WHITE PAPER

# Linear VoltDB Scalability on a 30-node SGI® Rackable™ C1001-TY3 Cluster For Cloud Applications

Scaling database applications to support heavier, "web-scale" transaction (OLTP) workloads in the cloud, can be painful. To adapt, applications are typically migrated to expensive hardware and DBMS software and require more extensive database tuning as time goes on. Or, to avoid these hard costs, some organizations will incur softer costs like the overhead of developing and maintaining complex caching or partitioning (sharding) schemes, or sacrificing transactional consistency for performance by using a NoSQL key-value store.

## VoltDB – Fast Performance 24 x7 and Painless Scalability

Created by DBMS R&D pioneer, Mike Stonebraker, VoltDB solves the OLTP performance and scalability dilemma. Not only has it been shown to run 45x faster than a current price/performance benchmark leader, it is the first OLTP SQL database designed to scale painlessly on clusters of inexpensive servers while providing transactional consistency (ACID) and automatic high availability and disaster recovery.

With VoltDB you can develop and deploy applications at low cost and then scale them to handle millions of transactions per minute, and without the costs of rip and replace DBMS or hardware upgrades, extensive tuning, or application redesigns. Just add additional servers to the VoltDB cluster.

To prove the point, a benchmark with a cloud application was run on a 30-node SGI® Rackable™ C1001-TY3 cluster showing how VoltDB performance scales linearly as servers are added to the cluster.

## The Benchmark Application

The benchmark "Voter" application that was used features high-frequency transactions simulating a phone-based voting contest (such as a televised pop idol singing competition). This application is representative of the massive workloads supported by social networking, ad-serving, gaming, financial trading, telecommunications, on-line businesses and other cloud applications with large, fast-growing transaction volumes. Voters, identified by their phone numbers, are allowed a set number of votes, and may attempt to vote for a non-existent contestant. A "vote" attempt is returned as "vote accepted;" "vote rejected, contestant does not exist;" or "vote rejected, voter over limit."

Voter Application Requirements:

- Must support processing > 100,000 votes per second
- Each contest contains between 2 and 12 contestants
- Voters are allowed a maximum of 10 votes
- A vote consists of a phone number and contestant number
- Successful votes must be kept for audit
- Application must provide real-time vote tallies per contestant at least once per second

Voter Application:

- Vote table (contestant_number, phone_number) partitioned on phone_number
- Contestants table (contestant_number, contestant_name) replicated to all partitions
- Transaction "vote" – single-partition, all votes for a caller's phone number are within the partition
- Transaction "results" – multi-partition by design
- Virtually unlimited number of partitions, popular contestants are spread over all partitions (CPU/RAM)

## The Benchmark Hardware

The benchmark configuration comprised a SGI® Rackable™ C1001-TY3 cluster with 30 server nodes and 6 client nodes, each with 2x Intel® Xeon® 5590 processors, 24GB@1333MHz memory and running Red Hat Enterprise Linux Server release 5.4 and SGI Management Center 1.2. Nodes were connected by an SMC Networks SMC8848M switch. An additional six of these servers were used as clients, each with two cables 802.3ad bonded to the switch. The benchmark application software was written in Java (Oracle Java HotSpot 1.6) accessing the VoltDB 1.1 database.

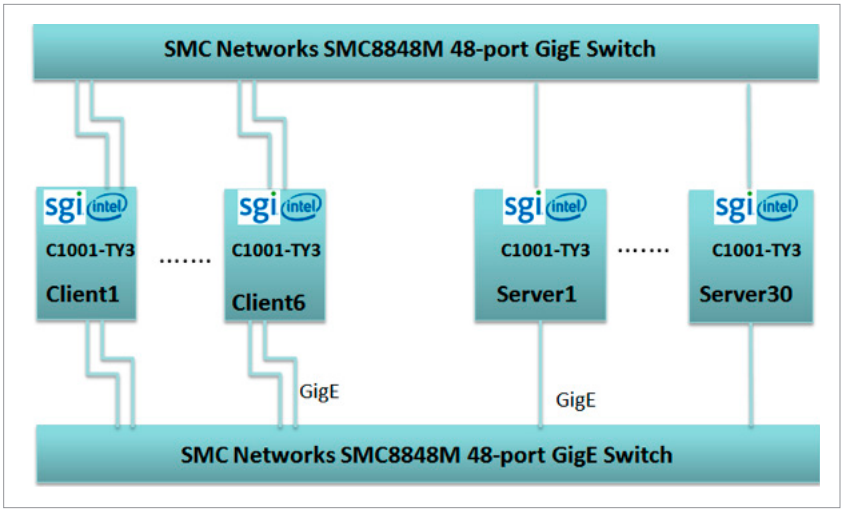


## The Benchmark Results

The benchmark application was run several times, each time with 5 more servers than the last run. The table below shows the results.

| Servers | Clients | Server Rate (TPS) |
|---|---|---|
| 5 | 2 | 640,000 |
| 10 | 3 | 1,200,000 |
| 15 | 4 | 1,740.000 |
| 20 | 5 | 2,340,000 |
| 25 | 6 | 2,860,000 |
| 30 | 6 | 3,370,000 |

Results Summary

- Single-node throughput: 120K transactions/sec
- 30-node cluster throughput: 3.4M transactions/sec
- Scaling: linear, scale factor of 0.86/node
- Real-time vote tallies up to 10/sec have no impact on the throughput

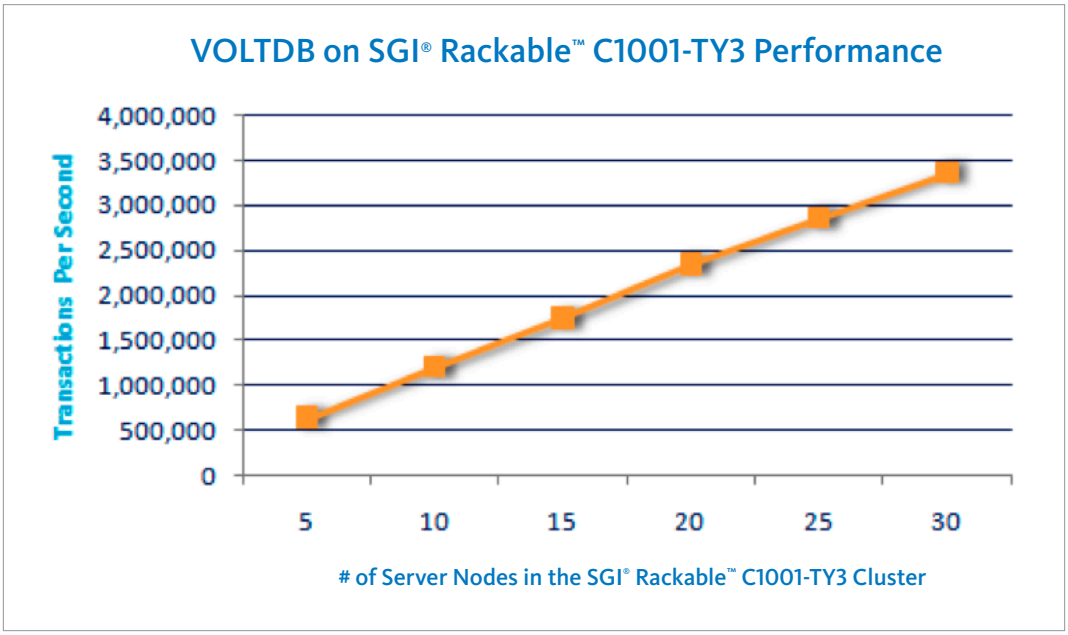


## What Makes VoltDB and SGI So Scalable

Most OLTP databases are based on decades-old designs conceived long before the Cloud era. The problem is, those databases waste lots of CPU cycles (as shown in the chart to the right) preparing to execute the transaction (to guarantee transactional consistency) instead of executing the transaction.

While necessary for data integrity, this overhead prevents those databases from running faster and scaling out on shared-nothing clusters of commodity hardware to support today's web scale workloads in the Cloud.

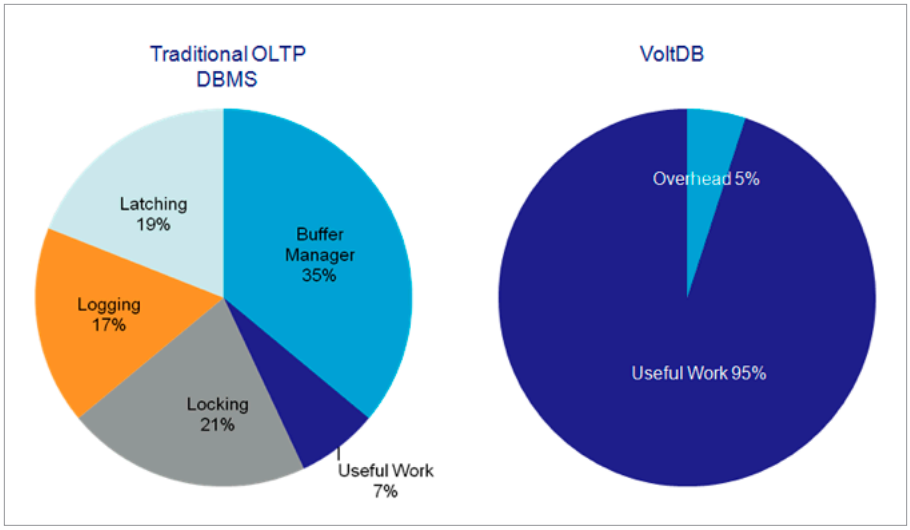

The innovative architecture of the VoltDB DBMS delivers breakthrough performance and near-linear scalability on the SGI platform. VoltDB automatically "shards" the database into data partitions and distributes each partition—along with a SQL engine to process it—to a different CPU core in the cluster. Each high-frequency Intel Nehalem-EP CPU core with 8MB L3 cache executes its VoltDB database transaction queue at in-memory speeds. VoltDB automatically maintains data consistency across all nodes in the cluster, and also replicates data throughout the SGI® Rackable™ C1001-TY3 cluster to provide 24x7x365 high availability year-round.

The benchmark with VoltDB proves that the energy-efficient SGI® Rackable™ C1001-TY3 cluster platform yields excellent performance at a breakthrough price. With an aggregate throughput of 3,370,000 transactions per second (TPS) across 30 server nodes, the cluster yields a Price/TPS of only 0.08 USD.†

## Try the Free Open-Source VoltDB Community Edition on SGI® Rackable™ servers

If you're curious to learn more about VoltDB, please visit community.voltdb.com for access to VoltDB product documentation, free community version downloads and developer support forums. The VoltDB download kit includes the Voter application source code used in this benchmark. To test VoltDB on SGI® Rackable™ servers, or to get more information about SGI® Rackable™ servers, please call SGI for expert consulting.

† *Pricing is available upon request*

**Corporate Office**
46600 Landing Parkway
Fremont, CA 94538
tel 510.933.8300
fax 408.321.0293
www.sgi.com

North America +1 800.800.7441
Latin America +55 11.5185.2860
Europe +44 118.912.7500
Asia Pacific +61 2.9448.1463