

Squeezing The Most Performance

from your VMware-based SQL Server

SQL Saturday Atlanta – May 18, 2013

David Klee – Principal Architect (@kleegeek)



© 2013 House of Brick Technologies, LLC



About HoB

- ▣ Founded in 1998
- ▣ Partner-Focused Strategy
- ▣ House of Brick Key Services



- ▣ **Virtualization and Cloud Computing — VBCA**
- ▣ **Replatforming and Data Migration**
- ▣ **Managed Services**

© 2013 House of Brick Technologies, LLC



About David





David Klee

 @kleegeek

 bit.ly/HoBSA

davidklee.net

 gplus.to/kleegeek


 linkedin.com/a/davidaklee



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

- SQL Server Principal Architect, practice lead
- Experience in Microsoft, VMware, Linux, networking, security, application development technologies

© 2013 House of Brick Technologies, LLC



Agenda

- So you finally virtualized! Great! *Now what?*
- Do you know how well things were running when physical?
- VMware Infrastructure Tuning
- SQL Server VM Template Tuning
- Demonstrating Equivalent Performance
- Ongoing Steady-State Performance Monitoring

© 2013 House of Brick Technologies, LLC


Assumptions

- ▣ The following tips and tricks are geared for **single instance** SQL Server on VMware.
- ▣ 2005+ Database Mirroring and 2012 AlwaysOn Availability Groups also follow these configuration guidelines.
- ▣ SQL Server Failover Clustering uses a somewhat **different** set of rules.

© 2013 House of Brick Technologies, LLC


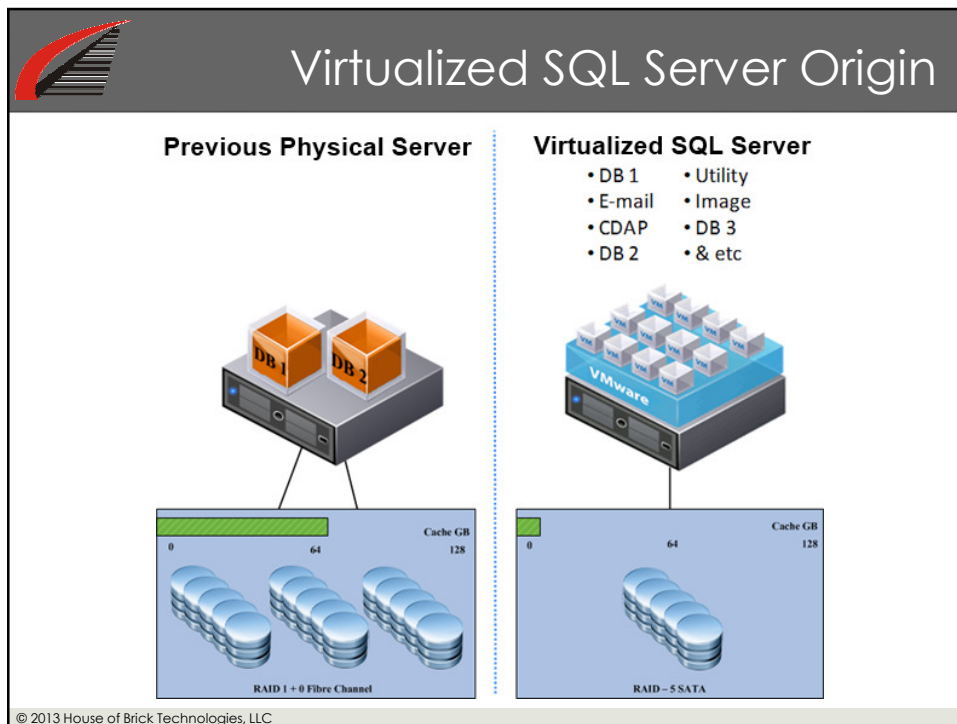
In The Beginning

- ▣ In the beginning of your virtualization journey...
- ▣ Me: "Did you performance test and compare against your established historical baselines to objectively demonstrate equivalent performance of the virtual machine?"
- ▣ Common Response: "Yeah, right. Sure. Uh-huh."



<http://www.flickr.com/photos/raggle/3617884814/>



© 2013 House of Brick Technologies, LLC





Benchmark - Storage

- ▣ Storage performance is my number one obstacle
- ▣ SQLIO & IOMeter



© 2013 House of Brick Technologies, LLC



Benchmark – Perfmon Stats

- ▣ Record every 5m, cycle log files nightly
- ▣ Collect all individual stats, not rollups

Counters:

Memory – Pages/sec, Available Mbytes, Pages / sec, Page Faults / sec
 Network Interface – Bytes total/sec
 Physical Disk – Disk Transfers/sec, Disk Read Bytes/sec and Disk Write Bytes/sec
 Processor - % Processor Time, Utilization By Core
 SQLServer:Access Methods - Full Scans/sec
 SQLServer:Buffer Manager – Buffer Cache Hit Ratio
 SQLServer:Databases Application Database - Transactions/sec
 SQLServer:General Statistics - User Connections
 SQLServer:Latches – Average Latch Wait Time
 SQLServer:Locks - Average Wait Time, Lock Timeouts/sec, Number of Deadlocks/sec
 SQLServer:Memory Manager - Memory Grants Pending

© 2013 House of Brick Technologies, LLC



Benchmark – SQL Server Instance

- ▣ DVDStore
 - ▣ <http://linux.dell.com/dvdstore>

Threads	Core x GHz	VM	VM	Itanium Physical	Itanium Physical	8CPU VM to Avg. Itanium Improvement
		8x2GHz	32x2GHz	8x1.6GHz Run 1	8x1.6GHz Run 2	
	MaxDOP	HP Orders	HP Orders	Itanium Orders	Itanium Orders	
2	1	19277	13589	15612	12853	35.44%
2	2	19251	17858	16368	17553	13.50%
2	3	18841	17453	17214	18209	6.38%
2	4	15839	15640	15147	18306	-5.31%
2	5	15953	15779	10201	16866	17.88%
2	6	16263	16055		17596	-7.58%
8	1	76590	63910	62896	60789	23.85%
8	2	76592	70705	69392	71523	8.71%
8	3	75441	69335	65737	69184	11.83%
8	4	57508	61412	44123	71918	-0.88%
8	5	55021	61579	42442 (extrapolated)	67073	0.48%

© 2013 House of Brick Technologies, LLC



Benchmark – DB Performance

- ▣ Every environment is different!
- ▣ Reports / Jobs / Etc.
 - ▣ Average Runtimes
- ▣ Query Performance
 - ▣ CPU Impact
 - ▣ Memory Impact
 - ▣ Storage Impact
 - ▣ TempDB Impact
- ▣ Application owners should be involved in benchmarking process
- ▣ Tools
 - ▣ Perfmon
 - ▣ Extended Events
 - ▣ Simple things:
 - ▣ Set statistics io on/off
 - ▣ Set statistics time on/off

© 2013 House of Brick Technologies, LLC



Baselines

- ▣ Baseline = Averages and peaks during routine activities
- ▣ Can help predict growth and resource contention
- ▣ Establish performance thresholds and high / low water marks
- ▣ Update your baselines after major system changes and/or fixed period of time
- ▣ Be consistent in your approach.
- ▣ Helpful to create a baseline repository
 - ▣ <http://www.sql-server-performance.com/2010/baseline-repository/>

© 2013 House of Brick Technologies, LLC



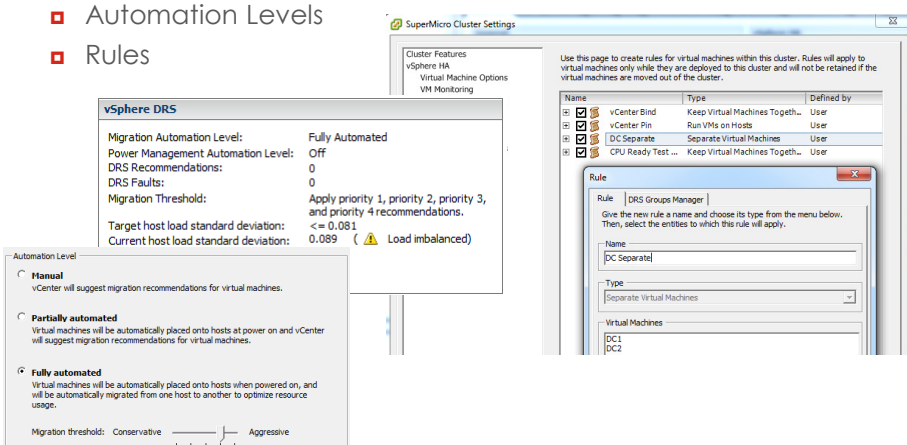
VMware Hardware Configuration

- ▣ 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)
- ▣ Multiple paths for everything (NIC, SAN, etc.)

© 2013 House of Brick Technologies, LLC

VMware Cluster Configuration

- VMware High Availability (HA)
- VMware Distributed Resource Scheduler (DRS)
 - Automation Levels
 - Rules

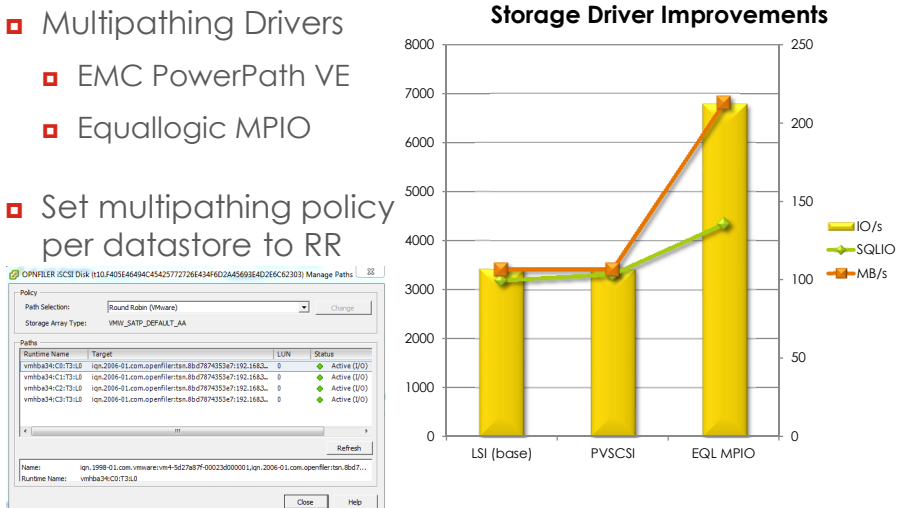


The screenshots show the vSphere DRS configuration interface. The first screenshot displays the 'vSphere DRS' settings, including Migration Automation Level (Fully Automated), Power Management Automation Level (Off), DRS Recommendations (0), DRS Faults (0), Migration Threshold (Apply priority 1, priority 2, priority 3, and priority 4 recommendations), Target host load standard deviation (<= 0.081), and Current host load standard deviation (0.089, Load imbalanced). The second screenshot shows the 'Automation Level' options: Manual, Partially automated, and Fully automated. The third screenshot shows the 'Rule' configuration for 'DC Separate', with the Type set to 'Separate Virtual Machines' and the Virtual Machines list containing 'DC1' and 'DC2'.

© 2013 House of Brick Technologies, LLC

VMware Storage - 1

- Paravirtual (PVSCSI) Driver
- Multipathing Drivers
 - EMC PowerPath VE
 - Equallogic MPIO
- Set multipathing policy per datastore to RR



The screenshots show the VMware storage configuration interface. The first screenshot displays the 'Policy' tab for a 'CINFLER (SCSI) Disk', showing the Path Selection (Round Robin (VMware)) and Storage Array Type (VMW_SATP_DEFAULT_AA). The second screenshot shows the 'Paths' table with columns for Runtime Name, Target, LUN, and Status. The third screenshot shows the 'Runtime Name' and 'Runtime Name' fields.

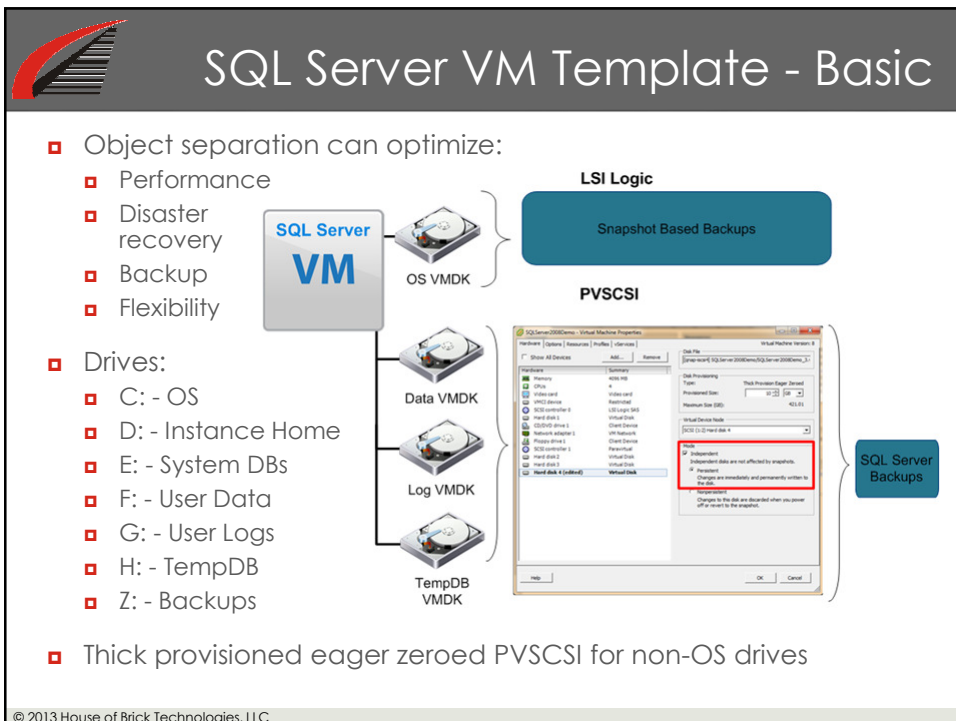
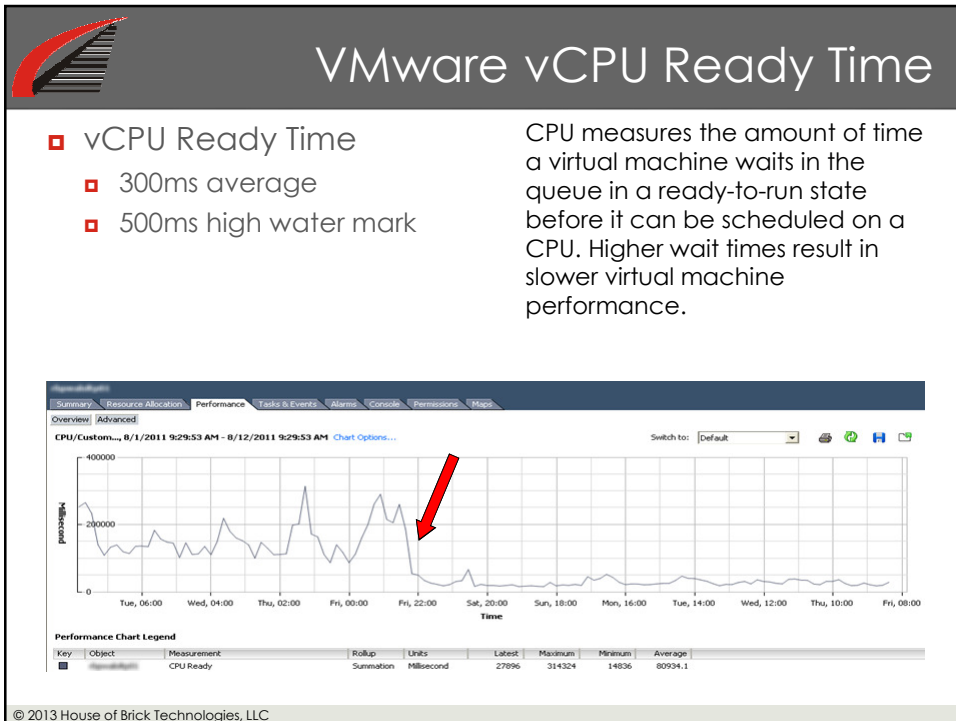
The bar chart, titled 'Storage Driver Improvements', compares the performance of three storage drivers: LSI (base), PVSCSI, and EQL MPIO. The chart shows three metrics: IO/s (yellow bars), SQLIO % (green line), and MB/s (orange line). The data is as follows:

Storage Driver	IO/s	SQLIO %	MB/s
LSI (base)	~3500	~100	~100
PVSCSI	~3500	~100	~100
EQL MPIO	~6800	~150	~200

© 2013 House of Brick Technologies, LLC

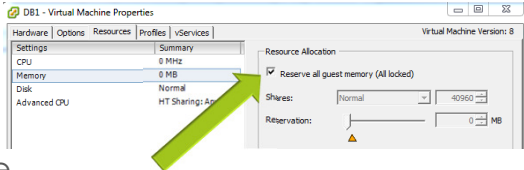
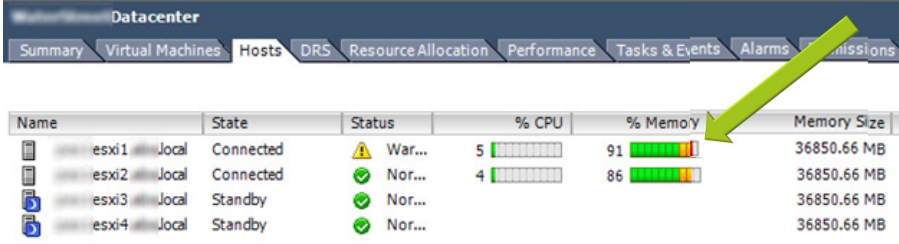
© 2013 House of Brick Technologies, LLC

- © 2013 House of Brick Technologies, LLC



VM Template - 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	5	91	36850.66 MB
esxi2	Local	Connected	4	86	36850.66 MB
esxi3	Local	Standby			36850.66 MB
esxi4	Local	Standby			36850.66 MB

© 2013 House of Brick Technologies, LLC


VM Template - Advanced

- Balance vCPU vCores / vSockets with physical architecture
- Configure vNUMA settings when < 8vCPUs
 - <http://bit.ly/13Hkob2>
- Set vCenter Alarms for high CPU Ready time
 - <http://bit.ly/XEnDKu>
- Resource Pools to ensure top-tier resource availability
- Want to *really* geek out? (Do not change these without reason)
 - <http://bit.ly/r726Wc> - VMware Tuning for Latency Sensitive Workloads




<http://www.flickr.com/photos/hyperxp/7698784450/>


© 2013 House of Brick Technologies, LLC



Configuring Windows for SQL Server

- VMware Tools installed and up to date
- WDDM video driver update
 - <http://bit.ly/9muzD7>
- 64KB NTFS allocation unit for non-OS volumes
 - <http://bit.ly/hWtxiA>
- Ensure proper disk alignment
 - No fancy partition creation tools!
- Antivirus engine rules
 - <http://bit.ly/B5qlE>
- Windows Power Options set to 'High Performance'

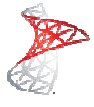





© 2013 House of Brick Technologies, LLC

Configuring a SQL Server Instance

- ▣ Enable Lock Pages in Memory *(weigh pros and cons)*
- ▣ Set "Max Server Memory" and "Min Server Memory"
 - ▣ <http://bit.ly/WVBfT>
- ▣ Enable Instant File Initialization
 - ▣ <http://bit.ly/XOKJOY>
- ▣ Use Large Pages – Trace Flag 834 (YMMV)
- ▣ Full VM RAM Reservation
- ▣ Enable Optimize for Ad-hoc Workloads
- ▣ Set TempDB data file count
- ▣ Set MaxDOP and Cost Threshold for Parallelism


 Microsoft
SQL Server 2008 R2


 Microsoft
SQL Server 2012

© 2013 House of Brick Technologies, LLC


Proper Maintenance

- ▣ Proper Maintenance is a Must!
- ▣ Fantastic database maintenance solution – ola.hallengren.com
 - ▣ Backups
 - ▣ Indexes / Statistics
 - ▣ Integrity Checks
 - ▣ Work / log file cleanup
- ▣ Quarterly Windows and SQL Server patches at a minimum!



<http://www.flickr.com/photos/25080113@N06/4758038882/>

© 2013 House of Brick Technologies, LLC




VMware Template

- ▣ Install your usual suite of programs and tools before you template
- ▣ Once done with your master server build, convert the VM to a VMware Template
- ▣ Patch periodically
- ▣ Use Windows Guest Customization Specification to:
 - ▣ Set machine name properly
 - ▣ Sysprep and set Windows product key
 - ▣ Configure network settings
 - ▣ Join to domain
 - ▣ Execute command(s) and/or script(s) post deployment
 - ▣ <http://sqlspade.codeplex.com>
- ▣ When you deploy SQL Server for production purposes, use 'Thick Provisioned Eager Zeroed' VMDK format

```
exec sp_dropserver 'OldserverName'
go

exec sp_addserver 'NewServerName',
'LOCAL'
go
```

© 2013 House of Brick Technologies, LLC



Monitoring Performance

- ▣ Perfmon / IOMeter / SQLIO / DVDStore
- ▣ SQL Server health checks
 - ▣ sqlserverperformance.wordpress.com
 - ▣ brentozar.com/blitz
- ▣ Benchmark and compare to baselines (physical and virtual)
- ▣ Remember to update your baselines when the configuration changes!
- ▣ No double standards, just more information
 - ▣ vCenter Statistics access - <http://bit.ly/SzGAiD>

© 2013 House of Brick Technologies, LLC



Conclusions

- ▣ Anyone *can* install SQL Server (click Next, Next, Finish).
- ▣ Not everyone can *squeeze* the most out of SQL Server, but now *you* can.
- ▣ You also now have the means to *demonstrate* the performance improvements.
- ▣ You can also use these to *objectively* demonstrate the performance differences between physical and virtual.
- ▣ Use these tips to ensure that your SQL Servers run as fast as possible, and that you can prove it whenever required.

© 2013 House of Brick Technologies, LLC



Questions



© 2013 House of Brick Technologies, LLC



HoB Contacts

- ▣ **Jim Ogborn (Northwest & Mid-South)**
VP, Client Solutions
jogborn@houseofbrick.com
402-445-0764 x104
- ▣ **Bob Lindquist (West & Southwest)**
VP, Client Solutions
blindquist@houseofbrick.com
480-414-9134
- ▣ **Jon Shields (Central)**
Director, Account Solutions
jshields@houseofbrick.com
847-507-1693
- ▣ **David Klee**
Principal Architect
Twitter @kleegeek
dklee@houseofbrick.com
402-445-0764 x133
- ▣ **KC Alvano (East)**
Director, Account Solutions
kcalvano@houseofbrick.com
703-395-3138

RATE ME!
<http://speakerrate.com/kleegeek>

© 2013 House of Brick Technologies, LLC