

# Alibaba Cloud Computing Ltd.

---

TPC Express Benchmark™ Big Bench (TPCx-BB)

Full Disclosure Report

for

Alibaba Cloud MaxCompute

(with 58x Compute Nodes,

9x Storage Nodes

3x Admin Nodes)

using

MaxCompute v3.37

and

Alibaba Group Enterprise Linux Server 7.2 (Paladin)

---

**First Edition**

**August 18, 2021**

**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.

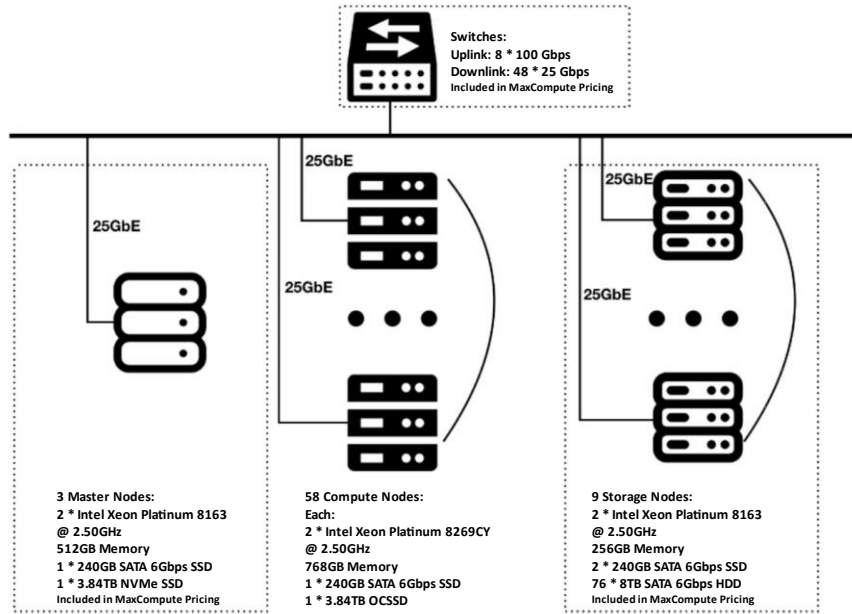
The Alibaba products, services or features identified in this document may not yet be available or may not be available in all areas and may be subject to change without notice. Consult your local Alibaba business contact for information on the products or services available in your area. You can find additional information via Alibaba's web site at [www.alibabacloud.com](http://www.alibabacloud.com). Actual performance and environmental costs of Alibaba products will vary depending on individual customer configurations and conditions.

**Copyright © 2021 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.

		<b>Alibaba Cloud MaxCompute</b>		TPCx-BB Rev. v1.5.0 TPC-Pricing Rev. v2.7.0	
				Report Date: August 18, 2021	
Total System Cost		TPCx-BB Performance Metric		Price/Performance	
<b>3,664,590 USD</b>		<b>46,059.27</b> BBQpm@100000		<b>79.57 USD</b> \$/BBQpm@100000	
Framework	Operating System	Other Software	Availability Date	Scale Factor	Streams
MaxCompute v3.37	Alibaba Group Enterprise Linux Server 7.2 (Paladin)	None	August 1, 2021	100000	4

### System Configuration



Physical Storage/Scale Factor: 57.25		Scale Factor/Physical Memory: 2.07	
Servers: 3x Admin Nodes / 58x Compute Node / 9x Storage Node Total Processors/Cores/Threads: 140/3,592/7,184			
<b>3x Master Node:</b> 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 ConnectX-4 Lx		<b>58x Compute Node (Compute):</b> 2x Intel(R) Xeon(R) Platinum 8269CY CPU @ 2.50GHz 768 GiB Onboard SATA Controller 1x 240 GB SATA 6 Gbps SSD 1x 3.84 TB OCSSD Mellanox MT27710 ConnectX-4 Lx	
		<b>9x Storage Node (Storage):</b> 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 ConnectX-4 Lx	
Connectivity:		Network Switch (8x 100 Gbps Up; 48x 25 Gbps Down)	



# Alibaba Cloud MaxCompute

TPCx-BB Rev. v1.5.0  
 TPC-Pricing Rev. v2.7.0

Report Date:  
 August 18, 2021

Description	Part Number	Source	Unit Price	Qty	Ext. Price	3-Year Maint.
<b>License Compute and Software Services</b>						
<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				58		
Intel® Xeon® Platinum 8269CY @ 2.50 GHz				2		
32 GB Memory				24		
240 GB SATA 6 Gbps SSD				1		
3.84 TB OCSSD				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						
<b>License Compute and Software Services Sub-Total</b>					<b>\$3,340,460.22</b>	<b>\$320,232.24</b>
<b>Other Components</b>						
13-inch MacBook Pro 1.4GHz (includes 2 spares)		2	\$1,299.00	3	\$3,897.00	
<b>Other Components Sub-Total</b>					<b>\$3,897.00</b>	<b>\$0.00</b>

Pricing: 1 = Alibaba; 2 = Apple.com  <sup>(1)</sup> 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.  <b>Audited by Doug Johnson, InfoSizing</b>	<b>Three-Year Cost of Ownership</b>	<b>\$3,664,590</b>
	<b>BBQpm@100000</b>	<b>46,059.27</b>
	<b>\$/BBQpm@100000</b>	<b>\$ 79.57</b>

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](mailto:pricing@tpc.org). Thank you.

### Numerical Quantities

Scale Factor	100000
Streams	4
SUT Validation Test	PASS

### Performance Run (Run 1)

Overall Run Start Time	2021-07-25 10:50:43.451
Overall Run End Time	2021-07-25 17:39:23.742
Overall Run Elapsed Time	24,520.291
Load Test Start Time	2021-07-25 10:50:43.452
Load Test End Time	2021-07-25 12:06:37.186
Load Test Elapsed Time	4,553.734
Power Test Start Time	2021-07-25 12:06:37.188
Power Test End Time	2021-07-25 14:14:07.715
Power Test Elapsed Time	7,650.527
Throughput Test Start Time	2021-07-25 14:14:07.715
Throughput Test End Time	2021-07-25 17:39:23.742
Throughput Test Elapsed Time	12,316.027
Performance Metric (BBQpm@ 100000)	46,059.27

### Repeatability Run (Run 2)

Overall Run Start Time	2021-07-25 20:12:43.107
Overall Run End Time	2021-07-26 02:46:40.529
Overall Run Elapsed Time	23,637.422
Load Test Start Time	2021-07-25 20:12:43.108
Load Test End Time	2021-07-25 21:13:40.809
Load Test Elapsed Time	3,657.701
Power Test Start Time	2021-07-25 21:13:40.811
Power Test End Time	2021-07-25 23:00:08.340
Power Test Elapsed Time	6,387.529
Throughput Test Start Time	2021-07-25 23:00:08.340
Throughput Test End Time	2021-07-26 02:46:40.528
Throughput Test Elapsed Time	13,592.188
Performance Metric (BBQpm@ 100000)	47,385.53

## Performance Run Report (Run 1)

\*\*\*\*\*

TPCx-BB

Result

v1.5.0

\*\*\*\*\*

INFO: T\_LOAD = 4553.734

INFO: T\_LD = 0.1 \* T\_LOAD: 455.3734

INFO: T\_PT = 3871.6012115828

INFO: T\_T\_PUT = 12316.027

INFO: T\_TT = 3079.00675

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 = 46059.2701226814

## Repeatability Run Report (Run 2)

\*\*\*\*\*

TPCx-BB

Result

v1.5.0

\*\*\*\*\*

INFO: T\_LOAD = 3657.701

INFO: T\_LD = 0.1 \* T\_LOAD: 365.7701

INFO: T\_PT = 3468.02375756023

INFO: T\_T\_PUT = 13592.188

INFO: T\_TT = 3398.047

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 = 47385.5399634586

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 .....	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.....	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 v1.5.0.

The test was conducted at a Scale Factor of 100000 with 70 nodes (58x Compute Node, 9x Storage Node) running MaxCompute v3.37 on Alibaba Group Enterprise Linux Server 7.2 (Paladin).

## Measured Configuration

Company Name	Cluster Node	Virtualization	Operating System
Alibaba Cloud Computing Ltd.	3x Admin Nodes 58x 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	46,059.27	79.57 USD	August 1, 2021



# 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](http://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](http://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.

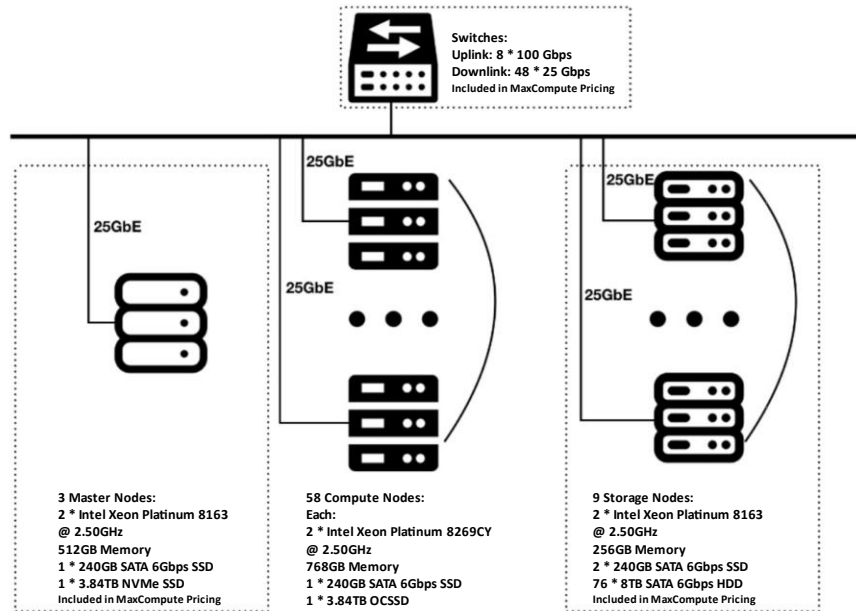
## 1.3 Configuration Diagrams

*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:	70
Total Processors/Cores/Threads:	140/3,592/7,184
Total Memory:	48,384 GiB
Total Number of Storage Devices:	824
Total Storage Capacity:	5,725,200

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

<b>3x Master Nodes each with:</b>	<b>58x Compute Node (Compute):</b>	<b>9x Storage Node (Storage):</b>
2/48/96 (processors/cores/threads)	2/52/104	2/48/96
2x Intel(R) Xeon(R) Platinum 8163 CPU @ 2.50GHz	2x Intel(R) Xeon(R) Platinum 8269CY CPU @ 2.50GHz	2x Intel(R) Xeon(R) Platinum 8163 CPU @ 2.50GHz
384 GiB	768 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 OCSSD	76x 8 TB SATA 6 Gbps HDD
Mellanox MT27710 ConnectX-4 Lx	Mellanox MT27710 ConnectX-4 Lx	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	s54a04374.cloud.eo166	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: 1 * 240GB SATA 6Gbps SSD Data Drive: 1 * 3.84TB NVMe SSD
Maxcompute Master Node2	Pangu master Fuxi Master Nuwa	1	N	s54c06379.cloud.eo166	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: 1 * 240GB SATA 6Gbps SSD Data Drive: 1 * 3.84TB NVMe SSD
Maxcompute Master Node3	Pangu master Fuxi Master Nuwa	1	N	s54d06393.cloud.eo166	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: 1 * 240GB SATA 6Gbps SSD Data Drive: 1 * 3.84TB NVMe SSD
Maxcompute Compute Nodes	Fuxi tubo	58	N	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB/nodelist.txt	2 * Intel Xeon Platinum 8269CY @ 2.50GHz 768GB Memory 1 * 240GB SATA 6Gbps SSD 1 * 3.84TB OCSSD 2 * 25Gb OS: Linux 3.10.0-327.ali2016.alios7.x86_64 Apsara: 1.11_u32	OS: 1 * 240GB SATA 6Gbps SSD Data Drive: 1 * 3.84TB OCSSD
Maxcompute Storage Nodes	Pangu ChunkServer	9	N	Support-Files-for-Alibaba-Maxcompute-58nodes-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: 2 * 240GB SATA 6Gbps SSD Data Drive: 76 * 8TB SATA 6Gbps HD

## 2.2 Distributed File System Implementation

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

MaxCompute v3.37.

## 2.3 Engine Implementation

*The Engine implementation and corresponding version must be disclosed.*

Component	Version
MaxCompute	3.37

## 2.4 Frameworks

*Frameworks and Engine used in the benchmark should be disclosed.*

Framework	Version
MaxCompute	3.37

## 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.*

<b>TPCx-BB Kit Version</b>
v1.5.0

## 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 (Performance Run)**

```
*****
TPCx-BB
Result
v1.5.0
*****
INFO: T_LOAD = 4553.734
INFO: T_LD = 0.1 * T_LOAD: 455.3734
INFO: T_PT = 3871.6012115828
INFO: T_T_PUT = 12316.027
INFO: T_TT = 3079.00675
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 = 46059.2701226814
```

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

```
*****
TPCx-BB
Result
v1.5.0
*****
INFO: T_LOAD = 3657.701
INFO: T_LD = 0.1 * T_LOAD: 365.7701
INFO: T_PT = 3468.02375756023
INFO: T_T_PUT = 13592.188
INFO: T_TT = 3398.047
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 = 47385.5399634586
```

### 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	49.832	63.381	47.941	151.966	56.732
	6	51.919	638.046	270.011	512.389	145.370
	7	55.495	52.381	223.356	59.509	67.616
	9	31.011	296.005	36.848	72.148	67.033
	11	32.626	32.838	137.781	164.304	32.942
	13	43.851	220.991	109.408	200.803	53.646
	14	28.277	34.729	44.908	46.197	230.905
	15	42.867	96.341	87.839	48.083	101.507
	16	125.668	116.268	584.072	157.788	115.741
	17	42.756	57.214	150.013	66.313	37.220
	20	280.459	279.513	287.472	424.734	355.789
	21	355.796	348.969	319.213	272.379	393.932
	22	27.679	32.161	25.285	26.600	34.571
	23	59.697	91.617	185.288	175.639	161.690
	24	58.433	54.511	122.603	65.603	72.711
25	272.662	598.340	396.297	286.027	520.726	
26	229.344	246.557	324.777	239.093	245.611	
29	100.079	558.357	267.596	408.100	93.723	
Semi-structured	2	652.291	1,055.789	1,084.737	921.003	890.984
	3	314.446	559.232	495.371	628.885	574.635
	4	411.225	463.559	631.493	431.292	2,096.173
	5	617.663	793.845	811.049	849.074	1,007.826
	8	181.181	304.903	365.933	252.119	331.530
	12	110.320	83.178	157.884	102.341	212.528
	30	1,002.061	2,325.037	585.203	2,279.878	870.044
Unstructured	10	189.245	213.204	399.328	229.697	195.658
	18	1,755.242	1,988.943	3,050.600	2,047.447	1,588.164
	19	343.731	448.940	494.302	409.210	343.663
	27	54.079	100.059	75.495	216.720	64.583
	28	130.551	161.104	142.746	145.193	316.804

### 3.5 Validation Test Output

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

<b>Query Number</b>	<b>Query Execution</b>	<b>Output Validation</b>
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/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 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)
61	1	240	14,640
61	1	3,840	234,240
9	2	240	4,320
9	76	8,000	5,472,000

Total Storage (GB)	5,725,200
Scale Factor	100000
Data Storage Ratio	57.25

## 4.5 Scale Factor to Memory Ratio

*The Scale Factor to memory ratio must be disclosed.*

Nodes	Memory (GB)	Total (GB)
3	512	1,536
58	768	44,544
9	256	2,304

Scale Factor	100000
Total Memory (GB)	48,384
SF / Memory Ratio	2.07

# 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 46,059.27

## 5.2 Repeatability Run Metric

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

Repeatability Run
BBQpm@100000 47,385.53

## 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 79.57

## 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.*

<b>Run</b>	<b>Elapsed Time</b>	<b>Seconds</b>
Run 1	00 06:48:40.291	24,520.291
Run 2	00 06:33:57.422	23,637.422

## 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.*

<b>Test</b>	<b>Performance Run</b>	<b>Repeatability Run</b>
Load Test	4,553.734	3,657.701
Power Test	7,650.527	6,387.529
Throughput Test	12,316.027	13,592.188

# 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](http://www.tpc.org).

A copy of the auditor's attestation letter is included in the next two pages.

Xiening Dai  
205 108<sup>th</sup> Ave NE, Suite 400  
Bellevue, WA 98004

August 13, 2021

I verified the TPC Express Benchmark™ BB v1.5.0 performance of the following configuration:

Platform: Alibaba Cloud MaxCompute  
(w/ 3x Master Nodes, 58x Compute Nodes, 9 Storage Nodes)  
Operating System: Alibaba Group Enterprise Linux Server 7.2 (Paladin)  
Framework: MaxCompute v3.37

The results were:

**Performance Metric**      **46,059.27 BBQpm@100000**  
Run Elapsed Time            00 06:48:40.291 (24,520.291 Seconds)

**Cluster**    **3x Master Nodes, 58x 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 8269CY (2.50 GHz, 18-core, 35.75 MB L3) (Compute nodes)		
Memory	512GiB (Master nodes), 768GiB (Compute nodes), 256GiB (Storage nodes)		
Storage	<b>Qty</b>	<b>Size</b>	<b>Type</b>
	1	240 GB	SATA 6Gbps SSD (Master, Compute nodes)
	2	240 GB	SATA 6Gbps SSD (Storage nodes)
	1	3.84 TB	NVMe SSD (Master nodes)
	1	3.84 TB	OCSSD (Compute nodes)
	76	8 TB	SATA 6Gbps 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 v1.5.0
- 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,

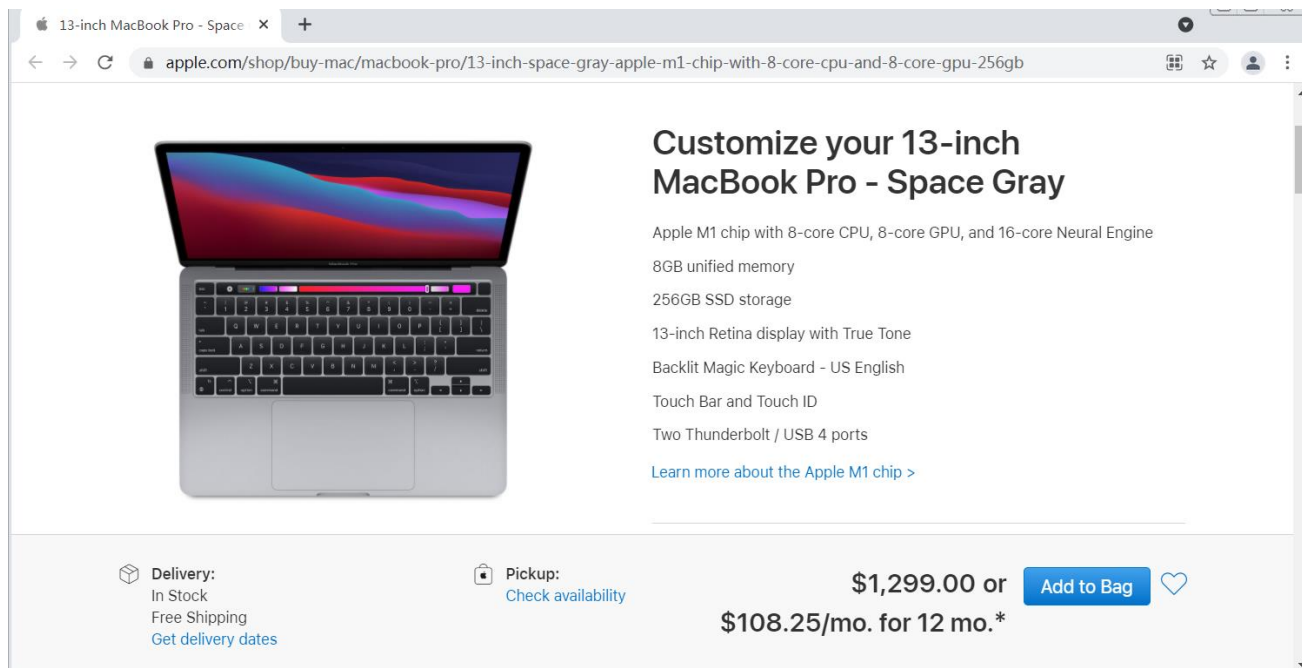
A handwritten signature in cursive script that reads "Doug Johnson". The signature is written in black ink and has a long, sweeping horizontal line extending to the right.

Doug Johnson, TPC Auditor

63 Lourdes Dr. | Leominster, MA 01453 | 978-343-6562 | [www.sizing.com](http://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
<b>Clause 1 - General Items</b>	
The Supporting Files Archive contains the parameters and options used to configure the components involved in this benchmark	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB
Validation Run Files	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB/Validation-Run-logs-20210725-084511-sql-sf100000
Performance Run Files	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB/Performance-Run-logs-20210725-175133-sql-sf100000
Repeatability Run Files	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB/Repeatability-Run-logs-20210726-030803-sql-sf100000
<b>Clause 3 - Workload Related Items</b>	
Benchmark Generic Parameters	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB/Performance-Run-logs-20210725-175133-sql-sf100000/bigBench-configs/conf/userSettings.conf
Query Parameters used in the benchmark execution Settings	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB/Performance-Run-logs-20210725-175133-sql-sf100000/bigBench-configs/sql/conf/queryParameters.sql
Benchmark Global Framework Parameters Settings	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB/Performance-Run-logs-20210725-175133-sql-sf100000/bigBench-configs/sql/conf/engineSettings.sql
Benchmark Global Framework Parameters Settings	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB/Performance-Run-logs-20210725-175133-sql-sf100000/bigBench-configs/sql/population/odpsCreateLoad.sql
Load Test script	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB/Performance-Run-logs-20210725-175133-sql-sf100000/bigBench-configs/sql/queries/q[01-30]/engineLocalSettings.conf
Queries specific optimization parameters settings	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB/Performance-Run-logs-20210725-175133-sql-sf100000/bigBench-configs/conf/userSettings.conf
Queries specific optimization parameters settings	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB/Performance-Run-logs-20210725-175133-sql-sf100000/bigBench-configs/sql/conf/queryParameters.sql
<b>Clause 4 - SUT Related Items</b>	
Data Redundancy report	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB/pangu_redundant_info_20210726-030925.txt
Benchmark execution script	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB/TPCxBB_FullBenchmark_sequence_run.sh
Hardware and Software Report from a representative master node	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB/envInfo-s54a04374.cloud.eo166/envInfo.log
Hardware and Software Report from a representative compute node	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB/envInfo-s37a09328.cloud.eo166/envInfo.log
Hardware and Software Report from a representative storage node	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB/envInfo-s39e12451.cloud.eo166/envInfo.log
<b>Clause 5 - Metric and Scale Factor Related Items</b>	
Benchmark Performance Report	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB/Performance-Run-logs-20210725-175133-sql-sf100000/run-logs/BigBenchResult.log
Validation Test Report	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB/Validation-Run-logs-20210725-084511-sql-sf100000/run-logs/BigBenchResult.log
<b>Clause 6 – Other</b>	
compute Nodes	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB/nodelist.txt
storage Nodes	Support-Files-for-Alibaba-Maxcompute-58nodes-100TB/storage-nodelist.txt