



**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

(w/ 18x HPE ProLiant DL380 Gen10, 3x HPE ProLiant DL360 Gen10)

using

Cloudera for Apache Hadoop (CDH) 5.16.1

and

Red Hat Enterprise Linux Server 7.5

First Edition

April 2, 2019

Hewlett Packard Enterprise Company (HPE), the Sponsor of this benchmark test, believes that the information in this document is accurate as of the publication date. The information in this document is subject to change without notice. The Sponsor assumes no responsibility for any errors that may appear in this document.

The pricing information in this document is believed to accurately reflect the current prices as of the publication date. However, the Sponsor provides no warranty of the pricing information in this document.

Benchmark results are highly dependent upon workload, specific application requirements, and system design and implementation. Relative system performance will vary 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.

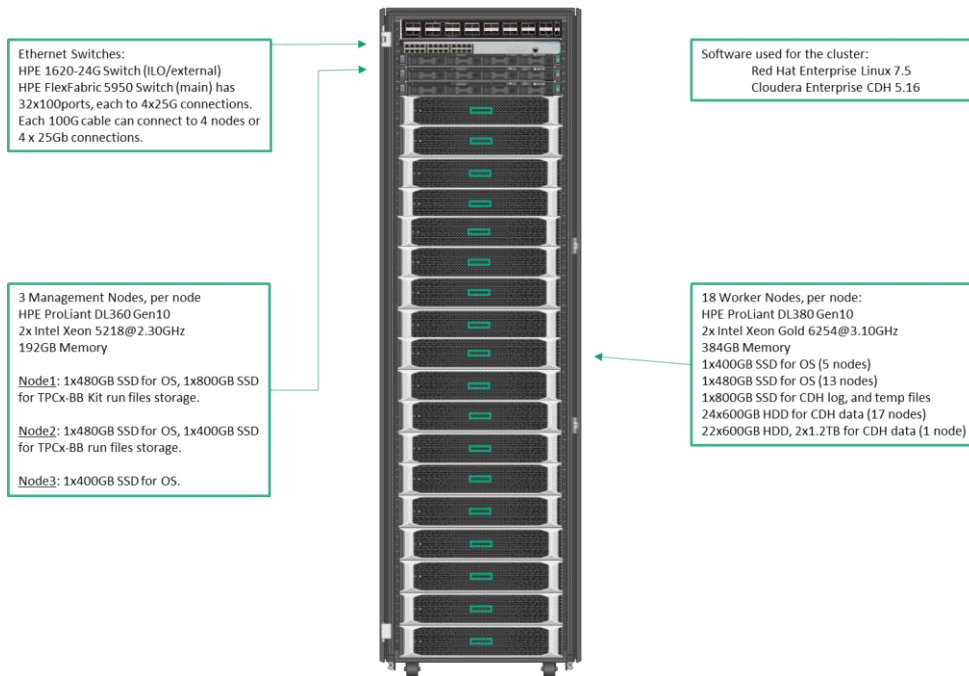
Copyright © 2019 Hewlett Packard Enterprise Company

All rights reserved. Permission is hereby granted to reproduce this document in whole or in part provided the copyright notice printed above is set forth in full text or on the title page of each item reproduced.

 Hewlett Packard Enterprise	Hewlett Packard Enterprise ProLiant DL for Big Data	TPCx-BB Rev. v1.2.0 TPC-Pricing Rev. v2.4.0
		Report Date: April 2, 2019
Total System Cost	TPCx-BB Performance Metric	Price/Performance
917,360 USD	1,759.64 BBQpm@10000	521.34 USD \$/BBQpm@10000

Framework	Operating System	Other Software	Availability Date	Scale Factor	Streams
Cloudera for Apache Hadoop (CDH) 5.16.1	Red Hat Enterprise Linux Server 7.5	None	April 2, 2019	10000	10

System Configuration



Physical Storage/Scale Factor: 28.56	Scale Factor/Physical Memory: 1.34
--------------------------------------	------------------------------------

Servers:	18x HPE ProLiant DL380 Gen10, 3x HPE ProLiant DL360 Gen10
Total Processors/Cores/Threads	42/744/1,488

Server Configuration:	Per HPE ProLiant DL380 Gen10:	Per HPE ProLiant DL360 Gen10:
Processors	2x Intel Xeon Gold 6254 @ 3.10GHz	2x Intel Xeon Gold 5218 @ 2.30GHz
Memory	384GB	192GB
Storage Controller	HPE Smart Array P408i-a SR	HPE Smart Array E208i-a SR
Storage Device	1x HPE 800GB (all nodes) 1x HPE 400GB, 24x HPE 600GB (5 nodes) 1x HPE 480GB, 24x HPE 600GB (12 nodes) 1x HPE 480GB, 24x HPE 600GB, 2x HPE 1.2TB (1 node)	1x HPE 480GB, 1x HPE 800GB Node1 1x HPE 400GB, 1x HPE 480GB Node2 1x HPE 400GB Node3
Network	HPE Eth 10/25Gb 2P 640SFP28	HPE Eth 10/25Gb 2P 640SFP28

Connectivity:	HPE 1620-24G Switch (ILO connection) HPE FlexFabric 5950 100/25G (main connection)
---------------	---



**Hewlett Packard
Enterprise**

**Hewlett Packard Enterprise
ProLiant DL for Big Data**

TPCx-BB Rev. v1.2.0
TPC-Pricing Rev. v2.4.0

Report Date:
April 2, 2019

Description	Price Key	Part Number	Unit Price	Qty	Extended Price	3 Yr Maint Price
Server Hardware						
HPE DL360 Gen10 8SFF CTO Server	1	867959-B21	\$2,099	3	\$6,297	
HPE DL360 Gen10 Xeon-G 5218 Kit	1	P02592-B21	\$2,019	3	\$6,057	
HPE DL360 Gen10 Xeon-G 5218 FIO	1	P02592-L21	\$2,019	3	\$6,057	
HPE 16GB 1Rx4 PC4-2666V-R Smart Kit	1	835955-B21	\$630	36	\$22,680	
HPE 480GB 6G SATA RI-2 SFF SC SSD	1	804593-B21	\$609	2	\$1,218	
HPE 400GB 6G SATA WI-2 SFF SC SSD	1	804665-B21	\$719	2	\$1,438	
HPE 800GB 6G SATA WI-2 SFF SC SSD	1	804671-B21	\$1,389	1	\$1,389	
HPE 800W CS Platinum Plus AC Power Supply	1	865414-B21	\$379	6	\$2,274	
HPE 96W Smart Storage Battery 145mm Cbl	1	875241-B21	\$99	3	\$297	
HPE Eth 10/25Gb 2P 640SFP28	1	817753-B21	\$779	3	\$2,337	
HPE 3Y FC 24x7 DL360 Gen10 SVC	1	H8QF0E	\$1,811	3		\$5,433
HPE iLO Adv incl 3yr TS U E-LTU	1	E6U64ABE	\$469	3		\$1,407
HP W2082a 20-inch LED Monitor (1 + 2 spare)	2	#L8K84AA	\$85	3	\$255	
HPE USB US Keyboard/Mouse Kit	1	631341-B21	\$29	3	\$87	
			Subtotal		\$50,386	\$6,840
HPE DL380 Gen10 24SFF CTO Server	1	868704-B21	\$2,759	18	\$49,662	
HPE DL380 Gen10 High Perf Fan Kit	1	867810-B21	\$239	18	\$4,302	
HPE DL380 Gen10 2SFF Bay Kit	1	826687-B21	\$249	18	\$4,482	
HPE DL380 Gen10 Xeon-G 6254 Kit	1	P02517-B21	\$6,065	18	\$109,170	
HPE DL380 Gen10 Xeon-G 6254 Kit	1	P02517-L21	\$6,065	18	\$109,170	
HPE 32GB 2Rx4 PC4-2666T-R Kit	1	815100-B21	\$1,170	216	\$252,720	
HPE 1.2TB SAS 10K SFF SC HDD	1	718162-B21	\$909	2	\$1,818	
HPE 600GB SAS 6G 10K SFF SC DS HDD	1	872477-B21	\$525	430	\$225,750	
HPE 480GB 6G SATA RI-2 SFF SC SSD	1	804593-B21	\$609	13	\$7,917	
HPE 400GB 6G SATA WI-2 SFF SC SSD	1	804665-B21	\$719	5	\$3,595	
HPE 800GB 6G SATA WI-2 SFF SC SSD	1	804671-B21	\$1,389	18	\$25,002	
HPE Smart Array P408i-a SR	1	804331-B21	\$599	18	\$10,782	
HPE 12Gb DL380 Gen10 SAS Expander Card	1	870549-B21	\$699	18	\$12,582	
HPE 1600W FS Plat Ht Plg LH Pwr Sply Kit	1	830272-B21	\$479	36	\$17,244	
HPE Eth 10/25Gb 2P 640SFP28	1	817753-B21	\$779	18	\$14,022	
HPE 3Y FC 24x7 DL380 Gen10 SVC	1	H8QP7E	\$2,448	18		\$44,064
HPE iLO Adv incl 3yr TS U E-LTU	1	E6U64ABE	\$469	18		\$8,442
			Subtotal		\$848,218	\$52,506

(continued on next page)



**Hewlett Packard
Enterprise**

**Hewlett Packard Enterprise
ProLiant DL for Big Data**

TPCx-BB Rev. v1.2.0
TPC-Pricing Rev. v2.4.0

Report Date:
April 2, 2019

(continued from previous page)

Description	Price Key	Part Number	Unit Price	Qty	Extended Price	3 Yr Maint Price	
Network							
HPE 1620-24G Switch	1	JG913A	\$315	3	\$299		
HPE FlexFabric 5950 32QSFP28 Switch	1	JH321A	\$27,290	1	\$27,290		
HPE 3Y FC 24x7 FF 5950 32Q28 Swch SVC	1	H2HK7E	\$8,003	1		\$8,003	
HPE 58x0AF 650W AC Power Supply	1	JC680A	\$785	2	\$1,570		
HPE X712 Fan Tray	1	JH389A	\$160	6	\$960		
CAT6 UTP Ethernet Network Cable 7ft (42 + 5 spare)	3	2302	\$2	47	\$94		
HPE 100Gb QSFP28 to 4x25Gb SFP28 3m DAC (6 + 2 spare)	1	845416-B21	\$699	8	\$5,592		
			Subtotal		\$35,805	\$8,003	
Rack							
HPE 42U 600x1075mm Adv G2 Kit Plt Rack	1	P9K07A	\$1,179	1	\$1,179		
HPE 24A High Voltage Core Only Corded PDU	1	252663-D74	\$259	2	\$518		
			Subtotal		\$1,697	\$0	
Server Software							
Cloudera Ent Data Eng Ed	1	P9V18AAE	\$12,000	21	\$252,000		
RHEL Svr 2 Sckt/2 Gst 3yr 24x7 E-LTU	1	G3J30AAE	\$3,702	21	\$77,742		
			Subtotal		\$329,742	\$0	
			Total Extended Price		\$1,265,848	\$67,349	
			Total Discounts		\$415,837	\$0	
Sales contact: HPE WW Headquarters, 3000 Hanover St., Palo Alto, CA 94304-1185 (650) 857-1501 or HPE: 855-472-5233					Grand Total	\$850,011	\$67,349

Pricing: 1 = HPE; 2 = store.hp.com; 3 = monoprice.com

⁽¹⁾ All discounts are based on US list prices and for similar quantities and configurations. A 35% discount was based on the overall specific components pricing from vendor 1 in this single quotation. Discounts for similarly sized configurations will be similar to those quoted here, but may vary based on the components in the configuration.

Audited by Doug Johnson of InfoSizing

Three-Year Cost of Ownership	\$917,360
BBQpm@10000	1,759.64
\$/BBQpm@10000	\$ 521.34

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	10
SUT Validation Test	PASS

Performance Run (Run 2)

Overall Run Start Time	2019-03-30 07:25:31.362
Overall Run End Time	2019-03-31 13:54:22.487
Overall Run Elapsed Time	109,731.125
Load Test Start Time	2019-03-30 07:25:31.363
Load Test End Time	2019-03-30 08:05:21.473
Load Test Elapsed Time	2,390.110
Power Test Start Time	2019-03-30 08:05:21.475
Power Test End Time	2019-03-30 12:41:46.982
Power Test Elapsed Time	16,585.507
Throughput Test Start Time	2019-03-30 12:41:46.984
Throughput Test End Time	2019-03-31 13:54:22.486
Throughput Test Elapsed Time	90,755.502
Performance Metric (BBQpm@ 10000)	1,759.64

Repeatability Run (Run 1)

Overall Run Start Time	2019-03-29 00:04:27.737
Overall Run End Time	2019-03-30 06:23:02.071
Overall Run Elapsed Time	109,114.334
Load Test Start Time	2019-03-29 00:04:27.737
Load Test End Time	2019-03-29 00:43:36.735
Load Test Elapsed Time	2,348.998
Power Test Start Time	2019-03-29 00:43:36.737
Power Test End Time	2019-03-29 05:20:44.390
Power Test Elapsed Time	16,627.653
Throughput Test Start Time	2019-03-29 05:20:44.391
Throughput Test End Time	2019-03-30 06:23:02.070
Throughput Test Elapsed Time	90,137.679
Performance Metric (BBQpm@ 10000)	1,764.71



Performance Run Report (Run 2)

TPCx-BB

Result

v1.2

INFO: T_LOAD = 2390.11

INFO: T_LD = 0.1 * T_LOAD: 239.011

INFO: T_PT = 10997.2657832701

INFO: T_T_PUT = 90755.502

INFO: T_TT = 9075.5502

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 = 1759.64806959625

Repeatability Run Report (Run 1)

TPCx-BB

Result

v1.2

INFO: T_LOAD = 2348.998

INFO: T_LD = 0.1 * T_LOAD: 234.8998

INFO: T_PT = 11016.7330817479

INFO: T_T_PUT = 90137.679

INFO: T_TT = 9013.7679

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 = 1764.71410297685

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

Table of Contents

ABSTRACT	9
PREFACE	10
CLAUSE 1: GENERAL ITEMS.....	11
1.1 TEST SPONSOR.....	11
1.2 PARAMETER SETTINGS	11
1.3 CONFIGURATION DIAGRAMS.....	11
CLAUSE 2: SOFTWARE COMPONENTS AND DATASET DISTRIBUTION	13
2.1 ROLES AND DATASET DISTRIBUTION.....	13
2.2 DISTRIBUTED FILE SYSTEM IMPLEMENTATION	14
2.3 ENGINE IMPLEMENTATION	14
2.4 FRAMEWORKS	14
2.5 APPLIED PATCHES	14
CLAUSE 3: WORKLOAD RELATED ITEMS	15
3.1 HARDWARE & SOFTWARE TUNABLE	15
3.2 KIT VERSION	15
3.3 RUN REPORT.....	15
3.4 QUERY ELAPSED TIMES	16
3.5 VALIDATION TEST OUTPUT.....	18
3.6 GLOBAL FRAMEWORK PARAMETERS.....	18
3.7 KIT MODIFICATIONS.....	18
CLAUSE 4: SUT RELATED ITEMS.....	20
4.1 SPECIALIZED HARDWARE/SOFTWARE	20
4.2 FRAMEWORK CONFIGURATION FILES	20
4.3 SUT ENVIRONMENT INFORMATION	20
4.4 DATA STORAGE TO SCALE FACTOR RATIO.....	20
4.5 SCALE FACTOR TO MEMORY RATIO	20
CLAUSE 5: METRICS AND SCALE FACTORS.....	21
5.1 PERFORMANCE RUN METRIC	21
5.2 REPEATABILITY RUN METRIC	21
5.3 PRICE-PERFORMANCE METRIC	21
5.4 SCALE FACTOR.....	21
5.5 STREAM COUNT.....	21
5.6 ELAPSED RUN TIMES.....	22
5.7 ELAPSED TEST TIMES.....	22
AUDITORS' INFORMATION AND ATTESTATION LETTER.....	23
THIRD PARTY PRICE QUOTES.....	26
STORE.HP.COM	26
MONOPRICE.COM.....	26
SUPPORTING FILE INDEX.....	27

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 v1.2.0.

The test was conducted at a Scale Factor of 10000 with 21 nodes ({Node Model}) running Cloudera for Apache Hadoop (CDH) 5.16.1 on Red Hat Enterprise Linux Server 7.5.

Measured Configuration

Company Name	Cluster Node	Virtualization	Operating System
Hewlett Packard Enterprise Company	18x HPE ProLiant DL380 Gen10 3x HPE ProLiant DL360 Gen10	n/a	Red Hat Enterprise Linux Server 7.5

TPC Express Benchmark© Big Bench Metrics

Total System Cost	BBQpm@10000	Price/Performance	Availability Date
917,360 USD	1,759.64	521.34 USD	April 2, 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 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.

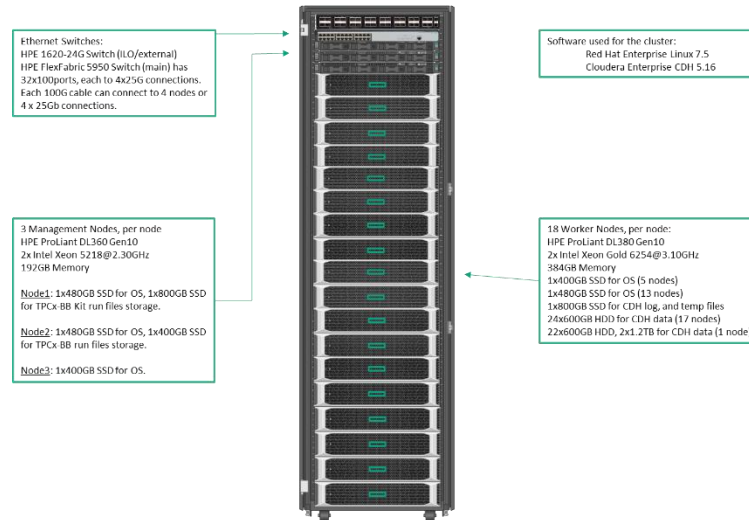
The Supporting Files Archive contains the parameters and options used to configure the components involved in this benchmark.

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: 21
- Total Processors/Cores/Threads: 42/744/1,488
- Total Memory: 7,488
- Total Number of Storage Drives/Devices: 473
- Total Storage Capacity: 285,600

Network connectivity detail:

- HPE 1620-24G Switch (ILO connection), HPE FlexFabric 5950 100/25G (main connection)

Server nodes details:

18x HPE ProLiant DL380 Gen10, each with:

- Processors/Cores/Threads: 2/36/72
- Processor Model: 2x Intel Xeon Gold 6254 @ 3.10GHz
- Memory: 384GB
- Controller: 1x HPE Smart Array P408i-a SR
- Drives:
 - 1x HPE 800GB (all nodes)
 - 1x HPE 400GB, 24x HPE 600GB (5 nodes)
 - 1x HPE 480GB, 24x HPE 600GB (12 nodes)
 - 1x HPE 480GB, 24x HPE 600GB, 2x HPE 1.2TB (1 node)
- Network: HPE Eth 10/25Gb 2P 640SFP28

3x HPE ProLiant DL360 Gen10, each with:

- Processors/Cores/Threads: 2/32/64
- Processor Model: 2 x Intel Xeon Gold 5218 @ 2.30GHz
- Memory: 192GB
- Controller: 1x HPE Smart Array E208i-a SR
- Drives:
 - 1x HPE 480GB, 1x HPE 800GB Node1
 - 1x HPE 400GB, 1x HPE 480GB Node2
 - 1x HPE 400GB Node3
- Network: HPE Eth 10/25Gb 2P 640SFP28

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/Hive Gateway/YARN Node Manager/Spark Gateway	18	N	pan-[04-21]	<ul style="list-style-type: none"> HPE DL380 Gen10 HW/SW Config (Intel Xeon Gold 6254, 2, 3.1GHz, 72) Memory: 384GB Storage: 24 x 600GB SAS HDD, 1 x 480GB or 400GB SSD, 1x800GB SSD Network: HPE Eth 10/25Gb 2P 640SFP28 OS: RHEL 7.5 Cloudera CDH 5.16.1 	OS: HPE 480GB or 400GB SSD, Intermediate/Shuffle/Temp Data/ Distributed FS: 1 x 800GB SSD, 24 x HPE 600GB 6G SAS 10k HDD (except 2 x 1.2TB HDD replacement)
Cloudera Manager Node #1	HDFS Balancer/HDFS Namenode/Hive Gateway/Cloudera Management Services Alert Publisher/Cloudera Management Services Event Server/Hive Metastore Server/Hue Server/Cloudera Management Services/YARN JobHistory Server/YARN ResourceManager/ ZooKeeper Server/Spark Gateway/Spark History	1	N	pan-01	<ul style="list-style-type: none"> HPE DL360 Gen10 Server HW/SW Config (Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz, 64) Memory: 192GB Storage: see storage Network: HPE Eth 10/25Gb 2P 640SFP28 OS: RHEL 7.5 Cloudera CDH 5.16.1 	1x480GB SSD for OS, 1x800GB SSD for TPCx-BB run files storage.
Cloudera Manager Node #2	Hive Gateway/HiveServer2/ZooKeeper Server	1	N	pan-02	<ul style="list-style-type: none"> HPE DL360 Gen10 Server HW/SW Config Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz, 64) Memory: 192GB Storage: see storage Network: HPE 560 SFP+10G NIC OS: RHEL 7.5 Cloudera CDH 5.16.1 	1x480GB SSD for OS, 1x400GB SSD for TPCx-BB run files storage.
Cloudera Manager Node #3	HDFS SecondaryNameNode/Hive Gateway/Cloudera Management Service Activity Monitor/ZooKeeper Server	1	N	pan-03	<ul style="list-style-type: none"> HPE DL360 Gen10 Server HW/SW Config (Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz, 64) Memory: 192GB Storage: 1 x 800GB SSD Network: HPE 560 SFP+10G NIC OS: RHEL 7.5 Cloudera CDH 5.16.1 	1x400GB SSD for OS

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) 5.16.1 (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	1.6.0
MapReduce	2.6.0
ZooKeeper	3.4.5

2.4 Frameworks

Frameworks and Engine used in the benchmark should be disclosed.

Framework	Version
CDH	5.16.1
Hive	1.1.0
HDFS	2.6.0
YARN	2.6.0
Spark	1.6.0
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
v1.2

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
v1.2
*****
INFO: T_LOAD = 2348.998
INFO: T_LD = 0.1 * T_LOAD: 234.8998
INFO: T_PT = 11016.7330817479
INFO: T_T_PUT = 90137.679
INFO: T_TT = 9013.7679
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 = 1764.71410297685
```

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

```
*****
TPCx-BB
Result
v1.2
*****
INFO: T_LOAD = 2390.11
INFO: T_LD = 0.1 * T_LOAD: 239.011
INFO: T_PT = 10997.2657832701
INFO: T_T_PUT = 90755.502
INFO: T_TT = 9075.5502
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 = 1759.64806959625
```

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
Structured	1	206.783	848.146	850.255	682.899	909.372	907.839
	6	522.363	2,647.955	2,757.042	2,646.768	2,501.265	2,622.719
	7	394.986	2,188.051	1,995.036	2,009.641	2,135.107	2,084.495
	9	389.429	2,028.372	2,096.322	1,970.654	2,042.025	2,005.895
	11	129.693	504.274	574.641	681.419	512.348	526.270
	13	198.594	1,090.801	1,091.220	923.583	1,061.172	1,013.568
	14	72.973	236.928	243.356	277.706	305.576	256.932
	15	135.949	357.490	358.411	398.525	333.583	307.240
	16	593.874	2,589.584	3,128.111	3,345.662	3,351.780	3,301.114
	17	316.378	937.812	897.824	846.586	886.321	1,008.856
	20	348.039	1,238.702	1,180.283	1,210.698	1,236.588	1,381.581
	21	596.080	2,891.522	2,948.704	3,170.966	2,842.726	2,905.214
	22	101.465	221.405	108.054	218.471	223.778	238.206
	23	141.533	386.995	293.945	360.501	399.226	362.305
	24	194.192	967.383	1,019.772	1,107.276	1,009.553	983.766
	25	494.210	2,764.301	2,285.165	2,533.306	2,518.367	2,447.438
26	288.802	942.728	746.623	769.179	795.607	729.400	
29	224.804	1,106.709	1,154.995	561.723	1,013.726	1,010.382	
Semi-structured	2	1,980.790	12,796.489	12,926.406	12,836.080	13,102.909	13,303.428
	3	1,009.961	6,557.042	6,799.878	6,106.766	6,821.534	6,863.327
	4	1,590.335	11,356.336	11,410.502	11,391.575	10,587.359	10,449.749
	5	597.136	3,197.330	3,485.120	3,412.552	3,151.035	3,342.347
	8	629.332	4,165.772	3,631.069	4,334.661	3,629.935	3,589.134
	12	332.012	1,682.442	1,713.946	1,936.932	2,027.313	1,587.461
	30	1,911.937	17,352.440	16,884.687	17,267.018	17,006.575	16,520.898
Unstructured	10	327.494	1,116.373	1,220.533	989.612	1,308.683	1,322.589
	18	1,885.520	5,496.997	5,888.388	5,601.941	5,497.672	5,671.143
	19	590.155	1,271.548	1,374.876	1,260.538	1,214.412	1,148.460
	27	65.099	212.842	177.107	198.746	199.378	173.818
	28	315.533	817.471	783.515	709.160	623.239	653.998

Type	Query	Stream 6	Stream 7	Stream 8	Stream 9	Stream 10
Structured	1	934.650	757.875	205.746	896.789	994.639
	6	2,625.496	2,483.108	2,847.379	2,642.648	2,631.321
	7	2,147.151	2,403.889	1,881.323	2,371.660	2,392.317
	9	1,930.524	1,800.302	1,834.492	1,924.360	2,120.748
	11	506.113	526.295	524.570	546.418	571.617
	13	998.127	930.998	1,126.839	946.521	924.660
	14	243.173	87.269	277.318	251.484	293.305
	15	409.338	315.697	348.291	363.925	390.098
	16	3,598.361	3,385.778	3,202.594	3,676.691	3,472.522
	17	754.919	406.854	762.549	864.824	890.636
	20	1,131.099	1,298.730	343.161	1,203.487	1,474.303
	21	2,888.959	2,732.516	2,916.490	2,445.786	2,896.164
	22	274.439	199.427	274.677	266.743	236.672
	23	351.011	317.485	146.461	349.978	351.765
	24	1,074.512	1,080.825	1,133.026	998.879	1,065.761
	25	2,874.250	2,922.959	2,760.857	2,283.055	2,656.151
26	825.354	721.533	709.259	790.439	790.832	
29	1,232.403	1,043.298	952.801	1,072.249	1,061.563	
Semi-structured	2	12,811.175	12,992.428	12,611.117	12,941.236	11,030.056
	3	6,833.531	6,735.104	6,936.484	6,713.224	6,419.288
	4	11,517.798	12,039.676	11,040.208	10,777.516	11,090.576
	5	2,902.829	3,124.297	3,349.956	3,157.362	3,396.590
	8	4,005.756	4,074.344	4,051.021	3,691.623	3,979.830
	12	1,498.994	1,784.477	1,815.630	1,828.441	1,987.629
	30	16,364.432	16,896.144	19,182.353	17,317.580	16,683.812
Unstructured	10	1,231.549	1,147.525	1,303.184	1,347.880	1,248.794
	18	5,738.848	6,025.930	6,080.508	5,799.559	5,448.098
	19	1,137.918	1,133.776	1,293.393	1,351.713	1,167.600
	27	179.999	211.681	200.720	168.344	262.496
	28	768.848	571.880	643.055	692.826	784.216

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/hive/conf/engineSettings.conf
- bigBench-configs/hive/conf/engineSettings.sql
- bigBench-configs/hive/queries/q01/engineLocalSettings.sql
- bigBench-configs/hive/queries/q02/engineLocalSettings.sql
- bigBench-configs/hive/queries/q03/engineLocalSettings.sql
- bigBench-configs/hive/queries/q04/engineLocalSettings.sql
- bigBench-configs/hive/queries/q05/engineLocalSettings.sql
- bigBench-configs/hive/queries/q06/engineLocalSettings.sql
- bigBench-configs/hive/queries/q07/engineLocalSettings.sql
- bigBench-configs/hive/queries/q08/engineLocalSettings.sql
- bigBench-configs/hive/queries/q09/engineLocalSettings.sql
- bigBench-configs/hive/queries/q10/engineLocalSettings.sql
- bigBench-configs/hive/queries/q11/engineLocalSettings.sql
- bigBench-configs/hive/queries/q12/engineLocalSettings.sql
- bigBench-configs/hive/queries/q13/engineLocalSettings.sql
- bigBench-configs/hive/queries/q14/engineLocalSettings.sql
- bigBench-configs/hive/queries/q15/engineLocalSettings.sql
- bigBench-configs/hive/queries/q16/engineLocalSettings.sql
- bigBench-configs/hive/queries/q17/engineLocalSettings.sql
- bigBench-configs/hive/queries/q18/engineLocalSettings.sql
- bigBench-configs/hive/queries/q19/engineLocalSettings.sql
- bigBench-configs/hive/queries/q20/engineLocalSettings.sql
- bigBench-configs/hive/queries/q21/engineLocalSettings.sql
- bigBench-configs/hive/queries/q22/engineLocalSettings.sql
- bigBench-configs/hive/queries/q23/engineLocalSettings.sql
- bigBench-configs/hive/queries/q24/engineLocalSettings.sql
- bigBench-configs/hive/queries/q25/engineLocalSettings.sql
- bigBench-configs/hive/queries/q26/engineLocalSettings.sql
- bigBench-configs/hive/queries/q27/engineLocalSettings.sql
- bigBench-configs/hive/queries/q28/engineLocalSettings.sql
- bigBench-configs/hive/queries/q29/engineLocalSettings.sql
- bigBench-configs/hive/queries/q30/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 from 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)
7	1	400	2,800
15	1	480	7,200
17	24	600	244,800
1	22	600	13,200
19	1	800	15,200
1	2	1,200	2,400
Total Storage (GB)			285,600
Scale Factor			10000
Data Storage Ratio			28.56

4.5 Scale Factor to Memory Ratio

The Scale Factor to memory ratio must be disclosed.

Nodes	Memory (GB)	Total (GB)
3	192	576
18	384	6,912
Scale Factor		10000
Total Memory (GB)		7,488
SF / Memory Ratio		1.34

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 1,759.64

5.2 Repeatability Run Metric

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

Repeatability Run
BBQpm@10000 1,764.71

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 521.34

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
10

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	01 06:18:34.334	109,114.334
Run 2	01 06:28:51.125	109,731.125

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	2,390.110	2,348.998
Power Test	16,585.507	16,627.653
Throughput Test	90,755.502	90,137.679

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.

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

April 1, 2019

I verified the TPC Express Benchmark™ BB v1.2.0 performance of the following configuration:

Platform: Hewlett Packard Enterprise ProLiant DL for Big Data
(w/ 18x HPE ProLiant DL380 Gen 10, 3x HPE ProLiant DL360 Gen10)
Operating System: Red Hat Enterprise Linux Server 7.5
Apache Hadoop Cloudera for Apache Hadoop (CDH) 5.16.1
Compatible Software:

The results were:

Performance Metric 1,759.64 BBQpm@1000GB
Run Elapsed Time 01 06:28:51.125 (109,731.125 Seconds)

**Cluster 18x HPE ProLiant DL380 Gen 10 (Data nodes),
3x HPE ProLiant DL360 Gen10 (Management nodes)**

CPU	2 x Intel Xeon Gold 6254 (3.10 GHz, 18-core, 24.75 MB L3) (Data nodes)		
	2 x Intel Xeon Gold 5218 (2.30 GHz, 16-core, 22 MB L3) (Mgmt. nodes)		
Memory	384GB (Data nodes), 192GB (Mgmt. nodes)		
Storage	Qty	Size	Type
	7	400GB	6G SATA SSD (OS; 2 Mgmt. nodes, 5 Data nodes)
	15	480GB	6G SATA SSD (OS/misc, 2 Mgmt. nodes, 13 Data nodes)
	430	600GB	6G SAS 10K HDD (Data, Data nodes)
	19	800GB	6G SATA SSD (Temp, Data nodes, 1 Mgmt. node)
	2	1.2TB	SAS 10K HDD (Data, 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 v1.2.0
- 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
- 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:

From the TPCx-BB Kit's README:

Q28 Depending on the Hadoop distribution version can fail automated Engine Validation due to empty space characters when the output is written to HDFS. Manually open the result file and validate the reference values and written values.

Query 28 failed automated Engine Validation. A manual validation was performed as part of this audit to confirm the only differences were due to white space.

Respectfully Yours,



Doug Johnson, TPC Auditor

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

Third Party Price Quotes


store.hp.com

Not secure | www8.hp.com/us/en/monitors/product-details/7734122#!8&tab=specs

hp Laptops & 2-in-1s Desktops Printers 3D Print Ink & toner Displays & accessories Business solutions Support

OVERVIEW SPECS REVIEWS ACCESSORIES SUPPORT WHERE TO BUY

MONITORS / HP W2082A 20-INCH LED BACKLIT MONITOR



HP W2082a 20-inch LED Backlit Monitor


★★★★★ (0) Write a review | #LBK84AA

Performance with Style

SEE ALL SPECS

\$84.99*

WHERE TO BUY



Product image may differ from actual product


monoprice.com

Monoprice, Inc. [US] | https://www.monoprice.com/product?p_id=2302&gclid=Cj0KCQjwhPfkBRD0ARIsAAcYrcEZnmrMRW_XsvHePlq5yn3S5rZH4leV6RKbdPAXNNz2DDvKs_n6t8aArBDEALw_wcB

MONOPRICE Find Amazing Products Help Business

Shop Featured Products My Account Quick Order My Cart

HOME / Cables / Networking Patch Cables / Cat6 Ethernet Cables



Monoprice Cat6 Ethernet Patch Cable - Snagless RJ45, Stranded, 550Mhz, UTP, Pure Bare Copper Wire, 24AWG, 7ft, Black

★★★★★ 3050 Reviews | 20 Questions, 27 Answers Product # 2302

\$1.99

Length: 0.5ft 1ft 2ft 3ft 5ft 7ft 10ft 14ft 20ft 25ft 30ft 50ft 75ft 100ft

Color: [Color Selection]

Qty: 1 ADD TO CART ADD TO WISHLIST

In Stock This item should ship today (3/29/2019) if ordered within 5 hours 31 minutes

(Log in to see Member Pricing)

Volume Pricing				
Qty: 1	Qty: 2-9	Qty: 10-19	Qty: 20-49	Qty: 50+
\$1.99	\$1.97	\$1.95	\$1.93	\$1.89

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-CLX-4-2019\
Validation Run Files	Supporting-Files-10TB-CLX-4-2019\logs-20190328-232831-hive-sf10000-Validation
Performance Run Files	Supporting-Files-10TB-CLX-4-2019\logs-20190331-135757-hive-sf10000-run2
Repeatability Run Files	Supporting-Files-10TB-CLX-4-2019\logs-20190330-062639-hive-sf10000-run1
Clause 3 - Workload Related Items	
Benchmark Generic Parameters	Supporting-Files-10TB-CLX-4-2019\logs-20190331-135757-hive-sf10000-run2\bigBench-configs\conf\userSettings.conf
Query Parameters used in the benchmark execution Settings	Supporting-Files-10TB-CLX-4-2019\logs-20190331-135757-hive-sf10000-run2\bigBench-configs\hive\conf\queryParameters.sql
Benchmark Global Framework Parameters Settings	Supporting-Files-10TB-CLX-4-2019\logs-20190331-135757-hive-sf10000-run2\bigBench-configs\hive\conf\engineSettings.sql
Benchmark Global Framework Parameters Settings	Supporting-Files-10TB-CLX-4-2019\logs-20190331-135757-hive-sf10000-run2\bigBench-configs\hive\conf\engineSettings.conf
Load Test script	Supporting-Files-10TB-CLX-4-2019\logs-20190331-135757-hive-sf10000-run2\bigBench-configs\hive\population\hiveCreateLoad.sql
Queries specific optimization parameters settings	Supporting-Files-10TB-CLX-4-2019\logs-20190331-135757-hive-sf10000-run2\bigBench-configs\hive\queries\q[01-30]\engineLocalSettings.conf
Queries specific optimization parameters settings	Supporting-Files-10TB-CLX-4-2019\logs-20190331-135757-hive-sf10000-run2\bigBench-configs\hive\queries\q[01-30]\engineLocalSettings.sql
Clause 4 - SUT Related Items	
Data Redundancy report	Supporting-Files-10TB-CLX-4-2019\data-redundancy-report.txt
Benchmark execution script	Supporting-Files-10TB-CLX-4-2019\run-all.sh
Hardware and Software Report from a representative node	Supporting-Files-10TB-CLX-4-2019\logs-20190331-135757-hive-sf10000-run2\logs-20190331-135757-hive-sf10000\run-logs\envInfo-pan-08
All Framework configuration files are included in the Supporting Files Archive	Supporting-Files-10TB-CLX-4-2019\logs-20190331-135757-hive-sf10000-run2\logs-20190331-135757-hive-sf10000\run-logs\envInfo-pan-08\envInfo-pan-08\hadoop
	Supporting-Files-10TB-CLX-4-2019\logs-20190331-135757-hive-sf10000-run2\logs-20190331-135757-hive-sf10000\run-logs\envInfo-pan-08\envInfo-pan-08\hive
	Supporting-Files-10TB-CLX-4-2019\logs-20190331-135757-hive-sf10000-run2\logs-20190331-135757-hive-sf10000\run-logs\envInfo-pan-08\envInfo-pan-08\spark
Clause 5 - Metric and Scale Factor Related Items	
Benchmark Performance Report	Supporting-Files-10TB-CLX-4-2019\logs-20190331-135757-hive-sf10000-run2\logs-20190331-135757-hive-sf10000\run-logs\BigBenchResult.log
Validation Test Report	Supporting-Files-10TB-CLX-4-2019\logs-20190328-232831-hive-sf10000-Validation\run-logs\BigBenchResult.log