ScaleDB
Scaling MySQL to New Heights

ScaleDB for MySQL

ScaleDB for MySQL is a plug-in storage engine for MySQL that delivers enterprise-class shared-disk clustering. With shared-disk clustering, each node in the cluster has full read/write ability across all of the data for the cluster. Because each node has full access to the entirety of the data, any node can address any database request. This enables unprecedented flexibility, scalability and availability while retaining a low total cost of ownership.

ScaleDB provides all of these enterprise capabilities under the covers of MySQL without changing your application. This future-proofs your application, eliminating the need to port to expensive commercial database management systems as your database requirements grow. In addition, you can add and remove cluster nodes without impacting operation of the database. This capability enables the elastic scalability that is a core principle of cloud computing or Database as a Service (DaaS).

ScaleDB Delivers High-Availability

ScaleDB employs a shared-disk architecture that delivers superior high-availability. While shared-nothing requires each node to store and manipulate a portion of the data—creating data silos and single points of failure—ScaleDB's shared-disk architecture enables any node to satisfy any database request. As a result, failure of a single node simply means that the other nodes continue sharing the database load until the failed node is revived. Until now, high-availability meant high-price, ScaleDB changes that.

ScaleDB Delivers Flexibility

ScaleDB easily accommodates shifting or evolving usage patterns. Shared-nothing databases must be tuned for a very specific usage pattern. However, OLTP applications, by their very nature involve variable usage patterns. Database workloads may vary by time of day (e.g. morning logins), day of the week (e.g. end-of-week reporting), day of the month (e.g. payroll days) and time of the year (e.g. tax time). Usage patterns can also evolve over time. While shared-nothing databases are partitioned and tuned for one specific usage pattern, shared-disk databases are far more flexible. This flexibility comes from the fact that any node in a shared-disk cluster can satisfy any database request. For example, in a 3-node cluster, you might find 2 or even 3 of the nodes handling logins in the morning, reporting at the end of the week, payroll on the 15th and taxes at the end of the year. This ability to share the database load at the cluster-level—known as cluster-level load balancing—delivers the dynamic flexibility you need in a rapidly changing business environment.

ScaleDB Delivers High-Performance

ScaleDB utilizes a unique and patented indexing mechanism that delivers superior performance. Overall database performance is addressed through a combination of means, but it all starts with the efficiency of the index. ScaleDB's Multi-Table Index

(MTI) is highly compressed, typically only 25% of the size of a B-Tree index. Smaller means faster. The MTI is also designed to minimize disk access. By reducing its reliance on the disk, the slowest part of the database, performance is improved in most scenarios. The MTI also inherently maintains referential integrity, ensuring database integrity without the performance penalty inherent in B-Tree. Finally, the MTI provides lighting fast table joins, similar to materialized views, but without the data synchronization headache. These and other advantages combine to provide unparalleled performance.

ScaleDB Delivers Seamless Scalability

Databases need the ability to scale to address growing, or shrinking, workloads, users and data sizes. These may result in the need to add or remove servers. In a shared-nothing database, this means an extensive process that involves shutting down the database and repartitioning it to accommodate a change in the number of servers. ScaleDB's shared-disk architecture allows you to simply add or remove servers on-the-fly, without any shut-down. You can even change your database schema while the database is running, without interruption. This ability to scale dynamically and without interruption is very valuable to mission-critical applications, and it is an absolute requirement for cloud computing.

ScaleDB Lowers Your Total Cost of Ownership (TCO)

By eliminating the need for slave servers, ScaleDB reduces your hardware and software costs. By leveraging open source MySQL, ScaleDB further reduces your software costs. By providing dynamic cluster-level load balancing, ScaleDB reduces the number of servers required to handle your database load, thereby further reducing your hardware costs. By eliminating the need for data partitioning, re-partitioning, data routing and also reducing your tuning efforts, ScaleDB significantly reduces your set-up and maintenance costs. This all adds up to significantly lower TCO, even when compared to completely free software.

ScaleDB for MySQL provides the following capabilities:

- Innovative high-speed indexing
- Plug-and-Cluster™ simplicity
- Automatic data recovery
- ACID compliance
- Shared-everything architecture
- High-performance transaction processing
- Graceful fault-tolerance
- Row-level locking
- Multi-node concurrency control
- Eliminates the requirement to partition data

ScaleDB
3723 Haven Avenue
Menlo Park, CA 94025
(650) 587-8787
(650) 587-1571
info@scaledb.com
www.scaledb.com