Hewlett Packard Enterprise

Hewlett Packard Enterprise Company

TPC Express Benchmark™ HS
Full Disclosure Report

HPE ProLiant DL325 Gen11

(with 1x ProLiant DL325 Gen11 Server; 16x ProLiant DL325 Gen11 Servers)

Running

CDP Private Cloud Base Edition 7.1.7

on

Red Hat Enterprise Linux 8.6

TPCx-HS Version	2.0.3
Report Edition	First
Report Submitted	March 23, 2023

First Edition - March 2023

Hewlett Packard Enterprise Company (HPE)*, 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.

HPE and the HPE Logo are trademarks of Hewlett Packard Enterprise Company 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 HPE and any other company

TPC Express Benchmark™ HS, TPCx-HS, and HSph, are registered certification marks of the Transaction Processing Performance Council.

The HPE 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 HPE business contact for information on the products or services available in your area. You can find additional information via HPE's web site at www.hpe.com. Actual performance and environmental costs of HPE products will vary depending on individual customer configurations and conditions.

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
HPE	ProLiant DL325 Gen11	CDP Private Cloud Base Edition 7.1.7	Red Hat Enterprise Linux 8.6

TPC Express Benchmark™ HS Metrics			
Total System Cost	HSph@3TB	Price/Performance	Availability Date
$1,886,046	49.83	$37,849.61	December 5, 2022

Executive Summary

The Executive Summary follows on the next several pages.

Hewlett Packard Enterprise	HPE ProLiant DL325 Gen11		TPCx-HS 2.0.3 TPC Pricing 2.8.0 Report Date Mar. 23, 2023
Availability Date	TPCx-HS Performance	Price/Performance	Total System Cost
December 5, 2022	**49.83 HSph@3TB**	**$37,849.61 $ / HSph@3TB**	**$1,886,046 USD**
System Under Test Configuration Overview			
Scale Factor	Hadoop Software	Operating System	Other Software
3	CDP Private Cloud Base Edition 7.1.7	Red Hat Enterprise Linux 8.6	N/A

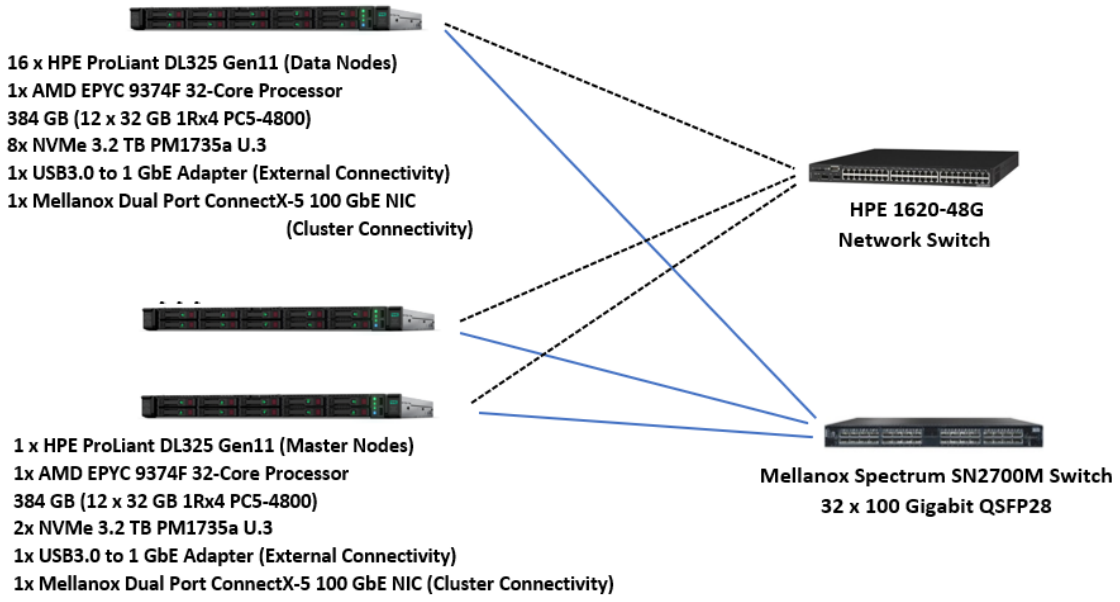

Physical Storage/Scale Factor: 138.67		Scale Factor/Physical Memory: 0.47
Total Number of Servers: Total Processors/Cores/Threads:		17 (1x ProLiant DL325 Gen11; 16x ProLiant DL325 Gen11) 17/544/1,088
Server Configuration: Processors Memory Storage Device Network	1x ProLiant DL325 Gen11 1x AMD EPYC 9374F 32-Core Processor 384 GiB 2x 3.2 TB NVMe Mellanox Dual Port 100 GbE	16x ProLiant DL325 Gen11 1x AMD EPYC 9374F 32-Core Processor 384 GiB 8x 3.2 TB NVMe Mellanox Dual Port 100 GbE
Connectivity: Total Rack Units:	1x Mellanox Spectrum SN2700M 32x100 GbE (cluster) 1x HPE 1620-48G (admin) 17 (1U) + 1 (1U) + 1 (1U) = 19U	

Hewlett Packard Enterprise	**HPE ProLiant DL325 Gen11**	TPCx-HS 2.0.3 TPC Pricing 2.8.0 Report Date Mar. 23, 2023

Description	Price Key	Part Number	Unit Price	Qty	Extended Price	3 Yr Maint Price
Server Hardware						
Master Node						
HPE DL325 Gen11 8SFF CTO Server	1	P54199-B21	$2,767	1	$2,767	
AMD EPYC 9374F 3.85GHz 32-core 320W Processor for HPE	1	P53710-B21	$11,087	1	$11,087	
HPE ProLiant DL3X5 Gen11 CPU Performance 1U Heat Sink Kit	1	P58457-B21	$430	1	$430	
HPE ProLiant DL3XX Gen11 1U Performance Fan Kit	1	P58462-B21	$305	7	$2,135	
HPE 32GB (1x32GB) Single Rank x4 DDR5-4800 EC8 Reg. Smart Memory	1	P50310-B21	$2,465	12	$29,580	
HPE DL3X5 Gen11 NS204i-u NVMe Hot Plug Boot Device Cable Kit	1	P57013-B21	$83	1	$83	
HPE 3.2TB NVMe MU SFF BC U.3 PM1735a SSD	1	P50230-B21	$6,476	2	$12,952	
HPE 1600W FS Plat Ht Plug LH PS Kit	1	P38997-B21	$685	2	$1,370	
Mellanox MCX623106AS-CDAT Ethernet 100Gb 2-port QSFP56 Adapter for HPE	1	P25960-B21	$4,373	1	$4,373	
HPE DL3XX Gen11 Easy Install Rail 2 Kit	1	P52351-B21	$141	1	$141	
HPE DL325Gen11 Standard Riser	1	stdDL325Gen11Riser		1		
Data Nodes						
HPE DL325 Gen11 8SFF CTO Server	1	P54199-B21	$2,767	16	$44,272	
AMD EPYC 9374F 3.85GHz 32-core 320W Processor for HPE	1	P53710-B21	$11,087	16	$177,392	
HPE ProLiant DL3X5 Gen11 CPU Performance 1U Heat Sink Kit	1	P58457-B21	$430	16	$6,880	
HPE ProLiant DL3XX Gen11 1U Performance Fan Kit	1	P58462-B21	$305	112	$34,160	
HPE 32GB (1x32GB) Single Rank x4 DDR5-4800 EC8 Reg. Smart Memory	1	P50310-B21	$2,465	192	$473,280	
HPE DL3X5 Gen11 NS204i-u NVMe Hot Plug Boot Device Cable Kit	1	P57013-B21	$83	16	$1,328	
HPE 3.2TB NVMe MU SFF BC U.3 PM1735a SSD	1	P50230-B21	$6,476	128	$828,928	
HPE 1600W FS Plat Ht Plug LH PS Kit	1	P38997-B21	$685	32	$21,920	
Mellanox MCX623106AS-CDAT Ethernet 100Gb 2-port QSFP56 Adapter for HPE	1	P25960-B21	$4,373	16	$69,968	
HPE DL3XX Gen11 Easy Install Rail 2 Kit		P52351-B21	$141	16	$2,256	
HPE DL325Gen11 Standard Riser		stdDL325Gen11Riser		16		
Other Hardware Components						
HPE 1620 48G Switch (no support available above 90 days, hence increase qty to 3)	1	JG914A	$630	3	$1,890	
HPE SN2700M 100GbE 32QSFP28 Switch with 3Y Tech Care Essential Service	1	Q2F21A	$34,510	1	$34,510	
HPE USB US Keyboard/Mouse Kit	1	631341-B21	$32	3	$96	
HPE 100Gb QSFP28 to QSFP28 3m Direct Attach Copper Cable	1	845406-B21	$685	20	$13,700	
Rack 48U 600mmx 1075mm G2 Advanced Pallet	1	P9K19A	$3,023	1	$3,023	
HPE C13 - JIS C8303 JP 100V 12Amp 2.0m Power Cord	1	AF572A	$54	37	$1,994	
HPE 3 Year Tech Care Essential DL325 GEN11 Service	1	H78S6E	$2,958	17		$50,286
			Subtotal		**$1,780,515**	**$50,286**
Other						
HP V22v G5 FHD Monitor	3	65P56AA	$130	3	$390	
USB3.0 to 1GbE Adapter	2		$14.99	20	$300	
			Subtotal		**$690**	**$0**

(Continued on next page.)

Hewlett Packard Enterprise

HPE ProLiant DL325 Gen11

TPCx-HS	2.0.3
TPC Pricing	2.8.0
Report Date	Mar. 23, 2023

(Continued from previous page.)

Description	Price Key	Part Number	Unit Price	Qty	Extended Price	3 Yr Maint Price
Server Software						
RHEL Svr Sckt/2 Gst 3yr 24x7 E-LTU	1	G3J30AAE	$3,702	17	$62,934	
Cloudera Data Platform Private Cloud Base Edition - Annual Subscription per Node for up to 16 Cores/128 GB RAM for compute or up to 48 TB for storage. Business-Level Support. (3 years support)	4	CDP-PVC-BASE-BUS	$37,200	17	$632,400	
			Subtotal		**$695,334**	**$0**
		Total Extended Price			**$2,476,539**	**$50,286**
		Total Discounts			**$623,180**	**$17,600**
		Grand Total			**$1,853,359**	**$32,686**

Pricing: 1 = HPE; 2 = BestBuy.com; 3 = hp.com; 4 = cloudera.com

* All discounts are based on US list prices and for similar quantities and configurations. A 35% discount was based on the overall specific components pricing from vendor 1 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: $1,886,046

HSph@3TB: 49.83

$ / HSph@3TB: $37,849.61

Sales contact: HPE WW Headquarters, 3000 Hanover St., Palo Alto, CA 94304-1185 (650) 857-1501 or HPE: 855-472-5233

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.

Hewlett Packard Enterprise	HPE ProLiant DL325 Gen11	TPCx-HS	2.0.3
		TPC Pricing	2.8.0
		Report Date	Mar. 23, 2023

Numerical Quantities

Performance Run – Run 2	
Scale Factor	3TB
Run Start Time	2023-02-28 22:32:21.000
Run End Time	2023-02-28 22:35:57.000
Run Elapsed Time	217.000
HSGen Start Time	2023-02-28 22:32:22.000
HSGen End Time	2023-02-28 22:33:10.000
HSGen Elapsed Time	48.849
HSSort Start Time	2023-02-28 22:33:12.000
HSSort End Time	2023-02-28 22:35:39.000
HSSort Elapsed Time	147.964
HSValidate Start Time	2023-02-28 22:35:42.000
HSValidate End Time	2023-02-28 22:35:57.000
HSValidate Elapsed Time	16.842

Repeatability Run – Run 1	
Scale Factor	3TB
Run Start Time	2023-02-28 22:28:35.000
Run End Time	2023-02-28 22:32:08.000
Run Elapsed Time	214.000
HSGen Start Time	2023-02-28 22:28:36.000
HSGen End Time	2023-02-28 22:29:23.000
HSGen Elapsed Time	47.860
HSSort Start Time	2023-02-28 22:29:25.000
HSSort End Time	2023-02-28 22:31:50.000
HSSort Elapsed Time	145.970
HSValidate Start Time	2023-02-28 22:31:53.000
HSValidate End Time	2023-02-28 22:32:08.000
HSValidate Elapsed Time	16.827

Run Reports

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

Test Run 2 Details	Total Time =	217
	Total Size =	30000000000
	Scale-Factor =	3

TPCx-HS Performance Metric (HSph@SF): 49.8338

==========

Run Report for Repeatability Run – Run 1
==========
TPCx-HS Performance Metric (HSph@SF) Report

Test Run 1 Details	Total Time =	214
	Total Size =	30000000000
	Scale-Factor =	3

TPCx-HS Performance Metric (HSph@SF): 50.5050

==========

Revision History

Date	Edition	Description
March 23, 2023	First	Initial Publication

Table of Contents

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 Hewlett Packard Enterprise Company.

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 Priced Configuration

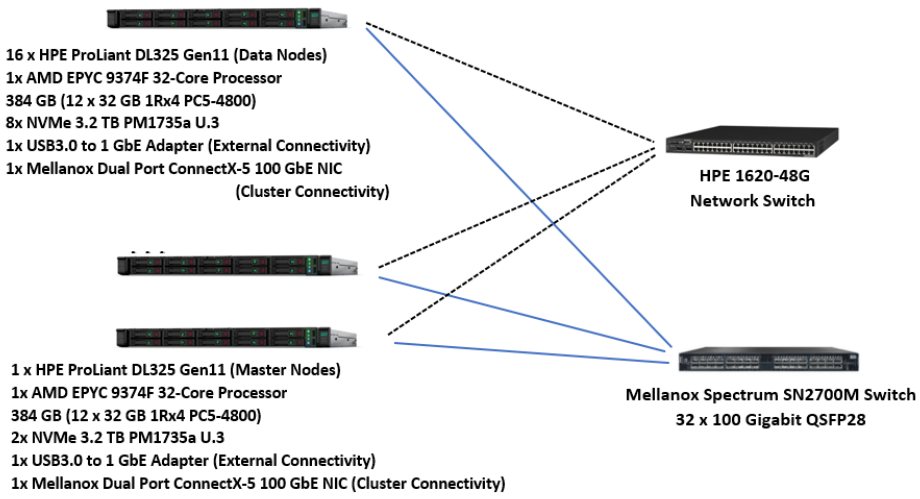

Figure 1-1 Priced Configuration

The priced configuration consists of:

- Total Nodes: 17 (1x ProLiant DL325 Gen11; 16x ProLiant DL325 Gen11)
- Total Processors/Cores/Threads: 17/544/1,088
- Total Memory: 6.38TiB
- Total Number of Storage Drives/Devices: 130
- Total Storage Capacity: 416.00TB

Server node details:

- 1x ProLiant DL325 Gen11 Servers, each with:
  - Processors/Cores/Threads: 1/32/64
  - Processor Model: AMD EPYC 9374F 32-Core Processor
  - Memory: 384 GiB
  - Drives: 2x 3.2 TB NVMe
  - Network: Mellanox Dual Port 100 GbE
- 16x ProLiant DL325 Gen11 Servers, each with:
  - Processors/Cores/Threads: 1/32/64
  - Processor Model: AMD EPYC 9374F 32-Core Processor
  - Memory: 384 GiB
  - 8rives: 2x 3.2 TB NVMe
  - Network: Mellanox Dual Port 100 GbE

Network connectivity detail:

- 1x Mellanox Spectrum SN2700M 32x100 GbE (cluster connectivity)
- 1x HPE 1620-48G (admin)

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

1.3.2 Measured Configuration

The measured configuration used 8x Mellanox Technologies MT27800 Family [ConnectX-5] MCX556A-ECAT QSFP28 adapters and 9x Mellanox Technologies MT28800 Family [ConnectX-5 Ex] MCX556A-EDAT QSFP28 adapters that were substituted with 17x Mellanox MCX623106AS-CDAT Ethernet 100Gb 2-port QSFP56 Adapter for HPE adapters in the priced 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	NVMe	nvme0n1, nvme1n1	Operating System, Root, swap, Hadoop Master
2-5	NVMe	nvme0n1	Operating System, Root, swap, Hadoop Master
2-5	NVMe	nvme0n1, nvme1n1, nvme2n1, nvme3n1, nvme4n1, nvme5n1, nvme6n1, nvme7n1	Data, Temp
6-17	NVMe	nvme0n1	Operating System, Root, swap, Hadoop Master
6-17	NVMe	nvme0n1, nvme1n1, nvme2n1, nvme3n1, nvme4n1, nvme5n1, nvme6n1, nvme7n1	Data, Temp

Table 1-1Dataset 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.

	Map/Reduce		HDFS		ZooKeeper	Spark
Node	Resource Manager	Node Manager	NameNode	DataNode	QuorumPeer	HistoryServer
1	X		X		X	X
2-5		X		X	X	
6-17		X		X		

Table 1-2 Software Component Distribution

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

CDP Private Cloud Base Edition 7.1.7 (fully HDFS compatible at the API level).

Map/Reduce implementation and corresponding version must be disclosed.

CDP Private Cloud Base Edition 7.1.7 (compatible equivalent to Hadoop 3.1.1.7.1.7.0-551).

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

Test Run 1 Details      Total Time =                  214
                        Total Size =          30000000000
                        Scale-Factor =                3

TPCx-HS Performance Metric (HSph@SF):             50.5050

=======================================================


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

Test Run 2 Details      Total Time =                  217
                        Total Size =          30000000000
                        Scale-Factor =                3

TPCx-HS Performance Metric (HSph@SF):             49.8338

=======================================================
```

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
HSDataCheck.sh	faeff3091759aac98080be4e39f7896a
TPCx-HS-master_Spark.jar	19f3ce092066e056b884a85ee92fb7fc
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)
2	3.2 TB	6.4
128	3.2 TB	409.6
Total Storage (TB)		**416.00**

Table 3-1 Storage Device Capacities

Scale Factor = 3

Data Storage Ratio = (Total Storage (TB) / SF) = **138.67**

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)) = **0.47**

Clause 4 – Metrics Related Items

4.1 HSGen Time

The HSGen time must be disclosed for Run1 and Run2.

	Run 1	Run 2
HSGen	47.860	48.849

Table 4-1 HSGen Times

4.2 HSSort Time

The HSSort time must be disclosed for Run1 and Run2.

	Run 1	Run 2
HSSort	145.970	147.964

Table 4-2 HSSort Times

4.3 HSValidate Time

The HSValidate time must be disclosed for Run1 and Run2.

	Run 1	Run 2
HSValidate	16.827	16.842

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@3TB	50.50	49.83

Table 4-5 Performance Metrics

Run 2 Price-Performance: 37,849.61 $/ HSph@3TB

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.

InfoSizing
The Right Metric For Sizing IT

TPC
Certified Auditor

Ankit Chouksey
Hewlett-Packard Enterprise
192 Mahadevapura,
Whitefield Road
Bangalore, India 560048

March 22, 2023

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

Platform:	16x HPE ProLiant DL325 Gen11 Servers (Data Nodes) 1x HPE ProLiant DL325 Gen11 Server (Master Node)
Operating System:	Red Hat Enterprise Linux 8.6
Apache Hadoop Compatible Software:	CDP Private Cloud Base Edition 7.1.7 (using Spark)

The results were:

Performance Metric	**49.83 HSph@3TB**
Run Elapsed Time	217.00 Seconds

Cluster	**16x ProLiant DL325 Gen11, 1x ProLiant DL325 Gen11 with:**		
CPUs	1x AMD® EPYC 9374F 32-Core Processor (all nodes)		
Memory	384 GiB (all nodes)		
Storage	***Qty***	***Size***	***Type***
	8	3.2 TB	NVMe (data nodes)
	2	3.2 TB	NVMe (master 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 3 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:

The measured configuration used 8x Mellanox Technologies MT27800 Family [ConnectX-5] MCX556A-ECAT QSFP28 adapters and 9x Mellanox Technologies MT28800 Family [ConnectX-5 Ex] MCX556A-EDAT QSFP28 adapters that were substituted with 17x Mellanox MCX623106AS-CDAT Ethernet 100Gb 2-port QSFP56 Adapter for HPE adapters in the priced configuration. The TPCx-HS Subcommittee approved this substitution and based on product specifications it is my opinion that this substitution has no significant effect on performance.

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

BestBuy.com

BEST BUY Menu Search Best Buy Compton Cart

Top Deals Deal of the Day Totaltech Membership Credit Cards More Account Recently Viewed Order Status Saved Items

Content creation made easy. logitech Shop now ›

Best Buy › Computers & Tablets › Computer Accessories & Peripherals › Cables & Connectors › USB Cables & Adapters

Best Buy essentials™ - USB to Ethernet Adapter - White

Model: BE-PA3U6E **SKU:** 6451327

4.8 (88 Reviews) | 3 Answered Questions

Highly rated by customers for: Connections, Price, Ethernet

A BEST BUY Brand

$14.99

15-DAY FREE & EASY RETURNS
If received today, the last day to return this item would be Jan 21.
Learn more >

Open-Box: from $11.99

Deals on related items:

3-Day Sale | USB Cables & Adapters deals | Outlet Deals

Get it today nearby

Pickup: Order now for pickup on Thu, Jan 12 at Compton
Available today at a location 6 miles away
See all pickup locations

FREE Shipping: Get it by Thu, Jan 12
See all shipping options for 90001

Add to Cart

Compare Save

Cardmember Offers

Get rewards

HP.com

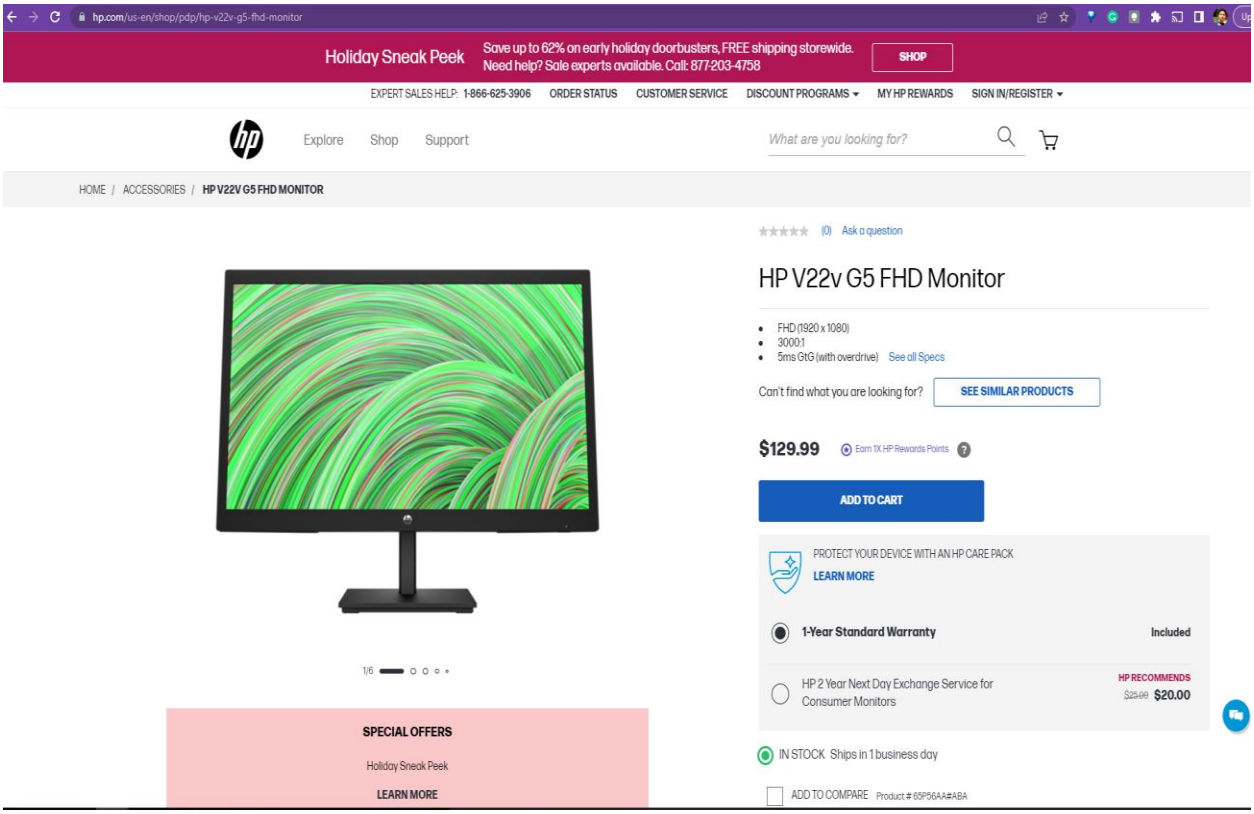
hp.com/us-en/shop/pdp/hp-v22v-g5-fhd-monitor
Holiday Sneak Peek
Save up to 62% on early holiday doorbusters, FREE shipping storewide.
Need help? Sale experts available. Call: 877-203-4758
SHOP
EXPERT SALES HELP: 1-866-625-3906
ORDER STATUS
CUSTOMER SERVICE
DISCOUNT PROGRAMS
MY HP REWARDS
SIGN IN/REGISTER
Explore
Shop
Support
What are you looking for?
HOME / ACCESSORIES / HP V22V G5 FHD MONITOR
(0) Ask a question
HP V22v G5 FHD Monitor
FHD (1920 x 1080)
3000:1
5ms GtG (with overdrive) See all Specs
Can't find what you are looking for?
SEE SIMILAR PRODUCTS
$129.99
Earn 1X HP Rewards Points
ADD TO CART
PROTECT YOUR DEVICE WITH AN HP CARE PACK
LEARN MORE
1-Year Standard Warranty
Included
HP 2 Year Next Day Exchange Service for Consumer Monitors
HP RECOMMENDS
$25.00 $20.00
IN STOCK Ships in 1 business day
ADD TO COMPARE
Product # 65P56AA#ABA
1/6
SPECIAL OFFERS
Holiday Sneak Peek
LEARN MORE

Cloudera.com

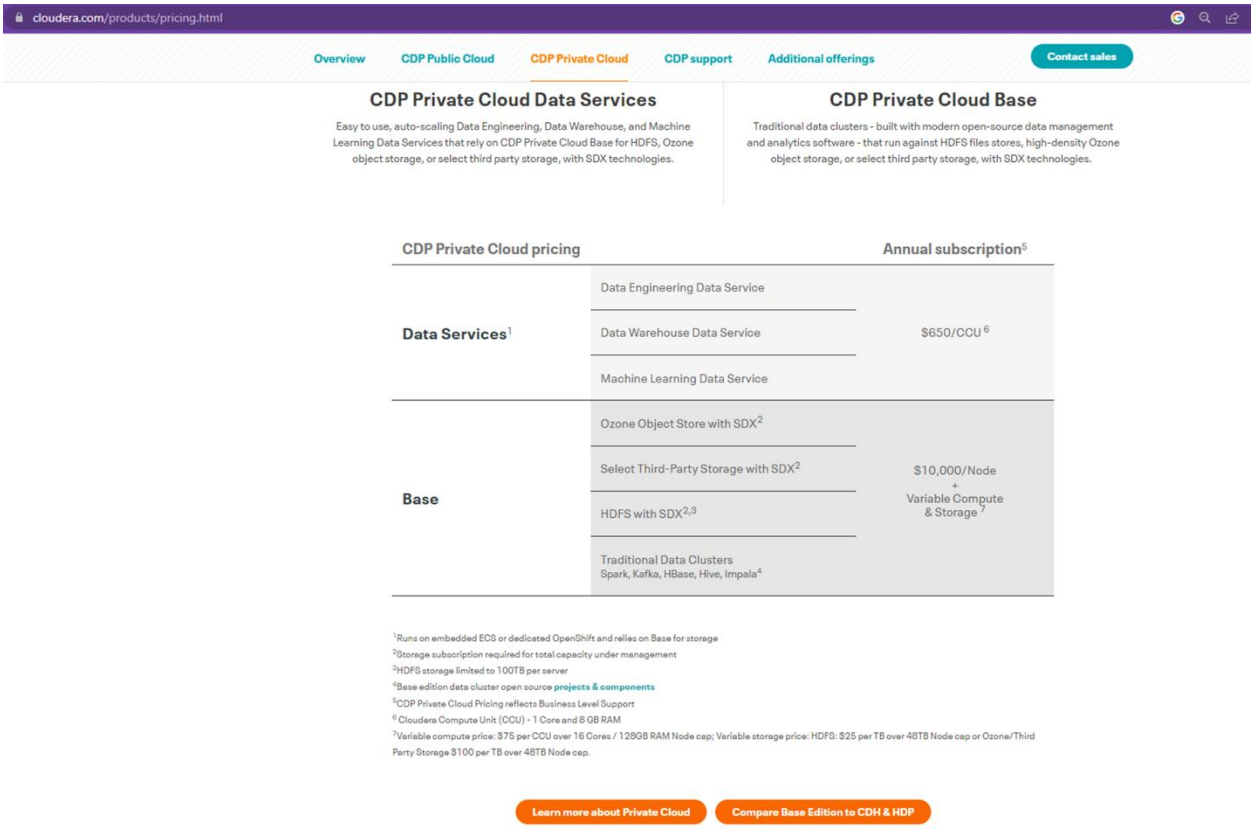

cloudera.com/products/pricing.html

Overview | CDP Public Cloud | CDP Private Cloud | CDP support | Additional offerings | Contact sales

CDP Private Cloud Data Services

Easy to use, auto-scaling Data Engineering, Data Warehouse, and Machine Learning Data Services that rely on CDP Private Cloud Base for HDFS, Ozone object storage, or select third party storage, with SDX technologies.

CDP Private Cloud Base

Traditional data clusters - built with modern open-source data management and analytics software - that run against HDFS files stores, high-density Ozone object storage, or select third party storage, with SDX technologies.

CDP Private Cloud pricing		Annual subscription[5]
Data Services[1]	Data Engineering Data Service	$650/CCU[6]
	Data Warehouse Data Service	
	Machine Learning Data Service	
Base	Ozone Object Store with SDX[2]	$10,000/Node + Variable Compute & Storage[7]
	Select Third-Party Storage with SDX[2]	
	HDFS with SDX[2,3]	
	Traditional Data Clusters Spark, Kafka, HBase, Hive, Impala[4]	

[1] Runs on embedded ECS or dedicated OpenShift and relies on Base for storage
[2] Storage subscription required for total capacity under management
[3] HDFS storage limited to 100TB per server
[4] Base edition data cluster open source projects & components
[5] CDP Private Cloud Pricing reflects Business Level Support
[6] Cloudera Compute Unit (CCU) - 1 Core and 8 GB RAM
[7] Variable compute price: $75 per CCU over 16 Cores / 128GB RAM Node cap; Variable storage price: HDFS: $25 per TB over 48TB Node cap or Ozone/Third Party Storage $100 per TB over 48TB Node cap.

Learn more about Private Cloud | Compare Base Edition to CDH & HDP