



Cisco Systems, Inc.

TPC Express Benchmark™ HS Full Disclosure Report

Cisco Data Intelligence Platform with All NVMe

(with 16x UCSC-C220-M5SN Servers; 1x UCSC-C220-M5SX Servers)

Running

Cloudera Data Platform Private Cloud Base

7.1.1

on

Red Hat Enterprise Linux Server 7.7

TPCx-HS Version	2.0.3
Report Edition	First
Report Submitted	November 9, 2020

First Edition - November 2020

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 because of these and other factors. Therefore, the TPC Express Benchmark™ V 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. 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 Express Benchmark™ HS, 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© 2020 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.

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 2.0.3.


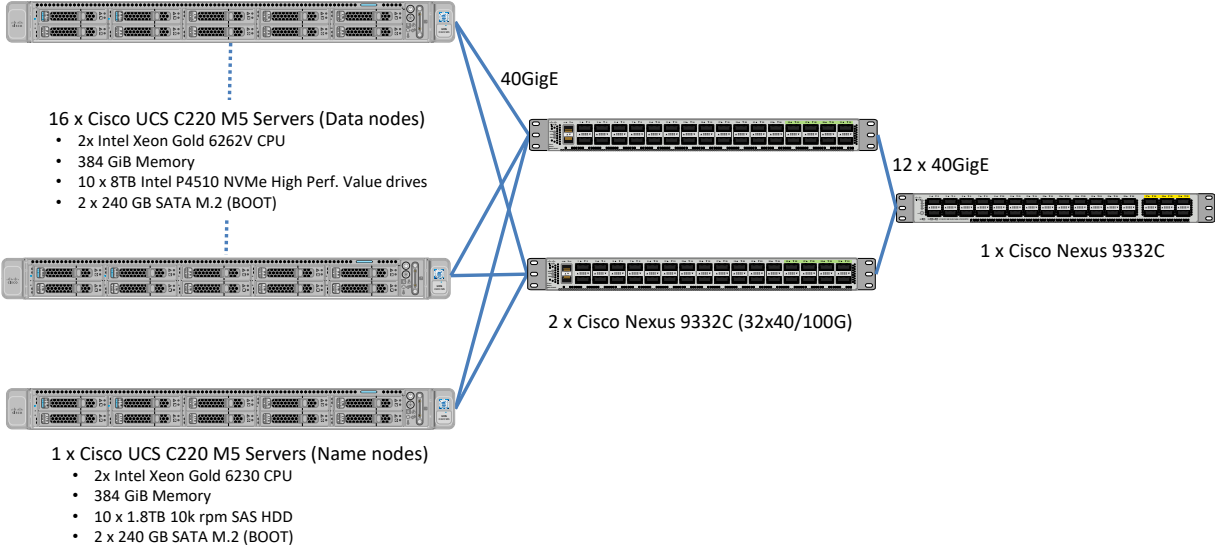
The benchmark results are summarized below.


Measured Configuration			
Company Name	Cluster Node	Hadoop Software	Operating System
Cisco	UCSC-C220-M5SN	Cloudera Data Platform Private Cloud Base 7.1.1	Red Hat Enterprise Linux Server 7.7


TPC Express Benchmark™ HS Metrics			
Total System Cost	HSph@10TB	Price/Performance	Availability Date
\$1,964,279	16.51	\$118,975.11	Currently Available


Executive Summary

The [Executive Summary](#) follows on the next several pages.

	<h2 style="text-align: center;">Cisco Data Intelligence Platform with All NVMe</h2>		TPCx-HS 2.0.3
			TPC Pricing 2.6.0
			Report Date Nov. 09, 2020
Availability Date Currently Available	TPCx-HS Performance 16.51 HSph@10TB	Price/Performance \$118,975.11 \$ / HSph@10TB	Total System Cost \$1,964,279 USD
System Under Test Configuration Overview			
Scale Factor 10	Hadoop Software Cloudera Data Platform Private Cloud Base 7.1.1	Operating System Red Hat Enterprise Linux Server 7.7	Other Software None
 <p>16 x Cisco UCS C220 M5 Servers (Data nodes)</p> <ul style="list-style-type: none"> • 2x Intel Xeon Gold 6262V CPU • 384 GiB Memory • 10 x 8TB Intel P4510 NVMe High Perf. Value drives • 2 x 240 GB SATA M.2 (BOOT) <p>1 x Cisco UCS C220 M5 Servers (Name nodes)</p> <ul style="list-style-type: none"> • 2x Intel Xeon Gold 6230 CPU • 384 GiB Memory • 10 x 1.8TB 10k rpm SAS HDD • 2 x 240 GB SATA M.2 (BOOT) <p>2 x Cisco Nexus 9332C (32x40/100G)</p> <p>1 x Cisco Nexus 9332C</p>			
Physical Storage/Scale Factor: 130.62		Scale Factor/Physical Memory: 1.57	
Total Number of Servers:		17 (16x UCSC-C220-M5SN; 1x UCSC-C220-M5SX)	
Total Processors/Cores/Threads:		34/808/1,616	
Server Configuration:	Per UCSC-C220-M5SN	Per UCSC-C220-M5SX	
Processors	2x Intel(R) Xeon(R) Gold 6262V CPU	2x Intel(R) Xeon(R) Gold 6230 CPU	
Memory	384 GiB	384 GiB	
Storage Controller	None	1x Cisco 12G SAS Modular RAID	
Storage Device	2x 240GB M.2 SATA SSD	2x 240GB M.2 SATA SSD	
	10x 8TB Intel P4510 NVMe	10x 1.8TB 12G SAS 10K RPM HDD	
Network	1x VIC 1387 Dual Port 40Gb MLOM	1x VIC 1387 Dual Port 40Gb MLOM	
Connectivity:	2x Cisco Nexus 9332C, 1x Cisco Nexus 9332C		
Total Rack Units:	(17*C220M5) + (2*N9332C) + (1*N9332C) = (17*2) + (2*1) + (1*1) = 37 RU		

	<h2 style="text-align: center;">Cisco Data Intelligence Platform with All NVMe</h2>		TPCx-HS		2.0.3	
			TPC Pricing		2.6.0	
			Report Date		Nov. 09, 2020	
Description	Part Number	Source	Unit Price	Qty	Extended Price	3 Yr. Maint. Price
Data Nodes						
UCS C220 M5 SFF 10 NVMe w/o CPU, mem, HD, PCIe, PSU	UCSC-C220-M5SN	1	\$5,418.15	16	\$86,690.40	
SNTC 24X7X40S UCS C220 M5 SFF NVMe 10 HD w/o CPU, mem, HD, P	CON-OSP-C220M5SN	1	\$2,227.05	16		\$35,632.80
32GB DDR4-2933-MHz RDIMM/2Rx4/1.2v	UCS-MR-X32G2RT-H	1	\$2,050.15	192	\$393,628.80	
8TB 2.5in U.2 Intel P4510 NVMe High Perf. Value Endurance	UCSC-NVMEHW-I8000	1	\$14,726.36	160	\$2,356,217.60	
Cisco VIC 1387 Dual Port 40Gb QSFP CNA MLOM	UCSC-MLOM-C40Q-03	1	\$2,199.67	16	\$35,194.72	
240GB SATA M.2	UCS-M2-240GB	1	\$536.87	32	\$17,179.84	
IMC SW (Recommended) latest release for C-Series Servers.	CIMC-LATEST	1	\$0.00	16	\$0.00	
Cisco UCS 1050W AC Power Supply for Rack Server	UCSC-PSU1-1050W	1	\$731.55	32	\$23,409.60	
Power Cord, 200/240V 6A North America	CAB-N5K6A-NA	1	\$0.00	32	\$0.00	
Ball Bearing Rail Kit for C220 & C240 M4 & M5 rack servers	UCSC-RAILB-M4	1	\$220.77	16	\$3,532.32	
Big Data and Analytics Platform (Hadoop/loT/ITOA/AI/ML)	UCS-SID-INFR-BD	1	\$0.00	16	\$0.00	
Big Data and Analytics (Hadoop/loT/ITOA)	UCS-SID-WKL-BD	1	\$0.00	16	\$0.00	
Mini Storage carrier for M.2 SATA/NVME (holds up to 2)	UCS-MSTOR-M2	1	\$0.00	16	\$0.00	
Heat sink for UCS C220 M5 rack servers 150W CPUs & below	UCSC-HS-C220M5	1	\$0.00	32	\$0.00	
Intel 6262V 1.9GHz/135W 24C/ 33MB DCP DDR4 2400 MHz	UCS-CPU-I6262V	1	\$10,396.26	32	\$332,680.32	
Name Node						
UCS C220 M5 SFF 10 HD w/o CPU, mem, HD, PCIe, PSU	UCSC-C220-M5SX	1	\$4,151.95	1	\$4,151.95	
SNTC 24X7X40S UCS C220 M5 SFF 10 HD w/o CPU, mem, HD, PCIe,	CON-OSP-C220M5SX	1	\$2,227.05	1		\$2,227.05
32GB DDR4-2933-MHz RDIMM/2Rx4/1.2v	UCS-MR-X32G2RT-H	1	\$2,050.15	12	\$24,601.80	
1.8TB 12G SAS 10K RPM SFF HDD (4K)	UCS-HD18TB10K4KN	1	\$1,924.71	10	\$19,247.10	
Cisco VIC 1387 Dual Port 40Gb QSFP CNA MLOM	UCSC-MLOM-C40Q-03	1	\$2,199.67	1	\$2,199.67	
240GB SATA M.2	UCS-M2-240GB	1	\$536.87	2	\$1,073.74	
IMC SW (Recommended) latest release for C-Series Servers.	CIMC-LATEST	1	\$0.00	1	\$0.00	
Cisco UCS 1050W AC Power Supply for Rack Server	UCSC-PSU1-1050W	1	\$731.55	2	\$1,463.10	
Power Cord, 200/240V 6A North America	CAB-N5K6A-NA	1	\$0.00	2	\$0.00	
Ball Bearing Rail Kit for C220 & C240 M4 & M5 rack servers	UCSC-RAILB-M4	1	\$220.77	1	\$220.77	
Performance Optimized setting for C220 M5 servers	UCSC-SW-C220M5-P01	1	\$0.00	1	\$0.00	
Big Data and Analytics Platform (Hadoop/loT/ITOA/AI/ML)	UCS-SID-INFR-BD	1	\$0.00	1	\$0.00	
Big Data and Analytics (Hadoop/loT/ITOA)	UCS-SID-WKL-BD	1	\$0.00	1	\$0.00	
Mini Storage carrier for M.2 SATA/NVME (holds up to 2)	UCS-MSTOR-M2	1	\$0.00	1	\$0.00	
Super Cap cable for UCSC-RAID-M5 on C240 M5 Servers	CBL-SC-MR12GM52	1	\$0.00	1	\$0.00	
Super Cap for UCSC-RAID-M5, UCSC-MRAID1GB-KIT	UCSC-SCAP-M5	1	\$0.00	1	\$0.00	
Heat sink for UCS C220 M5 rack servers 150W CPUs & below	UCSC-HS-C220M5	1	\$0.00	2	\$0.00	
Intel 6230 2.1GHz/125W 20C/27.50MB DCP DDR4 2933 MHz	UCS-CPU-I6230	1	\$6,522.75	2	\$13,045.50	
Cisco 12G Modular RAID controller with 2GB cache	UCSC-RAID-M5	1	\$2,163.55	1	\$2,163.55	
Network						
Nexus 9K ACI & NX-OS Spine, 32p 40/100G & 2p 10G	N9K-C9332C	1	\$36,126.00	3	\$108,378.00	
SNTC-24X7X4 Nexus 9K ACI NX-OS Spine, 32p 40/100G	CON-SNTP-N9KC9332	1	\$8,121.00	3		\$24,363.00
Dummy PID for mode selection	MODE-NXOS	1	\$0.00	3	\$0.00	
Nexus 9500, 9300, 3000 Base NX-OS Software Rel 9.3.5	NXOS-9.3.5	1	\$0.00	3	\$0.00	
Nexus 3K/9K Fixed Accessory Kit, 1RU front and rear removal	NXK-ACC-KIT-1RU	1	\$0.00	3	\$0.00	
Nexus AC 1100W PSU - Port Side Exhaust	NXA-PAC-1100W-PE2	1	\$0.00	6	\$0.00	
Power Cord, 125VAC 13A NEMA 5-15 Plug, North America	CAB-9K12A-NA	1	\$0.00	6	\$0.00	
Nexus Fan, 35CFM, port side exhaust airflow	NXA-FAN-35CFM-PE	1	\$0.00	15	\$0.00	
40GBASE Active Optical Cable, 3m	QSFP-H40G-AOC3M=	1	\$950.97	46	\$43,744.62	
(continued next page)						

	<h2 style="text-align: center;">Cisco Data Intelligence Platform with All NVMe</h2>	<p>TPCx-HS 2.0.3 TPC Pricing 2.6.0 Report Date Nov. 09, 2020</p>							
(continued from previous page)									
<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 50%;">Description</th> <th style="width: 10%;">Part Number</th> <th style="width: 10%;">Source</th> <th style="width: 10%;">Unit Price</th> <th style="width: 5%;">Qty</th> <th style="width: 10%;">Extended Price</th> <th style="width: 5%;">3 Yr. Maint. Price</th> </tr> </thead> </table>			Description	Part Number	Source	Unit Price	Qty	Extended Price	3 Yr. Maint. Price
Description	Part Number	Source	Unit Price	Qty	Extended Price	3 Yr. Maint. Price			
Infrastructure									
Cisco R42612 standard rack, w/side panels	RACK2-UCS2	1	\$6,262.84	1	\$6,262.84				
24A Metered Input 1-Phase 6x C19, 36x C13 - 0U PDU (spares)	RP208-30M1P-6-36	1	\$2,471.62	2	\$4,943.24				
Software									
Cloudera Data Platform Data Center with 3-Yr Gold Support	UCS-BD-CDPDC-GL-3Y	1	\$60,000.00	17		\$1,020,000.00			
price per CCU per year for compute in excess of 16 cores/128GB RAM per	UCS-BD-CDP-C-3Y	1	\$450.00	544		\$244,800.00			
price per TB per year for TB in excess of 48 TB per Node	UCS-BD-CDP-S-3Y	1	\$150.00	512		\$76,800.00			
Red Hat Enterprise Linux (1-2 CPU, 1-2 VN); 3-Yr Support Req	RHEL-2S2V-3A=	1	\$0.00	17	\$0.00				
ISV 24X7 RHEL Server 2Socket-OR-2Virtual; ANNUAL List Price	CON-ISV1-EL2S2V3A	1	\$3,897.00	17		\$66,249.00			
Acceptance of Terms, Standalone RHEL License for UCS Servers	UCS-RHEL-TERMS	1	\$0.00	17	\$0.00				
Misc									
Acer V246HQL - LED monitor		2	\$93.99	3	\$281.97				
Logitech MK120 USB Wired Keyboard/Mouse Set		2	\$17.99	3	\$53.97				
Large Purchase Discount 1									
61% for products		1			-\$2,122,817.98				
61% for Cloudera Data Platform						-\$818,376.00			
35% for service and RHEL		1				-\$44,965.14			
Totals					\$1,357,547.44	\$606,730.71			
Pricing: 1 = Cisco; 2 = CDW * Discount applies to all line items where Source = 1. Discount based upon total system cost as purchased by a regular customer.		Three-Year Cost of Ownership: \$1,964,279 HSph@10TB: 16.51 \$ / HSph@10TB: \$118,975.11							
Audited by Doug Johnson, InfoSizing									
<p><i>Prices used in TPC benchmarks reflect the actual prices a customer would pay for a one-time purchase of the stated Line Items. 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 Line Items. For complete details, see the pricing section of the TPC Benchmark Standard. If you find that the stated prices are not available according to these terms, please inform the TPC at pricing@tpc.org. Thank you.</i></p>									

	Cisco Data Intelligence Platform with All NVMe	TPCx-HS 2.0.3 TPC Pricing 2.6.0 Report Date Nov. 09, 2020
Numerical Quantities		
Performance Run – Run 1		
Scale Factor		10TB
Run Start Time		2020-10-20 07:10:48.000
Run End Time		2020-10-20 07:47:06.000
Run Elapsed Time		2,180.000
HSGen Start Time		2020-10-20 07:10:50.000
HSGen End Time		2020-10-20 07:19:16.000
HSGen Elapsed Time		507.427
HSSort Start Time		2020-10-20 07:19:18.000
HSSort End Time		2020-10-20 07:43:07.000
HSSort Elapsed Time		1,429.601
HSValidate Start Time		2020-10-20 07:43:10.000
HSValidate End Time		2020-10-20 07:47:06.000
HSValidate Elapsed Time		238.196
Repeatability Run – Run 2		
Scale Factor		10TB
Run Start Time		2020-10-20 07:48:11.000
Run End Time		2020-10-20 08:24:17.000
Run Elapsed Time		2,168.000
HSGen Start Time		2020-10-20 07:48:12.000
HSGen End Time		2020-10-20 07:56:36.000
HSGen Elapsed Time		505.173
HSSort Start Time		2020-10-20 07:56:38.000
HSSort End Time		2020-10-20 08:20:32.000
HSSort Elapsed Time		1,434.925
HSValidate Start Time		2020-10-20 08:20:35.000
HSValidate End Time		2020-10-20 08:24:17.000
HSValidate Elapsed Time		223.258

	Cisco Data Intelligence Platform with All NVMe	TPCx-HS	2.0.3
		TPC Pricing	2.6.0
		Report Date	Nov. 09, 2020

Run Reports

Run Report for Performance Run – Run 1

=====

TPCx-HS Performance Metric (HSph@SF) Report

Test Run 1 Details	Total Time =	2180
	Total Size =	100000000000
	Scale-Factor =	10

TPCx-HS Performance Metric (HSph@SF): 16.5152

=====

Run Report for Repeatability Run – Run 2

=====

TPCx-HS Performance Metric (HSph@SF) Report

Test Run 2 Details	Total Time =	2168
	Total Size =	100000000000
	Scale-Factor =	10

TPCx-HS Performance Metric (HSph@SF): 16.6057

=====


	Cisco Data Intelligence Platform with All NVMe	TPCx-HS 2.0.3 TPC Pricing 2.6.0 Report Date Nov. 09, 2020
Revision History		
Date November 9, 2020	Edition First	Description Initial Publication

Table of Contents

Abstract..... 3

Executive Summary 3

Table of Contents.....10

Clause 0 – Preamble11

 0.1 TPC Express Benchmark™ HS Overview.....11

Clause 1 – General Items12

 1.1 Test Sponsor12

 1.2 Parameter Settings12

 1.3 Configuration Diagrams12

 1.3.1 Measured Configuration.....13

 1.3.2 Priced Configuration14

 1.4 Dataset Distribution.....14

 1.5 Software Components Distribution.....14

Clause 2 – Workload Related Items.....15

 2.1 Hardware & Software Tunables15

 2.2 Run Report15

 2.3 Benchmark Kit Identification.....15

 2.4 Benchmark Kit Changes15

Clause 3 – SUT Related Items.....16

 3.1 Data Storage Ratio16

 3.2 Memory Ratio16

Clause 4 – Metrics Related Items17

 4.1 HSGen Time.....17

 4.2 HSSort Time17

 4.3 HSValidate Time17

 4.4 HSDataCheck Times17

 4.5 Performance & Price-Performance17

Auditor’s Information & Letter of Attestation18

Supporting Files Index21

Third-Party Price Quotes.....22

Clause 0 – Preamble

0.1 TPC Express Benchmark™ HS Overview

The 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. 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. 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-HS 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

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.*

1.3.1 Measured Configuration

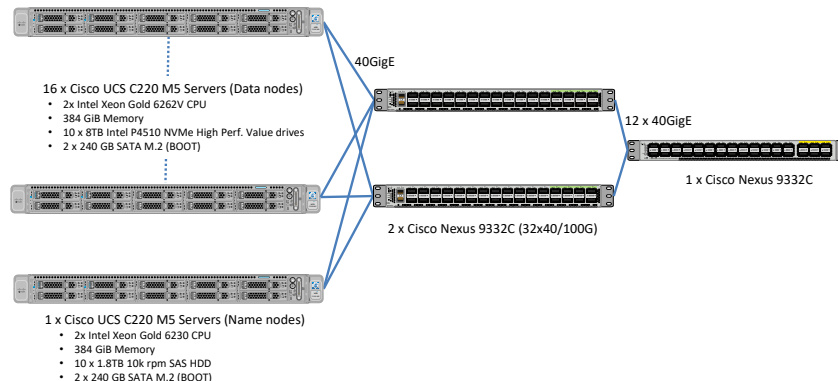


Figure 1-1 Measured Configuration

The measured configuration consisted of:

- Total Nodes: 17 (16x UCSC-C220-M5SN; 1x UCSC-C220-M5SX)
- Total Processors/Cores/Threads: 34/808/1,616
- Total Memory: 6.38TiB
- Total Number of Storage Drives/Devices: 204
- Total Storage Capacity: 1,306.16TB

Server node details:

- 16x UCSC-C220-M5SN Servers, each with:
 - Processors/Cores/Threads: 2/48/96
 - Processor Model: Intel(R) Xeon(R) Gold 6262V CPU
 - Memory: 384 GiB
 - Controller: None
 - Drives:
 - 2x 240GB M.2 SATA SSD
 - 10x 8TB Intel P4510 NVMe
 - Network: 1x VIC 1387 Dual Port 40Gb MLOM
- 1x UCSC-C220-M5SX Servers, each with:
 - Processors/Cores/Threads: 2/40/80
 - Processor Model: Intel(R) Xeon(R) Gold 6230 CPU
 - Memory: 384 GiB
 - Controller: 1x Cisco 12G SAS Modular RAID with 2 GB cache
 - Drives:
 - 2x 240GB M.2 SATA SSD
 - 10x 1.8TB 12G SAS 10K RPM HDD
 - Network: 1x VIC 1387 Dual Port 40Gb MLOM

Network connectivity detail:

- 2x Cisco Nexus 9332C, 1x Cisco Nexus 9332C

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

1.3.2 Priced Configuration

There are no differences between the priced configuration and the measured configuration.

1.4 Dataset Distribution

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

Table 1-1 describes the distribution of the dataset across all media in the system.

Server Node	Controller	Disk Drive	Description of Content
1	Cisco 12G Modular RAID controller w/ 2GB cache	1-10 (1.8TB 12G SAS 10K RPM HDD; RAID-1)	Data and Temp
1	Embedded RAID PCH SATA	0 (2x240GB SSD RAID-1)	Boot Disk; OS, Root, Swap, Hadoop Master
2-17	Embedded RAID PCH SATA	0 (2x240GB SSD RAID-1)	Boot Disk; OS, Root, Swap, Hadoop Master
2-17	NVMe	1-2 (8TB Intel P4510 NVMe)	Temp (NodeManager Local Directories, NodeManager Container Log Directories)
2-17	NVMe	3-10 (8TB Intel P4510 NVMe)	Data

Table 1-1 Dataset Distribution

1.5 Software Components Distribution

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

Table 1-2 Describes the distribution of the software components across the system.

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

Table 1-2 Software Component Distribution

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

Cloudera Data Platform Private Cloud Base 7.1.1 (fully HDFS compatible at the API level).

Map/Reduce implementation and corresponding version must be disclosed.

Cloudera Data Platform Private Cloud Base 7.1.1 (compatible equivalent to Hadoop 3.1.1).

Clause 2 – Workload Related Items

2.1 Hardware & Software Tunables

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 =                2180
                        Total Size =                100000000000
                        Scale-Factor =                10

TPCx-HS Performance Metric (HSph@SF):                16.5152
    
```

```

Run Report for Run 2 – Repeatability Run
=====
TPCx-HS Performance Metric (HSph@SF) Report

Test Run 2 Details      Total Time =                2168
                        Total Size =                100000000000
                        Scale-Factor =                10

TPCx-HS Performance Metric (HSph@SF):                16.6057
    
```

2.3 Benchmark Kit Identification

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

```

Kit Version                2.0.3

File                        MD5
-----
BigData_cluster_validate_suite.sh  57f7cd68251a9aba0feb6648630ff5da
HSDDataCheck.sh                faeff3091759aac98080be4e39f7896a
TPCx-HS-master_MR2.jar         492cbc51a1a60c28b43d96c79d08683d
TPCx-HS-master.sh              c619a0819571ecd00cd75d2b76ba8c64
    
```

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.

Quantity	Capacity	Total (TB)
32	0.24	7.68
160	8.00	1,280.00
2	0.24	0.48
10	1.80	18.00
Total Storage (TB)		1,306.16

Table 3-1 Storage Device Capacities

Scale Factor = 10

Data Storage Ratio = (Total Storage (TB) / SF) = 130.62

3.2 Memory Ratio

The Scale Factor to memory ratio must be disclosed.

Total Configured Memory (TiB) = 6.38

Scale Factor to Memory Ratio = (SF / Total Memory(TiB)) = 1.57

Clause 4 – Metrics Related Items

4.1 HSGen Time

The HSGen time must be disclosed for Run1 and Run2.

	Run 1	Run 2
HSGen	507.427	505.173

Table 4-1 HSGen Times

4.2 HSSort Time

The HSSort time must be disclosed for Run1 and Run2.

	Run 1	Run 2
HSSort	1,429.601	1,434.925

Table 4-2 HSSort Times

4.3 HSValidate Time

The HSValidate time must be disclosed for Run1 and Run2.

	Run 1	Run 2
HSValidate	238.196	223.258

Table 4-3 HSValidate Times

4.4 HSDataCheck Times

Both HSDataCheck times must be disclosed for Run1 and Run2.

	Run 1	Run 2
HSDataCheck (pre-sort)	2.000	2.000
HSDataCheck (post-sort)	3.000	3.000

Table 4-4 HSDataCheck Times

4.5 Performance & Price-Performance

The performance metric (HSph@SF) must be disclosed for Run 1 and Run 2. Price-performance metric (\$/HSph@SF) must be disclosed for the performance run.

	Run 1	Run 2
HSph@10TB	16.51	16.60

Table 4-5 Performance Metrics

Run 1 Price-Performance: 118,975.11 \$/ HSph@10TB

Auditor's Information & Letter of Attestation

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 Letter of Attestation follows.



Karthik Krishna
 Cisco Systems Inc.
 3800 Zanker Road
 San Jose, CA 95134

November 8, 2020

I verified the TPC Express Benchmark™ HS v2.0.3 performance of the following configuration:

Platform: Cisco Data Intelligence Platform with All NVMe
 16x Cisco UCSC C220-M5SN Servers (Data Nodes)
 1x Cisco UCSC C220-M5SX Server (Name Node)

Operating System: Red Hat Enterprise Linux Server 7.7

Apache Hadoop Cloudera Data Platform Private Cloud Base 7.1.1 using MapReduce

Compatible Software:

The results were:

Performance Metric 16.51 HSpH@10TB
 Run Elapsed Time 2,180.00 Seconds

Cluster	16x UCSC C220 M5SN, 1x UCSC C220 M5SX with:		
CPU	2x Intel® Xeon® Gold 6262V (1.90 GHz, 24-core, 33 MB L3) Data Nodes		
	2x Intel® Xeon® Gold 6230 (2.10 GHz, 20-core, 27.5 MB L3) Name Node		
Memory	384 GiB (all nodes)		
Storage	Qty	Size	Type
	2	240 GB	SATA M.2 (All nodes)
	10	8 TB	Intel P4510 NVMe (Data Nodes)
	10	1.8 TB	12G SAS 10K RPM 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 v2.0.3
- 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 10 TB

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

- 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, Certified TPC Auditor

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

Supporting Files Index

Clause	Description	Archive File Pathname
Clause 1	Parameters and options used to configure the system	SupportingFiles/Clause1
Clause 2	Configuration scripts and Run Report	SupportingFiles/Clause2
Clause 3	System configuration details	SupportingFiles/Clause3

Third-Party Price Quotes



5567906



Sign In

Cart (0)

Hardware Software Services IT Solutions Brands Tech Library

Home > Monitors & Projectors > LCD / LED Monitors



Roll over image to zoom



Share Save as Favorites



★★★★★

Acer V246HQL – LED monitor – Full HD (1080p) – 23.6"

Mfg.Part: UM.UV6AA.002 | CDW Part: 5567906 | UNSPSC: 43211902

Availability: • **In Stock**
Get it Monday, September 21 to **95101**

Warranties

- None
- Include: Acer – extended service agreement – 2 years – 4th/5th year – carry-in – \$98.99

[View All Warranties](#) [View Services](#)

\$93.99 Advertised Price

Claim Your 5% Discount
Create an account to get My CDW Advantage today.

1

Add to Cart

Product Details

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



Roll over image to zoom



★★★★★ (14)

Logitech MK120 USB Wired Keyboard/Mouse Set

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

BEST SELLER

Availability: • **In Stock**
Get it Monday, September 21 to **95101**

[View Accessories](#)

Was ~~\$49.99~~
\$17.99 Advertised Price

Claim Your 5% Discount
Create an account to get My CDW Advantage today.

1

Add to Cart