

Database Health and Performance (AKA I can *prove* it's not my system's fault!)

SQL Saturday Detroit – March 16, 2013

David Klee – Solutions 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 Me



vmware
 CERTIFIED
 ADVANCED
 PROFESSIONAL 5
 DATACENTER
 DESIGN

Microsoft
 CERTIFIED
 IT Professional

David Klee

Twitter: @kleegeek
 Blog: bit.ly/HoBSA
 davidklee.net
 G+: gplus.to/kleegeek

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

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


© 2013 House of Brick Technologies, LLC



Average Day

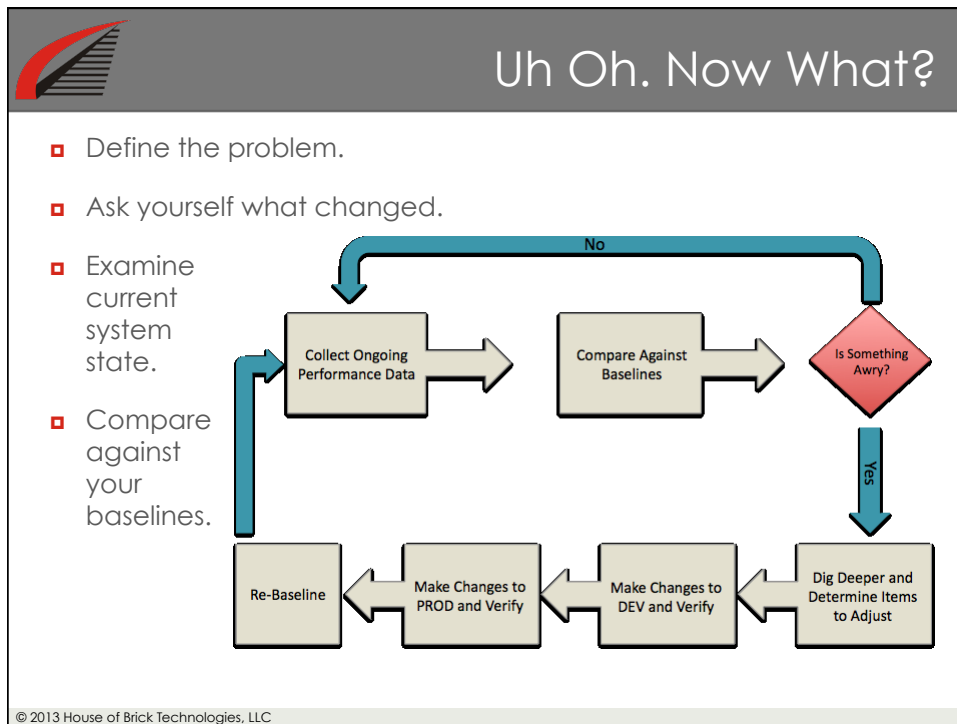
- ▣ 8:15 AM: "My app is running slow! What did you do to it?!"


(App owner and what he wants to do to your systems)



How can you prove that your systems are running optimally?


© 2013 House of Brick Technologies, LLC



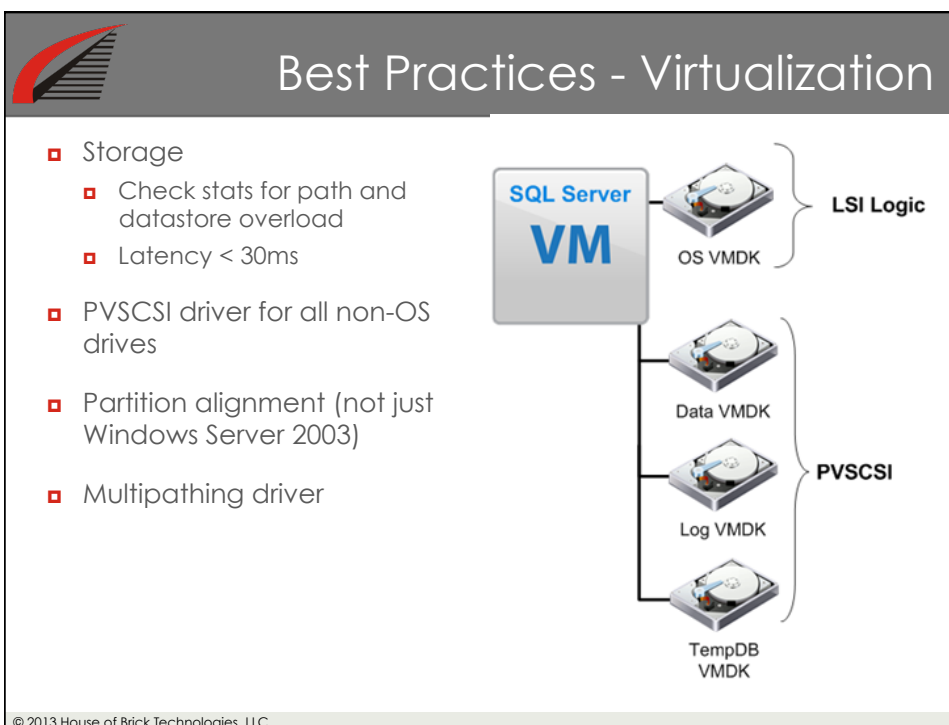
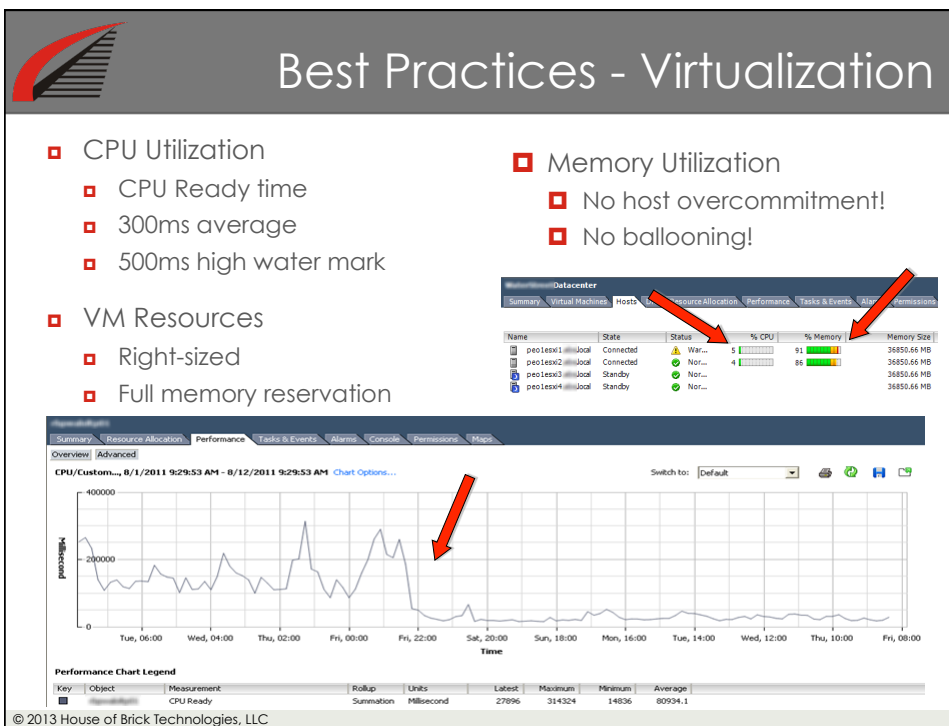



Best Practices - Infrastructure

- Hardware
 - Set power management to Maximum Performance
 - Up-to-date BIOS, hardware drivers
- Storage
 - IOMeter – At least 60MB/s in all tests
 - SQLIO – latency no greater than 25ms
- Interconnects
 - Fastest storage fiber and Ethernet you can get
 - Multiple paths a requirement
 - iperf to determine if you have problems




© 2013 House of Brick Technologies, LLC






Best Practices - OS

- ▣ Disks
 - ▣ NTFS 64KB allocation size
 - ▣ Check partition alignment (Windows 2003 especially)
 - ▣ Disable generation of 8.3 names (fsutil behavior set disable8dot3 1)
 - ▣ Disable last file access time tracking (fsutil behavior disableaccess 1)
- ▣ System
 - ▣ Antivirus exclusions for MDF, NDF, LDF, BAK set
(<http://support.microsoft.com/kb/309422>)
- ▣ SQL Server Services
 - ▣ Enable Lock Pages in Memory
 - ▣ Instant File Initialization
- ▣ Perfmon
 - ▣ Set to always collect perf counters every 5m
 - ▣ Rotate log files nightly




Windows Server® 2008 R2




Windows Server® 2012

© 2013 House of Brick Technologies, LLC




Best Practices – SQL Server

- ▣ Set min and max instance memory settings
- ▣ Optimize for Ad-hoc Workloads
- ▣ TempDB data files
 - ▣ More than one (# cores?)
 - ▣ Grow at same rate
- ▣ Watch for high log file VLF counts
- ▣ Agent system alerts – severities 17-25, 823-825
 - ▣ Default Operator set to distribution group



Microsoft®
SQL Server® 2008 R2



Microsoft®
SQL Server® 2012

© 2013 House of Brick Technologies, LLC



Best Practices – SQL Server


- ▣ Proper Maintenance is a Must!
- ▣ Fantastic database maintenance solution – ola.hallengren.com
 - ▣ Backups
 - ▣ Indexes / Statistics
 - ▣ Integrity Checks
 - ▣ Work file cleanup
- ▣ Configure email notifications and set default operator
- ▣ Schedule integrity checks!
- ▣ Demo (1)

© 2013 House of Brick Technologies, LLC

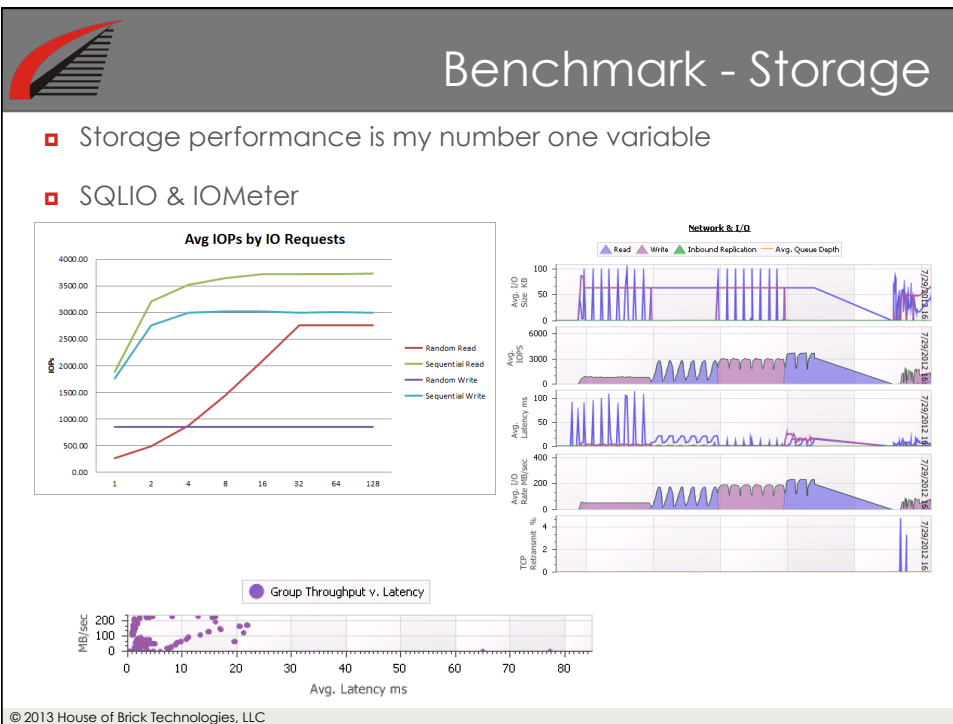


Getting Started - Benchmarking

- ▣ Must know how to benchmark so you can establish baselines
- ▣ Repeatable process to get point in time performance metrics
- ▣ Benchmarks affect the speed of the system during the test!
- ▣ What changes between tests / iterations?
- ▣ What to benchmark?
 - ▣ Subsystem speed
 - ▣ Objective SQL Server instance speed
 - ▣ Known process / job performance and runtimes
 - ▣ Query runtimes / impact
 - ▣ Application performance



© 2013 House of Brick Technologies, LLC




Benchmark – SQL Server

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

Threads	Core x GHz	VM		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
 - ▣ Demo (2)

© 2013 House of Brick Technologies, LLC



Baselines

- ▣ Creating a baseline != Tuning Time
- ▣ 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



Running Baselines - Perfmon

- ▣ Create your own constant, running system baseline
- ▣ Record every 5m, cycle log files nightly

Counters:

Memory – Pages/sec
 Network Interface – Bytes total/sec
 Physical Disk – Disk Transfers/sec
 Processor – % Processor Time
 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
 SQLServer:Locks – Lock Timeouts/sec
 SQLServer:Locks – Number of Deadlocks/sec
 SQLServer:Memory Manager – Memory Grants Pending

© 2013 House of Brick Technologies, LLC




Getting Started with Benchmarks & Baselines

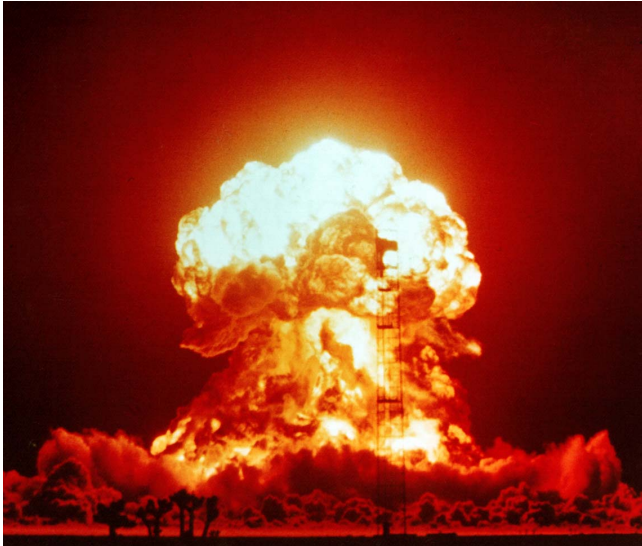
- ▣ **Easiest Places to Start:**
- ▣ Glenn Berry – SQL Server Diagnostics Queries
 - ▣ <http://sqlserverperformance.wordpress.com>
- ▣ Brent Ozar – SQL Blitz
 - ▣ <http://www.brentozar.com/blitz>
- ▣ You know your systems. Baseline your system performance frequently.
- ▣ What objective application metrics can you repetitively measure?




© 2013 House of Brick Technologies, LLC



Back to the Problem At Hand...




© 2013 House of Brick Technologies, LLC




Performance Triage

- ▣ Performance triage starts with comparing real-time stats against your baselines.
- ▣ Work Down the Stack
or
- ▣ Work Up the Stack
 - ▣ Hardware, Storage, and Virtualization
 - ▣ Operating System
 - ▣ SQL Server instance
 - ▣ Database(s)
 - ▣ Queries




© 2013 House of Brick Technologies, LLC




System Stack

- ▣ Raw Server Health
 - ▣ Hardware alarms / warnings / failures
 - ▣ Interconnect failures / dead paths
 - ▣ Networking overloaded or degraded network path
 - ▣ Sub-par raw storage performance
 - ▣ IOmeter & SQLIO
- ▣ Virtualization Health
 - ▣ Host CPU and memory states
 - ▣ High CPU Ready and/or memory ballooning
 - ▣ Host resource overcommittment
- ▣ Operating System
 - ▣ Check Windows Event Log



© 2013 House of Brick Technologies, LLC



Database Integrity

- ▣ You will experience database corruption at some point in your career. Guaranteed.
- ▣ Check frequency?
- ▣ Set DB Page Verify to CHECKSUM
- ▣ How to respond?
- ▣ MCM Readiness videos @ <http://bit.ly/ZAp5>
- ▣ `select * from msdb..suspect_pages`

```

use master;
go
dbcc checkdb(demofatalcorruption1) with all_errormsgs, no_infomsgs
go
  
```

100 %

Messages

Msg 8928, Level 16, State 6, Line 1
Object ID 0, index ID -1, partition ID 0, alloc unit ID 0 (type Unknown): Page (1:17) could not be processed.
CHECKDB found 1 allocation errors and 0 consistency errors not associated with any single object.
Msg 8906, Level 16, State 1, Line 1
Page (1:16) in database ID 13 is allocated in the SGAM (1:13) and PFS (1:1), but was not allocated in any IAM. P
Msg 2575, Level 16, State 1, Line 1
The Index Allocation Map (IAM) page (1:17) is pointed to by the next pointer of IAM page (0:0) in object ID 7.
Msg 7965, Level 16, State 2, Line 1
Table error: Could not check object ID 7, index ID 1, partition ID 458752, alloc unit ID 458752 (type In-row da
Msg 8906, Level 16, State 1, Line 1
Page (1:17) in database ID 13 is allocated in the SGAM (1:13) and PFS (1:1), but was not allocated in any IAM. P
Msg 8939, Level 16, State 5, Line 1
Table error: Object ID 7, index ID 1, partition ID 458752, alloc unit ID 458752 (type In-row data), page (1:17)
Msg 8939, Level 16, State 6, Line 1
Table error: Object ID 7, index ID 1, partition ID 458752, alloc unit ID 458752 (type In-row data), page (1:17)
Msg 8939, Level 16, State 5, Line 1
Table error: Object ID 7, index ID 1, partition ID 458752, alloc unit ID 458752 (type In-row data), page (1:17)
Msg 8939, Level 16, State 6, Line 1
Table error: Object ID 7, index ID 1, partition ID 458752, alloc unit ID 458752 (type In-row data), page (1:17)
Msg 8906, Level 16, State 1, Line 1
Page (1:24) in database ID 13 is allocated in the SGAM (1:13) and PFS (1:1), but was not allocated in any IAM. P
Msg 8906, Level 16, State 1, Line 1
Page (1:130) in database ID 13 is allocated in the SGAM (1:13) and PFS (1:1), but was not allocated in any IAM.
CHECKDB found 7 allocation errors and 3 consistency errors in table 'sys.sysallocunits' (object ID 7).
Msg 7995, Level 16, State 1, Line 1
Database 'DemofatalCorruption1': consistency errors in system catalogs prevent further DBCC checkdb processing.
CHECKDB found 0 allocation errors and 1 consistency errors in table '(Object ID 99)' (object ID 99).
CHECKDB found 8 allocation errors and 4 consistency errors in database 'DemofatalCorruption1'.

© 2013 House of Brick Technologies, LLC



CPU Bottlenecks

- ▣ First, is SQL Server the cause of high CPU activity? Task Manager!
- ▣ Raw CPU usage by instance is very high
- ▣ Signal wait statistic is greater than 25% of total waits
- ▣ Plan re-use is less than 90%
- ▣ Parallel wait statistic CXPACKET is greater than 10% of total waits
- ▣ High work queue count for long duration
- ▣ Demo (3)

Query Source: <http://blog.sqlauthority.com/2009/08/17/sql-server-measure-cpu-pressure-cpu-business/>

© 2013 House of Brick Technologies, LLC



Memory Pressure


- ▣ Does the OS have enough memory? Is > 150MB RAM free?
- ▣ Page life expectancy is low.
 - ▣ 300 is NOT the magic number!
 - ▣ One Suggestion: [Data Cache Size (GB) / 4GB] x 300
 - ▣ See <http://www.sqlskills.com/blogs/jonathan/post/Finding-what-queries-in-the-plan-cache-use-a-specific-index.aspx> for more details
- ▣ Buffer Cache Hit ratio < 90%
- ▣ High Checkpoint pages / sec & Lazy Writes / sec
- ▣ Ring Buffers
 - ▣ http://sqlskills.com/blogs/jonathan/post/Identifying-External-Memory-Pressure-with-dm_os_ring_buffers-and-RING_BUFFER_RESOURCE_MONITOR.aspx
- ▣ Demo (4)
- ▣ Read More:
 - ▣ <http://bit.ly/YnOgIM> and <http://bit.ly/X03Ghv>

© 2013 House of Brick Technologies, LLC



Network Problems

- ▣ You have high network latency
- ▣ You have dropped packets
- ▣ You have maxed out your network bandwidth
- ▣ Wait stats to check
 - ▣ ASYNC_IO_COMPLETION
 - ▣ ASYNC_NETWORK_IO
- ▣ Task manager / Perfmon on server
- ▣ Use iperf to check network paths



© 2013 House of Brick Technologies, LLC



I/O Performance

- ▣ Exceptionally high average disk seconds per read
- ▣ High volume of disk stalls and/or high disk stall wait times
- ▣ High wait stats around:
 - ▣ ASYNCH_IO_COMPLETION
 - ▣ IO_COMPLETION
 - ▣ LOGMGR
 - ▣ PAGEIOLATCH_*
 - ▣ WRITELOG
- ▣ Disk read / write stalls high
- ▣ Your triage benchmark falls well below your established baseline
- ▣ Get more reports from your storage administrator and investigate further



© 2013 House of Brick Technologies, LLC



Indexes and Statistics

- ▣ No index? Table scans = slow performance.
- ▣ No stats? The optimizer has to guess!
- ▣ Check for:
 - ▣ Index fragmentation is high
 - ▣ Statistics are out of date
 - ▣ High number of 'bad' indexes
 - ▣ Execution plan missing index warnings / bad stats
 - ▣ Is estimated number of rows WAY off from actual?
 - ▣ (Bad) DTA query index recommendations that **someone** applied
 - ▣ Is Auto Create Statistics set to off?
- ▣ Demo (5)
- ▣ Read More:
 - ▣ <http://www.simple-talk.com/sql/database-administration/brads-sure-guide-to-indexes/>
 - ▣ <http://www.simple-talk.com/sql/performance/sql-server-statistics-problems-and-solutions/>

© 2013 House of Brick Technologies, LLC



Bad Queries

- ▣ Tools: DMVs and Profiler
- ▣ Obvious places to start:
 - ▣ High CPU utilization
 - ▣ High I/O requirements
 - ▣ High memory consumption
 - ▣ Execution counts
- ▣ Missing / bad indexes and/or statistics
- ▣ Fetching more columns / rows than needed (*select * from dbo.x*)
- ▣ Bad execution plan

Include Actual Execution Plan
Ctrl+M

 - ▣ `select plan_handle, creation_time, last_execution_time, execution_count, qt.text
FROM sys.dm_exec_query_stats qs
CROSS APPLY sys.dm_exec_sql_text (qs.sql_handle)) AS qt`
 - ▣ `DBCC FREESYSTEMCACHE(plan handle)` to clear it

© 2013 House of Brick Technologies, LLC



Locking, Blocking, and Deadlocks

- ▣ A top wait statistic is LCK_M_XX AND the average time is high
 - ▣ Sys.dm_os_wait_stats
- ▣ High number of deadlocks
 - ▣ Send to SQL Server error log with trace flags 1204 and 1222
- ▣ Block process report reports long blocks
- ▣ High average row lock or latch waits
- ▣ Long running transactions
- ▣ Index contention
 - ▣ Sys.dm_db_index_operational_stats
- ▣ Read More:
 - ▣ <http://www.sqlskills.com/BLOGS/PAUL/post/Wait-statistics-or-please-tell-me-where-it-hurts.aspx>

© 2013 House of Brick Technologies, LLC



Database Origins

- ▣ Is the database that is performing poorly a third-party database, or one where you have NO control over the design?
- ▣ If third-party, will the vendor support you changing things?
 - ▣ Doubtful.
 - ▣ Will application support be dropped if you change anything?
- ▣ If custom, can you change things?
 - ▣ Maybe. It depends.
- ▣ If third-party, is your organization willing to 'void the warranty'?
- ▣ Plan Guide substitutions?



© 2013 House of Brick Technologies, LLC



Bad Database Design

- ❑ Too many table joins for frequent queries
- ❑ Table too wide (many null columns without sparse columns)
- ❑ Not enough appropriate indexes
- ❑ Too many inappropriate ones
- ❑ Indexes too huge – check for GUIDs in primary key
- ❑ No / bad normalization
- ❑ No foreign keys / primary keys / relationships
- ❑ Inappropriate isolation level


© 2013 House of Brick Technologies, LLC




Put Up or (Don't) Shut Up

- ❑ Document your findings thoroughly and accurately
- ❑ Don't just point another finger if it is not your system's fault
 - ❑ Provide evidence of the problem
 - ❑ Identify the system stack component at fault
 - ❑ Assist with remediation as needed
 - ❑ Provide validation when resolved
- ❑ If the problem is in your system, acknowledge it and continue
 - ❑ Remediate at first possible opportunity
 - ❑ Develop long term plan for solution, not just a patch
 - ❑ Append monitoring to try to proactively alert if it happens again
- ❑ Remember
 - ❑ This is an art AND a science. Know SQL Server internals. Know your environment, apps, and how they are used.

© 2013 House of Brick Technologies, LLC



Questions



© 2013 House of Brick Technologies, LLC



HoB Contacts

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

© 2013 House of Brick Technologies, LLC