

TPC Express Benchmark™ HS Full Disclosure Report

HPE DL325 Gen10

(with 1x ProLiant DL325 Gen10 Servers; 16x ProLiant DL325 Gen10 Servers)

Running

Hortonworks Data Platform, HDP 3.1

on

Red Hat Enterprise Linux Server 7.6

TPCx-HS Version
Report Edition
Report Submitted

2.0.3
First
August 7, 2019

First Edition - August 2019

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.

Copyright© 2019 Hewlett Packard Enterprise Company

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.


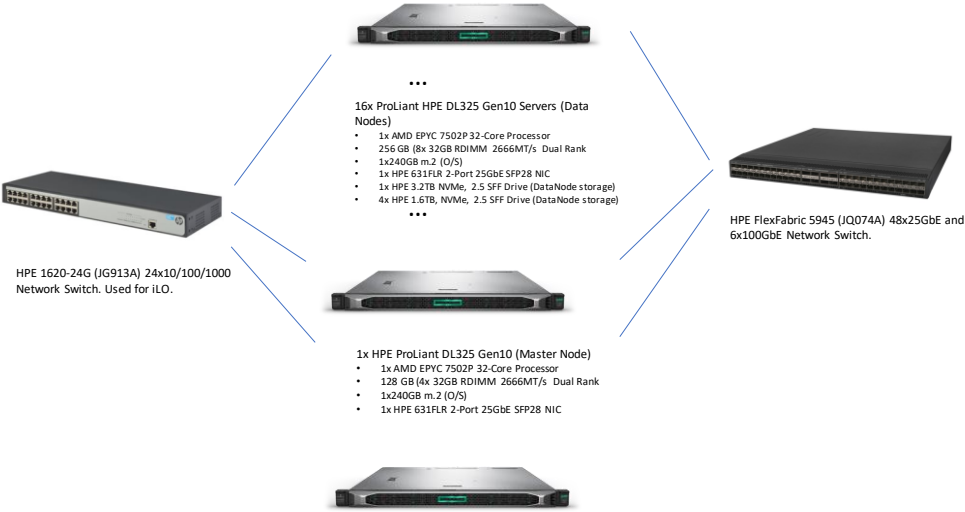
The benchmark results are summarized below.


Measured Configuration			
Company Name	Cluster Node	Hadoop Software	Operating System
HPE	ProLiant DL325 Gen10	Hortonworks Data Platform, HDP 3.1	Red Hat Enterprise Linux Server 7.6

TPC Express Benchmark™ HS Metrics			
Total System Cost	HSph@30TB	Price/Performance	Availability Date
\$704,230	25.47	\$27,649.40	August 26, 2019

Executive Summary


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

 Hewlett Packard Enterprise	<h1>HPE DL325 Gen10</h1>		TPCx-HS 2.0.3 TPC Pricing 2.4.0 Report Date Aug. 07, 2019
Availability Date August 26, 2019	TPCx-HS Performance 25.47 HSph@30TB	Price/Performance \$27,649.40 \$ / HSph@30TB	Total System Cost \$704,230 USD
System Under Test Configuration Overview			
Scale Factor 30	Hadoop Software Hortonworks Data Platform, HDP 3.1	Operating System Red Hat Enterprise Linux Server 7.6	Other Software N/A
 <p>16x ProLiant HPE DL325 Gen10 Servers (Data Nodes)</p> <ul style="list-style-type: none"> • 1x AMD EPYC 7502P 32-Core Processor • 256 GB (8x 32GB RDIMM 2666MT/s Dual Rank) • 1x240GB m.2 (O/S) • 1x HPE 631FLR 2-Port 25GbE SFP28 NIC • 1x HPE 3.2TB NVMe, 2.5 SFF Drive (DataNode storage) • 4x HPE 1.6TB, NVMe, 2.5 SFF Drive (DataNode storage) <p>1x HPE ProLiant DL325 Gen10 (Master Node)</p> <ul style="list-style-type: none"> • 1x AMD EPYC 7502P 32-Core Processor • 128 GB (4x 32GB RDIMM 2666MT/s Dual Rank) • 1x240GB m.2 (O/S) • 1x HPE 631FLR 2-Port 25GbE SFP28 NIC <p>HPE FlexFabric 5945 (IQ074A) 48x25GbE and 6x100GbE Network Switch.</p> <p>HPE 1620-24G (JG913A) 24x10/100/1000 Network Switch. Used for iLO.</p>			
Physical Storage/Scale Factor: 5.26		Scale Factor/Physical Memory: 7.27	
Total Number of Servers: Total Processors/Cores/Threads:		17x ProLiant DL325 Gen10 17/544/1,088	
Server Configuration: Processors Memory Storage Controller Storage Device		Per ProLiant DL325 Gen10 1x AMD EPYC 7502P 32-Core Processor 256 GiB (128 GiB Mgmt. node) Dual m.2 SATA, NVM Express 1x 240 GB m.2 SATA 4x HPE 1.6 TB NVMe (data nodes) 1x HPE 3.2 TB NVMe (data nodes) 1x HPE 631FLR 2-Port 25GbE SFP28 NIC	
Network Connectivity:		HPE FlexFabric 5945 Switch HPE 1620-24G 24x10/100/1000 Network Switch (used for iLO)	
Total Rack Units:		17 (1U) + 1 (1U) + 1 (1U) = 19U	

 Hewlett Packard Enterprise	<h1>HPE DL325 Gen10</h1>	TPCx-HS 2.0.3 TPC Pricing 2.4.0 Report Date Aug. 07, 2019
<p>(Continued on next page)</p>		

Description	Source	Part Number	Unit Price	Qty	Extended Price	3 Yr. Maint. Price
Server Hardware						
HPE DL325 Gen10 8SFF CTO Server	1	P04654-B21	\$1,350	17	\$22,950	
HPE DL325 Gen10 8SFF CTO Server, EPYC 1x7502P	1	P16639-L21	\$3,459	17	\$58,803	
HPE 32GB 2Rx4 PC4-2933-R	1	P00924-B21	\$1,170	132	\$154,440	
HPE 800W CS Platinum Plus AC Power Supply	1	865414-B21	\$379	17	\$6,443	
HPE DL325 Gen10 8SFF (NVMe backplane) CTO Server	1	P04662-B21	\$1,699	17	\$28,883	
HPE Dual M.2 SSD enablement option DL/ML	1	878783-B21	\$159	17	\$2,703	
HPE 240GB m.2 SATA SSD	1	P04556-B21	\$379	17	\$6,443	
HPE 3.2 TB NVMe MO003200KWVNC	1	P10224-B21	\$5,580	16	\$89,280	
HPE 1.6 TB NVMe, MO001600KWVNB	1	P10222-B21	\$3,120	64	\$199,680	
HPE Ethernet 10/25G Network Adapter 631FLR-SFP28	1	817709-B21	\$749	17	\$12,733	
HPE 3Y FC 24x7 DL325 Gen10 SVC	1	HB4G8E	\$1,355	17		\$23,035
HPE iLO Adv incl 3yr TS U E-LTU	1	E6U64ABE	\$469	17		\$7,973
HP V194 18.5" HD 1366x768 LED Monitor (1 + 2 spare)	3	V5E94A6#ABA	\$95	3	\$285	
HP PS/2 Keyboard And Mouse Bundle (1 + 2 spare)	3	H3C53AA#ABA	\$30	3	\$90	
		Subtotal			\$582,733	\$31,008
Network						
HPE 1620-24G Switch + 2 spares	1	JG913A	\$299	3	\$897	
5ft (1.5m) Cat6 Snagless Unshielded (UTP) PVC CM 17 + 2 spares	2	C6-UTPSGPVCBE	\$2	19	\$38	
HPE 5945 48SFP28 8QSFP28 Switch	1	JQ074A	\$30,970	1	\$30,970	
HPE 25Gb SFP28 to SFP28 3m DAC 34 + 10% spares	1	844477-B21	\$222	39	\$8,658	
HPE 3Y FC 24x7 FF 5945 Switch SVC	1	HB4S3E	\$8,003	1		\$8,003
		Subtotal			\$40,563	\$8,003
Rack						
HPE 42U 600x1075mm Adv G2 Kit Pltt Rack	1	P9K07A	\$1,179	1	\$1,179	
HPE 24A High Voltage Core Only Corded PDU + 2 spares	1	252663-D74	\$259	3	\$777	
		Subtotal			\$1,956	\$0

 Hewlett Packard Enterprise	<h1>HPE DL325 Gen10</h1>	TPCx-HS 2.0.3 TPC Pricing 2.4.0 Report Date Aug. 07, 2019																																																								
(Continued from previous page)																																																										
<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Description</th> <th style="text-align: left;">Source</th> <th style="text-align: left;">Part Number</th> <th style="text-align: right;">Unit Price</th> <th style="text-align: right;">Qty</th> <th style="text-align: right;">Extended Price</th> <th style="text-align: right;">3 Yr. Maint. Price</th> </tr> </thead> <tbody> <tr> <td colspan="7">Server Software</td> </tr> <tr> <td>Hortonworks 1yr 24x7</td> <td>1</td> <td>G7M27A</td> <td style="text-align: right;">\$2,000</td> <td style="text-align: right;">51</td> <td style="text-align: right;">\$102,000</td> <td></td> </tr> <tr> <td>RHEL Svr 2 Sckt/2 Gst 3yr 24x7 E-LTU</td> <td>1</td> <td>G3J30AAE</td> <td style="text-align: right;">\$3,702</td> <td style="text-align: right;">17</td> <td style="text-align: right;">\$62,934</td> <td></td> </tr> <tr> <td colspan="5" style="text-align: right;">Subtotal</td> <td style="text-align: right;">\$164,934</td> <td style="text-align: right;">\$0</td> </tr> <tr> <td colspan="5" style="text-align: right;">Total Extended Price</td> <td style="text-align: right;">\$790,186</td> <td style="text-align: right;">\$39,011</td> </tr> <tr> <td colspan="5" style="text-align: right;">Total Discounts</td> <td style="text-align: right;">\$124,967</td> <td style="text-align: right;">\$0</td> </tr> <tr> <td colspan="5" style="text-align: right;">Grand Total</td> <td style="text-align: right;">\$665,219</td> <td style="text-align: right;">\$39,011</td> </tr> </tbody> </table>			Description	Source	Part Number	Unit Price	Qty	Extended Price	3 Yr. Maint. Price	Server Software							Hortonworks 1yr 24x7	1	G7M27A	\$2,000	51	\$102,000		RHEL Svr 2 Sckt/2 Gst 3yr 24x7 E-LTU	1	G3J30AAE	\$3,702	17	\$62,934		Subtotal					\$164,934	\$0	Total Extended Price					\$790,186	\$39,011	Total Discounts					\$124,967	\$0	Grand Total					\$665,219	\$39,011
Description	Source	Part Number	Unit Price	Qty	Extended Price	3 Yr. Maint. Price																																																				
Server Software																																																										
Hortonworks 1yr 24x7	1	G7M27A	\$2,000	51	\$102,000																																																					
RHEL Svr 2 Sckt/2 Gst 3yr 24x7 E-LTU	1	G3J30AAE	\$3,702	17	\$62,934																																																					
Subtotal					\$164,934	\$0																																																				
Total Extended Price					\$790,186	\$39,011																																																				
Total Discounts					\$124,967	\$0																																																				
Grand Total					\$665,219	\$39,011																																																				
Pricing: 1 = HPE; 2 = fs.com; 3 = hp.factoryoutletstore.com * All discounts are based on US list prices and for similar quantities and configurations. A 20% 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. Sales contact: HPE WW Headquarters, 3000 Hanover St., Palo Alto, CA 94304-1185 (650) 857-1501 or HPE: 855-472-5233 Audited by Doug Johnson, InfoSizing	<table style="width: 100%;"> <tr> <td style="text-align: right;">Three-Year Cost of Ownership:</td> <td style="text-align: right;">\$704,230</td> </tr> <tr> <td style="text-align: right;">HSph@30TB:</td> <td style="text-align: right;">25.47</td> </tr> <tr> <td style="text-align: right;">\$ / HSph@30TB:</td> <td style="text-align: right;">\$27,649.40</td> </tr> </table>		Three-Year Cost of Ownership:	\$704,230	HSph@30TB:	25.47	\$ / HSph@30TB:	\$27,649.40																																																		
Three-Year Cost of Ownership:	\$704,230																																																									
HSph@30TB:	25.47																																																									
\$ / HSph@30TB:	\$27,649.40																																																									
<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>																																																										

 <p>Hewlett Packard Enterprise</p>	<h1>HPE DL325 Gen10</h1>	<p>TPCx-HS 2.0.3 TPC Pricing 2.4.0 Report Date Aug. 07, 2019</p>
--	--------------------------	--

Numerical Quantities

Performance Run – Run 2

Scale Factor	30TB
Run Start Time	2019-06-29 14:06:45.000
Run End Time	2019-06-29 15:17:21.000
Run Elapsed Time	4,240.000
HSGen Start Time	2019-06-29 14:06:47.000
HSGen End Time	2019-06-29 14:20:15.000
HSGen Elapsed Time	809.953
HSSort Start Time	2019-06-29 14:20:28.000
HSSort End Time	2019-06-29 15:07:09.000
HSSort Elapsed Time	2,802.153
HSValidate Start Time	2019-06-29 15:07:20.000
HSValidate End Time	2019-06-29 15:17:21.000
HSValidate Elapsed Time	602.329

Repeatability Run – Run 1

Scale Factor	30TB
Run Start Time	2019-06-29 12:46:11.000
Run End Time	2019-06-29 13:56:39.000
Run Elapsed Time	4,231.000
HSGen Start Time	2019-06-29 12:46:12.000
HSGen End Time	2019-06-29 12:59:50.000
HSGen Elapsed Time	819.318
HSSort Start Time	2019-06-29 13:00:02.000
HSSort End Time	2019-06-29 13:46:39.000
HSSort Elapsed Time	2,797.960
HSValidate Start Time	2019-06-29 13:46:50.000
HSValidate End Time	2019-06-29 13:56:39.000
HSValidate Elapsed Time	590.317

 Hewlett Packard Enterprise	<h1>HPE DL325 Gen10</h1>	TPCx-HS 2.0.3 TPC Pricing 2.4.0 Report Date Aug. 07, 2019
--	--------------------------	---

Run Reports

Run Report for Performance Run – Run 2

=====

TPCx-HS Performance Metric (HSph@SF) Report

Test Run 2 Details	Total Time =	4240
	Total Size =	300000000000
	Scale-Factor =	30

TPCx-HS Performance Metric (HSph@SF): 25.4733

=====

Run Report for Repeatability Run – Run 1

=====

TPCx-HS Performance Metric (HSph@SF) Report

Test Run 1 Details	Total Time =	4231
	Total Size =	300000000000
	Scale-Factor =	30

TPCx-HS Performance Metric (HSph@SF): 25.5275

=====


 Hewlett Packard Enterprise	<h1>HPE DL325 Gen10</h1>	TPCx-HS 2.0.3 TPC Pricing 2.4.0 Report Date Aug. 07, 2019
<h2>Revision History</h2>		
Date August 7, 2019	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 Configuration13

 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 Time.....17

 4.4 HSDataCheck Times17

 4.5 Performance & Price-Performance17

Auditor’s Information & Letter of Attestation18

Supporting Files Index21

Third-Party Price Quotes.....22

 fs.com.....22

 hp.factorystoreoutlet.com23

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 Measured Configuration

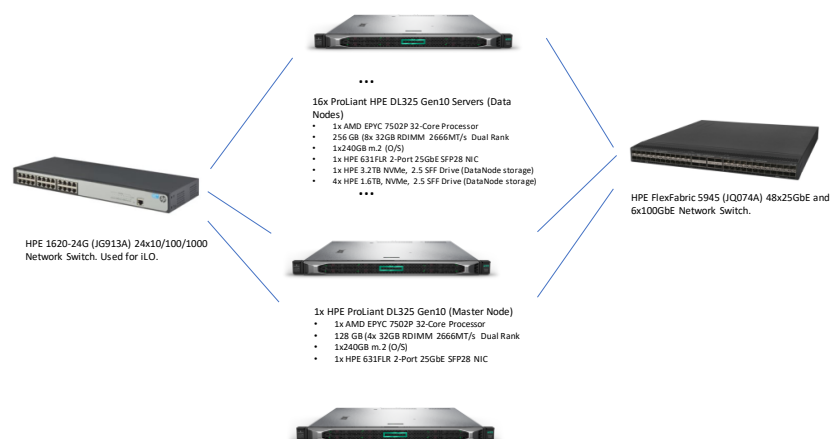


Figure 1-1 Measured Configuration

The measured configuration consisted of:

- Total Nodes: 17 (1x ProLiant DL325 Gen10; 16x ProLiant DL325 Gen10)
- Total Processors/Cores/Threads: 17/544/1,088
- Total Memory: 4.13TiB
- Total Number of Storage Drives/Devices: 97
- Total Storage Capacity: 157.68TB

Server node details:

- 17x ProLiant DL325 Gen10 Servers, each with:
 - Processors/Cores/Threads: 1/32/64
 - Processor Model: AMD EPYC 7502P 32-Core Processor
 - Memory: 256 GiB (128 GiB Mgmt. node)
 - Controller: Dual m.2 SATA NVM Express
 - Drives:
 - 1x 240 GB m.2 SATA
 - 4x HPE 1.6 TB NVMe (data nodes)
 - 1x HPE 3.2 TB NVMe (data nodes)
 - Network: 1x HPE 631FLR 2-Port 25GbE SFP28 NIC

Network connectivity detail:

- HPE FlexFabric 5945 Switch
- HPE 1620-24G 24x10/100/1000 Network Switch (used for iLO)

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	m.2	sda	Operating System, Root, Swap, Hadoop Master
2-3	m.2	sda	Operating System, Root, Swap, Hadoop Master
2-3	NVMe	nvme0n1, nvme1n1, nvme2n1, nvme3n1, nvme4n1	Data, Temp
4-17	m.2	sda	Operating System, Root, Swap, Hadoop Master
4-17	NVMe	nvme0n1, nvme1n1, nvme2n1, nvme3n1, nvme4n1	Data, Temp

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
	Resource Manager	Node Manager	NameNode	DataNode	QuorumPeer
1	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.

Hortonworks Data Platform, HDP 3.1 (fully HDFS compatible at the API level).

Map/Reduce implementation and corresponding version must be disclosed.

Hortonworks Data Platform, HDP 3.1 (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 – Repeatability Run
=====
TPCx-HS Performance Metric (HSph@SF) Report

Test Run 1 Details      Total Time =                4231
                        Total Size =                300000000000
                        Scale-Factor =                30

TPCx-HS Performance Metric (HSph@SF):                25.5275

=====
```

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

Test Run 2 Details      Total Time =                4240
                        Total Size =                300000000000
                        Scale-Factor =                30

TPCx-HS Performance Metric (HSph@SF):                25.4733

=====
```

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                bcf0b946a49d1249c9da174b5d9805f1
TPCx-HS-master_MR2.jar         492cbc51a1a60c28b43d96c79d08683d
TPCx-HS-master.sh              c619a0819571ecd00cd75d2b76ba8c64
```

2.4 Benchmark Kit Changes

HSDDataCheck.sh was modified to collect HDFS Erasure Coding data.

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)
17	240 GB	4.08
64	1.6 TB	102.40
16	3.2 TB	51.20
Total Storage (TB)		157.68

Table 3-1 Storage Device Capacities

Scale Factor = 30

Data Storage Ratio = (Total Storage (TB) / SF) = 5.26

3.2 Memory Ratio

The Scale Factor to memory ratio must be disclosed.

Total Configured Memory (TiB) = 4.13

Scale Factor to Memory Ratio = (SF / Total Memory(TiB)) = 7.27

Clause 4 – Metrics Related Items

4.1 HSGen Time

The HSGen time must be disclosed for Run1 and Run2.

	Run 1	Run 2
HSGen	819.318	809.953

Table 4-1 HSGen Times

4.2 HSSort Time

The HSSort time must be disclosed for Run1 and Run2.

	Run 1	Run 2
HSSort	2,797.960	2,802.153

Table 4-2 HSSort Times

4.3 HSValidate Time

The HSValidate time must be disclosed for Run1 and Run2.

	Run 1	Run 2
HSValidate	590.317	602.329

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)	12.000	13.000
HSDataCheck (post-sort)	11.000	11.000

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@30TB	25.52	25.47

Table 4-5 Performance Metrics

Run 2 Price-Performance: 27,649.40 \$/ HSph@30TB

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.



Mr. Craig A. Estep
 Hewlett Packard Enterprise
 11445 Compaq Center Dr West
 Houston, TX 77070

July 30, 2019

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

Platform: HPE DL 325 Gen10 (with 17x HPE DL 325 Gen10 Servers)
 Operating System: Red Hat Enterprise Linux Server 7.6
 Apache Hadoop Hortonworks Data Platform, HDP 3.1
 Compatible Software:

The results were:

Performance Metric 25.47 HSph@30TB
 Run Elapsed Time 4,240.00 Seconds

Cluster	17x HPE DL 325 Gen10 Servers, each node with:		
CPU	1 x AMD EPYC 7502P (2.50 GHz, 32-core, 16 MB L3)		
Memory	256 GiB (16 data nodes), 128 GiB (1 Mgmt. node)		
Storage	Qty	Size	Type
	1	240GB	SSD SATA (All nodes)
	4	1.6TB	NVMe (16 data nodes)
	1	3.2TB	NVME (16 data nodes)

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 30TB
- The generated dataset and the sorted dataset were erasure coded with a policy of RS-6-3-1024k

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

- 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 required data protection was provided by HDFS Erasure Coding rather than the default three-way data replication. A policy of RS-6-3-1024k was used. Therefore, each block group consisted of 6 data blocks and 3 parity blocks. Each block within a given block group was placed on a different node thus ensuring the required data protection.

In order to collect the necessary data for auditing, the HSDataCheck.sh script was modified. In accordance with the TPCx-HS Standard Specification, this change received prior approval from the TPCx-HS subcommittee.

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

fs.com

The screenshot shows the fs.com website interface. At the top, there are navigation links for 'All Categories', 'Solutions', 'Resources', 'About Us', and 'Support'. A search bar is located on the right side of the header. The main content area displays a product page for a '5ft (1.5m) Cat6 Snagless Unshielded (UTP) PVC CM Ethernet Network Patch Cable, Blue' (product ID #70725). The product is priced at 'US\$ 2.00'. Below the price, there are options for 'Length' (3ft, 6in, 1ft, 2ft, 4ft, 5ft, 6ft, 7ft, 8ft, 10ft, 12ft, 15ft, 20ft, 25ft, 35ft, 50ft, 75ft, 100ft, 150ft, 200ft) and 'Color' (Blue, Black, Green, Red, Yellow, Purple). The 'Size' is set to '1 piece' (40 pcs/case). The page also indicates '17928 in Stock, U.S. Warehouse' and 'Free Shipping via FedEx Ground® on orders over US\$ 79'. An 'Add to Cart' button is visible at the bottom of the product details.

hp.factorystoreoutlet.com

The screenshot shows a web browser window with multiple tabs. The active tab is 'HP Monitors | Elite'. The address bar shows the URL 'https://hp.factoryoutletstore.com/cat/64664/HP-Monitors.html'. The page features a navigation sidebar on the left with categories like 'HP Elite Display', 'HP Z Display', 'Accessories', 'Refurbished Computers', and 'Clearance Products'. A 'Partner First Gold' badge is visible. The main content area displays a product listing for 'HP Monitor V194 - 18.5" Display'. The product is marked as a 'Best Seller In Category'. Key features listed include: 'Business-class presentation features', 'Brand New Includes One Year Warranty', 'Returnable: 30 Day Guarantee', and 'Replaces Model 5RD65A8'. Technical specifications include: Display Size: 18.5", 1366 x 768 Resolution, Aspect Ratio: 16:9, LED Backlit, Brightness: 200 cd/m2, and Anti-Glare Screen. The product number is V5E94A6#ABA. The price is shown as \$81.95, with a crossed-out original price of \$94.99. A 'Ships Free' badge and 'BRAND NEW' label are present. An 'Add to Cart' button is visible, along with a 'Pay Over Time' option starting at \$3.79/month. A 'More Details' link is also present.

The screenshot shows a web browser window with multiple tabs. The active tab is 'Hewlett Packard H3C53AA#ABA'. The address bar shows the URL 'https://hp.factoryoutletstore.com/details/229933/hewlett-packard-h3c53aaaba.html?category_id=64694&catalogtermid=205490'. The page features a navigation sidebar on the left with categories like 'Hot Deals', 'Business Printers', 'Ink & Toner', 'Printer Accessories', 'Business Scanners', 'Batteries', 'Hard Drives', 'Laptops', 'Tablets', 'Desktops', 'Z Workstations', 'Monitors', 'Accessories', 'Batteries', 'Calculators', 'Cables', and 'Carrying Cases'. The main content area displays a product listing for 'HP C2500 Keyboard and Mouse Combo H3C53AA#ABA'. The product is marked as 'In Stock'. Key features listed include: 'Brand New Includes One Year Limited Warranty', 'Returnable Item: 30 Day Guarantee', and 'Ships Free'. The price is shown as \$24.95, with a crossed-out original price of \$29.95. A 'Ships Free' badge and 'BRAND NEW' label are present. An 'Add to Cart' button is visible, along with a 'Pay Over Time' option starting at \$3.99/month. A 'More Details' link is also present. A '120 DAY MONEY-BACK GUARANTEE' banner is visible at the top of the page.