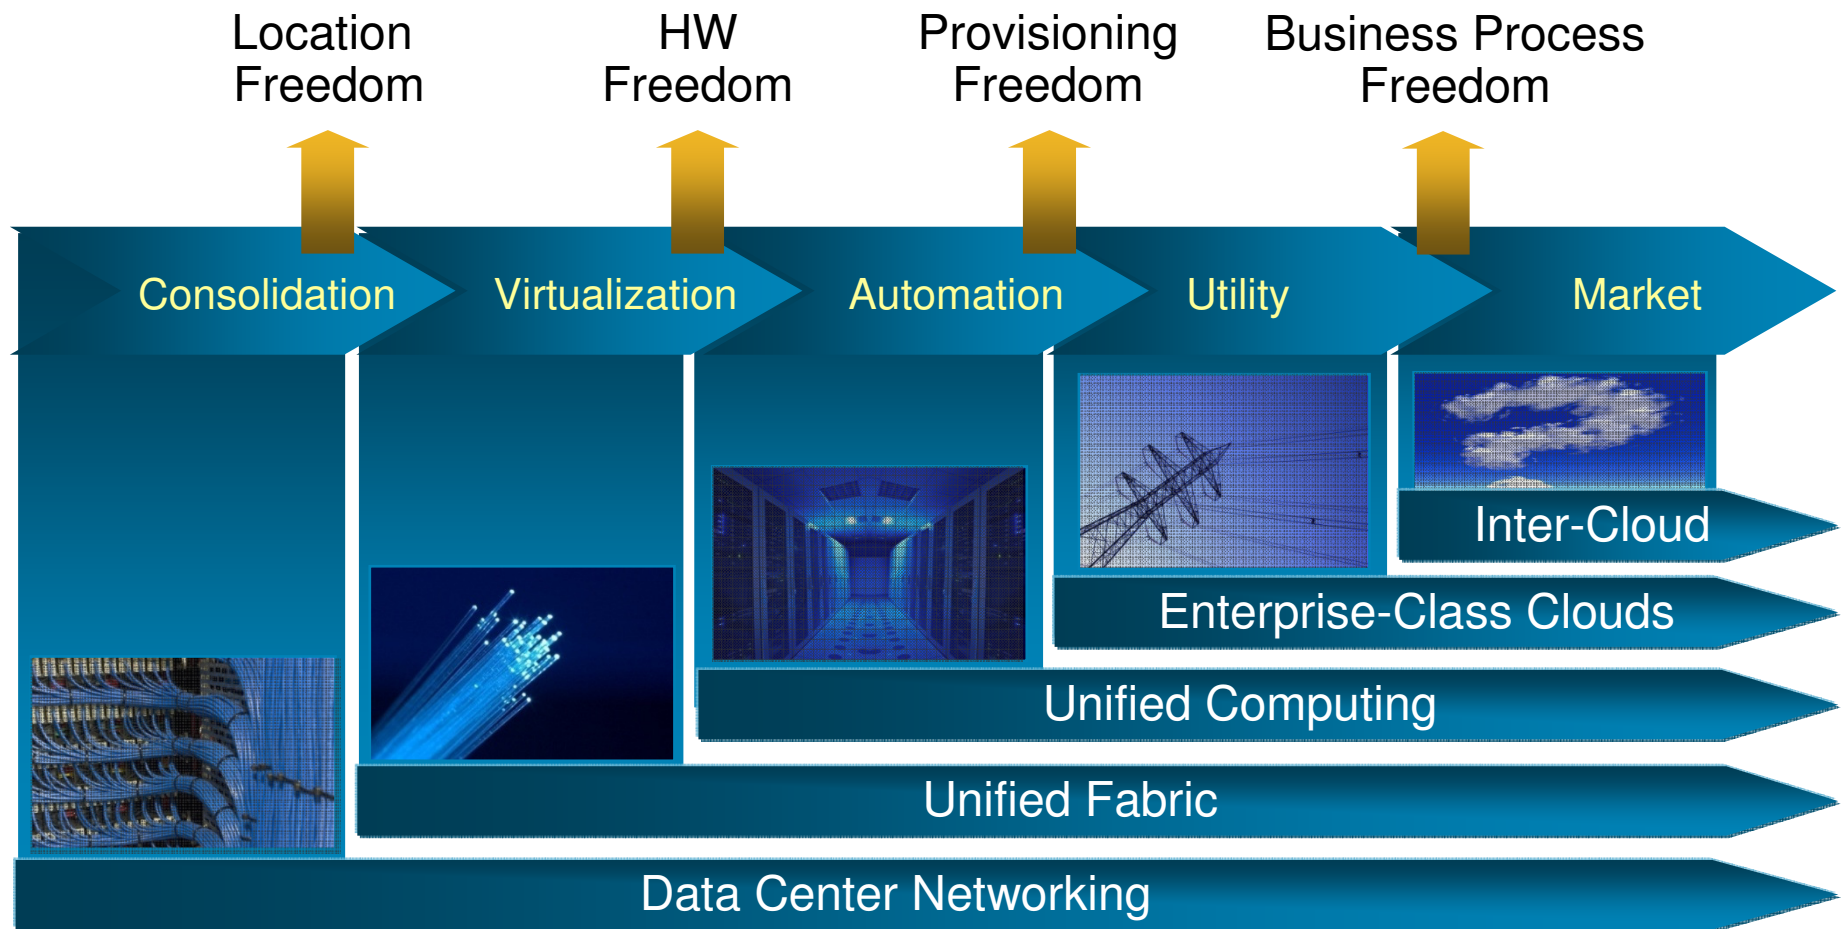


Data Center 3.0



Višnja Mllovanović, Systems Engineer
vmilovan@cisco.com

Data Center 3.0 Evolution Path



Agenda




- Today's Trends and Evolving Standards
- Unified Fabric and Nexus Switches
- Unified Computing System
- Architecture Design

Today's Trends and Evolving Standards




- **DC Consolidation**
- **Server Consolidation**
- **Server and DC Virtualization**
- **Green Field Data Centers**

Reducing Cable & Power Costs with Virtualization



‘ The Cable Mess ‘




Server Virtualization - Key DC Trend

Today




- Many under utilized servers
- Cable sprawl
- High power, cooling costs
- High CAPEX
- For every \$1 spent on server capex ~\$5 spent on opex

Virtualization Step1




GE




- Cable sprawl
- power, cooling costs
- Less number of access layer Ethernet ports

Virtualization Step2




10 GE




- GE to 10GE in access layer
- Less interfaces – reduced Cable sprawl
- Savings from power and cooling

Virtualization Step3



10 GE/FCOE




- Unified I/O - LAN & SAN consolidation
- Reduce NICs, HBAs,
- Reduce cabling
- More Savings from power and cooling
- Lower capex

Data Center Evolution

Virtualized, Flexible, Distributed Computing

- The Data Center is Changing
- Servers – CPU Density, Blade Systems, Clouds
- Storage – NAS/FCoE (Ethernet-connected) “Paradigm shift”
- Virtual Machines – changes everything in the DC
- Nexus – Enabling standards based virtualization at a network scale




**Virtualize at
Network Scale**

Data Center Architecture


Evolution of the Hierarchical Design

- The Data Center Architecture is based on a hierarchical design model
 - Aggregation block contains the access and aggregation layers
 - Core provides layer 3 boundary to the rest of the network
 - Dedicated service switches provide application load balancing, firewall, etc.
- Architecture is based on optimized design for control plane stability and scalability
- Need to understand how the design needs to evolve to accommodate server, applications and facilities requirements



Data Center Architecture

Network Requirements for Virtual Machines



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

Data Center Architecture

Unified Fabric – Ethernet Based Storage



- Storage Requirements are evolving
 - Virtual Machines increasing the number of SAN connected devices
 - Disk consolidation and Data De-duplication
- Unified Fabric provides efficiencies as the number of servers connected to network storage increases
- Unified Fabric provides one of the foundations for Unified Computing

Data Center Architecture

Changing Facilities Requirements

- Data Center environments responding to facilities demands and technology changes
- Technology changes will impact any cabling plant design
 - Migration to 10GE as the default LoM technology
 - Migration to 40GE and 100GE within 5 years
- Requirement for a power efficient and very flexible access architecture
- Not a choice 'between' EoR and ToR



End of the Row (Centralized Server to Network cabling)



Top of the Rack (De-centralized)

Transmission Media: SFP+ Copper Twin-ax



SFP+ Slot Characteristics:

10 Gbps bandwidth

PHY = 1.8 watt

Max Temp = 70 C

Technology	Cable	Distance	Power (each side)	Transceiver Latency (link)
SFP+ Cu Twin-Ax	Twinax	10m	~0.1W	~0.25µs
SFP+ USR ultra short reach	MM OM2 MM OM3	10m 100m	1W	~0
SFP+ SR short reach	MM 62.5µm MM 50µm	82m 300m	1W	~0
10GBASE-T	Cat6 Cat6a/7 Cat6a/7	55m 100m 30m	~8W ~8W ~4W	2.5µs 2.5µs 1.5µs

Data Center Architecture

Evolving Ethernet Capabilities - DCE

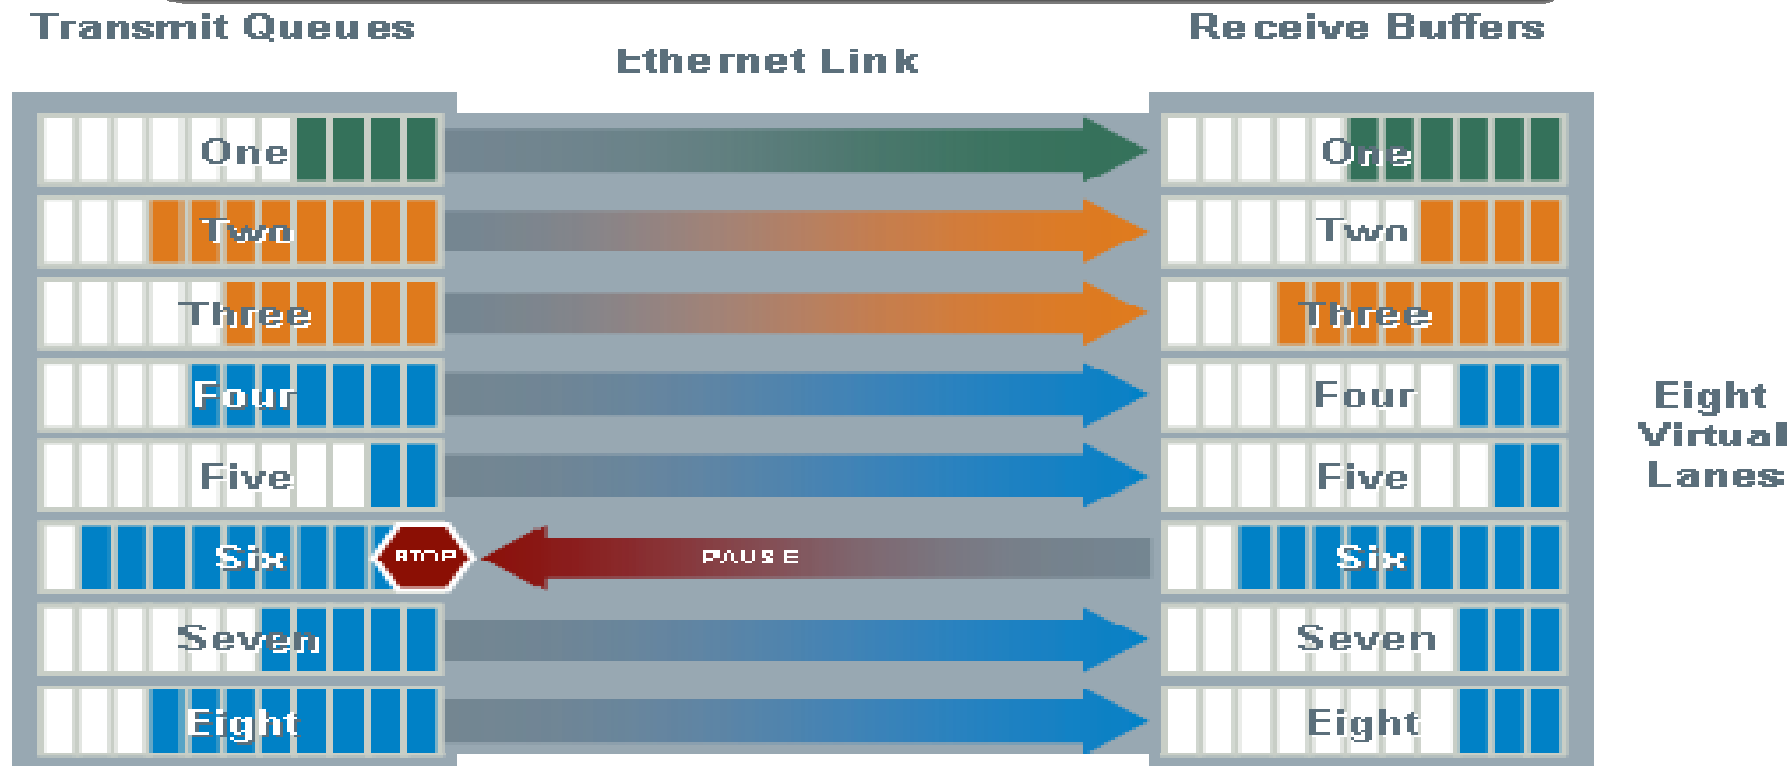


Feature / Standard	Benefit
Priority Flow Control (PFC) IEEE 802.1Qbb	Enable multiple traffic types to share a common Ethernet link without interfering with each other
Bandwidth Management IEEE 802.1Qaz	Enable consistent management of QoS at the network level by providing consistent scheduling
Congestion Management IEEE 802.1Qau	End-to-end congestion management for L2 network (future)
Data Center Bridging Exchange Protocol (DCBX)	Management protocol for enhanced Ethernet capabilities
L2 Multipath for Unicast and Multicast	Increase bandwidth, multiple active paths. No spanning tree (future)

Enabling Differentiated Services in an Ethernet Fabric

Priority Flow Control

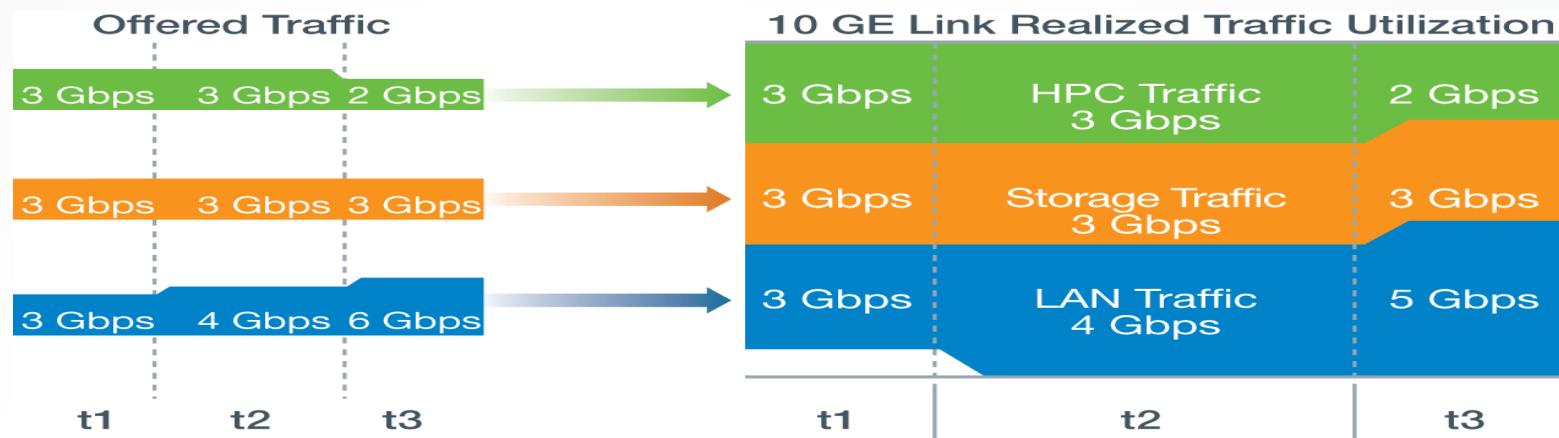
Priority based Flow Control



- Enables lossless behavior
for each class of service*
- *PAUSE sent per priority
when buffers limit exceeded*

Priority based bandwidth management

Priority based Bandwidth Management




***Enables Intelligent sharing of
bandwidth between traffic classes
control of bandwidth***

- 802.1Qaz Enhanced Transmission***

FCoE Enablers

- 10Gbps Ethernet
- Lossless Ethernet
 - Matches the lossless behavior guaranteed in FC by B2B credits
- Ethernet jumbo frames
 - Max FC frame payload = 2112 bytes

Normal ethernet frame, ethertype = FCoE



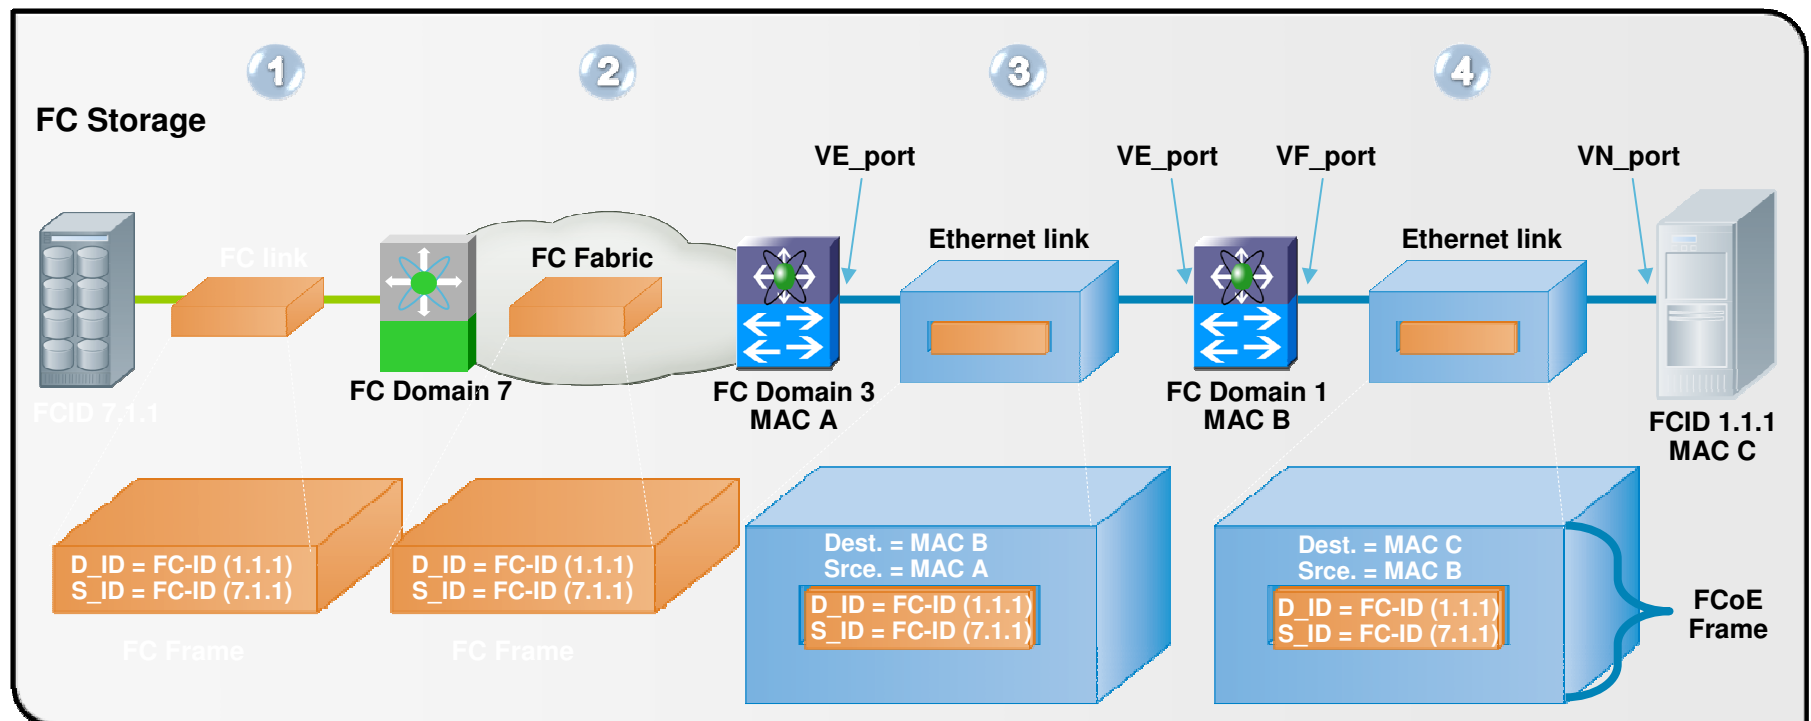
FCoE Specification progress

- Cisco submitted FCoE proposal on May 22, 2007 as a joint proposal among 16 companies
 - Adopted by ANSI T11 FC-BB5 in June 2007
 - Frame format agreed upon by T11 in August 2007
 - Support from entire storage and switching industry: EMC, HDS, HP, IBM, Sun, Brocade, NetApp, Cisco, Emulex, Qlogic, Nuova, Intel
- Fully ratified in June 2009 (www.t11.org/fcoe)

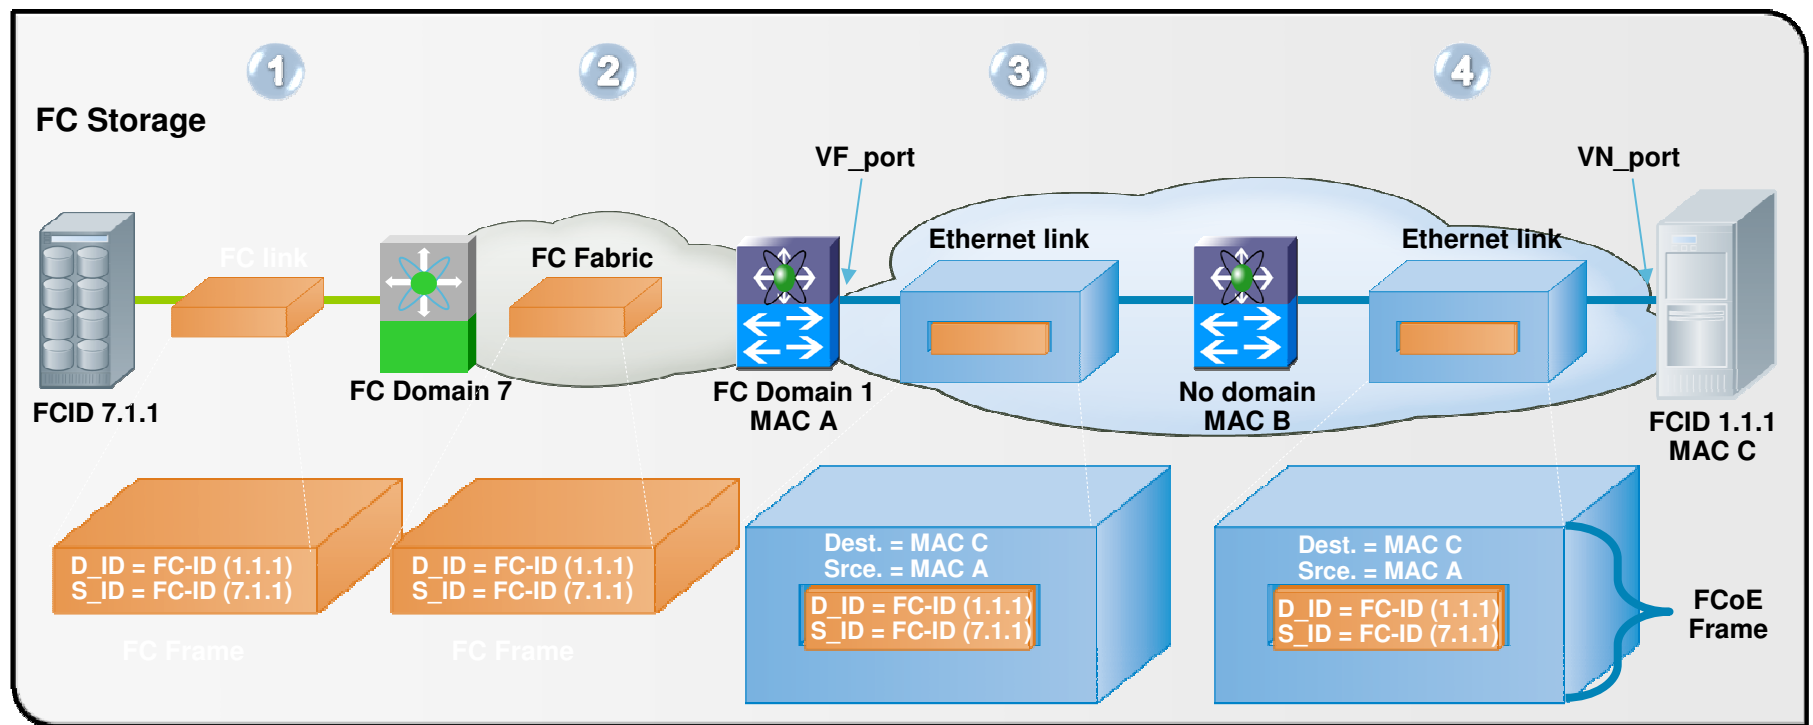


FCoE Forwarding (VE_ports)

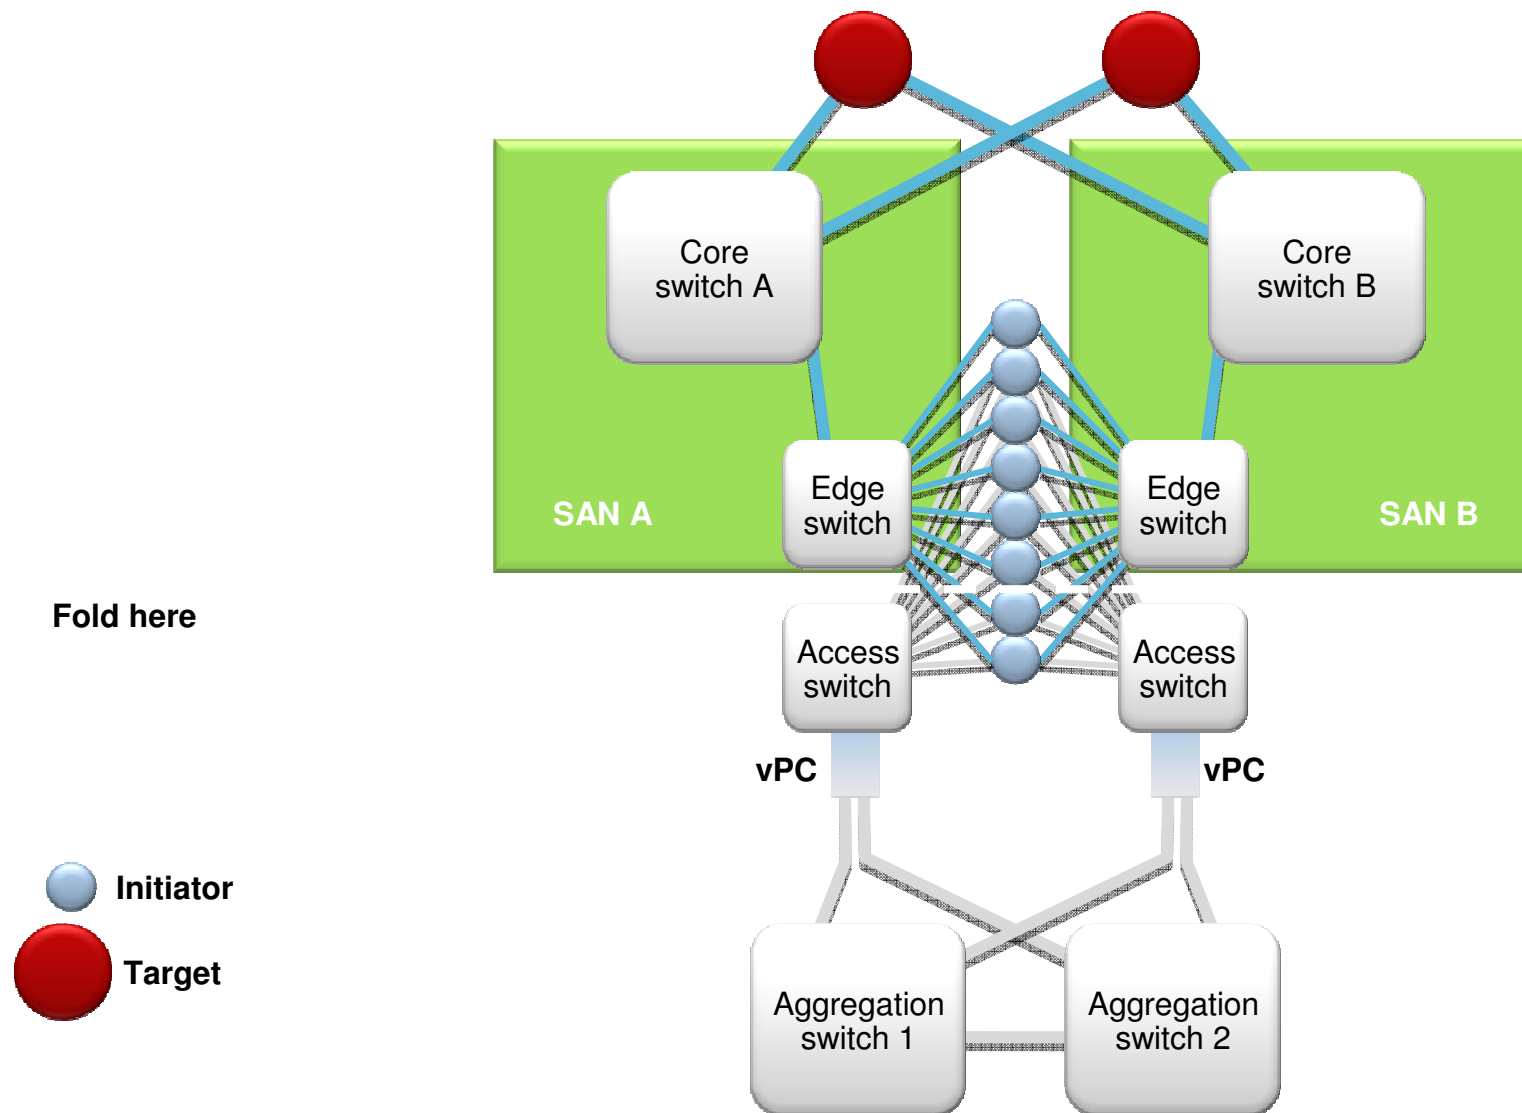
- FCoE frames have:
 - MAC addresses (hop-by-hop)
 - FC addresses (end-to-end)




FCoE Forwarding (FIP snooping)



Mapping to Ethernet, instructions



Core-edge design, mapped to Ethernet



Just a first cut.

FCoE and IP traffic

- Who says FCoE and IP ***must*** run on the same wire?

A choice is not an obligation

- Choose to share a wire when

It's simple

It saves more money than the alternatives

- Sharing a device is still Unified I/O

Even if some cables are dedicated to IP traffic, and some other are dedicated to FC traffic

Not ideal, agreed, but do what's best for the customer



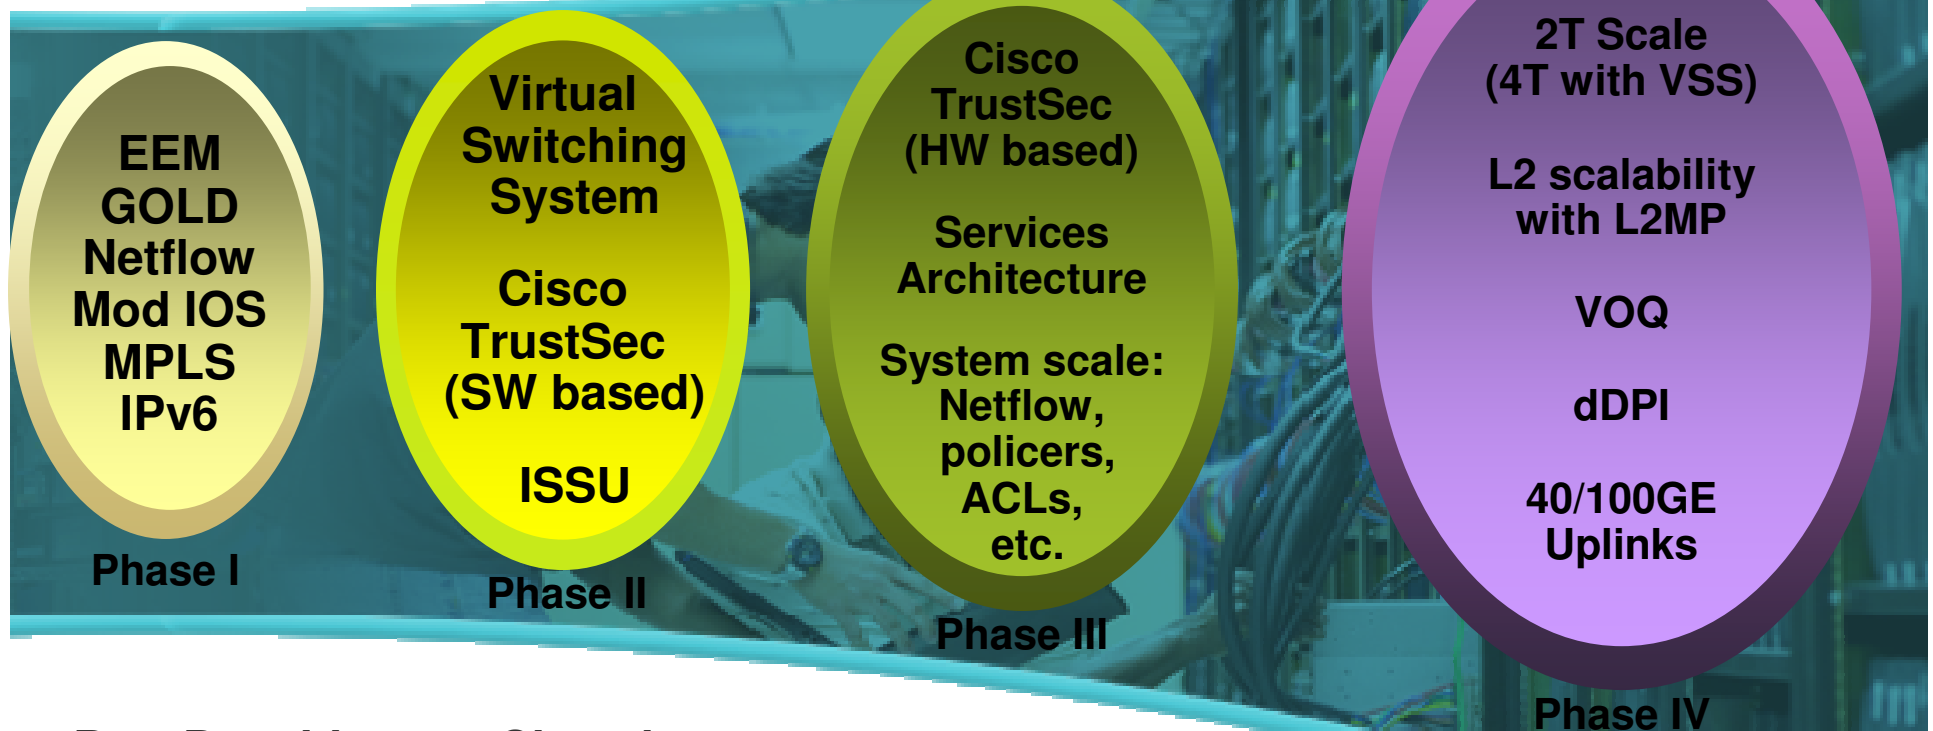
DC Switching and Nexus Series



- Cat 6500
- Nexus 7000
- Nexus 5000
- Nexus 2000
- Nexus 1000

Catalyst 6500

Data Center Technology Roadmap



Port Densities per Chassis

66 x 10G
384 x 1G

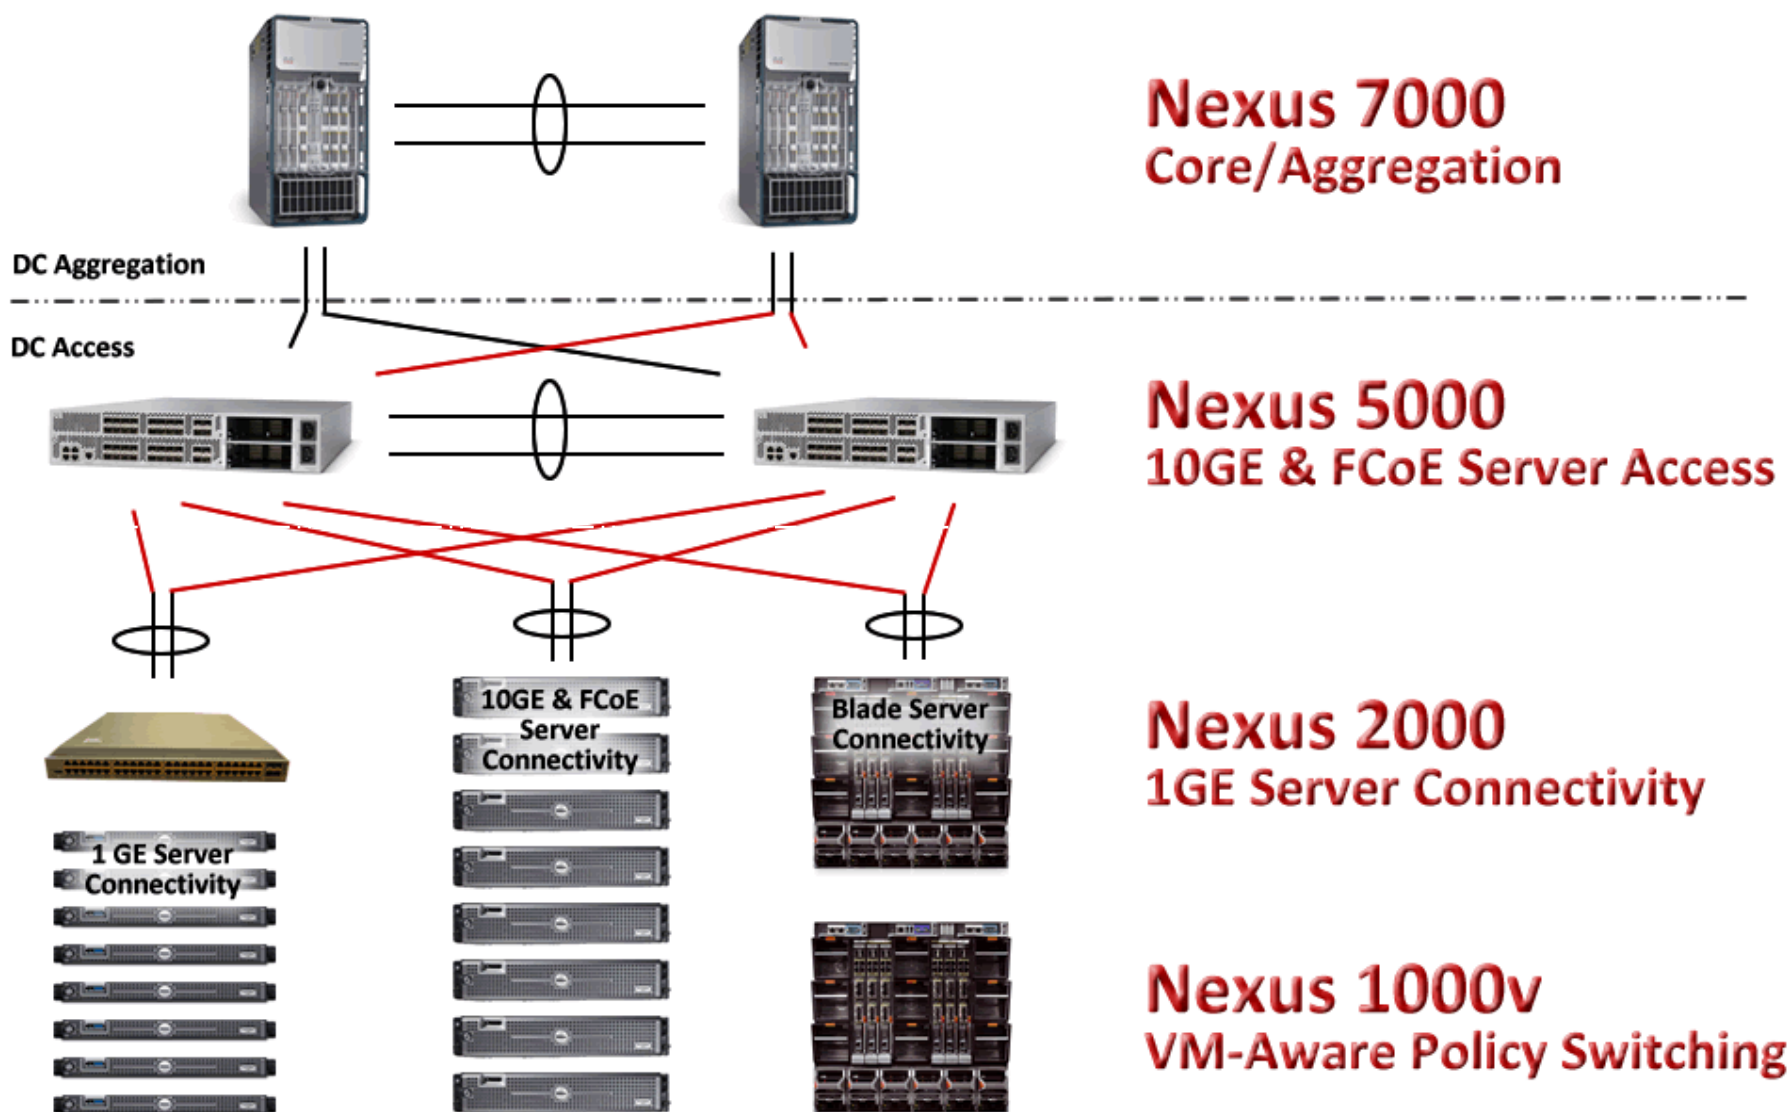
130 x 10
384 x 1G

130 x 10G
768 x 1G

>40 x 40G
>250 x 10G
>1,000 x 1G


Data Center Access Evolution

Virtual Access




Nexus 7000 Platform Overview

Industry's First Data Center Class Platform



Nexus 7000 and NX-OS

- 10 & 18 Slot versions
- 15+ Terabit System
- Unified Fabric Ready
- Modern, Modular OS
 - Cisco TrustSec
- Continuous Operations




Supervisor

10G Ethernet
• 32 Port SFP+ 10G

1G Ethernet
• 48 Port 1G-TX
• 48 Port 1G - SFP

10G DCE
• future

Line card
Modules



Cisco NX-OS Multi-protocol Operating System
Data Center Network Manager (DCNM)

Nexus 7000 Chassis

System status LEDs

Cable Management

Linecard slots (8)


Supervisor slots (2)

Module Ejector Release Buttons

Lockable Front Doors (opt)


Air Filters (opt)

Air intake
(front/bottom)



Air outlet
(rear/top)

21 RU 36.75" (934mm)



Fans
(2 system + 2 fabric)

Fabric modules
(up to 5)

Module Ejector
Release Buttons


Power Supplies
(up to 3)

Chassis depth 30" (762mm)

17.3" (439mm)



Front and Rear of 10 Slot Chassis

Bandwidth Scales with Each Fabric Module



Investment Protection and Unified Fabric

Consolidation with Nexus 7000



Lower Capital Expenditure

Consolidate multiple devices
Remove interconnect links
Logical separation of each Virtual Device Context



Reduce Operational Costs

Fewer number of devices to manage
Lowers overall data center power draw

Nexus 7000 High Availability

How to enable the level of High Availability required by the Data Center?

Highly Available Design

- **Structured Design** allows you to manage and understand:
 - Traffic flows
 - Network failure behavior
- **Modular design**
 - Allows for easier evolution and change to the network
- **Hierarchical design**
 - Provides for improved scalability
 - Separates network services into manageable building blocks

Operations











- **Best Practice**
- **Staff Well Trained**
- **Formal Network Documentation**
- **Network Prototypes**
- **Formal Procedure in case of Failures**

Technology: Nexus 7000

- **Fully Redundant Hardware:**
 - Redundant Supervisors
 - Redundant Fabric Modules
 - Redundant Fans and Power Supplies
- **Software Robustness**
 - Hitless ISSU
 - Hitless Supervisor Switchover
 - Hitless Process Restarts
 - Modularity
- **Features**
 - VDC, vPC, STP Enhancements, etc

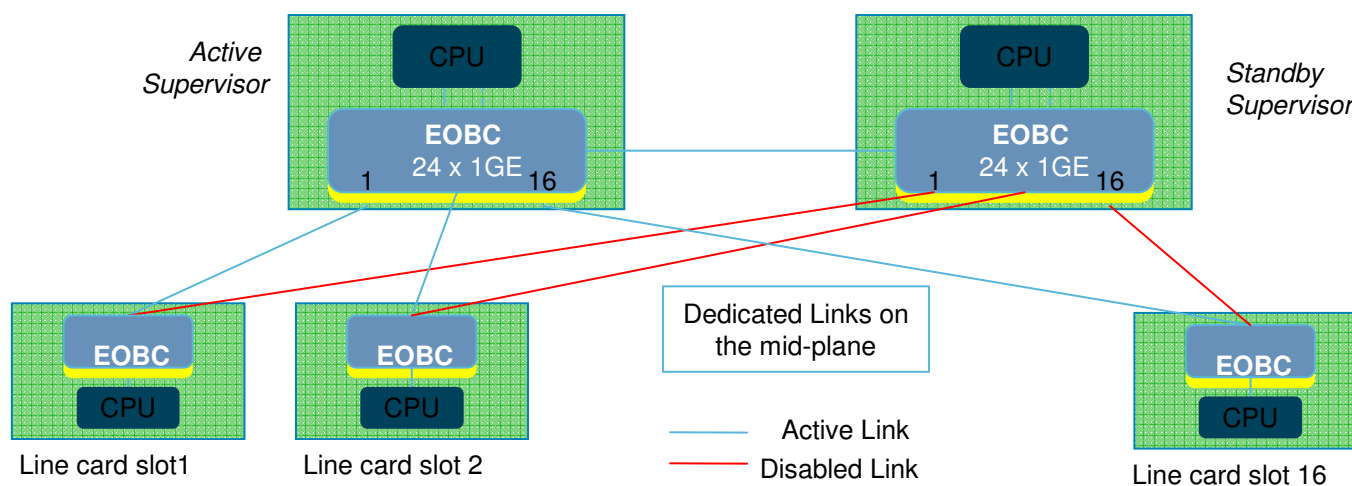
Nexus 7000 High Availability Dual vs. Single Supervisor

- Comparison from a **system level** prospective:

	Dual Supervisor	Single Supervisor
Hitless ISSU		-
Hitless Supervisor Switchover		-
Redundant EOBC		-
Redundant Arbiter		-
Hitless Fabric OIRs		
Stateful Process Restart		
Graceful Restart		

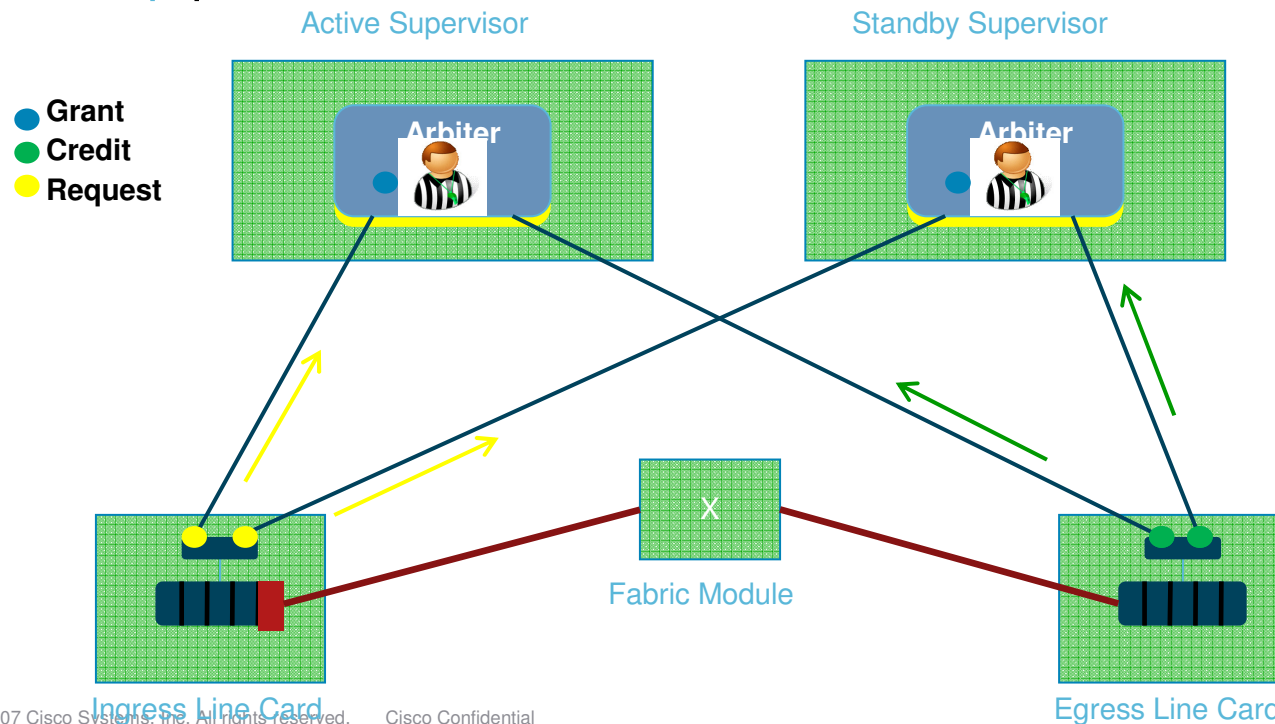
Nexus 7000 High Availability Supervisor – EOBC

- Ethernet **O**ut of **B**and **C**hannel:
 - The connectivity between the CPU on the Supervisor and the CPU on each line card is provided by the **EOBC**.
 - Given its role, the EOBC on the Nexus 7000 has being built to provide the **highest level of availability**.
- Advantages of the Nexus 7000 **EOBC**:
 - The Nexus 7000 EOBC is a **gigabit ethernet switched architecture** between line cards and Supervisors.

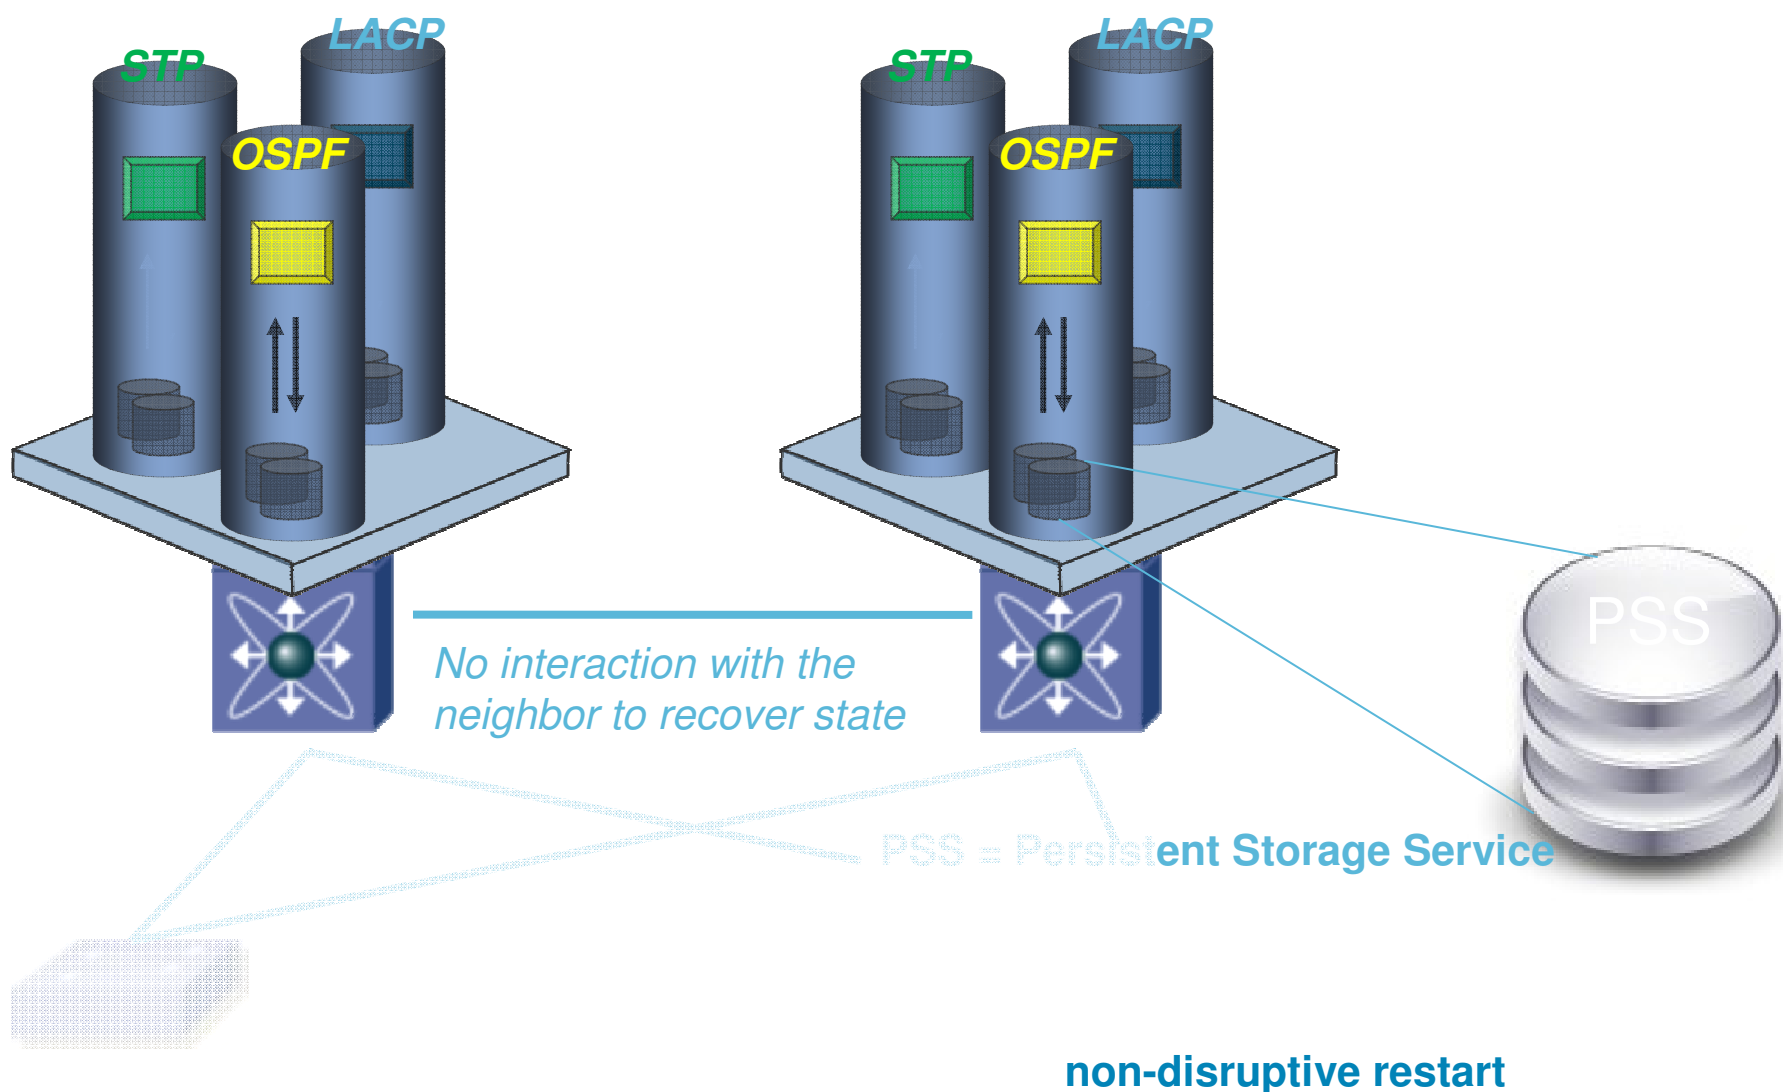


Nexus 7000 High Availability Supervisor – Packet Arbiter

- Advantages of the Nexus 7000 Arbitrated Fabric
 - Head of Line Blocking prevention
 - Fair Access to the Fabric
 - Dual-Sup provides redundant fabric arbiter




Nexus 7000 High Availability NX-OS – Stateful Process Restart



Nexus 7000 High Availability NX-OS – Stateful Process Restart

- Common question: “Which protocols support Stateful Process Restart?”
- The Protocols/Processes supporting Stateful Process Restart are:
 - Layer 3: **NETSTAK, OSPFv2, GLBP, VRRP, HSRP, ARP, DHCP_SNOOP, NTP, ICMPv6 and CTS**
 - Layer 2: **vPC, STP, LACP, CDP, IGMP and Dot1.X**
- In future releases we expect **EIGRP, OSPFv3** and **ISIS** to be supported.



Elastic Workload Mobility

VN-link + OTV synergies

N1Kv-VN-Link:

VM \leftrightarrow Network Communication:

Preflight notification

Move completion notification

Re-address notification

OTV & Routing Control Plane:


On receipt of VN-link notifications from the N1Kv, OTV can:

- Complete pre-flight network checks
- Pre-converge the Network for fast-reroute
- MAC address re-localization
- L3 IP localization
- Establish LAN extensions on demand
- Provide intelligent MAC address handling

MAC scaling

Elastic workload mobility

- anywhere
- on demand



Nexus 5000 Family Multi-Protocol Server Switch

Industry's First I/O Consolidation Virtualization Fabric for Enterprise Data Center

Switch Family


NX5010



28-Port L2 Switch

- 20 Ports 10GE fixed
- 1 Expansion Module

NX5020



56-Port L2 Switch

- 40 Ports 10GE fixed
- 2 Expansion Modules

Expansion Modules



Fibre Channel

- 8 Ports 1/2/4G FC



FC + Ethernet

- 4 Ports 10GE
- 4 Ports 1/2/4G FC



Ethernet

- 6 Ports 10GE

Partners



CNA

- 2 Port 10GE/FCoE



Server Adapter

- FCoE SW stack

OS


Cisco NX-OS

Cisco Fabric Manager and Cisco Data Center Network Manager

All 10GE switch/module ports are FCoE/Data Center Ethernet capable

Nexus 2000 Fabric Extender Virtual Chassis


Nexus 5000



Nexus 2000 Fabric Extender



Nexus 5000
Virtualized chassis

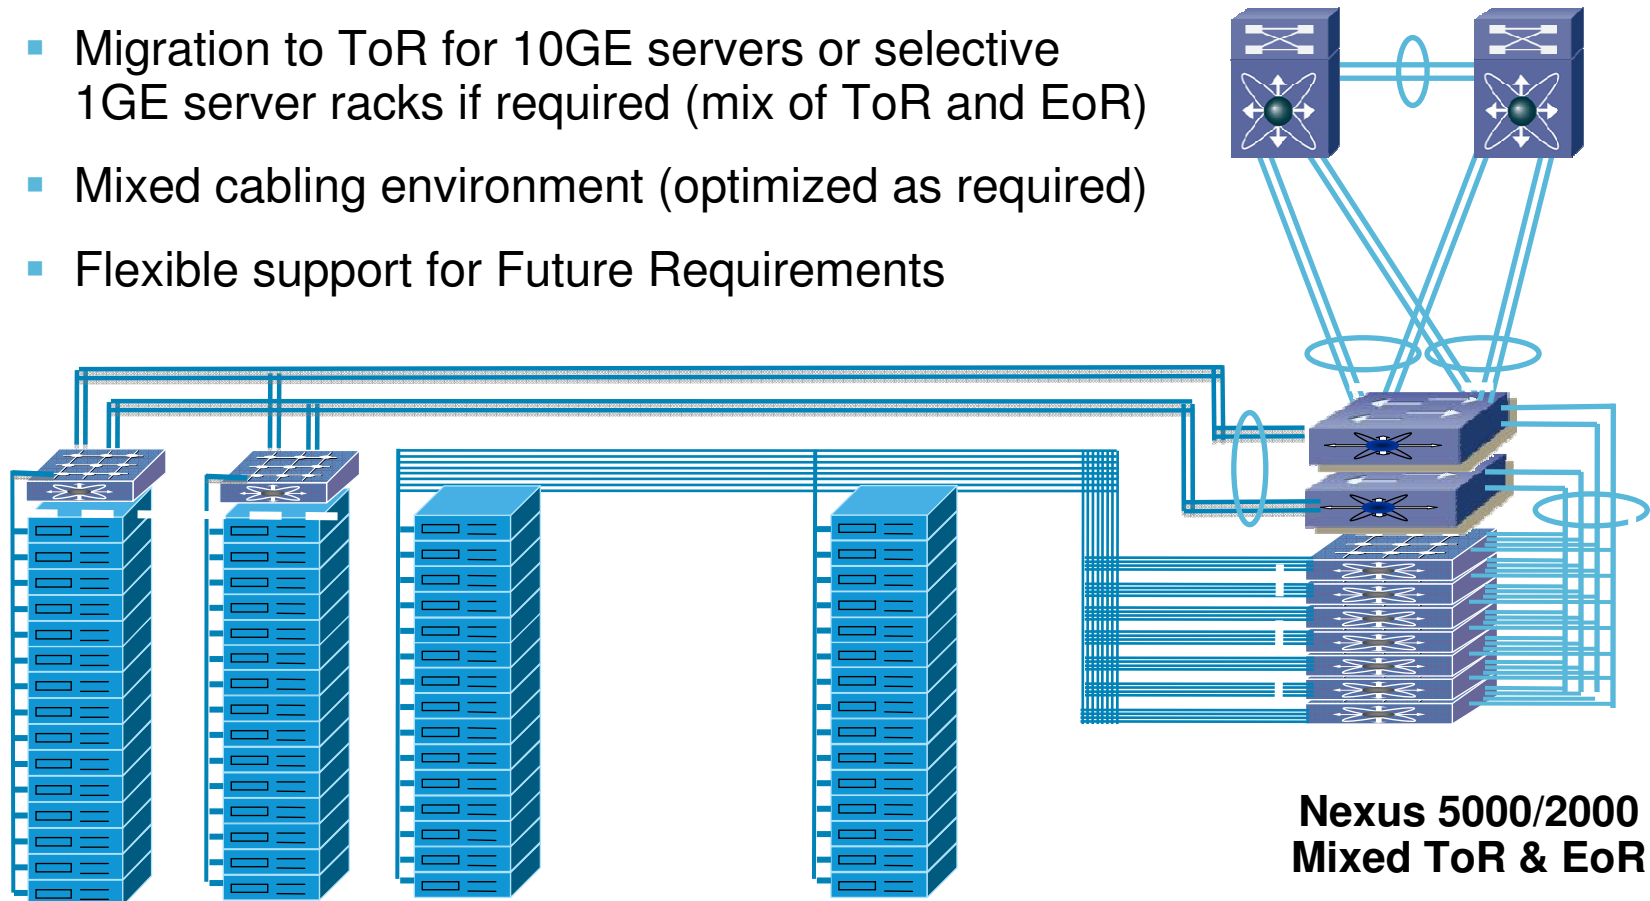


The Nexus 2000 Fabric Extender (FEX) acts as a remote linecard for the Nexus 5000, retaining all centralized management and configuration on the Nexus 5000, transforming it into a Virtualized Chassis

Data Center Access Architecture

N5K/N2K Advantages – Flexible Cabling

- Cisco Nexus Fabric Extender (FEX) and Nexus 5000 provide a Flexible Access Solution
- Migration to ToR for 10GE servers or selective 1GE server racks if required (mix of ToR and EoR)
- Mixed cabling environment (optimized as required)
- Flexible support for Future Requirements




Combination of EoR and ToR cabling

**Nexus 5000/2000
Mixed ToR & EoR**

Data Center Access Architecture

Virtualized Access Switch

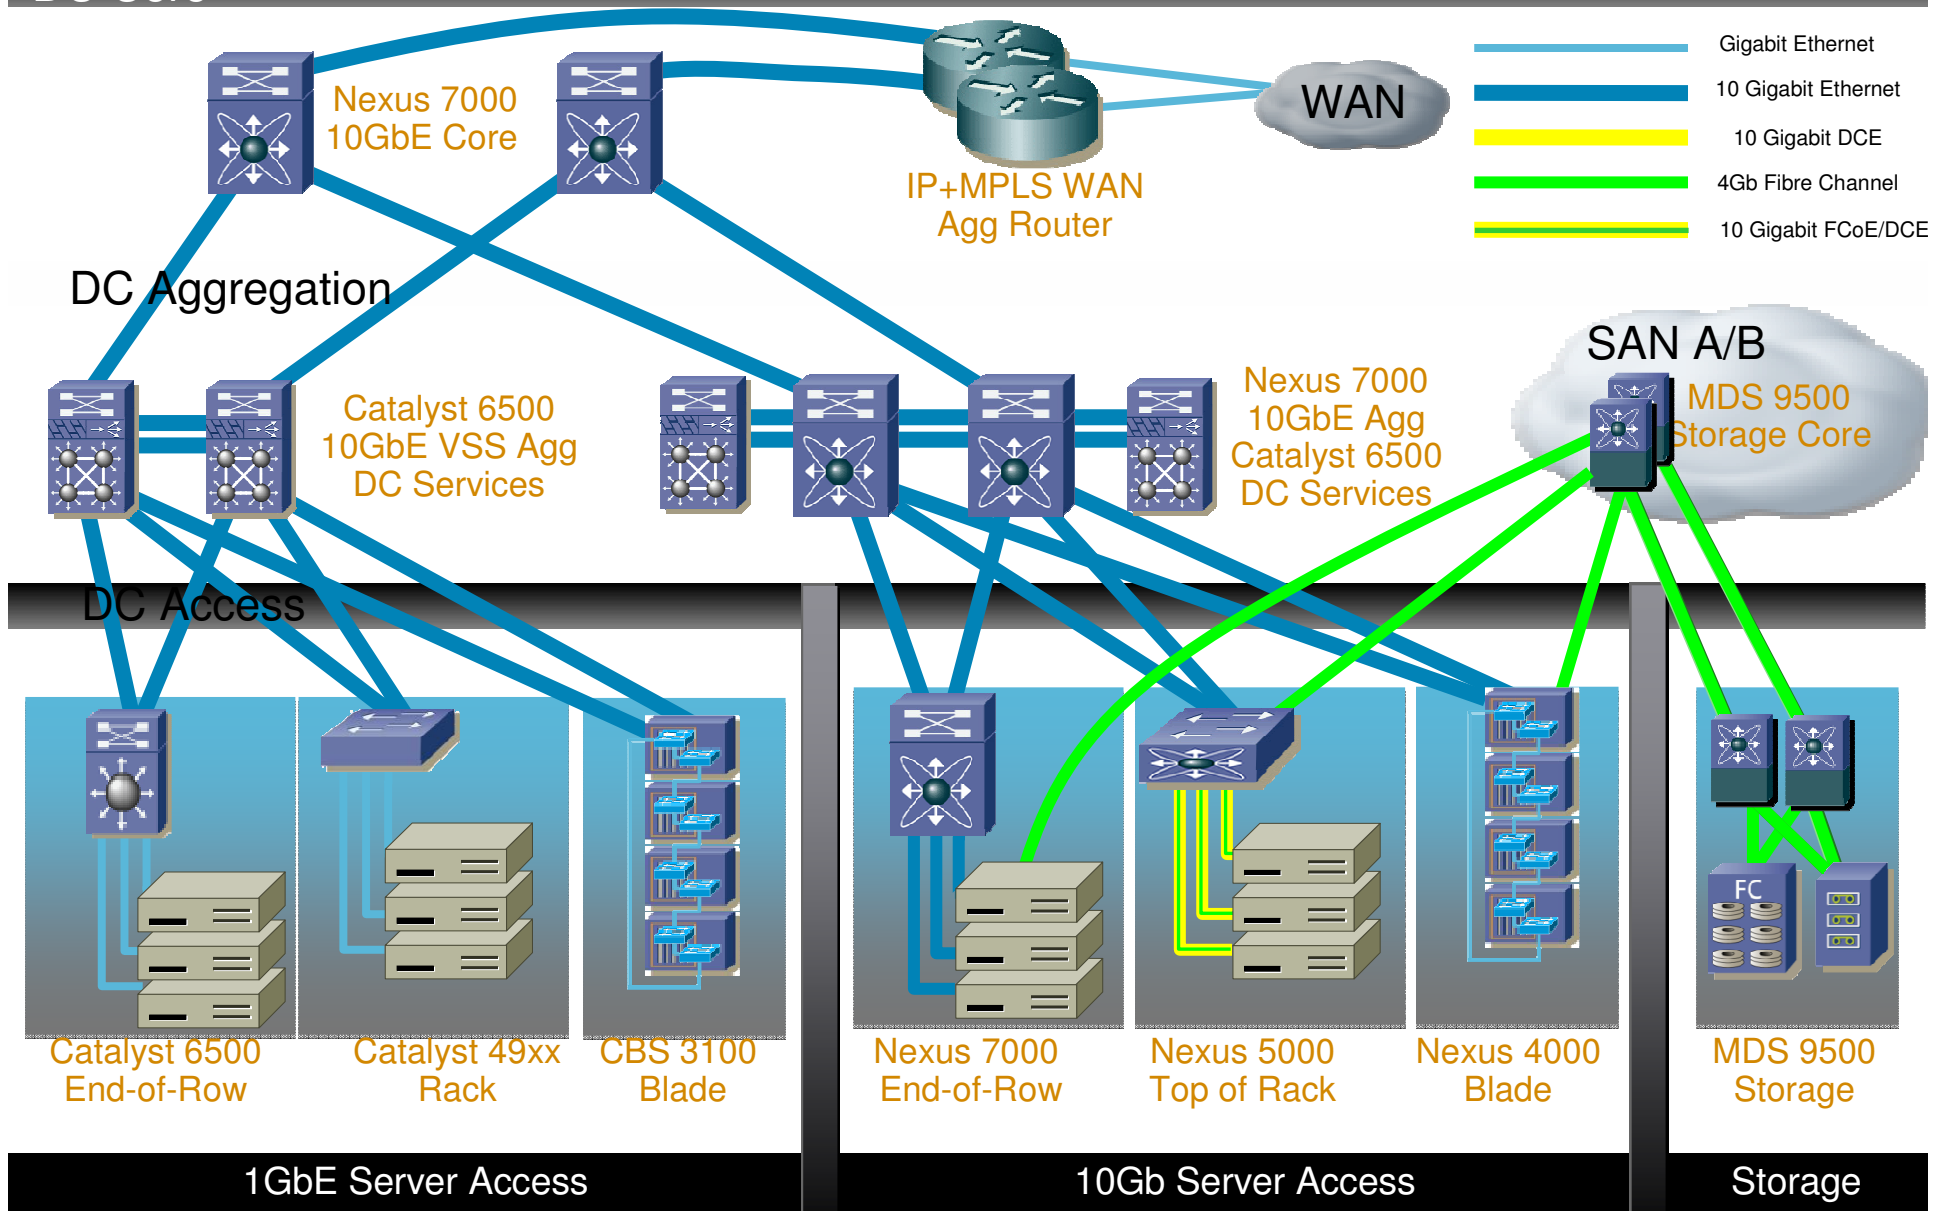
- Virtualized Access Layer
 - Architectural flexibility
 - 1GE to 10GE transition
- Consolidated and Consistent Operations
- Virtualization Optimized VN-Link
- Network Fabric
 - vPC
 - Unified Fabric (FCoE)
 - Lossless Ethernet (DCE)
 - L2 Multipathing (DCE)



Virtualized Access Switch

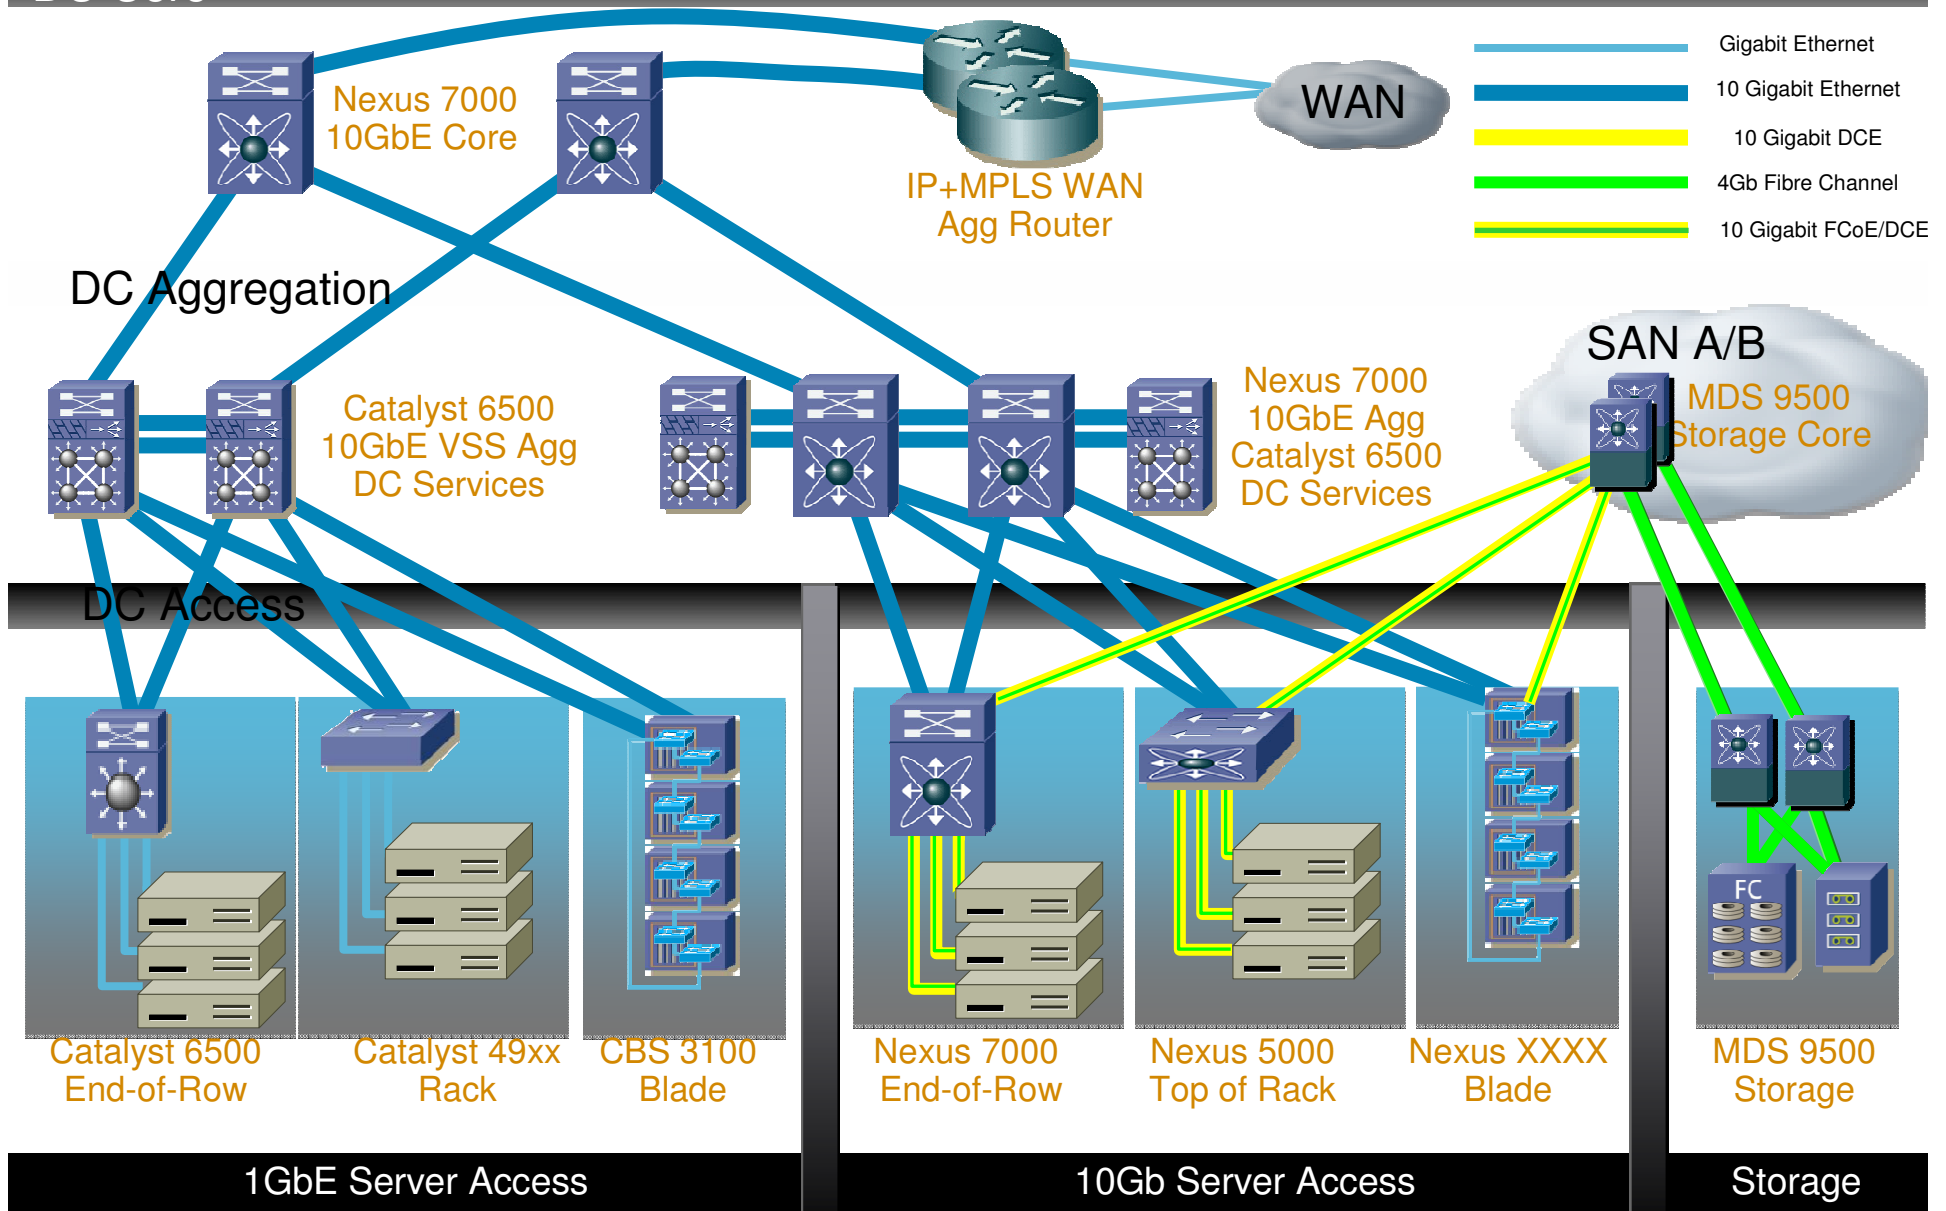
Nexus Deployments

DC Core



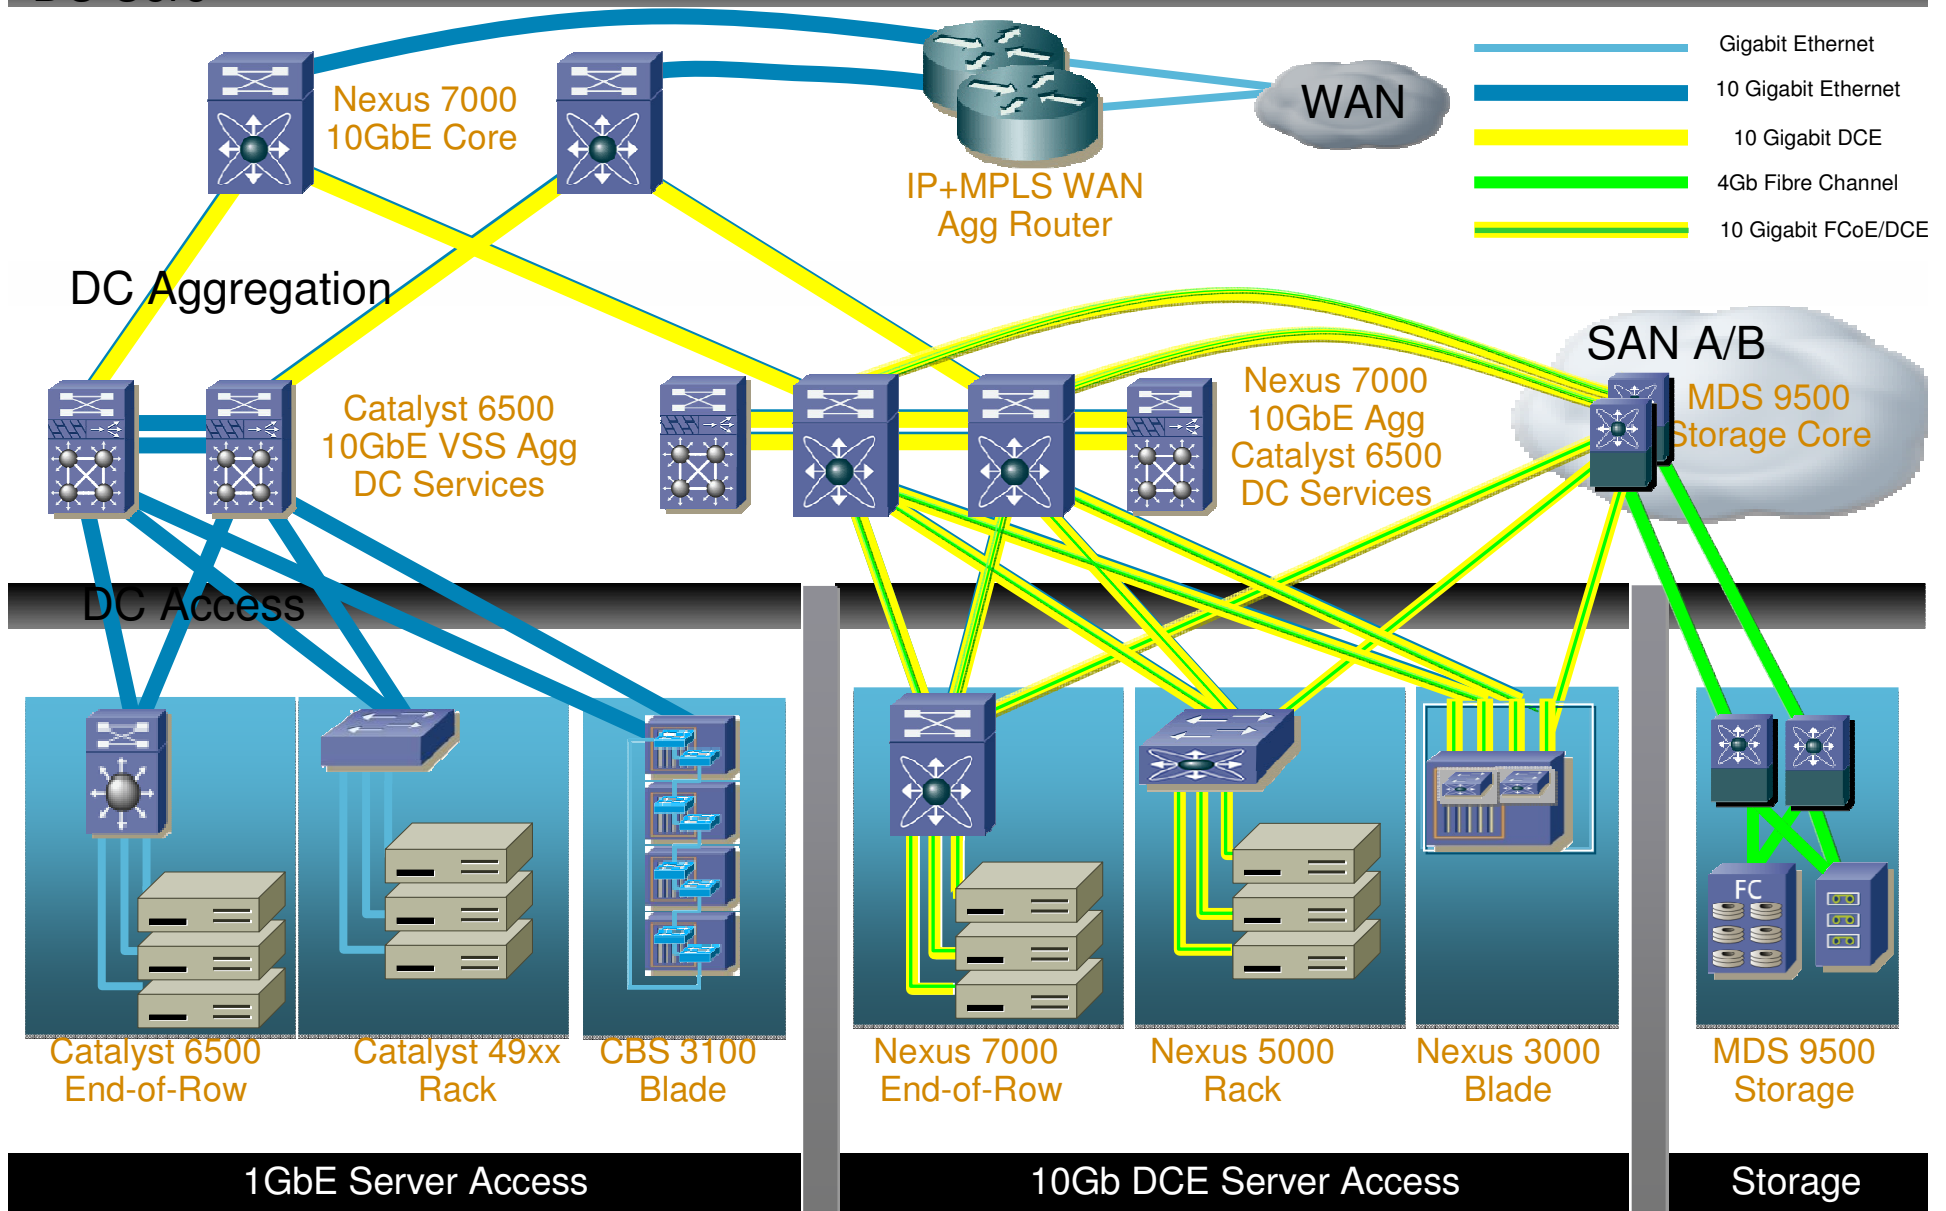
Unified Fabric Evolution

DC Core



Unified Fabric Evolution

DC Core



Cisco VN-Link

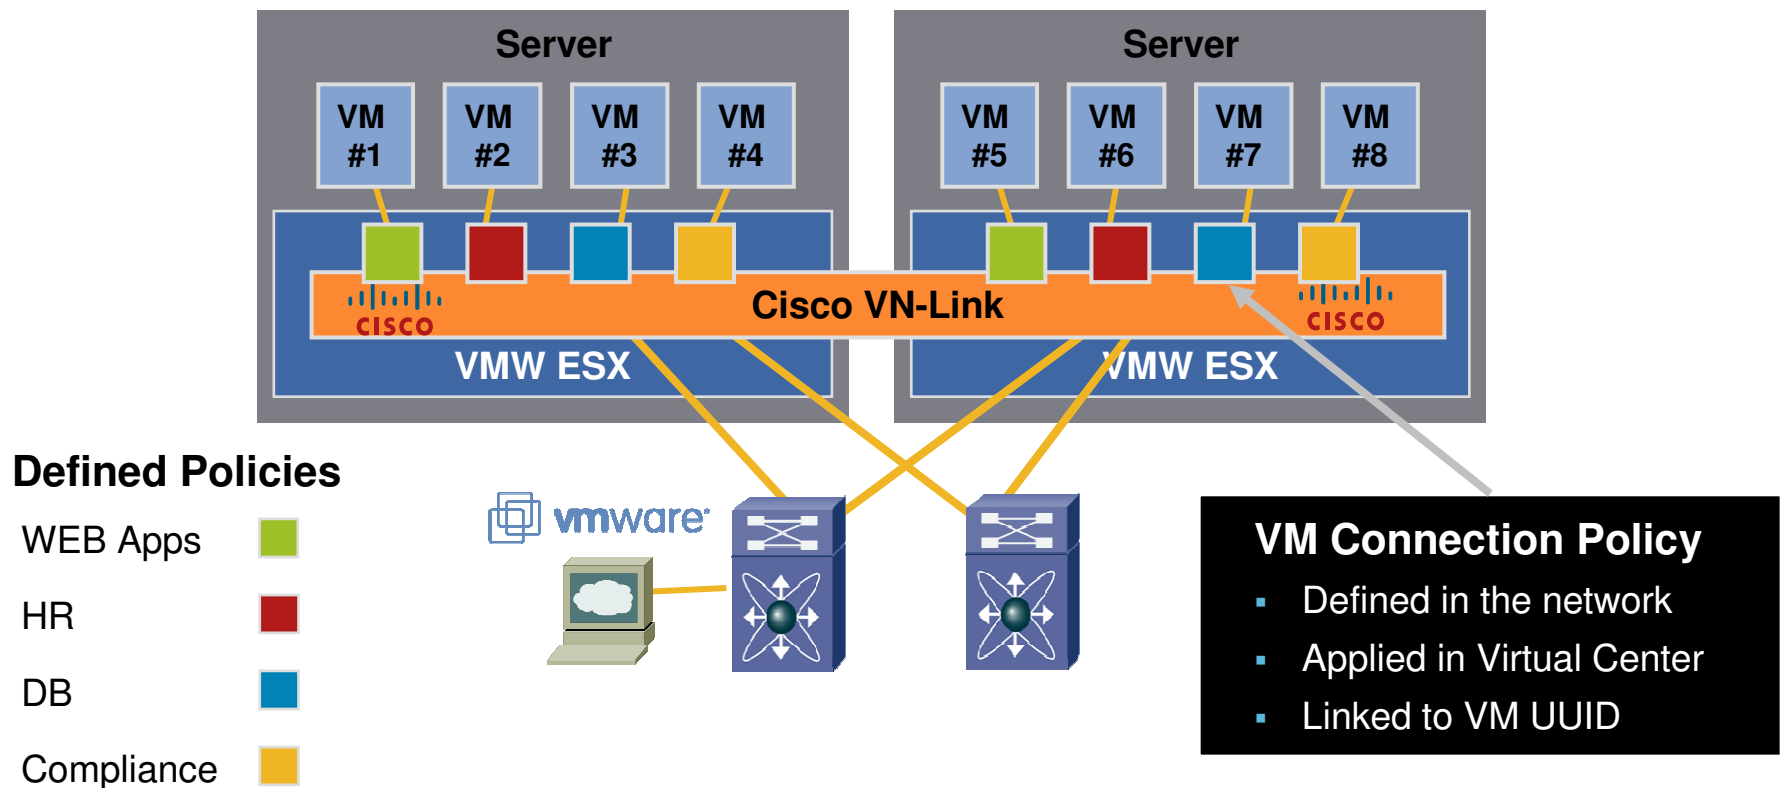
Faster VM Deployment

Cisco VN-Link—Virtual Network Link

**Policy-Based
VM Connectivity**

**Mobility of Network
& Security Properties**

**Non-Disruptive
Operational Model**

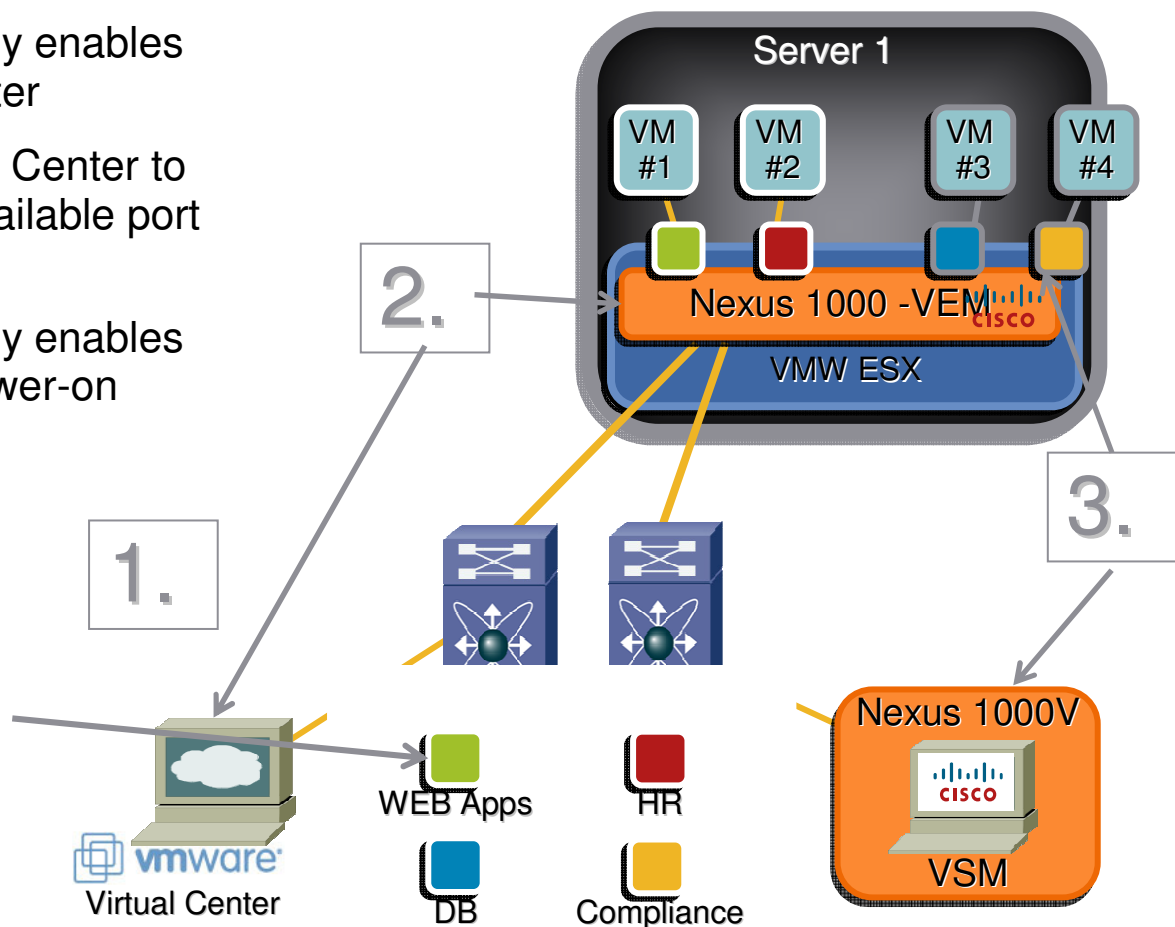


Policy Based VM Connectivity

Enabling Policy

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

- WEB Apps:**
- PVLAN 108, Isolated
 - Security Policy = Port 80 & 443
 - Rate Limit = 100 Mbps
 - QoS Priority = Medium
 - Remote Port Mirror = Yes



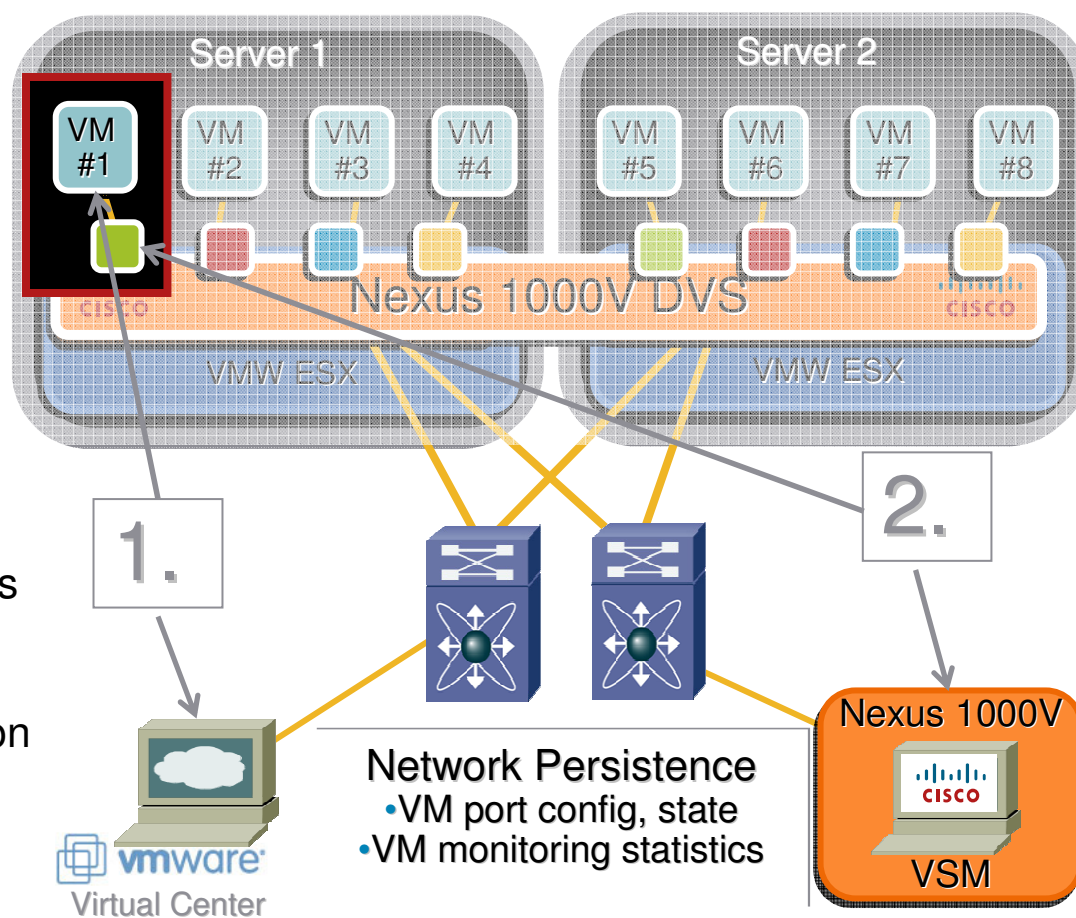
Mobility of Security & Network Properties

Following your VMs around

1. Virtual Center kicks off a Vmotion (manual/DRS) & notifies Nexus 1000V
2. During VM replication, Nexus 1000V copies VM port state to new host

Mobile Properties Include:

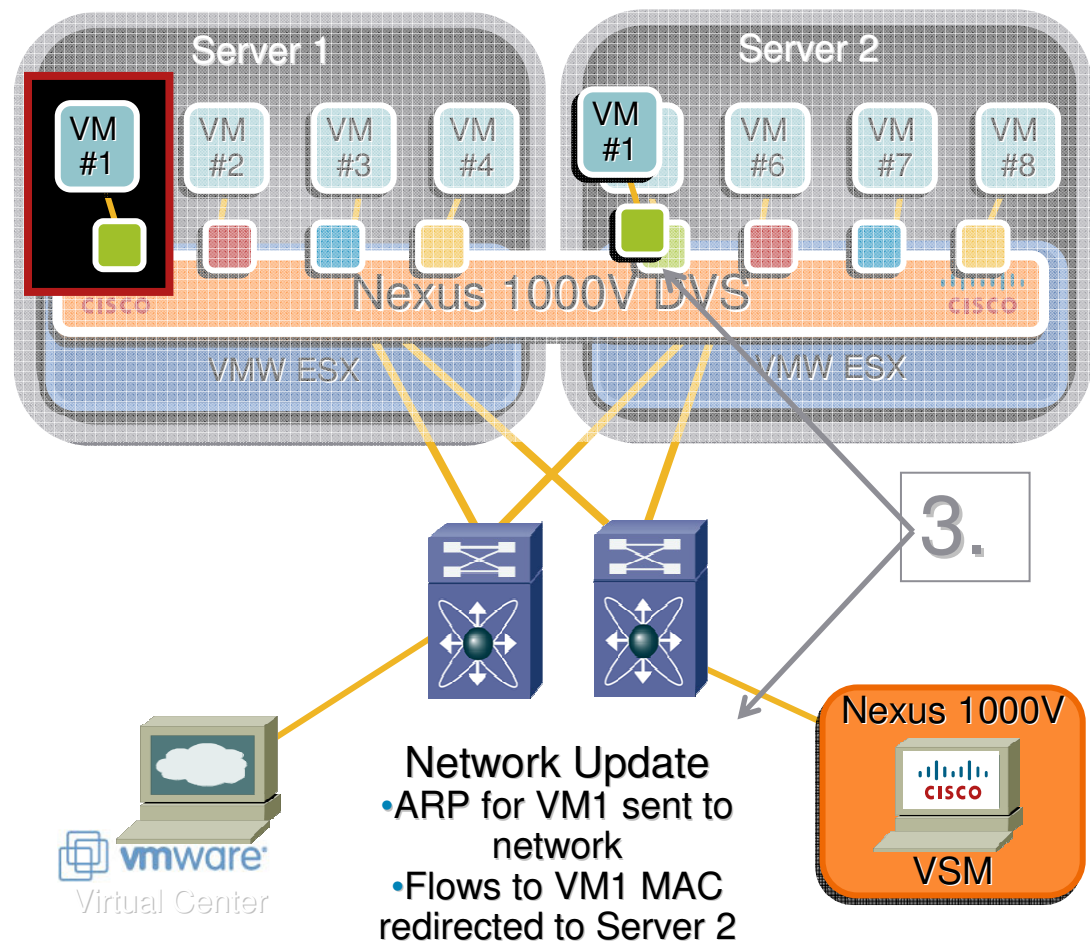
- Port Policy
- Interface State & Counters
 - Flow Statistics
- Remote Port Mirror Session



Mobility of Security & Network Properties

Following your VMs around


1. Virtual Center kicks off a Vmotion (manual/DRS) & notifies Nexus 1000V
2. During VM replication, Nexus 1000V copies VM port state to new host
3. Once VMotion completes, port on new ESX host is brought up & VM's MAC address is announced to the network



What Can A Profile Contain?

Policy definition supports:


- VLAN, PVLAN settings
- ACL, Port Security, ACL Redirect
- Cisco TrustSec (SGT)
- NetFlow Collection
- Rate Limiting
- QoS Marking (COS/DSCP)
- Remote Port Mirror (ERSPAN)



VN-Link with the Nexus 5000

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

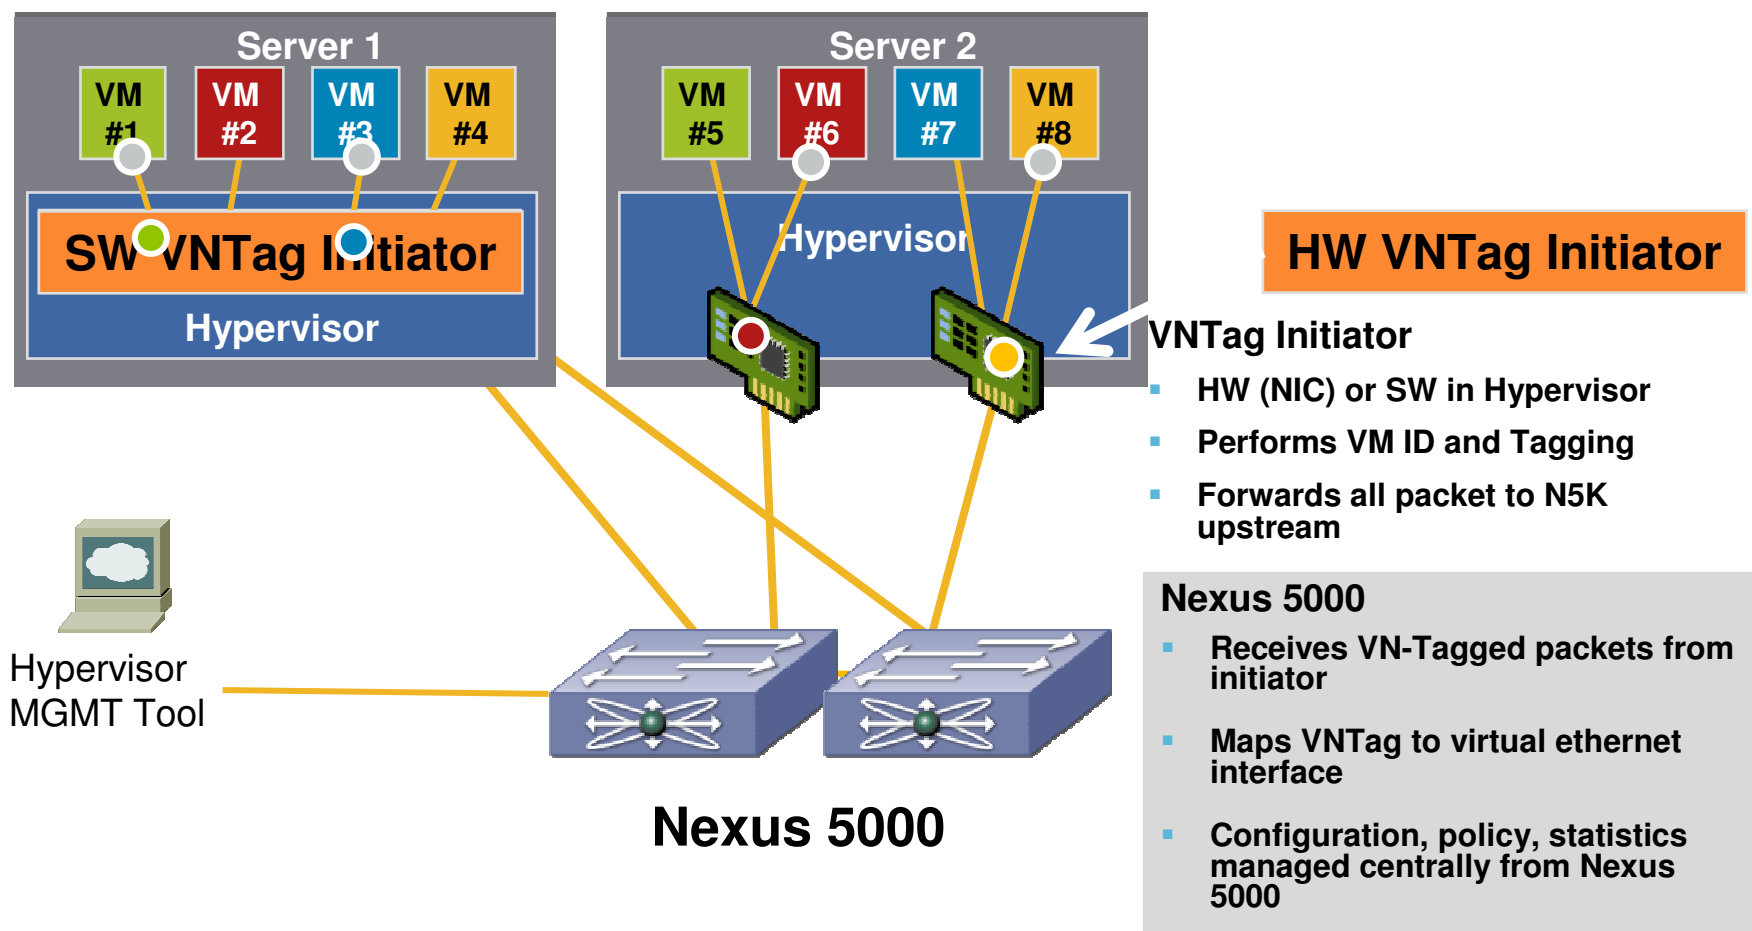


**Policy-Based
VM Connectivity**

**Mobility of Network
and Security Properties**


**Non-Disruptive
Operational Model**

Cisco Nexus 5000 VN-Link Architecture




VN-Link Deployment: Nexus 1000V

1. Works with all types of servers
2. Works with any type of switch
3. Requires External Management Appliance (VSM)
4. Requires Cisco SW running in the hypervisor (VEM)
5. Network stats, interface state, flow stats maintained in VEM, exposed through VSM



VN-Link Deployment: Nexus 5000


1. Focused on 10G attached servers
2. Requires the N5K upstream
3. Integrates virtual & physical interface management
4. Requires HW capable NIC or SW enabled hypervisor for host initiator (New NIC or N1K VEM)
5. Network stats, interface state, maintained in HW in the N5K




VN-Link

Complimentary Models for Evolving Requirements

Cisco Nexus 1000V (Software Based)



UCS VIC (Hardware Based)



**Policy-Based
VM Connectivity**

**Mobility of Network
and Security Properties**

**Non-Disruptive
Operational Model**


Unified Computing System

The Compute Platform for
Virtualized Data Centers




Industry Transformation in Motion

Consolidation of all platforms



Unified Computing

Harmonization efforts among the platforms



Unified Computing System

A single system that unifies

- Compute: Industry standard x86
- Network: Unified fabric
- Virtualization: Control, scale, performance
- Storage Access: Wire once for SAN, NAS, iSCSI


Embedded management

- Increase scalability without added complexity
- Dynamic resource provisioning
- Ability to integrate with broad partner ecosystem

Energy efficient

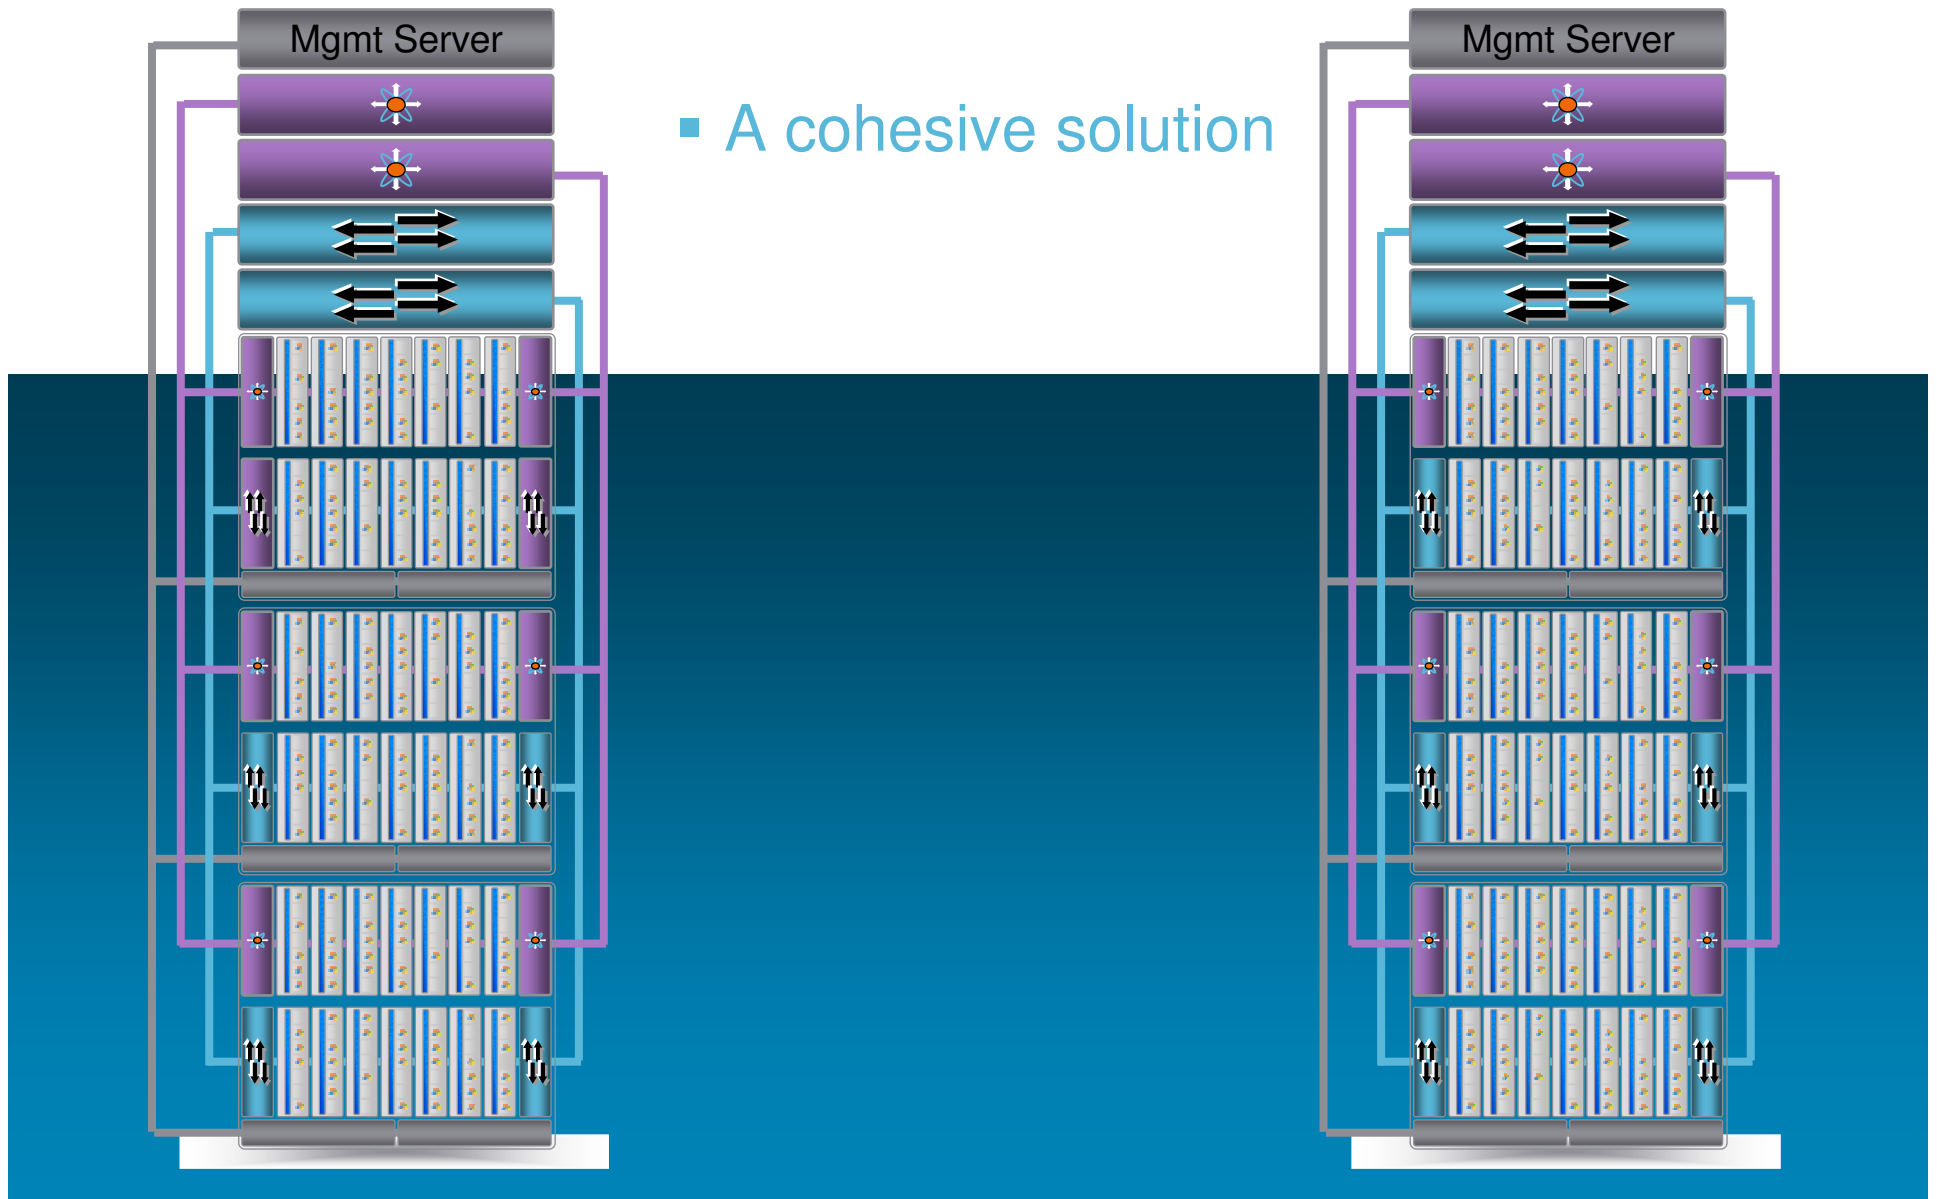
- Fewer servers, switches, adapters, cables
- Lower power and cooling requirements
- Increase compute efficiency by removing I/O and memory bottlenecks

Simplifying the Data Center




Simplifying the Data Center

- A cohesive solution





Simplifying the Data Center

- A cohesive solution
- Embed management




Simplifying the Data Center


- A cohesive solution
- Embed management
- Unify fabrics




Simplifying the Data Center





- A cohesive solution
- Embed management
- Unify fabrics
- Optimize virtualization



Simplifying the Data Center





- A cohesive solution
- Embed management
- Unify fabrics
- Optimize virtualization
- Remove unnecessary
Switches
Adapters
Management modules
- Less than 1/2 the support infrastructure for a given application




Cisco Unified Computing System

- UCS
Scale compute platform
Integrated virtualization
Natural aggregation point for network
- Unified embedded management




Our Solution: Unified Computing System

- Single, scalable integrated system
- Network + compute virtualization
- Dynamic resource provisioning



Overall System (Rear)



UCS Incorporates Nexus Technology

UCS Building Blocks

UCS Manager

Embedded in Fabric Interconnect



UCS Fabric Interconnect

20 Port 10Gb FCoE
40 Port 10Gb FCoE




UCS Fabric Extender

Logically part of Fabric Interconnect
Inserts into Blade Enclosure




UCS Blade Server Chassis

Flexible bay configurations
Logically part of Fabric Interconnect



UCS Blade Server

Different blade types
Mix blade types within enclosure



UCS Virtual Adapters

Three adapter options
Mix adapters within blade



Nexus Products



Nexus 5000
Unified Fabric



Nexus 2148
Fabric Extender



Nexus 1000V

CNAs with FCOE




Key Cisco Differentiation

- **Unified Fabric** – simplifies server and removes infrastructure
- **Embedded Management** – one management domain simplifies management framework
- **Large Memory Footprint** – Industry leading footprint drives up consolidation
- **Virtualization Adapter** – Improves performance and reduces NIC infrastructure
- **Service Profiles** – encompasses HW state for rapid deployment of hardware


Blade Overview

Half-width blade



4x the memory

Full-width blade



2x I/O bandwidth

Common Attributes

2 x Intel Nehalem-EP processors

2 x SAS hard drives (optional)

Blade Service processor

Blade and HDD hot plug support

Stateless blade design

10Gb CNA and 10GbE adapter options

Differences

Half-width blade

- 12 x DIMM slots
- 1 x dual port adapter

Full-width blade

- 48 x DIMM slots
- 2 x dual port adapters

Optimizing Memory with the Xeon 5500

Typical Memory




Fixed number of DIMMs can be addressed by the CPU



Intel Xeon 5500 Series

- 3 Channels per CPU
- Either
- 2 DIMMs @ 1066MHz
- 6/CPU, 12 per System
- 96GB Total Capacity
- Or
- 3 DIMMs @ 800MHz
- 9/CPU, 18 per System
- 144GB Total Capacity

Cisco UCS Memory



Each DIMM the CPU looks for is made of 4 standard DIMMs



Intel Xeon 5500 Series with UCS

- 3 Channels per CPU
- 8 DIMMs/channel @ 1066MHz
- 24/CPU, 48 per System
- 384GB Total Capacity

Benefit

- Higher memory capacity possible
- Lower costs for same capacity

Cisco UCS C-Series Rack-Mount Servers

Available in Q4 CY2009



Item	CPU	Size	Memory	Disks	Adaptor
UCS C250 M1 (memory intensive)	Intel Nehalem EP	2RU	48 DIMM 384 GB	8 SFF SAS/SATA Drives	5 PCIe
UCS C210 M1	Intel Nehalem EP	2RU	12 DIMM 96 GB	16 SFF SAS/SATA Drives	5 PCIe
UCS C200 M1	Intel Nehalem EP	1RU	12 DIMM 96GB	4 x 3.5" SAS/SATA Drives	2 PCIe

Adapter Offerings


Compatibility

Existing Driver Stacks

EMULEX


QLOGIC

10GbE/FCoE



Cost

“Free” SAN Access for any Ethernet equipped host




Software FCoE

Virtualization

VM I/O Virtualization and Consolidation
CISCO


10GbE/FCoE



Virtualization Adapter

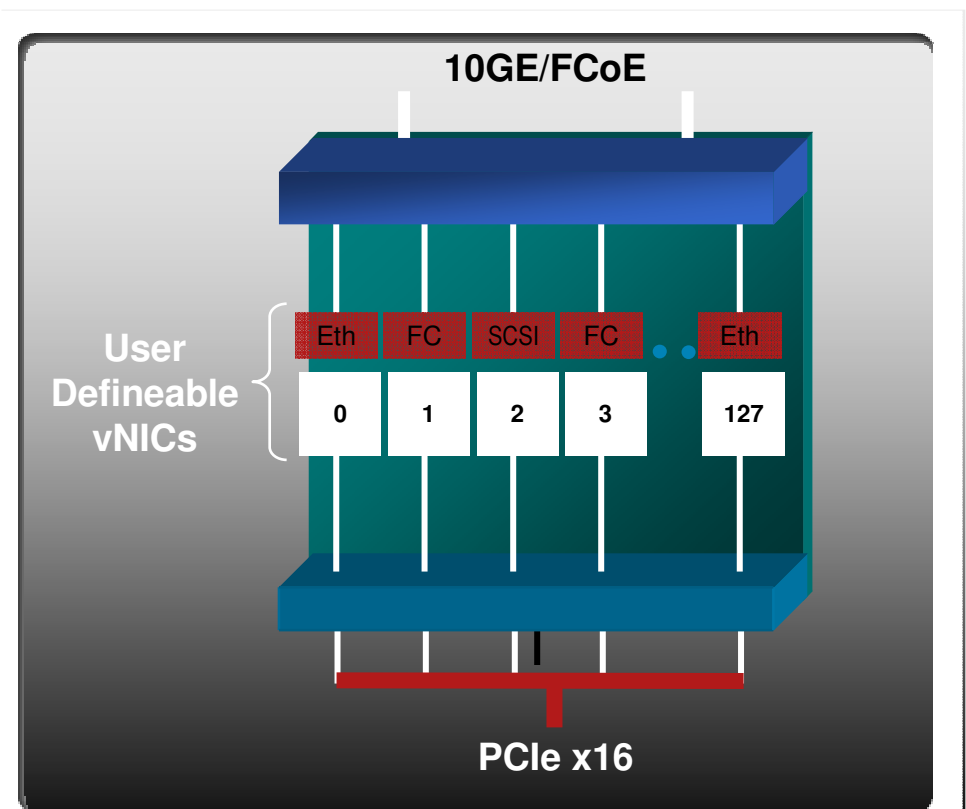
True wire once architecture – highly dynamic
Network policy and visibility brought to VMs

Hypervisor bypass (Pass-through) – increases performance
Reduce NIC and mezz card infrastructure



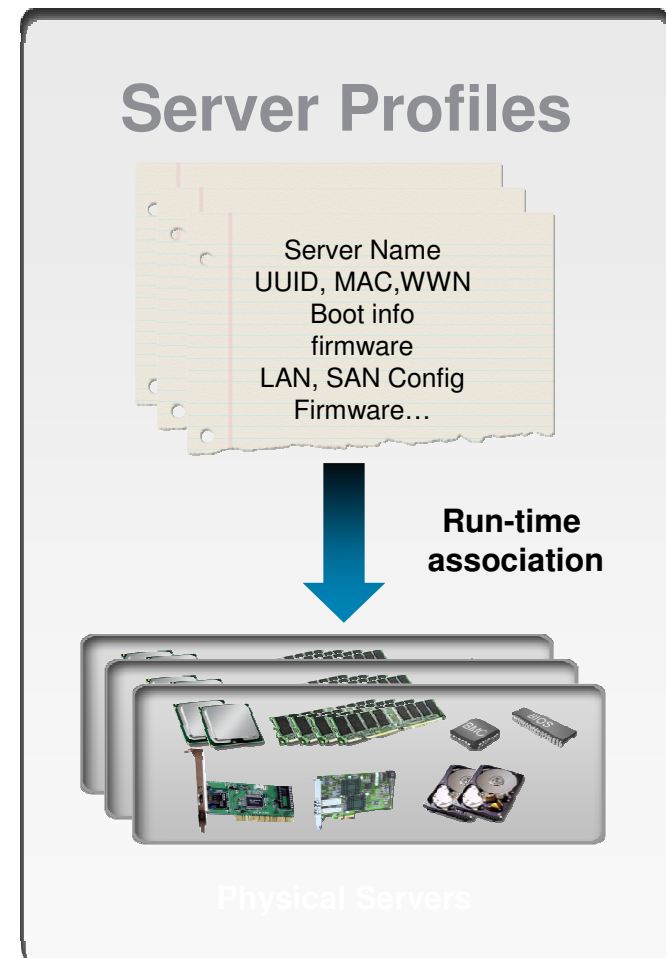
Cisco UCS Virtualized Adapter

- Virtualized adapter designed for both single-OS and VM-based deployments
- Provides mobility, isolation, and management from the network
 - Secure
 - Transparent to hosts
- Cut-through architecture
- High Performance
 - 2x 10Gb
 - Low latency
 - High BW IPC support
- 128 vNICs**
 - Ethernet, FC or SCSI
 - 500K IOPS
 - Initiator and Target mode




Dynamic Management

- Server profiles
 - Abstracts server characteristics from the physical server hardware
- Pre-defined and pre-created server identities
 - Default is shipped hardware
 - Stored in switch
- “Associated” with a physical server
 - Manual or policy-driven



Service Profile



•Storage

- Optional Disk usage
 - SAN settings
 - LUNs
- Persistent Binding
 - SAN settings
 - vSAN
- Firmware
 - Revisions

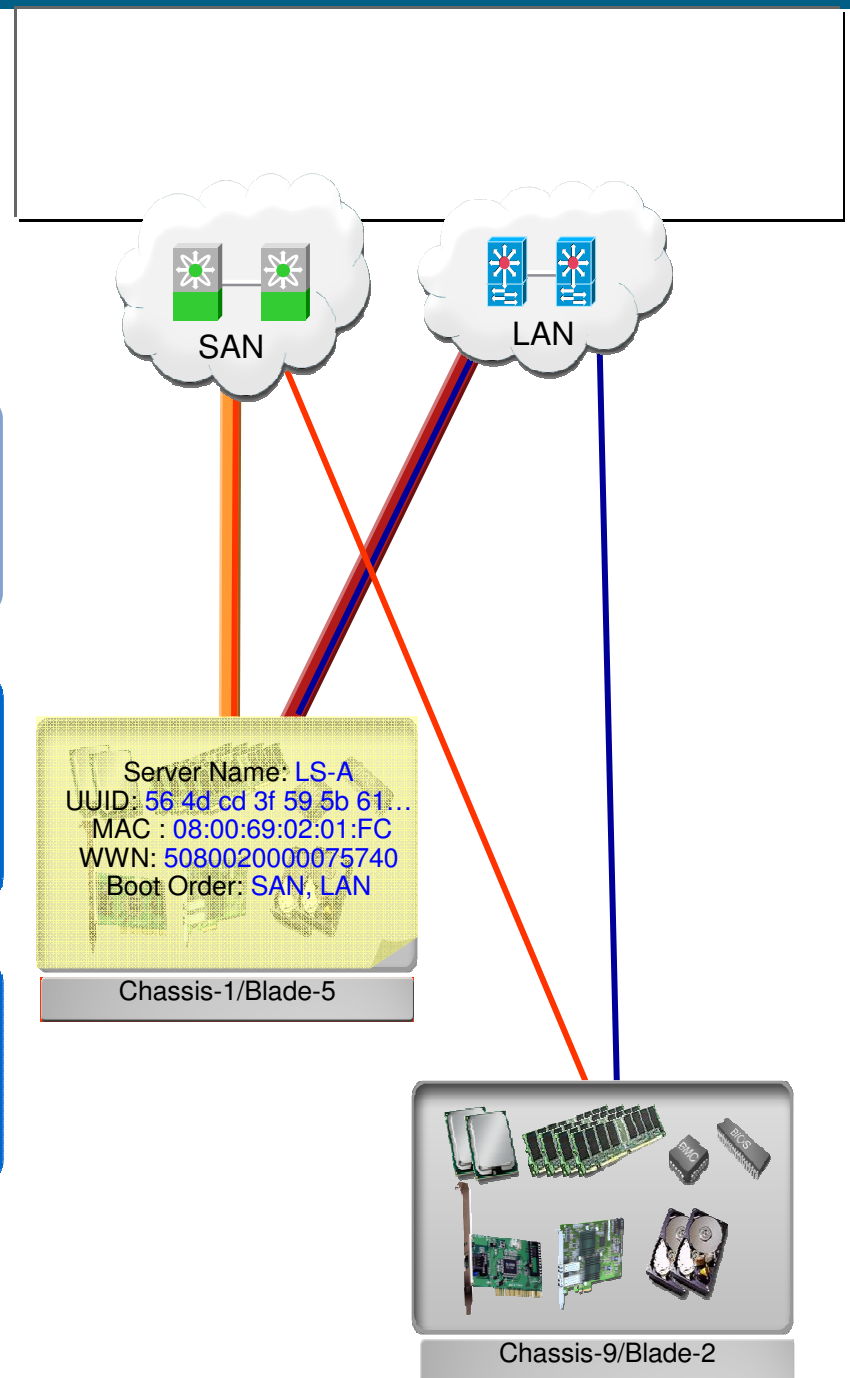

■Server

Identity (UUID)
 Adapters
 Number
 Type: FC, Ethernet
 Identity
 Characteristics
 Firmware
 Revisions
 Configuration settings

■Network

Uplinks
 LAN settings
 vLAN
 QoS
 etc...
 Firmware
 Revisions

Service Profiles in action

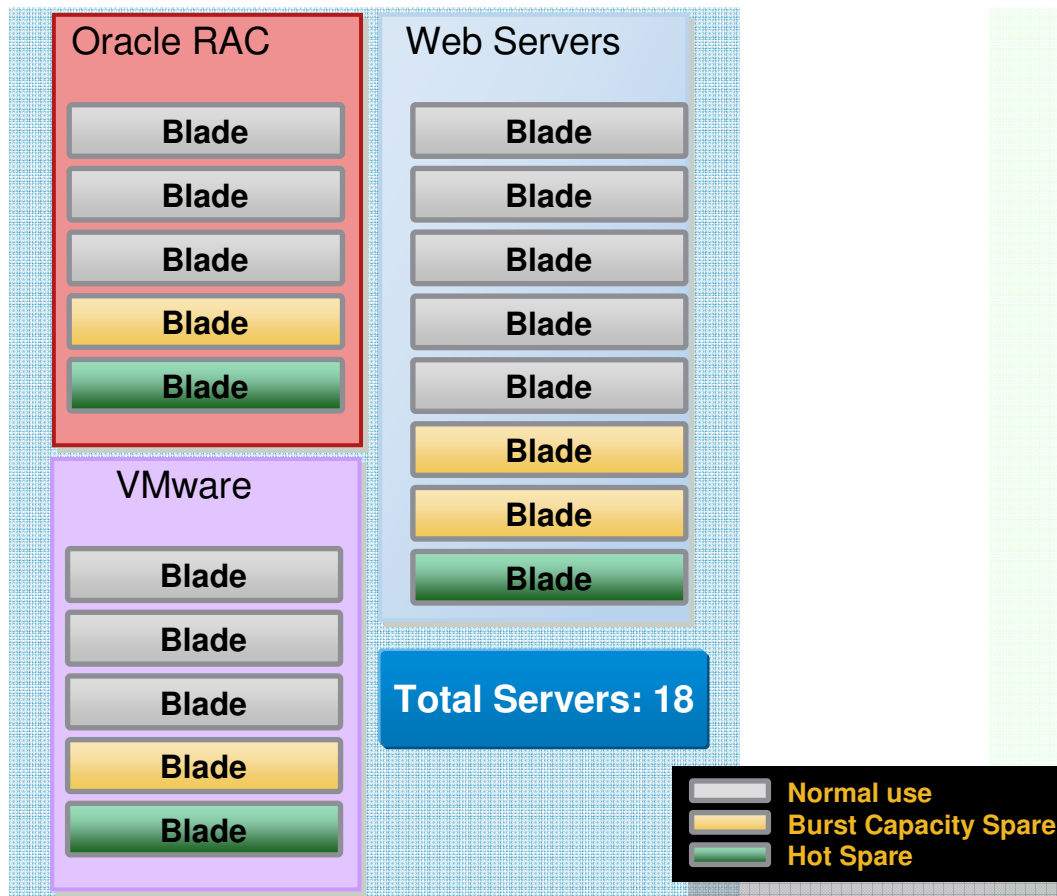


Service Profiles - Reduce Overall Server CAPEX

Today's Deployment:

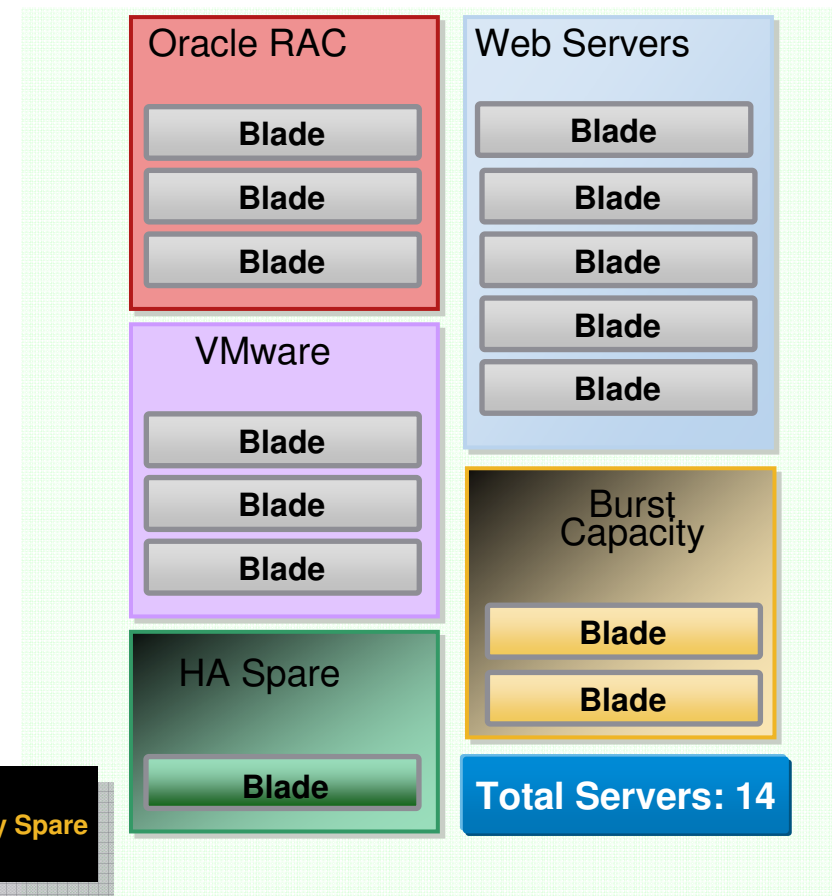
Provisioned for peak capacity

Spare node per workload



With Server Profiles:

- Resources provisioned as needed
- Same availability with fewer spares






Architecture Design







Physical Infrastructure and Network Topology

Mapping the Physical to the Logical




DAF POD



Distributed Access Fabric in Blade Environment

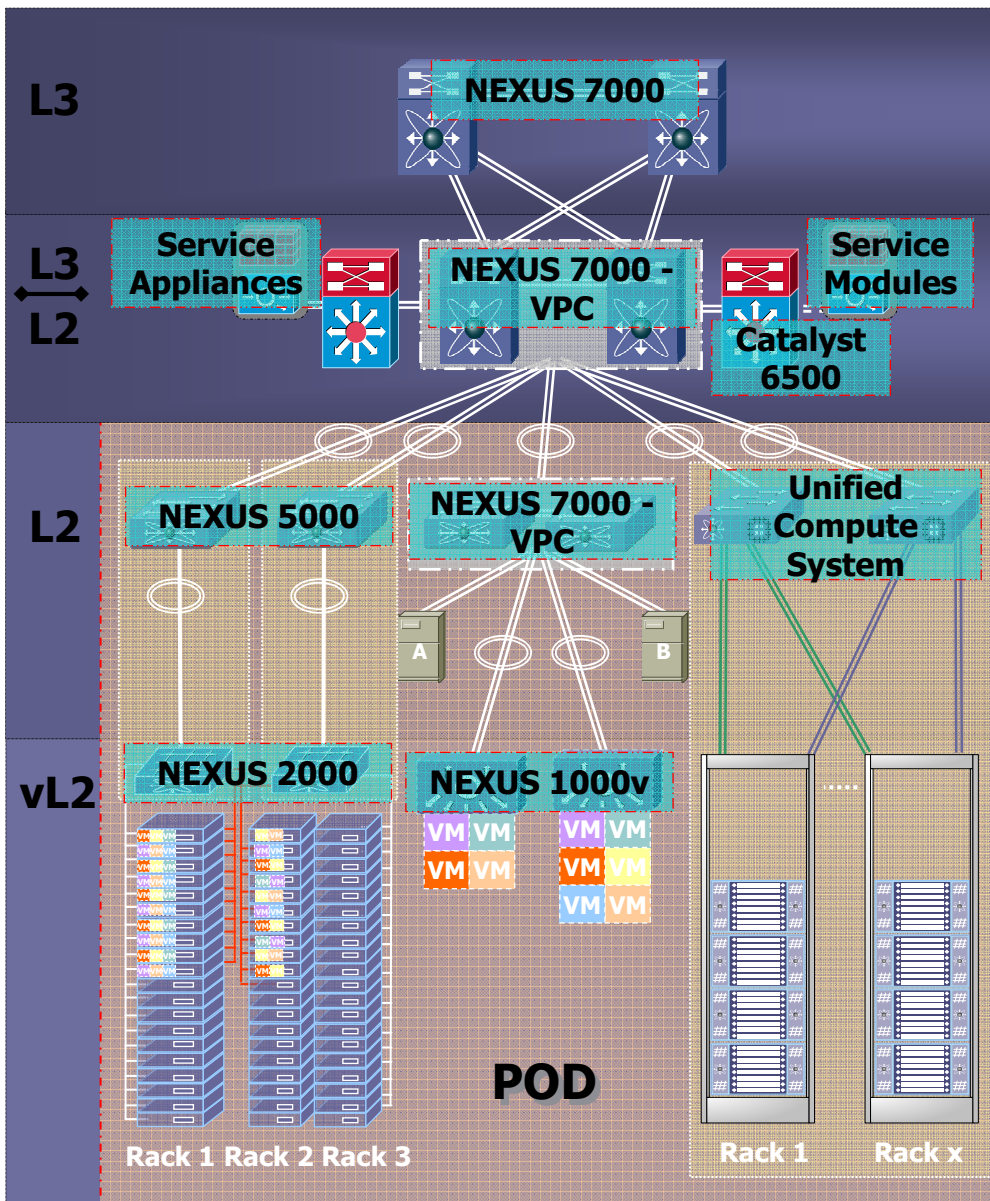
Distributed Access Fabric - DAF



Why Distributed Access Fabric?

- All chassis are managed by Fabric Interconnect: single config point, single monitoring point
Fabric instances per chassis present at rack level – reduced management point
Fabric instances are extensions of the fabric Interconnect: they are fabric Extenders
- Simplifies cabling infrastructure: Horizontal cabling choice: “what is available” Fiber or copper
CX1 cabling for brownfield installations – Fabric Interconnect Centrally located
USR for greenfield installations – Fabric Interconnect at End of Row near the cross connect
Vertical cabling: just an in-rack patch cable

The Unified DC Architecture



Core: L3 boundary to the DC network. Functional point for route summarization, the injection of default routes and termination of segmented virtual transport networks

Aggregation: Typical L3/L2 boundary. DC aggregation point for uplink and DC services offering key features: VPC, VDC, 10GE density and 1st point of migration to 40GE and 100GE

Access: Classic network layer providing non-blocking paths to servers & IP storage devices through VPC. It leverages Distributed Access Fabric Model (DAF) to centralize config & mgmt and ease horizontal cabling demands related to 1G and 10GE server environments

Virtual Access: A virtual layer of network intelligence offering access layer-like controls to extend traditional visibility, flexibility and mgmt into virtual server environments. Virtual network switches bring access layer switching capabilities to virtual servers without burden of topology control plane protocols. Virtual Adapters provide granular control over virtual and physical server IO resources

