

World-Record Results for Big Data Using Cisco UCS with Solid-State Disks and 40-Gbps Networking



New TPC Express Benchmark HS (TPCx-HS) Results for Cisco's Fourth-Generation Big Data and Analytics Platform

Performance Brief
December 2016

Highlights

Best Performance and Price/Performance at the 3-TB Scale Factor

- Our world-record-setting performance result improves on the previous record by 14.5 percent by moving to a 40-Gbps network fabric, from hard-disk drives (HDDs) to solid-state disks (SSDs), and to the Intel® Xeon® processor E5-2600 v4 family.

History of Benchmark Leadership

- The Cisco Unified Computing System™ (Cisco UCS®) has captured more than 120 world-record performance results since its introduction.
- We are the first to publish Transaction Processing Council Express Benchmark-HS (TPCx-HS) results using SSDs.

Industry-Leading Solution

- The benchmark was run on Cisco UCS Integrated Infrastructure for Big Data and Analytics, a platform widely adopted across industry segments.

Innovation, Integrated Management, and Scalability

- The tested configuration consists of servers connected through Cisco® SingleConnect technology managed by Cisco UCS Manager and can be scaled to support a larger cluster.

High Performance at Lower Cost

- With innovations in servers, storage, and fabric, we are the first to demonstrate the use of SSDs and 40-Gbps network connectivity.

Cisco UCS® Integrated Infrastructure for Big Data and Analytics harnessing the power of Intel® solid-state disks (SSDs) and 40-Gbps network bandwidth delivers the best results on the TPC Express Benchmark HS (TPCx-HS) at the 3-TB scale factor.

Capitalizing on the abundant information you have available can make the difference for your company in an increasingly competitive and time-sensitive world. That's why many companies are turning to big data solutions that can help them make better decisions faster. However, big data collection and analytics applications can overwhelm traditional IT infrastructure. With decreasing prices, SSDs are becoming a cost-effective alternative to hard-disk drives (HDDs). But how do you find the right solution that delivers the best performance at the right price?

Table 1 This New Result Demonstrates a 14.5 Percent Performance Improvement at Lower Cost Compared to Our Last World-Record-Setting Result

Number of Cisco UCS C240 M4 Rack Servers with Disk and CPU Type	Connectivity and Management	Performance and Price/Performance	Availability Date and Disclosure
16 data nodes with small form-factor (SFF) Enterprise Value 6-Gbps SATA SSDs, each and 2 Intel Xeon processor E5-2690 v4 CPUs	2 Cisco UCS 6332 16UP Fabric Interconnects Cisco UCS Manager	13.47 HSph@3TB US\$41,830.67/ HSph@3TB	December 20, 2016 Disclosure
16 data nodes with 24 1.2-TB 10,000-rpm SFF disk drives each and two Intel Xeon processor E5-2680 v3 CPUs	2 Cisco UCS 6296 96-Port Fabric Interconnects Cisco UCS Manager	11.76 HSph@3TB US\$44,052.98/ HSph@3TB	September 24, 2015 Disclosure

World-Record Results for Big Data Using Cisco UCS with Solid-State Disks and 40-Gbps Networking

Recently, we demonstrated the performance capabilities of Cisco UCS Integrated Infrastructure for Big Data and Analytics using the TPCx-HS benchmark, an industry-standard benchmark that gives you a vendor-neutral way to evaluate the performance and price-to-performance ratio of solutions. This benchmark evaluates the raw performance as well as the cost of introducing SSDs for Hadoop-based deployments (Table 1).

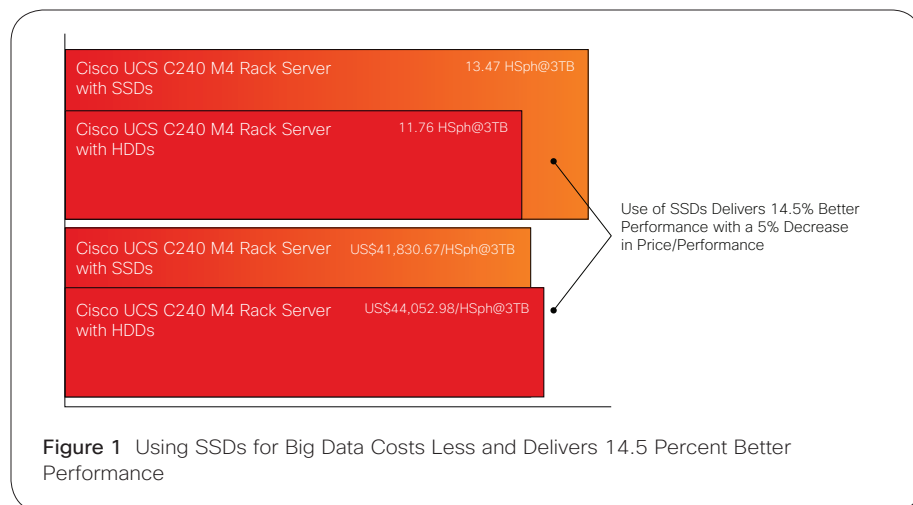
The performance of any Hadoop job is directly affected by underlying infrastructure, which includes computing, networking, and storage. With our new SSD-based configuration, we are addressing three common sources of performance bottlenecks: networking, I/O bandwidth, and processing power (Figure 1):

- Upgrading to the latest generation Cisco UCS fabric interconnects, we increased the underlying fabric from 10 to 40 Gbps.
- We replaced traditional spinning disks with SSDs.
- We incorporated the latest generation of Intel® Xeon® processors.

The result of these changes is a 14.5 percent improvement over our prior record-setting result at the 3-TB scale factor, with an associated 5 percent reduction in cost.

Big Results for Big Data Deployments

The fourth generation of Cisco UCS Integrated Infrastructure for Big Data



and Analytics captures this leading result using SSDs on TPCx-HS at a scale factor of 3 TB. It also achieved and continues to hold world-record results at scale factors of 1, 30, and 100 TB.

Configuration

The results reported here were achieved using Cisco UCS Integrated Infrastructure for Big Data and Analytics. The configuration, illustrated in Table 2, is centrally connected and managed by 2 Cisco UCS 6332-16UP Fabric Interconnects. Compute nodes consist of 16 Cisco UCS C240 M4 Rack Servers powered by the new Intel Xeon processor 2600 v4 family.

- **Cisco UCS 6332-16UP Fabric Interconnects** provide high-bandwidth, line-rate, lossless 40 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE), and Fibre Channel functions. They also provide low-latency connectivity for servers and embedded management

for all connected devices. Cisco UCS Manager enables rapid and consistent server integration using Cisco UCS service profiles, which automate ongoing system maintenance activities such as firmware update operations across the entire cluster as a single operation, advanced monitoring, and options to raise alarms and send notifications about the health of the entire cluster.

- **Cisco UCS C240 M4 Rack Servers** are powerful 2-rack-unit (2RU) servers. For the results reported in this document, they were configured with two Intel Xeon processor E5-2690 v4 CPUs, 256 GB of memory, and eight 1.6-TB 2.5-inch Enterprise Value SSDs. Cisco UCS virtual interface cards (VICs) provided high-bandwidth and low-latency cluster connectivity with support for up to 256 virtual devices.

World-Record Results for Big Data Using Cisco UCS with Solid-State Disks and 40-Gbps Networking

Quick Deployments

Cisco UCS Integrated Infrastructure for Big Data and Analytics combines these components into a unified, fabric-based architecture optimized for big data workloads and is available as a single part number solution bundle. You can quickly deploy and cost-effectively expand this efficient, scalable, high-performance solution. You can scale to thousands of servers with Cisco Nexus® 9000 Series Switches, and you can scale to hundreds of petabytes of storage, all controlled from a single pane with Cisco UCS management. Widely adopted across industry vertical markets, the solution gives you a proven, fast, and simple way to deploy big data environments, accelerate data analysis, and quickly deliver results to users.

For More Information

- Read the [Cisco® blog about the TPCx-HS publication](#).
- Review the [TPCx-HS benchmark specifications](#).
- Learn more about [Cisco UCS performance](#).

Table 2 Cisco UCS Integrated Infrastructure for Big Data and Analytics Solution Used to Achieve This Result

Component	Cisco UCS Integrated Infrastructure for Big Data and Analytics with SSDs
UCS-SPBD-C240M4-H2	High-performance configuration with SSDs
Connectivity	2 Cisco UCS 6332 Fabric Interconnects
Servers	17 Cisco UCS C240 M4 Rack Servers, each with: <ul style="list-style-type: none">• 2 Intel Xeon processor E5-2690 v4 CPUs 14 cores, 2.6 GHz• 256 GB of memory• 8 x 1.6-TB 6G Enterprise Value SATA SSD drives• 2 x 240-GB 6G Enterprise Value SATA SSD drives for boot• Cisco UCS Virtual Interface Card (VIC) 1387
Storage controller	Cisco 12-Gbps SAS Modular RAID Controller with 2-GB flash-backed write cache (FBWC)

- For more information about Cisco UCS big data solutions, please visit <http://www.cisco.com/go/bigdata>.
- Visit the Cisco® big data design zone at <http://www.cisco.com/go/bigdata/design>.
- Read the blog about our fourth-generation big data and analytics solutions at <http://blogs.cisco.com/datacenter/cpav4>.

Disclosures

The Transaction Processing Performance Council (TPC) is a nonprofit corporation founded to define transaction processing and database benchmarks, and to disseminate objective and verifiable performance data to the industry. TPC membership includes major hardware and software companies. The performance results described in this document are derived from detailed benchmark results available as of December 20, 2016, at http://www.tpc.org/tpcx-hs/results/tpcxhs_perf_results.asp



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.