Hortonworks

In-Memory Computing for Hadoop

SCALEOUT SOFTWARE

Hortonworks and ScaleOut Software make operational intelligence a reality. Now you obtain real-time insights from your live, fast-changing data by running Hadoop MapReduce applications and Hive queries using ScaleOut hServer®'s in-memory MapReduce engine. Breakthrough technology extends the industry-leading Hortonworks Data Platform to enable real-time analytics on live data using standard MapReduce applications and Hive queries.

ScaleOut hServer installs either stand-alone or on a Hadoop cluster and adds both a scalable, in-memory data grid and an in-memory Hadoop MapReduce engine designed for real-time analytics. MapReduce applications run unchanged with sub-second startup times and can access memory-based data sets or disk-based data sources, such as HDFS. Using YARN, ScaleOut hServer also accelerates Hive queries using Hortonworks' Hive distribution.

ScaleOut hServer opens the door to operational intelligence by enabling analysis of live, memory-based data and avoiding costly batch-processing overheads. Hortonworks customers now have an integrated solution for a full spectrum of use cases within a single, familiar Hadoop MapReduce platform combining real-time insights on live data with batch processing on historical data.

Achieve Operational Intelligence

- **Run in-memory MapReduce on live, operational data**

- **Accelerate Hive queries with in-memory MapReduce**

- **Consolidate batch & operational architectures**

- **Prototype Hadoop MapReduce applications**
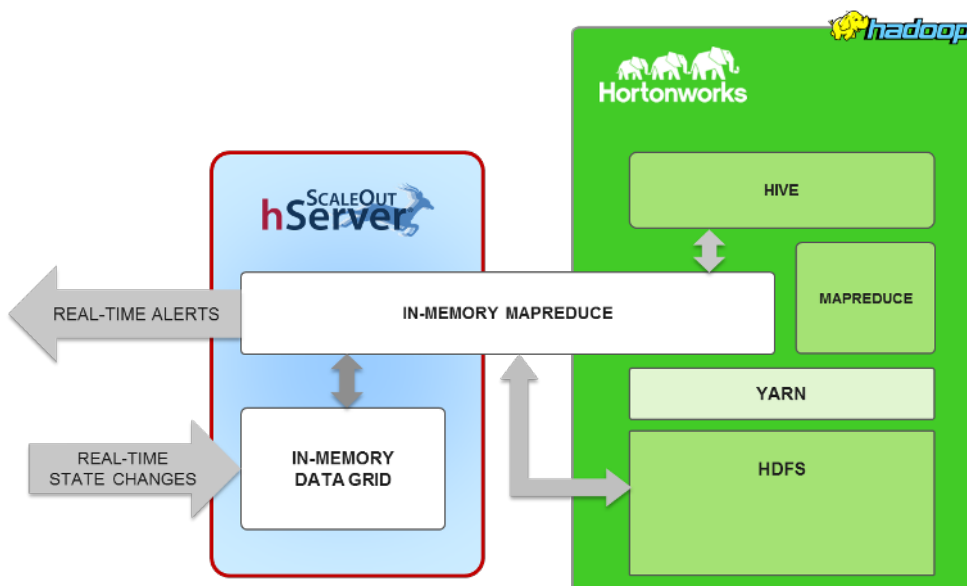
What Business Challenges are Solved?

- **Run MapReduce on live data**
  Now you can take the deep insights you've discovered in batch Hadoop processing and implement them for live, operational data. ScaleOut hServer is the only solution that delivers this capability without any change to your MapReduce code.

- **Run Hive queries on in-memory data**
  You can leverage your investment in Hive for use with live, in-memory data. ScaleOut hServer is the only platform that allows you to run Hive unchanged using in-memory MapReduce, and you can query live, operational data in addition to HDFS data sets.

- **Consolidate batch & operational clusters**
  ScaleOut hServer lets you consolidate your batch and operational environments into a single architecture, all under the Hortonworks Data Platform. You also can use ScaleOut hServer to perform ETL on data ingested into HDFS for batch processing.

ScaleOut Software - a Certified Technology Partner

As a pioneer in in-memory computing since 2005, ScaleOut Software develops in-memory computing products that help our customers solve performance challenges and gain a competitive edge for their businesses. ScaleOut Software is a Certified Technology Partner on Hortonworks Data Platform (HDP). The Hortonworks Certified Technology Program reviews and certifies technologies for architectural best practices, validated against a comprehensive suite of integration test cases, benchmarked for scale under varied workloads and comprehensively documented.

"Integrating with Hortonworks makes it possible for Hortonworks users to analyze live, fast-changing data sets using the familiar Hadoop MapReduce framework. Until recently, users were limited to batch processing on static data sets. Using our in-memory computing platform, they now are able to immediately react to patterns and trends in operational data," said Bill Bain, Founder and CEO of ScaleOut Software. "In industries where time is a competitive advantage, the benefits are enormous."

For additional questions, contact:

ScaleOut Software
sales@scaleoutsoftware.com
(503) 643-3422
www.scaleoutsoftware.com
Hortonworks
USA: (855) 8-HORTON (1 for sales)
www.hortonworks.com

Hortonworks

3460 West Bayshore Road
Palo Alto, CA 94303 USA

US: 1.855.846.7866
International: 1.408.916.4121
www.hortonworks.com