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 5500 Series switches.

This chapter includes the following sections:

Information About FabricPath

The FabricPath feature provides the following:

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

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

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

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 from 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 switch. The frames carry the same FTag that is assigned at ingress throughout the FabricPath network, and Layer 2 IS-IS allows all switches to have the same view of all the trees built by the software.

For multi-direction traffic, such as unknown unicast, multicast, and broadcast, the tree path is used only to avoid packet looping. FTag ensures that all switches forward packets on the same tree across the network, although FabricPath supports multiple trees. Known unicast packets are forwarded throughout the network using equal cost multipath (ECMP) and FTag trees are not used for these packets. Using ECMP and the trees, the software automatically load balances traffic throughout the FabricPath network.

Information About Conversational MAC Address Learning

Beginning with Cisco NX-OS Release 5.1(3)N1(1), you can use conversational MAC address learning or traditional MAC address learning.

To use conversational MAC address learning, you must do the following:

- Enable FabricPath.
- Ensure VLANs do not have switch virtual interface (SVI) enabled.

Conversational MAC address learning means that each switch learns only those MAC addresses for interested hosts, rather than all MAC addresses in the domain. Each switch learns only those MAC addresses that are actively speaking with it. 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.

Note CE VLANs support only traditional MAC address learning, where each switch learns the MAC addresses of all hosts in the network.

Layer 3 Routing Considerations for FabricPath

This section provides information you need to consider when you implement routing technologies for Layer 3 routing that take advantage of FabricPath to improve both unicast and multicast routing in your data center.

This section includes the following topics:

In addition to the above, FabricPath includes support for the **delay restore interface-vlan** command and the **ip arp synchronize** command. For more information about these commands, see the *Cisco Nexus 5500 Series NX-OS vPC Command Reference*.

Hot Standby Router Protocol Support

You can configure Hot Standby Router Protocol (HSRP) in Active/Active mode between pairs of switches in the vPC+ domains and access layer, and between pairs of switches in the aggregation layer.

If you configure a peer link between a pair of switches in the aggregation layer, HSRP runs in Active/Active mode for those switches. If you do not configure a peer link between the switches in the aggregation layer, HSRP runs in Active/Standby mode for those switches.

Note If a data center outage occurs and you enable HSRP before the vPC+ successfully comes up, traffic loss can occur. You must 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 5500 Series switches allows Layer 2 switching only after both timers expire.

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

No Support for the bind-vrf Command

One consequence of the single-DR implementation for FabricPath is that you do not need to bind a virtual routing and forwarding (VRF) instance to a vPC. As a result, the **vpc bind-vrf** command is not supported for FabricPath.

If you want to configure an existing vPC domain for FabricPath, you must first use the **no vpc bind-vrf** command to remove the static binding between the vPC and the VRF.

Peer Gateway Not Recommended on Aggregation Layer

We do not recommend that you configure a peer gateway on the aggregation layer, as Fabric Path does not require one.

The **peer-gateway** command is supported by vPC and vPC+ at the access. If traffic from a vPC for peer-1's MAC address is sent to peer-2, the peer gateway configuration ensures that peer-2 can route the packet on peer-1's behalf. However, a FabricPath topology does not have a port channel that can cause traffic for one peer switch to be sent to the other peer switch. Therefore, FabricPath does not require peer gateway configuration, even though there is a peer link provisioned for HSRP Active-Active support.

Prerequisites for FabricPath

FabricPath has the following prerequisites:

- You should have a working knowledge of Classical Ethernet Layer 2 functioning.
- You must be logged onto the switch.
- You must ensure an Enhanced Layer 2 license is installed on the switch.
- If you want to configure an existing vPC domain for FabricPath, you must first use the **no vpc bind-vrf** command to remove the static binding between the vPC and the VRF.

Guidelines and Limitations for FabricPath

FabricPath has the following configuration guidelines and limitations:

- FabricPath switches and interfaces carry only FabricPath-encapsulated traffic.
- You must install and enable FabricPath on each switch before you can view or access the commands. See the "Configuring FabricPath Switching" section on page 3-6 for more information.
- STP does not run inside a FabricPath network.

- 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 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 software 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.

- FabricPath 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.

- FabricPath does not support VTP when in the same VDC. You must disable VTP when the FabricPath feature set is enabled on the VDC.

Note For information about FabricPath and DHCP snooping, including guidelines and limitations, see the *Cisco Nexus 5500 Series NX-OS Security Configuration Guide, Release 6.0*.

Licensing Requirements for FabricPath

FabricPath requires an Enhanced Layer 2 license. You must install this license on every switch in a FabricPath network.

In addition, if you want to take advantage of the Layer 3 routing interactions with FabricPath, you must install a LAN Base Services license on all switches with a Layer 3 card.