Designing Your VMware Virtual Infrastructure for Optimal Performance, Resilience and Availability – *Straight from the Source*

Deji Akomolafe – VMware
David Klee – Heraflux Technologies
Cody Chapman – Heraflux Technologies

usenix LISA16

December 4–9, 2016 | Boston, MA
www.usenix.org/lisa16 #lisa16

Dèjì Akọ́mọláfẹ́

@dejify

Staff Solutions Architect, VMware Global Technical and Professional Services

- Microsoft Applications Virtualization Lead
- Member of VMware CTO Ambassador Program
- 20+ years IT experience, specializing in Microsoft technologies
- Former Microsoft MVP in multiple designations
  - Exchange Server
  - Directory Services
  - Windows Security
- Speaker at:
  - VMworld | EMCworld | VMUG | SQL Saturday

https://blogs.vmware.com/apps

http://www.dejify.com

http://bit.ly/2h3Rf53

About David Klee

Microsoft® MVP Most Valuable Professional

vmware® vEXPERT

@kleegeek

davidklee.net

heraflux.com

linkedin.com/in/davidaklee

- Performance Tuning & Troubleshooting
- Virtualization
- Cloud Enablement
- Infrastructure Architecture
- Health & Efficiency
- Capacity Management

HERAFLUX TECHNOLOGIES®

About Cody Chapman

@codyrchapman
heraflux.com
linkedin.com/in/codyrchapman

- Performance Tuning & Troubleshooting
- Virtualization
- Infrastructure Architecture
- Scripting and Automation
- Health & Efficiency

HERAFLUX TECHNOLOGIES®

Things we can all agree on

Virtualization is mainstream

You want to virtualize your applications

You care about the outcome

Your applications are Important

That is WHY we are here

Is the Application "Critical"?

- Do business processes depend on it?
- Is the outage impactful?
- Outage NOT easily survivable?
- Outage NOT easily recoverable?
- Will it/you be missed?

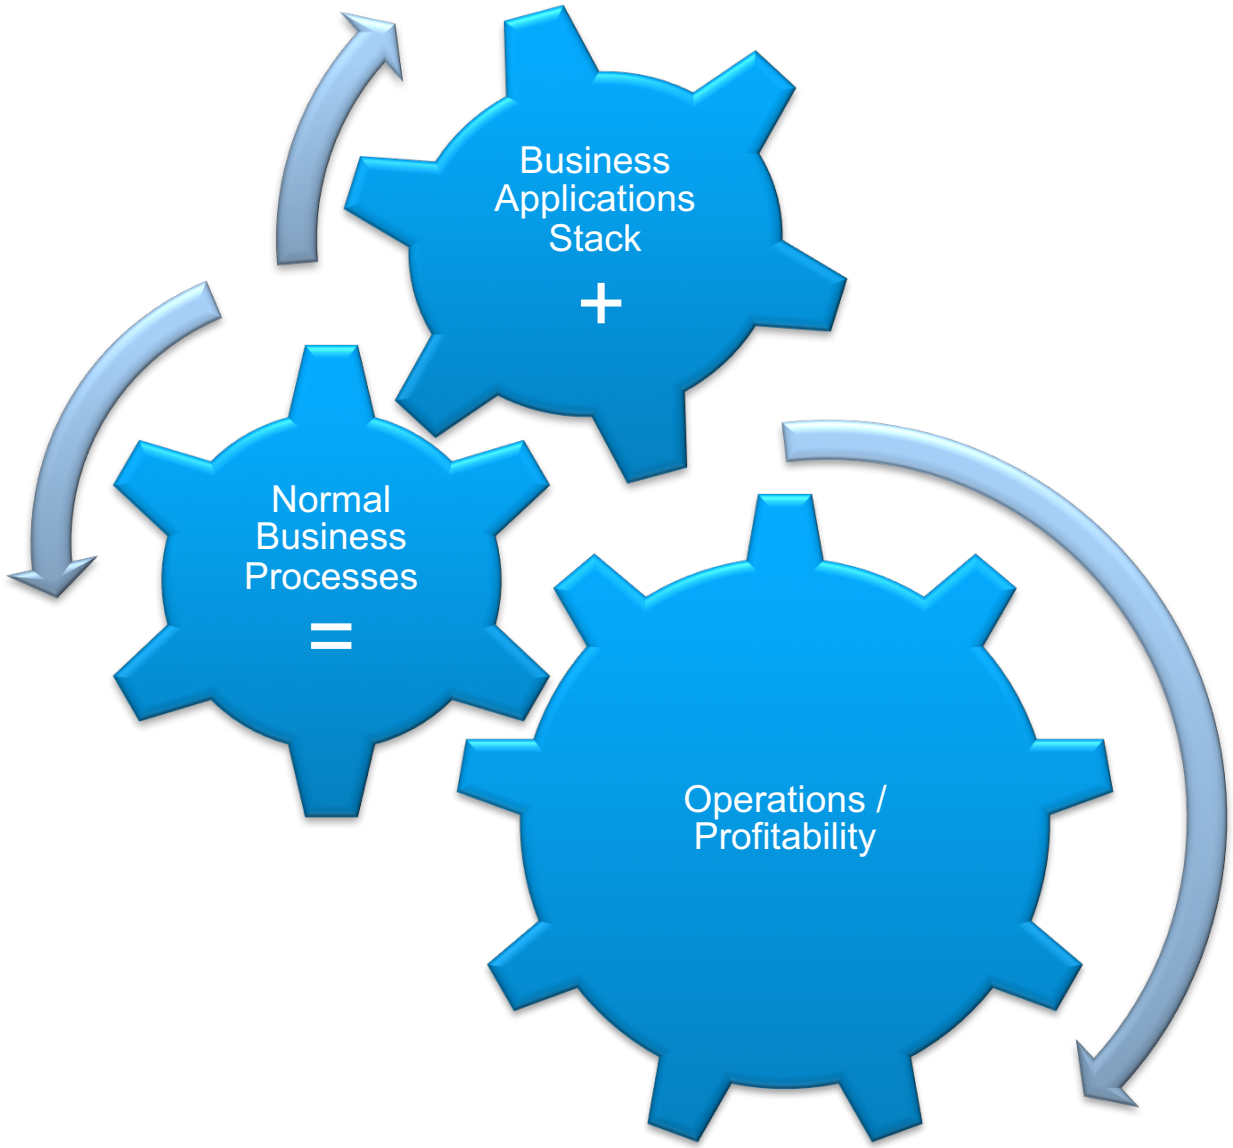

Business Critical Applications Characteristics

Why Virtualize Critical Applications

Resource Maximization
- Server resources increase too much for one application instance
- Virtualization improves resource utilization
- Reduces wastage

Enhanced Availability
- Native application HA features incomplete for most critical applications
- vSphere HA features complement native App HA features
- Result is improved availability

Dev Testing
- Virtualization improves adaptivity and elasticity
- Lifecycle management easier in virtual (provisioning/de-provisioning)

Rapid Provisioning And Scaling
- All the known and latent benefits of virtualization
- Project lifecycle considerably reduced

Job Security
- It's 2016, and all the cool kids have done it
- You can't get to the "Cloud" without virtualizing

Lower TCO
- Significant savings in power, cooling, and datacenter space, and administration

Common Objections
To Virtualizing Critical Applications

Common Objections to Virtualizing BCA

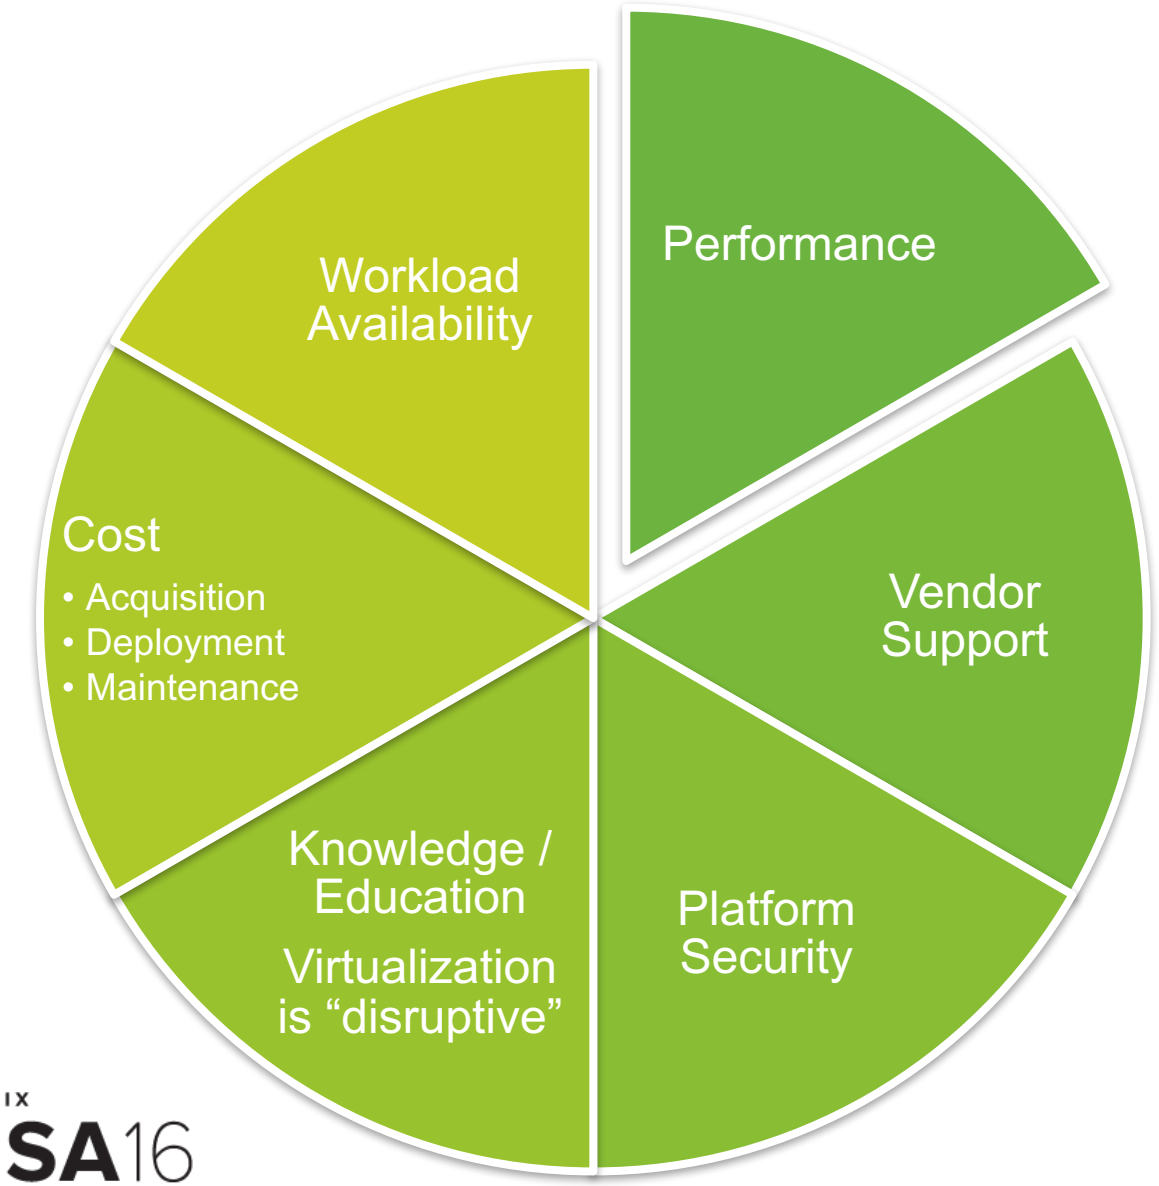

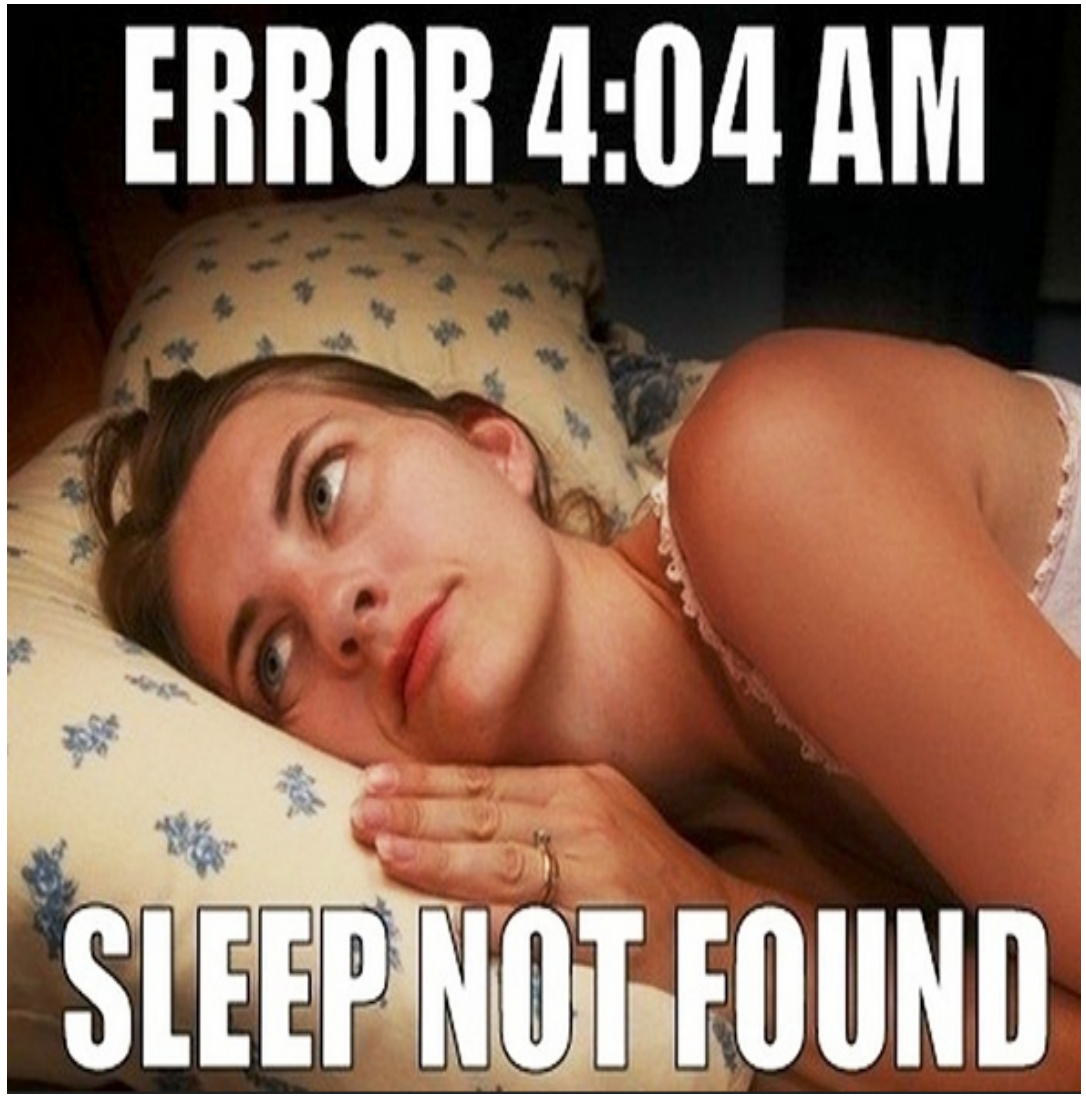

Common Objections to Virtualization - Vendor Support

Vendor	Reference
Everything Business Critical Applications on VMware vSphere • http://www.vmware.com/business-critical-apps/	http://vmw.re/15MO7oL
Microsoft Supports Virtualization of ALL its Critical Applications	http://bit.ly/1uvVRkk
• Exchange Server	http://bit.ly/1H1xYfu
• SQL Server	http://bit.ly/15MrBMy
Oracle mySupport (**Note 249212.1**)	http://bit.ly/15DrLW3
SAP General Support Statement for Virtual Environments (**Note 1492000**)	http://bit.ly/1Ctkd4T
• SAP on VMware	http://bit.ly/15NEiH4
• SAP Notes Related to VMware	http://bit.ly/1wyohKe

For when you are in a jam
http://www.tsanet.org

Common Objections to Virtualization - Security

I heard about a TPS Security Bug	Yes, we did, too, and we quashed it – http://vmw.re/1x95NBV
I have a Regulatory Compliance Requirement for "Hard" Separation	Multi-tenancy and "fencing" allowed Multi-tenancy is NOT a requirement
Privilege Escalation	vCenter privileges do NOT elevate guest operating system or application privileges
Deviates from our build standards	Virtualization improves standardization Use templates for optimization
The fear of the "stolen vmdk"	How about the "stolen server"? Or "stolen/copied backup tape"? We have a solution in just a few slides…

Stolen VMDK? Meet VM Encryption

The "Dye Pack" of Enterprise Virtualization

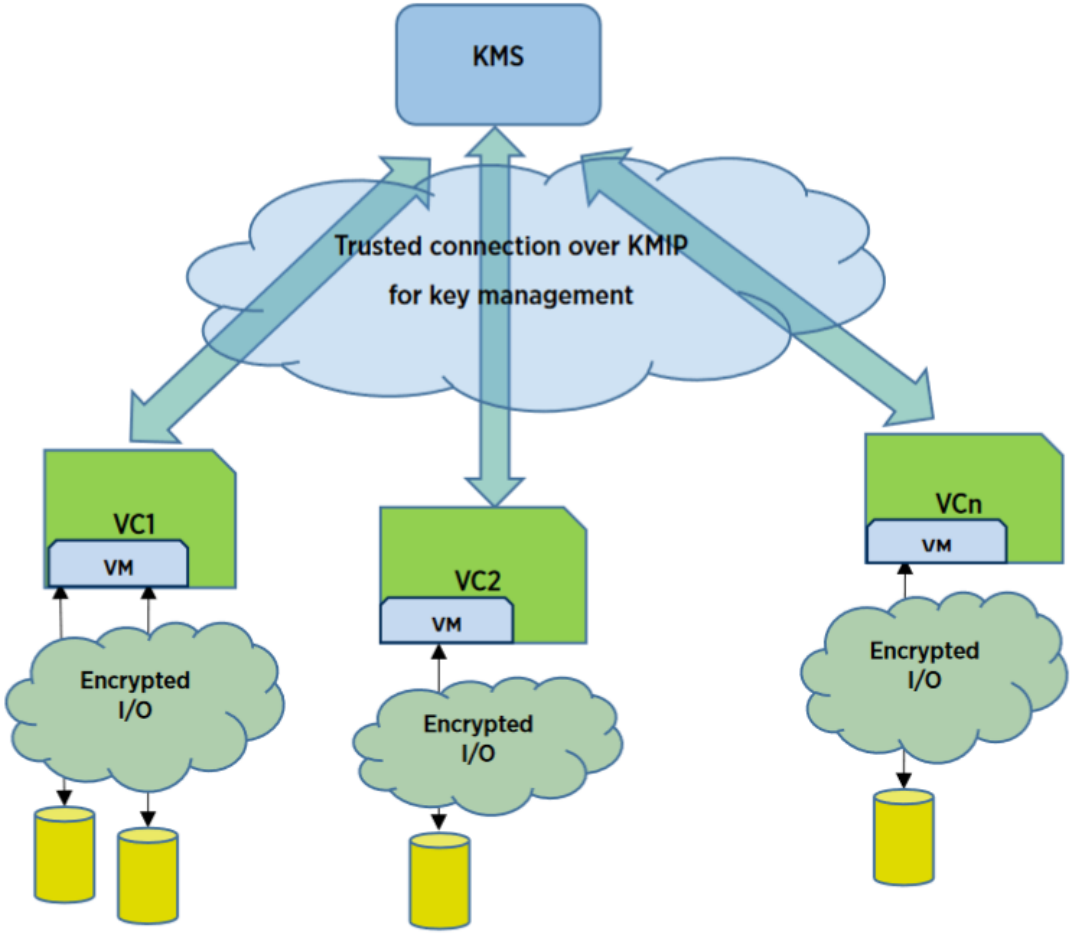

- Introduced in vSphere 6.5
- Secures Data in a VM's VMDK
- Uses vSphere APIs for I/O filtering (VAIO)
- VM Possesses Decryption Key
- vCenter Serves as Broker/Facilitator Only
- Data Meaningless to Unauthorized Entities
- No SPECIAL Hardware Required *

* AES-NI Capable Server Hardware Improves Performance

VM Encryption – How it Works

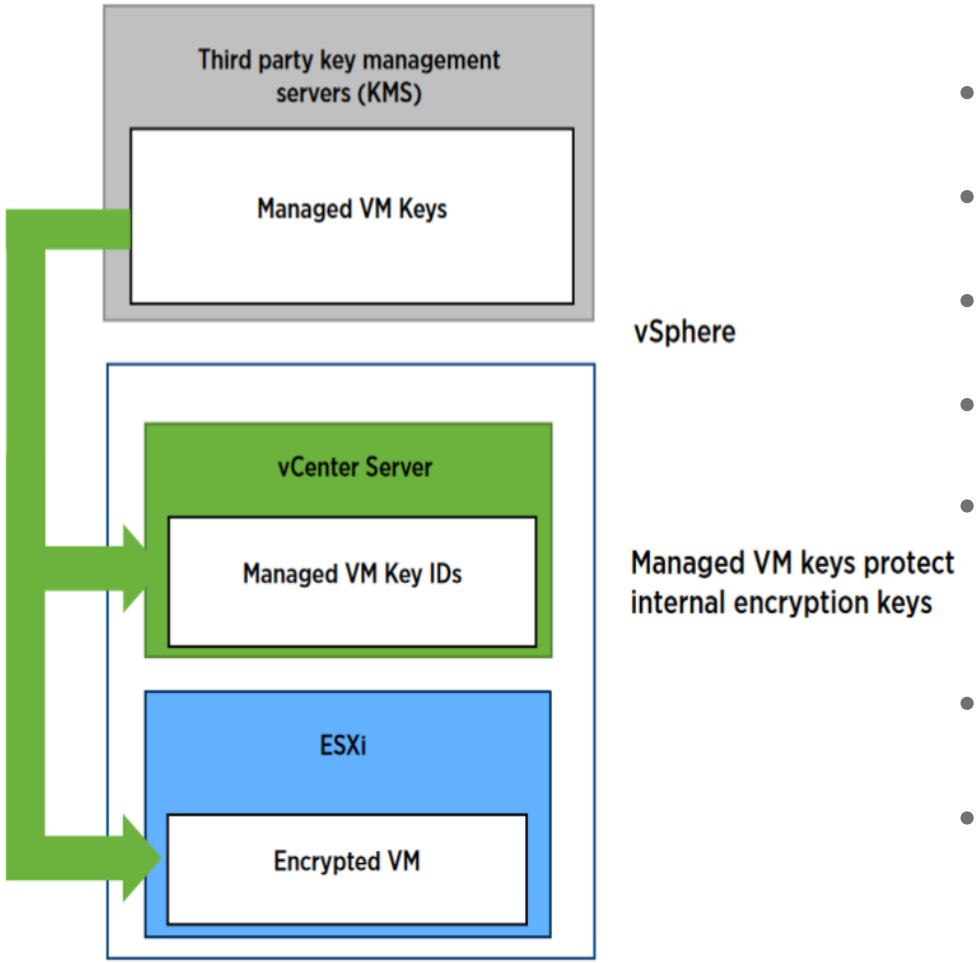

- Customer-Supplied Key Management Server (KMS)
  - Customer-owned and Operated
- Centralized Repository for Crypto Keys
  - No Special Requirement – KMIP 1.1-compliant
- KMS Clusters can be created
  - For Redundancy and Availability
- vCenter is Manually Enrolled in KMS
  - Establishing Trust
- vCenter Obtains Crypto KEKs from KMS
  - Distributes KEKs to ESXi
- ESXi Uses KEK to Generate DEK
  - Used for Encrypting/Decrypting VM Files
  - Encrypted DEKs Stored in VM Config Files
  - KEK for VMs Resides in ESXi's Memory
- IF ESXi Powered-Cycled (or Otherwise Unavailable), vCenter **Must Request** New KEK for Host
- If Encrypted VM Unregistered, vCenter **Must Request** KEK During Re-Registration

VM Unable to Power-On if Request Fails

Common Objections to Virtualization - Knowledge / Education

The Fear of Change…. Leads to inertia

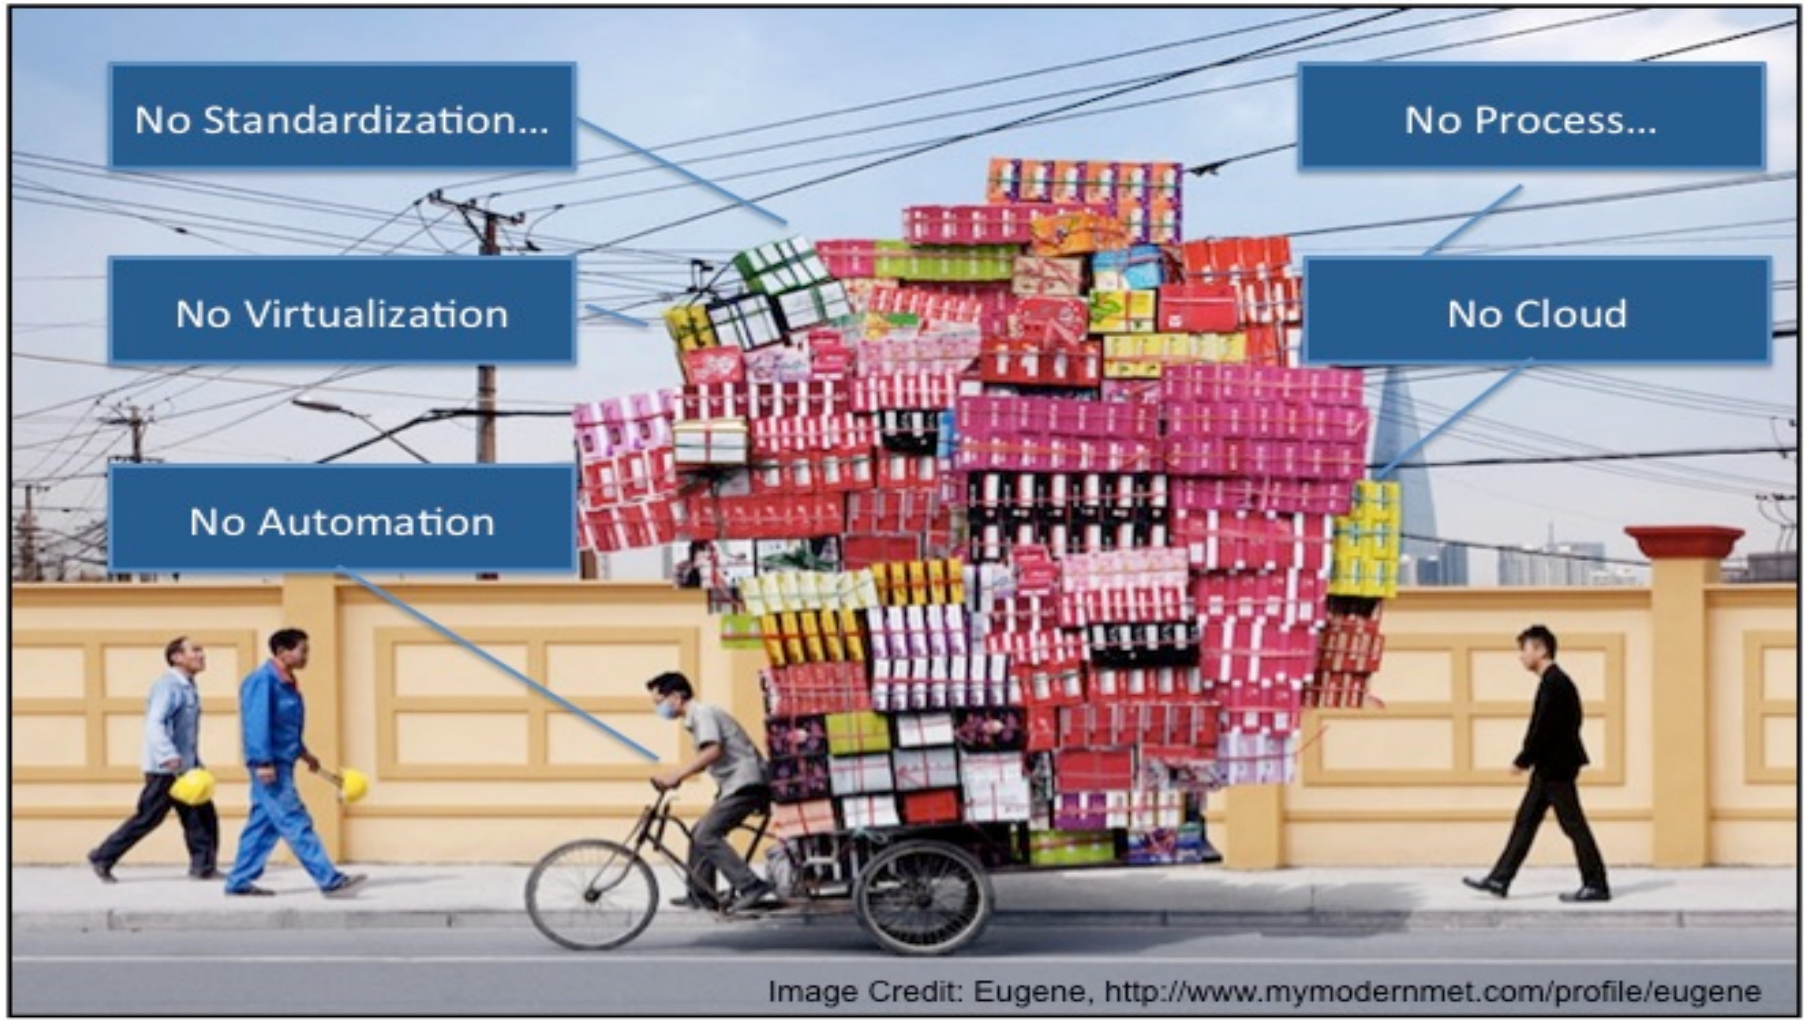

Virtualizing Applications for Performance and Scale

Can vSphere handle the load?

Configuration Item	ESXi 6.0	ESXi 6.5
Virtual CPUs per virtual machine (Virtual SMP)	128	128
RAM per virtual machine	4TB	6TB
Virtual machine swapfile size	4TB	6TB
Logical CPUs per host	480	576
Virtual CPUs per host	4096	4096
Virtual machines per host	1024	1024
Virtual CPUs per core	32	32
Virtual CPUs per FT virtual machine	4	4
FT Virtual machines per host	4	4
RAM per host	4TB	6TB
Hosts per cluster	64	64
Virtual Machines per cluster	8000	8000
LUNs per cluster/host	254	512
Paths per cluster/host	1024	2048
LUN / VMDK Size	62 TB	62 TB
Virtual NICs per virtual machine	10	10

Ensuring Application Performance on vSphere

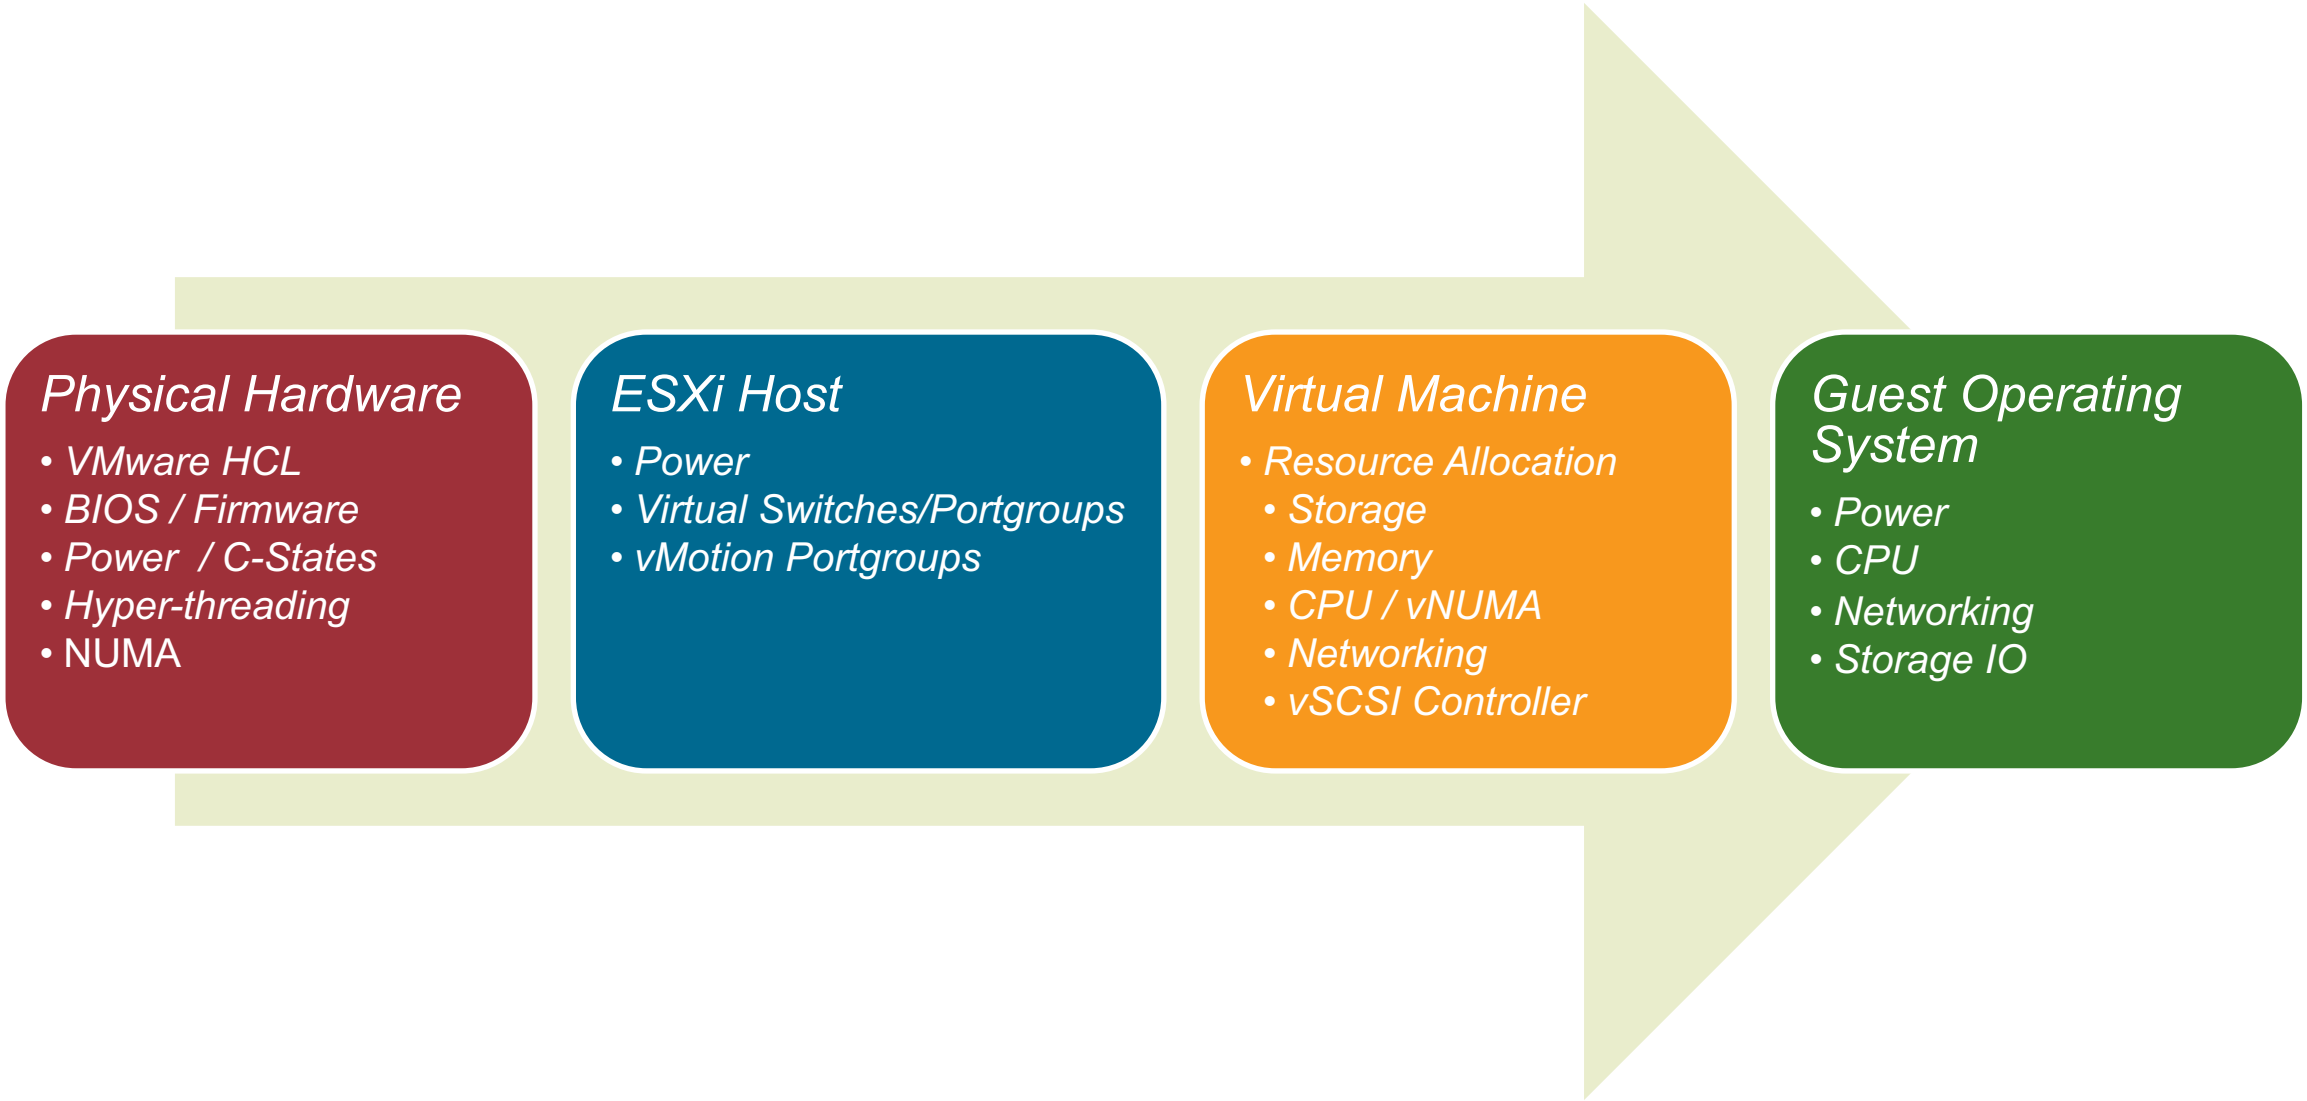

Designing to Requirements – Know the Constraints

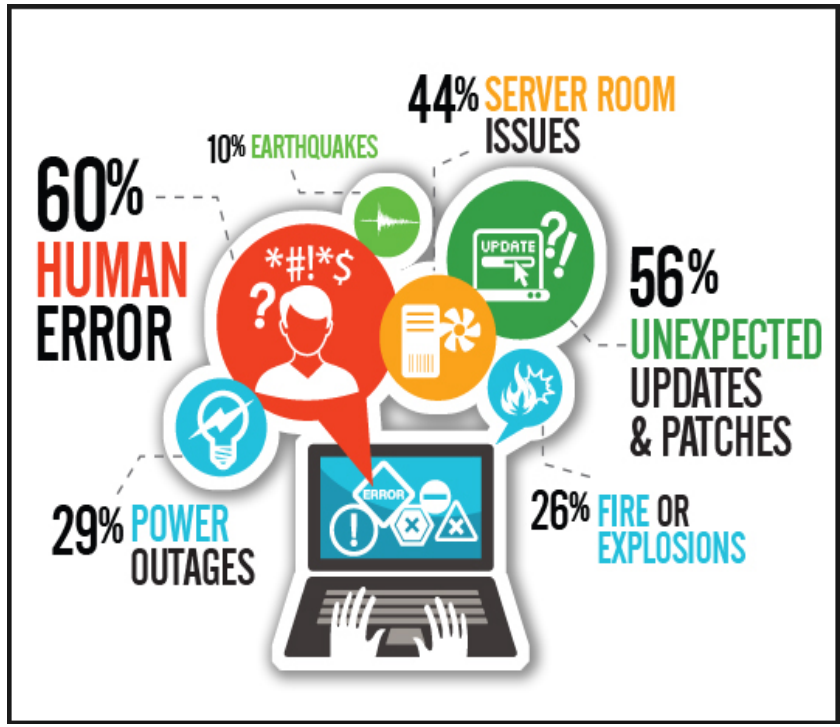

Performance and Scale

Availability and Reliability

Recoverability

Design Constraints

Budget	Personnel	Compliance	Storage	Networking	Server Hardware	vSphere	Windows	Application

Performance-based Designing Tenets

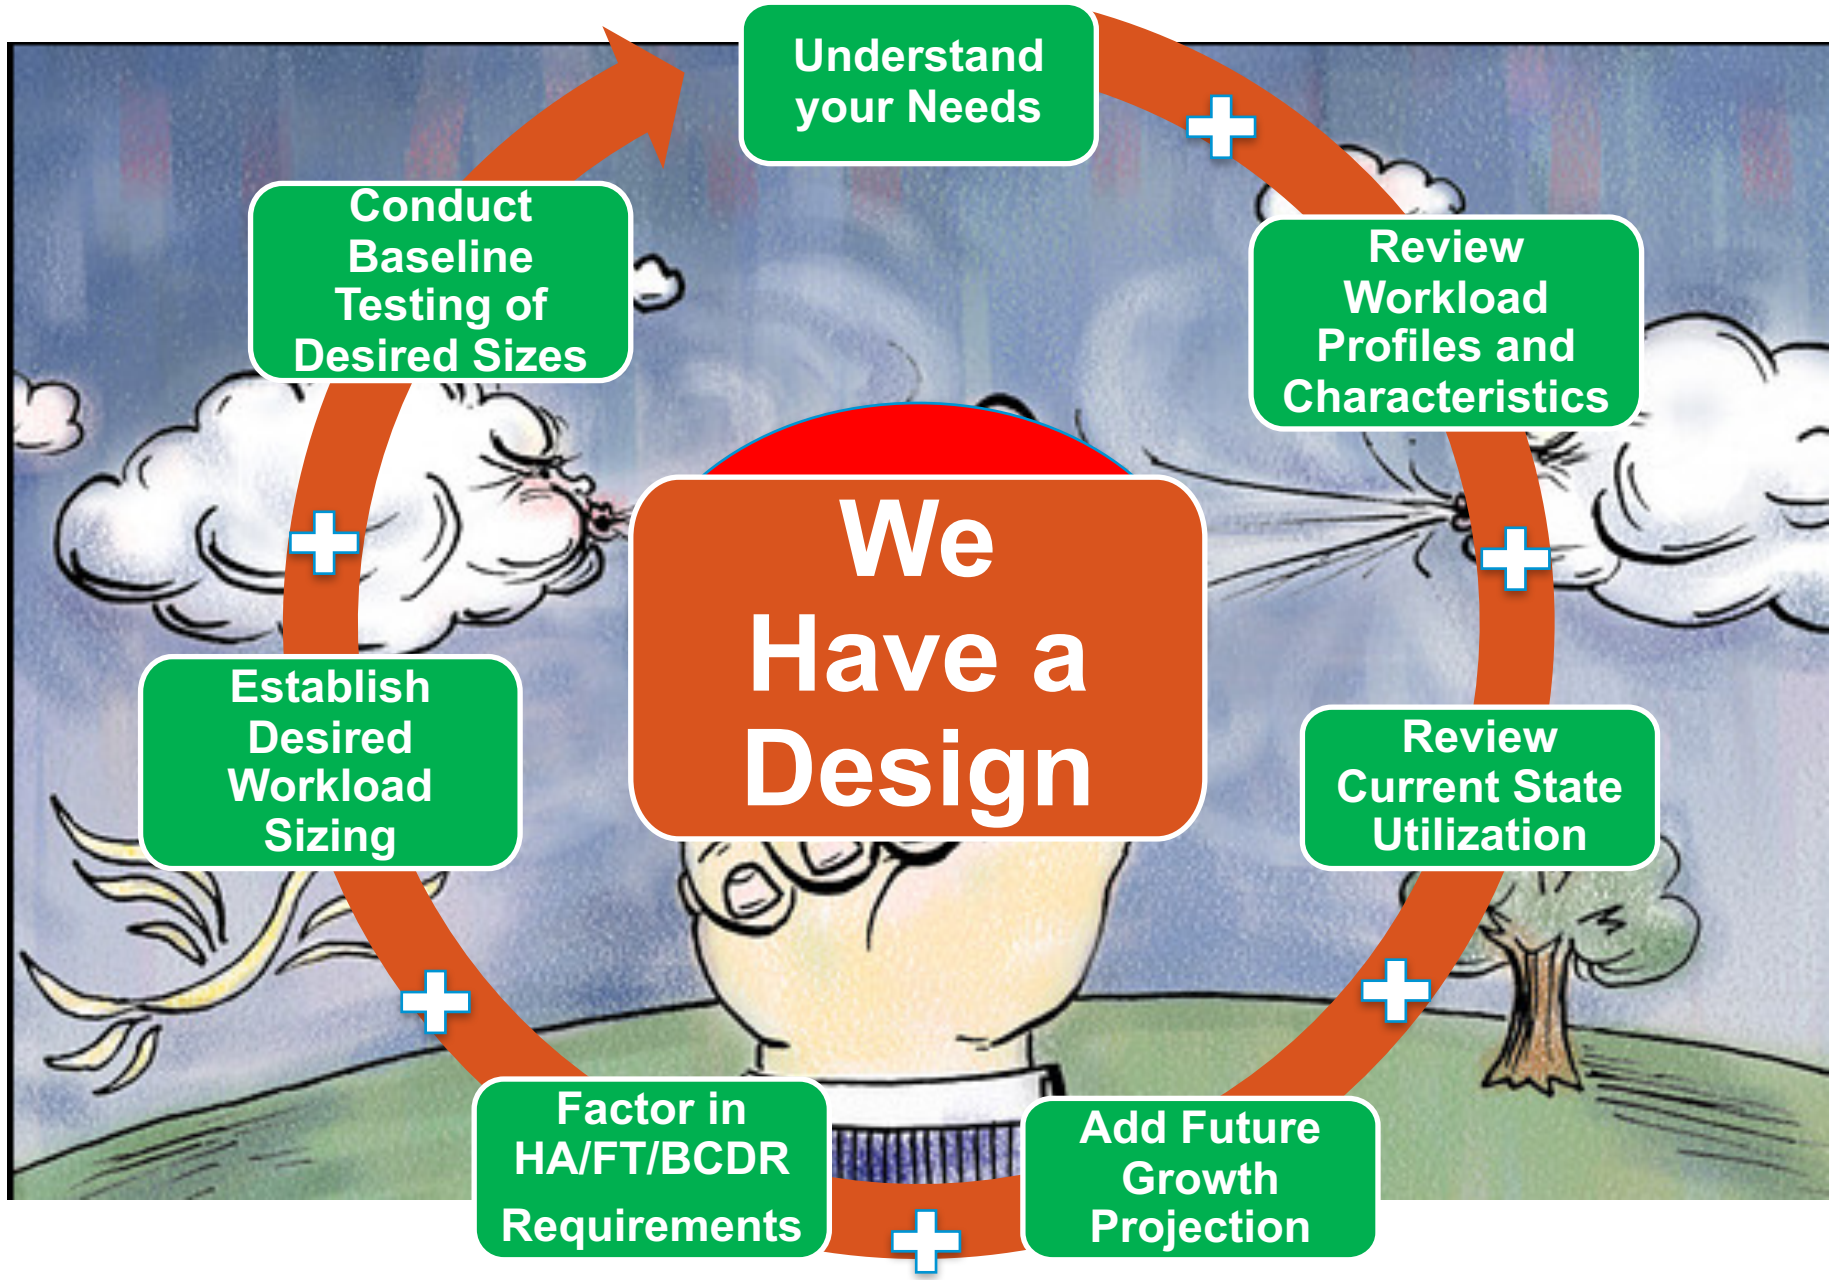

Everything rides on the physical hardware – E.V.E.R.Y.T.H.I.N.G

- Physical Hardware
  - Hardware MUST Be On VMware's HCL
    - Outdated drivers, firmware and BIOS Revs adversely impact virtualization
  - Always Disable unused physical hardware devices
  - Leave memory scrubbing rate in BIOS at default
    - Incorrect firmware, BIOS and Drivers Revs adversely impact virtualization
  - Default hardware Power Scheme unsuitable for virtualization
    - Change Power setting to **"OS controlled"**
    - Set ESXi Power Management Policy to **"High Performance"**
    - Enable Turbo Boost (or Equivalent)
    - Disable Processor C-states / C1E halt State
  - Enable All Cores – Don't let hardware turn off cores dynamically

VMware Compatibility Guide

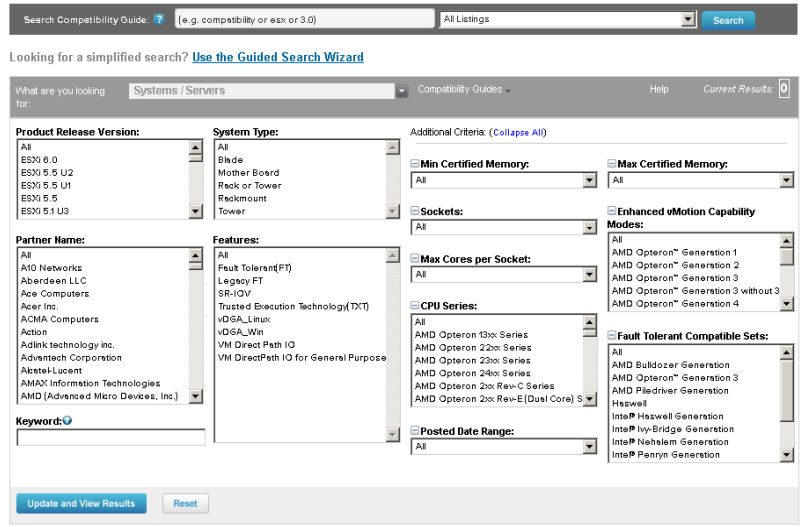

WRONG BIOS, FIRMWARE, AND DRIVERS REVS ADVERSELY IMPACT VIRTUALIZATION

Time-Keeping in your vSphere Infrastructure

Back in the Days…..

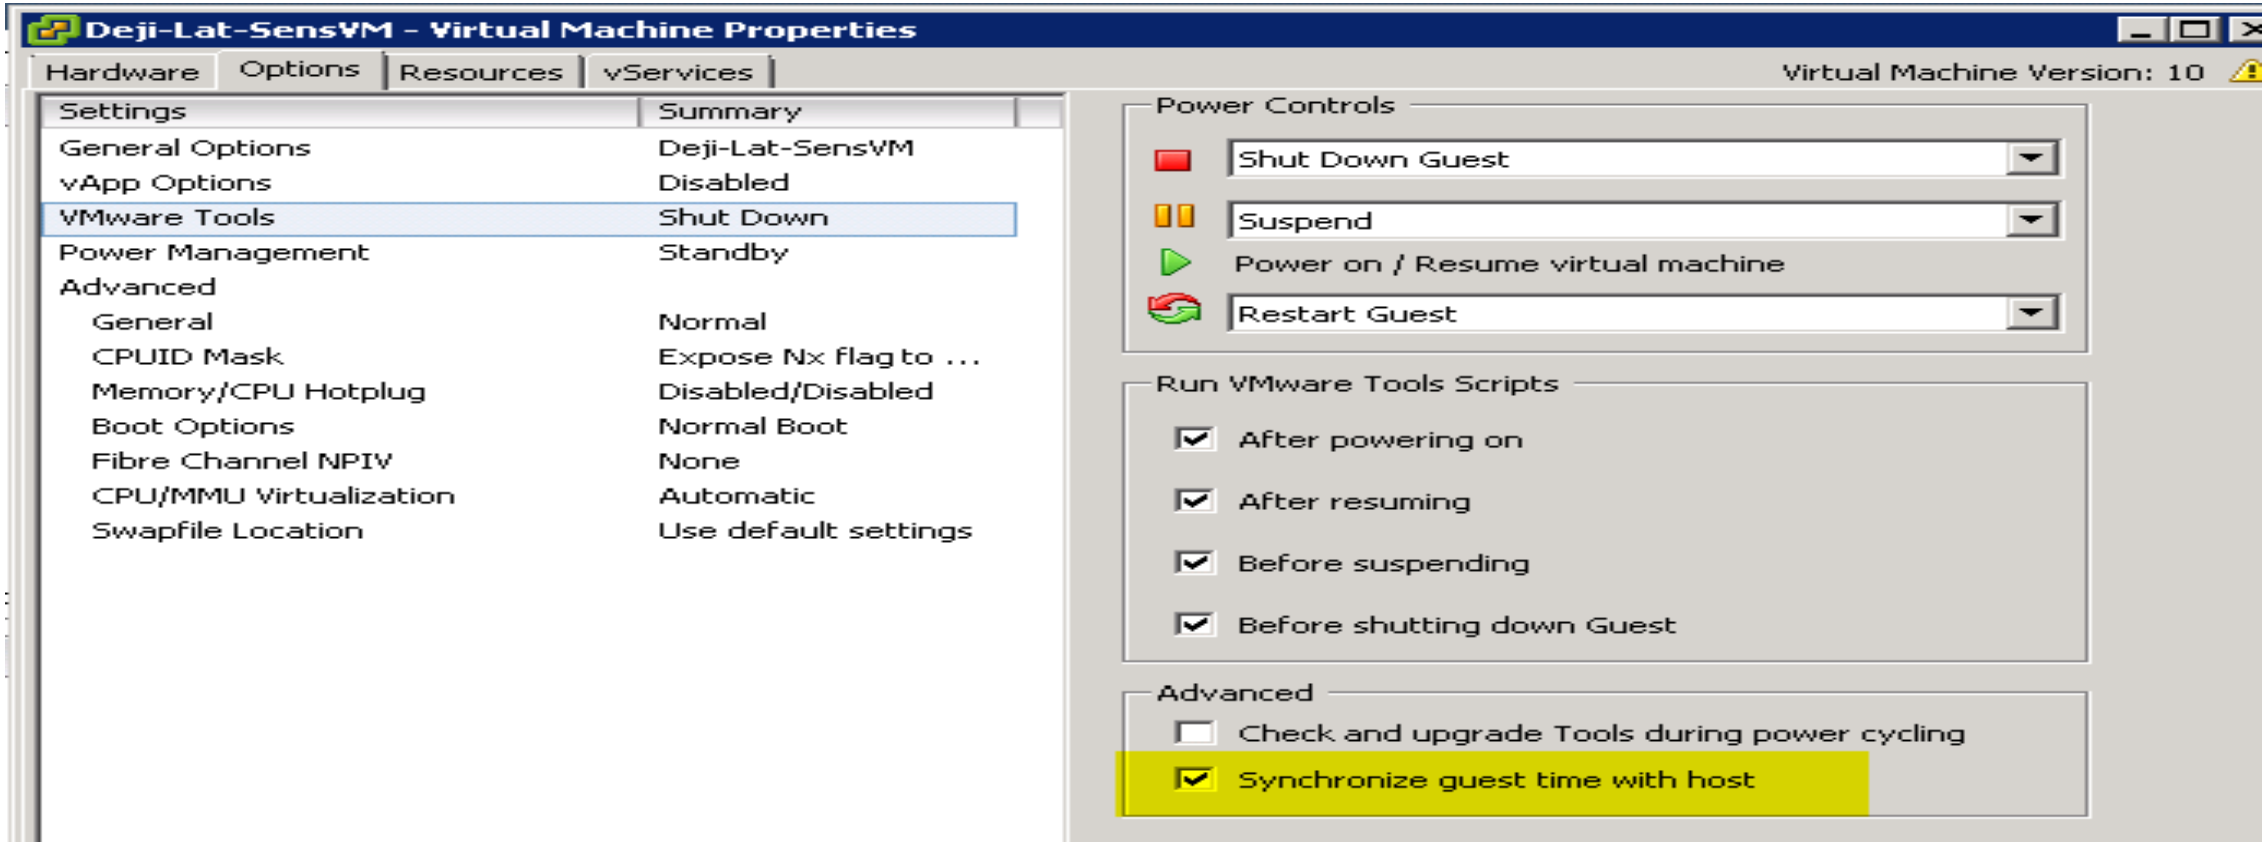

That was Problematic …..

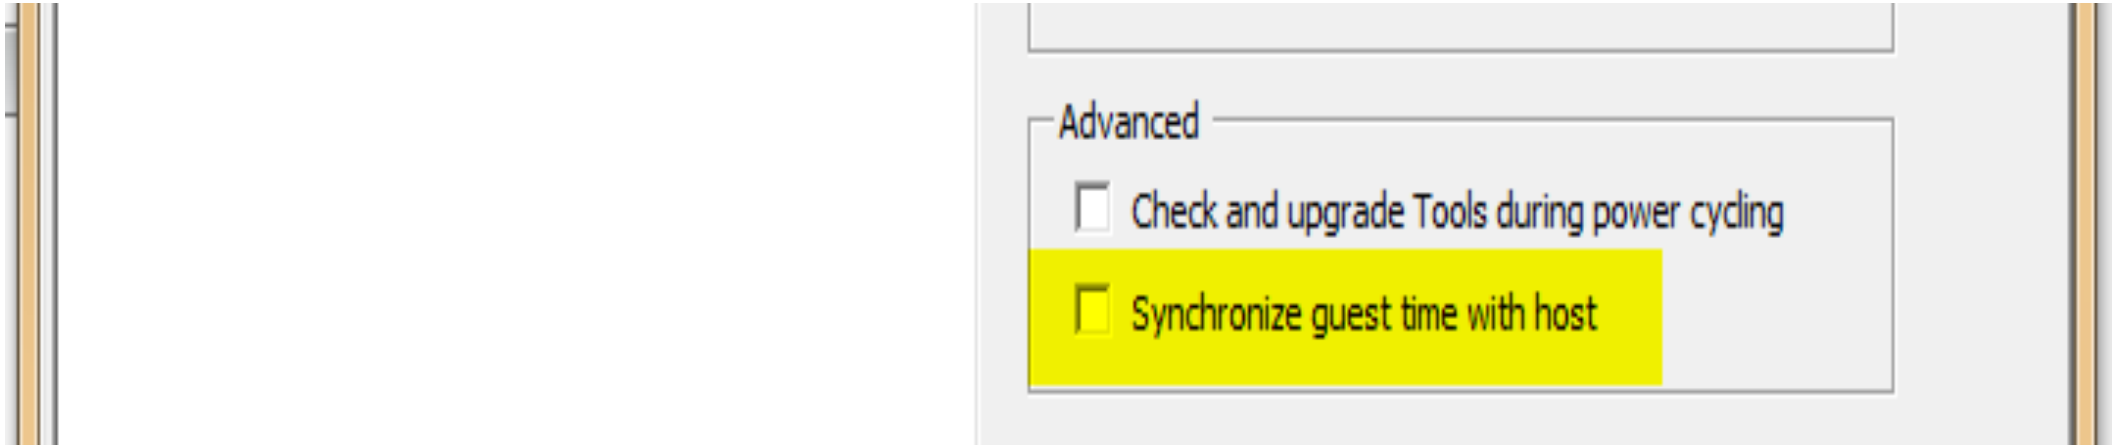

But, That, Too, Is Insufficient

Because Even When You Do THAT, We Still Do THIS

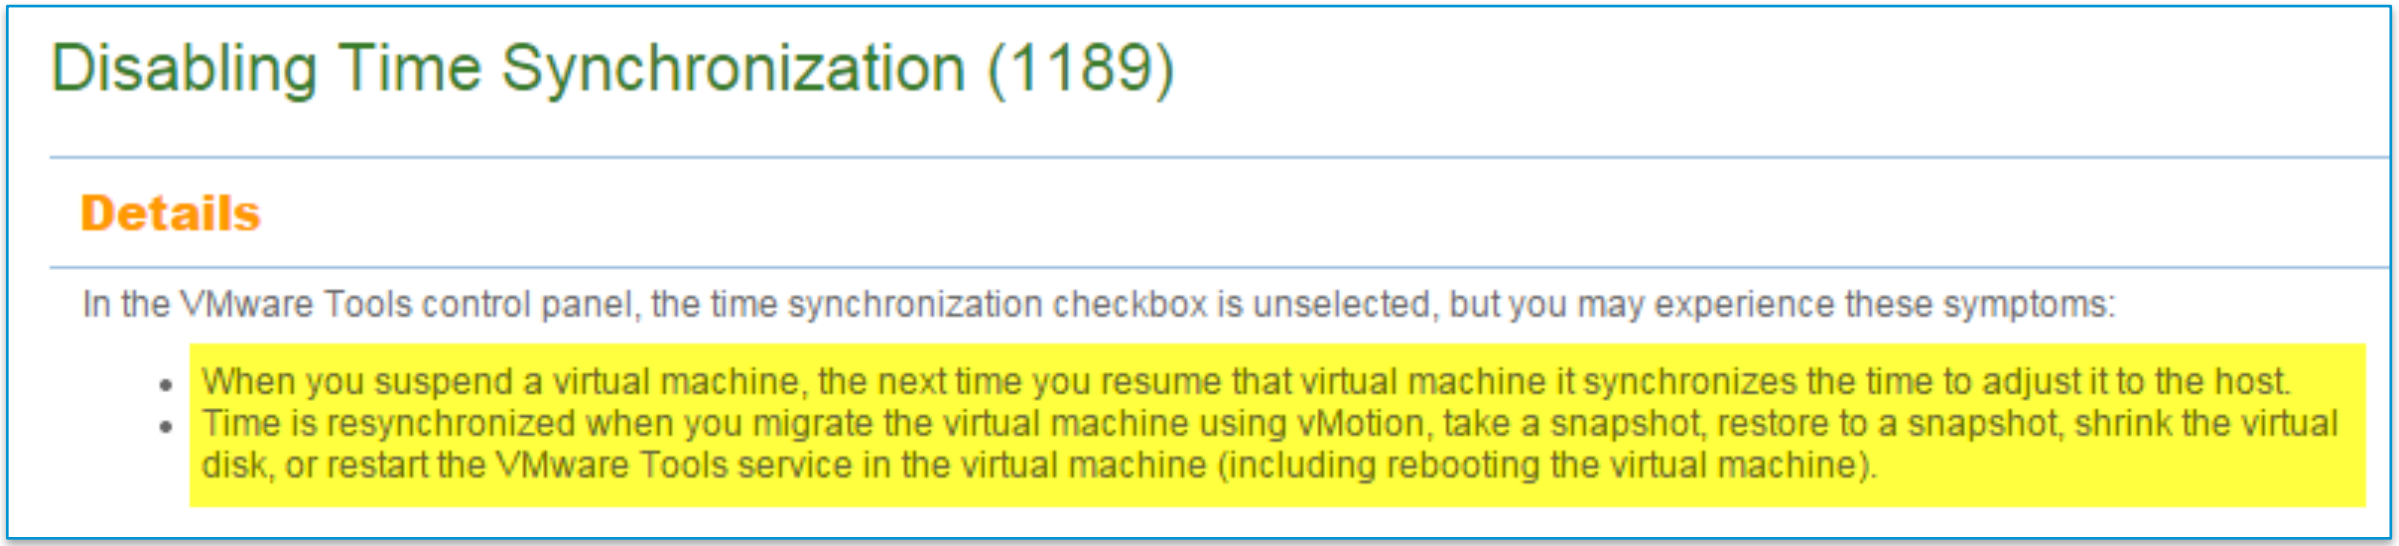

Disabling Time Synchronization (1189)

Details

In the VMware Tools control panel, the time synchronization checkbox is unselected, but you may experience these symptoms:

- When you suspend a virtual machine, the next time you resume that virtual machine it synchronizes the time to adjust it to the host.
- Time is resynchronized when you migrate the virtual machine using vMotion, take a snapshot, restore to a snapshot, shrink the virtual disk, or restart the VMware Tools service in the virtual machine (including rebooting the virtual machine).

Reference: http://kb.vmware.com/kb/1189

Preventing Bad Time Sync

- Ensure Hardware Clock on ESXi Hosts is **CORRECT**
- Configure Reliable NTP on ALL ESXi Hosts
- Configure in-Guest NTP Source
- IF Internal Authoritative Time Source Virtualized
  - (e.g.) Windows Active Directory PDC
  - Disable DRS for the VM
  - Use Host-Guest Affinity Rule for the VM
    - Helps you find it in Emergency

Completely Disabling Time Sync

Add the Following VM's Advanced Configuration Options to your VMs/Templates

tools.syncTime = "0"
time.synchronize.continue = "0"
time.synchronize.restore = "0"
time.synchronize.resume.disk = "0"
time.synchronize.shrink = "0"
time.synchronize.tools.startup = "0"
time.synchronize.tools.enable = "0"
time.synchronize.resume.host = "0"

To add these settings across multiple VMs at once, use VMware vRealize Orchestrator:

http://blogs.vmware.com/apps/2016/01/completely-disable-time-synchronization-for-your-vm.html

Designing for Performance

- NUMA
  - To enable or to not enable? Depends on the Workloads
  - More on NUMA later
- Sockets, Cores and Threads
  - Enable Hyper-threading
  - Size to physical cores, not logical hyper-threaded cores.
- Reservation, Limits, Shares and Resource Pools
  - Use reservation to guarantee resources – IF mixing workloads in clusters
  - Use limits CAREFULLY for non-critical workloads
    - Limits must never be less than Allocated Values *
  - Use Shares on Resource Pools
    - Only to contain non-critical Workload's consumption rate
    - Resource Pools must be continuously managed and reviewed
    - Avoid nesting Resource Pools – complicates capacity planning
- *Only possible with scripted deployment

Designing for Performance

- **Network**
  - Use VMXNET3 Drivers
    - VMXNET3 Template Issues in Windows 2008 R2 - kb.vmware.com\kb\1020078
    - Hotfix for Windows 2008 R2 VMs - http://support.microsoft.com/kb/2344941
    - Hotfix for Windows 2008 R2 SP1 VMs - http://support.microsoft.com/kb/2550978
    - Remember Microsoft's "Convenience Update"? https://support.microsoft.com/en-us/kb/3125574
  - Disable interrupt coalescing – at vNIC level
  - On 1GB network, use dedicated physical NIC for different traffic type

- **Storage**
  - Latency is king - Queue Depths exist at multiple paths (Datastore, vSCSI, HBA, Array)
  - Adhere to storage vendor's recommended multi-pathing policy
  - Use multiple vSCSI controllers, distribute VMDKS evenly
  - Disk format and snapshot
  - Smaller or larger datastores?
    - Determined by storage platform and workload characteristics (VVOL is the future)
  - IP Storage? - Jumbo Frames, if supported by physical network devices

The more you know…

It's the Storage, Stupid

- There is ALWAYS a Queue
- One-lane highway vs 4-Lane highway. More is better
- PVSCSI for all data ask volumes
- Ask Your Storage Vendor about multi-pathing policy

More is NOT Better

- Know your hardware NUMA boundary. Use it to guide your sizing
- Beware of the memory tax
- Beware of CPU fairness
- There is no place like 127.0.0.1 (VM's Home Node)

Don't Blame the vNIC

- VMXNET3 is NOT the problem
- Outdated VMware Tools **MAY** be the problem
- Check in-guest network tuning options – e.g. RSS
- Consider Disabling Interrupt Coalescing

Use Your Tools

- Virtualizing does NOT change OS/App administrative tasks
- ESXTop – Native to ESXi
- *Visualesxtop - https://labs.vmware.com/flings/visualesxtop*
- *Esxplot - https://labs.vmware.com/flings/esxplot*

Storage Optimization

Factors affecting storage performance

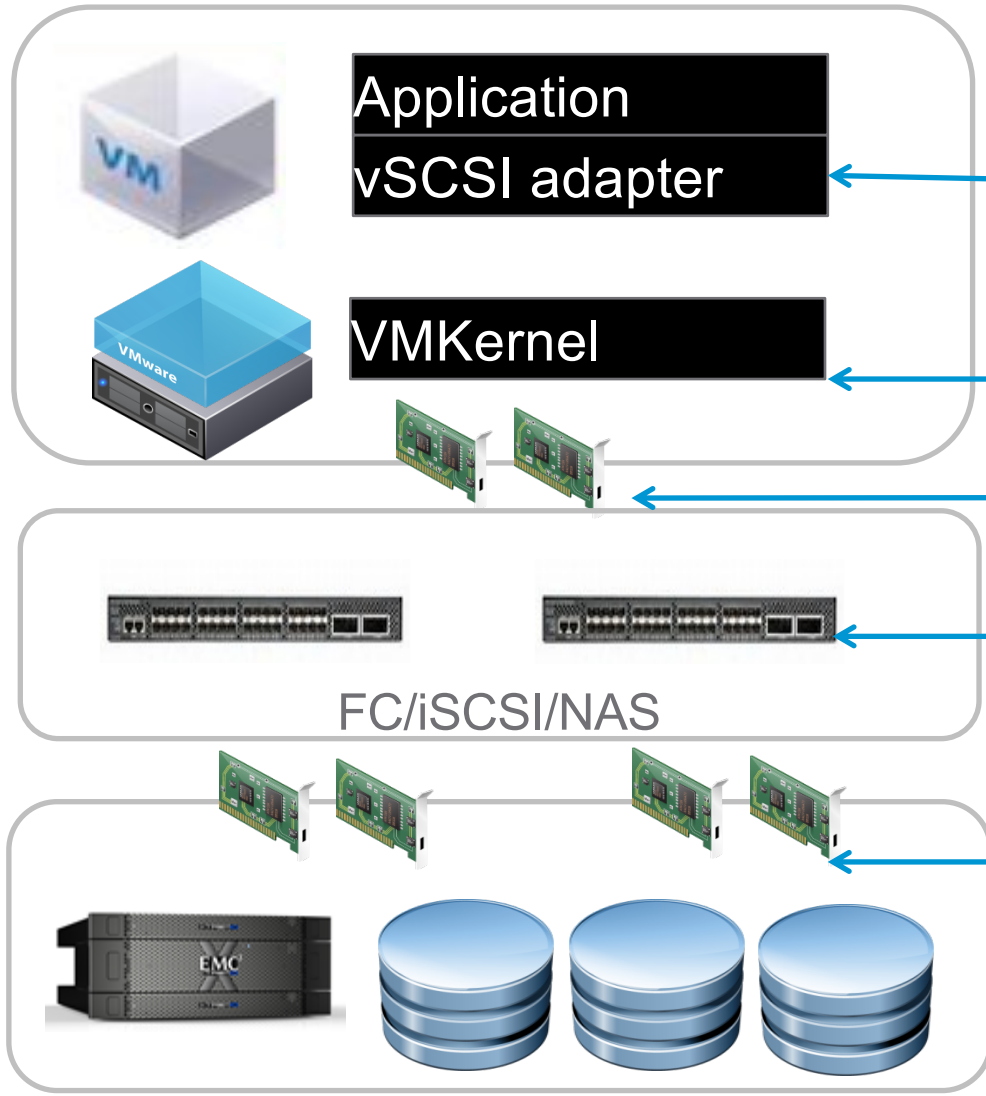

Adapter type
Number of virtual disks
Virtual adapter queue depth

VMKernel admittance (Disk.SchedNumReqOutstanding)

Per path queue depth
Adapter queue depth

Storage network (link speed, zoning, subnetting)

LUN queue depth
Array SPs
HBA target queues
Number of disks (spindles)

Nobody Likes Long Queues





Arriving Customers

Queue

Checkout

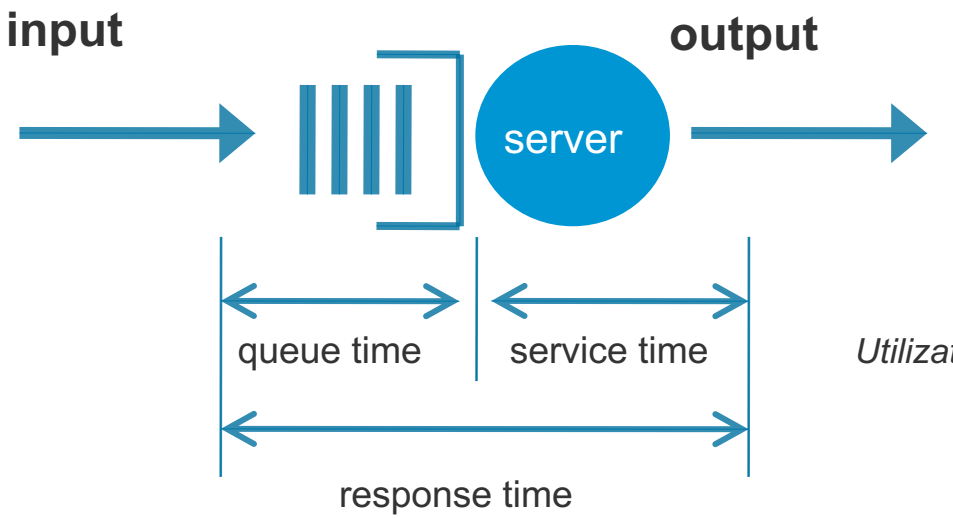

Utilization = busy-time at server / time elapsed

Additional vSCSI controllers improves concurrency

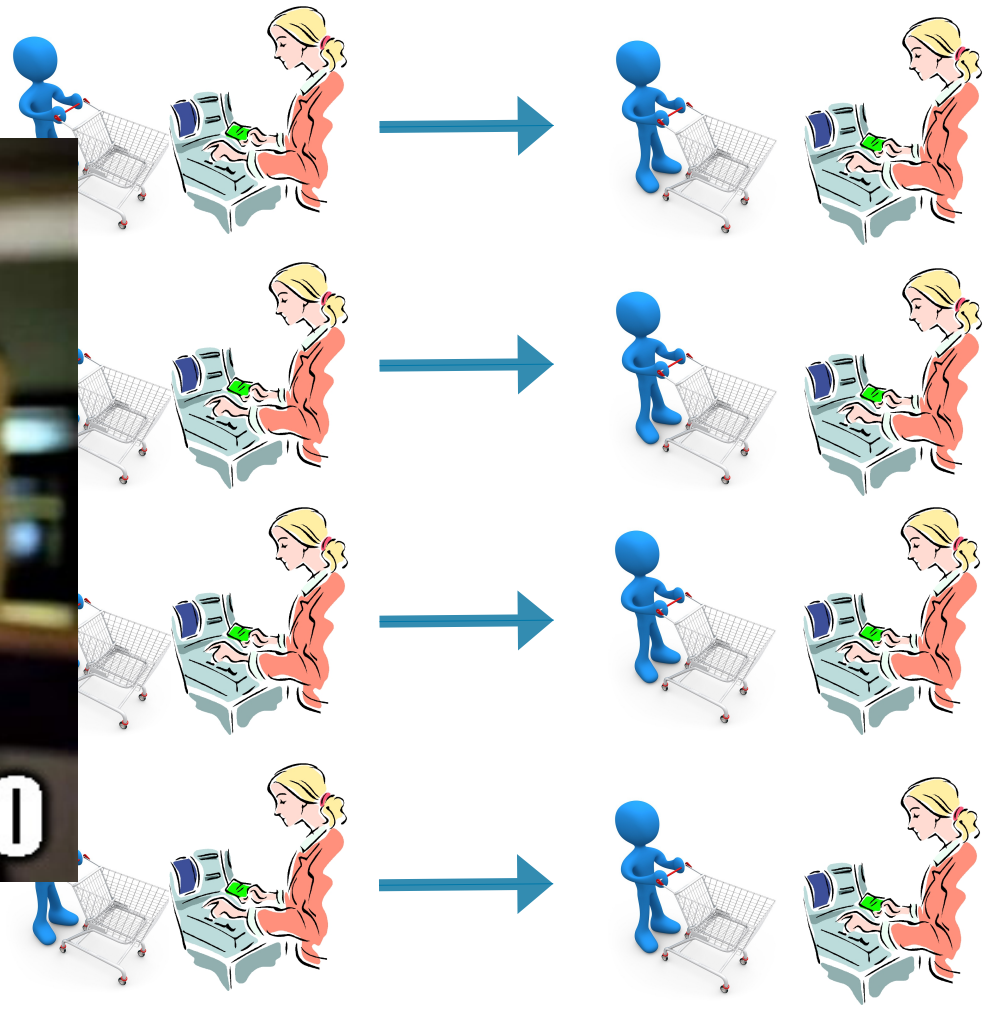

vSCSI Device

Storage Subsystem

Optimize for Performance – Queue Depth

- vSCSI Adapter
  - Be aware of per device/adapter queue depth maximums (KB 1267)
  - Use multiple PVSCSI adapters
- VMKernel Admittance
  - VMKernel admittance policy affecting shared datastore (KB 1268), use dedicated datastores for DB and Log Volumes
  - VMKernel admittance changes dynamically when SIOC is enabled (may be used to control IOs for lower-tiered VMs)
- Physical HBAs
  - Follow vendor recommendation on max queue depth per LUN (http://kb.vmware.com/kb/1267)
  - Follow vendor recommendation on HBA execution throttle
  - Be aware settings are global if host is connected to multiple storage arrays
  - Ensure cards are installed in slots with enough bandwidth to support their expected throughput
  - Pick the right multi-pathing policy based on vendor storage array design (ask your storage vendor)

Increase PVSCSI Queue Depth

- Just increasing LUN, HBA queue depths is NOT ENOUGH
- PVSCSI - http://KB.vmware.com/kb/2053145

- Increase PVSCSI Default Queue Depth (after consultation with array vendor)
  - **Linux:**
    - Add following line to /etc/modprobe.d/ or /etc/modprobe.conf file:
    - options vmw_pvscsi cmd_per_lun=254 ring_pages=32
  - OR, append these to the appropriate kernel boot arguments (*grub.conf or grub.cfg*)
    - vmw_pvscsi.cmd_per_lun=254
    - vmw_pvscsi.ring_pages=32

  - **Windows:**
    - **Key:** HKLM\SYSTEM\CurrentControlSet\services\pvscsi\Parameters\Device
      - **Value: DriverParameter | Value Data: "RequestRingPages=32,MaxQueueDepth=254“**

Optimize for Performance – Storage Network

- Link Type/Speed
  - FC vs. iSCSI vs. NAS
  - Latency suffers when bandwidth is saturated
- Zoning and Subnetting
  - Place hosts and storage on the same switch, minimize Inter-Switch Links
  - Use 1:1 initiator to target zoning or follow vendor recommendation
  - Enable jumbo frame for IP based storage (MTU needs to be set on all connected physical and virtual devices)
  - Make sure different iSCSI IP subnets cannot transmit traffic between them

“Thick” vs “Thin”

- Thin (Fully Inflated and Zeroed) Disk Performance = Thick Eager Zero Disk
- Performance impact due to zeroing, not result of allocation of new blocks
- To get maximum performance from the start, must use Thick Eager Zero Disks (think Business Critical Apps)
- Maximum Performance happens eventually, but when using lazy zeroing, zeroing needs to occur before you can get maximum performance

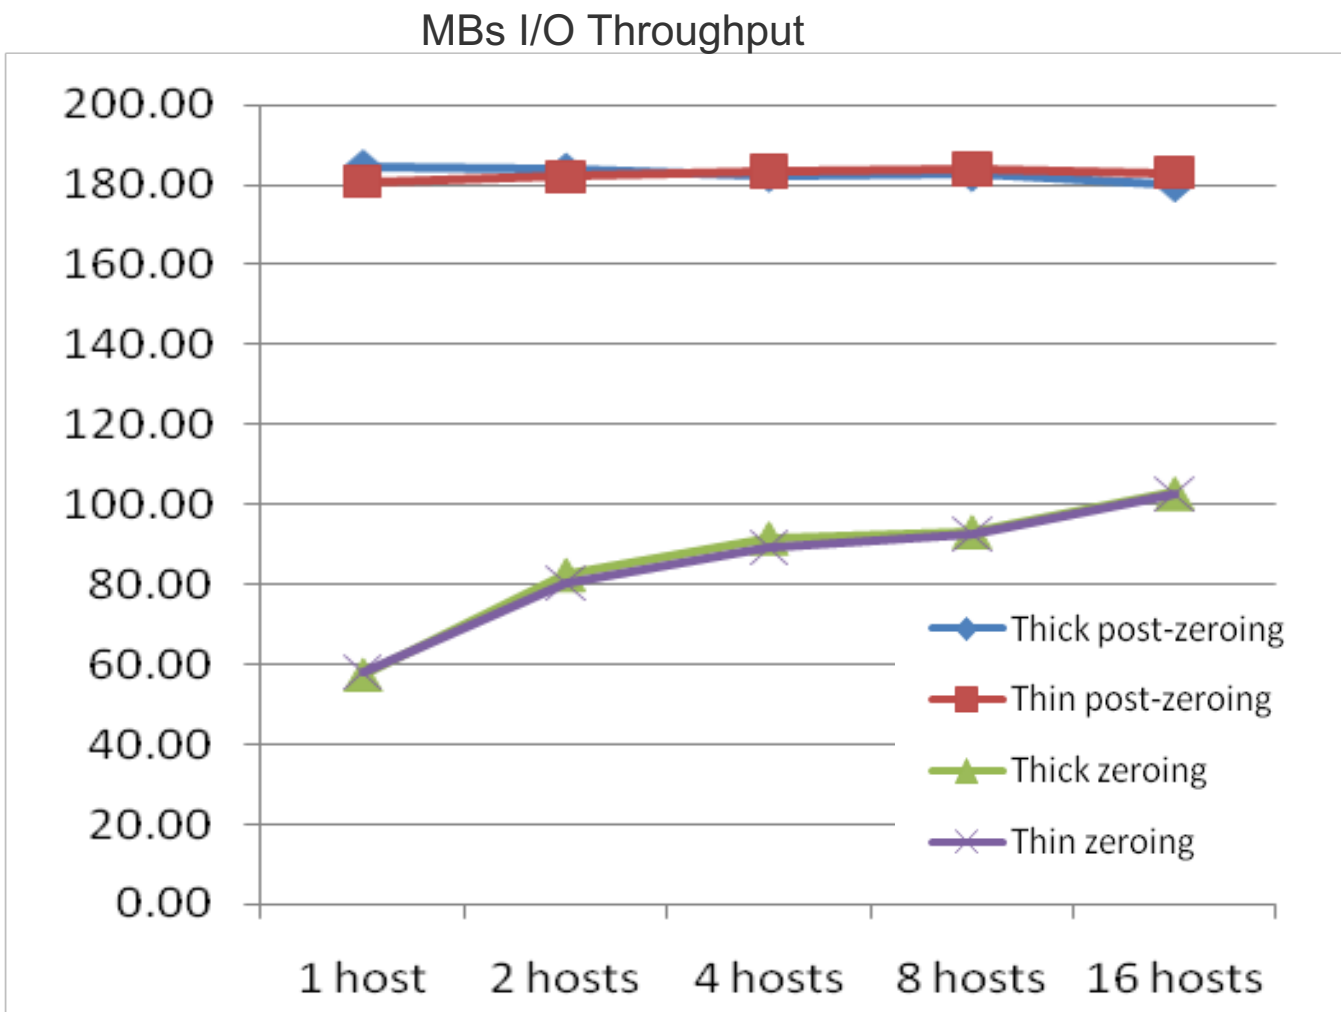

Choose Storage which supports VMware vStorage APIs for Array Integration (VAAI)

VMFS or RDM?

- Generally similar performance http://www.vmware.com/files/pdf/performance_char_vmfs_rdm.pdf
- vSphere 5.5 and later support up to 62TB VMDK files
- Disk size no longer a limitation of VMFS

VMFS	RDM
Better storage consolidation – multiple virtual disks/virtual machines per VMFS LUN. But still can assign one virtual machine per LUN	Enforces 1:1 mapping between virtual machine and LUN
Consolidating virtual machines in LUN – less likely to reach vSphere LUN Limit of 255	More likely to hit vSphere LUN limit of 255
Manage performance – combined IOPS of all virtual machines in LUN < IOPS rating of LUN	Not impacted by IOPS of other virtual machines

- When to use raw device mapping (RDM)
  - Required for shared-disk failover clustering
  - Required by storage vendor for SAN management tools such as backup and snapshots
- Otherwise use VMFS

Example Best Practices for VM Disk Layout (Microsoft SQL Server)

Characteristics:

- OS on shared DataStore/LUN
- 1 database; 4 equally-sized data files across 4 LUNs
- 1 TempDB; 4 (1/vCPU) equally-sized tempdb files across 4 LUNs
- Data, TempDB, and Log files spread across 3 PVSCSI adapters
  - Data and TempDB files share PVSCSI adapters
- ***Virtual Disks could be RDMs***

Advantages:

- Optimal performance; each Data, TempDB, and Log file has a dedicated VMDK/Data Store/LUN
- I/O spread evenly across PVSCSI adapters
- Log traffic does not contend with random Data/TempDB traffic

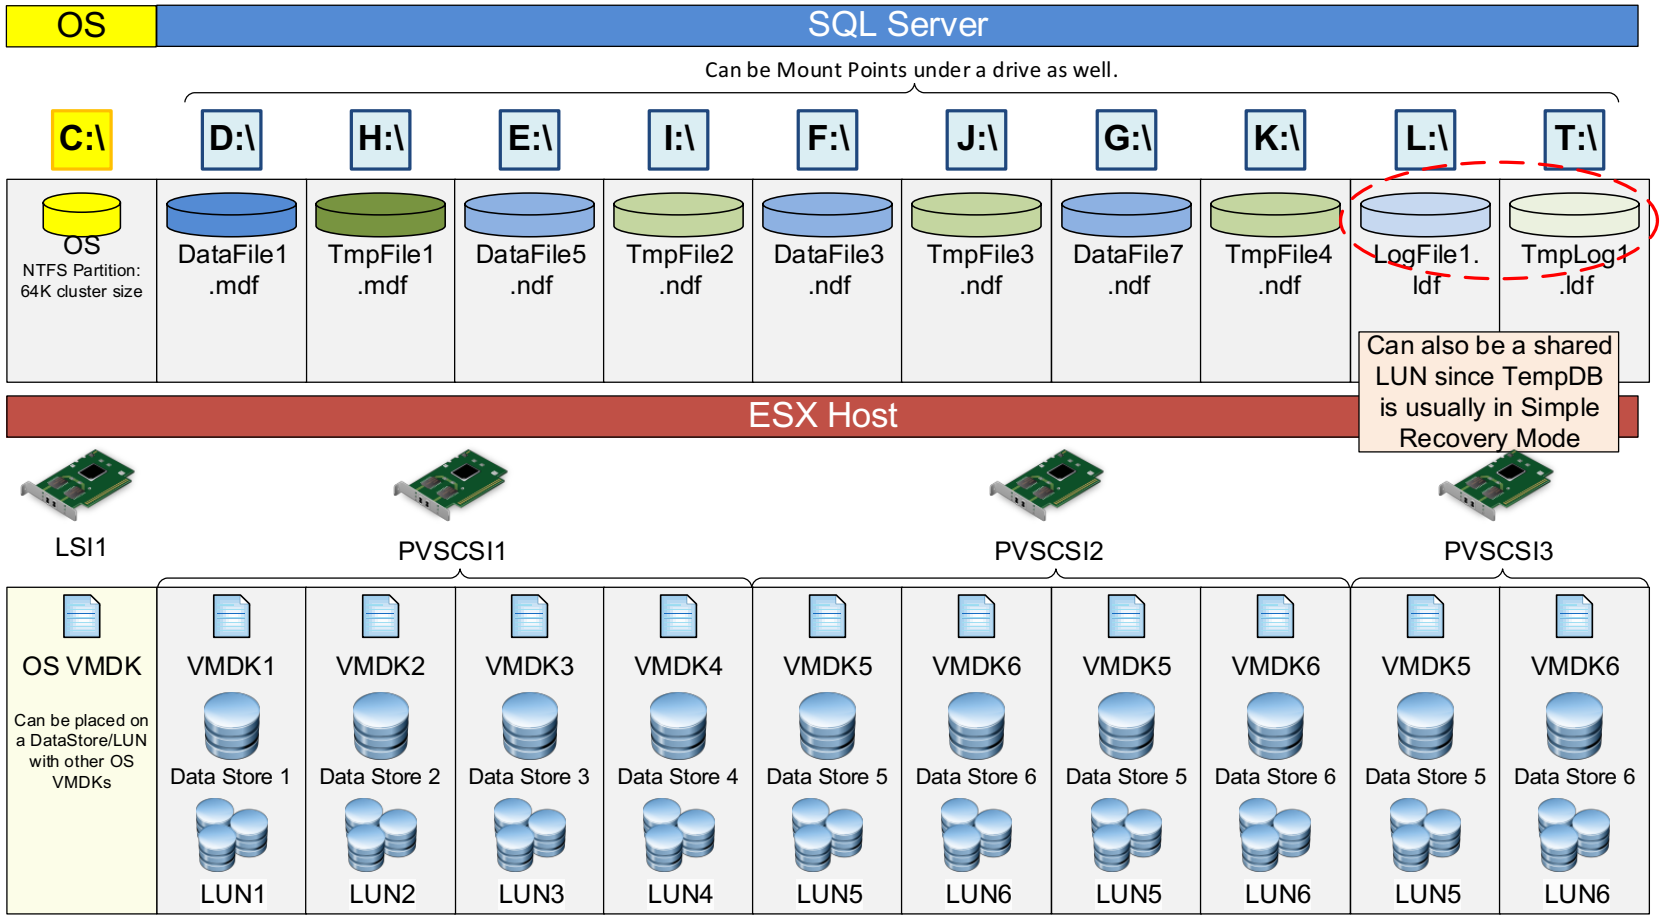

Disadvantages:

- You can quickly run out of Windows driver letters!
- More complicated storage management

Realistic VM Disk Layout (Microsoft SQL Server)

Characteristics:

- OS on shared DataStore/LUN
- 1 database; 8 Equally-sized data files across 4 LUNs
- 1 TempDB; 4 files (1/vCPU) evenly distributed and mixed with data files to avoid “hot spots”
- Data, TempDB, and Log files spread across 3 PVSCSI adapters
- ***Virtual Disks could be RDMs***

Advantages:

- Fewer drive letters used
- I/O spread evenly/TempDB hot spots avoided
- Log traffic does not contend with random Data/TempDB traffic

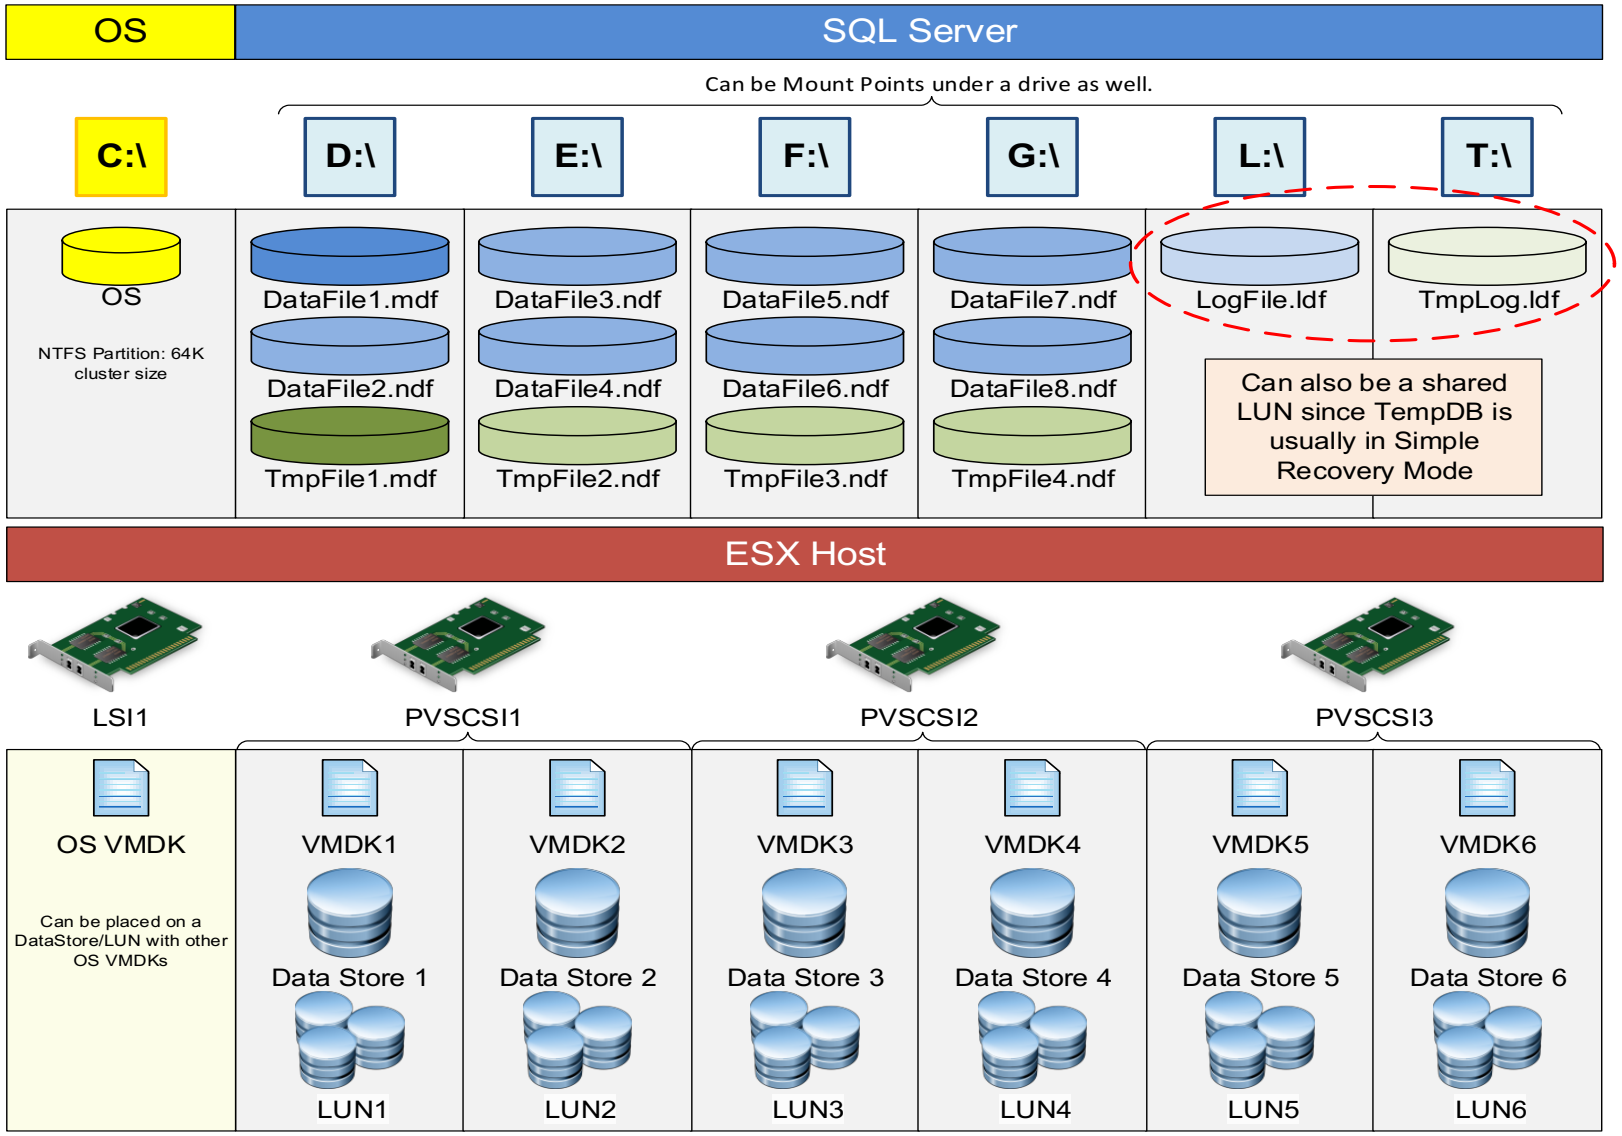

Lets talk about CPU, vCPUs and other Things

Optimizing Performance – Know Your NUMA

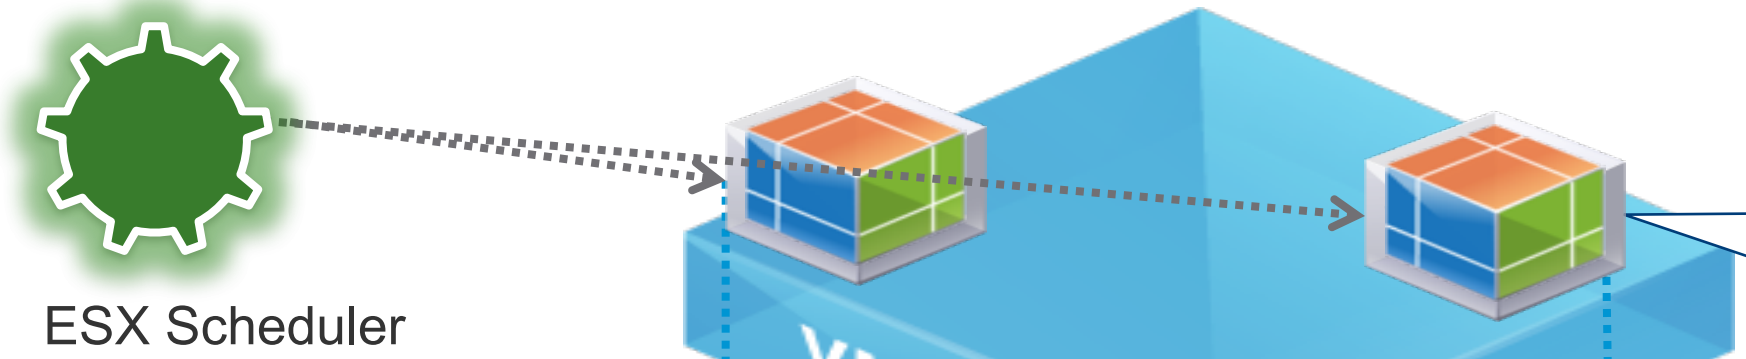

ESX Scheduler

8 vCPU VMs
less than
45GB RAM
on each VM

If VM is sized greater than 45GB or 8 CPUs, Then NUMA interleaving and subsequent migration occur and can cause 30% drop in memory throughput performance

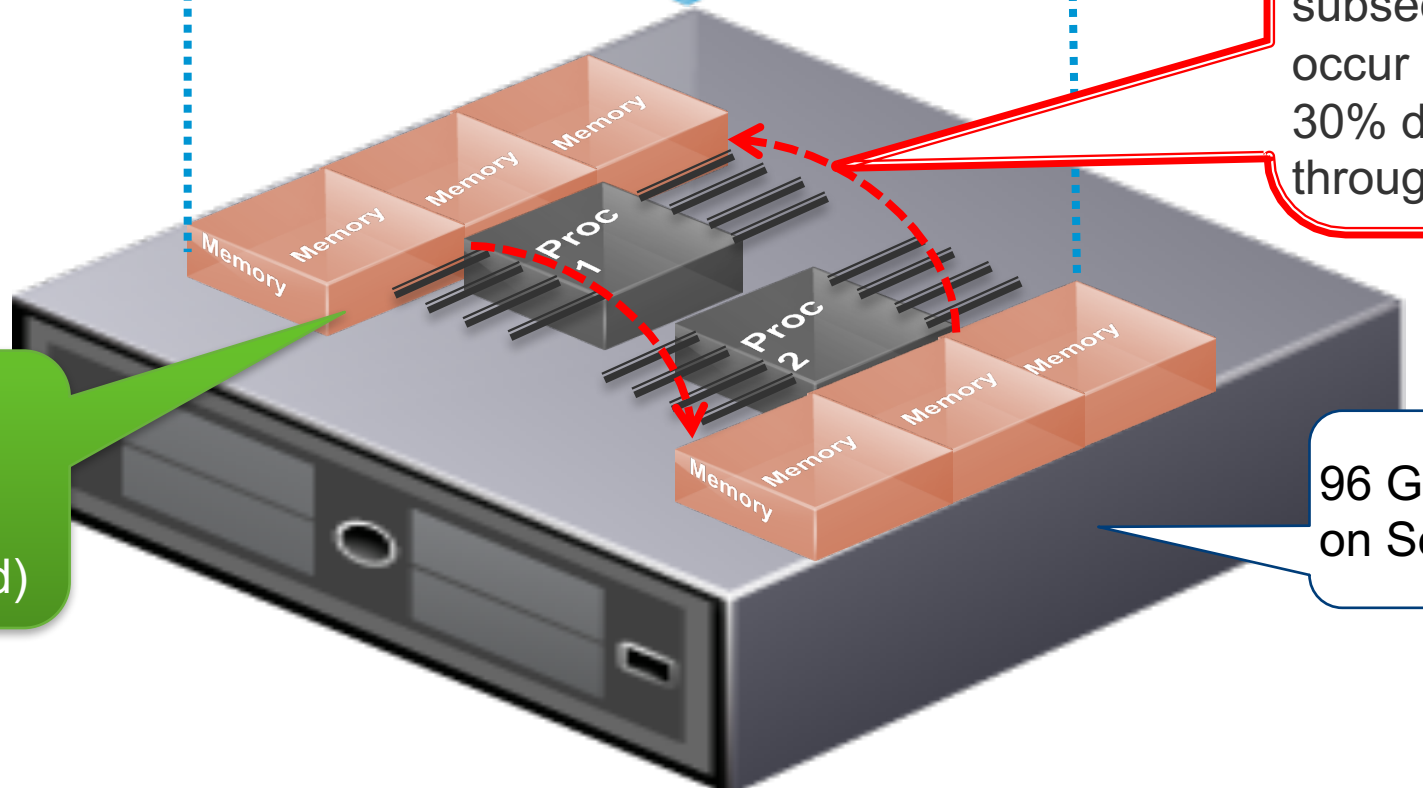

Each NUMA Node has 94/2 45GB (less 4GB for hypervisor overhead)

96 GB RAM
on Server

NUMA Local Memory with Overhead Adjustment

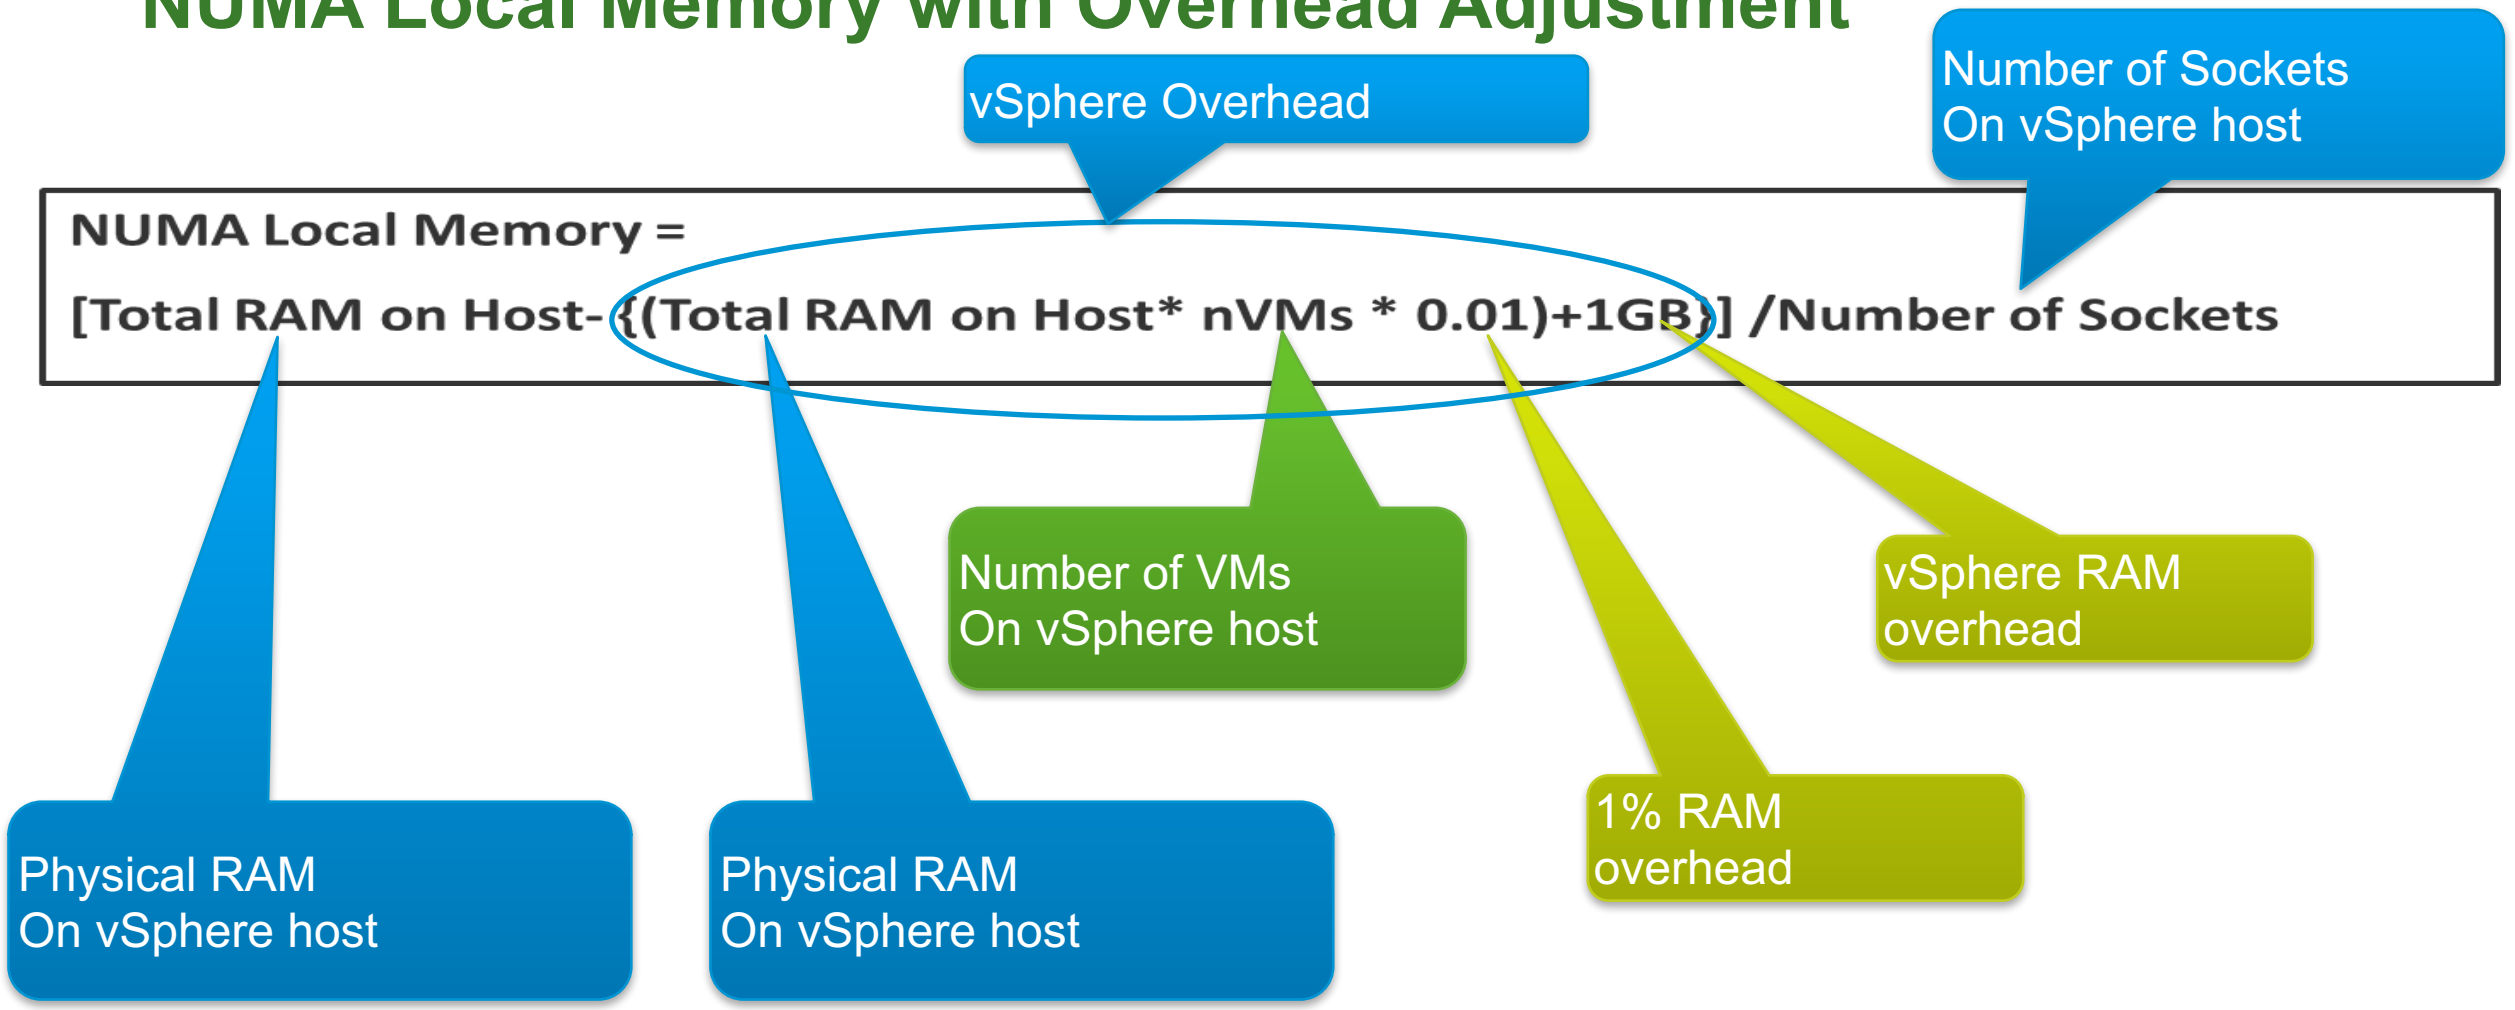

NUMA and vNUMA FAQ!

- Shall we Define NUMA Again? Nah…..
- Why VMware Recommends Enabling NUMA
  - Modern Operating Systems are NUMA-aware
  - Some applications are NUMA-aware (some are not)
  - vSphere Benefits from NUMA
- Use it, People
  - Enable Host-Level NUMA
    - Disable “Node Inter-leaving” in BIOS – on HP Systems
    - Consult Hardware Vendor for SPECIFIC Configuration
- Virtual NUMA
  - Auto-enabled on vSphere for Any VM with 9 or more vCPUs
  - Want to use it on Smaller VMs?
    - Set **“numa.vcpu.min”** to # of vCPUs on the VM
- CPU Hot-Plug **DISABLES** Virtual NUMA
- vSphere 6.5 changes vNUMA config

vSphere 6.5 vCPU Allocation Guidance

Physical Topology	vCPUs Required	VM Configuration		Resulting vNUMA Nodes Presented
		vSockets	vCores	
Hardware: **Intel** • 2 Sockets • 2 pNUMA Nodes • 14 Cores Per Socket • 28 Logical Threads	1	1	1	1
	2	1	2	1
	3	1	3	1
	4	1	4	1
	5	1	5	1
	6	1	6	1
	7	1	7	1
	8	1	8	1
	9	1	9	1
	10	1	10	1
	11	1	11	1
	12	1	12	1
	13	1	13	1
	14	1	14	1
	15	Sub-optimal		
	16	2	8	2
	17	Sub-optimal		
	18	2	9	2
	19	Sub-optimal		
	20	2	10	2
	21	Sub-optimal		
	22	2	11	2
	23	Sub-optimal		
	24	2	12	2
	25	Sub-optimal		
	26	2	13	2
	27	Sub-optimal		
	28	2	14	2

NUMA Best Practices

- Avoid Remote NUMA access
  - Size # of vCPUs to be <= the # of cores on a NUMA node (processor socket)
    - Where possible, align VMs with physical NUMA boundaries
  - For wide VMs, use a multiple or even divisor of NUMA boundaries
  - http://www.vmware.com/files/pdf/techpaper/VMware-vSphere-CPU-Sched-Perf.pdf
- Hyper-threading
  - Initial conservative sizing: set vCPUs equal to # of physical cores
  - HT benefit around 30-50%, < for CPU intensive batch jobs (based on OLTP workload tests)
- Allocate vCPUs by socket count
  - Default “Cores Per Socket” is set to “1”
  - Applicable to vSphere versions prior to 6.5. Not as relevant in 6.5
- ESXTOP to monitor NUMA performance in vSphere
  - Coreinfo.exe to see NUMA topology in Windows Guest
- vMotioning VMs between hosts with dissimilar NUMA topology leads to performance issues

Non-Wide VM Sizing Example (VM fits within NUMA Node)

- 1 vCPU per core with hyper-threading OFF
  - Must license each core for SQL Server
- 1 vCPU per thread with hyper-threading ON
  - 10%-25% gain in processing power
  - Same licensing consideration
  - HT does not alter core-licensing requirements

"numa.vcpu.preferHT" to true to force 24-way VM to be scheduled within NUMA node

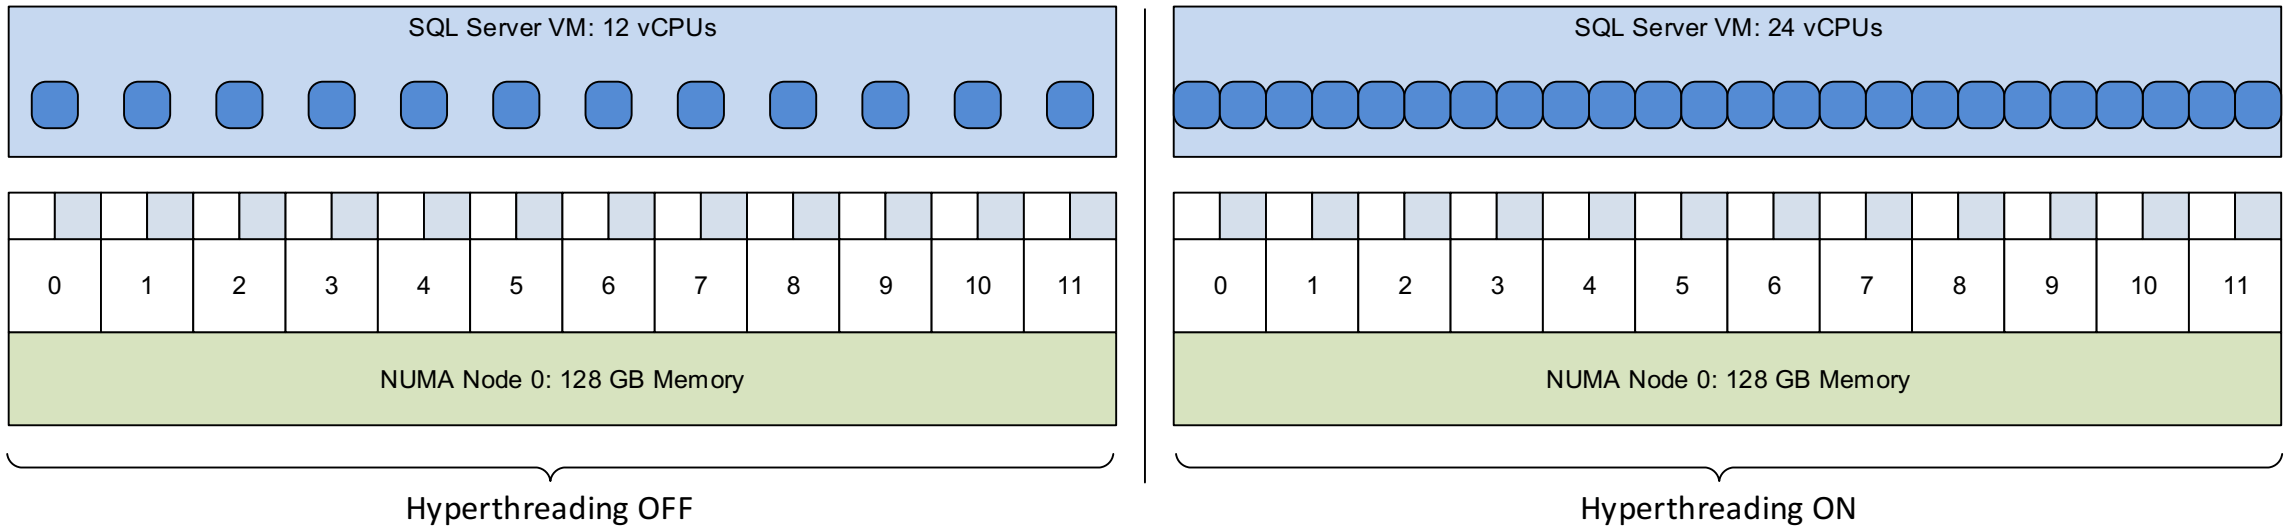

Wide VM Sizing Example (VM crosses NUMA Node)

- Extends NUMA awareness to the guest OS
- Enabled through multicore UI
  - On by default for 9+ vCPU multicore VM
  - Existing VMs are not affected through upgrade
  - For smaller VMs, enable by setting numa.vcpu.min=4
- Do NOT turn on CPU Hot-Add
- For wide virtual machines, confirm feature is on for best performance

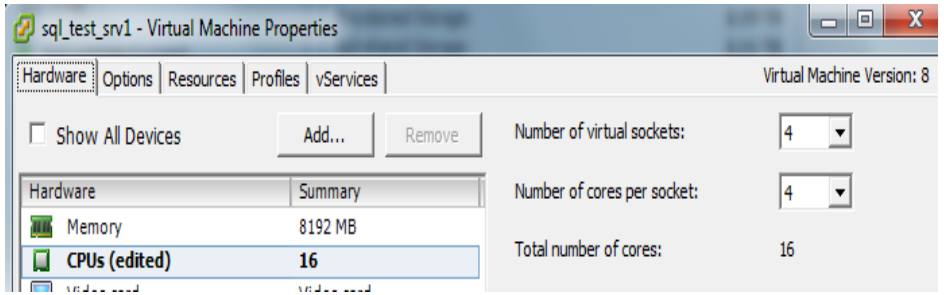

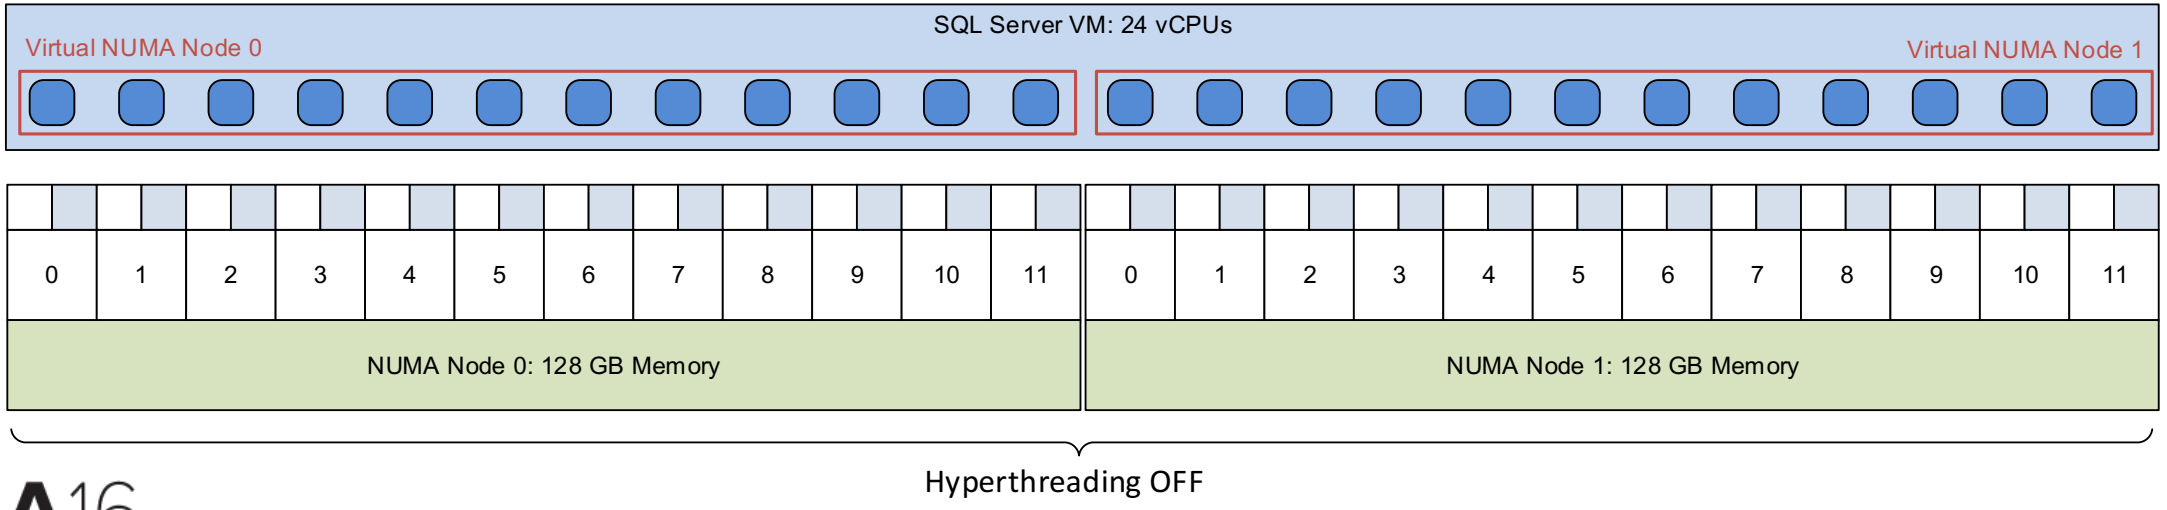

Designing for Performance

- The VM itself matters – In-guest optimization
  - Windows CPU Core Parking = BAD
    - Set Power to “High Performance” to avoid core parking
    - Relevant IF ESXi Host Power Setting NOT “High Performance”
  - Windows Receive Side Scaling settings impact CPU utilization
    - Must be enabled at NIC and Windows Kernel level
      - Use “netsh int tcp show global” to verify
      - More on this later
- Application-level tuning
  - Follow vendor’s recommendation
  - Virtualization does not change the consideration

Why Your Windows App Server Lamborghini Runs Like a Pinto

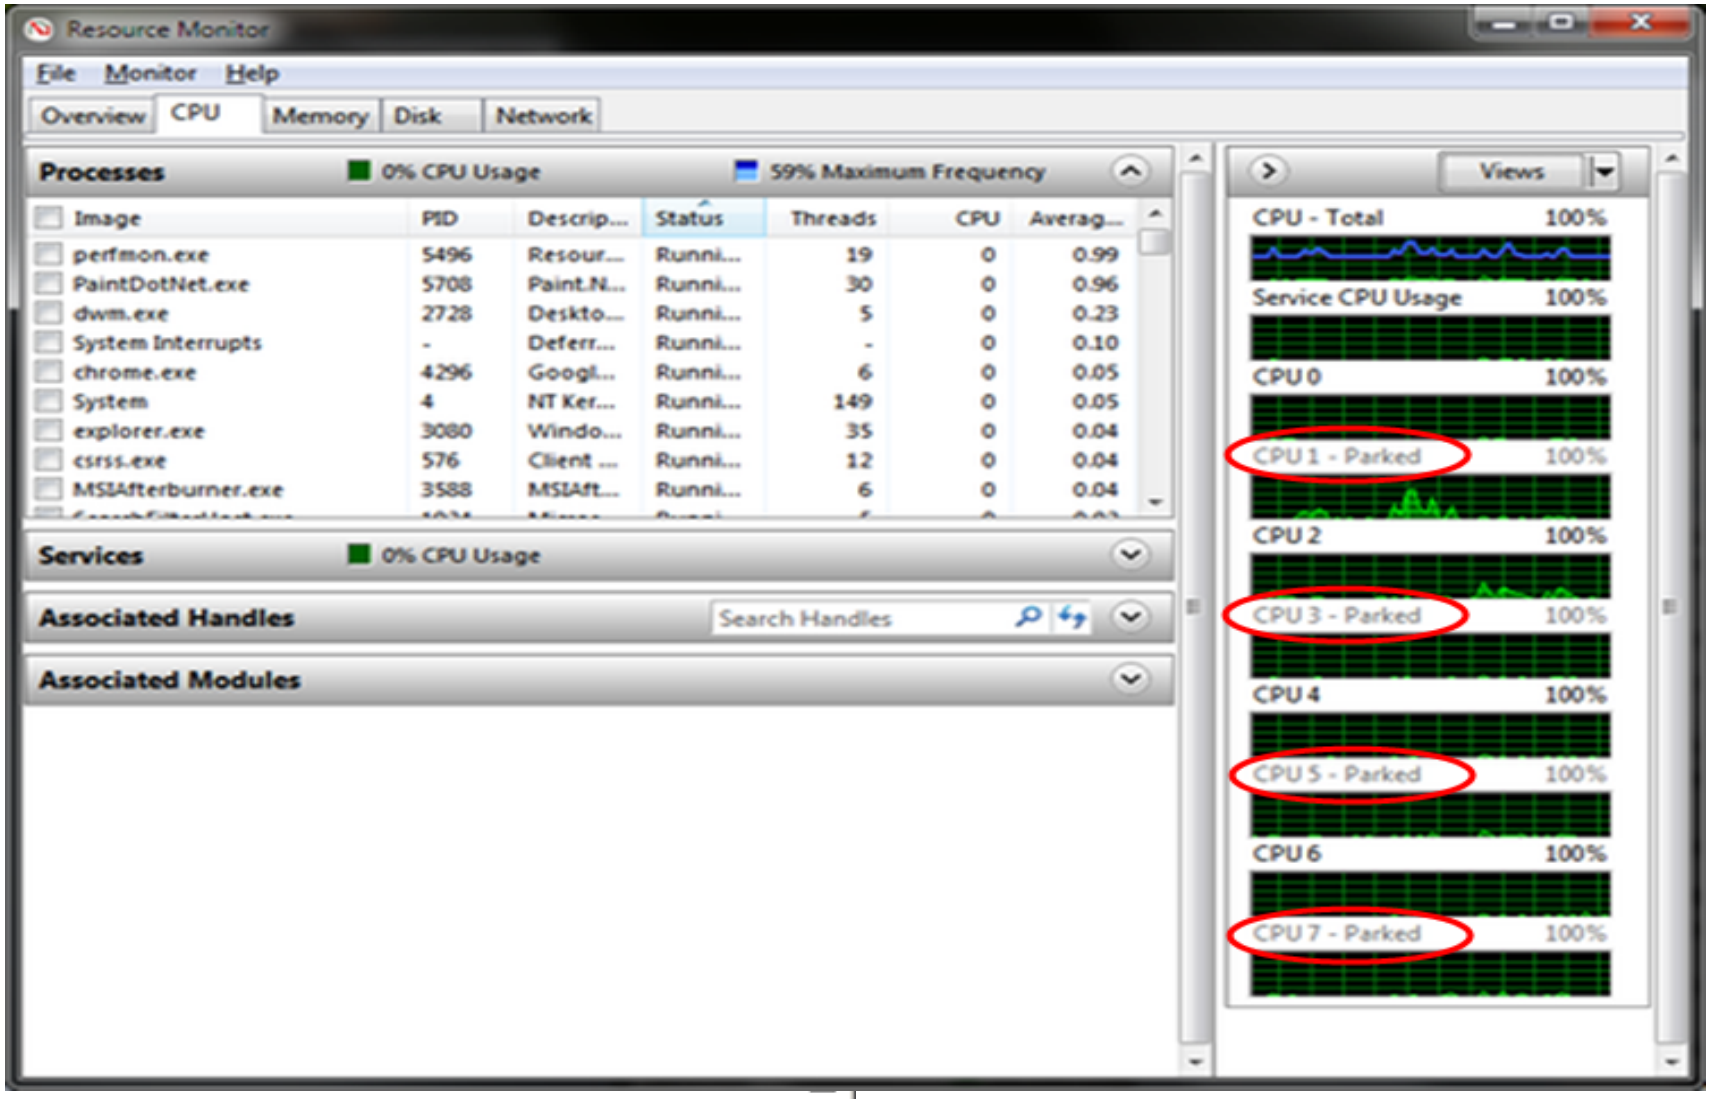

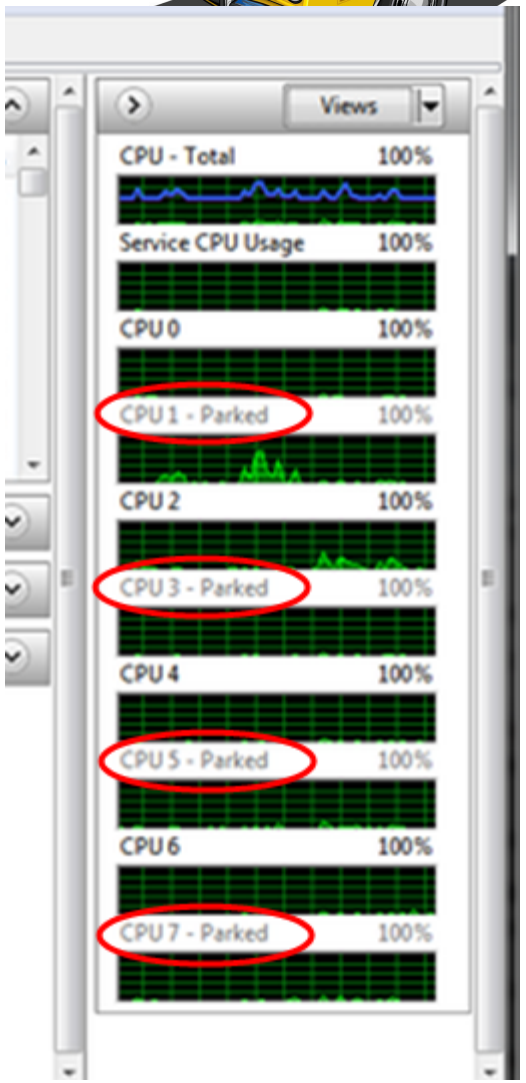

Memory Optimization

Memory Reservations

- Guarantees allocated memory for a VM
- The VM is only allowed to power on if the CPU and memory reservation is available (strict admission)
- If Allocated RAM = Reserved RAM, you avoid swapping
- Do NOT set memory limits for Mission-Critical VMs
- If using Resource Pools, Put Lower-tiered VMs in Resource Pools
- Some Applications Don't Support "Memory Hot-add"
  - E.g. Microsoft Exchange Server CANNOT use Hot-added RAM
- Don't use it on ESXi versions lower than 6.0
- Virtual:Physical memory allocation ratio **should** not exceed 2:1
- Remember NUMA? It's not just about CPU
  - Fetching remote memory is VERY expensive
  - Use "numa.vcpu.maxPerVirtualNode" to control memory locality

What about Dynamic Memory?

- Not Supported by Most Microsoft's Critical Applications
- Not a feature of VMware vSphere

Memory Reservations and Swapping on vSphere

- Setting a reservation creates zero (or near-zero) swap file

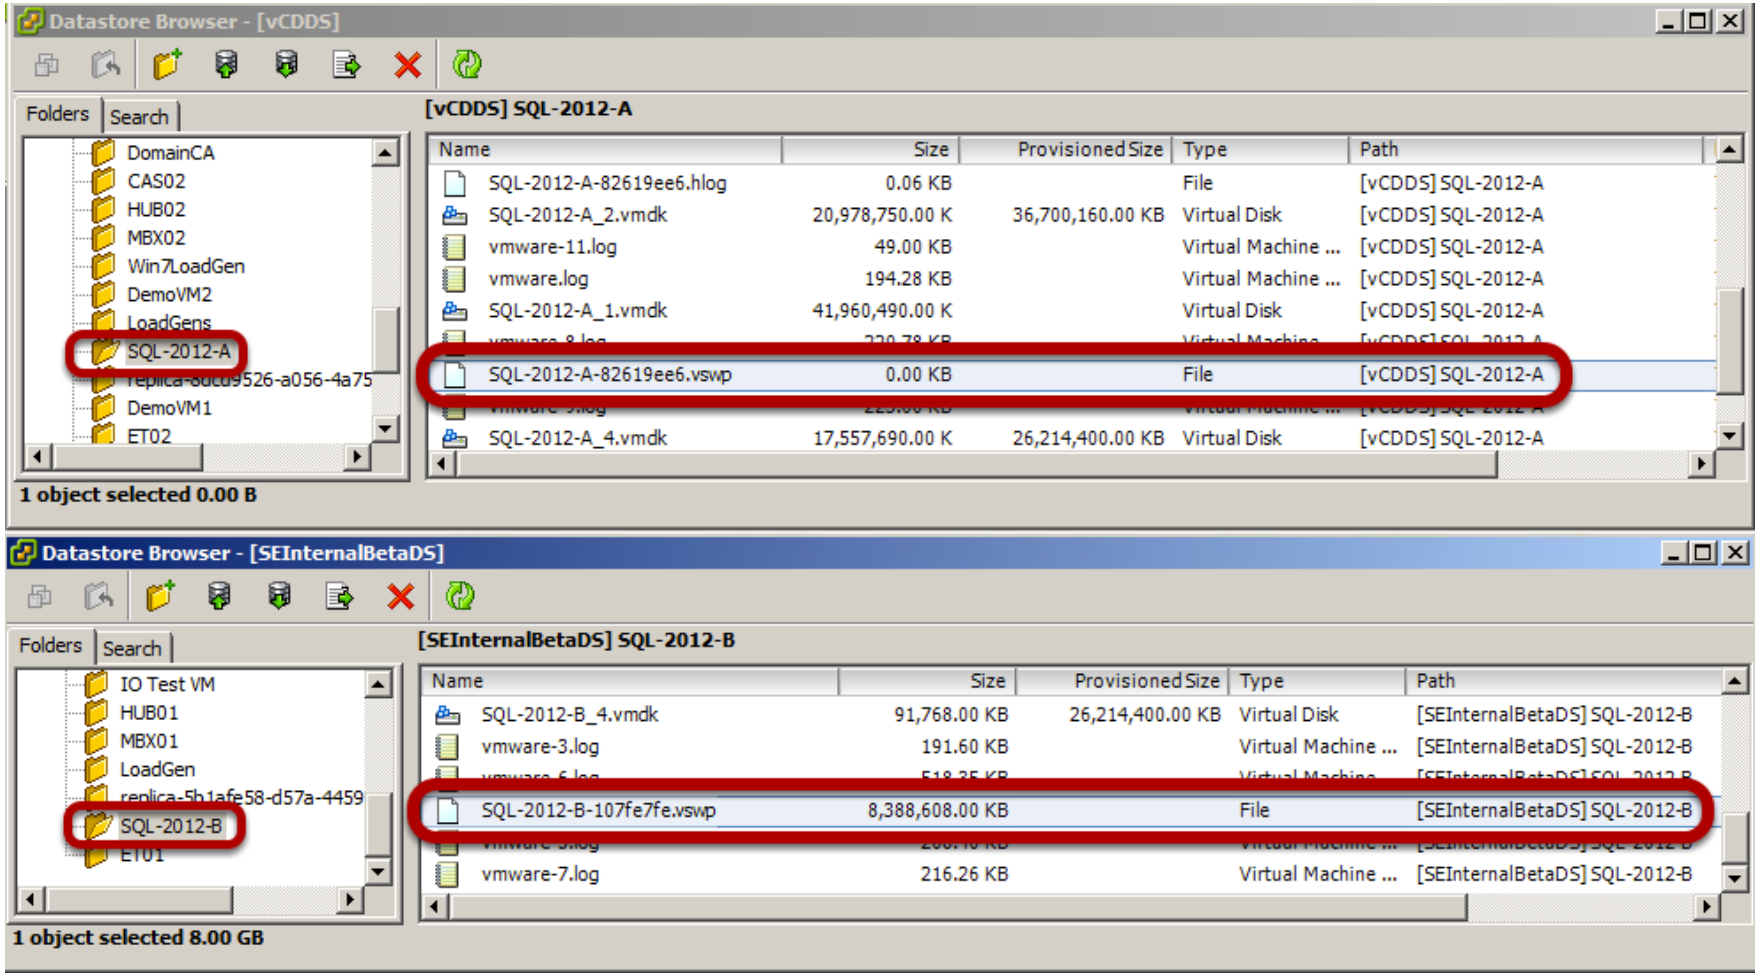

Network Optimization

vSphere Distributed Switch (VDS) Overview

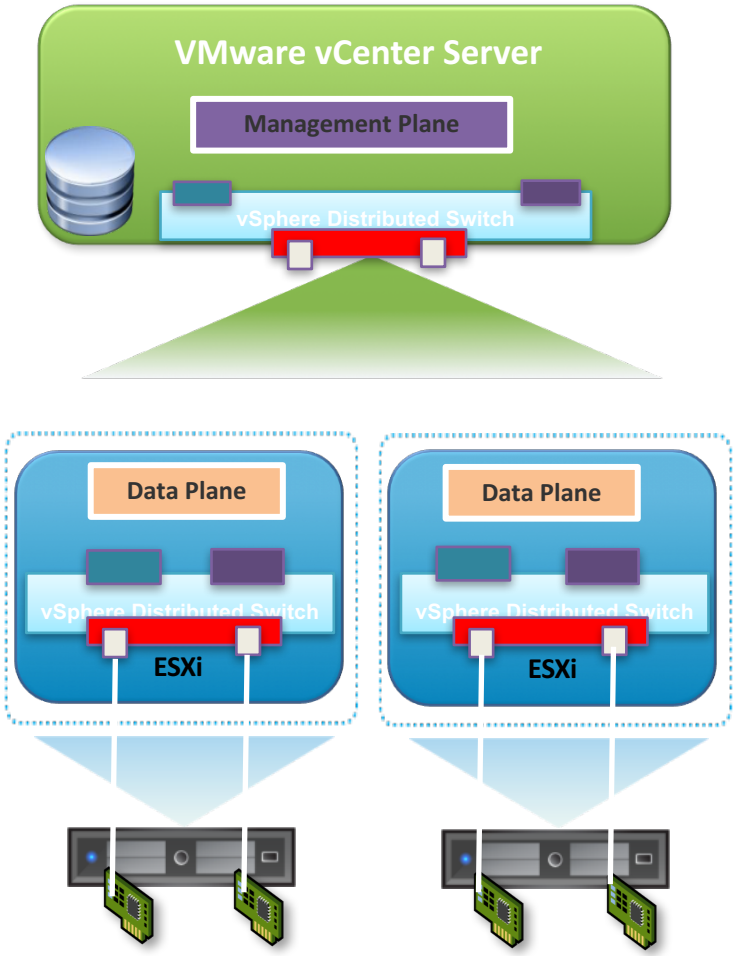

- **Unified network virtualization management**
  - Independent of physical fabric
- **vMotion aware : Statistics and policies follow the VM**
- **vCenter management plane independent of data plane**
- **Advanced Traffic Management features**
  - Load Based Teaming (LBT)
  - Network IO Control (NIOC)
- **Monitoring and Troubleshooting features**
  - NetFlow
  - Port Mirroring

Common Network Misconfiguration

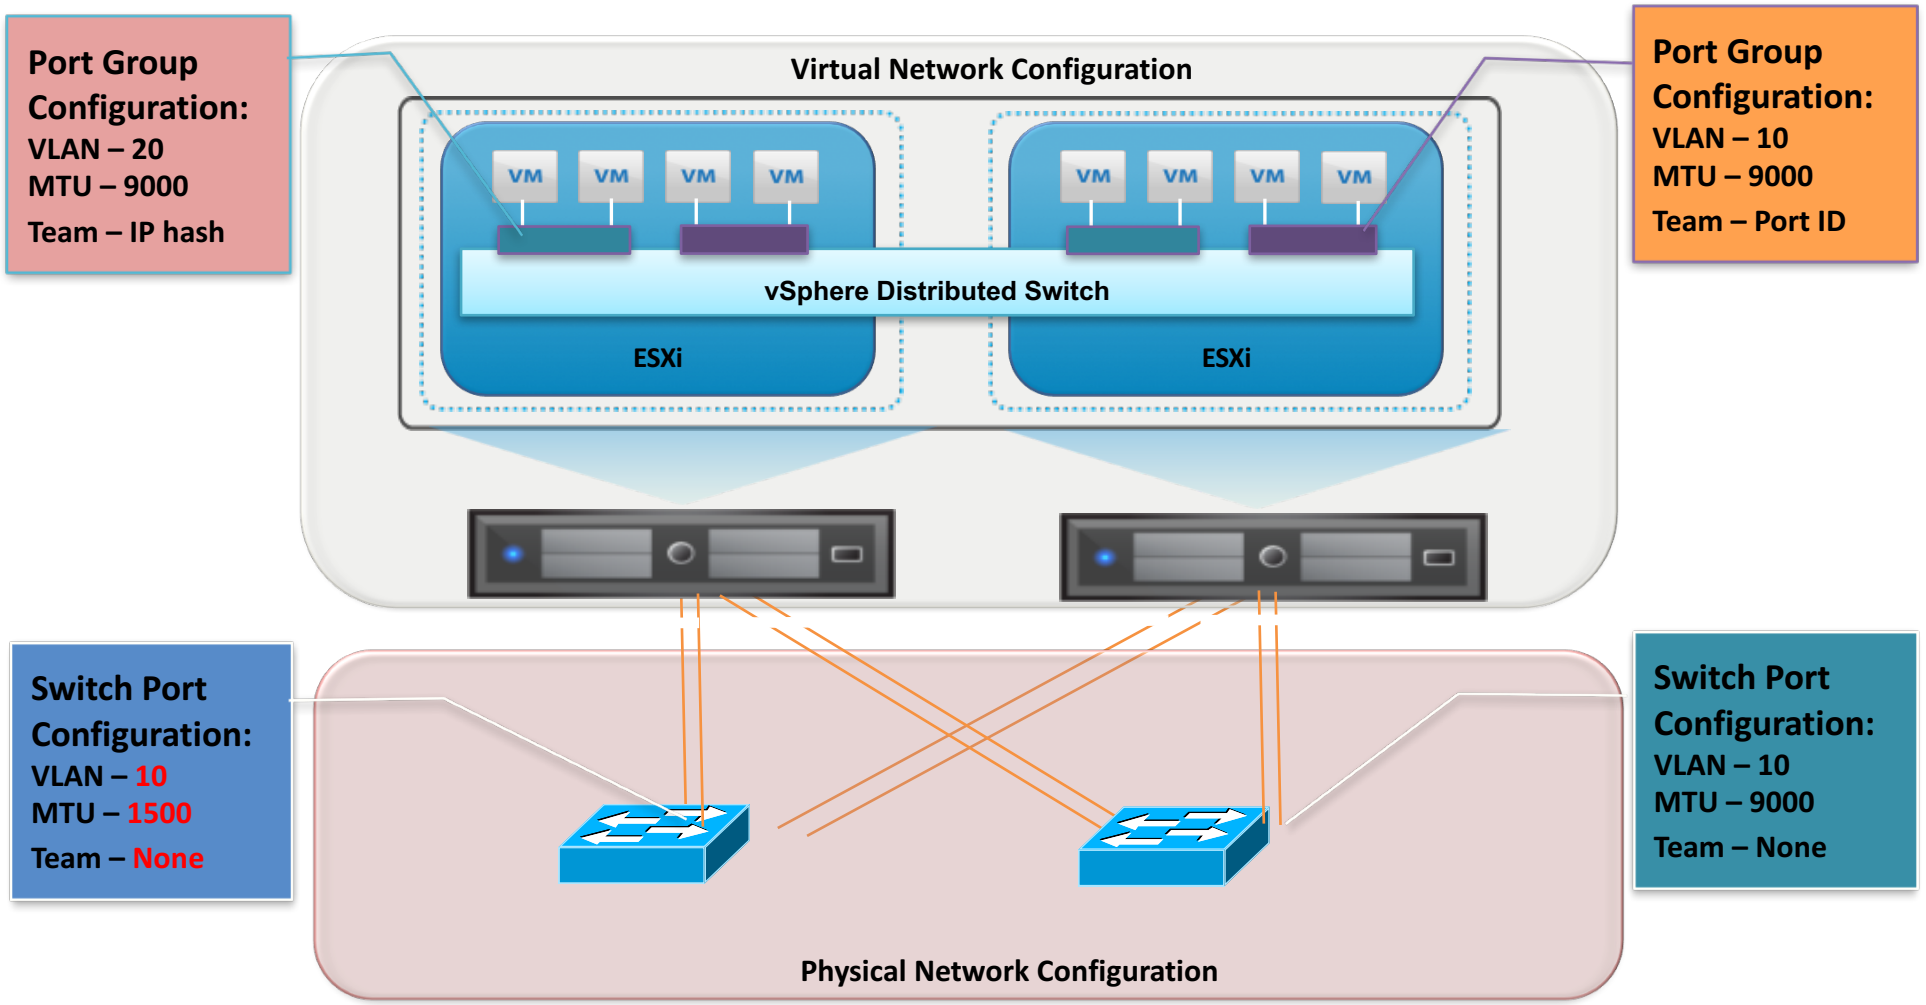

The network health check feature sends a probe packet every 2 mins

Misconfiguration of Management Network

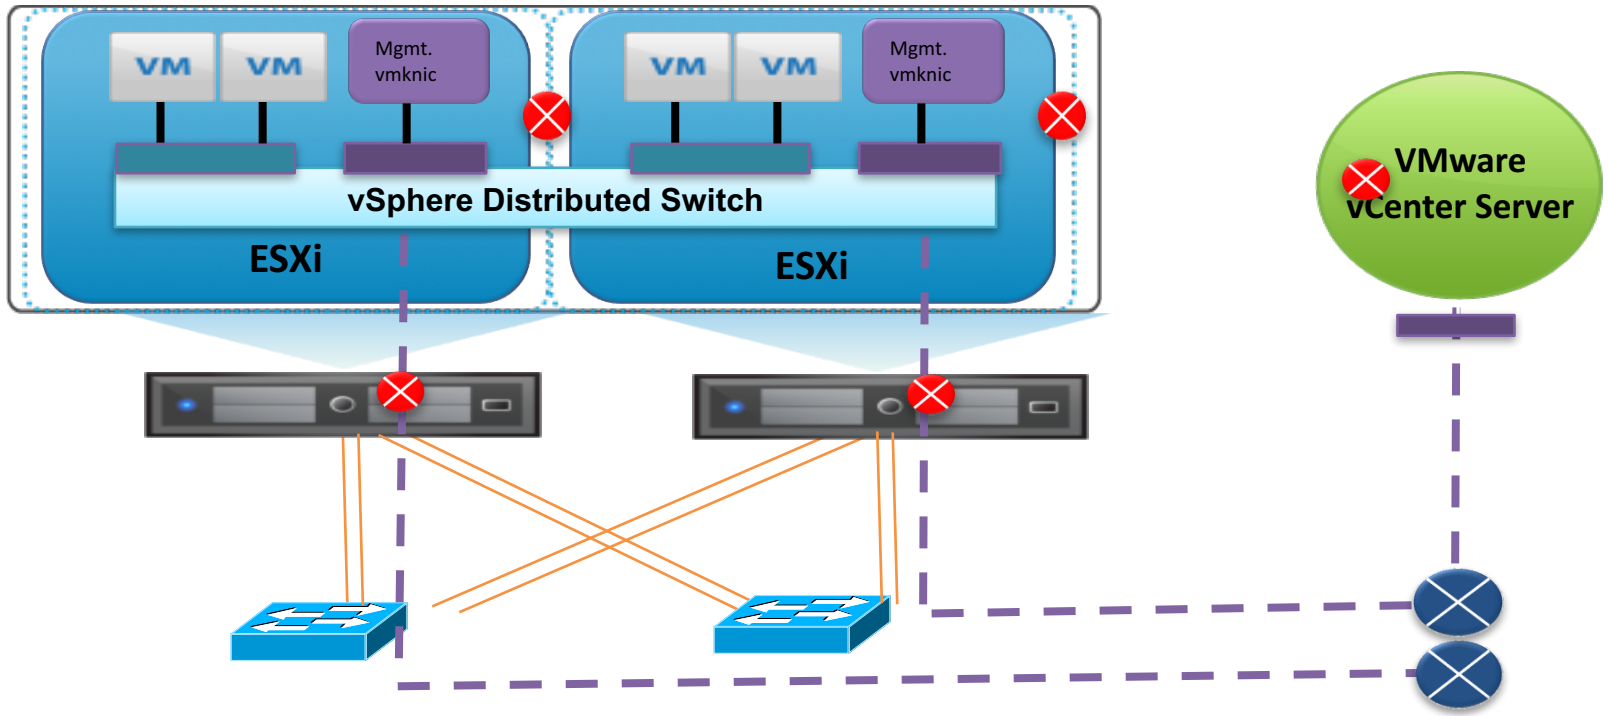

Two different updates that triggers rollback

- Host level Rollback gets triggered when there is change in the host networking configurations such as: Physical NIC speed change, Change in MTU configuration, Change in IP settings etc..
- VDS level rollback can happen after the user updates some VDS related objects such as port group or dvports.

Network Best Practices

- Allocate separate NICs for different traffic type
  - Can be connected to same uplink/physical NIC on 10GB network
- vSphere versions 5.0 and newer support multi-NIC, concurrent vMotion operations
- Use NIC load-based teaming (route based on physical NIC load)
  - For redundancy, load balancing, and improved vMotion speeds
- Have minimum 4 NICs per host to ensure performance and redundancy of network
- Recommend the use of NICs that support:
  - Checksum offload , TCP segmentation offload (TSO)
  - Jumbo frames (JF), Large receive offload (LRO)
  - Ability to handle high-memory DMA (i.e. 64-bit DMA addresses)
  - Ability to handle multiple Scatter Gather elements per Tx frame
  - NICs should support offload of encapsulated packets (with VXLAN)
- ALWAYS Check and Update Physical NIC Drivers
- Keep VMware Tools Up-to-Date - ALWAYS

Network Best Practices (continued)

- Use Virtual Distributed Switches for cross-ESX network convenience
- Optimize IP-based storage (iSCSI and NFS)
  - Enable Jumbo Frames
  - Use dedicated VLAN for ESXi host's vmknic & iSCSI/NFS server to minimize network interference from other packet sources
  - Exclude in-Guest iSCSI NICs from WSFC use
  - Be mindful of converged networks; storage load can affect network and vice versa as they use the same physical hardware; ensure no bottlenecks in the network between the source and destination
- Use VMXNET3 Para-virtualized adapter drivers to increase performance
  - NEVER use any other vNIC type, unless for legacy OSes and applications
  - Reduces overhead versus vlance or E1000 emulation
  - Must have VMware Tools to enable VMXNET3
- Tune Guest OS network buffers, maximum ports

Network Best Practices (continued)

- VMXNET3 can bite – but only if you let it
  - ALWAYS keep VMware Tools up-to-date
  - ALWAYS keep ESXi Host Firmware and Drivers up-to-date
  - Choose your physical NICs wisely
  - Windows Issues with VMXNET3
    - Older Windows versions
    - VMXNET3 template issues in Windows 2008 R2 - kb.vmware.com\kb\1020078
    - Hotfix for Windows 2008 R2 VMs - http://support.microsoft.com/kb/2344941
    - Hotfix for Windows 2008 R2 SP1 VMs - http://support.microsoft.com/kb/2550978

  - Disable interrupt coalescing – at vNIC level
    - **ONLY if ALL other options fail to remedy network-related performance Issue**

Your PC ran into a problem and needs to restart. We're just collecting some error info, and then we'll restart for you.

If you'd like to know more, you can search online later for this error: ALWAYS_LOOK_ON_THE_BRIGHT_SIDE_OF_LIFE

Networking – The changing landscape

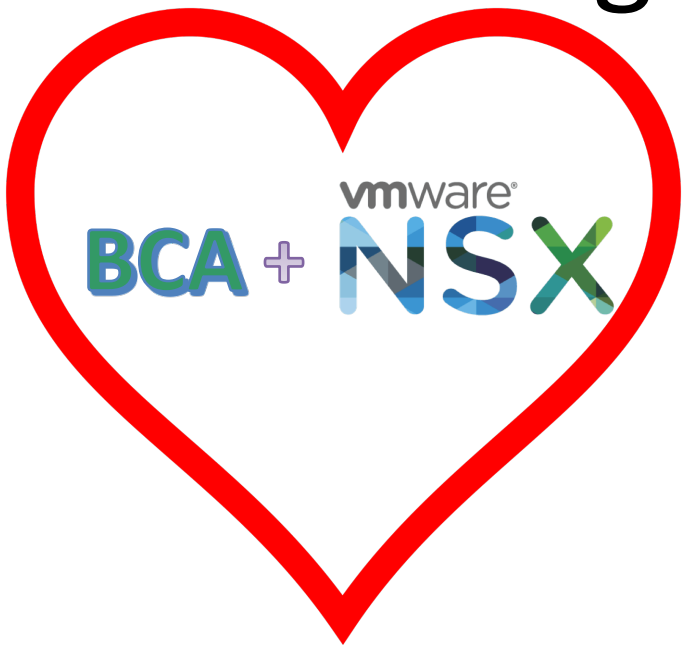

What is NSX?

- Network Overlay
- Logical networks
- Logical Routing
- Logical Firewall
- Logical Load Balancing
- Additional Networking services (NAT, VPN, more)
- Programmatically Controlled

NSX

IS Cool!!!

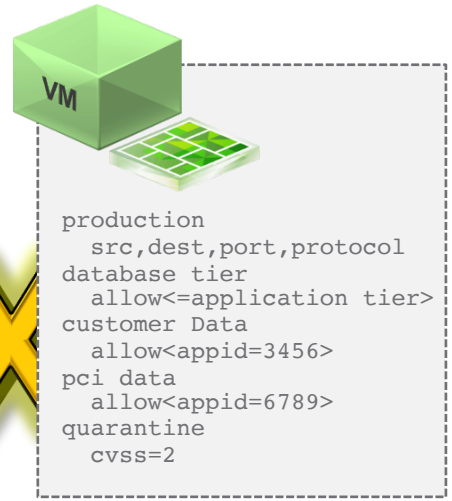

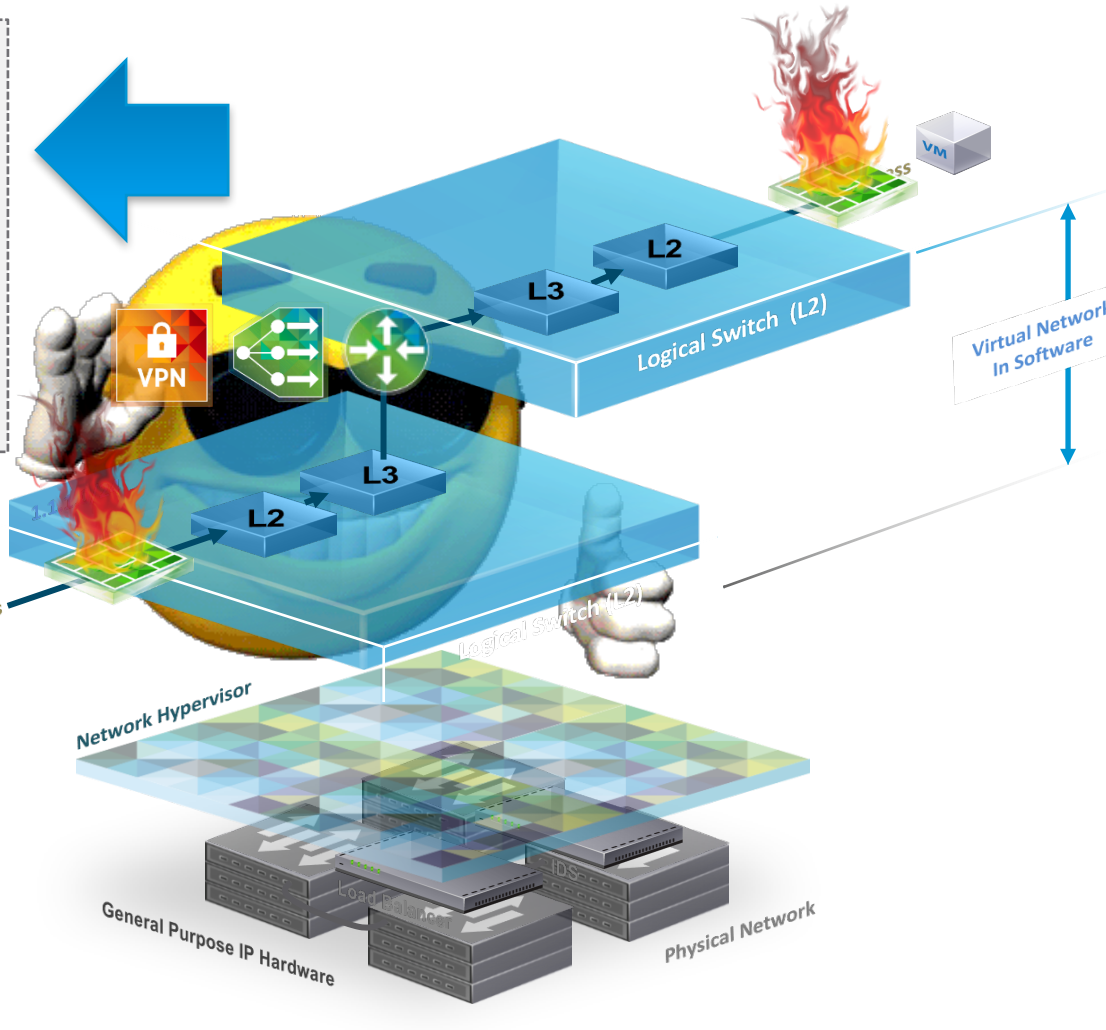

What is NSX?

- Network Overlay
- Logical networks
- Logical Routing
- Logical Firewall
- Logical Load Balancing
- Additional Networking services (NAT, VPN, more)
- Programmatically Controlled

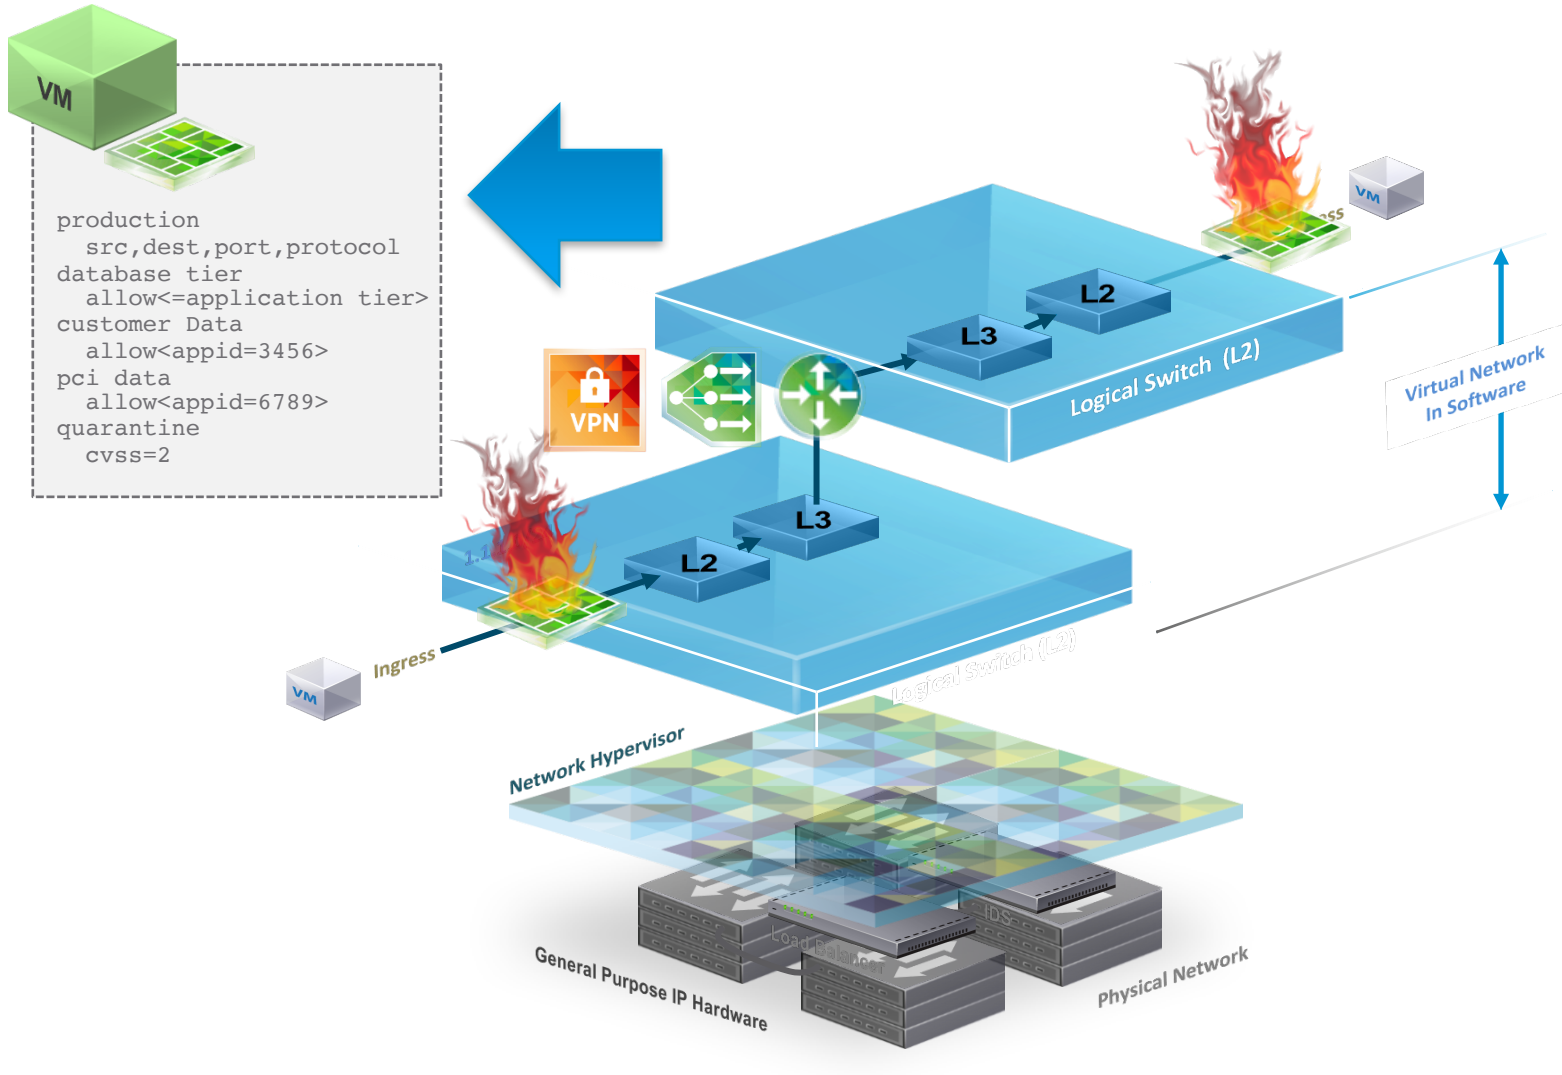

What do app owners care about?

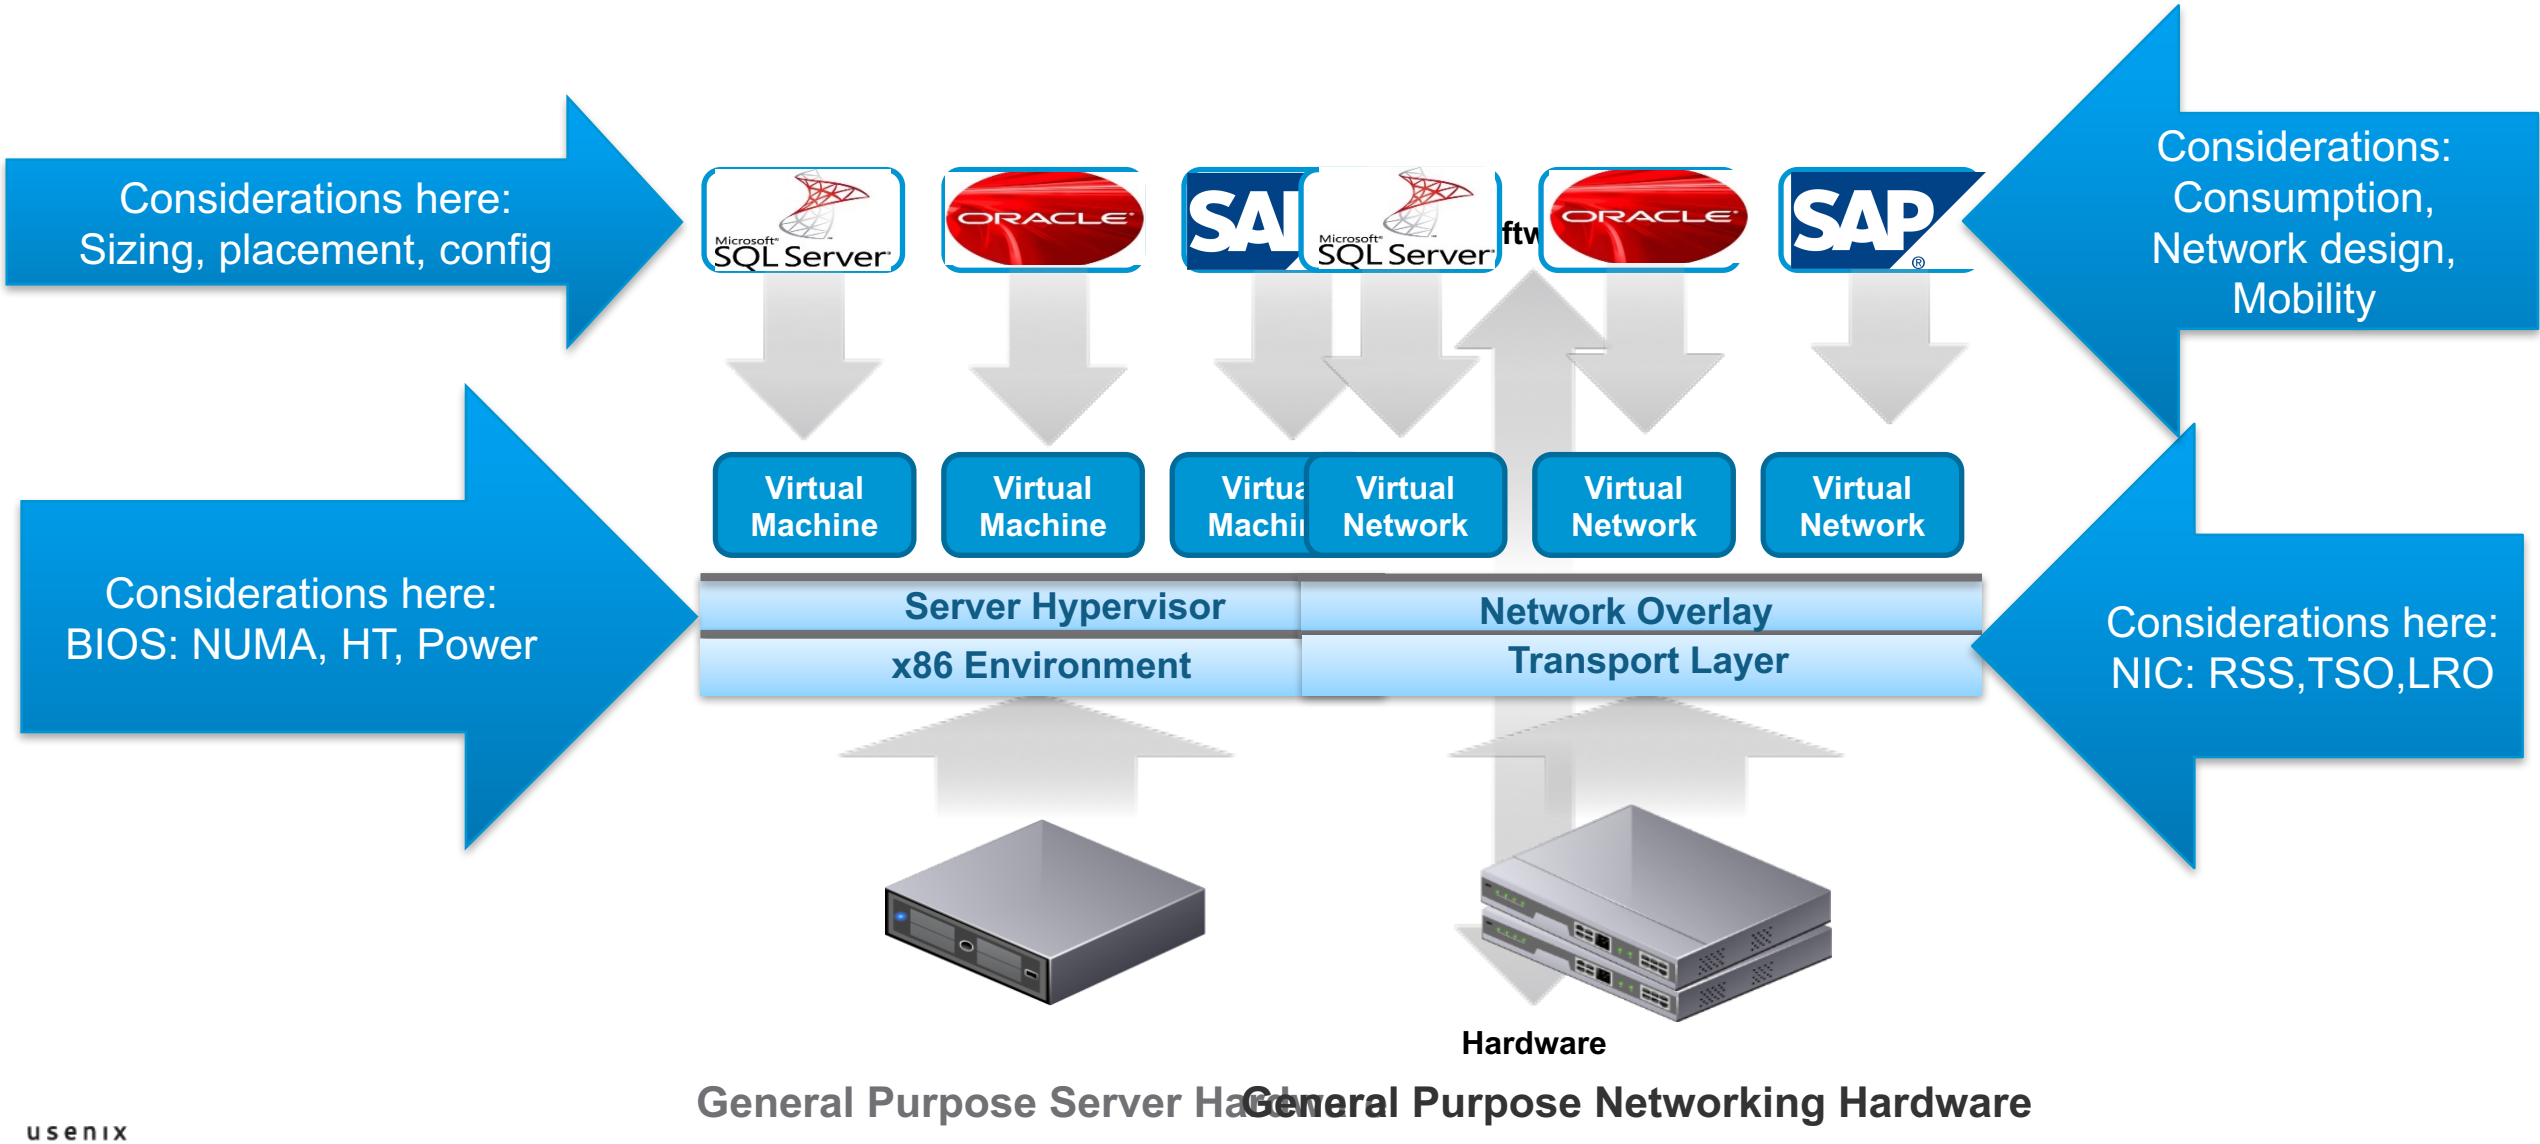

Performance Considerations

- All you need is IP connectivity between ESXi hosts
- The physical NIC and the NIC driver should support:
  - TSO - TCP Segmentation Offload = NIC divides larger data chunks into TCP segments
  - VXLAN offload – NIC encapsulates VXLAN instead of ESXi
  - RSS – Receive side scaling, allows the NIC to distribute received traffic to multiple CPU
  - LRO (Large Receive Offload) NIC reassembles incoming network packets

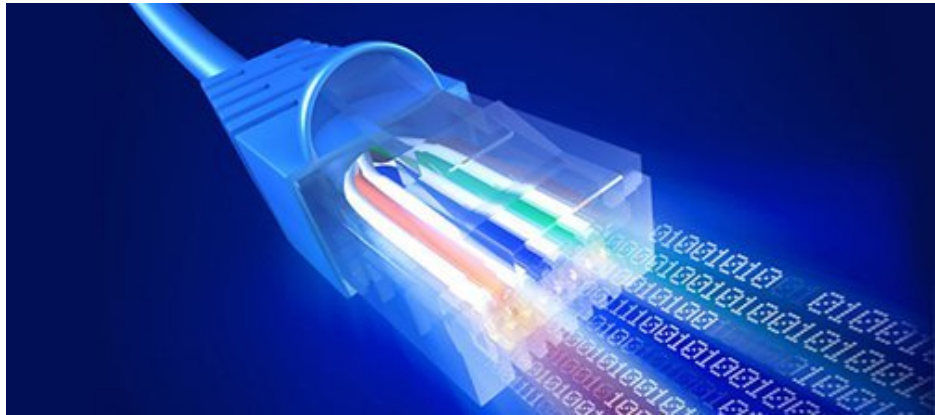

App owners say…

- So if the “Network hypervisor” fail does my app fail?
- What about NSX components dependencies?

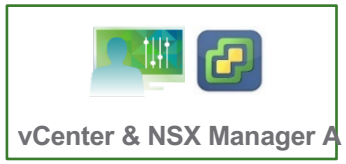

Management plane:
UI, API access
Not in the data path

Control plane:
Decouples virtual networks form physical topology
Not in Data Path
Highly Available

Data plane: Logical switches, Distributed Routers, Distributed Firewall, Edge devices

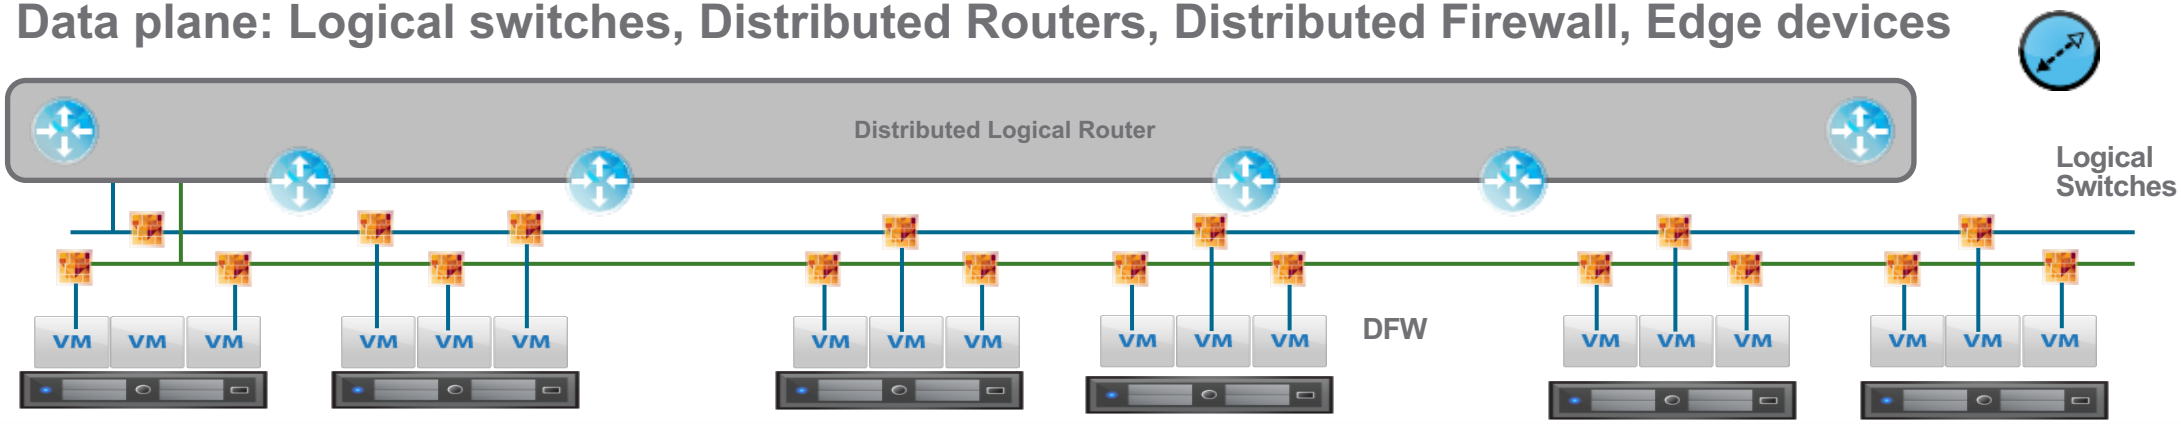

Connecting to the physical network

- Typical use case: 3-tier application, Web/App/DB, with non-virtualized DB tier.
- Option 1 – Route using an Edge device in HA mode:

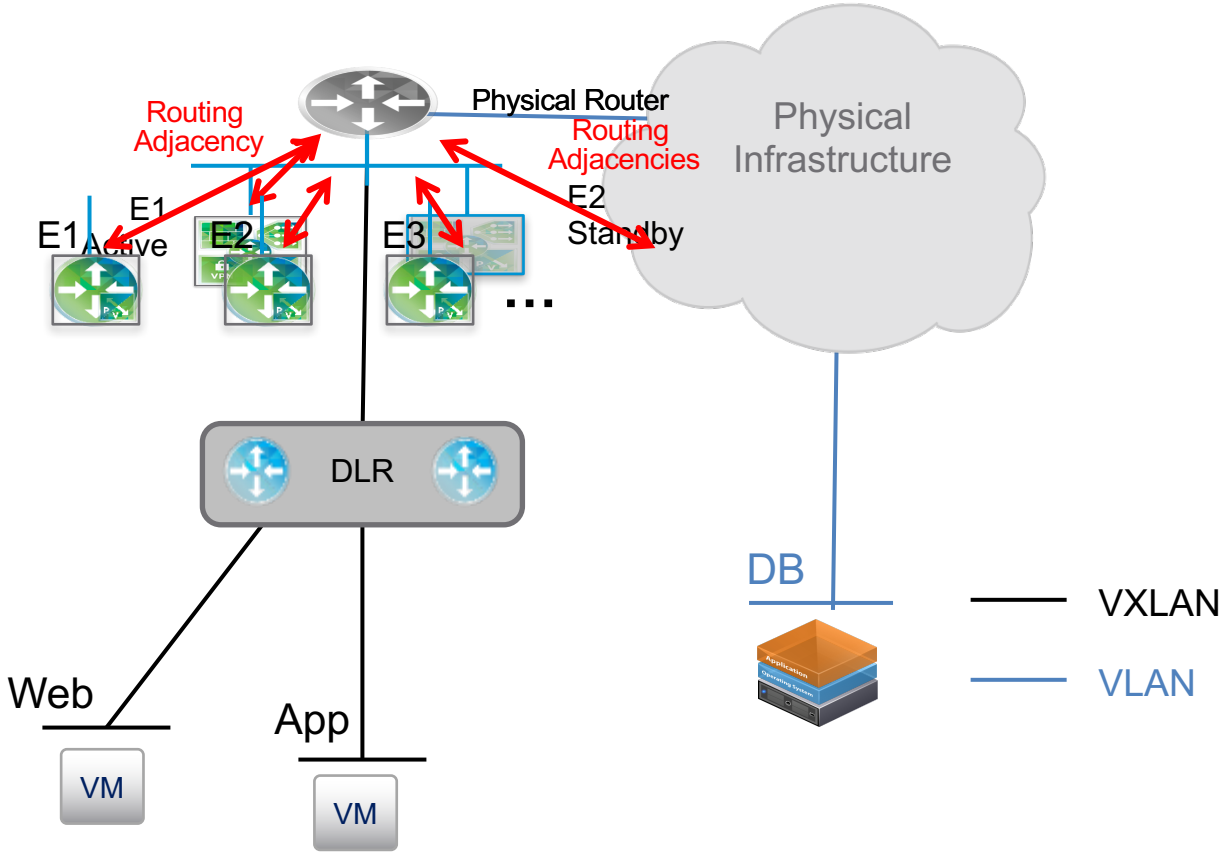

NSX Edge

Allows NO stateful services such as NAT, LB, VPN.

Limited throughput of 10Gbit (single ECMP node)

Failover takes a few seconds

Connecting to the physical network

- Typical use case: 3-tier application, Web/App/DB, with non-virtualized DB tier.
- Option 3 – Bridging L2 network using software or hardware GW:

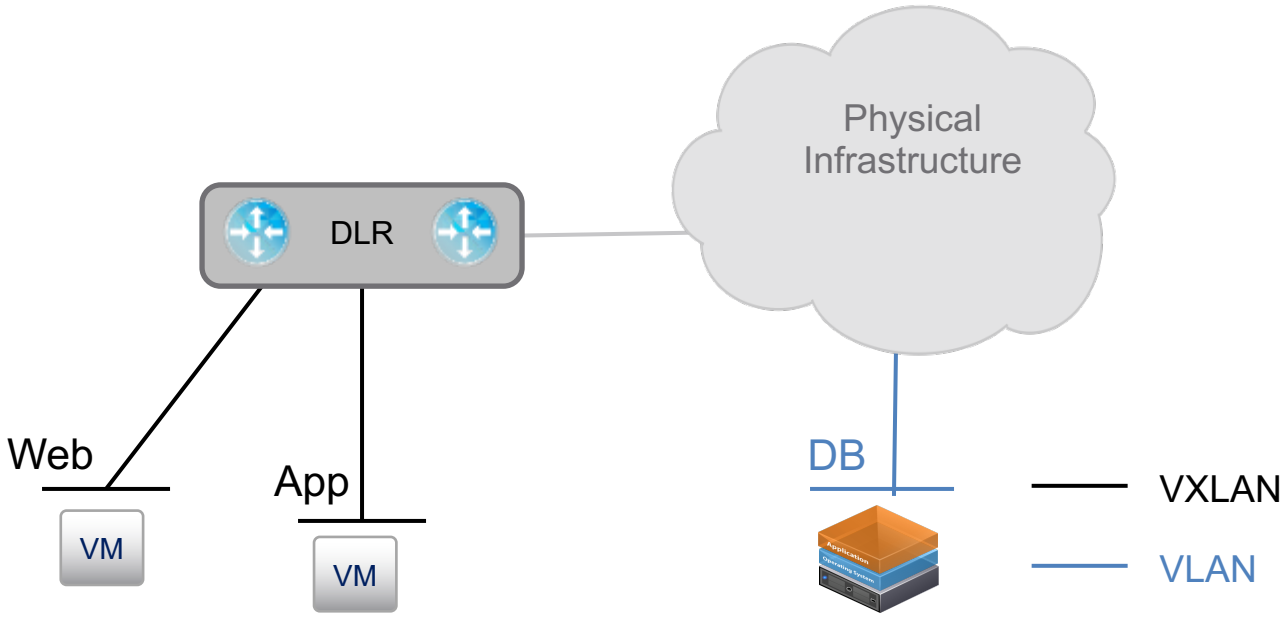

Straight from the ESXi kernel to the VLAN backed network

Lowest Latency
L2 adjacency between the tiers
Design complexity
Redundancy limitations

Designing for Availability

vSphere Native Availability Features

- **vSphere vMotion**
  - Can reduce virtual machine planned downtime
  - Relocates VMs without end-user interruption
    - Behavior COMPLETELY Configurable
  - Enables Admin to perform on-demand host maintenance without service interruption
- **vSphere DRS**
  - Monitors state of virtual machine resource usage
  - Can automatically and intelligently locate virtual machine
  - Can create a dynamically balanced Exchange Server deployment
  - Uses vMotion. Behavior COMPLETELY Configurable
- **vSphere High Availability (HA)**
  - HA Evaluates DRS Rules **BEFORE** Recovery – Just a checkbox operati
    - *** Now DEFAULT BEHAVIOR is vSphere 6.5**
  - Does not require Vendor-specific clustering solutions
  - **NOT** a replacement for app-specific native HA features
  - **COMPLEMENTS** and **ENHANCES** app-specific HA features
  - Automatically restarts failed virtual machine in minutes

vSphere Native Availability Feature Enhancements – vSphere 6.5

- **vCenter High Availability**
  - vCenter Server Appliance **ONLY**
  - Active, Passive, and Witness nodes – Exact Clones of existing vCenter Server.
  - Protects vCenter against Host, Appliance or Component Failures
  - 5-minute RTO at release

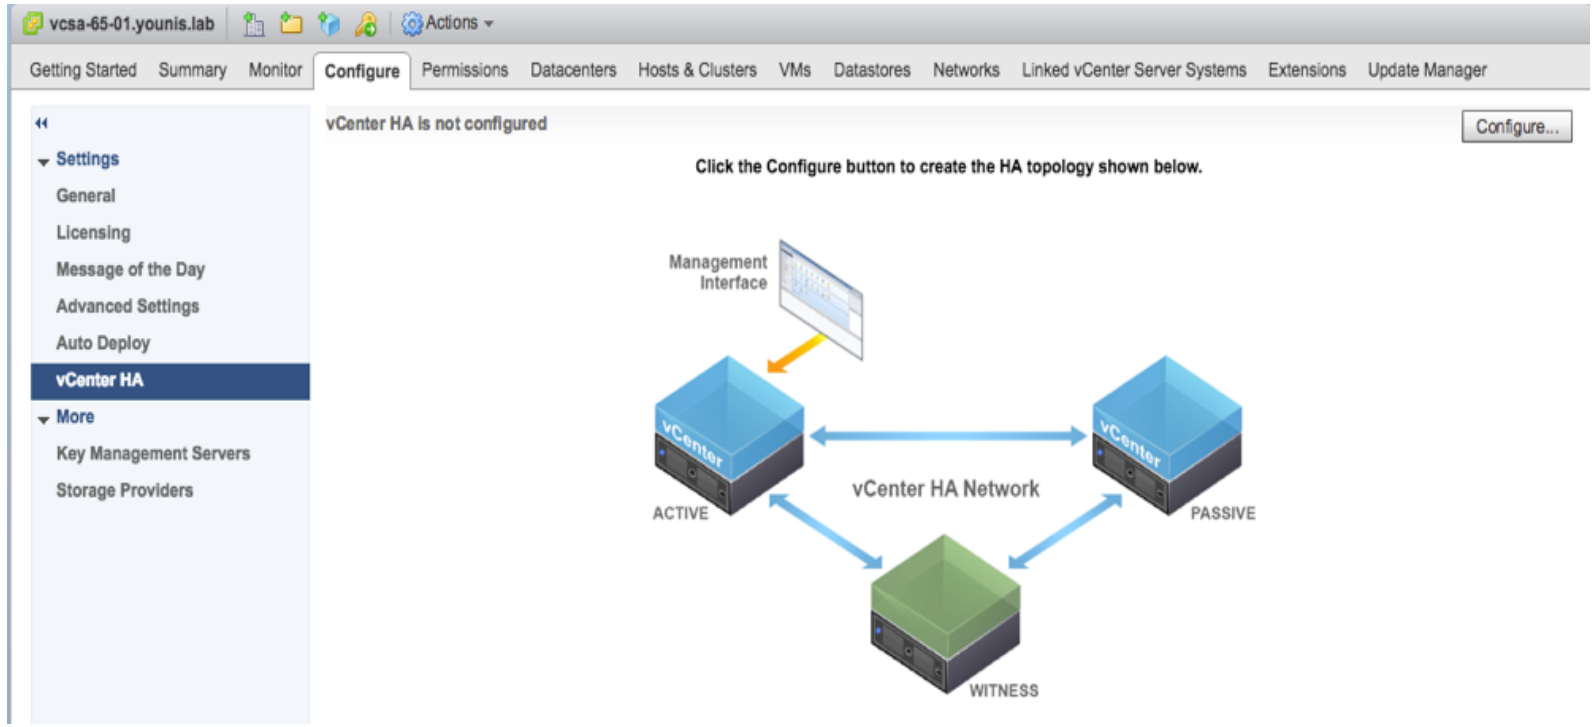

vSphere Native Availability Feature Enhancements – vSphere 6.5

- **Proactive High Availability**
  - Detects ESXi Host hardware failure or degradation
  - Leverage Hardware Vendor-provided plugin for monitoring
  - Reports Hardware state to vCenter
  - Unhealthy or failed hardware component is categorized bas
  - Puts impacted Hosts one of 2 states:
    - Quarantine Mode:
      - Existing VMs on Host not **IMMEDIATELY** evacuated.
      - Now new VM placed on Host
      - DRS attempts to remove Host if no performance impact to workloads in Cluster
    - Maintenance Mode:
      - Existing VMs on Host Evacuated
      - Host no longer participates in Cluster

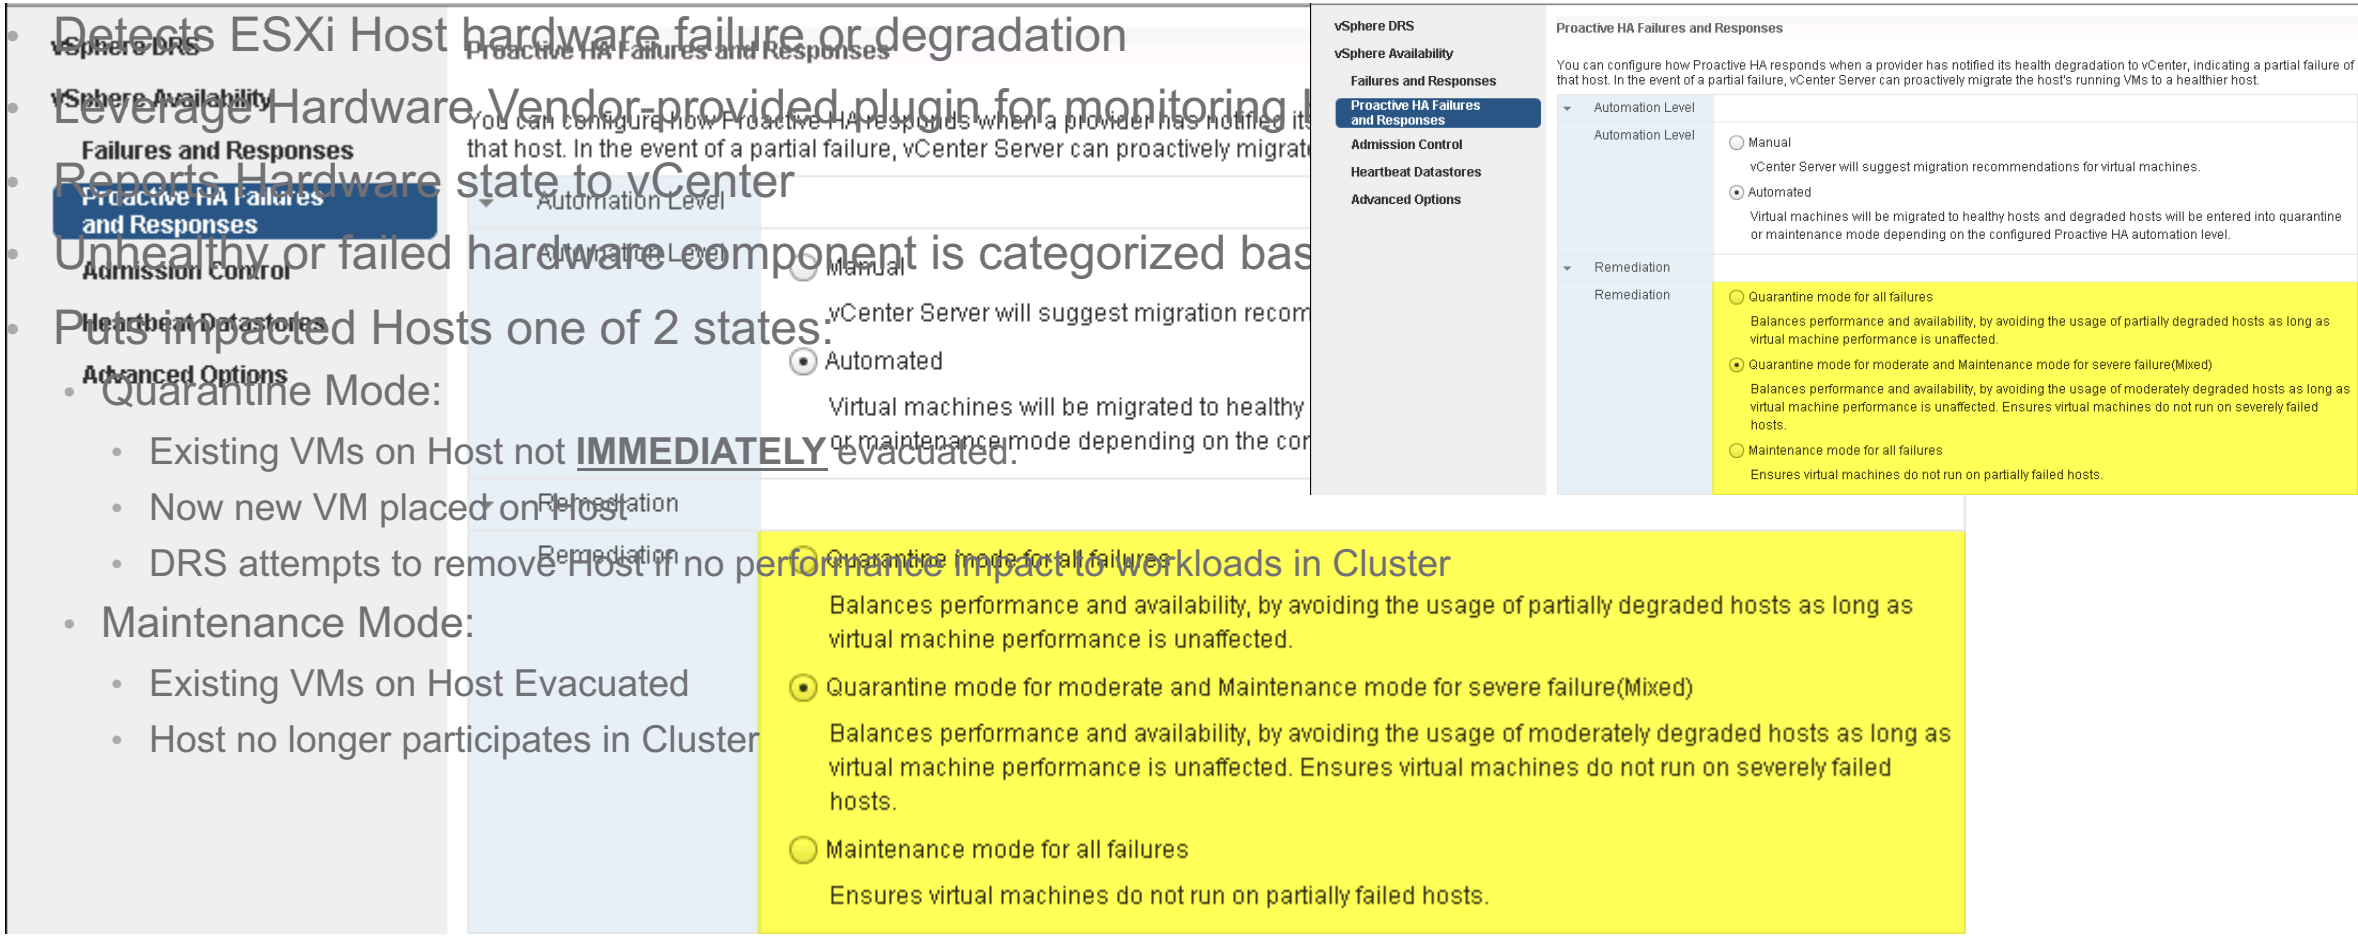

vSphere Native Availability Feature Enhancements – vSphere 6.5

- **Continuous VM Availability**
  - For when VMs MUST be up, even at the expense of **PERFORMANCE**

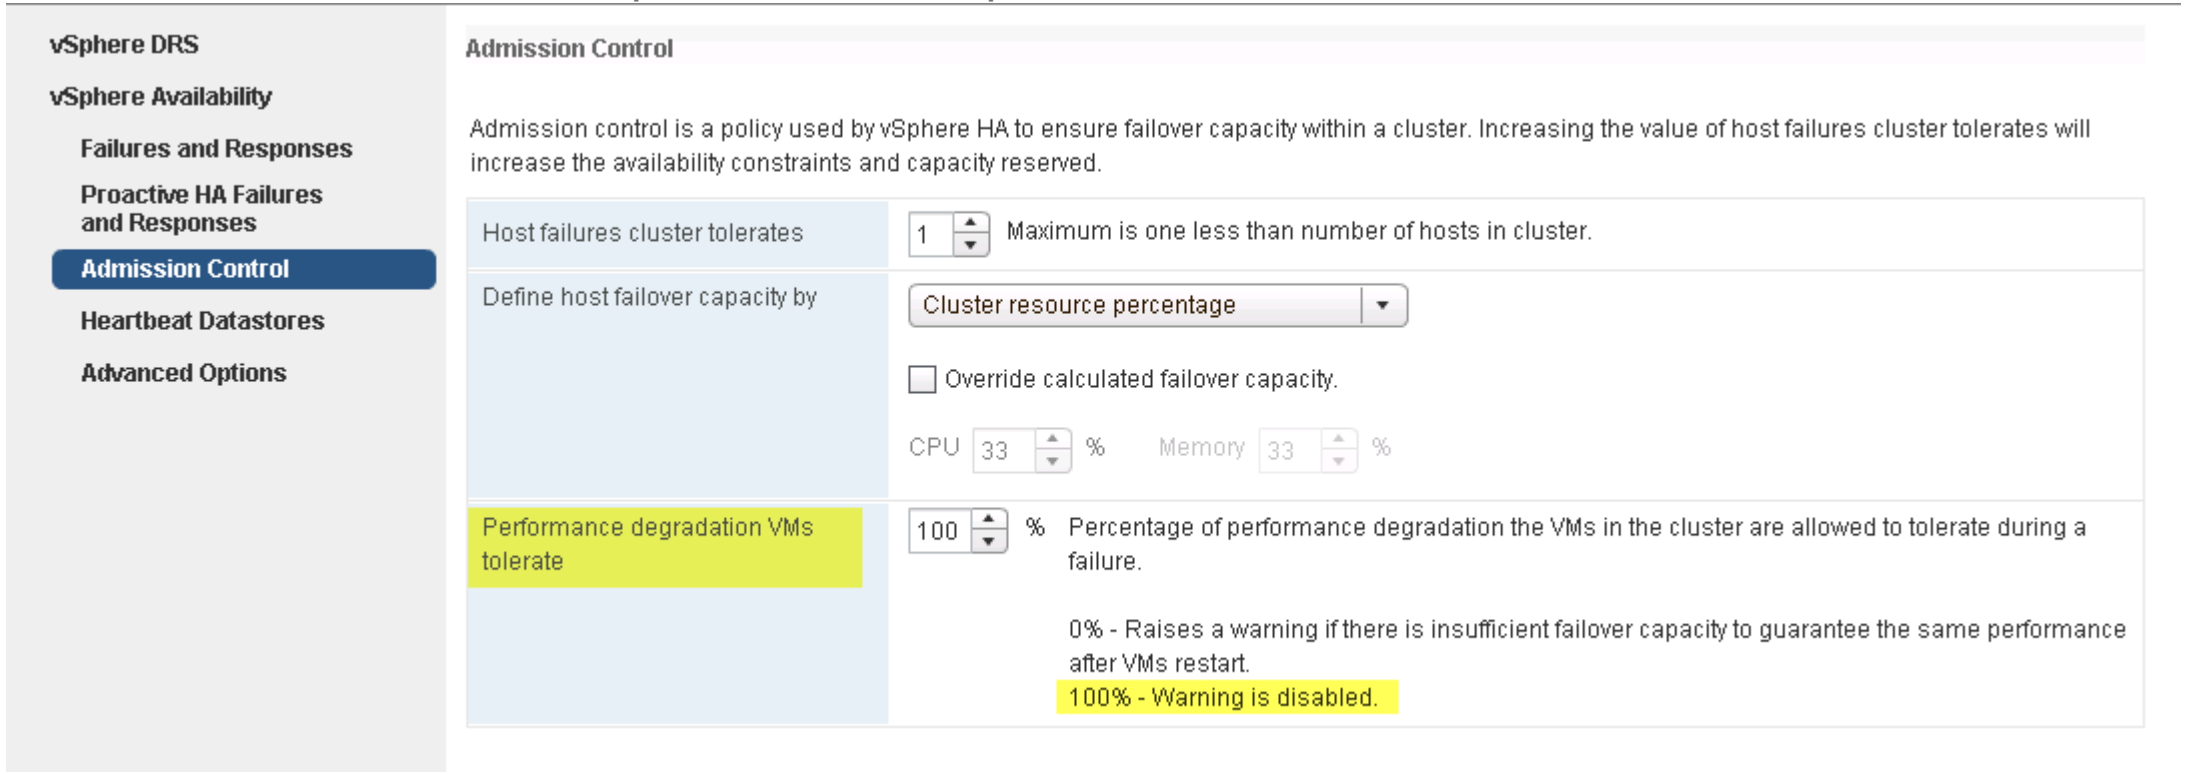

vSphere Native Availability Feature Enhancements – vSphere 6.5

- **vSphere DRS Rules**
  - Rules now includes “VM Dependencies”
  - Allows VMs to be recovered in order of PRIORITIES

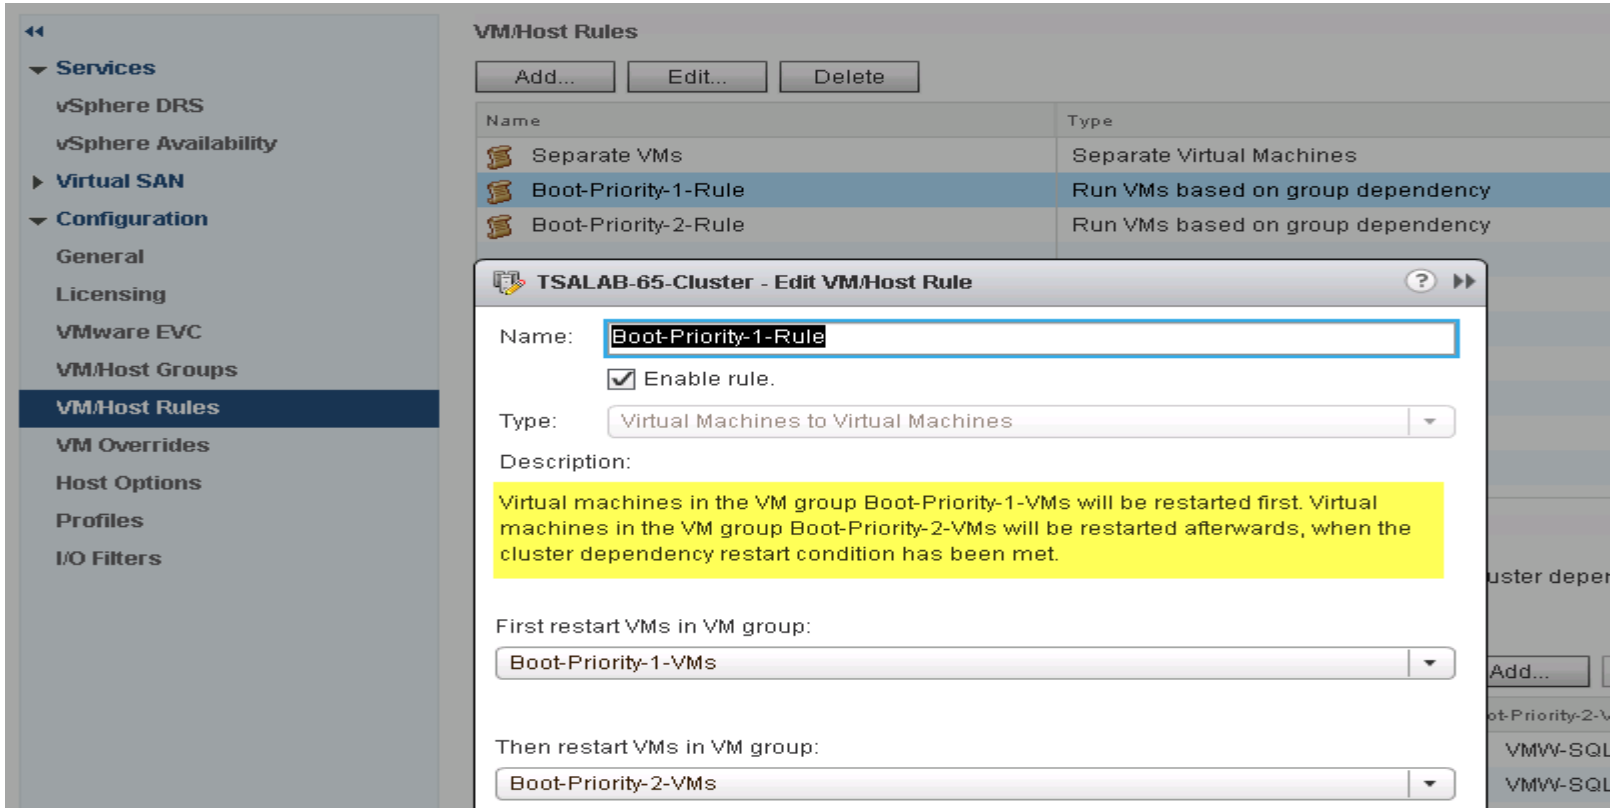

vSphere Native Availability Feature Enhancements – vSphere 6.5

- **Predictive DRS**
  - Integrated with VMware's vRealize Operations Monitoring Capabilities

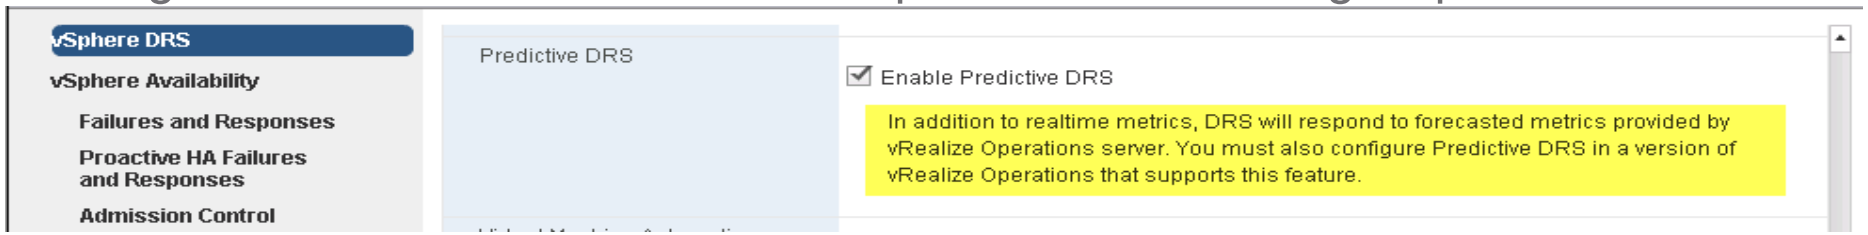

- **Network-Aware DRS**
  - Considers Host's Network Bandwith Utilization for VM Placement
  - Does NOT Evacuate VMs Based on Utilization

- **Simplified Advanced DRS Configuration Tasks**
  - Now just Checkbox options

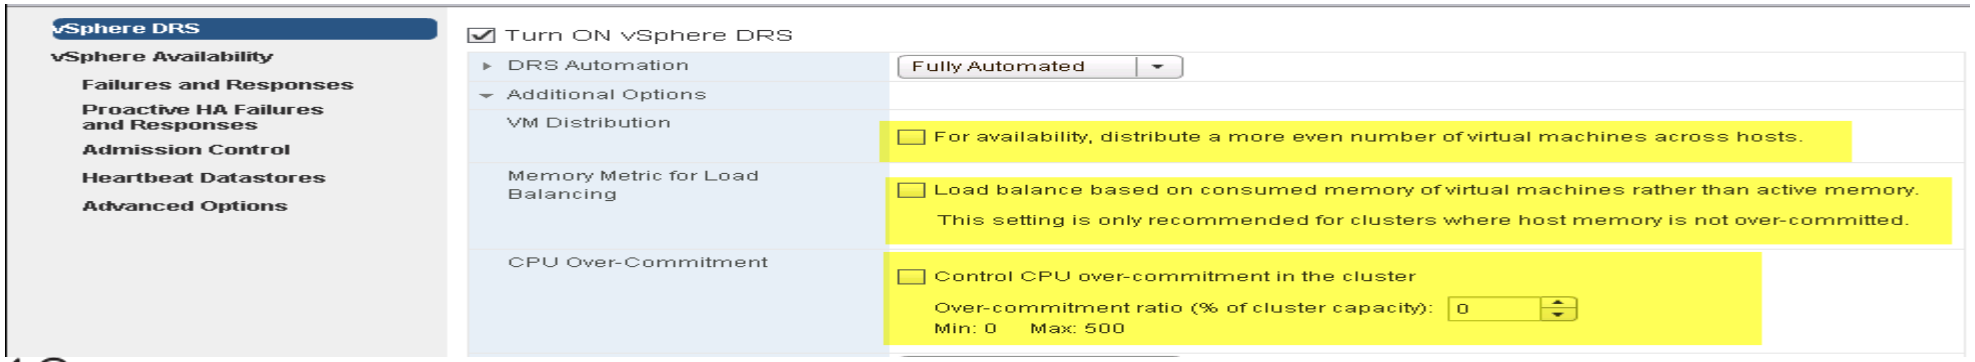

Combining Windows Applications HA with vSphere HA Features – The Caveats

Are You Going to Cluster THAT?

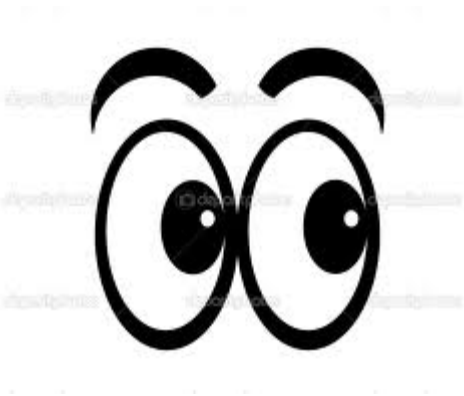

- **Do you NEED App-level Clustering?**
  - **Purely business and administrative decision**
  - **Virtualization does not preclude you from doing so**
- **Share-nothing Application Clustering?**
  - No “Special” requirements on vSphere
- **Shared-Disk Application Clustering (e.g. FCI / MSCS)**
  - You MUST use Raw Device Mapping (RDM) Disks Type for Shared Disks
  - MUST be connected to vSCSI controllers in PHYSICAL Mode Bus Sharing
  - Wonder why it’s called “Physical Mode RDM”, eh?
  - In Pre-vSphere 6.0, FCI/MSCS nodes CANNOT be vMotioned. Period
  - In vSphere 6.0 and above, you have vMotions capabilities under following conditions
    - **Clustered VMs are at Hardware Version > 10**
    - **vMotion VMKernel Portgroup Connected to 10GB Network**

vMotioning Clustered Windows Nodes – Avoid the Pitfall

- **Clustered Windows Applications Use Windows Server Failover Clustering (WSFC)**
  - WSFC has a **Default 5 Seconds Heartbeat Timeout Threshold**
  - vMotion Operations **MAY Exceed 5 Seconds** (During VM Quiescing)
    - Leading to Unintended and Disruptive Clustered Resource Failover Events
- **SOLUTION**
  - Use **MULTIPLE** vMotion Portgroups, where possible
  - **Enable jumbo frames on all vmkernel ports,** IF **PHYSICAL Network Supports it**
  - If jumbo frames is not supported, consider modifying default WSFC behaviors:
    - *(get-cluster).SameSubnetThreshold = 10*
    - *(get-cluster).CrossSubnetThreshold = 20*
    - *(get-cluster).RouteHistoryLength = 40*
    - **NOTES:**
      - You may need to ***"Import-Module FailoverClusters"*** first
- **Behavior NOT Unique to VMware or Virtualization**
  - If Your Backup Software Quiesces Exchange, You Experience Symptom
  - See Microsoft's "Tuning Failover Cluster Network Thresholds" – http://bit.ly/1nJRPs3

Monitoring and Identifying Performance Bottlenecks

Performance Needs Monitoring at Every Level

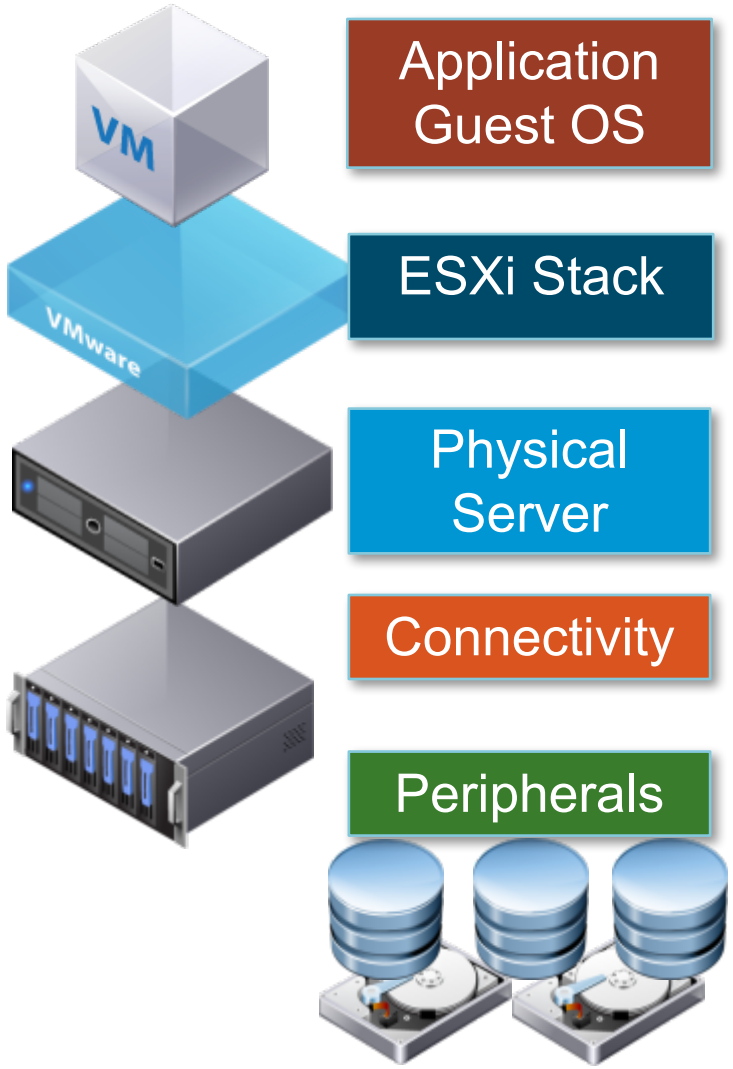

Application Level
App Specific Perf tools/stats

Guest OS
CPU Utilization, Memory Utilization, I/O Latency

Virtualization Level
vCenter Performance Metrics /Charts
Limits, Shares, Virtualization Contention

Physical Server Level
CPU and Memory Saturation, Power Saving

Connectivity Level
Network/FC Switches and data paths
Packet loss, Bandwidth Utilization

Peripherals Level
SAN or NAS Devices
Utilization, Latency, Throughput

START HERE

Host Level Monitoring

- VMware vSphere Client™
  - GUI interface, primary tool for observing performance and configuration data for one or more vSphere hosts
  - Does not require high levels of privilege to access the data

- Resxtop/ESXTop
  - Gives access to detailed performance data of a single vSphere host
  - Provides fast access to a large number of performance metrics
  - Runs in interactive, batch, or replay mode
  - ESXTop Cheat Sheet - http://www.running-system.com/vsphere-6-esxtop-quick-overview-for-troubleshooting/

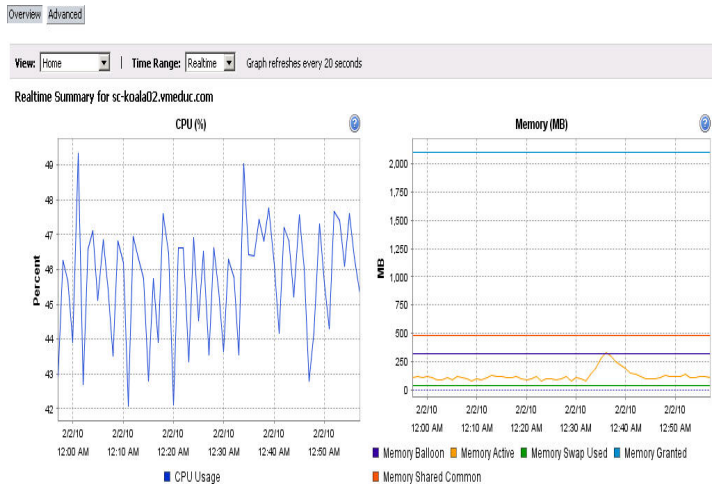

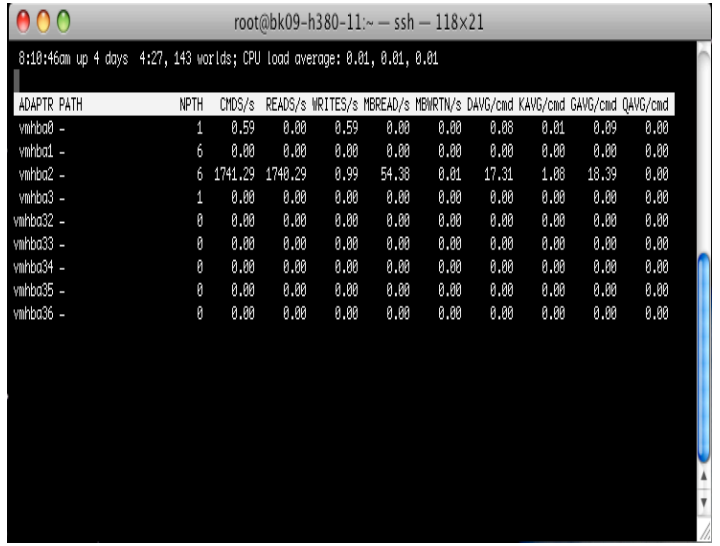

Key Metrics to Monitor for vSphere

Resource	Metric	Host / VM	Description
CPU	%USED	Both	CPU used over the collection interval (%)
	%RDY	VM	CPU time spent in ready state
	%SYS	Both	Percentage of time spent in the ESX Server VMKernel
Memory	Swapin, Swapout	Both	Memory ESX host swaps in/out from/to disk (per VM, or cumulative over host)
	MCTLSZ (MB)	Both	Amount of memory reclaimed from resource pool by way of ballooning
Disk	READs/s, WRITEs/s	Both	Reads and Writes issued in the collection interval
	DAVG/cmd	**Both**	**Average latency (ms) of the device (LUN)**
	KAVG/cmd	**Both**	**Average latency (ms) in the VMkernel, also known as "queuing time"**
	GAVG/cmd	**Both**	**Average latency (ms) in the guest. GAVG = DAVG + KAVG**
Network	MbRX/s, MbTX/s	Both	Amount of data transmitted per second
	PKTRX/s, PKTTX/s	Both	Packets transmitted per second
	%DRPRX, %DRPTX	Both	Drop packets per second

Key Indicators

CPU

- Ready (%RDY)
  - % time a vCPU was ready to be scheduled on a physical processor but couldn't’t due to processor contention
  - Investigation Threshold: 10% per vCPU
- Co-Stop (%CSTP)
  - % time a vCPU in an SMP virtual machine is “stopped” from executing, so that another vCPU in the same virtual machine could be run to “catch-up” and make sure the skew between the two virtual processors doesn't grow too large
  - Investigation Threshold: 3%
- Max Limited (%MLMTD)
  - % time VM was ready to run but wasn’t scheduled because it violated the CPU Limit set ; added to %RDY time
  - Virtual machine level – processor queue length

Key Performance Indicators

Memory

Balloon driver size (MCTLSZ)

the total amount of guest physical memory reclaimed by the balloon driver

Investigation Threshold: 1

Swapping (SWCUR)

the current amount of guest physical memory that is swapped out to the ESX kernel VM swap file.

Investigation Threshold: 1

Swap Reads/sec (SWR/s)

the rate at which machine memory is swapped in from disk.

Investigation Threshold: 1

Swap Writes/sec (SWW/s)

The rate at which machine memory is swapped out to disk.

Investigation Threshold: 1

Network

Transmit Dropped Packets (%DRPTX)

The percentage of transmit packets dropped.

Investigation Threshold: 1

Receive Dropped Packets (%DRPRX)

The percentage of receive packets dropped.

Investigation Threshold: 1

Logical Storage Layers: from Physical Disks to vmdks

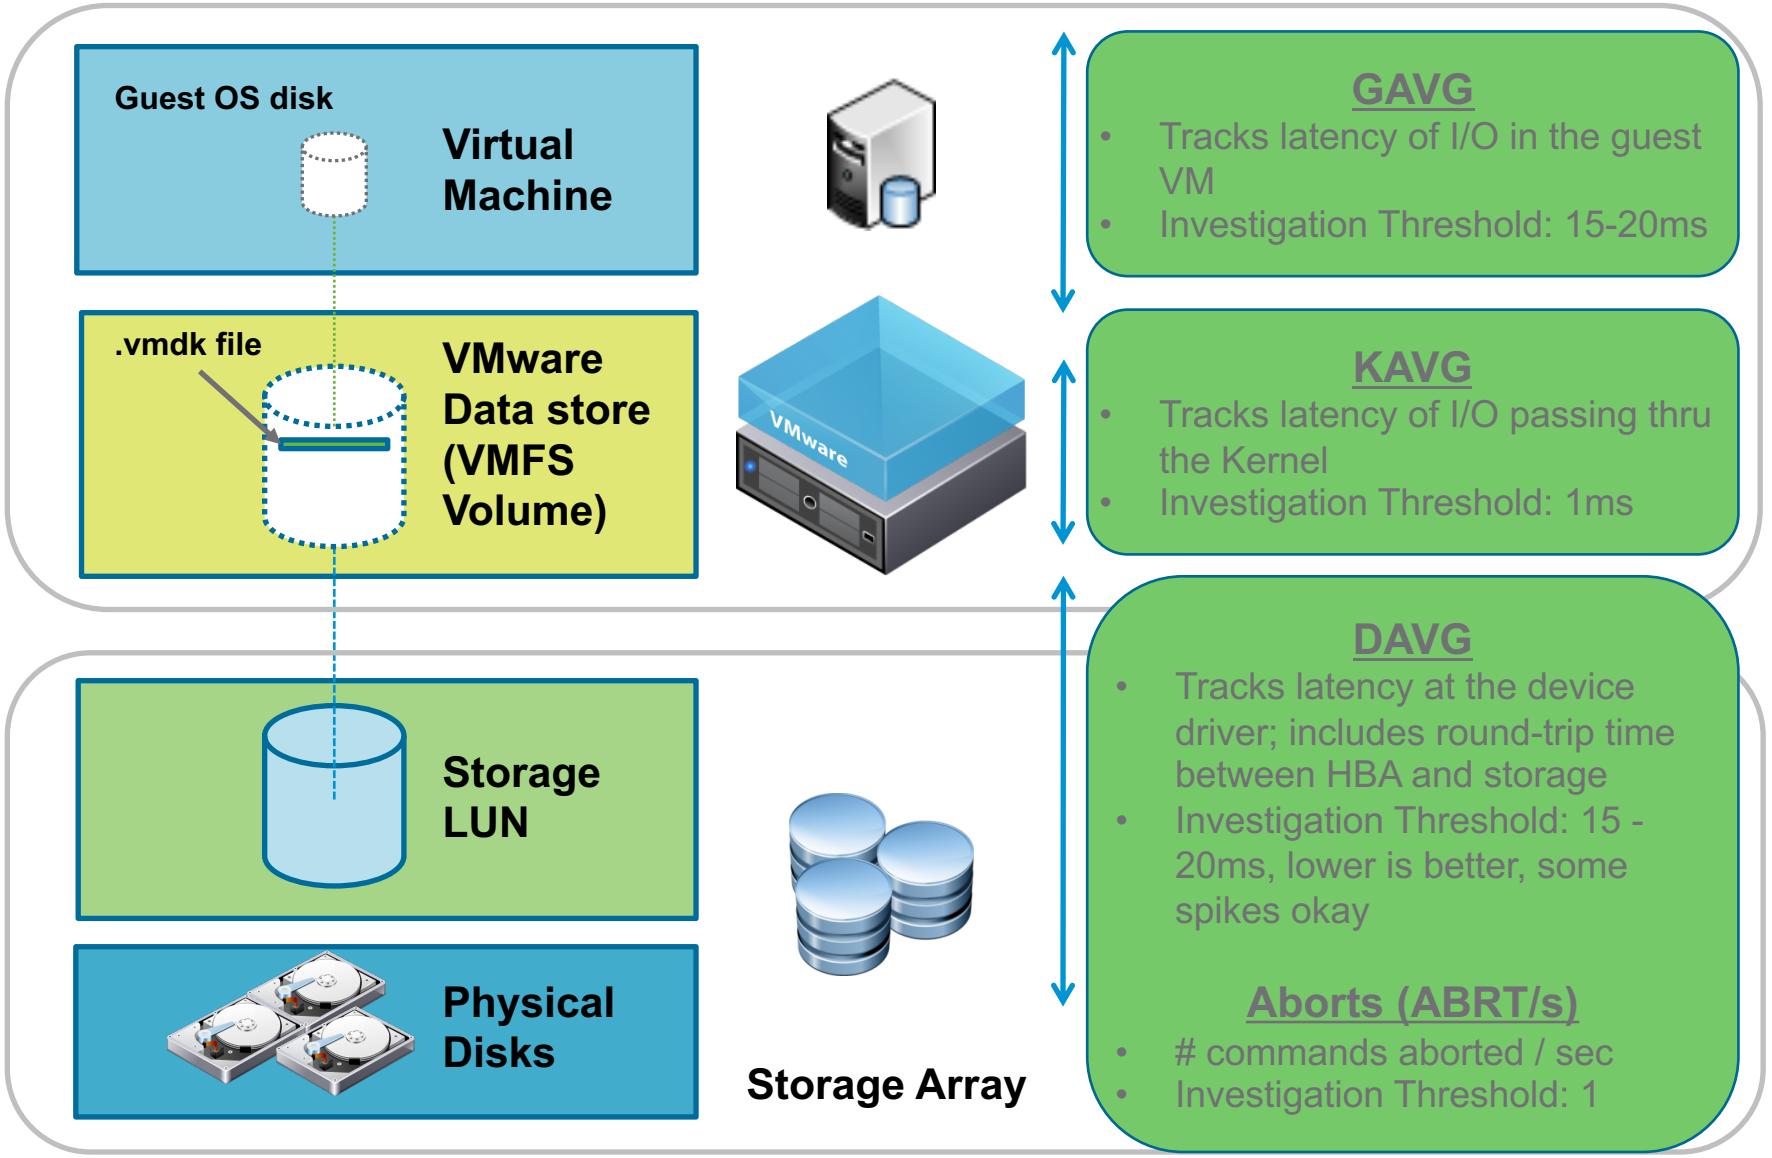

Key Indicators

Storage

- Kernel Latency Average (KAVG)
  - This counter tracks the latencies of IO passing thru the Kernel
  - Investigation Threshold: 1ms
- Device Latency Average (DAVG)
  - This is the latency seen at the device driver level. It includes the roundtrip time between the HBA and the storage.
  - Investigation Threshold: 15-20ms, lower is better, some spikes okay
- Aborts (ABRT/s)
  - The number of commands aborted per second.
  - Investigation Threshold: 1
- Size Storage Arrays appropriately for Total VM usage
  - > 15-20ms Disk Latency could be a performance problem
  - > 1ms Kernel Latency could be a performance problem or a undersized ESX device queue

Storage Performance Troubleshooting Tools

Storage Profiling Tips and Tricks

- Common IO Profiles (database, web, etc): http://blogs.msdn.com/b/tvoellm/archive/2009/05/07/useful-io-profiles-for-simulating-various-workloads.aspx

- Make Sure to Check / Try:
  - Load balancing / multi-pathing
  - Queue depth & outstanding I/Os
  - pvSCSI Device Driver

- Look out for:
  - I/O contention
  - Disk Shares
  - SIOC & SDRS
  - IOP Limits

vscsiStats – DEEP Storage Diagnostics

- vscsiStats characterizes IO for each virtual disk
  - Allows us to separate out each different type of workload into its own container and observe trends
- Histograms only collected if enabled; no overhead otherwise
- Metrics
  - I/O Size
  - Seek Distance
  - Outstanding I/Os
  - I/O Interarrival Times
  - Latency

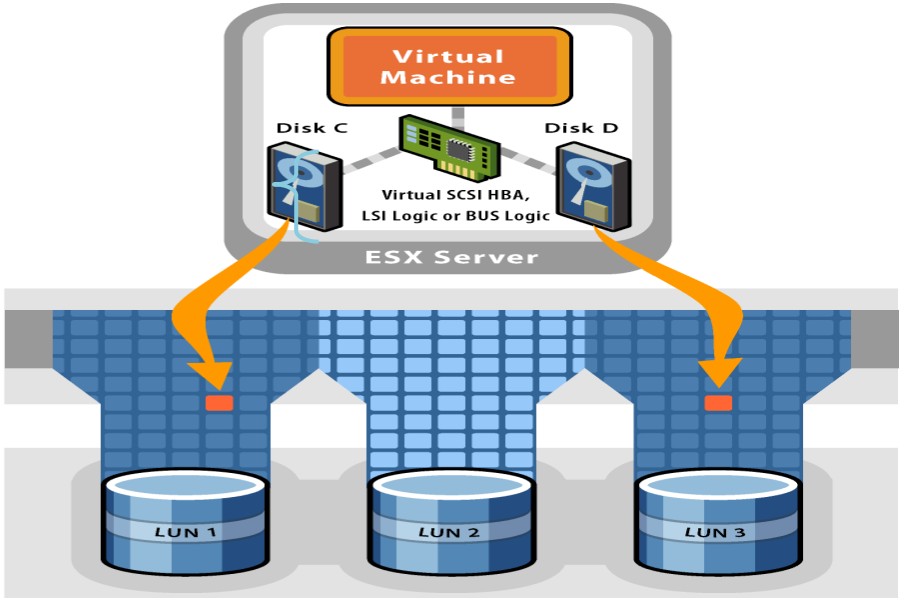


```
root@bk09-h380-11:/usr/lib/vmware/bin
[root@bk09-h380-11 bin]# ./vscsiStats -w 4968 -i 8208          \
                         -p seekDistance

Histogram: distance (in LBNs) between successive commands for vi
rtual machine worldGroupID : 4968, virtual disk handleID : 8208
{
 min : -16771839
 max : 807168
 mean : -36
 count : 429263
   {
      1                    (<                  -128)
      0                    (<=                 -64)
      0                    (<=                 -32)
      0                    (<=                 -1)
      0                    (<=                 0)
      429261               (<=                 1)
      0                    (<=                 32)
      0                    (<=                 64)
      1                    (>                  128)
   }
```

Monitoring Disk Performance with esxtop

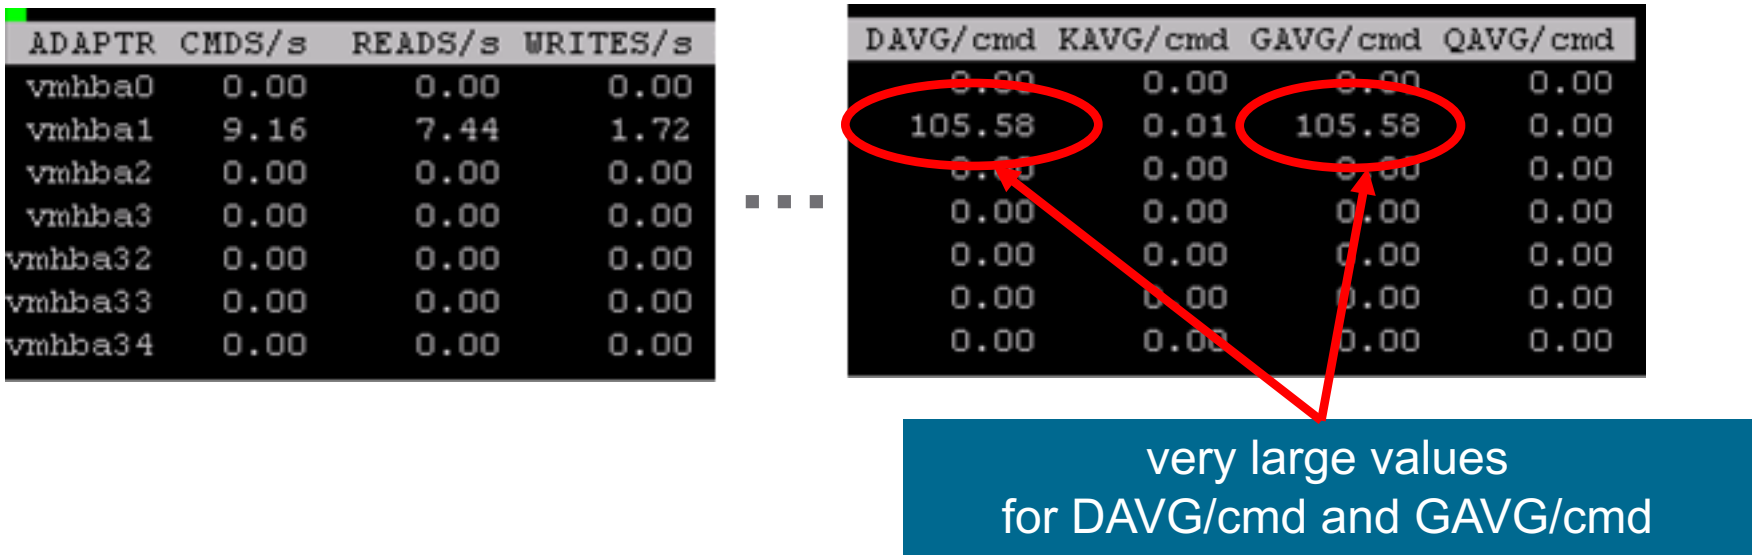

ADAPTR	CMDS/s	READS/s	WRITES/s
vmhba0	0.00	0.00	0.00
vmhba1	9.16	7.44	1.72
vmhba2	0.00	0.00	0.00
vmhba3	0.00	0.00	0.00
vmhba32	0.00	0.00	0.00
vmhba33	0.00	0.00	0.00
vmhba34	0.00	0.00	0.00

DAVG/cmd	KAVG/cmd	GAVG/cmd	QAVG/cmd
0.00	0.00	0.00	0.00
105.58	0.01	105.58	0.00
0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00

- Rule of thumb
  - GAVG/cmd > 20ms = high latency!
- What does this mean?
  - When command reaches device, latency is high
  - Latency as seen by the guest is high
  - Low KAVG/cmd means command is not queuing in VMkernel

Iometer

An I/O subsystem measurement and characterization tool for single and clustered systems.

Supports Windows and Linux

- Windows and Linux
- Free (Open Source)
- Single or Multi-server capable
- Multi-threaded

- Metrics Collected
  - Total I/Os per Sec.
  - Throughput (MB)
  - CPU Utilization
  - Latency (avg. & max)

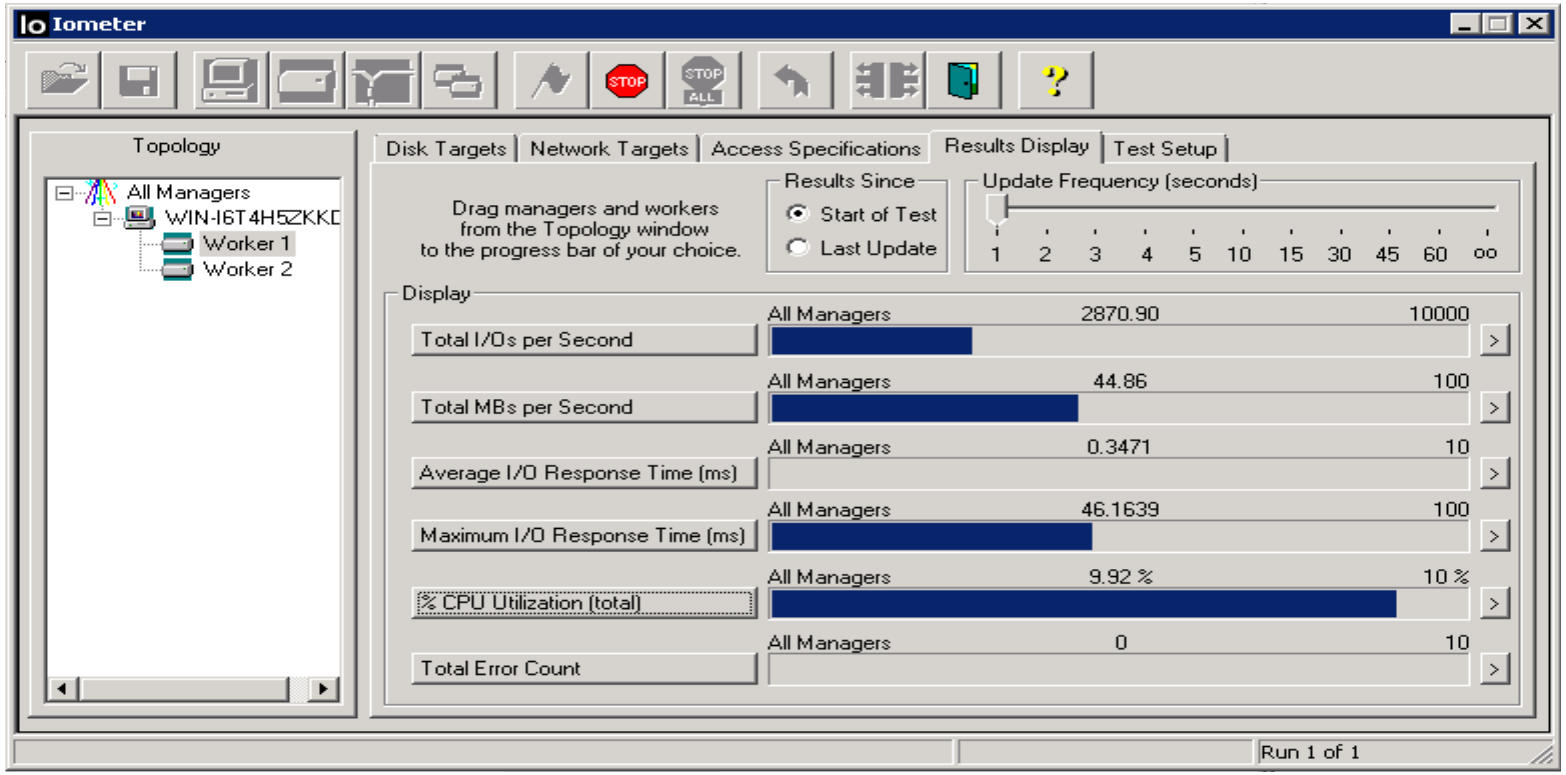

DiskSpd Utility: A Robust Storage Testing Tool (SQLIO)

https://gallery.technet.microsoft.com/DiskSpd-a-robust-storage-6cd2f223
http://hfxte.ch/diskspd

- Windows-based feature-rich synthetic storage testing and validation tool
- Replaces SQLIO and effective for baselining storage for MS SQL Server workloads
- Fine-grained IO workload characteristics definition
- Configurable runtime and output options
- Intelligent and easy-to-understand tabular summary in text-based output

```
Total IO
thread |       bytes     |     I/Os     |    MB/s   | I/O per s |  AvgLat  | LatStdDev |  file
------------------------------------------------------------------------------------------------------
     0 |         9101312 |         1111 |      4.34 |    555.48 |    7.193 |     5.687 | s:\diskspd\io.dat (50GB)
     1 |         9256960 |         1130 |      4.41 |    564.98 |    7.073 |     6.121 | s:\diskspd\io.dat (50GB)
     2 |         8814592 |         1076 |      4.20 |    537.98 |    7.437 |     5.880 | s:\diskspd\io.dat (50GB)
     3 |         9101312 |         1111 |      4.34 |    555.48 |    7.181 |     5.555 | s:\diskspd\io.dat (50GB)
------------------------------------------------------------------------------------------------------
total:          36274176 |         4428 |     17.30 |   2213.90 |    7.219 |     5.817

Read IO
thread |       bytes     |     I/Os     |    MB/s   | I/O per s |  AvgLat  | LatStdDev |  file
------------------------------------------------------------------------------------------------------
     0 |         7397376 |          903 |      3.53 |    451.48 |    8.614 |     5.384 | s:\diskspd\io.dat (50GB)
     1 |         7528448 |          919 |      3.59 |    459.48 |    8.480 |     5.954 | s:\diskspd\io.dat (50GB)
     2 |         7086080 |          865 |      3.38 |    432.48 |    9.018 |     5.498 | s:\diskspd\io.dat (50GB)
     3 |         7282688 |          889 |      3.47 |    444.48 |    8.740 |     5.136 | s:\diskspd\io.dat (50GB)
------------------------------------------------------------------------------------------------------
total:          29294592 |         3576 |     13.97 |   1787.92 |    8.708 |     5.508

Write IO
thread |       bytes     |     I/Os     |    MB/s   | I/O per s |  AvgLat  | LatStdDev |  file
------------------------------------------------------------------------------------------------------
     0 |         1703936 |          208 |      0.81 |    104.00 |    1.026 |     0.370 | s:\diskspd\io.dat (50GB)
     1 |         1728512 |          211 |      0.82 |    105.50 |    0.948 |     0.335 | s:\diskspd\io.dat (50GB)
     2 |         1728512 |          211 |      0.82 |    105.50 |    0.955 |     0.328 | s:\diskspd\io.dat (50GB)
     3 |         1818624 |          222 |      0.87 |    111.00 |    0.938 |     0.284 | s:\diskspd\io.dat (50GB)
------------------------------------------------------------------------------------------------------
total:          6979584 |          852 |      3.33 |    425.98 |    0.966 |     0.332
```

I/O Analyzer

A virtual appliance solution

Provides a simple and standardized way of measuring storage performance.

http://labs.vmware.com/flings/io-analyzer

- Readily deployable virtual appliance
- Easy configuration and launch of I/O tests on one or more hosts
- I/O trace replay as an additional workload generator
- Ability to upload I/O traces for automatic extraction of vital metrics
- Graphical visualization

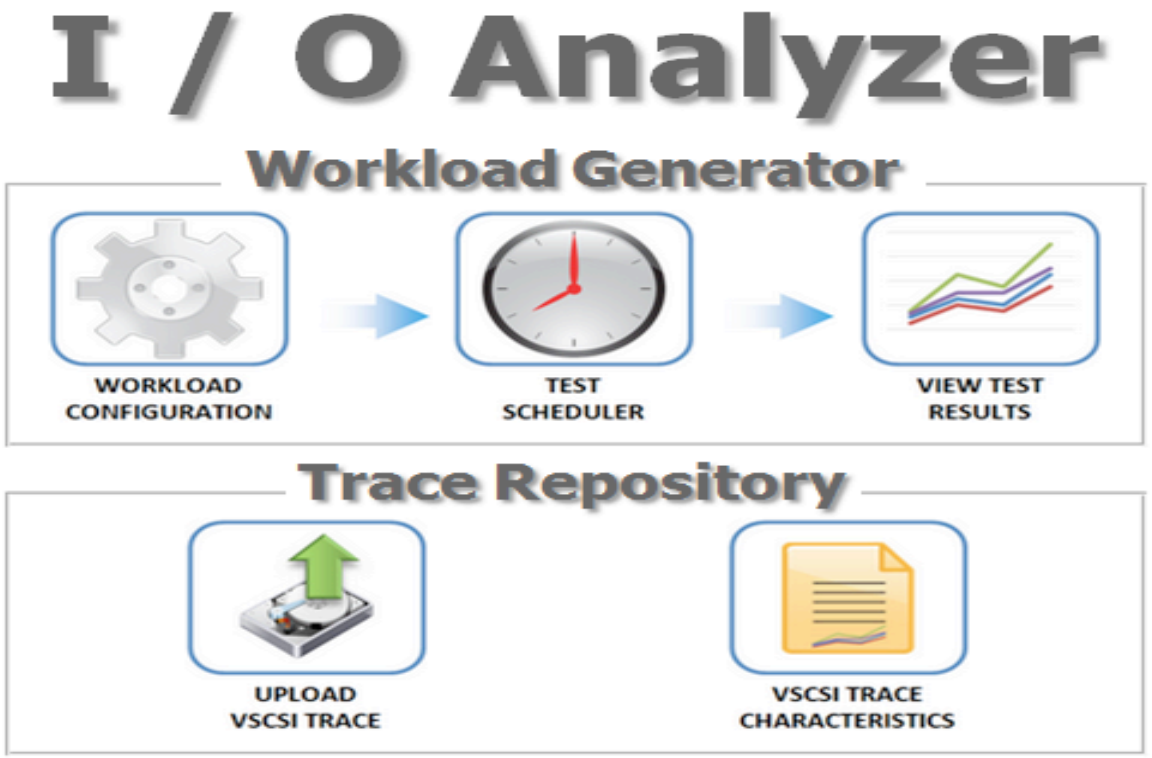

IO Blazer

Multi-platform storage stack micro-benchmark.

Supports Linux, Windows and OSX.

http://labs.vmware.com/flings/ioblazer

- Capable of generating a highly customizable workloads
- Parameters like: IO size, number of outstanding Ios, interarrival time, read vs. write mix, buffered vs. direct IO
- IOBlazer is also capable of playing back VSCSI traces captured using vscsiStats.
- Metrics reported are throughput and IO latency.

```
IOBlazer, the ultime IO benchmark

usage: ioblazer [options]

        Options      Description                                    Default

        -a <s|r>     IO access pattern: sequential or random        r
        -b <size>    Memory buffer size in MB                       1 MB
        -B           Buffered IOs (go thru FS buffer cache)         0
        -c           Computes pseudo-checksum on data               0
        -d <path>    Device or File path                            d:\tmp
        -f <size>    File  or Device size in MB                     1 MB
        -F           Fill the test file before starting IOs         0
        -g <gap>     Inter-burst time in microseconds (Avg)         0 us
        -G <f|u>     Inter-burst time pattern (fixed or uniform)    f
        -h           Print this message
        -i <size>    IO size in B (Avg)                             8192 B
        -I <f|u|b>   IO size pattern (fixed, uniform or bimodal)    f
        -l <thresh>  Latency Alarm Threshold in ms                  5000 ms
        -o <#OIOs>   Burst size [aka Otstanding IOs] (Avg)          32
        -O <f|u>     Burst size (fixed or uniform)                  f
        -p <c|f|z>   Printout format: csv, free, csv no header      f
        -P <file>    Playback VSCSI trace from 'file'               NULL
        -r <ratio>   Read ratio (1.0 = 100% read)                   1.0
        -R           Raw Device Access                              0
        -t <time>    Test duration in s                             180 s
        -w <num>     Number of worker threads                       1
```

Disaster Recovery with VMware Site Recovery Manager (SRM)

Architectural model #1 – Dedicated 1 to 1 Architecture

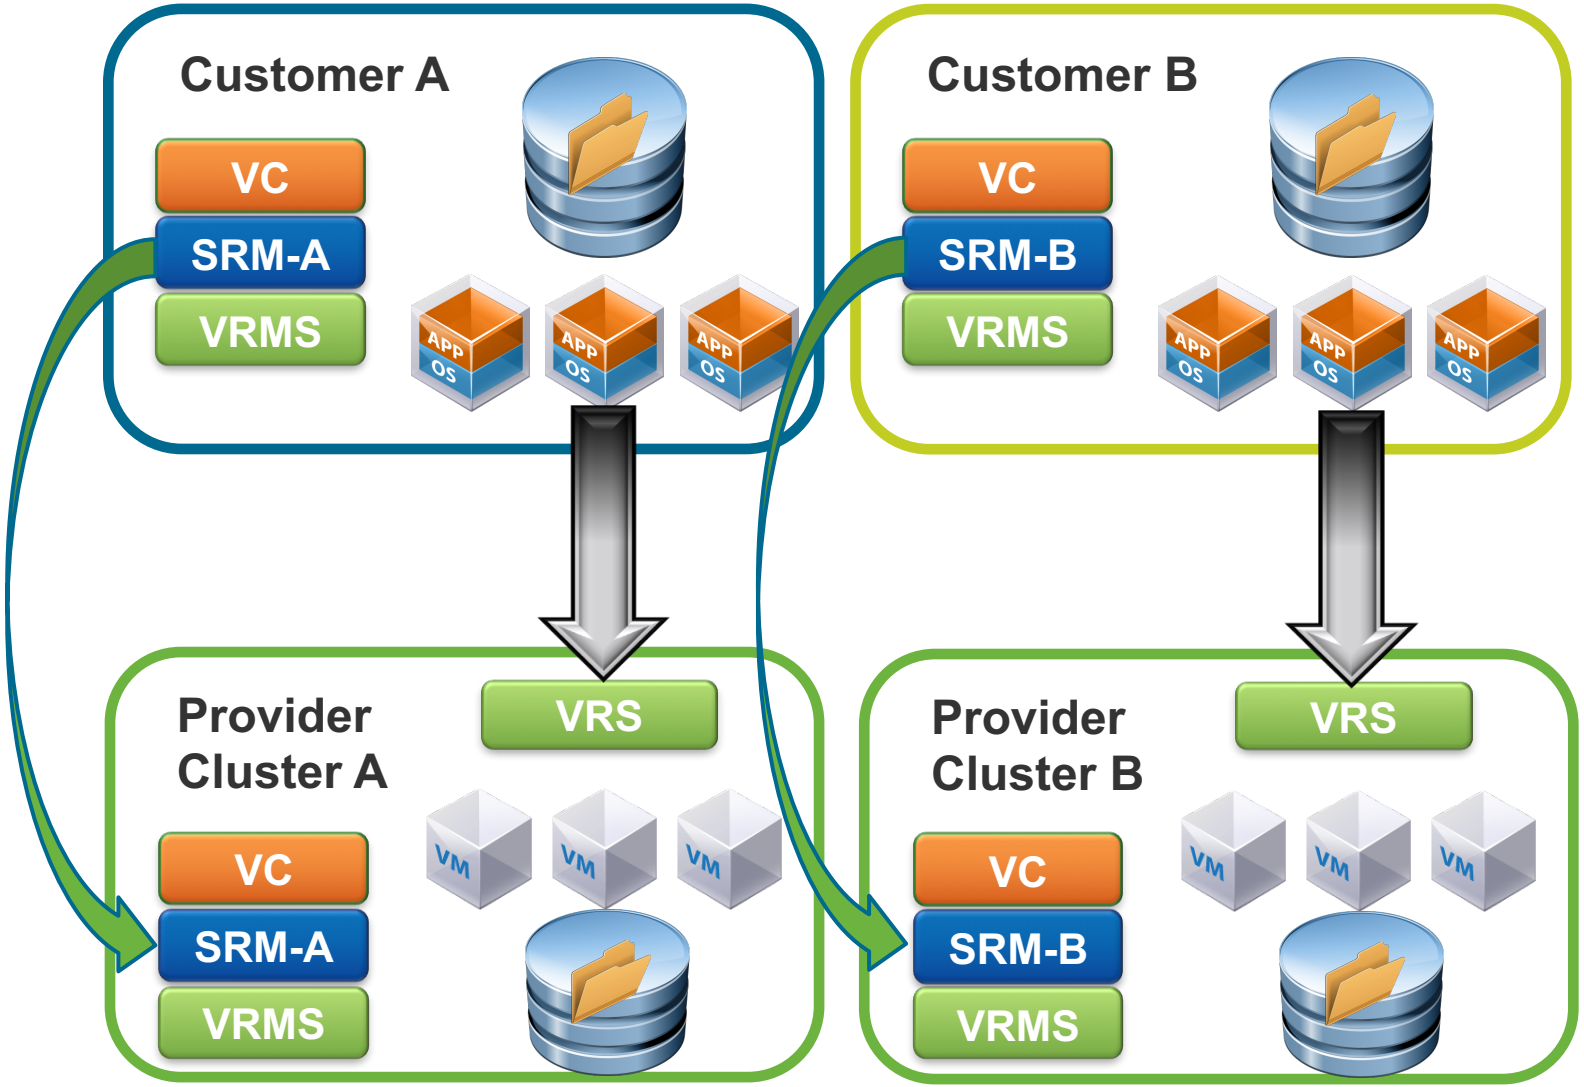

Pros and Cons of 1 to 1 paired architecture

Pros	Cons
Ensures customer isolation	Highest cost model
Dedicated resources per consumer	High level of ongoing management
Can provide full admin rights to consumers	Wasted resources during non-failover times
Easy self-service for consumers	
Well known and traditional model for configuration	
Easy upgrades	
Custom options allowable per consumer	

Use Case – Shared N to 1 Architecture

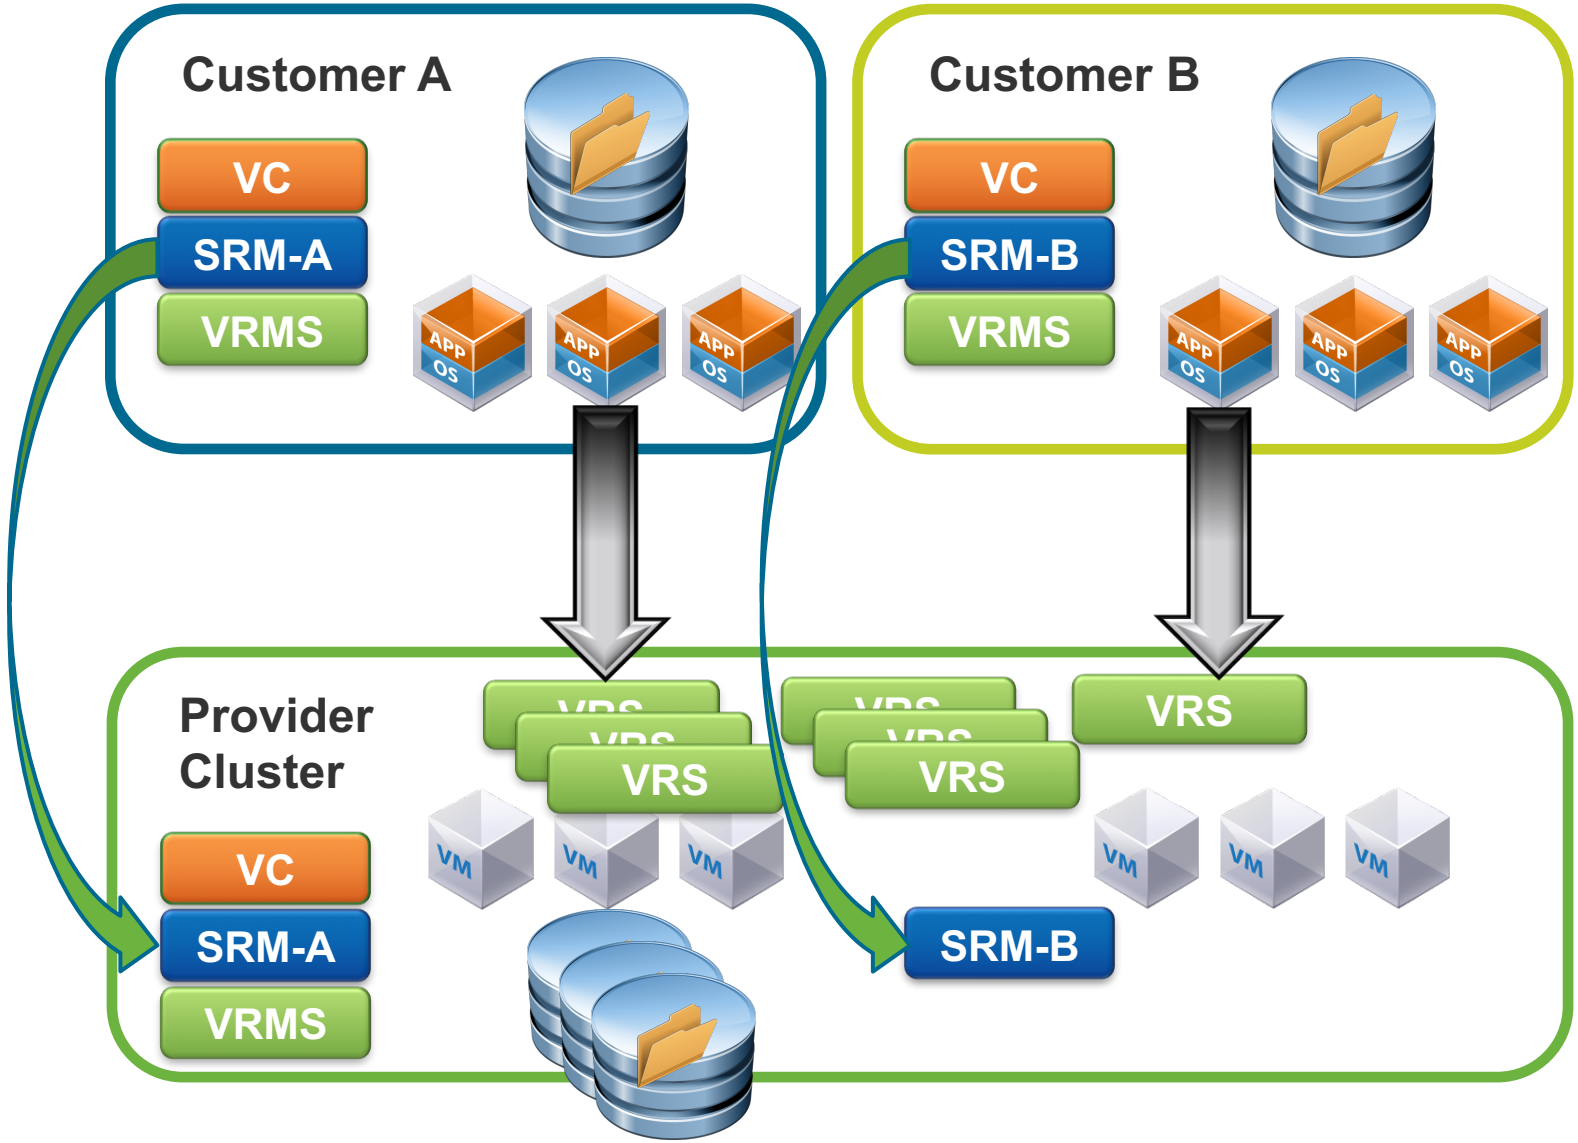

Use Case - DR as a Service (DRaaS) N:1 provider model layout

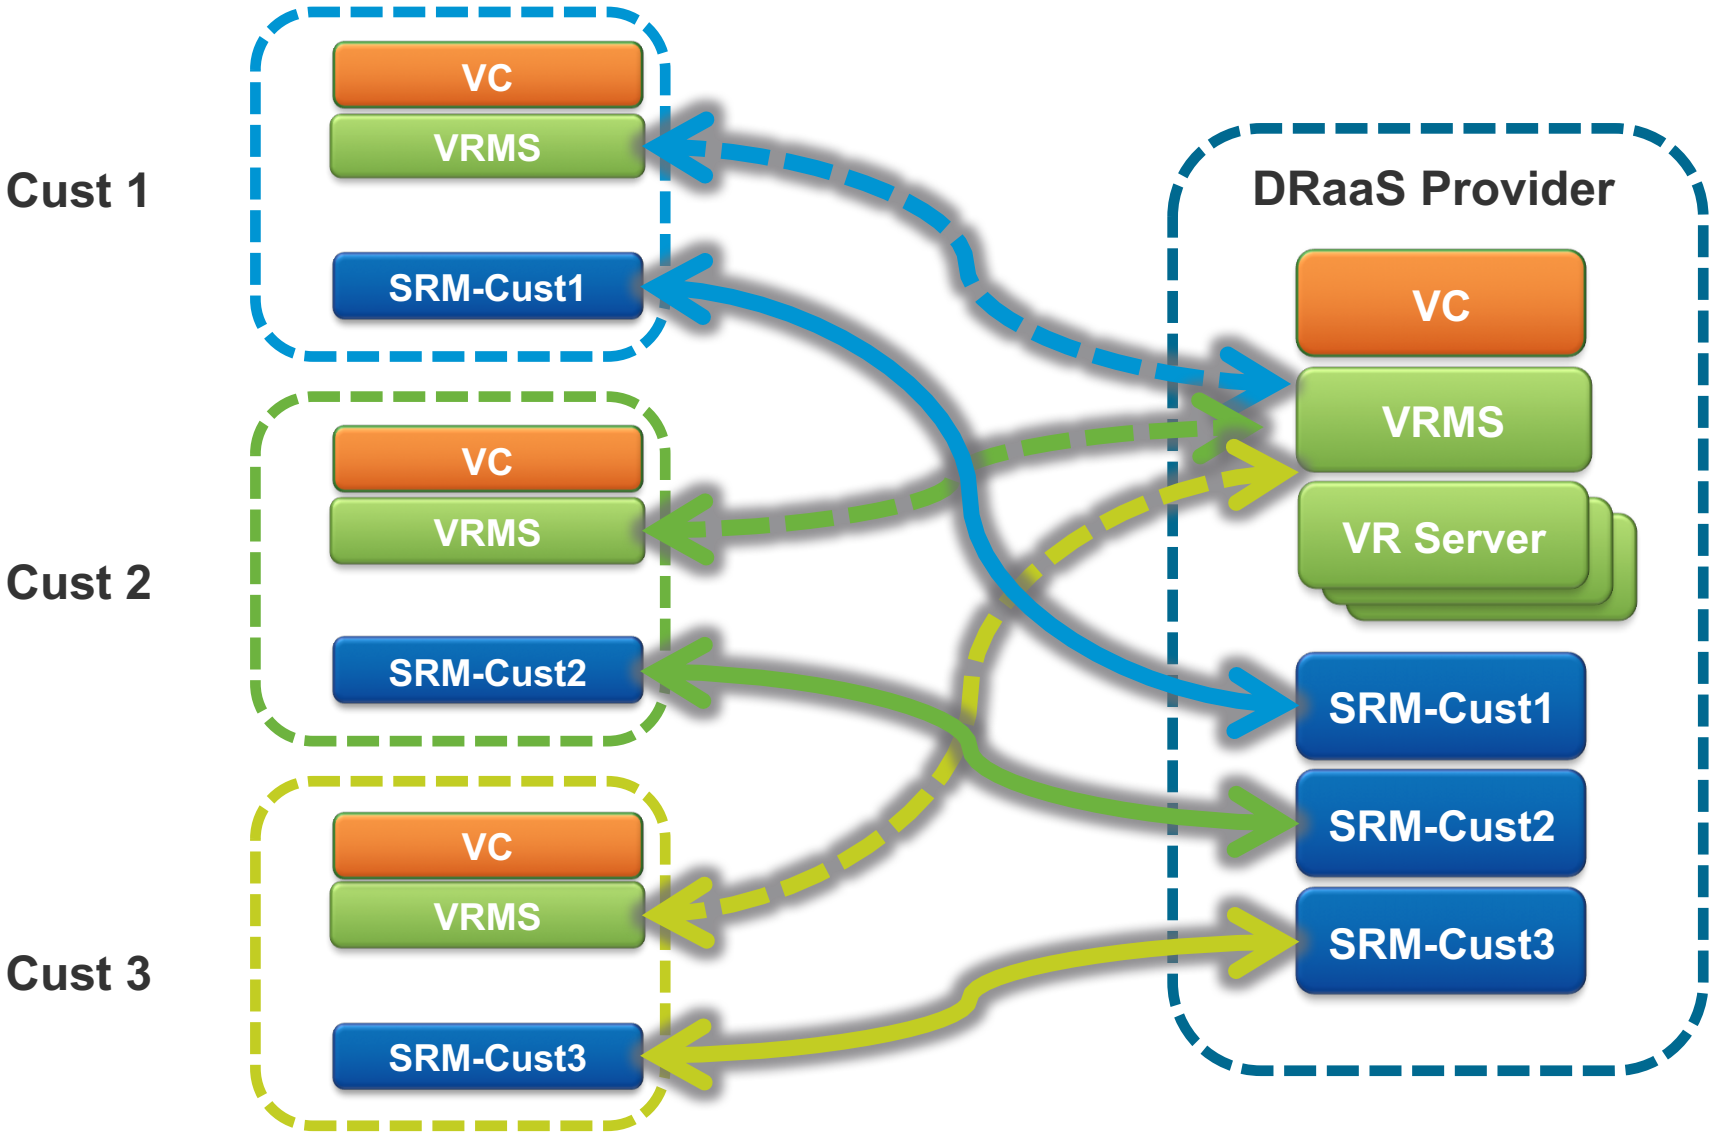

Use Case – DR as a Service (DRaaS) provider model

- **Minimum Component Requirements**
  - Same as site-to-site requirements
  - Remote customer site SRM “pairs” installed using SRM shared site option
  - Remote customer site VRMS connection paired to recovery site VRMS as per default VRMS setup
  - Typically provider runs whole solution as a managed service
  - Provider usually own / administer all component VM’s (SRM servers etc.) to reduce security complexities (i.e. user accounts / credentials)
  - Current targeted N:1 limit is 10:1 meaning for each vCenter at the provider site there can be up to 10 inbound customers. To go beyond this scale out by adding additional clusters with own dedicated vCenter/VRMS/VR components
  - Up to 500 VMs can be protected by VR under a single framework
- **Host Requirements**
  - Remote site VMs protected with vSphere Replication
    - You WILL need ESXi hosts to run those VMs on
    - Typically, provider will configure VR at customer site

Pros and Cons of Shared N to 1 architecture

Pros	Cons
Lower cost of infrastructure	Difficult coordinated upgrades
Ease of management	Difficult to isolate customer environments
Ease of scaling	Cluster-wide events affect every customer
Central management of customer environments	More difficult to provide self-service
	Requires extensive role and permission management
	Scalability limits of 10:1

Resources

VMware Hands-on Labs

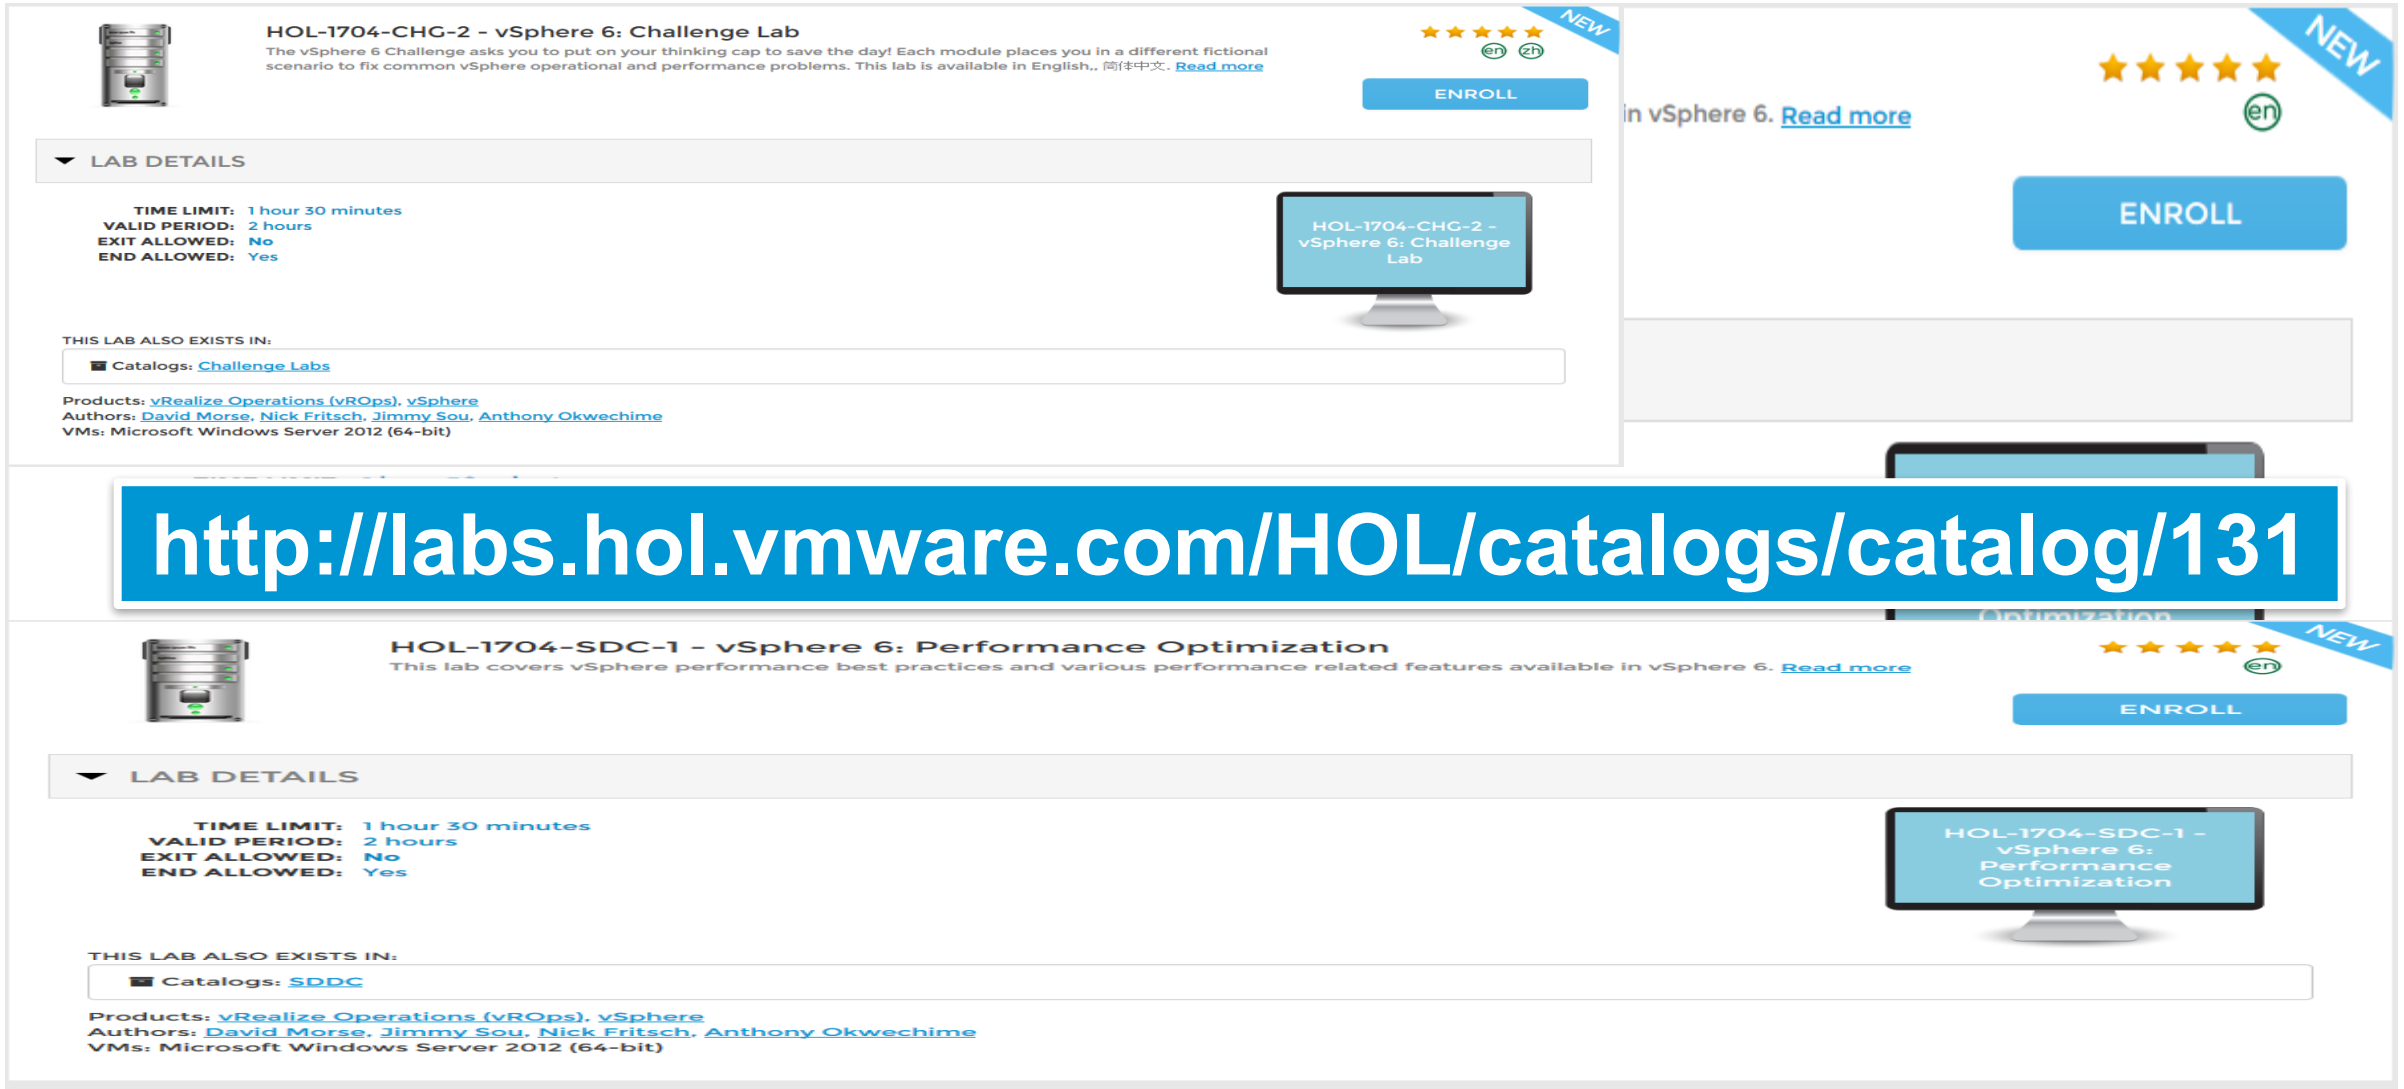

The Links are Free. Really

Virtualizing Business Critical Applications

- http://www.vmware.com/solutions/business-critical-apps/
- http://blogs.vmware.com/apps

VMware vSphere 6.5 Document

- https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html
- https://pubs.vmware.com/vsphere-65/index.jsp
- http://pubs.vmware.com/vsphere-65/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-65-setup-mscs.pdf

VMware's Performance – Technical Papers

- https://www.vmware.com/pdf/vsphere6/r65/vsphere-65-configuration-maximums.pdf
- http://www.vmware.com/files/pdf/techpaper/VMW-Tuning-Latency-Sensitive-Workloads.pdf
- http://pubs.vmware.com/vsphere-65/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-65-monitoring-performance-guide.pdf
- http://www.vmware.com/files/pdf/techpaper/VMware-PerfBest-Practices-vSphere6-0.pdf
- http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.5.pdf
- http://www.running-system.com/vsphere-6-esxtop-quick-overview-for-troubleshooting/ - ESXTop Cheat Sheet
- VMware vSphere Data Protection Documentation page

Questions? #rtfm

Thanks for attending