Huawei Technologies Co., LTD.

TPC Express Benchmark™ Big Bench (TPCx-BB)

Full Disclosure Report

for

Huawei FusionServer for Big Data

(with 16x RH2288H V3)

using

Cloudera for Apache Hadoop (CDH) 5.8

and

Red Hat Enterprise Linux Server 6.7

Second Edition

February 16, 2017

***Huawei Technologies Co., LTD. (Huawei)**, 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.

Huawei and the Huawei Logo are trademarks of Huawei Technologies Co., LTD. 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 Huawei and any other company.

TPC Benchmark™, TPCx-BB and BBQpm, are registered certification marks of the Transaction Processing Performance Council.

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

Copyright © 2017 Huawei Technologies Co., LTD.

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.

HUAWEI	**Huawei FusionServer for Big Data**	TPCx-BB Rev. 1.2.0 TPC-Pricing Rev. 2.0.0 Report Date: February 16, 2017
Total System Cost	TPCx-BB Performance Metric	Price/Performance
395,092 USD	**366.69** BBQpm@1000	**1,077.44 USD** $/BBQpm@1000

Framework	Operating System	Other Software	Availability Date	Scale Factor	Streams
Cloudera for Apache Hadoop (CDH) 5.8	Red Hat Enterprise Linux Server 6.7	None	December 7, 2016	1000	8

System Configuration

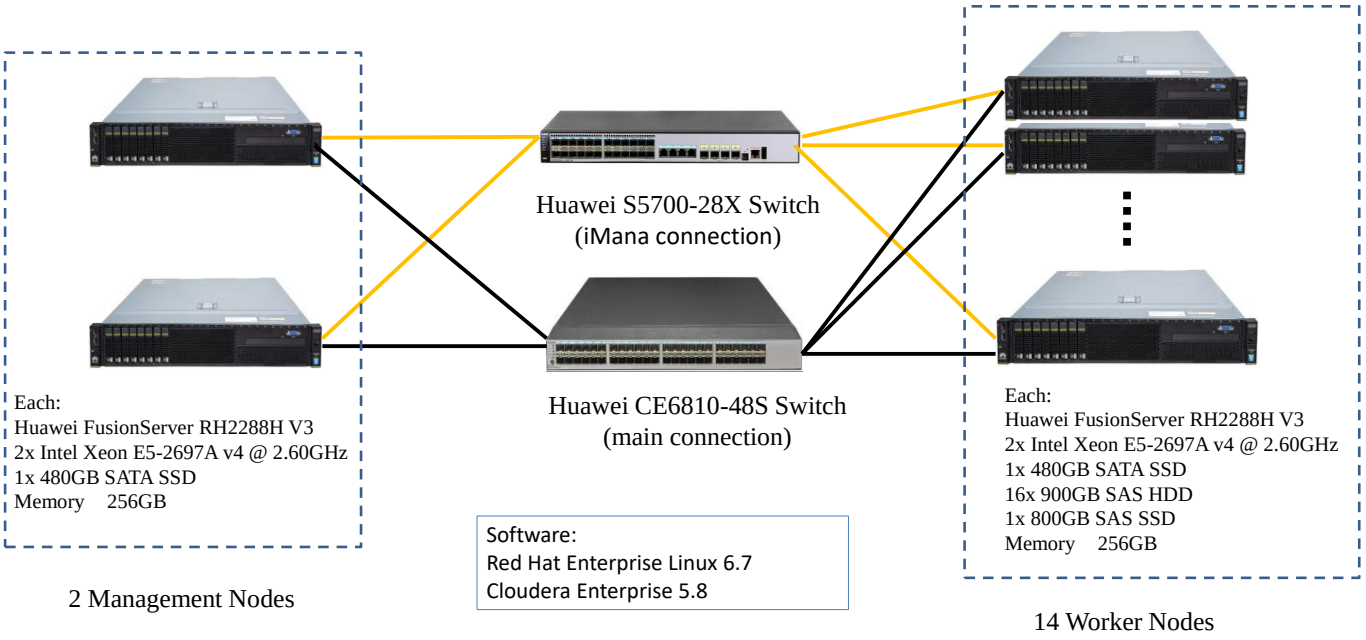

Physical Storage/Scale Factor: 220.48	Scale Factor/Physical Memory: 0.24

Servers: Total Processors/Cores/Threads	16x RH2288H V3 32/512/1024
Server Configuration: Processors Memory Storage Controller Storage Device Network	**Per RH2288H V3:** 2x Intel Xeon E5-2697A v4 @ 2.60GHz 256GB LSI 3108 RAID Card 1x 480GB SATA SSD (all nodes) 16x 900GB SAS HDD (worker nodes) 1x 800GB SAS HDD (worker nodes) Intel Ethernet 2X10GE(82599)-SFP+ Adapter
Connectivity:	Huawei CE6810-48S Switch (main connection), Huawei S5700-28X Switch (iMana connection)

Description	Part Number	Source	Unit Price	Qty	Extended Price	3 Year Maint. Price
Server Hardware - Management Nodes						
RH2288H V3 (25HDD EXP Chassis)H22H-03	02310YJW	1	$ 1,706.60	2	$ 3,413.20	
Onboard Network card,2X10GE Port(82599)	02311EUU	1	$ 1,139.50	2	$ 2,279.00	
8056 Plus Fan module	02310YKP	1	$ 45.58	8	$ 364.64	
3*8X Riser Card Module	02310YKQ	1	$ 126.14	2	$ 252.28	
2*2.5" Rear Hard Disk Backplane Module	02311DUP	1	$ 62.54	2	$ 125.08	
750W platinum AC power supply unit	02310QWX	1	$ 295.74	4	$ 1,182.96	$ 2,168.00
Intel Xeon E5-2697A v4(with Heatsink)	02311NFD	1	$ 6,857.14	4	$ 27,428.56	
32GB,288pin,0.83ns,2400000KHz,1.2V,ECC,2Rank(2G*4bit)	6200214	1	$ 634.94	16	$ 10,159.04	
480GB,SATA 6Gb/s, Mixed Use, 2.5inch,VE Series	02311LTP	1	$ 1,260.34	2	$ 2,520.68	
LSI 3108 RAID CARD-Support SuperCap	02310YMF	1	$ 694.30	2	$ 1,388.60	
LSI Flash Card,8GB,TFM,Supercap and 620mm Cable Moudle	02311BNX	1	$ 386.90	2	$ 773.80	
Network Card, Gigabit, RJ45 Copper, 2 ports, PCIE 2.0 x4-8086-1521-2	6310070	1	$ 60.92	2	$ 121.84	
			SubTotal		$ 50,009.68	$ 2,168.00
Server Hardware - Worker Nodes						
RH2288H V3 (25HDD EXP Chassis)H22H-03	02310YJW	1	$ 1,706.60	14	$ 23,892.40	
Onboard Network card,2X10GE Port(82599)	02311EUU	1	$ 1,139.50	14	$ 15,953.00	
8056 Plus Fan module	02310YKP	1	$ 45.58	56	$ 2,552.48	
3*8X Riser Card Module	02310YKQ	1	$ 126.14	14	$ 1,765.96	
2*2.5" Rear Hard Disk Backplane Module	02311DUP	1	$ 62.54	14	$ 875.56	
750W platinum AC power supply unit	02310QWX	1	$ 295.74	28	$ 8,280.72	
Intel Xeon E5-2697A v4(with Heatsink)	02311NFD	1	$ 6,857.14	28	$ 191,999.92	$ 15,176.00
32GB,288pin,0.83ns,2400000KHz,1.2V,ECC,2Rank(2G*4bit)	6200214	1	$ 634.94	112	$ 71,113.28	
480GB,SATA 6Gb/s,Mixed Use,2.5inch,VE Series	02311LTP	1	$ 1,260.34	14	$ 17,644.76	
900GB,SAS 12Gb/s,10000rpm,2.5inch	02311HAL	1	$ 592.54	224	$ 132,728.96	
SSD,800GB,SAS 12Gb/s,Write Intensive	02311HAG	1	$ 3,433.34	14	$ 48,066.76	
LSI 3108 RAID CARD-Support SuperCap	02310YMF	1	$ 694.30	14	$ 9,720.20	
LSI Flash Card,8GB,TFM,Supercap and 620mm Cable Moudle	02311BNX	1	$ 386.90	14	$ 5,416.60	
Network Card, Gigabit, RJ45 Copper, 2 ports, PCIE 2.0 x4-8086-1521-2	6310070	1	$ 60.92	14	$ 852.88	
			SubTotal		$ 530,863.48	$ 15,176.00

(continued next page)

Description	Part Number	Source	Unit Price	Qty	Extended Price	3 Year Maint. Price
(continued from previous page)						
Network						
Huawei S5700-28X-LI-AC Switch	2354215	1	$ 2,730.56	1	$ 2,730.56	
Huawei CE6810-48S-LI Switch	02350AQC	1	$ 8,776.80	1	$ 8,776.80	
Optical Transceiver,SFP+,10G,Multi-mode Module	2318169	1	$ 270.40	16	$ 4,326.40	$ 6,192.72
HUAWEI VTL6900 42U AC CABINET (RACK-42U-VTL)	2114327	1	$ 2,139.99	1	$ 2,139.99	
Patch cord-LC/PC-LC/PC-Multimode-A1b-2mm-10m-PVC-Orange	14130222	1	$ 14.40	20	$ 288.00	
			SubTotal		$ 18,261.75	$ 6,192.72
Monitor / Keyboard / Mouse (includes 2 spares each)						
Lenovo ThinkVision T2424p 23.8-inch FHD LED Backlit LCD Monitor	9SIAA0T56C4340	2	$ *198.19*	3	$ 594.57	
Lenovo 73P2620 Black 104 Normal Keys USB Wired Standard Keyboard	9SIA1N82KJ2787	2	$ *40.79*	3	$ 122.37	
lenovo 41U3074 Black 1 x Wheel USB Wired Laser 2000 dpi Mouse	N82E16826509009	2	$ *30.99*	3	$ 92.97	
			SubTotal		$ 809.91	$ -
Server Software						
Cloudera Ent Basic Ed 1yr 24x7	G7M 27A	1	$ 2,304.00	48	$ 110,592.00	Included
RHEL Svr 2 Sckt/2 Gst 3yr 24x7 E-LTU	G3J30AAE	1	$ 3,889.00	16	$ 62,224.00	Included
			SubTotal		$ 172,816.00	$ -
			Total Extended Price		$772,760.82	$ 23,536.72
			Total Discounts		$ 389,437.69	$ 11,768.36
			Totals		$383,323.13	$ 11,768.36

Pricing:1 = Huawei; 2 = newegg.com

(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	**$395,092**
BBQ pm@1000	**366.69**
$/BBQ pm@1000	**$1,077.44**

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	1000
Streams	8
SUT Validation Status	TRUE

Performance Run

Overall Run Start Time	2016-12-13 07:20:27.081
Overall Run End Time	2016-12-13 16:45:05.843
Overall Run Elapsed Time	33,878.762
Load Test Start Time	2016-12-13 07:20:27.082
Load Test End Time	2016-12-13 07:32:58.426
Load Test Elapsed Time	751.344
Power Test Start Time	2016-12-13 07:32:58.431
Power Test End Time	2016-12-13 10:33:27.374
Power Test Elapsed Time	10,828.943
Throughput Test Start Time	2016-12-13 10:33:27.375
Throughput Test End Time	2016-12-13 16:45:05.843
Throughput Test Elapsed Time	22,298.468
Performance Metric (BBQpm@1000)	366.6978534

Repeatability Run

Overall Run Start Time	2016-12-13 16:52:43.512
Overall Run End Time	2016-12-14 02:09:07.379
Overall Run Elapsed Time	33,383.867
Load Test Start Time	2016-12-13 16:52:43.514
Load Test End Time	2016-12-13 17:05:22.502
Load Test Elapsed Time	758.988
Power Test Start Time	2016-12-13 17:05:22.507
Power Test End Time	2016-12-13 20:01:27.468
Power Test Elapsed Time	10,564.961
Throughput Test Start Time	2016-12-13 20:01:27.470
Throughput Test End Time	2016-12-14 02:09:07.379
Throughput Test Elapsed Time	22,059.909
Performance Metric (BBQpm@1000)	370.5561682

Run Report – Run 1

```
************
TPCx-BB
Result
v1.2
************
INFO: T_LOAD = 751.344
INFO: T_LD = 0.1 * T_LOAD: 75.13440000000001
INFO: T_PT = 8381.956751650918
INFO: T_T_PUT = 22298.468
INFO: T_TT = 2787.3085
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@1000 = 366.697853361945
```

Run Report – Run 2

```
************
TPCx-BB
Result
v1.2
************
INFO: T_LOAD = 758.988
INFO: T_LD = 0.1 * T_LOAD: 75.89880000000001
INFO: T_PT = 8291.716624644496
INFO: T_T_PUT = 22059.909
INFO: T_TT = 2757.488625
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@1000 = 370.5561682041379
```

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 1.2.0.

The test was conducted at a Scale Factor of 1000 with 16 nodes (16x RH2288H V3) running Cloudera for Apache Hadoop (CDH) 5.8 on Red Hat Enterprise Linux Server 6.7.

Measured Configuration

Company Name	Cluster Node	Virtualization	Operating System
Huawei Technologies Co., LTD.	16x RH2288H V3	n/a	Red Hat Enterprise Linux Server 6.7

TPC Express Benchmark© Big Bench Metrics

Total System Cost	BBQpm@1000	Price/Performance	Availability Date
395,092 USD	366.69	1,077.44 USD	December 7, 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 Huawei Technologies Co., LTD.

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

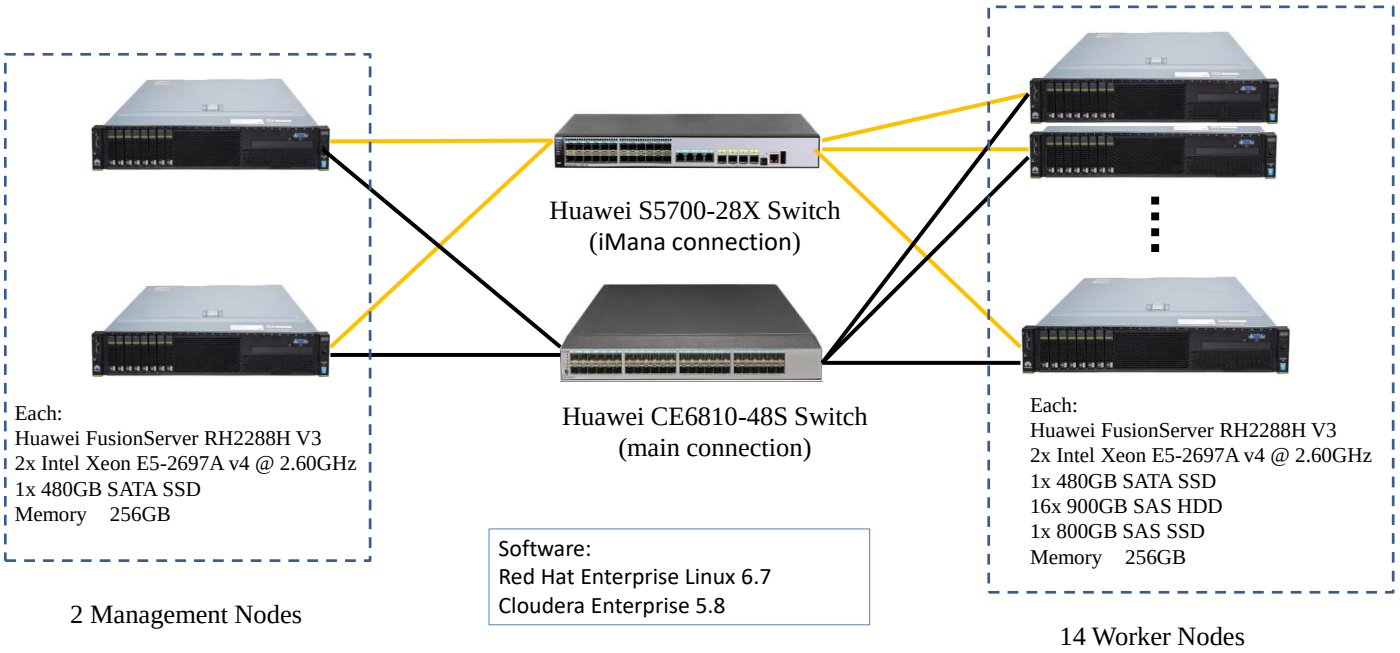

The measured configuration consisted of:

- Total Nodes: 16
- Total Processors/Cores/Threads: 32/512/1024
- Total Memory: 4,096GB
- Total Number of Storage Drives/Devices: 254
- Total Storage Capacity: 220,480GB

Network connectivity detail:

- Huawei CE6810-48S Switch (main connection), Huawei S5700-28X Switch (iMana connection)

16x RH2288H V3, each with:

- Processors/Cores/Threads: 2/32/64
- Processor Model: 2x Intel Xeon E5-2697A v4 @ 2.60GHz
- Memory: 256GB
- Controller: 1 x LSI 3108 RAID Card
- Drives:
  - 1 x 480GB SATA SSD (all nodes)
  - 16 x 900GB SAS HDD (worker nodes)
  - 1 x 800GB SAS HDD (worker nodes)
- Network: Intel Ethernet 2X10GE(82599)-SFP+ Adapter

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 Data Node/YARN Node Manager	13	N	test[3-15]	• Huawei FusionServer RH2288H V3 • HW/SW Config (Intel E5-2697Av4,2, 2.6GHz, 64) • Memory: 256GB • Storage: 16 x 900TB SAS HDD, 1 x 480GB SSD, 1 x 900GB SAS SSD • Network: Intel 82599 SFP+10G • OS: RHEL 6.7 • Cloudera CDH 5.8	OS: 480GB 6G SATA SSD, Intermedia te/Shuffle /Temp Da ta / Dis tributed FS: 1 x 900GB 12G SAS SSD, 16 x 900GB 12G SAS 10k HDD
Worker	HDFS Data Node/YARN Node Manager/ZooKeeper Server	1	N	test2	• Huawei FusionServer RH2288H V3 • HW/SW Config (Intel E5-2697Av4,2, 2.6GHz, 64) • Memory: 256GB • Storage: 16 x 900TB SAS HDD, 1 x 480GB SSD, 1 x 900GB SAS SSD • Network: Intel 82599 SFP+10G • OS: RHEL 6.7 • Cloudera CDH 5.8	OS: 480GB 6G SATA SSD, Intermedia te/Shuffle /Temp Da ta / Dis tributed FS: 1 x 900GB 12G SAS SSD, 16 x 900GB 12G SAS 10k HDD
Cloudera Manager Node #1	HDFSBalancer/HDFS Namenode/Hive Gateway/Hive Metastore Server/Hive Server2/Hue Server/Cloudera Management Services/YARN JobHistory Server/YARN ResourceManager/Zoo Keeper Server/Spark Gateway/Spark History	1	N	test0	• Huawei FusionServer RH2288H V3 • HW/SW Config (Intel E5-2697Av4, 2, 2.6GHz, 64) • Memory: 256GB • Storage: 1 x 480GB SSD • Network: Intel 82599 SFP+10G • OS: RHEL 6.7 • Cloudera CDH 5.8	OS: 480GB 6G SATA SSD
Cloudera Manager Node #2	HDFS Secondary NameNode /Hive Gateway /Spark Gateway /ZooKeeper Server	1	N	test1	• Huawei FusionServer RH2288H V3 • HW/SW Config (Intel E5-2697Av4, 2, 2.6GHz, 64) • Memory: 256GB • Storage: 1 x 480GB SSD • Network: Intel 82599 SFP+10G • OS: RHEL 6.7 • Cloudera CDH 5.8	OS: 480GB 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) 5.8 (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

2.4 Frameworks

Frameworks and Engine used in the benchmark should be disclosed.

Framework	Version
CDH	5.8.0
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
1.2

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 Performance Summary**

```
**********
TPCx-BB
Result
v1.2
**********
INFO: T_LOAD = 751.344
INFO: T_LD = 0.1 * T_LOAD: 75.13440000000001
INFO: T_PT = 8381.956751650918
INFO: T_T_PUT = 22298.468
INFO: T_TT = 2787.3085
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@1000 = 366.697853361945
```

- **Run2 Performance Summary**

```
**********
TPCx-BB
Result
v1.2
**********
INFO: T_LOAD = 758.988
INFO: T_LD = 0.1 * T_LOAD: 75.89880000000001
INFO: T_PT = 8291.716624644496
INFO: T_T_PUT = 22059.909
INFO: T_TT = 2757.488625
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@1000 = 370.5561682041379
```

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	Throughput							
		Stream 1	Stream 1	Stream 2	Stream 3	Stream 4	Stream 5	Stream 6	Stream 7	Stream 8
Structured	1	398.155	761.121	435.936	925.525	629.766	477.382	1,018.101	1,098.869	526.941
	6	225.010	598.115	330.700	499.869	537.440	222.719	354.508	239.914	390.382
	7	313.241	695.423	543.193	831.710	650.452	528.970	669.999	454.809	673.057
	9	123.744	282.068	157.183	166.312	293.199	302.536	236.149	237.889	289.138
	11	292.344	674.646	545.533	712.598	528.333	788.226	345.795	816.993	697.937
	13	181.069	319.467	337.535	381.579	183.286	485.831	219.335	218.331	349.030
	14	51.892	99.292	77.710	67.045	80.313	52.710	56.676	89.637	55.340
	15	99.304	195.513	112.954	120.791	205.608	101.444	170.478	111.630	146.083
	16	345.863	1,043.555	812.113	540.677	405.933	869.651	987.329	932.186	887.199
	17	206.418	466.832	421.287	306.620	250.545	511.444	442.669	321.729	357.165
	20	2,009.982	3,118.740	3,248.719	3,112.853	2,589.292	3,184.223	3,690.596	3,257.390	3,109.782
	21	582.390	1,236.402	869.102	1,102.063	1,243.417	818.936	766.278	1,195.979	1,331.877
	22	207.277	409.120	288.091	523.067	356.589	431.586	412.733	572.938	218.723
	23	313.238	562.577	802.456	780.846	742.658	320.317	565.913	814.186	665.836
	24	553.346	921.079	1,370.066	1,736.277	1,739.603	1,777.685	856.047	1,100.241	906.545
	25	180.843	440.238	463.328	289.303	431.950	305.534	368.783	362.278	177.657
	26	168.265	504.015	272.573	439.519	423.374	420.662	434.159	250.614	253.098
	29	264.164	358.942	649.588	532.936	376.528	359.093	536.148	304.020	608.537
Semi-structured	2	187.113	383.536	447.201	332.316	522.290	437.422	485.230	400.795	329.063
	3	519.505	1,558.482	1,582.119	1,005.272	1,516.920	1,576.463	1,248.556	1,153.868	1,729.648
	4	363.466	772.520	1,017.383	456.109	973.808	697.736	979.367	808.057	755.471
	5	421.239	689.662	922.483	1,350.790	943.777	1,134.719	1,359.101	976.196	586.526
	8	581.480	790.526	1,473.691	1,370.466	942.581	1,485.549	1,596.633	1,436.569	1,882.911
	12	292.211	581.664	644.035	787.972	475.633	383.282	333.477	671.878	734.700
	30	214.142	556.843	427.953	261.208	614.738	437.712	596.646	215.814	412.956
Unstructured	10	273.517	526.511	572.513	314.327	442.706	561.985	504.338	542.780	401.682
	18	101.365	172.882	239.516	143.653	199.020	113.246	160.384	246.857	188.951
	19	397.313	400.703	587.018	442.933	787.857	604.099	635.356	485.233	804.467
	27	251.969	491.929	602.336	480.618	450.375	631.458	459.624	473.207	412.239
	28	708.993	1,487.023	1,584.103	1,599.102	2,118.734	2,275.809	1,124.660	1,976.371	2,138.937

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/q10/engineLocalSettings.sql
- 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)
16	1	480	7,680
14	1	800	11,200
14	16	900	201,600

Total Storage (GB)	220,480
Scale Factor	1,000
Data Storage Ratio	220.48

4.5 Scale Factor to Memory Ratio

The Scale Factor to memory ratio must be disclosed.

Nodes	Memory (GB)	Total (GB)
16	256	4,096

Scale Factor	1,000
Total Memory (GB)	4,096
SF / Memory Ratio	0.24

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@1000	366.6978534

5.2 Repeatability Run Metric

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

Repeatability Run	
BBQpm@1000	370.5561682

5.3 Price-Performance Metric

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

Price-Performance	
$/BBQpm@1000	$1,077.44

5.4 Scale Factor

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

Scale Factor
1000

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	09:24:38.762	33,878.762
Run 2	09:16:23.867	33,383.867

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	751.344	758.988
Power Test	10,828.943	10,564.961
Throughput Test	22,298.468	22,059.909

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 Transaction Processing Performance Council
Certified Auditors

Mr. Long Qiu
President of the Server Product Line
Huawei Technologies Co. Ltd
Huawei Base D1-1A01
Bantian, Longgang District
Shenzhen, China 518129

December 28, 2016

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

Platform:	Huawei FusionServer for Big Data (with 16x RH2288H V3 Servers)
Operating System:	Red Hat Enterprise Linux Server 6.7
Apache Hadoop Compatible Software:	Cloudera for Apache Hadoop (CDH) 5.8

The results were:

Performance Metric	**366.69 BBQpm@1000GB**
Run Elapsed Time	9:24:38.762 (33,878.762 Seconds)

Cluster	**16x Huawei FusionServer RH2288H V3 Servers**		
CPUs	2 x Intel Xeon Processor E5-2697A v4 (2.60 GHz, 16-core, 40 MB L3)		
Memory	256GB		
Storage	***Qty***	***Size***	***Type***
	1	480GB	SATA SSD (all nodes)
	1	800GB	SAS SSD (worker nodes)
	16	900GB	SAS SSD (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 v1.2.0
- 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 (1000GB) was confirmed to be valid for publication
- All validation queries executed successfully and produced compliant results

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

- 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

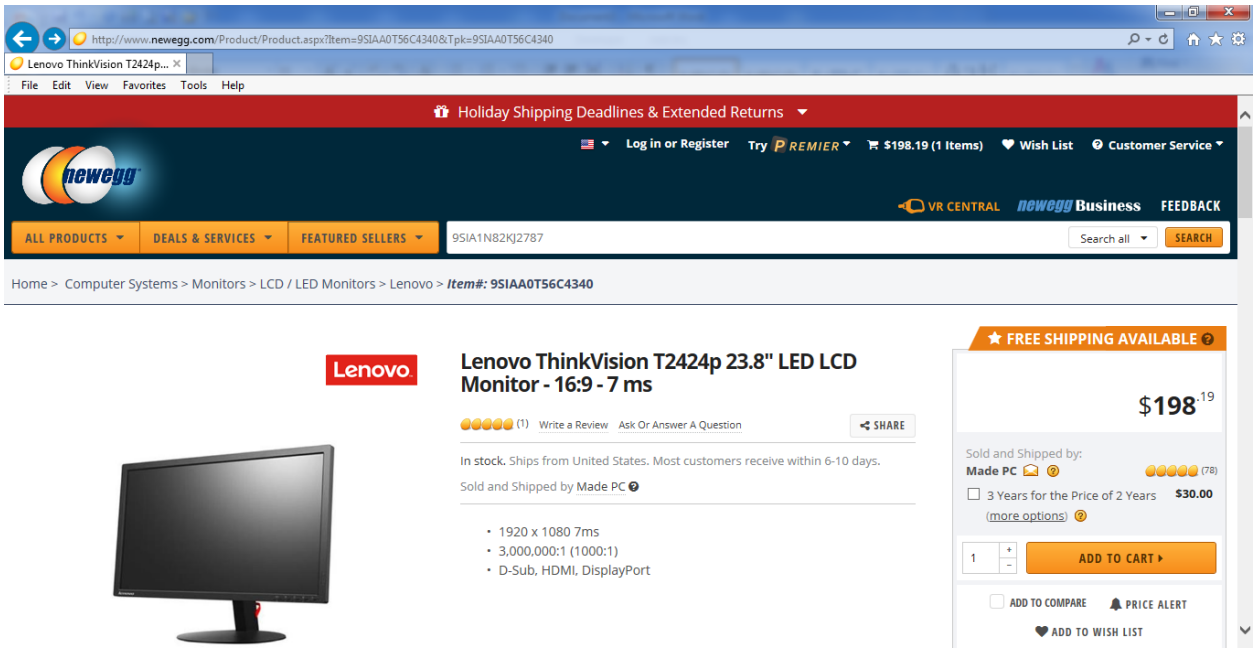

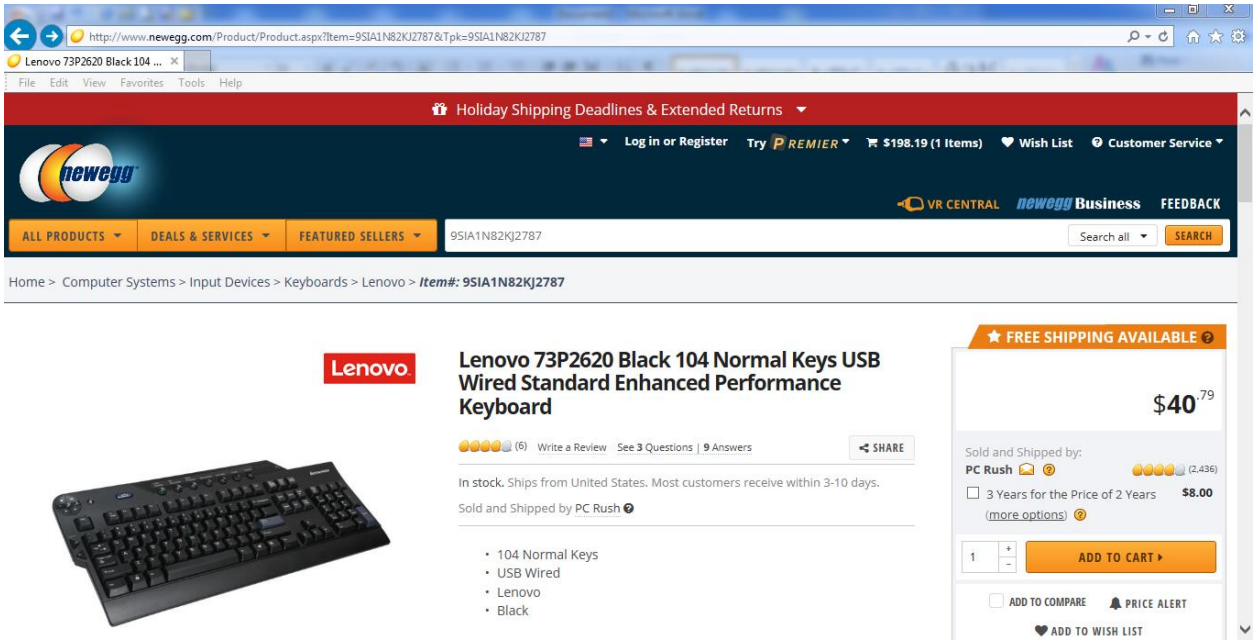

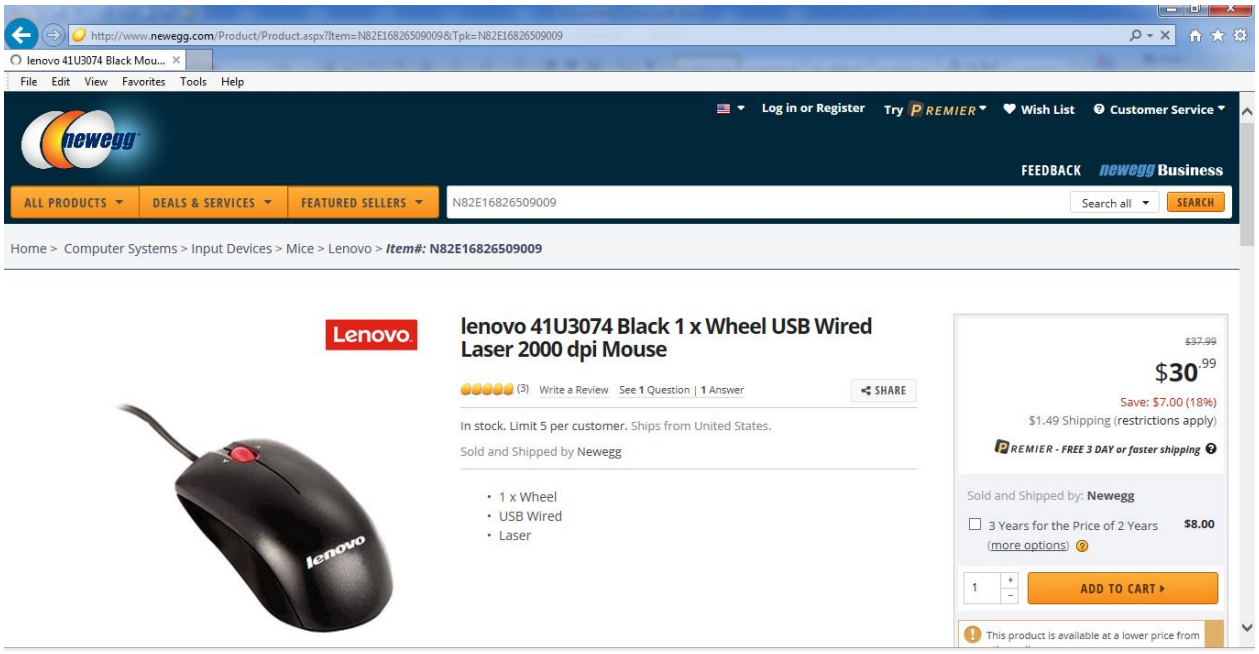
http://www.newegg.com/Product/Product.aspx?Item=N82E16826509009&Tpk=N82E16826509009
lenovo 41U3074 Black Mou...
File Edit View Favorites Tools Help
Log in or Register
Try PREMIER
Wish List
Customer Service
FEEDBACK
newegg Business
ALL PRODUCTS
DEALS & SERVICES
FEATURED SELLERS
N82E16826509009
Search all
SEARCH
Home > Computer Systems > Input Devices > Mice > Lenovo > Item#: N82E16826509009
Lenovo
lenovo 41U3074 Black 1 x Wheel USB Wired Laser 2000 dpi Mouse
(3) Write a Review See 1 Question | 1 Answer
SHARE
In stock. Limit 5 per customer. Ships from United States.
Sold and Shipped by Newegg
• 1 x Wheel
• USB Wired
• Laser
$37.99
$30.99
Save: $7.00 (18%)
$1.49 Shipping (restrictions apply)
PREMIER - FREE 3 DAY or faster shipping
Sold and Shipped by: Newegg
3 Years for the Price of 2 Years
$8.00
(more options)
1
ADD TO CART
This product is available at a lower price from

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-1TB-BDW-12-2016\
Validation Run Files	Supporting-Files-1TB-BDW-12-2016\Validation-run-logs-20161213-070004-hive-sf1000
Performance Run Files	Supporting-Files-1TB-BDW-12-2016\logs-20161213-164733-hive-sf1000
Repeatability Run Files	upporting-Files-1TB-BDW-12-2016\Repeatability-run-logs-20161214-021136-hive-sf1000
Clause 3 - Workload Related Items	
Benchmark Generic Parameters	Supporting-Files-1TB-BDW-12-2016\logs-20161213-164733-hive-sf1000\bigBench-configs\conf\userSettings.conf
Query Parameters used in the benchmark execution Settings	Supporting-Files-1TB-BDW-12-2016\logs-20161213-164733-hive-sf1000\bigBench-configs\hive\conf\queryParameters.sql
Benchmark Global Framework Parameters Settings	Supporting-Files-1TB-BDW-12-2016\logs-20161213-164733-hive-sf1000\bigBench-configs\hive\conf\engineSettings.sql
Benchmark Global Framework Parameters Settings	Supporting-Files-1TB-BDW-12-2016\logs-20161213-164733-hive-sf1000\bigBench-configs\hive\conf\engineSettings.conf
Load Test script	Supporting-Files-1TB-BDW-12-2016\logs-20161213-164733-hive-sf1000\bigBench-configs\hive\population\hiveCreateLoad.sql
Queries specific optimization parameters settings	Supporting-Files-1TB-BDW-12-2016\logs-20161213-164733-hive-sf1000\bigBench-configs\hive\queries\q[01-30]\engineLocalSettings.conf
Queries specific optimization parameters settings	Supporting-Files-1TB-BDW-12-2016\logs-20161213-164733-hive-sf1000\bigBench-configs\hive\queries\q[01-30]\engineLocalSettings.sql
Clause 4 - SUT Related Items	
Data Redundancy report	Supporting-Files-1TB-BDW-12-2016\hdfs-data-redundancy-report.txt
Benchmark execution script	Supporting-Files-1TB-BDW-12-2016\run-all.sh
Hardware and Software Report from a representative node	Supporting-Files-1TB-BDW-12-2016\logs-20161213-164733-hive-sf1000\run-logs\envInfo-test4\envInfo.log
All Framework configuration files are included in the Supporting Files Archive	Supporting-Files-1TB-BDW-12-2016\logs-20161213-164733-hive-sf1000\bigBench-configs\hadoop
	Supporting-Files-1TB-BDW-12-2016\logs-20161213-164733-hive-sf1000\bigBench-configs\hive
	Supporting-Files-1TB-BDW-12-2016\logs-20161213-164733-hive-sf1000\bigBench-configs\spark
Clause 5 - Metric and Scale Factor Related Items	
Benchmark Performance Report	Supporting-Files-1TB-BDW-12-2016\logs-20161213-164733-hive-sf1000\run-logs\BigBenchResult.log
Validation Test Report	Supporting-Files-1TB-BDW-12-2016\Validation-run-logs-20161213-070004-hive-sf1000\run-logs\BigBenchResult.log