



Solution Brief

Zeepabyte Cascade Engine Setting new landmark in Star Schema Benchmark

AT A GLANCE

Zeepabyte Cascade Engine is a big data analytics system designed to process massive data economically and efficiently.

Cascade performs complex execution plans in a distributed (multi-node and multi-partition) environment, maximizing available resources.

Its patent pending multidimensional data organization and processing technology enables real time speeds in analytic tasks.

A recent Star Schema Benchmark ran over a cluster connected by Mellanox Technologies end-to-end 10GbE intelligent interconnect solution demonstrated the **world's fastest performance**. Zeepabyte delivers subsecond query response times on billions of data records enabling new business models in many industries.

KEY BENEFITS

- ✓ Highest performance at the lowest cost.
- ✓ May function as analytical database, in memory or disk based.
- ✓ Can performed efficiently in both real-time, interactive and batch modes.
- ✓ Not tied to a particular platform, can run on Hadoop clusters or on devices like Raspberry Pi.
- ✓ Designed to handle SQL and ML workflows on large and dynamically changing data sets.
- ✓ Helps improve data ecology preventing data swamps.



What is Star Schema Benchmark (SSB)

SSB was created by UMB researchers upon request from Michael Stonebraker's Vertica. It derives directly from industry standard TPC benchmarks and reflects **workload typical for data warehouse analytics**.

What is measured

SSB ZEEPABYTE metrics are the same as in TPC-H:

- Performance: Query per hour (Qph) @ Scale Factor (SF)
- Cost of performance: Cost of the system / Qph @ Scale Factor (SF)

Learn more at zeepabyte.com/docs.php

Results - 2x Performance gain, 10x Cost reduction

| Scale factor | Number of records | Zeepabyte Cascade | | | Current best results in TPC-H | | |
|--------------|-------------------|-------------------|-----------|---------------|-------------------------------|-----------|--------------|
| | | Nodes used | Qph@SF | Cost/ Qph@SF | Nodes used | Qph@SF | Cost/ Qph@SF |
| 100 | 600M | 1 | 1,416,786 | \$.014 | 1 | 420,092 | \$.11 |
| 100 | 600M | 3 | 2,855,786 | \$.016 | 6 | 1,582,736 | \$.12 |
| 300 | 1.8B | 4 | 5,683,743 | \$.009 | 12 | 2,948,721 | \$.12 |
| 1000 | 6B | 5 | 9,366,072 | \$.007 | 20 | 5,246,338 | \$.14 |

Benchmark performed with **Mellanox Technologies**, leading supplier of intelligent interconnect solutions for hyper-converged infrastructure - www.mellanox.com

About Zeepabyte, Inc

Zeepabyte is a startup founded in July 2015 by experts in databases, business intelligence and service platforms, with primary focus on Big Data Analytics for BI and IoT. At the heart of Zeepabyte's technology offering is patent pending world's fastest Cascade Analytics System.

Contact us: info@zeepabyte.com; or visit us at zeepabyte.com