
TOMORROW starts here.
cisco
Cisco live!

Scalable Midsize Data Centre Designs

BRKDCT-2218

James Gilarte

Virtual Systems Engineer

Cisco live!

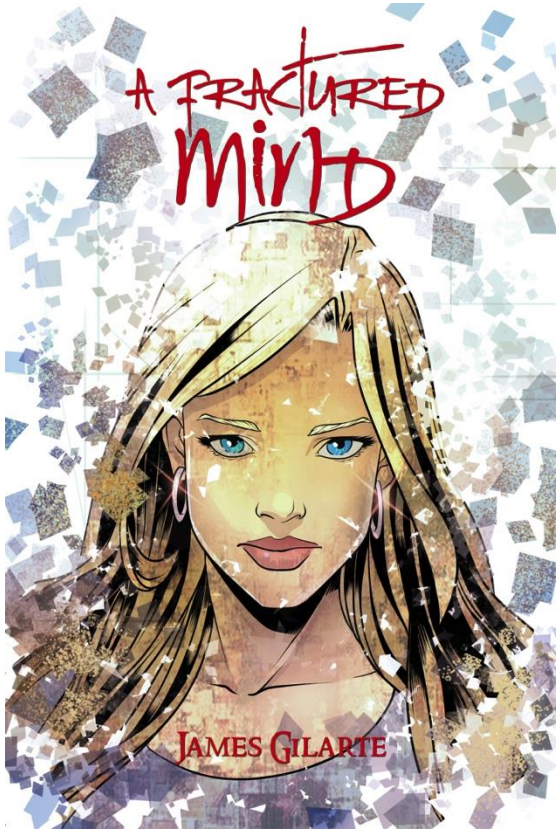

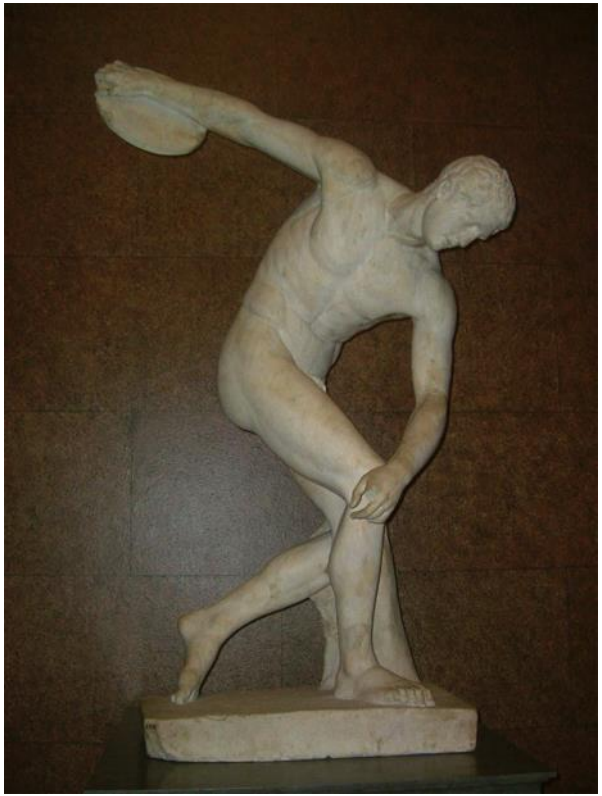

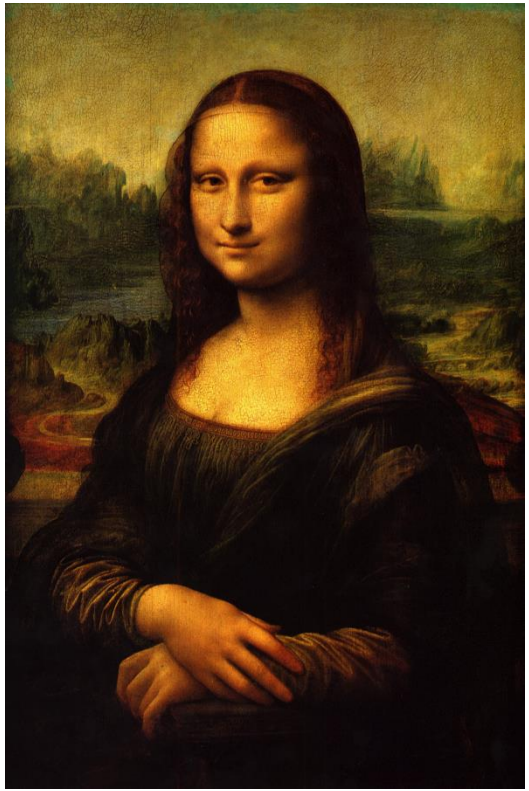

House Keeping Rules

- This is an interactive session. Feel free to ask questions – I will repeat them for the audience
- You can ask questions through the virtual on-line tool
- There will be a game played through out the session – so be prepared for questions ☺
- Let's have fun!

“Scalable Midsize Data Centre Designs”

A few key-words to consider

SCALABLE:

- Right-sizing the Data Centre, not just large scale.
- Using components that will also transition easily into larger designs.

MIDSIZE:

- Requiring a dedicated pair of DC switches.
- The transition point upwards from collapsed-core.
- Separate Layer 2/3 boundary, with DC-oriented feature set.
- Layer-2 edge switching for virtualisation.

DATA CENTRE DESIGNS:

- Topology options from single to dual-tier data centres.
- Tradeoffs of components to fill topology roles.

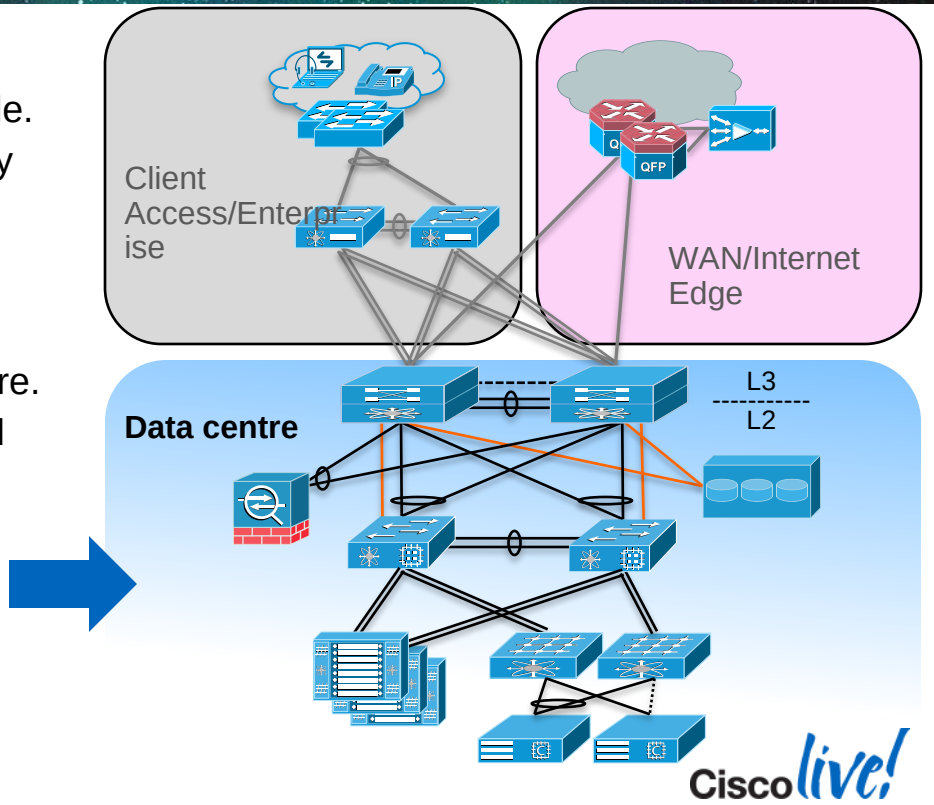

Cisco have the product line-up to build a DC of any Size and Scale

Core Message

Session Agenda

- Midsize Data Centre Requirements
  - Goals and Challenges
  - Fabric Requirements
  - Cisco Terminology, Product and Features Overview
- Starting Point: The Access Pod
  - Single Pod Design Examples
  - Nexus 5600, 6000, 7000-based
  - vPC Best Practices
- Moving to a Multi-Tier Fabric
  - Spine/Leaf Designs
  - Best Practices with FabricPath

Trends in The Data Centre

BUSINESS CHALLENGES

- Business Agility
- 24x7 Business
- Security and Compliance
- Reducing Operational Costs

OPERATIONAL CHALLENGES

- Managing Capacity
- Availability and Business Continuance
- Evolving Application Environment
- Changing IT Service Delivery Model

Growth with Investment Protection

Re-Use of key components as the design scales

Single layer expands to dual-tier with additional switch pair to form aggregation/access design

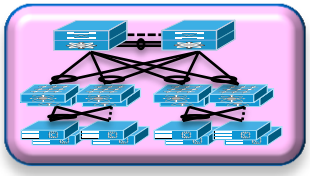

Single Layer DC Models

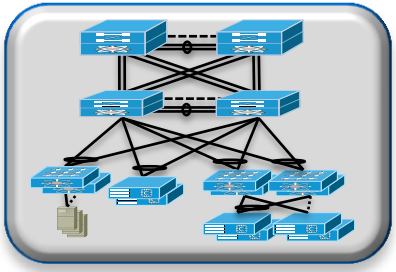

Dual Tier DC

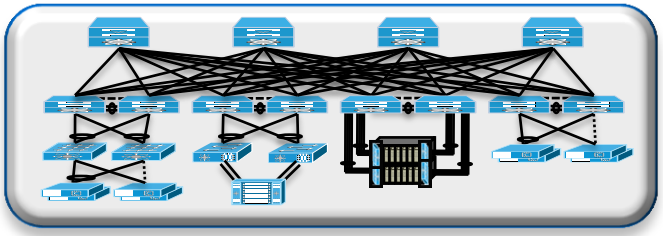

Scalable Spine/Leaf DC Fabric

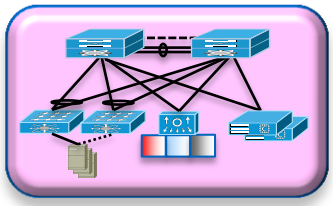

Easily scale the fabric by adding switches

- Add Spine switches to scale fabric bandwidth
- Add Leaf switches to scale edge port density

Server Edge Requirements Drive Design Choices

- Form Factor
  - Unified Computing Fabric
  - 3rd Party Blade Servers
  - Rack Servers (Non-UCSM)
- Storage Protocols
  - Fibre Channel
  - FCoE
  - iSCSI, NAS
- Virtualisation Requirements
  - vSwitch/DVS
  - Nexus 1000v
  - VM-FEX HW Switching
- NIC Connectivity Model
  - 10 or 1-GigE Server ports
  - Physical Interfaces per server
  - NIC Teaming models

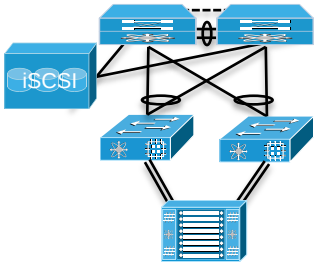

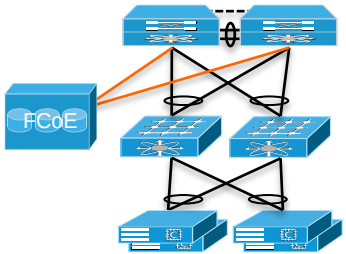

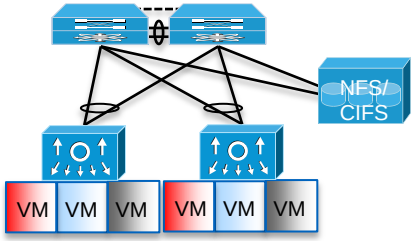

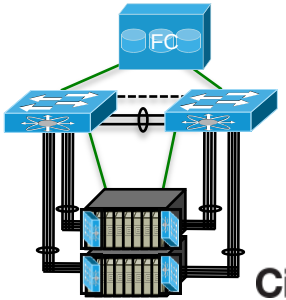

Communications Fabric Requirements

- Varied "North-South" communication needs with users and external entities.
- Increasing "East-West" communication needs: clustered applications and workload mobility.
- Flexibility to support multiple protocols and connectivity types.
- High Throughput and low latency requirements.
- Increasing availability requirements.

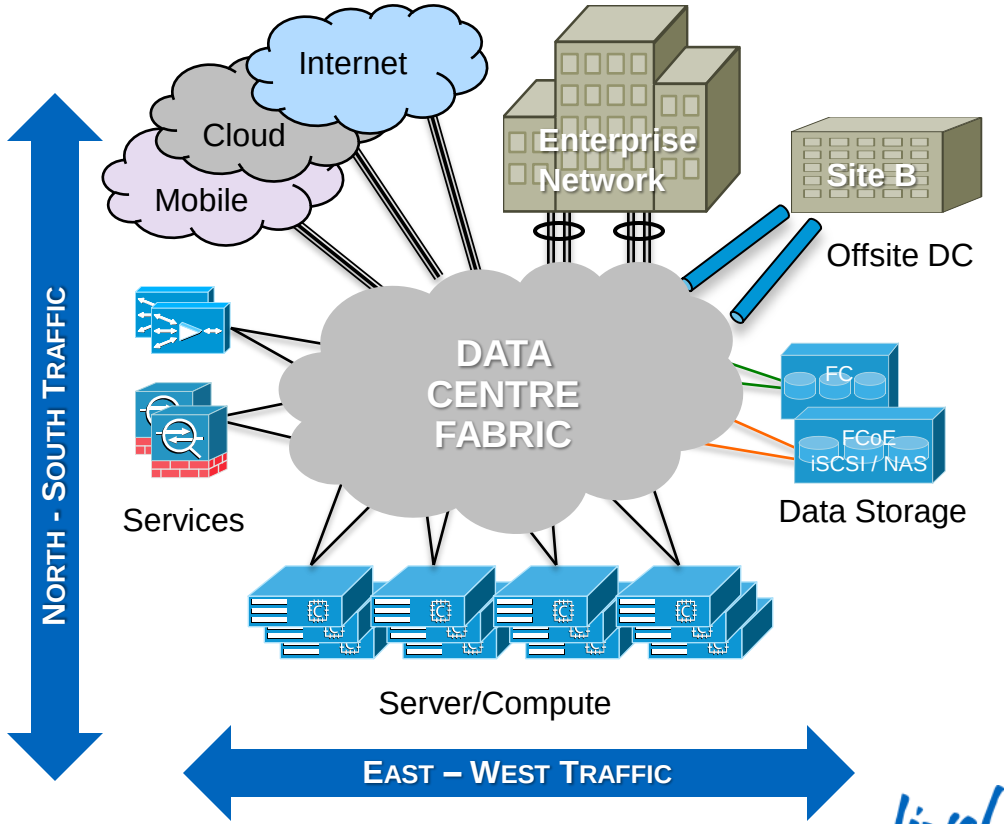

Session Agenda

- Midsize Data Centre Requirements
  - Goals and Challenges
  - Fabric Requirements
  - Cisco Terminology, Product and Features Overview
- Starting Point: The Access Pod
  - Single Pod Design Examples
  - Nexus 5600, 6000, 7000-based
  - vPC Best Practices
- Moving to a Multi-Tier Fabric
  - Spine/Leaf Designs
  - Best Practices with FabricPath

Mid-Market Size

One size does not fit all. There is no one answer!

- Australian Tax Office
  - Small Business In Australia: has to be <20 people.
  - Midsize Business In Australia: has to be <200 people.

- Cisco Categorisation: Customers with 1 to 999 Employees
  - Small = 1 to 99
  - Medium = 100 to 249
  - Mid Market = 250 to 999

Cisco Technologies – Let The Games Begin

Cisco Bingo				
Nexus	UCS	FEX	Network Services	Unified Ports
VDC's	FCoE	VPC	VPC+	EvPC
FabricPath	DFA	ACI	OTV	LISP

Expanding DC and Cloud Networking Portfolio

- 55K+ NX-OS customers
- 17K+ FEX customers
- 8.5K+ Nexus 1KV customers
- 3K+ FabricPath customers

APIC-EM (XNC)

Cisco Nexus 1000V

Cisco Nexus 2000

Cisco Nexus 3100

Cisco Nexus 3000

Cisco Nexus 5600

Cisco Nexus 5000

Cisco Nexus 6000

Cisco Nexus 7700

Cisco Nexus 7000

Application Centric Infrastructure (ACI)

Cisco Nexus 9000 ACI

OPEN
APIs/ Open Source/ Application Policy Model

HIGH PERFORMANCE FABRIC
1/10/40/100 GE

SCALABLE SECURE SEGMENTATION
VXLAN

DELIVERING TO YOUR DATA centre NEEDS

- Resilient, Scalable Fabric
- Workload Mobility Within/ Across DCs
- LAN/SAN Convergence
- Operational Efficiency
- Architectural Flexibility

Storage Networking

From separate networks to a converged fabric

- Flexibility to support dynamic connectivity requirements
- Block-level protocols:
  - Fibre Channel
  - Fibre Channel over Ethernet (FCoE)
  - iSCSI
- File-based protocols:
  - CIFS
  - NFS

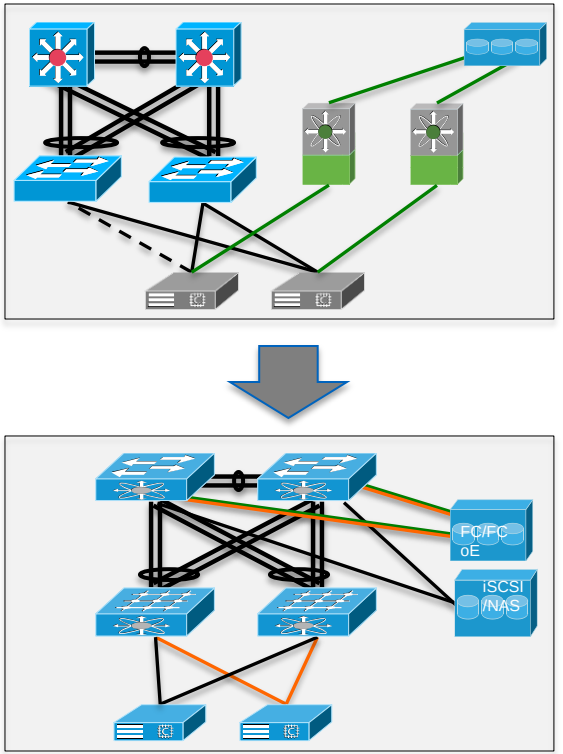

The Goal of a Unified Data Centre Fabric

Keep storage connectivity options open

- Flexibility to support a range of storage systems
- Connectivity for IP/Ethernet and Fibre Channel Endpoints
- ‘Any Server/RU to Any Storage Spindle’

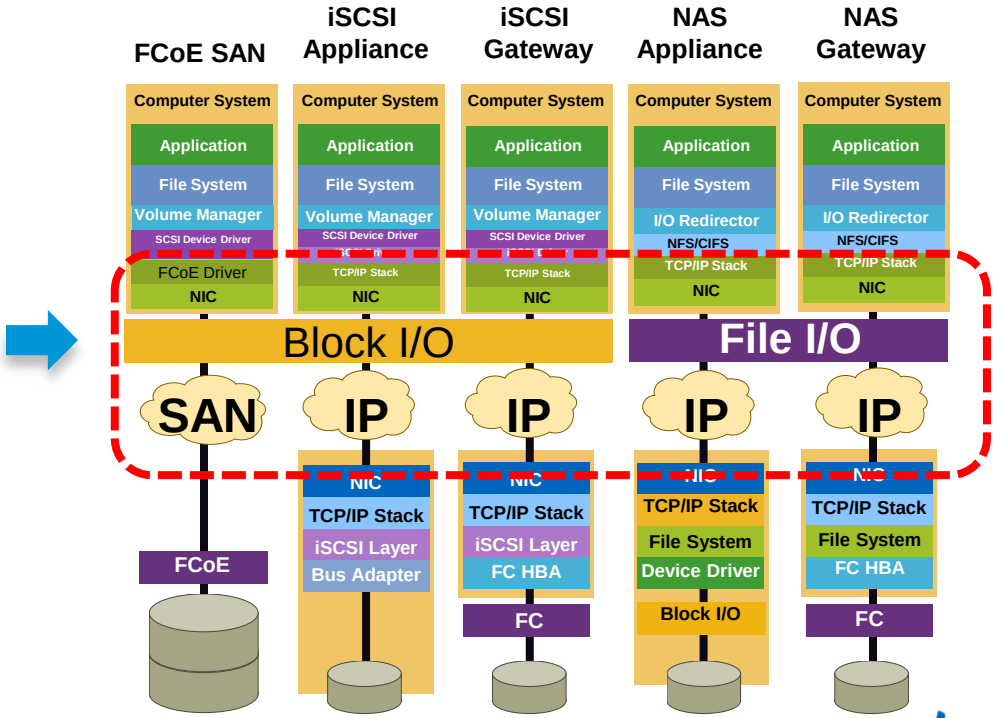

Virtual Port Channel - vPC

Increasing stability and throughput in the Layer-2 domain

- vPC allows a pair of switches to create a port-channel as if they were a single device
- Spanning Tree Protocol (STP) no longer providing the primary loop prevention mechanism
- Eliminates the STP blocked links in traditional topologies, increasing usable bandwidth
- BRKDCT-2048 – Deploying vPC

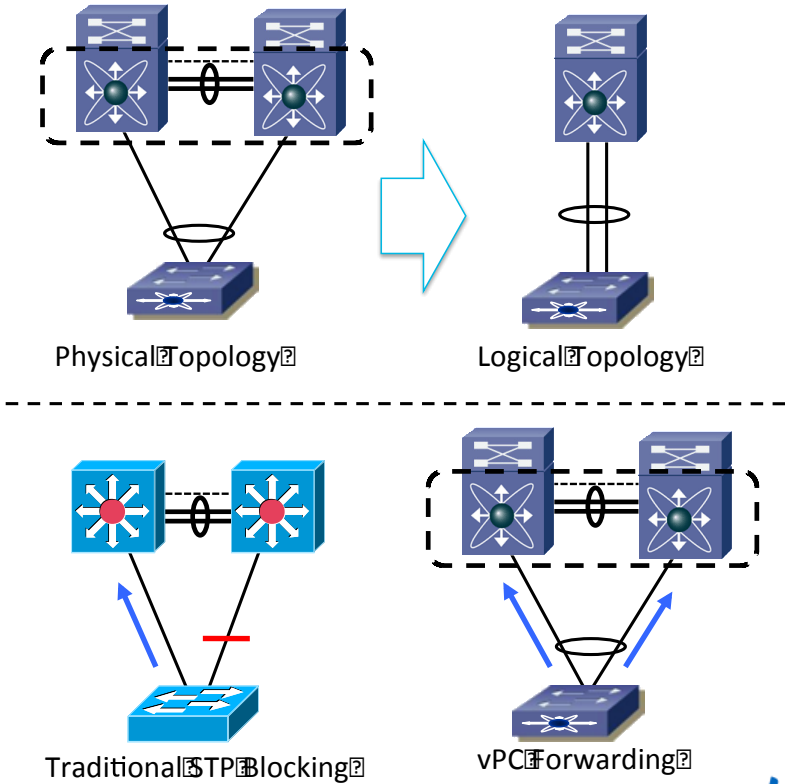

Fabric Extender - FEX

Changing the boundaries of the Ethernet switch

- Using FEX provides Top-of-Rack presence in more racks with reduced points of management, less cabling, and lower cost.
- In a “straight-through” FEX configuration, each Nexus 2000 FEX is only connected to one parent switch.
- Supported straight-through FEX parent switch is Nexus 7000, 6000 or 5000 Series.
- The same Fabric Extension technology is used between UCS FI and I/O Modules in blade chassis.
- See current platform FEX scale numbers on cisco.com under configuration guides.

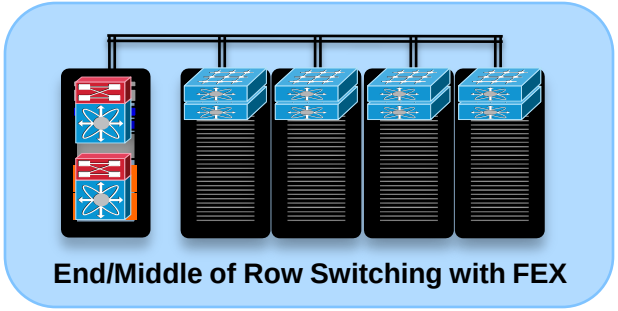

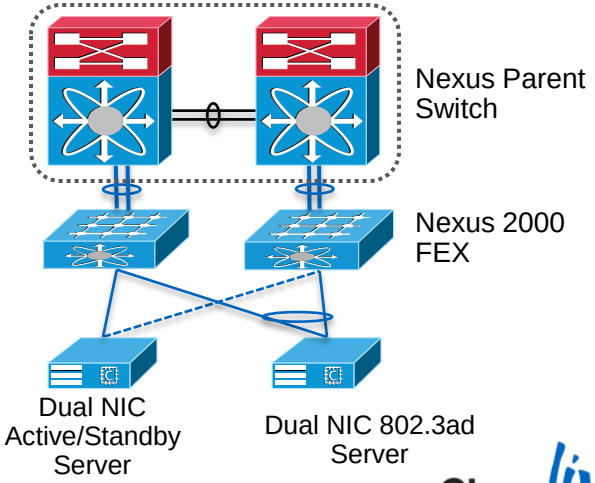

Access Pod Features: Enhanced vPC (EvPC)

- In an Enhanced vPC configuration any server NIC teaming configuration will be supported on any port. No ‘orphan ports’ in the design.
- All components in the network path are fully redundant.
- Supported dual-homed FEX parent switch is Nexus 6000 or 5500.
- Provides flexibility to mix all three server NIC configurations (single NIC, Active/Standby and NIC Port Channel).

Note, Port Channel to active/active server is standard port channel, not configured as “vPC”.

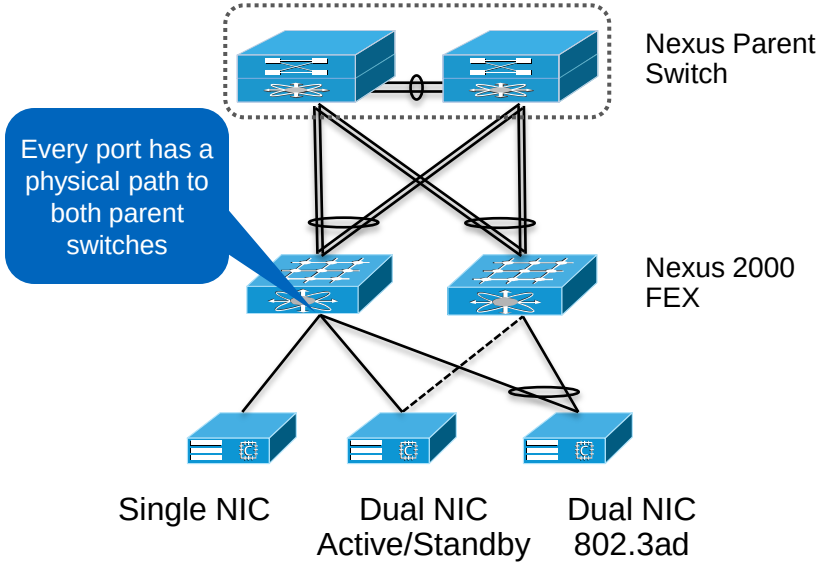

Access Pod Features: Unified Ports and FCoE

- Unified Port allows a physical port to be configured to support either native Fibre Channel or Ethernet.
- SFP+ optic needs to be chosen to support the setting of the port
- Fibre Channel over Ethernet (FCoE) allows encapsulation and transport of Fibre Channel traffic over a shared Ethernet network
- Traffic may be extended over Multi-Hop FCoE, or directed to an FC SAN
- SAN “A” / “B” isolation is maintained across the network

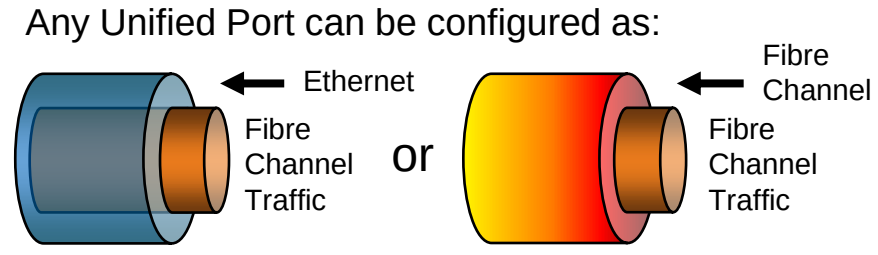

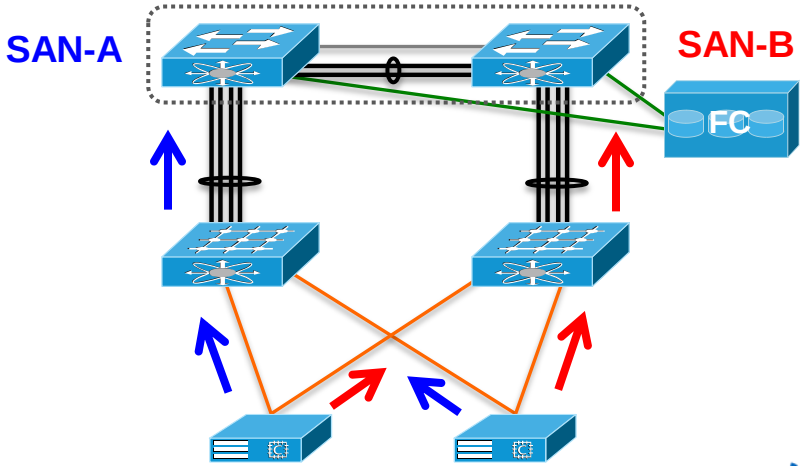

Cisco Adapter-FEX

Fabric Extender implemented in the server NIC

- Adapter-FEX presents standard PCIe virtual NICs (vNICs) to servers
- Adapter-FEX virtual NICs are configured on the server and managed via Nexus 5000
- Forwarding, Queuing, and Policy enforcement for vNIC traffic by Nexus 5000
- Adapter-FEX can be connected to Nexus 2000 Fabric Extender for a cascaded FEX-Link deployment
- Forwarding, Queuing, and Policy enforcement for vNIC traffic still done by Nexus 5000

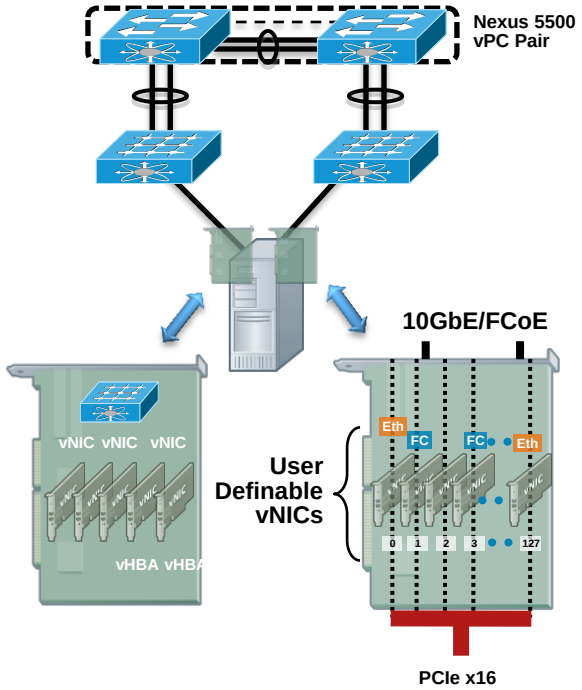

Cisco Nexus 1000v

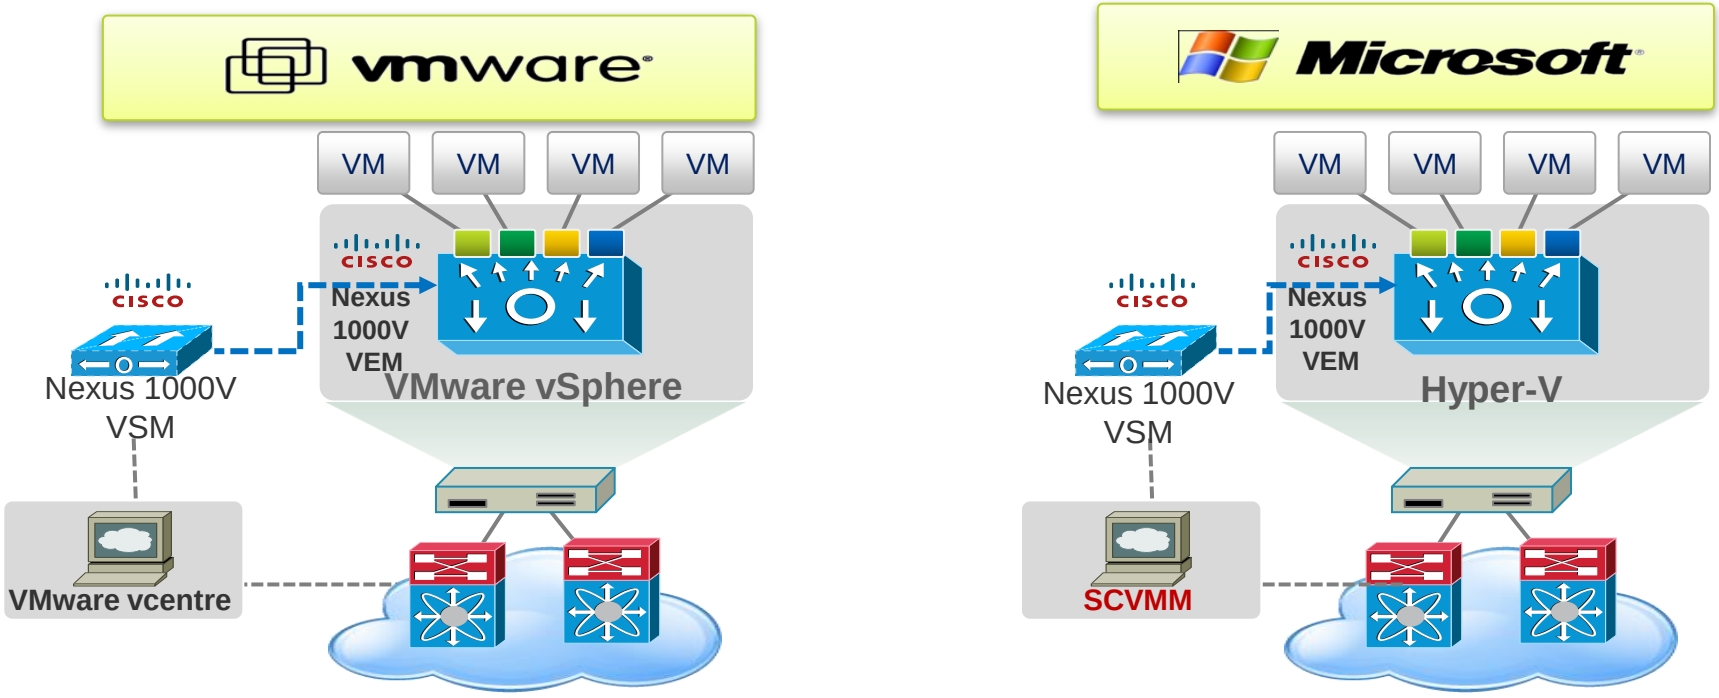

Consistent architecture, feature-set & network services ensures operational transparency across multiple hypervisors.

Cisco UCS Architecture

CONVERGENCE

AUTOMATION

INTELLIGENCE

Industry Standard APIs

UCS Manager

Fabric Interconnects

Fabric Extenders
(I/O modules)

Compute
Blade Form Factor

Nexus 7000 Virtualisation with VDCs

Single physical switch acting as multiple virtual devices

- VDC – Virtual Device Context
- Flexible separation/distribution of hardware resources and software components
- Complete data plane and control plane separation

- Complete software fault isolation
- Securely delineated administrative contexts
- Forwarding engine scalability with appropriate interface allocation

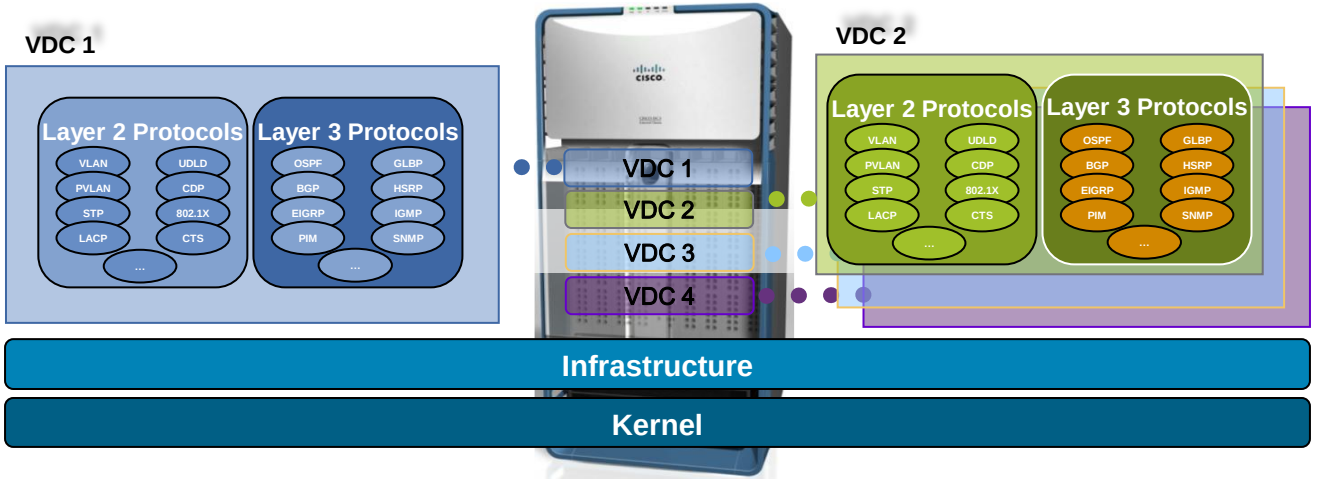

Switch Partitioning
Virtual Device Context—VDC

Separate Logical Entities Within a Switch—Available on Cisco Nexus® 7000

USE CASES

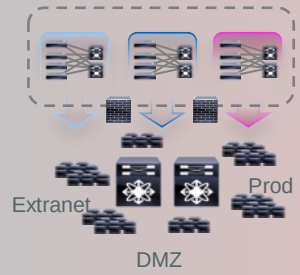

Isolating Security Domains

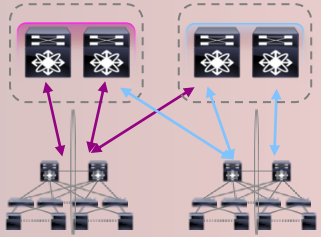

Separating Networks

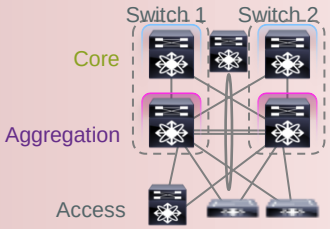

Collapsing Architecture Layers

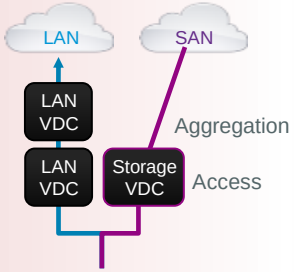

LAN/SAN Isolation

RESULTING IN

- Lower CapEx—Reduced number of physical switches
- Lower OpEx—Reduced power and management requirements
- Flexible separation/distribution of resources
- Hardware and software fault isolation

Overlay Transport Virtualisation (OTV)

O — ***Overlay*** **-** A solution that is *independent of the infrastructure technology* and services, flexible over various inter-connect facilities

T — ***Transport*** **-** Transporting services for *layer 2 and layer 3* Ethernet and IP traffic

V — ***Virtualisation*** **-** Provides *virtual* connections, *connections* that are in turn *virtualised and partitioned* into VPNs, VRFs, VLANs

OTV LAN Extensions

OTV delivers a virtual L2 transport

Overlay Transport Virtualisation

Technology Pillars

OTV is a “MAC in IP” technique for supporting Layer 2 VPNs **OVER ANY TRANSPORT**.

Uses Standard IS IS protocol

Dynamic Encapsulation

- No Pseudo-Wire State Maintenance
- Optimal Multicast Replication
- Multi-point Connectivity
- Point-to-Cloud Model

Protocol Learning

- Built-in Loop Prevention
- Preserve Failure Boundary
- Seamless Site Addition/Removal
- Automated Multi-homing

Physical Connection Choices

40G QSFP+ Interfaces	40GBASE-SR4 		40GBASE-SR4 	40G QSFP+ Interfaces
40G QSFP+ Interfaces				40G QSFP+ Interfaces
40GE SFP+ Interfaces	40GBASE-SR4 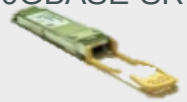			10GE SFP+ Interfaces
40GE SFP+ Interfaces				10GE SFP+ Interfaces

QSFP+ to QSFP+ 1,3,5,7,10m Copper, QSFP+ to SFP+ 7,10m Copper,
QSFP-SR4 (100m over OM3), QSFP-40G-CSR4 (300m over OM3)

Working with 40 Gigabit Ethernet

Nexus 6000 and 7000 support QSFP-based 40 Gigabit Ethernet interfaces.*

Nexus 6004 at FCS provides QSFP ports only, but splitter cables can be used to provision 4x10GigE ports out of 1 QSFP

40 Gigabit Ethernet cable types:

- **Direct-attach copper** [QSFP <-> QSFP] and [QSFP <-> 4 x SFP+]. Passive cables at 1/3/5m, active cables at 7 and 10m.
- **SR4** uses bit-spray over 4 fibre pairs within a 12 fibre MPO/MTP connector to reach up to 100/150m on multimode OM3/OM4
- **CSR4** is a higher powered SR4 optic with reach up to 300/400m on multimode OM3/OM4
- **LR4** uses CWDM to reach up to 10km on a single-mode fibre pair.

* Verify platform-specific support of specific optics/distances

QSFP-40G-CR4 direct-attach cables

QSFP+ to 4-SFP+ direct-attach cables (splitter)

QSFP-40G-SR4 with direct MPO and 4x10 MPO-to-LC duplex splitter fibre cables

F3-Series 40G Offering
QSFP BiDi support

Technology

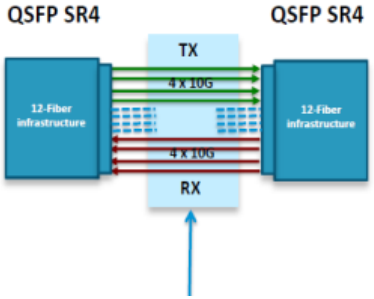

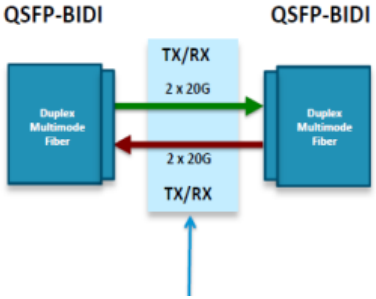

12-fibre MPO **Duplex LC**

Value Proposition

- Utilise existing duplex fibre commonly deployed in 10G environment today
- Reduce 40G transition cost by eliminating the need to upgrade fibre plant
- 75% average savings over parallel fibre for new deployments

QSFP References

- QSFP BiDi 40Gig Data Sheet

http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps13386/datasheet-c78-730160.html

- QSFP BiDi Compatibility Matrix

http://www.cisco.com/en/US/docs/interfaces_modules/transceiver_modules/compatibility/matrix/OL_24900.html

- QSFP 40Gig Datasheet

http://www.cisco.com/en/US/prod/collateral/modules/ps5455/data_sheet_c78-660083_ps11541_Products_Data_Sheet.html

- Platform Specific SFP + QSFP Support

http://www.cisco.com/en/US/prod/collateral/modules/ps5455/data_sheet_c78

- 40Gig Cabling WP's

http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps13386/white-paper-c11-729493.pdf

http://www.cisco.com/en/US/products/ps11708/index.html

Pop Quiz

What is one value of using a Fabric Extender (FEX)

A Simple Formula

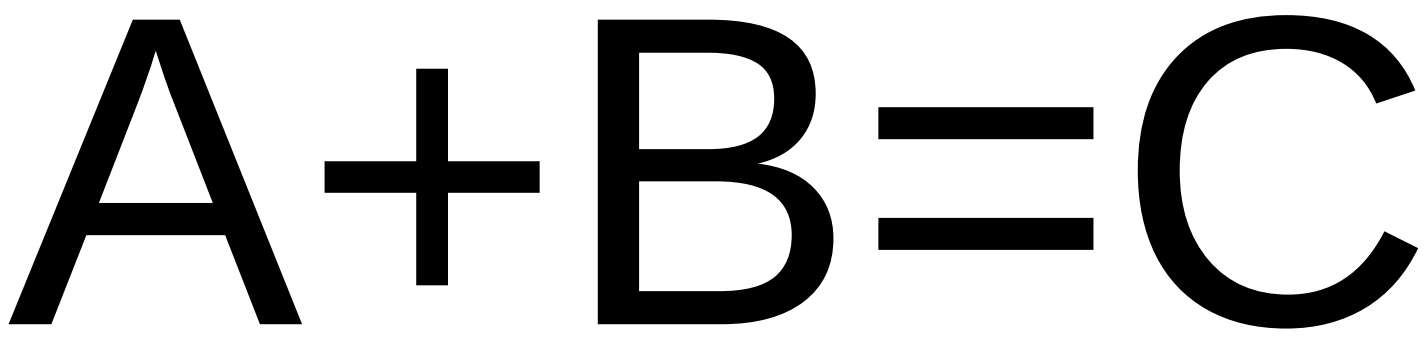

Session Agenda

- Midsize Data Centre Requirements
  - Goals and Challenges
  - Fabric Requirements
  - Cisco Terminology, Product and Features Overview
- Starting Point: The Access Pod
  - Single Pod Design Examples
  - Nexus 5600, 6000, 7000-based
  - vPC Best Practices
- Moving to a Multi-Tier Fabric
  - Spine/Leaf Designs
  - Best Practices with FabricPath

Access Pod Basics
Compute, Storage, and Network

“Different Drawing, Same Components”

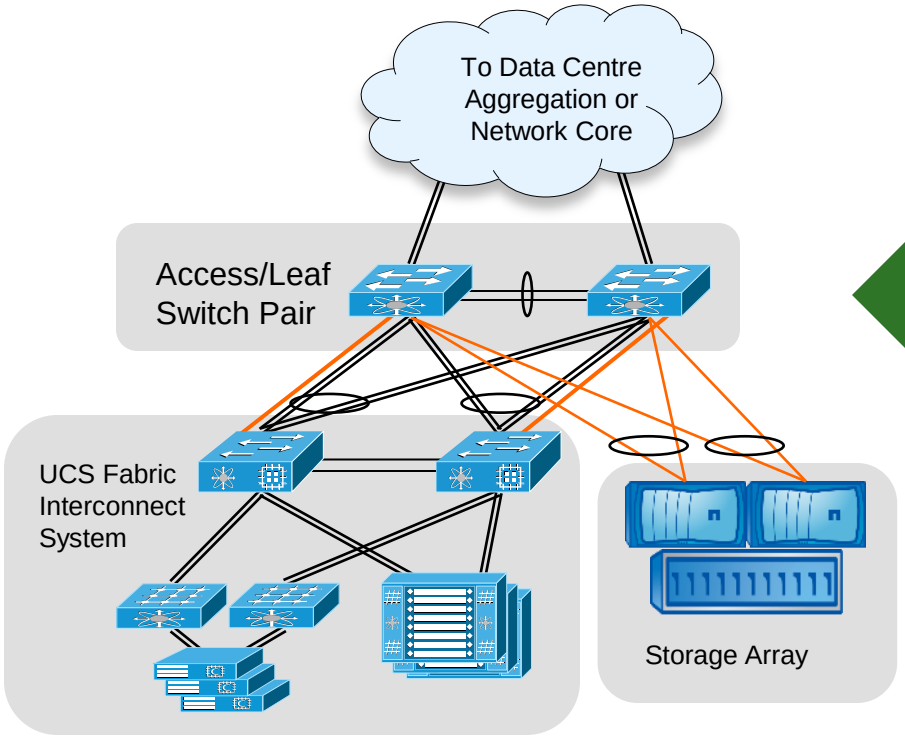

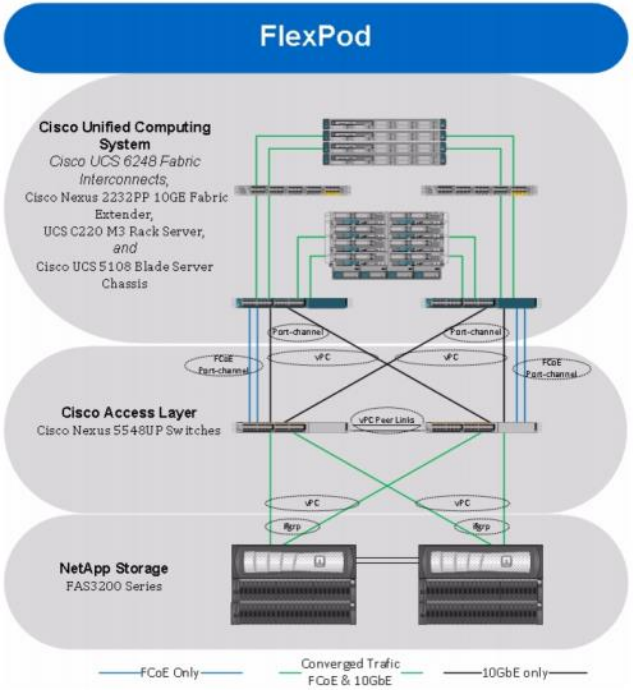

Nexus 5000/6000 Series Positioning

40G Innovations

Nexus 6000

High 40G Density
Low Latency
100G Uplinks

10G Innovations

Nexus 5600

High 10G Density
Low Latency
40G Uplinks

Nexus 5500

Flexibility
10G Uplinks

Nexus 5010/ 5020

Highest Density

Fabric Innovations | Programmability

VXLAN | Network Visibility

LAN/SAN Convergence | FabricPath

FEX Architecture

CUSTOMER VALUE

CISCO INNOVATION

Data Centre Service Integration Approaches

Data Centre Service Insertion Needs

- ✓ Firewall, Intrusion Prevention
- ✓ Application Delivery, Server Load Balancing
- ✓ Network Analysis, WAN Optimisation

Physical Service Appliances

- ➢ Typically introduced at Layer 2/3 Boundary or Data centre edge.
- ➢ Traffic direction with VLAN provisioning, Policy-Based Routing, WCCP.
- ➢ Use PortChannel connections to vPC.
- ➢ Static Routed through vPC, or transparent.

Virtualised Services

- ➢ Deployed in a distributed manner along with virtual machines.
- ➢ Traffic direction with vPath and Nexus 1000v.

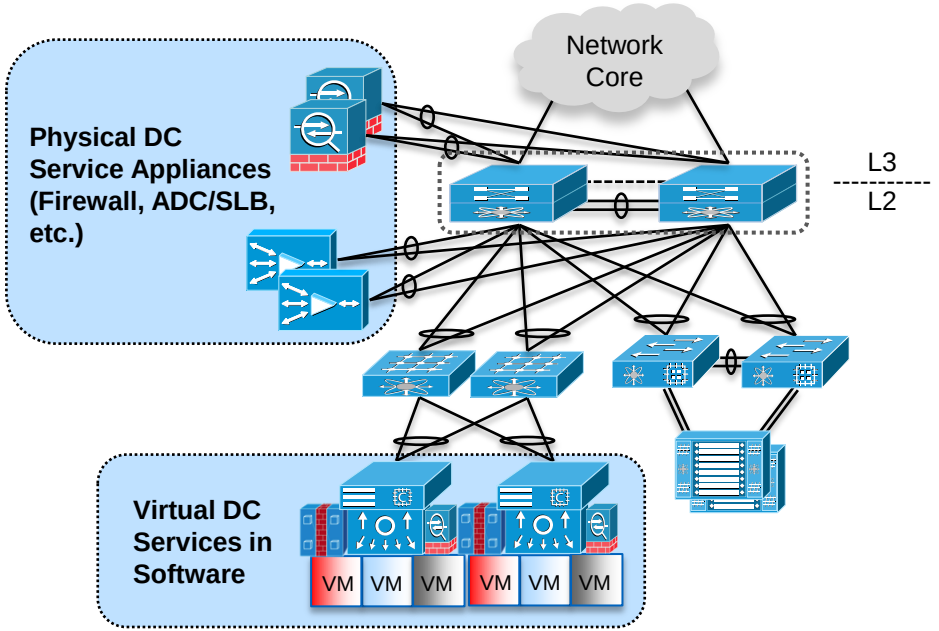

Virtualised Servers with Nexus 1000v and vPath

Challenge #1

We need a data centre design that achieves the following:

- High 10GB connectivity
- Layer 3
- FHRP
- Native Fibre Channel Storage along with IP Storage
- Connect legacy and new servers of various speeds together
- FEX

Single Layer Data Centre, Nexus 5600 Series

Dedicated Nexus 5600 Series DC Switch Pair

- Unified Ports support native Fibre Channel, FCoE, iSCSI or NAS Storage.
- Nexus 5600 supports physical FEX, Adapter-FEX, VM-FEX capabilities.
- 5600 integrated line rate L3
- 5600 VXLAN support

Notes:

- OTV/LISP DCI services may be provisioned on separate Nexus 7000 or ASR 1000 WAN Routers

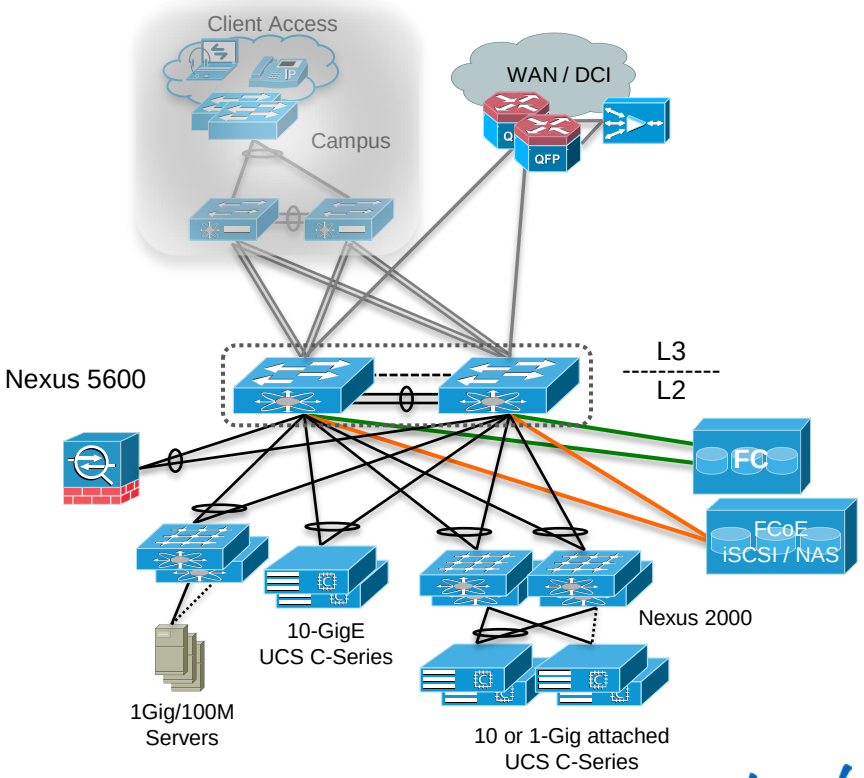

Single Layer Data Centre, Nexus 6001

Dedicated Nexus 6001-based Data Centre, switch pair

Nexus 6001 benefits:

- Integrated line-rate layer-3
- Native 40-Gig capability
- Low ~1us switch latency at Layer-2/3
- Up to 16 full-rate SPAN sessions
- Greater 10GigE port density (at 1-RU)

Example Design Components:

- 2 x Nexus 6001, Layer-3 and Storage Licenses
- 4 x Nexus 2232PP/TM-E (10-GigE FEX)
- 2 x Nexus 2248TP-E (1-GigE FEX)

Note: Native Fibre Channel storage support would require separate Nexus 5500/5600 or MDS SAN

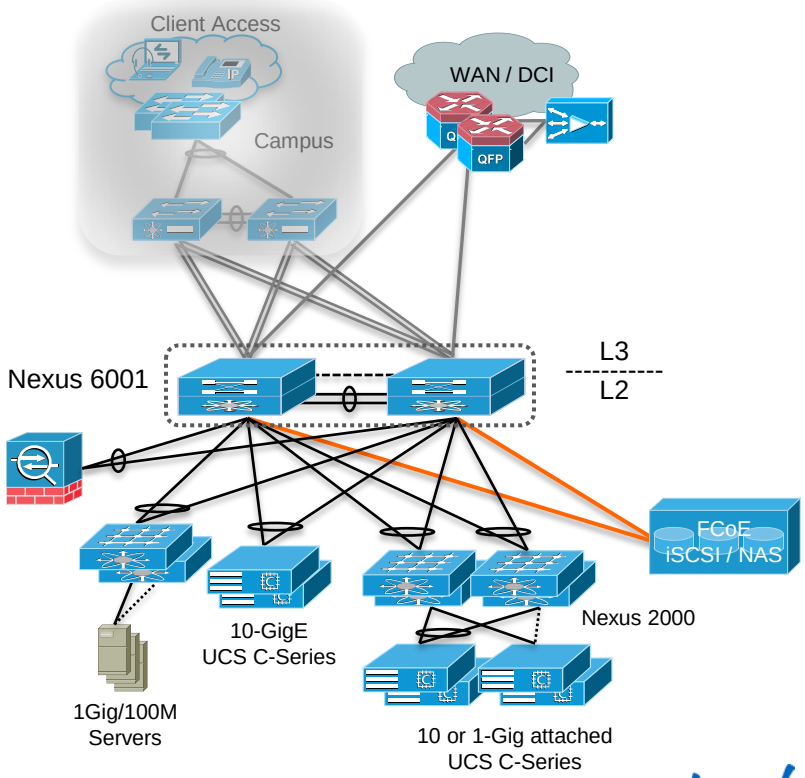

Single Layer Data Centre plus UCS Fabric

Alternate Server Edge 1: UCS Fabric Interconnects with Blade and Rack Servers

- UCS Fabric Interconnects support up to 20 UCS Chassis
- UCSM can also manage C-Series servers
- Dedicated FCoE uplinks from UCS FI to the Nexus 6001 for FCoE SAN Access
- Nexus switching layer provides inter-VLAN routing, upstream connectivity, and storage fabric services.
- Example DC Switching Components:
  - 2 x Nexus 6001
  - Layer- 3 and Storage Licensing
  - 2 x Nexus 2232PP/TM-E

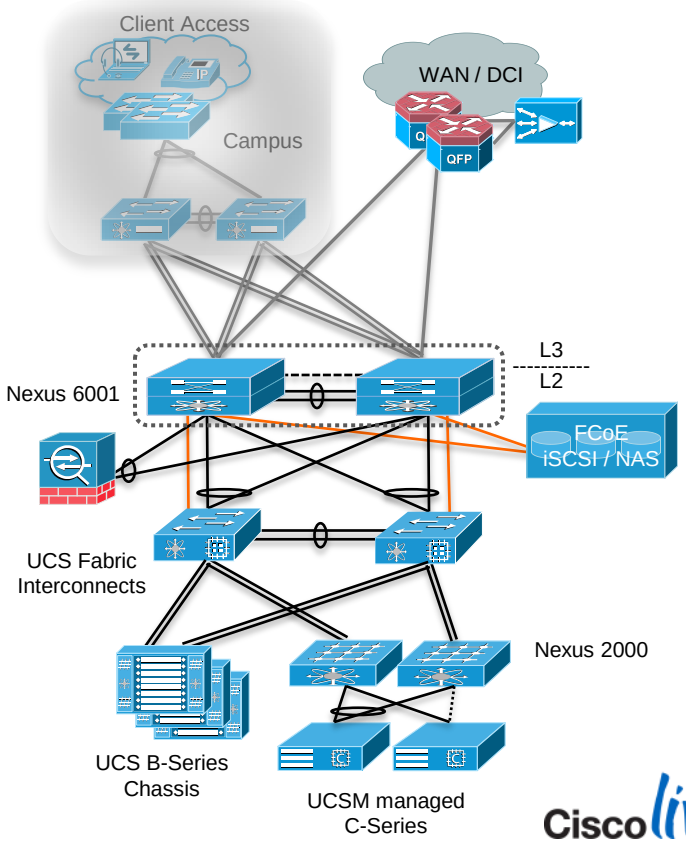

Challenge #2

We need a data centre design that achieves the following:

- Connect multiple PODs together
- High 40GB bandwidth
- Support for FEX
- Ability to dual home servers

Flexible Design with Access Pod Variants

Mix and match Layer-2 compute connectivity for migration or scale requirements

- Nexus switching and FEX provide operational consistency

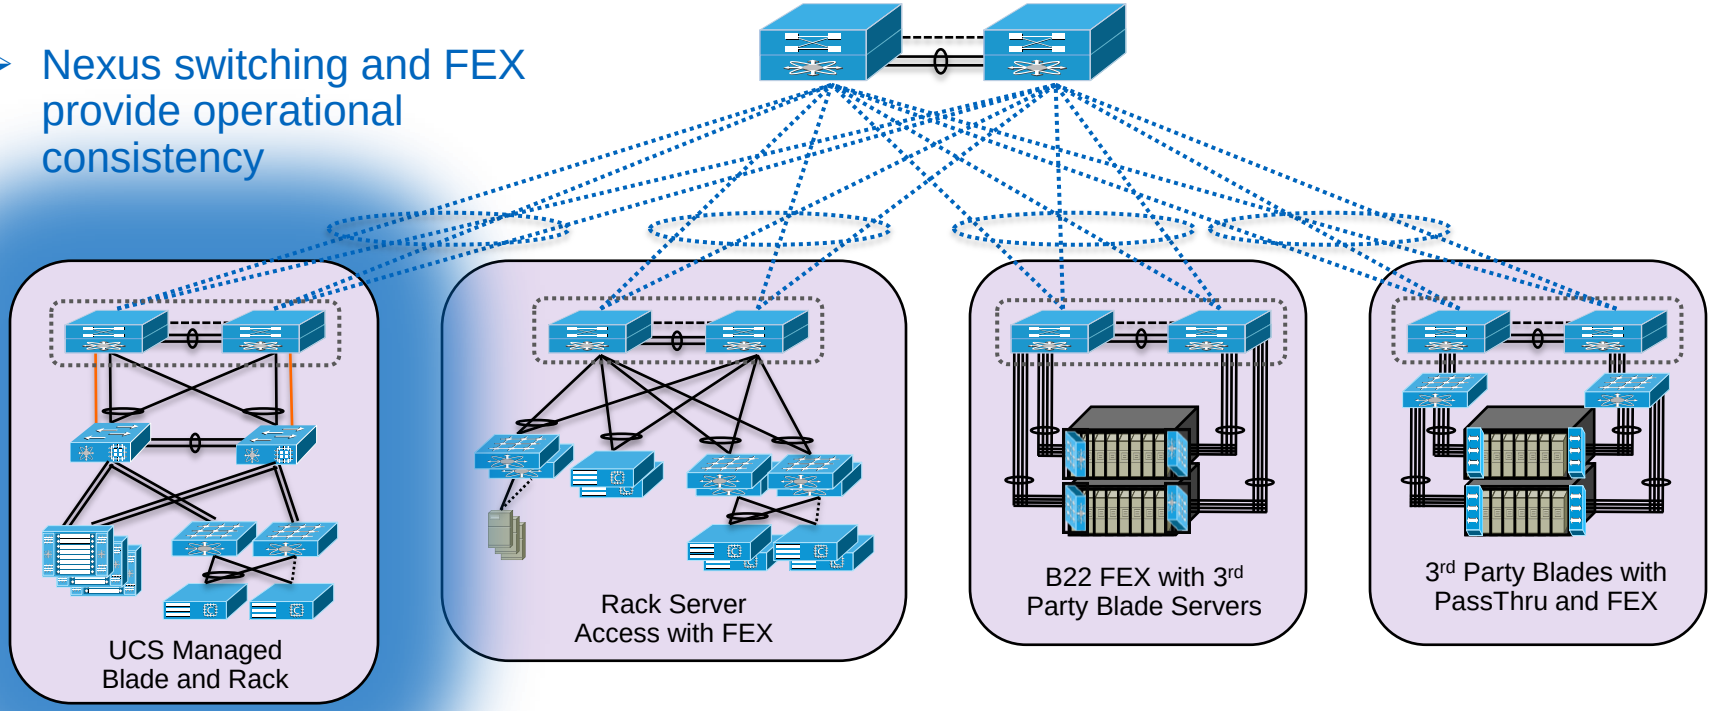

More features, highest value and physical consolidation

Single Layer Data Centre, Nexus 6004

Positioned for rapid scalability and a 40-GigE Fabric

Nexus 6004 Benefits:

- Includes 48 40-GigE QSFP or 192 10-GigE ports, up to 96 40-GigE or 384 10-GigE
- Integrated line-rate layer-3
- Native 40-Gig switch fabric capability
- Low ~1us switch latency at Layer-2/3
- Line-rate SPAN at 10/40 GigE

Example Components:

- 2 x Nexus 6004, 24 40G or 96 10G ports active
- L3 and Storage Licensing
- 8 x Nexus 2232PP/TM-E

Note: FCoE, iSCSI, NAS storage are supported on 6004 native FC module

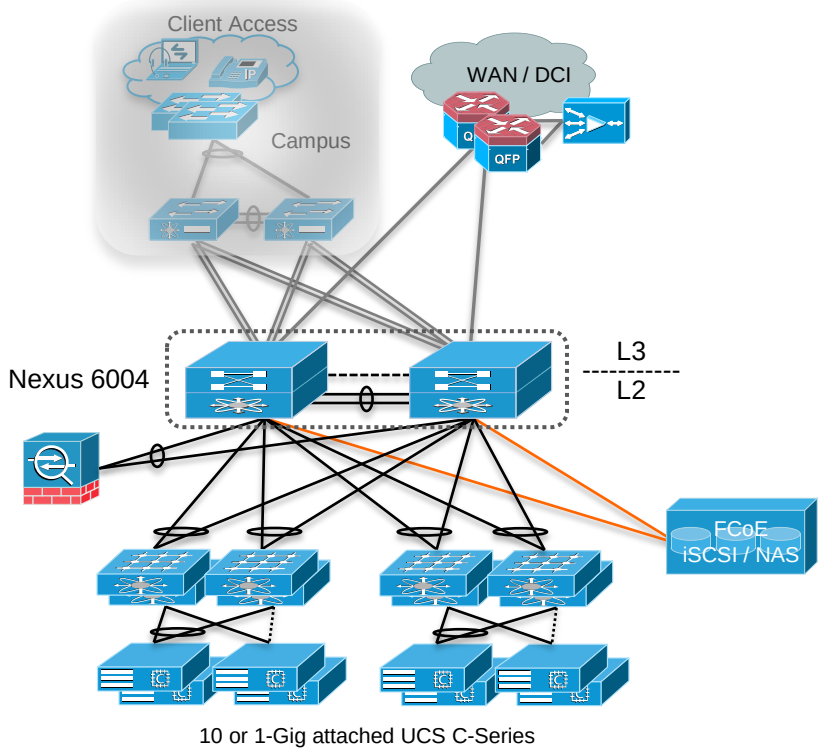

Challenge #3

We need a data centre design that achieves the following:

- High bandwidth connectivity supporting:
  - 10GB for hosts
  - 40GB and more for future requirements in
- High availability
- Scalable POD Design
- Option for interconnecting other data centres together (DCI)
- Multiple protocols (MPLS, VPLS etc)
- VDC’s
- Director class FCoE

Single Layer Data Centre, Nexus 7004/7706

Highly Available Virtualised Chassis Access/Aggregation Model

Benefits of Nexus 7004/7706 with F3 I/O Modules:

- Supervisor High Availability
- Layer-2/3 ISSU
- Virtual Device Contexts (VDC)
- PBR, WCCP for service integration
- FabricPath support for future expansion

Example Components:

- 2 x Nexus 7004/7706, dual SUP-2/2e
- Dual F3 line cards Layer-3 Licensing
- 4 x Nexus 2232PP

Note: For native Fibre Channel or FCoE add Nexus 5500 access layer or MDS SAN.

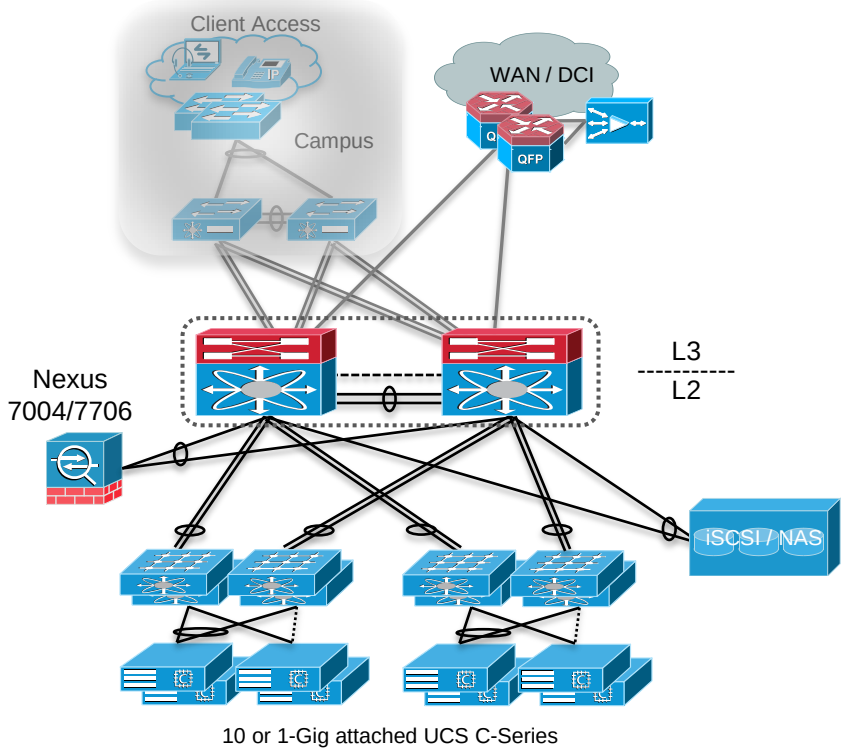

Configuration Best Practices Summary: vPC with Layer-2, Layer-3

Virtual Port Channel and Layer-2 Optimisations

What features to enable?

- Autorecovery: Enables a single vPC peer to bring up port channels after power outage scenarios
- Orphan Port Suspend: Allows non-vPC ports to fate-share with vPC, enables consistent behavior for Active/Standby NIC Teaming
- vPC Peer Switch: Allows vPC peers to behave as a single STP Bridge ID (not required with vPC+
- Unidirectional Link Detection (UDLD): Best practice for fibre port connectivity to prevent one-way communication (use “normal” mode)

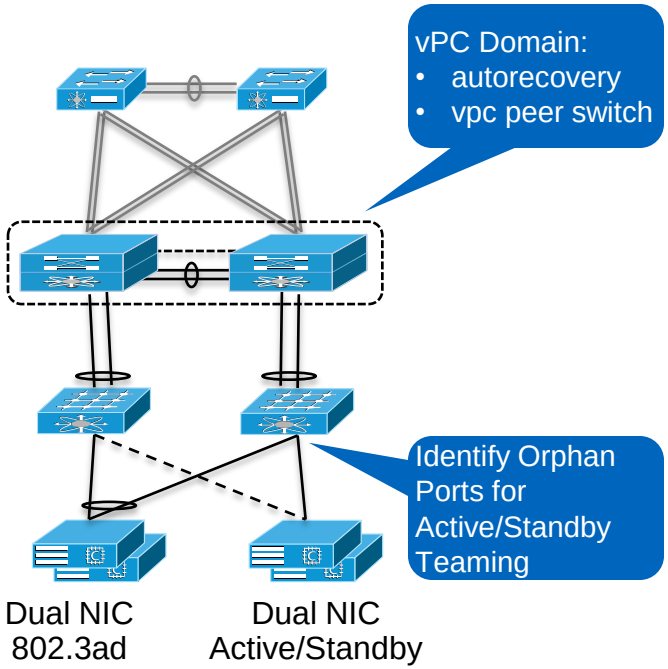

Virtual Port Channel and Layer-3 Optimisations

What features to enable?

- vPC and HSRP: Keep HSRP timers at defaults, vPC enables active/active HSRP forwarding
- vPC Peer Gateway: Allows the peers to respond to the HSRP MAC, as well as the physical MAC's of both peers.
- IP ARP Synchronise: Proactively synchronises the ARP table between vPC Peers over Cisco Fabric Services (CFS)
- Layer-3 Peering VLAN: Keep a single VLAN for IGP peering between N5k/6k vPC peers on the peer link. (On N7k can also use a separate physical link)
- Bind-VRF: Required on Nexus 5500, 6000 for multicast forwarding in a vPC environment. (Not required if using vPC+ with FabricPath)

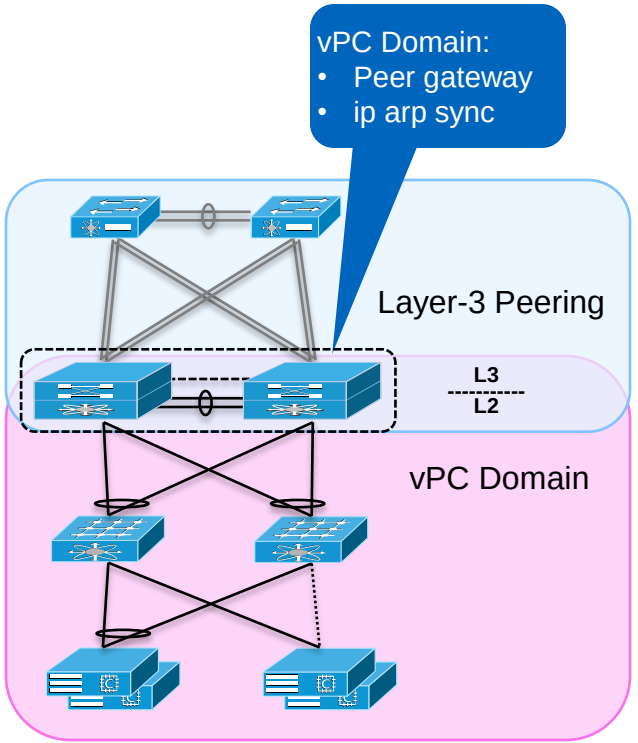

Pop Quiz

If I wanted a switch to handle both my LAN and FC storage traffic simultaneously, what switch would I need?

Session Agenda

- Midsize Data Centre Requirements
  - Goals and Challenges
  - Fabric Requirements
  - Cisco Terminology, Product and Features Overview
- Starting Point: The Access Pod
  - Single Pod Design Examples
  - Nexus 5600, 6000, 7000-based
  - vPC Best Practices
- Moving to a Multi-Tier Fabric
  - Spine/Leaf Designs
  - Best Practices with FabricPath

Migration from Single to Dual-layer Switching

Moving from single to dual-layer models:

- Larger switches (Nexus 7000, 6004) more suited to becoming spine/aggregation.
- Smaller switches (Nexus 5XXX, 6001) more suited to becoming leaf/access.
- Starting with larger switches eliminates need to move SVI's (Layer-3 gateway.)
- Aggregation switches can support access switch connections, FEX, and direct-attached servers during migration.

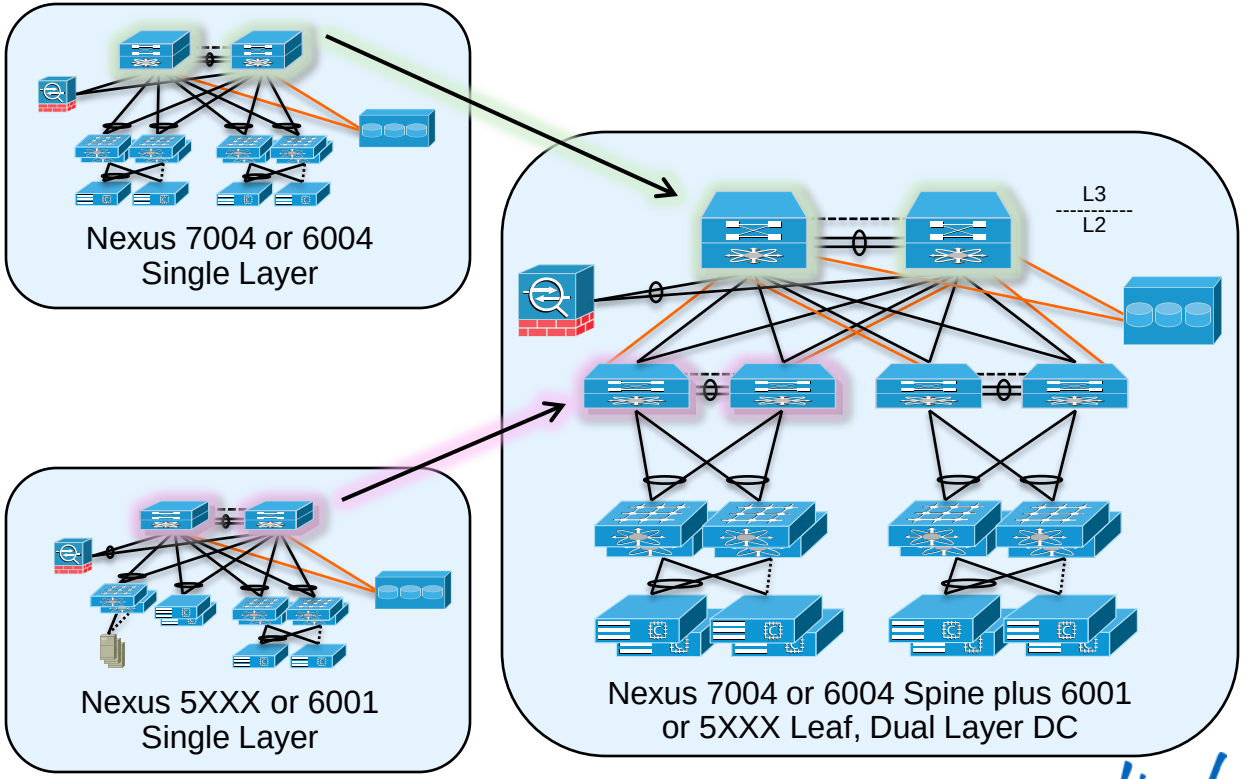

Oversubscription: Balancing Cost and Performance

Oversubscription:

- Most servers will not be consistently filling a 10 GigE interface.
- A switch may be line-rate non-blocking, but still introduce oversubscription into an overall topology by design.
- Consider Ethernet-based storage traffic when planning ratios, don't be overly aggressive.

Example device-oriented numbers with all ports active:

- Nexus 6001: 48x10Gig + 4x40Gig uplink, 48:16 or 3:1 topology-level oversubscription.
- Nexus 2232PP FEX: 32x10Gig + 8x10Gig uplink, 32:8 or 4:1 topology-level oversubscription.
- Actual oversubscription can be controlled by how many ports and uplinks are physically connected.

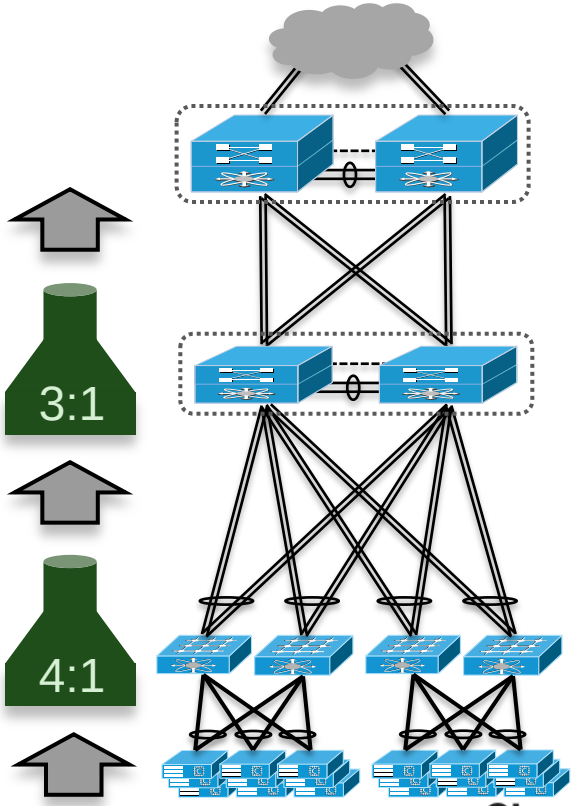

Cisco FabricPath

Scaling Layer 2 Domains—Available on Cisco Nexus® 7000, 6000 and 5000 Series

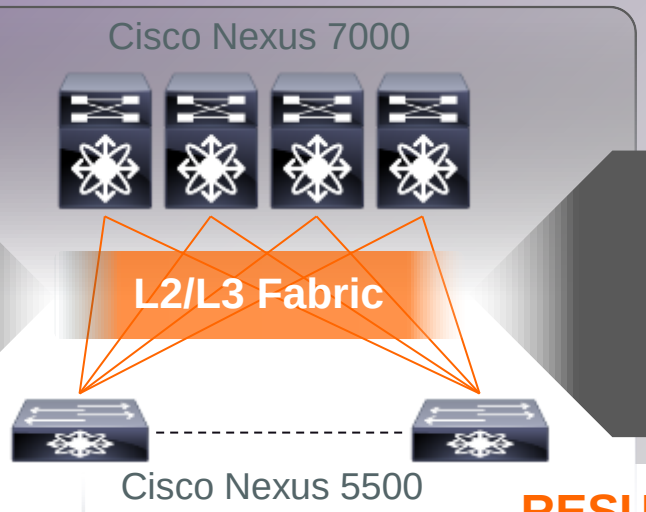

Layer 2 Strengths

- Simple configuration
- Flexible provisioning
- Low cost

Layer 3 Strengths

- All links active
- Fast convergence
- Highly scalable

USE CASES

- Scalable virtualisation pods
- Massively scalable L2 fabric
- L2 interconnect with dark fibre
- High-bandwidth HPC

RESULTING IN

- Simplicity
- High resiliency—Eliminate Spanning Tree Protocol (STP)
- Flexibility
- Scalability

FabricPath Configuration

FabricPath Terminology

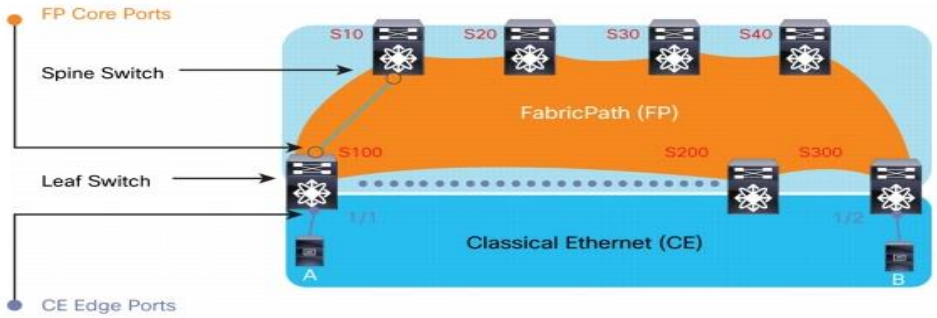

Cisco FabricPath Terminology	
Term	**Definition**
Cisco FabricPath domain	Layer 2 domain formed by interconnected Cisco FabricPath core interfaces and carrying Cisco FabricPath VLAN traffic: All traffic in the Cisco FabricPath domain is Cisco FabricPath encapsulated.
Cisco FabricPath core port	Interface connected to the Cisco FabricPath domain: The Cisco FabricPath core interface carries traffic encapsulated in Cisco FabricPath frames and can also be referred to as a FabricPath (FP) port. A Cisco FabricPath core port must be connected to another Cisco FabricPath core port. The Cisco FabricPath core port carries all Cisco FabricPath VLANs and, therefore, can be conceptually considered as a trunk port.
Cisco FabricPath edge port	Interface at the edge of the Cisco FabricPath domain: Cisco FabricPath edge interfaces carry traffic encapsulated in regular Ethernet frames and also can be referred as Classical Ethernet (CE) ports. Cisco FabricPath edge ports can be connected to any standard Ethernet port. Cisco FabricPath edge ports are used to attach any regular Ethernet device to the Cisco FabricPath domain.
Cisco FabricPath VLAN	VLAN allowed to cross a Cisco FabricPath domain.
Classical Ethernet (CE) VLAN	VLAN not allowed to cross a Cisco FabricPath domain, but allowed to exist on the edge interfaces of the Cisco FabricPath attached device.

FabricPath Configuration

FabricPath Designs – Multi POD (with FP Multi-Topology)

Layer 2 Classical Ethernet
Layer 2 FabricPath
Layer 3

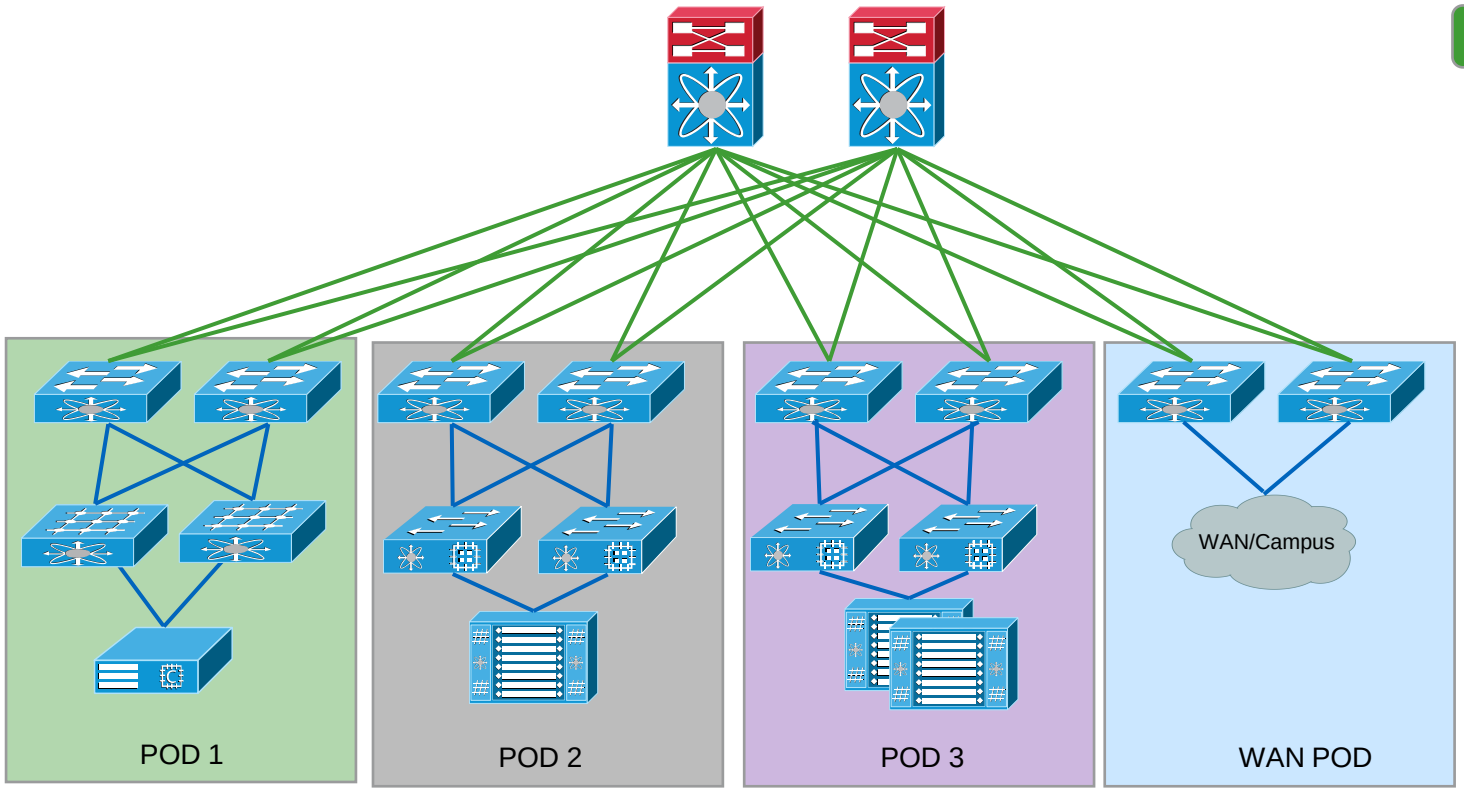

FabricPath Configuration

FabricPath Designs – Multi POD (with FP Multi-Topology)

Layer 2 Classical Ethernet
Layer 2 FabricPath
Layer 3

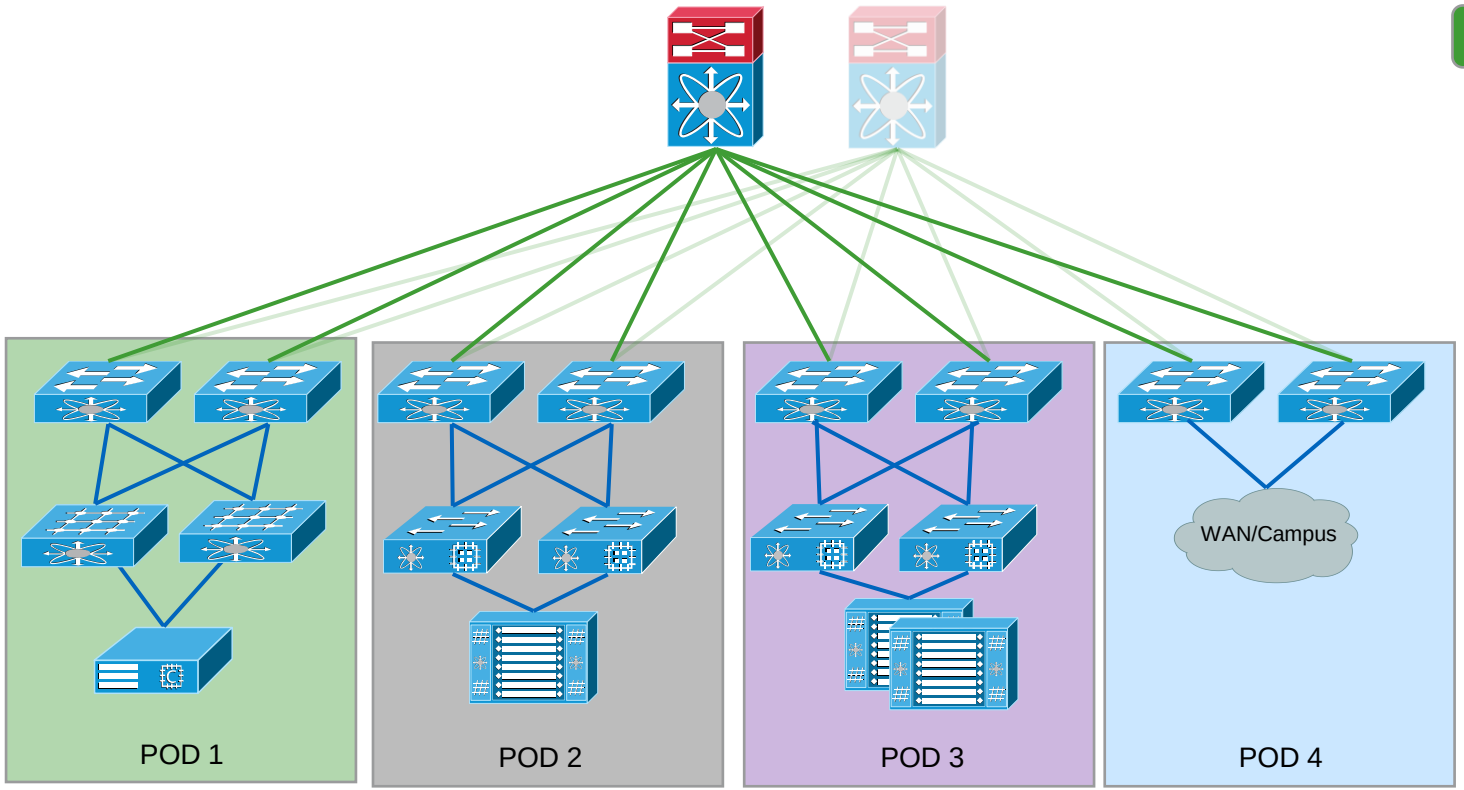

Value of FabricPath/vPC+ in Spine/Leaf Designs

Using FabricPath with a traditional DC Topology

- Aggregation becomes Spine
- Access becomes Leaf
- FabricPath Benefits:
  - ✓ Ease of configuration
  - ✓ **Completely** eliminates STP from running between Leaf and Spine
  - ✓ No Orphan Port isolation on Access (Leaf) switch vPC Peer-link loss
  - ✓ Improved Multicast support, no “bind-vrf” needed (N5500/6000), also adds PIM-SSM capability
  - ✓ Greater flexibility for future growth and change in the topology

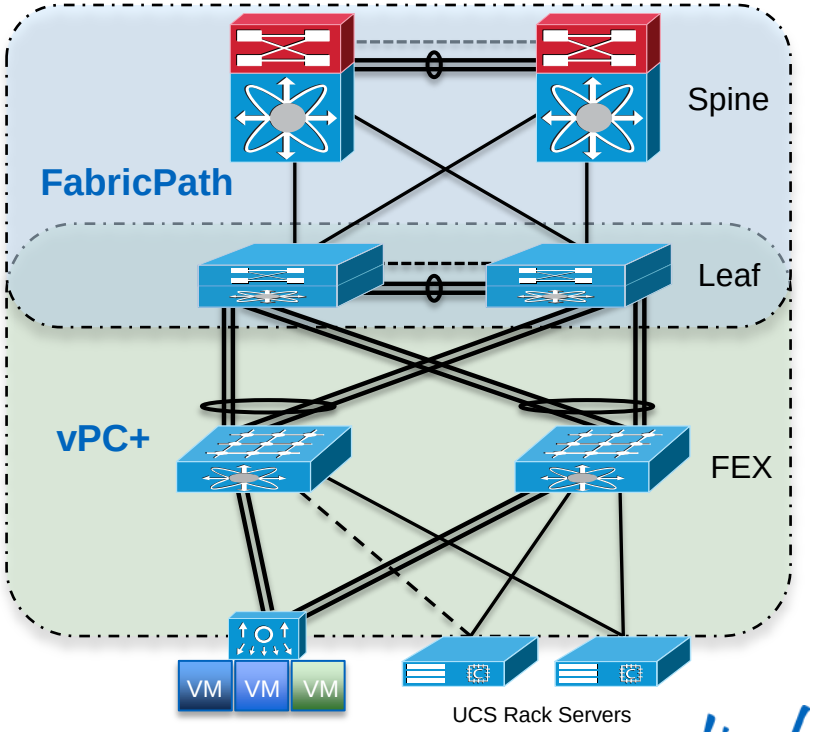

Dual Layer Nexus 6000 Data Centre

40-GigE Optimised low-latency switch fabric

Data Centre switching control plane distributed over Dual Layers.

- Aggregation (Spine): Layer-3 and services boundary, scale point to expand fabric.
- Access (Leaf): Physical TOR or FEX aggregation point, Layer-2 virtualisation services.

Multi-hop FCoE with dedicated links.

Example Components:

- 2 x Nexus 6004, 2 x Nexus 6001
- Layer-3 and Storage Licensing
- 12 x Nexus 2232PP/TM-E

Enable FabricPath between tiers for configuration simplicity and future scale.

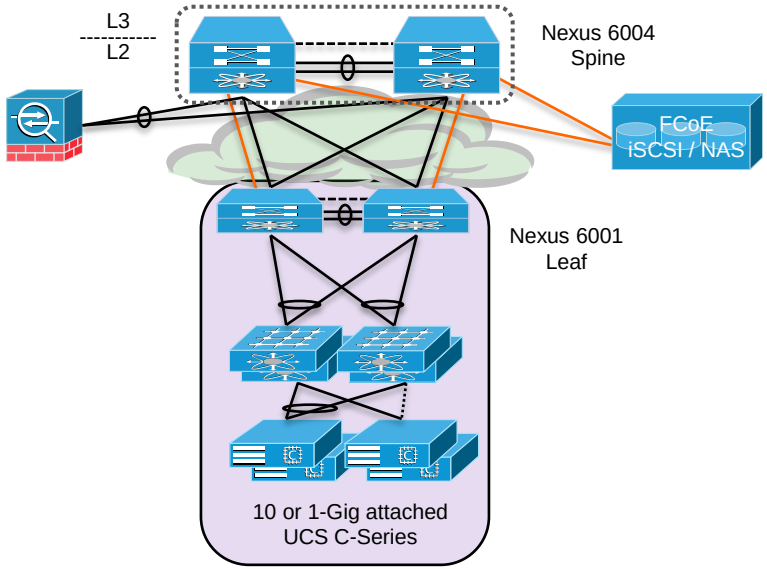

Dual Layer Nexus 6000 Data Centre, Expansion

40-GigE Optimised low-latency switch fabric

Data Centre switching control plane distributed over Dual Layers.

- Aggregation (Spine): Layer-3 and services boundary, scale point to expand fabric.
- Access (Leaf): Physical TOR or FEX aggregation point, Layer-2 virtualisation services.

Multi-hop FCoE with dedicated links.

Example Components:

- 2 x Nexus 6004, 4 x Nexus 6001
- Layer-3 and Storage Licensing
- 24 x Nexus 2232PP/TM-E

Enable FabricPath between tiers for configuration simplicity and future scale.

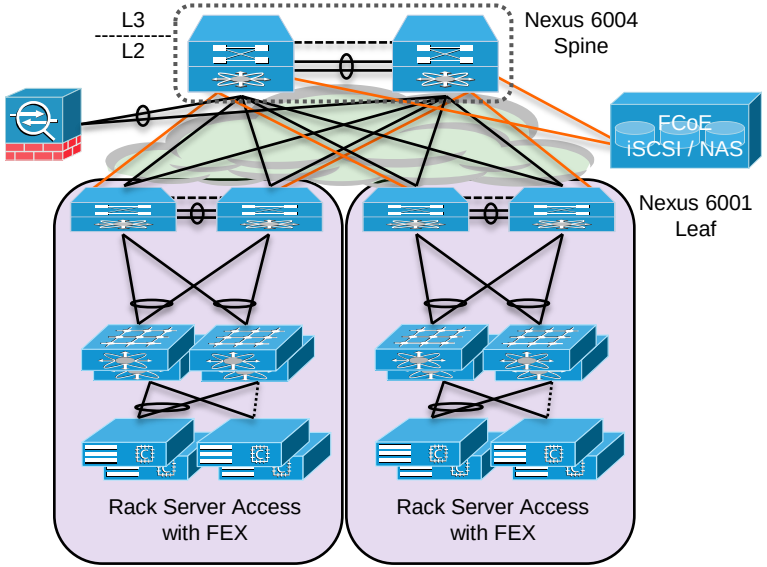

Dual Layer Nexus 7004/7706 Data Centre

High Availability switching fabric with chassis-based Spine

Data Centre switching control plane distributed over Dual Layers.

- Aggregation (Spine): Layer-3 and services boundary, scale point to expand fabric.
- Access (Leaf): Physical TOR or FEX aggregation point, Layer-2 virtualisation services.

FCoE support on dedicated links and VDC.

Example Components:

- 2 x Nexus 7004/7706, 4 x Nexus 5000/6000
- Layer-3, VDC and Storage Licensing
- 24 x Nexus 2232PP/TM-E

Enable FabricPath between tiers for configuration simplicity and future scale.

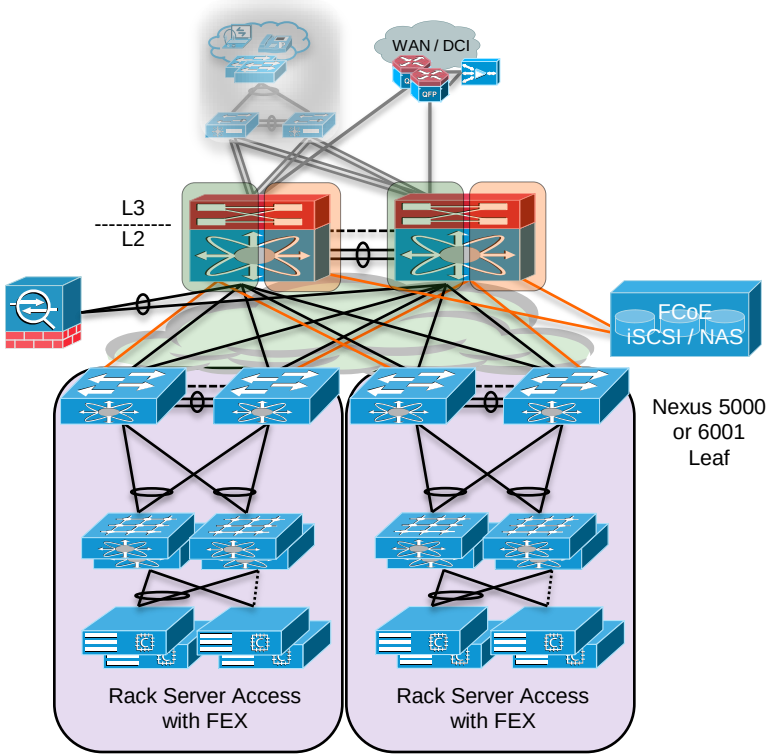

Dual Layer Nexus 7000/6000 Fabric with VDCs

Virtual Device Contexts partitioning the physical switch

Nexus 7009/7706 FabricPath Spine

- Highly Available design with dual-supervisor
- Add leaf pairs for greater end node connectivity
- Add spine nodes for greater fabric scale and HA
- FCoE support over dedicated links and VDC

Nexus 7000 Series Benefits:

- Integrated DCI support with OTV, LISP, and MPLS
- Feature-rich switching fabric with VDC's, FEX, vPC, FabricPath, FCoE
- Nexus 7000 Service Module capability starting with Network Analysis Module (6.2/Freetown)
- Investment protection of a chassis-based solution

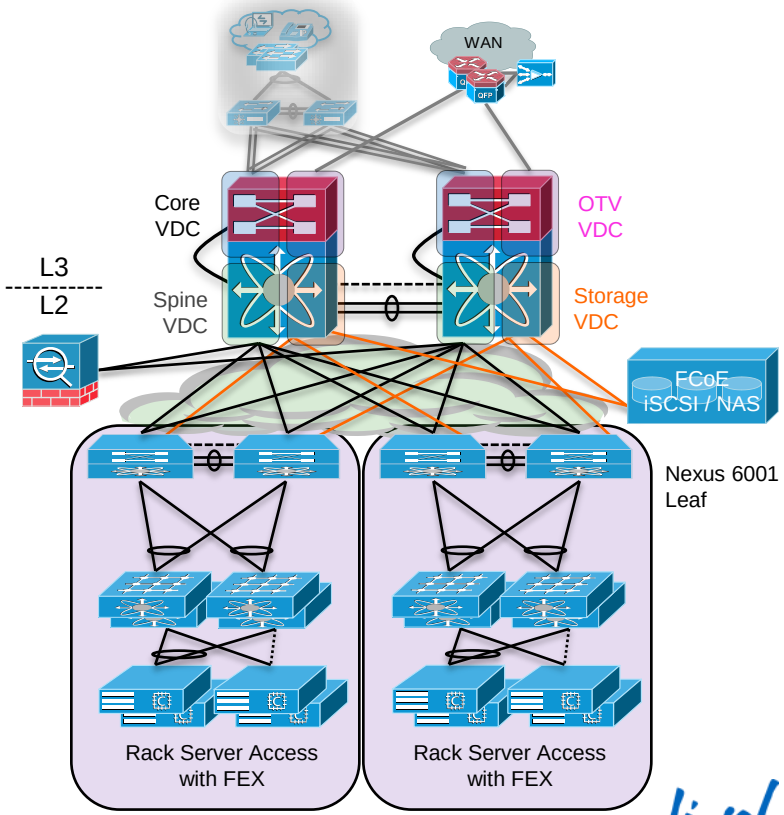

FabricPath with vPC+ Best Practices Summary

- Manually assign physical and emulated switch ID's to easily identify switches for operational support.
- Configure all leaf switches with STP root priority, or use pseudo-priority to control STP.
- Ensure all access VLANs are "mode fabricpath" to allow forwarding over the vPC+ peer-link which is a FabricPath link.
- Use vPC+ at the Leaf for port channels, and also at the Layer-2/3 spine to provide active/active HSRP.
- Set FabricPath root-priority on the Spine switches for multi-destination trees
- Convergence optimisations:
  - ✓ Set linkup-delay timer to 60 seconds
  - ✓ Set isis spf-interval 50 50 50, lsp-gen-interval 50 50 50

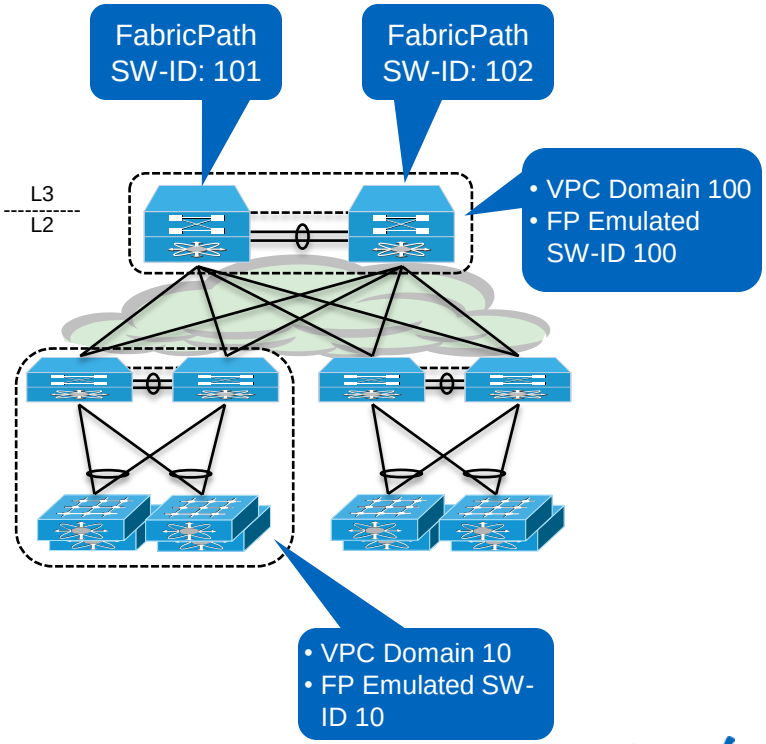

Evolution Of Fabric

Spanning-Tree → vPC → FabricPath → DFA

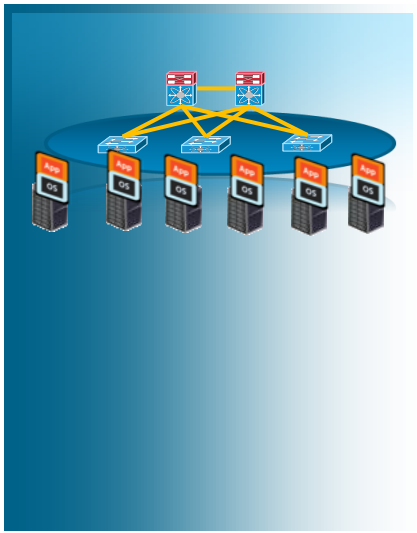

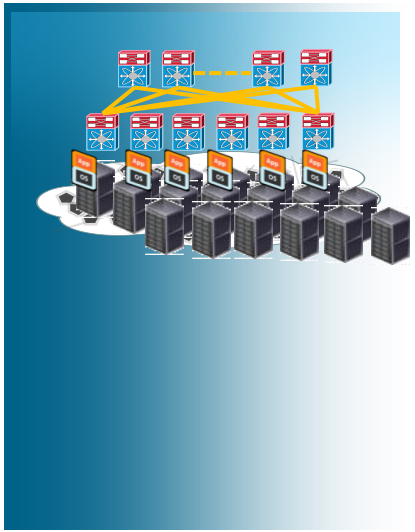

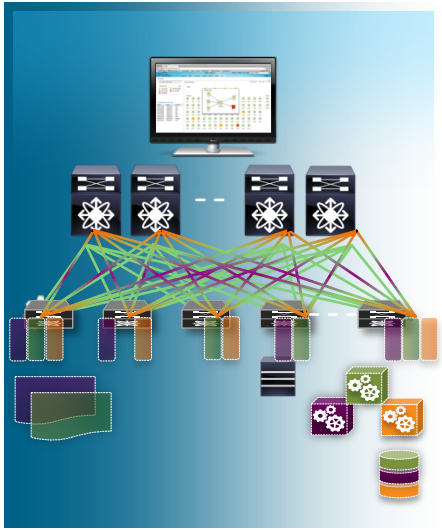

Layer 2 Scalability

Infrastructure Virtualisation and Capacity

Dynamic Fabric Automation (DFA)

Evolution Of The Unified Fabric

SIMPLIFY

AUTOMATE

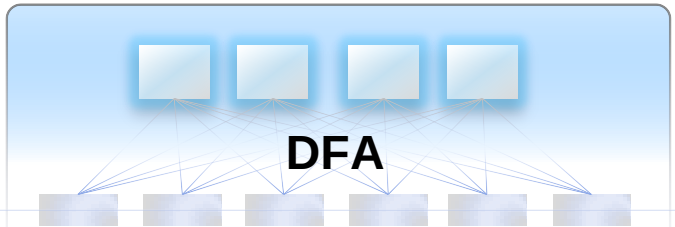

OPTIMISE

SCALE

CPOM

Integrate With Existing Nexus Infrastructure

Physical And Virtual Integration

Auto Detect Mis-Cabling

Fabric FCoE

Optimisation of Network Traffic

Innovation Design For Scalable Fabric

For More information: BRKDCT-2385 - Cisco Dynamic Fabric Automation Architecture

Pop Quiz

Which Cisco Layer 2 Multi-Pathing protocol allows for the easiest configuration and graceful growth when adding in new switches?

Pop Quiz

Name the design which FabricPath is based on:
Bonus point – name the elements of the design

Summary: Scalable Midsize Data Centre Designs

- Midsize Data Centres can benefit from the same technology advances as larger ones.
- Increase the stability of larger Layer-2 workload domains using vPC, FabricPath, and vPC+.
- Start with a structured approach that allows modular design as requirements grow.
- Evaluate Nexus switching options based on feature support, scale, and performance.
- Plan ahead for re-use of components in new roles as needs change.

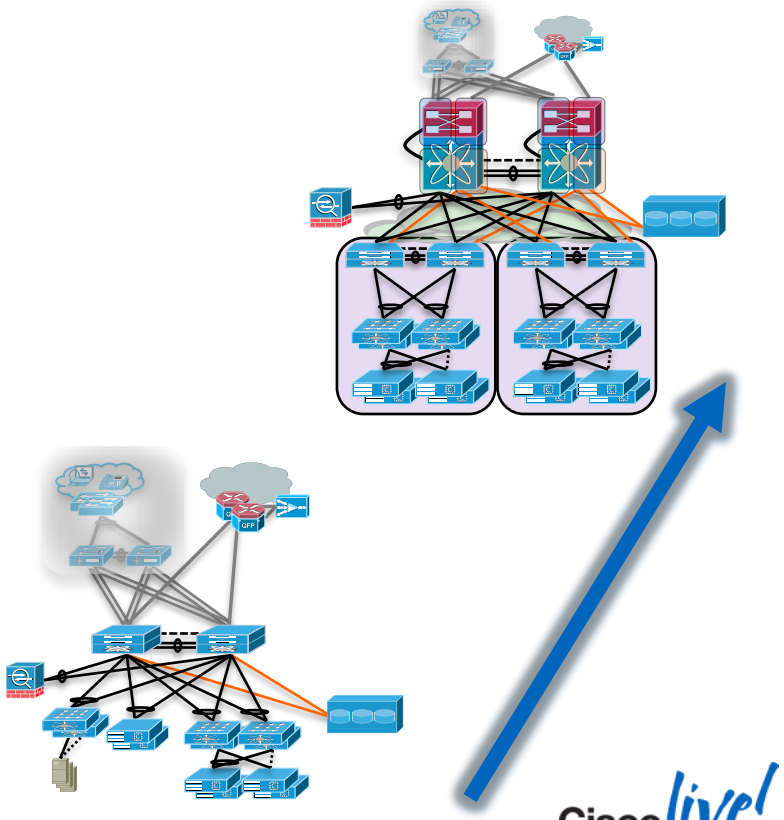

Other Cisco Live Presentations to Attend

- BRKDCT-2048 Deploying Virtual Port Channel in NXOS
- BRKCOM-1001 UCS Fabric Fundamentals
- BRKCOM-1005 UCS Architecture Overview
- BRKCOM-2001 UCS Management Deep Dive
- BRKCOM-2003 UCS Networking - Deep Dive
- BRKVIR-2012 Inside the Nexus 1000v
- BRKSEC-2205 Security and Virtualisation in the Data
- BRKDCT-2049 Overlay Transport Virtualisation
- BRKARC-3470 Cisco Nexus 7000 Switch Architecture
- BRKCOM-2640 UCS C-Series Deployment Options, Best Practice and UCSM Integration

Other Cisco Live Presentations to Attend

- BRKDCT-2081 Cisco FabricPath Technology and Advanced Fabric Architectures
- BRKDCT-2615 How To Achieve True Active-Active Data centre Infrastructures
- BRKAPP-9000 Introduction To Application Centric Infrastructure
- BRKCRT-8003 CCIE Data centre
- BRKDCT-2385 - Cisco Dynamic Fabric Automation Architecture




Q & A

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2014 Polo 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 www.ciscoliveaustralia.com/mobile
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 21 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
TM