

# Virtualizing Business Critical SQL Servers

Tips, Tricks, and Other Goodies to Ensure Your Success

SQL Saturday Washington DC – December 8, 2012

**David Klee** – Solutions Architect (@kleegeek)



© 2012 House of Brick Technologies, LLC



## 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



© 2012 House of Brick Technologies, LLC



About David





**David Klee**  
Twitter: @kleegeek  
Blog: davidklee.net



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

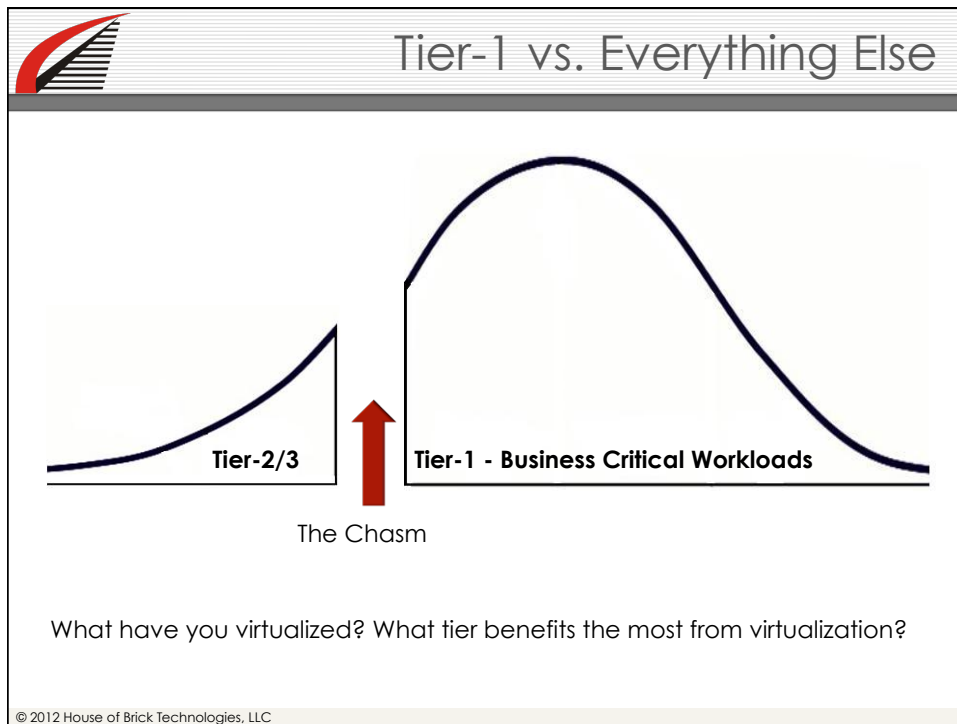
© 2012 House of Brick Technologies, LLC



Agenda

- ▣ Why Virtualize Business Critical SQL Servers?
- ▣ Physical Stack Fundamentals
- ▣ VM-Layer Fundamentals
- ▣ Designing the New Infrastructure
- ▣ Performance Baselineing, Benchmarking, and Monitoring
- ▣ Clustering and Alternatives

© 2012 House of Brick Technologies, LLC



## Myth – Cannot Run Tier-1 Virtualized

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

Production

VMware Test

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

RAID 1 + 0 Fibre Channel

RAID - 5 SATA

Applies to Apples Proofing Environment?

© 2012 House of Brick Technologies, LLC



## 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 TSANet
  - ❑ [http://vmware.com/support/policies/ms\\_support\\_statement.html](http://vmware.com/support/policies/ms_support_statement.html)
- ❑ Vendors says won't support virtualization?

**Microsoft KB 897615**

**FACT**

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

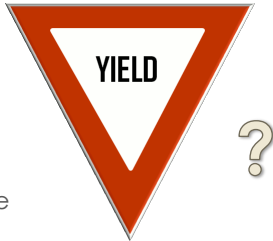
© 2012 House of Brick Technologies, LLC




## Database Size Myths

- ❑ Database size has no impact on performance. **Period.**
- ❑ Large database concerns
  - ❑ Backup/Recovery throughput
  - ❑ DR operations
  - ❑ One-time migration
- ❑ vSphere benchmarks
  - ❑ 2008 – 102K IOPs, 1.4% I/O wedge
  - ❑ 2010 – 1M IOPs, 100 microsecond wedge
  - ❑ Less than 1% CPU overhead
- ❑ No distinction between physical and virtual

**Databases in the TB's**



© 2012 House of Brick Technologies, LLC




## Licensing Concerns


Dedicated SQL Server vSphere cluster	vSphere SQL Server Sub-cluster
<ul style="list-style-type: none"> <li>Maximum consolidation</li> <li>Maximum SQL Server license optimization</li> <li>Maximum utilization of vSphere tooling</li> </ul>	<ul style="list-style-type: none"> <li>Frequently reduces SQL Server license requirements</li> <li>Consolidation ratios tend to be lower than dedicated SQL Server clusters</li> </ul>

Talk to your Microsoft licensing rep.

© 2012 House of Brick Technologies, LLC



## Virtual Storage Presentation

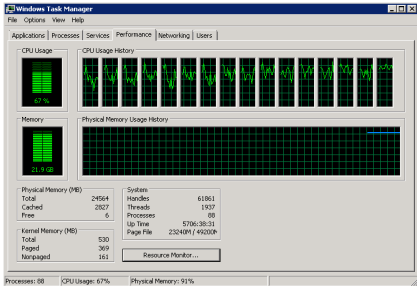
<ul style="list-style-type: none"> <li>Virtual Machine Disk (VMDK/VMFS)               <ul style="list-style-type: none"> <li>Preferred for Tier-1</li> <li>Maximum ESXi-level storage tooling</li> </ul> </li> </ul>	
<ul style="list-style-type: none"> <li>Raw Device Mapping-Virtual (RDM-V)               <ul style="list-style-type: none"> <li>More configuration and operations overhead</li> <li>Reduced ESXi-level storage tooling</li> </ul> </li> </ul>	
<ul style="list-style-type: none"> <li>RDM-P               <ul style="list-style-type: none"> <li>Maximizes SAN-level tooling transparency</li> <li>Even less ESXi-level storage tooling</li> <li>No snapshots or vMotion</li> </ul> </li> </ul>	
<ul style="list-style-type: none"> <li>Direct-mounted (In-guest iSCSI)               <ul style="list-style-type: none"> <li>No ESXi-level storage tooling</li> <li>vMotion works</li> </ul> </li> </ul>	

© 2012 House of Brick Technologies, LLC

## Tier-1 Analysis - Profiling CPU

- ▣ CPU – Count, Speed, Average, and Peak Utilization
- ▣ Windows Perfmon counters
- ▣ SQL Server health check

<http://sqlserverperformance.wordpress.com/tag/dmv-queries/>

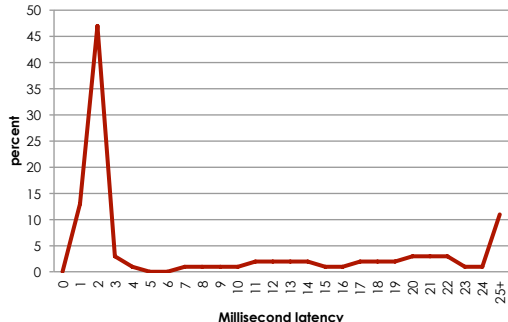


© 2012 House of Brick Technologies, LLC

## 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**



Millisecond latency	percent
0	0
1	15
2	48
3	5
4	2
5	1
6	1
7	1
8	1
9	1
10	1
11	1
12	1
13	1
14	1
15	1
16	1
17	1
18	1
19	1
20	1
21	1
22	1
23	1
24	1
25+	10

- ▣ Microsoft Download: SQLIO Disk Subsystem Benchmark Tool
- ▣ [tools.davidklee.net](http://tools.davidklee.net) – SQLIO Analyzer

© 2012 House of Brick Technologies, LLC




## Tier-1 Analysis – Storage

- IOMeter – I/O Performance Stressor
- [www.iometer.org](http://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


© 2012 House of Brick Technologies, LLC



## 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

© 2012 House of Brick Technologies, LLC



## Now What?

- ▣ Design your virtual environment carefully.
- ▣ Scale Up? More vHardware on less VMs.
- ▣ Scale Out?
  - ▣ More VMs?
  - ▣ More instances on less VMs?
- ▣ Plan for flexibility and agility.

© 2012 House of Brick Technologies, LLC

```

A problem has been detected and windows has been shut down to prevent damage
to your computer.

If this is the first time you've seen this stop error screen,
restart your computer. If this screen appears again, follow
these steps:

check to be sure you have adequate disk space. If a driver is
identified in the Stop message, disable the driver or check
with the manufacturer for driver updates. Try changing video
adapters.

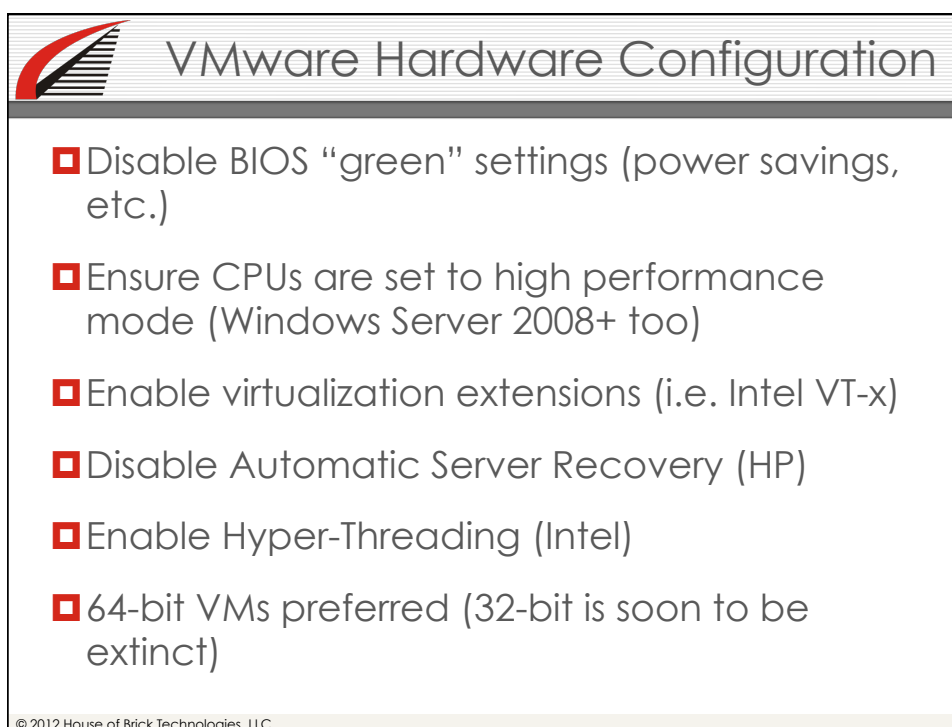
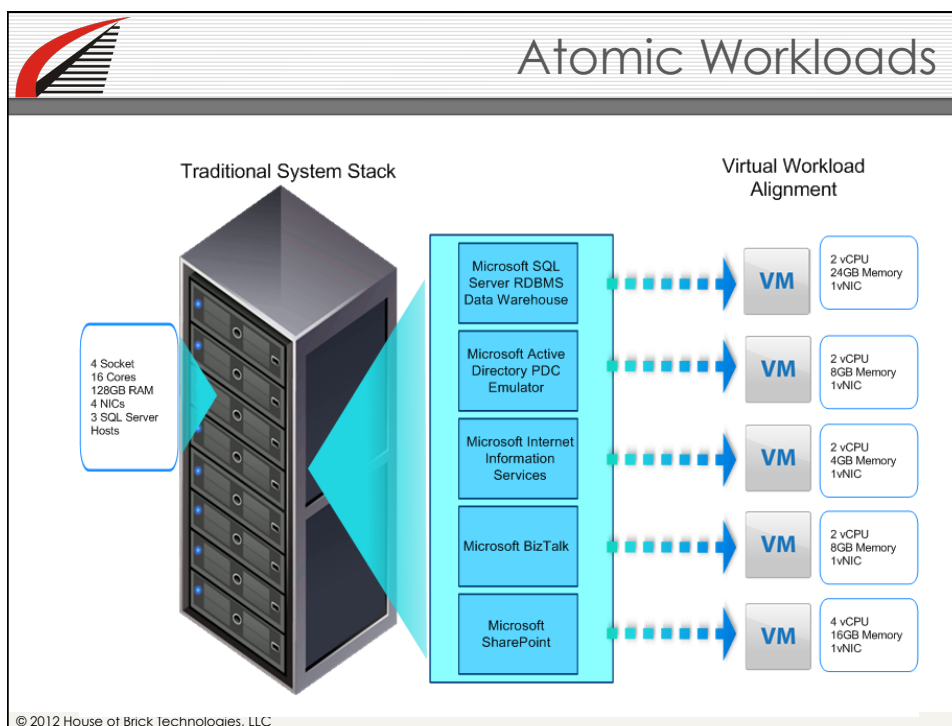
Check with your hardware vendor for any BIOS updates. Disable
BIOS memory options such as caching or shadowing. If you need
to use Safe Mode to remove or disable components, restart your
computer, press F8 to select Advanced Startup options, and then
select Safe Mode.

Technical information:

*** STOP: 0x0000008E (0xC0000420,0x82F883B2,0x82807BD4,0x00000000)

Collecting data for crash dump ...
Initializing disk for crash dump ...
  
```






## 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.



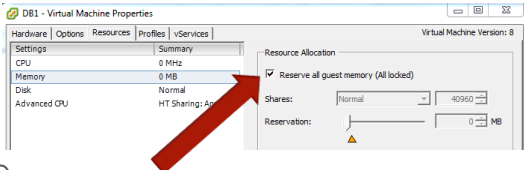
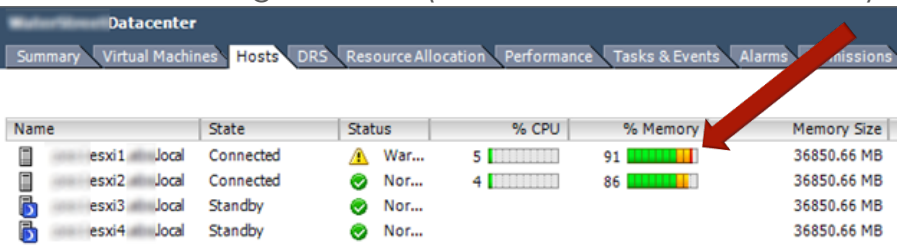
Performance Chart Legend

Key	Object	Measurement	Rollup	Units	Latest	Maximum	Minimum	Average
■	esxi1	CPU Ready	Summation	Millisecond	27896	314324	14836	80934.1

© 2012 House of Brick Technologies, LLC


## 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	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


© 2012 House of Brick Technologies, LLC



## Tier-1 - vStorage

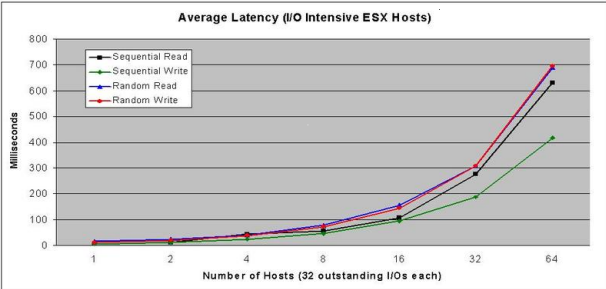
- ▣ Performance is top priority
  - ▣ RAID-10 or RAID-5? Cache?
  - ▣ *I'll repeat. Performance is **TOP** priority.*
- ▣ **Independence Rule.** Design for absolute workload object independence
  - ▣ *Monitor usage and split workloads if red lines are exceeded*
  - ▣ Can go down to one disk group per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

© 2012 House of Brick Technologies, LLC





## 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)



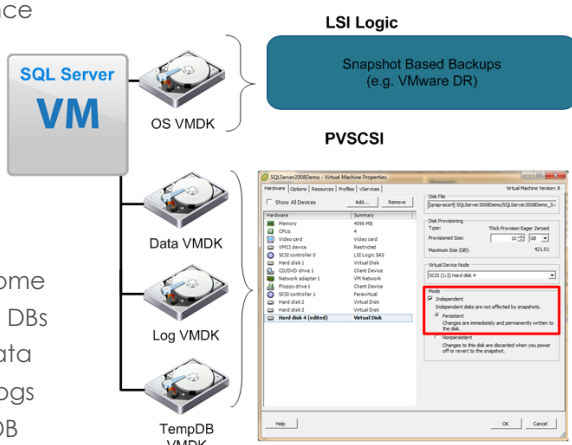
Number of Hosts (32 outstanding I/Os each)	Sequential Read (ms)	Sequential Write (ms)	Random Read (ms)	Random Write (ms)
1	~10	~10	~10	~10
2	~20	~20	~20	~20
4	~40	~40	~40	~40
8	~80	~80	~80	~80
16	~150	~150	~150	~150
32	~300	~200	~300	~300
64	~650	~400	~650	~650

© 2012 House of Brick Technologies, LLC



## Installing a SQL Server Instance

- ❑ Extending the atomic model...
- ❑ Object separation can optimize:
  - ❑ Performance
  - ❑ Disaster recovery
  - ❑ Backup
  - ❑ Flexibility
- ❑ Drives:
  - ❑ C: - OS
  - ❑ D: - SQL Home
  - ❑ E: - System DBs
  - ❑ F: - User Data
  - ❑ G: - User Logs
  - ❑ H: - TempDB



**LSI Logic**

Snapshot Based Backups  
(e.g. VMware DR)

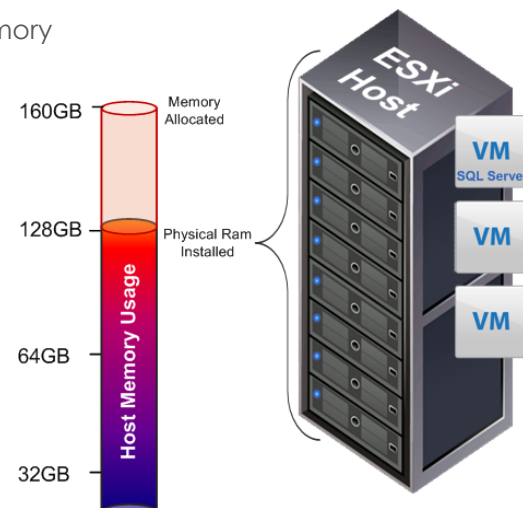
**PVSCSI**

**SQL Server Backups**


© 2012 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"
- ❑ Enable Instant File Init
- ❑ Use Large Pages – Trace Flag 834
- ❑ VM RAM Reservation
  - ❑ Memory Provisioned
  - ❑ SQL Server memory + OS + VM overhead
- ❑ Enable Optimize for Ad-hoc Workloads




© 2012 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

© 2012 House of Brick Technologies, LLC




## Failover Clustering vs. VMware HA

Four MFC 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

© 2012 House of Brick Technologies, LLC




## 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




© 2012 House of Brick Technologies, LLC




## Conclusions

- ▣ You *can* virtualize your business-critical SQL Servers
- ▣ You *should* virtualize your business-critical SQL Servers
- ▣ Ensure you baseline your physical servers, otherwise you have no means to compare performance
- ▣ Architect for atomicity, agility, performance, and scalability
- ▣ Get access to vCenter statistics and learn how to interpret them
- ▣ Follow these guidelines, and your virtualization initiative will succeed!


© 2012 House of Brick Technologies, LLC



Questions



© 2012 House of Brick Technologies, LLC



HoB Contacts

■ **Jim Ogborn (West & South)**  
VP, Client Solutions

jogborn@houseofbrick.com  
402-445-0764, ext 104

■ **Jon Shields (Central)**  
Director, Account Solutions

jshields@houseofbrick.com  
847-507-1693

■ **KC Alvano (East)**  
Director, Account Solutions

kcalvano@houseofbrick.com  
703-395-3138

■ **David Klee, Solutions Architect**

Twitter: @kleegeek  
dklee@houseofbrick.com  
402-445-0764, ext 133

© 2012 House of Brick Technologies, LLC