



Cisco Nexus 7000 Series NX-OS FabricPath Configuration Guide

First Published: 2016-11-24

Last Modified: 2022-08-08

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2016 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1

Preface 1

- Audience 1
- Document Conventions 1
- Related Documentation 3
- Documentation Feedback 3
- Communications, Services, and Additional Information 3

CHAPTER 2

New and Changed Information 5

- New and Changed Information 5

CHAPTER 3

Overview 9

- Licensing Requirements 9
- Information About FabricPath 9
- Information About Conversational MAC Address Learning 10
- Virtualization for FabricPath 11
- High Availability for FabricPath 11

CHAPTER 4

Configuring FabricPath Switching 13

- Finding Feature Information 13
- Information About FabricPath Switching 13
 - FabricPath Encapsulation 14
 - FabricPath Headers 14
 - Forwarding Tags (FTags) 15
 - Default IS-IS Behavior with FabricPath 15
 - Conversational MAC Address Learning 16
 - Core Port Learning 17

Switching Using FabricPath	18
FEX Support for an Emulated Switch	19
Inner MAC Address-based Classification	21
Conflict Resolution and Optional FabricPath Tunings	21
FabricPath Timers	22
Interoperation Between the M Series and the F Series Modules	23
High Availability	24
Virtual Device Contexts	24
Prerequisites for FabricPath	25
Guidelines and Limitations for FabricPath Switching	25
Default Setting for FabricPath Switching	26
Configuring FabricPath Switching	27
Enabling the FabricPath Feature Set on the VDC on the Device	27
Disabling the FabricPath Feature Set on the VDC	28
Configuring the MAC Learning Mode for CE VLANs (Optional)	29
Configuring the Remote MAC Learning Mode (Optional)	30
Configuring the MAC Learning Mode for Core Ports (Optional)	31
Configuring the Switch ID (Optional)	32
Configuring the FabricPath Timers (Optional)	33
Disabling the FabricPath Linkup-Delay (Optional)	34
Disabling FabricPath Graceful Merges (Optional)	35
Configuring TTL for Unicast and Multicast Packets (Optional)	36
Forcing the Links to Come Up (Optional)	37
Enabling Inner MAC Address-based Classification (Optional)	38
Verifying FabricPath Switching	39
Monitoring and Clearing FabricPath Switching Statistics	39
Configuration Example for FabricPath Switching	39
Feature History for Configuring FabricPath Switching	40

CHAPTER 5

Configuring FabricPath Interfaces	43
Finding Feature Information	43
Information About FabricPath Interfaces	43
FabricPath Interfaces	43
STP and the FabricPath Network	44

vPC+	46
Anycast HSRP	48
Designated Forwarder	49
High Availability	49
Virtual Device Contexts	49
Prerequisites for FabricPath	50
Guidelines and Limitations for FabricPath Interfaces	51
Configuring FabricPath Interfaces	52
Configuring FabricPath Interfaces	52
Configuring the STP Priority with Rapid PVST+	53
Configuring the STP Priority with MST	54
Configuring the STP Domain ID for STP Domains Connected to the Layer 2 Gateway Switch	55
Configuring a vPC+ Switch ID	56
vPC+ to vPC Configuration	57
Configuring an Anycast HSRP Bundle	58
Configuring an HSRP Group	58
Configuring an Anycast Bundle	59
Configuring Anycast Bundle Limits	61
Verifying FabricPath Interface Configuration	62
Monitoring FabricPath Interface Statistics	63
Configuration Example for FabricPath Interface	64
Feature History for Configuring FabricPath Interface	65

CHAPTER 6

Configuring FabricPath Forwarding 67

Finding Feature Information	67
Information About FabricPath Forwarding	67
FabricPath Forwarding Overview	67
FabricPath VLANs	68
Forwarding Known Unicast Packets Using ECMP	69
Forwarding Trees for Broadcast, Unknown Unicast, and Multicast Packets	70
Forwarding Multicast Packets	70
High Availability	71
Virtual Device Contexts	71
Load Balancing Using Port Channels	72

Unicast Static Routes in FabricPath	72
Prerequisites for FabricPath	72
Guidelines and Limitations for FabricPath Forwarding	73
Default Settings for FabricPath Forwarding	74
Configuring FabricPath Forwarding	74
Setting the VLAN Mode to FP or CE	75
Configuring FabricPath Unicast Load Balancing (Optional)	76
Configuring FabricPath Multicast Load Balancing (Optional)	78
Configuring FabricPath Increased Multicast Scalability (Optional)	81
Configuring FabricPath Unicast Static Routes	82
Verifying the FabricPath Configuration	83
Configuration Example for FabricPath Forwarding	85
Feature History for Configuring FabricPath Forwarding	86

CHAPTER 7
Advanced FabricPath Features 87

Finding Feature Information	87
Information About FabricPath Advanced Features	87
Information About Advanced FabricPath Layer 2 IS-IS Configurations	87
High Availability	88
Virtual Device Contexts	89
Multiple Topologies	89
Prerequisites for FabricPath	90
Guidelines and Limitations for FabricPath Advanced Features	90
Setting Advanced FabricPath Layer 2 IS-IS Parameters	91
Setting Advanced FabricPath Layer 2 IS-IS Parameters Globally (Optional)	91
Setting Advanced FabricPath Layer 2 IS-IS Parameters per Interface (Optional)	95
Clearing Advanced FabricPath Layer 2 IS-IS Counters	97
Configuring Multiple Topologies	98
Configuring FabricPath IS-IS Multiple Topologies	99
Verifying the FabricPath Advanced Configurations	100
Feature History for Configuring FabricPath Advanced Features	101

CHAPTER 8
Configuring FabricPath in Transit Mode 103

Information about FabricPath in Transit Mode	103
--	-----

Guidelines and Limitations	103
Configuring FabricPath in Transit Mode	104
Feature History for Configuring FabricPath in Transit Mode	104

APPENDIX A

Configuration Limits for Cisco NX-OS FabricPath	105
Configuration Limits for Cisco NX-OS FabricPath	105



CHAPTER 1

Preface

This preface describes the audience, organization, and conventions of the Book Title. It also provides information on how to obtain related documentation.

This chapter includes the following topics:

- [Audience, on page 1](#)
- [Document Conventions, on page 1](#)
- [Related Documentation, on page 3](#)
- [Documentation Feedback, on page 3](#)
- [Communications, Services, and Additional Information, on page 3](#)

Audience

This publication is for experienced network administrators who configure and maintain Cisco NX-OS on Cisco Nexus 7000 Series Platform switches.

Document Conventions



Note

- As part of our constant endeavor to remodel our documents to meet our customers' requirements, we have modified the manner in which we document configuration tasks. As a result of this, you may find a deviation in the style used to describe these tasks, with the newly included sections of the document following the new format.
- The Guidelines and Limitations section contains general guidelines and limitations that are applicable to all the features, and the feature-specific guidelines and limitations that are applicable only to the corresponding feature.

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.

Convention	Description
<i>Italic</i>	Italic text indicates arguments for which the user supplies the values.
[x]	Square brackets enclose an optional element (keyword or argument).
[x y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.
[x {y z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.
variable	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Examples use the following conventions:

Convention	Description
screen font	Terminal sessions and information the switch displays are in screen font.
boldface screen font	Information you must enter is in boldface screen font.
<i>italic screen font</i>	Arguments for which you supply values are in italic screen font.
<>	Nonprinting characters, such as passwords, are in angle brackets.
[]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:



Note

Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.



Caution

Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Related Documentation

Documentation for Cisco Nexus 7000 Series Switches is available at:

- Configuration Guides

<http://www.cisco.com/c/en/us/support/switches/nexus-7000-series-switches/products-installation-and-configuration-guides-list.html>

- Command Reference Guides

<http://www.cisco.com/c/en/us/support/switches/nexus-7000-series-switches/products-command-reference-list.html>

- Release Notes

<http://www.cisco.com/c/en/us/support/switches/nexus-7000-series-switches/products-release-notes-list.html>

- Install and Upgrade Guides

<http://www.cisco.com/c/en/us/support/switches/nexus-7000-series-switches/products-installation-guides-list.html>

- Licensing Guide

<http://www.cisco.com/c/en/us/support/switches/nexus-7000-series-switches/products-licensing-information-listing.html>

Documentation for Cisco Nexus 7000 Series Switches and Cisco Nexus 2000 Series Fabric Extenders is available at the following URL:

<http://www.cisco.com/c/en/us/support/switches/nexus-2000-series-fabric-extenders/products-installation-and-configuration-guides-list.html>

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to nexus7k-docfeedback@cisco.com. We appreciate your feedback.

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at [Cisco Profile Manager](#).
- To get the business impact you're looking for with the technologies that matter, visit [Cisco Services](#).
- To submit a service request, visit [Cisco Support](#).
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit [Cisco Marketplace](#).
- To obtain general networking, training, and certification titles, visit [Cisco Press](#).
- To find warranty information for a specific product or product family, access [Cisco Warranty Finder](#).

Cisco Bug Search Tool

[Cisco Bug Search Tool](#) (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.



CHAPTER 2

New and Changed Information

- [New and Changed Information](#), on page 5

New and Changed Information

Your software release might not support all the features in this document. 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.

Table 1: New and Changed FabricPath Features

Feature	Description	Changed in Release
Inner MAC Address-based Classification	Packets coming from a FabricPath interface can now be classified using inner MAC addresses.	8.4(1)
FabricPath Support on VDC having M3 and F3 modules	FabricPath feature is supported on a VDC that has M3 and F3 Series modules.	8.2(1)
FabricPath support on M3 modules	This feature was introduced. FabricPath support was added for M3 modules. FabricPath support is available on an M3 VDC.	8.1(1)
vPC+ to vPC configuration	Changed warning prompt message and added requirement for all vPCs to be deleted and reconfigured.	6.2(10)
HSRP Anycast	Added the ability to add or delete a VLAN to or from an existing VLAN range (for an HSRP Anycast bundle) without having to enter the complete VLAN range again.	6.2(10)
FabricPath timers	Added the linkup-delay always option.	6.2(2)

Feature	Description	Changed in Release
TTL for Unicast and Multicast Packets	Added TTL for unicast and multicast packets.	6.2(2)
Unicast static routes in FabricPath	Added support for unicast static routes.	6.2(2)
Advanced FabricPath Layer 2 IS-IS Parameters Globally	Added support for VLAN pruning.	6.2(2)
Advanced FabricPath Layer 2 IS-IS Parameters Globally	Added support for the overload bit.	6.2(2)
Advanced FabricPath Layer 2 IS-IS Parameters	Added support for route-map redistribution and mesh group.	6.2(2)
Multiple Topologies	Added support to create multiple topologies.	6.2(2)
Proxy Layer 2 Learning	Added Proxy Layer 2 learning that disables core port learning in a mixed chassis system.	6.2(2)
MAC Proxy	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.	6.2(2)
Anycast HSRP	Provided capability to create an anycast HSRP bundle to support scalability on the spine layer.	6.2(2)
Configuring more than 244 vPC+ port channels	Added support for configuring more than 244 vPC+ port channels with the no port-channel limit command.	6.1(3)
Configuring vPC+ with FEX ports	Added support for configuring vPC+ with FEX ports with the fabricpath multicast load-balance command.	6.1(3)
FEX Support for an Emulated Switch	Added support to emulate a switch using FEXs.	6.1(2)
Core Port Learning	Core port learning introduced to support forwarding for FEX with VPC+ on F2 cards.	6.1(1)
Load Balancing Using Port Channels	Load balancing to support F2 modules introduced.	6.0(1)

Feature	Description	Changed in Release
New default MAC learning address method for mixed chassis	Created new default method for learning MAC addresses in a chassis containing an F Series and an M Series module.	5.2(1)



CHAPTER 3

Overview

This chapter provides an overview of the FabricPath and conversational MAC address learning features that are supported by the Cisco NX-OS software for the Cisco Nexus 7000 Series devices.

- [Licensing Requirements, on page 9](#)
- [Information About FabricPath, on page 9](#)
- [Information About Conversational MAC Address Learning, on page 10](#)
- [Virtualization for FabricPath, on page 11](#)
- [High Availability for FabricPath, on page 11](#)

Licensing Requirements

For a complete explanation of Cisco NX-OS licensing recommendations and how to obtain and apply licenses, see the [Cisco NX-OS Licensing Guide](#).

Information About FabricPath

Beginning with the Cisco NX-OS Release 5.1 and when you use an F Series module, you can use the FabricPath feature.



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

The FabricPath feature provides the following:

- Allows Layer 2 multipathing in the FabricPath network.
- Provides built-in loop prevention and mitigation with no need to use the Spanning Tree Protocol (STP).
- Provides a single control plane for unknown unicast, unicast, broadcast, and multicast traffic.
- Enhances mobility and virtualization in the FabricPath network.

The system randomly assigns a unique switch ID to each device that is enabled with FabricPath.

When a frame enters the FabricPath network from a Classical Ethernet (CE) network, the ingressing interfaces encapsulate the frame with a FabricPath header. The system builds paths, called trees, through the FabricPath network and assigns a forwarding tag (FTag) by flow to all the traffic in the FabricPath network. When the frame leaves the FabricPath network to go to a CE network, the egressing interface decapsulates the frame and leaves the regular CE header.



Note Classical Ethernet is referred to as CE in this document.

The FabricPath network uses the Layer 2 Intermediate System-to-Intermediate System (IS-IS) protocol to forward traffic in the network using the FabricPath headers. Layer 2 IS-IS is different than Layer 3 IS-IS; the two protocols work independently. Layer 2 IS-IS requires no configuration and becomes operational when you enable FabricPath on the device. The frames carry the same FTag that is assigned at ingress throughout the FabricPath network, and Layer 2 IS-IS allows all devices to have the same view of all the trees built by the system. Known unicast traffic uses the Equal Cost Multipath Protocol (ECMP) to forward traffic throughout the network. Finally, using ECMP and the trees, the system automatically load balances traffic throughout the FabricPath network.

FabricPath provides configuration simplicity, scalability, flexibility, and resiliency within a Layer 2 domain.



Note Precision Time Protocol (PTP) over FabricPath is not supported.

Information About Conversational MAC Address Learning

Beginning with Cisco NX-OS Release 5.1 and when you use an F Series module, you can use conversational MAC address learning. You configure the type of MAC address learning—conversational or traditional—by VLAN.

Conversational MAC address learning means that each interface learns only those MAC addresses for interested hosts, rather than all MAC addresses in the domain. 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.

This selective learning, or conversational MAC address learning, allows you to scale the network beyond the limits of individual switch MAC address tables.

All FabricPath VLANs use conversational MAC address learning.

CE VLANs use traditional MAC address learning by default, but you can configure the CE VLANs to use conversational MAC learning.

Beginning with Cisco NX-OS Release 6.1, support for a Fabric Extender (FEX) with VPC+ on F2 cards is available. To support forwarding with this approach, core port learning is used.

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

Virtualization for FabricPath

You can create multiple virtual device contexts (VDCs). Each VDC is an independent logical device to which you can allocate interfaces. Once an interface is allocated to a VDC, you can only configure that interface if you are in the correct VDC. For more information on VDCs, see the *Virtual Device Context Configuration Guide, Cisco DCNM for LAN, Release 6.x*.

High Availability for FabricPath

FabricPath retains the configurations across ISSU.

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



CHAPTER 4

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 13](#)
- [Information About FabricPath Switching, on page 13](#)
- [Prerequisites for FabricPath, on page 25](#)
- [Guidelines and Limitations for FabricPath Switching, on page 25](#)
- [Default Setting for FabricPath Switching, on page 26](#)
- [Configuring FabricPath Switching, on page 27](#)
- [Verifying FabricPath Switching, on page 39](#)
- [Monitoring and Clearing FabricPath Switching Statistics, on page 39](#)
- [Configuration Example for FabricPath Switching, on page 39](#)
- [Feature History for Configuring FabricPath Switching, on page 40](#)

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 ingress edge port of the FabricPath network and decapsulates the frame on the egress 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 multideestination 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.

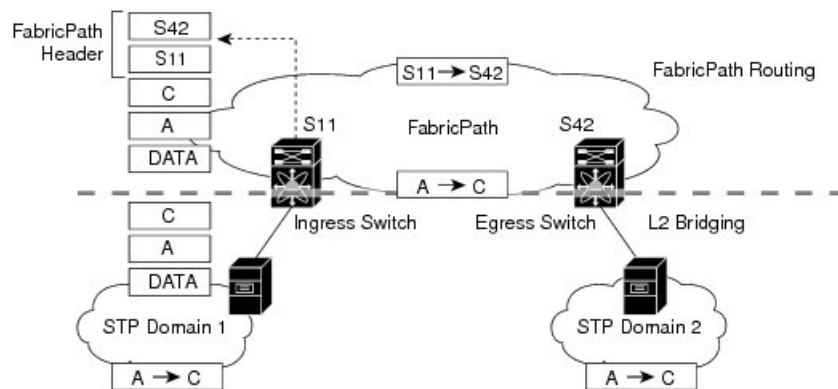
Beginning with Cisco NX-OS Release 6.1(2), you can disable MAC address learning on F2 Series modules. All the active or used ports on the port group must be FabricPath core ports.

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_number
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 ISSU 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 interoperability 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 2: Default FabricPath Parameters

Parameters	Default
FabricPath	Disabled
MAC address learning mode	<ul style="list-style-type: none"> • FP VLANs—Only conversational learning • 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. Note You must install the FabricPath feature set before you enable FabricPath on the switch. See <i>Configuring Feature-Set for FabricPath</i> 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 <i>vlan-id</i>	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. 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. 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 <i>module_number</i> {port-group <i>port_group</i>}	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. 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 <i>module</i>	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 <i>seconds</i> linkup-delay <i>seconds</i> linkup-delay always transition-delay <i>seconds</i> }	Specifies the FabricPath timer values. The range is from 1 to 1200 seconds for each of these timers. The default values are as follows: <ul style="list-style-type: none"> • allocate-delay—10 seconds • linkup-delay—10 seconds As a best practice, use a linkup-delay timer value of at least 60 seconds before

	Command or Action	Purpose
		<p>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.</p> <ul style="list-style-type: none"> • linkup-delay always <p>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.</p> <ul style="list-style-type: none"> • 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	<p>Enables and disables the port suspension protocol for conflict resolution. Enabled by default.</p> <p>The timer values take effect only when linkup-delay is enabled.</p> <p>Use the no form of this command to disable the linkup-delay feature.</p> <p>Note You should not disable the linkup-delay feature in networks with unknown or dynamically derived switch IDs.</p>
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 <i>numhops</i>	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 <i>numhops</i>	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. 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 <i>vlan-id</i> }	Displays the VLANs and the MAC address learning mode. Note MAC learning modes are available only on the F Series modules.
show fabricpath conflict {all [<i>detail</i>] link [<i>detail</i>] switch-id [<i>detail</i>] transitions [<i>detail</i>]}	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 { <i>mac-addr</i> }	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 3: 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.



CHAPTER 5

Configuring FabricPath Interfaces

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

- [Finding Feature Information, on page 43](#)
- [Information About FabricPath Interfaces, on page 43](#)
- [Prerequisites for FabricPath, on page 50](#)
- [Guidelines and Limitations for FabricPath Interfaces, on page 51](#)
- [Configuring FabricPath Interfaces, on page 52](#)
- [Verifying FabricPath Interface Configuration, on page 62](#)
- [Monitoring FabricPath Interface Statistics, on page 63](#)
- [Configuration Example for FabricPath Interface, on page 64](#)
- [Feature History for Configuring FabricPath Interface, on page 65](#)

Finding Feature Information

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

Information About FabricPath Interfaces



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

FabricPath Interfaces

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

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

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



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

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




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

STP and the FabricPath Network



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

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

Figure 5: STP Boundary Termination at FabricPath Network Border

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

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



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

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

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

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

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

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

- Flushes all remote MAC addresses for that STP domain and the MAC addresses on the designated port.
- Propagates the TCN to the other devices in the specified STP domain.

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

vPC+

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



Note

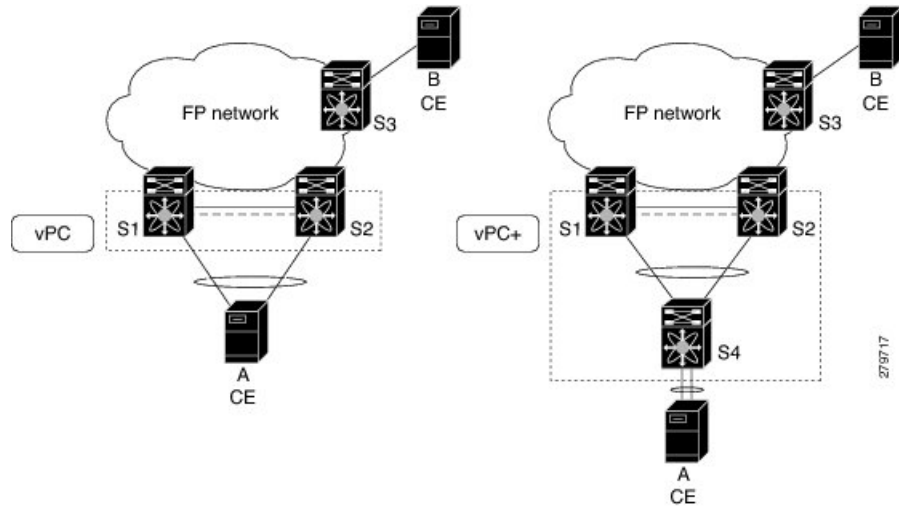
- vPC+ is an extension to virtual port channels (vPCs) that run CE only (see the “Configuring vPCs” chapter in the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide*). You cannot configure a vPC+ domain and a vPC domain in the same VDC.
- In a vPC+ system running 7.2(0)D1(0.444S4), the mroutes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

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

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

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

Figure 6: vPC/vPC+



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

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

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

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

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

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



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

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

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

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

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

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

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

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

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



Note If a data center outage occurs and you enable HSRP before the vPC+ successfully comes up, traffic loss can occur. You need to enable an HSRP delay to give the vPC time to stabilize. If you enable both an HSRP delay and a preemption delay, the Cisco Nexus 7000 Series devices allow Layer 2 switching only after both timers expire.

The delay option is available only with HSRP. If you use any other FHRP, traffic loss is still possible.

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

Anycast HSRP

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

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

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



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

Designated Forwarder

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

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

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

There are three designated forwarder states for a vPC port:

- All—If the local vPC leg is up and the peer vPC is not configured or down, the local switch is the designated forwarder for all RBHs/FTAGs for that vPC.
- Partial—If the vPC path is up on both sides, each peer is the designated forwarder for half the RBHs or FTAGs. For the latter, the vPC port allows only the active FTAGs on that peer.
- None—If the local vPC path is down or not configured, the local switch does not forward any multi destination packets from this vPC path.

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



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

High Availability

The FabricPath topologies retain their configuration through ISSU.

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

Virtual Device Contexts

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

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

Table 4: Port Pairs and Port Sets for F Series Modules

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

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

Prerequisites for FabricPath

FabricPath forwarding has the following prerequisites:

- You should have a working knowledge of Classical Ethernet Layer 2 functionality.
- You must install the FabricPath feature set in the default and nondefault VDC before you enable FabricPath on the switch. See the Configuring Feature Set for FabricPath for complete information on installing and enabling the FabricPath feature set.
- The FabricPath feature set operation might cause the standby supervisor to reload if it is in an unstable state, such as following a service failure or powering up.
- You are logged onto the device.

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

Guidelines and Limitations for FabricPath Interfaces

FabricPath switching has the following configuration guidelines and limitations:

- FabricPath interfaces carry only FabricPath-encapsulated traffic.
- You enable FabricPath on each device before you can view or access the commands. Enter the **feature-set fabricpath** command to enable FabricPath on each device. See *Configuring Feature-Set for FabricPath* for complete information on installing and enabling the FabricPath feature set.
- STP does not run inside a FabricPath network.
- Set the STP priority value on all FabricPath Layer 2 gateway devices to 8192.
- The F Series modules do not support multiple SPAN destination ports or virtual SPAN. If a port on an F Series module is in a VDC and that VDC has multiple SPAN destination ports, that SPAN session is not brought up.
- The following guidelines apply to private VLAN configuration when you are running FabricPath:
 - All VLANs in a private VLAN must be in the same VLAN mode; either CE or FabricPath. If you attempt to put different types of VLANs into a private VLAN, these VLANs will not be active in the private VLAN. The system remembers the configurations, and if you change the VLAN mode later, that VLAN becomes active in the specified private VLAN.
 - FabricPath ports cannot be put into a private VLAN.
- The system does not support hierarchical static MAC addresses. That is, you cannot configure static FabricPath ODAs or OSAs; you can only configure CE static MAC addresses.
- On the F Series modules, user-configured static MAC addresses are programmed on all forwarding engines (FEs) that have ports in that VLAN.
- Pruning does not occur in a virtual port channel (vPC) domain. In a vPC domain, all switches receive multicast traffic, but only one switch forwards the traffic to the receiver.
- A single vPC+ domain between two VDCs on the same physical Cisco Nexus 7000 device is not supported.
- At least one FabricPath interface must be operational on a device for multidestination traffic to be forwarded on vPC+ member ports.
- Support for more than 244 vPC+ port channels (per vPC+ domain) is enabled with the **no port-channel limit** command.
 - Only VDCs that have an F2 series module can support more than 244 vPC+ port channels.
 - The **fabricpath multicast load-balance** command must be entered before the **no port-channel limit** command.



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

- An anycast HSRP bundle provides the support for more than two nodes at the spine layer.
- An anycast HSRP bundle is supported only in HSRP version 2.
- Because of a limitation with an ASIC on the 32-port 1/10-Gigabit Ethernet F1 Series module, a packet that egresses from that module through both ports in FabricPath VLAN mode has an incorrect outer source address (OSA) if the first port is configured as a FabricPath edge port and the second port is configured as a FabricPath core port. To work around this issue, configure the first port as a FabricPath core port and the second port as a FabricPath edge port.
- Beginning with Cisco NX-OS Release 6.2(2), SSM is supported on virtual port channel+ (vPC+).
- When multicast routing is occurring on a FabricPath spine switch, the egress core ports towards the FabricPath leaf switches should not have a mix of F2e and F3 Series module ports. This may cause multicast traffic to be forwarded on both FTags, which can lead to duplicate multicast traffic received at the destination leaf switch, depending on the topology. This limitation only affects Layer-3 routed multicast traffic.

Configuring FabricPath Interfaces



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



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

Configuring FabricPath Interfaces

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



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

Before you begin

Ensure that you are working on an F Series module.

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

Procedure

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

Example

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

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

Configuring the STP Priority with Rapid PVST+

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

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

Before you begin

Ensure that you are working on an F Series module.

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

Procedure

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

Example

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

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

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

Configuring the STP Priority with MST

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

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

Before you begin

Ensure that you are working on an F Series module.

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

Procedure

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

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

Example

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

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

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

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

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

Before you begin

Ensure that you are working on an F Series module.

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

Procedure

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

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

Example

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

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

Configuring a vPC+ Switch ID



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

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



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



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

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

Before you begin

Ensure that you are working on an F Series module.

Ensure that you have enabled the vPC feature.

Ensure that you have enabled the FabricPath feature.

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

Procedure

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

Example

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

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

vPC+ to vPC Configuration

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

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# vpc domain <i>domain-id</i>	Enters the vpc-domain configuration mode for configuration purposes.
Step 3	switch(config-vpc-domain)# no fabricpath switch-id <i>switch-id</i>	Deconfigures the FabricPath switch ID.
Step 4	Perform one of the following: <ul style="list-style-type: none"> For Cisco NX-OS Release 6.2(10) or a later release, enter yes at the following prompt: <pre>Deconfiguring fabricpath switch id will flap vPCs. vPC+ to vPC transition needs reconfiguration of vPCs for this release, please refer to</pre> 	

	Command or Action	Purpose
	<p>configuration guide for more details. Continue (yes/no)? [no]</p> <ul style="list-style-type: none"> For releases prior to Cisco NX-OS Release 6.2(10), enter yes at the following prompt: <p>Deconfiguring fabricpath switch id will flap vPCs. Continue (yes/no)? [no]</p>	
Step 5	For Cisco NX-OS Release 6.2(10) or a later release, delete and reconfigure all vPCs.	

Configuring an Anycast HSRP Bundle

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



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

Configuring an HSRP Group

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

Before you begin

Ensure that you are working on an F Series module.

Ensure that you have enabled the FabricPath.

Ensure that you have enabled the HSRP feature.

Ensure that you have enabled the interface VLAN feature.

Procedure

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

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

Example

This example shows how to configure an HSRP group:

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

switch# show hsrp
```

Configuring an Anycast Bundle

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

Before you begin

Ensure that you are working on an F Series module.

Ensure that you have enabled the FabricPath.

Ensure that you have enabled the HSRP feature.

Ensure that you have enabled the interface VLAN feature.



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

Procedure

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

	Command or Action	Purpose
Step 2	switch(config)# [no] hsrp anycast bundle-id { ipv4 ipv6 both }	Configures an anycast bundle. The arguments and keywords are as follows: <ul style="list-style-type: none"> • bundle-id—Bundle ID. The range is from 1 to 4096. • ipv4—Specifies an IPv4 bundle. All the IPv4 groups in the interface are associated with this bundle. • ipv6—Specifies an IPv6 bundle. All the IPv6 groups in the interface are associated with this bundle. • both—Specifies an IPv4 and IPv6 bundle. This is the default. All the IPv4 and IPv6 groups in the interface are associated with this bundle.
Step 3	switch(config-anycast-bundle)# [no] force gateway-down	Enforces the anycast bundle to remain in the down state even if one invalid VLAN is configured for the bundle.
Step 4	switch(config-anycast-bundle)# [no] switch-id <i>asid</i>	Configures the switch ID for the anycast bundle.
Step 5	switch(config-anycast-bundle)# vlan <i>range</i>	Configures the VLAN range for the anycast bundle. <p>Note Beginning with Cisco NX-OS Release 6.2(10), you can add or delete a VLAN to or from an existing VLAN range for the anycast bundle without having to enter the complete VLAN range again.</p>
Step 6	switch(config-anycast-bundle)# [no] priority <i>priority_value</i>	Configures the priority for the anycast bundle. This value is used to elect a root for all the groups in the range. The range is from 1 to 127 and the default value is 100.
Step 7	switch(config-anycast-bundle)# [no] track <i>object_id</i>	Configures the tracking value that is used to track the anycast bundle. The range is from 1 to 500 and the default value is 0, which indicates that nothing is tracked.
Step 8	switch(config-anycast-bundle)# [no] timer <i>hello_interval</i>	Configures the timers for the groups using this anycast bundle. The default value is 3.
Step 9	switch(config-anycast-bundle)# [no] shutdown	Configures the group to take the switch out of the anycast bundle. To bring the switch into

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

Example

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

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

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

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

Configuring Anycast Bundle Limits

You can create limits for the anycast bundles.

Before you begin

Ensure that you are working on an F Series module.

Ensure that you have enabled the FabricPath.

Ensure that you have enabled the HSRP feature.

Ensure that you have enabled the interface VLAN feature.

Procedure

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

Example

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

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

Verifying FabricPath Interface Configuration

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

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

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

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

Monitoring FabricPath Interface Statistics

Use the following commands to display FabricPath statistics:

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

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

Configuration Example for FabricPath Interface

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

- Enable FabricPath on each device.
- Configure the interfaces that you want to designate as FabricPath interfaces.
- Set the STP priority device to 8192 on all FabricPath Layer 2 gateway devices.
- (Optional) Set the STP domain ID for each of the separate STP domains that are connected to the FabricPath network.
- (Optional) Configure a vPC+ switch ID.

To configure FabricPath interfaces, follow these steps:

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

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

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

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

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

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

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

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

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

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

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


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



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

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

1. In the vPC domain configuration mode, enter the **fabricpath switch-id** *switch-id* command.
2. On each of the vPC+ peer link interfaces in interface configuration mode, enter the **switchport mode fabricpath** command.
3. On each vPC+ peer link port channel, enter the **vpc peer-link** command.

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

1. On each vPC peer link port channel, enter the **shutdown** command.
2. In the vPC domain configuration mode, enter the **fabricpath switch-id** *switch-id* command.
3. On each of the vPC+ peer link interfaces in interface configuration mode, enter the **switchport mode fabricpath** command.
4. On each vPC+ peer link port channel, enter the **no shutdown** command.

Step 7: Save the configuration.

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

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

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

Feature History for Configuring FabricPath Interface

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

Table 5: Feature History for FabricPath Interface

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



CHAPTER 6

Configuring FabricPath Forwarding

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

- [Finding Feature Information, on page 67](#)
- [Information About FabricPath Forwarding, on page 67](#)
- [Prerequisites for FabricPath, on page 72](#)
- [Guidelines and Limitations for FabricPath Forwarding, on page 73](#)
- [Default Settings for FabricPath Forwarding, on page 74](#)
- [Configuring FabricPath Forwarding, on page 74](#)
- [Verifying the FabricPath Configuration, on page 83](#)
- [Configuration Example for FabricPath Forwarding, on page 85](#)
- [Feature History for Configuring FabricPath Forwarding, on page 86](#)

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 Forwarding



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

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 “Configuring FabricPath Switching,” for more information about 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 multideestination 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 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 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 the 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 multideestination 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, “Configuring FabricPath Interfaces,” for information about 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 7: Example FabricPath Topology and Classical Ethernet Hosts



The figure above shows a sample FabricPath topology with Classical Ethernet switches and FP/CE VLANs.

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.

For best practices, consistent VLAN configuration within a FabricPath topology is a good practice because FabricPath does not perform topology calculations on a per-VLAN basis. Therefore, if a VLAN is not defined on a particular Cisco FabricPath switch that belongs to a specific topology, the control plane will not be aware of it and may try to forward traffic for this VLAN through this particular switch, with the result that the traffic is black-holed. Note that with Cisco FabricPath, core ports forward traffic only for VLANs that are defined in the switch. The loss of traffic that is caused by lack of the required VLANs in the VLAN database is especially difficult to troubleshoot.

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 “Configuring FabricPath Switching,” for more information about 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.

Forwarding Trees for Broadcast, Unknown Unicast, and Multicast Packets

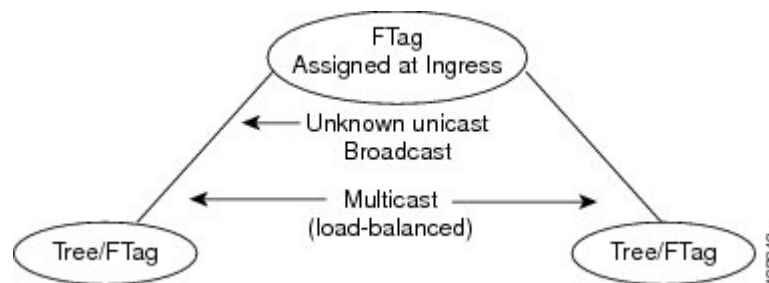
FabricPath introduces a new loop-free broadcast functionality that carries broadcast, unknown unicast, and multicast packets, or multdestination traffic. For each broadcast, unknown unicast, and multicast traffic flow, the system chooses the forwarding path from among multiple system-created paths or trees. The system creates two trees to forward the multdestination 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 multdestination 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 multdestination traffic flow follows through the FabricPath network. The system assigns the trees per flow.

The figure below shows these trees.

Figure 8: Trees for Forwarding Multdestination 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 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 the “Configuring FabricPath Increased Multicast Scalability (Optional)” section on for more information.

For Layer 2 multicast traffic, you do not need to run PIM 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* for more information on high availability.

Virtual Device Contexts

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

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

Table 6: Port Pairs and Port Sets for F Series Modules

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

Port Pairs for F1 Modules	Port Sets for F2 Modules
Ports 23 and 24	Ports 45, 46, 47, 48
Ports 25 and 26	
Ports 27 and 28	
Ports 29 and 30	
Ports 31 and 32	

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

Load Balancing Using Port Channels

The Cisco NX-OS software load balances traffic across all operational interfaces in a port channel by hashing the addresses in the frame to a numerical value that selects one of the links in the channel. Port channels provide load balancing by default. Port-channel load-balancing uses MAC addresses, IP addresses, or Layer 4 port numbers to select the link. Port-channel load balancing uses either source or destination addresses or ports, or both source and destination addresses or ports.

See the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide* for more information about load balancing.

Unicast Static Routes in FabricPath

FabricPath uses Layer 2 Integrated Intermediate System-to-System (IS-IS) as a link state protocol to compute unicast topologies. You can configure unicast static routes in the forwarding tables to ensure a predictable operation of the network.

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 Forwarding

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.
- Starting from Cisco NX-OS Release 8.2(2), configure the **no fabricpath load-balance multicast include-vlan** command on any VDC in which FabricPath is configured along with both F2e-Series and F3-Series I/O modules.
- 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.
- Because of a limitation with an ASIC on the 32-port 1/10-Gigabit Ethernet F1 Series module, a packet that egresses from that module through both ports in FabricPath VLAN mode has an incorrect outer source address (OSA) if the first port is configured as a FabricPath edge port and the second port is configured as a FabricPath core port. To work around this issue, configure the first port as a FabricPath core port and the second port as a FabricPath edge port.
- Beginning with Cisco NX-OS Release 6.2(2), FabricPath supports unicast static routes. It does not support multicast static routes.
- 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.
- Multicast traffic sent to a group with no receivers present might not be constrained to the router port optimized multicast flooding (OMF) entry for a VLAN. The OMF entry is maintained on a per-VDC basis, not on a per-VLAN basis, which means that if multiple ports are members of the OMF entry, the ports that forward the FTag also forward the multicast traffic.

- Use the **show fabricpath mroute vdc-omf** command to view all ports forwarding on the OMF entry.
- Use the **show fabricpath mroute omf resolved ftag [ftag]** command to view all resolved OMF entries on a per-FTag basis.
- 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.
- Extending the FabricPath VLANs over the VPLS infrastructure is not supported. Only regular Ethernet VLANs can be extended over VPLS.

Default Settings for FabricPath Forwarding

Table 7: Default FabricPath Parameters

Parameters	Default
FabricPath Topology	0
VLAN mode	CE

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 “Configuring 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.

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 enabled the FabricPath feature.

Ensure that you have created the VLANs

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	(config)# vlan <i>vlan-id</i>	Enters the VLAN configuration mode and identifies those VLANs that you want to carry FabricPath traffic.
Step 3	switch(config-vlan)# mode [ce fabricpath]	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	switch(config-vlan)# exit	Exits 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	switch(config)# exit	Exits global configuration mode.
Step 6	(Optional) switch# show fabricpath topology vlans [active]	Displays information about all VLANs in the FabricPath topology.
Step 7	(Optional) switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to specify VLANs as FP VLANs:

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

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 enabled the FabricPath feature.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	<ul style="list-style-type: none"> switch(config)# [no] fabricpath load-balance {source source-destination xor destination symmetric} switch(config)# [no] port-channel load-balance [algorithm [module module]] 	<p>To configure source/destination/symmetric/src-dst algorithms for load-balancing FabricPath unicast traffic in vDCs that do not allow F2 resource types, use the fabricpath load-balance command.</p> <p>To configure source/destination/symmetric/src-dst algorithms for load-balancing FabricPath unicast traffic in vDCs that allow F2 resource types, use the port-channel load-balance command.</p>
Step 3	switch(config)# [no] fabricpath load-balance unicast [layer 3 layer 4 mixed] [rotate-amount rot_amt] [include-vlan]	<p>Configures options such as rotation-skew, VLAN inclusion, use of Layer 3/Layer 4 traffic parameters, for load-balancing FabricPath unicast traffic.</p> <p>Note To return to the default unicast load-balancing scheme, enter the no form of this command.</p>
Step 4	switch(config)# exit	Exits global configuration 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 FabricPath unicast load balancing for VDCs that do not allow F2 resource types:

```
switch# configure terminal
switch(config)# fabricpath load-balance unicast layer3
switch(config)#
```

This example shows how to configure FabricPath unicast load balancing for VDCs that allow F2 resource types:



Note The command in this example enables destination MAC-based selection for port-channel hash for ingress modules in the chassis.

```
switch# configure terminal
switch(config)# port-channel load-balance dst mac

switch(config)# show port-channel load-balance
Port Channel Load-Balancing Configuration:
System: dst mac
Port Channel Load-Balancing Addresses Used Per-Protocol:
Non-IP: dst mac
IP: dst mac
```



Note For FabricPath unicast traffic (ECMP selection)—These commands include a mixed preference of Layer 3 and Layer 4 parameters, a rotation of 14 bytes, a VLAN that is included in hash calculations, and a destination-based selection for all modules in the F2 FabricPath-enabled VDC

```
switch(config)# fabricpath load-balance unicast include-vlan
switch(config)# show fabricpath load-balance
ECMP load-balancing configuration:
L3/L4 Preference: Mixed
Rotate amount: 14 bytes
Use VLAN: TRUE
Ftag load-balancing configuration:
Rotate amount: 3 bytes
Use VLAN: TRUE
```

This example shows how to configure F2 VDC FabricPath unicast load balancing:



Note The command in this example enables source IP-VLAN and MAC-based selection for port-channel hash for ingress module 4. All other modules in the chassis retain destination MAC-based selection.

```
switch(config)# port-channel load-balance src ip-vlan module 4
switch(config)# show port-channel load-balance module 4
Port Channel Load-Balancing Configuration:
Module 4: src ip-vlan
```

```
Port Channel Load-Balancing Addresses Used Per-Protocol:
Non-IP: src mac
IP: src ip-vlan
```



Note For FabricPath unicast traffic (ECMP selection)—These commands include a mixed preference of Layer 3 and Layer 4 parameters, a rotation of 9 bytes, a VLAN that is excluded in hash calculation with source based selection for module 4, and a destination based selection for other modules in the F2 FabricPath-enabled VDC.

```
switch(config)# fabricpath load-balance unicast mixed rotate-amount 0x9
switch(config)# show fabricpath load-balance
ECMP load-balancing configuration:
L3/L4 Preference: Mixed
Rotate amount: 9 bytes
Use VLAN: FALSE
Ftag load-balancing configuration:
Rotate amount: 2 bytes
Use VLAN: FALSE
```

This example shows how to configure FabricPath unicast load balancing for VDCs that allow F2 resource types:



Note The command in this example enables source-destination IP-L4PORT-VLAN and MAC-based selection for port-channel hash for ingress module 4. All other modules in the chassis retain the destination MAC-based selection. For FabricPath unicast traffic (ECMP selection), these commands include a mixed preference of Layer 3 and Layer 4 parameters, a rotation of 9 bytes, and a VLAN that is excluded in the hash calculation with a source-based selection for module 4, source-destination based selection for module 10, and destination-based selection for other modules in the F2 FabricPath-enabled VDC.

```
switch(config)# port-channel load-balance src-dst ip-l4port-vlan module 10
switch(config)# show port-channel load-balance module 10
Port Channel Load-Balancing Configuration:
Module 10: src-dst ip-l4port-vlan
Port Channel Load-Balancing Addresses Used Per-Protocol:
Non-IP: src-dst mac
IP: src-dst ip-l4port-vlan
```

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 enabled the FabricPath feature.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	<ul style="list-style-type: none"> switch(config)# [no] fabricpath load-balance {source source-destination xor destination symmetric} switch(config)# [no] port-channel load-balance [algorithm [module module]] 	<p>To configure source/destination/symmetric/src-dst algorithms for load-balancing FabricPath multicast traffic in vDCs that do not allow F2 resource types, use the fabricpath load-balance command.</p> <p>To configure source/destination/symmetric/src-dst algorithms for load-balancing FabricPath multicast traffic in vDCs that allow F2 resource types, use the port-channel load-balance command.</p>
Step 3	switch(config)# [no] fabricpath load-balance multicast [rotate-amount <i>rot_amt</i>] [include-vlan]	<p>Configures options such as rotation-skew, VLAN inclusion, use of Layer 3/Layer 4 traffic parameters, for load-balancing FabricPath multicast traffic.</p> <p>Note To return to the default unicast load-balancing scheme, enter the no form of this command.</p>
Step 4	switch(config)# exit	Exits global configuration 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 FabricPath multicast load balancing:

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

This example shows how to configure FabricPath multicast load balancing for VDCs that allow F2 resource types:



Note The command in this example enables destination MAC-based selection for port-channel hash for ingress modules in the chassis.

```
switch# configure terminal
switch(config)# port-channel load-balance dst mac

switch(config)# show port-channel load-balance
Port Channel Load-Balancing Configuration:
System: dst mac
```

```
Port Channel Load-Balancing Addresses Used Per-Protocol:
Non-IP: dst mac
IP: dst mac
```



Note For FabricPath unicast traffic (forwarding tree selection)—These commands include a rotation of 3 bytes and a VLAN that is included in hash calculations.

```
switch(config)# fabricpath load-balance multicast rotate-amount 0x3 include-vlan
switch(config)# show fabricpath load-balance
ECMP load-balancing configuration:
L3/L4 Preference: Mixed
Rotate amount: 14 bytes
Use VLAN: TRUE
Ftag load-balancing configuration:
Rotate amount: 3 bytes
Use VLAN: TRUE
```

This example shows how to configure FabricPath multicast load balancing for VDCs that allow F2 resource types:



Note The command in this example enables source IP-VLAN and MAC-based selection for port-channel hash as well as FabricPath unicast load balancing for ingress module 4. All other modules in the chassis retain destination MAC-based selection.

```
switch(config)# port-channel load-balance src ip-vlan module 4
switch(config)# show port-channel load-balance module 4
Port Channel Load-Balancing Configuration:
Module 4: src ip-vlan
Port Channel Load-Balancing Addresses Used Per-Protocol:
Non-IP: src mac
IP: src ip-vlan
```



Note For FabricPath multicast traffic (forwarding tree selection)—These commands include a rotation of 2 bytes, a VLAN that is excluded in hash calculation with source-based selection for module 4, and destination-based selection for other modules in F2 FabricPath-enabled VDC.

```
switch(config)# fabricpath load-balance multicast rotate-amount 0x2
switch(config)# show fabricpath load-balance
ECMP load-balancing configuration:
L3/L4 Preference: Mixed
Rotate amount: 9 bytes
Use VLAN: FALSE
Ftag load-balancing configuration:
Rotate amount: 2 bytes
Use VLAN: FALSE
```

This example shows how to configure FabricPath multicast load balancing for VDCs that allow F2 resource types:



Note The command in this example enables source-destination IP-L4PORT-VLAN, MAC-based selection for port-channel hash for ingress module 10, and Source IPVLAN and MAC-based selection for port-channel hash for ingress module 4. All other modules in the chassis retain destination MAC-based selection. For FabricPath multicast traffic (forwarding tree selection), these commands include a rotation of 2 bytes, a VLAN that is excluded in hash calculation with source-based selection for module 4, source-destination based selection for module 10, and destination-based selection for other modules in the F2 FabricPath-enabled VDC.

```
switch(config)# port-channel load-balance src-dst ip-l4port-vlan module 10
switch(config)# show port-channel load-balance module 10
Port Channel Load-Balancing Configuration:
Module 10: src-dst ip-l4port-vlan
Port Channel Load-Balancing Addresses Used Per-Protocol:
Non-IP: src-dst mac
IP: src-dst ip-l4port-vlan
```

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 enabled the FabricPath feature.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# fabricpath multicast aggregate-routes [exclude ftag-id]	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 I3 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 ftag argument.
Step 3	(Optional) switch(config)# show I2 multicast ftag ftag	Displays the configuration that you just applied to the FTag for route programming.
Step 4	switch(config)# exit	Exits global configuration mode.

	Command or Action	Purpose
Step 5	(Optional) switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to configure increased FabricPath multicast scalability:

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

Configuring FabricPath Unicast Static Routes

You can configure unicast static routes to override the routes computed by dynamic protocols such as IS-IS in FabricPath. For example, you might want to route traffic to a particular device using a specific link to ensure better load balancing or to route traffic through a firewall in the network.

Before you begin

Ensure that you are working on an F Series module.

Ensure that you have enabled the FabricPath feature.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	(Optional) switch(config)# fabricpath topology topology-number	Enters FabricPath topology configuration mode (config-fp-topology). Note Enter this command to configure unicast static routes for a specific FabricPath topology (other than the default). If you want to configure unicast static routes for the default topology, skip Step 2 and go to Step 3.
Step 3	[no] fabricpath route switch-id switch-id nh_if_range	Configures a unicast static route and specifies the device and interfaces through which to send the traffic. You can enter a range of Ethernet ports or port channels. The interfaces specified must be in the same VDC where the FabricPath feature set is enabled. This command can be run in two modes:

	Command or Action	Purpose
		<ul style="list-style-type: none"> • Within a specific FabricPath topology configuration mode (config-fp-topology). • Within the global configuration mode. <p>To delete the static route, enter the no form of the command specifying the static route switch ID. To delete the association between the interfaces and the static route, enter the no form of the command specifying the interface ranges.</p> <p>When the last association is deleted, the static route is deleted.</p> <p>Repeat this step to specify additional interfaces for the static route.</p>
Step 4	switch(config)# exit	Exits global configuration mode.
Step 5	(Optional) switch# show fabricpath static route	Displays the static routes within the FabricPath configuration.
Step 6	(Optional) switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to configure a unicast static route for the default topology:

```
switch# configure terminal
switch(config)# fabricpath route switch-id 25 ethernet 1/2
```

This example shows how to configure a unicast static route for a specific topology:

```
switch# configure terminal
switch(config)# fabricpath topology 2
switch(config-fp-topology)# fabricpath route switch-id 221 ethernet 1/2
switch(config-fp-topology)# fabricpath route switch-id 221 port-channel 1
```

This example shows how to delete a unicast static route.

```
switch# configure terminal
switch(config)# no fabricpath route switch-id 221
```

Verifying the FabricPath Configuration

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

Command	Purpose
show feature-set	Displays whether FabricPath is enabled or not.
show {l2 fabricpath} route [switchid <i>switch-id</i>] [detail] [hex]	Displays unicast routes.
show {l2 fabricpath} mroute {[vdc_omf] vlan <i>vlan-id</i> [{omf flood} {source {<i>srcaddr</i> <i>v6srcaddr</i>}}] [group {<i>groupaddr</i> <i>v6grpaddr</i>}}] [resolved] [ftag <i>ftag-id</i>] [hex]	Displays multicast routes.
show fabricpath topology [detail]	Displays information about all FabricPath topologies.
show fabricpath topology interface	Displays information about all FabricPath topology interfaces.
show fabricpath topology vlan [active]	Displays information about all FabricPath topology VLANs.
show fabricpath topology ftag [active] [multicast] [unicast]	Displays information about all FabricPath topology FTags.
show running-config fabricpath	Displays the running configuration for FabricPath.
show fabricpath load-balance unicast forwarding-path ftag <i>ftag-id</i> switchid <i>switch-id</i> flow-type {l2 [{dst-mac <i>dst-mac</i> source-mac <i>src-mac</i>} ether-type <i>ether-type</i>}] {l3 {dst-ip <i>dst-ip</i> src-ip <i>src-ip</i> dst-ipv6 <i>dst-ipv6</i> srcipv6 <i>src-ipv6</i>}} {l4 {l4-src-port <i>l4-src-port</i> l4-dst-port <i>l4-dst-port</i> dst-ip <i>dst-ip</i> src-ip <i>src-ip</i> dst-ipv6 <i>dst-ipv6</i> srcipv6 <i>src-ipv6</i>}} } {vlan <i>vlan-id</i>} {module <i>mod-no</i>}	Displays FabricPath unicast load-balancing information.
show fabricpath load-balance multicast ftag-selected flow-type {l2 [{dst-mac <i>dst-mac</i> source-mac <i>src-mac</i>} ether-type <i>ether-type</i>}] {l3 {dst-ip <i>dst-ip</i> src-ip <i>src-ip</i> dst-ipv6 <i>dst-ipv6</i> srcipv6 <i>src-ipv6</i>}} {l4 {l4-src-port <i>l4-src-port</i> l4-dst-port <i>l4-dst-port</i> dst-ip <i>dst-ip</i> src-ip <i>src-ip</i> dst-ipv6 <i>dst-ipv6</i> srcipv6 <i>src-ipv6</i>}} } {vlan <i>vlan-id</i>} {module <i>mod-no</i>}	Displays FabricPath multicast load-balancing information.
show vlan	Displays information on all FP and CE VLANs.
show fabricpath static route	Displays the static routes within the FabricPath configuration.

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
```

Configuration Example for FabricPath Forwarding

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 “Configuring FabricPath Interfaces,” for information about 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# configure 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.

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

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

Feature History for Configuring FabricPath Forwarding

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

Table 8: Feature History for FabricPath Forwarding

Feature Name	Release	Feature Information
Unicast static routes	6.2(2)	Unicast static routes were introduced.
Load Balancing Using Port Channels	6.0(1)	Load balancing to support F2 modules introduced.
Additional FabricPath topologies	5.2(1)	This feature was introduced.
FabricPath	5.1(1)	These features were introduced.



CHAPTER 7

Advanced FabricPath Features

This chapter describes how to configure advanced FabricPath features, such as using the Intermediate System-to-Intermediate System (IS-IS) protocol on Cisco NX-OS devices.

- [Finding Feature Information, on page 87](#)
- [Information About FabricPath Advanced Features, on page 87](#)
- [Prerequisites for FabricPath, on page 90](#)
- [Guidelines and Limitations for FabricPath Advanced Features, on page 90](#)
- [Setting Advanced FabricPath Layer 2 IS-IS Parameters, on page 91](#)
- [Verifying the FabricPath Advanced Configurations, on page 100](#)
- [Feature History for Configuring FabricPath Advanced Features, on page 101](#)

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 Advanced Features



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

Information About Advanced FabricPath Layer 2 IS-IS Configurations



Note See “Configuring FabricPath Switching,” for information on the default Layer 2 IS-IS behavior with FabricPath.

We recommend that you run the FabricPath network using the default Layer 2 IS-IS configurations.

Optionally, you can also change many of the IS-IS settings. You change these settings as follows:

- Globally on the entire device and on each device in the FabricPath network
- On specified FabricPath interfaces within the FabricPath network

If you do change any of the FabricPath Layer 2 IS-IS settings, ensure that you make the same changes for those global parameters on every device in the FabricPath network and for those interface parameters on every applicable FabricPath interface in the network.

Layer 2 IS-IS is based on Layer 3 IS-IS with enhancements to run on Layer 2. The commands for Layer 2 IS-IS and Layer 3 IS-IS are not the same. Layer 2 IS-IS is the control plane in FabricPath and a single protocol controls all unicast and multicast traffic. From a forwarding standpoint, FabricPath Layer 2 IS-IS forwards traffic for unicast, unknown unicast, broadcast, and multicast frames. Using Layer 2 IS-IS, the system maintains loop-free paths throughout the FabricPath network (see “Configuring FabricPath Switching,” for information on default FabricPath Layer 2 IS-IS behavior and “Configuring FabricPath Forwarding,” for information on FabricPath forwarding.)

You can use these advanced FabricPath Layer 2 IS-IS configurations to fine-tune the operation of the FabricPath network.

Beginning with Cisco Nexus Release 6.2(2), the following features for advanced FabricPath Layer 2 IS-IS are available:

- Overload bit—You can configure the overload bit for FabricPath IS-IS. You achieve consistent routing behavior in conditions where a node reboots or gets overloaded.
- VLAN pruning—The switch will only attract data traffic for the VLANs that have active Classic Ethernet (CE) ports on an F1 Series module, F2 Series module, or switch virtual interfaces (SVIs) for those VLANs.
- Route-map and mesh group—You can use a route-map to control the routes that are redistributed into the FabricPath IS-IS topology. The mesh group reduces flooding for parallel links and mesh topologies. For the parallel links, the blocked mode stops flooding after an initial exchange. For the mesh topologies, the group mode groups the links to stop the link-state packet (LSP) flooding back to the same link in the group where the LSP is received.

**Note**

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

High Availability

The FabricPath topologies retain their configuration through ISSU.

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

Virtual Device Contexts

You must install the FabricPath feature set before you enable FabricPath on the switch. See 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.

Multiple Topologies

In the FabricPath paradigm, a network can be divided into multiple topologies. Within each topology, one or more trees can be computed for forwarding of broadcast and multicast traffic. A tree is a subset of links of an acyclic graph, and a graph is a collection of Layer 2 multipath (L2MP) nodes and links that forms an acyclic topology. The L2MP IS-IS component supports multiple topologies that run in the same process, which reduces CPU usage when compared with using one process per VLAN.

You can have multiple pods (small Layer 2 blocks) in the same Layer 2 domain, but all the pods must have the same set of VLANs configured. Without FabricPath, each pod could have some VLANs used as local VLANs and the traffic on those VLANs are localized to the switches in the pod. To restrict local VLAN traffic to the pod, different FabricPath topologies are configured for the local VLANs. Each pod must be configured with a unique set of local VLANs. The broadcast and multicast traffic on the local VLANs might go through the spine switches and other pods based on the multicast tree.

The L2MP network might have multiple topologies. Each topology has multiple graphs that are associated with them. However, not all graphs can be used until a trigger is received from the Dynamic Resource Allocation Protocol (DRAP). On receipt of the trigger, the graphs are activated. When the topology changes, to maintain loop-free properties of these graphs, triggers are sent to set the hardware states of the ports. The L2MP IS-IS component requests redistribution of the multicast routes from other protocols. All routes that are populated to the multicast Layer 2 routing information base (M2RIB) are redistributed by L2MP IS-IS in its group membership (GM) link state protocols (LSP).

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 Advanced Features

FabricPath 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 Classical Ethernet 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.

Setting Advanced FabricPath Layer 2 IS-IS Parameters



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

Although the Layer 2 IS-IS protocol works automatically once you enable FabricPath, you can optionally configure parameters. Some FabricPath Layer 2 IS-IS parameters you configure globally and some you configure per interface.

Setting Advanced FabricPath Layer 2 IS-IS Parameters Globally (Optional)

Although the FabricPath Layer 2 IS-IS protocol works automatically once you enable FabricPath, you can optionally configure the global parameters.

Before you begin

Ensure that you are working on an F Series module.

Ensure that you have enabled the FabricPath feature.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# fabricpath domain default	Enters global FabricPath Layer 2 IS-IS configuration mode.
Step 3	(Optional) switch(config-fabricpath-isis)# authentication-check	Configures an authentication check on a PDU reception. To turn the authentication check off, enter the no form of this command. The default is ON.
Step 4	(Optional) switch(config-fabricpath-isis)# authentication key-chain <i>auth-key-chain-name</i>	Configures the authentication key chain. To clear this parameter, enter the no form of this command. An example of key chain creation is as follows: key chain trees key 0

	Command or Action	Purpose
		<pre>key-string cisco01 accept-lifetime 07:00:00 Sep 20 2011 infinite send-lifetime 07:00:00 Sep 20 2011 infinite</pre> <p>See the <i>Cisco Nexus 7000 Series NX-OS Security Configuration Guide</i>, for information about key chains.</p>
Step 5	(Optional) switch(config-fabricpath-isis)# authentication type {cleartext md5}	Configures the authentication type. To clear this parameter, enter the no form of this command.
Step 6	(Optional) switch(config-fabricpath-isis)# log-adjacency-changes	Sets the device to send a log message when the state of a FabricPath Layer 2 IS-IS neighbor changes. To stop the log messages, enter the no form of this command. The default is off.
Step 7	(Optional) switch(config-fabricpath-isis)# lsp-gen-interval <i>lsp-max-wait</i> [<i>lsp-initial-wait</i> <i>lsp-second-wait</i>]	<p>Configures the LSP generation interval. To return to the default values, enter the no form of this command. The optional arguments are as follows:</p> <ul style="list-style-type: none"> • <i>lsp-max-wait</i>—The initial wait between the trigger and LSP generation. The range is from 50 to 12000 milliseconds, and the default value is 8000 milliseconds. • <i>lsp-initial-wait</i>—The initial wait between the trigger and LSP generation. The range is from 50 to 12000 milliseconds, and the default value is 50 milliseconds. • <i>lsp-second-wait</i>—The second wait used for LSP throttle during backoff. The range is from 50 to 12000 milliseconds, and the default value is 50 milliseconds.
Step 8	(Optional) switch(config-fabricpath-isis)# lsp-mtu <i>mtu</i>	Sets the LSP MTU. To return to the default values, enter the no form of this command. The range is from 128 to 4352, and the default value is 1492.
Step 9	(Optional) switch(config-fabricpath-isis)# max-lsp-lifetime <i>secs</i>	Sets the maximum LSP lifetime in seconds. To return to the default values, enter the no form of this command. The range is from 128 to 4352, and the default value is 1492.
Step 10	(Optional) switch(config-fabricpath-isis)# maximum-paths <i>max-paths</i>	Sets the maximum number of paths per destination. To return to the default values,

	Command or Action	Purpose
		enter the no form of this command. The range is from 1 to 16, and the default value is 16.
Step 11	(Optional) switch(config-fabricpath-isis)# reference-bandwidth { <i>ref-mbps</i> [Mbps] <i>ref-gbps</i> [Gbps]}	Configures the reference bandwidth, which is used to assign the FabricPath Layer 2 IS-IS cost. The default value is 400000 Mbps. To return to the default values, enter the no form of this command. The optional arguments are as follows: <ul style="list-style-type: none"> • <i>ref-mbps</i>—The range is from 1 to 400000, and the default value is 400000. • <i>ref-gbps</i>—The range is from 1 to 4000, and the default value is 400.
Step 12	(Optional) switch(config-fabricpath-isis)# spf-interval <i>spf-max-wait</i> [<i>spf-initial-wait</i> <i>spf-second-wait</i>]	Configures the interval between LSA arrivals. To return to the default values, enter the no form of this command. The optional keywords are as follows: <ul style="list-style-type: none"> • <i>spf-max-wait</i>—The maximum wait between the trigger and SPF computation. The range is from 50 to 120000 milliseconds, and the default value is 8000 milliseconds. • <i>spf-initial-wait</i>—The initial wait between the trigger and SPF computation. The range is from 50 to 120000 milliseconds, and the default value is 50 milliseconds. • <i>spf-second-wait</i>—The second wait used for SPF computation during backoff. The range is from 50 to 120000 milliseconds, and the default value is 50 milliseconds.
Step 13	(Optional) switch(config-fabricpath-isis)# graceful-restart [t3 manual <i>secs</i>]	Enables graceful restart for the FabricPath Layer 2 IS-IS protocol. To disable graceful restart, enter the no form of this command. Use the t3 manual keyword to set the graceful-restart timer; the range is from 30 to 65535, and the default value is 60. This feature is on by default.
Step 14	(Optional) switch(config-fabricpath-isis)# redistribute filter route-map <i>map-name</i>	Configures the route map to control the routes that are redistributed into the FabricPath IS-IS topology.

Setting Advanced FabricPath Layer 2 IS-IS Parameters Globally (Optional)

	Command or Action	Purpose
		Note See <i>Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide</i> for more information on configuring route maps.
Step 15	(Optional) switch(config-fabricpath-isis)# hostname dynamic	Enables dynamic hostname for the FabricPath Layer 2 IS-IS protocol. To disable the dynamic hostname, enter the no form of this command.
Step 16	(Optional) switch(config-fabricpath-isis)# root-priority <i>value</i>	Configures the priority for which node becomes the Layer 2 IS-IS protocol root in the FabricPath network. The highest numerical value for the priority is likely to become the root. To return to the default values, enter the no form of this command. The range is from 1 to 255, and the default value is 64.
Step 17	(Optional) switch(config-fabricpath-isis)# [no] set-overload-bit { always on-startup <i>seconds</i> }	<p>Configures the overload bit for the system. To disable the overload bit enter the no form of this command. The optional keywords are as follows:</p> <ul style="list-style-type: none"> • always—The overload bit is always on. • on-startup—The overload bit is set upon system startup and remains set for the specified number of seconds.
Step 18	(Optional) switch(config-fabricpath-isis)# [no] vlan-pruning enable	Configures the VLAN pruning for the system. To disable VLAN pruning, enter the no form of this command.
Step 19	switch(config-fabricpath-isis)# exit	Exits global FabricPath Layer 2 IS-IS configuration mode.
Step 20	switch(config)# exit	Exits global configuration mode.
Step 21	(Optional) switch# show running-config	Displays the running configuration.
Step 22	(Optional) switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

What to do next

See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide* for more information about IS-IS commands.

Setting Advanced FabricPath Layer 2 IS-IS Parameters per Interface (Optional)

Although the FabricPath Layer 2 IS-IS protocol works automatically once you enable FabricPath, you can optionally configure the interface parameters.

Before you begin

Ensure that you are working on an F Series module.

Ensure that you have enabled the FabricPath feature.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	interface { <i>ethernet mod/slot</i> port-channel <i>channel-number</i> }	Enters interface configuration mode.
Step 3	(Optional) switch(config-if)# fabricpath isis authentication-check	Enables authentication checking on incoming FabricPath Layer 2 IS-IS hello PDUs. The default is ON. To disable authentication, enter the no form of the command.
Step 4	(Optional) switch(config-if)# fabricpath isis authentication key-chain <i>auth-key-chain-name</i>	<p>Assigns a password to authentication hello PDUs. To remove this password, enter the no form of the command.</p> <p>Note A level specification is not required.</p> <p>An example of key chain creation is as follows:</p> <pre>key chain trees key 0 key-string cisco01 accept-lifetime 07:00:00 Sep 20 2011 infinite send-lifetime 07:00:00 Sep 20 2011 infinite</pre> <p>See the <i>Cisco Nexus 7000 Series NX-OS Security Configuration Guide</i>, for information about key chains.</p>
Step 5	(Optional) switch(config-if)# fabricpath isis authentication-type { <i>cleartext</i> <i>md5</i> }	<p>Specifies the authentication type for an interface for FabricPath Layer 2 IS-IS hello PDUs. To remove this type, enter the no form of the command.</p> <p>Note A level specification is not required.</p>
Step 6	(Optional) switch(config-if)# fabricpath isis csnp-interval <i>seconds</i>	Specifies the interval between CSNP PDUs sent on the interface. To return to the default value, enter the no form of this command. The

Setting Advanced FabricPath Layer 2 IS-IS Parameters per Interface (Optional)

	Command or Action	Purpose
		range is from 1 to 65535, and the default value is 10.
Step 7	(Optional) switch(config-if)# fabricpath isis hello-interval <i>seconds</i>	Sets the hello interval between PDUs sent on the interface. To return to the default value, enter the no form of this command. The range is from 1 to 65535, and the default value is 10. Note A level specification is not required.
Step 8	(Optional) switch(config-if)# fabricpath isis hello-multiplier <i>multiplier</i>	Specifies the multiplier used to calculate the interval within which hello PDUs must be received or adjacency goes down. To return to the default value, enter the no form of this command. The range is from 3 to 1000. The default is 3. Note A level specification is not required.
Step 9	(Optional) switch(config-if)# fabricpath isis hello-padding	Enables padding on the hello PDUs. The default is on. To disable authentication, enter the no form of the command. If you enter the always keyword with the no form of this command, the padding is always on.
Step 10	(Optional) switch(config-if)# fabricpath isis lsp-interval <i>milliseconds</i>	Sets the interval in milliseconds between LSPs sent on this interface during flooding. To return to the default value, enter the no form of this command. The range is from 10 to 65535. The default is 33.
Step 11	(Optional) switch(config-if)# fabricpath isis mesh-group <i>group-number</i>	Specifies the mesh-group state and sets the mesh-group attribute on the interface.
Step 12	(Optional) switch(config-if)# fabricpath isis metric <i>metric</i>	Configures the FabricPath Layer 2 IS-IS metric for this interface. The range is from 0 to 16777215. To return to the default value, enter the no form of this command. The default values are as follows (the default interface for the F Series module is 10 GB): <ul style="list-style-type: none"> • 1 GB—400 • 10 GB—40
Step 13	(Optional) switch(config-if)# fabricpath isis retransmit-interval <i>seconds</i>	Sets the interval between initial LSP retransmissions. To return to the default value, enter the no form of this command. The range is from 1 to 65535. The default is 5.

	Command or Action	Purpose
Step 14	(Optional) switch(config-if)# fabricpath isis retransmit-throttle-interval <i>milliseconds</i>	Sets the interval between subsequent LSP retransmissions. To return to the default value, enter the no form of this command. The range is from 20 to 65535. The default is 66.
Step 15	switch(config-if)# exit	Exits interface configuration mode.
Step 16	switch(config)# exit	Exits global configuration mode.
Step 17	(Optional) switch# show running-config	Displays the running configuration.
Step 18	(Optional) switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

What to do next

See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide* for more information about IS-IS commands.

Clearing Advanced FabricPath Layer 2 IS-IS Counters

You can clear the FabricPath Layer 2 IS-IS counters.

Before you begin

Ensure that you are working on an F Series module.

Ensure that you have enabled the FabricPath feature.

Procedure

	Command or Action	Purpose
Step 1	(Optional) switch# clear fabricpath isis adjacency [* system-id interface { ethernet <i>mod/slot</i> port-channel <i>channel-number</i> }]	Clears the FabricPath Layer 2 IS-IS adjacency state. Note If you enter the * variable, you affect forwarding which might interrupt traffic; this command tears down all adjacencies.
Step 2	(Optional) switch# clear fabricpath isis statistics *	Clears all FabricPath Layer 2 IS-IS protocol statistics.
Step 3	(Optional) switch# clear fabricpath isis traffic [* interface { ethernet <i>mod/slot</i> port-channel <i>channel-number</i> }]	Clears FabricPath Layer 2 IS-IS traffic information.

Configuring Multiple Topologies

You can create a topology, map VLANs to the topology, and add an interface to the topology.

Before you begin

Ensure that you are working on an F Series module.

Ensure that you have enabled the FabricPath feature.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# fabricpath topology id	Creates a new fabricpath topology and enters the FabricPath topology configuration mode.
Step 3	switch(config-fp-topology)# member vlan range	Configures the VLANs for the topology. The range of the VLAN ID is from 1 to 4094.
Step 4	switch(config-fp-topology)# exit	Exits FabricPath topology configuration mode.
Step 5	switch(config)# interface port-channel number	Configures a port-channel interface and enters interface configuration mode. You can configure any of the available interfaces.
Step 6	switch(config-if)# fabricpath topology-member id	Adds the interface to the topology.
Step 7	switch(config-if)# exit	Exits interface configuration mode.
Step 8	switch(config)# exit	Exits global configuration mode.
Step 9	(Optional) switch# show fabricpath topology vlan	Displays information about the VLANs in the Layer 2 topology.
Step 10	(Optional) switch# show fabricpath isis topology summary	Displays information about the IS-IS summary topology.
Step 11	(Optional) switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to create a topology, map VLANs to the topology, and add an interface to the topology.

```
switch# configure terminal
switch(config)# fabricpath topology 1
switch(config-fp-topology)# member vlan 7-19
switch(config-fp-topology)# exit
switch(config)# interface port-channel 1
```

```

switch(config-if)# fabricpath topology-member 1
switch(config-if)# exit
switch(config)# show fabricpath topology vlan
switch(config)# show fabricpath isis topology summary

```

Configuring FabricPath IS-IS Multiple Topologies

You can configure FabricPath IS-IS multiple topologies.

Before you begin

Ensure that you are working on an F Series module.

Ensure that you have enabled the FabricPath feature.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# fabricpath domain default	Enters global FabricPath Layer 2 IS-IS configuration mode.
Step 3	switch(config-fabricpath-isis)# topology id	Enters the Layer 2 topology for IS-IS configuration mode.
Step 4	(Optional) switch(config-fabricpath-isis-topo)# maximum-paths max-paths	Configures the maximum paths per destination on the switch for the Layer 2 topology.
Step 5	(Optional) switch(config-fabricpath-isis-topo)# reference-bandwidth {ref-mbps mbps ref-gbps gbps}	Configures the reference bandwidth for setting the interface metrics on the switch for the Layer 2 topology.
Step 6	(Optional) switch(config-fabricpath-isis-topo)# root-priority priority	Configures the priority with which nodes become root on the switch for the Layer 2 topology.
Step 7	switch(config-fabricpath-isis-topo)# exit	Exits FabricPath IS-IS topology configuration mode.
Step 8	(Optional) switch# show fabricpath topology vlan	Displays information about the VLANs in the Layer 2 topology.
Step 9	(Optional) switch# show fabricpath isis topology summary	Displays information about the IS-IS summary topology.
Step 10	(Optional) switch# copy running-config startup-config	Copies the running configuration to the startup configuration.

Example

This example shows how to configure FabricPath IS-IS multiple topologies:

```

switch# configure terminal
switch(config)# fabricpath domain default
switch(config-fabricpath-isis)# topology 5
switch(config-fabricpath-isis-topo)# maximum-paths 5
switch(config-fabricpath-isis-topo)# reference-bandwidth ref-mbps 100
switch(config-fabricpath-isis-topo)# root-priority 1
switch(config-fabricpath-isis-topo)# exit
switch(config-fabricpath-isis)# show fabricpath topology vlan
switch(config-fabricpath-isis)# show fabricpath isis topology summary

```

Verifying the FabricPath Advanced Configurations

To display FabricPath information for advanced configurations perform one of the following tasks:

Command	Purpose
show fabricpath isis adjacency [interface { ethernet <i>mod/slot</i> port-channel <i>channel-number</i> } system-id detail summary]	Displays the FabricPath Layer 2 IS-IS adjacency database.
show fabricpath isis database [<i>level</i>] [mgroup] [detail summary] [<i>lid</i>] { zero-seq router-id adjacency } [<i>SID.XX-XX</i>]	Displays the FabricPath Layer 2 IS-IS database.
show fabricpath isis hostname [detail]	Displays the FabricPath Layer 2 IS-IS dynamic hostname exchange information.
show fabricpath isis interface [ethernet <i>mod/slot</i> port-channel <i>channel-number</i>] [brief]	Displays the FabricPath Layer 2 IS-IS related interface information.
show fabricpath isis route [summary detail]	Displays the FabricPath Layer 2 IS-IS routing table for unicast routes.
show fabricpath isis spf-log [detail]	Displays the FabricPath Layer 2 IS-IS SPF calculation statistics.
show fabricpath isis [statistics]	Displays the FabricPath Layer 2 IS-IS event counters.
show fabricpath isis ftag [multidestination <i>tree_id</i>]	Displays the FTag values associated with the trees in the topology.
show fabricpath isis vlan-range	Displays the congruent VLAN-set to topology mapping.
show fabricpath isis trees [multidestination <i>tree_id</i>]	Displays the nodes in the trees.
show fabricpath isis switch-id	Displays the switch IDs and reachability information for the topology.
show fabricpath isis ip redistribute mroute [vlan [group [source]]]	Displays the locally learned multicast routes.

Command	Purpose
show fabricpath isis ip mroute [vlan <i>vlan-id</i> [group <i>group-id</i> [source <i>source-id</i>]]]	Displays the multicast routes learned from neighbors.
show fabricpath isis [protocol]	Displays the FabricPath Layer 2 IS-IS process level information.
show fabricpath isis rrm [gm] interface {ethernet <i>mod/slot</i> port-channel <i>channel-number</i> }	Displays the FabricPath Layer 2 IS-IS retransmit-routing-message information.
show fabricpath isis srm [gm] interface {ethernet <i>mod/slot</i> port-channel <i>channel-number</i> }	Displays the FabricPath Layer 2 IS-IS send-routing-message information.
show fabricpath isis topology summary	Displays the FabricPath Layer 2 IS-IS topology database.
show fabricpath isis traffic [interface {ethernet <i>mod/slot</i> port-channel <i>channel-number</i> }]	Displays the FabricPath Layer 2 IS-IS traffic information.
show fabricpath isis ssn [gm] interface {ethernet <i>mod/slot</i> port-channel <i>channel-number</i> }	Displays the FabricPath Layer 2 IS-IS send-sequence-number information.
show fabricpath isis mesh-group	Displays the FabricPath IS-IS mesh-group information.

Feature History for Configuring FabricPath Advanced Features

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

Table 9: Feature History for Advanced FabricPath Features

Feature Name	Release	Feature Information
Multiple topologies	6.2(2)	This feature was introduced.
Advanced FabricPath Layer 2 IS-IS Parameters per Interface	6.2(2)	Route-map and mesh group were introduced.
Advanced FabricPath Layer 2 IS-IS Parameters Globally	6.2(2)	Overload bit and VLAN pruning for FabricPath IS-IS were introduced.
Advanced FabricPath features	5.1(1)	These features were introduced.



CHAPTER 8

Configuring FabricPath in Transit Mode

This chapter contains the following sections:

- [Information about FabricPath in Transit Mode, on page 103](#)
- [Guidelines and Limitations, on page 103](#)
- [Configuring FabricPath in Transit Mode, on page 104](#)
- [Feature History for Configuring FabricPath in Transit Mode, on page 104](#)

Information about FabricPath in Transit Mode

In transit mode, all the incoming traffic is mapped to one internal bridge domain (VLAN). A pure Layer 2 spine node does not learn the MAC addresses; however, the packet is VDC based on the VDC ID lookup table in the FabricPath header.

In a FabricPath network, the Layer 2 spine node can be configured in Transit mode. After the spine is configured in the Transit mode, uni-destination traffic received on the spine is forwarded based on the FTAG (forwarding tag) and destination VDC ID. For multi-destination traffic, the traffic on the spine is forwarded based on FTAG, which is the identifier of the multi-destination traffic that is selected by the ingress switch.

Guidelines and Limitations

The following are the guidelines and limitations for configuring FabricPath in Transit mode:

- Ensure that a minimum of one VLAN is enabled in FabricPath mode. The default VLAN “VLAN1” can be enabled as FabricPath mode. However, we recommend you set another VLAN as FabricPath mode.
- When you enable FabricPath transit mode, local IPs will be unreachable and the local IP destination will be dropped on ingress F3 card, this affects VPC keepalive. Make sure that the fabricpath transit mode must not be used on FP edge VDCs.
- Switches for which Fabricpath in transit mode is configured are doing only Layer 2 switching.
- Fabricpath mode transit will be activated only after reloading of VDC.

Configuring FabricPath in Transit Mode

To configure the FabricPath in transit mode, perform the following.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# install feature-set fabricpath	Installs FabricPath feature set on the switch.
Step 3	switch(config)# feature-set fabricpath	Enables the FabricPath feature set on the switch.
Step 4	switch(config)# fabricpath mode transit	Configures the FabricPath in transit mode.
Step 5	switch(config)# copy running-config startup-config	Copy and save the running configuration to the startup configuration.
Step 6	switch(config)# reload	Reload the spine for transit mode to take effect.
Step 7	(Optional) switch(config)# show fabricpath mode	Displays configuration details for FabricPath in Transit Mode.

Example

This example shows a running configuration, followed by a verification command that displays the configuration details for Cisco FabricPath in Transit Mode.

```
switch(config)# fabricpath mode transit
switch(config)# show fabricpath mode
switch(config)# copy running-config startup-config
switch(config)# reload
switch(config)# show fabricpath mode
```

Feature History for Configuring FabricPath in Transit Mode

The following table lists the release history for these features.

Table 10: Feature History for FabricPath in Transit Mode

Feature Name	Release	Feature Information
FabricPath in Transit Mode	8.1(1)	This feature was introduced.



APPENDIX **A**

Configuration Limits for Cisco NX-OS FabricPath

- [Configuration Limits for Cisco NX-OS FabricPath, on page 105](#)

Configuration Limits for Cisco NX-OS FabricPath

The configuration limits are documented in the [Cisco Nexus 7000 Series NX-OS Verified Scalability Guide](#).

