



Cisco Systems, Inc.

TPC Express Benchmark™ HS (TPCx-HS)
Full Disclosure Report
for
Cisco UCS Integrated Infrastructure for Big Data
(with 17 Cisco UCS C240M4 Servers)
using
Transwarp Data Hub V4.6
and
Red Hat Enterprise Linux Server Release 6.7

First Edition

June 14, 2017

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 © 2017 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. 1.4.2
TPC-Pricing Rev. 2.1.1

Report Date:
June 14, 2017

Total System Cost	TPCx-HS Performance Metric	Price/Performance		
3,831,761 RMB	12.18 HSph@10TB	314,594.50 RMB ¥/HSph@10TB		
Scale Factor	Apache Hadoop Compatible Software	Operating System	Other Software	Availability Date
10	Transwarp Data Hub V4.6	Red Hat Enterprise Linux Server Release 6.7	n/a	June 14, 2017

System Configuration



Physical Storage/Scale Factor: 70.66

Scale Factor/Physical Memory: 2.35

Servers:	17 x Cisco UCS C240M4 Server
Total Processors/Cores/Threads	34/476 /952
Server Configuration:	Per node:
Processors	2 x Intel® Xeon® Processor E5-2680 v4, 2.40 GHz, 35 MB L3
Memory	256
Storage Controller	1 x Cisco 12G SAS Modular Raid Controller
Storage Device	24 x 1.8 TB 10K RPM SAS HDD (data nodes) 4 x 1.8 TB 10K RPM SAS HDD (name node)
	2 x 240 GB 2.5" 6G SATA SSD (all nodes, boot disk)
Network	1 x Cisco VIC 1227 Dual Port 10Gb SFP+
Connectivity:	1 x Cisco Nexus 9372PX Switch 2 x Cisco UCS 6296UP 96-Port Fabric Interconnect
Total Rack Units:	$(17 * C240M4) + (2 * FI) + (1 * 9372PX) = (17 * 2RU) + (2 * 2RU) + (1 * 1RU) = 39RU$



Cisco UCS Integrated Infrastructure for Big Data

TPCx-HS Rev. 1.4.2
TPC-Pricing Rev. 2.1.1

Report Date:
June 14, 2017

Description	Part Number	Source	Unit Price	Qty	Extended Price	3 Year Maint. Price
UCS BD PERF 2x6296,16xC240M4S w /2xE52680v4,256G,24x1.8TB HDD	UCS-SL-CPA4-P2		1 ¥ -	1 ¥	-	
UCS C240M4SX w /2xE52680v4,16x16GB mem,24x1.8TB SAS,VIC1227	UCS-SPBD-C240M4-P2		1 ¥704,806.03	16 ¥	11,276,896.48	
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 ¥	-	
1.8 TB 12G SAS 10K RPM SFF HDD (4K)	UCS-HD18TB10KS4K		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 Pow er Supply for 2U C-Series Servers	UCSC-PSU2V2-1200W		1 ¥ -	32 ¥	-	
AC Pow er Cord - 250V , 10A - PRC	CAB-250V-10A-CN		1 ¥ -	32 ¥	-	
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 Fl w /18p LIC,16xCables	UCS-SA-BD-F196		1 ¥327,847.09	2 ¥	655,694.18	
UCS 6296UP Pow er Supply/100-240VAC	UCS-PSU-6296UP-AC		1 ¥ -	4 ¥	-	
10GBASE-CU SFP+ Cable 3 Meter	SFP-H10GB-CU3M		1 ¥ -	32 ¥	-	
UCS Manager v3.1	N10-MGT014		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 ¥	-	
AC Pow er Cord - 250V , 10A - PRC	CAB-250V-10A-CN		1 ¥ -	4 ¥	-	
UCS C240 M4 SFF 24 HD w /o CPU,mem,HD,PCle,PS,railkt w /expndr	UCSC-C240-M4SX		1 ¥ 58,754.67	1 ¥	58,754.67	
2.40 GHz E5-2680 v4/120W 14C/35MB Cache/DDR4 2400MHz	UCS-CPU-E52680E		1 ¥ 77,344.38	2 ¥	154,688.76	
16GB DDR4-2400-MHz RDIMM/PC4-19200/single rank/x4/1.2v	UCS-MR-1X161RV-A		1 ¥ 10,147.87	16 ¥	162,365.92	
240 GB 2.5 inch Enterprise Value 6G SATA SSD (boot)	UCS-SD240GBKS4-EB		1 ¥ 13,486.37	2 ¥	26,972.74	
1.8 TB 12G SAS 10K RPM SFF HDD (4K)	UCS-HD18TB10KS4K		1 ¥ 26,869.78	4 ¥	107,479.12	
Cisco 12Gbps SAS 2GB FBWC Cache module (Raid 0/1/5/6)	UCSC-MRAID12G-2GB		1 ¥ 20,663.41	1 ¥	20,663.41	
Cisco 12G SAS Modular Raid Controller	UCSC-MRAID12G		1 ¥ 9,647.83	1 ¥	9,647.83	
Cisco UCS VIC1227 VIC MLOM - Dual Port 10Gb SFP+	UCSC-MLOM-CSC-02		1 ¥ 22,045.87	1 ¥	22,045.87	
1200W / 800W V2 AC Pow er Supply for 2U C-Series Servers	UCSC-PSU2V2-1200W		1 ¥ 11,015.58	2 ¥	22,031.16	
AC Pow er Cord - 250V , 10A - PRC	CAB-250V-10A-CN		1 ¥ -	2 ¥	-	
Right PCI Riser Bd (Riser 1) 2onbd SATA bootdrvs+ 2PCI slts	UCSC-PCI-1C-240M4		1 ¥ 2,176.64	1 ¥	2,176.64	
Heat sink for UCS C240 M4 rack servers	UCSC-HS-C240M4		1 ¥ -	2 ¥	-	
Supercap cable 250mm	UCSC-SCCBL240		1 ¥ -	1 ¥	-	
Ball Bearing Rail Kit for C220 M4 and C240 M4 rack servers	UCSC-RAILB-M4		1 ¥ 3,235.55	1 ¥	3,235.55	
Cisco Smart Net 24X7X4 3Y UCS C240 M4S BD SP Server	CON-OSP-C240V4SP		1 ¥ 21,712.50	17		¥ 369,112.50



Cisco UCS Integrated Infrastructure for Big Data

TPCx-HS Rev. 1.4.2
TPC-Pricing Rev. 2.1.1

Report Date:
June 14, 2017

Description	Part Number	Source	Unit Price	Qty	Extended Price	3 Year Maint. Price
2rd Gen FI License to connect C-direct only	UCS-L-6200-10G-C		¥ 13,604.02	30	¥ 408,120.60	
10GBASE-CU SFP+ Cable 3 Meter	SFP-H10GB-CU3M	1	¥ 1,470.71	16	¥ 23,531.36	
Cisco R42610 standard rack w/side panels	RACK-UCS2	1	¥ 76,108.99	1	¥ 76,108.99	
Cisco Smart Net 24X7X4 3Y UCS 6296UP 2RU Fabrc Int/2 PSU/4 Fans	CON-SNTP-SDBFIV96	1	¥ 53,794.65	2		¥ 107,589.30
Cisco Nexus 9372PX Sw itch	N9K-C9372PX	1	¥411,797.44	1	¥ 411,797.44	
Cisco SNTP Cisco Nexus 9372PX Sw itch, 36 24x7	CON-SNTP-9372PX	1	¥ 52,877.22	1		¥ 52,877.22
Red Hat Enterprise Linux Server, 3Y 24x7	CON-ISV1-EL2S2V3A	1	¥ 35,457.60	17		¥ 602,779.20
					Total	¥1,132,358.22
Large Purchase Discount ¹	82% for products and 52% for service	1			¥ -11,022,612.79	¥ -588,826.27
Acer V206HQL - LED monitor - 20" (Inc 2 spares)	UM.1V6AA.A02	2	¥ 516.00	3	¥ 1,548.00	
Logitech USB Corded Keyboard/Mouse Combo MK120 (Inc 2 spares)	920-002565	2	¥ 69.90	3	¥ 209.70	
Transw arp Support for Hadoop	TDH	3	¥ 67,990.00	17	¥ 1,155,830.00	
Transw arp Support Discount (25%)	TDH	3			¥ -288,957.50	

Pricing: 1 = Cisco, 2 = CDW, 3 = TDH

⁽¹⁾ All discounts are based on China 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, InfoSizing

Three-Year Cost of Ownership ¥3,831,761

HSph@10TB 12.18

¥/HSph@10TB ¥314,594.50

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.



Numerical Quantities

Performance Run – Run 2

Scale Factor	10TB
Run Start Time	2017-06-07 21:29:03.000
Run End Time	2017-06-07 22:18:15.232
Run Elapsed Time	2,954.000
HSGen Start Time	2017-06-07 21:29:04.353
HSGen End Time	2017-06-07 21:41:22.310
HSGen Elapsed Time	738.887
HSSort Start Time	2017-06-07 21:41:25.711
HSSort End Time	2017-06-07 22:12:26.548
HSSort Elapsed Time	1,861.789
HSValidate Start Time	2017-06-07 22:12:30.330
HSValidate End Time	2017-06-07 22:18:15.232
HSValidate Elapsed Time	346.592

Repeatability Run – Run 1

Scale Factor	10TB
Run Start Time	2017-06-07 20:30:06.000
Run End Time	2017-06-07 21:18:57.762
Run Elapsed Time	2,933.000
HSGen Start Time	2017-06-07 20:30:07.550
HSGen End Time	2017-06-07 20:42:19.275
HSGen Elapsed Time	732.695
HSSort Start Time	2017-06-07 20:42:22.765
HSSort End Time	2017-06-07 21:13:06.662
HSSort Elapsed Time	1,844.850
HSValidate Start Time	2017-06-07 21:13:10.365
HSValidate End Time	2017-06-07 21:18:57.762
HSValidate Elapsed Time	349.051



Run Reports

Run Report for Performance Run – Run 2

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

Test Run 2 Details	Total Time =	2954
	Total Size =	100000000000
	Scale-Factor =	10

TPCx-HS Performance Metric (HSph@SF): 12.1876

=====
Run Report for Repeatability Run – Run 1

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

Test Run 1 Details	Total Time =	2933
	Total Size =	100000000000
	Scale-Factor =	10

TPCx-HS Performance Metric (HSph@SF): 12.2744

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	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 HSDATA CHECK TIMES	16
4.5 PERFORMANCE & PRICE-PERFORMANCE	16
AUDITORS' INFORMATION AND ATTESTATION LETTER	17
SUPPORTING FILE INDEX	20
THIRD PARTY PRICE QUOTES	21

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 1.4.2.

The test was conducted at a Scale Factor of 10 with 17 Cisco UCS C240M4 Servers running Transwarp Data Hub V4.6 on Red Hat Enterprise Linux Server Release 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 Release 6.7

TPC Express Benchmark©HS Metrics

Total System Cost	HSph@10TB	Price/Performance	Availability Date
3,831,761 RMB	12.18	314,594.50 RMB	June 14, 2017

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 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: 17
- Total Processors/Cores/Threads: 34/476 /952
- Total Memory: 4.25TB
- Total Number of Storage Drives/Devices: 422
- Total Storage Capacity: 706.56TB

Server nodes details:

- 17 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, 35 MB L3
 - Memory: 256GB
 - Controller: 1 x Cisco 12G SAS Modular Raid Controller
 - Drives:
 - 24 x 1.8 TB 10K RPM SAS HDD (data nodes)
 - 4 x 1.8 TB 10K RPM SAS HDD (name node)
 - 2 x 240 GB 2.5" 6G SATA SSD (all nodes, boot disk)
 - Network: 1 x Cisco VIC 1227 Dual Port 10Gb SFP+

Network connectivity detail:

- 1 x Cisco Nexus 9372PX Switch
- 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
1-17	MegaRAID 3108	1-24 (HDD)	Data, Temp
1-17	Intel Chipset Embedded SATA RAID	0 (2 SSD, RAID-1, Internal)	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		X		X	
3		X		X	
4-6		X		X	
7-11		X		X	X
12-17		X		X	

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

Transwarp Data Hub V4.6 (fully HDFS compatible at the API level).

Map/Reduce implementation and corresponding version must be disclosed.

Transwarp Data Hub V4.6 (compatible equivalent to Hadoop 2.8).

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 =                2933
                        Total Size =            100000000000
                        Scale-Factor =                10

TPCx-HS Performance Metric (HSph@SF):                12.2744
=====
```

Run Report for Run 2 – Performance Run

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

Test Run 2 Details      Total Time =                2954
                        Total Size =            100000000000
                        Scale-Factor =                10

TPCx-HS Performance Metric (HSph@SF):                12.1876
=====
```

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)
388	1,800	698,400
34	240	8,160
Total Storage (TB)		706.56

Scale Factor = 10

Data Storage Ratio = (Storage / SF) = **70.66**

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) = **2.35**

Clause 4: Scale Factors and Metrics

4.1 HSGen Time

The HSGen time must be disclosed for Run1 and Run2.

	Run1	Run2
HSGen	732.695	738.887

4.2 HSSort Time

The HSSort time must be disclosed for Run1 and Run2.

	Run1	Run2
HSSort	1,844.850	1,861.789

4.3 HSValidate Time

The HSValidate time must be disclosed for Run1 and Run2.

	Run1	Run2
HSValidate	349.051	346.592

4.4 HSDataCheck Times

Both HSDataCheck times must be disclosed for Run1 and Run2.

	Run1	Run2
HSDataCheck (pre-Sort)	3.490	3.401
HSDataCheck (post-Sort)	3.703	3.782

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	12.27	12.18

¥/HSph@10TB	314,594.50 RMB
-------------	----------------

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, InfoSizing.

www.sizing.com
63 Lourdes Drive
Leominster, MA 10453
978-343-6562

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

June 14, 2017

I verified the TPC Express Benchmark™ HS v1.4.2 performance of the following configuration:

Platform: Cisco UCS Integrated Infrastructure for Big Data
(with 17 Cisco UCS C240M4 Servers)
Operating System: Red Hat Enterprise Linux Server 6.7
Apache Hadoop Transwarp Data Hub V4.6
Compatible Software:

The results were:

Performance Metric 12.18 HSph@10TB
Run Elapsed Time 2,954.00 Seconds

Cluster	17 Cisco UCS C240M4 Servers, each node with:		
CPUs	2 x Intel Xeon Processor E5-2680 v4 (2.40 GHz, 14-core, 35 MB L3)		
Memory	256 GB		
Storage	Qty	Size	Type
	24	1.8TB	10K RPM SAS HDD (Data Nodes)
	4	1.8TB	10K RPM SAS HDD (Name Node)
	2	240GB	2.5" 6G SATA SSD (All Nodes, 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 v1.4.2
- 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 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:

Testing for this audit was performed using v1.4.2 of the TPC-provided kit. However, because of a small issue in this kit, the version is erroneously reported as v1.4.1 in the log file.

Respectfully Yours,

A handwritten signature in black ink that reads "Doug Johnson". The signature is fluid and cursive, with a long horizontal flourish extending to the right.

Doug Johnson, Certified TPC Auditor

63 Lourdes Dr. | Leominster, MA 01453 | 978-343-6562 | www.sizing.com

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

CDW

Secure | https://item.gome.com.cn/A0005829203-pop8008748243.html

国美Plus 国美会员 登录 注册 我的订单 我的国美 企业采购 服务中心

国美在线 GOME.COM.CN 品质·低价·速达

国美手机K1 搜全站 搜本店

显示器 电脑显示器 AOC 明基 飞利浦 LG 宏碁

Lenovo 您身边的电脑专家 宏锦电脑专营店

首页 宏碁台式机 笔记本 一体机 平板电脑 投影仪 联想专区 显示器 私人定制

电脑 办公打印 文仪 > 电脑配件 > 显示器 > 宏碁 (ACER) > 宏碁(acer)V206...

宏锦电脑专营店 4.5分



宏碁(acer)V206HQL 商用办公19.5英寸LED背光宽屏液晶显示器 支持壁挂

缤纷清凉夏 6月礼不断！经典外观，商务办公轻松应对，稳重大气！eColor色彩管理技术！

国美价 **¥516.00** 降价通知

好评度100%
4人评价

促销 **满赠** 订单金额满999元送赠品1件

特色 **管家服务**

送至 北京市朝阳区朝外街道 有货 支持 免运费

服务 由 **宏锦电脑专营店** 发货并负责售后服务。

1 **加入购物车** 手机下单

item.gome.com.cn/A0005198656-pop8005596945.html

国美Plus 国美会员 登录 注册 我的订单 我的国美 企业采购 服务中心

国美在线 GOME.COM.CN 品质·低价·速达

冰箱 搜全站 搜本店

键盘 机械键盘 黑爵 魅影 达尔优 鼠标

正蓝数码专营店 官方授权 国际品质

正蓝电脑办公专营店 4.6分



Logitech/罗技MK120 USB有线鼠标键盘套装 电脑台式机键鼠套装

外形靓丽、手感舒适。高分辨率光学鼠标、有线套装，经久耐用的键鼠套装集舒适与时尚的结合！

国美价 **¥69.90** 降价通知

好评度98
120人评价

特色 **管家服务**

送至 北京市朝阳区朝外街道 有货 支持 免运费

服务 由 **正蓝电脑办公专营店** 发货并负责售后服务。

1 **加入购物车** 手机下单

温馨提示 1. 正品保障；支持7天无理由退货



Cisco

3800 Zanker Road,

San Jose, CA 95134

以下是用于 TPCx-HS 测试用的星环软件产品的价格:

所有的价格是以人民币(¥元)为单位:

产品代码: PID	产品描述: Product Description	单价: Unit Price	数量: quantity	价格: Price
TDH-BAS-SUB TDH-ICP-S 24 x 7 x4 support	星环标准版 4.7, 包含 3 年 产品维护 Standard edition 4.7, including maintenance for 3 years	¥ 67,990 每节 点, per node price 大批量折扣 25% : ¥ 50,990 每个节点 bulk discount of 25% or 50,990 per node	1	¥ 67,990 每节 点 per node 大批量折扣 25% : ¥ 50,990 每个节点 bulk discount 25% or 50990 per node