



Calpont Announces InfiniDB for Apache Hadoop

Integrates with Hadoop Distributed File System (HDFS) to Enable High-Performance Analytics Directly Within an Apache Hadoop Cluster. Fourth Generation InfiniDB Platform is Significantly Faster than Competing Solutions and Provides a Robust Library of SQL Syntax for Big Data Developers.

Frisco, TX and New York, NY – October 29, 2013 – Calpont Corporation, a leading provider of high-performance analytic databases today announced at the [Strata Conference + Hadoop World](#), InfiniDB® for Apache Hadoop™. InfiniDB for Apache Hadoop complements organizations' Hadoop infrastructure with a columnar architecture that is optimized for rapid and easy access to highly dimensional data and has the flexibility, scale and horsepower to work with massive data sets.

Hadoop and HDFS™ have been a great way for organizations to capture, process and efficiently store copious amounts of unstructured data. But there is a difference between collecting data and querying data. “InfiniDB for Apache Hadoop differentiates itself by empowering organizations to actively use, engage and execute high-performance analytics directly within an Apache Hadoop cluster against all of the data. InfiniDB for Apache Hadoop incorporates SQL query language and data structures for Hadoop deployments without constraining or bottlenecking performance. This is simply not possible with any other technology or platform on the market,” said Jim Tommaney, chief technology officer at Calpont.

Highlights of InfiniDB for Apache Hadoop include:

- InfiniDB for Apache Hadoop provides structured query performance for both large dimensional or non-dimensional data sets where speed and reliability are paramount to business execution.
- No need to custom code or learn new programming languages for Hadoop, Big Data developers can use the familiar and rich SQL syntax to develop high performance analytic queries.
- InfiniDB is the first MySQL storage engine to implement the SQL standard Analytical Functions to enable new in-database analytics horsepower and further leverage existing developers, analytic tools, and solutions.
- As part of the InfiniDB platform, organizations have the ability to run the same technology on their choice of local disk, GlusterFS, AWS, and HDFS without any compromise of performance, scale or flexibility.
- InfiniDB incorporates a library of embedded frameworks, job scheduling and cluster resource management features that are battle-proven for high performance Big Data analytics.

Shortcomings of Many Hadoop Implementations

More organizations are using Hadoop for large scale structured and unstructured data. However, the sheer volume of data collected often results in a Hadoop implementation that is no longer performing at the desired speed for analysis. Unlike parallel, column-oriented databases such as InfiniDB for Apache Hadoop, many Hadoop implementations don't distribute query processing and have to write intermediate query results to local data structures as part of the query execution process that can result in bottlenecks. In addition to slowing query response, this results in a mishmash of custom-coded infrastructure that can limit future technology scaling to support growing business requirements.

Proven, Fourth Generation InfiniDB is Optimized for Performance and is Faster

The InfiniDB for Apache Hadoop is optimized to do faster queries and scale as processing demands grow. InfiniDB is built for Big Data analytics, with a massively parallel processing (MPP), column-oriented architecture that is designed specifically for read intensive analytic applications and data warehouses that consume large amounts data. InfiniDB scales performance linearly with hardware, is easy to deploy and maintain, and requires no special database tuning, indexing or materialized views, making it ideally suited for Big Data analytics. Results of benchmarking tests that demonstrate the robust and consistent performance are available on the [InfiniDB for Apache Hadoop webpage](#).

Disruptive Technology with Licensing to Match

Recently, Calpont announced its core platform is open-source GPLv2.0 license. This is a perfect complement for organizations turning to Hadoop and other emerging technologies that are seeking to experience Big Data analytics by scaling InfiniDB across multiple processing cores, servers and huge data sets without software licensing limitations. See the announcement [Calpont Announces New Open Source Licensing for InfiniDB](#).

More Information on the InfiniDB Platform

Calpont is exhibiting at the [Strata Conference + Hadoop World](#) in New York City, October 28-30, where attendees can learn more about InfiniDB 4, InfiniDB for the Cloud, and InfiniDB for Apache Hadoop by visiting Calpont at booth #40. Organizations can also dive right in to using InfiniDB by downloading the software for at www.calpont.com.

About Calpont

Calpont empowers data superstars to solve problems and create new solutions with powerful Big Data analytics. The company's platform, InfiniDB, 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 that demand speed, scale and efficiency across their analytics platforms. InfiniDB products are licensed as GPL-2.0 with complementary consulting services, maintenance and support agreements available from Calpont.

For more information, visit www.calpont.com, join the community www.infinidb.org, or follow @Calpont and @InfiniDB.

###

Contact:

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

Apache Hadoop and HDFS are trademarks of the Apache Software Foundation.