Cisco and VMware: Virtualizing the Data Center

Virtualization Server Bootcamp

Europe 2011

Sponsors of the vBootcamp

Agenda morning

Introduction / Agenda	09:00
UCS Introduction & Highlights and best practices	09:10
VMware – What's New	10:00
Break	10:45
Nexus 1000v Intro und Setup	11:00
Intel Server Technology Update	11:30
Summary UCS	12:00
Lunch	12:15

Agenda Afternoon

Intro to the Labs 13:30

Lab 1:
Creation of Service Profile
Connect to KVM
Explore UCS-M

Lab 2:
Install ESXi
Connect to vCenter
Configure Host-Profiles and vMotion

Lab 3:
Install Nexus 1000v
Configure Nexus 1000v

Bootcamp Event stop 5pm

House-Keeping

No Smoking in whole building

Toilettes

Break

Lunch

Please turn off mobiles

Please reduce in-out during sessions

Presentation download follows Survey....

The UCS Platform – A Technical Overview

Agenda

Introduction to UCS

UCS System Components

- UCS Manager
- Interconnect
- Fabric Extender
- Enclosure
- Compute Node
- IO Adapters

UCS Differentiators

Server Deployment: Rackmounts

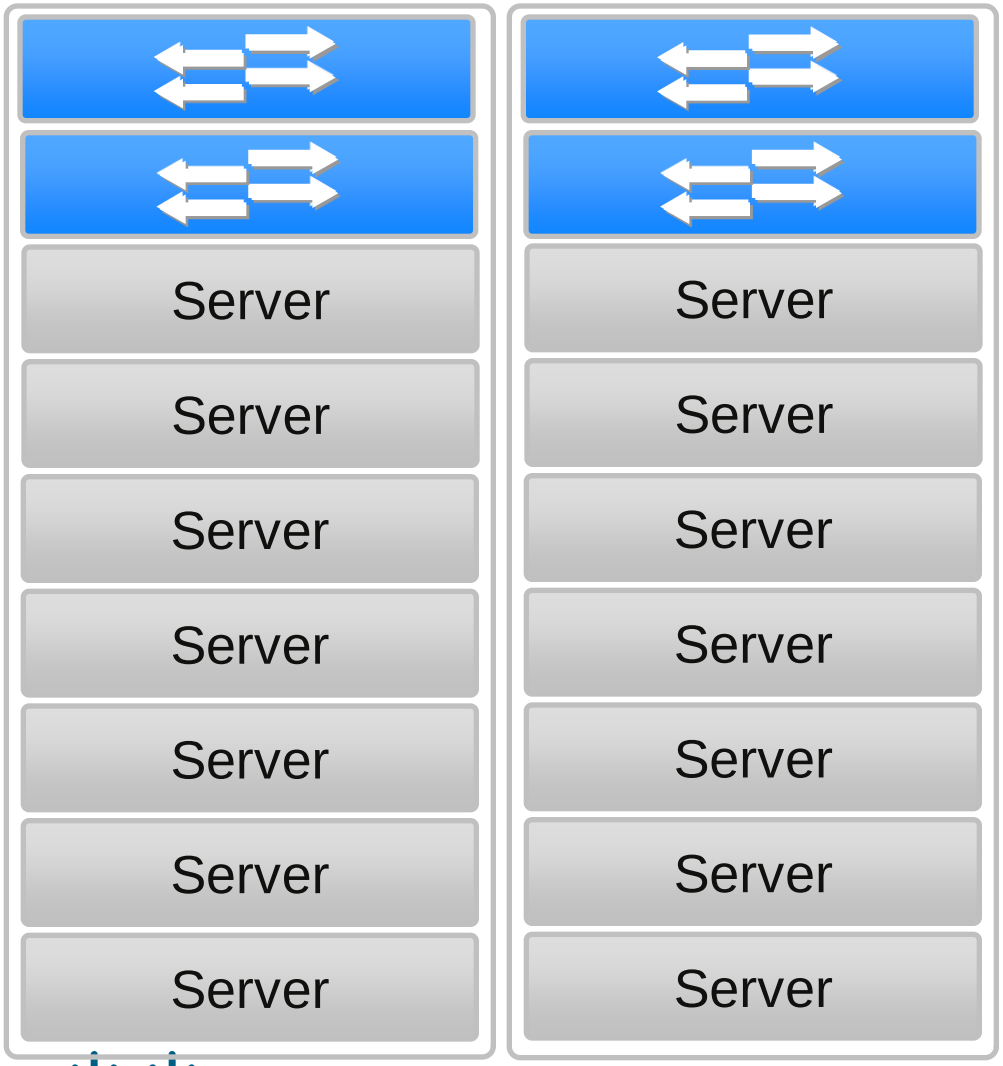

First generation

- Rack-optimized
- Top of Rack or End of Row switches
- Cables

Benefits

- Space utilization
- Highly flexible

Weakness

- Cabling
- Serviceability
- Power efficiency

Server Deployment: Blades

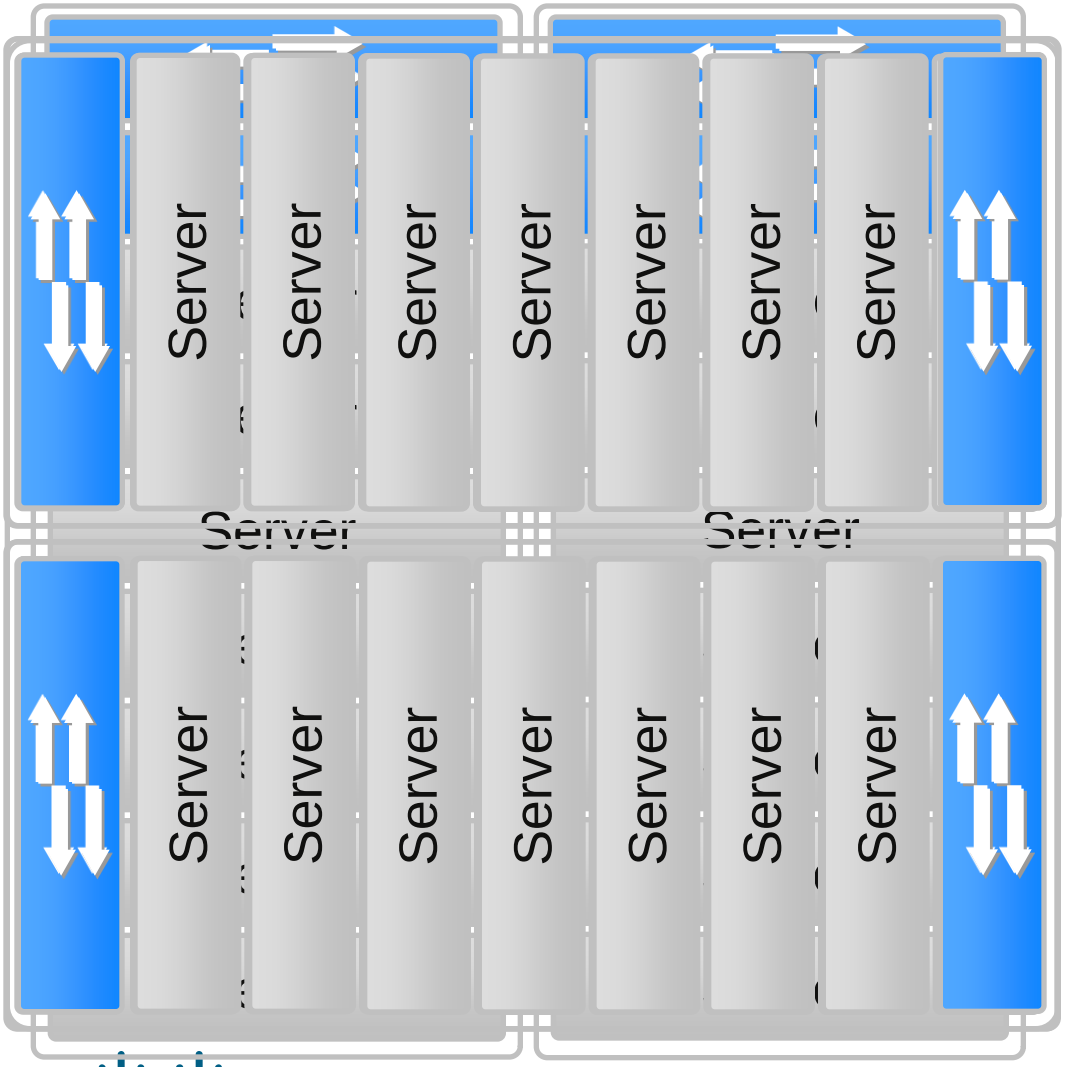

Server Deployment Management

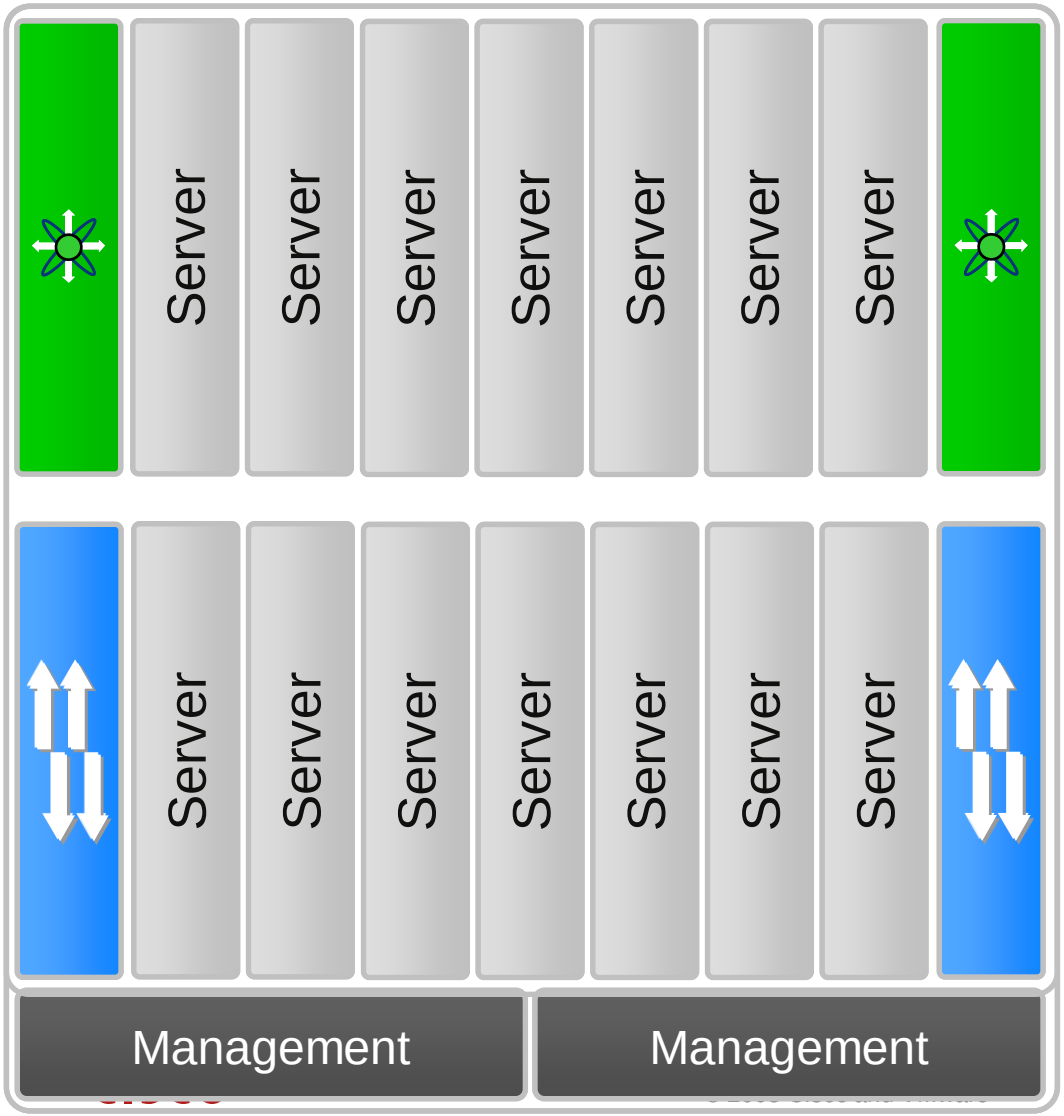

Chassis Management

- New management layer

Benefits

- Consistency in chassis
- Shared chassis infrastructure monitoring

Weakness

- Additional mgmt overhead
- Additional cost overhead
- Need chassis aggregation management
- Artificial aggregation point

Server Deployment: Applications

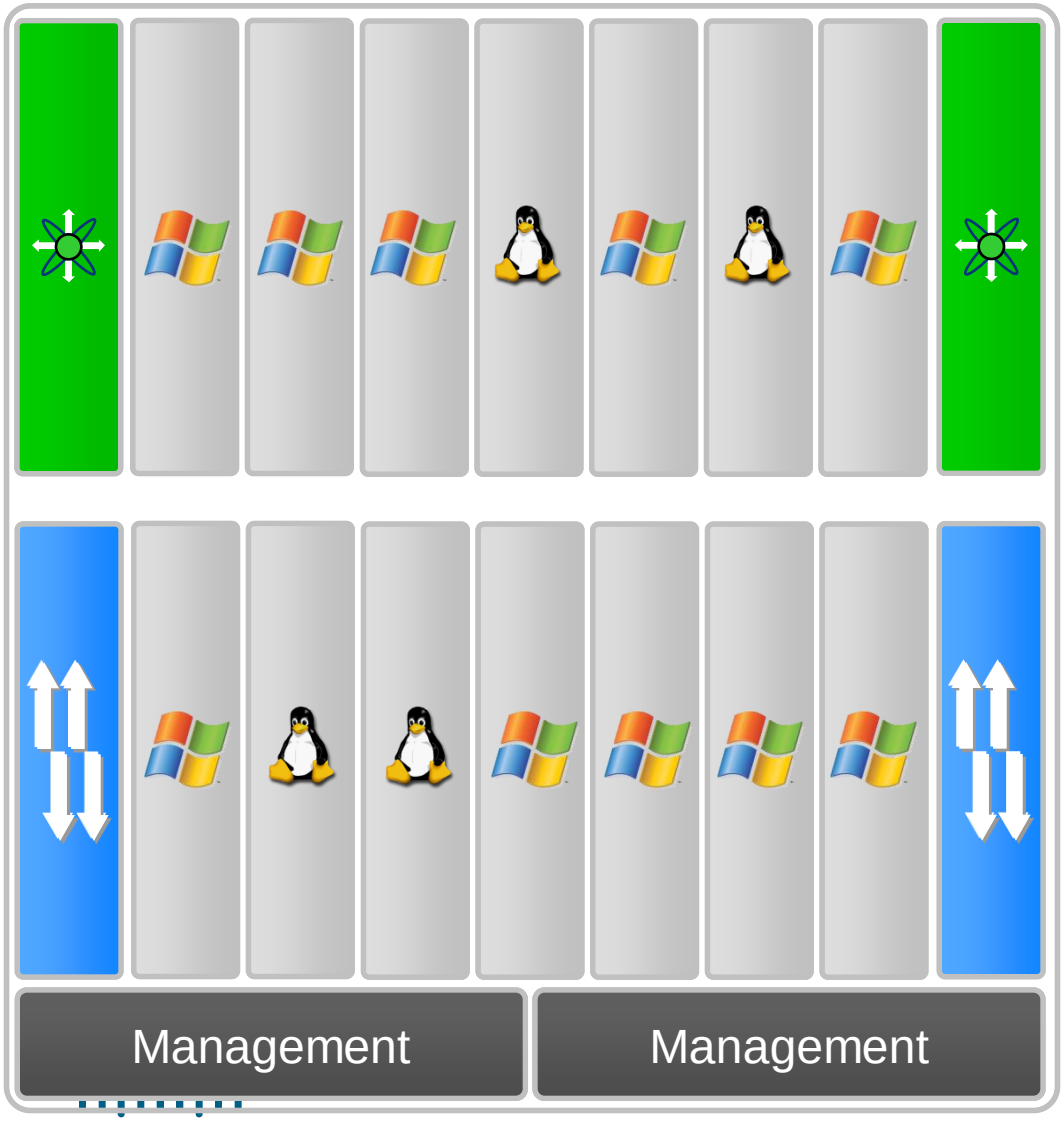

Application Deployment

- Single application per OS

Benefits

- Isolation
- Flexibility
- Simplicity

Weakness

- Low utilization
- Power & Cooling
- Server sprawl

Server Deployment

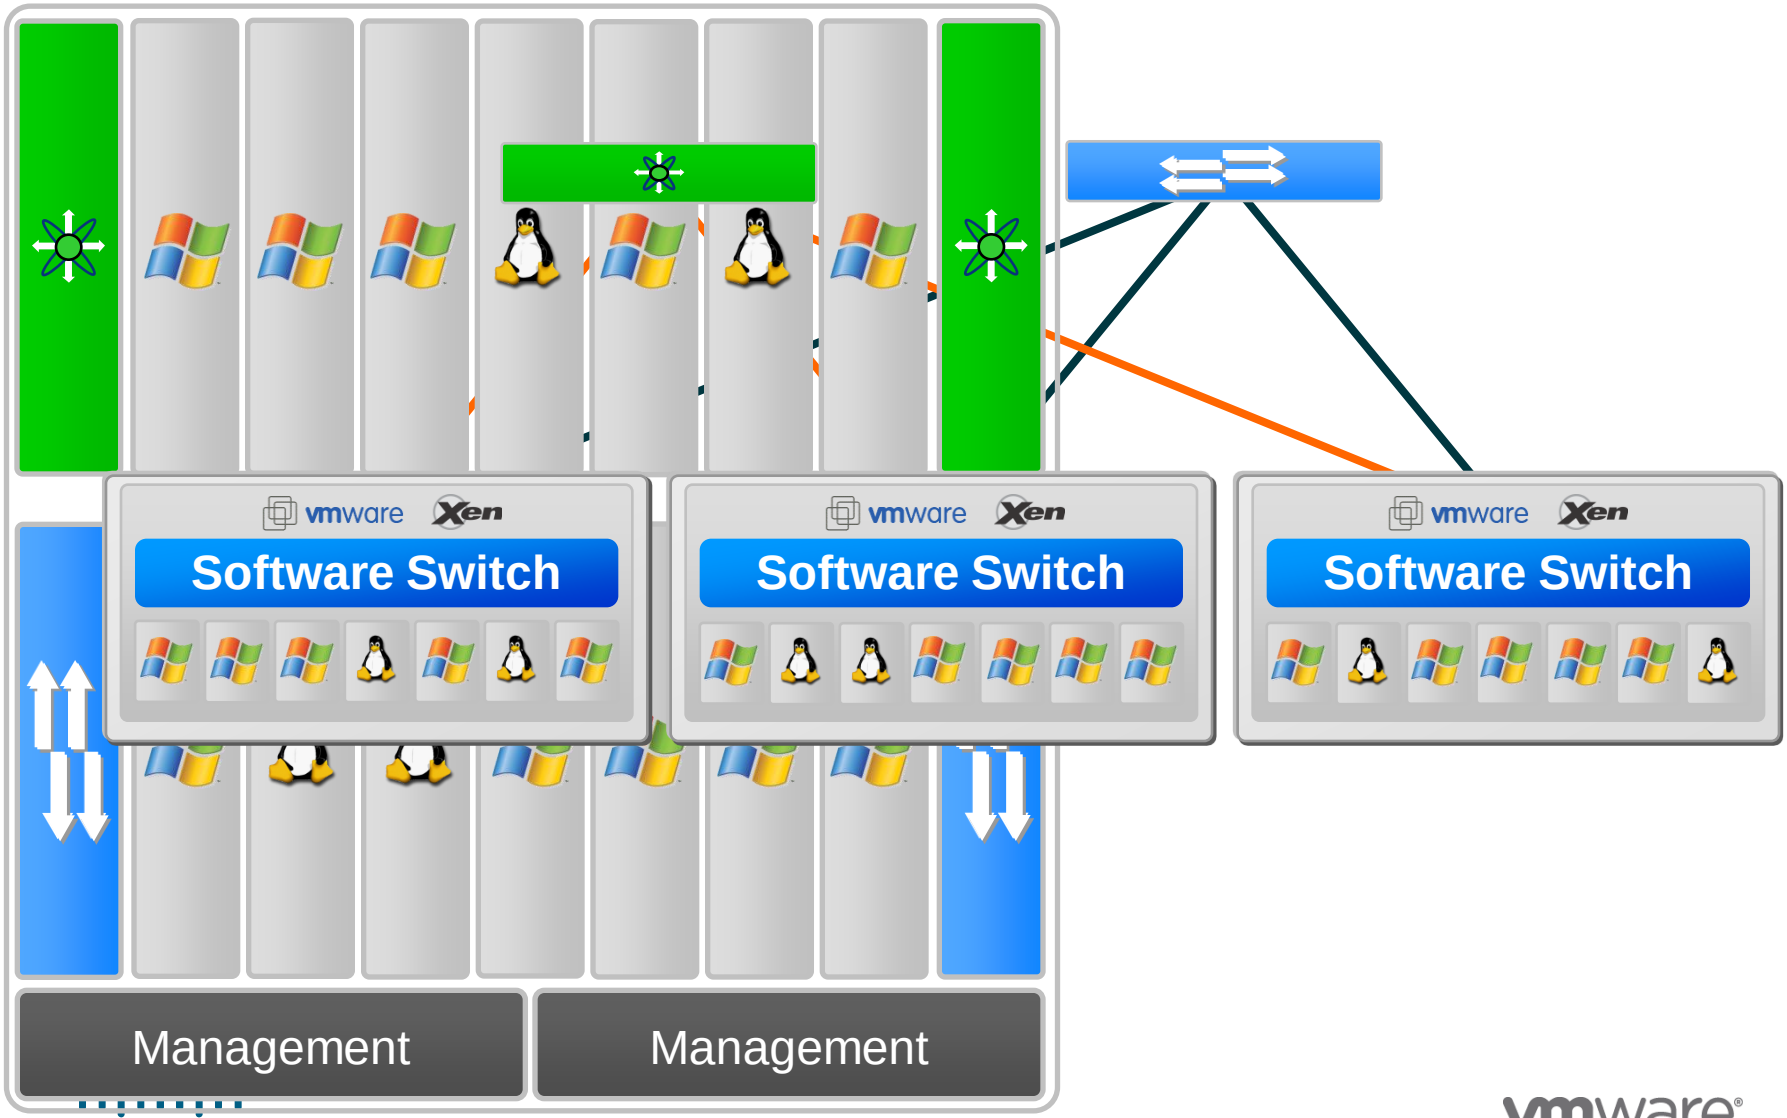

Server Deployment Virtualization

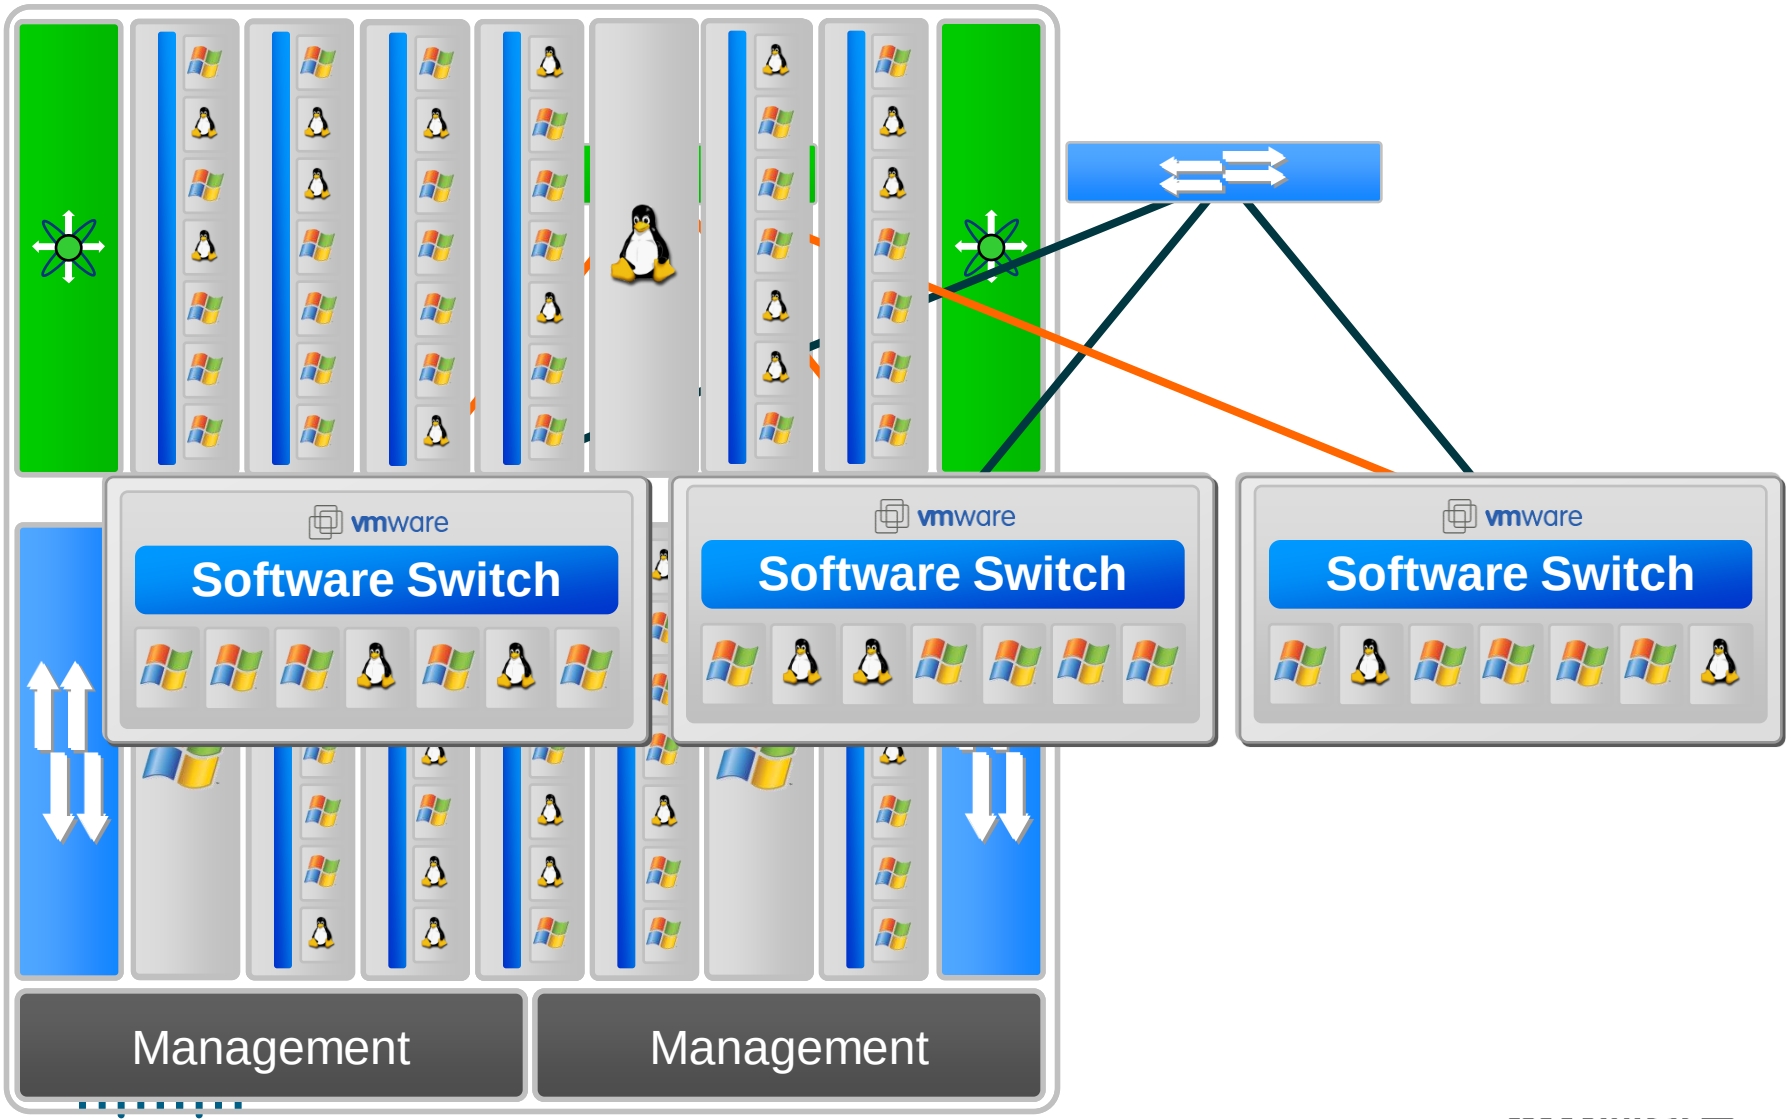

Server Deployment Today

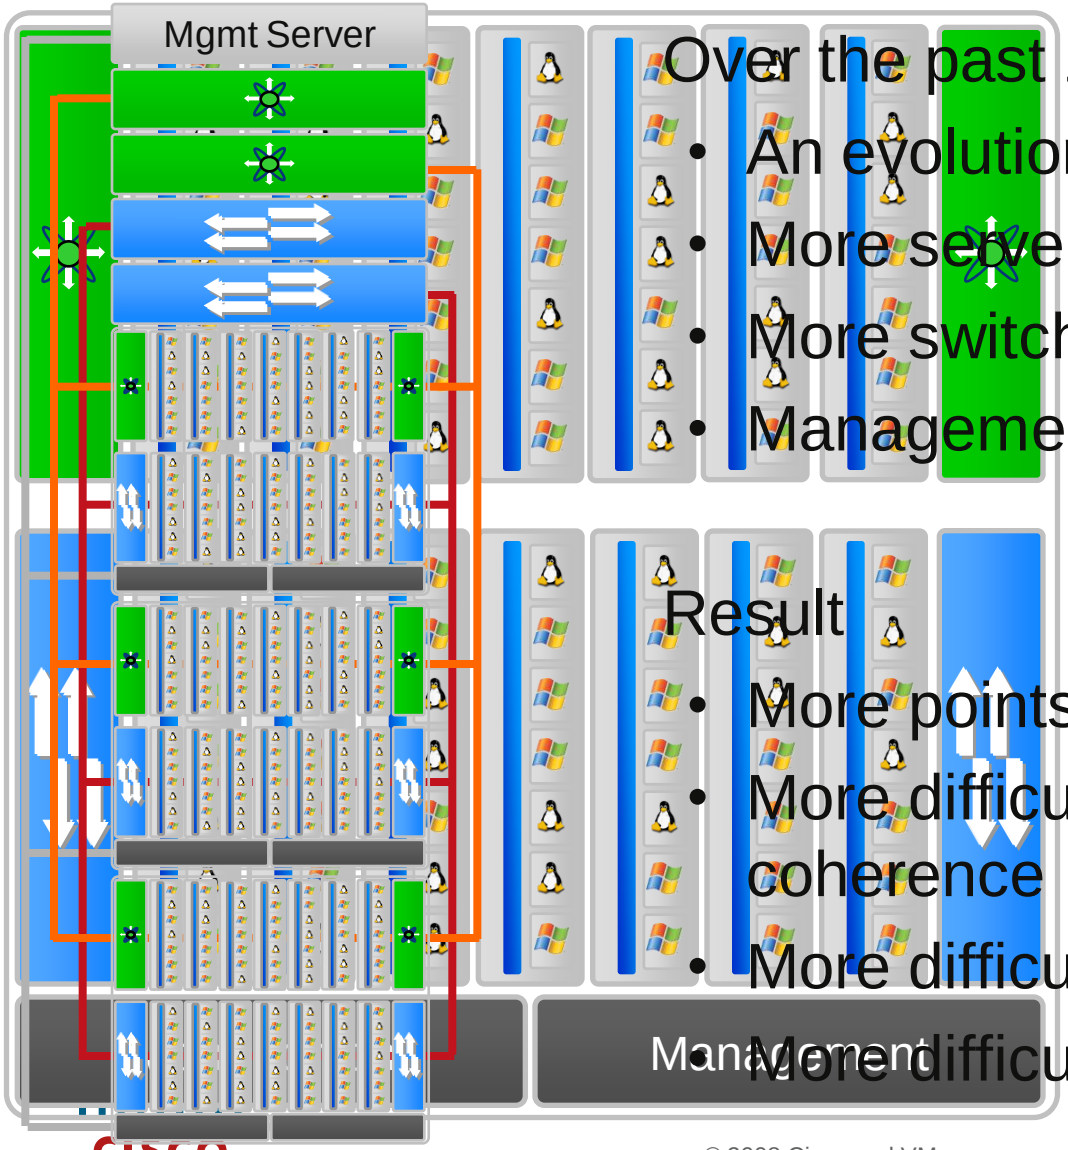

Over the past 10 years

- 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

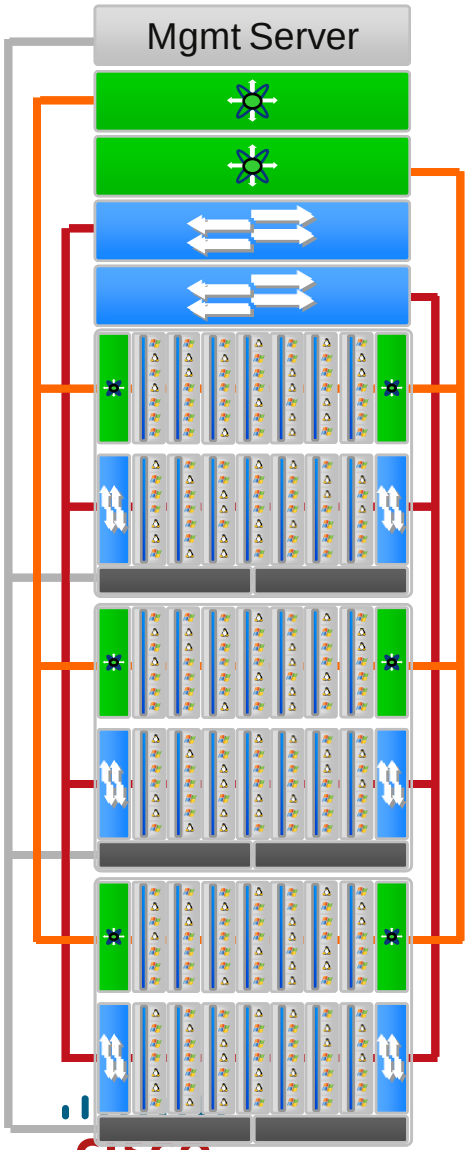

Embed management

Unify fabrics

Optimize virtualization

Remove unnecessary

- switches,
- adapters,
- management modules

Less than 1/3rd infrastructu

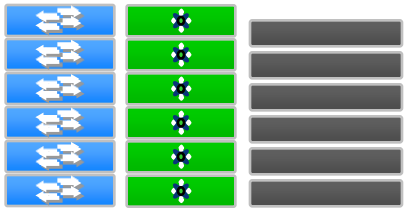

Cisco Unified Computing Solution

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

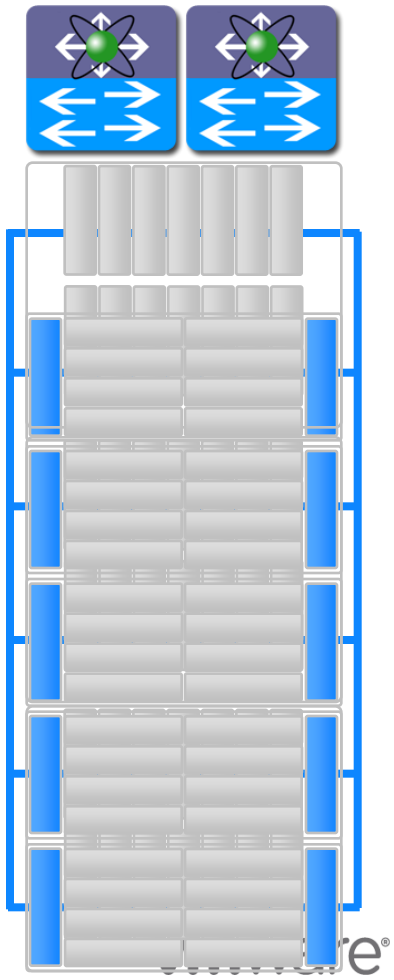

Building Blocks

UCS Manager
Embedded– manages entire system

UCS Fabric Interconnect
20 Port 10Gb FCoE
40 Port 10Gb FCoE

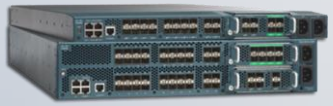

UCS Fabric Extender
Remote line card

UCS Blade Server Chassis
Flexible bay configurations

UCS Blade Server
Industry-standard architecture

UCS Virtual Adapters
Choice of multiple adapters

System Components

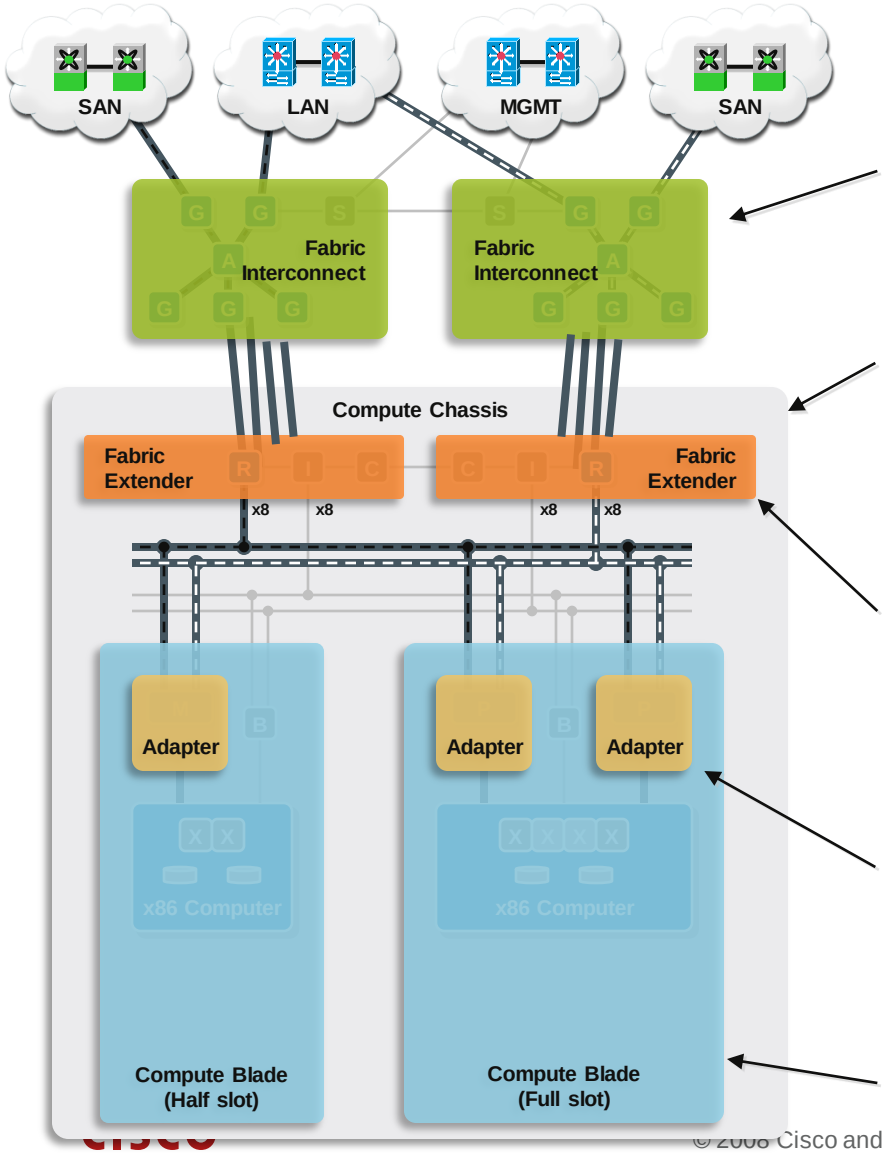

Fabric Interconnect

(40 or 20 10GE ports) + (2 or 1 GEM slots)

Chassis

Upto 8 half width blades or 4 full width blades

Fabric Extender

Host to uplink traffic engineering

Up to 80Gb Flexible bandwidth allocation (4 10G Links per FEX)

Adapter

Virtualized adapter for single OS and hypervisor systems

Compute Blade

Overall System Connectivity

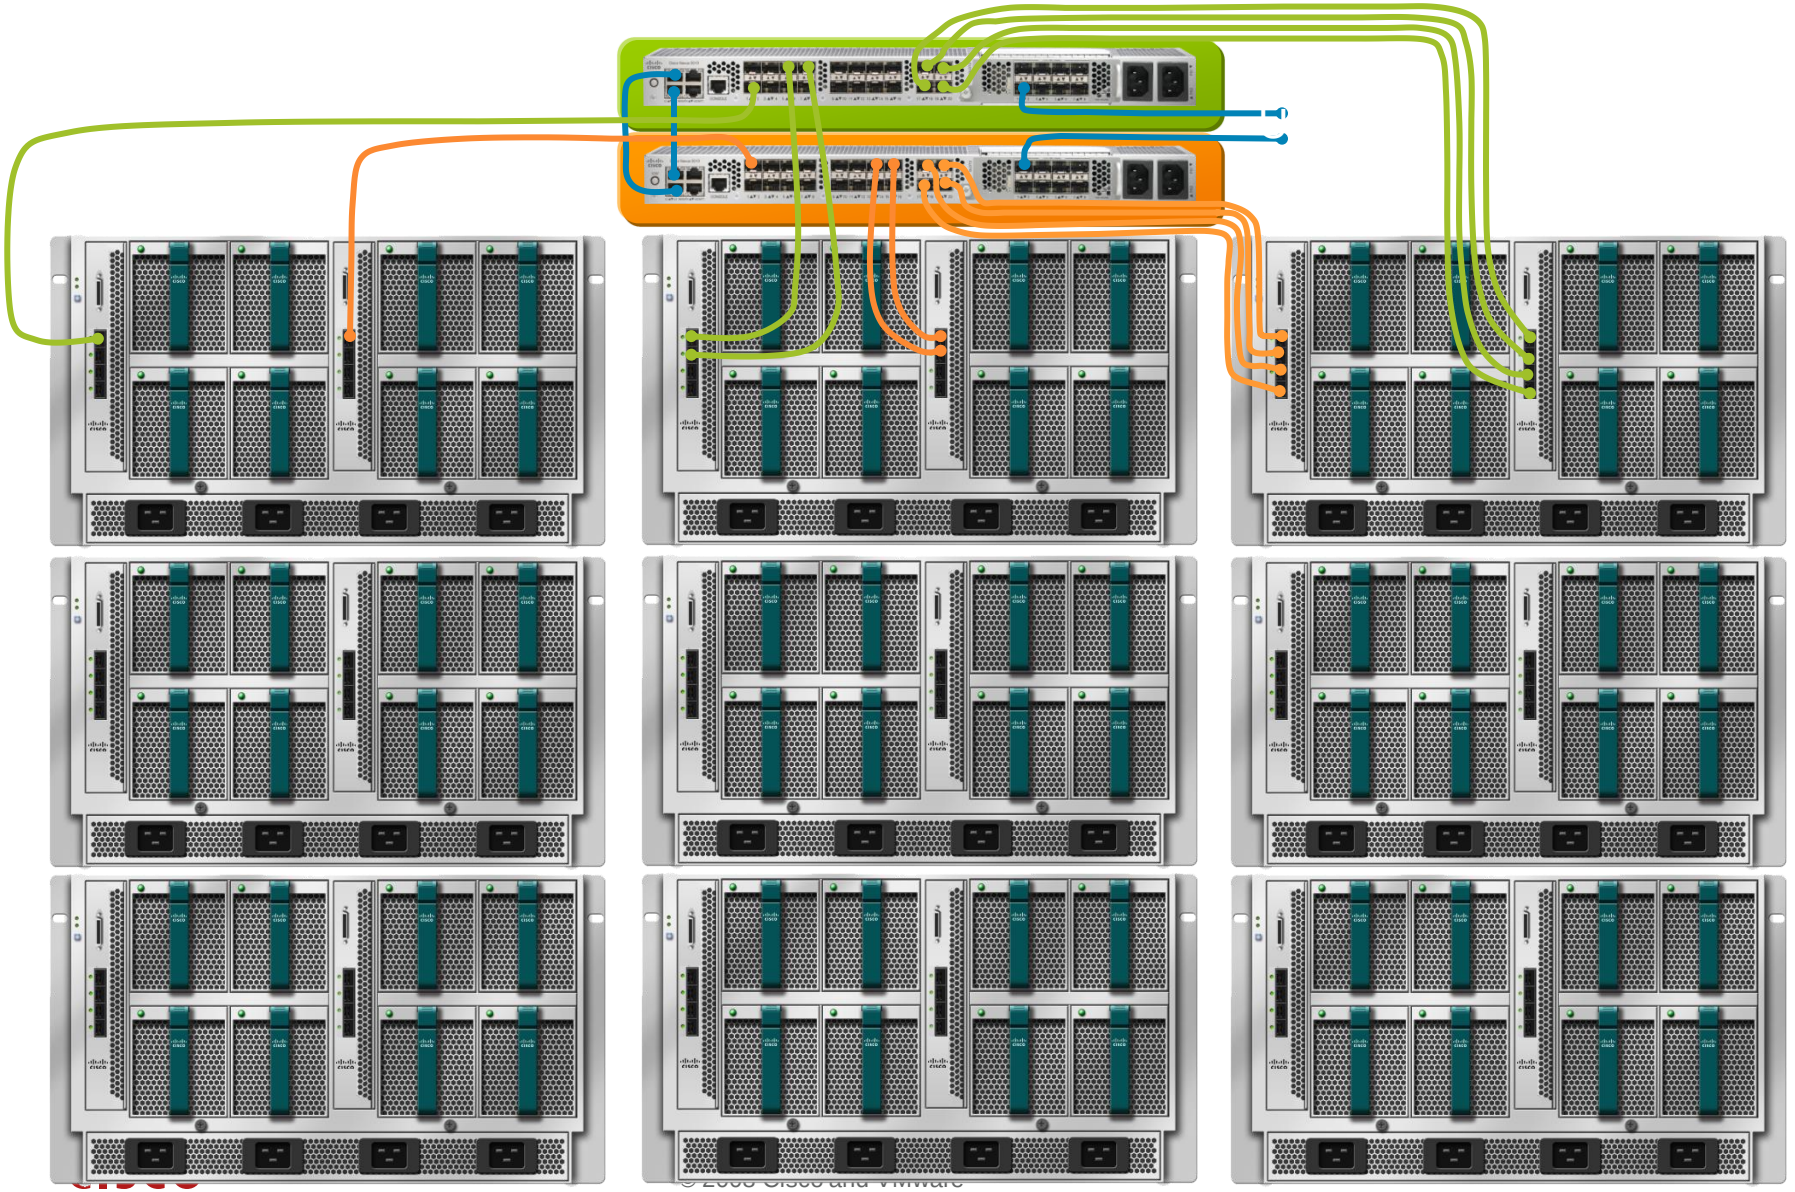

Agenda

Introduction to UCS

UCS System Components
- UCS Manager
- Interconnect
- Fabric Extender
- Enclosure
- Compute Node
- IO Adapters

UCS Differentiators

UCS System Components

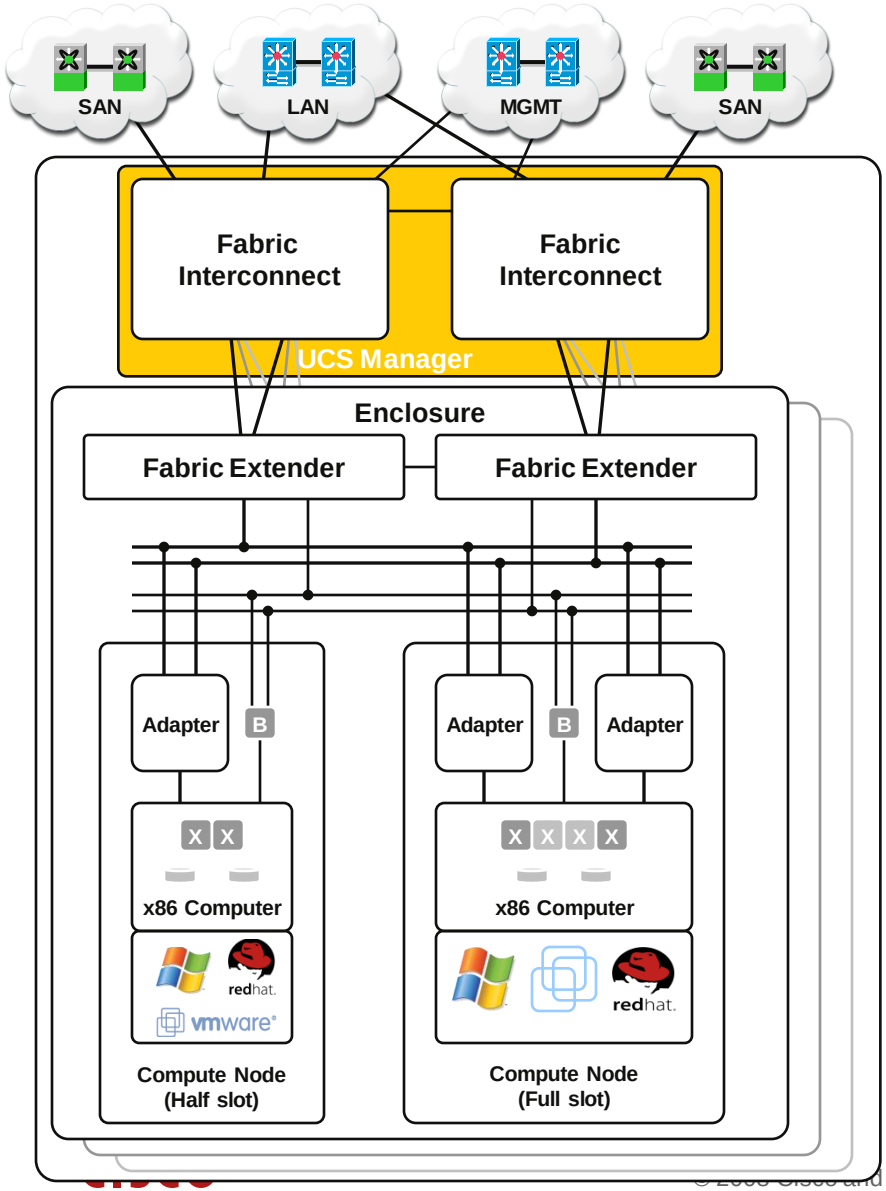

UCS

- **UCS Manager**
- Interconnect
- Fabric Extender
- Enclosure
- Compute Node
- IO Adapter(s)

Unified Computing System Manager

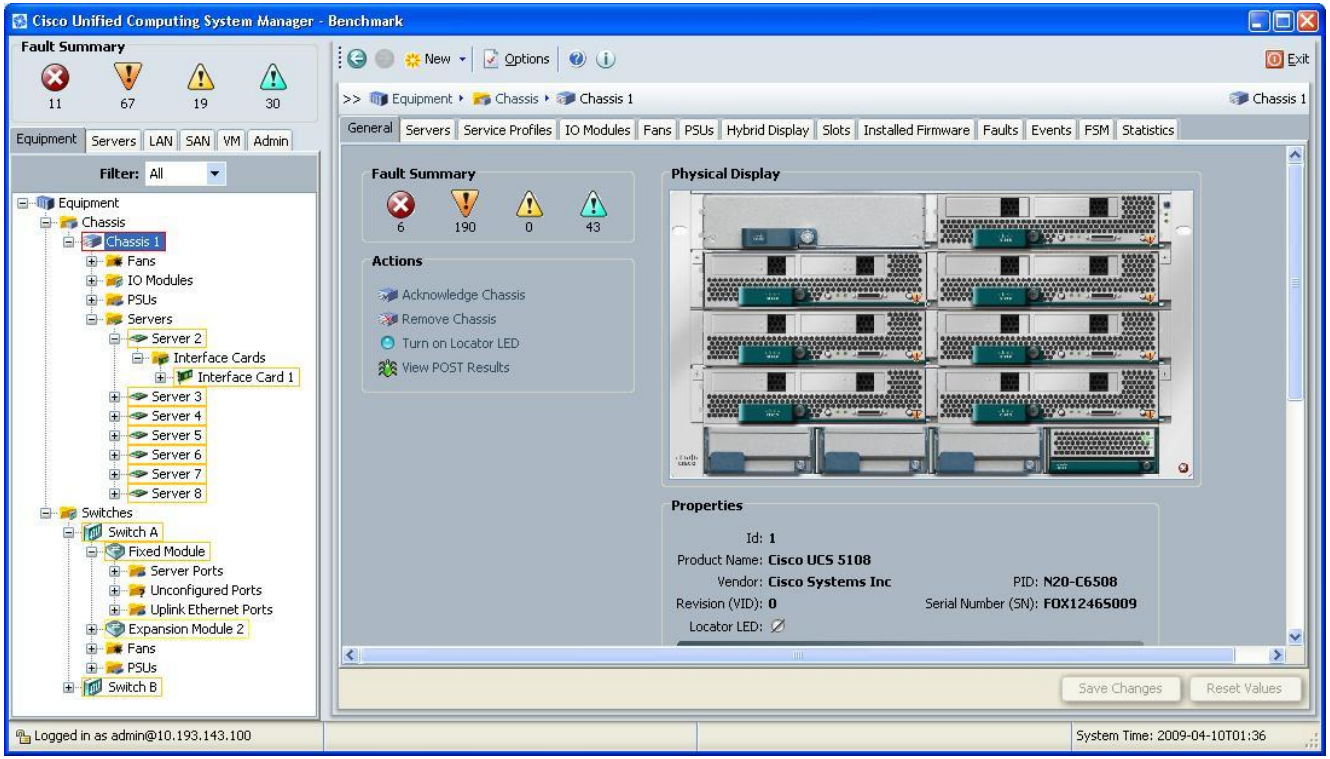

Embedded device manager for family of UCS components

Enables stateless computing via Service Profiles

Efficient scale: Same effort for 1 to 320 blades

APIs for integration with new and existing data center infrastructure

UCS Manager

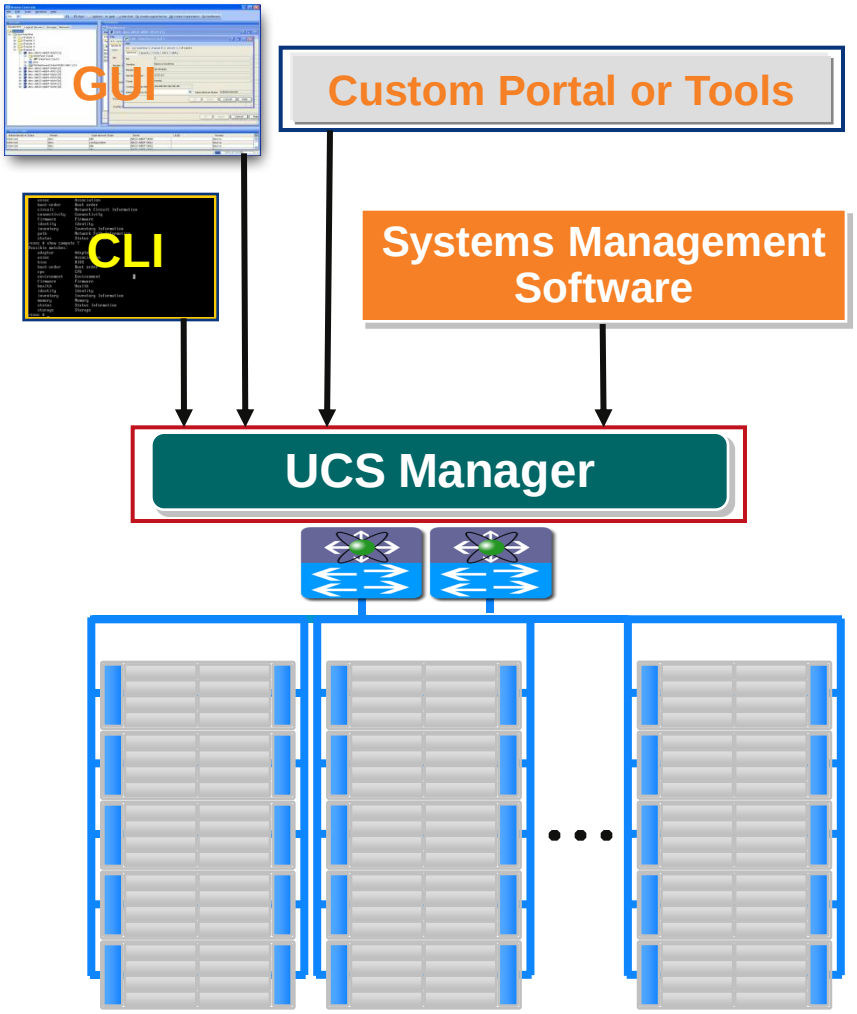

Single point of management for UCS system components

- Adapters, blades, chassis, fabric extenders, fabric interconnects

Embedded device manager

- Discovery, Inventory, Configuration, Monitoring, Diagnostics, Statistics Collection
- Coordinated deployment to managed endpoints

APIs for integration with new and existing data center infrastructure

- SMASH-CLP, IPMI, SNMP
- XML-based SDK for commercial & custom implementations

UCS System Components

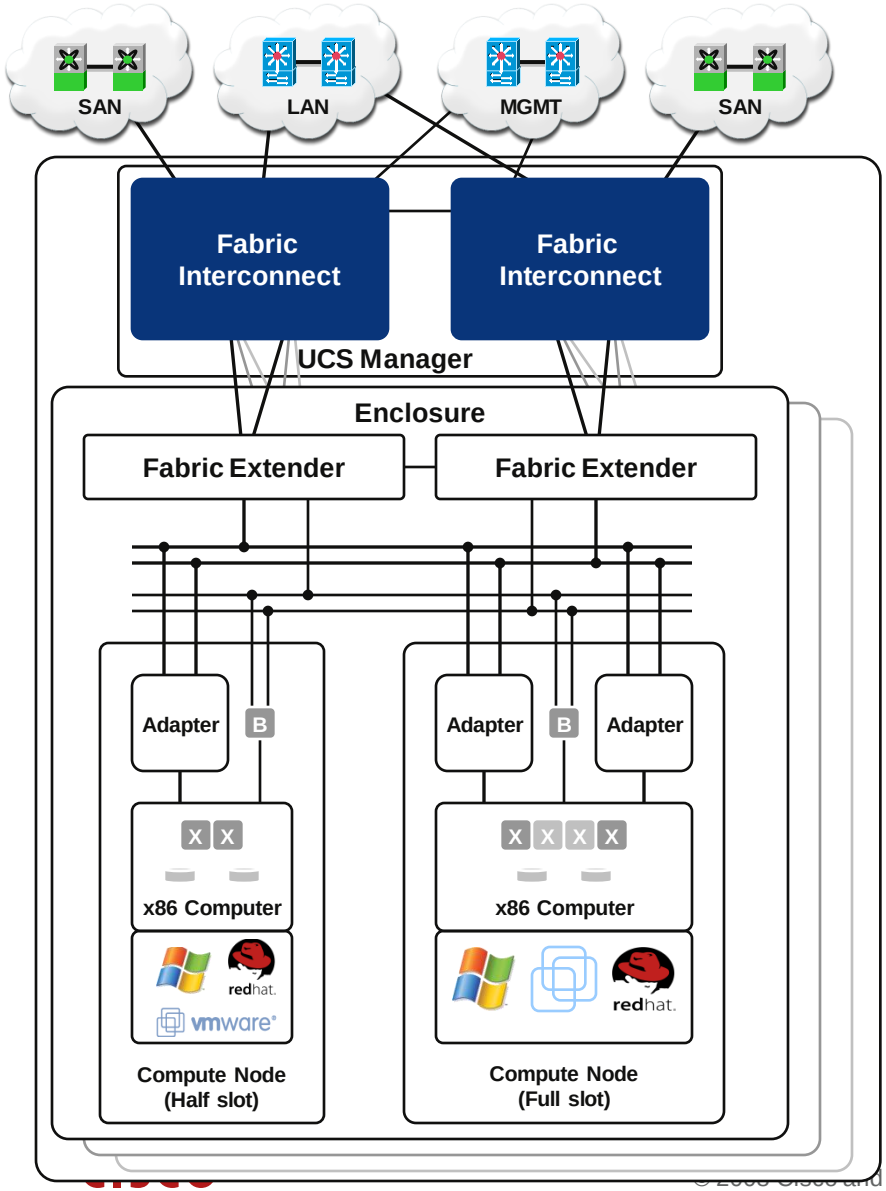

- UCS
  - California Manager
  - **Interconnect**
  - Fabric Extender
  - Enclosure
  - Compute Node
  - IO Adapter(s)

UCS 6100 Series Fabric Interconnects

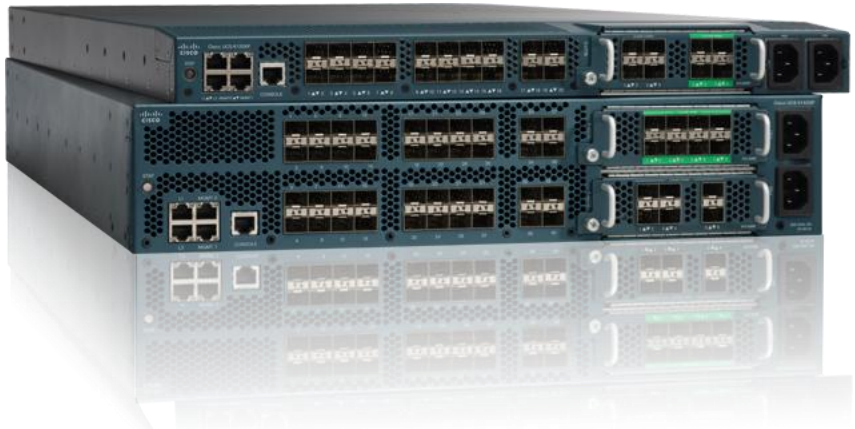

6100 Series Fabric Interconnects

10 Gigabit Ethernet, FCoE capable, SFP+ ports

20 and 40 fixed port versions with Expansion slots for additional Fiber Channel and 10 GE connectivity

Up to to 1.04 Tbps of throughput

Hot pluggable fan and power supplies

Hardware based support for Cisco VN-Link technology

Supports up to 40 chassis per UCS system

UCS Fabric Interconnect Portfolio

UCS 6100 Family

20-Port Fabric Interconnect
- 20 fixed ports 10GE/FCoE, fixed
- 1 Expansion Module

40-Port Fabric Interconnect
- 40 fixed ports 10GE/FCoE, 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
- 4 Ports 1/2/4G FC

Ethernet
- 6 Ports 10GE/FCoE

UCS System Components

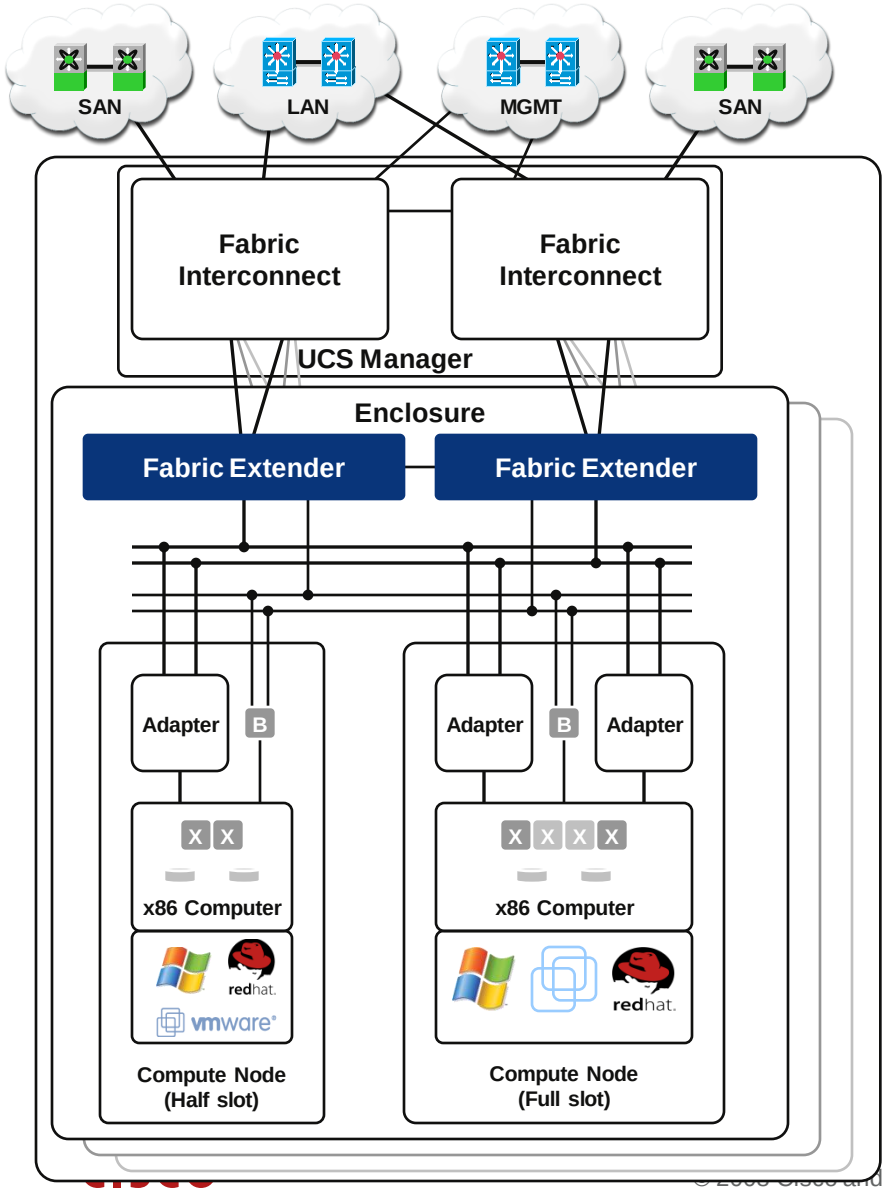

- UCS
  - UCS Manager
  - Interconnect
  - **Fabric Extender**
  - Enclosure
  - Compute Node
  - IO Adapter(s)

UCS 2100 Series Fabric Extenders

2104 Fabric Extender

Connects UCS blade chassis to the Fabric Interconnect

Four 10 Gigabit Ethernet, FCoE capable, SFP+ ports

Up to 2 Fabric Extenders per chassis for redundancy and up to 80 Gbps of bandwidth per chassis

Built-in chassis management functionality

Hardware based support for Cisco VN-Link technology

Fully managed by UCS Manager through Fabric Interconnect

UCS System Components

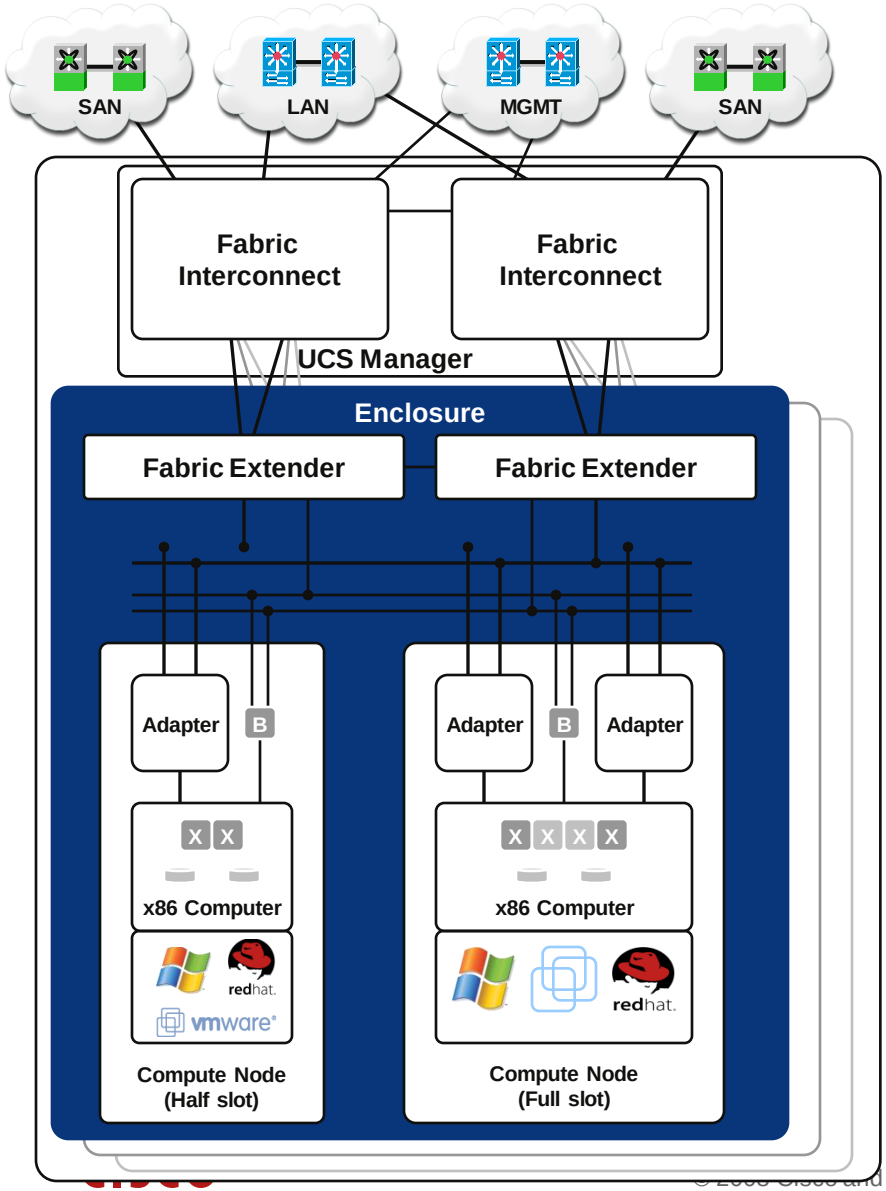

- UCS
  - UCS Manager
  - Interconnect
  - Fabric Extender
  - **Enclosure**
  - Compute Node
  - IO Adapter(s)
  - Virtualization

UCS 5108 Blade Chassis

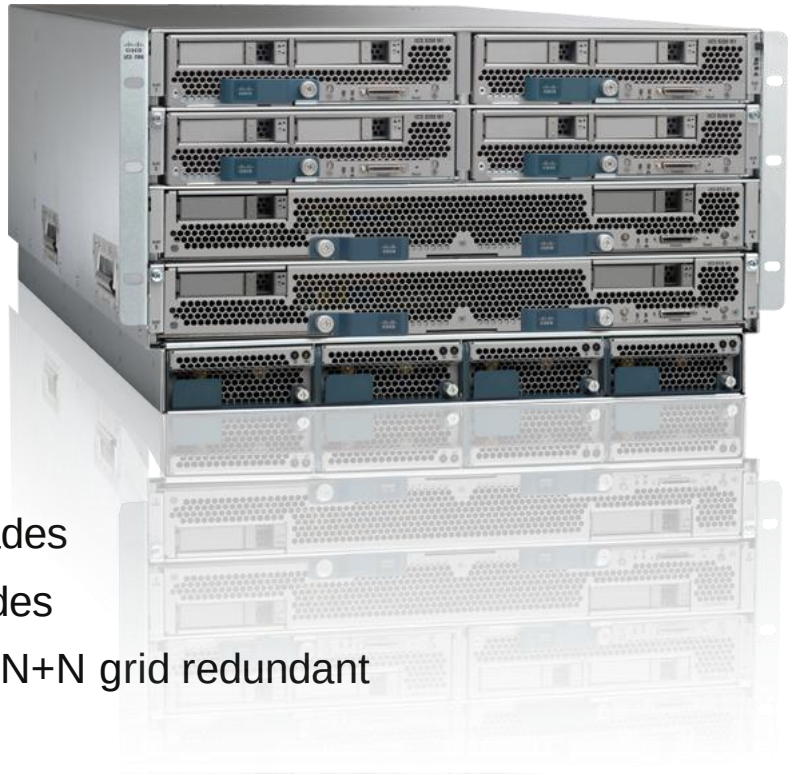

Chassis

- Up to 8 half slot blades
- Up to 4 full slot blades
- 4x power supplies, N+N grid redundant
- 8x fans included
- 2x UCS 2104 Fabric Extender
- All items hot-pluggable
- Up to 40 chassis per UCS system

UCS System Components

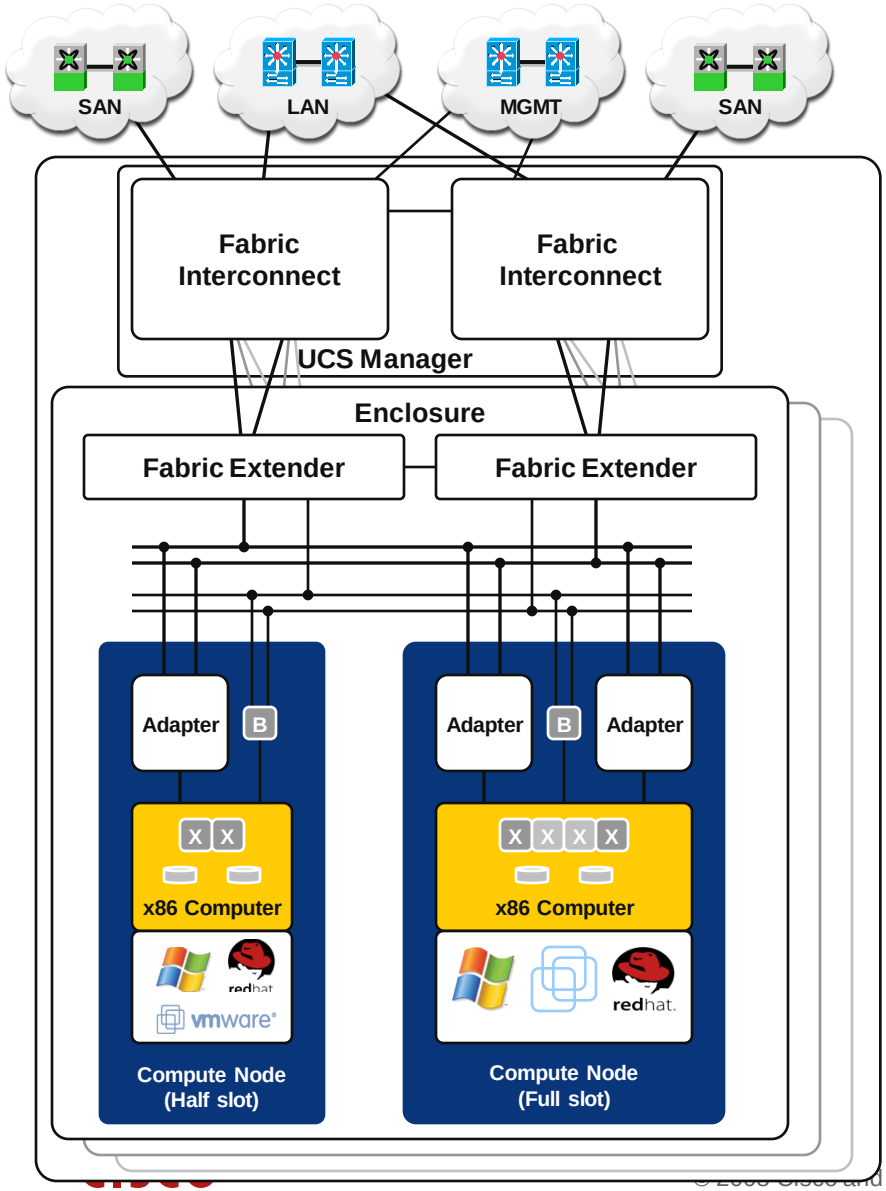

- UCS
  - UCS Manager
  - Interconnect
  - Fabric Extender
  - Enclosure
  - **Compute Node**
  - IO Adapter(s)

UCS B200 M1/M2 Blade

Blade Attributes

- Stateless design
- M1: 2x Intel Xeon 5500 Series Processors (4 Cores)
- M2: 2x Intel Xeon 5600 Series Processors (6 Cores)
- 12x DIMM slots - up to 96GB RAM
- 2x optional SAS hot-plug hard drives
- RAID 0, 1
- 1x 10Gb dual port mezzanine adapter
- Remote and local access to keyboard, video, mouse, serial
- Integrated with UCS Manager
- Up to 8 blades per UCS 5108 Blade Chassis

UCS B250 M1/M2 Blade

Blade Attributes

- Stateless design
- M1: 2x Intel Xeon 5500 Series Processors (4 Cores)
- M2: 2x Intel Xeon 5600 Series Processors (6 Cores)
- 48x DIMM slots - up to 384GB RAM
- 2x optional SAS hot-plug hard drives
- RAID 0, 1
- 2x 10Gb dual port mezzanine adapter
- Remote and local access to keyboard, video, mouse, serial
- Integrated with UCS Manager
- Up to 4 blades per UCS 5108 Blade Chassis

VMmark Benchmarks

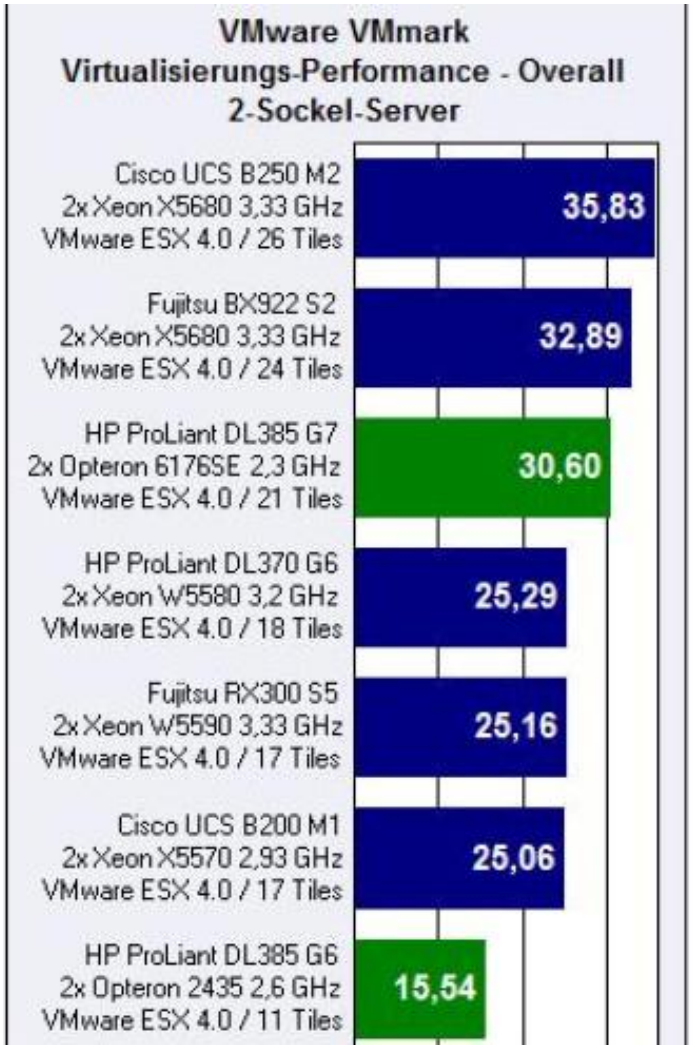

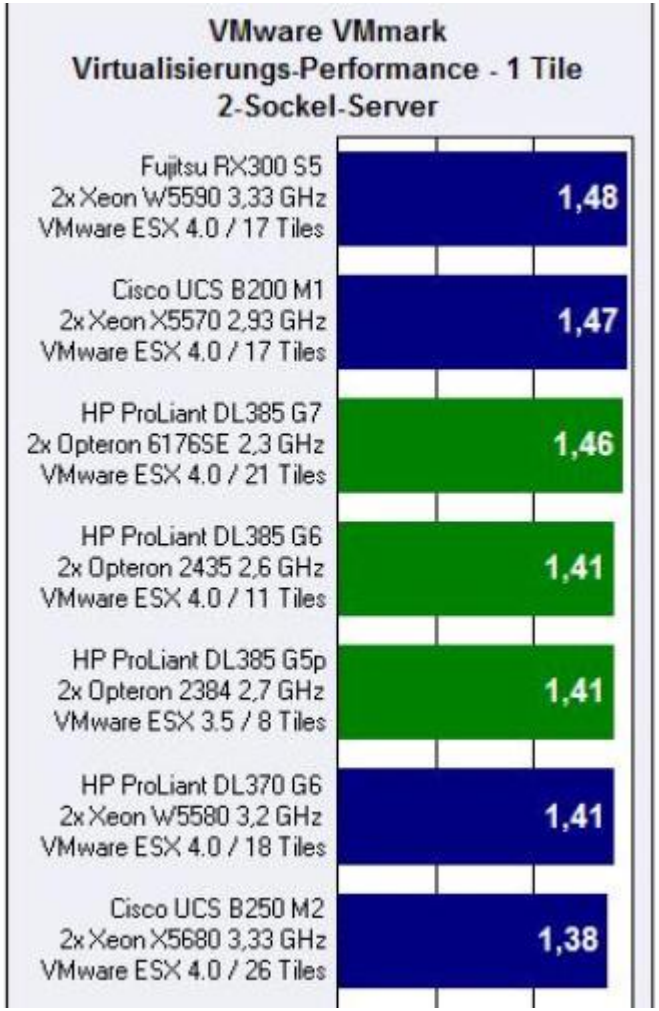

UCS B440 M1 Specification

4 x Intel Xeon 7500 Series Processors

32 x DIMM slots, up to 256GB RAM

4 x optional SFF SAS or SSD hot-plug hard drives

RAID 0, 1 standard

RAID 5, 6 optional

Battery backed cache (Optional)

2 x 10Gb dual port mezzanine adapter

Remote and local access to keyboard, video, mouse, serial

Integrated with UCS Manager

Up to 4 blades per
UCS 5108 Blade Chassis

New B-Series Blades: Westmere and Nehalem-EX

UCS B200 M2
General Purpose Blade Server

High-density server with balanced compute performance and I/O flexibility

UCS B250 M2
Extended Memory Blade Server

Memory-intensive server for virtualized and large-data-set workloads

UCS B440 M1 High-Performance Blade Server

Compute & memory-intensive server for enterprise-critical workloads

Item	Size	CPU Sockets/ Cores	CPU	Memory	Disks	I/O
UCS B200 M2	Half	2/6	Intel Xeon 5600	12 DIMM 96 GB	2 SFF SAS	1 Mezz
UCS B250 M2	Full	2/6	Intel Xeon 5600	48 DIMM 384 GB	2 SFF SAS	2 Mezz
UCS B440 M1	Full	4/8	Intel Xeon 7500	32 DIMM 256GB	4 SFF SAS/SATA	2 Mezz

Mapping Requirements to Systems

Key Application Drivers - Absolute Performance -		B250 & C250 Extended Memory Xeon 5600		B440 & C460 Xeon 7500	
Sensitive Apps (*Generalized*)	**Architectural Element**	**Per Processor**	**Per System**	**Per Processor**	**Per System**
Virtualization	Max Core Count	6	12	8	**32**
HPCC, Virtualization	Max Core Speed	**3.33GHz**	160 GFLOP/S	2.26GHz	**288 GFLOP/S**
	Max Memory Speed	**1333MHz**		1066MHz	
OLTP / Database	Max Memory Bandwidth	1X[1]		**1.15X[1]**	
Virtualization, Scalable DB, VDI, EDA	Max Memory Size	192	**384**	**64 (B440) 128 (C460)**	256 (B440) **512 (C460)**
	Silicon RAS features	Standard		Advanced (MCA recovery etc.)	

Find out which architectural element gates your customer's application

Requirements for **$/Performance/W** optimization

Requirements for **absolute performance** and/or **mission critical** RAS features

[1] Intel measurements

UCS System Components

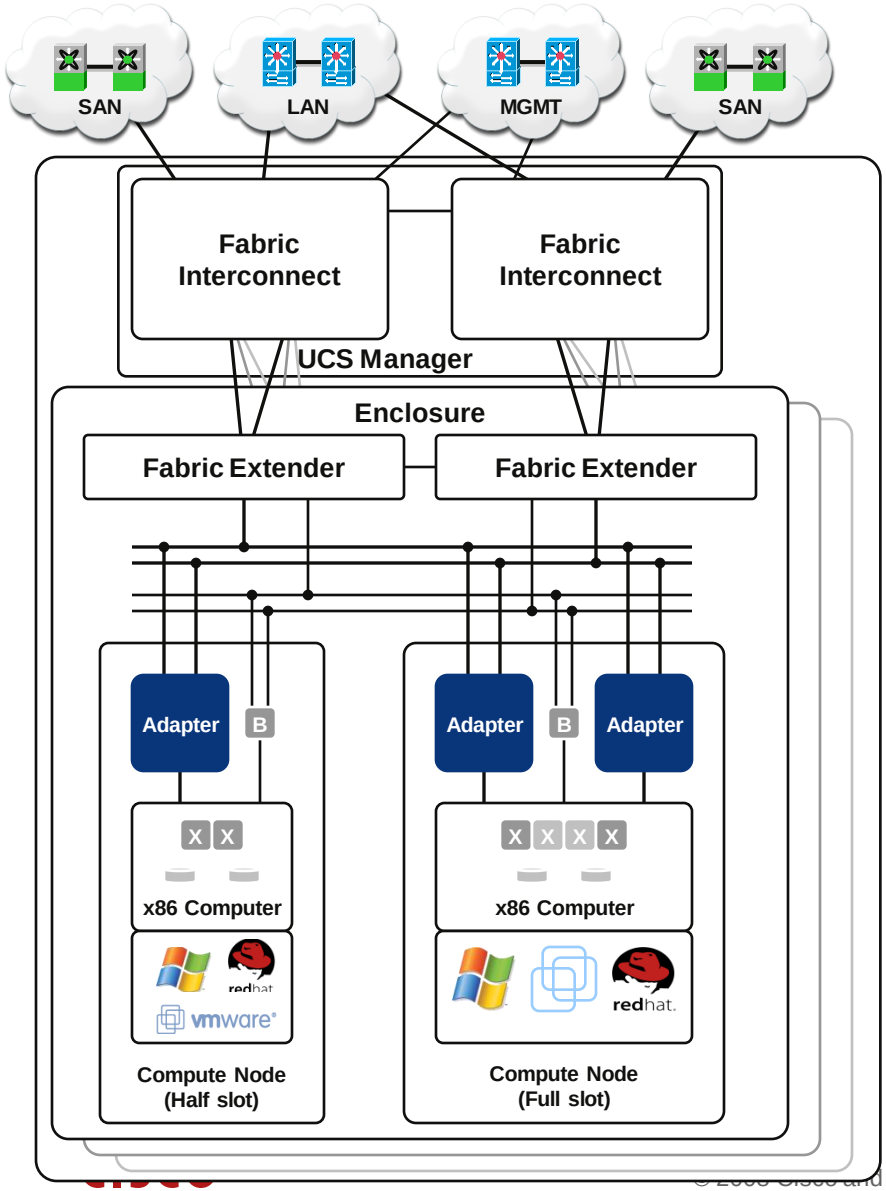

- UCS
  - UCS Manager
  - Interconnect
  - Fabric Extender
  - Enclosure
  - Compute Node
  - **IO Adapter(s)**

Three Pronged Adapter Strategy

Virtualization

Virtual Machine Aware: Virtualization and Consolidation

CISCO

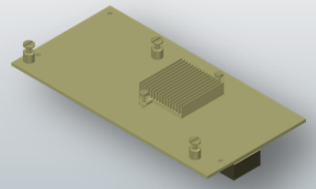

Compatibility

Existing Driver Stacks

EMULEX QLOGIC

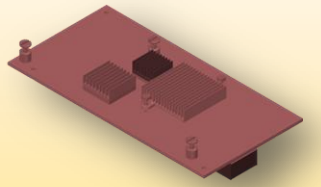

Cost

"Free" SAN Access for Any Ethernet Equipped Host

intel

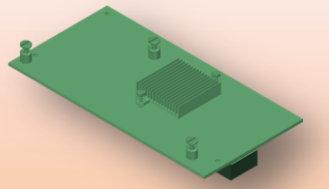

Converged network adapters (CNA)

Ability to mix and match adapter types within a system

Automatic discovery of component types

Qlogic/Emulex CNA

Cisco ASIC (Menlo)
- 7.7M gates
- 9.4Mb SRAM (including 512K of CPU SRAM)
- Embedded MIPS 24k at 350Mhz

Interfaces
- Two 10G to a 3rd party Ethernet NIC
- Two 1/2/4G to a 3rd party FC HBA
- Two 10G to an Ethernet network
- Other misc. interfaces

No changes to customer's software/drivers

I/O Consolidation, FCoE

Priority Flow Control

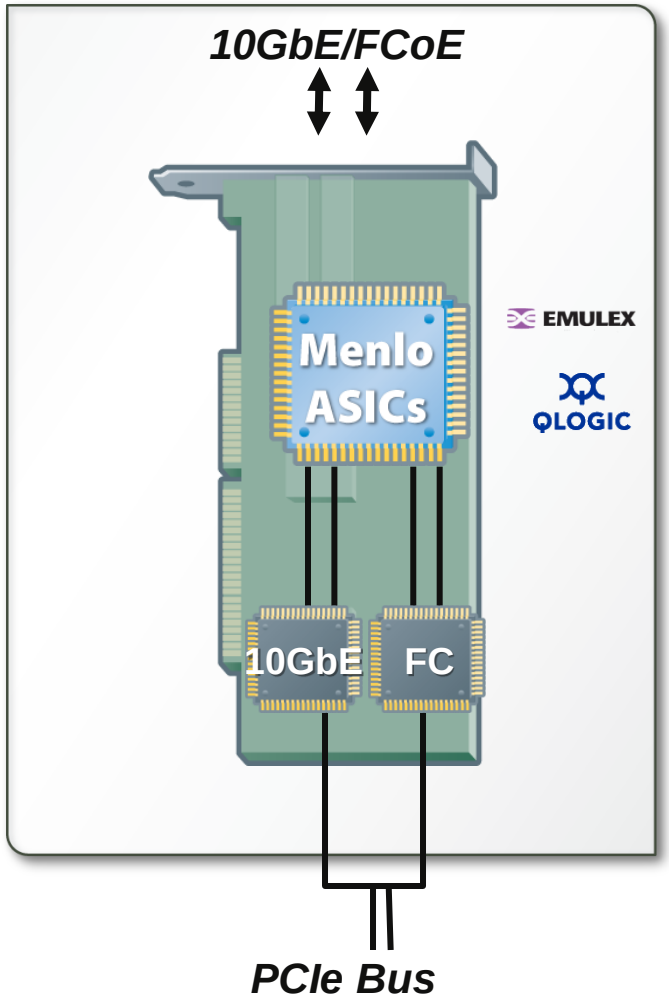

Cisco UCS M81KR VIC Overview "Palo"

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
- 600K 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

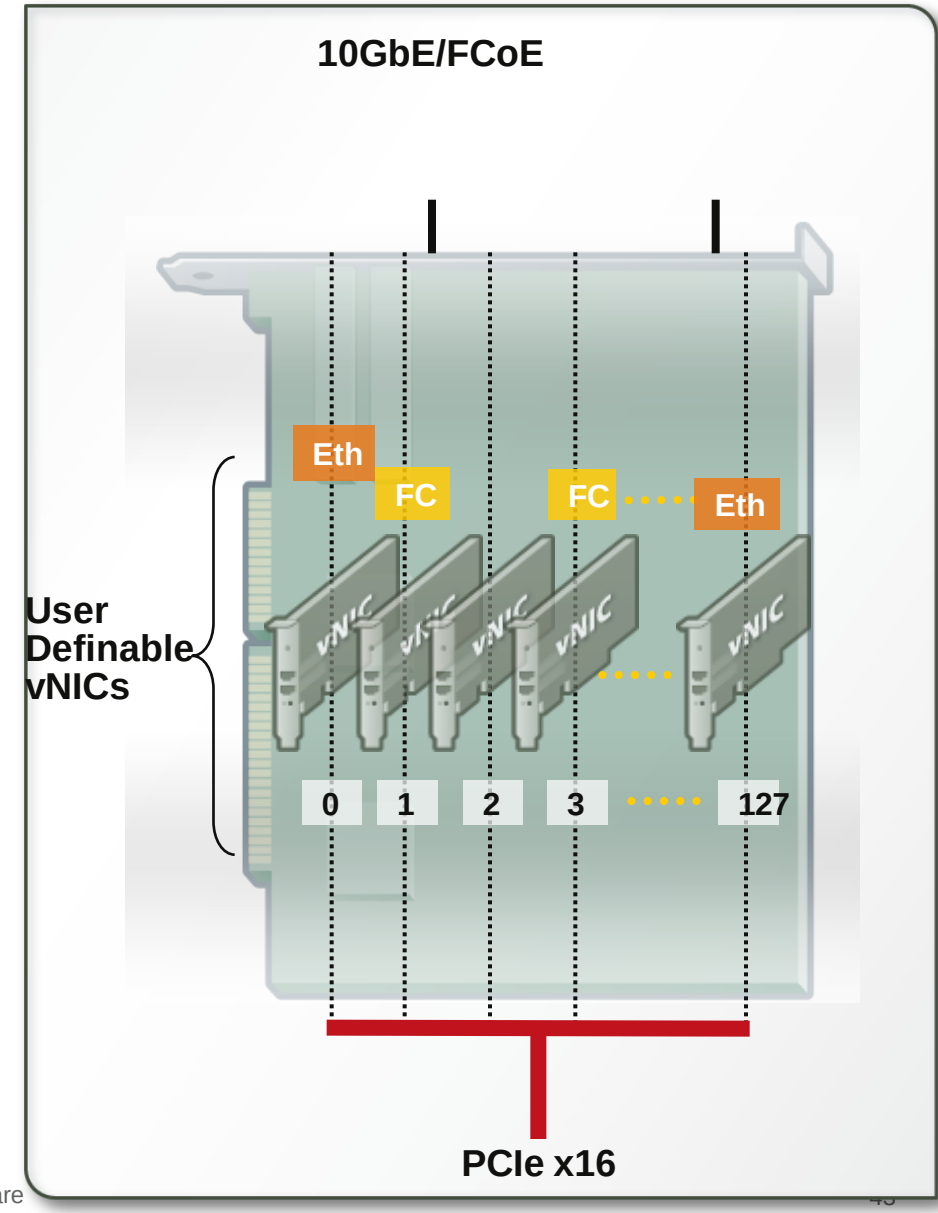

Cisco UCS VIC Overview

Multiple Separate Interfaces – Ideal for Certain Workloads

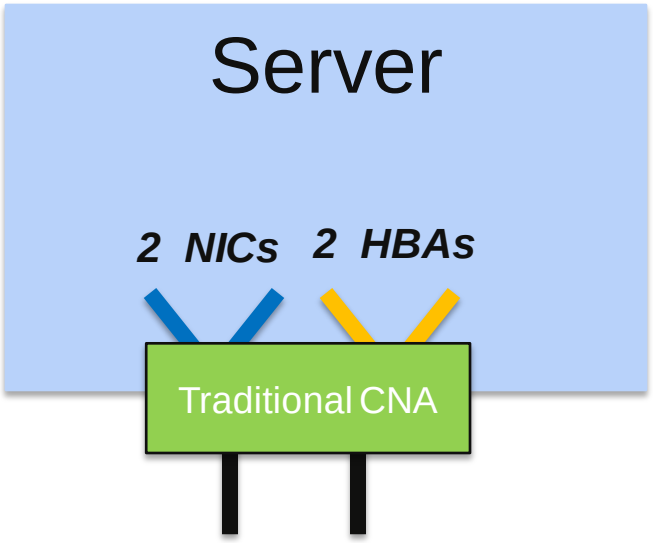

2 x 10G ports

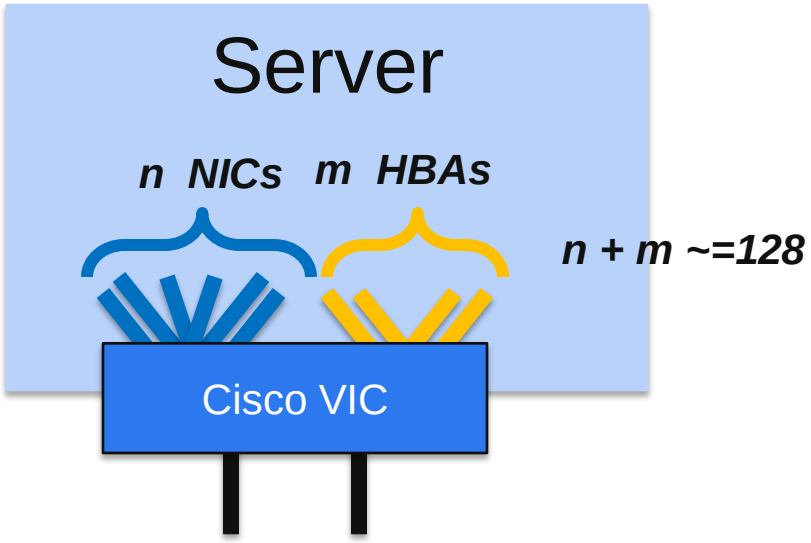

2 x 10G ports

- Ideal for workloads/applications that recommend multiple separate interfaces
- Applicable to both Single OS (e.g. Windows/RHEL) or Virtualized (ESX) environments
- Virtualization achieved using classical PCIe devices (no special OS support necessary)

Cisco VIC Offers Flexibility for Any Application

Example Use Cases

Web Servers	***Database Servers***	***Virtual Servers***
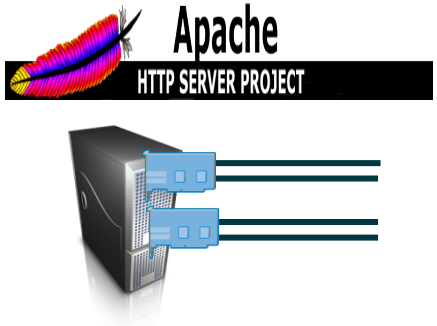	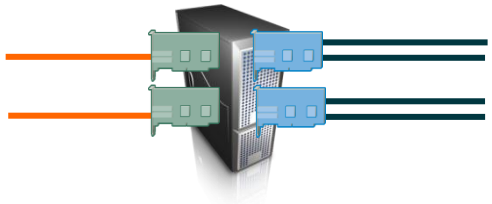	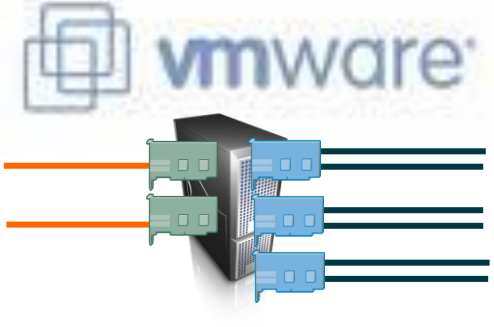
2-4 NICs	*4+ NICs & 2+ HBAs*	*4-8 NICs & 2+ HBAs*
Separate interfaces for http & Backup traffic	***SQL queries, Cluster heart beats or Cache Coherency***	***VM traffic, Service Console & Vmotion/Vmkernel***

- Ideal for workloads that recommend multiple separate interfaces as best practice
- Same Infrastructure can be used for any application – can create a true Stateless Server Farm

Agenda

- Introduction to UCS
- UCS System Components
  - UCS Manager
  - Interconnect
  - Fabric Extender
  - Enclosure
  - Compute Node
  - IO Adapters
- **UCS Differentiators**

Unified Computing System Key Differentiation

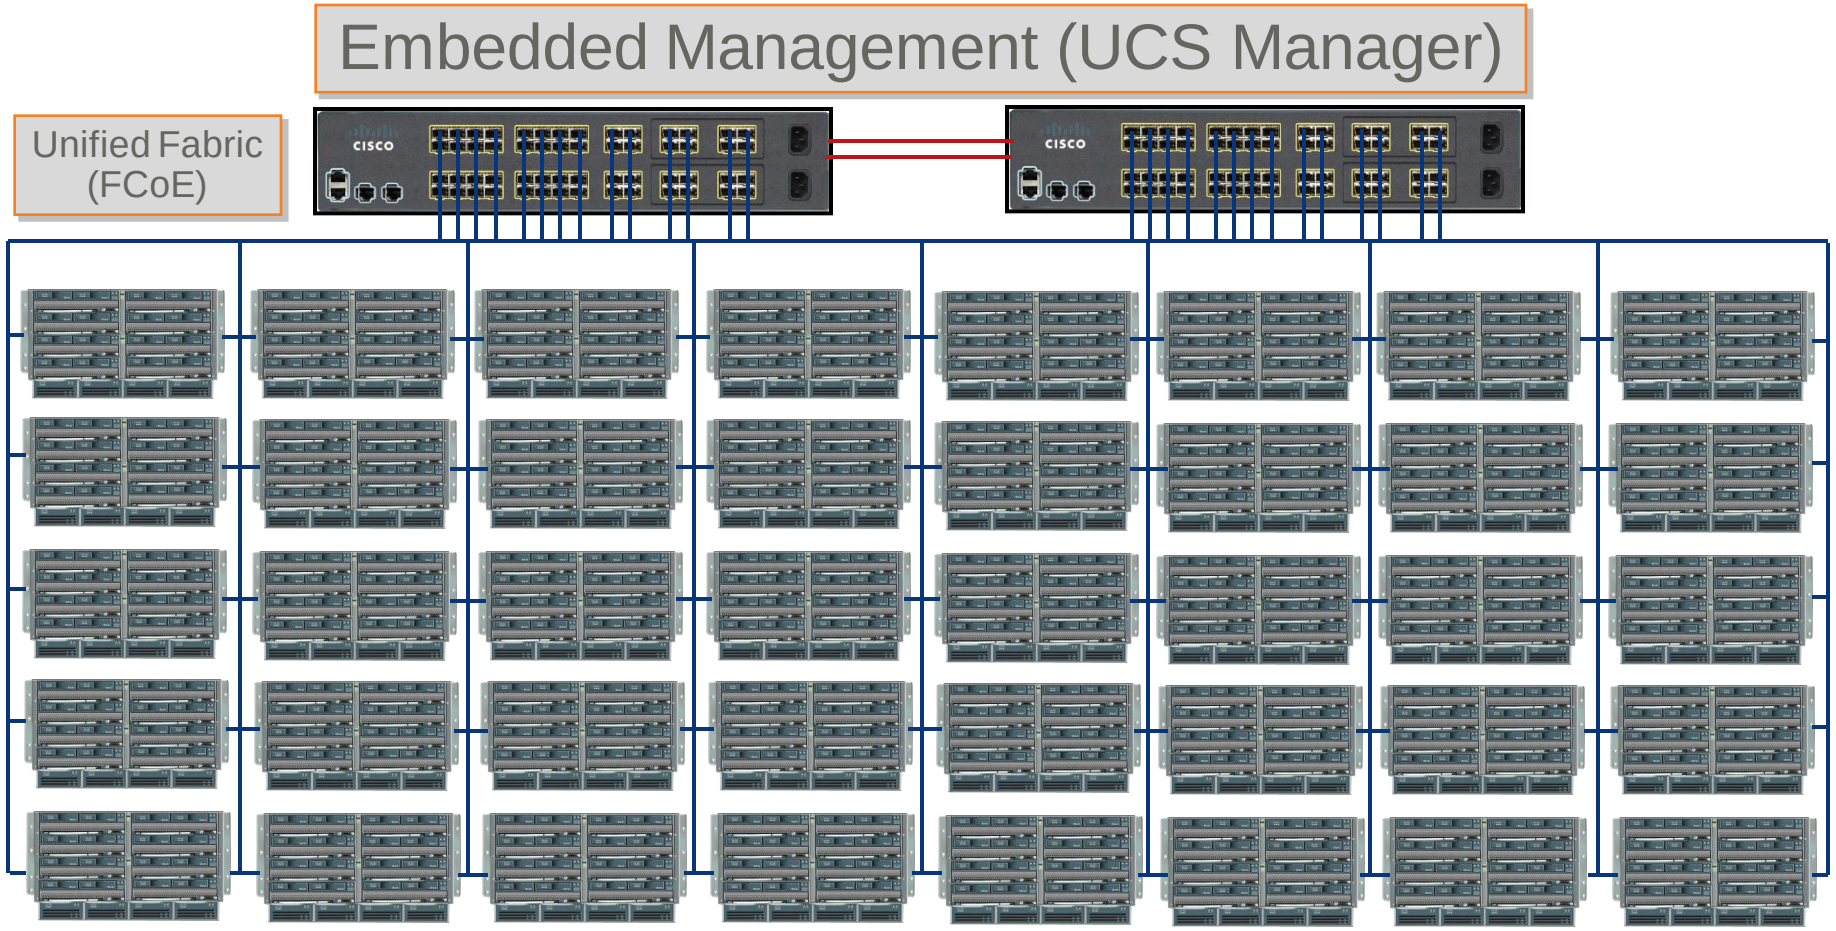

UCS Service Profiles
Hardware "State" Abstraction

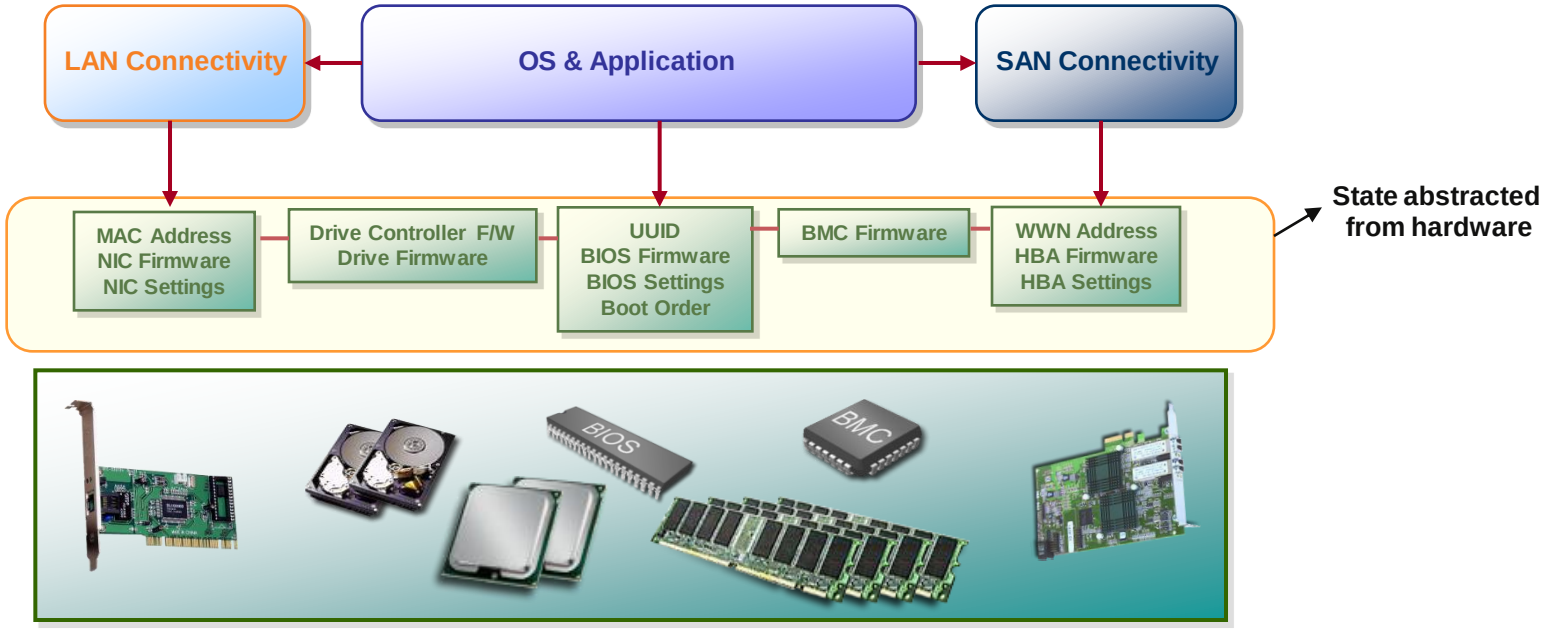

Separate firmware, addresses, and parameter settings from server hardware

Separate access port settings from physical ports

Physical servers become interchangeable hardware components

Easy to move OS & applications across server hardware

UCS Service Profiles

End to End Configure of Full UCS HW Stack

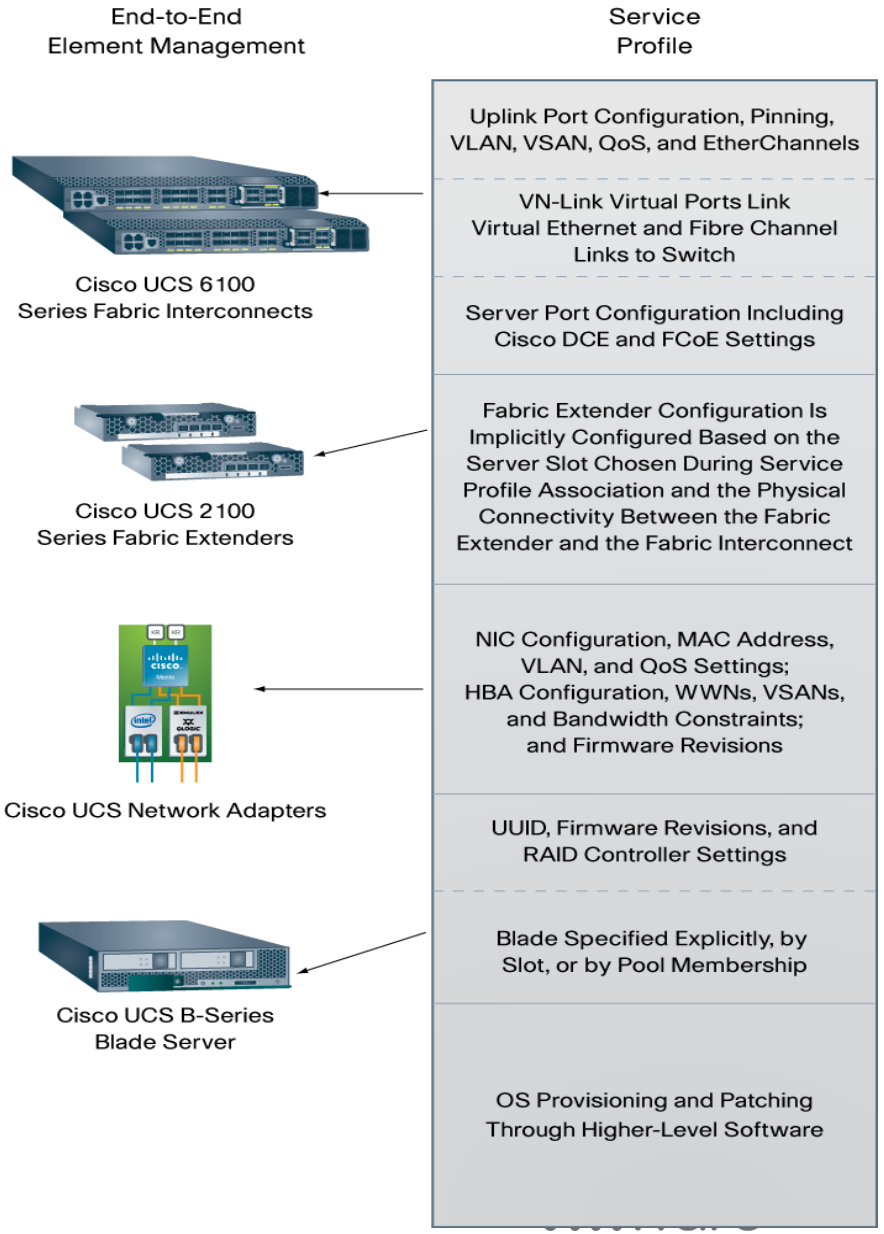

Don't I get this already from VMware? Hypervisors & Hardware State

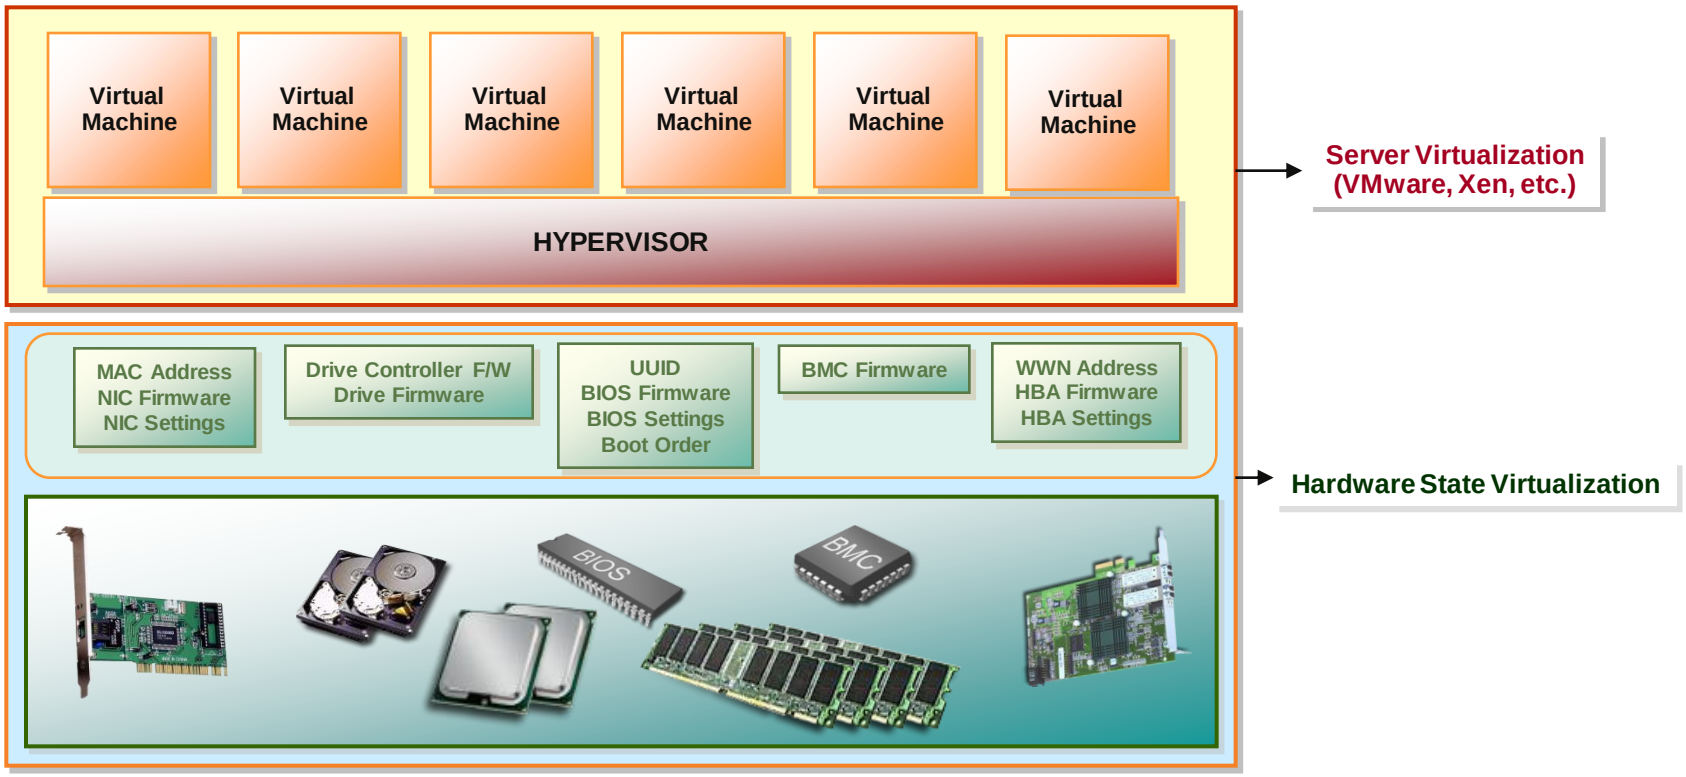

- Server virtualization & hardware state abstraction are independent of each other
- Hypervisor (or OS) is unaware of underlying hardware state abstraction

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

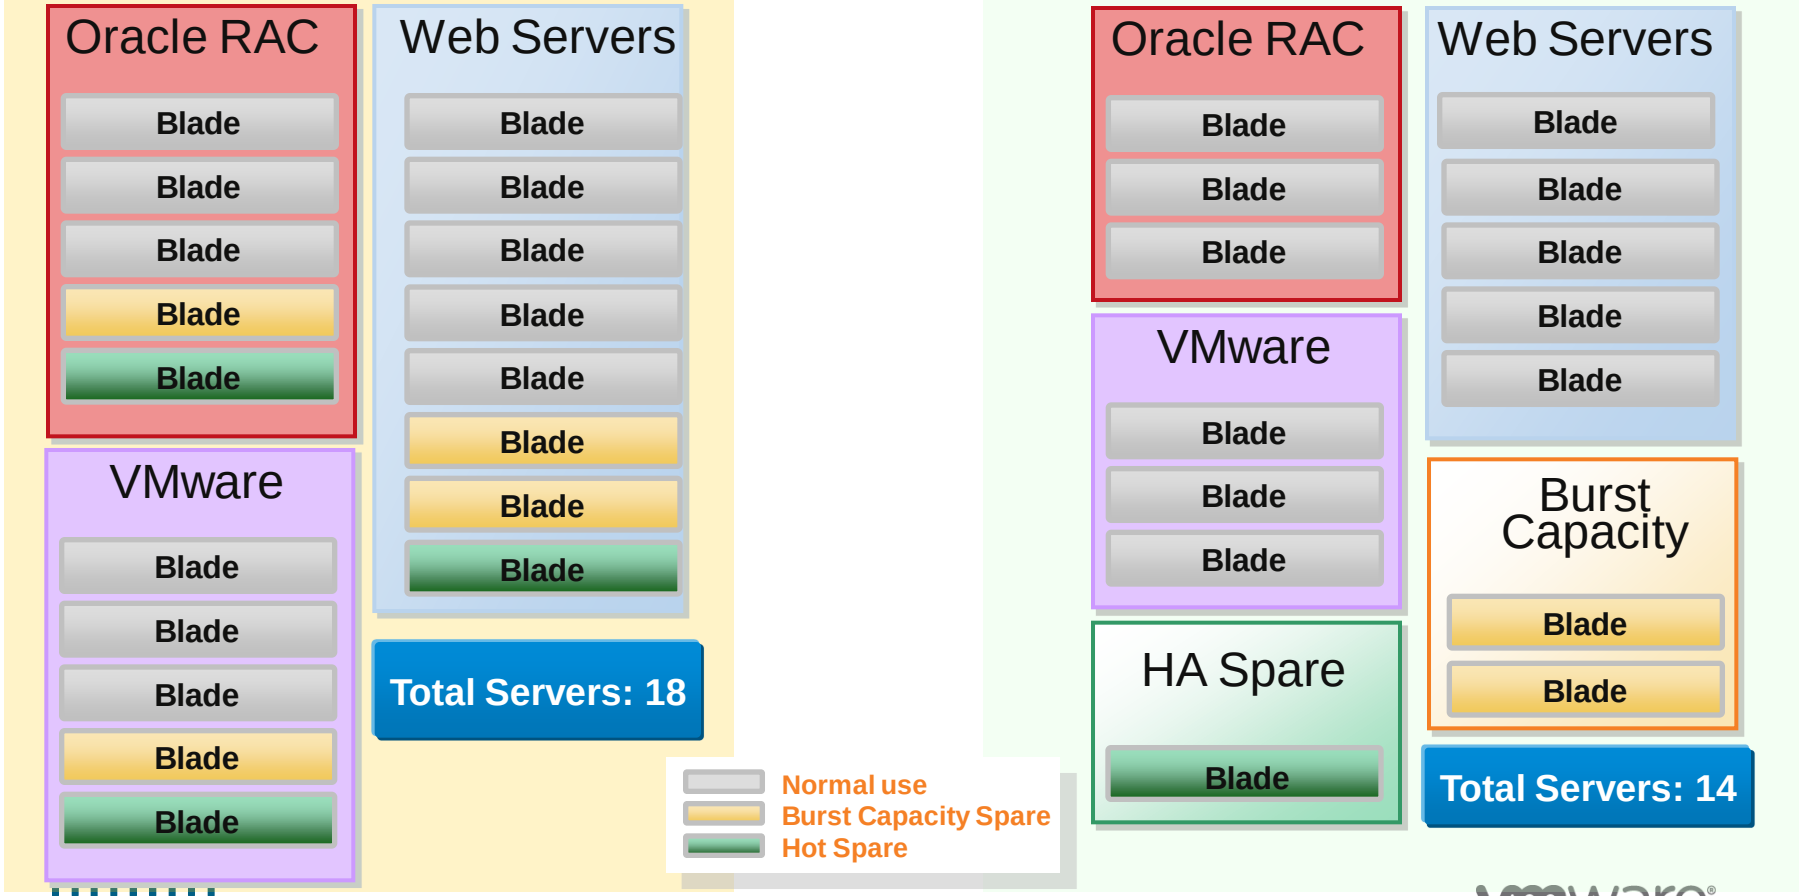

Scaling VMotion, DRS and DPM to 10GE

Intra-Cluster Mobility

Scenario: Moving all VMs from one host to another in a cluster

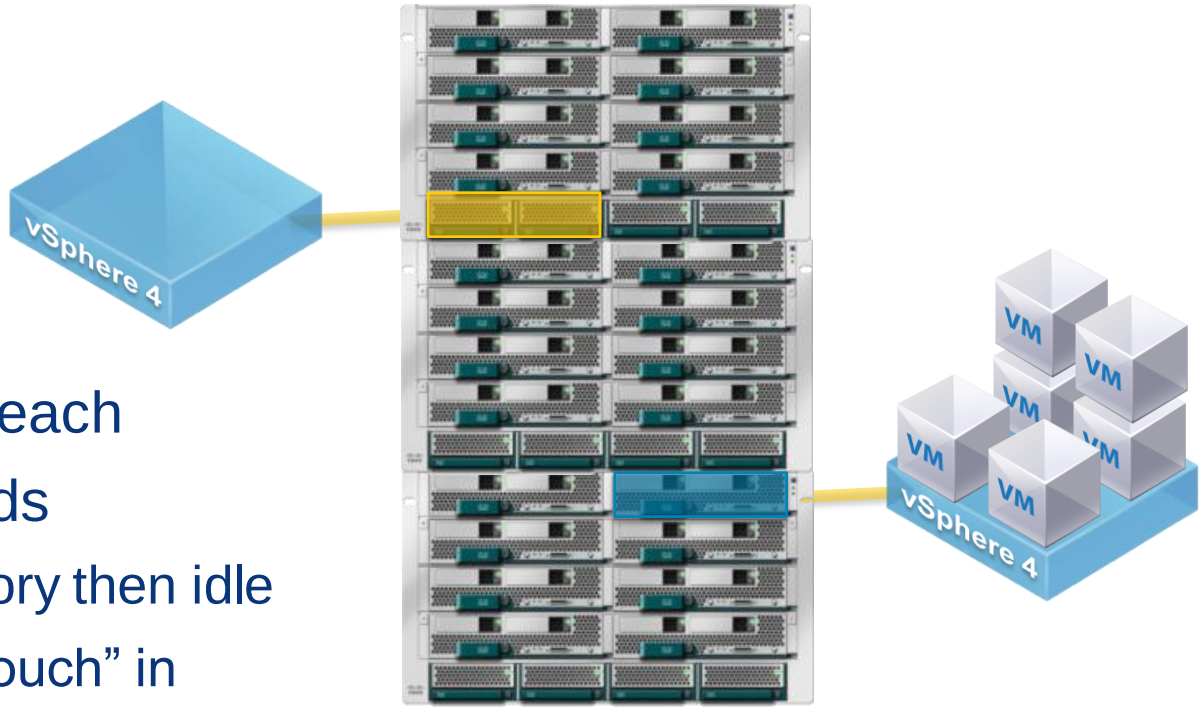

- 6 VMs, 8GB each
- Two workloads
  - Touch memory then idle
  - Run “mem touch” in background

Scaling VMotion, DRS and DPM to 10GE

Intra-Cluster Mobility

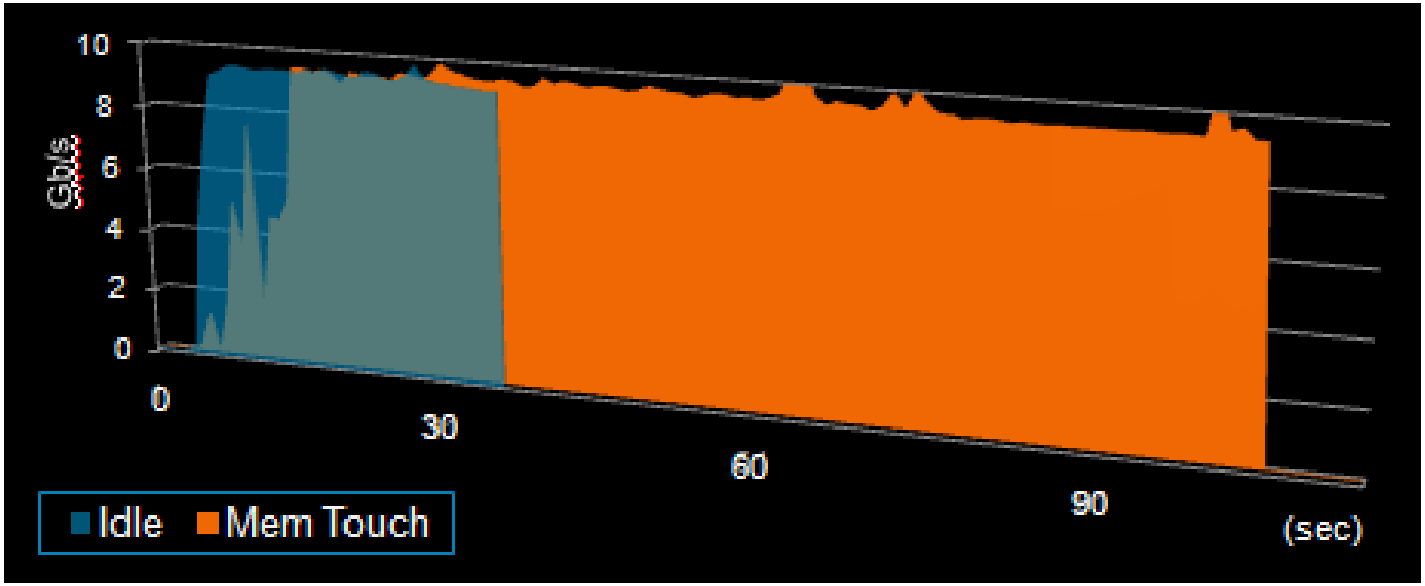

	Idle	Mem Touch
Total Completion Time	45 sec*)	103 sec**)
# of Pre-Copy Iterations	1	3
Total Data Transferred	41GB	112GB

Source: Cisco

*) 5 minutes, **) 15 minutes for 1G Link

53

UCS – Best Practise

UCS 6100 – Fabric Interconnect

•Fabric Interconnect Options

6120 – 20 10 Gbps ports fixed, 1 GEM
6140 – 40 10 Gbps ports fixed, 2 GEM's

•Ethernet ports can be "Server Ports" or "Uplink Ports"

Server Ports – FEX Discovery
Uplink Ports – Defined VLANs trunked

•Ethernet Switching Modes
End-Host Mode
Switch Mode

•Fiber Channel Switching Mode
NPV

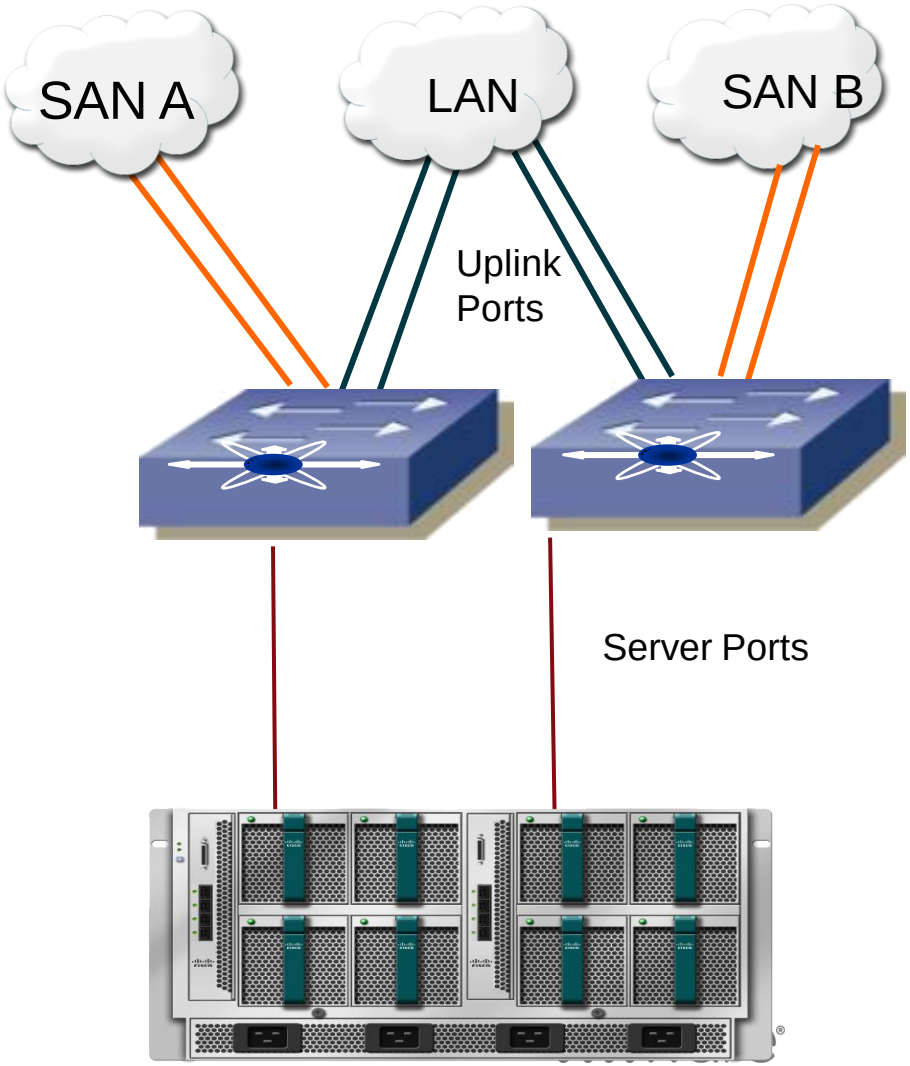

UCS 6100 Ethernet Switching Modes

End Host Mode (EHM)

- The external LAN sees the UCS 6100 as an end-host with multiple adapters. No Spanning Tree protocol on uplink ports.
- Active/Active use of uplinks by pinning.

Switch Mode

- The UCS 6100 acts like a traditional Ethernet switch with support for Spanning Tree protocol on the uplink ports.
- Links usage as per Spanning tree.

End Host Mode Operations

- Spanning Tree protocol is not run on both the uplink and the server ports.
- MAC learning only happens on the Server ports
- Traffic forwarding happens using the concept of pinning i.e a server port is mapped to an uplink port
- MAC aging does not happen for static MAC's.
- Active/Active use of links irrespective of the number of uplink switches – 3,4 etc.
- Highly scalable as Control plane is not occupied.
- All uplink ports should connect to the same L2 cloud.

Efficient use of uplink bandwidth

- Pingroups only applicable in EHM.
- Static pingroups used to define uplink port for a vNIC providing determistic path

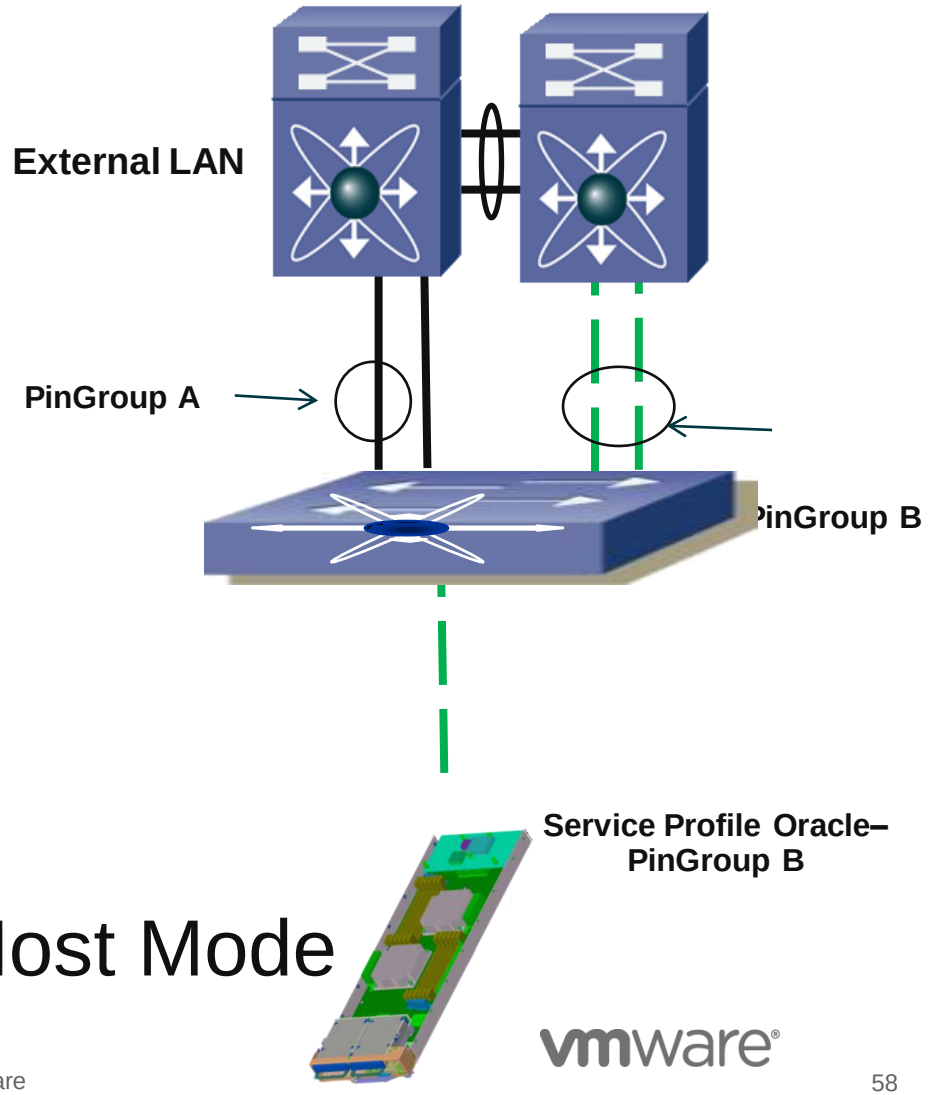

Recommendation: End Host Mode

End Host Mode Unicast Forwarding

- Server to server traffic on the same VLAN and Fabric (6100) is locally switched
- Uplink port to Uplink port traffic not switched
- Each server link is pinned to an uplink port / port-channel.
- Network to server unicast traffic is forwarded to server only if it arrives on pinned uplink port. This is termed as the Reverse Path Forwarding – (RPF) check.
- Packet with source MAC belonging to a server received on an uplink port is dropped (Deja-Vu Check)

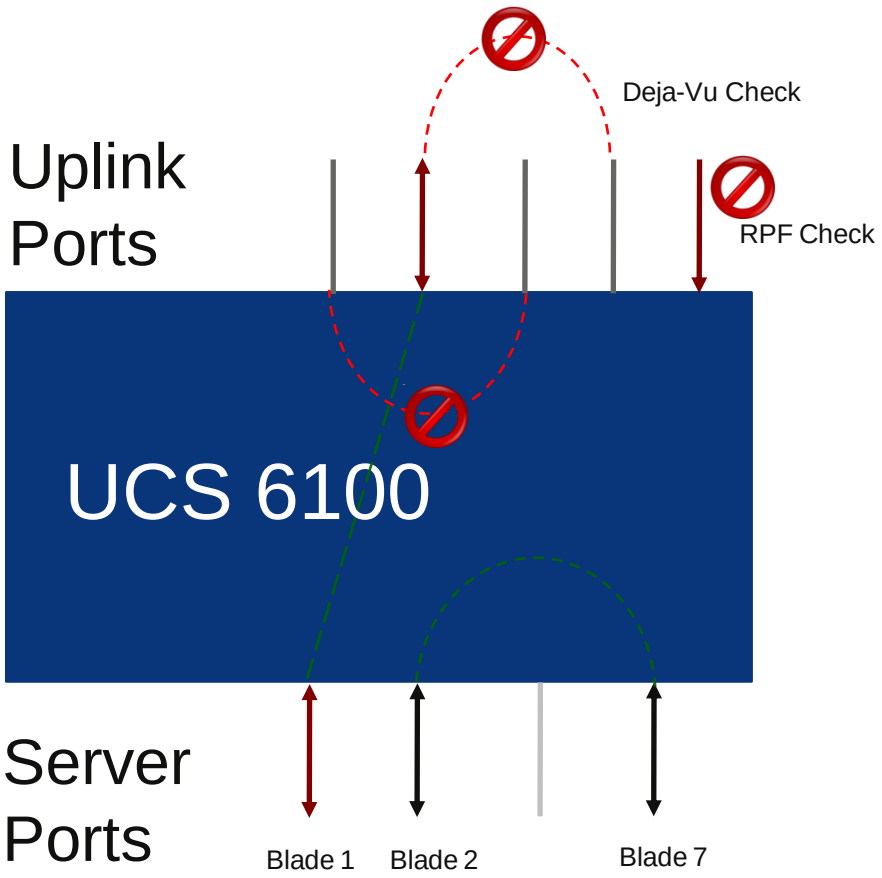

External LAN for Layer 2 Switching between Fabrics

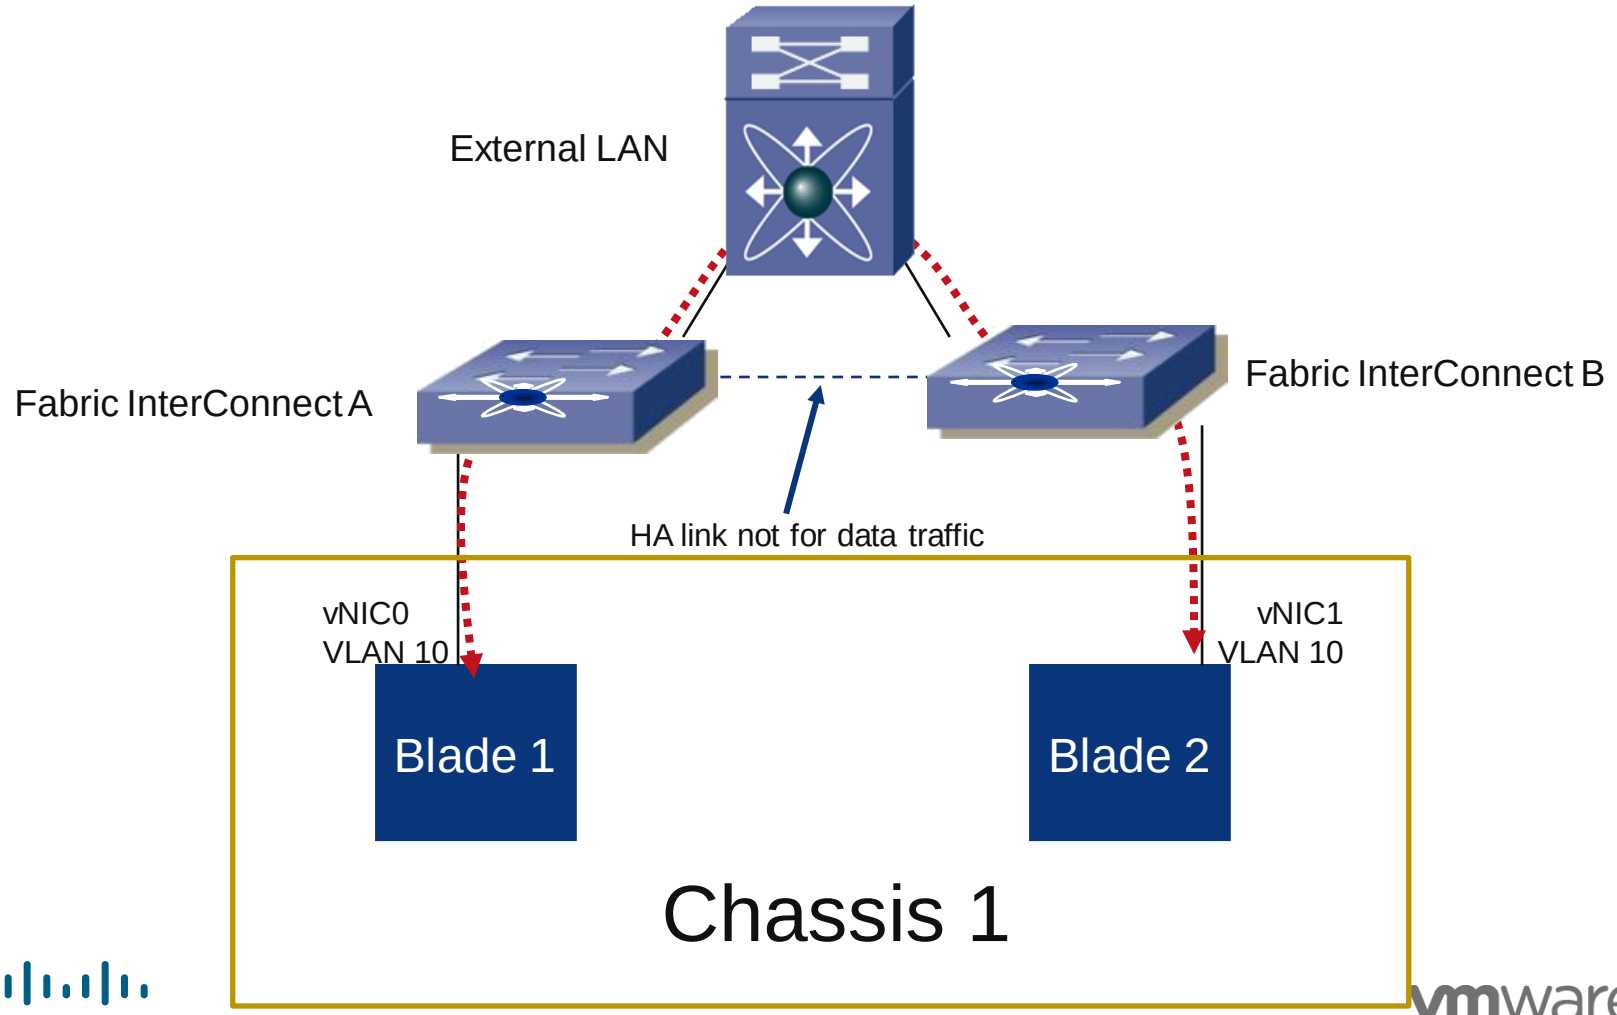

End Host Mode – Multicast Handling

- Broadcast traffic is pinned on exactly one uplink port i.e. it is dropped when received on other uplink ports.
- IGMP multicast groups are pinned based on IGMP snooping. Each group is pinned to exactly one uplink port.
- Server to server multicast traffic is locally switched.
- RPF and Deja-vu check also applies for multicast traffic.

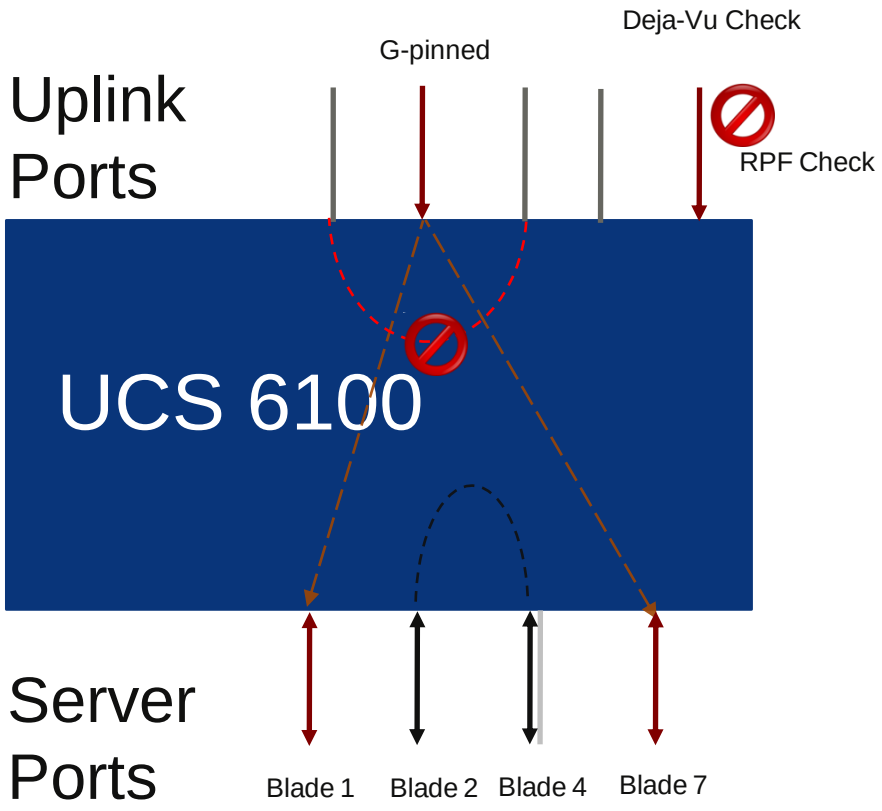

Active/Active use of Uplinks for EHM

End Host Mode

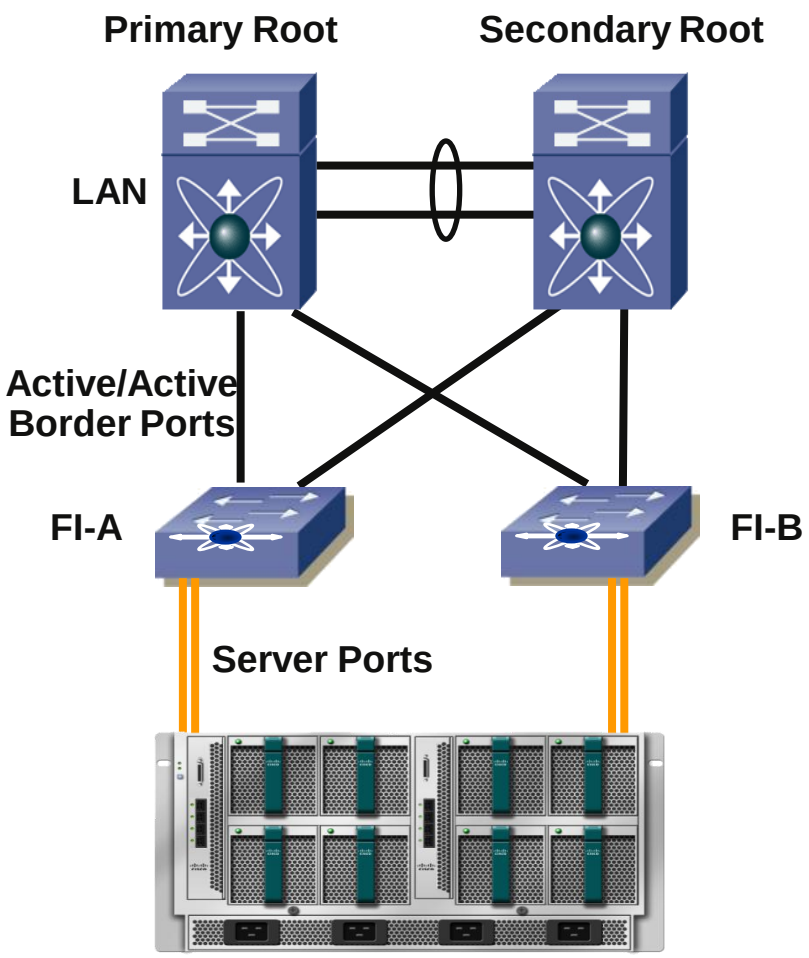

Switch Mode

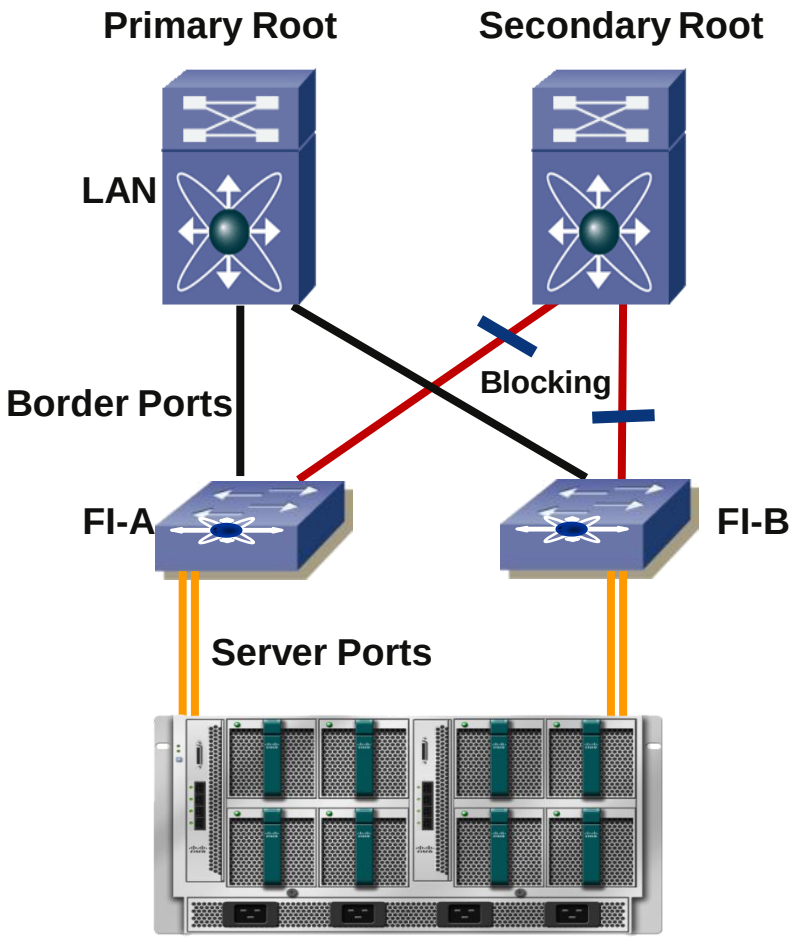

Recommendation: End Host Mode

Disjoint L2 Upstream

- EHM built on the premise that the L2 upstream is NOT disjoint.
- Incoming broadcast/multicast received only on 1 uplink for ALL VLANs

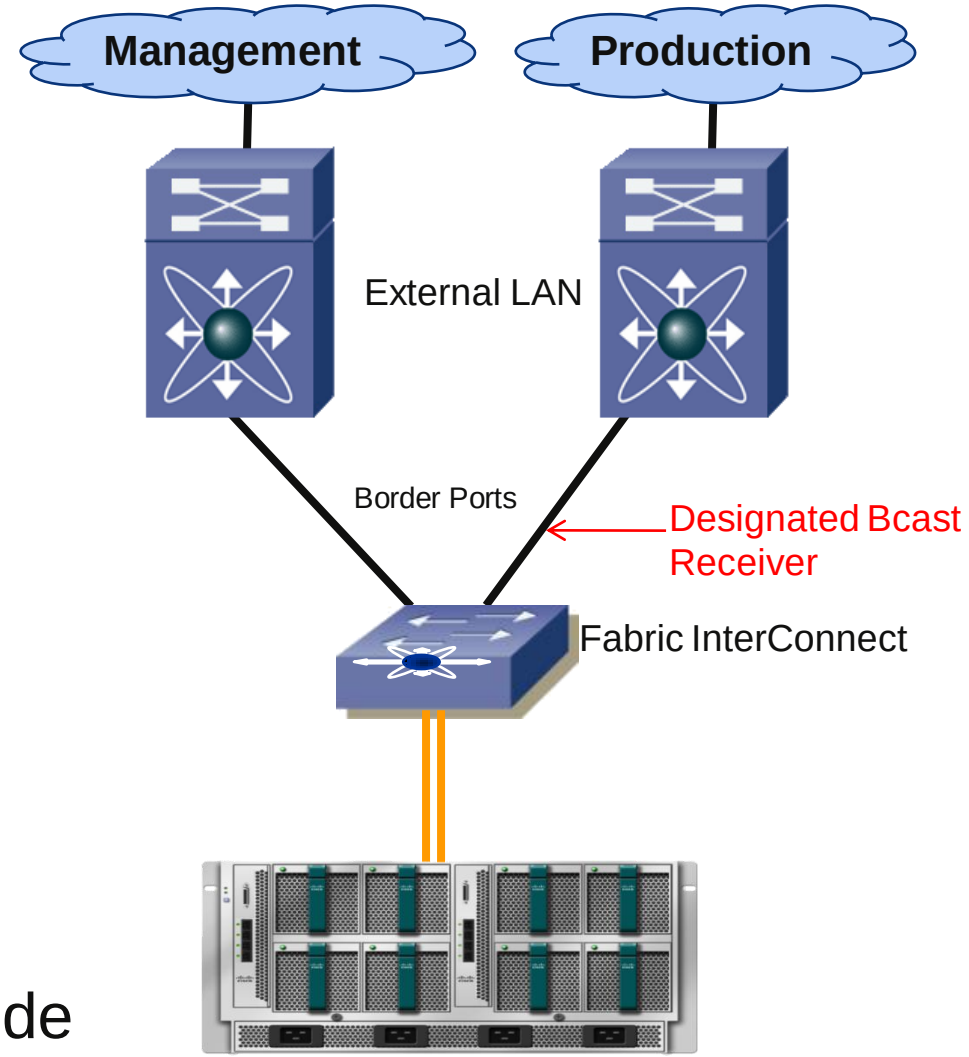

Recommendation: Switch Mode

Directly Connecting 3rd party devices like NAS

- Support for directly connecting NAS device will only work in Switch mode.
- Dependent on certifications

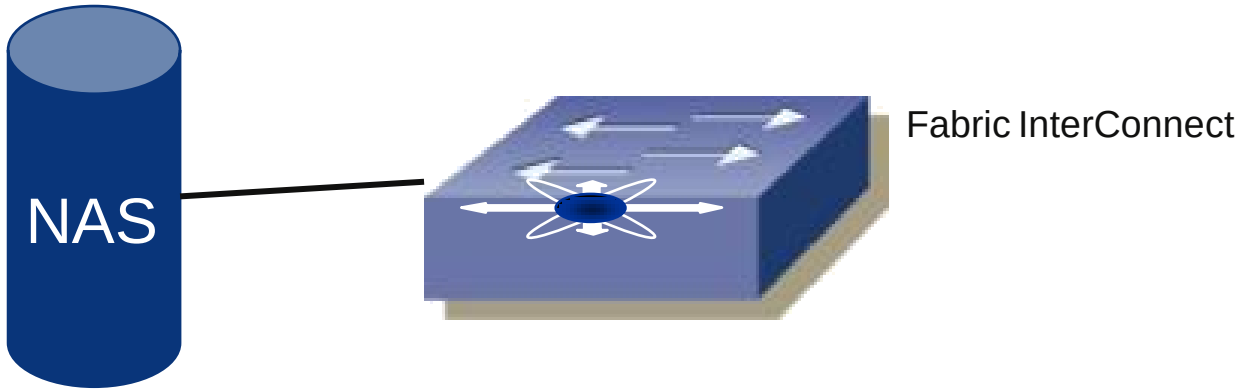

Recommendation: Switch Mode

N-Port Virtualization (NPV) mode

- UCS FI always work in NPV mode
  - Server-facing ports are regular F ports
  - Uplinks toward SAN core fabric are NP ports
- UCS distributes (relays) FCIDs to attached devices
  - No domain ID to maintain locally
- One VSAN per uplink on UCS Fabrics
  - No trunking or channelling of NP ports
- Zoning, FSPF, DPVM, etc are not configured on the UCS Fabrics
- Domain mgr, FSPF, zone server, fabric login server, name server
  - They do not run on UCS Fabrics
- No local switching
  - All traffic routed via the core SAN switches

N-Port Virtualization (NPV): An Overview

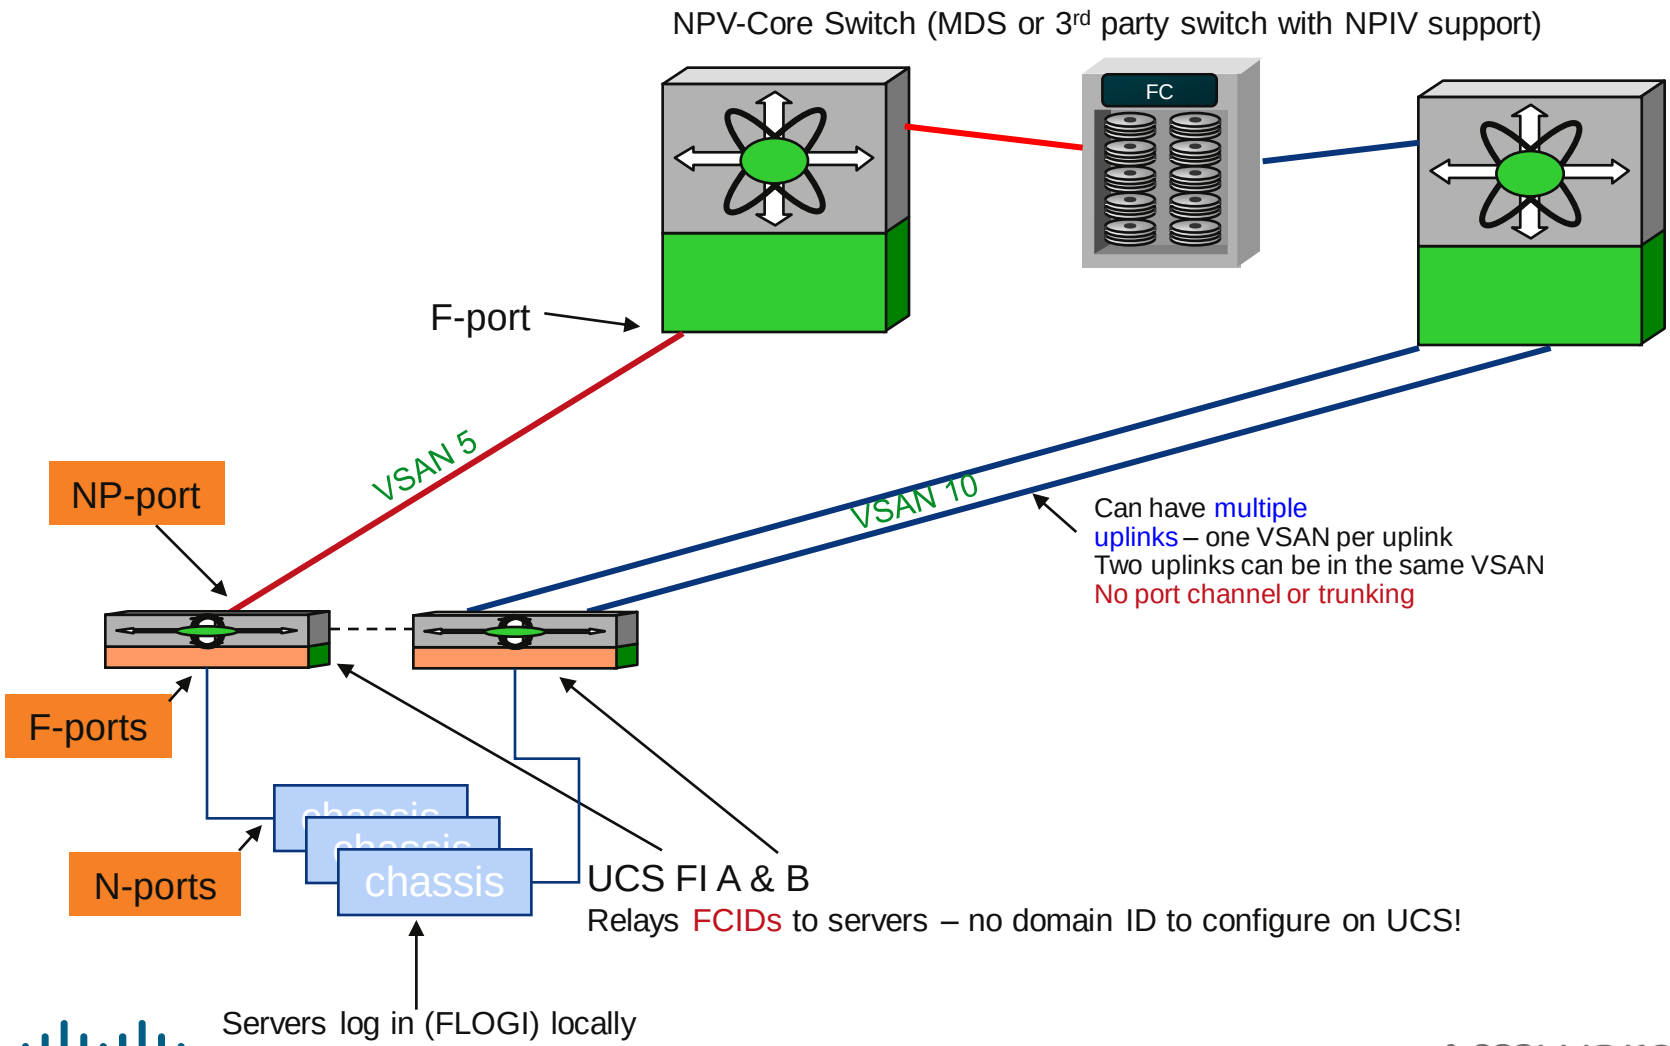

With UCS in End-Host Mode – design 1

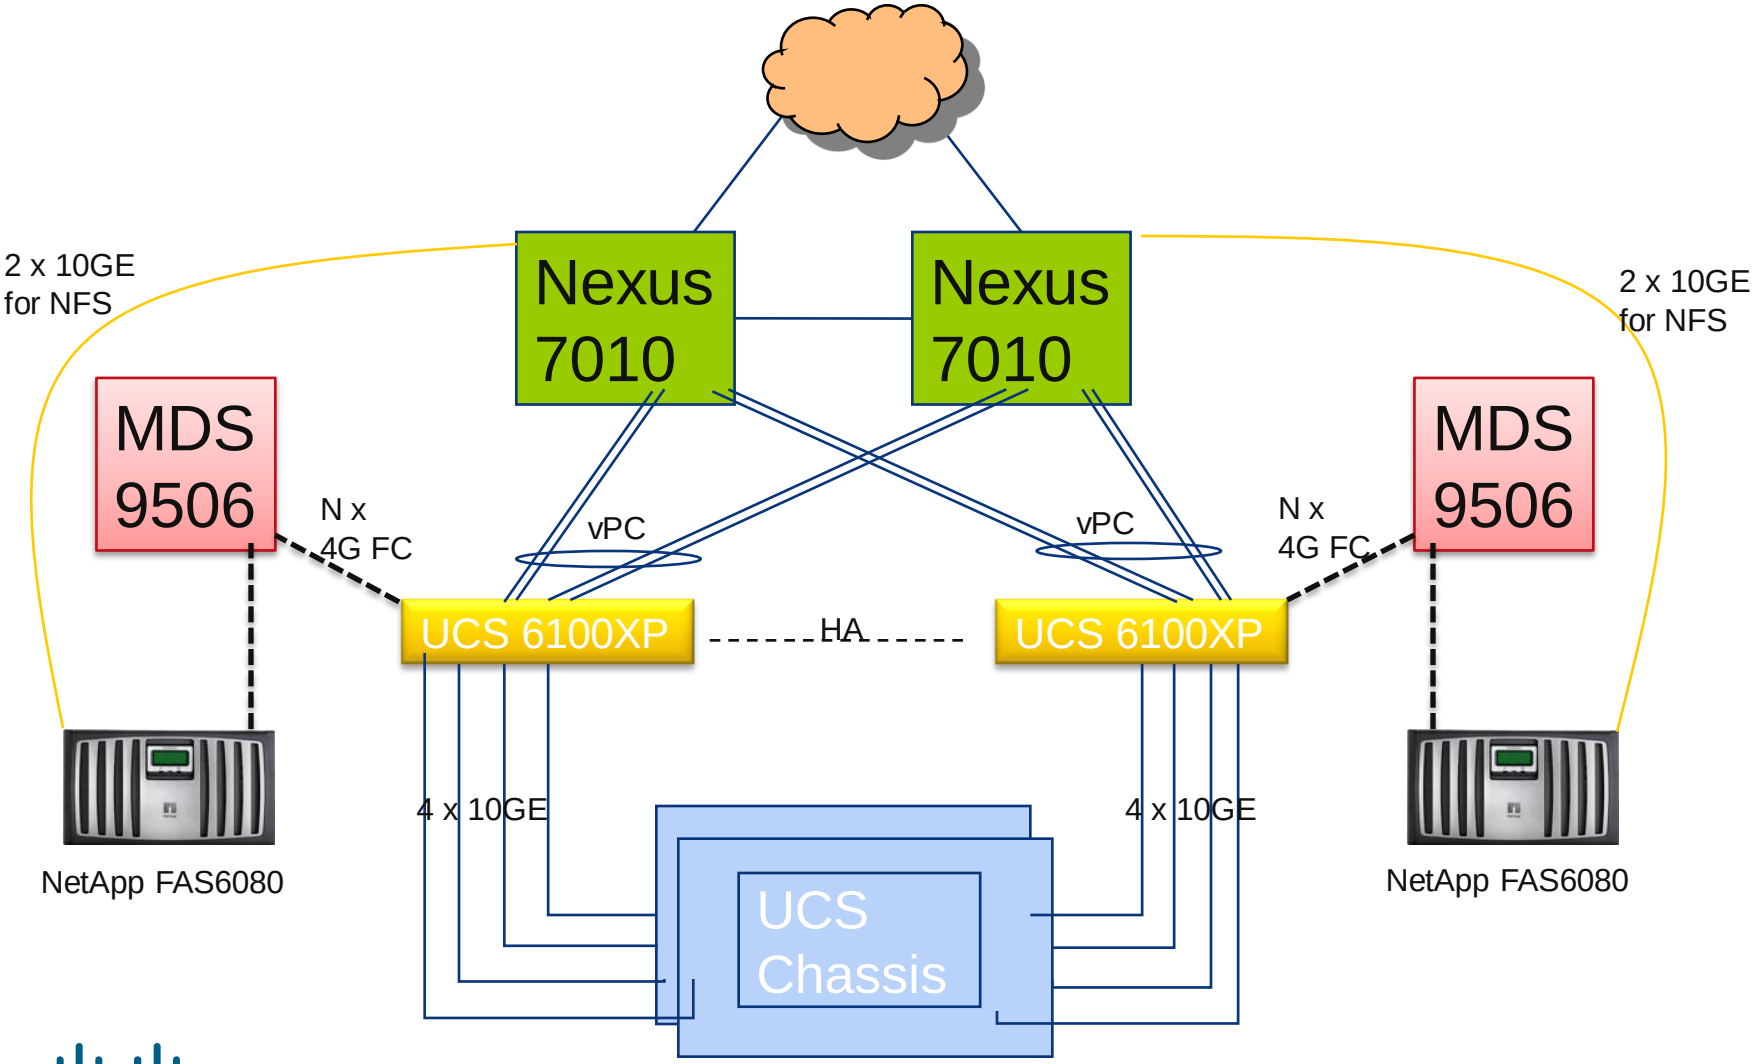

Cisco Unified Computing System

Unified Fabric

Fabric Extender
Virtualized Adapter

Scale Out

Extended Memory

Virtualization Optimization

- Fine-grained control, portability, and visibility of network, compute, and storage attributes
- More than double the memory capacity of competing systems

Industry Standard Servers

- Intel Xeon processor 5600 and 7500 series
- 150% generational performance increase
- Intelligent platform for performance and energy efficiency

Unified Fabric

- Wire once, low latency FC and Ethernet
- Virtualization aware
- Less than half the normal amount of adapters, switches, cables

Automated Provisioning

- Embedded single point of management and provisioning
- Visibility and control across datacenter organizations
- Infrastructure policy management and compliance


CISCO

vSphere 4.0 Overview

Introducing VMware vSphere™

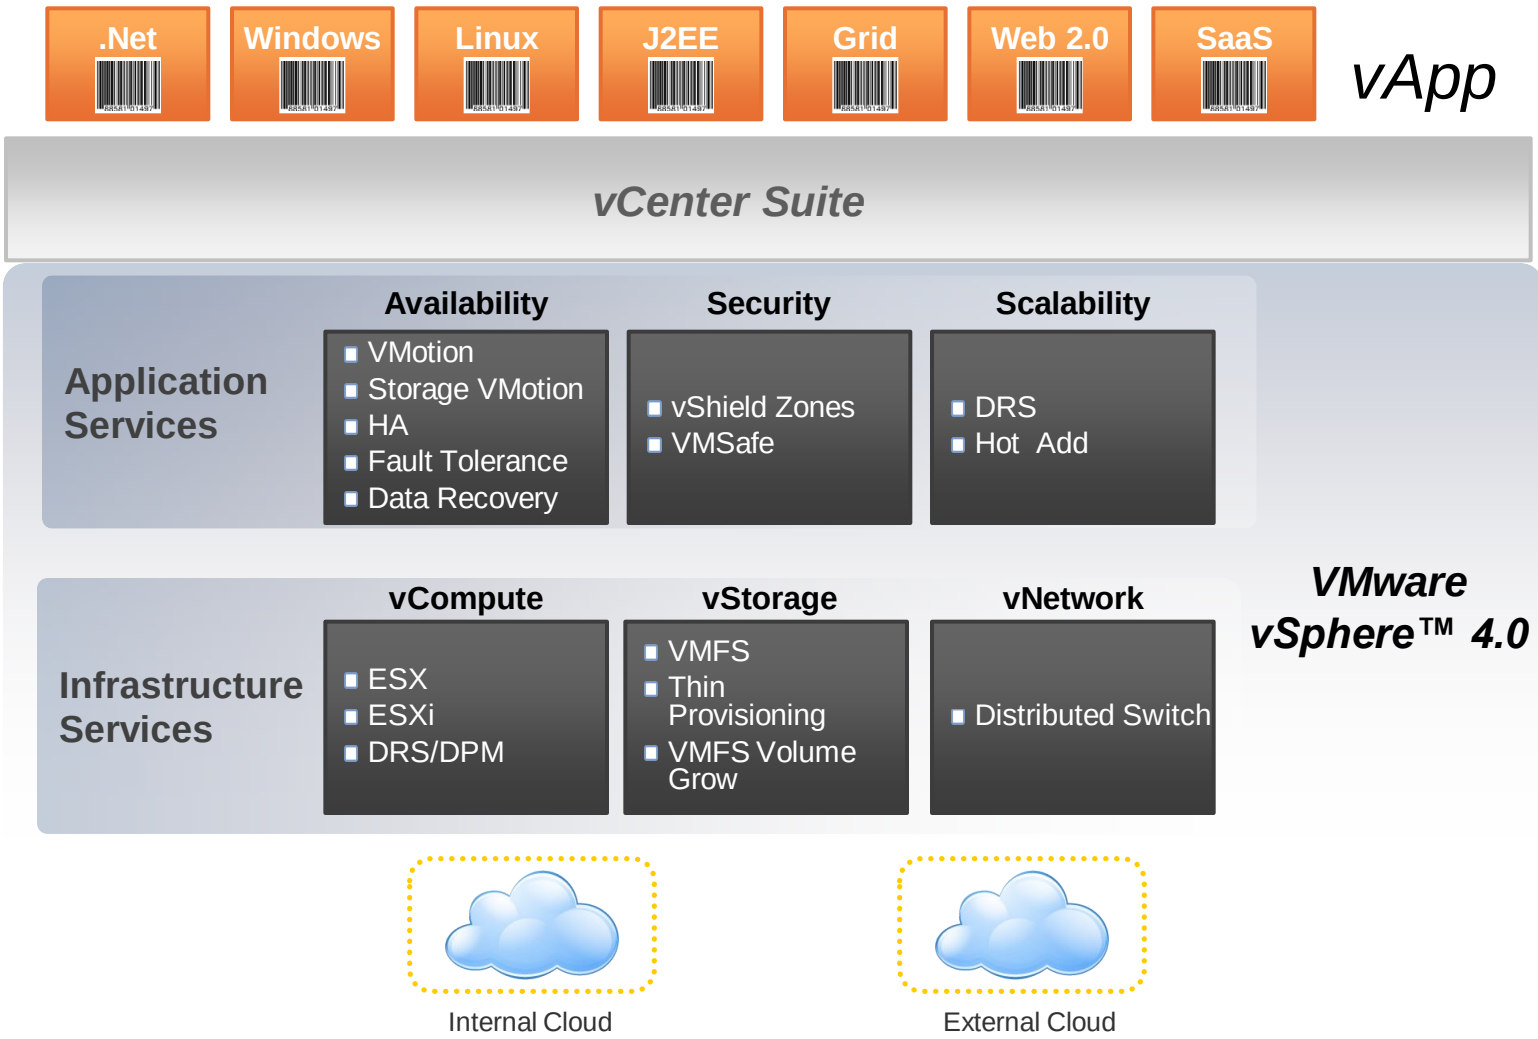

*Note vCenter Server and its components are a separate purchase

Optimization for the Highest Consolidation Ratios

vCompute | vStorage | vNetwork

Virtual Machines

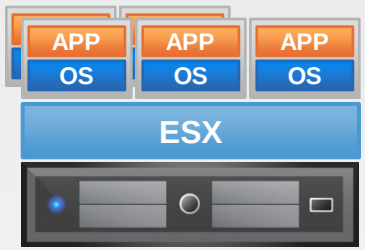

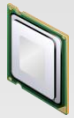

CPU

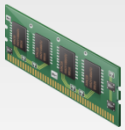

Memory

Networking

Storage

VM Scale Up	8-way vSMP and 255 GB of RAM per VM
Hardware Scale Up	64 cores and 1 TB of physical RAM
Hardware Assist Purpose Built Scheduler	Lowest CPU overhead
Hardware Assist Page Sharing Ballooning	Maximum memory efficiency
VMXNET3 VMDirectPath I/O	Wirespeed network access
Storage stack optimization VMDirectPath I/O	Greater than 350k iops per second Lower than 20 microsecond latency

Current NEW

VMware DPM Expanded Support

vCompute | vStorage | vNetwork

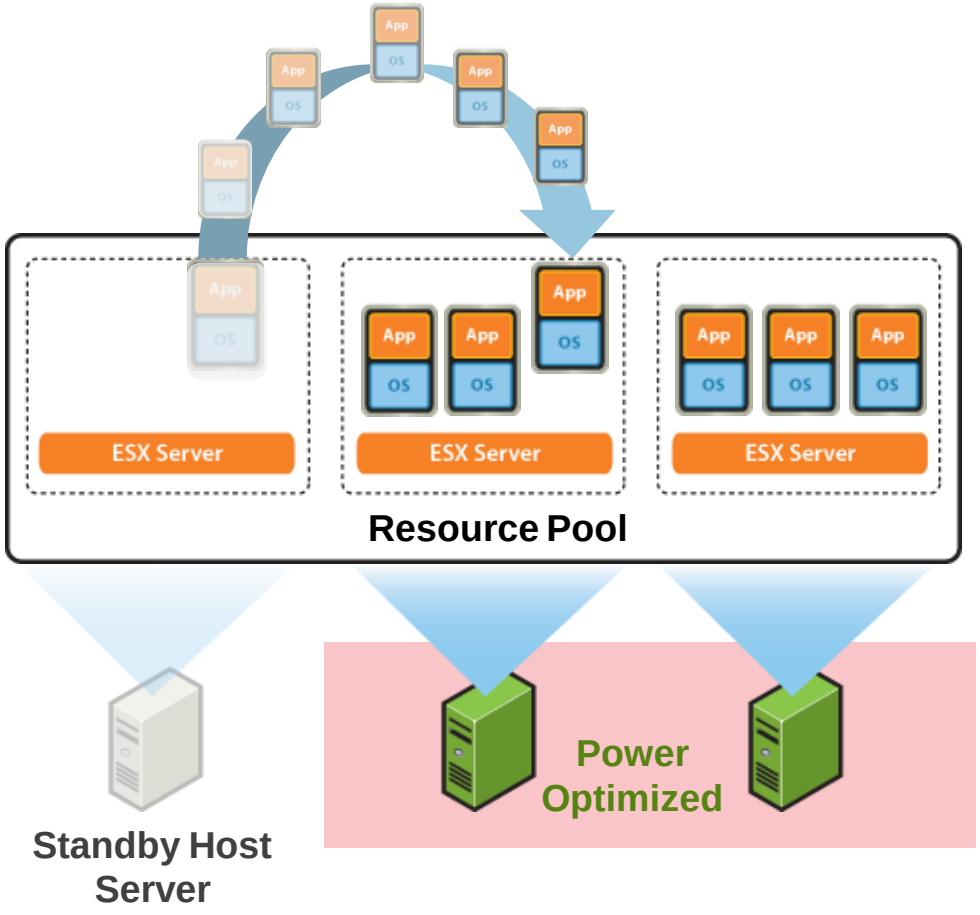

DPM consolidates workloads to reduce power consumption

- Cuts power and cooling costs
- Automates management of energy efficiency

Supports three wake protocols:

- Intelligent platform management interface (IPMI)
- Integrated Lights-Out (iLO)
- Wake-On-LAN (WOL)

Configure and test wake on every host in cluster

vStorage Thin Provisioning

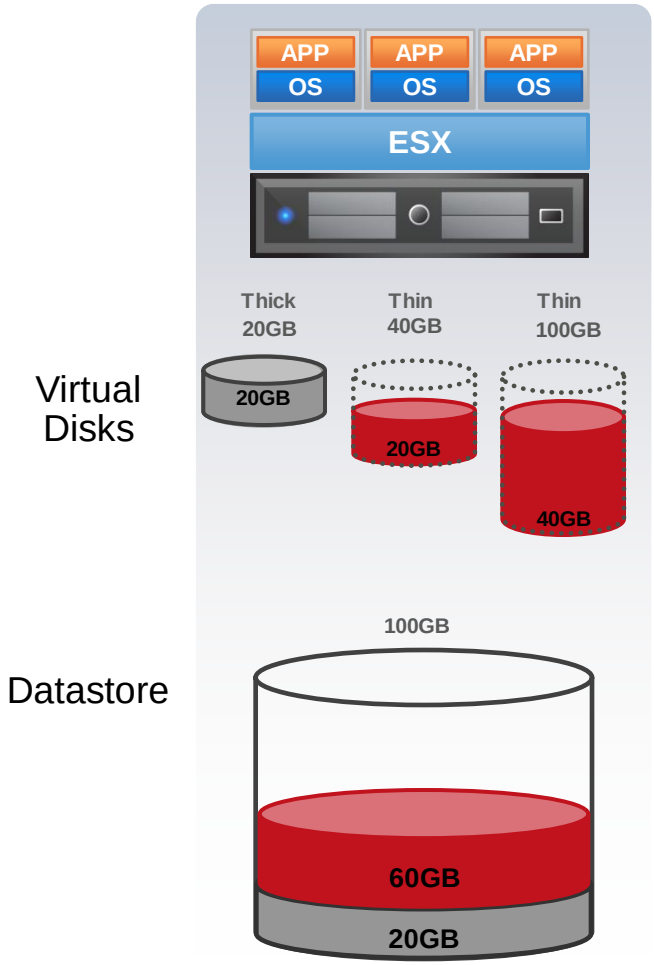

- Virtual machine disks consume only the amount of physical space in use
  - Virtual machine sees full logical disk size at all times
  - Full reporting and alerting on allocation and consumption
- Significantly improve storage utilization
- Eliminate need to over-provision virtual disks
- Reduce storage costs by up to 50%

Thin Disk Provisioning Operations

A thin-disk option is available when you:

- Create a virtual machine
- Clone to a template
- Clone a virtual machine
- Migrate virtual machine storage (Storage VMotion)

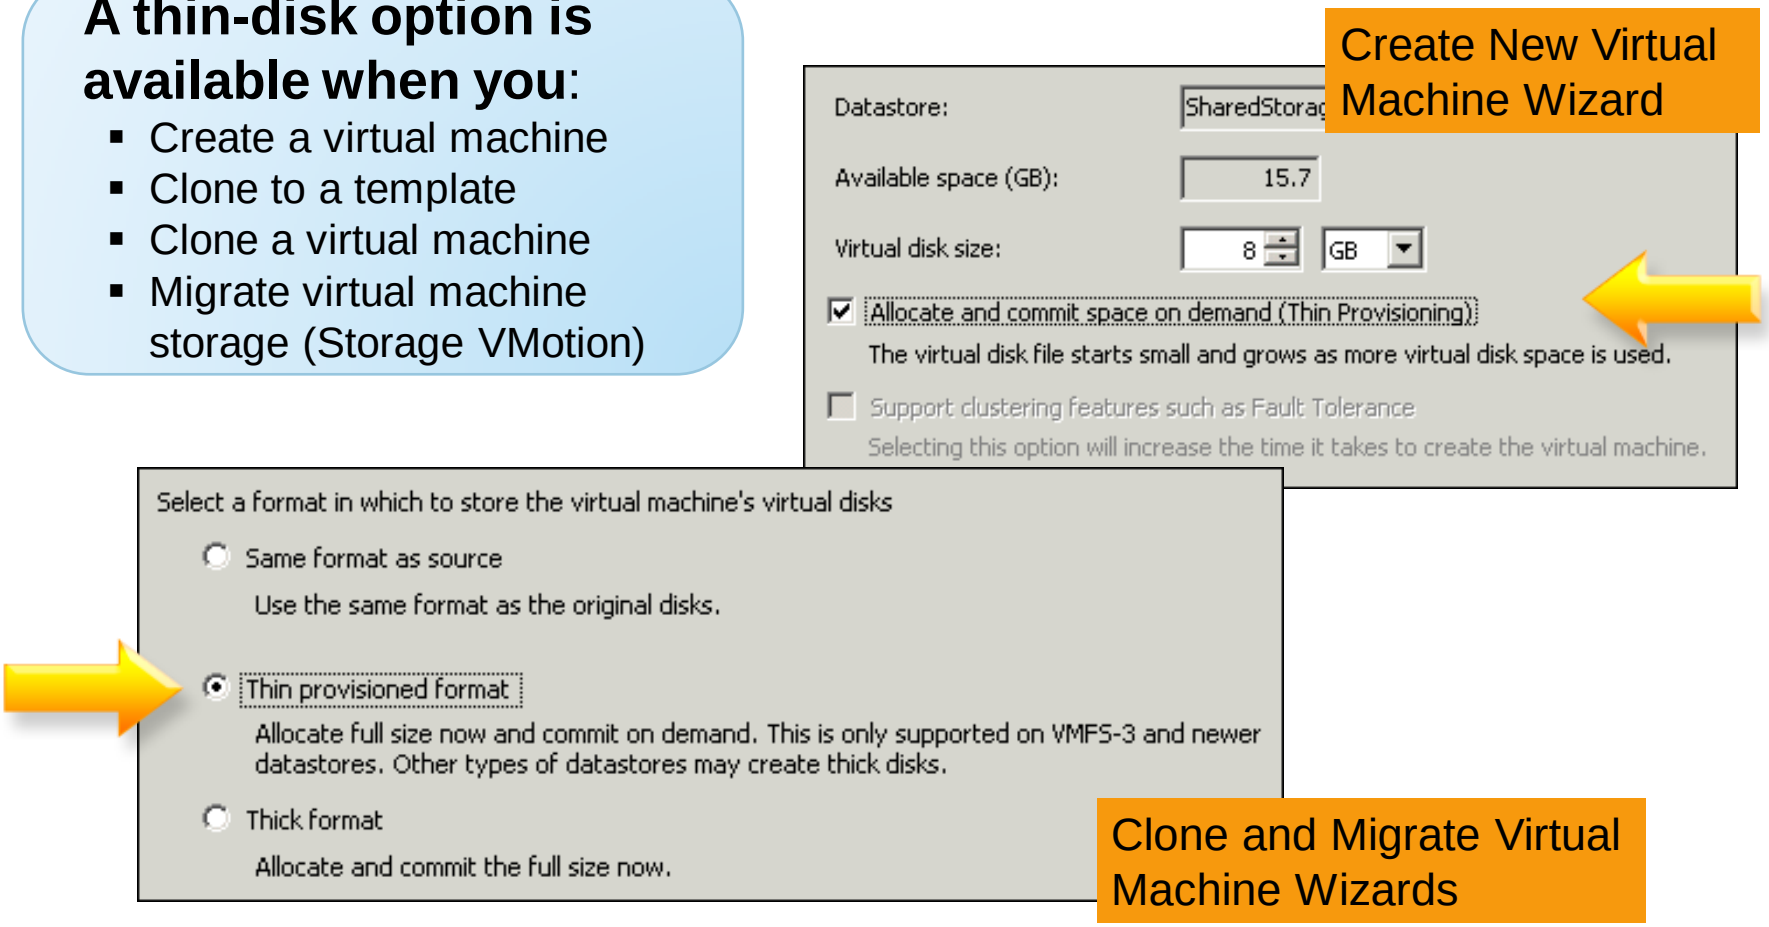

VMware Fault Tolerance

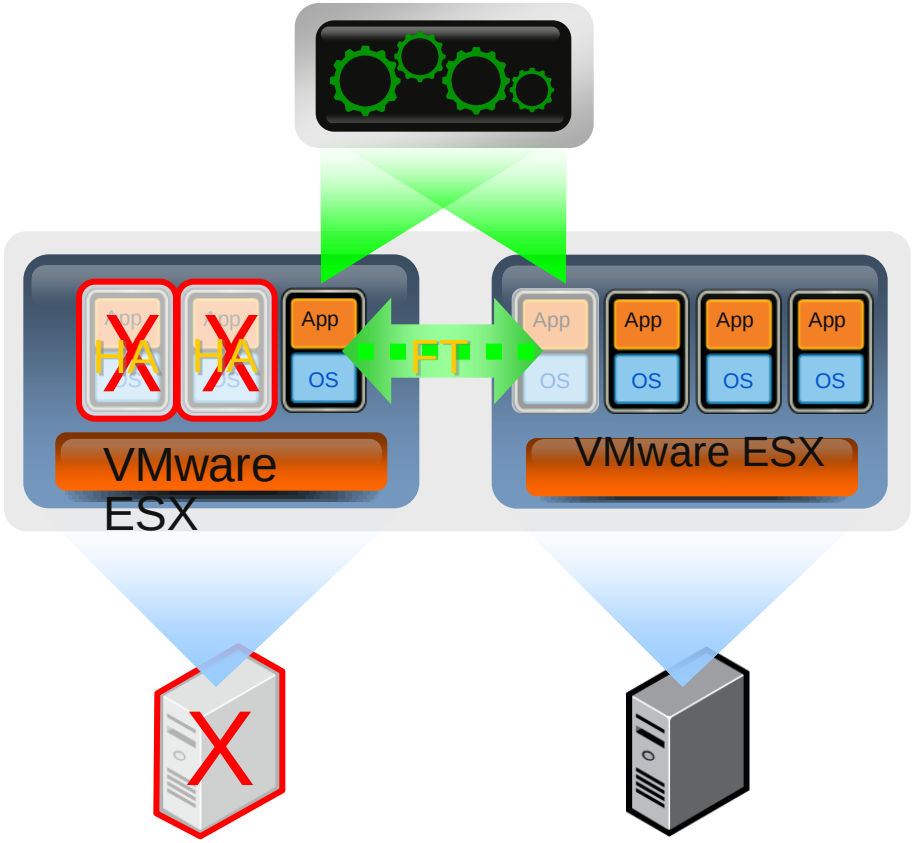

- Single identical VMs running in lockstep on separate hosts
- Zero downtime, zero data loss failover for all virtual machines in case of hardware failures
- Integrated with VMware HA/DRS

- Zero downtime, zero data loss
- No complex clustering or specialized hardware required
- Single common mechanism for all applications and OS-es

Enable Fault Tolerance with a Single Click

Availability | Security | Scalability

Primary Virtual Machine > Summary Tab

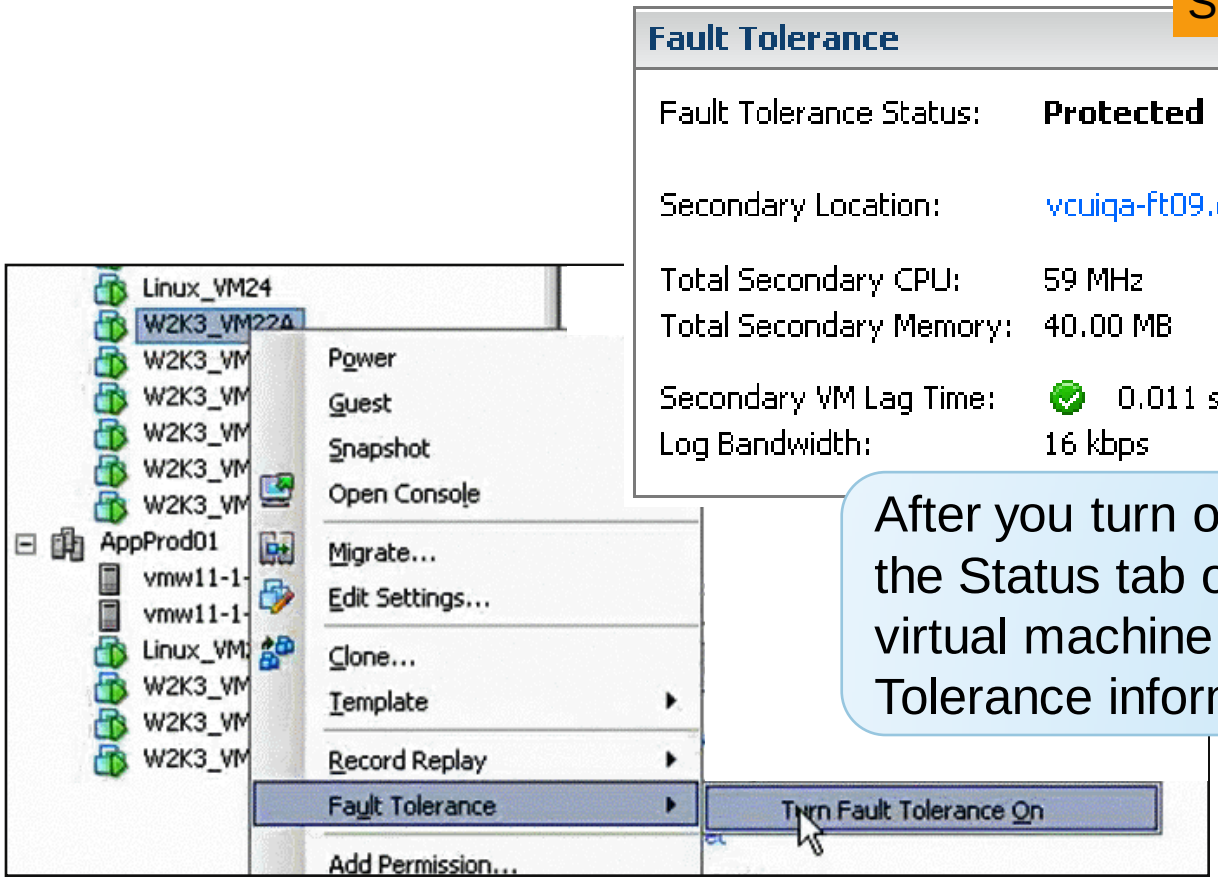

After you turn on Fault Tolerance, the Status tab on the primary virtual machine shows Fault Tolerance information.

vStorage APIs for Data Protection

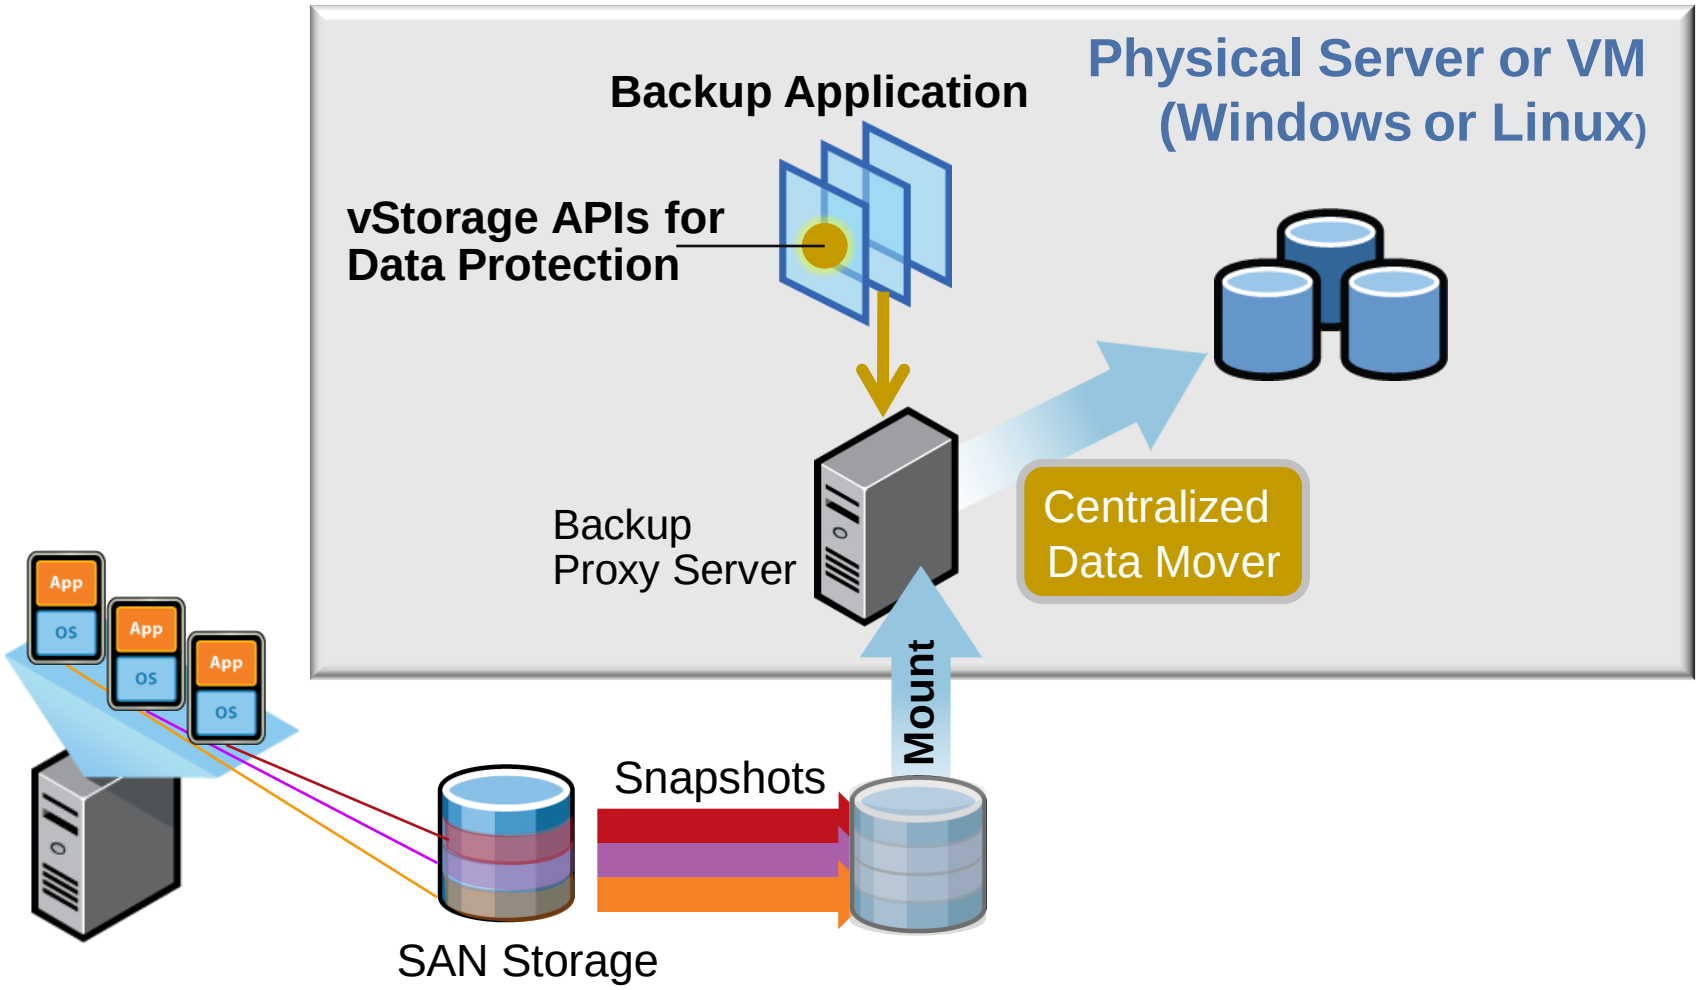

vCenter Data Recovery

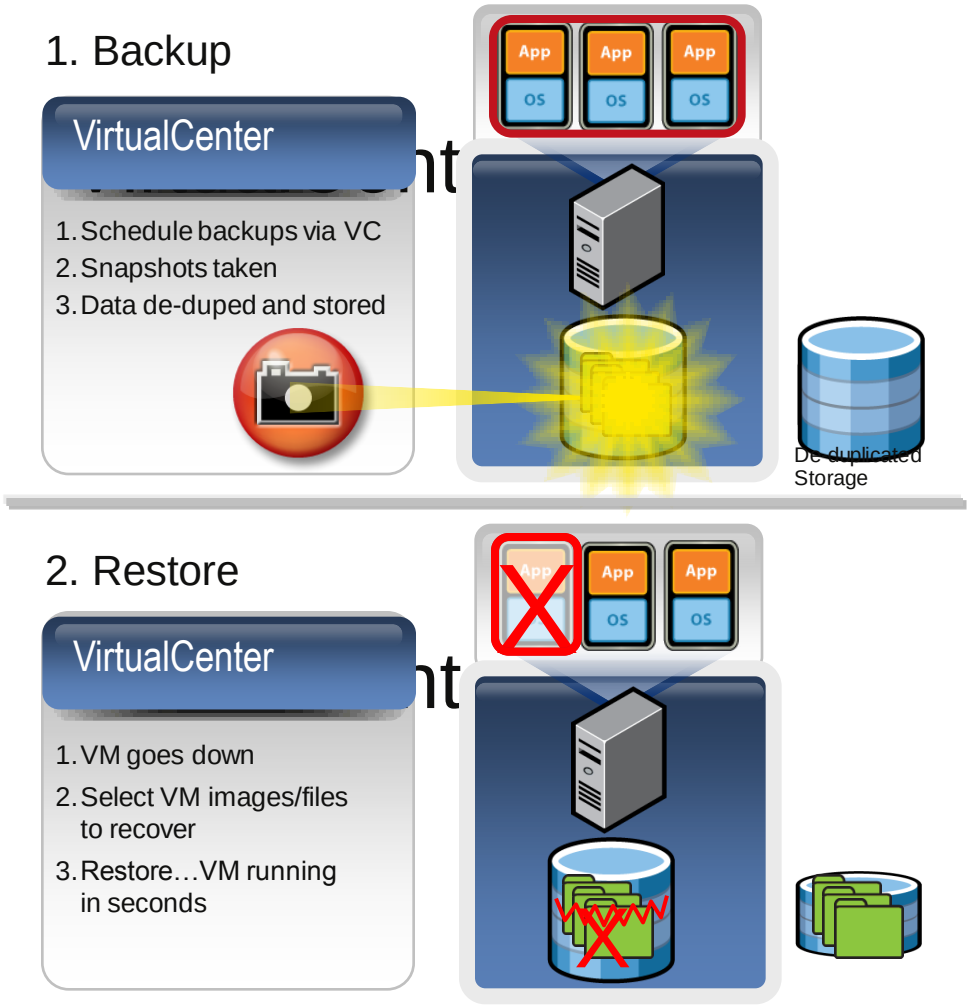

- Agent-less, disk-based backup and recovery of your VMs
- VM or file level restore
- Incremental backups and data de-dupe to save disk space

- Quick, simple and complete data protection for your VMs
- Centralized Management through VirtualCenter
- Cost Effective Storage Management

VMware Data Recovery

Availability | Security | Scalability

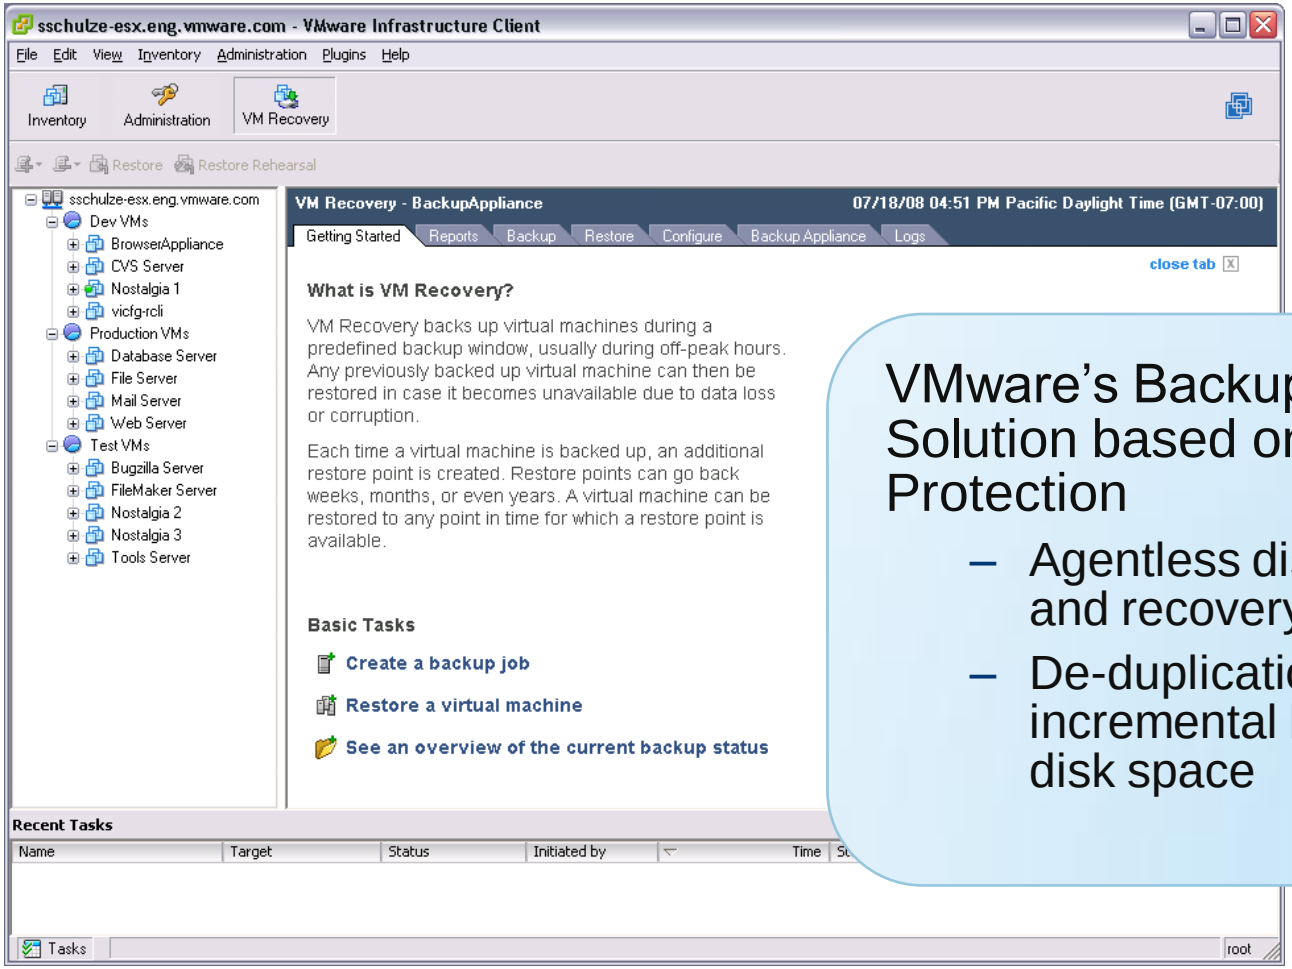

VMware's Backup/Recovery Solution based on APIs for Data Protection

- Agentless disk-based backup and recovery
- De-duplication and incremental backups to save disk space

vNetwork Distributed Switch

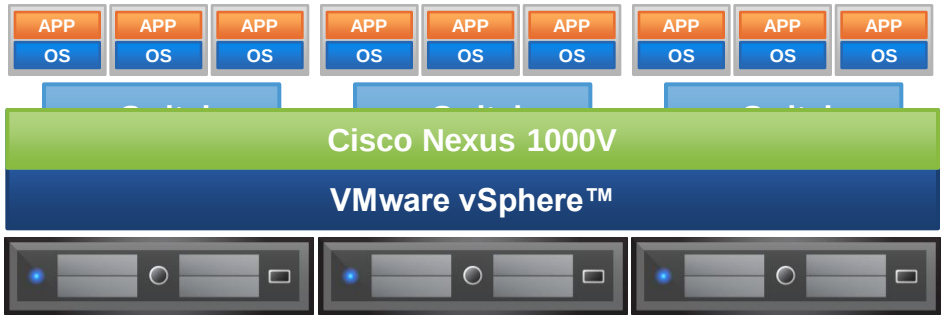

- Aggregated datacenter level virtual networking
- Simplified setup and change
- Easy troubleshooting, monitoring and debugging
- Enables transparent third party management of virtual environments

vApp – New Model for Describing and Deploying Applications

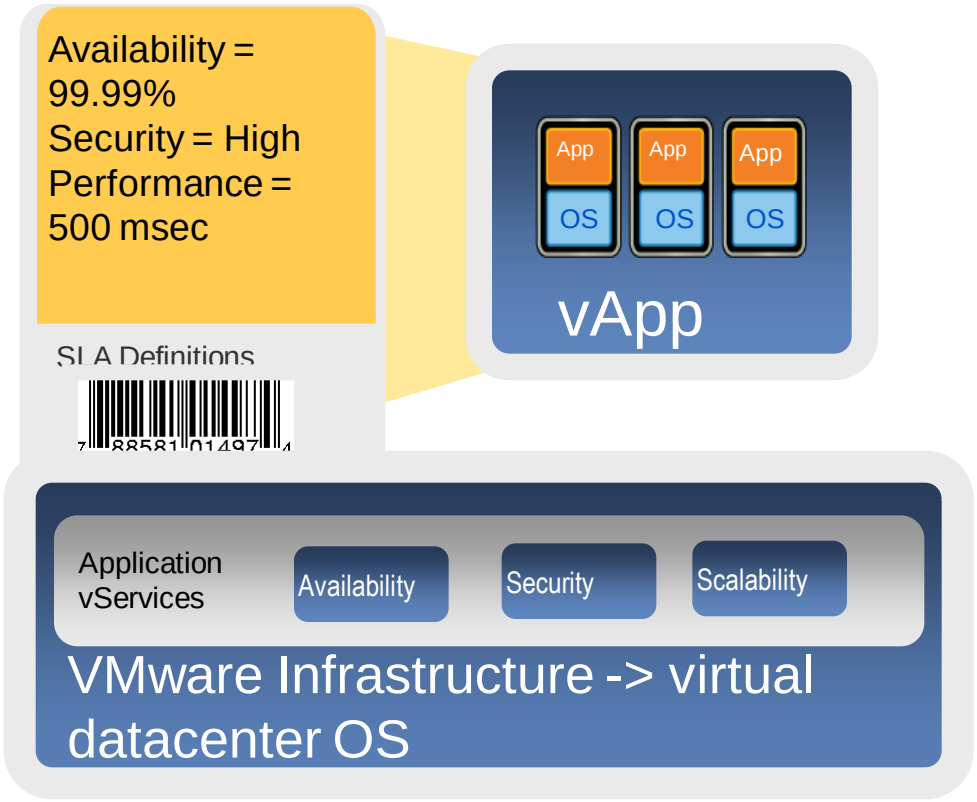

- Allows management of multi-tier applications as a single entity
- Utilizes industry standard OVF to provide instructions on how to deploy
- Templates, Clone and other operations execute at the vService level

- Simpler, application centric view of management
- Easier portability of applications
- Applications can now be written to monitor and scale themselves

VMware Solutions Maximize Uptime

	Prevent Planned Downtime	Minimize Unplanned Downtime
Site 	Site Recovery Manager	
Storage 	Storage vMotion	Consolidated Backup + backup software, Data Recovery
Server 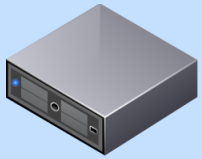	vMotion + DRS Maintenance Mode	HA, Fault Tolerance
Interconnect 	Network Redundancy	NIC & HBA Teaming

Enhanced VMotion Compatibility (EVC)

Availability | Security | Scalability

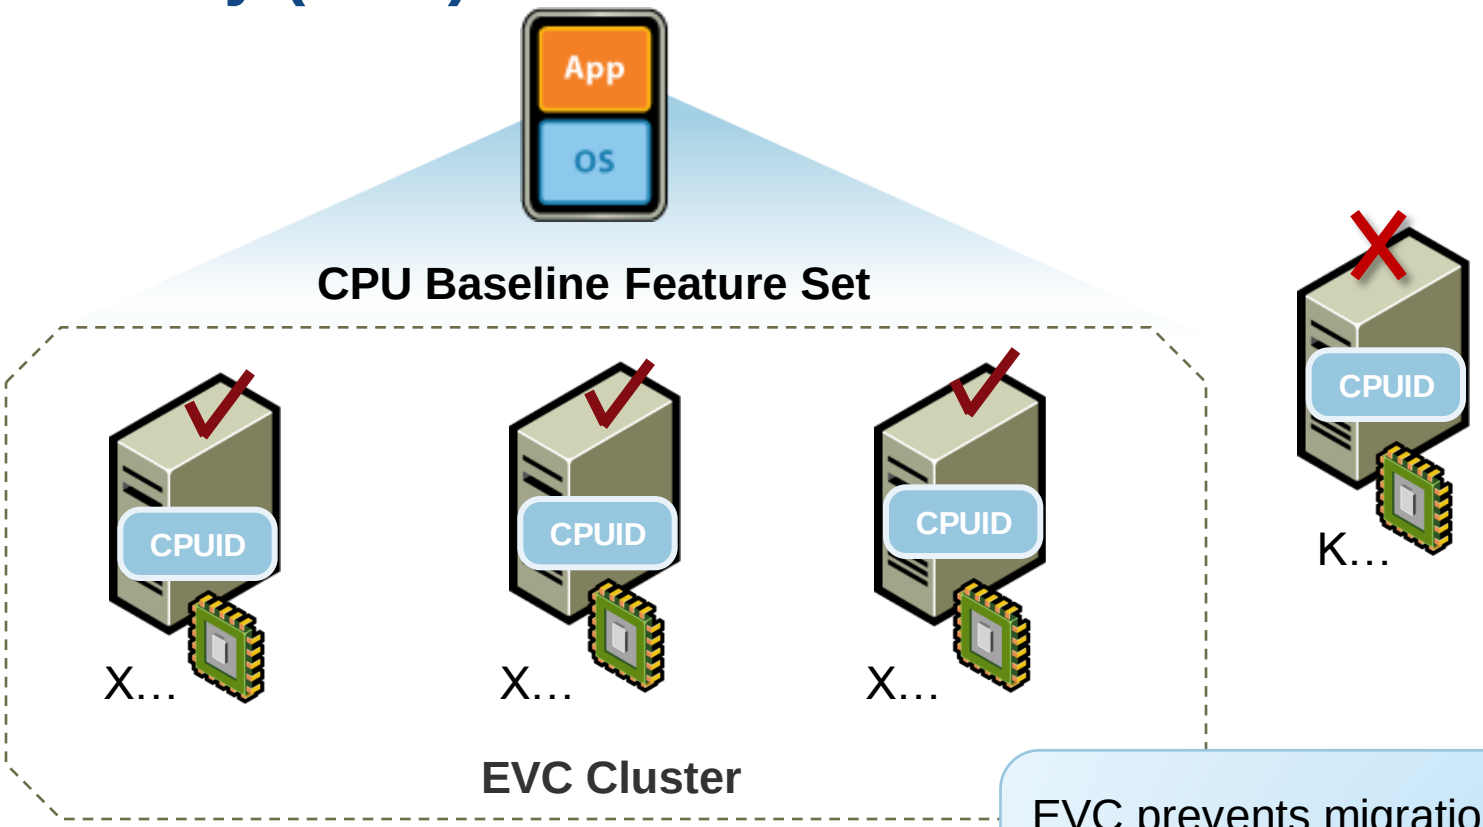

EVC prevents migrations with VMotion from failing due to incompatible CPUs.

Storage VMotion in vSphere 4

Availability | Security | Scalability

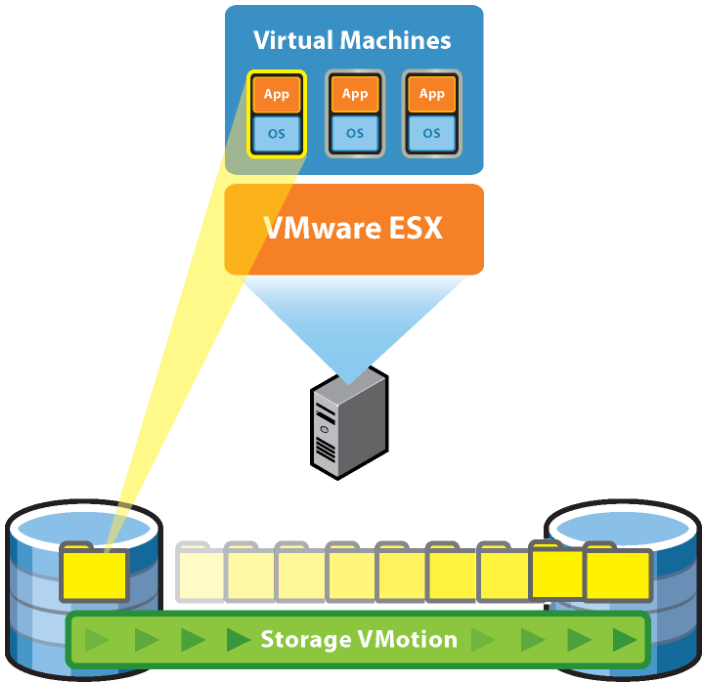

Enhancements

- Can administer via vSphere Client
- Supports NFS, Fibre Channel, and iSCSI
- No longer requires 2 x memory
- Supports moving VMDKs from thick to thin formats
- Can migrate RDMs to RDMs and RDMs to VMDKs (non-passthrough)
- Leverages new vSphere 4 features to speed migration

Limitations

- Virtual machine cannot include snapshots
- VM must be powered off to simultaneously migrate both host and datastore

vStorage Technologies and Interfaces

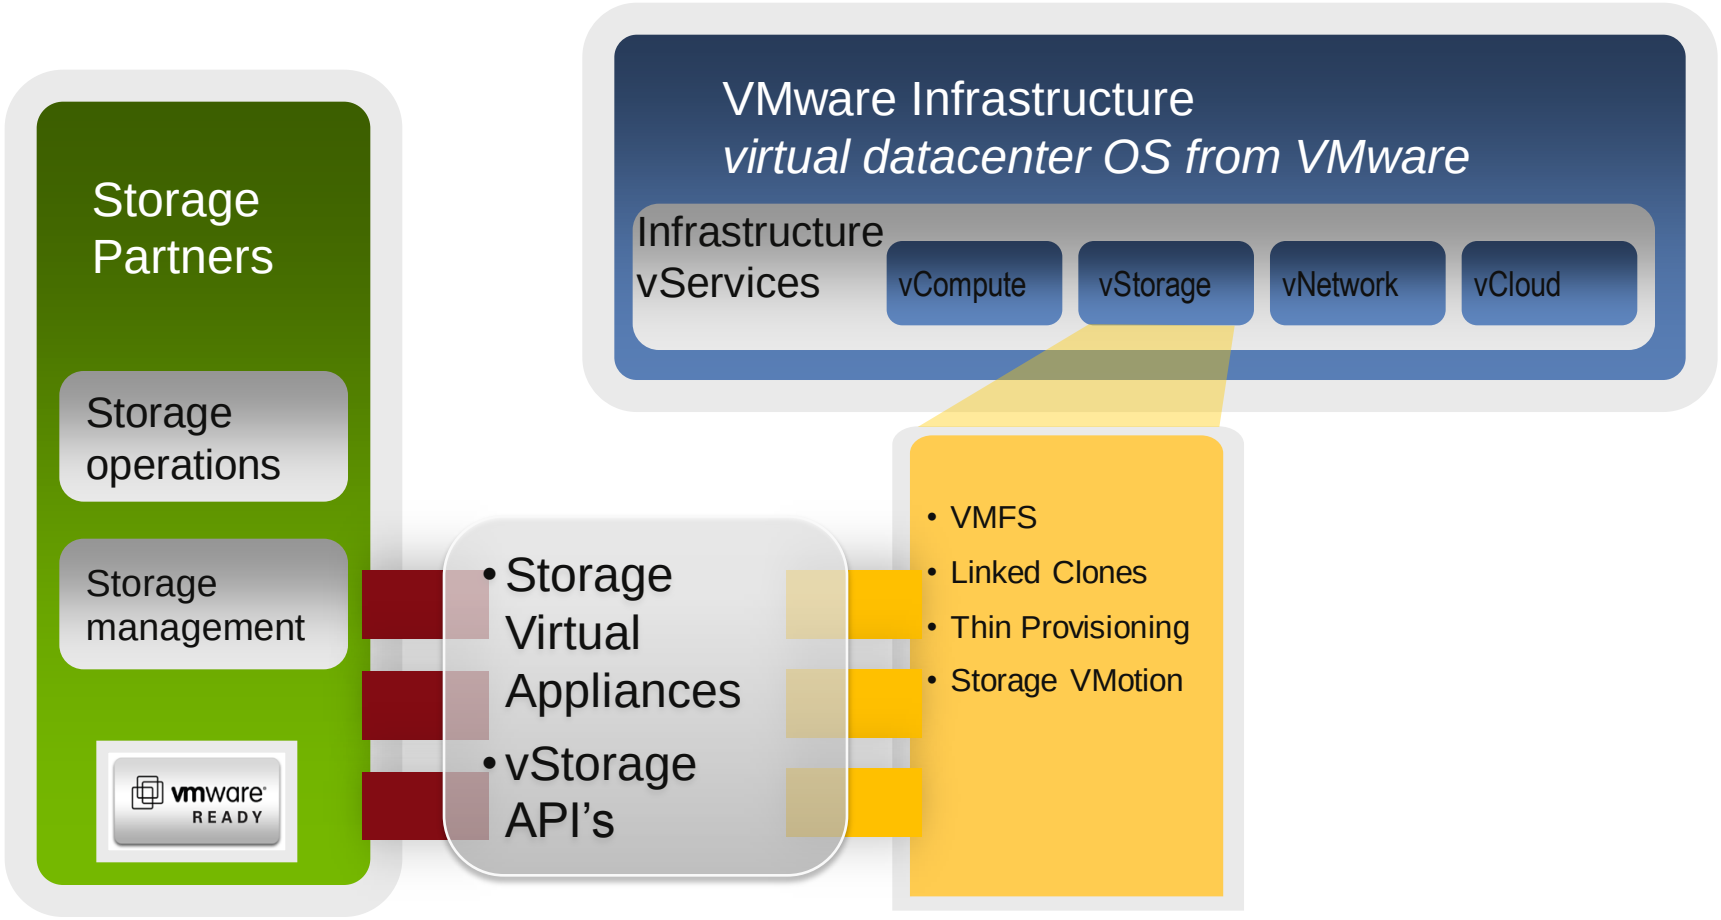

VMware VMsafe

Availability | **Security** | Scalability

- API that enables protection of VMs by inspection of virtual components in conjunction with hypervisor
- Isolation of protection engine from malware
- Broad ranging coverage of virtual machine CPU, memory, storage and network

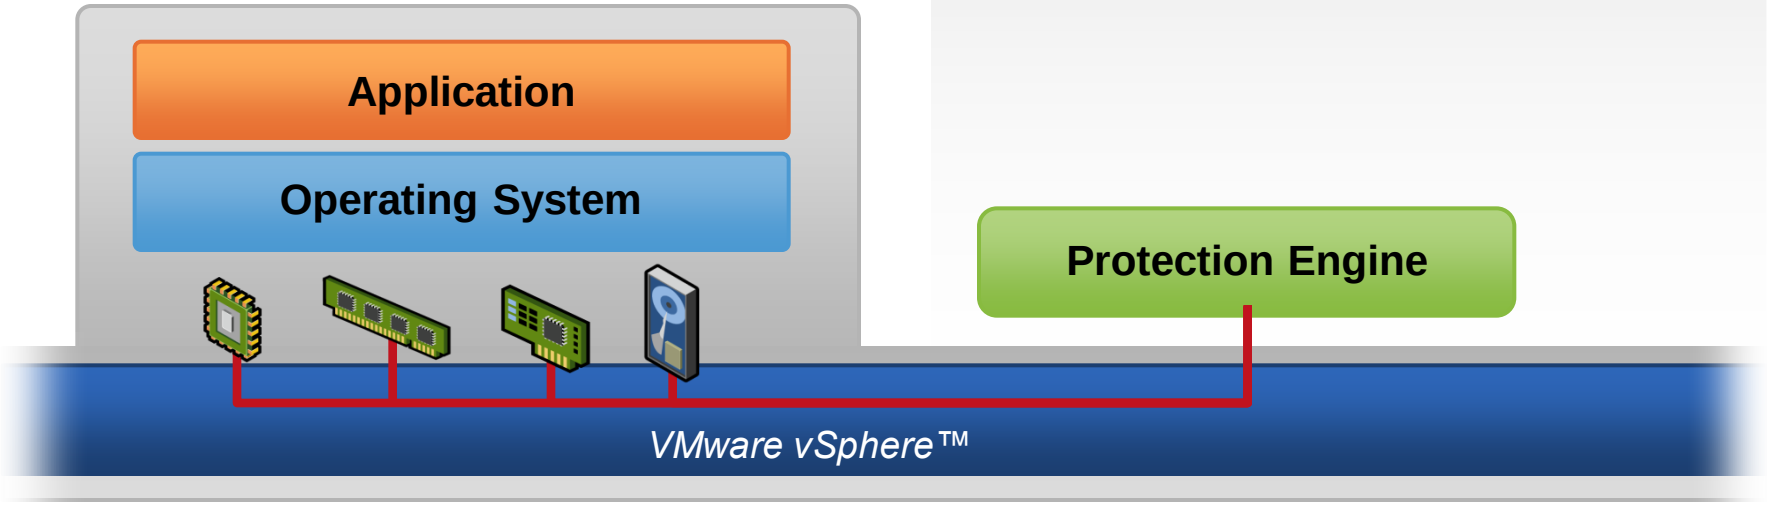

vShield Zones

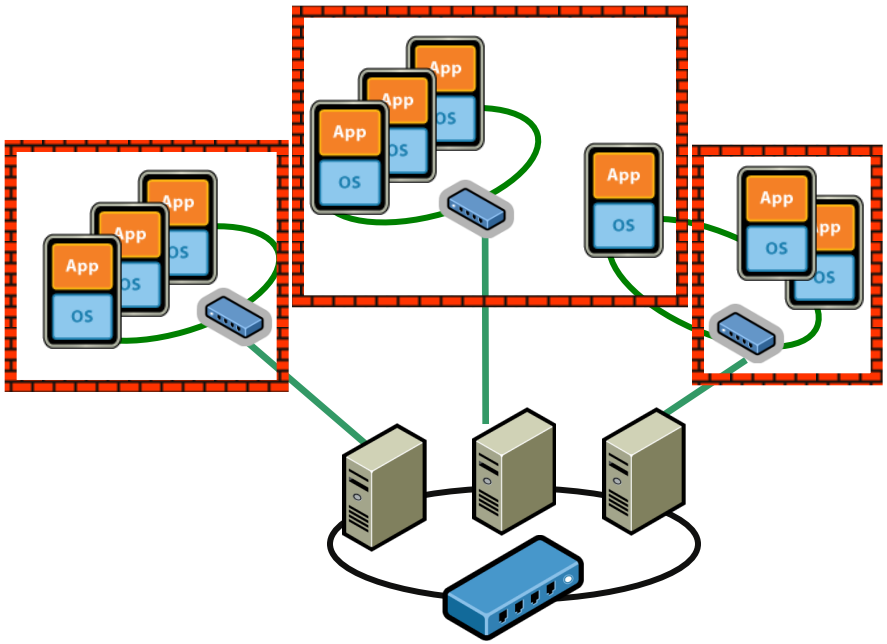

Capabilities

- Bridge, firewall, or isolate VM zones based on familiar VI containers
- Monitor allowed and disallowed activity by application-based protocols
- One-click flow-to-firewall blocks precise network traffic

Benefits

- Well-defined security posture within virtual environment
- Monitoring and assured policies, even through Vmotion and VM lifecycle events
- Simple zone-based rules reduces policy errors

Logical Hierarchy of Zones

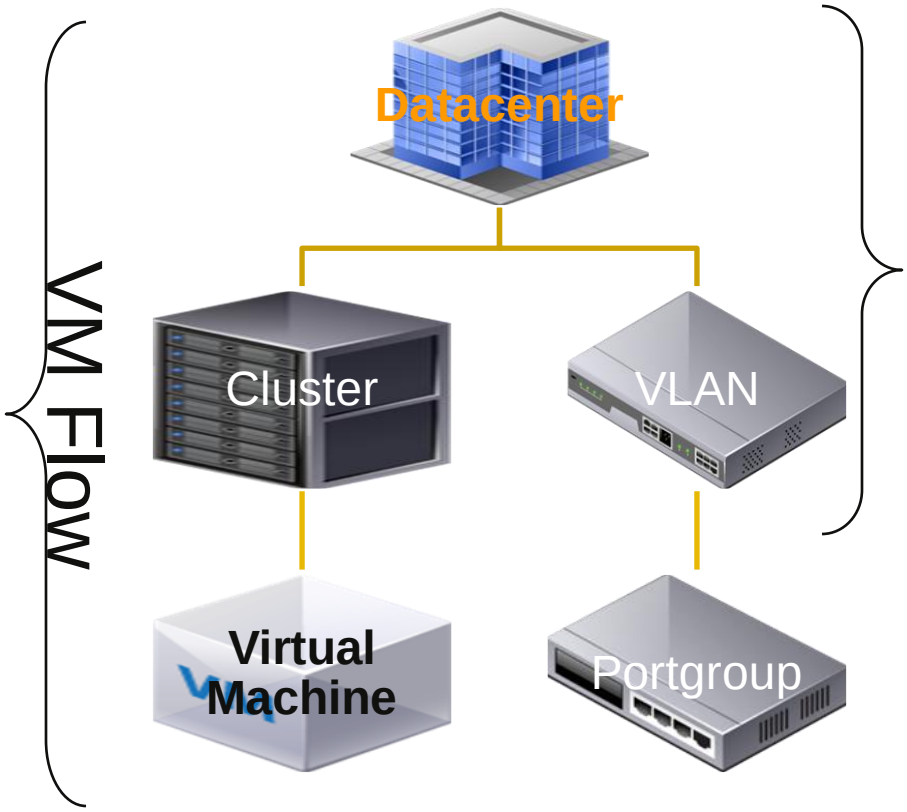

- VM Flow: Monitoring
  - Can monitor and audit traffic at all levels of hierarchy
- VM Wall: Blocking
  - Can define rules at Cluster, VLAN, Datacenter level
    - Network packet processed according to first match
    - Order of rule processing can be modified
  - Policy maintained automatically even for inventory changes, e.g.
    - FT failover
    - New VM provisioned
    - Host changes cluster
    - VLAN extended to new host

New HA Cluster Settings

Availability | Security | Scalability

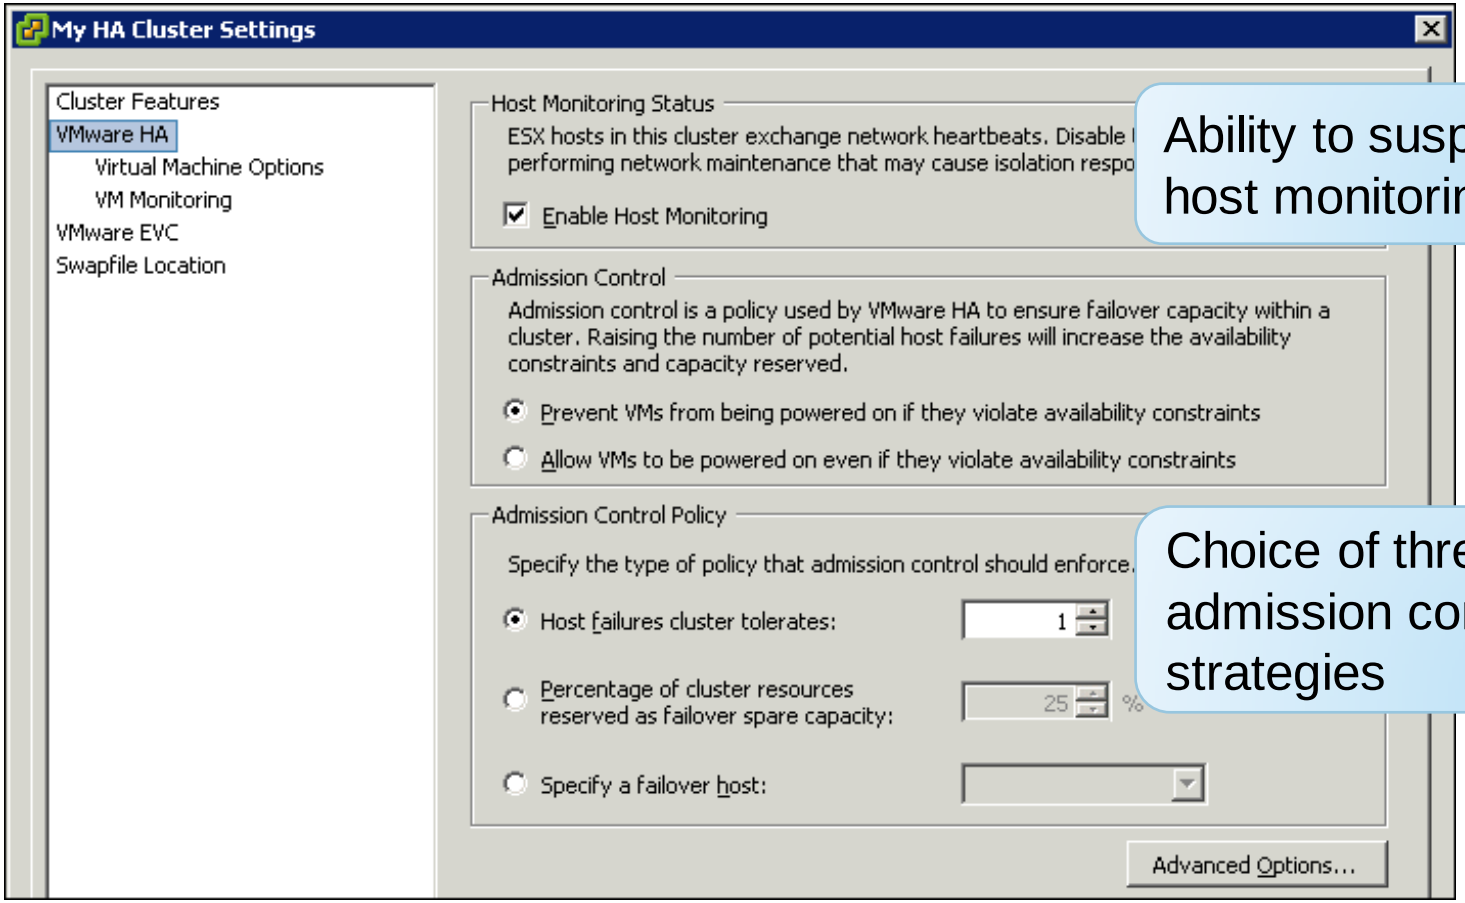

Ability to suspend host monitoring

Choice of three admission control strategies

Hot Add for Memory and CPU

Availability | Security | **Scalability**

Virtual Machine > Edit Settings > Options Tab > Memory/CPU Hotplug

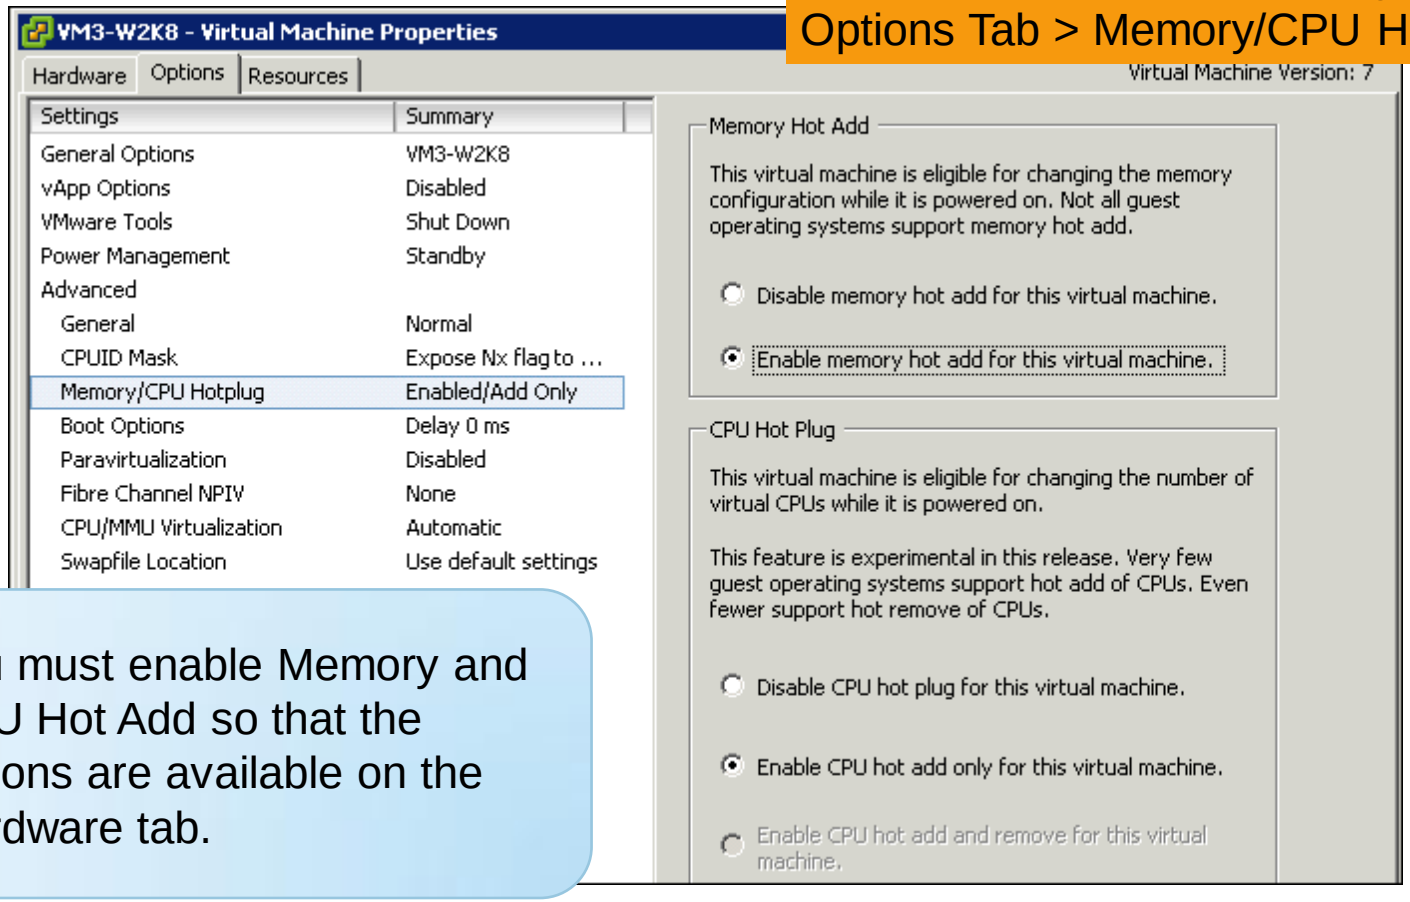

You must enable Memory and CPU Hot Add so that the options are available on the Hardware tab.

Host Profiles Overview

vCenter

Host profiles reduce setup time and allow you to manage configuration consistency and correctness.

Reference Host

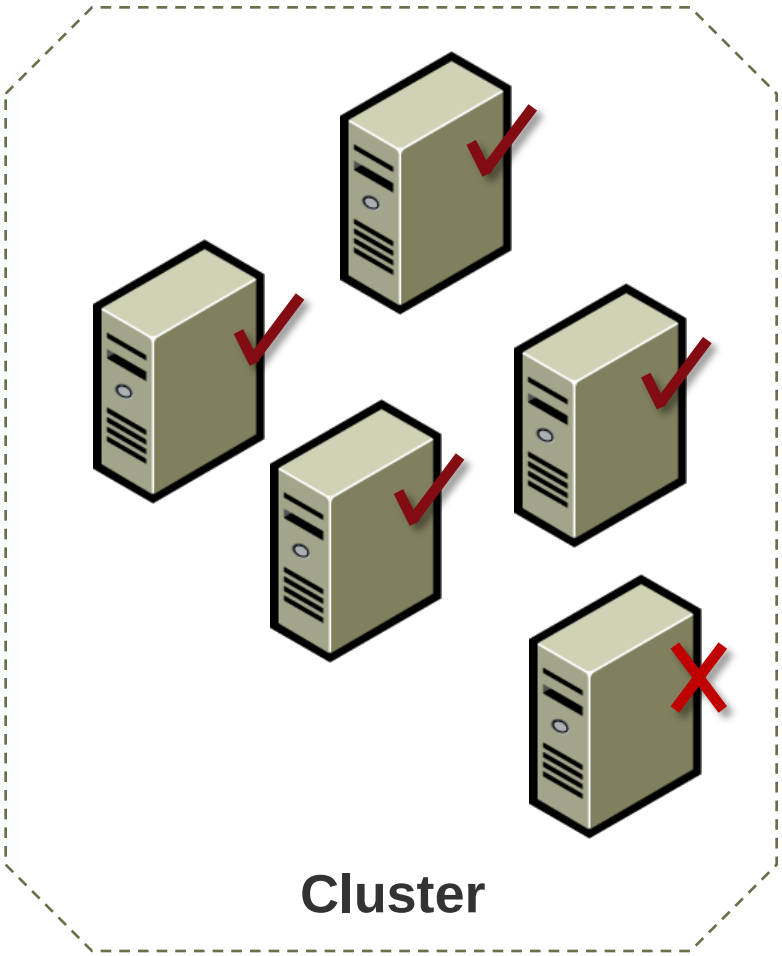

Cluster

Summary of VMware vSphere™

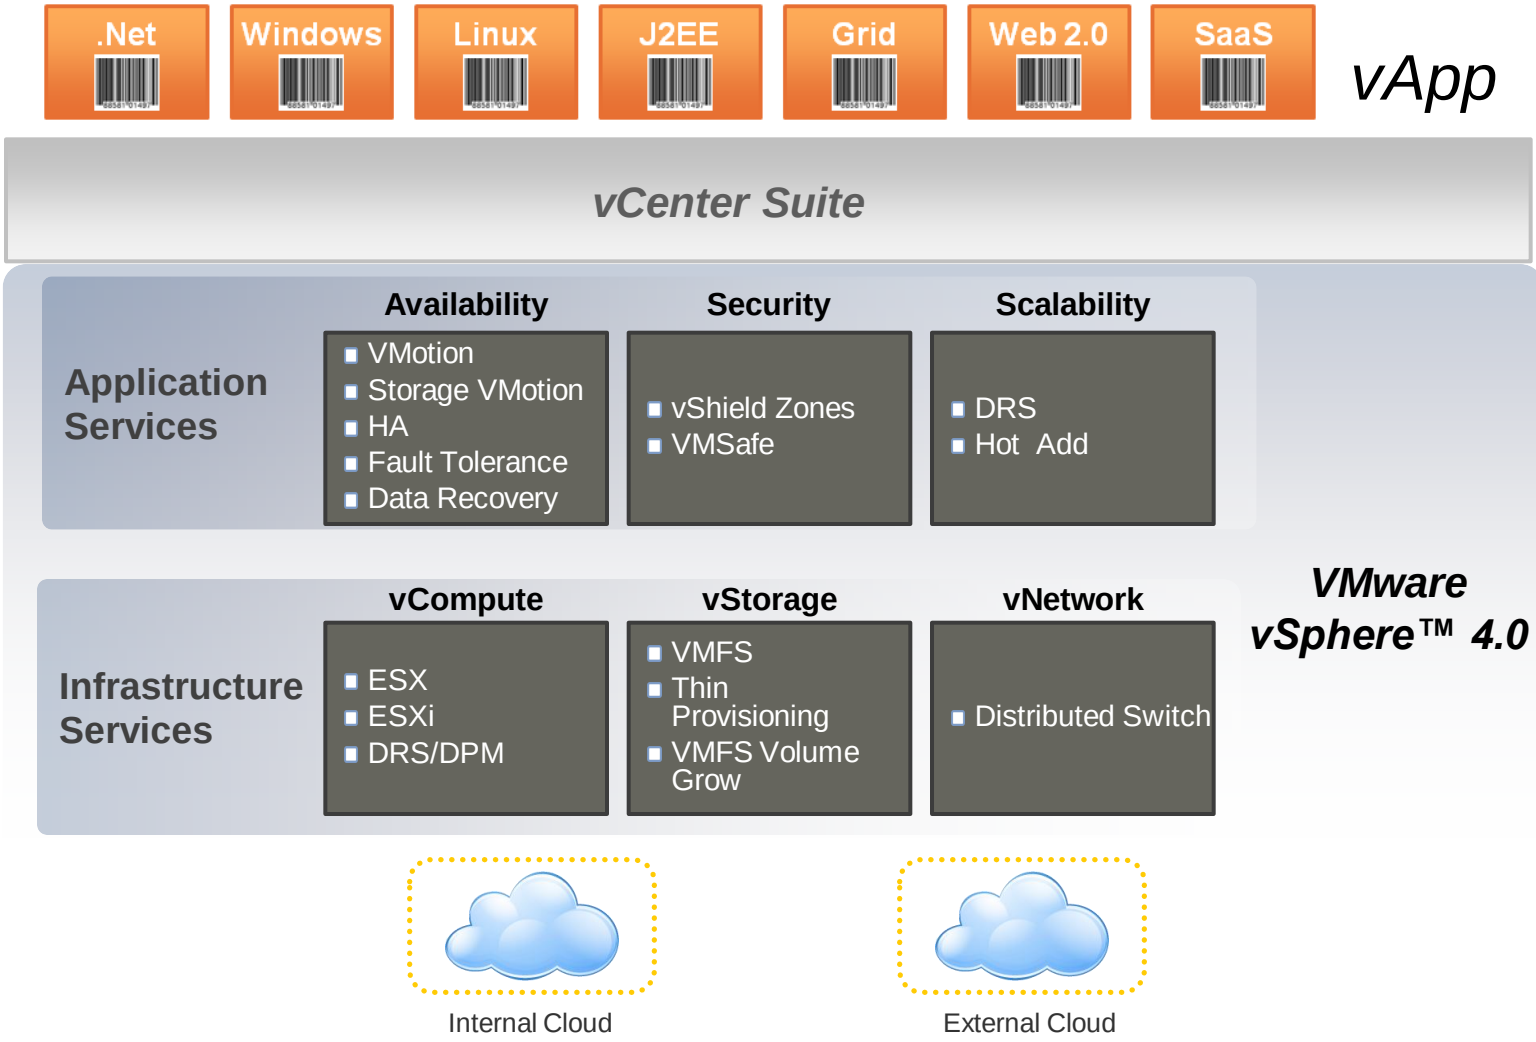

*Note vCenter Server and its components are a separate purchase

Thank You !


vmware®

Nexus 1000V
Intro and Setup

Agenda

Visibility in virtualized environments

N1k in more detail

VSM options

Connectivity options

Comparison vswitch/DVS/N1k

Challenges with Server-Virtualization

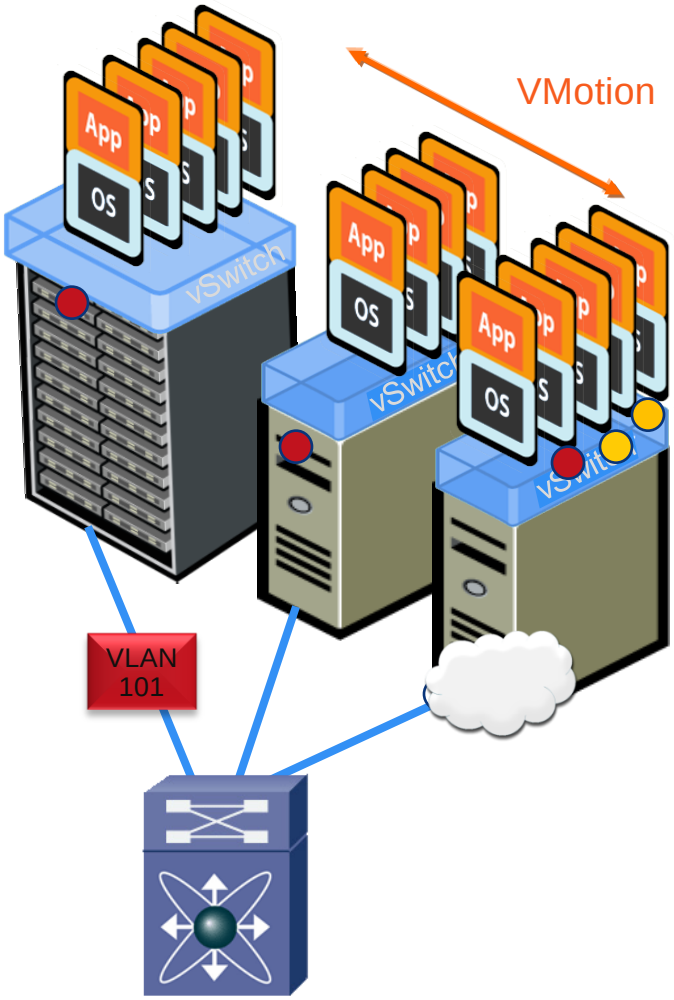

Problems:

- 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

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
```

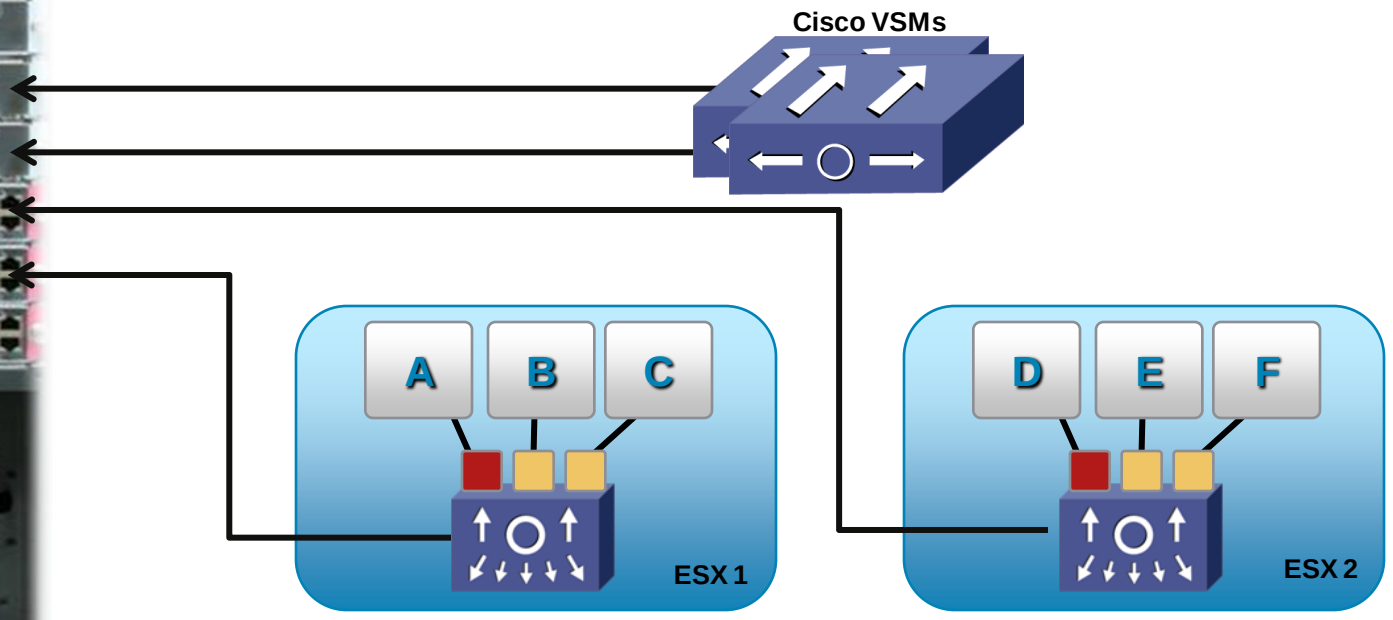

Cisco Nexus 1000V

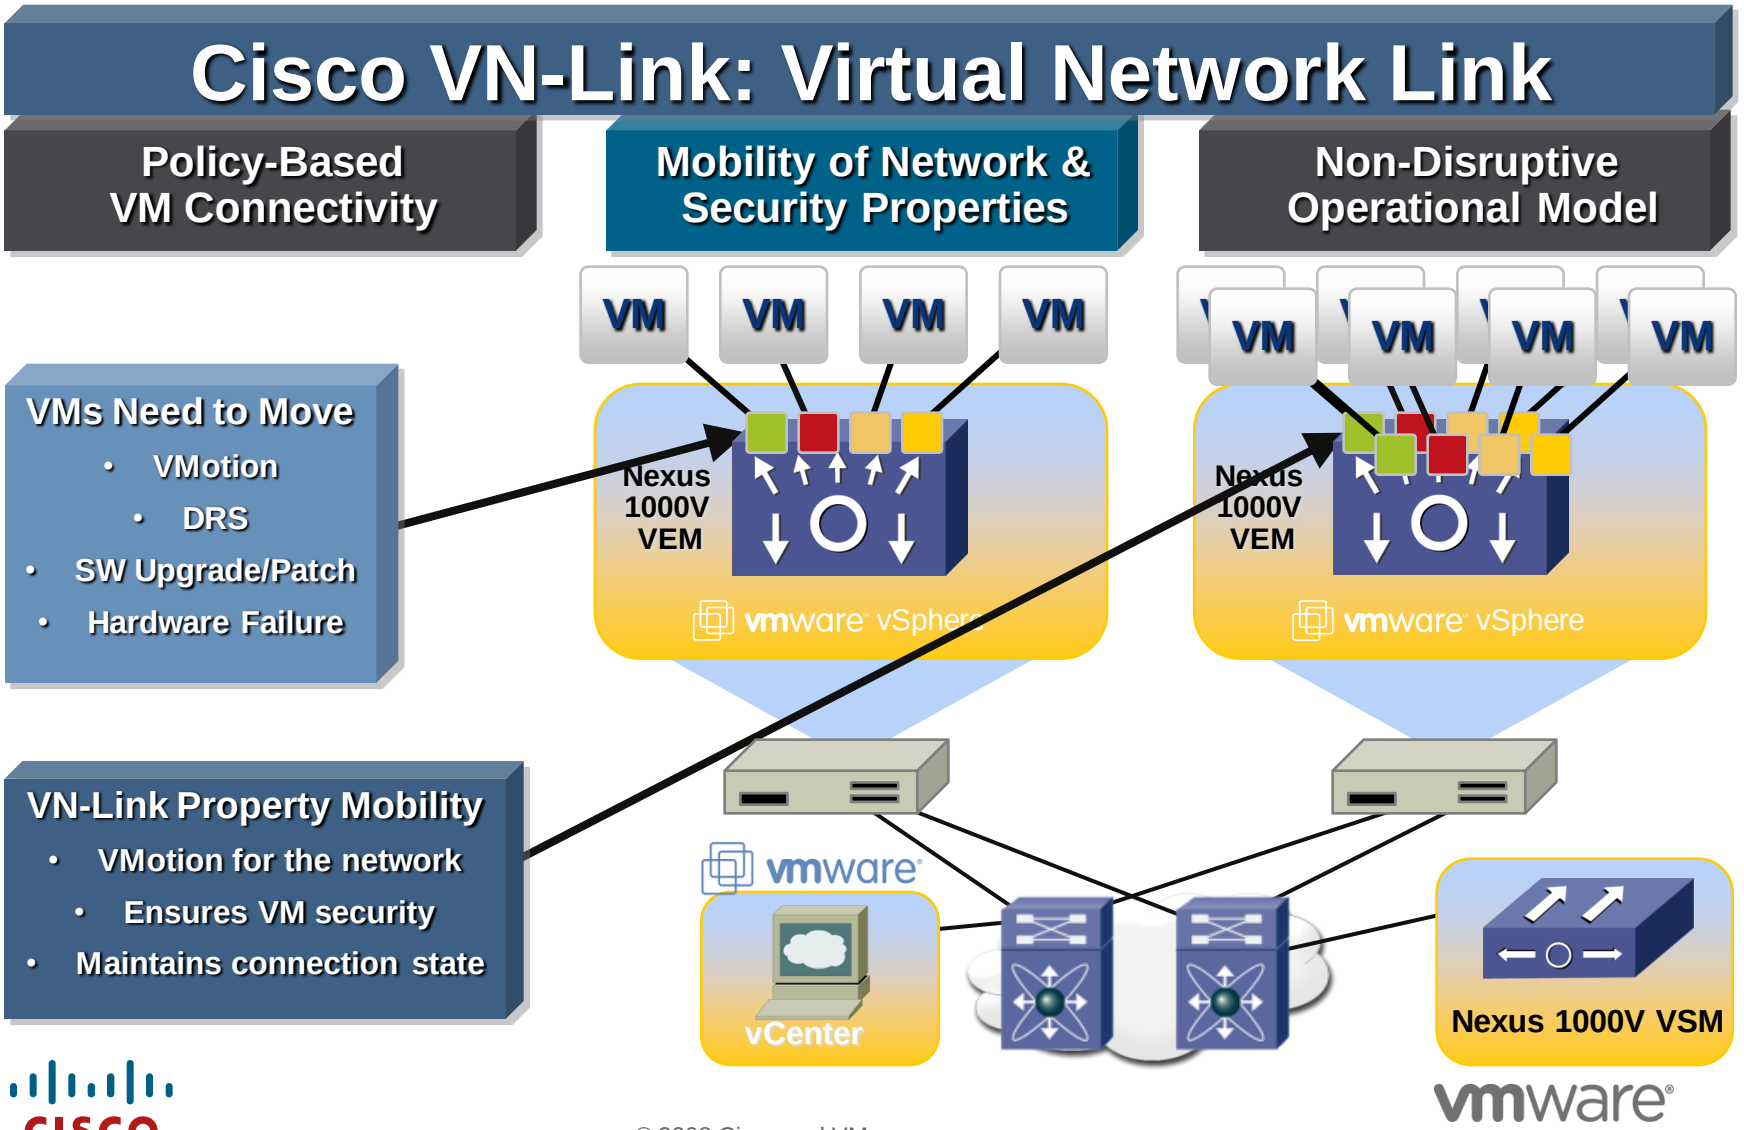

Cisco Nexus 1000V

Three New Features that Make a Difference

Encapsulated Remote SPAN (ERSPAN)

- Mirror VM interface traffic to a remote sniffer
- Identify root cause for connectivity issues
- No host-based sniffer virtual appliance to maintain
- Follows your VM with VMotion or DRS

NetFlow v.9 with Data Export

- View flow-based stats for individual VMs
- Captures multi-tiered app traffic inside a single ESX host
- Export aggregate stats to dedicated collector for DC-wide VM view
- Follows your VM with VMotion or DRS

Private VLANs (PVLANs)

- Great for mixed use ESX clusters
- Segment VMs w/o burning IP addresses
- Supports isolated, community and promiscuous trunk ports
- Follows your VM with VMotion or DRS

Visibility in virtualized environments

N1k in more detail

VSM options

Connectivity options

Comparison vswitch/DVS/N1k

Port Profile: Network Admin View

```
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

Port Profile: Server Admin View

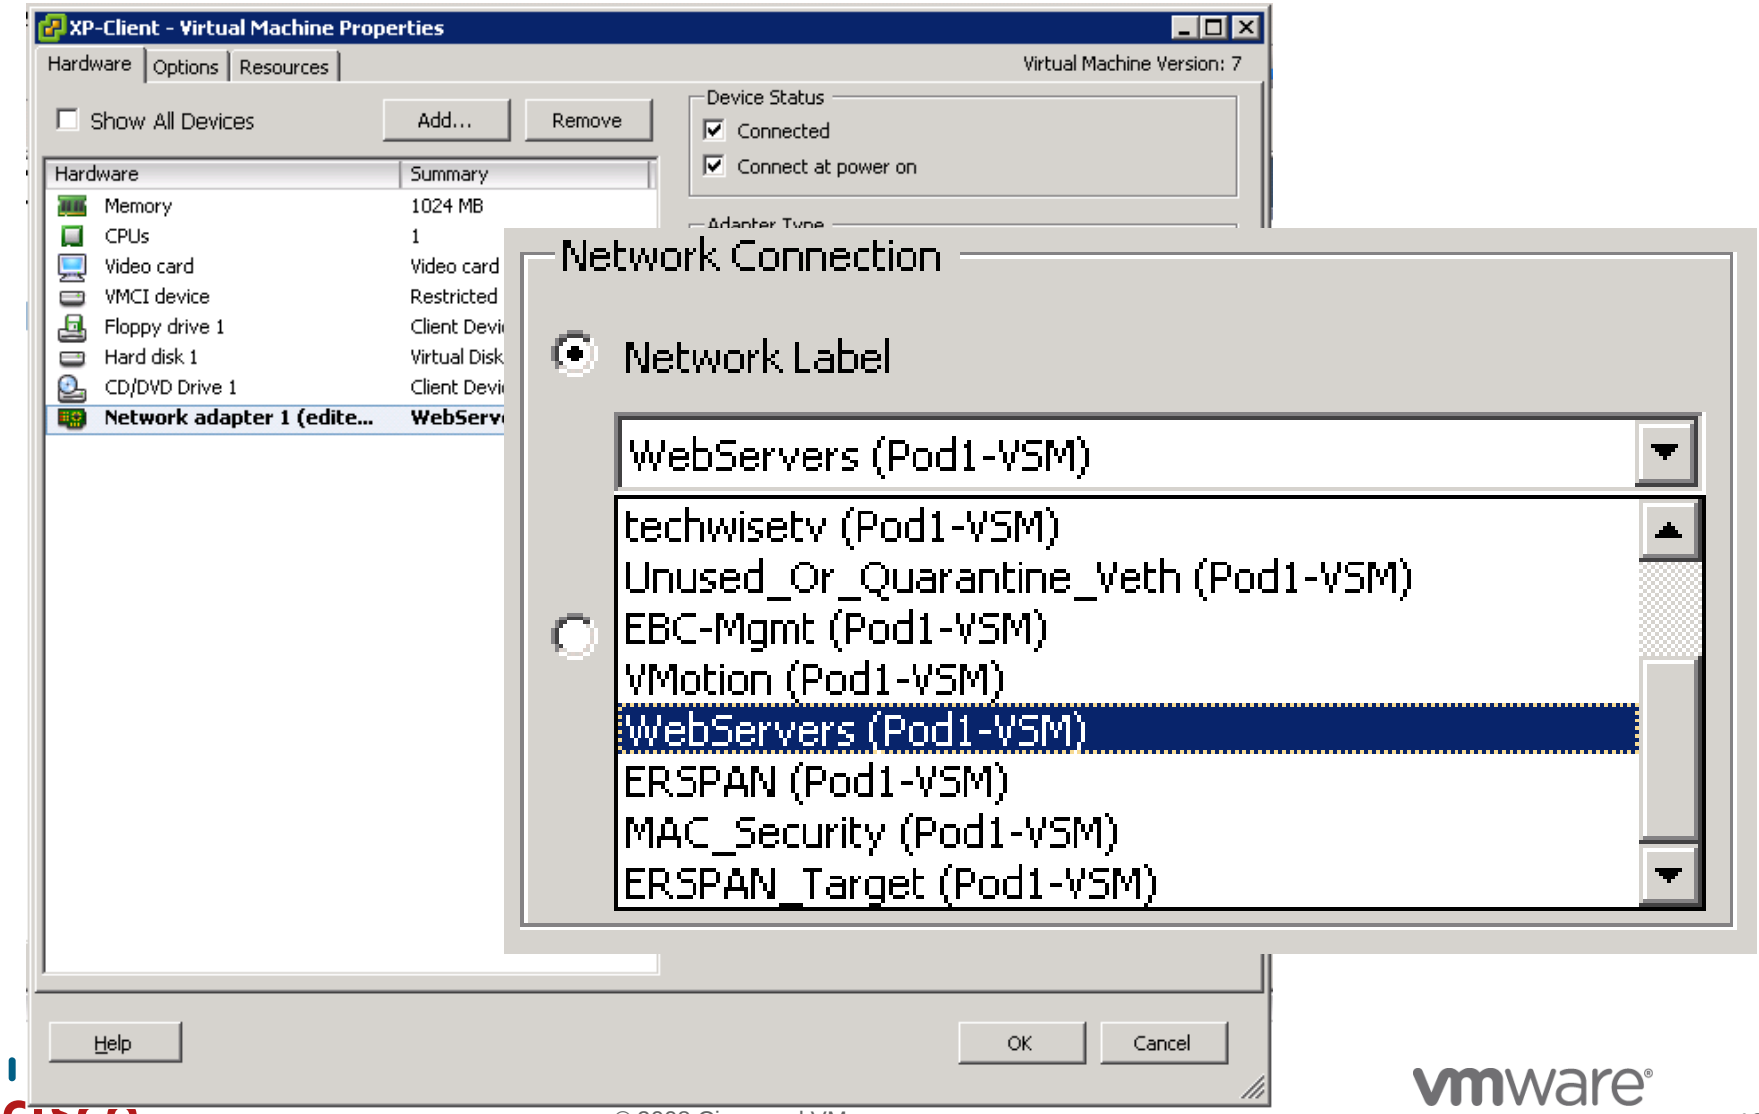

Visibility in virtualized environments

N1k in more detail

VSM options

Connectivity options

Comparison vswitch/DVS/N1k

Architecture Comparison

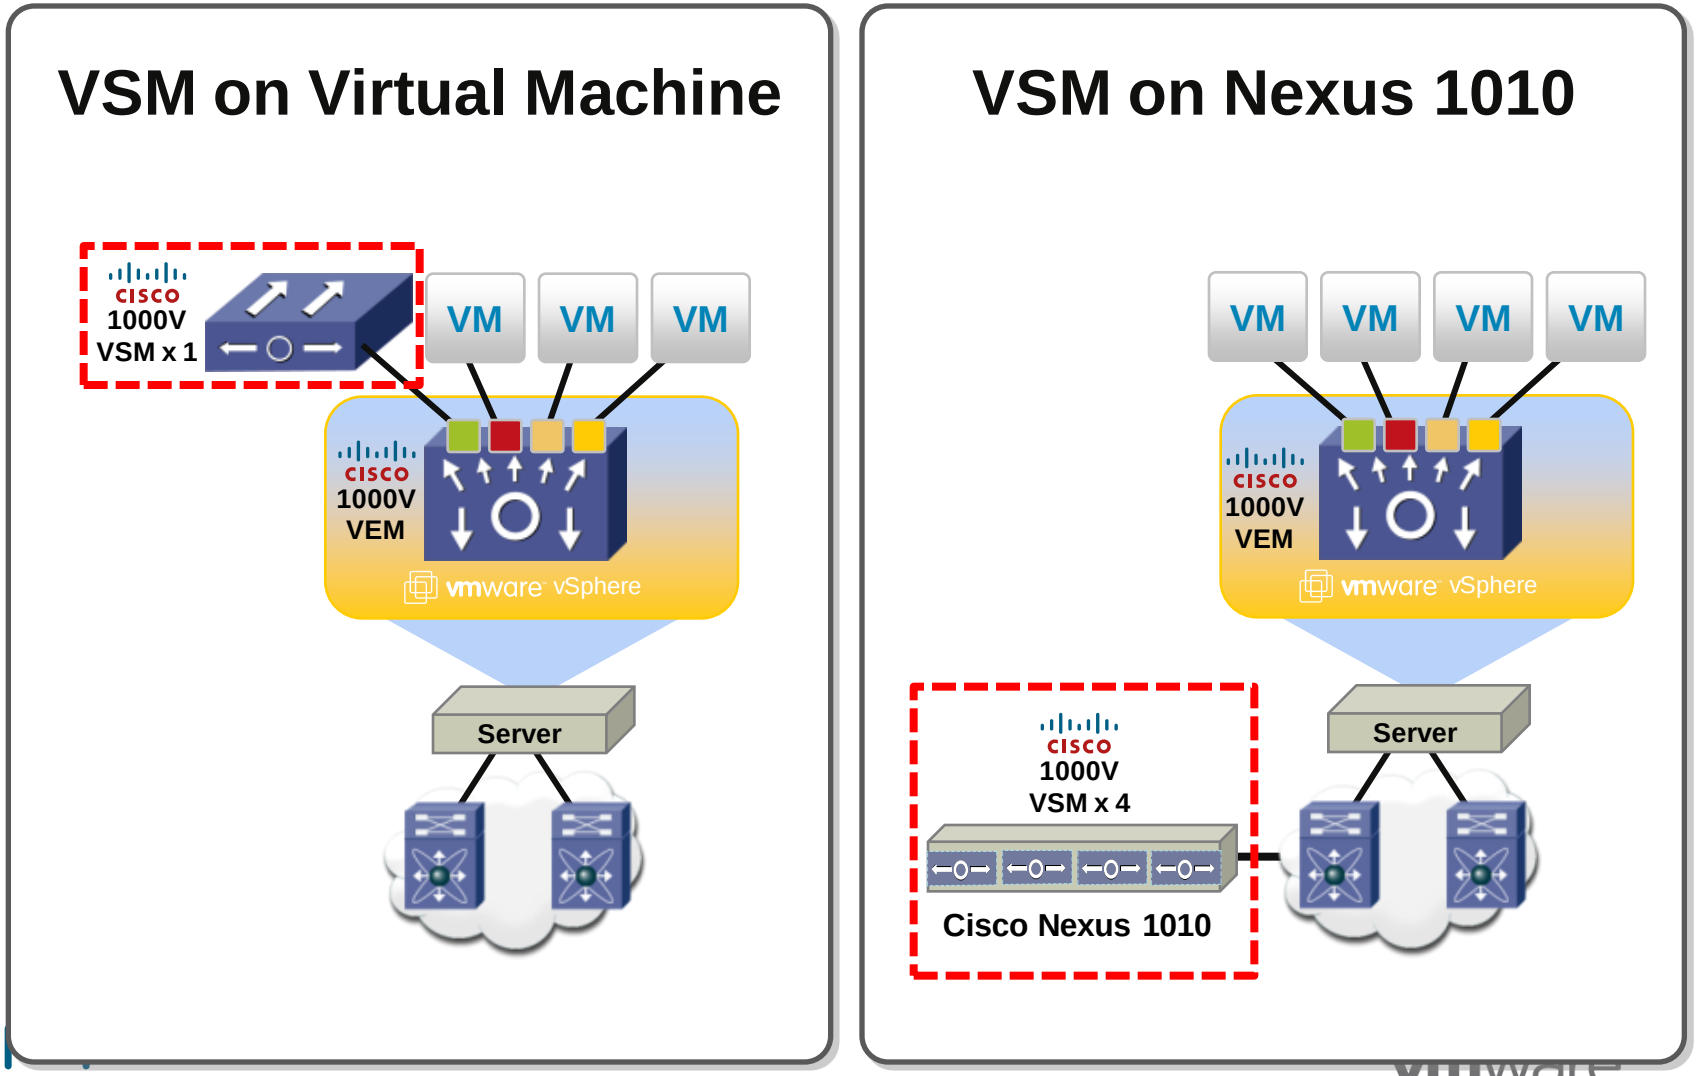

What is the Nexus 1010?

- The Nexus 1010 is a dedicated networking appliance to host four Nexus 1000V virtual supervisor modules (VSM)
- Allows network administrators to manage the Nexus 1000V supervisor like a standard Cisco switch, with all the same 1000V features
- $24,995 per appliance (inc. 32 1000V licenses till 12/31/10)
- Available April/May 2010

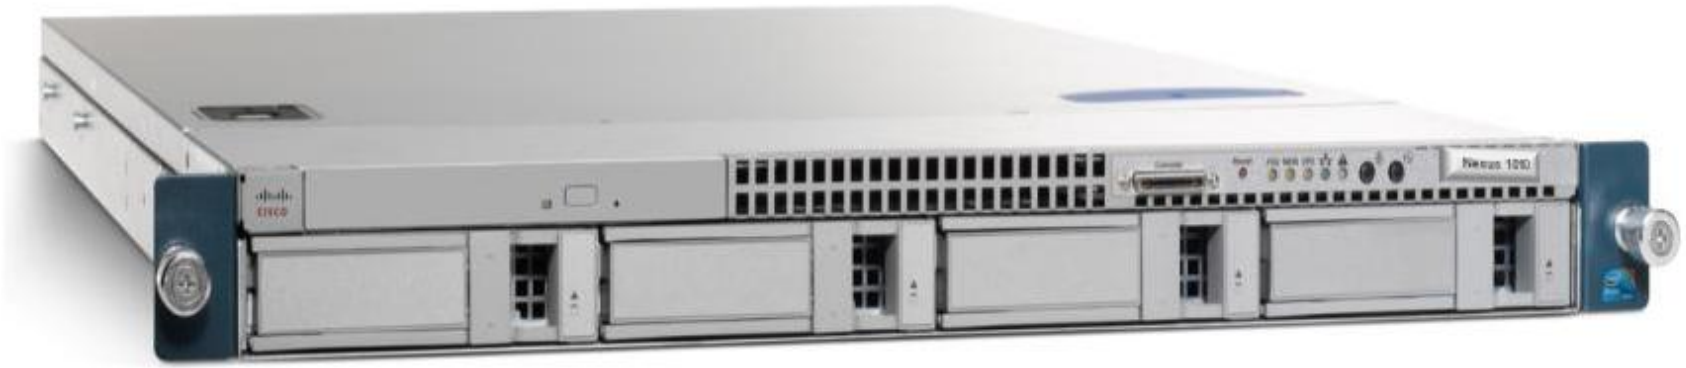

Feature Comparison

	Network Team manages the switch hardware
	Installation like a standard Cisco switch
NX-OS high availability of VSM	NX-OS high availability of VSM
VEM running on vSphere 4 Enterprise Plus	VEM running on vSphere 4 Enterprise Plus
Nexus 1000V features and scalability	Nexus 1000V features and scalability

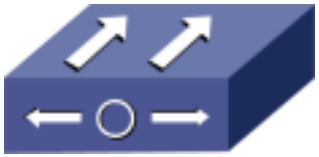

VSM on Virtual Machine

VSM on Nexus 1010

Optional: Network Analysis Module

- NAM is integrated with Nexus 1010 to provide:
  - Application and Traffic Performance Monitoring
  - Per-application, per-user, per-VM traffic analysis
  - Historical Reporting and Trending
- $ 3995 list price

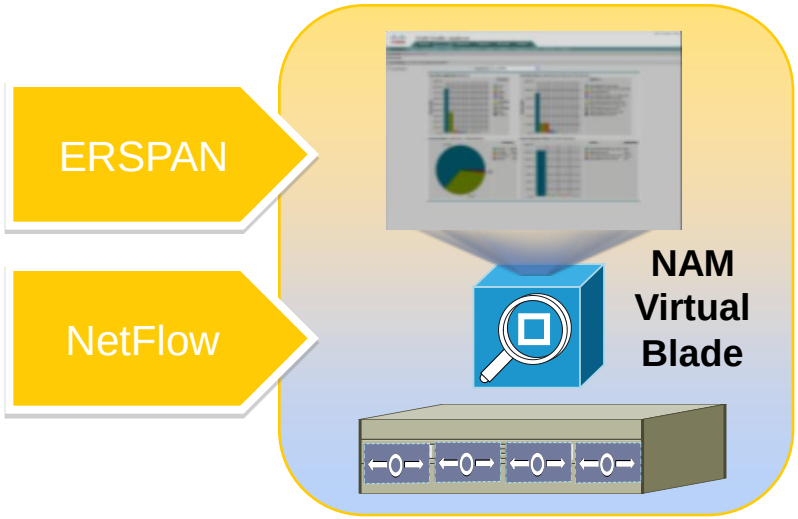

Visibility in virtualized environments

N1k in more detail

VSM options

Connectivity options

Comparison vswitch/DVS/N1k

Edge Node behavior

End-Host behavior

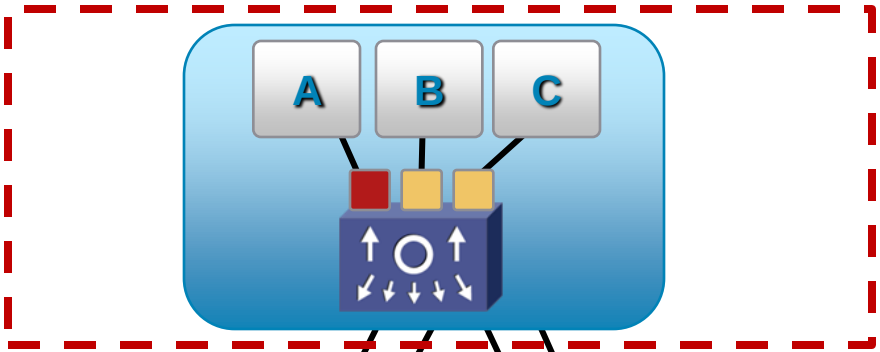

Typical Switch behavior

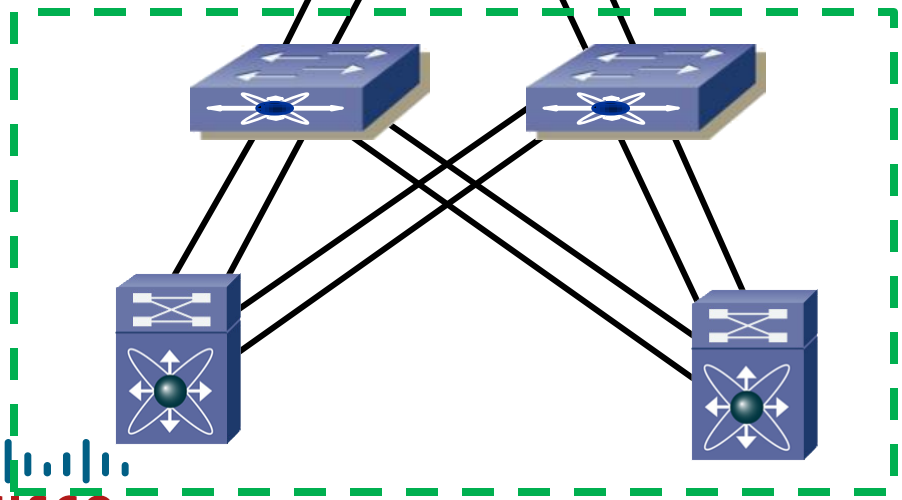

The Nexus 1000V is a end-host switch

An end-host switch can make some assumption to provide more features and an easier deployment

For example
-Port-channel on the Nexus 1000V doesn't require port-channel on the upstream switch

-Spanning tree not needed because a loop can be broken on the host itself

115

Loop Prevention without Spanning Tree

The Nexus 1000V being an end-host switch does NOT require spanning-tree to break loops

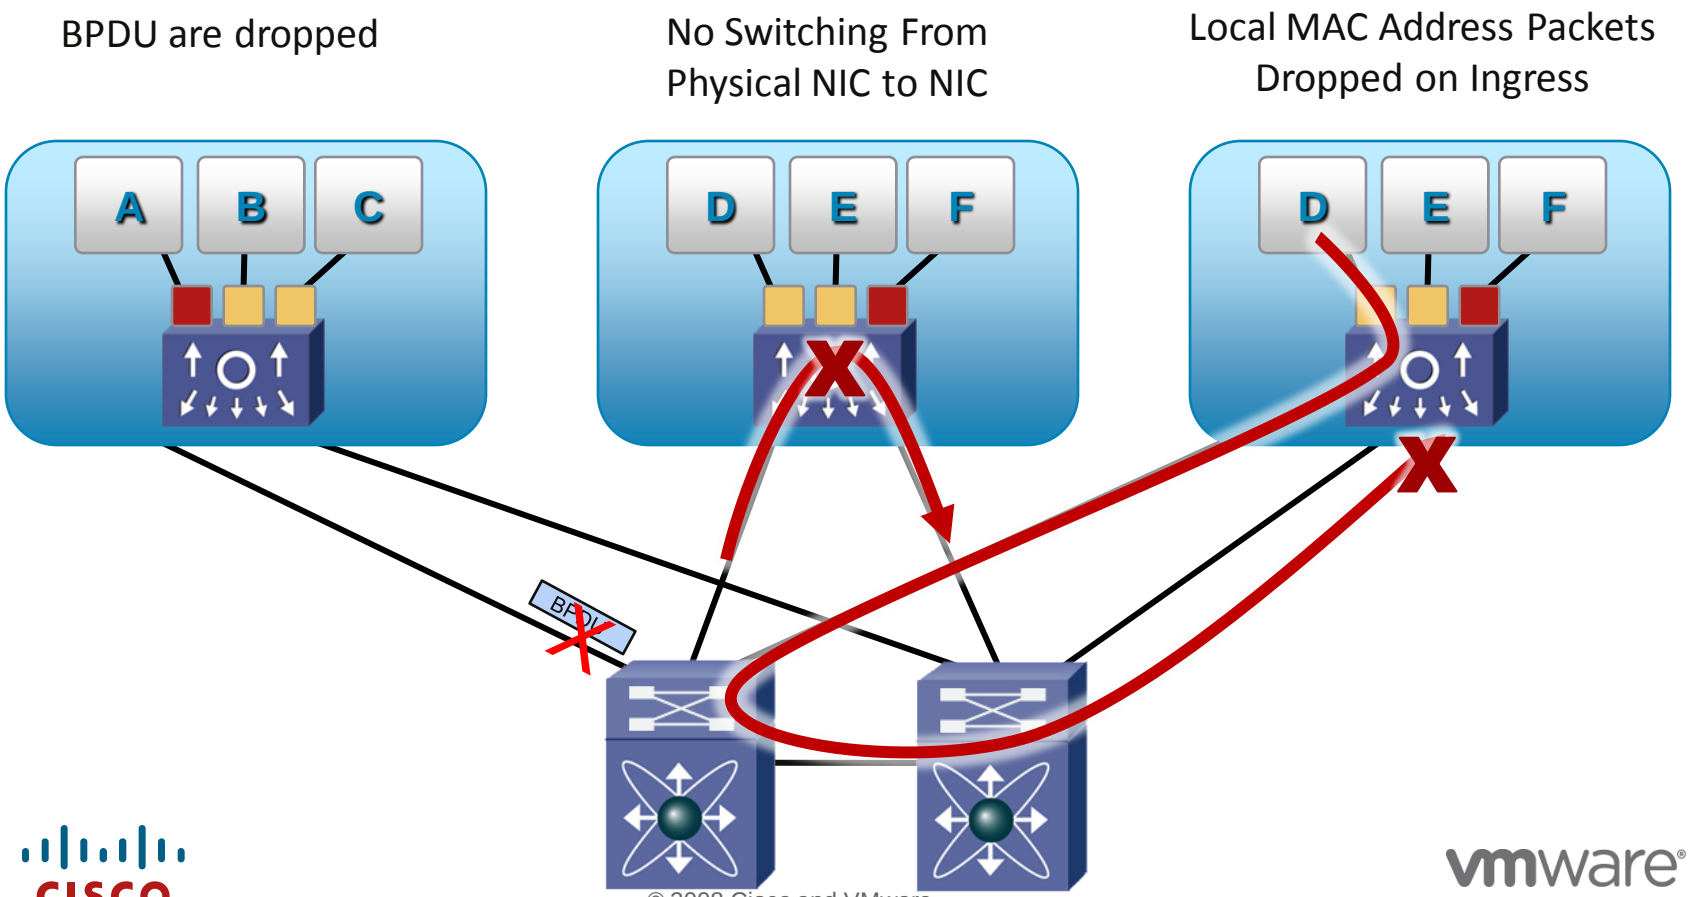

New Nexus 1000V Connectivity

The Nexus 1000V has different ways of traffic engineering between the VEM and the upstream switch:

- Virtual Port-Channel Host Mode
- Manual Configuration (of vPC HM, no CDP available)
- LACP
- MAC Pinning
- Pinned a port-profile to a sub-group

117

LACP with the Nexus 1000V

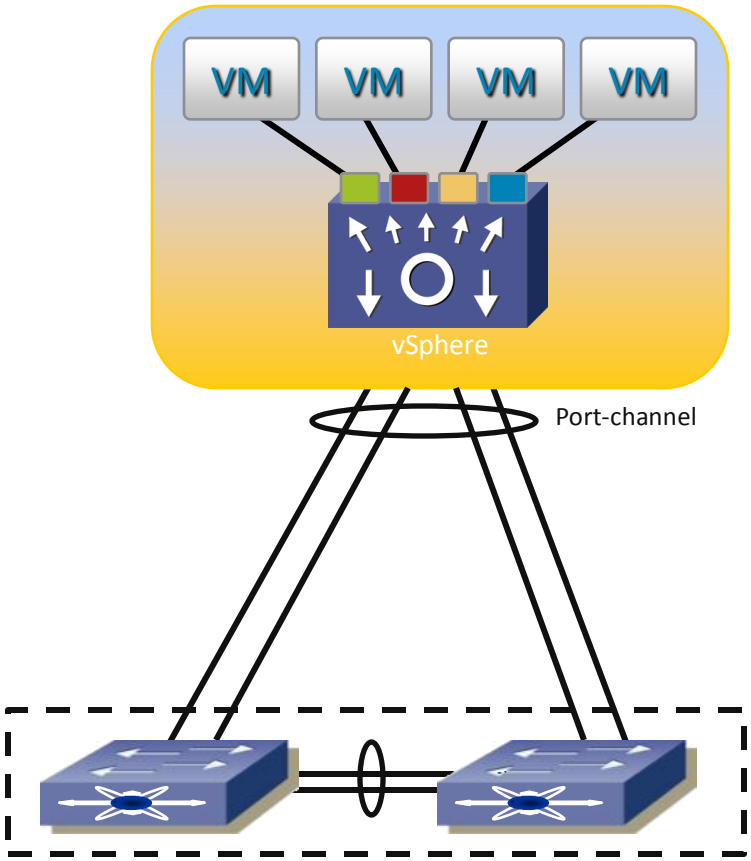

LACP is a port-channeling control mechanism to ensure the proper configuration of a port-channel

LACP allow the VMs and VMKernel Interfaces to utilize more than one link for its traffic

Allow faster VMotion and faster VM connectivity by using flow based hasing.

When using LACP with clustered upstream switch, VPC Host-mode will not be used

Upstream switch clustered (VPC,VSS,VBS,Stack…)

118

MAC Pinning

MAC Pinning provides the dynamism of VPC Host-Mode without requiring CDP Upstream

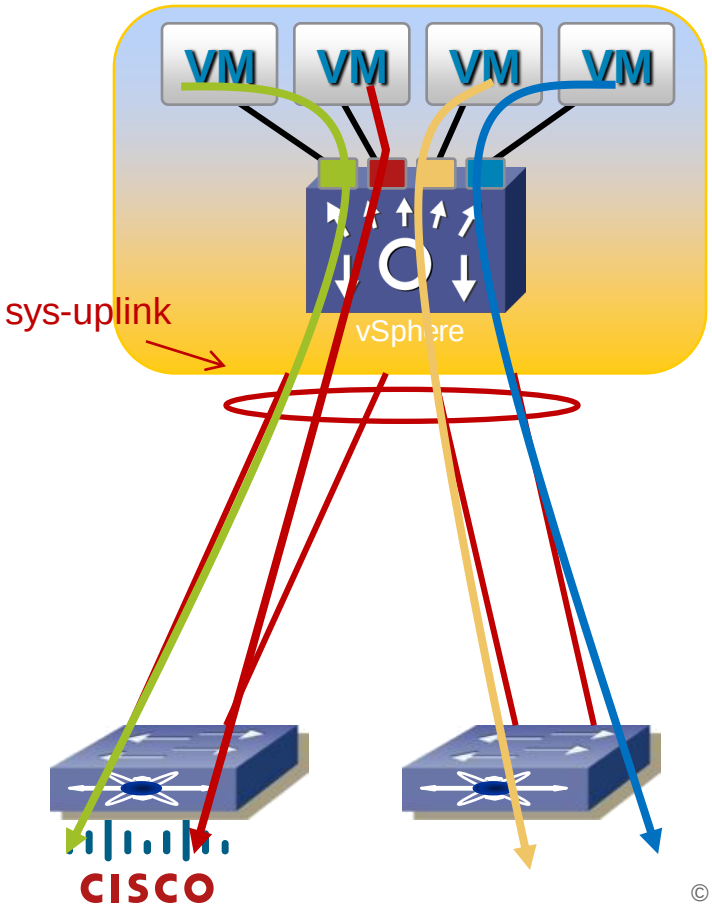

The MAC address of the VM will be used to select which link to use

```
Nexus1000(config)#port-profile sys-uplink
Nexus1000(config-port-prof)#no shut
Nexus1000(config-port-prof)#capability uplink
Nexus1000(config-port-prof)#channel-group auto mac-pinning
Nexus1000(config-port-prof)#switchport mode trunk
Nexus1000(config-port-prof)#switchport trunk allowed vlan 10-25
Nexus1000(config-port-prof)#state enabled
Nexus1000(config-port-prof)#vmware port-group
```


119

MAC Pinning

If a failover occurs, all the traffic pinned to an interface will be migrated to the other interfaces.

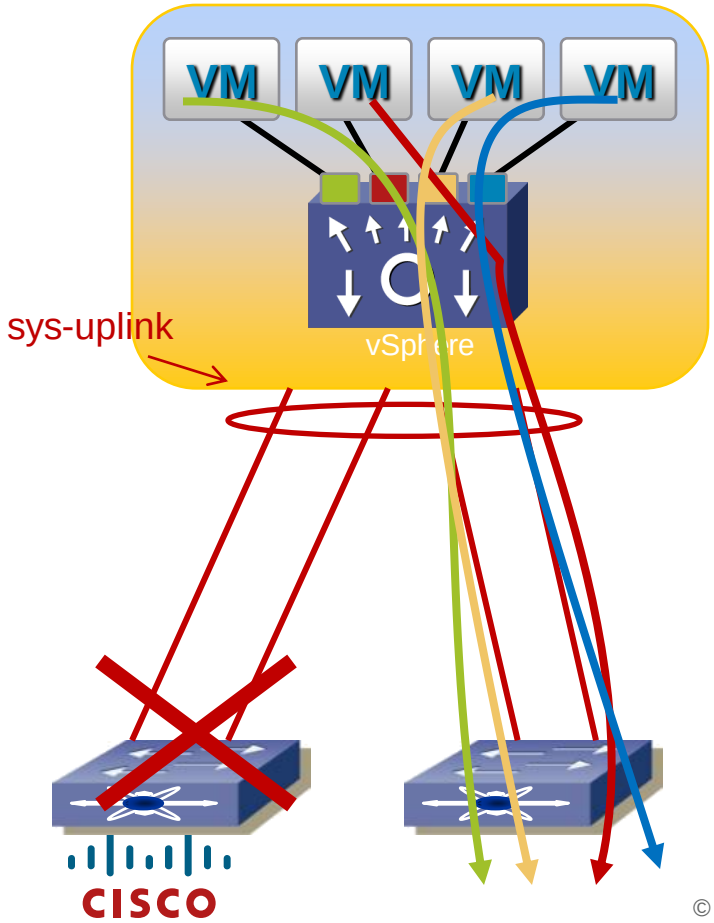


```
Nexus1000(config)#port-profile sys-uplink
Nexus1000(config-port-prof)#no shut
Nexus1000(config-port-prof)#capability uplink
Nexus1000(config-port-prof)#channel-group auto mac-pinning
Nexus1000(config-port-prof)#switchport mode trunk
Nexus1000(config-port-prof)#switchport trunk allowed vlan 10-25
Nexus1000(config-port-prof)#state enabled
Nexus1000(config-port-prof)#vmware port-group
```


120

Connectivity Best Practices

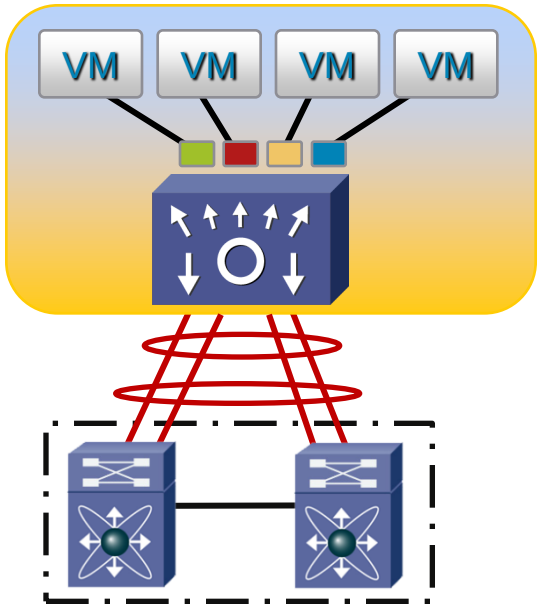

If the upstream switch can be clustered (VPC, VBS Stack, VSS) use LACP

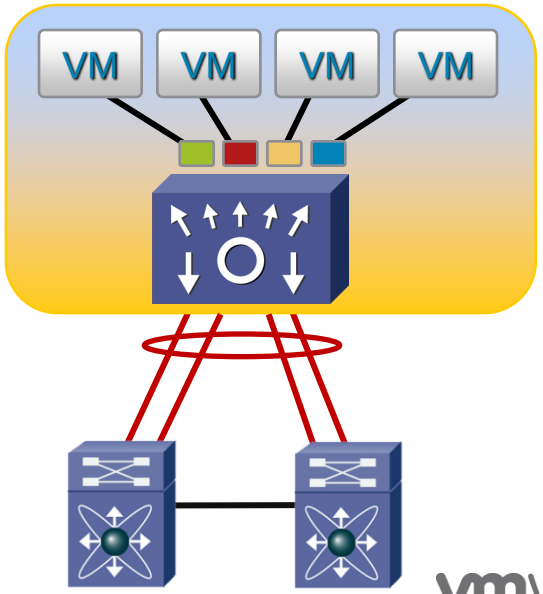

If the upstream switch can NOT be clustered use MAC-PINNING

124

Visibility in virtualized environments

N1k in more detail

VSM options

Connectivity options

Comparison vSwitch/DVS/N1k

Switch Feature Comparison 1

Feature	ESX 3.5: Standard vSwitch	ESX 4.0: vNetwork Standard Switch (U1)	ESX 4.0: vNetwork Distributed Switch (U1)	Nexus 1000V VSM on VM	Nexus 1000V VSM on Nexus 1010
Switching Features					
Layer 2 Forwarding	Yes	Yes	Yes	Yes	Yes
IEEE 802.1Q VLAN Tagging	Yes	Yes	Yes	Yes	Yes
Multicast Support (IGMP v2 and v3)	Yes	Yes	Yes	Yes	Yes
IGMPv3 Snooping	-	-	-	Yes	Yes
VMware VMotion Support	Yes	Yes	Yes	Yes	Yes
Network VMware VMotion (Network Policy)	-	-	Yes	Yes	Yes
Upstream Switch Connectivity					
Virtual MAC Pinning	Yes	Yes	Yes	Yes	Yes
EtherChannel	Yes	Yes	Yes	Yes	Yes
Virtual Port Channels	-	-	-	Yes	Yes
Link Aggregation Control Protocol (LACP)	-	-	-	Yes	Yes
Load Balancing Algorithms					
Virtual Switchport ID	Yes	Yes	Yes	Yes	Yes
Source MAC	Yes	Yes	Yes	Yes	Yes
Source and Destination IP	Yes	Yes	Yes	Yes	Yes
Source and Destination MAC	-	-	-	Yes	Yes
Source and Destination Port IP	-	-	-	Yes	Yes
Additional Hashing Options	-	-	-	Yes	Yes

Switch Feature Comparison 2

Feature	ESX 3.5: Standard vSwitch	ESX 4.0: vNetwork Standard Switch (U1)	ESX 4.0: vNetwork Distributed Switch (U1)	Nexus 1000V VSM on VM	Nexus 1000V VSM on Nexus 1010
Traffic Management Features					
Tx Rate Limiting (from virtual machine)	Yes	Yes	Yes	Yes	Yes
Rx Rate Limiting (from virtual machine)	-	-	Yes	Yes	Yes
iSCSI Multipathing	-	Yes	Yes	Yes	Yes
Quality-of-service (QoS) marking					
Differentiated Services Code Point (DSCP)	-	-	-	Yes	Yes
Type of Service	-	-	-	Yes	Yes
Class of Service	-	-	-	Yes	Yes
Security Features					
Port Security	Yes	Yes	Yes	Yes	Yes
VMware VMSafe compatible	Yes	Yes	Yes	Yes	Yes
Private VLANs (PVLANs)	-	-	Yes	Yes	Yes
Local PVLAN enforcement	-	-	-	Yes	Yes
Access Control Lists (ACL)	-	-	-	Yes	Yes
DHCP Snooping	-	-	-	Yes	Yes
IP Source Guard	-	-	-	Yes	Yes
Dynamic ARP Inspection	-	-	-	Yes	Yes
Virtual Service Domain	-	-	-	Yes	Yes

Switch Feature Comparison 3

Feature	ESX 3.5: Standard vSwitch	ESX 4.0: vNetwork Standard Switch (U1)	ESX 4.0: vNetwork Distributed Switch (U1)	Nexus 1000V VSM on VM	Nexus 1000V VSM on Nexus 1010
Management Features					
VMware vCenter Support	Yes	Yes	Yes	Yes	Yes
Third Party Accessible APIs	Yes	Yes	Yes	Yes	Yes
Network Policy Groups	Yes	Yes	Yes	Yes	Yes
VMware port mirroring (promiscuous)	Yes	Yes	Yes	-	-
Multi-Tier Policy Groups (inheritance)	-	-	-	Yes	Yes
SPAN	-	-	-	Yes	Yes
ERSPAN	-	-	-	Yes	Yes
Netflow v9	-	-	-	Yes	Yes
SNMP v3 Read/Write	-	-	-	Yes	Yes
CDP v1/v2	Yes	Yes	Yes	Yes	Yes
Syslog	**	**	**	Yes	Yes
Packet Capture & Analysis	-	-	-	Yes	Yes
Radius/TACACS+	-	-	-	Yes	Yes
Configuration and management console and interface	VI Client	VI Client	VI Client to VMware vCenter Server	VMware vCenter and Cisco CLI	VMware vCenter and Cisco CLI
IPv6 for Management	Yes	Yes	Yes	Yes	Yes
NX-OS XML API	-	-	-	Yes	Yes

** Virtual switch network syslog information is exported and included with VMware ESX Server events.

Nexus 1000V and Virtual Desktop Deployment

The campus and data center used to be separated by firewall, protection. However Virtual Desktop blurs that boundaries.

The Nexus 1000V by providing security features, like DHCP Snooping, Dynamic ARP Inspection, usually used to secure the campus protect the Virtual Desktop Deployment

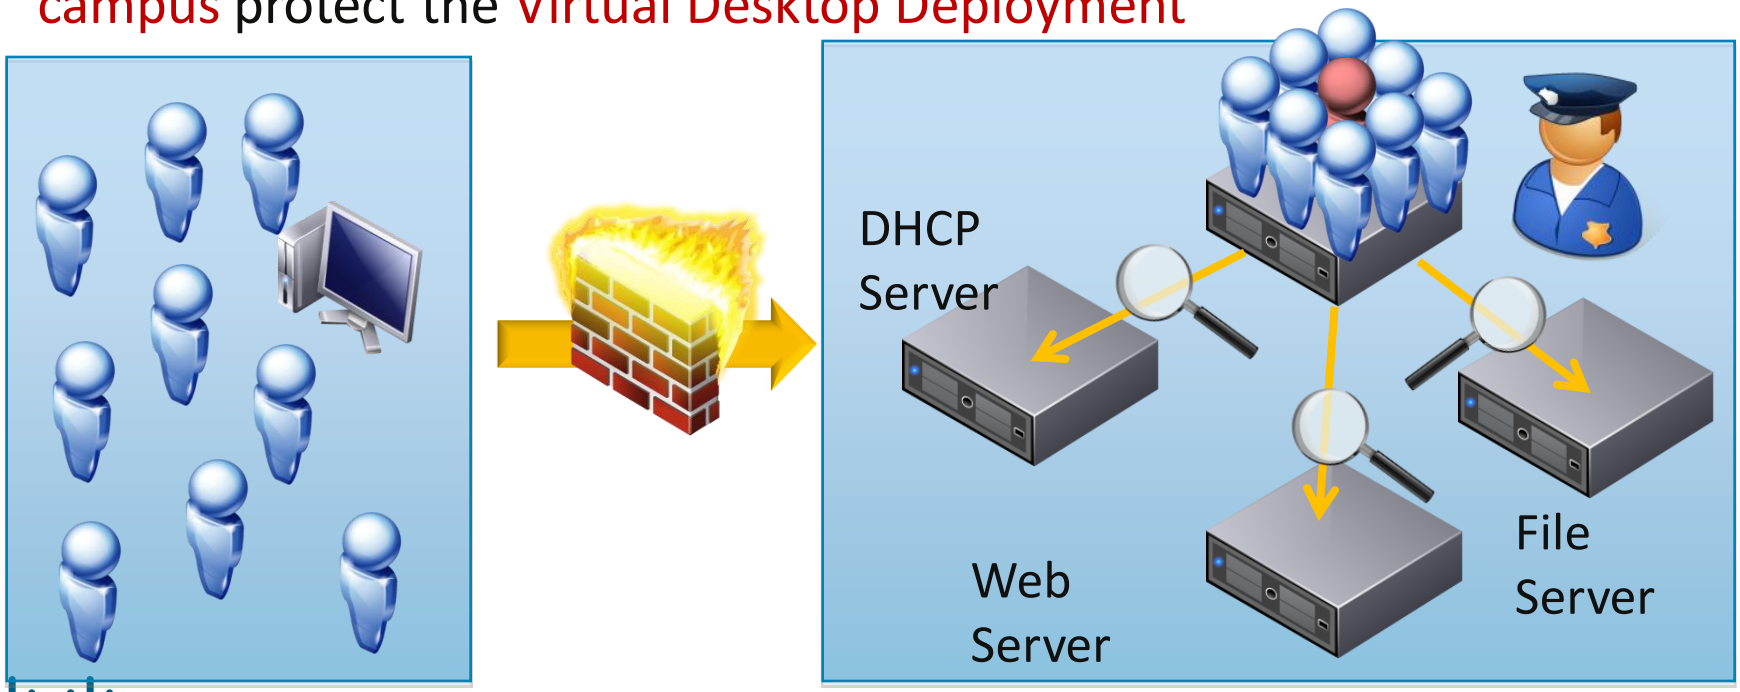

129

Nexus 1000V and DMZ Deployment

The virtualized DMZ takes advantage of virtualization technologies to reduce the DMZ footprint, thereby freeing valuable rack space, which in turn reduces power consumption and overall operating costs..

Maintaining Isolation and Protection

VLANs, Private VLANs, ACLs, Anti-Spoofing

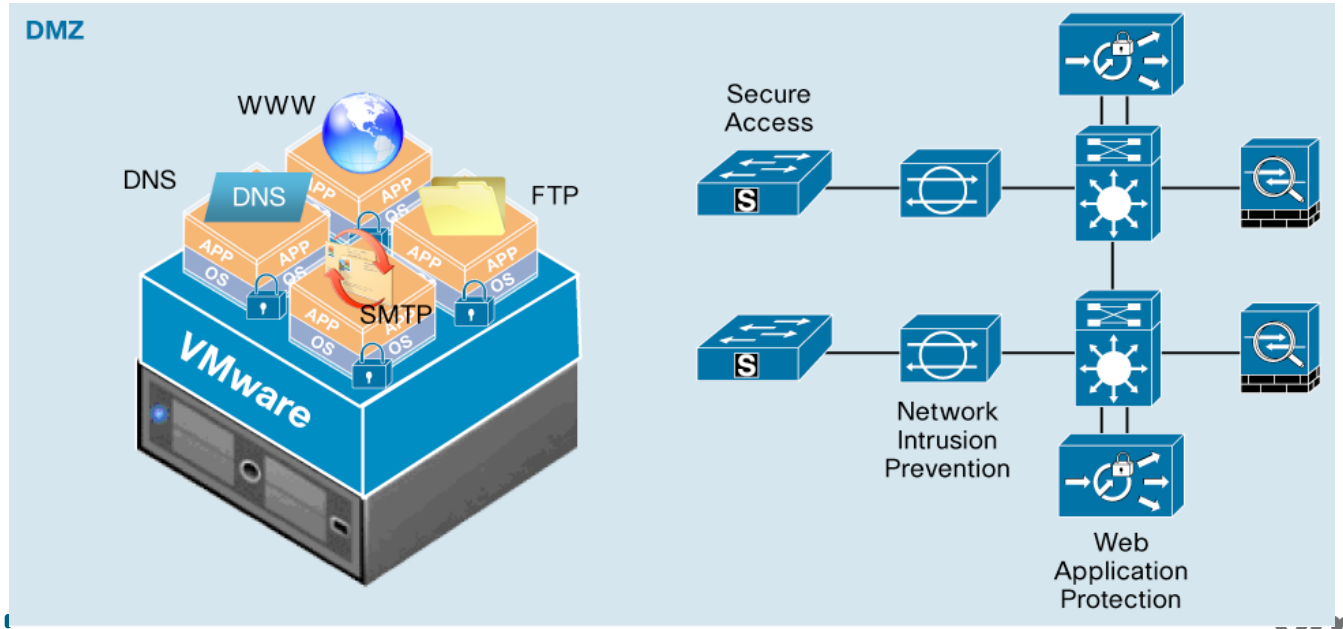

130

Evaluate

60 DAYS FREE

No-Charge Evaluation
www.cisco.com/go/1000veval

Bundle

Bundle $795 for Nexus 1000V
and vSphere Enterprise Plus

Summary

UCS Value Proposition

One System one Vendor

`trusted`Virtualization inside and Cloud ready

Service Profiles

Unified Fabric - FcoE

Extended Memory Blades

Embedded Management

Cabling

Cisco Unified Computing System

Undefined Computing System

86% cable reduction with the Cisco Unified Computing System

Bundle Pricing

Bundle 1
4 ESX vSphere B-200M2 hosts with 2 CPU (6 Cores) each
48 GB RAM, VIC with FCoE, 2*73GB disk
Nexus 1000v, UCS-Manager, 8G SAN all redundant
VMware Enterprise Plus License
1 Year 7*24*4 Onsite Service

Netto: 98‘000 USD no VAT included

Bundle 2
4 ESX vSphere B-250M2 hosts with 2 CPU (6 Cores) each
96 GB RAM, VIC with FCoE, 2*73GB disk
Nexus 1000v, UCS-Manager, 8G SAN all redundant
Enterprise Plus License
1 Year 7*24*4 Onsite Service

Netto: 179‘000 USD no VAT included

Check latest price bundle

UCS References

Visit our external reference site to find more:
http://www.cisco.com/en/US/netsol/ns944/index.html (scroll down to “Case Studies”)

LAB scenario

Virtualization Bootcamp 2011

Lunch

Intro LAB

Lab Setup

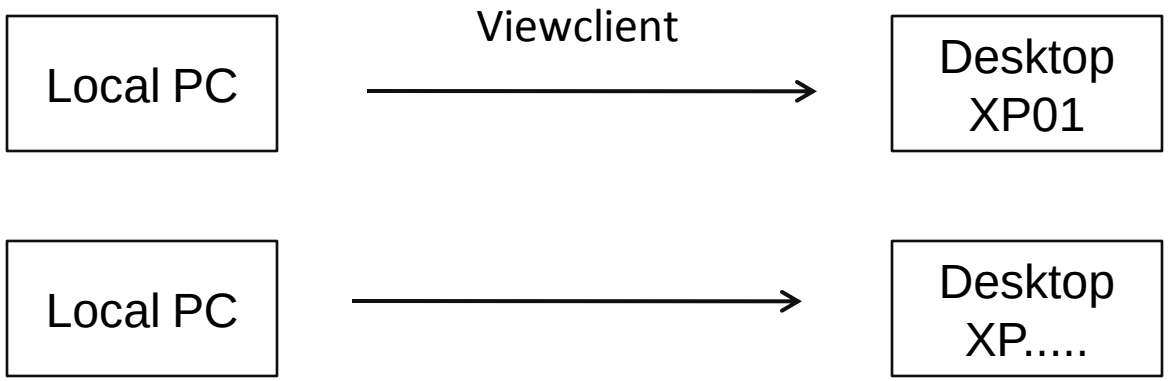

Desktop XP....
Shortcuts
UCS Manager
Vsphere Client

Download Viewclient:
http://ftpsite.vmware.com/download/view4.rar

Lab Setup

Vcenter01	XP01 XP02	
Vcenter02	XP03 XP04	
Vcenter03	XP05 XP06	
Vcenter04	XP07 XP08	
Vcenter05	XP09 XP10	
Vcenter06	XP11 XP12	
Vcenter07	XP13 XP14	
Vcenter08	XP15 XP16	
Vcenter09	XP17 XP18	
Vcenter10	XP19 XP20	
Vcenter11	XP21 XP22	
Vcenter12	XP23 XP24	
Vcenter13	XP25 XP26	

Lab 1: UCS Installation

Establish WLAN connection

Login to View4 server to get desktop

Create service profile on UCS-M

Assign service profile to pool

Create KVM Session to server

Mount ISO file

Explore UCS-M

Passwort View Client: see whiteboard

Lab2: Install and configure ESXi

Mount ESXi install ISO to server

Install ESXi on local disk

Configure ESXi

Integrate in vCenter

Configure Networking

Setup Datastore

Start first VM

Configure and execute vMotion

Lab 3: Install Nexus 1000v

Install VSM

Configure VSM on NX-OS console

Install Plugin in to vCenter

Connect VSM to vCenter

Install VEM on ESX host

Configure Nexus 1000v

Lab Setup

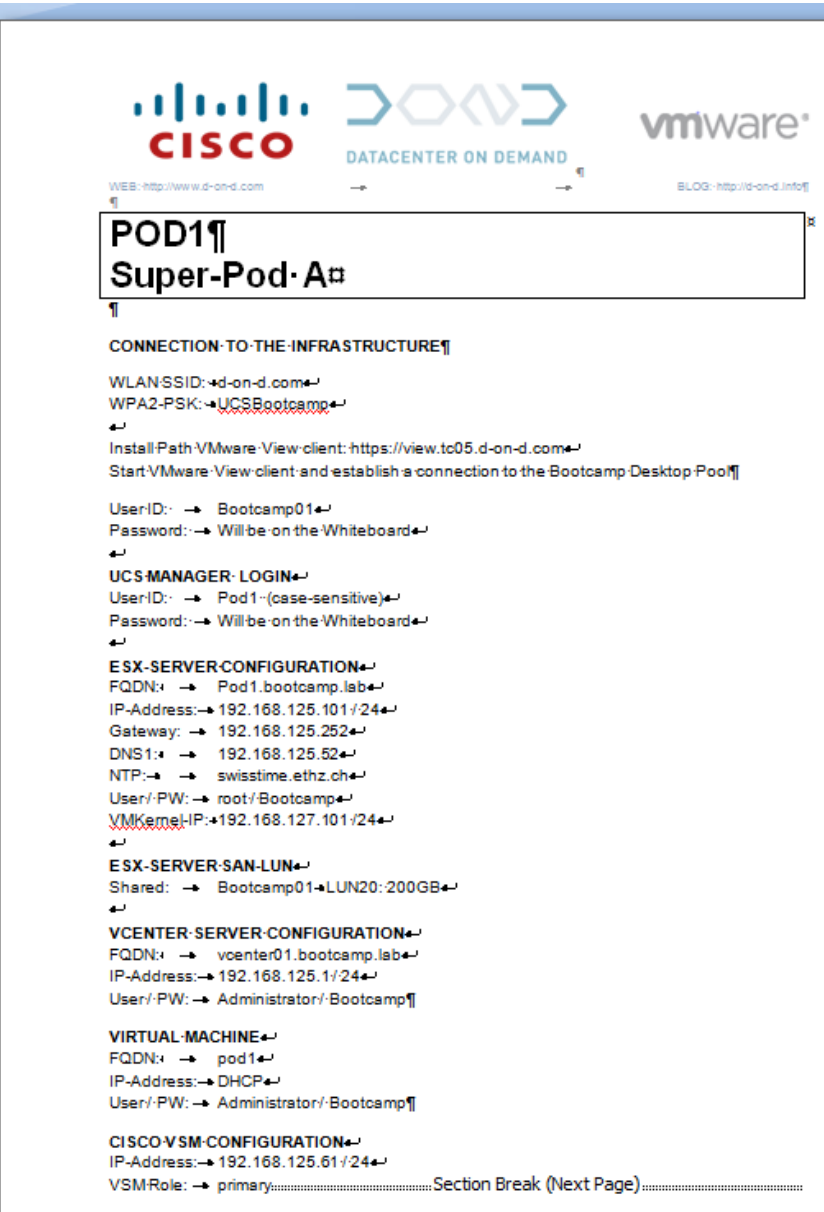

WEB: http://www.d-on-d.com → → BLOG: http://d-on-d.info/

POD1
Super-Pod A

CONNECTION TO THE INFRASTRUCTURE

WLAN SSID: d-on-d.com
WPA2-PSK: UCSBootcamp

Install Path VMware View client: https://view.tc05.d-on-d.com
Start VMware View client and establish a connection to the Bootcamp Desktop Pool

User ID: → Bootcamp01
Password: → Will be on the Whiteboard

UCS MANAGER LOGIN
User ID: → Pod1 (case-sensitive)
Password: → Will be on the Whiteboard

ESX-SERVER CONFIGURATION
FQDN: → Pod1.bootcamp.lab
IP-Address: → 192.168.125.101 / 24
Gateway: → 192.168.125.252
DNS1: → 192.168.125.52
NTP: → swisstime.ethz.ch
User / PW: → root / Bootcamp
VMKernel IP: 192.168.127.101 / 24

ESX-SERVER SAN-LUN
Shared: → Bootcamp01 LUN20: 200GB

VCENTER SERVER CONFIGURATION
FQDN: → vcenter01.bootcamp.lab
IP-Address: → 192.168.125.1 / 24
User / PW: → Administrator / Bootcamp

VIRTUAL MACHINE
FQDN: → pod1
IP-Address: → DHCP
User / PW: → Administrator / Bootcamp

CISCO VSM CONFIGURATION
IP-Address: → 192.168.125.61 / 24
VSM Role: → primary Section Break (Next Page)

WEB: http://www.d-on-d.com → → BLOG: http://d-on-d.info/

POD2
Super-Pod A

CONNECTION TO THE INFRASTRUCTURE

WLAN SSID: d-on-d.com
WPA2-PSK: UCSBootcamp

Install Path VMware View client: https://view.tc05.d-on-d.com
Start VMware View client and establish a connection to the Bootcamp Desktop Pool

User ID: → Bootcamp02
Password: → Will be on the Whiteboard

UCS MANAGER LOGIN
User ID: → Pod2 (case-sensitive)
Password: → Will be on the Whiteboard

ESX-SERVER CONFIGURATION
FQDN: → Pod2.bootcamp.lab
IP-Address: → 192.168.125.102 / 24
Gateway: → 192.168.125.252
DNS1: → 192.168.125.52
NTP: → swisstime.ethz.ch
User / PW: → root / Bootcamp
VMKernel IP: 192.168.127.102 / 24

ESX-SERVER SAN-LUN
Shared: → Bootcamp01 LUN20: 200GB

VCENTER SERVER CONFIGURATION
FQDN: → vcenter01.bootcamp.lab
IP-Address: → 192.168.125.1 / 24
User / PW: → Administrator / Bootcamp

VIRTUAL MACHINE
FQDN: → pod2
IP-Address: → DHCP
User / PW: → Administrator / Bootcamp

CISCO VSM CONFIGURATION
IP-Address: → 192.168.125.62 / 24
VSM Role: → standby Section Break (Next Page)


CISCO