

TPC Express Benchmark™ HS Full Disclosure Report

Dell PowerEdge R6415

(with 17x Dell PowerEdge R6415 Servers)

Running

Cloudera Enterprise Edition 6.0

on

Red Hat Enterprise Linux Server Release 7.5

Second Edition - November 2018

Dell Inc. (Dell), 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.

Dell and the Dell Logo are trademarks of Dell 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 Dell and any other company

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

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

Copyright© 2018 Dell 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 v2.0.3.






The benchmark results are summarized below.


Measured Configuration			
Company Name	Cluster Node	Hadoop Software	Operating System
Dell	Dell PowerEdge R6415	Cloudera Enterprise Edition 6.0	Red Hat Enterprise Linux Server Release 7.5


TPC Express Benchmark™ HS Metrics			
Total System Cost	HSph@1TB	Price/Performance	Availability Date
\$507,648	11.01	\$46,107.91	October 9, 2018


Executive Summary

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

	<h1>Dell PowerEdge R6415</h1>		TPCx-HS v2.0.3
			TPC Pricing v2.3.0
			Report Date Nov. 16, 2018
Availability Date October 9, 2018	TPCx-HS Performance 11.01 HSp@1TB	Price/Performance \$46,107.91 \$ / HSp@1TB	Total System Cost \$507,648 USD
System Under Test Configuration Overview			
Scale Factor 1	Hadoop Software Cloudera Enterprise Edition 6.0	Operating System Red Hat Enterprise Linux Server Release 7.5	Other Software Spark 2.2.0-cdh6.0
<div style="display: flex; justify-content: space-between; align-items: flex-start;"> <div style="width: 45%;">  <p>...</p> <p>16x Dell PowerEdge R6415 Servers (Data Nodes)</p> <ul style="list-style-type: none"> • 1x AMD EPYC 7551P 32-Core Processor • 256 GB (8x 32GB RDIMM 2666MT/s Dual Rank) • 1x 120GB SSD SATA 2.5in Hot-plug Drive • 1x Mellanox Dual Port 25GbE SFP28 NIC • 2x Dell 1.6TB, NVMe, 2.5 SFF Drive <p>...</p>  <p>...</p>  <p>1x Dell PowerEdge R6415 Server (Master Node)</p> <ul style="list-style-type: none"> • 1x AMD EPYC 7551P 32-Core Processor • 256 GB (8x 32GB RDIMM 2666MT/s Dual Rank) • 1x 120GB SSD SATA 2.5in Hot-plug Drive • 1x Mellanox Dual Port 25GbE SFP28 NIC </div> <div style="width: 50%; text-align: center;">  <p>Dell S-Series (S5148F-ON) 48x25GbE and 6x100GbE Network Switch.</p> </div> </div>			
Physical Storage/Scale Factor: 53.24		Scale Factor/Physical Memory: 0.24	
Total Number of Servers:		17x Dell PowerEdge R6415	
Total Processors/Cores/Threads:		17/544/1,088	
Server Configuration:		Per Dell PowerEdge R6415	
Processors		1x AMD EPYC 7551P 32-Core	
Memory		256 GiB	
Storage Controller		Perc H740P	
Storage Device		1x 120 GB SSD SATA (all nodes)	
		2x Dell 1.6TB NVMe (Data Nodes)	
Network		1x Mellanox Dual Port 25GbE SFP28 NIC	
Connectivity:		Dell S-Series (S5148F-ON) Network Switch	
Total Rack Units:		(17xR6415) + (1xS5148F) = (17x1) + (1x1) = 18RU	

	<h1 style="margin: 0;">Dell PowerEdge R6415</h1>				TPCx-HS	v2.0.3	
					TPC Pricing	v2.3.0	
					Report Date	Nov. 16, 2018	
Description		Part Number	Key	Unit Price	Qty	Extended Price	3 yr. Maint. Price
HARDWARE COMPONENTS							
PowerEdge R6415 Server		210-ANJO	1	\$23,989.00	16	\$383,824.00	
PowerEdge R6415/R7415 Motherboard		384-BBSR	1	\$0.00	16		
No Trusted Platform Module		461-AADZ	1	\$0.00	16		
2.5" Chassis with up to 10 Hard Drives, including up to 8 SAS/SATA or 9 NVME Drives		321-BDFY	1	\$0.00	16		
PowerEdge R6415 Shipping		340-BTFM	1	\$0.00	16		
PowerEdge R6415 x4 or x10 Drive Shipping Material		343-BBGL	1	\$0.00	16		
AMD EPYC™ 7551P 2.00GHz/2.55GHz, 32C/64T, 64M Cache (180W) DDR4-2666		338-BNCG	1	\$0.00	16		
Standard Heatsink		412-AALH	1	\$0.00	16		
2666MT/s RDIMMs		370-ADNU	1	\$0.00	16		
Performance Optimized		370-AAIP	1	\$0.00	16		
Unconfigured RAID		780-BCDS	1	\$0.00	16		
PERC H740P RAID Controller, 8GB NV Cache, Mini card		405-AAMS	1	\$0.00	16		
No Operating System		619-ABVR	1	\$0.00	16		
No Media Required		421-5736	1	\$0.00	16		
iDRAC9, Enterprise		385-BBKT	1	\$0.00	16		
iDRAC Group Manager, Disabled		379-BCQY	1	\$0.00	16		
iDRAC, Factory Generated Password		379-BCSF	1	\$0.00	16		
Riser Config 1, 2 x 16 LP		330-BBIV	1	\$0.00	16		
On-Board LOM		542-BBBP	1	\$0.00	16		
No Internal Optical Drive		429-AAIQ	1	\$0.00	16		
Dual, Hot Plug, Redundant Power Supply (1+1), 550W		450-AGZB	1	\$0.00	16		
No Bezel		350-BBBW	1	\$0.00	16		
Dell EMC Luggage Tag		350-BBME	1	\$0.00	16		
No Quick Sync		350-BBKR	1	\$0.00	16		
Performance BIOS Settings		384-BBBL	1	\$0.00	16		
UEFI BIOS Boot Mode with GPT Partition		800-BBDM	1	\$0.00	16		
ReadyRails Sliding Rails With Cable Management Arm		770-BCKT	1	\$0.00	16		
No Systems Documentation, No OpenManage DVD Kit		631-AACK	1	\$0.00	16		
US Order		332-1286	1	\$0.00	16		
Dell Hardware Limited Warranty Plus On-Site Service		816-0779	1	\$200.00	16		\$3,200.00
ProSupport Mission Critical: 7x24 HW / SW Technical Support and Assistance, 3 Years		816-0779	1	\$1,438.00	16		\$23,008.00
3 Years ProSupport and Mission Critical 4Hr Onsite Service		816-0780	1	\$262.00	16		\$4,192.00
Thank you choosing Dell ProSupport. For tech support, visit //www.dell.com/support or call 1-800- 945-3355		989-3439	1	\$0.00	16		
On-Site Installation Declined		900-9997	1	\$0.00	16		
32GB RDIMM 2666MT/s Dual Rank		370-ADNF	1	\$0.00	128		
120GB SSD SATA Boot 6Gbps 512n 2.5in Hot-plug Drive, 1 DWPD, 219 TBW		400-ASEG	1	\$0.00	16		
Dell 1.6TB, NVMe, Mixed Use Express Flash, 2.5 SFF Drive, U.2, PM1725a with Carrier		400-AUMP	1	\$0.00	32		
Mellanox ConnectX-4 Lx Dual Port 25GbE SFP28 Network Adapter, Low Profile		406-BBLD	1	\$0.00	16		
C13 to C14, PDU Style, 12 AMP, 6.5 Feet (2m) Power Cord, North America		492-BBDI	1	\$0.00	32		
Dell EMC S5148F-ON Switch, 48x 25GbE, 6x 100GbE QSFP28, 10 to PSU, 2 PSU, OS10		210-ANCK	1	\$25,279.00	1	\$25,279.00	
OS10 Enterprise S5148F-ON		634-BMIF	1	\$0.00	1		
Dell Networking, Cable, SFP28 to SFP28, 25GbE, Passive Copper Twinax Direct Attach Cable, 2 Meter		470-ACET	1	\$0.00	1		
Dell EMC S5148 Series User Guide		343-BBFV	1	\$0.00	1		
US No Canada Ship Charge		332-1286	1	\$0.00	1		
Force10, Power Cord, 125V, 15A, 10 Feet, NEMA 5-15/C13, S-Series		450-AAFH	1	\$0.00	1		
Dell Hardware Limited Warranty 1 Year		814-8268	1	\$0.00	1		\$422.00
Thank you choosing Dell ProSupport. For tech support, visit //www.dell.com/support or call 1-800- 945-3355		989-3439	1	\$0.00	1		\$0.00
Dell Limited Hardware Warranty Extended Year(s)		975-3461	1	\$0.00	1		\$0.00
Mission Critical Package: 4-Hour 7X24 On-Site Service with Emergency Dispatch, 2 Year Extended		814-8278	1	\$423.00	1		\$423.00
ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Year		814-8288	1	\$5,956.00	1		\$5,956.00
Info 3rd Party Software Warranty provided by Vendor		997-6306	1	\$0.00	1		\$0.00
Mission Critical Package: 4-Hour 7X24 On-Site Service with Emergency Dispatch, 1 Year		814-8277	1	\$111.00	1		\$111.00
On-Site Installation Declined		900-9997	1	\$0.00	1		\$0.00
APC NetShelter SX 24U 600mm x 1070mm Deep Enclosure		A7067508	1	\$1,254.98	1	\$1,254.98	
Rack PDU, Basic, Zero U, 15A, 120V, 5-15 input, (14) 5-15 output		A7541364	1	\$0.00	1		
Logitech MK120 Keyboard and Mouse		A6999510	1	\$13.99	1	\$13.99	
Dell 24 Monitor		210-AIWG	1	\$169.99	1	\$169.99	
Subtotal						\$410,541.96	\$37,312.00

	<h1 style="margin: 0;">Dell PowerEdge R6415</h1>		TPCx-HS	v2.0.3			
			TPC Pricing	v2.3.0			
			Report Date	Nov. 16, 2018			
Description		Part Number	Key	Unit Price	Qty	Extended Price	3 yr. Maint. Price
PowerEdge R6415 Server		210-ANJO	1	\$17,231.00	1	\$17,231.00	
PowerEdge R6415/R7415 Motherboard		384-BBSR	1	\$0.00	1		
No Trusted Platform Module		461-AADZ	1	\$0.00	1		
2.5" Chassis with up to 10 Hard Drives, including up to 8 SAS/SATA or 9 NVME Drives		321-BDFY	1	\$0.00	1		
PowerEdge R6415 Shipping		340-BTFM	1	\$0.00	1		
PowerEdge R6415 x4 or x10 Drive Shipping Material		343-BBGL	1	\$0.00	1		
AMD EPYC™ 7511P 2.00GHz/2.55GHz, 32C/64T, 64M Cache (180W) DDR4-2666		338-BNCG	1	\$0.00	1		
Standard Heatsink		412-AALH	1	\$0.00	1		
2666MT/s RDIMMs		370-ADNU	1	\$0.00	1		
Performance Optimized		370-AAIP	1	\$0.00	1		
Unconfigured RAID		780-BCDS	1	\$0.00	1		
PERC H740P RAID Controller, 8GB NV Cache, Mini card		405-AAMS	1	\$0.00	1		
No Operating System		619-ABVR	1	\$0.00	1		
No Media Required		421-5736	1	\$0.00	1		
iDRAC9,Enterprise		385-BBKT	1	\$0.00	1		
iDRAC Group Manager, Disabled		379-BCQY	1	\$0.00	1		
iDRAC,Factory Generated Password		379-BCSF	1	\$0.00	1		
Riser Config 1, 2 x 16 LP		330-BBIV	1	\$0.00	1		
On-Board LOM		542-BBBP	1	\$0.00	1		
No Internal Optical Drive		429-AAIQ	1	\$0.00	1		
Dual, Hot Plug, Redundant Power Supply (1+1), 550W		450-AGZB	1	\$0.00	1		
No Bezel		350-BBBW	1	\$0.00	1		
Dell EMC Luggage Tag		350-BBME	1	\$0.00	1		
No Quick Sync		350-BBKR	1	\$0.00	1		
Performance BIOS Settings		384-BBBL	1	\$0.00	1		
UEFI BIOS Boot Mode with GPT Partition		800-BBDM	1	\$0.00	1		
ReadyRails Sliding Rails With Cable Management Arm		770-BCKT	1	\$0.00	1		
No Systems Documentation, No OpenManage DVD Kit		631-AACK	1	\$0.00	2		
US Order		332-1286	1	\$0.00	1		
Dell Hardware Limited Warranty Plus On-Site Service		816-0779	1	\$200.00	1		\$200.00
ProSupport Mission Critical: 7x24 HW / SW Technical Support and Assistance, 3 Years		816-0779	1	\$1,438.00	1		\$1,438.00
3 Years ProSupport and Mission Critical 4Hr Onsite Service		816-0780	1	\$262.00	1		\$262.00
Thank you choosing Dell ProSupport. For tech support, visit //www.dell.com/support or call 1-800- 945-3355		989-3439	1	\$0.00	1		\$0.00
On-Site Installation Declined		900-9997	1	\$0.00	1		
32GB RDIMM 2666MT/s Dual Rank		370-ADNF	1	\$0.00	8		
120GB SSD SATA Boot 6Gbps 512n 2.5in Hot-plug Drive, 1 DWPD, 219 TBW		400-ASEG		\$0.00	1		
Mellanox ConnectX-4 Lx Dual Port 25GbE SFP28 Network Adapter, Low Profile		406-BBLD	1	\$0.00	1		
C13 to C14, PDU Style, 12 AMP, 6.5 Feet (2m) Power Cord, North America		492-BBDI	1	\$0.00	2		
Sub Total						\$17,231.00	\$1,900.00
HARDWARE COMPONENTS						Subtotal	\$427,772.96
SOFTWARE COMPONENTS							
Cloudera Enterprise Data Engineering Edition, Node License, 24x7, 3YR		CEDEN-GOLD	1	\$14,769.24	17		\$251,077.08
Red Hat Enterprise Linux Non Factory Install, x64,Reqs Lic &Sub Selection		421-4727	1	\$0.00	17	\$0.00	
Red Hat Enterprise Linux,1-2SKT,3yr Premium Subscription,1 Virtual Guest		421-5721	1	\$3,702	17		\$62,934.00
SOFTWARE COMPONENTS						Subtotal	\$0.00
Total						\$427,772.96	\$314,011.08
Large Purchase Discount (35%)*						-149,720.54	-123,628.08
Pricing: 1 = Dell		Three-Year Cost of Ownership:		\$507,648			
* Discount applies to all line items where Key = 1. Discount based upon total system cost as purchased by a regular customer.		HSph@1TB:		11.01			
Audited by Doug Johnson, InfoSizing		\$/ HSph@1TB:		\$46,107.91			
<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>							

	<h2 style="margin: 0;">Dell PowerEdge R6415</h2>	TPCx-HS	v2.0.3
		TPC Pricing	v2.3.0
		Report Date	Nov. 16, 2018

Numerical Quantities	
----------------------	--

Performance Run – Run 1	
Scale Factor	1TB
Run Start Time	2018-09-29 08:47:09.000
Run End Time	2018-09-29 08:52:34.000
Run Elapsed Time	327.000
HSGen Start Time	2018-09-29 08:48:33.000
HSGen End Time	2018-09-29 08:48:33.000
HSGen Elapsed Time	84.340
HSSort Start Time	2018-09-29 08:52:06.000
HSSort End Time	2018-09-29 08:52:06.000
HSSort Elapsed Time	211.093
HSValidate Start Time	2018-09-29 08:52:34.000
HSValidate End Time	2018-09-29 08:52:34.000
HSValidate Elapsed Time	26.644

Repeatability Run – Run 2	
Scale Factor	1TB
Run Start Time	2018-09-29 08:53:40.000
Run End Time	2018-09-29 08:59:05.000
Run Elapsed Time	327.000
HSGen Start Time	2018-09-29 08:55:05.000
HSGen End Time	2018-09-29 08:55:05.000
HSGen Elapsed Time	85.757
HSSort Start Time	2018-09-29 08:58:36.000
HSSort End Time	2018-09-29 08:58:36.000
HSSort Elapsed Time	208.992
HSValidate Start Time	2018-09-29 08:59:05.000
HSValidate End Time	2018-09-29 08:59:05.000
HSValidate Elapsed Time	26.776



Dell PowerEdge R6415

TPCx-HS	v2.0.3
TPC Pricing	v2.3.0
Report Date	Nov. 16, 2018

Run Reports

Run Report for Performance Run – Run 1

TPCx-HS Performance Metric (HSph@SF) Report

Test Run 1 Details	Total Time =	327
	Total Size =	10000000000
	Scale-Factor =	1

TPCx-HS Performance Metric (HSph@SF): 11.0132

Run Report for Repeatability Run – Run 2

TPCx-HS Performance Metric (HSph@SF) Report

Test Run 2 Details	Total Time =	327
	Total Size =	10000000000
	Scale-Factor =	1

TPCx-HS Performance Metric (HSph@SF): 11.0132


	Dell PowerEdge R6415	TPCx-HS v2.0.3 TPC Pricing v2.3.0 Report Date Nov. 16, 2018
Revision History		
Date	Edition	Description
October 9, 2018	First	Initial Publication
November 16, 2018	Second	Corrected Cloudera line item Unit Price. Corrected total number of storage devices.

Table of Contents

Abstract..... 3

Executive Summary 3

Table of Contents.....11

Clause 0 – Preamble12

 0.1 TPC Express Benchmark™ HS Overview.....12

Clause 1 – General Items13

 1.1 Test Sponsor13

 1.2 Parameter Settings13

 1.3 Configuration Diagrams13

 1.3.1 Measured Configuration.....14

 1.3.2 Priced Configuration15

 1.4 Dataset Distribution.....15

 1.5 Software Components Distribution.....15

Clause 2 – Workload Related Items.....16

 2.1 Hardware & Software Tunables16

 2.2 Run Report16

 2.3 Benchmark Kit Identification.....16

 2.4 Benchmark Kit Changes16

Clause 3 – SUT Related Items.....17

 3.1 Data Storage Ratio17

 3.2 Memory Ratio17

Clause 4 – Metrics Related Items18

 4.1 HSGen Time.....18

 4.2 HSSort Time18

 4.3 HSValidate Time.....18

 4.4 HSDataCheck Times18

 4.5 Performance & Price-Performance18

Auditor’s Information & Letter of Attestation19

Supporting Files Index22

Third-Party Price Quotes.....23

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 Dell 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

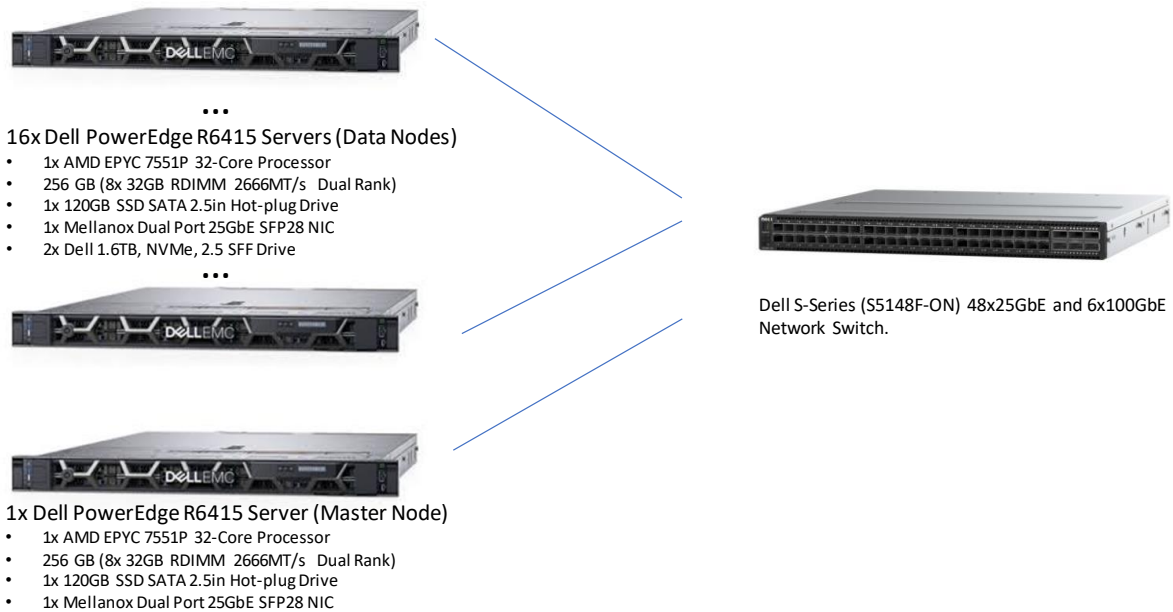


Figure 1-1 Measured Configuration

The measured configuration consisted of:

- Total Nodes: 17 (17x Dell PowerEdge R6415; 16x Dell PowerEdge R6415)
- Total Processors/Cores/Threads: 17/544/1,088
- Total Memory: 4.25TiB
- Total Number of Storage Drives/Devices: 49
- Total Storage Capacity: 53.24TB

Server node details:

- 17x Dell PowerEdge R6415 Servers, each with:
 - Processors/Cores/Threads: 1/32/64
 - Processor Model: AMD EPYC 7551P 32-Core
 - Memory: 256 GiB
 - Controller: Perc H740P
 - Drives:
 - 1x 120 GB SSD SATA (all nodes)
 - 2x Dell 1.6TB NVMe (Data Nodes)
 - Network: 1x Mellanox Dual Port 25GbE SFP28 NIC

Network connectivity detail:

- Dell S-Series (S5148F-ON) Network Switch

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	Perc H740p	sda	Operating System, Root, Swap, Hadoop Master
2-3	Perc H740p	sda	Operating System, Root, Swap, Hadoop Master
2-3	NVMe	nvme0n1, nvme1n1	Data, Temp
4-17	Perc H740p	sda	Operating System, Root, Swap, Hadoop Master
4-17	NVMe	nvme0n1, nvme1n1	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	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 Enterprise Edition 6.0 (fully HDFS compatible at the API level).

Map/Reduce implementation and corresponding version must be disclosed.

Cloudera Enterprise Edition 6.0 (compatible equivalent to Hadoop 3.0.0 and Spark 2.2.0).

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 =                327
                        Total Size =                10000000000
                        Scale-Factor =                1

TPCx-HS Performance Metric (HSph@SF):                11.0132

=====
```

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

Test Run 2 Details      Total Time =                327
                        Total Size =                10000000000
                        Scale-Factor =                1

TPCx-HS Performance Metric (HSph@SF):                11.0132

=====
```

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_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.

Qty	Capacity	Total	Total (TB)
17	120GB	2,040GB	2.04
32	1.6TB	51.2TB	51.20
Total Storage (TB)			53.24

Table 3-1 Storage Device Capacities

Scale Factor = 1

Data Storage Ratio = (Total Storage (TB) / SF) = 53.24

3.2 Memory Ratio

The Scale Factor to memory ratio must be disclosed.

Total Configured Memory (TiB) = 4.25

Scale Factor to Memory Ratio = (SF / Total Memory(TiB)) = 0.24

Clause 4 – Metrics Related Items

4.1 HSGen Time

The HSGen time must be disclosed for Run1 and Run2.

	Run 1	Run 2
HSGen	84.340	85.757

Table 4-1 HSGen Times

4.2 HSSort Time

The HSSort time must be disclosed for Run1 and Run2.

	Run 1	Run 2
HSSort	211.093	208.992

Table 4-2 HSSort Times

4.3 HSValidate Time

The HSValidate time must be disclosed for Run1 and Run2.

	Run 1	Run 2
HSValidate	26.644	26.776

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)	213.000	211.000
HSDataCheck (post-sort)	28.000	29.000

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@1TB	11.01	11.01

Table 4-5 Performance Metrics

Run 1 Price-Performance: 46,107.91 \$/ HSph@1TB

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. Nicholas Wakou
Dell Inc.
701 E. Parmer Ln. Bldg. 2
Austin, TX 78753

November 16, 2018

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

Platform: Dell PowerEdge R6415 (with 17x Dell PowerEdge R6415 Servers)
Operating System: Red Hat Enterprise Linux Server 7.5
Apache Hadoop Cloudera Enterprise Edition V6.0 (Spark)
Compatible Software:

The results were:

Performance Metric 11.01 HSph@1TB

Run Elapsed Time 327.00 Seconds

Cluster 17x Dell PowerEdge R6415 Servers, each node with:

CPU	1 x AMD EPYC 7551P (2.00 GHz, 32-core, 64 MB L3)		
Memory	256 GB		
Storage	Qty	Size	Type
	1	120GB	SSD SATA (All nodes)
	2	1.6TB	NVMe (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 1TB
- The generated dataset and the sorted dataset were replicated 3-ways

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 initial audit was completed on October 8, 2018. A correction to the Cloudera line item Unit Price was reviewed as part of the required FDR revision process.

Respectfully Yours,

A handwritten signature in cursive script that reads "Doug Johnson". The signature is written in black ink and has a long, sweeping horizontal line extending to the right.

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

All components are directly available through the Test Sponsor, Dell Inc.