

0.1.0

Initial Release

Release Date: 27 March 2017

Main Features:

1. **Competitively runs SSB and TPC-H benchmarks on a single node**

We have benchmarked Quickstep extensively against 2 popular RDMS benchmarks on standard hardware against other popular databases

2. **Block Oriented Pipelined Execution**

Allows for parallel execution of work units ("WorkOrders") in the pipeline. This makes good use of hardware resources.

3. **Storage type selection**

Storage can be specified as either *column* or *row* oriented. Both storage types can include compressed features. Storage type can be specified via the SQL interface on table creation and changed after table creation, but subject to constraints.

4. **Recent Execution Performance Enhancements**

LIP Filters improve join performance by providing parallel information passing. These filters make Quickstep more robust to poor optimizer join ordering choices.

Parallel GroupBy implementations. There are three implementations of parallel group-by aggregations where the optimizer can choose from: two-phase aggregation / partitioned aggregation / array-based aggregation. They provide good performance for a variety of workloads.