

Virtualizing Business Critical SQL Servers

Tips, Tricks, and Other Goodies to Ensure Your Success

SQL Saturday Lincoln – October 6, 2012

David Klee – Solutions Architect (@kleegeek)

Andy Galbraith – Technical Consultant (@DBA_Andy)



House of Brick[®]
T E C H N O L O G I E S



About House of Brick

- ❑ 14 year old Omaha-based company
 - ❑ Leader: Tier-1 VMware, Database Performance
 - ❑ Rock-solid reputation for optimizing the entire system stack to maximize Tier-1 performance
- 
- ❑ House of Brick key service value components
 - ❑ Hybrid/private cloud architectures for complex Tier-1 workloads
 - ❑ Legacy to virtualization, and private/hybrid cloud system replatforming
 - ❑ SQL Server and Oracle virtualization specialties
 - ❑ Short term assessments and proof-of-concept projects
 - ❑ Long-term project analysis, PM, implementation, & validation



About Us



Microsoft
CERTIFIED
IT Professional

Database Administrator 2008
Database Administrator on SQL Server® 2005
Database Developer 2008

David Klee

@kleegeek

- ❑ SQL Server on VMware team lead
- ❑ Experience in VMware, Microsoft, Linux, networking, security, application development technologies



Andy Galbraith

@DBA_Andy

- ❑ SQL Server Consultant
- ❑ Experience in SQL Server, clustering, performance tuning





- Why Virtualize Business Critical SQL Servers?
- Physical Stack Fundamentals
- VM-Layer Fundamentals
- Designing the New Infrastructure
- Performance Baselineing, Benchmarking, and Monitoring
- Disaster Recovery, High Availability



Your Trends

- ❑ Do you have a virtualization strategy?
- ❑ Have you tried to virtualize SQL Server?
- ❑ Was it successful?
- ❑ Are your production SQL Servers virtual?
- ❑ Why not?
- ❑ Do you have a plan for 100% virtual Tier-1?

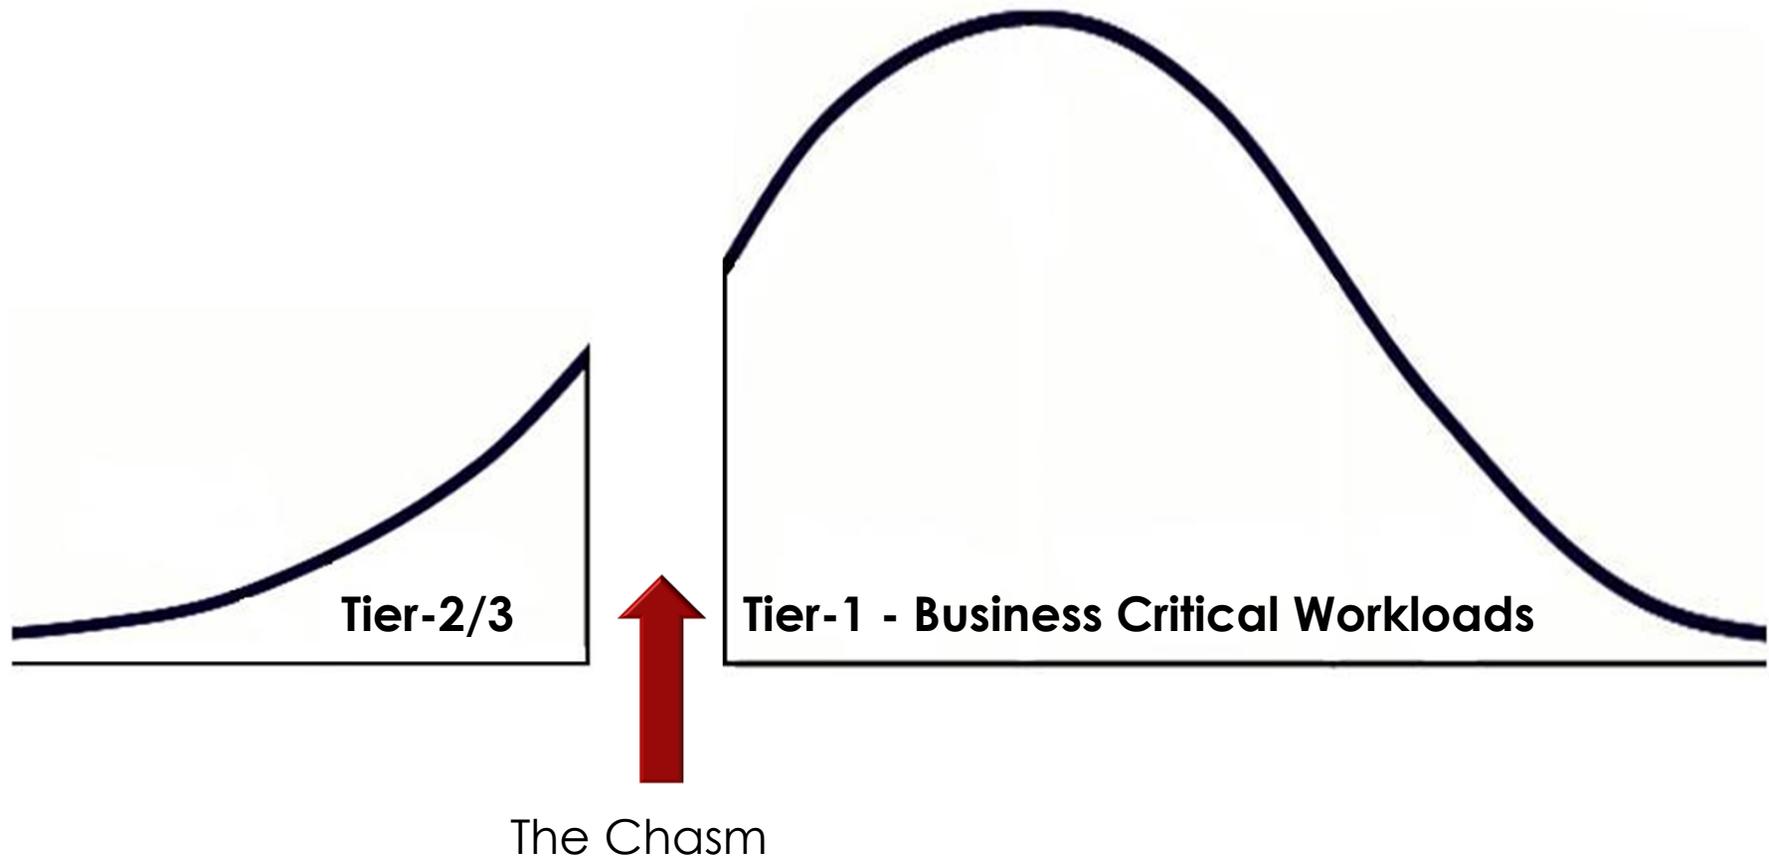


Six Reasons to Virtualize

- ❑ Consolidation
- ❑ Near 100% reliable Disaster Recoverability (DR)
- ❑ Product release cycle optimization
- ❑ High Availability (HA)
- ❑ Cost Reduction, including license optimization
- ❑ Security



Tier-1 vs. Everything Else



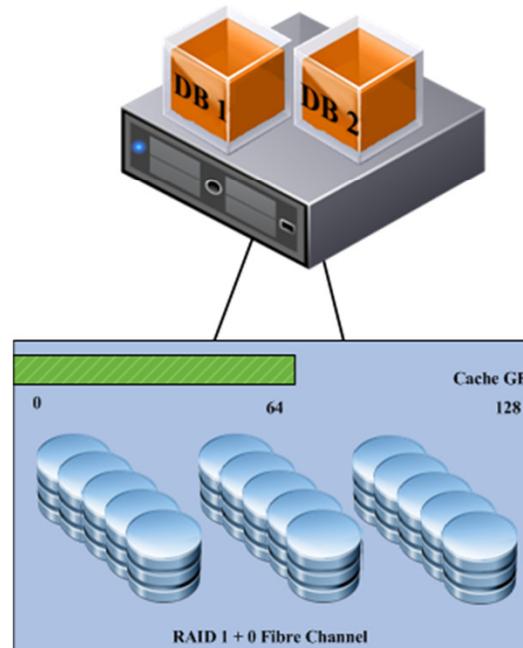


Myth – Cannot Run Tier-1 Virtualized

Common concerns

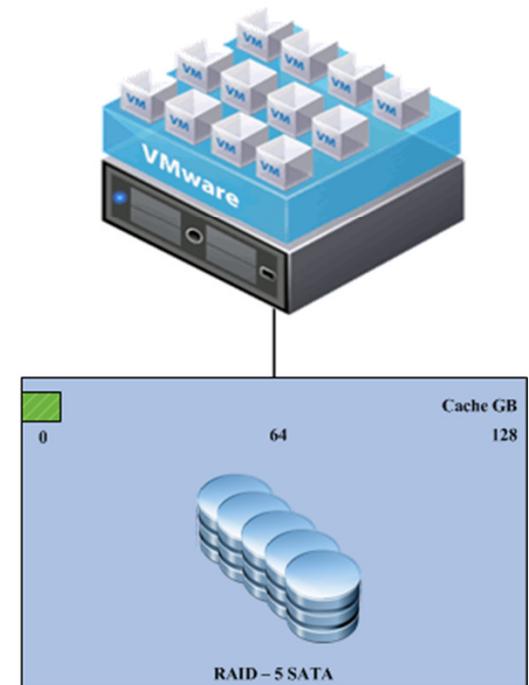
- Performance
- Licensing
- Support
- Database size
- Most problems: I/O
- Most Difficult Problem: Environmental, not technical

Production



VMware Test

- DB 1
- E-mail
- CDAP
- DB 2
- Utility
- Image
- DB 3
- & etc



Apples to Apples Proofing Environment?



Performance Concerns

□ Dev and QA traditionally isolated

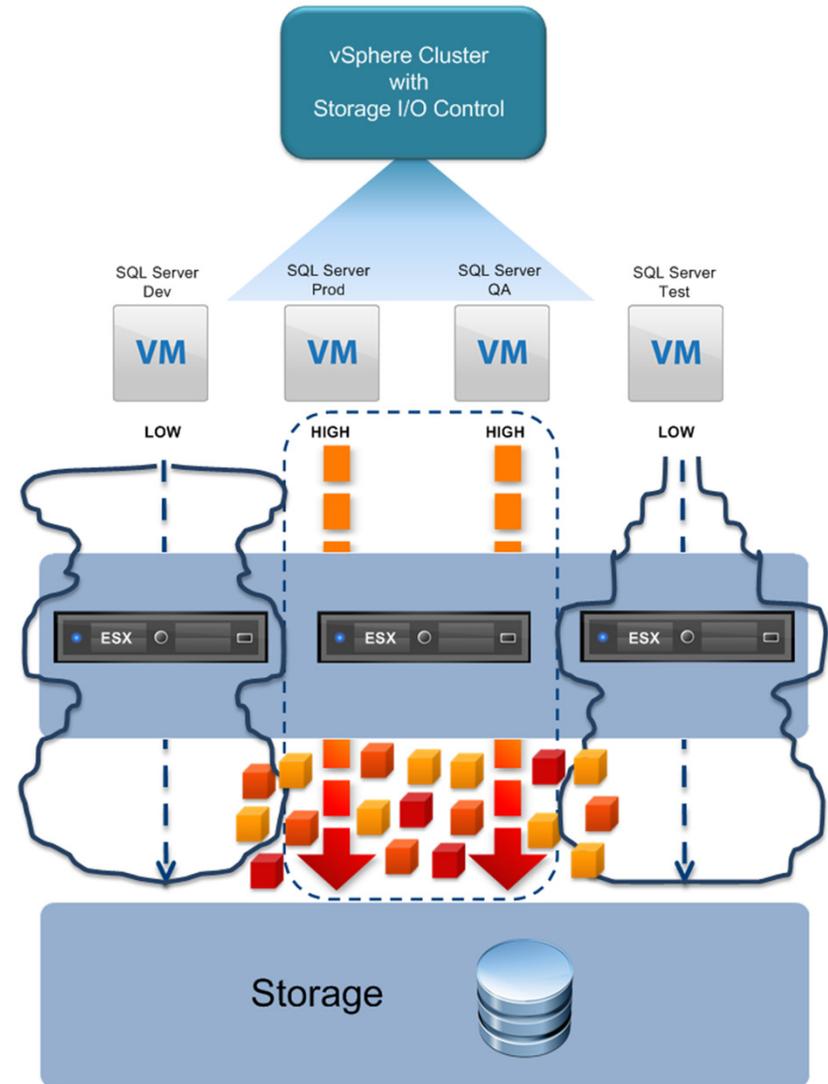
- Separate servers
- Hardware partitioning
- Avoid “Noisy Neighbor” problem

□ The new paradigm

- x86 commodity hardware
- vSphere DRS
- Lower CAPEX

□ vSphere benchmarks

- 2008 – 102K IOPs, 1.4% I/O wedge
- 2010 – 1M IOPs, 100 microsecond wedge





Support Concerns

- ❑ Policy for support of VMware published for years
- ❑ Officially supported via Server Virtualization Validation Program (SVVP)
- ❑ Known issues support
- ❑ VMware Customer Support Statement and TSA Net
- ❑ http://vmware.com/support/policies/ms_support_statement.html

Microsoft KB 897615



- Official support statement
- VMware offers full support and total ownership

Database Size Myths

❑ Database size has no impact on performance. *Period.*

❑ Database performance factors

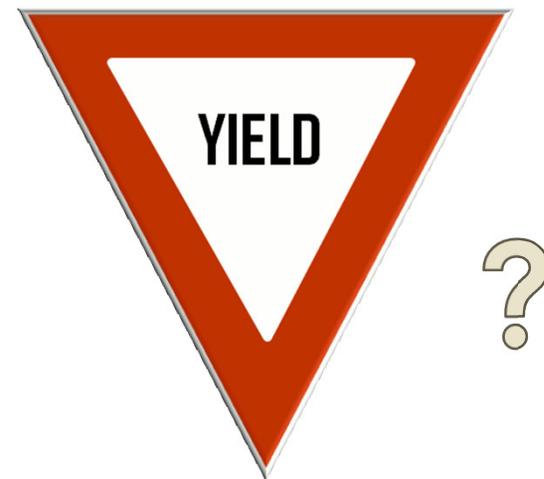
- ❑ Execution counts
- ❑ Concurrent connections
- ❑ SQL I/O access paths

❑ Large database concerns

- ❑ Backup/Recovery throughput
- ❑ DR operations
- ❑ One-time migration

❑ No distinction between physical and virtual

Databases in the TB's





Licensing Concerns

Dedicated SQL Server vSphere cluster

- Maximum consolidation
- Maximum SQL Server license optimization
- Maximum utilization of vSphere tooling

vSphere SQL Server Sub-cluster

- Frequently reduces SQL Server license requirements
- Consolidation ratios tend to be lower than dedicated SQL Server clusters



Licensing (SQL Server 2008)

- Consolidation Rules (SQL Server 2008/2008 R2)
 - Datacenter Edition – Unlimited VMs if all sockets in host are licensed
 - Enterprise Edition – 4 VMs per license if all sockets in host are licensed



Licensing (SQL Server 2012)

- ❑ Consolidation Rules (SQL Server 2012)
 - ❑ There is no more Datacenter Edition
 - ❑ Enterprise Edition
 - ❑ Unlimited consolidation if and only if:
 - ❑ All cores (not sockets) in host are licensed
 - ❑ Software assurance (SA)
 - ❑ Without SA, only 1 VM/core license
 - ❑ Minimum of 4 core licenses required per socket, even if only single or dual core
 - ❑ Unlimited VMs with SA if upgrading from SQL Server 2008 EE + SA



Licensing Details

- **When VMs move, licenses don't necessarily move with them:**
 - With Standard Edition, licenses can only proactively move **once every 90 days** in a **non-failover situation**.
 - With Enterprise and Datacenter Editions, you have unlimited license mobility rights
 - 2012 also require Software Assurance (SA) for unlimited mobility!

- **For more details on licensing SQL Server on vSphere, refer to Microsoft.com:**
 - 2008R2: SQL2008R2_LicensingQuickReference-updated.pdf
 - 2012: SQL Server 2012 Licensing Quick Reference Guide.pdf

- **NOTE: Licensing Individual VMs is different!**
 - Still must obey 4 core minimum!



Virtual Storage Presentation

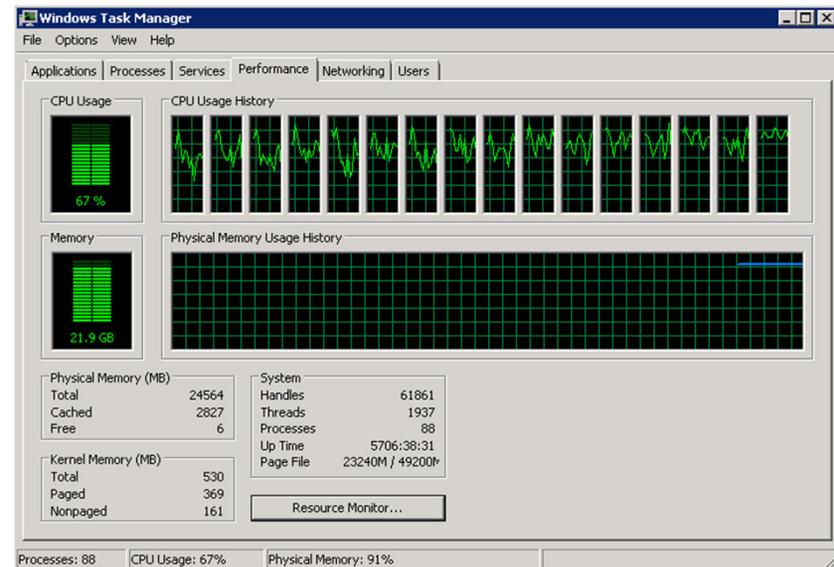
- ❑ Virtual Machine Disk (VMDK/VMFS)
 - ❑ Preferred for Tier-1
 - ❑ Maximum ESXi-level storage tooling
- ❑ Raw Device Mapping–Virtual (RDM-V)
 - ❑ More configuration and operations overhead
 - ❑ Reduced ESXi-level storage tooling
- ❑ RDM-P
 - ❑ Maximizes SAN-level tooling transparency
 - ❑ Even less ESXi-level storage tooling
 - ❑ No snapshots or vMotion
- ❑ Direct-mounted (In-guest iSCSI)
 - ❑ No ESXi-level storage tooling
 - ❑ vMotion works





Tier-1 Analysis - Profiling Physical

- ❑ CPU – Count, Speed, Average, and Peak Utilization
- ❑ Memory – Amount, Utilization, Active
- ❑ Windows Perfmon – CPU, RAM, Network, Disks statistics
- ❑ Disk Performance
- ❑ Storage Benchmarking
- ❑ Environmental Growth Planning
- ❑ SQL Server health check



<http://sqlserverperformance.wordpress.com>



Tier-1 Analysis – Storage

- SQLIO – I/O performance and latency benchmark

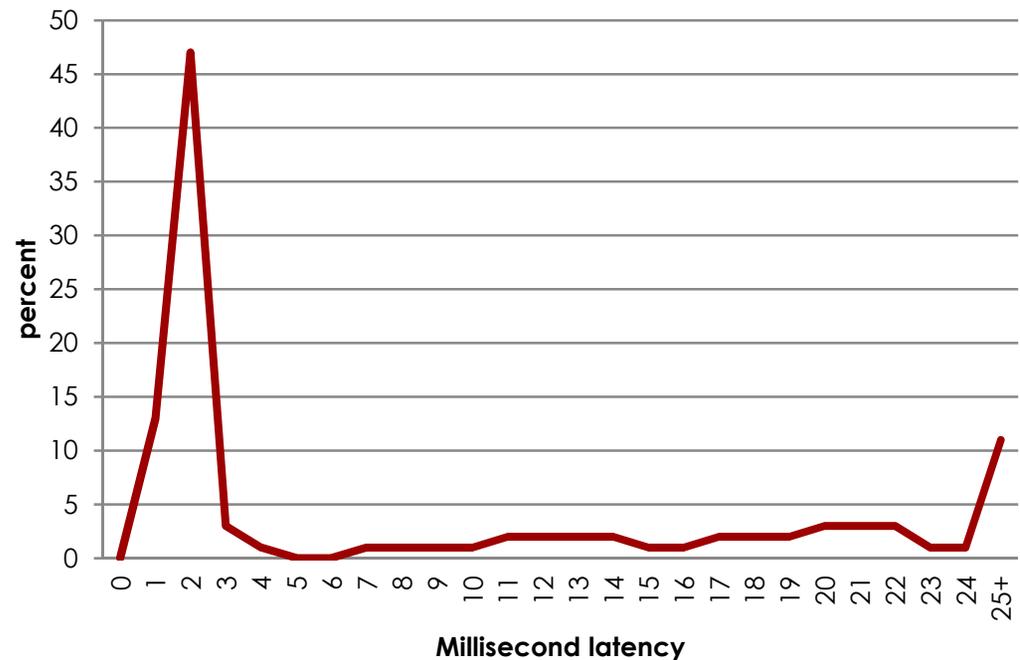
- Throughput metrics:

- IOs/sec
- MBs/sec

- Latency metrics:

- Avg. (ms): < 25

SQLIO Latency Histogram



- Microsoft Download: SQLIO Disk Subsystem Benchmark Tool

- tools.davidklee.net – SQLIO Analyzer



Tier-1 Analysis – Storage

□ IOMeter – I/O Performance Stressor

□ www.iometer.org

Read/Write %	Type	Block	Threads / # Outstanding per Thread	Similar to...
80/20	Random	8K	# cores / start low & increase until IOPs limit reached	Typical OLTP data files
0/100	Sequential	2K - 60K	1 / 1	Transaction Log
100/0	Sequential	64K - 512K	1 / 16	Table Scans
0/100	Sequential	256K	1 / 16	Bulk load
100/0	Random	32K	# cores / 1	SSAS Workload
100/0	Sequential	1MB	1 / 32	Backup
0/100	Random	8K-256K	# cores / start low & increase until IOPs limit reached	Checkpoints

Source: SQLPASS 2011 DBA-413-M Slide 48



Tier-1 Analysis – Perfmon

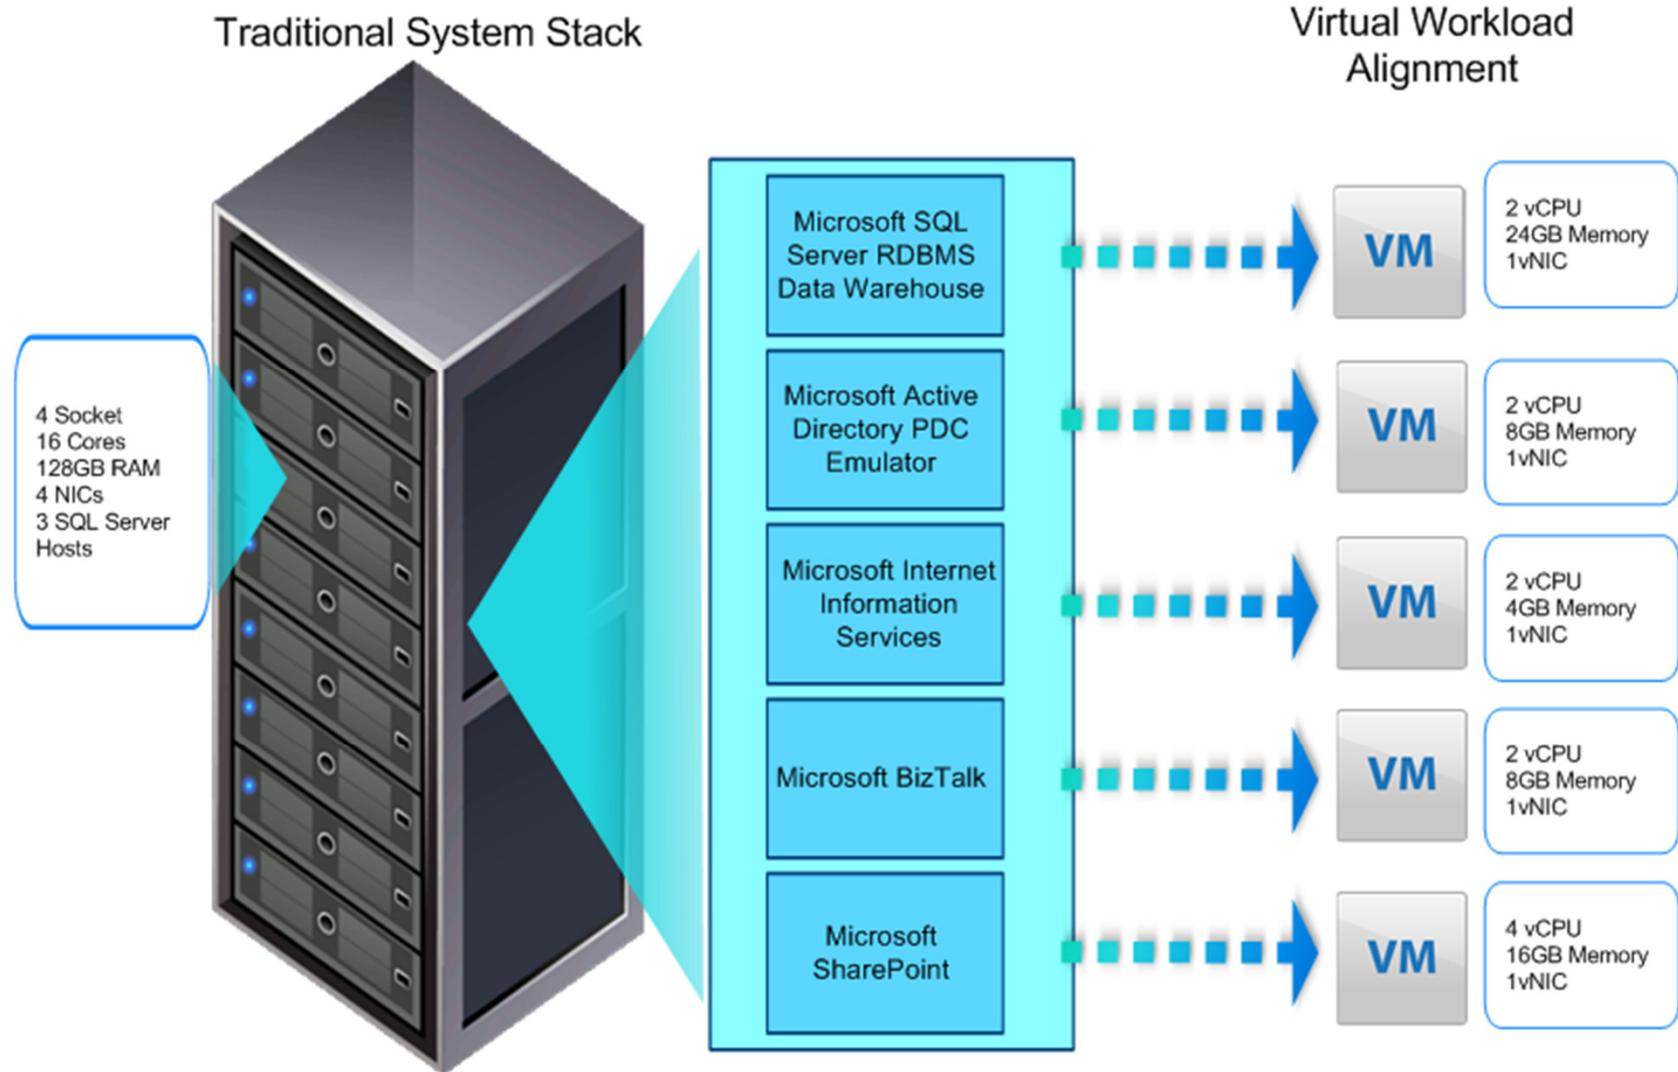
- ❑ Memory: Available Mbytes, Pages / sec, Page Faults / sec
- ❑ Processor: % Processor Time
- ❑ Physical Disk: Disk Read Bytes/sec and Disk Write Bytes/sec
- ❑ System: Processor Queue Length
- ❑ SQL Server: Databases: Transactions / sec
- ❑ SQL Server: SQL Statistics: SQL Compilations / sec, SQL Re-Compilations / sec, Batch Requests / sec



Now What?

- ▣ Design your virtual environment carefully.
- ▣ Scale Up? More vHardware on less VMs.
- ▣ Scale Out?
 - ▣ More VMs?
 - ▣ More instances on less VMs?

Atomic Workloads





VMware Hardware Configuration

- ❑ Disable BIOS “green” settings (power savings, etc.)
- ❑ Ensure CPUs are set to high performance mode
- ❑ Enable virtualization extensions (i.e. Intel VT-x)
- ❑ Disable Automatic Server Recovery (HP)
- ❑ Enable Hyper-Threading (Intel)
- ❑ 64-bit VMs preferred (32-bit is soon to be extinct)



Tier-1 - vCPU

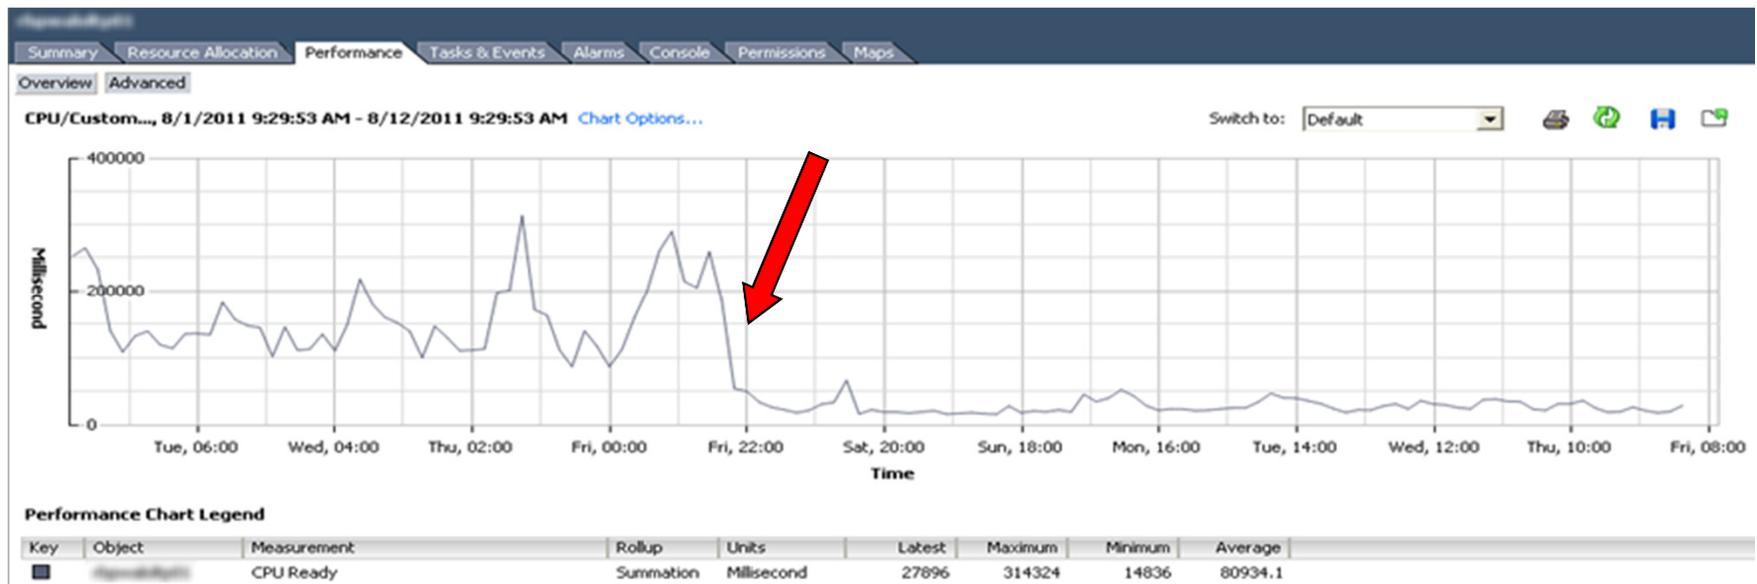
- Start conservative. Do not over-allocate vCPUs.

- vCPU Ready Time

 - 300ms average

 - 500ms high water mark

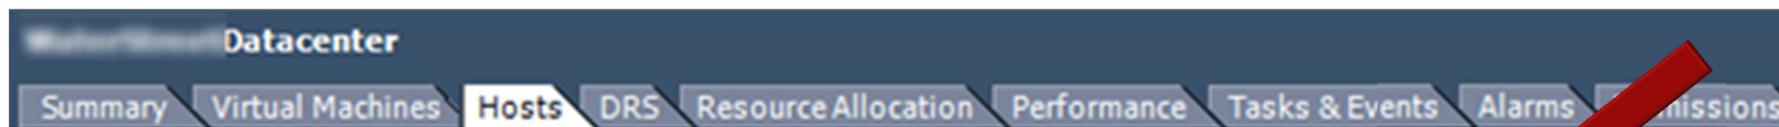
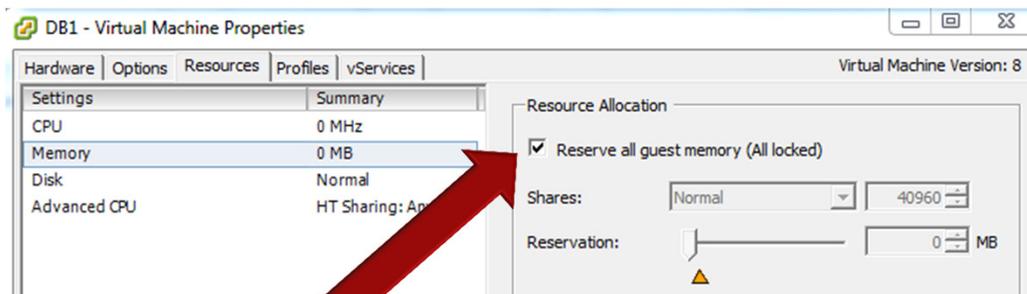
CPU measures the amount of time a virtual machine waits in the queue in a ready-to-run state before it can be scheduled on a CPU. Higher wait times result in slower virtual machine performance.





Tier-1 - vMemory

- Full RAM reservations for production Tier-1 workloads
- Do NOT oversubscribe
- Do NOT over-allocate host RAM
- No ballooning allowed! (Don't disable balloon driver)



Name	State	Status	% CPU	% Memory	Memory Size
esxi1	Local Connected	Warning	5	91	36850.66 MB
esxi2	Local Connected	Normal	4	86	36850.66 MB
esxi3	Local Standby	Normal			36850.66 MB
esxi4	Local Standby	Normal			36850.66 MB



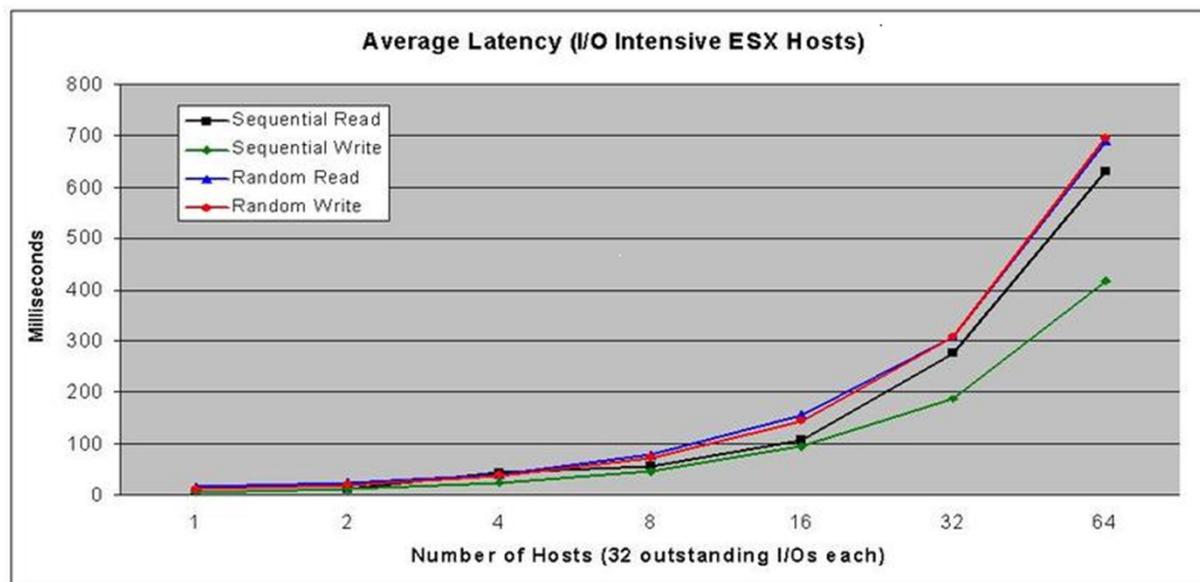
Tier-1 - vStorage

- Performance is top priority
 - RAID-10 or RAID-5
 - *I'll repeat. Performance is **TOP** priority.*
- **Independence Rule.** Design for absolute workload object independence
 - *Monitor usage and split workloads if red lines are exceeded*
 - Can go down to one disk group per LUN
 - Can go down to one LUN per VMware datastore
 - Can go down to one datastore per virtual hard drive file (VMDK)



Tier-1 - vStorage

- Storage performance attributes
 - <25ms average I/O latency times
 - <50% average spindle busy
 - >60MB/s sustained writes (at a bare minimum after piercing cache)

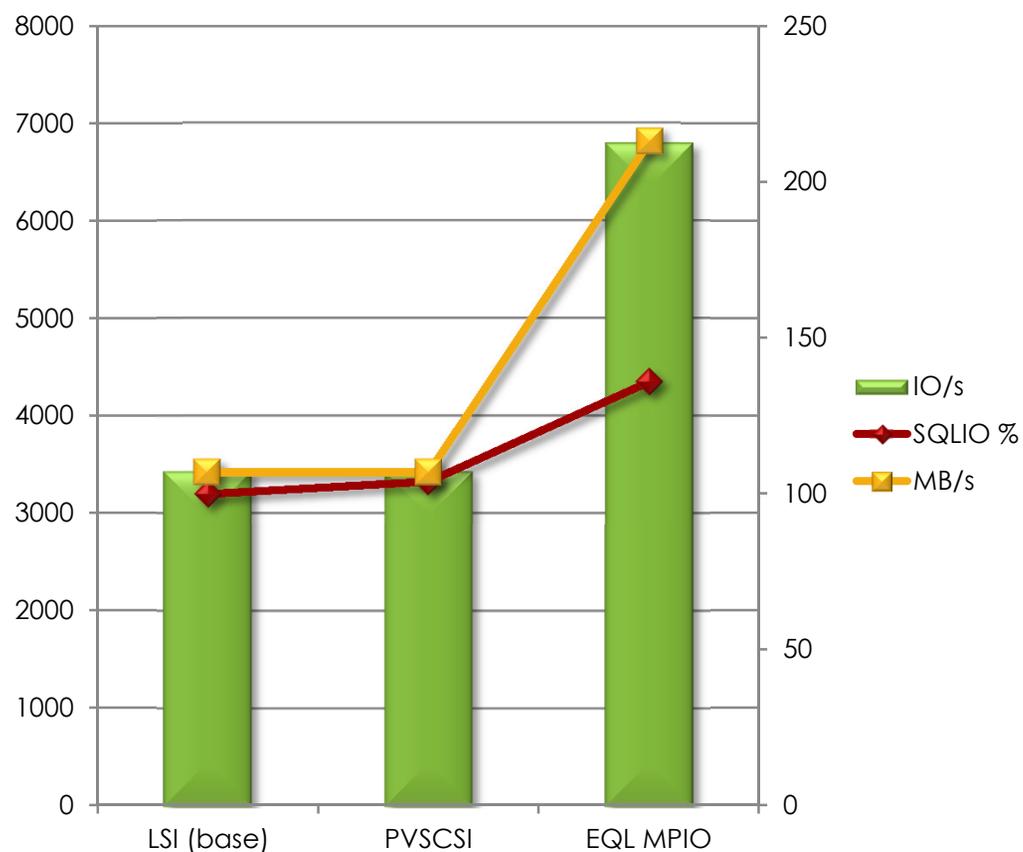




Tier-1 - vStorage

- ❑ Paravirtual (PVSCSI) Driver
- ❑ Multipathing Drivers
 - ❑ EMC PowerPath VE
 - ❑ Equallogic MPIO
- ❑ Profiling with benchmarking tools
 - ❑ SQLIO
 - ❑ IOMeter

Storage Driver Improvements

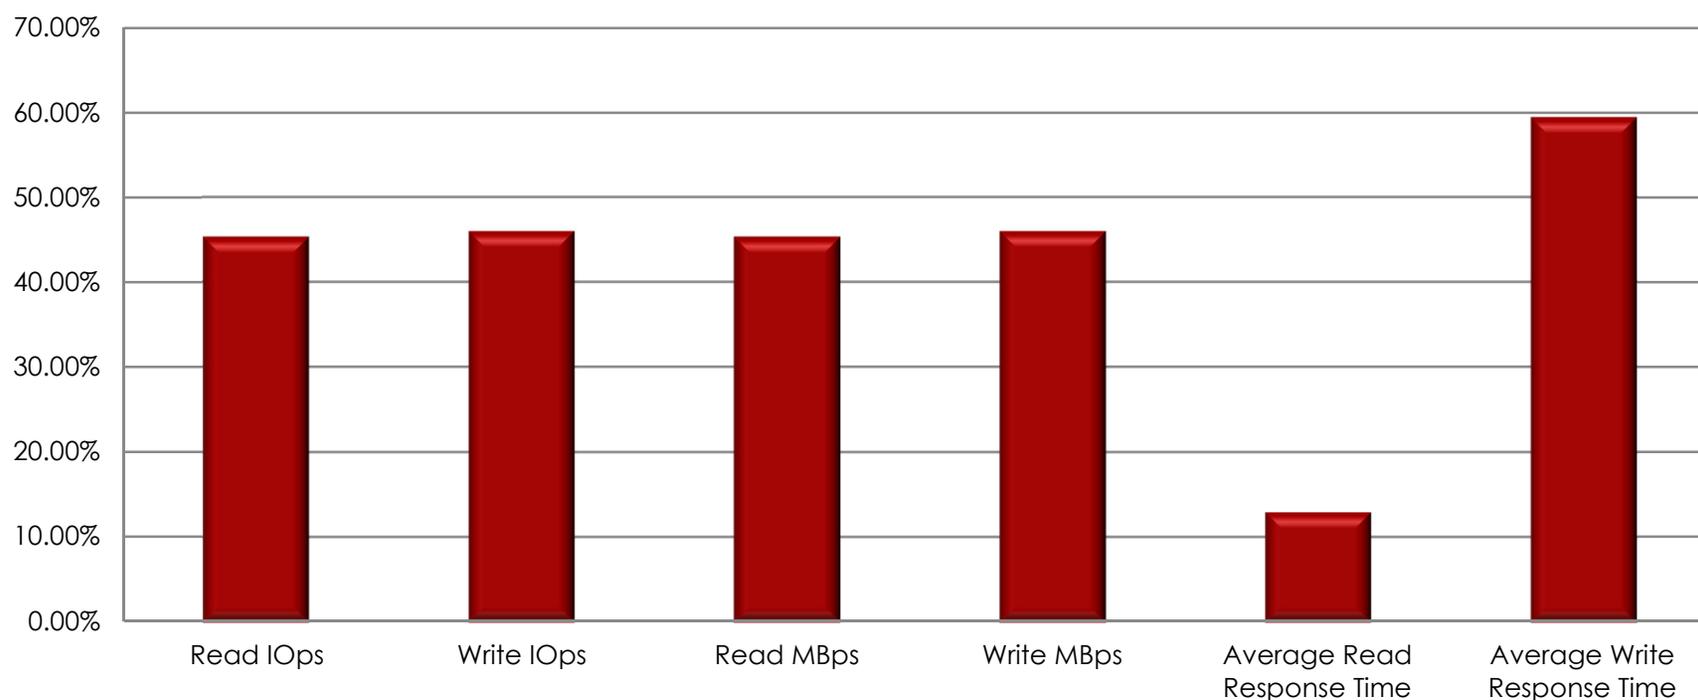




Tier-1 - vStorage

- ❑ Disk block alignment (not just a Windows 2003 problem)
 - ❑ (46% sequential I/O, 24% random I/O performance improvement)
 - ❑ 1MB offset usually works great

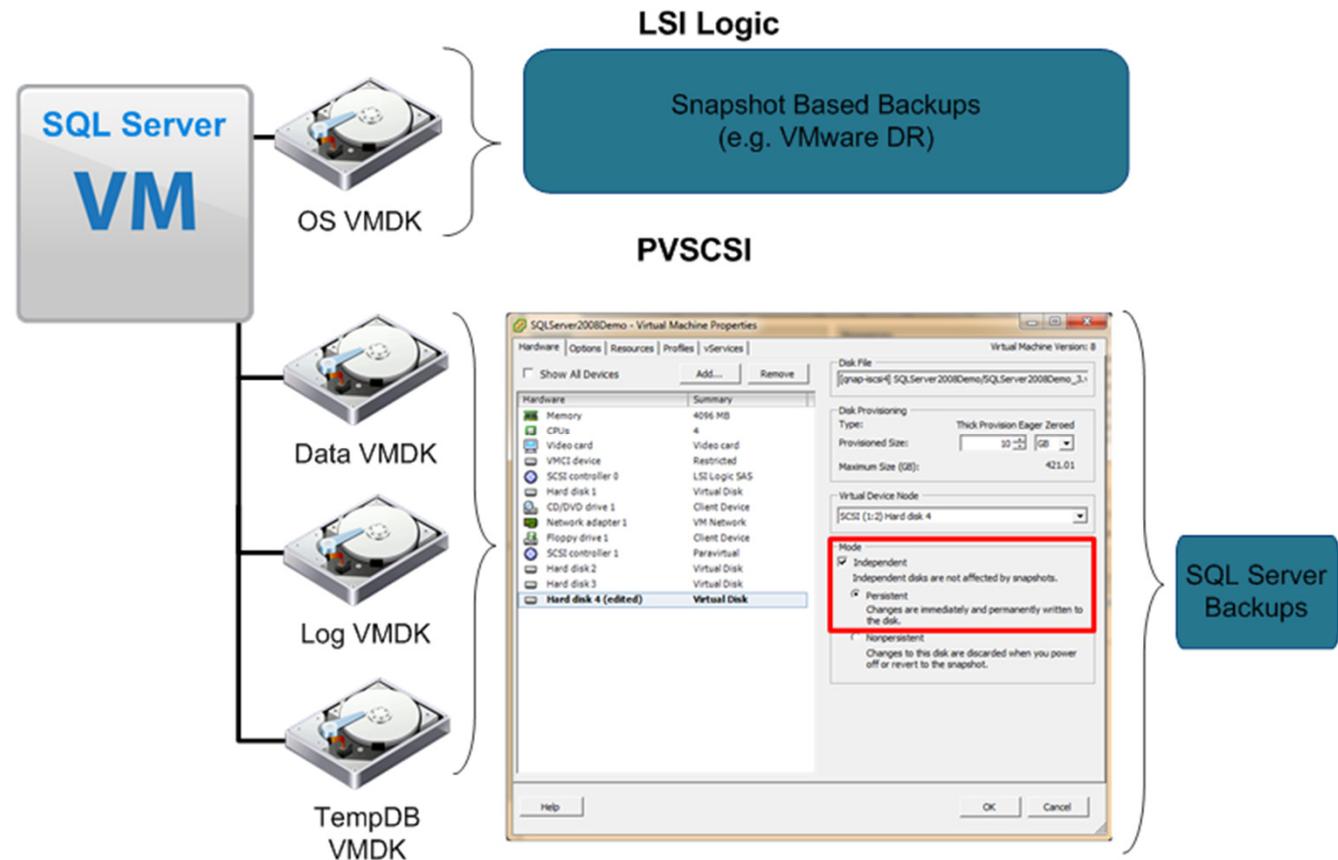
Improvements from Disk Partition Alignment





Installing a SQL Server Instance

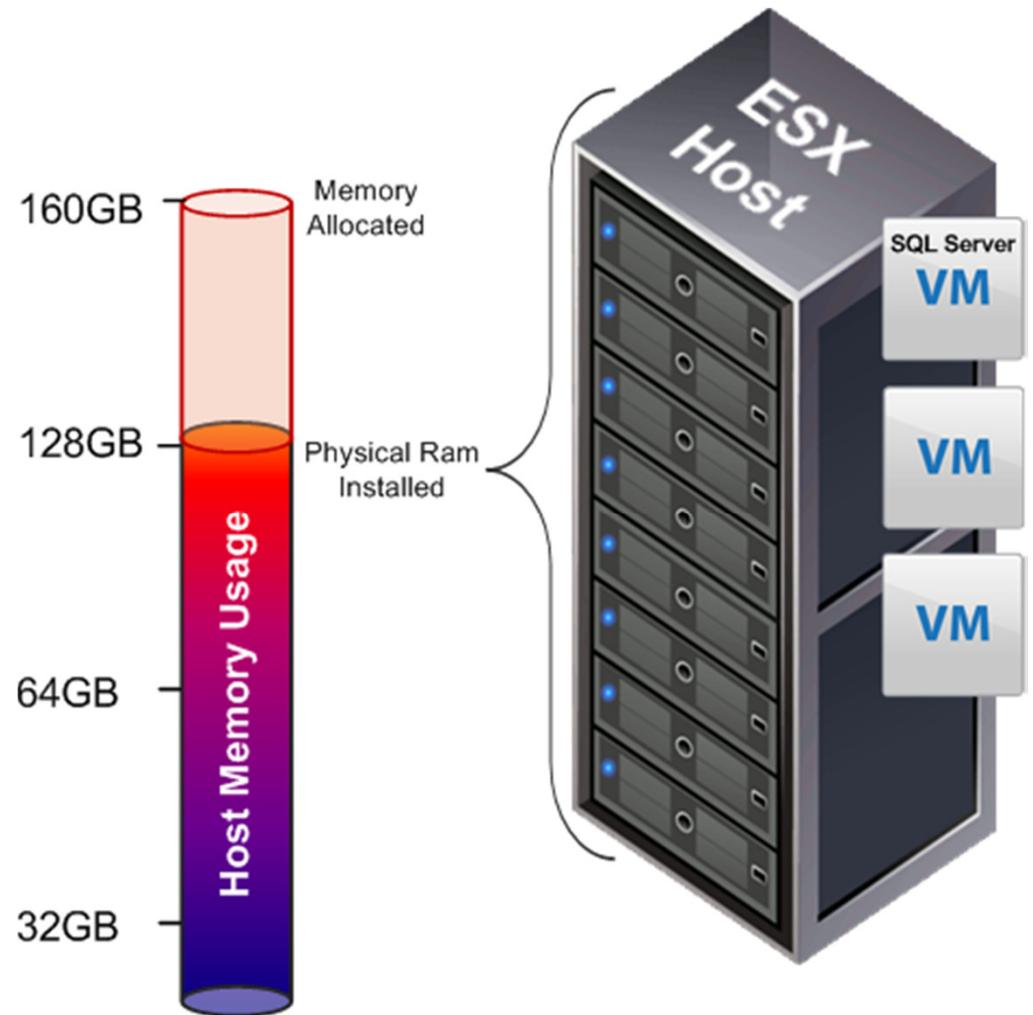
- Object separation can optimize:
 - Performance
 - Disaster recovery
 - Backup
 - Licensing





Configuring a SQL Server Instance

- ❑ Enable Lock Pages in Memory
- ❑ Enable Instant File Init
- ❑ Use Large Pages – Trace Flag 834
- ❑ VM RAM Reservation
 - ❑ Memory Provisioned
 - ❑ SQL Server memory + OS + VM overhead
- ❑ Set “Max Server Memory” and “Min Server Memory”
- ❑ Enable Optimize for Ad-hoc Workloads





Monitoring Performance

- ▣ Perfmon / IOMeter / SQLIO / DVDStore
- ▣ vCenter Statistics
- ▣ SQL Server health checks
 - ▣ sqlserverperformance.wordpress.com
 - ▣ brentozar.com/blitz
 - ▣ bradmckehee.com/healthcheck.zip
- ▣ Benchmark and compare to baselines (physical and virtual)
- ▣ Remember to update your baselines when the configuration changes!



MSFC Clustering vs. VMware HA

- ❑ Four MSFC evaluation criteria:
 - ❑ Less than four minute SLA?

then

- ❑ Rolling maintenance utilized
- ❑ Cluster-aware middle tier?
- ❑ Does technical expertise exist to support clustering?



Criteria answers = “Yes”

- Consider clustered SQL Server on VMware HA

Otherwise

- Other options exist



SQL Server 2012

- ▣ AlwaysOn + VMware = Complementary technologies
- ▣ Blurs line between HA and DR
- ▣ Best of MSFC and Mirroring
- ▣ Current best practices directly apply to 2012
- ▣ Watch your licensing





SQL Server on VMware Boot Camp

VIRTUALIZING SQL SERVER ON VMware BOOT CAMP

Tuesday November 6th, 2012

HELD AT SQL PASS SUMMIT

Aspen Room

Sheraton Seattle Hotel

1400 Sixth Avenue, Seattle, WA 98101

11:00 – 12:00 Registration, Lunch, and Networking

12:00 – 4:30 Boot Camp Intensive



BOOT CAMP COURSE OVERVIEW

Intro to SQL Server Virtualization

- Virtualization Trends
- Common Objections and Misconceptions

Physical Stack Fundamentals

- SQL Server Licensing Concepts
- Storage, vSphere Host, and Networking

Virtual Machine Layer

- VM and Guest Operating System Customization
- Virtual Storage Presentation Options
- Installation and Optimally Configure SQL Server

SQL Server on VMware Prototype

- Benchmarking and Baselineing Performance
- Workload Selection

Beyond the Prototype

- Disaster Recovery and HA Options
- SQL Server Clustering
- SQL Server 2012 AlwaysOn

House of Brick
E D U C A T I O N

More information at <http://bit.ly/Plne7f>. Register at <http://bit.ly/PUKkbO>.



Questions



SQL Server Virtualization

Tips, Tricks, and Other Goodies to Ensure Your Success

SQL Saturday Lincoln – October 6, 2012

David Klee – Solutions Architect (@kleegeek)

Andy Galbraith – Technical Consultant (@DBA_Andy)



House of Brick[®]
T E C H N O L O G I E S