Configuring vPCs

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

Note From 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.

From Cisco NX-OS Release 5.1(1), you can use any of the 10-Gigabit Ethernet (10GE) interfaces, or higher, on the F-series modules or the 10-Gigabit Ethernet interfaces, or higher, 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 9000 for more information about shared interfaces.

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.

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at https://tools.cisco.com/bugsearch/ and the release notes

for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information"chapter or the Feature History table in this chapter.

Feature History for Configuring vPCs

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

Table 1: Feature History for Configuring vPCs

Feature Name	Release	Feature Information
Dynamic Routing over vPC	8.4(1)	Added support for Dynamic Routing over vPC feature on Cisco Nexus F4 Series modules for IPv4 and IPv6 unicast traffic.
vPC support on M3 modules	7.3(0)DX(1)	Added support for vPCs on M3 modules.
Hitless vPC Role Change	7.3(0)D1(1)	Added support for switching vPC roles without impacting traffic flows.
vPC Shutdown	7.2(0)D1(1)	Added the **shutdown** command that shuts down the peer to isolate it for debugging, reloading, or physically removing it from the vPC complex, and enables the peer vPC switch to take over as the primary peer.
Physical Port vPC on F3	7.2(0)D1(1)	Added support for physical port vPCs for F3.
1500 host vPC for FEX (Physical Port vPC on FEX)	7.2(0)D1(1)	Added support for 1500 host vPC for FEX (Physical Port vPC on FEX).
vPC Configuration Synchronization	7.2(0)D1(1)	vPC Configuration Synchronization feature synchronizes the configurations of one switch automatically to other similar switches.
Layer 3 over vPC for F2E and F3 modules	7.2(0)D1(1)	Added support for this feature.
Physical Port vPC on F2	6.2(6)	Added support for physical port vPCs for F2.
LAN shutdown	6.2(6)	Added the **shutdown lan** command to support this feature.
FCoE over physical port vPCs	6.2(6)	Added support for this feature.
Physical port vPCs	6.2(6)	Added support for physical port vPCs on the physical interface of vPC peer devices.
vPCs	6.2(2)	Added the **mode auto** command to enable certain commands for vPCs simultaneously.

Feature Name	Release	Feature Information
vPCs	6.1(3)	Added the **multicast load-balance** command that allows two peers to be partially designated forwarders when both vPC paths are up.
vPCs	5.2(1)	Support increased to 528 vPCs.
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)	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 per-VLAN consistency checking so that only those VLANs with inconsistent configuration are suspended.
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	5.0(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	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	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	4.2(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 **peer-gateway** command to ensure that all packets use the gateway MAC address of the device.
vPCs	4.2(1)	Support increased to 256 vPCs.
vPCs	4.1(4)	Support increased to 192 vPCs.
vPCs	4.1(2)	These features were introduced.

Information About vPCs

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. 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 the bisectional bandwidth by enabling multiple parallel paths between nodes and allowing load balancing traffic.

vPC+

A virtual port channel+ (vPC+) is an extension to virtual port channels (vPCs) that run CE only. A vPC+ domain allows a classical Ethernet (CE) vPC domain and a Cisco FabricPath cloud to interoperate and also provides a First Hop Routing Protocol (FHRP) active-active capability at the FabricPath to Layer 3 boundary. A vPC+ domain enables Cisco Nexus 7000 Series enabled with FabricPath devices to form a single vPC+, which is a unique virtual switch to the rest of the FabricPath network. For more detailed information on vPC+ see the Cisco Nexus 7000 Series NX-OS FabricPath Configuration Guide.

Note You cannot configure a vPC+ domain and a vPC domain in the same VDC.

Figure 1: vPC Architecture

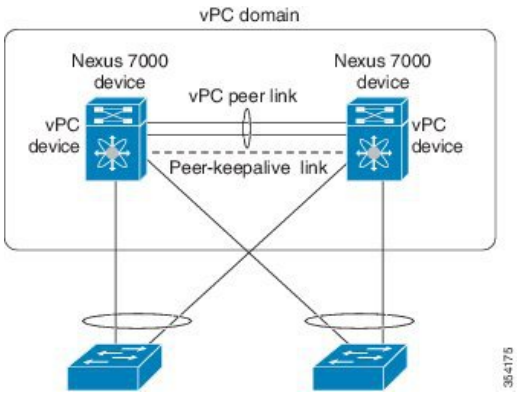

You can use only Layer 2 port channels in the vPC. A vPC domain is associated to a single Virtual Device Context (VDC), so all vPC interfaces belonging to a given vPC domain must be defined in the same VDC. You must have a separate vPC peer link and peer-keepalive link infrastructure for each VDC deployed. 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 at least 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, the F-series line cards can have 16 active links and M-series line cards can have 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, F-series card each device can have eight active links and eight standby links in a single port channel. (See the "vPC Interactions with Other Features" section for more information on using LACP and vPCs.)

You can use the **lacp graceful-convergence** command to configure port channel Link Aggregation Control Protocol (LACP) graceful convergence. You can use this command only on a port-channel interface that is in an administratively down state. You cannot configure (or disable) LACP graceful convergence on a port channel that is in an administratively up state.

You can use the **lacp suspend-individual** command to enable LACP port suspension on a port channel. LACP sets a port to the suspended state if it does not receive an LACP bridge protocol data unit (BPDU) from the peer ports in a port channel. This can cause some servers to fail to boot up as they require LACP to logically bring up the port.

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

From 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 for information about rollbacks 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 port mode. To ensure that you have the correct hardware to enable and run a vPC from Cisco NX-OS Release 4.1(5), enter the show hardware feature-capability command. If you see an X across from the vPC in your command output, your hardware cannot enable the vPC feature.

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

From 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.

You can 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 for information about 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.

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 interfaces that 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 virtual routing and forwarding (VRF) instance that is 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 the figure below).

Figure 2: Separate Switch Required to Connect Management Ports for vPC Peer-Keepalive Link

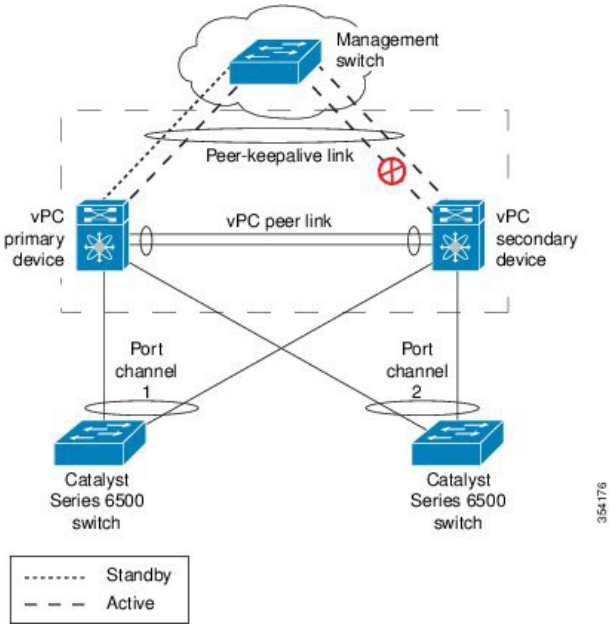

No data or synchronization traffic moves over the vPC peer-keepalive link; the only traffic on this link is a message that indicates that the originating switch is operating and running a vPC.

- 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 must be on 10-Gigabit Ethernet interfaces.

- Keeps both vPC peer switches synchronized for control plane information (such as the 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.

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

Note Starting from Cisco NX-OS Release 8.0(1) you cannot configure vPC peer-link on a port-channel with non-default MTU configuration. The following error message is displayed if you try to configure:

```
ERROR: Cannot configure peer-link since mtu is non-default
```

To configure peer-link, remove the non-default MTU configuration and re apply the **vpc peer-link** command. By default packets of all sizes are allowed in peer-link.

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.

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

vPC Peer Link and I/O Modules Support in Cisco NX-OS Release 6.2

You can configure F2e VDCs. The VDC type for two vPC peer devices must match when the F2 Series module and the F2e Series module are used in the same VDC or system. For an F2 Series module and an F2e Series module in the same topology, the features related to the F2 Series module will only apply.

After ISSU to Cisco NX-OS Release 6.2(2), F2 VDCs will automatically change to F2 F2e VDCs, regardless of the existence of an F2e Series module.

The table below displays the I/O modules that are supported on both sides of a vPC peer link in Cisco NX-OS Release 6.2.

Table 2: I/O Module Combinations Supported on Both Sides of a vPC Peer Link, Cisco NX-OS Release 6.2 and Later

vPC Primary	vPC Secondary
M1 I/O module	M1 I/O module
M2 I/O module	M2 I/O module
M3 I/O module	M3 I/O module
F2 I/O module	F2 I/O module
F2 I/O module	F2e I/O module
F2e I/O module	F2e I/O module
F2e I/O module	F2 I/O module
F3 I/O module	F3 I/O module

vPC Peer Link and I/O Modules Support in Cisco NX-OS Release 6.1 and Earlier Releases

In Cisco NX-OS Release 6.1 and earlier releases, 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. The table above displays the I/O modules that are supported on both sides of a vPC peer link.

While using port channels, we recommended that you use identical line cards on both sides.

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.

The figure below for invalid vPC peer configurations.

Figure 3: vPC Peer Configurations That Are Not Allowed

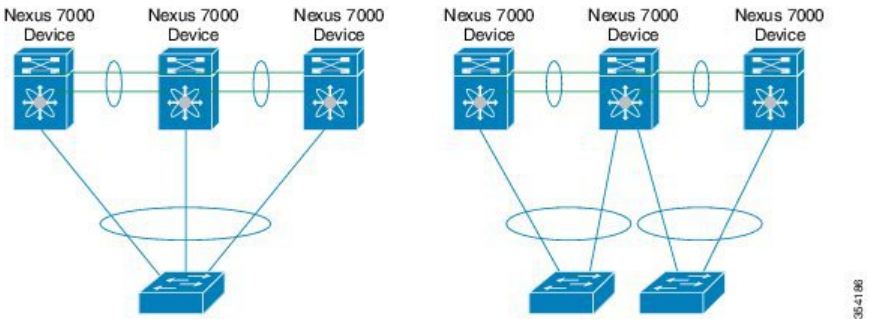

To make a valid configuration, you first configure a port channel on each device and then configure the vPC domain. You assign the port channel on each device as a peer link, using the same vPC domain ID. For redundancy, we recommend that you should configure at least two of the dedicated ports into 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). 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.

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.

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). 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 the 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 entering the **no shutdown** command.

We recommend that you use two different modules for redundancy on each vPC peer device on each vPC peer link.

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

Configuration information flows across the vPC peer links using the Cisco Fabric Services over Ethernet (CFSoE) protocol. (See the "Cisco Fabric Services Over Ethernet" section on page 7-30 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 "Cisco Fabric Services Over Ethernet" section on page 7-30 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.

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 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 for more information about vPCs and STP.
  - When the port-channel is designated as the vPC peer link, the spanning-tree port type network command is added and so the port-channel becomes the bridge assurance port.
  - We recommend that you configure Rapid per VLAN Spanning Tree plus (PVST+) so that the primary device is the root for all VLANs and configure Multiple Spanning Tree (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 Hot Standby Router Protocol (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 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 for information about configuring UDLD.

Configuring Layer 3 Backup Routes on a vPC Peer Link

You can use VLAN network interfaces on the vPC peer devices for such applications as HSRP and PIM. You can use a VLAN network interface for routing from the vPC peer devices.

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 about configuring VLAN network interfaces, see "Configuring Layer 3 Interfaces."

From Cisco NX-OS Release 6.2(2), if the vPC peer link is on an F2e-Series module in a mixed chassis with an M-Series module and an F2e-Series module, do not use the Layer 3 backup routing path over the vPC peer link; instead deploy a dedicated Layer 3 backup routing path using an additional inter-switch port channel.

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.

From 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 From Cisco NX-OS Release 7.2(0)D1(1), 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.

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 about configuring VRFs, see the Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide.

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. You can configure a hold-timeout and a timeout value simultaneously.

Hold-timeout value—The hold-timeout value range is between 3 to 10 seconds, with a default value of 3 seconds. This timer starts when the vPC peer link goes down. The purpose of the hold-timeout period is to prevent false-positive cases.

If you configure a hold-timeout value that is lower than the timeout value, then the vPC system ignores vPC peer-keepalive messages for the hold-timeout period and considers messages for the reminder of the timeout period. If no keepalive message is received for this period, the vPC secondary device takes over the role of the primary device. For example, if the hold-timeout value is 3 seconds and the timeout value is 5 seconds, for the first 3 seconds vPC keepalive messages are ignored (such as, when accommodating a supervisor failure for a few seconds after peer link failure) and keepalive messages are considered for the remaining timeout period of 2 seconds. After this period, the vPC secondary device takes over as the primary device, in case there is no keep alive message.

Timeout value—The timeout value range is between 3 to 20 seconds, with a default value of 5 seconds. This timer starts at the end of the hold-timeout interval. If you configure a timeout value that is lower than or equal to the hold-timeout value, then the timeout duration is initiated after the hold-timeout period. For example, if the timeout value is 3 seconds and the hold-timeout value is 5 seconds, the timeout period starts after 5 seconds

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 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.

This 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 for complete information about configuring trusted ports and precedence.

vPC Peer Gateway

From 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.

Note From Cisco NX-OS Release 6.2(2), you can use the mode auto command to automatically enable this feature. See the "Enabling Certain vPC Commands Automatically" section for more information about using this command.

Some network-attached storage (NAS) devices or load balancers might have features that help to optimize the performances of particular applications. These features enable the device to avoid a routing-table lookup when responding to a request that originated from a host that is not locally attached to the same subnet. Such devices might reply to traffic using the MAC address of the sender Cisco Nexus 7000 Series and Cisco Nexus 7700 Series devices rather than the common HSRP gateway. This behavior is noncompliant with some basic Ethernet RFC standards. Packets that reach a vPC device for the nonlocal 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 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 From 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 line card 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.

Packets that arrive at the peer-gateway vPC device have their Time to Live (TTL) decremented, so that packets carrying a TTL of 1 might get dropped in transit due to TTL expiration. You should take this situation into account when the peer-gateway feature is enabled and particular network protocols that source packets with a TTL of 1 operate on a vPC VLAN.

Dynamic Routing over vPC

Dynamic Routing over vPC feature is supported on F2E, F3, and M3 series modules (for IPv4 and IPv6 Unicast traffic). From Cisco NX-OS Release 8.4(1), the dynamic routing over vPC feature is supported on F4 Series modules.

This feature enables L3 routing protocols such as OPSF to form adjacency with the two vPC peer chassis. The equal routing cost matrices must be configured on applicable interface on each of the vPC peers, failure to do so can result in blocking the traffic. Asymmetric routing feature has to be implemented to address this issue and to configure Dynamic Routing over vPC. Additionally, when Dynamic Routing over vPC is enabled a warning log message is printed.

Layer 3 over vPC for F2E, F3 Modules

This section describes the Layer 3 over vPC for F2E, F3 and M3 Modules feature and how to configure it. Starting from Cisco NX-OS Release 7.2(0)D1(1), Layer 3 over vPC is available on F2E and F3 Series modules. Using this feature, a Layer 3 device can form peering adjacency between both the vPC peers in a vPC complex. vPC peers must have identical VLANs. The TTL of the traffic sent over a peer link does not decrement. The peer-gateway feature should be enabled on all I/O modules before configuring the Layer 3 over vPC feature. The peer-gateway feature allows the vPC peer (SVI-X) (refer the figure below) to forward packets on behalf of other peer (SVI-Y). This feature saves bandwidth by avoiding traffic over the peer link. You can set up peer adjacency between Layer 3 device and vPC peer without separate Layer 3 links. Both bridged and routed traffic can flow over the same link.

Routing adjacency between Layer 3 device and vPC peer is formed without a non-vPC VLAN. Adjacency is formed on the vPC VLAN. Routing adjacency between a Layer 3 device and a vPC peer is formed without Layer 3 inter-switch links between the vPC peers. Adjacency is formed on the vPC peer-link. There is faster convergence when a link or device fails for all traffic. vPC loop avoidance mechanism is available for all traffic.

Figure 4: Layer 3 Over vPC Solution

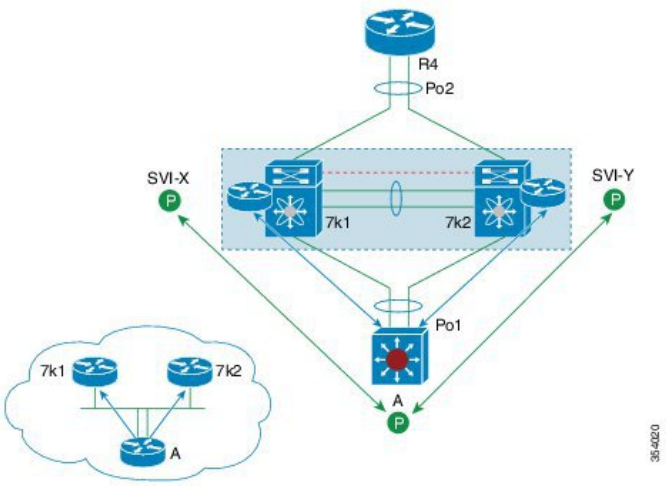

Layer 3 over VPC Support in Cisco NX-OS Release 7.2(0)D1(1)

The following are figures illustrates the Layer 3 over VPC Support in Cisco NX-OS Release 7.2(0)D1(1):

Figure 5: Supported: Peering Over a vPC Interconnection Where the Router Peers with Both the vPC Peers.

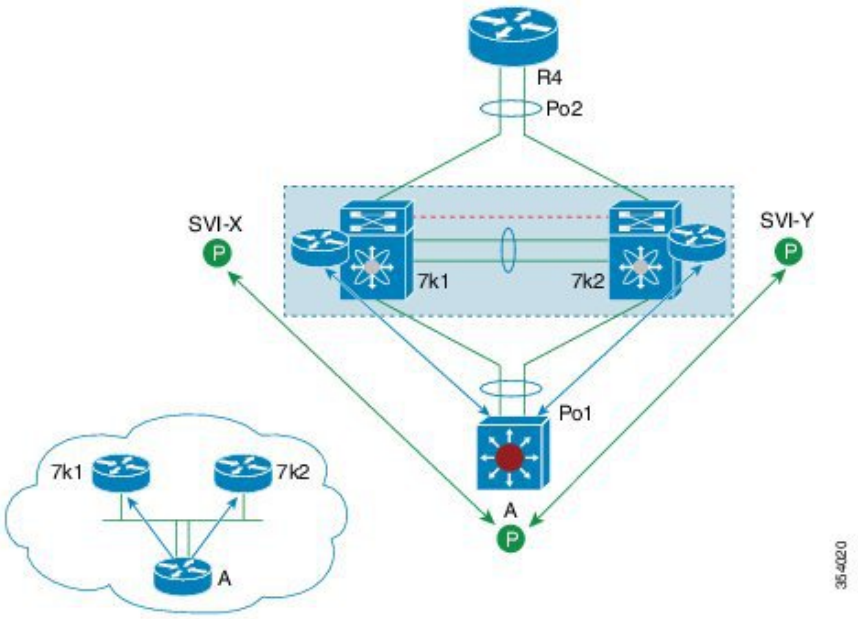

Figure 6: Supported: Peering Over an STP Interconnection Using a vPC VLAN Where the Router Peers with Both the vPC Peers.

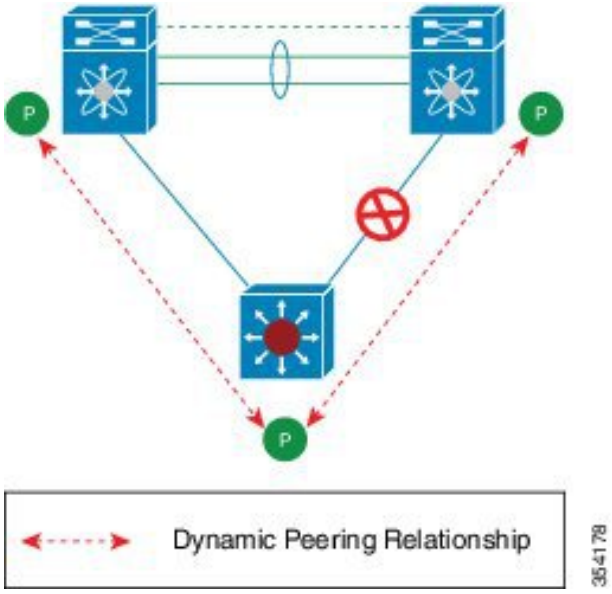

Figure 7: Supported: Peering Over an Orphan Device with Both the vPC Peers.

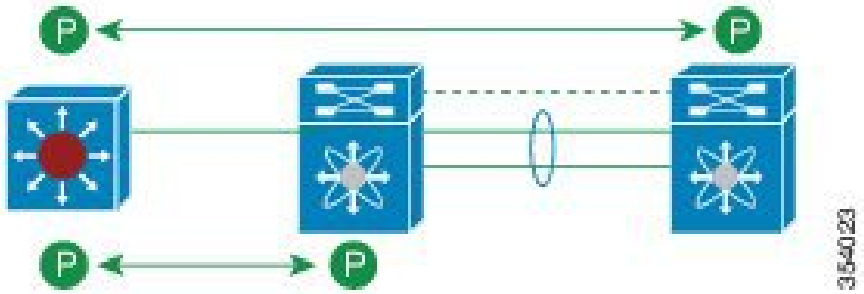

Figure 8: Supported: Peering Over a vPC Interconnection Where Each Nexus Device Peers with Two vPC Peers.

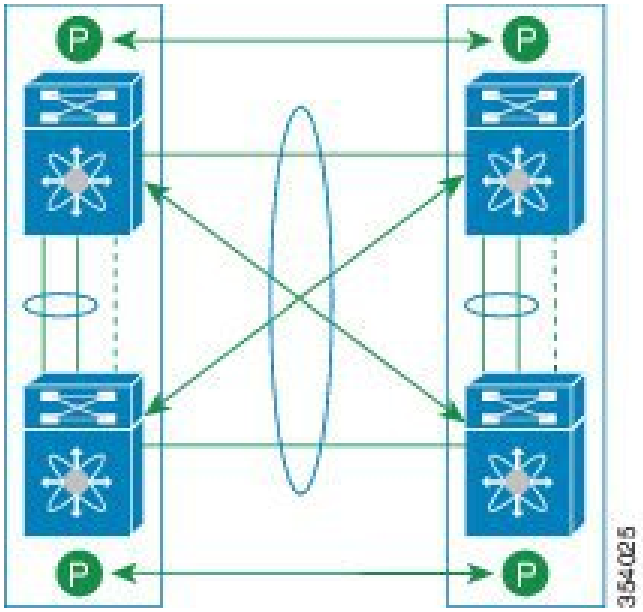

Figure 9: Supported: Peering with vPC Peers Over FEX vPC Host Interfaces

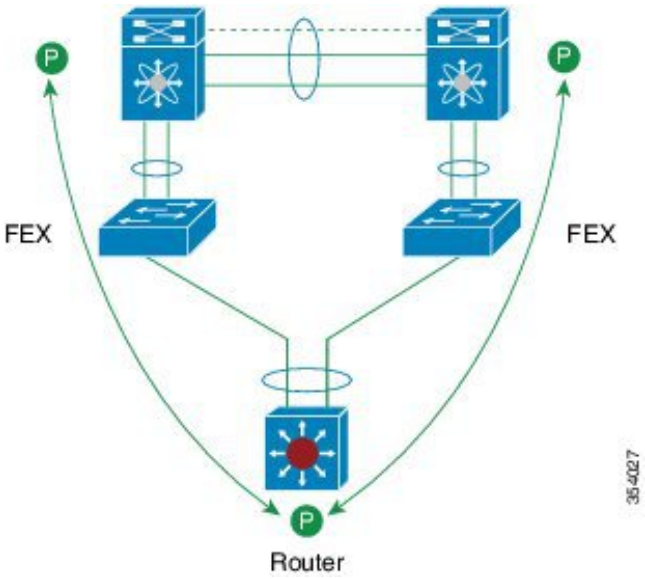

The FEX is connected to Nexus in straight-through topology. The router peers with both Nexus boxes over satellite ports. Layer 3 over vPC in FEX Active-Active mode vPC is not supported.

Figure 10: Unsupported: Peering Across vPC Interfaces with Unequal Layer 3 Metrics

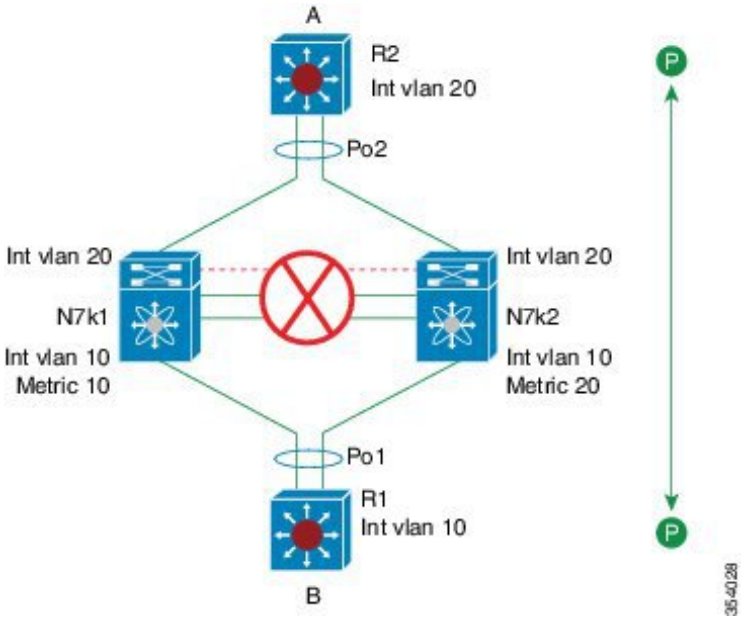

Figure 11: Unsupported: Peering Over vPC+ Interfaces in Cisco NX-OS 7.2(0)D1(1)

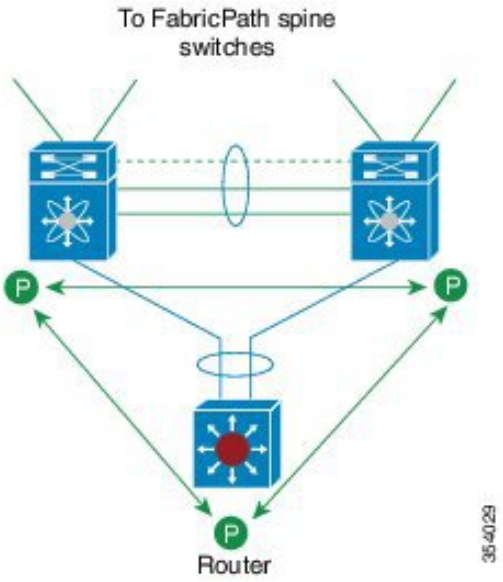

Peering with vPC peers over vPC+ interfaces is unsupported.

Figure 12: Unsupported: Peering with vPC+ Peers an STP Interconnection Using a vPC+ VLAN

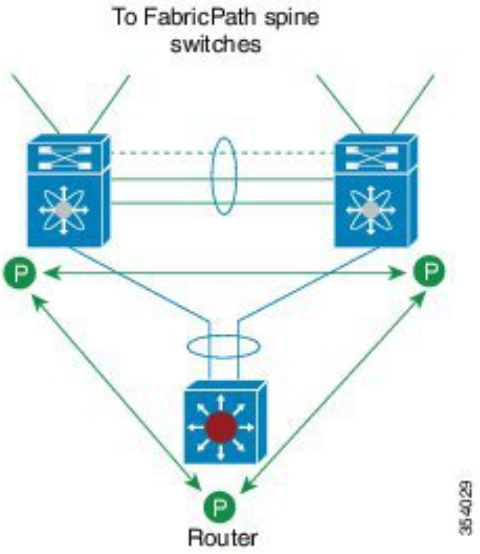

Figure 13: Unsupported: Route Peering with Orphan Device with Both the vPC+ Peers

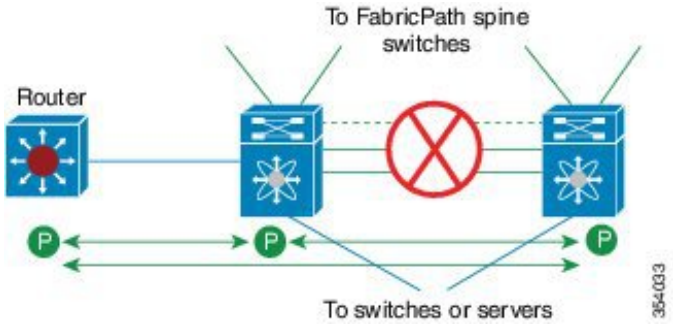

Figure 14: Unsupported: Peering Over PC Interconnection and Over vPC+ Peer Link Using vPC VLAN

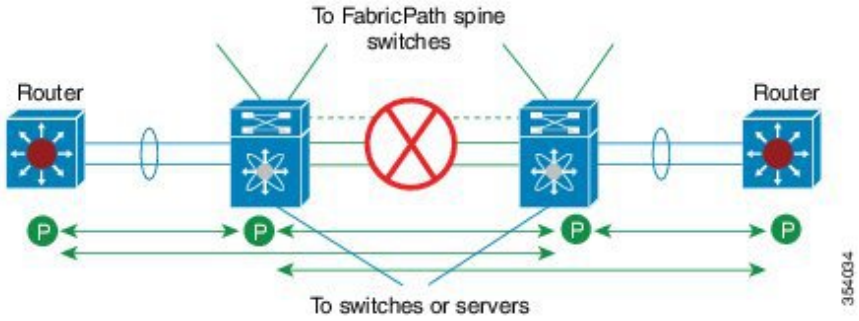

vPC Domain

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 other vPC peer link parameters rather than accept the default values. See the "Configuring vPCs" section for more information about 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 "Cisco Fabric Services Over Ethernet" section for more information about 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 does not come up.

vPC Topology

The figure below shows a basic configuration in which the Cisco Nexus 7000 Series device 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 15: Switch vPC Topology

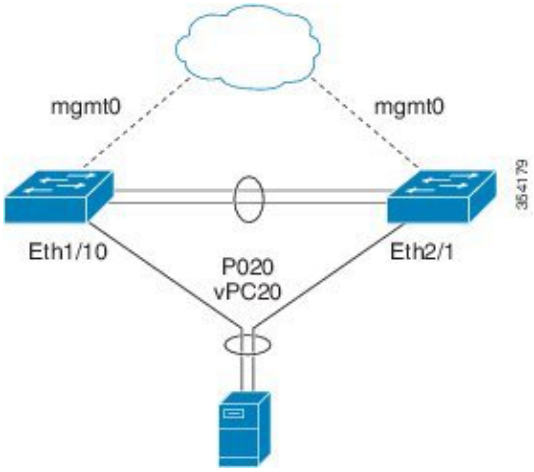

In the figure, vPC 20 is configured on port channel 20, which has Eth1/10 on the first device and Eth2/1 on the second as member ports.

From Cisco NX-OS Release 5.2(1), you can configure a vPC from the peer devices through Fabric Extenders (FEXs), as shown in the figure below.

Figure 16: FEX Straight-Through Topology (Host vPC)

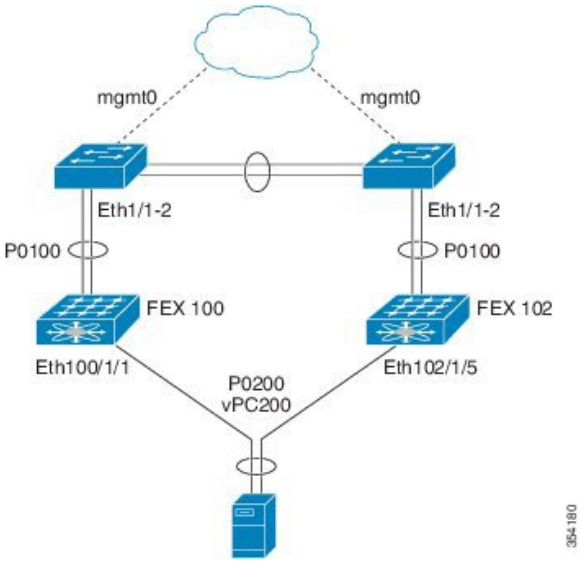

In the figure, each FEX is single-homed (straight-through FEX topology) with a Cisco Nexus 7000 Series device. The host interfaces on this FEX are configured as port channels and those port channels are configured as vPCs. Eth100/1/1 and Eth102/1/5 are configured as members of PO200, and PO200 is configured 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 Cisco Nexus 2000 Series Fabric Extender Software Configuration Guide for Cisco Nexus 7000 Series Switches, Release 7.x for more information about configuring FEX ports.

Physical Port vPCs

Physical port vPCs are vPCs configured on the physical interface of a vPC peer devices. Physical port vPCs can optionally run Link Aggregation Control Protocol (LACP) to the downstream device. Physical port vPCs are supported on F2 and F2E modules. The vPC configuration is applied directly on the member port. You can also enable LACP protocol on the physical interface configured with vPC. From Cisco NX-OS Release 7.2(0)D1(1), physical port vPCs are supported on F3 and FEX modules as well.

Physical Port vPCs for F2, F3, and FEX

This section describes Physical Port VPC for F2, F3, and FEX modules.

The Physical Port VPC for F2, F3, and FEX feature provides the following benefits:

- Enables simple configuration as the user does not create a port-channel to enable the vPC configuration. The vPC configuration is applied directly on the member port.
- Supports vPC setup that has only one 10 Gigabit Ethernet, 40 Gigabit Ethernet, or 100 Gigabit Ethernet port in each leg of the vPC. Creation of port-channel for a vPC setup in such case is not optimal. This feature is best suited for port-channel vPC with only one interface.
- Enhances scalability enabling future support for more physical ports.

- Provides accounting logs and system logs for the physical port, rather than the port-channel.
- Supports large FEX setups. This feature is best suited for port-channel vPC with only one interface.
- Expands the limits of vPC by decoupling the configuration and deployment from the port-channel constructs.
- Enables additional enhancement to extend FCOE support on physical port on the vPC, thus enabling multipathing for the Ethernet traffic while preserving existing constructs for FCOE support.

Note The **fabricpath multicast load-balance** command must be enabled before configuring Physical Port vPC+. This requirement applies to regular front panel and FEX ports.

Figure 17: Physical Port vPC Topology

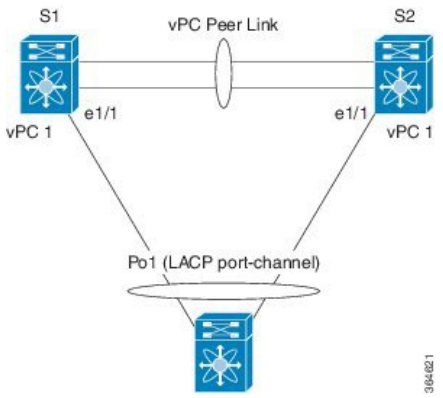

Compatibility Parameters for vPC Interfaces

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, Cisco Fabric Services (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 "Cisco Fabric Services Over Ethernet" section for more information about 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.

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.

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

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 (port-channel mode can, however, be active/passive on each side of the vPC peer)
- 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)

The following parameters were added in Cisco NX-OS Release 6.2(6) for physical port vPCs:

- Native VLAN
- Port mode
- Interface type
- VLAN xLT mapping
- vPC card type
- Shared mode

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 might cause undesirable behavior in the traffic flow:

- MAC aging timers
- Static MAC entries
- VLAN interface
- 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)
- 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 Mismatches

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.

From 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.

From Cisco NX-OS Release 5.2(1), the vPC remains operational, and only the inconsistent VLANs are 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 for more information about STP port types.

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 that connects to the downstream device from the vPC peer device must be identical on both vPC peer devices.

vPC Shutdown

The vPC Shutdown feature enables a user to isolate a switch from a vPC complex before it is debugged, reloaded, or even removed physically, so that the vPC traffic passing through the peer vPC switch in the vPC complex is not affected.

When the user executes the **shutdown** command, the MCEC module (MCECM) stops sending out-of-band (OOB) keep-alive messages and also brings down all the vPC ports, SVIs, and the peer-link. On detection of the peer-link going down and the non-availability of the keep-alive messages, the peer vPC switch takes over as the primary peer. As the keep-alive messages are not received, the peer vPC switch does not bring up the vPC peer-link even after a flap. The isolated vPC switch keeps all the vPCs down as the peer-link is down. The vPC orphan port suspends configured orphan ports.

When the user executes the **no** form of this command, the switch is brought back into the vPC complex with minimal disruption of the network traffic. Executing the **no** form of this command, starts the keepalives, brings up the peer links, and consecutively brings up all the vPCs.

When executed on the primary switch, the **shutdown** command dual-active status is established.

Orphan ports lose connectivity when the vPC **shutdown** command is executed.

Cisco NX-OS services saves the **shutdown** command in the persistent storage service (PSS). The command is restored when the switch reloads. The **shutdown** command is saved as vPC configuration. The **shutdown** command executed again along with the vPC configuration, if it has been copied to the startup configuration. The **shutdown** command is restored when the switch reloads

Version Compatibility Among vPC Switches After vPC shutdown Command

It is possible that the vPC operating version of an isolated vPC peer switch that comes up after debugging or after an ISSU, is different from that the peer switch. When the **no shutdown** command is applied, the vPC peer-link comes up with both the switches having as their versions the lower of the two versions.

Role of STP in vPC Shutdown

The STP synchronizes the port states to the vPC peer causing the new primary vPC peer to take over from the current state, when the role switchover happens. If the MCECM take more than 6 seconds to detect the role change and notify the STP, then the STP bridge protocol data units (BPDUs) that are sent on the vPC are timed out. To avoid this, it is recommended to configure STP peer switch feature so that both vPC switches send BPDUs over the vPC ports.

vPC shutdown Command for a Switch in FEX Active-Active Mode

If you configure the **shutdown** command on the switch to which a dual-homed FEX is connected in a vPC, the FEX goes offline on that switch. An ISSU of the isolated switch does not update the software image on the FEX. You cannot use the vPC **shutdown** command to perform ISSU by isolating and upgrading each switch for FEX Active-Active.

Consider the following FEX Active-Active scenario where peers Peer 1 and Peer 2 are involved:

- The inactive peer, that is Peer 2, is offline because of reasons such as the VPC shutdown command
- An ISSU has been performed on the active peer, that is Peer 1, for upgrading from one software image version to a higher version

All line cards and the remote line cards, including FEX Active-Active, upgrade to higher version of the software image. This happens because the FEX Active-Active is offline on the inactive peer.

Consecutively, when the inactive peer becomes online due to the VPC no shutdown command, this peer will still run the lower version of the software image. In such as case, the status of FEX Active-Active toggles between AA version mismatch and Offline in this peer. This is because both the peers run different versions of the software image. To avoid this situation, the user should not bring up the Peer 2, or execute the VPC shutdown command on it, until the Peer 2 is also upgraded to higher version software image.

Role of the Layer 2 MCECM in vPC Shutdown

When you execute the **shutdown** command, the Multichassis EtherChannel Module (MCECM) stops the keep-alive messages and brings down the peer-link. If the vPC peer switch does not receive keep-alive messages in 5 seconds, it assumes the primary role.

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. 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.

From 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.

You should 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 for information about configuring object tracking and track lists.

See the Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide for information about configuring object tracking.

Note This example uses Boolean OR in the track list and forces all traffic to the vPC peer device only for a complete module failure. Note that the Boolean AND operation is not supported with vPC object tracking.

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 a vPC to the remote peer when all related interfaces on a single module fail, follow these steps:

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

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

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

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

3. Add this track object to the vPC domain:

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

4. Display the track object:

```
switch# 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:

```
switch# 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

vPC and LACP

LACP uses the system MAC address of the vPC domain to form the LACP Aggregation Group (LAG) ID for the vPC.

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, a 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, a vPC peer link and downstream vPC channels support up to 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 does not come up.

vPC Peer Links and STP

Although vPCs provide a loop-free Layer 2 topology, 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 a 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.

When the port-channel is designated as the vPC peer link, the spanning-tree port type network command is added and so the port-channel becomes the bridge assurance port. We recommend that you do not enable any of the STP enhancement features on vPC peer links. If the STP enhancements are already configured, they do not cause any problems for the vPC peer links.

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 for information about STP enhancement features and PVST simulation.

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 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 "Cisco Fabric Services Over Ethernet" section for information about 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 for information about 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 for an example.

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 for more information about STP port types.

Note If you bridge two VLANs on a Nexus 7000 peer-switch, with an Adaptive Security Appliance (ASA) in a transparent mode, the switch puts one of the VLAN in a STP dispute. To avoid this, disable peer-switch or STP on the ports.

vPC Peer Switch

The vPC peer switch feature is enabled on 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 the 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 vPC topology (non-hybrid), in which all the devices belong to the vPC topology.

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

See the Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide for information about STP enhancement features and Rapid PVST+.

vPC Peer Link's Designated Forwarder

From Cisco NX-OS Release 6.0, Cisco NX-OS provides a way to control two peers to be partially designated forwarders when both vPC paths are up. When this control is enabled, each peer can be the designated forwarder for multi-destination southbound packets for a disjoint set of RBHs/FTAGs (depending on the hardware). The designated forwarder is negotiated on a per-vPC basis. This control is enabled with the **fabricpath multicast load-balance** command which is configured under vPC domain mode, for example:

```
switch(config)# vpc domain 1
switch(config-vpc-domain)# fabricpath multicast load-balance
```

From Cisco NX-OS Release 6.2(2), this feature is automatically enabled when the **mode auto** command is used. See the "Enabling Certain vPC Commands Automatically" section for more information about using this command.

Note Only an F2-series module supports multicast load balancing. On an F1-series module, the configuration is supported, but load balancing does not occur.

Note The **fabricpath multicast load-balance** command is required for configuring vPC+ with FEX ports.

See the Cisco Nexus 7000 Series NX-OS FabricPath Configuration Guide for more detailed information on enabling designated forwarders on vPCs.

vPC and ARP or ND

A feature was added in the 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 overcomes the delay that occurs 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.

Note From Cisco NX-OS Release 6.2(2), you can use the mode auto command to automatically enable this feature. See the "Enabling Certain vPC Commands Automatically" section for information about using this command.

vPC Multicast—PIM, IGMP, and IGMP Snooping

Note The Cisco NX-OS software for the Nexus 7000 Series devices does not support Product Independent Multicast (PIM), Source-Specific Multicast(SSM) or Bidirectional (BIDR) on a vPC. The Cisco NX-OS software fully supports PIM Any Source Multicast (ASM) on a vPC.

A PIM adjacency between an Switched Virtual Interface (SVI) on a vPC VLAN (a VLAN that is carried on a vPC Peer-Link) and a downstream device is not supported; this configuration can result in dropped multicast packets. If a PIM neighbor relationship is required with a downstream device, a physical Layer 3 interface must be used on the Nexus switches instead of a vPC SVI.

For SVIs on vPC Vlans, only one PIM adjacency is supported - which is with the vPC Peer Switch. PIM adjacencies over the VPC Peer-Link with devices other than the VPC Peer Switch for the vPC-SVI are NOT supported.

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 might 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 might 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.

Overall multicast convergence times are scale and vPC role change / PIM restart duration dependent.

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 for information about commands that display information on a vPC and multicast.

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

- vPC PIM—The PIM process in vPC mode ensures that only one vPC peer device forwards multicast traffic. The PIM process in vPC mode synchronizes the source state with both vPC peer devices and elects which vPC peer device forwards the traffic.
- vPC IGMP/IGMP snooping—The IGMP process in vPC mode synchronizes the designated router (DR) information on both vPC peer devices. Dual DRs are available for IGMP when you are in vPC mode. Dual DRs are not available when you are not in vPC mode, because both vPC peer devices maintain the multicast group information between the peers.

Note A PIM neighbor relationship between a vPC VLAN (a VLAN that is carried on a vPC peer link) and a downstream vPC-attached Layer 3 device is not supported, which can result in dropped multicast packets. If a PIM neighbor relationship is required with a downstream Layer 3 device, a physical Layer 3 interface must be used instead of a vPC interface.

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 for more information about 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 outgoing interface (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 DRs, both vPC devices have an identical (*,G) entry with respect to the vPC SVI OIFs as shown in this 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 toward the source as shown in this 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 the case of a failure (for example, a Layer 3 Reverse Path Forwarding(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 toward the source. The downside is that the link bandwidth upstream from the nonforwarder gets used for the traffic that is ultimately dropped by it. The benefits that result with better convergence far outweigh the link bandwidth usage. Therefore, we recommend that you use this command if you use vPCs.

PIM DUAL DR and IP PIM PRE-BUILD SPT with VPC Peer Link on F2 Modules

In the vPC implementation in F2-mode, because of a hardware limitation, the PIM dual DR mode is disabled. As a result, only the PIM DR adds the OIF, and the states are shown in this example:

```
Case 1: One OIF
===============
VPC Device1 (say this is PIM DR on oif1):
----------------------------------------
(*,G)
  oif1 (igmp)

VPC Device2:
------------
(*,G) will not be created.
```

When the source traffic is received, only vPC Device 1 adds the (S,G) route.

```
VPC Device1 (say this is PIM DR on oif1):
```

```
----------------------------------------
(*,G)
  oif1 (igmp)
(S,G)
  oif1 (mrib)

VPC Device2:
------------
(*, G) will not be created.
(S, G) will not be created.
```

In this case (with F2 mode), even if you enter the **ip pim pre-build-spt** command, no value is added because the corresponding (S,G) route is not created in the first place.

```
Case 2: Two OIFs
================
VPC Device1 (say this is PIM DR on oif1):
----------------------------------------
(*,G)
  oif1 (igmp)

VPC Device2 (say this is PIM DR on oif2):
----------------------------------------
(*,G)
  oif2 (igmp)
```

When the source traffic is received, associated OIFs are inherited by the (S,G) routes as shown in this example:

```
VPC Device1 (say this is PIM DR on oif1):
----------------------------------------
(*,G)
  oif1 (igmp)

(S,G)
  oif1 (mrib)

VPC Device1 (say this is PIM DR on oif2):
----------------------------------------
(*,G)
  oif2 (igmp)

(S,G)
  oif2 (mrib)
```

In the case of a vPC peer link with F2 modules, you do not need to enter the **ip pim pre-build-spt** command because PIM sends (S,G) joins upstream because associated routes have a non-NULL oiflist.

Note Do not enter the **ip pim pre-build-spt** command if the vPC feature is enabled in F2 mode.

vPC Peer Links and Routing

The First Hop Routing Protocols (FHRPs) 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 loses Layer 3 uplinks, the 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 a backup routing path in the following ways:

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

We do not recommend that you configure the burnt-in MAC address option (use-bia) for HSRP or manually configure virtual MAC addresses for any FHRP protocol in a vPC environment because these configurations can adversely affect vPC load balancing. The HSRP use-bia option is not supported on vPCs. When you are configuring custom MAC addresses, you must configure the same MAC address on both vPC peer devices.

From Cisco NX-OS Release 4.2(1), you can use the **delay restore** command to configure a restore timer that delays the vPC coming back up until after the peer adjacency forms and the VLAN interfaces are back up. This feature enables you to avoid packet drops when the routing tables might 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 for more information about FHRPs and routing.

Cisco Fabric Services Over Ethernet

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

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

The Cisco FSoE transport is local to each VDC.

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

Note Do not enter the **no cfs eth distribute** or the **no cfs distribute** command. You must enable Cisco FSoE 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 Cisco FSoE.

Cisco Fabric Service also transports data over TCP/IP. See the Cisco Nexus 7000 Series NX-OS System Management Configuration Guide for more information about Cisco Fabric Service over IP.

Note The software does not support Cisco Fabric Service 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.

If a peer link failure or restoration occurs, an orphan port's connectivity might 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 loses any connections through the primary peer if a peer link failure occurs and the vPC ports are suspended by the secondary peer. If the secondary peer were to also suspend the active orphan port, the device's standby port becomes active, provides a connection to the primary peer, and restores connectivity. From 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.

Fibre Channel over Ethernet over Physical Port vPCs

The Fibre Channel over Ethernet (FCoE) over Physical Port Virtual Port Channels (vPCs) feature extends the shared model for physical Ethernet interfaces to vPC interfaces.

Each Ethernet interface that forms a vPC leg is shared between the storage virtual device context (VDC) and the Ethernet VDC. The shared Ethernet interface carries both FCoE and LAN traffic. Mutually exclusive FCoE and LAN VLANs are allocated to carry the traffic on the vPC leg; FCoE traffic is carried by the FCoE VLAN and LAN traffic is carried by the LAN VLAN.

Shutdown LAN

Certain configuration and network parameters must be consistent across peer switches in order for physical port vDCs to work. If an inconsistency impacting the network (Type 1) is detected, the secondary vPC leg (the physical link between the access switch and the host) is brought down. With FCoE over physical port vPC, vPC legs carry both FCoE and LAN traffic so that the FCoE and LAN link are both brought down. The shutdown LAN feature enables you to shut down or bring up only the LAN VLANs on an Ethernet interface.

vPC Recovery After an Outage

In a data center outage, both of the Cisco Nexus 7000 Series devices that include a vPC get reloaded. Occasionally only one peer 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 might be available to allow vPC services to use only the local ports of the functional peer.

Restore on Reload

Note From 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.

From 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

From 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 reinitialize 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 the STP primary regardless of its role priority and also acts as the master for LACP port roles.

From Cisco NX-OS Release 6.2(2), you can use the **mode auto** command to automatically enable this feature. See the "Enabling Certain vPC Commands Automatically" section for information about using this command.

From Cisco NX-OS Release 7.2(0)D1(1), the secondary device assumes primary role, if the primary peer is down and 15 keep-alives messages are lost.

From Cisco NX-OS Release 7.2(0)D1(1), to enable the secondary peer to take over as the primary peer if the secondary peer misses 15 keep-alives from primary peer, you can configure **auto-recovery** command. When the switch reloads, the auto-recovery timer starts, and the switch takes on the primary STP role if the peer switch does not respond to it.

When vPC shutdown command is configured, auto-recovery is blocked.

From Cisco NX-OS Release 6.2.(2), for auto recovery to occur during the initial boot, the logical peer link must be down, and no peer keepalive messages must be received. In earlier releases, auto recovery did not occur if peer kepalive messages were not received and the physical peer link was set to Up status.

vPC Peer Roles After a 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 by using CFS messaging over the vPC communications channel. Only one device at a time is upgraded. When the first device completes 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 temporarily run different releases of Cisco NX-OS, however the system functions correctly because of its backward compatibility support.

See the Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide for complete information about high-availability features.

Hitless vPC Role Change

The vPC hitless role change feature provides a framework to switch vPC roles between vPC peers without impacting traffic flows. The vPC role swapping is done based on the role priority value of the device under the vPC domain. A vPC peer device with lower role priority is selected as the primary vPC device when the **vpc role preempt** command is executed.

Use Case Scenario for Hitless vPC Role Change

The hitless vPC role change feature can be used in the following scenarios:

- Role change request—When you want to change the roles of the peer devices in a vPC domain.
- Primary switch reload—When the devices comes up after a reload and roles are defined, you can use the hitless vPC role change feature to restore the roles. For example, after a reload if the primary device takes the role of operational secondary and the secondary device takes the role of primary operational, you can change the vPC peer roles to their original defined roles using the **vpc role preempt** command.

Note Always check the existing device role priority before configuring the **vpc role preempt** command. Configure **no port-channel limit** under the vpc domain command before configuring the **vpc role preempt** command.

- Dual-active recovery—In a dual-active recovery scenario, the vPC primary switch continues to be (operational) primary, but the vPC secondary switch becomes the targeted primary switch and keeps its vPC member ports up. You can use the vPC hitless feature and restore the device roles. After the Dual-active recovery, if one side is operational primary and the other side operational secondary, then you can use the the **vpc role preempt** command to restore the device roles to be primary and secondary.

vPC Configuration Synchronization

Virtual port channels (vPC) topologies require identical configurations on peer switches. As a result, you must repeat configurations on both peer switches. This process, which can cause errors due to misconfigurations or omissions, can result in additional service disruptions because of mismatched configurations. Configuration synchronization eliminates these problems by allowing you to configure one switch and automatically synchronize the configuration on the peer switch.

In a vPC topology, each Cisco Nexus 7000 Series switch must have some matching parameters. You can use a vPC consistency check to verify that both Cisco Nexus 7000 Series switches have the same configuration (Type 1 or Type 2). If they do not match, depending on whether it is a global (for example, spanning-tree port mode), a port-level (for example, speed, duplex, or channel-group type), or even a port-channel interface, the vPC can go into a suspended state or a VLAN can go into a blocking state on both peer switches. As a result, you must ensure that the configuration from one switch is copied identically to the peer switch.

Configuration synchronization allows you to synchronize the configuration between a pair of switches in a network. Configuration synchronization and vPCs are two independent features and configuration synchronization does not eliminate vPC consistency checks. The checks will continue. If there is a configuration mismatch, the vPC can still go into a suspended state.

In a FEX Active-Active setup:

- All the Host Interfaces (HIFs) ports are mapped to the internal vPC.
- The vPC Config-Sync feature listens to the internal vPC creation notification and triggers a merge of the HIF port configuration.
- All the future HIF configuration are synchronized with the peer switch, if the merge is successful.
- The status of HIF is marked as "peer out of synchronization" and the configuration of the interface is not synchronized, if the merge fails.
- We recommend that you disable **vpc-config-sync** command before starting ASCII configuration. After the ASCII configuration is completed, enable **config-sync** command for regular operation.

Note

- vPC peer-link should be configured and up state.
- You cannot chose which commands are synchronized.

Benefits of vPC Configuration Synchronization

Configuration synchronization benefits are as follows:

- Provides a mechanism to synchronize configuration from one switch to another switch.
- Merges configurations when connectivity is established between peers.
- Provides mutual exclusion for commands.
- Supports existing session and port profile functionality.
- Provides minimal user intervention.

- Minimizes the possibility of user error.

Supported Commands for vPC Configuration Synchronization

The following types of commands are enabled for configuration synchronization:

Note The **show vpc config-sync cli syntax** command lists all the commands that are enabled for configuration synchronization. You cannot choose which commands are synchronized. For more information, see the Cisco Nexus 7000 Series NX-OS Interfaces Command Reference.

- Type-1 configurations:
  - Global configurations
  - vPC member port-channel configurations
- vPC configurations.

Note The configurations can be given on either of the vPC peer switches.

Guidelines and Limitations for vPCs

vPCs have the following configuration guidelines and limitations:

- Enable vPCs before you configure them.
- Configure the peer-keepalive link and messages before the system can form the vPC peer link.
- Routing over vPC is supported only on F2E and F3 modules prior to Cisco NX-OS Release 8.1(1). Starting from Cisco NX-OS Release 8.1(1), routing over vPC is also supported on M3 series modules for IPv4 unicast traffic. Starting from Cisco NX-OS Release 8.2(1), routing over vPC is also supported on M3 series modules for IPv6 unicast traffic. Routing over vPC is supported on F4 series modules from Cisco NX-OS Release 8.4(1).
- Configure a separate Layer 3 link for routing from the vPC peer devices, rather than using a VLAN network interface for this purpose.
- All ports for a given vPC must be in the same VDC.
- Physical port vPC is not supported with VDCs containing F4 modules. If you have a mixed VDC with F3 and F4 modules, physical port vPC is not supported even when the FEXs are connected to F3 modules.
- Assign a unique vPC domain ID for each respective vPC to configure multilayer (back-to-back) vPCs.
- DHCP Relay is supported.
- We recommend that you configure all the port channels in the vPC using LACP with the interfaces in active mode.

- When a pair of Cisco Nexus 7000 series switches is connected to a downstream device in a vPC setup, and the vPC domain Id is changed, the LACP port channel configuration on one of the switches might go in hot stand-by mode. To avoid the above scenario, we recommend that you remove the vPC configurations and reconfigure the vPC configurations.
- Configure both vPC peer devices; the configuration is not sent from one device to the other.
- Only Layer 2 port channels can be in vPCs.
- vPC peers can operate dissimilar versions of NX-OS software only during the upgrade or downgrade process.
- Different versions of NX-OS software on vPC peer switches is not supported.
- IPv6 multicast on a vPC is not supported.
- Back-to-back, multilayer vPC topologies require unique domain IDs on each respective vPC.
- When using vPCs, 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.
- Configure **vpc orphan-ports suspend** command on all non-vPC-interfaces (port channel or ethernet) that carry vPC peer-link VLAN traffic. During vPC shutdown, vPC manager brings down vPC interfaces, vPC interface VLANs and non-vPC interfaces with **vpc orphan-ports suspend** configuration.
- If you configure open shortest path first (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 for further details about OSPF.
- When you configure a static MAC address on a vPC switch, ensure to configure a corresponding static MAC address on the other vPC switch. If you configure the static MAC address only on one of the vPC switches, the other vPC switch will not learn the MAC address dynamically.
- In a vPC topology, when a Multichassis EtherChannel Trunk (MCT) link is shut down on a vPC primary switch, and is followed by the vPC primary switch reload, the vPC secondary switch's ports do not come up immediately. This may cause a drop in traffic.
- Check that the necessary configuration parameters are compatible on both sides of the vPC peer link. See the "Compatibility Parameters for vPC Interfaces" section for information about compatibility recommendations.
- From Cisco NX-OS Release 7.2(0)D1(1), 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 the 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.
- From Cisco NX-OS Release 8.1(x), in a vPC topology, non-MAC-in-MAC-encapsulated traffic can be lost if all the following conditions are met:
  - The non-MAC-in-MAC-encapsulated traffic that is routed through FabricPath enabled VLANs.

- The packets have to hit the vPC switch from a non-core interface (an orphan port or from one of the hosts hanging off the vPC leg).
- The packet must be destined to one of the hosts hanging off the vPC leg. It has to be an Layer 3 routing case.
- The **no port-channel limit** command is configured under vPC.
- The vPC leg connecting to the vPC host is down and the traffic is routed through the vPC peer link.
- The vPC peer link is on M3 line card modules.

In such a scenario, we recommend that you do not configure the **no port-channel limit** command under vPC.

- The STP port cost is fixed to 200 in a vPC environment.
- You might experience minimal traffic disruption while configuring vPCs.
- Jumbo frames are enabled by default on the vPC peer link.
- Routing protocol adjacency over a fabric path VLAN is not supported.
- 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.
- The software does not support CFS regions.
- Port security is not supported on port channels.
- BFD for HSRP is not supported in a vPC environment.
- A single vPC domain between two VDCs on the same physical Cisco Nexus 7000 device is not supported.
- When Layer 3 over vPC feature is enabled using the **layer3 peer-router** command, BFD enabled with echo function is not supported on a switched virtual interface (SVIs) using vPC VLANs that are part of a vPC peer-link.

Auto recovery has the following limitations and guidelines:

- In Cisco NX-OS Release 6.2(2) and later releases, auto recovery is enabled by default. If you already enabled auto recovery in an earlier release and you upgrade to Release 6.2(2) or a later release, auto recovery will remain enabled after the upgrade. If you want to disable auto recovery in Release 6.2(2) or a later release, you must use the **auto-recovery disable** command to explicitly disable auto recovery.
- From Cisco NX-OS Release 6.2.(2), for auto recovery to occur during the initial boot, the logical peer link must be down and no peer keepalive messages must be received. In releases earlier than 6.2.2, if peer keepalive messages were not received and the physical peer link was set to UP status, auto recovery did not occur.

Physical port vPCs have the following guidelines and limitations:

- Physical port vPCs are supported only on Nexus F2, F2e, and F3 Series modules.
- Physical port vPC is not supported with VDCs containing M3 modules.
- Physical port vPC is supported with vPC+ only on Nexus F2, F2e, and F3 Series modules.

- Physical port vPC is supported on a Fabric Extender (FEX) interface.
- Physical port vPC peer-link must be configured on Cisco Nexus F2, F2E, or F3 Series modules. It cannot be configured on a M Series module.
- Link Aggregation Control Protocol (LACP) cannot be enabled on a physical port without vPC.
- Same vPC configuration cannot be applied to multiple physical ports.
- Physical port vPC does not support ASCII-replay. When ASCII-replay occurs during a non-ISSU upgrade or downgrade between incompatible images, the physical port vPCs on the peer that is not undergoing upgrade will also go down temporarily.
- STP port-type network is not supported for vPC port-channels and STP port-type network is not supported, when **vpc role preempt** is configured on vPC port-channels.

FCoE over physical port vPC has the following guidelines and limitations:

- FCoE is supported only on trunk ports.
- FCoE is supported only for shared interfaces.
- FCoE is not supported on port channel vPCs.
- FCoE over a physical port vPC is supported in storage VDCs of type F2 only.
- FCoE over a physical port vPC is not supported in storage VDCs because Layer 2 multipathing over physical port vPCs are supported only for LAN.
- FCoE over a VPC+ is not supported.
- The shutdown LAN configuration is supported on shared interfaces only.
- The Link Layer Discovery Protocol (LLDP) must be enabled in the Ethernet VDC for shutdown LAN.

Hitless vPC role change feature has the following guidelines and limitations:

- vPC STP hitless role change feature is supported only from Cisco Nexus 7.3(0)D1(1) release onwards.
- vPC role change can be performed from either of the peer devices.
- If the original secondary device has higher role priority value than the original primary device, role swapping cannot be performed. Change the role priority on either vPC device so that the value of the original secondary device is lower than the original primary one. To view the existing role of a device, use the **show vpc role** command on local and peer switch.
- On vPC+, enable the **fabricpath multi path load-balance** command before configuring the vPC hitless role change feature. The Forwarding Tag (FTag) scheme is used in vPC+ to seamlessly configure the role change. To ensure FTag scheme is used, you need to enable the **no port channel limit** command on vPC+ as it has dependencies on the **fabricpath multi path load-balance** command.
- Enable the **no port channel limit** command on vPC+ before configuring the vPC hitless role change feature. If this command is not enabled, vPC hitless role change cannot be configured and an error message is displayed. Configure this command on both the vPC devices.

Note Always check the existing configured role priority before configuring vPC hitless role change feature.

- In a vPC domain, enable the **peer-switch** command, where both vPC peers have same STP priorities, and ensure it is operational before issuing a role change. If you do not enable the **peer-switch** command, it can lead to convergence issues.
- vPC hitless role change cannot be performed if there are any Type 1 inconsistencies on the peer devices.
- When the peer-switch feature is enabled under a vPC domain, ensure that the vPC pair is configured as spanning-tree root for all the vPC VLANs.

Configuring vPCs

Enabling vPCs

Before you begin

- You must enable the vPC functionality before you can configure and use vPCs.
- Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **feature vpc**	Enables vPCs on the device.
Step 3	switch(config)# **exit**	Exits global configuration mode.
Step 4	(Optional) switch# **show feature**	Displays which features are enabled on the device.
Step 5	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

This example shows how to enable the vPC feature:

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

Disabling vPCs

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

Before you begin

Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **no feature vpc**	Disables vPCs on the device.
Step 3	switch(config)# **exit**	Exits global configuration mode.
Step 4	(Optional) switch# **show feature**	Displays which features are enabled on the device.
Step 5	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

This example shows how to disable the vPC feature:

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

Creating a vPC Domain and Entering 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 vpc-domain command mode.

Before you begin

- Ensure that you are in the correct VDC (if you are not in the correct VDC, use the **switchto vdc** command).
- Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	Creates a vPC domain on the device, and enters vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# **exit**	Exits vpc-domain configuration mode.
Step 4	(Optional) switch# **show vpc brief**	Displays brief information about each vPC domain.
Step 5	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

This example shows how to create a vPC domain:

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

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)# exit
```

Configuring a 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 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 about creating and configuring VRFs, see the Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide. 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

- Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.
- Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	Creates a vPC domain on the device, and enters vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# **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**}]	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. Ensure that you either use IPv4 address to configure the 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 about creating and configuring VRFs, see the Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide.
Step 4	switch(config-vpc-domain)# **exit**	Exits vpc-domain configuration mode.
Step 5	(Optional) switch# **show vpc statistics**	Displays information about the configuration for the keepalive messages.
Step 6	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

For more information about configuring VRFs, see the Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide.

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

Creating a 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

- Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.
- Ensure that you have enabled the vPC feature.
- Ensure that you are using a Layer 2 port channel.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface port-channel** *channel-number*	Selects the port channel that you want to use as the vPC peer link for this device, and enters interface configuration mode.
Step 3	(Optional) switch(config-if)# **switchport mode trunk**	Configures this interface in trunk mode.
Step 4	(Optional) switch(config-if)# **switchport trunk allowed vlan** *vlan-list*	Configures the permitted VLAN list.
Step 5	switch(config-if)# **vpc peer-link**	Configures the selected port channel as the vPC peer link, and enters vpc-domain configuration mode.

	Command or Action	Purpose
		Note When the port-channel is designated as the vPC peer link, the **spanning-tree port type network** command is added, so the port-channel becomes the bridge assurance port.
Step 6	switch(config-vpc-domain)# **exit**	Exits vpc-domain configuration mode.
Step 7	(Optional) switch# **show vpc brief**	Displays information about each vPC, including information about the vPC peer link.
Step 8	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

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)# exit
```

Configuring Physical Port vPC on F2, F3, and FEX

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface** *name number*	Specifies the interface that you want to add to a physical port, and enters the interface configuration mode.
Step 3	switch(config-if)# **switchport**	Configures the interface as a Layer 2 switching port.
Step 4	switch(config-if)# **vpc** *number*	Configures the selected physical interface into the vPC to connect to the downstream device, and enters interface vPC configuration mode. You can use any module in the device for the physical interface. The range is from 1 and 4096.

	Command or Action	Purpose
		Note The vPC number that you assign to the physical interface connecting to the downstream device from the vPC peer device must be identical on both vPC peer devices.
Step 5	Required: switch(config-if-vpc)# **lacp mode active**	Enables LACP on the physical port. **Note** Static mode can also be used.
Step 6	Required: switch(config-if-vpc)# **exit**	Exits the interface vPC configuration mode.
Step 7	Required: switch(config-if)# **exit**	Exits the interface configuration mode.
Step 8	Required: switch(config)# **exit**	Exits the global configuration mode.
Step 9	(Optional) switch# **show running-config interface** *name number*	Displays information about the interface.

Example

This example shows how to configure Physical Port vPC on F2, F3, and FEX modules:

```
switch# configure terminal
switch(config)# interface ethernet 1/1
switch(config-if)# switchport
switch(config-if)# vpc 10
switch(config-if-vpc)# lacp mode active
switch(config-if-vpc)# exit
switch(config-if)# exit
switch(config)# exit
switch# show running-config interface
```

This example shows how to verify the LACP mode:

```
switch# show running-config interface

   Interface Ethernet1/1
   no shutdown
   Switchport
    vpc 1
      lacp mode active
```

Creating VLAN on vPC

vPC VLAN is a VLAN that is allowed on vPC member port and vPC peer-link. When configuring large number of VLANs in a vPC environment, it is recommended to configure the VLANs simultaneously by specifying the range of VLANs, instead of configuring one VLAN at a time.

Before you begin

- Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.
- Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vlan 200-299**	Configures VLANs in the range 200 to 299 and enters the VLAN configuration mode.
Step 3	switch(config-vlan)# **exit**	Exits the VLAN configuration mode.

Example

This example shows how to configure 100 VLANs and name each of them:

```
switch# configure terminal
switch(config)# vlan 200-299
switch(config-vlan)# exit
switch(config)# vlan 201
switch(config-vlan)# name finance
switch(config-vlan)# exit
```

Configuring Layer 3 over vPC for F2E, F3 Modules

Before you begin

- Ensure that the peer-gateway in enabled and configured on both the peers and both the peers are running image that supports Layer 3 over vPC feature. If you enter the **layer3 peer-router** command without enabling the peer-gateway feature, a syslog message is displayed recommending you to enable the peer-gateway feature.
- Ensure that the peer link is up

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	Creates a vPC domain on the device, and enters vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000.

	Command or Action	Purpose
Step 3	switch(config-vpc-domain)# **layer3 peer-router**	Enables the Layer 3 device to form peering adjacency with both peers. **Note** Configure this command in both the peers.
Step 4	switch(config-vpc-domain)# **peer-gateway**	Enables Layer 3 forwarding for packets destined for the peer's gateway MAC address.
Step 5	switch(config-vpc-domain)# **exit**	Exits vpc-domain configuration mode.
Step 6	(Optional) switch# **show vpc brief**	(Optional) Displays brief information about each vPC domain.
Step 7	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

This example shows how to configure a Layer 3 over vPC for F2E, F3 modules:

```
switch# configure terminal
switch(config)# vpc domain 5
switch(config-vpc-domain)# layer3 peer-router
switch(config-vpc-domain)# peer-gateway
switch(config-vpc-domain)# exit
```

This example shows how to verify if the Layer 3 over vPC for F2E, F3 modules feature is configured:

```
switch# show vpc brief
vPC domain id : 1
Peer status : peer adjacency formed ok
vPC keep-alive status : peer is alive
Configuration consistency status : success
Per-vlan consistency status : failed
Type-2 consistency status : success
vPC role : secondary
Number of vPCs configured : 2
Peer Gateway : Enabled
Peer gateway excluded VLANs : -
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
Auto-recovery status : Enabled (timeout = 240 seconds)
Operational Layer3 Peer : Enabled
```

Configuring a vPC Peer Gateway

From 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.

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

- Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.
- Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	Creates a vPC domain on the device, and enters vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# **peer-gateway**	Enables Layer 3 forwarding for packets destined to the peer's gateway MAC address.
Step 4	(Optional) switch(config-vpc-domain)# **peer-gateway exclude-vlan** *backup-vlan-id*	From Cisco NX-OS Release 5.1(3), avoids software switching of transit VLAN traffic in a mixed chassis mode. See the "vPC Peer Gateway" section for more information.
Step 5	switch(config-vpc-domain)# **exit**	Exits vpc-domain configuration mode.
Step 6	(Optional) switch# **show vpc brief**	Displays brief information about each vPC, including information about the vPC peer link..
Step 7	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Configuring a Graceful Consistency Check

From 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 for information about consistent configurations on the vPCs.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# **vpc domain** *domain-id*	Creates a vPC domain on the device, and enters vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000.
Step 3	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	switch(config-vpc-domain)# **exit**	Exits vpc-domain configuration mode.
Step 5	(Optional) switch# **show vpc brief**	Displays brief information about each vPC domain.

Example

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
switch(config-vpc-domain)# exit
```

Configuring vPC Shutdown

From Cisco NX-OS Release 7.2(0)D1(1), you can use the **shutdown** command to isolate a switch from a vPC complex before it is debugged, reloaded, or even removed physically, so that the vPC traffic passing through the peer vPC switch in the vPC complex is not affected.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	Creates a vPC domain on the device, and enters vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# **shutdown**	Shuts down the peer to isolate it for debugging, reloading, or physically removing it from the vPC complex, and enables the peer vPC switch to take over as the primary peer.

	Command or Action	Purpose
		Use the **no** form of this command to disable the feature.
Step 4	switch(config-vpc-domain)# **exit**	Exits vPC-domain configuration mode.

Example

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

```
switch# configure terminal
switch(config)# vpc domain 1
switch(config-vpc-domain)# shutdown
switch(config-vpc-domain)# exit
```

Configuring vPC Config Synchronization

Enabling vPC Configuration Synchronization

Before you begin

- You must create identical vPC domain IDs on both vPC peer switches.
- Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	Creates a vPC domain on the device, and enters vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# **config-sync**	Enables vPC configuration synchronization. **Note** This command must be configured on both the primary and secondary switch.

The table below shows the process of configuration synchronization on switch 1 and switch 2:

Primary Switch	Secondary Switch

`switch-1# configure terminal`
`switch-1(config)# vpc domain 300`
`switch-1(config-vpc-domain)# config-sync`
	`switch-2# configure terminal`
	`switch-2(config)# vpc domain 300`
	`switch-2(config-vpc-domain)# config-sync`

Configuration synchronization is enabled on both switches in the same vPC domain.	
`switch-1# configure terminal` `switch-1(config)# spanning-tree mode mst`	
The above configuration is applied on the primary switch and is configuration synchronized to the secondary switch. The configuration is either successfully applied to both switches or will be failed on both.	

`switch-1# show running-config`
`...`
`spanning-tree mode mst`
`...`
	`switch-2# show running-config`
	`...`
	`spanning-tree mode mst`
	`...`

	`switch-2# configure terminal`
	`switch-2(config)# spanning-tree port type`
	`switch-2 default`

The configuration is applied on the secondary switch and is configuration synchronized to the primary switch. **Note** The configuration can be applied to either switch.	

`switch-1# show running-config`
`...`
`spanning-tree port type network default`
`...`
	`switch-2# show running-config`
	`...`
	`spanning-tree port type network default`
	`...`

Synchronizing Configuration for a Physical Port vPC

Before you begin

Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface** *type slot/port*	Specifies the vPC physical port, and enters interface configuration mode.

	Command or Action	Purpose
Step 3	switch(config-if)# **vpc** *vpc-id* [**sync** {**export** | **import**}]	Moves port channel into a vPC and enters interface vPC configuration mode. The range is from 1 to 4096. • **sync export** enables the primary switch configuration to be exported to the secondary switch. • **sync import** enables the secondary switch configuration to be imported to primary switch.
Step 4	(Optional) switch(config-if)# **show running-config interface ethernet** *slot/port*	Displays the running configuration for the physical port.

Asymmetric Mapping

The table below shows the process of enabling configuration synchronization (asymmetric mapping) on the vPC physical port on the primary and the secondary switch:

Primary Switch	Secondary Switch

`switch-1# configure terminal`
`switch-1(config)# interface eth1/1`
`switch-1(config-if)# vpc 100`	

The physical port (ethernet1/1) is added to the vPC 100 domain on the primary switch. vPC 100 is not configured on the secondary switch. The configuration will not be synchronized until vPC 100 is added to the secondary switch.	

	`switch-2# configure terminal`
	`switch-2(config)# interface eth2/3`
	`switch-2(config-if)# vpc 100`

Following the configuration of vPC 100 to the secondary switch, the physical ports (interface ethernet2/3 on the secondary switch and interface ethernet1/1 on the primary switch) will be configuration synchronized.	

Symmetric Mapping

The table below shows the process of enabling configuration synchronization (symmetric mapping) on the vPC physical port on the primary and the secondary switch:

Primary switch	Secondary switch

`switch-1# configure terminal`
`switch-1(config)# interface eth1/1`
`switch-1(config-if)# vpc 100 symmetric`
	`switch-2# configure terminal`
	`switch-2(config)# interface eth1/1`

Primary switch	Secondary switch
The physical port (ethernet1/1) is added to the vPC 100 domain on the primary switch. The physical port (ethernet 1/1) is also present on the secondary switch. The configuration of the physical port on both the primary and secondary switch will be kept in synchronization.	

`switch-1# show running-config interface eth1/10`
 	
`interface ethernet1/1`
`  switchport`
`  switchport mode trunk`
`  vpc 100`
	`switch-2# show running-config interface eth1/10`
	
	`interface ethernet1/1`
	`  switchport`
	`  switchport mode trunk`
	`  vpc 100`

Synchronizing Configuration of vPC Member Port Channel

Before you begin

Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface port-channel** *channel-number*	Selects the port channel that you want to use as the vPC peer link for this device, and enters interface configuration mode.
Step 3	switch(config-if)# **switchport**	Configures the interface as a Layer 2 switching port.
Step 4	switch(config-if)# **vpc** *vpc-id* [**sync** {**export** | **import**}]	Moves port channel into a vPC and enters interface vPC configuration mode. The range is from 1 to 4096. • **sync export** enables the primary switch configuration to be exported to the secondary switch. • **sync import** enables the secondary switch configuration to be imported to primary switch.
Step 5	(Optional) switch(config-if)# **show running-config interface port-channel** *channel-number*	Displays the running configuration for the port channel.

The table below shows the process of enabling configuration synchronization under port channel 10 on the primary and the secondary switch:

Primary Switch	Secondary Switch

`switch-1# configure terminal`
`switch-1(config)# interface port-channel 10`
`switch-1(config-if)# switchport`
`switch-1(config-if)# vpc 10`	

The configuration under port-channel 10 is configuration synchronized to the secondary switch. **Note** The **vpc** *number* command can be given first on either the primary or secondary switch.	

	`switch-2# show running-config interface po10`
	
	`interface port-channel10`
	`  switchport`
	`  vpc 10`

The configuration is applied on the secondary switch and is configuration synchronized to the primary switch. **Note** The configuration can be applied to either switch.	

	`switch-2# configure terminal`
	`switch-2(config)# interface port-channel 10`
	`switch-2(config-if)# switchport mode trunk`

The **show running-config interface port-channel** *channel-number* command shows that the configuration synchronization for port channel 10 is successful:	

`switch-1# show running-config interface port-channel 10`
 	
`interface port-channel10`
`  switchport`
`  switchport mode trunk`
`  vpc 10`
	`switch-2# show running-config interface port-channel 10`
	
	`interface port-channel10`
	`  switchport`
	`  switchport mode trunk`
	`  vpc 10`

Verifying vPC Configuration Synchronization

To verify vPC configuration synchronization, perform one of the following tasks:

Command	Purpose
show running-config vpc-config-sync	Displays whether config-sync is available or not.
show vpc config-sync cli syntax	Displays the list of commands that are able to be configuration synchronized.
show vpc config-sync database	Displays the configuration synchronization database.

Command	Purpose
show vpc config-sync merge status	Displays the merge status of the switch and of each vPC interface.
show vpc config-sync status	Displays the status of the last 10 operations of the vPC configuration synchronization process. • Displays merge status (success/failure). • Displays the last action done by the vPC configuration synchronization process and the result of that action.

Checking 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 for information about consistent configurations on the vPCs.

Procedure

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

Example

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
```

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. 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

- Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.
- Ensure that you are using a Layer 2 port channel.
- Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface port-channel** *channel-number*	Selects the port channel that you want to use as the vPC peer link for this device, and enters interface configuration mode.
Step 3	switch(config-if)# **vpc** *number*	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	switch(config-vpc-domain)# **exit**	Exits vpc-domain configuration mode.
Step 5	(Optional) switch# **show vpc brief**	Displays brief information about each vPC domain.
Step 6	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

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

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

Enabling Certain vPC Commands Automatically

From Cisco NX-OS Release 6.2(2), you can automatically and simultaneously enable the following commands using the **mode auto** command: **peer-gateway**, **auto-recovery**, **fabricpath multicast load-balance**, **ip arp synchronize**, and **ipv6 nd synchronize**.

Note From Cisco NX-OS Release 6.2(2) and later releases, auto recovery is enabled by default. If you want to disable auto recovery in Release 6.2(2) and later releases, you must use the **no auto-recovery** command to explicitly disable auto recovery.

Before you begin

- Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.
- Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **feature vpc**	Enables vPCs on the device.
Step 3	switch(config)# **vpc domain** *domain-id*	Creates a vPC domain on the device, and enters vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000.
Step 4	switch(config-vpc-domain)# [**no**] **mode auto**	Enables the following commands simultaneously: **peer-gateway**, **auto-recovery**, **fabricpath multicast load-balance**, **ip arp synchronize**, and **ipv6 nd synchronize**. Use the **no** form of this command to disable the feature.
Step 5	switch(config-vpc-domain)# **exit**	Exits vpc-domain configuration mode.
Step 6	switch(config)# **exit**	Exits global configuration mode.
Step 7	(Optional) switch# **show running-config vpc**	Displays information about the vPC, including the commands that are enabled.
Step 8	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

This example shows how to simultaneously enable the following commands: **peer-gateway**, **auto-recovery**, **fabricpath multicast load-balance**, **ip arp synchronize**, and **ipv6 nd synchronize**.

```
switch# configure terminal
switch# feature vpc
switch(config)# vpc domain 1
switch(config-vpc-domain)# mode auto

The following commands are executed:
peer-gateway ;
auto-recovery ;
ip arp synchronize ;
ipv6 nd synchronize ;
fabricpath multicast load-balance ;

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# show running-config vpc

!Command: show running-config vpc
!Time: Thu Feb 18 12:31:42 2013

version 6.2(2)
feature vpc

vpc domain 1
peer-gateway
auto-recovery
fabricpath multicast load-balance
ip arp synchronize
ipv6 nd synchronize
```

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 might choose to configure the vPC domain MAC address manually.

Before you begin

- Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.
- Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# **vpc domain** *domain-id*	Creates a vPC domain on the device, and enters vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# **system-mac** *mac-address*	Enters the MAC address that you want for the specified vPC domain in the following format: aaaa.bbbb.cccc.
Step 4	switch(config-vpc-domain)# **exit**	Exits vpc-domain configuration mode.
Step 5	(Optional) switch# **show vpc role**	Displays the vPC system MAC address.
Step 6	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

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

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

Manually Configuring 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 does not come up.

Before you begin

- Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.
- Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# **vpc domain** *domain-id*	Creates a vPC domain on the device, and enters vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# **system-priority** *priority*	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	switch(config-vpc-domain)# **exit**	Exits vpc-domain configuration mode.
Step 5	(Optional) switch# **show vpc role**	Displays the vPC system MAC address.
Step 6	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

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
switch(config-vpc-domain)# exit
```

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 might 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.

vPCs do 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

- Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.
- Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# **vpc domain** *domain-id*	Creates a vPC domain on the device, and enters vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# **role priority** *priority*	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 that this switch has a better chance of being the primary vPC.
Step 4	switch(config-vpc-domain)# **exit**	Exits vpc-domain configuration mode.
Step 5	(Optional) switch# **show vpc role**	Displays the vPC system priority.
Step 6	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

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
switch(config-vpc-domain)# exit
```

Configuring the Tracking Feature on a Single-Module vPC

From 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

- Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.
- Ensure that you have enabled the vPC feature.
- 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.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	Creates a vPC domain on the device, and enters vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# **track** *track-object-id*	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 for information about configuring object tracking and track lists.
Step 4	switch(config-vpc-domain)# **exit**	Exits vpc-domain configuration mode.
Step 5	(Optional) switch# **show vpc brief**	Displays information about the tracked objects.
Step 6	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

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
switch(config-vpc-domain)# exit
```

Configuring for Recovery After an Outage

If an outage occurs, the vPC waits for a 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 From 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 an Autorecovery" section.

From 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

- Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.
- Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	Creates a vPC domain on the device, and enters vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# **reload restore** [**delay** *time-out*]	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 its default settings.
Step 4	switch(config-vpc-domain)# **exit**	Exits vpc-domain configuration mode.
Step 5	(Optional) switch# **show running-config vpc**	Displays information about the vPC, specifically the reload status.
Step 6	(Optional) switch# **show vpc consistency-parameters interface port-channel** *number*	Displays information about the vPC consistency parameters for the specified interface.
Step 7	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

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 an Autorecovery

From 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.

Note From Cisco NX-OS Release 6.2(2) and later releases, auto recovery is enabled by default. If you want to disable auto recovery in Release 6.2(2) or a later release, you must use the **no auto-recovery** command to explicitly disable auto recovery.

Before you begin

- Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.
- Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	Creates a vPC domain on the device, and enters vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000.

	Command or Action	Purpose
Step 3	switch(config-vpc-domain)# **auto-recovery** [**reload-delay** *time*]	Configures the vPC to assume its peer is not functional and to bring up the vPC, and specifies the time to wait after a 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 its default settings.
Step 4	switch(config-vpc-domain)# **exit**	Exits vpc-domain configuration mode.
Step 5	(Optional) switch# **show running-config vpc**	Displays information about the vPC, specifically the reload status.
Step 6	(Optional) switch# **show vpc consistency-parameters interface port-channel** *number*	Displays information about the vPC consistency parameters for the specified interface.
Step 7	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

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. From Cisco NX-OS Release 5.2(1), you can explicitly declare physical 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 From Cisco NX-OS Release 6.2 and earlier, configure the vPC orphan-port command on all the member ports and bundle them into the port channel. For later releases, configure the command directly on the port-channel interfaces.

Before you begin

- Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.
- Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	(Optional) switch(config)# **show vpc orphan-ports**	Displays a list of the orphan ports.
Step 3	switch(config)# **interface port-channel** *channel-number*	Selects the port channel that you want to use as the vPC peer link for this device, and enters interface configuration mode.
Step 4	switch(config-if)# **vpc orphan-ports suspend**	Configures the selected interface as a vPC orphan port to be suspended by the secondary peer in the case of a vPC failure.
Step 5	switch(config-if)# **exit**	Exits interface configuration mode.
Step 6	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

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

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

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

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

Note When using a non-VPC dedicated trunk link between the VPC peers, the non-VPC VLANs should have a different global priority on the peers to prevent STP from blocking the VLANs.

Before you begin

- Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.
- Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	Creates a vPC domain on the device, and enters vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000.
Step 3	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	switch(config-vpc-domain)# **spanning-tree vlan** *vlan-range* **priority** *value*	Configures the bridge priority of the VLAN. Valid values are multiples of 4096. The default value is 32768.
Step 5	switch(config-vpc-domain)# **exit**	Exits vpc-domain configuration mode.
Step 6	(Optional) switch# **show spanning-tree summary**	Displays a summary of the spanning tree port states including the vPC peer switch.
Step 7	(Optional) switch# **copy running-config startup-config**	Copies the running configuration to the startup configuration.

Example

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)# spanning-tree vlan 1 priority 8192
switch(config-vpc-domain)# exit
```

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 (for more information, see the Cisco Nexus 7000 Series NX-OS Layer 2 Switching Command Reference) 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.

Note When using a non-VPC dedicated trunk link between the VPC peers, the non-VPC VLANs should have a different pseudo root priority on the peers to prevent STP from blocking the VLANs.

Before you begin

- Ensure that you are in the correct VDC (if you are not in the correct VDC, use the **switchto vdc** command).
- Ensure that you have enabled the vPC feature.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **spanning-tree pseudo-information**	Configures the spanning tree pseudo information.
Step 3	switch(config-pseudo)# **vlan** *vlan-range* **designated priority** *value*	Configures the designated bridge priority of the VLAN. Valid values are multiples of 4096 from 0 to 61440.
Step 4	switch(config-pseudo)# **vlan** *vlan-range* **root priority** *value*	Configures the root bridge priority of the VLAN. Valid values are multiples of 4096 from 0 to 61440.
Step 5	switch(config)# **vpc domain** *domain-id*	Creates a vPC domain on the device, and enters vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000.
Step 6	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	switch(config-vpc-domain)# **exit**	Exits vpc-domain configuration mode.

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

Example

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)# vpc domain 5
switch(config-vpc-domain)# peer-switch
switch(config-vpc-domain)# exit
```

Enabling Distribution for vPC

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vpc domain** *domain-id*	Creates a vPC domain on the device, and enters vpc-domain configuration mode for configuration purposes. There is no default; the range is from 1 to 1000.
Step 3	switch(config-vpc-domain)# **config-sync**	Enables the vPC config-sync on the switch and registers with the CFS for physical-ethernet (CFSoE). **Note** Repeat the configuration of the **config-sync** command on the other vPC peer as well.
Step 4	switch(config-vpc-domain)# **exit**	Exits vPC-domain configuration mode.

	Command or Action	Purpose
Step 5	switch(config-vpc-domain)# **vpc config-sync re-emerge** [**sync** { **export**|**import**}]	(Optional) Triggers the merging of configuration with the peer switch if the current merge has failed. **Note** You can use the **sync export** option to apply the local switch configuration to the peer switch. You can use the **sync import** option to apply the remote switch configuration to the local switch.
Step 6	switch(config-vpc-domain)# **vpc config-sync re-emerge interface port-channel** *channel-name* [**sync** { **export**|**import**}]	(Optional) Triggers the merging of **interface port-channel** configuration with the peer switch if the current merge has failed. **Note** You can use the **sync export** option to apply the local **interface port-channel channel-number** command configuration with the peer switch. You can use the **sync import** option to apply the remote **interface port-channel channel-number** command configuration to the local switch.
Step 7	switch(config-vpc-domain)# **vpc config-sync re-emerge interface** *type slot/port* [**sync** { **export**|**import**}]	(Optional) Triggers the merging of **interface ethernet** with the peer switch if the current merge has failed. **Note** You can use the **sync export** option to apply the local **interface ethernet slot/port** command configuration with the peer switch. You can use the **sync import** option to apply the remote **interface ethernet slot/port** command configuration to the local switch.
Step 8	switch(config-vpc-domain)# **exit**	Exits vPC domain configuration mode.
Step 9	switch(config)# **exit**	Exits global configuration mode.
Step 10	switch(config)# **show vpc config-sync merge status**	Displays the status of the configuration merge with the peer switch.

Example

```
switch# configure terminal
switch(config)# vpc domain 1
switch(config-vpc-domain)# config-sync
switch(config-vpc-domain)# vpc config-sync re-merge sync export
```

```
switch(config)# vpc config-sync re-merge interface port-channel 1 sync export
switch(config)# vpc config-sync re-merge interface ethernet 1/1 sync export  import
switch(config)# exit
switch(config)# show vpc config-sync merge status
```

Configuring FCoE Over a Physical Port vPC

Configure Physical Port vPC Interfaces

Perform the following task to configure a physical port vPC interface in the Ethernet VDC. Repeat this task to configure the peer VDC.

Before you begin

- Ensure that you have enabled the vPC feature.
- Ensure that you have configured the per link port channel and port channel members.
- Before you configure this feature for the entire system, ensure that you are in the correct VDC. To change the VDC, use the **switchto vdc** command.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface ethernet** *slot/port-list*	Specifies an Ethernet interface and enters interface configuration mode. The range is from 1 to 253 for the slot and from 1 to 128 for the port.
Step 3	switch(config-if)# **switchport**	Configures the interface as a Layer 2 switching port.
Step 4	switch(config-if)# **switchport mode trunk**	Specifies the trunking VLAN interface in Layer 2. A trunk port can carry traffic in one or more VLANs (based on the trunk allowed VLAN list configuration) on the same physical link.
Step 5	switch(config-if)# **switchport trunk allowed vlan** *vlan-list*	Configures a list of allowed VLANs on the trunking interface.
Step 6	switch(config-if)# **spanning-tree port type network**	Configures the interface that connects to a Layer 2 switch as a network spanning tree port.
Step 7	switch(config-if)# **vpc** *number*	Moves port channels into a vPC and enters interface vPC configuration mode. The range of the number argument is from 1 to 4096.

	Command or Action	Purpose
Step 8	switch(config-if-vpc)# **lapc mode active**	Enables LAPC on the peer link member interfaces on which you configured the **channel group mode active** command.
Step 9	switch(config-if-vpc)# **no shutdown**	Brings the port administratively up.

Example

These examples show how to configure a physical port vPC in an Ethernet VDC:

```
switch-eth(config)# feature vpc

switch-eth(config)# interface port-channel 1
switch-eth(config-if)# switchport
switch-eth(config-if)# switchport mode trunk
switch-eth(config-if)# switchport trunk allowed vlan 10-20
switch-eth(config-if)# spanning-tree port type network
switch-eth(config-if)# vpc peer-link

switch-eth(config)# interface Ethernet3/21
switch-eth(config-if)# switchport
switch-eth(config-if)# switchport mode trunk
switch-eth(config-if)# switchport trunk allowed vlan 10-20
switch-eth(config-if)# channel group 1 mode active
switch-eth(config-if)# no shutdown

switch-eth(config)# interface Ethernet3/1
switch-eth(config-if)# switchport
switch-eth(config-if)# switchport mode trunk
switch-eth(config-if)# switchport trunk allowed vlan 10-20
switch-eth(config-if)# vpc 10
switch-eth(config-if-vpc)# lacp mode active
switch-eth(config-if-vpc)# no shutdown
```

These examples show how to configure a physical port vPC in the peer VDC:

```
switch-eth(config)# feature vpc

switch-eth(config)# interface port-channel 1
switch-eth(config-if)# switchport
switch-eth(config-if)# switchport mode trunk
switch-eth(config-if)# switchport trunk allowed vlan 10-20
switch-eth(config-if)# spanning-tree port type network
switch-eth(config-if)# vpc peer-link

switch-eth(config)# interface Ethernet4/21
switch-eth(config-if)# switchport
switch-eth(config-if)# switchport mode trunk
switch-eth(config-if)# switchport trunk allowed vlan 10-20
switch-eth(config-if)# channel group 1 mode active
switch-eth(config-if)# no shutdown

switch-eth(config)# interface Ethernet4/1
switch-eth(config-if)# switchport
switch-eth(config-if)# switchport mode trunk
switch-eth(config-if)# switchport trunk allowed vlan 10-20
switch-eth(config-if)# vpc 10
switch-eth(config-if-vpc)# lacp mode active
switch-eth(config-if-vpc)# no shutdown
```

Configuring Hitless vPC Role Change

Before you begin

- Enable the vPC feature
- Ensure vPC peer link is up
- Verify the role priority of devices

Procedure

Step 1 Enable hitless vPC role change feature.

switch# **vpc role preempt**

Step 2 (Optional) Verify hitless vPC role change feature.

switch# **show vpc role**

Configuring Hitless vPC Role Change

This example on how to configure hitless vPC role change:

```
! The following is an output from the show vpc role command before the
vPC hitless feature is configured !

switch# show vpc role

vPC Role status
----------------------------------------------------
vPC role                        : secondary
vPC system-mac                  : 00:23:04:ee:be:01
vPC system-priority             : 32667
vPC local system-mac            : 8c:60:4f:03:84:41
vPC local role-priority         : 32668
vPC peer system-mac             : 8c:60:4f:03:84:43
vPC peer role-priority          : 32667

! Configure vPC hitless role change on the device!

switch# vpc role preempt

! The following is an output from the show vpc role command after the
vPC hitless feature is configured !

switch# show vpc role

vPC Role status
----------------------------------------------------
vPC role                        : primary
vPC system-mac                  : 00:00:00:00:00:00
vPC system-priority             : 32667
vPC local system-mac            : 8c:60:4f:03:84:41
vPC local role-priority         : 32666
```

```
vPC peer system-mac             : 8c:60:4f:03:84:43
vPC peer role-priority          : 32667
```

Upgrading Line Card Modules for vPC

To upgrade to a new line card module for a virtual port channel (vPC), use one of the following methods:

- Upgrade line card modules using the ISSU method.
- Upgrade line card modules using the reload method.

Upgrading a Line Card Module Using the ISSU Method

In this task, the primary switch is Switch A, and the secondary switch is Switch B.

Note

- Traffic outage might occur on orphan ports when a vPC peer is isolated.
- Multicast receivers behind the vPC might experience traffic outages.
- Ensure that there are alternate paths from core routes to each vPC peer.
- Ensure that the new line card module has the same slot ID and number as the old line card module.

Before you begin

Before you upgrade a line card module, refer to the Cisco Nexus 7000 Series NX-OS Release Notes document, to see the supported Cisco NX-OS release version for a line card module.

Procedure

Step 1 Perform an ISSU upgrade to a supported Cisco NX-OS release version for a new line card module on both the switches. Perform this task one at a time on both the switches. For information on supported release version for a line card module type, see the Cisco Nexus 7000 Series NX-OS Release Notes document. For information on how to perform an ISSU upgrade, see the Cisco Nexus 7000 Series NX-OS Software Upgrade and Downgrade Guide.

Step 2 On both the switches, move the peer-keepalive link out of the existing module, and use the management interface for the peer-keepalive link.

Example:

```
switch# configure terminal
switch(config)# vpc domain <domain-id>
switch(config-vpc-domain)# peer-keepalive destination <peer-switch management-ip>
```

Step 3 Enable the hidden commands on both the switches, one at a time.

Example:

```
switch# configure terminal
switch(config)# vpc domain <domain-id>
switch(config-vpc-domain)# bypass module-check
```

Step 4 Copy the running configuration to the startup configuration on both the switches.

Example:

```
switch# copy running-config startup-config vdc-all
```

Step 5 On the secondary switch (Switch B), shut down the vPC legs. Perform this action in batches and wait until all the traffic is converged. All traffic is now on the primary switch (Switch A).

Example:

```
switch(config)# interface port-channel <channel-number>
switch(config-if)# shutdown
```

Step 6 On the secondary switch (Switch B), shut down all the ports going to core devices. Perform this action in batches and wait until all the traffic is converged.

Step 7 On the secondary switch (Switch B), shut down the vPC peer link.

Step 8 On the secondary switch (Switch B), save the running configuration to a file on bootflash.

Example:

```
switch# copy running-config bootflash:run-cfg-SwitchB.txt vdc-all
```

Step 9 On the secondary switch (Switch B), edit the saved configuration file to change the Virtual Device Context (VDC) type from an existing module to a new module.

For more information on Cisco NX-OS release support for a module type, see the Cisco Nexus 7000 Series NX-OS Release Notes document.

This example shows that the VDC type has changed from an existing module (F2 or F2e) to a new module (F3):

```
Edit { vdc <xyx>
      limit-resource module-type “f3” }
```

Step 10 On the secondary switch (Switch B), replace the old line card with the new line card module.

Step 11 On the secondary switch (Switch B), reconnect the vPC leg ports to the new module. Ensure that all the ports have the same number as the old line card module.

Step 12 On the secondary switch (Switch B), reconfigure the respective ports on the new module using the saved configuration file on bootflash. Ensure that vPC leg ports are in shut state.

Example:

```
switch# copy bootflash:run-cfg-SwitchB.txt running-config
```

Step 13 On the secondary switch, copy the running configuration to the startup configuration on the admin VDC.

Example:

```
switch# copy running-config startup-config vdc-all
```

Step 14 On the secondary switch (Switch B), bring up the vPC peer link. Ensure that the vPC peer link speed is the same on both the switches.

Ensure that vPC is up and Switch A is the primary switch and Switch B is the secondary switch.

Step 15 On the secondary switch (Switch B), bring up the vPC leg ports. Perform this task in batches and wait for all the traffic to converge.

Step 16 On the secondary switch (Switch B), bring up all the ports going to the core device. Perform this task in batches and wait for all the traffic to converge.

Step 17 On the secondary switch (Switch B), clear all the dynamic MAC entries from the MAC address table.

Example:

```
switch# clear mac address-table dynamic
switch# test l2fm dump smac
```

Migration to the new module on the secondary switch is completed.

Step 18 On the primary switch (Switch A), shut down the vPC legs. Perform this action in batches and wait until all the traffic is converged.

Example:

```
switch(config)# interface port-channel <channel-number>
switch(config-if)# shutdown
```

All the traffic is now on the secondary switch (Switch B).

Step 19 On the secondary switch (Switch B), change the vPC role priority to match the primary switch.

Example:

```
switch(config)# vpc-domain <domain-id>
switch(config-vpc-domain)# role priority <priority-id>
```

Step 20 On the primary switch (Switch A), shut down all the ports going to the core devices. Perform this action in batches and wait until all the traffic is converged. All traffic is now on the secondary switch (Switch B).

Step 21 On the primary switch (Switch A), reconfigure the vPC peer-keepalive link by configuring a dummy IP address.

Example:

```
switch# configure terminal
switch(config-if)# vpc domain <domain-id>
switch(config-if)# peer-keepalive destination <dummy-ip>
```

Step 22 On the primary switch (Switch A), shut down the vPC peer link.

vPC role change takes place without any disruption because of the sticky bit feature on the Switch B.

Step 23 On Switch A, save the running configuration to a file on bootflash.

Example:

```
switch# copy running-config bootflash:run-cfg-SwitchA.txt vdc-all
```

Step 24 Edit the saved configuration file to change the VDC type from the existing module to the new module.

For information on Cisco NX-OS release support for a module type, see the Cisco Nexus 7000 Series NX-OS Release Notes document.

Example:

This example shows that the VDC type is changed from F2 to F3 module.

```
Edit { vdc <xyx>
      limit-resource module-type “f3” }
```

Step 25 On the primary switch (Switch A), replace the old line card with the new line card module.

Step 26 On the primary switch (Switch A), reconnect the vPC leg ports to the new module. Ensure that all the ports have the same number as the old line card module.

Step 27 On the primary switch (Switch A), reconfigure the respective ports on the new module using the saved configuration file on bootflash.

Example:

```
switch# copy bootflash:run-cfg-SwitchA.txt running-config
```

Note Ensure that all the vPC leg ports are in shut state.

Step 28 On the primary switch (Switch A), copy the running configuration to the startup configuration on the Admin virtual device context (VDC).

Example:

```
switch# copy running-config startup-config vdc-all
```

Step 29 On the primary switch (Switch A), bring up the vPC peer-keepalive link by configuring the peer-keepalive destination address back to the management IP of Switch B.

Example:

```
switch# configure terminal
switch(config-if)# vpc domain <domain-id>
switch(config-if)# peer-keepalive destination <management-ip peer-device
```

Step 30 On the primary switch (Switch A), bring up the vPC peer link.

Note Ensure that the vPC peer-link speed configuration is same on both the switches.

All the traffic is on the secondary switch (Switch B).

Step 31 On the primary switch (Switch A), bring up the vPC leg ports. Perform this task in batches and wait for all the traffic to converge.

All the traffic is load balanced on both the switches.

Step 32 On the primary switch (Switch A), bring up all the ports going to the core device. Perform this task in batches and wait for all the traffic to converge.

Step 33 Disable the hidden commands on both the switches. Perform this step one at a time on both the switches.

Example:

```
switch# configure terminal
switch(config)# vpc-domain <domain-id>
switch(config-vpc-domain)# no bypass module-check
```

Step 34 On both the switches, reconfigure the peer-keepalive link on the new card modules.

Step 35 Copy the running configuration to the startup configuration on the Admin VDC on both the switches.

Example:

```
switch# copy running-config startup-config vdc-all
```

Step 36 On the primary switch (Switch A), clear all the dynamic MAC entries from the MAC address table.

Example:

```
switch# clear mac address-table dynamic
switch# test l2fm dump smac
```

Step 37 On the secondary switch (Switch B), run the **test l2fm dump smac** command to view any errors.

Example:

```
switch# test l2fm dump smac
```

Migration to the new module on the primary switch is completed.

Migration from existing line card module to a new module is completed on both the switches.

Upgrading Line Card Modules Using the Reload Method

To upgrade from an existing line card module to a new line card module for vPC using the reload method, perform the following tasks:

1. Install Cisco NX-OS image on vPC peers
2. Install line card module using the reload method

Before you plan to upgrade a line card module, refer the Cisco Nexus 7000 Series NX-OS Release Notes document, to see the supported Cisco NX-OS release version for a line card module.

Installing a Cisco Image on vPC Peers

Perform this task on all the vPC peer devices. Switch A is the primary switch, and Switch B is the secondary switch in this task.

Note Traffic outage might occur on orphan ports when a vPC peer is isolated. Multicast receivers behind the vPC might experience traffic outages (30 to 40 seconds).

Before you begin

- Before you upgrade a line card module, refer to the Cisco Nexus 7000 Series NX-OS Release Notes document to see the supported Cisco NX-OS release version for a line card module.
- Ensure that the feature vPC is enabled on both the primary switch and the secondary switch.
- Ensure that there are alternate paths from core routes to each of the vPC peers.

Procedure

Step 1 Set equal vPC role priority on both the vPC peer devices.

Example:

```
switch# configure terminal
switch(config)# vpc domain <domain-id>
switch(config-vpc-domain)# role priority <priority-id>
```

Step 2 Set the **auto-recovery reload-delay** value, in seconds, to maximum delay time on both the switches.

Example:

```
switch(config-vpc-domain)# auto-recovery reload-delay 84600
```

Step 3 Change the system boot parameters to boot the system from the Cisco NX-OS release verison that is supported on the new module on both the switches.

Example:

This example shows that the Cisco NX-OS 6.2(16) image is used for the Cisco Nexus F3 module:

```
switch(config)# no boot kickstart
switch(config)# no boot system
switch(config)# boot kickstart bootflash://n7000-s2-kickstart.6.2.16.bin
switch(config)# boot system bootflash://n7000-s2-dk9.6.2.16.bin
```

For information on the supported release version for a module type, see the Cisco Nexus 7000 Series NX-OS Release Notes document.

Step 4 On the secondary switch (Switch B), shut down the vPC legs. Perform this action in batches and wait until all the traffic is converged.

Example:

```
switch(config)# interface port-channel <channel-number>
switch(config-if)# shutdown
```

All the traffic is now on the primary switch (Switch A).

Step 5 On the secondary switch (Switch B), copy the running configuration to the start up configuration for an Admin VDC.

Example:

```
switch# copy running-config startup-config vdc-all
```

Step 6 On the secondary switch (Switch B), reboot the system with the new Cisco NX-OS image. Wait for the system to boot up and for the Layer 3 links to come up.

Example:

```
switch# reload
```

Step 7 On the secondary switch (Switch B), bring the vPC legs up again. Perform this action in batches and wait until all the traffic is converged.

Example:

```
switch(config)# interface port-channel <channel-number>
switch(config-if)# no shutdown
```

Step 8 On the primary switch (Switch A), shut down the vPC legs. Perform this action in batches and wait until all the traffic is converged.

Example:

```
switch(config)# interface port-channel <channel-number>
switch(config-if)# shutdown
```

Step 9 On the primary switch (Switch A), copy the running configuration to the start up configuration for an Admin VDC.

Example:

```
switch# copy running-config startup-config vdc-all
```

Step 10 On the primary switch (Switch A), reboot the system with the new Cisco NX-OS image. Wait for the system to boot up and for the Layer 3 links to come up.

Example:

```
switch# reload
```

Step 11 On the primary switch (Switch A), bring the vPC legs up again. Perform this action in batches and wait until all the traffic is converged.

Example:

```
switch(config)# interface port-channel <channel-number>
switch(config-if)# no shutdown
```

Traffic is load balanced between the primary switch (Switch A) and the secondary switch (Switch B).

Switch B takes on the role of the operational primary, and Switch A takes on the role of the operational secondary.

Installing a Line Card Module on a vPC Peer Using the Reload Method

Before you begin

- Ensure that you have installed a compatible Cisco NX-OS release version on the vPC peers. For more information, on how to install a Cisco NX-OS release version using the reload method, see Installing a Cisco Image on vPC Peers, on page 85. For more information on the compatible Cisco NX-OS release version for a line card module type, refer to the Cisco Nexus 7000 Series NX-OS Release Notes document.
- Ensure that the new line card module has the same slot ID and number as the old line card module.

Note In this task, Switch A is the operational secondary, and Switch B is the operational primary switch.

Procedure

Step 1 Set equal vPC role priority on both the switches.

Example:

```
switch(config)# vpc-domain <domain-id>
switch(config-vpc-domain)# role priority <priority-id>
```

Step 2 Set the **auto-recovery reload-delay** value , in seconds, to maximum delay time on both the switches.

Example:

```
switch(config-vpc-domain)# auto-recovery reload-delay 86400
```

Step 3 Enable the hidden commands on both the switches, one at a time.

Example:

```
switch# configure terminal
switch(config)# vpc domain <domain-id>
switch(config-vpc-domain)# bypass module-check
```

Step 4 Copy the running configuration to the startup configuration on the Admin VDC on both the switches.

Example:

```
switch# copy running-config startup-config vdc-all
```

Step 5 On the operational secondary (Switch A) switch, shut down the vPC legs. Perform this action in batches and wait until all the traffic is converged.

Example:

```
switch(config)# interface port-channel <channel-number>
switch(config-if)# shutdown
```

All the traffic is on Switch B.

Step 6 Save the running configuration to a file on bootflash and transfer the configuration file outside the switch (Switch A).

Example:

```
switch# copy running-config bootflash:run-cfg-SwitchA.txt vdc-all
switch# copy bootflash:run-cfg-SwitchA.txt tftp://server/run-cfg-SwitchA.txt vrf management
```

Step 7 On the operational secondary switch, edit the saved configuration file to change the VDC type from an existing module to a new module. Copy the configuration file back to the switch (Switch A).

Example:

This example show that the VDC type is changed from F2 to F3 module:

```
Edit { vdc <xyx>
      limit-resource module-type "f3" }

switch# copy tftp://server/ run-cfg-SwitchA.txt bootflash:run-cfg-SwitchA.txt vrf management
```

For information on the Cisco NX-OS release support for a module type, see the Cisco Nexus 7000 Series NX-OS Release Notes document.

Step 8 Power off the operational secondary switch (Switch A) and physically replace the existing module with the new module on the switch.

Step 9 Power on the switch (Switch A) and wait for the system to go online.

Ensure that the Admin VDC is active. On the Admin VDC, reconfigure the new module ports using the saved configuration file. Ensure that all the ports have the same number as the old line card module.

Ensure that all the vPC leg ports are in shut state, and the vPC peer link and the Layer 3 links are up.

Example:

```
switch# copy bootflash:run-cfg-SwitchA.txt running-config
```

Step 10 Bring up the vPC legs on the operational secondary (Switch A). Perform this task in batches and wait for all the traffic to converge.

Example:

```
switch# interface port-channel <channel-number>
Switch# no shutdown
```

Step 11 On the operational primary (Switch B) switch, shut down the vPC legs. Perform this action in batches and wait until all the traffic is converged.

Example:

```
switch(config)# interface port-channel <channel-number>
switch(config-if)# shutdown
```

All the traffic is on Switch A.

Step 12 Save the running configuration to a file on bootflash and transfer the configuration file outside the switch (Switch B).

Example:

```
switch# copy running-config bootflash:run-cfg-SwitchB.txt vdc-all
switch# copy bootflash:run-cfg-SwitchA.txt tftp://server/run-cfg-SwitchB.txt vrf management
```

Step 13 On the operational primary switch (Switch B), edit the saved configuration file to change the VDC type from an existing module to a new module. Copy the configuration file back to the switch (Switch B).

Example:

This example shows that the VDC type is changed from F2 to F3 module:

```
Edit { vdc <xyx>
      limit-resource module-type "f3" }

switch# copy tftp://server/ run-cfg-SwitchB.txt bootflash:run-cfg-SwitchB.txt vrf management
```

For information on the Cisco NX-OS release support for a module type, see the Cisco Nexus 7000 Series NX-OS Release Notes document.

Step 14 Power off the operational primary switch (Switch B) and physically replace the existing module with the new module on the switch.

Step 15 Power on the switch (Switch B) and wait for the system to go online.

Note Ensure that the Admin VDC is active. On the Admin VDC, reconfigure the new module ports using the saved configuration file. Ensure that all the ports have the same number as the old line card module.

Ensure that all the vPC leg ports are in shut state, and the vPC peer link and the Layer 3 links are up.

Example:

```
switch# copy bootflash:run-cfg-SwitchB.txt running-config
```

Step 16 Bring up the vPC legs on the operational primary (Switch B). Perform this task in batches and wait for all the traffic to converge.

Switch A resumes the role of a primary switch and Switch B takes on the role of a secondary switch. Traffic is load balanced between both the switches.

Example:

```
switch# interface port-channel <channel-number>
Switch# no shutdown
```

Step 17 Disable the hidden commands on both the switches. Perform this step one at a time on both the switches.

Example:

```
switch# configure terminal
switch(config)# vpc-domain <domain-id>
switch(config-vpc-domain)# no bypass module-check
```

Step 18 Copy the running configuration to the startup configuration on the Admin VDC on both the switches.

Example:

```
switch# copy running-config startup-config vdc-all
```

Migration from existing line card module to a new module is completed on both the switches.

Verifying the vPC Configuration

Use the information in the following table to verify the vPC configuration:

Table 3: Verifying the vPC Configuration

Command	Purpose
show feature	Displays whether the vPC is enabled or not.
show vpc brief	Displays brief information about 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 about the vPCs.
show vpc peer-keepalive	Displays information about 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.

Verifying Physical Port vPC on F2, F3, and FEX

Use the information in the following table to verify the physical port vPC on F2, F3, and FEX:

Table 4: Verifying Physical Port vPC on F2, F3, and FEX

Command	Purpose
show vpc brief	Displays brief information about the vPCs.

Command	Purpose
show lacp port-vpc summary	Displays the LACP status for the physical port VPC, such as the vPC ID, physical port, and the LACP port state details.
show lacp counters	Displays the LACP counters for port-channels and physical port vPC interfaces.
show lacp counters interface *name number*	Displays the LACP counters on a physical interface or port-channel interface depending on the interface name.
show lacp neighbor	Displays LACP neighbor information for the port.
show lacp neighbor interface *name number*	Displays the neighbors of ports that are configured on a physical interface.

This example shows how to verify brief information about the vPCs:

```
switch# show vpc brief

vPC status
----------------------------------------------------------------------------------------
id   Port          Status     Consistency   Reason           Active    vlans
----------------------------------------------------------------------------------------
1    Ethernet1/1    up        success       - - - -          200-250,  900-1000
```

This example shows how to verify LACP status for the physcial port VPC, such as the vPC ID, physical port, and the LACP port state details:

```
switch# show lacp port-vpc summary

Flags:         D - Down                    P - up
               s - Suspended              H - Hot-standby (LACP only)

VPC-Id                  Member Port
1                       Ethernet 1/1(P)
2                       Ethernet 1/2(H)
3                       Ethernet 1/3(s)
```

This example shows how to verify LACP counters for port-channel and physical port vPC interfaces:

```
switch# show lacp counters

                    LACPDUs         Marker      Marker Response    LACPDUs
Port              Sent   Recv     Sent   Recv     Sent   Recv      Pkts Err
--------------------------------------------------------------------------------
Ethernet2/1
Ethernet2/1        1677   1804     0      0        0      0         0

port-channel2
Ethernet2/2        1677   1808     0      0        0      0         0
```

This example shows how to verify the LACP counters on a physical interface:

```
switch# show lacp counters  interface ethernet 1/1

LACPDUs         Marker      Marker Response    LACPDUs
Port              Sent   Recv     Sent   Recv     Sent   Recv      Pkts Err
--------------------------------------------------------------------------
Ethernet1/1
Ethernet1/1        17466  17464    0      0        0      0        0
```

This example shows how to verify the neighbors of ports that are configured both as a vPC and as a port-channel member:

```
switch# show lacp neighbor

Flags:  S - Device is sending Slow LACPDUs F - Device is sending Fast LACPDUs
        A - Device is in Active mode       P - Device is in Passive mode
Eth1/1 neighbors
Partner's information
            Partner               Partner                     Partner
Port        System ID             Port Number     Age         Flags
Eth1/1      32768,2-0-0-0-0-66     0x2402          41595       SA

            LACP Partner          Partner                     Partner
            Port Priority         Oper Key                    Port State
            32768                 0x91                        0x3d
```

This example shows how to verify the neighbors of ports that are configured on the physical interface:

```
switch# show lacp neighbor interface ethernet 1/1

Flags:  S - Device is sending Slow LACPDUs F - Device is sending Fast LACPDUs
        A - Device is in Active mode       P - Device is in Passive mode
Eth1/1 neighbor
Partner's information
            Partner               Partner                     Partner
Port        System ID             Port Number     Age         Flags
Eth1/1      32768,0-26-98-14-e-c1  0x207           13          SA

            LACP Partner          Partner                     Partner
            Port Priority         Oper Key                    Port State
            32768                 0x0                         0x3d
```

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

This example shows how to configure vPC on device A as shown in the figure below:

Figure 18: vPC Configuration Example

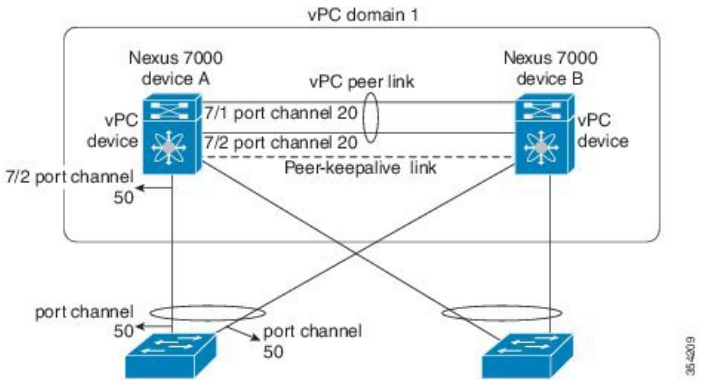

1. Enable vPC and LACP:

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

2. (Optional) Configure one of the interfaces that you want to be a peer link in the dedicated port 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)#
```

3. (Optional) Configure the second, redundant interface that you want to be a peer link in the dedicated port 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)#
```

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
```

5. Create and enable the VLANs:

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

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
```

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
```

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)#
```

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)#
```

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.

Related Documents

Table 5: Related Documents

Related Topic
Cisco Nexus 7000 Series NX-OS Interfaces Command Reference
Cisco Nexus 7000 Series NX-OS System Management Configuration Guide
Cisco Nexus 7000 Series NX-OS High Availability and Redundancy Guide
Cisco Nexus 2000 Series NX-OS Fabric Extender Software Configuration Guide for Cisco Nexus 7000 Series Switches, Release 6.x
Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide
Cisco NX-OS Licensing Guide
VLANs, MAC address tables, private VLANs, and the Spanning Tree Protocol. Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide
Cisco Nexus 7000 Series NX-OS FabricPath Command Reference
Cisco Nexus 7000 Series NX-OS FabricPath Configuration Guide
Cisco Nexus 7000 Series NX-OS Release Notes

Standards

Table 6: Standards

Standards	Title
IEEE 802.3ad	—

MIBs

Table 7: MIBs

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