



Data Analytics Firm Accelerates Business Expansion

Cisco UCS supports the SPINS deployment of the MapR Distribution for Apache Hadoop to provide a faster business intelligence service

Executive Summary

Challenge

- Increase scale of services provided to customers
- Speed up report generation time
- Control costs during expansion

Solution

- MapR M7: Enterprise Database Edition
- Cisco UCS CPA for Big Data

Results

- Improved data processing speed by an order of magnitude
- Reduced troubleshooting time from days to hours
- Gained additional scalability at minimal cost
- Simplified Big Data cluster management

“It’s well known in the Hadoop community that MapR on Cisco UCS architecture is the fastest out there. The server platform is faster, cheaper, more scalable, and innovative compared with the competitors.”

Jim Scott, Senior Vice President, Information Technology and Operations, SPINS

Challenge

Founded in 1998, SPINS is a data analytics and business intelligence company focused on the natural, organic, and specialty products industries. The company provides consumer insights, retail measurement, product libraries, analytics reporting, and consulting services for retail and manufacturing clients. In 2004, the company installed a relational database as a single-source solution to store and analyze data.

But as SPINS attracted more customers, requiring ever more detailed data analyses, the system started taking longer to respond to requests for reports. A report might take several minutes or even hours, sometimes giving the impression that the service was down, causing frustration and reducing customers’ ability to react quickly and decisively to the very rapid changes that affect their highly dynamic markets.

The logical response was to scale the existing data center, but the price suggested by the incumbent provider was considered excessive. In addition, the database was better suited to generating a significant number of short reports, whereas SPINS was increasingly offering very large reports, several millions of rows long.

Solution

In December 2013, SPINS decided to replace its existing database with MapR M7 Enterprise Database Edition, which is a distribution of Apache Hadoop. Jim Scott, senior vice president of Information Technology and Operations at SPINS, says: “MapR is the best distribution for Hadoop out there. With MapR I get raw performance and a platform that can support innovation in business intelligence. There is a magnitude of difference in performance.”

Scott says various hardware vendors were considered when it came to the selection of servers that could offer the best platform for the continued growth of the company. In the end, however, the company chose a Cisco® Unified Computing System® (UCS™) Common Platform Architecture (CPA) for Big Data.

Business Benefits

Scott says: “The innovative Cisco UCS architecture is recognized by the Hadoop community to be faster, cheaper, more scalable, and innovative, compared with the competitors. Cisco’s innovation in the server platform took a step further, making it easier to manage, simpler to troubleshoot, and a more cost-effective solution. MapR on Cisco UCS set the TeraSort record on a tenth of the hardware.”

SPINS purchased the UCS Big Data High Capacity bundle, that includes Cisco UCS C240M3 Rack Servers, Cisco UCS 6296UP 96-Port Fabric Interconnects and Cisco Nexus® 2232PP 10GE Fabric Extenders. SPINS now stores its data in MapR M7, a complete distribution for Apache Hadoop that provides enterprise-grade reliability and performance optimizations for Apache HBase applications. MapR lets customers query data quickly using tools such as JasperSoft and QlikView.

“When you have Cisco UCS and MapR technology that simplifies scalability, you can do things differently,” Scott says. “We don’t have to follow the same prescribed path. It gives us the ability to go in many directions. The combination of MapR and Cisco UCS CPA delivers a high-performance and scalable big data solution”

Results

SPINS has transformed its business to provide analytic solutions that will allow for consistent query processing, over larger volumes of data, for its growing customer base. As its customer base increases, scalability and linear performance assurance provided by Cisco UCS CPA and MapR is vital to its continued customer satisfaction.

Scott reports order-of-magnitude improvements in data processing speed, with resource-intensive analytics tasks that took an hour previously being achieved in six minutes on Hadoop. “The speed of the MapR database is a key differentiator for business intelligence. When you can deliver a product or data faster, you get an edge,” says Scott.

The system also supports enterprise features such as high availability and disaster recovery to enable higher uptime. Having all of its data in one central repository enables SPINS to work on larger datasets and to ask new and better questions on the entire dataset. It can keep more data longer to provide extended histories so customers get a performance view over time.

“Anyone who has dealt with a lot of data knows you have to ask a lot of questions to figure out the right question,” explains Scott. “MapR on Cisco UCS CPA platform gives you the ability to realize you’ve asked the wrong question faster.”

System troubleshooting times have been reduced from days to hours, thanks to the Cisco UCS Manager’s ability to isolate problems between individual servers rather than SPINS having to capture network data in order to resolve issues.

SPINS expects to continue to save money once its big data operation scales up further, as additional servers for the new system do not require extra staff, allowing the company to significantly expand its operation at minimal cost.

By removing limits on performance and scalability, the Cisco UCS CPA for Big Data and MapR is helping SPINS provide new and enhanced services to its customers and open up opportunities to take the business in entirely new directions.

Product List

Data Center

- Cisco Unified Computing System (UCS) CPA for Big Data
 - Cisco UCS C240 M3 Rack Servers
 - Cisco UCS 6296UP 96-Port Fabric Interconnects
 - Cisco Nexus 2232PP 10GE Fabric Extenders

Software

- MapR M7 Enterprise Database Edition (a distribution for Apache Hadoop)

For More Information

To learn more about the Cisco architectures and solutions featured in this case study go to:

- www.cisco.com/go/ucs