

SharePoint, SQL, and Storage - Make
it work right.

Eric Ku
iTrain Malaysia



**Proper SQL and Storage
design is critical to
overall SharePoint health!**

SQL Health = SharePoint Health!

**Sub-optimal SQL
performance will radiate
to other components in
the farm.**

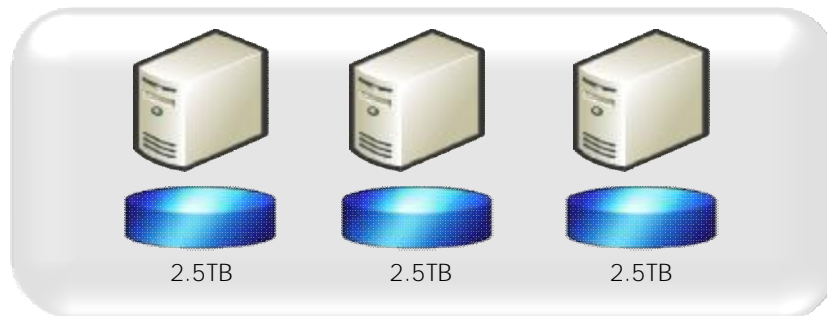


Scaling SQL Up and Out

Question #1

- How much content do you have now?
- How much is it growing per day, week, month, year?
- How will your projections change?

Scaling SQL




SCALE OUT →

Scaling SQL - Out

- More SQL servers = More flexibility
- There aren't really any physical barriers
- SharePoint will allow you to place 100 databases on 100 different SQL instances
- The real barriers are manageability and cost.
 - More servers = more money
 - More servers = more management
 - \$\$ + > management = \$\$\$\$

Scaling SQL



Scaling SQL - Up



- **Design is Paramount!**
- **Consider the following:**
 - Overall SQL Throughput (transactions/sec)
 - Disk throughput (IOPS)
 - Network throughput (MB/sec)
 - Disk backup throughput (MB/sec)
 - Network based backup throughput (MB/sec)
 - Length of maintenance windows (hours -> minutes)
 - SharePoint upgrade throughput

SQL: Scale Out VS. Scale Up

Scale Out	Scale Up
Advantages	
Better Performance	Easier to Manage
Better Flexibility	Cheaper
Disadvantages	
More Expensive	System Design is Critical
Harder to Manage	Single Point of Failure

Walkthrough: Scale Up VS. Out



How to design a 5TB
SharePoint SQL
Deployment



Consider the Organization

- Will the SharePoint SQL Servers be self managed?
- What experience does the team managing SQL have?
- Do they have:
 - Monitoring?
 - Standard Maintenance Procedures?
 - Standard Maintenance Windows?
 - Standard SQL Builds?
 - What are the break/fix and standard SLA's?

Dollars and Cents




Management Overhead

– Zen and The Art of Server Administration

MO SERVERS = MO PROBLEMS
MO PROBLEMS = MO MONEY




Scaling Direction: Effect on Replication Strategy



Scaling Direction Affect on Maintenance

0 Hours to Backup



Scaling SQL – The Bottom Line

- Don't scale SQL instances beyond comfort zones!
- Do measure system throughput – Know All of your bottlenecks!
- Scaling out is more flexible but scaling up is more cost effective.
Find a balance between scaling up and out and stick to it. (5TB per instance for example)

Storage basics

Calculating Disk IO

Spindle Count

X Spindle IOPS

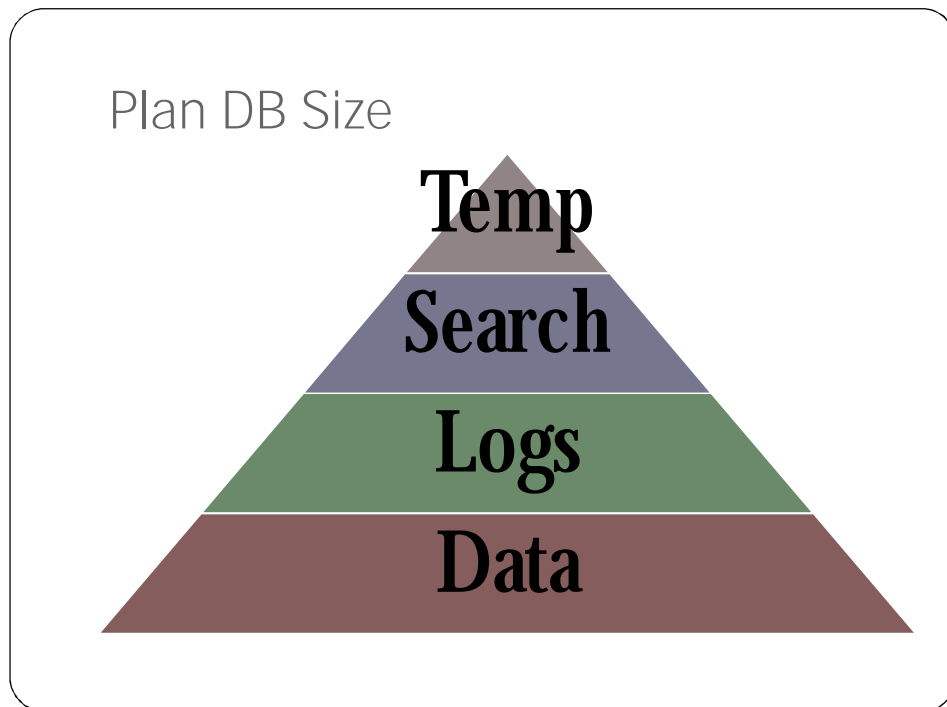
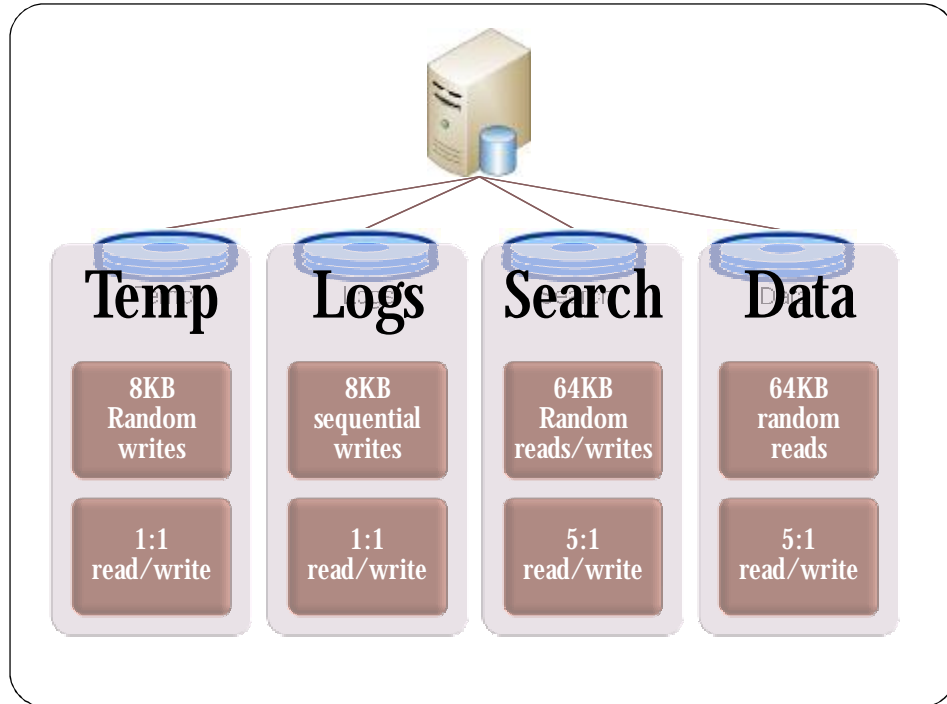
= Total IOPS capability

Calculating Disk IO – An Example

10 spindles

X 200 IOPS per spindle

= 2000 IOPS



Best Practices – SQL Disk IO

- ☐ Allocate separate and dedicated disks with the following specifications:

Workload	Size	Raid Type	IOPS
TempDB	300GB or 10%	Raid 1*	800 or 2 IOPS/GB
Logs	25% of data storage	Raid 1*	2 IOPS/GB
Search	**	Raid 1*	2000 IOPS ***
Data	110% of stored blobs	Raid 1* or 5****	.75 to 1 IOPS/GB

* Raid 1 or variants (0+1, 1+0)

** Depends on type and amount of content being indexed

*** 2000 IOPS minimum. Plan on 1500 IOPS per simultaneous crawl. (e.g. 3 crawls = 4500 IOPS)

**** Use Raid 5 when redundancy needs are met with replication

Disk IO – Understanding Raid Overhead

- ☐ **Raid 0** has no overhead but cannot tolerate any disk failure.
- ☐ **Raid 1** mirrors writes to another disk. 50% overhead (1 disk for capacity, 1 disk for redundancy)
- ☐ **Raid 0+1, 1+0** combines striping with mirroring for increased capacity, redundancy, and performance. 50% overhead (2 disks for capacity and 2 disks for redundancy)
- ☐ **Raid 5** stripes data across disks and dedicates 1 disk for parity. 3 disk set (2 disks capacity/1 for parity) 10 disk set (9 disks capacity/1 for parity)
- ☐ **Raid 6** stripes data across disks and dedicates 2 disks for parity. 4 disks set (2 disks capacity/2 for parity)

What about RAID 5?

- **Efficient use of space, but....**
- **Slow for writes**
- **Risky**
- **Recovery takes longer**

Use Raid 5 when

- **Storing content +**
- **Content is replicated +**
- **SQL failover is automatic
between content**

Random storage BP

- **Align partitions**
- **Use 64KB or larger multiple for RAID stripe size**
- **Dedicate storage for SQL**

Align partitions

What about SSD?

Good

- **SLC not MLC**

Bad

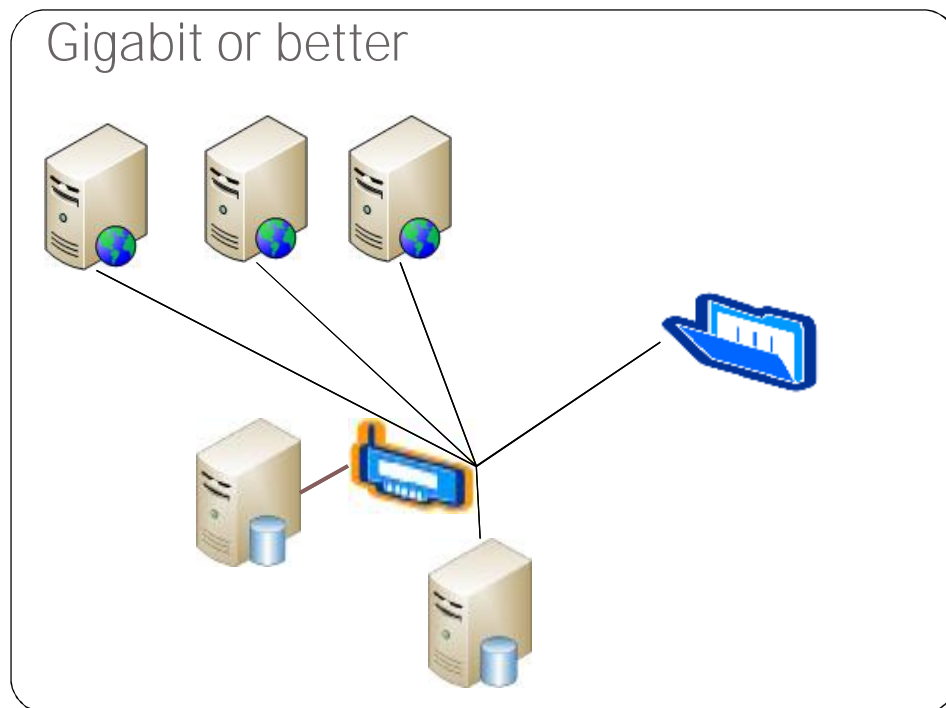
- **New, not proven.**

Ugly

- **Expensive**

DAS verses SAN – Which is Better?

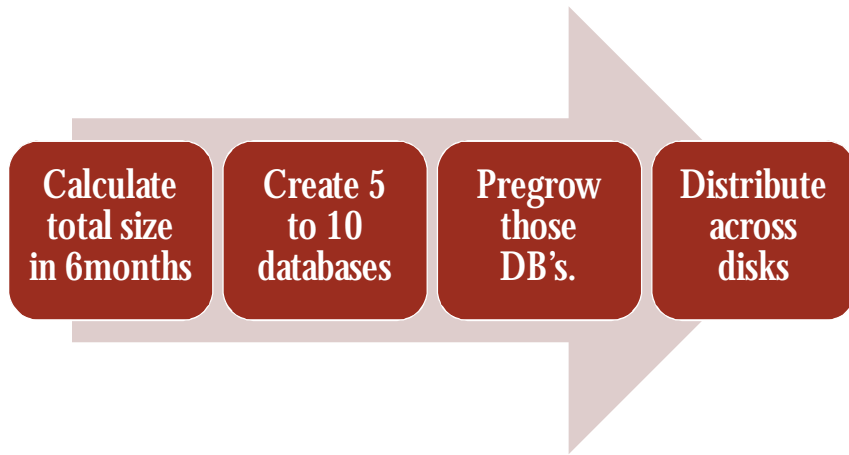
- Both type of storage can scale, perform, and serve a multi-TB farm.
- So... where the difference?
 - Ease of management.
 - Growth potential.
 - Advanced capabilities (snapshots, remote replication).
 - The ability to share with other applications.
 - Price...



Isolate workloads

- **Centralized SSP/Search/Profiles**
 - >50K user profiles
 - >10million search items
 - Consuming from multiple farms
- **Portal**
 - High performance – SLA, dedicated
- **MySites**
 - Important? – requires heavy CPU/IO

Pregrow databases and never
autogrow




Use
64bit!

Content DB Size Limitation 100GB?

- **Exceeding 100GB? Keep in mind:**
 - Backup/restore/maintenance will be harder.
 - Use differential backup.
 - All sites share the same tables. Isolate large sites.
 - Use multiple data files
 - Defrag regularly.

*Your experience may vary: H/W and usage profile dependant.

Don't exceed 5 web/app servers per
SQL instance



Virtualize SQL?

Best case	OK	Worst case
<ul style="list-style-type: none">• Don't	<ul style="list-style-type: none">• Dedicate same HW• Careful design• Scale out	<ul style="list-style-type: none">• Inadequate host• Vhd• Single proc• <8GB memory

Monitoring and Troubleshooting Perfmon

- **Physical Server**
 - Keep CPU on the low-end of 50%-75% in off-peak times.
 - Monitor Processor Queue length for CPU bottlenecks. (no more than 10 per proc)
 - Always keep at least 20% of the RAM available.
 - Use the SQL Server SP2 Performance Reports!

Summary

- **SQL is extremely important to SharePoint health and Performance**
- **Put SQL on 64bit.**
- **Think IOPS when designing disk arrays.**
- **Always separate work loads with the following priority: temp, log, search, content.**
- **SQL scales up and out. Don't push the limits upward, but keep manageability and costs in mind when scaling out.**
- **Designing enterprise services with great care. Separate SSP and Search when possible.**

