CCIE in 8 Weeks
Learn. Practice. Achieve

ALL-IN-ONE

CCIE DATA CENTER V2.1

WRITTEN EXAM 400-151 CERT GUIDE

Learn, Practice and Ace the New CCIE Written Exams with Test Prep Resources from CCIEin8Weeks™

PAUL ADAM, CCIE®

CCIEin8Weeks.com

3RD EDITION

All-in-One CCIE Data Center V2.1

400-151 Written Exam Cert Guide

3rd Edition

Contents at a Glance

Chapter 22: Analyze and troubleshoot logical fabric elements

Chapter 23: Design, implement, and troubleshoot virtual networking

Part 6 Evolving Technologies V1.1

Chapter 24: Cloud

Chapter 25: Network Programmability

Chapter 26: Internet of Things (IoT)

Table of Contents

PART 1 CISCO DATA CENTER L2/L3 TECHNOLOGIES

Chapter 1: Design, implement, and troubleshoot Layer 2 technologies

Chapter 2: Design, implement, and troubleshoot overlays

Chapter 3: Design, implement, and troubleshoot routing protocols and features

Chapter 4: Design, implement, and troubleshoot multicast protocols

Chapter 5: Describe inter-fabric connectivity

Chapter 6: Design, implement, and troubleshoot external fabric connectivity

Chapter 7: Design, implement, and troubleshoot traffic management

CHAPTER 1: DESIGN, IMPLEMENT, AND TROUBLESHOOT LAYER 2 TECHNOLOGIES

This chapter covers the following exam topics from Cisco's official 400-151 V2.1 written exam blueprint.

- Link aggregation
- Tagging and Trunking
- Spanning Tree Protocol (STP)

Link Aggregation

Link aggregation allows multiple ports to merge logically in the form of a single link. Link aggregation offers higher aggregate bandwidth to traffic-heavy servers and reroute capability in case of a single port or cable failure. Link aggregation provides increased bandwidth and redundancy. It is sometimes also known as port trunking or link bonding.

Link Aggregation Control Protocol (LACP) is part of an IEEE specification (802.3ad) that allows you to bundle several physical ports together to form a single logical channel. LACP allows a switch to negotiate an automatic bundle by sending LACP packets to the peer. It performs a similar function as Port Aggregation Protocol (PAgP) with Cisco EtherChannel. LACP allows a network device to negotiate an automatic bundling of links by sending LACP packets to the peer (directly connected device that also implements LACP).

vPC is a virtualization technology that presents both Cisco Nexus 7000 Series paired devices as a unique Layer 2 logical node to access layer devices or endpoints. vPC belongs to Multichassis EtherChannel [MCEC] family of technology.

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 to a third device. The third device can be a switch, server, or any other networking device that supports link aggregation technology.

- vPC provides the following technical benefits:
- Eliminates Spanning Tree Protocol (STP) blocked ports
- Uses all available uplink bandwidth
- Allows dual-homed servers to operate in active-active mode

- Provides fast convergence upon link or device failure
- Offers dual active/active default gateways for servers

vPC also leverages native split horizon/loop management provided by port-channeling technology: a packet entering a port-channel cannot immediately exit that same port-channel.

By using vPC, users get the immediate operational and architectural advantages:

- Simplifies network design
- Build highly resilient and robust Layer 2 network
- Enables seamless virtual machine mobility and server high-availability clusters
- Scales available Layer 2 bandwidth, increasing bisectional bandwidth
- Grows the size of the Layer 2 network

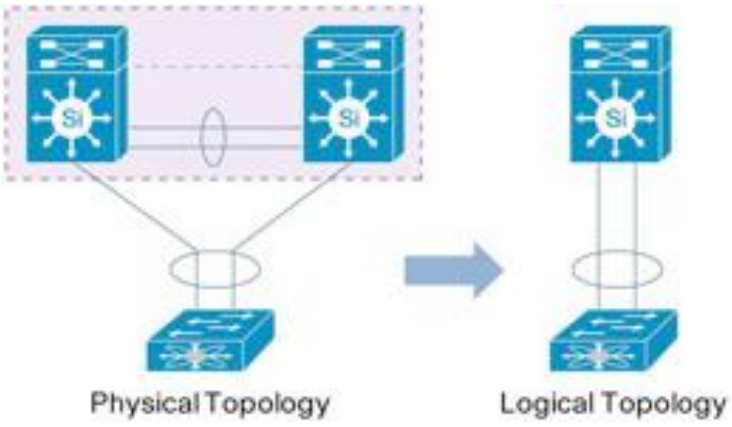

(Source: Cisco.com)

- vPC leverages both hardware and software redundancy aspects:

- vPC uses all port-channel member links available so that in case an individual link fails, hashing algorithm will redirect all flows to the remaining links.

- vPC domain is composed of two peer devices. Each peer device processes half of the traffic coming from the access layer. In case a peer device fails, the other peer device will absorb all the traffic with minimal convergence time impact.

- Each peer device in the vPC domain runs its own control plane, and both devices work independently. Any potential control plane issues stay local to the peer device and does not propagate or impact the other peer device.

From a Spanning-Tree standpoint, vPC eliminates STP blocked ports and uses all available uplink bandwidth. Spanning-Tree is used as a fail-safe mechanism and does not dictate L2 path for vPC-attached devices.

Within a vPC domain, user can connect access devices in multiple ways: vPC-attached connections leveraging active/active behavior with port-channel, active/standby connectivity using spanning-tree, single attachment without spanning-tree running on the access device.

NX-OS Version Requirement for vPC

vPC technology is supported since NX-OS 4.1.3. (i.e. since the inception of Nexus 7000 platform).

NX-OS appropriate version depends on line cards configuration (M1, F1 or F2), chassis type (7010, 7018 or 7009) and Fabric Module generation (FM generation 1 [46 Gbps per module] or generation 2 [110 Gbps per module]).

NX-OS License Requirement for vPC

vPC feature is included in the base NX-OS software license.

- Hot Standby Router Protocol (HSRP), Virtual Router Redundancy Protocol (VRRP), Link Aggregation Control Protocol (LACP) are also included in this base license.
- Layer 3 features like Open Shortest Path First (OSPF) protocol or Intermediate-System-to-Intermediate System (IS-IS) protocol require LAN_ENTERPRISE_SERVICES_PKG license.
- Virtual device contexts (VDCs) requires LAN_ADVANCED_SERVICES_PKG license.

Components of vPC

Table below lists important terms you need to know to understand vPC technology. These terms are used throughout this guide.

Term	Description
vPC	The combined port-channel between the vPC peers and the downstream device. A vPC is a L2 port type: switchport mode trunk or switchport mode access
vPC peer device	A vPC switch (one of a Cisco Nexus 7000 Series pair).
vPC domain	Domain containing the 2 peer devices.

	Only 2 peer devices max can be part of same vPC domain.
vPC member port	One of a set of ports (that is, port-channels) that form a vPC (or port-channel member of a vPC).
vPC peer-link	Link used to synchronize the state between vPC peer devices. It must be a 10-Gigabit Ethernet link. vPC peer-link is a L2 trunk carrying vPC VLAN.
vPC peer-keepalive link	The keepalive link between vPC peer devices; this link is used to monitor the liveness of the peer device.
vPC VLAN	VLAN carried over the vPC peer-link and used to communicate via vPC with a third device. As soon as a VLAN is defined on vPC peer-link, it becomes a vPC VLAN
non-vPC VLAN	A VLAN that is not part of any vPC and not present on vPC peer-link.
Orphan port	A port that belong to a single attached device. vPC VLAN is typically used on this port.

Cisco Fabric Services (CFS) protocol	Underlying protocol running on top of vPC peer-link providing reliable synchronization and consistency check mechanisms between the 2 peer devices.

Virtual PortChannels (vPCs) allow links that are physically connected to two different Cisco® switches to appear to a third downstream device to be coming from a single device and as part of a single PortChannel. The third device can be a switch, a server, or any other networking device that supports IEEE 802.3ad PortChannels. vPC allows the creation of Layer 2 PortChannels that span two switches. At the time of this writing, vPC is implemented on the Cisco Nexus 7000 and 5000 Series platforms (with or without Cisco Nexus 2000 Series Fabric Extenders).

vPCs consist of two vPC peer switches connected by a peer link. Of the vPC peers, one is primary and one is secondary. The system formed by the switches is referred to as a vPC domain.

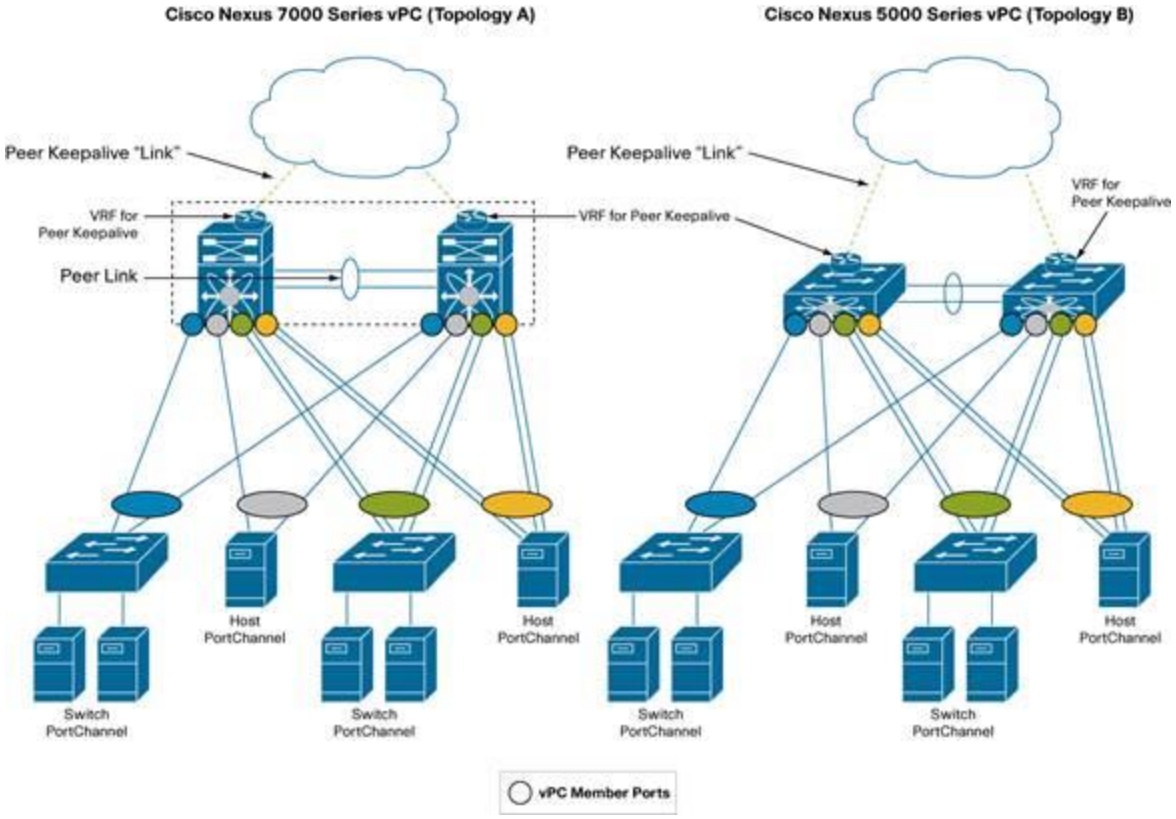

(Source: Cisco.com)

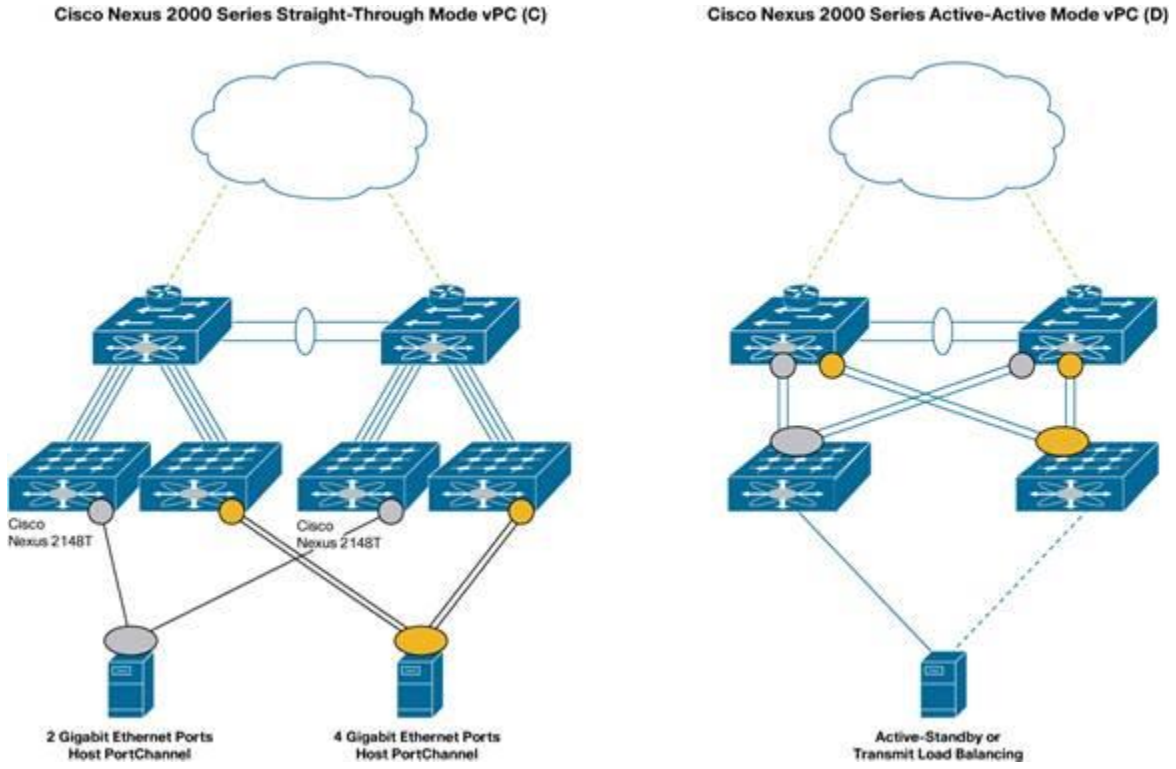

(Source: Cisco.com)

Following is a list of some possible Cisco Nexus vPC topologies, refer to diagrams above:

- vPC on the Cisco Nexus 7000 Series (topology A): This topology consists of access layer switches dual-homed to the Cisco Nexus 7000 Series with a switch PortChannel with Gigabit Ethernet or 10 Gigabit Ethernet links. This topology can also consist of hosts connected with virtual PortChannels to each Cisco Nexus 7000 Series Switch.
- vPC on Cisco Nexus 5000 Series (topology B): This topology consists of switches dual-connected to the Cisco Nexus 5000 Series with a switch PortChannel with 10 Gigabit Ethernet links, with one or more links to each Cisco Nexus 5000 Series Switch. Like topology A, topology B can consist of servers connected to each Cisco Nexus 5000 Series Switch via virtual PortChannels.
- vPC on the Cisco Nexus 5000 Series with a Cisco Nexus 2000 Series Fabric Extender single-homed (also called straight-through mode) (topology C): This topology consists of a Cisco Nexus 2000 Series Fabric Extender single-homed with one to eight 10 Gigabit Ethernet links (depending on the fabric extender model) to a single Cisco Nexus 5000 Series Switch, and of Gigabit Ethernet or 10 Gigabit Ethernet-connected servers that form virtual PortChannels to the fabric extender devices. Note that each fabric extender connects to a single Cisco Nexus 5000 Series Switch and not to both, and that the virtual PortChannel can be formed only by connecting the server network interface cards (NICs) to two fabric extenders, where fabric extender 1 depends on Cisco Nexus 5000 Series Switch 1 and fabric extender 2 depends on Cisco Nexus 5000 Series Switch 2. If both fabric extender 1 and fabric extender 2 depend on switch 1 or both of them depend on switch2, the PortChannel cannot be formed.
- Dual-homing of the Cisco Nexus 2000 Series Fabric Extender (topology D): This topology is also called Cisco Nexus 2000 Series Fabric Extender (FEX for brief) Active/Active. In this topology each FEX is connected to each Cisco Nexus 5000

Series device with a virtual PortChannel. With this topology, the server cannot create a PortChannel split between two fabric extenders. The servers can still be dual-homed with active-standby or active-active transmit-load-balancing (TLB) teaming.

The vPC peer switches are connected through a link called a peer link, also known as a multi-chassis EtherChannel trunk (MCT).

vPC Peer Link

The vPC peer link is the most important connectivity element in the vPC system. This link is used to create the illusion of a single control plane by forwarding Bridge Protocol data units (BPDUs) or Link Aggregation Control Protocol (LACP) packets to the primary vPC switch from the secondary vPC switch. The peer link is used to synchronize MAC addresses between aggregation groups 1 and 2, to synchronize IGMP entries for the purpose of IGMP snooping, it provides the necessary transport for multicast traffic and for the communication of orphaned ports. The term “orphaned ports” refers to switch ports connected to single-attached hosts, or vPC ports whose members are all connected to a single vPC peer.

In the case of a vPC device that is also a Layer 3 switch, the peer link also carries Hot Standby Router Protocol (HSRP) frames. For a vPC to forward a VLAN, that VLAN must exist on the peer link and on both vPC peers, and it must appear in the allowed list of the switch port trunk for the vPC itself. If either of these conditions is not met, the VLAN is not displayed when you enter the command show vpc brief, nor is it a vPC VLAN. When a PortChannel is defined as a vPC peer link, Bridge Assurance is automatically configured on the peer link.

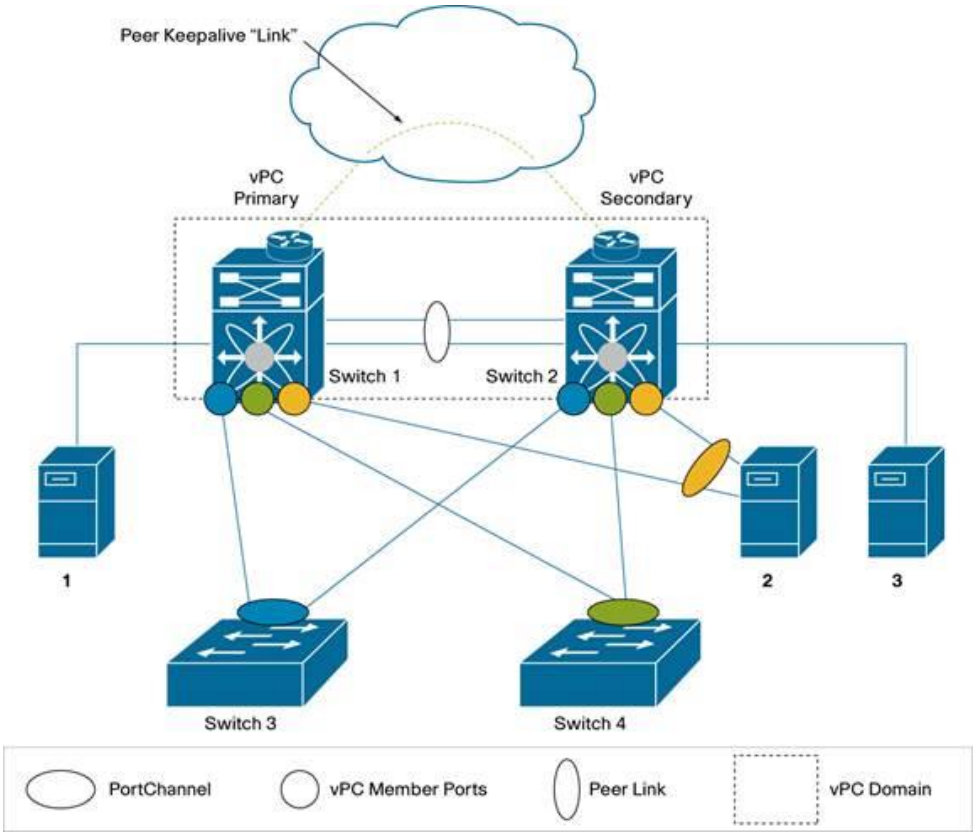

(Source: Cisco.com)

vPC Peer-Keepalive

A routed "link" (it is more accurate to say "path") is used to resolve dual-active scenarios in which the peer link connectivity is lost. This link is referred to as a vPC peer-keepalive or fault-tolerant link. The peer-keepalive traffic is often transported over the management network through the management 0 port of the Cisco Nexus 5000 Series Switch or the management 0 ports on each Cisco Nexus 7000 Series supervisor. The peer-keepalive traffic is typically routed over a dedicated Virtual Routing and Forwarding (VRF) instance (which could be the management VRF, for example). The keepalive can be carried over a routed infrastructure; it does not need to be a direct point-to-point link, and, in fact, it is desirable to carry the peer-keepalive traffic on a different network instead of on a straight point-to-point link.

vPC Ports, and Orphaned Ports

A vPC port is a port that is assigned to a vPC channel group. The ports that form the virtual PortChannel are split between the vPC peers and are referred to as vPC member ports.

A non-vPC port, also known as an orphaned port, is a port that is not part of a vPC.

To summarize, a vPC system consists of the following components:

● Two peer devices: the vPC peers, of which one is primary, and one is secondary; both are part of a vPC domain

● A Layer 3 Gigabit Ethernet link called a peer-keepalive link to resolve dual-active scenarios

● A redundant 10 Gigabit Ethernet PortChannel called a peer link which is used to carry traffic from one system to the other when needed and to synchronize forwarding tables

● vPC member ports forming the virtual PortChannel

Traffic Flows

vPC configurations are optimized to help ensure that traffic through a vPC-capable system is symmetric. In Figure 6, for example, the flow on the left (in blue) reaching a Cisco Nexus switch (Agg1 in the figure) from the core is forwarded toward the access layer switch (Acc1 in the figure without traversing the peer Cisco Nexus switch device (Agg2). Similarly, traffic from the server directed to the core reaches a Cisco Nexus Switch (Agg1), and the receiving Cisco Nexus Switch routes this traffic directly to the core without unnecessarily passing it to the peer Cisco Nexus device. This process occurs regardless of which Cisco Nexus device is the primary HSRP device for a given VLAN.

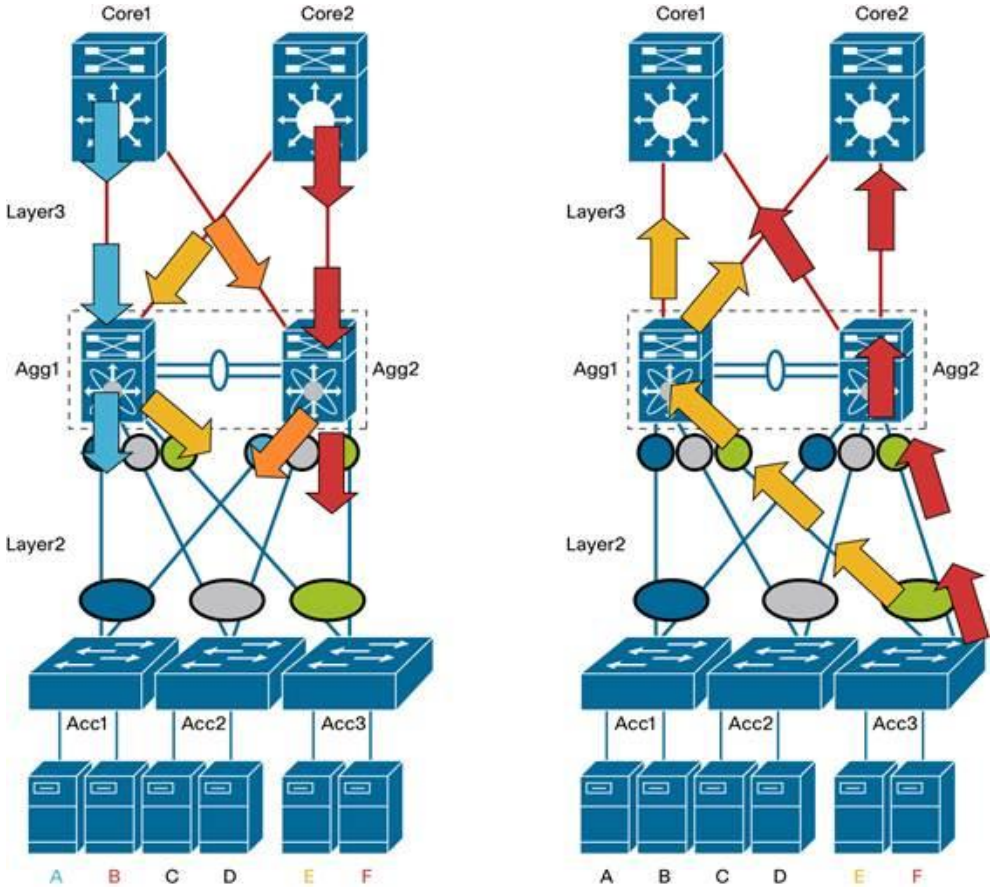

(Source: Cisco.com)

Dual-Control Plane with Single Layer 2 Node Behavior

While still operating with two separate control planes, vPC helps ensure that the neighboring devices connected in vPC mode see the vPC peers as a single spanning-tree and LACP entity. For this to happen, the system has to perform IEEE 802.3ad control-plane operations in a slightly modified way (which is not noticeable to the neighbor switch).

Link Aggregation Group Identifier

IEEE 802.3ad specifies the standard implementation of PortChannels. PortChannel specifications provide LACP as a standard protocol, which enables negotiation of port bundling.

LACP makes misconfiguration less likely, because if ports are mismatched, they will not form a PortChannel.

LACP discerns that the only ports that can be bundled are port 1 going to port 4. According to the IEEE specifications, to allow LACP to determine whether a set of links connect to the same system and to determine whether those links are compatible with aggregation, you need to be able to establish two types of identification:

- You need a globally unique identifier for each system that participates in link aggregation (that is, the switch itself needs to be unique. This number is referred to as the system ID and is composed of a priority and a MAC address that uniquely identifies the switch. Figure 7 illustrates the system ID.
- You need a means of identifying a link aggregation group.

Dual-Control Plane with Single Layer 2 Node Behavior

While still operating with two separate control planes, vPC helps ensure that the neighboring devices connected in vPC mode see the vPC peers as a single spanning-tree and LACP entity. For this to happen, the system has to perform IEEE 802.3ad control-plane operations in a slightly modified way (which is not noticeable to the neighbor switch).

Link Aggregation Group Identifier

IEEE 802.3ad specifies the standard implementation of PortChannels. PortChannel specifications provide LACP as a standard protocol, which enables negotiation of port bundling.

LACP makes misconfiguration less likely, because if ports are mismatched, they will not form a PortChannel.

ACP makes misconfiguration less likely, because if ports are mismatched, they will not form a PortChannel. Consider example A in the figure below, in which switch 1 connects

to switch 2. Port 1 on switch 1 connects to port 4 on switch 2, and port 2 on switch 1 connects to port 6 on switch 2.

Now imagine that the administrator configured a PortChannel on switch 1 between ports 1 and 2, while on switch 2 the PortChannel is configured between ports 5 and 3. Without LACP, the ports could be configured in channel-group mode, and you would not discover that this is a misconfiguration until you notice that traffic has dropped. LACP discerns that the only ports that can be bundled are port 1 going to port 4. According to the IEEE specifications, to allow LACP to determine whether a set of links connect to the same system and to determine whether those links are compatible with aggregation, you need to be able to establish two types of identification:

- You need a globally unique identifier for each system that participates in link aggregation (that is, the switch itself needs to be unique. This number is referred to as the system ID and is composed of a priority and a MAC address that uniquely identifies the switch.
- You need a means of identifying a link aggregation group.

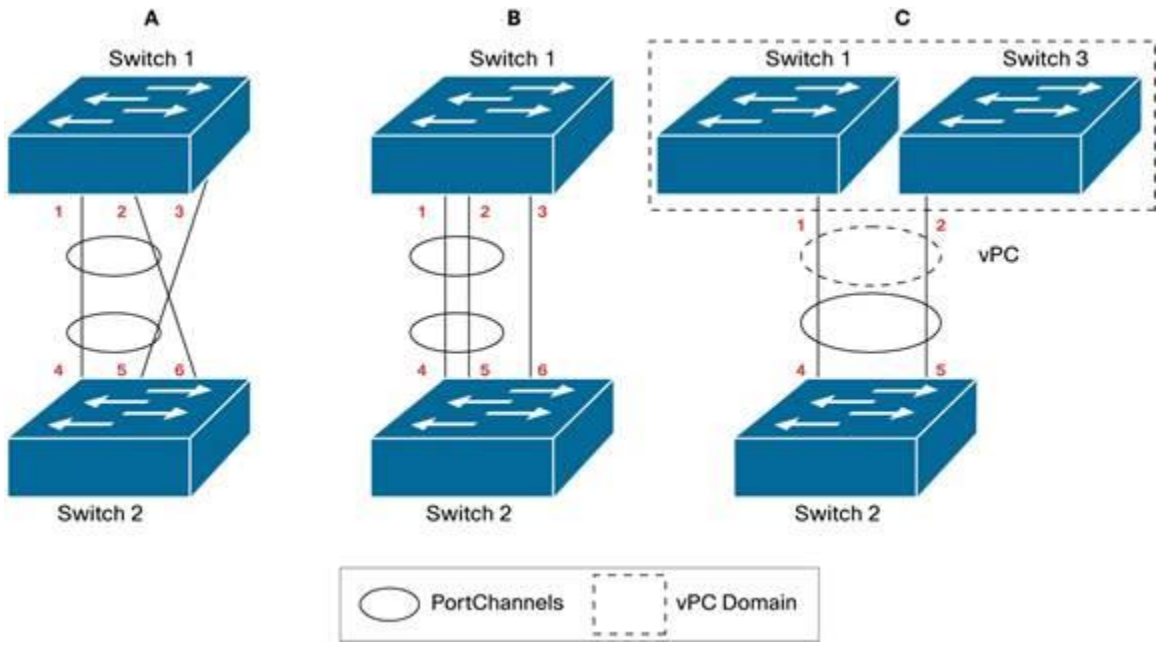

(Source: Cisco.com)

System ID in a vPC System

Spanning tree and LACP use the switch MAC address for, respectively, the bridge ID field in the spanning-tree BPDU and as part of LACP LAGID. In a single chassis, they use the systemwide MAC address for this purpose. For systems that use vPCs, use of the systemwide MAC address would not work because the vPC peers needs to appear as a single entity as shown in example C in Figure 8. To meet this requirement, vPC offers both an automatic configuration and a manual configuration of the system ID for the vPC peers.

The automatic solution implemented by vPC consists of generation of a system ID composed of a priority and a MAC address, with the MAC derived from a reserved pool of MAC addresses combined with the domain ID specified in the vPC configuration. The domain ID is encoded in the last octet and the trailing 2 bits of the previous octet of the MAC address.

By configuring domain IDs to be different on adjacent vPCs complexes (and to be identical on each vPC peer complex), you will help ensure the uniqueness of the system ID for LACP negotiation purposes. You also help ensure that the spanning-tree BPDUs use a MAC address that is representative of the vPC complex.

You can override the automatic generation of the system ID by using the command-line interface (CLI) and configuring the system ID on both vPC peers manually, as follows: (config-vpc-domain)#system-mac <mac>

Primary and Secondary vPC Roles

In a vPC system, one vPC switch is defined as primary and one is defined as secondary, based on defined priorities. The lower number has higher priority, so it wins. Also, these roles are non-preemptive, so a device may be operationally primary, but secondary from a configuration perspective. To understand the operational role of a

vPC member, you need to consider the status of the peer-keepalive link and the peer link.

When the two vPC systems are joined to form a vPC domain, the priority decides which device is the vPC primary and which is the vPC secondary. If the primary device were to reload, when the system comes back online and connectivity to the vPC secondary device (now the operational primary) is restored, the operational role of the secondary device (operational primary) will not change, to avoid unnecessary disruptions. This behavior is achieved with a sticky-bit method, whereby the sticky information is not saved in the startup configuration, thus making the device that is up and running win over the reloaded device. Hence, the vPC primary becomes the vPC operational secondary.

If the peer link is disconnected but the vPC peers are still connected through the vPC peer keepalive link, the vPC operational roles stay unchanged.
If both the peer link and peer-keepalive link are disconnected, both vPC peers become operational primary, but upon reconnection of the peer-keepalive link and the peer link, the vPC secondary device (operational primary) keeps the primary role, and the vPC primary becomes the operational secondary device.

Spanning Tree

vPC modifies the way in which spanning tree works on the switch to help ensure that the vPC peers in a vPC domain appear as a single spanning-tree entity on vPC ports. Also, vPC helps ensure that devices can connect to a vPC domain in a non-vPC fashion with classic spanning-tree topology. vPC is designed to support hybrid topologies. Depending on the Cisco NX-OS Software release, this can be achieved in slightly different ways.
In all Cisco NX-OS releases, the peer link is always forwarding because of the need to maintain the MAC address tables and Internet Group Management Protocol (IGMP) entries synchronized.

vPC by default ensures that only the primary switch forwards BPDUs on vPCs. This modification is strictly limited to vPC member ports. As a result, the BPDUs that may be received by the secondary vPC peer on a vPC port are forwarded to the primary vPC peer through the peer link for processing. Non-vPC ports operate like regular spanning-tree ports.

The special behavior of the primary vPC member applies uniquely to ports that are part of a vPC.
Starting from Cisco NX-OS Releases 4.2(6) and 5.0(2), vPC allows the user to choose the peer-switch option. This option optimizes the behavior of spanning tree with vPC as follows:

● The vPC primary and secondary are both root devices and both originate BPDUs
● The BPDUs originated by both the vPC primary and the vPC secondary have the same designated bridge ID on vPC ports
● The BPDUs originated by the vPC primary and secondary on non-vPC ports maintain the local bridge ID instead of the vPC bridge ID and advertise the Bridge ID of the vPC system as the root

The peer-switch option has the following advantages:

● It reduces the traffic loss upon restoration of the peer link after a failure.
● It reduces the disruption associated with a dual-active failure (whereby both vPC members become primary). Both devices keep sending BPDUs with the same bridge ID information on vPC member ports, which prevents errdisable from potentially disabling the PortChannel for an attached device.
● It reduces the potential loss of BPDUs if the primary and secondary roles change.

Cisco Fabric Services over Ethernet Synchronization Protocol

The vPC peers use the Cisco Fabric Services protocol to synchronize forwarding-plane information and implement necessary configuration checks.
vPC peers must synchronize the Layer 2 forwarding table - that is, the MAC address information between the vPC peers. This way, if one vPC peer learns a new MAC

address, that MAC address is also programmed on the Layer 2 forwarding table of the other peer device.
The Cisco Fabric Services protocol travels on the peer link and does not require any configuration by the user.

To help ensure that the peer link communication for the Cisco Fabric Services over Ethernet protocol is always available, spanning tree has been modified to keep the peer-link ports always forwarding. The Cisco Fabric Services over Ethernet protocol is also used to perform compatibility checks to validate the compatibility of vPC member ports to form the channel, to synchronize the IGMP snooping status, to monitor the status of the vPC member ports, to synchronize the Address Resolution Protocol (ARP) table (starting from Cisco NX-OS 4.2(6) and future Release 5.0 releases). If the peer link is disconnected between two vPC peers, the synchronization between vPC peers is interrupted, which may lead to traffic drop for multicast traffic and to flooding for unicast traffic.

vPC Configuration Changes When the Peer Link Fails

The correct sequence for setting up vPCs requires that the two participating vPC switches see each other over the peer link and that they can communicate over the vPC peer keepalive link.
Figure below illustrates this fundamental concept: the configuration depicted in Figure 9 (2) requires starting from Figure (1). If you try to configure a vPC like in Figure (2) without having established vPC peer-link and vPC peer keepalive connectivity, vPC ports won’t go into forwarding state.

Once you are in the state depicted in Figure (1) (which is a fully formed vpc domain) the peer keepalive connectivity is not strictly required in order to create or modify vPCs. In other words, you can configure the vPC in Figure (2) even if there’s a loss of vPC peer keepalive connectivity after the configuration depicted in Figure (1).

It is necessary and recommended to have functional vPC peer keepalive connectivity for the correct behavior of vPC in presence of failures, but from a traffic forwarding and configuration purpose, the temporary failure of the peer keepalive link doesn't have any impact.

The vPC peer-link connectivity is more important for the correct traffic forwarding operations as well as for the ability to create or modify vPCs.

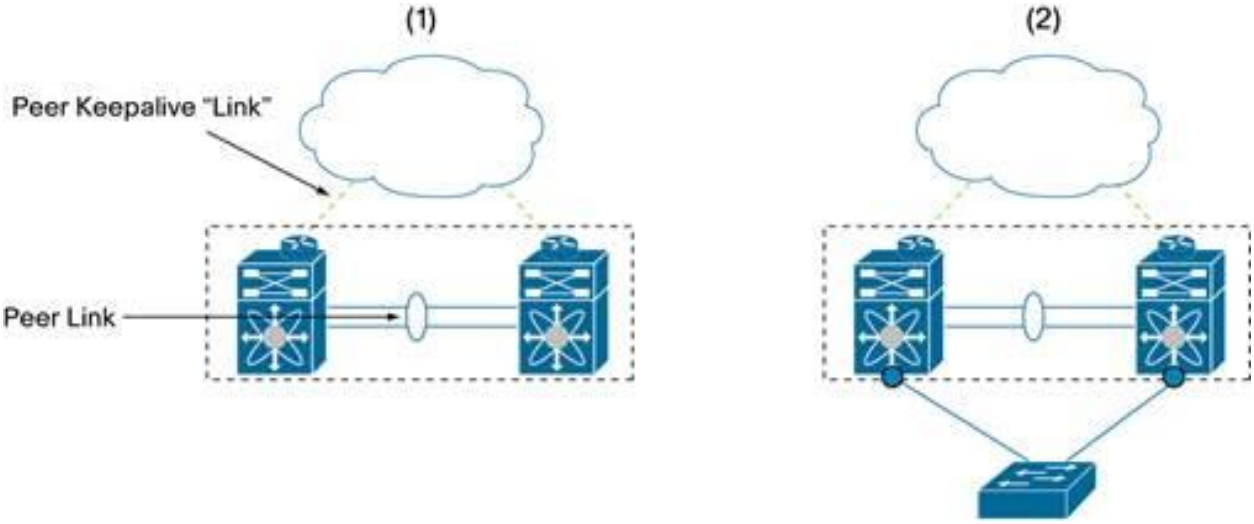

(Source: Cisco.com)

The initial implementation of vPC did not allow configuration changes when the peer link was disconnected to avoid, upon reconnection of the peer switch, inconsistencies that could bring these links down. Thus, if the peer-link was lost, the interfaces on the vPC primary could not be flapping (i.e. if they went down, they stayed down up until the peer-link was restored), and while the peer link was down, a new vPC member port could not be activated.

Example A in Figure shows the case of a single vPC peer. The user cannot activate any vPC member port until a vPC peer switch is present (example B in Figure).

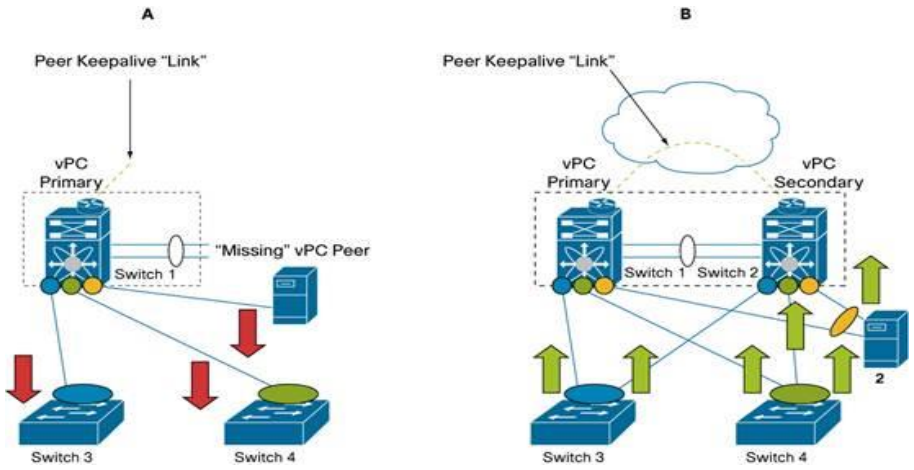

(Source: Cisco.com)

Because of this behavior, if the peer-link connection is lost, by default the user cannot add any vPC ports and activate them, nor can an interface flap. If a vPC interface flaps, the port will stay down after flapping. For example, imagine a vPC setup with PortChannel 8 configured as vPC 8:

vPC status

---

id Port Status Consistency Reason Active vlans

------ ----------- ------ ----------- ------------------------- -----------

8 Po8 up success success 23,50

After the peer-link failure only the primary keeps the vPC interfaces up. If the interface associated with PortChannel 8 flaps, it never goes up again.

vPC status

---

id Port Status Consistency Reason Active vlans

------ ----------- ------ ----------- ------------------------- -----------

8 Po8 down failed Peer-link is down

To restore connectivity, you need to first restore the peer link, and then enter a shut/no shut command under the interface on the primary device. Similarly, if you want to create a new vPC interface and to activate it, a vPC peer needs to be present and connected

through the peer link. While a peer link is not connected, vPC prevents activation of any new vPC member port.

This vPC behavior may manifest itself in several scenarios:

● A vPC pair in which the peer link is lost but the peer-keepalive is still connected (shown as case 1 in Figure)

● A vPC pair in which peer link and peer-keepalive links are lost (split brain)

● A switch that was part of a vPC but has been reloaded; upon coming online, the vPC peer is unavailable (shown as case 2 in Figure)

● A vPC switch that has never been part of a vPC domain because it has just been powered up and configured for vPC, but no peer-link and peer-keepalive connectivity has been established yet

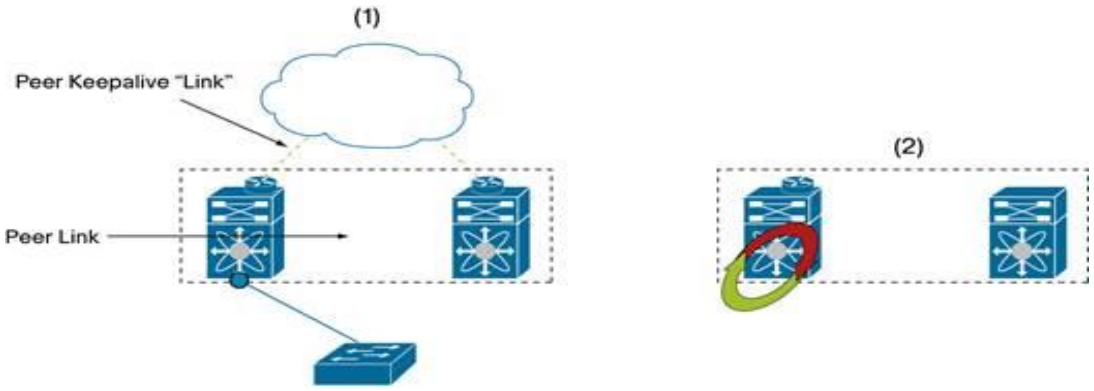

(Source: Cisco.com)

Case 1 has been addressed by the Cisco Nexus 7000 Series with Cisco NX-OS 4.2(3). This solution allows the vPC device to shut/no shut existing vPC member ports as long as the vPC secondary device can still be reached through the peer-keepalive link. This same behavior on the Cisco Nexus 5000 Series requires configuration of the keyword peer-check-config-bypass under the vPC domain configuration, but this process is superseded with NXOS 5.0(2)N1(1).

For case 1, to add a new vPCs you need the reload restore command of the Cisco Nexus 7000 Series and on the Cisco Nexus 5000 Series running NXOS 5.0(2)N1(1) or

higher. You can achieve the same results with peer-check-config-bypass for the Cisco Nexus 5000 Series if running an earlier version of code.

The case of complete disconnection between the vPC primary and secondary devices (split brain) is addressed in the Cisco Nexus 7000 Series or Cisco Nexus 5000 Series running NXOS 5.0(2)N1(1) or higher (for existing vPC member ports) and reload restore (for new vPC member ports). The equivalent command on the Cisco Nexus 5000 Series for earlier releases is peer-check-config-bypass.

The third case, reload of a vPC device, is addressed with the vPC reload restore command.
To override the default vPC behavior, which prevents activation of new vPC member ports when the peer link is down, you can use the command peer-config-check-bypass under the vPC domain configuration (on the Cisco Nexus 5000 Series only).

As an example:
vpc domain 2
role priority 100
peer-keepalive destination 10.51.35.18
peer-config-check-bypass

The peer-config-check-bypass command was introduced in Cisco NX-OS 4.1(3)N2(1), and it will be superseded by the vPC reload restore command when that command is available.
With this command in place, even if the Cisco Nexus 5000 Series peer link is down, a vPC member port can flap, and you can also create new vPC member ports and activate them.

vPC Reload Restore

The vPC reload restore feature was introduced with Cisco NX-OS 5.0(2). With vPC reload restore, when the peer link is down, or when the both peer link and peer-keepalive links are lost, the vPC primary can activate new vPCs ports.

If the vPC peer switch is brought online or connected to the existing vPC peer switch, the configuration checks are performed; if inconsistencies are found, the vPC ports are shut down.

In addition, with reload restore a vPC peer can reload, and after a reload, if no vPC peer exists (a condition that is detected by a timer), the user can create vPCs on the standalone vPC peer.

Upon reload, Cisco NX-OS starts a user-configurable timer (with a default of 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 have been received, Cisco NX-OS assumes the primary role. The software reinitializes the vPCs, activating its local ports. Because there are no peers, the consistency check is bypassed for the local vPC ports.

The timer is user configurable and defines how long the standalone vPC device waits to detect a vPC peer. If at the timer expiration no peer-keepalive- or peer-link-up packets have been received, the software reinitializes the vPCs, activating its local ports. Because there are no peers, the consistency check is bypassed for the local vPC ports.

The following output shows the status of a virtual PortChannel configured on a standalone vPC system with restore reload:

```
--------------------------------------------------------------------------
id Port Status Consistency Reason Active vlans
-- ---- ------ ----------- ------------------------ ------------
51 Po51 up success Type checks were bypassed 10-14, 21-24
for the vPC ,50,60
```