



InfiniDB Speeds Ahead of Other Open-Source SQL-on-Hadoop Query Engines in Benchmark Testing

Analysis Reveals the Importance of Speed, Scale, and SQL Capability When Organizations Look to Unlock The Business Value of Big Data

Frisco, TX – April 3, 2014 – [InfiniDB](#), a leading provider of high-performance analytic data platforms, today announced results of a new, independent benchmark from Radiant Advisors that looked at the performance comparison of leading open-source SQL-on-Hadoop query engines. As more companies look to unlock business value from volumes of data stored in the Hadoop Distributed File System (HDFS), demand for SQL access to Hadoop has been skyrocketing. Within the last year, several Hive alternatives have gained acceptance in the Hadoop ecosystem. The benchmark, “Open-Source SQL-on-Hadoop Performance,” was conducted in Q1 of 2014 by Radiant Advisors.

Radiant Advisors selected open-source and free for download SQL engines in various data file encodings in three Hadoop clusters on identical Amazon EC2 servers. Piwik was the application workload selected for the performance benchmark from the open-source web analytics application, piwik.org. The Piwik schema included three tables representing user visits and actions. A base of three years of initial data was provided, and then was duplicated to represent 100x user activities over the same three years. In total, this schema represented approximately one terabyte of user data.

The benchmark and analysis used distributions from Hortonworks, Cloudera, and [InfiniDB](#) that included Hive 0.11/0.12, Impala 1.2, Presto 0.57, and InfiniDB for Hadoop 4.0. After loading data files, all queries were manually run independently at least 10 times in each cluster, SQL engine, and file encoding.

Highlights of the benchmark and testing include:

- The InfiniDB engine delivered up to 40x performance gains over Hive on reporting queries and ran ad-hoc and analytic queries “interactively” that were not run to completion on Hive due to the excessive time required.
- The InfiniDB engine delivered 20-40x performance advantage over Presto. Presto was the runner-up finisher in terms of query coverage and able to complete 9 of 10 queries.
- The InfiniDB engine performed similarly to Impala for reporting queries, but showed significant performance gains on ad-hoc queries and was able to run analytic queries that are unsupported by Impala.
- InfiniDB was the only SQL engine that completed all reporting, adhoc, and analytic queries in the benchmark, and was the fastest in 6 out of 10 queries.

“This benchmark is intended to be insightful and pragmatic for organizations faced with performance challenges in their Hadoop environments,” said John O’Brien, Principal Analyst and CEO, Radiant Advisors. “Additionally, an analytic culture requires high-performance SQL engines to unlock the latent value in Hadoop environments to drive discovery and interactivity.”

Bob Wilkinson, chief operation officer at InfiniDB added, “Organizations are storing and increasing the amount of information in Hadoop, but now they are struggling with how to use that data in a strategic and useful manner. InfiniDB is available as a quick time-to-value solution because it is SQL accessible for

directly querying HDFS. While the InfiniDB analytic platform provides immediate insights, its flexible data structure, deployment options, and scalability ensure long-term value for organizations that need to conquer evolving Big Data problems.”

To download your free copy of the report visit <http://www.infiniDB.co>. A full presentation of the testing process, benchmark results, and interactive discussion is scheduled for Wednesday, April 23. To register and receive a copy of the report go to: <https://www3.gotomeeting.com/register/961338942> .

For more information about InfiniDB and the InfiniDB platform, visit <http://www.infiniDB.co> and keep updated by following [@InfiniDB](https://twitter.com/InfiniDB) .

Tweet this: News: New @RadiantAdvisors benchmark evaluates #SQL on #Hadoop query engines for #analytics. @InfiniDB = FAST #opensource #hdfs

About InfiniDB

InfiniDB empowers organizations to solve problems and create new solutions with powerful Big Data analytics. The company’s platform is a fourth-generation massive parallel processing (MPP) column-oriented data technology that is known for its rapid implementation, simplicity and extraordinary value. InfiniDB, InfiniDB for the Cloud, and InfiniDB for Apache™ Hadoop® are built for today’s growing enterprise. These organizations demand speed, scale and efficiency in their analytics platforms where leveraging traditional and emerging data technologies, structures and architectures are required. InfiniDB products are licensed as GPL-2.0 with complementary consulting services, maintenance and support agreements.

For more information, to join the community, and download software, visit www.InfiniDB.co and follow @InfiniDB.

About Radiant Advisors

Radiant Advisors is a leading strategic research and advisory firm that delivers innovative, cutting-edge research and thought-leadership to transform today’s organizations into tomorrow’s data-driven industry leaders.

To learn more, visit <http://www.RadiantAdvisors.com> or follow on Twitter @RadiantAdvisors.

InfiniDB and the InfiniDB logo are registered trademarks of InfiniDB, Inc. Other product names and logos may be trademarks or registered trademarks of their respective owners.

###

Contact:

Mark Peterson
Peterson Communications for InfiniDB
831.626.4400
mark@petersoncom.com