



*Cisco Systems, Inc.*

---

TPC Express Benchmark™ HS (TPCx-HS)  
Full Disclosure Report  
for  
Cisco UCS Integrated Infrastructure for Big Data  
(with 8 Cisco UCS S3260 2-Node Servers)  
using  
Cloudera Enterprise Edition  
and  
Red Hat Enterprise Linux Server Release 6.7

---

**First Edition**

**November 11, 2016**

**Cisco Systems, Inc. (Cisco)**, the Sponsor of this benchmark test, believes that the information in this document is accurate as of the publication date. The information in this document is subject to change without notice. The Sponsor assumes no responsibility for any errors that may appear in this document.

The pricing information in this document is believed to accurately reflect the current prices as of the publication date. However, the Sponsor provides no warranty of the pricing information in this document.

Benchmark results are highly dependent upon workload, specific application requirements, and system design and implementation. Relative system performance will vary as a result of these and other factors. Therefore, the TPC Express Benchmark™ HS should not be used as a substitute for a specific customer application benchmark when critical capacity planning and/or product evaluation decisions are contemplated.

All performance data contained in this report was obtained in a rigorously controlled environment. Results obtained in other operating environments may vary significantly. No warranty of system performance or price/performance is expressed or implied in this report.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at [www.cisco.com/go/trademarks](http://www.cisco.com/go/trademarks). Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company.

TPC Benchmark™, TPCx-HS and HSpH, are registered certification marks of the Transaction Processing Performance Council.

The Cisco products, services or features identified in this document may not yet be available or may not be available in all areas and may be subject to change without notice. Consult your local Cisco business contact for information on the products or services available in your area. You can find additional information via Cisco's web site at [www.cisco.com](http://www.cisco.com). Actual performance and environmental costs of Cisco products will vary depending on individual customer configurations and conditions.

**Copyright © 2016 Cisco Systems, Inc.**

All rights reserved. Permission is hereby granted to reproduce this document in whole or in part provided the copyright notice printed above is set forth in full text or on the title page of each item reproduced.



## Cisco UCS Integrated Infrastructure for Big Data

TPCx-HS Rev. 1.4.1  
TPC-Pricing Rev. 2.0.0

Report Date:  
**November 11, 2016**

Total System Cost

TPCx-HS Performance Metric

Price/Performance

**572,353 USD**

**11.34**  
HSph@300TB

**50,472.05 USD**  
\$/HSph@300TB

Scale Factor

Apache Hadoop  
Compatible Software

Operating System

Other Software

Availability Date

300TB

Cloudera Enterprise  
Edition

Red Hat Enterprise Linux  
Server Release 6.7

None

November 11,  
2016

### System Configuration



8 x Cisco UCS S3260 Dual-node Servers (16 DataNodes)  
Per Node:

- 2 x Intel Xeon E5-2680 v4
- 256 GB Memory
- 28 x 8TB LFF HDD
- 2 x 480 GB 2.5 inch Enterprise Value SSD (BOOT)



1 x Cisco UCS C240 M4 Server (NameNode)  
2 x Intel Xeon E5-2690 v3  
256 GB Memory  
2 x 1.8TB 12G SAS 10K rpm SFF HDD(4K)  
2 x 120GB 2.5 inch Enterprise Value SSD (BOOT)



2 x Cisco UCS 6332 32-Port Fabric  
Interconnects

40GigE

1\*40GigE

Upstream

Physical Storage/Scale Factor: 12.01

Scale Factor/Physical Memory: 70.59

Servers:	8 x Cisco UCS S3260 2-Node Servers (DataNodes), 1 x Cisco UCS C240 M4 Server (NameNode)
Total Processors/Cores/Threads	34/472 /944
Server Configuration:	Per Node:
Processors	2 x Intel® Xeon® Processor E5-2680 v4, 2.40 GHz, 35 MB L3 (DataNodes) 2 x Intel® Xeon® Processor E5-2690 v3, 2.60 GHz, 30 MB L3 (NameNode)
Memory	256GB
Storage Controller	1 x Cisco UCS C3000 RAID Controller M4 Server w 4G RAID Cache (DataNodes) 1 x Cisco 12Gbps Modular SAS HBA (NameNode)
Storage Device	28 x 8TB LFF HDD (DataNodes) 2 x 480GB 2.5in. Enterprise Value SSD (DataNodes, boot disk) 2 x 1.8TB 12G SAS 10K rpm SFF HDD (NameNode) 2 x 120GB 2.5 Enterprise Value SSD (NameNode, boot disk)
Network	Cisco VIC 1300
Connectivity:	2 x Cisco UCS 6332 32-Port Fabric Interconnect
Total Rack Units:	(8*S3260)+(1*C240)+(2*FI) = (8*4RU)+(1*2RU)+(2*1RU) = 32+2+2 = 36



## Cisco UCS Integrated Infrastructure for Big Data

TPCx-HS Rev. 1.4.1  
TPC-Pricing Rev. 2.0.0

Report Date:  
November 11, 2016

Description	Part Number	Source	Unit Price	Qty	Extended Price	3 Year Maint. Price
<b>Cisco S3260 2 Node Bundle</b>						
UCS C3000 M4 Server Node for Intel E5-2600 v4	UCSS-SP-S3260-BV1	1	\$97,360.20	8	\$ 778,881.60	
2.40 GHz E5-2680 v4/120W 14C/35MB Cache/DDR4 2400MHz	UCSC-C3K-M4SRB	1	\$ -	2	\$ -	
16GB DDR4-2400-MHz RDIMM/PC4-19200/single rank/x4/1.2v	UCS-CPU-E52680E	1	\$ -	2	\$ -	
Cisco UCS C3000 RAID Controller M4 Server w 4G RAID Cache	UCS-MR-1X161RV-A	1	\$ -	16	\$ -	
Cisco UCS C3X60 Server Node CPU Heatsink	UCS-C3K-M4RAID	1	\$ -	1	\$ -	
UCSC C3X60 480GB Boot SSD (Gen 2)	UCSC-HS-C3X60	1	\$ -	2	\$ -	
Cisco UCS C3260 System IO Controller w ith VIC 1300 incl.	UCS-C3X60-G2SD48	1	\$ -	2	\$ -	
UCSC 3X60 8TB NL-SAS 7.2KHelium HDD w ith HDD Carrier	UCSC-C3260-SIOC	1	\$ -	2	\$ -	
Cisco UCS C3260 Base Chassis w /4x PSU, SSD, Railkit	UCSC-C3X60-HD8TB	1	\$ -	48	\$ -	
Power Cord Jumper, C13-C14 Connectors, 2 Meter Length	UCSC-C3260	1	\$ -	1	\$ -	
Cisco UCS C3160 System Bezel	CAB-C13-C14-2M	1	\$ -	4	\$ -	
UCS C3X60 Rack Rails Kit	UCSC-C3160-BEZEL	1	\$ -	1	\$ -	
UCS C3X60 1050W Power Supply Unit	UCSC-C3X60-RAIL	1	\$ -	1	\$ -	
8 TB 12G SAS 7.2K RPM LFF HDD (4K)	UCSC-PSU1-1050W	1	\$ -	4	\$ -	
	UCS-HD8T7KL4K=	1	\$ 2,719.00	64	\$ 174,016.00	
UCS C240 M4 SFF 24 HD w/o CPU mem HD PCIe PS railkit w /expndr	UCSC-C240-M4SX	1	\$ 3,995.00	1	\$ 3,995.00	
2.60 GHz E5-2690 v3/135W 12C/30MB Cache/DDR4 2133MHz	UCS-CPU-E52690D	1	\$ 6,307.00	2	\$ 12,614.00	
16GB DDR4-2133-MHz RDIMM/PC4-17000/dual rank/x4/1.2v	UCS-MR-1X162RU-A	1	\$ 600.00	16	\$ 9,600.00	
1.8 TB 12G SAS 10K RPM SFF HDD (4K)	UCS-HD18TB10KS4K	1	\$ 1,827.00	2	\$ 3,654.00	
Cisco VIC 1387 Dual Port 40Gb QSFP CNA MLOM	UCSC-MLOM-C40Q-03	1	\$ 2,192.00	1	\$ 2,192.00	
Cisco 12Gbps Modular SAS HBA	UCSC-SAS12GHBA	1	\$ 656.00	1	\$ 656.00	
120 GB 2.5 inch Enterprise Value 6G SATA SSD (boot)	UCS-SD120GBKS4-EB	1	\$ 567.00	2	\$ 1,134.00	
Ball Bearing Rail Kit for C220 M4 and C240 M4 rack servers	UCSC-RAILB-M4	1	\$ 220.00	1	\$ 220.00	
1200W / 800W V2 AC Power Supply for 2U C-Series Servers	UCSC-PSU2V2-1200W	1	\$ 749.00	2	\$ 1,498.00	
Power Cord 125VAC 13A NEMA 5-15 Plug North America	CAB-9K12A-NA	1	\$ -	2	\$ -	
Right PCI Riser Bd (Riser 1) 2onbd SATA bootdrvs+ 2PCI slts	UCSC-PCI-1C-240M4	1	\$ 148.00	1	\$ 148.00	
UCS 6332 1RU F/No PSU/32 QSFP+	UCS-SP-FI6332	1	\$22,000.00	2	\$ 44,000.00	
3rd Gen FI Per port License to connect C-direct only	UCS-LIC-6300-40GC	1	\$ 1,388.00	20	\$ 27,760.00	
40GBASE-CR4 Passive Copper Cable, 3m	QSFP-H40G-CU3M-RF	1	\$ 150.00	34	\$ 5,100.00	
Cisco R42610 standard rack w /side panels	RACK-UCS2	1	\$ 3,429.00	1	\$ 3,429.00	
3YR 24X7X4 Support UCS S3260	CON-OSP-S3260BSE	1	\$ 7,872.00	8		\$ 62,976.00
3YR 24X7X4 Support UCS C240 M4 SFF	CON-OSP-C240V4SP	1	\$ 1,284.99	1		\$ 1,284.99
3YR 24x7x4 Support UCS6332	CON-OSPT-FI6332	1	\$ 3,330.25	2		\$ 6,660.50
Red Hat Enterprise Linux Server, 3Y 24x7	CON-ISV1-RH2SUG3A	1	\$ 2,397.00	17	\$ 40,749.00	
Cloudera Enterprise Edition	UCS-BD-CEBN-GD=	1	\$14,057.00	17	\$ 238,969.00	
<b>Total</b>					<b>1348615.6</b>	<b>70921.49</b>

(continued next page)



## Cisco UCS Integrated Infrastructure for Big Data

TPCx-HS Rev. 1.4.1  
TPC-Pricing Rev. 2.0.0

Report Date:  
November 11, 2016

Description	Part Number	Source	Unit Price	Qty	Extended Price	3 Year Maint. Price
(continued from previous page)						
				Total	\$1,348,615.60	\$ 70,921.49
Large Purchase Discount 1	61% for products and 35% for service	1			\$ (822,655.52)	\$(24,822.52)
Acer V206HQLAbd - LED monitor - 20" (Inc 2 spares)	UM.1V6AA.A02	2	\$ 79.99	3	\$ 239.97	
Logitech USB Corded Keyboard/Mouse Combo MK120 (Inc 2 spares)	920-002565	2	\$ 17.99	3	\$ 53.97	

Pricing: 1 = Cisco, 2 = CDW.com

<sup>(1)</sup> All discounts are based on US list prices and for similar quantities and configurations. The discounts are based on the overall specific components pricing from respective vendors in this single quotation. Discounts for similarly sized configurations will be similar to those quoted here, but may vary based on the components in the configuration.

**Audited by Doug Johnson of InfoSizing**

**Three-Year Cost of Ownership      \$572,353**

**HSph@300TB      11.34**

**\$/HSph@300TB      \$50,472.05**

Prices used in TPC benchmarks reflect the actual prices a customer would pay for a one-time purchase of the stated components. Individually negotiated discounts are not permitted. Special prices based on assumptions about past or future purchases are not permitted. All discounts reflect standard pricing policies for the listed components. For complete details, see the pricing sections of the TPC benchmark specifications. If you find that the stated prices are not available according to these terms, please inform at [pricing@tpc.org](mailto:pricing@tpc.org). Thank you.



## Cisco UCS Integrated Infrastructure for Big Data

TPCx-HS Rev. 1.4.1  
TPC-Pricing Rev. 2.0.0

Report Date:  
November 11, 2016

### Performance Run

Scale Factor 300TB

Run Start Time 2016-10-30 07:44:10

Run End Time 2016-10-31 10:10:17

Run Elapsed Time 95,169.000

HSGen Start Time 2016-10-30 07:44:11

HSGen End Time 2016-10-30 14:06:36

HSGen Elapsed Time 22,946.110

HSSort Start Time 2016-10-30 14:06:38

HSSort End Time 2016-10-31 07:48:08

HSSort Elapsed Time 63,690.014

HSValidate Start Time 2016-10-31 07:48:11

HSValidate End Time 2016-10-31 10:10:17

HSValidate Elapsed Time 8,528.144

### Repeatability Run

Scale Factor 300TB

Run Start Time 2016-10-29 05:11:50

Run End Time 2016-10-30 07:27:19

Run Elapsed Time 94,531.000

HSGen Start Time 2016-10-29 05:11:52

HSGen End Time 2016-10-29 11:31:05

HSGen Elapsed Time 22,755.357

HSSort Start Time 2016-10-29 11:31:08

HSSort End Time 2016-10-30 05:11:13

HSSort Elapsed Time 63,605.884

HSValidate Start Time 2016-10-30 05:11:16

HSValidate End Time 2016-10-30 07:27:19

HSValidate Elapsed Time 8,164.399



**Cisco UCS Integrated Infrastructure  
for Big Data**

TPCx-HS Rev. 1.4.1  
TPC-Pricing Rev. 2.0.0

Report Date:  
November 11, 2016

Run Report for Performance Run - Run 2

---

---

TPCx-HS Performance Metric (HSph@SF) Report

Test Run 1 Details	Total Time =	95169
	Total Size =	3000000000000
	Scale-Factor =	300

TPCx-HS Performance Metric (HSph@SF): 11.3482

---

---

Run Report for Repeatability Run - Run 1

---

---

TPCx-HS Performance Metric (HSph@SF) Report

Test Run 2 Details	Total Time =	94531
	Total Size =	3000000000000
	Scale-Factor =	300

TPCx-HS Performance Metric (HSph@SF): 11.4248

---

---

# Table of Contents

---

ABSTRACT .....	9
PREFACE .....	10
CLAUSE 1: GENERAL ITEMS .....	11
1.1 TEST SPONSOR .....	11
1.2 PARAMETER SETTINGS .....	11
1.3 CONFIGURATION DIAGRAMS .....	11
1.4 DATASET DISTRIBUTION .....	13
1.5 SOFTWARE COMPONENTS DISTRIBUTION .....	13
CLAUSE 2: WORKLOAD RELATED ITEMS .....	14
2.1 HARDWARE & SOFTWARE TUNABLE .....	14
2.2 RUN REPORT .....	14
2.3 BENCHMARK KIT IDENTIFICATION .....	15
2.4 BENCHMARK KIT CHANGES .....	15
CLAUSE 3: SUT RELATED ITEMS .....	16
3.1 DATA STORAGE RATIO .....	16
3.2 MEMORY RATIO .....	16
CLAUSE 4: SCALE FACTORS AND METRICS .....	17
4.1 HSGEN TIME .....	17
4.2 HSSORT TIME .....	17
4.3 HSVALIDATE TIME .....	17
4.4 HSDATA CHECK TIMES .....	17
4.5 PERFORMANCE & PRICE-PERFORMANCE .....	17
AUDITORS' INFORMATION AND ATTESTATION LETTER .....	18
SUPPORTING FILE INDEX .....	21
THIRD PARTY PRICE QUOTES .....	22



# Abstract

This document contains the methodology and results of the TPCExpress Benchmark™ HS (TPCx-HS) test conducted in conformance with the requirements of the TPCx-HS Standard Specification, Revision 1.4.1.

The test was conducted at a Scale Factor of 300TB with 8 Cisco UCS S3260 2-Node Servers running Cloudera Enterprise Edition on Red Hat Enterprise Linux Server Release 6.7.

## Measured Configuration

Company Name	Cluster Node	Virtualization	Operating System
Cisco Systems, Inc.	Cisco UCS S3260 2-Node Server	n/a	Red Hat Enterprise Linux Server Release 6.7

## TPC Express Benchmark© HS Metrics

Total System Cost	HSph@300TB	Price/Performance	Availability Date
572,353 USD	11.34	50,472.05 USD	November 11, 2016

# Preface

## TPC Express Benchmark™ HS Overview

*TPC Express Benchmark™ HS (TPCx-HS) was developed to provide an objective measure of hardware, operating system and commercial Apache Hadoop File System API compatible software distributions, and to provide the industry with verifiable performance, price-performance and availability metrics. The benchmark models a continuous system availability of 24 hours a day, 7 days a week.*

*Even though the modeled application is simple, the results are highly relevant to hardware and software dealing with Big Data systems in general. The TPCx-HS stresses both hardware and software including Hadoop runtime, Hadoop File-system API compatible systems and MapReduce layers. This workload can be used to assess a broad range of system topologies and implementation of Hadoop clusters. The TPCx-HS can be used to assess a broad range of system topologies and implementation methodologies in a technically rigorous and directly comparable and vendor-neutral manner.*

*The TPCx-HS kit is available from the TPC (See [www.tpc.org/tpcx-hs](http://www.tpc.org/tpcx-hs) for more information). Users must sign-up and agree to the TPCx-HS User Licensing Agreement (ULA) to download the kit. Re-distribution of the kit is prohibited. All related work (such as collaterals, papers, derivatives) must acknowledge the TPC and include TPCx-HS copyright. The TPCx-H Kit includes: TPCx-HS Specification document, TPCx-HS Users Guide documentation, shell scripts to set up the benchmark environment and Java code to execute the benchmark load.*

*The purpose of TPC benchmarks is to provide relevant, objective performance data to industry users. To achieve that purpose, TPC benchmark specifications require that benchmark tests be implemented with systems, products, technologies and pricing that:*

- *Are generally available to users;*
- *Are relevant to the market segment that the individual TPC benchmark models or represents (e.g., TPCx-HS models and represents Hadoop run-time and Hadoop File-system API compatible systems);*
- *Would plausibly be implemented by a significant number of users in the market segment the benchmark models or represents.*

*The use of new systems, products, technologies (hardware or software) and pricing is encouraged so long as they meet the requirements above. Specifically prohibited are benchmark systems, products, technologies or pricing (hereafter referred to as "implementations") whose primary purpose is performance optimization of TPC benchmark results without any corresponding applicability to real-world applications and environments. In other words, all "benchmark special" implementations that improve benchmark results but not real-world performance or pricing, are prohibited.*

*The rules for pricing are included in the TPC Pricing Specification and rules for energy measurement are included in the TPC Energy Specification.*

*Further information is available at [www.tpc.org](http://www.tpc.org)*

# Clause 1: General Items

## 1.1 Test Sponsor

*A statement identifying the benchmark sponsor(s) and other participating companies must be provided.*

This benchmark was sponsored by Cisco Systems, Inc.

## 1.2 Parameter Settings

*Settings must be provided for all customer-tunable parameters and options which have been changed from the defaults found in actual products, including but not limited to:*

- *Configuration parameters and options for server, storage, network and other hardware component incorporated into the pricing structure;*
- *Configuration parameters and options for operating system and file system component incorporated into the pricing structure;*
- *Configuration parameters and options for any other software component incorporated into the pricing structure;*
- *Compiler optimization options.*

*Comment 1: In the event that some parameters and options are set multiple times, it must be easily discernible by an interested reader when the parameter or option was modified and what new value it received each time.*

*Comment 2: This requirement can be satisfied by providing a full list of all parameters and options, as long as all those that have been modified from their default values have been clearly identified and these parameters and options are only set once.*

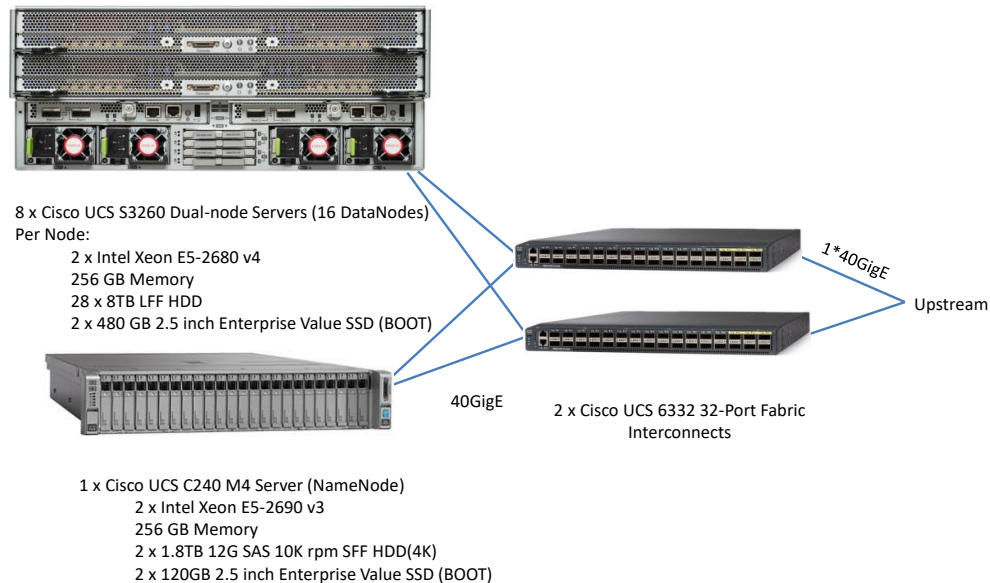
The supporting files contain the parameters and options used to configure the components involved in this benchmark.

## 1.3 Configuration Diagrams

*7.4.4 Diagrams of both measured and priced configurations must be provided, accompanied by a description of the differences. This includes, but is not limited to:*

- *Total number of nodes used;*
- *Total number and type of processors used/total number of cores used/total number of threads used (including sizes of L2 and L3 caches);*
- *Size of allocated memory, and any specific mapping/partitioning of memory unique to the test;*
- *Number and type of disk units (and controllers, if applicable);*
- *Number of channels or bus connections to disk units, including their protocol type;*
- *Number of LAN (e.g., Ethernet) connections and speed for switches and other hardware components physically used in the test or are incorporated into the pricing structure;*
- *Type and the run-time execution location of software components.*

## Measured Configuration



The measured configuration consisted of:

- Total Nodes: 17
- Total Processors/Cores/Threads: 34/472/944
- Total Memory: 4.25TB
- Total Number of Storage Drives/Devices: 484
- Total Storage Capacity: 3,603.2TB

Server nodes details:

- 8 x Cisco UCS S3260 2-Node Servers (DataNodes), 1 x Cisco UCS C240 M4 Server (NameNode) each node with:
  - Processors/Cores/Threads: 2/28 /56 (DataNodes); 2/24/48 (NameNode)
  - Processor Model:
    - 2 x Intel® Xeon® Processor E5-2680 v4, 2.40 GHz, 35 MB L3 (DataNodes)
    - 2 x Intel® Xeon® Processor E5-2690 v3, 2.60 GHz, 30 MB L3 (NameNode)
  - Memory: 256GB
  - Controller:
    - 1 x Cisco UCS C3000 RAID Controller M4 Server w 4G RAID Cache (DataNodes)
    - 1 x Cisco 12Gbps Modular SAS HBA (NameNode)
  - Drives:
    - 24 x 8TB LFF HDD (DataNodes)
    - 2 x 480GB 2.5in. Enterprise Value SSD (DataNodes, boot disk)
    - 2 x 1.8TB 12G SAS 10K rpm SFF HDD (NameNode)
    - 2 x 120GB 2.5 Enterprise Value SSD (NameNode, boot disk)
  - Network: Cisco VIC 1300

Network connectivity detail:

- 2 x Cisco UCS 6332 32-Port Fabric Interconnect

The distribution of software components over server nodes is detailed in section 1.5.

## Priced Configuration

There are no differences between the priced and measured configurations.

## 1.4 Dataset Distribution

*The distribution of dataset across all media must be explicitly described.*

Table 1.4 describes the distribution of the dataset across all media in the system.

**Table 1.4: Dataset Distribution**

Server Node	Controller	Disk Drive	Description of Content
1	Cisco 12Gbps Modular SAS HBA	1-2 (HDD)	Data, Temp
1	Intel Chipset Embedded SATA RAID	0 (2 SSD, RAID 1)	Operating System, Root, Swap, Hadoop Master
2-17	Cisco UCS C3000 RAID Controller M4 Server w 4G RAID Cache	1-28 (HDD)	Data, Temp
2-17	Cisco UCS C3000 RAID Controller M4 Server w 4G RAID Cache	0 (2 SSD, RAID-1)	Operating system, Root, Swap, Hadoop Master

## 1.5 Software Components Distribution

*The distribution of various software components across the system must be explicitly described.*

Table 1.5 describes the distribution of the software components across the system.

**Table 1.5: Software Component Distribution**

Node	Map/Reduce		HDFS		ZooKeeper
	Resource Manager	Node Manager	NameNode	DataNode	QuorumPeer
1	X		X		X
2-3		X		X	X
4-17		X		X	X

*Distributed file system implementation and corresponding Hadoop File System API version must be disclosed.*

Cloudera Enterprise Edition (fully HDFS compatible at the API level).

*Map/Reduce implementation and corresponding version must be disclosed.*

Cloudera Enterprise Edition (compatible equivalent to Hadoop 2.7.1).

# Clause 2: Workload Related Items

## 2.1 Hardware & Software Tunable

*Script or text used to set for all hardware and software tunable parameters must be reported.*

The Supporting File Archive contains all configuration scripts.

## 2.2 Run Report

*The run report generated by TPCx-HS benchmark kit must be reported.*

The Supporting File Archive contains the full run report. Following are extracts from the run report that lists the performance summary for both runs.

### Run Report for Run 1 - Repeatability Run

---

---

#### TPCx-HS Performance Metric (HSph@SF) Report

Test Run 1 Details	Total Time =	94531
	Total Size =	3000000000000
	Scale-Factor =	300

TPCx-HS Performance Metric (HSph@SF): 11.4248

---

---

### Run Report for Run 2 - Performance Run

---

---

#### TPCx-HS Performance Metric (HSph@SF) Report

Test Run 1 Details	Total Time =	95169
	Total Size =	3000000000000
	Scale-Factor =	300

TPCx-HS Performance Metric (HSph@SF): 11.3482

---

---

## 2.3 Benchmark Kit Identification

*Version number of TPCx-HS kit and checksum for HSGen, HSSort and HSValidate Programs must be reported.*

Kit Version	1.4.1
File	MD5
BigData_cluster_validate_suite.sh	58c13ddb98a2d1228f2df10f4a087a71
TPCx-HS-master.jar	a7310f65339708afe92af0029960a2cc
TPCx-HS-master.sh	f0d6c7361870983740ff65956a9cbe2e

## 2.4 Benchmark Kit changes

No modifications were made to the TPC-provided kit.

# Clause 3: SUT Related Items

## 3.1 Data Storage Ratio

The data storage ratio must be disclosed.

Table 3.1 describes the details of the storage devices configured on the system and their capacity.

Table 3.1: Storage Device Capacity

Qty	Capacity (GB)	Total (GB)
448	8,000	3,584,000
32	480	15,360
2	1,800	3,600
2	120	240
<b>Total Storage (TB)</b>		<b>3,603.2</b>

Scale Factor = 300TB

**Data Storage Ratio** = (Storage / SF) = **12.01**

## 3.2 Memory Ratio

The Scale Factor to memory ratio must be disclosed.

Total Configured Memory = 4.25TB

**Scale Factor to Memory Ratio** = (SF / Memory) = **70.59**



# Clause 4: Scale Factors and Metrics

## 4.1 HSGen Time

The HSGen time must be disclosed for Run1 and Run2.

	Run1	Run2
HSGen	22,755.357	22,946.110

## 4.2 HSSort Time

The HSSort time must be disclosed for Run1 and Run2.

	Run1	Run2
HSSort	63,605.884	63,690.014

## 4.3 HSValidate Time

The HSValidate time must be disclosed for Run1 and Run2.

	Run1	Run2
HSValidate	8,164.399	8,528.144

## 4.4 HSDataCheck Times

Both HSDataCheck times must be disclosed for Run1 and Run2.

	Run1	Run2
HSDataCheck (pre-Sort)	3.000	2.000
HSDataCheck (post-Sort)	3.000	3.000

## 4.5 Performance & Price-Performance

The performance metric (HSph@SF) must be disclosed for Run1 and Run2. Price-performance metric (\$/HSph@SF) must be disclosed for the performance run.

	Run1	Run2
HSph@300TB	11.42	11.34

\$/HSph@300TB	50,472.05 USD
---------------	---------------

# Auditors' Information and Attestation Letter

*The auditor's agency name, address, phone number, and Attestation letter must be included in the full disclosure report. A statement should be included specifying who to contact in order to obtain further information regarding the audit process.*

This benchmark was audited by Doug Johnson for InfoSizing

www.sizing.com  
20 Kreg Lane  
Manitou Springs, CO 80829  
719-473-7555.

This benchmark's Full Disclosure Report (FDR) can be downloaded from [www.tpc.org](http://www.tpc.org).

A copy of the auditor's attestation letter is included in the next two pages.

Raghunath Nambiar  
Cisco Systems Inc.  
3800 Zanker Road  
San Jose, CA 95134

November 11, 2016

I verified the TPC Express Benchmark™ HS v1.4.1 performance of the following configuration:

Platform: Cisco UCS Integrated Infrastructure for Big Data  
(with 8 Cisco UCS S3260 Dual-node Servers)  
Operating System: Red Hat Enterprise Linux Server 6.7  
Apache Hadoop Cloudera Enterprise Edition  
Compatible Software:

The results were:

**Performance Metric 11.34 HSph@30TB**

Run Elapsed Time 95,169.00 Seconds

**Cluster** **8 Cisco UCS S3260 Dual-node Servers (Data Nodes),**  
**1 Cisco UCS C240 M4 Server (Name Node)**  
**each node with:**

CPU	(Data Nodes)	2 x Intel Xeon Processor E5-2680 v4 (2.40 GHz, 14-core, 35 MB L3)
	(Name Node)	2 x Intel Xeon Processor E5-2690 v3 (2.60 GHz, 12-core, 30 MB L3)
Memory		256 GB
Storage		
		<b>Qty Size Type</b>
		24 8 TB LFF HDD (Data Nodes)
		2 480GB SSD (Data Nodes, boot disk)
		2 1.8TB 12G SAS 10K rpm SFF HDD (Name Node)
		2 120GB SSD (Name Node, boot disk)

In my opinion, these performance results were produced in compliance with the TPC requirements for the benchmark.

The following verification items were given special attention:

- All TPC-provided components were verified to be v1.4.1
- No modifications were made to any of the Java code
- Any and all modifications to shell scripts were reviewed for compliance

- All checksums were validated for compliance
- The generated dataset was properly scaled to 300TB
- The generated dataset and the sorted dataset were replicated 3-ways
- The elapsed times for all phases and runs were correctly measured and reported
- The Storage and Memory Ratios were correctly calculated and reported
- The system pricing was verified for major components and maintenance
- The major pages from the FDR were verified for accuracy

Additional Audit Notes:

None.

Respectfully Yours,



Doug Johnson, Auditor



François Raab, President

# Supporting File Index

The following index outlines the information included in the supporting files archive.

Clause	Description	Archive File Pathname
Clause 1	Parameters and options used to configure the system	SupportingFilesArchive\Clause1
Clause 2	Configuration scripts & Run report	SupportingFilesArchive\Clause2
Clause 3	System configuration details	SupportingFilesArchive\Clause3

# Third Party Price Quotes



Home > Computer Accessories > Keyboards & Key pads > Keyboard and Mouse Bundle



## Logitech MK120 USB Wired Keyboard/Mouse Set – \$4.50 svgs while supplies last

**logitech** Mfg. Part: 920-002565 | CDW Part: 2124292 | UNSPSC: 43211706

★★★★★ [Read 1 review](#) | [Write a review](#)

Availability: **In Stock** Ships today if ordered within 2 hrs 20 mins

1

~~\$19.99~~

**\$17.99**

Advertised Price

[Add to Cart](#)



- Keyboard
- mouse
- Wired
- PC compatible
- USB interface
- standard version

[View More](#)

[Log On to Email this page or Save as Favorite](#)



Home > Monitors & Projectors > LCD / LED Monitors



## Acer V206HQL – LED monitor – 20"

**acer** Mfg. Part: UM.IV6AA.A02 | CDW Part: 3051875 | UNSPSC: 43211902

Availability: **In Stock** Ships today if ordered within 2 hrs 19 mins

1

~~\$105.00~~

**\$79.99**

Advertised Price

[Add to Cart](#)

[View All Warranties and Services](#)



- LED monitor
- 20" ( 19.5" viewable )
- 1600 x 900
- TN
- 200 cd/m2
- 5 ms
- DVI

[View More](#)

### Recommended Warranty and Services

[Acer Two-way freight for depot – extended service agreement – 3 years – pic](#)

0

**\$10.99**

Advertised Price