
TOMORROW starts here.
cisco
Cisco live!

Versatile architecture of using Nexus 7000 with F and M-series I/O modules to deliver FEX, FabricPath and Multihop FCoE all at the same time

BRKDCT-2237

David Klebanov

Technical Solutions Architect

CCIE #13791


klebanov@cisco.com


@DavidKlebanov

Umar Shafiq

Technical Solutions Architect

CCIE #7119


ushafiq@cisco.com

Cisco live!

Recommended Pre-requisites

- The following sessions are recommended to provide technical background
  - **BRKARC-3470** Cisco Nexus 7000 Switch Architecture
  - **BRKARC-3452** Cisco Nexus 5000/5500 and 2000 Switch Architecture
  - **BRKDCT-2121** Virtual Device Context (VDC) Design and Implementation Considerations with Nexus 7000
  - **BRKDCT-2048** Deploying Virtual Port Channel in NXOS
  - **BRKDCT-2081** Cisco FabricPath Technology and Design
  - **BRKDCT-1044** FCoE for the IP Network Engineer
  - **BRKSAN-2047** FCoE - Design, Operations and Management Best Practices
  - **BRKMPL-2107** Deployments with MPLS on NX-OS
  - **BRKRST-3045** LISP – A Next Generation Networking Architecture

- Don’t worry… we will provide technology refreshers throughout the presentation

© 2013 Cisco and/or its affiliates. All rights reserved.

Icons and Color Coding

Network (FabricPath)

Network (Classical Ethernet)

FEX Fabric Links

FCoE

Logical Links

Cisco Nexus 7000 Switch

Cisco Nexus 5500 Switch

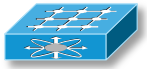

Cisco Nexus 2200 Fabric Extender (1Gb)

Cisco Nexus 2200 Fabric Extender (1/10Gb)

F1 M1
F2/E M

Cisco Nexus 7000 I/O Modules (with color coding)

ABC

Command Line Interface Configuration Snippet

© 2013 Cisco and/or its affiliates. All rights reserved.

What We Want

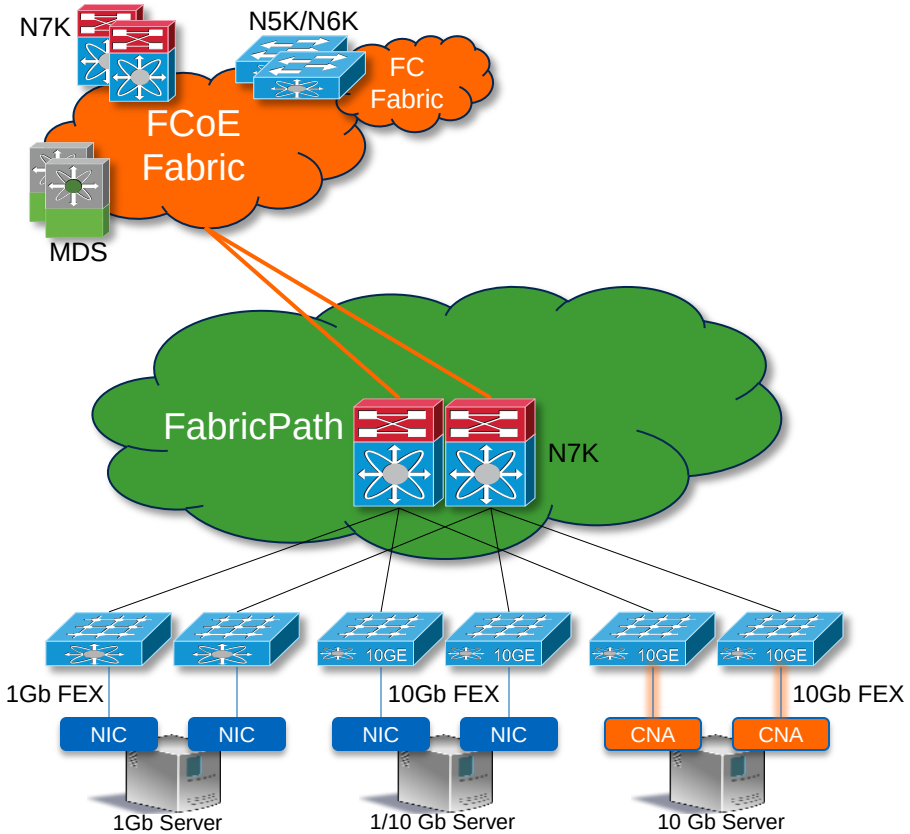

- *Consolidated* yet *feature-rich* hardware deployment
  - Cost savings
- F and M-series I/O modules
  - Hardware investment protection
  - Design versatility
- Fabric Extenders (FEXs)
  - Top of Rack server connectivity model
  - Large scale network access
  - Efficient cabling
  - Simplified management
- FabricPath
  - Seamless large scale Layer 2 connectivity model
  - Easy server provisioning
  - Workload mobility

© 2013 Cisco and/or its affiliates. All rights reserved.

What We Want

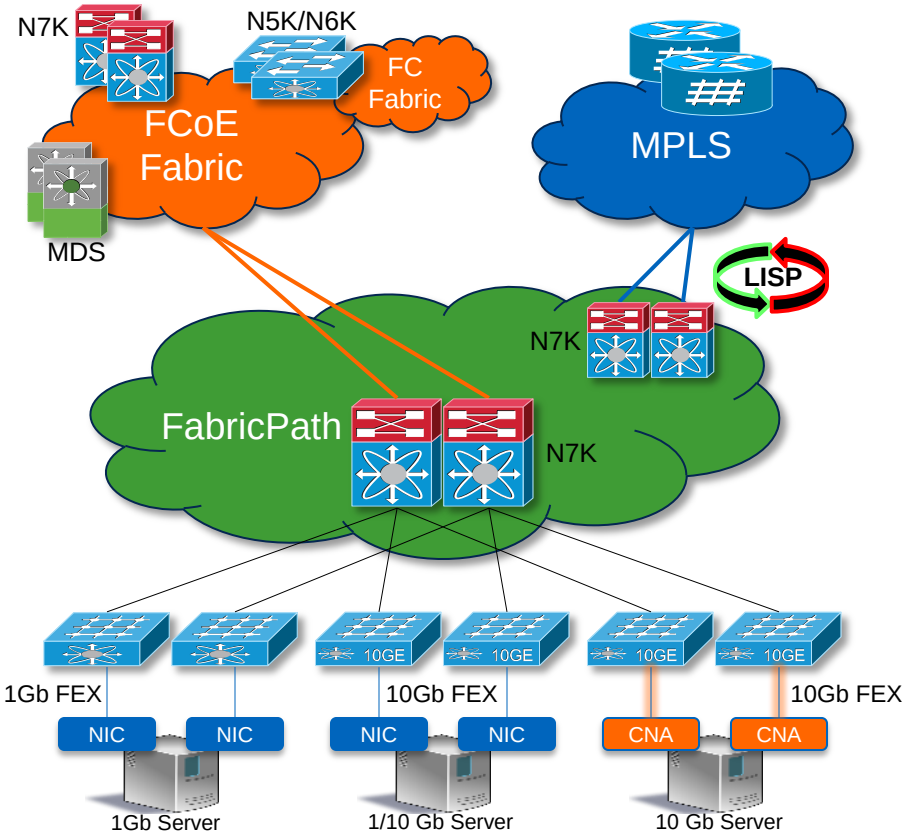

- Fibre Channel over Ethernet
  - Server cabling simplification
  - Cost reduction
  - End-to-end Unified Fabric
- MPLS
  - Extend connectivity beyond a single DC
  - Support for multi-tenancy in the fabric
- Locator ID Separation Protocol (LISP)
  - Virtual Machine mobility aware routing

© 2013 Cisco and/or its affiliates. All rights reserved.

A Trivia Thought…

Can't we just deploy it like this?

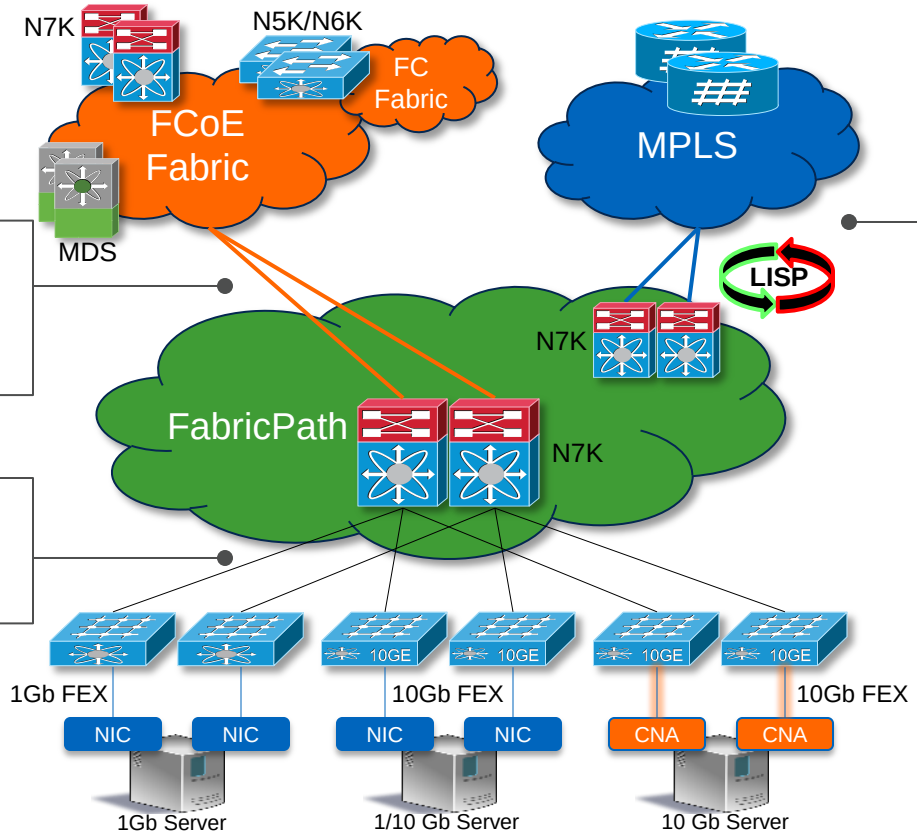

FCoE + FEX

- F2/F2E support both, but not at the same time
- F1 support FCoE, but not FEX
- M support FEX, but not FCoE

FabricPath + FEX

- F2/F2E support both at the same time
- F1 support FabricPath, but not FEX
- M support FEX, but not FabricPath

FabricPath + MPLS + LISP

- F1/F2/F2E support FabricPath, but not MPLS/LISP
- M support MPLS/LISP, but not FabricPath
  - N7K-M132XP-12L

How do we solve it today?

© 2013 Cisco and/or its affiliates. All rights reserved.

Solution Overview

Leveraging Nexus 7000 Virtual Device Contexts

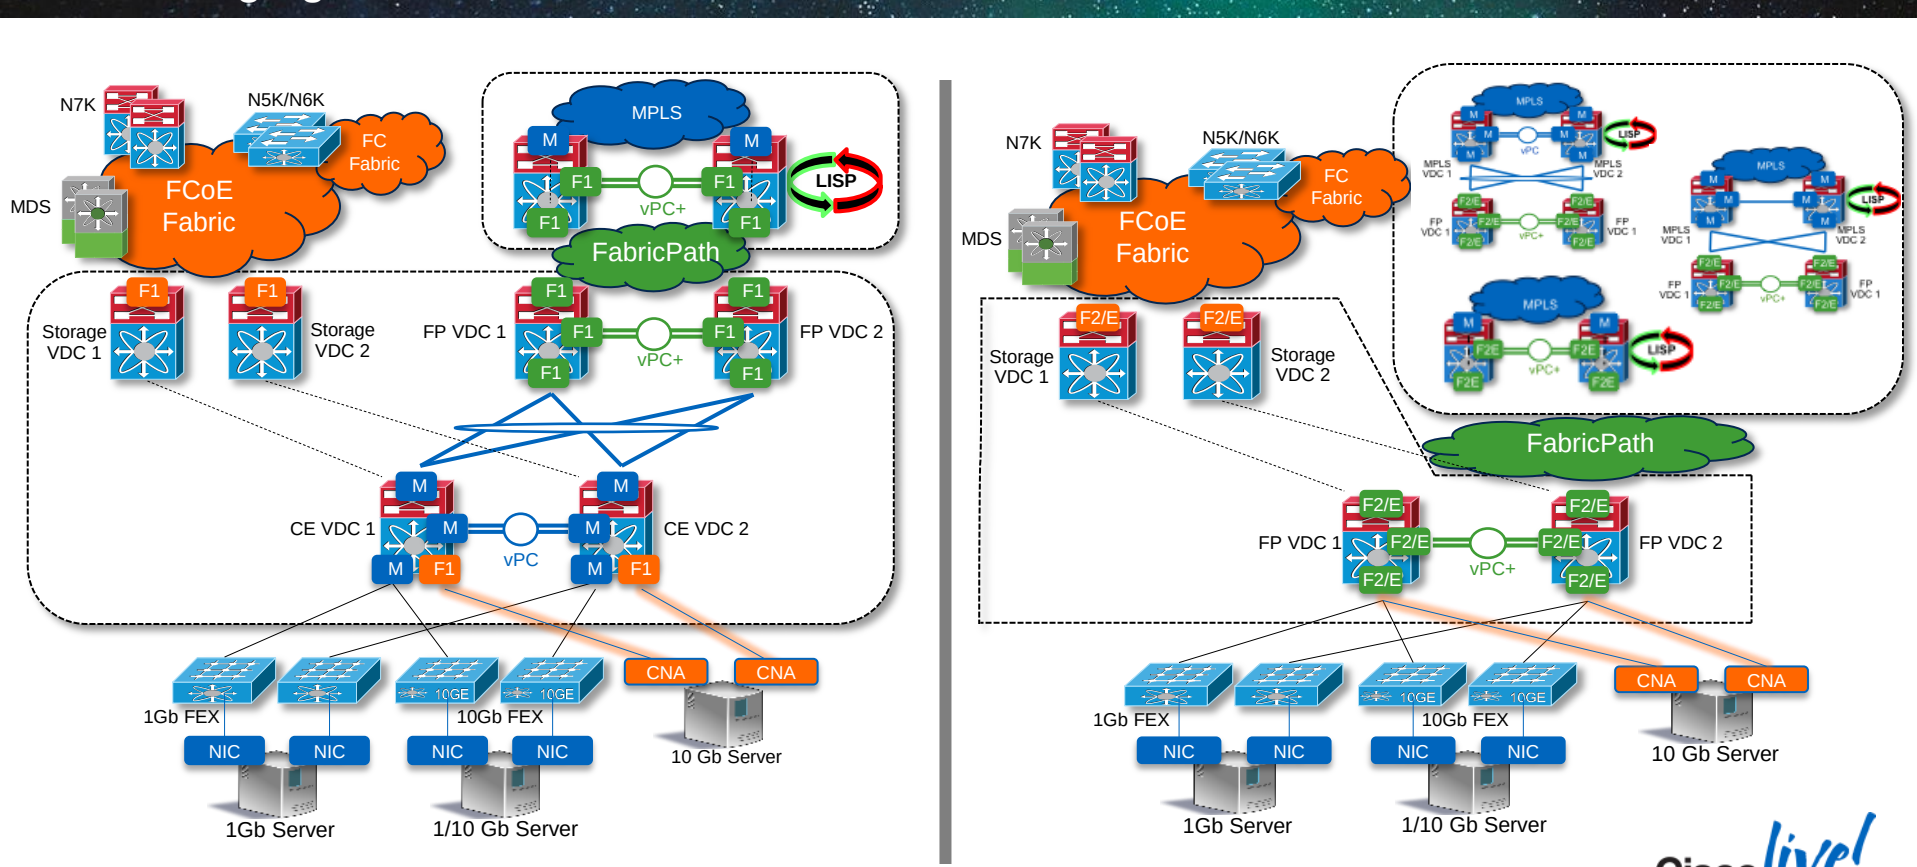

© 2013 Cisco and/or its affiliates. All rights reserved.

Discussion Agenda

- **Fabric Extenders attachment to Nexus 7000 switches**
- Fabric Extenders and Cisco FabricPath on Nexus 7000 switches
- Cisco FabricPath and MPLS/LISP on Nexus 7000 switches
- Multihop FCoE on Nexus 7000 switches

© 2013 Cisco and/or its affiliates. All rights reserved.

FEX Attachment to Nexus 7000

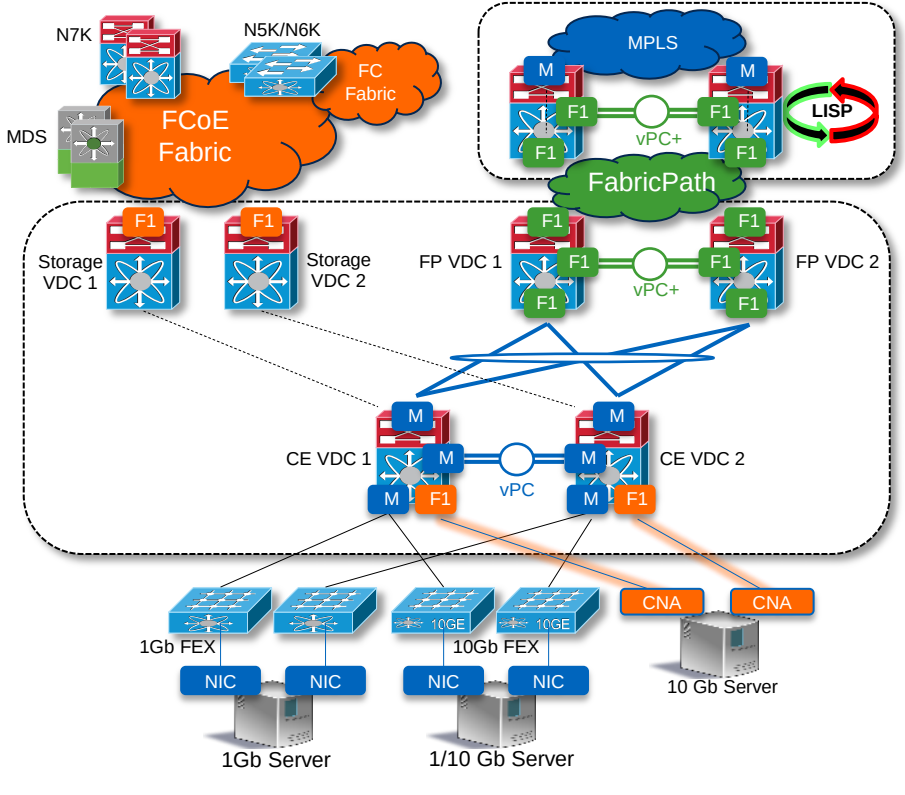

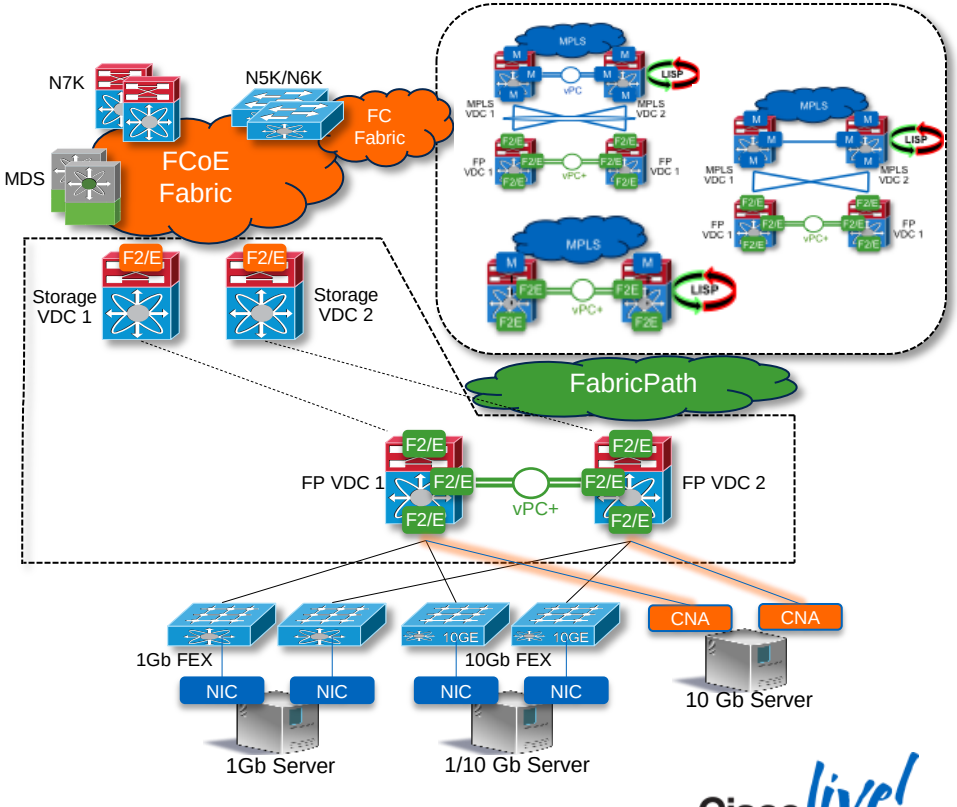

© 2013 Cisco and/or its affiliates. All rights reserved.

Switching and Routing Component
Cisco Nexus 7000

Nexus 7018
- 18 slot system
- 550Gb per slot
- 768 wire-rate 10Gb ports
- 96 wire-rate 40Gb porta
- 32 wire-rate 100Gb ports

Nexus 7010
- 10 slot system
- 550Gb per slot
- 384 wire-rate 10Gb ports
- 48 wire-rate 40Gb ports
- 16 wire-rate 100Gb porta

Any Supported N7K Can Be Used for Our Solution

Nexus 7009
- 9 slot system
- 550Gb per slot
- 336 wire-rate 10Gb ports
- 42 wire-rate 40Gb ports
- 14 wire-rate 100Gb ports

Nexus 7004
- 4 slot system
- 440Gb per slot
- 96 near wire-rate 10Gb ports
- 12 wire-rate 40Gb ports
- 4 wire-rate 100Gb ports

http://www.cisco.com/go/nexus7000

© 2013 Cisco and/or its affiliates. All rights reserved.

Port Extension Component
Cisco Nexus 2000


N2248TP (-E)

- 48 Port 100Mb/1Gb Host Interfaces (RJ-45)
- 4 x 10Gb Uplinks (SFP+/Twinax/FET)

N2224TP

- 24 Port 100Mb/1Gb Host Interfaces (RJ-45)
- 2 x 10Gb Uplinks (SFP+/Twinax/FET)

N2232PP

- 32 Port 1/10Gb FCoE Host Interfaces (SFP/SFP+)
- 8 x 10Gb Uplinks (SFP+/Twinax/FET)

Any Supported FEX Can Be Used for Our Solution

N2232TM (-E)

- 32 Port 1/10Gb 10G-BaseT Host Interfaces (RJ-45)
- 8 x 10Gb Uplinks (SFP+/Twinax/FET)
- *N2232TM-E is not currently supported on Nexus 7000*

N2248PQ

- 48 Port 1Gb/10Gb Host Interfaces (SFP/SFP+)
- 4 x 40Gb Uplinks (QSFP+/Twinax40)
- *Not currently supported on Nexus 7000*

http://www.cisco.com/go/nexus2000

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Nexus 7000-2000 Distributed Switch

Top of Rack Cabling with End of Row Management

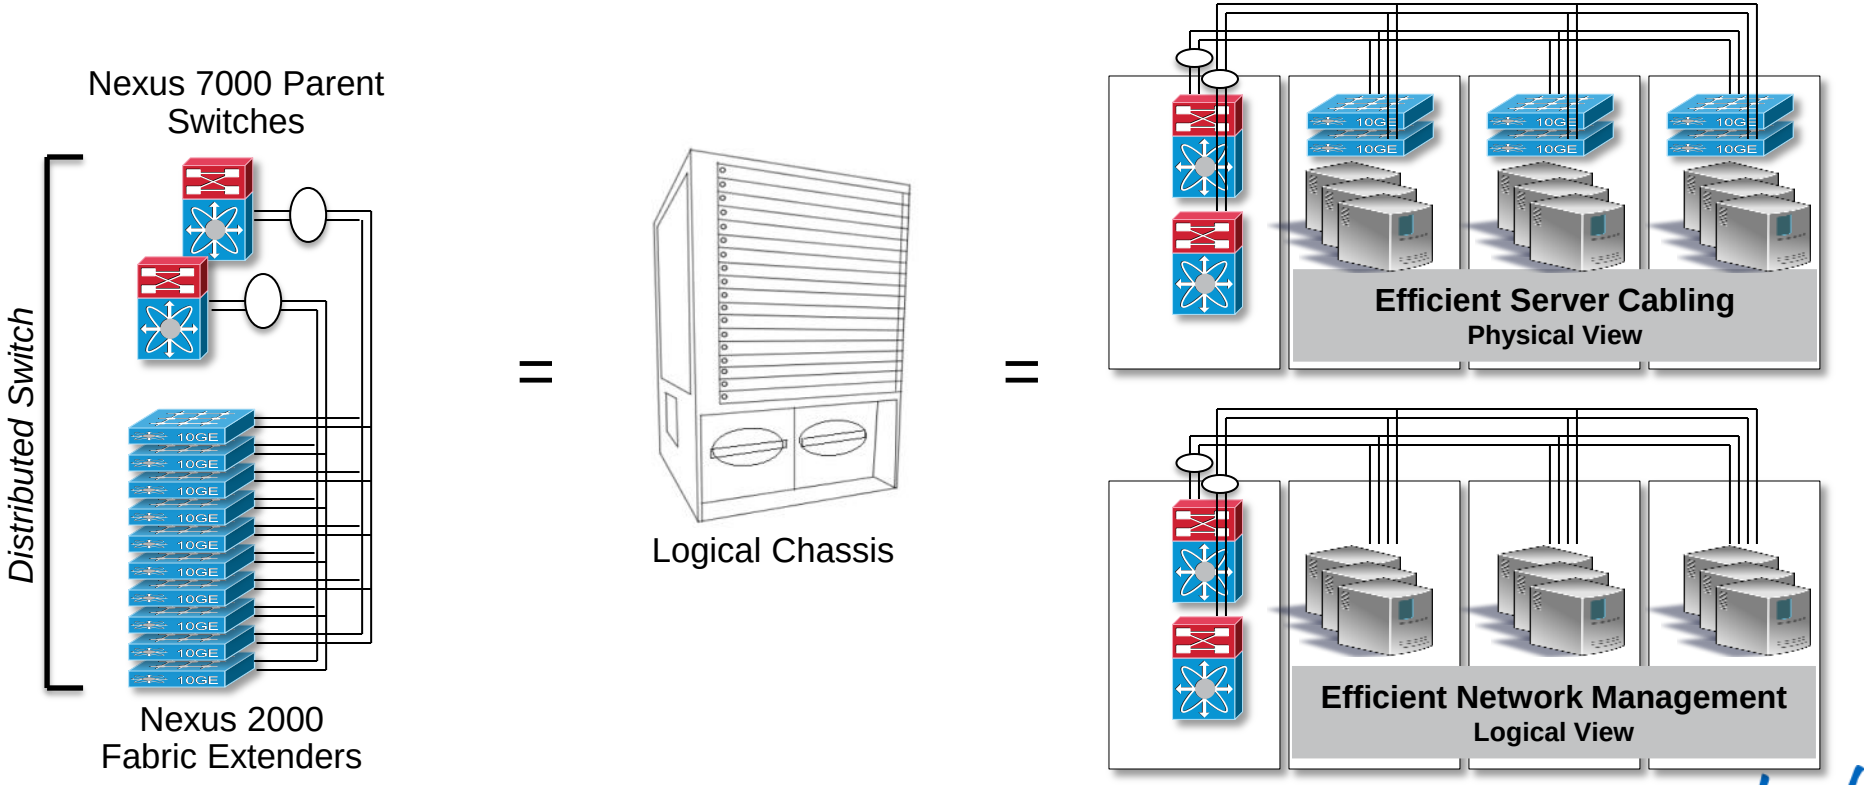

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Nexus 7000 Parent Switch Requirements

Fabric Extender Attachment

Memory

Sup 1

N7K-SUP1-8GBUPG

Sup 2/2E

Parent I/O Module

N7K-M132XP-12L

N7K-M224XP-23L

N7K-F248XP-25
N7K-F248XP-25E

* N7K-M132XP-12 is EOS

CLI

- Install and activate the feature
  - No special license is required

```
N7K(config)# install feature-set fex
N7K(config)# feature-set fex
```

Note: Please refer to cisco.com for full configuration

Software

NX-OS

- Sup 2/2E requires NX-OS 6.1
- N7K-SUP1-8GBUPG requires NX-OS 5.1
- N7K-F248XP-25 requires NX-OS 6.0
- N7K-F248XP-25E requires NX-OS 6.1
- N7K-M224XP-23L requires NS-OS 6.1
- N2248TP-E requires NX-OS 6.1
- N2224TP requires NX-OS 5.2
- N2232PP requires NX-OS 5.2
- N2232TM requires NX-OS 6.1

© 2013 Cisco and/or its affiliates. All rights reserved.

FEX Integration Topology

Connectivity into Nexus 7000 Parent Switches

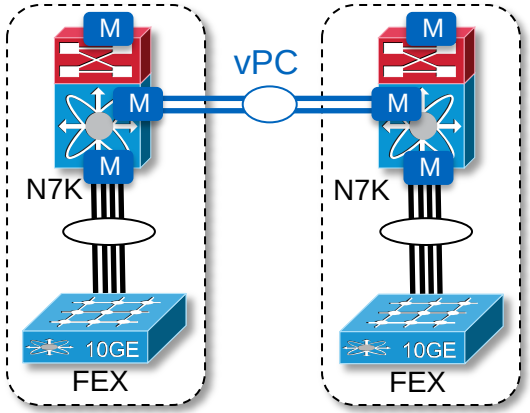

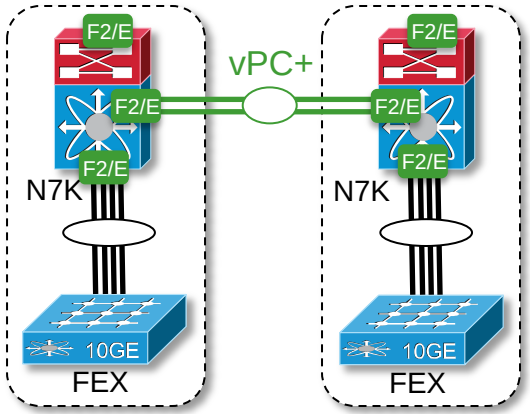

- Same topology for M and F2/F2E
- Fabric Extenders can currently be attached in straight-through topology only*
  - Each parent Nexus 7000 switch is fully redundant
  - Different from the N5K/N6K

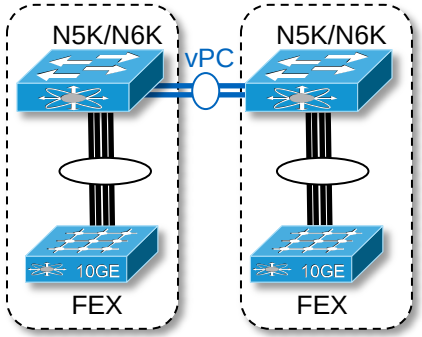

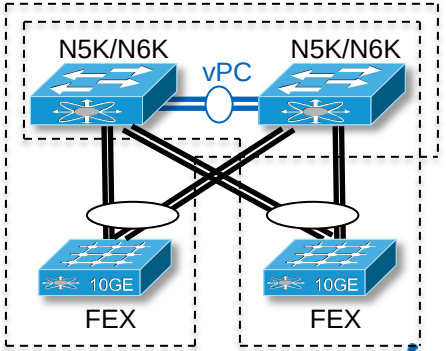

* Check documentation for the latest information

© 2013 Cisco and/or its affiliates. All rights reserved.

FEX Integration Topology

Connectivity into Nexus 7000 Parent Switches

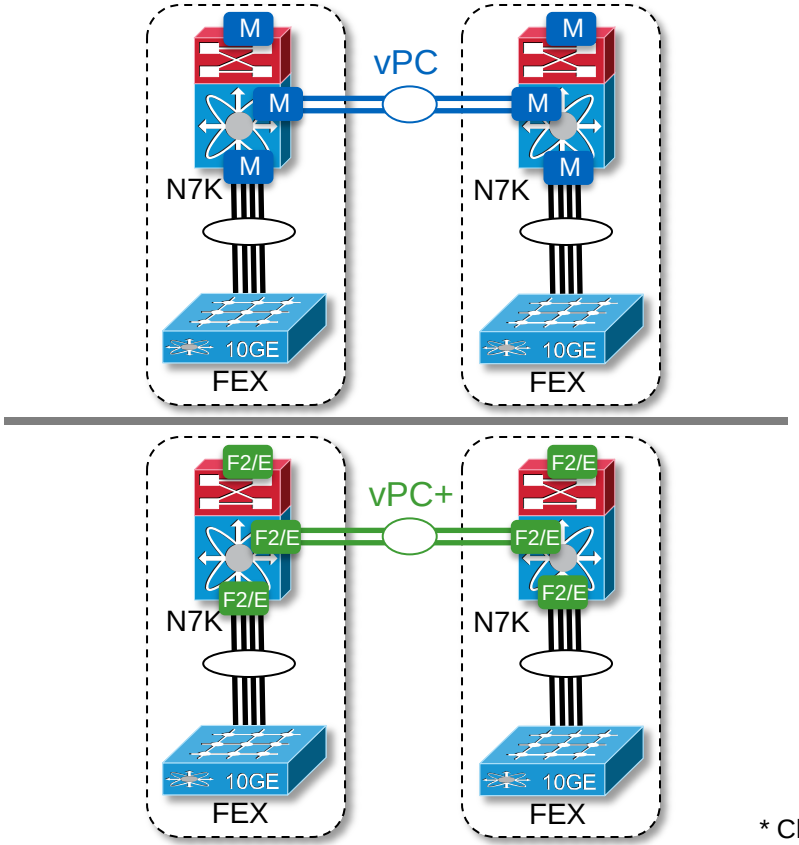

- Same topology for M and F2/F2E
- Fabric Extenders can currently be attached in straight-through topology only*
  - Each parent Nexus 7000 switch is fully redundant
- Each Cisco Nexus 7000 currently supports attaching **48 Fabric Extenders***
  - Sup 1/2 supports 32 Fabric Extenders
- Mix and match any supported FEX type on the same Nexus 7000 parent switch

* Check documentation for the latest information

© 2013 Cisco and/or its affiliates. All rights reserved.
Cisco Public

FEX Integration Topology

Connectivity into Nexus 7000 Parent Switches using F2/F2E I/O Modules

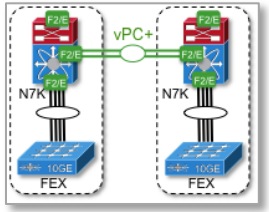

F2/F2E

SoC SoC

1 3 4 7

2 4 6 8

Fabric Extender

Port-Group {2,4} {6,8}

≠

F2/F2E

SoC SoC

1 3 4 7

2 4 6 8

Fabric Extender

Port-Group {2,4} {6,7}

- F2/F2E SoC is a Switch on Chip and it defines a port-group
- FEX uplinks must have symmetric connectivity into ports across port-groups
- Port-group can be dispatched on different F2/F2E Modules
- Multiple FEX can share a same port-group

© 2013 Cisco and/or its affiliates. All rights reserved.

Server Connectivity Topologies

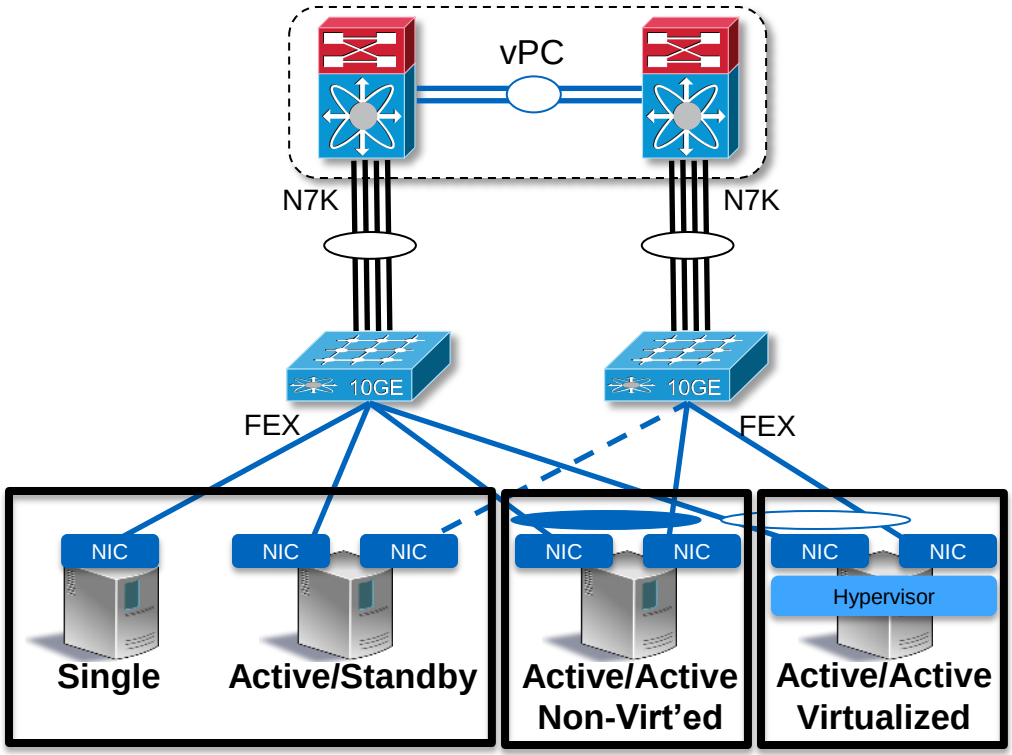

- Active/Active *non-virtualized* servers
  - Port channeling on server NICs
  - Leverage vPC/vPC+ on Nexus 7000
  - Supported on N7K starting from NX-OS 5.2(1)
- Active/Active *virtualized* servers
  - Same as non-virtualized **OR**
  - No port channeling on server NICs
    - ○ Virtual Switch Load Balancing
    - X Can result in traffic blackholing during vPC Dual-Active condition
- Single and Active/Standby
  - X Can result in traffic blackholing during vPC Dual-Active condition

© 2013 Cisco and/or its affiliates. All rights reserved.

vPC Dual-Active Condition Explained

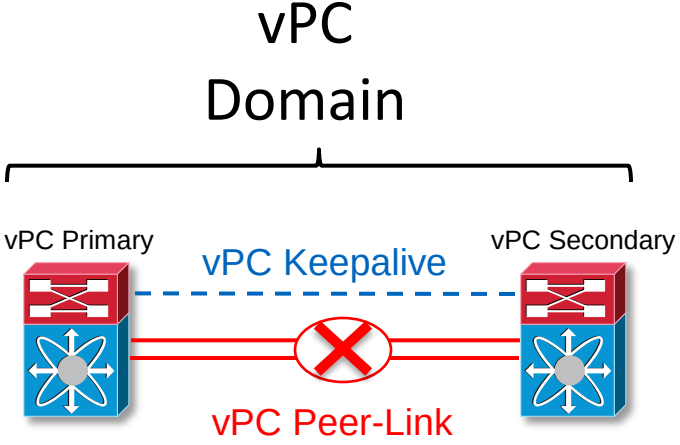

- vPC Dual-Active is triggered when
  a. vPC Peer-Link goes down
  b. vPC Keepalive stays up
- vPC Secondary shuts down all vPC member ports and vPC VLAN SVIs

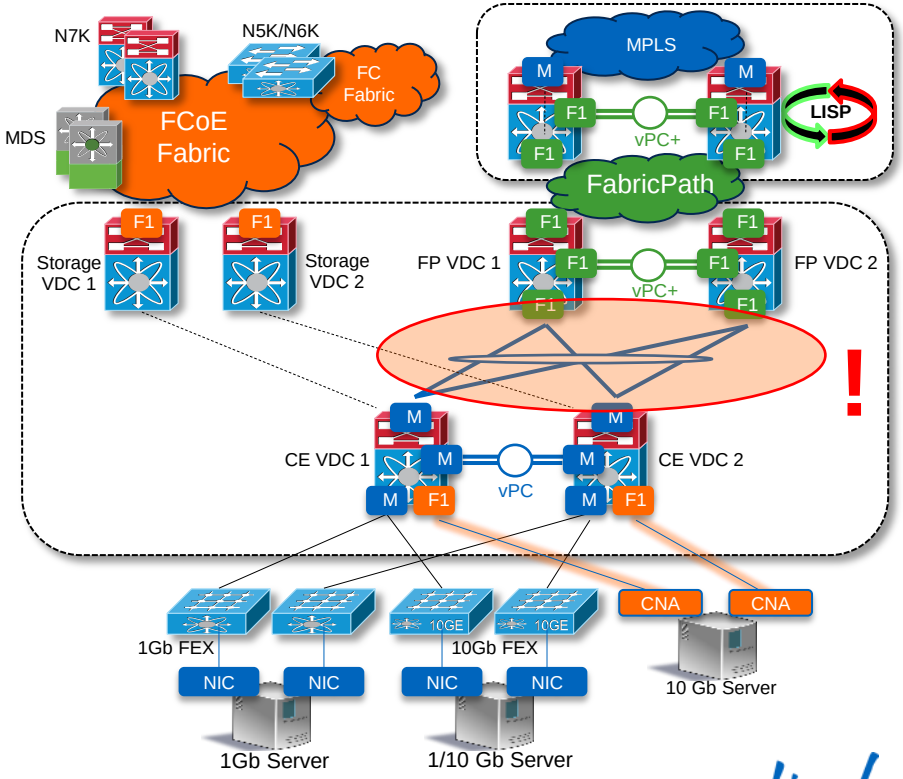

Note: Please refer to cisco.com for full vPC configuration

© 2013 Cisco and/or its affiliates. All rights reserved.

Active/Standby Servers
vPC Dual Active and Orphan Ports (M I/O Modules)

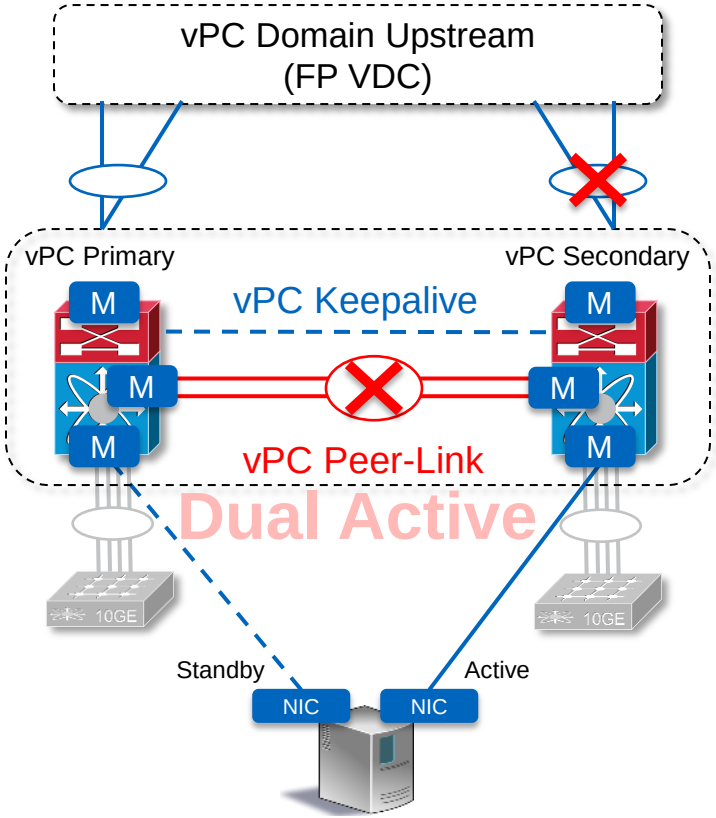

- Active server NIC is connected to vPC Secondary (*important*)
  - Network admin cannot predict which one will become active
  - This is an Orphan Port*
- During vPC Dual-Active condition vPC Secondary will shut down all vPC member ports
- Uplink ports towards FP VDC are vPC member ports, they will be shutdown

* Port that belongs to non-vPC attached device

© 2013 Cisco and/or its affiliates. All rights reserved.

Active/Standby Servers
Traffic Blackholing (M I/O Modules)

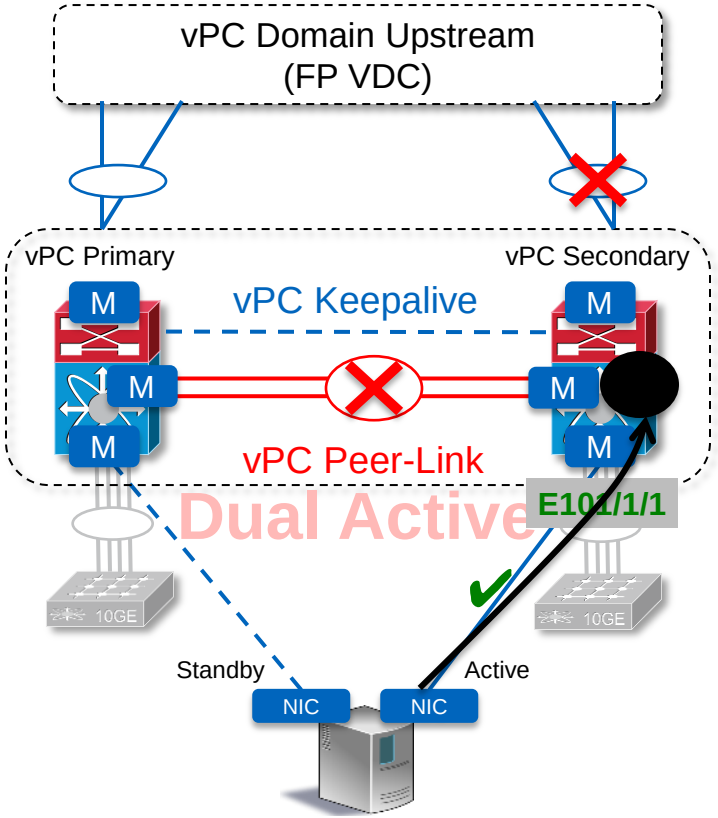

- By default, vPC Dual-Active does **not** shut down Orphan Ports, which is where Active/Standby servers are connected
  - Reason is Orphan Ports are not vPC member ports
- Active/Standby servers will keep on sending traffic towards vPC Secondary, which now has nowhere to send it to
  - vPC Peer-Link is down
  - Uplink was shutdown by vPC Dual-Active condition
- Server traffic is blackholed

© 2013 Cisco and/or its affiliates. All rights reserved.

Active/Standby Servers

Solving Traffic Blackholing using Orphan Port Suspend Feature (M I/O Modules)

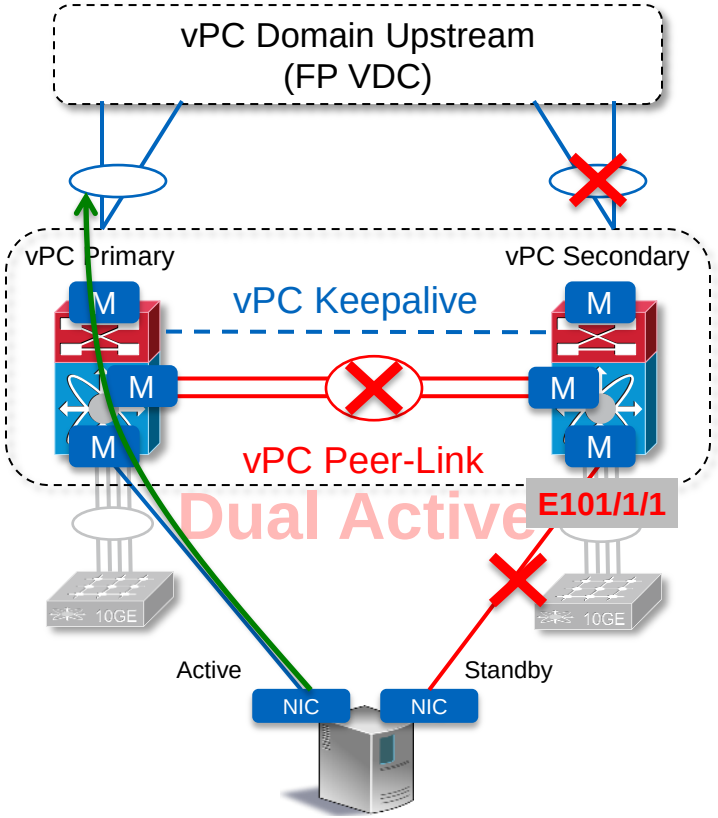

- Solution is to have vPC shut down Orphan Ports on vPC Secondary during Dual-Active condition
  - Orphan Port suspend feature

```
N7K(config)# interface ethernet 101/1/1
N7K(config-if)# vpc orphan-ports suspend
```

- Active/Standby server NIC will perform switchover and carry the traffic along the active path
  - Relies on healthy server operation
- **Recommendation**: Always prefer Active/Active port-channel server connectivity into vPC domain

© 2013 Cisco and/or its affiliates. All rights reserved.

Singly Connected Servers
vPC Dual Active Traffic Blackholing (M I/O Modules)

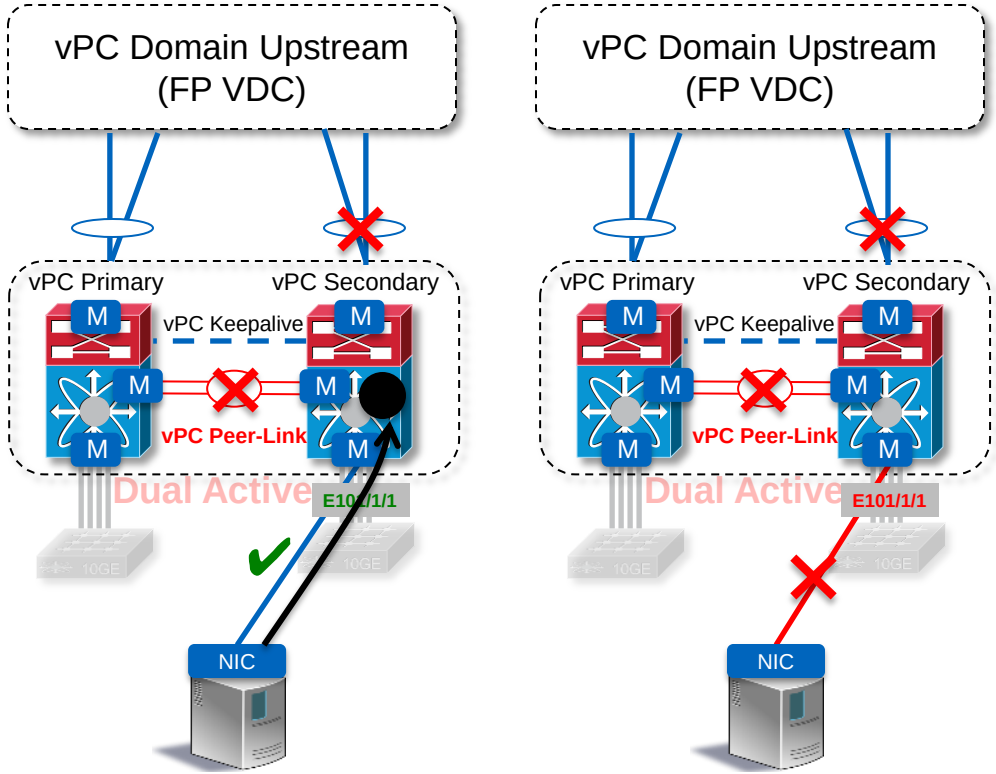

vPC "orphan-port suspend" feature OFF

vPC "orphan-port suspend" feature ON

- If singly connected servers are connected to vPC Secondary, traffic will be blackholed whether "orphan-port suspend" feature is configured on the interface or not
  - Might make a difference if applications or scripts running on the server can take action based on NIC status
  - With "orphan-port suspend" feature being OFF, server will rely on transmission timeouts to determine loss of connectivity
- Singly connected servers need to be physically re-plugged to the vPC Primary or stay down until vPC domain fully recovers from Dual-Active condition

© 2013 Cisco and/or its affiliates. All rights reserved.

Active/Active Virtualized Servers w/o Port Channel

Solving Traffic Blackholing using Orphan Port Suspend Feature (M I/O Modules)

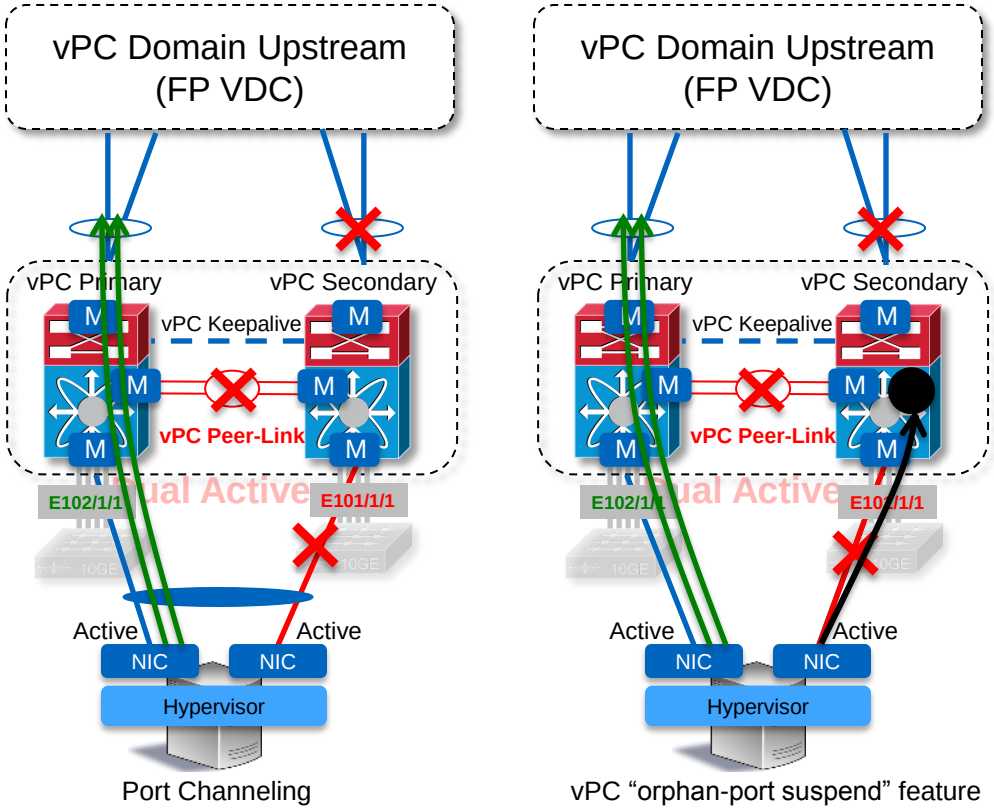

- No Port Channeling on server
  - Traffic is pinned to an uplink by the virtual switch
  - Ports are considered Orphan Ports
- During vPC Dual-Active VM traffic sent over server NIC connected to vPC Secondary will be blackholed
- Solution is to use Orphan Port suspend feature
- … or even better use Port Channeling to achieve Active/Active traffic forwarding
  - No need for Orphan Port suspend feature

© 2013 Cisco and/or its affiliates. All rights reserved.

vPC+ Dual-Active Condition and FabricPath

Design Leveraging F2/F2E I/O Modules

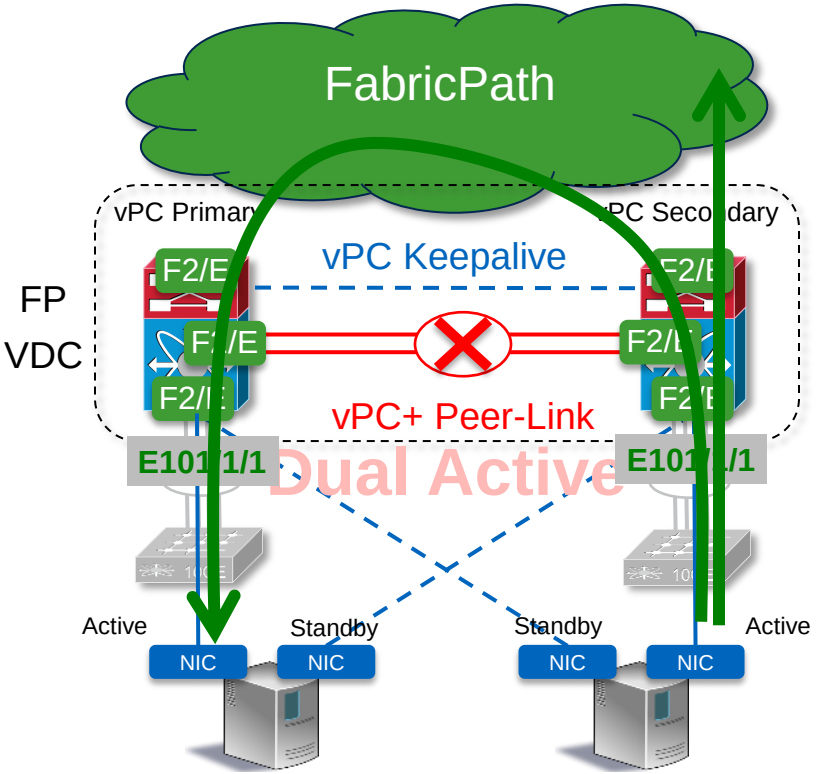

- Traffic blackholing does not apply to design when utilizing F2/F2E I/O modules for FEX attachment into Cisco Nexus 7000 switches
- *No vPC domain upstream*, no vPC member ports to shut down in case vPC Dual-Active condition is triggered
- During vPC Dual-Active traffic is forwarded through FabricPath network
  - No need for vPC Orphan Port Suspend feature

© 2013 Cisco and/or its affiliates. All rights reserved.

FEX Attachment to Nexus 7000 Switches
Progress Checkpoint

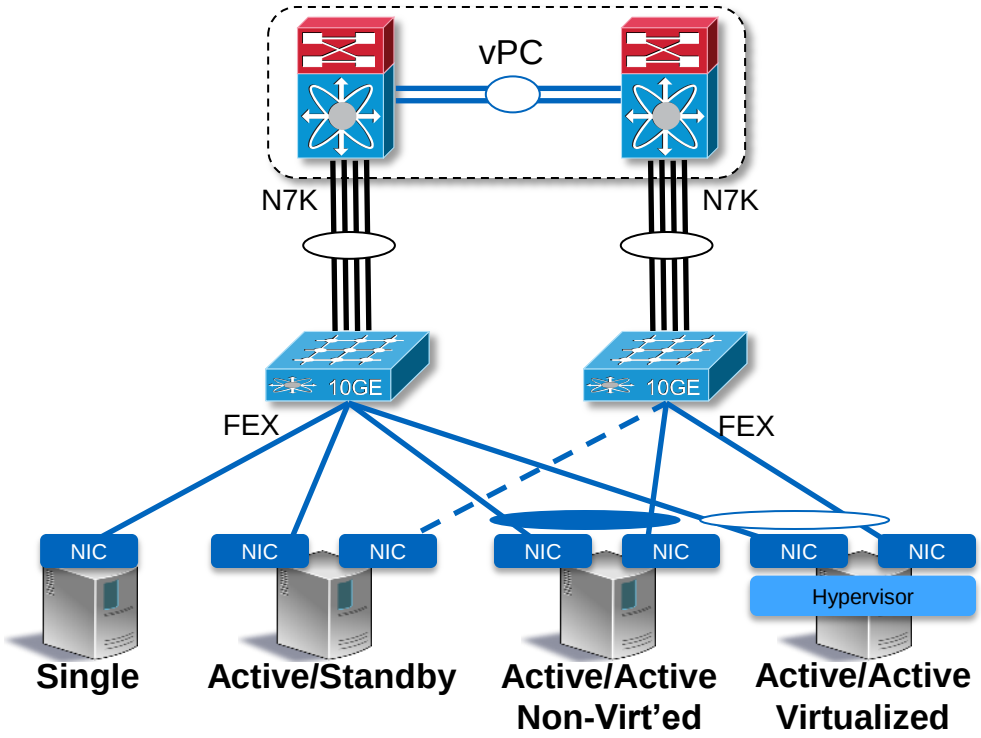

- All design conditions had been satisfied!
  - ✓ Servers are connected to Fabric Extenders
  - ✓ Fabric Extenders are attached to Cisco Nexus 7000 parent switches
  - ✓ Possible traffic blackholing scenarios had been addressed

© 2013 Cisco and/or its affiliates. All rights reserved.

Discussion Agenda

- Fabric Extenders attachment to Nexus 7000 switches
- **Fabric Extenders and Cisco FabricPath on Nexus 7000 switches**
- Cisco FabricPath and MPLS/LISP on Nexus 7000 switches
- Multihop FCoE on Nexus 7000 switches

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco FabricPath and FEXs on Cisco Nexus 7000

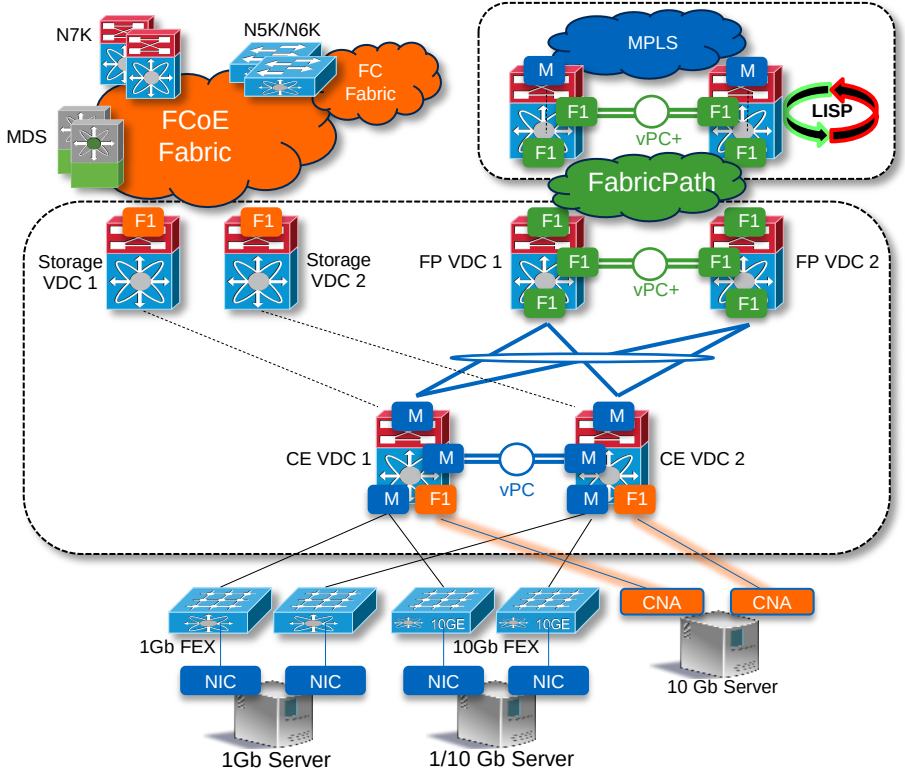

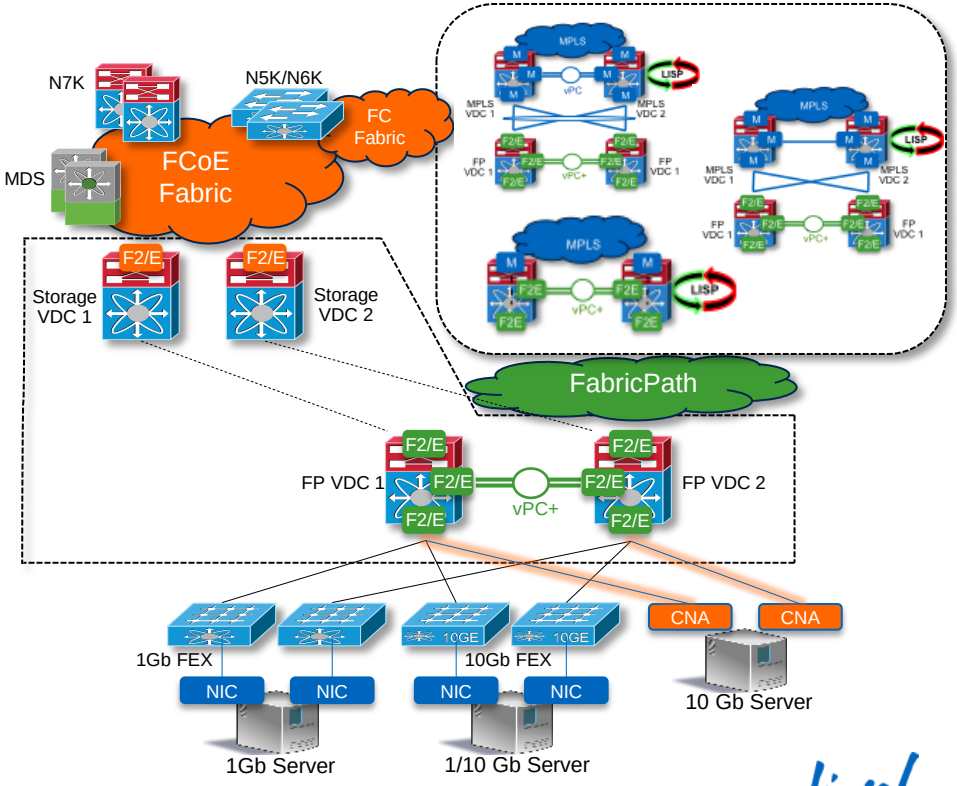

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco FabricPath

Selected Features and Characteristics

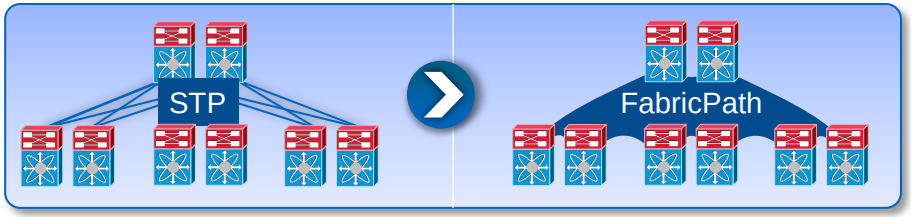

- From PODs to Fabric

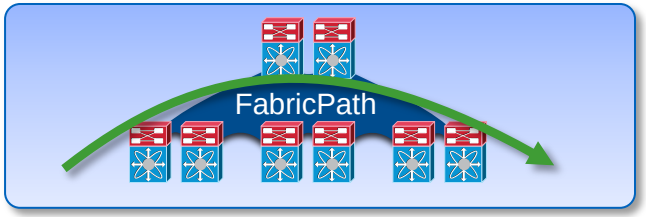

- Optimum Low Latency Switching
- Conversational Learning

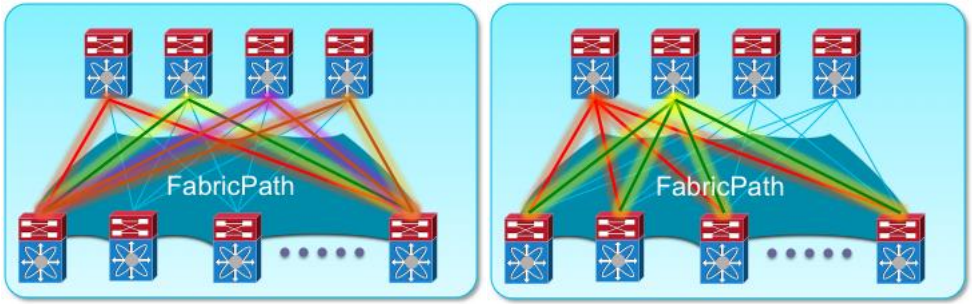

- Unicast and Multicast MultiPathing
  - High bandwidth and High Resilience

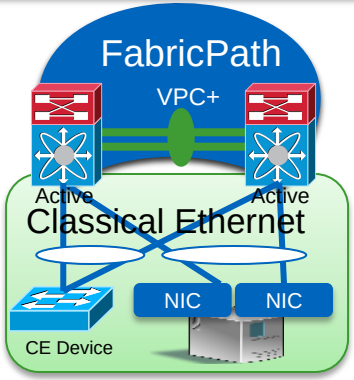

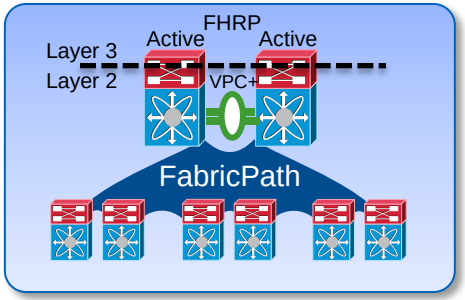

- Classical Ethernet Integration
- Active/Active Default Gateway
- vPC+

*** http://www.cisco.com/go/fabricpath

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Nexus 7000-2000 and FabricPath Edge

Using F2 and F2E I/O Modules

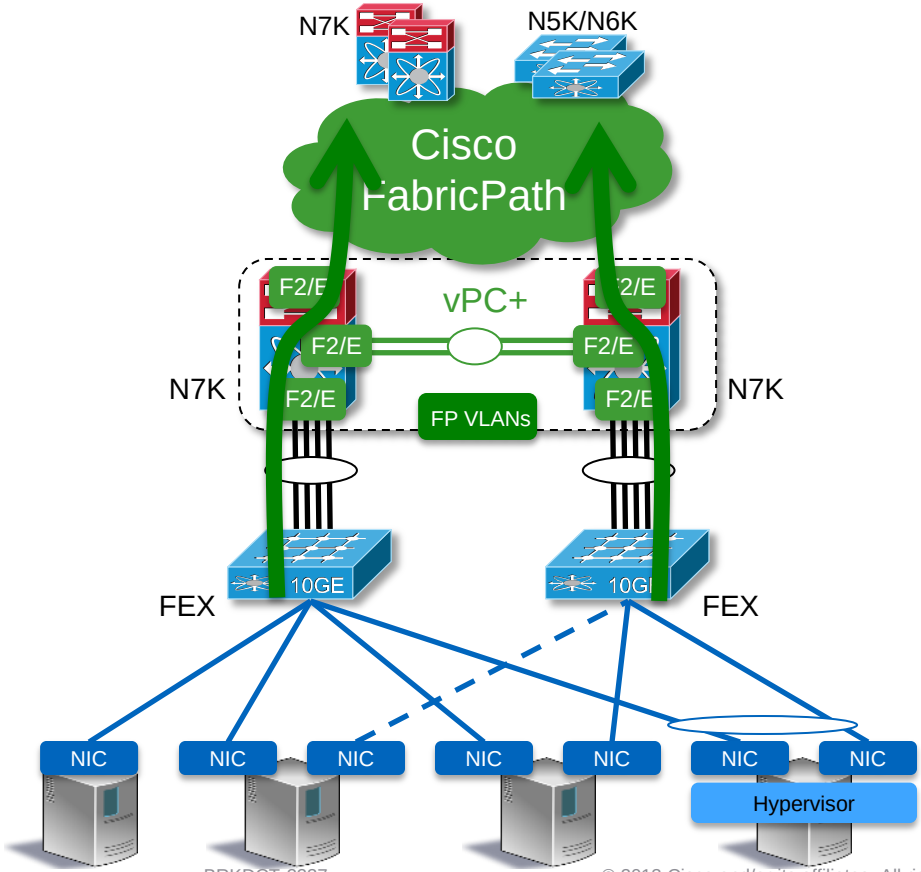

- F2 and F2E I/O modules support Fabric Extender attachment
- F2 and F2E I/O modules support FabricPath
- Server traffic is forwarded through Fabric Extenders, Nexus 7000 switches and on to the rest of the FabricPath environment

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Nexus 7000-2000 and FabricPath Edge

Using F1 and M-series I/O Modules

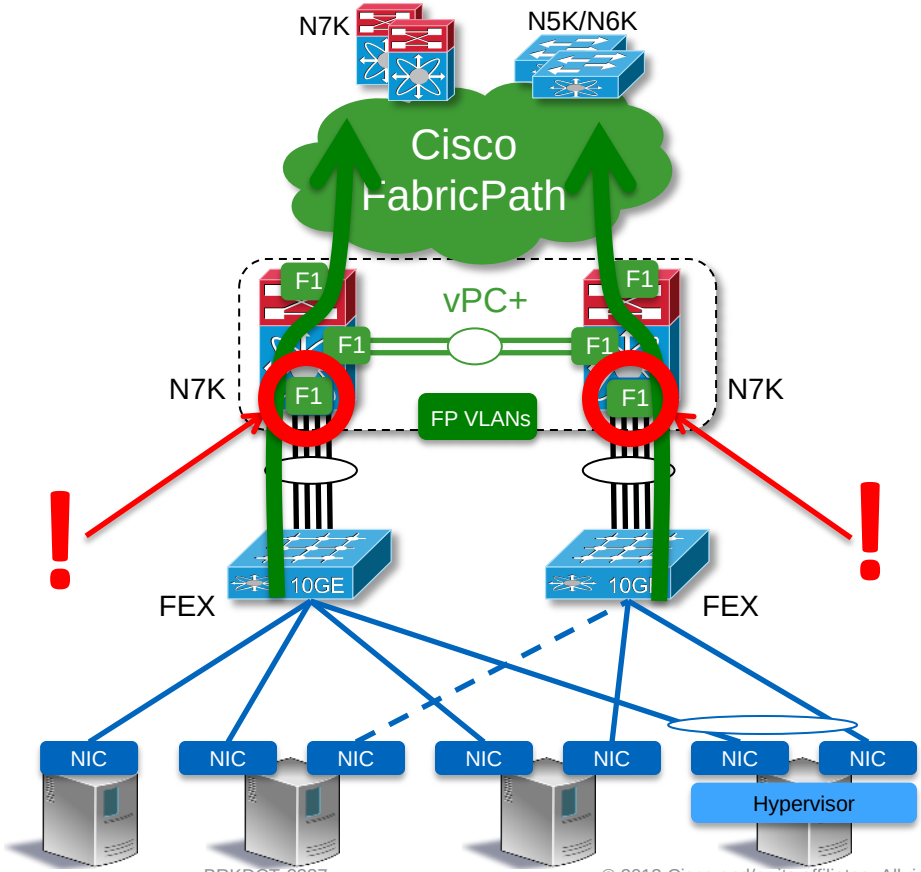

- Server VLANs extended over the FabricPath network must be set to mode FabricPath
- FabricPath mode VLANs cannot exist on M-series I/O modules
- Fabric Extenders **cannot** be attached to F1 ports
- Solution is to use Virtual Device Contexts for design utilizing F1 and M-series I/O modules

© 2013 Cisco and/or its affiliates. All rights reserved.

Nexus 7000 Device Virtualization
VDC Port Assignment

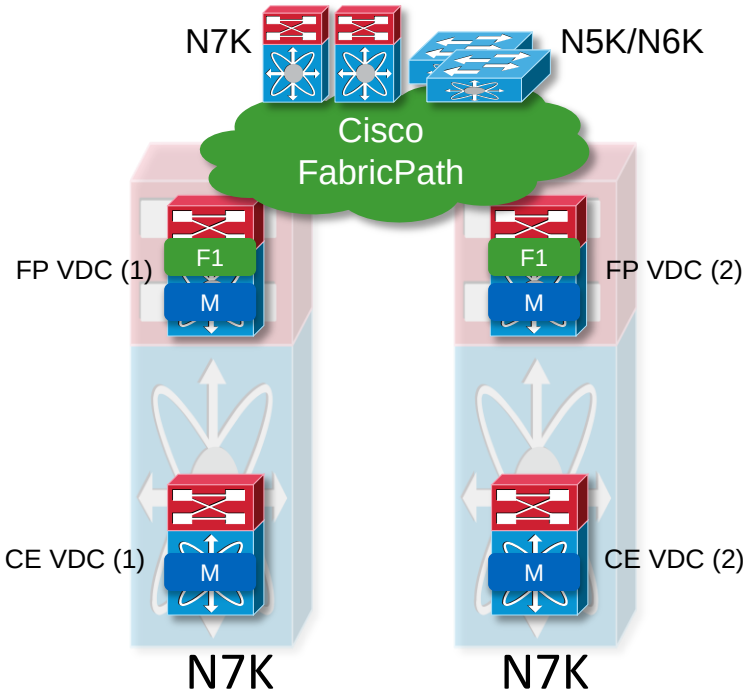

- Nexus 7000 is virtualized into two VDCs
  - Classical Ethernet VDC
  - FabricPath VDC
- Ports on M-series I/O modules are assigned to Classical Ethernet VDC
  - Must observe port ASIC boundaries
- Ports on M-series I/O modules are assigned to FabricPath VDC when routing is required on that VDC
  - Must observe port ASIC boundaries
  - More on FabricPath VDC routing in a short while
- Ports on F1 I/O modules are assigned to FabricPath VDC
  - Must observe SOC boundaries

© 2013 Cisco and/or its affiliates. All rights reserved.

FP VDC and CE VDC Interconnect

Virtual Port Channels

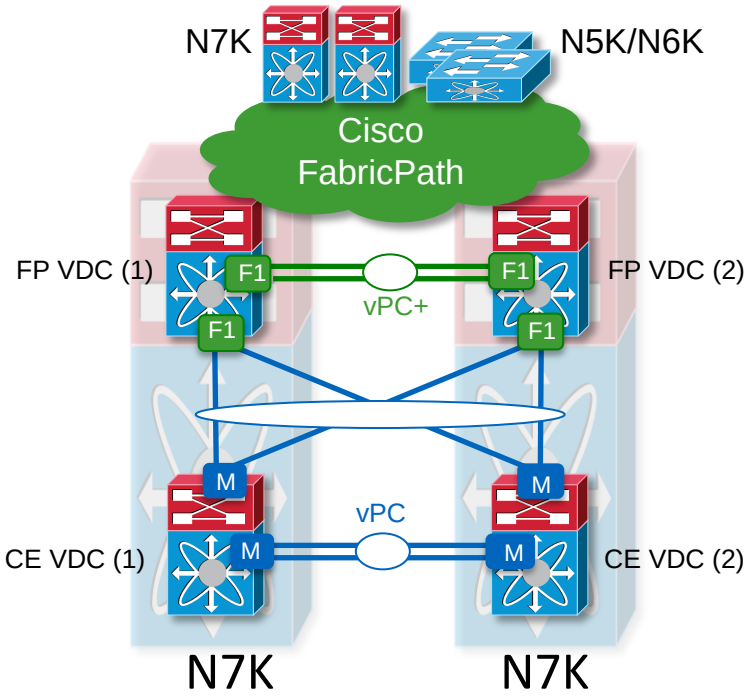

- VDCs are interconnected using external cables
  - Can use TwinAx
  - F1 ports are used on FP VDC side
  - M ports are used on CE VDC side
- Classical Ethernet VDC participates in vPC domain
- FabricPath VDC participates in vPC+ Domain
- Back-to-back vPC design

© 2013 Cisco and/or its affiliates. All rights reserved.

FEX Attachment and FabricPath Connectivity

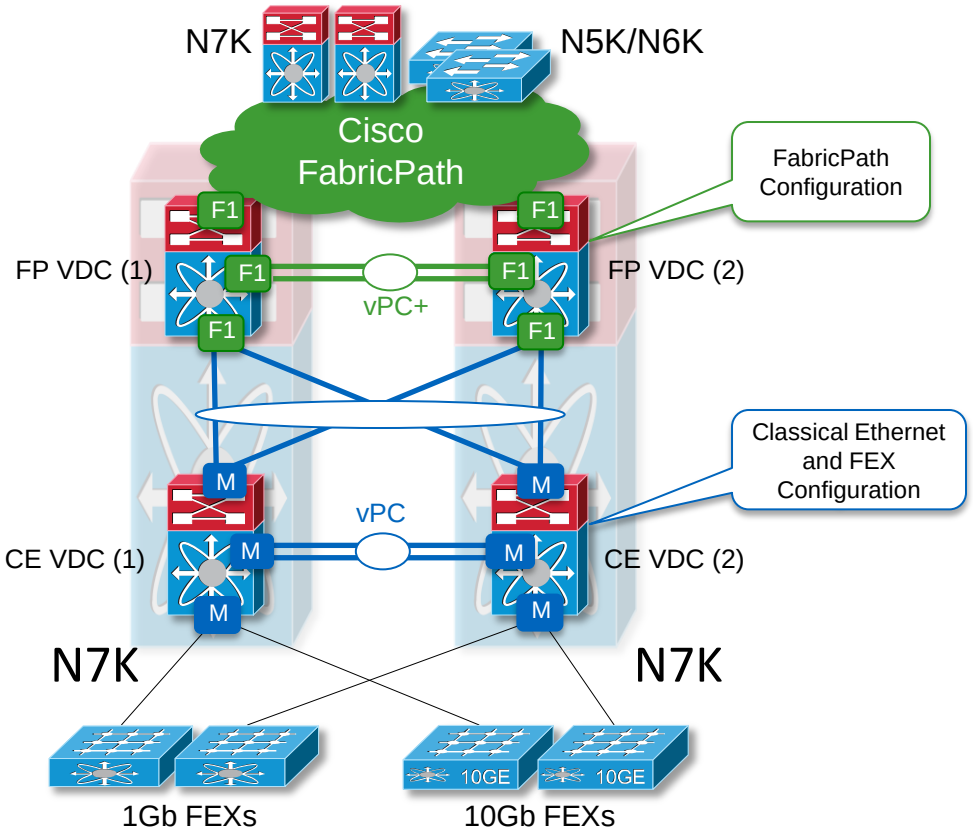

- F1 ports in FP VDC are used to connect to FabricPath network
- Fabric Extenders are attached to M-series I/O modules in a CE VDC
- All Classical Ethernet and FEX configurations are done in the Classical Ethernet VDC
- All FabricPath configurations are done in FabricPath VDC

© 2013 Cisco and/or its affiliates. All rights reserved.
Cisco Public

Extending FabricPath VLANs Between VDCs

VLAN Mode Setting and Trunking

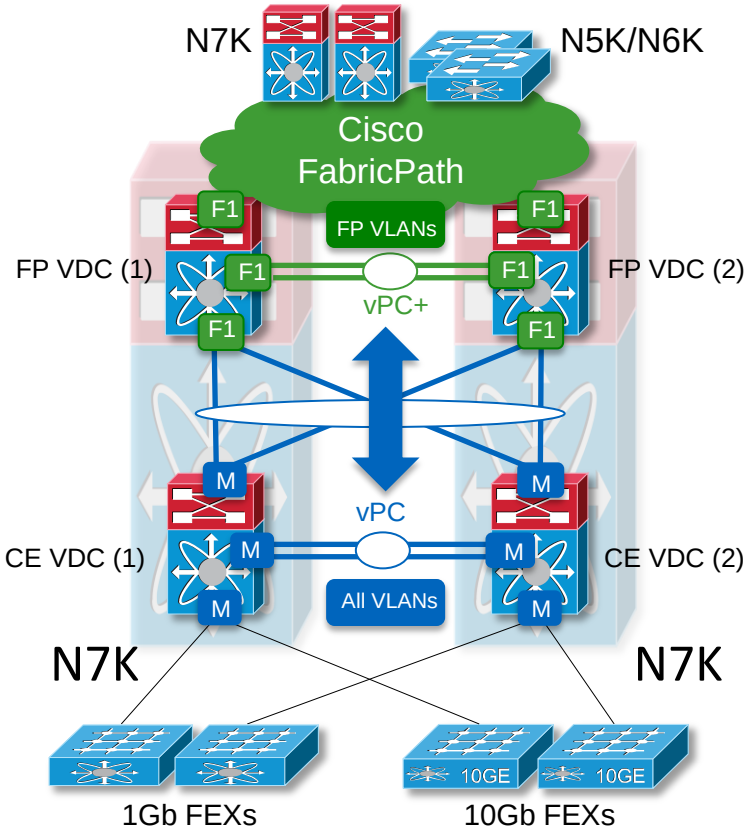

- All VLANs are operating in Classical Ethernet mode in CE VDC
  - Classical Ethernet mode is a default VLAN mode
- VLANs requiring FabricPath transport are trunked between Classical Ethernet VDC and FabricPath VDC over the vPC/vPC+
  - Defined as FabricPath mode VLANs in FabricPath VDC

© 2013 Cisco and/or its affiliates. All rights reserved.

Extending FabricPath VLANs Between VDCs

FabricPath VDC and STP Root Priority

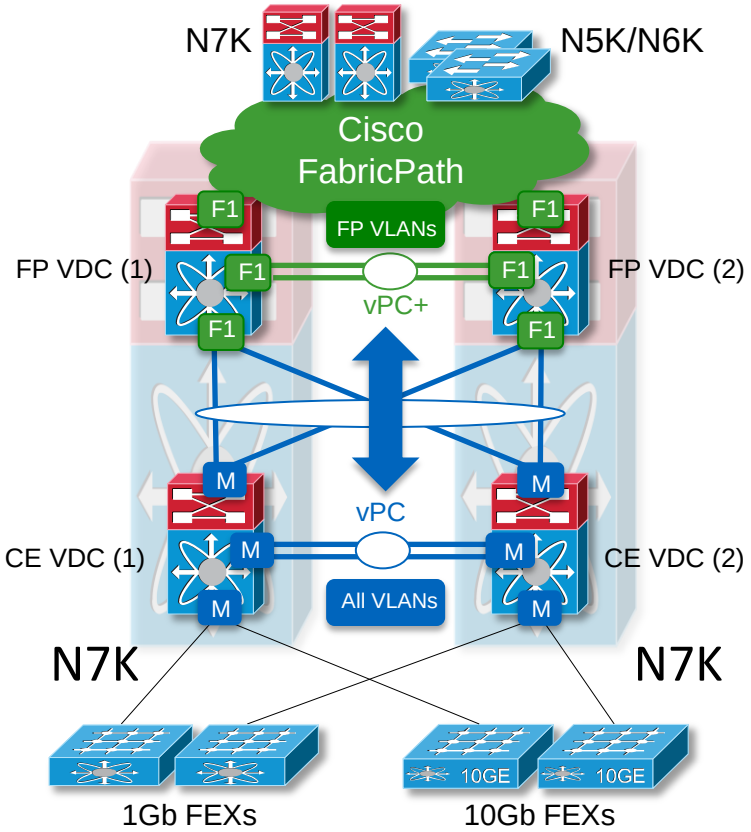

- Useful Reminder… Make sure that FabricPath VDC has lower STP root priority
  - FabricPath network **must be** the root of local STP domain
  - FabricPath edge ports run RootGuard and will place VLANs in root-inconsistent state if “better” BPDUs are received from the local Classical Ethernet environment

n7k(config)# spanning-tree vlan < vlan> priority 4096

© 2013 Cisco and/or its affiliates. All rights reserved.




Server Traffic in the FabricPath Network

Applies to both F2/F2E and F1/M Designs

Global VLANs

All VLANs are forwarded over FabricPath Network

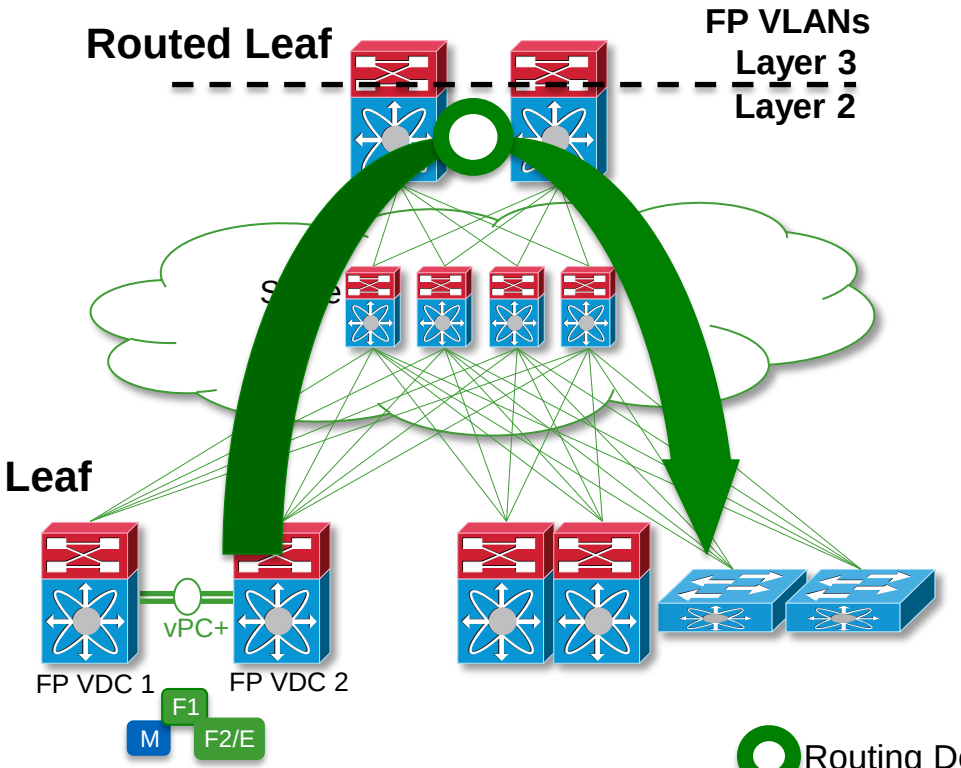

- Any VLAN anywhere
- All VLANs are forwarded over the FabricPath network
- Routing functionality is provided by the Routed Leaf nodes
  - Please refer to the FabricPath design sessions for more in-depth discussion about fabric routing functionality
- Routed Leaf nodes perform inter-VLAN routing sending traffic between source and destination server subnets

© 2013 Cisco and/or its affiliates. All rights reserved.
Cisco Public

Local VLANs

Some VLANs are forwarded over FabricPath Network

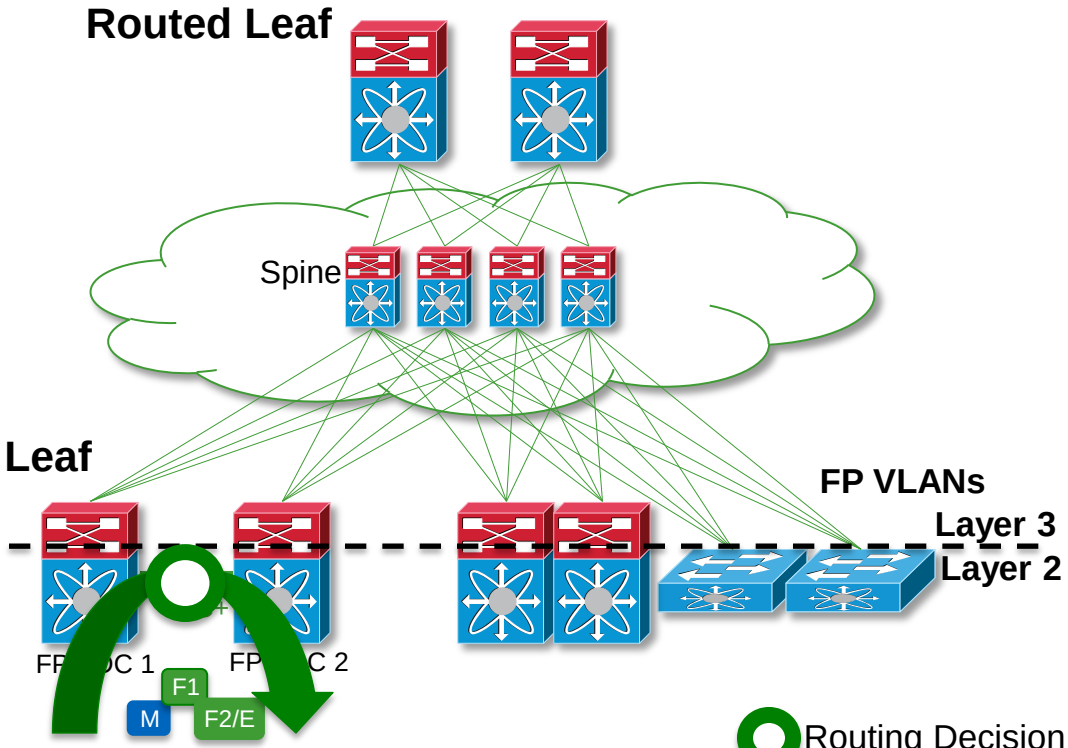

- Maybe not all VLANs need to be extended over the FabricPath network
  - Some servers need to communicate only within their own Leaf Layer node block
  - E.g. no workload mobility across blocks
- VLANs that are not carried over the FabricPath network have their Layer 3 at the local Leaf nodes
- Leaf nodes Nexus 7000 switches perform inter-VLAN routing in FabricPath VDC sending traffic between source and destination servers

© 2013 Cisco and/or its affiliates. All rights reserved.

Interconnecting Global and Local VLANs

Routing Between Leaf and Spine

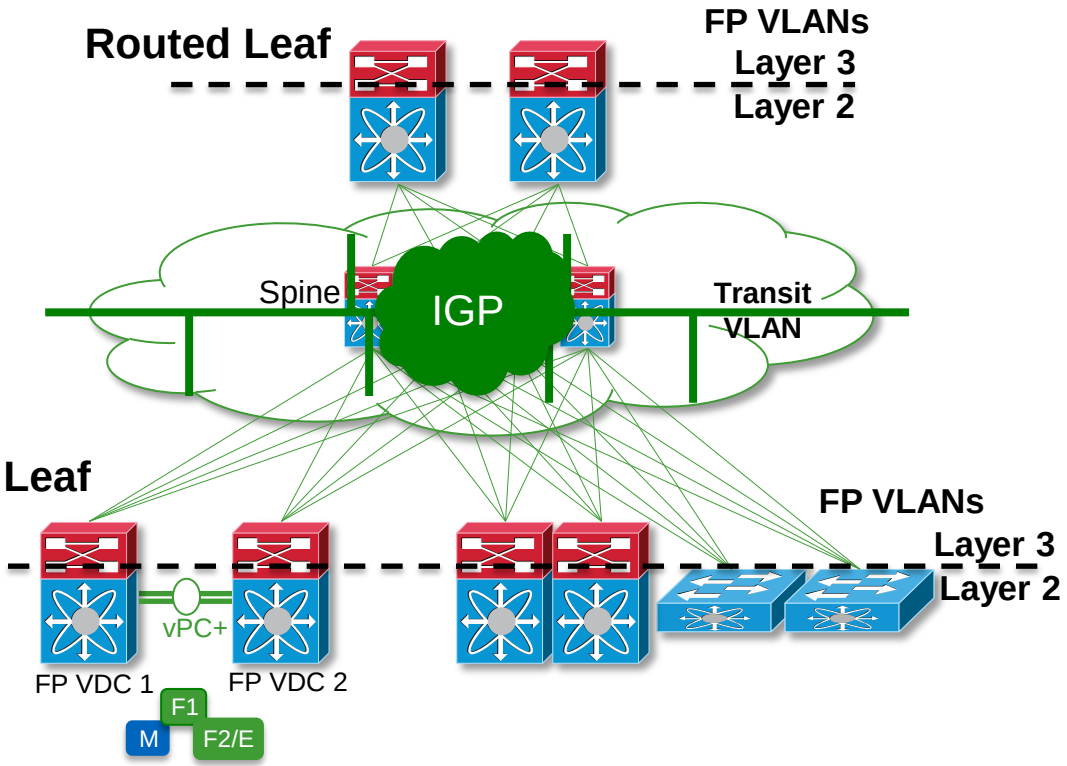

- What if VLANs extended over the FabricPath network (L3 at Routed Leaf nodes) need to communicate with VLANs that aren't (L3 at local Leaf nodes)
  - Inter-VLAN routing
- Special **Transit VLAN** is used to route traffic between Routed Leaf nodes and local Leaf nodes
- Routed Leaf nodes and local Leaf nodes establish routing protocol adjacency across the Transit VLAN

© 2013 Cisco and/or its affiliates. All rights reserved.

Interconnecting Global and Local VLANs

Routing Between Leaf and Spine

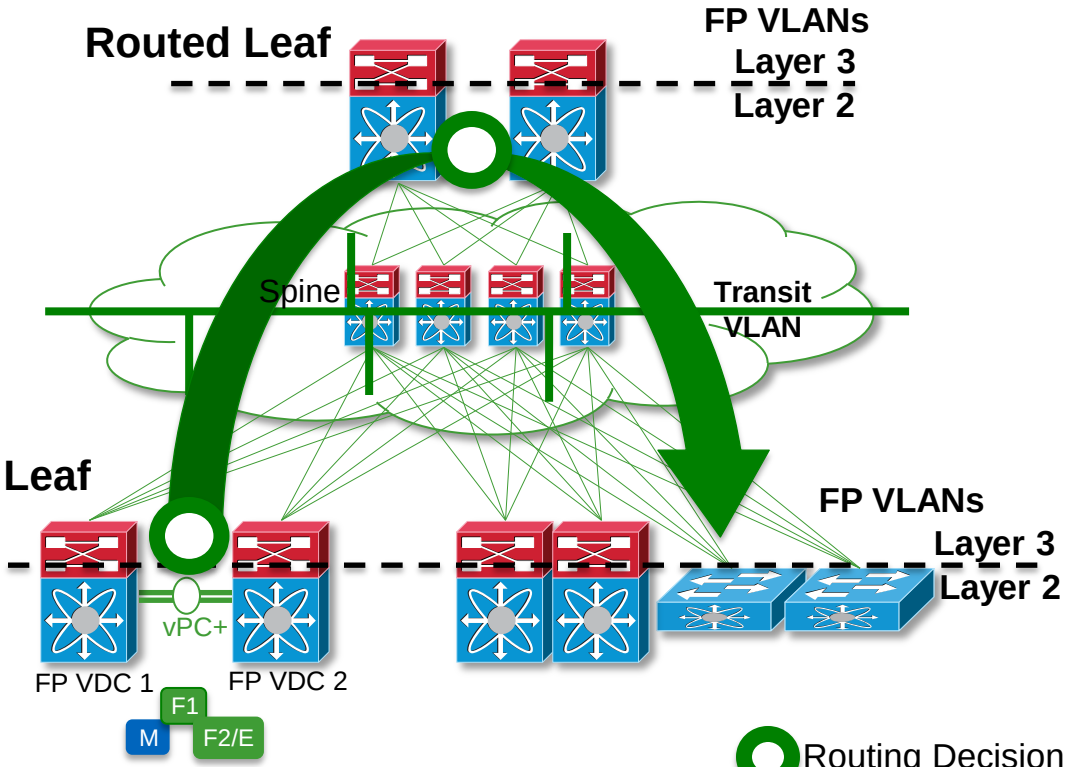

- Layer 3 local Leaf nodes make routing decision sending the traffic towards the Routed Leaf nodes across the Transit VLAN
- Routed Leaf nodes make routing decision sending the traffic towards the Layer 2 destination Leaf nodes
- Destination Leaf nodes forward the traffic to the destination servers
  - No routing decisions are made on destination Leaf nodes

© 2013 Cisco and/or its affiliates. All rights reserved.

Interconnecting Global and Local VLANs

Routing Between Leaf and Spine

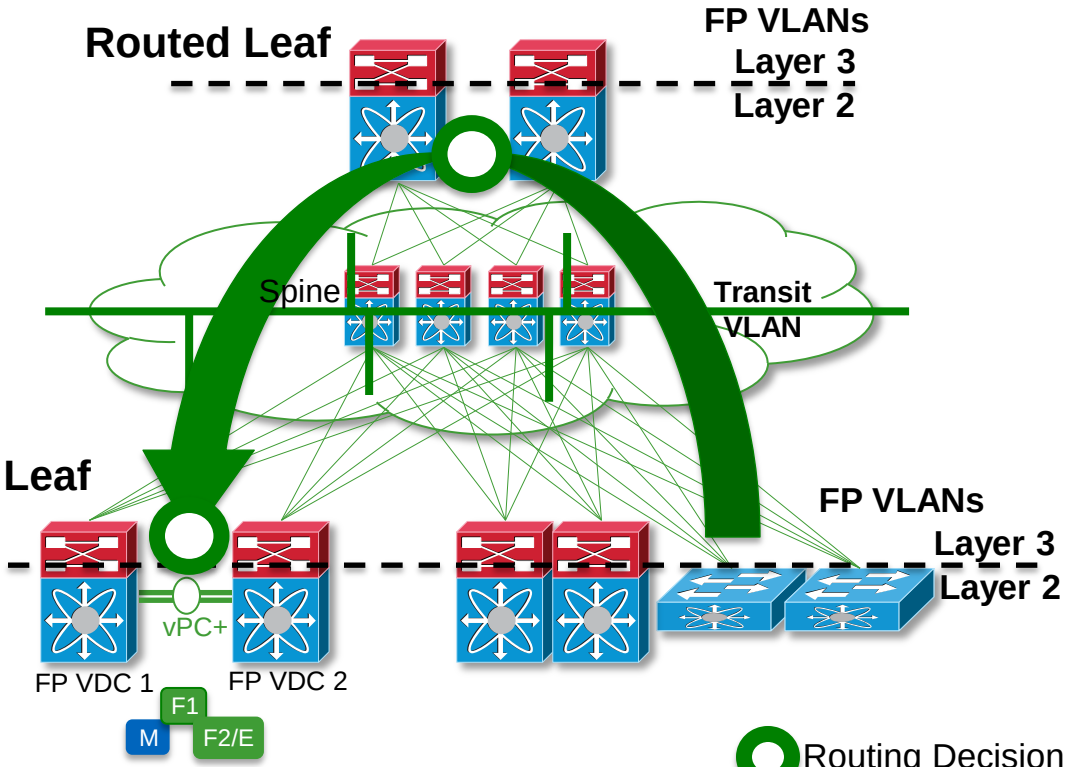

- In the opposite direction local Leaf nodes send the Global VLAN traffic towards the Routed Layer nodes
- Routed Leaf nodes make routing decision sending the traffic towards the Layer 3 destination Leaf nodes across Transit VLAN
- Destination Leaf nodes route the traffic to the destination servers

© 2013 Cisco and/or its affiliates. All rights reserved.

Interconnecting Local VLANs
Routing Between Leaf Nodes

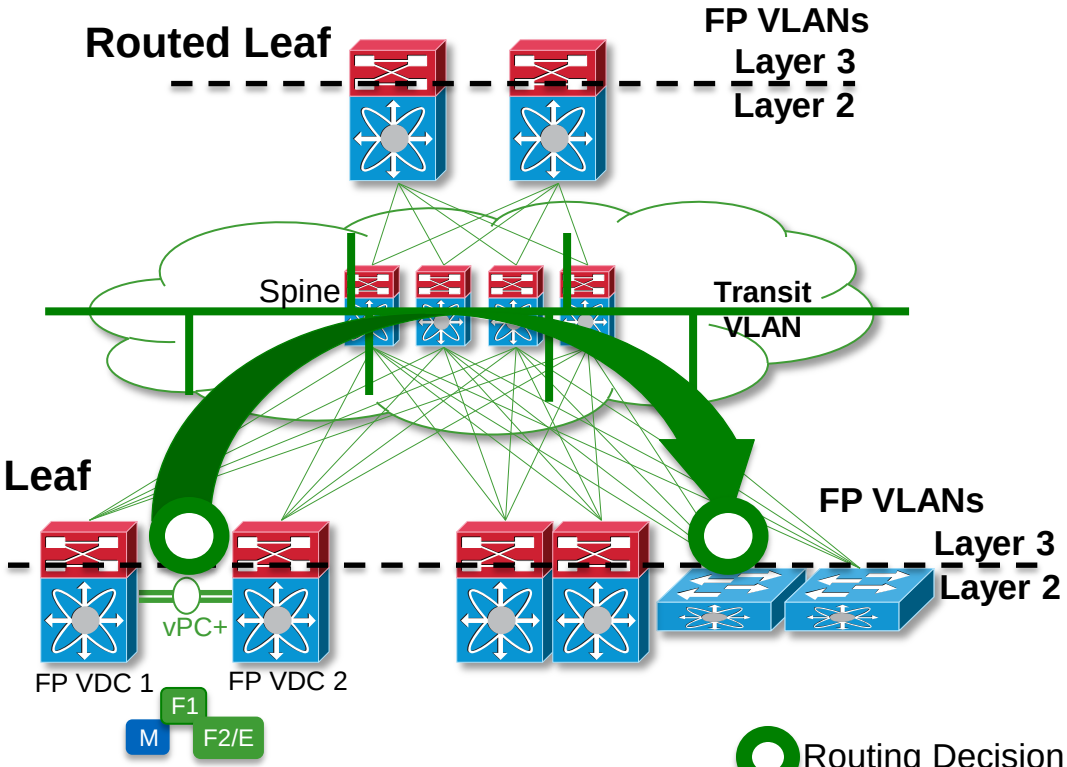

- Local to Local VLANs on a different pairs of Leaf Nodes communicate across the Transit VLAN
  - IGP adjacency is established over the Transit VLAN
- Traffic does not cross Routed Leaf nodes

© 2013 Cisco and/or its affiliates. All rights reserved.

FEX and FabricPath on Nexus 7000 switches
Progress Checkpoint

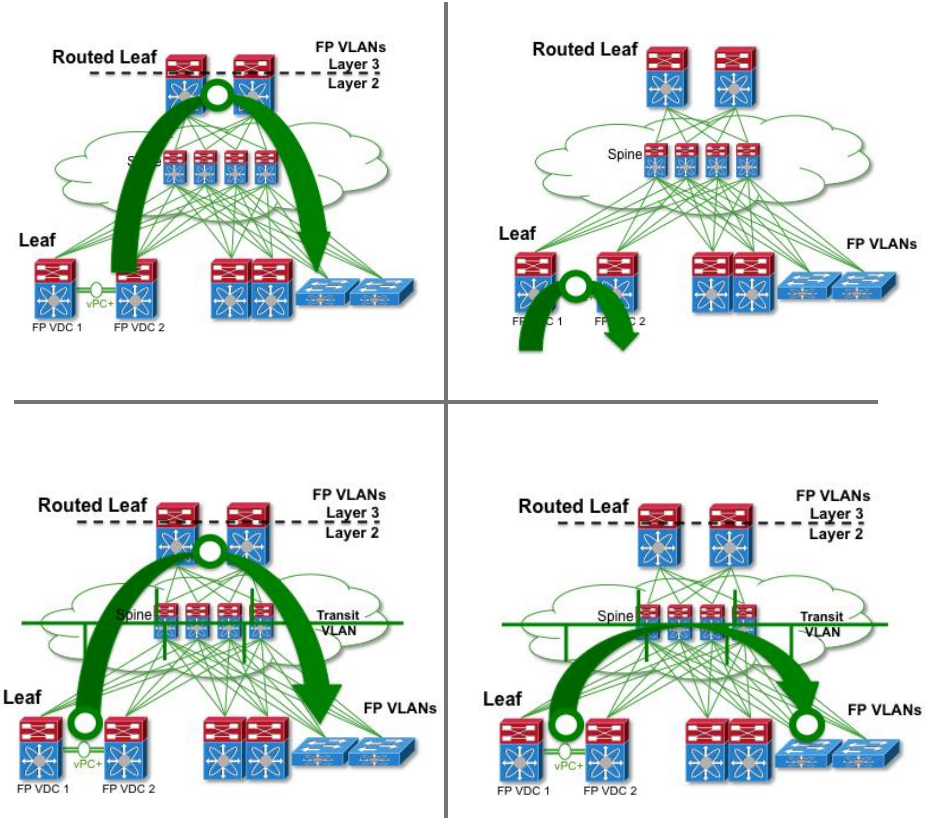

- All design conditions had been satisfied!

  ✓ Servers connected to Fabric Extenders can communicate across FabricPath network

  ✓ Global VLANs can be forwarded across FabricPath network

  ✓ Global and Local VLANs can communicate with each other

© 2013 Cisco and/or its affiliates. All rights reserved.

Discussion Agenda

- Fabric Extenders attachment to Nexus 7000 switches
- Fabric Extenders and Cisco FabricPath on Nexus 7000 switches
- **Cisco FabricPath and MPLS/LISP on Nexus 7000 switches**
- Multihop FCoE on Nexus 7000 switches

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco FabricPath and MPLS/LISP on Nexus 7000

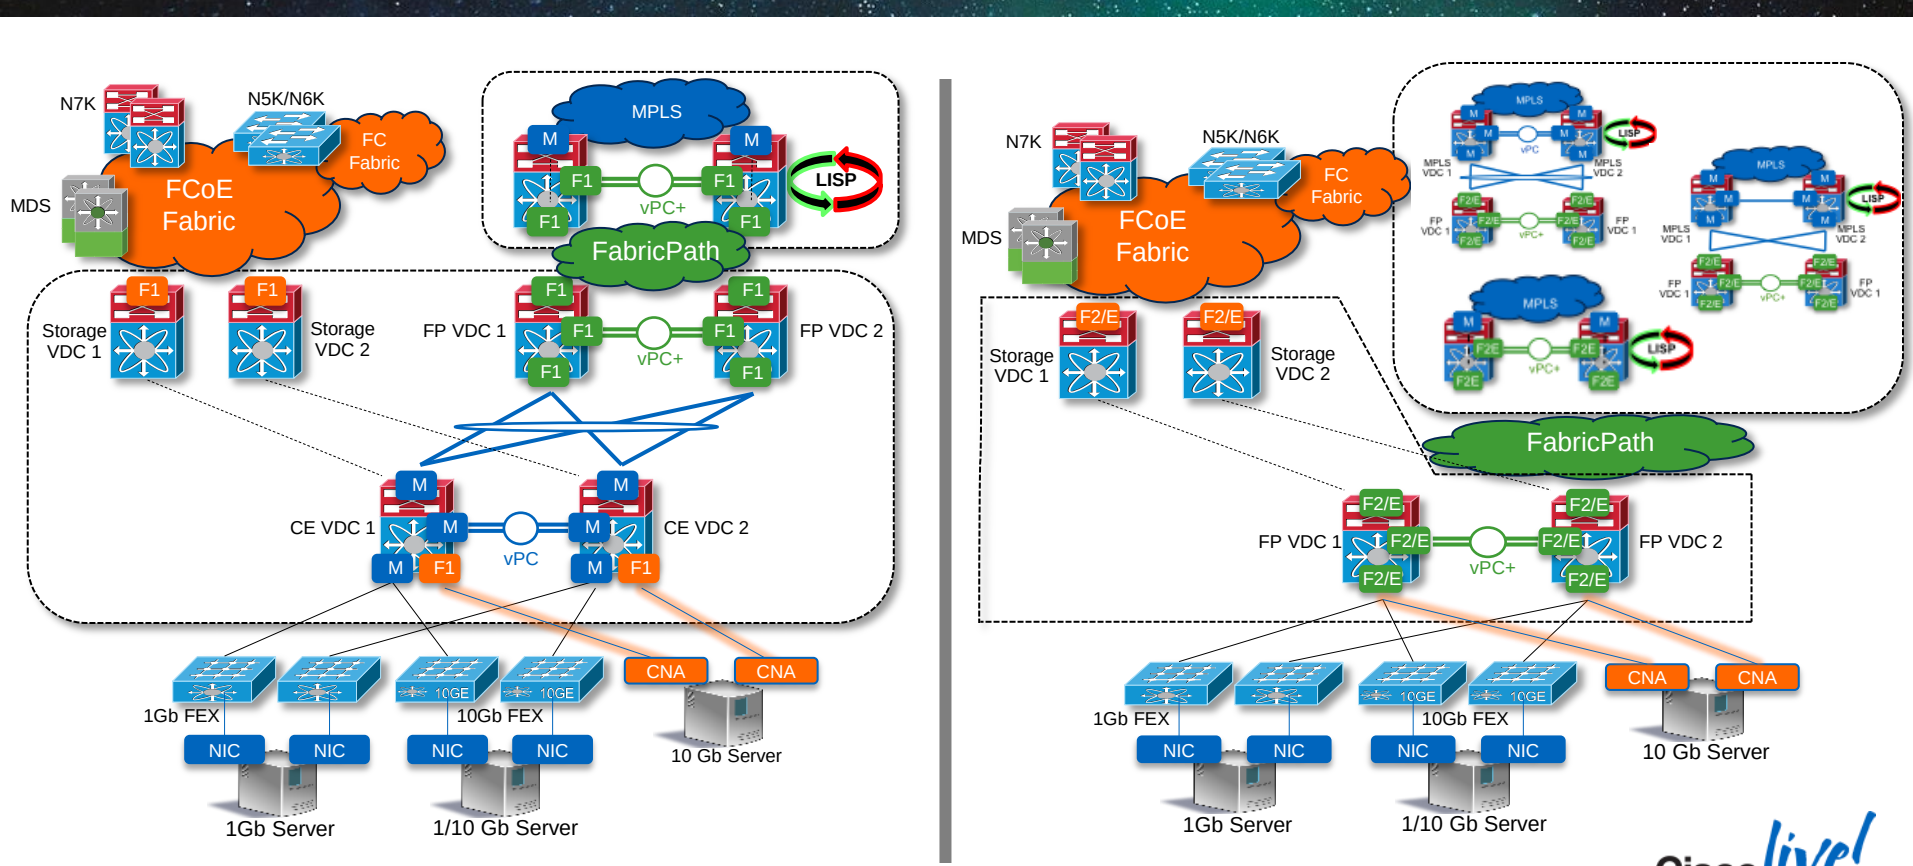

© 2013 Cisco and/or its affiliates. All rights reserved.

Why MPLS

What's important in our case:

- Provides convenient and proven connectivity method beyond single Data Center fabric
- Maintains end-to-end multi-tenancy properties and traffic segregation/segmentation
- Supported on the Cisco Nexus 7000 switches utilized in our design

© 2013 Cisco and/or its affiliates. All rights reserved.

FabricPath and MPLS

Option1: Two Switches/VDCs with Layer 3 Interconnect

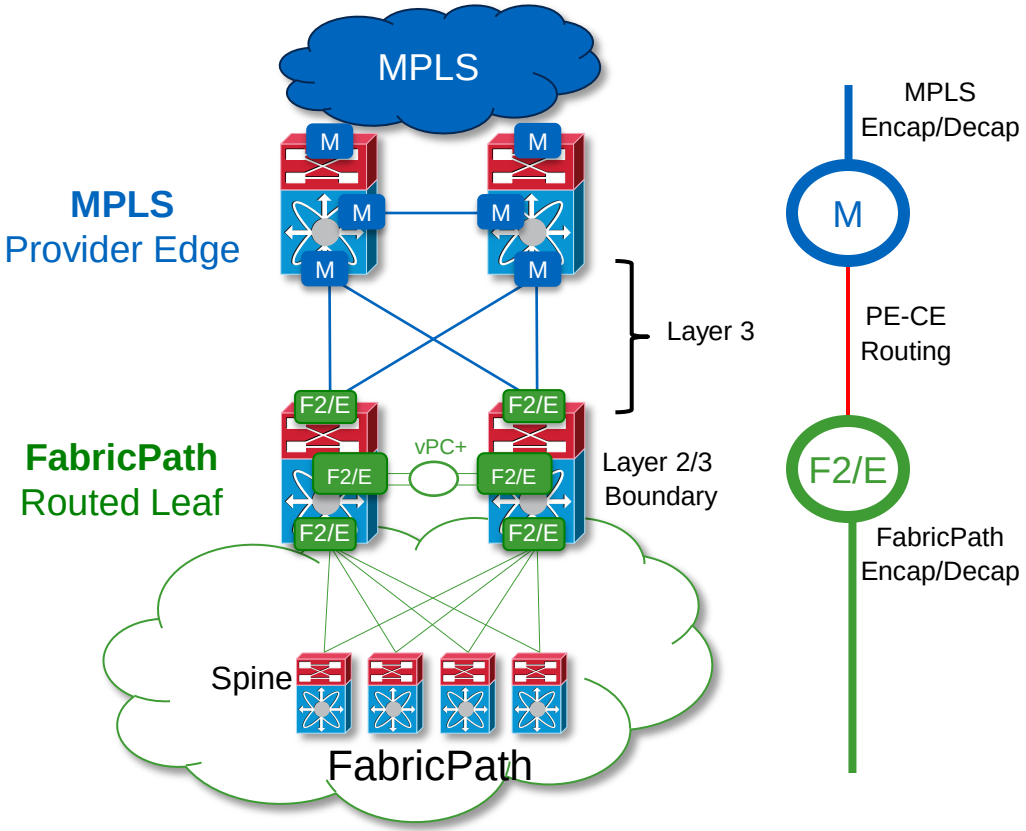

- MPLS PE and FabricPath Routed Leaf are delivered on two separate Cisco Nexus 7000 switches/VDCs
  - F2/F2E support FabricPath, but not MPLS
  - M-series support MPLS, but not FabricPath
- The interconnect is Classical Ethernet Layer 3 routing
  - Routing over point-to-point VLANs between FabricPath Routed Leaf nodes and MPLS PE nodes for multi-tenancy
  - Cannot leverage F1 modules on Routed Leaf, they do not support Layer 3 function
- Layer 2/Layer 3 boundary is on the FabricPath Routed Leaf nodes

© 2013 Cisco and/or its affiliates. All rights reserved.

FabricPath and MPLS Multitenancy

Option1: Two Switches/VDCs with Layer 3 Interconnect

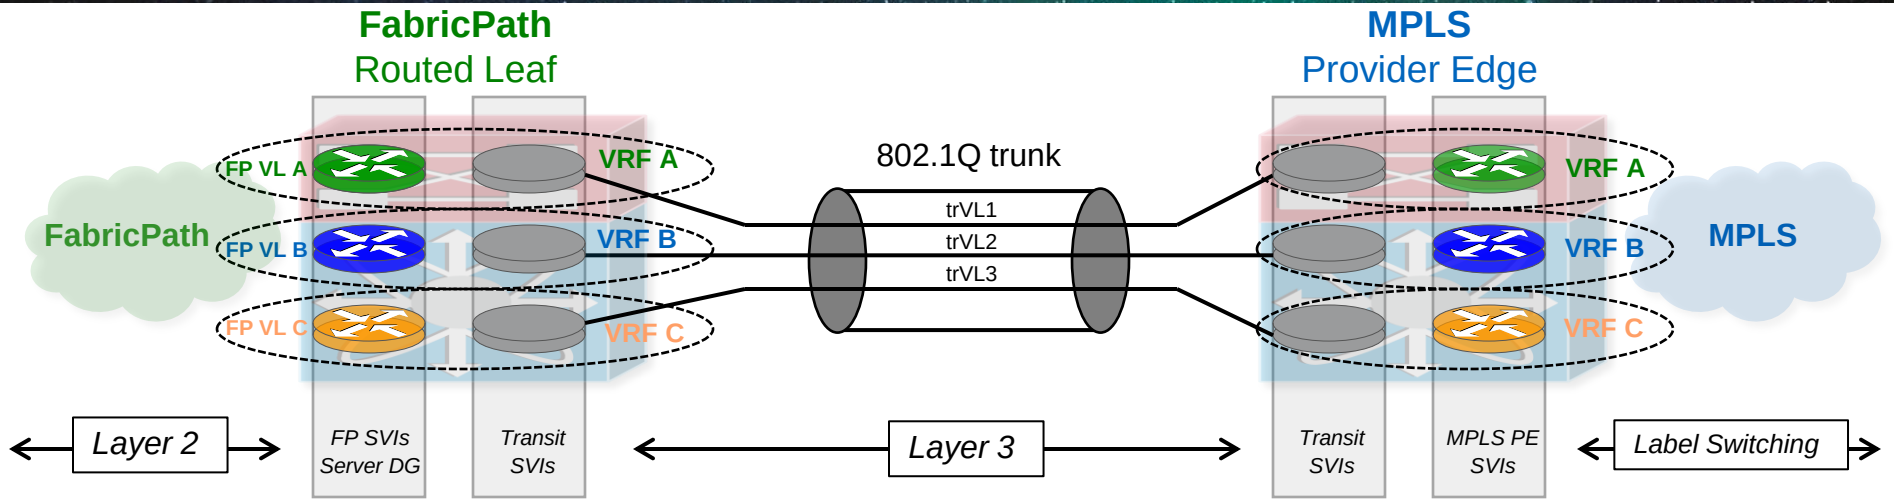

Multi-tenancy

- VRF-Lite between MPLS PE and FabricPath Routed Leaf
- *On FabricPath Routed Leaf:* FP VLAN SVIs and Transit VLAN SVI are mapped to the same VRF for any given tenant
- *On MPLS PE:* Transit VLAN SVI and MPLS PE SVIs are mapped to the same VRF for any given tenant
- FabricPath VLANs are extended across MPLS WAN

© 2013 Cisco and/or its affiliates. All rights reserved.

FabricPath and MPLS

Option2: Two Switches/VDCs with Layer 2 Interconnect

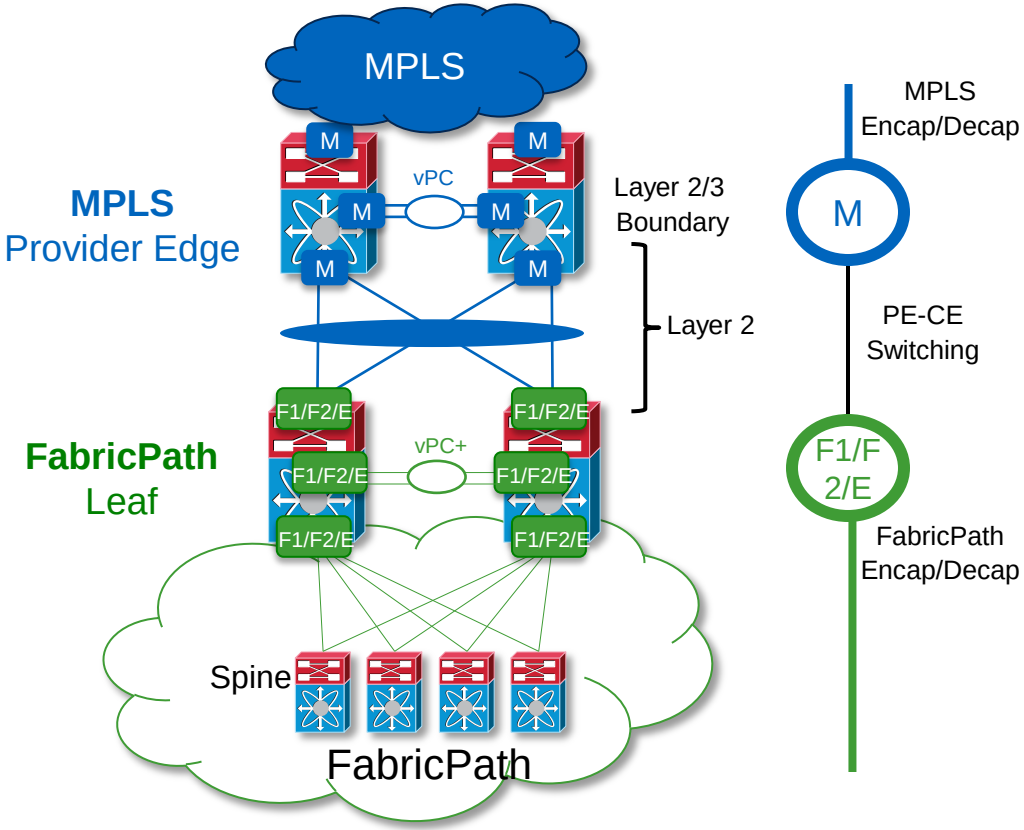

- MPLS PE and FabricPath Leaf are delivered on two separate Cisco Nexus 7000 switches/VDCs
  - F1/F2/F2E support FabricPath, but not MPLS
  - M-series support MPLS, but not FabricPath
- The interconnect is Classical Ethernet Layer 2 switching
  - FabricPath VLANs are extended between FabricPath Leaf nodes and MPLS PE nodes for multi-tenancy
  - Can leverage F1 modules as well
- Layer 2/Layer 3 boundary is on the MPLS PE nodes

© 2013 Cisco and/or its affiliates. All rights reserved.

FabricPath and MPLS Multitenancy

Option2: Two Switches/VDCs with Layer 2 Interconnect

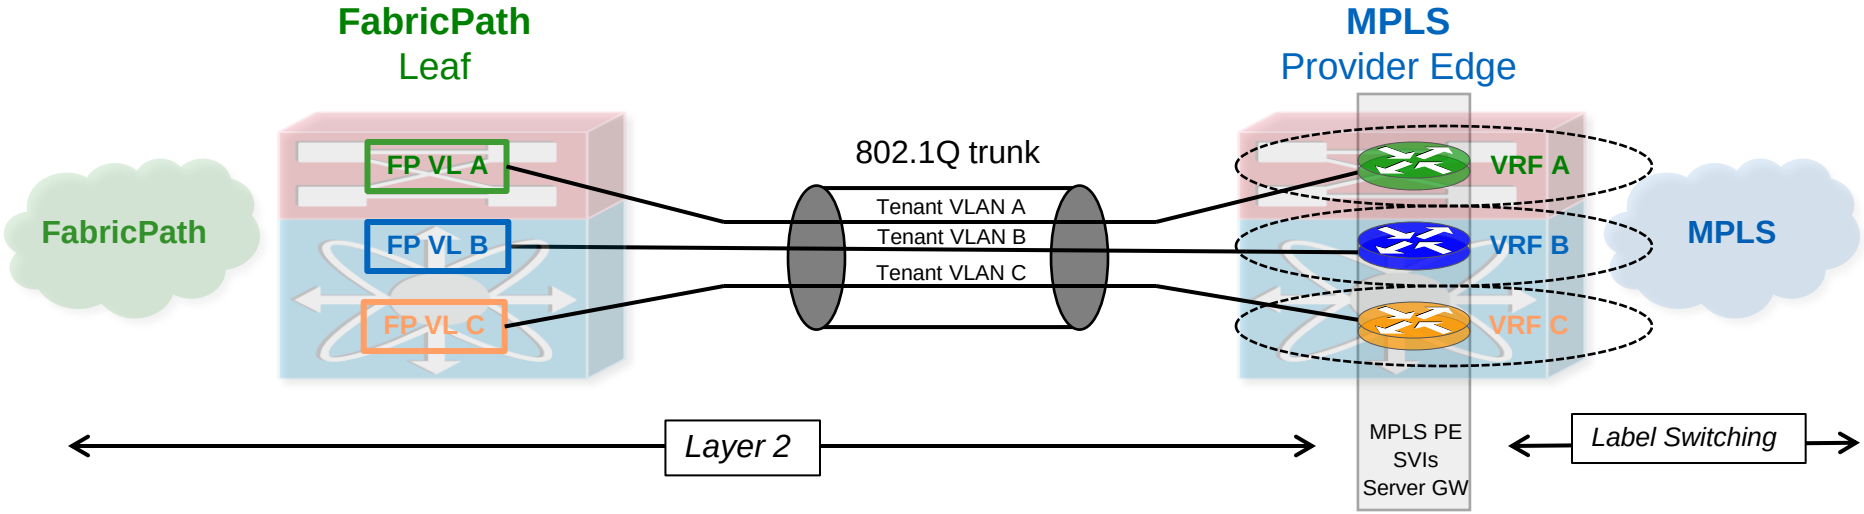

Multi-tenancy

- VLAN trunking between MPLS PE and FabricPath Routed Leaf
- *On FabricPath Leaf:* FP VLANs are mapped to any given tenant
- *On MPLS PE:* MPLS PE SVIs are mapped to VRF for any given tenant
- FabricPath VLANs are extended across MPLS WAN

© 2013 Cisco and/or its affiliates. All rights reserved.

Proxy-Routing

Cisco Nexus 7000 with F1/F2E and M-series I/O modules in the same VDC

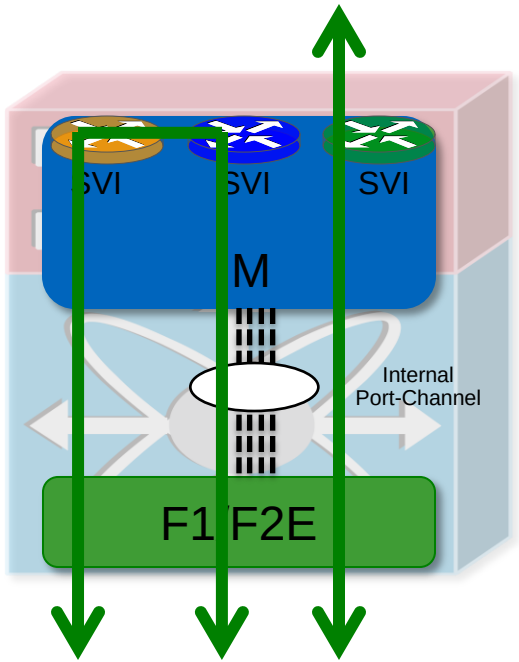

* F2E-M proxy-routing requires NX-OS 6.2

- F1 or F2E I/O modules are placed in the same VDC as the M-series I/O modules and Layer 3 (SVIs) is configured
  - F2 I/O modules cannot co-exists with any other module in the same VDC
- Internal Port-Channel is formed between the F1/F2E and M-series I/O modules
- Traffic received on the F1/F2E I/O modules requiring routing function is forwarded to the M-series I/O modules over the Internal Port-Channel links

© 2013 Cisco and/or its affiliates. All rights reserved.

FabricPath and MPLS

Option3: Single Switch/VDC with Proxy Routing

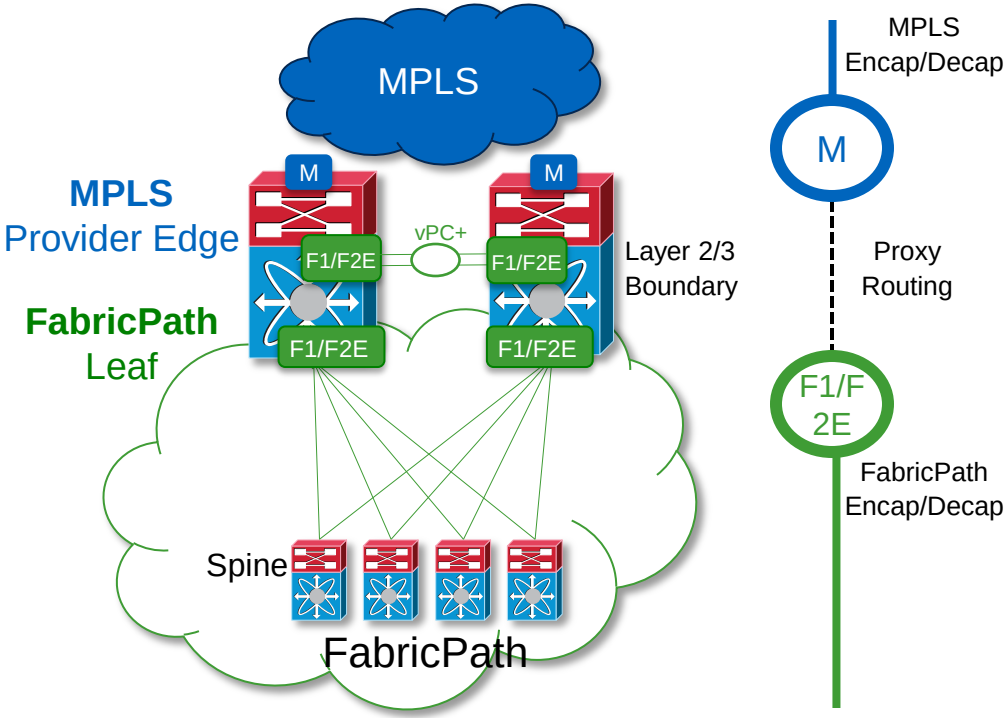

- MPLS PE and FabricPath Routed Leaf are delivered on the same Cisco Nexus 7000 switch/VDC
- Leverage proxy-routing between F1/F2E and M-series I/O modules
- FabricPath functionality is delivered on the F1/F2E I/O modules
- MPLS and Routing functionality is delivered on the M-series I/O module

© 2013 Cisco and/or its affiliates. All rights reserved.

FabricPath and MPLS Multitenancy

Option3: Single Switch/VDC with Proxy Routing

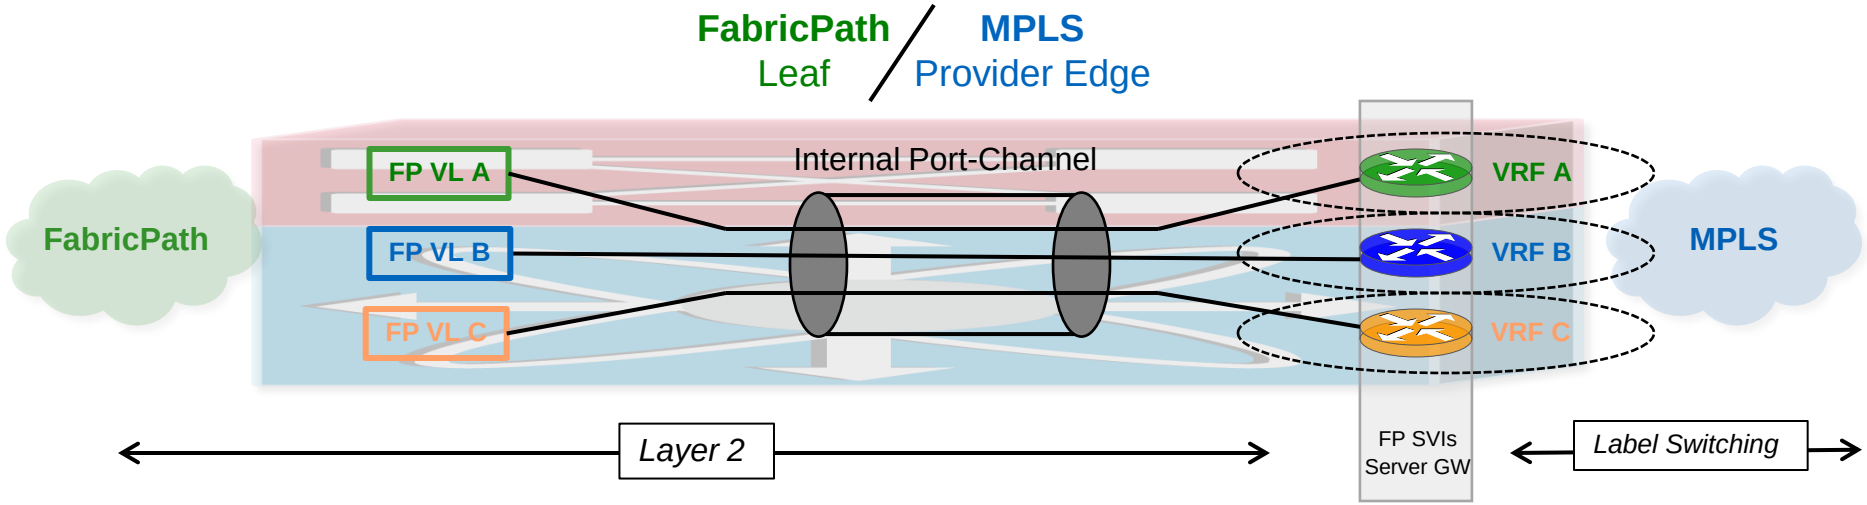

Multi-tenancy

- VLANs forwarded across the internal port-channel
- FP VLANs are mapped to any given tenant
- FP SVIs are mapped to VRF for any given tenant
- FabricPath VLANs are extended across MPLS WAN

© 2013 Cisco and/or its affiliates. All rights reserved.

Why LISP

What's important in our case:

- Provide Virtual Machine mobility aware routing beyond a single Data Center
- Supported on the Cisco Nexus 7000 switches utilized in our design

© 2013 Cisco and/or its affiliates. All rights reserved.

LISP Refresher
Traffic Forwarding

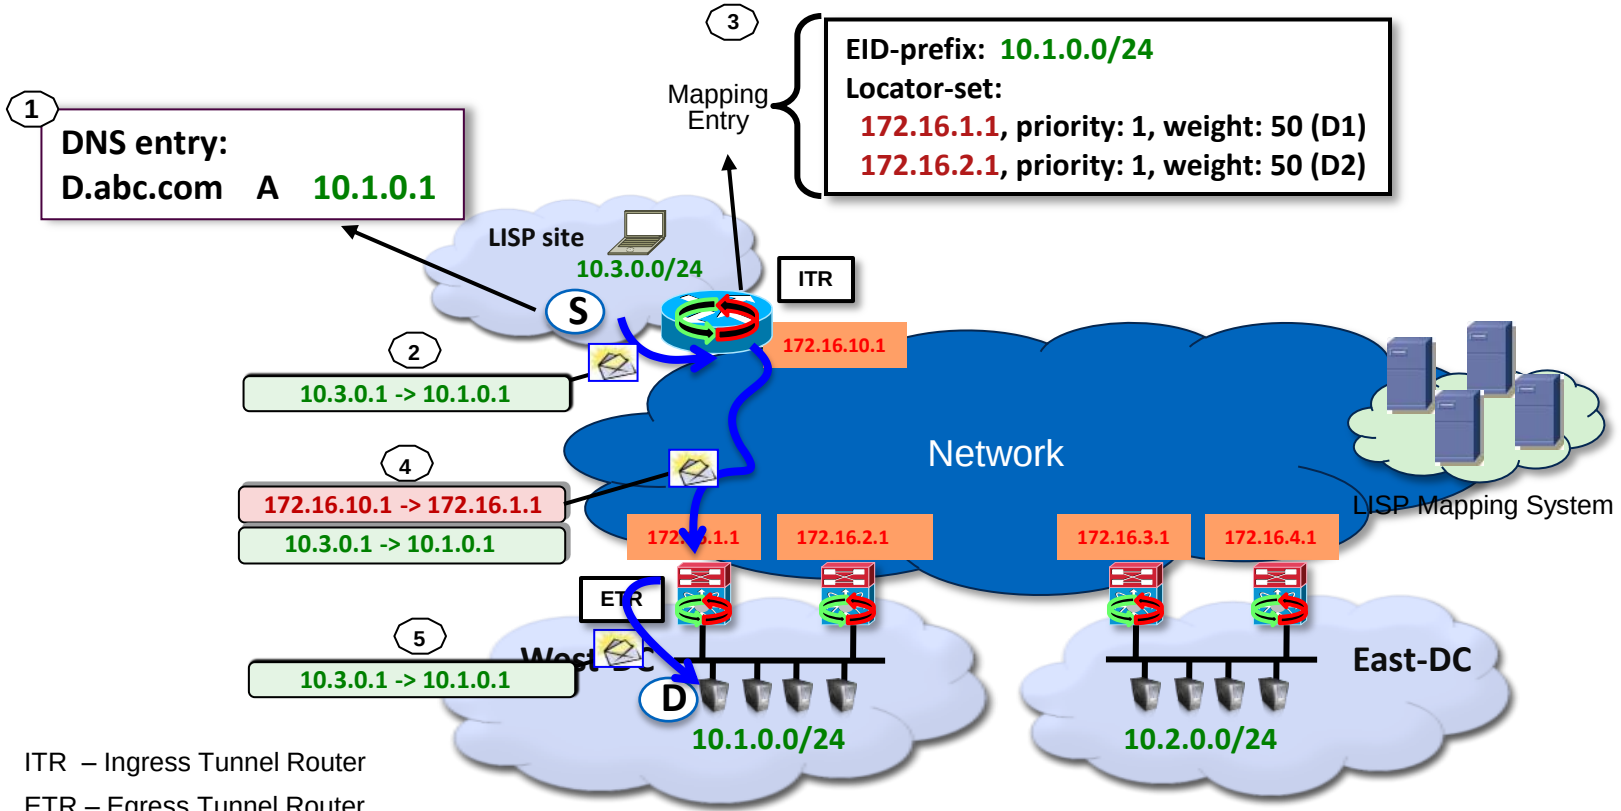

ITR – Ingress Tunnel Router

ETR – Egress Tunnel Router

© 2013 Cisco and/or its affiliates. All rights reserved.

LISP Host Mobility
Multi-Data Center

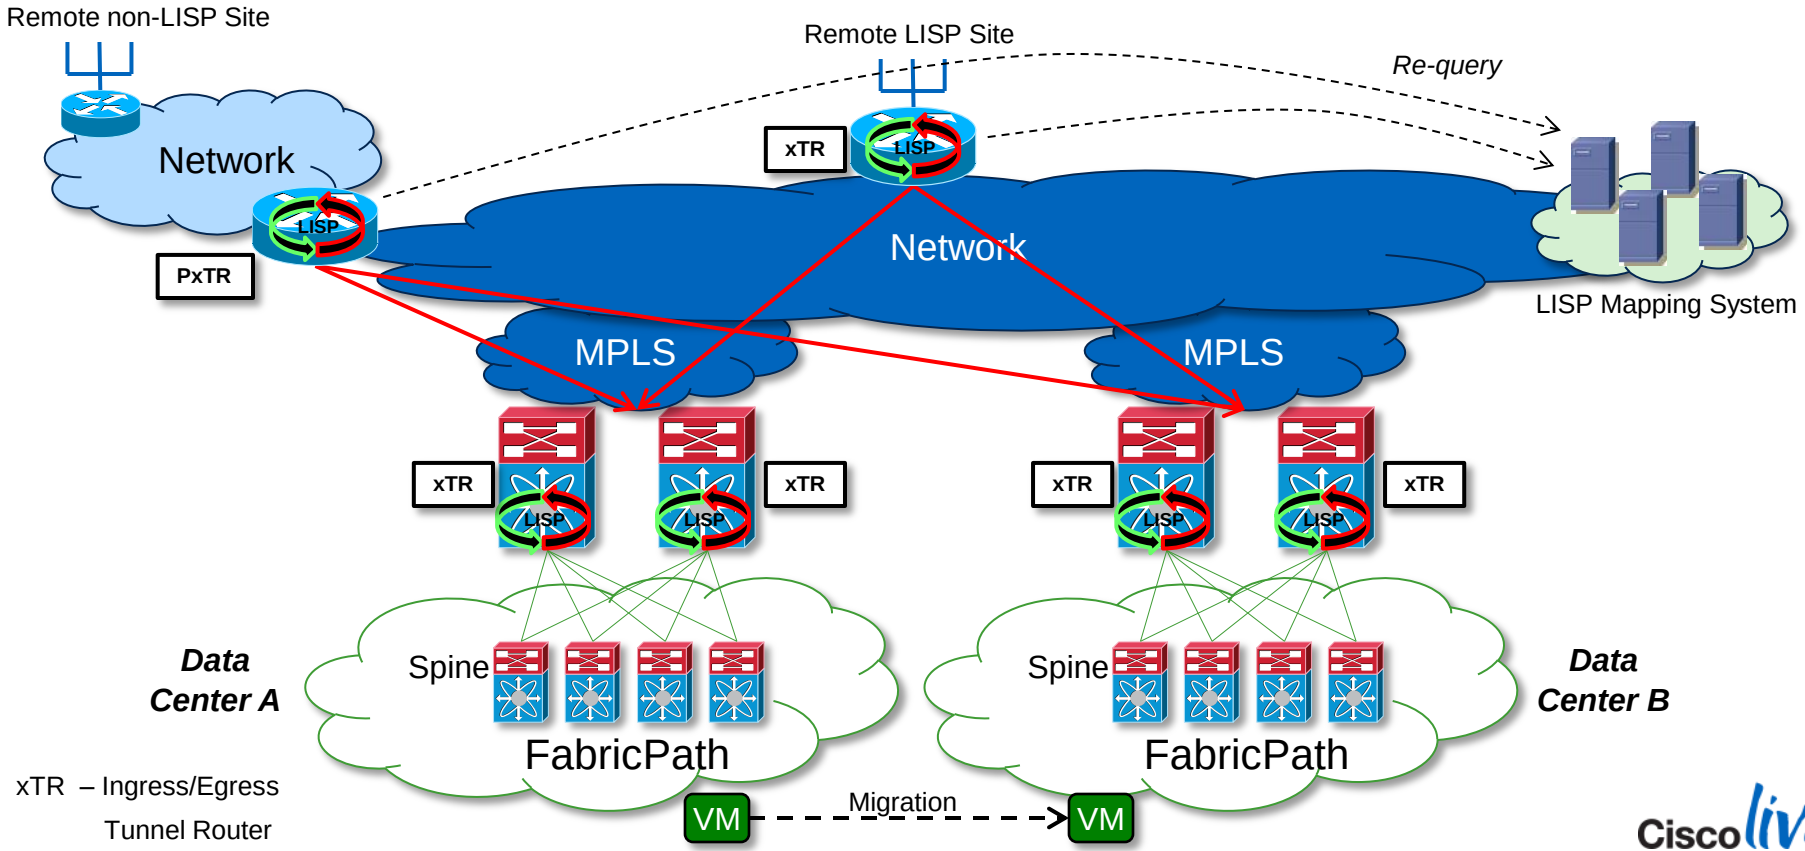

xTR – Ingress/Egress Tunnel Router

© 2013 Cisco and/or its affiliates. All rights reserved.

FabricPath, MPLS and LISP

Option1: Two Switches/VDCs with Layer 2 Interconnect

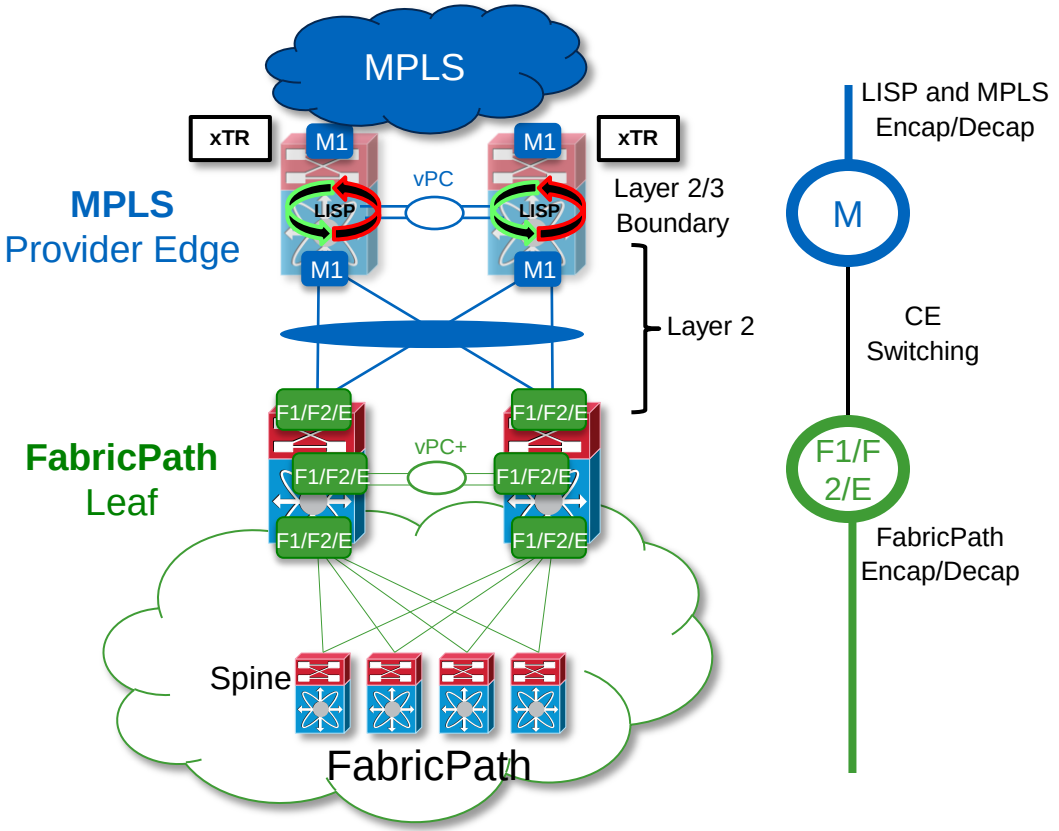

- MPLS PE and FabricPath Leaf are delivered on two separate Cisco Nexus 7000 switches/VDCs
  - F1/F2/F2E support FabricPath, but not MPLS/LISP
  - M-series support MPLS/LISP, but not FabricPath (requires N7K-M132XP-12L line card)
- LISP host mobility requires xTR to be L3 adjacent to the host
- LISP xTR is delivered on MPLS PE Nexus 7000 switch/VDC
- The interconnect is Classical Ethernet Layer 2 switching

© 2013 Cisco and/or its affiliates. All rights reserved.

FabricPath, MPLS and LISP

Option2: Single Switch/VDC with Proxy Routing

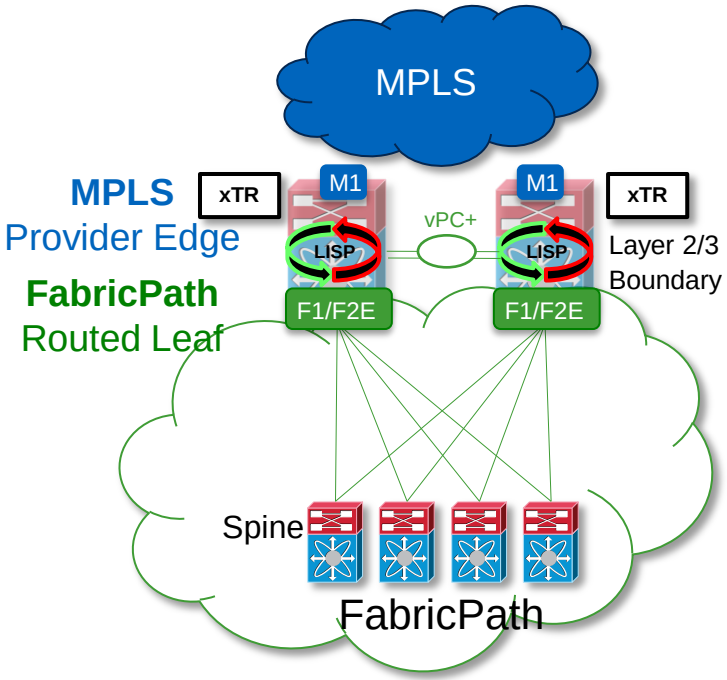

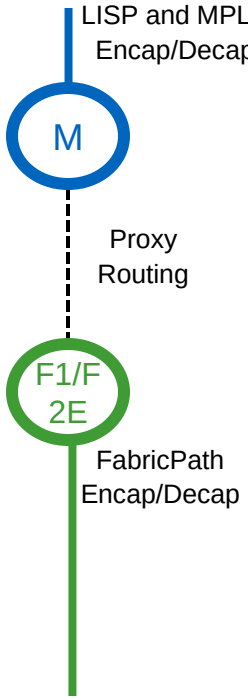

- FabricPath Leaf, MPLS PE and LISP xTR are delivered on a single Cisco Nexus 7000 switch/VDC
- Leverage proxy-routing between F1/F2E and M-series I/O modules
- FabricPath functionality is delivered on the F1/F2E I/O modules
- LISP, MPLS and Routing functionality is delivered on the M-series I/O module
  - LISP requires N7K-M132XP-12L line card

© 2013 Cisco and/or its affiliates. All rights reserved.

FabricPath, MPLS and LISP
Progress Checkpoint

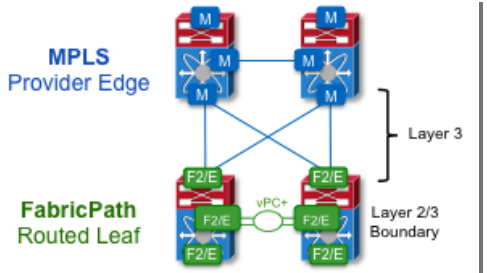

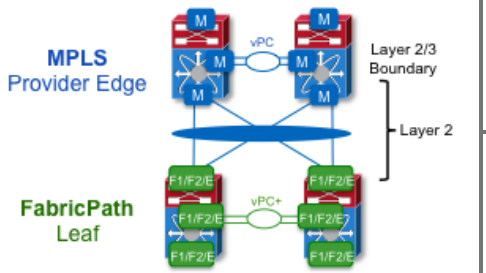

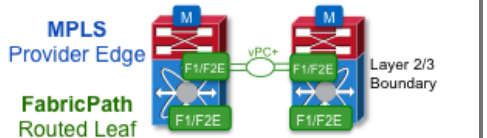

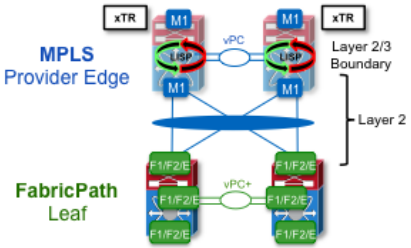

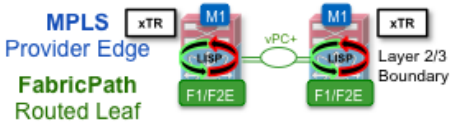

- All design conditions had been satisfied!
  - ✓ Connectivity is extended beyond a single Data Center
  - ✓ Multi-tenancy segregation is preserved
  - ✓ Routing is dynamically adjusted for Virtual Machine mobility

© 2013 Cisco and/or its affiliates. All rights reserved.

Discussion Agenda

- Fabric Extenders attachment to Nexus 7000 switches
- Fabric Extenders and Cisco FabricPath on Nexus 7000 switches
- Cisco FabricPath and MPLS/LISP on Nexus 7000 switches
- **Multihop FCoE on Nexus 7000 switches**

© 2013 Cisco and/or its affiliates. All rights reserved.

MultiHop FCoE on Cisco Nexus 7000

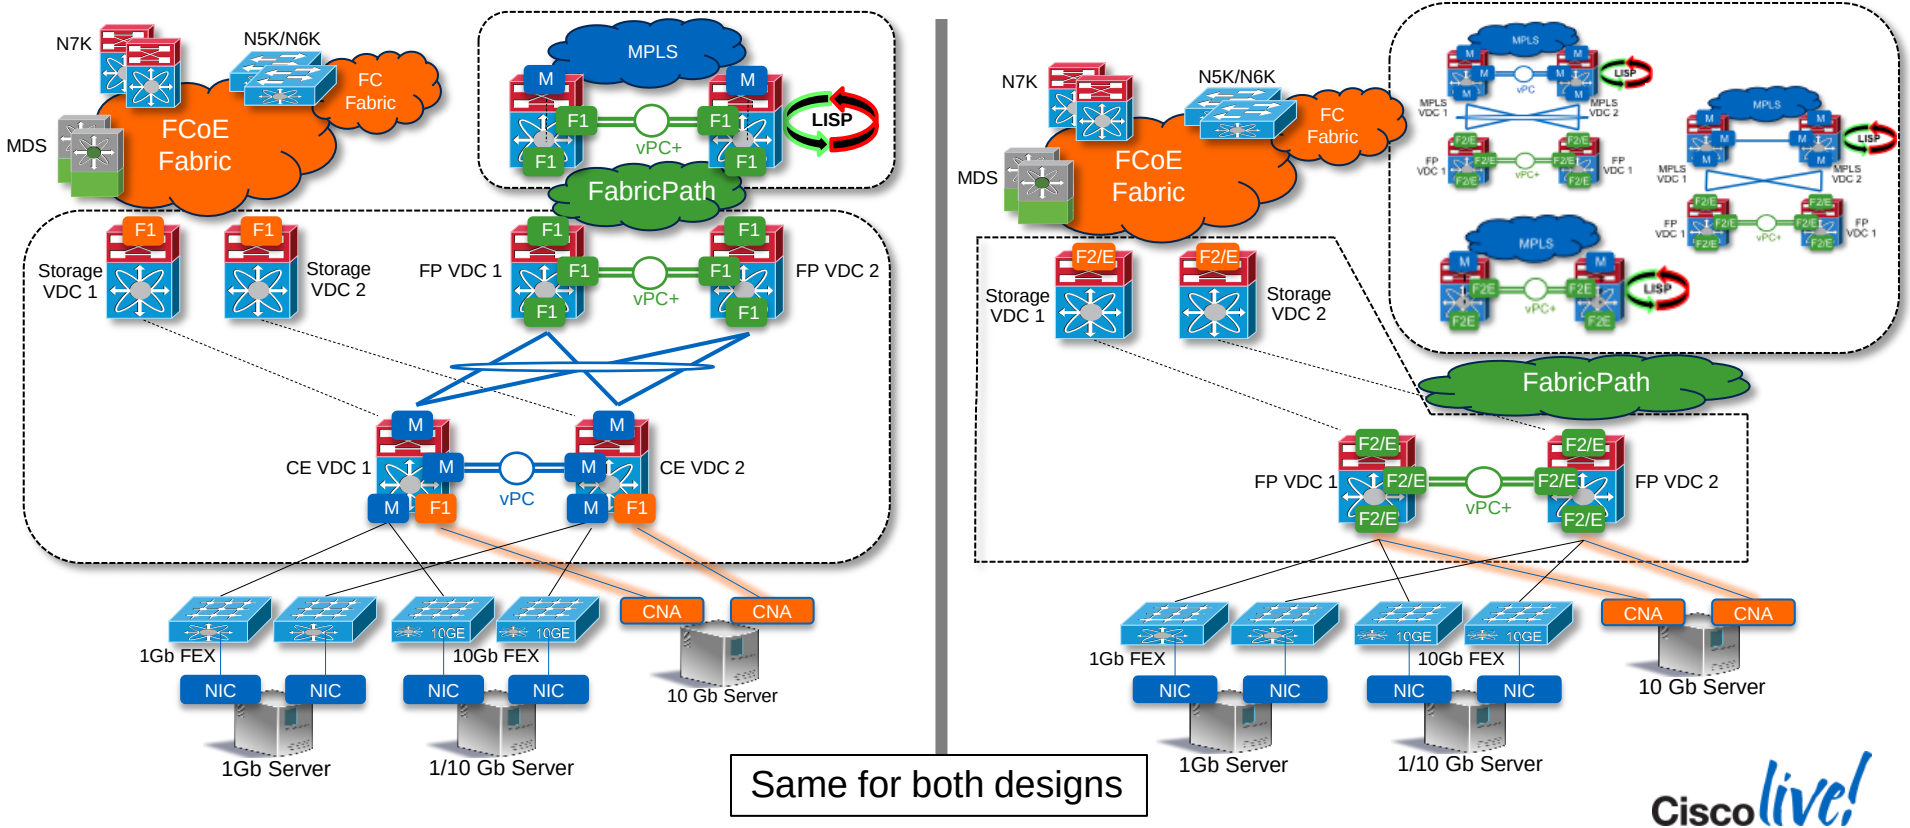

© 2013 Cisco and/or its affiliates. All rights reserved.

Storage Transport Technologies Landscape

FCoE and DCB

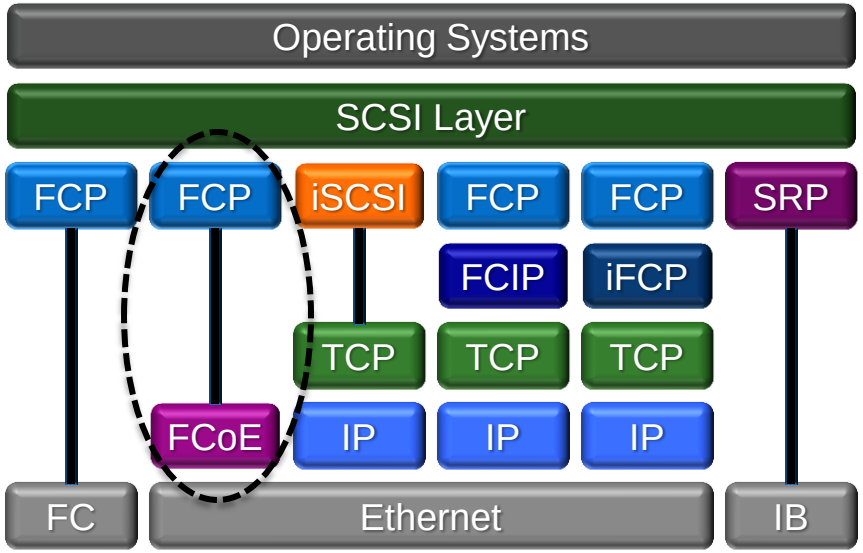

- Our solution is based on mapping of Fibre Channel frames over Ethernet transport, aka FCoE
- FCoE must run on a lossless Ethernet fabric, aka **Data Center Bridging**

Data Center Bridging

Priority Flow Control 802.1Qbb

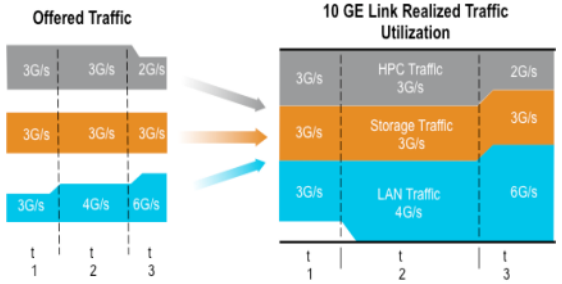

Enhanced Transmission Selection 802.1Qaz

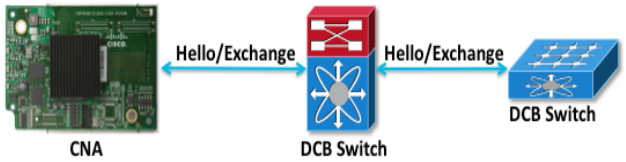

Data Center Bridging Exchange

© 2013 Cisco and/or its affiliates. All rights reserved.
Cisco Public

FCoE Protocol Organization

Control and Data Plane

- FCoE Initialization Protocol
  - FCoE VLAN Discovery
  - Fibre Channel Forwarder Discovery
  - Fabric Login
- FC Commands

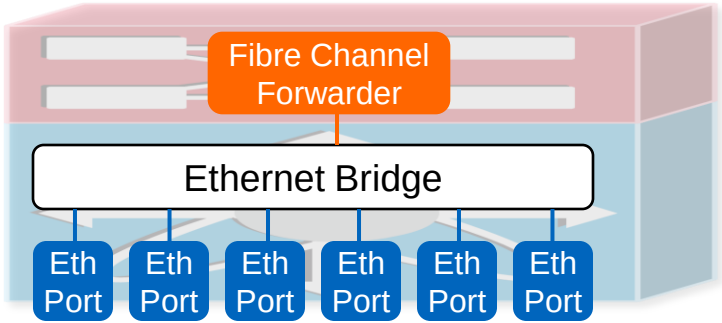

- FCF is a logical Fibre Channel switch inside an FCoE (Ethernet) switch
- Fibre Channel login happens at FCF
- FCoE encapsulation/decapsulation happens at FCF

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Nexus 7000 Fibre Channel Forwarder

Storage VDC

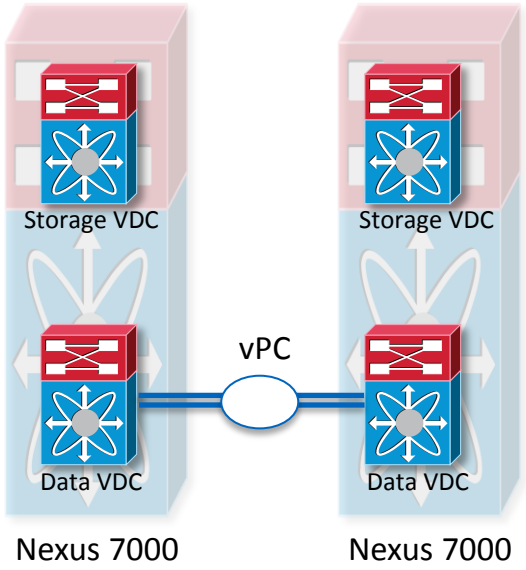

- FCF on Cisco Nexus 7000 is delivered in a form of Storage VDC
  - Requires Storage License
- Storage VDC acts as Director Class SAN switch
- All FCoE related configurations are carried out in Storage VDC
- Storage VDC is counted against total number of VDCs currently supported on Cisco Nexus 7000 switches
  - Currently up to 8 VDCs with Sup2E
  - Single Storage VDC per N7K

© 2013 Cisco and/or its affiliates. All rights reserved.

FCoE Server Connectivity
Unified Wire and Shared Interfaces

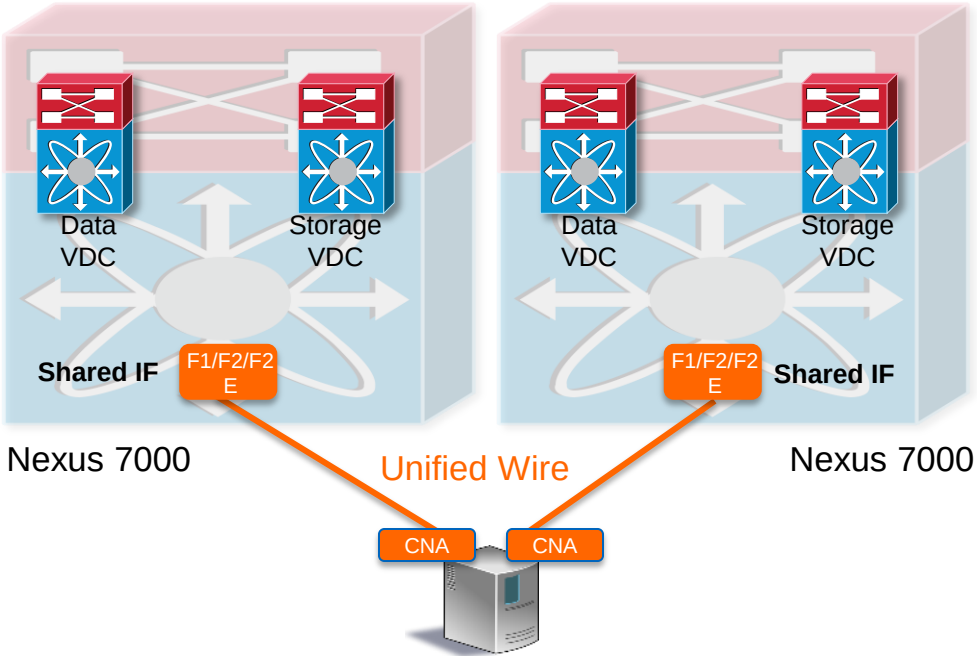

- Servers utilize the concept of Unified Wire
- Unified Wire plugs directly into an F1/F2/F2E interface, which can accommodate both network data and FCoE traffic
  - Support for Fabric Extenders is upcoming
- Interfaces are *shared* between Data and Storage VDC
  - Shared Interface

```
N7K(config)# vdc fcoe type storage
N7K(config-vdc)# allocate shared interface e2/1
```


© 2013 Cisco and/or its affiliates. All rights reserved.

Server Traffic Over Unified Wire

Traffic Forwarding over Shared Interfaces

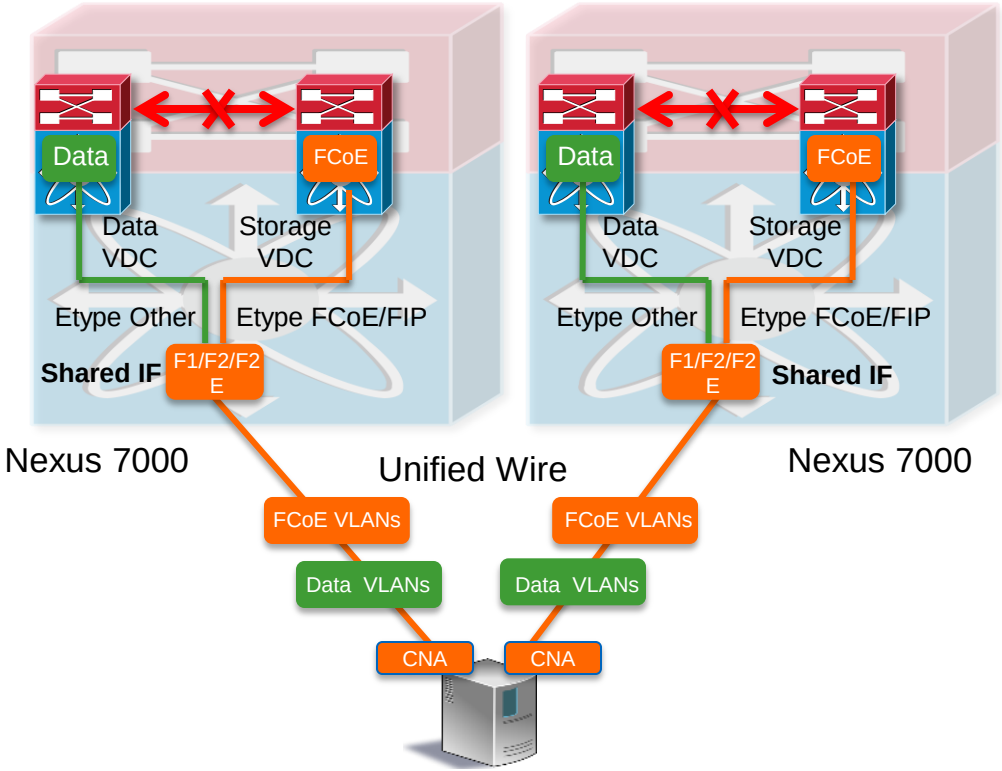

- Shared Ports utilize 802.1Q trunking to carry both network data and FCoE VLANs
- Traffic with FCoE and FIP ethertypes received over shared ports appears in Storage VDC for processing
- Traffic with other ethertype values appears in Data VDC
- *There is no communication between data VDC and Storage VDC!*

© 2013 Cisco and/or its affiliates. All rights reserved.

Server Traffic Over Inter-Switch Links and Uplinks

Dedicated Wire Traffic Forwarding

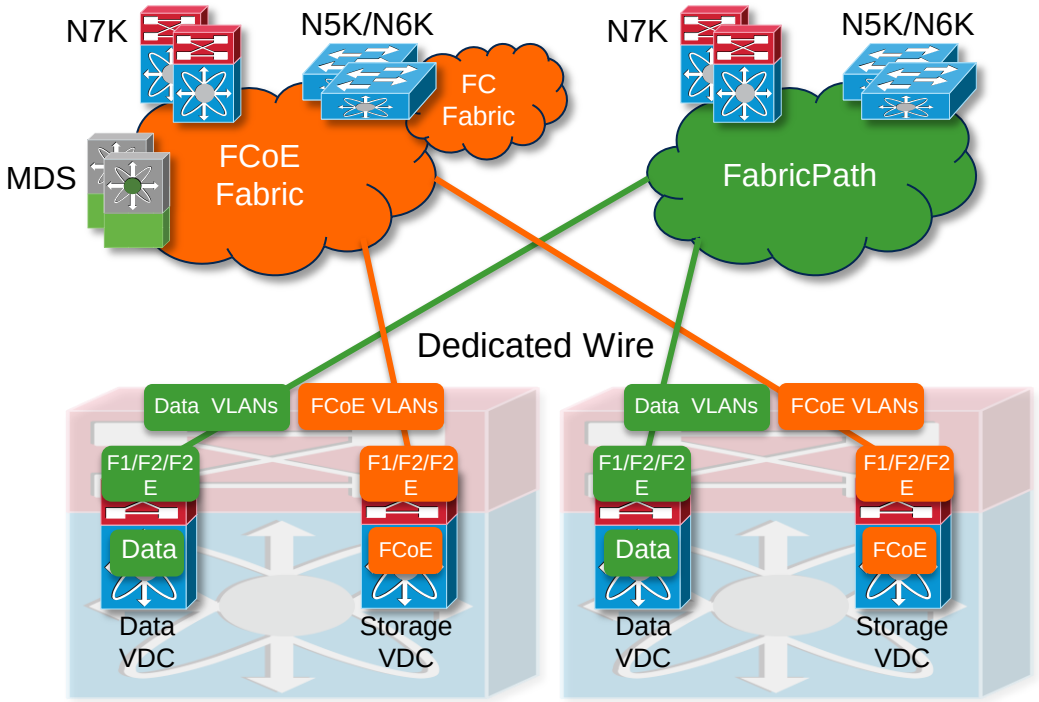

- Inter-Switch Links use Dedicated Wire where network data and FCoE traffic are carried over separate physical links
- Data VLANs are forwarded towards the FabricPath network
- FCoE VLANs are forwarded towards the FCoE fabric

© 2013 Cisco and/or its affiliates. All rights reserved.

End-to-End Unified Fabric

Inter-Switch Links and Multihop FCoE

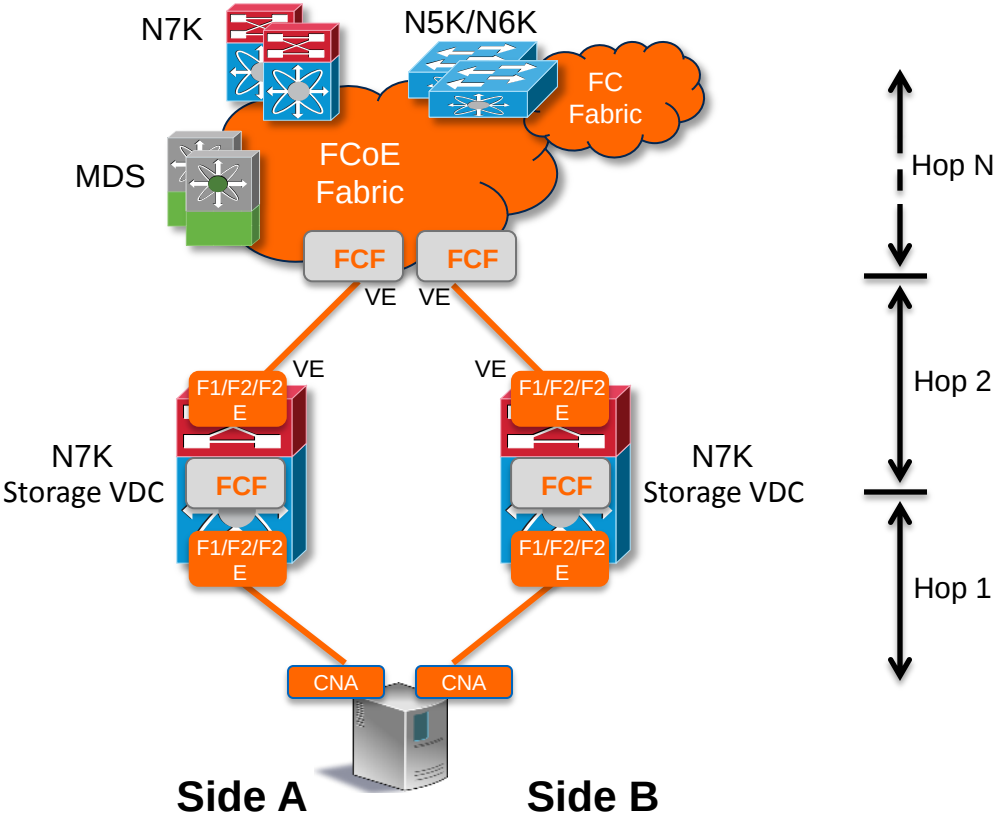

- Nexus 7000 can operate only in FC switching mode
  - Cannot be in NPV mode
  - Need to assign FC Domain ID and participate in fabric zoning
- Consider the following alternative

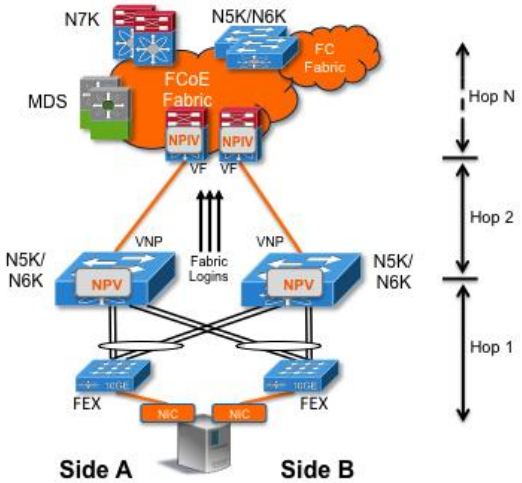

Nexus 5K/6K

- NPV
- No Domain ID
- No Zoning
- Leverage FEX

© 2013 Cisco and/or its affiliates. All rights reserved.

Multihop FCoE on Nexus 7000 switches
Progress Checkpoint… Last one ☺

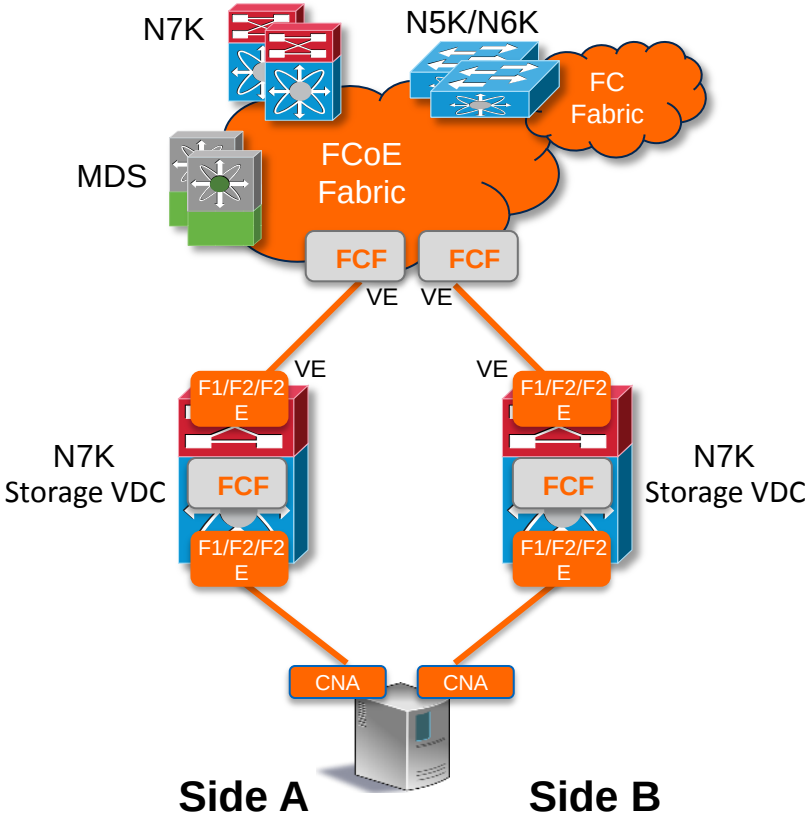

- All design conditions had been satisfied!
  - ✓ FCoE Servers are connected to Cisco Nexus 7000 switches
  - ✓ Cisco Nexus 7000 switches use device virtualization in a form of Storage VDC for all FCoE functionality
  - ✓ Multihop FCoE carries storage traffic across the Unified Fabric

© 2013 Cisco and/or its affiliates. All rights reserved.

Key Takeaways

- The use of Cisco Nexus 7000 device virtualization techniques allows deploying **diverse features in consolidated physical topology**
- Cisco Nexus 7000 provide **investment protection** by enabling innovative features on existing I/O modules
- Cisco Nexus family of products **lay a solid foundation** for the next generation Data Centers

© 2013 Cisco and/or its affiliates. All rights reserved.

Complete Your Online Session Evaluation

- Give us your feedback and you could win fabulous prizes. Winners announced daily.
- Receive 20 Cisco Daily Challenge points for each session evaluation you complete.
- Complete your session evaluation online now through either the mobile app or internet kiosk stations.

Maximize your Cisco Live experience with your free Cisco Live 365 account. Download session PDFs, view sessions on-demand and participate in live activities throughout the year. Click the Enter Cisco Live 365 button in your Cisco Live portal to log in.

© 2013 Cisco and/or its affiliates. All rights reserved.

Final Thoughts

- Get hands-on experience with the Walk-in Labs located in World of Solutions, booth 1042
- Come see demos of many key solutions and products in the main Cisco booth 2924
- Visit www.ciscoLive365.com after the event for updated PDFs, on-demand session videos, networking, and more!
- Follow Cisco Live! using social media:
  - Facebook: https://www.facebook.com/ciscoliveus
  - Twitter: https://twitter.com/#!/CiscoLive
  - LinkedIn Group: http://linkd.in/CiscoLI

© 2013 Cisco and/or its affiliates. All rights reserved.


CISCO
TM