Cisco Connect

Praha, hotel Clarion
10. – 11. dubna 2013

Kompaktní a přesto agilní datové centrum

MM2 / L2

Tomáš Horák, Systems Engineer, Cisco
Luděk Šafář, Systems Engineer Manager, EMC

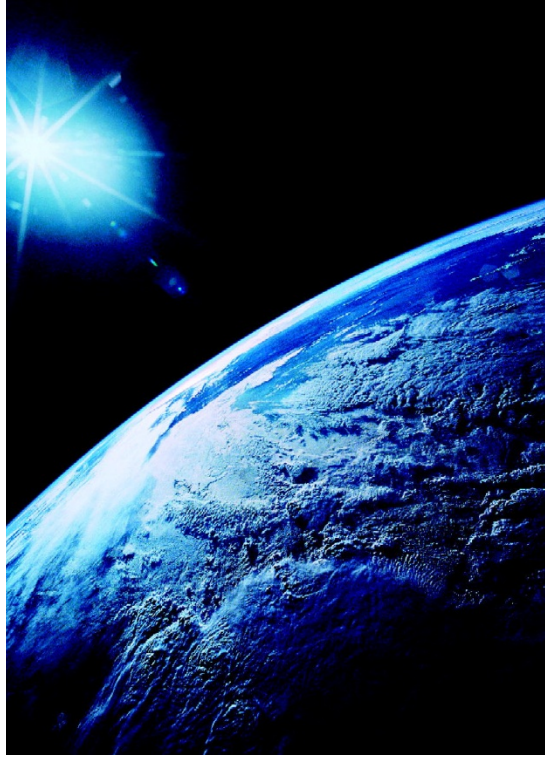

Agenda

- Data Center networking requirements
  - Characteristics of network design unique to the Data Center
- Data Center networking features and solutions
  - Elements of Cisco feature innovations, and progress in the standards bodies
- Reference Topologies
  - Entry-level models: sub-100 to 300 server ports
- EMC Storage Solution
  - VSPEX
- Conclusion

Data Center networking requirements

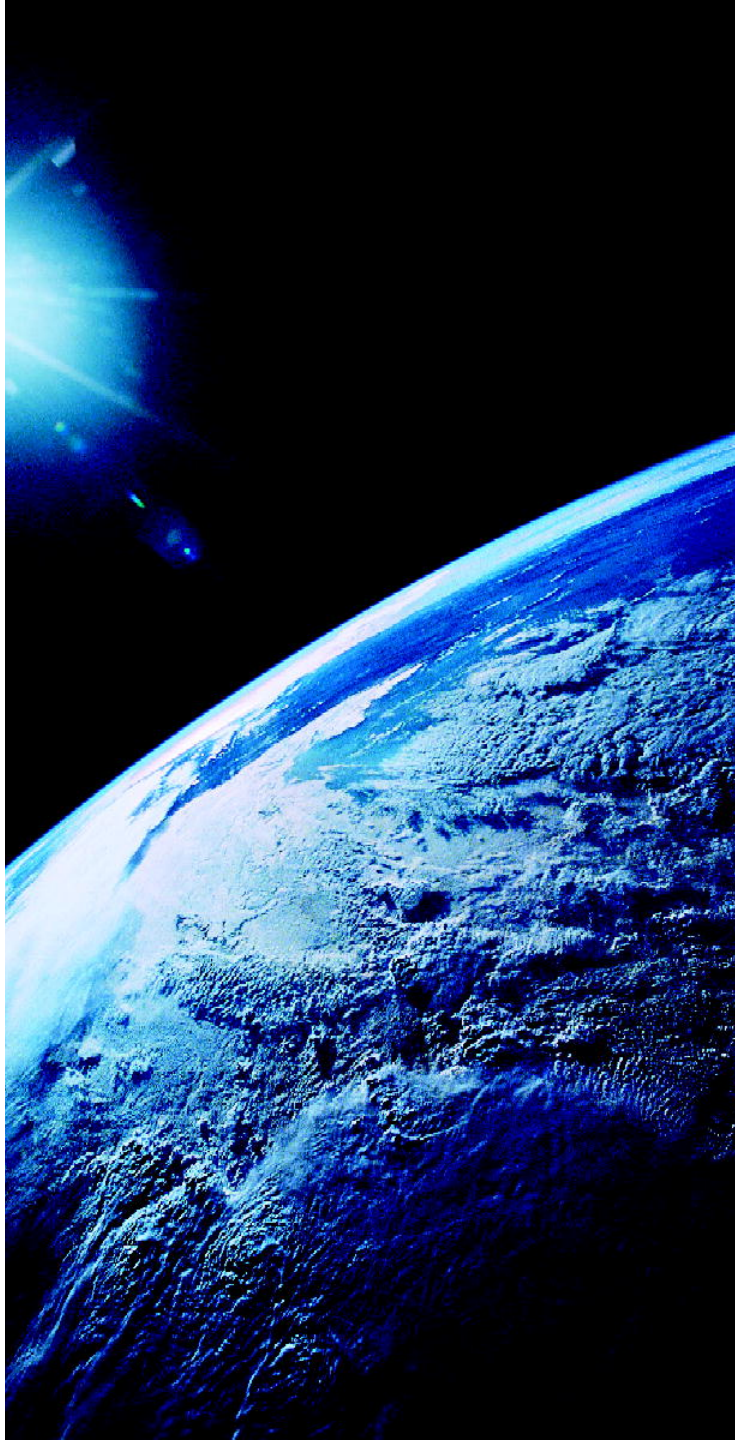

Traditional Data Center Conceptual Layers

Classic Enterprise Model terminology for reference

- **Core**: Connection to the "rest of the network" at a main location
- **Aggregation**: Boundary between Layer-2 and Layer-3 switching; servers IP default gateway lives here.
- **Access**: Provides physical connections for physical servers (or now logical ports for virtual servers)
- Mainly North/South traffic patterns

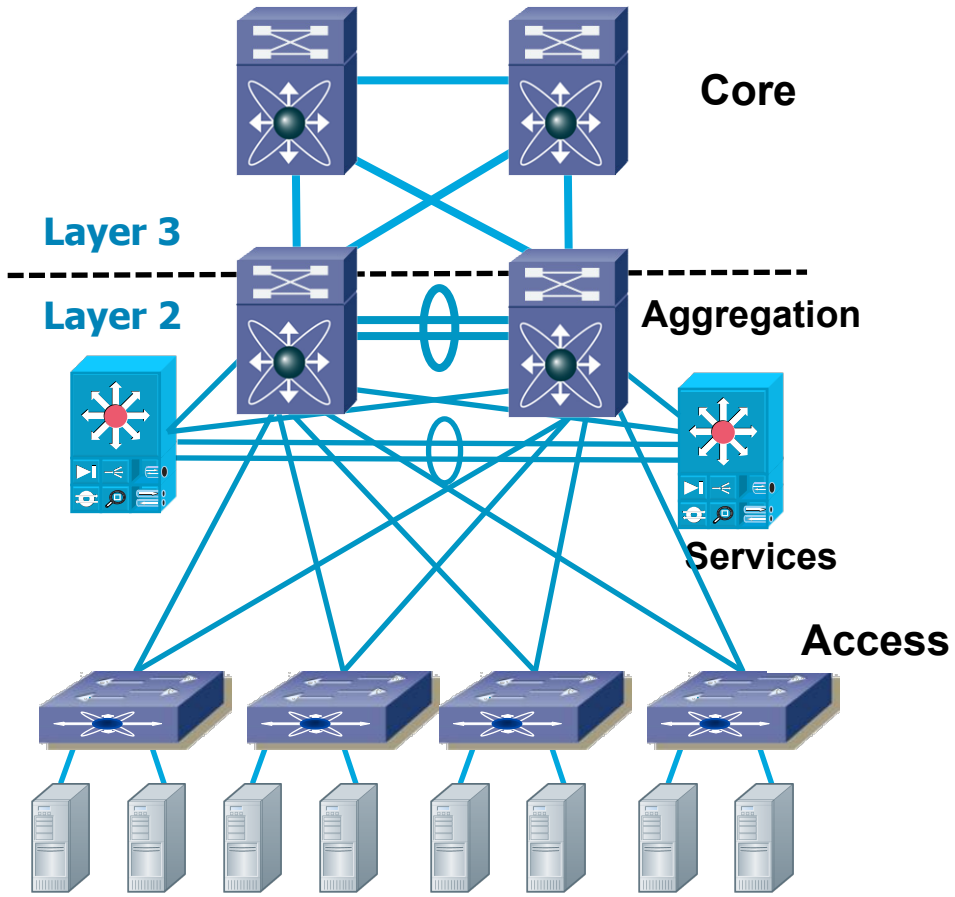

Hypervisor-Based Server Virtualization

De-coupling applications from physical hardware

Virtualization benefits

- Better hardware utilization and application availability
- Workload Mobility, within or across DC's
- Reduced provisioning time/effort

Increasing demands on the network:

- Larger Layer-2 domains
- Focus on shared storage systems
- L2 extension between physical DC's
- Software switching redefining the network edge

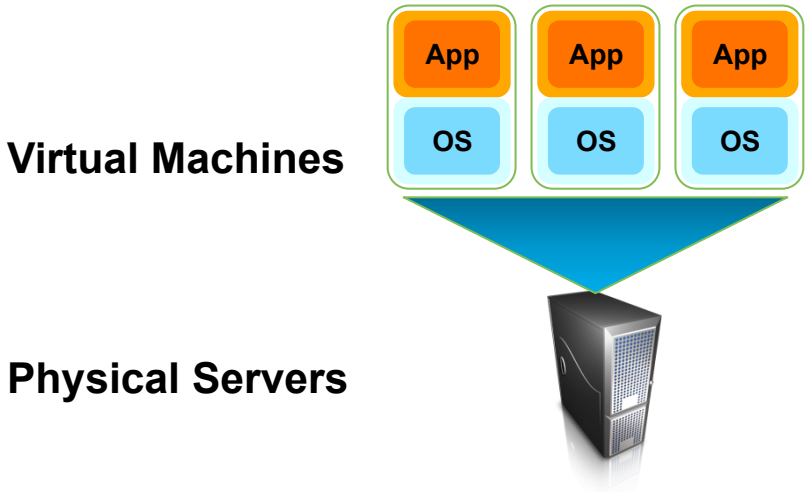

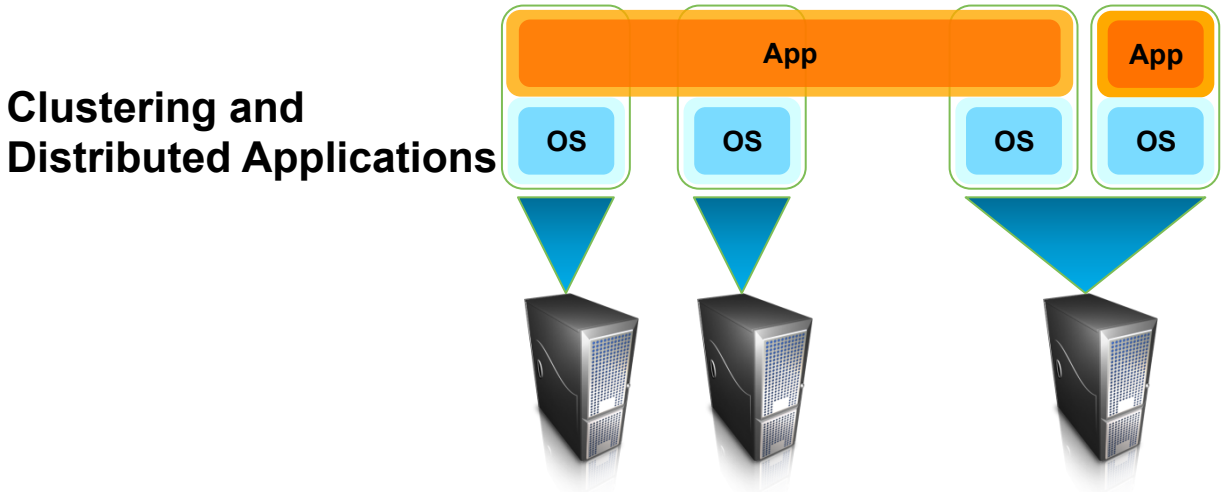

Virtualization and Software Switching

Redefining the “network edge”

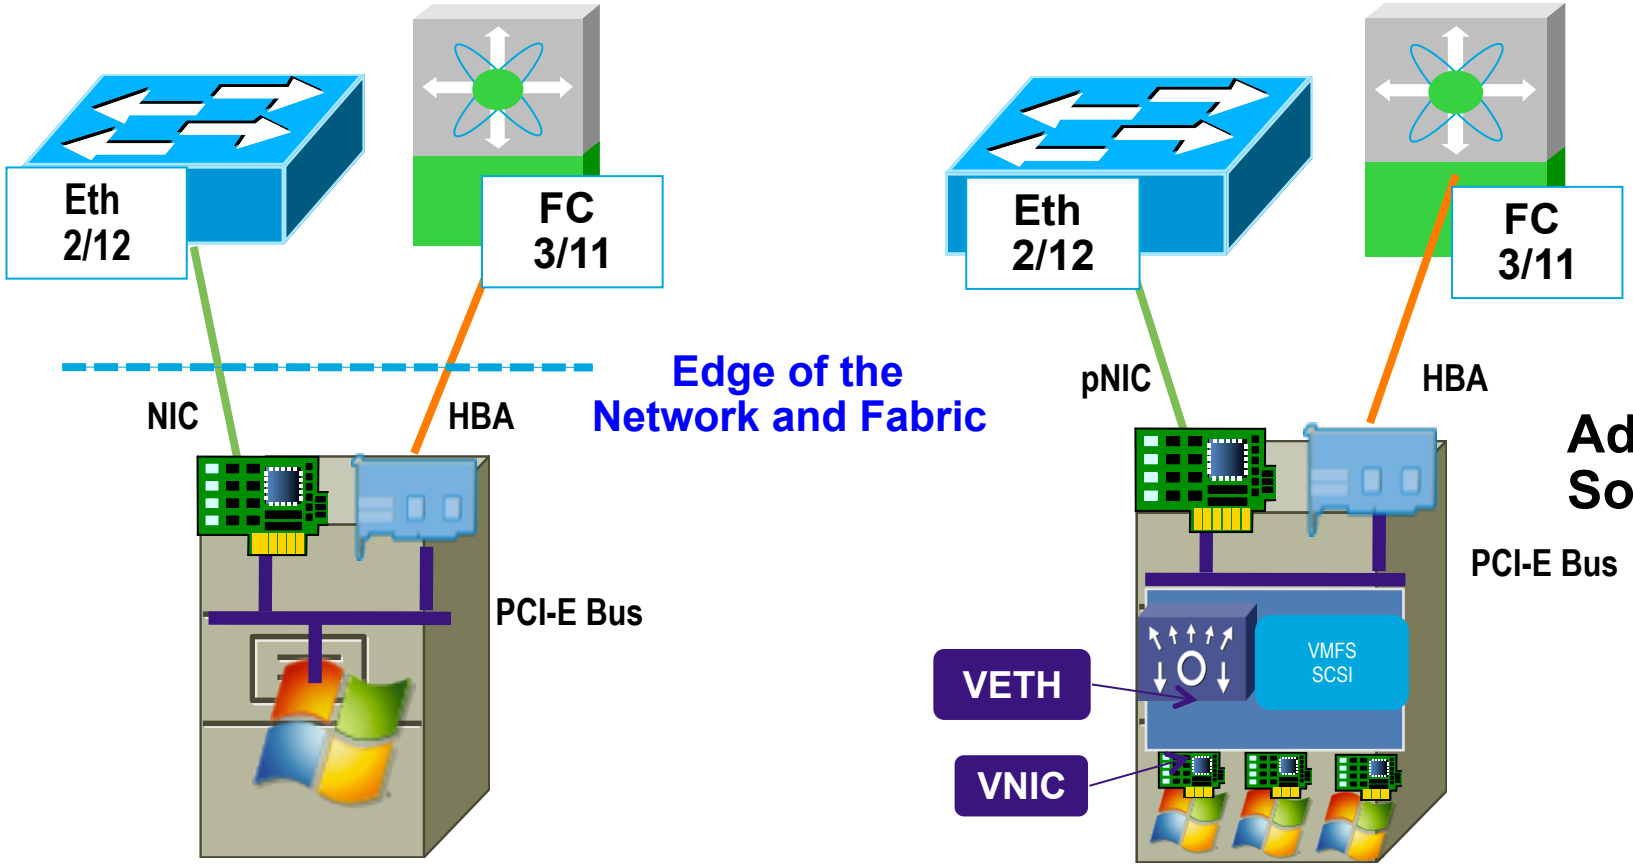

Traditional Model

Add Hypervisor Software Switch

Operating System
and Device Drivers

Hypervisor provides
virtualization of PCI-E resources

Compute and Fabric Edge are Merging

Virtualization and Software Switching

Redefining the "network edge"

Add CNA for virtual interface density

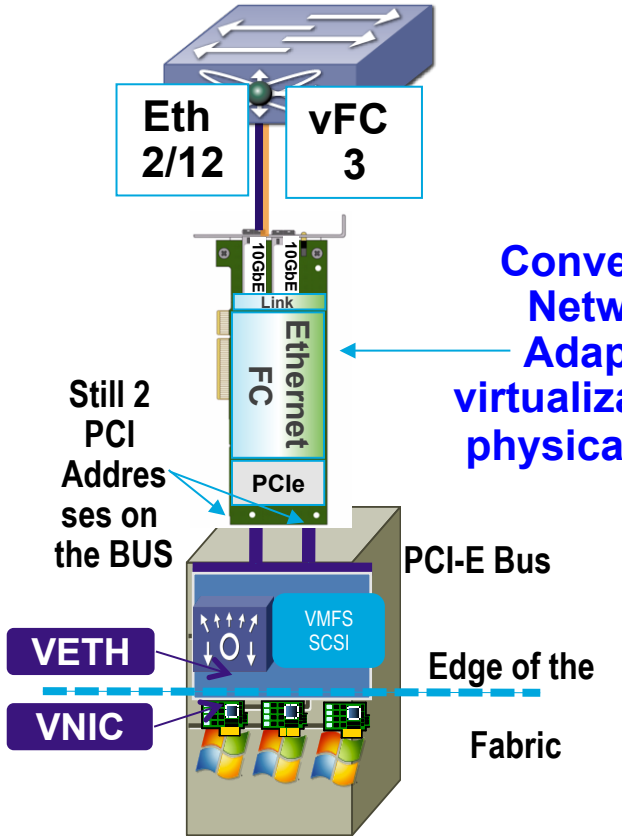

Converged Network Adapter: virtualization of physical Media

Move switching function back upstream

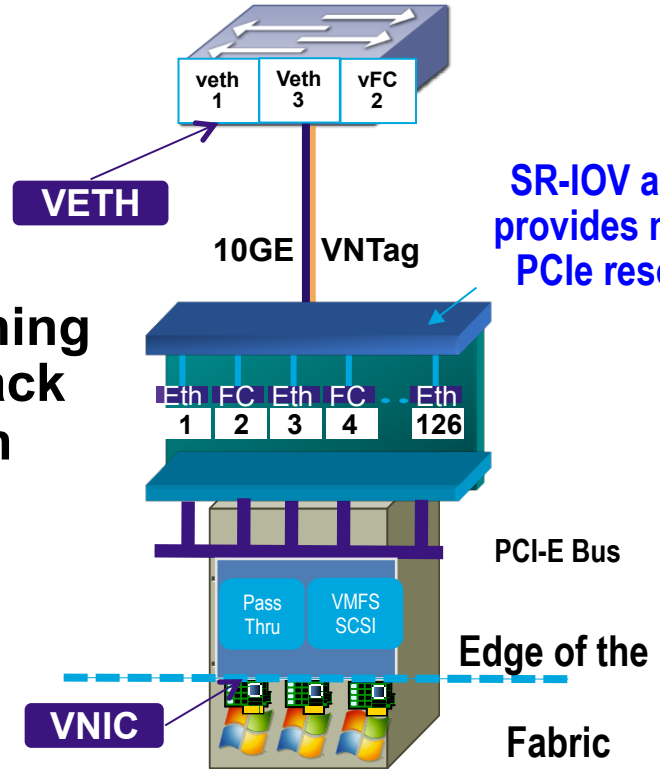

SR-IOV adapter provides multiple PCIe resources

Compute and Fabric Edge are Merging

Storage Networking

From separate networks to a converged fabric

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

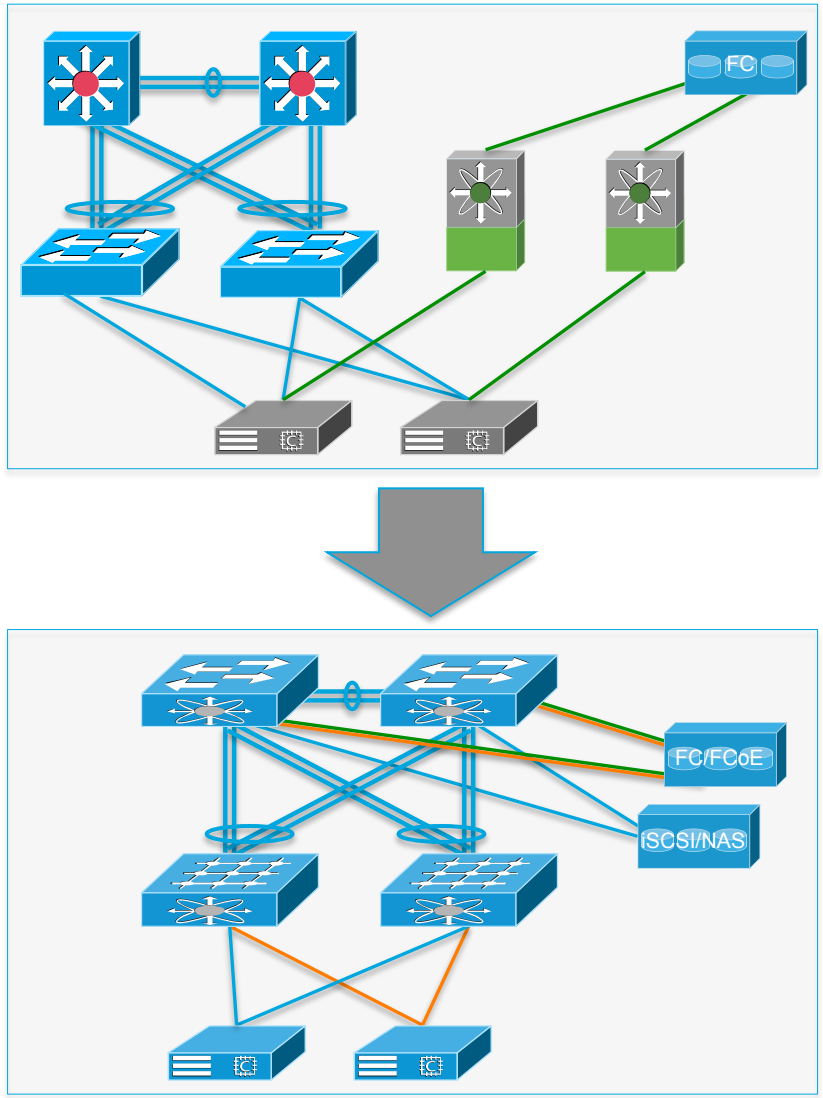

The Goal of a Unified Data Center Fabric

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’

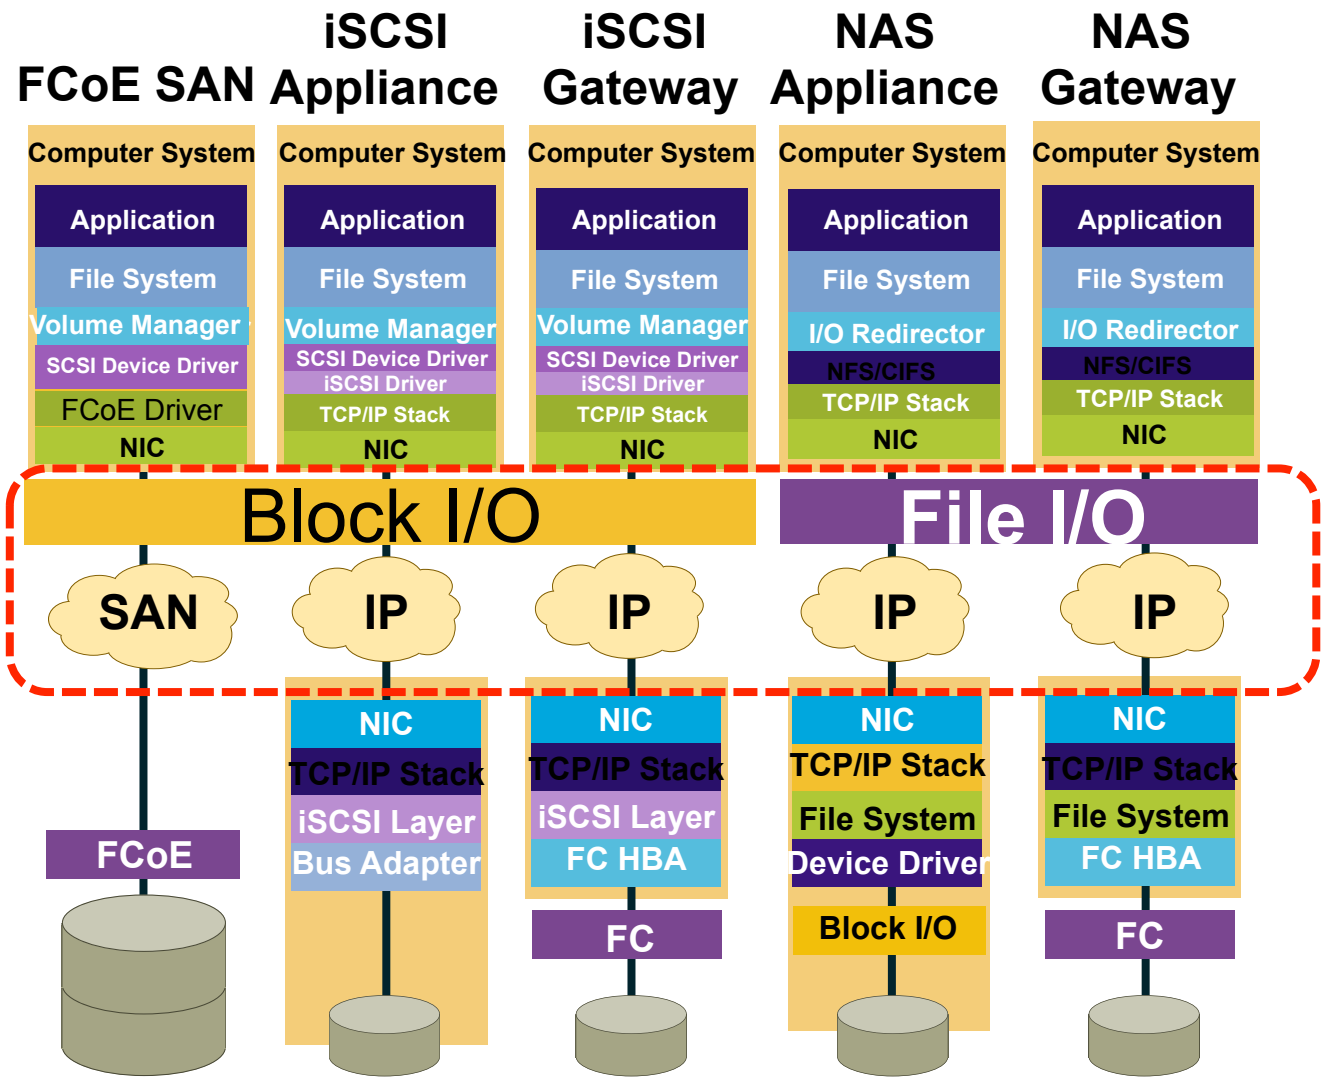

Data Center Network Services

Integrating service functions into a virtualized fabric

- In the non-virtualized model, services were inserted into the Data Path at network 'choke points'
- The Logical Topology matched the Physical Topology
- Migration to higher server speeds push the limits of physical service appliances
- Virtualized workloads may require a re-evaluation of where the services are applied and how they are scaled

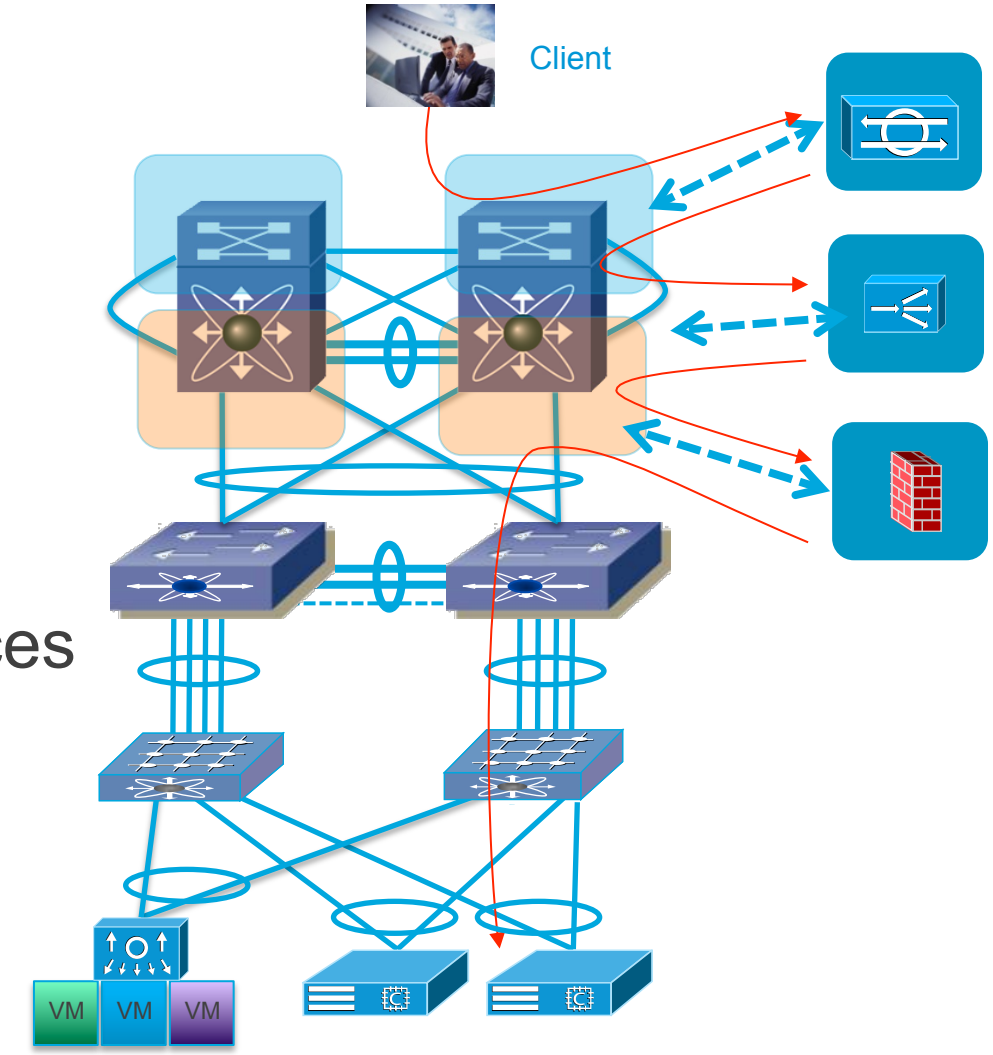

Data Center Interconnect

Active/Active Workloads and Disaster Recovery

- Virtualization and workload mobility open up greater flexibility in using multiple locations
- Active/Active workload sharing can be constrained by application latency requirements and storage replication
- Managing Layer-2 extension and Layer-3 address mobility presents many of the same challenges facing larger customers
- Advanced networking features are required to address these needs.

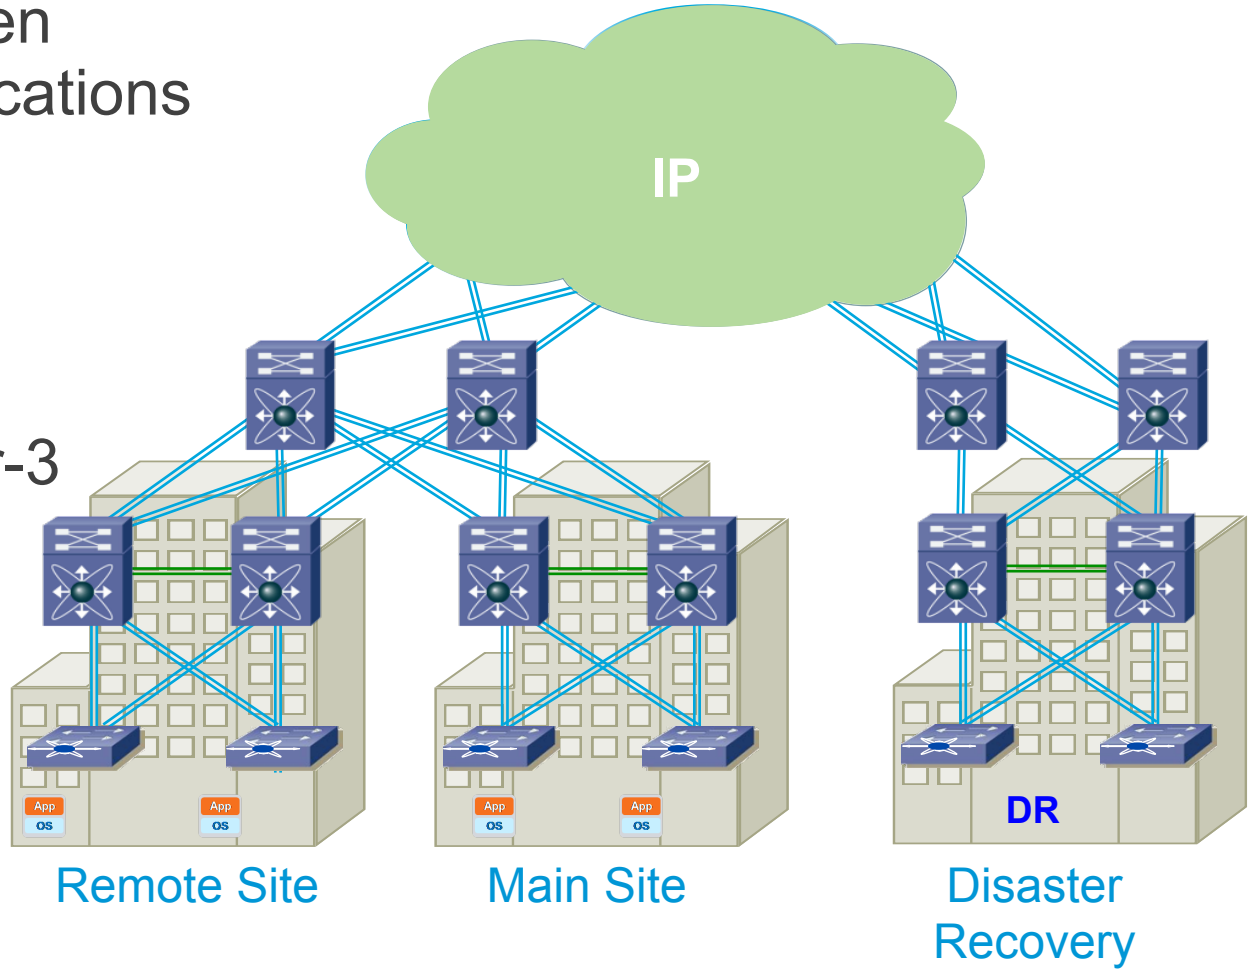

Goals for Small and Medium DC Design

- Take advantage of current data center feature sets while targeting a smaller port count
- Consolidate conceptual layers into fewer physical devices where possible
- Control costs through platform selection and virtualized switching
- Consider future needs and design to allow the network to scale as business requirements grow

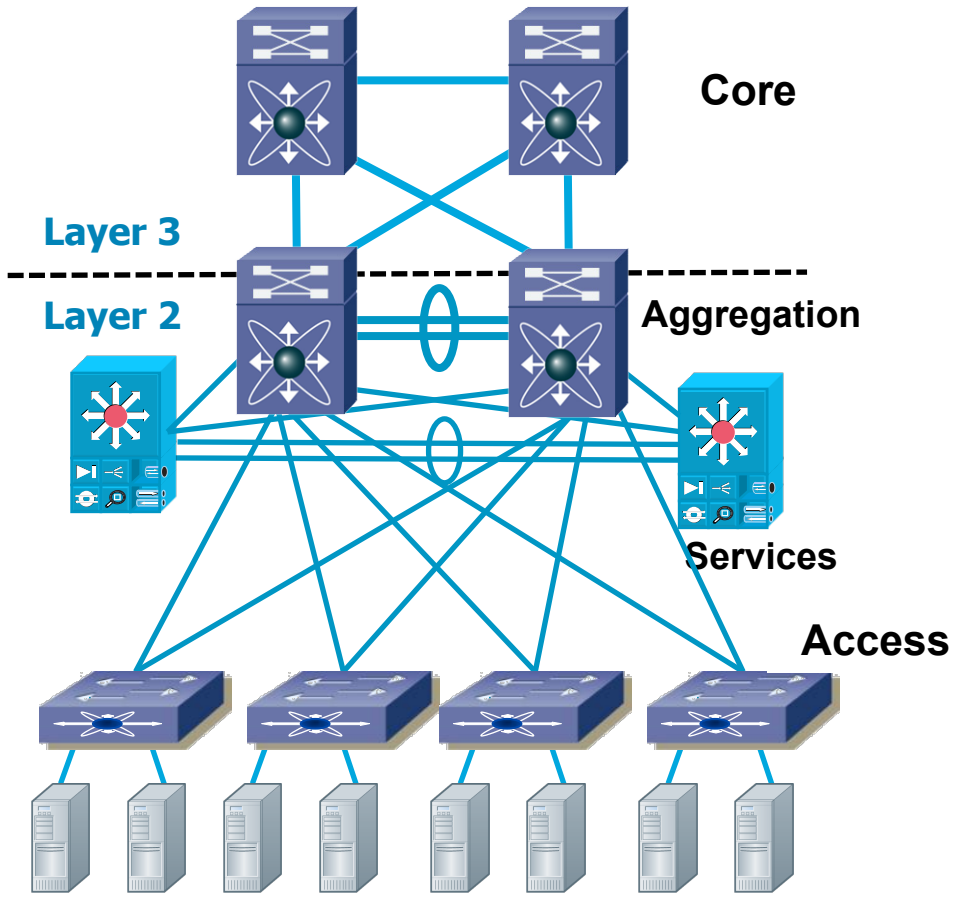

Data Center networking features and solutions

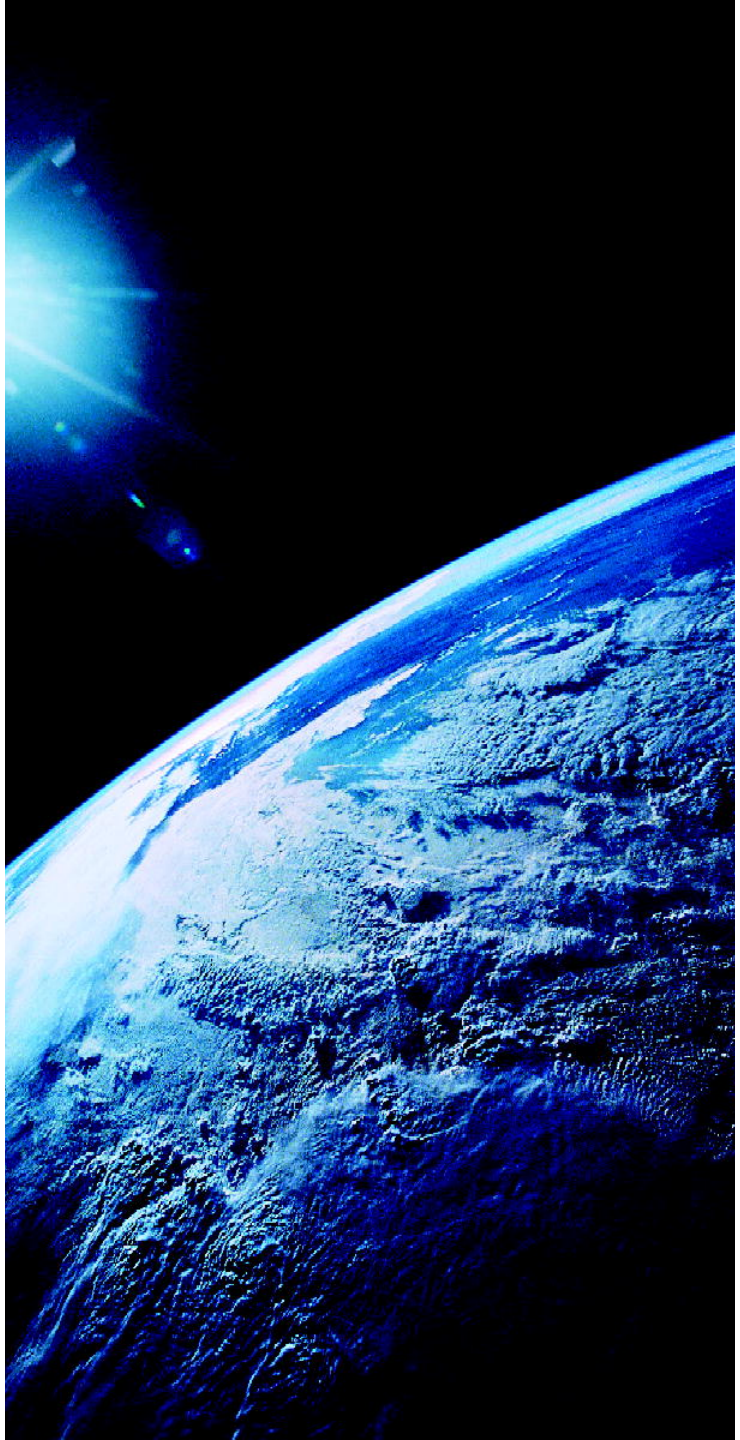

Virtual Port Channel – vPC

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
- vPC maintains independent control planes and switch management, switches are “vPC peers”

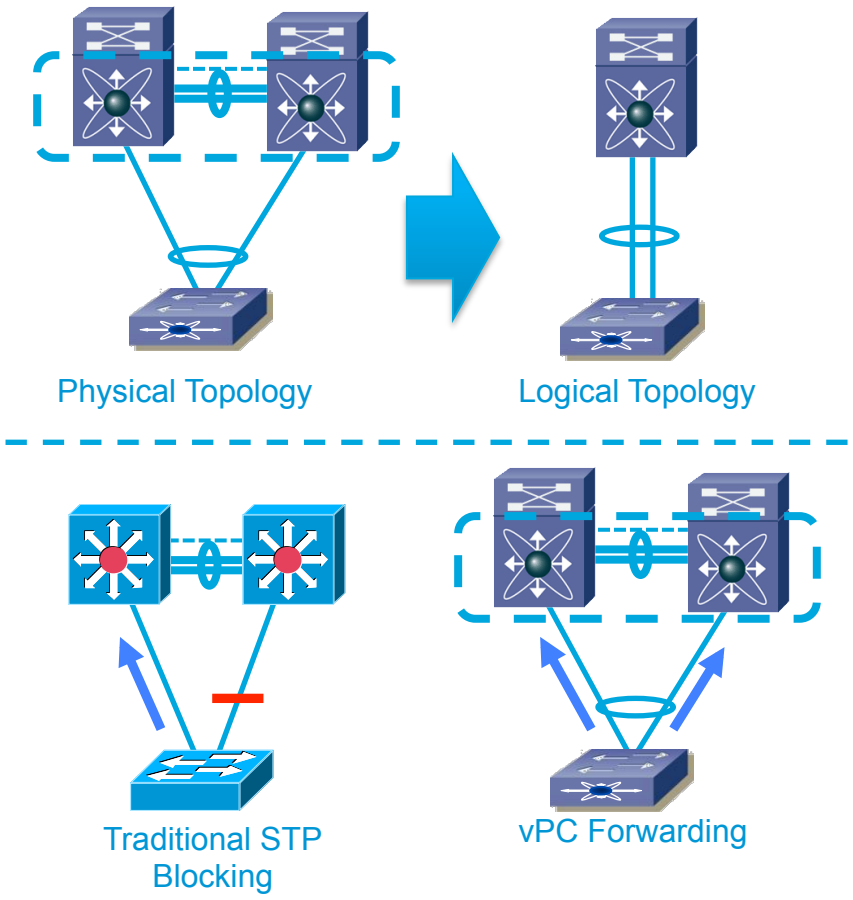

Fibre Channel over Ethernet and vPC

SAN “A” / “B” Support

- LAN and SAN utilize different High Availability Models
- SAN is dual fabric, LAN is fully meshed fabric
- vPC enables ‘both’ architectures at the edge (single device models not acceptable to SAN engineers)
- Campus technologies VSS and Stackwise are not SAN-aware

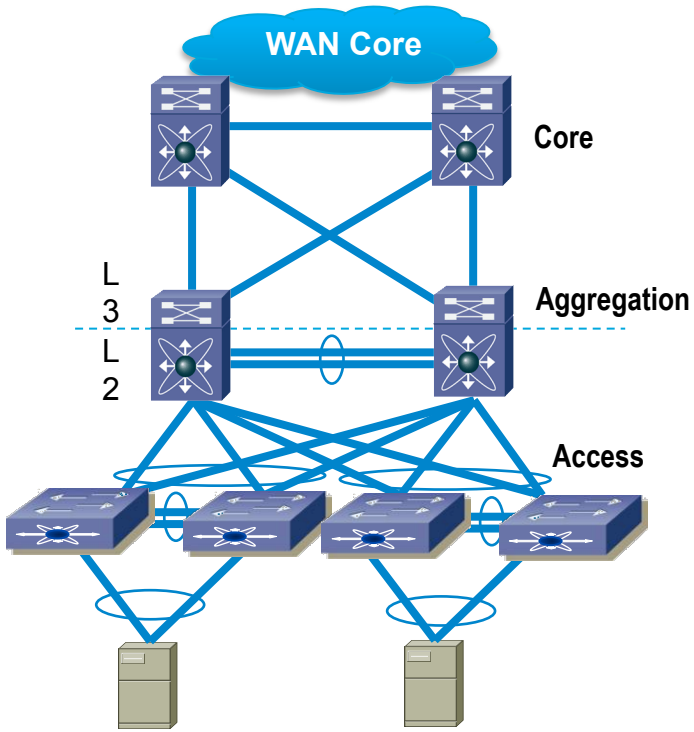

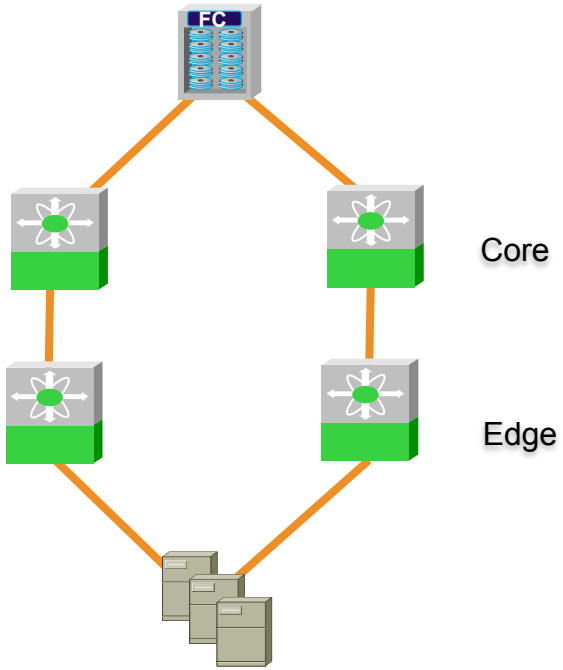

FEX: Virtualized Access Switching

Changing the boundaries of the Ethernet switch

- De-coupling of the Layer-1 and Layer-2 topologies
- Line Card Portability (N2K supported with Multiple Parent Switches – N5K, N6K, N7K, FI 6200)
- Unified access for any server (100M→1GE→10GE→ FCoE): Scalable Ethernet, unified fabric or virtualized servers
- Simplified Management Model
- Plug and play provisioning
- Centralized configuration

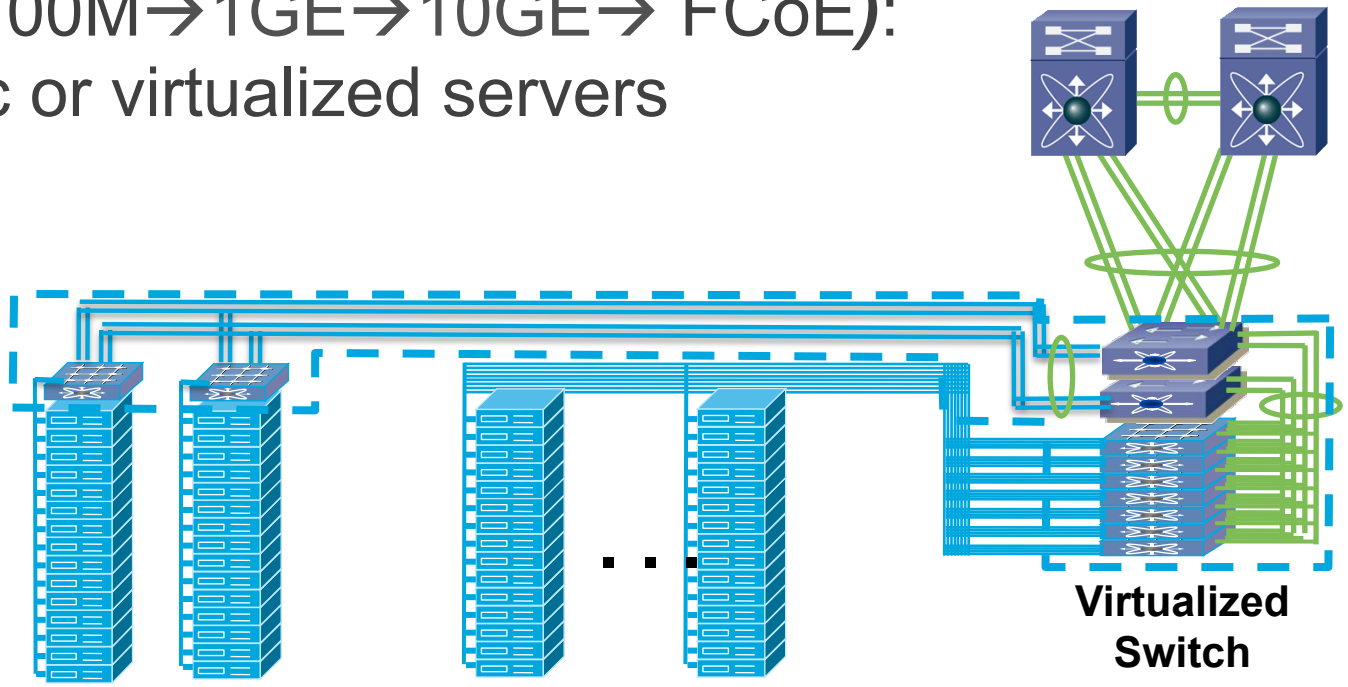

FEX and vPC Virtualized Access

Nexus 5000/5500 Topologies prior to 5.1.(3)N1

Straight Through (Same for Nexus 7000)

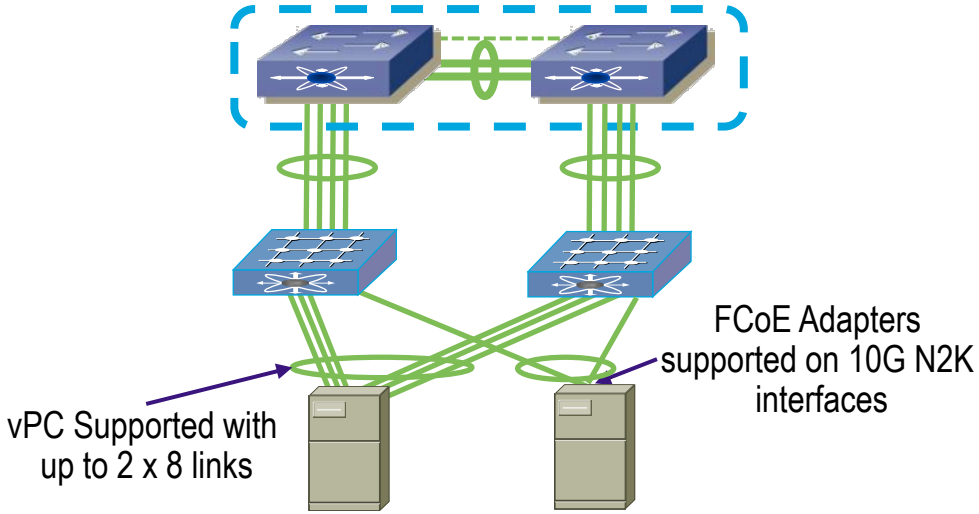

- Redundancy model – Dual Switch with redundant fabric
- Provides FCoE isolation for Storage topologies (SAN ‘A’ and ‘B’)
- Port Channel and Pinning supported for Fabric Link

Dual Homed

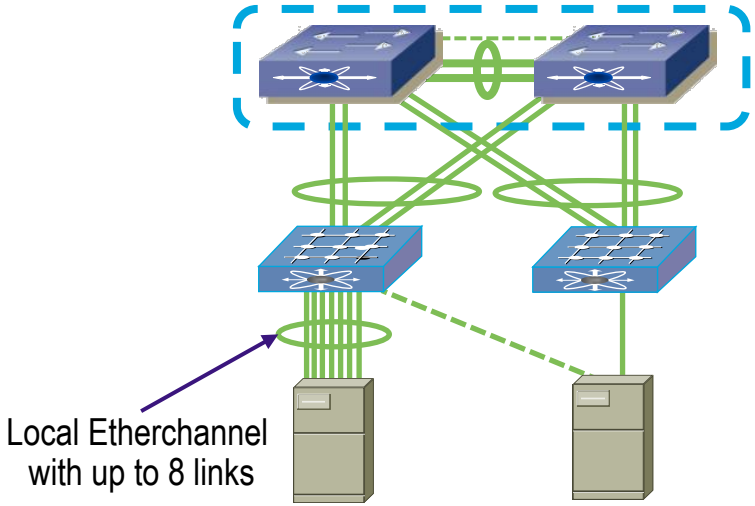

- Redundancy model – Single switch with dual ‘supervisor’ for fabric, data control & management planes
- No storage SAN ‘A’ and ‘B’ isolation
- No active/active dual-home was allowed

Enhanced Virtual Port Channel (EvPC)

- In an Enhanced vPC configuration any and all server NIC teaming configurations will be supported on any port.
- All components in the network path are fully redundant.
- Provides flexibility to mix all three server NIC configurations (single NIC, Active/Standby and NIC Port Channel)

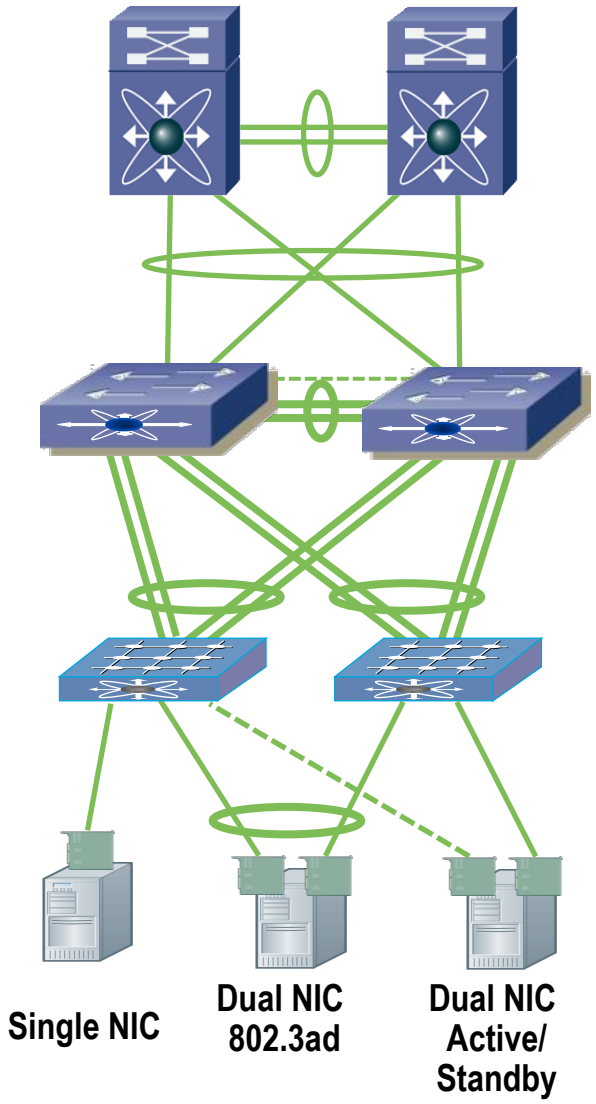

EvPC and FCoE SAN Traffic

Maintaining SAN A / B isolation

- In an Enhanced vPC (EvPC) SAN ‘A/B’ isolation is configured by associating each FEX with either SAN ‘A’ or SAN ‘B’ Nexus 5500
- FCoE & FIP traffic is forwarded only over the links connected to the specific parent switch
- Ethernet is hashed over all FEX fabric links

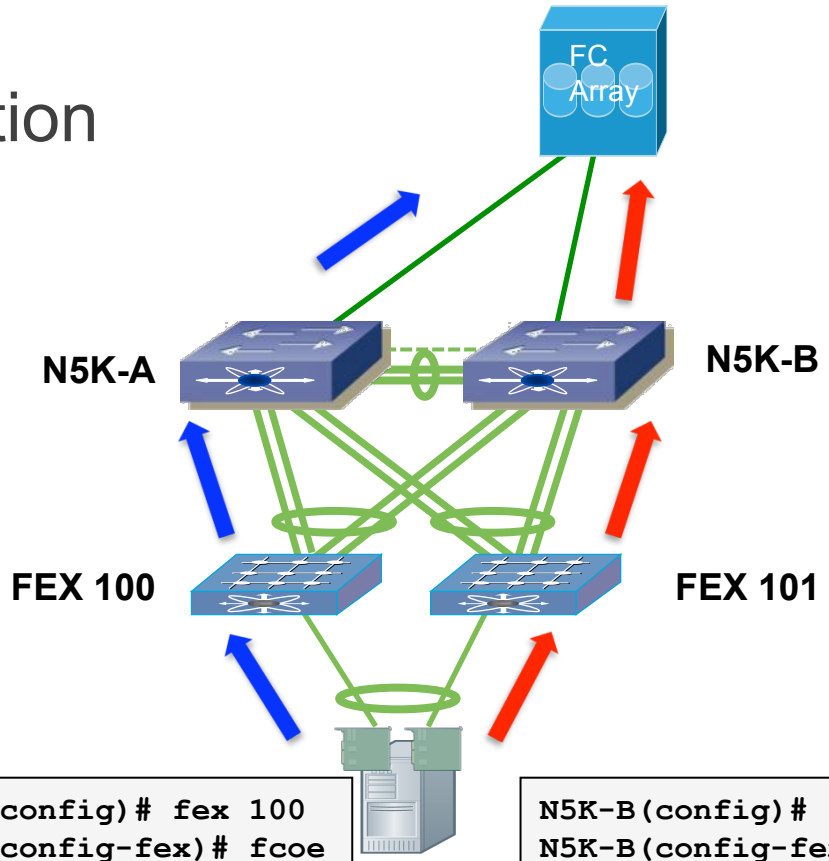


```
N5K-A(config)# fex 100
N5K-A(config-fex)# fcoe

N5K-A(config)# fex 101
```

```
N5K-B(config)# fex 101
N5K-B(config-fex)# fcoe

N5K-B(config)# fex 100
```

Fabric Extender Evolution

Cisco Innovation and 802.1BR

- The FEXLink Architecture provides the ability to extend the bridge (switch) interface to downstream devices
- FEXLink associates the Logical Interface (LIF) to a Virtual Interface (VIF)

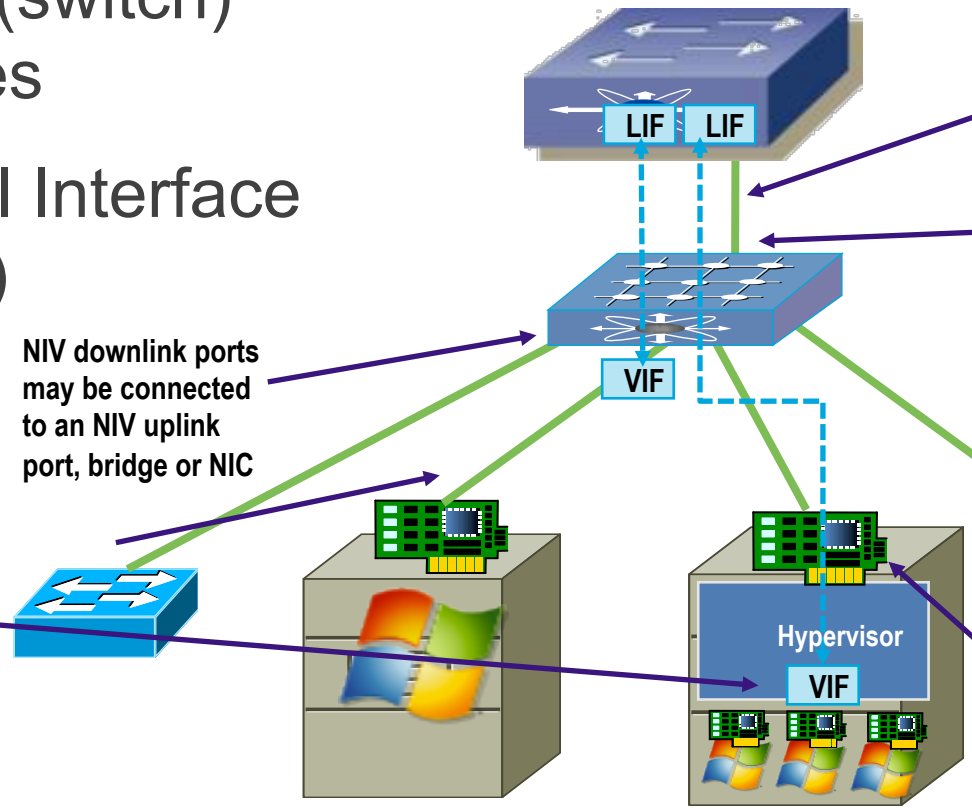

Bridges that support Interface Virtualization ports must support VNTag and the VIC protocol

NIV uplink ports must connect to an NIV capable bridge or an NIV Downlink

NIV downlink ports may be connected to an NIV uplink port, bridge or NIC

NIV may be cascaded extending the port extension one additional level

NIV downlink ports are assigned a virtual identifier (VIF) that corresponds to a virtual interface on the bridge and is used to forward frames through NIV's

NIV capable adapters with port extension

Note: Not All Designs Supported in the FEXLink Architecture Are Currently Implemented

FCoE and Unified Ports

Nexus 5500 flexible definition of port function

- 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 a FC SAN
- SAN “A” / “B” isolation is maintained across the network
- Unified Ports may be configured to support either native Fibre Channel or Ethernet

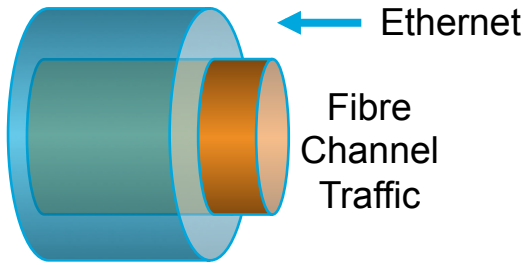

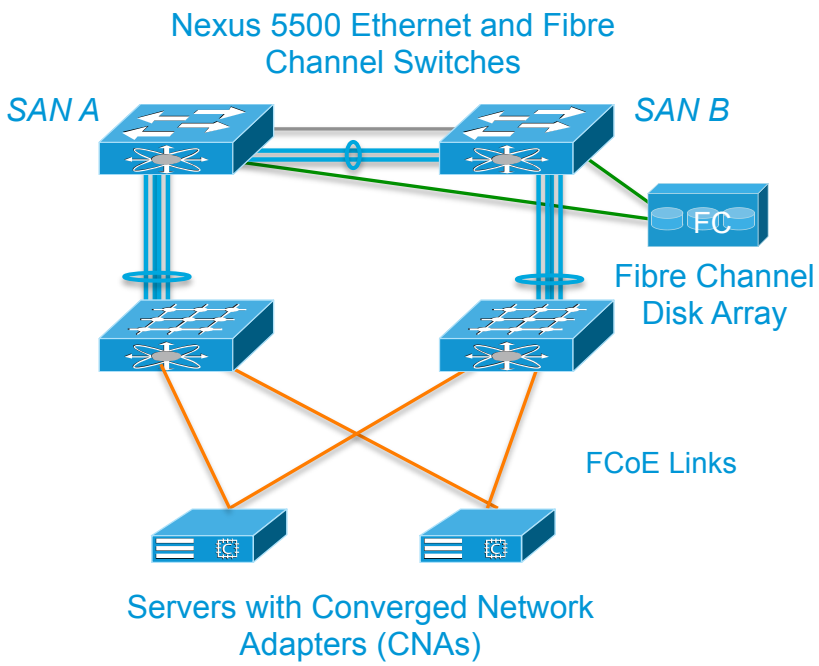

Cisco Adapter-FEX

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 5500
- Forwarding, Queuing, and Policy enforcement for vNIC traffic by Nexus 5500
- 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 5500

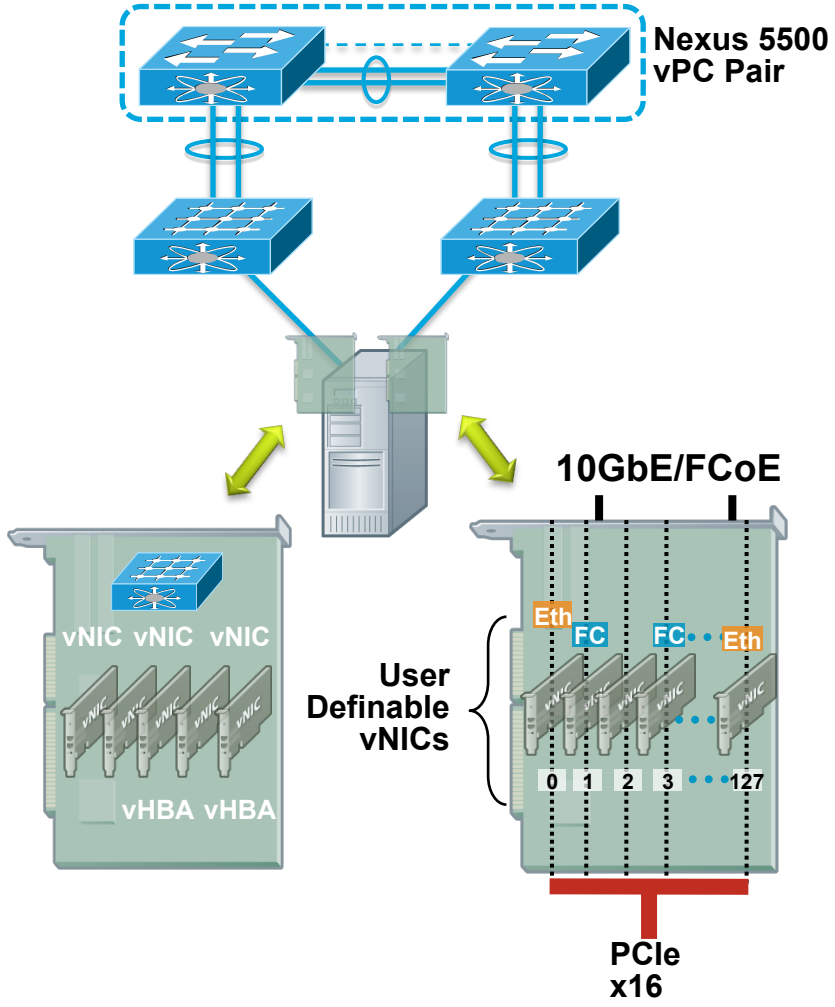

Cisco VM-FEX
Virtual Machine Fabric Extender

- Allows Nexus 5500 pair to register as a DVS in vCenter
- Extends Cisco Adapter FEX technology to the Virtual Machine with vMotion support
- Consolidated management of network, server, and virtual server interfaces
- Offload I/O processing from Server CPU with Cisco Nexus hardware performance
- Technology from the UCS Blade system extended to UCS Rack servers and Nexus 5500

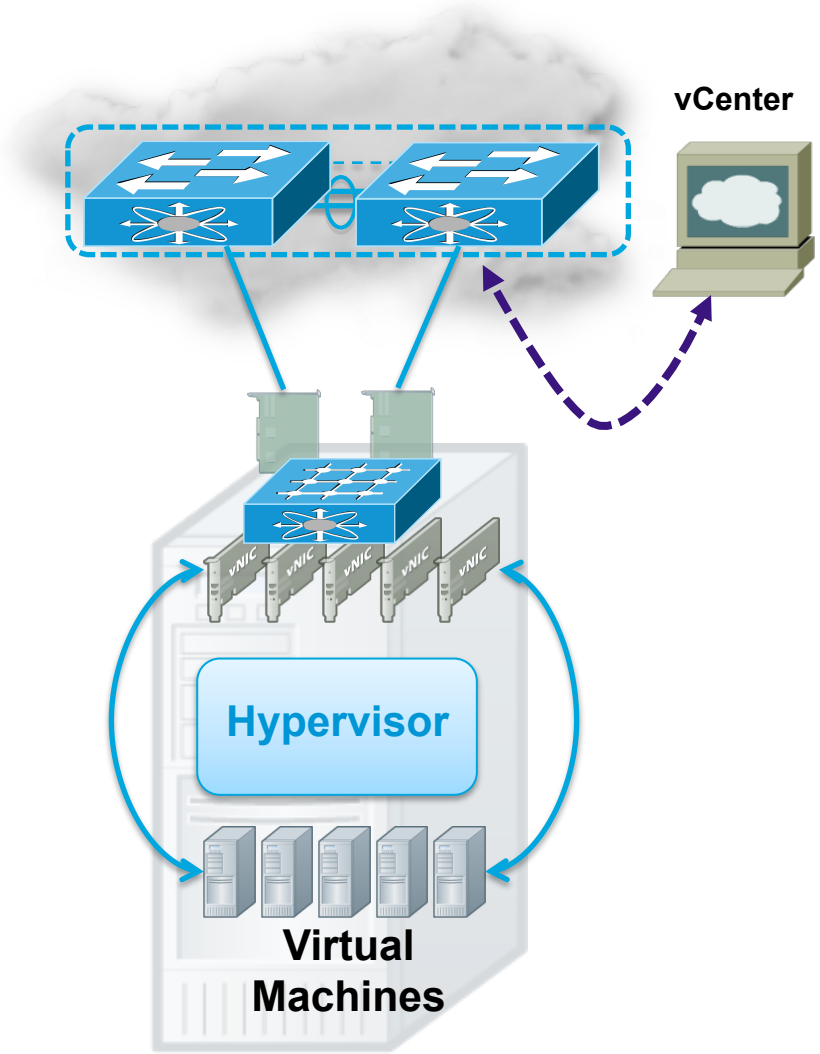

Nexus 1000v and vPATH

Traffic Interception for Data Center services

- Intelligence built into Virtual Ethernet Module (VEM) of Cisco Nexus 1000V virtual switch
- vPATH performs traffic interception and redirection for multiple virtual service nodes:
  - Virtual Security Gateway (VSG)
  - Virtual WAAS (vWAAS)
  - Virtual Network Analysis Module (vNAM)
  - Virtual ASA (vASA)
- vPATH is Multi-tenant Aware.
- Leveraging vPATH can enhance the service performance by moving the processing to hypervisor.

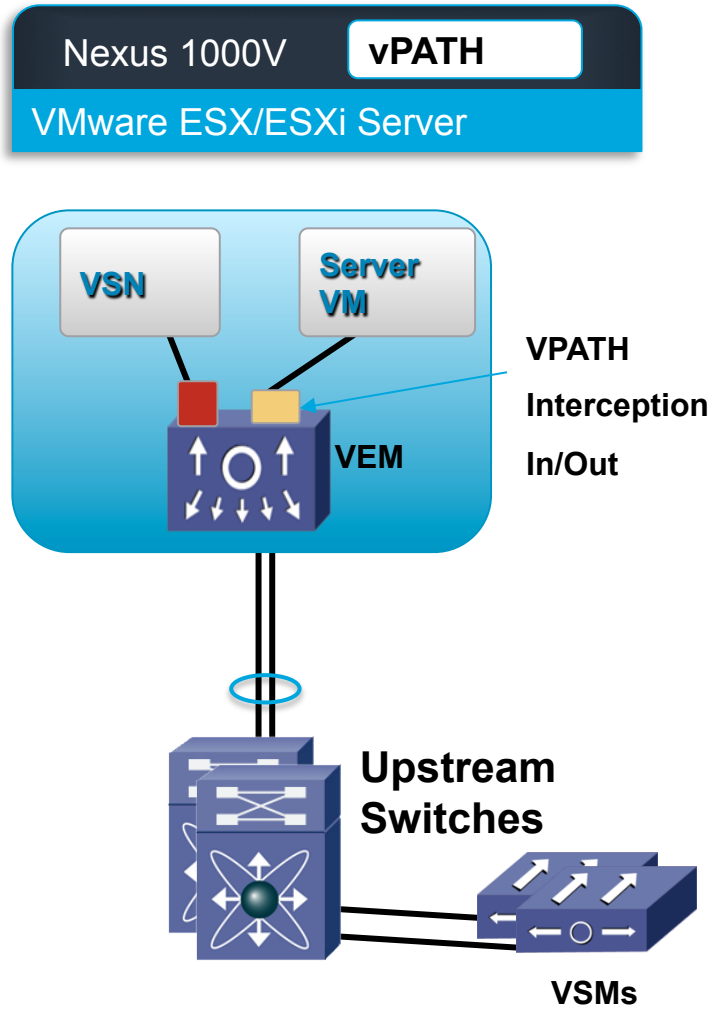

Cisco FabricPath

The best characteristics of both Layer-2, **and** Layer-3 switching

- **Easy Configuration**
- **Plug & Play**
- **Provisioning Flexibility**

Switching

- **Multi-pathing (ECMP)**
- **Fast Convergence**
- **Highly Scalable**

Routing

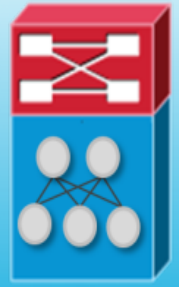

FabricPath

- FabricPath brings Layer 3 routing benefits to flexible Layer 2 bridged Ethernet networks.
- Key features such as vPC+ and multiple topologies make FabricPath a deployable solution; support will be extended to TRILL as it matures.

Building the Data Center Interconnect
Complementary Innovations

- FabricPath: Scalable Fabrics for Application Deployment Flexibility
- OTV : Layer 2 extensions over Layer 3 for Distributed Clustered Applications
- LISP: IP mobility, optimized routing and segmentation within the flexible Fabric

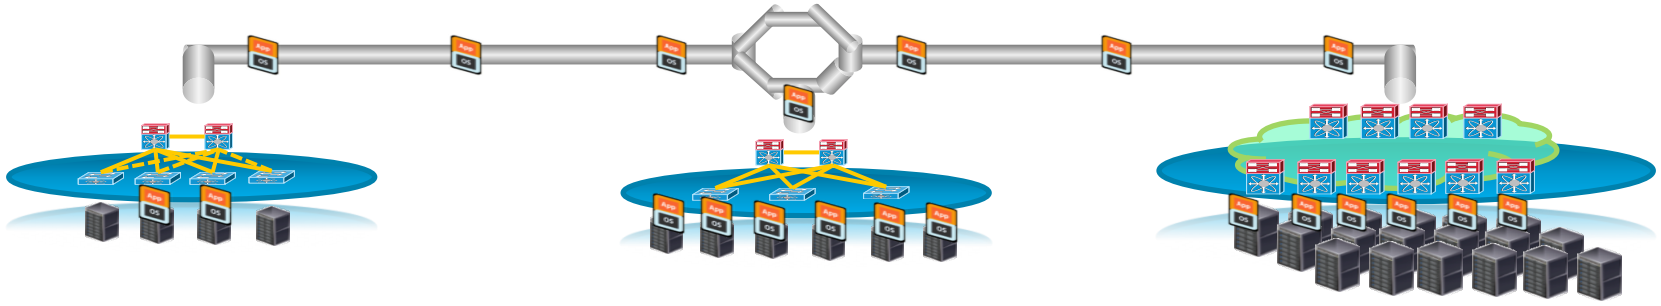

Classical Pod
Spanning Tree Protocol

Scalable Pod
vPC & FEX

Highly Scalable Pod
FabricPath

Technology White Papers
Available on cisco.com

For Your Reference

Building a Unified DC CVD

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/VMDC/2.2/design_guide/vmdcDesign22.html

Scaling Data Centers with Fabric Path

http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9402/white_paper_c11-605488.html

Extending L2 Networks with OTV

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DCI/whitepaper/DCI3_OTV_Intro_WP.pdf

Adapter FEX

http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/data_sheet_c78-657397.html

VM FEX

http://www.cisco.com/en/US/netsol/ns1124/index.html

Reference Topologies

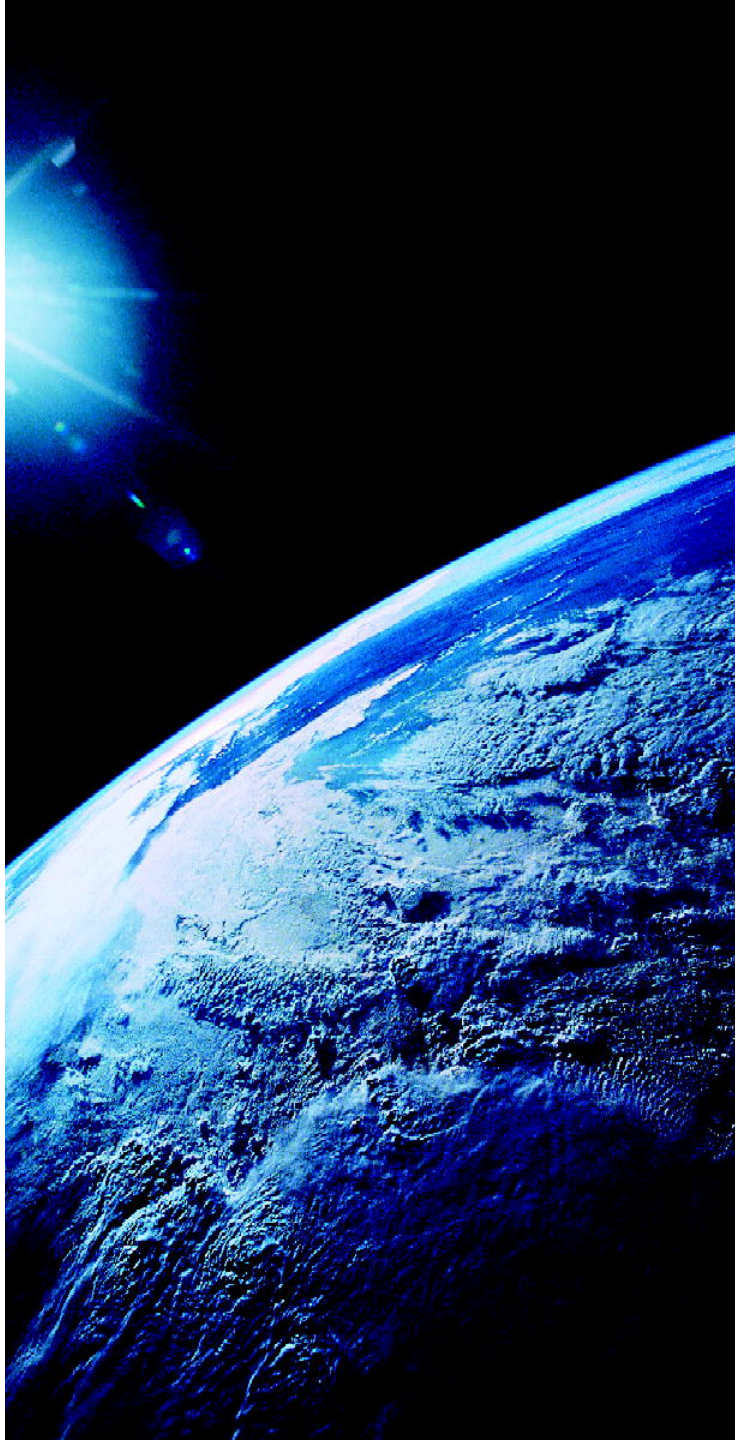

Cisco Nexus Switching Family

Building blocks for today's unified data center fabric

- **Nexus 7000 Series**: Chassis-based, highly scalable multilayer NX-OS switch
- **Nexus 6000 Series:** Modular 1-4RU high density 10G/40G/100G NX-OS switch
- **Nexus 5500 Series**: Modular 1-2 RU Ethernet, Fibre Channel, and FCoE multilayer NX-OS switch
- **Nexus 3000 Series**: Ultra low-latency, high-performance multilayer NX-OS switch
- **Nexus 2000 Fabric Extenders**: 10Gbps, 1Gbps, 100Mbps virtual line cards for 7000/5000 Series
- **Nexus 1000v**: Virtual machine NX-OS switching for hypervisor environments with service integration

Entry-Level Data Center Models

Collapsed Core and Virtualized Data Center Switch sub-100 server ports

- Nexus 5500 virtual chassis combined core and DC pair
- Data Center specific feature set:
  - FEX (Physical ports, Adapter-FEX, VM-FEX)
  - Virtual Port Channel (vPC)
  - Storage networking with FC, FCoE, Unified Ports and DCB
  - Nexus 1000v with vPATH
  - FabricPath support
  - WAN router required for DCI and other transport services

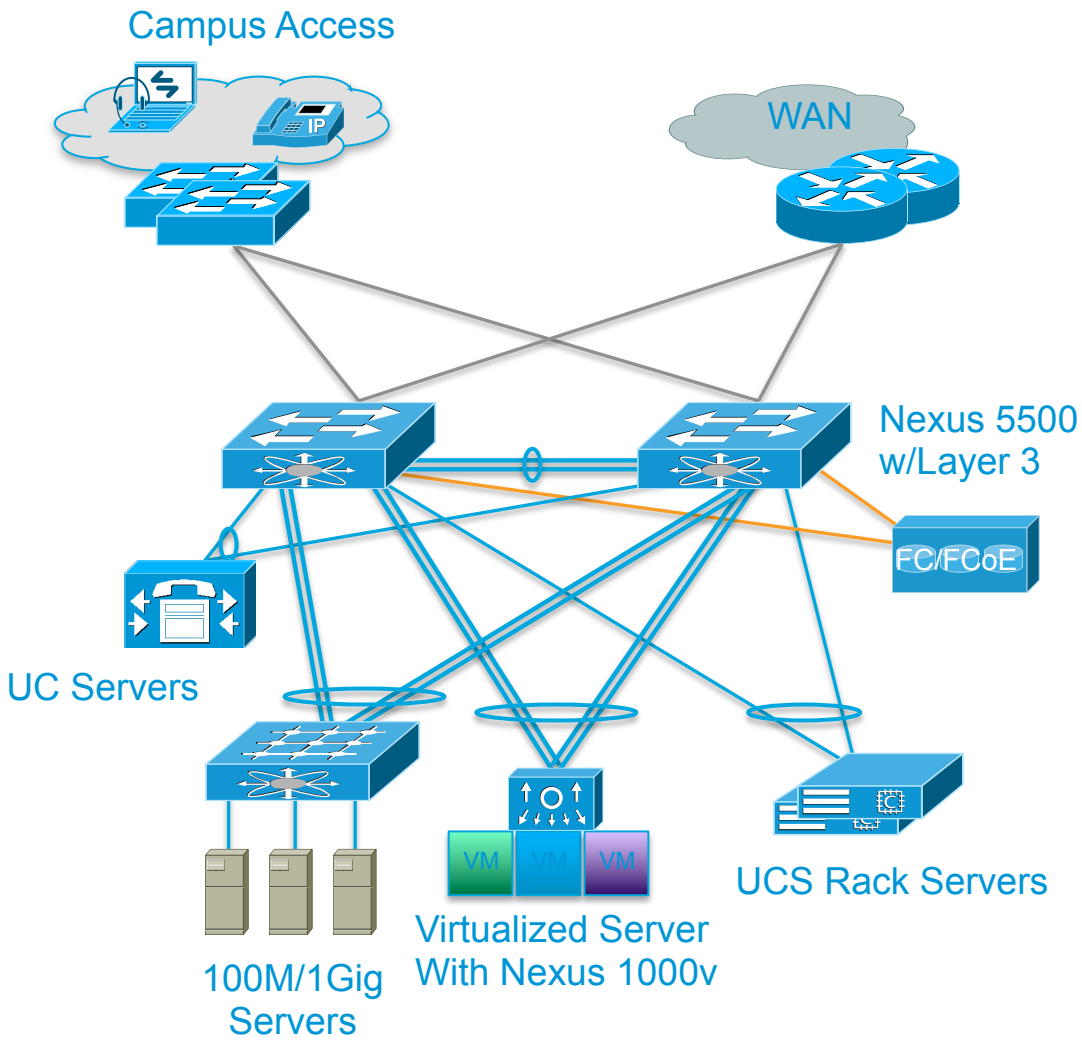

Entry-Level Data Center Models

Nexus 5500 Layer 3 DC, FEX, and Storage support; 100 to 300 server ports

- Nexus 5500 virtual chassis
  - Scale port count with additional FEX virtual line cards
- Data Center specific feature set:
  - FEX (Physical ports, Adapter-FEX, VM-FEX)
  - Virtual Port Channel (vPC)
  - Storage networking with FC, FCoE, Unified Ports, DCB
  - Nexus 1000v with vPATH
  - FabricPath support

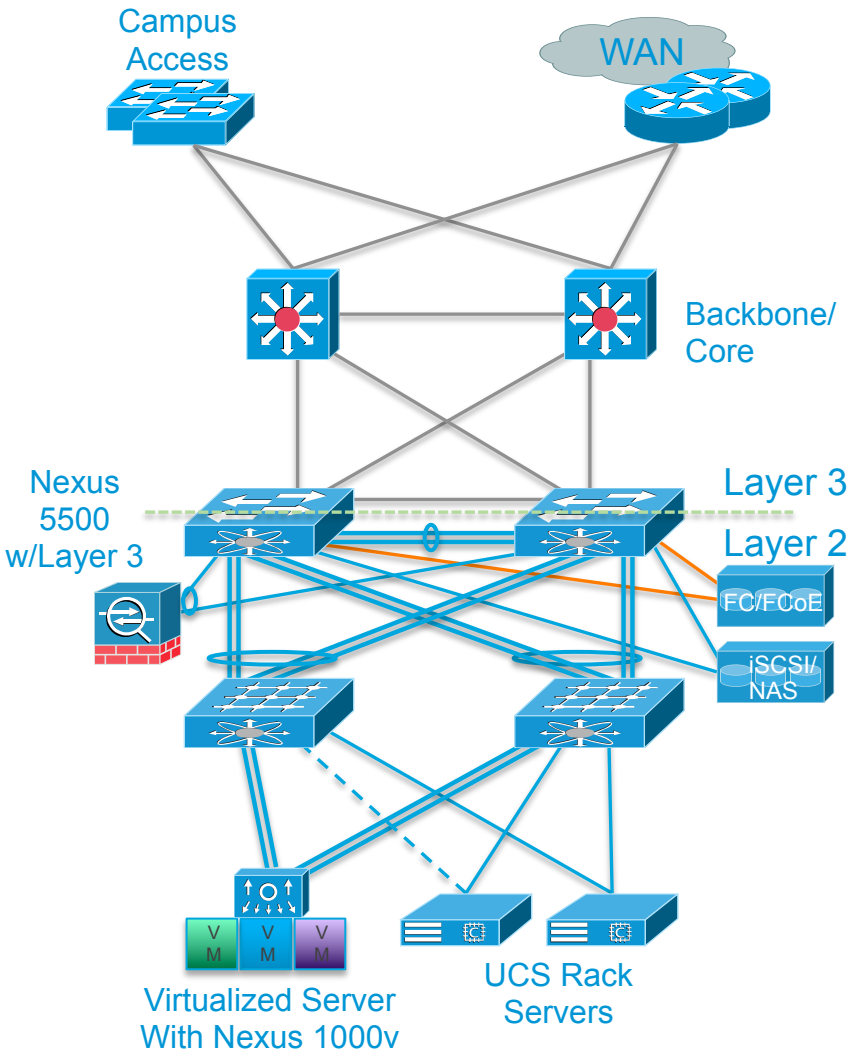

Entry-Level Data Center Models

Nexus 5500 Layer-3 Configuration Considerations

- Ensure a separate link is configured for Layer-3 peering outside of the vPC links.
- Set COS value with QoS policy to ensure that queuing is correct for traffic passing to and from the L3 module
- Base Layer-3 license is limited to 256 OSPF routes, or EIGRP Stub-only
  - Enterprise Layer-3 license for greater OSPF scale or full EIGRP
- Turn on vPC enhancements; peer-gateway, arp table sync, auto-recovery

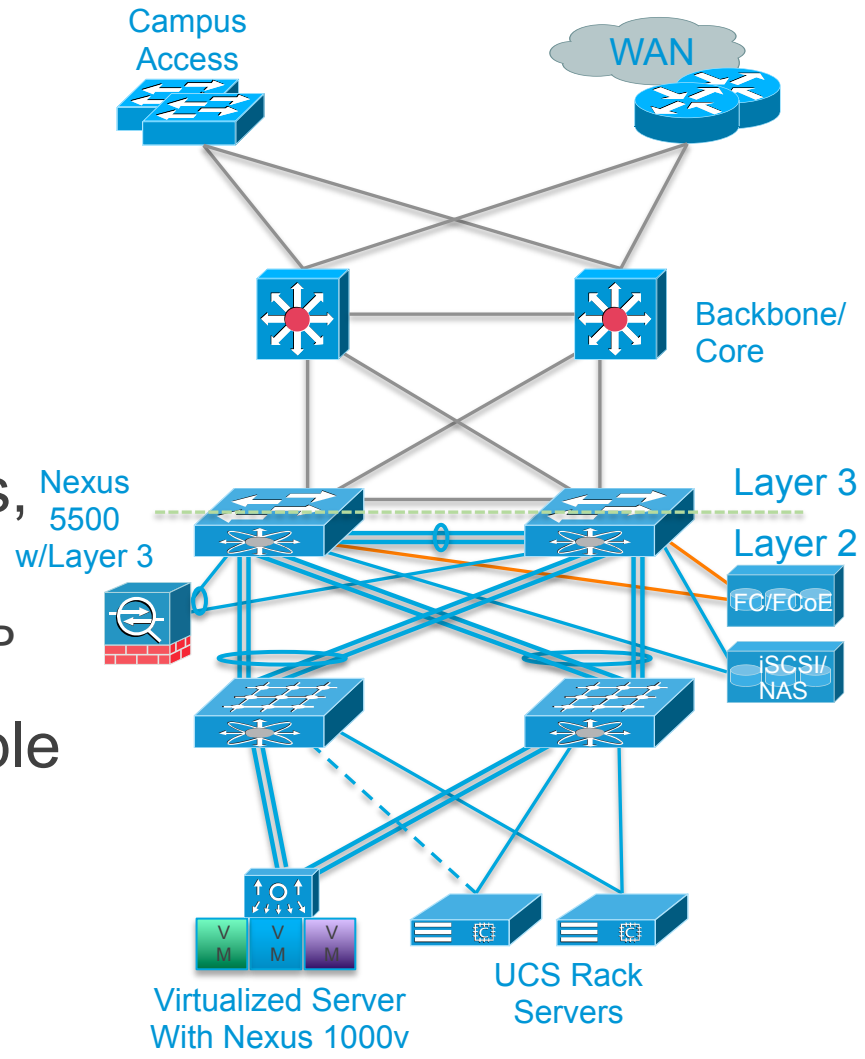

Entry-Level Data Center Models

Nexus 5500 Layer-3 Scaling Considerations

- Support for up to 16 FEX when Layer-3 enabled in Nexus 5500.
- Nexus 5500 Layer-3 module version 1 up to 8000 IPv4 hosts, 4000 Multicast routes.
  - 16000 IPv4 hosts / 8000 Multicast routes on V2 Layer-3 module with updated software.
  - Layer-3 + FEX + vPC designs limited to 1000 multicast groups.
- Operational consideration, In-Service Software Upgrade (ISSU) not supported with Layer-3 on the 5500.

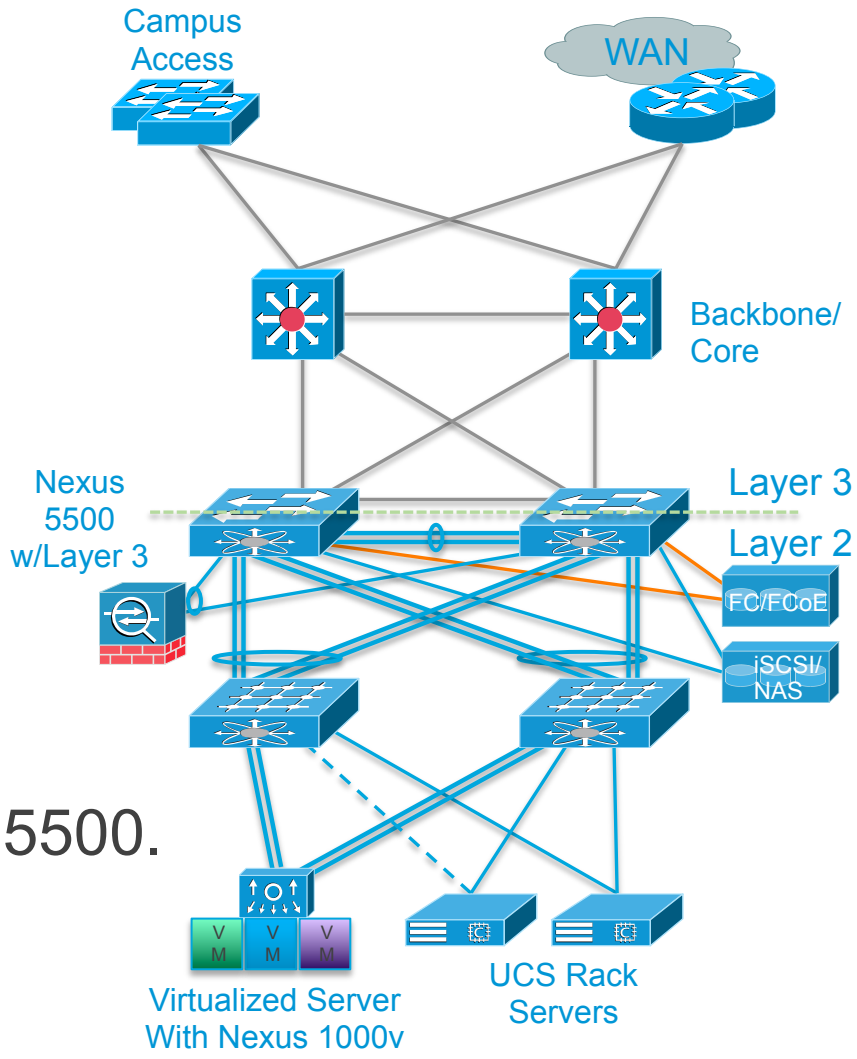

Scaling Designs with Access Block Variants

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

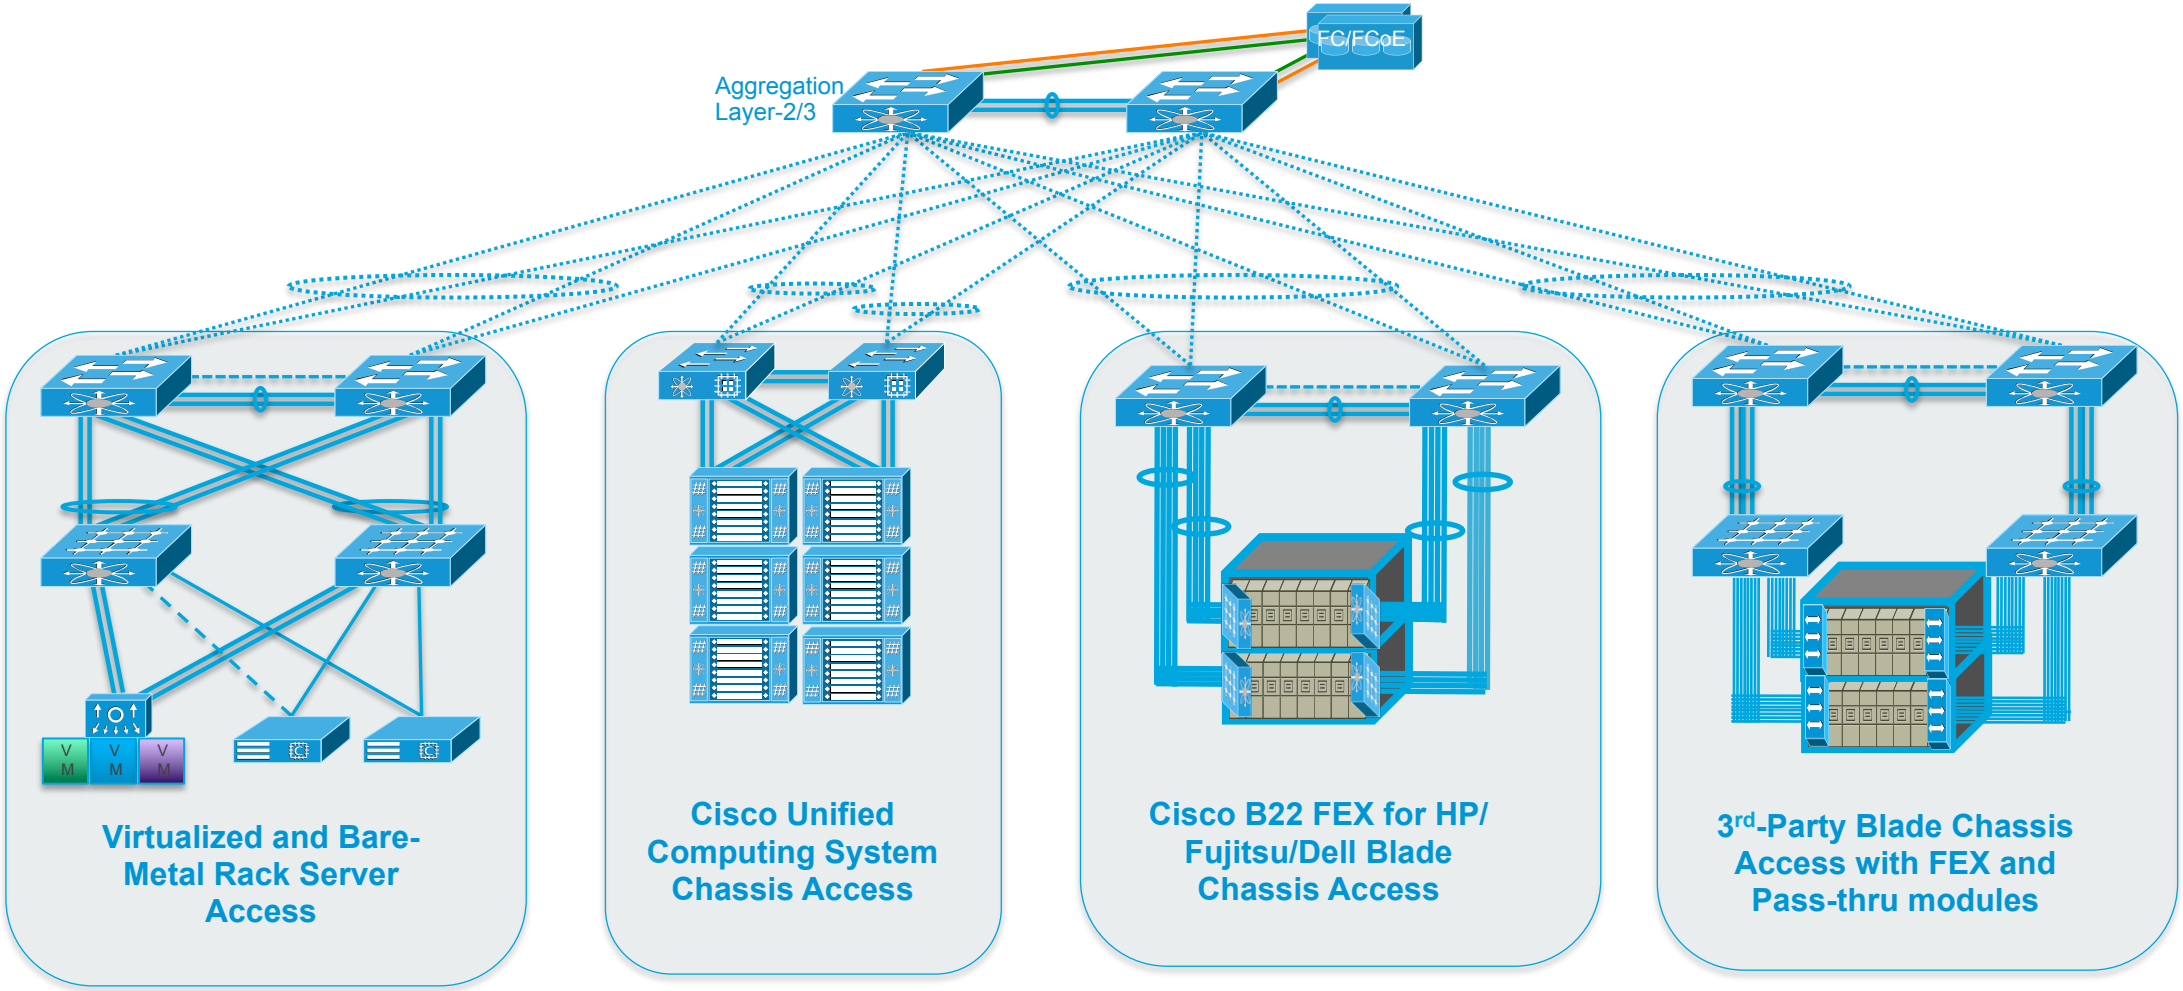

Kompaktní datové centrum

EMC²®

VSPEX

EMC: #1 storage for **applications**

Oracle	SAP	Exchange	VDI	Home Directory	SharePoint	Analytics
EMC² #1	EMC² #1	EMC² #1	EMC² #1	EMC² #1	EMC² #1	EMC² #1

Source : IDC

Three Paths To Private Cloud

Build Your Own...

Best Of Breed
Infrastructure Components

EMC VSPEX

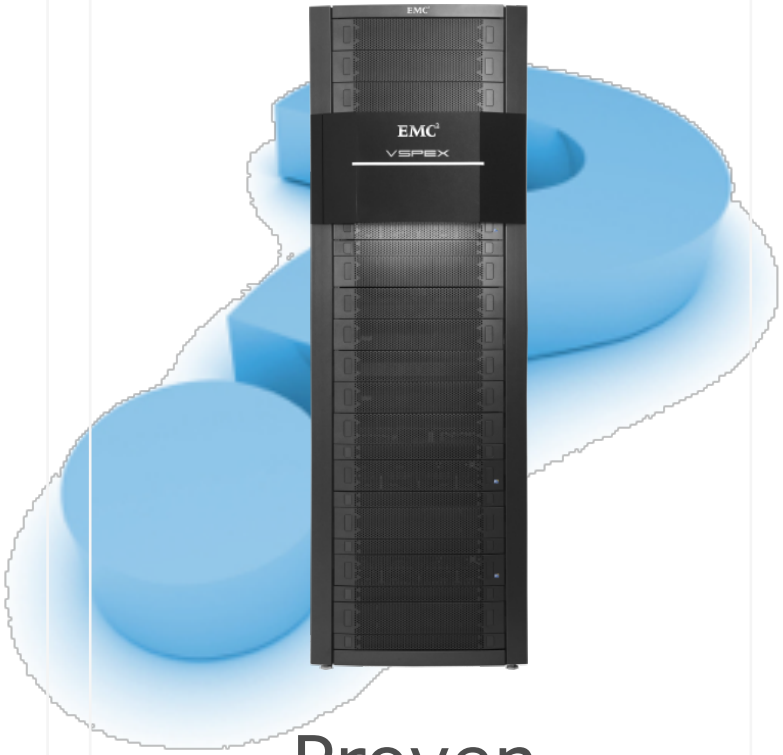

Proven
Infrastructure

VCE VBLOCK

Converged
Infrastructure

There is a clear Market for every flavor of Converged Infrastructure (CI).

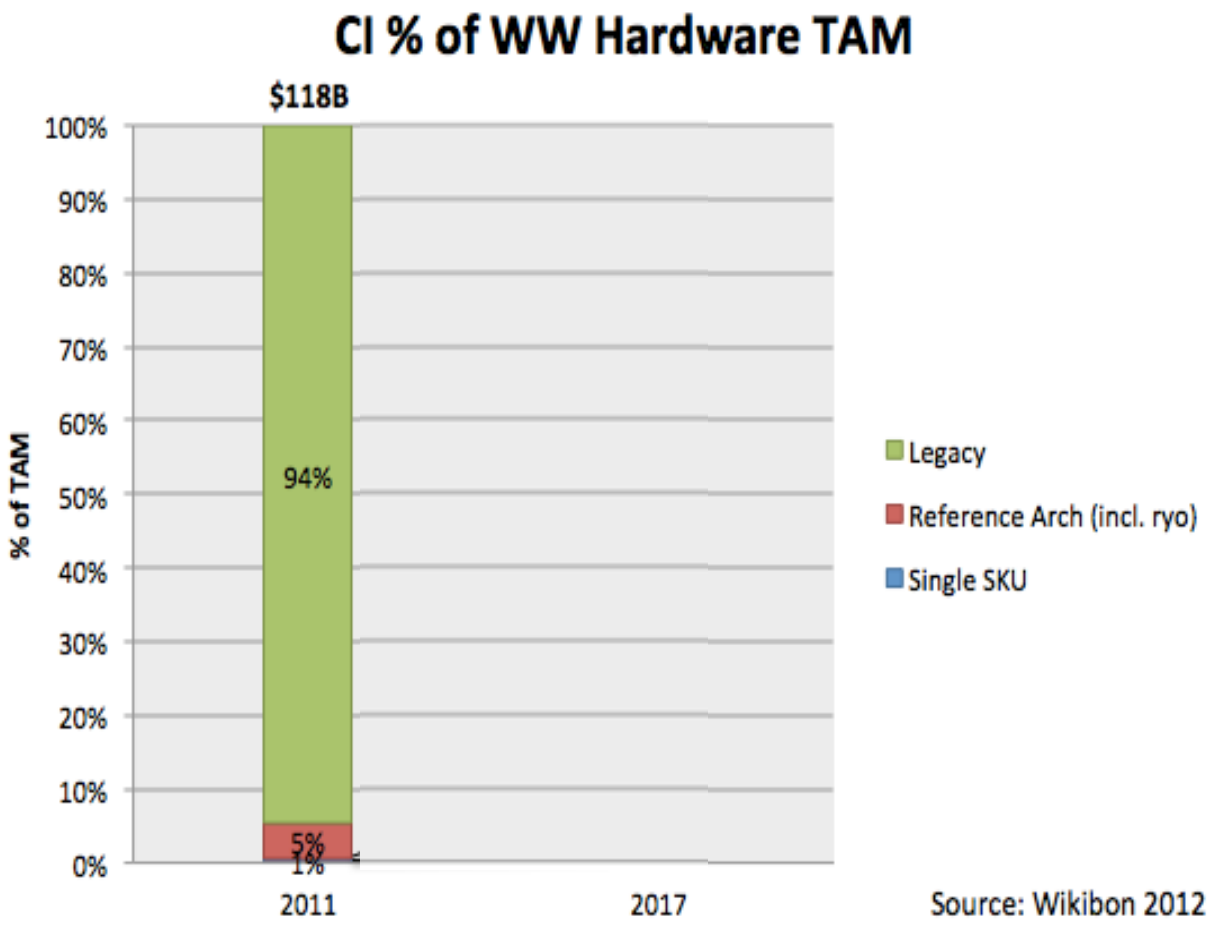

- The Problem
  - Today we sell Cloud infrastructure a-la-carte We pitch point products and product strenghts.
  - The market for this approach is dramatically shrinking. It will go down to ~1/3 over a few years.

Source : WIKIBON - David Vellante in Wikibon on February 28, 2012 – "New Data Shows EMC and Cisco Doubling Down on VCE"

EMC²

There is a clear Market for every flavor of Converged Infrastructure (CI).

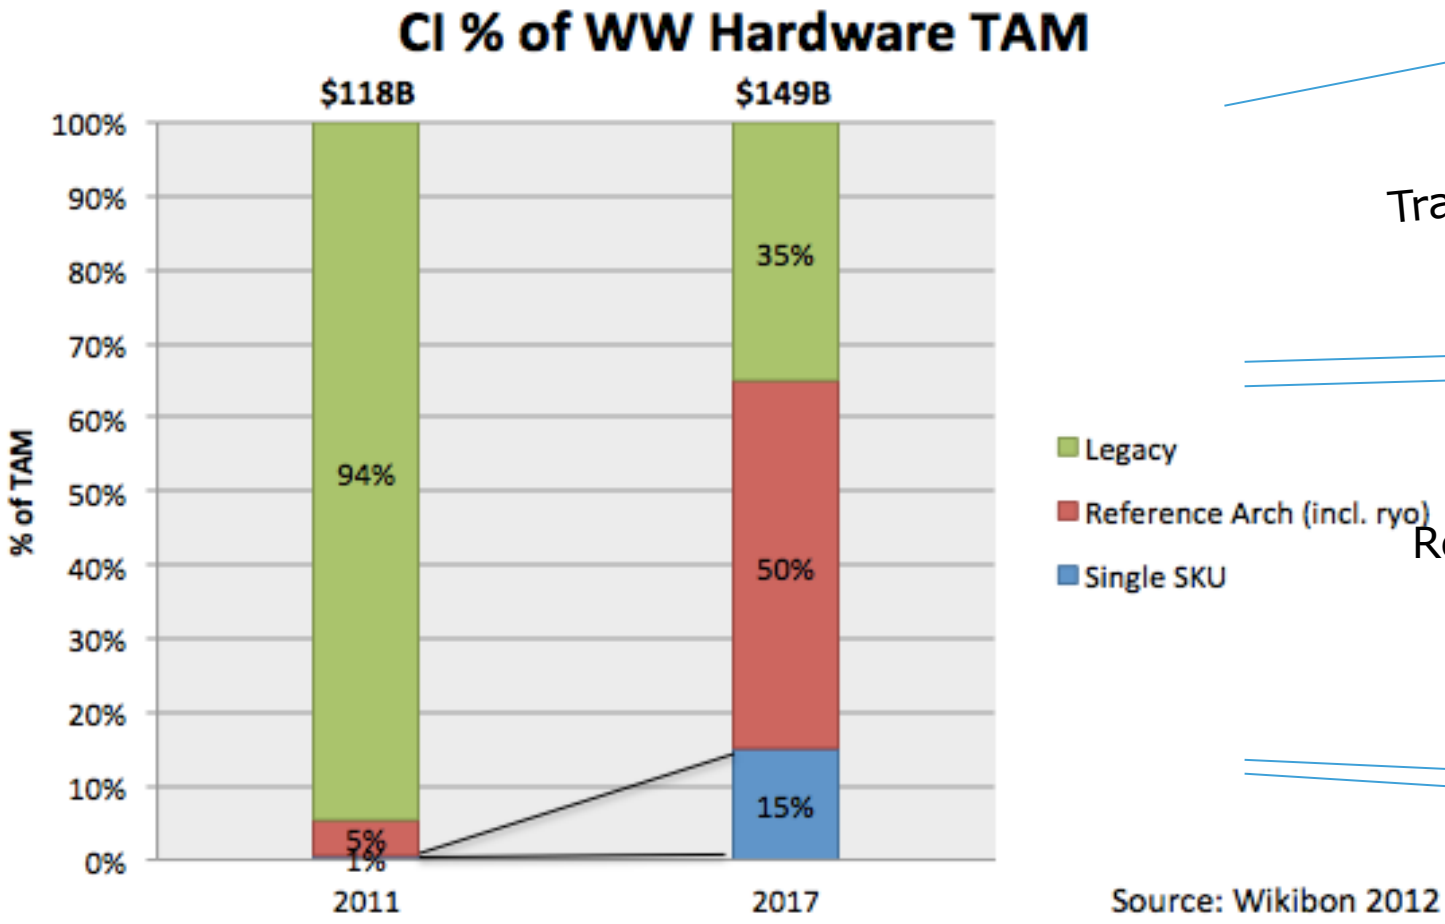

Source: Wikibon 2012

Traditional IT (BYO)

Ref Architecture

VSPEX

Single SKU

Source : WIKIBON - David Vellante in Wikibon on February 28, 2012 – "New Data Shows EMC and Cisco Doubling Down on VCE"

EMC²

EMC VSPEX

SIMPLE. EFFICIENT. FLEXIBLE.

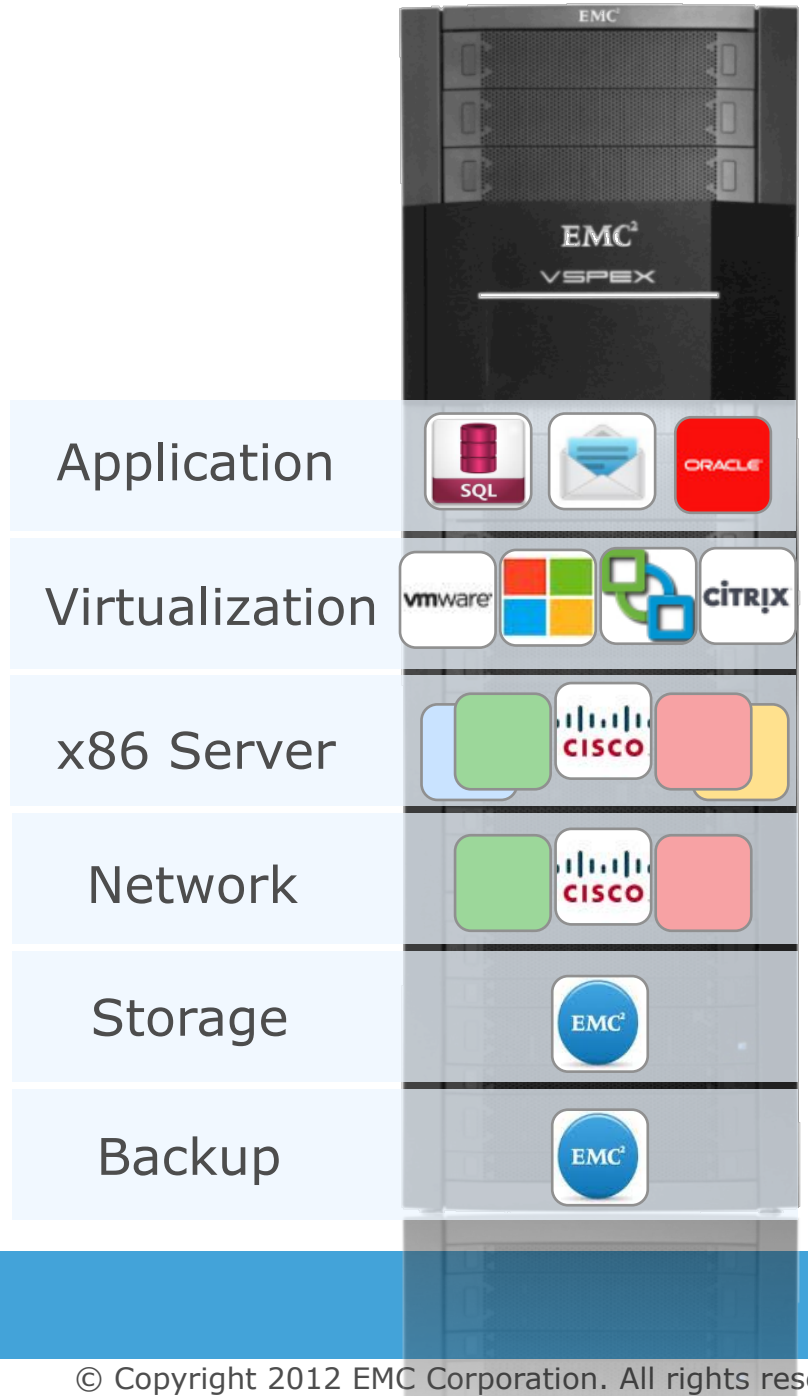

- Low Risk, Modular, & Scalable
- RAs for Repeatable Deployments
- Investment Protection
    - Choice in Virtualization
    - Choice of Server
    - Choice of Network
    - Re-use existing HW/SW/ Licensing Assets
- EMC Storage
- EMC Next Generation Backup
- EMC Global Financing Available

EMC²

VSPEX Proven by EMC

Proven means Tested, Validated, and Confidence.

01 **Proven & Extensively Tested**

- VM defined based on real world configurations
- VDI validation with LoginVSI (simulates desktop and user load).
- Install quicker, with confidence and with ease.

02 **Validated Design**

- Scale-Up and Steady State testing
- High Availability Testing
- Failure Simulations
- Backup and Recovery Testing
- Add VSPEX configs easily as needed

03 **Confident Set Up**

- Design Guides
- Implementation Guides

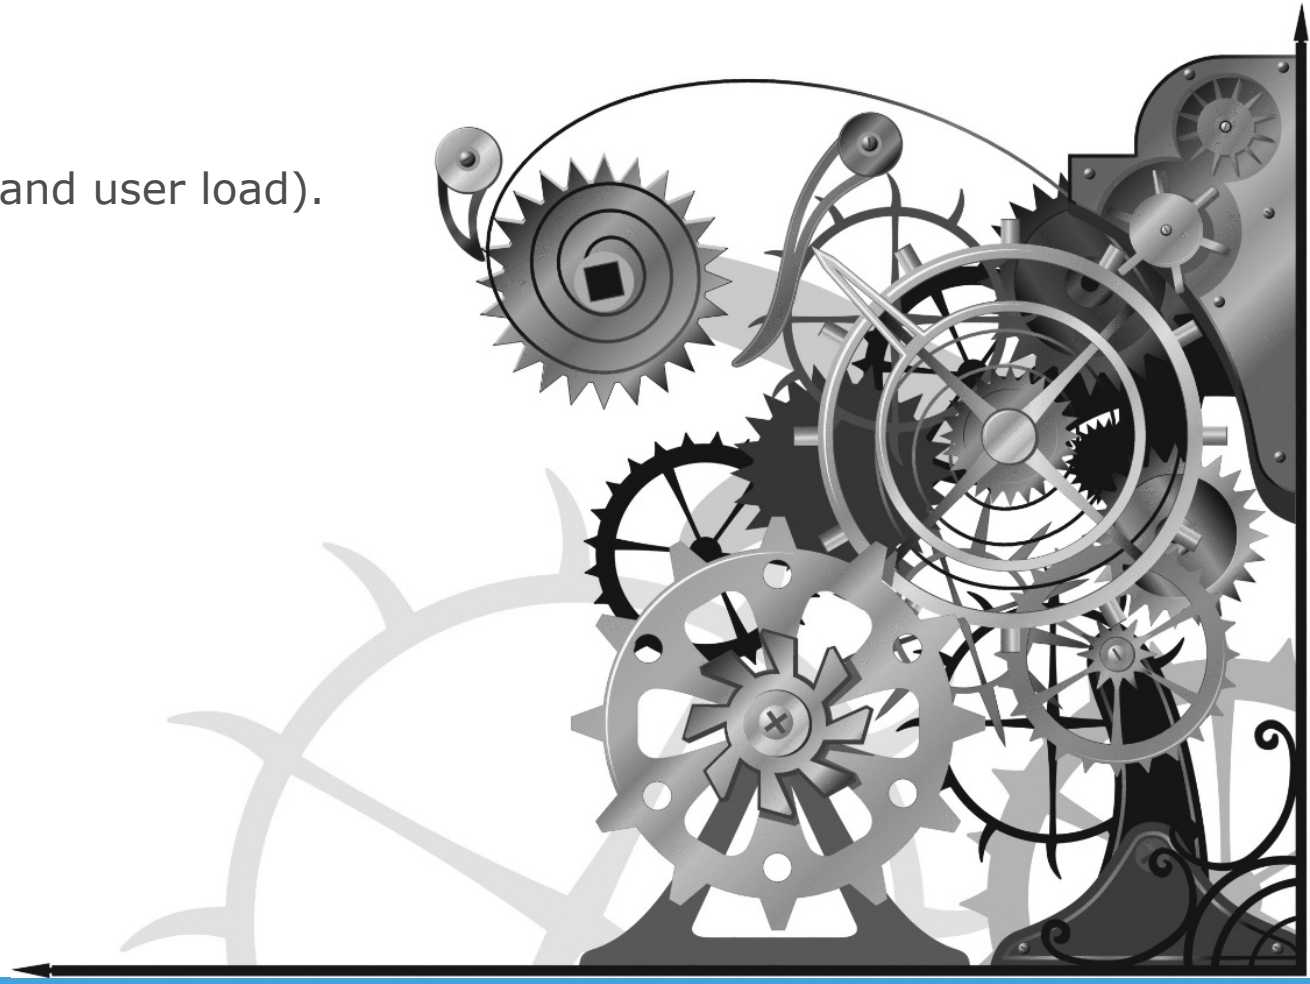

VSPEX Partner Program

VSPEX Proven Infrastructure for Applications

Virtualizing Applications on VSPEX Private Cloud

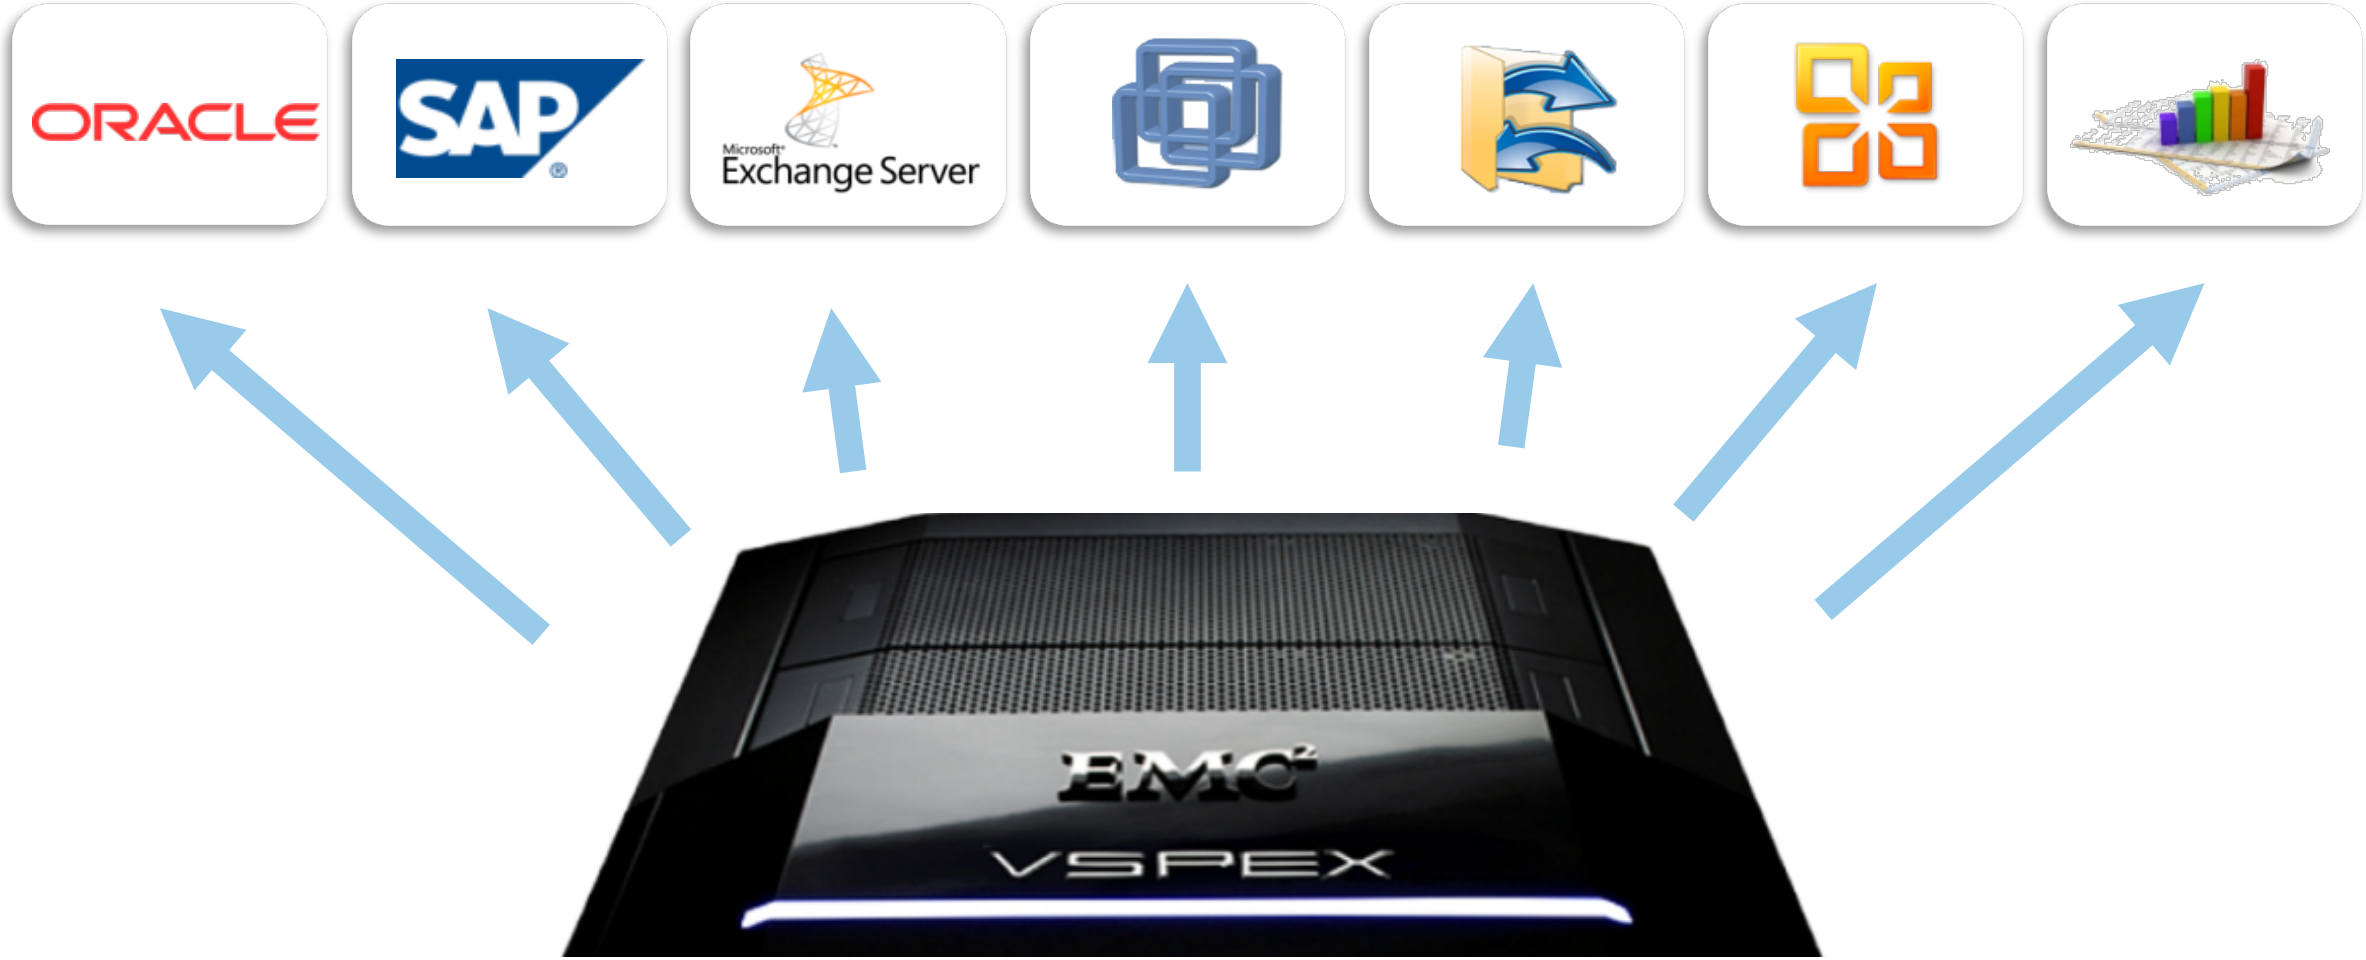

Conclusion

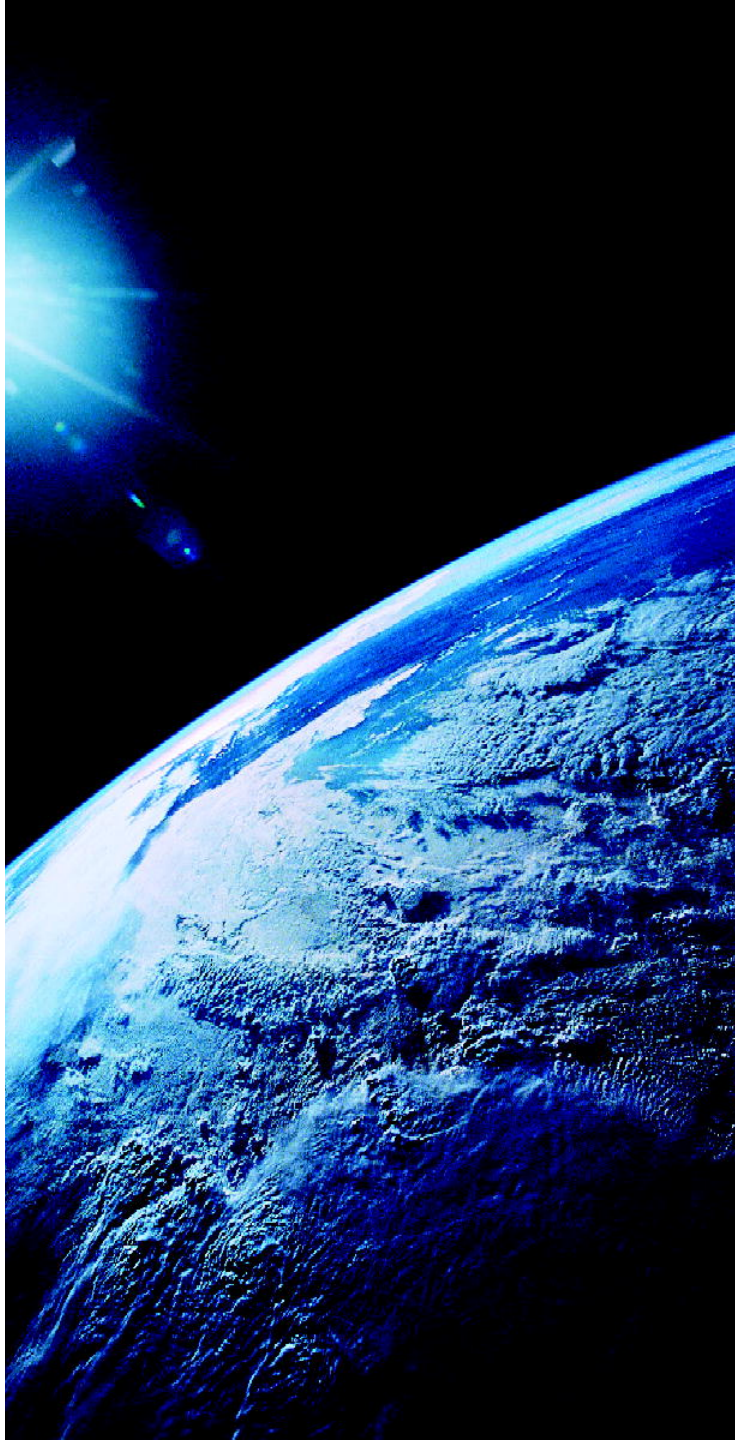

Small and Medium Data Center Design

Summary and key takeaway points

- Data center design for small to mid-market customers requires solving similar challenges that larger customers face.
- Data center specific features from Cisco are available on the Nexus family of switching platforms to address these challenges in varying levels of scale.
- Design elements of the example models covered today allow customers to scale their data center footprint from entry level through hundreds of server ports.
- Cisco has extended many of the capabilities of the Unified Computing System into the Nexus switching family to allow a single switching fabric to support all variants of compute and storage connectivity needs.
- EMC VSPEX

Děkujeme za pozornost.