CHAPTER 7

Configuring vPCs

This chapter describes how to configure virtual port channels (vPCs) on Cisco NX-OS devices.

Note Beginning with Cisco NX-OS Release 5.1(1), vPCs have been enhanced to interoperate with FabricPath. To configure vPCs with FabricPath networks, see the *Cisco Nexus 7000 Series NX-OS FabricPath Configuration Guide, Release 5.x*.

Beginning with Cisco NX-OS Release 5.1(1), you can use any of the 10-Gbps Ethernet (10GE) interfaces on the F series modules or the 10GE interfaces on the M series modules for the vPC peer link on an individual switch, but you cannot combine member ports on an F module with ports on an M module into a single port channel on a single switch. The port channel compatibility parameters must be the same for all the port channel members on the physical switch.

You cannot configure shared interfaces to be part of a vPC. See the *Cisco NX-OS FCoE Configuration Guide for Cisco Nexus 7000 and Cisco MDS 9500* for more information on shared interfaces.

Note The port channel compatibility parameters must also be the same for all vPC member ports on both peers and therefore you must use the same type of module in each chassis.

This chapter includes the following sections:

Note For information about configuring port channels and the Link Aggregation Control Protocol (LACP), see Chapter 6, “Configuring Port Channels.”

Information About vPCs

This section includes the following topics:

vPC Overview

A virtual port channel (vPC) allows links that are physically connected to two different Cisco Nexus 7000 Series devices to appear as a single port channel by a third device (see Figure 7-1). The third device can be a switch, server, or any other networking device that supports port channels. A vPC can provide Layer 2 multipathing, which allows you to create redundancy and increase bisectional bandwidth by enabling multiple parallel paths between nodes and allowing load balancing traffic.

Figure 7-1 vPC Architecture

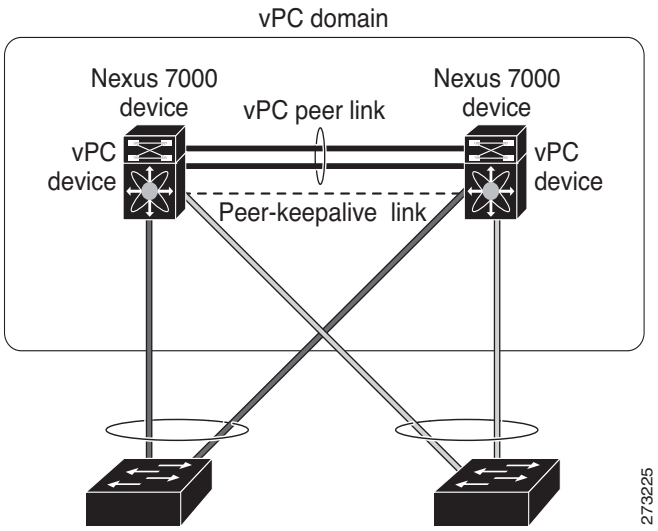

You can use only Layer 2 port channels in the vPC. A vPC domain is associated to a single VDC, so all vPC interfaces belonging to a given vPC domain must be defined in the same VDC. You must have a separate vPC peer link and peer-keepalive link infrastructure for each VDC deployed. Consolidating a vPC pair (two vPC peer devices of the same domain) in two VDCs of the same physical device is not supported. The vPC peer link must use 10-Gigabit Ethernet ports for both ends of the link or the link will not form.

You configure the port channels by using one of the following:

- No protocol
- Link Aggregation Control Protocol (LACP)

When you configure the port channels in a vPC—including the vPC peer link channel—without using LACP, each device can have up to 8 active links in a single port channel. When you configure the port channels in a vPC—including the vPC peer link channels—using LACP, each device can have 8 active links and 8 standby links in a single port channel. (See the "vPC Interactions with Other Features" section on page 7-21 for more information on using LACP and vPCs.)

Note You must enable the vPC feature before you can configure or run the vPC functionality.

Beginning with Cisco NX-OS Release 4.2, the system automatically takes a checkpoint prior to disabling the feature, and you can roll back to this checkpoint. See the *Cisco Nexus 7000 Series NX-OS System Management Configuration Guide, Release 5.x*, for information on roll backs and checkpoints.

After you enable the vPC functionality, you create the peer-keepalive link, which sends heartbeat messages between the two vPC peer devices.

You can create a vPC peer link by configuring a port channel on one Cisco Nexus 7000 Series chassis by using two or more 10-Gigabit Ethernet ports in dedicated mode. To ensure that you have the correct hardware to enable and run vPC beginning with Cisco NX-OS Release 4.1(5), enter the **show hardware feature-capability** command. If you see an X across from vPC, your hardware cannot enable the vPC feature.

We recommend that you configure the vPC peer link Layer 2 port channels as trunks. Then, on another Cisco Nexus 7000 Series chassis, you configure another port channel again using two or more 10-Gigabit Ethernet ports in dedicated mode. Connecting these two port channels creates a vPC peer link in which

the two linked Cisco Nexus devices appear as one device to a third device. The third device, or downstream device, can be a switch, server, or any other networking device that uses a regular port channel to connect to the vPC. If you are not using the correct module, the system displays an error message.

Note We recommend that you configure the vPC peer links on dedicated ports of different modules to reduce the possibility of a failure. For the best resiliency scenario, use at least two modules.

Beginning with Cisco NX-OS Release 4.2, if you must configure all the vPC peer links and core-facing interfaces on a single module, you should configure a track object that is associated with the Layer 3 link to the core and on all the links on the vPC peer link on both vPC peer devices. Once you configure this feature and if the primary vPC peer device fails, the system automatically suspends all the vPC links on the primary vPC peer device. This action forces all the vPC traffic to the secondary vPC peer device until the system stabilizes.

Create a track object and apply that object to all links on the primary vPC peer device that connect to the core and to the vPC peer link. See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 5.x*, for information on the **track interface** command.

The vPC domain includes both vPC peer devices, the vPC peer-keepalive link, the vPC peer link, and all of the port channels in the vPC domain connected to the downstream device. You can have only one vPC domain ID on each device.

In this version, you can connect each downstream device to a single vPC domain ID using a single port channel.

Note Always attach all vPC devices using port channels to both vPC peer devices.

A vPC (see Figure 7-2) provides the following benefits:

- Allows a single device to use a port channel across two upstream devices
- Eliminates Spanning Tree Protocol (STP) blocked ports
- Provides a loop-free topology
- Uses all available uplink bandwidth
- Provides fast convergence if either the link or a device fails
- Provides link-level resiliency
- Assures high availability

Figure 7-2 vPC Interfaces in One VDC

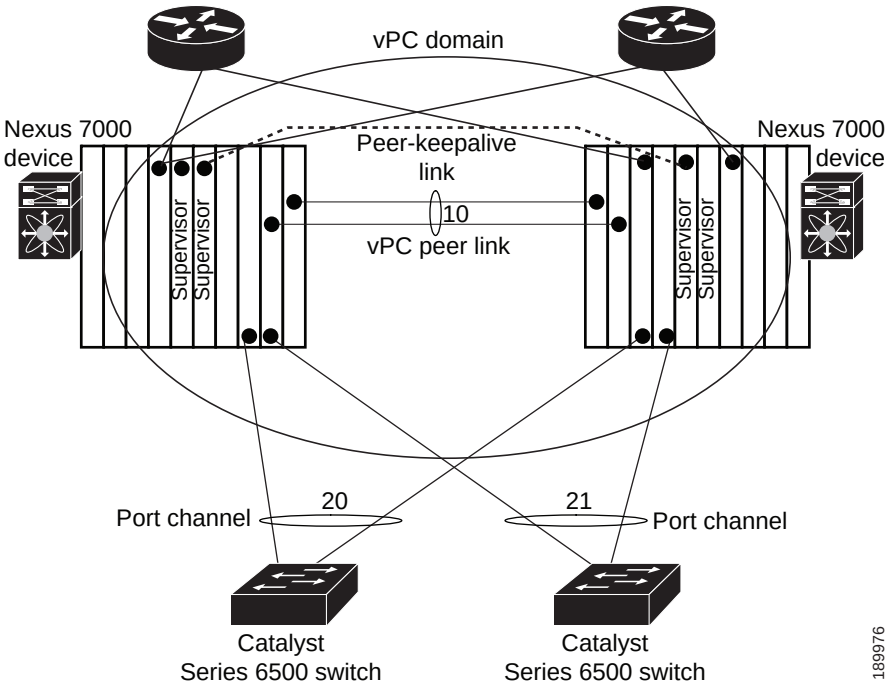

For more information on VDCs, see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 5.x*.

vPC Terminology

The terminology used in vPCs is as follows:

- vPC—The combined port channel between the vPC peer devices and the downstream device.
- vPC peer device—One of a pair of devices that are connected with the special port channel known as the vPC peer link.
- vPC peer link—The link used to synchronize states between the vPC peer devices. Both ends must be on 10-Gigabit Ethernet interfaces.
- vPC member port—An interface that belongs to a vPC.
- Host vPC port—A Fabric Extender host interfacesthat belongs to a vPC.
- vPC domain—This domain includes both vPC peer devices, the vPC peer-keepalive link, and all of the port channels in the vPC connected to the downstream devices. It is also associated to the configuration mode that you must use to assign vPC global parameters.
- vPC peer-keepalive link—The peer-keepalive link monitors the vitality of a vPC peer Cisco Nexus 7000 Series device. The peer-keepalive link sends configurable, periodic keepalive messages between vPC peer devices.

  We recommend that you associate a peer-keepalive link to a separate VRF mapped to a Layer 3 interface in each vPC peer device. If you do not configure a separate VRF, the system uses the management VRF by default. However, if you use the management interfaces for the peer-keepalive link, you must put a management switch connected to both the active and standby management ports on each vPC peer device (see Figure 7-3).

Figure 7-3 Separate Switch Required to Connect Management Ports for vPC Peer-Keepalive Link

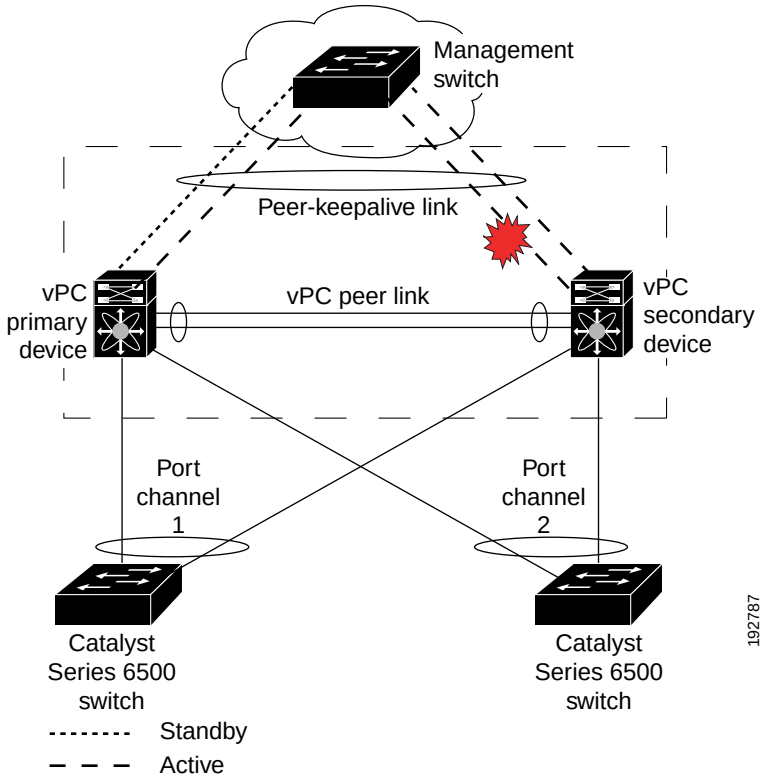

No data or synchronization traffic moves over the vPC peer-keepalive link; the only traffic on this link is a message that indicates that the originating switch is operating and running vPC.

- vPC member port—Interfaces that belong to the vPCs.
- Dual-active— Both vPC peers act as primary. This situation occurs when the peer-keepalive and peer-link go down when both the peers are still active. In this case, the secondary vPC assumes that the primary vPC is inactive and acts as the primary vPC.
- Recovery—When the peer-keepalive and the peer-link come up, one switch becomes the secondary vPC. On the switch that becomes the secondary vPC, the vPC links go down and come back up.

vPC Peer Links

A vPC peer link is the link that is used to synchronize the states between the vPC peer devices. Both ends of the link much be on 10-Gigabit Ethernet interfaces.

Since vPC includes a split control plane working on two switches, vPC uses the peer link as follows:

- Keeps both vPC peer switches in sync for control plane info (such as vPC state, consistency parameters, and MAC addresses).
- Forwards data packets to the vPC peer switch, when the local vPC is down.
- A single vPC domain between two VDCs on the same physical Cisco Nexus 7000 device is not supported.

This section describes the vPC peer link and includes the following topics:

- vPC Peer Link Overview, page 7-8

- Features That You Must Manually Configure on the Primary and Secondary Devices, page 7-10
- Configuring Layer 3 Backup Routes on a vPC Peer Link, page 7-11

Note You must configure the peer-keepalive link before you configure the vPC peer link or the peer link will not come up. (See the "Peer-Keepalive Link and Messages" section on page 7-11 for information on the vPC peer-keepalive link and messages.)

You can configure a vPC peer link to configure two devices as vPCs peers. You must use the module in order to configure a vPC peer link.

Note We recommend that you use the dedicated port mode when you configure a vPC peer link. For information about the dedicated port mode, see Chapter 2, "Configuring Basic Interface Parameters."

vPC Peer Link and I/O Modules Support

Only identical I/O modules on either side of a vPC peer link are supported. Using different I/O modules on either side of a vPC peer link is not supported. Mixing I/O modules on the same side of a port channel is also not supported. Table 7-1 displays the I/O modules that are supported on both sides of a vPC peer link. Figure 7-4 to Figure 7-7 show the configurations that are supported and not supported on both sides of a vPC peer link.

Table 7-1 I/O Module Combinations Supported on Both Sides of a vPC Peer Link

vPC Primary	vPC Secondary
F1 I/O module	F1 I/O module
M1 I/O module	M1 I/O module

Figure 7-4 Supported—Two F1 I/O Modules on Either Side of a vPC Peer Link

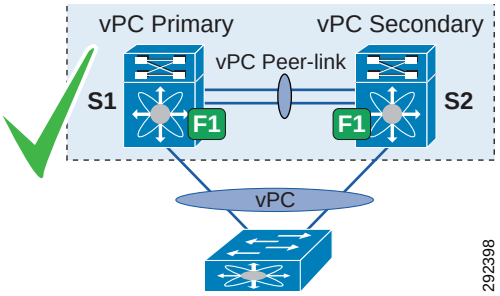

Figure 7-5 Supported—Two M1 I/O Modules on Either Side of a vPC Peer Link

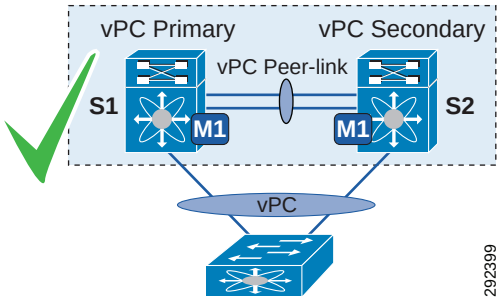

Figure 7-6 Not Supported—Two Different I/O Modules on Either Side of a vPC Peer Link (M1 and F1)

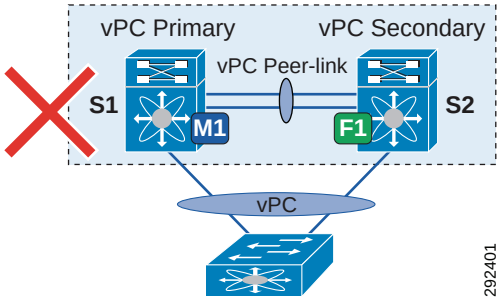

Figure 7-7 Not Supported—Mixing I/O Modules on the Same Side of a Peer Link

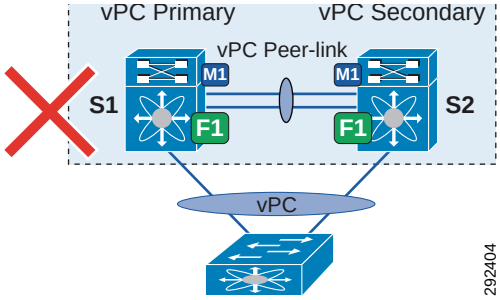

vPC Peer Link Overview

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

See Figure 7-8 for invalid vPC peer configurations.

Figure 7-8 vPC Peer Configurations That Are Not Allowed

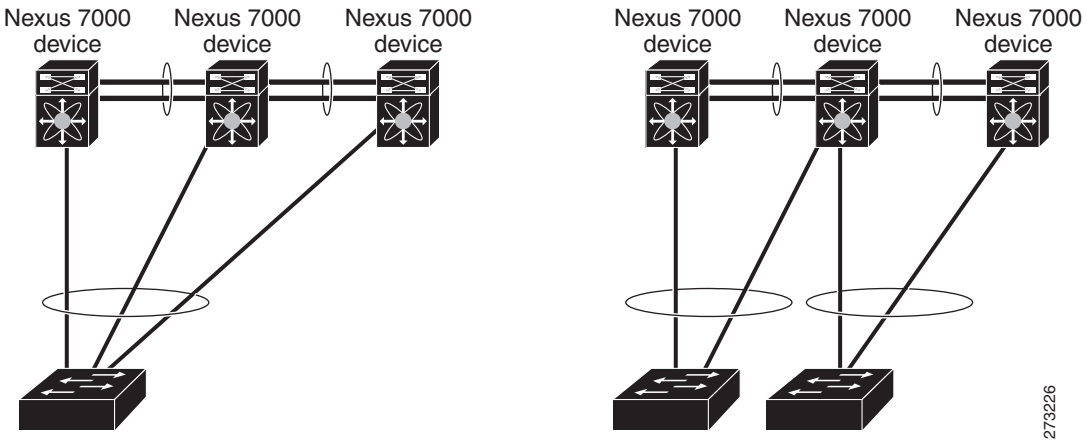

To make a valid configuration, you first configure a port channel on each device and then configure the vPC domain. You assign the port channel on each device as a peer link, using the same vPC domain ID. For redundancy, we recommend that you should configure at least two of the dedicated ports into the port channel because if one of the interfaces in the vPC peer link fails, the device automatically falls back to use another interface in the peer link.

Note We recommend that you configure the Layer 2 port channels in trunk mode.

Many operational parameters and configuration parameters must be the same in each device connected by a vPC peer link (see the "Compatibility Parameters for vPC Interfaces" section on page 7-15). Because each device is completely independent on the management plane, you must ensure that the devices are compatible on the critical parameters. vPC peer devices have separate control planes. After configuring the vPC peer link, you should display the configuration on each vPC peer device to ensure that the configurations are compatible.

Note You must ensure that the two devices connected by the vPC peer link have certain identical operational and configuration parameters. For more information on required configuration consistency, see the "Compatibility Parameters for vPC Interfaces" section on page 7-15.

When you configure the vPC peer link, the vPC peer devices negotiate that one of the connected devices is the primary device and the other connected device is the secondary device (see the "Configuring vPCs" section on page 7-31). The Cisco NX-OS software uses the lowest MAC address to elect the primary device. The software takes different actions on each device—that is, the primary and secondary—only in certain failover conditions. If the primary device fails, the secondary device becomes the new primary device when the system recovers, and the previously primary device is now the secondary device.

You can also configure which of the vPC devices is the primary device. Changing the priority of the vPC peer devices can cause the interfaces in your network to go up and down. If you want to configure the role priority again to make one vPC device the primary device, configure the role priority on both the primary vPC device with a lower priority value, and secondary vPC device with the higher value. Then, shut down the port channel that is the vPC peer link on both devices by entering the **shutdown command**, and finally reenable the port channel on both devices by enter the **no shutdown** command.

Note We recommend that you use two different modules for redundant on each vPC peer device vPC peer links.

The software keeps all traffic that forwards across the vPC peer devices as local traffic. A packet that ingresses the port channel uses one of the local links rather than moving across the vPC peer link. Unknown unicast, multicast, and broadcast traffic (including STP BPDUs) are flooded across the vPC peer link. The software keeps the multicast forwarding state synchronized on both of the vPC peer devices.

You can configure any of the standard load-balancing schemes on both the vPC peer link devices and the downstream device (see Chapter 6, "Configuring Port Channels," for information on load balancing).

Configuration information flows across the vPC peer links using the Cisco Fabric Service over Ethernet (CFSoE) protocol. (See the "CFSoE" section on page 7-27 for more information about CFSoE.)

All MAC addresses for those VLANs configured on both devices are synchronized between vPC peer devices. The software uses CFSoE for this synchronization. (See the "CFSoE" section on page 7-27 for information about CFSoE.)

If the vPC peer link fails, the software checks the status of the remote vPC peer device using the peer-keepalive link, which is a link between vPC peer devices that ensures that both devices are up. If the vPC peer device is up, the secondary vPC device disables all vPC ports on its device, to prevent loops and disappearing or flooding traffic. The data then forwards down the remaining active links of the port channel.

Note We recommend that you create and configure a separate VRF and configure a Layer 3 port on each vPC peer device in that VRF for the vPC peer-keepalive link. The default ports and VRF for the peer-keepalive are the management ports and VRF.

The software learns of a vPC peer device 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 devices. 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 device. The keepalive messages are used only when all the links in the peer link fail. See the "Peer-Keepalive Link and Messages" section on page 7-11 for information about the keepalive message.

Features That You Must Manually Configure on the Primary and Secondary Devices

You must manually configure the following features to conform to the primary/secondary mapping of each of the vPC peer devices:

- STP root—Configure the primary vPC peer device as the STP primary root device and configure the vPC secondary device to be the STP secondary root device. See the "vPC Peer Links and STP" section on page 7-21 for more information on vPC and STP.
  - We recommend that you configure the vPC peer link interfaces as STP network ports so that Bridge Assurance is enabled on all vPC peer links
  - We recommend that you configure Rapid PVST+ so that the primary device is the root for all VLANs and configure MST so that the primary device is the root for all instances.
- Layer 3 VLAN network interface—Configure Layer 3 connectivity from each vPC peer device by configuring a VLAN network interface for the same VLAN from both devices.

- HSRP active—If you want to use HSRP and VLAN interfaces on the vPC peer devices, configure the primary vPC peer device with the HSRP active highest priority. Configure the secondary device to be the HSRP standby. And ensure that you have VLAN interfaces on each vPC device that are in the same administrative and operational mode. (See the "vPC Peer Links and Routing" section on page 7-26 for more information on vPC and HSRP.)

While you configure Unidirectional Link Detection (UDLD), note the following recommendations:

- If LACP is used as port-channel aggregation protocol, UDLD is not required in a vPC domain.
- If LACP is not used as the port-channel aggregation protocol (static port-channel), use UDLD in normal mode on vPC member ports.
- If STP is used without Bridge Assurance and if LACP is not used, use UDLD in normal mode on vPC orphan ports.

See the "Configuring the UDLD Mode" section on page 2-39 for information on configuring UDLD.

Configuring Layer 3 Backup Routes on a vPC Peer Link

You can use VLAN network interfaces on the vPC peer devices to link to Layer 3 of the network for such applications as HSRP and PIM. However, we recommend that you configure a separate Layer 3 link for routing from the vPC peer devices, rather than using a VLAN network interface for this purpose.

Note Ensure that you have a VLAN network interface configured on each peer device and that the interface is connected to the same VLAN on each device. Also, each VLAN interface must be in the same administrative and operational mode. For more information on configuring VLAN network interfaces, see Chapter 4, "Configuring Layer 3 Interfaces."

If a failover occurs on the vPC peer link, the VLAN interfaces on the vPC peer devices are also affected. If a vPC peer link fails, the system brings down associated VLAN interfaces on the secondary vPC peer device.

Beginning with Cisco NX-OS Release 4.2(1), you can ensure that specified VLAN interfaces do not go down on the vPC secondary device when the vPC peer link fails.

Use the **dual-active exclude interface-vlan** command to configure this feature.

Note When you attach a Layer 3 device to a vPC domain, the peering of routing protocols using a VLAN also carried on the vPC peer link is not supported. If routing protocol adjacencies are needed between vPC peer devices and a generic Layer 3 device, you must use physical routed interfaces for the interconnection. Use of the vPC peer-gateway feature does not change this requirement.

Peer-Keepalive Link and Messages

The Cisco NX-OS software uses the peer-keepalive link between the vPC peers to transmit periodic, configurable keepalive messages. You must have Layer 3 connectivity between the peer devices to transmit these messages; the system cannot bring up the vPC peer link unless the peer-keepalive link is already up and running.

Note We recommend that you associate the vPC peer-keepalive link to a separate VRF mapped to a Layer 3 interface in each vPC peer device. If you do not configure a separate VRF, the system uses the management VRF and management ports by default. Do not use the peer link itself to send and receive vPC peer-keepalive messages. For more information on configuring VRFs, see the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 5.x*.

If one of the vPC peer devices fails, the vPC peer device on the other side of the vPC peer link senses the failure by not receiving any peer-keepalive messages. The default interval time for the vPC peer-keepalive message is 1 second, and you can configure the interval between 400 milliseconds and 10 seconds.

You can configure a hold-timeout value with a range of 3 to 10 seconds; the default hold-timeout value is 3 seconds. This timer starts when the vPC peer link goes down. During this hold-timeout period, the secondary vPC peer device ignores vPC peer-keepalive messages, which ensures that network convergence occurs before vPC action takes place. The purpose of the hold-timeout period is to prevent false-positive cases.

You can also configure a timeout value with a range of 3 to 20 seconds; the default timeout value is 5 seconds. This timer starts at the end of the hold-timeout interval. During the timeout period, the secondary vPC peer device checks for vPC peer-keepalive hello messages from the primary vPC peer device. If the secondary vPC peer device receives a single hello message, that device disables all vPC interfaces on the secondary vPC peer device.

The difference between the hold-timeout and the timeout parameters is as follows:

- During the hold-timeout, the vPC secondary device does not take any action based on any keepalive messages received, which prevents the system taking action when the keepalive might be received just temporarily, such as if a supervisor fails a few seconds after the peer link goes down.
- During the timeout, the vPC secondary device takes action to become the vPC primary device if no keepalive message is received by the end of the configured interval.

See the "Configuring vPCs" section on page 7-31 for information on configuring the timer for the keepalive messages.

Note Ensure that both the source and destination IP addresses used for the peer-keepalive messages are unique in your network and these IP addresses are reachable from the VRF associated with the vPC peer-keepalive link.

Use the command-line interface (CLI) to configure the interfaces you are using the vPC peer-keepalive messages as trusted ports. Leave the precedence at the default (6) or configure it higher. The following is an example of configuring an interface as a trusted port:

```
(config)# class-map type qos match-all trust-map
(config-cmap-qos)# match cos 4-7

(config)# policy-map type qos ingresspolicy
(config-pmap-qos)# class trust-map

(config)# interface Ethernet8/11
(config-if)# service-policy type qos input ingresspolicy
```

See the *Cisco Nexus 7000 Series NX-OS Quality of Service Configuration Guide, Release 5.x*, for complete information on configuring trusted ports and precedence.

vPC Peer-Gateway

Beginning with Cisco NX-OS Release 4.2(1), you can configure vPC peer devices to act as the gateway even for packets that are destined to the vPC peer device's MAC address.

Use the **peer-gateway** command to configure this feature.

Some network-attached storage (NAS) devices or load-balancers may have features aimed to optimize the performances of particular applications. Essentially these features avoid performing a routing-table lookup when responding to a request that originated from a host not locally attached to the same subnet. Such devices may reply to traffic using the MAC address of the sender Cisco Nexus 7000 Series device rather than the common HSRP gateway. Such behavior is non-complaint with some basic Ethernet RFC standards. Packets reaching a vPC device for the non-local router MAC address are sent across the peer link and could be dropped by the built in vPC loop avoidance mechanism if the final destination is behind another vPC.

The vPC peer-gateway capability allows a vPC switch to act as the active gateway for packets that are addressed to the router MAC address of the vPC peer. This feature enables local forwarding of such packets without the need to cross the vPC peer link. In this scenario, the feature optimizes use of the peer link and avoids potential traffic loss.

Configuring the peer-gateway feature must be done on both primary and secondary vPC peers and is nondisruptive to the operations of the device or to the vPC traffic. The vPC peer-gateway feature can be configured globally under the vPC domain submode.

When you enable this feature, Cisco NX-OS automatically disables IP redirects on all interface VLANs mapped over a vPC VLAN to avoid generation of IP redirect messages for packets switched through the peer gateway router.

Note With Cisco NX-OS Release 5.1(3) and above, when a VLAN interface is used for Layer 3 backup routing on the vPC peer devices and an F1 linecard is used as the peer link, the VLAN must be excluded from the peer-gateway feature, if enabled, by running the **peer-gateway exclude-vlan** *vlan-number* command. For more information about backup routes, see the "Configuring Layer 3 Backup Routes on a vPC Peer Link" section on page 7-11.

Packets arriving at the peer-gateway vPC device will have their TTL decremented, so packets carrying TTL = 1 may be dropped in transit due to TTL expire. This needs to be taken into account when the peer-gateway feature is enabled and particular network protocols sourcing packets with TTL = 1 operate on a vPC VLAN.

vPC Domain

You can use the vPC domain ID to identify the vPC peer links and the ports that are connected to the vPC downstream devices.

The vPC domain is also a configuration mode that you use to configure the keepalive messages, and configure other vPC peer link parameters rather than accept the default values. See the "Configuring vPCs" section on page 7-31 for more information on configuring these parameters.

To create a vPC domain, you must first create a vPC domain ID on each vPC peer device using a number from 1 to 1000. You can have only one vPC domain per VDC.

You must explicitly configure the port channel that you want to act as the peer link on each device. You associate the port channel that you made a peer link on each device with the same vPC domain ID to form a single vPC domain. Within this domain, the system provides a loop-free topology and Layer 2 multipathing.

You can only configure these port channels and vPC peer links statically. All ports in the vPC on each of the vPC peer devices must be in the same VDC. You can configure the port channels and vPC peer links either using LACP or no protocol. We recommend that you use LACP with the interfaces in active mode to configure port channels in each vPC, which ensures an optimized, graceful recovery in a port-channel failover scenario and provides configuration checks against configuration mismatches among the port channels themselves.

The vPC peer devices use the vPC domain ID that you configure to automatically assign a unique vPC system MAC address. Each vPC domain has a unique MAC address that is used as a unique identifier for the specific vPC-related operations, although the devices use the vPC system MAC addresses only for link-scope operations, such as LACP. We recommend that you create each vPC domain within the contiguous Layer 2 network with a unique domain ID. You can also configure a specific MAC address for the vPC domain, rather than having the Cisco NX-OS software assign the address.

See the "CFSoE" section on page 7-27 for more information on displaying the vPC MAC table.

After you create a vPC domain, the Cisco NX-OS software creates a system priority for the vPC domain. You can also configure a specific system priority for the vPC domain.

Note When manually configuring the system priority, you must ensure that you assign the same priority value on both vPC peer devices. If the vPC peer devices have different system priority values, vPC will not come up.

vPC Topology

Figure 7-3 shows a basic configuration in which the Cisco Nexus 7000 Series Switch ports are directly connected to another switch or host and are configured as part of a port channel that becomes part of a vPC.

Figure 7-9 Switch vPC Topology

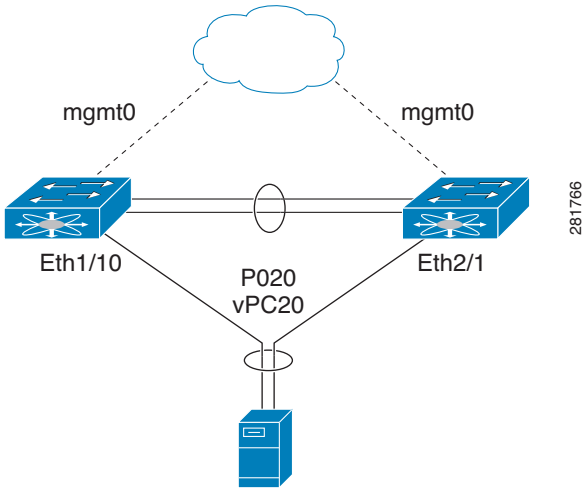

In the figure, vPC 20 is configured on port channel 20, which has Eth1/10 on the first switch and Eth2/1 on the second as member ports.

Beginning with Cisco NX-OS Release 5.2(1), you can configure a vPC from the peer devices through Fabric Extenders (FEXs), as shown in Figure 7-10.

Figure 7-10 FEX Straight-Through Topology (Host vPC)

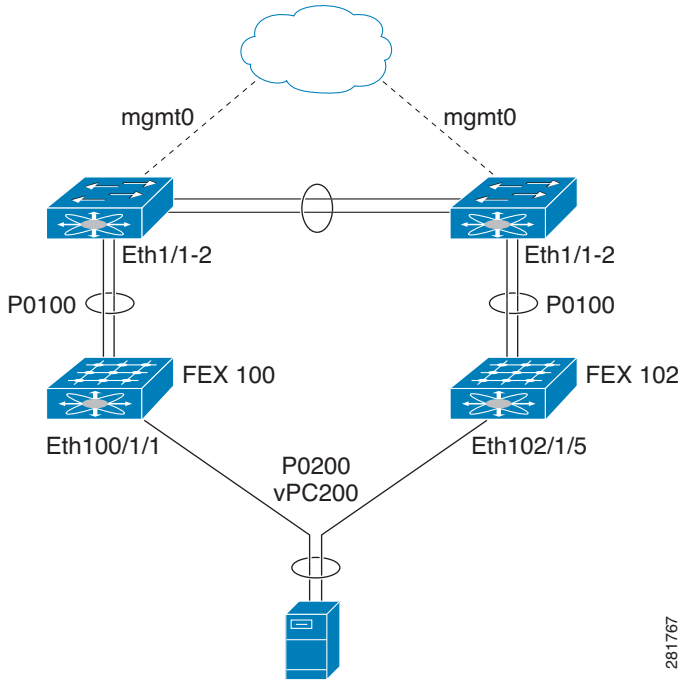

In the figure, each FEX is single-homed (straight-through FEX topology) with a Cisco Nexus 7000 Series Switch. The host interfaces on this FEX are configured as port channels and those port channels are configured as vPCs. Eth100/1/1 and Eth102/1/5 are configured as members of PO200, and PO200 is configured for vPC 200.

In both topologies, port channels P020 and P0200 must be configured identically on the peer switches and configuration synchronization is used to synchronize the configurations of the vPC switches.

See the *Configuring the Cisco Nexus 2000 Series Fabric Extender* for more information on configuring FEX ports.

Compatibility Parameters for vPC Interfaces

Many configuration and operational parameters must be identical on all interfaces in the vPC. We recommend that you configure the Layer 2 port channels that you use for the vPC peer link in trunk mode.

After you enable the vPC feature and configure the peer link on both vPC peer devices, CFS messages provide a copy of the configuration on the local vPC peer device configuration to the remote vPC peer device. The system then determines whether any of the crucial configuration parameters differ on the two devices. (See the "CFSoE" section on page 7-27 for more information on CFS.)

Note Enter the **show vpc consistency-parameters** command to display the configured values on all interfaces in the vPC. The displayed configurations are only those configurations that would limit the vPC peer link and vPC from coming up.

The compatibility check process for vPCs differs from the compatibility check for regular port channels. See Chapter 6, "Configuring Port Channels," for information on regular port channels.

This section includes the following sections:

Configuration Parameters That Must Be Identical

The configuration parameters in this section must be configured identically on both devices of the vPC peer link; otherwise, the vPC moves fully or partially into a suspended mode.

Caution When spanning tree is disabled for any VLANs irrespective of whether it is a VPC VLAN or not, it will cause a Global type 1 inconsistency causing all VPCs VLANs to be suspended. This issue is also observed when disabling spanning tree for a non-existent VLAN.

Note You must ensure that all interfaces in the vPC have the identical operational and configuration parameters listed below.

Note Enter the **show vpc consistency-parameters** command to display the configured values on all interfaces in the vPC. The displayed configurations are only those configurations that would limit the vPC peer link and vPC from coming up.

The devices automatically check for compatibility for some of these parameters on the vPC interfaces. The per-interface parameters must be consistent per interface, and the global parameters must be consistent globally:

- Port-channel mode: on, off, or active
- Link speed per channel
- Duplex mode per channel
- Trunk mode per channel:
  - Native VLAN
  - VLANs allowed on trunk
  - Tagging of native VLAN traffic
- Spanning Tree Protocol (STP) mode
- STP region configuration for Multiple Spanning Tree
- Enable/disable state per VLAN

- STP global settings:
  - Bridge Assurance setting
  - Port type setting
  - Loop Guard settings
- STP interface settings:
  - Port type setting
  - Loop Guard
  - Root Guard
- Maximum Transmission Unit (MTU)

If any of these parameters are not enabled or defined on either device, the vPC consistency check ignores those parameters.

Note To ensure that none of the vPC interfaces are in the suspend mode, enter the **show vpc brief** and **show vpc consistency-parameters** commands and check the syslog messages.

Configuration Parameters That Should Be Identical

When any of the following parameters are not configured identically on both vPC peer devices, a misconfiguration may cause undesirable behavior in the traffic flow:

- MAC aging timers
- Static MAC entries
- VLAN interface—Each device on the end of the vPC peer link must have a VLAN interface configured for the same VLAN on both ends and they must be in the same administrative and operational mode. Those VLANs configured on only one device of the peer link do not pass traffic using the vPC or peer link. You must create all VLANs on both the primary and secondary vPC devices, or the VLAN will be suspended.
- All ACL configurations and parameters
- Quality of Service (QoS) configuration and parameters
- STP interface settings:
  - BPDU Filter
  - BPDU Guard
  - Cost
  - Link type
  - Priority
  - VLANs (Rapid PVST+)
- Port security
- Cisco Trusted Security (CTS)
- Dynamic Host Configuration Protocol (DHCP) snooping
- Network Access Control (NAC)
- Dynamic ARP Inspection (DAI)

- IP source guard (IPSG)
- Internet Group Management Protocol (IGMP) snooping
- Hot Standby Routing Protocol (HSRP)
- Protocol Independent Multicast (PIM)
- Gateway Load-Balancing Protocol (GLBP)
- All routing protocol configurations

To ensure that all the configuration parameters are compatible, we recommend that you display the configurations for each vPC peer device once you configure the vPC.

Consequences of Parameter Mismatch

In releases earlier than Cisco NX-OS Release 5.2(1), when a consistency check detects a mismatch in a parameter from the list of parameters that must be identical, the vPC peer link and vPC are prevented from coming up. If a parameter mismatch is configured after the vPC is already established, the vPC moves into suspend mode and no traffic flows on the vPC.

Beginning with Cisco NX-OS Release 5.2(1), you can configure the graceful consistency check feature, which suspends only the links on the secondary peer device when a mismatch is introduced in a working vPC. This feature is configurable only in the CLI and is enabled by default.

Use the **graceful consistency-check** command to configure this feature.

As part of the consistency check of all parameters from the list of parameters that must be identical, the system checks the consistency of all VLANs. In releases earlier than Cisco NX-OS Release 5.2(1), if the configuration of any enabled VLAN is inconsistent across the peer devices, the vPC is prevented from establishing or moves into a suspended mode.

Beginning with Cisco NX-OS Release 5.2(1), the vPC will remain operational, and only the inconsistent VLANs will be brought down. This per-VLAN consistency check feature cannot be disabled and does not apply to Multiple Spanning Tree (MST) VLANs.

vPC Number

Once you have created the vPC domain ID and the vPC peer link, you create port channels to attach the downstream device to each vPC peer device. That is, you create one port channel to the downstream device from the primary vPC peer device and you create another port channel to the downstream device from the secondary peer device.

Note We recommend that you configure the ports on the downstream devices that connect to a host or a network device that is not functioning as a switch or a bridge as STP edge ports. See the *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide, Release 5.x*, for more information on STP port types.

Finally, working on each vPC peer device, you assign a vPC number to the port channel that connects to the downstream device. You will experience minimal traffic disruption when you are creating vPCs. To simplify the configuration, you can assign the vPC ID number to every port channel to be the same as the port channel itself (that is, vPC ID 10 for port channel 10).

Note The vPC number that you assign to the port channel connecting to the downstream device from the vPC peer device *must* be identical on *both* vPC peer devices.

Moving Other Port Channels into a vPC

Note You must attach a downstream device using a port channel to both vPC peer devices.

To connect to the downstream device, you create a port channel to the downstream device from the primary vPC peer device and you create another port channel to the downstream device from the secondary peer device. Finally, working on each vPC peer device, you assign a vPC number to the port channel that connects to the downstream device. You will experience minimal traffic disruption when you are creating vPCs.

Configuring vPC Peer Links and Links to the Core on a Single Module

Note We recommend that you configure the vPC peer links on dedicated ports of different modules to reduce the possibility of a failure. For the best resiliency scenario, use at least two modules.

Beginning with Cisco NX-OS Release 4.2, if you must configure all the vPC peer links and core-facing interfaces on a single module, you should configure, using the command line interface, a track object and a track list that is associated with the Layer 3 link to the core and on all vPC peer links on both vPC peer devices. You use this configuration to avoid dropping traffic if that particular module goes down because when all the tracked objects on the track list go down, the system does the following:

- Stops the vPC primary peer device sending peer-keepalive messages which forces the vPC secondary peer device to take over.
- Brings down all the downstream vPCs on that vPC peer device, which forces all the traffic to be rerouted in the access switch toward the other vPC peer device.

Once you configure this feature and if the module fails, the system automatically suspends all the vPC links on the primary vPC peer device and stops the peer-keepalive messages. This action forces the vPC secondary device to take over the primary role and all the vPC traffic to go to this new vPC primary device until the system stabilizes.

Create a track list that contains all the links to the core and all the vPC peer links as its object. Enable tracking for the specified vPC domain for this track list. Apply this same configuration to the other vPC peer device. See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 5.x*, for information about configuring object tracking and track lists.

See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 5.x*, for information on configuring object tracking.

Note This example uses boolean OR in the track list and will force all traffic to the vPC peer device only for a complete module failure. If you want to trigger a switchover when any core interface or peer-link goes down, use a boolean AND in the track list below.

A vPC deployment with a single Cisco Nexus 7000 Series M132XP-12 module or M108XP-12 module, where the L3 core uplinks and vPC peer-link interfaces are localized on the same module, is vulnerable to access layer isolation if the 10-Gbps module fails on the primary vPC (vPC member ports are defined on both 1-Gbps line cards and on 10-Gbps line card).

To configure a track list to switch over vPC to the remote peer when all related interfaces on a single module fail, follow these steps:

Step 1 Configure track objects on an interface (Layer 3 to core) and on a port channel (vPC peer link).

```
n7k-1(config-if)# track 35 interface ethernet 8/35 line-protocol
n7k-1(config-track)# track 23 interface ethernet 8/33 line-protocol
n7k-1(config)# track 55 interface port-channel 100 line-protocol
```

Step 2 Create a track list that contains all the interfaces in the track list using the Boolean OR to trigger when all objects fail.

```
n7k-1(config)# track 44 list boolean OR
n7k-1(config-track)# object 23
n7k-1(config-track)# object 35
n7k-1(config-track)# object 55
n7k-1(config-track)# end
```

Step 3 Add this track object to the vPC domain:

```
n7k-1(config)# vpc domain 1
n7k-1(config-vpc-domain)# track 44
```

Step 4 This example shows how to display the track object using the **show vpc brief** command.

```
n7k-1# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link
vPC domain id                     : 1
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status: success
vPC role                          : secondary
Number of vPCs configured         : 52
Track object                      : 44

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po100  up     1-5,140

vPC status
----------------------------------------------------------------------------
id   Port   Status Consistency Reason                     Active vlans
--   ----   ------ ----------- -------------------------- ------------
1    Po1    up     success     success                    1-5,140
```

This example shows how to display information about the track objects using the **show track brief** command.

```
n7k-1# show track brief
Track Type         Instance                                Parameter      State  Last
Change
23    Interface    Ethernet8/33                            Line Protocol  UP     00:03:05
35    Interface    Ethernet8/35                            Line Protocol  UP     00:03:15
```

```
44    List          -----                                          Boolean
or     UP     00:01:19
55    Interface     port-channel100                      Line Protocol  UP    00:00:34
```

vPC Interactions with Other Features

This section includes the following topics:

vPC and LACP

LACP uses the system MAC address of the vPC domain to form the LACP Aggregation Group (LAG) ID for the vPC. (See Chapter 6, “Configuring Port Channels,” for information on LAG-ID and LACP.)

You can use LACP on all the vPC port channels, including those channels from the downstream device. We recommend that you configure LACP with active mode on the interfaces on each port channel on the vPC peer devices. This configuration allows you to more easily detect compatibility between devices, unidirectional links, and multihop connection, and provides dynamic reaction to run-time changes and link failures.

With M- series modules and LACP, vPC peer link supports 16 LACP interfaces: 8 active links and 8 hot standby links. You can configure 16 LACP links on the downstream vPC channel: 8 active links and 8 hot standby links. If you configure the port channels without using LACP, you can have only 8 links in each channel. With F-Series line-cards, vPC peer link and downstream vPC channels support upto 16 active LACP links. You can have 16 links in each channel even if the port channels are not configured using LACP.

We recommend that you manually configure the system priority on the vPC peer link devices to ensure that the vPC peer link devices have a higher LACP priority than the downstream connected devices. A lower numerical value system priority means a higher LACP priority.

Note When manually configuring the system priority, you must ensure that you assign the same priority value on both vPC peer devices. If the vPC peer devices have different system priority values, vPC will not come up.

vPC Peer Links and STP

Although vPCs provides a loop-free Layer 2 topology, Spanning Tree Protocol (STP) is still required to provide a “fail-safe” mechanism to protect against any incorrect or defective cabling or possible misconfiguration. When you first bring up vPC, STP reconverges. STP treats the vPC peer link as a special link and always includes the vPC peer link in the STP active topology.

We recommend that you set all the vPC peer link interfaces to the STP network port type so that Bridge Assurance is automatically enabled on all vPC peer links. We also recommend that you do not enable any of the STP enhancement features on vPC peer links. It will not cause any problems if the STP enhancements are already configured, but you need not configure these.

When you are running both MST and Rapid PVST+, ensure that the PVST simulation feature is correctly configured.

See the *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide, Release 5.x*, for information on STP enhancement features and PVST simulation.

Note You must configure a list of parameters to be identical on the vPC peer devices on both sides of the vPC peer link. See the "Compatibility Parameters for vPC Interfaces" section on page 7-15 for information about these required matched settings.

STP is distributed; that is, the protocol continues running on both vPC peer devices. However, the configuration on the vPC peer device elected as the primary device controls the STP process for the vPC interfaces on the secondary vPC peer device.

The primary vPC device synchronizes the STP state on the vPC secondary peer device using Cisco Fabric Services over Ethernet (CFSoE). See the "CFSoE" section on page 7-27 for information on CFSoE.

The STP process for vPC also relies on the periodic keepalive messages to determine when one of the connected devices on the peer link fails. See the "Peer-Keepalive Link and Messages" section on page 7-11 for information on these messages.

The vPC manager performs a proposal/handshake agreement between the vPC peer devices that set the primary and secondary devices and coordinates the two devices for STP. The primary vPC peer device then controls the STP protocol on both the primary and secondary devices. We recommend that you configure the primary vPC peer device as the STP primary root device and configure the secondary VPC device to be the STP secondary root device.

If the primary vPC peer device fails over to the secondary vPC peer device, there is no change in the STP topology.

The BPDUs uses the MAC address set for the vPC for the STP bridge ID in the designated bridge ID field. The vPC primary device sends these BPDUs on the vPC interfaces.

You must configure both ends of vPC peer link with the identical STP configuration for the following parameters:

- STP global settings:
  - STP mode
  - STP region configuration for MST
  - Enable/disable state per VLAN
  - Bridge Assurance setting
  - Port type setting
  - Loop Guard settings
- STP interface settings:
  - Port type setting
  - Loop Guard
  - Root Guard

Note If any of these parameters are misconfigured, the Cisco NX-OS software suspends all interfaces in the vPC. Check the syslog and enter the **show vpc brief** command to see if the vPC interfaces are suspended.

Ensure that the following STP interface configurations are identical on both sides of the vPC peer links or you may see unpredictable behavior in the traffic flow:

- BPDU Filter
- BPDU Guard
- Cost
- Link type
- Priority
- VLANs (PVRST+)

Note Display the configuration on both sides of the vPC peer link to ensure that the settings are identical.

You can use the **show spanning-tree** command to display information about the vPC, when that feature is enabled. See the *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide, Release 5.x,* for an example.

Note We recommend that you configure the ports on the downstream devices as STP edge ports. You should configure all host ports connected to a switch as STP edge ports. (See the *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide, Release 5.x*, for more information on STP port types.)

Note When migrating from vPC+ to vPC topology, the **spanning-tree bridge-assurance** command should be enabled on vPC peer-link for the transition to work.

vPC Peer Switch

The vPC peer switch feature was added to Cisco NX-OS Release 5.0(2) to address performance concerns around STP convergence. This feature allows a pair of Cisco Nexus 7000 Series devices to appear as a single STP root in the Layer 2 topology. This feature eliminates the need to pin the STP root to the vPC primary switch and improves vPC convergence if the vPC primary switch fails.

To avoid loops, the vPC peer link is excluded from the STP computation. In vPC peer switch mode, STP BPDUs are sent from both vPC peer devices to avoid issues related to STP BPDU timeout on the downstream switches, which can cause traffic disruption.

This feature can be used with the pure peer switch topology in which the devices all belong to the vPC.

Note Peer-switch feature is supported on networks that use vPC and STP-based redundancy is not supported. If the vPC peer-link fail in a hybrid peer-switch configuration, you can lose traffic. In this scenario, the vPC peers use the same STP root ID as well same bridge ID. The access switch traffic is split in two with half going to the first vPC peer and the other half to the second vPC peer. With the peer link failed, there is no impact on north/south traffic but east-west traffic will be lost (black-holed).

See the *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide, Release 5.x*, for information on STP enhancement features and Rapid PVST+.

vPC and ARP or ND

A feature was added to Cisco NX-OS Release 4.2(6) to address table synchronization across vPC peers using the reliable transport mechanism of the Cisco Fabric Service over Ethernet (CFSoE) protocol. You must enable the **ip arp synchronize** and **ipv6 nd synchronize** commands to support faster convergence of address tables between the vPC peers. This convergence is designed to overcome the delay involved in ARP table restoration for IPv4 or ND table restoration for IPv6 when the peer link port channel flaps or when a vPC peer comes back online.

vPC Multicast—PIM, IGMP, and IGMP Snooping

Note The Cisco NX-OS software for the Nexus 7000 Series devices does not support PIM SSM or BIDR on vPC. The Cisco NX-OS software fully supports PIM ASM on vPC.

The software keeps the multicast forwarding state synchronized on both of the vPC peer devices. The IGMP snooping process on a vPC peer device shares the learned Group information with the other vPC peer device through the vPC peer link; the multicast states are always synchronized on both vPC peer devices. The PIM process in vPC mode ensures that only one of the vPC peer devices forwards the multicast traffic to the receivers.

Each vPC peer is a Layer 2 or Layer 3 device. Multicast traffic flows from only one of the vPC peer devices. You may see duplicate packets in the following scenarios:

- Orphan hosts
- When the Source and Receivers are in the Layer 2 vPC cloud in different VLANs with multicast routing enabled and a vPC member link goes down.

You may see negligible traffic loss in the following scenarios:

- When you reload the vPC peer device that is forwarding the traffic.
- When you restart PIM on the vPC peer device that is forwarding the traffic.

Ensure that you dual-attach all Layer 3 devices to both vPC peer devices. If one vPC peer device goes down, the other vPC peer device continues to forward all multicast traffic normally.

See the *Cisco Nexus 7000 Series NX-OS Interfaces Command Reference, Release 5.x*, for information on commands that display information on vPC and multicast.

The following discusses vPC PIM and vPC IGMP/IGMP snooping:

- vPC PIM—The PIM process in vPC mode ensures that only one of the vPC peer devices forwards multicast traffic. The PIM process in vPC mode synchronizes the source state with both vPC peer devices and elects with vPC peer device forwards the traffic.
- vPC IGMP/IGMP snooping—The IGMP process in vPC mode synchronizes the DR information on both vPC peer devices. There is a dual-DR concept for IGMP when you are in vPC mode, which is not available when not in vPC mode, that has both vPC peer devices maintain the multicast group information between the peers.

You should enable or disable IGMP snooping identically on both vPC peer devices, and all the feature configurations should be identical. IGMP snooping is on by default.

Note The following commands are not supported in vPC mode:

```
ip pim spt-threshold infinity
ip pim use-shared-tree-only
```

See the *Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide, Release 5.x,* for more information on multicasting.

Multicast PIM Dual DR (proxy DR)

By default, a multicast router sends PIM joins upstream only if it has interested receivers. These interested receivers can either be IGMP hosts (they communicate through IGMP reports) or other multicast routers (they communicate through PIM joins).

In the Cisco NX-OS vPC implementation (in non-F2 mode), PIM works in Dual designated router (DR) mode. That is, if a vPC device is a DR on a vPC SVI OIF, its peer automatically assumes the proxy DR role. IGMP adds an OIF (the report is learned on that OIF) to the forwarding if the OIF is a DR. With Dual DR, both vPC devices have an identical (*,G) entry with respect to the vPC SVI OIFs, as displayed in the following example:

```
VPC Device1:
------------
(*,G)
   oif1 (igmp)

VPC Device2:
------------
(*,G)
   oif1 (igmp)
```

IP PIM PRE-BUILD SPT

When the multicast source is in a Layer 3 cloud (outside the vPC domain), one vPC peer is elected as the forwarder for the source. This forwarder election is based on the metrics to reach the source. If there is a tie, the vPC primary is chosen as the forwarder. Only the forwarder has the vPC OIFs in its associated (S,G) and the nonforwarder (S,G) has 0 OIFs. Therefore, only the forwarder sends PIM (S,G) joins towards the source, as displayed in the following example:

```
VPC Device1 (say this is Forwarder for Source 'S'):
------------
(*,G)
   oif1 (igmp)

(S,G)
   oif1 (mrib)

VPC Device2:
------------
(*,G)
   oif1 (igmp)

(S,G)
   NULL
```

In case of a failure (for example, a Layer 3 RPF link on the forwarder becomes inoperational or the forwarder gets reloaded), if the current nonforwarder ends up becoming the forwarder, it has to start sending PIM joins for (S,G) toward the source to pull the traffic. Depending upon the number of hops to reach the source, this operation might take some time (PIM is a hop by hop protocol).

To eliminate this issue and get better convergence, use the **ip pim pre-build-spt** command. This command enables PIM send joins even if the multicast route has 0 OIFs. In a vPC device, the nonforwarder sends PIM (S,G) joins upstream towards the source. The downside is that the link bandwidth upstream from the nonforwarder gets used for the traffic that is ultimately dropped by it, but the benefits that result with better convergence far outweigh the link bandwidth usage. Therefore, we recommend that all vPC customers use this command.

vPC Peer Links and Routing

The First Hop Routing Protocols (FHRP) interoperate with vPCs. The Hot Standby Routing Protocol (HSRP), Gateway Load Balancing Protocol (GLBP), and Virtual Router Redundancy Protocol (VRRP) all interoperate with vPCs. We recommend that you dual-attach all Layer 3 devices to both vPC peer devices.

The primary FHRP device responds to ARP requests, even though the secondary vPC device forwards the data traffic.

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 link is in standby or in listen state. You can configure lower and upper thresholds to prevent the interface from going up and down.

VRRP acts similarly to HSRP when running on vPC peer devices. You should configure VRRP the same way that you configure HSRP. For GLBP, the forwarders on both vPC peer devices forward traffic.

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

We recommend that you configure routing adjacency between the two vPC peer devices to act as a backup routing path. If one vPC peer device looses Layer 3 uplinks, vPC can redirect the routed traffic to the other vPC peer device and leverage its active Layer 3 uplinks.

You can configure the inter-switch link for backup routing path in the following ways:

- Create a Layer 3 link between the 2 vPC peer devices.
- Use the non-VPC VLAN trunk with a dedicated VLAN interface.
- Use vPC peer link with a dedicated VLAN interface.

We do not recommend configuring the burnt-in MAC address option (use-bia) for HSRP or manually configuring 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 on vPCs. When you are configuring custom MAC addresses, you must configure the same MAC address on both vPC peer devices.

Beginning with Cisco NX-OS Release 4.2(1), you can configure a restore timer that will delay the vPC coming back up until after the peer adjacency forms and the VLAN interfaces are back up. This feature avoids packet drops when the routing tables may not be converged before the vPC is once again passing traffic.

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

To delay the VLAN interfaces on the restored vPC peer device from coming up, use the **interfaces-vlan** option to the **delay restore** command.

See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 5.x*, for more information on FHRPs and routing.

CFSoE

The Cisco Fabric Services over Ethernet (CFSoE) is a reliable state transport mechanism that is used to synchronize the actions of the vPC peer devices. CFSoE carries messages and packets for many features linked with vPC, such as STP and IGMP. Information is carried in CFS/CFSoE protocol data units (PDUs).

When you enable the vPC feature, the device automatically enables CFSoE, and you do not have to configure anything. CFSoE distributions for vPCs do not need the capabilities to distribute over IP or the CFS regions. You need not configure anything for the CFSoE feature to work correctly on vPCs.

The CFSoE transport is local to each VDC.

You can use the **show mac address-table** command to display the MAC addresses that CFSoE synchronizes for the vPC peer link.

Note Do not enter the **no cfs eth distribute** or the **no cfs distribute** command. You must enable CFSoE for vPC functionality. If you do enter either of these commands with vPC enabled, the system displays an error message.

When you enter the **show cfs application** command, the output displays "Physical-eth," which shows the applications that are using CFSoE.

Cisco Fabric Services also transports data over TCP/IP. See the *Cisco Nexus 7000 Series NX-OS System Management Configuration Guide, Release 5.x*, for more information on CFS over IP.

Note The software does not support CFS regions.

vPC and Orphan Ports

When a device that is not vPC-capable connects to each peer, the connected ports are known as orphan ports because they are not members of a vPC. The device's link to one peer will be active (forwarding) and the other link will be standby (blocking) due to STP.

In case of a peer link failure or restoration, an orphan port's connectivity may be bound to the vPC failure or restoration process. For example, if a device's active orphan port connects to the secondary vPC peer, the device will lose any connections through the primary peer upon a peer link failure and the resulting suspending of vPC ports by the secondary peer. If the secondary peer were to also suspend the active orphan port, the device's standby port would become active and would provide a connection to the primary peer, restoring connectivity. Beginning with Cisco NX-OS Release 5.2(1), you can configure in the CLI that specific orphan ports are suspended by the secondary peer when it suspends its vPC ports, and are restored when the vPC is restored.

Virtualization Support

Note You must provision separate vPC peer link and peer-keepalive link infrastructures for each vPC domain deployed in a given VDC. Consolidating a vPC pair (two vPC peer devices of the same domain) in two VDCs of the same physical device is not supported.

All ports in a given vPC must be in the same VDC. This version of the software supports only one vPC domain per VDC. You can use the numbers from 1 to 4096 in each VDC to number the vPC and you can reuse these vPC numbers in a different VDC.

Note See the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 5.x*, for complete information on VDCs and assigning resources.

vPC Recovery After an Outage

In a data center outage, both of the Cisco Nexus 7000 Series devices that comprise a vPC get reloaded. Occasionally only one of the peers can be restored. With no functioning peer-keepalive or peer link, the vPC cannot function normally, but depending on your Cisco NX-OS release, a method may be available to allow vPC services using only the local ports of the functional peer.

Restore on Reload

Note Beginning with Cisco NX-OS Release 5.2(1), the **reload restore** command and method is deprecated. We recommend that you use the **auto-recovery** command and method.

Beginning with Cisco NX-OS Release 5.0(2), you can configure the Cisco Nexus 7000 Series device to restore vPC services when its peer fails to come online by using the **reload restore** command. You must save this setting in the startup configuration. On reload, the Cisco NX-OS software starts a user-configurable timer (the default is 240 seconds). If the peer link port comes up physically or if the peer-keepalive is functional, the timer is stopped and the device waits for the peer adjacency to form.

If at timer expiration no peer-keepalive or peer link up packets were received, the Cisco NX-OS software assumes the primary STP role and the primary LACP role. The software reinitializes the vPCs, bringing up its local ports. Because there are no peers, the consistency check is bypassed for the local vPC ports. The device elects itself to be STP primary regardless of its role priority, and also acts as the master for LACP port roles.

Autorecovery

Beginning with Cisco NX-OS Release 5.2(1), you can configure the Cisco Nexus 7000 Series device to restore vPC services when its peer fails to come online by using the **auto-recovery** command. You must save this setting in the startup configuration. On reload, if the peer link is down and three consecutive peer-keepalive messages are lost, the secondary device assumes the primary STP role and the primary LACP role. The software reinitializes the vPCs, bringing up its local ports. Because there are no peers, the consistency check is bypassed for the local vPC ports. The device elects itself to be STP primary regardless of its role priority and also acts as the master for LACP port roles.

vPC Peer Roles After Recovery

When the other peer device completes its reload and adjacency forms, the following process occurs:

1. The first vPC peer maintains its current role to avoid any transition reset to other protocols. The peer accepts the other available role.
2. When an adjacency forms, consistency checks are performed and appropriate actions are taken.

High Availability

During an In-Service Software Upgrade (ISSU), the software reload process on the first vPC device locks its vPC peer device using CFS messaging over the vPC communications channel. Only one device at a time is upgraded. When the first device has completed its upgrade, it unlocks its peer device. The second device then performs the upgrade process, locking the first device as it does so. During the upgrade, the two vPC devices will temporarily be running different releases of Cisco NX-OS, however the system functions correctly because of its backward compatibility support.

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

Licensing Requirements for vPCs

The following table shows the licensing requirements for this feature:

Product	License Requirement
Cisco NX-OS	vPC requires no license. Any feature not included in a license package is bundled with the Cisco NX-OS system images and is provided at no extra charge to you. For a complete explanation of the Cisco NX-OS licensing scheme, see the *Cisco NX-OS Licensing Guide*.

Guidelines and Limitations

vPC has the following configuration guidelines and limitations:

- VPC peers can only operate dissimilar versions of NX-OS software during the upgrade or downgrade process.
- VPC peers running different version out of upgrade/downgrade period is not supported.
- All ports for a given vPC must be in the same VDC.
- You must enable vPCs before you can configure them.
- You must configure the peer-keepalive link and messages before the system can form the vPC peer link.
- Only Layer 2 port channels can be in vPCs.
- You must configure both vPC peer devices; the configuration is not sent from one device to the other.
- To configure Multi-layer (back-to-back) vPCs, you must assign unique vPC domain ID for each respective vPC.
- Check that the necessary configuration parameters are compatible on both sides of the vPC peer link. See the "Compatibility Parameters for vPC Interfaces" section on page 7-15 for information on compatibility recommendations.
- You may experience minimal traffic disruption while configuring vPCs.
- The software does not support BIDR PIM or SSM on vPCs.
- The software does not support DHCP snooping, DAI, or IPSG in a vPC environment; DHCP Relay is supported.

- The software does not support CFS regions.
- Port security is not supported on port channels.
- We recommend that you configure all the port channels in the vPC using LACP with the interfaces in active mode.
- Configure a separate Layer 3 link for routing from the vPC peer devices, rather than using a VLAN network interface for this purpose.
- Back-to-back, multi-layer vPC topologies require unique Domain IDs on each respective vPC.
- When using vPC, we recommend that you use default timers for FHRP (HSRP, VRRP, GLBP), and PIM configurations. There is no advantage in convergence times when using aggressive timers in vPC configurations.
- If you configure OSPF in a vPC environment, use the following timer commands in router configuration mode on the core switch to ensure fast OSPF convergence when a vPC peer link is shut down:

```
switch (config-router)# timers throttle spf 1 50 50
switch (config-router)# timers lsa-arrival 10
```

  See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 5.x*, for further details about OSPF.
- BFD for HSRP is not supported in a vPC environment.
- The STP port cost is fixed to 200 in a vPC environment.
- A single vPC domain between two VDCs on the same physical Cisco Nexus 7000 device is not supported.
- Jumbo frames are enabled by default on the vPC peer link.

Default Settings

Table 7-2 lists the default settings for vPC parameters.

Table 7-2 Default vPC Parameters

Parameters	Default
vPC system priority	32667
vPC peer-keepalive message	Disabled
vPC peer-keepalive interval	1 second
vPC peer-keepalive timeout	5 seconds
vPC peer-keepalive UDP port	3200

Note

Configuring vPCs

Note You must use these steps on both devices on either side of the vPC peer link. You configure both of the vPC peer devices with these procedures.

This section describes how to configure vPCs using the command-line interface (CLI) and includes the following topics:

Note If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

Enabling vPCs

You must enable the vPC functionality before you can configure and use vPCs.

BEFORE YOU BEGIN

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

SUMMARY STEPS

1. **configure terminal**
2. **feature vpc**
3. **exit**

4. **(Optional) show feature**
5. **(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	**feature vpc** **Example:** switch(config)# feature vpc	Enables vPCs on the device.
Step 3	**exit** **Example:** switch(config)# exit switch#	Exits configuration mode.
Step 4	**show feature** **Example:** switch# show feature	(Optional) Displays which features are enabled on the device.
Step 5	**copy running-config startup-config** **Example:** switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to enable the vPC feature:

```
switch# configure terminal
switch(config)# feature vpc
switch(config)#
```

Disabling vPCs

Note When you disable the vPC functionality, the device clears all the vPC configurations.

BEFORE YOU BEGIN

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

SUMMARY STEPS

1. **configure terminal**
2. **no feature vpc**
3. **exit**
4. **(Optional) show feature**

5. **(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 feature vpc** **Example:** switch(config)# no feature vpc	Disables vPCs on the device.
Step 3	**exit** **Example:** switch(config)# exit switch#	Exits configuration mode.
Step 4	**show feature** **Example:** switch# show feature	(Optional) Displays which features are enabled on the device.
Step 5	**copy running-config startup-config** **Example:** switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to disable the vPC feature:

```
switch# configure terminal
switch(config)# no feature vpc
switch(config)#
```

Creating a vPC Domain and Entering the vpc-domain Mode

You can create a vPC domain and put the vPC peer link port channels into the identical vPC domain on both vPC peer devices. Use a unique vPC domain number throughout a single VDC. This domain ID is used to automatically to form the vPC system MAC address.

You can also use this command to enter the vpc-domain command mode.

BEFORE YOU BEGIN

Ensure that you have enabled the vPC 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. **exit**
4. **(Optional) show vpc brief**
5. **(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	**vpc domain** *domain-id* **Example:** switch(config)# vpc domain 5 switch(config-vpc-domain)#	Creates a vPC domain on the device, and enters the vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1,000.
Step 3	**exit** **Example:** switch(config-vpc-domain)# exit switch(config)#	Exits the vpc-domain configuration mode.
Step 4	**show vpc brief** **Example:** switch# show vpc brief	(Optional) Displays brief information about each vPC domain.
Step 5	**copy running-config startup-config** **Example:** switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to create a vPC domain:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)#
```

This example shows how to enter the vpc-domain command mode to configure an existing vPC domain:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)#
```

Configuring the vPC Keepalive Link and Messages

Note You must configure the vPC peer-keepalive link before the system can form the vPC peer link.

You can configure the destination IP for the peer-keepalive link that carries the keepalive messages. Optionally, you can configure other parameters for the keepalive messages.

Note We recommend that you configure a separate Virtual Routing and Forwarding (VRF) instance and put a Layer 3 port from each vPC peer device into that VRF for the vPC peer-keepalive link. Do not use the peer link itself to send vPC peer-keepalive messages. For information on creating and configuring VRFs, see the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 5.x*.
Ensure that both the source and destination IP addresses use for the peer-keepalive message are unique in your network.

The management port and management VRF are the defaults for these keepalive messages.

BEFORE YOU BEGIN

Ensure that you have enabled the vPC 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. **peer-keepalive destination** *ip address* [**hold-timeout** *secs* | **interval** *msecs* {**timeout** *secs*} | {**precedence** {*prec-value* | **network** | **internet** | **critical** | **flash-override** | **flash** | **immediate** | **priority** | **routine**}} | {**tos** {*tos-value* | **max-reliability** | **max-throughput** | **min-delay** | **min-monetary-cost** | **normal**}} | **tos-byte** *tos-byte-value*} | **source** *ipaddress* | **udp-port** *number* | **vrf** {*name* | **management** | **vpc-keepalive**}]
4. **exit**
5. **(Optional) show vpc statistics**
6. **(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	`vpc domain` *domain-id*
	
	Example:
	`switch(config)# vpc domain 5`
	`switch(config-vpc-domain)#`
		Creates a vPC domain on the device, and enters the vpc-domain configuration mode for configuration purposes.

	Command	Purpose

Step 3	**peer-keepalive destination** *ipaddress* [**hold-timeout** *secs* | **interval** *msecs* {**timeout** *secs*} | {**precedence** {*prec-value* | **network** | **internet** | **critical** | **flash-override** | **flash** | **immediate** **priority** | **routine**}} | **tos** {*tos-value* | **max-reliability** | **max-throughput** | **min-delay** | **min-monetary-cost** | **normal**}} |**tos-byte** *tos-byte-value*} | **source** *ipaddress* | **vrf** {*name* | **management** **vpc-keepalive**}]
	
	Example:
	`switch(config-vpc-domain)# peer-keepalive destination 172.28.230.85`
	`switch(config-vpc-domain)#`
		Configures the IPv4 address for the remote end of the vPC peer-keepalive link.
		
		Note The system does not form the vPC peer link until you configure a vPC peer-keepalive link,
		
		The management ports and VRF are the defaults
		
		Note We recommend that you configure a separate VRF and use a Layer 3 port from each vPC peer device in that VRF for the vPC peer-keepalive link. For more information on creating and configuring VRFs, see the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 5.x*.

Step 4	**exit**
	
	Example:
	`switch(config-vpc-domain)# exit`
	`switch(config)#`
		Exits the vpc-domain configuration mode.

Step 5	**show vpc statistics** **Example:** `switch# show vpc statistics`	(Optional) Displays information about the configuration for the keepalive messages.
Step 6	**copy running-config startup-config** **Example:** `switch# copy running-config startup-config`	(Optional) Copies the running configuration to the startup configuration.

For more information on configuring VRFs, see the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 5.x*.

This example shows how to configure the destination and source IP address and VRF for the vPC-peer-keepalive link:

```
switch# configure terminal
switch(config)# feature vpc
switch(config)# vpc domain 100
switch(config-vpc-domain)# peer-keepalive destination 172.168.1.2 source 172.168.1.1 vrf
vpc-keepalive
```

Creating the vPC Peer Link

You create the vPC peer link by designating the port channel that you want on each device as the peer link for the specified vPC domain. We recommend that you configure the Layer 2 port channels that you are designating as the vPC peer link in trunk mode and that you use two ports on separate modules on each vPC peer device for redundancy.

BEFORE YOU BEGIN

Ensure that you have enabled the vPC feature.

Ensure that you are using a Layer 2 port channel.

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

SUMMARY STEPS

1. **configure terminal**
2. **interface port-channel** *channel-number*
3. (Optional) **switchport mode trunk**
4. (Optional) **switchport trunk allowed vlan** *vlan-list*
5. **vpc peer-link**
6. **exit**
7. **(Optional) show vpc brief**
8. **(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	**interface port-channel** *channel-number* **Example:** switch(config)# interface port-channel 20 switch(config-if)#	Selects the port channel that you want to use as the vPC peer link for this device, and enters the interface configuration mode.
Step 3	**switchport mode trunk** **Example:** switch(config-if)# switchport mode trunk	(Optional) Configures this interface in trunking mode.
Step 4	**switchport trunk allowed vlan** *vlan-list* **Example:** switch(config-if)# switchport trunk allowed vlan 1-120,201-3967	(Optional) Configures the allowed VLAN list.
Step 5	**vpc peer-link** **Example:** switch(config-if)# vpc peer-link switch(config-vpc-domain)#	Configures the selected port channel as the vPC peer link, and enters the vpc-domain configuration mode.
Step 6	**exit** **Example:** switch(config-vpc-domain)# exit switch(config)#	Exits the vpc-domain configuration mode.

	Command	Purpose
Step 7	**show vpc brief** **Example:** switch# show vpc brief	(Optional) Displays information about each vPC, including information about the vPC peer link.
Step 8	**copy running-config startup-config** **Example:** switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to configure a vPC peer link:

```
switch# configure terminal
switch(config)# interface port-channel 20
switch(config-if)# switchport mode
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk allowed vlan 1-120,201-3967
switch(config-if)# vpc peer-link
switch(config-vpc-domain)#
```

Configuring the vPC Peer-Gateway

Beginning with Cisco NX-OS Release 4.2(1) and later releases, you can configure vPC peer devices to act as the gateway for packets that are destined to the vPC peer device's MAC address.

Note When you attach a Layer 3 device to a vPC domain, the peering of routing protocols using a VLAN also carried on the vPC peer-link is not supported. If routing protocol adjacencies are needed between vPC peer devices and a generic Layer 3 device, you must use physical routed interfaces for the interconnection. Use of the vPC peer-gateway feature does not change this requirement.

BEFORE YOU BEGIN

Ensure that you have enabled the vPC 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. **peer-gateway**
4. **peer-gateway exclude-vlan** *backup-vlan-id*
5. **exit**
6. **(Optional) show vpc brief**
7. **(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	**vpc domain** *domain-id* **Example:** switch(config-if)# vpc domain 5 switch(config-vpc-domain)#	Creates a vPC domain if it doesn't already exist, and enters the vpc-domain configuration mode.
Step 3	**peer-gateway** **Example:** switch(config-vpc-domain)# peer-gateway Note: ------:: Disable IP redirects on all interface-vlans of this vPC domain for correct operation of this feature ::------	Enables Layer 3 forwarding for packets destined to the peer's gateway MAC address.
Step 4	**peer-gateway exclude-vlan** *backup-vlan-id* **Example:** switch(config-vpc-domain)# peer-gateway exclude-vlan 2	(Optional) Beginning with Cisco NX-OS Release 5.1(3), avoids software switching of transit VLAN traffic in a mixed chassis mode. **Note** See the "vPC Peer-Gateway" section on page 7-13 for more information.
Step 5	**exit** **Example:** switch(config-vpc-domain)# exit switch(config)#	Exits the vpc-domain configuration mode.
Step 6	**show vpc brief** **Example:** switch# show vpc brief	(Optional) Displays information about each vPC, including information about the vPC peer link.
Step 7	**copy running-config startup-config** **Example:** switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Configuring a Graceful Consistency Check

Beginning with Cisco NX-OS Release 5.2(1), you can configure the graceful consistency check feature, which is enabled by default. Unless this feature is enabled, the vPC is completely suspended when a mismatch in a mandatory compatibility parameter is introduced in a working vPC. When this feature is enabled, only the links on the secondary peer device are suspended. See the "Compatibility Parameters for vPC Interfaces" section on page 7-15 for information on consistent configurations on the vPCs.

SUMMARY STEPS

1. **configure terminal**
2. **vpc domain** *domain-id*

3. **graceful consistency-check**
4. **exit**
5. **(Optional) show vpc brief**

DETAILED STEPS

	Command	Purpose
Step 1	**configure terminal** **Example:** switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	**vpc domain** *domain-id* **Example:** switch(config-if)# vpc domain 5 switch(config-vpc-domain)#	Creates a vPC domain if it does not already exist, and enters the vpc-domain configuration mode.
Step 3	**graceful consistency-check** **Example:** switch(config-vpc-domain)# graceful consistency-check	Specifies that only the links on the secondary peer device are suspended when a mismatch is detected in a mandatory compatibility parameter. Use the **no** form of this command to disable the feature.
Step 4	**exit** **Example:** switch(config-vpc-domain)# exit switch(config)#	Exits the vpc-domain configuration mode.
Step 5	**show vpc brief** **Example:** switch# show vpc brief	(Optional) Displays information on the vPCs.

This example shows how to enable the graceful consistency check feature:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)# graceful consistency-check
```

Checking the Configuration Compatibility on a vPC Peer Link

After you have configured the vPC peer link on both vPC peer devices, check that the configurations are consistent on all vPC interfaces. See the "Compatibility Parameters for vPC Interfaces" section on page 7-15 for information on consistent configurations on the vPCs.

SUMMARY STEPS

1. **configure terminal**
2. **(Optional) show vpc consistency-parameters** {**global** | **interface port-channel** *channel-number*}

DETAILED STEPS

	Command	Purpose
Step 1	**configure terminal** **Example:** switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	**show vpc consistency-parameters** {**global** | **interface port-channel** *channel-number*} **Example:** switch(config)# show vpc consistency-parameters global switch(config)#	(Optional) Displays the status of those parameters that must be consistent across all vPC interfaces.

This example shows how to check that the required configurations are compatible across all the vPC interfaces:

```
switch# configure terminal
switch(config)# show vpc consistency-parameters global
switch(config)#
```

Note Messages regarding the vPC interface configuration compatibility are also logged to the syslog.

Moving Other Port Channels into a vPC

Note We recommend that you attach the vPC domain downstream port channel to two devices for redundancy.

To connect to the downstream device, you create a port channel from the downstream device to the primary vPC peer device and you create another port channel from the downstream device to the secondary peer device. Finally, working on each vPC peer device, you assign a vPC number to the port channel that connects to the downstream device. You will experience minimal traffic disruption when you are creating vPCs.

BEFORE YOU BEGIN

Ensure that you have enabled the vPC feature.

Ensure that you are using a Layer 2 port channel.

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

SUMMARY STEPS

1. **configure terminal**
2. **interface port-channel** *channel-number*
3. **vpc** *number*
4. **exit**

5. **(Optional) show vpc brief**
6. **(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	**interface port-channel** *channel-number*
	
	Example:
	`switch(config)# interface port-channel 20`
	`switch(config-if)#`
		Selects the port channel that you want to put into the vPC to connect to the downstream device, and enters the interface configuration mode.

Step 3	**vpc** *number*
	
	Example:
	`switch(config-if)# vpc 5`
	`switch(config-vpc-domain)#`
		Configures the selected port channel into the vPC to connect to the downstream device. You can use any module in the device for these port channels. The range is from 1 and 4096.
		
		Note The vPC number that you assign to the port channel connecting to the downstream device from the vPC peer device *must* be identical on *both* vPC peer devices.

Step 4	**exit**
	
	Example:
	`switch(config-vpc-domain)# exit`
	`switch(config)#`
		Exits the vpc-domain configuration mode.

Step 5	**show vpc brief** **Example:** `switch# show vpc brief`	(Optional) Displays information on the vPCs.
Step 6	**copy running-config startup-config** **Example:** `switch# copy running-config startup-config`	(Optional) Copies the running configuration to the startup configuration.

This example shows how to configure a port channel that will connect to the downstream device:

```
switch# configure terminal
switch(config)# interface port-channel 20
switch(config-if)# vpc 5
switch(config-if
```

Manually Configuring a vPC Domain MAC Address

When you create a vPC domain, the Cisco NX-OS software automatically creates a vPC system MAC address, which is used for operations that are confined to the link-scope, such as LACP. However, you may choose to configure the vPC domain MAC address manually.

BEFORE YOU BEGIN

Ensure that you have enabled the vPC 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. **system-mac** *mac-address*
4. **exit**
5. **(Optional) show vpc role**
6. **(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	**vpc domain** *domain-id*
	
	Example:
	`switch(config)# vpc domain 5`
	`switch(config-vpc-domain)#`
		Enters the vPC domain number that you want to configure. The system enters the vpc-domain configuration mode.

Step 3	**system-mac** *mac-address*
	
	Example:
	`switch(config-vpc-domain)# system-mac 23fb.4ab5.4c4e`
	`switch(config-vpc-domain)#`
		Enters the MAC address that you want for the specified vPC domain in the following format: aaaa.bbbb.cccc.

Step 4	**exit**
	
	Example:
	`switch(config-vpc-domain)# exit`
	`switch(config)#`
		Exits the vpc-domain configuration mode.

Step 5	**show vpc role** **Example:** `switch# show vpc brief`	(Optional) Displays the vPC system MAC address.
Step 6	**copy running-config startup-config** **Example:** `switch# copy running-config startup-config`	(Optional) Copies the running configuration to the startup configuration.

This example shows how to manually configure the vPC domain MAC address:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)# system-mac 13gb.4ab5.4c4e
```

Manually Configuring the System Priority

When you create a vPC domain, the system automatically creates a vPC system priority. However, you can also manually configure a system priority for the vPC domain.

Note We recommend that you manually configure the vPC system priority when you are running LACP to ensure that the vPC peer devices are the primary devices on LACP. When you manually configure the system priority, ensure that you configure the same priority value on both vPC peer devices. If these values do not match, vPC will not come up.

BEFORE YOU BEGIN

Ensure that you have enabled the vPC 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. **system-priority** *priority*
4. **exit**
5. **(Optional) show vpc role**
6. **(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	**vpc domain** *domain-id*
	
	Example:
	`switch(config)# vpc domain 5`
	`switch(config-vpc-domain)#`
		Enters the vPC domain number that you want to configure. The system enters the vpc-domain configuration mode.

Step 3	**system-priority** *priority*
	
	Example:
	`switch(config-vpc-domain)# system-priority 4000`
	`switch(config-vpc-domain)#`
		Enters the system priority that you want for the specified vPC domain. The range of values is from 1 to 65535. The default value is 32667.

Step 4	**exit**
	
	Example:
	`switch(config-vpc-domain)# exit`
	`switch(config)#`
		Exits the vpc-domain configuration mode.

	Command	Purpose
Step 5	**show vpc role** **Example:** switch# show vpc role	(Optional) Displays the vPC system priority.
Step 6	**copy running-config startup-config** **Example:** switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to manually configure the vPC domain system priority:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)# system-priority 4000
```

Manually Configuring the vPC Peer Device Role

By default, the Cisco NX-OS software elects a primary and secondary vPC peer device after you configure the vPC domain and both sides of the vPC peer link. However, you may want to elect a specific vPC peer device as the primary device for the vPC. Then, you would manually configure the role value for the vPC peer device that you want as the primary device to be lower than the other vPC peer device.

vPC does not support role preemption. If the primary vPC peer device fails, the secondary vPC peer device takes over to become operationally the vPC primary device. However, the original operational roles are not restored if the formerly primary vPC comes up again.

BEFORE YOU BEGIN

Ensure that you have enabled the vPC 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. **role priority** *priority*
4. **exit**
5. **(Optional) show vpc role**
6. **(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	**vpc domain** *domain-id* **Example:** switch(config)# vpc domain 5 switch(config-vpc-domain)#	Enters the vPC domain number that you want to configure. The system enters the vpc-domain configuration mode.
Step 3	**role priority** *priority* **Example:** switch(config-vpc-domain)# role priority 4 switch(config-vpc-domain)#	Enters the role priority that you want for the vPC system priority.The range of values is from 1 to 65636, and the default value is 32667. A lower value means this switch has a better chance of being the primary vPC.
Step 4	**exit** **Example:** switch(config-vpc-domain)# exit switch(config)#	Exits the vpc-domain configuration mode.
Step 5	**show vpc role** **Example:** switch# show vpc role	(Optional) Displays the vPC system priority.
Step 6	**copy running-config startup-config** **Example:** switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to manually configure the role priority of the vPC peer device:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)# role priority 4
```

Configuring the Tracking Feature on a Single-Module vPC

Beginning with Cisco NX-OS Release 4.2, if you must configure all the vPC peer links and core-facing interfaces on a single module, you should configure a track object and a track list that is associated with the Layer 3 link to the core and on all the links on the vPC peer link on both primary vPC peer devices. Once you configure this feature and if the primary vPC peer device fails, the system automatically suspends all the vPC links on the primary vPC peer device. This action forces all the vPC traffic to the secondary vPC peer device until the system stabilizes.

You must put this configuration on both vPC peer devices. Additionally, you should put the identical configuration on both vPC peer devices because either device can become the operationally primary vPC peer device.

BEFORE YOU BEGIN

Ensure that you have enabled the vPC feature.

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

Ensure that you have configured the track object and the track list. Ensue that you assign all interfaces that connect to the core and to the vPC peer link to the track-list object on both vPC peer devices.

SUMMARY STEPS

1. **configure terminal**
2. **vpc domain** *domain-id*
3. **track** *track-object-id*
4. **exit**
5. **(Optional) show vpc**
6. **(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	**vpc domain** *domain-id*
	
	Example:
	`switch(config)# vpc domain 5`
	`switch(config-vpc-domain)#`
		Enters the vPC domain number that you want to configure. The system enters the vpc-domain configuration mode.

Step 3	**track** *track-object-id*
	
	Example:
	`switch(config-vpc-domain)# track object 23`
	`switch(config-vpc-domain)#`
		Adds the previously configured track-list object with its associated interfaces to the vPC domain. See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 5.x*, for information on configuring object tracking and track lists.

Step 4	**exit**
	
	Example:
	`switch(config-vpc-domain)# exit`
	`switch(config)#`
		Exits the vpc-domain configuration mode.

Step 5	**show vpc brief** **Example:** `switch# show vpc brief`	(Optional) Displays information on the tracked objects.
Step 6	**copy running-config startup-config** **Example:** `switch# copy running-config startup-config`	(Optional) Copies the running configuration to the startup configuration.

This example shows how to put the previously configured track-list object into the vPC domain on the vPC peer device:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)# track object 5
```

Configuring for Recovery After an Outage

In the event of an outage, vPC waits for peer adjacency to form on a switch reload. This situation can result in an unacceptably long service disruption. You can configure the Cisco Nexus 7000 Series device to restore vPC services when its peer fails to come on line.

Configuring Reload Restore

Note Beginning with Cisco NX-OS Release 5.2(1), the **reload restore** command and procedure described in this section is deprecated. We recommend that you use the **auto-recovery** command and procedure described in the "Configuring Autorecovery" section on page 7-50.

Beginning with Cisco NX-OS Release 5.0(2), you can configure the Cisco Nexus 7000 Series device to restore vPC services when its peer fails to come online by using the **reload restore** command.

BEFORE YOU BEGIN

Ensure that you have enabled the vPC 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. **reload restore** [**delay** *time-out*]
4. **exit**
5. **(Optional) show running-config vpc**
6. **(Optional) show vpc consistency-parameters interface port-channel** *number*
7. **(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	**vpc domain** *domain-id* **Example:** switch(config)# vpc domain 5 switch(config-vpc-domain)#	Enters the vPC domain number that you want to configure. The system enters the vpc-domain configuration mode.
Step 3	**reload restore** [**delay** *time-out*] **Example:** switch(config-vpc-domain)# reload restore	Configures the vPC to assume its peer is not functional and to bring up the vPC. The default delay is 240 seconds. You can configure a time-out delay from 240 to 3600 seconds. Use the **no** form of the command to reset the vPC to the standard behavior.
Step 4	**exit** **Example:** switch(config-vpc-domain)# exit switch(config)#	Exits the vpc-domain configuration mode.
Step 5	**show running-config vpc** **Example:** switch# show running-config vpc	(Optional) Displays information about the vPC, specifically the reload status.
Step 6	**show vpc consistency-parameters interface port-channel** *number* **Example:** switch# show vpc consistency-parameters interface port-channel 1	(Optional) Displays information about the vPC consistency parameters for the specified interface.
Step 7	**copy running-config startup-config** **Example:** switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration. **Note** To ensure the reload feature is enabled, you should perform this step.

This example shows how to set the vPC reload restore feature and save it in the switch startup configuration:

```
switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# vpc domain 5
switch(config-vpc-domain)# reload restore
Warning:
 Enables restoring of vPCs in a peer-detached state after reload, will wait for 240
seconds (by default) to determine if peer is un-reachable
switch(config-vpc-domain)# exit
switch(config)# exit
switch# copy running-config startup-config
switch# show running-config vpc
```

```
!Command: show running-config vpc
!Time: Wed Mar 24 18:43:54 2010

version 5.0(2)
feature vpc

logging level vpc 6
vpc domain 5
  reload restore
```

This example shows how to examine the consistency parameters:

```
switch# show vpc consistency-parameters interface port-channel 1
    Legend:
        Type 1 : vPC will be suspended in case of mismatch
Name                        Type  Local Value   Peer Value
-------------               ----  -----------   ---------------
STP Port Type                 1    Default              -
STP Port Guard                1    None                 -
STP MST Simulate PVST         1    Default              -
mode                          1    on                   -
Speed                         1    1000 Mb/s            -
Duplex                        1    full                 -
Port Mode                     1    trunk                -
Native Vlan                   1    1                    -
MTU                           1    1500                 -
Allowed VLANs                 -    1-3967,4048-4093
Local suspended VLANs         -    -                    -
```

Configuring Autorecovery

Beginning with Cisco NX-OS Release 5.2(1), you can configure the Cisco Nexus 7000 Series device to restore vPC services when its peer fails to come online by using the **auto-recovery** command.

BEFORE YOU BEGIN

Ensure that you have enabled the vPC 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. **auto-recovery** [**reload-delay** *time*]
4. **exit**
5. **(Optional) show running-config vpc**
6. **(Optional) show vpc consistency-parameters interface port-channel** *number*
7. **(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	**vpc domain** *domain-id* **Example:** switch(config)# vpc domain 5 switch(config-vpc-domain)#	Enters the vPC domain number that you want to configure. The system enters the vpc-domain configuration mode.
Step 3	**auto-recovery** [**reload-delay** *time*] **Example:** switch(config-vpc-domain)# auto-recovery	Configures the vPC to assume its peer is not functional and to bring up the vPC, and specifies the time to wait after reload to restore the vPC. The default delay is 240 seconds. You can configure a delay from 240 to 3600 seconds. Use the **no** form of the command to reset the vPC to the standard behavior.
Step 4	**exit** **Example:** switch(config-vpc-domain)# exit switch(config)#	Exits the vpc-domain configuration mode.
Step 5	**show running-config vpc** **Example:** switch# show running-config vpc	(Optional) Displays information about the vPC, specifically the reload status.
Step 6	**show vpc consistency-parameters interface port-channel** *number* **Example:** switch# show vpc consistency-parameters interface port-channel 1	(Optional) Displays information about the vPC consistency parameters for the specified interface.
Step 7	**copy running-config startup-config** **Example:** switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration. **Note** To ensure the autorecovery feature is enabled, you should perform this step.

This example shows how to set the vPC autorecovery feature and save it in the switch startup configuration:

```
switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# vpc domain 5
switch(config-vpc-domain)# auto-recovery
Warning:
 Enables restoring of vPCs in a peer-detached state after reload, will wait for 240
seconds to determine if peer is un-reachable
switch(config-vpc-domain)# exit
switch(config)# exit
switch# copy running-config startup-config
```

Configuring the Suspension of Orphan Ports

When a device that is not vPC-capable connects to each peer, the connected ports are known as orphan ports because they are not members of a vPC. Beginning with Cisco NX-OS Release 5.2(1), you can explicitly declare phyiscal interfaces as orphan ports to be suspended (shut down) by the secondary peer when it suspends its vPC ports in response to a peer link or peer-keepalive failure. The orphan ports are restored when the vPC is restored.

Note You can configure vPC orphan port suspension only on physical ports, not on port channel member ports.

BEFORE YOU BEGIN

Ensure that you have enabled the vPC feature.

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

SUMMARY STEPS

1. **configure terminal**
2. **show vpc orphan-ports**
3. **interface** *type slot/port*
4. **vpc orphan-ports suspend**
5. **exit**
6. **(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	**show vpc orphan-ports** **Example:** `switch# show vpc orphan-ports`	(Optional) Displays a list of the orphan ports.

Step 3	**interface** *type slot/port*
	
	Example:
	`switch(config)# interface ethernet 3/1`
	`switch(config-if)#`
		Specifies an interface to configure, and enters interface configuration mode.

Step 4	**vpc orphan-ports suspend** **Example:** `switch(config-if)# vpc orphan-ports suspend`	Configures the selected interface as a vPC orphan port to be suspended by the secondary peer in case of vPC failure.

	Command	Purpose
Step 5	**exit** **Example:** switch(config-if)# exit switch(config)#	Exits the interface configuration mode.
Step 6	**copy running-config startup-config** **Example:** switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to configure an interface as a vPC orphan port to be suspended by the secondary peer in case of vPC failure:

```
switch# configure terminal
switch(config)# interface ethernet 3/1
switch(config-if)# vpc orphan-ports suspend
```

Configuring the vPC Peer Switch

You can configure the Cisco Nexus 7000 Series device to make a pair of vPC devices appear as a single STP root in the Layer 2 topology. This section includes the following topics:

- Configuring a Pure vPC Peer Switch Topology, page 7-53
- Configuring a Hybrid vPC Peer Switch Topology, page 7-55

Configuring a Pure vPC Peer Switch Topology

You can configure a pure vPC peer switch topology using the **peer-switch** command and then you set the best possible (lowest) spanning tree bridge priority value.

BEFORE YOU BEGIN

Ensure that you have enabled the vPC 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. **peer-switch**
4. **spanning-tree vlan** *vlan-range* **priority** *value*
5. **exit**
6. **(Optional) show spanning-tree summary**
7. **(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	**vpc domain** *domain-id* **Example:** switch(config)# vpc domain 5 switch(config-vpc-domain)#	Enters the vPC domain number that you want to configure. The system enters the vpc-domain configuration mode.
Step 3	**peer-switch** **Example:** switch(config-vpc-domain)# peer-switch	Enables the vPC switch pair to appear as a single STP root in the Layer 2 topology. Use the **no** form of the command to disable the peer switch vPC topology.
Step 4	**spanning-tree vlan** *vlan-range* **priority** *value* **Example:** switch(config)# spanning-tree vlan 1 priority 8192	Configures the bridge priority of the VLAN. Valid values are multiples of 4096. The default value is 32768.
Step 5	**exit** **Example:** switch(config-vpc-domain)# exit switch(config)#	Exits the vpc-domain configuration mode.
Step 6	**show spanning-tree summary** **Example:** switch# show spanning-tree summary	(Optional) Displays a summary of the spanning tree port states including the vPC peer switch.
Step 7	**copy running-config startup-config** **Example:** switch# copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

This example shows how to configure a pure vPC peer switch topology:

```
switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# vpc domain 5
switch(config-vpc-domain)# peer-switch
2010 Apr 28 14:44:44 switch %STP-2-VPC_PEERSWITCH_CONFIG_ENABLED: vPC peer-switch
configuration is enabled. Please make sure to configure spanning tree "bridge" priority as
per recommended guidelines to make vPC peer-switch operational.
switch(config-vpc-domain)# exit
switch(config)# spanning-tree vlan 1 priority 8192
```

Configuring a Hybrid vPC Peer Switch Topology

You can configure a hybrid vPC and non-vPC peer switch topology by using the **spanning-tree pseudo-information** command (see the *Cisco Nexus 7000 Series NX-OS Layer 2 Switching Command Reference, Release 5.x*.) to change the designated bridge ID so that it meets the STP VLAN-based load-balancing criteria and then change the root bridge ID priority to a value that is better than the best bridge priority. You then enable the peer switch.

BEFORE YOU BEGIN

Ensure that you have enabled the vPC feature.

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

SUMMARY STEPS

1. **configure terminal**
2. **spanning-tree pseudo-information**
3. **vlan** *vlan-range* **designated priority** *value*
4. **vlan** *vlan-range* **root priority** *value*
5. **vpc domain** *domain-id*
6. **peer-switch**
7. **exit**
8. **(Optional) show spanning-tree summary**
9. **(Optional) opy running-config startup-config**

DETAILED STEPS

	Command	Purpose

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

Step 2	**spanning-tree pseudo-information**
	
	Example:
	`switch(config)# spanning-tree pseudo-information`
	`switch(config-pseudo)#`
		Configures the spanning tree pseudo information.

Step 3	**vlan** *vlan-id* **designated priority** *priority* **Example:** `switch(config-pseudo)# vlan 1 designated priority 8192`	Configures the designated bridge priority of the VLAN. Valid values are multiples of 4096 from 0 to 61440.
Step 4	**vlan** *vlan-id* **root priority** *priority* **Example:** `switch(config-pseudo)# vlan 1 root priority 4096`	Configures the root bridge priority of the VLAN. Valid values are multiples of 4096 from 0 to 61440.

	Command	Purpose

Step 5	**vpc domain** *domain-id*
	
	Example:
	`switch(config)# vpc domain 5`
	`switch(config-vpc-domain)#`
		Enters the vPC domain number that you want to configure. The system enters the vpc-domain configuration mode.

Step 6	**peer-switch** **Example:** `switch(config-vpc-domain)# peer-switch`	Enables the vPC switch pair to appear as a single STP root in the Layer 2 topology. Use the **no** form of the command to disable the peer switch vPC topology.

Step 7	**exit**
	
	Example:
	`switch(config-vpc-domain)# exit`
	`switch(config)#`
		Exits the vpc-domain configuration mode.

Step 8	**show spanning-tree summary** **Example:** `switch# show spanning-tree summary`	(Optional) Displays a summary of the spanning tree port states including the vPC peer switch.
Step 9	**copy running-config startup-config** **Example:** `switch# copy running-config startup-config`	(Optional) Copies the running configuration to the startup configuration.

This example shows how to configure a hybrid vPC peer switch topology:

```
switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# spanning-tree pseudo-information
switch(config-pseudo)# vlan 1 designated priority 8192
switch(config-pseudo)# vlan 1 root priority 4096
switch(config-pseudo)# exit
switch(config)# vpc domain 5
switch(config-vpc-domain)# peer-switch
```

Verifying the vPC Configuration

To display vPC configuration information, perform one of the following tasks:

Command	Purpose
show feature	Displays whether the vPC is enabled or not.
show vpc brief	Displays brief information on the vPCs.
show vpc consistency-parameters	Displays the status of those parameters that must be consistent across all vPC interfaces.
show running-config vpc	Displays running configuration information for vPCs.
show port-channel capacity	Displays how many port channels are configured and how many are still available on the device.
show vpc statistics	Displays statistics on the vPCs.

Command	Purpose
show vpc peer-keepalive	Displays information on the peer-keepalive messages.
show vpc role	Displays the peer status, the role of the local device, the vPC system MAC address and system priority, and the MAC address and priority for the local vPC device.

For detailed information about the fields in the output from these commands, see the *Cisco Nexus 7000 Series NX-OS Interfaces Command Reference, Release 5.x*.

Monitoring vPCs

Use the **show vpc statistics** command to display vPC statistics.

Note This command displays the vPC statistics only for the vPC peer device that you are working on.

Configuration Examples for vPCs

The following example shows how to configure vPC on device A as shown in Figure 7-11:

Figure 7-11 vPC Configuration Example

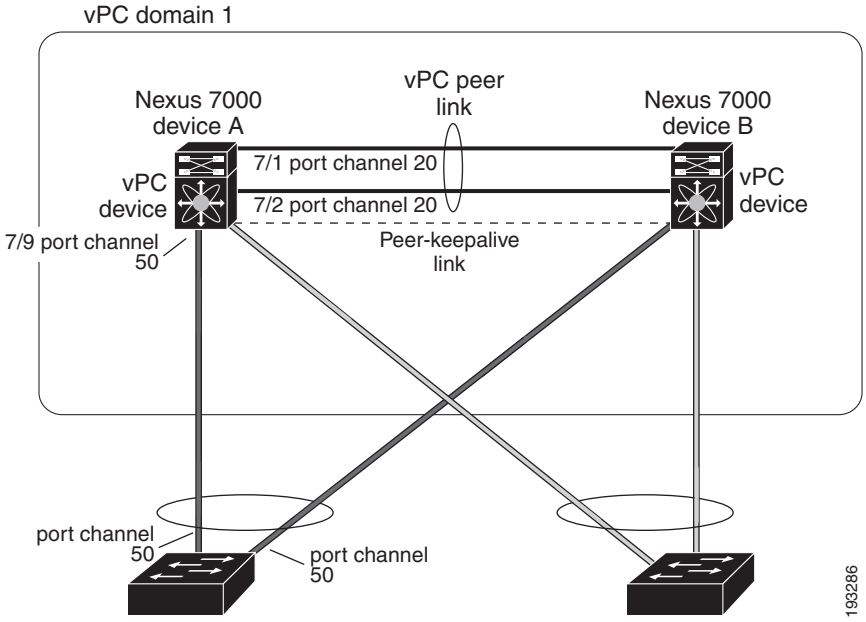

Step 1 Enable vPC and LACP.

```
switch# configure terminal
```

```
switch(config)# feature vPC
switch(config)# feature lacp
```

Step 2 (Optional) Configure one of the interfaces that you want to be peer link in dedicated mode.

```
switch(config)# interface ethernet 7/1, ethernet 7/3, ethernet 7/5. ethernet 7/7
switch(config-if)# shutdown
switch(config-if)# exit
switch(config)# interface ethernet 7/1
switch(config-if)# rate-mode dedicated
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)#
```

Step 3 (Optional) Configure the second, redundant interface that you want to be peer link in dedicated mode.

```
switch(config)# interface ethernet 7/2, ethernet 7/4, ethernet 7/6. ethernet 7/8
switch(config-if)# shutdown
switch(config-if)# exit
switch(config)# interface ethernet 7/2
switch(config-if)# rate-mode dedicated
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)#
```

Step 4 Configure the two interfaces (for redundancy) that you want to be in the peer link to be an active Layer 2 LACP port channel.

```
switch(config)# interface ethernet 7/1-2
switch(config-if)# switchport
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk allowed vlan 1-50
switch(config-if)# switchport trunk native vlan 20
switch(config-if)# channel-group 20 mode active
switch(config-if)# exit
```

Step 5 Create and enable the VLANs.

```
switch(config)# vlan 1-50
switch(config-vlan)# no shutdown
switch(config-vlan)# exit
```

Step 6 Create a separate VRF for the vPC peer-keepalive link and add a Layer 3 interface to that VRF.

```
switch(config)# vrf context pkal
switch(config-vrf)# exit
switch(config)# interface ethernet 8/1
switch(config-if)# vrf member pkal
switch(config-if)# ip address 172.23.145.218/24
switch(config-if)# no shutdown
switch(config-if)# exit
```

Step 7 Create the vPC domain and add the vPC peer-keepalive link.

```
switch(config)# vpc domain 1
switch(config-vpc-domain)# peer-keepalive destination 172.23.145.217 source 172.23.145.218
vrf pkal
switch(config-vpc-domain)# exit
```

Step 8 Configure the vPC peer link.

```
switch(config)# interface port-channel 20
switch(config-if)# switchport mode trunk
```

```
switch(config-if)# switchport trunk allowed vlan 1-50
switch(config-if)# vpc peer-link
switch(config-if)# exit
switch(config)#
```

Step 9 Configure the interface for the port channel to the downstream device of the vPC.

```
switch(config)# interface ethernet 7/9
switch(config-if)# switchport mode trunk
switch(config-if)# allowed vlan 1-50
switch(config-if)# native vlan 20
switch(config-if)# channel-group 50 mode active
switch(config-if)# exit
switch(config)# interface port-channel 50
switch(config-if)# vpc 50
switch(config-if)# exit
switch(config)#
```

Step 10 Save the configuration.

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

Note If you configure the port channel first, ensure that it is a Layer 2 port channel.

Additional References

For additional information related to implementing vPC, see the following sections:

- Related Documents, page 7-60
- Standards, page 7-60
- MIBs, page 7-60

Related Documents

Related Topic	Document Title
Configuring Port Channels	Chapter 6, "Configuring Port Channels,"
Configuring Layer 2 interface	Chapter 3, "Configuring Layer 2 Interfaces,"
Configuring Layer 3 interfaces	Chapter 4, "Configuring Layer 3 Interfaces,"
Shared and dedicated ports	Chapter 2, "Configuring Basic Interface Parameters,"
Command reference	*Cisco Nexus 7000 Series NX-OS Interfaces Command Reference, Release 5.x*
Interfaces	*Cisco DCNM Interfaces Configuration Guide, Release 5.x*
System management	*Cisco Nexus 7000 Series NX-OS System Management Configuration Guide, Release 5.x*
High availability	*Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide, Release 5.x*
VDCs	*Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 5.x*
Release Notes	*Cisco Nexus 7000 Series NX-OS Release Notes, Release 5.x*

Standards

Standards	Title
IEEE 802.3ad	—

MIBs

MIBs	MIBs Link
• IEEE8023-LAG-CAPABILITY • CISCO-LAG-MIB	To locate and download MIBs, go to the following URL: http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

Feature History for Configuring vPCs

Table 7-3 lists the release history for this feature.

Table 7-3 Feature History for Configuring vPCs

Feature Name	Releases	Feature Information
vPCs	5.2(1)	Added the **graceful consistency-check** command to enable the vPC primary device to forward traffic when inconsistent configuration is detected between the peers.
vPCs	5.2(1)	Added per-VLAN consistency checking so that only those VLANs with inconsistent configuration are suspended.

Table 7-3 Feature History for Configuring vPCs (continued)

Feature Name	Releases	Feature Information
vPCs	5.2(1)	Added the **auto-recovery** command to improve speed and reliability of vPC recovery after an outage. The **reload restore** command is deprecated.
vPCs	5.2(1)	Added the **vpc orphan-ports suspend** command to suspend orphan ports on the vPC secondary device when the vPC fails.
vPCs	5.2(1)	Support increased to 528 vPCs.
vPCs	5.0(2)	Added the **reload restore** command to configure the vPC switch to assume its peer is not functional and to bring up the vPC.
vPCs	5.0(2)	Added the **peer-switch** command to enable the vPC switch pair to appear as a single STP root in the Layer 2 topology.
vPCs	4.2(1)	Support increased to 256 vPCs.
vPCs	4.2(1)	Added the **peer-gateway** command to ensure that all packets use the gateway MAC address of the device.
vPCs	4.2(1)	Added the **dual-active exclude interface-vlan** command to ensure that VLAN interfaces remain up if the vPC peer link fails.
vPCs	4.2(1)	Added the **delay restore** command to delay the bringup of the vPC secondary device after reload until the routing table can converge.
vPCs	4.1(4)	Support increased to 192 vPCs.
vPCs	4.1(2)	These features were introduced.