



Configuring FabricPath Forwarding

This chapter describes how to configure FabricPath forwarding on Cisco Nexus 5500 Series switches.

This chapter includes the following sections:

- [Information About FabricPath Forwarding, page 4-1](#)
- [Configuring FabricPath Forwarding, page 4-5](#)
- [Verifying the FabricPath Configuration, page 4-9](#)
- [Configuration Example for FabricPath Forwarding, page 4-10](#)
- [Default Settings for Fabric Path Forwarding, page 4-11](#)
- [Feature History for Configuring FabricPath, page 4-11](#)



Note

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

Information About FabricPath Forwarding

This section includes the following sections:

- [FabricPath Forwarding Overview, page 4-1](#)
- [FabricPath VLANs, page 4-2](#)
- [Forwarding Known Unicast Packets Using ECMP, page 4-3](#)
- [Forwarding Trees for Broadcast, Unknown Unicast, and Multicast Packets, page 4-4](#)

FabricPath Forwarding Overview

FabricPath provides a multipath Layer 2 domain that does not require STP for a loop-free environment. Using the Intermediate System-to-Intermediate System (IS-IS) protocol, the switch provides multiple paths for Layer 2 packets.

Each FabricPath interface can learn multiple parallel paths to the other nodes in the FabricPath network. Because you do not need to use STP, all the paths are available for forwarding traffic. The switch assigns the optimal path per flow.

The flow for known unicast packets is determined by the hierarchical FabricPath outer destination address (ODA) and the outer source address (OSA) value (see [Chapter 2, “Configuring FabricPath Switching,”](#) for more information on FabricPath hierarchical encapsulation). The software uses IS-IS equal cost multipath (ECMP) to choose the forwarding path for these flows.

For multidestination traffic (unknown unicast, broadcast, and multicast), the software creates two paths or trees. The broadcast and unknown unicast traffic flows through one of these trees. The software distributes the multicast traffic between the two trees based on a hash. The software load balances multicast traffic in the FabricPath network (see the [“Forwarding Trees for Broadcast, Unknown Unicast, and Multicast Packets”](#) section on page 4-4 for more information).

FabricPath Layer 2 IS-IS defines the trees. The highest switch ID is chosen for the root and the first tree is built with that switch as the root. (If needed, you can configure the switch ID to specify which switch has the highest switch ID and is, therefore, the root for the first tree.) The root switch for the first tree elects the root of the second tree, based on the switch ID, and the second tree flows from that root switch. All of this information is advertised to the FabricPath network using Layer 2 IS-IS, so all the switches in the network have the same information.

The software assigns the path at ingress and encodes that path in the FTag portion of the FabricPath header. The software assigns one FTag per tree. Once decided and tagged, the packet uses the same tree throughout the entire FabricPath network. All the nodes in the FabricPath network forward traffic based on this same information because all nodes have the same information using Layer 2 IS-IS.

The FabricPath frame has a Reverse Path Forwarding (RPF) function for multidestination packets, which verifies that the packet is arriving on an interface that leads to the source switch. RPF drops the packet if it is received from an interface that is not part of the tree.

The FabricPath Layer 2 IS-IS protocol floods the link-state information across the FabricPath network. Each switch sends hello packets on each FabricPath link and discovers its neighbors. When a neighbor is discovered, the software creates an IS-IS adjacency. Each switch also sends advertisements and updates to the link-state database through all the existing adjacencies.

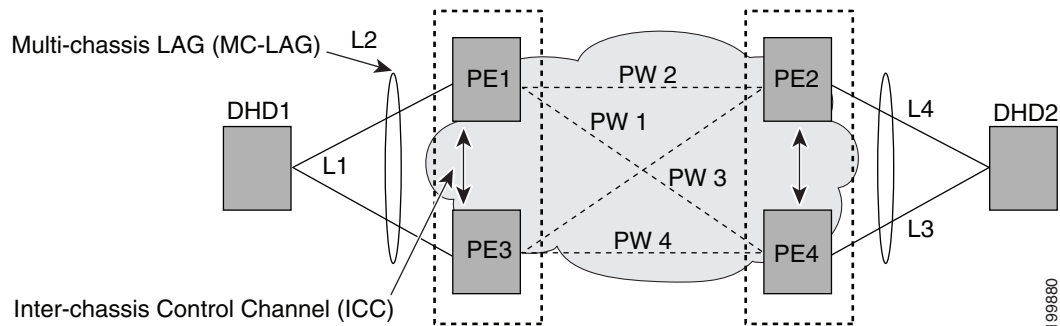
FabricPath VLANs

To interact with the Classical Ethernet (CE) network, you set VLANs to either CE or FabricPath mode. The CE VLANs carry traffic from the CE hosts to the FabricPath interfaces, and the FabricPath VLANs carry traffic throughout the FabricPath topology. Only the active FabricPath VLANs configured on a switch are advertised as part of the topology in the Layer 2 Intermediate System-to-Intermediate System (IS-IS) messages.

The software automatically assigns all FabricPath interfaces and FabricPath VLANs to the FabricPath topology. No added configuration is required. All the FabricPath VLANs and interfaces belong to that topology.

Figure 4-1 shows an example of the FabricPath topology with Classical Ethernet switches and FabricPath/CE VLANs.

Figure 4-1 Example FabricPath Topology and Classical Ethernet Hosts



The default VLAN mode on the switch is the CE VLAN mode. The FabricPath interfaces carry traffic only on the FabricPath VLANs; the CE VLANs do not come up on these interfaces.

You must exit the VLAN configuration mode for the VLAN mode change to take effect.



Note

After you configure the VLANs and interfaces, no further configuration is required. The software automatically creates and assigns the paths, and provides load balancing.

Forwarding Known Unicast Packets Using ECMP

The switch forwards unicast traffic per flow using the ODA field in the FabricPath header for known unicast traffic. The FabricPath-enabled switch assigns the switch ID and the ODA for all encapsulated traffic at the ingress switch. (See [Chapter 2, “Configuring FabricPath Switching,”](#) for more information on FabricPath encapsulation.)

Once the switch assigns the ODA, the FabricPath switch uses the FabricPath Layer 2 IS-IS ECMP to forward known unicast traffic. FabricPath Layer 2 IS-IS has up to 16 active Layer 2 paths for all known unicast packets. The Layer 2 IS-IS messages used by FabricPath are separate and distinct from the Layer 3 IS-IS messages used by the routing protocols and the Overlay Transport Virtualization (OTV).

The switches within the FabricPath network exchange topology information using IS-IS adjacencies and forward the traffic along those paths for known unicast traffic flows. Each node in the FabricPath network looks at the FabricPath header for each traffic flow and makes an ECMP forwarding choice based on the available next hops.

Forwarding Trees for Broadcast, Unknown Unicast, and Multicast Packets


FabricPath has a loop-free broadcast functionality that carries broadcast, unknown unicast, and multicast packets, or multidestination traffic. For each broadcast, unknown unicast, and multicast traffic flow, the software chooses the forwarding path from among two software-created paths or trees. The software creates two trees to forward multidestination traffic.

For the FabricPath network, the software creates a broadcast tree that carries broadcast traffic, unknown unicast traffic through the FabricPath network. The software also creates another tree; all the multicast traffic flows are load balanced across these two trees for each flow. Each tree is identified in the FabricPath network by a unique value or FTag. Within the FabricPath network, the software elects a root node that becomes root for the broadcast tree. That node also identifies another bridge to become root for the second multidestination tree, which load balances the multicast traffic.

The FTag is assigned by the ingress switch, along with the ODA and OSA, as part of the FabricPath encapsulation. The FTag determines which loop-free tree the multidestination traffic flow follows through the FabricPath network. The software assigns the trees per flow.

Figure 4-2 shows these trees.

Figure 4-2 Trees for Forwarding Multidestination FabricPath Flows for a Given Flow



Each node in the FabricPath network shares the same view of the forwarding trees for a given FTag.

Forwarding Multicast Packets

Using FabricPath, you can configure Layer 2 multicast multipathing. FabricPath uses a hash-based system to assign each of the multicast flows to one of the two designated trees to ensure that the multicast traffic is load balanced.

The software uses FabricPath Layer 2 IS-IS and Classical Ethernet IGMP snooping to learn the multicast group information at the boundaries of the FabricPath and Classical Ethernet network. The software carries that information through the FabricPath network using a new Layer 2 IS-IS LSP called Group Membership LSP (GM-LSP). GM-LSPs carry multicast group/source membership information. This information is carried across the FabricPath network. All FabricPath switches maintain multicast routing information and forward multicast data packets only to switches that have interested receivers. Each node in the FabricPath topology shares the same view and contains the same information.

The multicast traffic uses VLAN information to allocate traffic to one or the other of the two multicast trees. This software constrains multicast based on the group IP address.

IGMP snooping and FabricPath IS-IS, using GM-LSP, work together to build per-VLAN multicast group-based trees across the FabricPath network. IGMP snooping on edge switches learns of receivers and routers and builds an edge-port multicast state. FabricPath Layer 2 IS-IS propagates this group

information through the FabricPath network using GM-LSPs and builds the state in the FabricPath network. These GM-LSPs originate from the switches at the edge of the FabricPath network that have multicast groups.

For Layer 2 multicast traffic, you do not to run PIM at all when using FabricPath.

For Layer 3 multicast packets, the software sets the ODA to a special multicast group that identifies all IP routers for that group and forwards the traffic along the tree for that group.

Configuring FabricPath Forwarding



Tip

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

You must make these configurations on every switch in the FabricPath network.

Only those VLANs that are configured as FabricPath VLANs can belong to the FabricPath topology. By default, all FabricPath VLANs are assigned to the base topology (topology 0). However, you can configure the FabricPath VLAN can belong to a different topology.

FabricPath has two topologies, you need to set the VLAN mode for those VLANs that you want to traverse the FabricPath network to FabricPath VLAN.

After you configure the VLAN modes and interfaces, the software automatically creates the required paths. You do not need to perform any additional configuration for FabricPath.

This section includes the following topics:

- [Setting the VLAN Mode to FabricPath or CE, page 4-5](#)
- [Mapping FabricPath VLANs to a Topology, page 4-6](#)

Setting the VLAN Mode to FabricPath or CE

The default VLAN mode is CE.

You designate those VLANs that you want to carry FabricPath traffic on the network by configuring them as FabricPath VLANs.

BEFORE YOU BEGIN

Ensure that you have created the VLANs.

SUMMARY STEPS

1. **configure terminal**
2. **vlan *vlan-id***
3. **mode [ce | fabricpath]**
4. **exit**
5. **(Optional) show fabricpath topology vlans [active]**
6. **(Optional) copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	vlan <i>vlan-id</i> Example: switch(config)# vlan 1-10 switch(config-vlan)#	Enters the VLAN configuration mode and identifies those VLANs that you want to carry FabricPath traffic. The VLAN range is from 1 to 4094. Note You cannot modify VLAN 1 or any of the internally allocated VLANs.
Step 3	mode [ce fabricpath] Example: switch(config-vlan)# mode fabricpath switch(config-vlan)#	Configures the VLANs as FabricPath VLANs. The default VLAN mode is CE. Note A VLAN must be either a CE or an FabricPath VLAN on the FabricPath switch.
Step 4	exit Example: switch(config-vlan)# exit switch(config)#	Exits the VLAN configuration mode. Note As with all VLANs, you must exit the VLAN configuration mode for the VLAN mode (CE or FabricPath) to take effect.
Step 5	show fabricpath topology vlans [active] Example: switch# show fabricpath topology vlans	(Optional) Displays information about all active VLANs in the FabricPath topology.
Step 6	copy running-config startup-config Example: switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to specify VLANs as FabricPath VLANs:

```
switch# configure terminal
switch(config)# vlan 1-10
switch(config-vlan)# mode fabricpath
switch(config-vlan)# exit
switch(config)# exit
```

Mapping FabricPath VLANs to a Topology

The default VLAN mode is CE.

You designate those VLANs that you want to carry FabricPath traffic on the network by configuring them as FabricPath VLANs. By default, all FabricPath VLANs belong to default topology 0. Also FabricPath VLANs can be a member of only one topology.

BEFORE YOU BEGIN

Ensure that you have created the VLANs.

SUMMARY STEPS

1. **configure terminal**
2. **fabricpath topology *topology-number***
3. **member vlan *vlan-id***
4. **(Optional) show fabricpath topology vlans [active]**
5. **(Optional) copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	fabricpath topology <i>topology-number</i> Example: switch(config)# fabricpath topology 1 switch(config-fp-topology)#	Specifies the FabricPath topology number. Note Cisco Nexus devices support two FabricPath topologies, the default or base topology (topology 0), and another optional topology (for example, topology 1).
Step 3	member vlan <i>vlan-id</i> Example: switch(config-fp-topology)# member vlan 501-600 switch(config-vlan)#	Specifies the VLANs to assign to FabricPath topology.
Step 4	show fabricpath topology vlans [active] Example: switch# show fabricpath topology vlans	(Optional) Displays information about all active VLANs in the FabricPath topology.
Step 5	copy running-config startup-config Example: switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to map a VLAN to a FabricPath topology:

```
switch# configure terminal
switch(config)# fabricpath topology 1
switch(config-fp-topology)# member vlan 501-600
switch(config-fp-topology)# exit
```

Configuring FabricPath Unicast Load Balancing (Optional)

The FabricPath network automatically balances unicast traffic when multiple paths are available. You can configure specific load balancing for the unicast traffic. The default is to use all options.

BEFORE YOU BEGIN

Ensure that you have enabled the FabricPath feature set.

SUMMARY STEPS

1. **configure terminal**
2. **fabricpath load-balance unicast {include-vlan | layer2 | layer3 | layer4 | mixed}**
3. **exit**
4. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	fabricpath load-balance unicast {include-vlan layer2 layer3 layer4 mixed} Example: switch(config)# fabricpath load-balance unicast source switch(config-vlan)#	Configures load balancing for FabricPath unicast traffic along the following parameters: <ul style="list-style-type: none"> • include-vlan—Specifies that the load-balancing parameters use VLAN. • layer2—Specifies that the load-balancing parameters only include Layer 2 parameters. • layer3—Specifies that the load-balancing parameters only include Layer 3 parameters. • layer4—Specifies that the load-balancing parameters only include Layer 4 parameters. • mixed—Specifies that the load-balancing parameters include a mix of Layer 3 and Layer 4 parameters. This is the default configuration. To return to the default unicast load-balancing scheme, enter the no form of this command.

	Command	Purpose
Step 3	exit Example: switch(config)# exit switch(config)#	Exits configuration mode.
Step 4	copy running-config startup-config Example: switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to configure FabricPath unicast load balancing with source parameters:

```
switch# configure terminal
switch(config)# fabricpath load-balance unicast include-vlan
switch(config)#
```

Verifying the FabricPath Configuration

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



Tip

For more information about these commands and the options available, see the *Cisco Nexus 5500 Series NX-OS FabricPath Command Reference*.

Command	Purpose
show feature-set	Displays whether FabricPath is enabled or not.
show {I2 fabricpath} route [switchid switch-id] [detail] [hex]	Displays unicast routes.
show {I2 fabricpath} mroute {[vdc_omf] [vlan vlanid] {[omf] [flood] [source {srcaddr ipv6srcaddr}] [group {groupaddr ipv6groupaddr}]} [resolved] [ftag ftag-id] [hex]}	Displays multicast routes.
show fabricpath topology [detail]	Displays information on the FabricPath topology.
show fabricpath topology interface	Displays information on all FabricPath topology interfaces.
show fabricpath topology vlan [active]	Displays information on all FabricPath topology VLANs.
show fabricpath topology ftag [active] [multicast] [unicast]	Displays information on all FabricPath topology FTags.
show running-config fabricpath	Displays the running configuration for FabricPath.

Command	Purpose
show fabricpath load-balance unicast forwarding-path ftag <i>f-tag-ID</i> switchid <i>switch-ID</i> src-mac <i>MAC-addr</i>	<p>Displays FabricPath unicast load-balancing information.</p> <p>This command contains the following options:</p> <ul style="list-style-type: none"> <i>f-tag-ID</i>—FTag for which you want to view load balancing. The FTag value is from 0 to 4,294,967,295. <i>switch-ID</i>—The ID of the switch for which you want to view load balancing. The switch ID is from 0 to 4,294,967,295.
show fabricpath load-balance multicast ftag-selected vlan <i>vlan-id</i> macg <i>macg</i>	<p>Displays FabricPath multicast load-balancing information. The output from this command is valid for multicast traffic ingressing on non-vPC ports. For multicast traffic ingressing on vPC ports, traffic selects the active FTag only for that switch.</p> <p>This command contains the following options:</p> <ul style="list-style-type: none"> <i>vlan-id</i>—A VLAN ID in a range between 1 and 4094. <i>macg</i>—The multicast group MAC address.
show vlan	Displays information on all FabricPath and CE VLANs.

Configuration Example for FabricPath Forwarding

To configure the basic FabricPath network, you must accomplish the following tasks on each switch after you have configured the FabricPath interfaces:

- Enable the FabricPath feature set on each switch.
- Configure the FabricPath interfaces. (See [Chapter 3, “Configuring FabricPath Interfaces,”](#) for information on configuring FabricPath interfaces.)
- Configure the FabricPath VLANs. The default is CE VLANs.
- Enter the **show running-config fabricpath** command to make sure that your FabricPath configuration is correct.

To configure the FabricPath topology, follow these steps:

Step 1 Enable the FabricPath feature set.

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

Step 2 Set the VLAN modes for those VLANs that you want in the FabricPath topology to FabricPath.

```
switch# configure terminal
switch(config)# vlan 11-20
switch(config-vlan)# mode fabricpath
switch(config-vlan)# exit
```

```
switch(config)
```

Step 3 Display the configuration to ensure that you have the correct configuration.

```
switch(config)# show running-config fabricpath
switch(config)#
```

Step 4 Save the configuration.

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

Default Settings for Fabric Path Forwarding

Table 4-1 lists the default settings for the FabricPath forwarding parameters.

Table 4-1 Default FabricPath Parameters

Parameters	Default
FabricPath topology	0
VLAN mode	FabricPath

Feature History for Configuring FabricPath

Table 4-2 lists the release history for this feature.

Table 4-2 Feature History for FabricPath

Feature Name	Releases	Feature Information
FabricPath Forwarding	5.1(3)N1(1)	This feature was introduced

