# Configuring FabricPath Switching

**Note** You must have an F Series module installed in your Cisco Nexus 7000 Series chassis in order to run FabricPath and conversational learning.

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


- Finding Feature Information, on page 1
- Information About FabricPath Switching, on page 1
- Prerequisites for FabricPath, on page 13
- Guidelines and Limitations for FabricPath Switching, on page 13
- Default Setting for FabricPath Switching, on page 14
- Configuring FabricPath Switching, on page 15
- Verifying FabricPath Switching, on page 27
- Monitoring and Clearing FabricPath Switching Statistics, on page 27
- Configuration Example for FabricPath Switching, on page 27
- Feature History for Configuring FabricPath Switching, on page 28


## 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 Switching

FabricPath switching allows multipath networking at the Layer 2 level. The FabricPath network still delivers packets on a best-effort basis (which is similar to the Classical Ethernet [CE] network), but the FabricPath network can use multiple paths for Layer 2 traffic. In a FabricPath network, you do not need to run the Spanning Tree Protocol (STP) with its blocking ports. Instead, you can use FabricPath across data centers, some of which have only Layer 2 connectivity, with no need for Layer 3 connectivity and IP configurations.

The FabricPath encapsulation facilitates MAC mobility and server virtualization, which means that you can physically move the Layer 2 node but retain the same MAC address and VLAN association for the virtual machine. FabricPath also allows LAN extensions across data centers at Layer 2, which is useful in disaster recovery operations, as well as clustering applications such as databases. Finally, FabricPath is very useful in high-performance, low-latency computing.

With FabricPath, you use the Layer 2 intermediate System-to-Intermediate System (IS-IS) protocol for a single control plane that functions for unicast, broadcast, and multicast packets. There is no need to run the Spanning Tree Protocol (STP); it is a purely Layer 2 domain. This FabricPath Layer 2 IS-IS is a separate process than Layer 3 IS-IS.

Beginning in the Cisco NX-OS Release 5.1 and when you use the F Series module, Cisco supports the conversation-based MAC learning schema. Conversational learning can be applied to both FabricPath (FP) and CE VLANs. Using FabricPath and conversational MAC address learning, the device has to learn far fewer MAC addresses, which results in smaller, more manageable MAC tables.

# FabricPath Encapsulation

## FabricPath Headers

When a frame enters the FabricPath network, the system encapsulates the Layer 2 frame with a new FabricPath header. The switch IDs that the system assigns to each FabricPath device as it enters the FabricPath network is used as the outer MAC destination address (ODA) and outer MAC source address (OSA) in the FabricPath header. The figure below shows the FabricPath header encapsulating the classical Ethernet (CE) frame.

***Figure 1: FabricPath Frame Encapsulation***

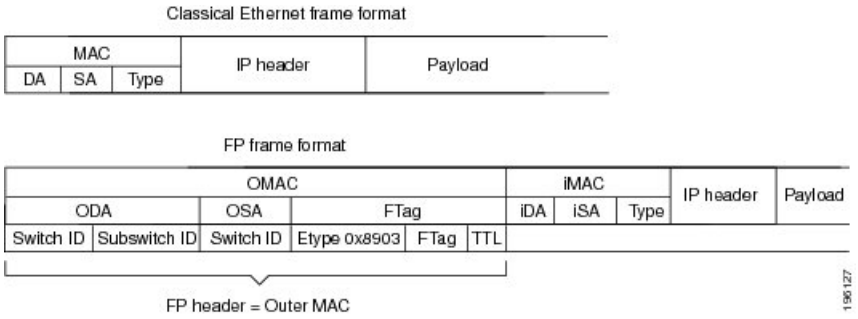


The system applies the encapsulation on the ingressing edge port of the FabricPath network and decapsulates the frame on the egressing edge port of the FabricPath network; all the ports within the FabricPath network are FabricPath ports that use only the hierarchical MAC address (see Chapter 3, “Configuring FabricPath Interfaces,” for more information on configuring FabricPath interfaces). This feature greatly reduces the size of the MAC tables in the core of the FabricPath network.

The system automatically assigns each device in the FabricPath network with a unique switch ID. Optionally, you can configure the switch ID for the FabricPath device.

The outer source address (OSA) is the FabricPath switch ID of the device where the frame ingresses the FabricPath network, and the outer destination address (ODA) is the FabricPath switch ID of the device where the frame egresses the FabricPath network. When the frame egresses the FabricPath network, the FabricPath device strips the FabricPath header, and the original CE frame continues on the CE network. The FabricPath network uses only the OSA and ODA, with the Layer 2 IS-IS protocol transmitting the topology information. Both the FabricPath ODA and OSA are in a standard MAC format (xxxx.xxxx.xxxx).

The FabricPath hierarchical MAC address carries the reserved EtherType 0x8903.

When the frame is originally encapsulated, the system sets the Time to Live (TTL) to 32. Optionally, you can configure the TTL value for multicast and unicast traffic. On each hop through the FabricPath network, each switch decrements the TTL by 1. If the TTL reaches 0, that frame is discarded. This feature prevents the continuation of any loops that may form in the network.

## Forwarding Tags (FTags)

The Forwarding Tag (FTag) in the FabricPath header specifies which one of multiple paths that the packet traverses throughout the FabricPath network. The system uses the FTag-specified paths for multidestination packets that enter the FabricPath network. The FTag is a fixed route that the software learns from the topology. The FTag is a 10-bit field with the values from 1 to 1023 (see "Configuring FabricPath Forwarding," for more information on topologies and multiple paths).

This FTag is assigned on the edge port as the frame ingresses the FabricPath network and is honored by all subsequent FabricPath switches in that FabricPath network. Each FTag is unique within one FabricPath topology.

# Default IS-IS Behavior with FabricPath

The interfaces in a FabricPath network run only the FabricPath Layer 2 IS-IS protocol; you do not need to run STP in the FabricPath network because FabricPath Layer 2 IS-IS discovers topology information dynamically.

FabricPath Layer 2 IS-IS is a dynamic link-state routing protocol that detects changes in the network topology and calculates loop-free paths to other nodes in the network. Each FabricPath device maintains a link-state database (LSDB) that describes the state of the network; each device updates the status of the links that are adjacent to the device. The FabricPath device sends advertisements and updates to the LSDB through all the existing adjacencies. FabricPath Layer 2 IS-IS protocol packets do not conflict with standard Layer 3 IS-IS packets because the FabricPath packets go to a different Layer 2 destination MAC address than that used by standard IS-IS for IPv4/IPv6 address families.

The system sends hello packets on the FabricPath core ports to form adjacencies. After the system forms IS-IS adjacencies, the FabricPath unicast traffic uses the equal-cost multipathing (ECMP) feature of Layer 2 IS-IS to forward traffic, which provides up to 16 paths for unicast traffic.

Within the FabricPath network, you use a single control plane protocol, Layer 2 IS-IS, for all unicast, multicast, and broadcast traffic. To use the basic FabricPath functionality, you do not need to configure Layer 2 IS-IS because you can use the default topology. The control plane Layer 2 IS-IS comes up and runs automatically when you enable FabricPath on the device.

The loop-free Layer 2 IS-IS protocol builds two trees for the topology. One tree carries unknown unicast, broadcast, and multicast traffic, and the second tree carries load-balanced multicast traffic. The system load balances multicast traffic across both trees (see "Configuring FabricPath Forwarding," for more information about trees and topology).

FabricPath Layer 2 IS-IS is based on the standard IS-IS protocol with the following extensions for the FabricPath environment:

- FabricPath has a single IS-IS area with no hierarchical Layer 1/Layer 2 routing as prescribed within the IS-IS standard. All devices within the FabricPath network are in a single Layer 1 area.
- Multiple instances of IS-IS can be run, one per set of VLANs/topology.
- The system uses a MAC address that is different from the MAC address used for Layer 3 IS-IS instances.

- The system adds a new sub-TLV that carries switch ID information, which is not in standard IS-IS. This feature allows Layer 2 information to be exchanged through the existing IS-IS protocol implementation.
- Within each FabricPath Layer 2 IS-IS instance, each device computes its shortest path to every other device in the network by using the shortest-path first (SPF) algorithm. This path is used for forwarding unicast FabricPath frames. FabricPath Layer 2 IS-IS uses the standard IS-IS functionality to populate up to 16 routes for a given destination device. The system uses multiple equal-cost available parallel links that provide equal-cost multipathing (ECMP).
- FabricPath IS-IS introduces certain modifications to the standard IS-IS in order to support the construction of broadcast and multicast trees (identified by the FTags). Specifically, using FabricPath, the system constructs two loop-free trees for forwarding multidestination traffic.

Once the adjacency is established among the devices in the FabricPath network, the system sends update information to all neighbors.

By default, you can run Layer 2 IS-IS with FabricPath with no configuration, however, you can fine-tune some of the Layer 2 IS-IS parameters (see "Advanced FabricPath Features," for information about configuring optional IS-IS parameters).

Additionally, FabricPath IS-IS helps to ensure that each switch ID in steady-state is unique within the FabricPath network. If FabricPath networks merge, switch IDs might collide. If the IDs are all dynamically assigned, FabricPath IS-IS ensures that this conflict is resolved without affecting any FabricPath traffic in either network.

# Conversational MAC Address Learning

**Note** You must be working on the F Series module in your Cisco Nexus 7000 Series chassis to use conversational MAC learning.

In traditional MAC address learning, each host learns the MAC address of every other device on the network. When you configure a VLAN for conversational learning, the associated interfaces learn only those MAC addresses that are actively speaking to them. Not all interfaces have to learn all the MAC addresses on an F Series module, which greatly reduces the size of the MAC address tables.

Beginning with Cisco NX-OS Release 5.1 when you use the F Series module, you can optimize the MAC learning process. Conversational MAC learning is configured per VLAN. All FabricPath VLANs always use conversational learning; you can configure CE VLANs for conversational learning on this module also. (See "Configuring FabricPath Forwarding," for more information about CE and FabricPath VLANs.)

The F Series modules have 16 forwarding engines (FEs), and the MAC learning takes place on only one of these FEs. Each FE performs MAC address learning independently of the other 15 FEs on the module. An interface only maintains a MAC address table for the MACs that ingress or egress through that FE; the interface does not have to maintain the MAC address tables on the other 15 FEs on the module.

Conversational MAC address learning and the 16 forward engines (FEs) on each F Series module result in MAC address tables that are much smaller for FabricPath.

The MAC address learning modes available on the F Series modules are the traditional learning and conversational learning. The learning mode is configurable and is set by VLAN mode.

The following VLAN modes have the following MAC learning modes:

- FabricPath (FP) VLANs—Only conversational MAC learning.

- CE VLANs—Traditional learning by default; you can configure CE VLANs on the F Series module for conversational learning.

With conversational MAC learning, the interface learns only the source MAC address of an ingressing frame if that interface already has the destination MAC address present in the MAC address table. If the source MAC address interface does not already know the destination MAC address, it does not learn that MAC address. Each interface learns only those MAC addresses that are actively speaking with the interface. In this way, conversational MAC learning consists of a three-way handshake. The interface learns the MAC address only if that interface is having a bidirectional conversation with the corresponding interface. Unknown MAC address are forwarded, or flooded, throughout the network.

This combination of conversational MAC address learning and multiple FEs on each F Series module produces smaller MAC address tables on each F Series module.

For CE VLANs, you can configure conversational learning per VLAN on the F Series module by using the command-line interface (CLI). CE VLANs use traditional MAC address learning by default. Traditional MAC learning is not supported on FabricPath VLANs with Cisco Release NX-OS 5.1 or later releases.

The figure below shows the allowed FabricPath and CE ports on the M and F Series modules and the allowed FP and CE VLANs.

***Figure 2: FP and CE VLAN Examples***

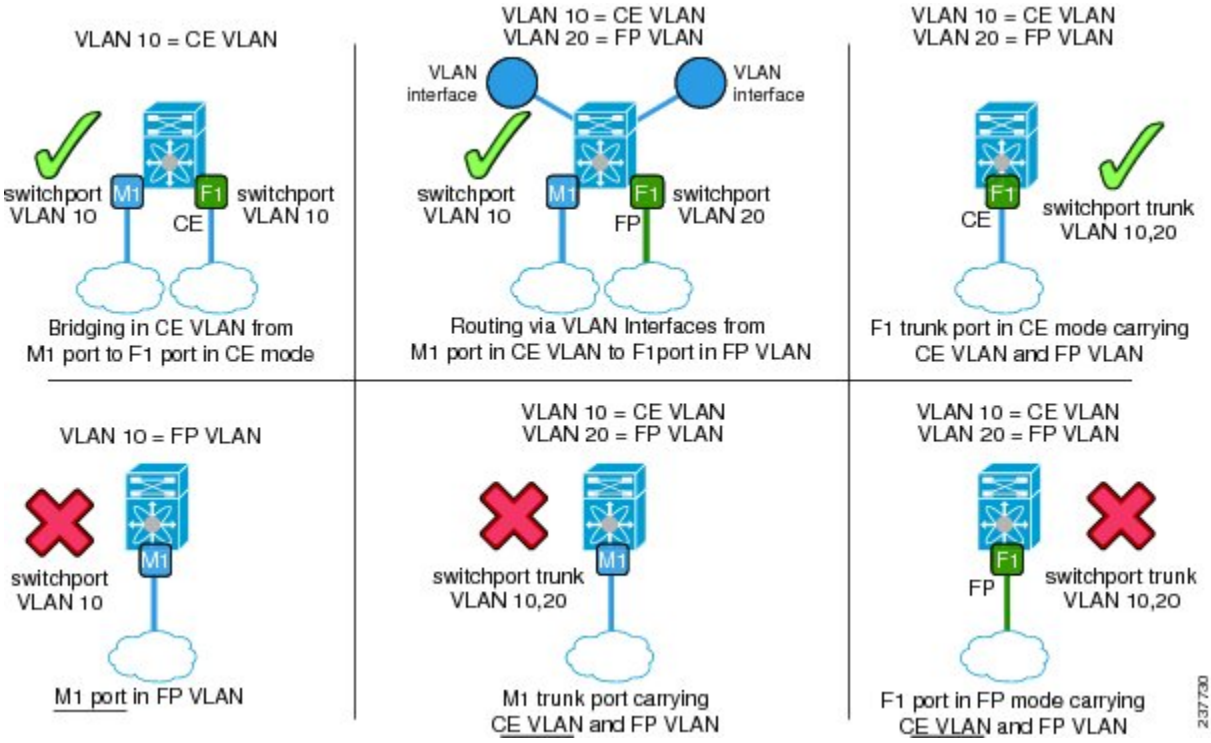


## Core Port Learning

Beginning with Cisco NX-OS Release 6.1, support for a Fabric Extender (FEX) with a virtual port channel + (VPC+) on F2 cards is available. FEX VPCs do not have unique subswitch IDs assigned and use the core port learning mode for forwarding.

With the core port learning mode, all local MACs are copied to the core port forwarding engines (FEs) and the MAC address table for the F2 module displays locally learned MAC addresses that are populated on core ports.

The core port learning mode is enabled by default on F2 VDCs.

Beginning with Cisco NX-OS Release 6.1(2), you can disable MAC address learning on F2 Series modules. All the active or used ports on the port group must be FabricPath core ports.

For VLANs where an SVI exists, the F2 module learns the source MAC addresses from the broadcast frames on the FabricPath core ports, whether the MAC learning is enabled or not. For any port group with MAC learning disabled, the F2 module does not learn the source MAC addresses from the broadcast frames in all the VLANs to which the port group belongs.

# Switching Using FabricPath

The FabricPath hierarchical MAC address scheme and conversational learning result in much smaller, conversational learning MAC tables within the FabricPath network. Within the FabricPath network, the system uses Layer 2 IS-IS to transmit topology information. The interfaces on the edge of the network, which use conversational MAC address learning, do not have to learn all the MAC addresses in the network (see the figure below).

***Figure 3: FabricPath Ports Use Only the FabricPath Header to Switch Frames***

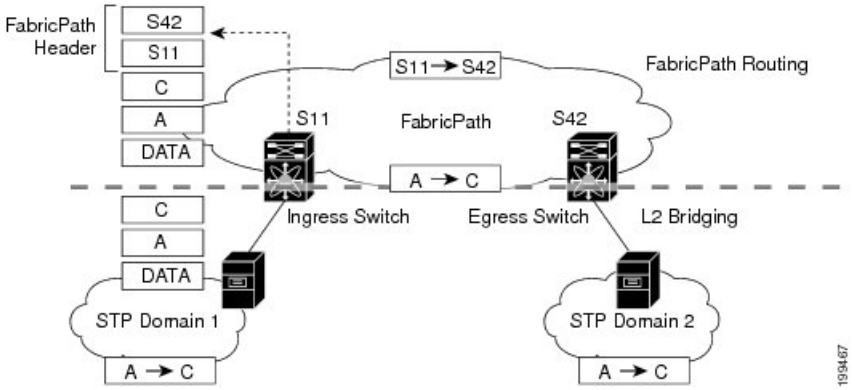


MAC mobility is expedited using the FabricPath hierarchical MAC addresses. That is, when you want to move a host and keep its same MAC address and VLANs, only the interfaces at the edge of the FabricPath network track this change. Within the FabricPath network, the FabricPath interfaces update their tables with only the outer MAC addresses (ODA and OSA) that have changed from the FabricPath encapsulation.

The interface on the edge of the FabricPath network encapsulates the original frame inside the FabricPath header. Once the frame reaches the last, or directly connected, FabricPath switch, the egress interface strips the FabricPath header and forwards the frame as a normal CE frame.

The ports on an F Series module at the edge of a FabricPath network can use conversational learning to learn only those MAC addresses that the specified edge port is having a bidirectional conversation with. Every edge interface does not have to learn the MAC address of every other edge interface; it just learns the MAC addresses of the speakers.

As the frame traverses the FabricPath network, all the devices work only with the FabricPath header. So, the FabricPath interfaces work only with the ODAs and OSAs; they do not need to learn the MAC address for any of the CE hosts or other devices attached to the network. The hierarchical MAC addressing provided by the FabricPath headers results in much smaller MAC tables in the FabricPath network, which are proportional

to the number of devices in that network. The interfaces in the FabricPath network only need to know how to forward frames to another FabricPath switch so they can forward traffic without requiring large MAC address lookup tables in the core of the network.

The switches in the FabricPath network decrement the TTL in the FabricPath header by 1 at each hop. When the TTL reaches 0, the packet is dropped. This process prevents the continuation of any loops that might form in the network.

## FEX Support for an Emulated Switch

Beginning with Cisco NX-OS Release 6.1, support for a FEX with a VPC+ on F2 cards is available. Using VPC+, an emulated switch can be configured using two FEXs.

**Note** For more information about FEXs, see the *Configuring the Cisco Nexus 2000 Series Fabric Extender*.

An example topology of two FEXs acting as an emulated switch is shown in the figure below.

***Figure 4: Two FEXs as an Emulated Switch.***

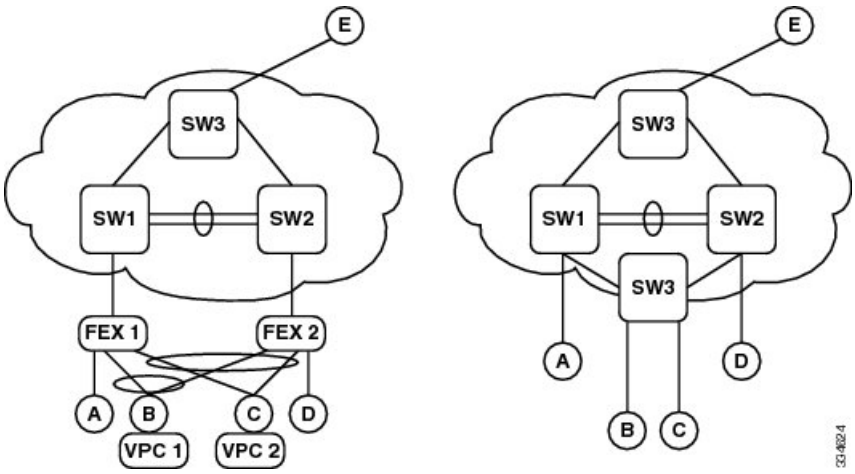


**Note** All the VPC+s of the same FEX have the same outer source address (OSA).

Because a FEX with VPC+ on F2 cards requires core port learning, the subswitch ID and flood ID fields of the outer source MAC addresses are reserved values and are not used.

**Note** Core port learning is enabled by default on F2 VDCs.

FEX orphan ports have the outer source MAC address of the physical switch to which it is connected.

### Partial Mode for FEX with VPC+

To allow a FEX with a VPC+ to function properly, the switch must operate in a partial FTag pruning mode. Traditionally, VPC+ environments operate in an all or none pruning mode where a physical switch is designated

as a primary forwarder. The peer acts as the secondary forwarder if the primary path is down. However, in a FEX with a VPC+ configuration, one switch acts as a designated forwarder for half the available FTags and the other switch forwards the other half. If one of the VPC+ paths is down, the packet is forwarded by the peer switch.

**Note** To configure the FEX port with VPC+, use the **fabricpath multi-cast load balance** command.

## Configuration Example: FEXs with VPC+ for an Emulated Switch

This example shows how to configure FEXs with VPC+ for an emulated switch. The following steps must be executed on both VPC peers.

Before you begin the configuration steps, ensure the following:

- Enable the FabricPath feature set.
- Enable the FEX feature set.

To configure the emulated switch, perform these steps:

1. In the VPC domain configuration mode, enable partial DF mode with the **fabricpath multicast load-balance** command.
2. In the VPC domain configuration mode, configure the emulated switch ID.

```
switch# configure terminal
switch(config)# interface port-channel channel-number
switch(config-if)# vpc domain ID
switch(config-vpc-domain)# fabricpath switch-id emulated switch-id
```

3. Configure a FEX.

```
switch# configure terminal
switch(config)# interface port-channel channel
switch(config-if)# switchport
switch(config-if)# switchport mode fex-fabric
switch(config-if)# fex associate FEX-number
switch(config-if)# no shutdown
switch(config-if)# exit
switch# show interface port-channel channel fex-intf
```

4. Create a FEX Layer 2 host interface (HIF) port channel.

```
switch# configure terminal
switch(config)# interface ethernet FEX-number/1/satellite_port_numnber
switch(config-if)# channel-group id/1001
switch(config-if)# no shutdown
```

5. Configure the VPC ID on the FEX Layer 2 host interface (HIF) port channel.

```
switch# configure terminal
switch# interface port-channel 1001
switch(config-if)# switchport
switch(config-if)# vpc vpcid
switch(config-if)# no shutdown
```

# Inner MAC Address-based Classification

Classification of packets based on Layer 2 header information, such as Destination MAC (DMAC), Source MAC (SMAC) and EtherType, creates a Layer 2 packet flow.. If the packets are coming from a FabricPath interface, classification of packets is based on outer MAC addresses that are present in the FabricPath header. Starting from Cisco NX-OS Release 8.4(1), you can classify packets coming from a FabricPath interface using the inner MAC address. This feature is supported on M3-, F3- and F4-Series I/O modules. The inner MAC address is the actual MAC address of the packet. Use the **flow exclude fabricpath header** command to enable classification of packets using inner MAC addresses and to disable classification of packets using the FabricPath header. Use the no flow exclude fabricpath header to disable classification of packets using inner MAC addresses. The **[no] flow exclude fabricpath header** command is configured on a per-VDC basis. By default, the Inner MAC Address-based Classification feature is disabled.

After you perform an ISSU to the Cisco NX-OS Release 8.4(1), use the **flow exclude fabricpath header** command to enable classification of packets using inner MAC addresses and to disable classification of packets using the FabricPath header. After you perform an ISSD to any release prior to Cisco NX-OS Release 8.4(1), the Inner MAC Address-based Classification feature is disabled. The packets are then classified using outer MAC addresses that are present in the FabricPath header.

# Conflict Resolution and Optional FabricPath Tunings

After you enable FabricPath in all devices, the system automatically assigns a random switch ID to each FabricPath device. The switch ID is a 12-bit value that is dynamically assigned to every switch in the FabricPath network, with each switch being a unique value in that FabricPath network. Optionally, you can configure a specific switch ID. If any of the switch IDs in the FabricPath network are not unique, the system provides automatic conflict resolution.

The FabricPath system chooses a random value for the switch ID and sets this value as tentative during a period when the system waits to hear if this value is already in use. If this value is being used by another device in the network, the system begins a conflict resolution process. The switch with the lower system ID keeps the specified value and the other switch gets a new value for its switch ID.

In the case of a single switch joining an existing FabricPath network, the single switch changes the switch ID value rather than any switches in the existing switches in the network changing values. If the specified value is not in use by another device or after the conflict is resolved, the switch ID is marked as confirmed.

Graceful migration provides that there is no traffic disruption if a conflict arises in the resources, such as two switches that temporarily have the same switch ID.

**Note** The FabricPath interfaces will come up, but they are not operational until the switch checks for FabricPath conflicts and resolves those conflicts.

The FabricPath resource timers have default values, but you can also change the timer values. You can tune the device to wait longer or shorter periods to check the conflicts.

Some of the important processes of the FabricPath network are as follows:

- Achieves a conflict-free allocation of switch IDs and FTags
- Provides graceful resource migration during network merges or partition healing

- Supports static switch IDs
- Provides fast convergence during link bringup or network merge

FabricPath uses the Layer 2 IS-IS protocol to transport the database to all switches in the network. The information is distributed among the FabricPath network devices using an IS-IS TLV. Each switch sends its version of the database that contains information about all the switches. The system allocates the FabricPath values, guarantees their uniqueness within the FabricPath network, and deletes the value from the database once that resource is no longer needed.

**Note** When you manually configure static switch IDs for the device, the automatic conflict resolution process does not work and the network does not come up. You will see syslog messages about the conflict and must manually change one or more switch IDs of the devices in the network.

## FabricPath Timers

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

You can change the following FabricPath timers:

- allocate-delay—Configures the delay for a new switch ID to be propagated throughout the network before that value becomes available and permanent.
- linkup-delay—Configures the link bringup delay to detect conflicts in the switch ID. If the system does find a conflict, the system takes some time to resolve the conflict and bring FabricPath to an operational state. When redundant links are brought up to connect to known networks, the default behavior is to speed up the link bringup. The timer is not used in this case as the network is already known.
- linkup-delay always—Configures the link bringup delay to enforce the timer to be honored in all scenarios.
- transition-delay—Configures the delay for propagating a transitioned value in the network; during this period, all old and new switch ID values exist in the network. This situation occurs only when the link comes up and the system checks to see if the network has two identical switch IDs.

Conflicts that occur with user-configured switch IDs are not resolved. Warning messages are displayed for conflicts of this type. To avoid incorrect traffic forwarding, we recommend that you set the linkup-delay high enough for Intermediate System-to-Intermediate System (IS-IS) to gather neighbor information while changing the topology. A high linkup-delay setting allows the timely detection of conflicts. Links are held down until conflicts are resolved by user intervention or until the expiration of the link-state packet (LSP) of the conflicting switch IDs.

This configuration of timers takes effect only if the link leads to a node that is not yet identified as reachable by the routing protocol. If other equal cost multipaths already exist in the forwarding state and the new link creates another new equal cost multipath, the linkup process might be expedited when the timer configuration is skipped for such links. The timer configuration is used only as a hold time for the routing protocol to gather network information. When networks are known to the routing protocol, you might observe that the timer is not getting used.

The linkup-delay timer is enabled by default If the linkup-delay timer has already been configured when you enable or re-enable this feature, the switch uses the configured timer value. In the absence of a configured linkup-delay timer, the switch uses the default value, which is 10 seconds.

Beginning with Cisco NX-OS Release 6.2(8), you can disable the link-up delay feature using the command line interface (CLI). After you disable the linkup-delay timer, the links are no longer suspended. If the switch detects a conflict, the switch either dynamically resolves this conflict or sends a warning on the system logs, while the links are still operationally up. You can disable the linkup-delay feature to speed up the link bring-up in known networks with statically configured switch IDs. In such networks, there is a guarantee that no conflict in switch IDs will arise and the link suspension is no longer needed for conflict detection.

**Note** Cisco strongly recommends not disabling the linkup-delay feature in networks with dynamically added or unknown switch IDs.

# Interoperation Between the M Series and the F Series Modules

Beginning with Cisco NX-OS Release 8.2(1), FabricPath feature is supported on a VDC that has M3 and F3 Series modules.

Beginning with Cisco NX-OS Release 8.1(1), FabricPath is supported on M3 line cards. FabricPath support is available on an M3 VDC, and not on an M3-F3 mixed VDC.

Beginning with Cisco NX-OS Release 6.2(2), when you have an M Series module and an F Series module in the same Cisco Nexus 7000 Series chassis, you can see the following:

- For an M Series module and an F2e Series module—When talking to the router MAC addresses, MAC address learning occurs on the core ports of the F2e Series modules. This problem is an F2e ASIC limitation and support is provided to disable MAC address learning. See the “Configuring the MAC Learning Mode for Core Ports (Optional)” section. Core and edge ports should not be on the same ASIC or forwarding engine in this scenario because MAC learning is disabled.
- For an M Series module and an F2e Series module—To support F1 access switches in ISSU that do not copy local MAC addresses to the core ports, the M Series and F2e Series modules learn all the remote MAC addresses by default. Support is provided to disable remote MAC address learning. See the “Configuring the Remote MAC Learning Mode (Optional)” section. When all the switches in the FabricPath topology are moved to Cisco NX-OS Release 6.2(2), remote MAC address learning can be disabled.
- For an M Series module and an F2e Series module—To enable proxy learning for Layer 2 on the M Series module, you must disable MAC address learning on the F2e Series module. See the “Configuring the MAC Learning Mode for Core Ports (Optional)” section. You also must disable remote MAC address learning. See the “Configuring the Remote MAC Learning Mode (Optional)” section.
- For an M Series module and an F1 Series module—When talking to all the remote MAC addresses, MAC address learning occurs. After an ISSU to Cisco NX-OS Release 6.2(2) for F1 Series core ports, you can disable remote MAC address learning on the F1 Series core ports. See the "Configuring the Remote MAC Learning Mode (Optional)" section.

Beginning with Cisco NX-OS Release 6.2(2), MAC address learning occurs on M Series module pointing to a gateway port channel (GPC). This scenario occurs in both an M Series module with an F1 Series module and an M Series module with an F2E Series module.

Beginning with Cisco NX-OS Release 6.2(2), when you route using a switch virtual interface (SVI) on an M Series module and that F2e operates in a Layer 2-only mode, the large MAC address table of the M Series module can address up to 128,000 hosts in the FabricPath network.

Beginning with Cisco Release 5.2(1) for the Nexus 7000 Series devices, the MAC learning for the F Series FabricPath-enabled modules when an M Series module is present in the chassis has changed. In this configuration, the FabricPath switches copy all locally learned MAC address entries onto the core port, which is the default learning mode in a chassis that contains both F Series and M Series modules.

When you have an M Series module and an F Series module in the same Cisco Nexus 7000 Series chassis, the FabricPath interface on the F Series modules also learns the MAC addresses that traverse that port from the M Series module. The FabricPath interface provides proxy learning for the MAC addresses on the M Series module in the mixed chassis.

Because M Series modules cannot enable FabricPath, those FabricPath-enabled interfaces that coexist in the same Cisco Nexus 7000 Series chassis do have to learn the MAC addresses of the packets that are traversing the FabricPath-enabled F Series interfaces from the M Series interfaces. The FabricPath interface provides proxy learning for the MAC addresses on the M Series module in the mixed chassis.

See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide* and the *Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide* for more information about interoperation between the F1 Series and M Series modules.

# High Availability

The FabricPath topologies retain their configuration through an in-service software upgrade (ISSU).

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

# Virtual Device Contexts

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

Because of the multiple 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*, 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 Switching

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.
- 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.
- 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 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.
- A maximum of 128 switch IDs can be supported in a FabricPath network.
- FabricPath does not support VTP when in the same VDC. You must disable VTP when the FabricPath feature set is enabled on the VDC.
- On F1, F2e, and F3 series modules, configure FabricPath core ports and CE/vPC+ member ports on separate ASIC instances. If you configure a FabricPath core port and a CE/vPC+ member port on the same ASIC instance, it can result in MAC learning issues. In certain forwarding scenarios, this leads to unicast flooding and traffic blackholing.
- 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.

# Default Setting for FabricPath Switching

***Table 1: Default FabricPath Parameters***

| Parameters | Default |
|---|---|
| FabricPath | Disabled |
| MAC address learning mode | • FP VLANs—Only conversational learning<br>• CE VLANs—Traditional (nonconversational) learning; can be configured for conversational learning on F Series modules |
| allocate-delay timer | 10 seconds |
| linkup-delay timer | 10 seconds |
| transition-delay timer | 10 seconds |
| linkup-delay | Enabled |

| Parameters | Default |
|---|---|
| graceful merge | Enabled |

# Configuring FabricPath Switching

After you enable FabricPath switching on each device, the encapsulation, default IS-IS, and learning occur automatically.

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

Instead of using the default values, you can optionally configure the following FabricPath features manually:

- The MAC learning mode for Classical Ethernet (CE) VLANs:
  - Conversational learning is the only MAC learning mode available for FabricPath (FP) VLANs.
- Various values that the system uses for conflict resolution and other tunings:
  - Switch ID for the device that is used globally in the FabricPath network
  - Timers
  - Graceful merge of FabricPath networks. (Enabled by default. You might experience traffic drops if the feature is disabled.)
  - A one-time forcing of the links to come up

## Enabling the FabricPath Feature Set on the VDC on the Device

You must enable the FabricPath feature set before you can access the commands that you use to configure the feature.

**Note** You must enable the FabricPath feature set on the default VDC, as well as separately on any other VDCs that are running FabricPath. See Configuring Feature-Set for FabricPath for complete information about installing and enabling the FabricPath feature set.

### Before you begin

Ensure that you have installed an F Series module.

### Procedure

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

| | Command or Action | Purpose |
|---|---|---|
| **Step 2** | switch(config)# **feature-set fabricpath** | Enables the FabricPath feature set in the VDC.<br>**Note** You must install the FabricPath feature set before you enable FabricPath on the switch. See *Configuring Feature-Set for FabricPath* for complete information on installing and enabling the FabricPath feature set. Also, you must enable the FabricPath feature set on the default VDC, as well as separately on any other VDCs that are running FabricPath. |
| **Step 3** | switch(config)# **exit** | Exits global configuration mode. |
| **Step 4** | (Optional) switch# **show feature-set** | Displays which feature sets are enabled on the device. |
| **Step 5** | (Optional) switch# **copy running-config startup-config** | Copies the running configuration to the startup configuration. |

### Example

This example shows how to enable the FabricPath feature on the VDC:

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

# Disabling the FabricPath Feature Set on the VDC

**Note** When you disable the FabricPath functionality, the device clears all the FabricPath configurations.

When you disable the FabricPath functionality, you will not see any of the CLI commands that you need to configure FabricPath. The system removes all the FabricPath configurations when you disable the feature set.

**Note** If your FabricPath configuration is large (multiple megabytes in size), disabling the FabricPath functionality may take some time to complete.

### Before you begin

Ensure that you have installed an F Series module.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | switch# **configure terminal** | Enters global configuration mode. |
| **Step 2** | switch(config)# **no feature-set fabricpath** | Disables the FabricPath feature in the VDC. |
| **Step 3** | switch(config)# **exit** | Exits global configuration mode. |
| **Step 4** | (Optional) switch# **show feature-set** | Displays which feature sets are enabled on the device. |
| **Step 5** | (Optional) switch# **copy running-config startup-config** | Copies the running configuration to the startup configuration. |

### Example

This example shows how to disable the FabricPath feature:

```
switch# configure terminal
switch(config)# no feature-set fabricpath
switch(config)#
```

# Configuring the MAC Learning Mode for CE VLANs (Optional)

CE VLANs use traditional learning mode by default. However, you can configure CE VLANs on the F Series modules to use conversational MAC address learning.

**Note** You cannot configure FP VLANs to use traditional MAC address learning; these VLANs use only conversational learning.

### Before you begin

Ensure that you have installed an F Series module.

Ensure that you are working with CE VLANs.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | switch# **configure terminal** | Enters global configuration mode. |
| **Step 2** | switch(config)# **mac address-table learning-mode conversational vlan** *vlan-id* | Configures the specified CE VLAN(s) on F Series modules for conversational MAC learning. Enter the no form of the command to return to traditional (or nonconversational learning) MAC learning mode. The default |

| | Command or Action | Purpose |
|---|---|---|
| | | MAC learning mode for CE VLANs is traditional.<br>**Note** You cannot configure FP VLANs for the traditional MAC address learning mode. |
| **Step 3** | switch(config)# **exit** | Exits global configuration mode. |
| **Step 4** | (Optional) switch# **show mac address-table learning-mode** {**vlan** *vlan-id*} | Displays the VLANs and the MAC learning mode. |
| **Step 5** | (Optional) switch# **copy running-config startup-config** | Copies the running configuration to the startup configuration. |

### Example

This example shows how to configure conversational MAC address learning on specified CE VLANs on the F Series module:

```
switch# configure terminal
switch(config)# mac address-table learning-mode conversational vlan 1-10
switch(config)#
```

# Configuring the Remote MAC Learning Mode (Optional)

By default, the MAC address learning mode is enabled. You can disable or enable remote MAC address learning for a mixed chassis that contains an M Series module and an F2e Series module (M-F2e) or an M Series module and an F1 Series module (M-F1).

### Before you begin

Ensure that you have installed an F Series module.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | switch# **configure terminal** | Enters global configuration mode. |
| **Step 2** | switch(config)# [**no**] **mac address-table fabricpath remote-learning** | Enables the remote MAC address learning mode. To disable the remote MAC address learning mode, enter the no form of this command.<br>**Note** Ensure that all active or used ports in the module or port group are core ports. |
| **Step 3** | switch(config)# **exit** | Exits global configuration mode. |

| | Command or Action | Purpose |
|---|---|---|
| **Step 4** | (Optional) switch# **show system internal l2fm info detail** | Displays the Layer 2 feature manager detailed information. |
| **Step 5** | (Optional) switch# **copy running-config startup-config** | Copies the running configuration to the startup configuration. |

### Example

This example shows how to enable the MAC address learning mode:

```
switch# configure terminal
switch(config)# mac address-table fabricpath remote-learning
```

This example shows how to disable the MAC address learning mode:

```
switch# configure terminal
switch(config)# no mac address-table fabricpath remote-learning
```

# Configuring the MAC Learning Mode for Core Ports (Optional)

By default, the MAC address learning mode is enabled. You can disable or enable MAC address learning on F2 modules. You can also disable or enable MAC address learning for a mixed chassis that contains an M Series module and an F2e Series module. The command is available only in the default or admin VDC.

### Before you begin

Ensure that you have installed an F Series module.

Ensure that you are working in the default VDC.

### Procedure

| | Command or Action | Purpose |
|---|---|---|
| **Step 1** | switch# **configure terminal** | Enters global configuration mode. |
| **Step 2** | switch(config)# [**no**] **hardware fabricpath mac-learning module** *module_number* {**port-group** *port_group*} | Enables the MAC address learning mode for core ports within the specified module. To disable the MAC address learning mode, enter the no form of this command.<br>**Note** Ensure that all active or used ports in the module and port group are core ports. |
| **Step 3** | switch(config)# **exit** | Exits global configuration mode. |
| **Step 4** | (Optional) switch# **show hardware fabricpath mac-learning module** *module* | Displays the module's hardware MAC learning mode. |

| | Command or Action | Purpose |
|---|---|---|
| **Step 5** | (Optional) switch# **show system internal l2fm info detail** | Displays the Layer 2 feature manager detailed information. |
| **Step 6** | (Optional) switch# **copy running-config startup-config** | Copies the running configuration to the startup configuration. |

### Example

This example shows how to enable the MAC learning mode for the given port group on the specified module:

```
switch# configure terminal
switch(config)# hardware fabricpath mac-learning module 4 port-group 1-4
```

This example shows how to disable the MAC learning mode on the specified module:

```
switch# configure terminal
switch(config)# no hardware fabricpath mac-learning module 4
```

# Configuring the Switch ID (Optional)

**Note** You will not lose any traffic during switch ID changes.

By default, FabricPath assigns each FabricPath device with a unique switch ID after you enable FabricPath on the devices. However, you can manually configure the switch ID.

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

### 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)# **fabricpath switch-id** *value* | Specifies the switch ID. The range is from 1 to 4094. There is no default value. |
| **Step 3** | switch(config)# **exit** | Exits global configuration mode. |

| | Command or Action | Purpose |
|---|---|---|
| **Step 4** | (Optional) switch# **show fabricpath switch-id** | Displays information about the switch IDs. |
| **Step 5** | (Optional) switch# **copy running-config startup-config** | Copies the running configuration to the startup configuration. |

### Example

This example shows how to manually configure a device to have the FabricPath switch ID of 25:

```
switch# configure terminal
switch(config)# fabricpath switch-id 25
switch(config)#
```

# Configuring the FabricPath Timers (Optional)

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

You can change the following FabricPath timers:

- allocate-delay
- linkup-delay
- linkup-delay always
- transition-delay

### 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)# **fabricpath timers** {**allocate-delay** *seconds* \| **linkup-delay** *seconds* \| **linkup-delay always** \| **transition-delay** *seconds*} | Specifies the FabricPath timer values. The range is from 1 to 1200 seconds for each of these timers. The default values are as follows:<br>• **allocate-delay**—10 seconds<br>• **linkup-delay**—10 seconds<br>As a best practice, use a linkup-delay timer value of at least 60 seconds before |

| | Command or Action | Purpose |
|---|---|---|
| | | introducing or joining nodes that are statically configured (directly or indirectly) in the network. This setting avoids incorrect traffic forwarding that might result from conflicts between switch IDs.<br>• **linkup-delay always**<br>As a best practice, you should avoid using the **linkup-delay always** keywords in steady state to speed up link bringups. Use this setting to decrease the traffic loss after you reload modules that provide redundant paths to known networks.<br>• **transition-delay**—10 seconds |
| **Step 3** | switch(config)# **exit** | Exits global configuration mode. |
| **Step 4** | (Optional) switch# **show fabricpath timers** | Displays information about FabricPath timers. |
| **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 allocation-delay FabricPath value to 600 seconds:

```
switch# configure terminal
switch(config)# fabricpath timers allocate-delay 600
switch(config)#
```

# Disabling the FabricPath Linkup-Delay (Optional)

**Note** You must make this configuration on each switch that you want to participate in the FabricPath network.

You can disable the linkup-delay feature to speed up the link bring-up in known networks with statically configured switch IDs. In such networks, there is a guarantee that no conflict in switch IDs will arise and the link suspension is no longer needed for conflict detection.

**Note** Cisco strongly recommends not disabling the linkup-delay feature in networks with dynamically added or unknown switch IDs.

### 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)# [**no**] **fabricpath linkup-delay** | Enables and disables the port suspension protocol for conflict resolution. Enabled by default.<br>The timer values take effect only when linkup-delay is enabled.<br>Use the **no** form of this command to disable the linkup-delay feature.<br>**Note** You should not disable the linkup-delay feature in networks with unknown or dynamically derived switch IDs. |
| **Step 3** | switch(config)# **exit** | Exits global configuration mode. |
| **Step 4** | (Optional) switch# **show fabricpath timers** | Displays information about FabricPath timers. |
| **Step 5** | (Optional) switch# **copy running-config startup-config** | Copies the running configuration to the startup configuration. |

### Example

This example shows how to re-enable the linkup-delay on FabricPath:

```
switch# configure terminal
switch(config)# fabricpath linkup-delay
switch(config)#
```

# Disabling FabricPath Graceful Merges (Optional)

**Note** You must make this configuration on each switch that you want to participate in the FabricPath network.

By default, graceful-merge is enabled; you can disable this aspect of the FabricPath feature.

**Note** You might experience traffic drops if you disable this feature.

### 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)# [**no**] **fabricpath graceful-merge disable** | Disables graceful merge of the FabricPath feature. To reenable this feature, enter the no form of the command. |
| **Step 3** | switch(config)# **exit** | Exits global configuration mode. |
| **Step 4** | (Optional) switch# **show running-config** | Displays information about the configuration running on the switch. |
| **Step 5** | (Optional) switch# **copy running-config startup-config** | Copies the running configuration to the startup configuration. |

### Example

This example shows how to disable the graceful merge aspect of the FabricPath feature:

```
switch# configure terminal
switch(config)# fabricpath graceful-merge disable
switch(config)#
```

# Configuring TTL for Unicast and Multicast Packets (Optional)

By default, FabricPath assigns a time to live (TTL) value for unicast and multicast traffic. However, you can overwrite this value.

**Note** The TTL is applied when the packets ingress on edge ports. The TTL value in the packet is only decremented when the packet travels across core ports.

### 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. |

| | Command or Action | Purpose |
|---|---|---|
| **Step 2** | switch(config)# [**no**] **fabricpath ttl unicast** *numhops* | Configures the TTL value for the unicast traffic in the VDC. The range is from 1 to 64 and the default value is 32. |
| **Step 3** | switch(config)# [**no**] **fabricpath ttl multicast** *numhops* | Configures the TTL value for the multicast traffic in the VDC. The range is from 1 to 64 and the default value is 32. |
| **Step 4** | switch(config)# **exit** | Exits global configuration mode. |
| **Step 5** | (Optional) switch# **show fabricpath ttl** | Displays the current TTL configuration for the unicast and multicast traffic. |
| **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 TTL values for multicast and unicast traffic:

```
switch# configure terminal
switch(config)# fabricpath ttl unicast 20
switch(config)# fabricpath ttl multicast 10
switch(config)# exit
switch#
```

# Forcing the Links to Come Up (Optional)

**Note** We do NOT recommend that you use the **fabricpath force link-bringup** command.

As a one-time event, you can force the FabricPath network links to connect if they are not coming up because of switch ID conflicts or other problems in the network.

**Note** You must make this configuration on each switch that you want to participate in the FabricPath network.

**Note** This configuration is not saved when you enter the **copy running-config startup-config** command.

### 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# **fabricpath force link-bringup** | Forces the FabricPath network links to come up as a one-time event.<br>**Note** This command is not saved when you enter the **copy running-config startup-config** command. |

### Example

This example shows how to force the FabricPath network links to come up one time:

```
switch# fabricpath force link-bringup
switch#
```

# Enabling Inner MAC Address-based Classification (Optional)

### Before you begin

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)# **flow exclude fabricpath header** | Enables classification of packets using inner MAC addresses for packets coming from a FabricPath interface. |
| **Step 3** | switch(config)# **exit** | Exits global configuration mode. |
| **Step 4** | (Optional) switch# **copy running-config startup-config** | Copies the running configuration to the startup configuration. |

### Example

This example shows how to enable classification of packets using inner MAC addresses for packets coming from a FabricPath interface:

```
switch# configure terminal
switch(config)# flow exclude fabricpath header
switch(config)#
```

# Verifying FabricPath Switching

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

| Command | Purpose |
|---|---|
| **show feature-set** | Displays whether FabricPath is enabled or not. |
| **show mac address-table learning-mode** {**vlan** *vlan-id*} | Displays the VLANs and the MAC address learning mode.<br>**Note** MAC learning modes are available only on the F Series modules. |
| **show fabricpath conflict** {**all** [*detail*] \| **link** [*detail*] \| **switch-id** [*detail*] \| **transitions** [*detail*]} | Displays information on conflicts in the FabricPath network. |
| **show fabricpath switch-id** [**local**] | Displays information on the FabricPath network by switch ID. |
| **show fabricpath system-id** {*mac-addr*} | Displays information on the FabricPath network by system ID. |
| **show fabricpath timers** | Displays settings for the allocate-delay, linkup-delay, and transition-delay timers for the FabricPath network. |

See “Advanced FabricPath Features,” for more commands that display FabricPath switching functionality.

# Monitoring and Clearing FabricPath Switching Statistics

Use the following commands to display FabricPath switching statistics:

- **clear counters** [*interface*]
- **load-interval** {**interval** *seconds* {**1** | **2** | **3**}}
- **show interface counters** [**module** *module*]
- **show interface counters detailed** [**all**]
- **show interface counters errors** [**module** *module*]

See the *Cisco Nexus 7000 Series NX-OS Interfaces Command Reference* for information about these commands.

# Configuration Example for FabricPath Switching

After installing the feature set (see Configuring Feature-Set for FabricPath for complete information on installing and enabling the FabricPath feature set), you must enable the FabricPath functionality on all the VDCs that you are using.

**Note** You must have an F Series module installed in your Cisco Nexus 7000 Series chassis in order to run FabricPath.

To configure FabricPath switching, follow these steps:

Step 1: Enable FabricPath on all the devices.

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

Step 2 (Optional): Configure MAC address learning mode.

```
switch(config)# mac address learning-mode conversational vlan 1-10
switch(config)# show mac address-table learning-mode
switch(config)# exit
```

Step 3 (Optional): Manually configure a switch ID for the FabricPath device.

```
switch# configure terminal
switch(config)# fabricpath switch-id 25
switch(config)#
```

Step 4: Save the configuration.

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

# Feature History for Configuring FabricPath Switching

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 Switching***

| Feature Name | Release | Feature Information |
|---|---|---|
| Inner MAC Address-based Classification | 8.4(1) | This feature was introduced. Packets coming from a FabricPath interface can be classified using inner MAC addresses. |
| FabricPath support on M3 line cards | 8.1(1) | Added FabricPath support for M3 line cards. FabricPath support is available on an M3 VDC. |
| Linkup-delay | 6.2(8) | Ability to disable the linkup-delay feature. |
| Proxy Layer 2 learning | 6.2(2) | Ability to disable MAC address learning. |
| MAC Proxy | 6.2(2) | Added support for leveraging the MAC address table of an M Series module in order to address up to 128,000 hosts in the FabricPath network. |

| Feature Name | Release | Feature Information |
|---|---|---|
| FabricPath timers | 6.2(2) | Linkup-delay always option introduced. |
| TTL for unicast and multicast packets | 6.2(2) | This feature was introduced. |
| Core port learning | 6.1(1) | This feature was introduced. |
| New default MAC address learning mode in chassis containing both F Series and M Series modules | 5.2(1) | This feature was introduced. |
| FabricPath | 5.1(1) | These features were introduced. |