



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



## **Cisco Nexus 7000 Series NX-OS FabricPath Configuration Guide**

December 7, 2011

### **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

Text Part Number: OL-22842-01

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 NONINFRINGEMENT 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.

CCDE, CCENT, CCSI, Cisco Eos, Cisco Explorer, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Nurse Connect, Cisco Pulse, Cisco SensorBase, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco TrustSec, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flipshare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), Cisco:Financed (Stylized), Cisco Store, Flip Gift Card, and One Million Acts of Green are service marks; and Access Registrar, Aironet, AllTouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Lumin, Cisco Nexus, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, EtherFast, EtherSwitch, Event Center, Explorer, Follow Me Browsing, GainMaker, iLNNX, IOS, iPhone, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, PCNow, PIX, PowerKEY, PowerPanels, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROSA, SenderBase, SMARTnet, Spectrum Expert, StackWise, WebEx, and the WebEx logo are registered trademarks of Cisco and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website 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. (1002R)

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



## CONTENTS

### New and Changed Information i-vii

### Preface ix

Audience ix

Organization ix

Document Conventions x

Related Documentation xi

Obtaining Documentation and Submitting a Service Request xii

---

## CHAPTER 1

### Overview 1-1

Information About FabricPath 1-1

Information About Conversational MAC Address Learning 1-2

Virtualization for FabricPath 1-2

High Availability for FabricPath 1-2

Licensing Requirements for FabricPath 1-2

---

## CHAPTER 2

### Configuring FabricPath Switching 2-1

Information About FabricPath Switching 2-1

FabricPath Encapsulation 2-2

FabricPath Headers 2-2

Forwarding Tags (FTags) 2-3

Default IS-IS Behavior with FabricPath 2-4

Conversational MAC Address Learning 2-5

Switching Using FabricPath 2-6

Conflict Resolution and Optional FabricPath Tunings 2-7

Interoperation between the M Series and the F Series Modules 2-8

High Availability 2-8

Virtual Device Contexts 2-8

Licensing Requirements for FabricPath Switching 2-9

Prerequisites for FabricPath Switching 2-9

Guidelines and Limitations for FabricPath Switching 2-10

Configuring FabricPath Switching 2-10

Enabling the FabricPath Feature Set on the VDC on the device 2-11

Disabling the FabricPath Feature Set on the VDC 2-12

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

Configuring the MAC Learning Mode for CE VLANs (Optional)	2-13
Configuring the Switch ID (Optional)	2-15
Configuring the FabricPath Timers (Optional)	2-16
Disabling FabricPath Graceful Merges (Optional)	2-17
Forcing the Links to Come Up (Optional)	2-18
Verifying FabricPath Switching	2-19
Displaying and Clearing FabricPath Switching Statistics	2-20
FabricPath Switching Example Configuration	2-20
Default Settings	2-21
Feature History for Configuring FabricPath Switching Using the CLI	2-21

## CHAPTER 3

### **FabricPath Interfaces 3-1**

Information About FabricPath Interfaces	3-1
FabricPath Interfaces	3-2
STP and the FabricPath Network	3-2
vPC+	3-4
High Availability	3-6
Virtual Device Contexts	3-6
Licensing Requirements for FabricPath Interfaces	3-7
Prerequisites for FabricPath Interfaces	3-7
Guidelines and Limitations for FabricPath Interfaces	3-8
Configuring FabricPath Interfaces	3-8
Configuring FabricPath Interfaces	3-9
Configuring the STP Priority with Rapid PVST+	3-10
Configuring the STP Priority with MST	3-11
(Optional) Configuring the STP Domain ID for STP Domains Connected to the Layer 2 Gateway Switch	3-13
Configuring a vPC+ Switch ID	3-14
Verifying FabricPath Interface Configuration	3-15
Displaying FabricPath Interface Statistics	3-16
FabricPath Interface Example Configuration	3-16
Feature History for Configuring FabricPath Using the CLI	3-18

## CHAPTER 4

### **FabricPath Forwarding 4-1**

Information About FabricPath Forwarding	4-1
FabricPath Forwarding Overview	4-2
FabricPath VLANs	4-2
Forwarding Known Unicast Packets Using ECMP	4-3

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

Forwarding Trees for Broadcast, Unknown Unicast, and Multicast Packets	4-4
Forwarding Multicast Packets	4-4
High Availability	4-5
Virtual Device Context	4-5
Licensing Requirements for FabricPath Forwarding	4-6
Prerequisites for FabricPath Forwarding	4-6
Guidelines and Limitations for FabricPath Forwarding	4-6
Configuring FabricPath Forwarding	4-7
Setting the VLAN Mode to FP or CE	4-7
Configuring FabricPath Unicast Load Balancing (Optional)	4-9
Configuring FabricPath Multicast Load Balancing (Optional)	4-10
Configuring FabricPath Increased Multicast Scalability (Optional)	4-11
Verifying the FabricPath Configuration	4-12
FabricPath Forwarding Example Configuration	4-13
Default Settings	4-14
Feature History for Configuring FabricPath Using the CLI	4-14

## CHAPTER 5

### Advanced FabricPath Features 5-1

Information About FabricPath Advanced Features	5-1
Information About Advanced FabricPath Layer 2 IS-IS Configurations	5-1
High Availability	5-2
Virtual Device Contexts	5-2
Licensing Requirements for FabricPath Advanced Features	5-3
Prerequisites for FabricPath Advanced Features	5-3
Guidelines and Limitations for FabricPath Advanced Features	5-3
Setting Advanced FabricPath Layer 2 IS-IS Parameters	5-4
Setting Advanced FabricPath Layer 2 IS-IS Parameters Globally (Optional)	5-4
Setting Advanced FabricPath Layer 2 IS-IS Parameters per Interface (Optional)	5-8
Clearing Advanced FabricPath Layer 2 IS-IS Counters	5-11
Verifying the FabricPath Advanced Configurations	5-12
Feature History for Configuring FabricPath Advanced Features Using the CLI	5-13

## APPENDIX A

### Configuration Limits for Cisco NX-OS FabricPath A-1

## INDEX

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



## New and Changed Information

This chapter provides release-specific information for each new and changed feature in the *Cisco Nexus 7000 Series NX-OS FabricPath Configuration Guide*. The latest version of this document is available at the following Cisco website:

[http://www.cisco.com/en/US/products/ps9402/products\\_installation\\_and\\_configuration\\_guides\\_list.html](http://www.cisco.com/en/US/products/ps9402/products_installation_and_configuration_guides_list.html)

To check for additional information about Cisco NX-OS Release 5.2, see the *Release Notes* available at the following Cisco website:

[http://www.cisco.com/en/US/products/ps9402/prod\\_release\\_notes\\_list.html](http://www.cisco.com/en/US/products/ps9402/prod_release_notes_list.html)

**Table 1** summarizes the new and changed features for the *Cisco Nexus 7000 Series NX-OS FabricPath Configuration Guide* and tells you where they are documented.

**Table 1**      ***New and Changed Features for Release 5.2***

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)

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





## Preface

---

This preface describes the audience, organization, and conventions of the *Cisco Nexus 7000 Series NX-OS FabricPath Configuration Guide*. It also provides information on how to obtain related documentation.

This chapter includes the following sections:

- [Audience, page ix](#)
- [Organization, page ix](#)
- [Document Conventions, page x](#)
- [Related Documentation, page xi](#)
- [Obtaining Documentation and Submitting a Service Request, page xii](#)

## Audience

This publication is for experienced users who configure and maintain Cisco NX-OS devices.

## Organization

This reference is organized as follows:

Chapter and Title	Description
<a href="#">Chapter 1, “Overview”</a>	Describes an overview of the FabricPath feature.
<a href="#">Chapter 2, “Configuring FabricPath Switching”</a>	Describes FabricPath encapsulation, Layer 2 IS-IS, and MAC address conversational learning. Also describes how to change timers for FabricPath and resolve switch ID conflicts.
<a href="#">Chapter 3, “FabricPath Interfaces”</a>	Describes FabricPath interfaces and VLANs as well as how the feature interacts with Classical Ethernet networks and STP.
<a href="#">Chapter 4, “FabricPath Forwarding”</a>	Describes how the FabricPath network uses multiple paths to forward unicast, unknown unicast, broadcast, and multicast frames.

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

Chapter and Title	Description
<a href="#">Chapter 5, “Advanced FabricPath Features”</a>	Describes how to configure more advanced parameters for Layer 2 IS-IS.
<a href="#">Appendix A, “Configuration Limits for Cisco NX-OS FabricPath”</a>	Describes the configuration limits for FabricPath.

## Document Conventions

Command descriptions use these conventions:

Convention	Description
<b>boldface font</b>	Commands and keywords are in boldface.
<i>italic font</i>	Arguments for which you supply values are in italics.
{ }	Elements in braces are required choices.
[ ]	Elements in square brackets are optional.
[ x   y   z ]	Optional alternative keywords are grouped in brackets and separated by vertical bars.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Screen examples use these conventions:

<code>screen font</code>	Terminal sessions and information that the switch displays are in screen font.
<b><code>boldface screen font</code></b>	Information you must enter is in boldface screen font.
<i><code>italic screen font</code></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.



### Tip

Means *the following information will help you solve a problem*.

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

## Related Documentation

Cisco NX-OS includes the following documents:

### Release Notes

*Cisco Nexus 7000 Series NX-OS Release Notes, Release 5.x*

### NX-OS Configuration Guides

*Cisco Nexus 7000 Series NX-OS Configuration Examples, Release 5.x*

*Configuring the Cisco Nexus 2000 Series Fabric Extender*

*Cisco Nexus 7000 Series NX-OS FabricPath Configuration Guide*

*Configuring Feature Set for FabricPath*

*Cisco NX-OS FCoE Configuration Guide for Cisco Nexus 7000 and Cisco MDS 9500*

*Cisco Nexus 7000 Series NX-OS Fundamentals Configuration Guide, Release 5.x*

*Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide, Release 5.x*

*Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 5.x*

*Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide, Release 5.x*

*Cisco Nexus 7000 Series NX-OS LISP Configuration Guide*

*Cisco Nexus 7000 Series NX-OS MPLS Configuration Guide*

*Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide, Release 5.x*

*Cisco Nexus 7000 Series NX-OS OTV Configuration Guide*

*Cisco Nexus 7000 Series OTV Quick Start Guide*

*Cisco Nexus 7000 Series NX-OS Quality of Service Configuration Guide, Release 5.x*

*Cisco Nexus 7000 Series NX-OS SAN Switching Configuration Guide*

*Cisco Nexus 7000 Series NX-OS Security Configuration Guide, Release 5.x*

*Cisco Nexus 7000 Series NX-OS System Management Configuration Guide, Release 5.x*

*Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 5.x*

*Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 4.2*

*Cisco Nexus 7000 Series NX-OS Virtual Device Context Quick Start, Release 5.x*

### NX-OS Command References

*Cisco Nexus 7000 Series NX-OS Command Reference Master Index*

*Cisco Nexus 7000 Series NX-OS FabricPath Command Reference*

*Cisco NX-OS FCoE Command Reference for Cisco Nexus 7000 and Cisco MDS 9500*

*Cisco Nexus 7000 Series NX-OS Fundamentals Command Reference*

*Cisco Nexus 7000 Series NX-OS High Availability Command Reference*

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

*Cisco Nexus 7000 Series NX-OS Interfaces Command Reference*  
*Cisco Nexus 7000 Series NX-OS Layer 2 Switching Command Reference*  
*Cisco Nexus 7000 Series NX-OS LISP Command Reference*  
*Cisco Nexus 7000 Series NX-OS MPLS Command Reference*  
*Cisco Nexus 7000 Series NX-OS Multicast Routing Command Reference*  
*Cisco Nexus 7000 Series NX-OS OTV Command Reference*  
*Cisco Nexus 7000 Series NX-OS Quality of Service Command Reference*  
*Cisco Nexus 7000 Series NX-OS SAN Switching Command Reference*  
*Cisco Nexus 7000 Series NX-OS Security Command Reference*  
*Cisco Nexus 7000 Series NX-OS System Management Command Reference*  
*Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference*  
*Cisco Nexus 7000 Series NX-OS Virtual Device Context Command Reference*

#### **Other Software Documents**

*Cisco NX-OS Licensing Guide*  
*Cisco Nexus 7000 Series NX-OS MIB Quick Reference*  
*Cisco Nexus 7000 Series NX-OS Software Upgrade and Downgrade Guide, Release 5.x*  
*Cisco NX-OS System Messages Reference*  
*Cisco Nexus 7000 Series NX-OS Troubleshooting Guide*  
*Cisco NX-OS XML Interface User Guide*

## **Obtaining Documentation and Submitting a Service Request**

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

<http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html>

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.



# CHAPTER 1

## 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.

This chapter includes the following sections:

- [Information About FabricPath, page 1-1](#)
- [Information About Conversational MAC Address Learning, page 1-2](#)
- [Virtualization for FabricPath, page 1-2](#)
- [High Availability for FabricPath, page 1-2](#)

## Information About FabricPath



### Note

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

Beginning with the Cisco NX-OS Release 5.1 and when you use an F Series module, you can use the FabricPath feature. This new 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 ingress 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 egress interface decapsulates the frame and leaves the regular CE header.



### Note

Classical Ethernet is referred to as CE in this document.

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

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 build 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.

## Information About Conversational MAC Address Learning

Beginning the 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.

## 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 5.x*.

## High Availability for FabricPath

FabricPath retains the configurations across ISSU.

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

## Licensing Requirements for FabricPath

FabricPath requires the Enhanced Layer 2 license. You must install this license on every system that enables FabricPath networks.

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

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





## CHAPTER 2

# 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.

This chapter includes the following sections:

- [Information About FabricPath Switching, page 2-1](#)
- [Licensing Requirements for FabricPath Switching, page 2-9](#)
- [Prerequisites for FabricPath Switching, page 2-9](#)
- [Guidelines and Limitations for FabricPath Switching, page 2-10](#)
- [Configuring FabricPath Switching, page 2-10](#)
- [Verifying FabricPath Switching, page 2-19](#)
- [Displaying and Clearing FabricPath Switching Statistics, page 2-20](#)
- [FabricPath Switching Example Configuration, page 2-20](#)
- [Feature History for Configuring FabricPath Switching Using the CLI, page 2-21](#)

## 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.

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

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.

This section includes the following topics:

- [FabricPath Encapsulation, page 2-2](#)
- [Default IS-IS Behavior with FabricPath, page 2-4](#)
- [Conversational MAC Address Learning, page 2-5](#)
- [Switching Using FabricPath, page 2-6](#)
- [Conflict Resolution and Optional FabricPath Tunings, page 2-7](#)
- [Interoperation between the M Series and the F Series Modules, page 2-8](#)
- [High Availability, page 2-8](#)
- [Virtual Device Contexts, page 2-8](#)

## FabricPath Encapsulation

This section includes the following topics:

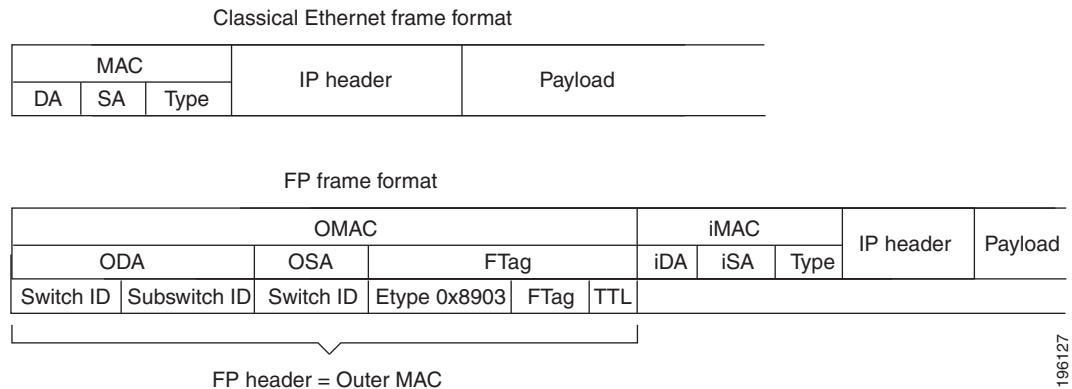
- [FabricPath Headers, page 2-2](#)
- [Forwarding Tags \(FTags\), page 2-3](#)

## 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. [Figure 2-1](#) shows the FabricPath header encapsulating the classical Ethernet (CE) frame.

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

**Figure 2-1 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, “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 TTL to 16; 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 1024 (see [Chapter 4, “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.

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

## 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 2 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 [Chapter 4, “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 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 [Chapter 5, “Advanced FabricPath Features,”](#) for information on configuring optional IS-IS parameters).

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

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 [Chapter 4, “FabricPath Forwarding”](#) for more information on 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.

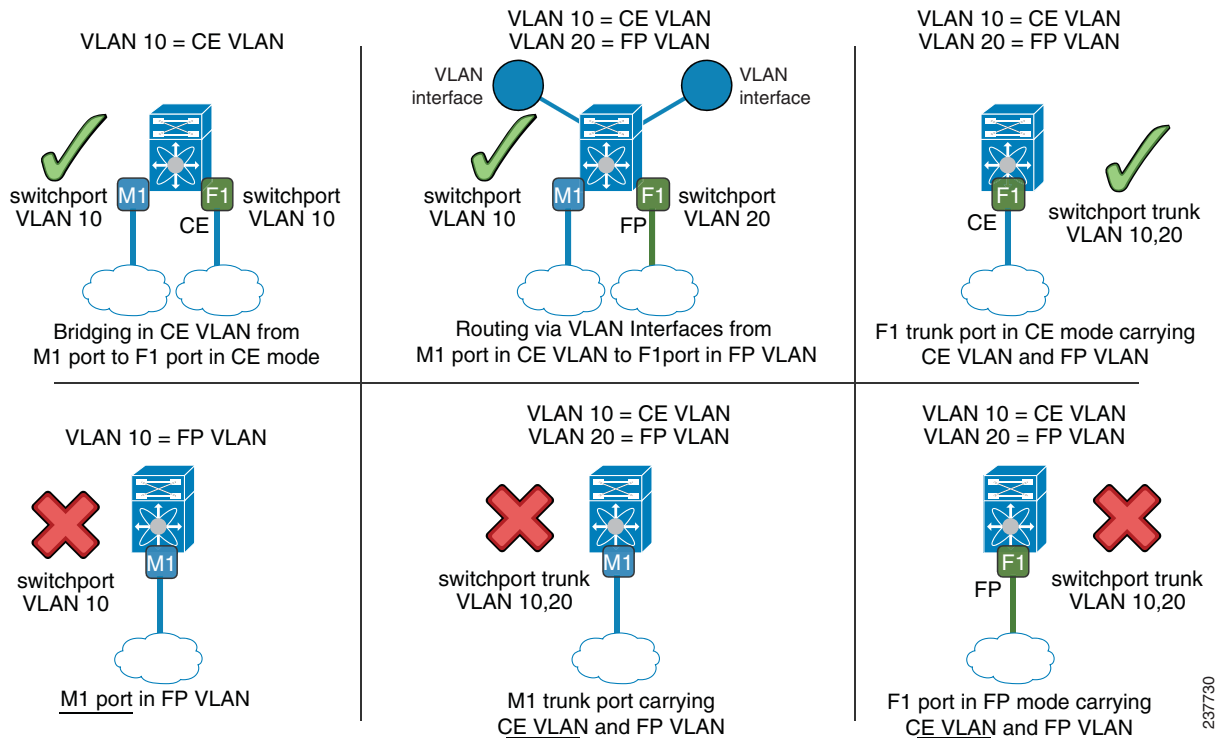
This combination of conversational MAC address learning and multiple FEs on each F Series module produces greatly reduced MAC address tables on each F Series module.

For CE VLANs, you can configure conversational learning per VLAN on the F Series module 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 higher.

[Figure 2-2](#) shows the allowed FabricPath and CE ports on the M and F Series modules and the allowed FP and CE VLANs.

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

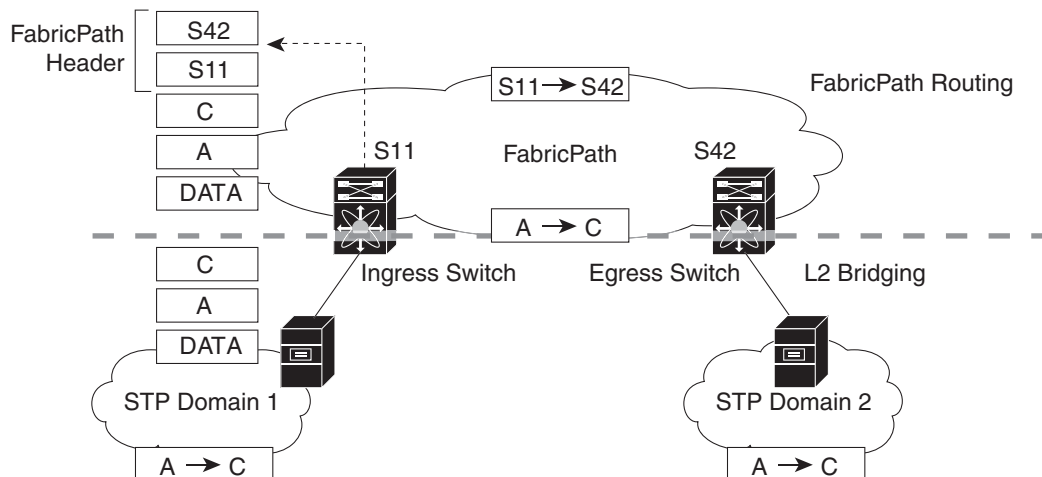
**Figure 2-2** *FP and CE VLAN Examples*



## 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 Figure 2-3).

**Figure 2-3** *FabricPath Ports Use Only the FabricPath Header to Switch Frames*



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

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.

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

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.

## 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 resources 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.



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

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.

## Interoperation between the M Series and the F Series Modules

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, Release 5.x*, and *Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide, Release 5.x*, for more information on interoperation between the F Series and M Series modules.

## High Availability

The FabricPath topologies retain their configuration through ISSU.

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

## Virtual Device Contexts

All the interfaces for one FabricPath network on the same device must be in the same virtual device context (VDC).



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

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, Release 5.x*, for more information on VDCs.

## Licensing Requirements for FabricPath Switching

The following table shows the licensing requirements for this feature:

## Prerequisites for FabricPath Switching

FabricPath switching has the following prerequisites:

- You should have a working knowledge of Classical Ethernet Layer 2 functioning.
- You must install the FabricPath feature set on the default and nondefault VDC before you enable FabricPath on the switch. See *Configuring Feature Set for FabricPath* for information on installing the FabricPath feature set.
- You are logged onto the device.
- Ensure that you have installed the Enhanced Layer 2 license.
- You are in the correct virtual device context (VDC). A VDC is a logical representation of a set of system resources. You can use the **switchto vdc** command with a VDC number.
- You are working on the F Series module.

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

## 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 may cause the standby supervisor to reload if it is in an unstable state, such as following a service failure or powering up.
- STP does not run inside an 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.

## 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 for 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

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

This section includes the following topics:

- [Enabling the FabricPath Feature Set on the VDC on the device, page 2-11](#)
- [Disabling the FabricPath Feature Set on the VDC, page 2-12](#)
- [Configuring the MAC Learning Mode for CE VLANs \(Optional\), page 2-13](#)
- [Configuring the Switch ID \(Optional\), page 2-15](#)
- [Configuring the FabricPath Timers \(Optional\), page 2-16](#)
- [Disabling FabricPath Graceful Merges \(Optional\), page 2-17](#)
- [Forcing the Links to Come Up \(Optional\), page 2-18](#)

## 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 on installing and enabling the FabricPath feature set.

### BEFORE YOU BEGIN

Ensure that you have installed the Enhanced Layer 2 license.

Ensure that you have installed an F Series module.

### SUMMARY STEPS

1. **config terminal**
2. **feature-set fabricpath**
3. **exit**
4. (Optional) **show feature-set**
5. (Optional) **copy running-config startup-config**

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

## DETAILED STEPS

	Command	Purpose
Step 1	<b>config terminal</b>  <b>Example:</b> switch# config terminal switch(config)#	Enters global configuration mode.
Step 2	<b>feature-set fabricpath</b>  <b>Example:</b> switch(config)# feature-set fabricpath	Enables the FabricPath feature set in the VDC.  <b>Note</b> 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	<b>exit</b>  <b>Example:</b> switch(config)# exit switch#	Exits configuration mode.
Step 4	<b>show feature-set</b>  <b>Example:</b> switch# show feature-set	(Optional) Displays which feature sets are enabled on the device.
Step 5	<b>copy running-config startup-config</b>  <b>Example:</b> switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

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

```
switch# config 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. And, the system removes all the FabricPath configurations when you disable the feature set.



### Note

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

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

## BEFORE YOU BEGIN

Ensure that you have installed the Enhanced Layer 2 license.

Ensure that you have installed an F Series module.

## SUMMARY STEPS

1. **config terminal**
2. **no feature-set fabricpath**
3. **exit**
4. (Optional) **show feature-set**
5. (Optional) **copy running-config startup-config**

## DETAILED STEPS

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

This example shows how to disable the FabricPath feature:

```
switch# config 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.

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

**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 the Enhanced Layer 2 license.

Ensure that you have installed an F Series module.

Ensure that you are working with CE VLANs.

**SUMMARY STEPS**

1. **config terminal**
2. **mac address-table learning-mode conversational vlan *vlan-id***
3. **exit**
4. (Optional) **show mac address-table learning-mode {vlan *vlan-id*}**
5. (Optional) **copy running-config startup-config**

**DETAILED STEPS**

	Command	Purpose
Step 1	<b>config terminal</b>  <b>Example:</b> switch# config terminal switch(config)#	Enters global configuration mode.
Step 2	<b>mac address-table learning-mode conversational vlan <i>vlan-id</i></b>  <b>Example:</b> switch(config)# mac address-table learning-mode conversational vlan 1-10 switch(config)#	Configures the specified CE VLAN(s) on F Series modules for conversational MAC learning. Enter the <b>no</b> form of the command to return to traditional (or non-conversational learning) MAC learning mode. The default MAC learning mode for CE VLANs is traditional.  <b>Note</b> You cannot configure FP VLANs for the traditional MAC address learning mode.
Step 3	<b>exit</b>  <b>Example:</b> switch(config)# exit switch#	Exits configuration mode.
Step 4	<b>show mac address-table learning-mode {vlan <i>vlan-id</i>}</b>  <b>Example:</b> switch# show mac address-table learning-mode	(Optional) Displays the VLANs and the MAC learning mode.
Step 5	<b>copy running-config startup-config</b>  <b>Example:</b> switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

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

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

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

## 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.

Ensure that you have installed the Enhanced Layer 2 license.

## SUMMARY STEPS

1. **config terminal**
2. **fabricpath switch-id** *value*
3. **exit**
4. (Optional) **show fabricpath switch-id**
5. (Optional) **copy running-config startup-config**

## DETAILED STEPS

	Command	Purpose
Step 1	<b>config terminal</b>  <b>Example:</b> switch# config terminal switch(config)#	Enters global configuration mode.
Step 2	<b>fabricpath switch-id</b> <i>value</i>  <b>Example:</b> switch(config)# fabricpath switch-id 25 switch(config)#	Specifies the switch ID. The range is from 1 to 4094. There is no default value.

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

	Command	Purpose
Step 3	<b>exit</b>  <b>Example:</b> switch(config)# exit switch)#	Exits the configuration mode.
Step 4	<b>show fabricpath switch-id</b>  <b>Example:</b> switch# show fabricpath switch-id	(Optional) Displays information on the switch IDs.
Step 5	<b>copy running-config startup-config</b>  <b>Example:</b> switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

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

```
switch# config 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**—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.
- **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.

### BEFORE YOU BEGIN

Ensure that you are working on an F Series module.

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

Ensure that you have installed the Enhanced Layer 2 license.

### SUMMARY STEPS

1. **config terminal**
2. **fabricpath timers {allocate-delay seconds | linkup-delay seconds | transition-delay seconds}**
3. **exit**
4. (Optional) **show fabricpath timers**



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

5. (Optional) **copy running-config startup-config**

## DETAILED STEPS

	Command	Purpose
Step 1	<b>config terminal</b>  <b>Example:</b> switch# config terminal switch(config)#	Enters global configuration mode.
Step 2	<b>fabricpath timers {allocate-delay seconds   linkup-delay seconds   transition-delay seconds}</b>  <b>Example:</b> switch(config)# fabricpath allocate-delay 600 switch(config)#	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"> <li>• <b>allocate-delay</b>—10 seconds</li> <li>• <b>linkup-delay</b>—10 seconds</li> <li>• <b>transition-delay</b>—10 seconds</li> </ul>
Step 3	<b>exit</b>  <b>Example:</b> switch(config)# exit switch)#	Exits the configuration mode.
Step 4	<b>show fabricpath timers</b>  <b>Example:</b> switch# show fabricpath timers	(Optional) Displays information on the FabricPath timers.
Step 5	<b>copy running-config startup-config</b>  <b>Example:</b> switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to configure the allocation-delay FabricPath value to 600 seconds:

```
switch# config terminal
switch(config)# fabricpath timer allocate-delay 600
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 may 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.

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

Ensure that you have installed the Enhanced Layer 2 license.

SUMMARY STEPS

1. **config terminal**
2. **[no] fabricpath graceful-merge disable**
3. **exit**
4. (Optional) **show running-config**
5. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	<b>config terminal</b>  <b>Example:</b> switch# config terminal switch(config)#	Enters global configuration mode.
Step 2	<b>[no] fabricpath graceful-merge disable</b>  <b>Example:</b> switch(config)# fabricpath graceful-merge disable switch(config)#	Disables graceful merge of the FabricPath feature. To reenale this feature, enter the <b>no</b> form of the command.
Step 3	<b>exit</b>  <b>Example:</b> switch(config)# exit switch)#	Exits configuration mode.
Step 4	<b>show running-config</b>  <b>Example:</b> switch# show running-config	(Optional) Displays information on the configuration running on the switch.
Step 5	<b>copy running-config startup-config</b>  <b>Example:</b> switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

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

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

# 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.

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

**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.

Ensure that you have installed the Enhanced Layer 2 license.

**SUMMARY STEPS**

**Step 1** **fabricpath force link-bringup**

**DETAILED STEPS**

	Command	Purpose
<b>Step 1</b>	<b>fabricpath force link-bringup</b>  <b>Example:</b> switch# fabricpath force link-bringup switch#	Forces the FabricPath network links to come up as a one-time event.  <b>Note</b> This command is not saved when you enter the <b>copy running-config startup-config</b> command.

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

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

## Verifying FabricPath Switching

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

Command	Purpose
<b>show feature-set</b>	Displays whether FabricPath is enabled or not.
<b>show mac address-table learning-mode {vlan vlan-id}</b>	Displays the VLANs and the MAC address learning mode.  <b>Note</b> MAC learning modes are available only on the F Series modules.
<b>show fabricpath conflict {all [detail]   link [detail]   switch-id [detail]   transitions [detail]}</b>	Displays information on conflicts in the FabricPath network.
<b>show fabricpath switch-id [local]</b>	Displays information on the FabricPath network by switch ID.

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

Command	Purpose
<b>show fabricpath system-id</b> { <i>mac-addr</i> }	Displays information on the FabricPath network by system ID.
<b>show fabricpath timers</b>	Displays settings for the allocate-delay, linkup-delay, and transition-delay timers for the FabricPath network.

See [Chapter 5, “Advanced FabricPath Features,”](#) for more commands that display FabricPath switching functionality.

# Displaying 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 on these commands.

# FabricPath Switching Example Configuration

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 Fabric Path on all the devices.

```
switch# config 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# config terminal
```

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

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

**Step 4** Save the configuration.

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

## Default Settings

Table 2-1 lists the default settings for FabricPath switching parameters.

**Table 2-1** Default FabricPath Parameters

Parameters	Default
FabricPath	Disabled.
MAC address learning mode	<ul style="list-style-type: none"><li>FP VLANs—Only conversational learning</li><li>CE VLANs—Traditional (non-conversational) learning; can be configured for conversational learning on F Series modules</li></ul>
allocate-delay timer	10 seconds
linkup-delay timer	10 seconds
transition-delay timer	10 seconds
graceful merge	Enabled

## Feature History for Configuring FabricPath Switching Using the CLI

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

**Table 2-2** Feature History for FabricPath Switching

Feature Name	Releases	Feature Information
FabricPath	5.1(1)	These features were introduced.
New default MAC address learning mode in chassis containing both F Series and M Series modules	5.2(1)	This feature was introduced.

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



## CHAPTER 3

# FabricPath Interfaces

---

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

This chapter includes the following sections:

- [Information About FabricPath Interfaces, page 3-1](#)
- [Licensing Requirements for FabricPath Interfaces, page 3-7](#)
- [Prerequisites for FabricPath Interfaces, page 3-7](#)
- [Guidelines and Limitations for FabricPath Interfaces, page 3-8](#)
- [Configuring FabricPath Interfaces, page 3-8](#)
- [Verifying FabricPath Interface Configuration, page 3-15](#)
- [Displaying FabricPath Interface Statistics, page 3-16](#)
- [FabricPath Interface Example Configuration, page 3-16](#)
- [Feature History for Configuring FabricPath Using the CLI, page 3-18](#)

## Information About FabricPath Interfaces

This section includes the following sections:

- [FabricPath Interfaces, page 3-2](#)
- [STP and the FabricPath Network, page 3-2](#)
- [vPC+, page 3-4](#)
- [High Availability, page 3-6](#)
- [Virtual Device Contexts, page 3-6](#)



### Note

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

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

## 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 [Chapter 4, “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 *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 [Figure 3-1](#)).



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

**Figure 3-1 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 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 [Figure 3-1](#)).

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 may also be many separate STP domains (see [Figure 3-1](#)). 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

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

domain that connects to the same FabricPath network. The Layer 2 IS-IS messages carry the TCNs across the FabricPath network. Only those FabricPath Layer 2 gateway switches in the same STP domain as the TCN message need to act and propagate the message to connected CE 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 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, Release 5.x*). You cannot configure a vPC+ domain and a vPC domain on the same Cisco Nexus 7000 Series device.

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.



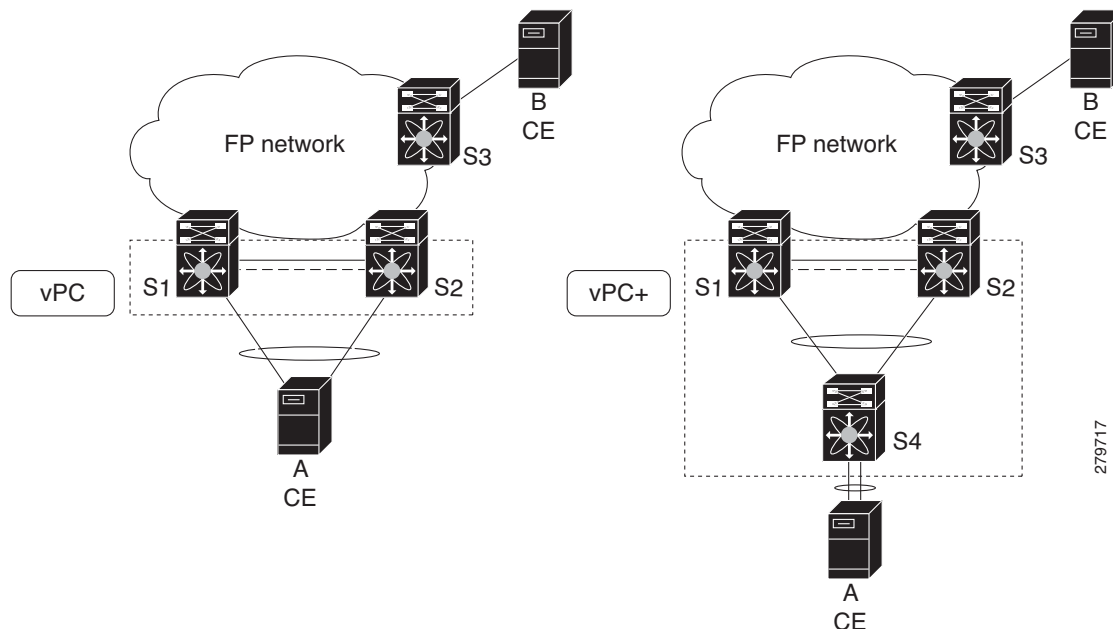
### Note

We do not recommend enabling the vPC peer switch feature when you are using vPC+.

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+ provides the solution by creating a unique virtual switch to the FabricPath network (see [Figure 3-2](#)).

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

**Figure 3-2 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 [Chapter 2, “Configuring FabricPath Switching”](#), for information on 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, Release 5.x* for information on proxy routing with the F1 Series modules).

The First Hop Routing Protocols (FHRP) 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.

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

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 for when a group state enabled on a vPC+ peer is in standby or in listen state. You can configure lower and upper thresholds to prevent the group state flap, if there is an interface flap (this feature is useful when there is more than one tracking object per group).

When the primary vPC+ peer 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 will allow Layer 2 switching only after both timers expire.

The delay option is only available 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, Release 5.x*, for more information on FHRPs and routing.

## 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, Release 5.x*, for more information on high availability.

## Virtual Device Contexts

All the interfaces for one FabricPath network on the same device must be in the same VDC.

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 following port pairs must be in the same VDC:

- Ports 1 and 2
- Ports 3 and 4

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

- Ports 5 and 6
- Ports 7 and 8
- Ports 9 and 10
- Ports 11 and 12
- Ports 13 and 14
- Ports 15 and 16
- Ports 17 and 18
- Ports 19 and 20
- Ports 21 and 22
- Ports 23 and 24
- Ports 25 and 26
- Ports 27 and 28
- Ports 29 and 30
- Ports 31 and 32

See the *Virtual Device Context Configuration Guide, Cisco DCNM for LAN, Release 5.x*, for more information on VDCs.

## Licensing Requirements for FabricPath Interfaces

The following table shows the licensing requirements for this feature:

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

## Prerequisites for FabricPath Interfaces

FabricPath interfaces have the following prerequisites:

- You should have a working knowledge of Classical Ethernet Layer 2 functioning.
- You must install the FabricPath feature set on the default and non-default VDC before you enable FabricPath on the switch. 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.
- You are logged onto the device.
- Ensure that you have installed the Enhanced Layer 2 license.
- You are in the correct virtual device context (VDC). A VDC is a logical representation of a set of system resources. You can use the **switchto vdc** command with a VDC number.
- You are working on the F Series module.

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

## Guidelines and Limitations for FabricPath Interfaces

FabricPath has the following configuration guidelines and limitations:

- FabricPath interfaces carry only FabricPath-encapsulated traffic.
- You must enable FabricPath on each device before you can view or access the commands. Enter the **feature-set fabricpath** command to enable FabricPath on each device.
- 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 the F Series module is in a VDC and that VDC has multiple SPAN destination ports, that SPAN session is not brought up.
- The following guidelines apply to private VLAN configurations when you are running FabricPath:
  - All VLANs in a private VLAN must be in the same VLAN mode; either CE or FP. If you attempt to put different types of VLANs into a private VLAN, these VLANs will not be active in the private VLAN. The system remembers the configurations, and if you change the VLAN mode later, that VLAN becomes active in the specified private VLAN.
  - FabricPath ports cannot be put into a private VLAN.
- The system does not support hierarchical static MAC addresses. 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.

## 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.

This section includes the following topics:

- [Configuring FabricPath Interfaces, page 3-9](#)
- [Configuring the STP Priority with Rapid PVST+, page 3-10](#)
- [Configuring the STP Priority with MST, page 3-11](#)
- [\(Optional\) Configuring the STP Domain ID for STP Domains Connected to the Layer 2 Gateway Switch, page 3-13](#)
- [Configuring a vPC+ Switch ID, page 3-14](#)



### Note

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

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

## Configuring FabricPath Interfaces



### Note

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

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

Ensure that you have installed the Enhanced Layer 2 license.

### SUMMARY STEPS

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

### DETAILED STEPS

	Command	Purpose
Step 1	<b>config terminal</b>  <b>Example:</b> switch# config terminal switch(config)#	Enters global configuration mode.
Step 2	<b>interface</b> [ethernet <i>slot/port</i>   <b>port-channel</b> <i>channel-no</i> ]  <b>Example:</b> switch(config)# interface ethernet 2/11-15 switch(config-if)#	Specifies the interfaces that you want to configure as FabricPath.
Step 3	<b>[no] switchport mode fabricpath</b>  <b>Example:</b> switch(config-if)# switchport mode fabricpath switch(config-if)#	Specifies interfaces as FabricPath port.  <b>Note</b> The <b>no</b> keyword returns the interface to the default CE access interface. The FabricPath ports carry traffic only for those VLANs that are configured as FabricPath VLANs.
Step 4	<b>system default switchport fabricpath</b>  <b>Example:</b> switch(config-if)# system default switchport fabricpath switch(config-if)#	(Optional) Converts all CE interfaces on the F Series module to FabricPath interfaces simultaneously.

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

	Command	Purpose
Step 5	<b>exit</b>  <b>Example:</b> switch(config-if)# exit switch(config)#	Exits interface configuration mode.
Step 6	<b>show interface</b>  <b>Example:</b> switch(config)# show interface	(Optional) Displays information on all interfaces.
Step 7	<b>copy running-config startup-config</b>  <b>Example:</b> switch(config)# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

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

```
switch# config 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.

Ensure that you have installed the Enhanced Layer 2 license.

### SUMMARY STEPS

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



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

## DETAILED STEPS

	Command	Purpose
Step 1	<b>config terminal</b>  <b>Example:</b> switch# config terminal switch(config)#	Enters global configuration mode.
Step 2	<b>spanning-tree vlan [vlan-id] priority [value]</b>  <b>Example:</b> switch(config)# spanning-tree vlan 11-20 priority 8192 switch(config)#	You must configure all the Rapid PVST+ VLANs on all the FabricPath Layer 2 gateway interfaces to a lower STP priority. We recommend that you configure the priority to be 8192.
Step 3	<b>exit</b>  <b>Example:</b> switch(config-if)# exit switch(config)#	Exits the interface configuration mode.
Step 4	<b>show spanning-tree summary</b>  <b>Example:</b> switch# show spanning-tree summary	(Optional) Displays information on STP.
Step 5	<b>copy running-config startup-config</b>  <b>Example:</b> switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

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

```
switch# config 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 on 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.

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

Ensure that you have installed the Enhanced Layer 2 license.

## SUMMARY STEPS

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

## DETAILED STEPS

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

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

```
switch# config 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 on this command.

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

## (Optional) 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.

Ensure that you have installed the Enhanced Layer 2 license.

### SUMMARY STEPS

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

### DETAILED STEPS

	Command	Purpose
Step 1	<b>config terminal</b>  <b>Example:</b> switch# config terminal switch(config)#	Enters global configuration mode.
Step 2	<b>spanning-tree domain <i>domain-id</i></b>  <b>Example:</b> switch(config)# spanning-tree domain 5 switch(config)#	Assigns an STP domain ID to the different STP domains attached to FabricPath Layer 2 gateway switches that are connected to a single FabricPath network. The range is from 1 to 1023.
Step 3	<b>exit</b>  <b>Example:</b> switch(config)# exit switch#	Exits configuration mode.
Step 4	<b>show spanning-tree summary</b>  <b>Example:</b> switch# show spanning-tree summary	(Optional) Displays information on STP.
Step 5	<b>copy running-config startup-config</b>  <b>Example:</b> switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

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

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

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

## 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).

See the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 5.x*, for complete information on 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).

## SUMMARY STEPS

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

## DETAILED STEPS

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

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

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

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
```

## Verifying FabricPath Interface Configuration

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

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

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

Command	Purpose
<b>show vpc statistics</b>	Displays statistics on the vPC+ domains.
<b>show running-config vpc</b>	Displays running configuration information for vPCs and vPC+ domains.

For information on 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*.

## Displaying FabricPath Interface Statistics

Use the following command 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*]

## FabricPath Interface Example Configuration

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# config 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# config terminal
```

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

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

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

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

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

```
switch# config 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, Release 5.x*, for information on 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.

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

## Feature History for Configuring FabricPath Using the CLI

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

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

Feature Name	Releases	Feature Information
FabricPath Interfaces	5.1(1)	These features were introduced.





## CHAPTER 4

# FabricPath Forwarding

---

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

This chapter includes the following sections:

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

## Information About FabricPath Forwarding



### Note

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

This section includes the following sections:

- [FabricPath Forwarding Overview, page 4-2](#)
- [FabricPath VLANs, page 4-2](#)
- [Forwarding Known Unicast Packets Using ECMP, page 4-3](#)
- [Forwarding Trees for Broadcast, Unknown Unicast, and Multicast Packets, page 4-4](#)
- [High Availability, page 4-5](#)
- [Virtual Device Context, page 4-5](#)

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

## FabricPath Forwarding Overview

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

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

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

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

FabricPath Layer 2 IS-IS defines the trees. The highest system ID is chosen for the root and the tree flows from that. The system ID is configurable per device. The second tree is the same but with a different root priority. After the system chooses the root switch, the tree is built with that as the root for the first tree. Then, the root switch for the first tree elects the root of the second tree, again based on system ID, and the second tree flows from that root switch. All of this information is advertised to the FabricPath network using Layer 2 IS-IS, so all the devices in the network have the same information.

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

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

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

## FabricPath VLANs


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

The system automatically assigns all FabricPath interfaces and FP VLANs to the topology. So, there is no added configuration required. (See [Chapter 3, “FabricPath Interfaces”](#), for information on FabricPath interfaces.) All the FP VLANs and FabricPath interfaces belong to that same topology. All ports on the same device in the same topology must be in the same virtual device context (VDC).

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

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

**Figure 4-1 Example FabricPath Topology and Classical Ethernet Hosts**



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

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



**Note**

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

## Forwarding Known Unicast Packets Using ECMP

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

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

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

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

## Forwarding Trees for Broadcast, Unknown Unicast, and Multicast Packets

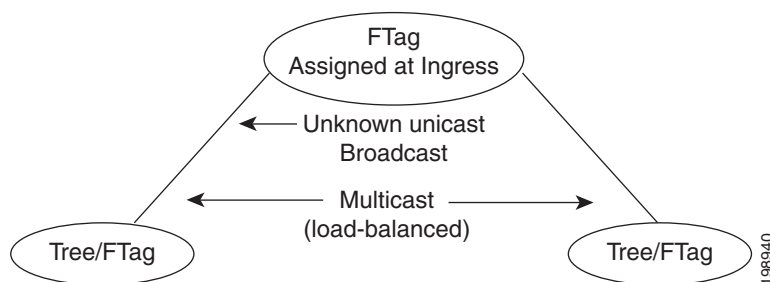
FabricPath introduces a new loop-free broadcast functionality that carries broadcast, unknown unicast, and multicast packets, or multidestination traffic. For each broadcast, unknown unicast, and multicast traffic flow, the system chooses the forwarding path from among multiple system-created paths or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

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

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

Figure 4-2 shows these trees.

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



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

## Forwarding Multicast Packets

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

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

The multicast traffic uses the per-VLAN source, multicast group, and flow information to allocate traffic to one or the other of the two trees. This system constrains multicast based on the group IP address.

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

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

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

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

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

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

## High Availability

The FabricPath topologies retain their configuration through ISSU.

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

## Virtual Device Context

All the interfaces for one FabricPath network on the same device must be in the same VDC.

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

Because of the multiple forwarding engines (FEs) on the F Series modules, the following port pairs must be in the same VDC:

- Ports 1 and 2
- Ports 3 and 4
- Ports 5 and 6
- Ports 7 and 8
- Ports 9 and 10
- Ports 11 and 12
- Ports 13 and 14
- Ports 15 and 16
- Ports 17 and 18
- Ports 19 and 20
- Ports 21 and 22
- Ports 23 and 24
- Ports 25 and 26
- Ports 27 and 28
- Ports 29 and 30
- Ports 31 and 32

See the *Virtual Device Context Configuration Guide, Cisco DCNM for LAN, Release 5.x*, for more information on VDCs.

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

## Licensing Requirements for FabricPath Forwarding

The following table shows the licensing requirements for this feature:

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

## Prerequisites for FabricPath Forwarding

FabricPath forwarding has the following prerequisites:

- You should have a working knowledge of Classical Ethernet Layer 2 functioning.
- You must install the FabricPath feature set in the default and non-default VDC before you enable FabricPath on the switch. See the *Configuring Feature Set for FabricPath* for complete information on installing and enabling the FabricPath feature set.
- You are logged onto the device.
- Ensure that you have installed the Enhanced Layer 2 license.
- You are in the correct virtual device context (VDC). A VDC is a logical representation of a set of system resources. You can use the **switchto vdc** command with a VDC number.
- You are working on the F Series module.

## Guidelines and Limitations for FabricPath Forwarding

FabricPath has the following configuration guidelines and limitations:

- FabricPath interfaces carry only FabricPath-encapsulated traffic.
- You must enable FabricPath on each device before you can view or access the commands. Enter the **feature-set fabricpath** command to enable FabricPath on each device. See the *Configuring Feature Set for FabricPath* for complete information on installing and enabling the FabricPath feature set.
- The FabricPath feature set operation may cause the standby supervisor to reload if it is in an unstable state, such as following a service failure or powering up.
- STP does not run inside a FabricPath network.
- The F Series modules do not support multiple SPAN destination ports or virtual SPAN. If a port on the F Series module is in a VDC and that VDC has multiple SPAN destination ports, that SPAN session is not brought up.
- The following guidelines apply to private VLAN configuration when you are running FabricPath:
  - All VLANs in a private VLAN must be in the same VLAN mode; either CE or FP. If you attempt to put different types of VLANs into a private VLAN, these VLANs will not be active in the private VLAN. The system remembers the configurations, and if you change the VLAN mode later, that VLAN becomes active in the specified private VLAN.
  - FabricPath ports cannot be put into a private VLAN.
- The system does not support hierarchical static MAC addresses or static routes.

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

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

## Configuring FabricPath Forwarding

**Note**

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

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

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

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

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

**Note**

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

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

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

## Setting the VLAN Mode to FP or CE

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

**Note**

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

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

### BEFORE YOU BEGIN

Ensure that you are working on an F Series module.

Ensure that you have installed the Enhanced Layer 2 license.

Ensure that you have enabled the FabricPath feature set.

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

Ensure that you have created the VLANs.

## SUMMARY STEPS

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

## DETAILED STEPS

	Command	Purpose
Step 1	<b>config terminal</b>  <b>Example:</b> switch# config terminal switch(config)#	Enters global configuration mode.
Step 2	<b>vlan <i>vlan-id</i></b>  <b>Example:</b> switch(config)# vlan 1-10 switch(config-vlan)#	Enters the VLAN configuration mode and identifies those VLANs that you want to carry FabricPath traffic.
Step 3	<b>mode [ce   fabricpath]</b>  <b>Example:</b> switch(config-vlan)# mode fabricpath switch(config-vlan)#	Configures the VLANs as FP VLANs. The default VLAN mode is CE.  <b>Note</b> A VLAN must be either a CE or an FP VLAN on the FabricPath device.
Step 4	<b>exit</b>  <b>Example:</b> switch(config-vlan)# exit switch(config)#	Exits the VLAN configuration mode.  <b>Note</b> As with all VLANs, you must exit the VLAN configuration mode for the VLAN mode (CE or FP) to take effect.
Step 5	<b>show fabricpath topology vlans [active]</b>  <b>Example:</b> switch# show fabricpath topology vlans	(Optional) Displays information about all VLANs in the FabricPath topology.
Step 6	<b>copy running-config startup-config</b>  <b>Example:</b> switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to specify VLANs as FP VLANs:

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

You can add interfaces to a non-default topology:



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

## Configuring FabricPath Unicast Load Balancing (Optional)

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

### BEFORE YOU BEGIN

Ensure that you are working on an F Series module.

Ensure that you have installed the Enhanced Layer 2 license.

Ensure that you have enabled the FabricPath feature set.

### SUMMARY STEPS

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

### DETAILED STEPS

	Command	Purpose
Step 1	<b>config terminal</b>  <b>Example:</b> switch# config terminal switch(config)#	Enters global configuration mode.
Step 2	<b>fabricpath load-balance unicast</b> {{{source   source-destination   xor   destination   symmetric}} [layer 3   layer4   mixed]} [rotate-amount <i>rot_amt</i> ] [include-vlan]}  <b>Example:</b> switch(config)# fabricpath load-balance unicast source switch(config-vlan)#	Configures load balancing for FabricPath unicast traffic along specified parameters.  <b>Note</b> To return to the default unicast load-balancing scheme, enter the <b>no</b> form of this command.
Step 3	<b>exit</b>  <b>Example:</b> switch(config)# exit switch(config)#	Exits configuration mode.
Step 4	<b>copy running-config startup-config</b>  <b>Example:</b> switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to configure FabricPath unicast load balancing:

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

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

## Configuring FabricPath Multicast Load Balancing (Optional)

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

### BEFORE YOU BEGIN

Ensure that you are working on an F Series module.

Ensure that you have installed the Enhanced Layer 2 license.

Ensure that you have enabled the FabricPath feature set.

### SUMMARY STEPS

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

### DETAILED STEPS

	Command	Purpose
Step 1	<b>config terminal</b>  <b>Example:</b> switch# config terminal switch(config)#	Enters global configuration mode.
Step 2	<b>fabricpath load-balance multicast</b> {{{source   source-destination   xor   destination   symmetric}}} [rotate-amount <i>rot_amt</i> ] [include-vlan]}  <b>Example:</b> switch(config)# fabricpath load-balance multicast source switch(config-vlan)#	Configures load balancing for FabricPath multicast traffic along specified parameters.  <b>Note</b> To return to the default multicast load-balancing scheme, enter the <b>no</b> form of this command.
Step 3	<b>exit</b>  <b>Example:</b> switch(config)# exit switch(config)#	Exits configuration mode.
Step 4	<b>copy running-config startup-config</b>  <b>Example:</b> switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to configure FabricPath multicast load balancing:

```
switch# config terminal
```

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

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

## Configuring FabricPath Increased Multicast Scalability (Optional)

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

### BEFORE YOU BEGIN

Ensure that you are working on an F Series module.

Ensure that you have installed the Enhanced Layer 2 license.

Ensure that you have enabled the FabricPath feature set.

### SUMMARY STEPS

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

### DETAILED STEPS

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

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

	Command	Purpose
Step 4	<b>exit</b>  <b>Example:</b> switch(config)# exit switch(config)#	Exits configuration mode.
Step 5	<b>copy running-config startup-config</b>  <b>Example:</b> switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to configure increased FabricPath multicast scalability:

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

## Verifying the FabricPath Configuration

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

Command	Purpose
<b>show feature-set</b>	Displays whether FabricPath is enabled or not.
<b>show {l2   fabricpath} route [switchid switch-id] [detail] [hex]</b>	Displays unicast routes.
<b>show {l2   fabricpath} mroute {[vdc_omf]   vlan vlan-id { {omf   flood   [source {srcaddr   v6srcaddr}] [group {groupaddr   v6grpaddr}] } [resolved] [ftag ftag-id] [hex]}</b>	Displays multicast routes.
<b>show fabricpath topology [detail]</b>	Displays information on all FabricPath topologies.
<b>show fabricpath topology interface</b>	Displays information on all FabricPath topology interfaces.
<b>show fabricpath topology vlan [active]</b>	Displays information on all FabricPath topology VLANs.
<b>show fabricpath topology ftag [active] [multicast] [unicast]</b>	Displays information on all FabricPath topology FTags.
<b>show running-config fabricpath</b>	Displays the running configuration for FabricPath.
<b>show fabricpath load-balance unicast forwarding-path ftag ftag-id switchid switch-id flow-type {l2 { {dst-mac dst-mac   source-mac src-mac} ether-type ether-type} }   {l3 {dst-ip dst-ip   src-ip src-ip   dst-ipv6 dst-ipv6   srcip6 src-ipv6} }   {l4 {l4-src-port l4-src-port   l4-dst-port l4-dst-port   dst-ip dst-ip   src-ip src-ip   dst-ipv6 dst-ipv6   srcip6 src-ipv6} } {vlan vlan-id} {module mod-no}</b>	Displays FabricPath unicast load-balancing information.

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

Command	Purpose
<b>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> }</b>	Displays FabricPath multicast load-balancing information.
<b>show vlan</b>	Displays information on all FP and CE VLANs.

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

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

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

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

## FabricPath Forwarding Example Configuration

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

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

To configure the default FabricPath topology, follow these steps:

### Step 1 Enable the FabricPath feature set.

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



#### Note

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

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

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

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

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

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

**Step 4** Save the configuration.

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

## Default Settings

Table 4-1 lists the default settings for FabricPath parameters

**Table 4-1**      **Default FabricPath Parameters**

Parameters	Default
FabricPath Topology	0
VLAN mode	CE

## Feature History for Configuring FabricPath Using the CLI

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

**Table 4-2**      **Feature History for FabricPath**

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



## CHAPTER 5

# 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.

This chapter includes the following sections:

- [Information About FabricPath Advanced Features, page 5-1](#)
- [Licensing Requirements for FabricPath Advanced Features, page 5-3](#)
- [Prerequisites for FabricPath Advanced Features, page 5-3](#)
- [Guidelines and Limitations for FabricPath Advanced Features, page 5-3](#)
- [Setting Advanced FabricPath Layer 2 IS-IS Parameters, page 5-4](#)
- [Verifying the FabricPath Advanced Configurations, page 5-12](#)
- [Feature History for Configuring FabricPath Advanced Features Using the CLI, page 5-13](#)

## Information About FabricPath Advanced Features



**Note**

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

This section includes the following sections:

- [Information About Advanced FabricPath Layer 2 IS-IS Configurations, page 5-1](#)
- [High Availability, page 5-2](#)
- [Virtual Device Contexts, page 5-2](#)

## Information About Advanced FabricPath Layer 2 IS-IS Configurations



**Note**

See [Chapter 2, “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:

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

- 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 [Chapter 2, “Configuring FabricPath Switching”](#), for information on default FabricPath Layer 2 IS-IS behavior and [Chapter 4, “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.

## High Availability

The FabricPath topologies retain their configuration through ISSU.

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

## Virtual Device Contexts

All the interfaces for one FabricPath network on the same device must be in the same Virtual Device Context (VDC).

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



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

- Ports 29 and 30
- Ports 31 and 32

See the *Virtual Device Context Configuration Guide, Cisco DCNM for LAN, Release 5.x*, for more information on VDCs.

## Licensing Requirements for FabricPath Advanced Features

The following table shows the licensing requirements for this feature:

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

## Prerequisites for FabricPath Advanced Features

FabricPath switching has the following prerequisites:

- You should have a working knowledge of Classical Ethernet Layer 2 functioning.
- You must install the FabricPath feature set in the default and non-default VDC before you enable FabricPath on the switch. See *Configuring Feature-Set for FabricPath* for complete information on installing and enabling the FabricPath feature set.
- The FabricPath feature set operation may cause the standby supervisor to reload if it is in an unstable state, such as following a service failure or powering up.
- You are logged onto the device.
- Ensure that you have installed the Enhanced Layer 2 license.
- You are in the correct virtual device context (VDC). A VDC is a logical representation of a set of system resources. You can use the **switchto vdc** command with a VDC number.
- You are working on the F Series module.

## Guidelines and Limitations for FabricPath Advanced Features

FabricPath has the following configuration guidelines and limitations:

- FabricPath interfaces carry only FabricPath-encapsulated traffic.
- You must enable FabricPath on each device before you can view or access the commands. Enter the **feature-set fabricpath** command to enable FabricPath on each device.
- STP does not run inside a FabricPath network.
- The F Series modules do not support multiple SPAN destination ports or virtual SPAN. If a port on the F Series module is in a VDC and that VDC has multiple SPAN destination ports, that SPAN session is not brought up.
- The following guidelines apply to private VLAN configuration when you are running FabricPath:

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

- All VLANs in a private VLAN must be in the same VLAN mode; either CE or FP. If you attempt to put different types of VLANs into a private VLAN, these VLANs will not be active in the Private VLAN. The system remembers the configurations, and if you change the VLAN mode later, that VLAN will now become 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. This section includes the following topics:

- [Setting Advanced FabricPath Layer 2 IS-IS Parameters Globally \(Optional\), page 5-4](#)
- [Setting Advanced FabricPath Layer 2 IS-IS Parameters per Interface \(Optional\), page 5-8](#)
- [Clearing Advanced FabricPath Layer 2 IS-IS Counters, page 5-11](#)

## 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 installed the Enhanced Layer 2 license.

Ensure that you have enabled the FabricPath feature set.

### SUMMARY STEPS

1. **config terminal**
2. **fabricpath domain default**
3. (Optional) **authentication-check**
4. (Optional) **authentication key-chain** *auth-key-chain-name*
5. (Optional) **authentication type** { **cleartext** | **md5** }
6. (Optional) **log-adjacency-changes**
7. (Optional) **lsp-gen-interval** *msecs* [*msecs msecs*]
8. (Optional) **lsp-mtu** *mtu*

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

9. (Optional) **max-lsp-lifetime** *secs*
10. (Optional) **maximum-paths** *max-paths*
11. (Optional) **reference-bandwidth** {*ref-mbps* [Mbps] | *ref-gbps* [Gbps]}
12. (Optional) **spf-interval** *msecs* [*msecs msecs*]
13. (Optional) **graceful-restart** [**t3 manual** *secs*]
14. (Optional) **hostname** *dynamic*
15. (Optional) **root-priority** *value*
16. **exit**
17. **exit**
18. (Optional) **show running-config**
19. (Optional) **copy running-config startup-config**

## DETAILED STEPS

Command	Purpose
<b>config terminal</b>  <b>Example:</b> switch#config terminal	Enters global configuration mode.
<b>fabricpath domain default</b>  <b>Example:</b> switch(config)# fabricpath domain default switch(config-fabricpath-isis)#	Enters the global FabricPath Layer 2 IS-IS configuration mode.
<b>authentication-check</b>  <b>Example:</b> switch(config-fabricpath-isis)# authentication-check switch(config-fabricpath-isis)#	(Optional) Configures authentication check on PDU reception. To turn the authentication check off, enter the <b>no</b> form of this command.  <b>Note</b> The default is ON.
<b>authentication key-chain</b> <i>auth-key-chain-name</i>  <b>Example:</b> switch(config-fabricpath-isis)# authentication key-chain trees switch(config-fabricpath-isis)#	(Optional) Configures the authentication keychain. To clear this parameter, enter the <b>no</b> form of this command. See the <i>Security Configuration Guide, Cisco DCNM for LAN, Release 5.x</i> , for information about key chains.
<b>authentication-type</b> { <i>cleartext</i>   <i>md5</i> }  <b>Example:</b> switch(config-fabricpath-isis)# authentication-type md5 switch(config-fabricpath-isis)#	(Optional) Configures the authentication type. To clear this parameter, enter the <b>no</b> form of this command.
<b>log-adjacency-changes</b>  <b>Example:</b> switch(config-fabricpath-isis)# log-adjacency-changes switch(config-fabricpath-isis)#	(Optional) 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 <b>no</b> form of this command. The default is Off.

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

Command	Purpose
<b>lsp-gen-interval</b> <i>msecs</i> [ <i>msecs msecs</i> ]  <b>Example:</b> <pre>switch(config-fabricpath-isis)# lsp-gen-interval 100 switch(config-fabricpath-isis)#</pre>	<p>(Optional) Configures the LSP generation interval. To return to the default values, enter the <b>no</b> form of this command. The optional parameters are as follows:</p> <ul style="list-style-type: none"> <li>• <b>max-wait</b>—The initial wait between the trigger and LSP generation. The range is from 50 to 12000 milliseconds, and the default value is 8000 milliseconds.</li> <li>• <b>lsp-initial-wait</b>—The initial wait between the trigger and LSP generation. The range is from 50 to 12000 milliseconds, and the default value is 50 milliseconds.</li> <li>• <b>lsp-second-wait</b>—The second wait used for LSP throttle during backoff. The range is from 50 to 12000 milliseconds, and the default value is 50 milliseconds.</li> </ul>
<b>lsp-mtu</b> <i>mtu</i>  <b>Example:</b> <pre>switch(config-fabricpath-isis)# lsp-mtu 2000 switch(config-fabricpath-isis)#</pre>	<p>(Optional) Sets the LSP MTU. To return to the default values, enter the <b>no</b> form of this command. The range is from 128 to 4352, and the default value is 1492.</p>
<b>max-lsp-lifetime</b> <i>secs</i>  <b>Example:</b> <pre>switch(config-fabricpath-isis)# max-lsp-lifetime 1000 switch(config-fabricpath-isis)#</pre>	<p>(Optional) Sets the maximum LSP lifetime in seconds. To return to the default values, enter the <b>no</b> form of this command. The range is from 128 to 4352, and the default value is 1492.</p>
<b>maximum-paths</b> <i>max-paths</i>  <b>Example:</b> <pre>switch(config-fabricpath-isis)# maximum-paths 4 switch(config-fabricpath-isis)#</pre>	<p>(Optional) Sets the maximum number of paths per destination. To return to the default values, enter the <b>no</b> form of this command. The range is from 1 to 16, and the default value is 16.</p>
<b>reference-bandwidth</b> { <i>ref-mbps</i> [ <b>Mbps</b> ]   <i>ref-gbps</i> [ <b>Gbps</b> ]}  <b>Example:</b> <pre>switch(config-fabricpath-isis)# reference-bandwidth 200000 switch(config-fabricpath-isis)#</pre>	<p>(Optional) 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 <b>no</b> form of this command. The optional parameters are as follows:</p> <ul style="list-style-type: none"> <li>• <b>ref-mbps</b>—The range is from 1 to 400000, and the default value is 400000.</li> <li>• <b>ref-gbps</b>—The range is from 1 to 4000, and the default value is 400.</li> </ul>

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

Command	Purpose
<b>spf-interval</b> <i>msecs</i> [ <i>msecs msecs</i> ]  <b>Example:</b> switch(config-fabricpath-isis)# spf-interval 10000 switch(config-fabricpath-isis)#	(Optional) Configures the interval between LSA arrivals. To return to the default values, enter the <b>no</b> form of this command. The optional parameters are as follows: <ul style="list-style-type: none"> <li>• <b>spf-max-wait</b>—The maximum wait between the trigger and SPF computation. The range is from 50 to 120000 milliseconds, and the default value is 8000 milliseconds.</li> <li>• <b>spf-initial-wait</b>—The initial wait between the trigger and SPF computation. The range is from 50 to 120000 milliseconds, and the default value is 50 milliseconds.</li> <li>• <b>spf-second-wait</b>—The second wait used for SPF computation during backoff. The range is from 50 to 120000 milliseconds, and the default value is 50 milliseconds.</li> </ul>
<b>graceful-restart</b> [ <b>t3 manual secs</b> ]  <b>Example:</b> switch(config-fabricpath-isis)# graceful-restart switch(config-fabricpath-isis)#	(Optional) Enables graceful restart for the FabricPath Layer 2 IS-IS protocol. To disable graceful restart, enter the <b>no</b> form of this command. Use the <b>t3 manual</b> parameter 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.
<b>hostname dynamic</b>  <b>Example:</b> switch(config-fabricpath-isis)# hostname-dynamic switch(config-fabricpath-isis)#	(Optional) Enables dynamic hostname for the FabricPath Layer 2 IS-IS protocol. To disable the dynamic hostname, enter the <b>no</b> form of this command.
<b>root-priority</b> <i>value</i>  <b>Example:</b> switch(config-fabricpath-isis)# root-priority 100 switch(config-fabricpath-isis)#	(Optional) Sets 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 <b>no</b> form of this command. The range is from 1 to 255, and the default value is 64.
<b>exit</b>  <b>Example:</b> switch(config-fabricpath-isis)# exit switch(config)#	Exits global FabricPath Layer 2 IS-IS configuration mode.
<b>exit</b>  <b>Example:</b> switch(config)# exit switch#	Exits global configuration mode.

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

Command	Purpose
<b>show running-config</b>  <b>Example:</b> switch# show running-config switch#	(Optional) Displays the running configuration.
<b>copy running-config startup-config</b>  <b>Example:</b> switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 5.x*, for more information on 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 installed the Enhanced Layer 2 license.

Ensure that you have enabled the FabricPath feature set.

### SUMMARY STEPS

1. **config terminal**
2. **interface { ethernet *mod/slot* | port-channel *channel-number* }**
3. (Optional) **fabricpath isis authentication-check**
4. (Optional) **fabricpath isis authentication key-chain *auth-key-chain-name***
5. (Optional) **fabricpath isis authentication type { cleartext | md5 }**
6. (Optional) **fabricpath isis csnp-interval *seconds***
7. (Optional) **fabricpath isis hello-interval *seconds***
8. (Optional) **fabricpath isis hello-multiplier *multiplier***
9. (Optional) **fabricpath isis hello-padding**
10. (Optional) **fabricpath isis lsp-interval *milliseconds***
11. (Optional) **fabricpath isis metric *metric***
12. (Optional) **fabricpath isis retransmit-interval *seconds***
13. (Optional) **fabricpath isis retransmit-throttle- interval *milliseconds***
14. **exit**
15. (Optional) **show running-config**
16. (Optional) **copy running-config startup-config**

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

## DETAILED STEPS

Command	Purpose
<b>config terminal</b>  <b>Example:</b> <pre>switch# config switch(config)#</pre>	Enters configuration mode.
<b>interface</b> { <b>ethernet</b> <i>mod/slot</i>   <b>port-channel</b> <i>channel-number</i> }  <b>Example:</b> <pre>switch(config)# interface ethernet 5/2 switch(config-if)#</pre>	Enters interface configuration mode.
<b>fabricpath isis authentication-check</b>  <b>Example:</b> <pre>switch(config-if)# fabricpath isis authentication-check switch(config-if)#</pre>	(Optional) Enables authentication checking on incoming FabricPath Layer 2 IS-IS hello PDUs. The default is ON. To disable authentication, enter the <b>no</b> form of the command.  <b>Note</b> Level specification is not required.
<b>fabricpath isis authentication key-chain</b> <i>auth-key-chain-name</i>  <b>Example:</b> <pre>switch(config-if)# fabricpath isis authentication key-chain trees switch(config-if)#</pre>	(Optional) Assigns a password to authentication hello PDUs. To remove this password, enter the <b>no</b> form of the command.  <b>Note</b> Level specification is not required. See the <i>Security Configuration Guide, Cisco DCNM for LAN, Release 5.x</i> , for information about key chains.
<b>fabricpath isis authentication-type</b> { <b>cleartext</b>   <b>md5</b> }  <b>Example:</b> <pre>switch(config-if)# fabricpath isis authentication-type md5 switch(config-if)#</pre>	(Optional) Specifies the authentication type for an interface for FabricPath Layer 2 IS-IS hello PDUs. To remove this type, enter the <b>no</b> form of the command.  <b>Note</b> Level specification is not required.
<b>fabricpath isis csnp-interval</b> <i>seconds</i>  <b>Example:</b> <pre>switch(config-if)# fabricpath isis csnp-interval 60 switch(config-if)#</pre>	(Optional) Specifies the interval between CSNP PDUs sent on the interface. To return to the default value, enter the <b>no</b> form of this command. The range is from 1 to 65535, and the default value is 10.
<b>fabricpath isis hello-interval</b> <i>seconds</i>  <b>Example:</b> <pre>switch(config-if)# fabricpath isis hello-interval 20 switch(config-if)#</pre>	(Optional) Sets the hello interval between PDUs sent on the interface. To return to the default value, enter the <b>no</b> form of this command. The range is from 1 to 65535, and the default value is 10.  <b>Note</b> Level specification is not required.

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

Command	Purpose
<b>fabricpath isis hello-multiplier</b> <i>multiplier</i>  <b>Example:</b> <pre>switch(config-if)# fabricpath isis hello-multiplier 20 switch(config-if)#</pre>	(Optional) 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 <b>no</b> form of this command. The range is from 3 to 1000. The default is 3.  <b>Note</b> Level specification is not required.
<b>fabricpath isis hello-padding</b>  <b>Example:</b> <pre>switch(config-if)# fabricpath hello-padding switch(config-if)#</pre>	(Optional) Enables padding on the hello PDUs. The default is ON. To disable authentication, enter the <b>no</b> form of the command.  <b>Note</b> If you enter the <b>always</b> keyword with the <b>no</b> form of this command, the padding is always on.
<b>fabricpath isis lsp-interval</b> <i>milliseconds</i>  <b>Example:</b> <pre>switch(config-if)# fabricpath isis lsp-interval 100 switch(config-if)#</pre>	(Optional) Sets the interval in milliseconds between LSPs sent on this interface during flooding. To return to the default value, enter the <b>no</b> form of this command. The range is from 10 to 65535. The default is 33.
<b>fabricpath isis metric</b> <i>metric</i>  <b>Example:</b> <pre>switch(config-if)# fabricpath isis metric 100 switch(config-if)#</pre>	(Optional) 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 <b>no</b> 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"> <li>• 1 GB—400</li> <li>• 10 GB—40</li> </ul>
<b>fabricpath isis retransmit-interval</b> <i>seconds</i>  <b>Example:</b> <pre>switch(config-if)# fabricpath isis retransmit-interval 100 switch(config-if)#</pre>	(Optional) Sets the interval between initial LSP retransmissions. To return to the default value, enter the <b>no</b> form of this command. The range is from 1 to 65535. The default is 5.
<b>fabricpath isis retransmit-throttle-interval</b> <i>milliseconds</i>  <b>Example:</b> <pre>switch(config-if)# fabricpath isis retransmit-throttle-interval 100 switch(config-if)#</pre>	(Optional) Sets the interval between subsequent LSP retransmissions. To return to the default value, enter the <b>no</b> form of this command. The range is from 20 to 65535. The default is 66.
<b>exit</b>  <b>Example:</b> <pre>switch(config-if)# exit switch(config)#</pre>	Exits interface configuration mode.



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

Command	Purpose
<b>show running-config</b>  <b>Example:</b> switch(config)# show running-config switch(config)#	(Optional) Displays the running configuration.
<b>copy running-config startup-config</b>  <b>Example:</b> switch(config)# copy running-config startup-config switch(config)#	(Optional) Copies the running configuration to the startup configuration.

See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 5.x*, for more information on 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 installed the Enhanced Layer 2 license.

Ensure that you have enabled the FabricPath feature set.

### SUMMARY STEPS

- (Optional) **clear fabricpath isis adjacency** [\* | system-id | interface {ethernet mod/slot | port-channel channel-number}]
- (Optional) **clear fabricpath isis statistics** \*
- (Optional) **clear fabricpath isis traffic** [\* | interface {ethernet mod/slot | port-channel channel-number}]

### DETAILED STEPS

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

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

Command	Purpose
<b>clear fabricpath isis statistics *</b>  <b>Example:</b> switch# clear fabricpath statistics * switch#	(Optional) Clears all FabricPath Layer 2 IS-IS protocol statistics.
<b>clear fabricpath isis traffic [*   interface {ethernet mod/slot   port-channel channel-number}]</b>  <b>Example:</b> switch# clear fabricpath traffic switch#	(Optional) Clears FabricPath Layer 2 IS-IS traffic information.

## Verifying the FabricPath Advanced Configurations

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

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

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

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

## Feature History for Configuring FabricPath Advanced Features Using the CLI

Table 5-1 lists the release history for these features.

**Table 5-1** Feature History for FabricPath

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

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

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



## APPENDIX **A**

# Configuration Limits for Cisco NX-OS FabricPath

---

The features supported by Cisco NX-OS have maximum configuration limits.

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

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



## INDEX

---

### C

configuration limits  
    description (table) [1-1](#)

---

### D

documentation  
    additional publications [1-viii](#)  
    related documents [1-viii](#)

---

### F

FabricPath  
    displaying [2-14](#)  
Fabtic Path  
    displaying [4-14, 5-33](#)  
field descriptions  
    vPCs [2-14, 3-16](#)

---

### I

IGMP snooping  
    prerequisites [2-8, 3-7, 4-6, 5-6](#)  
interfaces  
    tunnel [3-15](#)

---

### L

L2MP  
    configuring [4-8, 5-7, 5-23](#)  
    disabling [2-11](#)  
    displaying [3-16](#)

enabling [2-10](#)  
troubleshooting [3-9, 4-8, 5-7, 5-18, 5-23, 5-30](#)  
VLANs [2-9, 3-8, 4-7, 4-8, 5-6](#)

#### LACP

vPCs [2-14, 3-16](#)

#### limits

description (table) [1-1](#)

---

### M

mixed chassis [2-7](#)  
M Series modules [2-7](#)

---

### P

Port Channel pane  
    icon [2-14, 3-16](#)

---

### R

related documents [1-viii](#)

---

### S

switching [2-1](#)

---

### T

topology  
    licensing requirements [2-8, 3-7, 4-6, 5-5](#)  
troubleshooting [2-5, 3-4, 3-8, 4-8](#)  
    vPCs [2-14, 3-16](#)  
tunnels

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

creating [3-15](#)

---

## V

### vPC

configuring primary peer device [2-14, 3-16](#)

configuring secondary peer device [2-14, 3-16](#)

downstream port channels [2-14, 3-16](#)

role priority [2-14, 3-16](#)

system MAC address [2-14, 3-16](#)

system priority [2-14, 3-16](#)

### vPC peer link

synchronize configurations [2-14, 3-16](#)

### vPCs

Auto-Resolve option [2-14, 3-16](#)

compatible configurations [2-14, 3-16](#)

creating [2-14, 3-16](#)

deleting [2-14, 3-16](#)

disabling [2-14, 3-16](#)

enabling [2-14, 3-16](#)

field descriptions [2-14, 3-16](#)

icons [2-14, 3-16](#)

keepalive messages [2-14, 3-16](#)

LACP [2-14, 3-16](#)

modifying [2-14, 3-16](#)

peer-keepalive link

    configuring [2-14, 3-16](#)

port channels [2-14, 3-16](#)

required configuration consistency [2-14, 3-16](#)

role priorities [2-14, 3-16](#)

statistics [2-14, 3-16](#)

synchronizing configurations [2-14, 3-16](#)

synchronizing the vPC peer links [2-14, 3-16](#)

synchronizing the vPCs [2-14, 3-16](#)

Topology tab [2-14, 3-16](#)

troubleshooting [2-14, 3-16](#)

vPC system priority

    configuring [2-14, 3-16](#)

vPC Wizard [2-14, 3-16](#)