

Alibaba Cloud Computing Ltd.

TPC Express Benchmark™ Big Bench (TPCx-BB)

Full Disclosure Report

for

Alibaba Cloud MaxCompute

(with 47x Compute Nodes;

9x Storage Node

3x Master Nodes)

using

MaxCompute v3.42

and

Alibaba Group Enterprise Linux Server 7.2 (Paladin)

First Edition

August 18, 2022

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.

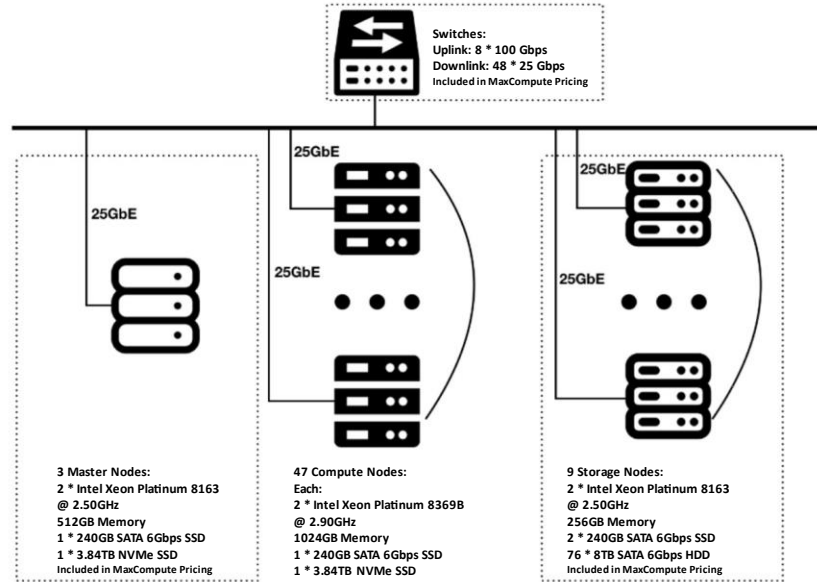
The Alibaba products, services or features identified in this document may not yet be available or may not be available in all areas and may be subject to change without notice. Consult your local Alibaba business contact for information on the products or services available in your area. You can find 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 © 2022 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. v1.5.2 TPC-Pricing Rev. v2.8.0	
				Report Date: August 18, 2022	
Total System Cost		TPCx-BB Performance Metric		Price/Performance	
3,664,590 USD		64,580.63 BBQpm@100000		56.75 USD \$/BBQpm@100000	
Framework	Operating System	Other Software	Availability Date	Scale Factor	Streams
MaxCompute v3.42	Alibaba Group Enterprise Linux Server 7.2 (Paladin)	None	August 18, 2022	100000	4

System Configuration



Physical Storage/Scale Factor: 56.80		Scale Factor/Physical Memory: 1.92	
Servers:		59x Master Node / Compute Node / Storage Node	
Total Processors/Cores/Threads		118/3,584/7,168	
3x Master Nodes: 2x Intel® Xeon® Platinum 8163 CPU @ 2.50GHz 512 GiB Onboard SATA Controller 1x 240 GB SATA 6 Gbps SSD 1x 3.84 TB NVMe SSD Mellanox MT27710 2-port	47x Compute Nodes: 2x Intel(R) Xeon(R) Platinum 8369B CPU @ 2.90GHz 1,024 GiB Onboard SATA Controller 1x 240 GB SATA 6 Gbps SSD 1x 3.84 TB NVMe SSD Mellanox MT28800 2-port	9x Storage Nodes: 2x Intel(R) Xeon(R) Platinum 8163 CPU @ 2.50GHz 256 GiB Onboard SATA Controller 2x 240 GB SATA 6 Gbps SSD 76x 8 TB SATA 6 Gbps HDD Mellanox MT27710 2-port	
Connectivity:		Network Switch (8x 100 Gbps Up; 48x 25 Gbps Down)	



Alibaba Cloud MaxCompute

TPCx-BB Rev. v1.5.2
TPC-Pricing Rev. v2.8.0

Report Date:
August 18, 2022

Description	Part Number	Source	Unit Price	Qty	Ext. Price	3-Year Maint.
License Compute and Software Services						
<u>MaxCompute Annual Subscription (6,000 CU)</u>	Asia Pacific SE 1 (Singapore)	1	\$1,584,000.00	3	\$4,752,000.00	
Master Node				3		
Intel® Xeon® Platinum 8163 @ 2.50 GHz				2		
32 GB Memory				16		
240 GB SATA 6 Gbps SSD				1		
3.84 TB NVMe SSD				1		
Compute Node				47		
Intel® Xeon® Platinum 8369B @ 2.90 GHz				2		
64 GB Memory				16		
240 GB SATA 6 Gbps SSD				1		
3.84 TB NVMe SSD				1		
Storage Node				9		
Intel® Xeon® Platinum 8163 @ 2.50 GHz				2		
32 GB Memory				8		
240 GB SATA 6 Gbps SSD				2		
8 TB SATA 6 Gbps HDD				76		
Network Switches (8x100Gbps Up; 48x25Gbps Down)				NA		
1-Year Annual Subscription Discount (30%)			-\$475,200.00	3	-\$1,425,600.00	
<u>MaxCompute Storage for 1 year</u>		1	\$4,686.74	3	\$14,060.22	
100000 Scale Factor (20.90 TB compressed)						
<u>MaxCompute Enterprise Service for 1 year</u>		1	\$106,744.08	3		\$320,232.24
24x7, 4 hour response						
License Compute and Software Services Sub-Total					\$3,340,460.22	\$320,232.24
Other Components						
13-inch MacBook Pro M1 Chip (includes 2 spares)		2	\$1,299.00	3	\$3,897.00	
Other Components Sub-Total					\$3,897.00	\$0.00
3-Year Cost of Ownership					\$3,664,590	

Pricing Source: 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 \$3,664,590

BBQpm@100000 64,580.63

\$/BBQpm@100000 \$ 56.75

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 1)

Overall Run Start Time	2022-07-31 00:14:59.316
Overall Run End Time	2022-07-31 05:22:41.848
Overall Run Elapsed Time	18,462.532
Load Test Start Time	2022-07-31 00:14:59.317
Load Test End Time	2022-07-31 00:44:08.199
Load Test Elapsed Time	1,748.882
Power Test Start Time	2022-07-31 00:44:08.200
Power Test End Time	2022-07-31 02:08:59.126
Power Test Elapsed Time	5,090.926
Throughput Test Start Time	2022-07-31 02:08:59.127
Throughput Test End Time	2022-07-31 05:22:41.848
Throughput Test Elapsed Time	11,622.721
Performance Metric (BBQpm@100000)	64,580.63

Repeatability Run (Run 2)

Overall Run Start Time	2022-07-31 08:57:59.970
Overall Run End Time	2022-07-31 13:49:10.932
Overall Run Elapsed Time	17,470.962
Load Test Start Time	2022-07-31 08:57:59.970
Load Test End Time	2022-07-31 09:27:31.462
Load Test Elapsed Time	1,771.492
Power Test Start Time	2022-07-31 09:27:31.464
Power Test End Time	2022-07-31 10:55:14.324
Power Test Elapsed Time	5,262.860
Throughput Test Start Time	2022-07-31 10:55:14.325
Throughput Test End Time	2022-07-31 13:49:10.932
Throughput Test Elapsed Time	10,436.607
Performance Metric (BBQpm@100000)	67,983.71

Performance Run Report (Run 1)

TPCx-BB

Result

v1.5.2

INFO: T_LOAD = 1748.882

INFO: T_LD = 0.1 * T_LOAD: 174.8882

INFO: T_PT = 2348.58710288479

INFO: T_T_PUT = 11622.721

INFO: T_TT = 2905.68025

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 = 64580.6300328222

Repeatability Run Report (Run 2)

TPCx-BB

Result

v1.5.2

INFO: T_LOAD = 1771.492

INFO: T_LD = 0.1 * T_LOAD: 177.1492

INFO: T_PT = 2339.29932051281

INFO: T_T_PUT = 10436.607

INFO: T_TT = 2609.15175

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 = 67983.7113235733

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	12
2.3 ENGINE IMPLEMENTATION	12
2.4 FRAMEWORKS	12
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.....	16
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
THIRD PARTY PRICE QUOTES.....	24
APPLE.COM	24
SUPPORTING FILE INDEX.....	25

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 v1.5.2.

The test was conducted at a Scale Factor of 100000 with 59 nodes (3x Master Nodes, 47x Compute Nodes, 9x Storage Nodes) running MaxCompute v3.42 on Alibaba Group Enterprise Linux Server 7.2 (Paladin).

Measured Configuration

Company Name	Cluster Node	Virtualization	Operating System
Alibaba Cloud Computing Ltd.	3x Master Nodes 47x Compute Node 9x Storage Node	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
3,664,590 USD	64,580.63	56.75 USD	August 18, 2022

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.

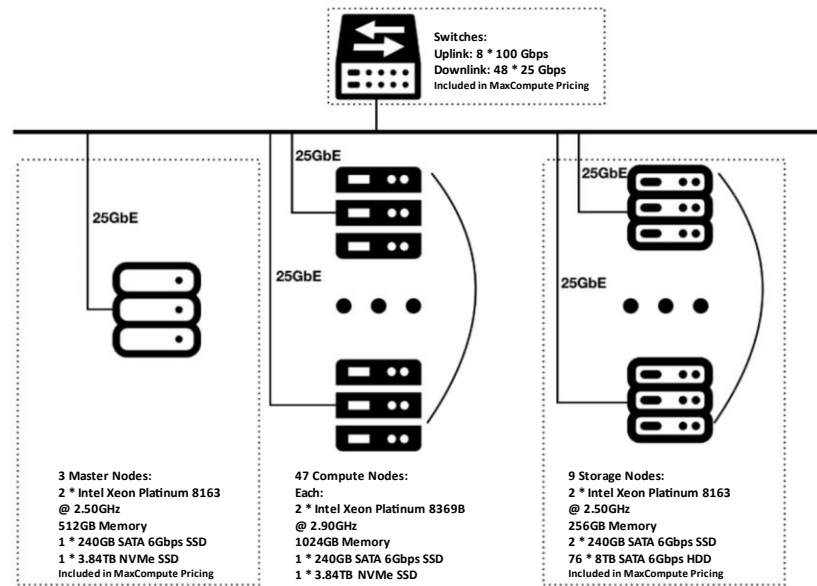
1.3 Configuration Diagrams

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

Alibaba Cloud



The measured configuration consisted of:

Total Nodes:	59
Total Processors/Cores/Threads:	118/3,584/7,168
Total Memory:	51,968
Total Number of Storage Devices:	802
Total Storage Capacity:	5,680,320

Network: Network Switch (8x 100 Gbps Up; 48x 25 Gbps Down)

3x Master Nodes:	47x Compute Nodes:	9x Storage Nodes:
2/48/96 (processors/cores/threads)	2/64/128	2/48/96
2x Intel(R) Xeon(R) Platinum 8163 CPU @ 2.50GHz	2x Intel(R) Xeon(R) Platinum 8369B CPU @ 2.90GHz	2x Intel(R) Xeon(R) Platinum 8163 CPU @ 2.50GHz
512 GiB	1,024 GiB	256 GiB
Onboard SATA Controller	Onboard SATA Controller	Onboard SATA Controller
1x 240 GB SATA 6 Gbps SSD	1x 240 GB SATA 6 Gbps SSD	2x 240 GB SATA 6 Gbps SSD
1x 3.84 TB NVMe SSD	1x 3.84 TB NVMe SSD	76x 8 TB SATA 6 Gbps HDD
Mellanox MT27710 2-port	Mellanox MT28800 2-port	Mellanox MT27710 2-port

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	Host Names	HW/SW Configuration	Storage Setup
MaxCompute Master Nodes	Pangu master Fuxi Master Nuwa	3	See Support-Files-for-Alibaba-Maxcompute-47nodes-100TB/master-nodelist.txt	2 * Intel Xeon Platinum 8163 @ 2.50GHz 512GB Memory 1 * 240GB SATA 6Gbps SSD 1 * 3.84TB NVMe SSD 2 * 25Gb OS: Linux 3.10.0-327.ali2014.alios7.x86_64 Apsara: 1.11_u32	OS: 1x 240GB SATA 6Gbps SSD Data Drive: 1x 3.84TB NVMe SSD
MaxCompute Compute Nodes	Fuxi tubo	47	See Support-Files-for-Alibaba-Maxcompute-47nodes-100TB/nodelist.txt	2 * Intel Xeon Platinum 8369B @ 2.90GHz 1,024GB Memory 1 * 240GB SATA 6Gbps SSD 1 * 3.84TB NVMe SSD 2 * 25Gb OS: Linux 3.10.0-327.ali2016.alios7.x86_64 Apsara: 1.11_u32	OS: 1x 240GB SATA 6Gbps SSD Data Drive: 1x 3.84TB NVMe SSD
MaxCompute Storage Nodes	Pangu ChunkServer	9	See Support-Files-for-Alibaba-Maxcompute-47nodes-100TB/storage-nodelist.txt	2 * Intel Xeon Platinum 8163 @ 2.50GHz 256GB Memory 2 * 240GB SATA 6Gbps SSD 76 * 8TB SATA 6Gbps HDD OS: Linux 3.10.0-327.ali2010.rc7.alios7.x86_64 Apsara: 1.11_u32	OS: 2x 240GB SATA 6Gbps SSD Data Drive: 76x 8TB SATA 6Gbps HDD

2.2 Distributed File System Implementation

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

MaxCompute v3.42.

2.3 Engine Implementation

The Engine implementation and corresponding version must be disclosed.

Component	Version
MaxCompute SQL Task	3.41

2.4 Frameworks

Frameworks and Engine used in the benchmark should be disclosed.

Framework	Version
MaxCompute	3.42

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

v1.5.2

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 (Performance Run)**

```
*****
TPCx-BB
Result
v1.5.2
*****
INFO: T_LOAD = 1748.882
INFO: T_LD = 0.1 * T_LOAD: 174.8882
INFO: T_PT = 2348.58710288479
INFO: T_T_PUT = 11622.721
INFO: T_TT = 2905.68025
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 = 64580.6300328222
```

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

```
*****
TPCx-BB
Result
v1.5.2
*****
INFO: T_LOAD = 1771.492
INFO: T_LD = 0.1 * T_LOAD: 177.1492
INFO: T_PT = 2339.29932051281
INFO: T_T_PUT = 10436.607
INFO: T_TT = 2609.15175
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 = 67983.7113235733
```

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	23.441	41.585	479.953	474.042	41.849
	6	32.727	118.490	72.263	315.612	58.374
	7	21.949	582.591	119.565	22.839	186.983
	9	41.738	140.731	464.454	39.766	49.140
	11	39.588	33.994	83.641	77.933	35.704
	13	26.814	227.774	26.848	28.896	25.647
	14	14.985	18.352	36.492	15.713	121.482
	15	19.323	43.071	18.758	20.697	23.329
	16	73.619	110.679	93.987	98.339	131.841
	17	23.616	34.895	29.534	397.724	22.807
	20	84.556	102.406	104.166	104.300	423.871
	21	278.987	411.876	671.235	623.380	409.867
	22	9.648	17.930	11.778	112.742	17.659
	23	28.132	86.560	29.869	31.797	453.671
	24	20.066	185.556	216.998	20.302	108.903
25	97.892	106.890	122.293	106.387	115.509	
26	93.534	109.214	290.907	136.676	113.469	
29	63.002	98.922	84.505	497.045	397.498	
Semi-structured	2	419.161	515.016	572.251	629.698	685.636
	3	289.091	367.840	454.722	462.237	677.717
	4	359.800	425.415	764.045	447.913	704.690
	5	449.392	963.474	677.531	1,021.952	758.593
	8	111.058	472.409	390.054	650.614	893.189
	12	34.379	371.234	172.935	174.714	40.316
	30	460.162	1,179.365	489.786	1,143.755	595.330
Unstructured	10	161.731	370.992	976.419	271.056	276.796
	18	1,279.306	1,664.111	3,488.115	1,789.935	2,550.765
	19	298.158	347.496	402.818	425.908	366.414
	27	63.919	1,250.512	99.974	173.815	327.467
	28	171.113	689.082	176.809	667.373	178.595

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/sql/conf/engineSettings.conf
- bigBench-configs/sql/conf/engineSettings.sql
- bigBench-configs/sql/conf/queryParameters.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)
50	1	240	12,000
50	1	3,840	192,000
9	2	240	4,320
9	76	8,000	5,472,000

Total Storage (GB)	5,680,320
Scale Factor	100000
Data Storage Ratio	56.80

4.5 Scale Factor to Memory Ratio

The Scale Factor to memory ratio must be disclosed.

Nodes	Memory (GiB)	Total (GiB)
3	512	1,536
47	1,024	48,128
9	256	2,304

Scale Factor	100000
Total Memory (GiB)	51,968
SF / Memory Ratio	1.92

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 64,580.63

5.2 Repeatability Run Metric

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

Repeatability Run
BBQpm@100000 67,983.71

5.3 Price-Performance Metric

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

Price / Performance
\$BBQpm@100000 56.75

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 05:07:42.532	18,462.532
Run 2	00 04:51:10.962	17,470.962

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	1,748.882	1,771.492
Power Test	5,090.926	5,262.860
Throughput Test	11,622.721	10,436.607

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.



Jing Sun
Alibaba Cloud Computing Co., Ltd.
No.12, Zhuantang
Science and Technology Economic Zone
Xihu District, Hangzhou City
Zhejiang Province China

August 17, 2022

I verified the TPC Express Benchmark™ BB v1.5.2 performance of the following configuration:

Platform: Alibaba Cloud MaxCompute
(w/ 3x Master Nodes, 47x Compute Nodes, 9 Storage Nodes)
Operating System: Alibaba Group Enterprise Linux Server 7.2 (Paladin)
Framework: MaxCompute v3.42

The results were:

Performance Metric 64,580.63 BBQpm@100000
Run Elapsed Time 00 05:07:43.532 (18,462.532 Seconds)

Cluster 3x Master Nodes, 47x Compute Nodes, 9x Storage Nodes

CPU	2x Intel® Xeon® Platinum 8163 (2.50 GHz, 24-core, 33 MB L3) (Master, Storage nodes)		
	2x Intel® Xeon® Platinum 8369B (2.90 GHz, 32-core, 48 MB L3) (Compute nodes)		
Memory	512GiB (Master nodes), 1,024GiB (Compute nodes), 256GiB (Storage nodes)		
Storage	Qty	Size	Type
	1	240 GB	SATA 6 Gbps SSD (Master, Compute nodes)
	2	240 GB	SATA 6 Gbps SSD (Storage nodes)
	1	3.84 TB	NVMe SSD (Master, Compute nodes)
	76	8 TB	SATA 6 Gbps HDD (Storage 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 v1.5.2
- 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,

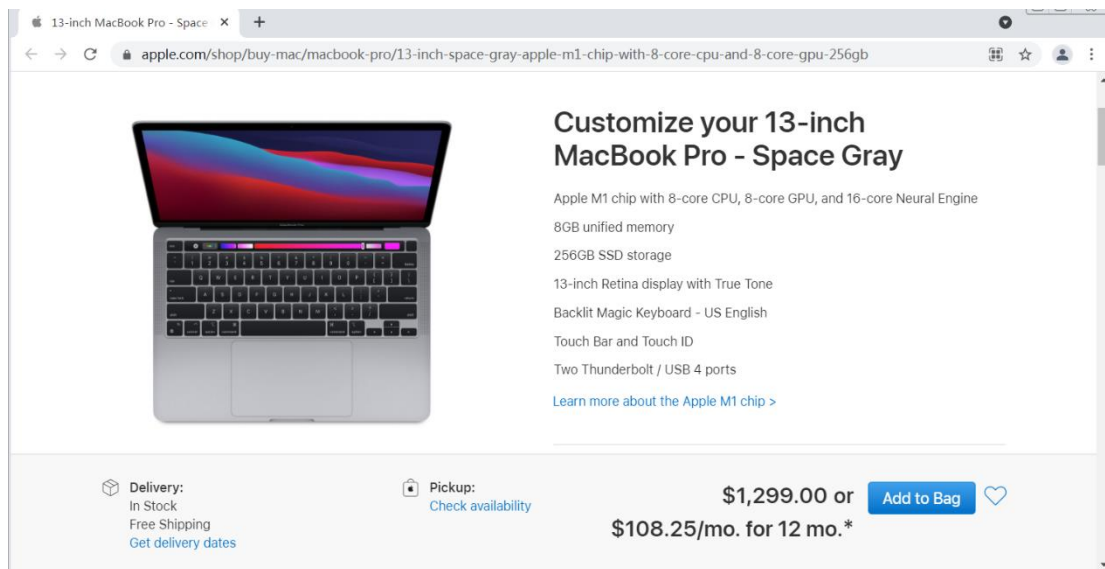
A handwritten signature in black ink that reads "Doug Johnson". The signature is written in a cursive style with a long horizontal flourish extending to the right.

Doug Johnson, TPC Auditor

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

Third Party Price Quotes

Apple.com



13-inch MacBook Pro - Space

apple.com/shop/buy-mac/macbook-pro/13-inch-space-gray-apple-m1-chip-with-8-core-cpu-and-8-core-gpu-256gb


Customize your 13-inch MacBook Pro - Space Gray

Apple M1 chip with 8-core CPU, 8-core GPU, and 16-core Neural Engine
8GB unified memory
256GB SSD storage
13-inch Retina display with True Tone
Backlit Magic Keyboard - US English
Touch Bar and Touch ID
Two Thunderbolt / USB 4 ports

[Learn more about the Apple M1 chip >](#)

Delivery:
In Stock
Free Shipping
[Get delivery dates](#)

Pickup:
[Check availability](#)

\$1,299.00 or [Add to Bag](#) 
\$108.25/mo. for 12 mo.*

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-47nodes-100TB
Validation Run Files	Support-Files-for-Alibaba-Maxcompute-47nodes-100TB/Validation-Run-logs-20220730-210451-sql-sf100000
Performance Run Files	Support-Files-for-Alibaba-Maxcompute-47nodes-100TB/Performance-Run-logs-20220731-054151-sql-sf100000
Repeatability Run Files	Support-Files-for-Alibaba-Maxcompute-47nodes-100TB/Repeatability-Run-logs-20220731-140656-sql-sf100000
Clause 3 - Workload Related Items	
Benchmark Generic Parameters	Support-Files-for-Alibaba-Maxcompute-47nodes-100TB/Performance-Run-logs-20220731-054151-sql-sf100000/bigBench-configs/conf/userSettings.conf
Query Parameters used in the benchmark execution Settings	Support-Files-for-Alibaba-Maxcompute-47nodes-100TB/Performance-Run-logs-20220731-054151-sql-sf100000/bigBench-configs/sql/conf/queryParameters.sql
Benchmark Global Framework Parameters Settings	Support-Files-for-Alibaba-Maxcompute-47nodes-100TB/Performance-Run-logs-20220731-054151-sql-sf100000/bigBench-configs/sql/conf/engineSettings.sql
Load Test script	Support-Files-for-Alibaba-Maxcompute-47nodes-100TB/Performance-Run-logs-20220731-054151-sql-sf100000/bigBench-configs/sql/population/odpsCreateLoad.sql
Queries specific optimization parameters settings	Support-Files-for-Alibaba-Maxcompute-47nodes-100TB/Performance-Run-logs-20220731-054151-sql-sf100000/bigBench-configs/sql/queries/q[01-30]/engineLocalSettings.conf
Clause 4 - SUT Related Items	
Data Redundancy report	Support-Files-for-Alibaba-Maxcompute-47nodes-100TB/pangu_redundant_info_20220731-140754.txt
Benchmark execution script	Support-Files-for-Alibaba-Maxcompute-47nodes-100TB/TPCxBB_FullBenchmark_sequence_run.sh
Hardware and Software Report from a representative master node	Support-Files-for-Alibaba-Maxcompute-47nodes-100TB/envInfo-s54c06379.cloud.eo166/envInfo.log
Hardware and Software Report from a representative compute node	Support-Files-for-Alibaba-Maxcompute-47nodes-100TB/envInfo-s78c12077.cloud.eo166/envInfo.log
Hardware and Software Report from a representative storage node	Support-Files-for-Alibaba-Maxcompute-47nodes-100TB/envInfo-s54c10320.cloud.eo166/envInfo.log
All Framework configuration files are included in the Supporting Files Archive	Support-Files-for-Alibaba-Maxcompute-47nodes-100TB/odps_config.ini Support-Files-for-Alibaba-Maxcompute-47nodes-100TB/session_config.ini
Clause 5 - Metric and Scale Factor Related Items	
Benchmark Performance Report	Support-Files-for-Alibaba-Maxcompute-47nodes-100TB/Performance-Run-logs-20220731-054151-sql-sf100000/run-logs/BigBenchResult.log
Validation Test Report	Support-Files-for-Alibaba-Maxcompute-47nodes-100TB/Validation-Run-logs-20220801-071510-sql-sf30000/run-logs/BigBenchResult.log