Cisco Expo
2009

Data Center Virtualization

***René Raeber** CE Datacenter*
Central Consulting Advanced Technologies/DC

Welcome to the Human Network.

cisco

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 !**

Agenda Datacenter Virtualization

- **Data Center Virtualization Overview**
- **Front End DC Virtualization**
- **Server Virtualization**
- **Back-End Virtualization**
- **Conclusion & Direction Q&A**

Cisco Expo
2009

Virtualization Overview

Welcome to the Human Network.

cisco

Cisco Expo 2009

The "Virtual Data Center" Approach

Abstracting Server Hardware From Software together with Consolidation

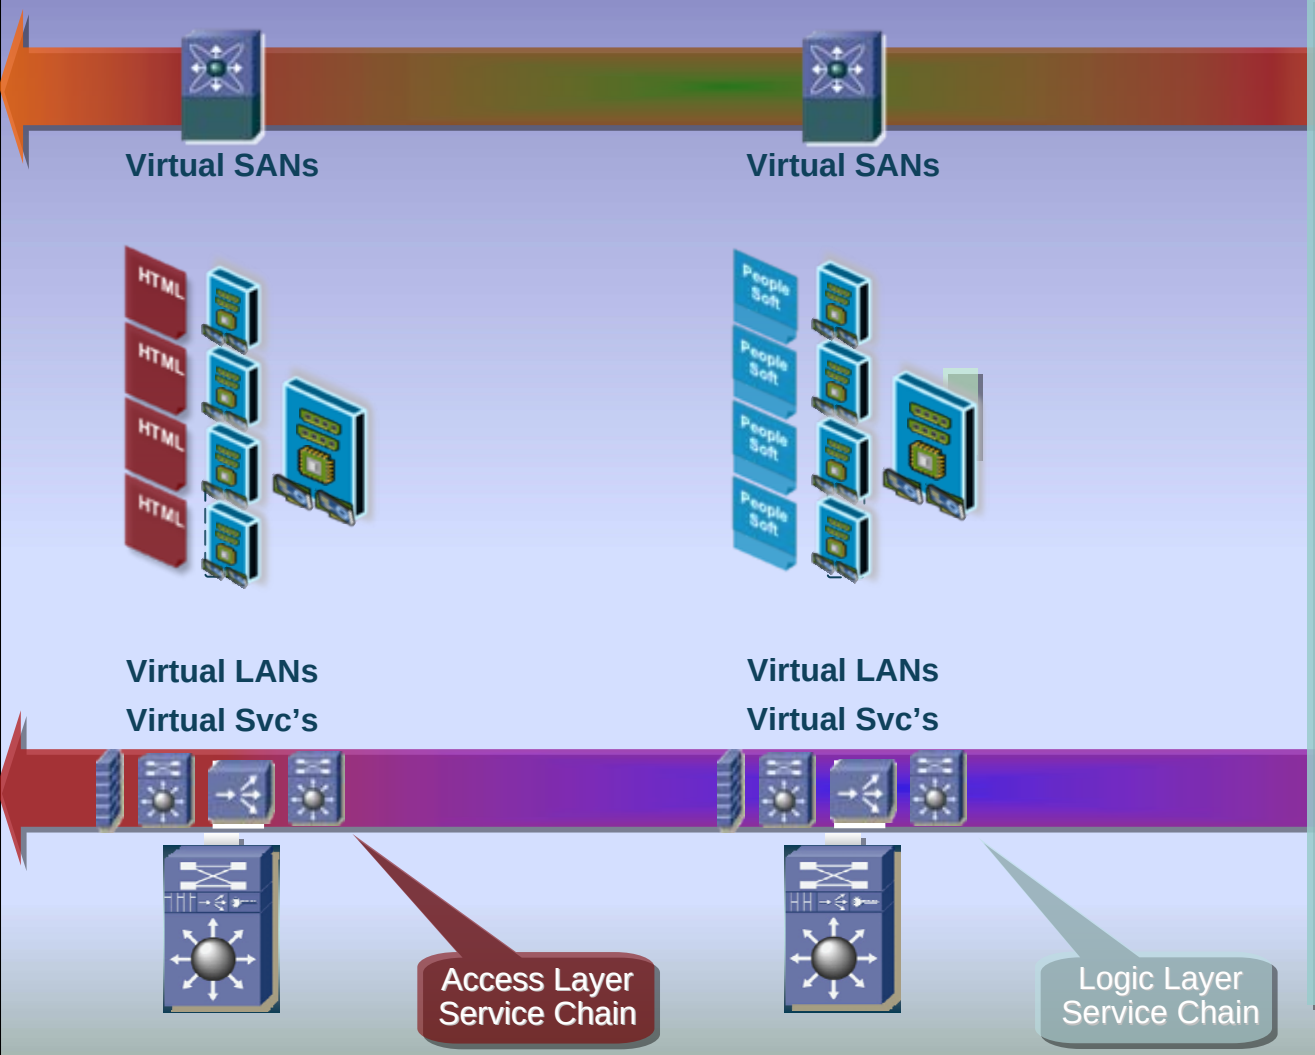

- Existing Service Chains are still aligned to the instances of Virtual Servers running in place of physical servers.
- VLANs at the Virtual Machine (Hypervisor) level, map to VLANs at the Network Switch Layer.
- Storage LUN's are similarly directly mapped to the VM's in the same way they would map to physical servers.

The Flexibility of Virtualization

VM's Mobility Across Physical Server Boundaries and Keeping Services

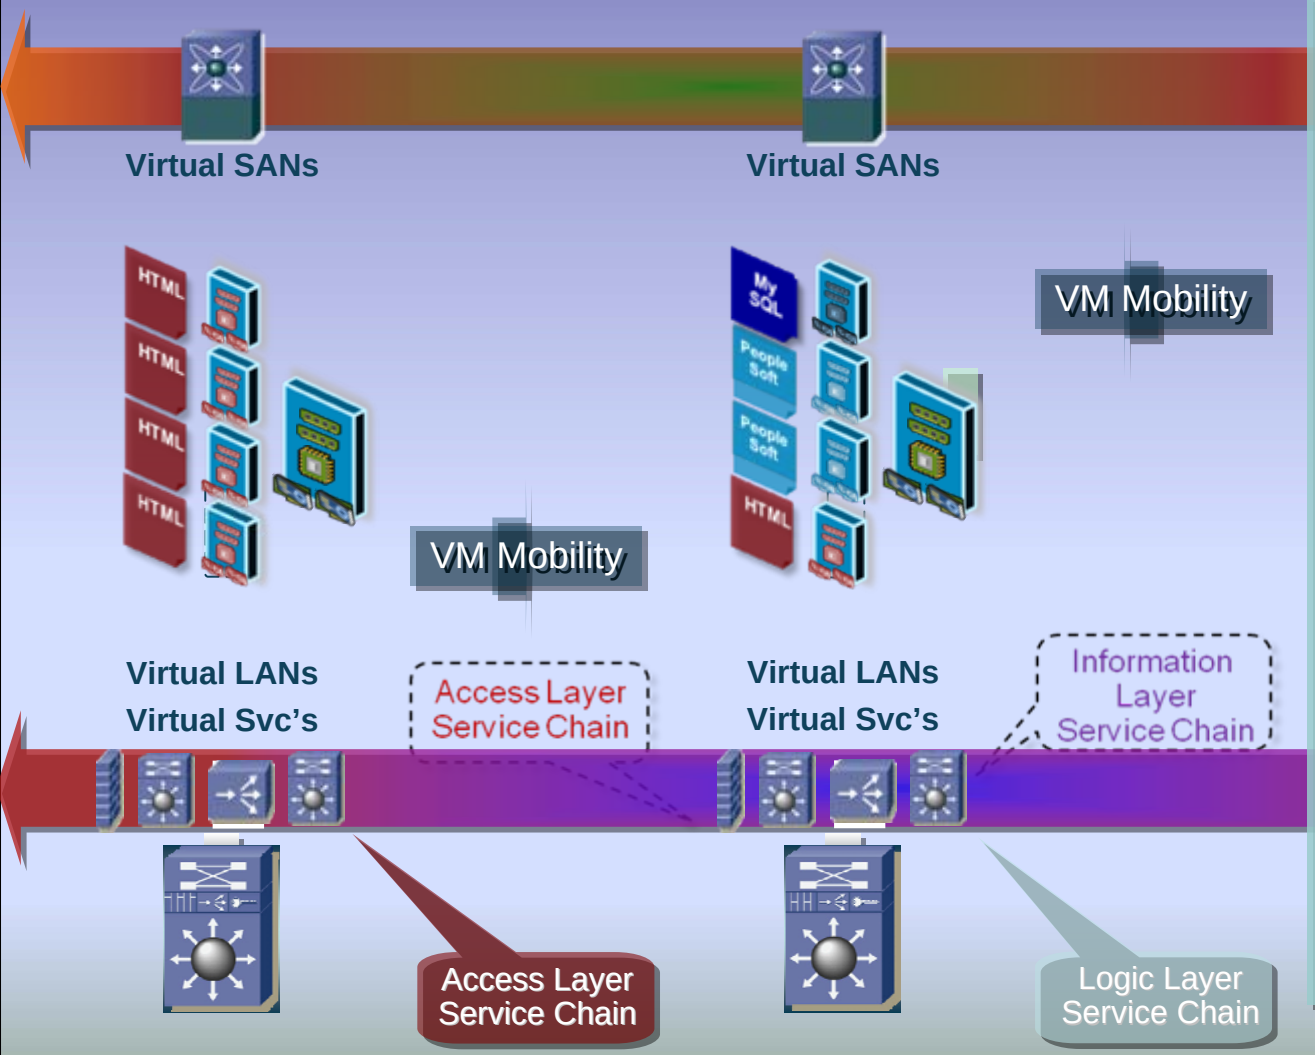

- **VM Mobility is capable of moving Virtual Machines across Physical Server**
- **The Application Services provided by the Network need to respond and be aligned to meet the new geometry of the VMs**
- **Close interaction required between the assets provisioning virtualized infrastructure and the Application Services supporting the Virtual Machines.**

Moving to a Unified Fabric

Moving to a fully Virtualized Data Center, with Any To Any Connectivity

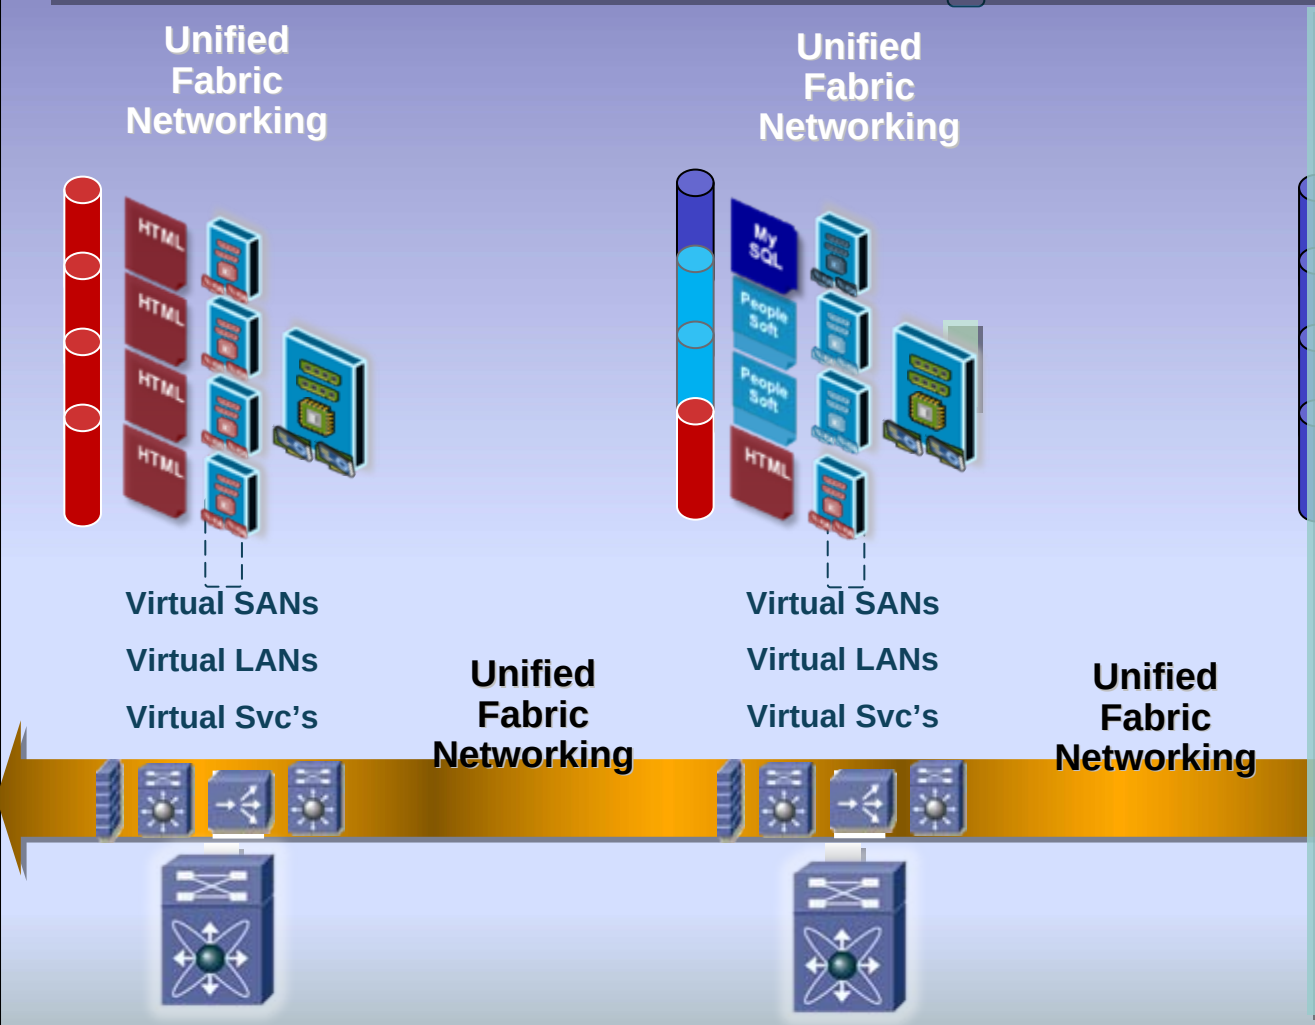

- Fully unified I/O delivers the following characteristics:
  - Ultra High Capacity 10Gbps+
  - Low latency
  - Loss Free (FCoE)
- True "Any to Any" Connectivity is possible as all devices are connected to all other devices.
- We can now simplify management, operations and enhance power and cooling efficiencies

Network Virtualization Building Blocks

Device Partitioning

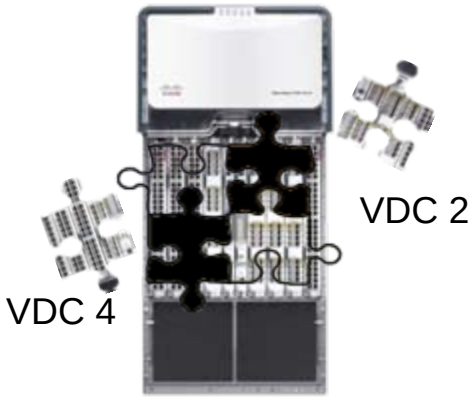

VDCs FW,ACE context

VLANs VRFs

1 : n

Virtualized Interconnect

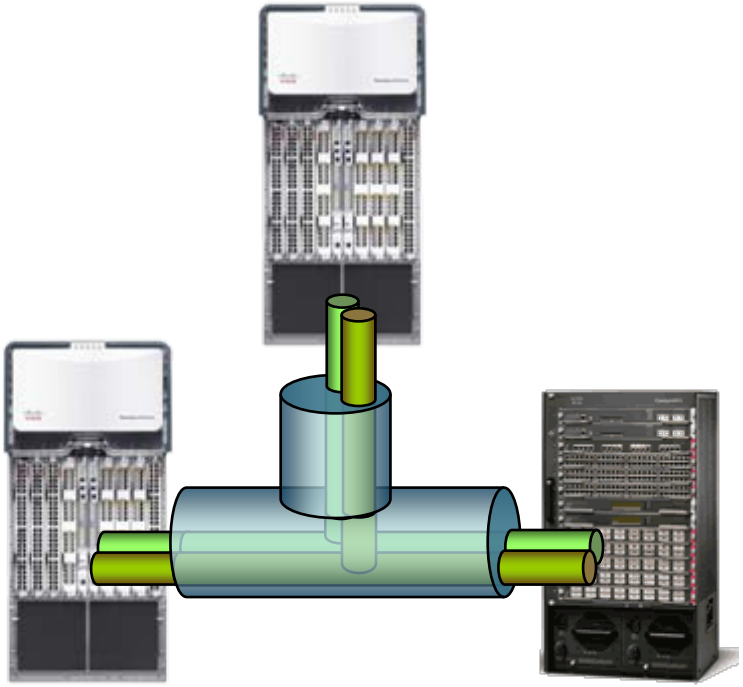

L3 VPNs – MPLS VPNs, GRE, VRF-Lite, etc.

L2 VPNs - AToM, Unified I/O, VLAN trunks, PW, etc.

n : m

Device Pooling

VSS, Stackwise, VBS,

Virtual Port Channel (vPC)

HSRP/GLBP

n : 1

Cisco Expo 2009

Virtualized Data Center Infrastructure

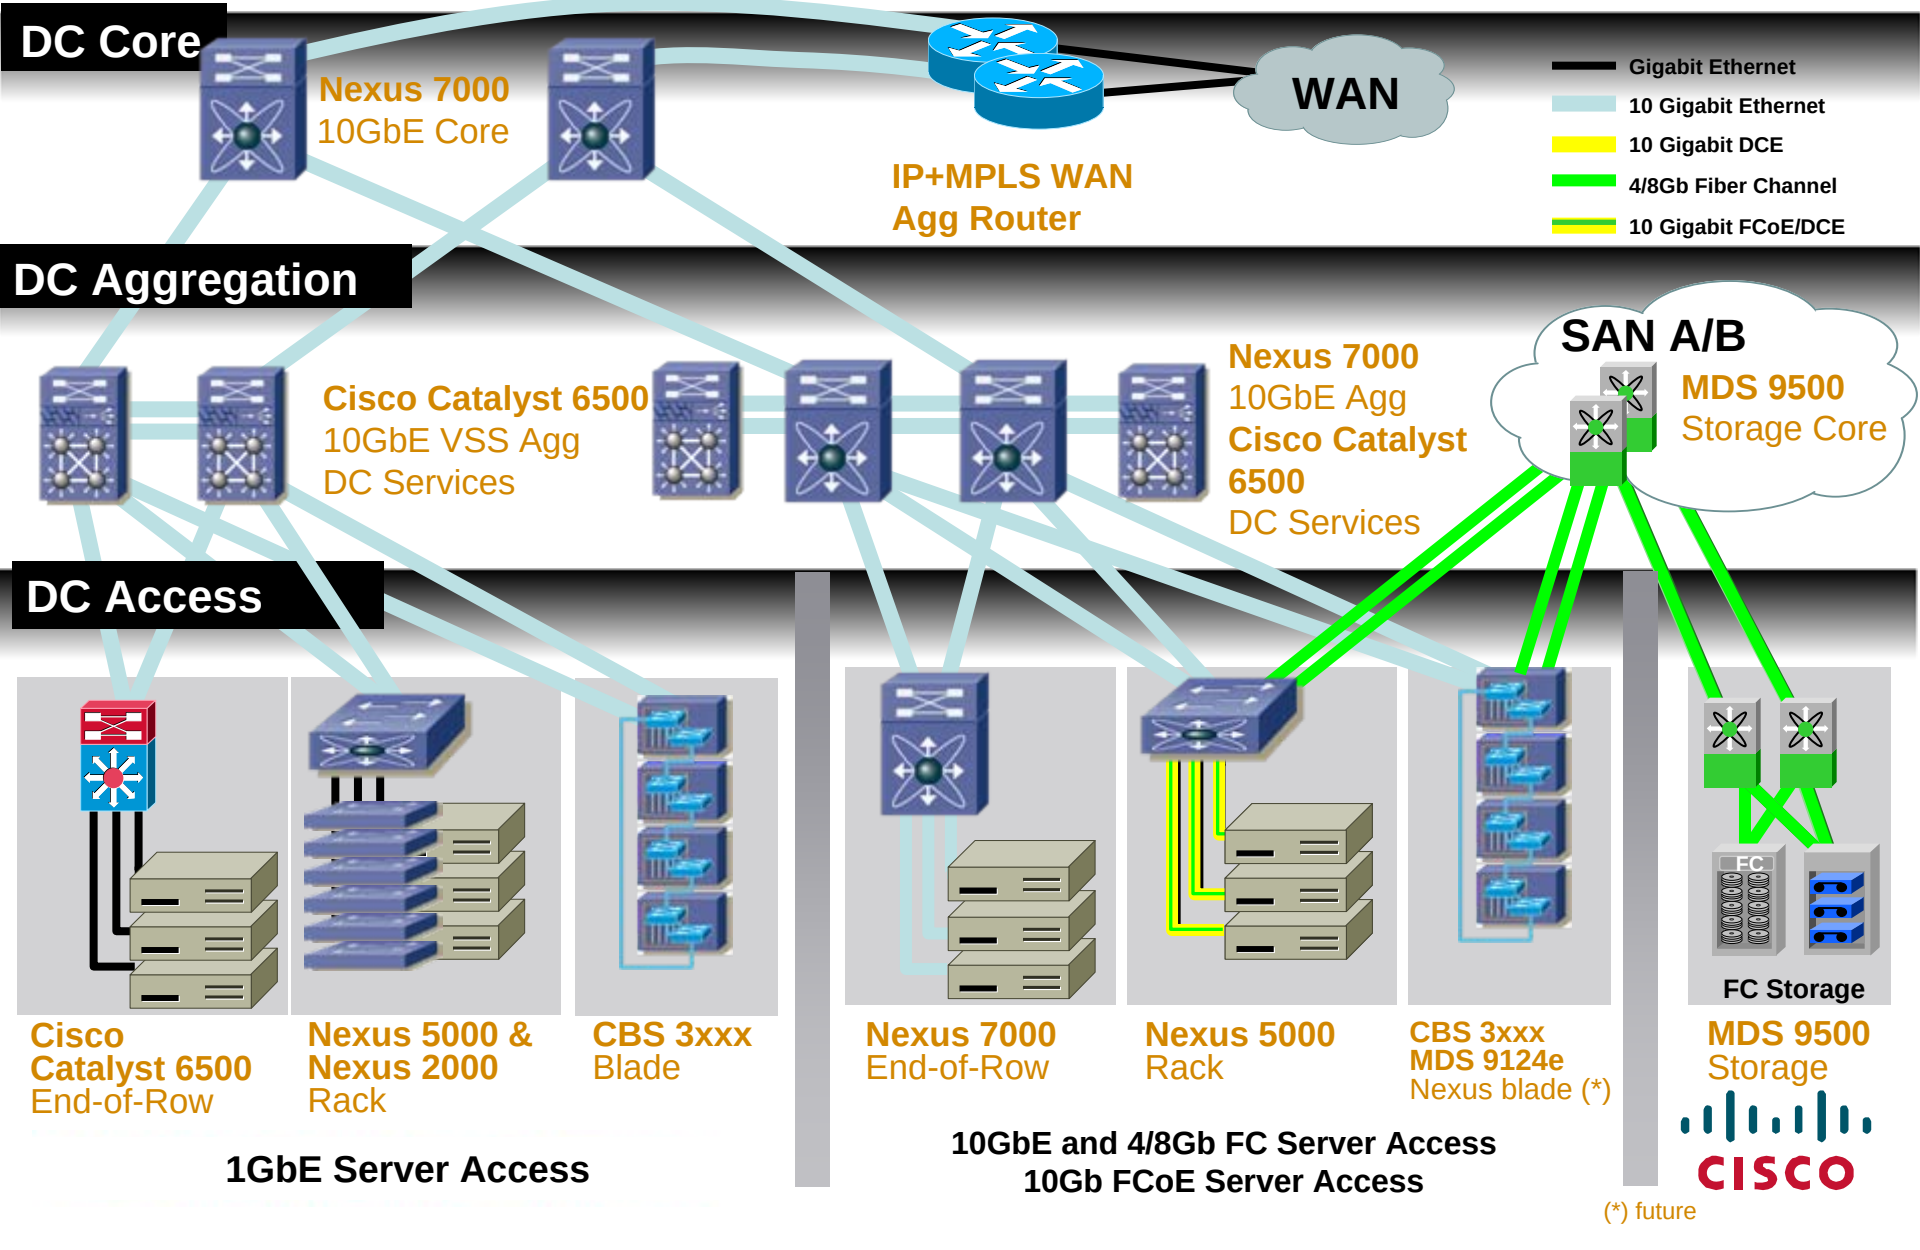

Cisco Expo
2009

Front-End Virtualization

Welcome to the Human Network.

cisco

Virtual Device Contexts at Nexus 7000

VDC Architecture

Virtual Device Contexts Provides Virtualization at the Device Level Allowing Multiple Instances of the Device to Operate on the Same Physical Switch at the Same Time

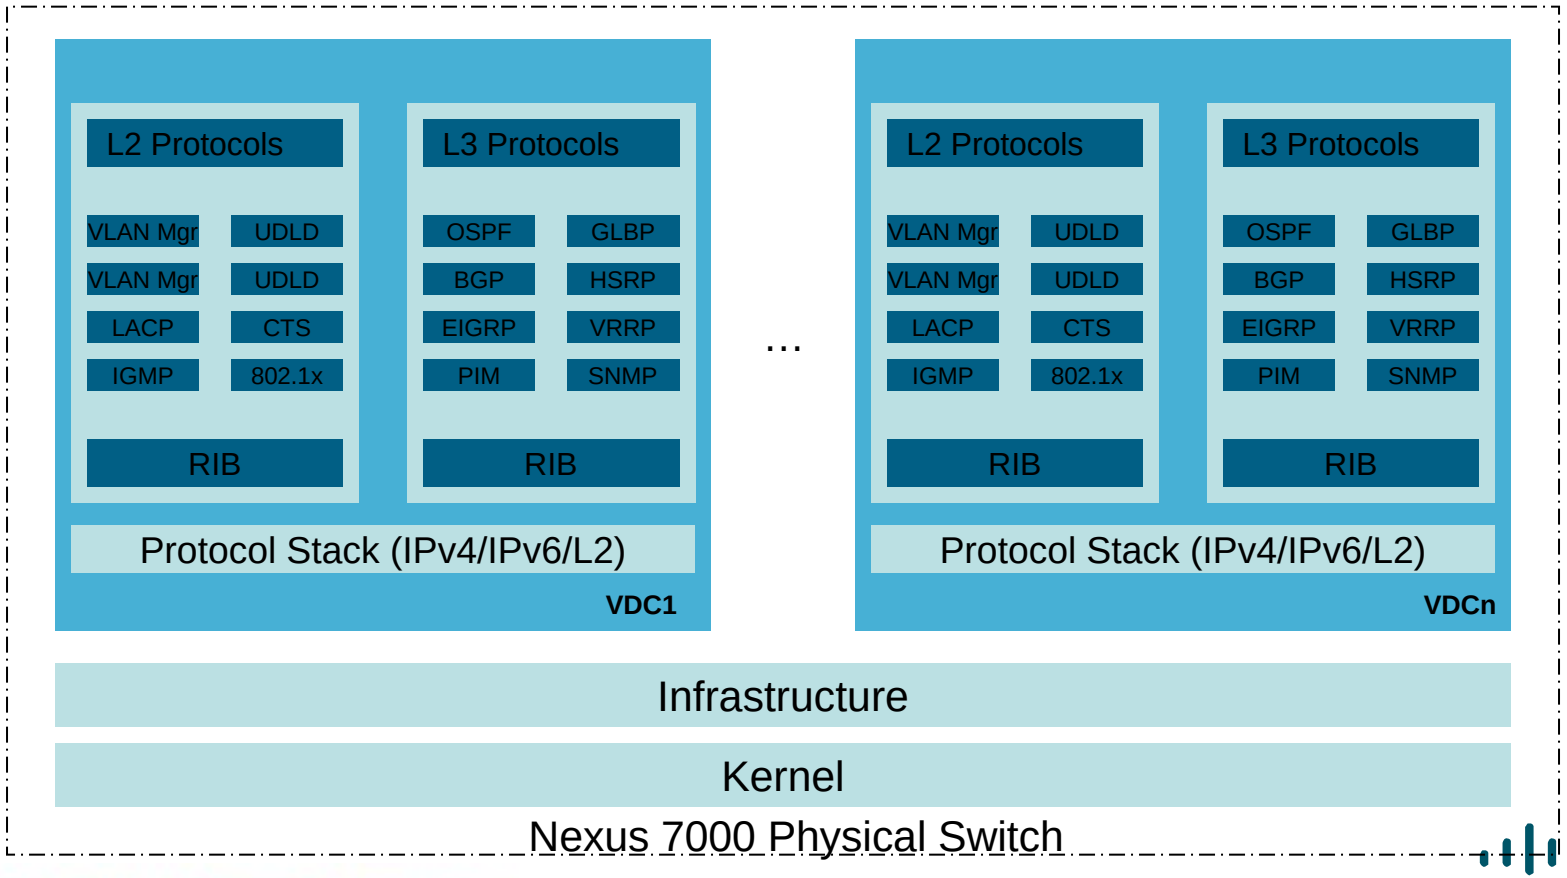

Virtual Device Contexts

VDC Fault Domain

A VDC Builds a Fault Domain Around All Running Processes Within That VDC—Should a Fault Occur in a Running Process, It Is Truly Isolated from Other Running Processes and They Will Not Be Impacted

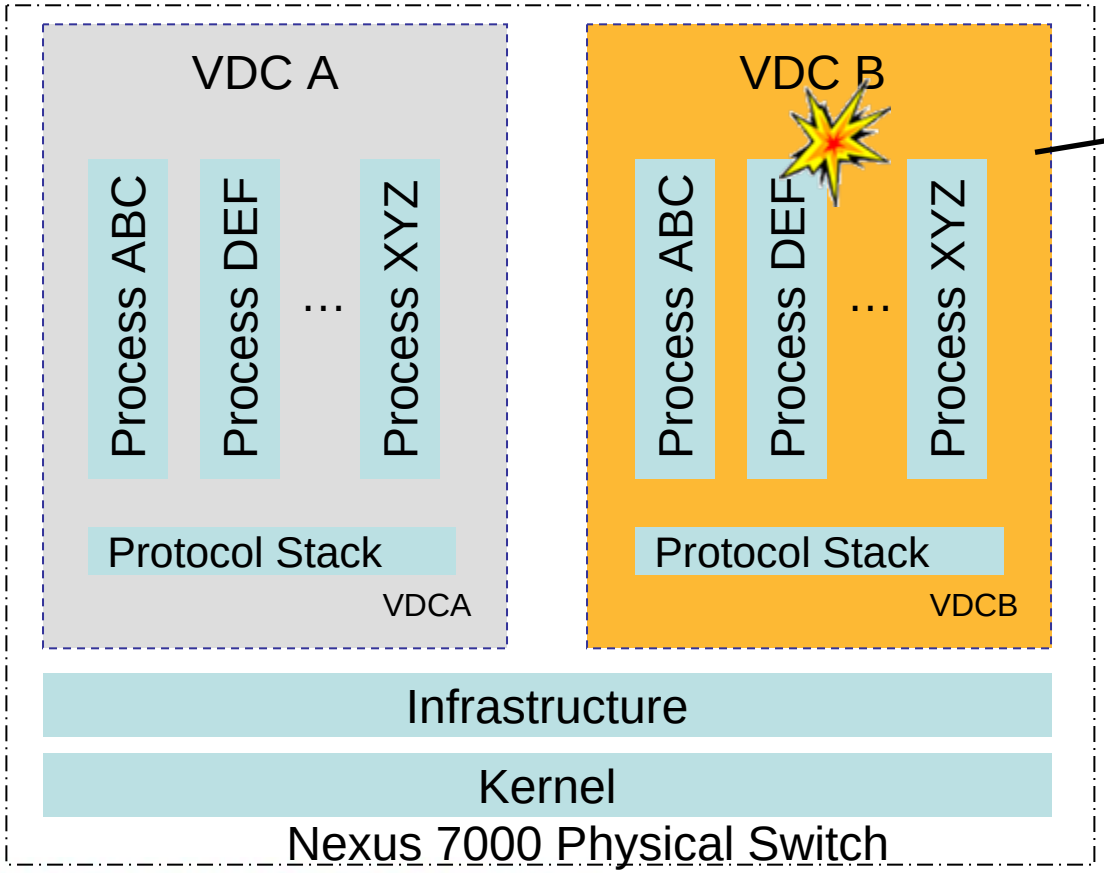

Process "DEF" in VDC B Crashes

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

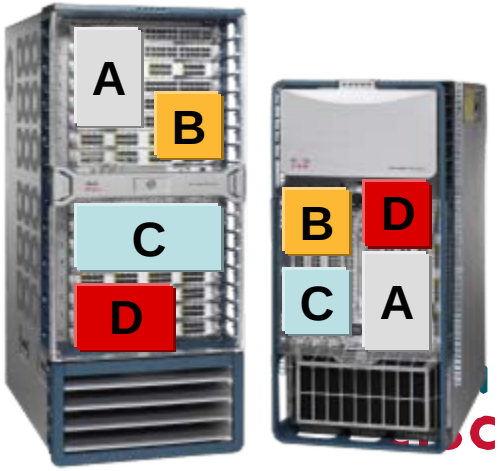

Virtual Device Contexts

VDC and Interface Allocation

VDC A

VDC C

Ports Are Assigned on a per VDC Basis and Cannot Be Shared Across VDCs

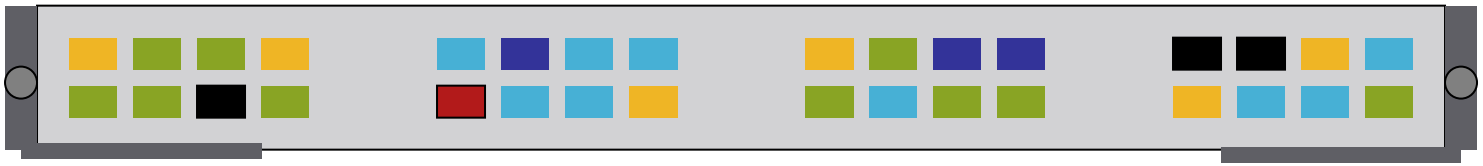

32-Port 10GE Module

VDC B

VDC C

Once a Port Has Been Assigned to a VDC, All Subsequent Configuration Is Done from Within That VDC…

VDC Use Case Examples
Security Partitioning

- Some Infosec departments are still reluctant about collapsed infrastructure
- Concerns around change management
- Infrastructure misconfiguration could bypass policies

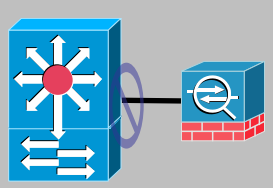
Appliance Model

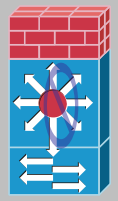
Service Module Model

- Ideally they want to have physically separately infrastructure.
- Not cost effective in larger deployments.

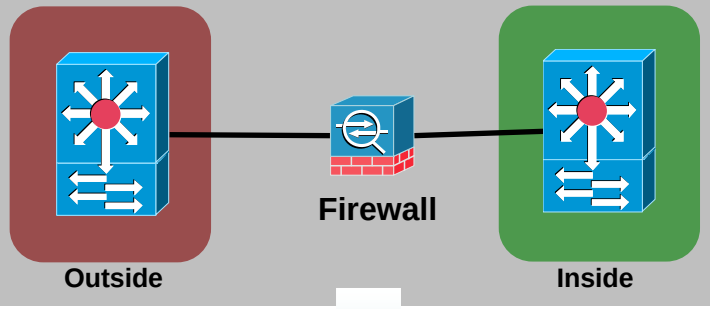

- VDCs provide logical separation simulating air gap
- Extremely low possibility of configuration bypassing security path – Must be physically bypassed
- Model can be applied for any DC services

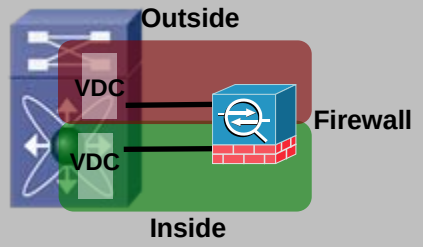

VDC Use Case Examples
Horizontal Consolidation

- Preface: **Lead with separate physical boxes** as they provide the most scalable solution. VDCs are useful in certain situations!
- Objective: Consolidate lateral infrastructure that delivers similar roles for separate operational or administrative domains.
- Benefits: Reduced power and space requirements, can maximize density of the platform, easy migration to physical separation for future growth
- Considerations: Number of VDCs (4), Four VDCs != Four CPU Does not significantly reduce cabling or interfaces needed.

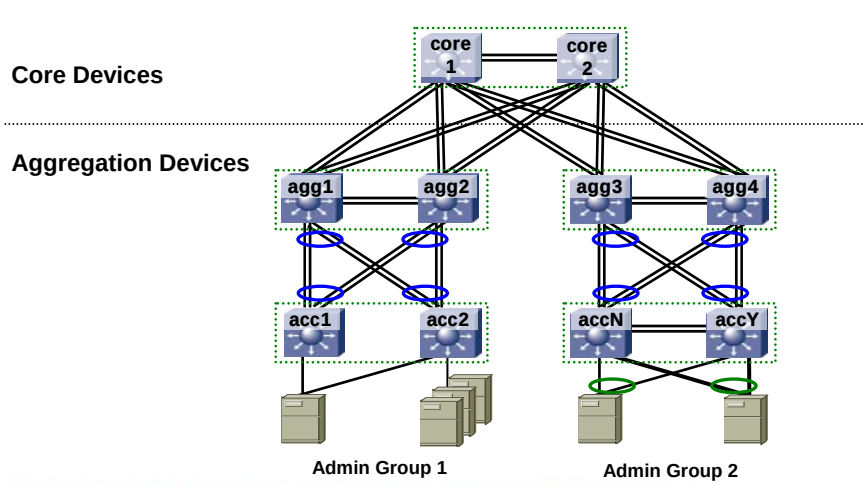

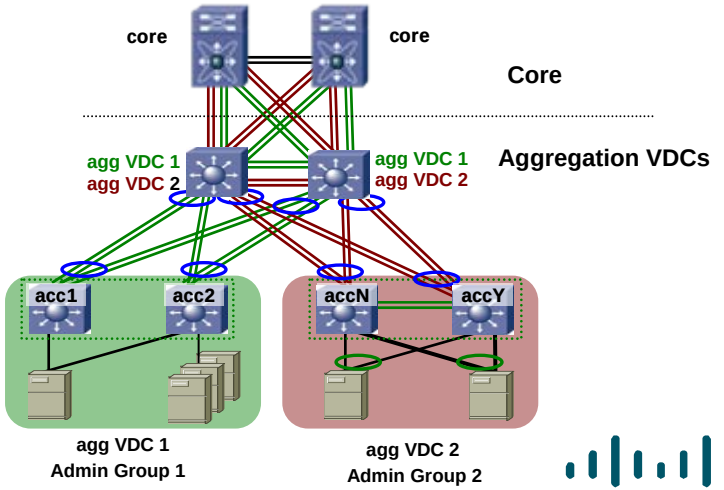

VDC Use Case Examples
Vertical Consolidation

- Preface: Lead with separate physical boxes as they provide the most scalable solution.
  - –Large Three Tier designs should remain physical.
  - –Smaller Two Tier designs can leverage VDCs for common logical design with three tier.
- Objective: Consolidate vertical infrastructure that delivers orthogonal roles to the same administrative or operational domain.
- Benefits: Reduced power and space requirements, can maximize density of the platform, provides smooth growth path, easy migration to physical separation in future
- Considerations: Number of VDCs (4), Four VDCs != Four CPU Intra-Nexus7000 cabling needed for connectivity between layers.

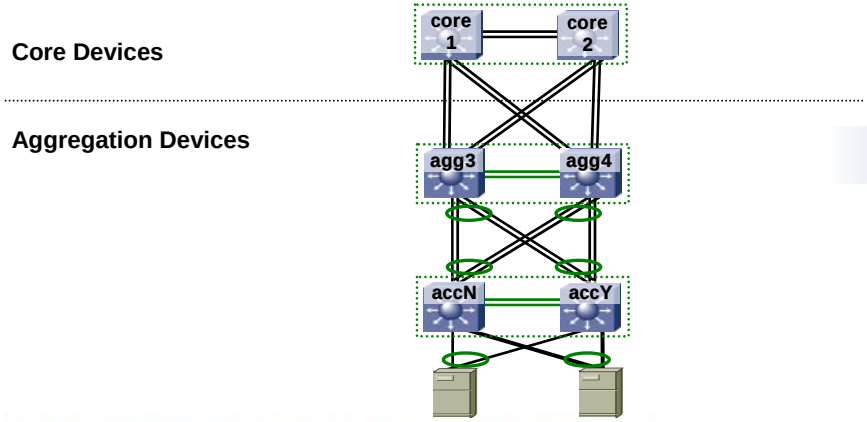

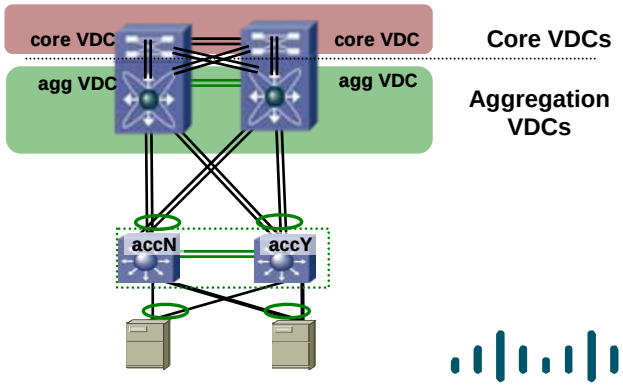

Core Virtualization

Virtual Port-Channel (vPC)
Feature Overview

- Allow a single device to use a port channel across two upstream switches
- Separate physical switches independent control and data plane
- Eliminate STP blocked ports. Uses all available uplink bandwidth
- Dual-homed server operate in active-active mode
- Provide fast convergence upon link/device failure
- Available in NX-OS 4.1 for Nexus 7000. Nexus 5000 availability planned for CY09.

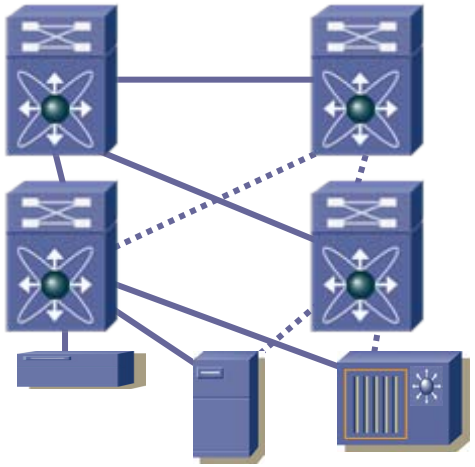

Logical Topology without vPC

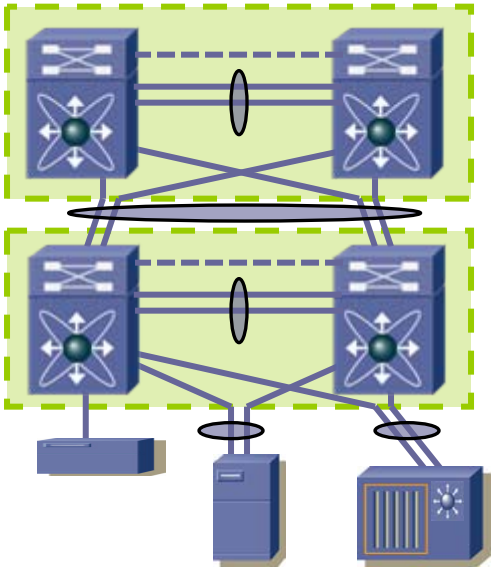

Logical Topology with vPC

Multi-level vPC

Physical View

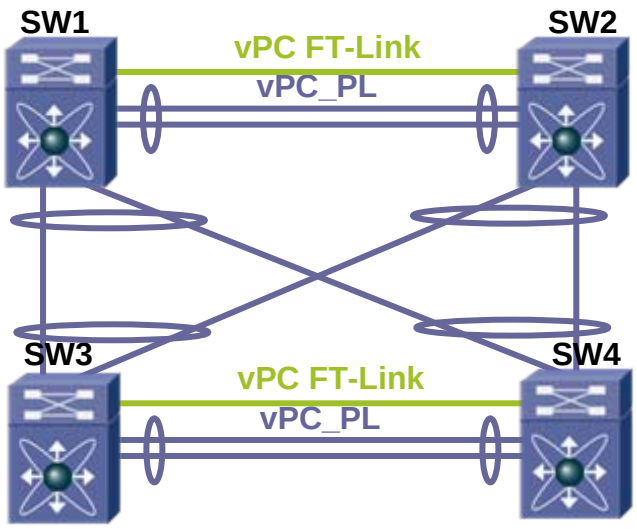

Logical View

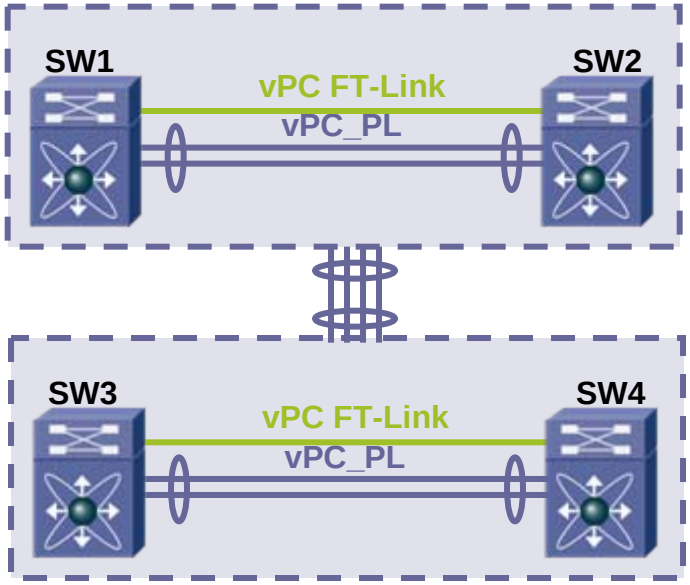

- Up to 16 links between both sets of switches: 4 ports from sw1-sw3, sw1-sw4, sw2-sw3, sw2-sw4
- Provides maximum non-blocking bandwidth between sets of switch peers
- Is not limited to one layer, can be extended as needed

Aggregation Virtualization

Cisco Expo 2009

Aggregation Services Design Options

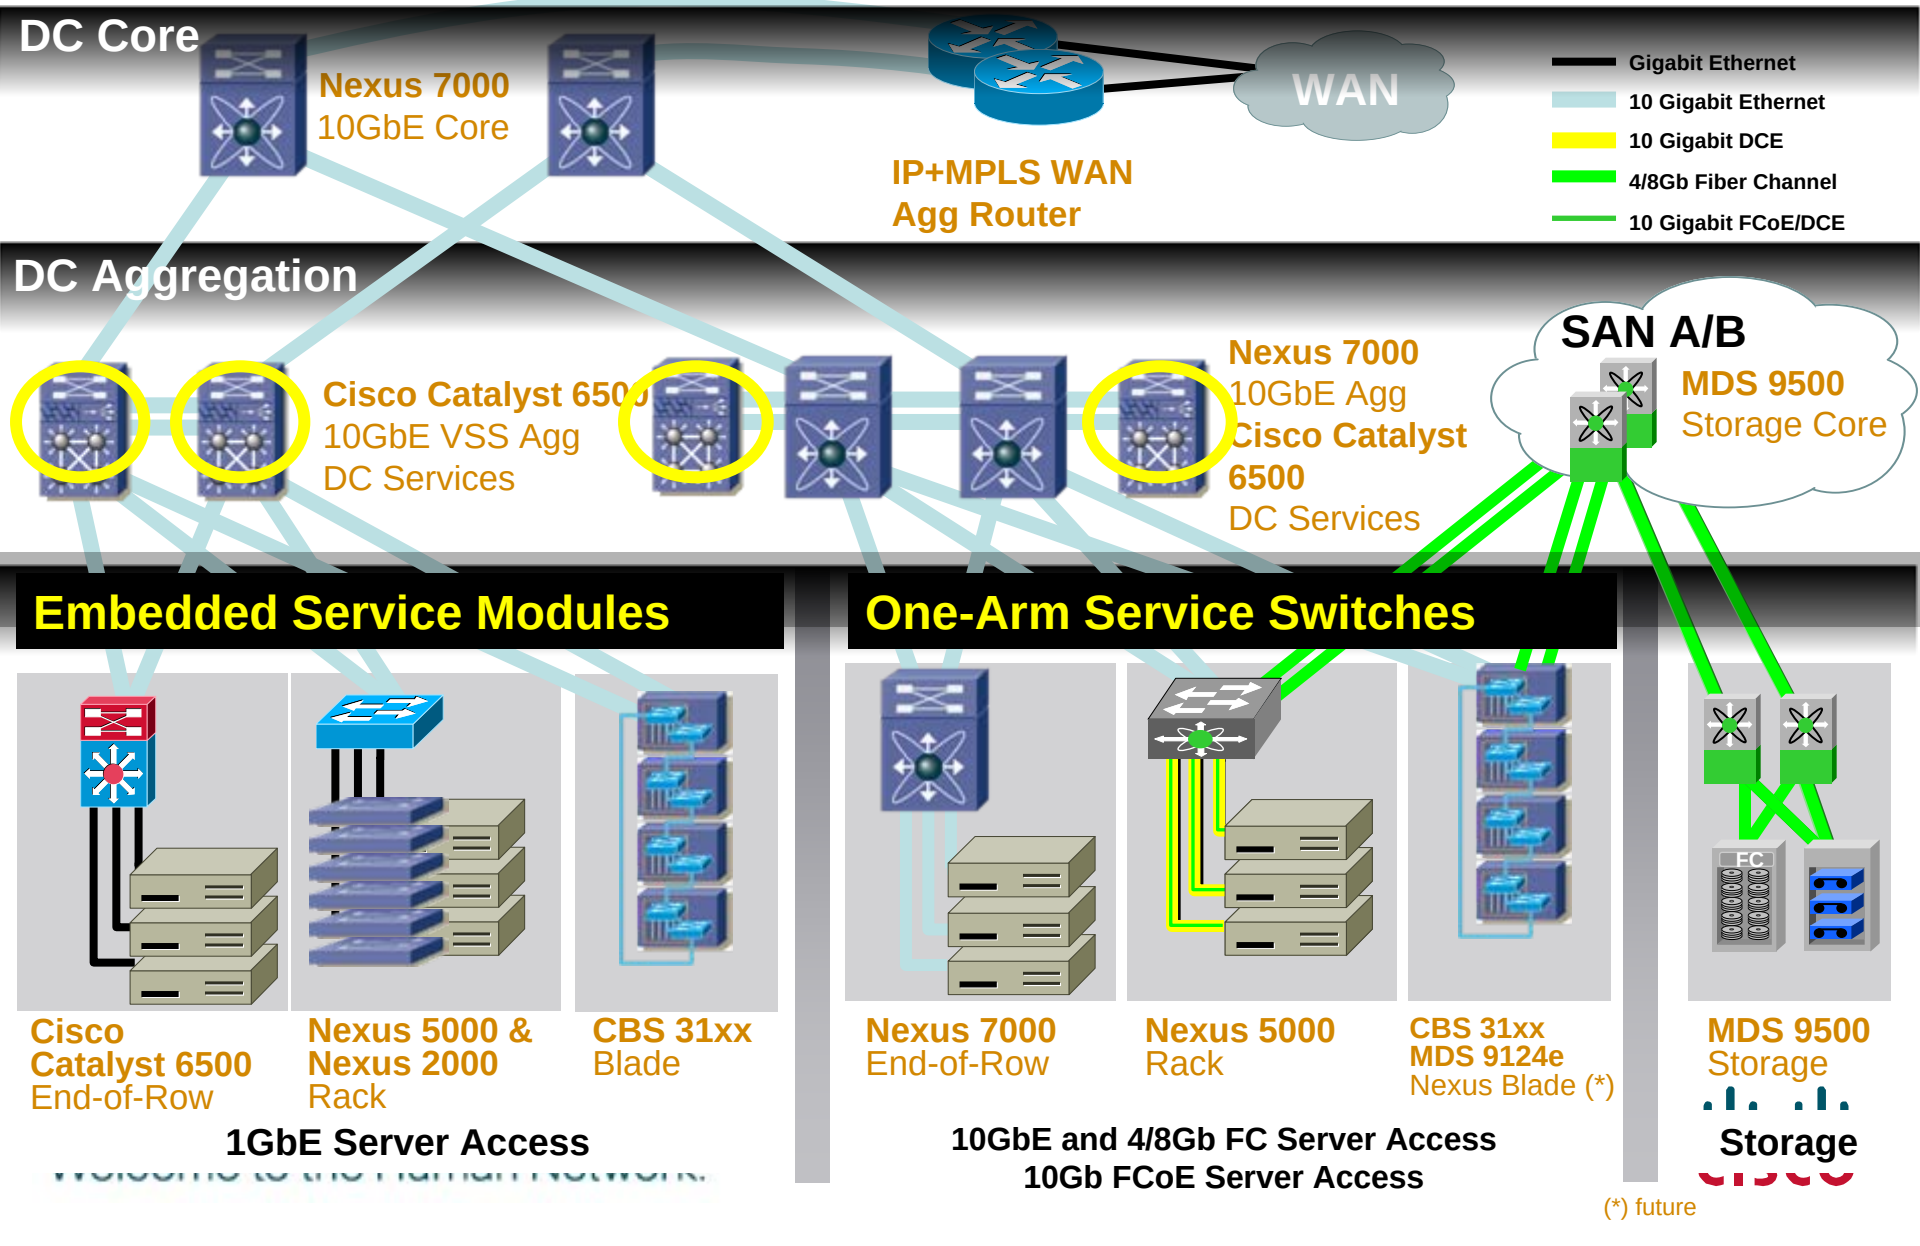

(*) future

Virtual Switch System (VSS)

Concepts

Virtual Switch System Is a Technology Break Through for the Cisco Catalyst 6500 Family

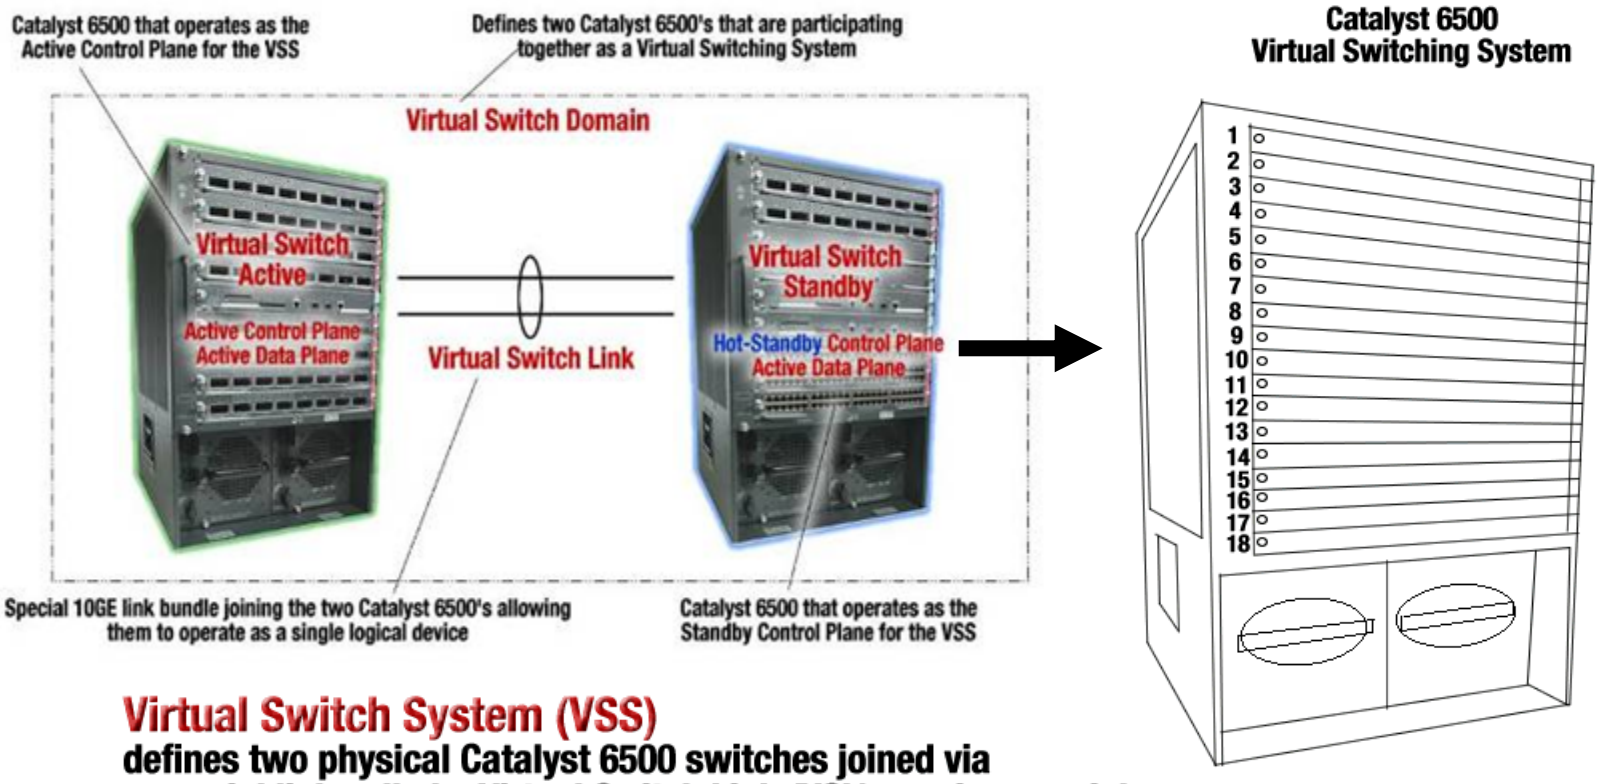

Virtual Switch System (VSS)
defines two physical Catalyst 6500 switches joined via a special link called a Virtual Switch Link (VSL) running special hardware and software that allows them to operate as a single logical switch

EtherChannel Concepts

Multichassis EtherChannel (MEC)

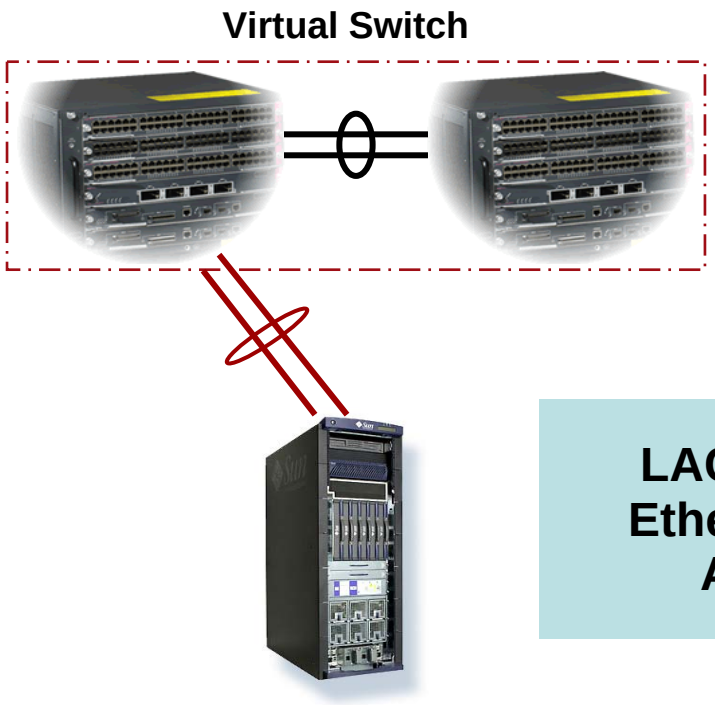

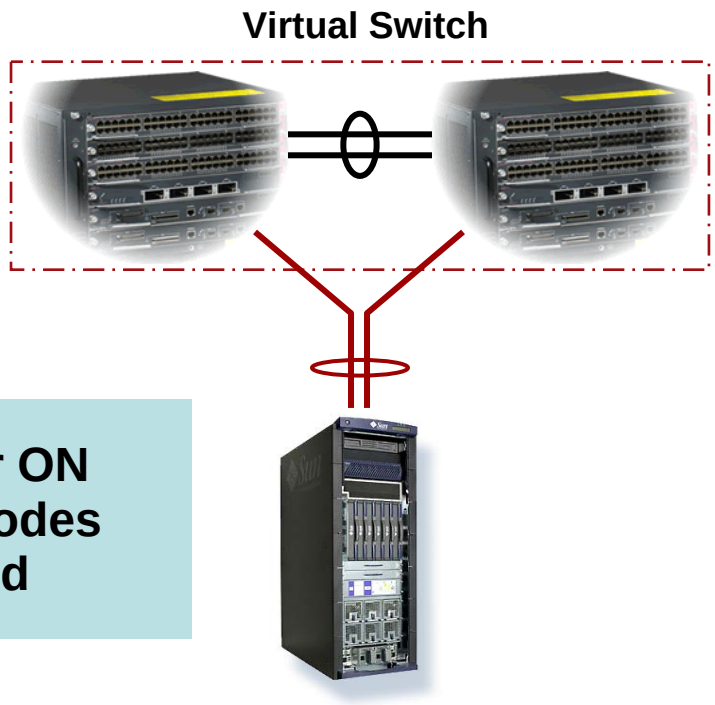

LACP, PAGP, or ON EtherChannel Modes Are Supported

Regular EtherChannel on Single Chassis

Multichassis EtherChannel (MEC) Across Two VSL-Enabled Chassis

ACE Module: Virtual Partitioning

One Physical Device

Multiple Virtual Systems (Dedicated Control and Data Path)

Traditional Device

- Single configuration file
- Single routing table
- Limited RBAC
- Limited resource allocation

Cisco Application Infrastructure Control

- Distinct context configuration files
- Separate routing tables
- RBAC with contexts, roles, domains
- Management and data resource control
- Independent application rule sets
- Global administration and monitoring
- Supports routed and bridged contexts at the same time

Firewall Service Module (FWSM)

Virtual Firewalls

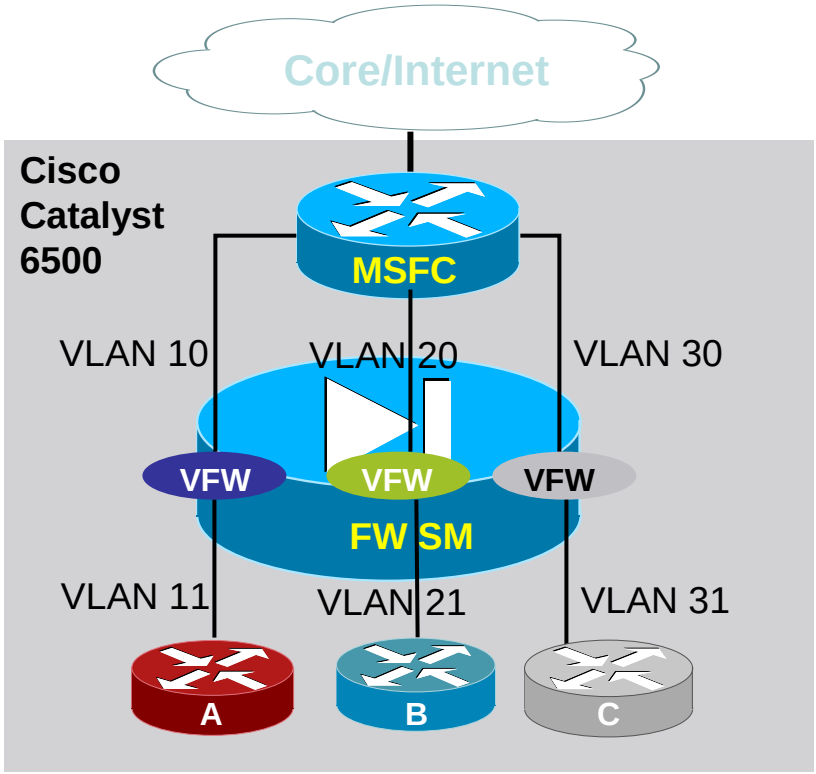

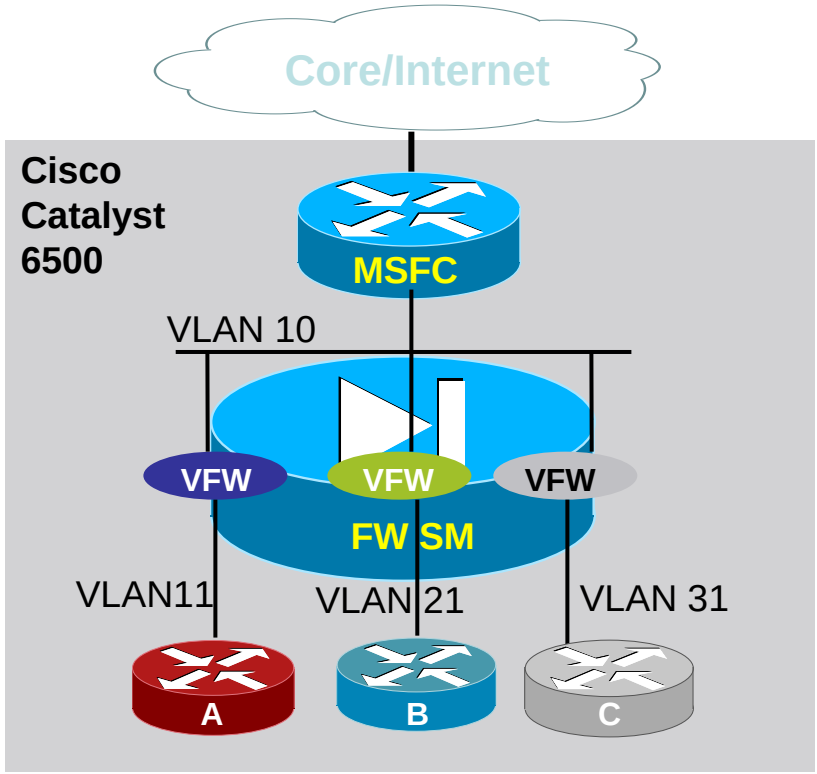

- e.g., Three customers → three security contexts—scales up to 250
- VLANs can be shared if needed (VLAN 10 on the right-hand side example)
- Each context has its own policies (NAT, access-lists, inspection engines, etc.)
- FWSM supports routed (Layer 3) or transparent (Layer 2) virtual firewalls at the same time

Data Center Virtualized Services Combination Example

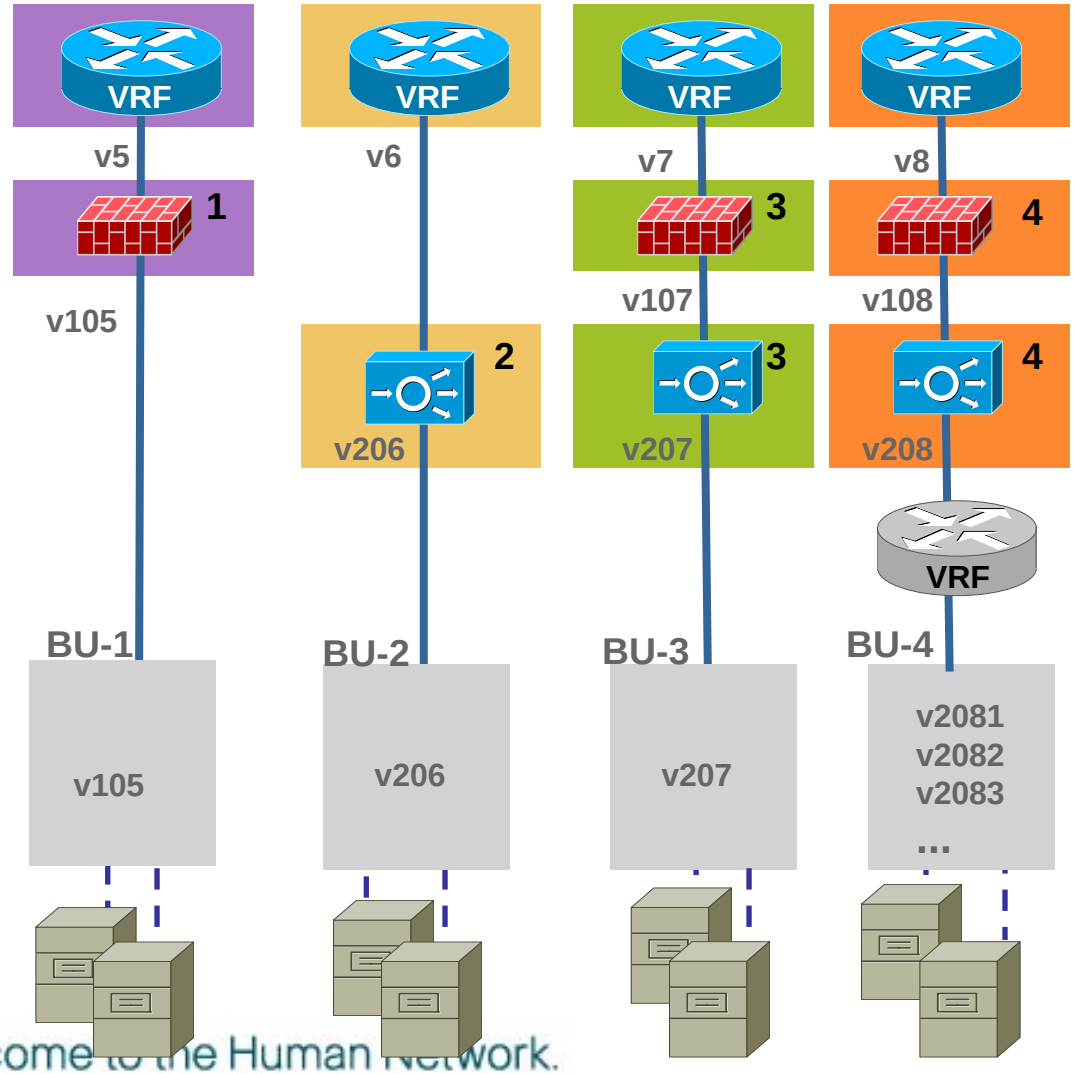

"Front-End" VRFs (MSFC)

Firewall Module Contexts

ACE Module Contexts

"Back-End" VRFs (MSFC)

Server Side VLANs

* vX = VLAN X
**BU = Business Unit

VSS with ACE and FWSM Modules

Active / Standby Pair

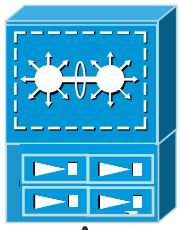

Virtual Switch System (VSS)

Switch-1 (VSS Active)

- Control Plane Active
- Data Plane Active
- ACE Active
- FWSM Standby

Failover/State sync Vlan

Switch-2 (VSS Standby)

- Control Plane Hot Standby
- Data Plane Active
- ACE Standby
- FWSM active

Combining vPC with VSS for Services

- Services can be…
  - attached using EtherChannel
  - Appliance based
  - Services-chassis based (standalone or VSS)

Nexus 7000 with vPC

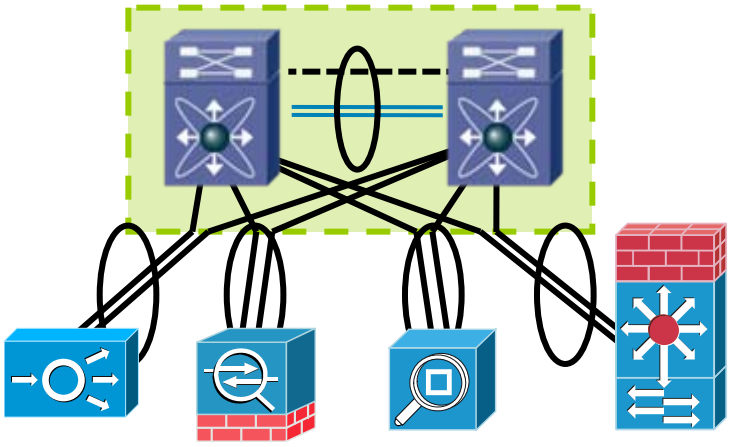

ACE Appliance | ASA | NAM Appliance | Services Chassis

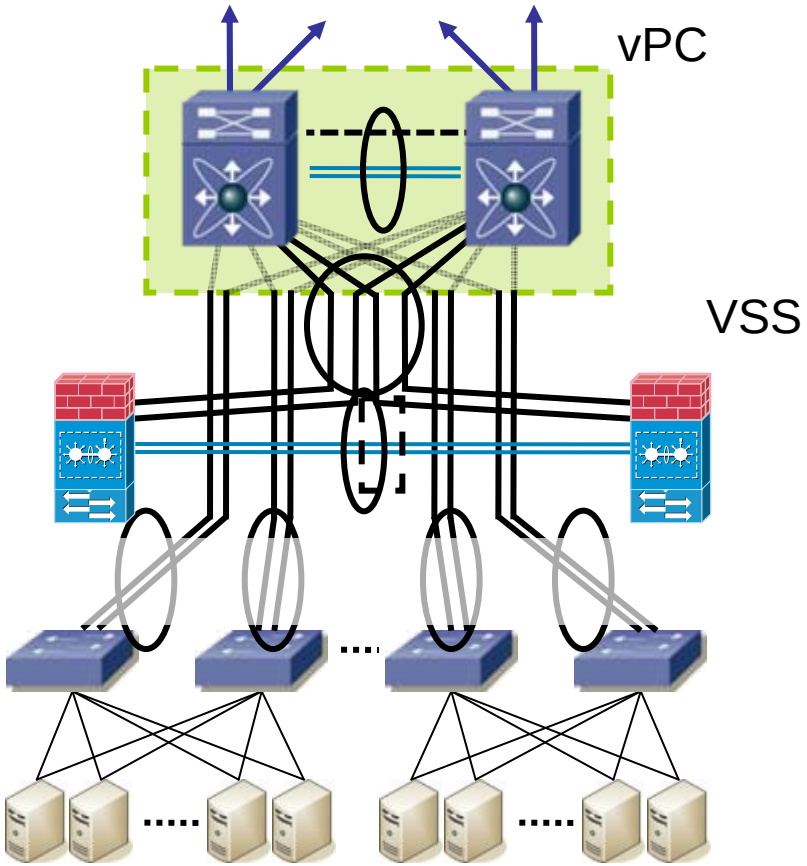

Cisco Expo
2009

Access Layer Virtualization

Welcome to the Human Network.

cisco

Data Center Access Layer Options

Top of Rack (ToR)

- Typically 1-RU servers
- 1-2 GE LOMs
- Mostly 1, sometimes 2 ToR switches
- Copper cabling stays within rack
- Low copper density in ToR
- Higher chance of *East-West* traffic hitting aggregation layer
- Drives higher STP logical port count for aggregation layer
- Denser server count

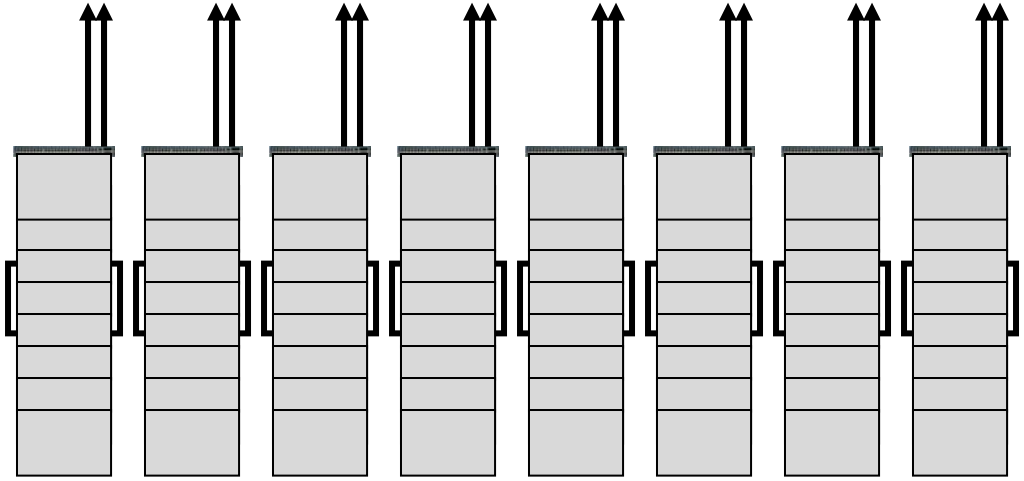

Middle of Row (MoR) (or End of Row)

- May be 1-RU or multi-RU servers
- Multiple GE or 10GE NICs
- Horizontal copper cabling for servers
- High copper cable density in MoR
- Larger portion of East-West traffic stays in access
- Larger subnets → less address waste
- Keeps agg. STP logical port count low (more EtherChannels, fewer trunk ports)
- Lower # of network devices to manage

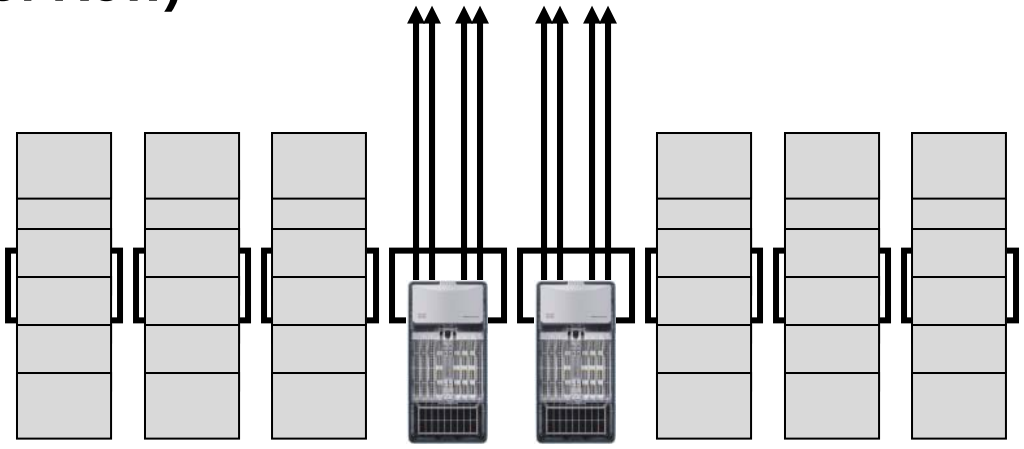

Middle of Row (MoR) (or End of Row) Virtual Switch (Nexus 7000 or Catalyst 6500)

Catalyst 6500

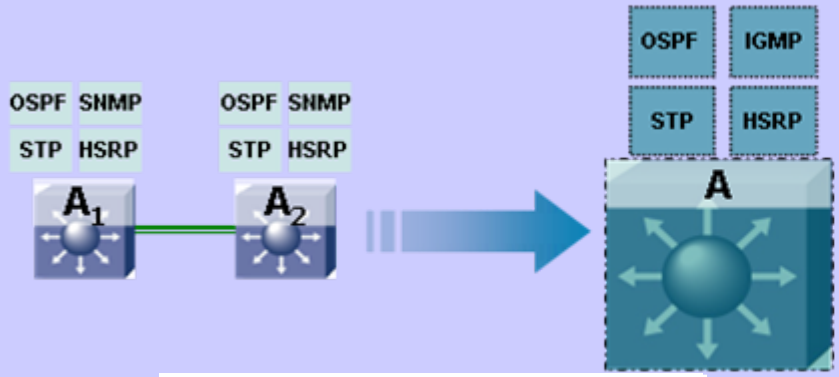

VSS and MEC

Many to 1 Virtualization
Service Modules
Single Control Plane

Nexus 7000

VDC and vPC

1 to Many Virtualization
High Density (10/100/1000 & 10GE)
Distinct control planes while virtualized

ToR @ 1GE:
Nexus 2000, the Nexus 5000 “virtual” linecard

- Nexus 2000 combines benefits of both ToR and EoR architectures
  - Physically resides on the top of each rack but
  - Logically acts like an end of row access device
- Nexus 2000 deployment benefits
  - Reduces cable runs
  - Reduce management points
  - Ensures feature consistency across hundreds of servers
  - Enable Nexus 5000 to become a high density 1GE access layer switch
  - VN-Link capabilities

Nexus 2000 (Fabric Extender - FEX)

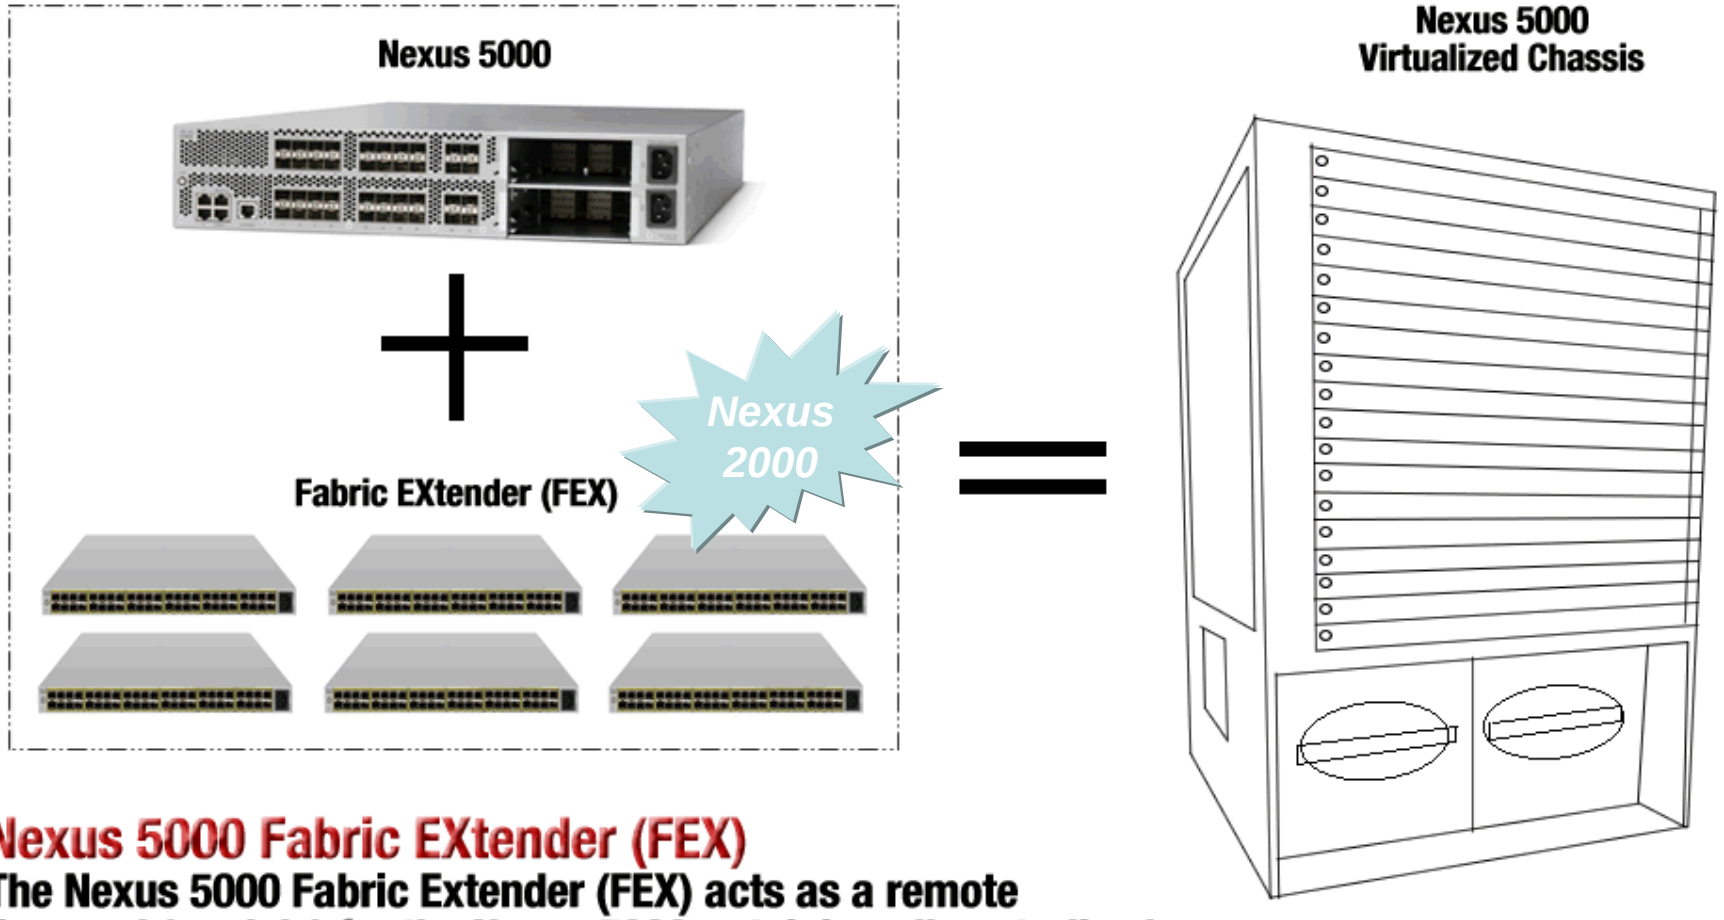

Nexus 5000 Fabric EXtender (FEX)

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

Nexus 2000 implementation example

•Physical Topology

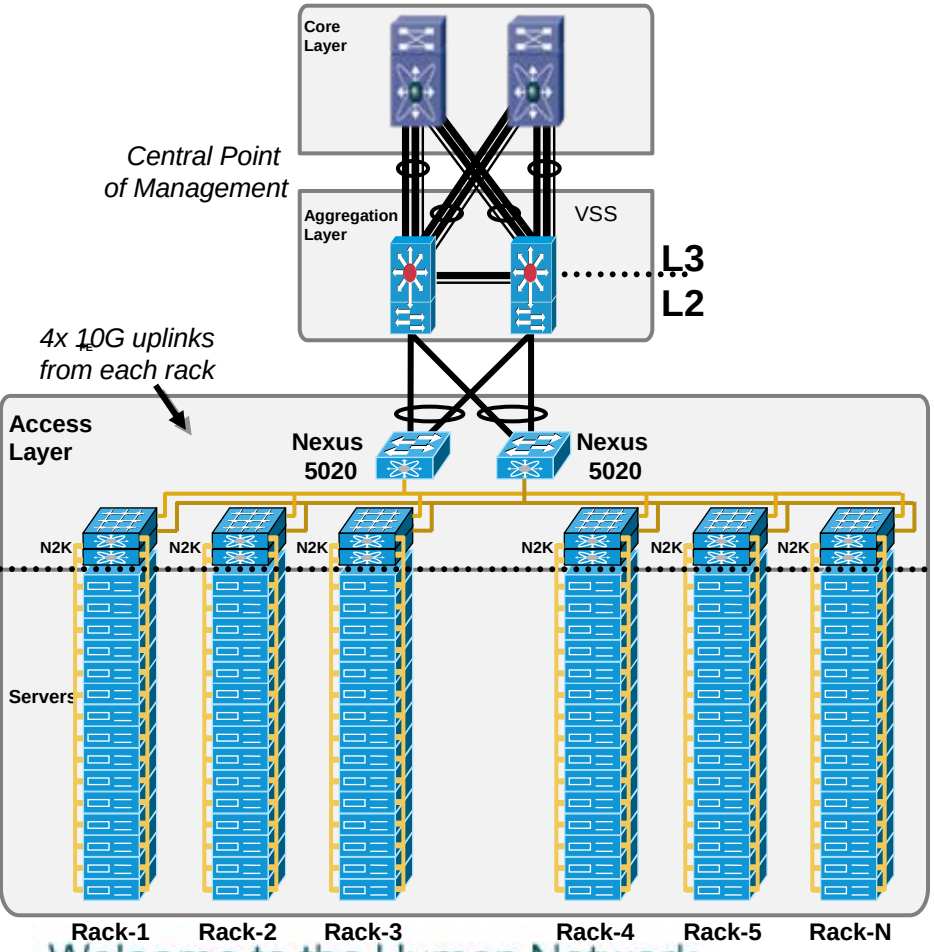

•Logical Topology

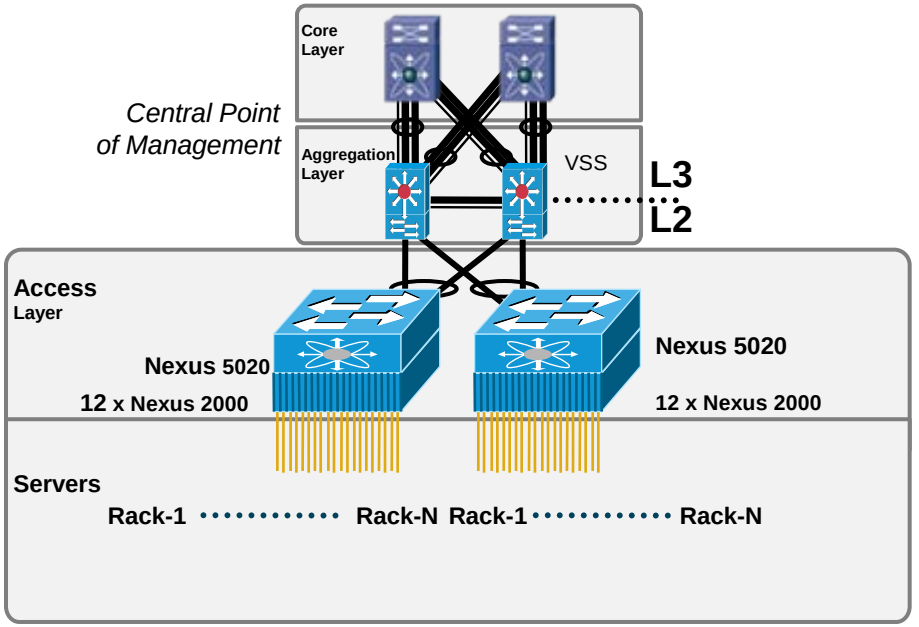

Blades: Cisco Virtual Blade Switching (VBS)

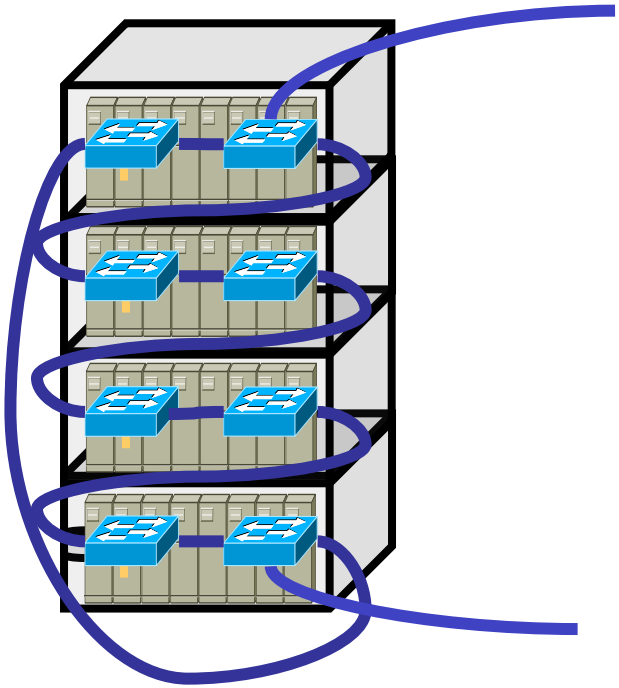

- ***Up to 8 Switches acts as Single VBS Switch***
  - ***–Distributed L2/ MAC learning***
  - ***–Centralized L3 learning***
- ***Each switch consists of***
  - ***– Switch Fabric***
  - ***– Port Asics (downlink & uplink ports)***
  - ***–***
- ***One Master Switch per VBS***
  - ***–1:N Resiliency for Master***
  - ***– L2/L3 reconvergence is sub 200 msec***
- ***High Speed VBS Cable (64 Gbps)***
- **Example Deployment:**
  - **–16 servers per enclosure X**
  - **– 2 GE ports per server X**
  - **– 4 enclosures per rack = 128GE**
  - **– 2 x 10GE uplinks = 20GE**
  - **– 128GE / 20GE = 6.4:1 oversubscription**

Cisco Catalyst Virtual Blade Switch (VBS)
with Non-vPC Aggregation

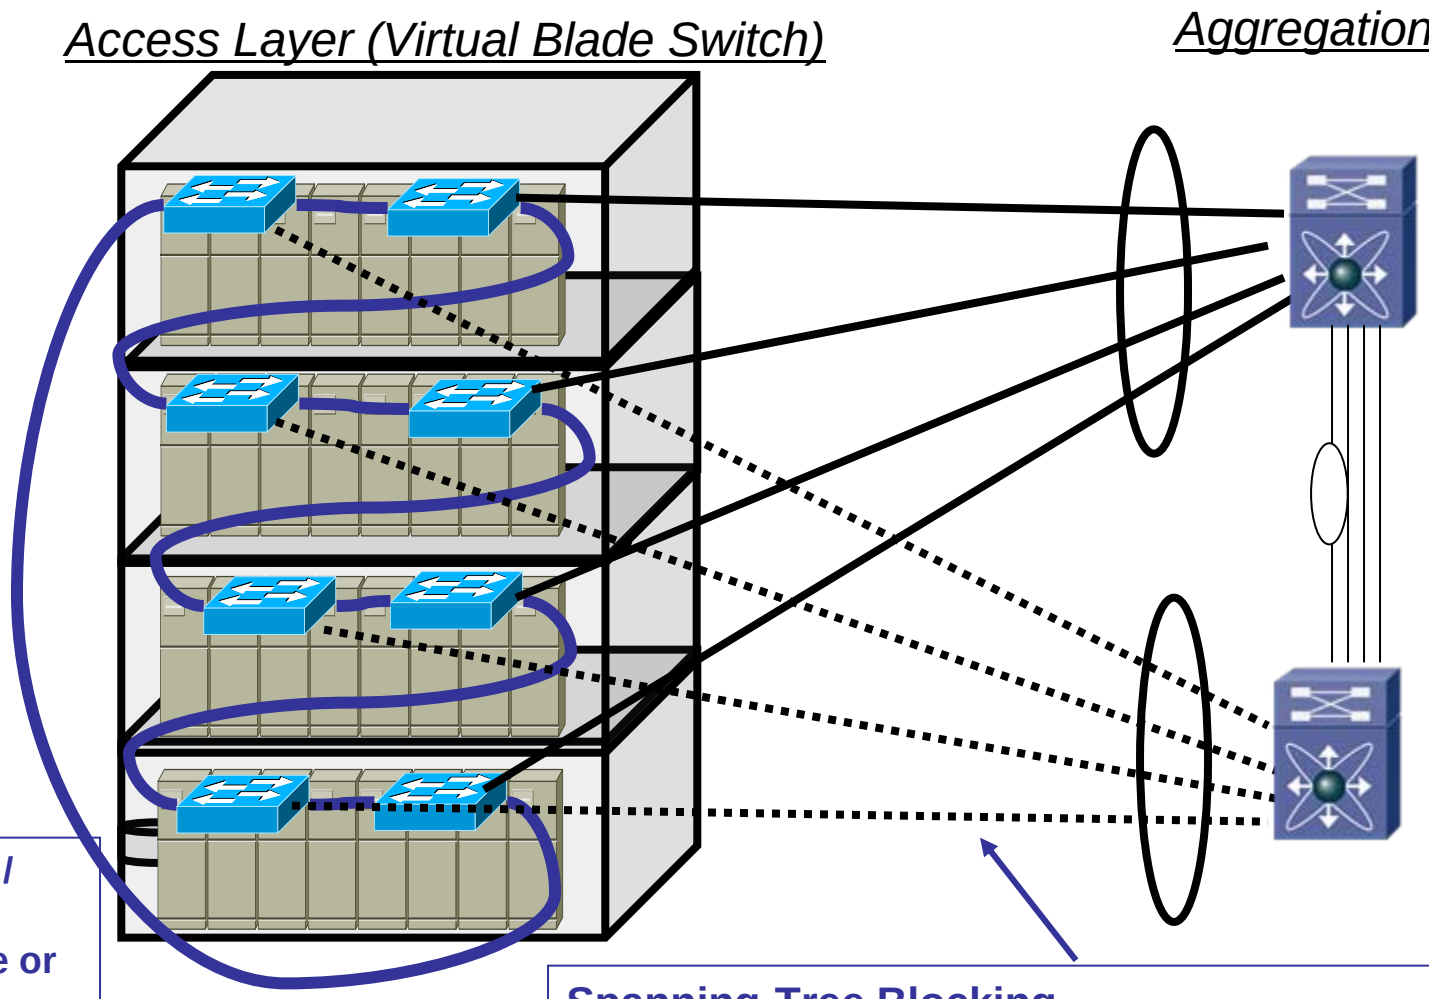

Cisco Catalyst Virtual Blade Switch (VBS)
with Non-vPC Aggregation

Aggregation Layer

Access Layer (Virtual Blade Switch)

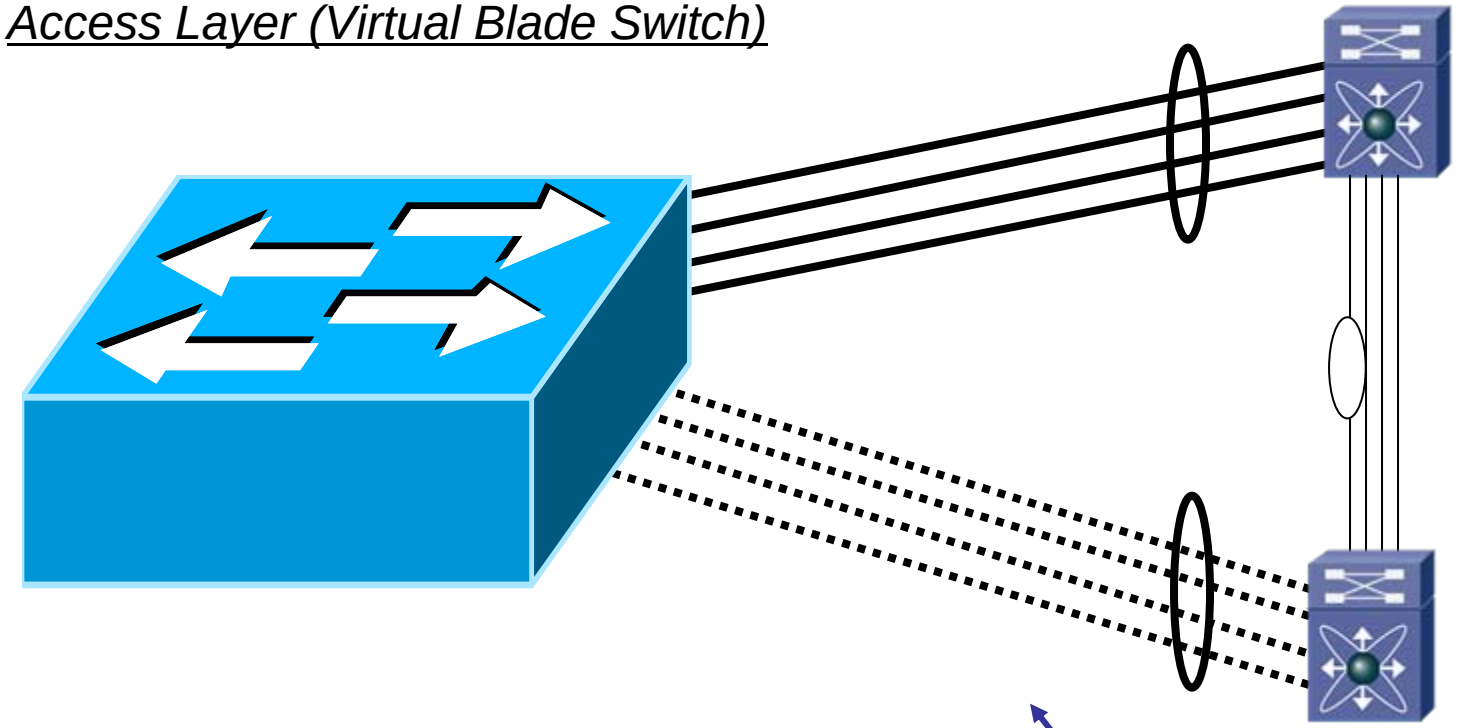

Single Switch / Node (for Spanning Tree or Layer 3 or Management)

Spanning-Tree Blocking

Cisco Expo 2009

Cisco Catalyst Virtual Blade Switch (VBS)
with Nexus vPC Aggregation

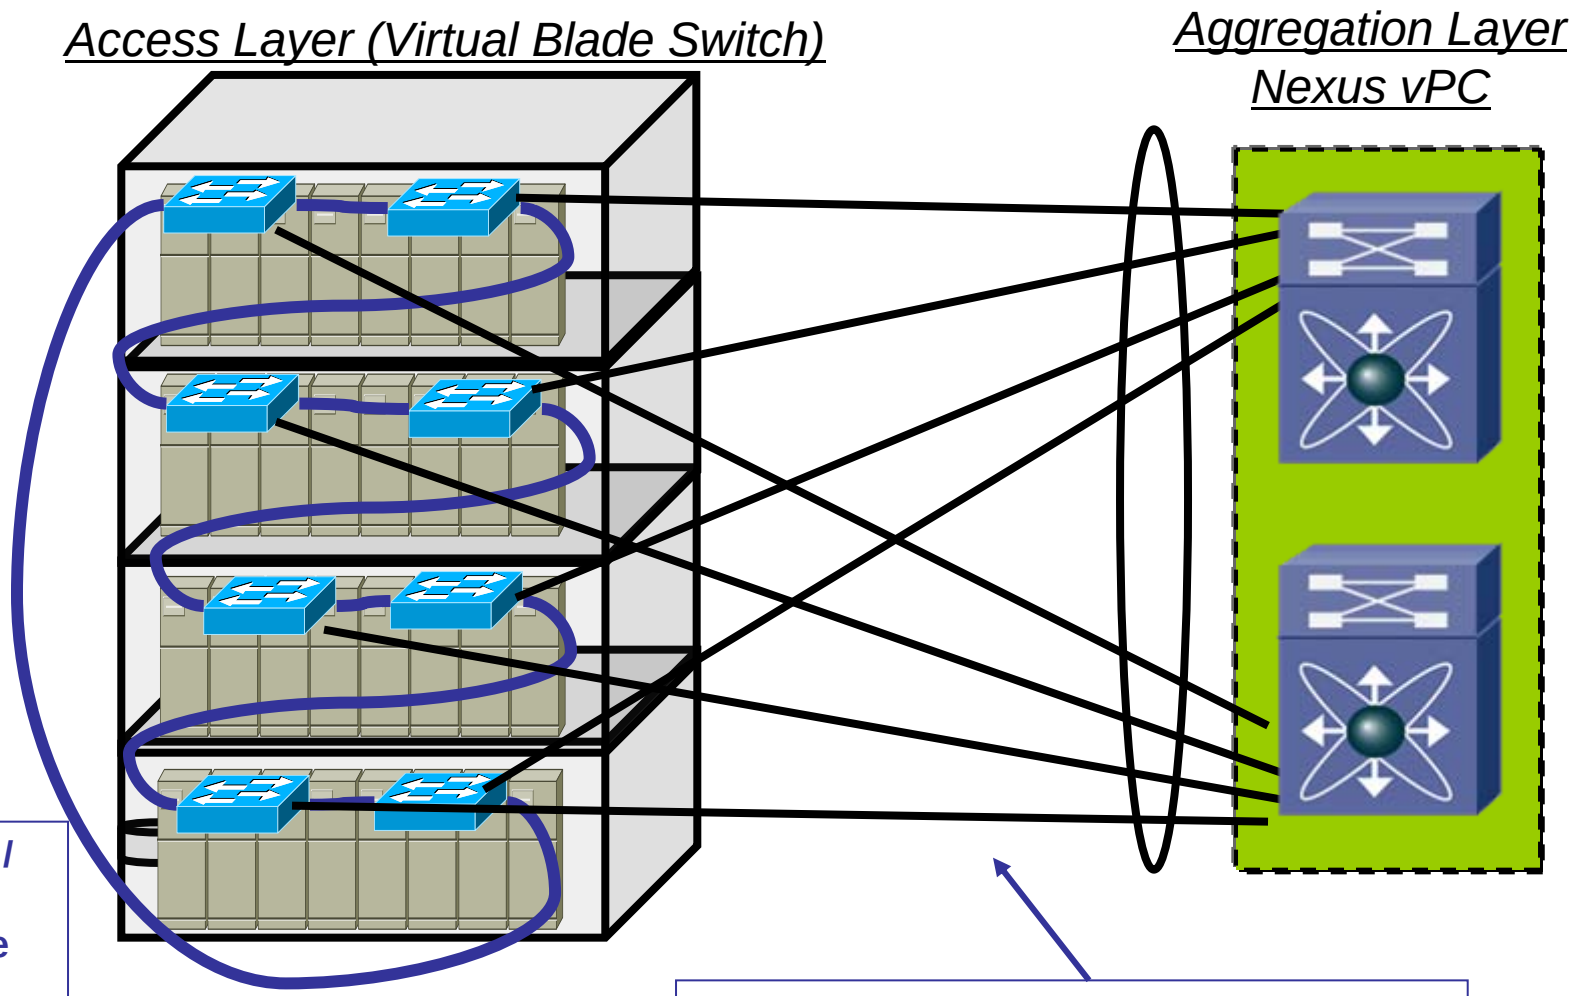

CISCO

Cisco Catalyst Virtual Blade Switch (VBS)
with Nexus vPC Aggregation

Aggregation Layer (Nexus vPC)

Access Layer (Virtual Blade Switch)

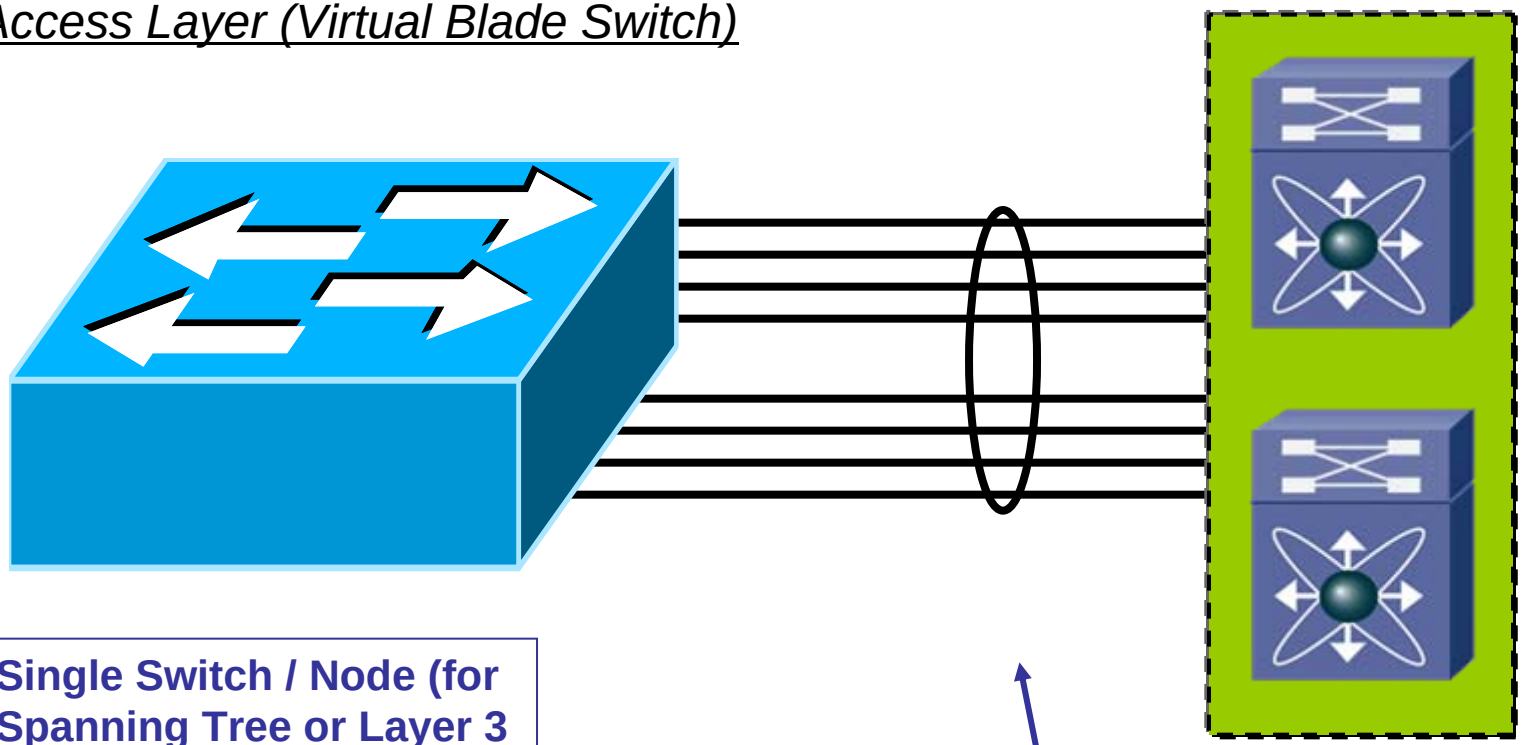

Single Switch / Node (for Spanning Tree or Layer 3 or Management)

All Links Forwarding

Server Virtualization

VMware ESX 3.x Networking Components

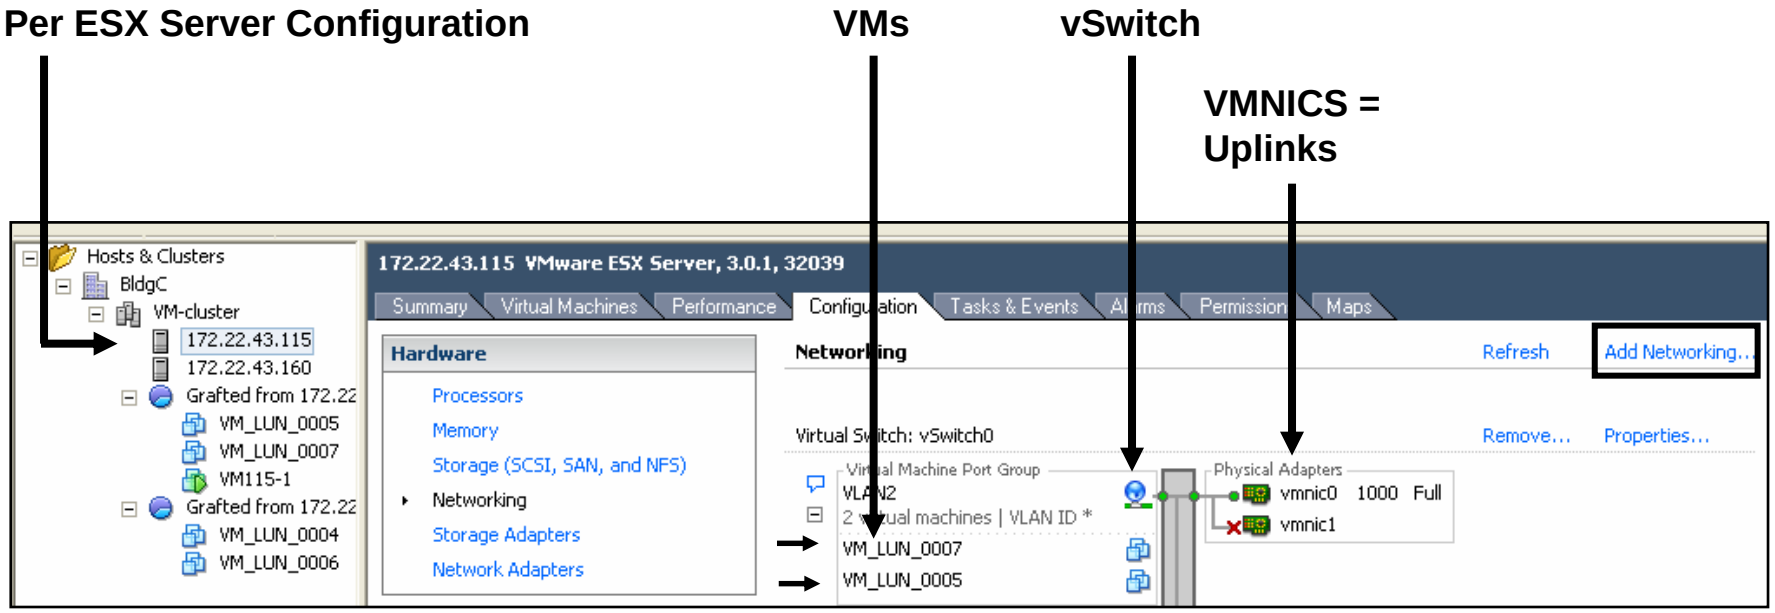

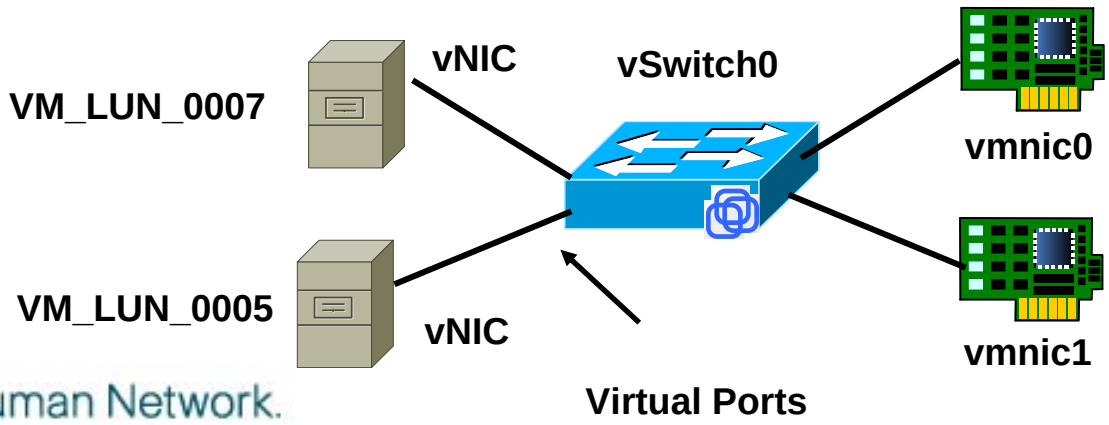

Cisco VN-Link

- VN-Link (or Virtual Network Link) is a term which describes a new set of features and capabilities that enable VM interfaces to be individually identified, configured, monitored, migrated and diagnosed.
  - The term literally refers to a VM specific link that is created between the VM and Cisco switch. It is the logical equivalent & combination of a NIC, a Cisco switch interface and the RJ-45 patch cable that hooks them together.

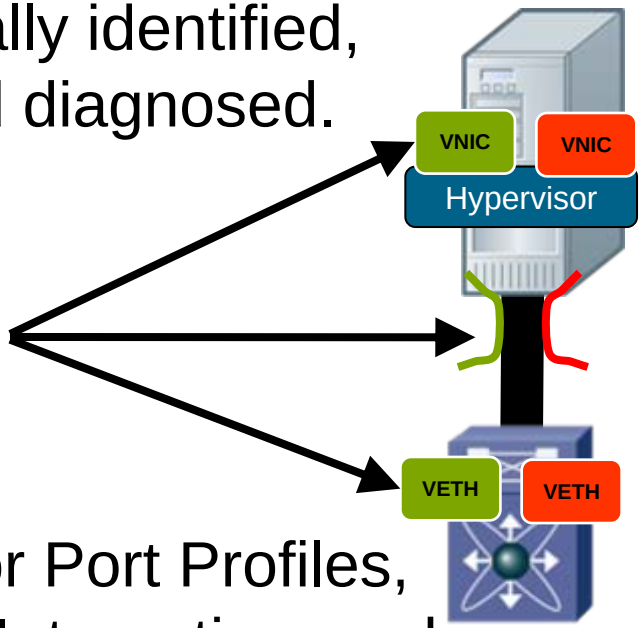

- VN-Link requires platform support for Port Profiles, Virtual Ethernet Interfaces, vCenter Integration, and Virtual Ethernet mobility.

Server Virtualization & VN-Link

VN-Link Brings VM Level Granularity

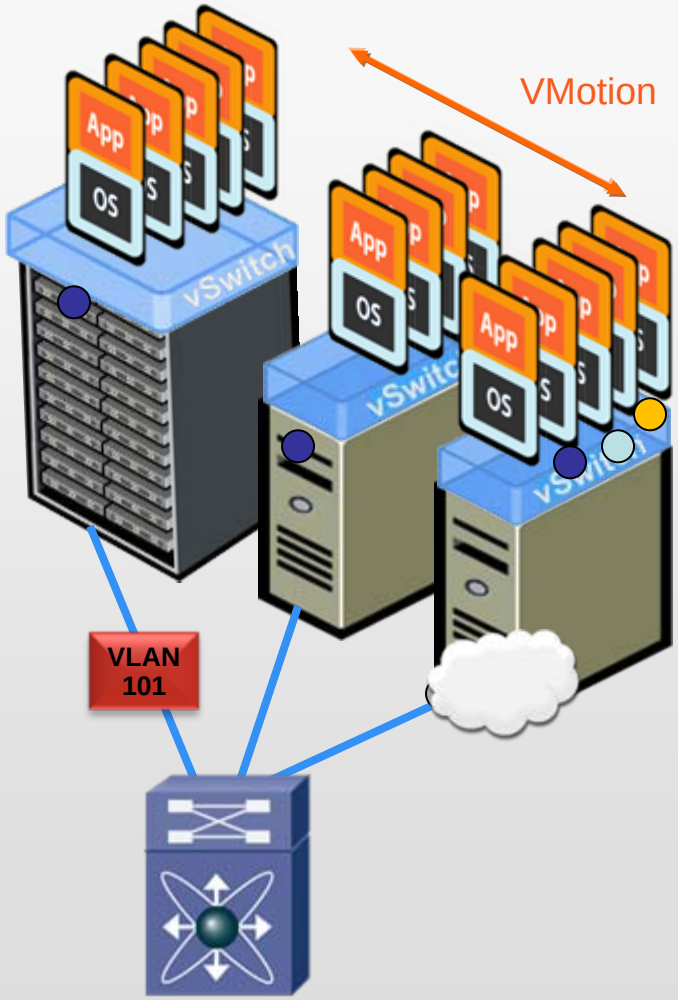

Problems:

- VMotion may move VMs across physical ports—policy must follow
- Impossible to view or apply policy to locally switched traffic
- Cannot correlate traffic on physical links—from multiple VMs

VN-Link:

- Extends network to the VM
- Consistent services
- Coordinated, coherent management

VN-Link With the Cisco Nexus 1000V

Cisco Nexus 1000V
Software Based

- Industry's first third-party ESX switch
- Built on Cisco NX-OS
- Compatible with switching platforms
- Maintain vCenter provisioning model unmodified for server administration but also allow network administration of Nexus 1000V via familiar Cisco NX-OS CLI

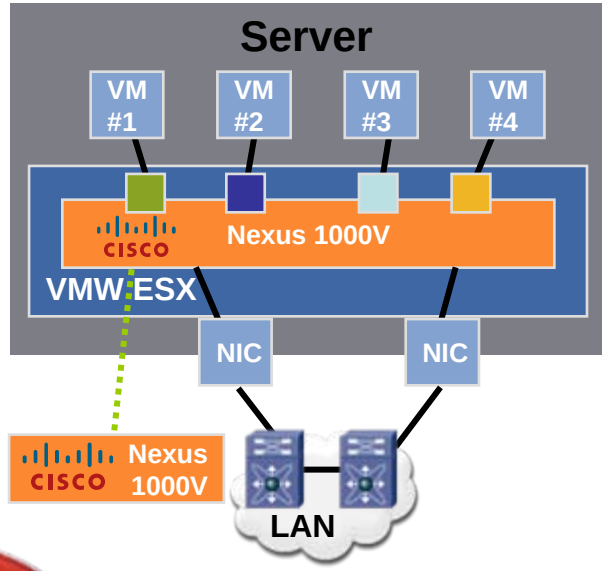

Announced 09/2008 Shipping H1CY09)

Policy-Based VM Connectivity

Mobility of Network and Security Properties

Non-Disruptive Operational Model

VN-Link with Network Interface Virtualization (NIV)

Nexus Switch with VN-Link
Hardware Based

- Allows scalable hardware-based implementations through hardware switches
- Standards-based initiative: Cisco & VMware proposal in IEEE 802 to specify "Network Interface Virtualization"
- Combines VM and physical network operations into one managed node
- Future availability

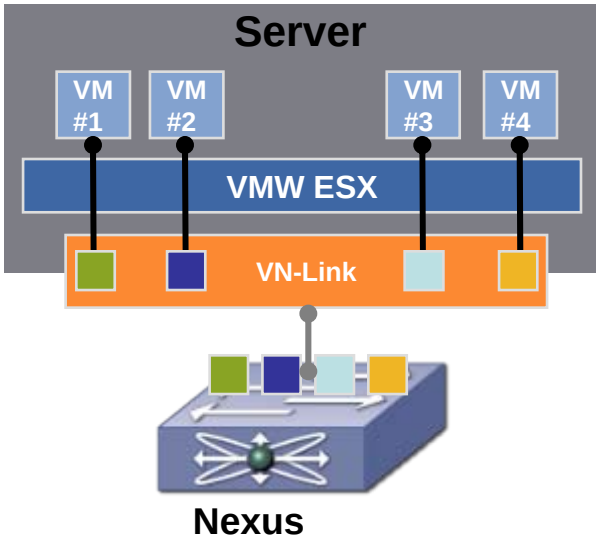

IEEE CISCO vmware

http://www.ieee802.org/1/files/public/docs2008/new-dcb-pelissier-NIC-Virtualization-0908.pdf

Policy-Based VM Connectivity

Mobility of Network and Security Properties

Non-Disruptive Operational Model

Cisco Nexus 1000V

Industry First 3rd Party Distributed Virtual Switch

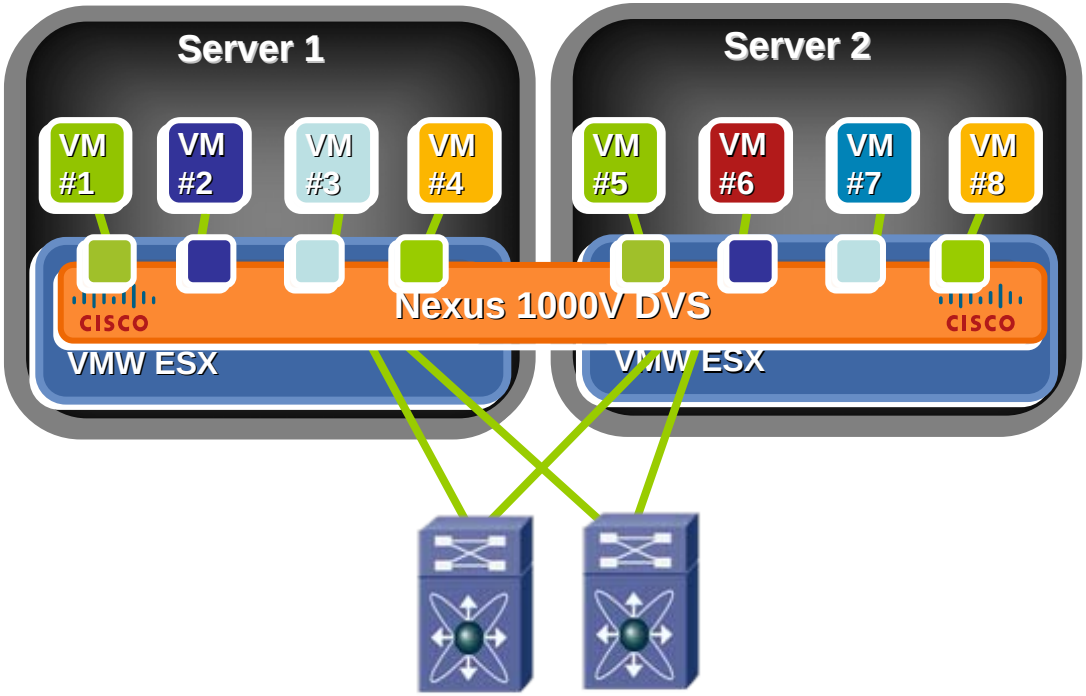

- **Nexus 1000V provides enhanced VM switching for VMware ESX**
- **Features Cisco VN-Link:**
  - Policy Based VM Connectivity
  - Mobility of Network & Security Properties
  - Non-Disruptive Operational Model
- **Ensures proper visibility & connectivity during VMotion**

Enabling Acceleration of Server Virtualization Benefits

Cisco Nexus 1000V Architecture

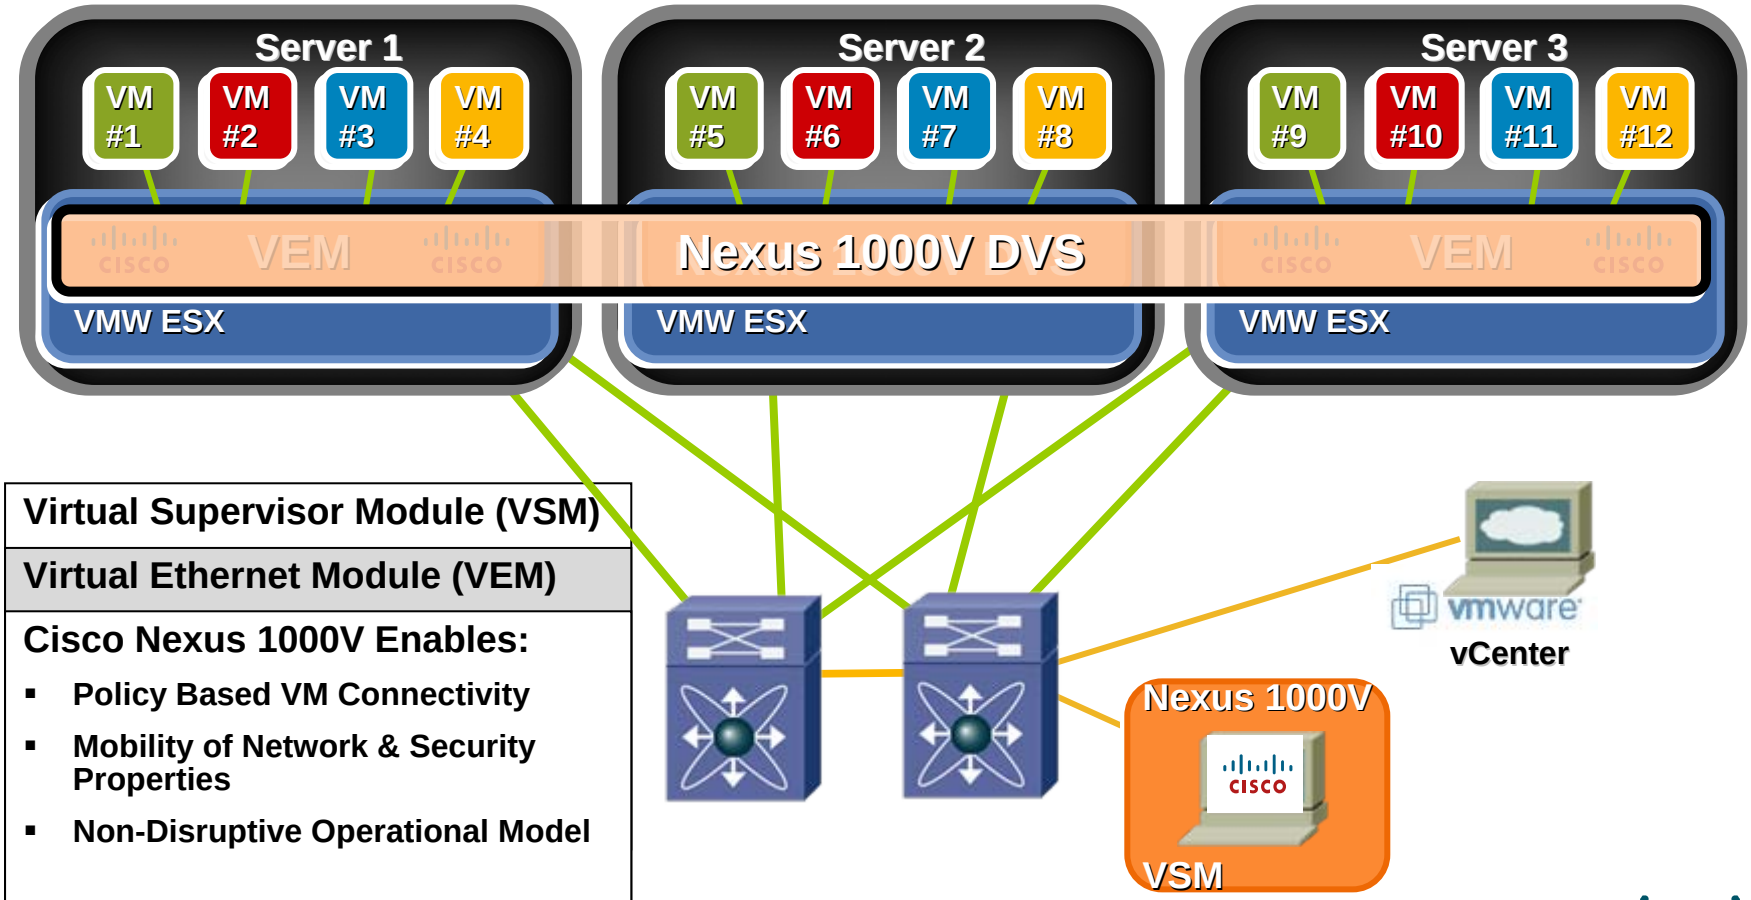

Virtual Supervisor Module (VSM)

Virtual Ethernet Module (VEM)

Cisco Nexus 1000V Enables:

- **Policy Based VM Connectivity**
- **Mobility of Network & Security Properties**
- **Non-Disruptive Operational Model**

Back-End Virtualization

End-to-End Back-End Virtualization

- Optimizes resource utilization
- Increases flexibility and agility
- Simplifies management
- Reduces TCO

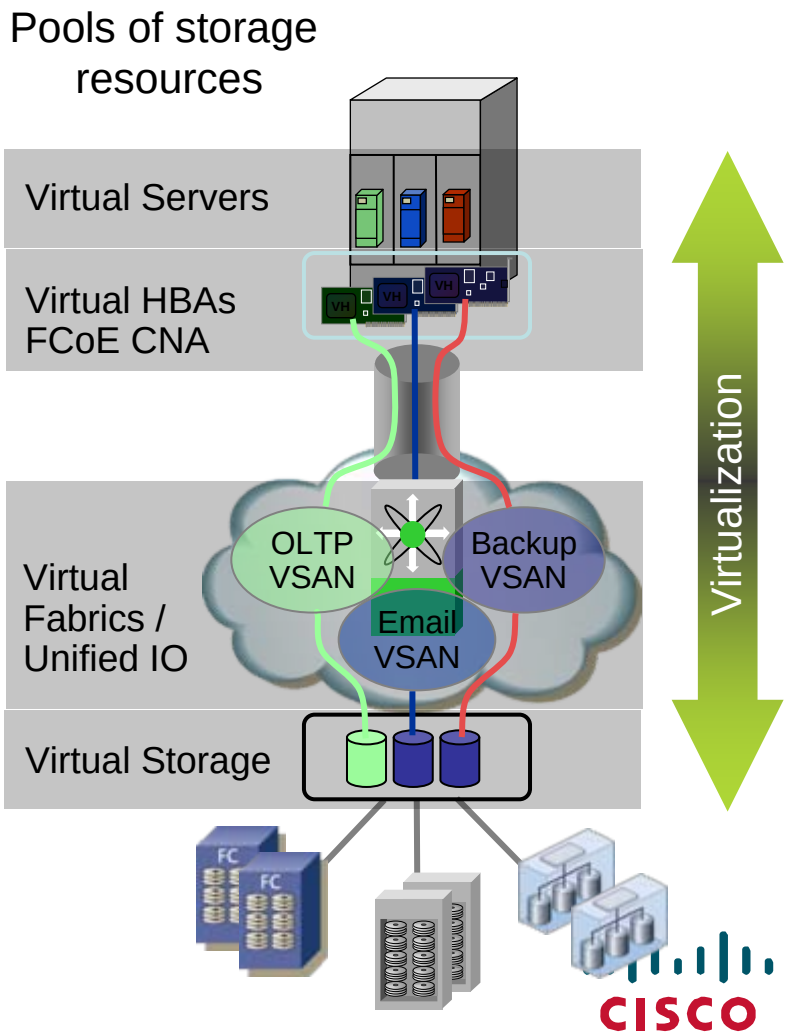

Virtual Storage Area Network (VSAN) Deployment

- 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)

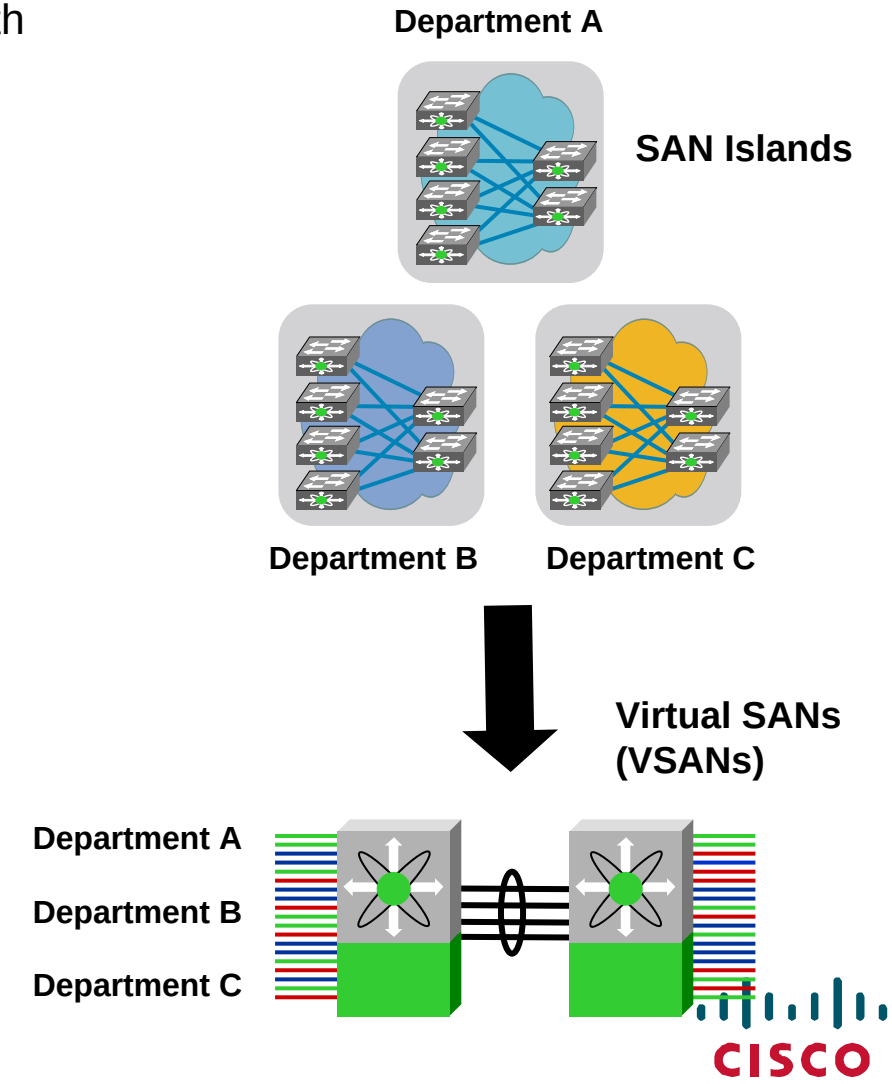

VSAN Technology

The Virtual SANs Feature Consists of Two Primary Functions

- Hardware-based isolation of tagged traffic belonging to different VSANs
- Create independent instance of fiber channel services for each newly created VSAN—services include:

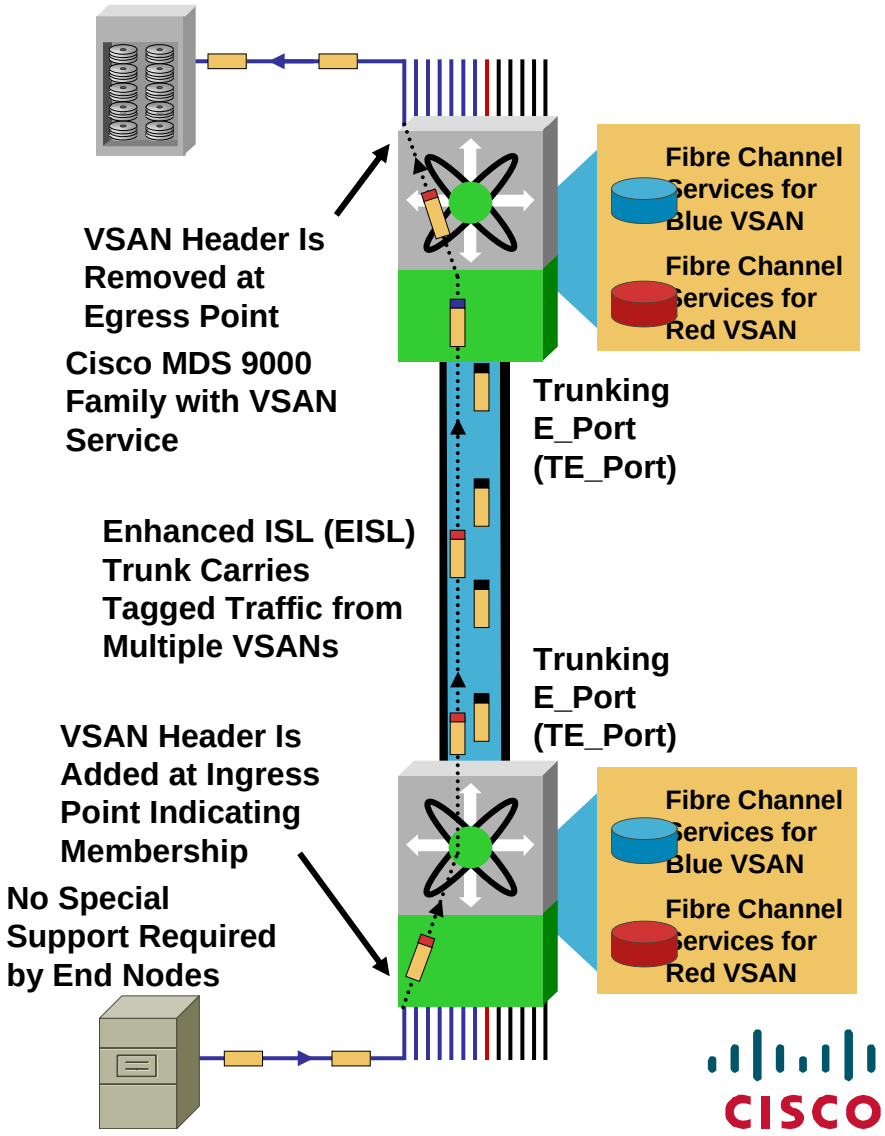

N-Port ID Virtualization (NPIV)

- Mechanism to assign multiple N_Port_IDs to a single N_Port
- Allows all the access control, zoning, port security (PSM) be implemented on application level
- Multiple N_Port_IDs are so far allocated in the same VSAN

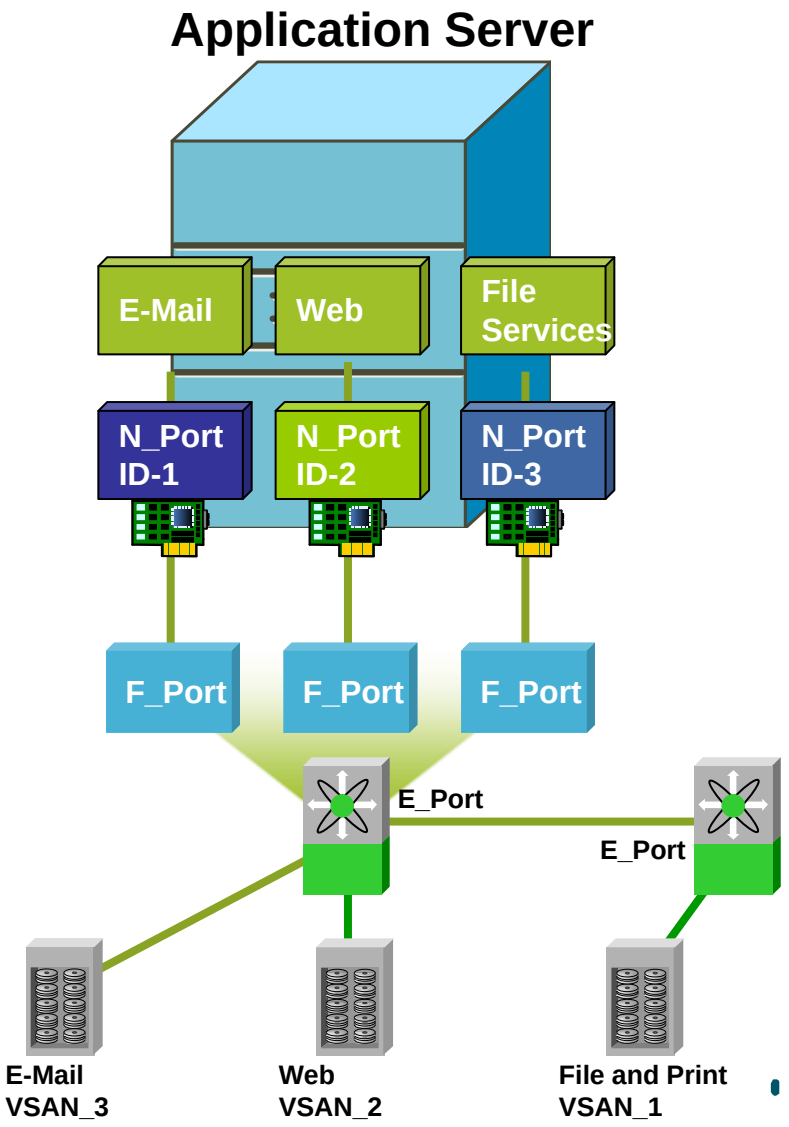

NPIV Usage Examples

Virtual Machine Aggregation

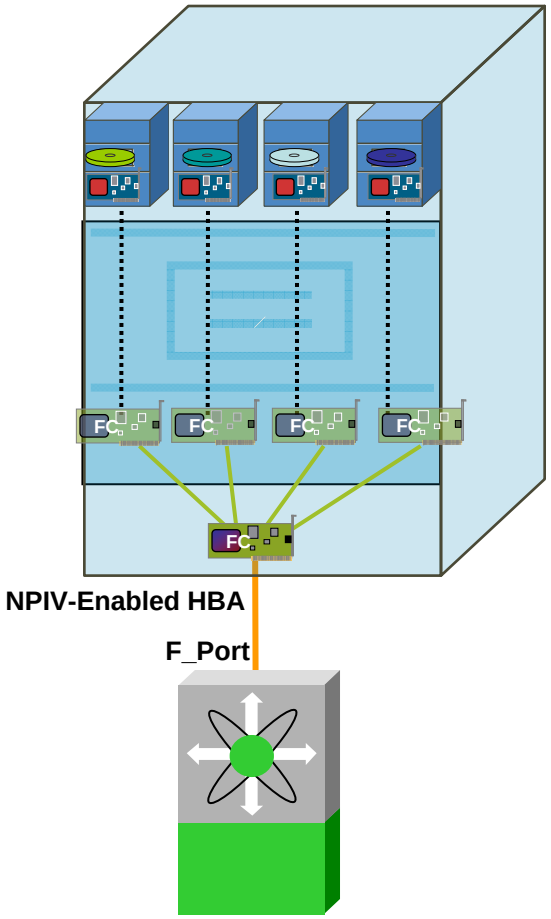

‘Intelligent Pass-Thru’

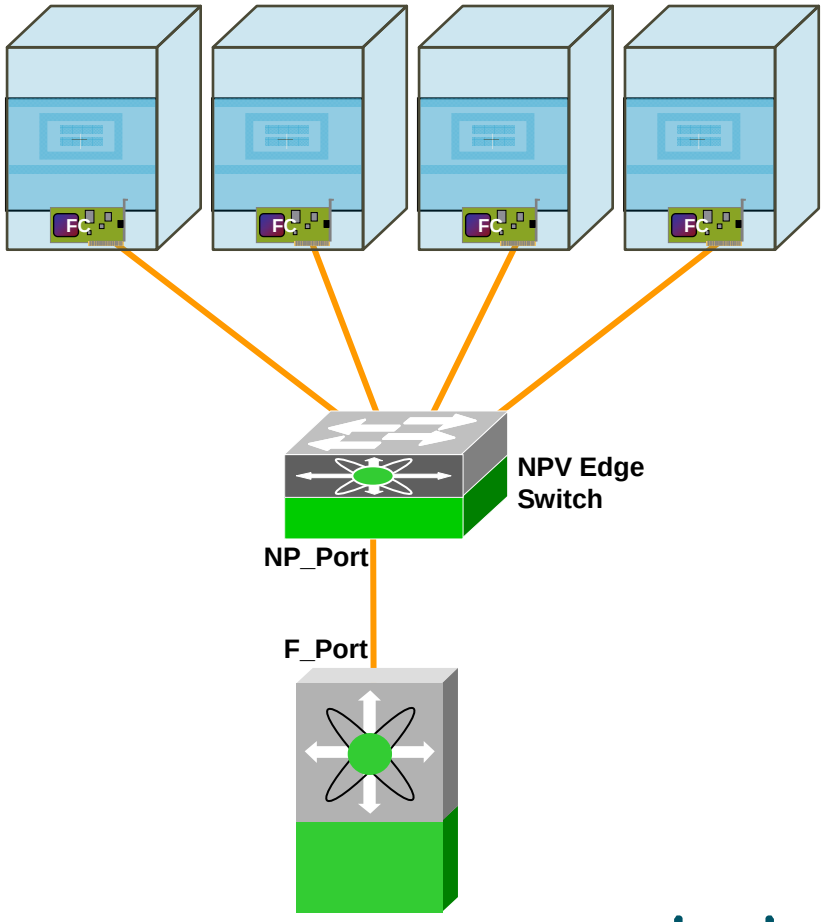

Virtual Servers Share a Physical HBA

- A zone includes the physical HBA and the storage array
- Access control is demanded to storage array “LUN masking and mapping”, it is based on the physical HBA pWWN and it is the same for all VMs
- The hypervisor is in charge of the mapping, errors may be disastrous

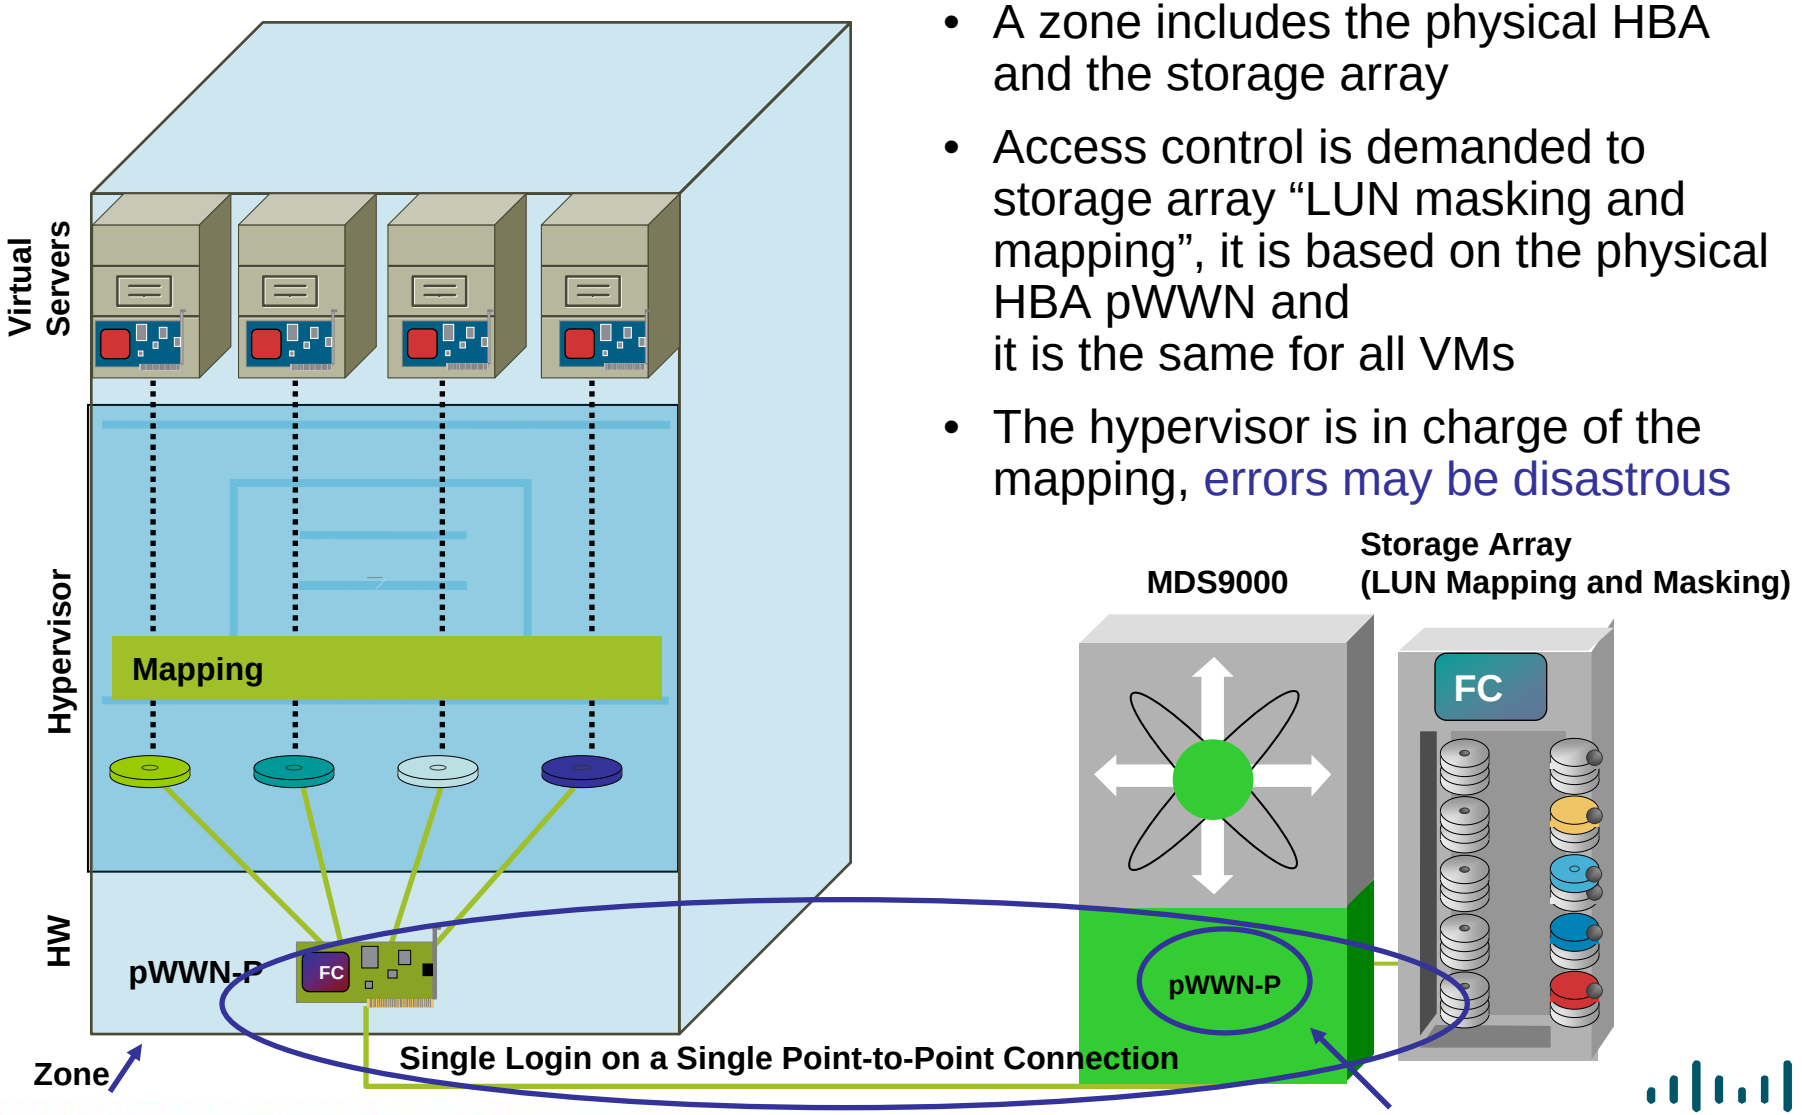

Virtual Server Using NPIV and Storage Device Mapping

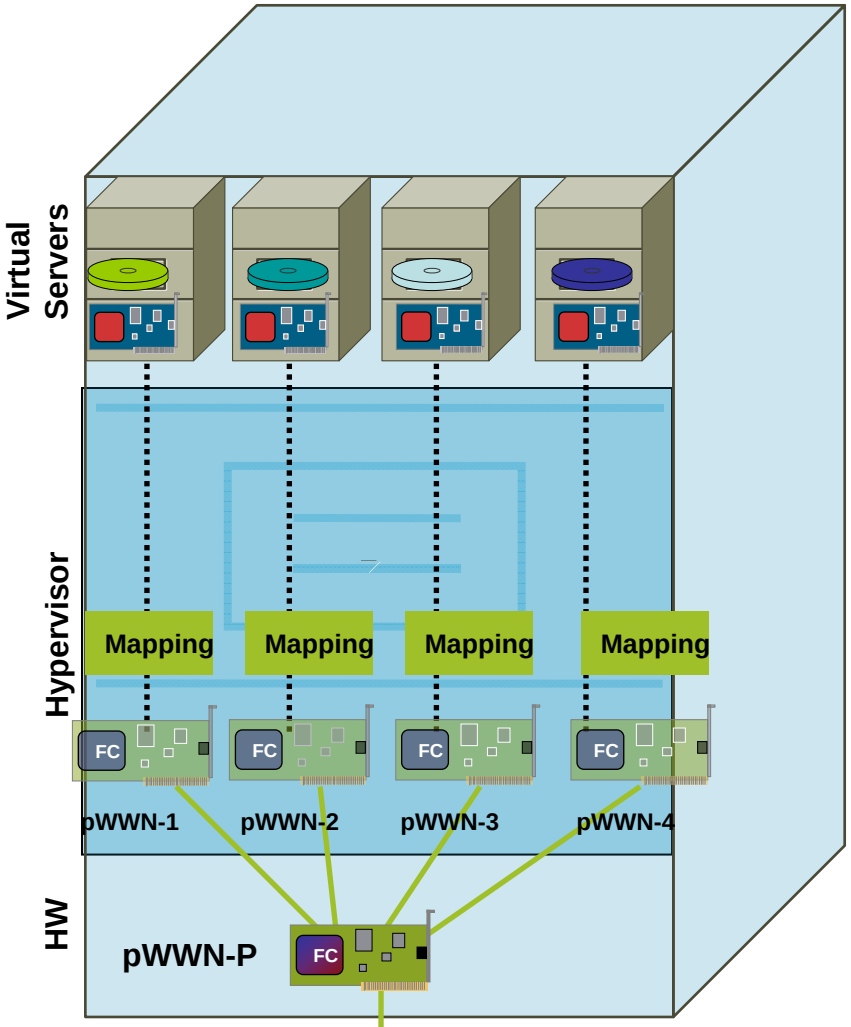

- Virtual HBAs can be zoned individually
- "LUN masking and mapping" is based on the virtual HBA pWWN of each VMs
- Very safe with respect to configuration errors
- Only supports RDM
- Available in ESX 3.5

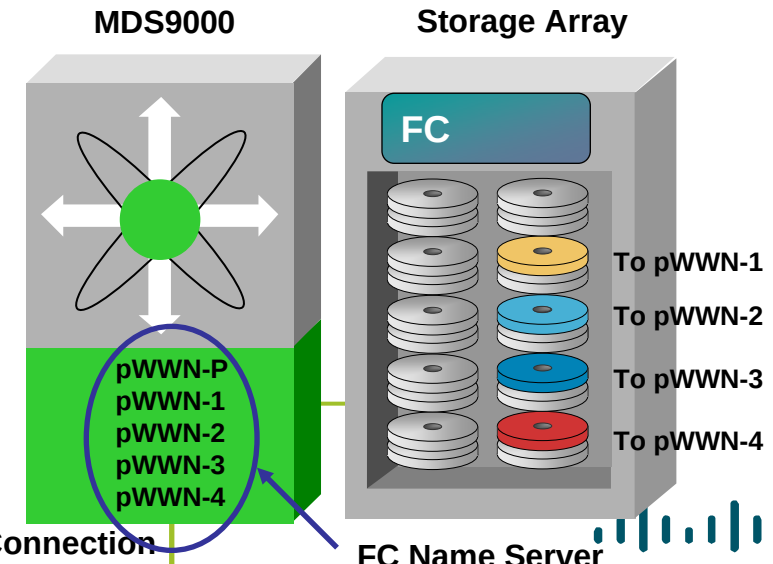

VMotion LUN Migration without NPIV

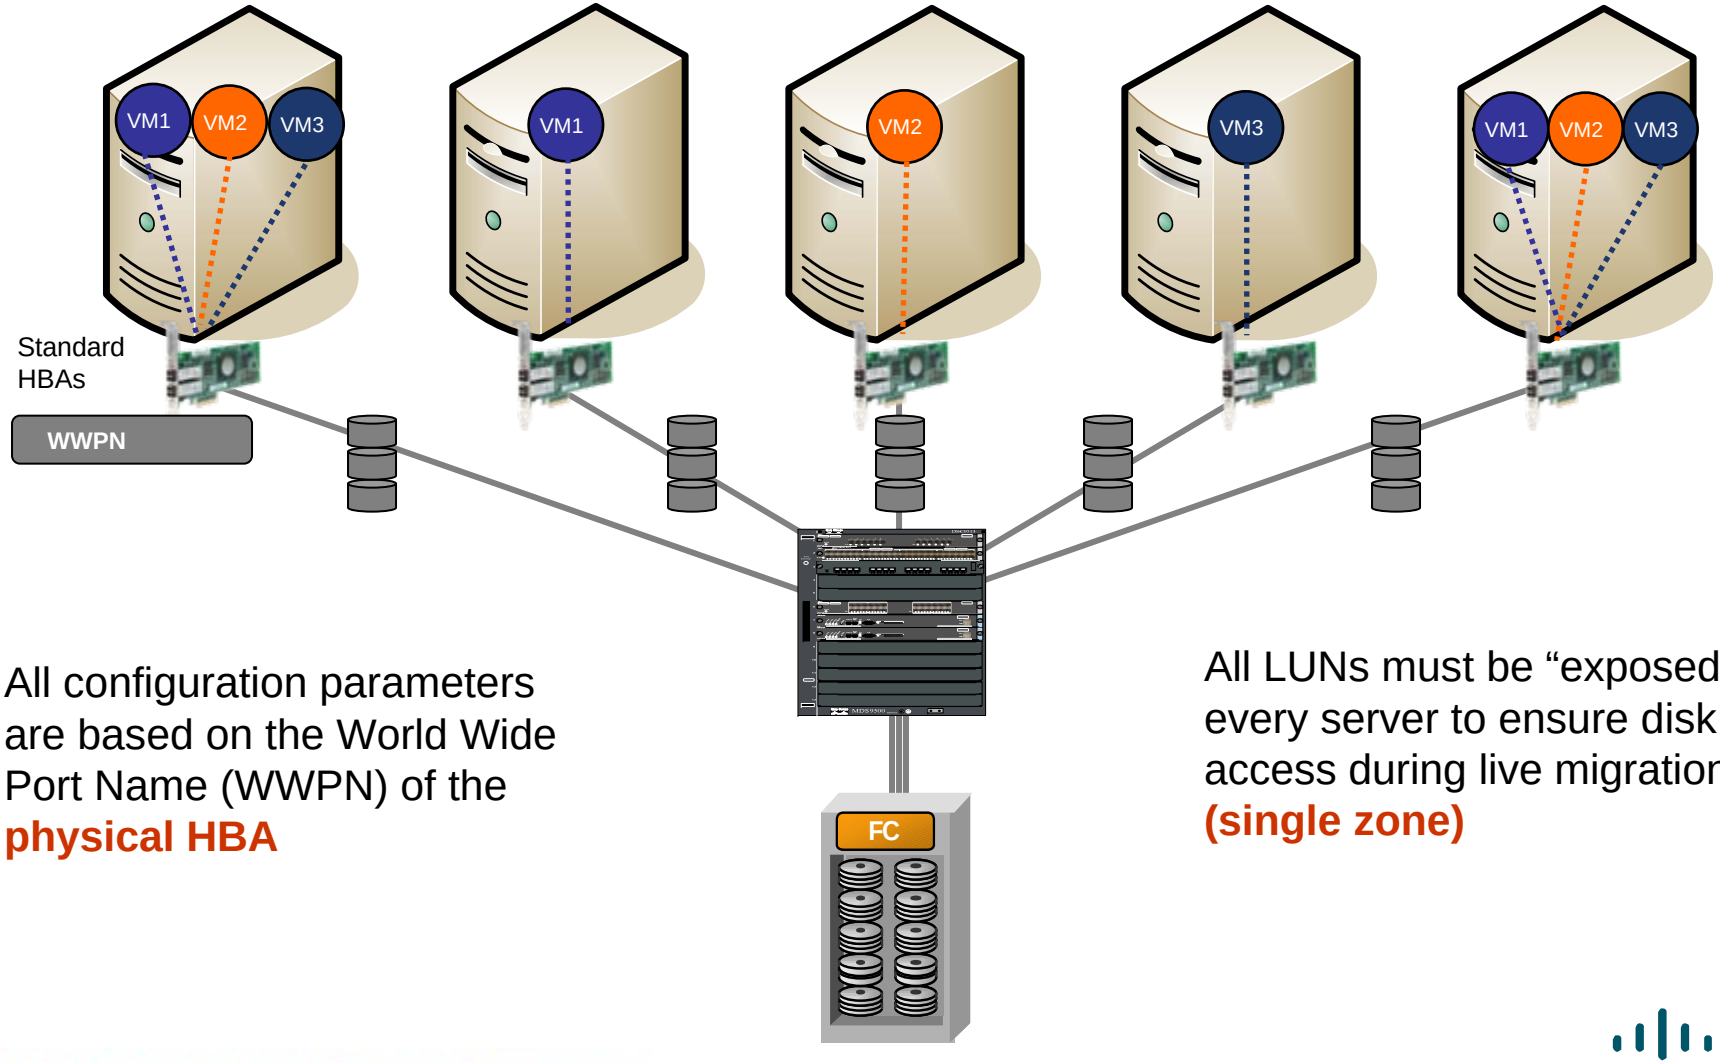

All configuration parameters are based on the World Wide Port Name (WWPN) of the **physical HBA**

All LUNs must be “exposed” to every server to ensure disk access during live migration **(single zone)**

VMotion LUN Migration with NPIV

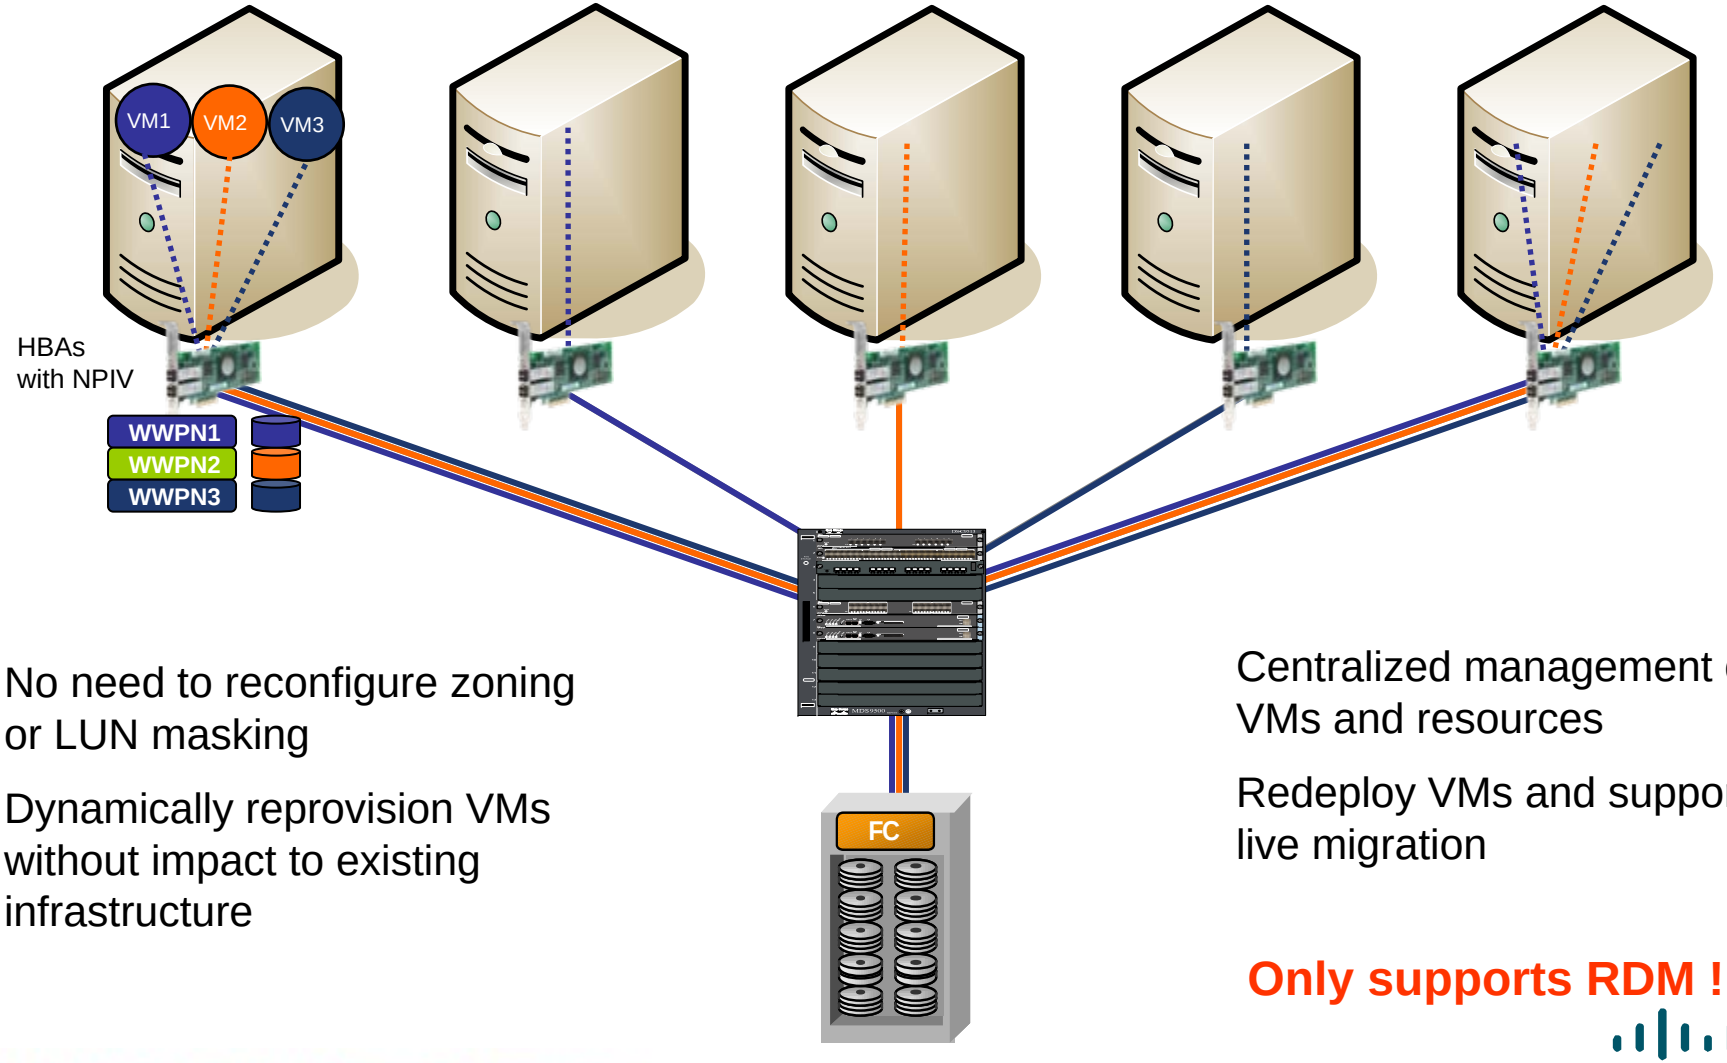

No need to reconfigure zoning or LUN masking

Dynamically reprovision VMs without impact to existing infrastructure

Centralized management of VMs and resources

Redeploy VMs and support live migration

Only supports RDM !

NPIV Usage Examples

Virtual Machine Aggregation

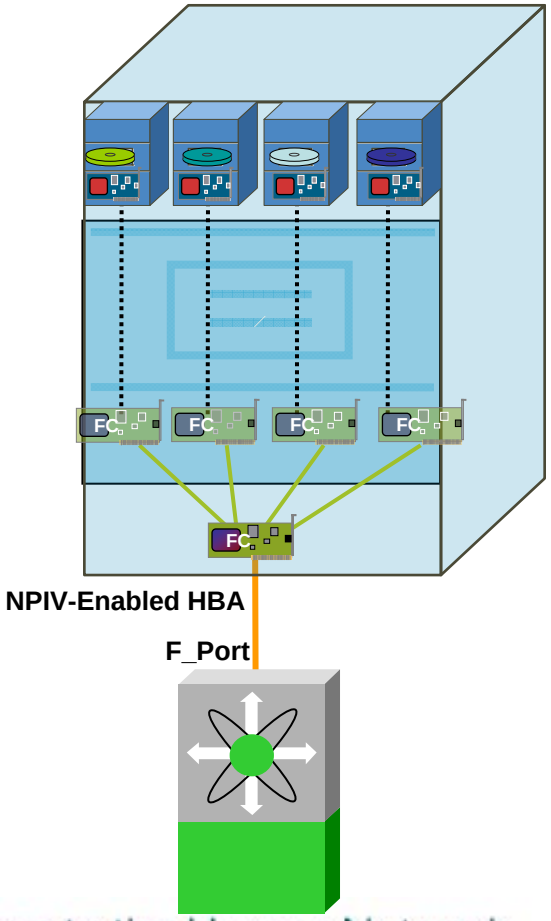

‘Intelligent Pass-Thru’

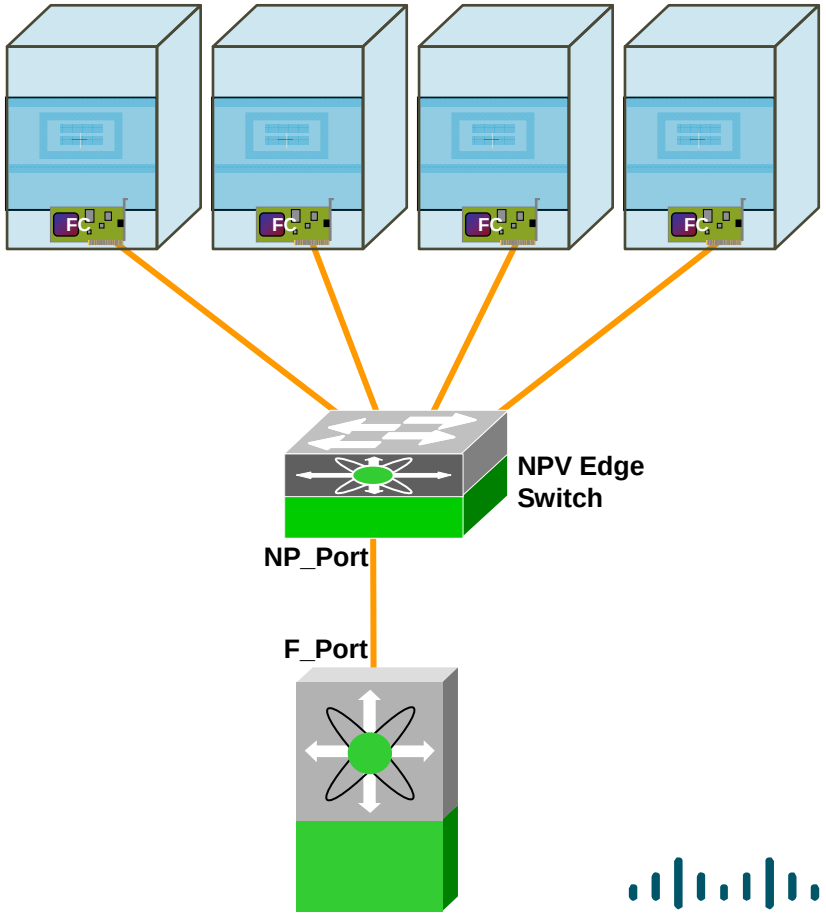

Blade Switch/Top-of-Rack Domain ID Explosion

- Domain ID used for addressing, routing, and access control
- One domain ID per SAN switch
- Theoretically 239 domain ID, practically much less supported
- Limits SAN fabric scalability

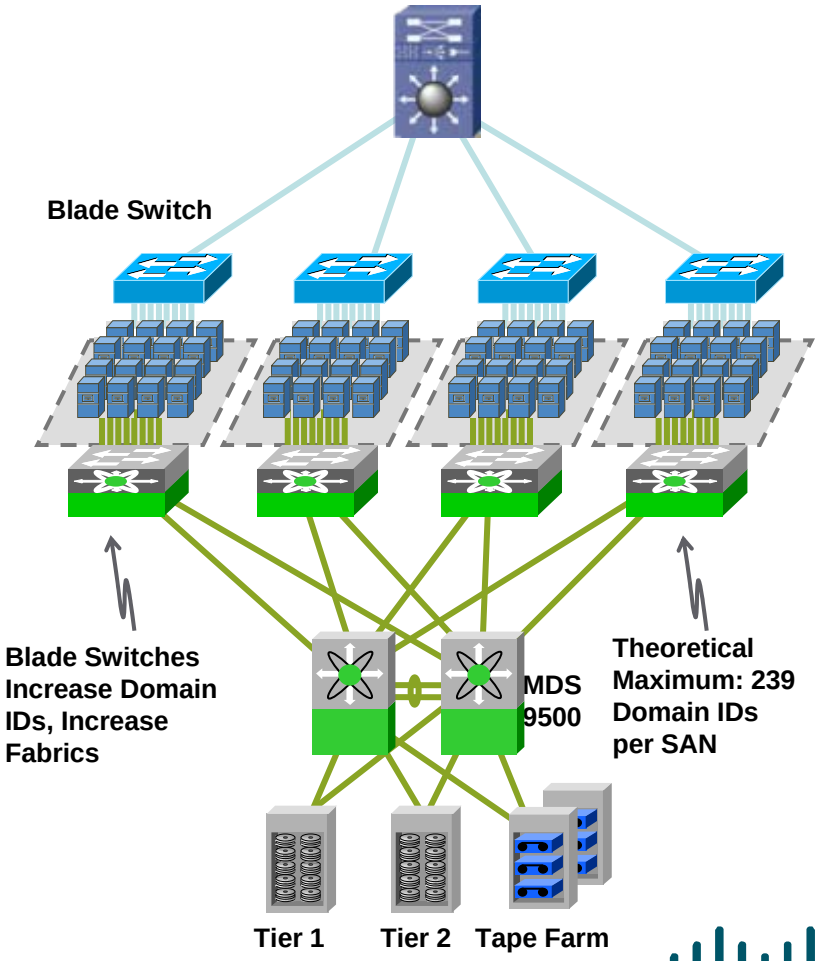

Cisco MDS Network Port Virtualization (NPV)

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

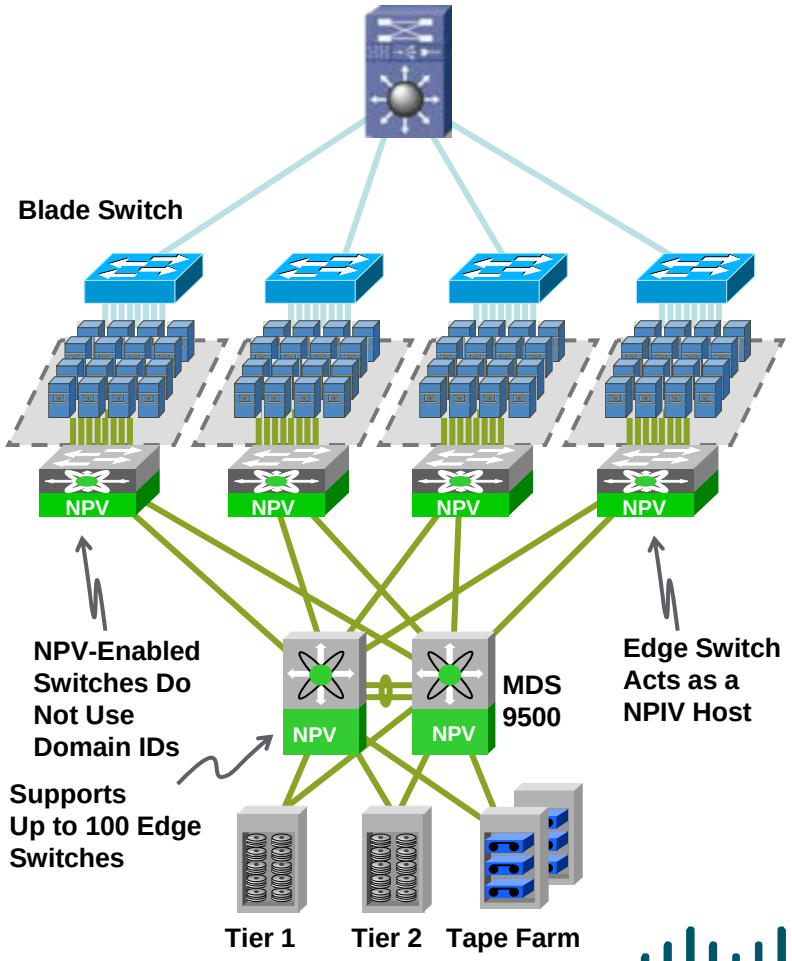

Flex Attach (Virtual PWWN)

- Assign virtual PWWN on NPV switch port
- Zone vPWWN to storage
- LUN masking is done on vPWWN
- Reduce operational overhead
  - –Enables server or physical HBA replacement
  - –No need for zoning modification
  - –No LUN masking change
- Automatic link to new PWWN
  - –No manual relinking to new PWWN is needed

Before

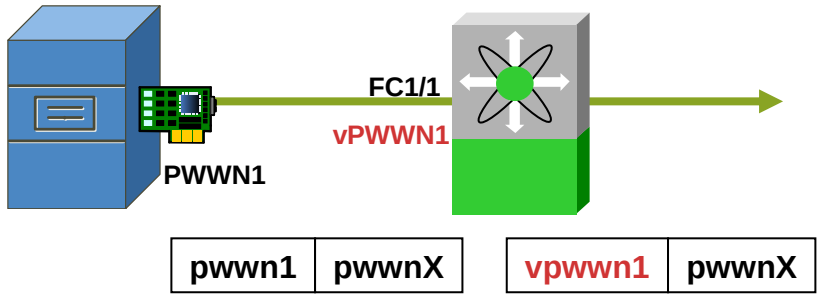

After

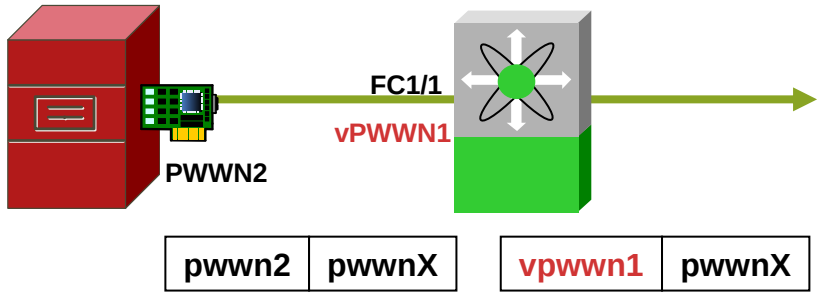

Storage Volume Virtualization

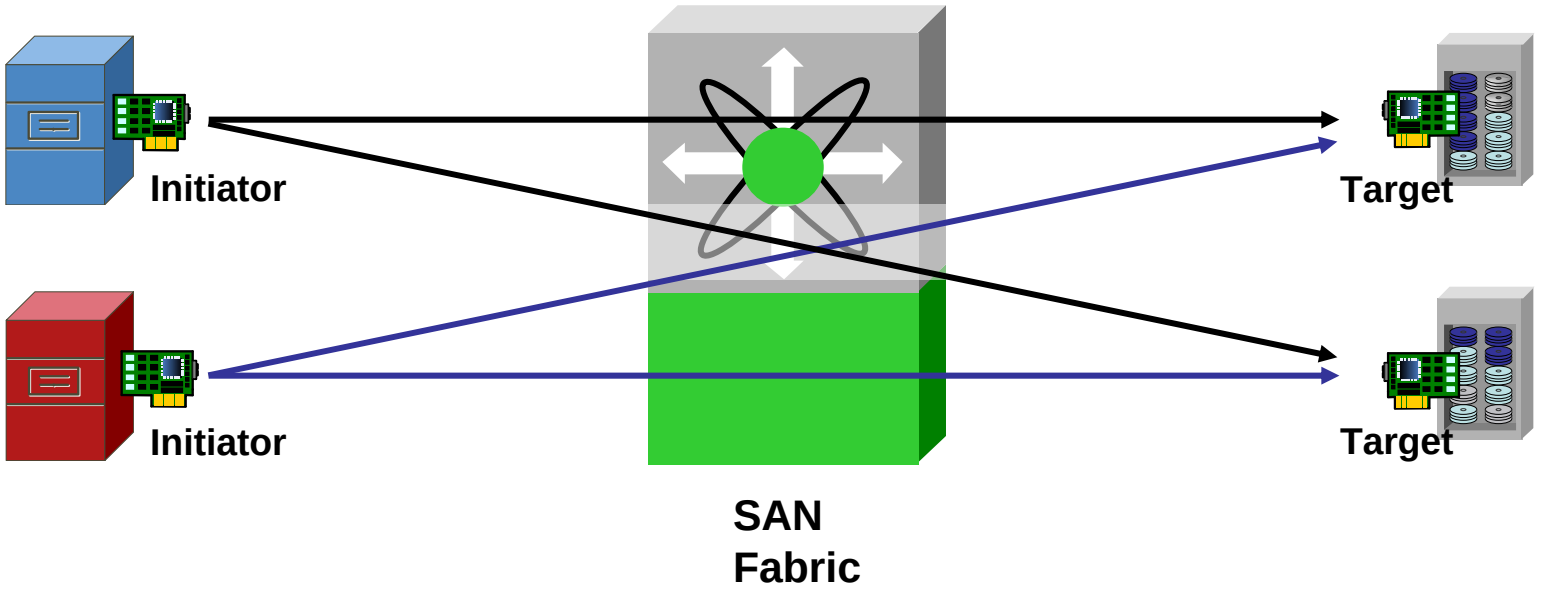

- Adding more storage requires administrative changes
- Administrative overhead, prone to errors
- Complex coordination of data movement between arrays

Storage Volume Virtualization

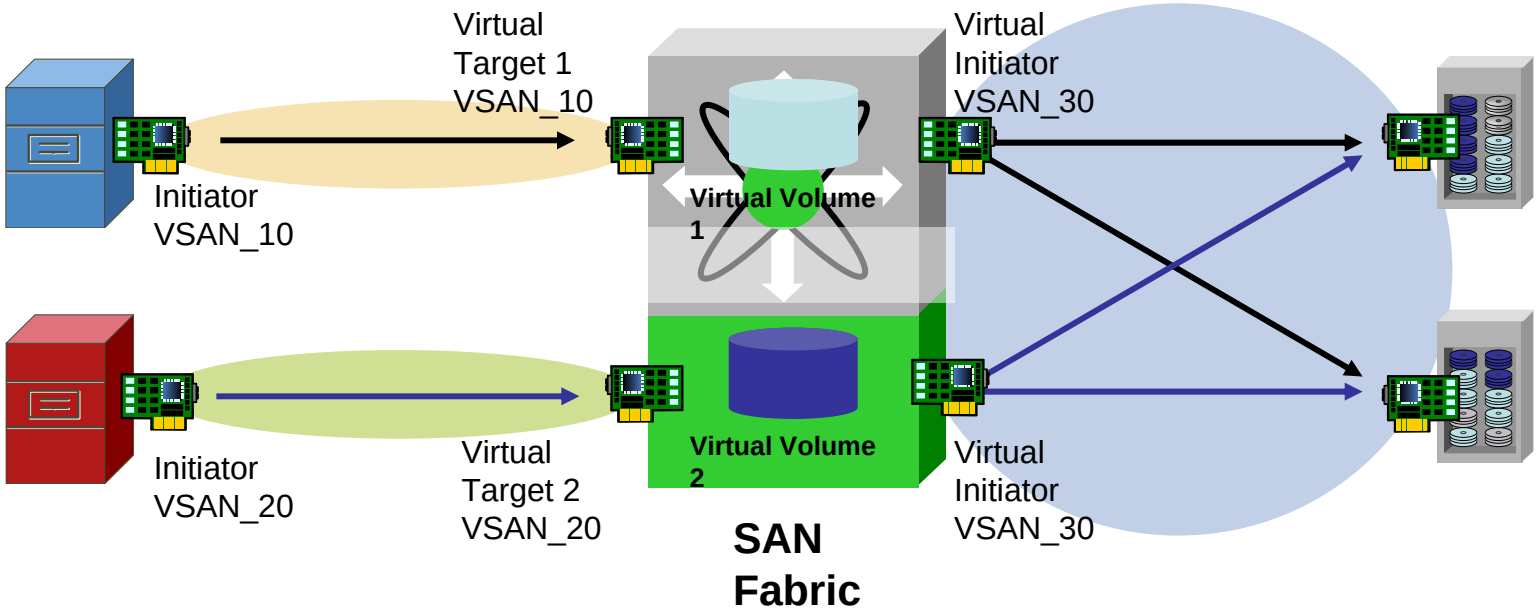

- A SCSI operation from the host is mapped in one or more SCSI operations to the SAN-attached storage
- Zoning connects real initiator and virtual target or virtual initiator and real storage
- Works across heterogeneous arrays

Sample Use: Seamless Data Mobility

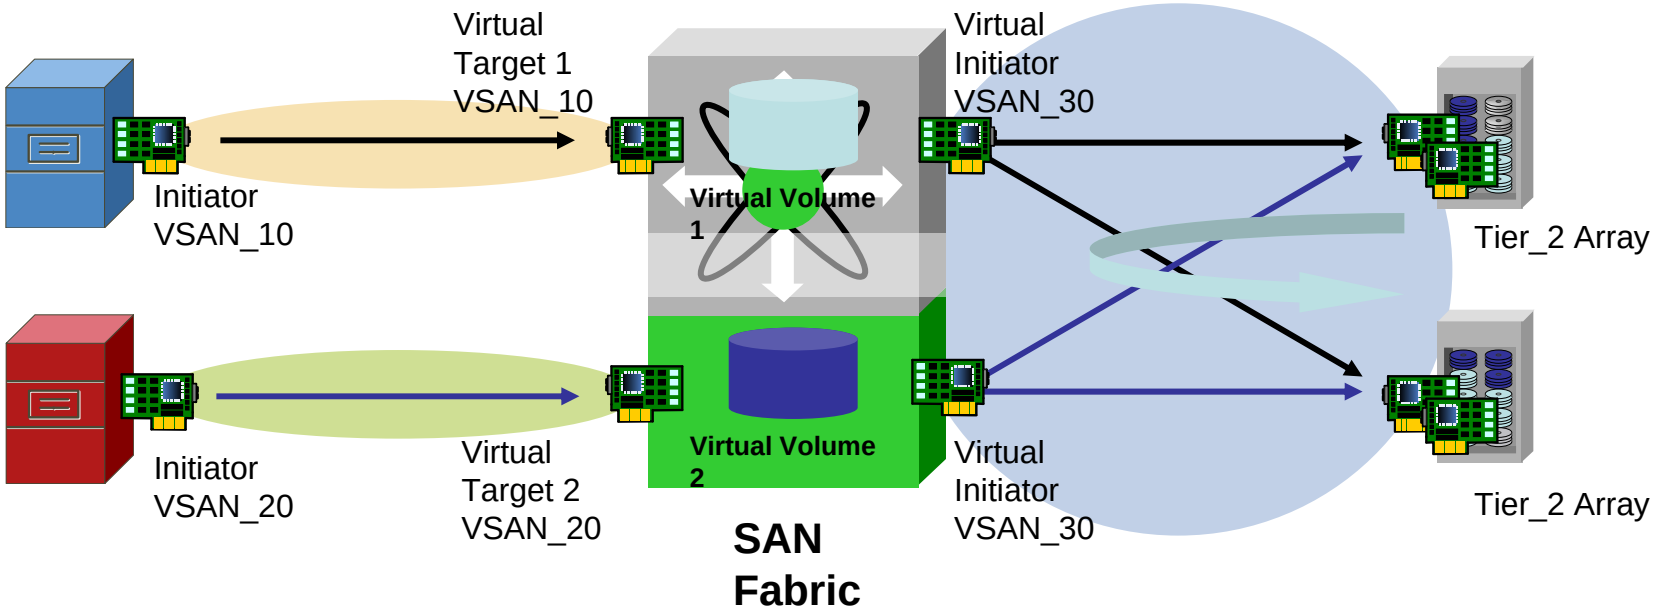

- Works across heterogeneous arrays
- Nondisruptive to application host
- Can be utilized for “end-of-lease” storage migration
- Movement of data from one tier class to another tier

Your session feedback is valuable

Please take the time to complete the breakout evaluation form and hand it to the member of staff by the door on your way out

Thank you!

Recommended Reading


CISCO