



White Paper

Jethro for Tableau: Achieving Interactive Speed on Multi-billion-row Tables on Tableau

November 2015



Jethro for Tableau is an Acceleration Server that Enables Tableau to Live Connect to **Multi-billion-row** Tables at *Interactive Speed*.

Technology Overview

Jethro for Tableau is an innovative solution that combines caching and indexing to accelerate performance of Tableau's Live Connect to large tables. It starts with identifying target tables and then loading them into the Jethro server. Jethro will then fully index the tables (every column is indexed!) and save the resulting data on a local drive, or in a central storage system. Queries use one or multiple indexes to surgically access only the data they need from the Jethro cache. As users drill down, queries become more and more selective, and response becomes even faster.

Easy To Get Started

Jethro for Tableau can be installed on one or multiple virtual servers through a simple install script. The target tables' data is extracted and loaded into Jethro via a simple script. Tableau Desktop and Server users can then connect via ODBC driver and query the data seamlessly.

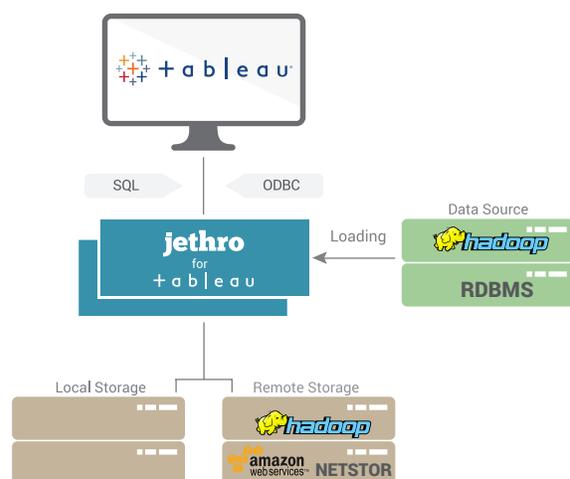
Unique Indexing Technology

A unique index-based approach is the key to Jethro for Tableau's superior performance. The indexes are automatically created for every column, thus enabling

full flexibility for users who can now filter and order data by any column. Jethro's indexing technology is implemented as multi-hierarchy compressed bitmaps that are both storage and memory efficient.

Multi-layered Intelligent Cache

Jethro for Tableau utilizes multiple caching systems to speed up queries. Metadata and highly accessed components are always cached in memory. Hot data file fragments are cached in the server when remote storage is used. Lastly, Jethro caches query results of common queries. This is highly valuable when users



share dashboards as the exact queries are often repeated. As new data is added, the query result cache is not invalidated. Instead, queries are executed on the incremental data and the results are merged with the saved cached query results.

Incremental Updates

As new data is added to the original tables, the same data can be applied to the cached tables in Jethro. Loading incremental data into Jethro typically takes only seconds to complete and can be run as frequently as needed – even every few minutes. During the incremental load, Jethro indexes are appended (instead of updated), ensuring that the performance of current queries is not negatively affected.

Shielding Users from “Noise”

Often times Tableau dashboard and workbook performance is affected by the overall load on the database system. At the same time, activity by Tableau users can affect other workloads sharing the database system. With Jethro for Tableau, all related Tableau activity is executed on the Jethro servers and is completely shielded from any workloads on other systems.

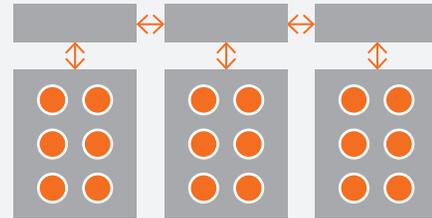
Scalability & High-Availability

Jethro for Tableau’s servers are stateless and highly elastic, allowing easy scale-out to meet concurrency requirements. When a new Jethro server is launched it access all the saved data from the central storage system and can be immediately available for queries. Similarly, as user concurrency load subsides, Jethro servers can be terminated with no interruption.

View a live demo of Jethro for Tableau at jethrodata.com.

Full Scan / Brute Force Other Tableau Acceleration Solutions

1. Reads entire dataset. Every time.



Impact on Performance

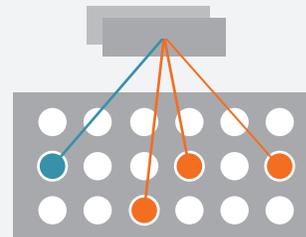
Massive number of unnecessary I/O
Extensive I/O, CPU and memory usage

VS

Index Access

Jethro

1. Analyzes indexes
2. Fetches only relevant data



Impact on Performance

Drastically lower load
Minimal I/O, CPU and memory usage

About Jethro

Jethro was founded by a team of industry veterans committed to making big data analytics work in real time. Our passion is solving big problems, in this case building the technology that lets non-technical users interactively explore data on Hadoop and get immediate answers, using standard SQL or common BI tools.

In June 2015, Jethro closed an \$8.1 million Series B financing round led by Square Peg Capital and existing investor Pitango Venture Capital.

Jethro is headquartered in New York City with an R&D office in Israel.