Cisco Nexus 5000 and 6000 Series NX-OS FabricPath Operations Guide, Release 6.0(2)N1(1)

March 12, 2013

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-28442-02

Cisco Nexus 5000 and 6000 Series NX-OS FabricPath Operations Guide, Release 6.0(2)N1(1)

Preface 3

Audience 3

Document Conventions 3

Related Documentation 5

Documentation Feedback 5

Obtaining Documentation and Submitting a Service Request 5

5

CHAPTER 1 **Using FabricPath** 1-1

Information About FabricPath 1-1

FabricPath Versus Classical Ethernet Networks 1-2

vPC+ Environment Migration 1-4

FabricPath Link Metrics 1-5

FabricPath Switch IDs 1-7

Conversational MAC Learning 1-7

Guidelines and Limitations of FabricPath 1-9

CE and FabricPath VLANs 1-10

Trees 1-11

Enabling FabricPath 1-12

Verifying the FabricPath Configuration 1-14

Migrating to a vPC+ Environment 1-14

Preface

This preface describes the audience, organization, and conventions of the *Cisco Nexus 5000 and 6000 Series NX-OS FabricPath Operations Guide, Release 6.0(2)N1(1)* . It also provides information on how to obtain related documentation.

This chapter includes the following topics:

- Audience, page 3
- Document Conventions, page 3
- Related Documentation, page 5
- Obtaining Documentation and Submitting a Service Request, page 5

Audience

This publication is for experienced network administrators who configure and maintain Cisco NX-OS on Cisco Nexus 5000 and Cisco Nexus 6000 Series switches.

Document Conventions

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
Italic	Italic text indicates arguments for which the user supplies the values.
[x]	Square brackets enclose an optional element(keyword or argument).
[x | y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x | y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.

[x {y | z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.
variable	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Screen examples use the following conventions::

Convention	Description
screen font	Terminal sessions and information the switch displays are in screen font.
boldface screen font	Information you must enter is in boldface screen font.
italic screen font	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.

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

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

Related Documentation

Documentation for Cisco Nexus 5000 Series Switches and Cisco Nexus 2000 Series Fabric Extenders is available at the following URL:

http://www.cisco.com/en/US/products/ps9670/tsd_products_support_series_home.html

Documentation for the Cisco Nexus 6000 Series Switch is available at the following URL:

http://www.cisco.com/en/US/products/ps12806/tsd_products_support_series_home.html

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to nexus5k-docfeedback@cisco.com or nexus6k-docfeedback@cisco.com. We appreciate your feedback.

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:

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

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

Using FabricPath

This chapter describes how to configure FabricPath on the Cisco Nexus 5500 Series and the Cisco Nexus 6000 Series switches.

This chapter includes the following sections:

- Information About FabricPath, page 1-1
- FabricPath Versus Classical Ethernet Networks, page 1-2
- vPC+ Environment Migration, page 1-4
- FabricPath Link Metrics, page 1-5
- FabricPath Switch IDs, page 1-7
- Conversational MAC Learning, page 1-7
- Guidelines and Limitations of FabricPath, page 1-9
- CE and FabricPath VLANs, page 1-10
- Trees, page 1-11
- Enabling FabricPath, page 1-12
- Verifying the FabricPath Configuration, page 1-14
- Migrating to a vPC+ Environment, page 1-14

Information About FabricPath

Note FabricPath switching is not supported on the Cisco Nexus 5000 Series switches.

FabricPath switching allows multipath networking at the Layer 2 level without using the Spanning Tree Protocol (STP) (see Figure 1-1). 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 run STP with its blocking ports. Instead, you use FabricPath across datacenters, some of which have only Layer 2 connectivity with no need for Layer 3 connectivity and IP configurations.

FabricPath encapsulation facilitates MAC address 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 datacenters at Layer 2, which is useful in disaster recovery operations, and clustering applications such as databases.

Finally, FabricPath is useful in high-performance, low-latency computing. With FabricPath, you use the Layer 2 Intermediate System-to-Intermediate System (IS-IS) protocol for a single control plane that functions for unicast, broadcast, and multicast packets. There is no need to run STP because the domain is purely Layer 2. This FabricPath Layer 2 IS-IS is a separate process from Layer 3 IS-IS.

Figure 1-1 FabricPath Topology Overview

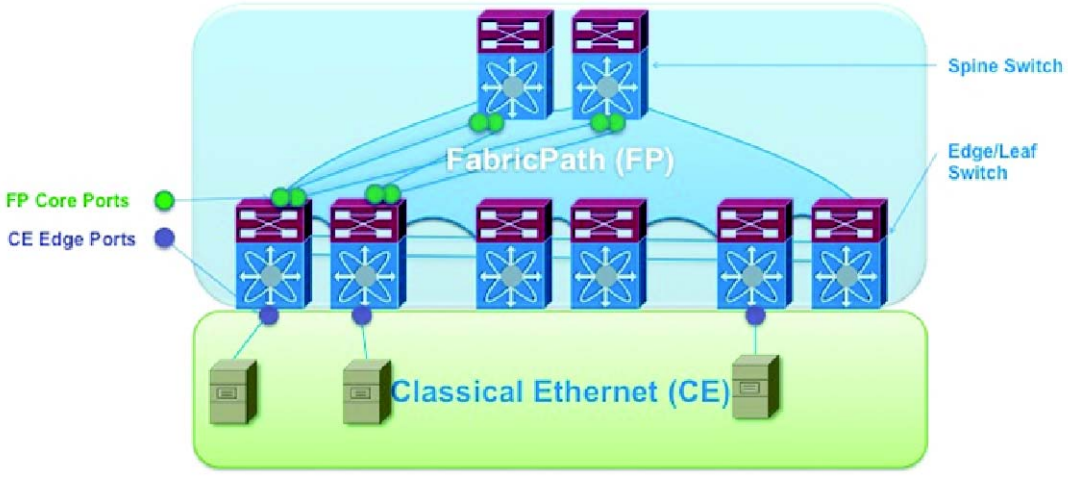

FabricPath Versus Classical Ethernet Networks

FabricPath and CE networks use two different protocols: FabricPath networks use Intermediate System-to-Intermediate System (ISIS) and CE networks use STP to construct a forwarding topology. In both networks, broadcast and unknown unicast traffic are flooded along a loop-free computed graph. However, because two different protocols govern the forwarding graph, a mechanism is needed to control the interaction between the two clouds when FabricPath and CE networks are interconnected forming a physical loop.

STP bridge protocol data units (BPDUs) are not carried through the FabricPath network. CE interfaces continue to run STP and exchange BDPUs (see Figure 1-2).

Layer 2 Gateway Spanning Tree Protocol (L2G-STP) builds a loop-free tree topology. However, it has some limitations. One limitation is that the STP root must always (virtually) be in the FabricPath cloud. For example, it is not possible to have two FabricPath networks connected through a CE cloud. A bridge ID for STP consists of a MAC address and bridge priority. When running in FabricPath mode, the system automatically assigns the edge switches with the MAC address c84c.75fa.6000 from a pool of reserved MAC addresses. As a result, each switch has the same MAC address used for the Bridge ID.

Figure 1-2 CE and FabricPath Example

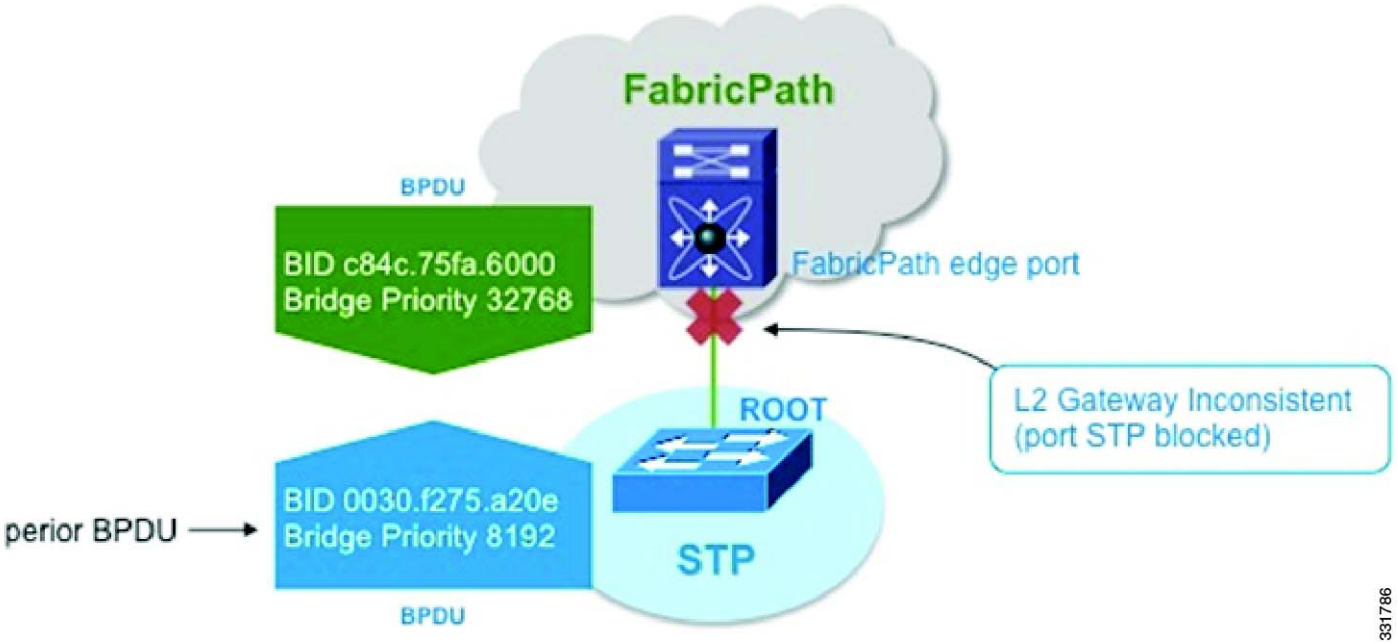

The switches that are in both the FabricPath domain and CE domain are considered to be edge switches or gateway switches. Edge ports have a FabricPath root guard-like function enabled implicitly. If a superior BPDU is received on an edge port, the port is placed in the Layer 2 Gateway inconsistent state until the condition is cleared.

```
%STP-2-L2GW_BACKBONE_BLOCK: L2 Gateway Backbone port inconsistency blocking port
port-channel100 on VLAN0010.
```

As a best practice, you should configure all edge switches with the lowest STP priority of all switches in the STP domain to which they are attached. By setting all of the edge switches to be the root bridge, the entire FabricPath domain looks like one virtual bridge to the CE domain. The same recommendation applies to a virtual port channel+ (vPC+) domain; you must configure each switch (primary and secondary) as the root.

You configure all FabricPath edge switches by manually setting the bridge priority lower than any STP bridge or by entering these commands:

```
sw7-vpc(config)# spanning-tree vlan <x> root primary
sw7-vpc(config)# spanning-tree vlan 1-50 root primary
```

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 switches. The STP domains do not cross into the FabricPath network. If multiple STP domains are defined, BPDUs and topology change notifications (TCNs) are localized to the domain. If a connected STP domain is multihomed to the FabricPath domain, a TCN must be able to reach to all switches in the STP domain through the FabricPath domain. As a result, the TCN is sent to the FabricPath domain through the IS-IS protocol data unit (PDU) by default.

vPC+ Environment Migration

The virtual port channel (vPC) feature was introduced on the Cisco Nexus 5000 Series platforms to provide two active paths, to eliminate the need to run the STP protocol, and to have Active-Active redundancy. vPCs are mainly used for servers that can do port channeling as well as connect to Cisco Nexus 2000 Fabric Extenders. vPCs are deployed in the CE domain. When you migrate to a FabricPath network, switches evolve from a vPC to a vPC+ design.

The vPC+ feature was introduced to allow interoperability between FabricPath and vPCs. The functionality and behavior of a vPC+ and a vPC is identical. The same rules apply in both technologies. That is, both require peer link and peer keepalive messages, the configurations must match between the vPC peers, and consistency checks still take place. In a vPC+ domain, a unique FabricPath switch ID is configured and the peer link is configured as a FabricPath core port. This FabricPath switch ID under the vPC+ domain is called the Emulated switch ID. The Emulated switch ID must be the same between the two peers and must be unique per vPC+.

The benefits of using a vPC+ at the edge of a domain are as follows:

- Allows you to attach servers to the switch using Link Aggregation Control Protocol (LACP) uplinks
- Allows you to attach other CE switches in vPC mode
- Allows you to attach Cisco Nexus 2000 Fabric Extenders in Active/Active mode
- Prevents orphan ports in a failure scenario. When a peer link fails in the vPC+ domain, the orphan port still has FabricPath uplinks for communication.
- Provides numerous paths

Note Moving from an existing vPC to vPC+ is disruptive to the performance of your network during the migration period. We recommend that you schedule a maintenance window when doing this migration.

A switch can be part of a VPC domain or VPC+ domain, but not both domains. When the peer link is a FabricPath core port, all VLANs that traverse the peer link must be FabricPath VLANs (see Figure 1-3).

Figure 1-3 FabricPath Migration Example

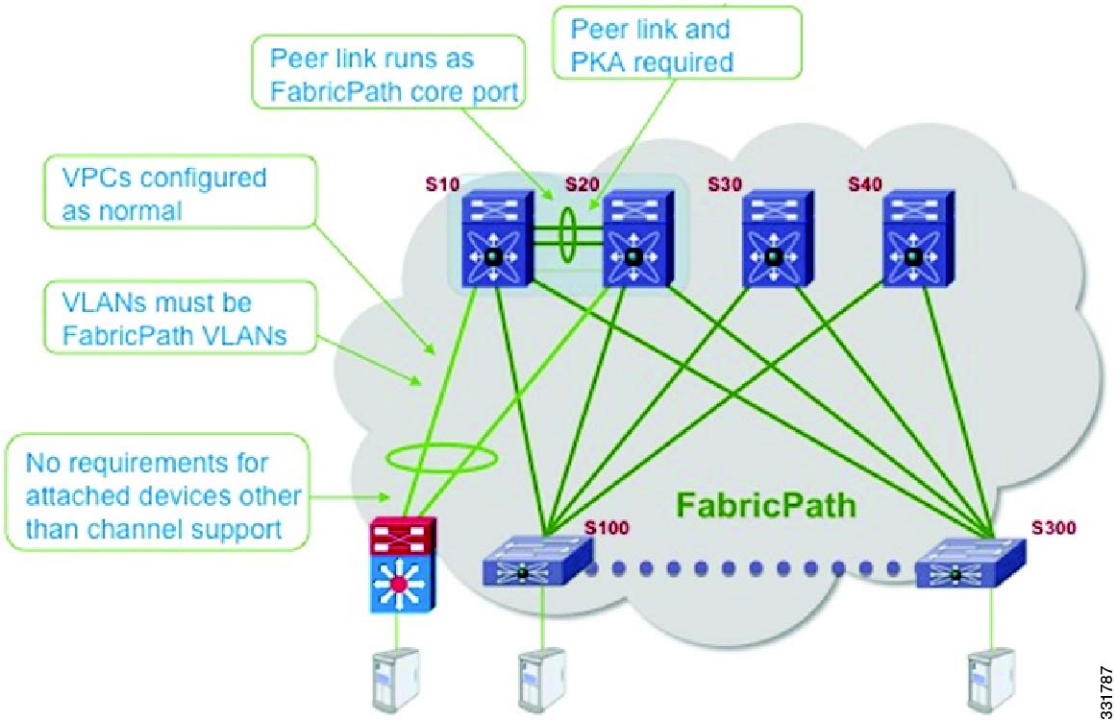

FabricPath Link Metrics

This example shows how to display the FabricPath switch ID table. A switch's switch ID is shown in addition to the emulated switch ID. The same system ID is displayed twice for each switch: one is associated with the switch ID and one is associated with the emulated switch ID.

```
sw7-vpc# show fabricpath switch id

FABRICPATH SWITCH-ID TABLE
Legend: '*' - this system
================================================================
SWITCH-ID       SYSTEM-ID       FLAGS          STATE    STATIC  EMULATED
-------------------------------------+---------------+-----------+-----------+---------
-----------
1            0022.5579.b1c1    Primary    Confirmed      Yes     No
2             0022.5579.b1c2    Primary    Confirmed      Yes     No
3             001b.54c2.7f41    Primary     Confirmed      Yes     No
4             001b.54c2.7f42    Primary     Confirmed      Yes     No
5             0005.73b1.f0c1    Primary     Confirmed      Yes     No
*6           0005.73af.08bc    Primary     Confirmed      Yes     No
7             0005.73b2.0fbc    Primary     Confirmed      Yes     No
8            0005.73af.0ebc    Primary     Confirmed      Yes     No
101        0005.73af.0ebc    Primary     Confirmed      No      Yes
101         0005.73b2.0fbc    Primary     Confirmed      No      Yes
```

After the peer link has been configured with switchport mode fabricpath, it becomes part of the FabricPath topology (see Figure 1-4). After a link is detected in a FabricPath topology, the metric of the link is verified and used towards the unicast routing table and the calculation of trees for multidestination

traffic. To take advantage of the equal cost multipath (ECMP) paths that are available from edge to spine in the FabricPath topology, we recommend that you increase the IS-IS metric of the peer link to be lower so that it is not added as part of the multidestination tree.

The preferred path to any switch ID is calculated based on the metric to any given destination. The metric is as follows:

- 1-Gbps links have a cost of 400.
- 10-Gigabit links have a cost of 40.
- 20-Gbps have a cost of 20.

Figure 1-4 FabricPath Preferred Paths

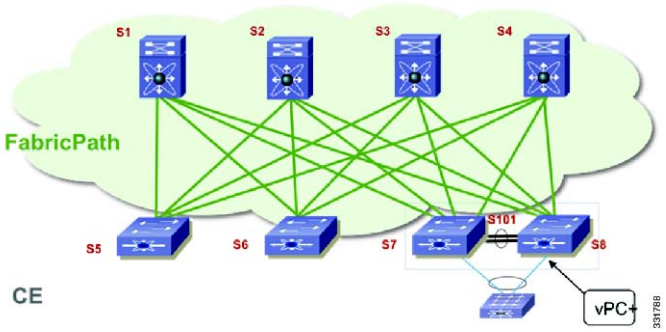

You must verify the metric of the links.

This example shows how to display the FabricPath interface information:

```
sw7-vpc# show fabricpath isis interface brief
Fabricpath IS-IS domain: default
Interface    Type  Idx State        Circuit   MTU  Metric  Priority  Adjs/AdjsUp
--------------------------------------------------------------------------------
port-channel1 P2P 2     Up/Ready   0x01/L1   1500 20      64          1/1
Ethernet1/7  P2P   4     Up/Ready   0x01/L1   1500 40      64          1/1
Ethernet1/8  P2P   1     Up/Ready   0x01/L1   1500 40      64          1/1
Ethernet1/9  P2P   3     Up/Ready   0x01/L1   1500 40      64          1/1
```

Because the peer link is a port channel, the metric will be the lowest cost. As a best practice, you should increase the metric so it is higher than the rest of the ECMP links that are part of the FabricPath cloud.

This example shows how to display the FabricPath metric:

```
sw7-vpc(config-if)# fabricpath isis metric 100
sw7-vpc(config-if)# show fabricpath isis interface brief
Fabricpath IS-IS domain: default
Interface    Type  Idx State        Circuit   MTU  Metric  Priority  Adjs/AdjsUp
--------------------------------------------------------------------------------
port-channel1 P2P 2 Up/Ready   0x01/L1   1500 100     64          1/1
Ethernet1/7  P2P   4  Up/Ready   0x01/L1   1500 40      64          1/1
Ethernet1/8  P2P   1  Up/Ready   0x01/L1   1500 40      64          1/1
Ethernet1/9  P2P   3  Up/Ready   0x01/L1   1500 40      64          1/1
```

FabricPath Switch IDs

When FabricPath is enabled globally, each switch is atomatically assigned with a switch ID (12 bits) You have the option to manually configure the switch ID, but you must ensure all switches in the FabricPath domain have unique values. The switch ID is encoded in the outer MAC addresses of the FabricPath MAC-in-MAC frames.

The Dynamic Resource Allocation Protocol (DRAP) automatically assigns a switch ID and ensures that no duplicate IDs exist in the FabricPath domain. The FabricPath network automatically detects conflicting switch IDs and prevents the data path initialization on the FabricPath interface. As a best practice, we recommend that you manually set the switch IDs.

The emulated switch ID is used in a VPC+ to identify the VPC+ bundle. The emulated switch ID must be unique within each VPC+ virtual switch domain. In a vPC+ domain, three switch ID's will be used—one unique switch ID for each vPC peer and one emulated switch ID that is common between both vPC peers.

This example shows how to display and manually configure the switch ID:

```
sw5# show fabricpath switch-id
                       FABRICPATH SWITCH-ID TABLE
Legend: '*' - this system
==============================================================
========
SWITCH-ID      SYSTEM-ID       FLAGS         STATE    STATIC  EMULATED
----------+---------------+-----------+-----------+--------------------
*3428   0005.73b1.f0c1Primary     Confirmed    No      No
```

Conversational MAC Learning

In conversational MAC learning the switch learns MAC addresses only if there is an active conversation between a local MAC and a remote MAC (bidirectional traffic) (see Figure 1-5). By default, conversational MAC learning is enabled for all FabricPath VLANs. All CE VLANs learn MAC addresses the traditional (CE) way. The default MAC address aging timers for both CE and FabricPath VLANs is 300 seconds on the Cisco Nexus 5500 Series switches. The default MAC address aging timers are 1800 seconds on the Cisco Nexus 6000 Series and Cisco Nexus 7000 Series switches.

If you have a mix of Cisco Nexus platforms in the FabricPath environment, set the MAC address aging timer on the Cisco Nexus 5500 switch to 1800 seconds to match the value on the Cisco Nexus 6000 and Cisco Nexus 7000 Series switches. This will avoid any unnecessary flooding.

By default, if Layer 3 is enabled, the ARP aging timer is 1500 seconds. When Layer 3 is enabled, you should set the MAC address aging timer to a higher value than the ARP table to avoid any unnecessary flooding.

Figure 1-5 MAC Address Learning with FabricPath

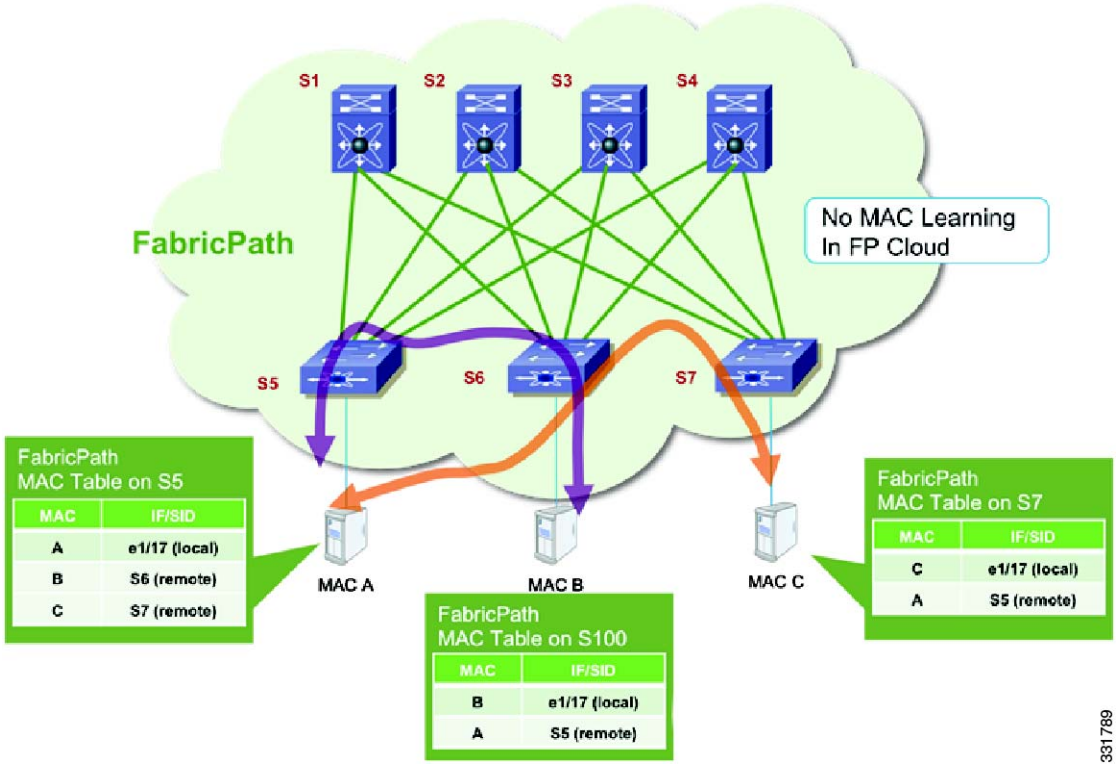

Note When you enable a switch virtual interface (SVI) (regardless if SVI is used for management or routing purposes), conversational MAC learning is disabled for that particular VLAN. As a result, when you enable the Hot Standby Router Protocol (HSRP) in a vPC+ environment, conversational MAC learning is disabled for that VLAN. Conversational MAC learning is only disabled on the particular VLAN on the Cisco Nexus 5500 Series or Cisco Nexus 6000 Series switches that terminates the SVI.

This example shows how to display the dynamic MAC address table for the S5 switch:

```
S5# show mac address-table dynamic
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
5         0000.0000.000c    dynamic   0          F    F  1:0:7
5         0000.0000.000a    dynamic   0          F    F  Eth1/17
5         0000.0000.000b    dynamic   10         F    F  1:0:6
```

This example shows how to display the dynamic MAC address table for the S6 switch:

```
S6# show mac address-table dynamic
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
5      0000.0000.000a    dynamic   0          F    F  1:0:5
5      0000.0000.000b    dynamic   0          F    F  Eth1/17
```

This example shows how to display the dynamic MAC address table for the S7 switch:

```
S7# show mac address-table dynamic
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
5     0000.0000.000c    dynamic   0          F    F  Eth1/17
5     0000.0000.000a    dynamic   0          F    F  1:0:5
```

Guidelines and Limitations of FabricPath

FabricPath has the following configuration guidelines and limitations:

- An in-service software upgrade (ISSU) is supported when the Cisco Nexus 5500 Series and Cisco Nexus 6000 switch are in Layer 2 mode running FabricPath and/or a vPC. The edge switch that is undergoing and ISSU is in both the CE and FabricPath cloud. The same rules that are applicable in an ISSU environment on the Cisco Nexus 5500 Series and the Cisco Nexus 6000 Series switches are applied in a FabricPath design.
- The spanning-tree configuration cannot have any designated port, with the exception of ports that are configured as a spanning-tree port-type edge with bridge protocol data unit (BPDU) Filtering. Bridge assurance must be disabled on all ports except on the peer link where it is possible to keep bridge assurance enabled. Bridge assurance operates only when a port is configured as a spanning-tree port type network. Ports that are configured as default or normal ports do not run bridge assurance.
- The Cisco Nexus 5500 Series and the Cisco Nexus 6000 Series switches cannot be the STP root bridge or have any designated nonedge ports in the STP topology.
- The Cisco Nexus 5500 Series switches, the Cisco Nexus 6000 Series switches and the Cisco Nexus 2000 Fabric Extenders that are undergoing an ISSU must be a leaf on the spanning tree.

Note The term “leaf” refers to a switch that connects servers in a data center fabric and the term “spine” refers to a switch that connects the leaf switches.

- The CE and FabricPath topology should be in a stable state before undergoing an ISSU. In the FabricPath cloud, no additional switches, links or switch IDs should be added or removed during an ISSU. During an ISSU process, there should not be any broadcast or multicast root change.
- When the switch undergoes an ISSU, it takes about 80 seconds for the control plane to restart. During this time, the switch that is undergoing an ISSU increases its ISSU timer to 100 seconds and informs its neighbors by sending IS-IS hellos. The timer is increased only between the switch that is undergoing the ISSU and the neighbors that are directly connected to it. After the switch completes the ISSU, the default timer starts sending out IS-IS hello times again.

For a complete list of ISSU guidelines, see the *Cisco Nexus 5000 Series NX-OS Software Upgrade and Downgrade Guide*.

CE and FabricPath VLANs

The CE VLANs carry traffic from the CE hosts to the FabricPath interfaces, and the FabricPath VLANs carry traffic throughout the FabricPath topology. Only the active FabricPath VLANs that are configured on a switch are advertised as part of the topology in the Layer 2 IS-IS messages. The switch automatically assigns all FabricPath interfaces and FabricPath VLANs to the default topology (Topology 0). Therefore, no additional configuration is required. If a VLAN is a CE VLAN only, it cannot traverse the FabricPath cloud. In order for traffic to transit across the FabricPath cloud, you must specify the VLAN as a FabricPath VLAN.

When configuring a port to be a FabricPath port, enter the **switchport mode fabricpath** command to put the interface in FabricPath mode and forward all FabricPath VLANs. With FabricPath, you do not need to enter the **switchport trunk allowed vlan** command. All of the VLANs that are defined you entered the **mode fabricpath** command are automatically carried on interfaces when you entered the **switchport mode fabricpath** command. Because all FabricPath VLANs forward on a FabricPath port, you do not need to use the **switchport trunk allow vlan x** command.

You must exit the VLAN configuration mode for the VLAN mode change to take effect. When running a vPC+, the peer link is configured as a FabricPath core interface. To forward the VLANs downstream on a vPC, you must configure them as FabricPath VLANs and they must be reachable across the peer link.

This example shows how to display the configuration of the FabricPath VLANs:

```
sw7-vpc# show vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                  : 100
vPC+ switch id                  : 101
Peer status                     : peer adjacency formed ok
vPC keep-alive status           : peer is alive
vPC fabricpath status           : peer is reachable through fabricpath
Configuration consistency status: success
Per-vlan consistency status     : success
Type-2 consistency status       : success
vPC role                        : secondary, operational primary
Number of vPCs configured       : 2
Peer Gateway                    : Disabled
Dual-active excluded VLANs      : -
Graceful Consistency Check      : Enabled

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po1    up     5

vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason       Active vlans vPC+ Attrib
--     ----------  ------ ----------- ------       ------------ -----------
111    Po111       up     success     success      -            DF: Partial
```

Trees

An Ethernet domain, always have two kinds of traffic, unicast traffic and multidestination traffic. In a FabricPath topology, unicast traffic occurs when the traffic is sent to another host (1:1), and the source and destination are known. For unicast traffic, FabricPath uses the routing table to identify the next hop. If there is one best hop, the protocol chooses the individual link. If there is equal cost multipathing (ECMP), the unicast traffic is load balanced across the core interfaces.

In the current release of FabricPath, up to 16 equal cost paths are possible. The default load-balancing scheme for ECMP is a mixed mode (Layer 2, Layer 3 and Layer 4 ports).

This example shows how to display the FabricPath load-balancing configuration:

```
sw5# show fabricpath load-balance
ECMP load-balancing configuration:
L3/L4 Preference: Mixed (L2, L3 and L4)
Hash Control: Symmetric
Use VLAN: TRUE
sw5#
```

You can modify the variables by using the **fabricpath load-balance unicast** command.

This example shows how to configure all of the arguments that are available with this command:

```
sw5(config)# fabricpath load-balance?
destination         Include destination parameters
source              Include source parameters
source-destination  Include source and destination parameters
unicast            Unicast

sw5(config)# fabricpath load-balance unicast ?
  <CR>
  include-vlan  Use vlan
  layer2        Layer-2 parameters considered
  layer3        Only Layer-3 parameters considered
  layer4        Only Layer-4 parameters considered
  mixed         Mix of Layer-2, Layer-3 and Layer-4 paramaters (default)

sw5(config)# fabricpath load-balance unicast
```

FabricPath introduces a loop-free broadcast functionality that carries broadcast, unknown unicast, and multicast packets (also known as multidestination traffic). For each broadcast, unknown unicast, and multicast traffic flow, the switch chooses the forwarding path from multiple system-created paths or trees.

The switch creates two trees to forward the multidestination traffic for each topology. Each tree is identified in the FabricPath network by a unique value or FTag. For the FabricPath network, the switch creates Tree 1 (FTag1) that carries broadcast traffic, unknown unicast traffic, and multicast traffic through the FabricPath network. The swtich also creates a second tree, Tree 2 (FTag 2). All of the multicast traffic flows are load balanced across these two trees for each flow.

Within the FabricPath network, the swtich elects a root node that becomes the root for the broadcast tree. That node also identifies another bridge to become the root for the second multidestination tree, which load balances the multicast traffic. In a unicast-only environment, Tree 1 or FTag1 is always used and is seen in all **show** commands.

The FabricPath network elects a single root switch for the first (broadcast) multidestination tree in the topology. All FabricPath switches announce their root priority in the Router Capability TLV. The switch with the highest priority value becomes the root for the tree. In the event of a tie, the FabricPath network chooses the switch with the highest system ID, and if there is still a tie, then it uses the switch with the

highest switch ID. The broadcast root determines the roots of any additional multicast trees and announces them in the Router Capability TLV. Multicast roots spread among available switches to balance the load. The selection is based on the same criteria as described above.

As a best practice, we recommend that you manually define the spine switches that are the root of each tree.

This example shows how to configure two spine switches at the root of two trees:

```
Spine 1:
fabricpath domain default
  root-priority 255

Spine 2:
fabricpath domain default
  root-priority 254
```

This example shows how to display the multidestination trees for ftag 1 and 2:

```
spine# show fabricpath isis topology summary
Fabricpath IS-IS domain: default FabricPath IS-IS Topology Summary
MT-0
  Configured interfaces:  Ethernet7/1  Ethernet7/2  Ethernet7/3  Ethernet7/4
  Number of trees: 2
    Tree id: 1, ftag: 1, root system: 0022.5579.b1c1, 1
    Tree id: 2, ftag: 2, root system: 0022.5579.b1c2, 2

spine# show fabricpath isis trees multidestination 1
Fabricpath IS-IS domain: default
Note: The metric mentioned for multidestination tree is from the root of that tree to that
switch-id

MT-0
Topology 0, Tree 1, Swid routing table
2, L1
 via Ethernet7/4, metric 40
3, L1
 via Ethernet7/1, metric 80
4, L1
 via Ethernet7/1, metric 80
5, L1
 via Ethernet7/2, metric 40
6, L1
 via Ethernet7/1, metric 40
7, L1
 via Ethernet7/3, metric 40
8, L1
 via Ethernet7/3, metric 60
101, L1
 via Ethernet7/3, metric 60
```

Enabling FabricPath

BEFORE YOU BEGIN

- Make sure the appropriate Cisco Nexus 5550 Series or Cisco Nexus 6000 Series switch is used. Only the Cisco Nexus 5500 Series and the Cisco Nexus 6000 platforms support FabricPath; the first generation Cisco Nexus 5000 Series switches do not support FabricPath.
- Download the correct version of Cisco NX-OS software.

- Obtain the Enhanced Layer 2 License.

PROCEDURE

Step 1 Install the Enhanced Layer 2 License.

```
sw7-vpc(config)# install license bootflash:///enhanced_layer2_pkg.lic
```

Step 2 Install the FabricPath feature-set.

```
sw7-vpc(config)# install feature-set fabricpath
```

Step 3 Enable the FabricPath feature-set.

Note Step 2 and Step 3 are independent steps and must be done to enable FabricPath successfully.

```
sw5(config)# vlan 5-20
sw5(config-vlan)# mode fabricpath
sw5(config-vlan)# exit
```

Step 4 Enable FabricPath mode on the DCE core ports connected to the spine switches.

```
sw5(config)# int ether 1/8-10
sw5(config-if-range)# switchport mode fabricpath
```

Figure 1-6 FabricPath Topology Example

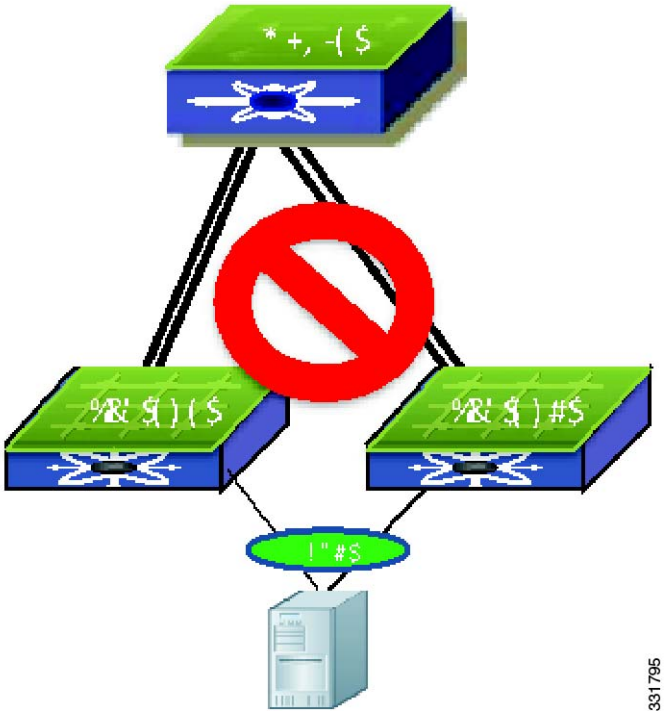

Verifying the FabricPath Configuration

PROCEDURE

Step 1 Verify that FabricPath is enabled on the core interfaces by using the **show vlan id** command.

Step 2 Verify that the VLAN is in FabricPath mode.

```
sw5# show vlan id 5

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
5    VLAN0005                         active    Po101, Eth1/5, Eth1/6, Eth1/7
                                                Eth1/8, Eth1/9
VLAN Type  Vlan-mode
---- ----- ----------
5    enet  FABRICPATH
```

Step 3 Verify that the ISIS adjacency is enabled on FabricPath ports.

```
sw5# show fabricpath isis adjacency
Fabricpath IS-IS domain: default Fabricpath IS-IS adjacency database:
System ID       SNPA            Level  State  Hold Time  Interface
vveerapp-7k3    N/A             1      UP     00:00:31   Ethernet1/7
vveerapp-7k4    N/A             1      UP     00:00:26   Ethernet1/8
vveerapp-7k1    N/A             1      UP     00:00:31   Ethernet1/9
vveerapp-7k2    N/A             1      UP     00:00:32   Ethernet1/10
```

Migrating to a vPC+ Environment

BEFORE YOU BEGIN

Install the FabricPath feature.

PROCEDURE

Step 1 Create a Switch ID under the vPC domain

```
sw5# vpc domain 100
  peer-keepalive destination 172.25.204.86 source 172.25.204.85
  fabricpath switch-id 101
```

Note

- When you configure the peer keepalive link for vPC+, we recommend that you use the mgmt0 interface and management virtual routing and forwarding (VRF) instance. If a dedicated port is used, we recommend that you configure the front panel ports as part of a dedicated VRF (that is, when Layer 3 is enabled on the Cisco Nexus 5500 Series switch) with a dedicated CE VLAN.
- The vPC+ keepalive messages must be transported between the primary vPC+ and secondary vPC+ across a dedicated link instead of going through the FabricPath cloud. In addition, if the Layer 3 functionality is enabled on the Cisco Nexus 5500 Series switch, the keyword management must be configured under the corresponding SVI to allow the SVI to remain up if the Layer 3 module fails.

C

Classical Ethernet
versus FabricPath 1-2
classical Ethernet VLANs 1-9

E

enabling
FabricPath 1-12

F

FabricPath
enabling 1-12
information about 1-1
ISSU 1-9
link metrics 1-5
MAC learning 1-7
switch ID 1-7
trees 1-11
versus Classical Ethernet 1-2
VLANs 1-9
FabricPath configuration
verifying 1-13, 1-14

I

information about
FabricPath 1-1
ISSU
FabricPath 1-9

M

MAC learning
FabricPath 1-7
metrics
FabricPath links 1-5
migrating
vPC+ environment 1-14
migration
vPC+ environment 1-4

S

switch IDs
FabricPath 1-7

T

trees
FabricPath 1-11

V

verifying
FabricPath configuration 1-13, 1-14
VLANs
classical Ethernet 1-9
FabricPath 1-9
vPC+ environment migration 1-4
vPC environment
migrating 1-14