

TOMORROW starts here.



Cisco *live!*

Deployment Best Practices for vSphere and Hyper-V on UCS

BRKVIR-2640

Doron Chosnek

Technical Solutions Architect

Introductions



- VMware/Hyper-V server virtualisation
- Desktop virtualisation
- Cisco UCS
- Area of expertise
 - Networking
 - Storage
 - Compute
- Job function
 - Architecture
 - Operations
 - Management

What We Will Cover

Deployment Best Practices for vSphere and Hyper-V on UCS

Plug-ins



Networking



Compute



Rapid Deployment



Graphics






Plug-ins for vSphere and Hyper-V



Cisco UCS vCenter Plugin

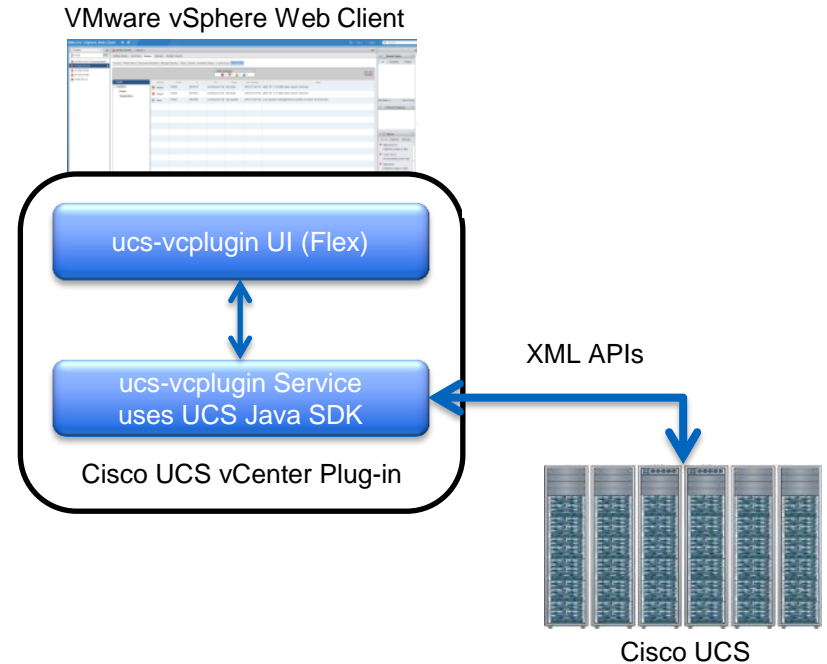
Version 0.9(1) beta

- Extension for vSphere Web Client 5.1
- Requires vCenter 5.0 or higher
- Enables admins to view, manage and monitor UCS physical infrastructure
- Free at <http://developer.cisco.com/web/unifiedcomputing/vmware>



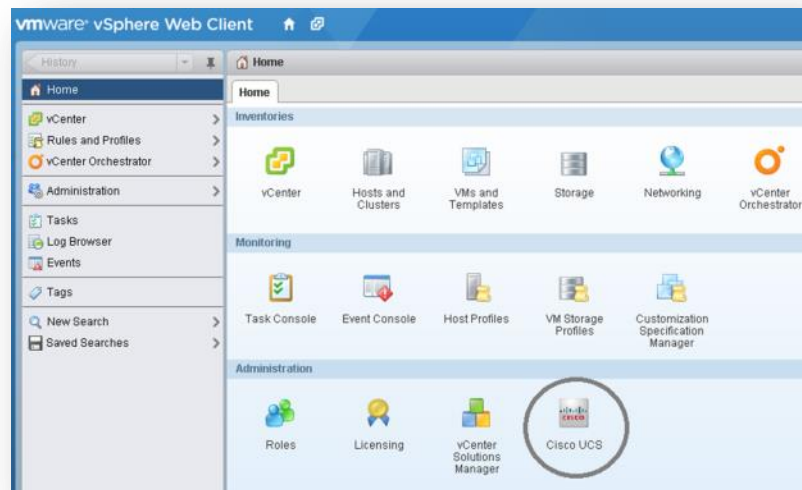
Cisco UCS vCenter Plug-in Architecture

- Two major components
 - UCS-vcplugin UI is Flex-based extension to VMware web client
 - UCS-plugin services uses XML API to get data from UCS
- v0.9(1) supports only one registration per UCS domain



Cisco UCS vCenter Plug-in Capabilities

- Physical hierarchy view
- Health of servers and infrastructure
- Maps physical servers to ESX hosts
- For each ESX host:
 - Inventory
 - Firmware
 - Faults
 - Power/temp statistics
- KVM launch
- Toggle blue locator LED



UCS Domains Shown in Inventory



The screenshot displays the VMware vSphere Web Client interface. On the left, the 'Inventory Lists' section is expanded to show 'Cisco UCS Domains' with a count of 3. A purple arrow points to this entry. The main content area shows a 'Getting Started' guide for vCenter, including a 'What is vCenter?' section and a list of steps to get started with virtual infrastructure.

vmware vSphere Web Client

Home vCenter Home

vCenter

- vCenter Home
- Inventory Trees**
 - Hosts and Clusters
 - VMs and Templates
 - Storage
 - Networking
- Inventory Lists**
 - vCenter Servers 1
 - Datacenters 2
 - Hosts 5
 - Clusters 0
 - Resource Pools 0
 - Datastores 7
 - Datastore Clusters 0
 - Standard Networks 1
 - Distributed Switches 0
 - Cisco UCS Domains**
 - UCS Domains 3

Getting Started

What is vCenter?

The vCenter inventory is where you find all the objects associated with vCenter Server systems, such as datacenters, hosts, clusters, networking, storage, and virtual machines.

The Inventory Lists allow you to view an aggregated list of these objects across vCenter Server systems. These flat lists enable easier batch operations.

The inventory tree is still available under Inventory Trees. Here objects are listed hierarchically in four categories: Hosts and Clusters, VMs and Templates, Storage, and Networking.

There are three steps to get started with your virtual infrastructure:

1. Create a datacenter
2. Add hosts to the datacenter
3. Create virtual machines on the hosts

Select an Inventory item

Explore Further

- What is vCenter Server?
- Networking in vSphere
- Storage in vSphere
- What are Inventory Views?

UCS Domain Inventory

Overview of domain components available



vmware vSphere Web Client

UCS Domains

- savbu-tpi01
 - Chassis 1
 - Rack Mounts 1
 - Fabric Interconnects 2**

Fabric Interconnects

- switch-A
- switch-B

switch-A Actions

Summary Related Objects

Dn sys/switch-A
Ucs savbu-tpi01

Fan


Name	Operability	Performan...	Power	Temperature
fan-1	operable	not-suppor	on	ok
fan-2	operable	not-suppor	on	ok

PSU

Name	Model	Operability	Performance	Power	Temperature	Voltage
psu-1		unknown	unknown	offduty	ok	unknown
psu-2	N10-PAC1-	operable	unknown	on	unknown	unknown

UCS Chassis View

ESX and non-ESX blades shown



Chassis details



chassis-1 Actions

Summary Related Objects

Name chassis-1
Dn sys/chassis-1
UCS Domain savbu-tpi01

Fan

Module	Name	Operability	Performance	Power	Temperature
Module-1	fan-1	operable	ok	on	ok
Module-1	fan-2	operable	ok	on	ok

PSU

Name	Model	Operability	Performance	Power	Temperature	Voltage
psu-1	N20-PAC5-2500W	operable	ok	on	ok	ok
psu-2	N20-PAC5-2500W	unknown	unknown	off	ok	unknown

IO Module

Id	Vendor	Model	If Role	If Type	Overall Status	Type	...
IO Module-1							
IO Module	Cisco Systems I	N20-I6584			operable		
Backplane Port							
Backplane port 1/2			server	physical	up	lan	dce

ESX Server Extension

Summary view

- Service profile name
- Server location
- Hardware inventory

The screenshot displays the VMware vSphere Web Client interface. The main content area shows the Summary view for a server extension with IP address 10.105.219.90. The summary card includes a server icon and the following details:

- Model: Cisco Systems Inc N20-B6620-1
- Processor Type: Intel(R) Xeon(R) CPU E5520 @ 2.27GHz
- Logical Processors: 16
- NICs: 2
- Virtual Machines: 5
- State: Connected
- Uptime: 123 days

A detailed configuration panel for the Cisco UCS server is highlighted with a purple border. It contains the following information:

UCS	Savbu-tpi01
Server Location	1/5
UUID	E9e92891-a54d-11de-be99-000bab01c0fb
Service Profile	Org-rootfs-Blade5
Serial	QC1133701J7
Overall Status	Ok

Below the configuration panel, there are buttons for "Launch KVM", "Launch UCSM", and "Turn Locator LED On". The Hardware section below shows:

- Manufacturer: Cisco Systems Inc
- Model: N20-B6620-1
- CPU: 8 CPUs x 2.26 GHz
- Memory: 24,304 MB / 28,605 MB

The right-hand side of the interface shows the Configuration and Tags sections. The Configuration section includes:

- ESX Version: VMware ESXi, 5.0
- Image Profile: ESXi-5.0.0-2012
- vMotion Enabled: No
- vSphere HA State: N/A
- Host Configured for FT: No
- EVC Mode: Disabled

The Tags section is currently empty.

ESX Server Extension

- Manage view

The screenshot shows the VMware vSphere Web Client interface for a Cisco UCS chassis. The 'Manage' tab is active, and the 'Cisco UCS' sub-tab is selected. The 'Inventory' section is expanded, showing a table of hardware components. A purple box highlights the 'Running Version', 'Startup Version', and 'Backup Version' columns for the BIOS component.

Type	Name	Model	Running Version	Startup Version	Backup Version	Update Status	Activate Status
CIMC Controller	CIMC Controller	Cisco UCS B200 M1	2.0(1f)	2.0(1f)	2.0(0.130)	ready	ready
Adaptors	adaptor-1	Cisco UCS M71KR-Q	2.0(1f)	2.0(1f)	1.4(0.433p)	ready	ready
Storage Controllers	storage-SAS-1	SAS1064E PCI-Express Fu	01.28.03.00	01.28.03.00 06.2	N/A	N/A	ready
BIOS	BIOS	Cisco UCS B200 M1	S5500.1.3.1b.0.0511201016	S5500.1.3.1b.0.0	N/A	N/A	ready

Inventory and firmware summary information

ESX Server Extension

- Monitor view

The screenshot displays the VMware vSphere Web Client interface for a Cisco UCS chassis. The left navigation pane shows a tree view with categories like Top Level Objects, Virtual Machines, vApps, Datastores, Networks, Distributed Switches, Chassis, and Rack Mounts. The 'Faults' section is expanded, showing 'Statistics' with sub-items for 'Power' and 'Temperature'. The main content area features a 'Fault Summary' widget with icons for error (0), warning (0), and info (2) states. Below this is a table of faults:

Severity	Code	Id	Dn	Cause	Last Transition	Descr
(Warning)	F4525253	3499472	org-root/ls-Blade	named-policy-un	2012-03-16T14:...	Policy reference so/PolicyName 'default' does not resolve to named policy
(Info)	F0461	3513290	sys/chassis-1/bl	log-capacity	2012-03-16T18:...	Log capacity on Management Controller on server 1/5 is very-low
(Warning)	F4528596	7300365	org-root/ls-Blade	named-policy-un	2012-09-26T14:...	Policy reference maint/PolicyName 'TestQA' does not resolve to named policy

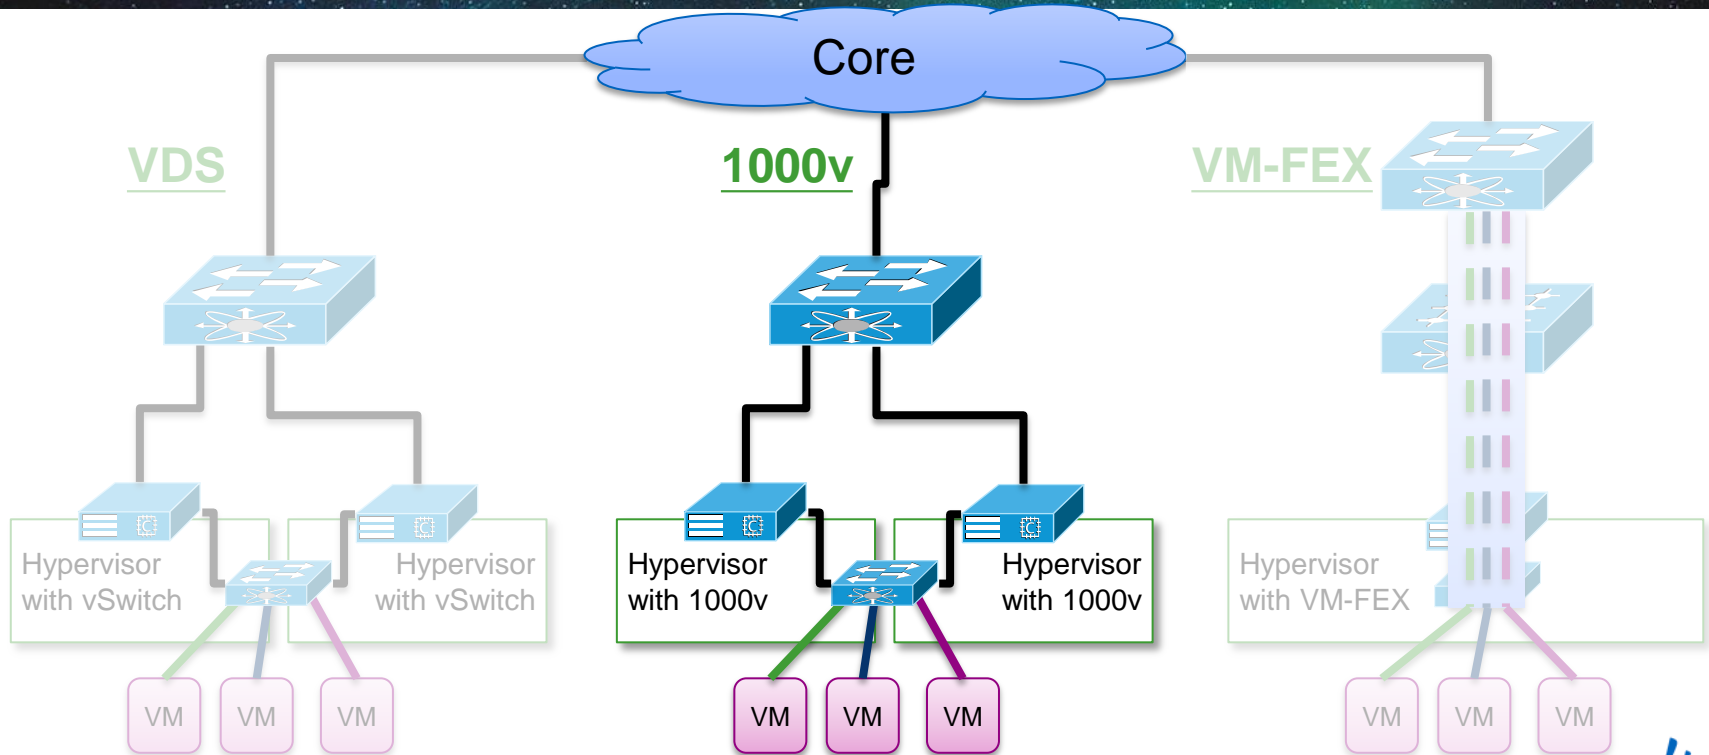
Faults, power statistics and temperature statistics



Networking

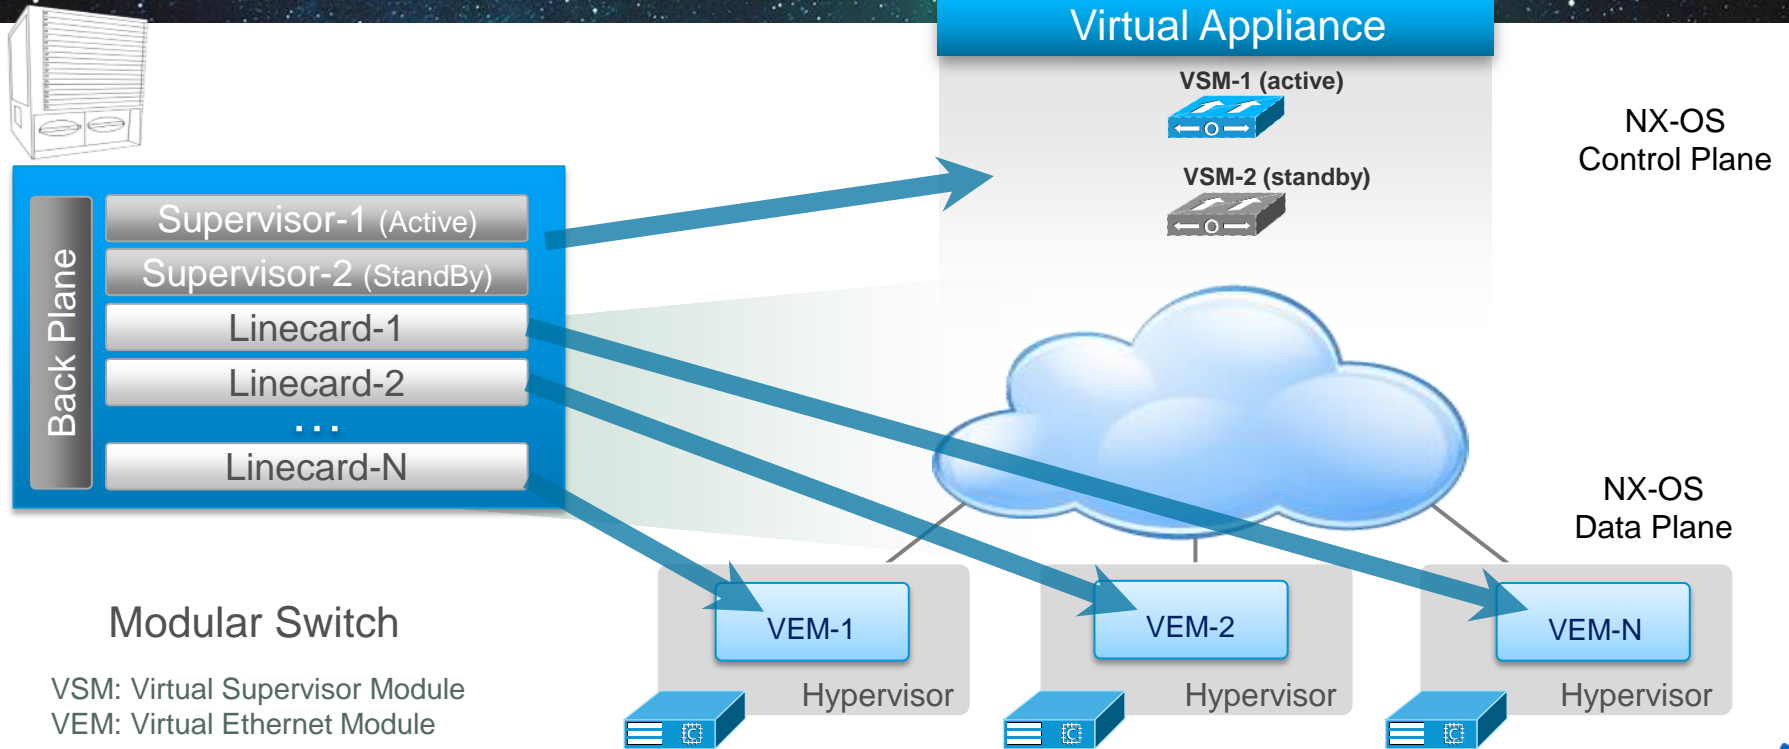
Three Ways to Provide IO to a VM

VMware VDS, Cisco 1000v, and Cisco VM-FEX



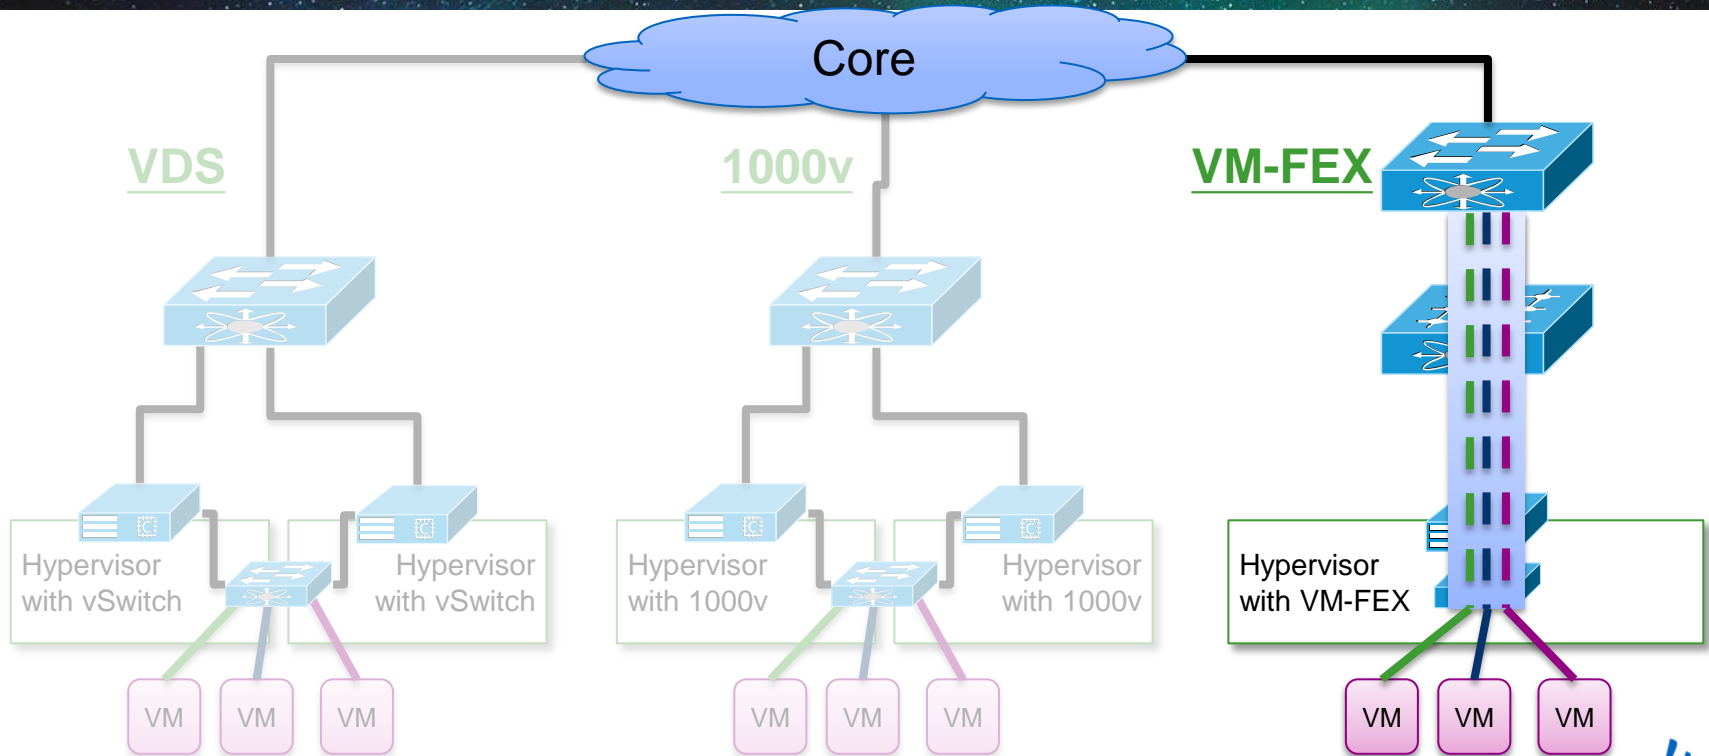
Nexus 1000V Architecture

A virtual modular switch with familiar NX-OS



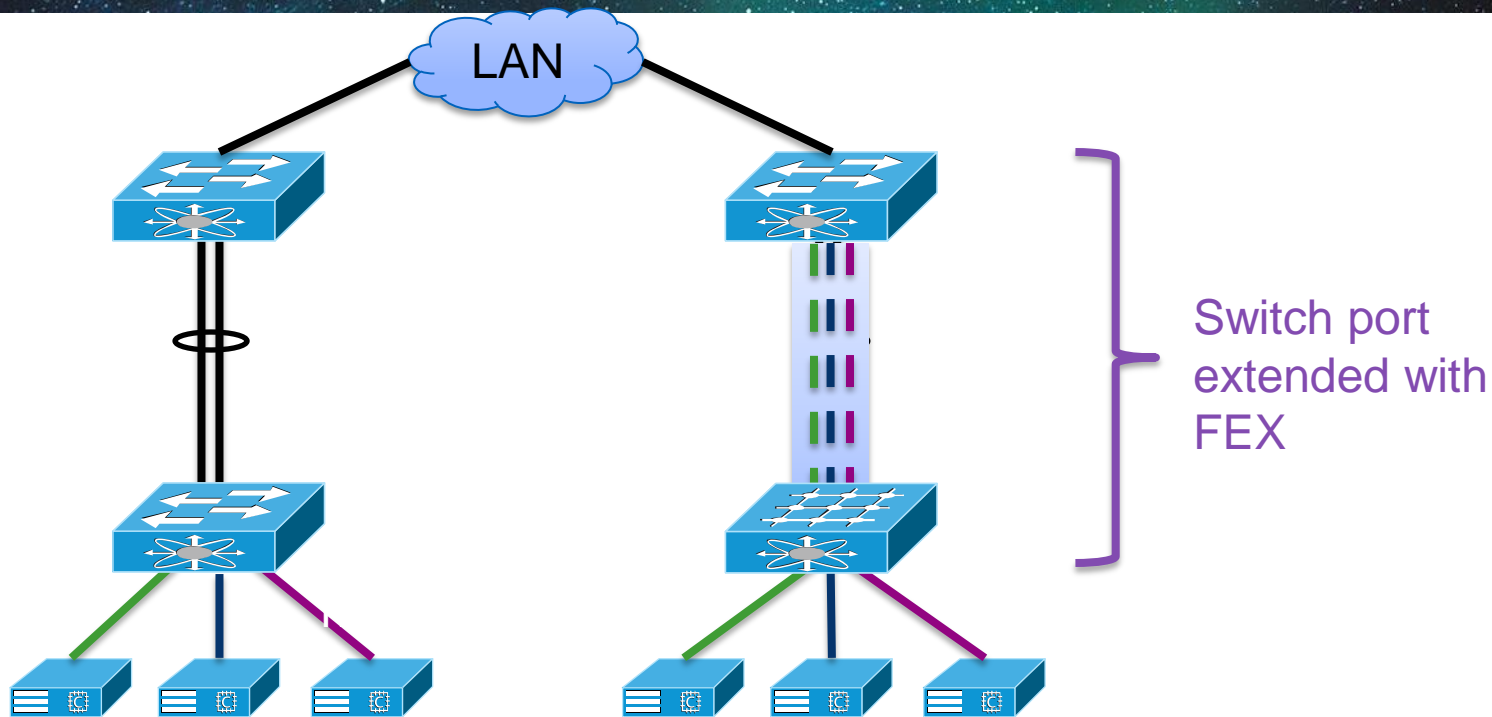
Three Ways to Provide IO to a VM

VMware VDS, Cisco 1000v, and Cisco VM-FEX



Fabric Extension (FEX)


Virtualising the network port



Nexus 5K Family with FEX

FEX puts switch ports closer to the load


Modular System




Distributed Modular System

- Nexus 2000 FEX is a Virtual Line Card to the Nexus 5000
- Nexus 5000 maintains all management & configuration
- No Spanning Tree between FEX & Nexus 5000

Distributed Modular System




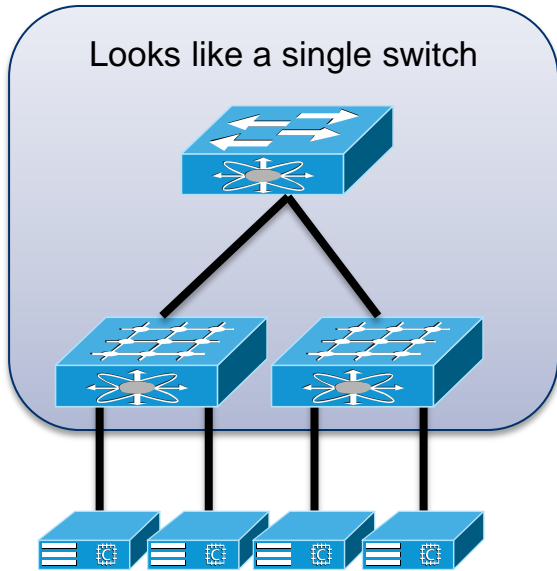

Distributed Modular System



20K customers over 5 years
25 million Nexus ports deployed

Cisco UCS Blade Chassis


FEX technology simplifies management



UCS 5108 blade chassis = 8 blades + 2 FEX


Cisco UCS Blade Chassis

Up to 160 blades can be part of one system



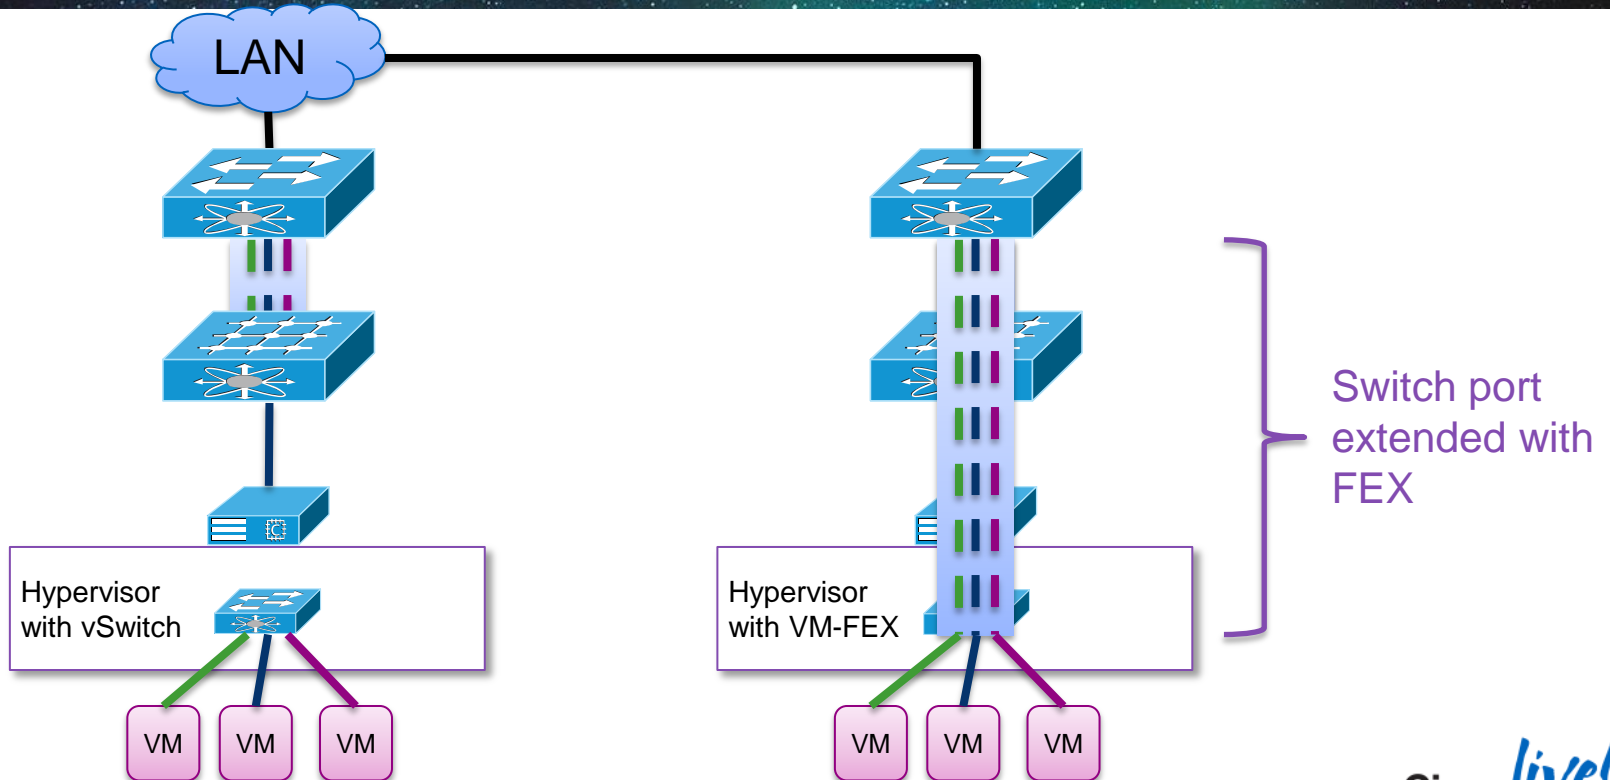
Key Architectural Component 1: VNTAG

Intra-chassis bus header



Extending FEX Architecture to VMs


FEX inside the server



Key Architectural Component 2: UCS VIC


UCS Virtual Interface Card family

- 256 PCIe devices
 - Devices can be vNICs or vHBAs
 - Each device has a corresponding switch interface
- Bandwidth 2x4x10 Gb
 - Uses 4x10 Ether Channel, HW 40Gb Capable
 - vNICs/vHBAs NOT limited to 10Gb
- PCIe Gen-2 x 16
- Mezzanine and PCIe



VM-FEX Modes of Operation

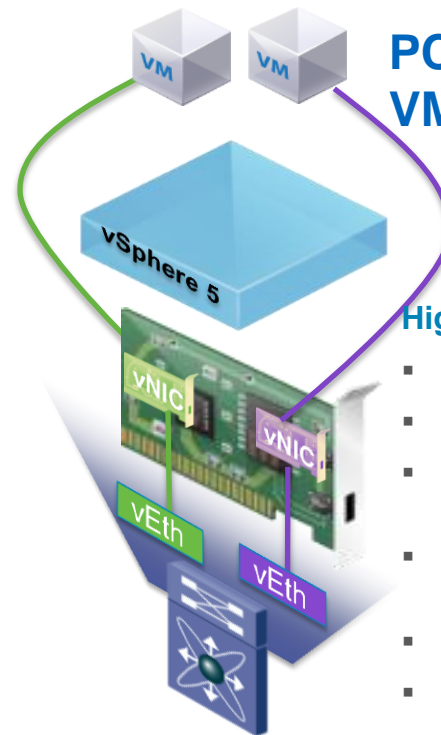
Enumeration vs. Hypervisor Bypass



Emulated Mode

Standard (Emulated) Mode

- Each VM gets a dedicated PCIe device
- ~12%-15% CPU performance improvement
- Appears as distributed virtual switch to hypervisor
- LiveMigration supported



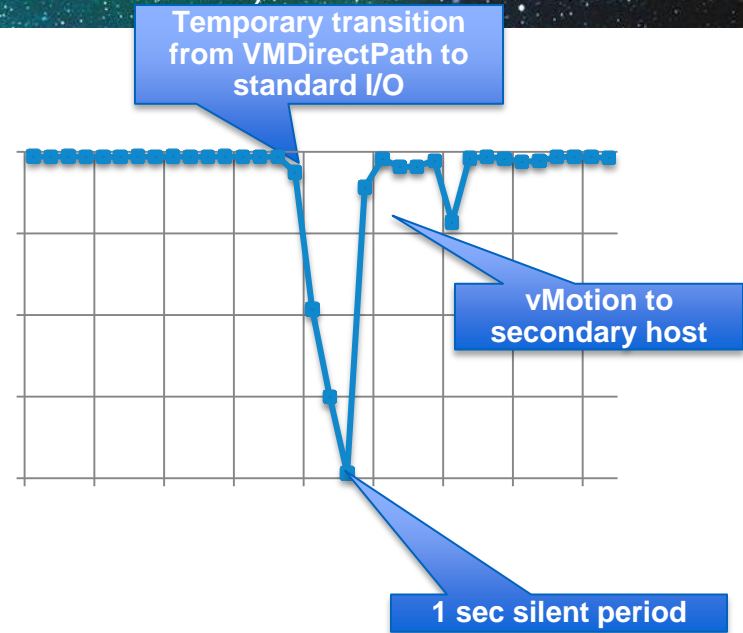

PCIe Pass-Thru or VMDirectPath

High Performance Mode

- Co-exists with Standard mode
- Bypasses Hypervisor layer
- ~30% improvement in I/O performance
- Appears as distributed virtual switch to hypervisor
- Supported with ESX 5.0 or later
- vMotion supported

VM-FEX Operational Model

vMotion with Hypervisor Bypass (VMDirectPath with VM-FEX)




- VM Sending TCP stream (1500MTU)
- UCS B200 M2 blades with UCS VIC card

VM-FEX Performance Review

Test of 10 VM's running HTTP and FTP Gets with IxLoad



Virtual Switch, CPU at ~65%



VM-FEX, CPU at ~ 37%



Compute

Processor Improvements


- More integrated components
 - Memory controllers
 - PCIe busses
- More QPI links
- More memory channels
- More memory
- More cores
- More cache



More **virtual machines** per host

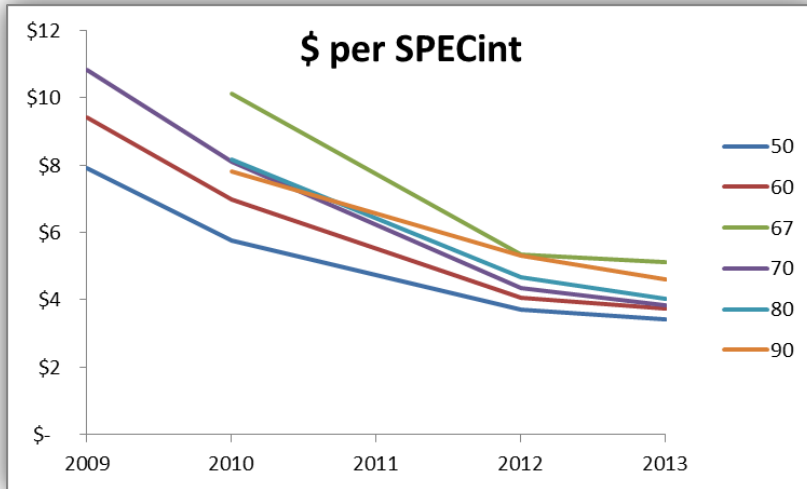
2-Socket CPU History

Performance goes up

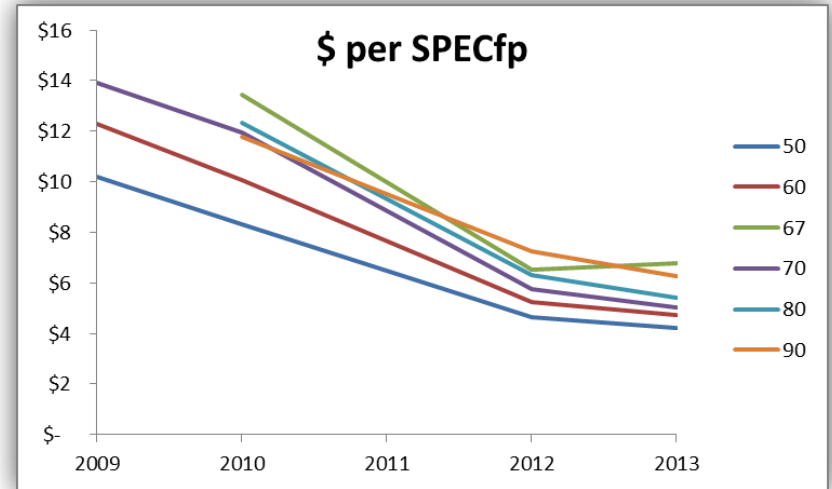


2-Socket CPU History

Price per performance comes down











Integer Performance



Floating Point Performance

40-60% reduction in CPU cost per workload

Top Bin E7 Versus Top Bin E5

	E7-2870	E5-2690
Power (W)	130 	135
Cores	10 	8
Cache	30M 	20M
Speed (GHz)	2.40	2.90 
QPI	6.4GT/s	8.0GT/s 
SPECint	267.5	348.5 
SPECfp	184.5	254.5 
Recommended Cost	\$4,227	\$2,057 

Higher performance at half the price. Which one do you prefer?

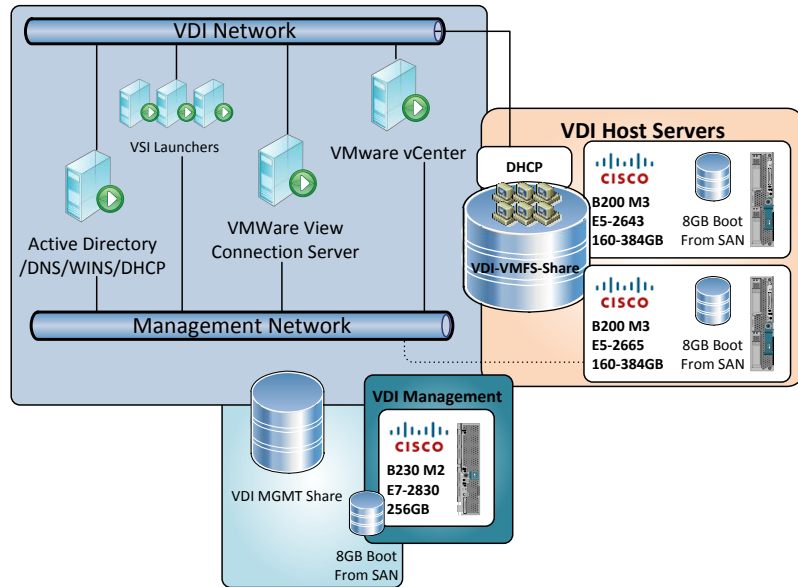
E7 Versus E5 Proc Comparison

To be fair...

- Most comparisons are not as extreme
- Most sizing exercises show E7 systems to be superior
- E5-2600 systems **substantially** less expensive

Lessons Learned From VDI Testing

Compute trends relevant to general server virtualisation



Logical test environment






- Various UCS B200 M3 configurations
- Login Consultants' Virtual Session Indexer (Login VSI) 3.6.1 benchmark
- Login VSI's Medium with Flash workload
- VMware View 5.1.1
- Microsoft Windows 7 SP1 32-bit virtual desktops

Configuration Settings for Tests

Configuration	Setting(s)
Login VSI Configuration Settings	<ul style="list-style-type: none">• Medium with Flash workload generator• 4,000ms response cut off
Windows 7 Configuration	<ul style="list-style-type: none">• 1.5GB memory for all tests• 32-bit Windows 7 SP1 and Windows updates through September 1, 2012
ESX Host Configuration	<ul style="list-style-type: none">• Power management set to High Performance in vSphere• BIOS settings have C1E disabled in UCSM policy• ESXi-5.0.0-623860-custom (Cisco specific build of ESXi with drivers)
VM Configuration	<ul style="list-style-type: none">• 1vCPU and 2vCPU configurations
View Configuration	<ul style="list-style-type: none">• Linked clones• View Optimisation* registry changes on all Virtual desktops• Did not use profile management• Did not use host Cache for View

<http://www.vmware.com/files/pdf/vmware-view-optimizationguidewindows7-en.pdf>



Processor Comparison

	Intel E5-2643	Intel E5-2665
Number of Cores	4	8 
Clock Speed	3.3 GHz 	2.4 GHz
Max Turbo Frequency	3.5 GHz	3.1 GHz
Max TDP	130W	115W
Cache	20M	20M
Recommended Customer Price (Tray)	\$885 	\$1440
SPEC CINT2006 Rate	187.5	305 
SPEC CFP2006 Rate	167.5	233.5
SPEC Blend/Core	44.38 	33.6

Question 1



Which is more important, core count or core speed?

- E5-2665 achieves 60-72% better density
- E5-2643 has lower latency at small desktop counts
- When density is the goal, core count is more important



Question 2

How well does a system scale with multi-vCPU VMs?




- Going from 1vCPU to 2vCPU yielded **66-72% as many desktops**
- Increasing vCPU count impacts scale, even with **the same workload**
- Advice: don't give a VM more vCPUs unless it needs it

Question 2 Continued

How well does a system scale with multi-vCPU VMs?


- Purpose of multi vCPU machine is better performance
- Graph shows better performance (lower latency)
- Performance benefits taper off after 45 desktops



Question 3

How does memory speed affect scalability?


- At 1vCPU (blue lines), only 4% difference between 1600MHz and 1066MHz
- At 2vCPU (red lines), only 3% difference between 1600MHz and 1066MHz
- In both cases, the difference is negligible



Question 3 Continued

How does memory speed affect scalability?


- At 2vCPU, performance difference is negligible
- At 1vCPU, E5-2643 saw huge 23% performance difference
- What is the cause? A higher VM/core ratio?



Question 4

How does memory capacity affect VDI scalability

- At 160GB system memory
 - Theoretical limit 98 desktops
 - Actual test reached 106 desktops
 - VMware memory overcommit gave ~12%
- At 256GB system memory
 - Theoretical limit 157 desktops
 - Actual test reached 130 desktops
 - System ran out of CPU
- Lesson: understand system constraints



Theoretical

98 desktops X (1.5GB + 130MB overhead) = **160GB**

157 desktops X (1.5GB + 130MB overhead) = **256GB**

More VDI Testing Analysis

11-part blog series at blogs.cisco.com

The screenshot shows a Cisco blog post interface. At the top, there's a blue header with the Cisco logo and 'Cisco Blogs'. Below that is a search bar and a 'Data Center and Cloud' banner. The main content area features a profile picture of Tony Paikeday, the article title 'VDI - The Questions You Didn't Ask (But Really Should)', and the date 'January 25, 2013 at 8:06 am PST'. The article text discusses the challenges of finding VDI content and the importance of intelligent decisions around CPU, memory, and network resources. On the left side, there are social sharing buttons for Twitter (75), Facebook (60), LinkedIn (90), and Google+ (2). On the right side, there are options to 'Subscribe Now', 'RSS Feed', and 'Follow Us'. At the bottom, there's a 'Cisco Social Rewards' section and a 'Blogs at a Glance' table.

75
Tweet
60
Like
90
Share
2
+1
Kindle

Cisco Blog > Data Center and Cloud
VDI - The Questions You Didn't Ask (But Really Should)
Tony Paikeday | January 25, 2013 at 8:06 am PST
(1 Comment)

There's no shortage of content out there (a quick Google search easily confirms this) when it comes to looking for vendor-originated material touting the latest server performance benchmarks for hosted virtual desktops. Being part of that community, I'm pretty sure I have my fingerprints on more than one such piece of collateral - and I'm constantly reminded of this, when we run into questions along the lines of "yeah, [xxx] desktops on a blade is great, but c'mon, you and I both know we'd never do that in practice". It's a balancing act of demonstrating solution performance, intersected with the practical reality of what IT managers would reasonably support in a production environment.

So what really matters? If I'm implementing VDI for the 1st time, and I'm trying to make intelligent decisions around CPU, memory speed, IOPS, etc., where do I go? VDI is unique in its consumption of compute, storage and network resources, when compared to other workloads hosted in the data center. Much of the performance benchmarking info put out by server manufacturers is not specific to VDI performance, or how user experience might be impacted by simple decisions like choice of clock speed or # of vCPU.

Thankfully, there are folks in my company that care a LOT about such questions. So much so, that a small, VDI-proficient group of them took it upon themselves to design and build an in-house lab

BRKVIR-2640

Subscribe Now
Enter email address +
RSS Feed Add
Follow Us
Cisco Social Rewards
Get credit for all your Cisco site participation with Social Rewards
Join Today Learn More
Already a rewards member? Log In
Blogs at a Glance
Most Recent Most Commented Popular

<http://blogs.cisco.com/datacenter/vdi-the-questions-you-didnt-ask-but-really-should/>



Rapid Deployment and Automation



Stateless Computing

What does it mean to you?



Stateless: UCS Service Profiles

- RAID settings
- Disk scrub actions
- Number of vHBAs
- HBA WWN assignments
- FC Boot Parameters
- HBA firmware
- FC Fabric assignments for HBAs
- QoS settings
- Border port assignment per vNIC
- NIC Transmit/Receive Rate Limiting
- VLAN assignments for NICs
- VLAN tagging config for NICs
- Number of vNICs
- PXE settings\
- NIC firmware
- Advanced feature settings
- Remote KVM IP settings
- Call Home behaviour
- Remote KVM firmware
- Server UUID
- Serial over LAN settings
- Boot order
- IPMI settings
- BIOS scrub actions
- BIOS firmware
- BIOS Settings



Stateless: UCS Service Profiles

- NIC MACs
- HBA WWNs
- Server UUID
- VLAN Assignments
- VLAN Tagging
- FC Fabrics Assignments
- FC Boot Parameters
- Number of vNICs
- Boot order
- PXE settings
- IPMI Settings
- Number of vHBAs
- QoS
- Call Home
- Template Association
- Org & Sub Org Assoc.
- Server Pool Association
- Statistic Thresholds
- BIOS scrub actions
- Disk scrub actions
- BIOS firmware
- Adapter firmware
- BMC firmware
- RAID settings
- Advanced NIC settings
- Serial over LAN settings
- BIOS Settings



Creating and Moving Service Profiles

- **State** is decoupled from **hardware**
- Easy to redeploy HW for other applications
- Dynamic provisioning
 - Speed/agility
 - Consistency
- Simplified management



Service Profile: **ESX_5**
Network1: www_prod
Network1 QoS: Gold
MAC : 08:00:69:10:78:ED
Boot Order: LOCAL
FW: WebServerBundle



Quickly Recover From Hypervisor Host Failure

Local boot media breaks stateless computing

- No local disk
- No USB flash drive
- No SD card



Local boot media is something an admin has to touch or move

Quickly Recover From Hypervisor Host Failure

Boot from SAN

Hypervisor Host
(WWN set by UCSM)



UCSM moves the WWN




Hypervisor Host
(WWN set by UCSM)

RESULT

WWN does not change,
zoning does not change,
LUN masking does not change

Quickly Recover From ESX Host Failure

VMware Auto Deploy (boot from Ethernet)



DHCP Server

set option 66 NextServer = IP of TFTP server
set option 67 FileName = name of PXE boot file
provides host with IP & FileName for boot

TFTP Server

gPXE boot image downloaded from vCenter
provides host gPXE boot image to perform HTTP boot
from Auto Deploy server

Auto Deploy Server

rules engine and profiles configured with PowerCLI
uses host MAC, IP, etc. in rules engine to assign a host profile



What happens when this host fails?

Use UCS Manager to move the MAC to spare server.

Auto Deploy Methods in vSphere 5.1

All three modes work well with UCS stateless computing

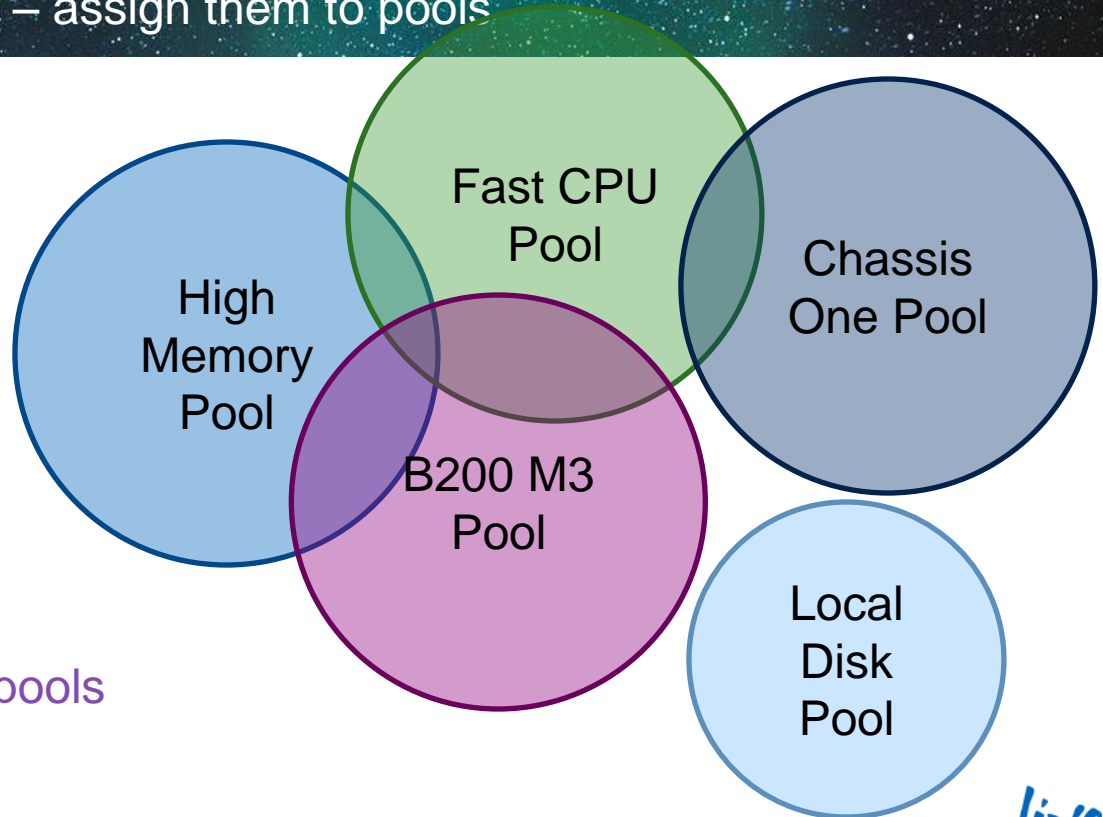
- Stateless
 - Host boot order: PXE only
 - Loads image and configuration from AD server on every reboot
- Stateless caching
 - Host boot order: PXE then local disk
 - Caches ESXi image to local disk; used only if PXE fails
- Stateful Install
 - Host boot order: local disk then PXE
 - Uses AD to install to disk

Using Server Pools

Don't assign profiles to servers – assign them to pools

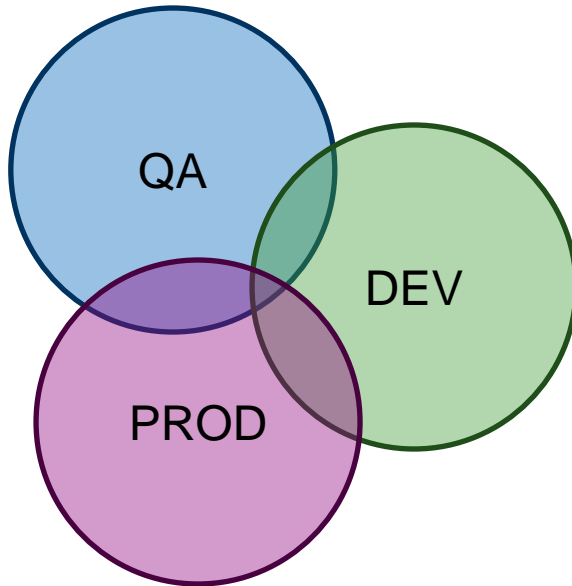
- Pools can be static/manual
- Pools can be dynamic
- Pools can overlap

The true power of pools

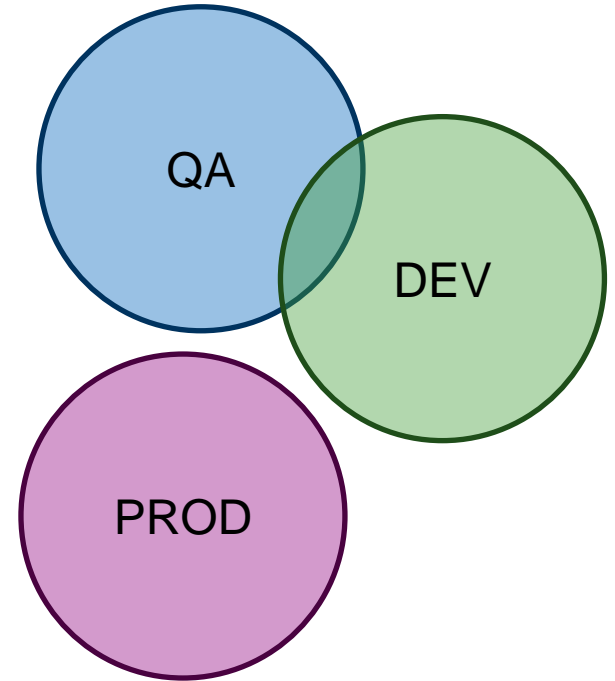


Beware That Pools Can Overlap

Don't accidentally take resources away from one pool



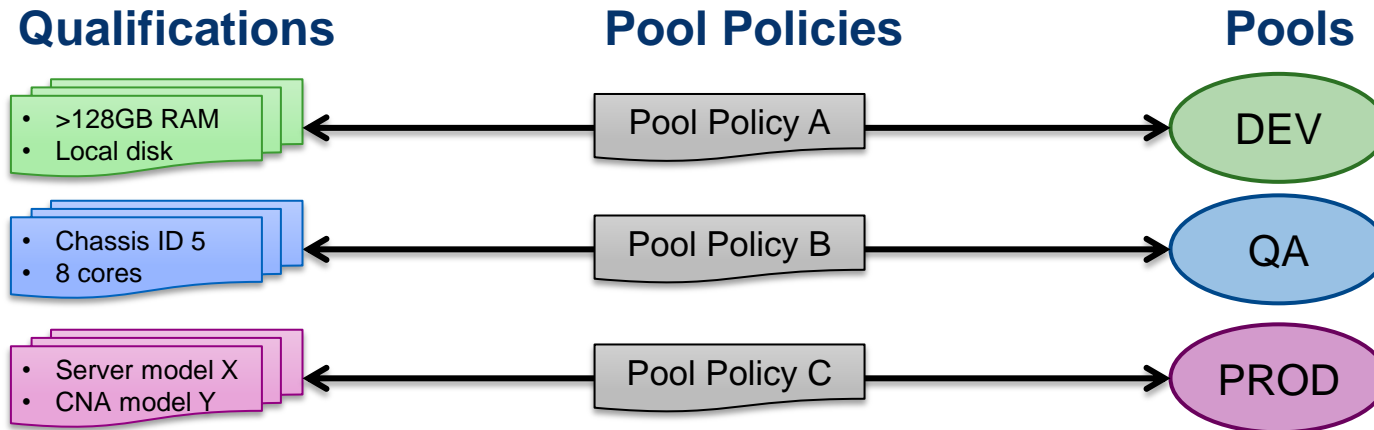
-OR-



Creating Dynamic Pools


Relationship between pools and policies

1. Create **Server Pool Policy Qualifications**
2. Create empty **Server Pools**
3. Link each **Server Pool** to **Qualification(s)** with a **Pool Policy**



Dynamic Pool Advanced Scenario 1

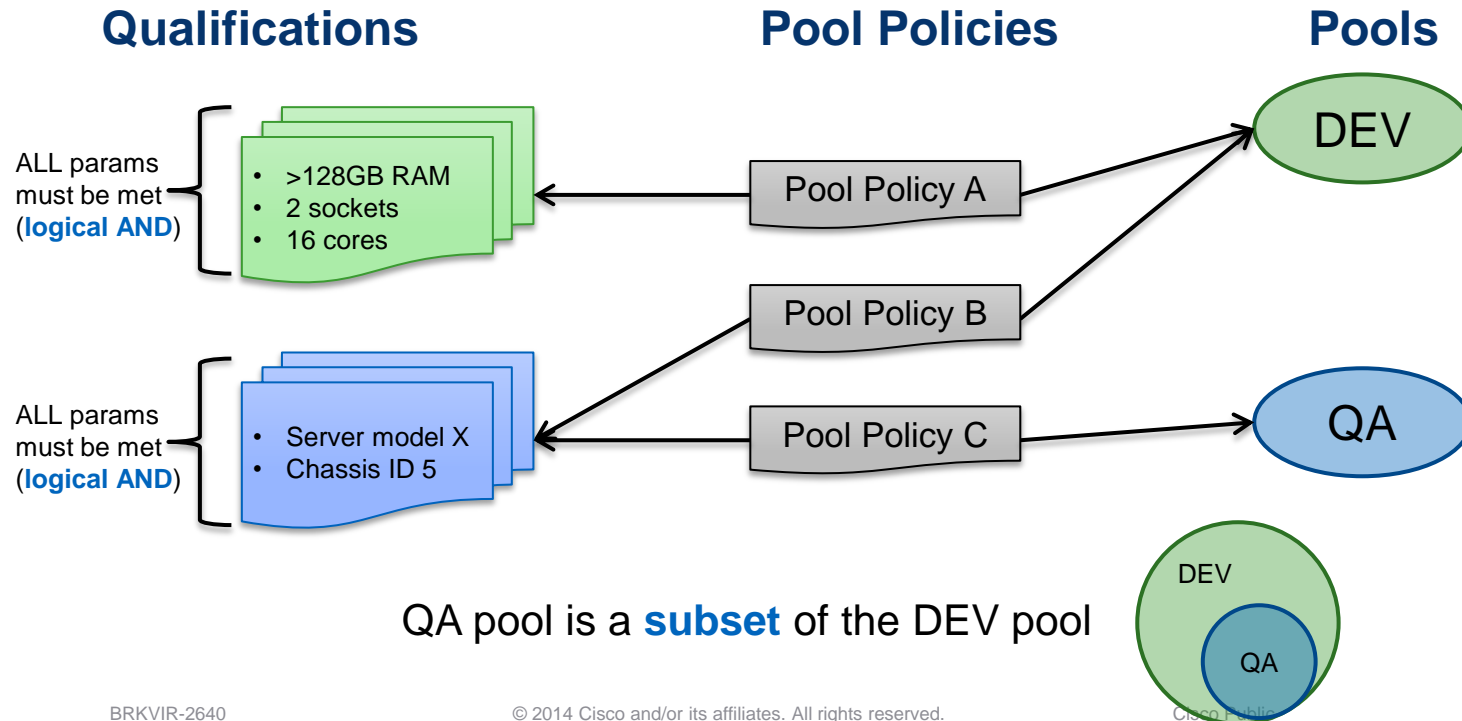
Pools can be fed by more than one qualification policy



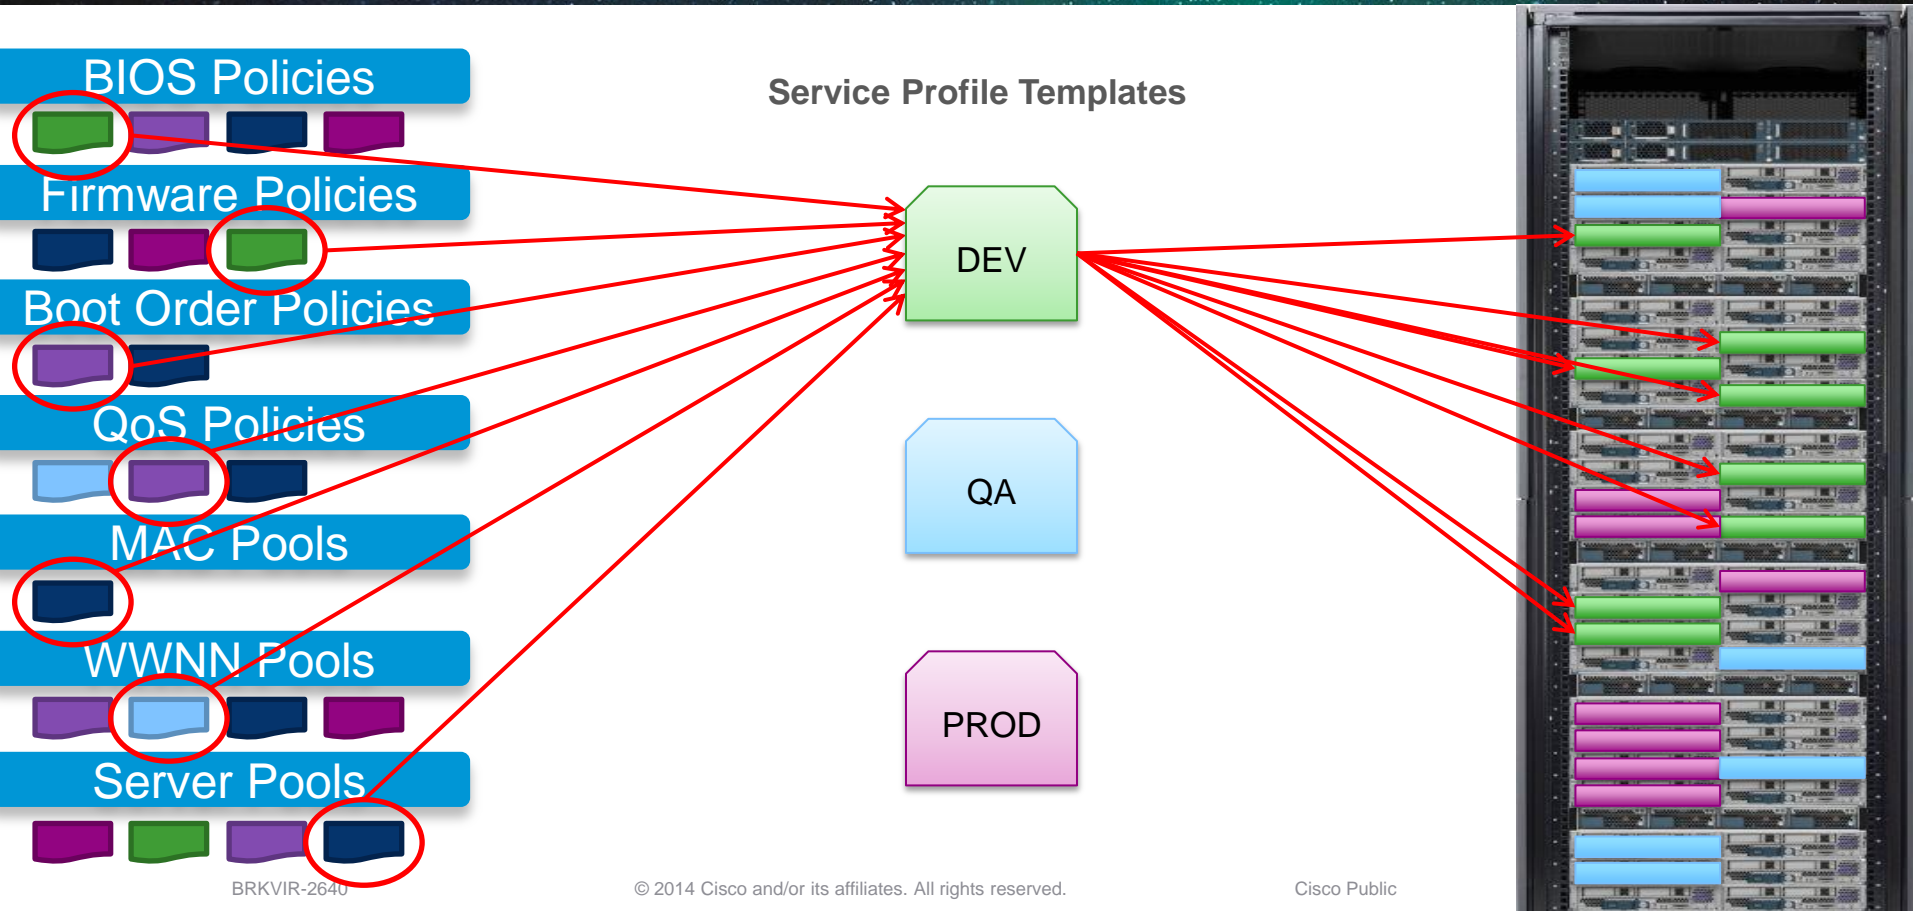
IF (Qual A is TRUE) OR (Qual B is TRUE) THEN add server to DEV pool

Dynamic Pool Advanced Scenario 2

One pool can be a subset of another

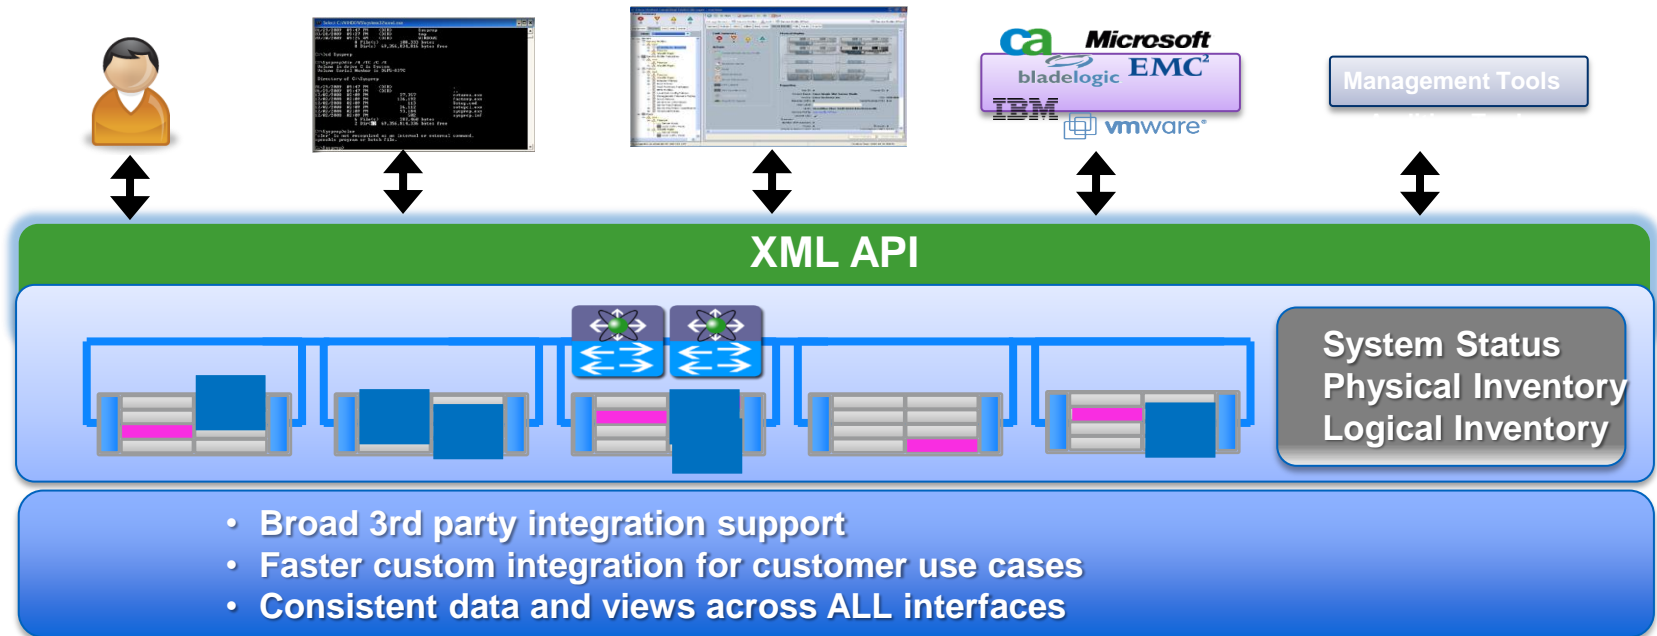


Templates Promote Agility, Consistency



Programmatic Infrastructure


- Comprehensive XML API, standards-based interfaces
- Bi-Directional access to physical & logical internals



UCS XML API Overview

UCS XML API Features

- Communicates over HTTP / HTTPS
- XML Based, Transactional
- XML Transactions are Order Agnostic
- Standard Request / Response cycle
- Role Based Authentication
- Object Model Hierarchy
- Built-in Object Browser
- Published XML Schema
- Java Doc Style documentation
- High Availability
- Event Stream

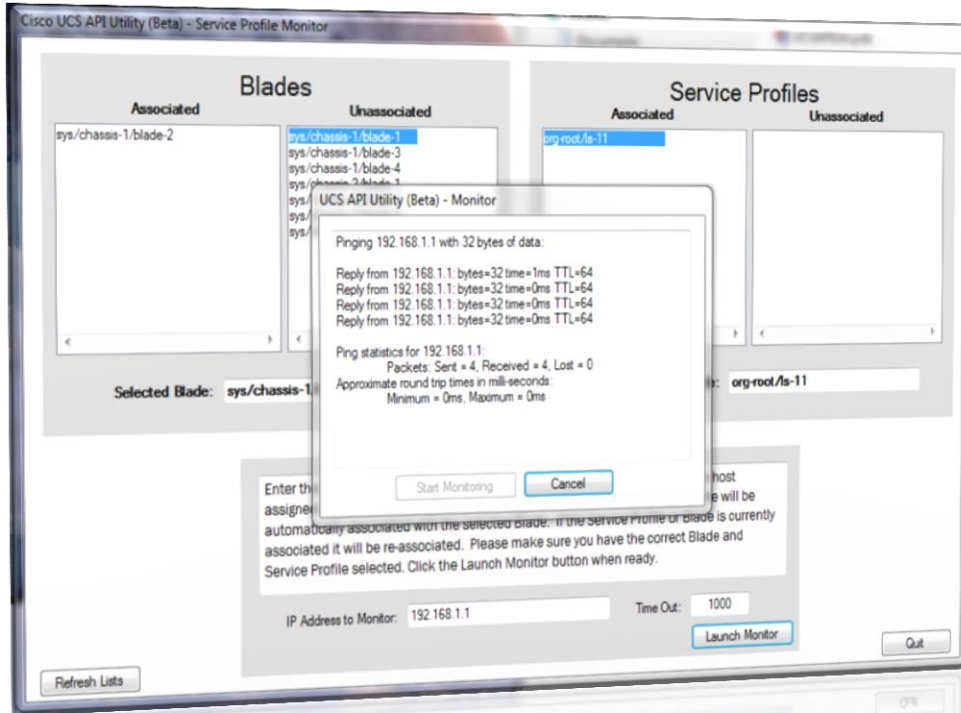


A Very Simple XML Query

pod	cookie	response	class Id	ack	cause	code	created	descr	dn	highest Severity	id	last Transition	lc	occur	orig Severity	prev Severity	rule	severity	tags	type	fault Inst
PFUCSPOD 1																					
	128076624071 16218753a36- 4485-86d2- b3c235c5e2b	yes	faultInst																		
PHUCSPOD 1																					
	128076549877 794a5d-d94c- 4a7b-3d05- 0ba73ca5910	yes	faultInst	no	equipment- inoperable	F0185	2010-07- 30T16:14:24.6 47	DIMM 1/7 on server 10/7 operability: inoperable	sys/chassis- 10/blade- 7/blade/mem- ray-1/mem- 7/fault-F0185	major	715935	2010-07- 30T16:14:24.6 47		1	major	major	memory-uni- inoperable	major	server	equipment	
PHUCSPOD 2																					
	128076623678 4c6328-f146- 4622-af13- 17c545daa15	yes	faultInst	no	thermal problem	F0411	2010-07- 29T18:38:30.3 98	Temperature on chassis 6 is upper non- recoverable	sys/chassis- 6/fault-F0411	critical	601433	2010-07- 29T18:38:30.3 98		1	critical	critical	equipment- chassis- thermal- threshold non- recoverable	critical	server	environmental	
	128076623678 4c6328-f146- 4622-af13- 17c545daa15	yes	faultInst	no	configuration- failure	F0327	2010-07- 29T11:04:21.2 78	Service profile SAP:PNFAS0 606-1 configuration failed due to compute- unavailable,ins ufficient resources	org-root/org- profile/ SAP:PNFAS0 606-1/fault- F0327	major	601488	2010-07- 29T11:04:21.2 78		1	major	major	is-server- config-failure	major	server	server	
	128076623678 4c6328-f146- 4622-af13- 17c545daa15	yes	faultInst	no	execute-peer- failed	F77959	2010-07- 28T11:05:03.7 62	[PSM STAGE REMOTE ERROR]: Result: understand- fail Code: ERR-0505- IBMC-For- retrieval-error Message: Could not get Fru from 7929607. dn=fabric/ser- v/chassis- 6/kit-7 (sametime:Fab ricComputeSto rEgIdentifyEx ecutePeer)	condition	601565	2010-08- 02T12:20:16.2 61		512	condition	cleared	fsm-identify- remote-iv	condition		fsm		
	128076623678 4c6328-f146- 4622-af13- 17c545daa15	yes	faultInst	no	fsm-failed	F99909	2010-07- 29T11:21:26.1 98	[PDM] FAILED to identify a blade in P/7 (PDM summary: FabricComp onditionIdentify)	fabric/serve/c hassis-6/kit- 7/fault- F99909	critical	616634	2010-07- 29T11:21:26.1 98		1	critical	critical	fsm-identify- fsm-fail	critical		fsm	
PHUCSPOD 3																					
	128076623676 c5684e-c243- 4ec1-93d0- 0c288cb9a06	yes	faultInst																		
PHUCSPOD 4																					
	128076623779 8a8b135-c81a- 4091-007a- 5074904a1299	yes	faultInst	no	equipment- inoperable	F0185	2010-07- 19T03:26:31.4 95	DIMM 1/6 on server 6/2 operability: inoperable	sys/chassis- server/6/2/ blade- operability: ray-1/mem- 6/fault-F0185	major	1306742	2010-07- 19T03:26:31.4 95		1	major	major	memory-uni- inoperable	major	server	equipment	
PHUCSPOD 5																					
	128076623770 Edmvtz 0387	yes	faultInst																		

Write Your Own App

Expose only features you want exposed OR create new functionality



Purpose-built console for Service Profile monitoring

Cisco Developer Network

The screenshot displays the Cisco Developer Network website. The main header includes the Cisco logo and the text "Cisco Developer Network". Below the header is a navigation menu with links for Home, Membership, Technologies, Community, Technology Partners, News & Events, and My C. The main content area is titled "UCS Manager" and has sub-links for UCS Manager, Forums, Blogs, and Documentation. The "UCS Manager" section contains a paragraph describing the Cisco Unified Computing System (UCS) and its XML API, followed by a section titled "What Is It?" with an "Overview" sub-section. Below this are three columns: "How Do I Get Started?" with a "Getting Started" sub-section, "What Resources Are Available?" with a "Resources" sub-section, "Developer Forums" with a "View All Forums" link, "Recent Bloggers" with a table of bloggers, and "Recent Downloads" with a list of download links.

UCS Manager

The Cisco Unified Computing System (UCS) includes an innovative XML API which offers you a programmatic way to integrate or interact with any of the over 9,000 managed objects in UCS. Managed objects are abstractions of UCS physical and logical components such as adaptors, chassis, blade servers, and fabric interconnects.

Developers can use any programming language to generate XML documents containing UCS API methods. The complete and standard structure of the UCS XML API makes it a powerful tool that is simple to learn and implement.

User	Posts	Date
Sheryl Sage	3	8/27/10
John McDonough	3	7/14/10
Pramod Borkar	1	6/23/10

Web based developer community

<http://developer.cisco.com/web/unifiedcomputing/home>

Downloads:

- UCS Platform Emulator (UCSPE)
- goUCS Automation Tool
- XML API, Perl, PowerShell code samples (44 and counting)
- Microsoft – PowerShell library, SCOM MP
- HP Software – HPOM, HPOO integration modules

Documentation:

- Programming & developer guides
- White papers
- Reference guides (XML model, Faults)

Collaboration:

- Blogs
- Peer to peer forums
- Videos
- Access to Cisco subject matter experts

UCS Platform Emulator (UCSPE)

No UCS hardware needed for code development



- Full featured emulator
- Installs as a Virtual Machine
- Support for all XML API calls
- Object Browser to view the UCSM model
- Import & replicate existing live UCS Manager physical inventory
- Share saved inventories among UCS Platform Emulators
- Drag-and-drop hardware builder to create custom physical inventory

UCS Platform Emulator: Great DEV Tool

Browse XML
objects

UCS
Manager
GUI

Change
emulated HW

UCSPE Control Panel - Cis x

192.168.39.132/config

Cisco UCS Platform Emulator™ Control Panel

UCS MANAGER

- UCS Manager Home
- Managed Object Browser
- API Model Documentation
- XML Examples & XSD Schema
- GoUCS Tool Download
- UCS PowerTool Download
- UCS CDN
- Online Help

HARDWARE INVENTORY

- EMULATOR SETTINGS
- RESTART

Cisco UCS Manager - 2.1(1a)

Single point of device management for the Cisco Unified Computing System

UCS Manager requires Java Runtime Environment 1.6 or higher. If it is not already installed, please [download](#) and install it on your system.

Launch UCS Manager

Launch KVM Manager

Version 2.1(1a)

[Click here](#) to obtain Cisco UCS SNMP MIB files

UCS Platform Emulator: Great DEV Tool

Import inventory
from live UCS
system

Build and replicate
custom blades

The screenshot displays the Cisco UCS Platform Emulator Control Panel interface. The main window is titled "Startup Hardware Inventory" and shows a configuration page for a UCS system. The interface includes a left-hand navigation menu with "UCS MANAGER" and "HARDWARE INVENTORY" sections. The "HARDWARE INVENTORY" section is expanded to show "Start-up Inventory" and "Hardware Catalog".

Key features and components visible in the interface:

- Startup Hardware Inventory:** A central area for configuring hardware. It includes a "Stash (unused servers: 0)" section with a "New Server" button and a "drop new server here" field. Below this are "Rack Servers" sections for "Fabric-A" and "Fabric-B", each with a "New Server" button and a "drop new server rack here" field.
- Chassis Configuration:** Three chassis are shown: "Chassis 1: chassis 1", "Chassis 2: chassis 2", and "Chassis 3: chassis 3". Each chassis has a "Model" (N20-C6508), "UCSM Chassis ID", "Serial" number, and "Fabric Extender" (UCS-IOM-2208XP). Below these details are counts for "Servers", "Fans", and "PSU".
- Hardware Catalog:** A table at the bottom of the interface lists available hardware components. The "Blades" tab is selected, showing a list of server models with their descriptions, vendors, part numbers, and PIDs.

Item	Description	Vendor	Part No	PID
Cisco UCS B200 M1 2 Socket Blade Server	2 Socket, Single slot Blade Server, 12 DIMMs, 2 SFF HDDs, Intel Xeon 5500 series, 1 Mezz. Slot	Cisco Systems, Inc.	74-5390-01	N20-B6620-1
Cisco UCS B440 M1 4 Socket, Extended Memory Blade Server	4 Socket, Dual slot Blade Server, 32 DIMMs, 4 SFF HDDs, Intel Xeon 7500 series, 2 Mezz. Slots	Cisco Systems, Inc.	68-3516-06	N20-B6740-2

PowerShell is Common Scripting Method



- Plain English cmdlets
- Pipeline commands
- Strong library/module/snapin support
- Strong support in Microsoft environments
- Free development tools



PowerShell



PowerTool



PowerCLI

Scripting with PowerShell

```
PS C:\> Get-Command -Module Hyper-V | Measure-Object
```

```
Count      : 177
```

Over 170 cmdlets for Hyper-V

```
PS C:\> Get-Command -Module CiscoUcsPS | Measure-Object
```

```
Count      : 1706
```

Over 1700 cmdlets for UCS

```
PowerCLI C:\> Get-Command -PSSnapin (Get-PSSnapin | ? {$_.Name -match "VMware"}) | Measure-Object
```

```
Count      : 400
```

400 cmdlets for VMware

Real World Scripting Example

How many and what kind of DIMM is in every server I own?

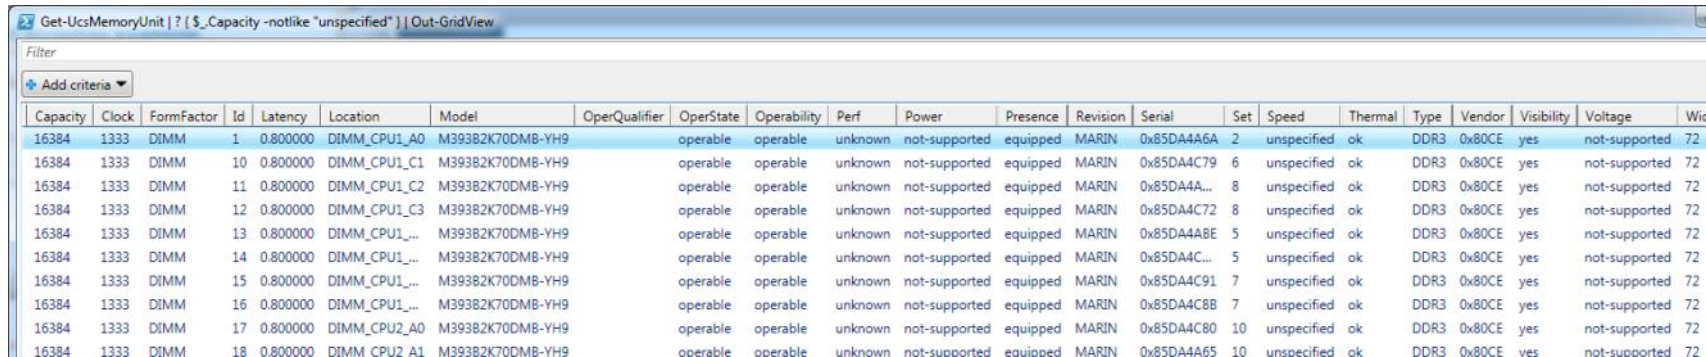
```
Get-UcsMemoryUnit | ? {$_.Capacity -notlike "unspecified"} | ft Bank, Capacity, Dn, Id, Location, Model, Serial, Type, Vendor -AutoSize
```

Or dump it to Excel...

```
Get-UcsMemoryUnit | ? {$_.Capacity -notlike "unspecified"} | Export-Csv -Path "mem.csv"
```

Or dump it to the screen...

```
Get-UcsMemoryUnit | ? {$_.Capacity -notlike "unspecified"} | Out-GridView
```



The screenshot shows a PowerShell window titled "Get-UcsMemoryUnit | ? {\$_.Capacity -notlike 'unspecified'} | Out-GridView". The window displays a table with 22 columns and 15 rows of data. The columns are: Capacity, Clock, FormFactor, Id, Latency, Location, Model, OperQualifier, OperState, Operability, Perf, Power, Presence, Revision, Serial, Set, Speed, Thermal, Type, Vendor, Visibility, Voltage, and Width. The data rows show various DIMM configurations across different CPU slots (A0, C1, C2, C3, ..., A1).

Capacity	Clock	FormFactor	Id	Latency	Location	Model	OperQualifier	OperState	Operability	Perf	Power	Presence	Revision	Serial	Set	Speed	Thermal	Type	Vendor	Visibility	Voltage	Width
16384	1333	DIMM	1	0.800000	DIMM_CPU1_A0	M393B2K70DMB-YH9		operable	operable	unknown	not-supported	equipped	MARIN	0x85DA4A6A	2	unspecified	ok	DDR3	0x80CE	yes	not-supported	72
16384	1333	DIMM	10	0.800000	DIMM_CPU1_C1	M393B2K70DMB-YH9		operable	operable	unknown	not-supported	equipped	MARIN	0x85DA4C79	6	unspecified	ok	DDR3	0x80CE	yes	not-supported	72
16384	1333	DIMM	11	0.800000	DIMM_CPU1_C2	M393B2K70DMB-YH9		operable	operable	unknown	not-supported	equipped	MARIN	0x85DA4A...	8	unspecified	ok	DDR3	0x80CE	yes	not-supported	72
16384	1333	DIMM	12	0.800000	DIMM_CPU1_C3	M393B2K70DMB-YH9		operable	operable	unknown	not-supported	equipped	MARIN	0x85DA4C72	8	unspecified	ok	DDR3	0x80CE	yes	not-supported	72
16384	1333	DIMM	13	0.800000	DIMM_CPU1_...	M393B2K70DMB-YH9		operable	operable	unknown	not-supported	equipped	MARIN	0x85DA4A8E	5	unspecified	ok	DDR3	0x80CE	yes	not-supported	72
16384	1333	DIMM	14	0.800000	DIMM_CPU1_...	M393B2K70DMB-YH9		operable	operable	unknown	not-supported	equipped	MARIN	0x85DA4C...	5	unspecified	ok	DDR3	0x80CE	yes	not-supported	72
16384	1333	DIMM	15	0.800000	DIMM_CPU1_...	M393B2K70DMB-YH9		operable	operable	unknown	not-supported	equipped	MARIN	0x85DA4C91	7	unspecified	ok	DDR3	0x80CE	yes	not-supported	72
16384	1333	DIMM	16	0.800000	DIMM_CPU1_...	M393B2K70DMB-YH9		operable	operable	unknown	not-supported	equipped	MARIN	0x85DA4C8B	7	unspecified	ok	DDR3	0x80CE	yes	not-supported	72
16384	1333	DIMM	17	0.800000	DIMM_CPU2_A0	M393B2K70DMB-YH9		operable	operable	unknown	not-supported	equipped	MARIN	0x85DA4C80	10	unspecified	ok	DDR3	0x80CE	yes	not-supported	72
16384	1333	DIMM	18	0.800000	DIMM_CPU2_A1	M393B2K70DMB-YH9		operable	operable	unknown	not-supported	equipped	MARIN	0x85DA4A65	10	unspecified	ok	DDR3	0x80CE	yes	not-supported	72

Real World Scripting Example

Which profile is associated with a specific MAC address?

```
$mac = "4f"  
Get-UcsVnic | ? {$_.Addr -match $mac} | Format-Table Ucs, Dn, Addr -AutoSize
```

Ucs	Dn	Addr
---	--	----
OPK-SELAB-Area51	org-root/org-jamarche/ls-jmESXTest01/ether-jmESXvmnic1	00:25:B5:AA:E0:4F
OPK-SELAB-Area51	org-root/org-VDI_TEST/ls-M3_ESXi_2/ether-vnic0	00:25:B5:10:00:4F
OPK-SELAB-Area51	org-root/org-VDI_TEST/ls-VI_MGMT/ether-vNIC0	00:25:B5:20:10:4F

Objects

An object is a collection of parts and how to use them

Parts


Front Wheel

Back Wheel

Pedals

Saddle

Frame



How to use

Pedal

Steer Left

Steer Right

Apply Front Brake

Apply Rear Brake

Objects

~~Parts~~
Properties

DisplayName

Status

RequiredServices



Windows
Service

~~How to use~~
Methods

Stop()

Start()

Pause()



Graphics

The BIG DEAL with GPU in VDI Now

Windows XP/Office 2003 EOS – **April 8th 2014**
Windows 7 is the new corporate standard

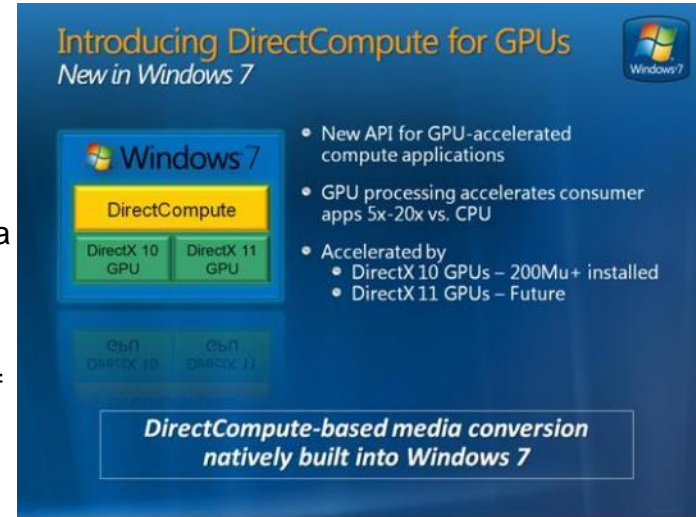
Its not just about CAD/3D graphics - GPU acceleration is native to Win7/Win8 and Office 2013!

Did you know?

GPU enhanced applications offload the CPU for better performance and scalability

GPU enhanced applications use a mix of System RAM and GPU Ram

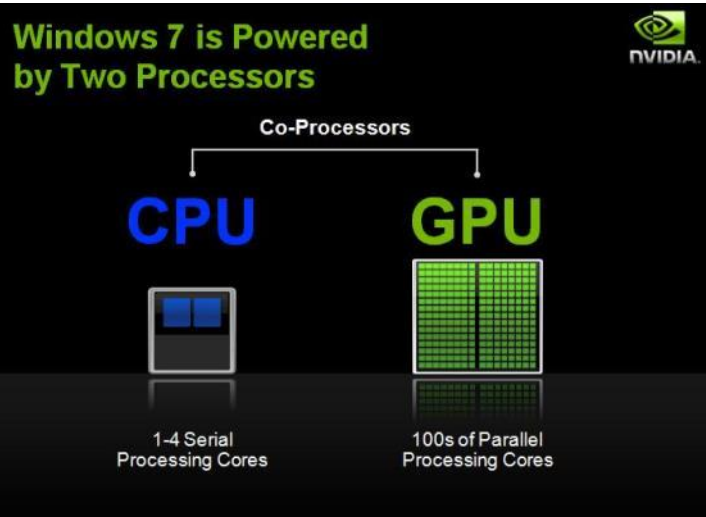
Office 2013: **20% CPU penalty** if GPU not available



Introducing DirectCompute for GPUs
New in Windows 7

- New API for GPU-accelerated compute applications
- GPU processing accelerates consumer apps 5x-20x vs. CPU
- Accelerated by
 - DirectX 10 GPUs – 200Mu+ installed
 - DirectX 11 GPUs – Future

DirectCompute-based media conversion natively built into Windows 7



Windows 7 is Powered by Two Processors

Co-Processors

CPU

1-4 Serial Processing Cores

GPU

100s of Parallel Processing Cores

NVIDIA

Why VDI Implementations Fail

Infrastructure costs and choices:

1. Spinning media doesn't scale
2. CPU selection impacts scale and experience



End User: “It does not look or feel like my old desktop. I don't like it.”

All Users Benefit From GPU



Designer

Graphics & Media
Professionals, Design
Engineers

CATIA, CS6, Inventor



Power User

Financial Analysts,
Traders, Design
Reviewers

PLM, Solidworks, Adobe
Dreamweaver, Medical
Imaging
Showcase



Knowledge Worker

Office workers,
productivity & line-of-
business workers

MS Office, Photoshop



NVIDIA GRID Cards



These 2 video cards are actually
7 different video cards?

	GRID K1	GRID K2
Number of GPUs	4 x entry Kepler GPUs	2 x high-end Kepler GPUs
Total NVIDIA CUDA cores	768 @ 891 MHz	3072 @ 745 MHz
Total memory size	16 GB DDR3 @ 891 MHz	8 GB GDDR5 @ 2,500 MHz
Max power	130 W	225 W
Board length	10.5"	10.5"
Board height	4.4"	4.4"
Board width	Dual slot	Dual slot
Display IO	None	None
Aux power	6-pin connector	8-pin connector
PCIe	x16	x16
PCIe generation	Gen3 (Gen2 compatible)	Gen3 (Gen2 compatible)
Cooling solution	Passive	Passive
Technical Specifications	GRID K1 Board Specifications	GRID K2 Board Specifications

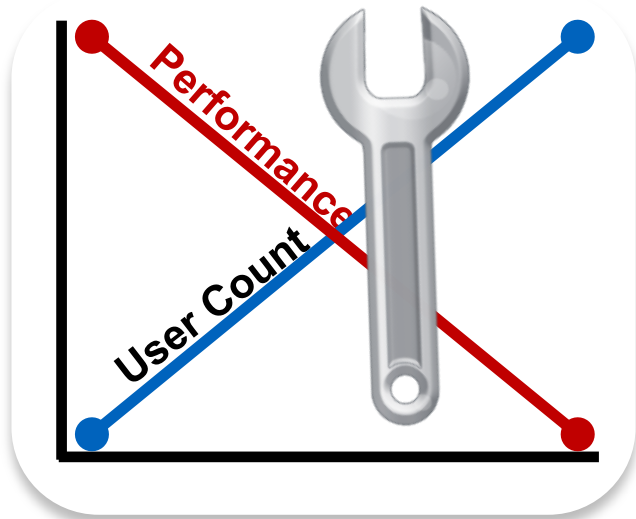
NVIDIA GRID Graphics Board	Virtual GPU Profile	Application Certifications	Graphics Memory	Max Displays Per User	Max Resolution Per Display	Max Users Per Graphics Board	Use Case
GRID K2	K260Q	✓	2,048 MB	4	2560x1600	4	Designer/Power User
	K240Q	✓	1,024 MB	2	2560x1600	8	Designer/Power User
	K200		256 MB	2	1900x1200	16	Knowledge Worker
GRID K1	K140Q	✓	1,024 MB	2	2560x1600	16	Power User
	K100		256 MB	2	1900x1200	32	Knowledge Worker
GRID K1	PassThru	✓	4,096 MB	4	2560x1920	4	Designer/Power User
GRID K2	PassThru	✓	4,096 MB	4	2560x1920	2	Designer/Power User

Modes of GPU Implementation

Shared GPU = Software virtualisation of the GPU

GPU Pass-through = 1:1 dedicated GPU to user

Virtual GPU = Hardware virtualisation of the GPU through NVIDIA GRID software



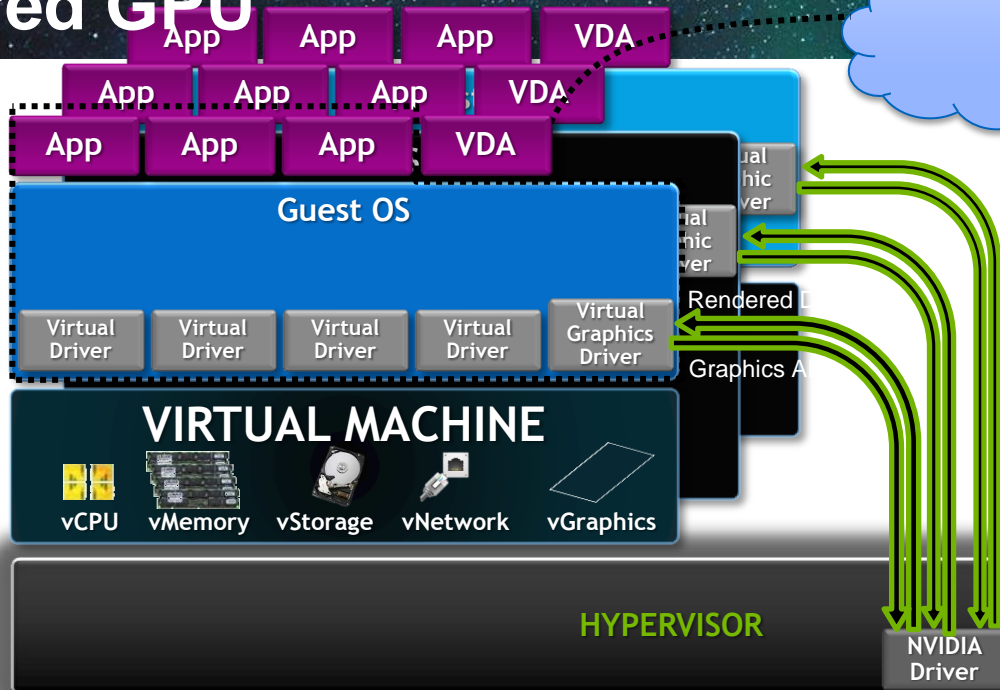
User count goes up and performance goes down, but **not at the same magnitude** in each mode.

Shared GPU

Software


Virtualisation

Hardware

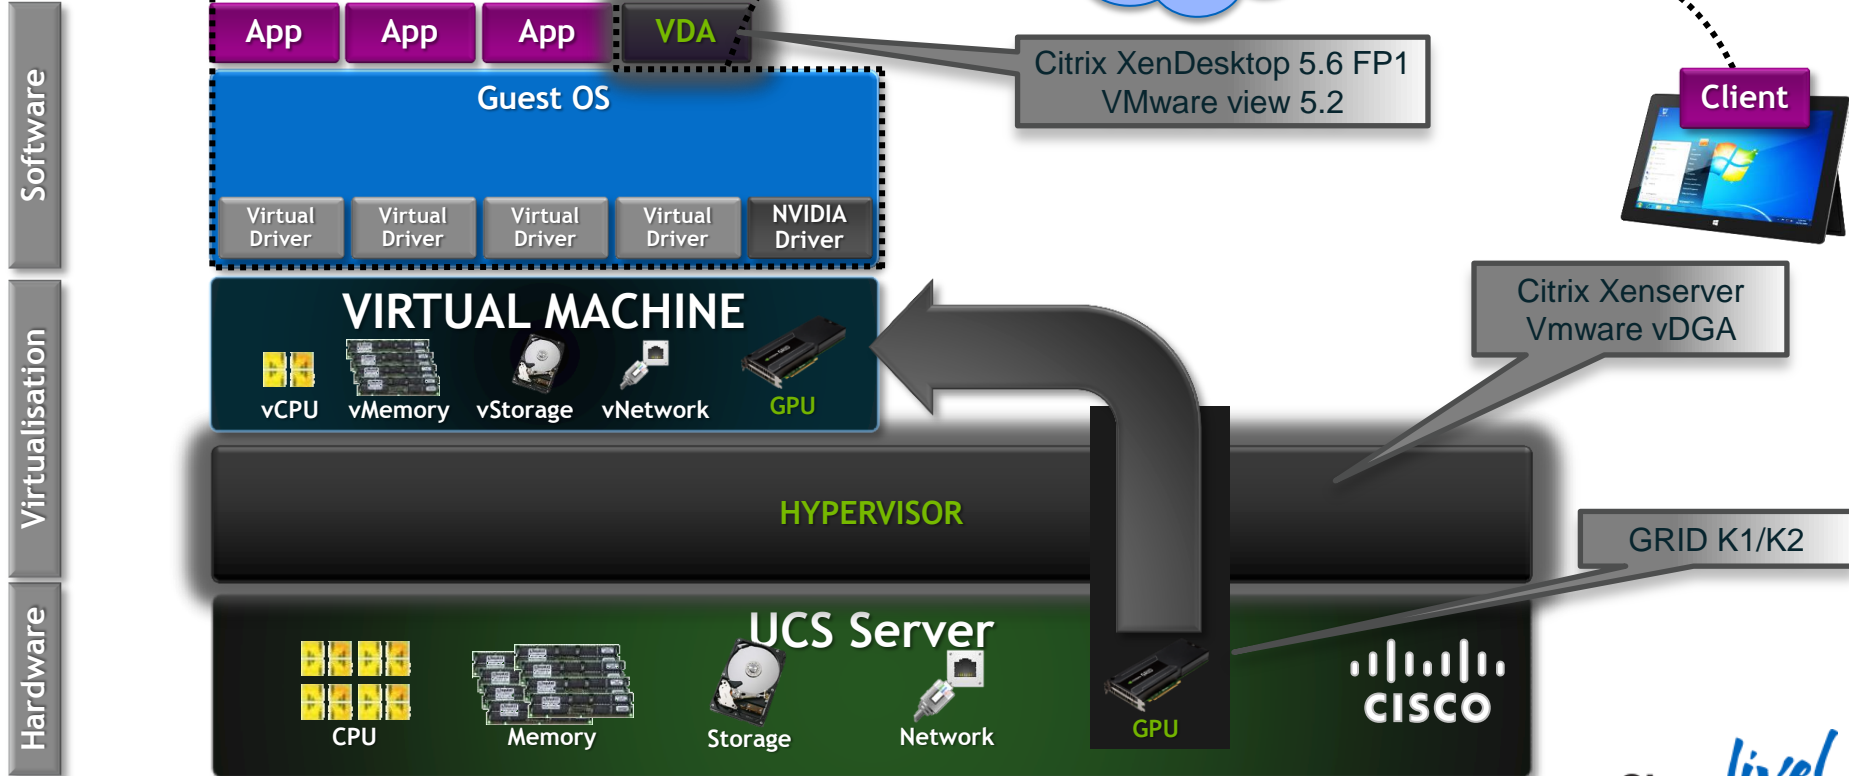


RemoteFX
Microsoft Server 2008 R2 – DX9
Microsoft Server 2012 – DX9, 10, 11
VMware vSGA - DX9, OGL2.1

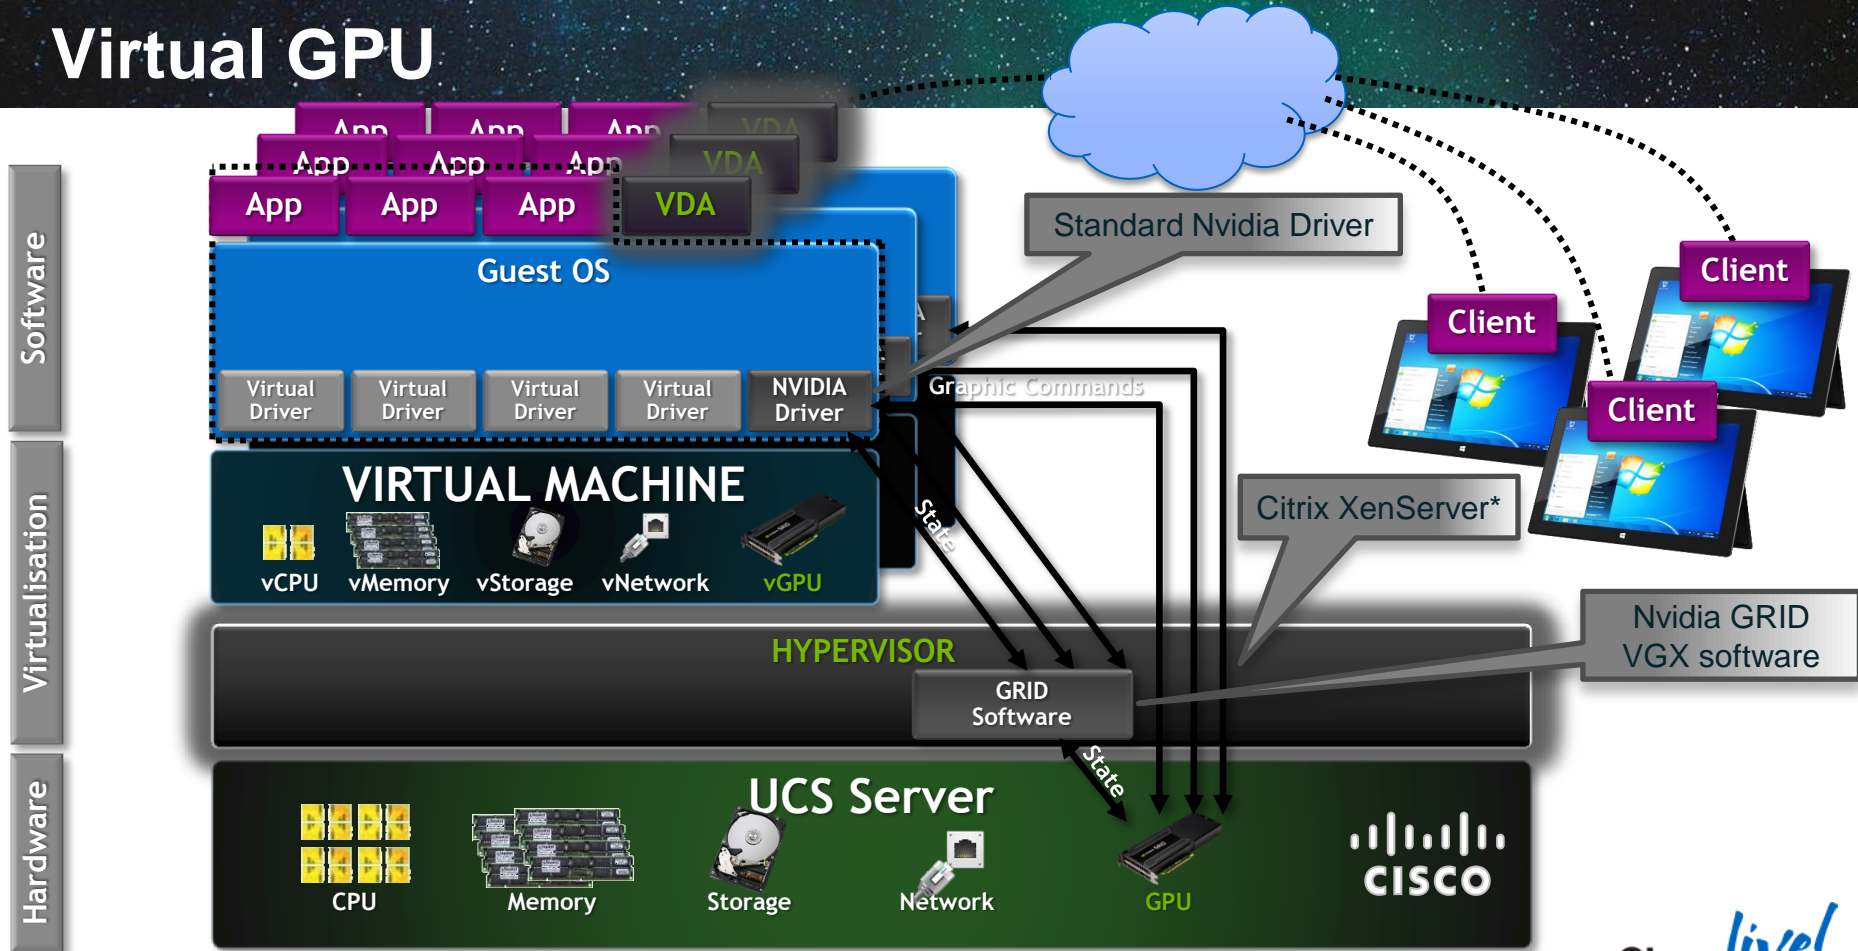
GRID K1/K2






GPU Pass-Through



Virtual GPU



GPU Support for VDI Profile

Vendor	GPU Pass-Through	GPU Sharing	Virtual GPU
	✓	XenApp 6.5 on Windows Server	✓
	✓ (vDGA)	✓ (vSGA)	✗
	✓	✓ (RemoteFX)	✗

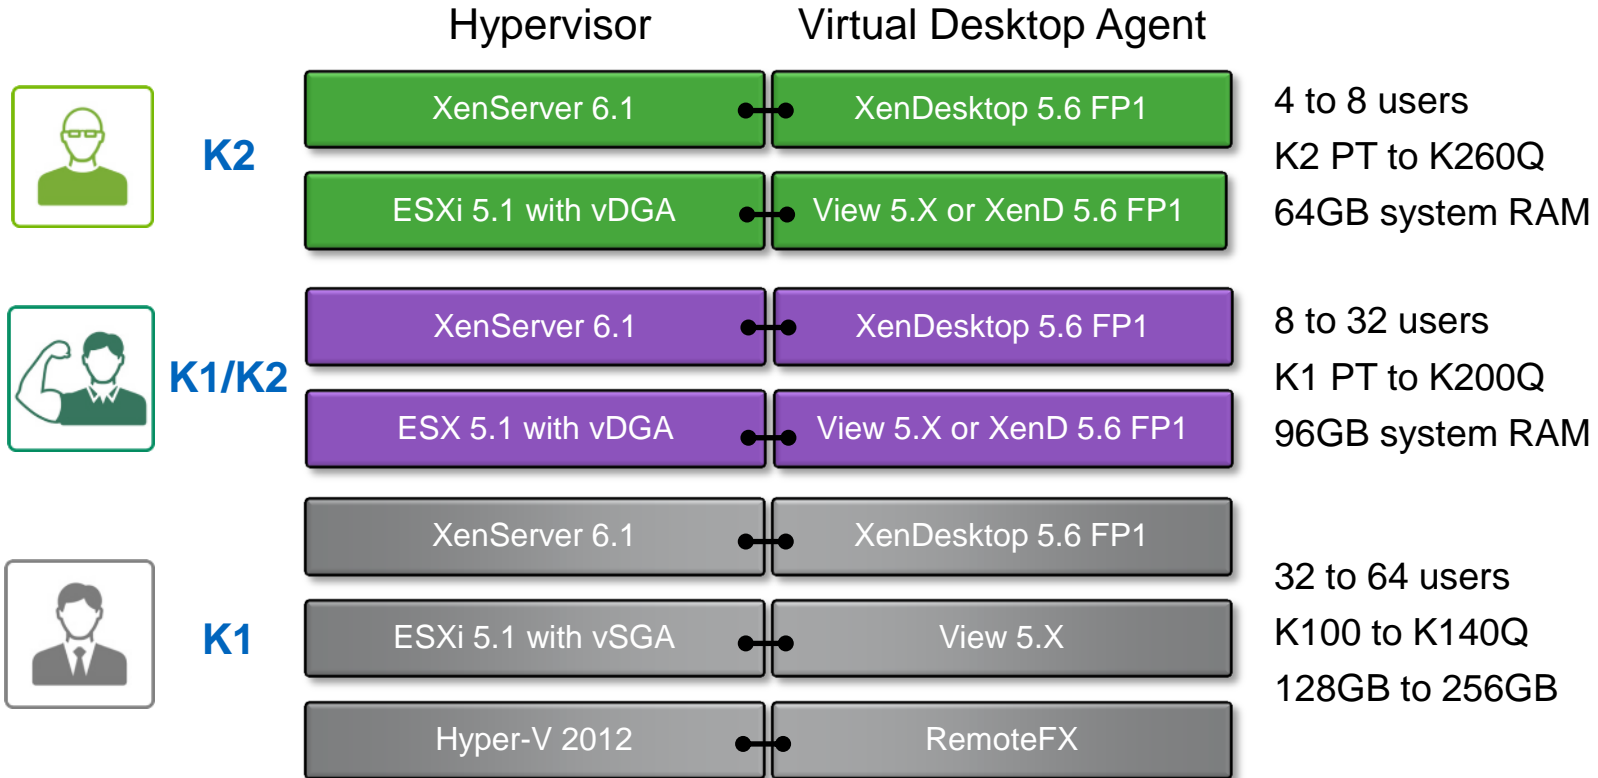
vGPU Unigine Results



NVIDIA GRID Graphics Board	Virtual GPU Profile	Application Certifications	Graphics Memory	Max Displays Per User	Max Resolution Per Display	Max Users Per Graphics Board	Use Case
GRID K2	K260Q	✓	2,048 MB	4	2560x1600	4	Designer/Power User
	K240Q	✓	1,024 MB	2	2560x1600	8	Designer/Power User
GRID K1	K200		256 MB	2	1900x1200	16	Knowledge Worker
	K140Q	✓	1,024 MB	2	2560x1600	16	Power User
	K100		256 MB	2	1900x1200	32	Knowledge Worker
GRID K1	PassThru	✓	4,096 MB	4	2560x1920	4	Designer/Power User
GRID K2	PassThru	✓	4,096 MB	4	2560x1920	2	Designer/Power User

Card	Mode	GPU Speed	Mem Speed	Temp	Benchmark	Type	Min FPS	Max FPS	Avg FPS	Score
K2	K200	NA	NA	NA	Heaven	Basic	6.4	45.4	22	554
K2	K200	NA	NA	NA	Heaven	Extreme	1.5	5.8	2.5	63
K2	K200	NA	NA	NA	Valley	Basic	5.1	44.8	16.4	684
K2	K200	NA	NA	NA	Valley	Extreme	1.1	3.9	1.8	74
K2	K240Q	NA	NA	NA	Heaven	Basic	8.6	66.1	53.1	1337
K2	K240Q	NA	NA	NA	Heaven	Extreme	7.6	57	29.6	746
K2	K240Q	NA	NA	NA	Valley	Basic	8.9	65.4	33.9	1418
K2	K240Q	NA	NA	NA	Valley	Extreme	8.9	49.8	28.3	1183
K2	K260Q	NA	NA	NA	Heaven	Basic	8.6	66.6	53.3	1343
K2	K260Q	NA	NA	NA	Heaven	Extreme	7.5	57.3	30.5	768
K2	K260Q	NA	NA	NA	Valley	Basic	8.1	59.3	33.2	1390
K2	K260Q	NA	NA	NA	Valley	Extreme	9.8	49.1	29.3	1225
K2	Pass Through	745 MHz	2,500 MHz	62C	Heaven	Basic	8.9	155.2	87.7	2209
K2	Pass Through	745 MHz	2,500 MHz	62C	Heaven	Extreme	7.4	78.2	30.7	774
K2	Pass Through	745 MHz	2,500 MHz	62C	Valley	Basic	15	111.7	59.5	2489
K2	Pass Through	745 MHz	2,500 MHz	62C	Valley	Extreme	12.3	66.9	34.3	1436
K1	K100	NA	NA	NA	Heaven	Basic	6.1	24	12.6	316
K1	K100	NA	NA	NA	Heaven	Extreme	1.3	3.2	1.8	CRASH 24/26
K1	K100	NA	NA	NA	Valley	Basic	4.2	14.4	7.7	322
K1	K100	NA	NA	NA	Valley	Extreme	1.2	3.7	1.7	73
K1	K140Q	NA	NA	NA	Heaven	Basic	6.5	27.8	17.7	446
K1	K140Q	NA	NA	NA	Heaven	Extreme	3.2	12.4	5.8	145
K1	K140Q	NA	NA	NA	Valley	Basic	5.7	23.7	12.3	513
K1	K140Q	NA	NA	NA	Valley	Extreme	3.7	11.8	6.8	284
K1	Pass Through	849 MHz	891 MHz	38C	Heaven	Basic	6.5	27.7	17.6	444
K1	Pass Through	849 MHz	891 MHz	38C	Heaven	Extreme	3.3	13.9	6.3	158
K1	Pass Through	849 MHz	891 MHz	38C	Valley	Basic	8.8	27.3	14.6	609
K1	Pass Through	849 MHz	891 MHz	38C	Valley	Extreme	5.3	13.5	7.8	326

GRID Reference Architecture



Getting Started



- Cisco UCS C240 with two K1 cards orderable!
- Cisco UCS C240 with two K2 cards orderable!
- Working on C240 with K1/K2 Mix for “GPU Evaluation”

Cisco UCS supports single-wire management of rack servers.
Manage a rack mount just like a blade!



Q & A

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2014 Polo Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site www.ciscoliveaustralia.com/mobile
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 21 March 12:00pm - 2:00pm



Learn online with Cisco Live!

Visit us online after the conference for full access to session videos and presentations.

www.CiscoLiveAPAC.com



CISCO™