vPC

Terminology

vPC Domain containing the 2 peer devices. Each peer device processes half of the traffic coming from the access layer. Only 2 peer devices max can be part of same vPC domain

vPC Domain ID used by peer devices to assign a unique vPC system MAC address, Domain ID should be same for all vPC pair defined in a contiguous layer 2 domain but should be different in double-sided(back-to-back) vPC topology.
vPC system MAC identifies the logical switch

vPC Peer a vPC switch (one of a Nexus 7K/6K/5K Series pair)

vPC The Combined port-channel between the vPC peers and the downstream device.vPC port can be trunk or access mode

Enhanced vPC & vPC+ term used when nexus 2K (FEX) is controlled by two parents (dual-homed FEX) & in vPC+ peer link will be in mode fabric path instead of 802.1q trunk/access

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 Normal STP VLAN that is not part of any vPC and not carried over the peer-link

CFS Cisco Fabric Services Protocol, underlying protocol running on top of vPC peer-link and provide synchronization and consistency check between the 2 peer devices

Orphan Device A device that connected to only one vPC peer and is on a vPC VLAN

Orphan Port an interface which connects to an orphan device

vPC Role there are two defined vPC roles: primary & secondary.vPC defines which of the two vPC peer devices processes Bridge Protocol Data Units (BPDU) and respond to ARP request.

vPC Components

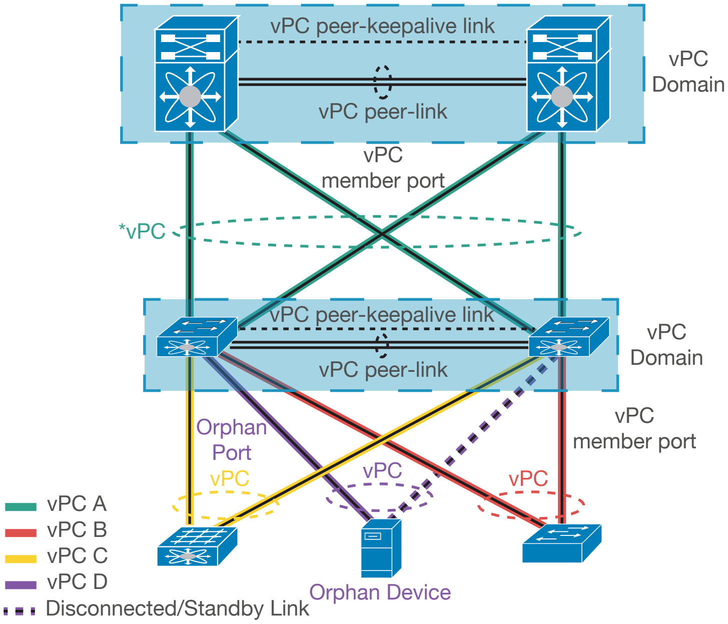

* vPC A requires two seperate vPC on each domain (Double-Sided vPC)

vPC Port Type

vPC Peer-link layer 2 trunk link used to sync control plane (CAM Table, ARP Cache, IGMP Snooping, HSRP, etc.) between vPC peers, normally not used for the data plane but able to carry vPC vlan traffic

vPC Peer-keepalive Link layer 3 link used as heartbeat in the control plane to prevent Active/Active or split brain (Dual Active) vPC roles,it can be Mgmt0 port, L3 VLAN Interface, point-to-point routed link

vPC Member Ports one of a set of ports (from port-channel) that form a vPC towards downstream neighbor

VSS vs vPC

	Multi-Chassis Port Channel	Control Plane	Layer 3 Port-Channels	Switch Configuration	Maximum Physical Node
vPC	Yes	Two independent Nodes, both active	No	Common Configs (Consistency Checker)	2
VSS	Yes	Single Logical Node	Yes	Combined Configs	2

vPC Configuration

```
# 1-Enable vPC & LACP feature globally
feature vpc/lacp

# 2-Configure peer keepalive interface (Mgmt or L3)
interface mgmt0 or eth 1/10
  no switchport
  vrf memeber MANAGEMENT
  ip address 1.1.1.1/24

# 3-Create vPC domain & define peer keepalive address
vpc domain 10
  peer-keepalive destination 1.1.1.2 source
  1.1.1.1 vrf MANAGEMENT
```

```
# 4-Establish port channel for vPC peer link
interface po 1
  switch mode trunk
  vpc peer-link
interface eth 1/1-2
  channel-group 1 mode active force

# 5-Verify vPC consistency parameters and create vPC
interface port-channel 10
  switchport mode trunk
  spanning-tree port type edge trunk
  vpc 10
interface eth 1/3
  channel-group 10 mode active force
```

vPC

vPC Failure Scenarios

vPC Member Port Fails if one vPC member port fails, the host detects a link failure on one of the port-channel members, it then redirect the traffic over remaining port channel members. after the failure all traffic points to secondary member port. This is one of the scenarios where a vPC peer link will used to carry data traffic. Enough bandwidth is required on vPC peer link

vPC Peer Link Failure if one vPC peer link goes down,the vPC secondary switch shuts down all of its vPC member ports if it can still receive keepalive messages from vPC primary switch and vPC primary switch keeps all of its interfaces up, as a best practice vPC peer link should be at least two physical 10G ports as the vPC peer link

vPC Peer Switch Failure when one peer switch fails, half of the network bandwidth is lost and the remaining vPC switch maintains the network connectivity.If failure occurs on a primary switch,the secondary switch becomes the operational primary switch. If the primary switch comes back again it will take the role of vPC operational secondary

vPC Peer Keepalive Link Failure keepalive link carries the heartbeat message between two vPC peer switches. During keepalive link failure there is no impact on traffic flow and vPC operation. Keepalive link should be up to avoid a dual active scenario

vPC Keepalive Link Failure Followed by Peer Link Failure if keepalive link fail first and then peer link fails,the vPC secondary switch assumes the primary role and keeps its vPC member ports up.If both switch are still operational but only the link between them fails, in this situation, both vPC switches claim the primary switch role and keep the vPC member ports up, this is known as a split-brain scenario.

vPC Consistency Check

CFS runs on the vPC peer link to synchronize the control plane, includes advertisement of “consistency parameters” that must match for vPC to form successfully. E.g line card type (M or F) , Speed, Duplex, Trunking, LACP, STP, etc

Type 1 Puts peer device or interface into a suspended state to prevent invalid packet forwarding, with graceful consistency check only the secondary peer device will be suspended. Parameters such as STP Configuration, STP global setting (BA, loop gaurd, BPDU filter,...), Port-channel LACP mode, link speed, duplex, MTU, allowed VLANs, native VLAN

Type 2 peer device or interface still forward traffic. However they are subject to undesired packet forwarding behavior. Parameters such as configuration of MAC aging timers, static MAC entries, VLAN interface (SVI), ACL, QoS, VLAN database, port security, Cisco TrustSec, DHCP & IGMP snooping, FHRP, PIM, routing protocol configuration

Graceful Consistency-Check any type 1 inconsistency leads all member ports on both vPC legs to goes down, with graceful consistency check, only secondary vPC member ports will be shut and primary device will still work (enable by default)

vPC Peer-Switch

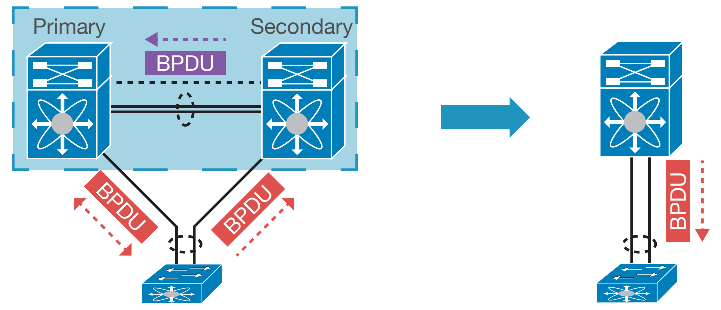

* vPC secondary device proxies BPDU’s to primary

vPC enhancement which allows a pair of vPC peer devices to appear as a single STP root in layer 2 topology (they have same bridge ID). Without it, vPC secondary peer device proxies any BPDU messages from access switch to primary vPC peer device

vPC Auto-Recovery

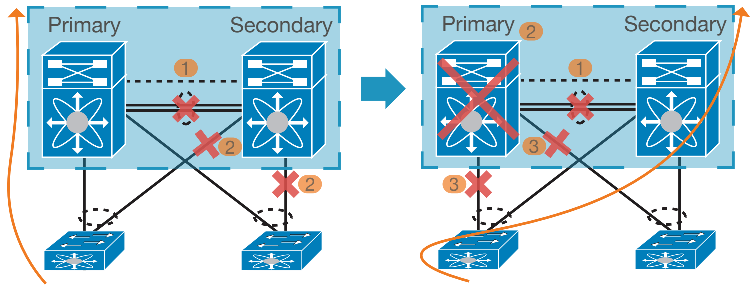

Provide a backup mechanism in case of vPC peer-link failure followed by vPC primary peer device failure and also to handle a specific case where both vPC peer devices reload but only one comes back to life

vPC Peer-Gateway

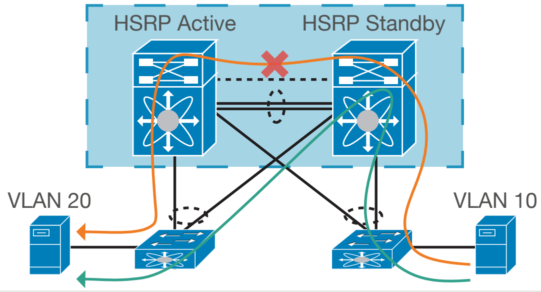

Allows a standby vPC peer device to act as the active gateway for packet addressed to other peer device, keeps the traffic forwarding local and avoids use of peer-link. Also address the problem of some device such loadbalancer and NAS that do not perform a typical default gateway ARP request at bootup

vPC Orphan-Port Suspend

Was developed for single-attached devices to vPC domain or working in active/standby mode (server, firewall or load-balancer).When a vPC peer-link goes down, the vPC secondary peer device shuts all of it’s vPC member ports, but it does not shut down orphan ports, with Orphan-Port suspend, an orphan port is also shutdown along with the vPC member ports