



*TOMORROW
starts here.*

Cisco *live!*



Deploying Virtual Port Channel (vPC) in NX-OS

BRKDCT-2048

Sutharsan Sivapalan

Customer Support Engineer

#clmel

Cisco *live!*

Session Abstract

- This session is targeted at Network Engineers, Network Architects and IT Administrators who have deployed or are considering the deployment of vPC to improve Layer 2 scalability and network operational efficiency.
- Session introduces basic concepts and terminology of virtual Port-Channels & also covers designs and best practices of the vPC technology.
- vPC in relation to FabricPath, VXLAN and ACI will be briefly discussed in this session.
- Nexus 2000 (FEX) will only be addressed from a vPC standpoint.
- vPC troubleshooting will not be covered in the session.
- The presentation includes hidden and reference slides.



For Your
Reference

Agenda

- Feature Overview
- vPC Design Guidance and Best Practices
- vPC Enhancements
- vPC and Fabric Technologies



Agenda

- **Feature Overview**
 - vPC Concept & Benefits
 - How does vPC help with STP?
 - vPC Terminology
 - Data-Plane Loop Avoidance with vPC
- vPC Design Guidance and Best Practices
- vPC Enhancements
- vPC and Fabric Technologies

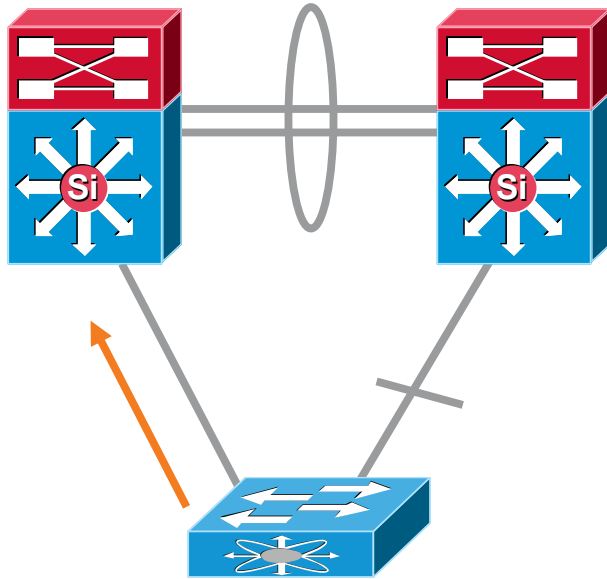


A long-exposure photograph of a city street at night. The foreground is filled with vibrant, multi-colored light trails from moving vehicles, creating a sense of motion. In the background, a pedestrian bridge spans the street, and various city buildings are illuminated with lights. The overall scene is a dynamic urban environment.

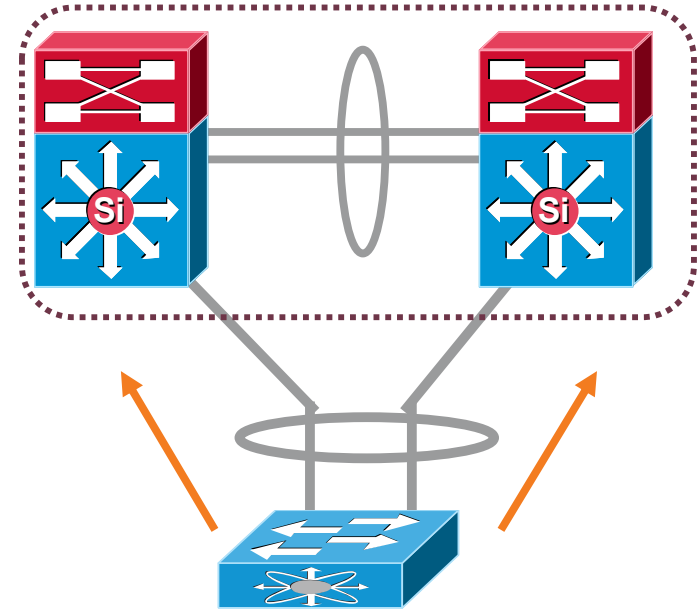
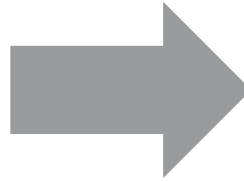
vPC Feature Overview

vPC Feature Overview

vPC Concept & Benefits



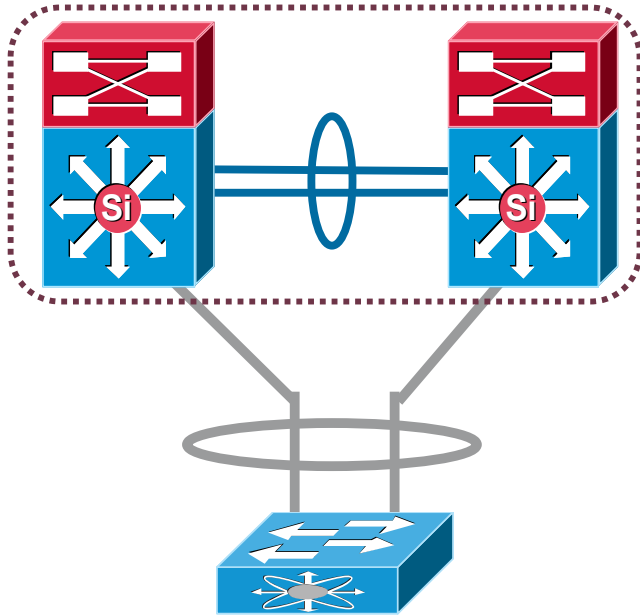
Non-vPC



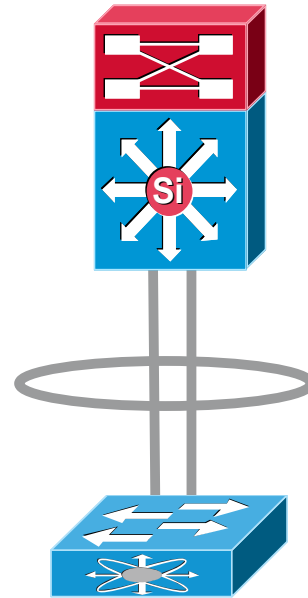
vPC

vPC Feature Overview

vPC Concept & Benefits



Physical Topology



Logical Topology

- Simplicity
- Redundancy
- Efficiency
- Fast-Convergence

Feature Overview

How does vPC help with STP? (1 of 2)

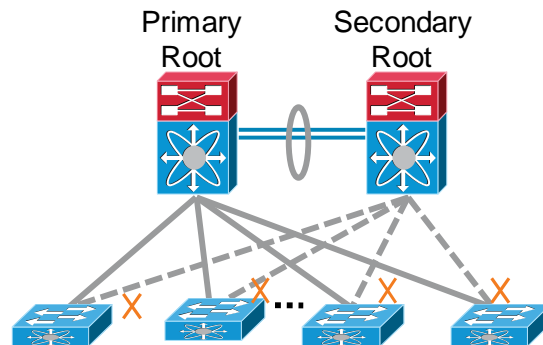
➤ Before vPC

STP blocks redundant uplinks

VLAN based load balancing

Loop Resolution relies on STP

Protocol Failure → 💣



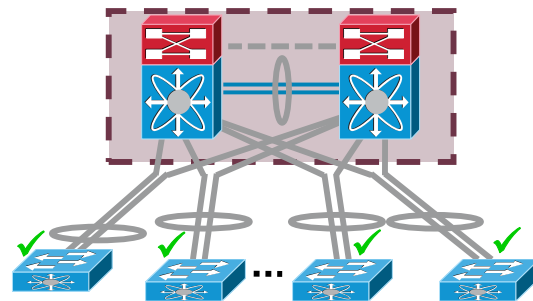
➤ With vPC

No blocked uplinks

EtherChannel load balancing (hash)

Loop Free Topology

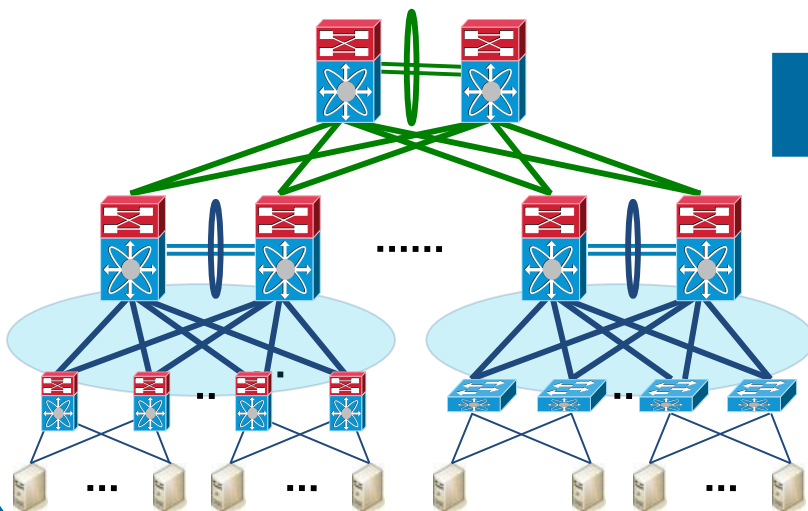
Lower oversubscription



Feature Overview

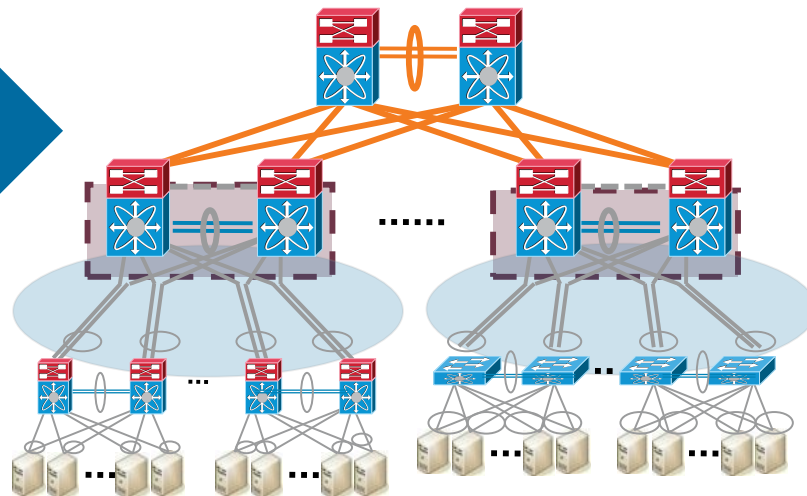
How does vPC help with STP? (2 of 2)

Reuse existing infrastructure



Smooth
Migration

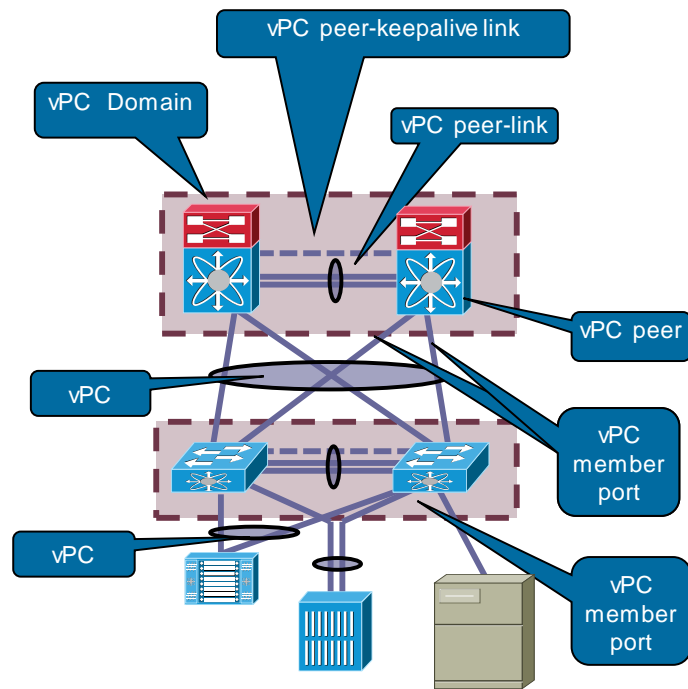
Build Loop-Free Networks



Feature Overview

vPC Terminology (1 of 2)

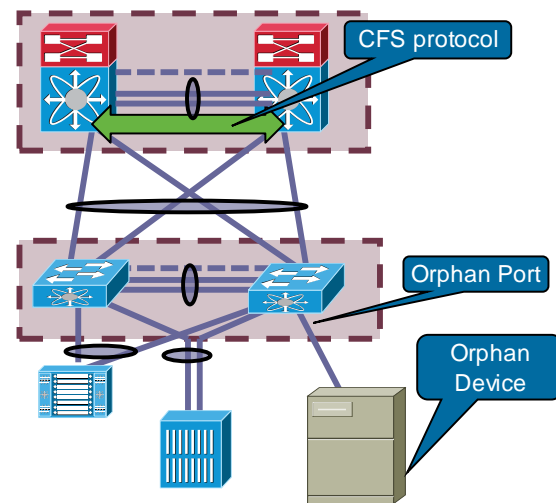
- **vPC Domain** - A pair of vPC switches in a vPC system
- **vPC Peer** - A vPC switch, one of a pair
- **vPC member port** - One of a set of ports (port channels) that form a vPC
- **vPC** - Port channel between the vPC peer and the downstream device
- **vPC peer-link** - Link used to synchronise state between vPC peer devices
- **vPC peer-keepalive link** - Link used to monitor liveness of the peer device



Feature Overview

vPC Terminology (2 of 2)

- **vPC VLAN** - VLAN carried over the peer-link and used to communicate via vPC with a peer device
- **Non-vPC VLAN** - STP VLAN, not carried over the peer-link
- **CFS** - Cisco Fabric Services protocol, used for state synchronisation and configuration validation between peer devices
- **Orphan Device** – A device which is on a vPC vlan but only connected to one vPC peer
- **Orphan Port** – An interface which connects to an orphan device

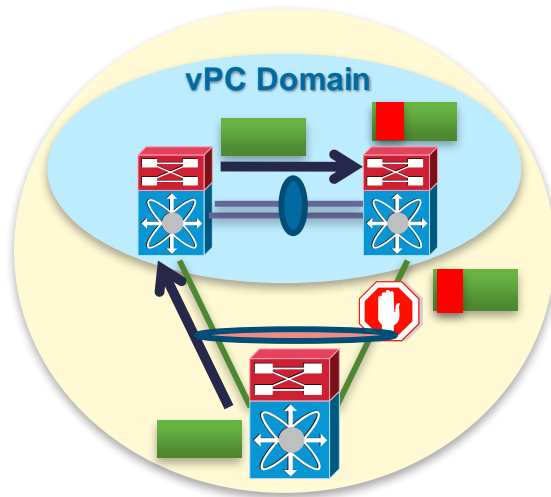
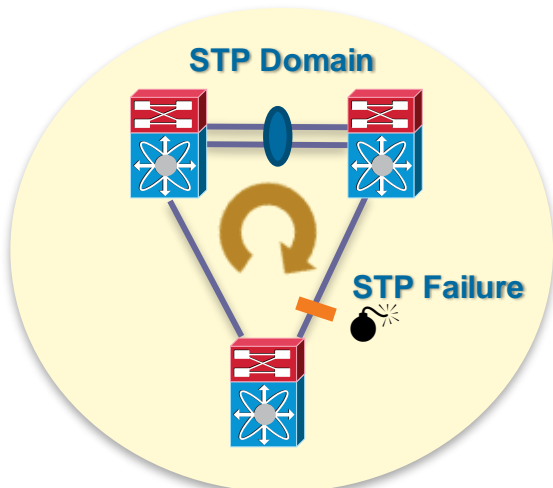


Feature Overview

Data-Plane Loop Avoidance with vPC (1 of 2)

Data-Plane vs. Control-Plane Loop Avoidance

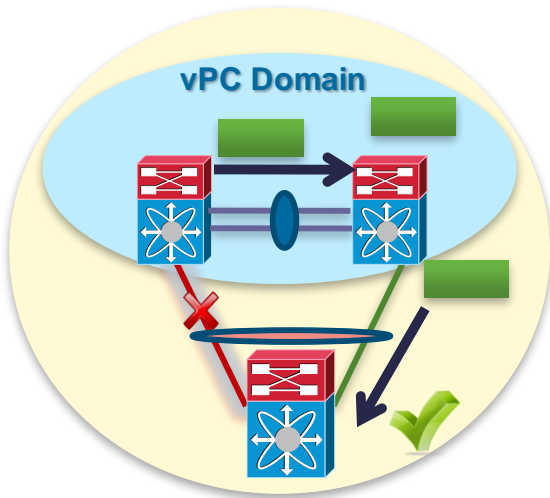
- vPC peers can forward all traffic locally
- Peer-link typically does not forward data packets (control plane extension)
- Traffic on the Peer-link is marked and not allowed to egress on a vPC



Feature Overview

Data-Plane Loop Avoidance with vPC (2 of 2)

- Exception for single-sided vPC failures
- Peer-link used as backup path for optimal resiliency





vPC Design Guidance and Best Practices

Agenda

- Feature Overview
- vPC Design Guidance and Best Practices
 - Building a vPC Domain
 - Mixed Chassis Mode
 - Attaching to a vPC Domain
 - Layer 3 and vPC
 - Spanning Tree Recommendations
 - FHRP with vPC
 - vPC / FEX Supported Topologies
- vPC Enhancements
- vPC and Fabric Technologies



Configuration Steps

Following steps are needed to build a vPC
(Order does Matter!)

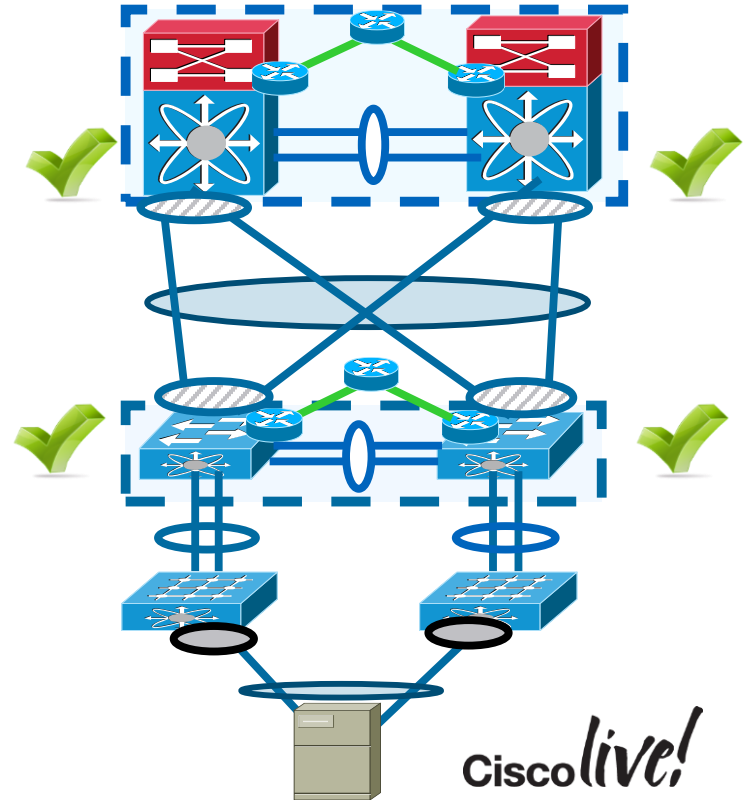
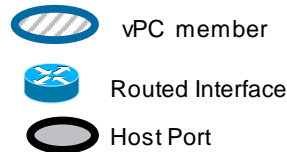
Define domains*

Establish Peer Keepalive connectivity

Create Peer link

Reuse port-channels and Create vPCs

Make Sure Configurations are Consistent



Building a vPC Domain

vPC Domains

- vPC Domain defines the pair of switches participating in the vPC
- Provides for definition of global vPC system parameters
- vPC domain ID used to assign unique vPC system MAC address
- You **MUST** utilise *unique* domain ID for each vPC pair

```
! Configure the vPC Domain ID - It should be unique within the layer 2 domain
NX-1(config)# vpc domain 20
```

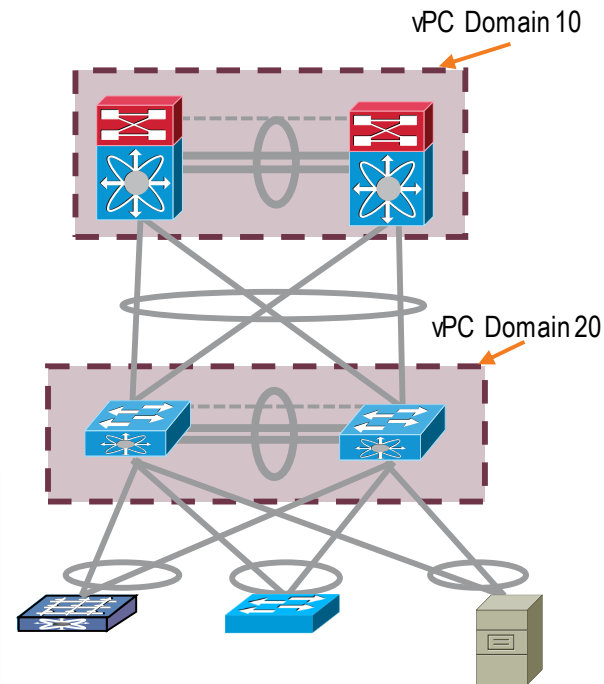
```
! Check the vPC system MAC address
```

```
NX-1# show vpc role
```

```
<snip>
```

```
vPC system-mac           : 00:23:04:ee:be:14
```

vPC System MAC identifies the Logical Switch in the network topology



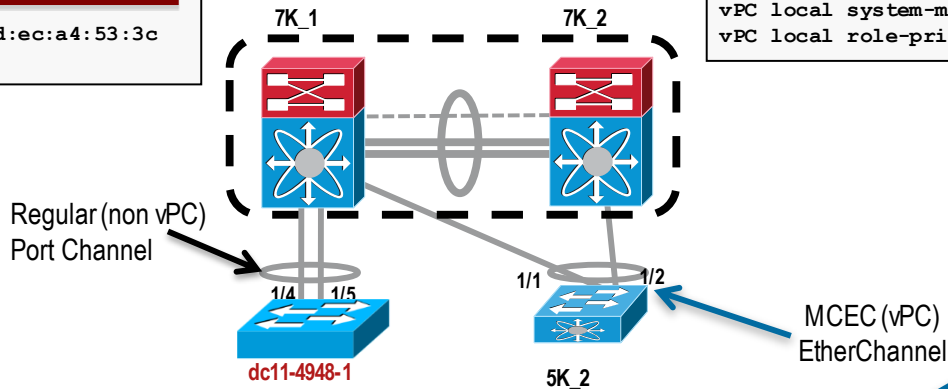
Building a vPC Domain

Independent Control Plane + Synchronised L2 State

- LACP neighbour sees the same System ID from both vPC peers
- The vPC 'system-mac' is used by both vPC peers

```
7K_1# sh vpc role
<snip>
vPC system-mac          : 00:23:04:ee:be:14
vPC system-priority      : 1024
vPC local system-mac     : 00:0d:ec:a4:53:3c
vPC local role-priority  : 1024
```

```
7K_2 # sh vpc role
<snip>
vPC system-mac          : 00:23:04:ee:be:14
vPC system-priority      : 1024
vPC local system-mac     : 00:0d:ec:a4:5f:7c
vPC local role-priority  : 32667
```



```
5K_2#sh lacp neighbor
<snip>
```

Port	Flags	LACP port	Priority	Dev ID	Age	Admin key	Oper Key	Port Number	Port State
E1/1	SA	32768	0023.04ee.be14	9s	0x0	0x801E	0x4104	0x3D	
E1/2	SA	32768	0023.04ee.be14	21s19	0x0	0x801E	0x104	0x3D	

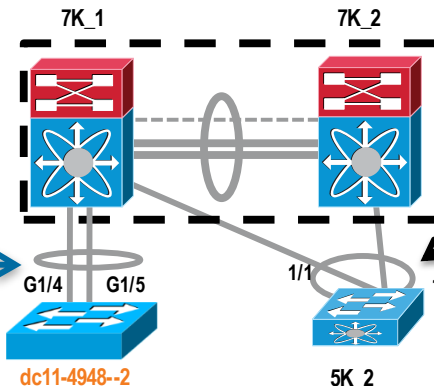
Building a vPC Domain

Independent Control Plane + Synchronised L2 State

- vPC peers function as independent devices as well as peers
- Local** 'system-mac' is used for all **non vPC** PDUs (LACP, STP, ...)

```
7k_1 # sh vpc role
<snip>
vPC system-mac          : 00:23:04:ee:be:14
vPC system-priority     : 1024
vPC local system-mac    : 00:0d:ec:a4:53:3c
vPC local role-priority : 1024
```

Regular (non vPC)
Port Channel



MCEC (vPC)
EtherChannel

```
dc11-4948-2#sh lacp neighbor
<snip>
```

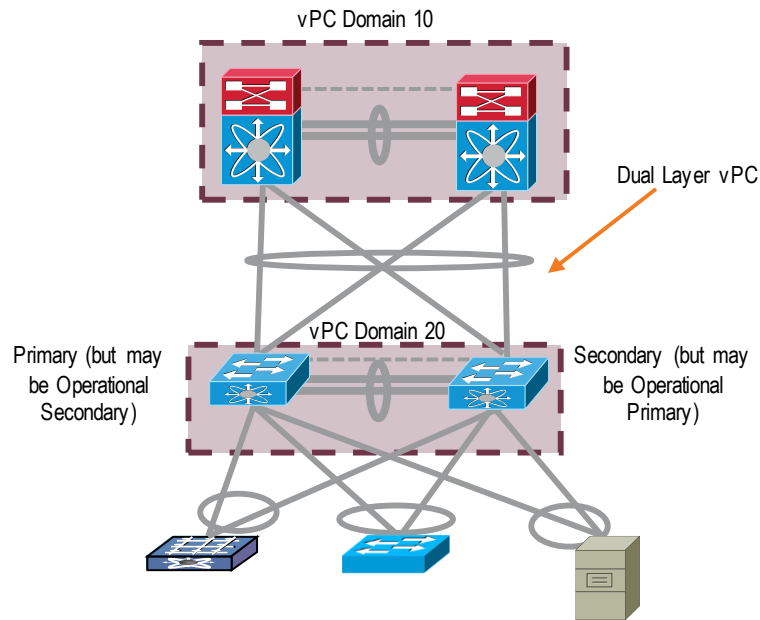
Port	Flags	LACP port Priority	Dev ID	Age	Admin key	Oper Key	Port Number	Port State
Gi1/4	SA	32768	000d.eca4.533c	8s	0x0	0x1D	0x108	0x3D
Gi1/5	SA	32768	000d.eca4.533c	8s	0x0	0x1D	0x108	0x3D

Building a vPC Domain

vPC Roles

- vPC primary switch election is based on role priority
- Lower priority wins, if tie, lower system MAC wins
- Role is non-preemptive, so operational role is what matters and not configured role
- vPC role defines which of the two vPC peers processes BPDUs
- Role matters for the behaviour with peer-link failures!

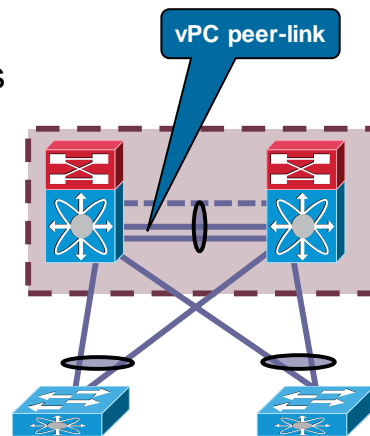
```
dc11-5020-3(config-vpc-domain)# role priority ?  
<1-65535> Specify priority value  
  
dc11-5020-3# sh vpc  
<snip>  
vPC role                : secondary, operational primary
```



Building a vPC Domain

vPC Peer-Link (a.k.a vPC PL a.k.a MCT)

- **Definition:**
 - Standard 802.1Q Trunk which carries CFS (Cisco Fabric Services) messages
 - Carries flooded traffic from vPC peer, STP BPDUs, HSRP Hellos, IGMP updates, etc.
- **Peer-Link member ports must be 10/40/100GE interfaces**
- **Peer-Link must be a point-to-point link**



Recommendations (strong ones!)

- Minimum: 2 x 10GE ports
- Use 2 separate cards for best resiliency
- 10GE ports in dedicated mode for oversubscribed modules

The *peer link* is **always** forwarding for any VLAN that is a member !

Building a vPC Domain

vPC Peer-Keepalive Link (aka vPC PKL)

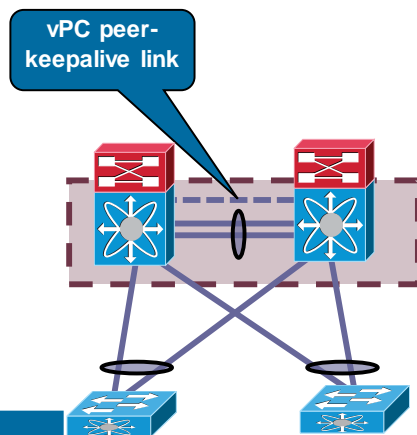
Definition:

- Heartbeat between vPC peers
- Active/Active detection (in case vPC Peer-Link is down)

Packet Structure:

- UDP message on port 3200, 96 bytes long (32 byte payload), includes version, time stamp, local & remote Ips and domain ID
- Default timers: interval 1sec / timeout 5sec

*vPC PKL messages should
NOT
be routed over the vPC PL !*



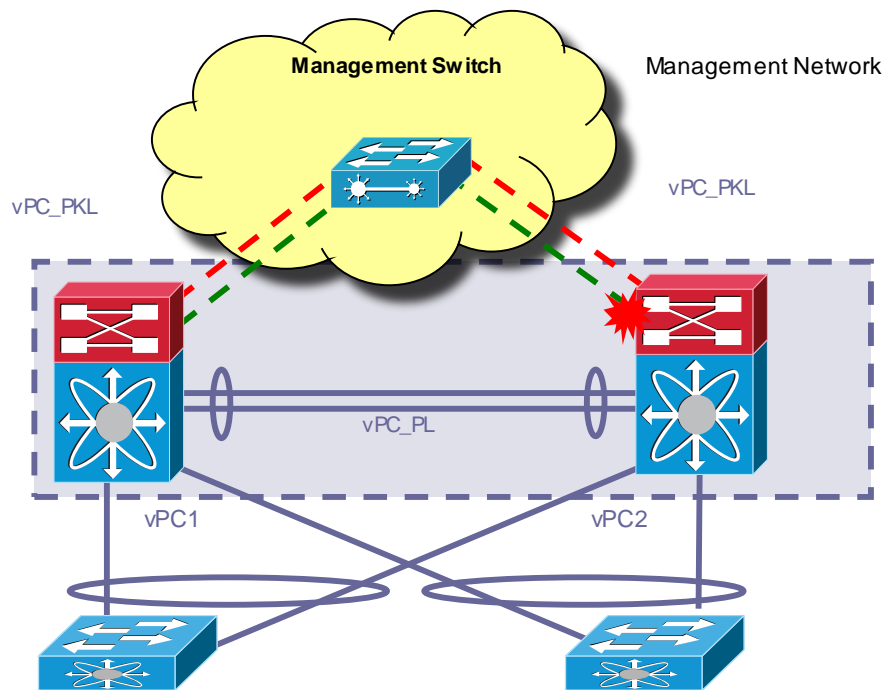
Recommendations
(in order of preference):



NEXUS 7000 /Nexus 7700	NEXUS 5000/5500/Nexus 6000
1- Dedicated link(s) (1GE/10GE LC)	1- mgmt0 interface (along with management traffic)
2- mgmt0 interface (along with management traffic)	2- Dedicated link(s) (1/10GE front panel ports)
3- As last resort, can be routed over L3 infrastructure	3 - As last resort, can be routed over L3 infrastructure

Building a vPC Domain

vPC Peer-Keepalive Link – Dual Supervisors on Nexus 7000

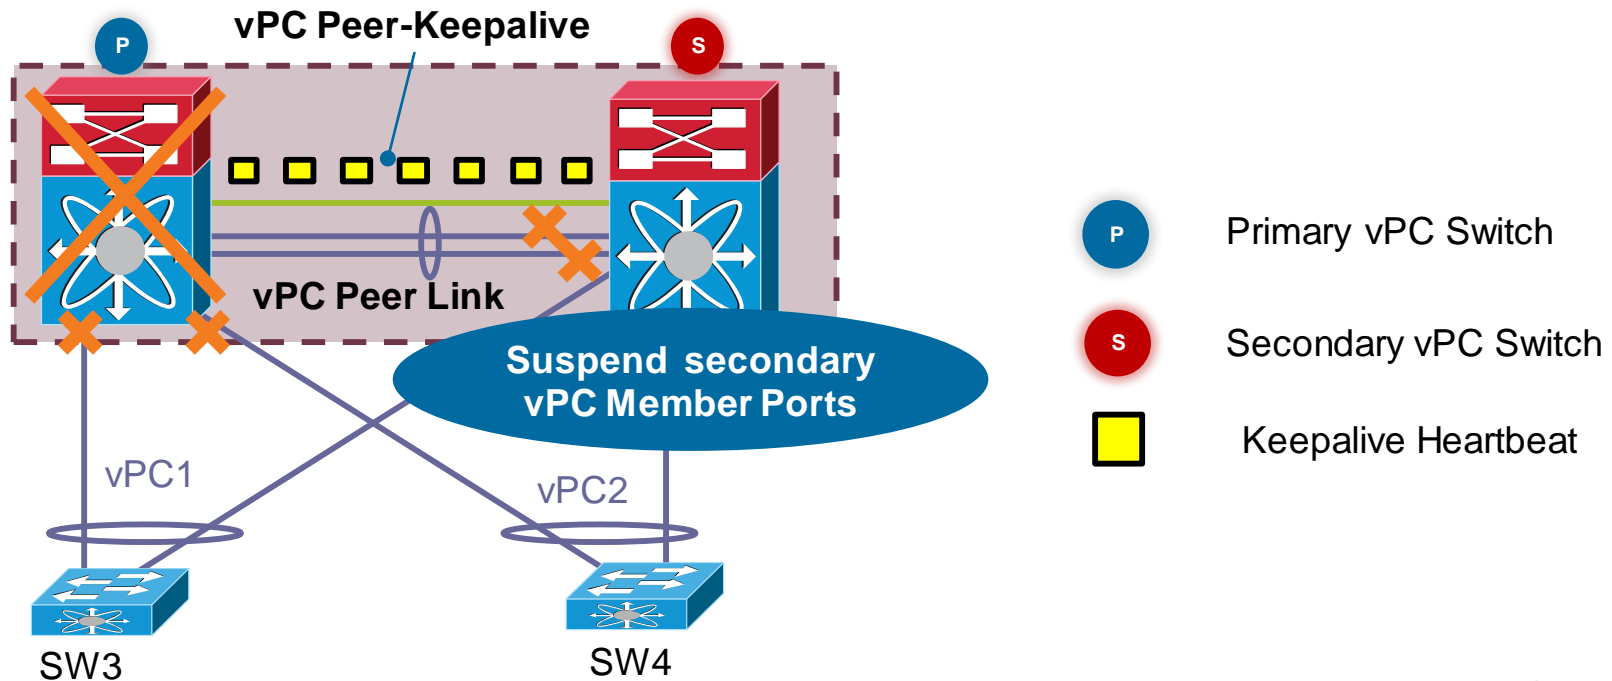


Do **NOT** use back-to-back mgmt0 connections on Nexus 7000 with Dual Supervisors

- Standby Management Interface
- Active Management Interface

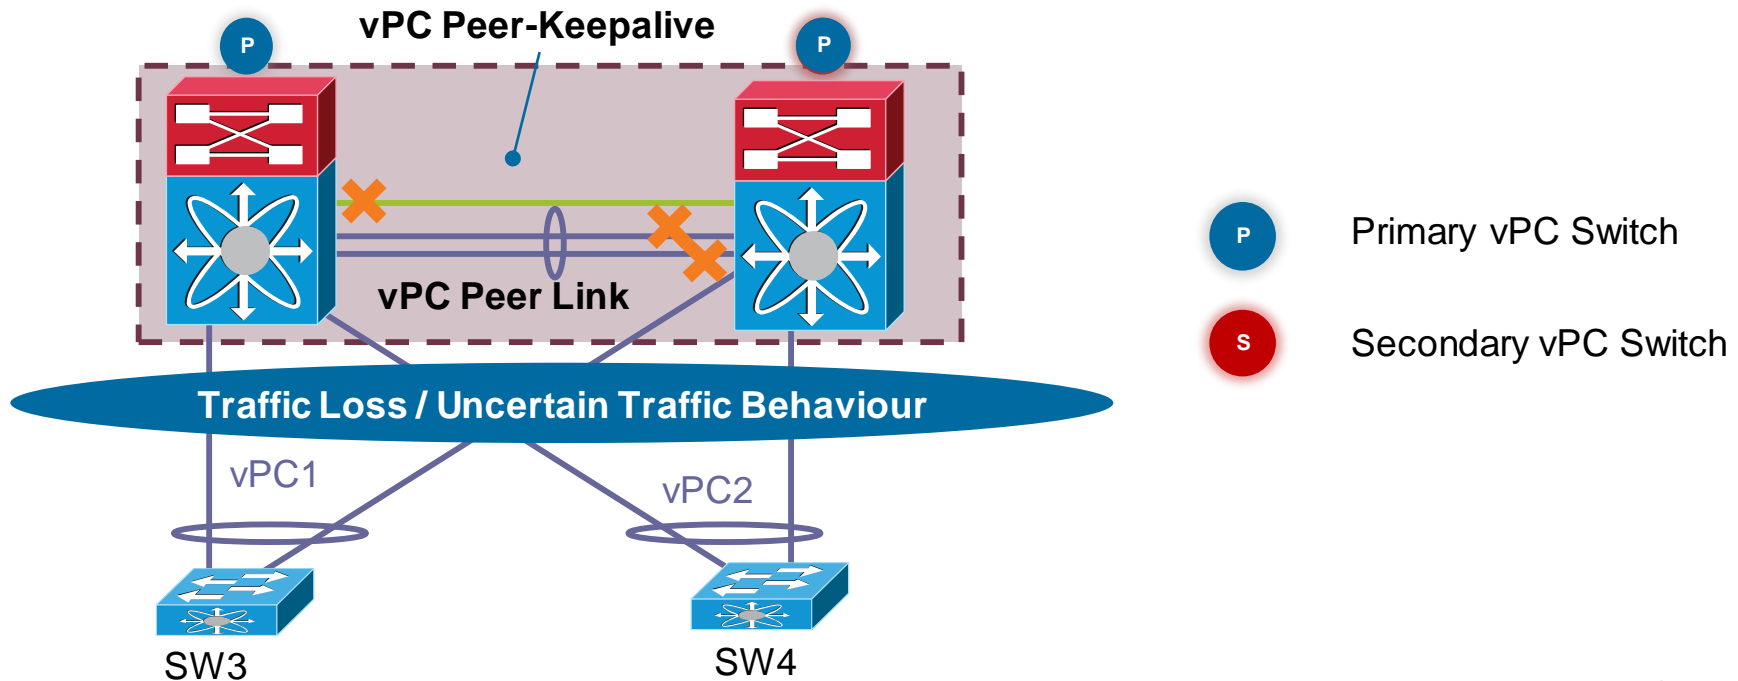
vPC Failure Scenario

vPC Peer-Keepalive Link up & vPC Peer-Link down



vPC Failure Scenario

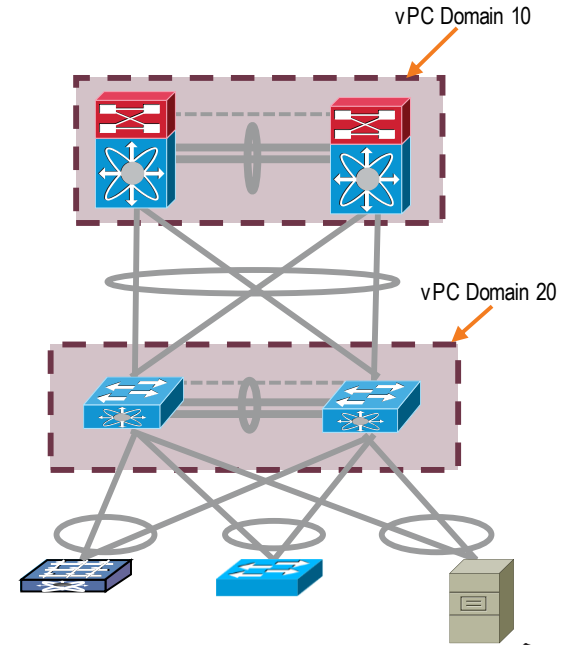
vPC Peer-Keepalive Link up & vPC Peer-Link down



vPC Configuration Consistency

vPC Control Plane - Consistency Check

- Both switches in the vPC Domain maintain distinct control planes
- CFS provides protocol state sync between peers (MAC Address table, IGMP state, ...)
- System configuration must also be kept in sync
- Two types of interface consistency checks:
 - Type 1** – Will put interfaces into suspend state to prevent incorrect forwarding of packets. With Graceful Consistency Check (5.2 & later), we only suspend on secondary peer.
 - Type 2** – Error messages to indicate potential for undesired forwarding behaviour.



vPC Configuration Consistency

vPC Control Plane – Type 1 Consistency Check

- **Type 1 Consistency Checks** are intended to prevent network failures
- Incorrect forwarding of traffic
- Physical network incompatibilities

```
dc11-5020-1# sh run int po 201
```

```
interface port-channel201
 switchport mode trunk
 switchport trunk native vlan 100
 switchport trunk allowed vlan 100-105
 vpc 201
 spanning-tree port type network
```

```
dc11-5020-2# sh run int po 201
```

```
interface port-channel201
 switchport mode trunk
 switchport trunk native vlan 100
 switchport trunk allowed vlan 100-105
 vpc 201
 spanning-tree port type network
 spanning-tree guard root
```

```
dc11-5020-2# show vpc brief
```

Legend:

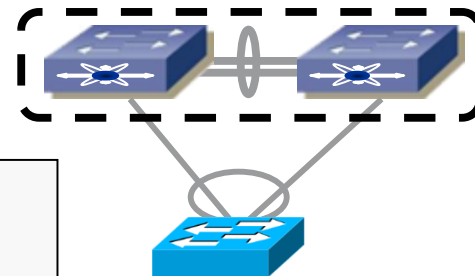
(*) - local vPC is down, forwarding via vPC peer-link

<snip>

vPC status

id	Port		Status	Consistency Reason	Active vlans

201	Po201	up	failed	vPC type-1 configuration incompatible - STP interface port guard - Root or loop guard inconsistent	-



**“vPC will be
suspended”**

vPC Configuration Consistency

vPC Control Plane – Type 2 Consistency Check

- Type 2 Consistency Checks are intended to prevent undesired forwarding
- vPC will be modified in certain cases (e.g. VLAN mismatch)

```
dc11-5020-1# sh run int po 201
version 4.1(3)N1(1)

interface port-channel201
 switchport mode trunk
 switchport trunk native vlan 100
 switchport trunk allowed vlan 100-105
 vpc 201
 spanning-tree port type network
```

```
dc11-5020-2# sh run int po 201
version 4.1(3)N1(1)

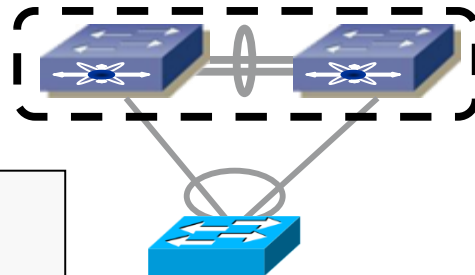
interface port-channel201
 switchport mode trunk
 switchport trunk native vlan 100
 switchport trunk allowed vlan 100-104
 vpc 201
 spanning-tree port type network
```

```
dc11-5020-1# show vpc brief vpc 201
```

vPC status

id	Port	Status	Consistency	Reason	Active vlans
201	Po201	up	success	success	100-104

2009 May 17 21:56:28 dc11-5020-1 %ETHPORT-5-IF_ERROR_VLANS_SUSPENDED: VLANs 105 on Interface port-channel201 are being suspended. (Reason: Vlan is not configured on remote vPC interface)



**“Inconsistent
config will be
disabled”**

Virtual Port Channel (vPC)

vPC Member Port

- **Definition:**

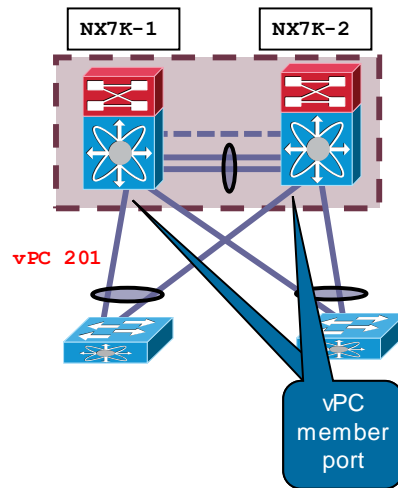
- Port-channel member of a vPC

- **Requirements**

- Configuration needs to match vPC peer
- In case of an inconsistency, a VLAN or the entire port-channel may be suspended (e.g. MTU mismatch)
- Up to 16 active ports between both vPC peers with M series LC
- Up to 32 active ports between both vPC peers with F series LC

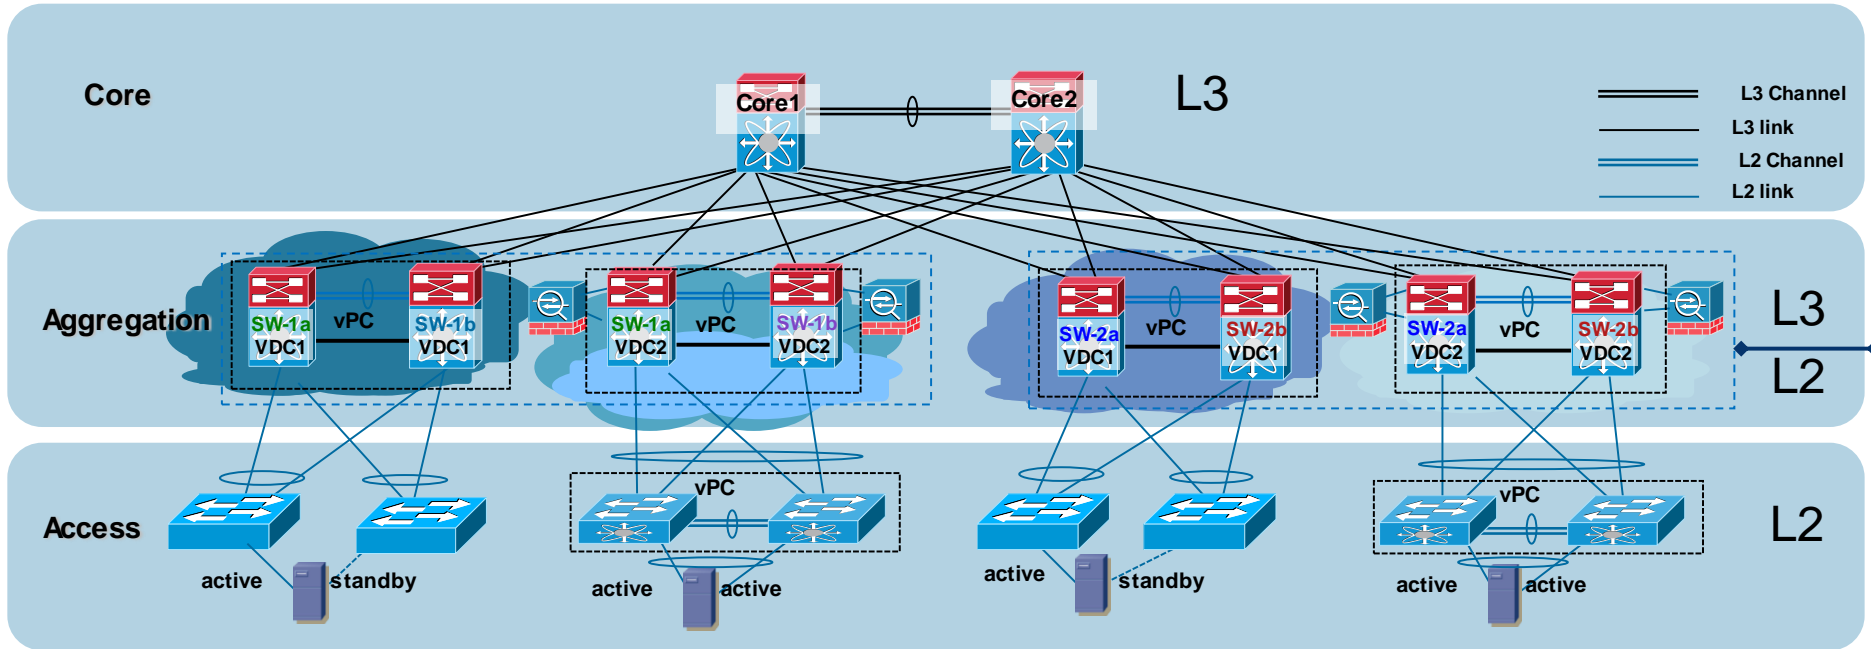
```
NX7K-1 :  
interface port-channel201  
  switchport mode trunk  
  switchport trunk native vlan 100  
  switchport trunk allowed vlan 100-105  
vpc 201
```

```
NX7K-2 :  
interface port-channel201  
  switchport mode trunk  
  switchport trunk native vlan 100  
  switchport trunk allowed vlan 100-105  
vpc 201
```



Virtual Port Channel (vPC)

VDC Interaction

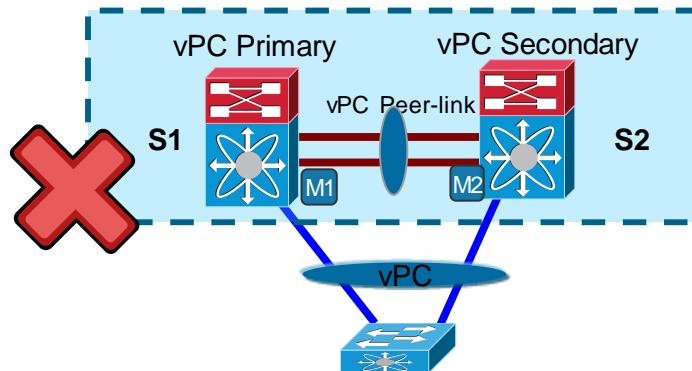
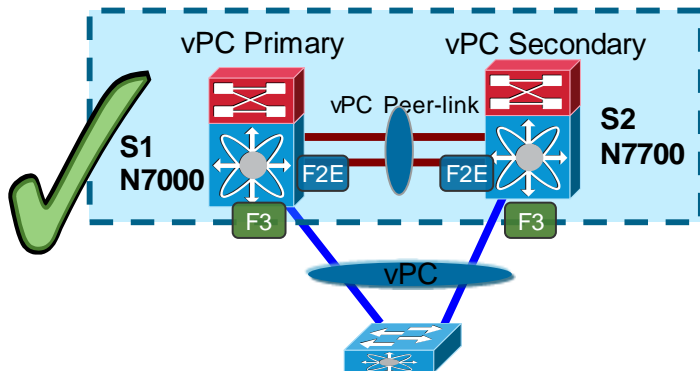


Mixed Chassis Mode

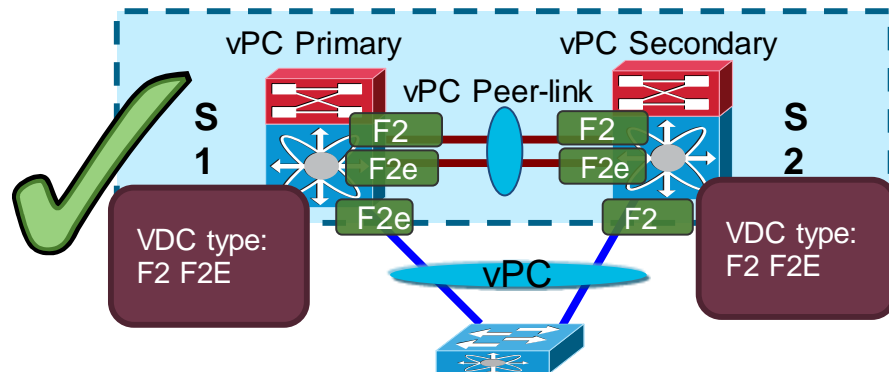
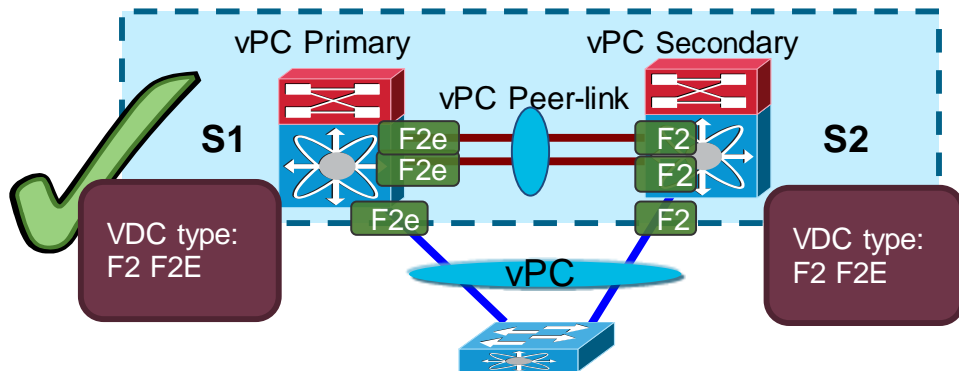
Rule of Thumb!

Always use identical line cards on either sides of the peer link and vPC legs!

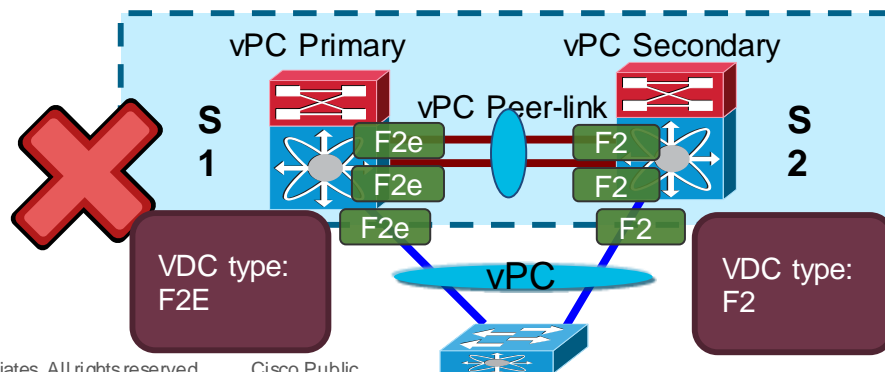
Examples



F2/F2E VDC – NX-OS 6.2(2) and Onwards



Always use identical
VDC type on both vPC
peer devices

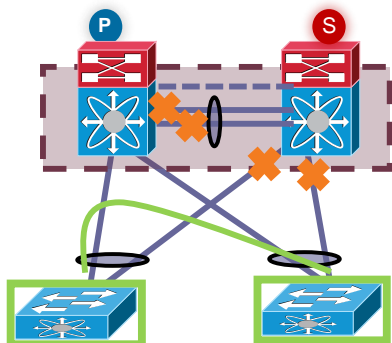


Attaching to a vPC Domain

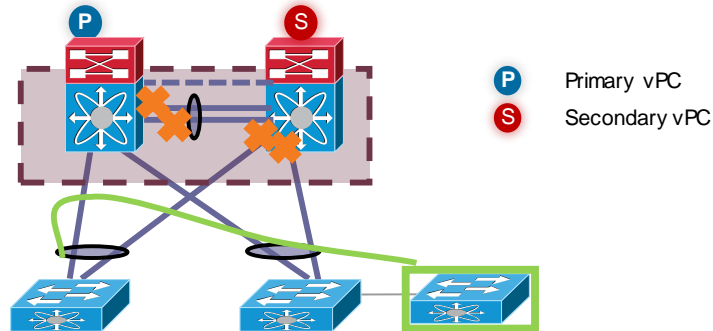
The Most Important Rule...

**Always Dual Attach Devices
to a vPC Domain!!!**

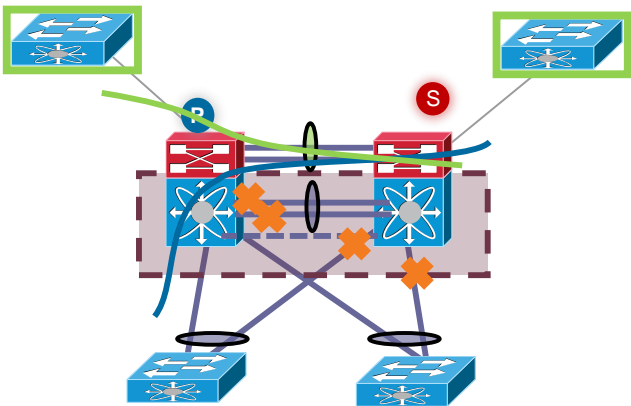
Attaching to a vPC Domain



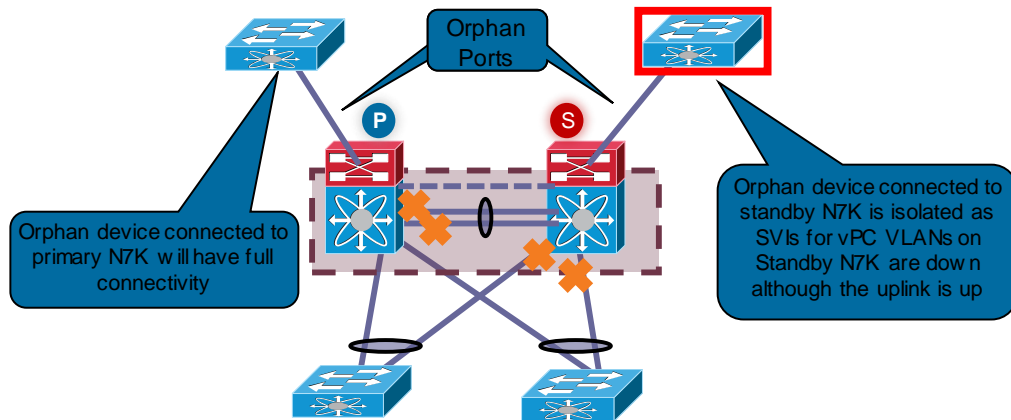
1. Dual Attached



2. Attached via VDC / Secondary Switch



3. Secondary inter switch Port-Channel (non-vPC VLAN)

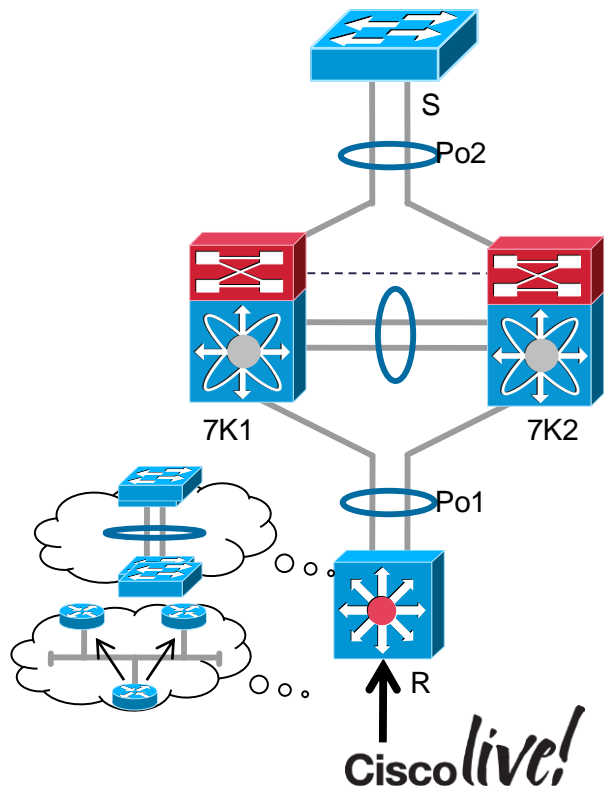


4. Single Attached to vPC Device

Layer 3 and vPC Interactions

Router Interconnection: Forwarding Sequence

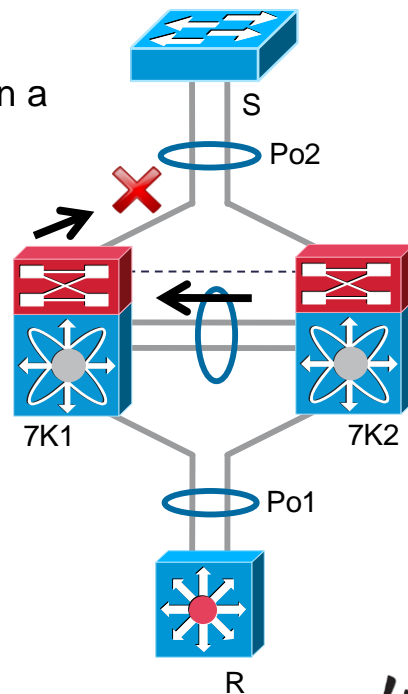
- 1) Packet arrives at R with a destination address of S
- 2) R does lookup in routing table and sees 2 equal paths going north (to 7K1 & 7K2)
- 3) Assume it chooses 7K1 (ECMP decision)
- 4) R now has rewrite information to which router it needs to go (router MAC 7K1 or 7K2)
- 5) L2 lookup happens and outgoing interface is Port-channel 1
- 6) Hashing determines which port-channel member is chosen (say to 7K2)
- 7) Packet is sent to 7K2
- 8) 7K2 sees that it needs to send it over the peer-link to 7K1 based on MAC address



Layer 3 and vPC Interactions

Router Interconnection: Forwarding Sequence (continued)

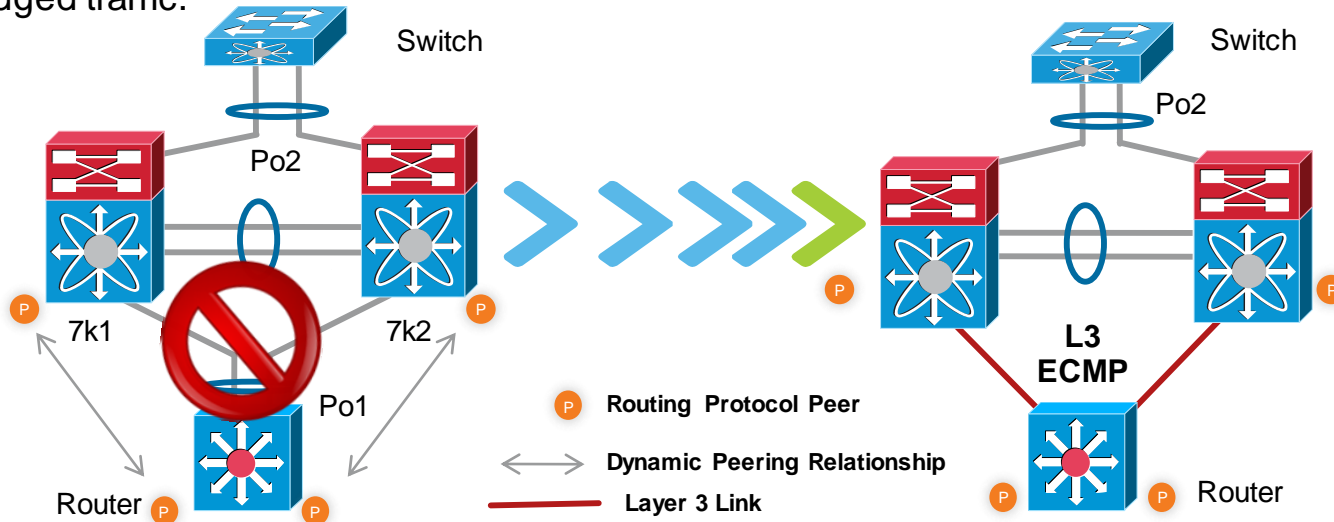
9. 7K1 performs lookup and sees that it needs to send to S
10. 7K1 performs check for if the frame came over peer link & is going out on a vPC
 - Frame will ONLY be forwarded if:
 - Outgoing interface is NOT a vPC, or
 - Outgoing vPC doesn't have active interface on other vPC peer (in our example, 7K2)



N7K Layer 3 and vPC Designs

Layer 3 and vPC Design Recommendation

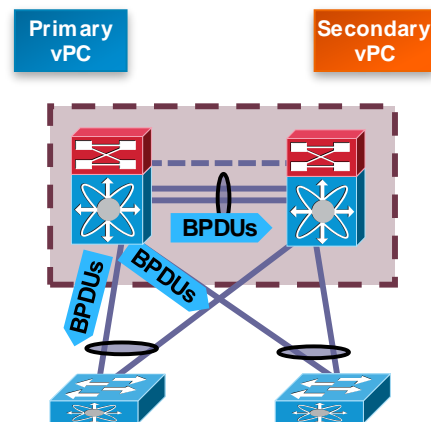
- Use L3 links to hook up routers and peer with a vPC domain.
- Don't use L2 port-channel to attach routers to a vPC domain.
- If both, routed and bridged traffic is required, use individual L3 links for routed traffic and L2 port-channel for bridged traffic.



Spanning Tree with vPC

vPC and STP BPDUs

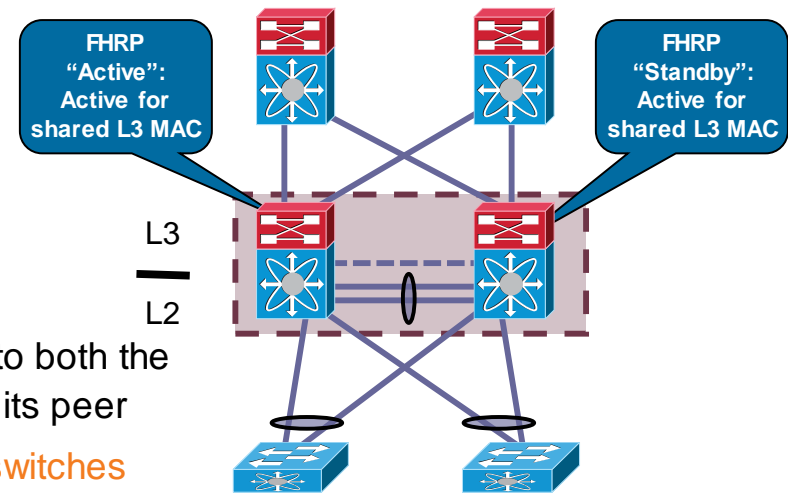
- STP for vPCs is controlled by the vPC operational primary switch and it is this device which sends BPDUs on STP designated ports.
- This occurs irrespective of where the designated STP Root is located.
- The vPC operational secondary device proxies STP BPDU messages from access switches toward the primary vPC peer.



FHRP with vPC

HSRP / VRRP / GLBP Active/Active

- Support for all FHRP protocols in Active/Active mode with vPC
- No additional configuration required
- Standby device communicates with vPC manager to determine if vPC peer is “Active” FHRP peer
- When running active/active, aggressive timers can be relaxed (i.e. 2-router vPC case)
- **‘peer-gateway’** command allows a vPC peer to respond to both the FHRP virtual and the real MAC address of both itself and its peer
- Recommendation is to use default FHRP timers as both switches are active

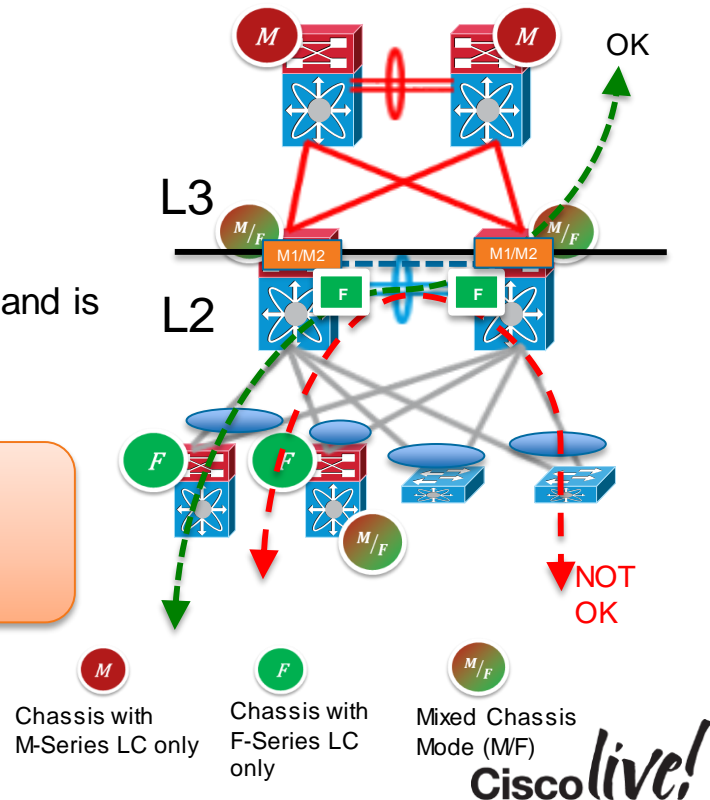


Proxy Routing Design Considerations

Dual Proxy Line Card for Redundancy

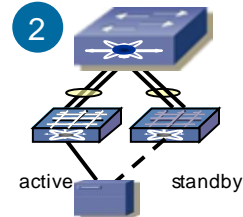
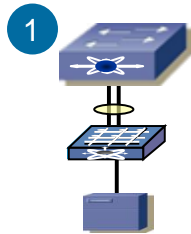
- If M1/M2 LC fails on one of the N7Ks:
 1. Inter-VLAN traffic (vPC -> FHRP -> vPC) : is dropped due to vPC loop avoidance rule
 2. Upstream traffic (vPC -> FHRP -> L3) : is bridged on vPC peer-link to other N7K FHRP vMAC and is then routed to L3 point

Recommendation is to use at least 2 M1/M2 LCs in mixed mode chassis (M/F) in order to provide redundancy for Proxy L3 Routing.

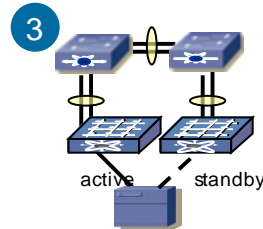


vPC Supported Topologies

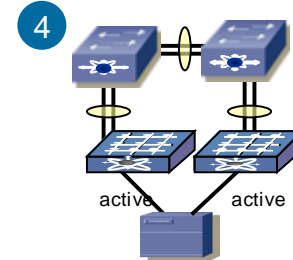
Nexus 7000 and 5000 / 5500



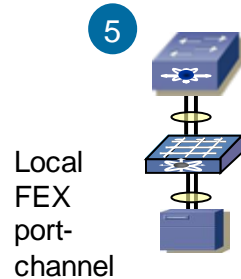
**server: active/standby
NIC teaming**



**server: active/standby
NIC teaming**

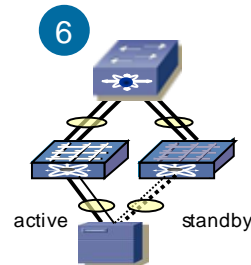


**server: active/active
no NIC teaming**

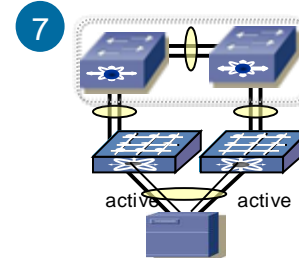


Local
FEX
port-
channel

**server: NIC teaming
(active-active)**



**server: active/standby
NIC teaming**

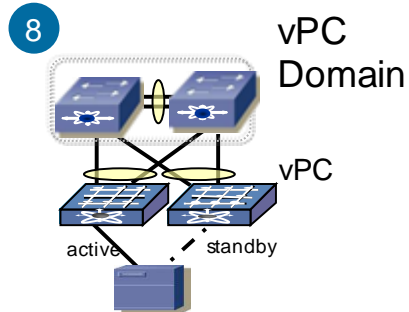


**server: NIC teaming
(active-active)**

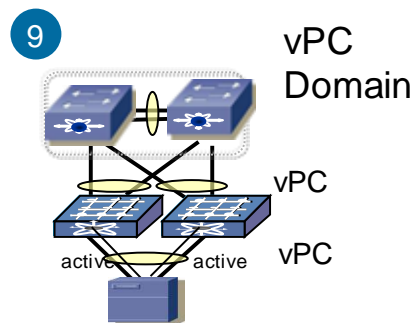
Port-Channel on HIF
(Host Interfaces supported)
- vPC to Host supported

vPC Supported Topologies

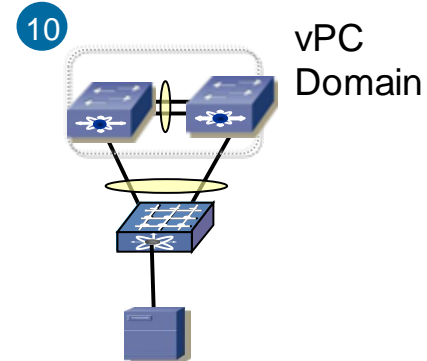
Nexus 5000 / 5500 / 6000 Only



**Dual-homed FEX
w/ A-S Server**

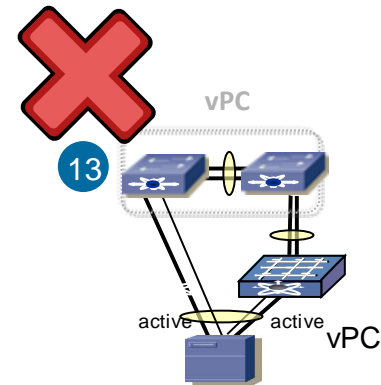
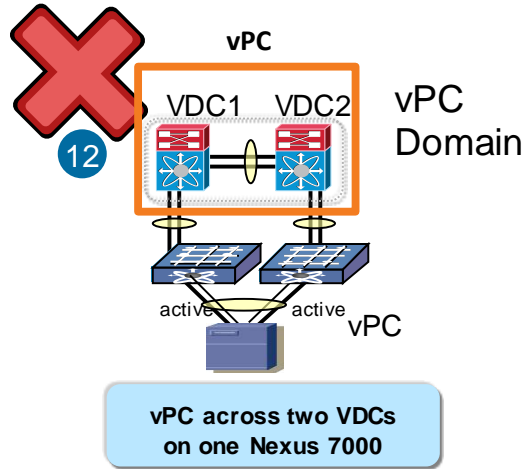
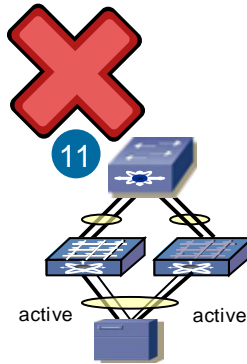


**Enhanced vPC
N5500/N6000 only**



**Dual-homed FEX
w/ Single NIC Server**

vPC Unsupported Topologies





vPC Enhancements

Agenda

- Feature Overview
- vPC Design Guidance and Best Practices
- vPC Enhancements
 - vPC Peer-Gateway
 - vPC Peer Switch
 - Graceful Consistency-Check
 - Orphan-Port Suspend
 - Redundancy with Enhanced vPC
- vPC and Fabric Technologies

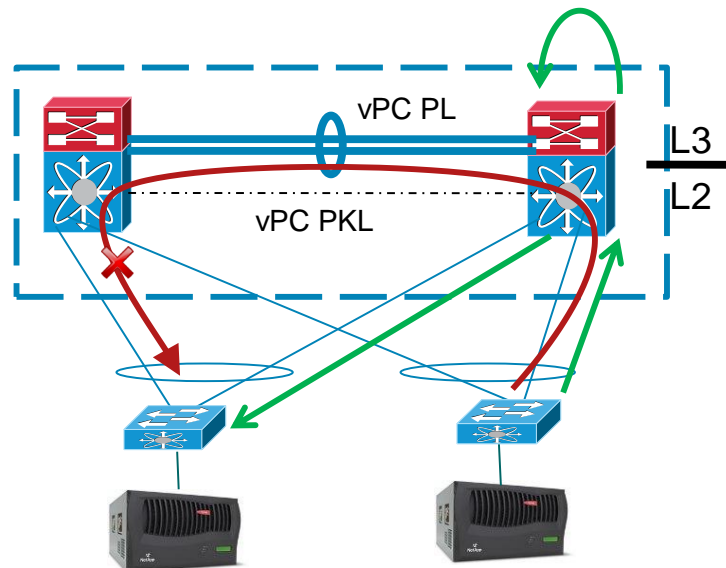


vPC Peer-Gateway

NX-OS
N7K - 4.2(1)
N5K - 5.0(3)N1(1)

Interoperability with non-RFC compliant NAS & LB

- Allows vPC device to act as active gateway for peer's router MAC.
- Avoids use of peer-link.
- No impact on traffic and existing functionality when activating this feature.



```
N7K(config-vpc-domain) # peer-gateway
```

Always enable!

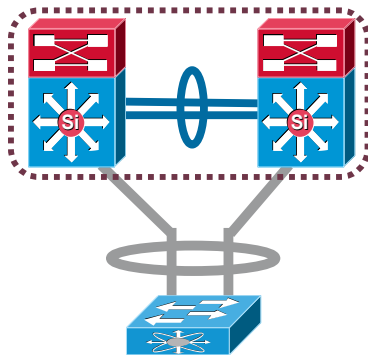
vPC Peer Switch

NX-OS
4.2(6)

Simplifies STP configuration

- Single STP root in Layer 2 topology [same bridge ID].
- Sub-second convergence for north-south flows post vPC primary failure & recovery.
- Supports single-attached devices.

```
N7K(config-vpc-domain)# peer-switch
```



vPC Peer Switch



STP Logical Topology

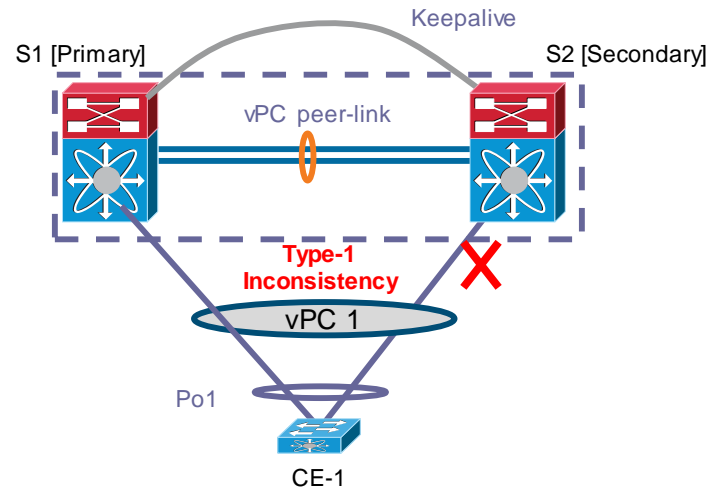
BPDUs

*Configure same STP
priority for vPC
VLANs.*

Graceful Consistency-Check

NX-OS
N7K - 5.2(1)
N5K - 5.0(2)N1(1)

- vPC member ports on S1 and S2 should have identical parameters (MTU, speed, ...)
- Any inconsistency in such parameters is Type-1 and leads to all VLANs on both vPC legs being brought down.
- With Graceful Consistency Check, only secondary vPC member ports are shut.
- Enabled by default.



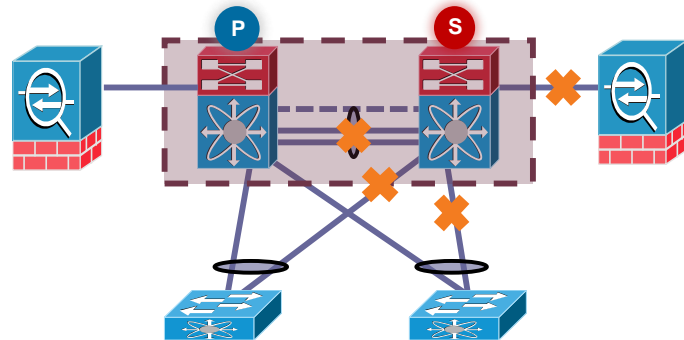
```
N7K(config-vpc-domain)# graceful consistency-check
```

Orphan-Port Suspend

vPC Active/Standby NIC Teaming Support

NX-OS
N7K - 5.2
N5K - 5.0(3)N2

- Prior to release 5.0(3)N2 on Nexus 5000/5500 and 5.2 on Nexus 7000, an orphan port was **not** shut down on loss of the vPC peer-link.
- Potential to cause black-holing of traffic at the Nexus.
- Orphan ports on vPC secondary peer can now be shut down, triggering NIC teaming recovery for all teaming configurations.
- Configuration is applied at the physical port.*



```
N7K(config)# int eth 1/1  
N7K(config-if) # vpc orphan-ports suspend
```

* prior to 6.1(2) release, 'vPC orphan-ports suspend' command may not work with FEX interface for a FEX connected to N7K, due to CSCua35190
* prior to 6.2 release, 'vPC orphan-ports suspend' command may not be applied to port-channel interface, due to CSCua37491

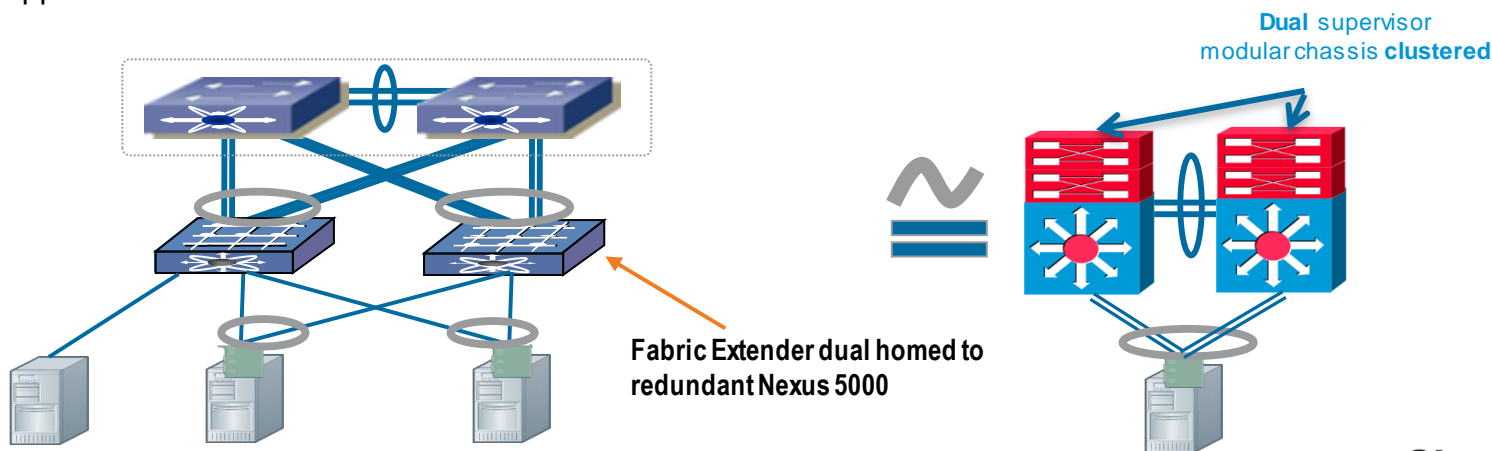
Redundancy with Enhanced vPC

Data, Control and Management Plane Redundancy

Suited for all types of servers

Port-channel connectivity to dual-homed FEXs

- From the server perspective, upstream is a single access switch with port-channel support.
- Full redundancy for supervisor or linecard via vPC and cable or NIC failure **via port-channeling**.
- Logically a similar HA model to that currently provided by dual supervisor based modular switch.
- Supports FCoE.



A long-exposure photograph of a city street at night. The foreground is filled with vibrant, multi-colored light trails from moving vehicles, creating a sense of motion and energy. In the background, modern city buildings are illuminated with various lights, and a pedestrian bridge spans across the street. The overall scene is a dynamic urban environment.

vPC and Fabric Technologies

Agenda

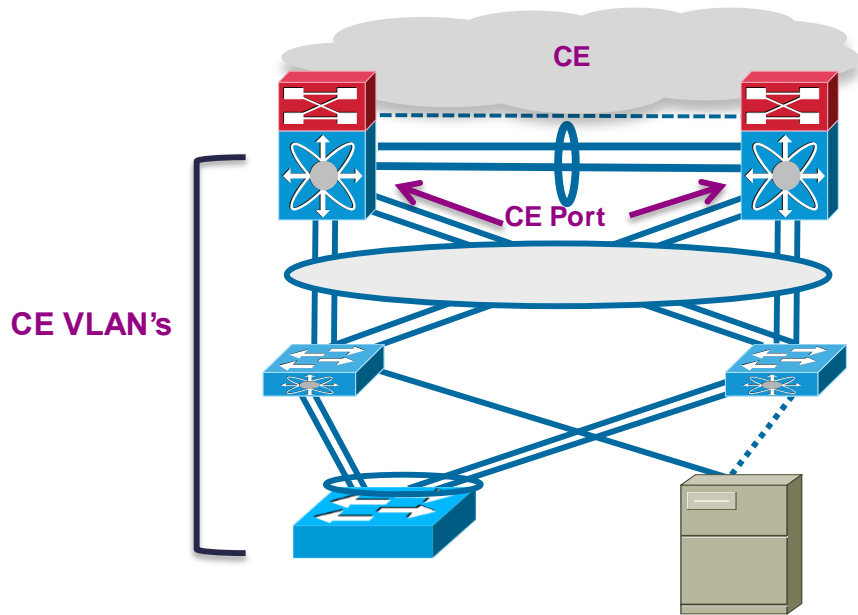
- Feature Overview
- vPC Design Guidance and Best Practices
- vPC Enhancements
- vPC and Fabric Technologies
 - vPC vs. vPC+
 - vPC with FCoE
 - vPC and VXLAN
 - vPC and ACI



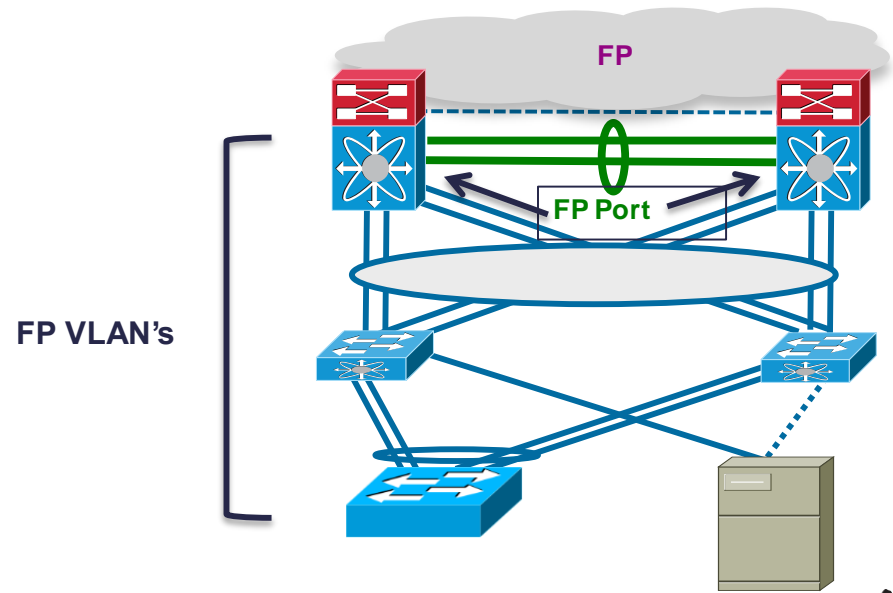
vPC vs. vPC+

Architecture of vPC and FabricPath with vPC+

vPC

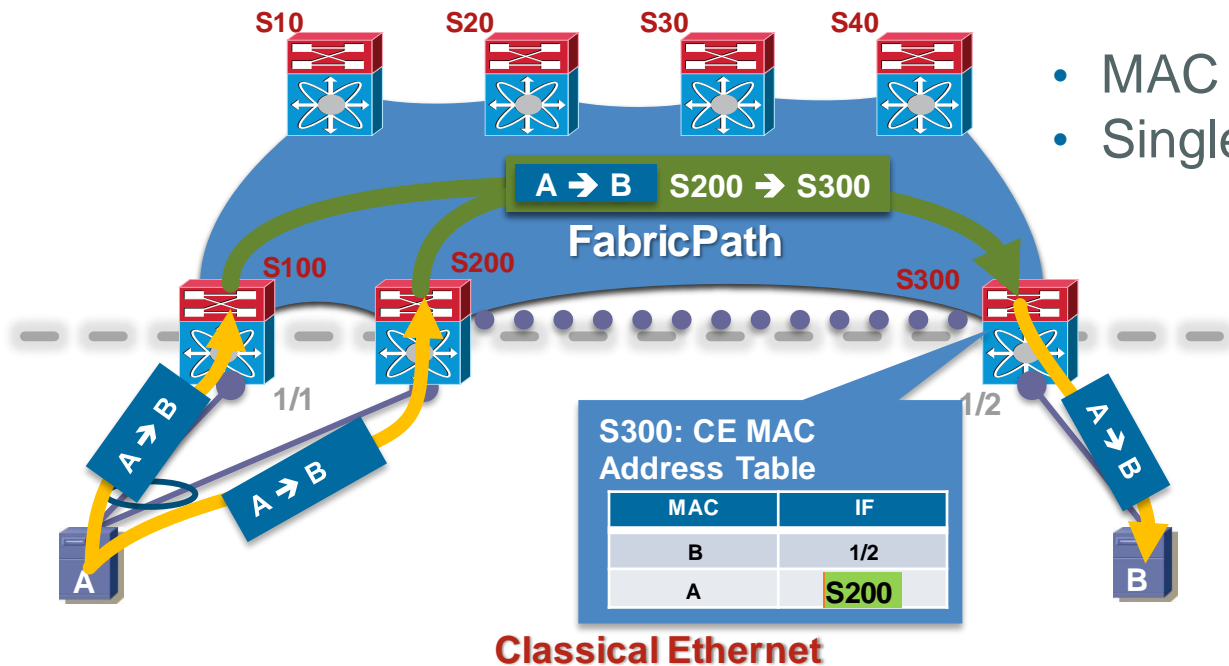


vPC+



Technical Challenges

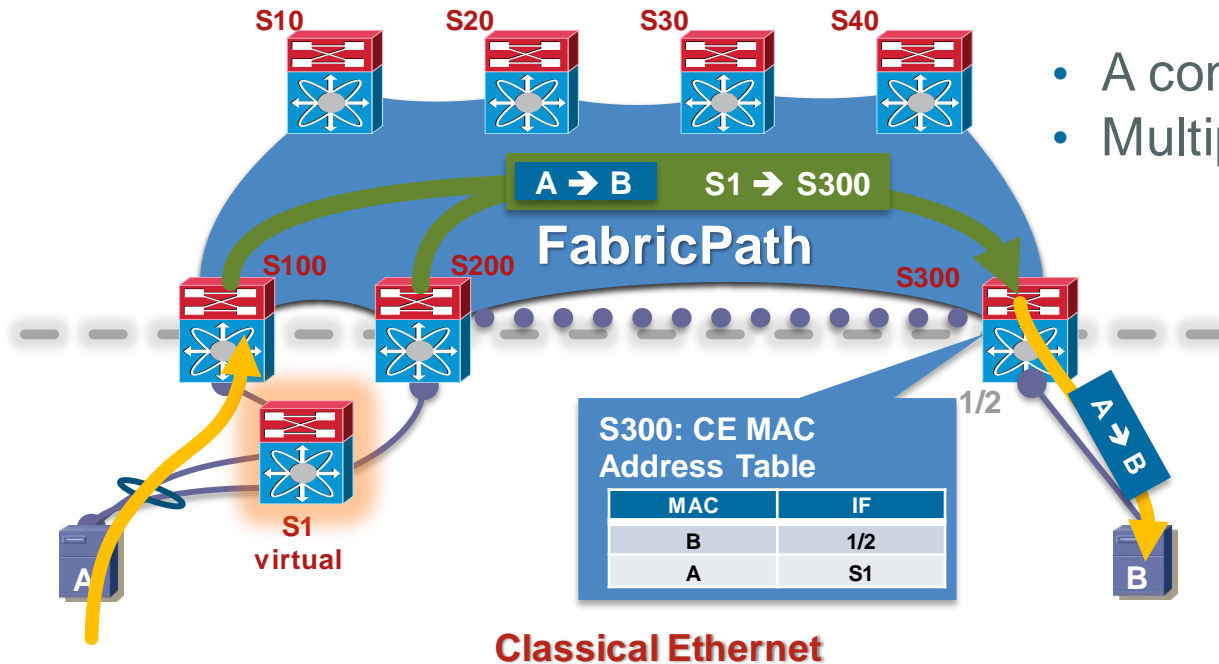
vPC vs. vPC+



- MAC address flapping on S300
- Single path to A

vPC+ Virtual Switch

vPC vs vPC+

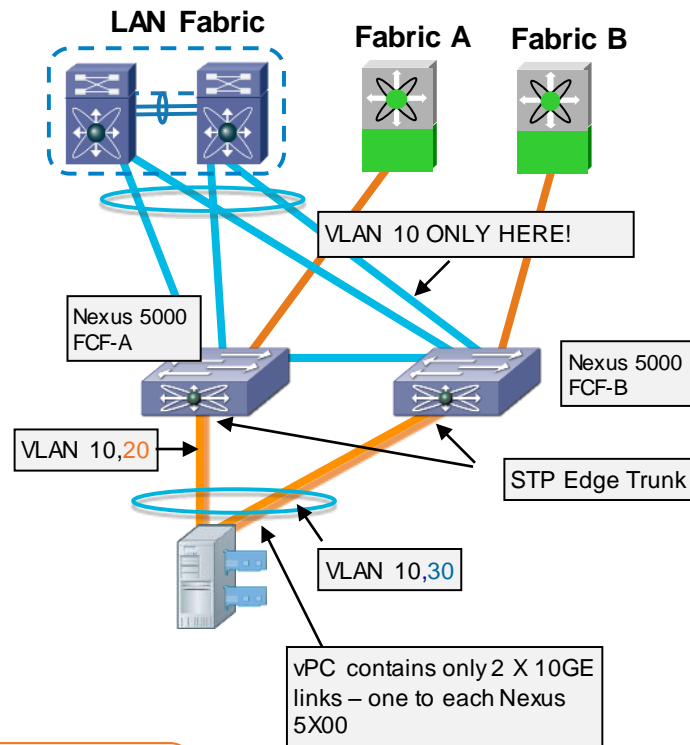


- A consistently associated to S1
- Multipathing to A

vPC with FCoE

Unified Fabric Design

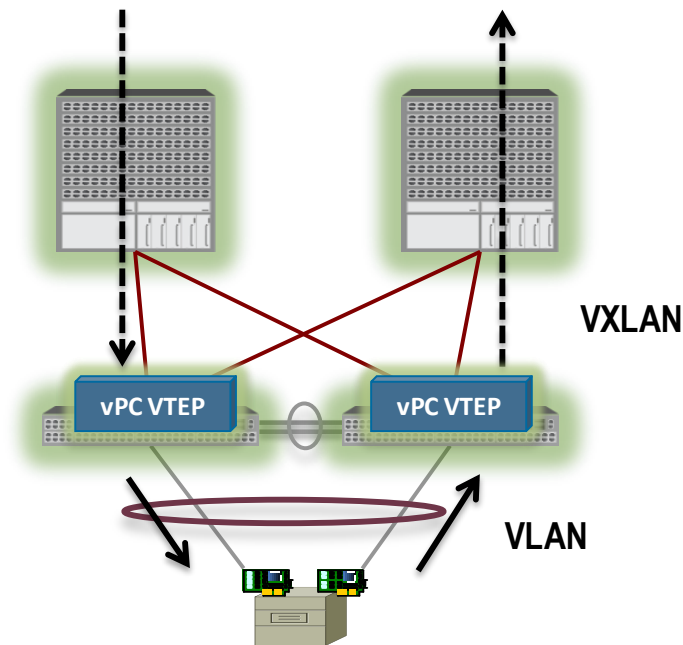
- vPC with FCoE is ONLY supported between hosts and N5K/N6K or N5K/N6K & N2232 pairs.
- Must follow specific rules:
 - A 'vfc' interface can only be associated with a single-port port-channel.
 - While the port-channel configurations are the same on both switches, the FCoE VLANs are different.
- FCoE VLANs are 'not' carried on the vPC peer-link (automatically pruned):
 - FCoE and FIP ethertypes are 'not' forwarded over the vPC peer link.
- vPC carrying FCoE between two FCF's is NOT supported.
- Best Practice:** Use static port channel rather than LACP with vPC and boot from SAN.
[If NX-OS is prior to 5.1(3)N1(1)]



vPC and VXLAN

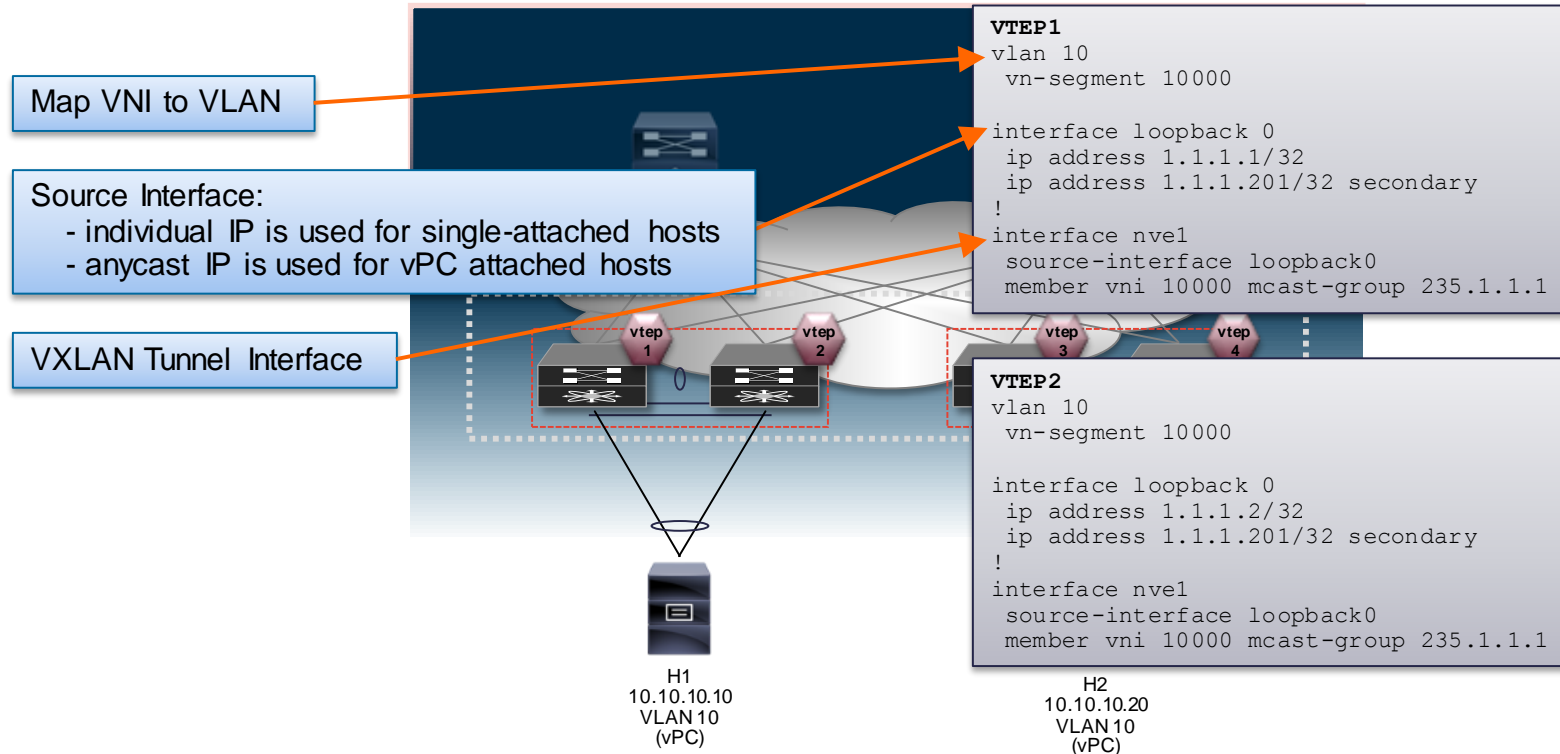
vPC VTEP

- When vPC is enabled an anycast VTEP address is programmed on both vPC peers.
- Symmetrical forwarding behaviour on both peers
- Multicast topology prevents broadcast, unknown unicast and multicast traffic being sent to the same IP address across the L3 network (prevents duplication of flooded packets).
- vPC peer-gateway feature must be enabled on both peers.
- VXLAN header is not carried on the vPC Peer link.



vPC and VXLAN

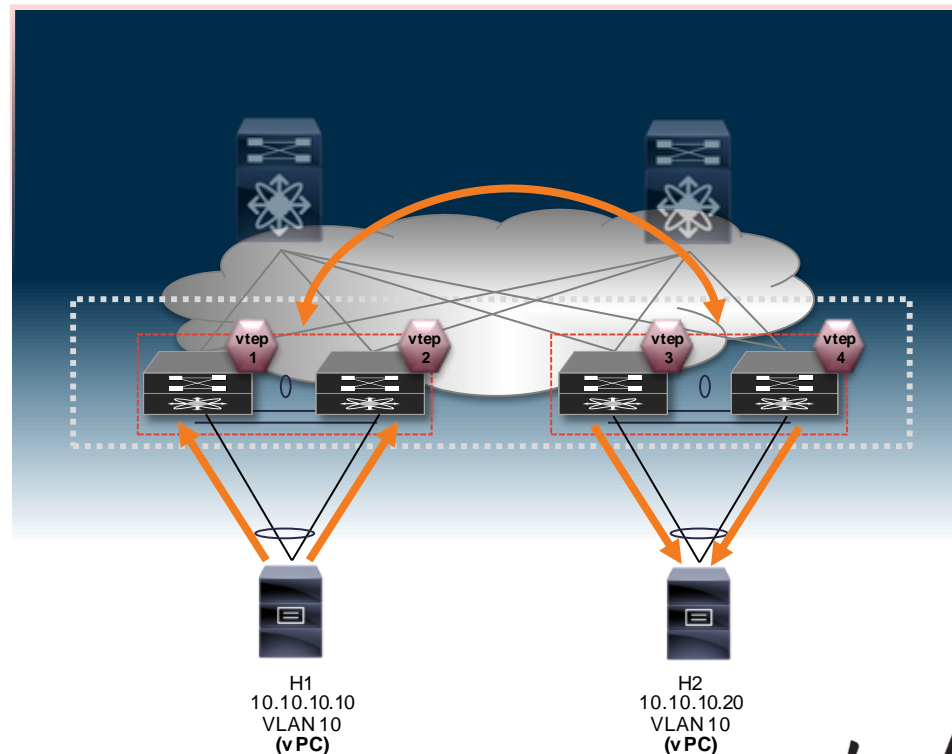
vPC Configuration



vPC and VXLAN

Dual-attached Host to dual-attached Host (Layer 2)

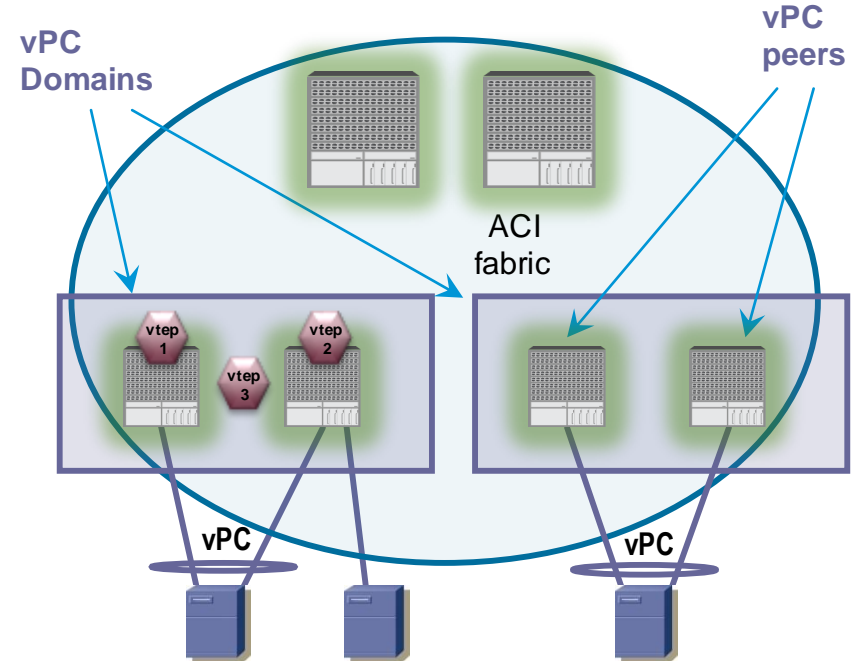
- Host 1 (H1) and Host 2 (H2) are **dual** connected to a VPC domain.
- As H1 is behind a vPC interface, the anycast VTEP IP is the source for the VXLAN encapsulation.
- As H2 is behind a vPC interface, the anycast VTEP IP is the destination.



vPC and ACI

ACI fabric utilised for control-plane

- No dedicated peer-link between vPC peers:
 - Fabric itself serves as the MCT.
- No out-of-band mechanism to detect peer liveness:
 - Due to rich fabric-connectivity (leaf-spine), it is very unlikely that peers will have no active path between them.
- CFS (Cisco Fabric Services) is replaced by Zero Message Queue (ZMQ).
- As ACI fabric is VXLAN-based, an anycast VTEP is shared by both leaf switches in a vPC domain.



Key Takeaways

NX-OS vPC Key Takeaways

- vPC is a very popular feature which makes it possible to **use all available bandwidth** while providing redundancy in L2 environments.
- Leverage vPC technology to obtain the **benefits of high availability** and to **avoid convergence** in Layer 2 Networks.
- Follow the **design guidelines and best practices** to successfully deploy vPC architecture.
- Use recent **enhancements** to optimise vPC behaviour.
- Use the recommended NX-OS release to leverage convergence, scalability & stability optimisations. Cisco N7K NX-OS recommended release page:
http://www.cisco.com/en/US/docs/switches/datacenter/sw/nx-os/recommended_releases/recommended_nx-os_releases.html



Q & A

Cisco *live!*

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2015 T-Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site
<http://showcase.genie-connect.com/clmelbourne2015>
- Visit any Cisco Live Internet Station located throughout the venue

T-Shirts can be collected in the World of Solutions on Friday 20 March 12:00pm - 2:00pm



Learn online with Cisco Live!

Visit us online after the conference for full access to session videos and presentations. www.CiscoLiveAPAC.com

Cisco *live!*



Thank you.

Cisco *live!*



CISCO

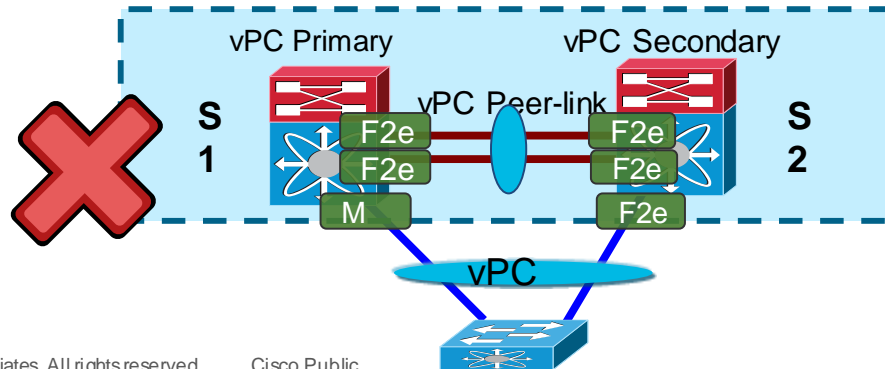
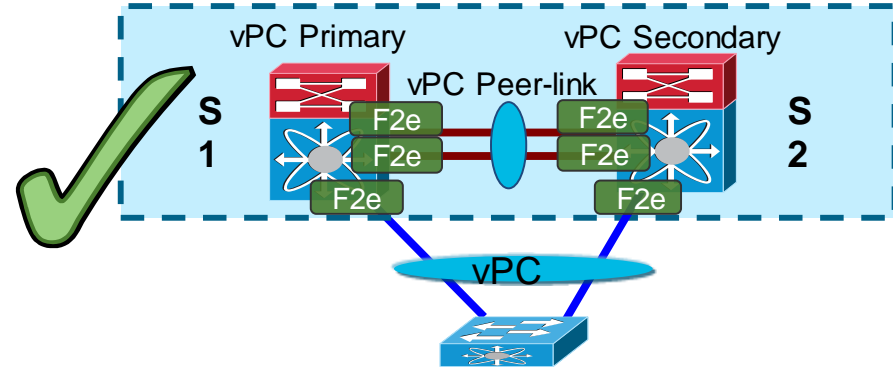
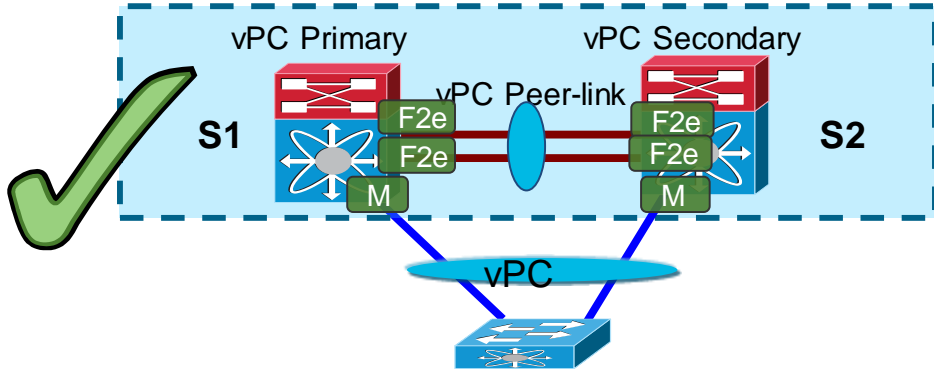
A long-exposure photograph of a city street at night. The foreground is filled with vibrant, multi-colored light trails from moving vehicles, creating a sense of motion. In the background, a pedestrian bridge spans the street, and modern city buildings with lit windows and signage are visible under a dark sky.

Appendix I

M1/M2 - F1/F2/F2e LC Design Considerations

Cisco *live!*

M/F2E – NX-OS 6.2(2) and Onwards

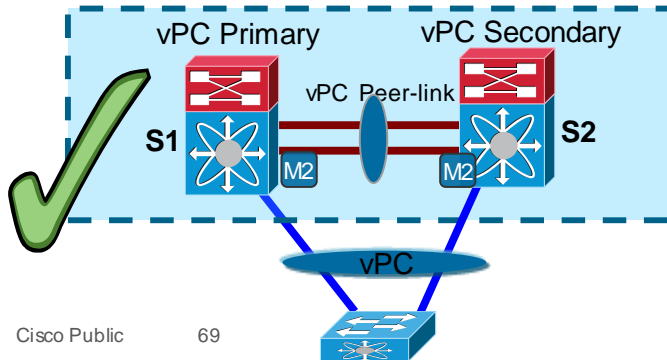
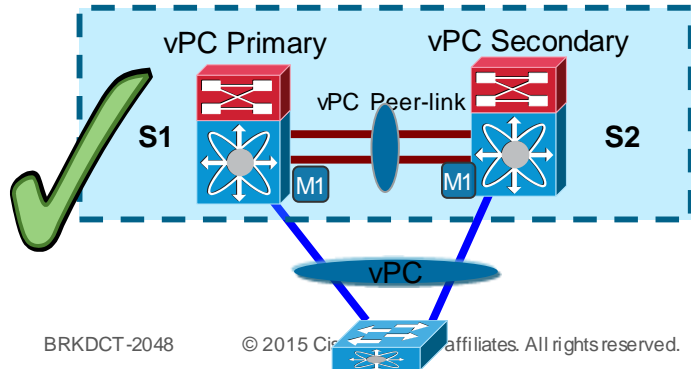
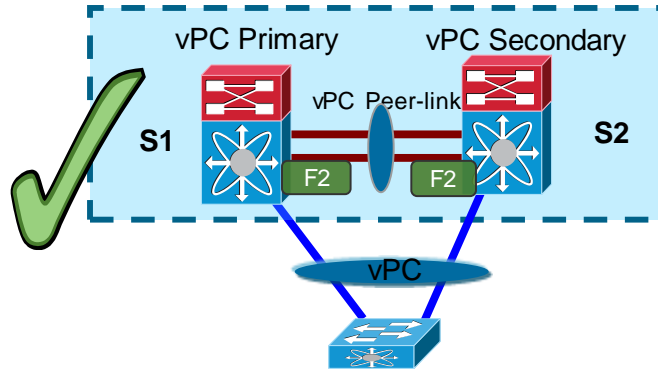
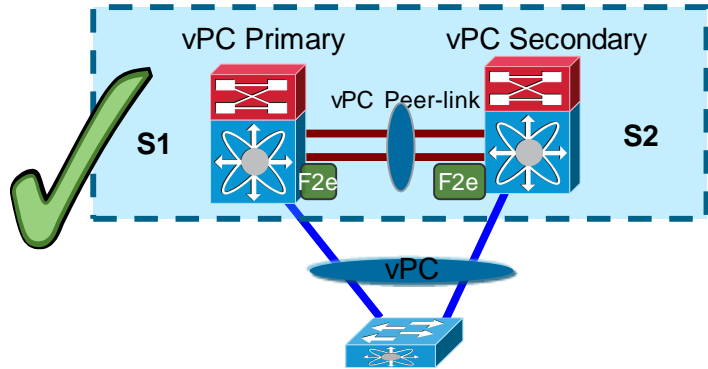


Always use identical line cards on either sides of the vPC Peer Link and vPC legs !

M1/M2 - F1/F2/F2e LC Design Considerations

vPC / Port-channel

Always use identical line cards on either sides of the peer link !

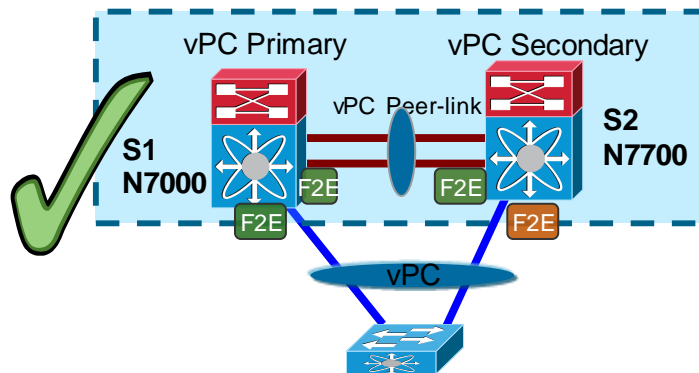
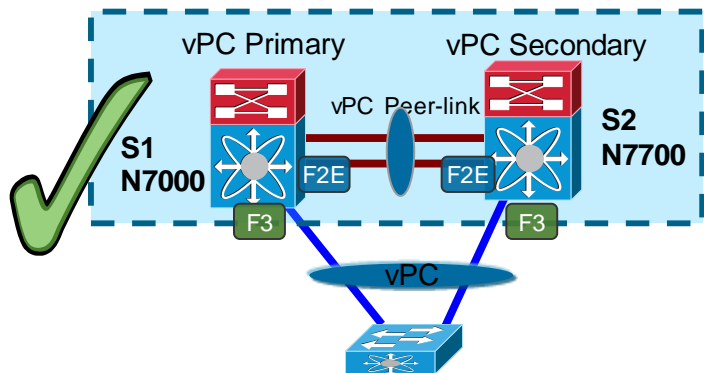


N7000 – N7700 VPC Design Considerations

vPC / Port-channel

- N7000 and N7700 in same VPC Construct
- VDC type should match on both peer device

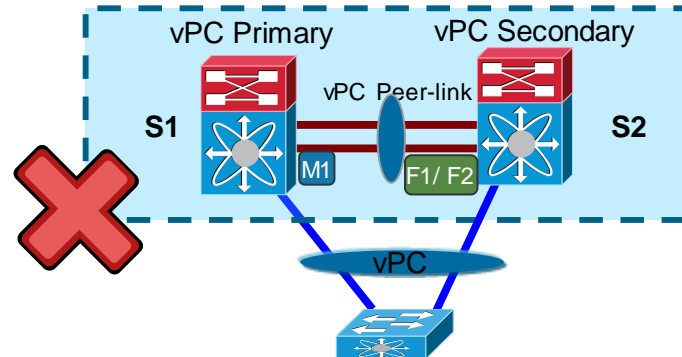
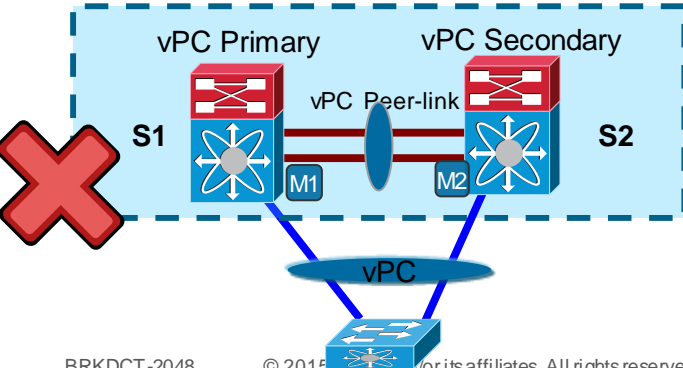
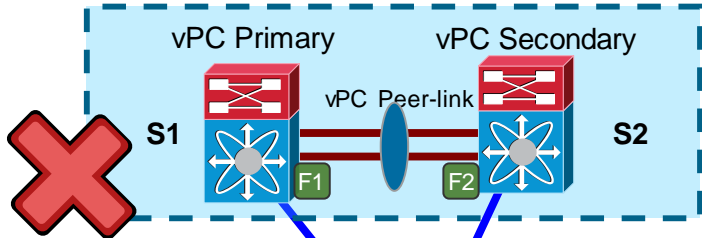
Always use identical line cards on either sides of the peer link !



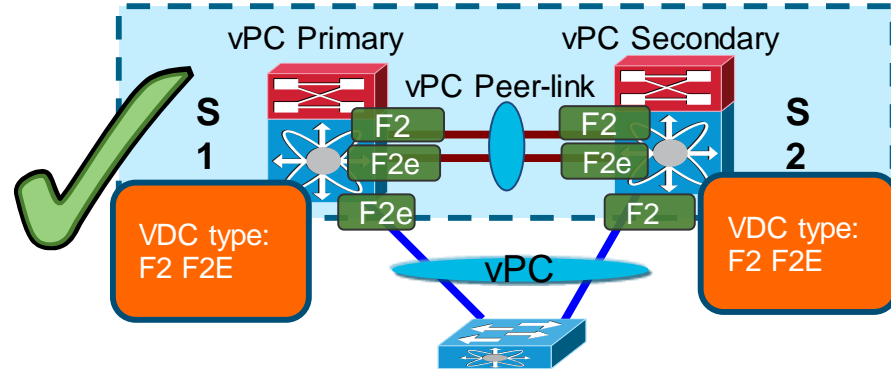
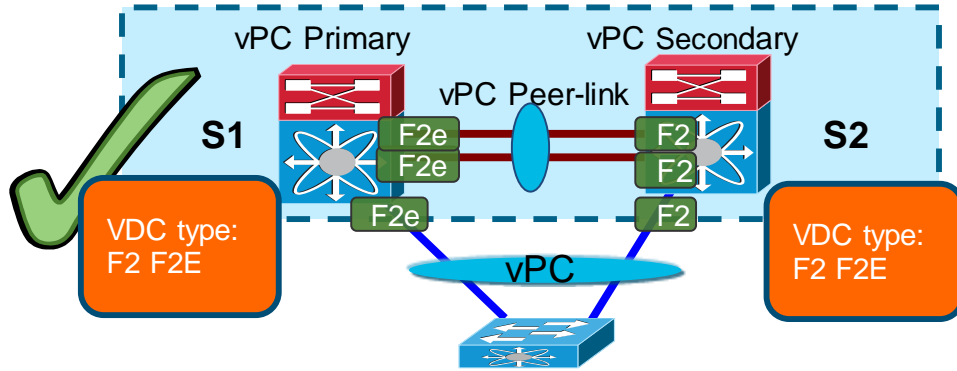
M1/M2 - F1/F2/F2e LC Design Considerations

vPC / Port-channel

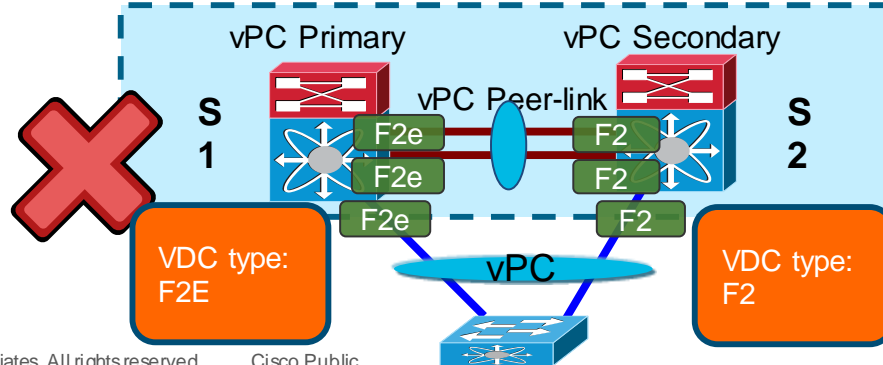
Always use identical line cards on either sides of the peer link !



F2/F2E VDC – NX-OS 6.2(2) and Onwards



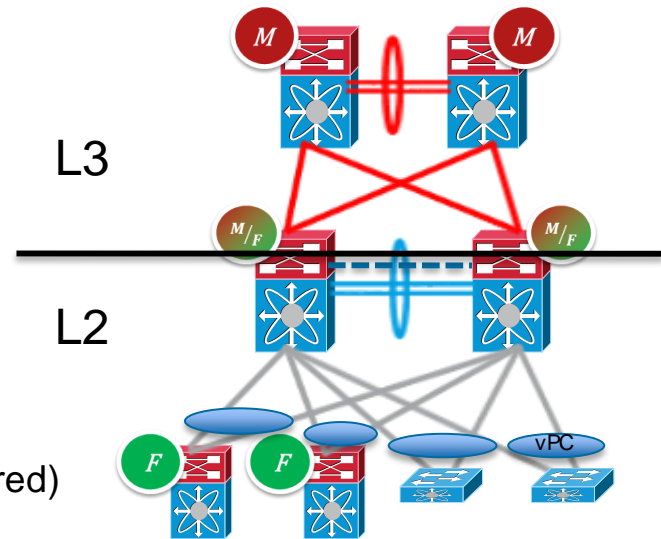
Always use identical
VDC type on both vPC
peer devices





M1/M2 – F2E LC Design Considerations


Mixed Chassis Mode

- M-Series provides scalable L2 and L3 tables
- F2E-Series provides high-density cost-effective Layer 2 10GbE connectivity
- Mixing M-Series and F2E-Series in aggregation switch (mixed chassis mode) provides several benefits :
 - Bridging traffic remain in F2E-Series LC
 - Internal proxy routing via M-Series LC for routed traffic coming from F-Series LC
- NEXUS 7000 with F2E-series LC only can be used as modular access switch (no need for M-Series LC if L3 function is not required)



 Chassis with
M-Series LC
only

 Chassis with
F2E-Series LC
only

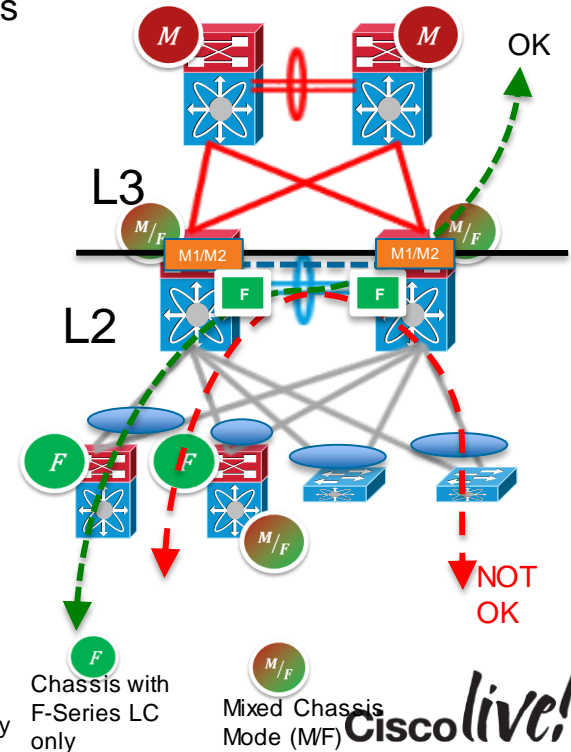
 Mixed Chassis
Mode (M/F)
Cisco live!

M1/M2 – F2E LC Design Considerations

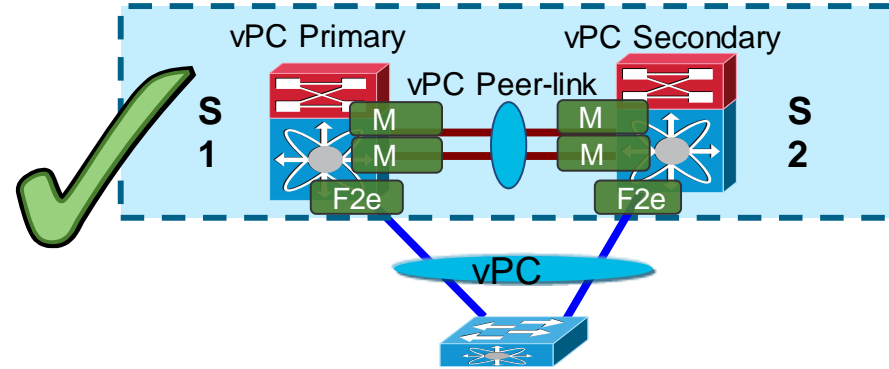
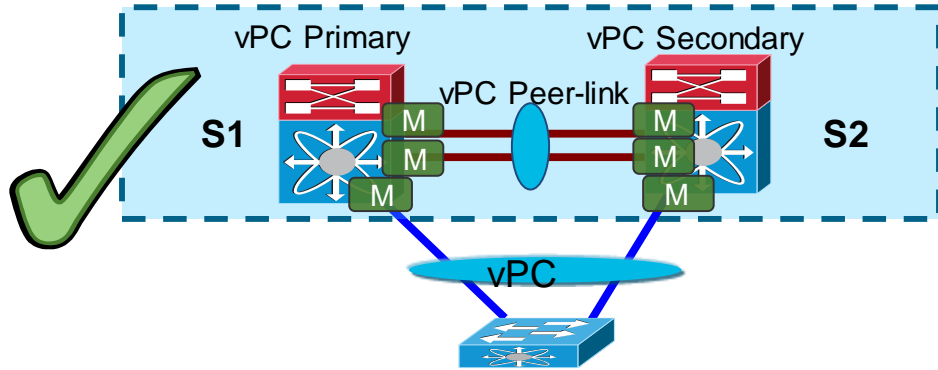
Mixed Chassis (M/F) Mode – 1 M1/M2 LC only ; Peer-link on F2E Ports

- Mixed chassis (M/F) can operate with 1 M1/M2 LC per chassis but it is not recommended:
- M1/M2 LC will provide all local routing capabilities for the chassis : inter-vlan routing and L3 northbound traffic
- When M1/M2 LC fails down on one of the N7Ks:
 - Inter-VLAN traffic (vPC -> FHRP -> vPC) : traffic gets dropped because of vPC loop avoidance rule
 - Upstream traffic (vPC -> FHRP -> L3) : traffic gets bridged on vPC peer-link to other NEXUS 7000 FHRP vMAC and then routed to L3 point

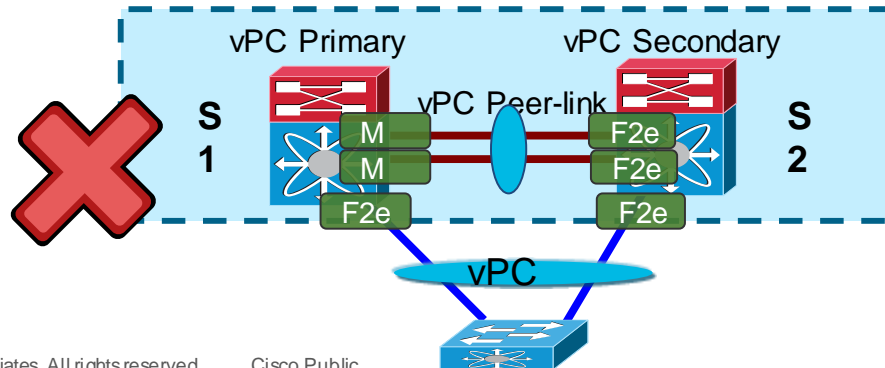
Recommendation is to use at least 2 M1/M2 LC in mixed mode chassis (M/F) in order to provide redundancy for Proxy L3 Routing.



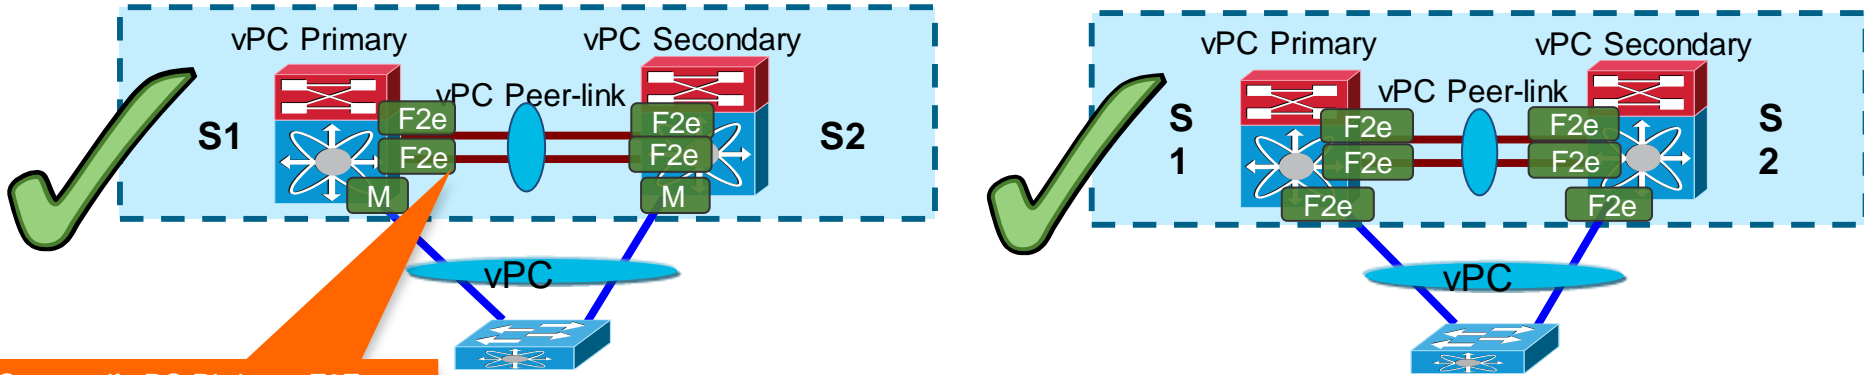
M1/M2 - F1/F2/F2e LC Design Considerations



Always use *identical* line cards on either sides of the vPC Peer Link and vPC legs !

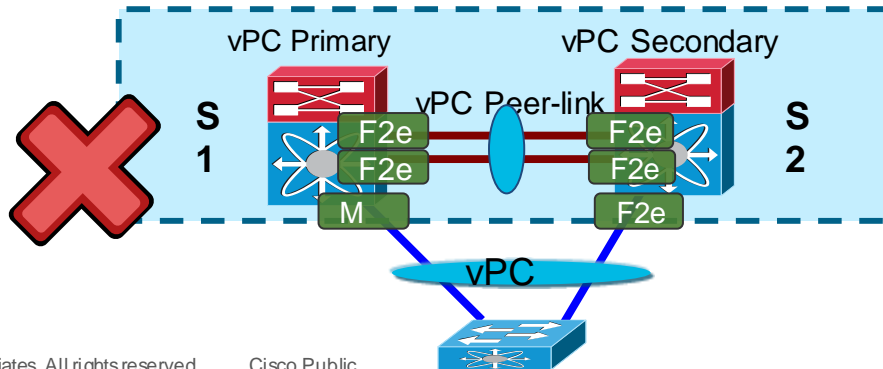


M/F2E – NX-OS 6.2(2) and Onwards

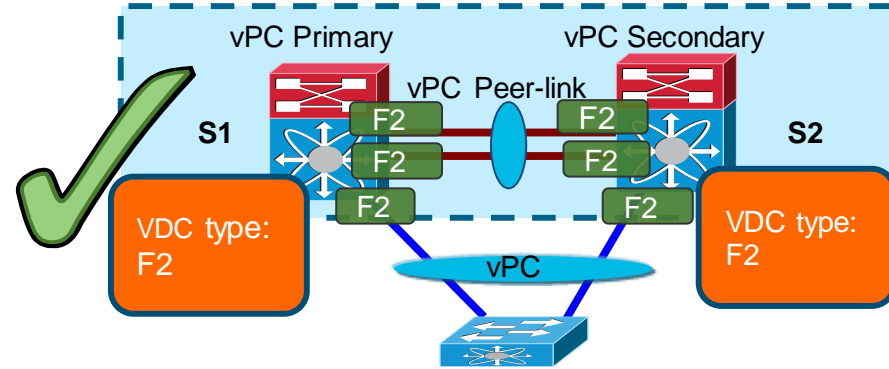
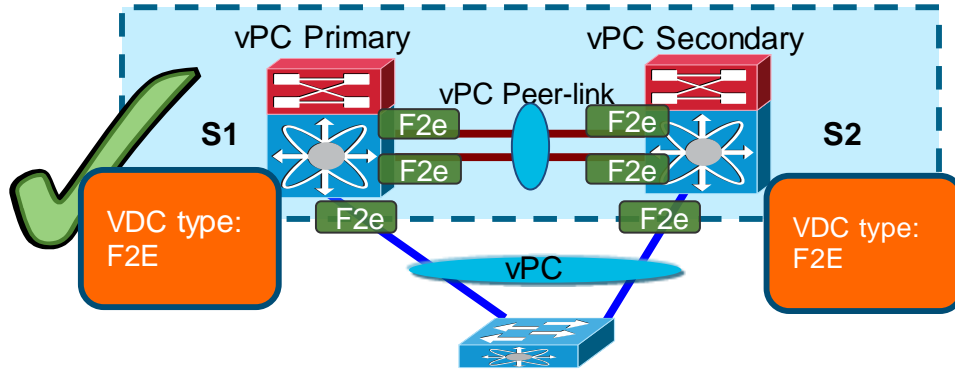


Caveat: If vPC PL is on F2E, doNOT use L3 backup routing path over vPC PL (deploy dedicated L3 backup routing path using additional inter-switch port-channel link)

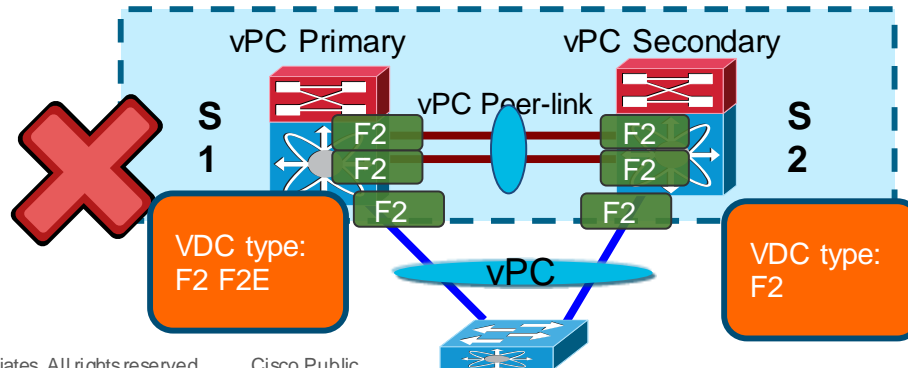
Always use identical line cards on either sides of the vPC Peer Link and vPC legs !



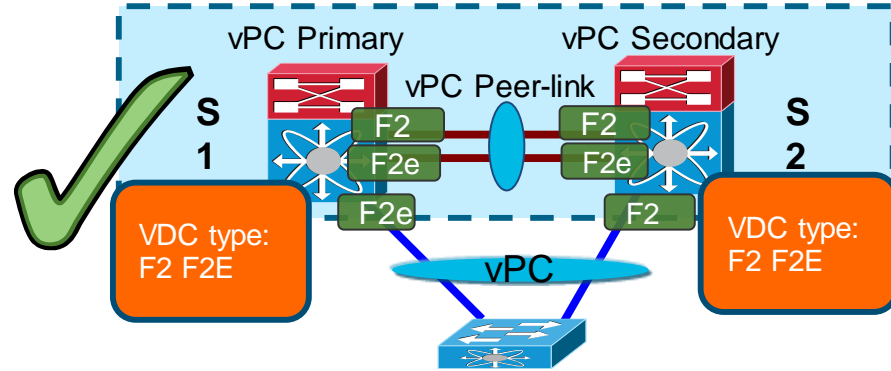
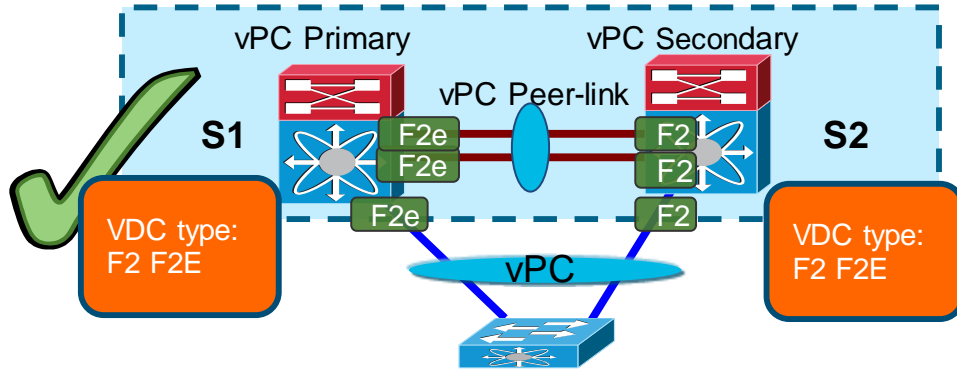
F2/F2E VDC – NX-OS 6.2(2) and Onwards



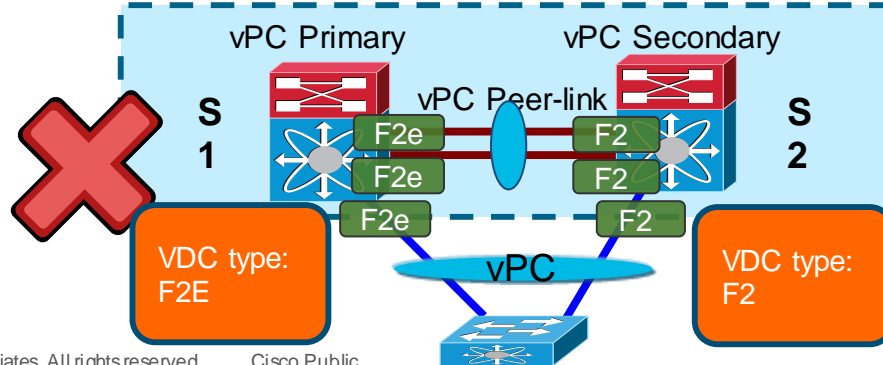
Always use identical
VDC type on both
vPC peer devices



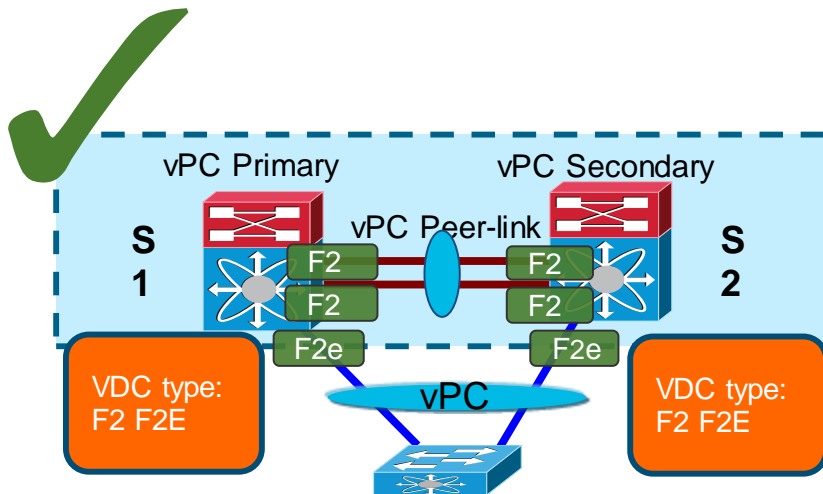
F2/F2E VDC – NX-OS 6.2(2) and Onwards



Always use identical
VDC type on both vPC
peer devices



F2/F2E – NX-OS 6.2(2) and Onwards

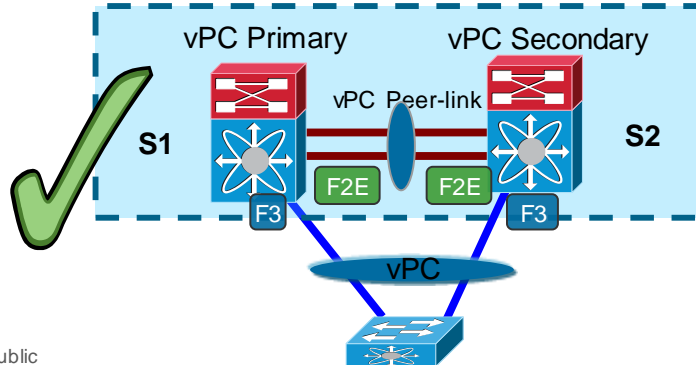
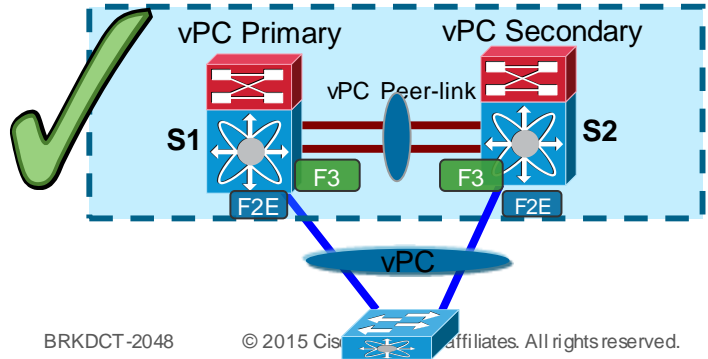
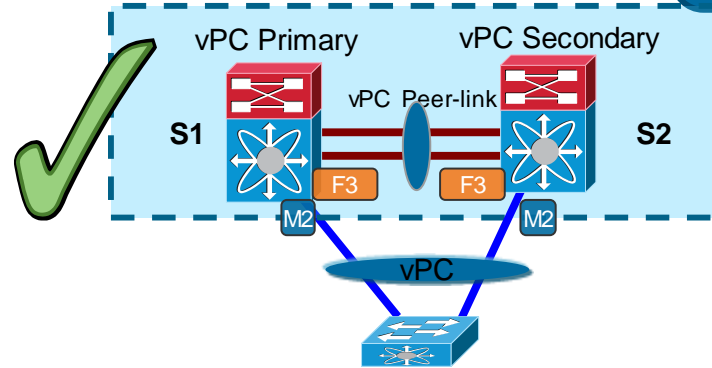
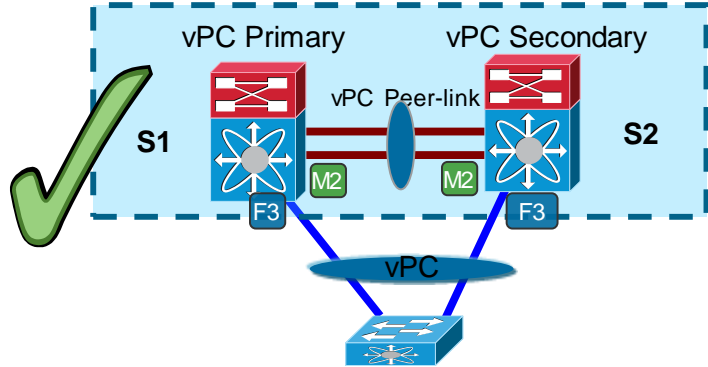


*Always use identical
VDC type on both
vPC peer devices*

F3 – M2 / F2E LC Design Considerations

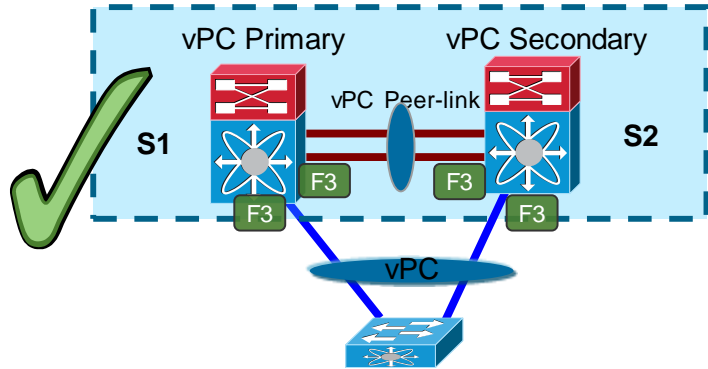
vPC / Port-channel

Always use identical line cards on either sides of the peer link!



F3 – M2 / F2E LC Design Considerations

vPC / Port-channel



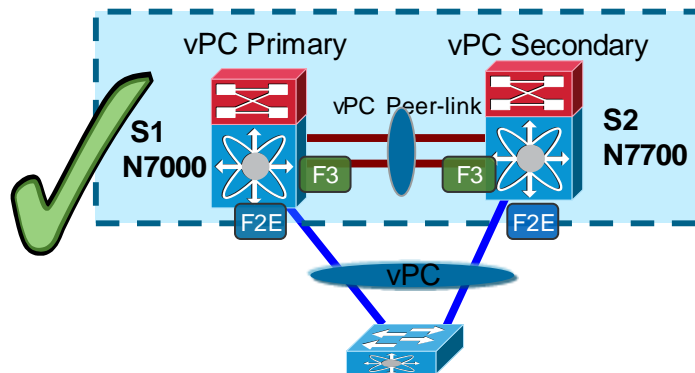
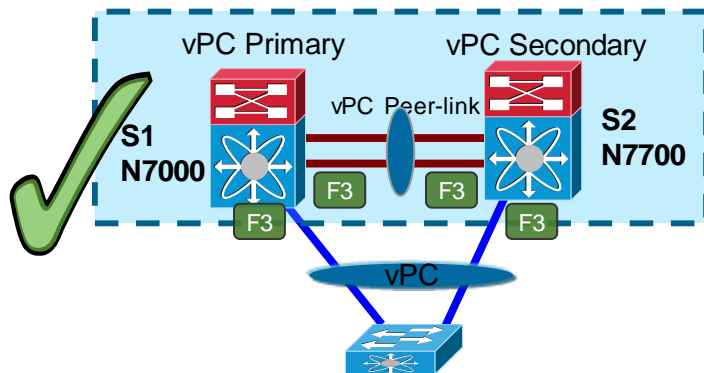
Always use identical line cards on either sides of the peer link !

N7000 – N7700 VPC Design Considerations

vPC / Port-channel

- N7000 and N7700 in same VPC Construct
- VDC type should match on both peer device

Always use identical line cards on either sides of the peer link !

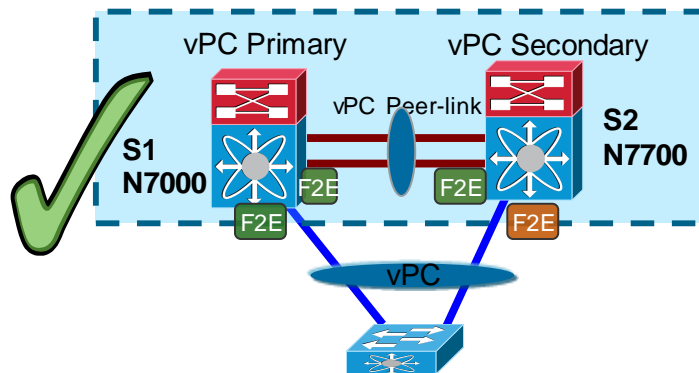
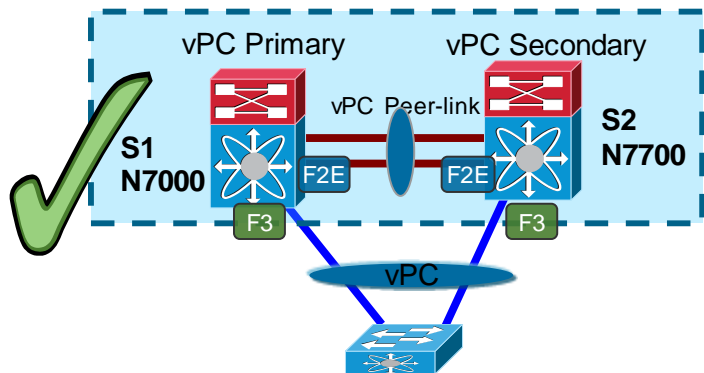


N7000 – N7700 VPC Design Considerations

vPC / Port-channel

- N7000 and N7700 in same VPC Construct
- VDC type should match on both peer device

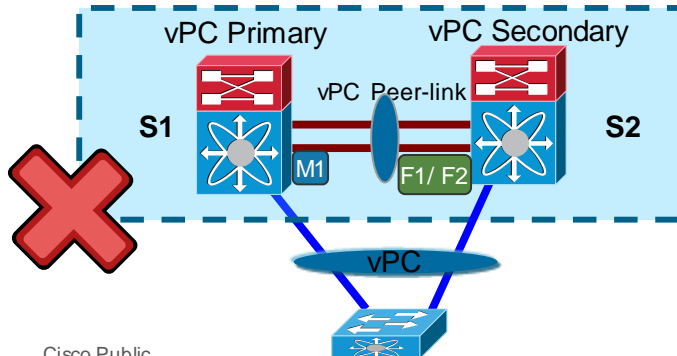
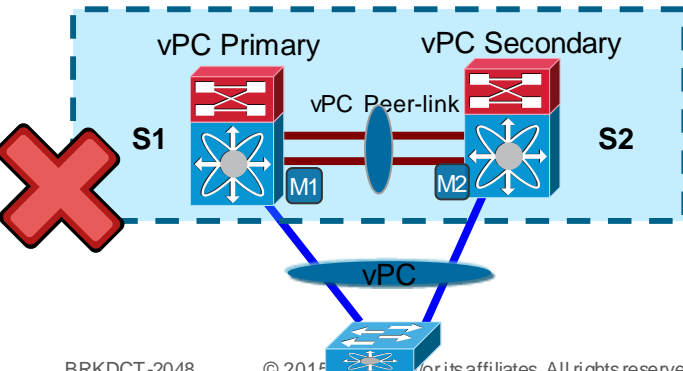
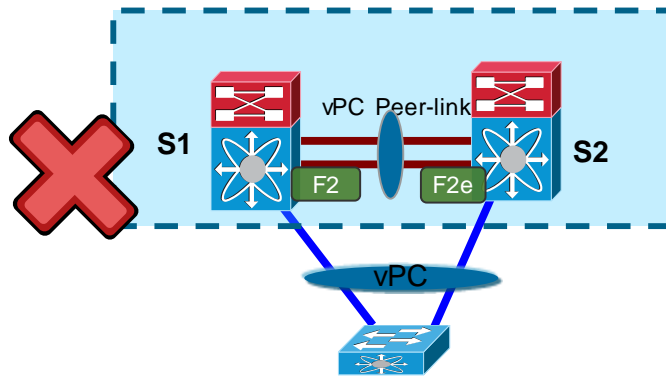
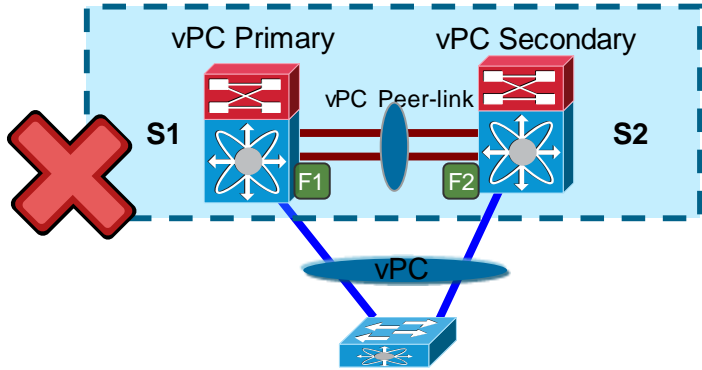
Always use identical line cards on either sides of the peer link !



M1/M2 - F1/F2/F2e LC Design Considerations

vPC / Port-channel

Always use identical line cards on either sides of the peer link !





Appendix II







vPC Hardware Support

Cisco *live!*

vPC - Supported Hardware







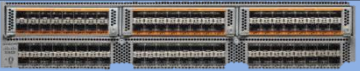


Nexus 7000

I/O Module		vPC Peer-link (10 GE Only)	VPC Interfaces
N7K-M132XP-12 N7K-M132XP-12L		✓	✓
N7K-M148GT-11L N7K-M148GS-11L		✗	✓
N7K-M108X2-12L		✓	✓
N7K-F132XP-15		✓	✓
N7K-F248XP-25 N7K-F248XP-25E N7K-F248XT-25E N77-F248XP-23E		✓	✓
N7K-M224XP-23L N7K-M206FQ-23L N7K-M202CF-22L		✓	✓

vPC - Supported Hardware








Nexus 5000/5500




Part Number / Chassis		vPC Peer-link (10 GE Only)	VPC Member Port
N5K-C5010P-BF		✓	✓
N5K-C5020P-BF		✓	✓
N5K-C5548P-FA		✓	✓
N5K-C5548UP-FA		✓	✓
N5K-C5596UP-FA		✓	✓
N5K-C5596T-FA		✓	✓
Nexus 6000		✓	✓

vPC - Supported Hardware

Nexus 2000 Platform



FEX	vPC Peerlink	VPC Member Port	
		NEXUS 5000/5500 parent switch	NEXUS 7000 parent switch
N2K-C2148T-1GE 	X	✓	X
N2K-C2224TP-1GE / N2K-C2248TP-1GE 	X	✓	✓
N2K-C2232PP-10GE 	X	✓	✓
N2K-C2232TM-10GE 	X	✓	✓
N2K-C2248TP-E-1GE 	X	✓	✓

FEX	vPC Peerlink	VPC Member Port	
		NEXUS 5000/5500 parent switch	NEXUS 7000 parent switch
N2K-B22-HP 	X	✓	6.2.2
N2K-C2248PQ 	X	✓	6.2.2
N2K-C2232TM-E 	X	✓	6.2.2



Appendix III

Convergence & Scalability

Cisco *live!*

vPC Scalability

For Latest Scalability numbers please refer to the scalability limits pages for the platform

- **Nexus 7000/7700:**

N7K Verified Scalability Guide :

http://www.cisco.com/en/US/docs/switches/datacenter/sw/verified_scalability/b_Cisco_Nexus_7000_Series_NX-OS_Verified_Scalability_Guide.html

- **Nexus 5000 /5500**

http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/configuration_limits/limits_521/nexus_5000_config_limits_521.html

- **Nexus 6000**

http://www.cisco.com/en/US/docs/switches/datacenter/nexus6000/sw/configuration_limits/b_N6000_Verified_Scalability_602N11.pdf

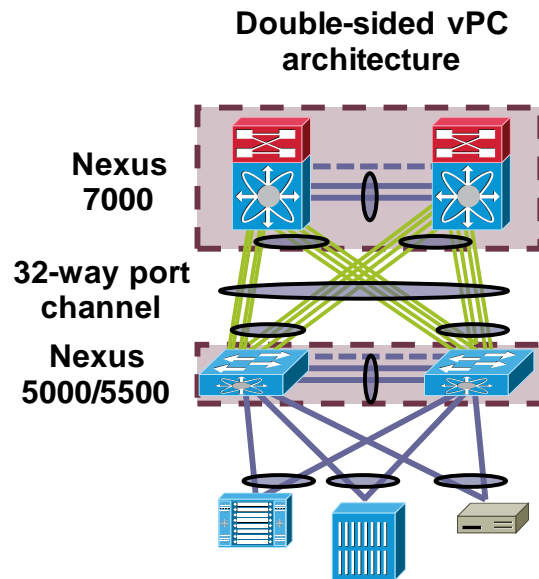
- **Nexus 3000**

http://www.cisco.com/en/US/docs/switches/datacenter/nexus3000/sw/configuration_limits/503_u5_1/b_Nexus3k_Verified_Scalability_503U51.html

Attaching to a vPC Domain

Up to 32-Way Port-Channel – Double-sided VPC

- Multilayer vPC can join eight active member ports of the port-channels in a unique 16-way port-channel*
- vPC peer load-balancing is LOCAL to the peer device
- Each vPC peer has only eight active links, but the pair has 16 active load balanced links (M-series LC)
- F-series Nexus 7000 line cards support 16 way active port-channel load balancing, providing for a 32 way vPC port channel



* Possible with Any Device Supporting vPC/MCEC and Eight-Way Active Port-Channels



Appendix IV

Layer 3 and vPC

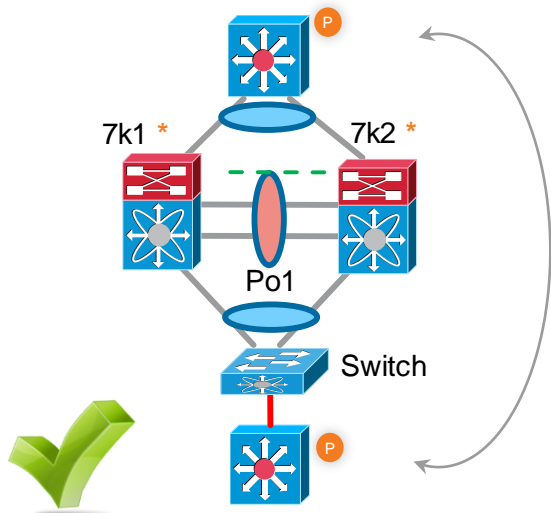
Cisco *live!*

N7K Layer 3 and vPC Designs

Layer 3 and vPC Interactions: Supported Designs

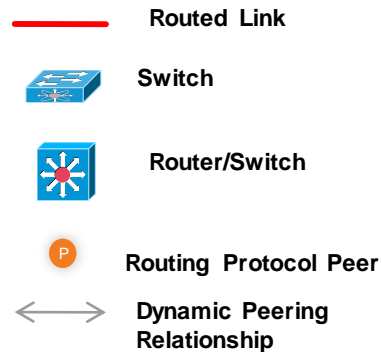
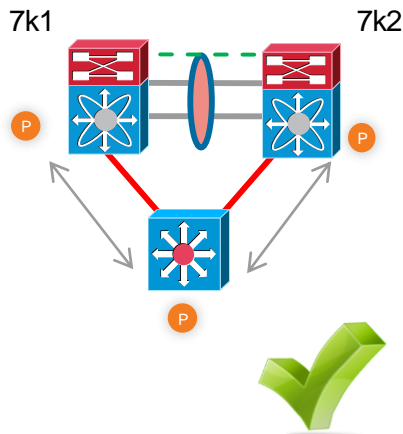


1. Peering between Routers



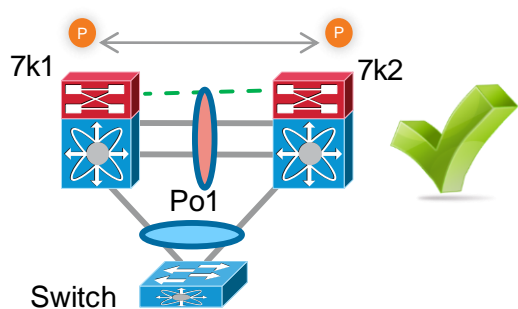
* Nexus 7000 configured for L2 Transport only

2. Peering with an external Router on Routed ports inter-connection

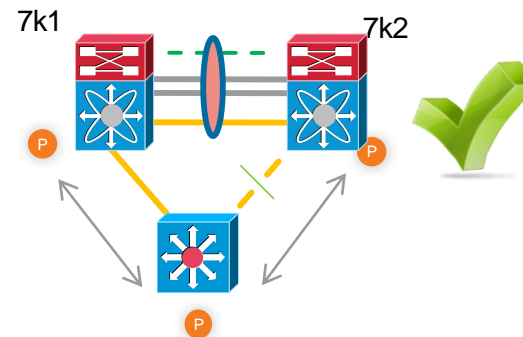


N7K Layer 3 and vPC Designs

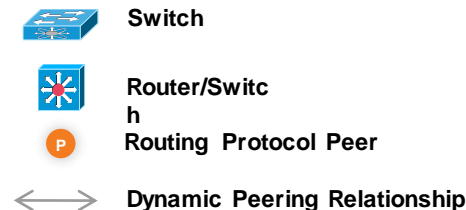
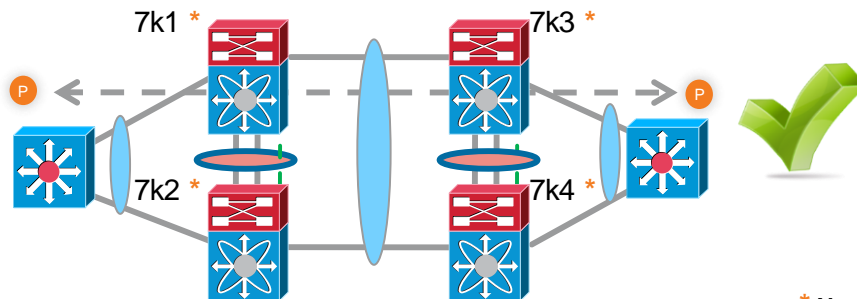
Layer 3 and vPC Interactions: Supported Designs



1. Peering between vPC Device



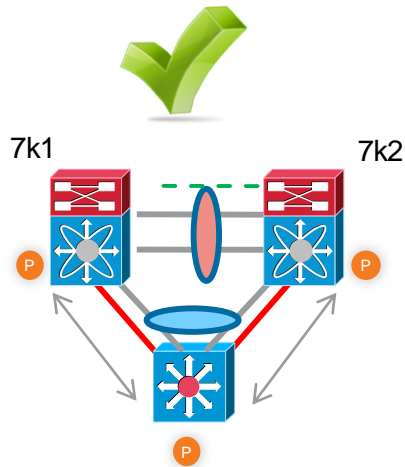
2. Peering over an STP inter-connection **NOT using** a vPC VLAN (Orange VLANs/Links)



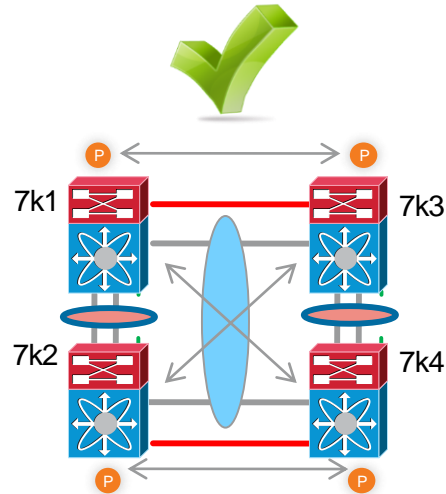
* Nexus 7000 configured for L2 Transport only

N7K Layer 3 and vPC Designs

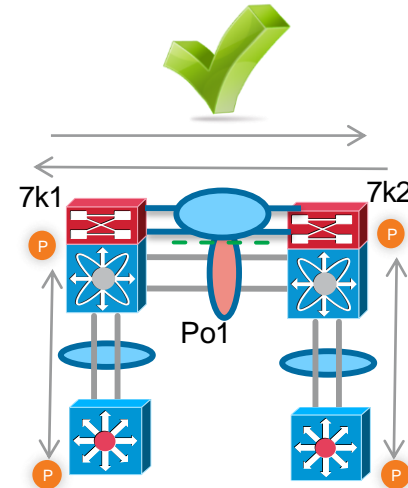
Layer 3 and vPC Interactions: Supported Designs



1. Peering with an external Router on parallel Routed ports inter-connection



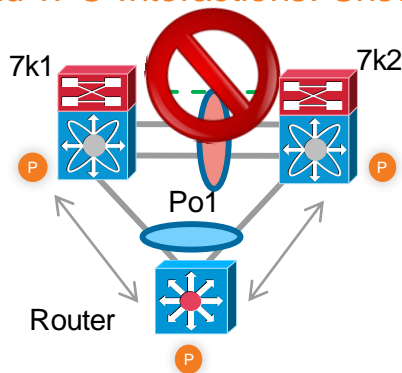
2. Peering over a vPC inter-connection (DCI case) on parallel Routed ports inter-connection



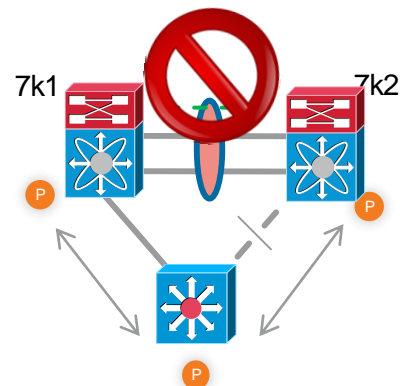
3. Peering over PC inter-connection and dedicated inter-switch link using **non-vPC VLAN**

N7K Layer 3 and vPC Designs

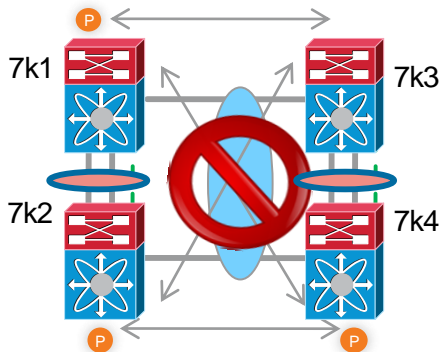
Layer 3 and vPC Interactions: Unsupported Designs



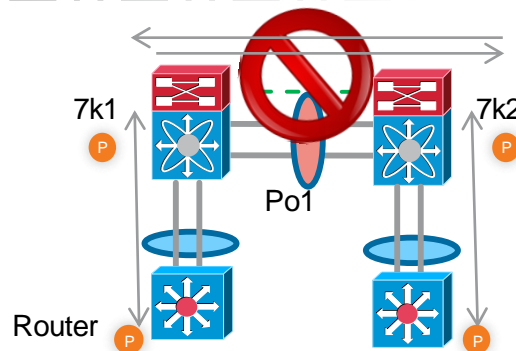
1. Peering over a vPC inter-connection



2. Peering over an STP inter-connection using a vPC VLAN



3. Peering over a vPC inter-connection (DCI case)



4. Peering over PC inter-connection and over vPC peer-link using vPC VLAN



Appendix V

Convergence

Cisco *live!*



Convergence

Agenda

- Feature Overview
- vPC Design Guidance and Best Practices
- vPC Enhancements
- Convergence

vPC Convergence

vPC Convergence Results



- Measured Unicast vPC failover and recovery times
- Convergence time is measured in the following scenarios*
 - vPC link member failure → Sub-second
 - vPC port-channel failover → Sub-Second
 - vPC Peer-Link Failure → Sub-Second
 - vPC Peer-Keepalive Link Failure → Hitless
 - vPC primary/secondary device failure → Sub-Second
 - vPC Supervisor Failover/Switchover → Hitless
 - vPC ISSU device Upgrade/Downgrade → Hitless

*NOTE: Convergence numbers may vary depending on the specific configuration (i.e. scaled number of VLANs/SVIs or HSRP groups) and traffic patterns (i.e. L2 vs. L3 flows).

A long-exposure photograph of a city street at night. The foreground is filled with vibrant, multi-colored light trails from moving vehicles, creating a sense of motion. In the background, a pedestrian bridge spans the street, and tall buildings with lit windows and signage line the street. The overall scene is a dynamic urban environment.

Appendix VI

Reference Material

Cisco *live!*

Reference Material



- **vPC white Paper:**
http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9402/white_paper_c11-516396.html
- **vPC design guides:**
http://www.cisco.com/en/US/partner/products/ps9670/products_implementation_design_guides_list.html
- **vPC and VSS Interoperability white Paper:**
http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/white_paper_c11_589890.html
- **Data Centre Design—IP Network Infrastructure:**
http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DC_3_0/DC-3_0_IPInfra.html
- **Layer 2 Extension Between Data Centres:**
http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/white_paper_c11_493718.html
- **Implementing Nexus 7000 in the Data Centre Aggregation Layer with Services:**
http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/nx_7000_dc.html
- **VPC Best Practices Design Guide:**
http://www.cisco.com/en/US/docs/switches/datacenter/sw/design/vpc_design/vpc_best_practices_design_guide.pdf
- **VPC Software Upgrade Technical Note**
http://www.cisco.com/en/US/docs/switches/datacenter/sw/nx-os/tech_note/vpc_upgrade.html





CISCO