



Cisco Expo  
2010

# Data Center Virtualization Power Session



# Agenda

## **Cisco Unified Computing System**

- Mihai Mincu – Cisco Romania
- Ciprian Tone – Cisco Romania

## **EMC Solutions for VCE**

- Ionut Lungu – EMC Romania
- Dumitru Taraianu – EMC Romania
- Puiu Leontescu – EMC Romania

## **Nexus 1000V – Distributed Virtual Switch**

- Mihai Mincu – Cisco Romania

## **Nexus 5000 & 2000 with Unified Fabric**

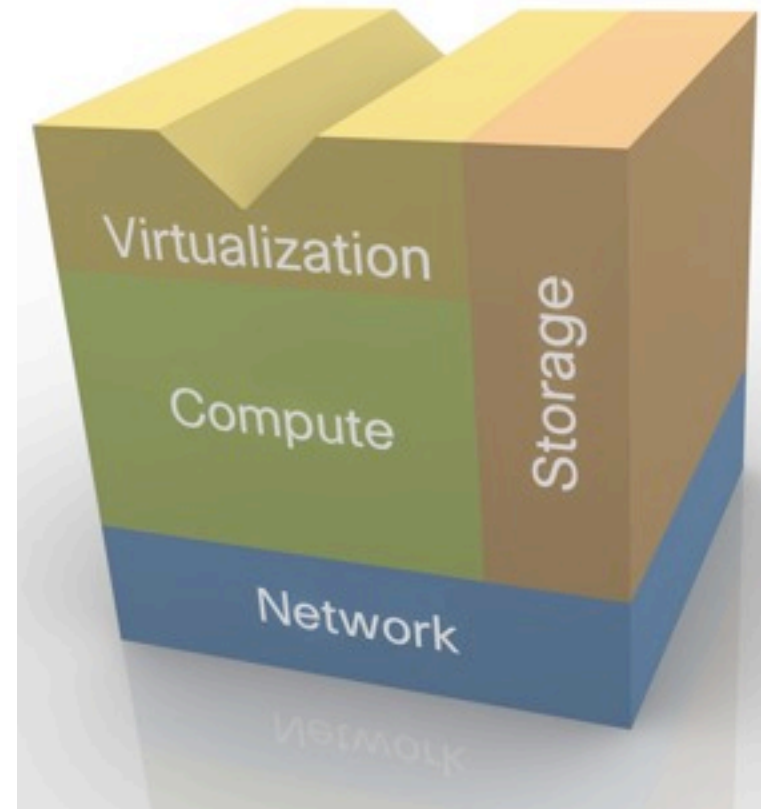
- Ciprian Tone – Cisco Romania



## Vblock Infrastructure Packages A New Way of Delivering IT

### Benefits:

- Accelerate the journey to pervasive virtualization and private cloud computing while lowering risk and operating expenses
- Pre-integrated and validated solutions reduce total cost of ownership
- Ensure security and minimize risk with certification paths
- Support and manage Service Level Agreements
  - Resource metering and reporting
  - Configuration and provisioning
  - Resource utilization
- Vblock is a validated platform that enables seamless extension of the environment



**Secure, Extensible, SLA-driven, Infrastructure**

# vBlock Infrastructure Packages

Definitions for a Virtualized Computing Environment

**Modular, Standardized, Scalable**

Three solution offerings addressing different market segments

- **Type 0- vBlock for SMB & Commercial**

- Positioned as an entry point small size DC and for POC testing
- Available 1HCY10

- **Type 1- vBlock for Commercial and Enterprise**

- Optimized for cost and positioned to support medium intensity application environments
- A mid-sized configuration - broad range of IT capabilities for organizations of all sizes
- Typical use case:  
Shared services – Email, File & Print, Virtual Desktops, etc  
Business critical ERP, CRM systems

- **Type 2- vBlock for Large Enterprise and Service Provider**

- Optimized for performance and positioned to support high intensity application environments - extensible to meet the most demanding IT needs
- Typical Use case:  
Shared services – Email, File & Print, Virtual Desktops, etc  
Business critical ERP, CRM systems

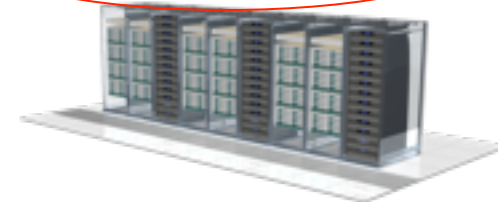
**Type-0 vBlock  
SMB & Commercial**



**Type-1 vBlock  
Commercial, Enterprise**



**Type-2 vBlock  
Large Enterprise and SP**



# Vblock X Components

- Compute
  - Cisco UCS B-series
- Network
  - Cisco Nexus 1000V
  - Cisco MDS 9000 Series
- Storage
  - EMC CLARiiON CX4
  - EMC Symmetrix V-Max
- Hypervisor
  - VMware vSphere 4
- Management
  - EMC Ionix Unified Infrastructure Manager (Optional)
  - VMware vCenter
  - EMC NaviSphere
  - EMC PowerPath
  - Cisco UCS Manager
  - Cisco Fabric Manager



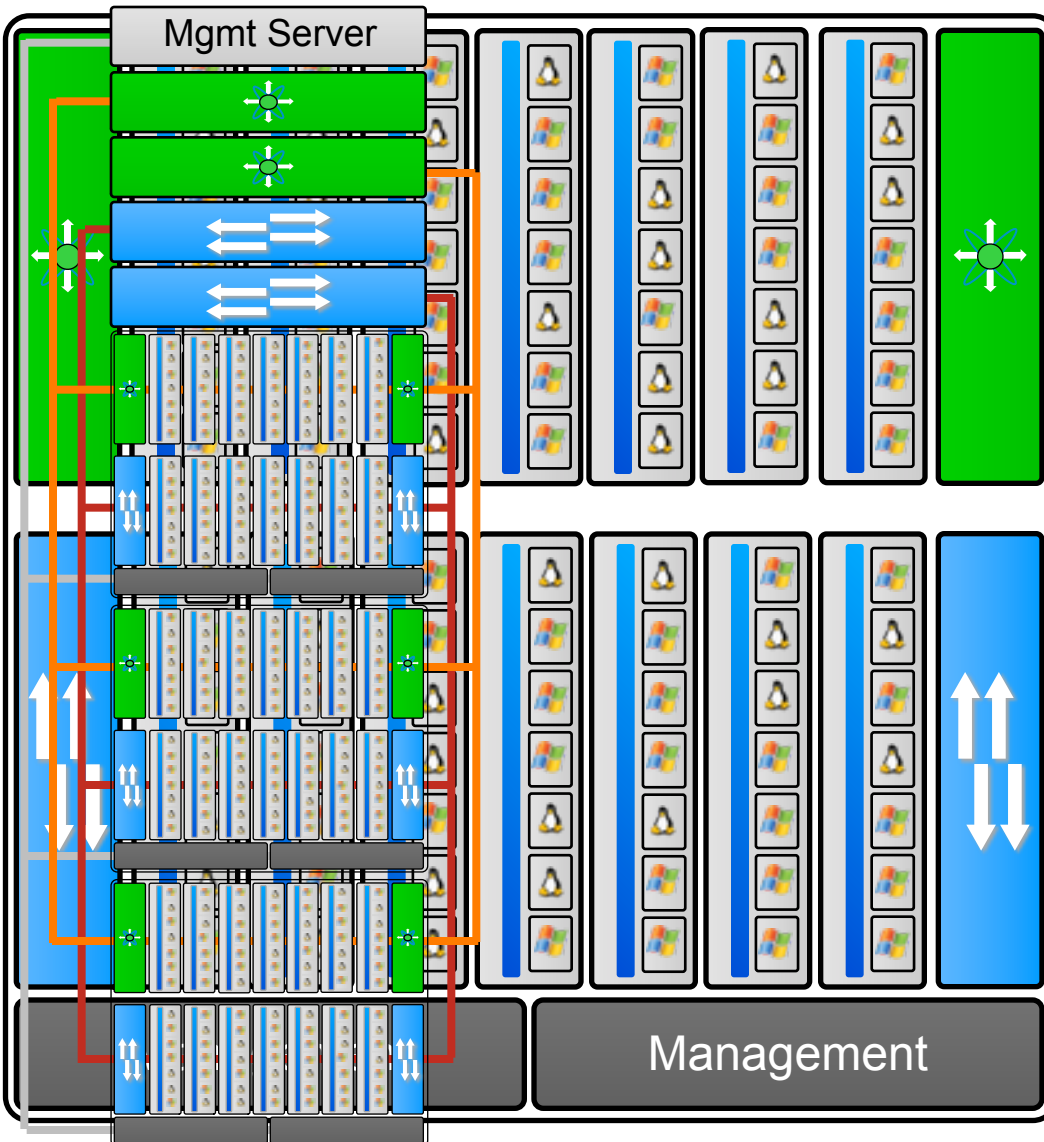


Cisco Expo  
2010

# Cisco Unified Computing System



# Server Deployment Today



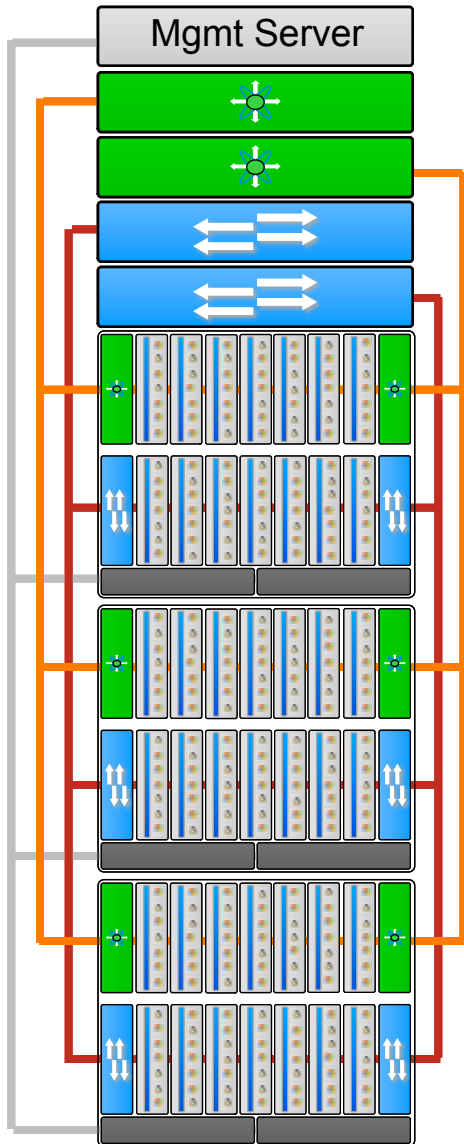
## Over the past 10 years

- An evolution of size, not system
- More servers & switches than ever
- More switches per server
- Management applied, not integrated

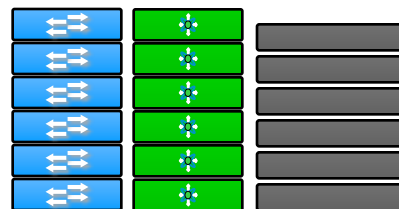
## Result

- More points of management
- More difficult to maintain policy coherence
- More difficult to secure
- More difficult to scale

# Our Solution



- Embed management
- Unify fabrics
- Optimize virtualization
- Remove unnecessary switches, adapters, management modules
- Less than 1/3rd infrastructure



# Cisco Unified Computing Solution

A single system that encompasses:

- Network: Unified fabric

- Compute: Industry standard x86

- Storage: Access options

- Virtualization optimized

Unified management model

- Dynamic resource provisioning

Efficient Scale

- Cisco network scale & services

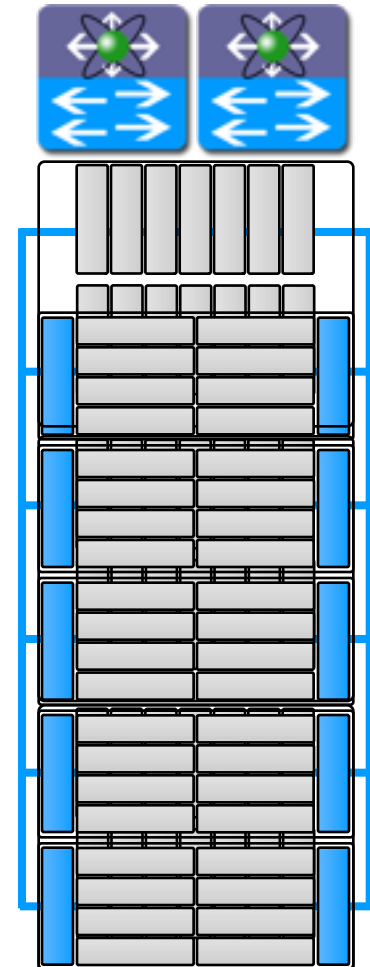
- Fewer servers with more memory

Lower cost

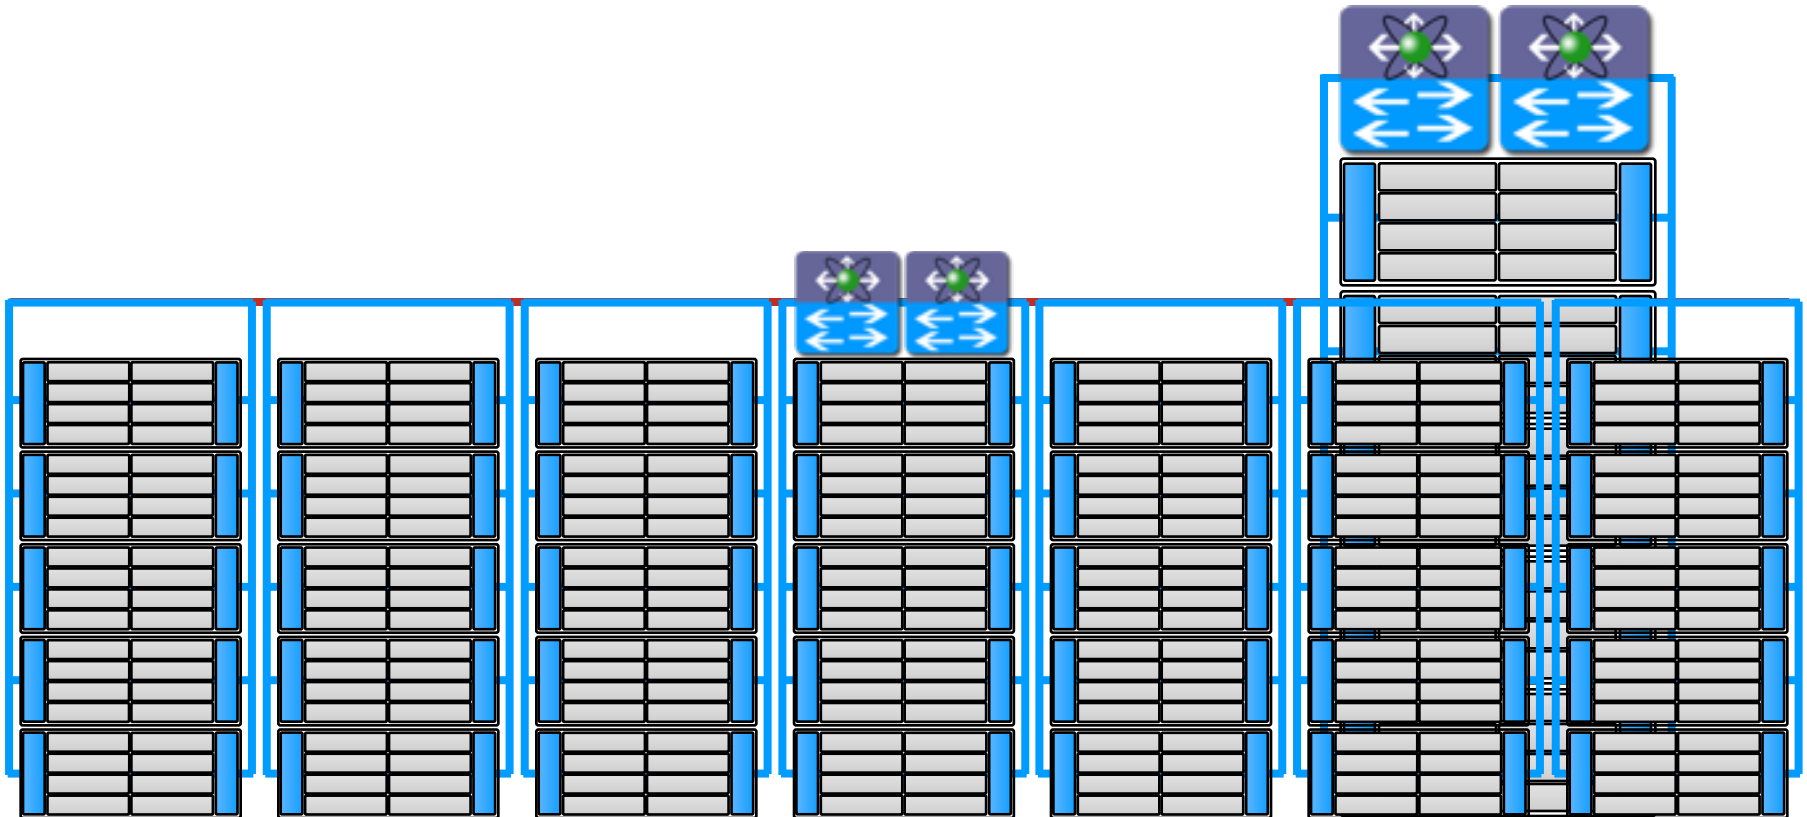
- Fewer servers, switches, adapters, cables

- Lower power consumption

- Fewer points of management



# Cisco Unified Computing Solution

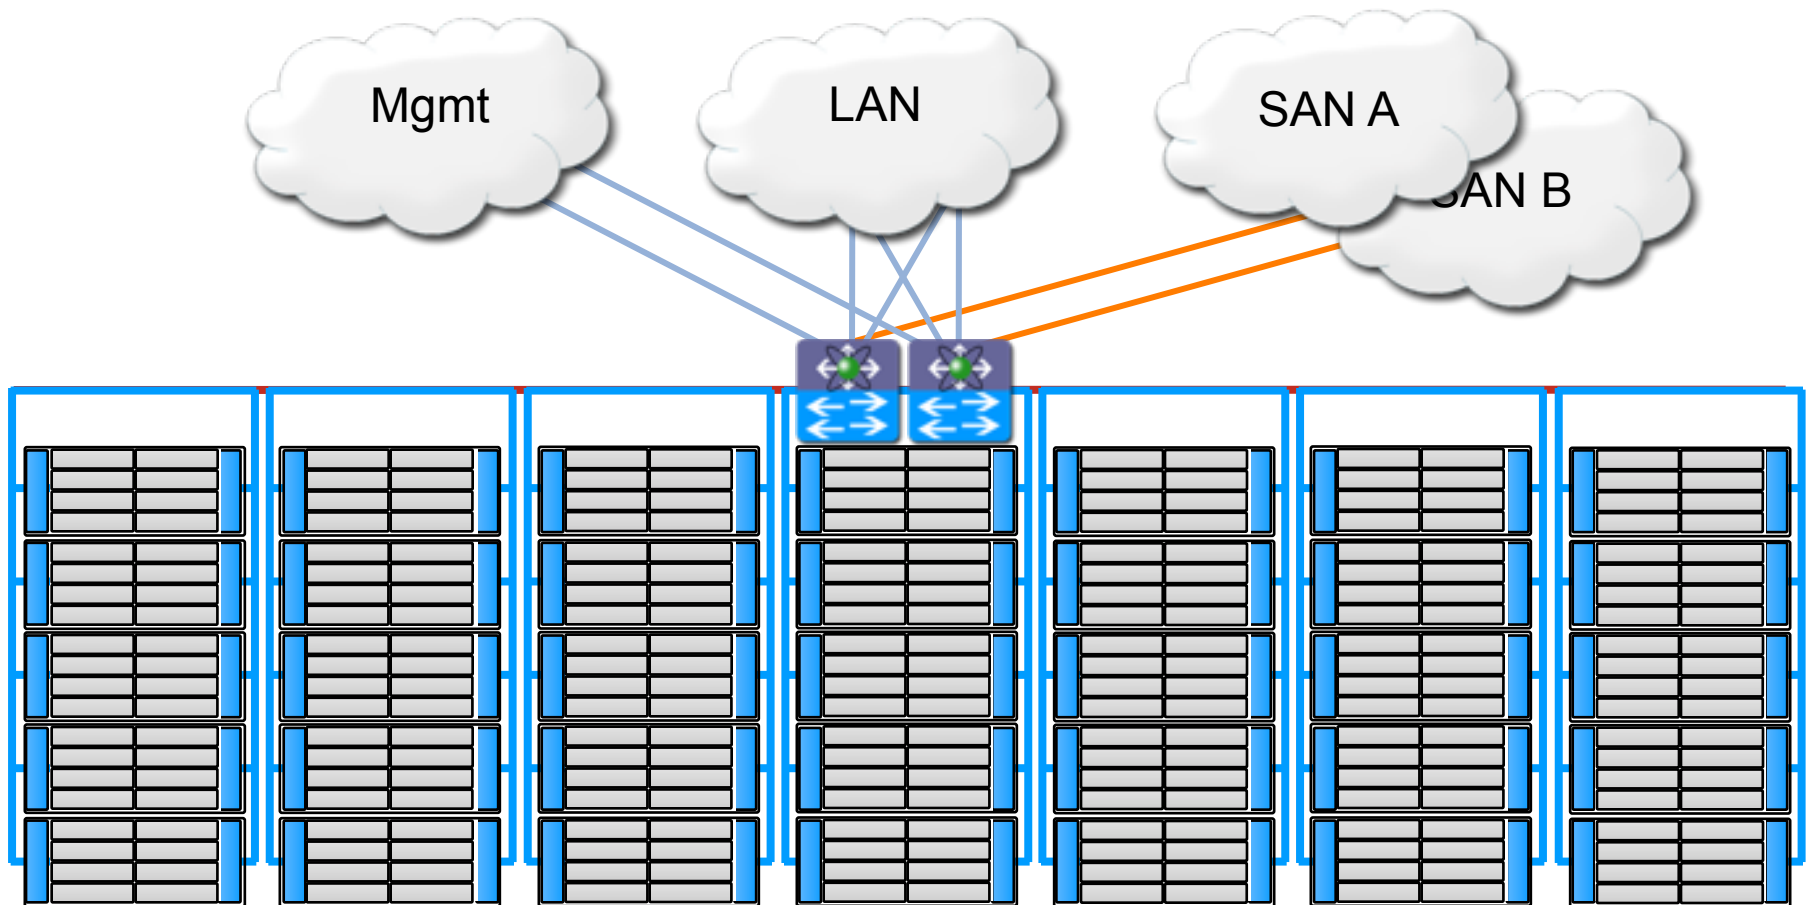


# Cisco Unified Computing Solution

Single, scalable integrated system

Network + compute virtualization

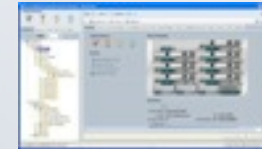
Dynamic resource provisioning



# Building Blocks

## **UCS Manager**

Embedded– manages entire system



## **UCS Fabric Interconnect**

20 Port 10Gb FCoE

40 Port 10Gb FCoE



## **UCS Fabric Extender**

Remote line card



## **UCS Blade Server Chassis**

Flexible bay configurations



## **UCS Blade Server**

Industry-standard architecture

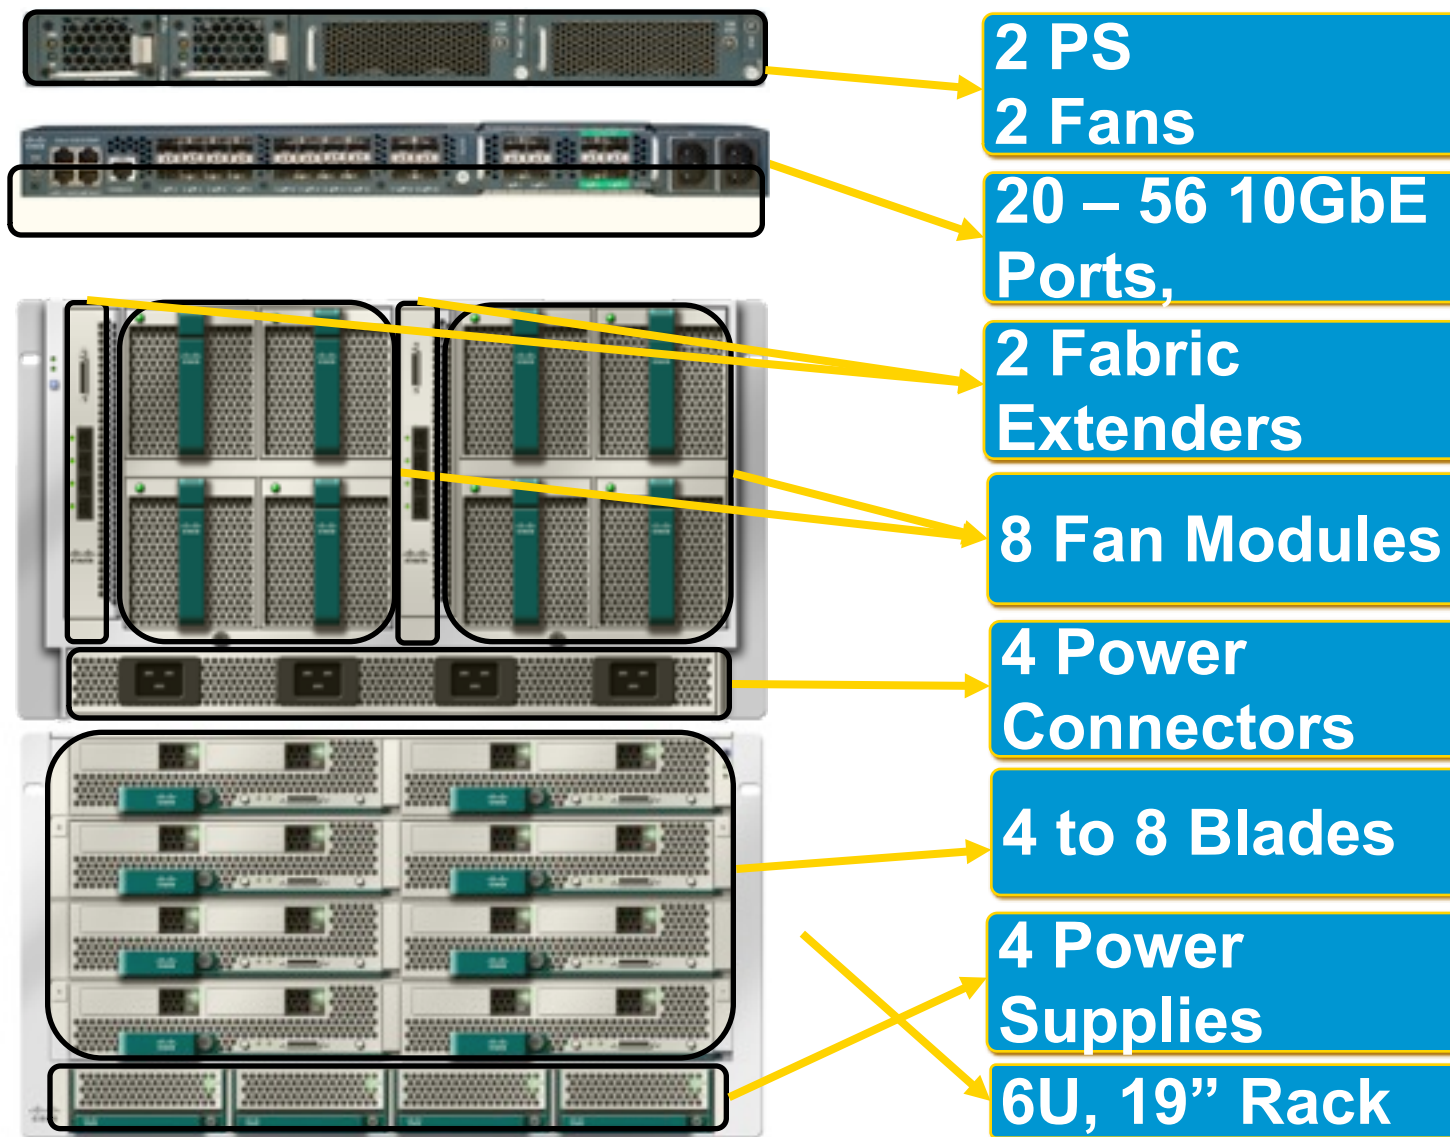


## **UCS Virtual Adapters**

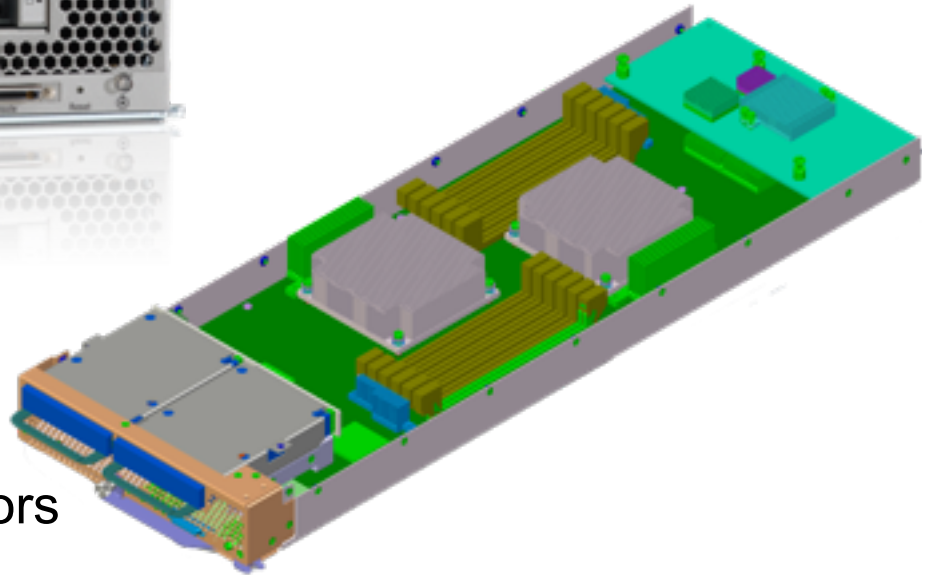
Choice of multiple adapters



# Simplified Infrastructure



# UCS B200 M1 Blade

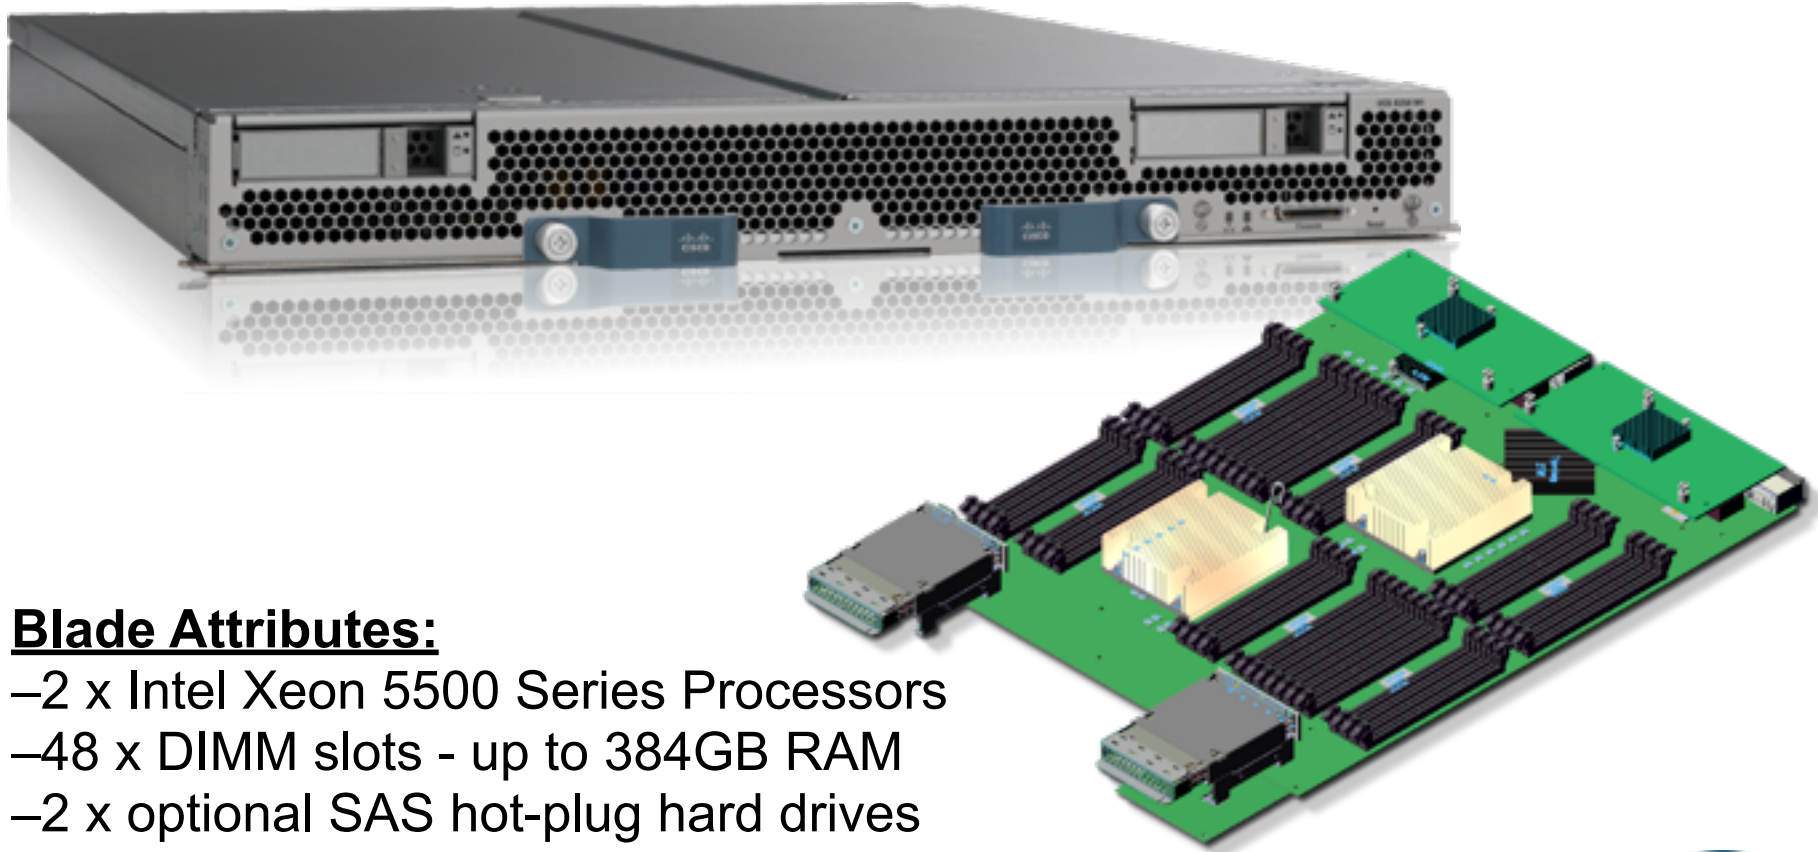


## **Blade Attributes:**

- 2 x Intel Xeon 5500 Series Processors
- 12 x DIMM slots - up to 96GB RAM
- 2 x optional SAS hot-plug hard drives
- RAID 0, 1, 0+1
- 1 x 10Gb dual port mezzanine adapter
- Remote and local access to keyboard, video, mouse, serial
- Integrated with UCS Manager



# UCS B250 M1 Blade



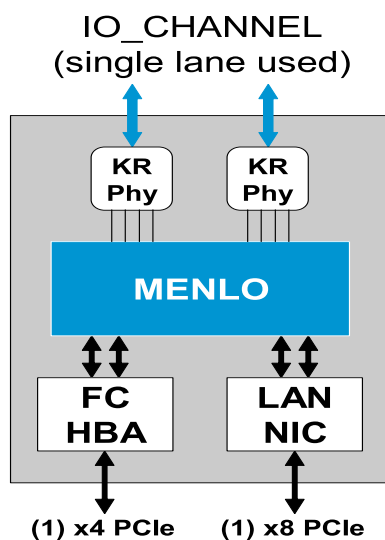
## **Blade Attributes:**

- 2 x Intel Xeon 5500 Series Processors
- 48 x DIMM slots - up to 384GB RAM
- 2 x optional SAS hot-plug hard drives
- RAID 0, 1
- 2 x 10Gb dual port mezzanine adapter
- Remote and local access to keyboard, video, mouse, serial
- Integrated with UCS Manager

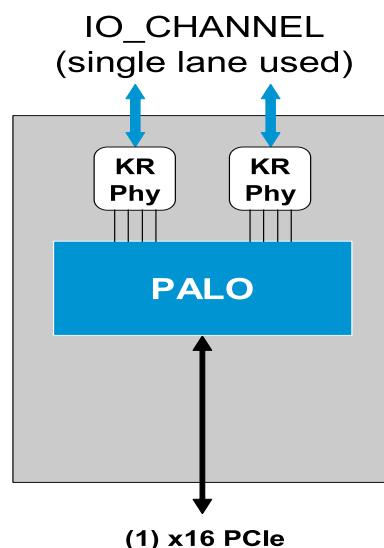


# Mezzanine Cards (Server I/O Options)

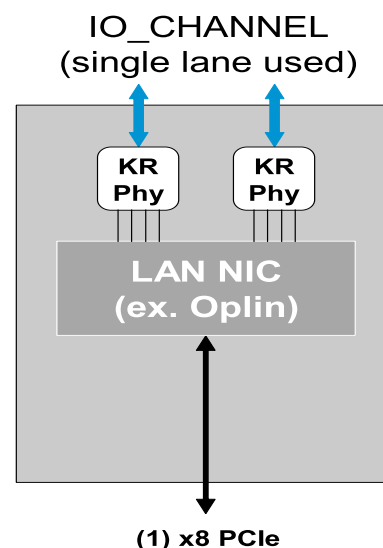
- Menlo
  - Intel 82598 (Oplin)
  - Qlogic ISP2432 or Emulex Zephyr
- PALO
  - Instantiates Fibre Channel, Ethernet/LAN, virtual NICs  
(in conjunction with NIV support from Central Switch )
- Oplin (Intel 82598)
  - 3<sup>rd</sup> party based



***MENLO based  
MEZZ\_CARD***

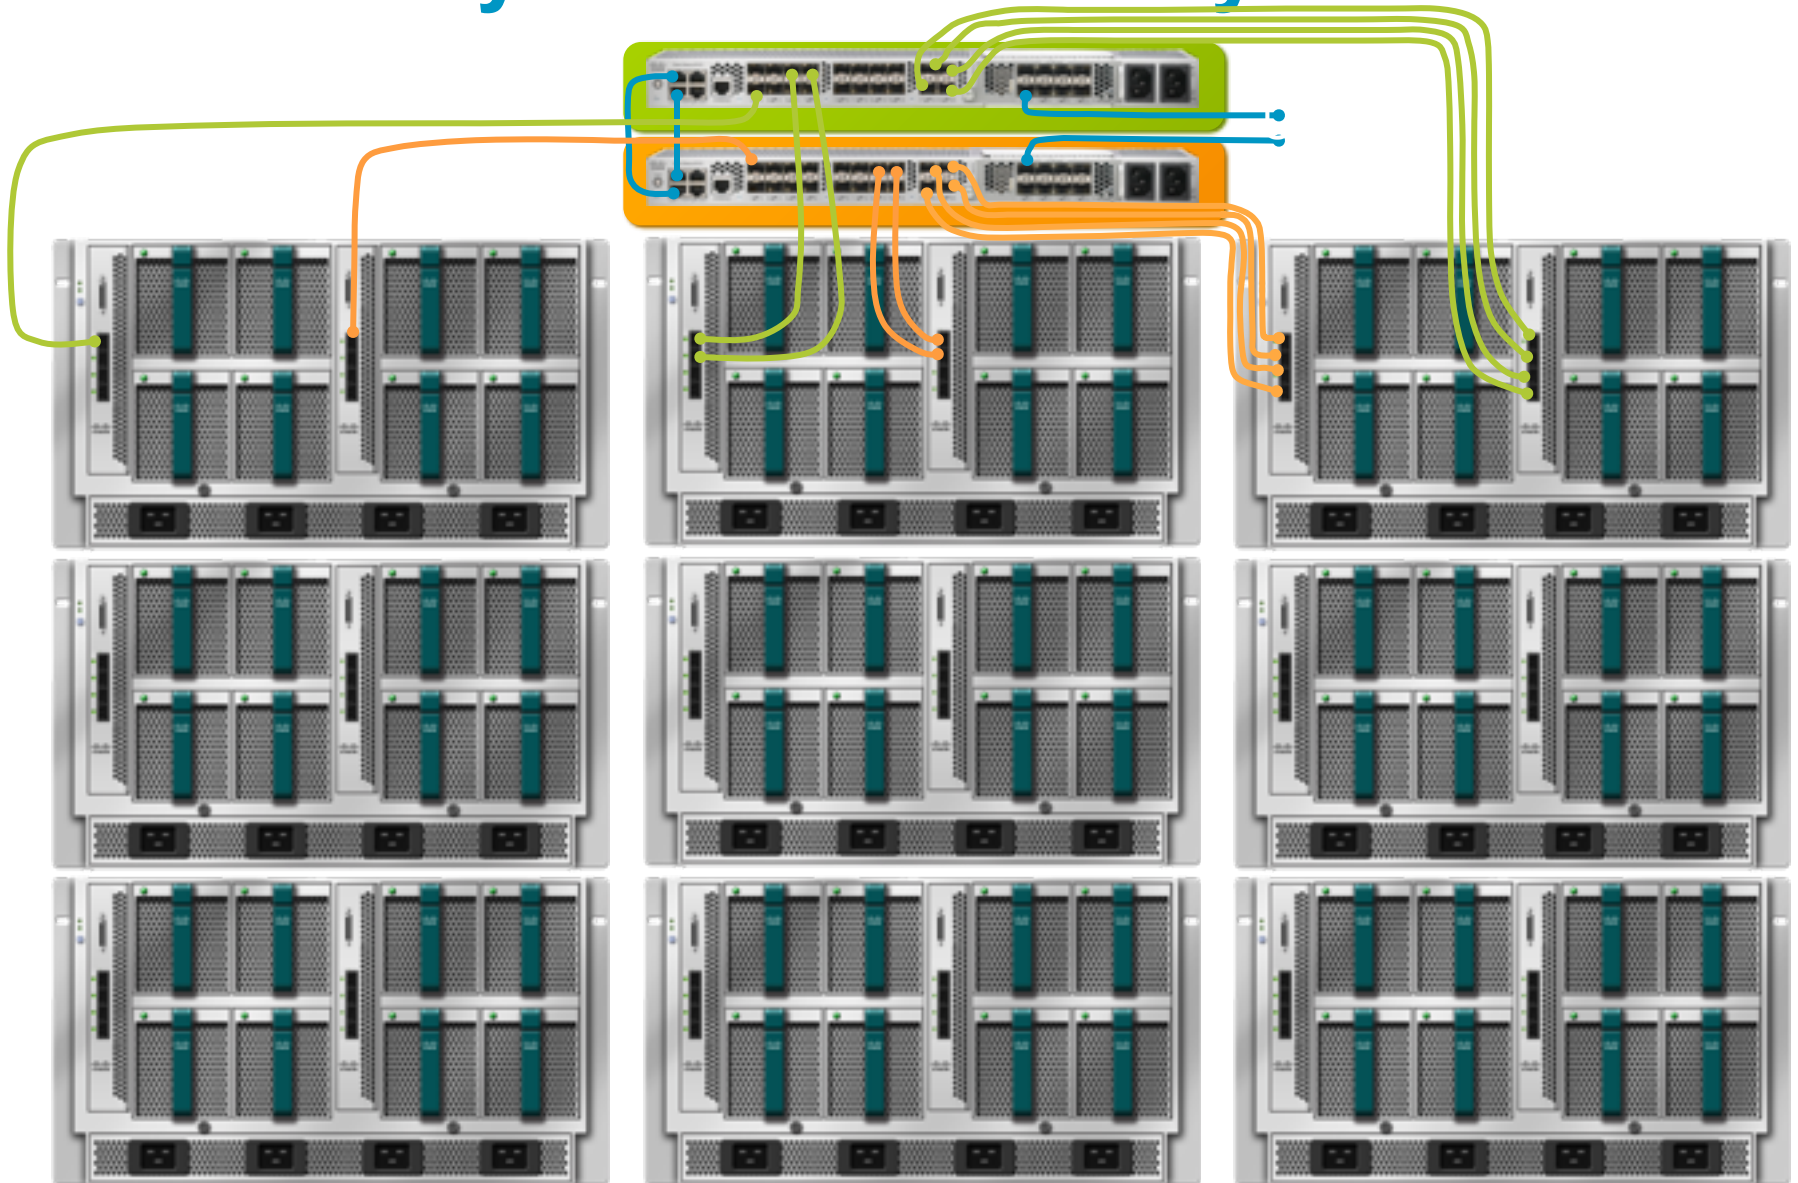


***PALO based  
MEZZ\_CARD***



***3<sup>rd</sup> party based  
MEZZ\_CARD***

# Overall System Connectivity



# UCS Key Value Propositions:

## *Drivers for use cases*



Hardware State Abstraction – Service Profiles

Unified Fabric - FCOE

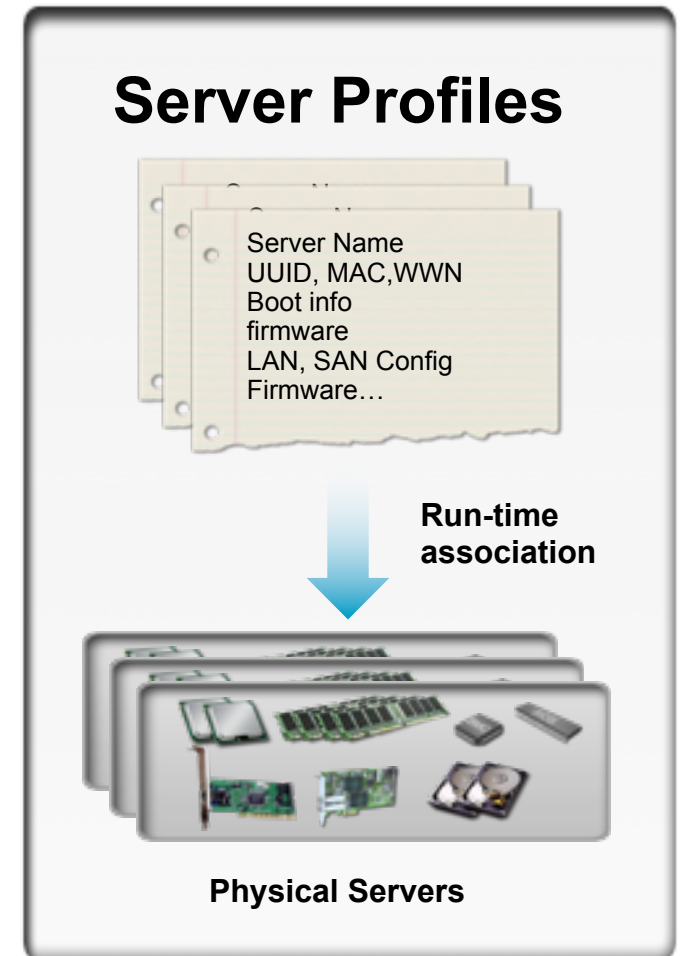
Virtualized Adapter - Palo

Expanded Memory Server – Full Height Blade

Unified Management - CAM

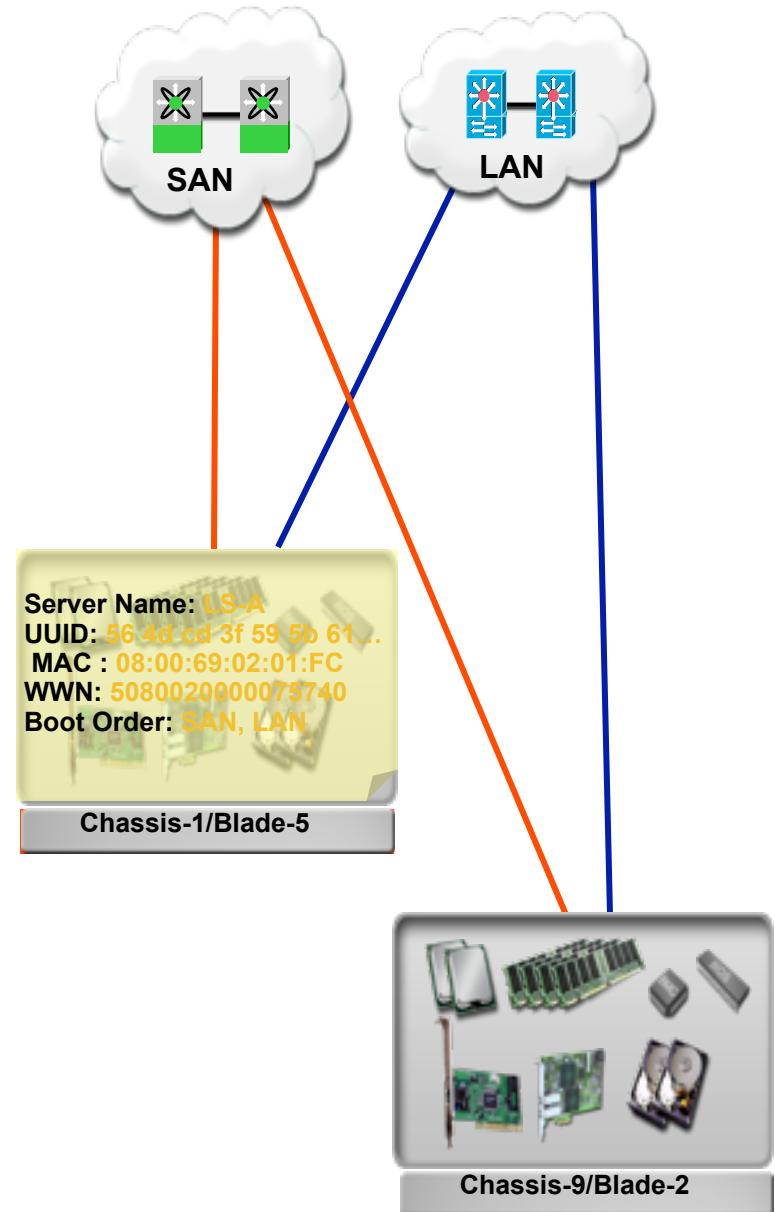
# Dynamic Management

- 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



# Stateless Computing

- Server attributes no longer tied to physical hardware
  - Not just identity
  - Seamless server mobility
  - Within switch domain
- Network boot (LAN or SAN)
  - Boot order and devices are part of server profile
  - Local disks can be used for temp, swap, etc.
    - Scrubbed between use (optional)



# UCS Key Value Propositions:

## *Drivers for use cases*

Hardware State Abstraction – Service Profiles



Unified Fabric - FCOE

Virtualized Adapter - Palo

Expanded Memory Server – Full Height Blade

Unified Management - CAM

# UCS Key Value Propositions:

*Drivers for use cases*

Hardware State Abstraction – Service Profiles

Unified Fabric - FCOE



Virtualized Adapter - Palo

Expanded Memory Server – Full Height Blade

Unified Management - CAM

# Cisco UCS VIC Overview

## Mezzanine Card for B-Series

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

- Virtualize in Hardware
- PCIe compliant

## High Performance

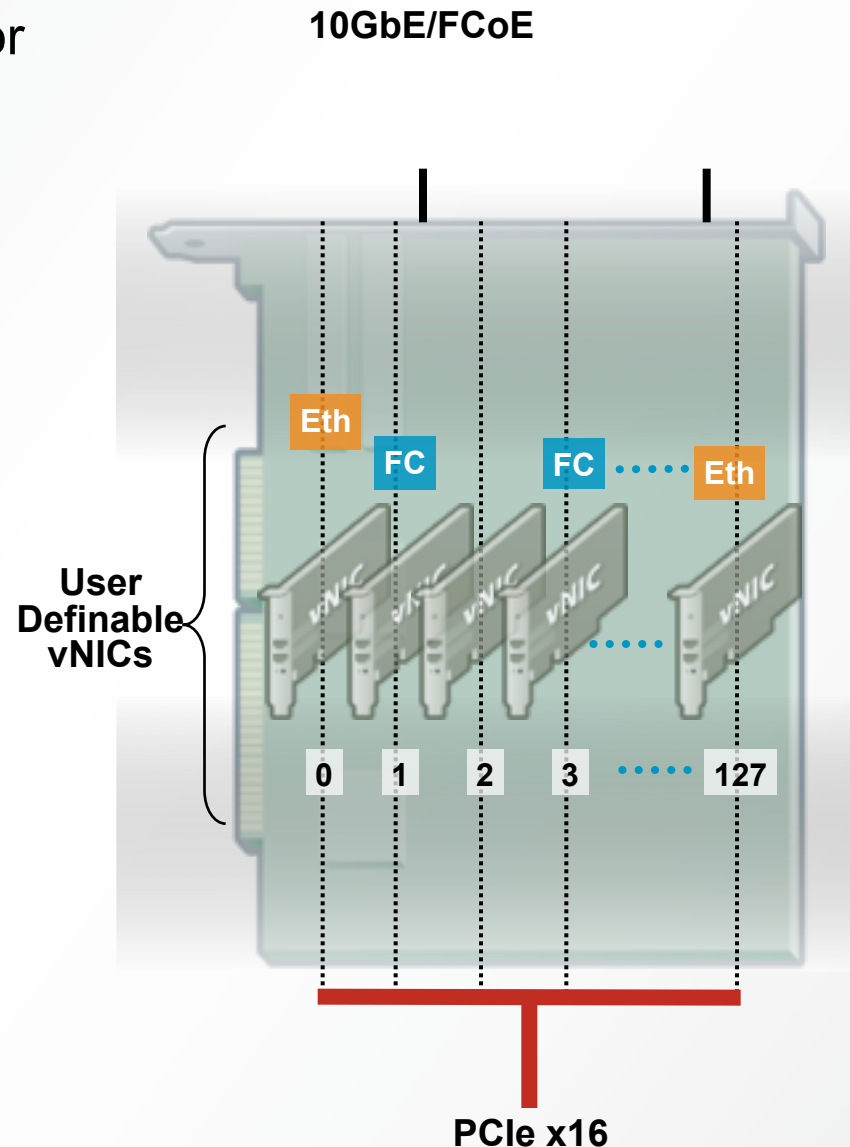
- 2x 10Gb
- >500K IOPS

The OS/Hypervisor sees up to ~128 distinct PCIe devices

- Ethernet vNIC and FC vHBA
- Management from the network

VN-Link in Hardware – Ideal for Virtualization Environments

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



# UCS Key Value Propositions:

## *Drivers for use cases*

Hardware State Abstraction – Service Profiles

Unified Fabric - FCOE

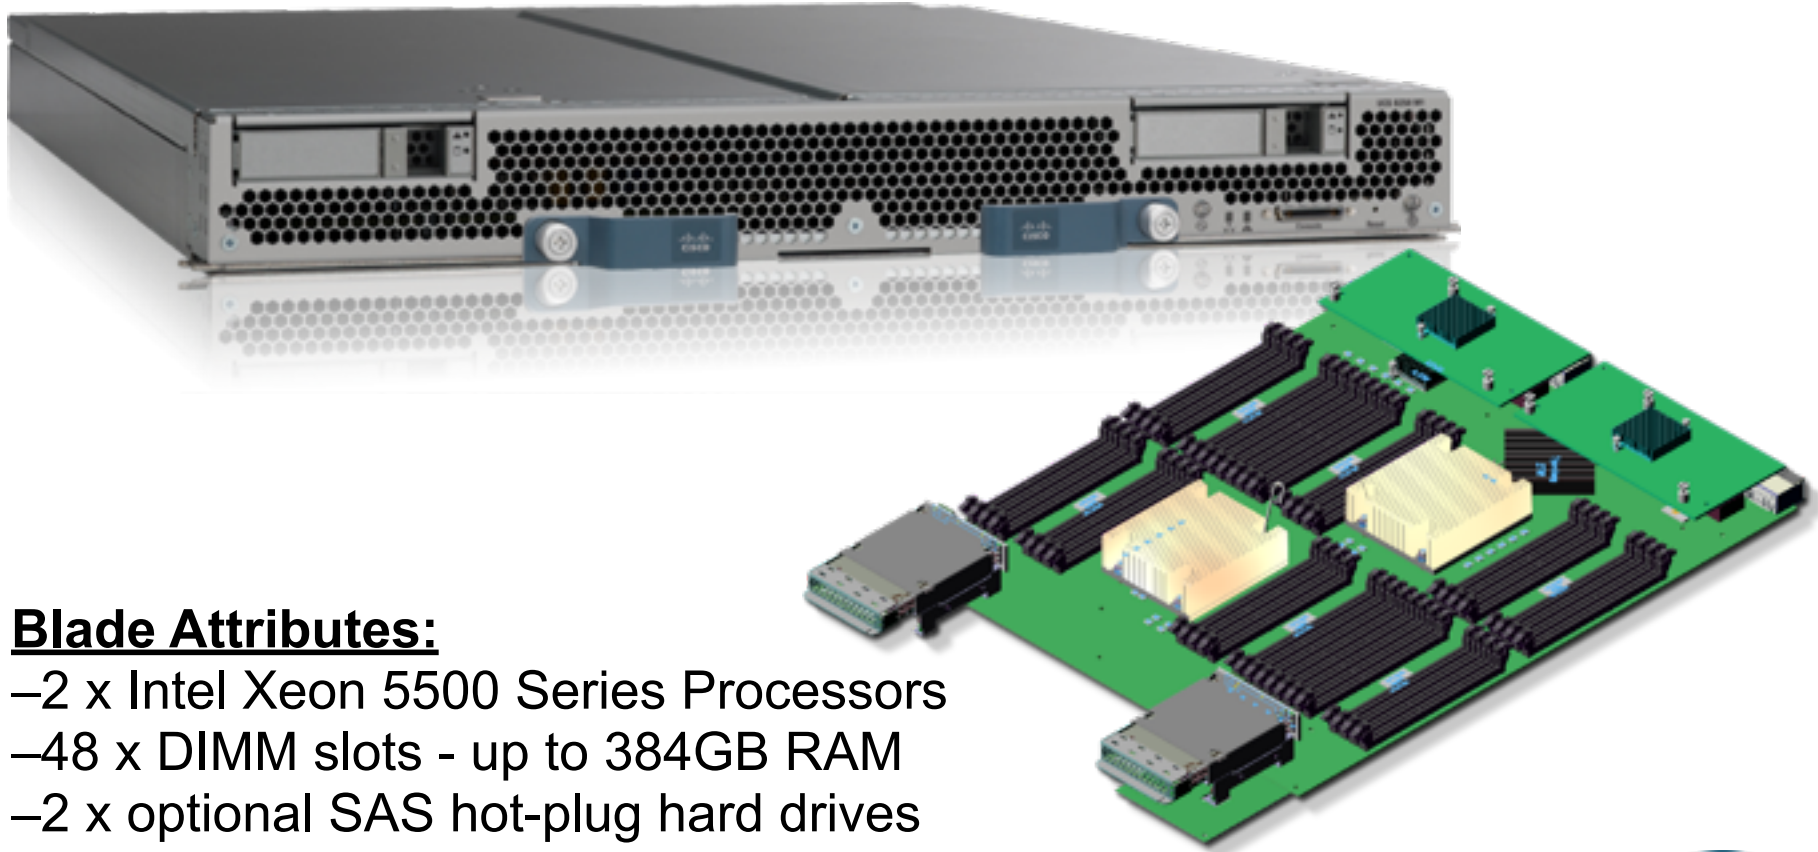
Virtualized Adapter - Palo



Expanded Memory Server – Full Height Blade

Unified Management - CAM

# UCS B250 M1 Blade



## **Blade Attributes:**

- 2 x Intel Xeon 5500 Series Processors
- 48 x DIMM slots - up to 384GB RAM
- 2 x optional SAS hot-plug hard drives
- RAID 0, 1
- 2 x 10Gb dual port mezzanine adapter
- Remote and local access to keyboard, video, mouse, serial
- Integrated with UCS Manager



# California Key Value Propositions:

## *Drivers for use cases*

Hardware State Abstraction – Service Profiles

Unified Fabric - FCOE

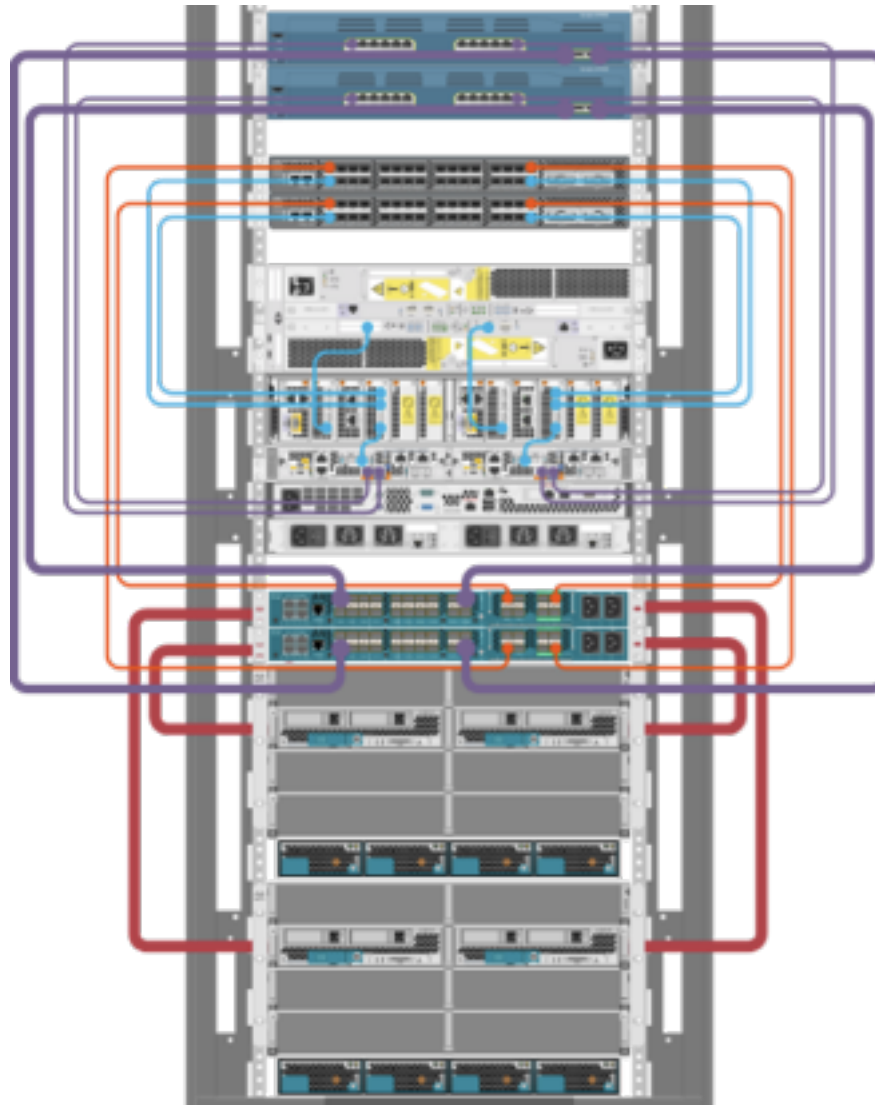
Virtualized Adapter - Palo

Expanded Memory Server – Full Height Blade



Unified Management - CAM

# System overview



# EMC features for VCE

- Load balancing for SAN in VMware env.
- Datastore provisioning / Thin provisioning
- VMware – EMC storage integration
- NAS storage provisioning
- Rapid deployment
- Backup in virtual environments

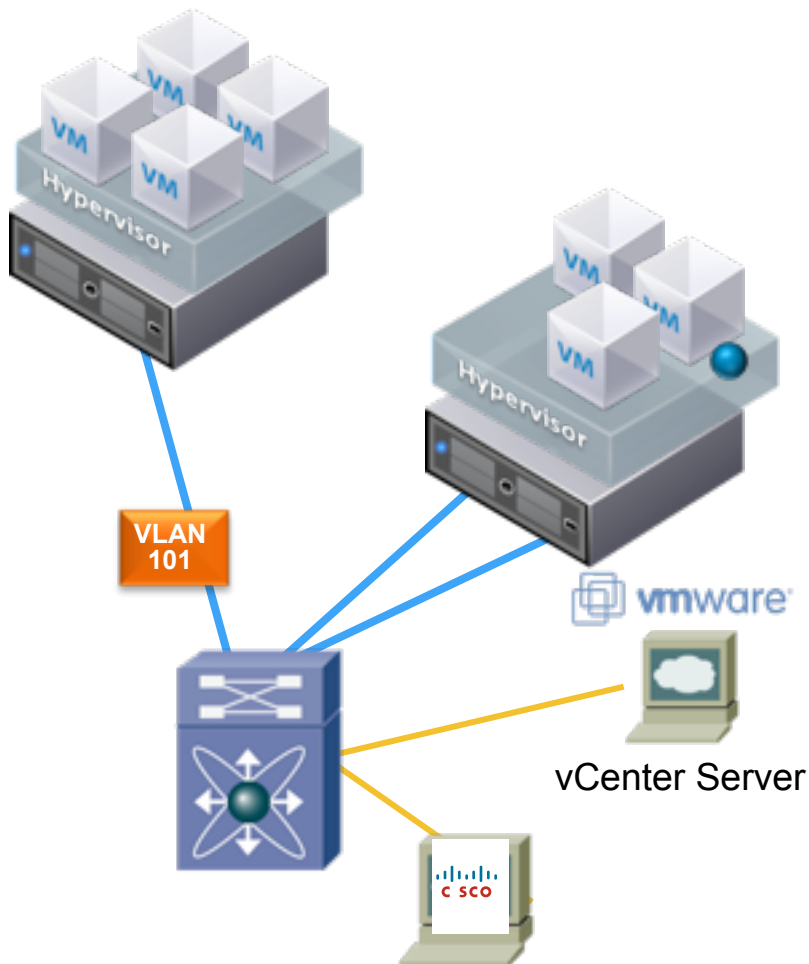
# Nexus 1000 V



# Networking for Server Virtualization

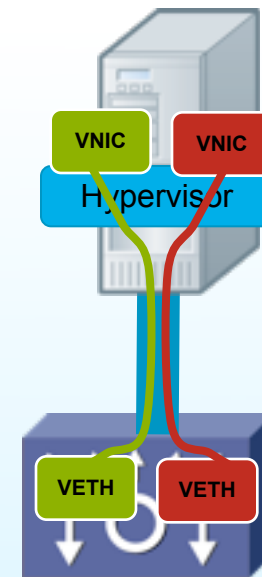
## Problems:

- VMotion may move VMs across physical ports—policy must follow
- Impossible to view or apply policy to locally switched traffic
- Need collaboration between network and virtualization admin



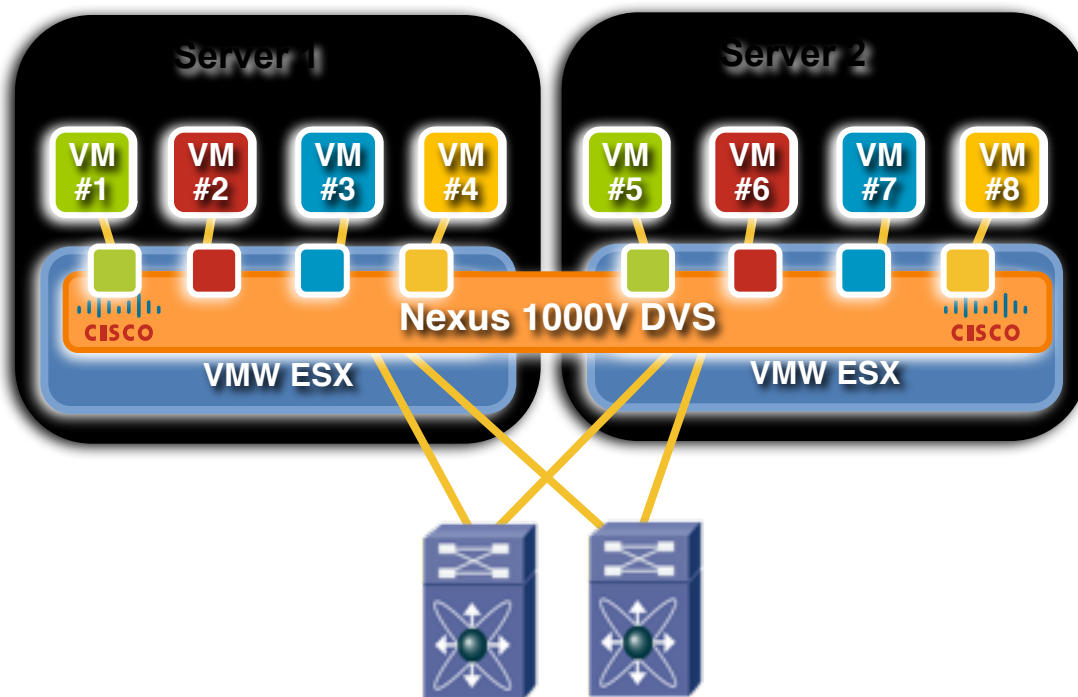
# Cisco Virtual Network Link – VN-Link

- Virtual Network Link (VN-Link) is about:
  - VM-level network granularity
  - Mobility of network and security properties (follow the VM)
  - Policy-based configuration of VM interfaces (Port Profiles)
  - Non-disruptive operational model
- VN-Link refers to a literal link between a VM VNIC & a Cisco VN-Link Switch
- VN-Link with Nexus 1000V
  - Replaces Hypervisor switch with Cisco modular switch (software)



# Cisco Nexus 1000V

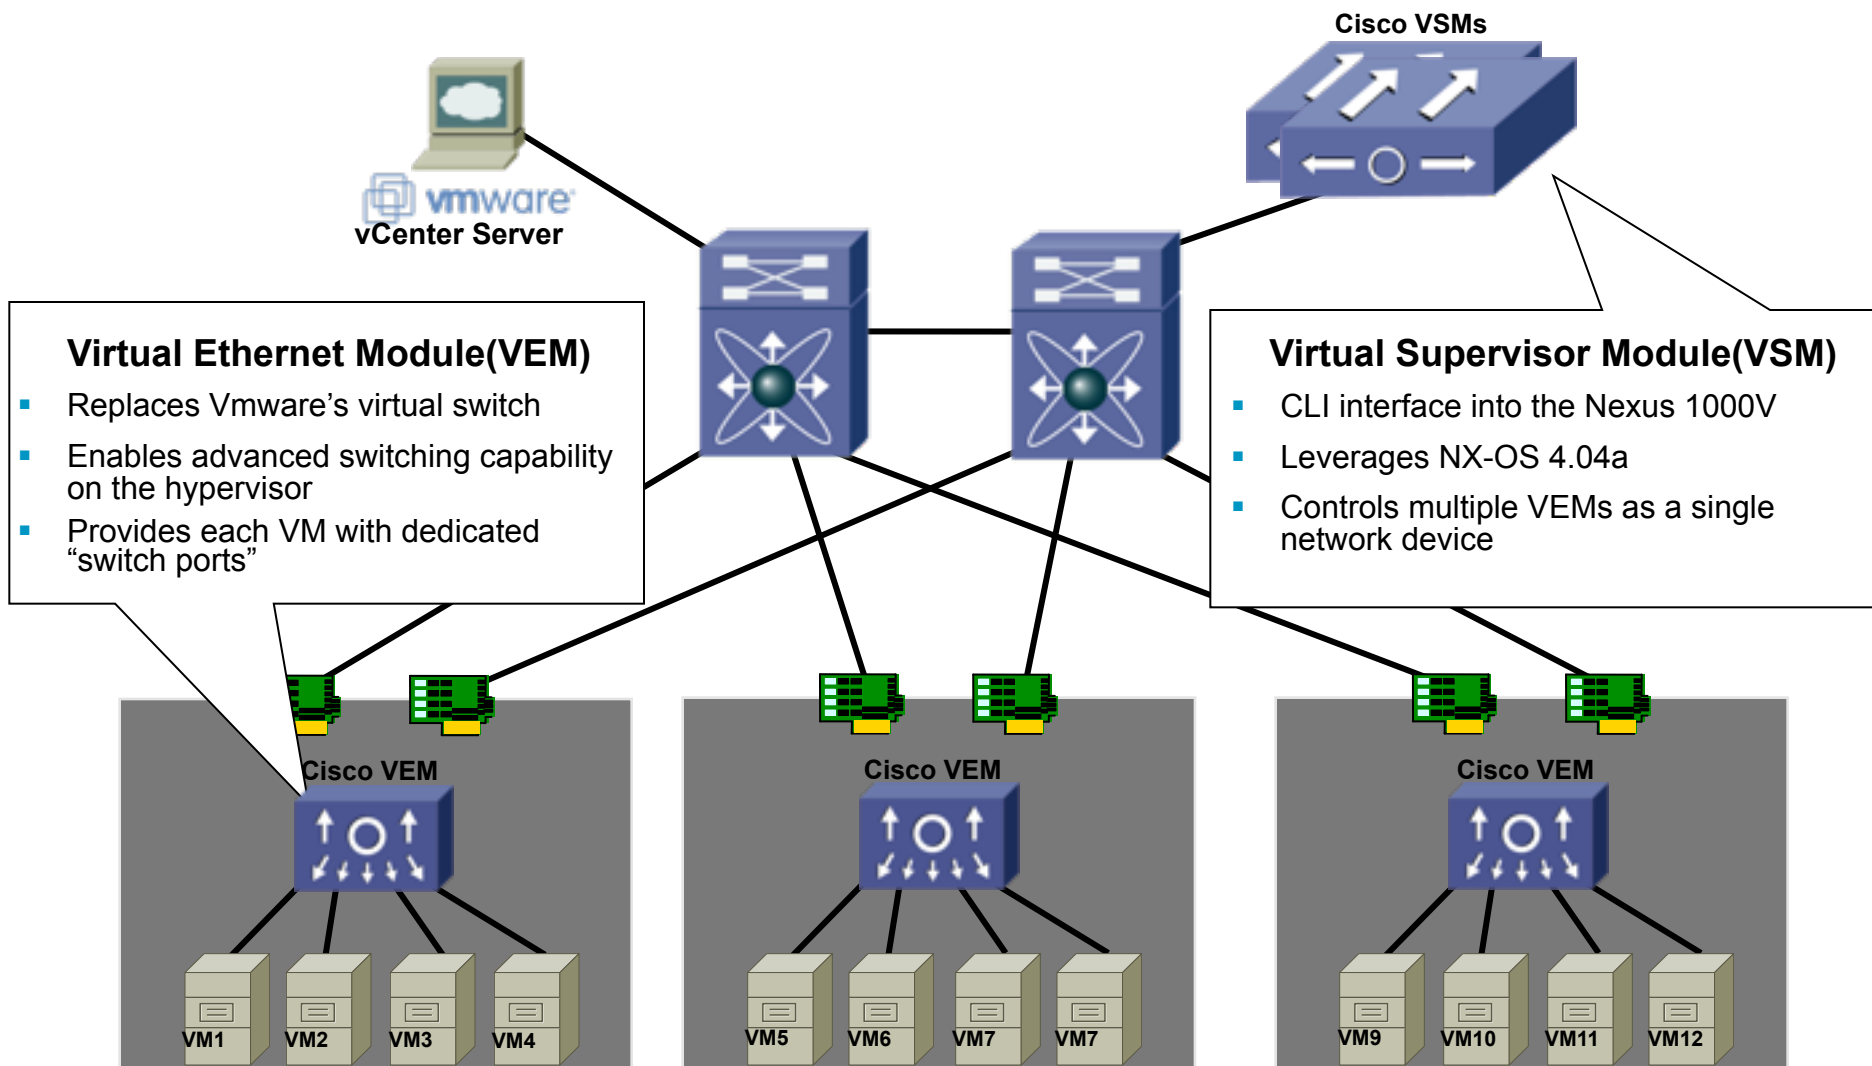
*Industry First 3<sup>rd</sup> Party Distributed Virtual Switch*



- 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
- Built on NX-OS

*Enabling Acceleration of Server Virtualization Benefits*

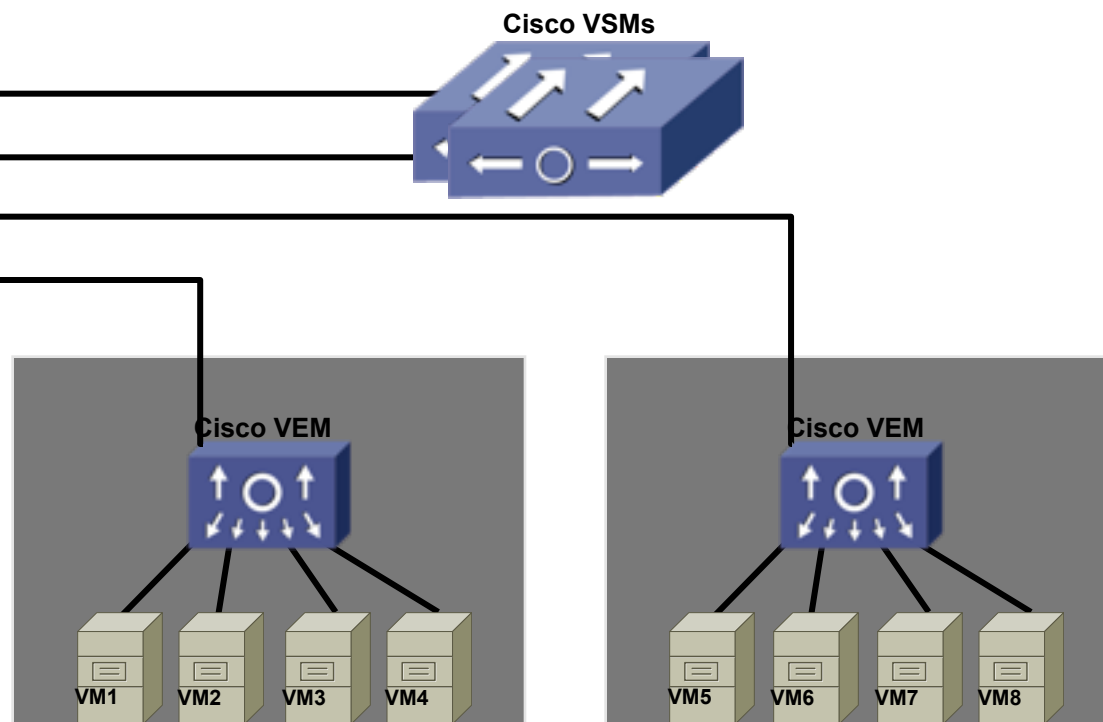
# Cisco Nexus 1000V Components



# Cisco Nexus 1000V 'Virtual Chassis'



pod5-vsm# show module				
Mod	Ports	Module-Type	Model	Status
1	0	Virtual Supervisor Module	Nexus1000V	active *
2	0	Virtual Supervisor Module	Nexus1000V	ha-standby
3	248	Virtual Ethernet Module	NA	ok



# Distributed Data Plane

- Each Virtual Ethernet Module forwards packets independent of each other.

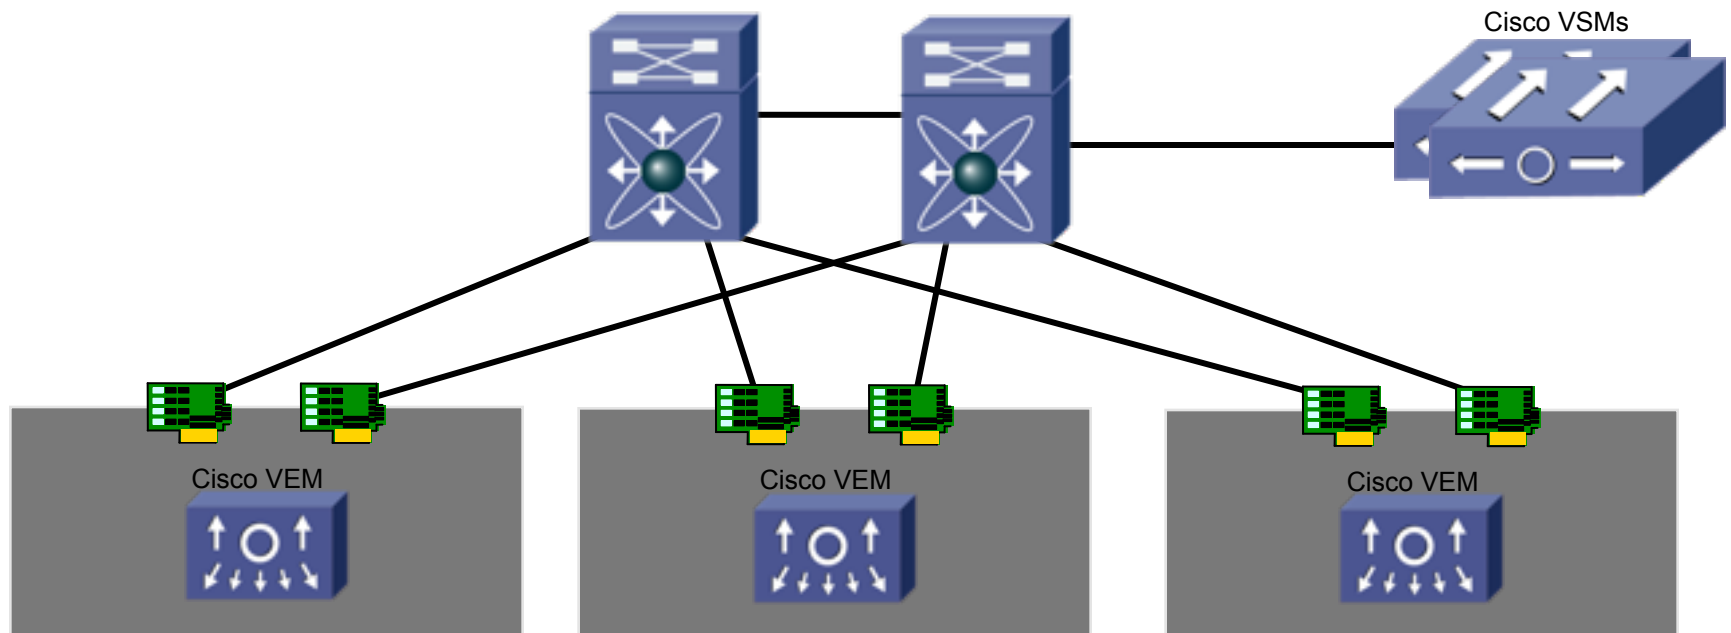
*No address learning/synchronization across VEMs*

*No concept of Crossbar/Fabric between the VEMs*

*Virtual Supervisor Module is NOT in the data path*

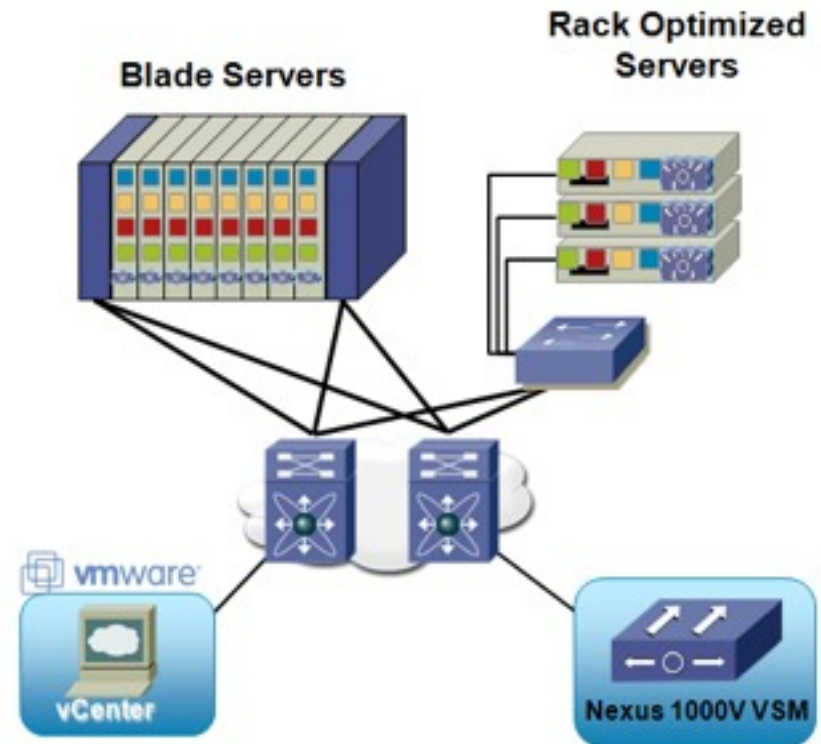
*No concept of forwarding from an ingress linecard to an egress linecard (another server)*

*No Etherchannel across VEMs*



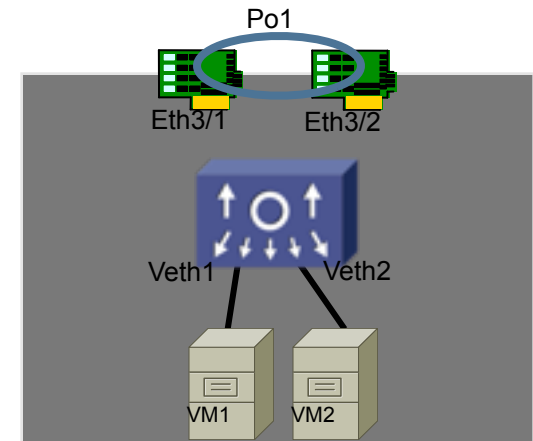
# Flexible Deployment Options

- Any type of physical switch (Cisco & other vendors)
- 1G & 10G NICs
- All types of servers



# Cisco Nexus 1000V Switch Interfaces

- Ethernet Port (eth)
  - 1 per physical NIC interface
  - Specific to each module
  - vmnic0 = ethx/1
  - Up to 32 per host
- Port Channel (po)
  - Aggregation of Eth ports
  - Up to 8 Port Channels per host
- Virtual Ethernet Port (veth)
  - 1 per VNIC (including SC and VMK)
  - Notation is Veth(port number).
  - No module number is assigned to enable consistent naming when moved
  - 216 per host



# What is a Port-Profile?

- A port-profile is a container used to define a common set of configuration commands for multiple interfaces
- Define once and apply many times
- Simplifies management by storing interface configuration
- Key to collaborative management of virtual networking resources
- Why is it not like a template or SmartPort macro?

Port-profiles are 'live' policies

Editing an enabled profile will cause config changes to propagate to all interfaces using that profile (unlike a static one-time macro)

# Port Profile Configuration

```
n1000v# show port-profile name WebProfile
  port-profile WebProfile
    description:
    status: enabled
    capability uplink: no
    system vlans:
    port-group: WebProfile
    config attributes:
      switchport mode access
      switchport access vlan 110
      no shutdown
    evaluated config attributes:
      switchport mode access
      switchport access vlan 110
      no shutdown
    assigned interfaces:
      Veth10
```

## Support Commands Include:

- ✓ Port management
- ✓ VLAN
- ✓ PVLAN
- ✓ Port-channel
- ✓ ACL
- ✓ Netflow
- ✓ Port Security
- ✓ QoS

WebServer-1 - Virtual Machine Properties

Hardware | Options | Resources

Virtual Machine Version: 7

☐ Show All Devices

Add... Remove

Hardware	Summary
Memory	1024 MB
CPUs	1
Video card	Video card
VMCI device	Restricted
Floppy drive 1	Client Device
CD/DVD Drive 1	Client Device
<b>Network adapter 1 (edit...)</b>	<b>WebServers (pod5-vsm)</b>
SCSI controller 0	LSI Logic Parallel
Hard disk 1	Virtual Disk

Device Status

☒ Connected

☒ Connect at power on

Adapter Type

Current adapter: E1000

MAC Address

00:50:56:00:36:c2

☒ Automatic ☐ Manual

Network Connection

☒ Network Label

WebServers (pod5-vsm)

Control

Management

Packet

VM Network

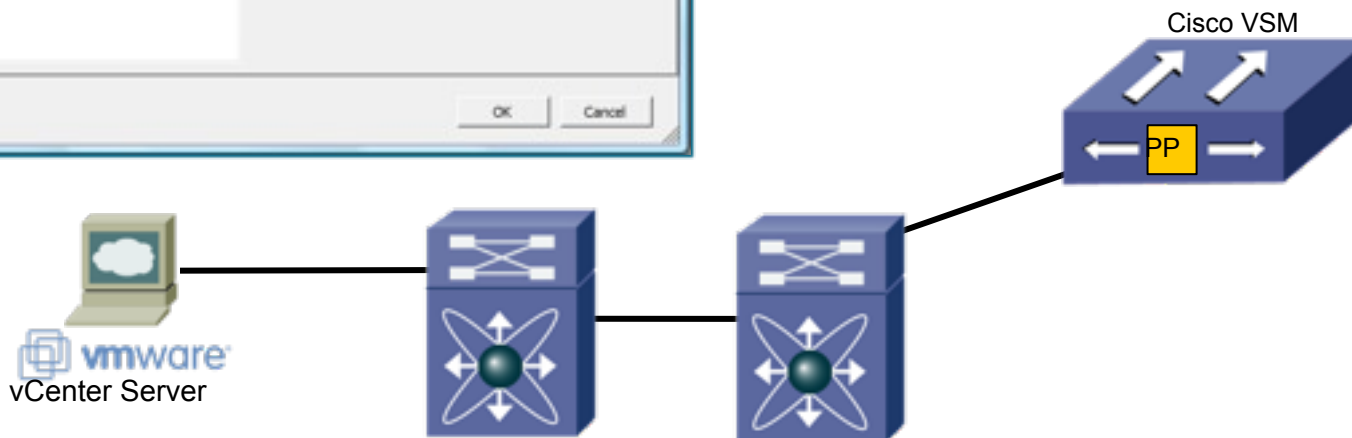
VMware-Web

Unused Or Quarantine Veth (pod5-vsm)

WebServers (pod5-vsm)

Help OK Cancel

```
n1000v(config)# port-profile WebServers
n1000v(config-port-prof)# switchport mode access
n1000v(config-port-prof)# switchport access vlan 100
n1000v(config-port-prof)# no shut
```



# Virtualization in Data Center



**Nexus 5000/ Nexus 2000 /FCoE -**

**Ciprian Tone**

**[ctone@cisco.com](mailto:ctone@cisco.com)**

# FCoE: Fibre Channel over Ethernet



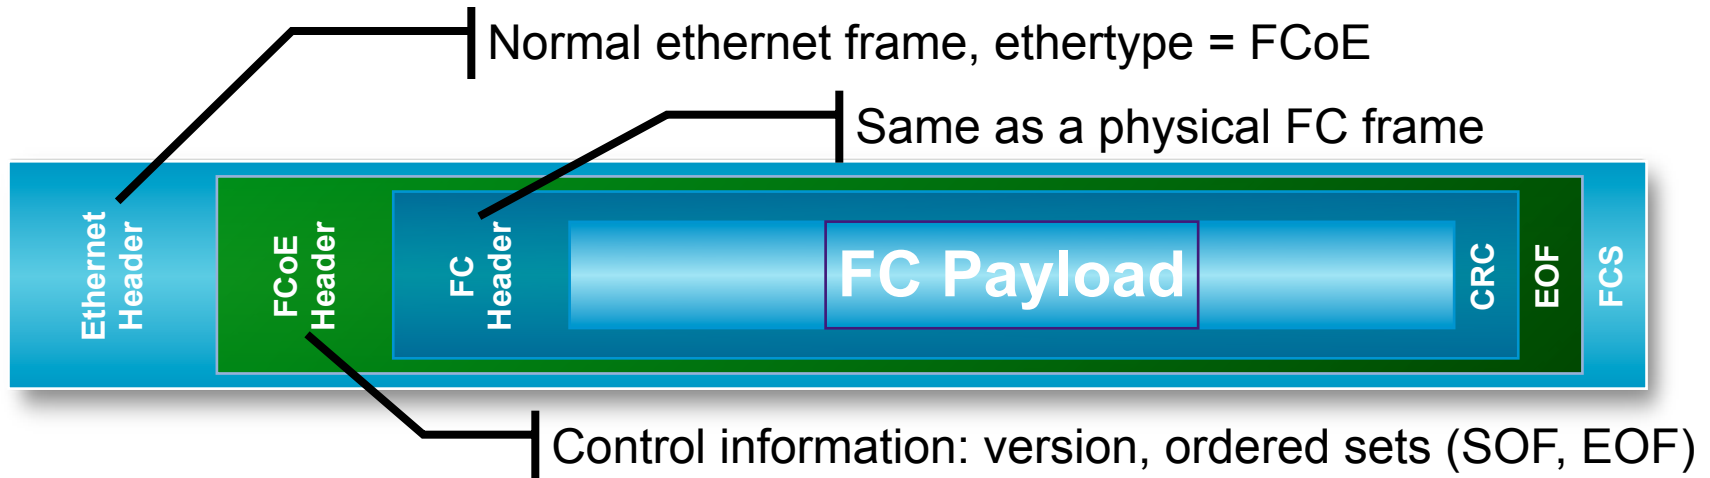
# What is Fibre Channel over Ethernet?

- From a Fibre Channel standpoint it's  
FC connectivity over a new type of cable called... an Ethernet cloud
- From an Ethernet standpoints it's  
Yet another ULP (Upper Layer Protocol) to be transported, but...  
a challenging one!
- And technically...

***FCoE is an extension of Fibre Channel  
onto a Lossless Ethernet fabric***

# FCoE Enablers and Encapsulation

- 10Gbps Ethernet
- Lossless Ethernet
  - Matches the lossless behavior guaranteed in FC by B2B credits
- Ethernet jumbo frames
  - Max FC frame payload = 2112 bytes
  - Total max frame size = 2180 bytes

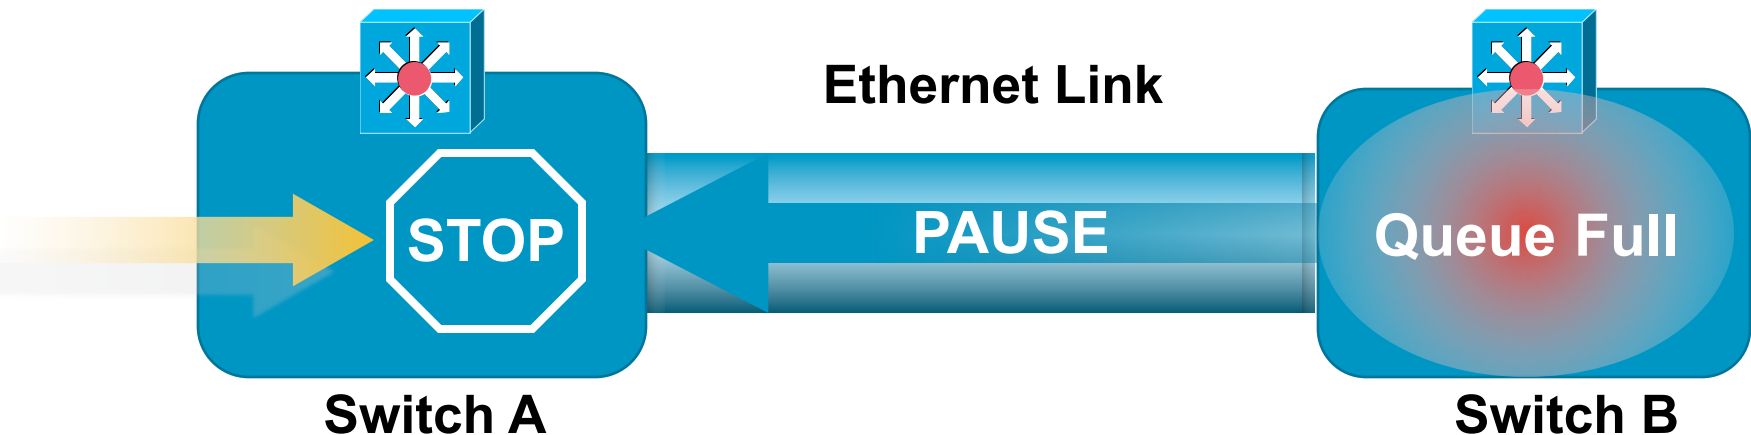


# Why Are Frames Lost?

Collision	Transmission Error	Congestion
<ul style="list-style-type: none"><li>▪ No longer present in full duplex Ethernet</li></ul>	<ul style="list-style-type: none"><li>▪ Very rare in the data center</li></ul>	<ul style="list-style-type: none"><li>▪ Most common cause</li><li>▪ Congestion is a switch issue, not a link issue<ul style="list-style-type: none"><li>A full duplex IEEE 802.3 link does not lose frames</li></ul></li><li>▪ It must be dealt with in the bridge/switch<ul style="list-style-type: none"><li>By IEEE 802.1</li></ul></li></ul>

# Can Ethernet Be Lossless?

- Yes, with Ethernet PAUSE Frame

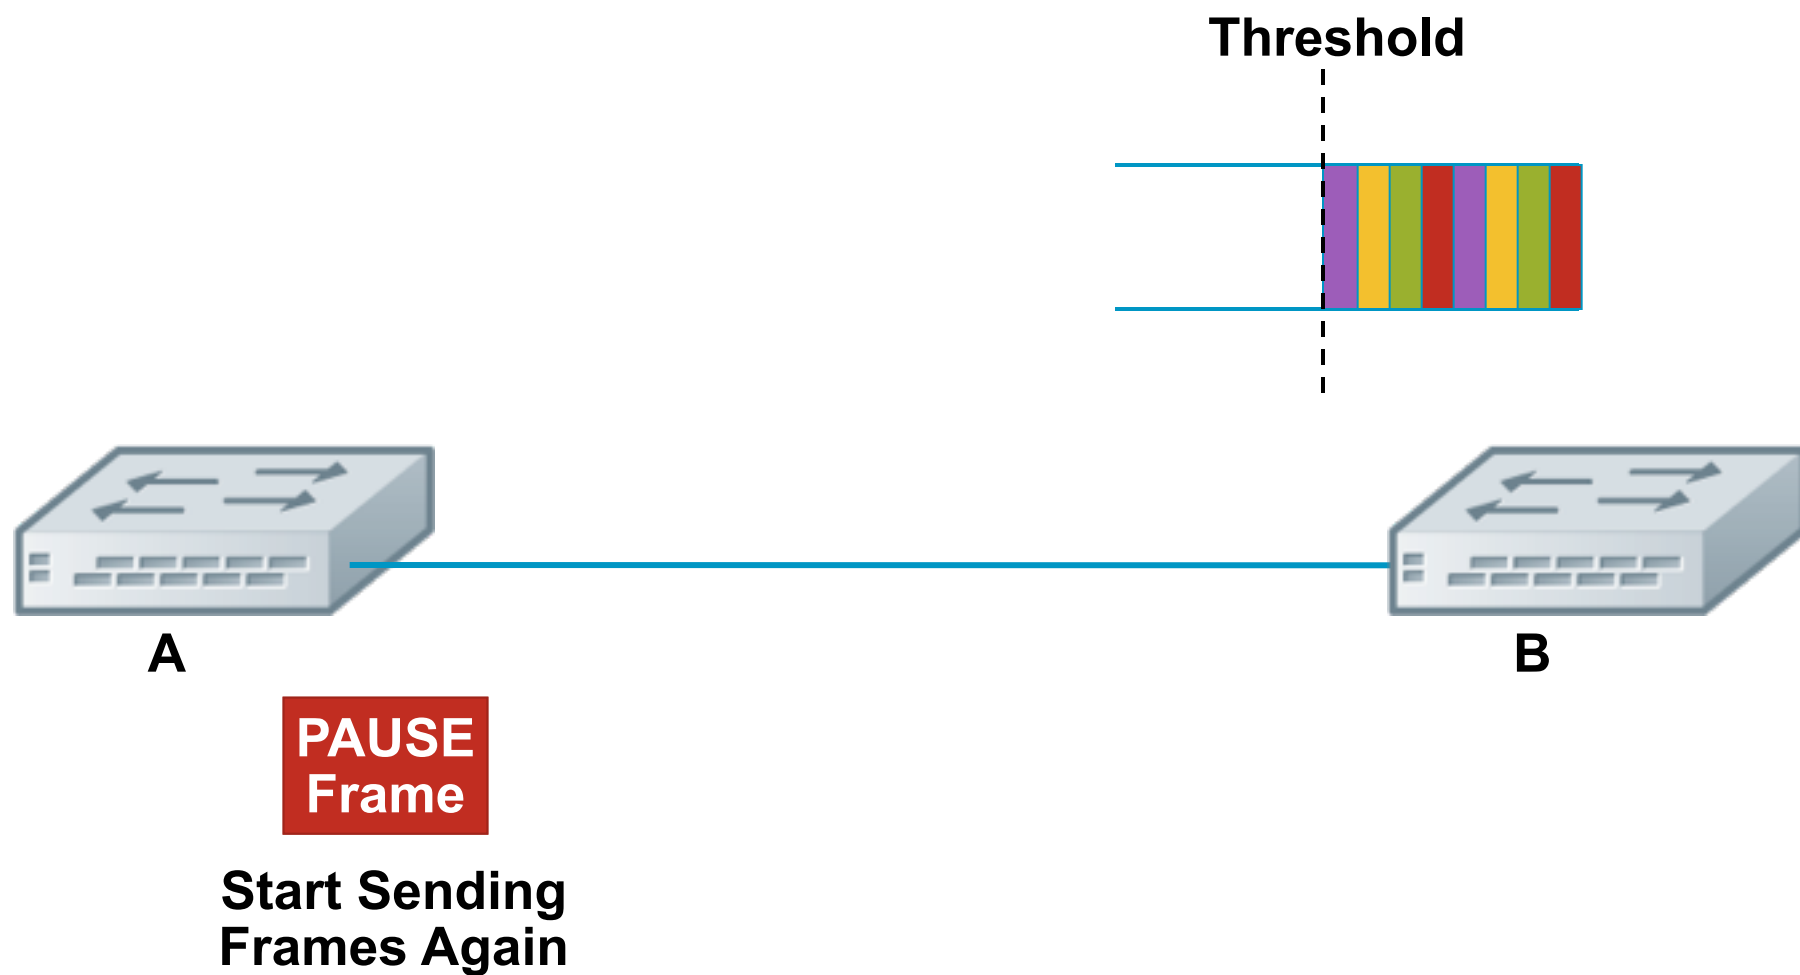


1. Defined in IEEE 802.3—Annex 31B

The PAUSE operation is used to inhibit transmission of data frames for a specified period of time

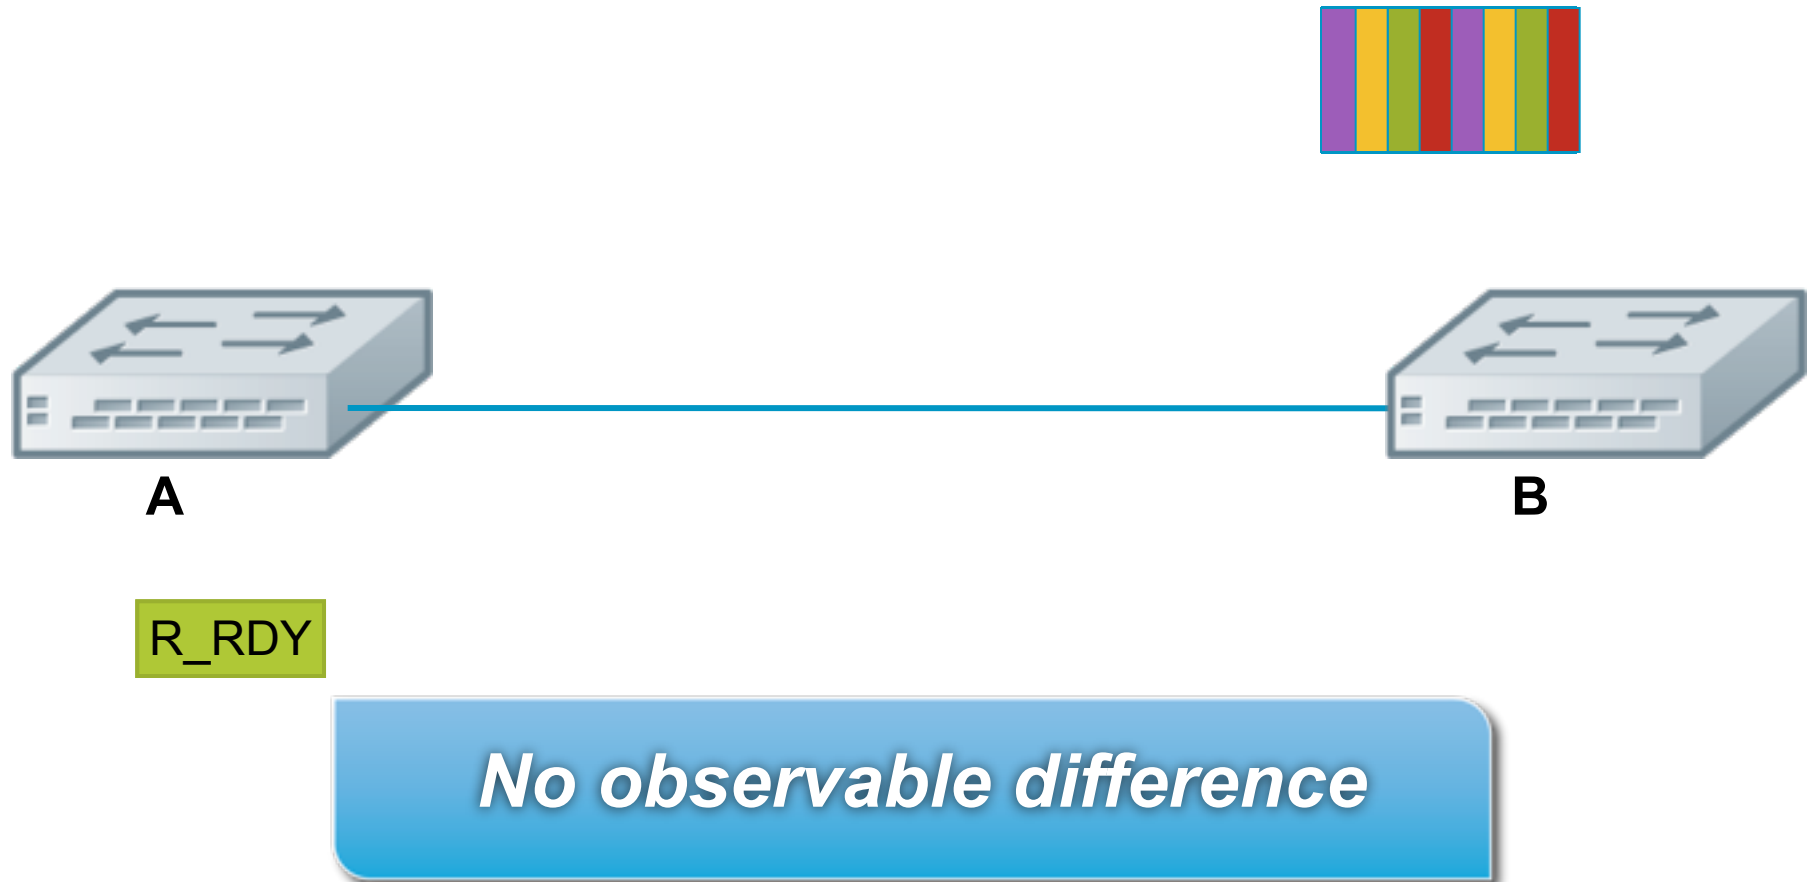
2. Ethernet PAUSE transforms Ethernet into a lossless fabric

# How PAUSE Works



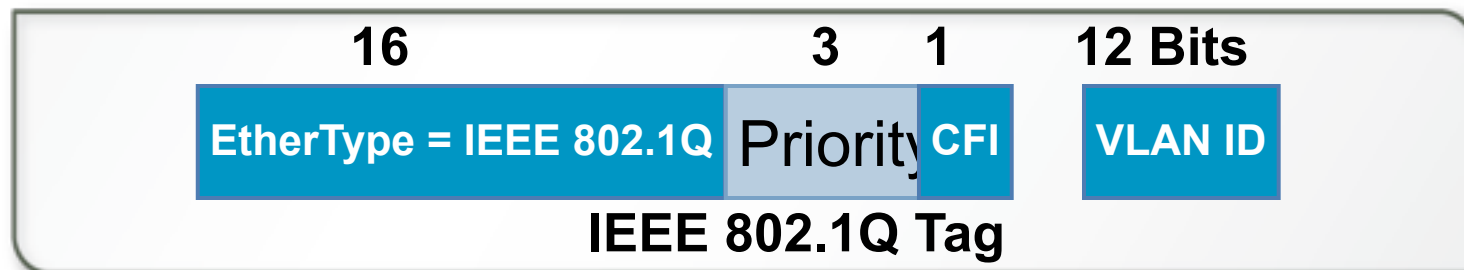
# Let's Compare PAUSE with FC Buffer to Buffer Credit

- Eight credits preagreed



# Priority Flow Control (PFC)

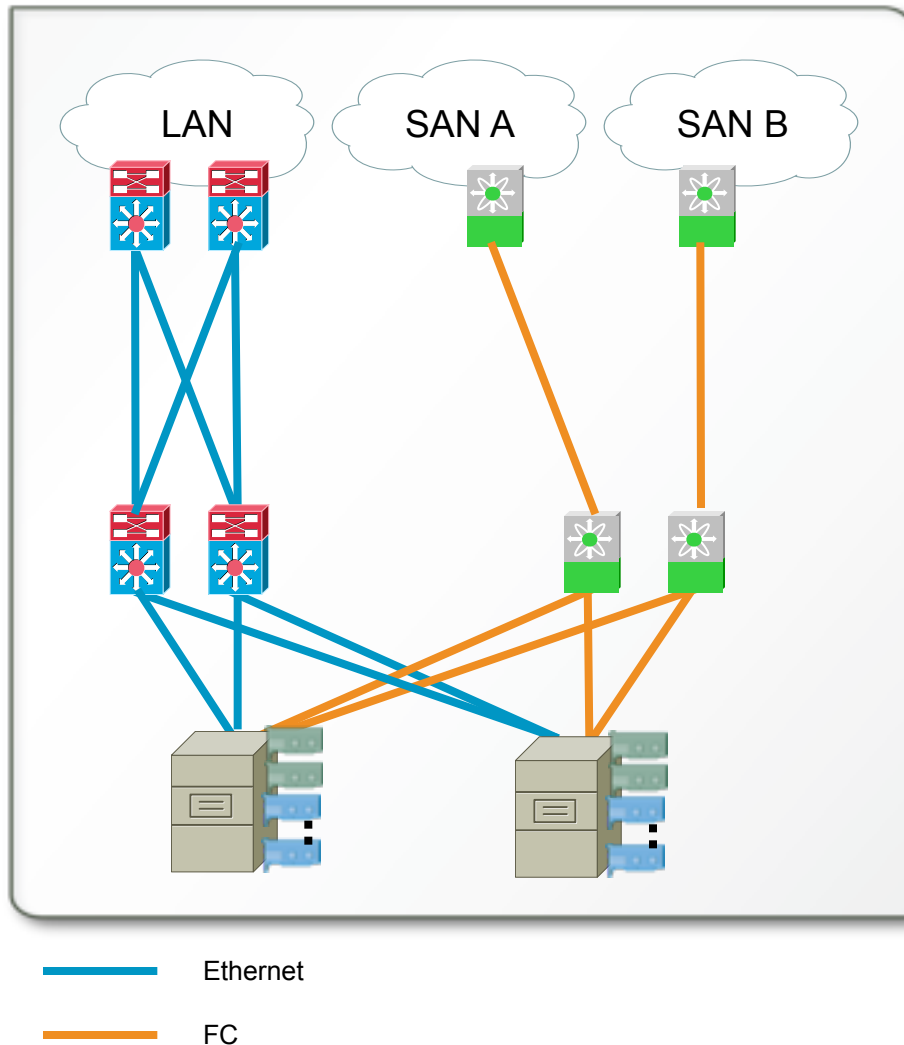
- a.k.a. PPP (Per Priority Pause)
- PFC enables PAUSE functionality per Ethernet priority
  - IEEE 802.1Q defines eight priorities
  - Traffic classes are mapped to different priorities:
    - No traffic interference
      - IP traffic may be paused while storage traffic is being forwarded
      - Or, vice versa
    - Requires independent resources per priority (buffers)
- High level of industry support
  - Cisco distributed proposal
  - Standard track in IEEE 802.1Qbb



# FCoE usage in Unified Fabric



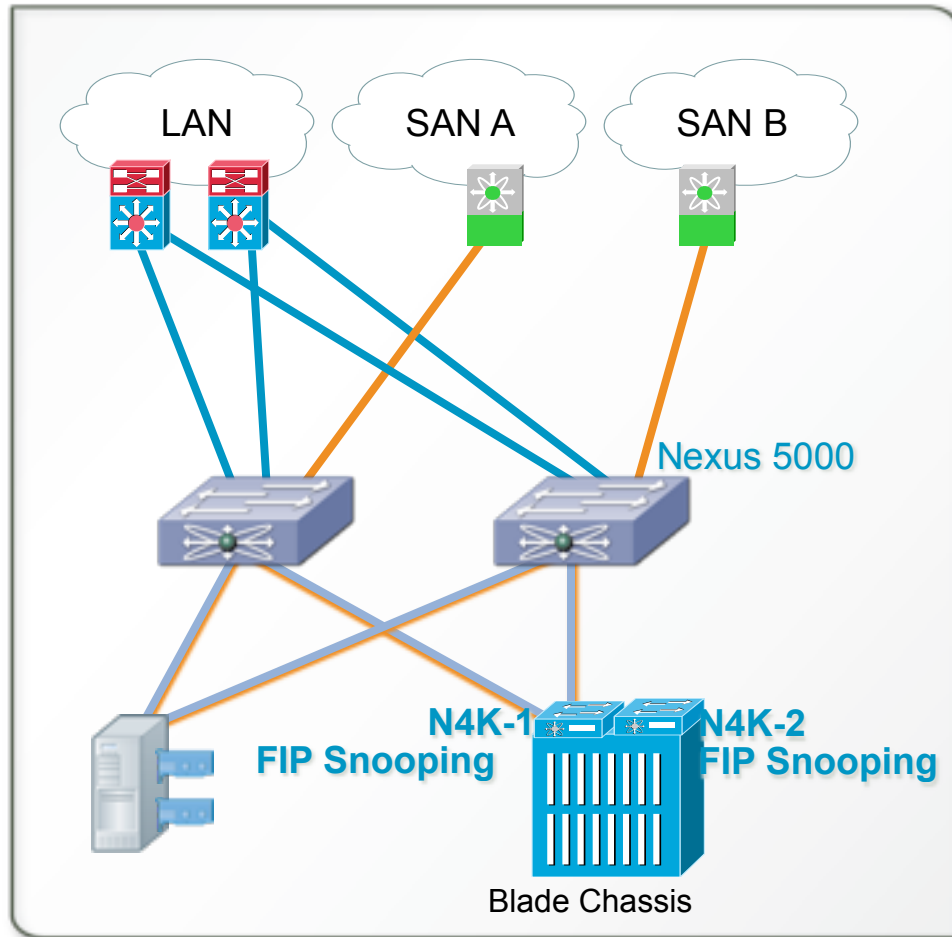
# Before I/O Consolidation



## Today

- Parallel LAN/SAN Infrastructure
- Inefficient use of Network Infrastructure
- 5+ connections per server—higher adapter and cabling costs
  - Adds downstream port costs; cap-ex and op-ex
  - Each connection adds additional points of failure in the fabric
- Longer lead time for server provisioning
- Multiple fault domains—complex diagnostics
- Management complexity

# I/O Consolidation



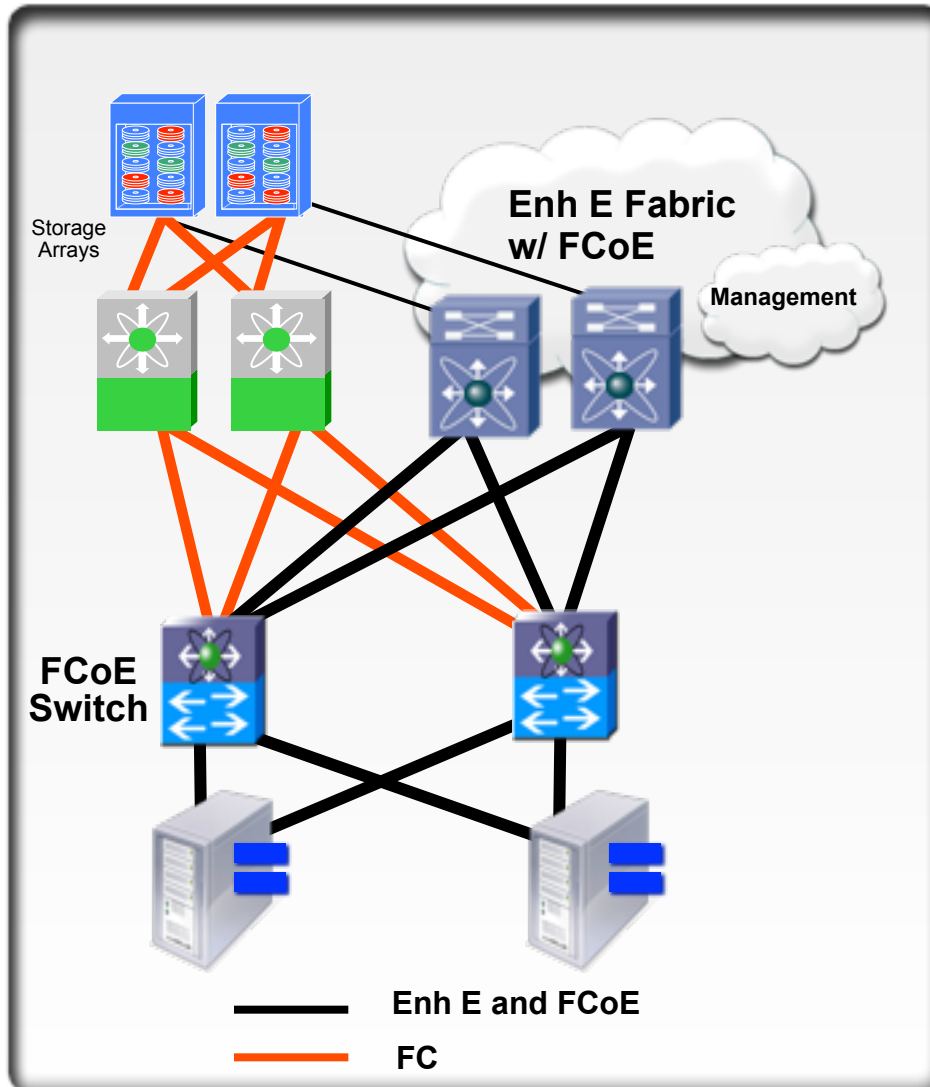
Enhanced Ethernet and FCoE

Ethernet FC

## Unified I/O Phase 1

- Reduction of server adapters
- Simplification of access layer and cabling
- Gateway free implementation—fits in installed base of existing LAN and SAN
- L2 Multipathing Access—Distribution
- Lower total cost of ownership
- Fewer cables
- Investment protection (LANs and SANs)
- Consistent operational model

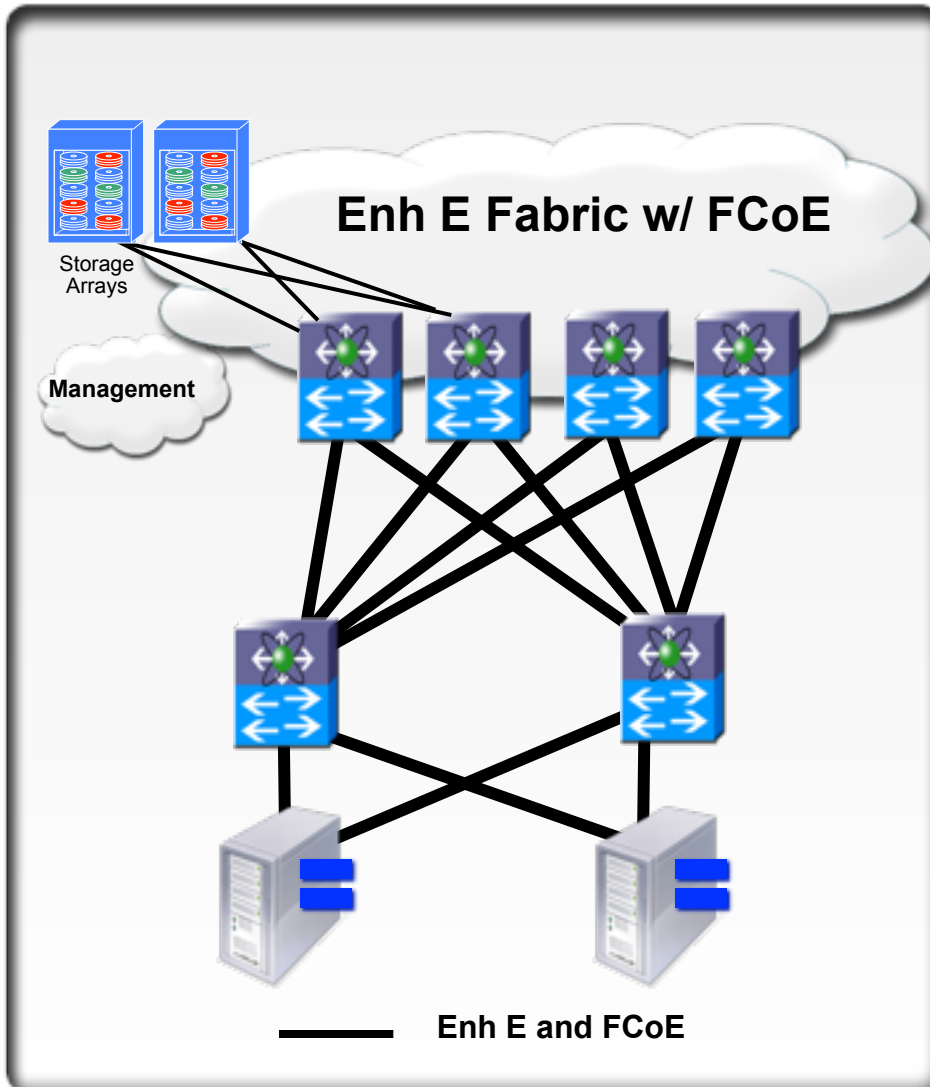
# I/O consolidation - next steps



## Unified I/O Phase 2

- Elimination of parallel network infrastructure
- L2/L3 Multipathing end to end
- Faster infrastructure provisioning
- Lower TCO
- Disk array access via Enhanced Ethernet or Native FC

# Unified I/O - next steps



## Unified I/O Phase 3

- Datacenter wide Unified Fabric for LAN and SAN
- L2/L3 Multipathing end to end
- Consistent network policies across datacenter
- Lower TCO

# Nexus 5000 family



# Nexus 5000 Systems Product Portfolio

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

### Nexus 5000 Switch Family



#### Nexus 5010 - 28-Port L2 Switch

- 20 Ports 10GE/FCoE/DCE, fixed
- 1 Expansion Module



#### Nexus 5020 - 56-Port L2 Switch

- 40 Ports 10GE/FCoE/DCE, fixed
- 2 Expansion Modules

### Expansion Modules



#### Fibre Channel

- 8 Ports 1/2/4G FC



#### Fibre Channel

- 6 Ports 2/4/8G FC



#### FC + Ethernet

- 4 Ports 10GbE/FCoE/DCE
- 4 Ports 1/2/4G FC



#### Ethernet

- 6 Ports 10GE/FCoE/DCE

### Partners



2x10GE/DCE/FCoE



SW FCoE/DCE + 2x10GE



2x10GE

### OS

Cisco NX-OS

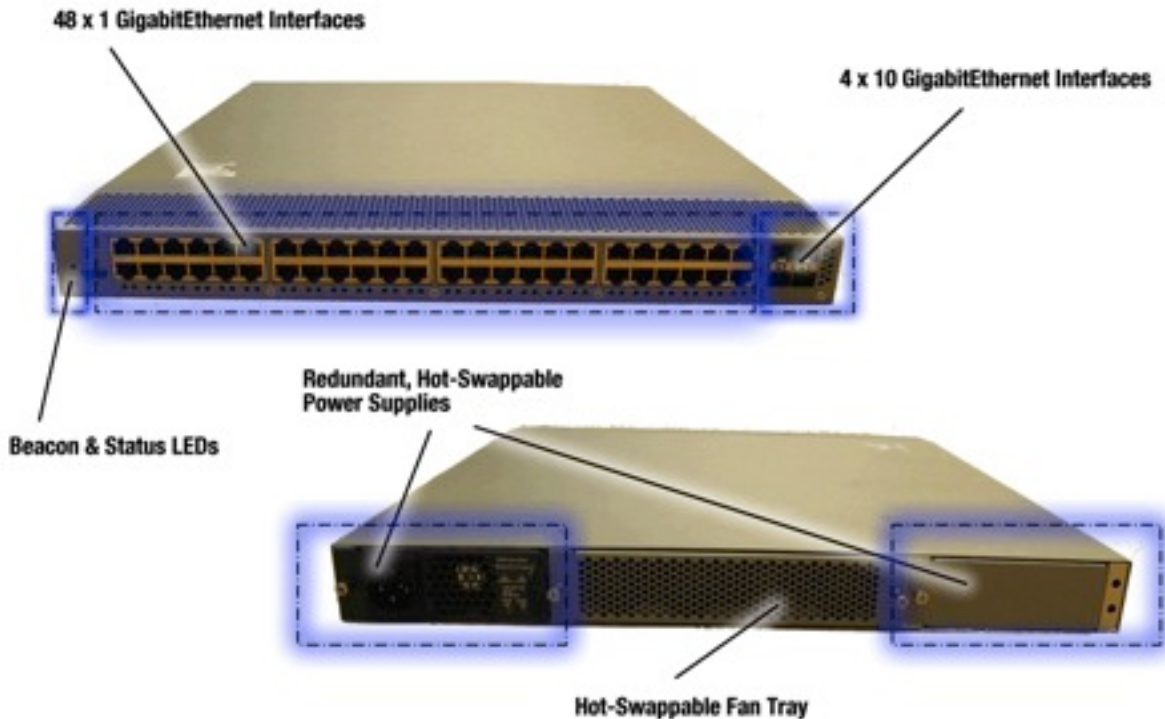
### Mgmt

Cisco Fabric Manager and Cisco Data Center Network Manager

# Nexus 2148T Fabric Extender

## 1GE Connectivity for servers

- "Remote linecard" for Nexus 5000
- Nexus 5000 will provide "supervisor module" functionality
- No local switching
- Not manageable entity
- Planned support for other platforms



# Demo