



Cisco *live!*

End-to-End Data Centre Virtualisation

BRKVIR-2931



Carlos Pereira
Data Centre Distinguished Systems Engineer
carlos.pereira@cisco.com

Setting the Stage:

What's the Meaning of Virtual?

- If you can see it and it is there

It's **real**



- If you can't see it but it is there

It's **transparent**



- If you can see it and it is not there

It's **virtual**



- If you can not see it and it is not there


It's **gone**



Today's Data Centre Virtualisation “Journey” ...




Today's Data Centre Virtualisation "Journey" ...



Our Journey in a Nutshell

 Applications team

 Servers team

 Networking team

 Storage team


 Facilities team




Front-End Virtualisation




Virtual Network Services



Access Layer – Physical and Virtual



Virtual Machines & IO Virtualisation



Virtual SANs / Unified IO




Virtual Storage







Ultimate Goal: Remove the Virtualisation ...




... so, please ...




Agenda

- 
 - Data Centre Virtualisation Overview
 - Front-End Data Centre Virtualisation
 - Core Layer
 - Aggregation Layer
 - Networking Services
 - Access Layer
- 
 - Server Virtualisation
 - Hypervisors
 - Virtual Access Layer
 - Virtualised Services
 - Server IO Virtualisation
- 
 - Back-End Virtualisation
 - Virtual HBA & NPV
 - Unified IO & FCoE
 - SAN & Storage
- 
 - Implementation Examples
 - vBlock
 - Secure Cloud
 - FlexPOD
 - Q&A



Key Trends Impacting IT & Data Centres




The need to reduce costs and/or maximize profits




IT as business enabler




Applications availability



Drive for Green—power, cooling and space




Server Virtualisation — higher performance



LAN and Storage convergence




VM-Level awareness



Workload provisioning

Virtualisation touches half (at least ...)




The need to reduce costs and/or maximize profits




IT as business enabler




Applications availability



Drive for Green—power, cooling and space




Server Virtualisation — higher performance



LAN and Storage convergence



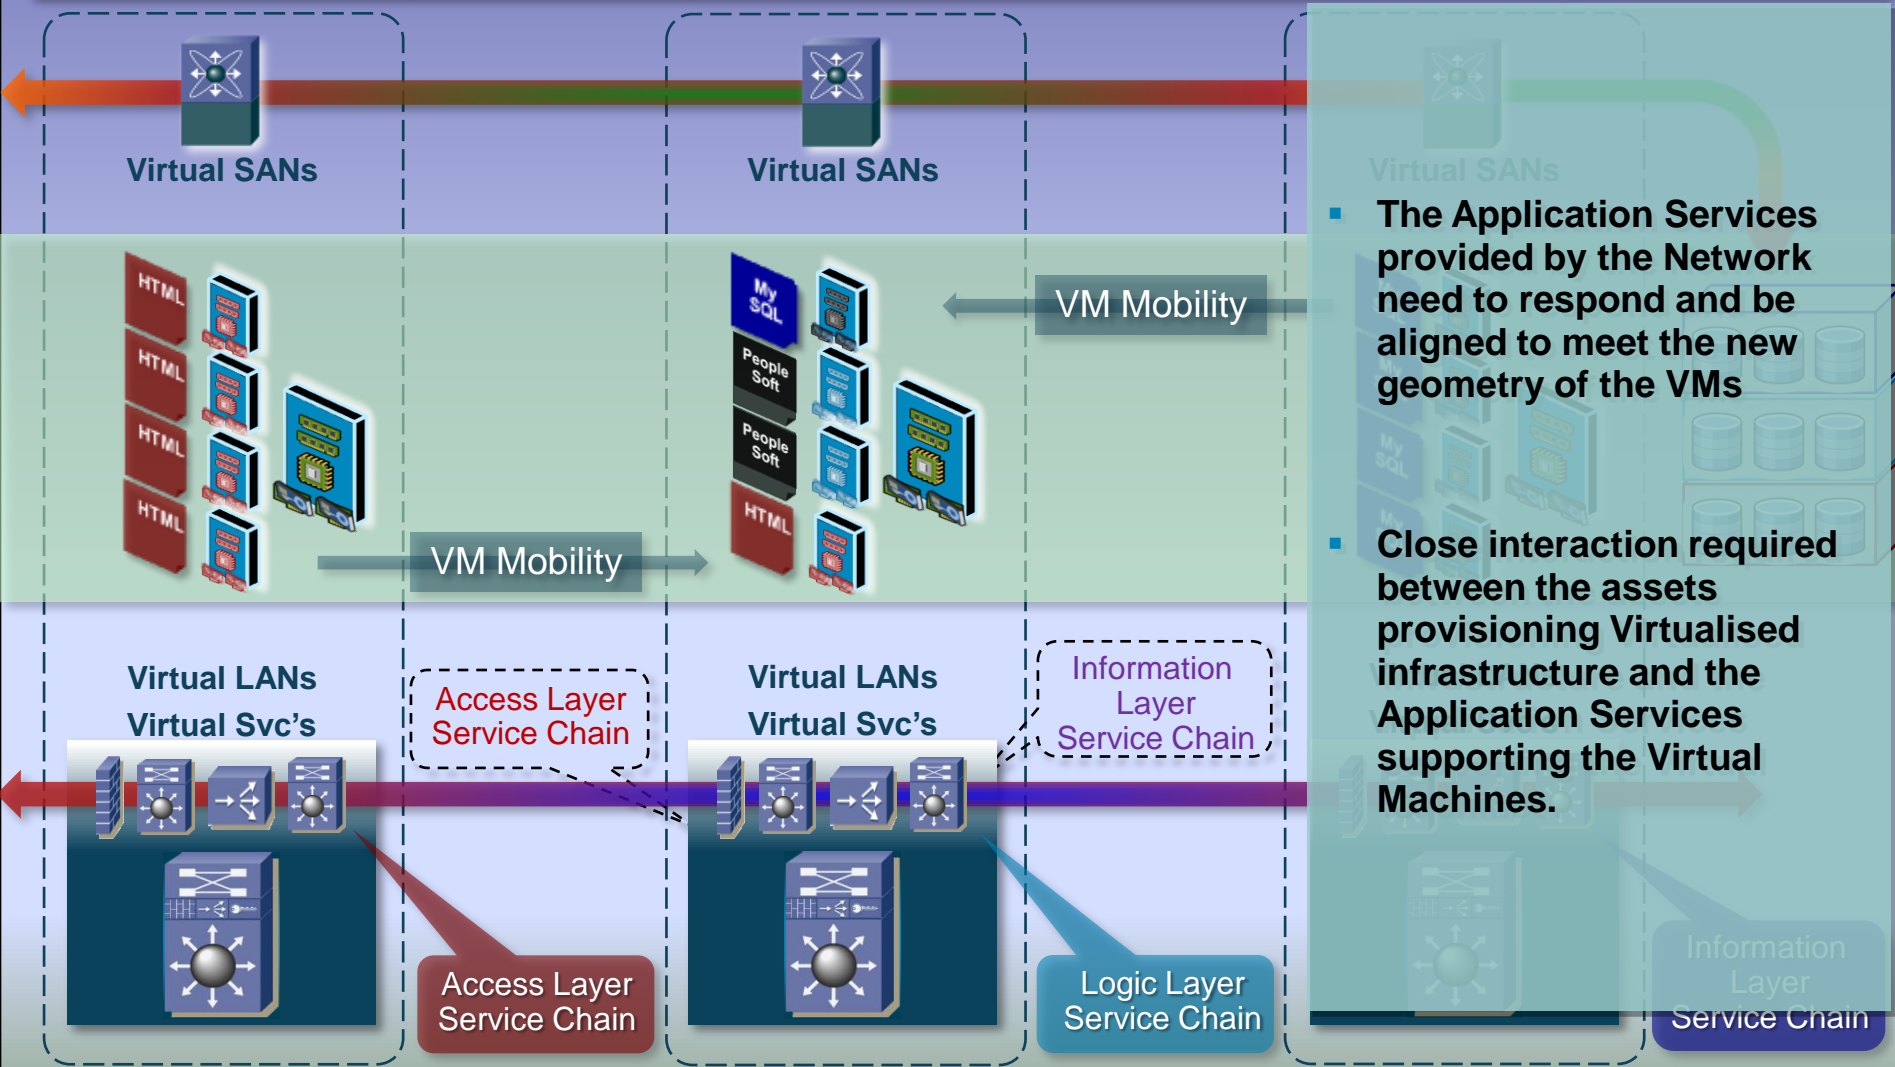
VM-Level awareness



Workload provisioning

The “Virtual Data Centre” Approach


VM's Mobility Across Physical Server Boundaries and Keeping Services




Moving to a Unified Fabric

Moving to a fully Virtualised Data Centre, with Any To Any Connectivity


Unified Fabric Networking




Virtual SANs
Virtual LANs
Virtual Svc's



Unified Fabric Networking



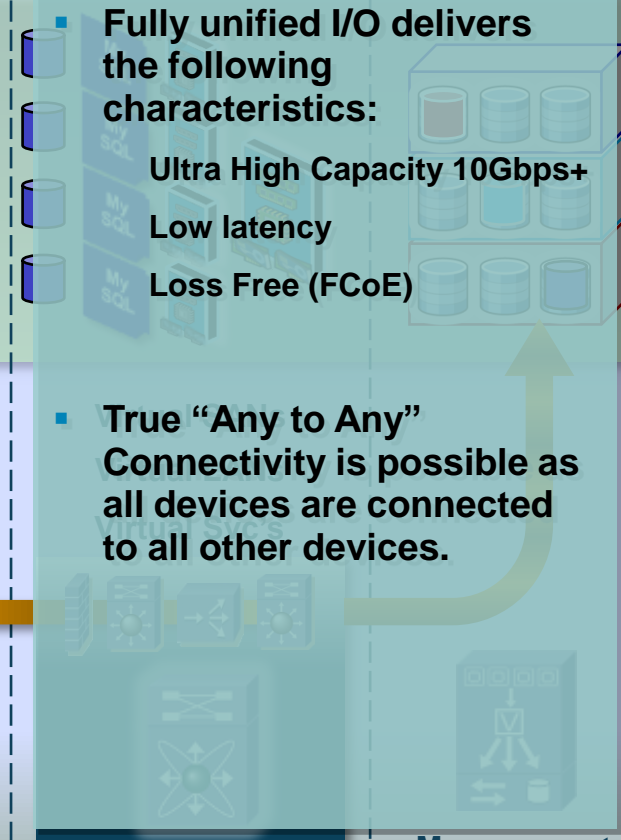
Virtual SANs
Virtual LANs
Virtual Svc's



Unified Fabric Networking

Unified Fabric Networking

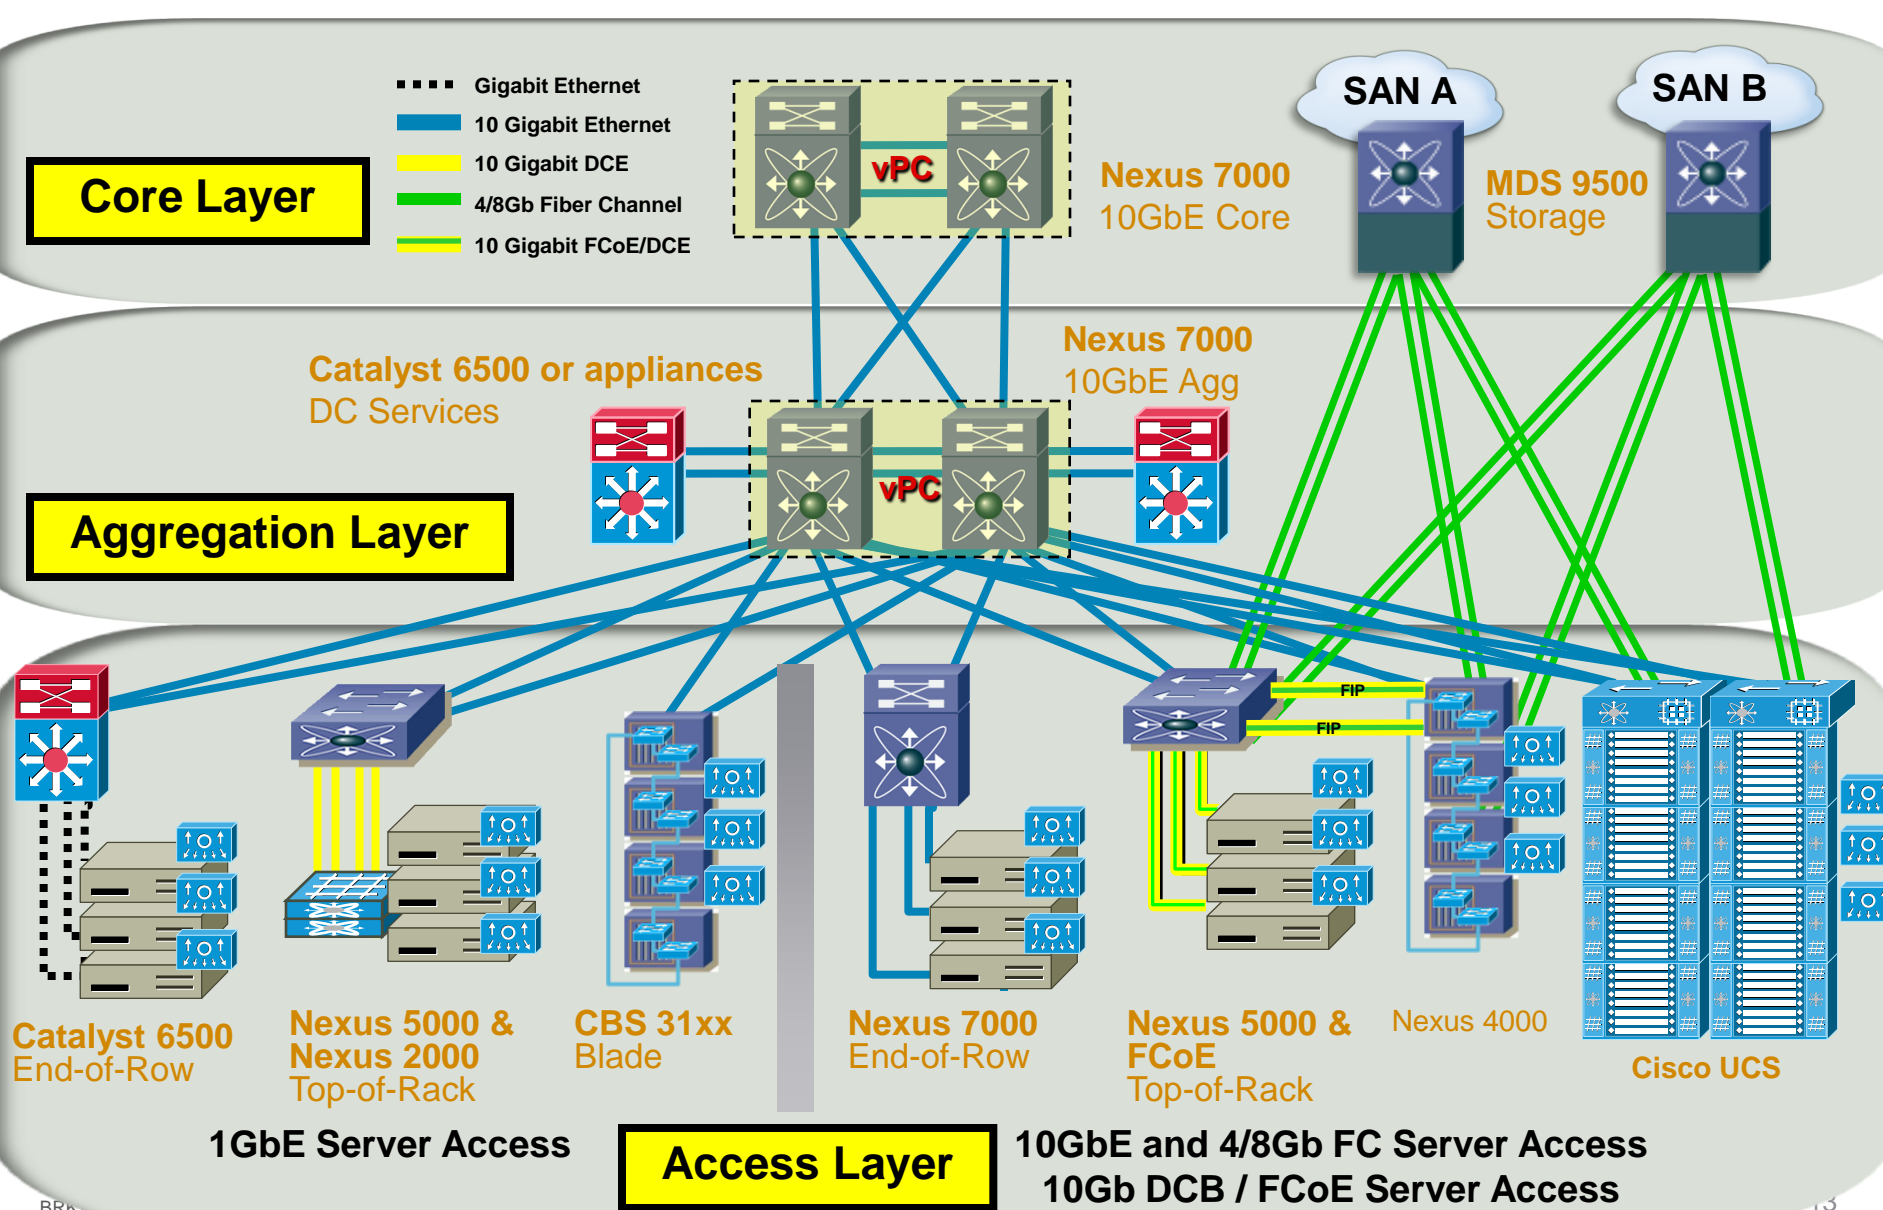
Unified Fabric Networking



- Fully unified I/O delivers the following characteristics:
 - Ultra High Capacity 10Gbps+
 - Low latency
 - Loss Free (FCoE)
- True "Any to Any" Connectivity is possible as all devices are connected to all other devices.

Management

Virtualised Data Centre Infrastructure



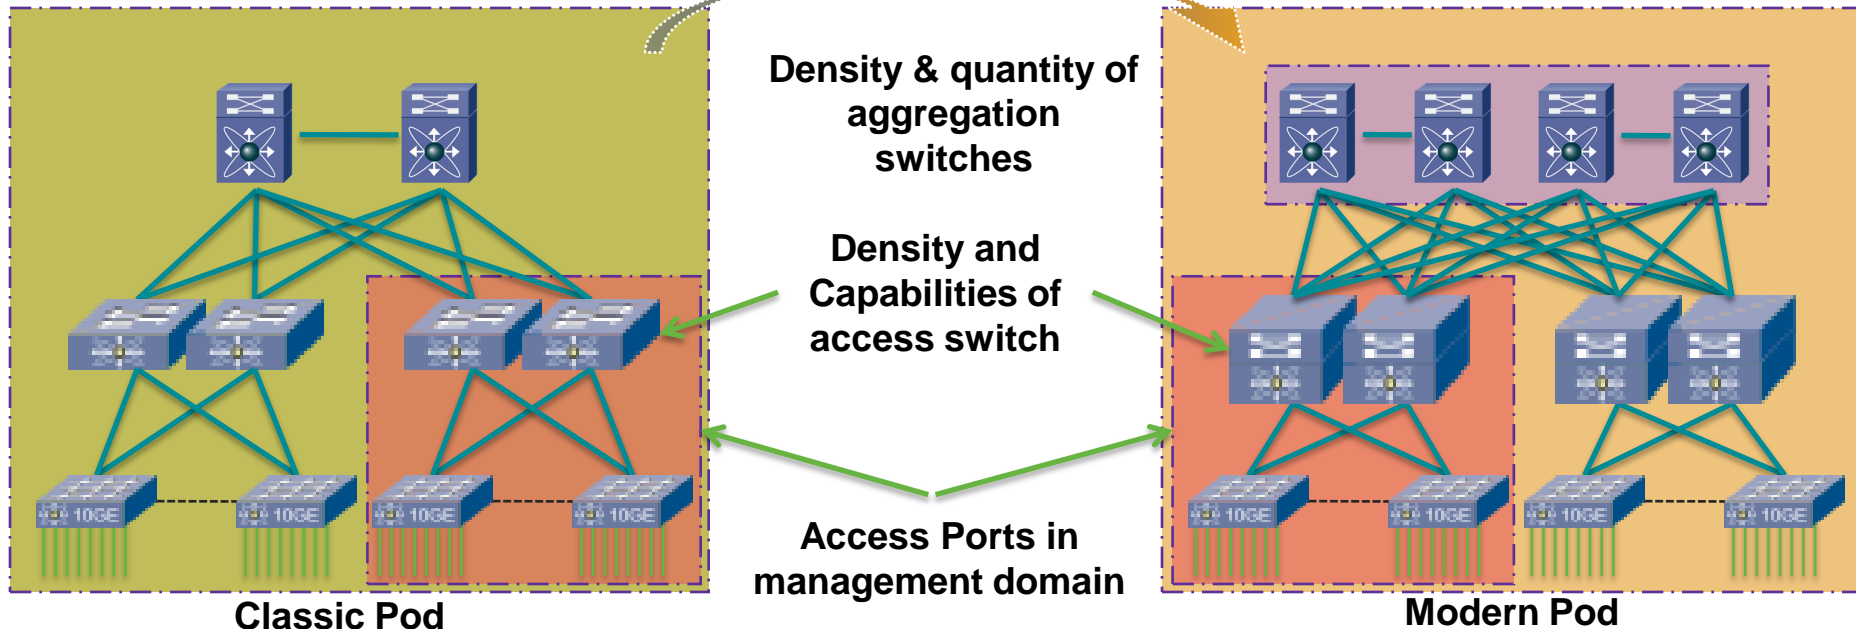
Virtualised Network Infrastructure: Evolution & Considerations

Typical DC Challenges

- L2 Fate-sharing
- VLAN Location
- L2 Adjacency
- Higher Scale
- L3 Access
- App Environments




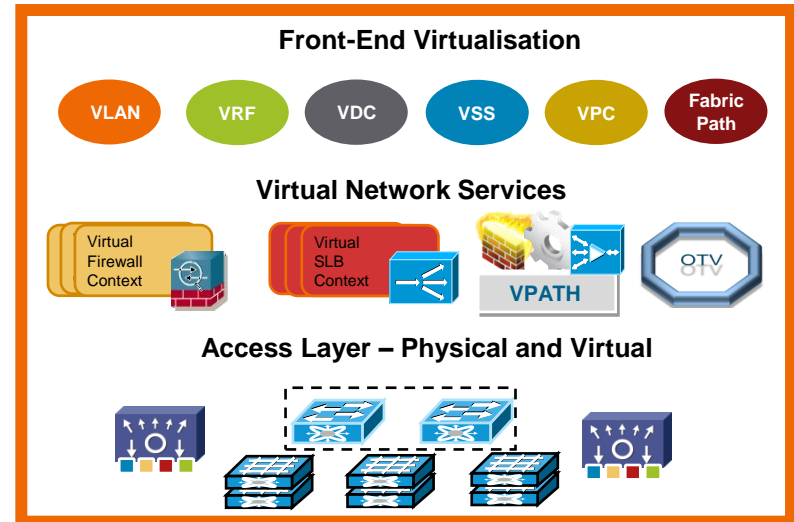
What are the implications...

- IS-IS as an L2 routing protocol
- Any VLAN anywhere resonates well
- Lower oversubscription
- Larger subnet sizes
- Global VLANs
- Specific app environments Designs

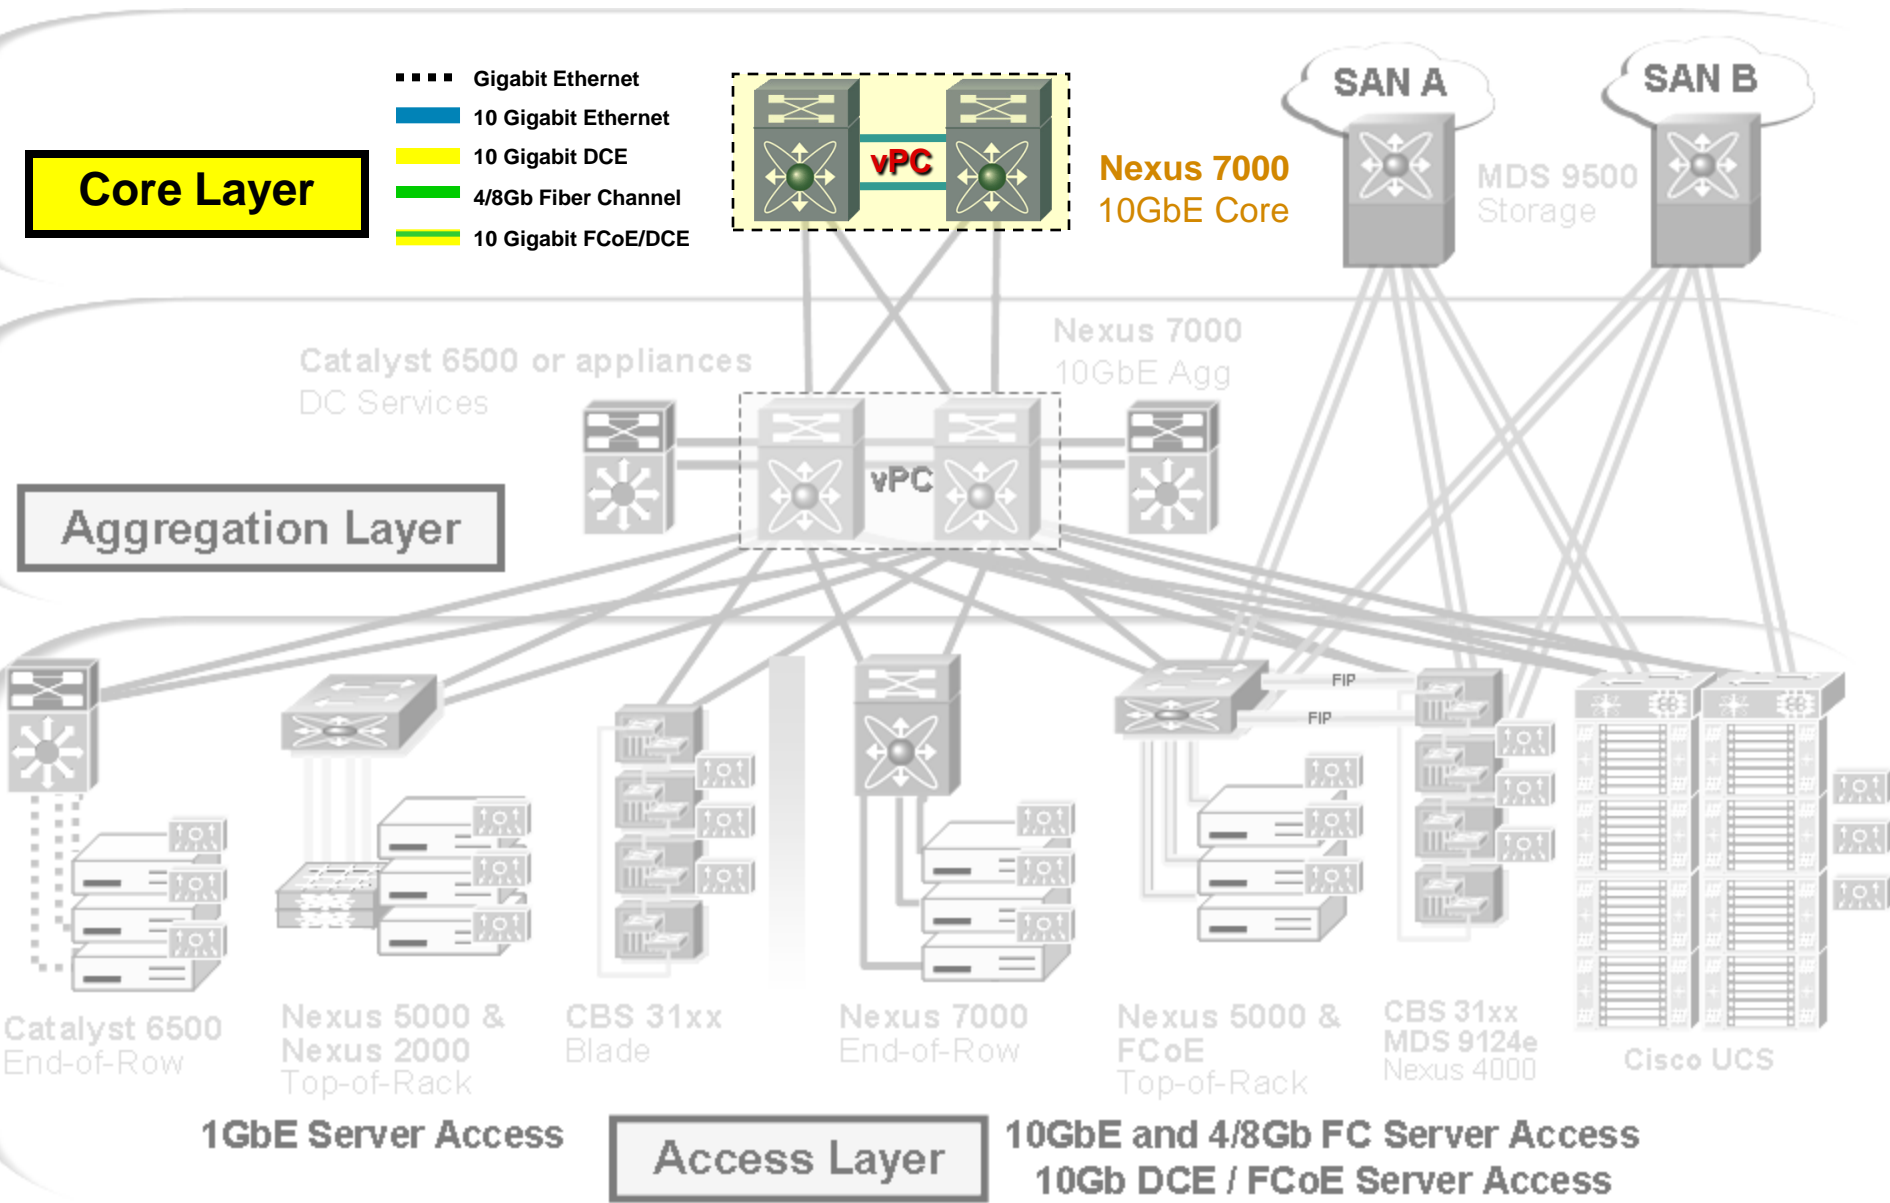


Agenda

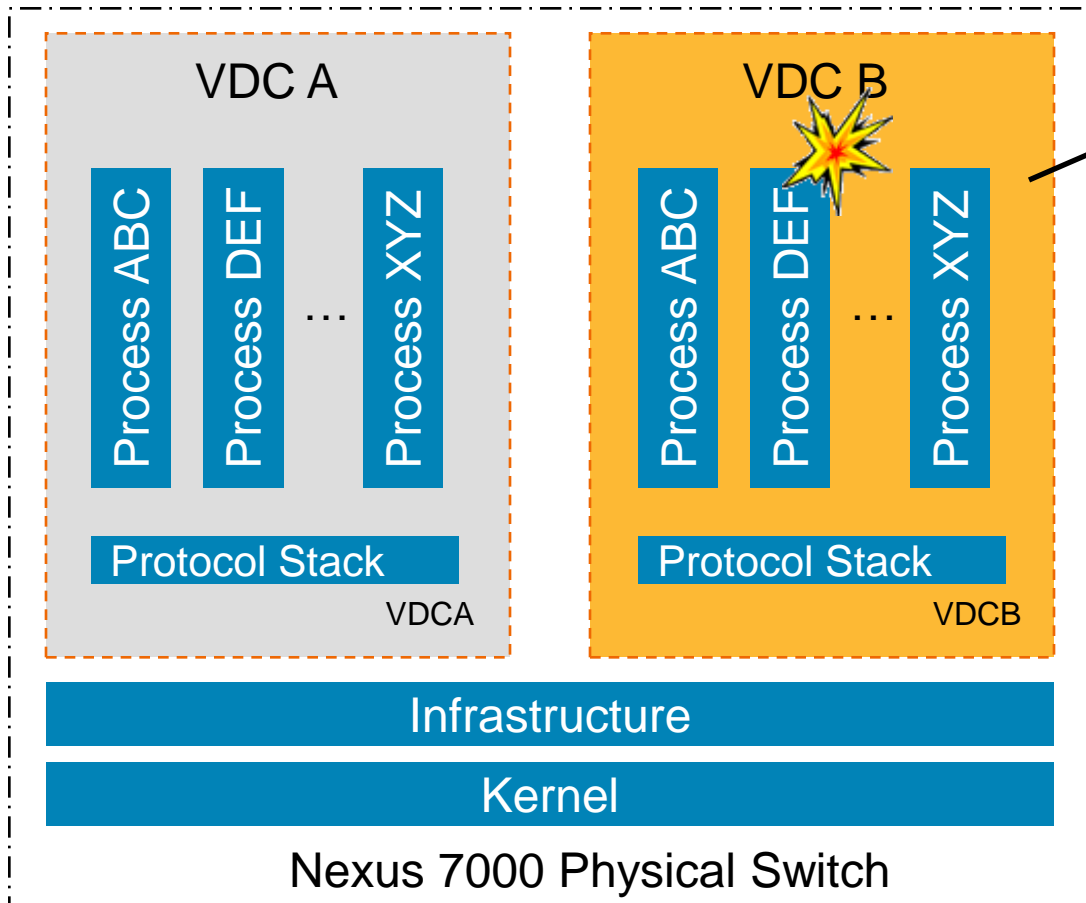
- Data Centre Virtualisation Overview
- **Front-End Data Centre Virtualisation**
 - Core Layer
 - Aggregation Layer
 - Networking Services
 - Access Layer
- Server Virtualisation
 - Hypervisors
 - Virtual Access Layer
 - Virtualised Services
 - Server IO Virtualisation
- Back-End Virtualisation
 - Virtual HBA & NPV
 - Unified IO & FCoE
 - SAN & Storage
- Implementation Examples
 - vBlock
 - Secure Cloud
 - FlexPOD
- Q&A



Front-End: Core Layer




Virtual Device Contexts @ Nexus 7000



Process "DEF" in VDC B Crashes

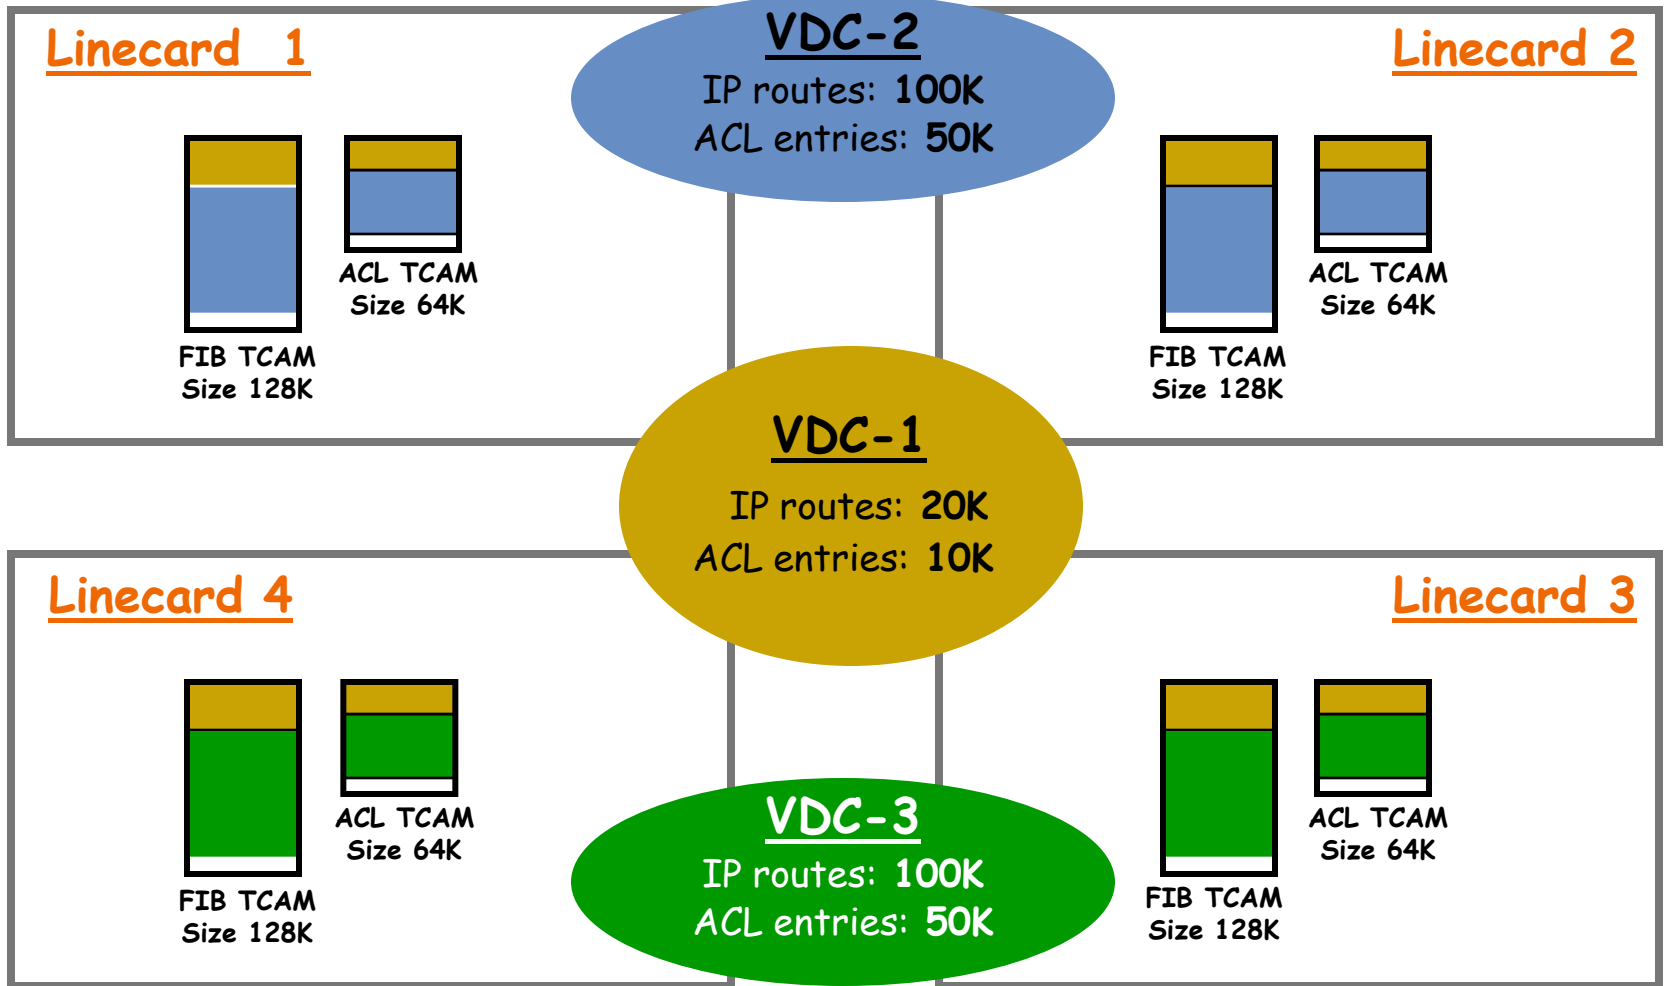
Process DEF in VDC A Is Not Affected and Will Continue to Run Unimpeded



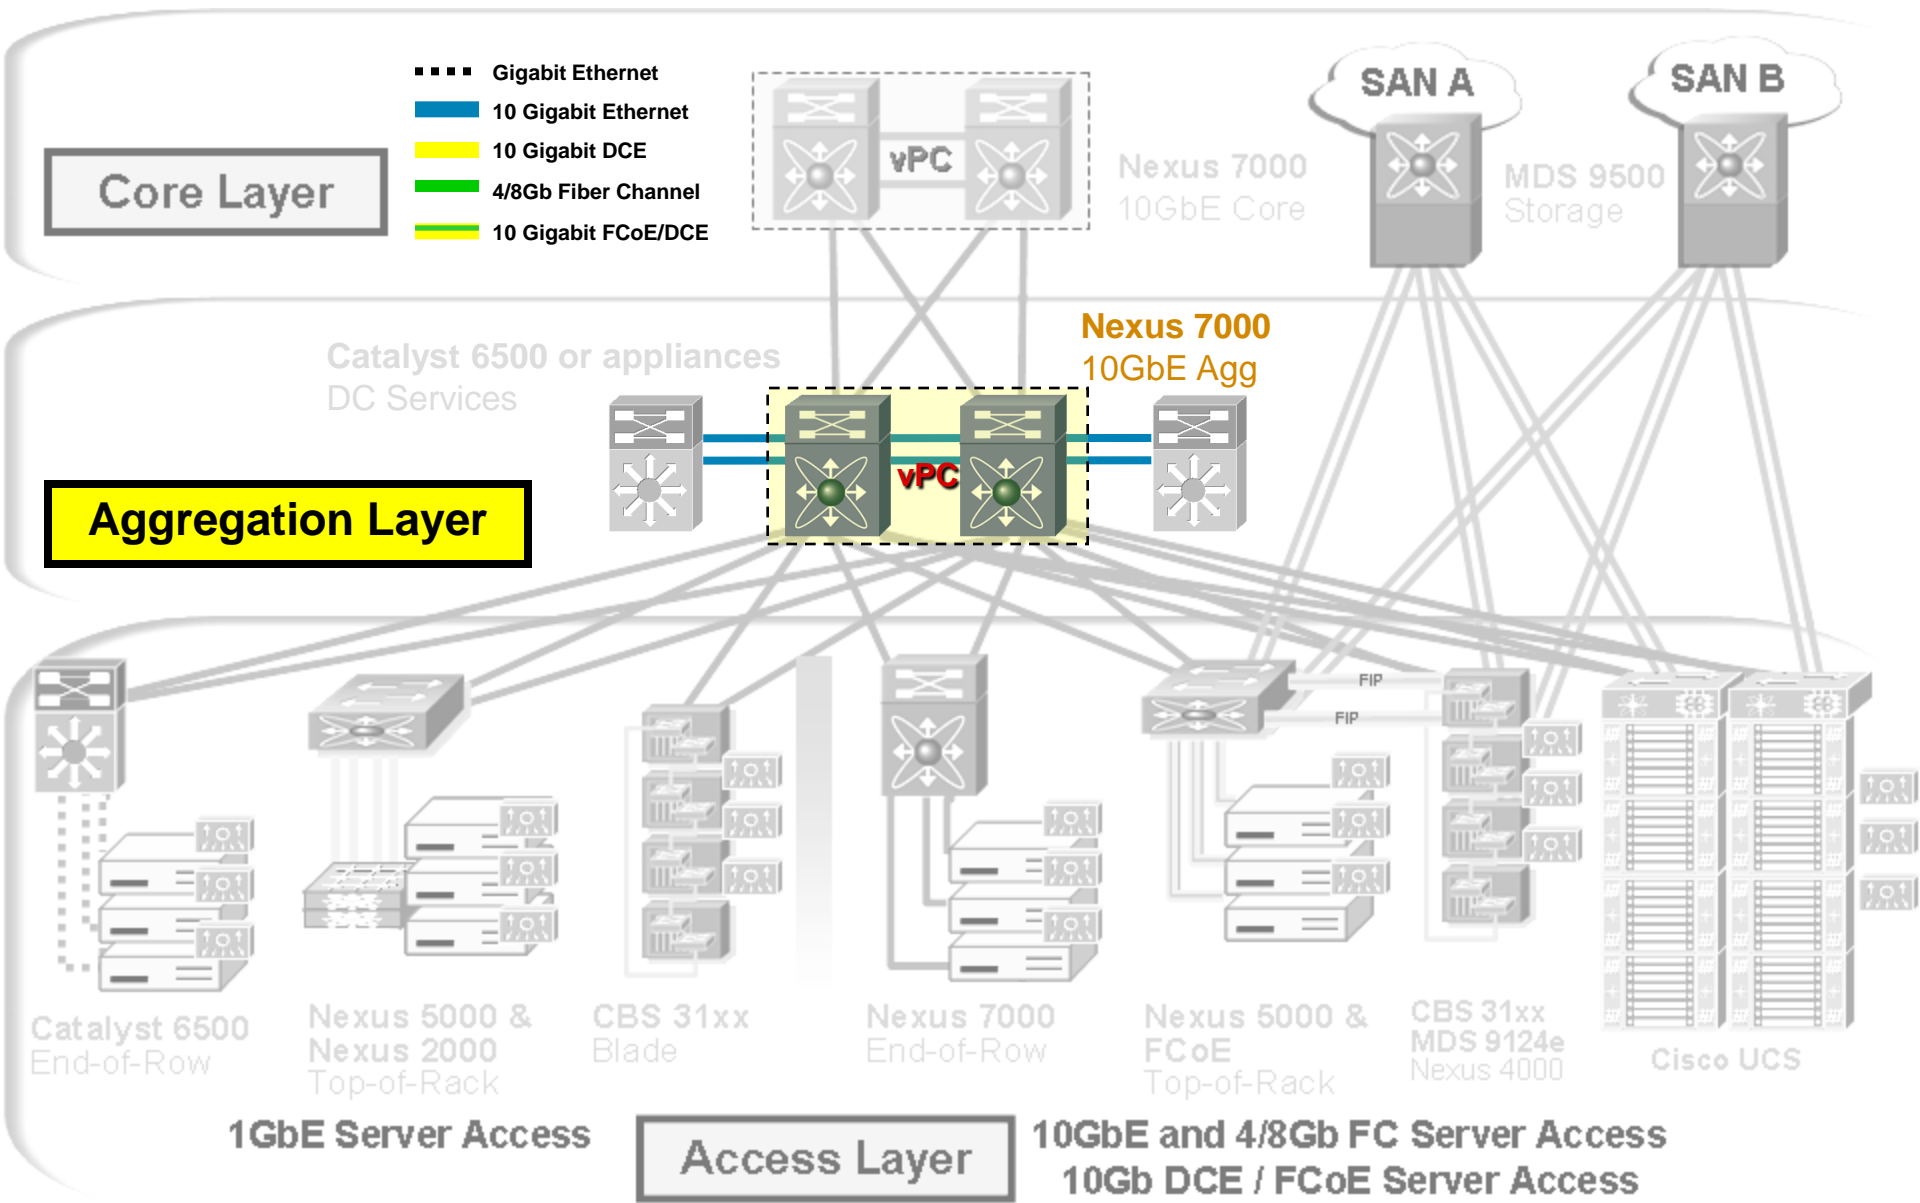
Virtual Device Contexts

Separate Resource Allocation Domains (Layer 3)

1 : N



Front-End: Aggregation Layer





Network Planes of Operation

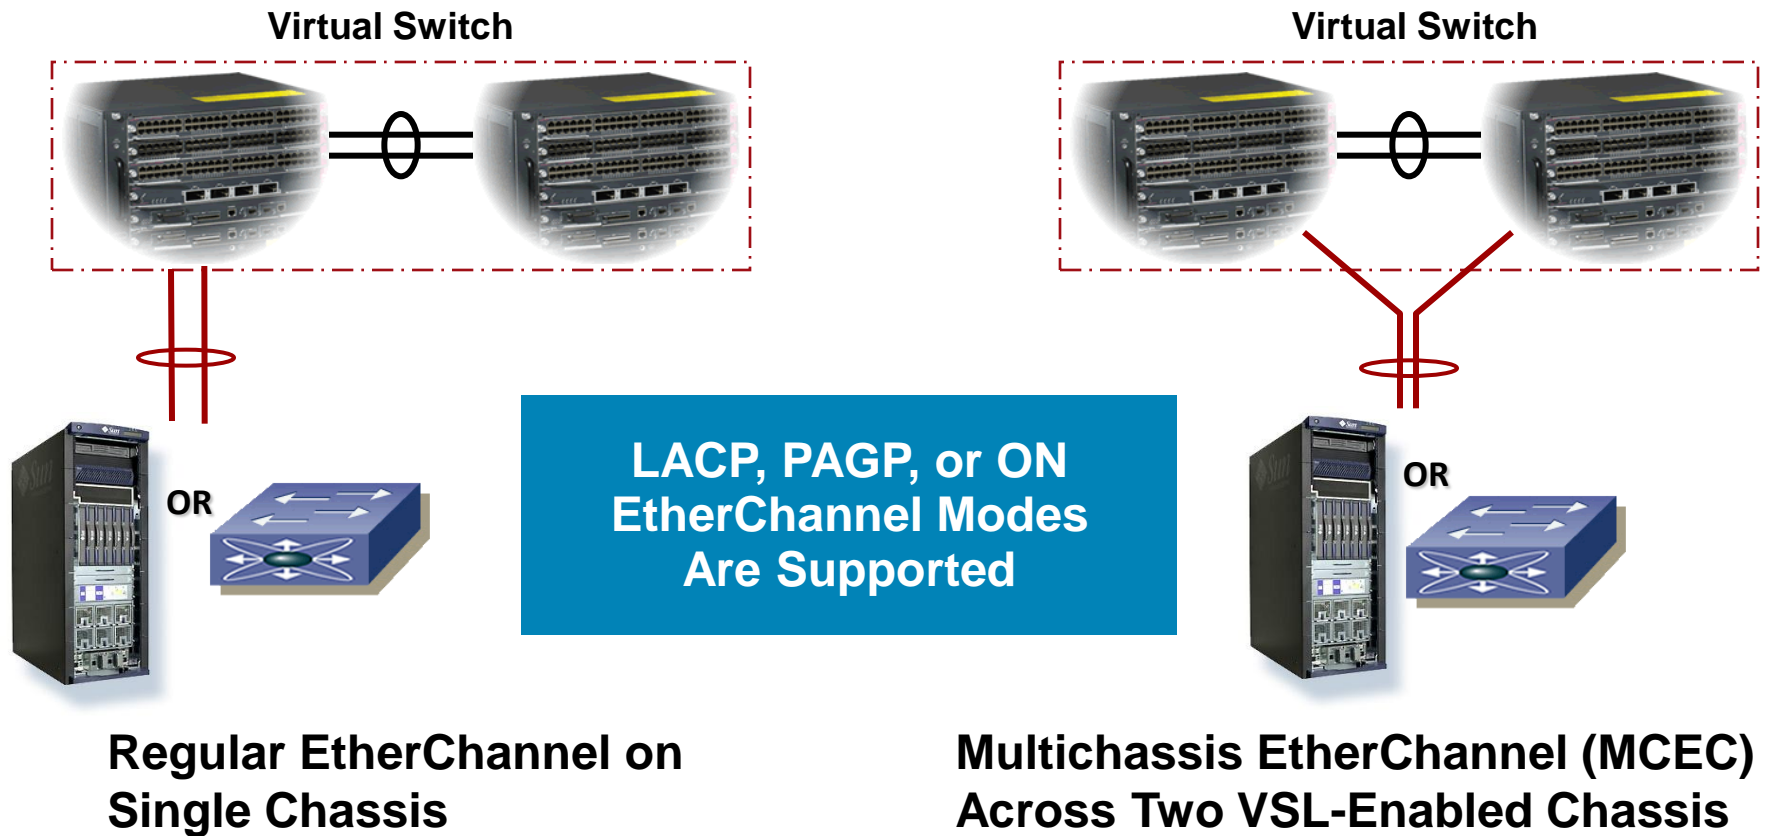
Policy Plane	The business glue of the network. Rules execution, decision making, Service Manager and all the other components to make a productize service.
Services Plane	Overlay “Layer 7” application flow built on the foundation of the other layers. Dependent on the other layers.
Management Plane	The management plane is the logical path of all traffic related to the system management of the platform.
Control Plane	It’s the brain of any networking platform and the technical glue of the network. The control plane is where all routing, switching, other protocols and control information are exchanged
Data Plane	The data plane receives, processes, and transmits network data between network elements, and represents the bulk of network traffic that passes to and through the gear.

Catalyst 6500 Virtual Switch System (VSS)

Forwarding Operation




Multichassis EtherChannel (MCEC)



Virtual Port Channel (vPC)

Active – Active Layer 2 Links



Bi-sectional BW with vPC

vPC is a Port-channeling concept extending link aggregation to two separate physical switches

Allows the creation of resilient L2 topologies based on Link Aggregation.

Eliminates the need for STP in the access-distribution Layer

Enable seamless VM Mobility, Server HA Clusters

Scale Available Layer 2 Bandwidth


Dual-homed server operate in active-active mode

Simplify Network Design


Available on Nexus 7000 and Nexus 5000

Cisco FabricPath

Scaling and Simplifying Layer 2 Ethernet Networks



Traditional Spanning Tree Based Network
-Blocked Links



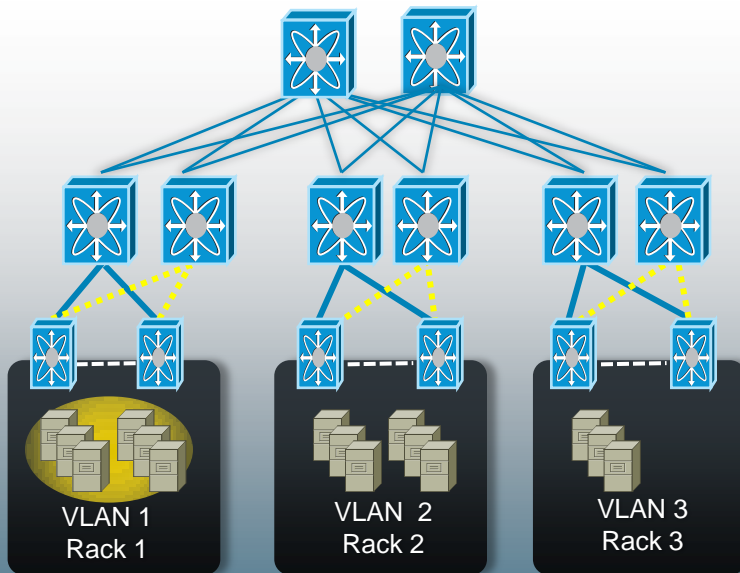
Cisco FabricPath Network
-All Links Active

Up to 16 Agg
switches
160+ Tbps
switching capacity

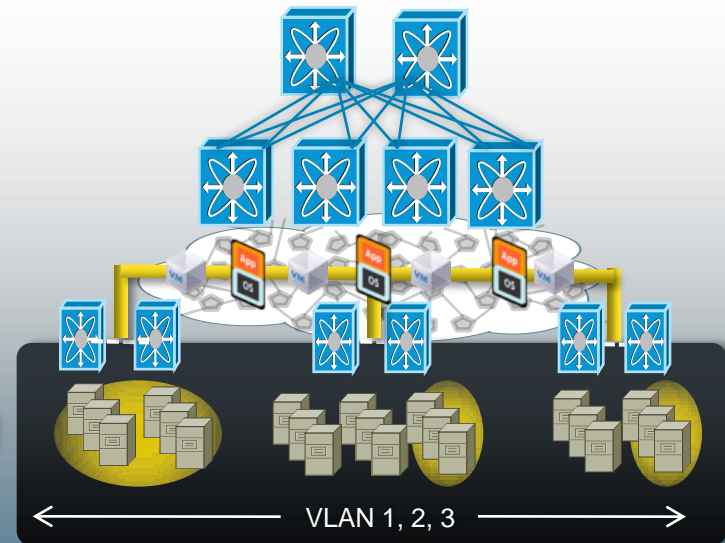
- Eliminate Spanning tree limitations
- Multi-pathing across all links, high cross-sectional bandwidth
- High resiliency, faster network re-convergence
- Any VLAN, any where in the fabric eliminate VLAN Scoping

Example: Handling Application Growth

Scenario: Application grows beyond currently compute capacity and allocated rack space causing network disruptions and physical changes






- Adding additional server capacity while maintaining layer 2 adjacencies in same VLAN
- Disruptive - Requires physical move to free adjacent rack space



- VLAN Extensibility – any VLAN any where!
- Location independence for workloads
- Consistent, predictable bandwidth and latency with Cisco FabricPath.

Cisco's Nexus Architecture Flexibility



Active Paths

Single

Dual

16 Way

POD Bandwidth

Up to 10 Tbps

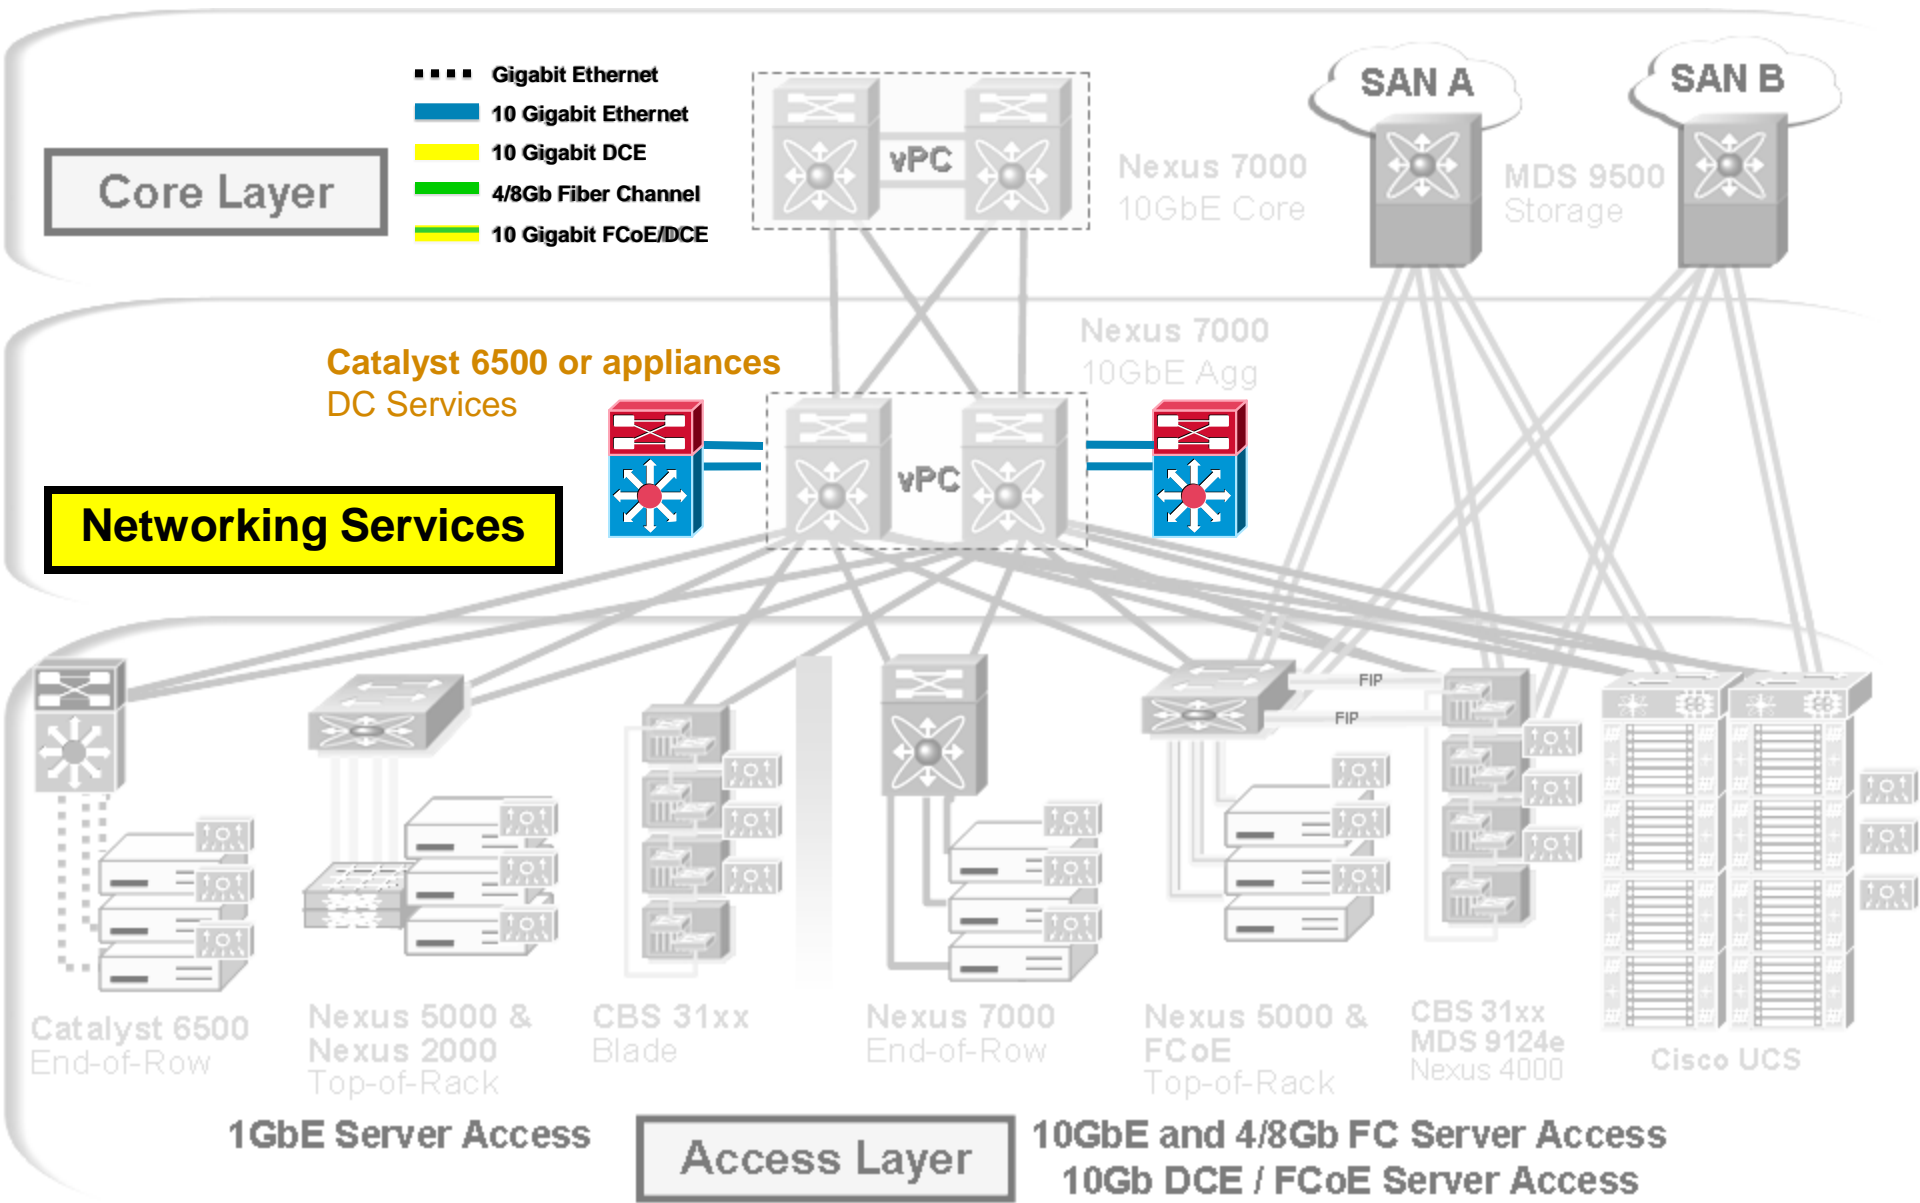
Up to 20 Tbps

Up to 160 Tbps

Layer 2 Scalability

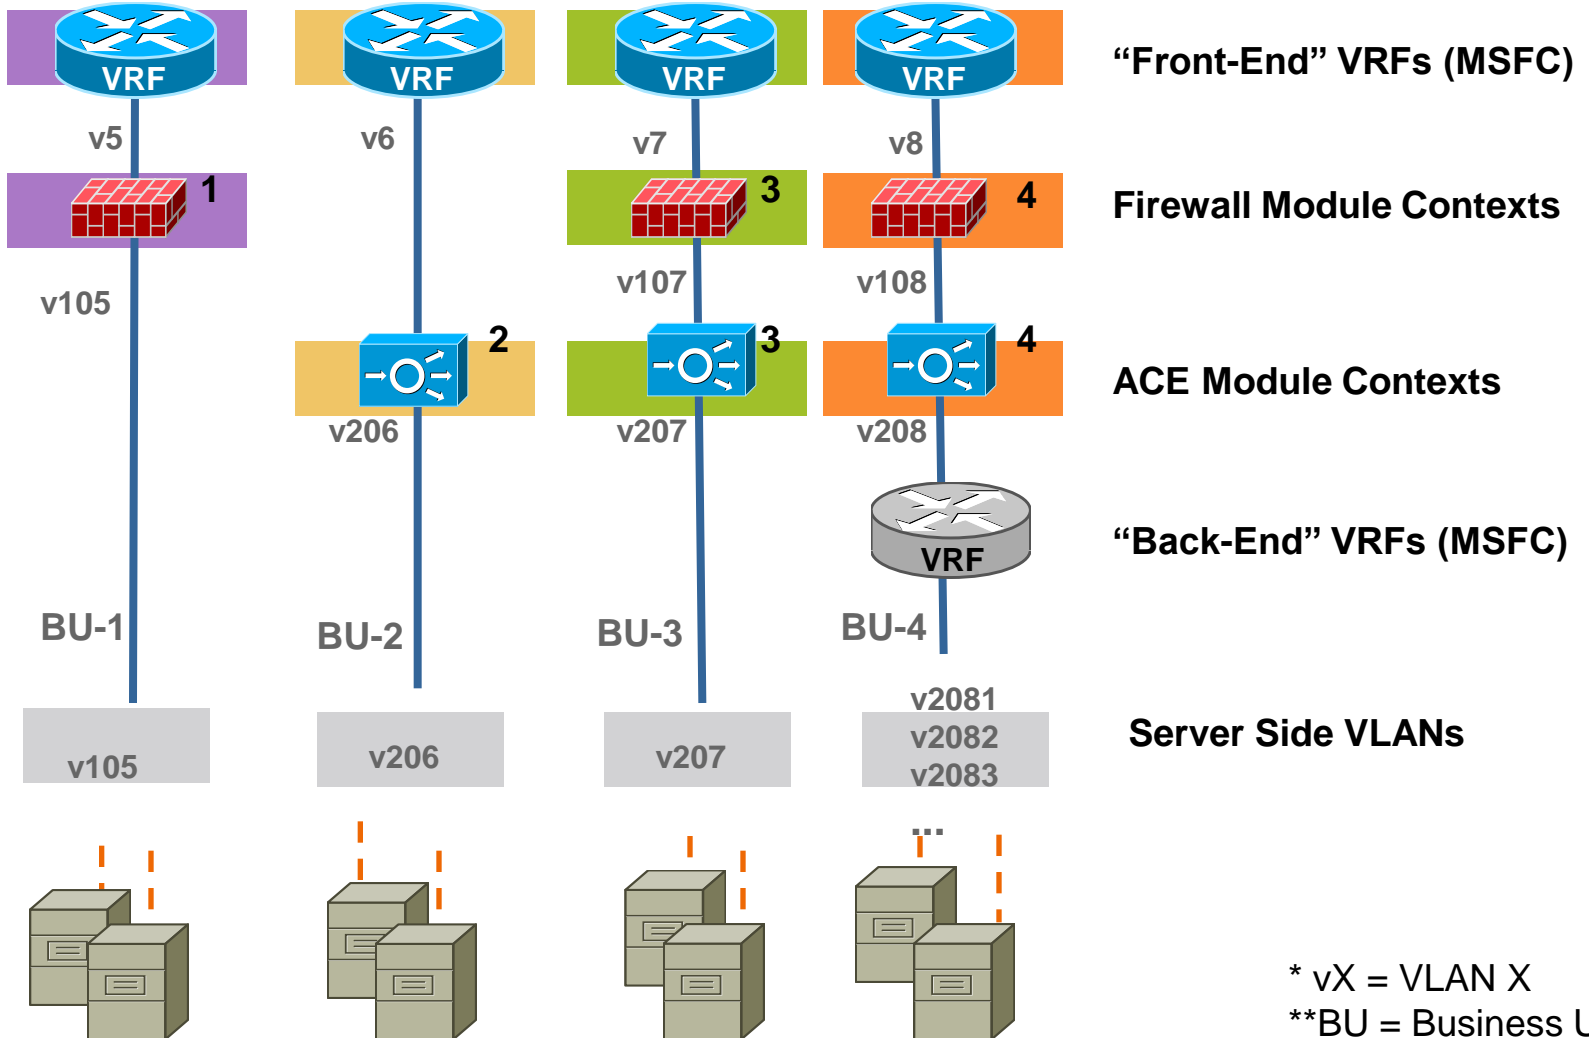
Infrastructure Virtualisation and Capacity

Front-End: Networking Services



Data Centre Virtualised Services

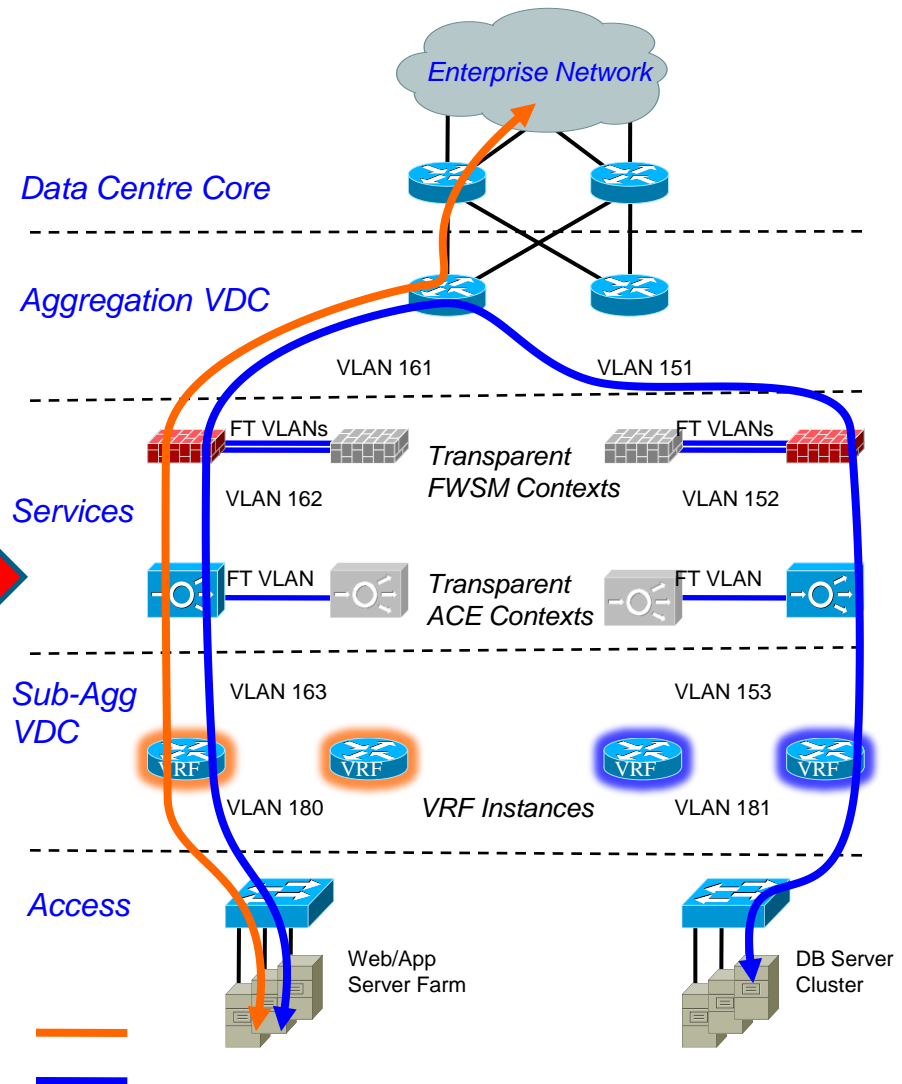
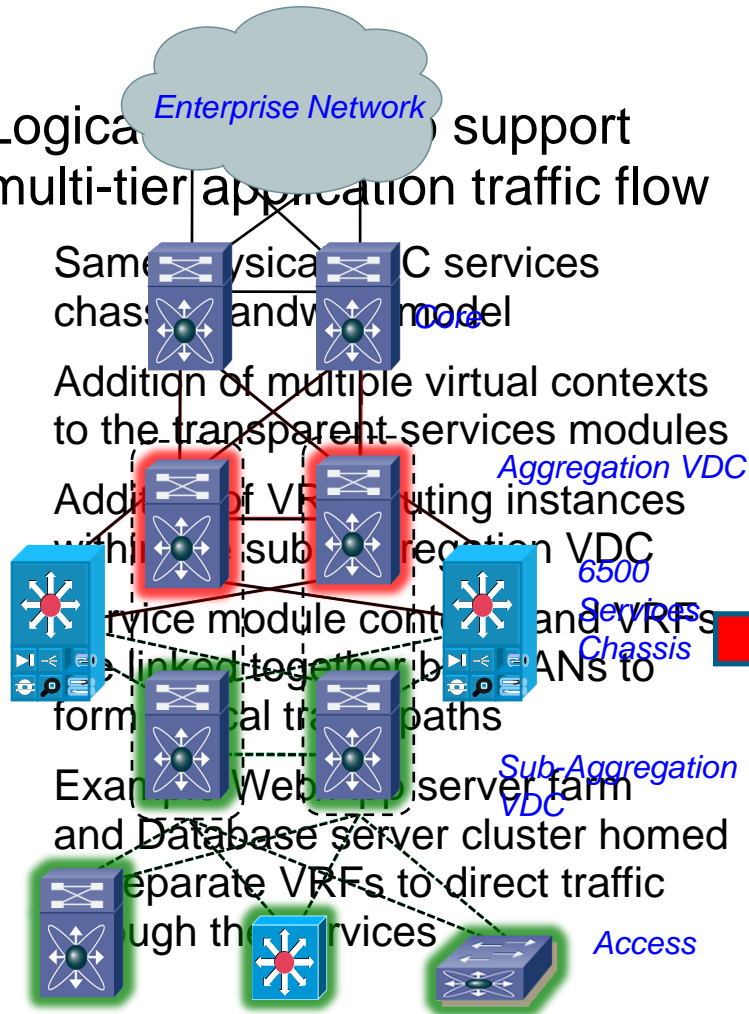
Combination Example



* vX = VLAN X
 **BU = Business Unit

Using Virtualisation and Service Insertion

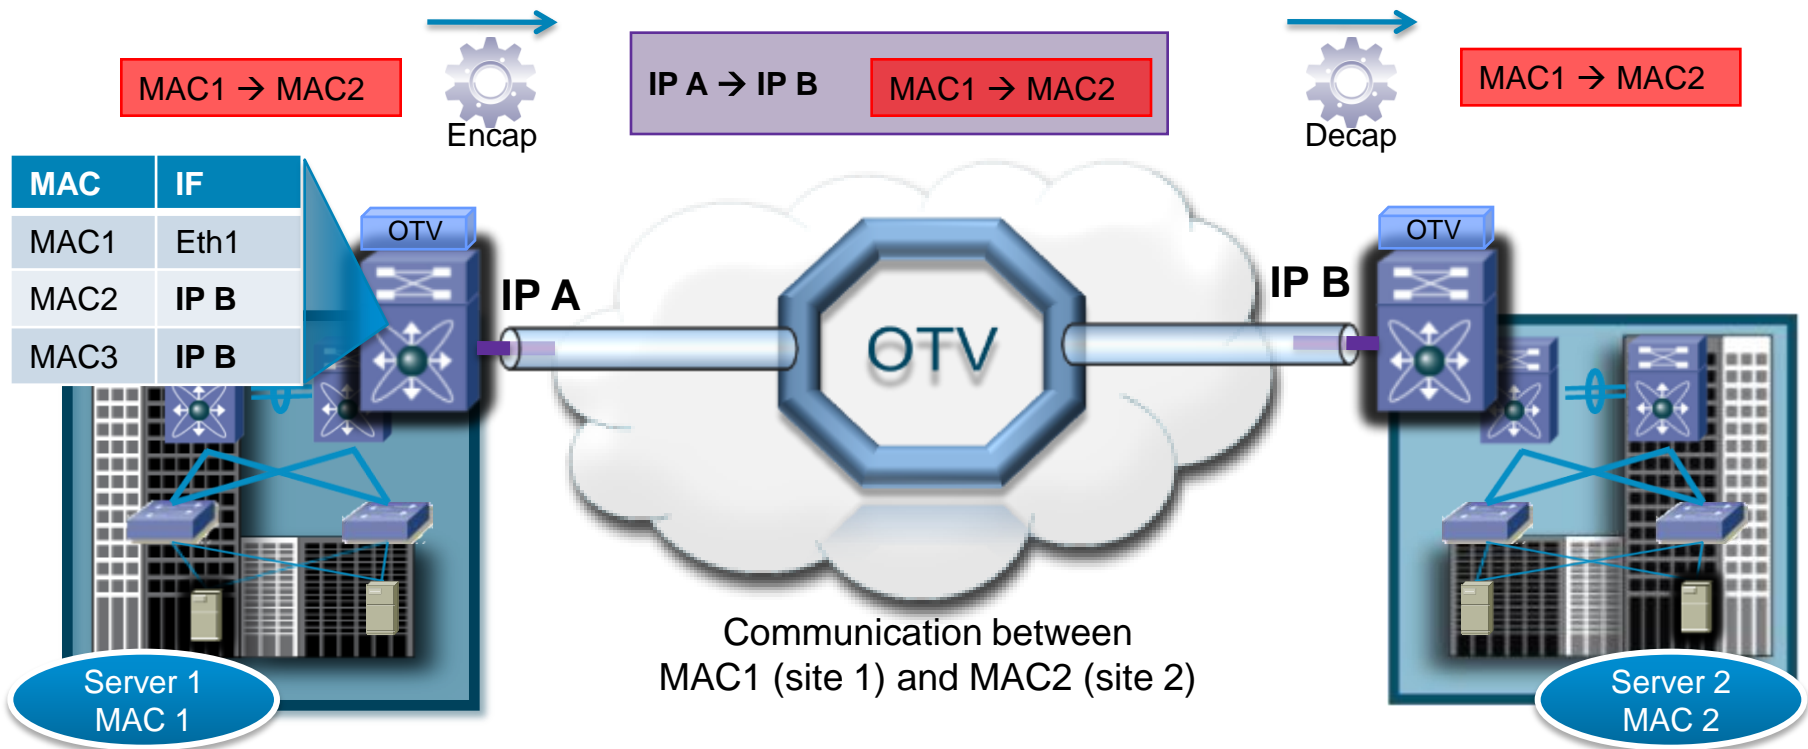
- Logical Enterprise Network support multi-tier application traffic flow



Overlay Transport Virtualisation (OTV)

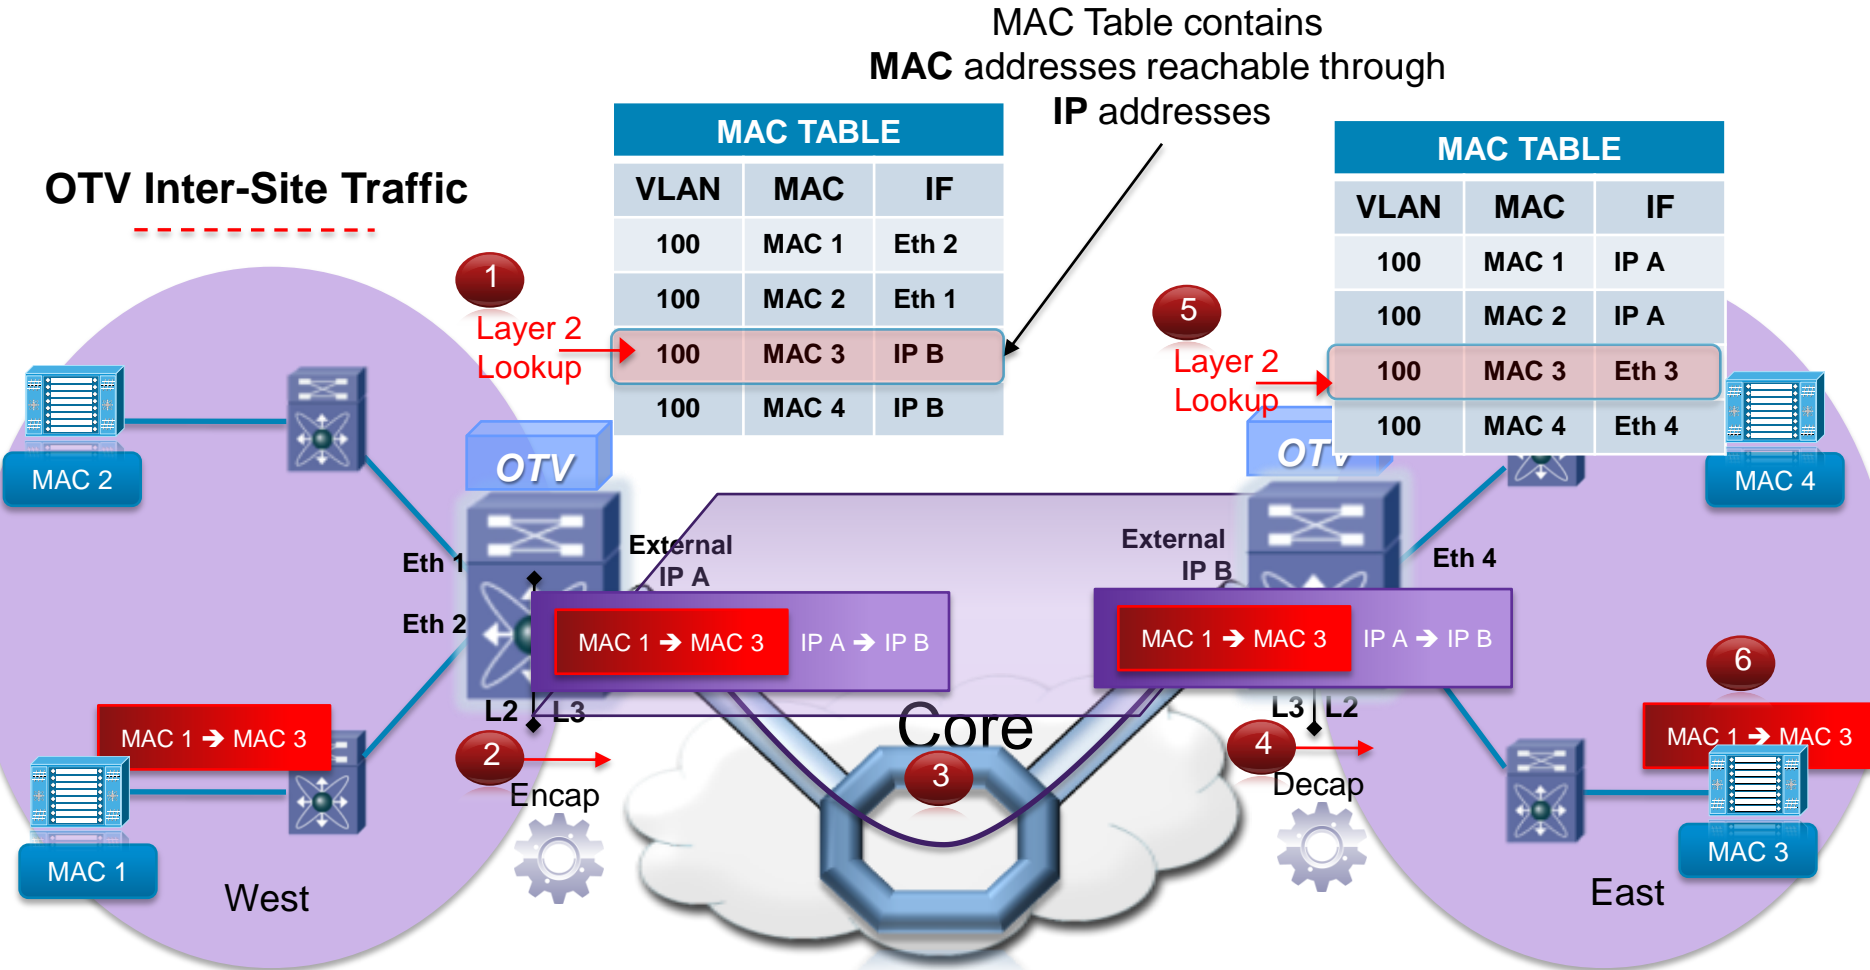
OTV at a Glance

- Ethernet traffic between sites is encapsulated in IP: “MAC in IP”
- Dynamic encapsulation based on MAC routing table
- No Pseudo-Wire or Tunnel state maintained



OTV Data Plane: Unicast

OTV Inter-Site Traffic



- No Pseudo-Wire state is maintained.
- The encapsulation is done based on a Layer 2 destination lookup.


Fabric Scalability beyond the POD

Common Infrastructure for flexible deployment models

Cisco Innovation towards an end-end Fabric:

- Cisco FabricPath: Scalable Fabric for Application Deployment Flexibility
- OTV : Layer 2 extensions over Layer 3 for Distributed Clustered Applications
- LISP: IP mobility, optimized routing

Data Centre Interconnect Extension Overlay Transport Virtualisation (OTV)

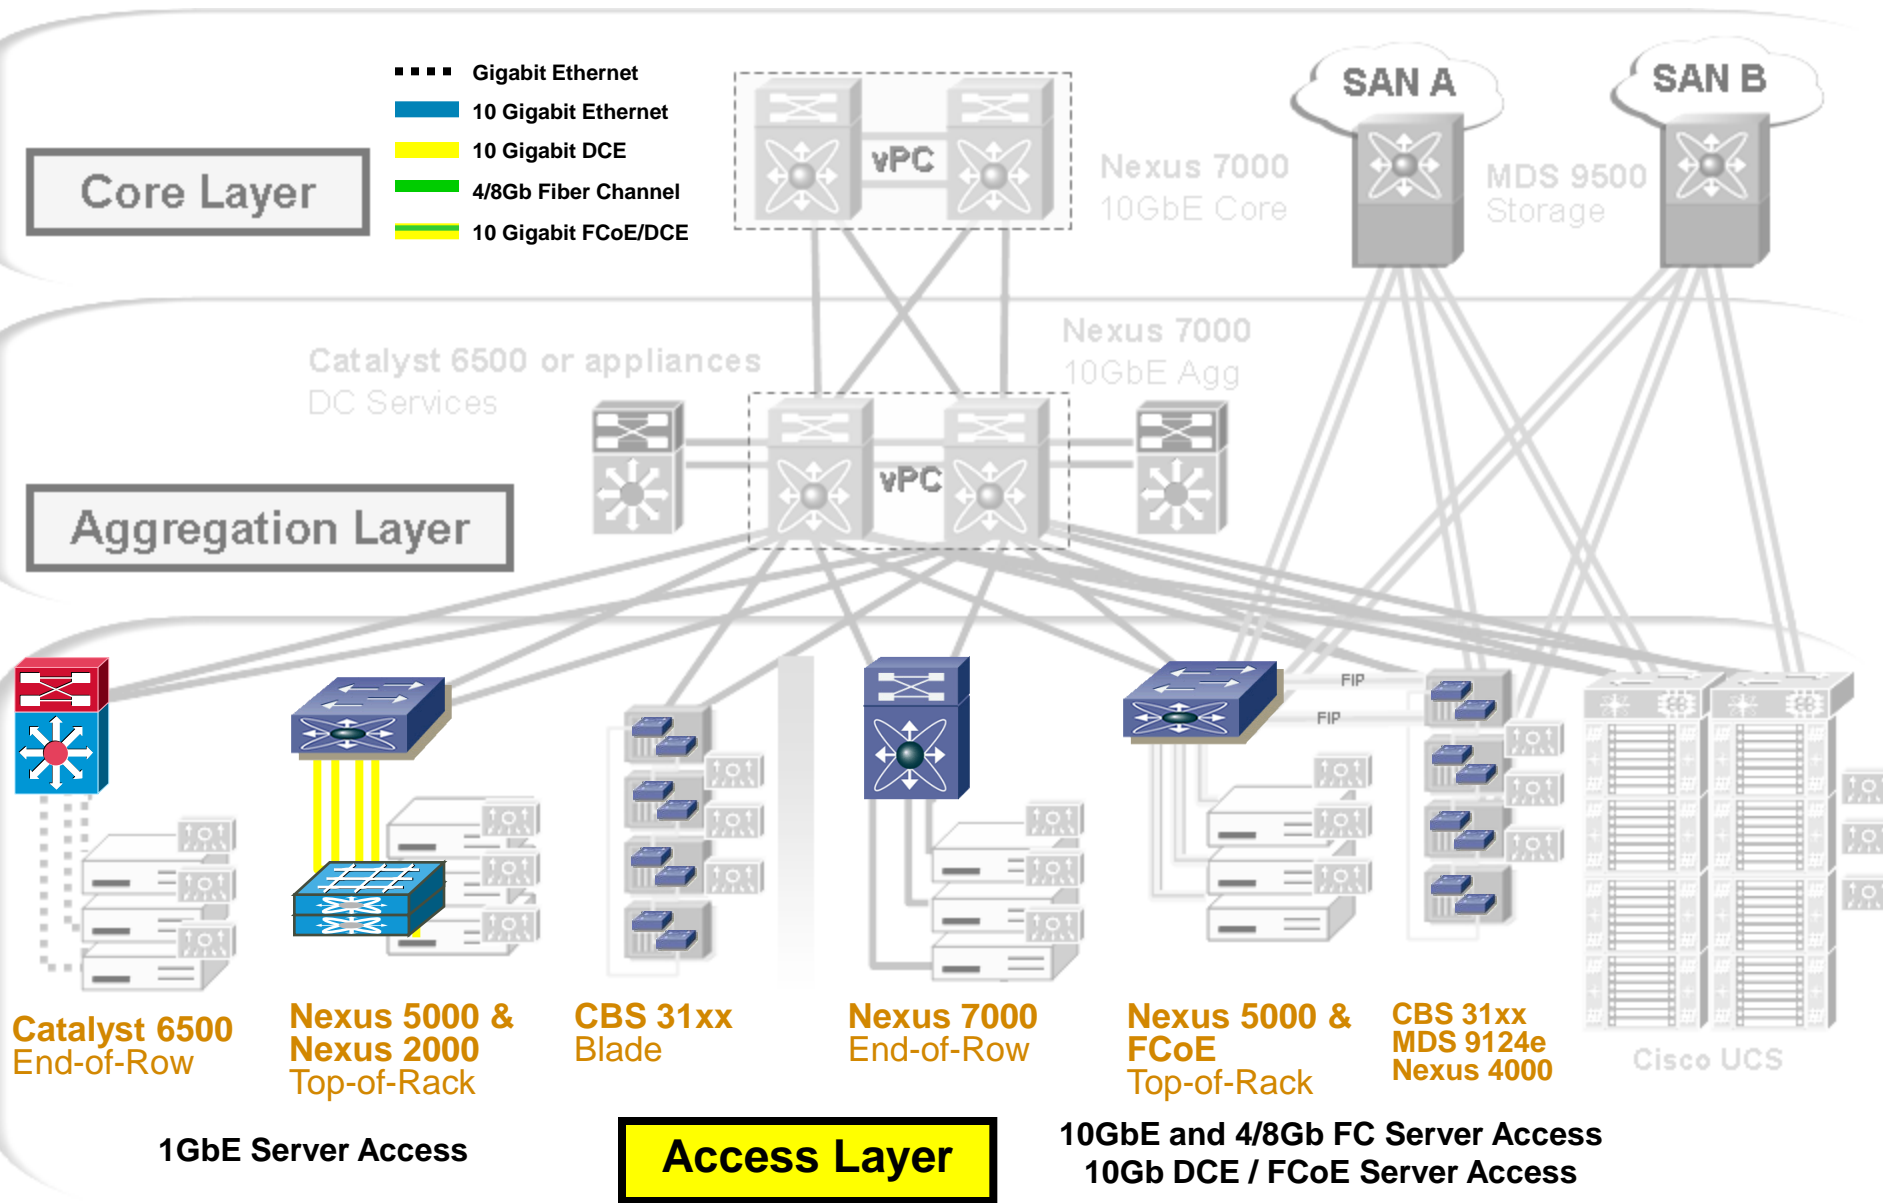


Classical Pod
Spanning Tree Protocol


Scalable Pod
vPC & FEX

Highly Scalable Pod
Cisco FabricPath

Front-End: Access Layer



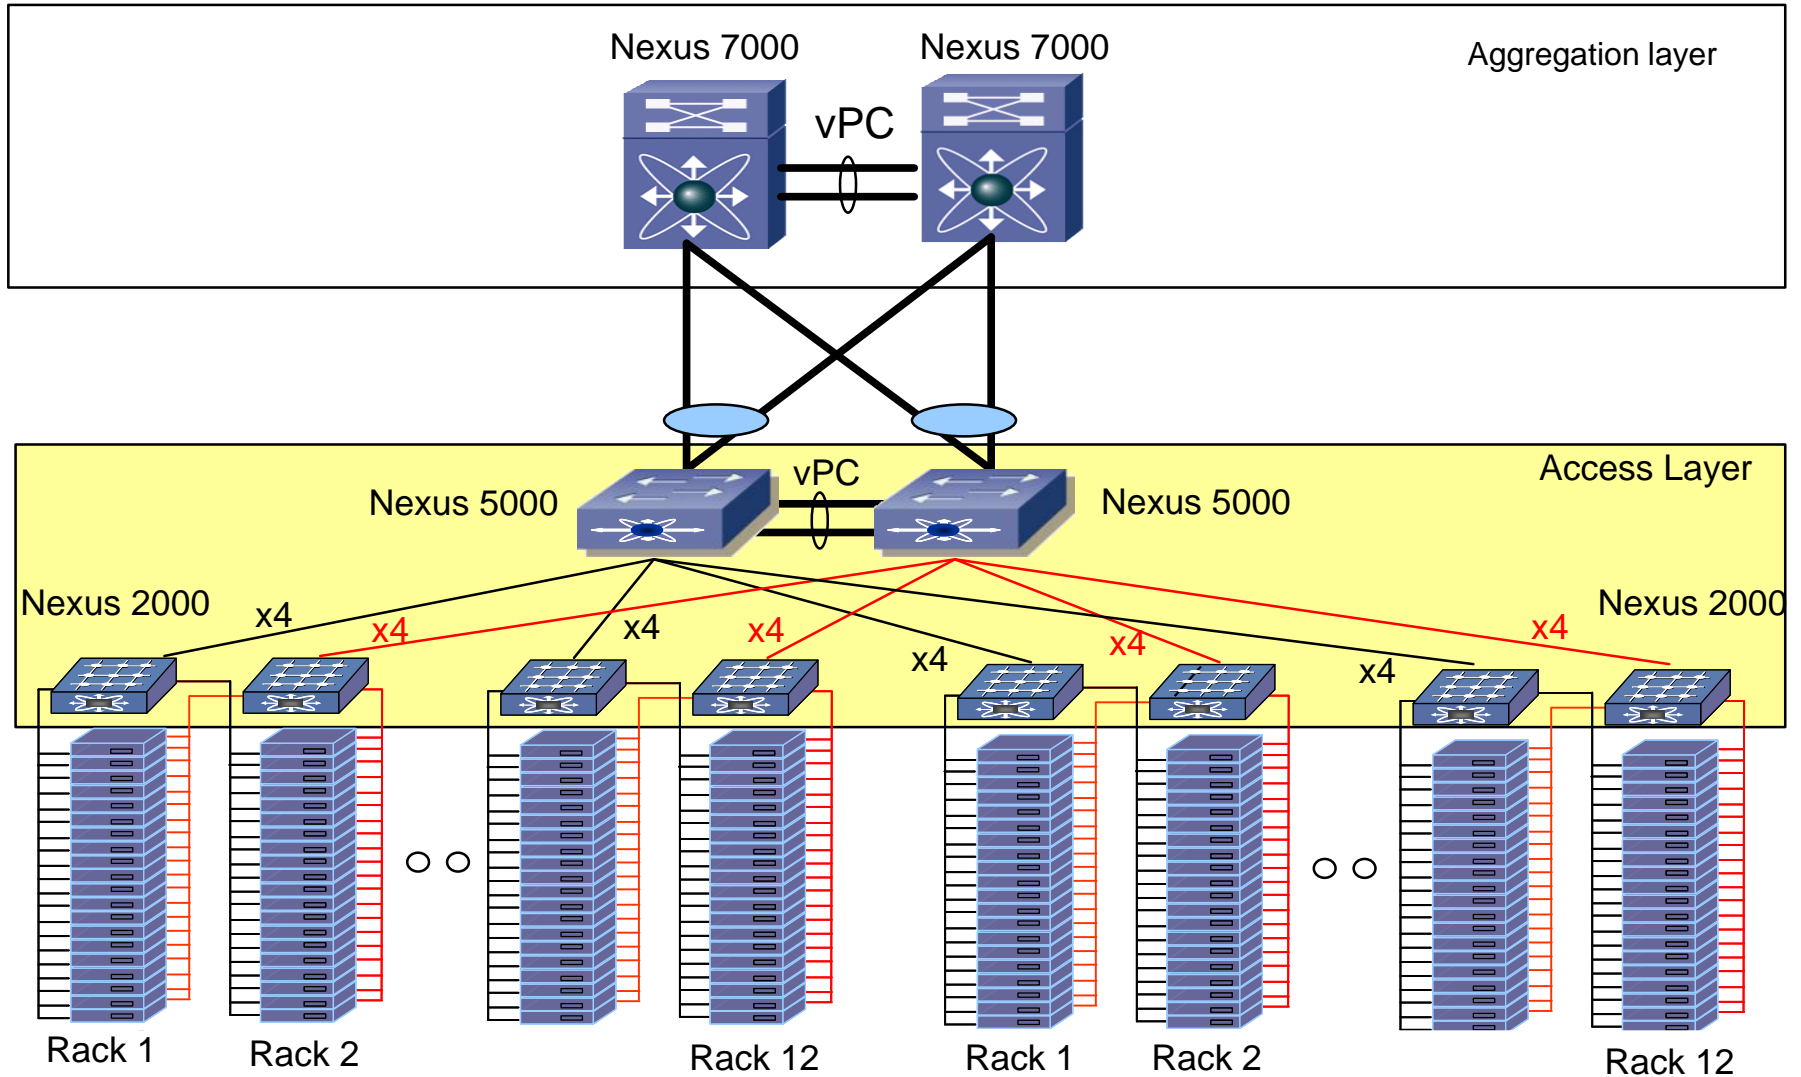
Top of Rack @ 1/10GE: Nexus 2000 (Fabric Extender - FEX)



Nexus 5000 Fabric Extender (FEX)

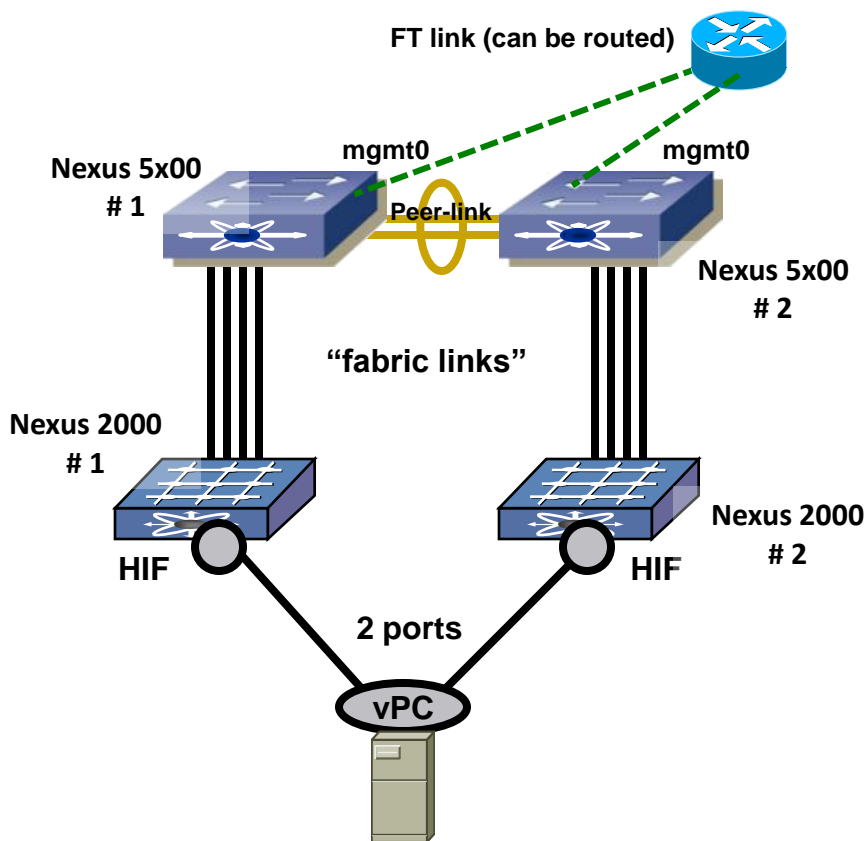
The Nexus 5000 Fabric Extender (FEX) acts as a remote line card (module) for the Nexus 5000, retaining all centralized management and configuration on the Nexus 5000, transforming it to a Virtualized Chassis

To2R: Nexus 2000 Deployment example



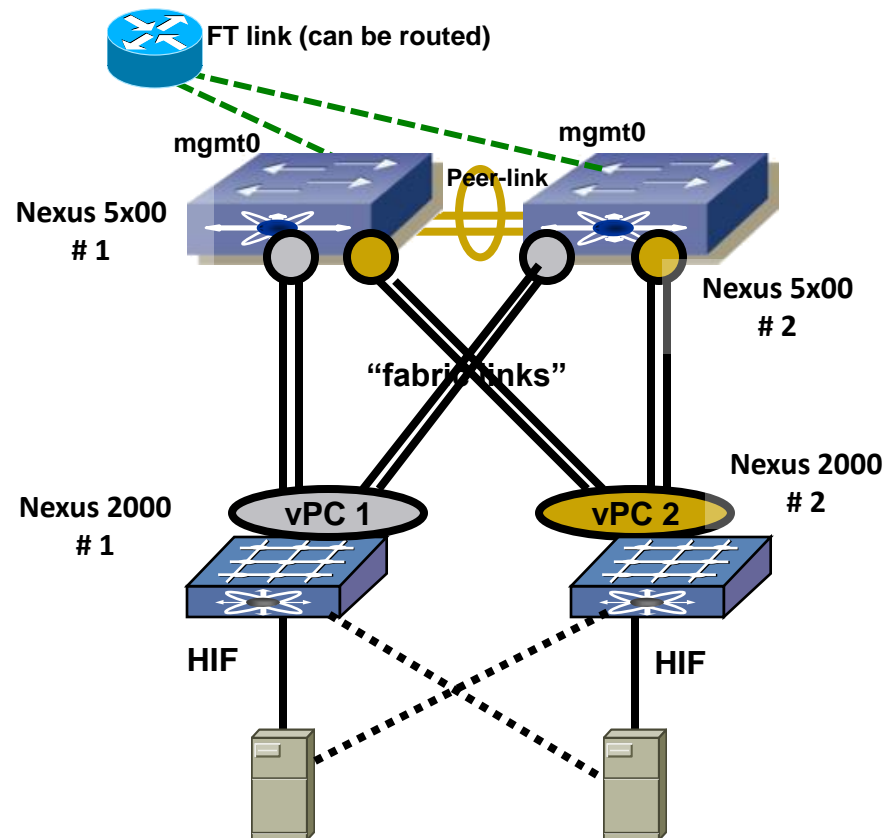
Nexus 2000 with vPC on Nexus 5000/5500

Nexus 2000 **Single-homed** vPC




2-GigE ports host port channel

Nexus 2000 **active/active (or dual homed)**




vPC between Nexus 7000 and Nexus 5x00 and Nexus 2000 – Active/Active




Virtualised Access Switch POD

- Cisco Nexus 5x00 and 2000 represent a virtual access switch POD
- Nexus 7000 at Aggregation Layer

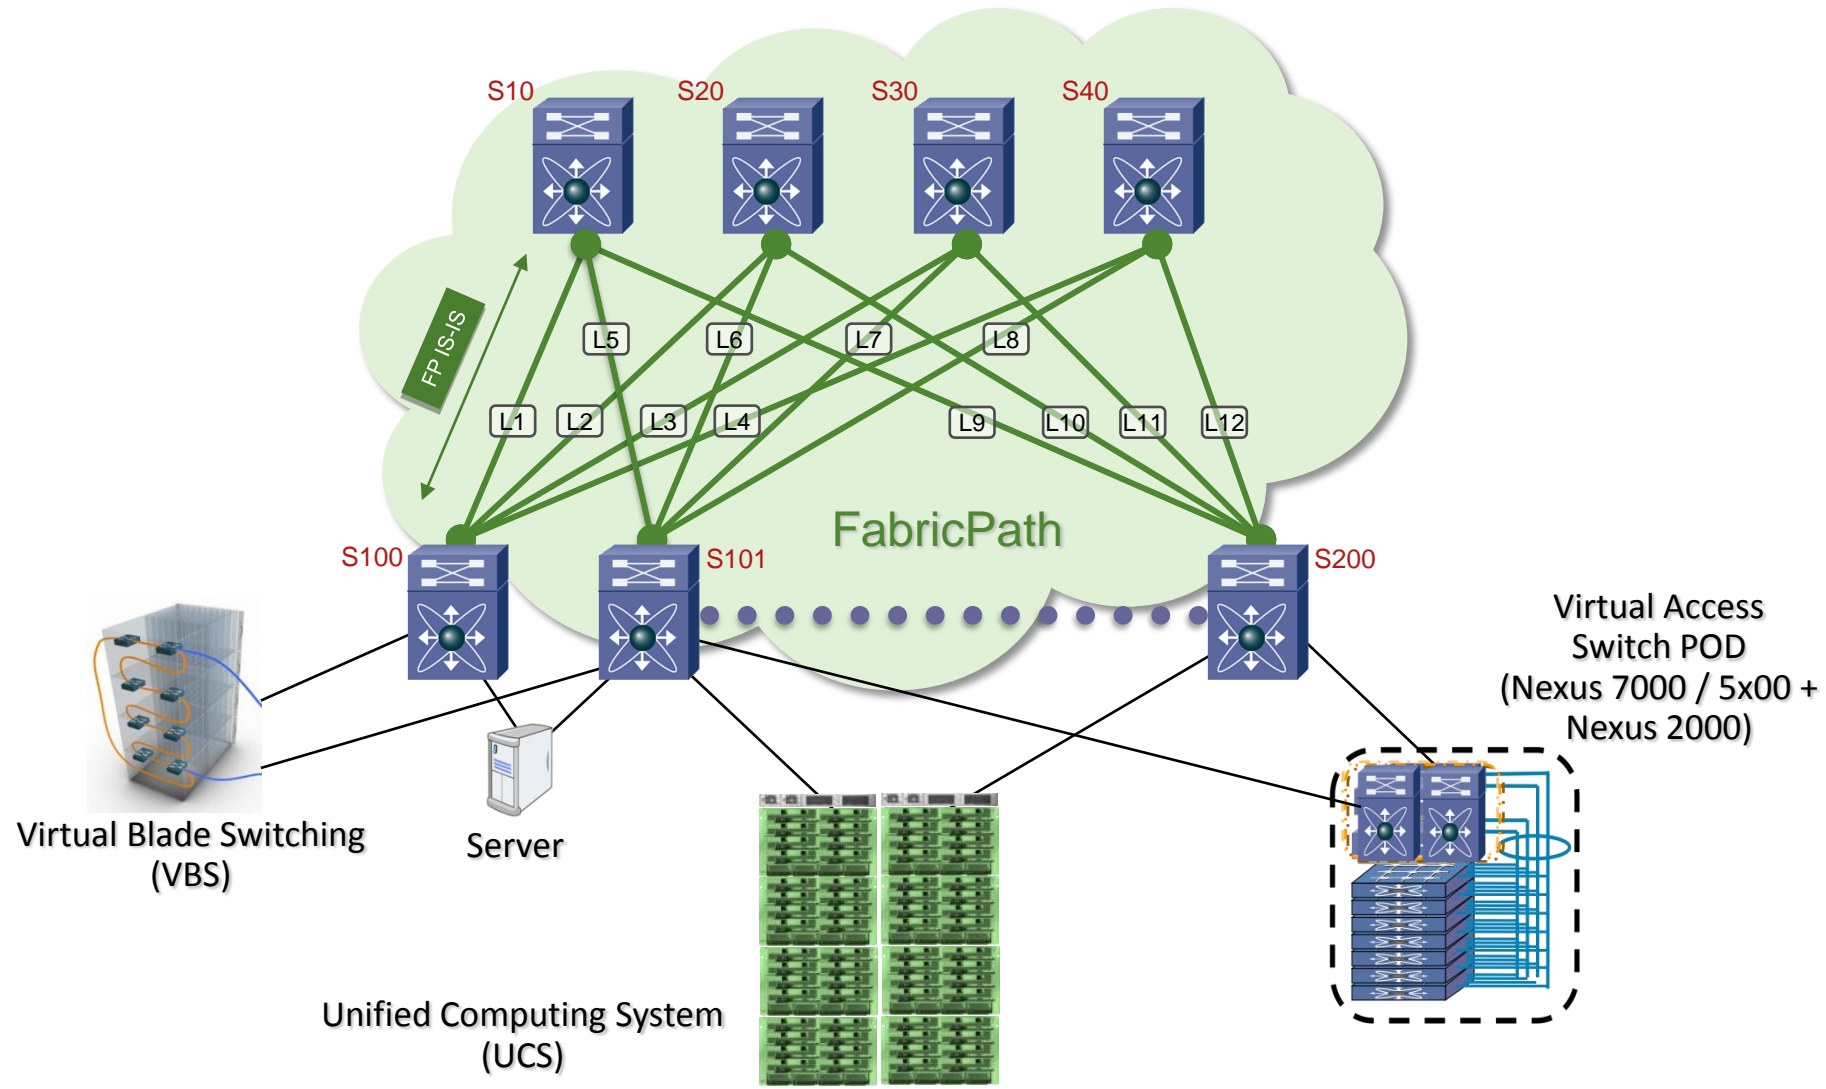


**Nexus 5x00/2000
Virtualised Access
Switch PODs**

Logical View: Star Topology without L2 loops

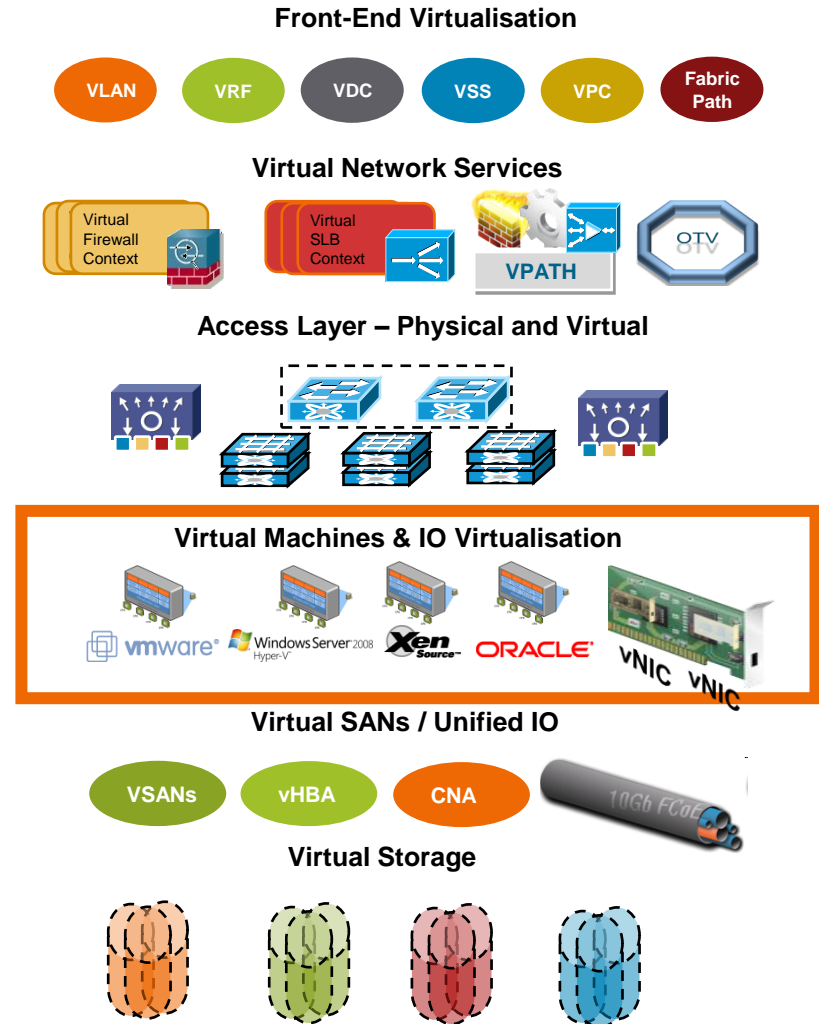


Logical View with FabricPath: Distributed Topology without L2 loops

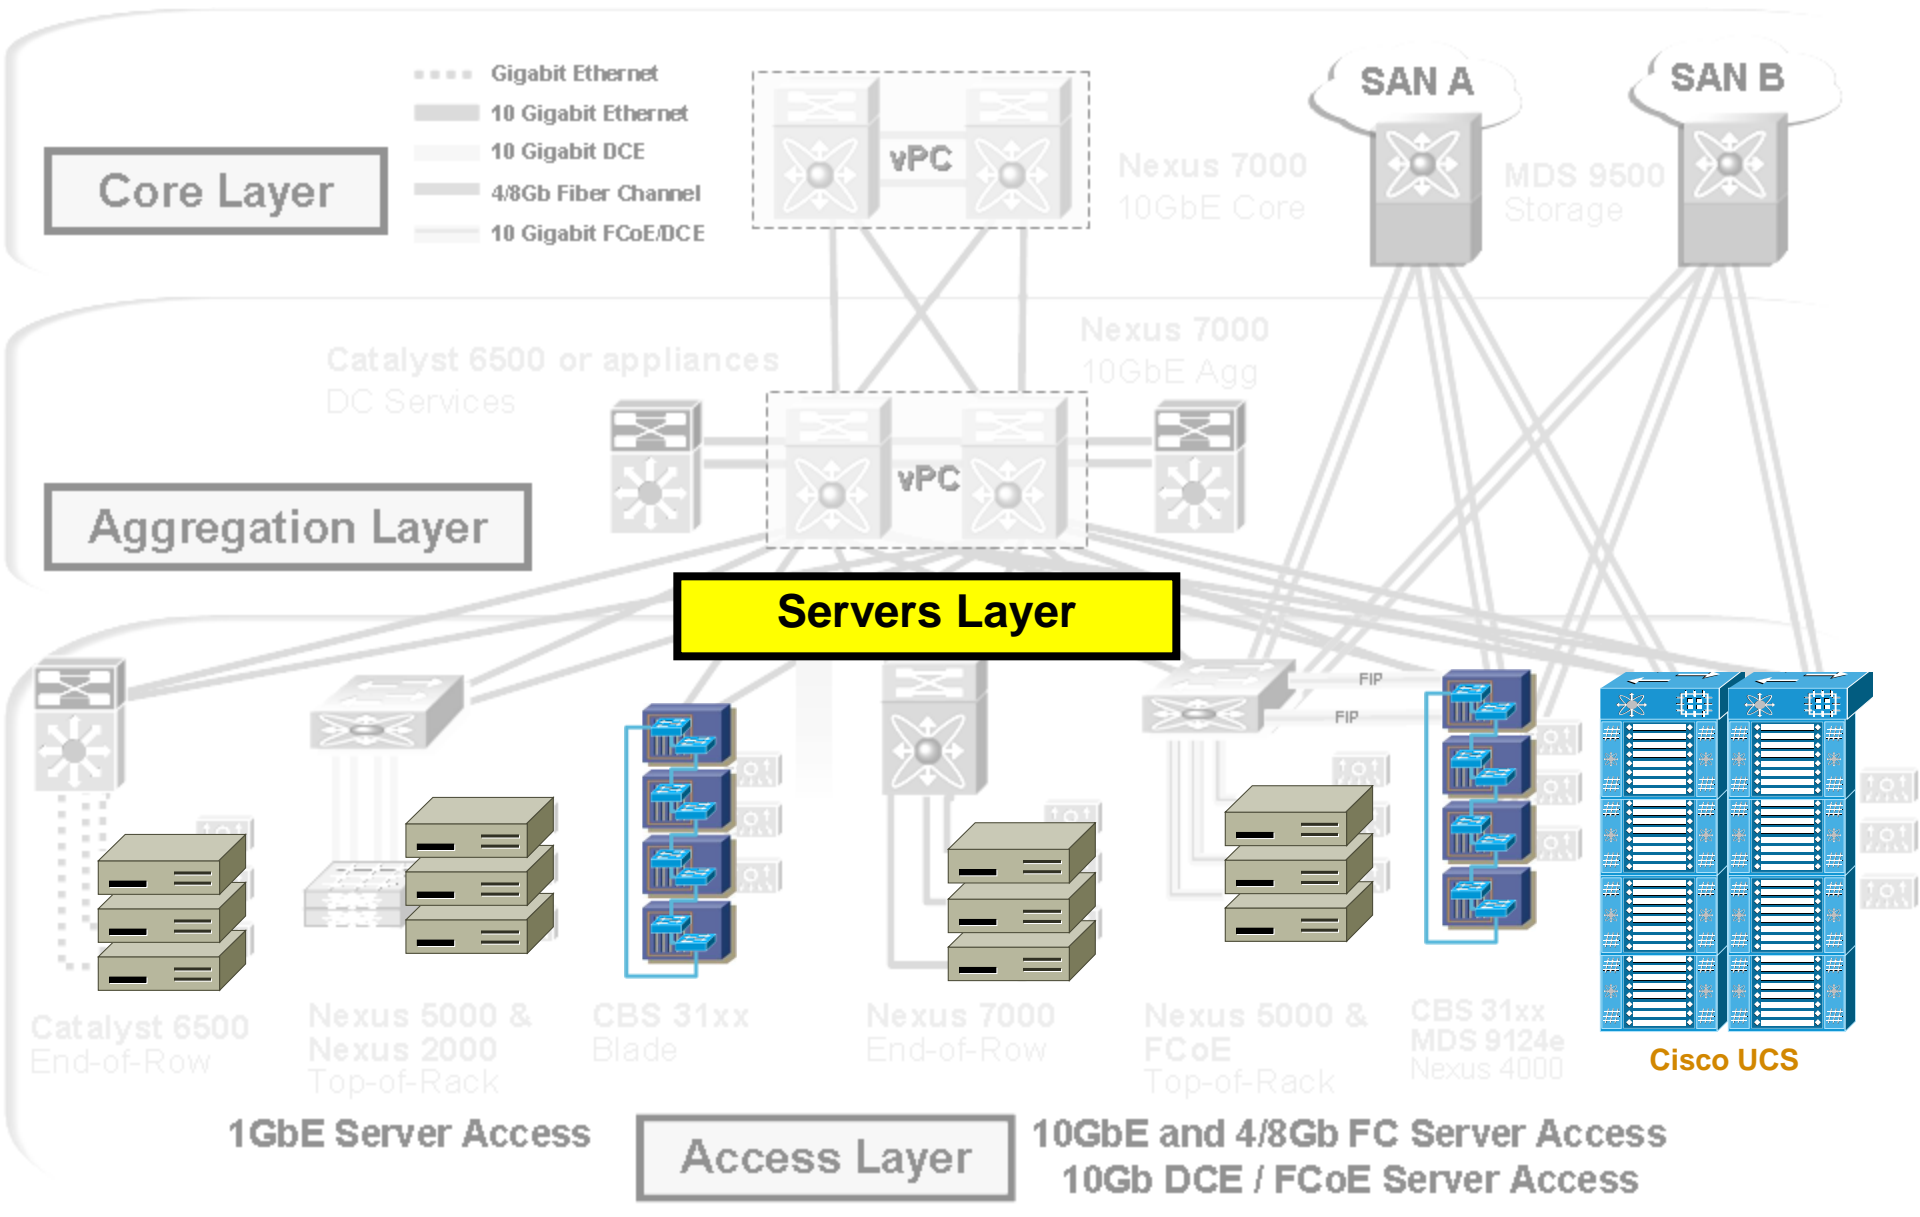


Agenda

- Data Centre Virtualisation Overview
- Front-End Data Centre Virtualisation
 - Core Layer
 - Aggregation Layer
 - Networking Services
 - Access Layer
- **Server Virtualisation**
 - Hypervisors**
 - Virtual Access Layer**
 - Virtualised Services**
 - Server IO Virtualisation**
- Back-End Virtualisation
 - Virtual HBA & NPV
 - Unified IO & FCoE
 - SAN & Storage
- Implementation Examples
 - vBlock
 - Secure Cloud
 - FlexPOD
- Q&A

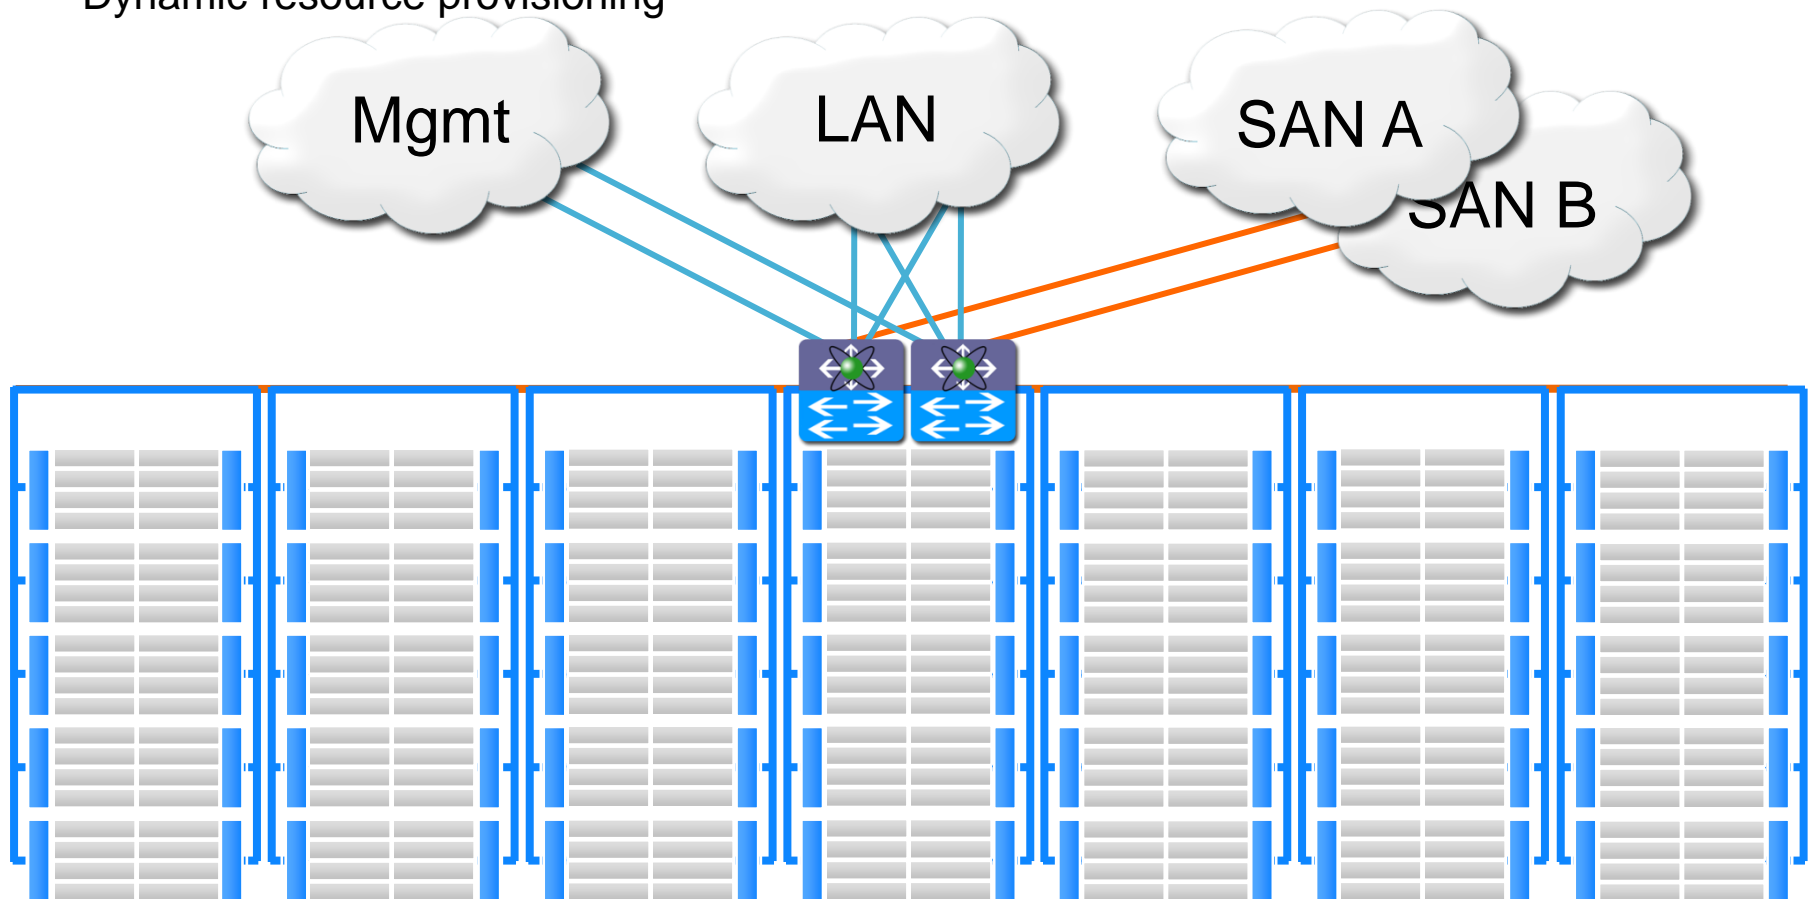


Servers Layer



What is Cisco UCS ?

- UCS = Unified Computing System
- Single, scalable integrated system
- Network + compute Virtualisation
- Dynamic resource provisioning



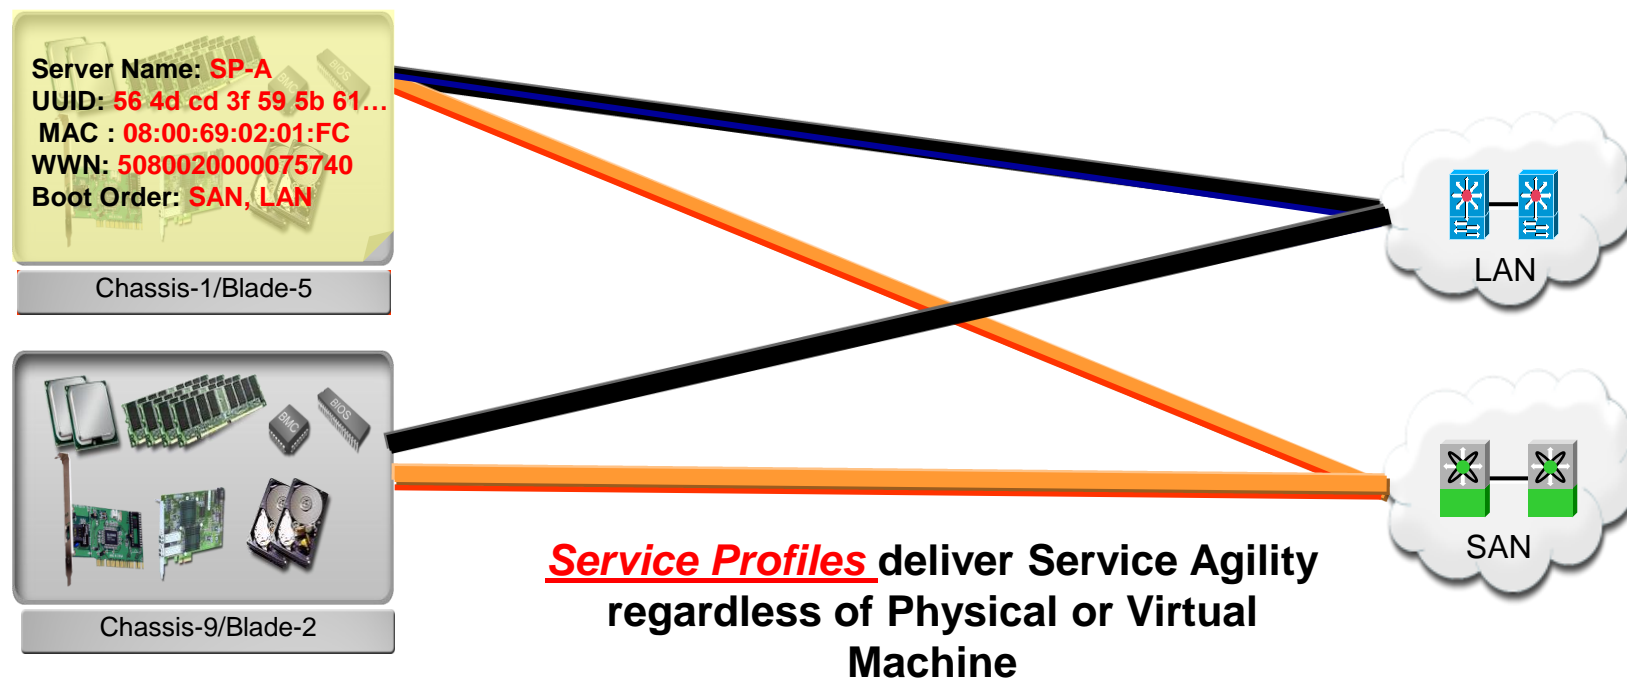
Network has Complete Visibility to Servers

- UCS Service Profiles Capture more than MAC & WWN

MAC, WWN, Boot Order, Firmware, network & storage policy

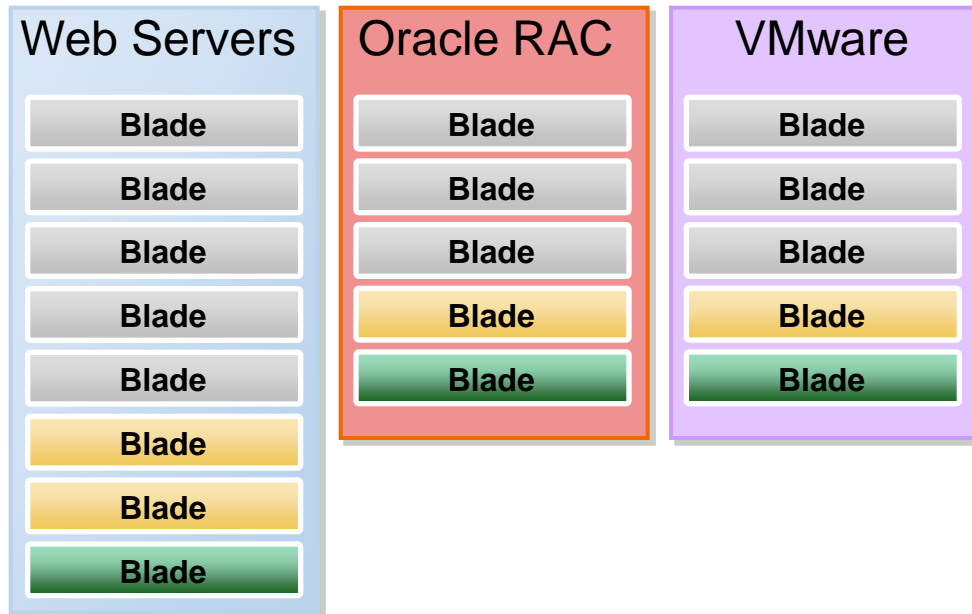
- Stateless compute where network & storage see all movement

Better diagnostics and QoS from network to blade, policy follows



Flexible Data Centre with Service Profiles

Workload	Server Capacity Needed				Server HW HA	Total Servers
	Oct	Nov	Dec	Jan		
Web Servers	5	7	6	5	1 hot spare	8
Oracle RAC	3	3	3	4	1 hot spare	5
VMware	3	3	4	4	1 hot spare	5



Today's Deployment:

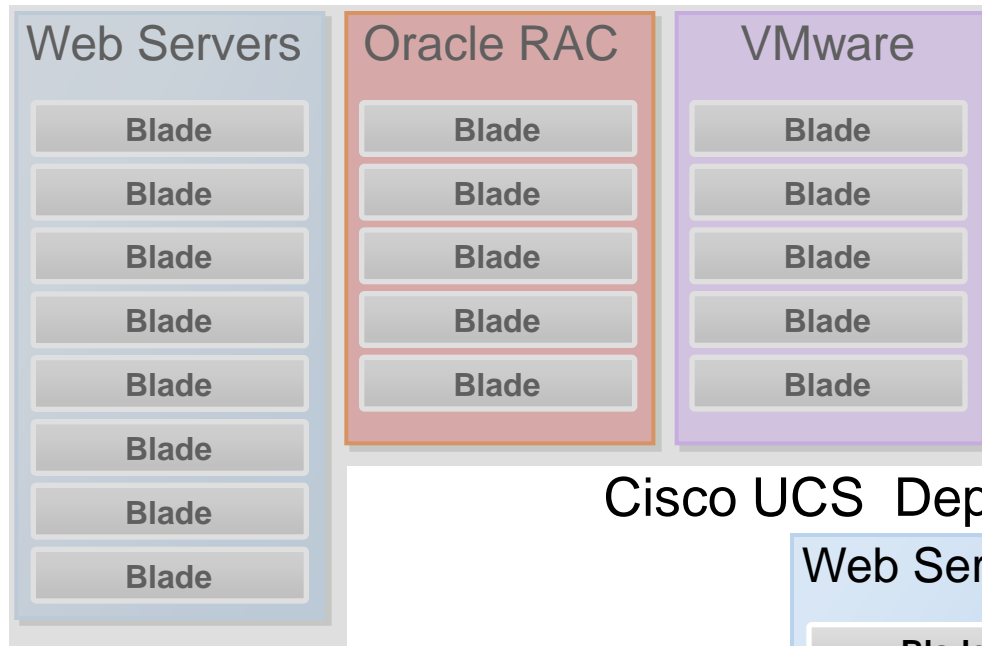
Provisioned for peak capacity

Spare node per workload

**Total Server Deployment
18 Servers**

Stateless Computing @ UCS

Old Deployment:

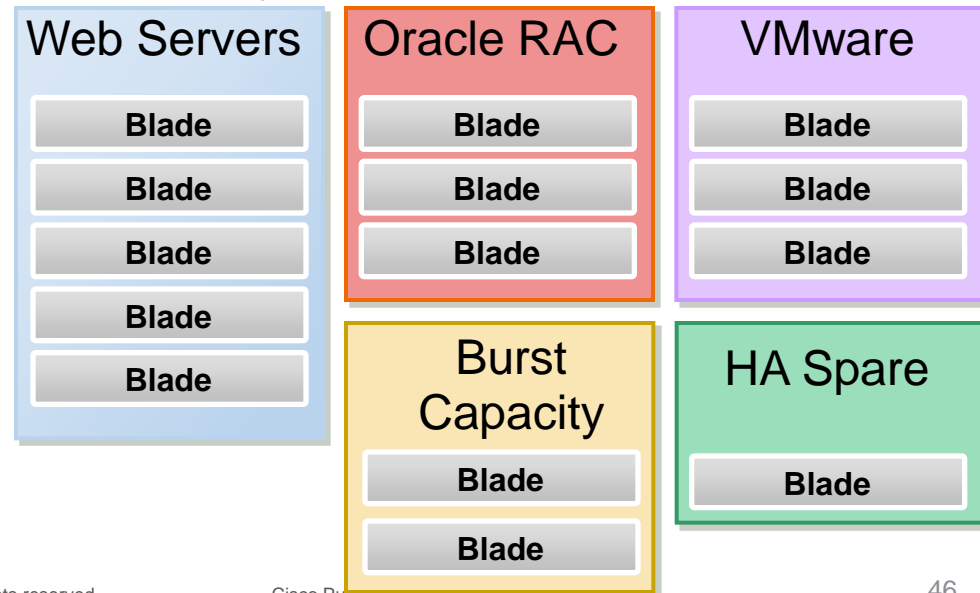


**Total Server Deployment
14 Servers**
**Reduction of 4 Servers
22% CapEx Savings**

Cisco UCS Deployment: (still 18 Service Profiles)

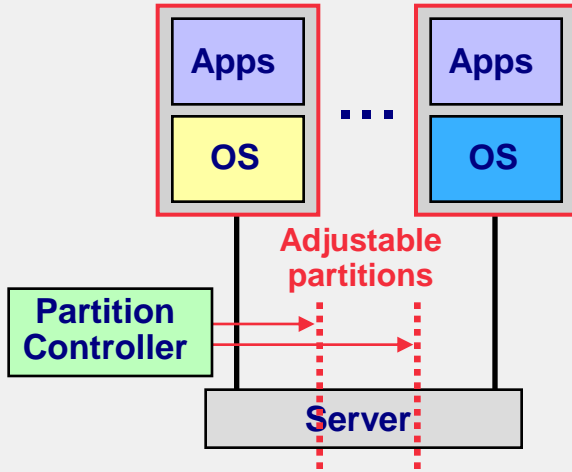
Cisco's Deployment:

- Resources provisioned based on business need
- Still HA with fewer spares



Typical Virtualisation Architectures

Hardware Partitioning




Server is subdivided into fractions each of which can run an OS

Physical partitioning
 IBM S/370 SI->PP & PP->SI,
 Sun Domains, HP nPartitions

Logical partitioning
 System p LPAR, HP vPartitions,
 Sun Logical Domains
 IBM System z LPAR

Dedicated Hypervisor




Hypervisor provides fine-grained timesharing of all resources

Hypervisor software/firmware runs directly on server

VMware ESX Server
Xen Hypervisor , KVM
Microsoft Hyper-V
Oracle VM

Hosted Hypervisor




Hypervisor uses OS services to do timesharing of all resources

Hypervisor software runs on a host operating system


VMware Server
Microsoft Virtual Server
HP Integrity VM
QEMU

What happens when we mix network and server virtualisation ?



Current View of the Access Layer with VMs

Boundary of network visibility




- Typically provisioned as trunk to the server running ESX
- No visibility to individual traffic from each VM
- Unable to troubleshoot, apply policy, address performance issues

Networking for Server Virtualisation


Problems:

- Dynamic Migration of VMs may move them across physical server ports—policy must follow
- Impossible to view or apply policy to locally switched traffic
- Need collaboration between network and Virtualisation admin




Cisco VN-Link

- VN-Link (or Virtual Network Link) is a term that refers to a VM specific link that is created between the VM and Cisco switch.
- Logical equivalent & combination of a NIC, a Cisco switch interface and the RJ-45 patch cable that hooks them together.



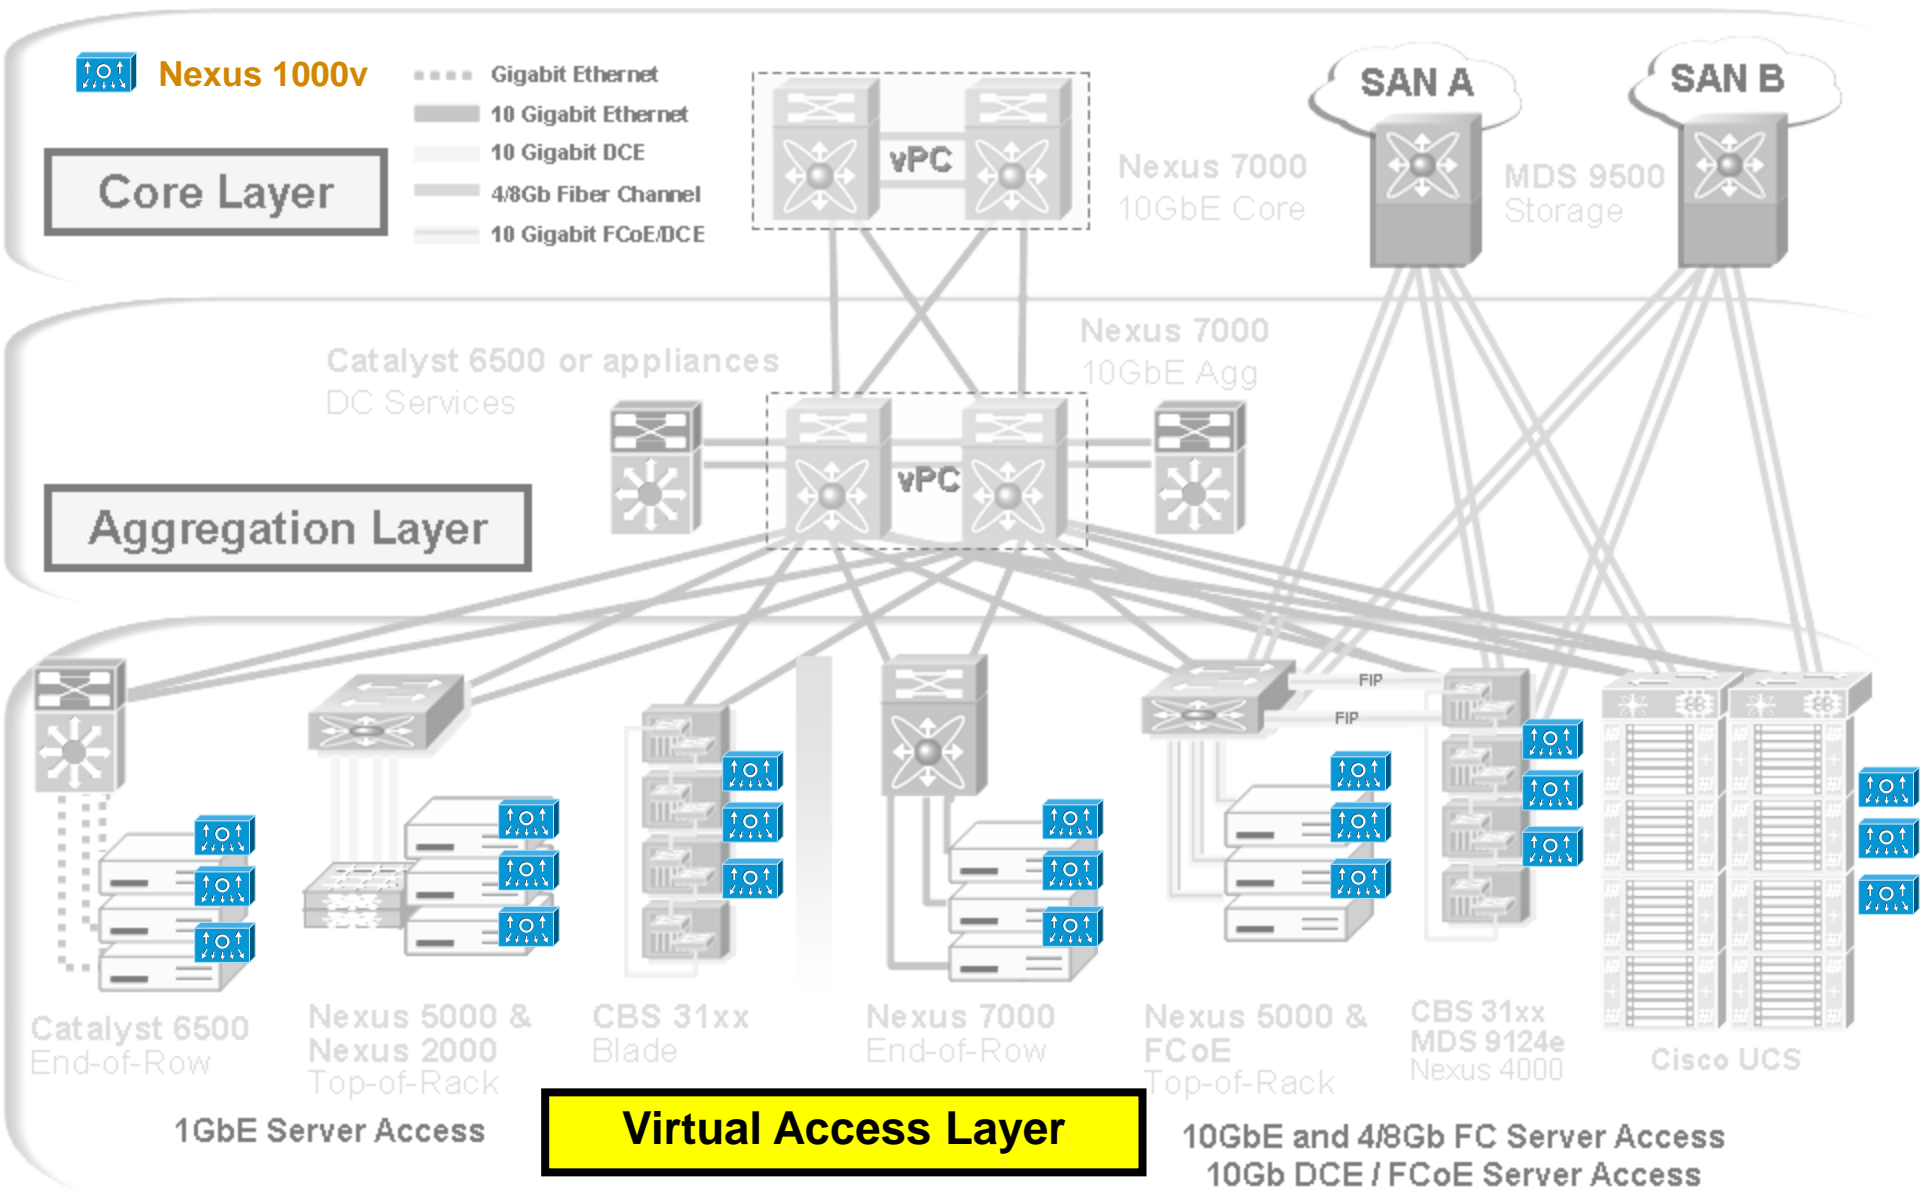
VN-Link View of the Access Layer



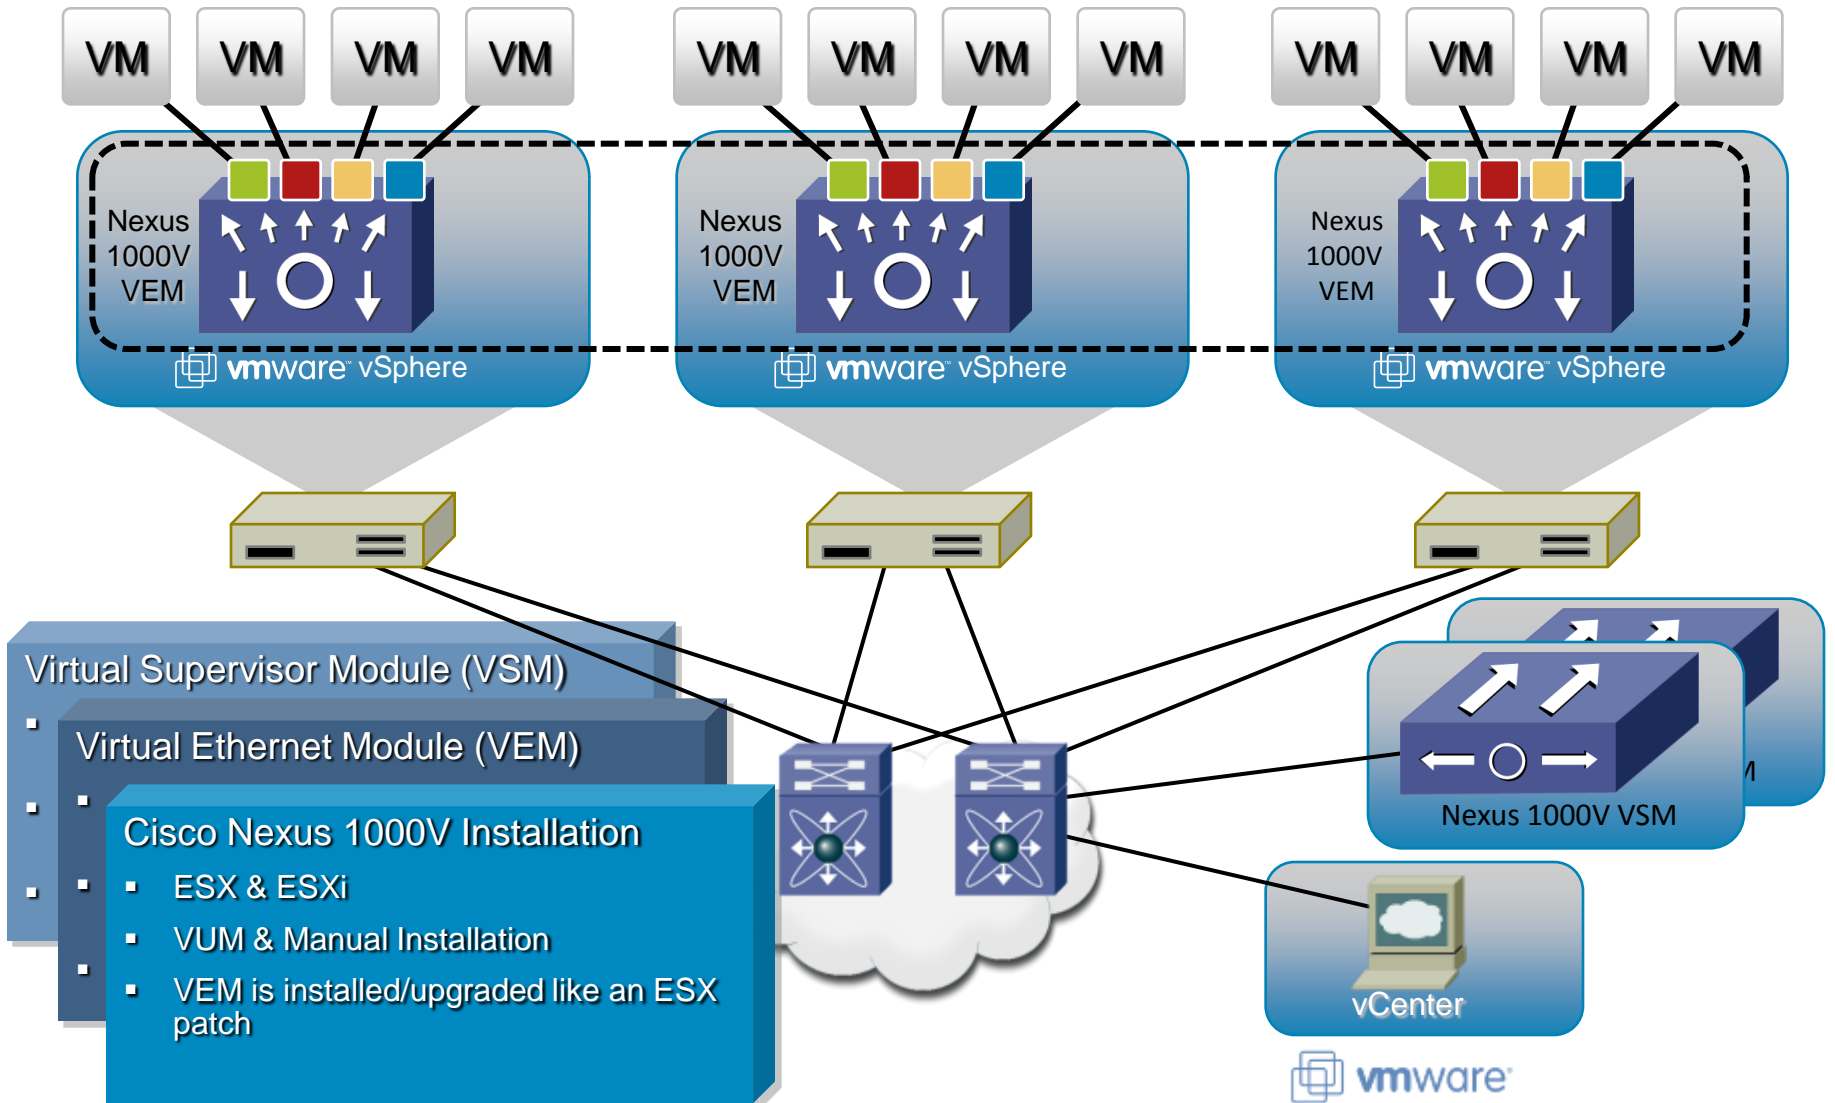
- Nexus 1000V and VN-Link provide visibility to the individual VMs
- Policy can be configured per-VM
- Policy is mobile within the ESX cluster

**Nexus 1000V
Distributed Virtual Switch**

Virtual Access Layer @ Virtualised Servers




Cisco Nexus 1000V Architecture



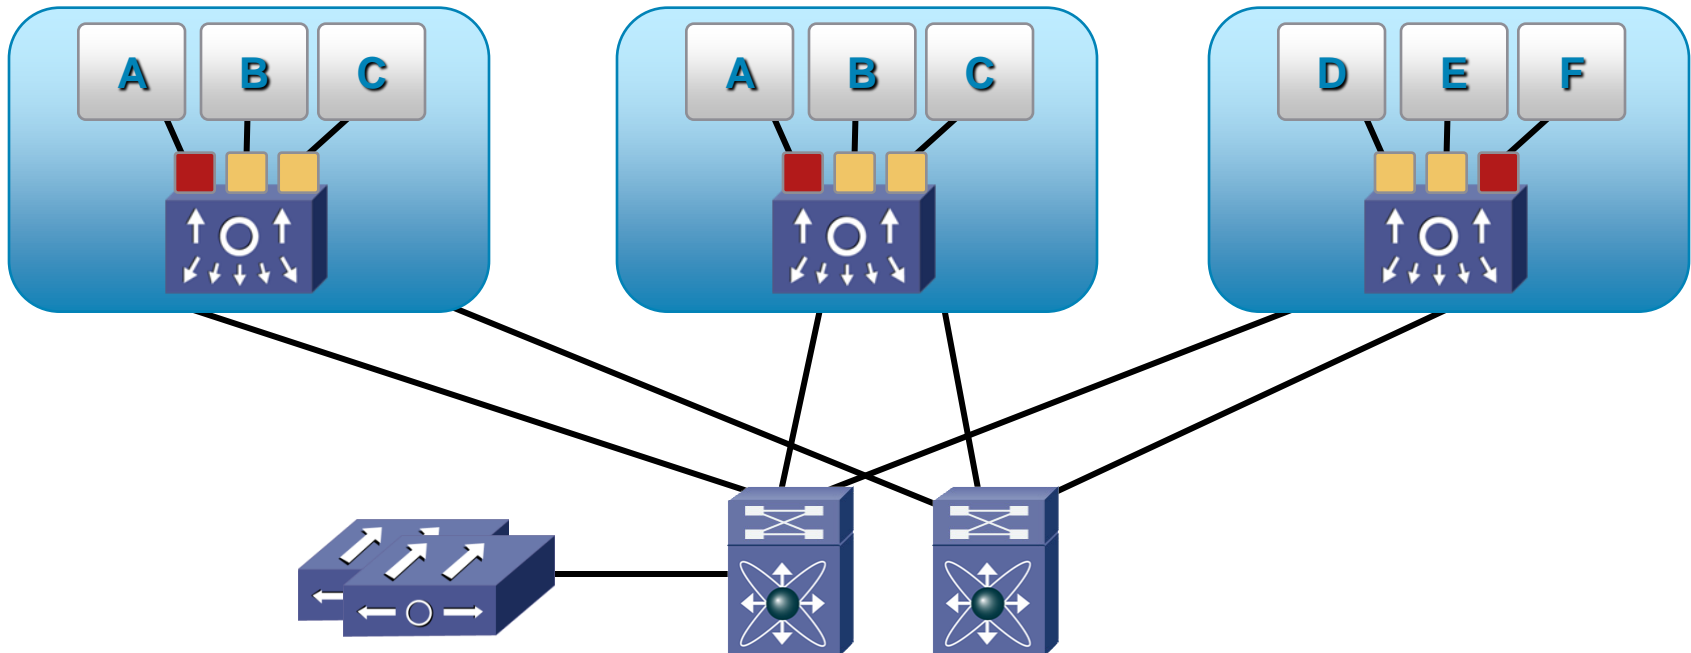
Single Control and Management Plane

- Even if the Nexus 1000V is a distributed switch. It looks like a single switch from control plane and management plane perspective
- Protocol like CDP, Netflow, SNMP are manage from one location the VSM




Distributed Data Plane

- The Virtual Ethernet Module (VEM) is in the Data path
- The Virtual Supervisor Module is only doing control plane and management function
- Each Virtual Ethernet Module forwards packets independent of each other

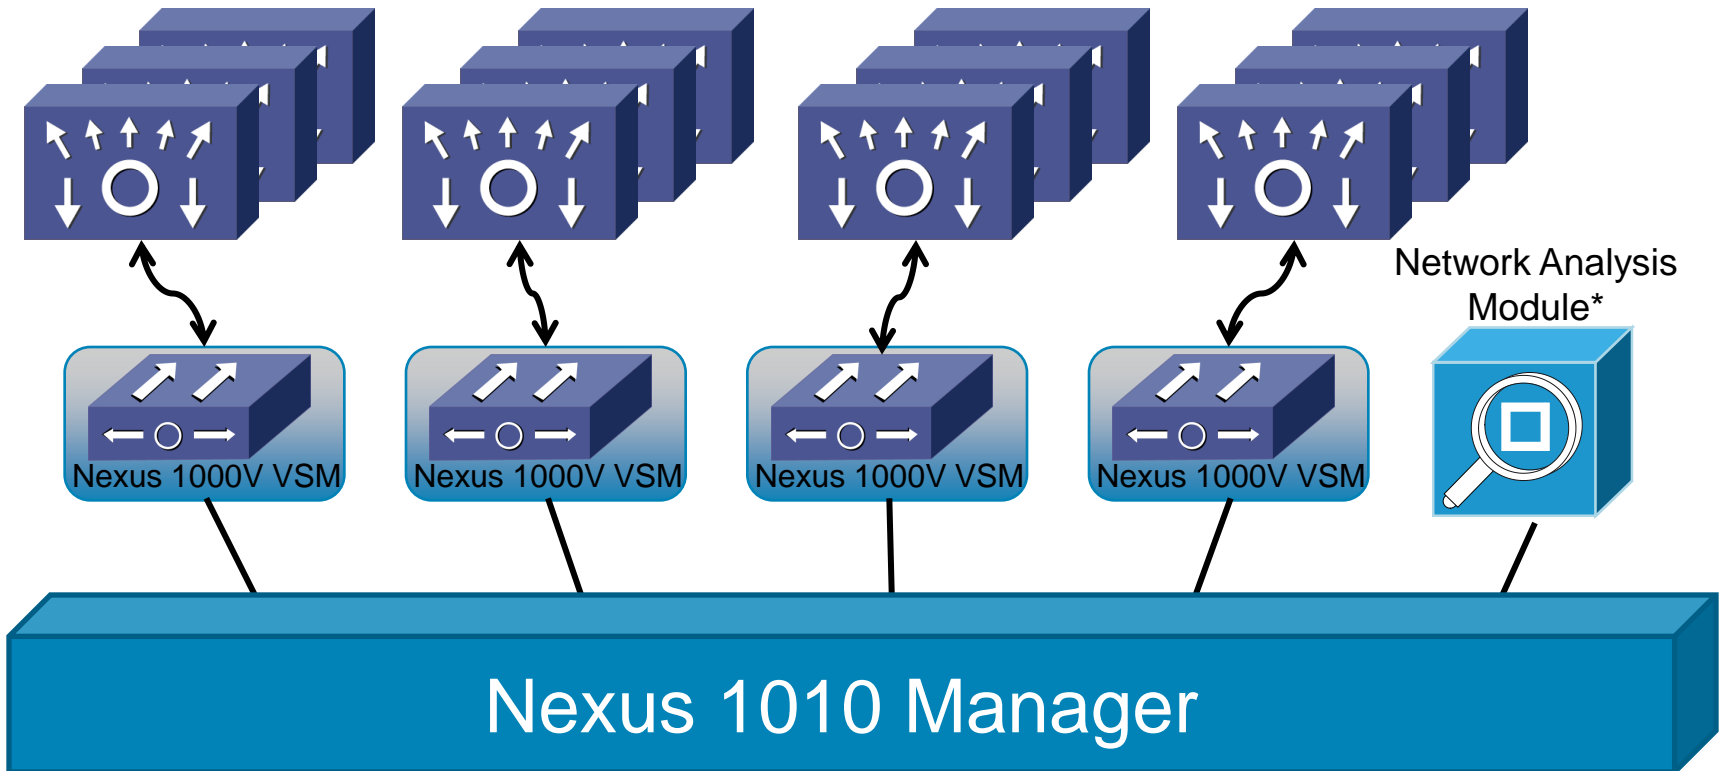


Port Profiles “how to”

1. Nexus 1000V automatically enables port groups in vCenter via API
2. Server Admin uses vCenter to assign vnic policy from available port groups
3. Nexus 1000V automatically enables VM connectivity at VM power-on



Nexus 1010: “Virtual Service Blade” Manager



Nexus 1010 Manager: Cisco management experience


Manages virtual service blades

Virtual Access Layer @ Virtualised Servers

 **Nexus 1000v**

- Gigabit Ethernet
- 10 Gigabit Ethernet
- 10 Gigabit DCE
- 4/8Gb Fiber Channel
- 10 Gigabit FCoE/DCE

Core Layer



Nexus 7000
10GbE Core



SAN A

MDS 9500
Storage



SAN B



vWAAS

Catalyst 6500 or appliances
DC Services




Virtual Service Gateway


Aggregation Layer




Nexus 7000
10GbE Agg




Catalyst 6500
End-of-Row




Nexus 5000 &
Nexus 2000
Top-of-Rack



CBS 31xx
Blade




Nexus 7000
End-of-Row



Nexus 5000 &
FCoE
Top-of-Rack



CBS 31xx
MDS 9124e
Nexus 4000



Cisco UCS

1GbE Server Access

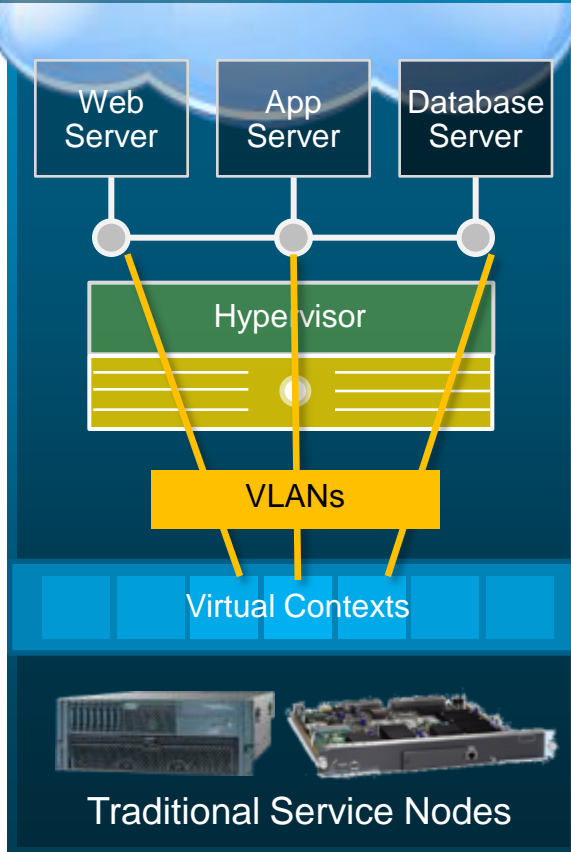
Virtualised Services (VSN)

10GbE and 4/8Gb FC Server Access
10Gb DCE / FCoE Server Access

Deployment options for VSN (Virtual Service Nodes)

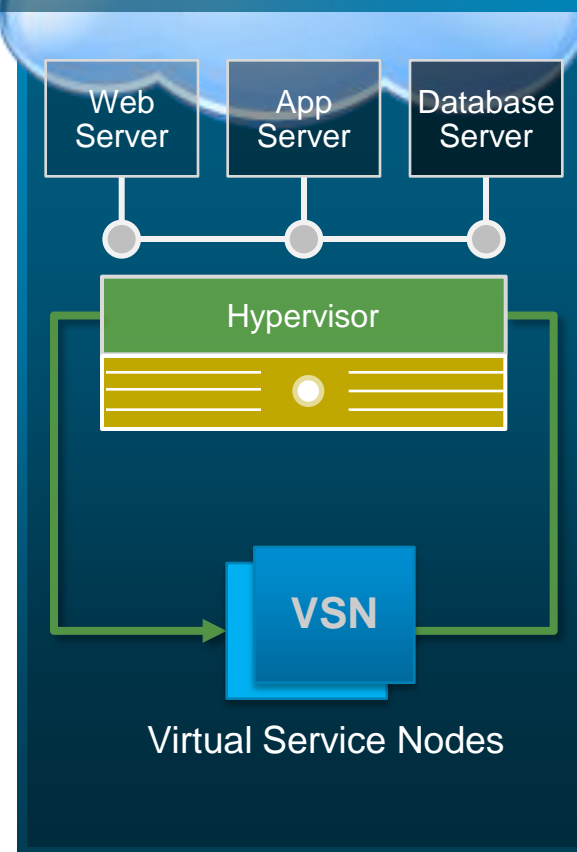
1

Redirect VM traffic via VLANs to external (physical) appliances




2

Apply hypervisor-based network services





vPATH Interception in Nexus 1000v

- vPATH Interception is configured on Server VM's Port Profile in both directions to redirect to a VSN
- Server traffic is intercepted by vPATH interception in VEM and redirected to a VSN
- VSN egress traffic forwarded without further vPATH interception.



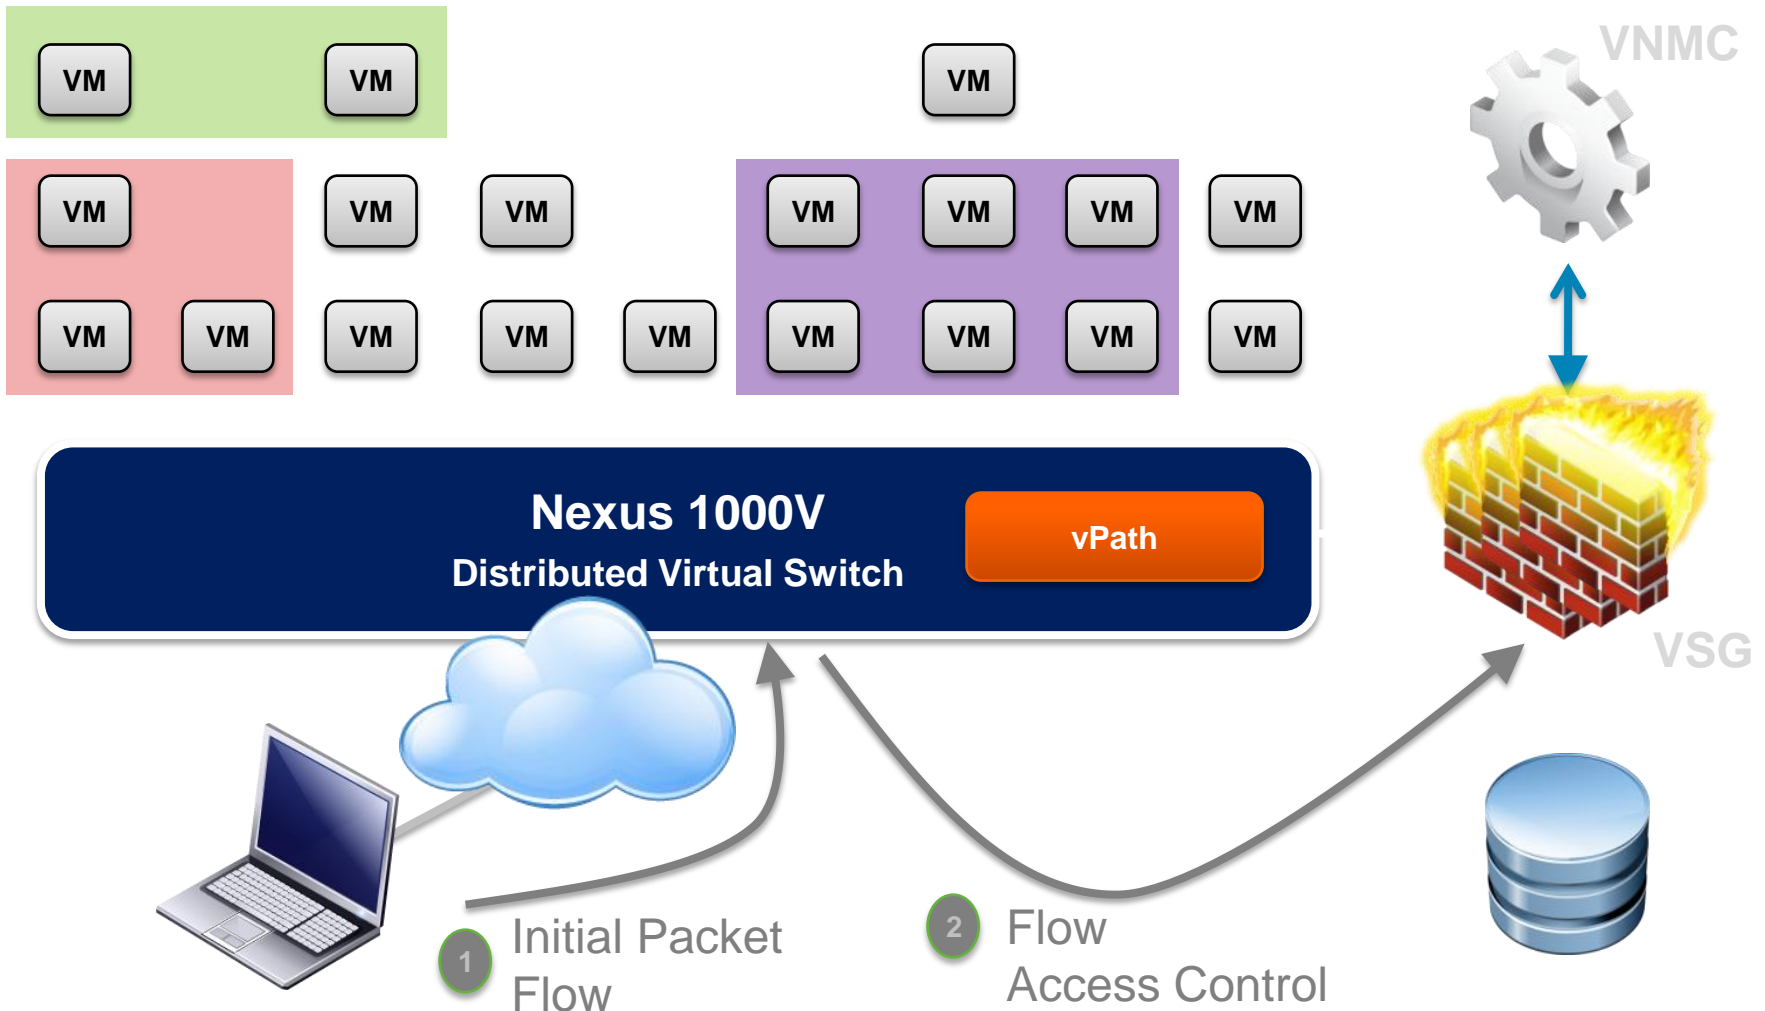
Cisco Virtual Security Gateway (VSG)

Intelligent Traffic Steering with vPath



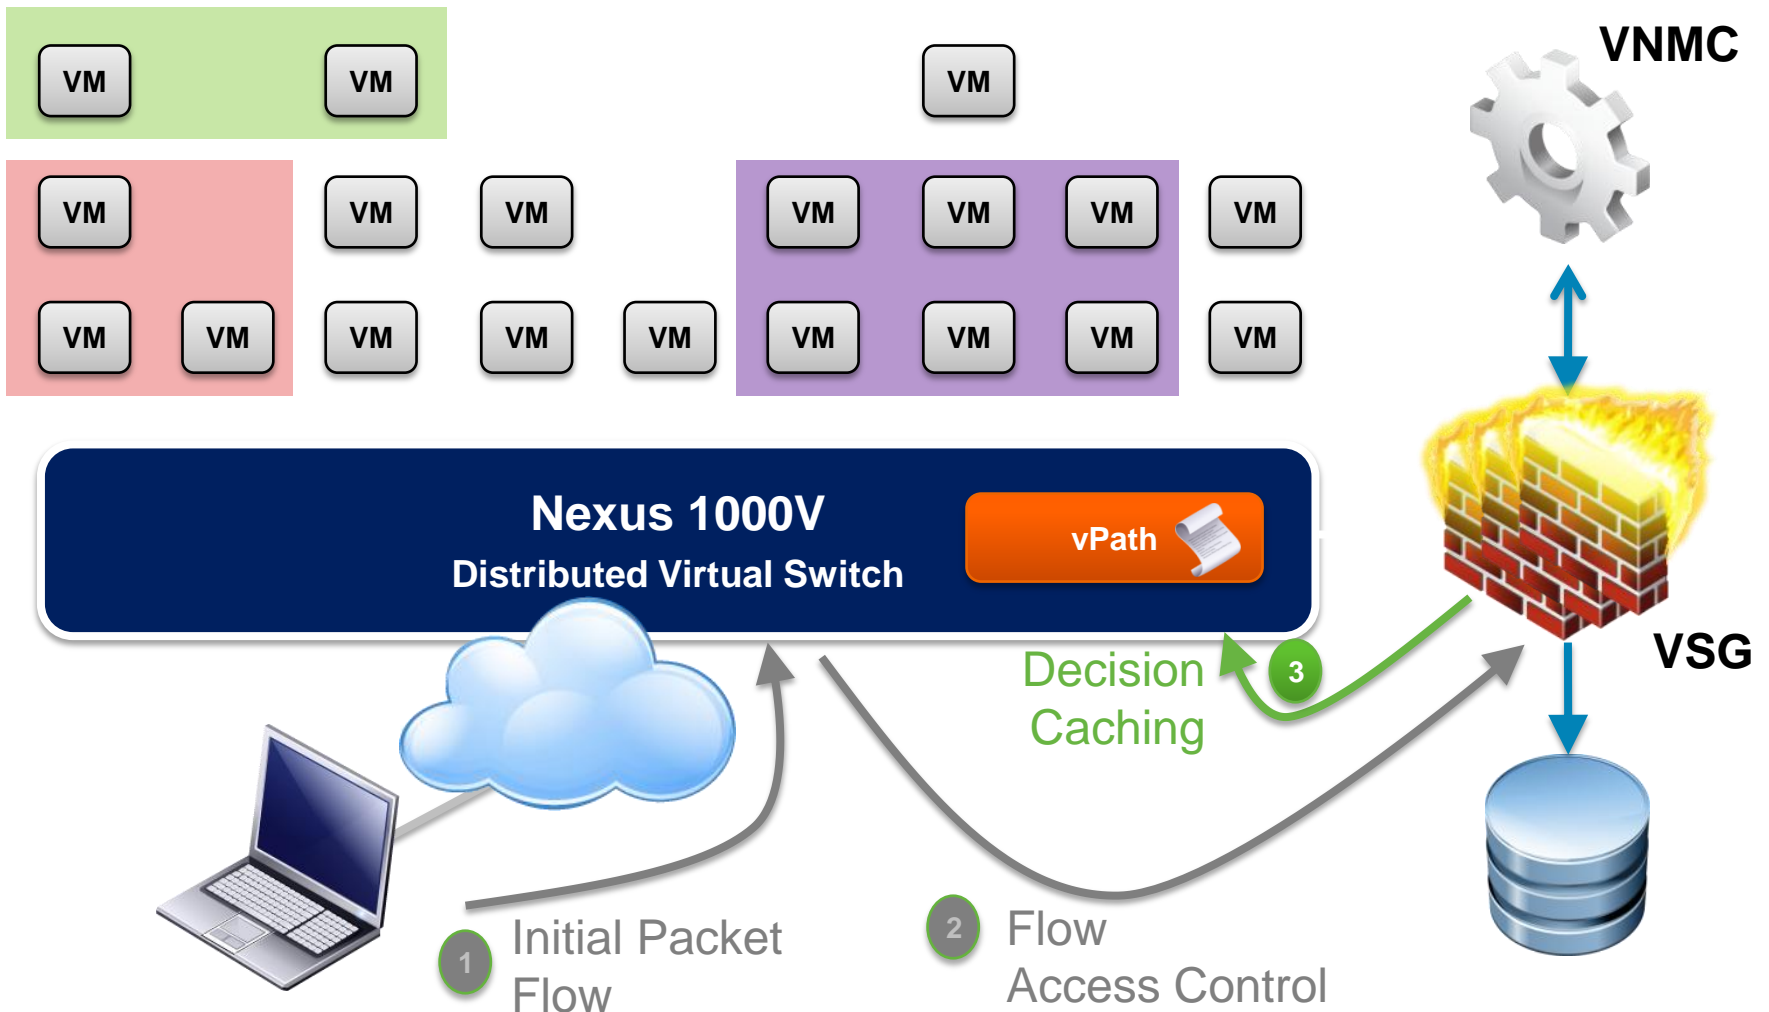
Cisco Virtual Security Gateway (VSG)

Intelligent Traffic Steering with vPath



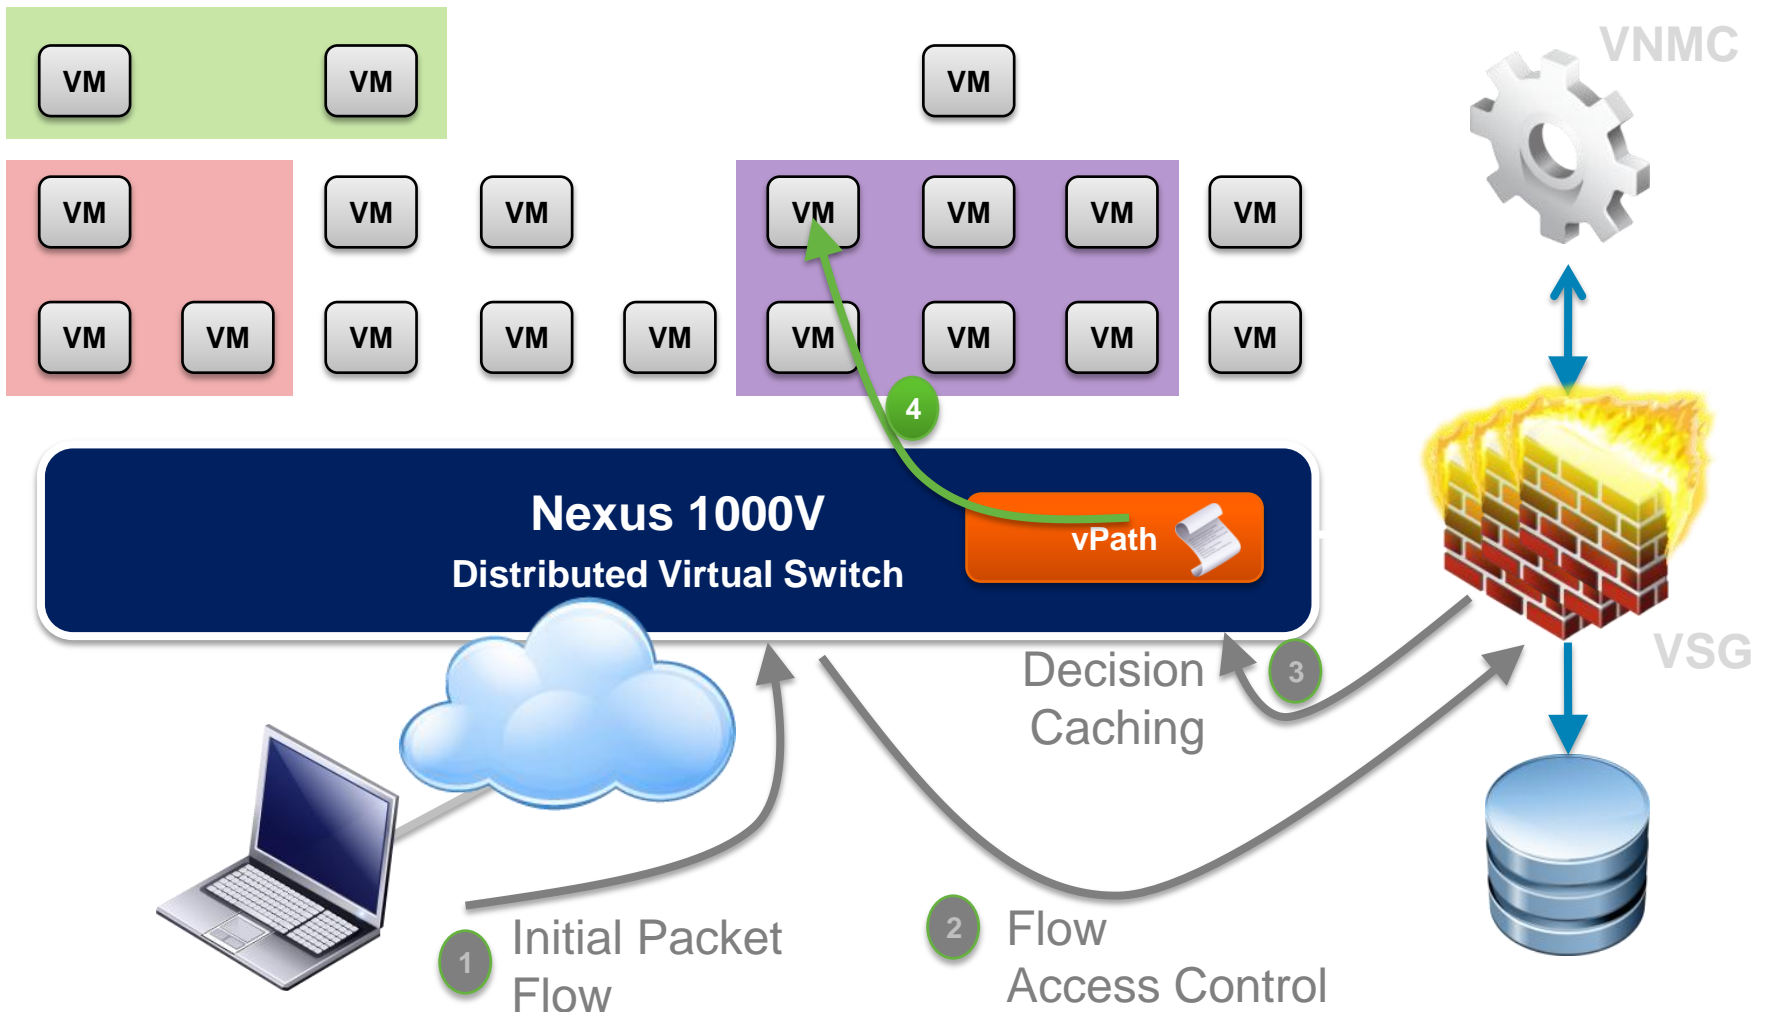
Cisco Virtual Security Gateway (VSG)

Intelligent Traffic Steering with vPath



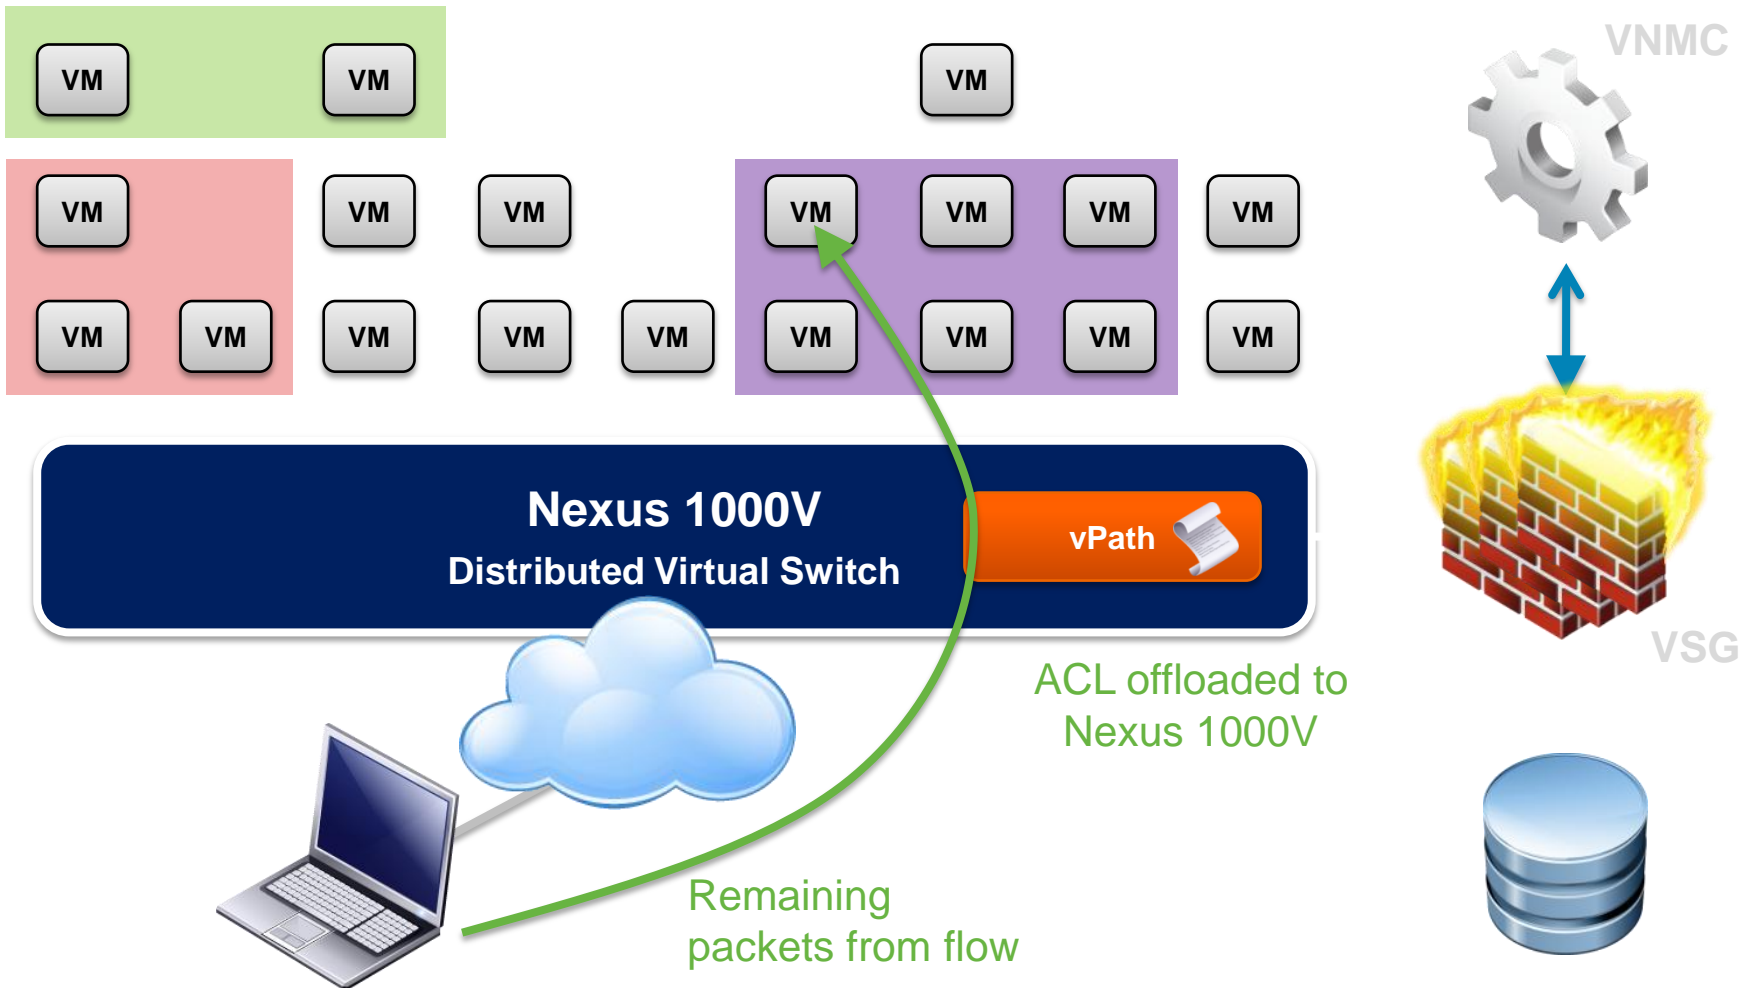
Cisco Virtual Security Gateway (VSG)

Intelligent Traffic Steering with vPath



Cisco Virtual Security Gateway (VSG)

Performance Acceleration with vPath



Cisco Virtual Security Gateway (VSG)

A Better Security Solution

Virtual Security Gateway (VSG)



Context aware Security

VM context aware rules

Zone based Controls

Establish zones of trust

Dynamic, Agile

Policies follow vMotion

Best-in-class Architecture

Efficient, Fast, Scale-out Software

Virtual Network Management Center (VNMC)



Non-Disruptive Operations

Security team manages security

Policy Based Administration


Central mgmt, scalable deployment, multi-tenancy

Designed for Automation

XML API, security profiles

vWAAS vPATH Interception

- Interception based on port-profile policy configured in Nexus 1000v
- **Bidirectional Interception** - (no IN/OUT configuration)
- **Pass-through traffic automatic bypass**




Cisco UCS x86 Server




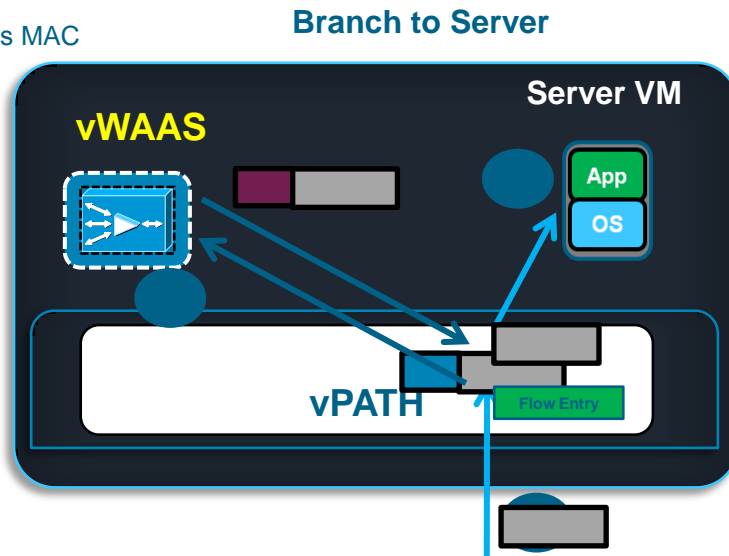

Nexus 1000v VSM



vCenter Server



Life of a Packet (with vPATH Interception)



When vPATH receives first packet from the WAN for optimized port-profile, it will lookup for any flow-entry relevant to the packet;

Should vPATH doesn't find any flow-entry for the packet, it encapsulates the packet with vWAAS MAC. The packet is forwarded to vWAAS though a MAC-in-MAC tunnel.

Upon received the packet from vWAAS, vPATH will create a flow-entry (including reverse flow) for the packet and will forward the packet – via normal layer 2 forwarding mechanisms - to the Server VM;

For pass-through (PT) traffic, vPATH creates PT flow entry and doesn't send subsequent packets to vWAAS after initial TCP SYN/SYN-ACK process.

vPATH – VM mobility awareness



Feature

1. vPATH aware of VM's movement from one host to another.
2. Traffic interception continue to work as-is without any disruption or changes required.

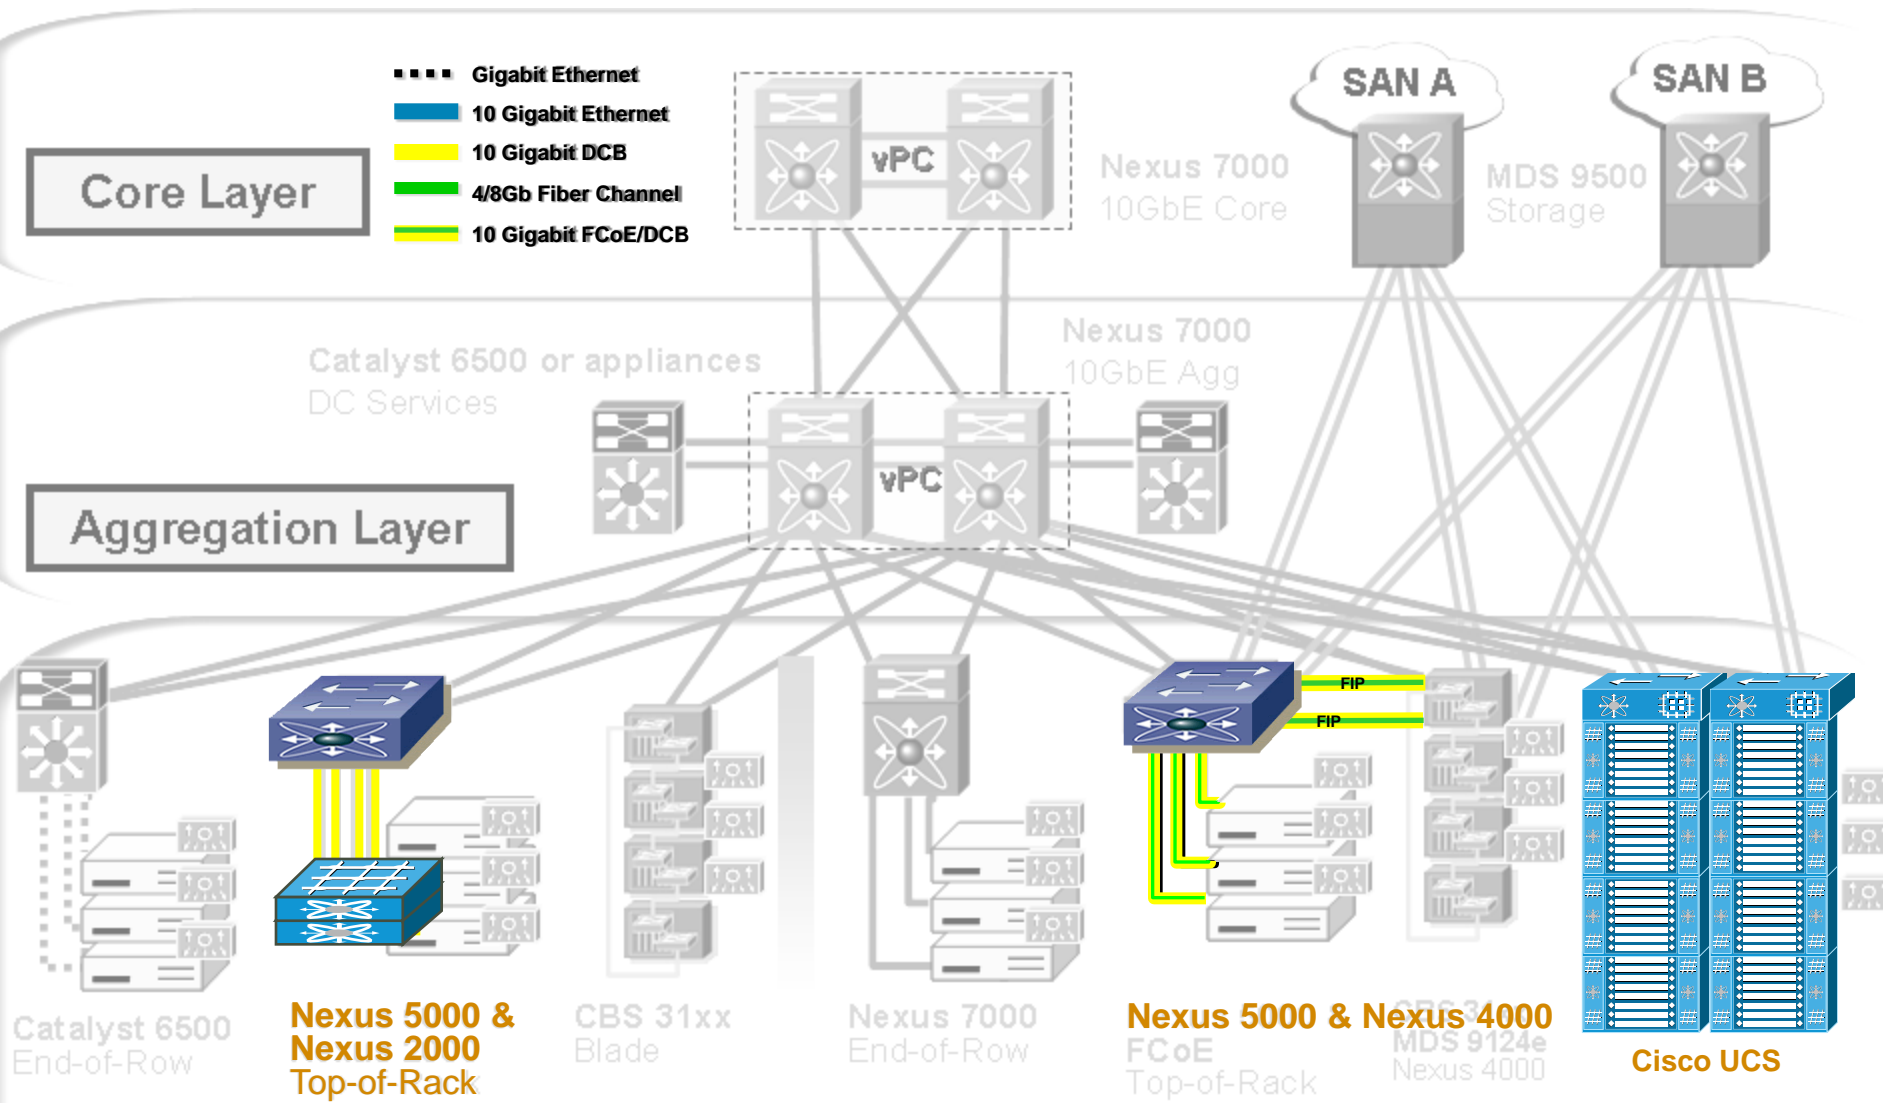
Benefit

No disruption in WAN optimization service if VM moves from one host to another.

Support VMware resources scheduling (DRS) and provides High availability

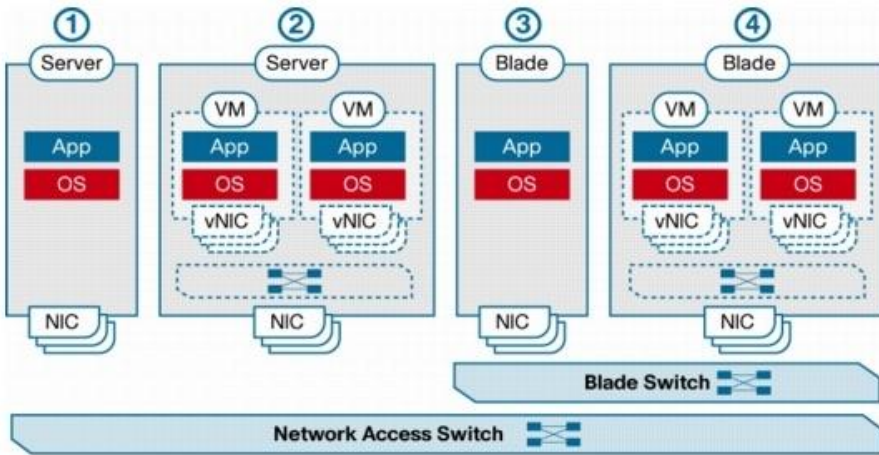


Server IO Virtualisation


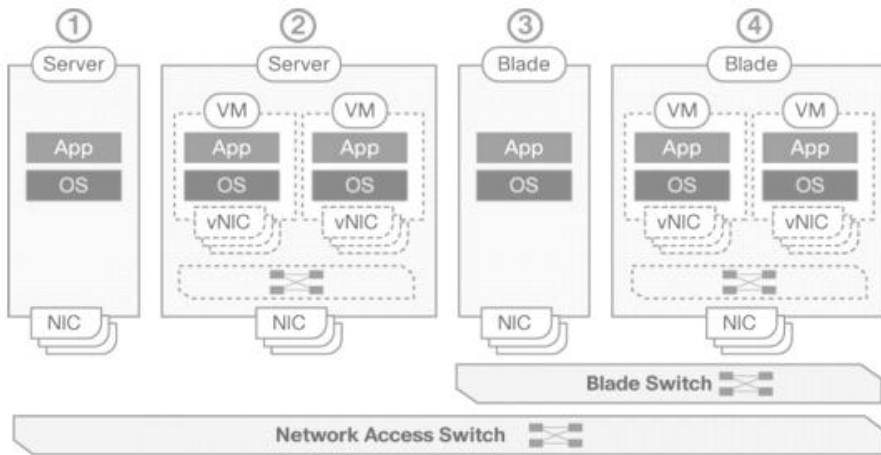


Server IO Virtualisation

Where does IO Virtualisation apply ?



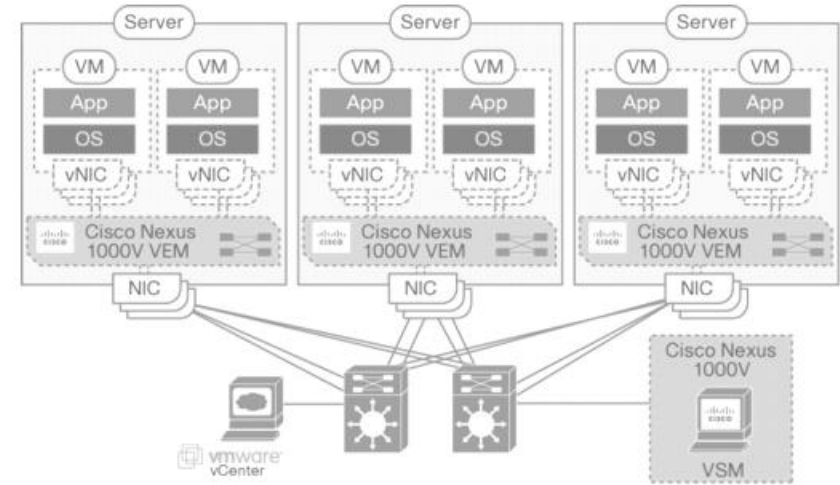
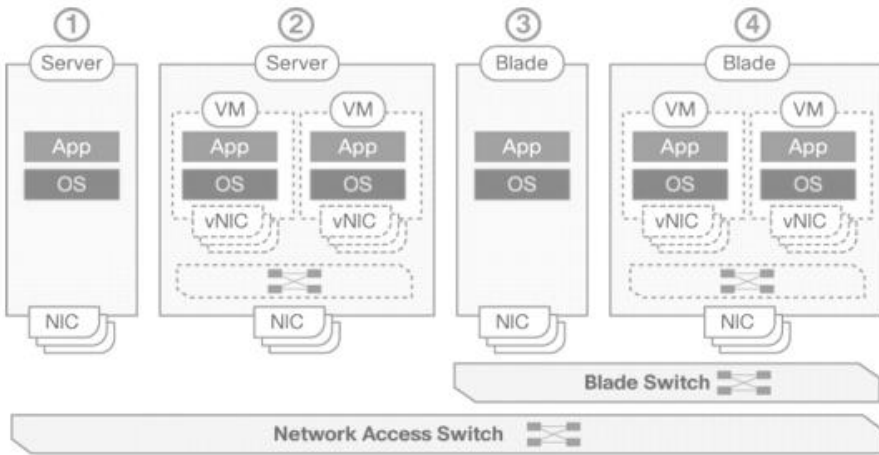
Where does IO Virtualisation apply ?



**OK. We've seen that already.
What about doing so in hardware ?**



Where does IO Virtualisation apply ?



SR-IOV @ PCIe



10GE / CNA



Intel VT-x



Cisco Nexus
5000

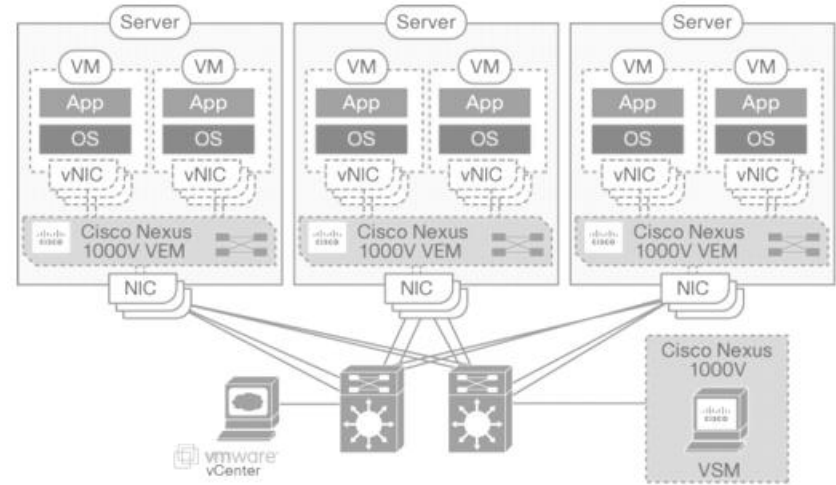



Cisco UCS




VN-Link & NIV

Where does IO Virtualisation apply ?



SR-IOV @ PCIe




10GE / CNA



Intel VT-x



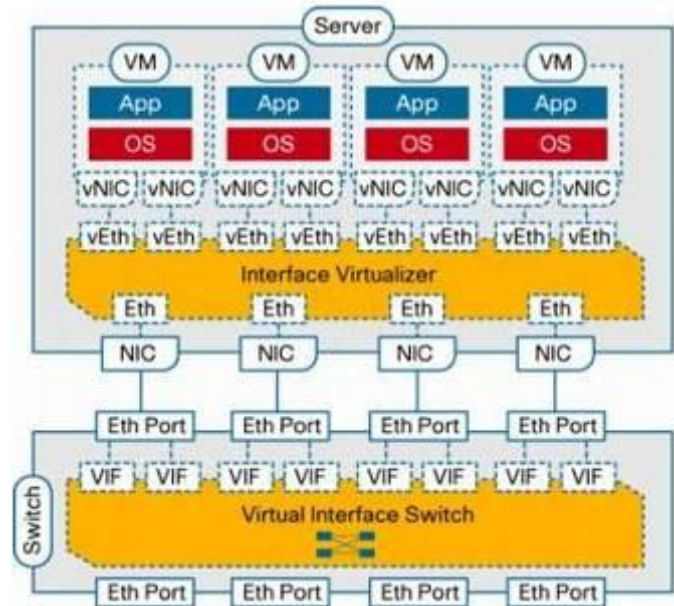
Cisco Nexus 5000



Cisco UCS

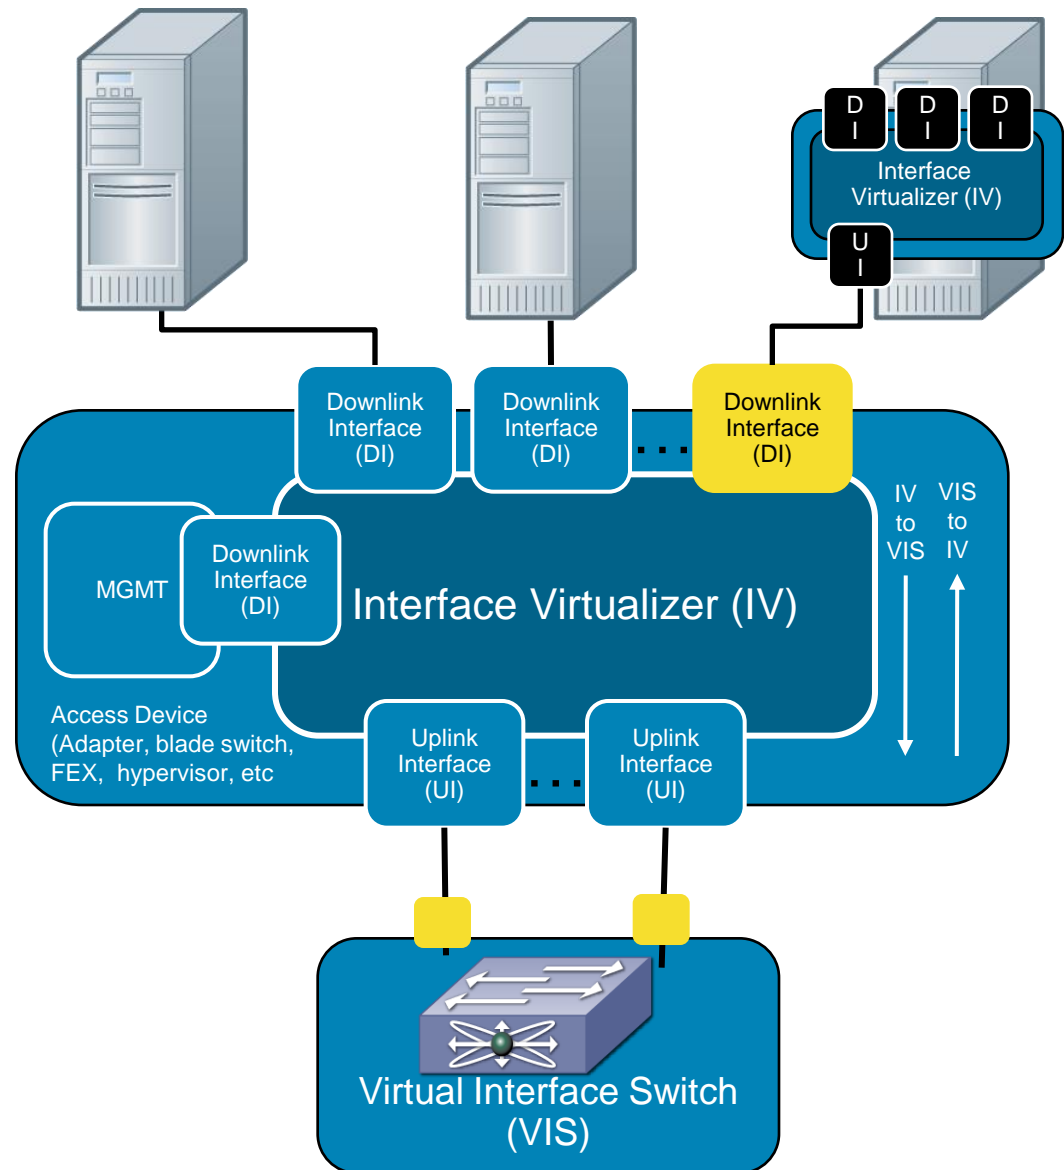


VN-Link & NIV



Network Interface Virtualisation (NIV) Functional Model

- NIV applies to both physical and virtual ports
- Can be cascaded through multiple layers
- Introduces new L2 tag & forwarding model that encompasses both unicast & multicast distribution.
- All frames flow from IV through VIS allowing consistent forwarding & policy enforcement for all traffic



Cisco UCS M81KR VIC Overview

Mezzanine Card for B-Series

Converged Network Adapter (CNA) designed for both single-OS and VM-based deployments

- Virtualised in Hardware
- PCIe compliant

High Performance


- 2x 10Gb
- 500K+ IOPS

The OS/Hypervisor sees up to ~58 distinct PCIe devices

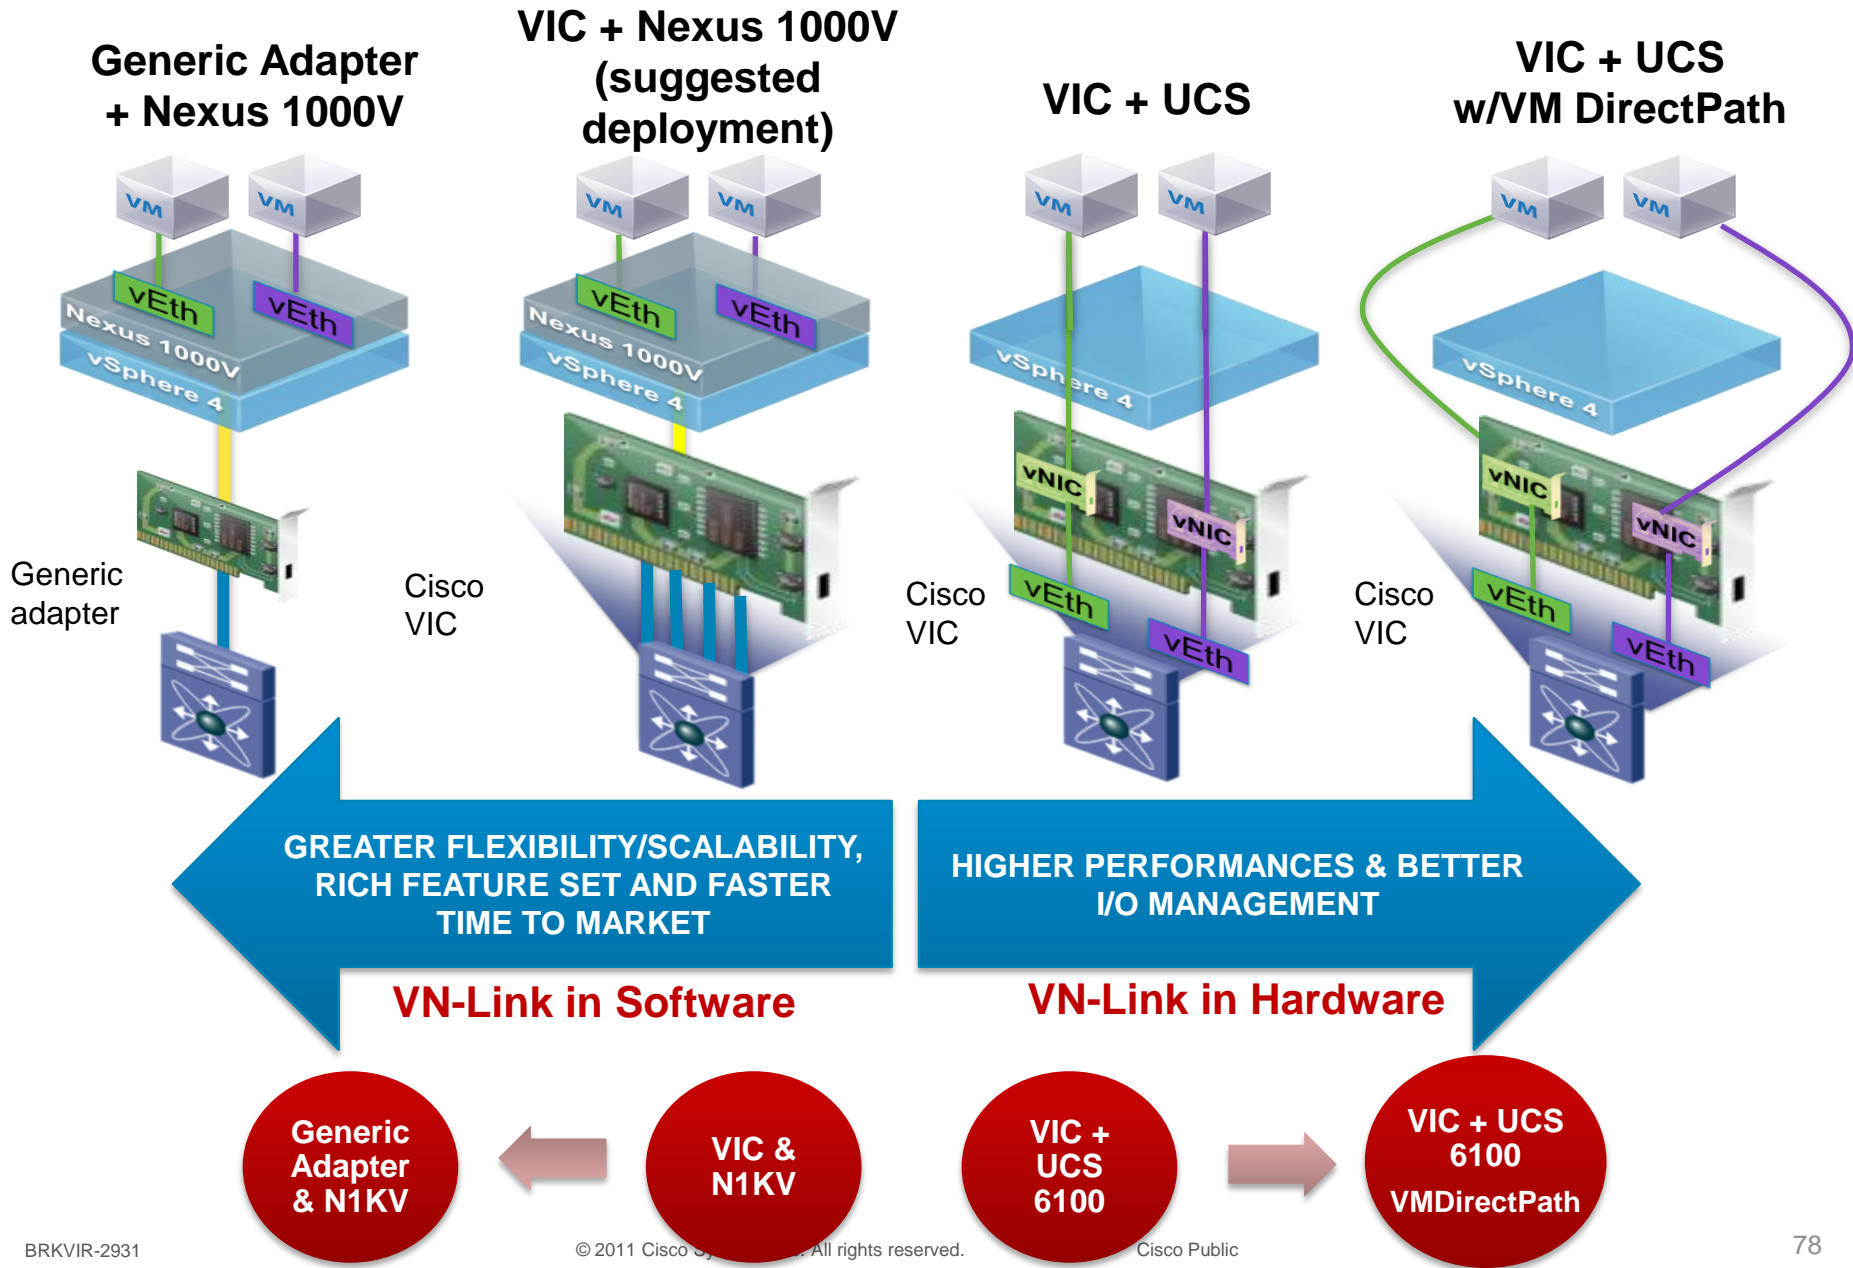
- Ethernet vNIC and FC vHBA
- Management from the network

VN-Link in Hardware – Ideal for Virtualisation Environments

- Bypass vSwitch to deliver VN-Link in hardware
- Tight integration with VMware vCenter

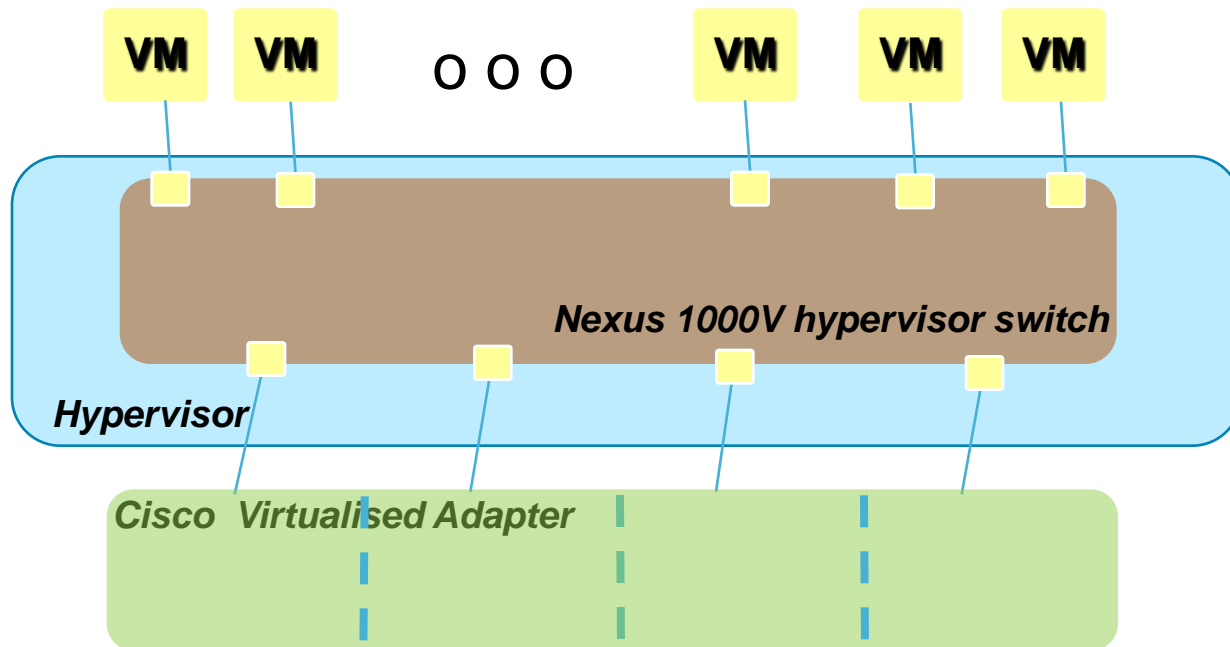


Summary of VN-Link Offerings



Optimize IO for Virtualised Environments

Scenario 2: VN-LINK in Software with VIC



VN-LINK in SW = Nexus 1000V

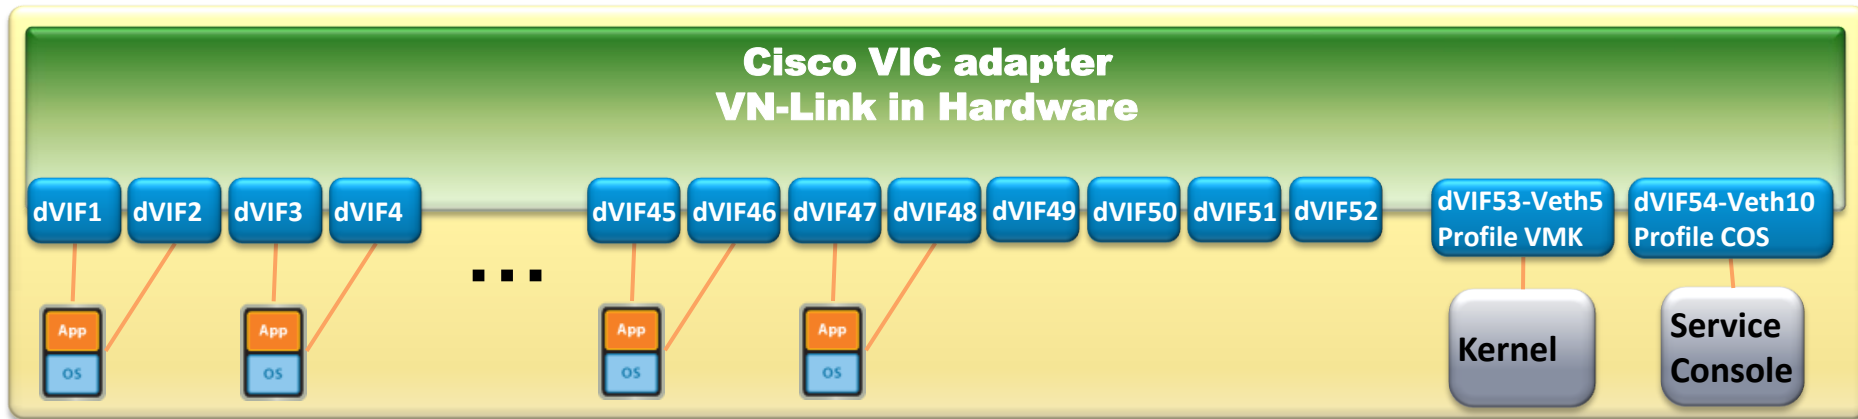
- Each VM vnic connects to Nexus 1000V hypervisor switch
- Nexus 1000V switch uplinks connect to multiple distinct Cisco virtual interfaces (VIFs)

Likely Use Case:

- Customer has already standardized on Nexus 1000V
- Customer deployment needs higher scalability and number of VMs

Optimize IO for Virtualised Environments

Scenario 3: VN-LINK in Hardware



VN-LINK in HW = Pass Thru Switching (PTS)

- Each VM vnic maps to a different virtual interface (VIF)
- IO to/from VM enters Cisco hypervisor switch module... it passes thru to Cisco VIF (switching not done on CPU)

Likely Use Case:

- Customer benefits from centralized Management through UCSM
- Customer needs higher performance

VN-Link in Hardware: UCS Service Profiles with VIC adapter

- Makes use of vCon first if separation by adapter is required
- UCS Administrator Sets the Dynamic vNIC Connection Policy
- 1 static vNIC to fabric A and 1 static vNIC to fabric B, for redundancy
- Actual VLAN, QoS, Pinning, etc. configuration done in port profiles

The screenshot displays the UCS Administrator interface for configuring a Service Profile. The left pane shows a tree view of the configuration hierarchy, with 'Service Profile PTS_PALO_1' selected. The right pane shows the configuration details for this profile, including the 'Network' tab.

Dynamic vNIC Connection Policy

Global vNIC Connection Policy

Name: **IOM_to_FI_4_Link**
 Number of Dynamic vNICs: **54**
 Adapter Policy: **VMWarePassThru**

vNIC/vHBA Placement Policy

Specific vNIC/vHBA Placement Policy

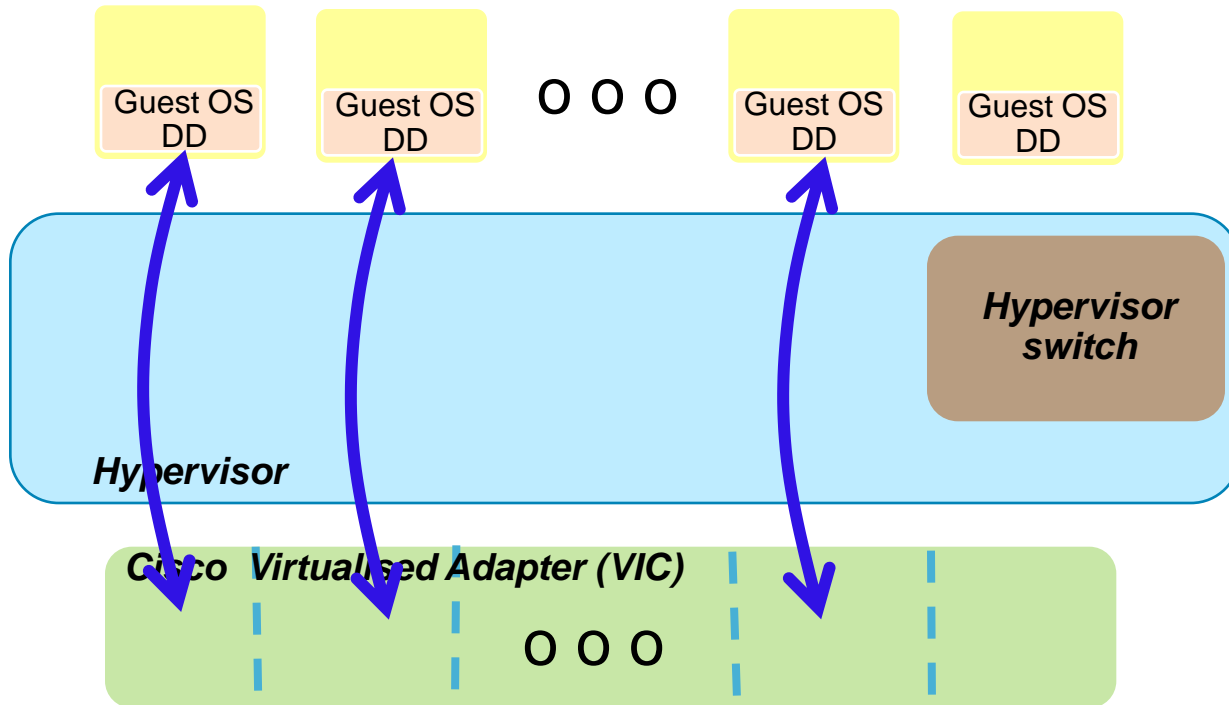
Virtual Slot	Selection Preference
1	all
2	all

vNICs

Name	MAC Address	Order	Fabric ID	vCon	Native VLAN
vNIC eth0	00:25:B5:00:01:FF	1	A	any	
vNIC eth1	00:25:B5:00:00:FF	2	B	any	
vNIC dynamic-prot-001	derived	3	A-B	any	
vNIC dynamic-prot-002	derived	4	B-A	any	
vNIC dynamic-prot-003	derived	5	A-B	any	
vNIC dynamic-prot-004	derived	6	B-A	any	

Optimize IO for Virtualised Environments

Scenario 4: VN-LINK in HW with VMDirectPath



VN-LINK in HW with VMDirectPath

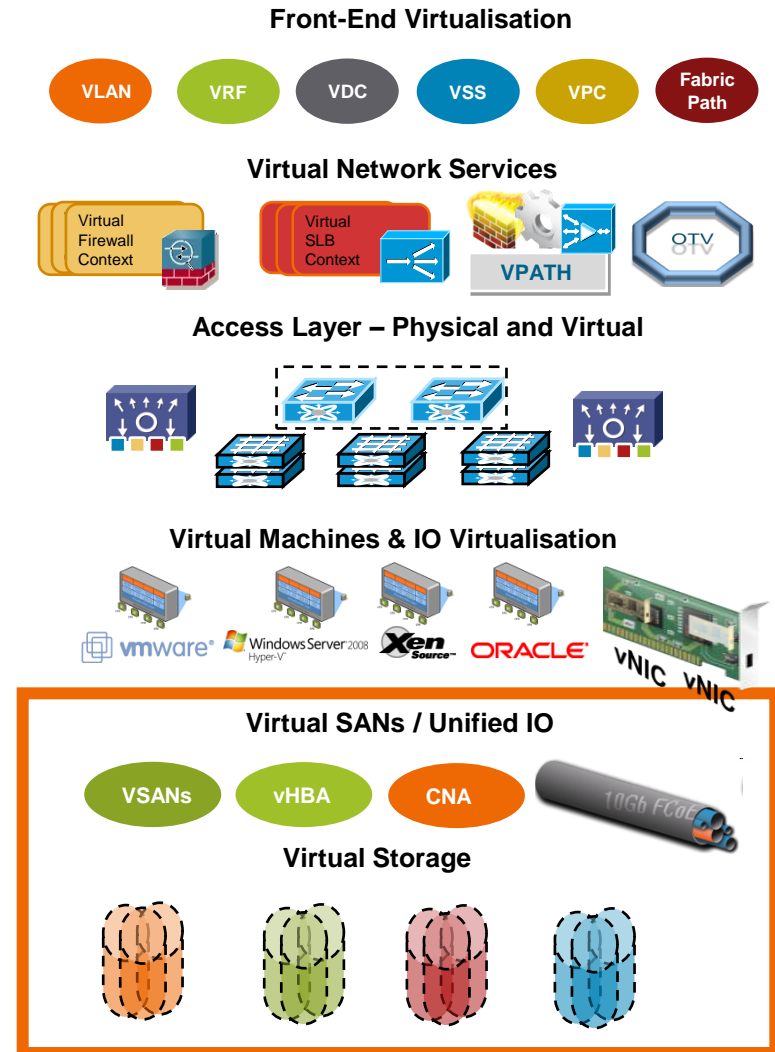
- Bypass of Hypervisor completely... Guest OS Device Driver talks directly to Cisco Virtualised adapter
- Much higher Performance (native HW performance)

Likely Use Case:

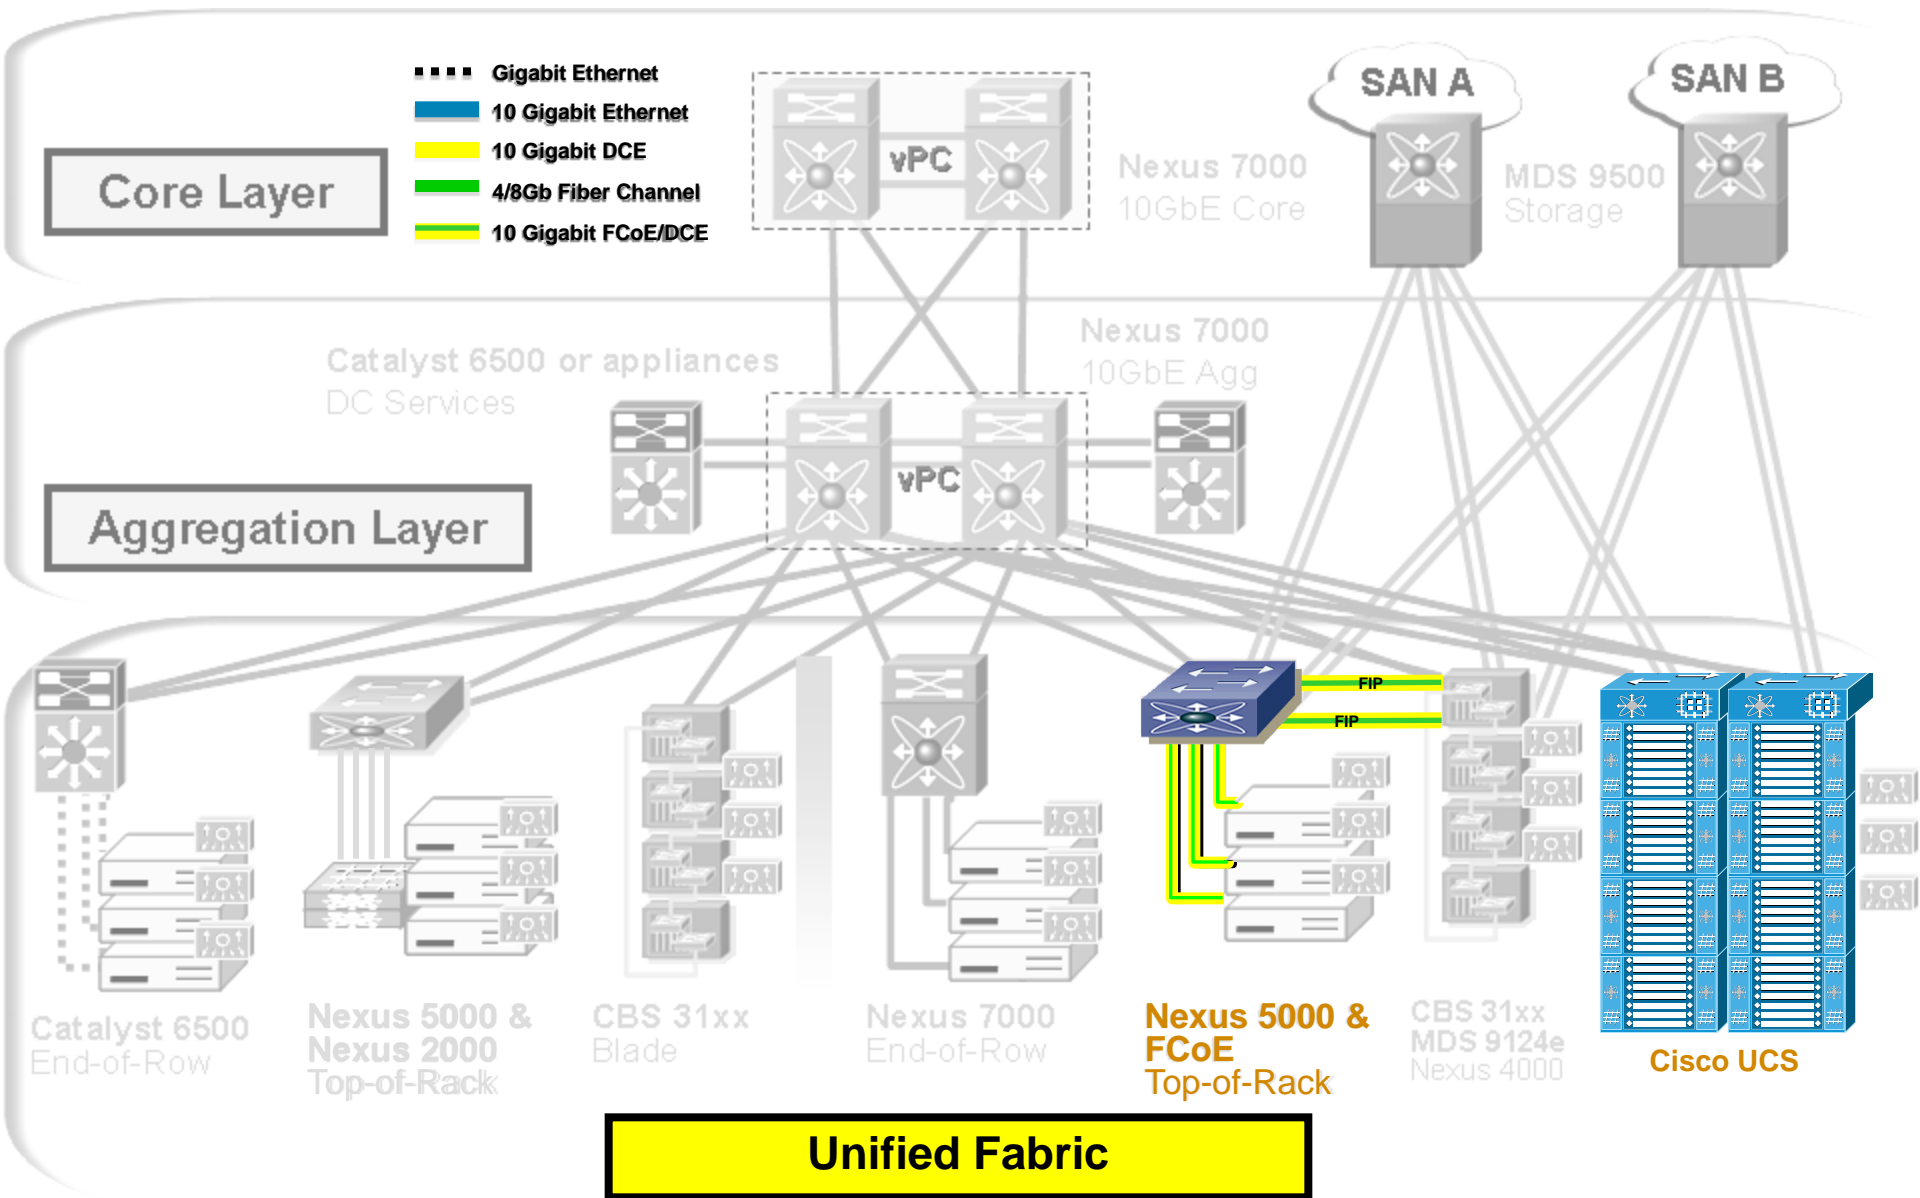
- High Performance Workloads (e.g. Appliances)
- Workloads that don't care about Vmotion

Agenda

- Data Centre Virtualisation Overview
- Front-End Data Centre Virtualisation
 - Core Layer
 - Aggregation Layer
 - Networking Services
 - Access Layer
- Server Virtualisation
 - Hypervisors
 - Virtual Access Layer
 - Virtualised Services
 - Server IO Virtualisation
- **Back-End Virtualisation**
 - Virtual HBA & NPV**
 - Unified IO & FCoE**
 - SAN & Storage**
- Implementation Examples
 - vBlock
 - Secure Cloud
- Q&A



Unified Fabric & IO



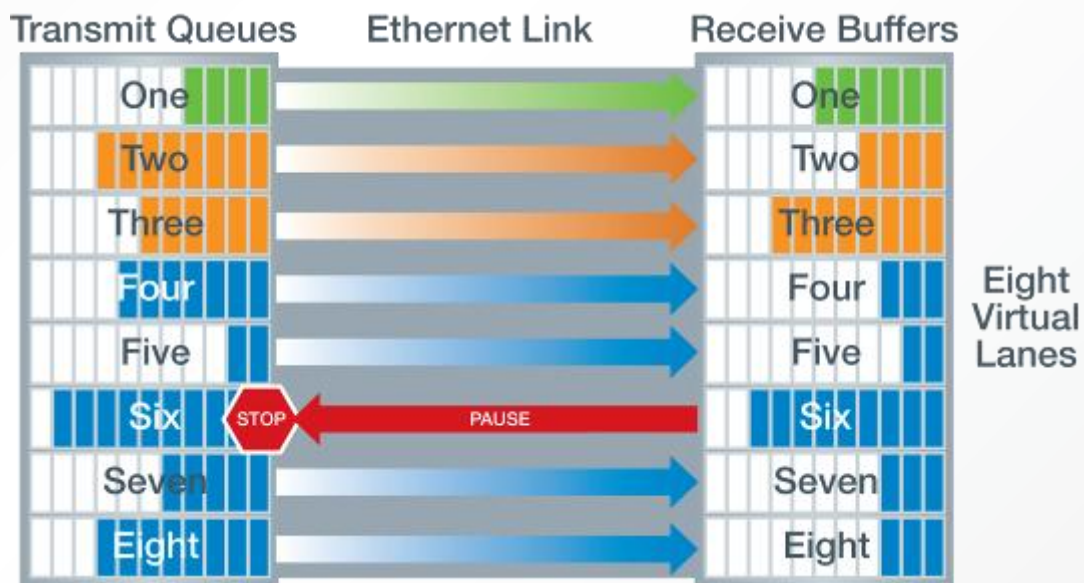
Network Behavior & Characteristics

- Ethernet is non-deterministic.
 - Flow control is destination-based
 - Relies on TCP drop-retransmission / sliding window
- Fibre-Channel is deterministic.
 - Flow control is source-based (B2B credits)
 - Services are fabric integrated (no loop concept)




Data Centre Bridging Features - PFC

Priority-based Flow Control (PFC)



- Enables lossless Fabrics for each class of service
- PAUSE sent per virtual lane when buffers limit exceeded
- Network resources are partitioned between VL's (E.g. input buffer and output queue)
- The switch behavior is negotiable per VL

DCB / FCoE Related Standards Status at a Glance



Standard		Investigation	Development	Approval	Publication
DCB	PFC	Completed	Completed	Completed	Awaiting
	ETS	Completed	Completed	Completing Approval Phase 3	
	DCBX	Completed	Completed	Completing Approval Phase 3	
FCoE	FC-BB-5	Completed	Completed	Completed	Completed

Technical Stability


Standard / Feature
IEEE 802.1Qbb - Priority-based Flow Control (PFC)
IEEE 802.3bd - Frame Format for PFC
IEEE 802.1Qaz Enhanced Transmission Selection (ETS) and Data Centre Bridging eXchange (DCBX)

CEE (Converged Enhanced Ethernet) is an informal group of companies that submitted initial inputs to the IEEE-DCB Working Group.

FC over Ethernet (FCoE)

FCoE

- Mapping of FC frames over Ethernet
- Enables FC to run on a lossless Data Centre Ethernet network





Benefits




- Wire Server Once
- Fewer cables and adapters
- Software Provisioning of I/O
- Interoperates with existing SANs
- No gateway—stateless
- Standard – June 3, 2009

Why FCoE ?

- Fewer CNAs (Converged Network adapters) instead of NICs, HBAs and HCAs
- Limited number of interfaces for Blade Servers



Converged Network Adapters (CNA)



Converged Network Adapter

Off-the-shelf NIC and HBA ASICs from QLogic and Emulex, providing Dual Active-Active 10GE and FCoE connection



Support for Native Drivers and Stacks to minimize disruption for existing environments

Replaces multiple adapters per server, consolidating 10GE and Fibre Channel on a single interface

Supports Priority Flow Control and DCBX

AVAILABLE TODAY!!

Unified I/O Architecture Consolidation



— Ethernet — FC — FCoE

Embedded FCoE at Cisco UCS



From ad hoc and inconsistent...

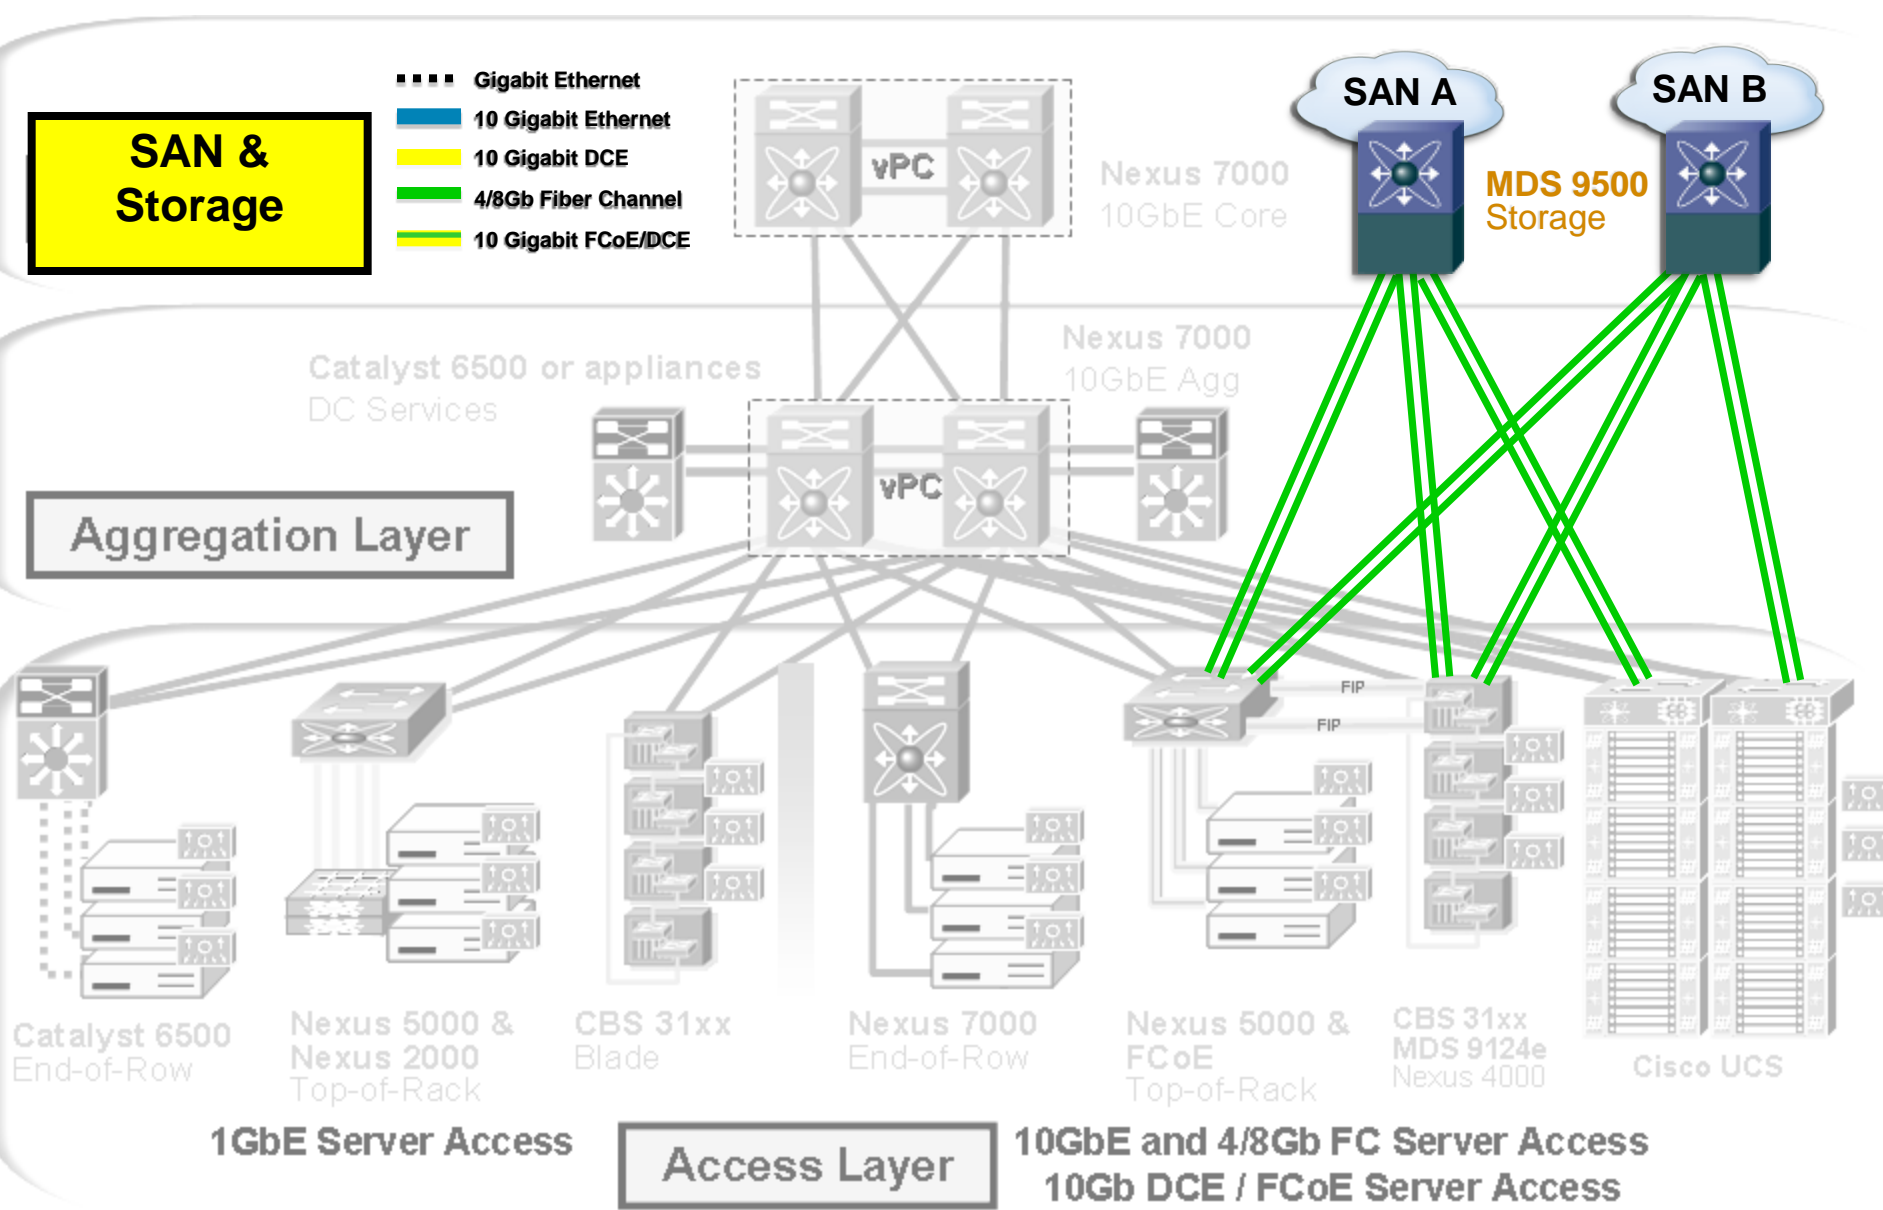


...to structured, but siloed, complicated and costly...




...to simple, optimized and automated

VSAN, NPIV, NPV and Storage Access

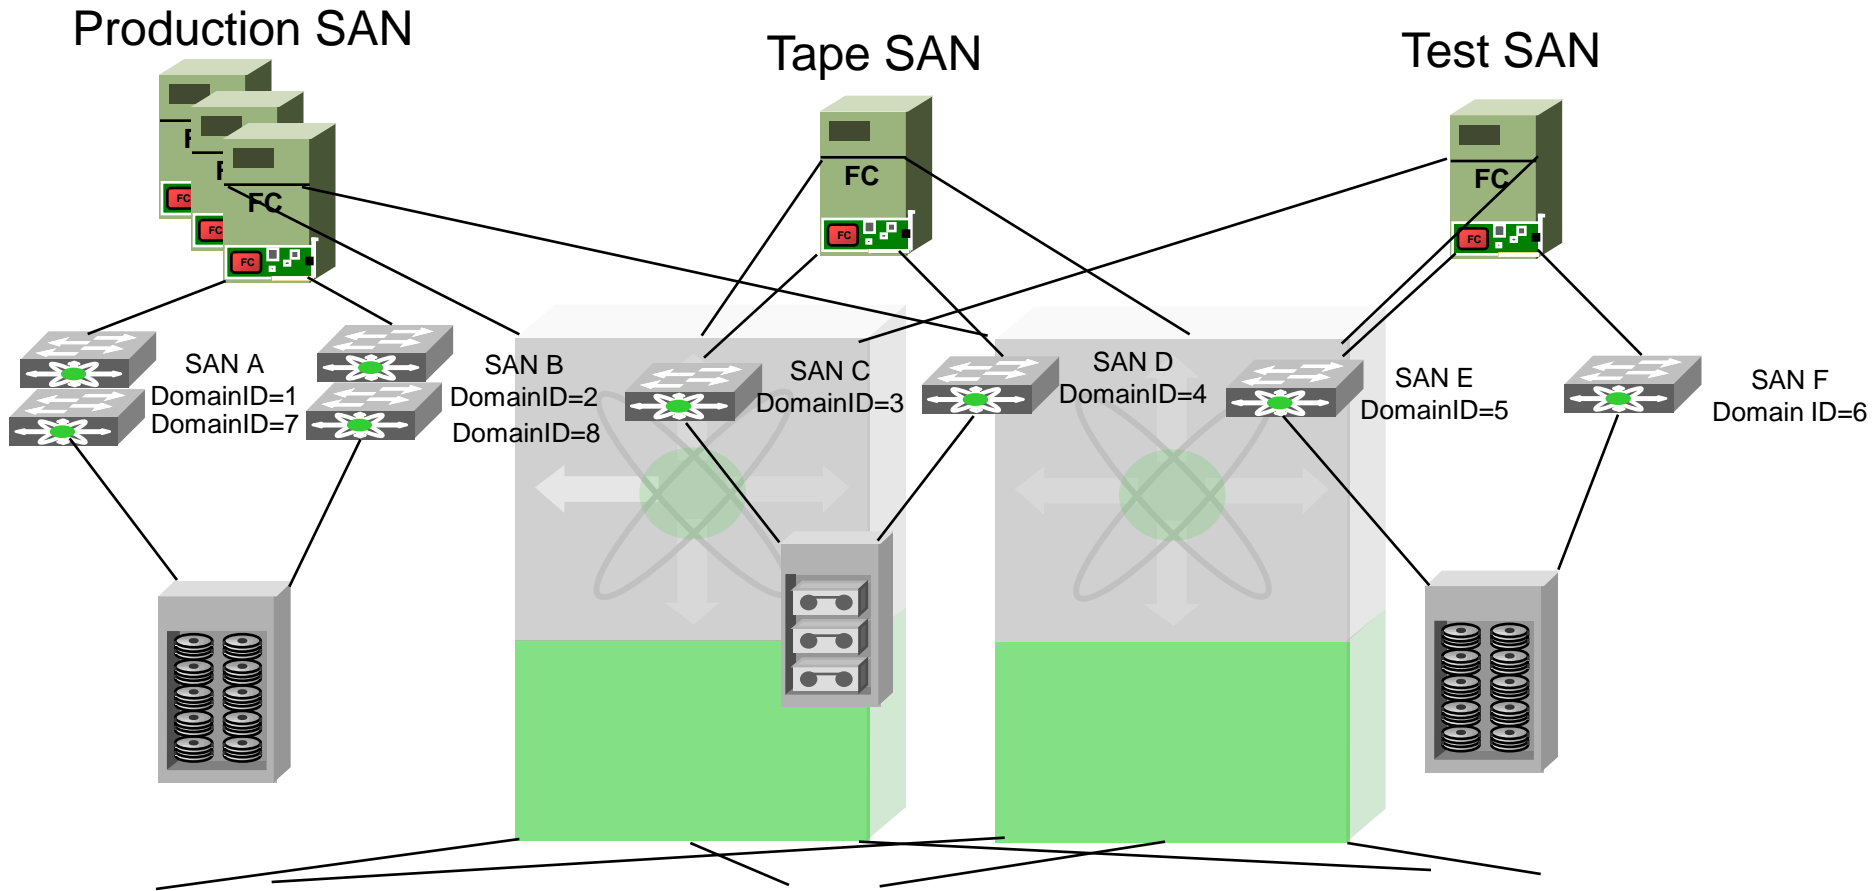


Virtual Storage Area Network (VSAN)

- Consolidation of SAN islands
 - Increased utilization of fabric ports with Just-In-Time provisioning
- Deployment of large fabrics
 - Dividing a large fabric in smaller VSANs
 - Disruptive events isolated per VSAN
 - RBAC for administrative tasks
 - Zoning is independent per VSAN
- Advanced traffic management
 - Defining the paths for each VSAN
 - VSANs may share the same EISL
 - Cost effective on WAN links
- Resilient SAN Extension
- Standard solution (ANSI T11 FC-FS-2 section 10)

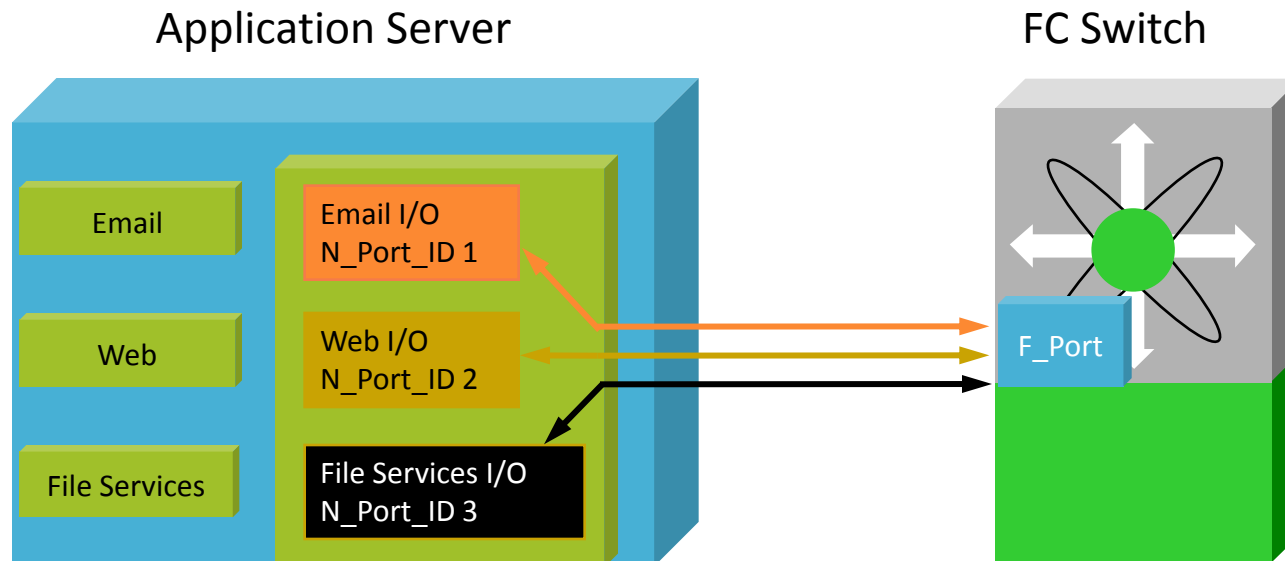


Understanding VSANs (or Virtual Fabrics)



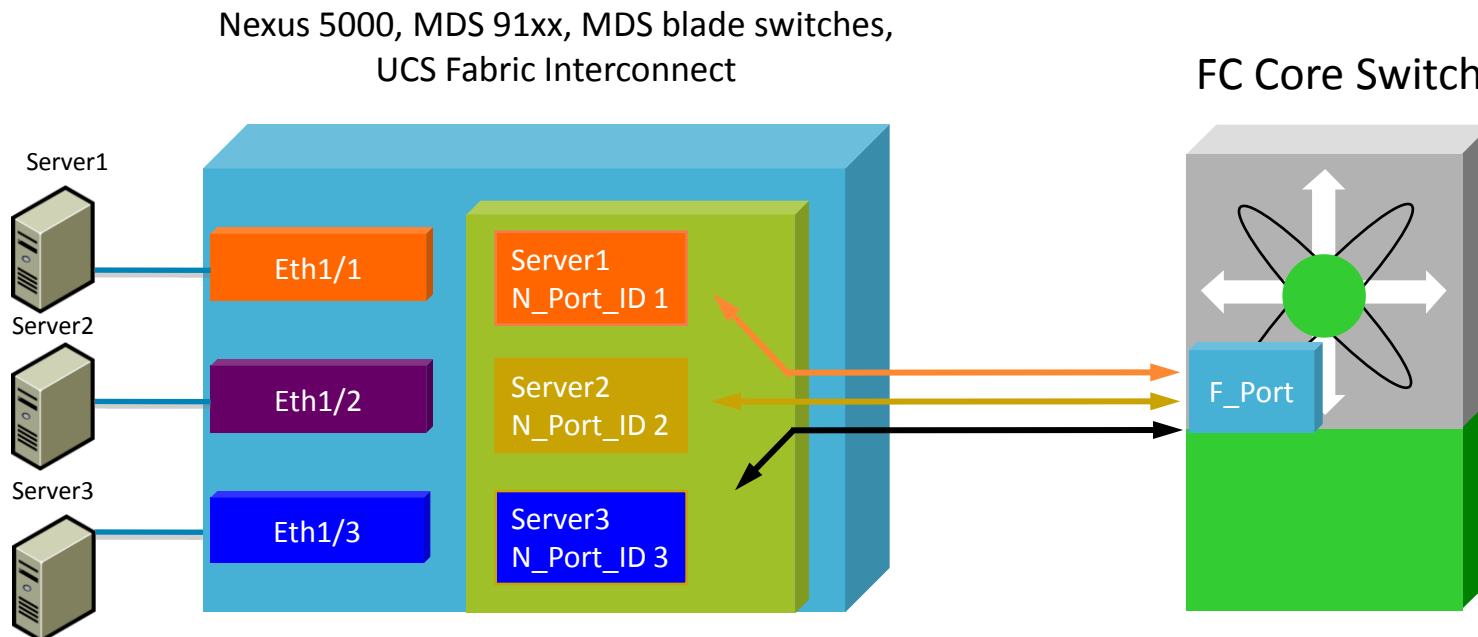
What is NPIV?

- N-Port ID Virtualisation (NPIV) provides a means to assign multiple FC IDs to a single N port.
- This feature was intended to allow multiple applications to share the same Fiber Channel HBA
- The use of different pWWN allows access control, zoning, and port security to be implemented at the application level.
- Usage applies to applications such as VMWare, MS Virtual Server and Citrix



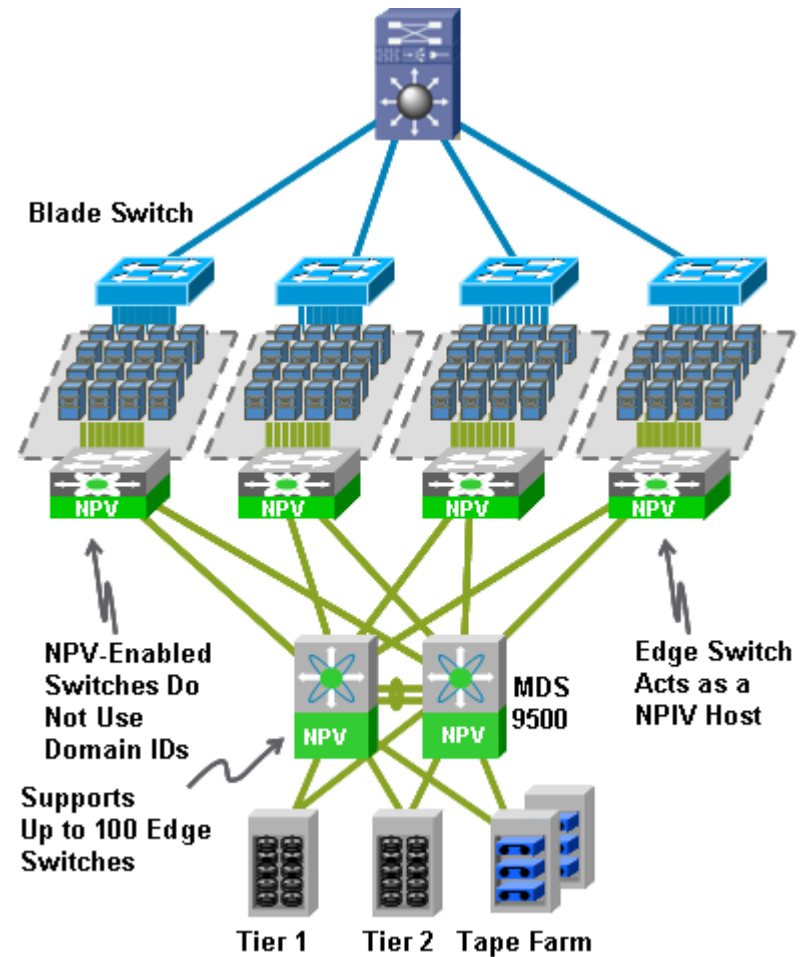
What is NPV?

- N-Port Virtualizer (NPV) utilizes NPIV functionality to allow a “switch” to act like a Server doing multiple logins through 1 physical link
- Real server connected (via CNAs) to Nexus 5000 do not login to Nexus 5000 but to upstream FC switch. The same applies to FC edge switches (ex.: MDS blade switches).
 - Physical uplink from Nexus 5000 to FC switch does actual “FLOGI”
 - Subsequent server logins are converted (proxy) to “FDISC” to login to upstream FC switch
- No local switching is done on an FC switch in NPV mode
- FC edge switch in NPV mode Does NOT take up a domain id
- Scalability will be dependent on FC “login” limitation (MDS is ~10K per fabric)



NPV at FC blade switches

- Eliminates edge switch Domain ID
- Edge switch acts as an NPIV host
- Simplifies server and SAN management and operations
- Increases fabric scalability

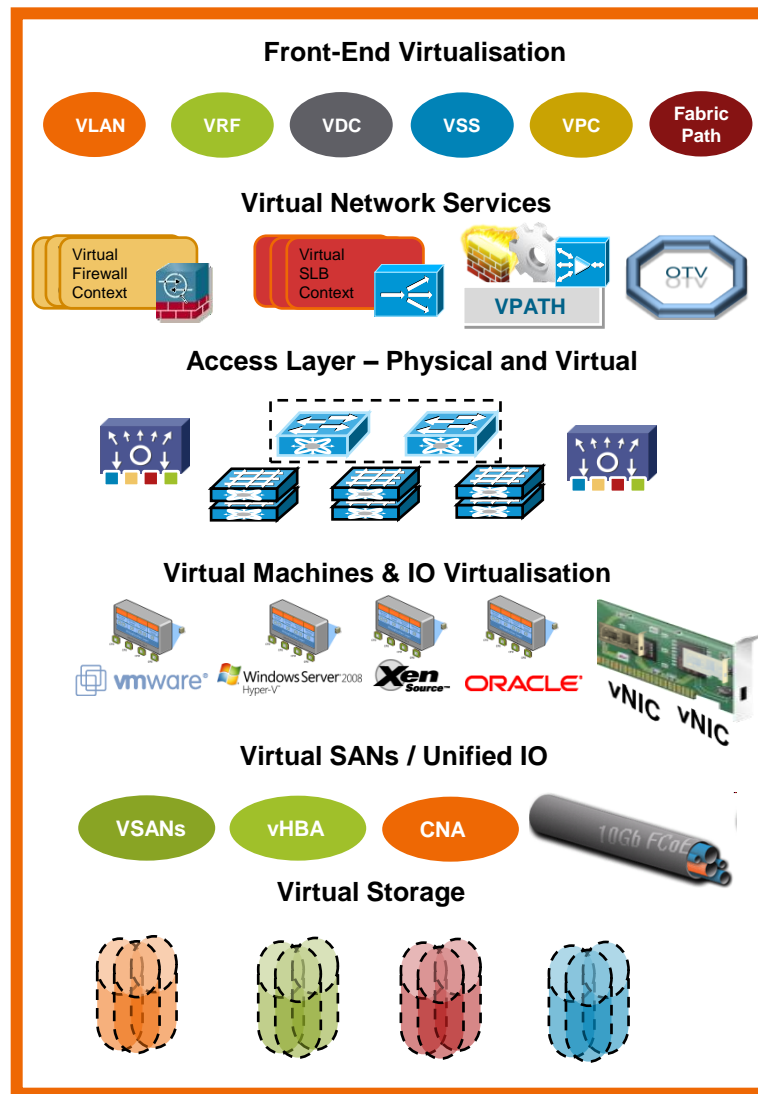


Switching – NPV Comparison for FCoE

	Switching	NPV
Scalability	Limited with number of supported physical switches in a fabric = Domain ID (typically by OSM)	Not dependent on number of physical switches but number of supported logins in a fabric (more scalable)
Interoperability	No issues with Cisco MDS. Some considerations with non-Cisco FC switches	No interoperability issues and only requires FC upstream switch to support NPIV (standards base feature)
Local Switching	Yes; Beneficial if native FCoE targets resides on Nexus 5000 or small scale switch deployments	No; All FCoE traffic must flow to upstream FC switch
Load-Balancing	Utilizes standard load-balancing (default is exchange-base on MDS)	Source/Destination
Redundancy	Utilizes normal MDS FC Port-Channel	In case of failure, server will logout and login to next available link for particular VSAN (F-Port Channeling solves it)
Trunking VSANs	Yes; Supports TE to MDS	F-Port Trunking solves it

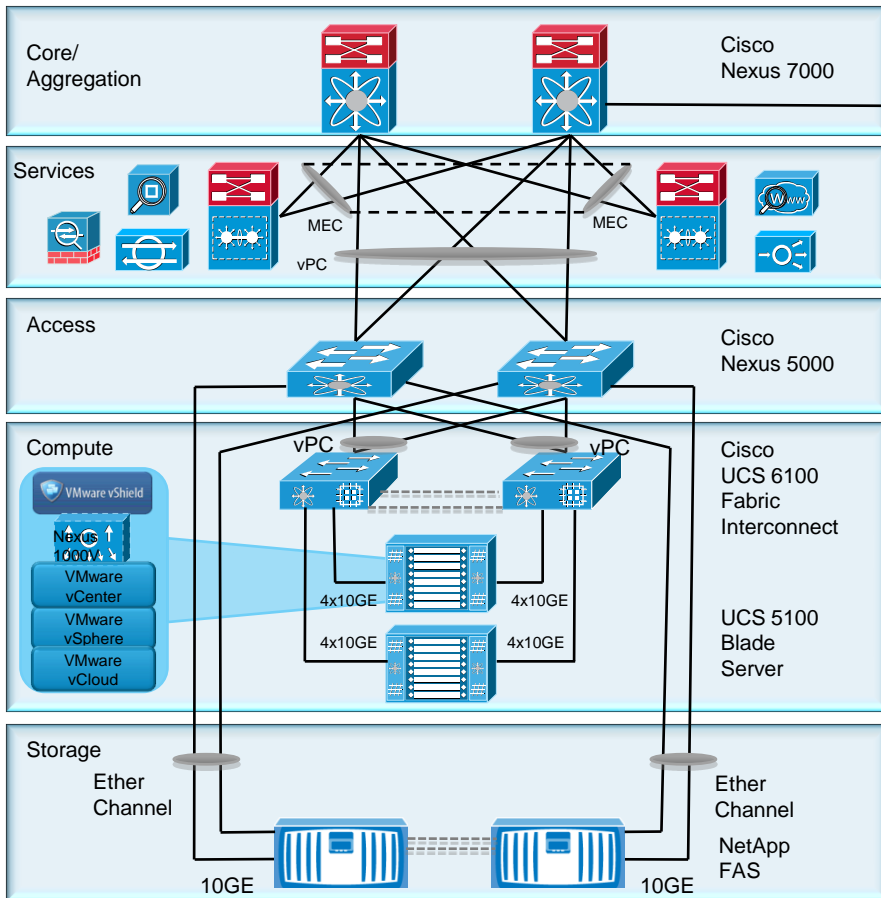
Agenda

- Data Centre Virtualisation Overview
- Front-End Data Centre Virtualisation
 - Core Layer
 - Aggregation Layer
 - Networking Services
 - Access Layer
- Server Virtualisation
 - Hypervisors
 - Virtual Access Layer
 - Virtualised Services
 - Server IO Virtualisation
- Back-End Virtualisation
 - Virtual HBA & NPV
 - Unified IO & FCoE
 - SAN & Storage
- **Implementation Examples**
 - vBlock**
 - Secure Cloud**
 - FlexPOD**
- **Q&A**

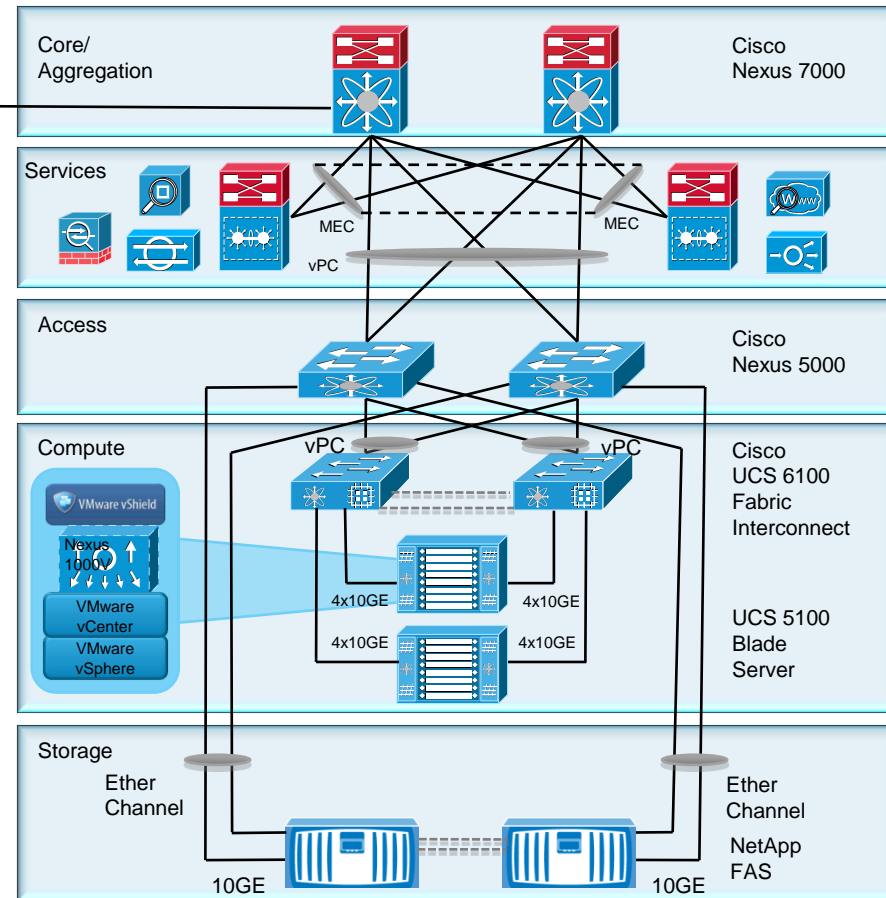


Ex.2: Enhanced SMT (E-SMT) Architecture

Primary DC Site

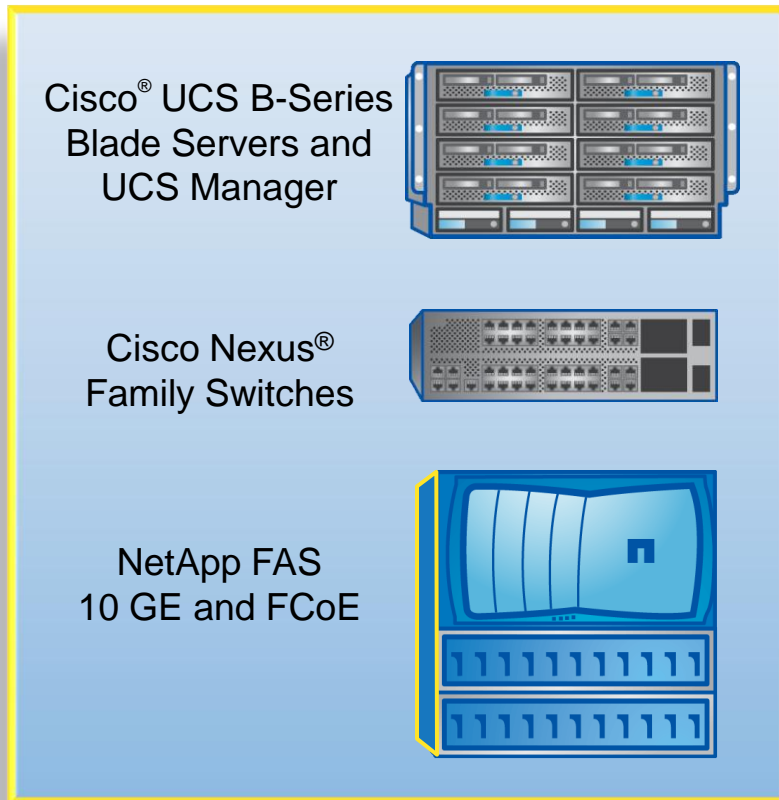


Second DC Site (DR)



40 KM

Ex. 3: FlexPod





- Standard, prevalidated, best-in-class infrastructure building blocks
- Flexible: One platform scales to fit many environments and mixed workloads
 - Add applications and workload
 - Scale up and out
- Simplified management and repeatable deployments
- Design and sizing guides
- Services: Facilitate deployment of different environments

Q&A; Q&Q; A&A; etc...



BRKVIR-2931

Recommended Reading





Network Virtualization

Provide secure network services to diverse user communities

Victor Moreno, CCIE® No. 6908
Kumar Reddy

ciscopress.com





Interconnecting Data Centers Using VPLS

Ensure Business Continuance on Virtualized Networks by Implementing Layer 2 Connectivity Across Layer 3

Nash Darukhanawalla, CCIE® No. 10332
Patrice Bellagamba

ciscopress.com



I/O Consolidation in the Data Center

A complete guide to Data Center Ethernet and Fibre Channel over Ethernet

Silvano Gai
Claudio DeSanti

ciscopress.com

Please complete your Session Survey

- We value your feedback - don't forget to complete your online session evaluations after each session. Complete 4 session evaluations & the Overall Conference Evaluation (available from Thursday) to receive your Cisco Networkers 20th Anniversary t-shirt.
- All surveys can be found on our onsite portal and mobile website: www.ciscoliveeurope.com/connect/mobi/login.ww
- You can also access our mobile site and complete your evaluation from your mobile phone:
 1. Scan the Access Code
(See <http://tinyurl.com/qrmelist> for software, alternatively type in the access URL)
 2. Login
 3. Complete and Submit the evaluation



