CHAPTER 33

Configuring Layer 2 VPN VPLS Dual-Homing with a vPC

This chapter describes how to configure dual-homing with a virtual port channel (vPC) to integrate Virtual Private LAN (VPLS) with the vPC functionality in active-standby mode and allow traffic from a customer edge (CE) device to be load balanced across both provider edge (PE) devices.

This chapter includes the following sections:

Finding Feature Information

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

Information about Layer 2 VPN VPLS Dual-Homing with a vPC

This section includes the following topics:

VPLS Integration with vPC

Virtual Private LAN Service (VPLS) provides a multipoint-to-multipoint Layer 2 service over a wide area network (WAN). VPLS is implemented by connecting all the nodes in a particular domain by using a full mesh of pseudowires (PWs).

The virtual port channel (vPC) functionality provides multichassis Ether channel support. Both the attachment circuit (AC) links in a vPC domain are in active mode, which increases the throughput of the network because all the interswitch links can be used to carry traffic.

In the VPLS integration with a vPC, a customer edge (CE) device is dual-homed to two provider edge (PE) devices. The PEs are part of the VPLS domain. One of the PEs in the VPLS domain is in Active state and forwards traffic, while the other PE is in Standby state. See the figure below.

Figure 33-1 VPLS Integration with a vPC

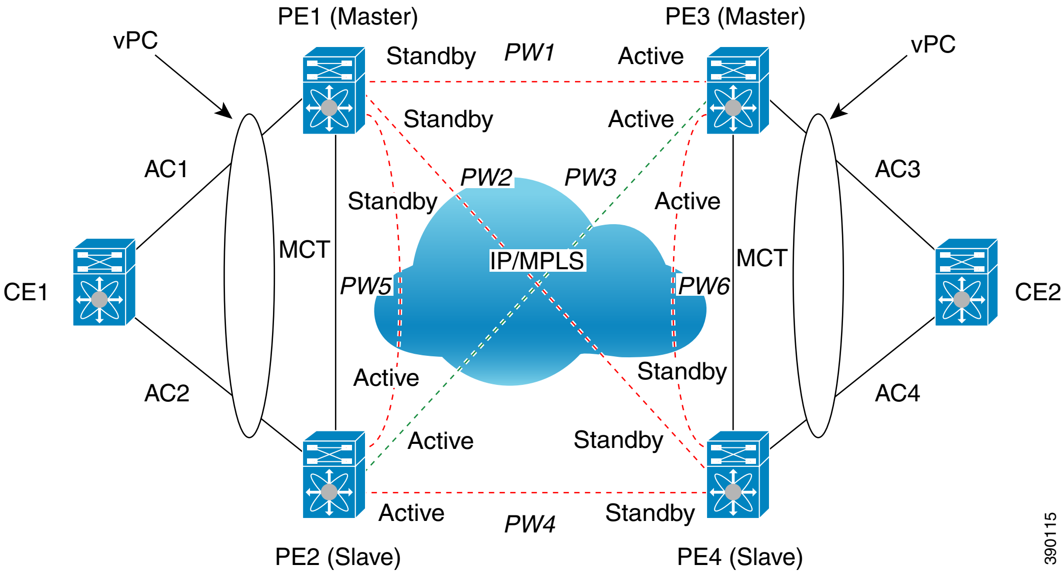

PE1 and PE2 belong to a vPC domain. PE3 and PE4 are part of another vPC domain. In addition to being part of their respective vPC domains, VPLS is configured on the PEs using a mesh of PWs.

In the above figure, the virtual forwarding instance (VFI) configured under a particular VLAN in PE2 is Active for vPC group (PE1, PE2) and VFI configured under a particular VLAN in PE3 is Active for vPC group (PE3, PE4). The Active VFI advertises the local status of Active on all the PWs. The Standby VFI advertises the local status of Standby on all the PWs. A PW is Active when both ends advertise the status of Active. Therefore, PW3 is Active between PE2 and PE3 and is used to carry traffic between CE1 and CE2.

The VPLS domain is configured in decoupled mode. As a result, the status of the AC links is not advertised to the PWs.

Overview of a vPC Peer Link

You can have only two switches as vPC peers; each switch can serve as a vPC peer to only one other vPC peer. The vPC peer switches can also have non-vPC links to other switches.

To create a valid configuration, you must configure an EtherChannel on each switch and then configure the vPC domain. You must assign the EtherChannel on each switch as a peer link. For redundancy, we recommend that you configure at least two dedicated ports into the EtherChannel; if one of the interfaces in the vPC peer link fails, the switch automatically falls back to use another interface in the peer link.

Note The two switches connected by the vPC peer link have certain identical operational and configuration parameters.

When you configure the vPC peer link, the vPC peer switches negotiate that one of the connected switches is the primary switch and the other connected switch is the secondary switch. By default, Cisco NX-OS software uses the lowest MAC address to elect the primary switch. The software takes different actions on each switch only in certain failover conditions. If the primary switch fails, the secondary switch becomes the operational primary switch when the system recovers, and the previously primary switch becomes the secondary switch.

MAC addresses that are learned over vPC links are synchronized between the peers. Configuration information flows across the vPC peer link using the Cisco Fabric Services over Ethernet (CFSoE) protocol. All MAC addresses for VLANs configured on both switches are synchronized between vPC peer switches. The software uses CFSoE for this synchronization.

If the vPC peer link fails, the software checks the status of the remote vPC peer switch by using the peer-keepalive link, which is a link between vPC peer switches, to ensure that both switches are up. If the vPC peer switch is up, the secondary vPC switch disables all vPC ports on its switch. The data then forwards the remaining active links of the EtherChannel. The software learns of a vPC peer switch failure when the keepalive messages are not returned over the peer-keepalive link. Use a separate link (vPC peer-keepalive link) to send configurable keepalive messages between the vPC peer switches. The keepalive messages on the vPC peer-keepalive link determines whether a failure is on the vPC peer link only or on the vPC peer switch. The keepalive messages are used only when all the links in the peer link fail.

Validating the Configuration Between Switches

Cisco NX-OS software validates the configuration between primary (master) and secondary (slave) switches. When a Type-1 mismatch occurs between the primary and secondary switches for a particular VLAN, the VLAN is suspended on both the switches. When a consistency check fails, only the secondary virtual port channel (vPC) switch is brought down. The VLAN remains up on the primary switch and Type-1 configurations can be performed without traffic disruption.

The virtual forwarding instance (VFI) on the vPC switch can be configured as primary or secondary, independent of the vPC state (master or slave) on the switch. Similarly, the VFI on the other vPC switch can be configured as primary or secondary; just not the same as the other vPC peer. If both the vPC peers are configured as primary or secondary or if no primary or secondary vPC peer is configured, a Type-1 error occurs. The table below summarizes the configurations to be avoided.

vPC Peer 1	vPC Peer 2
Primary	Primary
Secondary	Secondary
Primary	—
Secondary	—

Port, Link, and Node Failures

The VPLS Active-Standby Support with a vPC feature provides network resiliency by protecting against port, link, and node failures. These failures can be categorized into the following scenarios:

- Scenario A: Failure of the uplink port on the dual-homed device (DHD)
- Scenario B: Failure of the uplink (AC) of the DHD
- Scenario C: Failure of the port on a vPC peer
- Scenario D: Failure of the primary node in vPC
- Scenario E: Failure of the vPC peer link (MCT)
- Scenario F: Failure of the vPC node uplink towards the MPLS core

These failure points are shown in the figure below.

Figure 33-2 Port, Link, and Node Failures

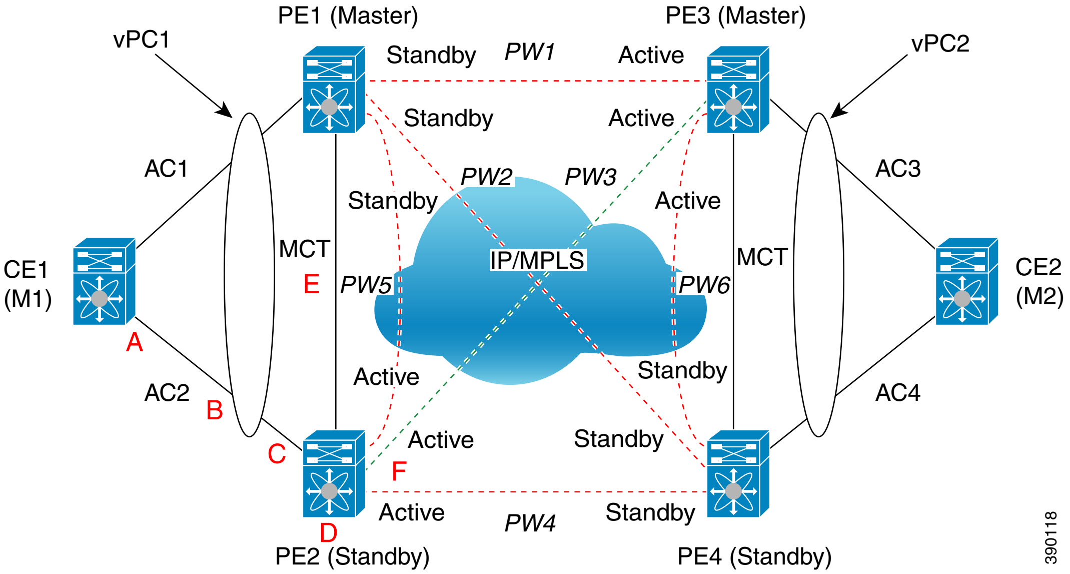

Failed Port or AC Link

The figure below shows the frame flow if a port on the DHD or vPC peer or the attachment circuit (AC) link fails (scenario A, B, or C).

Figure 33-3 Failure Scenario A, B, or C

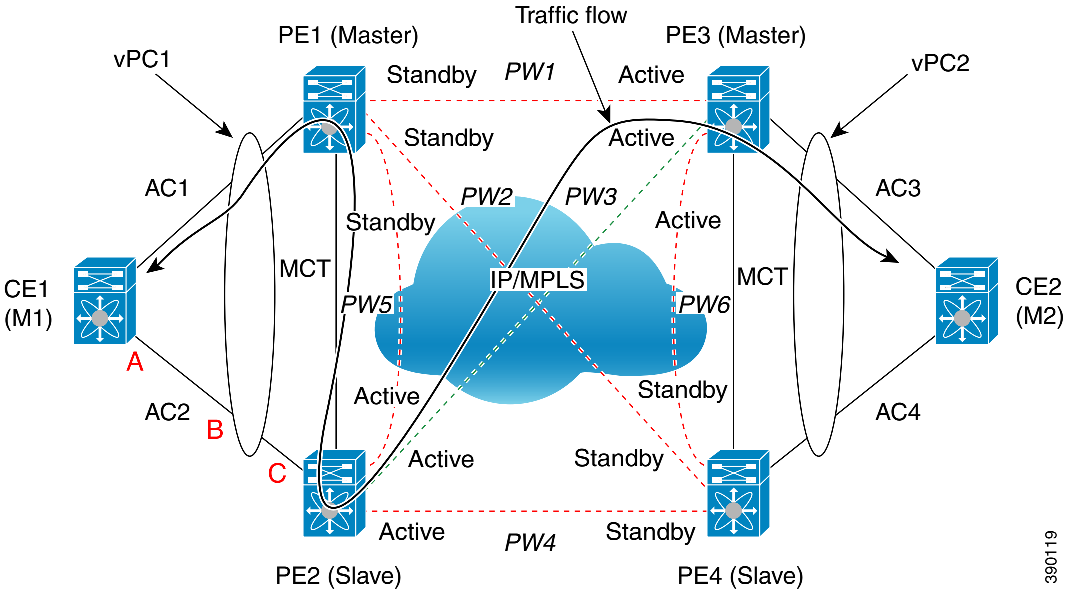

For A, B, or C failures, a vPC diverts the traffic destined from PE2 to AC2 toward multichassis trunk (MCT). In case of A, B, or C failures, PE2 sends a message to PE1 to forward the traffic that is received over MCT to vPC links in addition to sending it over non-vPC links. The traffic is forwarded on AC1 and is received or sent by CE1.

In all of these failure scenarios, because VPLS is configured in decoupled mode, PE2 continues to advertise the local status of Active on all its PWs, and PE1 continues to advertise the local status of Standby.

Failed Primary Node

The figure below shows the switch configured as the primary node fails (scenario D).

Figure 33-4 Failure Scenario D

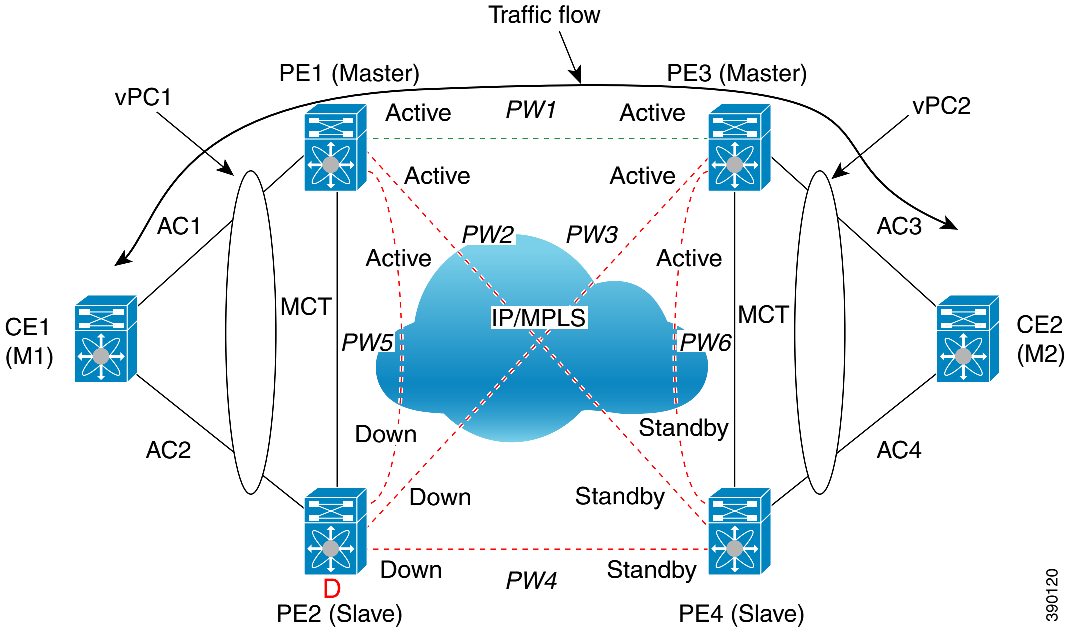

If the primary node in the vPC fails, the peer vPC node detects the failure and starts advertising its local status of Active on all core PWs. PW1 between PE1 and PE3 becomes the Active PW and traffic between vPC1 and vPC2 flows over PW1. If PE3 detects that PW3 to PE2 is down before PE1 has advertised the local status of Active, traffic is dropped by PE3 toward the core (that was using PW3) until PE3 receives the status of Active from PE1.

Failed vPC Peer Link

The figure below shows scenario E or the failure of vPC peer link, also known as multichassis trunk (MCT).

Figure 33-5 Failure Scenario E

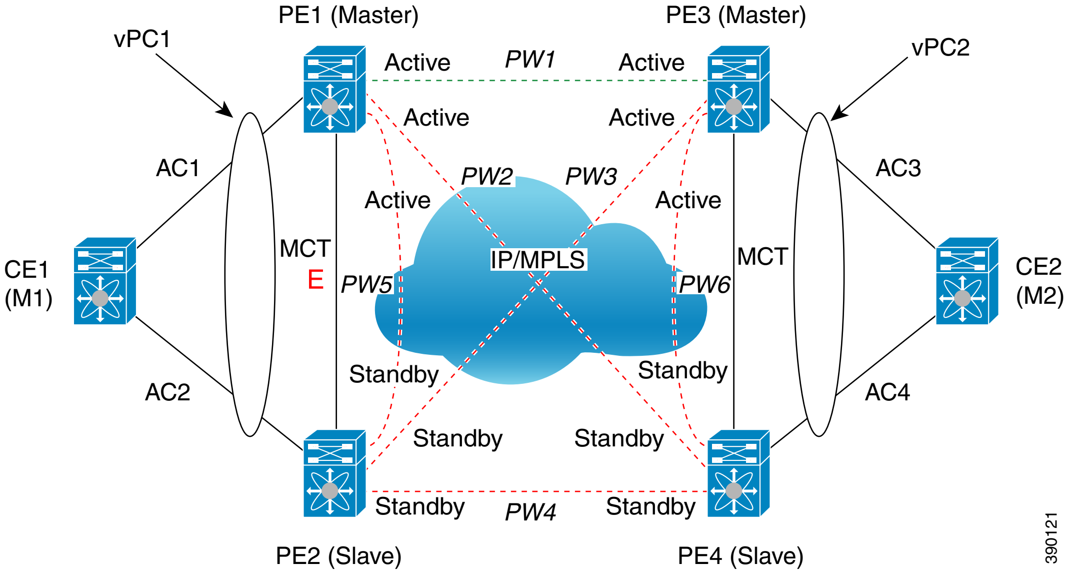

If MCT fails and both the nodes in the virtual port channel (vPC) are still up, the vPC master node keeps the AC link up and the vPC slave node brings its AC link down.

Because VPLS is configured in decoupled mode (the status of the AC link is not advertised to the core pseudowires), the Standby PE2 advertises the local status of Standby to all the core pseudowires (PWs) even though AC2 is down. Therefore, the traffic between vPC1 and vPC2 flows over PW1. The traffic between CE1 and PE1 flows only through the AC1 link.

You can configure vPC to detect a double fault if both MCT and the primary vPC node fail. You can configure vPC by using the out-of-band keepalive mechanism. In this scenario, the secondary vPC node keeps the AC link (AC) up and PE2 continues to advertise the status of Active to all core PWs. PW3 remains active.

Failed Core Pseudowires

The figure below shows the failure of core pseudowires (PWs) on the Active vPC node (scenario F).

Figure 33-6 Failure Scenario F

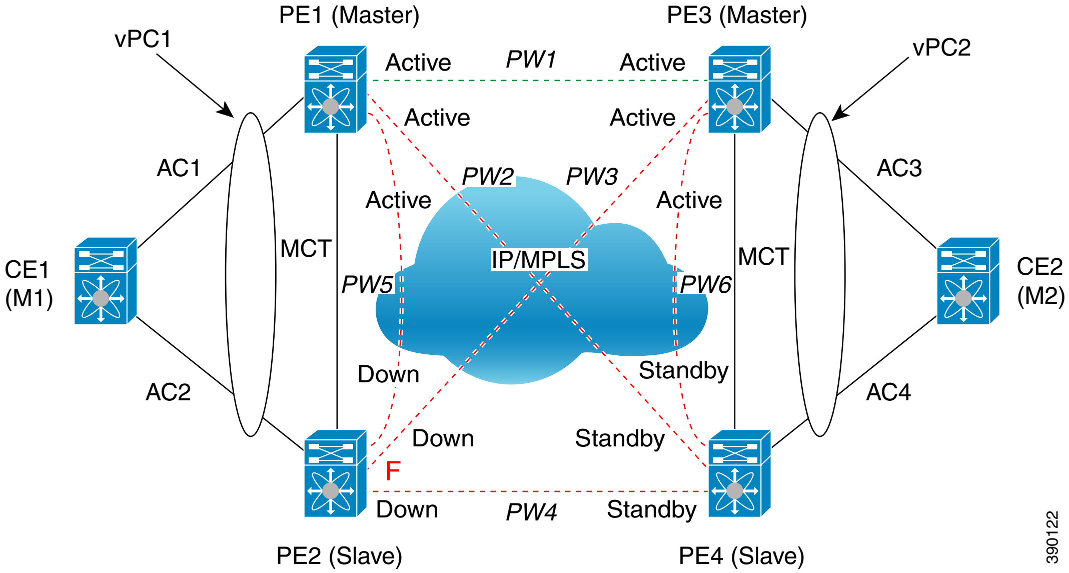

When all PWs in the core on the Active vPC node go down, the peer Standby vPC node changes its state to Active and advertises the local status of Active on all core PWs. PW 1 becomes Active.

When all core interfaces on a node go down, the PWs in the VPLS domain also go down.

Licensing for Layer 2 VPN VPLS Dual-Homing with a vPC

The following table shows the licensing requirements for this feature:

Product	License Requirement
Cisco NX-OS	Layer 2 MVPN requires an MPLS license. For a complete explanation of the Cisco NX-OS licensing scheme and how to obtain and apply licenses, see the *Cisco NX-OS Licensing Guide*.

Guidelines and Limitations for Layer 2 VPN VPLS Dual-Homing with a vPC

The VPLS Dual-Homing with a vPC feature is supported only on switch port links. Virtual port channel (vPC) is not supported on an Ethernet virtual circuit (EVC) port.

Configuring Layer 2 VPN VPLS Dual-Homing with a vPC

You can configure a vPC peer as the primary node in a dual-homed topology. Repeat this task to configure the other vPC peer as the secondary node.

Before You Begin

- Ensure that the Layer 2 VPN feature is enabled on the switch.
- Ensure that the dual-homed vPC domains are configured.
- Ensure that a VPLS domain with a mesh of core PWs that connect the PEs of the vPC domain is configured.

SUMMARY STEPS

1. **configure terminal**
2. [**no**] **l2vpn vfi context** *vfi-name*
3. (Optional) **description** *description*
4. **vpn id** *vpn-id*
5. **redundancy** {**primary** | **secondary**}
6. **member** *ip-address* **encapsulation mpls**
7. **exit**
8. [**no**] **bridge domain** *domain-id*
9. **member vfi** *vfi-name*
10. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose

Step 1	**`configure terminal`**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	`[`**`no`**`]` **`l2vpn vfi context`** *`vfi-name`*
	Example:
	`switch(config)# l2vpn vfi context vpls80`
	`switch(config-l2vpn-vfi)#`
		Establishes a Layer 2 VPN (L2VPN) Virtual Forwarding Interface (VFI) between two or more separate networks.
		• The *vfi-name* argument is a unique per-interface identifier for this VFI. The maximum range is 100 alphanumeric, case-sensitive characters.
		Note You can use the **no** form of this command to delete the VFI and the associated configuration.

Step 3	**`description`** *`description`*
	Example:
	`switch(config-l2vpn-vfi)# description VFIforDualHome`
		(Optional) Adds a description to the interface configuration.
		• The maximum range for the *description* argument is 254 alphanumeric characters.

	Command	Purpose
Step 4	**vpn** *vpn-id* **Example:** `switch(config-l2vpn-vfi)# vpn`	Configures a Virtual Private Network (VPN) ID on a VFI context. • The valid range is from 1 to 4294967295.

Step 5	**redundancy** {**primary** | **secondary**}
	Example:
	`switch(config-l2vpn-vfi)# redundancy primary`
		Configures this L2VPN VFI context as the primary or secondary node.

Step 6	**member** *ip-address* **encapsulation mpls** **Example:** `switch(config-l2vpn-vfi)# member 10.0.0.3 encapsulation mpls`	Specifies the devices that form a point-to-point L2VPN VFI connection.

Step 7	**exit**
	Example:
	`switch(config-l2vfi-vfi)# vpn 80`
	`switch (config)#`
		Exits Layer 2 VFI configuration mode.

Step 8	[**no**] **bridge-domain** *domain-id*
	Example:
	`switch(config)# bridge-domain 100`
	`switch(config-bdomain)#`
		Enters bridge-domain configuration mode and configures a bridge domain.
		• The *domain-id* argument is a unique identifier for the bridge domain and underlying VLAN to be created. The valid range is defined by the system bridge-domain configuration.
		Note You can use the **no** form of this command to remove the bridge-domain configuration including port associations. Removing the bridge-domain configuration does not remove the underlying VLAN.
		If a VLAN is associated with a bridge domain, you cannot remove the VLAN without first removing the bridge domain. To remove the underlying VLAN, use the **no vlan** command after you remove the bridge domain.

Step 9	**member vfi** *vfi-name* **Example:** `switch(config-bdomain)# member vfi vpls80`	(Optional) Binds a VFI to this bridge domain. • The *vfi-name* argument identifies the VFI to be bound. The maximum range is 100 alphanumeric, case-sensitive characters.
Step 10	**copy running-config startup-config** **Example:** `switch(config-if-srv)# copy running-config startup-config`	(Optional) Saves this configuration change.

Configuration Examples for Layer 2 VPN VPLS Dual-Homing with a vPC

PE1

```
l2vpn vfi context vpls-80
  vpn id 80
  redundancy primary
  member 10.0.0.4 encapsulation mpls
!
bridge-domain 80
  member vfi vpls-80
```

PE2

```
l2vpn vfi context vpls-80
  vpn id 80
  redundancy secondary
  member 10.0.0.4 encapsulation mpls
!
bridge-domain 80
  member vfi vpls-80
```

Additional References for Layer 2 VPN VPLS Dual-Homing with a vPC

Related Documents

Related Topic	Document Title
Interface commands	*Cisco Nexus 7000 Series NX-OS Interfaces Command Reference*
VLAN commands	*Cisco Nexus 7000 Series NX-OS Layer 2 Switching Command Reference*
Virtual port channels	“Configuring vPCs” chapter of the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide*

Feature History for Layer 2 VPN VPLS Dual-Homing with a vPC

Table 33-1 lists the history for this feature.

Table 33-1 Feature History for Layer 2 VPN VPLS Dual-Homing wit vPC

Feature Name	Releases	Feature Information
VPLS Dual-Homing with a vPC	6.2(2)	The VPLS Dual-Homing with a vPC feature integrates Virtual Private LAN (VPLS) with the virtual port channel (vPC) functionality in active-standby mode. This feature allows traffic from a customer edge (CE) device to be load-balanced across both provider edge (PE) devices. The active PE can then forward the traffic to the core. Similarly, traffic from the core can be received by the active PE and sent to the attached CE.
IP tunnels in VDC other than default	4.2(1)	This features was introduced.