


## Replace your Oracle® and MySQL™ databases with a modern RDBMS

Leverage in-memory and scale-out technologies and increase performance by 10-20x at one-fourth the cost


### Standard ANSI SQL

Splice Machine is an ANSI SQL-compliant database built on Hadoop and Spark that enables companies to leverage their existing SQL-trained resources, as well as SQL-based Business Intelligence platforms, ETL pipelines and existing applications.




### Horizontal Scale Out

The Splice Machine RDBMS utilizes the proven auto-sharding capability in HBase and the distributed datasets of Spark to provide massive scalability across commodity hardware, even up to dozens of petabytes.




### Real-Time Updates with ACID Transactions

Splice Machine provides ACID (Atomicity, Consistency, Isolation, Durability) transactions across multiple rows and tables to enable concurrent, real-time updates without data loss or corruption. This is ACID for highly concurrent applications, not just big batch updates.




### Massively Parallel Architecture

Splice Machine delivers massive parallelization by pushing computation down to each local distributed data shard. Each distributed node embeds HBase, Spark and the Splice parser, planner, and optimizer.



### Flexible, Hybrid Database Platform

With in-memory technology from Spark and scale-out capabilities from Hadoop, the Splice Machine RDBMS provides exceptional performance for simultaneous OLAP and OLTP workloads.



## Splice Machine: The Advantages of SQL, the Scale Out of NoSQL, the Performance of In-Memory

Splice Machine delivers a unique hybrid of state-of-the-art technologies:

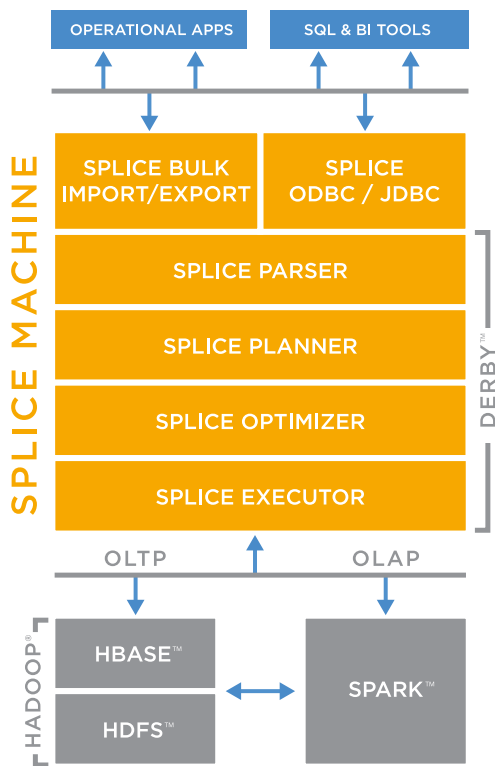
- **Transactional SQL** – support existing OLTP applications without rewrites or retraining of existing SQL-based analysts
- **Scale Out** – cost-effectively scale out on commodity hardware with proven technology from Hadoop and Spark
- **In-Memory** – experience outstanding performance for OLAP queries with in-memory technology from Spark

### Cost-Effective, Hybrid Architecture

The Splice Machine RDBMS is an innovative hybrid of in-memory computing from Spark and disk-based storage from Hadoop. Unlike in-memory-only databases, Splice Machine does not force companies to put all of their data in-memory, which can become prohibitively expensive as data volumes grow.

### Mixed OLTP/OLAP Workloads

Splice Machine started with Apache Derby, an ANSI SQL Java database, and replaced its storage layer with HBase/Hadoop and Spark. The Splice Machine optimizer automatically evaluates each query and sends it to the right computation engine: OLTP queries (i.e., small read/writes, range queries) go to HBase/Hadoop, OLAP queries (i.e., large joins or aggregations) go to Spark.



## Splice Machine Powers Big Data Solutions

Splice Machine can be used anywhere an operational RDBMS, like Oracle or MySQL, is used, but with the advantage of scaling out on commodity hardware accelerated by in-memory technology. Customers use Splice Machine to power a broad range of operational applications and real-time analytics.

### Use Cases for Splice Machine include:

- ETL Acceleration
- Digital Marketing
- Operational Data Lakes
- Data Warehouse Offloads
- Internet of Things Applications
- Web, Mobile, and Social Applications
- Operational Applications