

EITPA's Professionals Meeting
Best Practices in Software Architecture

Selecting the Best Persistent Storage Mechanism

Michael Abbey

*Enterprise Best Practices in Software
Architecture*

November 28, 2009


Questions

- Cost
- Scalability
- SQL compliance
- The user community
- Common myths
- Recommendations

Selecting the Best Persistent Storage Mechanism

Cost

- If you don't have XYZ dollars, then you have no choice
- An easy trap to fall into → less expensive does not necessarily mean "better"
- Savings today can translate into future costs



Selecting the Best Persistent Storage Mechanism

Free offerings

MySQL

- ~~database size limit~~
- ~~memory limit~~
- ~~CPU limit~~

Oracle

- database size limit = 4Gb
- memory limit = 1 Gb
- CPU limit = 1

SQL Server

- database size limit = 4Gb
- memory limit = 1 Gb
- CPU limit = 1

Free offerings

DB2 Express-C

- ~~• database size limit~~
- memory limit = 2 Gb
- CPU limit = 1

No SQL

- ~~• database size limit~~
- ~~• memory limit~~
- ~~• CPU limit~~

**Are the above
comparisons enough?**



wise



foolish

**Saving ETB 1,000,000
today may lead to**

Scalability

- Can the software handle an operating system change
- Can the software deal with increases in data volume
- Are there features (with a cost) of the software that may assist business apps as they grow
- Can data be partitioned as volume increases

Scalability

- Can the vendor handle distributed processing transparently
- Can nodes of a cluster interact with shared storage ... 1 database / multiple clustered servers
- Can the RDBMS leverage the addition of additional computing resources

SQL compliance

- All SQL database vendors are close to 100% compliant
- De facto standard now is SQL:2008
- Sixth revision of the SQL query language standard
- <http://en.wikipedia.org/wiki/SQL>
- Compliance makes code portable
- Is it still the standard

SQL compliance

- Do the vendors' extensions render the code *not-portable*
 - is information readily available from the vendor as a guide
 - what mechanism is delivered to test compliance
- *SQL Server* → set fipsflagger on
- *Oracle* → alter session set flagger = TRUE;
- *MySQL* → set fipsflagger on

The user community

- Is it static or dynamic; if dynamic, growth projections
- Local or remote
- Distributed updates or just reporting
- Do they require 24x7
- Concurrent users
- When may the increasing user base require hardware upgrade

Common myths

- We don't need support, it's too expensive
- Release we use is just fine; why bother keeping current with vendor's release
- Database management in a virtual machine is easier
- Virtualization will automatically solve all your performance issues
- A very large data cache will assist throughput

Common myths

- Parallelizing operations will speed up processing
- Index access is always faster and more efficient than full table scans
- Objects occupying a large number of extents will hurt performance

Recommendations

- Do not base your purchase on cost and cost alone
- Free is not always better
- Try to find out users' impressions of
 - The vendor's electronic support
 - Whether the vendor even has support
- Install base and market share should not drive decision

Q&A

Selecting the Best Persistent Storage Mechanism