mode, port VSAN, allowed VSAN, and port security). You can explicitly change those settings as required.

Deleting SAN Port Channels

This section explains how to delete a SAN port channel.

Step 1 switch# **configure terminal**

Enters global configuration mode.

Step 2 switch(config)# **no interface san-port-channel** *channel-number*

Deletes the specified port channel, its associated interface mappings, and the hardware associations for this SAN port channel.

Example

The following example demonstrates how to delete a SAN port channel:

```
switch(config)# no interface san-port-channel 1
```

The SAN port channel 1 is deleted and all its members are disabled. Please do the same operation on the switch at the other end of the SAN port channel.

Interfaces in a SAN Port Channel

You can add or remove a physical Fibre Channel interface (or a range of interfaces) to an existing SAN port channel. The compatible parameters on the configuration are mapped to the SAN port channel. Adding an interface to a SAN port channel increases the channel size and bandwidth of the SAN port channel. Removing an interface from a SAN port channel decreases the channel size and bandwidth of the SAN port channel.



Note Virtual Fibre Channel interfaces cannot be added to SAN port channels.

Adding an Interface to a SAN Port Channel

This section explains how to add an interface to a SAN port channel.

Step 1

switch# **configure terminal**

Enters global configuration mode.

Step 2

switch(config)# **interface** *type slot / port / BO port*

Enters configuration mode for the specified interface.

Step 3

switch(config-if)# **channel-group** *channel-number*

Adds the Fibre Channel interface to the specified channel group. If the channel group does not exist, it is created. The port is shut down.

Break out (BO) port option for Fibre Channel (FC) interfaces is required only for the Cisco Nexus N9K-C9336C-FX2-E platform switch.

Example

The following example adds an interface to a SAN port channel:

```
switch(config)# interface fc9/10
switch(config-if)# channel-group 15
```

fc9/10 is added to san-port-channel 15 and is disabled. Please do the same operation on the switch at the other end of the san-port-channel, then do “no shutdown” at both ends to bring them up

FC support in N9K-9336C-FX2-E switch is available only from port 9 to 36.

Forcing an Interface Addition

You can force the port configuration to be overwritten by the SAN port channel. In this case, the interface is added to a SAN port channel.



Note When SAN port channels are created from within an interface, the **force** option cannot be used.

This section explains how to force the addition of a port to a SAN port channel.

Step 1

switch# **configure terminal**

Enters global configuration mode.

Step 2

switch(config)# **interface** *type slot / port / BO port*

Enters configuration mode for the specified interface.

Step 3

switch(config-if)# **channel-group** *channel-number* **force**

Forces the addition of the interface into the specified channel group. The E port is shut down.

Example

The following example adds an interface to a SAN port channel:

```
switch(config)# interface fc9/10
switch(config-if)# channel-group 15 force
```

fc9/10 added to san-port-channel 15 and disabled. Please do the same operation on the switch at the other end of the san-port-channel, then perform a **no shutdown** at both ends to bring them up

FC support in N9K-9336C-FX2-E switch is available only from port 9 to 36.

About Interface Deletion from a SAN Port Channel

When a physical interface is deleted from the SAN port channel, the channel membership is automatically updated. If the deleted interface is the last operational interface, then the port channel status is changed to a down state. Deleting an interface from a SAN port channel decreases the channel size and bandwidth of the SAN port channel.

Deleting an Interface from a SAN Port Channel

This section explains how to delete a physical interface (or a range of physical interfaces) from a SAN port channel.

- Step 1** `switch(config)# interface type slot /port / BO port`
Enters configuration mode for the specified interface.
- Step 2** `switch(config)#shut`
Shuts down the interface before removing the specified channel group.
- Step 3** `switch(config)#no channel-group channel-number`
Deletes the physical Fibre Channel interface from the specified channel group.
Break out (BO) port option for Fibre Channel (FC) interfaces is required only for the Cisco Nexus N9K-C9336C-FX2-E platform switch.
-

Example

The following example deletes an interface from a SAN port channel:

```
switch(config)# interface fc9/10
switch(config-if)# shut
switch(config-if)# no channel-group 15
```

fc9/10 is removed from the SAN port-channel 15 and disabled.

Please do the same operation on the switch at the other end of the san-port-channel
FC support in N9K-9336C-FX2-E switch is available only from port 9 to 36.

Verifying SAN Port Channel Configurations

You can view specific information about existing SAN port channels at any time from EXEC mode. The following **show** commands provide further details on existing SAN port channels.

The **show san-port-channel summary** command displays a summary of SAN port channels within the switch. A one-line summary of each SAN port channel provides the administrative state, the operational state, the number of attached and active interfaces (up), and the first operational port (FOP), which is the primary operational interface selected in the SAN port channel to carry control-plane traffic (no load-balancing). The FOP is the first port that comes up in a SAN port channel and can change if the port goes down. The FOP is also identified by an asterisk (*).

To display VSAN configuration information, perform one of the following tasks:

-
- Step 1** switch# **show san-port-channel summary** | **database** | **consistency** [*details*] | **usage** | **compatibility-parameters**
Displays SAN port channel information.
- Step 2** switch# **show san-port-channel database interface san-port-channel** *channel-number*
Displays information for the specified SAN port channel.
- Step 3** switch# **show interface type slot / port / BO port**
Displays VSAN configuration information for the specified Fibre Channel interface.
Break out (BO) port option for Fibre Channel (FC) interfaces is required only for the Cisco Nexus N9K-C9336C-FX2-E platform switch.
-

Example

The following example shows how to display a summary of SAN port channel information:

```
switch# show san-port-channel summary
-----
Interface      Total Ports Oper Ports First Oper Port-
-----
san-port-channel    7         2         0         -
san-port-channel    8         2         0         -
san-port-channel    9         2         2
```

The following example shows how to display SAN port channel consistency:

```
switch# show san-port-channel consistency
Database is consistent
```

The following example shows how to display details of the used and unused port channel numbers:

```
switch# show san-port-channel usage
Totally 3 port-channel numbers used
```

```
=====
Used : 77 - 79
Unused: 1 - 76, 80 - 256
```

FLOGI Operation

When an NP port becomes operational, the switch first logs itself in to the core switch by sending a FLOGI request (using the port WWN of the NP port).

After completing the FLOGI request, the switch registers itself with the fabric name server on the core switch (using the symbolic port name of the NP port and the IP address of the edge switch).

The following table identifies port and node names in the edge switch used in FC NPV mode.

Table 1: Edge Switch FLOGI Parameters

| Parameter | Derived From |
|--------------------|---|
| pWWN | The fWWN of the NP port on the edge switch. |
| nWWN | The VSAN-based sWWN of the edge switch. |
| symbolic port name | The edge switch name and NP port interface string. Note If no switch name is available, the output will read "switch." For example, switch: fc 1/5. |
| IP address | The IP address of the edge switch. |
| symbolic node name | The edge switch name. |

NPV Traffic Management

Automatic Uplink Selection

NPV supports automatic selection of NP uplinks. When a server interface is brought up, the NP uplink interface with the minimum load is selected from the available NP uplinks in the same VSAN as the server interface.

When a new NP uplink interface becomes operational, the existing load is not redistributed automatically to include the newly available uplink. Server interfaces that become operational after the NP uplink can select the new NP uplink.

Traffic Maps

FC NPV supports traffic maps. A traffic map allows you to specify the NP uplinks that a server interface can use to connect to the core switches.



Note When an FC NPV traffic map is configured for a server interface, the server interface must select only from the NP uplinks in its traffic map. If none of the specified NP uplinks are operational, the server remains in a non-operational state.

The FC NPV traffic map feature provides the following benefits:

- Facilitates traffic engineering by allowing configuration of a fixed set of NP uplinks for a specific server interface (or range of server interfaces).
- Ensures correct operation of the persistent FC ID feature; this is because a server interface will always connect to the same NP uplink (or one of a specified set of NP uplinks) after an interface reinitialization or switch reboot.

Disruptive Auto Load Balancing of Server Logins across NP Links

FC NPV supports disruptive load balancing of server logins. When disruptive load balancing is enabled, FC NPV redistributes the server interfaces across all available NP uplinks when a new NP uplink becomes operational. To move a server interface from one NP uplink to another NP uplink, FC NPV forces reinitialization of the server interface so that the server performs a new login to the core switch.

In Release 7.0(3)I7(2) and later software releases, FC NPV supports disruptive load balancing of server logins. When disruptive load balancing is enabled, FC NPV redistributes the server interfaces across all available NP uplinks when a new NP uplink becomes operational. To move a server interface from one NP uplink to another NP uplink, FC NPV forces reinitialization of the server interface so that the server performs a new login to the core switch.

Only server interfaces that are moved to a different uplink are reinitialized. A system message is generated for each server interface that is moved.



Note Redistributing a server interface causes traffic disruption to the attached end devices. Adding a member to the existing port-channel does not trigger disruptive auto load-balance.

To avoid disruption of server traffic, you should enable this feature only after adding a new NP uplink, and then disable it again after the server interfaces have been redistributed.

If disruptive load balancing is not enabled, you can manually reinitialize some or all of the server interfaces to distribute server traffic to new NP uplink interfaces.

FC NPV Traffic Management Guidelines

When deploying FC NPV traffic management, follow these guidelines:

- Use FC NPV traffic management only when automatic traffic engineering does not meet your network requirements.
- You do not need to configure traffic maps for all server interfaces. By default, FC NPV will use automatic traffic management.

- Server interfaces configured to use a set of NP uplink interfaces cannot use any other available NP uplink interfaces, even if none of the configured interfaces are available.
- When disruptive load balancing is enabled, a server interface may be moved from one NP uplink to another NP uplink. Moving between NP uplink interfaces requires FC NPV to relogin to the core switch, causing traffic disruption.
- To link a set of servers to a specific core switch, associate the server interfaces with a set of NP uplink interfaces that all connect to that core switch.
- Configure Persistent FC IDs on the core switch and use the Traffic Map feature to direct server interface traffic onto NP uplinks that all connect to the associated core switch.

FC NPV Guidelines and Limitations

When configuring FC NPV, note the following guidelines and limitations:

- Fibre Channel N-port Virtualization (NPV) can co-exist with VXLAN on different fabric uplinks but on same or different front panel ports on the Cisco Nexus 93180YC-FX, N9K-C9336C-FX2-E, and N9K-C93360YC-FX2 switches. VXLAN can only exist on the Ethernet front panel ports but not on the FC front panel ports.
- In-order data delivery is not required in FC NPV mode because the exchange between two end devices always takes the same uplink from the edge switch to the core. Upstream of the edge switch, core switches will enforce in-order delivery if configured.
- You can configure zoning for end devices that are connected to edge switches using all the available member types on the core switch. However, the preferred way of zoning servers connected to any switch in NPV mode is via PWWN, device-alias and fealias. You must place multiple servers in the same zone only when using smart zoning. For more information about smart zoning on Cisco MDS switches, refer the chapter *Configuring and Managing Zones* in *Cisco MDS 9000 Series Fabric Configuration Guide*.
- Port tracking is not supported in FC NPV mode.
- Port security is supported on the core switch for devices logged in through the FC NPV switch. Port security is enabled on the core switch on a per-interface basis. To enable port security on the core switch for devices that log in through an FC NPV switch, you must adhere to the following requirements:
 - The internal FLOGI must be in the port security database; in this way, the port on the core switch will allow communications and links.
 - All the end device pWWNs must also be in the port security database.
- Edge switches can connect to multiple core switches. In other words, different NP ports can be connected to different core switches.
- If a server interface goes down and then returns to service, the interface is not guaranteed to be assigned to the same NP uplink.
- The server interface is only operational when its assigned NP uplink is operational.
- Both servers and targets can be connected to the switch when in FC NPV mode.
- Fibre Channel switching is not performed in the edge switch; all traffic is switched in the core switch.
- FC NPV supports NPIV-capable servers. This capability is called nested NPIV.

- Connecting two Cisco FC NPV switches together is not supported.
- Only F and NP ports are supported in FC NPV mode.
- **Speed auto-negotiation** is supported only for Cisco Nexus 93180YC-FX and N9k-C93360YC-FX2 switches. The default speed is set to auto.
- Speed auto-negotiation is not supported in N9K-C9336C-FX2-E switch and default speed is set to 32G.
- Nexus 9000 only supports the IDLE fill pattern on 8 Gbps Fibre Channel interfaces. For Nexus 9000 FC interface to operate at 8 Gbps, peer device must be configured to use a matching IDLE fill pattern. Most server and target FC interfaces do not support this and thus cannot connect to Nexus 9000 at 8 Gbps. To interoperate with other Fibre Channel switches at 8 Gbps ensure the peer switch FC interface also uses a matching IDLE fill pattern. For Cisco MDS switches, configure using the **switchport fill-pattern** interface configuration command. To connect to a peer Nexus 9000 at 8 Gbps, use no fill pattern configuration, as both devices use matching IDLE fill patterns by default.
- The default port-speed for all FC interfaces is auto for Cisco Nexus N9k-C93180YC-FX and N9k-C93360YC-FX2 switches.
- The default port-speed of all FC ports is 32G for Cisco Nexus N9K-C9336C-FX2-E switch.
- The receive B2B credit value is 64 in N9K-C93180YC-FX and 32 in N9K-C93360YC-FX2 and N9K-C9336C-FX2-E. This is not configurable.
- When a san-port channel is created, it is created in **channel mode active** by default; **channel mode on** is not supported for NPV switch.
- vFC flap may be required to bring up N Port vFC interfaces after changing the FCoE FC map.
- FC-NPV (up to 32G) and FCoE-NPV are supported on N9K-C93180YC-FX, N9K-C9336C-FX2-E, and N9k-C93360YC-FX2 switches both as NP uplink and F host port.
- Beginning with Cisco NX-OS Release 10.2(2)F, FC-NPV is supported on Cisco Nexus N9K-C9336C-FX2-E switch.
- Operating speed and member addition to san-po limitation on N9K-C9336C-FX2-E:
 1. Speed change of fc-breakout
 - Default speed is 32G
 - Speed change cannot be done on a single fc-breakout interface level
 - Speed change of fc-breakout is done on range of fc-breakout interface level
 - The range should contain full set of fc-breakout corresponds to a front panel port - for any partial range, speed config throws ERR_01 error.
 - The range should not contain any fc-breakout which is part of san-po - if range has any san-po member, speed config throws ERR_02 error.
 - The range can have fc-breakout ports corresponds to multiple front panel ports.
 2. Speed change of san-po
 - Default speed is 32G.

- Speed change of san-po is allowed only if its members includes all fc-breakout ports corresponds to a front panel port - if san-po has partial set fc-breakout ports corresponds a front panel port, speed change throws ERR_03 error.
- Speed change of san-po can be done by giving range of san-po interfaces.

3. Speed config in running config

- Speed config(not default one) will be displayed in fc-breakout interface range level; it will not be displayed under individual fc-breakout interface for the "sh runn" command. Speed config(not default one) will be displayed in the "show interface fc<int no>" command.

4. Member addition to san-po (channel-group x)

- The interface range should contain full set of fc-breakout corresponds to a front panel port - for any partial range, though the channel addition is successful, warningWARN_01 message will be thrown.
- The range can have fc-breakout ports corresponds to multiple front panel ports.

• Error and Warning messages:

- ERR_01:

if-range contains partial set of fc1/18/1-4 fc-breakout ports

- ERR_02:

if-range contains fc1/21/1-4 ports; some are part sanpo

- ERR_03:

san-port-channel21 does not contain full set of fc1/22/1-4 fc-breakout ports

- WARN_01:

Warning: if-range contains partial set of fc1/22/1-4 fc-breakout ports

Licensing Requirements for FC NPV

The following table shows the licensing requirements for FC NPV.

| Product | Product ID | License Requirement |
|-------------|--|--|
| Cisco NX-OS | <ul style="list-style-type: none"> • N93-16Y-SSK9 • N93-48Y-SSK9 • ACI-STRG | <p>FC NPV requires both the following licenses:</p> <ul style="list-style-type: none"> • SAN_ENTERPRISE_PKG—A feature license to activate FC and FCoE NPV. • FC_PORT_ACTIVATION_PKG—The number of ports to activate for FC. This comes in two variants (16 port and 48 port). <p>Note For a complete explanation of the Cisco NX-OS licensing scheme and how to obtain and apply licensees, see the Cisco NX-OS Licensing Guide .</p> |

Configuring NPV

Enabling FC NPV

FC NPV is enabled when **feature-set fcoe-npv** is installed and enabled.

To enable **fcoe-npv**, perform this task:



Note This enables both FC and FCoE NPV mode.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **install feature-set fcoe-npv**
3. switch(config-npv)# **feature-set fcoe-npv**

DETAILED STEPS

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | switch# configure terminal | Enters configuration mode. |
| Step 2 | switch(config)# install feature-set fcoe-npv | Installs the FC and FCoE NPV feature set. |
| Step 3 | switch(config-npv)# feature-set fcoe-npv | Enables FC and FCoE NPV. |

Converting Ethernet Ports to Fibre Channel

This section explains how to convert Ethernet ports to fibre channel ports.

Before you begin

This task requires installing and enabling the port license.

Step 1 Perform TCAM carving.

Example:

```
Switch(config)# hardware access-list tcam region ing-racl 1536
Switch(config)# hardware access-list tcam region ing-ifacl 256
Switch(config)# hardware access-list tcam region ing-redirect 256
```

Step 2 Confirm that **feature-set fcoe-npv** is installed and enabled.

Example:


```
Switch(config)# install feature-set fcoe-npv
Switch(config)# feature-set fcoe-npv
```

Step 3 Convert the port(s) to FC.

Example:

In this example, an Ethernet interface is being converted to FC interface on Cisco Nexus 9300-FX switches.

```
Switch(config)# slot 1
Switch(config)# port 1-4,45-48 type fc
Port type is changed. ACTION REQUIRED: Please save configurations and reload the switch
```

Note You must convert all the four front panel ports in a column to FC/Ethernet together.

In this example, an Ethernet interface is being converted to FC interface on Cisco Nexus N9K-93360YC-FX2 switches. In this switch, four ports form a port group. For example the first port group will be 1,2,49,50; the second port group will be 3,4,51,52 and likewise.

```
Switch(config)# slot 1
Switch(config)# port 1-2, 49-50 type fc
Switch(config)# port 3-4, 51-52 type fc
Port type is changed. ACTION REQUIRED: Please save configurations and reload the switch
```

In this example, an Ethernet interface is being converted to FC interface on Cisco Nexus N9K-9336C-FX2-E switches.

```
Switch(config)# slot 1
Switch(config)# port 9,12,33 type fc breakout
Port type is changed. ACTION REQUIRED: Please save configurations and reload the switch
```

Step 4 Convert the FC interface back to Ethernet port(s).

Example:

In this example, an FC interface is being converted back to Ethernet interface on Cisco Nexus 9300-FX switches.

```
Switch(config)# slot 1
Switch(config)# port 1-4,45-48 type eth
Port type is changed. ACTION REQUIRED: Please save configurations and reload the switch
```

In this example, an FC interface is being converted back to Ethernet interface on a Cisco Nexus N9K-93360YC-FX2 switch.

```
Switch(config)# slot 1
Switch(config)# port 1-2, 49-50 type eth
Port type is changed. ACTION REQUIRED: Please save configurations and reload the switch
```

In this example, an FC interface is being converted back to Ethernet interface on Cisco Nexus N9K-9336CFX2-E switches.

Note Port 1-8 cannot be converted to FC in N9K-C9336C-FX2-E.

```
Switch(config)# slot 1
Switch(config)# port 9,12,33 type eth
Port type is changed. ACTION REQUIRED: Please save configurations and reload the switch
```

After the conversion, save the configuration and reload the switch.

Note In Cisco Nexus 93180YC-FX, ports can be converted only in groups (sequential) of 4 (in multiples of 4).

Enabling the Fibre Channel Port License

This section explains how to enable the licensing for FC NPV.

Before you begin

To enable the port license, you must shut down the fibre channel (FC) ports.

Enable the port license.

Example:

```
Switch(config)# int fc1/1
Switch(config-if)# port-license acquire
```

Note This step is required during bring-up of native FC port.

Configuring FC NPV Interfaces

After you enable FC NPV, you should configure the NP uplink interfaces and the server interfaces.

Configuring FC NP Interfaces

To configure an NP uplink interface, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** { *fc slot/port/BO port* | **san-port-channel** <number> }
3. switch(config-if)# **switchport speed** *speed*
4. switch(config-if)# **switchport mode NP**
5. switch(config-if)# **no shutdown**

DETAILED STEPS

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | switch# configure terminal | Enters global configuration mode. |
| Step 2 | switch(config)# interface { <i>fc slot/port/BO port</i> san-port-channel <number> } | Selects an interface (Fibre Channel or SAN port channel) that will be connected to the core FC NPV switch. |
| Step 3 | switch(config-if)# switchport speed <i>speed</i> | Sets the speed, which can be 4G, 8G, 16G, 32G or auto. |

| | Command or Action | Purpose |
|---------------|--|---|
| | | <p>Note For 8G NP link, the fill-pattern should be set to IDLE on the core switch.</p> <p>4G and auto speed is not supported on Cisco N9K-C9336C-FX2-E switches.</p> <p>The following is an example of configuring IDLE fill pattern on a Cisco MDS switch:</p> <pre>Switch(config)# int fc2/3 Switch(config)# switchport fill-pattern IDLE speed 8000 Switch(config)# sh run int fc2/3 interface fc2/3 switchport speed 8000 switchport mode NP switchport fill-pattern IDLE speed 8000 no shutdown</pre> |
| Step 4 | switch(config-if)# switchport mode NP | Configures the interface as an NP port. |
| Step 5 | switch(config-if)# no shutdown | Brings up the interface. |

Configuring a Server Interface

To configure a server interface, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface fc slot/port**
3. switch(config-if)# **switchport speed speed**
4. switch(config-if)# **switchport mode F**
5. switch(config-if)# **no shutdown**

DETAILED STEPS

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | switch# configure terminal | Enters global configuration mode. |
| Step 2 | switch(config)# interface fc slot/port | Creates an interface that connects the server to the NPV switch. |
| Step 3 | switch(config-if)# switchport speed speed | <p>Sets the speed, which can be 4G, 8G, 16G, 32G or auto.</p> <p>Note 8G speed is not supported for server and target interfaces.</p> |
| Step 4 | switch(config-if)# switchport mode F | Configures the interface as an F port. |
| Step 5 | switch(config-if)# no shutdown | Brings up the interface. |

Configuring NPV Traffic Management

Configuring NPV Traffic Maps

An NPV traffic map associates one or more NP uplink interfaces with a server interface. The switch associates the server interface with one of these NP uplinks.



Note To map the server interface to a different uplink, the server interface must be shut down before configuring the traffic map.

To configure a traffic map, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **npv traffic-map server-interface** {fc slot/port | vfc vfc-id} **external-interface** { fc slot/port | san-port-channel <number> | vfc vfc-id | vfc-port-channel vfc-port-channel-id }
3. switch(config)# **no npv traffic-map server-interface** {fc slot/port | vfc vfc-id} **external-interface** { fc slot/port | san-port-channel <number> | vfc vfc-id | vfc-port-channel vfc-port-channel-id }

DETAILED STEPS

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | switch# configure terminal | Enters global configuration mode. |
| Step 2 | switch(config)# npv traffic-map server-interface {fc slot/port vfc vfc-id} external-interface { fc slot/port san-port-channel <number> vfc vfc-id vfc-port-channel vfc-port-channel-id } | Configures a mapping between a server interface (or range of server interfaces) and an NP uplink interface (or range of NP uplink interfaces). Note To map the server interface to a different uplink, the server interface must be shut down before configuring the traffic map. |
| Step 3 | switch(config)# no npv traffic-map server-interface {fc slot/port vfc vfc-id} external-interface { fc slot/port san-port-channel <number> vfc vfc-id vfc-port-channel vfc-port-channel-id } | Removes the mapping between the specified server interfaces and NP uplink interfaces. |

Enabling Disruptive Load Balancing

If you configure additional NP uplinks, you can enable the disruptive load-balancing feature to distribute the server traffic load evenly among all the NP uplinks.

To enable disruptive load balancing, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **npv auto-load-balance disruptive**



Note For each server interface, the External Interface value displays the assigned NP uplink.

To display the status of the server interfaces and the NP uplink interfaces, enter the **show npv status** command:

```
switch# show npv status

npiv is enabled

disruptive load balancing is disabled

External Interfaces:
=====
Interface: fcl/47, State: Down
Interface: san-port-channel 200, State: Trunking
  VSAN: 1, State: Up
  VSAN: 200, State: Up
  VSAN: 201, State: Up
  VSAN: 202, State: Up, FCID: 0xea0020
  VSAN: 100, State: Up
  VSAN: 55, State: Up
Interface: vfc-pol49, State: Trunking
  VSAN: 201, State: Up
  VSAN: 202, State: Up, FCID: 0xea0260
  VSAN: 100, State: Up
Interface: vfc-po4090, State: Trunking
  VSAN: 201, State: Up
  VSAN: 202, State: Up, FCID: 0xea0220
  VSAN: 100, State: Up
Interface: vfc1/9, State: Trunking
  VSAN: 201, State: Up
  VSAN: 202, State: Up, FCID: 0xea0240
  VSAN: 100, State: Up

Number of External Interfaces: 5

Server Interfaces:
=====
Interface: fcl/38, VSAN: 100, State: Up
Interface: fcl/39, VSAN: 202, State: Up
Interface: fcl/40, VSAN: 4094, State: Down
Interface: vfc100, VSAN: 4094, State: Down
Interface: vfc151, VSAN: 4094, State: Down
Interface: vfc1/14, VSAN: 100, State: Up

Number of Server Interfaces: 6
```



Note To view fcns database entries for FC NPV edge switches, you must enter the **show fcns database** command on the core switch.

To view all the FC NPV edge switches, enter the **show fcns database** command on the core switch:

```
core-switch# show fcns database
```

For additional details (such as IP addresses, switch names, interface names) about the FC NPV edge switches that you see in the **show fcns database** output, enter the **show fcns database detail** command on the core switch:

```

core-switch# show fcns database detail
=====
VSAN:100   FCID:0xe101c0
-----
port-wwn (vendor)      :50:0a:09:82:ad:0d:86:37 (NetApp)
node-wwn               :50:0a:09:80:8d:0d:86:37
class                 :3
node-ip-addr          :0.0.0.0
ipa                   :00 00 00 00 1e 22 a0 00
fc4-types:fc4_features :scsi-fcp:target
symbolic-port-name    :NetApp FC Target Adapter (8112) lab-D-netapp01:3b
symbolic-node-name    :NetApp FAS3240 (lab-D-netapp01)
port-type             :N
port-ip-addr          :0.0.0.0
fabric-port-wwn      :21:61:00:2a:6a:5b:da:00
hard-addr             :0x0000000
permanent-port-wwn (vendor) :50:0a:09:82:ad:0d:86:37 (NetApp)
connected interface   :vfc6/33
switch name (IP address) :MDS9706 (10.105.188.173)
-----
VSAN:100   FCID:0xe101ef
-----
port-wwn (vendor)      :50:06:01:6b:08:60:7c:71 (Clariion)
node-wwn               :50:06:01:60:88:60:7c:71
class                 :3
node-ip-addr          :0.0.0.0
ipa                   :ff ff ff ff ff ff ff ff
fc4-types:fc4_features :scsi-fcp:both
symbolic-port-name    :CLARiION:::SPB23::FC:::
symbolic-node-name    :CLARiION:::SPB::FC:::
port-type             :N
port-ip-addr          :0.0.0.0
fabric-port-wwn      :20:19:00:2a:6a:5b:da:00
hard-addr             :0x0000000
permanent-port-wwn (vendor) :50:06:01:6b:08:60:7c:71 (Clariion)
connected interface   :fcl/25
switch name (IP address) :MDS9706 (10.105.188.173)

core-switch# show interface fc 1/1
fc1/1 is trunking
  Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
  Port WWN is 20:01:2c:d0:2d:50:d2:a0
  Admin port mode is NP, trunk mode is on
  snmp link state traps are enabled
  Port mode is TNP
  Port vsan is 201
  Speed is 16 Gbps
  Transmit B2B Credit is 500
  Receive B2B Credit is 64
  Receive data field Size is 2112
  Beacon is turned off
  Belongs to san-port-channel 200
  Trunk vsans (admin allowed and active) (1,55,100,200-202,204)
  Trunk vsans (up) (100,202)
  Trunk vsans (isolated) (204)
  Trunk vsans (initializing) (1,55,200-201)
  5 minutes input rate 0 bits/sec,0 bytes/sec, 0 frames/sec
  5 minutes output rate 0 bits/sec,0 bytes/sec, 0 frames/sec
    406 frames input,40164 bytes
      0 discards,0 errors
      0 invalid CRC/FCS,0 unknown class
      0 too long,0 too short
    192 frames output,14364 bytes

```

```

0 discards,0 errors
1 input OLS,1 LRR,5 NOS,0 loop inits
3 output OLS,1 LRR, 4 NOS, 0 loop inits
500 transmit B2B credit remaining
0 low priority transmit B2B credit remaining
Last clearing of "show interface" counters :never

```

Verifying FC NPV Traffic Management

To display the FC NPV traffic map, enter the **show npv traffic-map** command.

```

switch# show npv traffic-map
NPV Traffic Map Information:
-----
Server-If      External-If(s)
-----
fc1/3          fc1/10,fc1/11
fc1/5          fc1/1,fc1/2
-----

```

To display the FC NPV internal traffic details, enter the **show npv internal info traffic-map** command.

Verifying Disruptive Load Balancing

To display the disruptive load-balancing status, enter the **show npv status** command:

```

switch# show npv status
npiv is enabled
disruptive load balancing is enabled
External Interfaces:
=====
Interface: fc1/1, VSAN: 2, FCID: 0x1c0000, State: Up
...

```

FC NPV Core Switch and FC NPV Edge Switch Configuration Example

Before you begin

This section demonstrates how to configure FC NPV core and edge switches.

Step 1 Procure and install the SAN_ENTERPRISE_PKG and PORT_ACTIVATION_PKG licenses.

Note The license file is in the .lic format and has to be copied to the switch and installed using the following command:

```
Switch# install license bootflash:Switch_port_lic_48.lic
```

Step 2 Check out the license:

```
Switch(config)# install feature-set fcoe-npv
Switch(config-vdc)# feature-set fcoe-npv
```

Step 3 Configure the needed features on the NPV:


```
Switch(config)# feature telnet
Switch(config)# feature lacp
Switch(config)# feature lldp
```

Step 4 Convert the FC port:

```
Switch(config)# slot 1
Switch(config-slot)# port 13-36 type fc
Port type is changed. ACTION REQUIRED: Please save configurations and reload the switch
```

Step 5 Configure service policies:

```
Switch(config)# system qos
Switch(config-sys-qos)# service-policy type network-qos default-fcoe-8q-nq-policy
Switch(config-sys-qos)# service-policy type queuing output default-fcoe-8q-out-policy
```

Step 6 Configure TCAM carving:

```
Switch(config-vrf)# hardware access-list tcam region ing-racl 1536
Warning: Please save config and reload the system for the configuration to take effect
Switch(config)# hardware access-list tcam region ing-redirect 256
Warning: Please save config and reload the system for the configuration to take effect
```

Step 7 Copy the running configuration to startup:

```
Switch(config)# copy running-config startup-config
[#####] 100%
```

Step 8 (Mandatory) Reload the switch so that the port conversion is applied and TCAMS are carved properly:

```
Switch(config)# reload
This command will reboot the system. (y/n)? [n] y
2017 Sep 14 10:12:19 Switch %PLATFORM-2-PFM_SYSTEM_RESET: Manual system restart from Command Line Interface
```

Step 9 Configure VLAN-VSAN mappings:

```
Switch(config)# vlan 1,20,30,40,1000,1002,1010
Switch(config-vlan)# vlan 20
Switch(config-vlan)# fcoe vsan 200
Switch(config-vlan)# vlan 30
Switch(config-vlan)# fcoe vsan 300
Switch(config-vlan)# vlan 40
Switch(config-vlan)# fcoe vsan 300
Switch(config)# vsan database
Switch(config-vsan-db)# vsan 40
Switch(config-vsan-db)# vsan 200
Switch(config-vsan-db)# vsan 300
```

Step 10 Configure the port license for FC ports:

```
Switch(config)# interface fc1/6
Switch(config-if)# port-license acquire
```

Note Checks out the port license for FC ports

Step 11 Configure the FC NP interface-facing core (this same configuration must be applied on the core switch with **switchport mode F** or **auto** for the FC interface):

```
Switch(config-if)# interface fc1/6
Switch(config-if)# switchport mode NP
Switch(config-if)# no shutdown
```

Step 12 Configure the virtual FC NP interface-facing core (this same configuration must be applied on the core switch with **switchport mode F** or **auto** for the virtual FC interface):

a) Configure the physical Ethernet interface:

```
Switch(config-if)# interface Ethernet1/7
Switch(config-if)# switchport
Switch(config-if)# switchport mode trunk
Switch(config-if)# service-policy type qos input default-fcoe-in-policy
Switch(config-if)# mtu 9216
Switch(config-if)# no shutdown
```

Note The steps *MTU* and *service-policy* are required only when a Cisco Nexus N9K-C93180YC-F, N9K-C9336C-FX2-E, or N9K-C93360YC-FX2 switch is used as the core switch.

b) Configure the virtual FC interface:

```
Switch(config-if)# interface vfc17
Switch(config-if)# bind interface ethernet1/7
Switch(config-if)# switchport mode NP
Switch(config-if)# no shutdown
```

Step 13 Configure the SAN port channel interface-facing core (This same configuration must be applied on the core switch with **switchport mode F** or **auto** for the port-channel interface. The SAN port-channel number can be different.):

a) Configure the SAN port channel:

```
Switch(config)# interface san-port-channel 250
Switch(config-if)# channel mode active
Switch(config-if)# switchport mode NP
Switch(config-if)# switchport trunk mode on
```

b) Add a member to the SAN port channel:

```
Switch(config-if)# interface fc1/13
Switch(config-if)# port-license acquire (this checks out the port license for FC ports)
Switch(config-if)# switchport trunk mode on
Switch(config-if)# channel-group 250 force
fc1/13 added to port-channel 250 and disabled
Please do the same operation on the switch at the other end of the port-channel,
then do "no shutdown" at both ends to bring it up
Switch(config-if)# no shutdown
```

Step 14 Configure the vFC port channel interface-facing core (This same configuration must be applied on the core switch with **switchport mode F** or **auto** for the virtual FC port-channel interface. The vFC port-channel number can be different):

a) Configure the Ethernet port-channel interface:

```
Switch(config)# interface port-channel500
Switch(config-if)# switchport
Switch(config-if)# switchport mode trunk
Switch(config-if)# mtu 9216
Switch(config-if)# service-policy type qos input default-fcoe-in-policy
```

Note The steps *MTU* and *service-policy* are required only when a Cisco Nexus N9K-C93180YC-FX, N9K-C9336C-FX2-E, or N9K-C93360YC-FX2 switch is used as the core switch.

- b) Add a member to the Ethernet port channel:

```
Switch(config-if)# interface Ethernet1/4
Switch(config-if)# channel-group 500 mode active
Switch(config-if)# no shutdown
```

- c) Create a virtual FC port-channel interface:

```
Switch(config)# interface vfc-po500 (this creates a vFC)
Switch(config-if)# bind interface port-channel500
Switch(config-if)# switchport mode NP
Switch(config-if)# switchport trunk mode on
```

Step 15

Configure the FCoE server interface-facing server:

- a) Configure the physical Ethernet interfaces:

```
Switch(config-if)# interface Ethernet1/6
Switch(config-if)# switchport
Switch(config-if)# switchport mode trunk
Switch(config-if)# service-policy type qos input default-fcoe-in-policy
Switch(config-if)# mtu 9216
Switch(config-if)# no shutdown
```

- b) Configure a virtual FC interface:

```
Switch(config-if)# interface vfc6
Switch(config-if)# bind interface ethernet1/6
Switch(config-if)# switchport trunk mode on
Switch(config-if)# no shutdown
```

- c) Assigning a port VSAN for the virtual FC interface:

```
Switch(config-if)# vsan database (this assigns the port vsan) (config-vsan-db)
Switch(config-vsan-db)# vsan 40 interface vfc6
```

Step 16

Configuring FC server interface

- a) Configure FC interface in F mode:

```
Switch(config)# interface fc1/39
Switch(config-if)# switchport mode F
```

- b) Apply port vsan for the FC interface:

```
Switch(config)# vsan database  
Switch(config-if)# vsan 100 interface fc1/39
```
