



Offload Operational Databases to Hortonworks with Splice Machine



Partner Brief



“Our partnership with Hortonworks is able to deliver to our clients 5-10x faster performance and over 75% reduction in TCO over traditional scale-up databases. With Splice Machine’s SQL-based transactional processing engine, our clients are able to migrate legacy database applications without application rewrites.”

Monte Zweben
Chief Executive Officer
Splice Machine

Facing increased data growth and cost pressures, scale-out technology has become very popular as more businesses become frustrated with their costly scale-up RDBMSs. With Hadoop emerging as the de facto scale-out file system, a Hadoop RDBMS is a natural choice to replace traditional relational databases like Oracle and IBM DB2, which struggle with cost or scaling issues.

Designed to meet the needs of real-time, data-driven businesses, Splice Machine is the only Hadoop RDBMS. Splice Machine offers an ANSI-SQL database with support for ACID transactions on the distributed computing infrastructure of Hadoop. Like Oracle and MySQL, it is an operational database that can handle operational (OLTP) or analytical (OLAP) workloads, while scaling out cost-effectively from terabytes to petabytes on inexpensive commodity servers.

Splice Machine, a technology partner with Hortonworks, chose HBase and Hadoop as its scale-out architecture because of their proven auto-sharding, replication, and failover technology. This partnership now allows businesses the best of all worlds: a standard SQL database, the proven scale-out of Hadoop, and the ability to leverage current staff, operations, and applications without specialized hardware or significant application modifications.

What business challenges are solved?

Leverage Existing SQL Tools

Leveraging the proven SQL processing of Apache Derby, Splice Machine is a true ANSI SQL database on Hadoop. Splice Machine is SQL-99 compliant and supports existing applications with minimal migration effort.

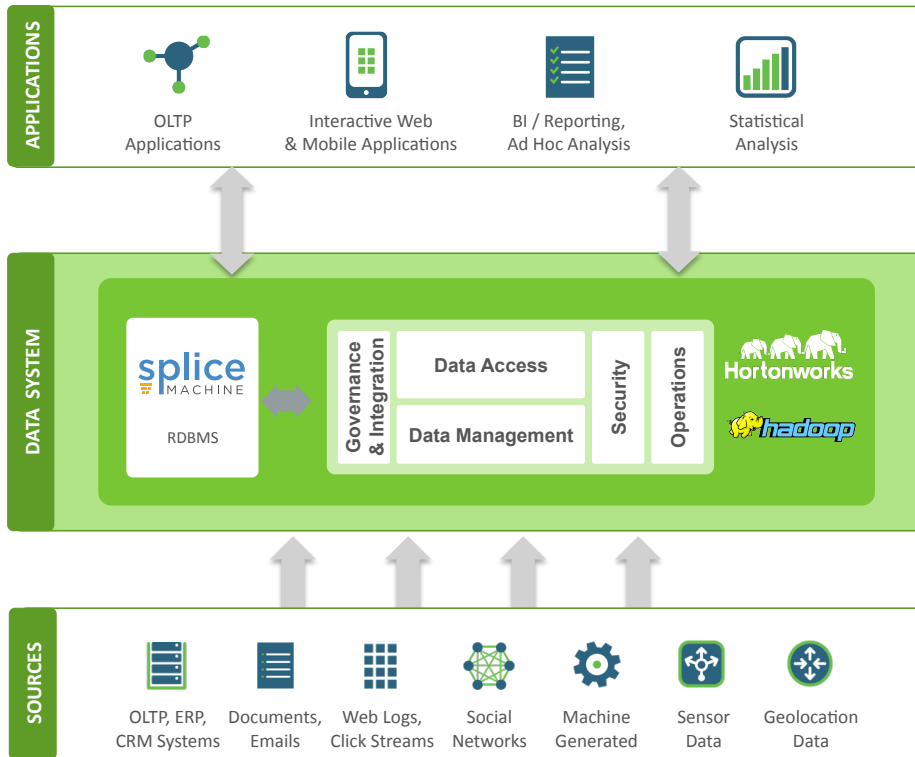
Cost-Effective Scaling

Splice Machine leverages the proven auto-sharding of HBase to scale with commodity servers. This auto-sharding occurs automatically without requiring changes to applications powered by Splice Machine.

Real-Time Updates

Splice Machine provides full ACID transactions across rows and tables by using lockless snapshot isolation. Its state-of-the-art design allows very high throughput and avoids troublesome deadlock conditions.

Splice Machine in the Modern Data Architecture



Features and Benefits of Combined Solution

Splice Machine delivers affordable scale-out with Hortonworks and Hadoop.

Splice Machine is the only Hadoop RDBMS to support full database functionality, such as:

- Transactions
- Joins
- Secondary indexes
- Aggregations
- Sub-queries
- User-defined functions (UDFs)
- Column-level security
- Primary Key constraints
- Window Functions
- Backup/Recovery
- LDAP support
- Native ODBC/JDBC

Seamless Integration

Using HBase co-processors, Splice Machine natively runs in the Hadoop stack, ensuring smooth integration to the ecosystem of Hadoop tools, such as YARN, ZooKeeper, Flume, Storm, Kafka, and Sqoop. HCatalog integration enables MapReduce, Pig, Hive, and Spark to execute queries that combine structured data in Splice Machine with unstructured data in HDFS.

Thanks to standards-based ODBC and JDBC drivers, Splice Machine provides seamless connectivity to BI tools, such as Tableau and MicroStrategy and SQL tools, such as Toad and DbVisualizer.

Splice Machine is a Certified Technology Partner

The Hortonworks Certified Technology Program reviews and certifies technologies for architectural best practices, validated against a comprehensive suite of integration test cases, benchmarked for scale under varied workloads and comprehensively documented.



Hortonworks. We do Hadoop.

Hortonworks is a leading commercial vendor of Apache Hadoop, the open source platform for storing, managing and analyzing Big Data. Hortonworks Data Platform, our distribution of Apache Hadoop, provides an open and stable foundation for enterprises and a growing ecosystem to build and deploy Big Data solutions.

Hortonworks is the trusted source for information on Hadoop, and together with the Apache community, Hortonworks is making Hadoop an enterprise data platform. Hortonworks provides unmatched technical support, training and certification programs for enterprises, systems integrators and technology vendors.

For additional questions, contact:

- Splice Machine
www.splicemachine.com
(415) 857-2111
- Hortonworks
www.hortonworks.com
(855) 8-HORTON