

Cisco UCS Demonstrates Leading TeraSort Benchmark Performance



Performance Brief
August 2013

Highlights

Industry Leading Performance

- Cisco's results for TeraSort with Cisco Unified Computing System™ (Cisco UCS®) have set a new record for performance and scalability.

Scalability for Big Data Workloads

- Cisco UCS Common Platform Architecture (CPA) for Big Data offers linear scalability along with essential operation simplification for single-rack and multirack deployments.

A Choice of Hadoop Distributions

- Cisco UCS supports leading Hadoop distributions, including Cloudera, HortonWorks, Intel, MapR, and Pivotal

Performance and scalability are critical to the success of enterprise big data analytics. In recent testing using the popular TeraSort benchmark, Cisco UCS Common Platform Architecture for Big Data demonstrated leading performance and scalability.

The need for agile big data solutions is growing faster than ever before, since effectively capitalizing on information can make the critical difference between success and failure. At the same time, even as more companies are faced with the prospect of deploying big data infrastructure, data sets are becoming larger. By using the TeraSort benchmark as a popular workload for stress-testing Hadoop clusters, it is finally possible to evaluate big data capabilities between vendors.

Cisco UCS Common Platform Architecture for Big Data

Through its unique architecture, Cisco Unified Computing System™ (Cisco UCS) has been very successful in unifying computing, networking, management, virtualization, and storage access resources into a single integrated architecture. This innovative approach enables end-to-end server visibility, management, and control, and it offers automated configuration for faster deployment and rapid scaling. By extending these critical innovations to big data, Cisco can provide industry-leading infrastructure for enterprise-class Hadoop deployments.

Cisco UCS Common Platform Architecture (CPA) for Big Data extends the strengths of Cisco UCS by offering computing and network scalability, performance, management, and monitoring that yields essential operation simplification, modularity, risk reduction, and lower total cost of ownership (TCO). Integral Cisco UCS fabric interconnects offer a common management plane for scale-out designs with both single-rack and multirack form factors. Cisco UCS virtual interface cards (VICs) enable unified fabrics for single-wire management and direct SAN access. When applied to big data workloads, the architecture yields compelling performance advantages and significantly lower operating costs. As a result, Cisco CPA for Big Data is gaining considerable momentum and is being installed in major industries including finance, healthcare, insurance, and government.

TeraGen and TeraSort

TeraGen and TeraSort have become a popular choice to benchmark and stress test Hadoop clusters. The TeraGen benchmark uses the MapReduce model to generate data for subsequent processing. TeraSort samples and sorts the input data across a Hadoop cluster, combining testing of the Hadoop Distributed File System (HDFS) and MapReduce layers. With multiple vendors now publishing TeraSort benchmark results, organizations can make performance comparisons between vendor solutions, using similar data-set sizes as guides.

Industry-Leading Performance

Cisco engineers recently conducted TeraGen and TeraSort benchmark testing on a Cisco UCS CPA for Big Data rack.¹ As shown in Figure 1, the architecture not only performed well but scaled in a near-linear fashion with increasing data set size. Sort rate also remained consistent at 3.2 GBytes per per second across the entire range.

These results demonstrate remarkable performance and scalability over a range of data set sizes when using Cisco UCS CPA for Big Data. Cisco also




Figure 1. TeraGen and TeraSort on Cisco UCS CPA for Big Data.

offers distinct performance advantages over competitive published results from other vendors. For example, Cisco's demonstrated 10-terabyte (TB) result is 40 percent faster than HP's published results for an 18-node HP ProLiant Generation 8 (Gen8) DL380 cluster (10 TB TeraSort completing in 5,128 seconds, http://www.hp.com/hpinfo/newsroom/press_kits/2012/HPDiscover2012/Hadoop_Appliance_Fact_Sheet.pdf).

Conclusion

Cisco UCS Common Platform Architecture for Big Data represents an excellent deployment platform for

enterprise-class big data deployments. This well-designed architecture delivers significant performance advantages and lower cost of ownership along with essential scaling capabilities across a broad range of Hadoop-based applications.

For More Information

For more information about big data, please and Cisco UCS performance, please visit:

- blogs.cisco.com/datacenter/cpa
- www.cisco.com/go/bigdata
- www.cisco.com/go/ucsatwork

1. Testing was performed using the Apache Hadoop open source distribution running on 16 high-density Cisco UCS C240 M3 Rack Servers equipped with two Intel Xeon E5-2665 processors, 256 GB of memory, Cisco UCS VIC 1225, and LSI 9266 8i with 24 1-TB SATA 7200-rpm disk drives.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.