Hewlett Packard Enterprise

Hewlett Packard Enterprise Company

TPC Express Benchmark™ Big Bench (TPCx-BB)

Full Disclosure Report

for

Hewlett Packard Enterprise ProLiant DL for Big Data

(with 9x HPE ProLiant DL380 Gen9, 3x HPE ProLiant DL360 Gen9)

using

Cloudera for Apache Hadoop (CDH) 5.6

and

Red Hat Enterprise Linux Server 6.7

First Edition

March 31, 2016

***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 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.

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 Benchmark™, TPCx-BB and BBQpm, 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.

Hewlett Packard Enterprise	**Hewlett Packard Enterprise ProLiant DL for Big Data**	TPCx-BB Rev. 1.0.1 TPC-Pricing Rev. 1.7.0 Report Date: March 31, 2016
Total System Cost	TPCx-BB Performance Metric	Price/Performance
371,977 USD	**337.26** BBQpm@3000	**1,102.94 USD** $/BBQpm@3000

Framework	Operating System	Other Software	Availability Date	Scale Factor	Streams
Cloudera for Apache Hadoop (CDH) 5.6	Red Hat Enterprise Linux Server 6.7	None	March 31, 2016	3000	2

System Configuration

Ethernet Switch:
HPE 1620-24G Switch (ILO connection)
HP5900AF-48XGT-4QSFP+ (main connection)

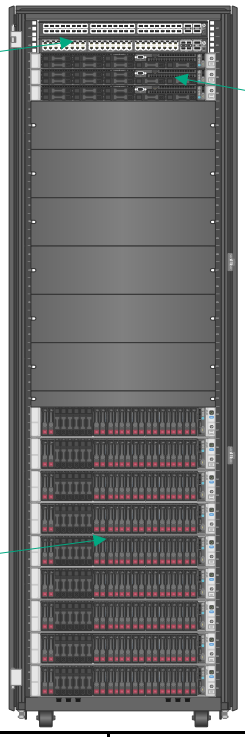

3 Management Nodes
Each:
HPE ProLiant DL360 Gen9
2x Intel E5-2640 v4 2.40GHz
1x HPE 800GB SSD
256GB Memory

9 Worker Nodes
Each:
HPE ProLiant DL380 Gen9
2x Intel E5-2697A v4 2.60GHz
1x HPE 480GB SSD
16x HPE 1TB HDD
1x HPE 1.2TB SDD
192GB Memory

Software:
Red Hat Enterprise Linux 6.7
Cloudera Enterprise 5.6

Physical Storage/Scale Factor: 53.84	Scale Factor/Physical Memory: 1.20

Servers: Total Processors/Cores/Threads	9x HPE ProLiant DL 380 Gen 9, 3x HPE ProLiant DL360 Gen9 24/348/696	
Server Configuration: Processors Memory Storage Controller Storage Device Network	**Per HPE ProLiant DL 380 Gen 9:** 2x Intel Xeon E5-2697A v4 @ 2.60GHz 192GB HPE Smart Array P840/4G 1x HPE 480GB SSD 16x HPE 1 TB HDD 1x HPE 1.2 TB SSD HPE Ethernet 10Gb 2P 560FLR-SFP+ Adptr	**Per HPE ProLiant DL360 Gen9:** 2x Intel Xeon E5-2640 v4 @ 2.40GHz 256GB Smart HBA H240ar 1x HPE 800GB SSD HPE Ethernet 10Gb 2P 560FLR-SFP+ Adptr
Connectivity:	HPE 1620-24G Switch, HPE 5900AF-48XG-4QSFP+ Switch	

Hewlett Packard Enterprise

Hewlett Packard Enterprise ProLiant DL for Big Data

TPCx-BB Rev. 1.0.1
TPC-Pricing Rev. 1.7.0

Report Date:
March 31, 2016

Description	Price Key	Part Number	Unit Price	Qty	Extended Price	3 Yr Maint Price
Server Hardware						
HPE DL360 Gen9 8SFF CTO Server	1	755258-B21	$1,797	3	$5,391	
HPE DL360 Gen9 E5-2640v4 Kit	1	842978-B21	$1,359	3	$4,077	
HPE DL360 Gen9 E5-2640v4 FIO Kit	1	842978-L21	$1,359	3	$4,077	
HPE 32GB 2Rx4 PC4-2400T-R Kit	1	805351-B21	$759	24	$18,216	
HPE H240ar 12Gb 2-ports Int FIO Smart Host Bus Adapter	1	749976-B21	$249	3	$747	
HPE 500W FS Plat Ht Plg Pwr Supply Kit	1	720478-B21	$309	6	$1,854	
HPE 800GB 6G SATA MU-2 SFF SC SSD	1	804671-B21	$3,038	3	$9,114	
HPE Ethernet 10Gb 2P 560FLR-SFP+ Adptr	1	665243-B21	$679	3	$2,037	
HPE 3Y FC 24x7 DL360 Gen9 SVC	1	U7AL9E	$1,404	3		$4,212
HPE iLO Adv incl 3yr TS U E-LTU	1	E6U64ABE	$469	3		$1,407
HP W1972a 18.5-In LED Monitor (1 + 2 spare)	1	B7M13A8#	$79	3	$238	
HP PS/2 Keyboard And Mouse Bundle (1 + 2 spare)	1	B1T13AA#	$27	3	$82	
			Subtotal		**$45,832**	**$5,619**
HPE DL380 Gen9 24SFF CTO Server	1	767032-B21	$2,107	9	$18,963	
HPE DL380 Gen9 High Perf Fan Kit	1	719079-B21	$239	9	$2,151	
HPE DL380 Gen9 2SFF Bay Kit	1	724864-B21	$149	9	$1,341	
HPE DL380 Gen9 Secondary Riser	1	719073-B21	$99	9	$891	
HPE XL2x0 Gen9 E5-2697Av4 Kit	1	841179-B21	$4,359	9	$39,231	
HPE XL2x0 Gen9 E5-2697Av4 FIO Kit	1	841179-L21	$4,359	9	$39,231	
HPE 32GB 2Rx4 PC4-2400T-R Kit	1	805351-B21	$759	54	$40,986	
HPE 1TB 6G SATA 7.2k 2.5in SC MDL HDD	1	655710-B21	$499	144	$71,856	
HPE 480GB 6G SATA RI-2 SFF SC SSD	1	804593-B21	$1,039	9	$9,351	
HPE 1.2TB 6G SATA WI-2 SFF SC SSD	1	804677-B21	$4,557	9	$41,013	
HPE Smart Array P840/4G Controller	1	726897-B21	$1,249	9	$11,241	
HPE 12Gb DL380 Gen9 SAS Expander Card	1	727250-B21	$699	9	$6,291	
HPE 800W FS Ti Ht Plg Pwr Supply Kit	1	720482-B21	$409	18	$7,362	
HPE Ethernet 10Gb 2P 560FLR-SFP+ Adptr	1	665243-B21	$679	9	$6,111	
HPE 3Y FC 24x7 DL380 Gen9 SVC	1	U7AE5E	$1,821	9		$16,389
HPE iLO Adv incl 3yr TS U E-LTU	1	E6U64ABE	$469	9		$4,221
			Subtotal		**$296,019**	**$20,610**
Network						
HPE 1620-24G Switch	1	JG913A	$299	1	$299	
HPE 5900AF-48XG-4QSFP+ Switch	1	JC772A	$20,990	1	$20,990	
HPE 1m Multi-mode OM3 LC/LC FC Cable	1	AJ834A	$70	12	$840	
HPE A58x0AF 650W AC Power Supply	1	JC680A	$749	4	$2,996	
HPE 58x0AF Bck(pwr)-Frt(ports) Fan Tray	1	JC682A	$179	4	$716	
			Subtotal		**$25,841**	**$0**
Rack						
HPE Universal Rack 11642 1075mm Shock Rack	1	H6J66A	$1,699	1	$1,699	
HPE 24A High Voltage Core Only Corded PDU	1	252663-D74	$259	2	$518	
			Subtotal		**$2,217**	**$0**

Hewlett Packard Enterprise

Hewlett Packard Enterprise ProLiant DL for Big Data

TPCx-BB Rev. 1.0.1
TPC-Pricing Rev. 1.7.0

Report Date:
March 31, 2016

Description	Price Key	Part Number	Unit Price	Qty	Extended Price	3 Yr Maint Price
Server Software						
Cloudera Ent Basic Ed 1yr 24x7	1	G7M27A	$2,304	36	$82,944	
RHEL Svr 2 Sckt/2 Gst 3yr 24x7 E-LTU	1	G3J30AAE	$3,889	12	$46,668	
			Subtotal		**$129,612**	**$0**
		Total Extended Price			**$499,521**	**$26,229**
		Total Discounts			**$152,171**	**$1,603**
Sales contact: HPE WW Headquarters, 3000 Hanover St., Palo Alto, CA 94304-1185 (650) 857-1501 or H-P direct: 800-203-6748		**Grand Total**			**$347,350**	**$24,626**

Pricing:1 = HPE

(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, Inc.

Three-Year Cost of Ownership	**$371,977**
BBQpm@3000	**337.26**
$/BBQpm@3000	**$1,102.94**

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	3000
Streams	2
SUT Validation Status	PASS

Performance Run

Overall Run Start Time	2016-03-22 16:10:23.675
Overall Run End Time	2016-03-23 08:06:00.543
Overall Run Elapsed Time	57,336.868
Load Test Start Time	2016-03-22 16:10:23.675
Load Test End Time	2016-03-22 16:46:01.997
Load Test Elapsed Time	2,138.322
Power Test Start Time	2016-03-22 16:46:01.999
Power Test End Time	2016-03-22 23:22:44.273
Power Test Elapsed Time	23,802.274
Throughput Test Start Time	2016-03-22 23:22:44.275
Throughput Test End Time	2016-03-23 08:06:00.542
Throughput Test Elapsed Time	31,396.267
Performance Metric (BBQpm@3000)	337.26

Repeatability Run

Overall Run Start Time	2016-03-23 08:27:54.443
Overall Run End Time	2016-03-24 00:21:31.933
Overall Run Elapsed Time	57,217.490
Load Test Start Time	2016-03-23 08:27:54.444
Load Test End Time	2016-03-23 09:03:43.324
Load Test Elapsed Time	2,148.880
Power Test Start Time	2016-03-23 09:03:43.326
Power Test End Time	2016-03-23 15:36:38.693
Power Test Elapsed Time	23,575.367
Throughput Test Start Time	2016-03-23 15:36:38.694
Throughput Test End Time	2016-03-24 00:21:31.933
Throughput Test Elapsed Time	31,493.239
Performance Metric (BBQpm@3000)	337.52

Run Report – Run 1

```
**********
TPCx-BB
Result
v1.0.1
**********
INFO: T_LOAD = 2138.322
INFO: T_LD = 0.1 * T_LOAD: 213.83220000000003
INFO: T_PT = 15897.004100627742
INFO: T_T_PUT = 31396.267
INFO: T_TT = 15698.1335
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@3000 = 337.26627326432424
```

Run Report – Run 2

```
**********
TPCx-BB
Result
v1.0.1
**********
INFO: T_LOAD = 2148.88
INFO: T_LD = 0.1 * T_LOAD: 214.88800000000003
INFO: T_PT = 15821.423876979647
INFO: T_T_PUT = 31493.239
INFO: T_TT = 15746.6195
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@3000 = 337.5239352171591
```

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 1.0.1.

The test was conducted at a Scale Factor of 3000 with 12 nodes (9x HPE ProLiant DL380 Gen9, 3x HPE ProLiant DL360 Gen9) running Cloudera for Apache Hadoop (CDH) 5.6 on Red Hat Enterprise Linux Server 6.7.

Measured Configuration

Company Name	Cluster Node	Virtualization	Operating System
Hewlett Packard Enterprise Company	9x HPE ProLiant DL380 Gen9, 3x HPE ProLiant DL360 Gen9	n/a	Red Hat Enterprise Linux Server 6.7

TPC Express Benchmark© Big Bench Metrics

Total System Cost	BBQpm@3000	Price/Performance	Availability Date
371,977 USD	337.26	1,102.94 USD	March 31, 2016

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 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 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

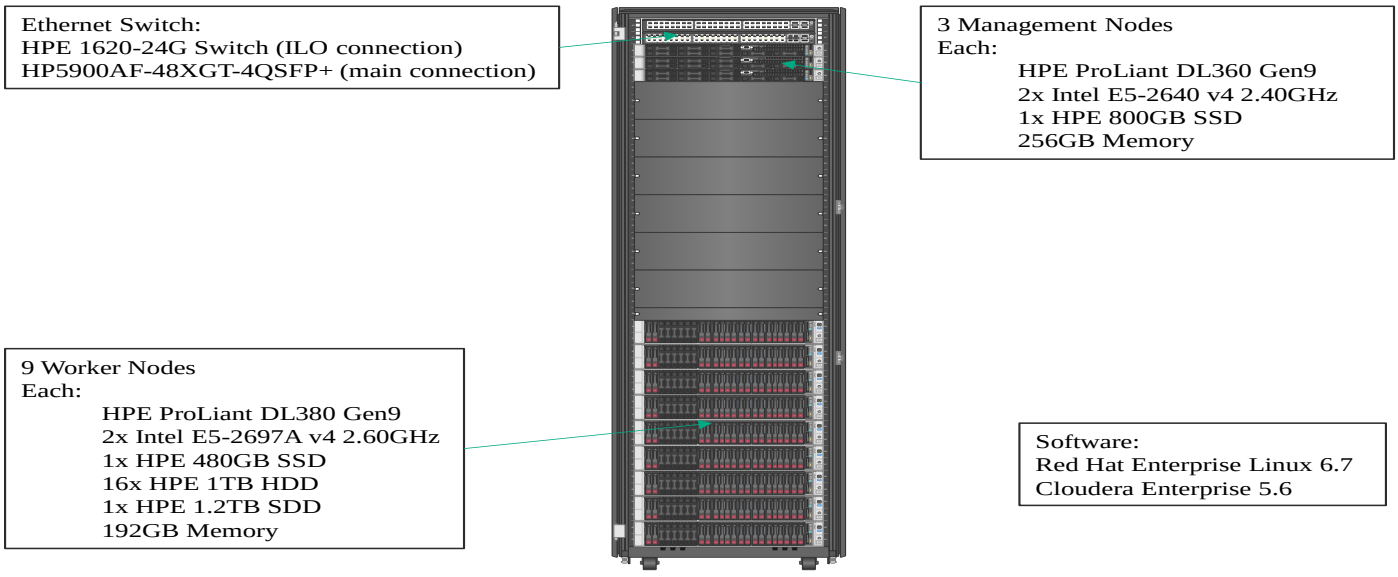

The measured configuration consisted of:

- Total Nodes: 12
- Total Processors/Cores/Threads: 24/348/696
- Total Memory: 2,496GB
- Total Number of Storage Drives/Devices: 165
- Total Storage Capacity: 161,520GB

Network connectivity detail:

- HP5900AF-XG-4QSFP+Switch, HPE 1620-24G Switch

Server nodes details:

9x HPE ProLiant DL380 Gen9, each with:

- Processors/Cores/Threads: 2/32/64
- Processor Model: 2x Intel Xeon E5-2697A v4 @ 2.60GHz
- Memory: 192GB
- Controller: 1 x HPE Smart Array P840/4G
- Drives:
  - 1 x HPE 480GB SSD
  - 16 x HPE 1TB SSD
  - 1 x HPE 1.2TB SSD
- Network: HPE Ethernet 10Gb 2P 560FLR-SFP+ Adptr

3x HPE ProLiant DL360 Gen9, each with:

- Processors/Cores/Threads: 2/20/40
- Processor Model: 2 x Intel Xeon E5-2640 v4 2.40GHz
- Memory: 256GB
- Controller: Smart HBA H240ar
- Drives:
  - 1 x HPE 800GB SSD
- Network: HPE Ethernet 10Gb 2P 560FLR-SFP+ Adptr

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 Name(s)	HW/SW Configuration	Storage Setup
Worker	HDFS DataNode/Hive Gateway/YARN Node Manager	9	N	bdw[04-12]	• HPE DL380 Gen9. • HW/SW Config (Intel E5-2697Av4, 2, 2.6GHz, 64) • Memory: 192GB • Storage: 16 x 1TB SATA HDD, 1 x 480GB SSD, 1x1.2TB SSD • Network: HPE 560 SFP+10G NIC • OS: RHEL 6.7 • Cloudera CDH 5.6	OS: HPE 480GB 6G SATA SSD, Intermediate/Shuffle /Temp Data/ Distributed FS: 1 x 1.2TB 6G SATA SSD, 16 x HPE 1TB 6G SATA 7.2k HDD
Cloudera Manager Node #1	Cloudera Manager/HDFS Balancer/HDFS Namenode/Hive Gateway/Hive Metastore Server/Hue Server/Cloudera Management Services/Oozie	1	N	bdw01	• HPE DL360 Gen9 Server • HW/SW Config (Intel E5-2640v4, 2, 2.4GHz, 40) • Memory: 256GB • Storage: 1 x 800GB SSD • Network: HPE 560 SFP+10G NIC • OS: RHEL 6.7 • Cloudera CDH 5.6	OS: HPE 800GB 6G SATA SSD
Cloudera Manager Node #2	Hive Gateway/HiveServer2/ZooKeeper Server	1	N	bdw02	• HPE DL360 Gen9 Server • HW/SW Config (Intel E5-2640v4, 2, 2.4GHz, 40) • Memory: 256GB • Storage: 1 x 800GB SSD • Network: HPE 560 SFP+10G NIC • OS: RHEL 6.7 • Cloudera CDH 5.6	OS: HPE 800GB 6G SATA SSD
Cloudera Manager Node #3	HDFS SecondaryNameNode/Hive Gateway/Cloudera Management Service Activity Monitor/ZooKeeper Server	1	N	bdw03	• HPE DL360 Gen9 Server • HW/SW Config (Intel E5-2640v4, 2, 2.4GHz, 40) • Memory: 256GB • Storage: 1 x 800GB SSD • Network: HPE 560 SFP+10G NIC • OS: RHEL 6.7 • Cloudera CDH 5.6	OS: HPE 800GB 6G SATA SSD

2.2 Distributed File System Implementation

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

Cloudera for Apache Hadoop (CDH) 5.6 (fully HDFS compatible at the API level).

2.3 Engine Implementation

The Engine implementation and corresponding version must be disclosed.

Component	Version
Hive	1.1.0
HDFS	2.6.0
YARN	2.6.0
Spark	1.5.0
MapReduce	2.6.0

2.4 Frameworks

Frameworks and Engine used in the benchmark should be disclosed.

Framework	Version
CDH	5.6.0
Hive	1.1.0
HDFS	2.6.0
YARN	2.6.0
Spark	1.5.0
MapReduce	2.6.0

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
1.0.1

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 Performance Summary**

```
**********
TPCx-BB
Result
v1.0.1
**********
INFO: T_LOAD = 2138.322
INFO: T_LD = 0.1 * T_LOAD: 213.83220000000003
INFO: T_PT = 15897.004100627742
INFO: T_T_PUT = 31396.267
INFO: T_TT = 15698.1335
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@3000 = 337.26627326432424
```

- **Run2 Performance Summary**

```
**********
TPCx-BB
Result
v1.0.1
**********
INFO: T_LOAD = 2148.88
INFO: T_LD = 0.1 * T_LOAD: 214.88800000000003
INFO: T_PT = 15821.423876979647
INFO: T_T_PUT = 31493.239
INFO: T_TT = 15746.6195
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@3000 = 337.5239352171591
```

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.

Query Type	Query Number	Power Stream 1	Throughput Stream 1	Throughput Stream 2
Structured	1	247.537	257.619	250.145
	6	703.207	988.188	984.363
	7	373.269	449.706	612.548
	9	367.574	435.588	459.640
	11	188.826	359.363	236.702
	13	327.513	438.327	420.693
	14	83.363	102.559	84.837
	15	126.917	227.329	126.085
	16	892.347	921.282	1,813.950
	17	347.767	509.892	393.250
	20	346.923	350.881	348.877
	21	545.608	618.651	669.811
	22	763.927	769.963	1,085.972
	23	287.463	578.074	275.175
	24	253.615	259.687	251.014
	25	518.737	708.806	725.586
	26	447.177	1,316.671	689.032
	29	517.187	565.990	524.386
Semi-structured	2	1,696.372	3,049.036	1,989.110
	3	901.898	1,273.406	909.371
	4	1,625.996	2,952.027	2,911.395
	5	1,901.160	2,210.128	1,794.893
	8	676.600	921.697	678.503
	12	572.280	554.697	671.150
	30	2,483.757	2,517.748	3,940.520
Unstructured	10	1,819.082	2,489.395	1,834.174
	18	3,705.999	4,017.053	5,173.764
	19	490.883	949.586	488.478
	27	123.783	125.485	314.394
	28	465.442	477.385	529.825

3.5 Validation Test Output

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

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

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/bigBench.properties
- bigBench-configs/conf/userSettings.conf
- bigBench-configs/hive/queries/q28/engineLocalSettings.sql

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 include 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)
3	1	800	2,400
9	1	480	4,320
9	1	1,200	10,800
9	16	1,000	144,000

Total Storage (GB)	161,520
Scale Factor	3,000
Data Storage Ratio	53.84

4.5 Scale Factor to Memory Ratio

The Scale Factor to memory ratio must be disclosed.

Nodes	Memory (GB)	Total (GB)
3	256	768
9	192	1,728

Scale Factor	3,000
Total Memory (GB)	2,496
SF / Memory Ratio	1.20

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@3000	337.26

5.2 Repeatability Run Metric

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

Repeatability Run	
BBQpm@3000	337.52

5.3 Price-Performance Metric

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

Price-Performance	
$/BBQpm@3000	$1,102.94

5.4 Scale Factor

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

Scale Factor
3000

5.5 Stream Count

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

Streams
2

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
Performance	15:55:36.868	57,336.868
Repeatability	15:53:37.490	57,217.490

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	2,138.322	2,148.880
Power Test	23,802.274	23,575.367
Throughput Test	31,396.267	31,493.239

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 for InfoSizing, Inc.

www.sizing.com
20 Kreg Lane
Manitou Springs, CO 80829
719-473-7555.

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 Transaction Processing Performance Council
Certified Auditors

Paul Cao
Hewlett Packard Enterprise
11445 Compaq Center Dr West
Houston, TX 77070

March 29, 2016

I verified the TPC Express Benchmark™ BB v1.0.1 performance of the following configuration:

Platform:	Hewlett Packard Enterprise ProLiant DL for Big Data (with 9x HPE ProLiant DL380 Gen9 and 3x HPE ProLiant DL360 Gen9 Servers)
Operating System:	Red Hat Enterprise Linux Server 6.7
Apache Hadoop Compatible Software:	Cloudera for Apache Hadoop (CDH) 5.6

The results were:

Performance Metric	**337.26 BBQpm@3000GB**
Run Elapsed Time	15:55:36.868 (73,955.222 Seconds)

Cluster		**9x HPE DL380 Gen9, 3x HPE DL360 Gen9 Servers**		
CPUs	DL380s	2 x Intel Xeon Processor E5-2697A v4 (2.60 GHz, 16-core, 40 MB L3)		
	DL360s	2 x Intel Xeon Processor E5-2640 v4 (2.40 GHz, 10-core, 25 MB L3)		
Memory	DL380s	192GB		
	DL360s	256GB		
Storage		***Qty***	***Size***	***Type***
	DL380s	1	480GB	SATA SSD (all nodes)
		1	1.2TB	1.2TB SATA SSD
		16	1TB	7.2K rpm SATA HDD (DL380 nodes)
	DL360s	1	800GB	SATA SSD (1 DL360 node)

In my opinion, these performance results were produced in compliance with the TPC requirements for the benchmark.

The following verification items were given special attention:

- All TPC-provided components were verified to be v1.0.1
- 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 (3000GB) was confirmed to be valid for publication

20 Kreg Lane • Manitou Springs, CO 80829 • 719-473-7555 • 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, Auditor

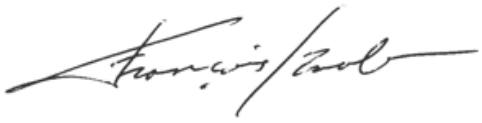

François Raab, President

20 Kreg Lane • Manitou Springs, CO 80829 • 719-473-7555 • www.sizing.com

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-3TB-BDW-3-2016\
Validation Run Files	Supporting-Files-3TB-BDW-3-2016\Validation-run-logs-20160322-155451-hive-sf3000
Performance Run Files	Supporting-Files-3TB-BDW-3-2016\logs-20160323-080841-hive-sf3000
Repeatability Run Files	upporting-Files-3TB-BDW-3-2016\Repeatability-run-logs-20160324-002413-hive-sf3000
Clause 3 - Workload Related Items	
Benchmark Generic Parameters	Supporting-Files-3TB-BDW-3-2016\logs-20160323-080841-hive-sf3000\bigBench-configs\conf\userSettings.conf
Query Parameters used in the benchmark execution Settings	Supporting-Files-3TB-BDW-3-2016\logs-20160323-080841-hive-sf3000\bigBench-configs\hive\conf\queryParameters.sql
Benchmark Global Framework Parameters Settings	Supporting-Files-3TB-BDW-3-2016\logs-20160323-080841-hive-sf3000\bigBench-configs\hive\conf\engineSettings.sql
Benchmark Global Framework Parameters Settings	Supporting-Files-3TB-BDW-3-2016\logs-20160323-080841-hive-sf3000\bigBench-configs\hive\conf\engineSettings.conf
Load Test script	Supporting-Files-3TB-BDW-3-2016\logs-20160323-080841-hive-sf3000\bigBench-configs\hive\population\hiveCreateLoad.sql
Queries specific optimization parameters settings	Supporting-Files-3TB-BDW-3-2016\logs-20160323-080841-hive-sf3000\bigBench-configs\hive\queries\q[01-30]\engineLocalSettings.conf
Queries specific optimization parameters settings	Supporting-Files-3TB-BDW-3-2016\logs-20160323-080841-hive-sf3000\bigBench-configs\hive\queries\q[01-30]\engineLocalSettings.sql
Clause 4 - SUT Related Items	
Data Redundancy report	Supporting-Files-3TB-BDW-3-2016\hdfs-data-redundancy-report.txt
Benchmark execution script	Supporting-Files-3TB-BDW-3-2016\run-all.sh
Hardware and Software Report from a representative node	Supporting-Files-3TB-BDW-3-2016\logs-20160323-080841-hive-sf3000\run-logs\envInfo-hsw04\envInfo.log
All Framework configuration files are included in the Supporting Files Archive	Supporting-Files-3TB-BDW-3-2016\logs-20160323-080841-hive-sf3000\bigBench-configs\hadoop
	Supporting-Files-3TB-BDW-3-2016\logs-20160323-080841-hive-sf3000\bigBench-configs\hive
	Supporting-Files-3TB-BDW-3-2016\logs-20160323-080841-hive-sf3000\bigBench-configs\spark
Clause 5 - Metric and Scale Factor Related Items	
Benchmark Performance Report	Supporting-Files-3TB-BDW-3-2016\logs-20160323-080841-hive-sf3000\run-logs\BigBenchResult.log
Validation Test Report	Supporting-Files-3TB-BDW-3-2016\Validation-run-logs-20160322-155451-hive-sf3000\run-logs\BigBenchResult.log