Configuring FabricPath Interfaces

This chapter describes how to configure the FabricPath interfaces on the Cisco NX-OS devices.

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at https://tools.cisco.com/bugsearch/ and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information"chapter or the Feature History table in this chapter.

Information About FabricPath Interfaces

Note   You must have an F Series module installed in the Cisco Nexus 7000 Series device to run FabricPath.

FabricPath Interfaces

After you enable FabricPath on the devices that you are using, you can configure an Ethernet interface or a port-channel interface as a FabricPath interface. If one member of the port channel is in FabricPath mode, all the other members will be in FabricPath mode. After you configure the interface as a FabricPath interface, it automatically becomes a trunk port, capable of carrying traffic for multiple VLANs. You can also configure all the ports on the F Series module as FabricPath interfaces simultaneously.

The following interface modes carry traffic for the following types of VLANs:

- Interfaces on the F Series modules that are configured as FabricPath interfaces can carry traffic only for FP VLANs.

- Interfaces on the F Series modules that are not configured as FabricPath interfaces carry traffic for the following:

  - FP VLANs

  - Classical Ethernet (CE) VLANs

- Interfaces on the M Series modules carry traffic only for CE VLANs.

Note See "Configuring FabricPath Forwarding," for information about FP and CE VLANs.

The FabricPath interfaces connect only to other FabricPath interfaces within the FabricPath network. These FabricPath ports operate on the information in the FabricPath headers and Layer 2 Intermediate System-to-Intermediate System (IS-IS) only, and they do not run STP. These ports are aware only of FP VLANs; they are unaware of any CE VLANs. By default, all VLANs are allowed on a trunk port, so the FabricPath interface carries traffic for all FP VLANs.

Note You cannot configure FabricPath interfaces as shared interfaces. See the *Cisco NX-OS FCoE Configuration Guide* for Cisco Nexus 7000 and Cisco MDS 9500 for information on shared interfaces.

STP and the FabricPath Network

Note The Layer 2 gateway switches, which are on the edge between the CE and the FabricPath network, must be the root for all STP domains that are connected to a FabricPath network.

The Spanning Tree Protocol (STP) domains do not cross into the FabricPath network (see the figure below).

Figure 1: STP Boundary Termination at FabricPath Network Border
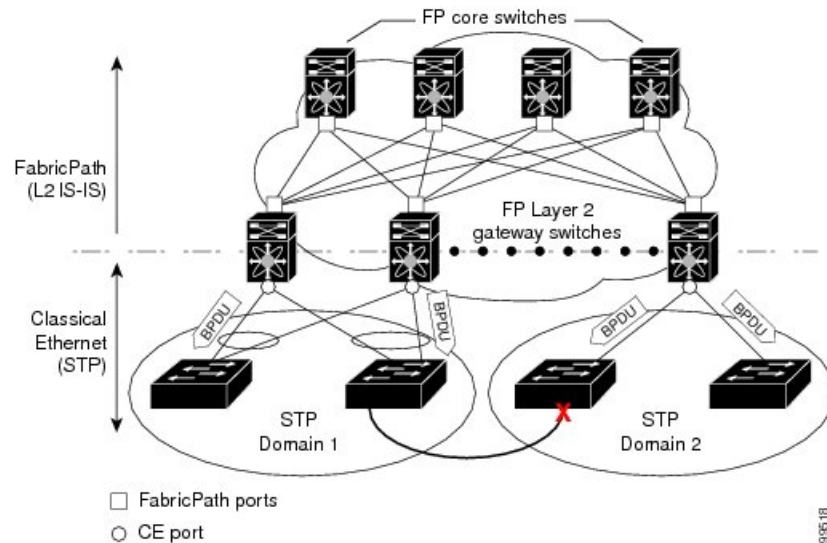
You must configure the FabricPath Layer 2 gateway device to have the lowest STP priority of all the devices in the STP domain to which it is attached. You must also configure all the FabricPath Layer 2 gateway devices that are connected to one FabricPath network to have the same priority. The system assigns the bridge ID for the Layer 2 gateway devices from a pool of reserved MAC addresses.

To have a loop-free topology for the CE/FabricPath hybrid network, the FabricPath network automatically displays as a single bridge to all connected CE devices.

Note   You must set the STP priority on all FabricPath Layer 2 gateway switches to a value low enough to ensure that they become root for any attached STP domains.

Other than configuring the STP priority on the FabricPath Layer 2 gateway switches, you do not need to configure anything for the STP to work seamlessly with the FabricPath network. Only connected CE devices form a single STP domain. Those CE devices that are not interconnected form separate STP domains (see the figure above).

All CE interfaces should be designated ports, which occurs automatically, or they are pruned from the active STP topology. If the system does prune any port, the system returns a syslog message. The system clears the port again only when that port is no longer receiving superior BPDUs.

The FabricPath Layer 2 gateway switch also propagates the Topology Change Notifications (TCNs) on all its CE interfaces.

The FabricPath Layer 2 gateway switches terminate STP. The set of FabricPath Layer 2 gateway switches that are connected by STP forms the STP domain. Because there can be many FabricPath Layer 2 gateway switches attached to a single FabricPath network, there might also be many separate STP domains (see the figure above). The devices in the separate STP domains need to know the TCN information only for the domain to which they belong. You can configure a unique STP domain ID for each separate STP domain that connects to the same FabricPath network. The Layer 2 Intermediate System-to-Intermediate System (IS-IS) messages carry the TCNs across the FabricPath network. Only those FabricPath Layer 2 gateway switches in the same STP domain as the TCN message need to act and propagate the message to connected CE devices.

When a FabricPath Layer 2 gateway switch receives a TCN for the STP domain it is part of, it takes the following actions:

- Flushes all remote MAC addresses for that STP domain and the MAC addresses on the designated port.

- Propagates the TCN to the other devices in the specified STP domain.

The devices in the separate STP domains need to receive the TCN information and then flush all remote MAC addresses that are reachable by the STP domain that generated the TCN information.

vPC+

A virtual port channel+ (vPC+) domain allows a classical Ethernet (CE) vPC domain and a Cisco FabricPath cloud to interoperate. A vPC+ also provides a First Hop Routing Protocol (FHRP) active-active capability at the FabricPath to Layer 3 boundary.

> **Note**
> - vPC+ is an extension to virtual port channels (vPCs) that run CE only (see the "Configuring vPCs" chapter in the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide*). You cannot configure a vPC+ domain and a vPC domain in the same VDC.
>
> - In a vPC+ system running 7.2(0)D1(0.444S4), the mroutes (both local and remote) between the two vPC+ peers do not sync as vPC+ does not support dual DR.

A vPC+ domain enables Cisco Nexus 7000 Series enabled with FabricPath devices to form a single vPC+, which is a unique virtual switch to the rest of the FabricPath network. You configure the same domain on each device to enable the peers to identify each other and to form the vPC+. Each vPC+ has its own virtual switch ID.

Enabling the vPC peer switch feature is not necessary when you are using vPC+. All FabricPath edge switches use a common reserved bridge ID (BID c84c.75fa.6000) when sending BPDUs on CE edge ports.

A vPC+ must still provide active-active Layer 2 paths for dual-homed CE devices or clouds, even though the FabricPath network allows only 1-to-1 mapping between the MAC address and the switch ID. vPC+ creates a unique virtual switch to the FabricPath network (see the figure below).

Figure 2: vPC/vPC+
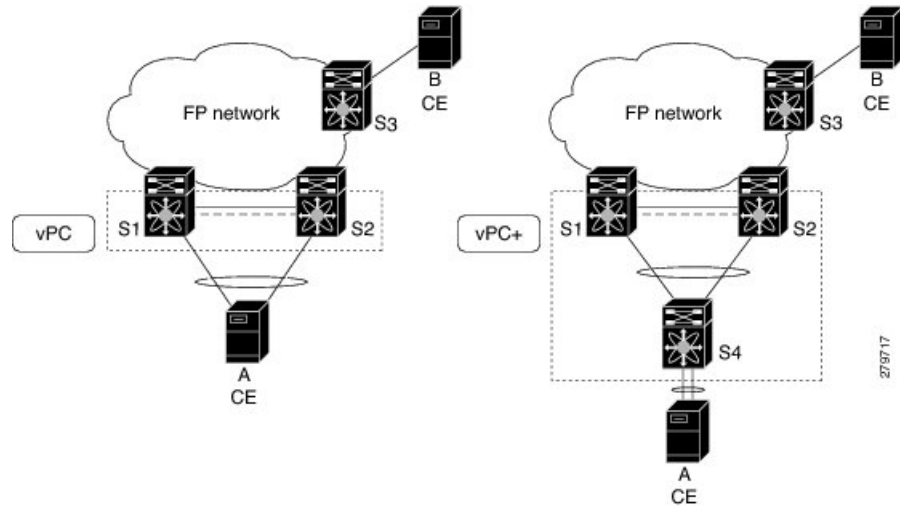
The FabricPath switch ID for the virtual switch becomes the outer source MAC address (OSA) in the FabricPath encapsulation header. Each vPC+ domain must have its own virtual switch ID.

Layer 2 multipathing is achieved by emulating a single virtual switch. Packets forwarded from host A to host B are tagged with the MAC address of the virtual switch as the transit source, and traffic from host B to host A is now load balanced.

You must have all interfaces in the vPC+ peer link as well as all the downstream vPC+ links on an F Series module with FabricPath enabled. The vPC+ downstream links will be FabricPath edge interfaces, which connect to the CE hosts.

The vPC+ virtual switch ID is used to assign the FabricPath Outer Source Address (OSA) to the FabricPath vPC+ peer devices (see "Configuring FabricPath Switching," for information about FabricPath encapsulation). You must assign the same switch ID to each of the two vPC+ peer devices so the peer link can form.

The F1 Series modules have only Layer 2 interfaces. To use routing with a vPC+, you must have an M Series module inserted into the same Cisco Nexus 7000 Series chassis. The system then performs proxy routing using both the N7K-F132-15 module and the M Series modules in the chassis (see the Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide for information on proxy routing with the F1 Series modules).

The First Hop Routing Protocols (FHRPs) and the Hot Standby Routing Protocol (HSRP) interoperate with a vPC+. You should dual-attach all Layer 3 devices to both vPC+ peer devices.

Note You must enable the Layer 3 connectivity from each vPC+ peer device by configuring a VLAN network interface for the same VLAN from both devices.

The primary FHRP device responds to ARP requests, even though the secondary vPC+ device also forwards the data traffic. Both the primary and secondary vPC+ devices forward traffic, but only the primary FHRP device responds to ARP requests.

To simplify initial configuration verification and vPC+/HSRP troubleshooting, you can configure the primary vPC+ peer device with the FHRP active router highest priority.

In addition, you can use the **priority** command in the if-hsrp configuration mode to configure failover thresholds when a group state that is enabled on a vPC+ peer is in standby or in listen state. You can configure lower and upper thresholds to prevent the group state flap, if there is an interface flap (this feature is useful when there is more than one tracking object per group).

When the primary vPC+ peer device fails over to the secondary vPC+ peer device, the FHRP traffic continues to flow seamlessly.

You should configure a separate Layer 3 link for routing from the vPC+ peer devices, rather than using a VLAN network interface for this purpose.

We do not recommend that you configure the burnt-in MAC address option (use-bia) for Hot Standby Router Protocol (HSRP) or manually configure virtual MAC addresses for any FHRP protocol in a vPC+ environment because these configurations can adversely affect the vPC+ load balancing. The HSRP use-bia is not supported with a vPC+. When you are configuring custom MAC addresses, you must configure the same MAC address on both vPC+ peer devices.

You can configure a restore timer that delays the vPC+ coming back up until after the peer adjacency forms and the VLAN interfaces are back up. This feature allows you to avoid packet drops if the routing tables do not converge before the vPC+ is once again passing traffic.

Use the **delay restore** command to configure this feature.

> **Note** If a data center outage occurs and you enable HSRP before the vPC+ successfully comes up, traffic loss can occur. You need to enable an HSRP delay to give the vPC time to stabilize. If you enable both an HSRP delay and a preemption delay, the Cisco Nexus 7000 Series devices allow Layer 2 switching only after both timers expire.
>
> The delay option is available only with HSRP. If you use any other FHRP, traffic loss is still possible.

See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide*, for more information about FHRPs and routing.

Anycast HSRP

Beginning with Release 6.2(2), Cisco NX-OS provides a way to facilitate further scalability at the spine layer giving support for more than two nodes. You can create an anycast bundle that is an association between a set of VLANs and an anycast switch ID. An anycast switch ID is the same as an emulated switch ID except the anycast switch ID is shared across more than two gateways. The set of VLANs or HSRP group elects an active router and a standby router. The remaining routers in the group are in listen state.

The active HSRP router advertises the anycast switch ID as the source switch ID in FabricPath IS-IS. The leaf switches learn that the anycast switch ID is reachable by all of the routers in the group.

For Release 6.2(2), Cisco NX-OS supports only four gateways. All the first-hop gateways at the spine layer must function in active-active forwarding mode. IP packets are received by any of the spine switches with the destination set as the gateway MAC address and these packets are terminated and locally forwarded.

Note   Prior to Cisco NX-OS Release 6.2(8), FabricPath Layer 2 IS-IS advertised the anycast switch ID even with the overload bit set, which would incur longer convergence times for selected nodes. Beginning with Cisco NX-OS Release 6.2(8), the system does not advertise the configured anycast switch ID while the overload bit is set, which effectively improves the convergence times.

Designated Forwarder

Beginning with Release 6.0, Cisco NX-OS provides a way to control two peers to be partial designated forwarders when both vPC paths are up. When this control is enabled, each peer can be the designated forwarder for multi destination southbound packets for a disjoint set of RBHs/FTAGs (depending on the hardware). The designated forwarder is negotiated on a per-vPC basis.

This control is enabled with the **fabricpath multicast load-balance** CLI command. This command is configured under vPC domain mode. For example:

```
switch(config)# vpc domain 1
switch(config-vpc-domain)# fabricpath multicast load-balance
```

There are three designated forwarder states for a vPC port:

- All—If the local vPC leg is up and the peer vPC is not configured or down, the local switch is the designated forwarder for all RBHs/FTAGs for that vPC.

- Partial—If the vPC path is up on both sides, each peer is the designated forwarder for half the RBHs or FTags. For the latter, the vPC port allows only the active FTags on that peer.

- None—If the local vPC path is down or not configured, the local switch does not forward any multi destination packets from this vPC path.

Only an F2 series module supports multicast load balancing. On an F1 series module, the configuration is supported, but load balancing does not occur.

Note   The **fabricpath multicast load-balance** command is required for configuring vPC+ with FEX ports.

High Availability

The FabricPath topologies retain their configuration through ISSU.

See the *Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide* for more information on high availability.

Virtual Device Contexts

You must install the FabricPath feature set before you enable FabricPath on the switch. See Configuring Feature Set for FabricPath for information on installing the FabricPath feature set.

Because of the multiple forwarding engines (FEs) on the F Series modules, the table below lists the port pairs and port sets that must be in the same VDC.

Table 1: Port Pairs and Port Sets for F Series Modules

Port Pairs for F1 Modules	Port Sets for F2 Modules
Ports 1 and 2	Ports 1, 2, 3, 4
Ports 3 and 4	Ports 5, 6, 7, 8
Ports 5 and 6	Ports 9, 10, 11, 12
Ports 7 and 8	Ports 13, 14, 15, 16
Ports 9 and 10	Ports 17, 18, 19, 20
Ports 11 and 12	Ports 21, 22, 23, 24
Ports 13 and 14	Ports 25, 26, 27, 28
Ports 15 and 16	Ports 29, 30, 31, 32
Ports 17 and 18	Ports 33, 34, 35, 36
Ports 19 and 20	Ports 37, 38, 39, 40
Ports 21 and 22	Ports 41, 42, 43, 44
Ports 23 and 24	Ports 45, 46, 47, 48
Ports 25 and 26	
Ports 27 and 28	
Ports 29 and 30	
Ports 31 and 32	

See the *Virtual Device Context Configuration Guide, Cisco DCNM for LAN*, for more information about VDCs.

Prerequisites for FabricPath

FabricPath forwarding has the following prerequisites:

- You should have a working knowledge of Classical Ethernet Layer 2 functionality.

- You must install the FabricPath feature set in the default and nondefault VDC before you enable FabricPath on the switch. See the Configuring Feature Set for FabricPath for complete information on installing and enabling the FabricPath feature set.

- The FabricPath feature set operation might cause the standby supervisor to reload if it is in an unstable state, such as following a service failure or powering up.

- You are logged onto the device.

- You are in the correct virtual device context (VDC). A VDC is a logical representation of a set of system resources. You can use the **switchto vdc** command with a VDC number.

- You are working on the F Series module.

Guidelines and Limitations for FabricPath Interfaces

FabricPath switching has the following configuration guidelines and limitations:

- FabricPath interfaces carry only FabricPath-encapsulated traffic.

- You enable FabricPath on each device before you can view or access the commands. Enter the **feature-set fabricpath** command to enable FabricPath on each device. See *Configuring Feature-Set for FabricPath* for complete information on installing and enabling the FabricPath feature set.

- STP does not run inside a FabricPath network.

- Set the STP priority value on all FabricPath Layer 2 gateway devices to 8192.
- The F Series modules do not support multiple SPAN destination ports or virtual SPAN. If a port on an F Series module is in a VDC and that VDC has multiple SPAN destination ports, that SPAN session is not brought up.

- The following guidelines apply to private VLAN configuration when you are running FabricPath:

    - All VLANs in a private VLAN must be in the same VLAN mode; either CE or FabricPath. If you attempt to put different types of VLANs into a private VLAN, these VLANs will not be active in the private VLAN. The system remembers the configurations, and if you change the VLAN mode later, that VLAN becomes active in the specified private VLAN.

    - FabricPath ports cannot be put into a private VLAN.

- The system does not support hierarchical static MAC addresses. That is, you cannot configure static FabricPath ODAs or OSAs; you can only configure CE static MAC addresses.

- On the F Series modules, user-configured static MAC addresses are programmed on all forwarding engines (FEs) that have ports in that VLAN.

- Pruning does not occur in a virtual port channel (vPC) domain. In a vPC domain, all switches receive multicast traffic, but only one switch forwards the traffic to the receiver.

- A single vPC+ domain between two VDCs on the same physical Cisco Nexus 7000 device is not supported.

- At least one FabricPath interface must be operational on a device for multidestination traffic to be forwarded on vPC+ member ports.

- Support for more than 244 vPC+ port channels (per vPC+ domain) is enabled with the **no port-channel limit** command.

    - Only VDCs that have an F2 series module can support more than 244 vPC+ port channels.

    - The **fabricpath multicast load-balance** command must be entered before the **no port-channel limit** command.

Note   The **no port-channel limit** command is not applicable with a FEX. A FEX can support more than 244 vPC+ port channels

- An anycast HSRP bundle provides the support for more than two nodes at the spine layer.

- An anycast HSRP bundle is supported only in HSRP version 2.

- Because of a limitation with an ASIC on the 32-port 1/10-Gigabit Ethernet F1 Series module, a packet that egresses from that module through both ports in FabricPath VLAN mode has an incorrect outer source address (OSA) if the first port is configured as a FabricPath edge port and the second port is configured as a FabricPath core port. To work around this issue, configure the first port as a FabricPath core port and the second port as a FabricPath edge port.

- Beginning with Cisco NX-OS Release 6.2(2), SSM is supported on virtual port channel+ (vPC+).

- When multicast routing is occurring on a FabricPath spine switch, the egress core ports towards the FabricPath leaf switches should not have a mix of F2e and F3 Series module ports. This may cause multicast traffic to be forwarded on both FTags, which can lead to duplicate multicast traffic received at the destination leaf switch, depending on the topology. This limitation only affects Layer-3 routed multicast traffic.

Configuring FabricPath Interfaces

Note   You must have an F Series module in the chassis and enabled FabricPath on all the devices before you can see the FabricPath commands on the devices.

Note   You must make these configurations on each switch that you want to participate in the FabricPath network.

Configuring FabricPath Interfaces

You configure the interfaces for the FabricPath network to be FabricPath interfaces.

Note   By default, all the interfaces on the N7K-F132XP-15 module are Layer 2 access interfaces.

Before you begin

Ensure that you are working on an F Series module.

Ensure that you have enabled the FabricPath feature on all devices.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface** [**ethernet** *slot/port* | **port-channel** *channel-no*]	Enters interface configuration mode.
Step 3	switch(config-if)# [**no**] **switchport mode fabricpath**	Specifies interfaces as FabricPath ports. **Note** The **no** keyword returns the interface to the default CE access interface. The FabricPath ports carry traffic only for those VLANs that are configured as FabricPath VLANs.
Step 4	(Optional) switch(config-if)# **system default switchport fabricpath**	Converts all CE interfaces on the F Series module to FabricPath interfaces simultaneously.
Step 5	switch(config-if)# **exit**	Exits interface configuration mode.
Step 6	switch(config)# **exit**	Exits global configuration mode.
Step 7	(Optional) switch# **show interface**	Displays information about all interfaces.
Step 8	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

This example shows how to configure specified interfaces as FabricPath interfaces:

```
switch# configure terminal
switch(config)# interface ethernet 2/11-15
switch(config-if)# switchport mode fabricpath
switch(config-if)#
```

Configuring the STP Priority with Rapid PVST+

All Layer 2 gateway devices must have the same bridge priority when they are in the same STP domain. Make sure that the STP priority configured for the Layer 2 gateway devices on a FabricPath network is the lowest value in the Layer 2 network. Additionally, the priorities must match.

We recommend that you configure the STP priority on all FabricPath Layer 2 gateway devices to 8192.

Before you begin

Ensure that you are working on an F Series module.

Ensure that you have enabled the FabricPath feature on all devices.

Procedure

	Command or Action	**Purpose**
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **spanning-tree vlan** [*vlan-id*] **priority** [*value*]	Configures all the Rapid PVST+ VLANs on all the FabricPath Layer 2 gateway interfaces to a lower STP priority. We recommend that you configure the priority to be 8192.
Step 3	switch(config)# **exit**	Exits global configuration mode.
Step 4	(Optional) switch# **show spanning-tree summary**	Displays information about STP.
Step 5	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

This example shows how to configure the Rapid PVST+ VLANs on the FabricPath Layer 2 gateway devices to have an STP priority of 8192:

```
switch# configure terminal
switch(config)# spanning-tree vlan 11-20 priority 8192
switch(config)#
```

See the *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Command Reference* for more information about this command.

Configuring the STP Priority with MST

All Layer 2 gateway devices must have the same bridge priority when they are in the same STP domain. Make sure that the STP priority configured for the Layer 2 gateway devices on a FabricPath network is the lowest value in the Layer 2 network. Additionally, the priorities must match.

You configure the STP priority for all Multiple Spanning-Tree (MST) instances on all FabricPath Layer 2 gateway devices to 8192.

Before you begin

Ensure that you are working on an F Series module.

Ensure that you have enabled the FabricPath feature on all devices.

Procedure

	Command or Action	**Purpose**
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **spanning-tree mst** [*instance-id*] **priority** [*value*]	Configures all the MST VLANs on all the FabricPath Layer 2 gateway interfaces to a

	Command or Action	Purpose
		lower STP priority. We recommend that you configure the priority to be 8192.
Step 3	switch(config)# **exit**	Exits global configuration mode.
Step 4	(Optional) switch# **show spanning-tree summary**	Displays information about STP.
Step 5	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

This example shows how to configure the MST instances on the FabricPath Layer 2 gateway devices to have an STP priority of 8192:

```
switch# configure terminal
switch(config)# spanning-tree mst 1-5 priority 8192
switch(config)#
```

See the *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Command Reference* for more information about this command.

Configuring the STP Domain ID for STP Domains Connected to the Layer 2 Gateway Switch

Because there can be many FabricPath Layer 2 gateway switches attached to a single FabricPath network, there are also many separate STP domains that are each connected to a Layer 2 gateway switch. You can configure a unique STP domain ID in the FabricPath network to propagate TCNs across all the STP domains that are connected to the FabricPath network to ensure that all the MAC addresses are flushed when the system receives a TCN.

Before you begin

Ensure that you are working on an F Series module.

Ensure that you have enabled the FabricPath feature on all devices.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **spanning-tree domain** *domain-id*	Assigns an STP domain ID to the different STP domains attached to FabricPath Layer 2 gateway switches that are connected to a single FabricPath network. The range is from 1 to 1023.

	Command or Action	Purpose
Step 3	switch(config)# **exit**	Exits global configuration mode.
Step 4	(Optional) switch# **show spanning-tree summary**	Displays information about STP.
Step 5	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

This example shows how to configure the STP domain ID attached to the FabricPath Layer 2 gateway device:

```
switch# configure terminal
switch(config)# spanning-tree domain 5
switch(config)# exit
```

Configuring a vPC+ Switch ID

Note   All the peer link and downstream links in the virtual private channel (vPC+) must be on the F Series module.

You configure the vPC+ switch ID by using the **fabricpath switch-id** command.

Note   You cannot configure a vPC+ domain and a vPC domain in the same virtual device context (VDC).

Note   No two vPC+ domains should have identical vPC+ domain IDs and matching emulated switch IDs. If a vPC+ has a domain ID and the configured emulated switch ID is identical then no other switch within the network is allowed to have the same set of IDs.

See the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide* for complete information about configuring vPCs.

Before you begin

Ensure that you are working on an F Series module.

Ensure that you have enabled the vPC feature.

Ensure that you have enabled the FabricPath feature.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	Creates a vPC+ domain on the device, and enters the vpc-domain configuration mode for configuration purposes.
Step 3	switch(config)# **fabricpath switch-id** *switch-id*	Assigns a static vPC+ ID to the vPC+ peer. The range is from 0 to 4094. This static ID is the virtual switch ID for FabricPath encapsulation. **Note** You must assign the same vPC+ switch ID to each of the two vPC+ peer devices before they can form an adjacency.

Example

This example shows how to configure a vPC+ switch ID on each vPC+ peer device:

```
switch# configure terminal
switch(config)# vpc domain 1
switch(config-vpc-domain)# fabricpath switch-id 1
```

vPC+ to vPC Configuration

You can switch from a vPC+ configuration to a standard vPC configuration.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	Enters the vpc-domain configuration mode for configuration purposes.
Step 3	switch(config-vpc-domain)# **no fabricpath switch-id** *switch-id*	Deconfigures the FabricPath switch ID.

Step 4	Perform one of the following:
	
	• For Cisco NX-OS Release 6.2(10) or a later release, enter **yes** at the following prompt:
	
	`Deconfiguring fabricpath switch id will flap vPCs. vPC+ to vPC transition needs reconfiguration of vPCs`
	`for this release, please refer to`	

	Command or Action	Purpose
| | ```
configuration guide for more
details. Continue (yes/no)? [no]
```<br><br>• For releases prior to Cisco NX-OS Release 6.2(10), enter **yes** at the following prompt:<br><br>```
Deconfiguring fabricpath switch id
will flap vPCs. Continue (yes/no)?
[no]
``` | |
| **Step 5** | For Cisco NX-OS Release 6.2(10) or a later release, delete and reconfigure all vPCs. | |

Configuring an Anycast HSRP Bundle

Beginning with Cisco Release 6.2(2), you can create an anycast Hot Standby Router Protocol (HSRP) bundle for a VLAN range that provides active-active forwarding on all nodes.

Note For more information about HSRP, see the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide*.

Configuring an HSRP Group

You can configure a HSRP group or a set of VLANs.

Before you begin

Ensure that you are working on an F Series module.

Ensure that you have enabled the FabricPath.

Ensure that you have enabled the HSRP feature.

Ensure that you have enabled the interface VLAN feature.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface vlan** *interface_number*	Configures the VLAN interface number and enters interface configuration mode.
Step 3	switch(config-if)# **hsrp version 2**	Specifies HSRP version 2. Anycast is supported only in HSRP version 2.
Step 4	switch(config-if)# [**no**] **hsrp** *group_number* {**ipv4** | **ipv6**}	Configures an HSRP group and enters HSRP configuration mode. The HSRP group can be either an IPv4 or an IPv6 group.

	Command or Action	**Purpose**
Step 5	switch(config-if-hsrp)# **ip** *ip_address*	Configures the virtual IP address of the HSRP group.
Step 6	switch(config-if-hsrp)# **exit**	Exits HSRP configuration mode.
Step 7	switch(config-if)# **exit**	Exits interface configuration mode.
Step 8	switch(config)# **exit**	Exits global configuration mode.
Step 9	(Optional) switch# **show hsrp**	Displays HSRP group information.
Step 10	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

This example shows how to configure an HSRP group:

```
switch# configure terminal
switch(config)# interface vlan 2
switch(config-if)# hsrp version 2
switch(config-if)# hsrp 1 ipv4
switch(config-if-hsrp)# ip 1.1.1.1

switch# show hsrp
```

Configuring an Anycast Bundle

You can create an anycast bundle that is an association between a set of VLANs and an anycast switchID.

Before you begin

Ensure that you are working on an F Series module.

Ensure that you have enabled the FabricPath.

Ensure that you have enabled the HSRP feature.

Ensure that you have enabled the interface VLAN feature.

Note In NX-OS versions prior to 6.2(10), if the VLAN range corresponding to the anycast HSRP bundle includes a partially configured or unconfigured SVI, the whole anycast bundle is brought down.

Procedure

	Command or Action	**Purpose**
Step 1	switch# **configure terminal**	Enters global configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# [**no**] **hsrp anycast** *bundle-id* {**ipv4**	**ipv6**	**both**}	Configures an anycast bundle. The arguments and keywords are as follows: *bundle-id*—Bundle ID. The range is from 1 to 4096. **ipv4**—Specifies an IPv4 bundle. All the IPv4 groups in the interface are associated with this bundle. **ipv6**—Specifies an IPv6 bundle. All the IPv6 groups in the interface are associated with this bundle. **both**—Specifies an IPv4 and IPv6 bundle. This is the default. All the IPv4 and IPv6 groups in the interface are associated with this bundle.
Step 3	switch(config-anycast-bundle)# [**no**] **force gateway-down**	Enforces the anycast bundle to remain in the down state even if one invalid VLAN is configured for the bundle.
Step 4	switch(config-anycast-bundle)# [**no**] **switch-id** *asid*	Configures the switch ID for the anycast bundle.
Step 5	switch(config-anycast-bundle)# **vlan** *range*	Configures the VLAN range for the anycast bundle. **Note** Beginning with Cisco NX-OS Release 6.2(10), you can add or delete a VLAN to or from an existing VLAN range for the anycast bundle without having to enter the complete VLAN range again.
Step 6	switch(config-anycast-bundle)# [**no**] **priority** *priority_value*	Configures the priority for the anycast bundle. This value is used to elect a root for all the groups in the range. The range is from 1 to 127 and the default value is 100.
Step 7	switch(config-anycast-bundle)# [**no**] **track** *object_id*	Configures the tracking value that is used to track the anycast bundle. The range is from 1 to 500 and the default value is 0, which indicates that nothing is tracked.
Step 8	switch(config-anycast-bundle)# [**no**] **timer** *hello_interval*	Configures the timers for the groups using this anycast bundle. The default value is 3.
Step 9	switch(config-anycast-bundle)# [**no**] **shutdown**	Configures the group to take the switch out of the anycast bundle. To bring the switch into

	Command or Action	Purpose
		the anycast bundle, enter the **no** form of the command.
Step 10	switch(config-anycast-bundle)# **exit**	Exits anycast configuration mode.
Step 11	switch(config)# **exit**	Exits global configuration mode.
Step 12	(Optional) switch# **show hsrp anycast bundle** [*bundle_id* **ipv4** | **ipv6** | **both**]	Displays anycast HSRP bundle information.
Step 13	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

This example shows how to configure an anycast bundle for a selection of VLANs:

```
switch# configure terminal
switch(config)# hsrp anycast 1 ipv4
switch(config-anycast-bundle)# force gateway-down
switch(config-anycast-bundle)# switch-id 1300
switch(config-anycast-bundle)# vlan 1,20-30
switch(config-anycast-bundle)# priority 90
switch(config-anycast-bundle)# track 2
switch(config-anycast-bundle)# timer 15 25
switch(config-anycast-bundle)# shutdown
```

This example shows how to add VLAN 5 to an existing VLAN range of 1,20-30 in different Cisco NX-OS releases:

```
switch(config-anycast-bundle)# vlan 1,5,20-30 (Cisco NX-OS Release
6.2(8) and earlier releases)
switch(config-anycast-bundle)# vlan 5 (Cisco NX-OS Release 6.2(10) and later
releases)
```

Configuring Anycast Bundle Limits

You can create limits for the anycast bundles.

Before you begin

Ensure that you are working on an F Series module.

Ensure that you have enabled the FabricPath.

Ensure that you have enabled the HSRP feature.

Ensure that you have enabled the interface VLAN feature.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vdc switch**	Enters VDC configuration mode.
Step 3	(Optional) switch(config-vdc)# [**no**] **limit-resource anycast_switchid minimum** *min* **maximum** *max*	Configures the limits for the anycast bundles that are allowed in the system. To return the limits to the default values, enter the no form of the command. *min*—The minimum number of anycast bundles allowed is set as 0 and cannot be changed. *max*—The maximum number of anycast bundles allowed. The default value is 16. For Supervisor 1 and Supervisor 2, the maximum value is limited to 64. For Supervisor 2e and Supervisor 3, the maximum value is limited to 128.
Step 4	switch(config-vdc)# **exit**	Exits VDC configuration mode.
Step 5	switch(config)# **exit**	Exits global configuration mode.
Step 6	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

This example shows how to configure the limits for an anycast bundle:

```
switch# configure terminal
switch(config)# vdc switch
switch(config-vdc)# limit-resource anycast_switchid minimum 0 maximum 8
```

Verifying FabricPath Interface Configuration

To display FabricPath interfaces information, perform one of the following tasks:

Command	Purpose
show feature-set	Displays whether FabricPath is enabled on the device or not.
show interface brief	Displays information on all interfaces.
show interface switchport	Displays information, including access and trunk interface, for all the Layer 2 interfaces.

Command	Purpose
show interface type {*slot/port*	*channel-number*} [**trunk**]	Displays interface configuration information.
show interface capabilities	Displays information on the capabilities of the interfaces.
show interface status	Displays information on the status of the interfaces.
show spanning-tree summary	Displays STP information.
show fabricpath is-is database	Displays STP TCN information.
show vpc brief	Displays brief information on the vPC+ domains.
show vpc consistency-parameters	Displays the status of those parameters that must be consistent across all vPC+ domain interfaces.
show vpc peer-keepalive	Displays information on the peer-keepalive messages.
show vpc role	Displays the peer status, the role of the local device, the vPC+ domain's system MAC address and system priority, and the MAC address and priority for the local vPC+ domain's device.
show vpc statistics	Displays statistics on the vPC+ domains.
show running-config vpc	Displays running configuration information for vPCs and vPC+ domains.
show hsrp anycast bundle [*bundle_id* **ipv4**	**ipv6**] [**brief**]	Displays information for anycast bundles.
show hsrp anycast bundle brief	Displays information for anycast bundles.
show hsrp anycast interface vlan *interface*	Displays information about the interface in the anycast bundle.
show hsrp anycast summary	Displays the summary of anycast information.
show hsrp anycast internal info bundle [*bundle_id* **ipv4**	**ipv6**]	Displays all the data structures related to anycast.
show hsrp anycast remote-db [*bundle_id* **ipv4**	**ipv6**]	Displays the remote database for all the bundles.

For information about the above commands, see the *Cisco Nexus 7000 Series NX-OS Interfaces Command Reference* and the *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Command Reference*.

Monitoring FabricPath Interface Statistics

Use the following commands to display FabricPath statistics:

- **clear counters** [**ethernet** *slot/port* | **port-channel** *channel-no*]

- **show interface counters** [**module** *module*]

- **show interface counters detailed** [**all**]

- **show interface counters errors** [**module** *module*]

Configuration Example for FabricPath Interface

To configure FabricPath interfaces, perform the following tasks on each device:

- Enable FabricPath on each device.

- Configure the interfaces that you want to designate as FabricPath interfaces.

- Set the STP priority device to 8192 on all FabricPath Layer 2 gateway devices.

- (Optional) Set the STP domain ID for each of the separate STP domains that are connected to the FabricPath network.

- (Optional) Configure a vPC+ switch ID.

To configure FabricPath interfaces, follow these steps:

Step 1 (Optional): Enable FabricPath on each device.

```
switch# configure terminal
switch(config)# feature fabricpath
switch(config-lldp)# exit
switch(config)#
```

Step 2: After you enable FabricPath on the device, configure the specified interface as FabricPath interfaces.

```
switch(config)# interface ethernet 1/2
switch(config-if)# switchport mode fabricpath
switch(config-if)# exit
switch(config)#
```

Step 3: Configure the STP priority for all Rapid PVST+ VLANs as 8192.

```
switch# configure terminal
switch(config)# spanning-tree vlan 11-20 priority 8192
switch(config)#
```

Step 4: Configure the STP priority for all MST instances as 8192.

```
switch# configure terminal
switch(config)# spanning-tree mst 1-5 priority 8192
switch(config)#
```

Step 5 (Optional): Configure the STP domain ID on each FabricPath Layer 2 gateway switch attached to the FabricPath network.

```
switch# configure terminal
switch(config)# spanning-tree domain 5
switch(config)
```

Step 6 (Optional): Configure the vPC+ switch ID.

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)# fabricpath switch-id 100
switch(config-vpc-domain)# exit
switch(config)
```

Note   See the Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, for information about configuring vPC.

If you are configuring the vPC+ with no existing vPC+, follow these steps:

1. In the vPC domain configuration mode, enter the **fabricpath switch-id** *switch-id* command.

2. On each of the vPC+ peer link interfaces in interface configuration mode, enter the **switchport mode fabricpath** command.

3. On each vPC+ peer link port channel, enter the **vpc peer-link** command.

If you are changing an existing vPC configuration to a vPC+ on an F Series module, follow these steps:

1. On each vPC peer link port channel, enter the **shutdown** command.

2. In the vPC domain configuration mode, enter the **fabricpath switch-id** *switch-id* command.

3. On each of the vPC+ peer link interfaces in interface configuration mode, enter the **switchport mode fabricpath** command.

4. On each vPC+ peer link port channel, enter the **no shutdown** command.

Step 7: Save the configuration.

```
switch(config)# save running-config startup-config
switch(config)#
```

When you are configuring vPC+, and you see the following situations, you must enter the **shutdown** command and then the **no shutdown** command on all the peer-link interfaces:

  • There is no switchport mode FabricPath configuration on the peer-link interfaces, but the FabricPath switch ID is configured in the vPC domain.

  • The **switchport mode fabricpath** configuration is on the peer-link interfaces, but there is no FabricPath switch ID in the vPC domain.

Feature History for Configuring FabricPath Interface

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

Table 2: Feature History for FabricPath Interface

Feature Name	Release	Feature Information
vPC+ to vPC configuration	6.2(10)	Changed warning prompt message and added requirement for all vPCs to be deleted and reconfigured.
Anycast HSRP	6.2(10)	Added the ability to add or delete a VLAN to or from an existing VLAN range (for an HSRP Anycast bundle) without having to enter the complete VLAN range again.
Anycast HSRP and overload bit	6.2(8)	The anycast switch ID is no longer advertised when the FabricPath Layer 2 IS-IS overload-bit is set. Please see more details about the Fabricpath Layer IS-IS overload bit in the section "Configuring Advanced FabricPath Features."
Configuring an anycast HSRP bundle	6.2(2)	Added the ability to create an anycast HSRP bundle.
Configuring more than 244 vPC+ port channels	6.1(3)	Added support for configuring more than 244 vPC+ port channels with the **no port-channel limit** command.
Configuring vPC+ with FEX ports	6.1(3)	Added support for configuring vPC+ with FEX ports with the **fabricpath multicast load-balance** command.
FabricPath Interfaces	5.1(1)	This feature was introduced.