



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
IBM Open Platform (IBM IOP) 4.1
and
Red Hat Enterprise Linux Server Release 6.7

First Edition

November 10, 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

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

Report Date:
November 10, 2016

Total System Cost

TPCx-HS Performance Metric

Price/Performance

529,779 USD

13.51
HSph@30TB

39,213.85 USD
\$/HSph@30TB

Scale Factor

Apache Hadoop
Compatible Software

Operating System

Other Software

Availability Date

30TB

IBM Open Platform
(IBM IOP) 4.1

Red Hat Enterprise Linux
Server Release 6.7

None

November 10,
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
- 24 x 8TB LFF HDD
- 2 x 480 GB 2.5 inch Enterprise Value SSD (BOOT)



1 x Cisco UCS C220 M4 Server (NameNode)

- 2 x Intel Xeon E5-2680 v4
- 256 GB Memory
- 8 x 1.2TB 12G SAS 10K rpm SFF HDD

40GigE



2 x Cisco UCS 6332 32-Port Fabric
Interconnect

1*40GigE

Upstream

Physical Storage/Scale Factor: 103.23

Scale Factor/Physical Memory: 7.06

Servers:	8 x Cisco UCS S3260 2-Node Servers (DataNodes), 1 x Cisco UCS C220 M4 Server (NameNode)
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	256GB
Storage Controller	1 x Cisco UCS C3000 RAID Controller M4 Server w 4G RAID Cache (DataNodes) 1 x Cisco 12G SAS Modular Raid Controller (NameNode)
Storage Device	24 x 8TB LFF HDD (DataNodes) 2 x 480GB 2.5in. Enterprise Value SSD (DataNodes, boot disk) 8 x 1.2TB 12G SAS 10K rpm SFF HDD (NameNode)
Network	Cisco VIC 1300
Connectivity:	2 x Cisco UCS 6332 32-Port Fabric Interconnect
Total Rack Units:	(8*S3260)+(1*C220)+(2*FI) = (8*4RU)+(1*1RU)+(2*1RU) = 32+1+2 = 35



Cisco UCS Integrated Infrastructure for Big Data

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

Report Date:
November 10, 2016

Description	Part Number	Source	Unit Price	Qty	Extended Price	3 Year Maint. Price
Cisco S3260 2 Node Bundle						
UCS C3000 M4 Server Node for Intel E5-2600 v4	UCSS-SP-S3260-BV1		\$97,360.20	8	\$ 778,881.60	
2.40 GHz E5-2680 v4/120W 14C/35MB Cache/DDR4 2400MHz	UCSC-C3K-M4SRB		\$ -	2	\$ -	
16GB DDR4-2400-MHz RDIMM/PC4-19200/single rank/x4/1.2v	UCS-CPU-E52680E		\$ -	2	\$ -	
Cisco UCS C3000 RAID Controller M4 Server w 4G RAID Cache	UCS-MR-1X161RV-A		\$ -	16	\$ -	
Cisco UCS C3X60 Server Node CPU Heatsink	UCS-C3K-M4RAID		\$ -	1	\$ -	
UCSC C3X60 480GB Boot SSD (Gen 2)	UCSC-HS-C3X60		\$ -	2	\$ -	
Cisco UCS C3260 System IO Controller w with VIC 1300 incl.	UCS-C3X60-G2SD48		\$ -	2	\$ -	
UCSC 3X60 8TB NL-SAS 7.2KHelium HDD w with HDD Carrier	UCSC-C3260-SIOC		\$ -	2	\$ -	
Cisco UCS C3260 Base Chassis w/4x PSU, SSD, Railkit	UCSC-C3X60-HD8TB		\$ -	48	\$ -	
Power Cord Jumper, C13-C14 Connectors, 2 Meter Length	UCSC-C3260		\$ -	1	\$ -	
Cisco UCS C3160 System Bezel	CAB-C13-C14-2M		\$ -	4	\$ -	
UCS C3X60 Rack Rails Kit	UCSC-C3160-BEZEL		\$ -	1	\$ -	
UCS C3X60 1050W Power Supply Unit	UCSC-C3X60-RAIL		\$ -	1	\$ -	
	UCSC-PSU1-1050W		\$ -	4	\$ -	
UCS C220 M4 SFF w/o CPU, mem, HD, PCIe, PSU, rail kit	UCSC-C220-M4S		\$ 3,365.00	1	\$ 3,365.00	
2.40 GHz E5-2680 v4/120W 14C/35MB Cache/DDR4 2400MHz	UCS-CPU-E52680E		\$ 5,259.00	2	\$ 10,518.00	
16GB DDR4-2133-MHz RDIMM/PC4-17000/dual rank/x4/1.2v	UCS-MR-1X162RU-A		\$ 600.00	16	\$ 9,600.00	
1.2 TB 12G SAS 10K RPM SFF HDD	UCS-HD12TB10K12G		\$ 1,460.00	8	\$ 11,680.00	
Cisco VIC 1387 Dual Port 40Gb QSFP CNA MLOM	UCSC-MLOM-C40Q-03		\$ 2,192.00	1	\$ 2,192.00	
Cisco 12G SAS Modular Raid Controller	UCSC-MRAID12G		\$ 656.00	1	\$ 656.00	
Cisco 12Gbps SAS 2GB FBWC Cache module (Raid 0/1/5/6)	UCSC-MRAID12G-2GB		\$ 1,405.00	1	\$ 1,405.00	
Ball Bearing Rail Kit for C220 M4 and C240 M4 rack servers	UCSC-RAILB-M4		\$ 220.00	1	\$ 220.00	
1050W V2 -48 VDC Power Supply for C220M4 C-Series Server	UCSC-PSUV2-1050DC		\$ 1,199.00	2	\$ 2,398.00	
C-Series -48V DC PSU Power Cord, 3.5M, 3 Wire, 8AWG, 40A	CAB-48DC-40A-8AWG		\$ -	2	\$ -	
UCS 6332 1RU FI/No PSU/32 QSFP+	UCS-SP-F6332		\$22,000.00	2	\$ 44,000.00	
3rd Gen FI Per port License to connect C-direct only	UCS-LIC-6300-40GC		\$ 1,388.00	20	\$ 27,760.00	
40GBASE-CR4 Passive Copper Cable, 3m	QSFP-H40G-CU3M-RF		\$ 150.00	36	\$ 5,400.00	
Cisco R42610 standard rack w/side panels	RACK-UCS2		\$ 3,429.00	1	\$ 3,429.00	
3YR 24X7X4 Support UCS S3260	CON-OSP-S3260BSE		\$ 7,872.00	8		\$ 62,976.00
3YR 24X7X4 Support UCS C220 M4 SFF	CON-3SNTP-C220M4S		\$ 1,496.00	1		\$ 1,496.00
3YR 24x7x4 Support UCS6332	CON-OSPT-F6332		\$ 3,330.25	2		\$ 6,660.50
Red Hat Enterprise Linux Server, 3Y 24x7	CON-ISV1-RH2SUG3A		\$ 2,397.00	17	\$ 40,749.00	
				Total	\$ 942,253.60	\$ 71,132.50

(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 10, 2016

Description	Part Number	Source	Unit Price	Qty	Extended Price	3 Year Maint. Price
(continued from previous page)						
				Total	\$ 942,253.60	\$ 71,132.50
Large Purchase Discount 1	61% for products and 35% for service	1			\$ (574,774.70)	\$(24,896.38)
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	
IBM Elite Support for Apache Hadoop	IBM-D1G8TLL	3	\$ 6,810.00	17	\$ 115,770.00	
Pricing: 1 = Cisco, 2 = CDW.com, 3 = IBM ⁽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. Audited by Doug Johnson of InfoSizing					Three-Year Cost of Ownership \$529,779 HSph@30TB 13.51 \$/HSph@30TB \$39,213.85	
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

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

Report Date:
November 10, 2016

Performance Run

Scale Factor	30TB
Run Start Time	2016-10-29 06:24:50
Run End Time	2016-10-29 08:37:59
Run Elapsed Time	7,991.000
HSGen Start Time	2016-10-29 06:24:51
HSGen End Time	2016-10-29 06:52:01
HSGen Elapsed Time	1,631.350
HSSort Start Time	2016-10-29 06:52:03
HSSort End Time	2016-10-29 08:22:05
HSSort Elapsed Time	5,403.093
HSValidate Start Time	2016-10-29 08:22:09
HSValidate End Time	2016-10-29 08:37:59
HSValidate Elapsed Time	951.669

Repeatability Run

Scale Factor	30TB
Run Start Time	2016-10-29 08:40:05
Run End Time	2016-10-29 10:53:14
Run Elapsed Time	7,991.000
HSGen Start Time	2016-10-29 08:40:06
HSGen End Time	2016-10-29 09:07:32
HSGen Elapsed Time	1,647.728
HSSort Start Time	2016-10-29 09:07:35
HSSort End Time	2016-10-29 10:36:35
HSSort Elapsed Time	5,341.457
HSValidate Start Time	2016-10-29 10:36:38
HSValidate End Time	2016-10-29 10:53:14
HSValidate Elapsed Time	997.461



**Cisco UCS Integrated Infrastructure
for Big Data**

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

Report Date:
November 10, 2016

Run Report for Performance Run - Run 1

TPCx-HS Performance Metric (HSph@SF) Report

Test Run 1 Details	Total Time =	7991
	Total Size =	300000000000
	Scale-Factor =	30

TPCx-HS Performance Metric (HSph@SF): 13.5153

Run Report for Repeatability Run - Run 2

TPCx-HS Performance Metric (HSph@SF) Report

Test Run 2 Details	Total Time =	7991
	Total Size =	300000000000
	Scale-Factor =	30

TPCx-HS Performance Metric (HSph@SF): 13.5153

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 30TB with 8 Cisco UCS S3260 2-Node Servers running IBM Open Platform (IBM IOP) 4.1 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@30TB	Price/Performance	Availability Date
529,779 USD	13.51	39,213.85 USD	November 10, 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.

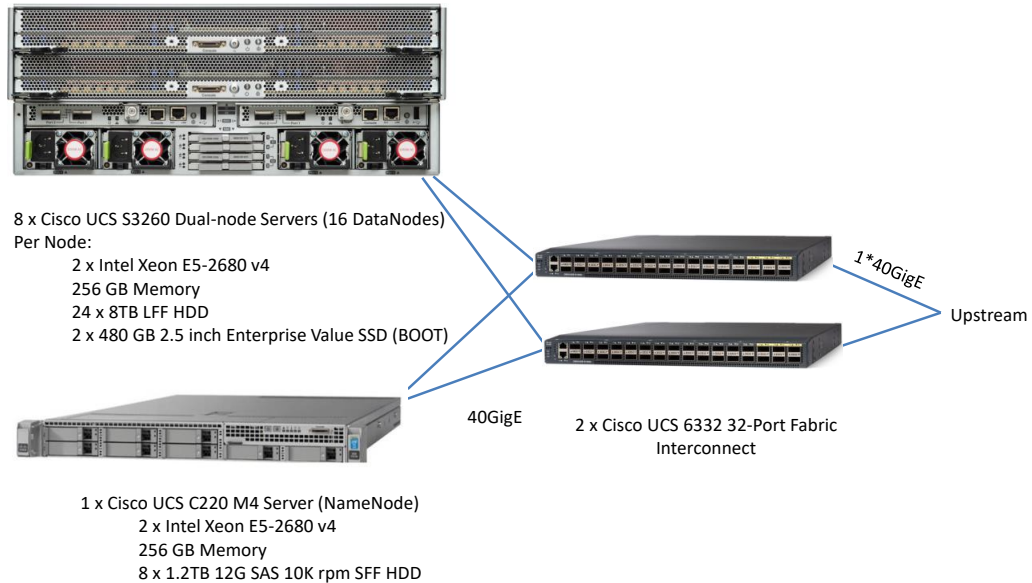
The supporting files contain the parameters and options used to configure the components involved in this benchmark.

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: 17
- Total Processors/Cores/Threads: 34/476/952
- Total Memory: 4.25TB
- Total Number of Storage Drives/Devices: 424
- Total Storage Capacity: 3,096.96TB

Server nodes details:

- 8 x Cisco UCS S3260 2-Node Servers (DataNodes), 1 x Cisco UCS C220 M4 Server (NameNode) each node 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 UCS C3000 RAID Controller M4 Server w 4G RAID Cache (DataNodes)
 - 1 x Cisco 12G SAS Modular Raid Controller (NameNode)
 - Drives:
 - 24 x 8TB LFF HDD (DataNodes)
 - 2 x 480GB 2.5in. Enterprise Value SSD (DataNodes, boot disk)
 - 8 x 1.2 TB 12G SAS 10K rpm SFF HDD (NameNode)
 - 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 12G SAS Modular Raid Controller	3-8 (HDD)	Data, Temp
1	Cisco 12G SAS Modular Raid Controller	1-2 (2 HDD, RAID-1)	Operating System, Root, Swap, Hadoop Master
2-17	Cisco UCS C3000 RAID Controller M4 Server w 4G RAID Cache	1-24 (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		X		X	X
3		X		X	X
4-17		X		X	

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

IBM Open Platform (IBM IOP) 4.1 (fully HDFS compatible at the API level).

Map/Reduce implementation and corresponding version must be disclosed.

IBM Open Platform (IBM IOP) 4.1 (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 - Performance Run

TPCx-HS Performance Metric (HSph@SF) Report

Test Run 1 Details	Total Time =	7991
	Total Size =	300000000000
	Scale-Factor =	30

TPCx-HS Performance Metric (HSph@SF): 13.5153

Run Report for Run 2 - Repeatability Run

TPCx-HS Performance Metric (HSph@SF) Report

Test Run 1 Details	Total Time =	7991
	Total Size =	300000000000
	Scale-Factor =	30

TPCx-HS Performance Metric (HSph@SF): 13.5153

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)
384	8,000	3,072,000
32	480	15,360
8	1,200	9,600
Total Storage (TB)		3,096.96

Scale Factor = 30TB

Data Storage Ratio = (Storage / SF) = **103.23**

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) = **7.06**

Clause 4: Scale Factors and Metrics

4.1 HSGen Time

The HSGen time must be disclosed for Run1 and Run2.

	Run1	Run2
HSGen	1,631.350	1,647.728

4.2 HSSort Time

The HSSort time must be disclosed for Run1 and Run2.

	Run1	Run2
HSSort	5,403.093	5,341.457

4.3 HSValidate Time

The HSValidate time must be disclosed for Run1 and Run2.

	Run1	Run2
HSValidate	951.669	997.461

4.4 HSDataCheck Times

Both HSDataCheck times must be disclosed for Run1 and Run2.

	Run1	Run2
HSDataCheck (pre-Sort)	2.000	3.000
HSDataCheck (post-Sort)	4.000	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@30TB	13.51	13.51

\$/HSph@30TB	39,213.85 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.

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 IBM Open Platform (IBM IOP) 4.1
Compatible Software:

The results were:

Performance Metric 13.51 HSph@30TB

Run Elapsed Time 7,991.00 Seconds

Cluster **8 Cisco UCS S3260 Dual-node Servers (Data Nodes),**
1 Cisco UCS C220 M4 Server (Name Node)
each node with:

CPU	2 x Intel Xeon Processor E5-2680 v4 (2.40 GHz, 14-core, 35 MB L3)		
Memory	256 GB		
Storage	Qty	Size	Type
	24	8 TB	LFF HDD (Data Nodes)
	2	480GB	SSD (Data Nodes, boot disk)
	8	1.2 TB	12G SAS 10K rpm SFF HDD (Name Node)

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 30TB
- 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

IBM

Date: 10/28/16

Cisco
Raghunath Nambiar

3800 Zanker Road

San Jose, CA 95134

Here is the information you requested regarding pricing for IBM products to be used in conjunction with your TPC-H benchmark testing.

All pricing shown is in US Dollars (\$).

Part Number	Description	Unit Price	Quantity	Price
D1G8TLL	IBM Elite Support for Apache Hadoop Unlimited Contacts SW Support 12 Months	\$2,580 per Virtual Server (US – SRP) Retail Discounted Price: \$2,270	1 Virtual Server per 1 Node - for 3 years	\$6,810

This quote is valid for the next 60 days.

IBM IOP Team



Search CDW...



HARDWARE SOFTWARE SOLUTIONS CLOUD BRANDS BLOG DEALS

800.800.4239

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](#)



Search CDW...



HARDWARE SOFTWARE SOLUTIONS CLOUD BRANDS BLOG DEALS

800.800.4235

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