



## Configuring N Port Virtualization

---

N Port virtualization (NPV) reduces the number of Fibre Channel domain IDs in SANs. Switches operating in the NPV mode do not join a fabric; rather, they pass traffic between NPV core switch links and end devices, which eliminates the domain IDs for these edge switches.

NPV is supported by the following Cisco MDS 9000 switches only:

- Cisco MDS 9124 Multilayer Fabric Switch
- Cisco MDS 9134 Fabric Switch
- Cisco Fabric Switch for HP c-Class BladeSystem
- Cisco Fabric Switch for IBM BladeCenter



**Note**

---

NPV is available on these switches only while in NPV mode; if in switch mode, NPV is not available.

---

This chapter includes the following sections:

- [About NPV, page 13-1](#)
- [NPV Guidelines and Requirements, page 13-6](#)
- [Configuring NPV, page 13-7](#)
- [Verifying NPV, page 13-10](#)

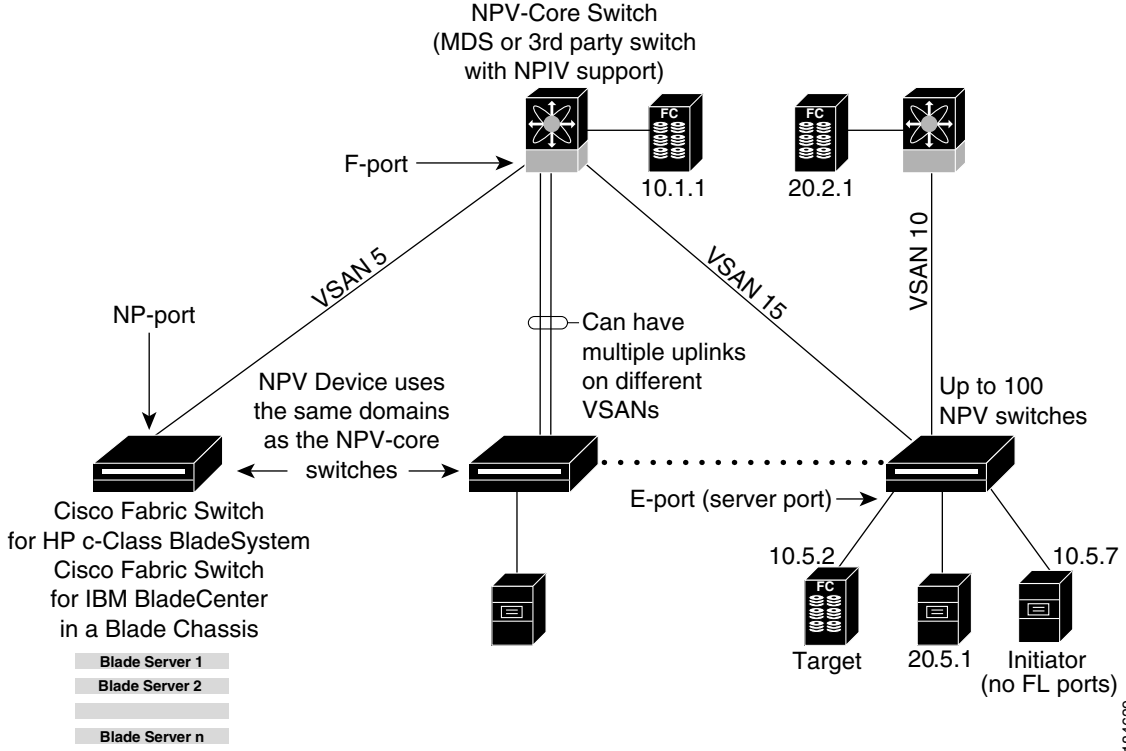
## About NPV

Typically, Fibre Channel networks are deployed using a core-edge model with a large number of fabric switches connected to core devices. However, as the number of ports in the fabric increases, the number of switches deployed also increases, and you can end up with a dramatic increase in the number of domain IDs (the maximum number supported is 239). This challenge becomes even more difficult when additional blade chassis are deployed in Fibre Channel networks.

NPV addresses the increase in the number of domain IDs needed to deploy a large number of the ports by making a fabric or module switch appear as a host to the core Fibre Channel switch, and as a Fibre Channel switch to the servers in the fabric or blade switch. NPV aggregates multiple locally connected N ports into one or more external NP links, which shares the domain ID of the NPV core switch among multiple NPV switches (see [Figure 13-1](#)). NPV also allows multiple devices to attach to the same port on the NPV core switch, thereby reducing the need for more ports on the core.

Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Figure 13-1 Cisco NPV Fabric Configuration




184639

While NPV is similar to N port identifier virtualization (NPIV), it does not offer exactly the same functionality. NPIV provides a means to assign multiple FC IDs to a single N port, and allows multiple applications on the N port to use different identifiers. NPIV also allows access control, zoning, and port security to be implemented at the application level. NPV makes use of NPIV to get multiple FCIDs allocated from the core switch on the NP port.

Figure 13-2 shows a more granular view of an NPV configuration at the interface level.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Figure 13-2 Cisco NPV Configuration–Interface View



## NPV Mode

A switch is in NPV mode after a user has enabled NPV and the switch has successfully rebooted. NPV mode applies to an entire switch. All end devices connected to a switch that is 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 interswitch 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.



### Note

In-order data delivery is not required in NPV mode because the exchange between two end devices always takes the same uplink to the core from the NPV device. For traffic beyond the NPV device, core switches will enforce in-order delivery if needed and/or configured.

After entering NPV mode, only the following commands are available:

<b>aaa</b>	Configure aaa functions
<b>arp</b>	[no] remove an entry from the ARP cache
<b>banner</b>	Configure banner message
<b>boot</b>	Configure boot variables
<b>callhome</b>	Enter the callhome configuration mode
<b>cli</b>	CLI configuration commands
<b>clock</b>	Configure time-of-day clock
<b>do</b>	EXEC command
<b>end</b>	Exit from configure mode
<b>exit</b>	Exit from configure mode
<b>fcanalyzer</b>	Configure cisco fabric analyzer
<b>fcxbbcredit</b>	Enable extended rx b2b credit configuration
<b>fips</b>	Enable/Disable FIPS mode
<b>hw-module</b>	Enable/Disable OBFL information
<b>interface</b>	Select an interface to configure
<b>ip</b>	Configure IP features
<b>ipv6</b>	Configure IPv6 features
<b>line</b>	Configure a terminal line
<b>logging</b>	Modify message logging facilities
<b>no</b>	Negate a command or set its defaults
<b>npv</b>	Config commands for FC N_port Virtualizer
<b>ntp</b>	NTP Configuration
<b>port-track</b>	Configure Switch port track config
<b>power</b>	Configure power supply
<b>poweroff</b>	Poweroff a module in the switch
<b>radius</b>	Configure RADIUS configuration
<b>radius-server</b>	Configure RADIUS related parameters

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

<b>rate-mode</b>	Configure rate mode oversubscription limit
<b>rmon</b>	Remote Monitoring
<b>role</b>	Configure roles
<b>snmp-server</b>	Configure snmp server
<b>ssh</b>	Configure SSH parameters
<b>switchname</b>	Configure system's network name
<b>system</b>	System config command
<b>tacacs+</b>	Enable tacacs+
<b>telnet</b>	Enable telnet
<b>username</b>	Configure user information.
<b>wwn</b>	Set secondary base MAC addr and range for additional WWNs

## NP Ports

An *NP port* (proxy N port) is a port on a device that is in NPV mode and connected to the NPV core switch using an F port. NP ports behave like N ports except that in addition to providing N port behavior, they also function as proxies for multiple, physical N ports.

## NP Links

An *NP link* is basically an NPIV uplink to a specific end device. NP links are established when the uplink to the NPV core switch comes up; the links are terminated when the uplink goes down. Once the uplink is established, the NPV switch performs an internal FLOGI to the NPV core switch, and then (if the FLOGI is successful) registers itself with the NPV core switch's name server. Subsequent FLOGIs from end devices in this NP link are converted to FDISCs. For more details refer to the [“Internal FLOGI Parameters”](#) section on page 13-4.

Server links are uniformly distributed across the NP links. All the end devices behind a server link will be mapped to only one NP link.

## Internal FLOGI Parameters

When an NP port comes up, the NPV device first logs itself in to the NPV core switch and sends a FLOGI request that includes the following parameters:

- The fWWN (fabric port WWN) of the NP port used as the pWWN in the internal login.
- The VSAN-based sWWN (switch WWN) of the NPV device used as nWWN (node WWN) in the internal FLOGI.

After completing its FLOGI request, the NPV device registers itself with the fabric name server using the following additional parameters:

- Switch name and interface name (for example, fc1/4) of the NP port is embedded in the symbolic port name in the name server registration of the NPV device itself.
- The IP address of the NPV device is registered as the IP address in the name server registration of the NPV device.



### Note

The BB\_SCN of internal FLOGIs on NP ports is always set to zero. The BB\_SCN is supported at the F-port of the NPV device.

[Figure 13-3](#) shows the internal FLOGI flows between an NPV core switch and an NPV device.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

**Figure 13-3 Internal FLOGI Flows**




Table 13-1 identifies the internal FLOGI parameters that appear in Figure 13-3.

**Table 13-1 Internal FLOGI parameters**

Parameter	Derived From
pWWN	The fWWN of the NP port.
nWWN	The VSAN-based sWWN of the NPV device.
fWWN	The fWWN of the F port on the NPV core switch.
symbolic port name	The switch name and NP port interface string. <b>Note</b> If there is no switch name available, then the output will simply read “switch.” For example, <code>switch: fc1/5</code> .
IP address	The IP address of the NPV device.
symbolic node name	The NPV switch name.

Although fWWN-based zoning is supported for NPV devices, it is not recommended because:

- Zoning is not enforced at the NPV device (rather, it is enforced on the NPV core switch).
- Multiple devices behind an NPV device log in via the same F port on the core (hence, they use same fWWN and 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.

## Default Port Numbers

Port numbers on NPV-enabled switches will vary depending on the switch model. For details about port numbers for NPV-eligible switches, see [Chapter 4, “On-Demand Port Activation Licensing.”](#)

## NPV Traffic Management

Before Cisco MDS SAN-OS release 3.3(1a), NPV supported automatic selection of external links. When a server interface is brought up, an external interface with the minimum load is selected from the available links. There is no manual selection on the server interfaces using the external links. Also, when

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

a new external interface was brought up, the existing load was not distributed automatically to the newly available external interface. This newly brought up interface is used only by the server interfaces that come up after this interface.

As in Cisco MDS SAN-OS release 3.3(1a), NPV supports traffic management by allowing you to select and configure the external interfaces that the server uses to connect to the core switches.



**Note**

---

When the NPV traffic management is configured, the server uses only the configured external interfaces. Any other available external interface will not be used.

---

The NPV traffic management feature provides the following benefits:

- Facilitates traffic engineering by providing dedicated external interfaces for the servers connected to NPV.
- Uses the shortest path by selecting external interfaces per server interface.
- Uses the persistent FC ID feature by providing the same traffic path after a link break, or reboot of the NPV or core switch.
- Balances the load by allowing the user to evenly distribute the load across external switches.

## NPV Guidelines and Requirements

Following are recommended guidelines and requirements when deploying NPV:

- NPV core switches must support NPIV.
- You can have up to 100 NPV devices.
- Nondisruptive upgrades are supported. See [Chapter 7, “Software Images.”](#)
- Port tracking is supported. See [Chapter 57, “Configuring Port Tracking.”](#)
- You can configure zoning for end devices that are connected to NPV devices using all available member types on the NPV core switch. If fWWN, sWWN, domain, or port-based zoning is used, then fWWN, sWWN or the domain/port of the NPV core switch should be used.
- Port security is supported on the NPV core switch for devices logged in via NPV.
- NPV uses a load balancing algorithm to automatically assign end devices in a VSAN to one of the NPV core switch links (in the same VSAN) upon initial login. If there are multiple NPV core switch links in the same VSAN, then you cannot assign a specific one to an end device.
- Both servers and targets can be connected to an NPV device.
- Remote SPAN is not supported.
- Local switching is not supported; all traffic is switched using the NPV core switch.
- NPV devices can connect to multiple NPV core switches. In other words, different NP ports can be connected to different NPV core switches.
- NPV supports NPIV-capable module servers (nested NPIV).
- Only F, NP, and SD ports are supported in NPV mode.



**Note**

---

In the case of servers that are booted over the SAN with NPV, if an NPV link failover occurs, servers will lose access to their boot LUN temporarily.

---

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

## NPV Traffic Management Guidelines

When deploying NPV traffic management, follow these guidelines:

- Use NPV traffic management only when the automatic traffic engineering by the NPV device is not sufficient for the network requirements.
- Do not configure traffic maps for all the servers. For non-configured servers, NPV will use automatic traffic engineering.
- Configure the Persisten FC ID on the core switch. Traffic engineering directs the associated server interface to external interfaces that lead to the core switch.
- Server interfaces configured to a set of external interfaces cannot use any other available external interfaces, even if the configured interfaces are not available.
- Do not configure non-disruptive load balancing because this involves moving a device from one external interface to another interface. Moving the device between external interfaces requires NPV relogin to the core switch through F port leading to traffic disruption.
- Link a set of servers to a core switch by configuring the server to a set of external interfaces that are linked to the core switch.

## Configuring NPV

When you enable NPV, your system configuration is erased and the system is rebooted with NPV mode enabled.



### Note

We recommend that you save your current configuration either bootflash or a TFTP server before NPV (if the configuration is required for later use). Use the following commands to save either your non-NPV or NPV configuration:

```
switch# copy running bootflash:filename
```

The configuration can be reapplied later using the following command:

```
switch# copy bootflash:filename running-config
```

### SUMMARY STEPS

1. Enable NPIV on the NPV core switch. Enable NPV on the NPV device.
2. Configure the interfaces connected to the NPV core switch as NP ports.
3. Configure NPV link as an F port on the NPV core switch.
4. Configure server link on the NPV device.



### Note

On the 91x4 platform, before you upgrade to 3.2(2b) or downgrade from 3.2(2b), shut the F-ports connected to NPIV capable hosts, and then disable the NPIV feature. After the upgrade or downgrade is complete, enable the NPIV feature and up the F-ports.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

On the 91x4 platform, before you downgrade from 3.2(2b) to prior versions, shut the F-port, enable and disable the FC domain persistency for that VSAN and then up the F-port. Use the following steps to configure NPV.

	Command	Purpose
Step 1	switch# <b>config t</b> switch(config)#	On the NPV core switch, enters configuration mode.
Step 2	switch(config)# <b>npiv enable</b> switch (config)#	Enables NPIV mode on the NPV core switch.
	switch (config)# <b>no npiv enable</b>	Disables NPIV mode on the NPV core switch.
Step 3	switch(config)# <b>interface fc2/1</b> switch(config-if)# <b>switchport mode F</b>	Configure the NPIV core switch port as an F port.
	switch(config-if)# <b>no shutdown</b>	Changes Admin status to bring up the interfaces.
Step 4	switch(config)# <b>npv enable</b>	Enables NPV mode on a NPV device (module, Cisco MDS 9124 or Cisco MDS 9134 Fabric Switch). The module or switch is rebooted, and when it comes back up, is in NPV mode.  <b>Note</b> A write-erase is performed during the reboot.
Step 5	switch(config)# <b>interface fc1/1</b> switch(config-if)# <b>switchport mode NP</b>	On the NPV device, select the interfaces that will be connected to the aggregator switch and configure them as NP ports.
	switch(config-if)# <b>no shutdown</b>	Changes Admin status to bring up the interfaces.
Step 6	switch(config-if)# <b>exit</b>	Exits interface mode for the port.
Step 7	switch(config)# <b>interface fc1/2</b> switch(config-if)# <b>switchport mode F</b>	Selects the remaining interfaces on the NPV-enabled device and configures them as F ports.
	switch(config-if)# <b>no shutdown</b>	Changes Admin status to bring up the interfaces.
Step 8	switch(config-npv)# <b>no npv enable</b> switch(config)#	Terminates session and disables NPV mode, which results in a reload of the NPV device.

## Configuring NPV Traffic Management

The NPV traffic management feature is enabled after configuring NPV. Configuring NPV traffic management involves configuring a list of external interfaces to the servers, and enabling or disabling disruptive load balancing.

### Configuring List of External Interfaces per Server Interfaces

A list of external interfaces are linked to the server interfaces when the server interface is down, or if the specified external interface list includes the external interface already in use.



**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

To configure list of external interfaces per server interfaces, perform the following tasks:

	Command	Purpose
Step 1	switch# <b>config t</b> switch(config)#	Enters configuration mode on the NPV.
Step 2	switch(config)# <b>npv traffic-map</b> <b>server-interface</b> svr-if-range <b>external-interface</b> ext-if-range switch (config)#	Allows you to configure a list of external interfaces per server interfaces by specifying the external interfaces in the <i>svr-if-range</i> . The server to be linked is specified in the <i>ext-if-range</i> .
	switch(config)# <b>no npv traffic-map</b> <b>server-interface</b> svr-if-range <b>external-interface</b> ext-if-range switch (config)#	Disables the NPV traffic management feature on the NPV.

### Enable or Disable the Global Policy for Disruptive Load Balancing

Disruptive load balancing allows you to review the load on all the external interfaces and balance the load disruptively. Disruptive load balancing is done by moving the servers using heavily loaded external interfaces, to the external interfaces running with fewer loads.

To enable or disable the global policy for disruptive load balancing, perform the following tasks:

	Command	Purpose
Step 1	switch# <b>config t</b> switch(config)#	Enters configuration mode on the NPV.
Step 2	switch(config)# <b>npv auto-load-balance</b> <b>disruptive</b> switch (config)#	Enables disruptive load balancing on the NPV core switch.
Step 3	switch (config)# <b>no npv auto-load-balance</b> <b>disruptive</b>	Disables disruptive load balancing on the NPV core switch.

## Multiple VSAN Support

By grouping devices into different NPV sessions based on VSANs, it is possible to support multiple VSANs at the NPV-enabled switch. The correct uplink must be selected based on the VSAN(s) that the uplink can carry.

## DPVM Configuration

When NPV is enabled, the following requirements must be met before you configure DPVM on the NPV core switch:

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)***

- You must explicitly configure the WWN of the internal FLOGI in DPVM. If DPVM is configured on the NPV core switch for an end device that is connected to the NPV device, then that end device must be configured to be in the same VSAN. Logins from a device connected to an NPV device will fail if the device is configured to be in a different VSAN. To avoid VSAN mismatches, ensure that the internal FLOGI VSAN matches the port VSAN of the NP port.
- The first login from an NP port determines the VSAN of that port. If DPVM is configured for this first login—which is the internal login of the NPV device—then the NPV core switch’s VSAN F port is located in that VSAN. Otherwise, the port VSAN remains unchanged.

For details about DPVM configuration, see [Chapter 21, “Creating Dynamic VSANs.”](#)

## NPV and Port Security

Port security is enabled on the NPV core switch on a per interface basis. To enable port security on the NPV core switch for devices logging in via NPV, you must adhere to the following requirements:

- The internal FLOGI must be in the port security database; in this way, the port on the NPV core switch will allow communications/links.
- All the end device pWWNs must also be in the port security database.

Once these requirements are met, you can enable port security as you would in any other context. For details about enabling port security, see [Chapter 37, “Configuring Port Security.”](#)

## Verifying NPV

To view all the NPV devices in all the VSANs that the aggregator switch belongs to, enter the **show fcns database** command.

```
switch# show fcns database

VSAN 1:
-----
FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE
-----
0x010000 N 20:01:00:0d:ec:2f:c1:40 (Cisco) npv
0x010001 N 20:02:00:0d:ec:2f:c1:40 (Cisco) npv
0x010200 N 21:00:00:e0:8b:83:01:a1 (Qlogic) scsi-fcp:init
0x010300 N 21:01:00:e0:8b:32:1a:8b (Qlogic) scsi-fcp:init

Total number of entries = 4
```

For additional details (such as IP addresses, switch names, interface names) about the NPV devices you see in the **show fcns database** output, enter the **show fcns database detail** command.

```
switch# show fcns database detail

-----
VSAN:1 FCID:0x010000
-----
port-wwn (vendor) :20:01:00:0d:ec:2f:c1:40 (Cisco)
node-wwn :20:00:00:0d:ec:2f:c1:40
class :2,3
node-ip-addr :172.20.150.38
ipa :ff ff ff ff ff ff ff ff
fc4-types:fc4_features :npv
symbolic-port-name :para-3:fc1/1
```

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)**

```

symbolic-node-name :para-3
port-type :N
port-ip-addr :0.0.0.0
fabric-port-wwn :20:01:00:0d:ec:04:99:40
hard-addr :0x000000
permanent-port-wwn (vendor) :20:01:00:0d:ec:2f:c1:40 (Cisco)

-----
VSAN:1 FCID:0x010001
-----
port-wwn (vendor) :20:02:00:0d:ec:2f:c1:40 (Cisco)
node-wwn :20:00:00:0d:ec:2f:c1:40
class :2,3
node-ip-addr :172.20.150.38
ipa :ff ff ff ff ff ff ff ff
fc4-types:fc4_features :npv
symbolic-port-name :para-3:fc1/2
symbolic-node-name :para-3
port-type :N
port-ip-addr :0.0.0.0
fabric-port-wwn :20:02:00:0d:ec:04:99:40
hard-addr :0x000000
permanent-port-wwn (vendor) :20:02:00:0d:ec:2f:c1:40 (Cisco)

```

If you need to contact support, enter the **show tech-support NPV** command and save the output so that support can use it to troubleshoot, if necessary.

To display a list of the NPV devices that are logged in, along with VSANs, source information, pWWNs, and FCIDs, enter the **show npv flogi-table** command.

```

switch# show npv flogi-table
-----
SERVER
INTERFACE VSAN FCID          PORT NAME          NODE NAME          EXTERNAL
-----
fc1/19     1      0xee0008 10:00:00:00:c9:60:e4:9a 20:00:00:00:c9:60:e4:9a fc1/9
fc1/19     1      0xee0009 20:00:00:00:0a:00:00:01 20:00:00:00:c9:60:e4:9a fc1/1
fc1/19     1      0xee000a 20:00:00:00:0a:00:00:02 20:00:00:00:c9:60:e4:9a fc1/9
fc1/19     1      0xee000b 33:33:33:33:33:33:33:33 20:00:00:00:c9:60:e4:9a fc1/1

```

Total number of flogi = 4.

To display the status of the different servers and external interfaces, enter the **show npv status** command.

```

switch# show npv status
npiv is enabled

External Interfaces:
=====
Interface: fc1/1, VSAN: 2, FCID: 0x1c0000, State: Up
Interface: fc1/2, VSAN: 3, FCID: 0x040000, State: Up

Number of External Interfaces: 2

Server Interfaces:
=====
Interface: fc1/7, VSAN: 2, NPIV: No, State: Up
Interface: fc1/8, VSAN: 3, NPIV: No, State: Up

Number of Server Interfaces: 2

```

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)*

## Verifying NPV Traffic Management

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

NPV Traffic Map Information:

```

-----
Server-If      External-If(s)
-----
fc1/3          fc1/10,fc1/11
fc1/5          fc1/1,fc1/2
-----

```

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

NPV Traffic Map Information:

```

-----
Server-If      External-If(s)
-----
fc1/3          fc1/10,fc1/11
fc1/5          fc1/1,fc1/2
-----

```