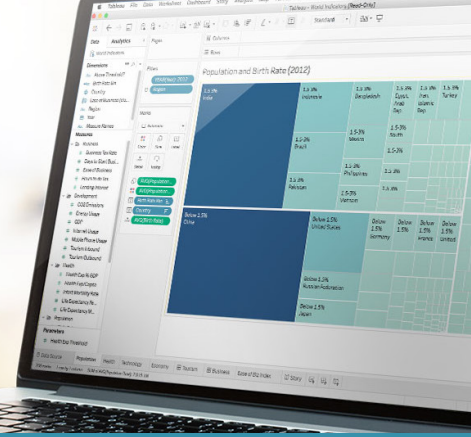


jethro, the BI ENGINE for BIG DATA



Jethro Delivers Interactive Business Intelligence at Hadoop Scale & Costs

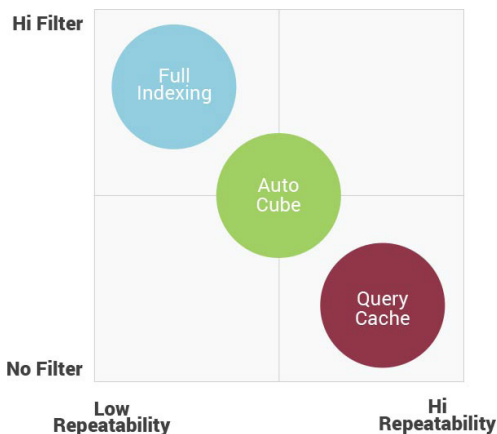
- ✓ BI dashboards respond in seconds on Big Data queries
- ✓ No manual data engineering
- ✓ Seamless integration—no changes to BI application or to underlying Hadoop data lake

Customers use Jethro to deliver Interactive BI service at attractive Hadoop costs. Jethro transparently supports all SQL queries, for thousands of concurrent users analyzing tens of billions of rows. All that with interactive response times measured in seconds.

Interactive Performance for All Types of Queries

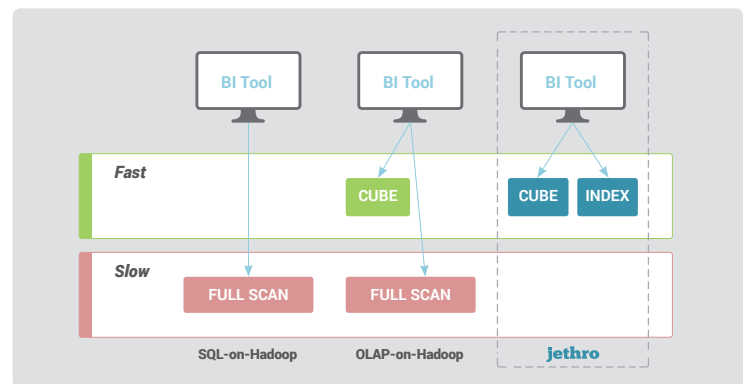
Jethro's Cost- Based Optimizer combines three different strategies to deliver interactive performance across all types of queries:

- **Full Indexing:** every column is automatically indexed
- **Auto Cubes:** every aggregation is automatically turned into a small cube
- **Result Cache:** every query result is automatically cached



Seamless Compatibility

BI: Tableau, Qlik, MicroStrategy, and almost any other BI tool
Hadoop: Any Hadoop distribution including CDH, HDP, EMR, MapR
Cloud: AWS, Azure



Jethro delivers Interactive BI

Jethro can support 1000s of concurrent users dicing and slicing 10s of billions of up-to-the-minute business data. Jethro delivers such high level of service by computing user queries in real time form Indexes, cubes, and query caches that are automatically maintained and kept current by background services. Jethro does not access the underlying data lake to answer user queries.

Jethro System Architecture

BI tools send live queries to Jethro via ODBC / JDBC connection. Jethro servers can dynamically scale out to support any level of concurrency. All data—columns, indexes, cubes, cache—is centrally stored in Hadoop, cloud, or NFS, and shared by all Jethro servers.