

Alibaba Cloud Computing Ltd.

TPC Express Benchmark™ Big Bench (TPCx-BB)

Full Disclosure Report

for

Alibaba Cloud MaxCompute

(with 100x Worker Nodes, 3x Master Nodes)

using

MaxCompute v3.31

and

Alibaba Group Enterprise Linux Server 7.2 (Paladin)

First Edition

September 18, 2019

Alibaba Cloud Computing Ltd. (Alibaba), 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.


Alibaba and the Alibaba Logo are trademarks of Alibaba Cloud Computing 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 Alibaba and any other company.

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


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

Copyright © 2019 Alibaba Cloud Computing 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.

		Alibaba Cloud MaxCompute		TPCx-BB Rev. v1.3.0 TPC-Pricing Rev. v2.4.0	
				Report Date: September 18, 2019	
Total System Cost		TPCx-BB Performance Metric		Price/Performance	
5,756,026 USD		25,641.21 BBQpm@100000		224.49 USD \$/BBQpm@100000	
Framework	Operating System	Other Software	Availability Date	Scale Factor	Streams
MaxCompute v3.31	Alibaba Group Enterprise Linux Server 7.2 (Paladin)	None	September 18, 2019	100000	4

System Configuration



Physical Storage/Scale Factor: 100.81		Scale Factor/Physical Memory: 2.50	
Servers: 100x Worker Node, 3x Master Node Total Processors/Cores/Threads: 206/4,896/9,792			
Server Configuration: Processors Memory Storage Controller Storage Device Network	Per Worker Node: 2x Intel® Xeon® Platinum 8163 @ 2.50GHz 384 GB Onboard SATA controller 1x 240GB SATA 6Gbps SSD 1x 3.84TB SATA 6Gbps SSD 12x 8TB SATA 6Gbps HDD Mellanox MT27710 ConnectX-4 Lx	Per Master Node: 2x Intel® Xeon® E5-2682 v4 @ 2.50GHz 512 GB Onboard SATA controller 1x 240GB SATA 6Gbps SSD 12x 2TB SATA 6Gbps SSD Mellanox MT27710 ConnectX-4 Lx	
Connectivity:		3x Network Switch (8x100Gbps Up; 48x25Gbps Down)	



Alibaba Cloud MaxCompute

TPCx-BB Rev. v1.3.0
 TPC-Pricing Rev. v2.4.0

Report Date:
 September 18, 2019

Description	Part Number	Source	Unit Price	Qty	Ext. Price	3-Year Maint.
License Compute and Software Services						
<u>MaxCompute Annual Subscription (9,500 CU)</u>	Asia Pacific SE 1 (Singapore)	1	\$2,508,000.00	3	\$7,524,000.00	
Worker Node				100		
Intel(R) Xeon(R) Platinum 8163 CPU @ 2.50GHz				2		
32 GB Memory				12		
Master Node				3		
Intel(R) Xeon(R) CPU E5-2682 v4 @ 2.50GHz				2		
32 GB Memory				16		
Network Switches (8x100Gbps Up; 48x25Gbps Down)				3		
1-Year Annual Subscription Discount (30%)			-\$752,400.00	3	-\$2,257,200.00	
<u>MaxCompute Storage for 1 year</u>		1	\$9,424.30	3	\$28,272.90	
100000 Scale Factor (20.89 TB compressed)						
<u>MaxCompute Enterprise Service for 1 year</u>		1	\$152,351.70	3		\$457,055.11
24x7, 4 hour response						
License Compute and Software Services Sub-Total					\$5,295,072.90	\$457,055.11
Other Components						
13-inch MacBook Pro 1.4GHz (includes 2 spares)		2	\$1,299.00	3	\$3,897.00	
Other Components Sub-Total					\$3,897.00	\$0.00

Pricing: 1 = Alibaba; 2 = Apple.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	\$5,756,026
	BBQpm@100000	25,641.21
	\$/BBQpm@100000	\$ 224.49

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	100000
Streams	4
SUT Validation Test	PASS

Performance Run (Run 2)

Overall Run Start Time	2019-08-23 23:40:09.740
Overall Run End Time	2019-08-24 13:29:17.630
Overall Run Elapsed Time	49,747.890
Load Test Start Time	2019-08-23 23:40:09.741
Load Test End Time	2019-08-24 00:26:31.482
Load Test Elapsed Time	2,781.741
Power Test Start Time	2019-08-24 00:26:31.483
Power Test End Time	2019-08-24 06:01:27.639
Power Test Elapsed Time	20,096.156
Throughput Test Start Time	2019-08-24 06:01:27.640
Throughput Test End Time	2019-08-24 13:29:17.630
Throughput Test Elapsed Time	26,869.990
Performance Metric (BBQpm@ 100000)	25,641.21

Repeatability Run (Run 1)

Overall Run Start Time	2019-08-23 06:35:38.228
Overall Run End Time	2019-08-23 20:16:48.806
Overall Run Elapsed Time	49,270.578
Load Test Start Time	2019-08-23 06:35:38.229
Load Test End Time	2019-08-23 07:20:05.721
Load Test Elapsed Time	2,667.492
Power Test Start Time	2019-08-23 07:20:05.723
Power Test End Time	2019-08-23 12:54:54.511
Power Test Elapsed Time	20,088.788
Throughput Test Start Time	2019-08-23 12:54:54.511
Throughput Test End Time	2019-08-23 20:16:48.806
Throughput Test Elapsed Time	26,514.295
Performance Metric (BBQpm@ 100000)	25,990.73

Performance Run Report (Run 2)

```

*****
TPCx-BB
Result
v1.3.0
*****
INFO: T_LOAD = 2781.741
INFO: T_LD = 0.1 * T_LOAD: 278.1741
INFO: T_PT = 6766.13889008518
INFO: T_T_PUT = 26869.99
INFO: T_TT = 6717.4975
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@100000 = 25641.214015581
    
```

Repeatability Run Report (Run 1)

```

*****
TPCx-BB
Result
v1.3.0
*****
INFO: T_LOAD = 2667.492
INFO: T_LD = 0.1 * T_LOAD: 266.7492
INFO: T_PT = 6689.15536971355
INFO: T_T_PUT = 26514.295
INFO: T_TT = 6628.57375
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@100000 = 25990.7348140493
    
```

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

Table of Contents

ABSTRACT	8
PREFACE	9
CLAUSE 1: GENERAL ITEMS.....	10
1.1 TEST SPONSOR.....	10
1.2 PARAMETER SETTINGS	10
1.3 CONFIGURATION DIAGRAMS.....	10
CLAUSE 2: SOFTWARE COMPONENTS AND DATASET DISTRIBUTION	12
2.1 ROLES AND DATASET DISTRIBUTION.....	12
2.2 DISTRIBUTED FILE SYSTEM IMPLEMENTATION	13
2.3 ENGINE IMPLEMENTATION	13
2.4 FRAMEWORKS	13
2.5 APPLIED PATCHES	13
CLAUSE 3: WORKLOAD RELATED ITEMS	14
3.1 HARDWARE & SOFTWARE TUNABLE	14
3.2 KIT VERSION	14
3.3 RUN REPORT.....	14
3.4 QUERY ELAPSED TIMES	15
3.5 VALIDATION TEST OUTPUT.....	16
3.6 GLOBAL FRAMEWORK PARAMETERS.....	16
3.7 KIT MODIFICATIONS.....	17
CLAUSE 4: SUT RELATED ITEMS.....	18
4.1 SPECIALIZED HARDWARE/SOFTWARE	18
4.2 FRAMEWORK CONFIGURATION FILES	18
4.3 SUT ENVIRONMENT INFORMATION	18
4.4 DATA STORAGE TO SCALE FACTOR RATIO.....	18
4.5 SCALE FACTOR TO MEMORY RATIO	18
CLAUSE 5: METRICS AND SCALE FACTORS.....	19
5.1 PERFORMANCE RUN METRIC	19
5.2 REPEATABILITY RUN METRIC	19
5.3 PRICE-PERFORMANCE METRIC	19
5.4 SCALE FACTOR.....	19
5.5 STREAM COUNT.....	19
5.6 ELAPSED RUN TIMES.....	20
5.7 ELAPSED TEST TIMES.....	20
AUDITORS' INFORMATION AND ATTESTATION LETTER.....	21
PROVISIONING COMPUTE SERVICES.....	24
THIRD PARTY PRICE QUOTES.....	25
APPLE.....	25
SUPPORTING FILE INDEX.....	26

Abstract

This document contains the methodology and results of the TPC Express Benchmark™ Big Bench (TPCx-BB) test conducted in conformance with the requirements of the TPCx-BB Standard Specification, Revision v1.3.0.

The test was conducted at a Scale Factor of 100000 with 100 nodes (Worker Nodes) running MaxCompute v3.31 on Alibaba Group Enterprise Linux Server 7.2 (Paladin).

Measured Configuration

Company Name	Cluster Node	Virtualization	Operating System
Alibaba Cloud Computing Ltd.	100x Worker Nodes 3x Master Nodes	n/a	Alibaba Group Enterprise Linux Server 7.2 (Paladin)

TPC Express Benchmark© Big Bench Metrics

Total System Cost	BBQpm@100000	Price/Performance	Availability Date
5,756,026 USD	25,641.21	224.49 USD	September 18, 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 Alibaba Cloud Computing 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.

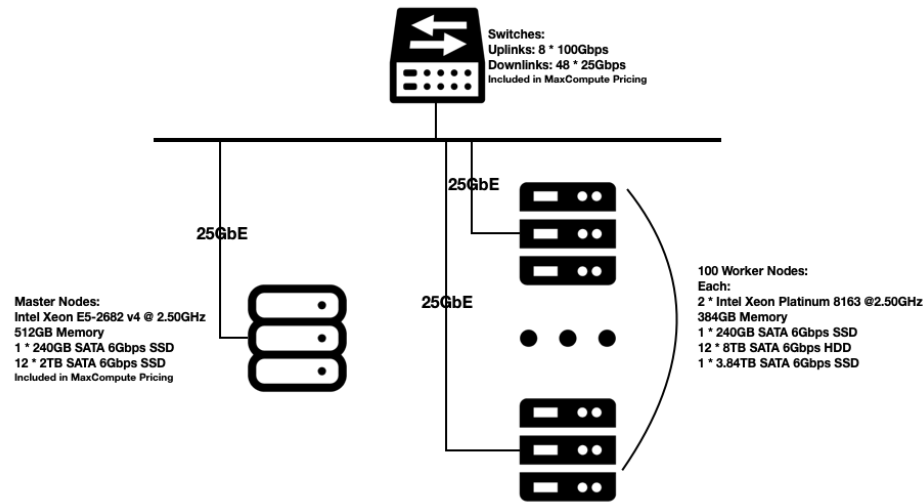
The Supporting Files Archive contains the parameters and options used to configure the components involved in this benchmark.

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: 103
- Total Processors/Cores/Threads: 206/4,896/9,792
- Total Memory: 39,936
- Total Number of Storage Drives/Devices: 1,439
- Total Storage Capacity: 10,080,720
- Network: 3x Network Switch (8x100Gbps Up; 48x25Gbps Down)

100x Worker Node, each with:

- Processors/Cores/Threads: 2/48/96
- Processor Model: 2x Intel® Xeon® Platinum 8163 @ 2.50GHz
- Memory: 384 GB
- Controller: Onboard SATA controller
- Drives:
 - 1x 240GB SATA 6Gbps SSD
 - 1x 3.84TB SATA 6Gbps SSD
 - 12x 8TB SATA 6Gbps HDD
- Network: Mellanox MT27710 ConnectX-4 Lx

3x Master Node, each with:

- Processors/Cores/Threads:
- Processor Model: 2x Intel® Xeon® E5-2682 v4 @ 2.50GHz
- Memory: 512 GB
- Controller: Onboard SATA controller
- Drives:
 - 1x 240GB SATA 6Gbps SSD
 - 12x 2TB SATA 6Gbps SSD
- Network: Mellanox MT27710 ConnectX-4 Lx

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
Maxcompute Master Node1	Pangu master Fuxi Master Nuwa	1	N	g90a111 78.cloud. et93	Intel Xeon E5-2682 v4 @ 2.50GHz 512GB Memory 1 * 240GB SATA 6Gbps SSD 12 * 2TB SATA 6Gbps SSD 2 * 25Gb OS: Linux 3.10.0-327.ali2008.odps.ali7.x86_64 Apsara: 1.11_u32	OS: 1 * 240GB SATA 6Gbps SSD Data Drive: 12 * 2TB SATA 6Gbps SSD
Maxcompute Master Node2	Pangu master Fuxi Master Nuwa	1	N	g90a073 74.cloud. et93	Intel Xeon E5-2682 v4 @ 2.50GHz 512GB Memory 1 * 240GB SATA 6Gbps SSD 12 * 2TB SATA 6Gbps SSD 2 * 25Gb OS: Linux 3.10.0-327.ali2008.odps.ali7.x86_64 Apsara: 1.11_u32	OS: 1 * 240GB SATA 6Gbps SSD Data Drive: 12 * 2TB SATA 6Gbps SSD
Maxcompute Master Node3	Pangu master Fuxi Master Nuwa	1	N	g90a111 77.cloud. et93	Intel Xeon E5-2682 v4 @ 2.50GHz 512GB Memory 1 * 240GB SATA 6Gbps SSD 12 * 2TB SATA 6Gbps SSD 2 * 25Gb OS: Linux 3.10.0-327.ali2008.odps.ali7.x86_64 Apsara: 1.11_u32	OS: 1 * 240GB SATA 6Gbps SSD Data Drive: 12 * 2TB SATA 6Gbps SSD
Maxcompute Worker Nodes	Fuxi tubo Pangu ChunkServer	100	N	See nodelist.txt in Suportin g Files	2 * Intel Xeon Platinum 8163 @2.50GHz 384GB Memory 1 * 240GB SATA 6Gbps SSD 12 * 8TB SATA 6Gbps HDD 1 * 3.84TB SATA 6Gbps SSD 2 * 25Gb OS: Linux 3.10.0-327.ali2008.odps.ali7.x86_64 Apsara: 1.11_u32	OS: 1 * 240GB SATA 6Gbps SSD Data Drive: 12 * 8TB SATA 6Gbps HDD 1 * 3.84TB SATA 6Gbps SSD

2.2 Distributed File System Implementation

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

MaxCompute v3.31.

2.3 Engine Implementation

The Engine implementation and corresponding version must be disclosed.

Component	Version
MaxCompute	3.31

2.4 Frameworks

Frameworks and Engine used in the benchmark should be disclosed.

Framework	Version
MaxCompute	3.31

2.5 Applied Patches

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

No additional patches were applied.

Clause 3: Workload Related Items

3.1 Hardware & Software Tunable

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

The Supporting Files Archive contains all configuration scripts.

3.2 Kit Version

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

TPCx-BB Kit Version

v1.3.0

3.3 Run Report

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

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

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

```
*****
TPCx-BB
Result
v1.3.0
*****
INFO: T_LOAD = 2667.492
INFO: T_LD = 0.1 * T_LOAD: 266.7492
INFO: T_PT = 6689.15536971355
INFO: T_T_PUT = 26514.295
INFO: T_TT = 6628.57375
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@100000 = 25990.7348140493
```

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

```
*****
TPCx-BB
Result
v1.3.0
*****
INFO: T_LOAD = 2781.741
INFO: T_LD = 0.1 * T_LOAD: 278.1741
INFO: T_PT = 6766.13889008518
INFO: T_T_PUT = 26869.99
INFO: T_TT = 6717.4975
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@100000 = 25641.214015581
```

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
Structured	1	52.009	69.030	64.869	48.211	49.755
	6	124.206	99.838	123.810	97.554	115.050
	7	77.475	82.040	84.934	121.572	77.382
	9	58.188	56.133	148.227	57.934	139.770
	11	45.749	45.082	106.369	46.142	58.387
	13	69.392	75.862	199.234	202.796	161.063
	14	55.309	353.383	62.701	110.507	48.624
	15	65.080	82.249	132.887	137.687	54.909
	16	176.209	178.632	613.550	295.362	353.431
	17	77.584	81.868	73.528	215.323	145.319
	20	262.585	368.079	338.420	261.155	256.475
	21	613.665	743.649	603.556	613.630	608.555
	22	37.830	43.633	92.907	39.258	59.481
	23	96.141	148.703	307.633	133.000	95.632
24	92.819	99.061	93.159	170.838	102.406	
25	287.193	454.526	294.946	326.505	290.813	
26	246.331	246.035	302.392	394.143	480.159	
29	136.410	220.919	161.679	129.090	484.857	
Semi-structured	2	2,205.440	2,985.137	2,819.812	2,579.806	3,290.000
	3	813.678	1,613.763	973.592	1,693.587	1,073.823
	4	1,991.261	2,974.876	1,890.079	3,149.863	2,930.048
	5	415.217	960.895	417.338	726.166	646.070
	8	478.231	690.487	949.035	525.207	1,010.237
	12	2,168.683	2,198.673	2,223.540	2,678.442	2,498.274
	30	2,320.643	3,375.661	2,620.401	3,045.896	1,910.259
Unstructured	10	119.650	218.088	278.863	117.550	228.046
	18	2,073.212	2,343.803	2,970.631	3,501.246	3,039.884
	19	4,795.368	5,079.337	4,519.712	5,165.139	4,854.389
	27	37.941	95.676	39.834	98.733	33.069
	28	102.621	247.682	97.510	187.634	138.747

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/bigBench.properties
- bigBench-configs/conf/userSettings.conf
- bigBench-configs/sql/conf/engineSettings.conf
- bigBench-configs/sql/conf/engineSettings.sql
- bigBench-configs/sql/queries/q01/engineLocalSettings.sql
- bigBench-configs/sql/queries/q02/engineLocalSettings.sql
- bigBench-configs/sql/queries/q03/engineLocalSettings.sql
- bigBench-configs/sql/queries/q04/engineLocalSettings.sql
- bigBench-configs/sql/queries/q05/engineLocalSettings.sql
- bigBench-configs/sql/queries/q06/engineLocalSettings.sql
- bigBench-configs/sql/queries/q07/engineLocalSettings.sql
- bigBench-configs/sql/queries/q08/engineLocalSettings.sql
- bigBench-configs/sql/queries/q09/engineLocalSettings.sql
- bigBench-configs/sql/queries/q10/engineLocalSettings.sql
- bigBench-configs/sql/queries/q11/engineLocalSettings.sql
- bigBench-configs/sql/queries/q12/engineLocalSettings.sql
- bigBench-configs/sql/queries/q13/engineLocalSettings.sql
- bigBench-configs/sql/queries/q14/engineLocalSettings.sql
- bigBench-configs/sql/queries/q15/engineLocalSettings.sql
- bigBench-configs/sql/queries/q16/engineLocalSettings.sql
- bigBench-configs/sql/queries/q17/engineLocalSettings.sql
- bigBench-configs/sql/queries/q18/engineLocalSettings.sql
- bigBench-configs/sql/queries/q19/engineLocalSettings.sql
- bigBench-configs/sql/queries/q20/engineLocalSettings.sql
- bigBench-configs/sql/queries/q21/engineLocalSettings.sql
- bigBench-configs/sql/queries/q22/engineLocalSettings.sql
- bigBench-configs/sql/queries/q23/engineLocalSettings.sql
- bigBench-configs/sql/queries/q24/engineLocalSettings.sql
- bigBench-configs/sql/queries/q25/engineLocalSettings.sql
- bigBench-configs/sql/queries/q26/engineLocalSettings.sql
- bigBench-configs/sql/queries/q27/engineLocalSettings.sql
- bigBench-configs/sql/queries/q28/engineLocalSettings.sql
- bigBench-configs/sql/queries/q29/engineLocalSettings.sql
- bigBench-configs/sql/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 included in the Supporting Files Archive.

4.4 Data Storage to Scale Factor Ratio

The data storage ratio must be disclosed.

Nodes	Disks	Size (GB)	Total (GB)
103	1	240	24,720
100	12	8,000	9,600,000
100	1	3,840	384,000
3	12	2,000	72,000

Total Storage (GB)	10,080,720
Scale Factor	100000
Data Storage Ratio	100.81

4.5 Scale Factor to Memory Ratio

The Scale Factor to memory ratio must be disclosed.

Nodes	Memory (GiB)	Total (GiB)
100	384	38,400
3	512	1,536

Scale Factor	100000
Total Memory (GiB)	39,936
SF / Memory Ratio	2.50

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@100000 25,641.21

5.2 Repeatability Run Metric

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

Repeatability Run
BBQpm@100000 25,990.73

5.3 Price-Performance Metric

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

Price / Performance
\$BBQpm@100000 224.49

5.4 Scale Factor

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

Scale Factor
100000

5.5 Stream Count

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

Streams
4

5.6 Elapsed Run Times

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

Run	Elapsed Time	Seconds
Run 1	00 13:41:10.578	49,270.578
Run 2	00 13:49:07.890	49,747.890

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	2,781.741	2,667.492
Power Test	20,096.156	20,088.788
Throughput Test	26,869.990	26,514.295

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.

Xiening Dai
Senior Staff Engineer
500 108th Ave NE, Suite 800
Bellevue, WA 98004

September 13, 2019

I verified the TPC Express Benchmark™ BB v1.3.0 performance of the following configuration:

Platform: Alibaba Cloud Cluster
(w/ 100x Worker Nodes, 3x Master Nodes)
Operating System: Alibaba Group Enterprise Linux Server 7.2 (Paladin)
Framework: MaxCompute v3.31

The results were:

Performance Metric **25,641.21 BBQpm@100000**
Run Elapsed Time 00 13:49:07.890 (49,747.890 Seconds)

Cluster **100x Worker Nodes, 3x Master Nodes**

CPU	2x Intel® Xeon® Platinum 8163 (2.50 GHz, 24-core, 33 MB L3) (Worker Nodes)		
	2x Intel® Xeon® E5-2682 v4 (2.50 GHz, 16-core, 40 MB L3) (Master Nodes)		
Memory	384GiB (Worker nodes), 512GiB (Master nodes)		
Storage	Qty	Size	Type
	1	240GB	6G SATA SSD (All nodes)
	1	3.84TB	6G SATA SSD (Worker Nodes)
	12	8TB	6G SATA HDD (Worker Nodes)
	12	2TB	6G SATA SSD (Master Nodes)

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

The following verification items were given special attention:

- All TPC-provided components were verified to be v1.3.0
- 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 (100000GB) 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:

None.

Respectfully Yours,

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

Doug Johnson, TPC Auditor

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

Provisioning Compute Services

发件人: 韦海青(海清) <haiqing.whq@alibaba-inc.com>
发送时间: 2019年9月11日(星期三) 19:55
收件人: 戴谢宁 <xiening.dai@alibaba-inc.com>; 路璐(杜撰) <lu.lu@alibaba-inc.com>
主 题: tpcxbb Maxcompute cu price quote

To Xiening Dai,

Here is the information you requested regarding pricing for Alibaba Cloud Maxcompute.
All pricing shown is in US Dollars(\$) and can be pre-paid.

Product Name	Region	Resource Quota(cu)	Unit Price(\$)	1 Year Price(\$)	discount	1 Year Price after discount (\$)
Alibaba Cloud MaxCompute CU	Asia Pacific SE 1 (Singapore)	9,500	22 \$/cu/month	2,508,000.0	752,400.0	1,755,600.0

- $2,508,000 = 9,500 \text{ (cu)} * 22 \text{ ($/cu/month)} * 12 \text{ (month)}$

Product Name	Region	Size(TB)	Unit Price(\$)	1 Year Price (\$)
Alibaba Cloud MaxCompute Storage	Asia Pacific SE 1 (Singapore)	20.89	25.82 \$/day	9424.30

- $25.82 \text{ ($/day)} = (100 - 1) * 0.0028 \quad // 1\text{GB} - 100\text{GB}$
+ $(1,024 - 100) * 0.0014 \quad // 100\text{GB} - 1\text{TB}$
+ $(10,240 - 1,024) * 0.0013 \quad // 1\text{TB} - 10\text{TB}$
+ $(20.89 * 1,024 - 10,240) * 0.0011 \quad // 10\text{TB} - 20.89\text{TB}$

If you have any questions about this price quote, please contact our sales via the link: https://common-buy-intl.aliyun.com/?commodityCode=odpsplus_intl#/buy

Haiqing Wei (Alibaba Colud Product Manager)
(haiqing.whq@alibaba-inc.com)

Third Party Price Quotes

Apple

MacBook Pro

[Overview](#) [macOS](#) [Tech Specs](#)

Add a trade-in

Get a refund of up to \$1400 when you trade in an eligible computer, or recycle it for free.*

[Get started](#)



Customize your 13-inch MacBook Pro - Space Gray

1.4GHz quad-core 8th-generation Intel Core i5 processor, Turbo Boost up to 3.9GHz

8GB 2133MHz LPDDR3 memory

128GB SSD storage

Retina display with True Tone

Intel Iris Plus Graphics 645

Two Thunderbolt 3 ports

Touch Bar and Touch ID

Backlit Keyboard - US English



Pickup:
Today at Apple Infinite Loop
[Check another store](#)



Order by 5pm, delivers:
Tomorrow – Fastest
Fri, Aug 30 – Free
[Delivery options for 01453**](#)

\$1,299.00

[Add to Bag](#)



[Get 3% Daily Cash with Apple Card, or get special financing >](#)

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	Support-Files-for-Alibaba-Maxcompute-100nodes-100TB
Validation Run Files	Support-Files-for-Alibaba-Maxcompute-100nodes-100TB/Validation-Run-logs-20190823-042321-sql-sf100000
Performance Run Files	Support-Files-for-Alibaba-Maxcompute-100nodes-100TB/Performance-Run-logs-20190823-204149-sql-sf100000
Repeatability Run Files	Support-Files-for-Alibaba-Maxcompute-100nodes-100TB/Repeatability-Run-logs-20190824-135332-sql-sf100000
Clause 3 - Workload Related Items	
Benchmark Generic Parameters	Support-Files-for-Alibaba-Maxcompute-100nodes-100TB/Performance-Run-logs-20190823-204149-sql-sf100000/bigBench-configs/conf/userSettings.conf
Query Parameters used in the benchmark execution Settings	Support-Files-for-Alibaba-Maxcompute-100nodes-100TB/Performance-Run-logs-20190823-204149-sql-sf100000/bigBench-configs/sql/conf/queryParameters.sql
Benchmark Global Framework Parameters Settings	Support-Files-for-Alibaba-Maxcompute-100nodes-100TB/Performance-Run-logs-20190823-204149-sql-sf100000/bigBench-configs/sql/conf/engineSettings.sql
Benchmark Global Framework Parameters Settings	Support-Files-for-Alibaba-Maxcompute-100nodes-100TB/Performance-Run-logs-20190823-204149-sql-sf100000/bigBench-configs/sql/conf/engineSettings.conf
Load Test script	Support-Files-for-Alibaba-Maxcompute-100nodes-100TB/Performance-Run-logs-20190823-204149-sql-sf100000/bigBench-configs/sql/population/odpsCreateLoad.sql
Queries specific optimization parameters settings	Support-Files-for-Alibaba-Maxcompute-100nodes-100TB/Performance-Run-logs-20190823-204149-sql-sf100000/bigBench-configs/sql/queries/q[01-30]/engineLocalSettings.conf
Queries specific optimization parameters settings	Support-Files-for-Alibaba-Maxcompute-100nodes-100TB/Performance-Run-logs-20190823-204149-sql-sf100000/bigBench-configs/sql/queries/q[01-30]/engineLocalSettings.sql
Clause 4 - SUT Related Items	
Data Redundancy report	Support-Files-for-Alibaba-Maxcompute-100nodes-100TB/pangu_redundant_info_20190823-213933.txt
Benchmark execution script	Support-Files-for-Alibaba-Maxcompute-100nodes-100TB/publication-runs.sh
Hardware and Software Report from a representative node	Support-Files-for-Alibaba-Maxcompute-100nodes-100TB/envInfo-g77k07229.cloud.et93/envInfo.log
Clause 5 - Metric and Scale Factor Related Items	
Benchmark Performance Report	Support-Files-for-Alibaba-Maxcompute-100nodes-100TB/Performance-Run-logs-20190823-204149-sql-sf100000/run-logs/BigBenchResult.log
Validation Test Report	Support-Files-for-Alibaba-Maxcompute-100nodes-100TB/Validation-Run-logs-20190823-042321-sql-sf100000/run-logs/BigBenchResult.log
Extra – worker node name list	Support-Files-for-Alibaba-Maxcompute-100nodes-100TB/nodelist.txt