The Open-Source SQL RDBMS Powered by Hadoop and Spark

Splice Machine started with a vision to disrupt the database market with technology to support applications that require hybrid transactional and analytical workloads. We have created a scale-out RDBMS that **is up to 10x faster than Oracle at a fraction of the cost** and that eliminates the need for ETL and expensive dedicated analytical appliances. Splice Machine incorporates the proven scalability of HBase, the standard ANSI SQL and ACID transactions of an enterprise RDBMS, and the in-memory performance of Spark. Splice Machine supports use cases such as financial services, digital marketing, life sciences and supply chain optimization, that are all characterized by extremely high data volumes, advanced analytics and "in-the-moment" decision making.

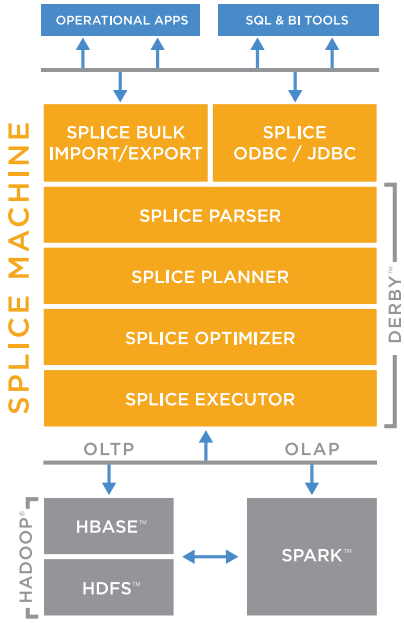

The Splice Machine RDBMS has:

- **ANSI SQL.** Splice Machine provides ANSI SQL-99 coverage, including full DDL and DML.
- **ACID Transactions.** Splice Machine provides fully ACID transactions with Snapshot Isolation semantics, which scale both to very small operational queries and large analytics.
- **In-Place Updates.** Updates in Splice Machine scale from a single row to millions of rows with a single transaction.
- **Secondary Indexing.** Splice Machine supports true secondary indices on data, in both unique and non-unique forms.
- **Referential Integrity.** Referential integrity such as Primary and Foreign key constraints can be enforced without requiring any behavior from the underlying application.
- **Joins.** Splice Machine supports inner, outer, cross and natural joins using join algorithms such broadcast, merge, merge sort, batch nested loop, and nested loop joins.
- **Resource Isolation.** The cost-based optimizer chooses a dataflow engine based on the estimation of the query plan – OLTP runs on HBase and OLAP runs on Spark.

For New Applications

Splice Machine is like a Lambda Architecture in-a-box. Developers can use standard SQL to ingest, access, update, and analyze the database without worrying about what compute engine to use because the Splice Machine optimizer picks the right integrated compute engine based on the query.

For Existing Applications

The Splice Machine RDBMS was designed to offload workloads from overwhelmed RDBMSs like Oracle, MySQL, IBM DB2 and Microsoft SQL Server that companies are finding too expensive to scale with standard SQL, while eliminating the need to rewrite existing applications.

Unprecedented Support for Simultaneous OLTP & OLAP Workloads

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

The Splice Machine RDBMS was designed to offload workloads from overwhelmed RDBMSs like Oracle, MySQL, IBM DB2 and Microsoft SQL Server that companies are finding too expensive to scale. Splice Machine provides cost-effective scale out on commodity hardware, but unlike NoSQL databases, it provides standard ANSI SQL, eliminating the need to rewrite existing applications.

Splice Machine automatically sends OLTP queries to HBase/Hadoop and OLAP queries to Spark. With separate processes for HBase and Spark, Splice Machine isolates the workloads and ensures that OLTP response times remain flat as OLAP loads increase.

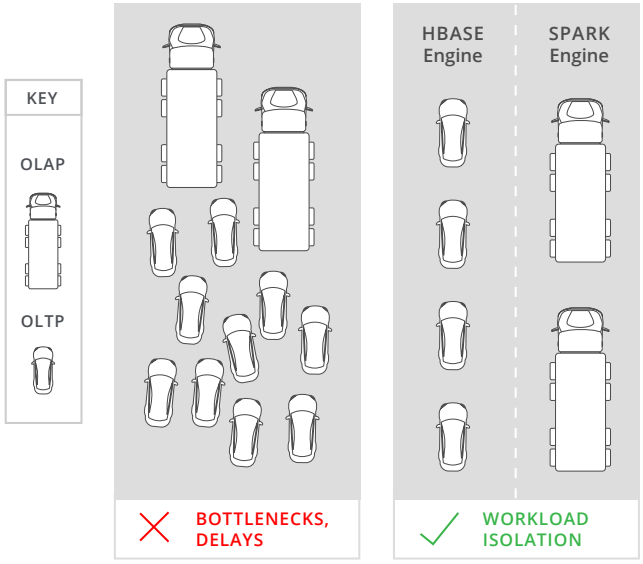

Advanced, In-Memory Technology – Spark Query Management

Splice Machine embeds Apache Spark™ – a fast, open source engine for large-scale data processing – to accelerate OLAP queries.

Visualization of active and completed queries

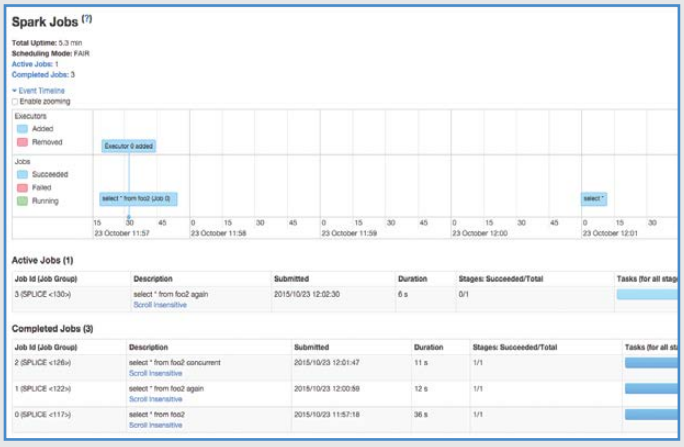

Visualization of stages for each query, plus kill function

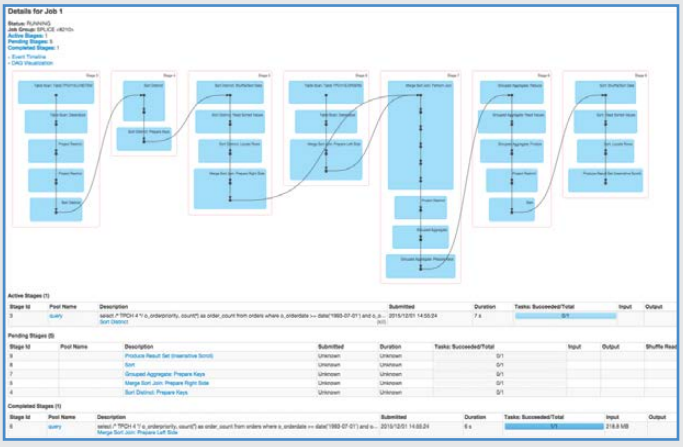

splice MACHINE

www.splicemachine.com
(415) 857-2111
info@splicemachine.com