



Cisco Systems, Inc.

TPC Express Benchmark™ HS (TPCx-HS)

Full Disclosure Report

for

Cisco UCS Integrated Infrastructure for Big Data
(Cisco UCS CPA v2)

using

Red Hat Enterprise Linux Server 6.4

and

MapR M5 Edition

First Edition

January 8, 2015

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 World Wide Web server at <http://www.cisco.com>. Actual performance and environmental costs of Cisco products will vary depending on individual customer configurations and conditions.

Copyright © 2015 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.



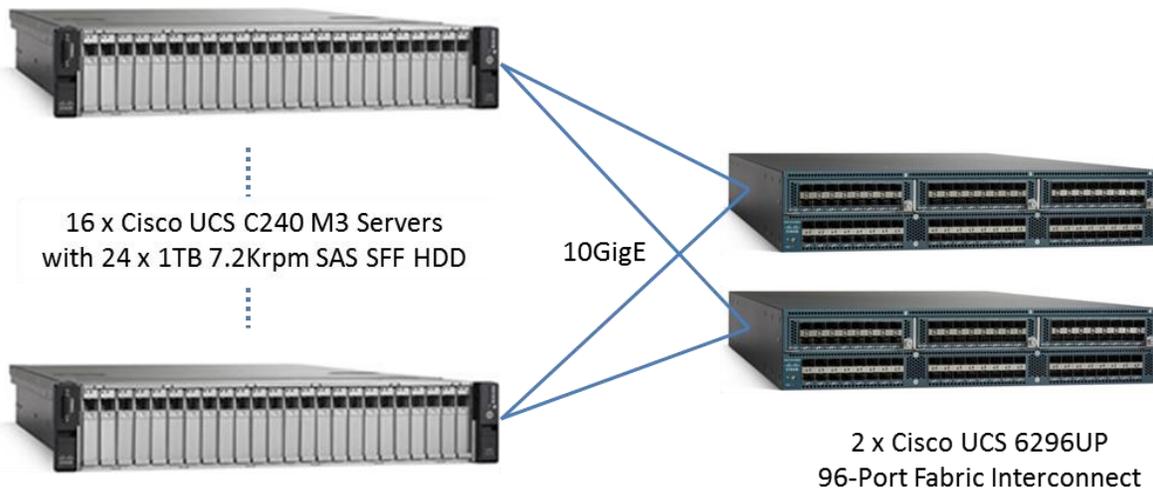
**Cisco UCS Integrated Infrastructure
for Big Data (Cisco UCS CPA v2)
with 16 Cisco UCS C240M3 Servers**

TPCx-HS Rev. 1.2.0
TPC-Pricing Rev. 1.7.0

Report Date:
January 8, 2015

Total System Cost		TPCx-HS Performance Metric		Price/Performance	
614,645 USD		5.77 HSph@10TB		106,524.27 USD \$/HSph@10TB	
Scale Factor	Apache Hadoop Compatible Software	Operating System	Other Software	Availability Date	
10TB	MapR M5 Edition	Red Hat Enterprise Linux Server 6.4	None	January 8, 2015	

System Configuration



Physical Storage/Scale Factor: 38.40

Scale Factor/Physical Memory: 2.50

Servers: 16 x Cisco UCS C240 M3 Server
 Total Processors/Cores/Threads 32/320/640
 Server Configuration:
 Processors 2 x Intel® Xeon® Processor E5-2660 v2, 2.20 GHz, 25M L3
 Memory 256GB
 Storage Controller 1 x MegaRAID 9271CV
 Storage Device 24 x 1TB 7.2K SAS SFF HDD (internal)
 Network 1 x Cisco VIC 1225 Dual Port 10Gb SFP+ CNA
 Connectivity: 2 x Cisco UCS 6296UP 96-Port Fabric Interconnect



**Cisco UCS Integrated Infrastructure
for Big Data (Cisco UCS CPA v2)
with 16 Cisco UCS C240M3 Servers**

TPCx-HS Rev. 1.2.0
TPC-Pricing Rev. 1.7.0

Report Date:
January 8, 2015

Description	Part Number	Source	Unit Price	Qty	Extended Price	3 Year Maint. Price
UCS C240 M3 SFF w/o CPU mem HD PCIe w/ rail kit expdr	UCSC-C240-M3S	1	3,600.00	16	\$57,600	
2.20 GHz E5-2660 v2/95W 10C/25MB Cache/DDR3 1866MHz	CON-SNTP-C240M3SF	1	3,949.00	32	\$126,368	
16GB DDR3-1866-MHz RDIMM/PC3-14900/dual rank/x4/1.5v	UCS-CPU-E52660B	1	741.00	256	\$189,696	
1TB SAS 7.2K RPM 2.5 inch HDD/hot plug/drive sled mounted	UCS-MR-1X162RZ-A	1	975.00	384	\$374,400	
Cisco VIC 1225 Dual Port 10Gb SFP+ CNA	UCS-HD1T7KS2-E	1	1,594.00	16	\$25,504	
1200W 2u Power Supply For UCS	UCSC-PCIE-CSC-02	1	652.00	32	\$20,864	
Power Cord 125VAC 13A NEMA 5-15 Plug North America	UCSC-PSU2-1200	1	0.00	32	\$0	
Full height PCIe filler for C-Series	CAB-9K12A-NA	1	0.00	48	\$0	
Heat Sink for UCS C240 M3 Rack Server	UCSC-PCIF-01F	1	0.00	32	\$0	
2U Rail Kit for UCS C-Series servers	UCSC-HS-C240M3	1	0.00	16	\$0	
Cisco M3 - v2 CPU asset tab ID label (Auto-Expand)	UCSC-RAIL-2U	1	0.00	16	\$0	
MegaRAID 9271CV with 8 internal SAS/SATA ports with Supercap	UCS-M3-V2-LBL	1	1,686.00	16	\$26,976	
SMARTNET 24X7X4 UCS C240 M3 SFF w/o	UCS-RAID9271CV-8I	1	1,521.00	16		\$24,336
UCS 6296UP 2RU Fabric Int/No PSU/48 UP/ 18p LIC	UCS-FI-6296UP-UPG	1	48,000.00	2	\$96,000	
10GBASE-CU SFP+ Cable 3 Meter	CON-SNTP-FI6296UP	1	210.00	32	\$6,720	
UCS 6296UP Chassis Accessory Kit	SFP-H10GB-CU3M	1	0.00	2	\$0	
UCS 6296UP Power Supply/100-240VAC	UCS-ACC-6296UP	1	2,000.00	4	\$8,000	
UCS Manager v2.2	UCS-PSU-6296UP-AC	1	0.00	2	\$0	
UCS 6200 Series Expansion Module Blank	N10-MGT012	1	0.00	6	\$0	
UCS 6296UP Fan Module	UCS-BLKE-6200	1	0.00	8	\$0	
Power Cord 125VAC 13A NEMA 5-15 Plug North America	UCS-FAN-6296UP	1	0.00	4	\$0	
SMARTNET 24X7X4 UCS 6296UP 2RU Fabric Int/2 PSU/4 Fans	CAB-9K12A-NA	1	5,505.00	2		\$11,010
Red Hat Enterprise Linux Server, 3Y 24x7	CON-ISV1-RH2SUG3A	1	2,397.00	16	\$38,352	Inc.
MapR M5 Edition, 3Y 24x7	UCS-BD-M5SL-3YR	1	25,083.00	16	\$401,328	Inc.
Cisco R42610 standard rack w/side panels	RACK-UCS2	1	3,429	1	\$3,429	
Total:					\$1,375,237	\$35,346
Large Purchase Discount		1			-\$783,885	-\$12,371
Acer - LED monitor - 18.5" (Inc 2 spares)	UM.XV6AA.A01	2	\$88	3	\$264	
Logitech USB KB/Mouse (Inc 2 spares)	920-002565	2	\$18	3	\$54	

1 = Cisco, 2= CDW.com

Three-Year Cost of Ownership: \$614,645

Discounts: 57% for products and 35% for service. All discounts are based on US list prices and for similar quantities and configurations. The discounts are based on the overall specific

HSph@10TB 5.77

Audited by Doug Johnson, InfoSizing Inc.

\$/HSph@10TB \$106,524.27

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 (Cisco UCS CPA v2)
with 16 Cisco UCS C240M3 Servers**

TPCx-HS Rev. 1.2.0
TPC-Pricing Rev. 1.7.0

Report Date:
January 8, 2015

Measurement Results for Performance Run

Scale Factor	10TB
Run Start Time	2014-12-19 19:09:05
Run End Time	2014-12-19 20:52:56
Run Elapsed Time	6,234.000
HSGen Start Time	2014-12-19 19:09:05
HSGen End Time	2014-12-19 19:22:38
HSGen Elapsed Time	813.984
HSSort Start Time	2014-12-19 19:22:44
HSSort End Time	2014-12-19 20:39:57
HSSort Elapsed Time	4,635.093
HSValidate Start Time	2014-12-19 20:40:01
HSValidate End Time	2014-12-19 20:52:56
HSValidate Elapsed Time	775.866

Measurement Results for Repeatability Run

Run Start Time	2014-12-19 17:35:06
Run End Time	2014-12-19 19:07:54
Run Elapsed Time	5,570.000
HSGen Start Time	2014-12-19 17:35:06
HSGen End Time	2014-12-19 17:42:51
HSGen Elapsed Time	465.887
HSSort Start Time	2014-12-19 17:42:56
HSSort End Time	2014-12-19 18:55:01
HSSort Elapsed Time	4,325.717
HSValidate Start Time	2014-12-19 18:55:05
HSValidate End Time	2014-12-19 19:07:54
HSValidate Elapsed Time	770.269



Run Report for Performance Run

```
=====  
TPCx-HS Performance Metric (HSph@SF) Report  
  
Test Run 2 details: Total Time = 6234  
                    Total Size = 100000000000  
                    Scale-Factor = 10.0000  
  
TPCx-HS Performance Metric (HSph@SF): 5.7750  
=====
```

Run Report for Repeatability Run

```
=====  
TPCx-HS Performance Metric (HSph@SF) Report  
  
Test Run 1 details: Total Time = 5570  
                    Total Size = 100000000000  
                    Scale-Factor = 10.0000  
  
TPCx-HS Performance Metric (HSph@SF): 6.4632  
=====
```

Table of Contents

ABSTRACT	8
PREFACE	9
CLAUSE 1: GENERAL ITEMS.....	10
1.1 TEST SPONSOR.....	10
1.2 PARAMETER SETTINGS	10
1.3 CONFIGURATION DIAGRAMS.....	10
1.4 DATASET DISTRIBUTION	12
1.5 SOFTWARE COMPONENTS DISTRIBUTION	12
CLAUSE 2: WORKLOAD RELATED ITEMS	13
2.1 HARDWARE & SOFTWARE TUNABLE	13
2.2 RUN REPORT.....	13
2.3 BENCHARK KIT IDENTIFICATION	13
2.4 BENCHARK KIT CHANGES	14
CLAUSE 3: SUT RELATED ITEMS.....	15
3.1 DATA STORAGE RATIO.....	15
3.2 MEMORY RATIO	15
CLAUSE 4: SCALE FACTORS AND METRICS.....	16
4.1 HSGEN TIME	16
4.2 HSSORT TIME.....	16
4.3 HSVVALIDATE TIME	16
4.4 HSDATACHECK TIMES	16
4.5 PERFORMANCE & PRICE-PERFORMANCE	16
AUDITORS' INFORMATION AND ATTESTATION LETTER.....	17
SUPPORTING FILE INDEX.....	20

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 1.2.0.

The test was conducted at a Scale Factor of 10TB on Cisco UCS® Integrated Infrastructure for Big Data (Cisco UCS CPA v2) with 16 Cisco UCS C240 M3 Servers running MapR Distribution including Apache Hadoop Enterprise Edition (MapR M5 Edition).

Measured Configuration

Company Name	Cluster Node	Virtualization	Operating System
Cisco Systems, Inc.	Cisco UCS C240 M3 Server	n/a	Red Hat Enterprise Linux Server 6.4

TPC Express Benchmark© HS Metrics

Total System Cost	HSph@10TB	Price/Performance	Availability Date
614,645 USD	5.77	106,524.27 USD	January 8, 2015

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 run-time, 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. with contribution from MapR.

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 by 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.

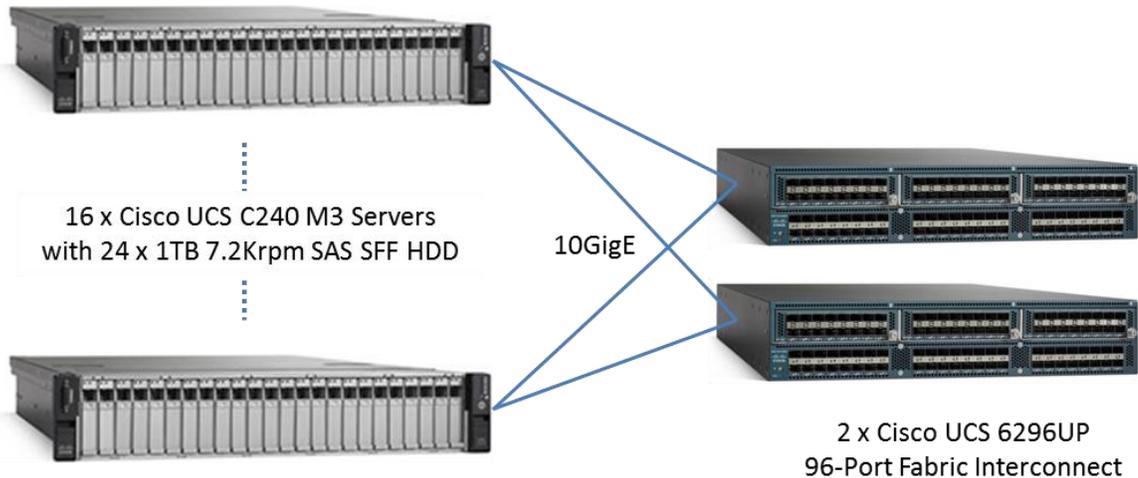
The supporting files contain the parameters and options used to configure the components involved in this benchmark.

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: 32/320/640
- Total Memory: 4,096GB
- Total Number of Storage Drives/Devices: 384
- Total Storage Capacity: 384TB

Server nodes details:

- 16 x Cisco UCS C240 M3 server nodes, each with:
 - Processors/Cores/Threads: 2/20/40
 - Processor Model: Intel® Xeon® E5-2660 v2 @ 2.20GHz with 25MB L3
 - Memory: 256GB
 - Controller: 1 x MegaRAID 9271CV
 - Drives: 24 x 1TB (930.4GB formatted) 7.2Krpm SAS SFF HDD
 - Network: 1 x Cisco VIC 1225 Dual Port 10Gb SFP+ CNA

Network Connectivity:

- 2 x UCS 6296UP Fabric Interconnects

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

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	Controller	Disk Drive	Description of Content
1 - 16	MR 9271	0	Operating system, root, swap, Hadoop Master
1 - 16	MR 9271	1 - 23	Data, Temp

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: Dataset Distribution

Server	Software Component
1	NodeManager ResourceManager JobHistoryServer
2	NodeManager Zookeeper
3	NodeManager CLDB
4 - 9	NodeManager
10	NodeManager Zookeeper
11	NodeManager
12	NodeManager Zookeeper
13 - 16	NodeManager

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

MapR-FS (fully HDFS compatible at the API level).

Map/Reduce implementation and corresponding version must be disclosed.

YARN M/R v2 on MapR 4.0.1 (compatible equivalent to Hadoop 2.4.1).

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.

- **Run1 Performance Summary**

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

Test Run 1 details: Total Time = 5570
                    Total Size = 100000000000
                    Scale-Factor = 10.0000

TPCx-HS Performance Metric (HSph@SF): 6.4632

=====
```

- **Run2 Performance Summary**

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

Test Run 2 details: Total Time = 6234
                    Total Size = 100000000000
                    Scale-Factor = 10.0000

TPCx-HS Performance Metric (HSph@SF): 5.7750

=====
```

2.3 Benchark 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 1.2.0. The checksums for the TPCx-HS kit files are:

```
3092422441 23671 TPCx-HS-master.jar
1487525362 4862 BigData_cluster_validate_suite.sh
3382017450 11873 TPCx-HS-master.sh
```

2.4 Benchark Kit changes

The file TPCx-HS-master.sh was modified to substitute the “hdfs dfs” commands in this shell script to “hadoop fs” commands, since MapR does not to support the hdfs command. The checksum of the modified file is as follows:

```
847410926 12199 TPCx-HS-master.sh
```

The file BigData_cluster_validate_suite.sh was modified as follows:

```
< psh $parg "$SUDO ip addr show"  
---  
> psh $parg 'ip addr show'
```

The checksum of the modified BigData_cluster_validate_suite.sh file is as follows:

```
2740716873 4868 BigData_cluster_validate_suite.sh
```

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	1,000	384,000
Total Storage (TB)		384.0

Scale Factor = 10TB

Data Storage Ratio = (Storage / SF) = 38.40

3.2 Memory Ratio

The Scale Factor to memory ratio must be disclosed.

Total Configured Memory = 4TB

Scale Factor to Memory Ratio = (SF / Memory) = 2.50

Clause 4: Scale Factors and Metrics

4.1 HSGen Time

The HSGen time must be disclosed for Run1 and Run2.

	Run1	Run2
HSGen	465.887	813.984

4.2 HSSort Time

The HSSort time must be disclosed for Run1 and Run2.

	Run1	Run2
HSSort	4325.717	4635.093

4.3 HSValidate Time

The HSValidate time must be disclosed for Run1 and Run2.

	Run1	Run2
HSValidate	770.269	775.866

4.4 HSDataCheck Times

Both HSDataCheck times must be disclosed for Run1 and Run2.

	Run1	Run2
HSDataCheck (pre-Sort)	5.000	6.000
HSDataCheck (post-Sort)	4.000	4.000

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@10TB	6.46	5.77
$\$/HSph@10TB$		106,524.27 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
531 Crystal Hills Blvd
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

January 8, 2015

I verified the TPC Express Benchmark™ HS v1.2.0 performance of the following configuration:

Platform: Cisco UCS Integrated Infrastructure for Big Data (Cisco UCS CPA v2)
with 16 Cisco UCS C240M3 Servers
Operating System: Red Hat Enterprise Linux Server 6.4
Apache Hadoop MapR M5 Edition
Compatible Software:

The results were:

Performance Metric 5.77 HSph@10TB
Run Elapsed Time 6,234 Seconds

<u>Cluster</u>	<u>16 x Cisco UCS C240M3 Servers (each with)</u>		
CPU	2 x Intel Xeon Processor E5-2660 v2 (2.20 GHz, 10-core, 25 MB L3)		
Memory	256 GB		
Storage	<i>Qty</i>	<i>Size</i>	<i>Type</i>
	24	1 TB	7.2K rpm SAS HDD

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 v1.2.0
- 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 10TB
- The generated dataset and the sorted dataset were replicated a minimum of 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	SupportingFilesArchive\Clause2
Clause 3	System configuration details	SupportingFilesArchive\Clause3
Clause 4	Run report	SupportingFilesArchive\Clause4