Dell Inc.

TPC Express Benchmark™ Big Bench (TPCx-BB)

Full Disclosure Report

for

Dell PowerEdge 14G R640/R740xd

(with 1x PowerEdge R640; 18x PowerEdge R740xd)

using

Horton Works HDP 3.0.1

and

Red Hat Enterprise Linux Server 7.6

First Edition

October 11, 2019

***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 as a result of these and other factors. Therefore, the TPC Express Benchmark™ BB 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 Benchmark™, TPCx-BB and BBQpm, 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.

DELL EMC	**Dell PowerEdge 14G R640/R740xd**	TPCx-BB Rev. v1.3.1 TPC-Pricing Rev. v2.4.0 Report Date: October 11, 2019
Total System Cost	TPCx-BB Performance Metric	Price/Performance
1,166,306 USD	**3,089.93** BBQpm@10000	**377.46 USD** $/BBQpm@10000

Framework	Operating System	Other Software	Availability Date	Scale Factor	Streams
Horton Works HDP 3.0.1	Red Hat Enterprise Linux Server 7.6	None	October 11, 2019	10000	8

System Configuration

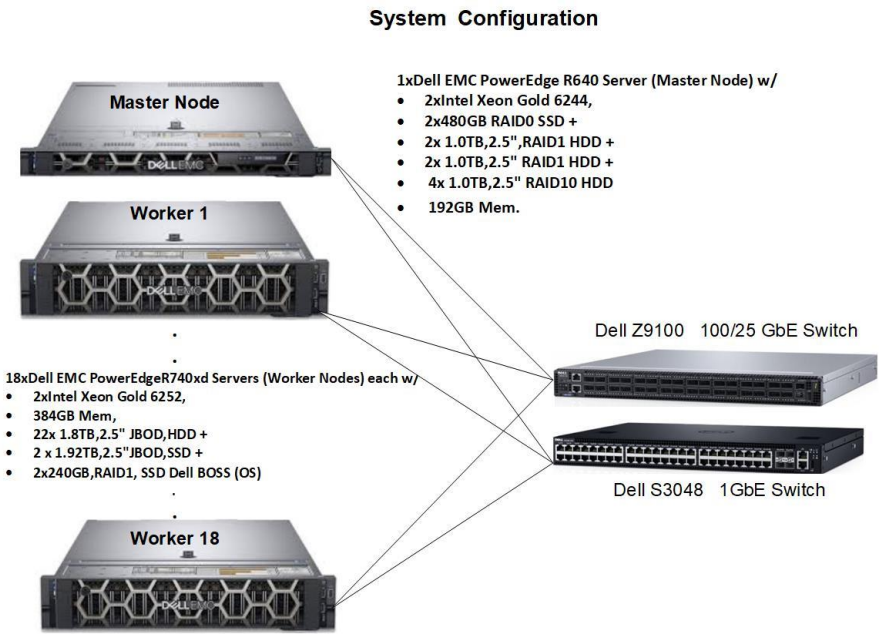

Physical Storage/Scale Factor: 79.95	Scale Factor/Physical Memory: 1.41

Servers: Total Processors/Cores/Threads	19x PowerEdge R640 / PowerEdge R740xd 38/880/1,760	
Server Configuration: Processors Memory Storage Controller Storage Device Network Controller	**1x PowerEdge R640 (Master):** 2x Intel(R) Xeon(R) Gold 6244 CPU @ 3.60GHz 192 GiB PERC H730P RAID Controller 2x 480 GB SSD SAS 8x 1 TB 7.2K RPM HDD Mellanox ConnectX-4 LX Dual Port 10/25GbE	**18x PowerEdge R740xd (Worker):** 2x Intel(R) Xeon(R) Gold 6252 CPU @ 2.10GHz 384 GiB HBA330 Controller Adapter 2x 240 GB RAID1, SSD Dell BOSS 2x 1.92 TB SSD SAS 22x 1.8 TB 10K RPM SAS HDD Mellanox ConnectX-4 LX Dual Port 10/25GbE
Connectivity:	1x Dell Z9100 100/25 GbE Switch, 1x Dell S3048 1 GbE Switch	

Description	Part Number	Key	Unit Price	Qty	Extended Price	3 yr. Maint. Price
HARDWARE COMPONENTS						
PowerEdge R740XD Server	210-AKZR	1	$59,137.51	18	$1,064,475.18	
PowerEdge R740/R740XD Motherboard	329-BEIK	1	$0.00	18		
No Trusted Platform Module	461-AADZ	1	$0.00	18		
Chassis with Up to 24 x 2.5" Hard Drives for 2CPU	321-BCPY	1	$0.00	18		
PowerEdge R740XD Shipping	340-BLBE	1	$0.00	18		
PowerEdge R740 Shipping Material	343-BBFU	1	$0.00	18		
Intel Xeon Gold 6252 2.1G, 24C/48T, 10.4GT/s, 35.75M Cache, Turbo, HT (150W) DDR4-2933	338-BSGU	1	$0.00	18		
Intel Xeon Gold 6252 2.1G, 24C/48T, 10.4GT/s, 35.75M Cache, Turbo HT (150W) DDR4-2933	338-BSGU	1	$0.00	18		
Additional Processor Selected	379-BDCO	1	$0.00	18		
Standard 2U Heatsink	412-AAIR	1	$0.00	18		
Standard 2U Heatsink	412-AAIR	1	$0.00	18		
2933MT/s RDIMMs	370-AEPP	1	$0.00	18		
Performance Optimized	370-AAIP	1	$0.00	18		
No RAID	780-BCDI	1	$0.00	18		
HBA330 Controller Adapter, Low Profile	405-AANK	1	$0.00	18		
BOSS controller card + with 2 M.2 Sticks 240G (RAID 1),FH	403-BBPT	1	$0.00	18		
Red Hat Enterprise Linux Non Factory Install, x64,Reqs Subscription Selection	421-4727	1	$0.00	18		
Red Hat Linux Registration Document, No Subscription	340-AVFG	1	$0.00	18		
iDRAC9,Enterprise	385-BBKT	1	$0.00	18		
iDRAC Group Manager, Enabled	379-BCQV	1	$0.00	18		
iDRAC,Legacy Password	379-BCSG	1	$0.00	18		
Riser Config 2, 3 x8, 1 x16 slots	330-BBHB	1	$0.00	18		
Mellanox ConnectX-4 LX Dual Port 10/25GbE SFP28, rNDC	406-BBLG	1	$0.00	18		
6 Performance Fans forR740/740XD	384-BBPZ	1	$0.00	18		
Dual, Hot-plug, Redundant Power Supply (1+1), 1100W	450-ADWM	1	$0.00	18		
PowerEdge 2U Standard Bezel	325-BCHU	1	$0.00	18		
PE R740XD Luggage Tag	389-BTTO	1	$0.00	18		
Quick Sync 2 (At-the-box mgmt)	350-BBJU	1	$0.00	18		
Power Saving Dell Active Power Controller	750-AABF	1	$0.00	18		
UEFI BIOS Boot Mode with GPT Partition	800-BBDM	1	$0.00	18		
ReadyRails Sliding Rails Without Cable Management Arm	770-BBBQ	1	$0.00	18		
No Systems Documentation, No OpenManage DVD Kit	631-AACK	1	$0.00	18		
US Order	332-1286	1	$0.00	18		
On-Site Installation Declined	900-9997	1	$0.00	18		
ProSupport 4Hr onSite and Mission Critical 24x7, 36 Month(s)	813-6068, 813-9295, 813-9296, 951-2015	1	$2,677.44	18		$48,193.92
Declined Remote Consulting Service	973-2426	1	$0.00	18		
32GB RDIMM, 2933MT/s, Dual Rank	370-AEQH	1	$0.00	216		
1.8TB 10K RPM SAS 12Gbps 512e 2.5in Hot-plug Hard Drive	400-ARXC	1	$0.00	396		
1.92TB SSD SAS Read Intensive 12Gbps 512e 2.5in	400-BBQQ	1	$0.00	36		
Red Hat Enterprise Linux,2SKT,1 Physical OR 2Guest,3Yr PREMIUM SUB,No Media	634-BJBO	1	$0.00	18		
NEMA 5-15P to C13 Wall Plug, 125 Volt, 15 AMP, 10 Feet (3m), Power Cord, North America	450-AALV	1	$0.00	36		
INFO QS, DATA ANALYTICS BUNDLE	379-BBWM	1	$0.00	18		

(continued on next page)

(continued from previous page)

PowerEdge R640 Server	210-AKWU	1	$32,108.64	1	$32,108.64	
PowerEdge R640 MLK Motherboard	329-BEIJ	1	$0.00	1		
No Trusted Platform Module	461-AADZ	1	$0.00	1		
2.5 Chassis with up to 10 Hard Drives and 3PCIe slots	321-BCQL	1	$0.00	1		
PowerEdge R640 Shipping	340-BKNE	1	$0.00	1		
PowerEdge R640 x4 and x10 Drive Shipping Material	340-BLUC	1	$0.00	1		
Intel Xeon Gold 6244 3.6G, 8C/16T, 10.4GT/s, 24.75M Cache, Turbo, HT (150W) DDR4-2933	338-BSHH	1	$0.00	1		
Intel Xeon Gold 6244 3.6G, 8C/16T, 10.4GT/s, 24.75M Cache, Turbo, HT (150W) DDR4-2933	338-BSHH	1	$0.00	1		
Additional Processor Selected	379-BDCO	1	$0.00	1		
DIMM Blanks for System with 2 Processors	370-ABWE	1	$0.00	1		
1U Pipe Low Profile Heatsink	412-AAIP	1	$0.00	1		
1U Pipe Low Profile Heatsink	412-AAIP	1	$0.00	1		
2933MT/s RDIMMs	370-AEPP	1	$0.00	1		
Performance Optimized	370-AAIP	1	$0.00	1		
Unconfigured RAID	780-BCDS	1	$0.00	1		
PERC H730P RAID Controller, 2GB NV Cache, Mini card	405-AANT	1	$0.00	1		
Red Hat Enterprise Linux Non Factory Install, x64,Reqs Subscription Selection	421-4727	1	$0.00	1		
Red Hat Linux Registration Document, No Subscription	340-AVFG	1	$0.00	1		
iDRAC9,Enterprise	385-BBKT	1	$0.00	1		
OpenManage Enterprise Advanced	528-BIYY	1	$0.00	1		
iDRAC Group Manager, Enabled	379-BCQV	1	$0.00	1		
iDRAC,Legacy Password	379-BCSG	1	$0.00	1		
Riser Config 4, 2x16 LP	330-BBGY	1	$0.00	1		
Mellanox ConnectX-4 LX Dual Port 10/25GbE SFP28, rNDC	406-BBLG	1	$0.00	1		
No Internal Optical Drive	429-AAIQ	1	$0.00	1		
8 Performance Fans for R640	384-BBQI	1	$0.00	1		
Dual, Hot-plug, Redundant Power Supply (1+1), 750W	450-ADWS	1	$0.00	1		
No Bezel	350-BBBW	1	$0.00	1		
Dell EMC Luggage Tag for x10	350-BBJT	1	$0.00	1		
Quick Sync 2 (At-the-box mgmt)	350-BBKC	1	$0.00	1		
Power Saving Dell Active Power Controller	750-AABF	1	$0.00	1		
UEFI BIOS Boot Mode with GPT Partition	800-BBDM	1	$0.00	1		
Energy Star	387-BBMK	1	$0.00	1		
ReadyRails Sliding Rails Without Cable Management Arm	770-BBBC	1	$0.00	1		
No Systems Documentation, No OpenManage DVD Kit	631-AACK	1	$0.00	1		
US Order	332-1286	1	$0.00	1		
Basic Next Business Day 36 Months	709-BBFM	1	$0.00	1		
On-Site Installation Declined	900-9997	1	$0.00	1		
ProSupport 4Hr onSite and Mission Critical 24x7, 36 Month(s)	813-9255,813-9259,813-9265,989-3439	1	$2,505.31	1		$2,505.31
16GB RDIMM, 2933MT/s, Dual Rank	370-AEQF	1	$0.00	12		
1TB 7.2K RPM NLSAS 12Gbps 512n 2.5in Hot-plug Hard Drive	400-ASHE	1	$0.00	8		
480GB SSD SAS Mixed use 12Gbps 512e 2.5in Hot-Plug PM5-V Drive, 3 DWPD, 2628 TBW	400-BCQG	1	$0.00	2		
Red Hat Enterprise Linux,2SKT,1 Physical OR 2Guest,3Yr PREMIUM SUB,No Media	634-BJBO	1	$0.00	1		
NEMA 5-15P to C13 Wall Plug, 125 Volt, 15 AMP, 10 Feet (3m), Power Cord, North America	450-AALV	1	$0.00	2		
INFO QS, DATA ANALYTICS BUNDLE	379-BBWM	1	$0.00	1		

(continued on next page)

(continued from previous page)

Dell Networking Z9100-ON Pod Switch		1	$52,477.00	1	$52,477.00	
Dell Networking Z9100-ON, 32x QSFP28 and 2x SFP+ fixed ports, PSU to IO airflow, 2x AC PSUs, OS9	210-AETC	1	$0.00	1		
OS9 installed on Z9100-ON, with entitlement to OS10 Enterprise	634-BPDD	1	$0.00	1		
Dell Networking Z9100-ON User Guide	631-AAPO	1	$0.00	1		
Software, Rights to use L3 on OS9, Z9100-ON	634-BEBM	1	$0.00	1		
US Order	332-1286	1	$0.00	1		
ProSupport Plus: Mission Critical 4-Hour 7x24 On-Site Service with Emergency Dispatch, 3 Year	803-9483,803-9474,803-9323,	1	$9,820.00	1		$9,820.00
ProSupport Plus: 7x24 HW/SW Tech Support and Assistance, 3 Year	803-9503	1	$0.00	1		
Thank you for choosing Dell ProSupport Plus. For tech support, visit //www.dell.com/contactdell	951-2015	1	$0.00	1		
Dell Limited Hardware Warranty Extended Year(s)	975-3461	1	$0.00	1		
Info 3rd Party Software Warranty provided by Vendor	997-6306	1	$0.00	1		
On-Site Installation Declined	900-9997	1	$0.00	1		
Dell Networking Cable,100GbE QSFP28 to 4 x SFP28 Passive Copper Breakout Cable, 2 Mete	470-ABOS	1	$0.00	5		
Dell Networking, Jumper Cord, 250V, 12A, 2 Meters, C13/C14, US	450-AASX	1	$0.00	1		
Dell Networking, Jumper Cord, 250V, 12A, 2 Meters, C13/C14, US	450-AASX	1	$0.00	1		
INFO QS, DATA ANALYTICS BUNDLE	379-BBWM	1	$0.00	1		

(continued on next page)

(continued from previous page)

Dell Networking S3048-ON iDRAC Switch		1	$11,750.37	1	$11,750.37	
Dell Networking S3048-ON, 48x 1GbE, 4x SFP+ 10GbE ports, Stacking, IO to PSU air, 1x AC PSU, DNOS 9	210-AEDM	1	$0.00	1		
Software, Rights to use L3 on OS9, S3048-ON	634-BDXE	1	$0.00	1		
OS9 installed on S3048-ON, with entitlement to OS10 Enterprise	528-BBSY	1	$0.00	1		
Dell Networking S3048-ON User Guide	634-BCXR	1	$0.00	1		
C2G 10ft Cat6 Ethernet Network Patch Cable	A5361988	1	$0.00	19		
US Order	332-1286	1	$0.00	1		
3 year ProSupport 4hr response, 24x7	802-7403, 802-7389,802-7394,802-7400	1	$1,495.33	1		$1,495.33
ProSupport: 7x24 HW / SW Tech Support and Assistance, 3 Years	802-7404	1	$0.00	1		
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		
Info 3rd Party Software Warranty provided by Vendor	997-6306	1	$0.00	1		
ProDeploy Dell Networking S Series 3XXX Switch - Deployment Verification	805-2399	1	$0.00	1		
ProDeploy Dell Networking S Series 3XXX Switch - Deployment	821-5792	1	$0.00	1		
Declined Remote Consulting Service	973-2426	1	$0.00	1		
Dell Networking, Jumper Cord, 250V, 12A, 2 Meters, C13/C14, US	450-AASX	1	$0.00	1		
INFO QS, DATA ANALYTICS BUNDLE	379-BBWM	1	$0.00	1		
Desktop MK120 Keyboard and Mouse	A6999510	1	$19.99	1	$19.99	
Dell 24 Monitor	210-AIWG	1	$169.99	1	$169.99	
Dell Netshelter SX 42U Rack - 600mm Wide x 1070mm Deep	A7545497	1	$1,299.99	1	$1,299.99	
HARDWARE COMPONENTS				**Subtotal**	$1,162,301.16	$62,014.56
SOFTWARE COMPONENTS						
Horton Works HDP Enterprise Plus, 24x7 1yr	HDP-ENT-1Y	1	10,000	57	$570,000.00	
SOFTWARE COMPONENTS				**Subtotal**	$570,000.00	$0.00
Total					**$1,732,301.16**	**$62,014.56**
Large Purchase Discount (35%)*					-$606,305.41	-$21,705.10

Pricing:1 = Dell

[1] All discounts are based on US list prices and for similar quantities and configurations. The discounts are based on the overall specific components pricing from respective vendors in this single quotation. Discounts for similarly sized configurations will be similar to those quoted here, but may vary based on the components in the configuration.

Audited by Doug Johnson of InfoSizing

Three-Year Cost of Ownership	**$1,166,306**
BBQpm@10000	**3,089.93**
$/BBQpm@10000	**$ 377.46**

Prices used in TPC benchmarks reflect the actual prices a customer would pay for a one-time purchase of the stated components. Individually negotiated discounts are not permitted. Special prices based on assumptions about past or future purchases are not permitted. All discounts reflect standard pricing policies for the listed components. For complete details, see the pricing sections of the TPC benchmark specifications. If you find that the stated prices are not available according to these terms, please inform at pricing@tpc.org. Thank you.

Numerical Quantities

Scale Factor	10000
Streams	8
SUT Validation Test	PASS

Performance Run (Run 2)

Overall Run Start Time	2019-09-24 05:58:54.356
Overall Run End Time	2019-09-24 22:33:57.344
Overall Run Elapsed Time	59,702.988
Load Test Start Time	2019-09-24 05:58:54.357
Load Test End Time	2019-09-24 06:43:26.019
Load Test Elapsed Time	2,671.662
Power Test Start Time	2019-09-24 06:43:26.020
Power Test End Time	2019-09-24 11:25:42.931
Power Test Elapsed Time	16,936.911
Throughput Test Start Time	2019-09-24 11:25:42.932
Throughput Test End Time	2019-09-24 22:33:57.343
Throughput Test Elapsed Time	40,094.411
Performance Metric (BBQpm@ 10000)	3,089.93

Repeatability Run (Run 1)

Overall Run Start Time	2019-09-23 13:02:17.656
Overall Run End Time	2019-09-24 05:25:10.734
Overall Run Elapsed Time	58,973.078
Load Test Start Time	2019-09-23 13:02:17.657
Load Test End Time	2019-09-23 13:46:54.671
Load Test Elapsed Time	2,677.014
Power Test Start Time	2019-09-23 13:46:54.673
Power Test End Time	2019-09-23 18:29:41.836
Power Test Elapsed Time	16,967.163
Throughput Test Start Time	2019-09-23 18:29:41.837
Throughput Test End Time	2019-09-24 05:25:10.734
Throughput Test Elapsed Time	39,328.897
Performance Metric (BBQpm@ 10000)	3,124.72

Performance Run Report (Run 2)

****
TPCx-BB
Result
v1.3.1
****
INFO: T_LOAD = 2671.662
INFO: T_LD = 0.1 * T_LOAD: 267.1662
INFO: T_PT = 6164.16941371755
INFO: T_T_PUT = 40094.411
INFO: T_TT = 5011.801375
INFO: === Checking validity of the final result ===
INFO: OK: All required BigBench phases were performed.
INFO: OK: All 30 queries were running in the power test.
INFO: OK: All 30 queries were running in the first throughput test.
INFO: OK: Pretend mode was inactive. All commands were executed.
INFO: === Final result ===
INFO: VALID BBQpm@10000 = 3089.93423799941

Repeatability Run Report (Run 1)

****
TPCx-BB
Result
v1.3.1
****
INFO: T_LOAD = 2677.014
INFO: T_LD = 0.1 * T_LOAD: 267.7014
INFO: T_PT = 6137.16833147717
INFO: T_T_PUT = 39328.897
INFO: T_TT = 4916.112125
INFO: === Checking validity of the final result ===
INFO: OK: All required BigBench phases were performed.
INFO: OK: All 30 queries were running in the power test.
INFO: OK: All 30 queries were running in the first throughput test.
INFO: OK: Pretend mode was inactive. All commands were executed.
INFO: === Final result ===
INFO: VALID BBQpm@10000 = 3124.72030342658

Summary details of the run reports are shown above. For the complete run reports, see the Support Files Archive.

Table of Contents

Abstract

This document contains the methodology and results of the TPC Express Benchmark™ Big Bench (TPCx-BB) test conducted in conformance with the requirements of the TPCx-BB Standard Specification, Revision v1.3.1.

The test was conducted at a Scale Factor of 10000 with 19 nodes (1x PowerEdge R640, 18x PowerEdge R740xd) running Horton Works HDP 3.0.1 on Red Hat Enterprise Linux Server 7.6.

Measured Configuration

Company Name	Cluster Node	Virtualization	Operating System
Dell Inc.	1x PowerEdge R640 18x PowerEdge R740xd	n/a	Red Hat Enterprise Linux Server 7.6

TPC Express Benchmark© Big Bench Metrics

Total System Cost	BBQpm@10000	Price/Performance	Availability Date
1,166,306 USD	3,089.93	377.46 USD	October 11, 2019

Preface

TPC Express Benchmark™ Big Bench Overview

Big data analytics is a growing field of research and business. The significant decrease in the overall cost of hardware, the emergence of Open Source based analytics frameworks, along with the greater depth of data mining capabilities allows new types of data sources to be correlated with traditional data sources. For example, online retailers used to record only successful transactions on their website, whereas modern systems are capable of recording every interaction. The former allowed for simple shopping basket analysis techniques, while the current level of detail in monitoring makes detailed user modeling possible. The growing demands on data management systems and the new forms of analysis have led to the development of a new type of ***Big Data Analytics Systems (BDAS)****.*

Similar to the advent of ***Database Management Systems****, there is a vastly growing ecosystem of diverse approaches to enabling Big Data Analytics Systems. This leads to a dilemma for customers of* ***BDAS****, as there are no realistic and proven measures to compare different* ***BDAS*** *solutions. To address this, TPC has developed TPCx-BB (BigBench), which is an express benchmark for comparing* ***BDAS*** *solutions. The TPCx-BB Benchmark was developed to cover essential functional and business aspects of big data use cases. The benchmark allows for an objective measurement of* ***BDAS*** *System under Test, and provides the industry with verifiable performance, price/performance, and availability metrics.*

The TPCx-BB kit is available from the TPC website (see www.tpc.org for more information). Users must sign-up and agree to the TPCx-BB End User Licensing Agreement (EULA) to download the kit. All related work (such as collaterals, papers, derivatives) must acknowledge the TPC and include the TPCx-BB copyright. The TPCx-BB kit includes: TPCx-BB Specification document (this document), TPCx-BB Users Guide documentation, shell scripts to set up the benchmark environment, Java code to execute the benchmark workload, Data Generator, ***Query*** *files, and Benchmark Driver.*

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-BB models and represents a Big Data Analytics System such as Hadoop ecosystem or 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 components used by the SUT.*
- *Configuration parameters and options for Operating System and file system components used by the SUT.*
- *Configuration parameters and options for any other software components (e.g compiler optimization options) used by the SUT.*

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 Archive contains the parameters and options used to configure the components involved in this benchmark.

1.3 Configuration Diagrams

7.4.4 Diagrams of both measured and priced configurations must be provided, accompanied by a description of the differences. This includes, but is not limited to:

- *Total number of nodes used;*
- *Total number and type of processors used/total number of cores used/total number of threads used (including sizes of L2 and L3 caches);*
- *Size of allocated memory, and any specific mapping/partitioning of memory unique to the test;*
- *Number and type of disk units (and controllers, if applicable;*
- *Number of channels or bus connections to disk units, including their protocol type;*
- *Number of LAN (e.g., Ethernet) connections and speed for switches and other hardware components physically used in the test or are incorporated into the pricing structure;*
- *Type and the run-time execution location of software components.*

Measured Configuration

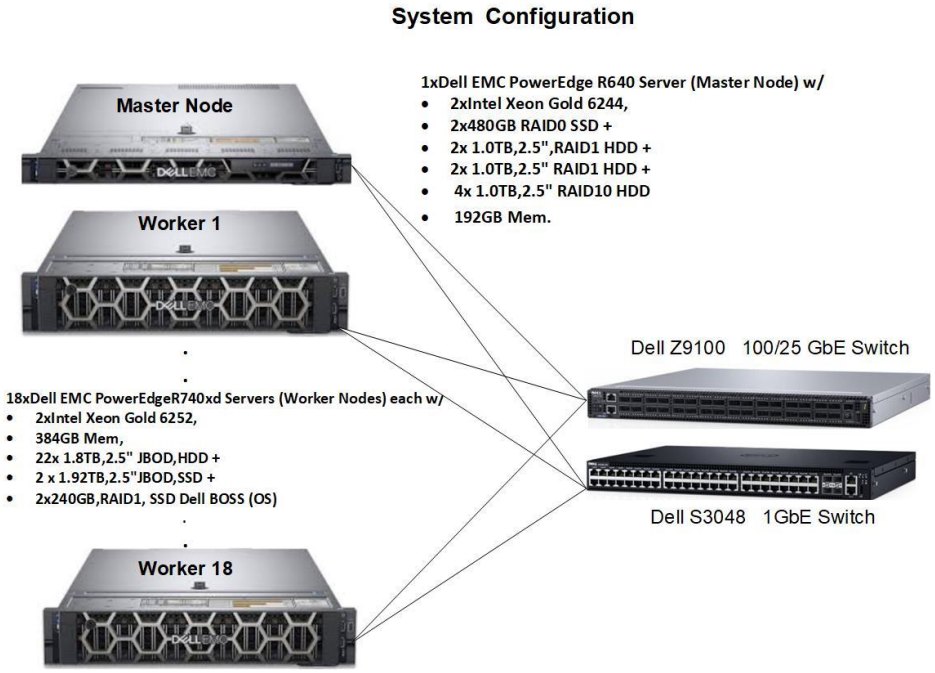

The measured configuration consisted of:

Total Nodes:	19
Total Processors/Cores/Threads:	38/880/1,760
Total Memory:	7,104
Total Number of Storage Devices:	478
Total Storage Capacity:	799,520
Network:	1x Dell Z9100 100/25 GbE Switch, 1x Dell S3048 1 GbE Switch

	1x PowerEdge R640 (Master):	**18x PowerEdge R740xd (Worker):**
Processors/Cores/Threads:	2/16/32	2/16/96
Processor Model:	2x Intel(R) Xeon(R) Gold 6244 CPU @ 3.60GHz	2x Intel(R) Xeon(R) Gold 6252 CPU @ 2.10GHz
Memory:	192 GiB	384 GiB
Storage Controller:	PERC H730P RAID Controller	HBA330 Controller Adapter
Storage Devices:	2x 480 GB SSD SAS 8x 1 TB 7.2K RPM HDD	2x 240 GB RAID1, SSD Dell BOSS 2x 1.92 TB SSD SAS 22x 1.8 TB 10K RPM SAS HDD
Network Controller:	Mellanox ConnectX-4 LX Dual Port 10/25GbE	Mellanox ConnectX-4 LX Dual Port 10/25GbE

The distribution of software components over server nodes is detailed in section 2.1.

Priced Configuration

There are no differences between the priced and measured configurations.

Clause 2: Software Components and Dataset Distribution

2.1 Roles and Dataset Distribution

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

Table 1.4 describes the distribution of the dataset across all media in the system.

Table 1.4: Software Components and Dataset Distribution

Server	Role(s)	Count	Virtual	Host Names	HW/SW Configuration	Storage Setup
Worker	HDFS DataNode HST Agent/SmartSense Ambari Metrics Monitor HDFS NFSGateway NodeManager HDFS Client Hive Client MapReduce2 Client Spark2 Client Tez Client YARN Client Zookeeper Client	18	N	r2xd[1-18] alias dn[1-18]	• Dell PowerEdge R740xd • Processor: 2x Intel Xeon Gold 6252 • Memory: 384 GB • Storage: 22x 1.8TB (Data), 240GB BOSS (OS), 2x1.92TB SSD • Network: Mellanox 25GbE 2p Connectx-4lx • OS: RHEL 7.6 • Horton Works HDP 3.0.1	OS: 240GB BOSS,SAS, HDD Intermediate/Shuffle /Temp Data Distributed FS: 22x 1.8 TB, JBOD, SAS
Master	SmartSense Activity Explorer YARN Timeline Service MapReduce2 History Server Hive Metastore HiveServer2 SmartSense HST Server Ambari Metrics Collector Grafana NameNode ResourceManager SNameNode Spark2 History Server YARN Timeline Service Reader YARN Registry DNS Zookeeper Server SmartSense HST Agent Ambari Metrics Monitor HDFS NFSGateway HDFS Client Hive Client MapReduce2 Client Spark2 Client Tez Client YARN Client Zookeeper Client	1	N	namenode2 alias nn2	• Dell PowerEdge R640 • Processor: 2x Intel Xeon Gold 6244 • Memory : 192 GB • Storage: 8x 1 TB (OS) • Network: Mellanox 25GbE 2p Connectx-4lx • OS: RHEL 7.6 • Horton Works HDP 3.0.1	OS: 2x1 TB RAID1, SAS, HDD Intermediate/Temp Data: 2x480GB, JBOD,SAS,SSD Metadata: 2x1 TB RAID1, SAS, HDD PostgreSQL DB: 4x 1.0TB RAID10,SAS, HDD

2.2 Distributed File System Implementation

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

Horton Works HDP 3.0.1 (fully HDFS compatible at the API level).

2.3 Engine Implementation

The Engine implementation and corresponding version must be disclosed.

Component	Version
HDFS	3.1.1
YARN	3.1.1
MapReduce2	3.1.1
Spark2	2.3.1
Hive	3.1.0

2.4 Frameworks

Frameworks and Engine used in the benchmark should be disclosed.

Framework	Version
HDP	3.0.1
HDFS	3.1.1
YARN	3.1.1
MapReduce2	3.1.1
Spark2	2.3.1
Hive	3.1.0

2.5 Applied Patches

Any additional vendor supported patches applied to the SUT should be disclosed.

No additional patches were applied.

Clause 3: Workload Related Items

3.1 Hardware & Software Tunable

Script or text used to set for all hardware and software tunable parameters must be reported.

The Supporting Files Archive contains all configuration scripts.

3.2 Kit Version

Version number of the TPCx-BB kit must be included in the Report.

TPCx-BB Kit Version
v1.3.1

3.3 Run Report

The run report generated by TPCx-BB benchmark kit must be included in the Report.

The Supporting File Archive contains the full run report. Following are summary extracts from both runs.

- **Run1 Report Summary (Repeatability Run)**

```
**********
TPCx-BB
Result
v1.3.1
**********
INFO: T_LOAD = 2677.014
INFO: T_LD = 0.1 * T_LOAD: 267.7014
INFO: T_PT = 6137.16833147717
INFO: T_T_PUT = 39328.897
INFO: T_TT = 4916.112125
INFO: === Checking validity of the final result ===
INFO: OK: All required BigBench phases were performed.
INFO: OK: All 30 queries were running in the power test.
INFO: OK: All 30 queries were running in the first throughput test.
INFO: OK: Pretend mode was inactive. All commands were executed.
INFO: === Final result ===
INFO: VALID BBQpm@10000 = 3124.72030342658
```

- **Run2 Report Summary (Performance Run)**

```
**********
TPCx-BB
Result
v1.3.1
**********
INFO: T_LOAD = 2671.662
INFO: T_LD = 0.1 * T_LOAD: 267.1662
INFO: T_PT = 6164.16941371755
INFO: T_T_PUT = 40094.411
INFO: T_TT = 5011.801375
INFO: === Checking validity of the final result ===
INFO: OK: All required BigBench phases were performed.
INFO: OK: All 30 queries were running in the power test.
INFO: OK: All 30 queries were running in the first throughput test.
INFO: OK: Pretend mode was inactive. All commands were executed.
INFO: === Final result ===
INFO: VALID BBQpm@10000 = 3089.93423799941
```

3.4 Query Elapsed Times

Elapsed times of all power and throughput Queries needs to be reported from the Performance Run, grouped respectively as Structured, semi-structured and unstructured buckets.

Type	Query	Power	Stream 1	Stream 2	Stream 3	Stream 4	Stream 5	Stream 6	Stream 7	Stream 8
Structured	1	50.060	65.056	86.608	105.688	77.300	98.663	236.833	74.794	61.564
	6	213.810	608.846	485.837	346.838	506.742	401.801	435.518	513.737	534.399
	7	50.855	57.492	83.135	83.342	73.347	107.943	181.210	90.030	83.399
	9	57.062	91.618	93.703	142.967	92.430	63.449	114.242	87.545	89.931
	11	41.589	82.344	59.759	78.000	58.255	86.640	74.107	141.343	53.318
	13	82.504	155.092	123.038	144.845	123.564	166.121	152.987	147.323	174.970
	14	41.450	91.205	59.706	85.175	47.956	72.348	78.738	45.604	63.298
	15	42.436	71.002	72.657	52.707	64.638	41.620	171.014	64.203	59.770
	16	339.154	361.931	644.863	638.409	501.482	613.663	740.345	677.573	808.888
	17	46.146	49.519	141.492	61.742	61.947	108.939	188.315	59.451	68.417
	20	131.125	188.464	263.908	247.496	277.825	248.717	343.489	219.144	180.052
	21	304.426	611.840	605.111	655.010	519.681	392.079	583.846	419.531	597.682
	22	142.729	196.211	131.707	178.126	206.547	215.734	211.942	191.893	233.197
	23	74.845	112.603	130.702	114.360	104.856	102.635	193.654	133.841	87.016
	24	52.750	63.640	116.211	78.492	75.633	87.076	313.884	112.069	80.141
	25	173.170	321.309	523.976	328.708	291.676	379.621	459.919	304.808	315.869
	26	115.325	131.492	163.533	172.281	174.400	193.480	192.035	176.901	181.646
	29	147.831	225.979	290.420	207.278	284.873	318.518	312.207	305.627	326.302
Semi-structured	2	3,548.893	11,283.031	10,670.366	10,443.934	11,195.059	11,309.618	10,324.562	9,852.016	10,960.396
	3	1,220.760	3,334.982	3,068.169	3,025.183	3,348.305	2,596.231	3,073.691	3,156.937	3,256.715
	4	1,890.452	5,897.913	5,240.957	5,340.582	5,308.315	4,837.549	5,820.342	5,854.963	4,551.566
	5	990.357	2,215.724	2,325.429	2,384.962	2,268.817	2,398.963	1,333.163	1,980.227	2,141.476
	8	384.688	668.751	842.168	775.770	746.900	916.720	950.375	1,263.857	809.690
	12	241.318	359.729	445.174	380.107	435.918	382.435	357.448	346.233	357.526
	30	2,592.345	6,551.249	6,303.641	6,614.376	6,081.812	6,812.430	5,984.097	6,382.101	6,378.116
Unstructured	10	335.877	481.433	519.548	646.715	510.906	609.126	452.111	547.940	520.385
	18	1,898.880	3,409.041	3,818.338	3,182.766	3,412.875	3,579.186	3,538.852	3,782.173	3,830.723
	19	1,555.959	2,121.831	2,446.076	2,405.583	2,017.212	2,523.339	2,905.867	2,336.152	2,756.516
	27	56.087	75.536	80.475	75.294	80.733	113.751	121.206	81.848	83.361
	28	114.004	150.119	189.689	236.746	192.496	220.107	248.401	254.019	204.818

3.5 Validation Test Output

Output report from successful SUT Validation test must be included in the Report.

Query Number	Query Execution	Output Validation
1	PASS	PASS
2	PASS	PASS
3	PASS	PASS
4	PASS	PASS
5	PASS	PASS
6	PASS	PASS
7	PASS	PASS
8	PASS	PASS
9	PASS	PASS
10	PASS	PASS
11	PASS	PASS
12	PASS	PASS
13	PASS	PASS
14	PASS	PASS
15	PASS	PASS
16	PASS	PASS
17	PASS	PASS
18	PASS	PASS
19	PASS	PASS
20	PASS	PASS
21	PASS	PASS
22	PASS	PASS
23	PASS	PASS
24	PASS	PASS
25	PASS	PASS
26	PASS	PASS
27	PASS	PASS
28	PASS	PASS
29	PASS	PASS
30	PASS	PASS

3.6 Global Framework Parameters

Global Framework parameter settings files must be included in the Report.

The Supporting File Archive contains the global framework parameter settings files.

3.7 Kit Modifications

Test Sponsor kit modifications files must be included in the Report.

The following files were modified by the Test Sponsor to facilitate system, platform and Framework differences.

- bigBench-configs/conf/userSettings.conf
- bigBench-configs/spark_sql/engineSettings.conf

Clause 4: SUT Related Items

4.1 Specialized Hardware/Software

Specialized Hardware/Software used in the SUT must be included.

No specialized hardware or software was used.

4.2 Framework Configuration Files

All Framework configuration files from SUT, for the performance run.

All Framework configuration files are included in the Supporting Files Archive.

4.3 SUT Environment Information

SUT environment info in form of envinfo.log from a representative worker node form every role in the server.

All envinfo.log files are included in the Supporting Files Archive.

4.4 Data Storage to Scale Factor Ratio

The data storage ratio must be disclosed.

Nodes	Disks	Size (GB)	Total (GB)
1	8	1,000	8,000
1	2	480	960
18	2	240	8,640
18	22	1,800	712,800
18	2	1,920	69,120

Total Storage (GB)	799,520
Scale Factor	10000
Data Storage Ratio	79.95

4.5 Scale Factor to Memory Ratio

The Scale Factor to memory ratio must be disclosed.

Nodes	Memory (GB)	Total (GB)
1	192	192
18	384	6,912

Scale Factor	10000
Total Memory (GB)	7,104
SF / Memory Ratio	1.41

Clause 5: Metrics and Scale Factors

5.1 Performance Run Metric

The Reported Performance Metric (BBQpm@SF for the Performance Run) must be disclosed in the Report.

Performance Run
BBQpm@10000 3,089.93

5.2 Repeatability Run Metric

The Performance Metric (BBQpm@SF) for the Repeatability Run must be disclosed in the Report.

Repeatability Run
BBQpm@10000 3,124.72

5.3 Price-Performance Metric

The Reported Performance Metric (BBQpm@SF for the Performance Run) must be disclosed in the Report.

Price / Performance
$BBQpm@10000 377.46

5.4 Scale Factor

The Scale Factor used for the Result must be disclosed in the Report.

Scale Factor
10000

5.5 Stream Count

The number of streams in the throughput run used for the Result must be disclosed in the Report.

Streams
8

5.6 Elapsed Run Times

The total elapsed time for the execution of the Performance Run and Repeatability Run must be disclosed in the Report.

Run	Elapsed Time	Seconds
Run 1	00 16:22:53.078	58,973.078
Run 2	00 16:35:02.988	59,702.988

5.7 Elapsed Test Times

The total time for each of the three tests must be disclosed for the Performance Run and the Repeatability Run.

Test	Performance Run	Repeatability Run
Load Test	2,671.662	2,677.014
Power Test	16,936.911	16,967.163
Throughput Test	40,094.411	39,328.897

Auditors' Information and Attestation Letter

The auditor's agency name, address, phone number, and Attestation letter must be included in the full disclosure report. A statement should be included specifying who to contact in order to obtain further information regarding the audit process.

This benchmark was audited by Doug Johnson, InfoSizing.

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

This benchmark's Full Disclosure Report (FDR) can be downloaded from www.tpc.org.

A copy of the auditor's attestation letter is included in the next two pages.

InfoSizing
The Right Metric For Sizing IT

TPC
Certified Auditor

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

October 10, 2019

I verified the TPC Express Benchmark™ BB v1.3.1 performance of the following configuration:

Platform:	Dell PowerEdge 14G R640/R740xd (with 1x R640, 18x R740xd)
Operating System:	Red Hat Enterprise Linux Server 7.6
Apache Hadoop Compatible Software:	Horton Works HDP 3.0.1

The results were:

Performance Metric	**3,089.93 BBQpm@10000GB**
Run Elapsed Time	00 16:35:02.988 (59,702.988 Seconds)

Cluster	**1x R640 (Master Node), 18x R740xd (Worker Nodes)**		
CPUs	2x Intel® Xeon® Gold 6244 (3.60 GHz, 8-core, 24.75 MB L3) (Master Node) 2x Intel® Xeon® Gold 6252 (2.10 GHz, 24-core, 35.75 MB L3) (Worker Nodes)		
Memory	192GB (Master Node); 384GB (Worker Nodes)		
Storage	***Qty***	***Size***	***Type***
	2	480 GB	SAS SSD (Master Node)
	8	1 TB	7.2K RPM HDD (Master Node)
	2	240 GB	SSD (Worker Nodes)
	2	1.92 TB	SAS SSD (Worker Nodes)
	22	1.8 TB	10K RPM SAS HDD (Worker 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 v1.3.1
- No modifications were made to any of the Java code
- Any and all modifications to shell scripts were reviewed for compliance
- The tested Scale Factor (10000GB) was confirmed to be valid for publication

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

- All validation queries executed successfully and produced compliant results
- No errors were reported during the run
- The elapsed times for all phases and runs were correctly measured and reported
- The Storage and Memory Ratios were correctly calculated and reported
- The system pricing was verified for major components and maintenance
- The major pages from the FDR were verified for accuracy

Additional Audit Notes:

None.

Respectfully Yours,

Doug Johnson, TPC Auditor

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

Third Party Price Quotes

None.

Supporting File Index

The following index outlines the information included in the supporting files archive.

Description	Archive File Pathname
Clause 1 - General Items	
The Supporting Files Archive contains the parameters and options used to configure the components involved in this benchmark	Supporting-Files-10TB-14G-xBB-10-2019\
Validation Run Files	Supporting-Files-10TB-14G-xBB-10-2019\Validation-run logs
Performance Run Files	Supporting-Files-10TB-14G-xBB-10-2019\Performance-run logs
Repeatability Run Files	Supporting-Files-10TB-14G-xBB-10-2019\Repeatability-run logs
Clause 3 - Workload Related Items	
Benchmark Generic Parameters	Supporting-Files-10TB-14G-xBB-10-2019\ Performance-run logs\bigBench-configs\conf\userSettings.conf
Query Parameters used in the benchmark execution Settings	Supporting-Files-10TB-14G-xBB-10-2019\ Performance-run logs\bigBench-configs\spark_sql\conf\queryParameters.sql
Benchmark Global Framework Parameters Settings	Supporting-Files-10TB-14G-xBB-10-2019\ Performance-run logs\bigBench-configs\spark_sql\conf\engineSettings.sql
Benchmark Global Framework Parameters Settings	Supporting-Files-10TB-14G-xBB-10-2019\ Performance-run logs\bigBench-configs\hive\spark_sql\engineSettings.conf
Load Test script	Supporting-Files-10TB-14G-xBB-10-2019\ Performance-run logs\bigBench-configs\hive\spark_sql\engineSettings.conf
Queries specific optimization parameters settings	Supporting-Files-10TB-14G-xBB-10-2019\ Performance-run logs\bigBench-configs \spark_sql\population\ sparkSqlCreateLoad.sql
Queries specific optimization parameters settings	Supporting-Files-10TB-14G-xBB-10-2019\ Performance-run logs\bigBench-configs\spark_sql\queries\q[01-30]\engineLocalSettings.conf
Clause 4 - SUT Related Items	
Data Redundancy report	Supporting-Files-10TB-14G-xBB-10-2019\Performance-run logs\data-redundancy-report.txt
Benchmark execution script	Supporting-Files-10TB-14G-xBB-10-2019\TPCxBB_Benchmarkrun.sh
Benchmark run script	Supporting-Files-10TB-14G-xBB-10-2019\ TPCxBB_FullBenchmark_sequence_run
Hardware and Software Report from a representative node	Supporting-Files-10TB-14G-xBB-10-2019\ Performance-run logs\run-logs\envInfo-r2xd1.ignition.dell.com\envInfo.log
All Framework configuration files are included in the Supporting Files Archive	Supporting-Files-10TB-14G-xBB-10-2019\Performance-run-logs\run-logs\envInfo-r2xd1.ignition.dell.com\hadoop
	Supporting-Files-10TB-14G-xBB-10-2019\Performance-run logs\run-logs\envInfo-r2xd1.ignition.dell.com\hive
Clause 5 - Metric and Scale Factor Related Items	
Benchmark Performance Report	Supporting-Files-10TB-14G-xBB-10-2019\Performance-run logs\run-logs\BigBenchResult.log
Validation Test Report	Supporting-Files-10TB-14G-xBB-10-2019\Validation-run logs\run-logs\BigBenchResult.log