



# Configuring FabricPath Interfaces

This chapter describes how to configure the FabricPath interfaces on the Cisco NX-OS 5500 Series switches.

This chapter includes the following sections:

- [Information About FabricPath Interfaces, page 3-1](#)
- [Guidelines and Limitations for FabricPath Interfaces, page 3-5](#)
- [Configuring FabricPath Interfaces, page 3-6](#)
- [Verifying the FabricPath Interface Configuration, page 3-15](#)
- [Monitoring FabricPath Interface Statistics, page 3-16](#)
- [Configuration Example for FabricPath Interface, page 3-16](#)
- [Feature History for Configuring FabricPath Interfaces, page 3-18](#)



**Note**

For information about prerequisites, guidelines and limitations, and licensing requirements for FabricPath, see [Chapter 1, “Overview.”](#)

## Information About FabricPath Interfaces

This section includes the following sections:

- [FabricPath Interfaces, page 3-1](#)
- [STP and the FabricPath Network, page 3-2](#)
- [vPC+, page 3-3](#)
- [Anycast HSRP, page 3-5](#)

## FabricPath Interfaces

After you enable FabricPath on the switches 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 switch as FabricPath interfaces simultaneously.

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

- Interfaces on the switch that are configured as FabricPath interfaces can carry traffic only for FabricPath VLANs.
- Interfaces on the switch that are not configured as FabricPath interfaces carry traffic for the following:
  - FabricPath VLANs
  - Classical Ethernet (CE) VLANs

See [Chapter 4, “Configuring FabricPath Forwarding,”](#) for information about FabricPath and CE VLANs.

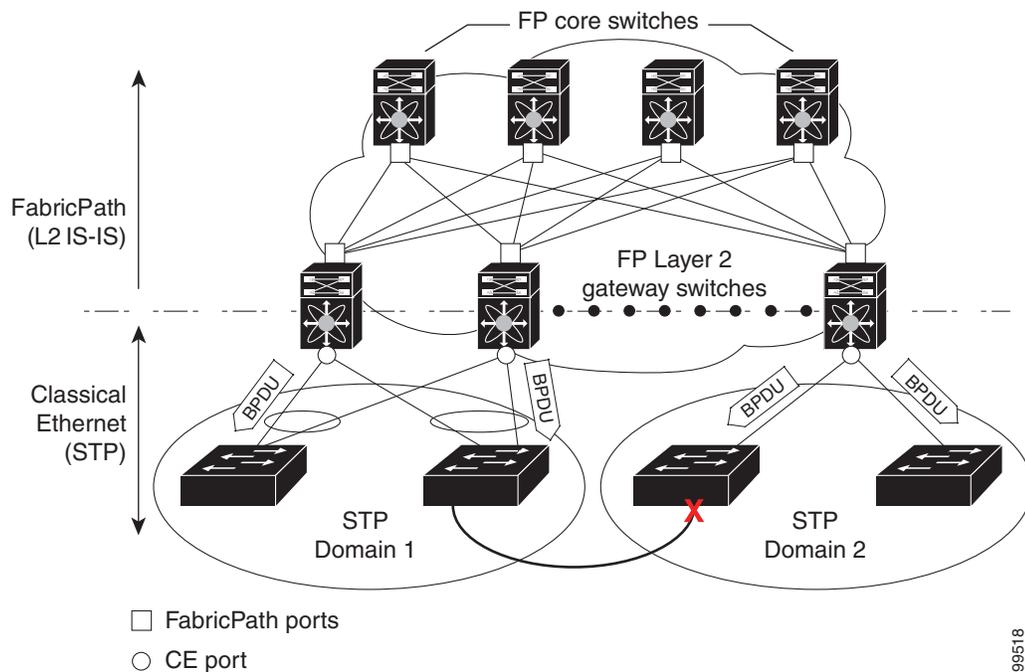
The FabricPath interfaces connect only to other FabricPath interfaces within the FabricPath network. These FabricPath ports operate on the information in the FabricPath headers and Layer 2 Intermediate System-to-Intermediate System (IS-IS) only, and they do not run STP. These ports are aware only of FabricPath 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 FabricPath VLANs.

## STP and the FabricPath Network

The Layer 2 gateway switches, which are on the edge between the CE and the FabricPath network, must be configured as the root for all Spanning Tree Protocol (STP) domains that are connected to a FabricPath network.

The STP domains do not cross into the FabricPath network (see [Figure 3-1](#)).

**Figure 3-1** STP Boundary Termination at FabricPath Network Border



199518

You must configure the FabricPath Layer 2 gateway switch to have the lowest STP priority of all the switches in the STP domain to which it is attached. You must also configure all the FabricPath Layer 2 gateway switches that are connected to one FabricPath network to have the *same* priority. The software assigns the bridge ID for the Layer 2 gateway switches from a pool of reserved MAC addresses.

To ensure a loop-free topology for the CE and FabricPath hybrid network, the FabricPath network automatically displays as a single bridge to all connected CE switches.

**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 switches form a single STP domain. Those CE switches that are not interconnected form separate STP domains (see [Figure 3-1](#)).

All CE interfaces should be designated ports, which occurs automatically, or they are pruned from the active STP topology. If the software prunes any port, the software returns a syslog message. The software 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 many FabricPath Layer 2 gateway switches might be attached to a single FabricPath network, there might also be many separate STP domains (see [Figure 3-1](#)). The switches 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 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 switches.

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 switches in the specified STP domain.

The switches in the separate STP domains need to receive the TCN information and then flush all remote MAC addresses 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 5500 Series NX-OS Interfaces Configuration Guide, Release 7.x*). You cannot configure a vPC+ domain and a vPC domain on the same Cisco Nexus 5500 Series switch.

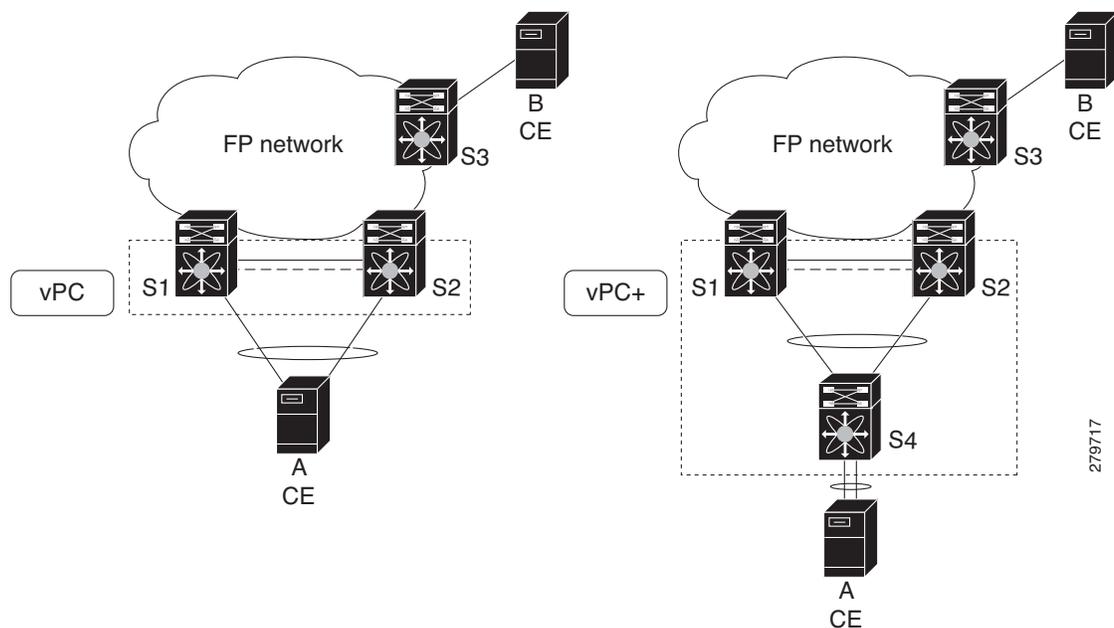
A vPC+ domain enables Cisco Nexus 5500 Series switches that have FabricPath enabled to form a single vPC+, which is a unique virtual switch to the rest of the FabricPath network. You configure the same domain on each switch to enable the peers to identify each other and to form the vPC+. Each vPC+ has its own virtual switch ID.

**Note**

We do not recommend that you enable the vPC peer switch feature when you are using a vPC+.

A vPC+ must still provide active-active Layer 2 paths for dual-homed CE switches or clouds, even though the FabricPath network allows only 1-to-1 mapping between the MAC address and the switch ID. vPC+ provides the solution by creating a unique virtual switch to the FabricPath network (see [Figure 3-2](#)).

**Figure 3-2 Comparison of vPC and 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 sent to the MAC address of the virtual switch as the transit source, and traffic from host B to host A is now load balanced.

The vPC+ downstream links are FabricPath edge interfaces that connect to the CE hosts.

The First Hop Routing Protocols (FHRPs), which include the Hot Standby Routing Protocol (HSRP) and the Virtual Router Redundancy Protocol (VRRP), interoperate with a vPC+. You should dual-attach all Layer 3 switches to both vPC+ peer switches.

**Note**

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

Both the primary and secondary vPC+ switches forward traffic, but only the primary FHRP switch responds to ARP requests.

To simplify initial configuration verification and vPC+/HSRP troubleshooting, you can configure the primary vPC+ peer switch 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 for when a group state 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 switch fails over to the secondary vPC+ peer switch, the FHRP traffic continues to flow seamlessly.

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

**Caution**

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 switches.

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.

See the *Cisco Nexus 5500 Series NX-OS Unicast Routing Configuration Guide, Release 7.x* for more information on FHRPs and routing.

## Anycast HSRP

Cisco NX-OS provides a way to facilitate further scalability at the spine layer providing 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.

Cisco NX-OS supports a maximum of 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.

## Guidelines and Limitations for FabricPath Interfaces

FabricPath has the following configuration guidelines and limitations:

- An anycast HSRP bundle provides support for a maximum of four nodes at the spine layer.
- An anycast HSRP bundle is supported only in HSRP version 2.

- A maximum of 10 anycast bundles are supported.

## Configuring FabricPath Interfaces

You must configure FabricPath interfaces on all switches that are enabled for FabricPath.



### Tip

If you cannot see any of these commands, verify that you have installed and enabled the FabricPath feature set.

This section includes the following topics:

- [Configuring FabricPath Interfaces, page 3-6](#)
- [Configuring the STP Priority with Rapid PVST+, page 3-8](#)
- [Configuring the STP Priority with MST, page 3-9](#)
- [Configuring the STP Domain ID for STP Domains Connected to the Layer 2 Gateway Switch \(Optional\), page 3-11](#)
- [Configuring a vPC+ Switch ID, page 3-12](#)
- [Configuring an Anycast HSRP Bundle, page 3-13](#)

## Configuring FabricPath Interfaces

You must configure the interfaces for the FabricPath network as FabricPath interfaces.

### BEFORE YOU BEGIN

Ensure that you have installed and enabled the FabricPath feature on all switches.

### SUMMARY STEPS

1. **configure terminal**
2. **interface** [*ethernet slot/port* | **port-channel** *channel-no*]
3. **fabricpath topology** *topology-number*
4. [**no**] **switchport mode fabricpath**
5. (Optional) **system default switchport fabricpath**
6. **exit**
7. (Optional) **show interface**
8. (Optional) **copy running-config startup-config**

## DETAILED STEPS

	Command	Purpose
Step 1	<p><b>configure terminal</b></p> <p><b>Example:</b>  switch# configure terminal  switch(config)#</p>	Enters global configuration mode.
Step 2	<p><b>interface</b> [ethernet <i>slot/port</i>    port-channel <i>channel-no</i>]</p> <p><b>Example:</b>  switch(config)# interface ethernet 2/11-15  switch(config-if)#</p>	<p>Enters interface configuration mode and specifies the interfaces that you want to configure as FabricPath.</p> <p>The slot can be from 1 to 3. The following list defines the slots available:</p> <ul style="list-style-type: none"> <li>• Slot 1 includes all the fixed ports. A Fabric Extender only has one slot.</li> <li>• Slot 2 includes the ports on the upper expansion module (if populated).</li> <li>• Slot 3 includes the ports on the lower expansion module (if populated).</li> </ul> <p>The port number within a particular slot can be from 1 to 128.</p> <p>The port channel number assigned to the EtherChannel logical interface can be from 1 to 4096.</p>
Step 3	<p><b>fabricpath topology</b> <i>topology-number</i></p> <p><b>Example:</b>  switch(config-if)# fabricpath topology 1  switch(config-if)#</p>	<p>Specifies the FabricPath topology number.</p> <p><b>Note</b> Cisco Nexus devices only supports two topologies; the default or base topology (topology 0) and another optional topology (for example, topology 1).</p>
Step 4	<p><b>[no] switchport mode fabricpath</b></p> <p><b>Example:</b>  switch(config-if)# switchport mode fabricpath  switch(config-if)#</p>	<p>Specifies interfaces as FabricPath ports.</p> <p><b>Note</b> The <b>no</b> 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.</p>
Step 5	<p><b>system default switchport fabricpath</b></p> <p><b>Example:</b>  switch(config-if)# system default switchport fabricpath  switch(config-if)#</p>	(Optional) Converts all CE interfaces on the switch to FabricPath interfaces simultaneously.
Step 6	<p><b>exit</b></p> <p><b>Example:</b>  switch(config-if)# exit  switch(config)#</p>	Exits interface configuration mode.

	Command	Purpose
Step 7	<b>show interface</b>  <b>Example:</b> switch(config)# show interface	(Optional) Displays information on all interfaces.
Step 8	<b>copy running-config startup-config</b>  <b>Example:</b> switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

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 switches 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 switches 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 switches to 8192.

### BEFORE YOU BEGIN

Ensure that you have installed and enabled the FabricPath feature on all switches.

### SUMMARY STEPS

1. **configure terminal**
2. **spanning-tree pseudo-information**
3. **spanning-tree vlan [vlan-id] priority [value]**
4. **exit**
5. **(Optional) show spanning-tree summary**
6. **(Optional) copy running-config startup-config**

## DETAILED STEPS

	Command	Purpose
Step 1	<b>configure terminal</b>  <b>Example:</b> switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	<b>spanning-tree pseudo-information</b>  Example: switch(config)# spanning-tree pseudo-information switch(config-pseudo)#	Enters spanning tree pseudo-information configuration mode.
Step 3	<b>spanning-tree vlan [vlan-id] priority [value]</b>  <b>Example:</b> switch(config-pseudo)# spanning-tree vlan 11-20 priority 8192 switch(config-pseudo)#	You must configure 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 4	<b>exit</b>  <b>Example:</b> switch(config-pseudo)# exit switch(config)#	Exits the interface configuration mode.
Step 5	<b>show spanning-tree summary</b>  <b>Example:</b> switch# show spanning-tree summary	(Optional) Displays STP information.
Step 6	<b>copy running-config startup-config</b>  <b>Example:</b> switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

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

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

See the *Cisco Nexus 5500 Series NX-OS Layer 2 Switching Configuration Guide, Release 7.x* for more information on this command.

## Configuring the STP Priority with MST

All Layer 2 gateway switches 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 switches 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 for all Multiple Spanning-Tree (MST) instances on all FabricPath Layer 2 gateway switches to 8192.

**BEFORE YOU BEGIN**

Ensure that you have installed and enabled the FabricPath feature on all switches.

**SUMMARY STEPS**

1. **configure terminal**
2. **spanning-tree pseudo-information**
3. **spanning-tree mst [instance-id] priority [value]**
4. **exit**
5. **(Optional) show spanning-tree summary**
6. **(Optional) copy running-config startup-config**

**DETAILED STEPS**

	<b>Command</b>	<b>Purpose</b>
<b>Step 1</b>	<b>configure terminal</b>  <b>Example:</b> switch# configure terminal switch(config)#	Enters global configuration mode.
<b>Step 2</b>	<b>spanning-tree pseudo-information</b>  <b>Example:</b> switch(config)# spanning-tree pseudo-information switch(config-pseudo)#	Enters spanning tree pseudo-information configuration mode.
<b>Step 3</b>	<b>spanning-tree mst [instance-id] priority [value]</b>  <b>Example:</b> switch(config-pseudo)# spanning-tree mst 1-5 priority 8192 switch(config-pseudo)#	Configures all the MST VLANs on all the FabricPath Layer 2 gateway interfaces to a lower STP priority. We recommend that you configure the priority to be 8192.
<b>Step 4</b>	<b>exit</b>  <b>Example:</b> switch(config-pseudo)# exit switch(config)#	Exits interface configuration mode.
<b>Step 5</b>	<b>show spanning-tree summary</b>  <b>Example:</b> switch# show spanning-tree summary	(Optional) Displays information on STP.
<b>Step 6</b>	<b>copy running-config startup-config</b>  <b>Example:</b> switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

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

```
switch# configure terminal
switch(config)# spanning-tree pseudo-information
```

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

See the *Cisco Nexus 5500 Series NX-OS Layer 2 Switching Configuration Guide, Release 7.x* for more information on this command.

## Configuring the STP Domain ID for STP Domains Connected to the Layer 2 Gateway Switch (Optional)

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. This configuration ensures that all MAC addresses are flushed when the software receives a TCN.

### BEFORE YOU BEGIN

Ensure that you have installed and enabled the FabricPath feature on all switches.

### SUMMARY STEPS

1. **configure terminal**
2. **spanning-tree domain** *domain-id*
3. **exit**
4. (Optional) **show spanning-tree summary**
5. (Optional) **copy running-config startup-config**

### DETAILED STEPS

	Command	Purpose
Step 1	<b>configure terminal</b>  <b>Example:</b> switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	<b>spanning-tree domain</b> <i>domain-id</i>  <b>Example:</b> switch(config)# spanning-tree domain 5 switch(config)#	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.
Step 3	<b>exit</b>  <b>Example:</b> switch(config)# exit switch#	Exits configuration mode.

	Command	Purpose
Step 4	<b>show spanning-tree summary</b>  <b>Example:</b> switch# show spanning-tree summary	(Optional) Displays STP information.
Step 5	<b>copy running-config startup-config</b>  <b>Example:</b> switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

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

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

## Configuring a vPC+ Switch ID

You can configure a vPC+ switch ID.



### Note

The vPC+ virtual SID cannot be the same as the FabricPath SID for the switch.

## BEFORE YOU BEGIN

Ensure that you have enabled the FabricPath feature.

## SUMMARY STEPS

1. **configure terminal**
2. **vPC domain** *domain-id*
3. **fabricpath switch-id** *switch-id*

## DETAILED STEPS

	Command	Purpose
Step 1	<b>configure terminal</b>  <b>Example:</b> switch# <b>configure terminal</b>	Enters global configuration mode.

	Command	Purpose
Step 2	<b>vpc domain</b> <i>domain-id</i>  <b>Example:</b> switch(config)# <b>vpc domain 1</b>	Creates a vPC+ domain on the switch, and enters the vPC-domain configuration mode for configuration purposes.
Step 3	<b>fabricpath switch-id</b> <i>switch-id</i>  <b>Example:</b> switch(config-vpc-domain)# <b>fabricpath switch-id 1</b>	Assigns a static vPC+ ID to the vPC+ peer. The range is from 1 to 4094. This static ID is the virtual switch ID for FabricPath encapsulation.  <b>Note</b> You must assign the same vPC+ switch ID to each of the two vPC+ peer switches before they can form an adjacency.

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

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

See the *Cisco Nexus 5500 Series NX-OS Interfaces Configuration Guide, Release 7.x* for complete information on configuring vPCs.

## Configuring an Anycast HSRP Bundle

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

### BEFORE YOU BEGIN

- Ensure that you have enabled the FabricPath feature.
- Ensure that you have enabled the HSRP feature.
- Ensure that you have enabled the interface VLAN feature.

## DETAILED STEPS

	Command	Purpose
Step 1	<code>switch# configure terminal</code>	Enters global configuration mode.
Step 2	<code>switch(config)# [no] hsrp anycast bundle-id {ipv4   ipv6   both}</code>	Configures an anycast bundle. The arguments and keywords are as follows: <ul style="list-style-type: none"> <li><i>bundle-id</i>—Bundle ID. The range is from 1 to 4096.</li> <li><b>ipv4</b>—Specifies an IPv4 bundle. All the IPv4 groups in the interface are associated with this bundle.</li> <li><b>ipv6</b>—Specifies an IPv6 bundle. All the IPv6 groups in the interface are associated with this bundle.</li> <li><b>both</b>—Specifies an IPv4 and IPv6 bundle. This is the default. All the IPv4 and IPv6 groups in the interface are associated with this bundle.</li> </ul>
Step 3	<code>switch(config-anycast-bundle)# [no] force-gateway down</code>	Enforces the anycast bundle to remain in the down state even if one invalid VLAN is configured for the bundle. The bundle will come up only when all configured VLANs are valid. This mode favors loss of traffic.  If you use the <b>no</b> form of the command, the bundle will come up even if one valid VLAN configuration exists. This mode favors loss of anycast gateway.
Step 4	<code>switch(config-anycast)# [no] switch-id asid</code>	Configures the switch ID for the anycast bundle. <ul style="list-style-type: none"> <li><i>asid</i>—Specifies the anycast switch to be used for the bundle.</li> </ul>
Step 5	<code>switch(config-anycast)# [no] vlan range</code>	Configures the VLAN range for the anycast bundle. <ul style="list-style-type: none"> <li><i>range</i>—Specifies the VLAN range.</li> </ul>
Step 6	<code>switch(config-anycast)# [no] priority priority_value</code>	Configures the priority for the anycast bundle. The <i>priority_value</i> 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	<code>switch(config-anycast)# [no] track object_id</code>	Configures the tracking value that is used to track the anycast bundle. The range of <i>object_id</i> is from 1 to 500. The default value is 0 which indicates that nothing is tracked.
Step 8	<code>switch(config-anycast)# [no] timer hello_interval</code>	Configures the timers for the groups using this anycast bundle. The default value of <i>hello_interval</i> is 3.
Step 9	<code>switch(config-anycast)# [no] shutdown</code>	Configures the group to take the switch out of the anycast bundle. To bring the switch into the anycast bundle, enter the <b>no</b> form of the command.

	Command	Purpose
Step 10	switch(config-anycast)# <b>show hsrp anycast</b> [ <i>bundle_id</i> <b>ipv4</b>   <b>ipv6</b>   <b>both</b> ]	(Optional) Displays anycast HSRP bundle information.
Step 11	switch(config-anycast)# <b>copy running-config startup-config</b>	(Optional) Copies the running configuration to the startup configuration.

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
switch(config-anycast-bundle)# no shutdown
```

## Verifying the FabricPath Interface Configuration

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

Command	Purpose
<b>show feature-set</b>	Displays whether FabricPath is enabled on the switch or not.
<b>show interface brief</b>	Displays information on all interfaces.
<b>show interface switchport</b>	Displays information, including access and trunk interface, for all the Layer 2 interfaces.
<b>show interface</b> <i>type</i> { <i>slot/port</i>   <i>channel-number</i> } [ <b>trunk</b> ]	Displays interface configuration information.
<b>show interface capabilities</b>	Displays information on the capabilities of the interfaces.
<b>show interface status</b>	Displays information on the status of the interfaces.
<b>show spanning-tree summary</b>	Displays STP information.
<b>show fabricpath is-is database</b>	Displays STP TCN information.
<b>show vPC brief</b>	Displays brief information on the vPC+ domains.
<b>show vPC consistency-parameters</b> {[ <b>vpc</b> <i>vpc-id</i>   <b>global</b>   <b>vlans</b>   <b>interface</b> {[ <b>ethernet</b> <i>mod/slot</i>   <b>port-channel</b> <i>channel-number</i> ]}]}	Displays the status of those parameters that must be consistent across all vPC+ domain interfaces.
<b>show vPC peer-keepalive</b>	Displays information on the peer-keepalive messages.
<b>show vPC role</b>	Displays the peer status, the role of the local switch, the vPC+ domain's MAC address and system priority, and the MAC address and priority for the local vPC+ domain's switch.

Command	Purpose
<code>show vPC statistics</code> {[peer-keepalive   peer-link   vpc vpc-id]}	Displays statistics on the vPC+ domains.
<code>show running-config vPC</code> [all]	Displays running configuration information for vPCs and vPC+ domains.
<code>show hsrp anycast</code> [bundle_id ipv4   ipv6][brief]	Displays information for anycast bundles.
<code>show hsrp anycast brief</code>	Displays information for anycast bundles.
<code>show hsrp anycast interface</code> vlan interface	Displays information about the interface in the anycast bundle.
<code>show hsrp anycast summary</code>	Displays the summary of anycast information.
<code>show hsrp anycast internal info</code> bundle_id ipv4   ipv6	Displays all the data structures related to anycast.
<code>show hsrp anycast remote-db</code> [bundle_id ipv4   ipv6]	Displays the remote database for all the bundles.

## Monitoring FabricPath Interface Statistics

Use the following commands to display or clear FabricPath interface statistics:

- `show interface counters` [module module]
- `show interface counters detailed` [all]
- `show interface counters errors` [module module]
- `clear counters` [interface {[all | ethernet mod/slot | port-channel channel-number | loopback virtual-interface-num | mgmt mgmt-interface-num | vlan vlan-id]} ]

## Configuration Example for FabricPath Interface

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

- Enable FabricPath on each switch.
- Configure the interfaces that you want to designate as FabricPath interfaces.
- Set the STP priority switch to 8192 on all FabricPath Layer 2 gateway switches.
- (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 switch.

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

**Step 2** After you enable FabricPath on the switch, 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)
```

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 the vPC+ peer link port channel, enter the **vPC peer-link** command.

If you are changing an existing vPC configuration to a vPC+, 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 the vPC+ peer link port channel, enter the **no shutdown** command.

**Step 7** Save the configuration.

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

---

If you are configuring a 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 Interfaces

Table 3-1 lists the release history for this feature.

**Table 3-1** Feature History for FabricPath

Feature Name	Releases	Feature Information
Anycast HSRP	7.1(0)N1(1)	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	7.1(0)N1(1)	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."
FabricPath Interfaces	5.1(3)N1(1)	This feature was introduced.