



Cisco Systems, Inc.

TPC Express Benchmark™ HS (TPCx-HS)

Full Disclosure Report

for

Cisco UCS Integrated Infrastructure for Big Data and Analytics

(with 16 Cisco UCS C240M4 Servers)

using

MapR Converged Community Edition, Version 5.0

and

Red Hat Enterprise Linux Server 6.7

First Edition

March 31, 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. 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. 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 and Analytics

TPCx-HS Rev. 1.3.0
TPC-Pricing Rev. 2.1.0

Report Date:
March 31, 2016

Total System Cost

TPCx-HS Performance Metric

Price/Performance

386,270 USD

10.12
HSph@1TB

38,168.98 USD
\$/HSph@1TB

Scale Factor

Apache Hadoop
Compatible Software

Operating System

Other Software

Availability Date

1TB

MapR Converged
Community Edition,
Version 5.0

Red Hat Enterprise Linux
Server 6.7

None

March 31, 2016

System Configuration



16 x Cisco UCS C240 M4 Servers (All DataNodes)
with 24 x 1.2TB 6G SAS 10K rpm SFF HDD and
2 x 240GB 2.5 inch Enterprise Value 6G SATA SSD (BOOT)



2 x Cisco UCS 6296UP
96-Port Fabric Interconnect

Physical Storage/Scale Factor: 468.48

Scale Factor/Physical Memory: 0.25

Total Processors/Cores/Threads Servers: 16 x Cisco UCS C240M4 Server
32/448/896

Server Configuration: Per node:
Processors 2 x Intel® Xeon® Processor E5-2680 v4, 2.40 GHz, 35MB L3
Memory 256GB
Storage Controller 1 x Cisco 12G SAS Modular Raid Controller, 2GB FBWC Cache
Storage Device 24 x 1.2TB 6G SAS 10Krpm SFF HDD
2 x 240GB 2.5in. Enterprise Value 6G SATA SSD (boot disks)
Network 1 x Cisco UCS VIC 1227 MLOM - Dual Port 10Gb SFP+

Connectivity: 2 x Cisco UCS 6296UP 96-Port Fabric Interconnect



Cisco UCS Integrated Infrastructure for Big Data and Analytics

TPCx-HS Rev. 1.3.0
TPC-Pricing Rev. 2.1.0

Report Date:
March 31, 2016

Description	Part Number	Source	Unit Price	Qty	Extended Price	3 Year Maint. Price
Hardware Components						
Cisco UCS Integrated Infrastructure for Big Data and Analytics Gen 4	UCS-SL-CPA4-P1	1	\$704,776.00	1	\$ 704,776.00	
(Not sold Standalone)UCS C240 M4SX w/E52680v4,256GB mem	UCS-SA-C240M4SX-P	1	\$ -	16	\$ -	
2.40 GHz E5-2680 v4/120W 14C/35MB Cache/DDR4 2400MHz	UCS-CPU-E52680E	1	\$ -	32	\$ -	
16GB DDR4-2400-MHz RDIMM/PC4-19200/single rank/x4/1.2v	UCS-MR-1X161RV-A	1	\$ -	256	\$ -	
240 GB 2.5 inch Enterprise Value 6G SATA SSD (boot)	UCS-SD240GBKS4-EB	1	\$ -	32	\$ -	
1.2 TB 6G SAS 10K rpm SFF HDD	UCS-HD12T10KS2-E	1	\$ -	384	\$ -	
Cisco 12Gbps SAS 2GB FBWC Cache module (Raid 0/1/5/6)	UCSC-MRAID12G-2GB	1	\$ -	16	\$ -	
Cisco 12G SAS Modular Raid Controller	UCSC-MRAID12G	1	\$ -	16	\$ -	
Cisco UCS VIC1227 VIC MLOM - Dual Port 10Gb SFP+	UCSC-MLOM-CSC-02	1	\$ -	16	\$ -	
Right PCI Riser Bd (Riser 1) 2onbd SATA bootdrvs+ 2PCI slts	UCSC-PCI-1C-240M4	1	\$ -	16	\$ -	
1200W / 800W V2 AC Power Supply for 2U C-Series Servers	UCSC-PSU2V2-1200W	1	\$ -	32	\$ -	
Power Cord 125VAC 13A NEMA 5-15 Plug North America	CAB-9K12A-NA	1	\$ -	32	\$ -	
Right PCI Riser Bd (Riser 1) 2onbd SATA bootdrvs+ 2PCI slts	UCSC-PCI-1C-240M4	1	\$ -	16	\$ -	
Heat sink for UCS C240 M4 rack servers	UCSC-HS-C240M4	1	\$ -	32	\$ -	
Supercap cable 250mm	UCSC-SCCBL240	1	\$ -	16	\$ -	
Ball Bearing Rail Kit for C220 M4 and C240 M4 rack servers	UCSC-RAILB-M4	1	\$ -	16	\$ -	
(Not sold standalone) UCS 6296UP 2RU Fibre/18p LIC,16xCables	UCS-SA-BD-FI96	1	\$ -	2	\$ -	
UCS 6296UP Power Supply/100-240VAC	UCS-PSU-6296UP-AC	1	\$ -	4	\$ -	
10GBASE-CU SFP+ Cable 3 Meter	SFP-H10GB-CU3M	1	\$ -	32	\$ -	
UCS Manager v2.2	N10-MGT012	1	\$ -	2	\$ -	
UCS 6296UP Chassis Accessory Kit	UCS-ACC-6296UP	1	\$ -	2	\$ -	
UCS 6200 Series Expansion Module Blank	UCS-BLKE-6200	1	\$ -	6	\$ -	
UCS 6296UP Fan Module	UCS-FAN-6296UP	1	\$ -	8	\$ -	
Power Cord 125VAC 13A NEMA 5-15 Plug North America	CAB-9K12A-NA	1	\$ -	4	\$ -	
UCSD Express for Big Data	CUIC-EBDS-LIC	1	\$ -	1	\$ -	
Cisco Smart Net 24X7X4 3Y UCS C240 M4S BD SP Server	CON-OSP-C240V4SP	1	\$ 1,284.99	16		\$20,560
Cisco R42610 standard rack w/side panels	RACK-UCS2	1	\$ 3,429.00	1	\$ 3,429.00	
Cisco Smart Net 24X7X4 3Y UCS 6296UP 2RU Fabric Int/2 PSU/4 Fans	CON-SNTP-FI6296UP	1	\$ 5,781.00	2		\$11,562



Cisco UCS Integrated Infrastructure for Big Data and Analytics

TPCx-HS Rev. 1.3.0
TPC-Pricing Rev. 2.1.0

Report Date:
March 31, 2016

Description	Part Number	Source	Unit Price	Qty	Extended Price	3 Year Maint. Price	
Software Components							
Red Hat Enterprise Linux Server, 3Y 24x7	CON-ISV1-RH2SUG3A	1	\$ 2,397.00	16	\$ 38,352.00	Inc.	
Cisco MapR Community Edi. Hadoop 3Y 24x7	UCS-BD-M3-AS-3YR	1	\$ 11,842.00	16	\$ 189,472.00	Inc.	
					Total	\$936,029.00	\$32,122
Large Purchase Discount ¹	61% for products and 35% for service		1		-\$570,977.69	-\$11,243	
Acer V206HQLAbd - LED monitor - 20" (Inc 2 spares)	UM.1V6AA.A02	2	\$ 94.99	3	\$ 284.97		
Logitech USB Corded Keyboard/Mouse Combo MK120 (Inc 2 spares)	920-002565	2	\$ 17.99	3	\$ 53.97		

Pricing: 1 = Cisco, 2 = CDW.com ⁽¹⁾ 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.	Three-Year Cost of Ownership \$386,270 HSph@1TB 10.12 \$/HSph@1TB \$38,168.98
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------

Audited by Doug Johnson of InfoSizing, Inc.

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. Thank you.



Cisco UCS Integrated Infrastructure for Big Data and Analytics

TPCx-HS Rev. 1.3.0
TPC-Pricing Rev. 2.1.0

Report Date:
March 31, 2016

Performance Run

Scale Factor 1TB

Run Start Time 2016-03-30 02:56:57

Run End Time 2016-03-30 03:02:52

Run Elapsed Time 356.000

HSGen Start Time 2016-03-30 02:56:58

HSGen End Time 2016-03-30 02:58:30

HSGen Elapsed Time 93.468

HSSort Start Time 2016-03-30 02:58:32

HSSort End Time 2016-03-30 03:02:13

HSSort Elapsed Time 221.790

HSValidate Start Time 2016-03-30 03:02:15

HSValidate End Time 2016-03-30 03:02:52

HSValidate Elapsed Time 38.175

Repeatability Run

Scale Factor 1TB

Run Start Time 2016-03-30 02:42:40

Run End Time 2016-03-30 02:48:34

Run Elapsed Time 355.000

HSGen Start Time 2016-03-30 02:42:41

HSGen End Time 2016-03-30 02:44:15

HSGen Elapsed Time 95.818

HSSort Start Time 2016-03-30 02:44:17

HSSort End Time 2016-03-30 02:47:55

HSSort Elapsed Time 218.533

HSValidate Start Time 2016-03-30 02:47:57

HSValidate End Time 2016-03-30 02:48:34

HSValidate Elapsed Time 38.057



**Cisco UCS Integrated Infrastructure
for Big Data and Analytics**

TPCx-HS Rev. 1.3.0
TPC-Pricing Rev. 2.1.0

Report Date:
March 31, 2016

Run Report for Performance Run - Run 2

TPCx-HS Performance Metric (HSph@1TB) Report

Test Run 2 Details	Total Time =	356
	Total Size =	10,000,000,000
	Scale-Factor =	1

TPCx-HS Performance Metric (HSph@1TB): 10.1214

Run Report for Repeatability Run - Run 1

TPCx-HS Performance Metric (HSph@1TB) Report

Test Run 1 Details	Total Time =	355
	Total Size =	10,000,000,000
	Scale-Factor =	1

TPCx-HS Performance Metric (HSph@1TB): 10.1419

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	14
2.4 BENCHMARK KIT CHANGES	14
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 TPC Express Benchmark™ HS (TPCx-HS) test conducted in conformance with the requirements of the TPCx-HS Standard Specification, Revision 1.3.0.

The test was conducted at a Scale Factor of 1TB with 16 Cisco UCS C240M4 Servers running MapR Converged Community Edition, Version 5.0 on Red Hat Enterprise Linux Server 6.7.

Measured Configuration

Company Name	Cluster Node	Virtualization	Operating System
Cisco Systems, Inc.	Cisco UCS C240M4 Server	n/a	Red Hat Enterprise Linux Server 6.7

TPC Express Benchmark©HS Metrics

Total System Cost	HSph @1TB	Price/Performance	Availability Date
386,270 USD	10.12	38,168.98 USD	March 31, 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 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

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.

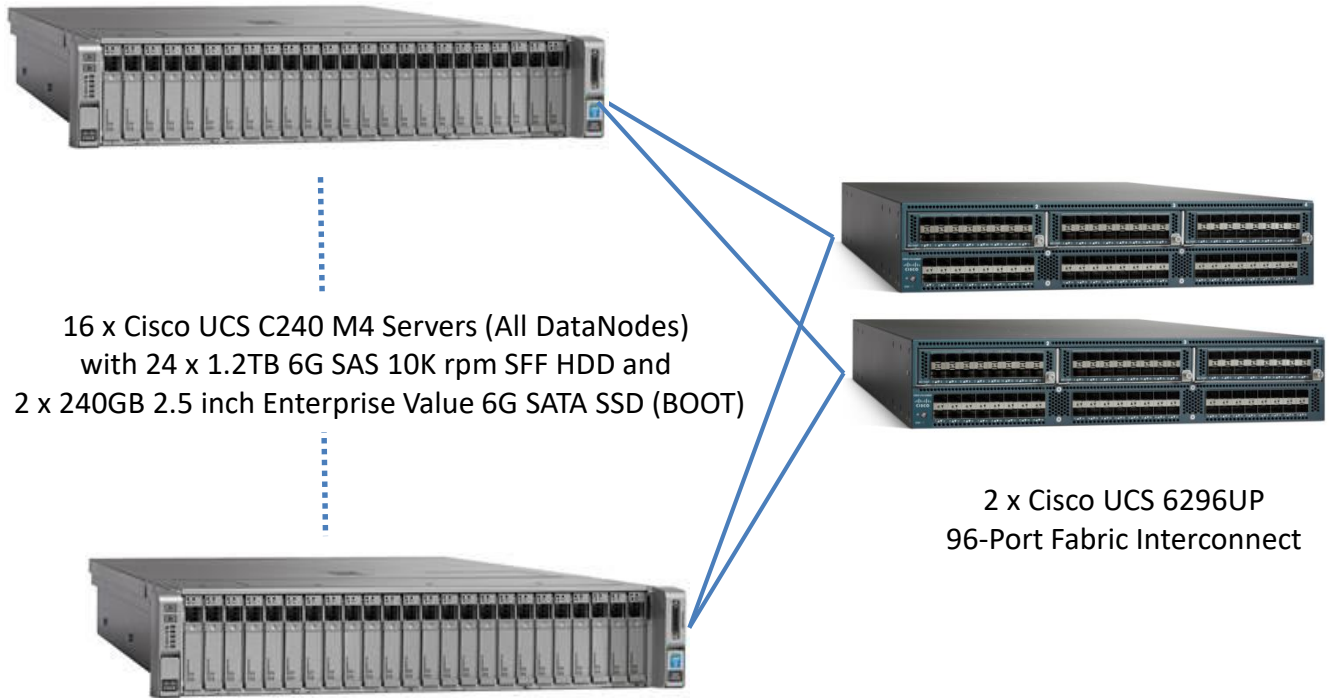
The supporting files contain the parameters and options used to configure the components involved in this benchmark.

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: 16
- Total Processors/Cores/Threads: 32/448/896
- Total Memory: 4TB
- Total Number of Storage Drives/Devices: 416
- Total Storage Capacity: 468.48TB

Server nodes details:

- 16 x Cisco UCS C240M4 Servers, each with:
 - Processors/Cores/Threads: 2/28/56
 - Processor Model: 2 x Intel® Xeon® Processor E5-2680 v4, 2.40 GHz, 35MB L3
 - Memory: 256GB
 - Controller: 1 x Cisco 12G SAS Modular Raid Controller, 2GB FBWC Cache
 - Drives:
 - 24 x 1.2TB 6G SAS 10Krpm SFF HDD
 - 2 x 240GB 2.5in. Enterprise Value 6G SATA SSD (boot disks)
 - Network: 1 x Cisco UCS VIC 1227 MLOM - Dual Port 10Gb SFP+

Network connectivity detail:

- 2 x Cisco UCS 6296UP 96-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
1-16	Cisco 12-Gbps SAS Modular Raid Controller with 2-GB FBWC	1-24 (HDD)	Data, Temp
1-16	Intel Chipset Embedded SATA RAID	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

	Map/Reduce		HDFS		ZooKeeper
Node	JobTracker	TaskTracker	NameNode	DataNode	QuorumPeer
1-16	3	1-16	1	1-16	1-3

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

MapR Converged Community Edition, Version 5.0 (fully HDFS compatible at the API level).

Map/Reduce implementation and corresponding version must be disclosed.

MapR Converged Community Edition, Version 5.0 (compatible equivalent to Hadoop 2.7.0).

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.

- **Run1 Performance Summary**

```
=====
TPCx-HS Performance Metric (HSph@1TB) Report

Test Run 1 Details      Total Time =                355
                        Total Size =           10,000,000,000
                        Scale-Factor =                1

TPCx-HS Performance Metric (HSph@1TB):           10.1419
=====
```

- **Run2 Performance Summary**

```
=====
TPCx-HS Performance Metric (HSph@1TB) Report

Test Run 2 Details      Total Time =                356
                        Total Size =           10,000,000,000
                        Scale-Factor =                1

TPCx-HS Performance Metric (HSph@1TB):           10.1214
=====
```

2.3 Benchmark Kit Identification

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

The version number of the TPCx-HS kit used is 1.3.0. The md5sum for the TPCx-HS kit files used during the benchmark are:

```
TPCx-HS-master.jar 4ceaefc51c698c0733b57244b7760808
BigData_cluster_validate_suite.sh 58c13ddb98a2d1228f2df10f4a087a71
TPCx-HS-master.sh 0df4c374f1671ed15ff0dc9352eb485e
```

2.4 Benchmark Kit changes

The file TPCx-HS-master.sh was modified. The diff output follows.

219,220c219

```
< #hdfs dfs -ls /user/"$HADOOP_USER"/TPCx-HS-benchmark/HSsort-input/* | tee -a ./TPCx-HS-result-"$prefix".log  
< hadoop dfs -ls /user/"$HADOOP_USER"/TPCx-HS-benchmark/HSsort-input/* | tee -a ./TPCx-HS-result-"$prefix".log  
---  
> hdfs dfs -ls /user/"$HADOOP_USER"/TPCx-HS-benchmark/HSsort-input/* | tee -a ./TPCx-HS-result-"$prefix".log
```

249,250c248

```
< #hdfs dfs -ls /user/"$HADOOP_USER"/TPCx-HS-benchmark/HSsort-output/* | tee -a ./TPCx-HS-result-"$prefix".log  
< hadoop dfs -ls /user/"$HADOOP_USER"/TPCx-HS-benchmark/HSsort-output/* | tee -a ./TPCx-HS-result-"$prefix".log  
---  
> hdfs dfs -ls /user/"$HADOOP_USER"/TPCx-HS-benchmark/HSsort-output/* | tee -a ./TPCx-HS-result-"$prefix".log
```

279,280c277

```
< #hdfs dfs -ls /user/"$HADOOP_USER"/TPCx-HS-benchmark/HSValidate/* | tee -a ./TPCx-HS-result-"$prefix".log  
< hadoop dfs -ls /user/"$HADOOP_USER"/TPCx-HS-benchmark/HSValidate/* | tee -a ./TPCx-HS-result-"$prefix".log  
---  
> hdfs dfs -ls /user/"$HADOOP_USER"/TPCx-HS-benchmark/HSValidate/* | tee -a ./TPCx-HS-result-"$prefix".log
```

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)
384	1,200	460,800
32	240	7,680
Total Storage (TB)		468.48

Scale Factor = 1TB

Data Storage Ratio = (Storage / SF) = **468.48**

3.2 Memory Ratio

The Scale Factor to memory ratio must be disclosed.

Total Configured Memory = 4TB

Scale Factor to Memory Ratio = (SF / Memory) = **0.25**

Clause 4: Scale Factors and Metrics

4.1 HSGen Time

The HSGen time must be disclosed for Run1 and Run2.

	Run1	Run2
HSGen	95.818	93.468

4.2 HSSort Time

The HSSort time must be disclosed for Run1 and Run2.

	Run1	Run2
HSSort	218.533	221.790

4.3 HSValidate Time

The HSValidate time must be disclosed for Run1 and Run2.

	Run1	Run2
HSValidate	38.057	38.175

4.4 HSDataCheck Times

Both HSDataCheck times must be disclosed for Run1 and Run2.

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

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@1TB	10.14	10.12

\$/HSph@1TB	38,168.98 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, Inc.

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.

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

March 29, 2016

I verified the TPC Express Benchmark™ HS v1.3.0 performance of the following configuration:

Platform: Cisco UCS Integrated Infrastructure for Big Data and Analytics
(with 16 Cisco UCS C240M4 Servers)
Operating System: Red Hat Enterprise Linux Server 6.7
Apache Hadoop MapR Converged Community Edition, Version 5.0
Compatible Software:

The results were:

Performance Metric **10.12 HSph@1TB**
Run Elapsed Time 356.00 Seconds

Cluster	16 Cisco UCS C240M4 Servers		
CPU	2 x Intel Xeon Processor E5-2680 v4 (2.40 GHz, 14-core, 35 MB L3)		
Memory	256 GB		
Storage	<u>Qty</u>	<u>Size</u>	<u>Type</u>
	24	1.2 TB	10K rpm SFF HDD
	2	240GB	SATA SSD (boot disks)

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 v1.3.0
- 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 1TB
- 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

Logitech MK120 USB Wired Keyboard & Mouse Set

CDW LLC [US] <https://www.cdw.com/shop/products/Logitech-MK120-USB-Wired-Keyboard-Mouse-Set/2124292.aspx?pfm=srh>

Account Log On or Create Account

CDW

Search CDW...

HARDWARE SOFTWARE SOLUTIONS CLOUD BRANDS BLOG DEALS


Home > Computer Accessories > Keyboards & Keypads > Keyboard and Mouse Bundle

Logitech MK120 USB Wired Keyboard & Mouse Set

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

★★★★★

Read 1 review | Write a review



Availability: **In Stock**

Ships same day if ordered before 4 PM CT

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

View More

Add All Items To Cart

Log On to Email this page or Save as Favorite

Acer V206HQL - LED monitor - 20"

CDW LLC [US] <https://www.cdw.com/shop/products/Acer-V206HQL-LED-monitor-20in/3051875.aspx?pfm=srh>


Home > Monitors & Projectors > LCD / LED Monitors

Acer V206HQL - LED monitor - 20"

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

★★★★★

Read all 47 reviews | Write a review



Availability: **In Stock**

Ships same day if ordered before 4 PM CT

- LED monitor
- 20"
- 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 - plc

0 \$10.99 Advertisised Price