



Fibre Channel NPV

This chapter contains the following sections:

- [Fibre Channel Connectivity Overview, on page 1](#)
- [NPV Traffic Management, on page 3](#)
- [SAN A/B Separation, on page 5](#)
- [SAN Port Channels, on page 6](#)
- [Fibre Channel N-Port Virtualization Guidelines and Limitations, on page 7](#)
- [Fibre Channel N-Port Virtualization Supported Hardware, on page 8](#)
- [Fibre Channel N-Port Virtualization Interoperability, on page 8](#)
- [Fibre Channel NPV GUI Configuration, on page 9](#)
- [Fibre Channel NPV NX-OS-Style CLI Configuration, on page 15](#)
- [Fibre Channel NPV REST API Configuration, on page 19](#)

Fibre Channel Connectivity Overview

Cisco ACI supports Fibre Channel (FC) connectivity on a leaf switch using N-Port Virtualization (NPV) mode. NPV allows the switch to aggregate FC traffic from locally connected host ports (N ports) into a node proxy (NP port) uplink to a core switch.

A switch is in NPV mode after enabling NPV. NPV mode applies to an entire switch. Each end device connected to an NPV mode switch 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.



Note In the FC NPV application, the role of the ACI leaf switch is to provide a path for FC traffic between the locally connected SAN hosts and a locally connected core switch. The leaf switch does not perform local switching between SAN hosts, and the FC traffic is not forwarded to a spine switch.

FC NPV Benefits

FC NPV provides the following:

- Increases the number of hosts that connect to the fabric without adding domain IDs in the fabric. The domain ID of the NPV core switch is shared among multiple NPV switches.


- FC and FCoE hosts connect to SAN fabrics using native FC interfaces.
- Automatic traffic mapping for load balancing. For newly added servers connected to NPV, traffic is automatically distributed among the external uplinks based on current traffic loads.
- Static traffic mapping. A server connected to NPV can be statically mapped to an external uplink.

FC NPV Mode

Feature-set `fcoe-npv` in ACI will be enabled automatically by default when the first FCoE/FC configuration is pushed.

FC Topology

The topology of various configurations supporting FC traffic over the ACI fabric is shown in the following figure:



- Server/storage host interfaces on the ACI leaf switch can be configured to function as either native FC ports or as virtual FC (FCoE) ports.

- An uplink interface to a FC core switch can be configured as any of the following port types:
 - native FC NP port
 - SAN-PO NP port
- An uplink interface to a FCF switch can be configured as any of the following port types:
 - virtual (vFC) NP port
 - vFC-PO NP port
- N-Port ID Virtualization (NPIV) is supported and enabled by default, allowing an N port to be assigned multiple N port IDs or Fibre Channel IDs (FCID) over a single link.
- Trunking can be enabled on an NP port to the core switch. Trunking allows a port to support more than one VSAN. When trunk mode is enabled on an NP port, it is referred to as a TNP port.
- Multiple FC NP ports can be combined as a SAN port channel (SAN-PO) to the core switch. Trunking is supported on a SAN port channel.
- FC F ports support 4/16/32 Gbps and auto speed configuration, but 8Gbps is not supported for host interfaces. The default speed is "auto."
- FC NP ports support 4/8/16/32 Gbps and auto speed configuration. The default speed is "auto."
- Multiple FDISC followed by Flogi (nested NPIV) is supported with FC/FCoE host and FC/FCoE NP links.
- An FCoE host behind a FEX is supported over an FCoE NP/uplink.
- Starting in the APIC 4.1(1) release, an FCoE host behind a FEX is supported over the Fibre Channel NP/uplink.
- All FCoE hosts behind one FEX can either be load balanced across multiple vFC and vFC-PO uplinks, or through a single Fibre Channel/SAN port channel uplink.
- SAN boot is supported on a FEX through an FCoE NP/uplink.
- Starting in the APIC 4.1(1) release, SAN boot is also supported over a FC/SAN-PO uplink.
- SAN boot is supported over vPC for FCoE hosts that are connected through FEX.

NPV Traffic Management

In most cases, Cisco recommends allowing all traffic to use all available uplinks. Use FC NPV traffic management only when automatic traffic engineering does not meet your network requirements.

Automatic Uplink Selection

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

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

Traffic Maps

FC NPV supports traffic maps. A traffic map allows you to specify the external (NP uplink) interfaces that a server (host) 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 external interfaces in its traffic map. If none of the specified external interfaces 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 external interfaces 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 external interface (or one of a specified set of external interfaces) by providing the same traffic path 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.

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.
- When initially configuring traffic map pinning, you must shut the server host port before configuring the first traffic map.
- If traffic mapping is configured for more than one uplink, when removing the traffic map through which a host has logged in, you must first shut the host before removing the traffic map.
- While configuring a traffic map for an FCoE host behind a FEX, you can map one host to either multiple FCoE NP/uplinks (VFC or VFC-PO) or to a single Fibre Channel/SAN port channel NP/uplink.

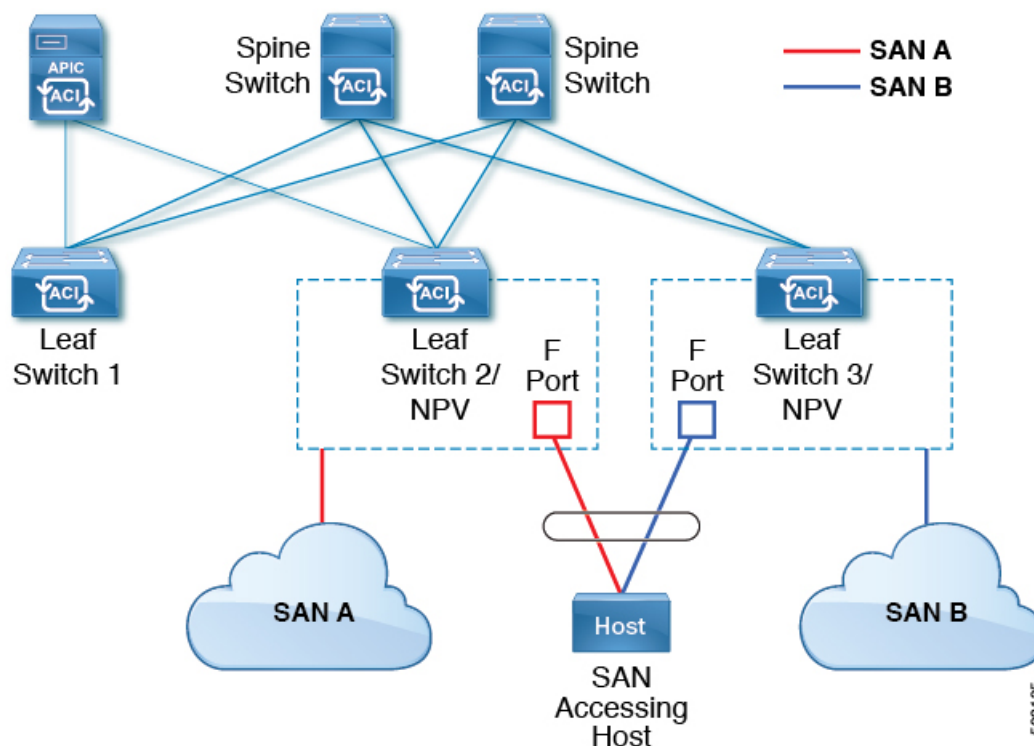
**Note**

When a server is statically mapped to an external interface, the server traffic is not redistributed in the event that the external interface becomes down for any reason.

SAN A/B Separation

SAN A and SAN B separation ensures that SAN connectivity is available even if one of the fabric components fails. SAN A and SAN B separation can be achieved physically or logically by separating the VSANs that are carried across the fabric.

Figure 1: SAN A/B Separation



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 ACI switches are used to connect to FC core switches and to provide optimal bandwidth utilization and transparent failover between the uplinks of a VSAN.

SAN Port Channel Guidelines and Limitations

- The maximum number of active port channels (SAN port channels plus VFC uplink/NP port channels) on the Cisco ACI switch is seven. Any additional configured port channels remain in the **errdisabled** state until you shut or delete one of the existing active port channels. After you shut/delete an existing active port channel, shut/no shut the **errdisabled** port channel to bring it up.
- 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 ACI switches for SAN port channels is **active**; this cannot be changed.
- When a SAN port channel is connected to a Cisco FC core switch, only channel mode active is supported. Channel mode active must be configured on the Cisco FC core switch.

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 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 is disabled. The active port channel mode allows automatic recovery without explicitly enabling and disabling the port-channel-member ports at either end.

Fibre Channel N-Port Virtualization Guidelines and Limitations

When configuring Fibre Channel N-Port Virtualization (NPV), note the following guidelines and limitations:

- Fibre Channel NP ports support trunk mode, but Fibre Channel F ports do not.
- On a trunk Fibre Channel port, internal login happens on the highest VSAN.
- On the core switch, the following features must be enabled:


```
feature npiv
feature fport-channel-trunk
```
- To use an 8G uplink speed, you must configure the IDLE fill pattern 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)# show run int fc2/3

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

- Fibre Channel NPV support is limited to the Cisco N9K-C93180YC-FX switch.
- You can use ports 1 through 48 for Fibre Channel configuration. Ports 49 through 54 cannot be Fibre Channel ports.
- If you convert a port from Ethernet to Fibre Channel or the other way around, you must reload the switch. Currently, you can convert only one contiguous range of ports to Fibre Channel ports, and this range must be a multiple of 4, ending with a port number that is a multiple of 4. For example, 1-4, 1-8, or 21-24.
- Fibre Channel Uplink (NP) connectivity to Brocade Port Blade Fibre Channel 16-32 is not supported when a Cisco N9K-93180YC-FX leaf switch port is configured in 8G speed.
- The selected port speed must be supported by the SFP. For example, because a 32G SFP supports 8/16/32G, a 4G port speed requires an 8G or 16G SFP. Because a 16G SFP supports 4/8/16G, a 32G port speed requires a 32G SFP.
- Speed autonegotiation is supported. The default speed is 'auto'.
- You cannot use Fibre Channel on 40G and breakout ports.
- FEX cannot be directly connected to FC ports.

- FEX HIF ports cannot be converted to FC.
- Reloading a switch after changing a switch's port profile configuration interrupts traffic through the data plane.

Fibre Channel N-Port Virtualization Supported Hardware

Fibre Channel N-Port Virtualization (FC NPV) is supported on the following switches:

- N9K-C93108TC-FX
- N9K-C93180YC-FX

The following Fibre Channel small form-factor pluggable (SFP) transceivers are supported:

- DS-SFP-FC8G-SW: 2/4/8G (2G is not a supported FC NPV port speed)
- DS-SFP-FC16G-SW: 4/8/16G (not compatible when FC NPV port speed is 32G)
- DS-SFP-FC32G-SW: 8/16/32G (not compatible when FC NPV port speed is 4G)

The supported NPIV core switches are the Cisco Nexus 5000 series, Nexus 6000 series, Nexus 7000 series (FCoE), and Cisco MDS 9000 series multilayer switches.

Fibre Channel N-Port Virtualization Interoperability

The following table lists third party products with which the Fibre Channel N-port virtualization (FC NPV) feature of Cisco Application Policy Infrastructure Controller (APIC) was tested for interoperability.

Table 1: Third Party Products That Are Supported With FC NPV

Third Party Switch Vendor	Brocade
Third Party Hardware Model	DS-6620B
Third Party Software Release	8.2.1a
Cisco NX-OS Release	14.1(1) and later
Cisco Nexus 9000 Model	N9K-C93180YC-FX
Interoperability Mode	NA (NPV)
Cisco SFP Module	DS-SFP-FC32G-SW
Third Party SFP Module	Brocade-32G

Fibre Channel NPV GUI Configuration

Configuring a Native Fibre Channel Port Profile Using the GUI

This procedure configures a set of native Fibre Channel (FC) F ports for connecting to Fibre Channel hosts, such as servers.

To simplify the configuration, this procedure uses the **Configure an Interface, PC, and vPC** wizard.

Procedure

Step 1 On the APIC menu bar, navigate to **Fabric > Access Policies > Quickstart** and click *Configure an interface, PC, and vPC*.

Step 2 In the **Configured Switch Interfaces** toolbar, click + to create a switch profile. Perform the following actions:

This switch profile configures your server host ports. Another switch profile configures your uplink ports.

a) From the **Switches** drop-down list, choose your NPV leaf switch.

This action automatically creates a leaf switch profile. You can accept or change the name of the leaf switch profile in the **Switch Profile Name** text box.

b) Click the large green + on the ports drawing to open more interface settings.

c) For **Interface Type**, select **FC** to specify Fibre Channel host interface ports (F ports).

d) For **Interfaces**, enter a port range for the FC ports.

Only one contiguous range of ports can be converted to FC ports. This range must be a multiple of 4 ending with a port number that is a multiple of 4 (for example, 1-4, 1-8, and 21-32 are valid ranges).

This action creates an interface selector policy. You can accept or change the name of the policy in the **Interface Selector Name** text box.

Note Port conversion from Ethernet to FC requires a reload of the switch. After the interface policy is applied, a notification alarm appears in the GUI, prompting you to reload the switch. During a switch reload, communication to the switch is interrupted, resulting in timeouts when trying to access the switch.

e) From the **Policy Group Name** drop-down list, select **Create FC Interface Policy Group**.

f) In the **Create FC Interface Policy Group** dialog box, type a name in the **Name** field.

g) In the **Fibre Channel Interface Policy** drop-down list, select **Create Fibre Channel Interface Policy**.

h) In the **Create Fibre Channel Interface Policy** dialog box, type a name in the **Name** field and configure the following settings:

Field	Setting
Port Mode	For host interfaces, select F .
Trunk Mode	For host interfaces, select trunk-off .
Speed	Select auto (default).

Field	Setting
Auto Max Speed	Auto Max Speed configuration is applicable only when speed is auto . Auto Max Speed is to limit maximum speed when speed is in auto mode.
Receive Buffer Credit	Select 64 .

- i) Click **Submit** to save the Fibre Channel interface policy and return to the **Create FC Interface Policy Group** dialog box.
 - j) From the **Attached Entity Profile** drop-down list, choose **Create Attachable Access Entity Profile**.
The attachable entity profile option specifies the interfaces where the leaf access port policy is deployed.
 - k) In the **Name** field, enter a name for the attachable entity policy.
 - l) In the **Domains (VMM, Physical, or External) To Be Associated To Interfaces** toolbar, click + to add a domain profile.
 - m) From the **Domain Profile** drop-down list, choose **Create Fibre Channel Domain**.
 - n) In the **Name** field, enter a name for the Fibre Channel domain.
 - o) From the **VSAN Pool** drop-down list, choose **Create VSAN Pool**.
 - p) In the **Name** field, enter a name for the VSAN pool.
 - q) In the **Encap Blocks** toolbar, click + to add a VSAN range.
 - r) In the **Create VSAN Ranges** dialog box, enter **From** and **To** VSAN numbers.
 - s) For **Allocation Mode**, select **Static Allocation** and click **OK**.
 - t) In the **Create VSAN Ranges** dialog box, click **Submit**.
 - u) In the **Create Fibre Channel Domain** dialog box, click **Submit**.
- Note** In the Fibre Channel Domain, when using native FC ports instead of FCoE, it is not necessary to configure a VLAN pool or VSAN attributes.
- v) In the **Create Attachable Access Entity Profile** dialog box, click **Update** to select the Fibre Channel domain profile and click **Submit**.
 - w) In the **Create FC Policy Group** dialog box, click **Submit**.
 - x) In the **Configure Interface, PC, and vPC** dialog box, click **Save** to save this switch profile for your server host ports.

Note Port conversion from Ethernet to FC requires a reload of the switch. After the interface policy is applied, a notification alarm appears in the GUI, prompting you to reload the switch. During a switch reload, communication to the switch is interrupted, resulting in timeouts when trying to access the switch.

Note When you change a port profile on a switch, for example reconfigure an uplink as downlink and reload the switch, communication to the switch is interrupted until the switch gets its configuration from Cisco APIC.

In **Fabric > Access Policies > Switches > Leaf Switches > Profiles > name**, the Fibre Channel port profile appears in the **Associated Interface Selector Profiles** list in the **Leaf Profiles** work pane.

What to do next

- Configure a Fibre Channel uplink connection profile.
- Deploy the server ports and uplink ports in a tenant to connect to a Fibre Channel core switch.

Configuring a Native FC Port Channel Profile Using the GUI

This procedure configures a native Fibre Channel port channel (FC PC) profile for an uplink connection to a Fibre Channel core switch.



Note This procedure can also be performed using the **Configure Interface, PC, and vPC** wizard.

Before you begin

Configure your uplink connections, including an attachable entity profile.

Procedure

Step 1 Expand **Fabric > Access Policies > Interfaces > Leaf Interfaces > Profiles**.

Step 2 Right click **Profiles** and click **Create Leaf Interface Profile**.

Step 3 In the **Create Leaf Interface Profile** dialog box, perform the following steps:

- In the **Name** field, enter a name for the leaf interface profile.
- In the **Interface Selectors** toolbar, click + to open the **Create Access Port Selector** dialog box.
- In the **Name** field, enter a name for the port selector.
- In the **Interface IDs** field, enter a port range for the FC PC ports.

The port channel can have a maximum of 16 ports.

Only one contiguous range of ports can be converted to FC ports. This range must be a multiple of 4 ending with a port number that is a multiple of 4 (for example, 1-4, 1-8, and 21-32 are valid ranges).

Note Port conversion from Ethernet to FC requires a reload of the switch. After the interface policy is applied, a notification alarm appears in the GUI, prompting you to reload the switch manually. During a switch reload, communication to the switch is interrupted, resulting in timeouts when trying to access the switch.

- From the **Interface Policy Group** drop-down list, choose **Create FC PC Interface Policy Group**.
- In the **Name** field, enter a name for the FC PC interface policy group.
- From the **Fibre Channel Interface Policy** drop-down list, choose **Create Fibre Channel Interface Policy**.
- In the **Name** field, enter a name for the FC PC interface policy.
- In the **Create Interface FC Policy** dialog box, type a name in the **Name** field and configure the following settings:

Field	Setting
Port Mode	For uplink interfaces, select NP .

Field	Setting
Trunk Mode	For uplink interfaces, select trunk-on .

- j) Click **Submit** to save the FC PC interface policy and return to the **Create FC PC Interface Policy Group** dialog box.
- k) From the **Port Channel Policy** drop-down list, choose **Create Port Channel Policy**.
- l) In the **Name** field, enter a name for the port channel policy.

The other settings in this menu can be ignored.

- m) Click **Submit** to save the port channel policy and return to the **Create FC PC Interface Policy Group** dialog box.
- n) From the **Attached Entity Profile** drop-down list, choose the existing attachable entity profile.
- o) Click **Submit** to return to the **Create Access Port Selector** dialog box.
- p) Click **OK** to return to the **Create Leaf Interface Profile** dialog box.
- q) Click **OK** to return to the **Leaf Interfaces - Profiles** work pane.

Step 4 Expand **Fabric > Access Policies > Switches > Leaf Switches > Profiles**.

Step 5 Right click the leaf switch profile that you created and click **Create Interface Profile**.

Step 6 In the **Create Interface Profile** dialog box, perform the following steps:

- a) From the **Interface Select Profile** drop-down list, choose the leaf interface profile that you created for the port channel.
- b) Click **Submit** to return to the **Leaf Interfaces - Profiles** work pane.

Note Port conversion from Ethernet to FC requires a reload of the switch. After the interface policy is applied, a notification alarm appears in the GUI, prompting you to reload the switch. During a switch reload, communication to the switch is interrupted, resulting in timeouts when trying to access the switch.

In **Fabric > Access Policies > Switches > Leaf Switches > Profiles > name**, the FC port channel profile appears in the **Associated Interface Selector Profiles** list in the work pane.

What to do next

Deploy the server ports and uplink ports in a tenant to connect to a Fibre Channel core switch.

Deploying Fibre Channel Ports

This procedure activates the Fibre Channel server host ports and uplink ports.

Before you begin

- Configure Fibre Channel (FC) server host port profiles (F ports).
- Configure FC uplink port profiles (NP or TNP ports).
- Configure a leaf switch profile that includes two associated interface selector profiles — one for host ports and one for uplink ports.

Procedure

- Step 1** Expand **Tenants > Tenant *name* > Application Profiles**
- If the tenant does not exist, you must create a tenant.
- Step 2** Right click **Application Profiles**, click **Create Application Profile**, and perform the following actions:
- In the **Name** field, enter a name for the application profile.
 - Click **Submit**.
- Step 3** Expand **Tenants > Tenant *name* > Application Profiles > *name* > Application EPGs**
- Step 4** Right click **Application EPGs** and click **Create Application EPG**, and perform the following actions:
- Step 5** In the **Create Application EPG** dialog box, perform the following actions:
- In the **Name** field, enter a name for the application EPG.
 - Configure the following settings:
- | Field | Setting |
|-------------------------------|----------------------------|
| Intra EPG Isolation | Select Unenforced . |
| Preferred Group Member | Select Exclude . |
| Flood on Encapsulation | Select Disabled . |
- From the **Bridge Domain** drop-down list, select **Create Bridge Domain**.
 - In the **Name** field, enter a name for the bridge domain.
 - For **Type**, select **fc** to specify a Fibre Channel bridge domain.
 - From the **VRF** drop-down list, select **Create VRF**.
 - In the **Name** field, enter a name for the VRF.
 - Click **Submit** to return to the **Create Bridge Domain** dialog box.
 - Click **Next**, then **Next**, then **Finish** to return to the **Create Application EPG** dialog box.
 - Click **Finish**.
- Step 6** Expand **Tenants > Tenant *name* > Application Profiles > *name* > Application EPGs > *name* > Domains (VMs and Bare-Metals)**.
- Step 7** Right click **Domains (VMs and Bare-Metals)** and click **Add Fibre Channel Domain Association**, and perform the following actions:
- From the **Fibre Channel Domain Profile** drop-down list, select the Fibre Channel domain that you created when you configured your host ports.
 - Click **Submit**.
- Step 8** Expand **Tenants > Tenant *name* > Application Profiles > *name* > Application EPGs > *name* > Fibre Channel (Paths)** and perform the following actions:
- This step deploys the server host ports.
- Right click **Fibre Channel (Paths)** and click **Deploy Fibre Channel**.
 - In the **Path Type** control, click **Port**.
 - From the **Node** drop-down list, choose the leaf switch.
 - From the **Path** drop-down list, choose the leaf switch port that is configured as a server host port.
 - In the **VSAN** field, enter the port VSAN.

- f) In the **VSAN Mode** control, click **Native**.
- g) Verify that the **Type** is fcoe.
- h) (Optional) If you require a traffic map, use the **Pinning Label** drop-down list.

Note If multiple uplink ports are available and you want this host port to always direct its FLOGI to a specific uplink, you can create a pinning profile (traffic map) to associate the host port to the uplink port. Otherwise, hosts are load-balanced among the available uplink ports.

- i) Click **Submit**.
- j) Repeat from **Step a** for each Fibre Channel host port.

Step 9 Expand **Tenants > Tenant name > Application Profiles > name > Application EPGs > name > Fibre Channel (Paths)** and perform the following actions:

This step deploys the uplink port channel.

- a) Right click **Fibre Channel (Paths)** and click **Deploy Fibre Channel**.
- b) In the **Path Type** control, click **Direct Port Channel**.
- c) From the **Path** drop-down list, choose the uplink port channel.
- d) In the **VSAN** field, enter the port default VSAN.
- e) In the **VSAN Mode** control, click **Native** for a port VSAN or **Regular** for a trunk VSAN.
- f) Verify that the **Type** is fcoe.
- g) Click **Submit**.
- h) Repeat from **Step a** for each Fibre Channel uplink port or port channel.

Configuring a Traffic Map for a Fibre Channel Port

In an application in which multiple uplink ports are available, server traffic by default is load-balanced among the available uplink ports. In some cases, it might be necessary to have a server send its login request (FLOGI) to one or more specific uplink ports or port channels. In such cases, you can create a pinning profile (traffic map) to associate the server port to those uplink ports or port channels.

This procedure assumes that you have already configured one or more server ports and one or more uplink ports or port channels. Because the server ports have already been configured, you must first shut (disable) any server port that is to be mapped to an uplink. After configuring the traffic map, re-enable the port.

Before you begin

This procedure assumes that the following items are already configured:

- Server ports (F ports) and uplink ports or port channels (NP ports)
- A tenant, including an application profile and application EPG



Note Before creating a pinning profile (traffic map), you must shut the server port that is to be mapped to an uplink.

Procedure

-
- Step 1** In the **Fabric > Inventory > Pod *n* > Leaf *n* > Interfaces > FC Interfaces** work pane, select and disable the server interface port that is to be mapped to an uplink.
- Step 2** Expand **Tenants > Tenant *name* > Application Profiles > *application profile name* > Application EPGs > EPG *name* > Fibre Channel (Paths)** and perform the following actions:
- Right click **Fibre Channel (Paths)** and click **Deploy Fibre Channel**.
 - In the **Path Type** control, click **Port**.
 - From the **Node** drop-down list, choose the leaf switch.
 - From the **Path** drop-down list, choose the server port that is to be mapped to a specific uplink port.
 - In the **VSAN** field, enter the port default VSAN.
 - In the **VSAN Mode** control, click **Native**.
 - Verify that the **Type** is **fcoe**.
 - From the **Pinning Label** drop-down list, choose **Create Pinning Profile**.
 - In the **Name** field, enter a name for the traffic map.
 - In the **Path Type** control, click **Port** to connect to a single NP uplink port or **Direct Port Channel** to connect to an FC port channel.
- If you choose **Port** for the path type, you must also choose the leaf switch from the **Node** drop-down list that appears.
- If you choose **Direct Port Channel** for the path type, you must also choose the FC PC you have defined in Interface Policy Group.
- From the **Path** drop-down list, choose the uplink port or port channel to which the server port will be mapped.
 - Click **Submit** to return to the **Deploy Fibre Channel** dialog box.
 - Click **Submit**.
- Step 3** In the **Fabric > Inventory > Pod *n* > Leaf *n* > Interfaces > FC Interfaces** work pane, select and re-enable the server interface port that is mapped to an uplink.
-

Fibre Channel NPV NX-OS-Style CLI Configuration

Configuring Fibre Channel Interfaces Using the CLI

On an NPV-enabled leaf switch, you can convert universal ports to Fibre Channel (FC) ports. The FC ports can be either F ports or NP ports, and NP ports can form a port channel.

Procedure

-
- Step 1** Convert a range of ports from Ethernet to Fibre Channel.
- Example:**

```
apicl(config)# leaf 101
apicl(config-leaf)# slot 1
apicl(config-leaf-slot)# port 1 12 type fc
```

This example converts ports 1/1-12 on leaf 101 to Fibre Channel ports. The **[no]** form of the **port type fc** command converts the ports from Fibre Channel back to Ethernet.

Note The conversion of ports takes place only after a reboot of the leaf switch.

Currently only one contiguous range of ports can be converted to FC ports, and this range must be a multiple of 4 ending with a port number that is a multiple of 4 (for example, 1-4, 1-8, or 21-24).

Step 2 Configure all Fibre channel interfaces.

Example:

```
apicl(config)# leaf 101
apicl(config-leaf)# interface fc 1/1
apicl(config-leaf-fc-if)# switchport mode [f | np]
apicl(config-leaf-fc-if)# switchport rxbbcredit <16-64>
apicl(config-leaf-fc-if)# switchport speed [16G | 32G | 4G | 8G | auto | unknown]
apicl(config-leaf-fc-if)# switchport trunk-mode [ auto | trunk-off | trunk-on | un-init]
apicl(config-leaf-fc-if)# switchport [trunk allowed] vsan <1-4093> tenant <name> \
                                application <name> epg <name>
```

Note FC host interfaces (F ports) do not support a speed configuration of 8Gbps.

A FC interface can be configured in access mode or trunk mode. To configure the FC port in access mode, use the following command format:

Example:

```
apicl(config-leaf-fc-if)# switchport vsan 2 tenant t1 application a1 epg e1
```

To configure a FC port in trunk mode, use the following command format:

Example:

```
apicl(config-leaf-fc-if)# switchport trunk allowed vsan 4 tenant t1 application a1 epg e1
```

To configure a FC port channel, configure a FC port interface template and apply it to FC interfaces that will be members of the FC port-channel.

The port channel can have a maximum of 16 members.

Example:

```
apicl(config)# template fc-port-channel my-fc-pc
apicl(config-fc-po-ch-if)# lacp max-links 4
apicl(config-fc-po-ch-if)# lacp min-links 1
apicl(config-fc-po-ch-if)# vsan-domain member dom1
apicl(config-fc-po-ch-if)# exit
apicl(config)# leaf 101
apicl(config-leaf)# interface fc 1/1-2
apicl(config-leaf-fc-if)# fc-channel-group my-fc-pc
apicl(config-leaf-fc-if)# exit
apicl(config-leaf)# interface fc-port-channel my-fc-pc
```



```

apic1(config-leaf-fc-pc) # switchport mode [f | np]
apic1(config-leaf-fc-pc) # switchport rxbbcredit <16-64>
apic1(config-leaf-fc-pc) # switchport speed [16G | 32G | 4G | 8G | auto | unknown]
apic1(config-leaf-fc-pc) # switchport trunkmode [ auto | trunk-off | trunk-on | un-init]

```

Configuring Fibre Channel NPV Policies Using the CLI

Before you begin

Leaf switch ports to be used in an NPV application have been converted to Fibre Channel (FC) ports.

Procedure

Step 1 Create a template of a Fibre Channel F port policy group.

Example:

```

apic1(config) # template fc-policy-group my-fc-policy-group-f-ports
apic1(config-fc-pol-grp-if) # vsan-domain member dom1
apic1(config-fc-pol-grp-if) # switchport mode f
apic1(config-fc-pol-grp-if) # switchport trunk-mode trunk-off

```

You can configure other switchport settings, such as speed.

Step 2 Create a template of a Fibre Channel NP port policy group.

Example:

```

apic1(config) # template fc-policy-group my-fc-policy-group-np-ports
apic1(config-fc-pol-grp-if) # vsan-domain member dom1
apic1(config-fc-pol-grp-if) # switchport mode np
apic1(config-fc-pol-grp-if) # switchport trunk-mode trunk-on

```

You can configure other switchport settings, such as speed.

Step 3 Create a fabric-wide Fibre Channel policy.

Example:

```

apic1(config) # template fc-fabric-policy my-fabric-fc-policy
apic1(config-fc-fabric-policy) # fctimer e-d-tov 1000
apic1(config-fc-fabric-policy) # fctimer r-a-tov 5000
apic1(config-fc-fabric-policy) # fcoe fcmapi 0E:FC:01

```

Step 4 Create a Fibre Channel port channel policy.

Example:

```

apic1(config) # template fc-port-channel my-fc-pc
apic1(config-fc-po-ch-if) # lacp max-links 4
apic1(config-fc-po-ch-if) # lacp min-links 1

```

```
apic1(config-fc-po-ch-if) # vsan-domain member dom1
```

Step 5 Create a leaf-wide Fibre Channel policy group.

Example:

```
apic1(config) # template fc-leaf-policy my-fc-leaf-policy
apic1(config-fc-leaf-policy) # npv auto-load-balance disruptive
apic1(config-fc-leaf-policy) # fcoe fka-adv-period 10
```

Note The policy commands that are shown here are only examples, and are not mandatory settings.

Step 6 Create a leaf policy group.

```
apic1(config) # template leaf-policy-group lpg1
apic1(config-leaf-policy-group) # inherit fc-fabric-policy my-fabric-fc-policy
apic1(config-leaf-policy-group) # inherit fc-leaf-policy my-fc-leaf-policy
```

The leaf policy group is created by inheriting FC-related policies.

Step 7 Create a leaf profile to apply a leaf-policy-group to a leaf-group.

Example:

```
apic1(config) # leaf-profile my-leaf-profile
apic1(config-leaf-profile) # leaf-group my-leaf-group
apic1(config-leaf-group) # leaf 101
apic1(config-leaf-group) # leaf-policy-group lpg1
```

This example applies fabric-wide FC policies and leaf-wide FC policies that are grouped into a leaf policy group lpg1 to leaf 101.

Step 8 Create a leaf interface profile and apply a fc-policy-group to a set of FC interfaces.

Example:

```
apic1(config) # leaf-interface-profile my-leaf-interface-profile
apic1(config-leaf-if-profile) # leaf-interface-group my-leaf-interface-group
apic1(config-leaf-if-group) # fc-policy-group my-fc-policy-group-f-ports
apic1(config-leaf-if-group) # interface fc 1/1-10
```

Configuring an NPV Traffic Map Using the CLI

This procedure maps traffic coming from a FC/FCoE server (host) interface to a FC/FCoE external (uplink) interface configured in NP mode.

Before you begin

All server interfaces must be F ports and all uplink interfaces must be NP ports.

Procedure

Example:

```
apic1(config)# leaf 101
apic1(config-leaf)# npv traffic-map server-interface \
    { vfc <slot/port> | vfc-po <po-name> | fc <slot/port> } \
    label <name> tenant <tn> app <ap> ep <ep>
apic1(config-leaf)# npv traffic-map external-interface \
    { vfc <slot/port> | vfc-po <po-name> | fc <slot/port> } \
    tenant <tn> label <name>
```

Example:

```
apic1(config)# leaf 101
apic1(config-leaf)# npv traffic-map server-interface vfc 1/1 label serv1 tenant t1 app ap1
    ep ep1
apic1(config-leaf)# npv traffic-map external-interface vfc-po my-fc-pc tenant t1 label ext1
```

Fibre Channel NPV REST API Configuration

Configuring FC Connectivity Using the REST API

You can configure FC-enabled interfaces and EPGs accessing those interfaces using the FC protocol with the REST API.

Procedure

- Step 1** To create a VSAN pool, send a post with XML such as the following example. The example creates VSAN pool myVsanPool1 and specifies the range of VSANs to be included as vsan-50 to vsan-60:

Example:

```
https://apic-ip-address/api/mo/uni/infra/vsanns-[myVsanPool1]-static.xml

<fvnsVsanInstP allocMode="static" name="myVsanPool1">
  <fvnsVsanEncapBlk from="vsan-50" name="encap" to="vsan-60"/>
</fvnsVsanInstP>
```

- Step 2** To create a Fibre Channel domain, send a post with XML such as the following example. The example creates Fibre Channel domain (VSAN domain) myFcDomain1 and associates it with the VSAN pool myVsanPool1:

Example:

```
https://apic-ip-address/api/mo/uni/fc-myFcDomain1.xml

<fcDomP name="myFcDomain1">
  <fcRsVsanNs tDn="uni/infra/vsanns-[myVsanPool1]-static"/>
</fcDomP>
```

```
</fcDomP>
```

Step 3 To create an Attached Entity Policy (AEP) for the FC ports, send a post with XML such as the following example. The example creates the AEP myFcAEP1 and associates it with the Fibre Channel domain myFcDomain1:

Example:

```
https://apic-ip-address/api/mo/uni.xml
```

```
<polUni>
<infraInfra>
  <infraAttEntityP name="myFcAEP1">
    <infraRsDomP tDn="uni/fc-myFcDomain1"/>
  </infraAttEntityP>
</infraInfra>
</polUni>
```

Step 4 To create a FC interface policy and a policy group for server host ports, send a post with XML. This example executes the following requests:

- Creates a FC interface policy myFcHostIfPolicy1 for server host ports. These are F ports with no trunking.
- Creates a FC interface policy group myFcHostPortGroup1 that includes the FC host interface policy myFcHostIfPolicy1.
- Associates the policy group to the FC interface policy to convert these ports to FC ports.
- Creates a host port profile myFcHostPortProfile.
- Creates a port selector myFcHostSelector that specifies ports in range 1/1-8.
- Creates a node selector myFcNode1 that specifies leaf node 104.
- Creates a node selector myLeafSelector that specifies leaf node 104.
- Associates the host ports to the leaf node.

Example:

```
https://apic-ip-address/api/mo/uni.xml
```

```
<polUni>
<infraInfra>
  <fcIfPol name="myFcHostIfPolicy1" portMode="f" trunkMode="trunk-off" speed="auto"/>

  <infraFuncP>
    <infraFcAccPortGrp name="myFcHostPortGroup1">
      <infraRsFcL2IfPol tnFcIfPolName="myFcHostIfPolicy1" />
    </infraFcAccPortGrp>
  </infraFuncP>
  <infraAccPortP name="myFcHostPortProfile">
    <infraHPortS name="myFcHostSelector" type="range">
      <infraPortBlk name="myHostPorts" fromCard="1" toCard="1" fromPort="1"
toPort="8" />
      <infraRsAccBaseGrp tDn="uni/infra/funcprof/fcaccportgrp-myFcHostPortGroup1"
/>
    </infraHPortS>
  </infraAccPortP>
  <infraNodeP name="myFcNode1">
    <infraLeafS name="myLeafSelector" type="range">
      <infraNodeBlk name="myLeaf104" from_="104" to_="104" />
    </infraLeafS>
  </infraNodeP>
</infraInfra>
</polUni>
```

```

        </infraLeafS>
        <infraRsAccPortP tDn="uni/infra/accportprof-myHostPorts" />
    </infraNodeP>
</infraInfra>
</polUni>

```

Note When this configuration is applied, a switch reload is required to bring up the ports as FC ports. Currently only one contiguous range of ports can be converted to FC ports, and this range must be multiple of 4 ending with a port number that is multiple of 4. Examples are 1-4, 1-8, or 21-24.

Step 5 To create a FC uplink port interface policy and a policy group for uplink port channels, send a post with XML. This example executes the following requests:

- Creates a FC interface policy myFcUplinkIfPolicy2 for uplink ports. These are NP ports with trunking enabled.
- Creates a FC interface bundle policy group myFcUplinkBundleGroup2 that includes the FC uplink interface policy myFcUplinkIfPolicy2.
- Associates the policy group to the FC interface policy to convert these ports to FC ports.
- Creates an uplink port profile myFcUplinkPortProfile.
- Creates a port selector myFcUplinkSelector that specifies ports in range 1/9-12.
- Associates the host ports to the leaf node 104.

Example:

<https://apic-ip-address/api/mo/uni.xml>

```

<polUni>
  <infraInfra>
    <fcIfPol name="myFcUplinkIfPolicy2" portMode="np" trunkMode="trunk-on" speed="auto"/>

    <infraFuncP>
      <infraFcAccBndlGrp name="myFcUplinkBundleGroup2">
        <infraRsFcL2IfPol tnFcIfPolName="myFcUplinkIfPolicy2" />
      </infraFcAccBndlGrp>
    </infraFuncP>
    <infraAccPortP name="myFcUplinkPortProfile">
      <infraHPortS name="myFcUplinkSelector" type="range">
        <infraPortBlk name="myUplinkPorts" fromCard="1" toCard="1" fromPort="9"
toPort="12" />
      <infraRsAccBaseGrp
tDn="uni/infra/funcprof/fcaccportgrp-myFcUplinkBundleGroup2" />
    </infraHPortS>
    </infraAccPortP>
    <infraNodeP name="myFcNode1">
      <infraLeafS name="myLeafSelector" type="range">
        <infraNodeBlk name="myLeaf104" from_="104" to_="104" />
      </infraLeafS>
      <infraRsAccPortP tDn="uni/infra/accportprof-myUplinkPorts" />
    </infraNodeP>
  </infraInfra>
</polUni>

```

Note When this configuration is applied, a switch reload is required to bring up the ports as FC ports. Currently only one contiguous range of ports can be converted to FC ports, and this range must be multiple of 4 ending with a port number that is multiple of 4. Examples are 1-4, 1-8, or 21-24.

Step 6 To create the tenant, application profile, EPG and associate the FC bridge domain with the EPG, send a post with XML such as the following example. The example creates a bridge domain myFcBD1 under a target tenant configured to support FC and an application EPG epg1. It associates the EPG with Fibre Channel domain myFcDomain1 and a Fibre Channel path to interface 1/7 on leaf switch 104. Each interface is associated with a VSAN.

Example:

`https://apic-ip-address/api/mo/uni/tn-tenant1.xml`

```
<fvTenant name="tenant1">
  <fvCtx name="myFcVRF"/>
  <fvBD name="myFcBD1" type="fc">
    <fvRsCtx tnFvCtxName="myFcVRF"/>
  </fvBD>
  <fvAp name="ap1">
    <fvAEPg name="epg1">
      <fvRsBd tnFvBDName="myFcBD1"/>
      <fvRsDomAtt tDn="uni/fc-myFcDomain1"/>
      <fvRsFcPathAtt tDn="topology/pod-1/paths-104/pathep-[fc1/1]" vsan="vsan-50"
vsanMode="native"/>
      <fvRsFcPathAtt tDn="topology/pod-1/paths-104/pathep-[fc1/2]" vsan="vsan-50"
vsanMode="native"/>
    </fvAEPg>
  </fvAp>
</fvTenant>
```

Step 7 To create a traffic map to pin server ports to uplink ports, send a post with XML such as the following example. The example creates a traffic map to pin server port vFC 1/47 to uplink port FC 1/7:

Example:

`https://apic-ip-address/api/mo/uni/tn-tenant1.xml`

```
<fvTenant name="tenant1">
  <fvAp name="ap1">
    <fvAEPg name="epg1">
      <fvRsFcPathAtt tDn="topology/pod-1/paths-104/pathep-[eth1/47]" vsan="vsan-50"
vsanMode="native">
        <fcPinningLbl name="label1"/>
      </fvRsFcPathAtt>
    </fvAEPg>
  </fvAp>
</fvTenant>
```

`https://apic-ip-address/api/mo/uni/tn-vfc_t1.xml`

```
<fvTenant name="tenant1">
  <fcPinningP name="label1">
    <fcRsPinToPath tDn="topology/pod-1/paths-104/pathep-[fc1/7]"/>
  </fcPinningP>
</fvTenant>
```

Note If traffic map pinning is configured for the first time, the server host port must be shut before configuring the first traffic map.
