



Configuring N Port Virtualization

This chapter describes how to configure N port virtualization (NPV) on Cisco Nexus 5000 Series switches.

This chapter includes the following sections:

- [Information About NPV, page 33-1](#)
- [Guidelines and Limitations, page 33-5](#)
- [Configuring NPV, page 33-5](#)
- [Verifying NPV, page 33-7](#)

Information About NPV

Switch operation in NPV mode is described in the following topics:

- [NPV Overview, page 33-1](#)
- [NPV Mode, page 33-2](#)
- [NP Uplinks \(External Interfaces\), page 33-3](#)
- [FLOGI Operation, page 33-3](#)
- [FLOGI Operation, page 33-3](#)
- [NPV Traffic Management, page 33-4](#)

NPV Overview

When a switch is in fabric switch mode, it provides standard Fibre Channel switching capability and features. By default, Cisco Nexus 5000 Series switches operate in fabric switch mode.

In fabric switch mode, each switch that joins a SAN is assigned a domain ID. Each SAN (or VSAN) supports a maximum of 239 domain IDs, so the SAN has a limit of 239 switches. In a SAN topology with a large number of edge switches, the SAN may need to grow beyond this limit. NPV alleviates the domain ID limit by sharing the domain ID of the core switch among multiple edge switches.


In NPV mode, the edge switch relays all traffic from server-side ports to the core switch. The core switch provides F port functionality (such as login and port security) and all the Fibre Channel switching capabilities.

Send feedback to nx5000-docfeedback@cisco.com

The edge switch appears as a Fibre Channel host to the core switch and as a regular Fibre Channel switch to its connected devices.

Figure 33-1 shows an interface-level view of an NPV configuration.

Figure 33-1 NPV Interface Configuration



NPV Mode

In NPV mode, the edge switch relays all traffic to the core switch, which provides the Fibre Channel switching capabilities. The edge switch shares the domain ID of the core switch.

To convert a switch into NPV mode, you should enable the NPV feature and then perform a switch initialization. You cannot configure NPV mode on a per-interface basis. NPV mode applies to the entire switch.

In NPV mode, only a subset of the switch mode CLI commands and functionality is supported. For example, commands related to FLOGI and FCNS are not required, because Fibre Channel traffic is only relayed through the edge switch in NPV mode. To display the FLOGI and FCNS databases, you must enter the **show flogi database** and **show fcns database** commands on the core switch.

Server 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. For additional information about NPIV, see the [“About N Port Identifier Virtualization”](#) section on page 31-14.

Server interfaces are uniformly distributed among the NP uplinks to the core switch. All the end devices connected to a server interface are mapped to the same NP uplink.

In Cisco Nexus 5000 Series switches, server interfaces can be physical or virtual Fibre Channel interfaces.

[Send feedback to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

NP Uplinks (External 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 performs an internal fabric login (FLOGI) to the core switch, and then (if the FLOGI is successful) it registers itself with the name server on the core switch. Subsequent FLOGIs from end devices connected to this NP uplink are converted to fabric discovery messages (FDISCs). For additional information about fabric login, see the “[Information About Fabric Login](#)” section on page 40-1.



Note

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

In Cisco Nexus 5000 Series switches, NP uplink interfaces must be native Fibre Channel interfaces.


FLOGI Operation

When an NP port becomes operational, the Cisco Nexus 5000 Series 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 Cisco Nexus 5000 Series 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).

[Figure 33-2](#) shows the internal FLOGI flows between the core switch and an edge switch.

Figure 33-2 Internal FLOGI Flows



Send feedback to nx5000-docfeedback@cisco.com

Table 33-1 identifies the internal FLOGI parameters that appear in Figure 33-2.

Table 33-1 Internal FLOGI Parameters

Parameter	Derived From
pWWN	The fWWN of the NP port.
nWWN	The VSAN-based sWWN of the edge switch.
fWWN	The fWWN of the F port on the core 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 NPV switch name.



Note

The BB_SCN of internal FLOGIs on NP ports is always set to zero. The BB_SCN is supported by the F port on the edge switch.

We do not recommend using fWWN-based zoning on the edge switch for the following reasons:

- Zoning is not enforced at the edge switch (rather, it is enforced on the core switch).
- Multiple devices attached to an edge switch log in through the same F port on the core, so they cannot be separated into different zones.
- The same device might log in using different fWWNs on the core switch (depending on the NPV link it uses) and may need to be zoned using different fWWNs.

For additional information about zoning, see the “[Information About Zoning](#)” section on page 37-1.

NPV Traffic Management

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. Only server interfaces that become operational after the new NP uplink can select this NP uplink.

In Release 4.0(0)N1(2a) and later software releases, NPV supports disruptive load balancing. When disruptive load balancing is enabled, 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, 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.

Send feedback to nx5000-docfeedback@cisco.com

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.

Guidelines and Limitations

When configuring NPV, note the following guidelines and limitations:

- In-order data delivery is not required in 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 available member types on the core switch. For fWWN, sWWN, domain, or port-based zoning, use the fWWN, sWWN, domain, or port of the core switch in the configuration commands.
- Port security is supported on the core switch for devices logged in through the 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 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.
- NPV uses a load-balancing algorithm to automatically assign end devices in a VSAN to one of the NP uplinks (in the same VSAN) upon initial login. If there are multiple NP uplinks in the same VSAN, you cannot assign an end device to a specific NP uplink.
- 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.
- Fibre Channel switching is not performed in the edge switch; all traffic is switched in the core switch.
- NPV supports NPIV-capable module servers. This capability is called nested NPIV.
- Only F, NP, and SD ports are supported in NPV mode.

Configuring NPV

When you enable NPV, the system configuration is erased and the switch is reinitialized with NPV mode enabled.



Note

We recommend that you save your current configuration either in boot flash memory or to a TFTP server before you enable NPV.

Send feedback to nx5000-docfeedback@cisco.com

To configure NPV, perform this task:

	Command	Purpose
Step 1	switch# configure terminal switch(config)#	Enters configuration mode.
Step 2	switch(config)# npv enable	Enables NPV mode. The switch is reinitialized, and it comes back up in NPV mode. Note A write-erase is performed during the initialization.
Step 3	switch(config)# interface fc slot/port	Selects an interface that will be connected to the core NPV switch.
Step 4	switch(config-if)# switchport mode NP switch(config-if)# no shutdown	Configures the interface as an NP port. Brings up the interface. (Repeat steps 3 and 4 for each NP uplink.)
Step 5	switch(config-if)# exit	Exits interface mode for the port.
Step 6	switch(config)# interface fc slot/port or switch(config)# interface vfc slot/port	Selects a server interface.
Step 7	switch(config-if)# switchport mode F switch(config-if)# no shutdown	Configures the interface as an F port. Brings up the interface. (Repeat steps 6 and 7 for each server interface.)
Step 8	switch(config-npv)# no npv enable switch(config)#	Disables NPV mode, which results in a reload of the switch.

Configuring NPV Traffic Management

Configure NPV traffic management after configuring the interfaces in NPV mode.

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

To enable or disable disruptive load balancing, perform this task:

	Command	Purpose
Step 1	switch# configure terminal switch(config)#	Enters configuration mode on the NPV.
Step 2	switch(config)# npv auto-load-balance disruptive switch (config)#	Enables disruptive load balancing on the switch.
Step 3	switch (config)# no npv auto-load-balance disruptive	Disables disruptive load balancing on the switch.

[Send feedback to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Verifying NPV

To display information about NPV, perform the following task:

Command	Purpose
switch# show npv flogi-table [all]	Displays the NPV configuration.

To display a list of devices that are logged in, along with VSANs, source information, pWWNs, and FCIDs, enter the **show npv flogi-table** command on the Cisco Nexus 5000 Series switch:

```
switch# show npv flogi-table
-----
SERVER
INTERFACE VSAN FCID                PORT NAME                NODE NAME                EXTERNAL
-----
vfc3/1    1    0xee0008 10:00:00:00:c9:60:e4:9a 20:00:00:00:c9:60:e4:9a fc2/1
vfc3/1    1    0xee0009 20:00:00:00:0a:00:00:01 20:00:00:00:c9:60:e4:9a fc2/2
vfc3/1    1    0xee000a 20:00:00:00:0a:00:00:02 20:00:00:00:c9:60:e4:9a fc2/3
vfc3/1    1    0xee000b 33:33:33:33:33:33:33:33 20:00:00:00:c9:60:e4:9a fc2/4
```

Total number of flogi = 4



Note

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

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

```
switch# show npv status
npiv is enabled

External Interfaces:
=====
Interface: fc2/1, VSAN: 2, FCID: 0x1c0000, State: Up
Interface: fc2/2, VSAN: 3, FCID: 0x040000, State: Up

Number of External Interfaces: 2

Server Interfaces:
=====
Interface: vfc2/1, VSAN: 2, NPIV: No, State: Up
Interface: vfc1/1, VSAN: 3, NPIV: No, State: Up

Number of Server Interfaces: 2
```



Note

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

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

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

Send feedback to nx5000-docfeedback@cisco.com

For additional details (such as IP addresses, switch names, interface names) about the 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
```

Disruptive load balancing is supported in Release 4.0(0)N1(2a) and later releases. To display the disruptive load-balancing status, enter the following command:

```
switch# show npv status
npiv is enabled
disruptive load balancing is enabled
External Interfaces:
=====
  Interface: fc2/1, VSAN: 2, FCID: 0x1c0000, State: Up
  ...
```