



CHAPTER 4

FabricPath Forwarding

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

This chapter includes the following sections:

- [Information About FabricPath Forwarding, page 4-1](#)
- [Licensing Requirements for FabricPath Forwarding, page 4-6](#)
- [Prerequisites for FabricPath Forwarding, page 4-6](#)
- [Guidelines and Limitations for FabricPath Forwarding, page 4-6](#)
- [Configuring FabricPath Forwarding, page 4-7](#)
- [Verifying the FabricPath Configuration, page 4-12](#)
- [FabricPath Forwarding Example Configuration, page 4-13](#)
- [Default Settings, page 4-14](#)
- [Feature History for Configuring FabricPath Using the CLI, page 4-14](#)

Information About FabricPath Forwarding



Note

You must have an F Series module in your chassis to run FabricPath.

This section includes the following sections:

- [FabricPath Forwarding Overview, page 4-2](#)
- [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](#)
- [High Availability, page 4-5](#)
- [Virtual Device Context, page 4-5](#)

Send document comments to nexus7k-docfeedback@cisco.com

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 device 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 device 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 system uses IS-IS Equal Cost Multipathing (ECMP) to choose the forwarding path for these flows, using FabricPath Layer 2 IS-IS.

For multidestination traffic (unknown unicast, broadcast, and multicast), the FabricPath system creates two paths or trees. The broadcast and unknown unicast traffic flows through one of these trees. The system distributes the multicast traffic between the two trees based on a hash. The system 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 system ID is chosen for the root and the tree flows from that. The system ID is configurable per device. The second tree is the same but with a different root priority. After the system chooses the root switch, the tree is built with that as the root for the first tree. Then, the root switch for the first tree elects the root of the second tree, again based on system 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 devices in the network have the same information.

The system assigns the path at ingress and encodes that path in the FTag portion of the FabricPath header. The system 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) mechanism 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 device sends hello packets on each FabricPath link and discovers its neighbors. When a neighbor is discovered, the system creates an IS-IS adjacency. Each device 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 (FP) mode. The CE VLANs carry traffic from the CE hosts to the FabricPath interfaces, and the FP VLANs carry traffic throughout the FabricPath topology. Only the active FP 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 system automatically assigns all FabricPath interfaces and FP VLANs to the topology. So, there is no added configuration required. (See [Chapter 3, “FabricPath Interfaces”](#), for information on FabricPath interfaces.) All the FP VLANs and FabricPath interfaces belong to that same topology. All ports on the same device in the same topology must be in the same virtual device context (VDC).

[Figure 4-1](#) shows a sample FabricPath topology with Classical Ethernet switches and FP/CE VLANs.

[Send document comments to nexus7k-docfeedback@cisco.com](mailto:nexus7k-docfeedback@cisco.com)

Figure 4-1 Example FabricPath Topology and Classical Ethernet Hosts



The default VLAN mode on the device is the CE VLAN mode. The FabricPath interfaces carry traffic only on the FP VLANs; the CE VLANs do not come up on these interfaces. The CE interfaces on the F Series modules carry traffic for both CE VLANs (traffic from the hosts) and FP VLANs.

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



Note

Once you configure VLANs and interfaces, no further configuration is required. The system automatically creates and assigns the paths, as well as provides load balancing.

Forwarding Known Unicast Packets Using ECMP

The system forwards unicast traffic per flow using the ODA field in the FabricPath header for known unicast traffic. The FabricPath-enabled system 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 system assigns the ODA, the FabricPath device uses the FabricPath Layer 2 IS-IS ECMP to forward known unicast traffic. FabricPath, using Layer 2 IS-IS, has up to 16 active Layer 2 paths. This feature provides up to 16-way ECMP at Layer 2 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 devices 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.

[Send document comments to nexus7k-docfeedback@cisco.com](mailto:nexus7k-docfeedback@cisco.com)

Forwarding Trees for Broadcast, Unknown Unicast, and Multicast Packets


FabricPath introduces a new 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 system chooses the forwarding path from among multiple system-created paths or trees. The system creates two trees to forward the multidestination traffic for each topology.

For the FabricPath network, the system creates a broadcast tree that carries broadcast traffic, unknown unicast traffic, and multicast traffic through the FabricPath network. The system also creates a second 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 system 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 loopfree tree that the multidestination traffic flow follows through the FabricPath network. The system 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 and an F Series module, 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 system uses FabricPath Layer 2 IS-IS and Classical Ethernet IGMP snooping to learn the multicast group information at the boundaries of the FabricPath/Classical Ethernet network. The system 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 each FabricPath topology shares the same view and has all the same information.

The multicast traffic uses the per-VLAN source, multicast group, and flow information to allocate traffic to one or the other of the two trees. This system 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 interfaces learns of receivers and routers and builds an edge-port multicast state. FabricPath Layer 2 IS-IS propagates this attached

Send document comments to nexus7k-docfeedback@cisco.com

group information through the FabricPath network using GM LSPs, building a state in the FabricPath network. Devices at the edge of the FabricPath network that have multicast groups originate the GM-LSP.

Beginning with Cisco Release 5.2(1), you can add a configuration to assist the device to quickly work with multiple multicast groups. See [“Configuring FabricPath Increased Multicast Scalability \(Optional\)” section on page 4-11](#) for more information.

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

For Layer 3 multicast packets, the system 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.

High Availability

The FabricPath topologies retain their configuration through ISSU.

See the *Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide, Release 5.x*, for more information on high availability.

Virtual Device Context

You must install the FabricPath feature set before you enable FabricPath on the switch. See the *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 following port pairs must be in the same VDC:

- Ports 1 and 2
- Ports 3 and 4
- Ports 5 and 6
- Ports 7 and 8
- Ports 9 and 10
- Ports 11 and 12
- Ports 13 and 14
- Ports 15 and 16
- Ports 17 and 18
- Ports 19 and 20
- Ports 21 and 22
- Ports 23 and 24
- 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, Release 5.x*, for more information on VDCs.

Send document comments to nexus7k-docfeedback@cisco.com

Licensing Requirements for FabricPath Forwarding

The following table shows the licensing requirements for this feature:

Product	License Requirement
Cisco NX-OS	FabricPath requires an Enhanced Layer 2 license. For a complete explanation of the Cisco NX-OS licensing scheme, see the <i>Cisco Nexus 7000 Series Licensing Guide</i> .

Prerequisites for FabricPath Forwarding

FabricPath forwarding has the following prerequisites:

- You should have a working knowledge of Classical Ethernet Layer 2 functioning.
- You must install the FabricPath feature set in the default and non-default 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.
- You are logged onto the device.
- Ensure that you have installed the Enhanced Layer 2 license.
- 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 Forwarding

FabricPath has the following configuration guidelines and limitations:

- FabricPath interfaces carry only FabricPath-encapsulated traffic.
- You must 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 the *Configuring Feature Set for FabricPath* for complete information on installing and enabling the FabricPath feature set.
- The FabricPath feature set operation may cause the standby supervisor to reload if it is in an unstable state, such as following a service failure or powering up.
- STP does not run inside a FabricPath network.
- The F Series modules do not support multiple SPAN destination ports or virtual SPAN. If a port on the 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 FP. 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 or static routes.

Send document comments to nexus7k-docfeedback@cisco.com

- On the F Series modules, user-configured static MAC addresses are programmed on all forwarding engines (FEs) that have ports in that VLAN.
- In order to have the VLAN mode take effect, you must exit the VLAN configuration mode after configuring the mode.

Configuring FabricPath Forwarding

**Note**

You must have FabricPath enabled on the F Series module and on all devices before you can see any of these commands.

Only those VLANs that are configured as FP VLANs can belong to the FabricPath topology. By default, all FP VLANs and interfaces are assigned to the FabricPath topology, FabricPath topo 0.

When you are using the default topology, you need only to set the VLAN mode for those VLANs that you want to traverse the FabricPath network to FP VLAN.

Because the system automatically creates the multiple paths once you specify the VLAN modes and interfaces, you are only required to configure these aspects of FabricPath.

(See [Chapter 3, “FabricPath Interfaces”](#), for information on FabricPath interfaces.)

**Note**

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

This section describes how to configure the FabricPath and includes the following topics:

- [Setting the VLAN Mode to FP or CE, page 4-7](#)
- [Configuring FabricPath Unicast Load Balancing \(Optional\), page 4-9](#)
- [Configuring FabricPath Multicast Load Balancing \(Optional\), page 4-10](#)
- [Configuring FabricPath Increased Multicast Scalability \(Optional\), page 4-11](#)

Setting the VLAN Mode to FP or CE

The default VLAN mode is CE on the F Series modules.

**Note**

You must have already created the VLANs before you can set the VLAN mode using FP.

You designate those VLANs that you want to carry FabricPath traffic on the network by configuring them as FP VLANs. By default, all FP VLANs and FabricPath interfaces are added to the default FabricPath topology, topo 0.

BEFORE YOU BEGIN

Ensure that you are working on an F Series module.

Ensure that you have installed the Enhanced Layer 2 license.

Ensure that you have enabled the FabricPath feature set.

Send document comments to nexus7k-docfeedback@cisco.com

Ensure that you have created the VLANs.

SUMMARY STEPS

1. **config 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	config terminal Example: switch# config 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.
Step 3	mode [ce fabricpath] Example: switch(config-vlan)# mode fabricpath switch(config-vlan)#	Configures the VLANs as FP VLANs. The default VLAN mode is CE. Note A VLAN must be either a CE or an FP VLAN on the FabricPath device.
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 FP) to take effect.
Step 5	show fabricpath topology vlans [active] Example: switch# show fabricpath topology vlans	(Optional) Displays information about all 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 FP VLANs:

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

You can add interfaces to a non-default topology:

[Send document comments to nexus7k-docfeedback@cisco.com](mailto:nexus7k-docfeedback@cisco.com)

Configuring FabricPath Unicast Load Balancing (Optional)

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

BEFORE YOU BEGIN

Ensure that you are working on an F Series module.

Ensure that you have installed the Enhanced Layer 2 license.

Ensure that you have enabled the FabricPath feature set.

SUMMARY STEPS

1. **config terminal**
2. **fabricpath load-balance unicast** {{{source | source-destination | xor | destination | symmetric}}] [{"layer3 | layer4 | mixed}]} [rotate-amount rot_amt] [include-vlan]}
3. **exit**
4. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config terminal Example: switch# config terminal switch(config)#	Enters global configuration mode.
Step 2	fabricpath load-balance unicast {{{source source-destination xor destination symmetric}}] [{"layer 3 layer4 mixed}]} [rotate-amount rot_amt] [include-vlan]} Example: switch(config)# fabricpath load-balance unicast source switch(config-vlan)#	Configures load balancing for FabricPath unicast traffic along specified parameters. Note To return to the default unicast load-balancing scheme, enter the no form of this command.
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:

```
switch# config terminal
switch(config)# fabricpath load-balance unicast source
switch(config)#
```

[Send document comments to nexus7k-docfeedback@cisco.com](mailto:nexus7k-docfeedback@cisco.com)

Configuring FabricPath Multicast Load Balancing (Optional)

Although the network automatically load balances the traffic, you can configure specific load balancing for the multicast traffic.

BEFORE YOU BEGIN

Ensure that you are working on an F Series module.

Ensure that you have installed the Enhanced Layer 2 license.

Ensure that you have enabled the FabricPath feature set.

SUMMARY STEPS

1. **config terminal**
2. **fabricpath load-balance multicast** {{{source | source-destination | xor | destination | symmetric}}} [rotate-amount *rot_amt*] [include-vlan]}
3. **exit**
4. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config terminal Example: switch# config terminal switch(config)#	Enters global configuration mode.
Step 2	fabricpath load-balance multicast {{{source source-destination xor destination symmetric}}} [rotate-amount <i>rot_amt</i>] [include-vlan]} Example: switch(config)# fabricpath load-balance multicast source switch(config-vlan)#	Configures load balancing for FabricPath multicast traffic along specified parameters. Note To return to the default multicast load-balancing scheme, enter the no form of this command.
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 multicast load balancing:

```
switch# config terminal
```

Send document comments to nexus7k-docfeedback@cisco.com

```
switch(config)# fabricpath load-balance multicast symmetric
switch(config)#
```

Configuring FabricPath Increased Multicast Scalability (Optional)

Beginning with Cisco Release 5/2(1), you can increase the FabricPath multicast scalability.

BEFORE YOU BEGIN

Ensure that you are working on an F Series module.

Ensure that you have installed the Enhanced Layer 2 license.

Ensure that you have enabled the FabricPath feature set.

SUMMARY STEPS

1. **config terminal**
2. **fabricpath multicast aggregate-routes** [**exclude** *ftag-id*]
3. (Optional) **show l2 multicast ftag** *ftag-id*
4. **exit**
5. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config terminal Example: switch# config terminal switch(config)#	Enters global configuration mode.
Step 2	fabricpath multicast aggregate-routes [exclude <i>ftag</i>] Example: switch(config)# fabricpath multicast aggregate-routes exclude 1 switch(config-vlan)#	Increases FabricPath multicast scalability. The default is to not aggregate FTag routes. To find the multicast FTag used for a given traffic that you want to exclude, enter the show fabricpath load-balance multicast ftag-selected flow-type l3 dst-ip x.x.x.x src-ip x.x.x.x vlan vlan-id module mod-num command. Note The no version of this command does not include the exclude <i>ftag</i> argument.
Step 3	show l2 multicast ftag <i>ftag</i> Example: switch(config)# show l2 multicast ftag 1 switch(config-vlan)#	(Optional) Displays the configuration that you just applied to the FTag for route programming.

Send document comments to nexus7k-docfeedback@cisco.com

	Command	Purpose
Step 4	exit Example: switch(config)# exit switch(config)#	Exits configuration mode.
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 configure increased FabricPath multicast scalability:

```
switch# config terminal
switch(config)# fabricpath multicast aggregate-routes
switch(config)#
```

Verifying the FabricPath Configuration

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

Command	Purpose
show feature-set	Displays whether FabricPath is enabled or not.
show {12 fabricpath} route [switchid switch-id] [detail] [hex]	Displays unicast routes.
show {12 fabricpath} mroute {[vdc_omf] vlan vlan-id {{omf flood [source {srcaddr v6srcaddr}] [group {groupaddr v6grpaddr}]}} [resolved] [ftag ftag-id] [hex]	Displays multicast routes.
show fabricpath topology [detail]	Displays information on all FabricPath topologies.
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.
show fabricpath load-balance unicast forwarding-path ftag ftag-id switchid switch-id flow-type {12 {{dst-mac dst-mac source-mac src-mac} ether-type ether-type}} {13 {dst-ip dst-ip src-ip src-ip dst-ipv6 dst-ipv6 srcipv6 src-ipv6}} {14 {14-src-port 14-src-port 14-dst-port 14-dst-port dst-ip dst-ip src-ip src-ip dst-ipv6 dst-ipv6 srcipv6 src-ipv6}} {vlan vlan-id} {module mod-no}	Displays FabricPath unicast load-balancing information.

Send document comments to nexus7k-docfeedback@cisco.com

Command	Purpose
show fabricpath load-balance multicast ftag-selected flow-type {I2 {{dst-mac dst-mac source-mac src-mac} ether-type ether-type}} {I3 {dst-ip dst-ip src-ip src-ip dst-ipv6 dst-ipv6 srcipv6 src-ipv6}} {I4 {I4-src-port I4-src-port I4-dst-port I4-dst-port dst-ip dst-ip src-ip src-ip dst-ipv6 dst-ipv6 srcipv6 src-ipv6}}} {vlan vlan-id} {module mod-no}	Displays FabricPath multicast load-balancing information.
show vlan	Displays information on all FP and CE VLANs.

The following is sample output from the **show fabricpath unicast load-balance** command:

```
switch# show fabricpath load-balance unicast forwarding-path ftag 1 switchid 2231
flow-type 13 src-ip 1.1.1.1 dst-ip 1.1.1.2 module 4
128b Hash Key generated : 00001010102010101010000000000000
This flow selects interface Po100
```

The following is sample output from the **show fabricpath multicast load-balance** command:

```
switch(config)# show fabricpath load-balance multicast ftag-selected flow-type 13 src-ip
1.1.1.1 dst-ip 1.1.1.2 vlan 2 module 4
128b Hash Key generated : 00 00 10 10 10 20 00 00 00 00 02 00 00 00 00 00
0x3
FTAG SELECTED IS : 1
```

FabricPath Forwarding Example Configuration

To configure the basic FabricPath network with a default topology, you must accomplish the following tasks on each device after you have configured the FabricPath interfaces:

- Enable the FabricPath feature set on each device.
- Configure the FabricPath interfaces. (See [Chapter 3, “FabricPath Interfaces”](#), for information on configuring FabricPath interfaces.)
- Configure the FP 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 default FabricPath topology, follow these steps:

Step 1 Enable the FabricPath feature set.

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



Note See the *Configuring Feature-Set for FabricPath* for complete information on installing and enabling the FabricPath feature set.

Send document comments to nexus7k-docfeedback@cisco.com

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

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

Default Settings

Table 4-1 lists the default settings for FabricPath parameters

Table 4-1 *Default FabricPath Parameters*

Parameters	Default
FabricPath Topology	0
VLAN mode	CE

Feature History for Configuring FabricPath Using the CLI

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

Table 4-2 *Feature History for FabricPath*

Feature Name	Releases	Feature Information
FabricPath	5.1(1)	These features were introduced.
Additional FabricPath topologies	5.2(1)	These features were introduced.