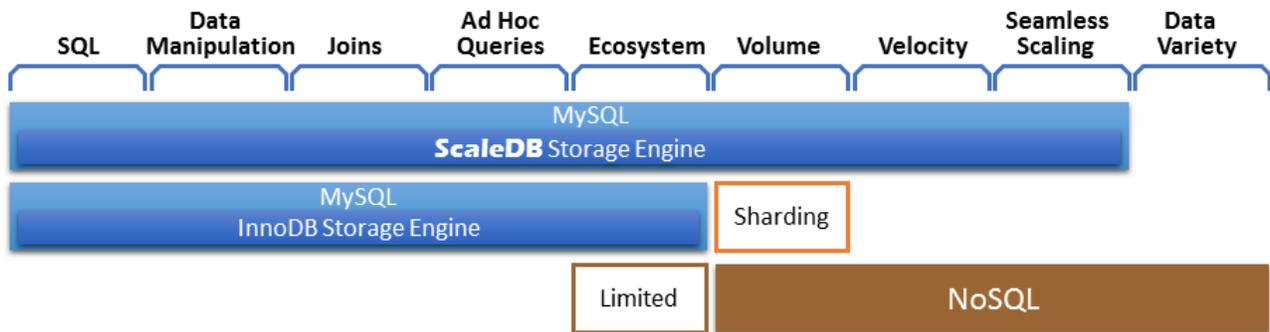


MySQL for High-Volume & High-Velocity Time-Series Data

ScaleDB extends MySQL for high-velocity high-volume data, enabling near real-time analytics across massive streams of data. Like Hadoop, ScaleDB leverages a cluster of commodity machines with low-cost spinning disks for storage, mitigating the typical trade-offs between data fidelity and the volume of data and time analyzed. NoSQL and Hadoop are excellent choices for unstructured data, or data variety challenges. However, if your data is structured time-series data, and generated at high-velocity and high-volume, ScaleDB is the tool for you. On a nine-node cluster, ScaleDB inserts millions of rows per second and queries up to 1000 times faster, more than enough speed for today's most demanding applications.

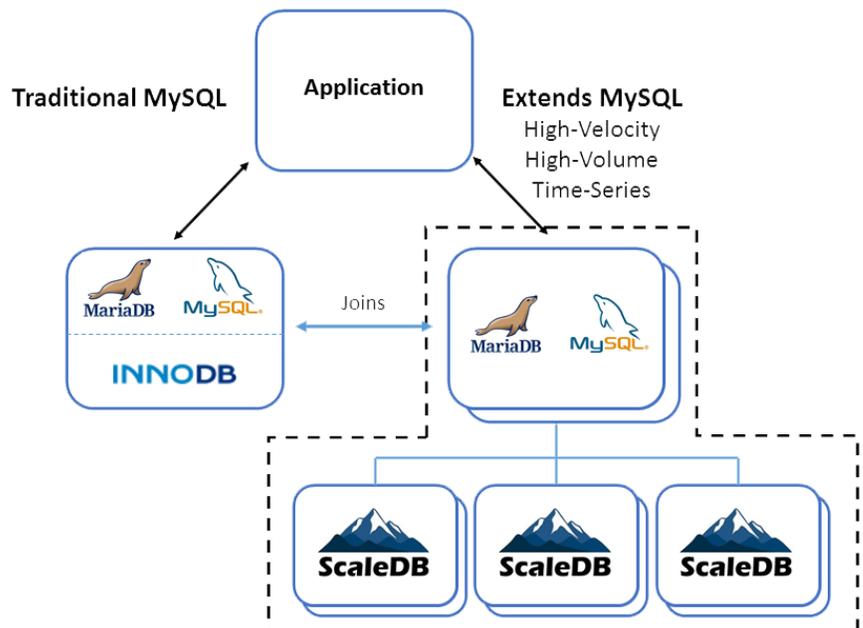
Features & Capabilities



Don't Change Your Database, Extend It

If you have an application using MySQL or MariaDB using the InnoDB storage engine, there is no need to change anything, just add ScaleDB alongside InnoDB. As the diagram below illustrates, you continue to use InnoDB, but for those tables where the data volume and velocity exceed the capabilities of InnoDB, you can use ScaleDB. Your application and tools still use MySQL or MariaDB, while ScaleDB handles the high-velocity and high-volume data.

Non-SQL solutions force trade-offs such as time-consuming batch processing, no ad-hoc queries or joins, uses difficult imperative programming languages (instead of SQL), new tools and much more. With ScaleDB, you no longer have to make these trade-offs, you get the same MySQL or MariaDB DBMS, but your ScaleDB tables deliver unprecedented performance across big/fast data.

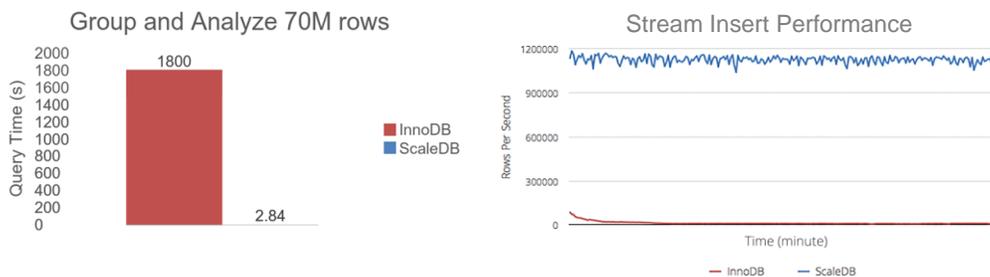


How It Works

ScaleDB plugs into the MySQL storage engine API (like InnoDB) turning MySQL into a high-performance cluster. Unlike expensive in-memory systems, ScaleDB uses commodity hardware, either on-premise or in the cloud. ScaleDB's modern clustering architecture spreads processes across the entire cluster and pushes the processing to the data, for optimal performance. All of your LAMP stack applications and tools continue to work, because you're still using MySQL. Unlike other Non-SQL tools, there is no need to build a whole new platform, with layers of nascent and unproven technologies. You continue to use the same proven and powerful MySQL platform you know and love, but it operates like it's on steroids.

How It Performs

ScaleDB's insert performance is typically 250-times that of MySQL and 3- to 5-times that of comparable NoSQL solutions. Query performance generally ranges from 20-60-times faster than InnoDB, up to 1000 times faster for queries with time predicates and 2-times that of NoSQL alternatives.



A Few Common Use-Cases

User Behavior Analysis: Collect and analyze every user action to improve sales, usability, content, security, marketing and more.

Real-Time Social Sentiment: Companies process and parse millions of tweets, likes and +1 in real-time enabling their customers to monitor and react to social sentiment.

Internet of Things (IoT): Devices and sensors create high-velocity streams of data that must be processed, acted upon and saved at speed.

Network Management: Everything from CDRs to packet monitoring.

Fraud Detection: Evaluate patterns across a massive flood of data and act on it preemptively to minimize the financial impact.

Online Gaming: Monitor player activity in real-time for support, in-game marketing, analytics and more.

Stream Processing: An ideal complement to stream processors like Storm, Kafka, Kinesis, etc. Instead of dumping data or stepping down data fidelity over time, store and process everything in ScaleDB.

NO COMPROMISES TIME-SERIES DATA

Imagine MySQL handling high-velocity, high-volume time-series data. How easy is that?

Velocity

- Millions of inserts/second
- Real-time analytics
- Scale-out writes

Volume

- Elastic scale
- Push queries to the data and process queries in parallel

Complete

- It is MySQL
 - Same BI Tools
 - Same Dev Tools
 - No New Training
- Joins
- Ad-hoc queries
- Save/query all source data at disk pricing

TCO

- LAMP
- Cloud or on-premise
- Commodity Hardware
- Disk-Based

Hardware Requirements

- Commodity servers with RAID-0
- Number depends on your volume/velocity