Alibaba Cloud Computing Ltd.

TPC Benchmark™ DS

Full Disclosure Report

for

Alibaba Cloud MaxCompute

(with 1600 MaxCompute Query Accelerator CU)

using

MaxCompute Query Accelerator 2.0 v5.57

and

Alibaba Cloud Linux 3 (Soaring Falcon)

First Edition

August 4, 2026

First Edition – August 4, 2026

Table of Contents

Abstract 5

Preface 10
TPC Benchmark™ DS Overview 10

General Items 11
0.1 Test Sponsor 11
0.2 Parameter Settings 11
0.3 Configuration Diagrams 11
Measured Configuration 11
Configuration of the Measured System 12
Priced Configuration 12

Clause 2: Logical Database Design Related Items 13
2.1 Database Definition Statements 13
Supporting Files Index 13
2.2 Physical Organization 13
2.3 Horizontal Partitioning 13
2.4 Replication 13

Clause 3: Scaling and Database Population 14
3.1 Initial Cardinality of Tables 14
3.2 Distribution of Tables and Logs Across Media 14
3.3 Mapping of Database Partitions and Replications 14
3.4 Implementation of RAID 15
3.5 Data Generation Tool Modifications 15
3.6 Database Load Time 16
3.7 Data Storage Ratio 17
3.8 Database Load Mechanism Details and Illustration 17
3.9 Qualification Database Configuration 18

Clauses 4–5: Query, Substitution and Data Maintenance 19
4.1 Query Language 19
4.2 Verifying Method of Random Number Generation 19
4.3 Generating Values for Substitution Parameters 19
4.4 Query Text and Output Data from Qualification Database 19
4.5 Query Substitution Parameters and Seeds Used 19
4.6 Refresh Setting 19
4.7 Source Code of Refresh Functions 19
4.8 Staging Area 20

Clause 6: Data Persistence Properties Related Items 21
6.1 Method Used to Satisfy the Data Accessibility Requirement 21
6.2 Data Accessibility Architecture 21
6.3 Public Data Accessibility Documentation 21
6.4 Supporting Evidence 21

Clause 7: Performance Metrics and Execution Rules Related Items 23
7.1 System Activity 23
7.2 Test Steps 23
7.3 Timing Intervals for Each Query and Refresh Function 23
7.4 Throughput Test Result 23
7.5 Time for Each Stream 23
7.6 Time for Each Refresh Function 23
7.7 Performance Metrics 23

Clause 8: SUT and Driver Implementation Related Items 24

8.1 Driver 24
8.2 Implementation Specific Layer (ISL) 24
8.3 Profile-Directed Optimization 24

Clause 9: Pricing Related Items 25
9.1 Hardware and Software Used 25
9.2 Availability Date 25
9.3 Country-Specific Pricing 25

Clause 11: Audit Related Items 26

Supporting Files Index 29

Appendix A: Purchase Page of Creating Alibaba Cloud MaxCompute with 1-Year Subscription 30

Appendix B: Aliyun MaxCompute Resource Utilization 31

Appendix C: Third Party Price Quotes 32

Abstract

This document contains the methodology and results of the TPC Benchmark™ DS (TPC-DS) test conducted in conformance with the requirements of the TPC-DS Standard Specification, Revision 4.0.0.

The test was conducted at a Scale Factor of 100000 GB with 1600 CUs running Alibaba Cloud MaxCompute Query Accelerator 2.0 v5.57 on Alibaba Cloud Linux 3 (Soaring Falcon).

Measured Configuration

Company Name	Cluster Node	Database Software	Operation System
Alibaba Cloud Computing Ltd.	Alibaba Cloud MaxCompute Query Accelerator 2.0 CU	Alibaba Cloud MaxCompute Query Accelerator 2.0 v5.57	Alibaba Cloud Linux 3 (Soaring Falcon)

TPC Benchmark™ DS Metrics

Total System Cost (CNY)	TPC-DS Throughput (QphDS@100000)	Price/Performance (CNY / kQphDS@100000)	Availability Date
¥2,717,787	**133,518,293**	**¥20.36**	**2026-08-04**

Alibaba Cloud	Alibaba Cloud MaxCompute		TPC-DS: 4.0.0 TPC-Pricing: 2.9.0 Report Date: Aug. 4, 2026	
Total System Cost	TPC-DS Throughput	Price / Performance	System Availability Date	
¥2,717,787 CNY	**133,518,293** QphDS@100000	**¥20.36** CNY/kQphDS@100000	**2026-08-04**	
Dataset Size[1]	Database Manager	Operation System	Other Software	Cluster
100,000 GB	Alibaba Cloud MaxCompute Query Accelerator 2.0 v5.57	Alibaba Cloud Linux 3 (Soaring Falcon)	N/A	Yes

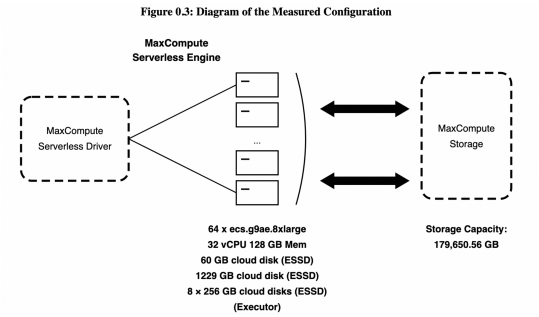

Benchmarked Configuration

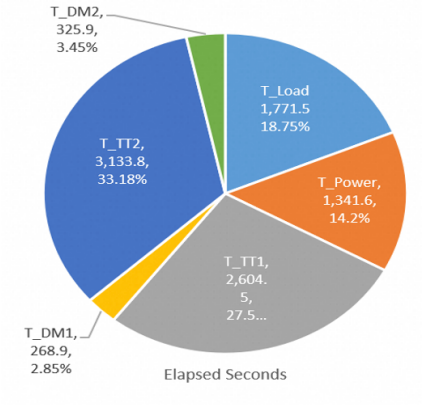

Elapsed Time

Load includes backup = No	RAID = No

System Configuration:	**Alibaba Cloud MaxCompute Query Accelerator 2.0 (MaxCompute Serverless Engine)**
Workspace Capacity[2]:	1600 CU
Servers:	64 × ecs.g9ae.8xlarge Executor Pods
Total Processors/Cores/Threads:	2,048 vCPUs
Total Memory:	8,192 GB
Total Storage[3]:	393,218.56 GB
Storage Ratio[4]:	3.94
Server Configuration:	**Per Executor Pod (64)**
Processors:	AMD EPYC 9T95 192-Core Processor
Memory:	128 GB
Network:	Bandwidth: 16 Gbps, Packet forwarding rate: 5,000,000
Storage Devices:	1 × 60 GB Cloud Disk (ESSD) 1 × 1,229 GB Cloud Disk (ESSD) 8 × 256 GB Cloud Disk (ESSD) 3,337 GB aggregate ESSD capacity per instance

1. Dataset Size includes only raw data (i.e., no temp, index, redundant storage space, etc.).
2. Each Executor Pod is provisioned with 32 vCPUs and 128 GB of gross memory. Under the MCQA product design, 25 vCPUs and 100 GB are customer-allocatable and correspond to 25 CUs. The remaining 7 vCPUs and 28 GB are reserved for the operating system and MCQA service runtime. Therefore, 64 Executor Pods provide a total customer-allocatable capacity of 1,600 CUs.
3. Total Storage = 64 × 3,337 GB Executor Storage + 179,650.56 GB MaxCompute Storage = 393,218.56 GB.
4. Storage Ratio = 393,218.56 GB / 100,000 GB = 3.9321856, rounded up to the nearest one-hundredth = 3.94.

Alibaba Cloud	**Alibaba Cloud MaxCompute**	TPC-DS: 4.0.0 TPC-Pricing: 2.9.0 Report Date: Aug. 4, 2026

Description	Part Number	Src	Unit Price (CNY)	Qty	Ext. Price (CNY)	3-Year Maint. (CNY)
Licensed Compute and Software Services						
MaxCompute,						
1600 CU, 1-year prepaid subscription (List Price)	(China East 2)	1	240,000.00	12	2,880,000.00	included
Standard 1-year subscription discount (15%) (See Note 1)	(China East 2)	1	-36,000.00	12	-432,000.00	0.00
Data Storage	(China East 2)	1	21,857.48	12	262,289.76	Included
Licensed Compute and Software Services Sub-Total					**2,710,289.76**	**0.00**
Other Components						
Lenovo P14H Laptop (includes 2 spares)		2	2,499.00	3	7,497.00	
Other Components Sub-Total					**7,497.00**	**0.00**

1 = Alibaba Cloud, 2 = Tmall.com

Note 1: The 15% discount is the standard discount for a 1-year prepaid subscription and is available to all customers purchasing the same 1600-CU configuration in China East 2. It is not individually negotiated.
Unit prices for recurring licensed services are monthly prices. Extended prices reflect 12 months of service. The MaxCompute 1600-CU subscription is prepaid in full for one year.

Audited by Doug Johnson, InfoSizing

1-Year Cost of Ownership:	**2,717,787**
QphDS@100000:	**133,518,293**
¥ /kQphDS@100000:	**20.36**

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.

Alibaba Cloud	Alibaba Cloud MaxCompute	TPC-DS: 4.0.0 TPC-Pricing: 2.9.0 Report Date: Aug. 4, 2026

Metrics Details:

Name	Value	Description / Unit
SF	100,000	Scale Factor 100TB
S	8	Total Throughput Streams
Sq	4	Streams / Throughput Test
Q	396	Queries / Throughput Test
T_LD	0.0197	hours@100000
T_PT	1.4907	hours@100000
T_TT	1.5940	hours@100000
T_DM	0.1653	hours@100000

Secondary Details:

Name	Value	Unit
T_Load	1,771.5	seconds@100000
T_Power	1,341.6	seconds@100000
T_TT1	2,604.5	seconds@100000
T_TT2	3,133.8	seconds@100000
T_DM1	268.9	seconds@100000
T_DM2	325.9	seconds@100000

Test Timeline:

Test	Start	End	Seconds	(hh:mm:ss)
Load	2026-08-04 01:07:21.728	2026-08-04 01:36:53.177	1,771.449	00:29:31
Audit/Admin	2026-08-04 01:36:53.000	2026-08-04 04:17:48.000	9,655.000	02:40:55
Power	2026-08-04 04:17:50.124	2026-08-04 04:40:11.675	1,341.551	00:22:22
TT-1	2026-08-04 04:40:22.965	2026-08-04 05:23:47.452	2,604.487	00:43:24
DM-1	2026-08-04 05:23:57.654	2026-08-04 05:28:26.539	268.885	00:04:29
TT-2	2026-08-04 05:29:06.141	2026-08-04 06:20:40.297	3,094.156	00:51:34
DM-2	2026-08-04 06:20:50.330	2026-08-04 06:26:16.147	325.817	00:05:26

Stream	Start	End	Seconds	(hh:mm:ss)
Power - 0	2026-08-04 04:17:50.124	2026-08-04 04:40:11.675	1,341.551	00:22:22
TT-1 - 1	2026-08-04 04:40:22.991	2026-08-04 05:23:26.868	2,583.877	00:43:04
TT-1 - 2	2026-08-04 04:40:22.982	2026-08-04 05:23:47.452	2,604.470	00:43:24
TT-1 - 3	2026-08-04 04:40:22.989	2026-08-04 05:23:28.372	2,585.383	00:43:05
TT-1 - 4	2026-08-04 04:40:22.965	2026-08-04 05:22:56.519	2,553.554	00:42:34
DM-1 - 1	2026-08-04 05:23:57.654	2026-08-04 05:25:48.576	110.922	00:01:51
DM-1 - 2	2026-08-04 05:26:24.729	2026-08-04 05:28:26.539	121.810	00:02:02
TT-2 - 5	2026-08-04 05:29:06.150	2026-08-04 06:19:55.215	3,049.065	00:50:49
TT-2 - 6	2026-08-04 05:29:06.141	2026-08-04 06:19:05.170	2,999.029	00:49:59
TT-2 - 7	2026-08-04 05:29:06.142	2026-08-04 06:20:40.297	3,094.155	00:51:34
TT-2 - 8	2026-08-04 05:29:06.154	2026-08-04 06:19:44.845	3,038.691	00:50:39
DM-2 - 3	2026-08-04 06:20:50.330	2026-08-04 06:23:12.311	141.981	00:02:22
DM-2 - 4	2026-08-04 06:23:48.819	2026-08-04 06:26:16.147	147.328	00:02:27

Timing Intervals for Each Query (In Seconds)

Query	Stream 0	Stream 1	Stream 2	Stream 3	Stream 4	Min	25th	Median	75th	Max	Stream 5	Stream 6	Stream 7	Stream 8	Min	25th	Median	75th	Max
1	5.0	25.1	8.6	12.5	25.4	8.60	11.53	18.80	25.18	25.40	9.5	6.9	7.2	9.0	6.90	7.13	8.10	9.13	9.50
2	14.1	26.7	24.4	20.8	34.0	20.80	23.50	25.55	28.53	34.00	21.4	21.6	28.3	32.1	21.40	21.55	24.95	29.25	32.10
3	4.3	3.8	14.1	7.2	6.5	3.80	5.83	6.85	8.93	14.10	5.7	7.9	16.2	6.2	5.70	6.08	7.05	9.98	16.20
4	33.0	46.8	44.4	63.9	73.9	44.40	46.20	55.35	66.40	73.90	73.9	70.7	55.0	89.1	55.00	66.78	72.30	77.70	89.10
5	11.8	14.1	15.9	15.8	16.3	14.10	15.38	15.85	16.00	16.30	29.2	26.2	31.1	20.2	20.20	24.70	27.70	29.68	31.10
6	4.1	27.6	4.7	24.4	14.0	4.70	11.68	19.20	25.20	27.60	4.8	6.1	19.6	10.9	4.80	5.78	8.50	13.08	19.60
7	7.8	13.9	30.8	22.5	26.0	13.90	20.35	24.25	27.20	30.80	21.5	27.2	12.6	25.2	12.60	19.28	23.35	25.70	27.20
8	4.9	18.3	15.0	7.6	13.8	7.60	12.25	14.40	15.83	18.30	25.5	13.2	13.2	7.0	7.00	11.65	13.20	16.28	25.50
9	26.5	42.0	74.4	48.7	27.6	27.60	38.40	45.35	55.13	74.40	48.6	38.4	39.2	40.8	38.40	39.00	40.00	42.75	48.60
10	7.0	22.7	12.9	36.3	11.9	11.90	12.65	17.80	26.10	36.30	10.0	10.0	24.7	33.1	10.00	10.00	17.35	26.80	33.10
11	16.3	36.5	25.1	48.7	20.5	20.50	23.95	30.80	39.55	48.70	27.6	33.5	38.2	58.4	27.60	32.03	35.85	43.25	58.40
12	3.7	5.9	3.5	6.5	3.6	3.50	3.58	4.75	6.05	6.50	21.2	21.7	6.9	6.5	6.50	6.80	14.05	21.33	21.70
13	22.0	44.0	25.4	25.6	24.0	24.00	25.05	25.50	30.20	44.00	64.3	45.7	47.6	47.0	45.70	46.68	47.30	51.78	64.30
14	40.5	38.2	37.2	15.5	102.3	15.50	31.78	37.70	54.23	102.30	15.8	85.4	94.2	126.5	15.80	68.00	89.80	102.28	126.50
15	5.8	9.9	11.6	18.7	4.5	4.50	8.55	10.75	13.38	18.70	20.4	29.7	11.1	7.0	7.00	10.08	15.75	22.73	29.70
16	27.4	22.0	58.6	57.8	46.8	22.00	40.60	52.30	58.00	58.60	25.6	9.5	49.1	83.2	9.50	21.58	37.35	57.63	83.20
17	8.2	8.9	25.9	20.2	27.0	8.90	17.38	23.05	26.18	27.00	10.4	59.4	11.0	18.9	10.40	10.85	14.95	29.03	59.40
18	9.1	13.3	27.2	18.1	16.3	13.30	15.55	17.20	20.38	27.20	38.0	33.9	20.8	16.9	16.90	19.83	27.35	34.93	38.00
19	4.9	11.3	10.2	13.7	12.3	10.20	11.03	11.80	12.65	13.70	22.8	10.8	28.4	37.5	10.80	19.80	25.60	30.68	37.50
20	3.6	18.0	13.7	12.0	19.8	12.00	13.28	15.85	18.45	19.80	8.8	13.8	25.4	17.6	8.80	12.55	15.70	19.55	25.40
21	2.2	16.8	3.8	3.0	6.0	3.00	3.60	4.90	8.70	16.80	9.4	9.8	5.7	18.1	5.70	8.48	9.60	11.88	18.10
22	2.8	3.7	14.7	8.6	11.6	3.70	7.38	10.10	12.38	14.70	6.1	16.2	3.0	6.5	3.00	5.33	6.30	8.93	16.20
23	105.4	52.6	41.9	133.6	130.6	41.90	49.93	91.60	131.35	133.60	133.2	160.6	98.1	79.6	79.60	93.48	115.65	140.05	160.60
24	39.6	77.6	54.7	62.2	74.0	54.70	60.33	68.10	74.90	77.60	73.5	64.2	69.5	99.2	64.20	68.18	71.50	79.93	99.20
25	11.0	18.2	13.2	21.0	26.9	13.20	16.95	19.60	22.48	26.90	18.3	52.1	38.9	17.3	17.30	18.05	28.60	42.20	52.10
26	5.8	33.7	21.8	7.1	5.6	5.60	6.73	14.45	24.78	33.70	10.5	33.9	12.7	11.6	10.50	11.33	12.15	18.00	33.90
27	6.3	17.1	7.4	12.3	16.6	7.40	11.08	14.45	16.73	17.10	21.3	9.4	7.8	9.1	7.80	8.78	9.25	12.38	21.30
28	22.3	39.7	30.9	49.4	40.8	30.90	37.50	40.25	42.95	49.40	45.2	41.7	79.2	30.9	30.90	39.00	43.45	53.70	79.20
29	10.3	40.0	25.1	25.6	11.5	11.50	21.70	25.35	29.20	40.00	15.8	17.5	36.3	18.2	15.80	17.08	17.85	22.73	36.30
30	6.4	7.8	16.2	9.0	15.0	7.80	8.70	12.00	15.30	16.20	30.0	20.9	50.6	8.6	8.60	17.83	25.45	35.15	50.60
31	7.8	25.3	22.7	18.2	26.9	18.20	21.58	24.00	25.70	26.90	40.0	29.7	21.6	26.3	21.60	25.13	28.00	32.28	40.00
32	4.0	9.7	18.5	3.8	6.8	3.80	6.05	8.25	11.90	18.50	17.6	5.8	14.4	37.8	5.80	12.25	16.00	22.65	37.80
33	8.8	20.2	23.2	16.8	35.8	16.80	19.35	21.70	26.35	35.80	15.3	33.4	13.8	12.2	12.20	13.40	14.55	19.83	33.40
34	39.4	54.9	57.6	60.7	40.1	40.10	51.20	56.25	58.38	60.70	45.7	44.7	45.2	48.2	44.70	45.08	45.45	46.33	48.20
35	9.8	39.2	18.3	22.3	30.7	18.30	21.30	26.50	32.83	39.20	42.8	14.2	27.0	31.2	14.20	23.80	29.10	34.10	42.80
36	9.3	16.3	14.5	16.6	13.3	13.30	14.20	15.40	16.38	16.60	39.3	14.1	73.2	11.4	11.40	13.43	26.70	47.78	73.20
37	9.2	11.9	19.6	9.9	9.6	9.60	9.83	10.90	13.83	19.60	29.5	10.7	24.6	25.7	10.70	21.13	25.15	26.65	29.50
38	17.1	36.8	32.1	38.9	37.9	32.10	35.63	37.35	38.15	38.90	38.0	49.0	65.1	58.8	38.00	46.25	53.90	60.38	65.10
39	9.9	37.0	59.0	9.3	12.8	9.30	11.93	24.90	42.50	59.00	24.1	10.2	28.3	20.0	10.20	17.55	22.05	25.15	28.30
40	6.8	5.6	11.7	8.1	41.3	5.60	7.48	9.90	19.10	41.30	7.4	31.4	11.1	12.2	7.40	10.18	11.65	17.00	31.40
41	1.6	21.5	1.6	20.8	14.1	1.60	10.98	17.45	20.98	21.50	13.7	1.7	1.5	7.9	1.50	1.65	4.80	9.35	13.70
42	3.0	10.8	23.1	36.9	2.9	2.90	8.83	16.95	26.55	36.90	14.2	17.3	11.6	4.2	4.20	9.75	12.90	14.98	17.30
43	4.3	5.9	22.3	21.3	7.0	5.90	6.73	14.15	21.55	22.30	50.3	18.9	20.1	22.5	18.90	19.80	21.30	29.45	50.30
44	13.5	18.0	17.6	53.1	23.5	17.60	17.90	20.75	30.90	53.10	21.0	18.3	20.5	45.9	18.30	19.95	20.75	27.23	45.90
45	5.3	25.9	12.0	24.3	17.3	12.00	15.98	20.80	24.70	25.90	51.2	7.9	12.5	9.2	7.90	8.88	10.85	22.18	51.20
46	10.1	13.1	34.2	18.7	17.1	13.10	16.10	17.90	22.58	34.20	38.7	58.9	76.8	24.5	24.50	35.15	48.80	63.38	76.80
47	8.8	25.1	12.9	19.1	18.8	12.90	17.33	18.95	20.60	25.10	26.6	22.9	18.7	61.4	18.70	21.85	24.75	35.30	61.40
48	6.4	53.8	6.4	18.7	10.4	6.40	9.40	14.55	27.48	53.80	34.5	8.4	42.1	27.1	8.40	22.43	30.80	36.40	42.10
49	9.0	9.7	25.9	15.1	8.8	8.80	9.48	12.40	17.80	25.90	29.6	20.6	14.0	40.8	14.00	18.95	25.10	32.40	40.80
50	25.1	25.8	32.6	23.9	24.6	23.90	24.43	25.20	27.50	32.60	31.9	37.7	31.3	43.7	31.30	31.75	34.80	39.20	43.70
51	6.3	11.4	14.3	21.5	9.4	9.40	10.90	12.85	16.10	21.50	14.7	13.8	31.9	32.2	13.80	14.48	23.30	31.98	32.20
52	2.7	11.5	21.8	3.0	8.0	3.00	6.75	9.75	14.08	21.80	14.5	21.5	10.8	17.4	10.80	13.58	15.95	18.43	21.50
53	5.6	7.6	7.3	37.1	9.9	7.30	7.53	8.75	16.70	37.10	23.8	19.5	11.8	7.6	7.60	10.75	15.65	20.58	23.80
54	38.3	90.5	65.6	55.4	38.3	38.30	51.13	60.50	71.83	90.50	61.1	50.8	78.5	47.3	47.30	49.93	55.95	65.45	78.50
55	3.3	10.7	17.2	3.9	18.3	3.90	9.00	13.95	17.48	18.30	7.9	4.1	6.6	4.2	4.10	4.18	5.40	6.93	7.90
56	7.1	24.0	14.4	16.8	16.4	14.40	15.90	16.60	18.60	24.00	66.2	41.7	13.9	21.4	13.90	19.53	31.55	47.83	66.20
57	6.5	35.2	9.9	7.6	21.6	7.60	9.33	15.75	25.00	35.20	27.8	31.9	11.3	13.4	11.30	12.88	20.60	28.83	31.90
58	4.9	24.4	20.5	26.2	13.2	13.20	18.68	22.45	24.85	26.20	33.3	7.5	28.1	11.4	7.50	10.43	19.75	29.40	33.30
59	18.7	37.5	28.4	36.7	44.5	28.40	34.63	37.10	39.25	44.50	30.6	88.7	37.3	40.8	30.60	35.63	39.05	52.78	88.70
60	12.9	19.0	28.5	29.5	14.7	14.70	17.93	23.75	28.75	29.50	34.4	25.8	28.3	26.0	25.80	25.95	27.15	29.83	34.40
61	10.9	23.8	17.9	12.4	20.5	12.40	16.53	19.20	21.33	23.80	28.3	52.5	22.1	15.3	15.30	20.40	25.20	34.35	52.50
62	5.0	20.5	21.3	4.8	14.6	4.80	12.15	17.55	20.70	21.30	8.1	10.7	21.3	7.9	7.90	8.05	9.40	13.35	21.30
63	5.7	12.2	8.0	5.1	32.0	5.10	7.28	10.10	17.15	32.00	28.9	9.9	8.1	7.1	7.10	7.85	9.00	14.65	28.90
64	62.0	92.9	72.6	89.0	90.8	72.60	84.90	89.90	91.33	92.90	78.0	100.2	79.2	95.9	78.00	78.90	87.55	96.98	100.20
65	8.0	13.2	10.5	17.8	21.0	10.50	12.53	15.50	18.60	21.00	13.0	19.4	21.9	83.1	13.00	17.80	20.65	37.20	83.10
66	7.3	24.1	7.0	17.4	8.6	7.00	8.20	13.00	19.08	24.10	12.9	11.0	10.6	19.3	10.60	10.90	11.95	14.50	19.30
67	22.5	26.1	24.8	33.9	39.3	24.80	25.78	30.00	35.25	39.30	95.5	40.0	101.8	26.2	26.20	36.55	67.75	97.08	101.80
68	5.8	9.2	19.3	14.6	12.7	9.20	11.83	13.65	15.78	19.30	17.6	10.9	10.7	29.9	10.70	10.85	14.25	20.68	29.90
69	7.4	16.8	27.0	33.0	27.2	16.80	24.45	27.10	28.65	33.00	39.1	17.4	22.3	11.0	11.00	15.80	19.85	26.50	39.10
70	10.0	12.0	22.4	45.0	22.5	12.00	19.80	22.45	28.13	45.00	28.2	33.6	23.0	34.3	23.00	26.90	30.90	33.78	34.30
71	15.3	31.2	47.6	19.5	24.2	19.50	23.03	27.70	35.30	47.60	51.5	31.8	45.2	37.4	31.80	36.00	41.30	46.78	51.50
72	26.7	37.3	21.7	35.3	21.4	21.40	21.63	28.50	35.80	37.30	30.8	51.3	28.3	56.5	28.30	30.18	41.05	52.60	56.50
73	4.3	9.1	6.5	17.4	21.2	6.50	8.45	13.25	18.35	21.20	11.2	8.1	22.5	29.6	8.10	10.43	16.85	24.28	29.60
74	17.6	30.9	28.2	30.2	48.5	28.20	29.70	30.55	35.30	48.50	33.8	25.1	46.1	38.4	25.10	31.63	36.10	40.33	46.10
75	37.1	39.2	46.5	37.4	70.9	37.40	38.75	42.85	52.60	70.90	53.0	68.5	44.5	63.1	44.50	50.88	58.05	64.45	68.50
76	20.6	38.9	33.2	43.3	24.4	24.40	31.00	36.05	40.00	43.30	67.8	42.6	37.6	43.7	37.60	41.35	43.15	49.73	67.80
77	6.2	16.4	35.0	12.2	19.4	12.20	15.35	17.90	23.30	35.00	20.1	36.4	10.9	10.5	10.50	10.80	15.50	24.18	36.40
78	48.6	53.8	56.9	59.9	54.7	53.80	54.48	55.80	57.65	59.90	79.1	73.4	81.1	78.0	73.40	76.85	78.55	79.60	81.10
79	12.5	73.6	18.1	30.0	16.5	16.50	17.70	24.05	40.90	73.60	15.3	40.1	31.0	65.9	15.30	27.08	35.55	46.55	65.90
80	9.5	33.3	35.2	9.7	24.1	9.70	20.50	28.70	33.78	35.20	14.4	28.8	35.7	15.3	14.40	15.08	22.05	30.53	35.70
81	6.9	11.4	21.7	27.0	8.2	8.20	10.60	16.55	23.03	27.00	8.7	26.8	37.7	18.8	8.70	16.28	22.80	29.53	37.70
82	10.7	33.9	28.8	19.8	17.1	17.10	19.13	24.30	30.08	33.90	15.5	31.2	15.0	14.1	14.10	14.78	15.25	19.43	31.20
83	4.7	12.4	59.0	20.6	15.3	12.40	14.58	17.95	30.20	59.00	29.3	18.9	9.2	26.5	9.20	16.48	22.70	27.20	29.30
84	6.3	10.7	24.6	32.3	31.1	10.70	21.13	27.85	31.40	32.30	10.5	7.3	10.2	24.9	7.30	9.48	10.35	14.10	24.90
85	8.3	29.4	71.6	23.2	25.3	23.20	24.78	27.35	39.95	71.60	19.9	30.3	66.9	23.0	19.90	22.23	26.65	39.45	66.90
86	5.1	37.6	32.3	6.7	5.6	5.60	6.43	19.50	33.63	37.60	12.5	16.7	27.6	17.2	12.50	15.65	16.95	19.80	27.60
87	20.2	33.5	60.2	34.7	38.5	33.50	34.40	36.60	43.93	60.20	27.8	68.4	53.1	28.8	27.80	28.55	40.95	56.93	68.40
88	20.1	23.8	23.0	40.0	22.5	22.50	22.88	23.40	27.85	40.00	47.5	60.6	42.0	33.0	33.00	39.75	44.75	50.78	60.60
89	5.5	8.2	17.9	6.4	32.9	6.40	7.75	13.05	21.65	32.90	20.7	8.0	13.0	8.0	8.00	8.00	10.50	14.93	20.70
90	5.4	8.9	42.7	11.7	18.0	8.90	11.00	14.85	24.18	42.70	29.9	6.8	9.3	27.0	6.80	8.68	18.15	27.73	29.90
91	5.2	9.9	12.2	14.4	21.6	9.90	11.63	13.30	16.20	21.60	30.1	6.0	15.2	7.0	6.00	6.75	11.10	18.93	30.10
92	4.2	36.8	19.7	9.2	18.2	9.20	15.95	18.95	23.98	36.80	9.3	9.1	47.8	19.9	9.10	9.25	14.60	26.88	47.80
93	32.0	37.9	57.9	68.9	46.1	37.90	44.05	52.00	60.65	68.90	71.5	66.7	63.1	83.3	63.10	65.80	69.10	74.45	83.30
94	17.8	20.7	34.4	65.0	39.2	20.70	30.98	36.80	45.65	65.00	41.3	30.5	31.9	49.2	30.50	31.55	36.60	43.28	49.20
95	26.0	100.9	38.3	32.3	48.7	32.30	36.80	43.50	61.75	100.90	65.0	63.7	48.4	33.3	33.30	44.63	56.05	64.03	65.00
96	6.8	5.7	25.9	30.6	16.9	5.70	14.10	21.40	27.08	30.60	10.2	25.9	16.5	5.9	5.90	9.13	13.35	18.85	25.90
97	14.1	17.0	14.3	15.5	54.7	14.30	15.20	16.25	26.43	54.70	23.8	21.9	26.5	31.7	21.90	23.33	25.15	27.80	31.70
98	8.8	12.8	13.5	32.7	10.7	10.70	12.28	13.15	18.30	32.70	24.2	22.0	25.0	22.2	22.00	22.15	23.20	24.40	25.00
99	7.4	28.2	44.7	8.5	14.4	8.50	12.93	21.30	32.33	44.70	36.0	12.5	23.3	16.1	12.50	15.20	19.70	26.48	36.00

Timing Intervals for Each Refresh Function (In Seconds)

Function	DM-1		DM-2						
ID	R-Run 1	R-Run 2	R-Run 3	R-Run 4	Min	25th	Median	75th	Max
DF_CS	91.8	102.4	121.9	127.9	91.80	99.75	112.15	123.40	127.90
DF_I	16.4	19.2	21.2	24.1	16.40	18.50	20.20	21.93	24.10
DF_SS	84.2	99.3	118.7	124.6	84.20	95.53	109.00	120.18	124.60
DF_WS	65.4	75.8	78.7	88.8	65.40	73.20	77.25	81.23	88.80
LF_CR	8.4	8.8	9.1	8.8	8.40	8.70	8.80	8.88	9.10
LF_CS	10.9	10.8	11.1	10.8	10.80	10.80	10.85	10.95	11.10
LF_I	5.5	5.5	6.2	5.9	5.50	5.50	5.70	5.98	6.20
LF_SR	8.1	8.1	8.7	7.8	7.80	8.03	8.10	8.25	8.70
LF_SS	9.3	9.2	9.4	9.1	9.10	9.18	9.25	9.33	9.40
LF_WR	7.7	8.2	8.1	7.9	7.70	7.85	8.00	8.13	8.20
LF_WS	10.3	10.5	11.0	11.4	10.30	10.45	10.75	11.10	11.40

Preface

TPC Benchmark™ DS Overview

The TPC Benchmark™ DS (TPC-DS) is a decision support benchmark that models several generally applicable aspects of a decision support system, including queries and data maintenance. The benchmark provides a representative evaluation of performance as a general purpose decision support system.

This benchmark illustrates decision support systems that:

- Examine large volumes of data;
- Give answers to real-world business questions;
- Execute queries of various operational requirements and complexities (e.g., ad-hoc, reporting, iterative OLAP, data mining);
- Are characterized by high CPU and IO load;
- Are periodically synchronized with source OLTP databases through database maintenance functions.
- Run on “Big Data” solutions, such as RDBMS as well as Hadoop/Spark based systems.

A benchmark result measures query response time in single user mode, query throughput in multi user mode and data maintenance performance for a given hardware, operating system, and data processing system configuration under a controlled, complex, multi-user decision support workload.

The purpose of TPC benchmarks is to provide relevant, objective performance data to industry users. To achieve that purpose, TPC benchmark specifications require benchmark tests be implemented with systems, products, technologies and pricing that:

a) Are generally available to users;
b) Are relevant to the market segment that the individual TPC benchmark models or represents (e.g., TPC-DS models and represents complex, high data volume, decision support environments);
c) Would plausibly be implemented by a significant number of users in the market segment modeled or represented by the benchmark.

In keeping with these requirements, the TPC-DS database must be implemented using commercially available data processing software, and its queries must be executed via SQL interface. 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, which improve benchmark results but not real-world performance or pricing, are prohibited.

TPC benchmark results are expected to be accurate representations of system performance. Therefore, there are specific guidelines that are expected to be followed when measuring those results. The approach or methodology to be used in the measurements are either explicitly described in the specification or left to the discretion of the test sponsor.

When not described in the specification, the methodologies and approaches used must meet the following requirements:

- The approach is an accepted engineering practice or standard;
- The approach does not enhance the result;
- Equipment used in measuring the results is calibrated according to established quality standards;
- Fidelity and candor is maintained in reporting any anomalies in the results, even if not specified in the benchmark requirements.

Further information is available at http://www.tpc.org/

General Items

0.1 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.

0.2 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 but not limited to:

- Database tuning options;
- optimizer/query execution options;
- query processing tool/language configuration parameters;
- recovery/commit options;
- consistency/locking options;
- operating system and configuration parameters;
- configuration parameters and options for any other software component incorporated into the pricing structure; and
- compiler optimization options.

This requirement can be satisfied by providing a full list of all parameters and options, as long as all those which have been modified from their default values have been clearly identified and these parameters and options are only set once.
The Supporting File Archive contains the Operating System and DBMS parameters used in this benchmark.

0.3 Configuration Diagrams

Diagrams of both measured and priced configurations must be provided, accompanied by a description of the differences. This includes the processors, memory, storage, connections and execution locations of software components used in the test or incorporated into the pricing structure.

Measured Configuration

Following is an illustration of the measured configuration.

Figure 0.3: Diagram of the Measured Configuration

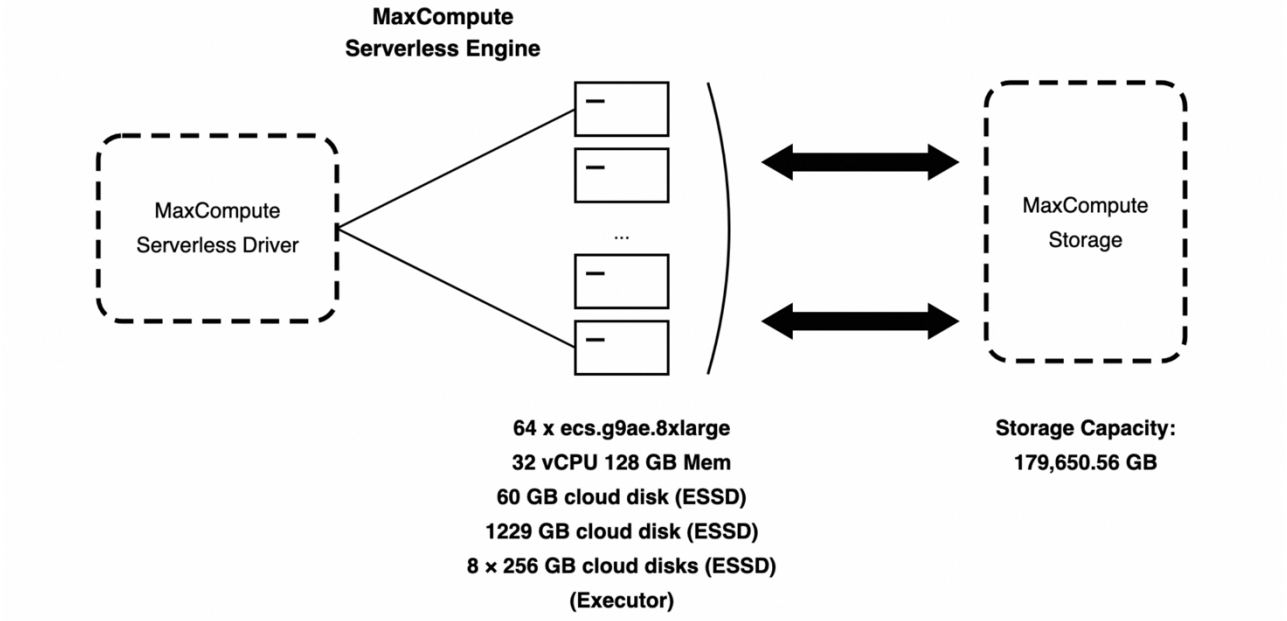

- The external TPC-DS benchmark driver submits the complete executable query text to MaxCompute, receives query output, collects timings, and maintains the audit logs. It does not participate in query execution.
- The product-internal MaxCompute Serverless Driver accepts the requests and dispatches the queries to the Executors.
- The Executors execute queries.
- The MaxCompute Storage stores and manages data.

The Driver is a shared infrastructure component of the MaxCompute Serverless platform and is included as part of the service offering. It is not separately purchased or priced by customers. Otherwise, there are no differences between the priced and measured configurations.

Configuration of the Measured System

The following table shows the configuration of the measured system.

Table 0.3: Configuration of the measured system.

Configuration	Executor
Instance Count	64
Instance Type	ecs.g9ae.8xlarge
vCPU per instance	32
Memory per instance	128 GB
Storage per instance	1 × 60 GB ESSD; 1 × 1,229 GB ESSD; and 8 × 256 GB ESSD (3,337 GB aggregate capacity per instance)
Processors	AMD EPYC 9T95 192-Core Processor
Network	Bandwidth: 16 Gbps; Packet forwarding rate: 5,000,000

Configuration	MaxCompute Storage
Storage Capacity	179, 650.56 GB

Priced Configuration

The Driver is a shared infrastructure component of the MaxCompute Serverless platform and is included as part of the service offering. It is not separately purchased or priced by customers.
The MaxCompute service configuration is identical in the priced and measured configurations. The measured benchmark used one Lenovo P14H laptop as the external TPC-DS benchmark driver. The priced configuration includes three Lenovo P14H laptops: the measured benchmark driver and two spares.

Clause 2: Logical Database Design Related Items

2.1 Database Definition Statements

Listings must be provided for the DDL scripts and must include all table definition statements and all other statements used to set up the test and qualification databases.

The Supporting File Archive contains the table definitions and all other statements used to set up the test and qualification databases.

Supporting Files Index

Clause	Description	Archive File Pathname
clause2	The table definitions and all other statements used to set up the test and qualification databases	support_files_20260805_081850/Clause_3/sqls/

2.2 Physical Organization

The physical organization of tables and indices within the test and qualification databases must be disclosed. If the column ordering of any table is different from that specified in Clause 2.3 or 2.4, it must be noted.

All 24 test database tables are hash clustered and sorted using the CLUSTERED BY and SORTED BY clauses. Nine tables use 1,600 buckets; the remaining fifteen tables use one bucket.

Table Name	Cluster Column	Sort Columns	Buckets
customer_address	ca_address_sk	ca_address_sk	1,600
customer	c_customer_sk	c_customer_sk	1,600
catalog_returns	cr_item_sk	cr_returned_date_sk, cr_returned_time_sk	1,600
catalog_sales	cs_item_sk	cs_sold_date_sk, cs_sold_time_sk	1,600
inventory	inv_item_sk	inv_date_sk, inv_item_sk	1,600
store_returns	sr_item_sk	sr_returned_date_sk, sr_return_time_sk	1,600
store_sales	ss_item_sk	ss_sold_date_sk, ss_sold_time_sk	1,600
web_returns	wr_item_sk	wr_returned_date_sk, wr_returned_time_sk	1,600
web_sales	ws_item_sk	ws_sold_date_sk, ws_sold_time_sk	1,600

The following fifteen tables use one bucket and are clustered and sorted by their respective surrogate key: call_center, catalog_page, customer_demographics, date_dim, household_demographics, income_band, item, promotion, reason, ship_mode, store, time_dim, warehouse, web_page, and web_site.

2.3 Horizontal Partitioning

If any directives to DDLs are used to horizontally partition tables and rows in the test and qualification databases, these directives, DDLs, and other details necessary to replicate the partitioning behavior must be disclosed.

The benchmark DDL does not define PARTITIONED BY clauses. The nine tables identified in Section 2.2 use hash-based bucketing with 1,600 buckets. Rows are distributed among buckets by hashing their respective cluster columns. The remaining fifteen tables use one bucket. The complete directives and DDL required to reproduce this organization are included in the Supporting Files Archive.

2.4 Replication

Any replication of physical objects must be disclosed and must conform to the requirements of Clause 2.5.3.

No replication was used.

Clause 3: Scaling and Database Population

3.1 Initial Cardinality of Tables

The cardinality (e.g., the number of rows) of each table of the test database, as it existed at the completion of the database load (see Clause 7.1.2) must be disclosed.

Table 3.1 lists the cardinality of each table at the completion of the database load.

Table 3.1 Initial Number of Rows

Table Name	Row Count
call_center	60
catalog_page	50,000
catalog_returns	14,403,369,483
catalog_sales	144,071,691,379
customer	100,000,000
customer_address	50,000,000
customer_demographics	1,920,800
date_dim	73,049
household_demographics	7,200
income_band	20
inventory	1,965,337,830
item	502,000
promotion	2,500
reason	75
ship_mode	20
store	1,902
store_returns	28,777,641,316
store_sales	288,018,711,538
time_dim	86,400
warehouse	30
web_page	5,004
web_returns	7,191,510,606
web_sales	71,977,978,409
web_site	96

3.2 Distribution of Tables and Logs Across Media

The source flat files produced by DSDGen were written to a MaxCompute Volume. During the measured load interval, external text tables were created over those files and used as the source for loading the test database.
The 24 test database tables were created as internal MaxCompute tables. The table metadata captured after creation identifies the tables as internal, transactional tables stored in AliORC format on the MaxCompute Standard storage tier. The DDL and captured metadata are included in Clause_3/sqls/create_tables.sql, Clause_3/logs/table_schema.out, and Clause_3/logs/table_schema_after.out.
Benchmark control scripts and the database population and validation logs were written to the benchmark driver file system and are included under Clause_3 in the Supporting Files.
The only storage charged according to usage was (1) the MaxCompute Volume containing the flat files generated directly by DSDGen, including the data-maintenance input files, and (2) the MaxCompute internal-table storage populated from those Volume files. The measured usage of both categories was combined in the Data Storage line item disclosed in the Executive Summary.
Other service-managed data, including shuffle data and data cache, is internal to the MaxCompute execution engine and does not incur a separate storage charge. No separate priced storage component is added for that engine-internal data. Benchmark scripts and logs stored on the benchmark driver are covered by the priced Lenovo P14H laptop.
For the data storage ratio, the disclosed Executor local-storage capacity is included as storage in the priced configuration. Engine-internal transient data does not add another capacity component beyond that disclosed Executor storage. The MaxCompute Volume and internal-table storage are represented by the combined 179, 650.56 GB MaxCompute Storage quantity. Consequently, the ratio calculation in Section 3.7 does not add shuffle data or data-cache contents as separate storage components.

3.3 Mapping of Database Partitions and Replications

The mapping of database partitions/replications must be explicitly described.

The benchmark DDL does not define table partitions, and it does not contain a user-configurable replication directive. All test tables are hash clustered and sorted. Nine tables use 1,600 buckets; the remaining fifteen tables use one bucket. Table 3.2 summarizes the 1,600-bucket tables. The complete DDL is included in Clause_3/sqls/create_tables.sql.

Table 3.2 Hash Clustering for 1,600-Bucket Tables

Table Name	Cluster Column	Sort Columns	Buckets
customer_address	ca_address_sk	ca_address_sk	1,600
customer	c_customer_sk	c_customer_sk	1,600
catalog_returns	cr_item_sk	cr_returned_date_sk, cr_returned_time_sk	1,600
catalog_sales	cs_item_sk	cs_sold_date_sk, cs_sold_time_sk	1,600
inventory	inv_item_sk	inv_date_sk, inv_item_sk	1,600
store_returns	sr_item_sk	sr_returned_date_sk, sr_return_time_sk	1,600
store_sales	ss_item_sk	ss_sold_date_sk, ss_sold_time_sk	1,600
web_returns	wr_item_sk	wr_returned_date_sk, wr_returned_time_sk	1,600
web_sales	ws_item_sk	ws_sold_date_sk, ws_sold_time_sk	1,600

The following tables use one bucket and are clustered and sorted by their respective surrogate key: call_center, catalog_page, customer_demographics, date_dim, household_demographics, income_band, item, promotion, reason, ship_mode, store, time_dim, warehouse, web_page, and web_site.

The benchmark scripts do not map a table, bucket, or replica to a specific physical device or availability zone. Any service-internal data placement or replication is managed by MaxCompute and is not controlled by the benchmark DDL.
The MaxCompute Standard storage tier used by the measured configuration applies the service-default three-replica policy described in Clause 6. Replica creation and placement across different machines and racks are managed internally by MaxCompute. Customers cannot configure or change the replica count or placement rules through benchmark DDL, session settings, or customer-visible storage configuration.
For pricing and capacity disclosure, the Data Storage price and the disclosed 179,650.56 GB MaxCompute Storage quantity represent the billable compressed storage quantity corresponding to one replica. The additional physical capacity used by the two service-managed redundant replicas is included in the MaxCompute service and is not separately billed or exposed as customer-configured capacity. Consequently, redundant-replica overhead is not added again to Total Storage or to the Data Storage Ratio.

3.4 Implementation of RAID

Implementations may use some form of RAID. The RAID level used must be disclosed for each device. If RAID is used in an implementation, the logical intent of its use must be disclosed

No user-configurable RAID level was selected or configured by the benchmark scripts. MaxCompute is a managed service and does not expose RAID configuration through benchmark DDL or session settings. Durability is provided by the service-default three-replica distributed-storage mechanism described in Clause 6, rather than by customer-visible, device-level RAID. Customers cannot configure the replica count or replica placement. The two redundant replicas are included in the service and are not added as separately priced or customer-configured capacity in Total Storage or the Data Storage Ratio.

3.5 Data Generation Tool Modifications

The version number (i.e., the major revision number, the minor revision number, and third tier number) of dsdgen must be disclosed. Any modifications to the dsdgen source code (see Appendix B:) must be disclosed. In the event that a program other than dsdgen was used to populate the database, it must be disclosed in its entirety.

TPC-DS DSDGen version 4.0.0 was used. The version output captured on the test date is included in Clause_3/logs/dsdgen_version.log.
One source patch was applied before the TPC-DS tools were compiled. The patch changes r_params.c as follows:

- PARAM_MAX_LEN is changed from 80 to PATH_MAX.
- The copy of a string-valued parameter is changed from strcpy to strncpy with PARAM_MAX_LEN as the maximum length.

The patch is limited to parameter string and path-length handling. It does not change the table-generation algorithms, data distributions, random-number generation, scale-factor processing, or row-generation logic. The complete patch is included in Clause_3/patches/tools/tpcds-strcpy.diff.
DSDGen was invoked by the Spark wrapper class com.aliyun.odps.spark.TPCDSSaveToVolume2. For the test database, the wrapper invoked DSDGen with scale factor 100,000 and parallel value 100,000 and wrote the generated flat files to a MaxCompute Volume. Data-maintenance input was generated for four streams, and validation data was generated with a validation count of 50. The invocation and orchestration are disclosed in Clause_3/datagen.sh.
The complete wrapper source is disclosed in tools/datagen/src/main/scala/com/aliyun/odps/spark/TPCDSSaveToVolume2.scala. The source creates one Spark partition for each requested DSDGen child, executes the official DSDGen executable with the disclosed command-line parameters, selects the resulting .dat or .vld files, creates the corresponding table directory in the MaxCompute Volume, and copies each generated file to that directory. It does not open, parse, transform, filter, aggregate, or rewrite the contents of any generated data file. Source inspection therefore confirms that the wrapper changes only execution parallelism and output placement; DSDGen remains the program that generates every record.
The wrapper is built by Maven from tools/datagen/pom.xml using mvn package, as invoked by tools/init_tool.sh. The Maven coordinates are com.aliyun.odps:spark-bech_2.12:1.0.0; therefore spark-bech_2.12-1.0.0.jar is the actual artifact name declared by

the source project. The build declares Scala 2.12.10 and Spark 3.3.0 dependencies, with Spark dependencies in provided scope, and uses scala-maven-plugin 3.3.2 and maven-shade-plugin 2.4.3. tools/datagen/dependency_tree.txt records the resolved dependency tree and a successful Maven build. At execution time, load/datagen.sh submits the JAR through the MaxCompute Cupid Spark 3.1.1-odps0.34.1 distribution and supplies the official dsdgen executable and tpcds.idx with --files.
The test-database invocation used 1,600 Spark executors and wrapper arguments --scale 100000 --parallel 100000; the wrapper mapped child numbers 1 through 100,000 to the official DSDGen options -scale 100000 -parallel 100000 -child n. Validation generation used --scale 100000 --vcount 50. The source archive tpc_benchmark.zip contains the wrapper source, POM, dependency tree, build script, and datagen.sh; the wrapper source, POM, and datagen.sh in that archive match the reviewed working-tree files byte for byte.
For data-maintenance generation, the reviewed and auditor-reviewed load/datagen.sh invocation used four Spark executors and wrapper arguments --scale 4 --update 4. The wrapper created four Spark partitions and mapped child numbers 1 through 4 to the official DSDGen options -scale 4 -update n. This is the actual implementation used by the reviewed script and is disclosed without alteration.
The query-template date-expression patch and the MaxCompute dialect template included elsewhere in the Supporting Files are not DSDGen source changes. They are disclosed with the query-generation material under Clause 4.

3.6 Database Load Time

The database load time for the test database (see Clause 7.4.3.7) must be disclosed.

The measured database load time was **1,771.449 seconds**.

3.7 Data Storage Ratio

The data storage ratio must be disclosed. It is computed by dividing the total data storage of the priced configuration (expressed in GB) by SF corresponding to the scale factor chosen for the test database as defined in Clause 3.1. The ratio must be reported to the nearest 1/100th, rounded up. For example, a system configured with 96 disks of 2.1 GB capacity for a 100GB test database has a data storage ratio of 2.02.

The data storage ratio is calculated from the storage disclosed in the current benchmark configuration divided by the scale factor of 100,000.

Table 3.3: Data Storage Ratio Calculation

Component	Quantity (GB)	Evidence
Executor Storage: 64 × 3,337 GB	213,568	Executive Summary benchmarked configuration
MaxCompute Storage	179,650.56	Executive Summary benchmarked configuration
Total storage	393,218.56	213,568 + 179,650.56
Scale factor	100,000	DSDGen command and configuration
Data storage ratio	3.94	393,218.56 / 100,000 = 3.9321856, rounded up to the nearest one-hundredth

The storage ratio reported for this configuration is **3.94**.

3.8 Database Load Mechanism Details and Illustration

The details of the database load must be disclosed, including a block diagram illustrating the overall process. Disclosure of the load procedure includes all steps, scripts, input and configuration files required to completely reproduce the test and qualification databases.

Figure 3.8: Block Diagram of Database Load Process

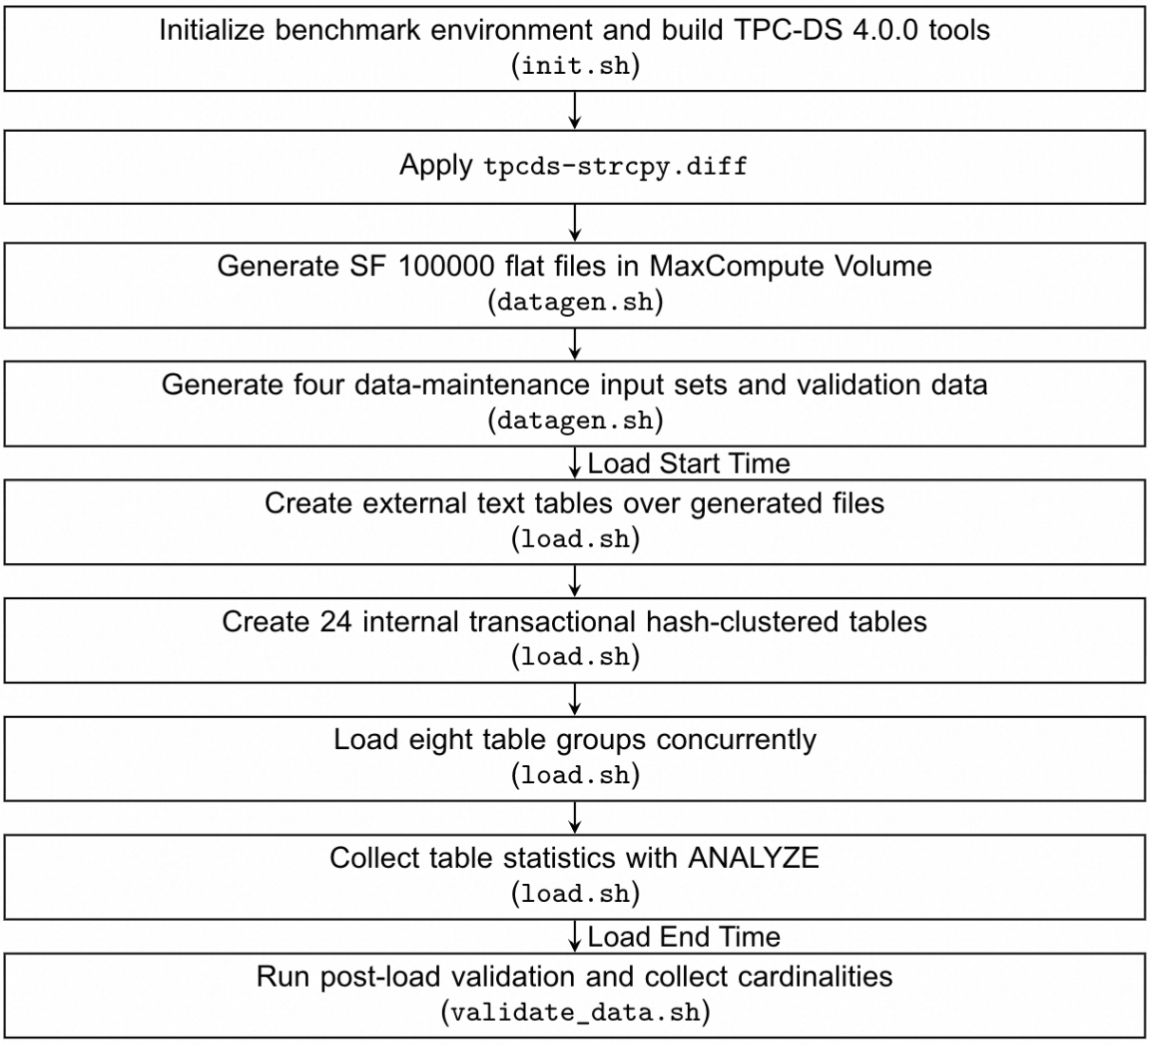

The detailed sequence was as follows:

1. init.sh initialized the benchmark environment. The TPC-DS 4.0.0 tools were built after applying the disclosed tpcds-strcpy.diff patch.
2. datagen.sh used the Spark wrapper to run DSDGen at scale factor 100,000 with parallel value 100,000. Generated flat files were written to a MaxCompute Volume.
3. The same script generated four data-maintenance input sets and validation data. Delete and inventory-delete input files required by the data-maintenance implementation were downloaded to the driver work area.
4. At the load start boundary, load.sh created external text tables over the generated files.
5. load.sh created the 24 internal, transactional, hash-clustered test tables.
6. The script ran the disclosed archive and merge control step and then loaded eight table groups concurrently: store_sales, catalog_sales, web_sales, catalog_returns, inventory, store_returns, web_returns, and the dimension-table group.
7. After each table group was loaded successfully, its associated ANALYZE statement collected table statistics. Completion of all load and statistics tasks established the load end boundary.
8. After the measured load interval, validate_data.sh created validation tables and executed Validate_Data.sql, Check_RI.sql, Check_Insert.sql, and count_tables.sql. The validation log records zero failures for Validate_Data.sql and zero failures for the referential-integrity checks. Check_Insert.sql was executed and its output was retained. The table cardinalities reported in Table 3.1 were then collected.

3.9 Qualification Database Configuration

Any differences between the configuration of the qualification database and the test database must be disclosed.

The qualification database was created in an independent MaxCompute project and schema at scale factor 1. It used the same TPC-DS 4.0.0 DSDGen executable, Spark data-generation wrapper, table DDL, clustering definitions, load SQL, ANALYZE statements, MaxCompute SQL engine, quota selection, and session initialization settings as the test database. The material differences were the scale factor, the independent project and schema, and the separate MaxCompute Volume used for the qualification data.
The qualification queries used the TPC-DS 4.0.0 Appendix B substitution parameters. Query-generation and qualification-result details are disclosed under Clause 4 and in the qualification material in the Supporting Files.

Clauses 4–5: Query, Substitution and Data Maintenance

4.1 Query Language

The query language used to implement the queries must be identified.

SQL was the query language used to implement the queries.

4.2 Verifying Method of Random Number Generation

The method of verification for the random number generation must be described unless the supplied dsdgen and dsqgen were used.

A Spark wrapper based on TPC-supplied DSDGen version 4.0.0 and DSQGen version 4.0.0 was used. In addition, the TPC-supplied tool source was patched to raise the command-line parameter length limit from 80 characters to PATH_MAX and to use a bounded string copy (r_params.c), so that long filesystem paths are accepted. This patch only affects command-line parameter handling; it does not affect data generation, query generation, or random number generation. The patch is included in the Supporting Files Archive.

4.3 Generating Values for Substitution Parameters

The method used to generate values for substitution parameters must be disclosed. The version number of DSQGen must be disclosed.

TPC-supplied DSQGen version 4.0.0 was used to generate the substitution parameters:

```
Executable: ./tools/DSGen-software-code-4.0.0/tools/dsqgen
-directory: ./build/patched_query_templates
-input: ./build/patched_query_templates/templates.lst
-scale: 100000
-streams: 9
-output_dir: ./build/queries
-dialect: maxcompute
-rngseed: $SEED
```

4.4 Query Text and Output Data from Qualification Database

The executable query text used for query validation must be disclosed along with the corresponding output data generated during execution against the qualification database. Any minor modifications applied to functional query definitions or approved variants must be disclosed and justified.

The Supporting Files Archive contains the actual query text and query output.
The following MQM is used:

- Use vendor-specific syntax of date expressions (MQM f.1) for Q5, Q12, Q16, Q20, Q21, Q32, Q37, Q40, Q72, Q77, Q80, Q82, Q92, Q94, Q95, and Q98.

No other minor query modifications were applied.

4.5 Query Substitution Parameters and Seeds Used

All query substitution parameters used during the performance test must be disclosed in tabular format, along with the seeds used to generate these parameters.

The Supporting Files Archive contains the query substitution parameters and seed used in the test.

4.6 Refresh Setting

All query and refresh session initialization parameters, settings and commands must be disclosed.

The Supporting Files Archive contains the query and scripts, along with initialization parameters and settings.

4.7 Source Code of Refresh Functions

The details of how the data maintenance functions were implemented must be disclosed, including source code of any non-commercial program used.

The Supporting Files Archive contains the source code implementing the refresh functions.

4.8 Staging Area

Any object created in the staging area used to implement the data maintenance functions must be disclosed. Any storage used for the staging area must be priced, and any mapping or virtualization of storage must be disclosed.

A staging area was used to implement the data maintenance functions.
For each refresh round from 1 through 4, ten MaxCompute external tables were created over the DSDGen-generated refresh flat files stored in a MaxCompute Volume. These external tables represented the applicable TPC-DS source-schema tables and were used by the Load Fresh data maintenance functions to populate the target TPC-DS tables.
The creation and use of the external tables occurred within the timed data maintenance functions. Their definitions, MaxCompute Volume locations, and the SQL statements that used them are disclosed in Clause_5/mtsqls_1 through Clause_5/mtsqls_4 in the Supporting Files Archive. The staging objects were not used to execute the Power Test or Throughput Tests.
MaxCompute Volume storage is charged under the same storage billing rules as MaxCompute internal-table storage. The storage occupied by the refresh data sets in the Volume was combined with the internal-table storage quantity and included in the priced storage amount.

Clause 6: Data Persistence Properties Related Items

6.1 Method Used to Satisfy the Data Accessibility Requirement

The results of the data accessibility tests must be disclosed along with a description of how the data accessibility requirements were met.

The Data Accessibility requirements were satisfied by Data Accessibility Documentation as permitted by Clause 6.1.2.2 of the TPC-DS Standard Specification, Revision 4.0.0. No induced durable-medium failure was used for this result.
The test database consisted of 24 internal MaxCompute tables. The DDL and post-load metadata in the Supporting Files show that these tables used MaxCompute transactional table format version 2 and the Standard storage tier. The benchmark did not use an OSS external table as a database-resident TPC-DS table. The refresh flat files in the MaxCompute Volume and the external-table definitions over those files formed the Clause 5 staging area; their storage was included in the priced MaxCompute storage quantity.

6.2 Data Accessibility Architecture

MaxCompute stores table data and metadata in its managed storage service. The public MaxCompute Storage Overview identifies internal tables as MaxCompute table data and identifies table schemas, partitioning and clustering specifications, table lifecycles, and permission configurations as MaxCompute metadata. It also states that MaxCompute storage uses three replicas by default.
The public MaxCompute Security White Paper describes the distributed-storage protection used by MaxCompute. It states:
"MaxCompute uses a distributed file system that automatically creates three replicas of stored data."
It further states:
"Add, modify, and delete operations are synchronized across all three replicas."
The same document explains that the replicas are placed on different physical machines and racks, preventing a single chunk-server or rack failure from making the data unavailable. This is the synchronous multi-copy write mechanism used for the database-resident TPC-DS data and the MaxCompute-managed metadata.
Alibaba Cloud's public Pangu architecture document identifies Pangu as a storage core used by MaxCompute. It describes a distributed metadata layer in which multiple MetaServers in each metadata partition use the RAFT protocol for high reliability and availability, and the metadata is persisted in ChunkServer storage. Thus the database metadata is not dependent on a single metadata server or a single durable storage device.
Alibaba Cloud's public description of Pangu data reliability states that Pangu keeps three copies, that losing one copy does not affect reads or writes, and that it performs back-end repair if one copy is not written. The public documentation for Alibaba Cloud's underlying triplicate-storage mechanism describes the repair operation as follows:
"When the system detects insufficient replicas, it selects a healthy storage node ... and copies data from an existing replica to restore the three-replica state."
The same paragraph states that this process is transparent to applications and requires no manual intervention. Together, these mechanisms are the documented basis for maintaining full access after the permanent failure of any one durable storage device containing a copy of a TPC-DS table or its metadata.

6.3 Public Data Accessibility Documentation

The following vendor documentation is publicly available and forms the Data Accessibility Documentation for this result:

1. MaxCompute Security White Paper - Multi-replica data storage - identifies the MaxCompute distributed file system, the default three replicas, placement across different machines and racks, and synchronized add, modify, and delete operations.
2. MaxCompute Storage Overview - identifies the managed MaxCompute table-data and metadata categories and the default three-replica storage configuration.
3. MaxCompute Storage Pricing - identifies the single-zone and multi-zone redundancy configurations and confirms that three-replica redundancy is included while billing is based on one compressed replica.
4. Pangu - The High Performance Distributed File System by Alibaba Cloud - identifies Pangu as a storage core used by MaxCompute and describes redundant, distributed metadata management and metadata persistence in ChunkServer storage.
5. How Table Store Implements High Reliability and High Availability - describes Pangu's three-copy operation, continued read/write access after one copy is lost, and back-end repair when one copy is not written.
6. Alibaba Cloud ESSD Cloud Disk Data Reliability - Fault recovery process - documents automatic reconstruction of an insufficient replica set without manual intervention for the underlying Alibaba Cloud triplicate-storage mechanism.

The preceding documentation applies to the MaxCompute managed storage used and priced for this benchmark. OSS durability percentages and OSS ZRS/LRS service claims are not used as evidence for this MaxCompute result because the test database was stored in internal MaxCompute tables rather than in OSS external tables.

6.4 Supporting Evidence

The Supporting Files Archive contains the following benchmark-specific evidence:

- Clause_3/sqls/create_tables.sql: DDL for all 24 internal transactional TPC-DS tables.

- Clause_3/logs/table_schema.out and Clause_3/logs/table_schema_after.out: captured table metadata, including internal-table status, table format, and Standard storage tier.
- Clause_3/logs/row_count.out: post-load row-count validation.
- Clause_5/mtsqls_1 through Clause_5/mtsqls_4: staging external-table definitions and Data Maintenance SQL.

No EADS were defined for the benchmark database. Hash clustering and sorting described in Clause 2 are physical organization of the internal tables and do not create a separately protected data service.

Clause 7: Performance Metrics and Execution Rules Related Items

7.1 System Activity

Any system activity on the SUT that takes place between the conclusion of the load test and the beginning of the performance test must be fully disclosed including listings of scripts or command logs.

Any system activity on the SUT that takes place between the conclusion of the load test and the beginning of the performance test must be fully disclosed, including listings of scripts or command logs.
The only activity between the end of the load test and the beginning of the performance test was the generation of the executable query text. The related scripts and command logs are included in the Supporting Files.

7.2 Test Steps

The details of the steps followed to implement the performance test must be disclosed.

The Supporting Files Archive contains the scripts and logs used to execute the performance test.

7.3 Timing Intervals for Each Query and Refresh Function

The timing intervals defined in Clause 7 must be disclosed.

The Supporting Files Archive contains the scripts, commands, and logs used to execute the Power Test, Throughput Tests, and Data Maintenance functions.

7.4 Throughput Test Result

For each Throughput Test, the minimum, the 25th percentile, the median, the 75th percentile, and the maximum times for each query shall be reported.

For each Throughput Test, the minimum, the 25th percentile, the median, the 75th percentile, and the maximum times for each query shall be reported.
See the Executive Summary at the beginning of this report.

7.5 Time for Each Stream

The start time and finish time for each query stream must be reported.

See the Executive Summary at the beginning of this report.

7.6 Time for Each Refresh Function

The start time and finish time for each data maintenance function in the refresh run must be reported for the Throughput Tests.

See the Executive Summary at the beginning of this report. Detailed Data Maintenance timing logs are included in the Supporting Files.

7.7 Performance Metrics

The computed performance metric, related numerical quantities and the price/performance metric must be reported.

QphDS@100000 = 133,518,293
See the Executive Summary at the beginning of this report for more detail.

Clause 8: SUT and Driver Implementation Related Items

8.1 Driver

A detailed textual description of how the driver performs its functions, how its various components interact and any product functionalities or environmental settings on which it relies must be provided. All related source code, scripts and configuration files must be disclosed. The information provided should be sufficient for an independent reconstruction of the driver.

The benchmark driver uses shell scripts for test sequencing and a Java benchmark client for Power and Throughput query execution. The Java client connects to Alibaba Cloud MaxCompute Query Accelerator 2.0 v5.57 through Alibaba Cloud MaxCompute JDBC driver version 3.9.4. It ran on Alibaba Cloud Linux 3 (Soaring Falcon) with OpenJDK Runtime Environment (Alibaba 8.17.21), build 1.8.0_345-b718.
Each query stream is executed by one Java process using one JDBC connection. The 99 queries within a stream are submitted sequentially in their generated order, and no query text is modified by the driver. The Power Test uses stream 0. The first Throughput Test runs streams 1 through 4 concurrently, and the second Throughput Test runs streams 5 through 8 concurrently. Data maintenance functions are submitted through the MaxCompute command-line client. The driver records query and data maintenance start times, finish times, elapsed times, returned query results, and execution logs. The benchmark kit and Supporting Files Archive contain the driver source code, build and configuration files, execution scripts, generated query streams, timing logs, and result files.

8.2 Implementation Specific Layer (ISL)

If an implementation specific layer is used, then a detailed description of how it performs its functions, how its various components interact and any product functionalities or environmental setting on which it relies must be provided. All related source code, scripts and configuration files must be disclosed. The information provided should be sufficient for an independent reconstruction of the implementation specific layer.

No implementation-specific layer was used.

8.3 Profile-Directed Optimization

If profile-directed optimization as described in Clause 7.2.10 is used, such use must be disclosed. In particular, the procedure and any scripts used to perform the optimization must be disclosed.

Profile-directed optimization was not used.

Clause 9: Pricing Related Items

9.1 Hardware and Software Used

A detailed list of hardware and software used in the priced system must be reported. The rules for pricing are included in the current revision of the TPC Pricing Specification located on the TPC website (https://www.tpc.org/).

The Data Storage line item shown in the Executive Summary is priced for the China East 2 region. The pricing source reports a displayed monthly value of CNY 21,857.48 and a 12-month extended price of CNY 262,289.76. The extended price is used in the priced configuration and Total System Cost; the displayed monthly value is rounded.

9.2 Availability Date

The System Availability Date must be the single availability date reported on the first page of the Executive Summary. The Full Disclosure Report must report Availability Dates individually for at least each category for which a pricing subtotal is reported. All required Availability Dates must be disclosed to a precision of one day.

The total system has been available since **2026-08-04**.

9.3 Country-Specific Pricing

Additional Clause 7 related items may be included in the full disclosure report for each country specific priced configuration.

The configuration is priced for the **Chinese market** in **CNY**, using the Alibaba Cloud **China East 2** region.

Clause 11: Audit Related Items

Auditor’s Information and Attestation Letter

The auditor's agency name, address, phone number, and attestation letter with a brief audit summary report indicating compliance must be included in the full disclosure report. A statement should be included specifying whom to contact in order to obtain further information regarding the audit process.

This benchmark was audited by: Doug Johnson, of InfoSizing.

InfoSizing
The Right Metric For Sizing IT

TPC
Certified Auditor

Xiaowen Zheng
Development Engineer
MaxCompute Team
Alibaba Cloud Computing Ltd.
Cloud Valley, No. 1008 Dengcai Street
Xihu District, Hangzhou, China

August 17, 2026

I verified the TPC Benchmark™ DS v4.0.0 performance of the following configuration:

Platform:	Alibaba Cloud MaxCompute with 1,600 MaxCompute Query Accelerator CUs
Operating System:	Alibaba Cloud Linux 3 (Soaring Falcon)
Database Manager:	Alibaba Cloud MaxCompute Query Accelerator 2.0 v5.57

The results were:

Performance Metric	**133,518,293**	**QphDS@100000**
Secondary Metrics		
T_{Load}	1,771.5	seconds@100000
T_{Power}	1,341.6	seconds@100000
T_{TT1}	2,604.5	seconds@100000
T_{TT2}	3,133.8	seconds@100000
T_{DM1}	268.9	seconds@100000
T_{DM2}	325.9	seconds@100000

System Under Test	**64x ecs.g9ae.8xlarge Executor Pods, each with:**		
CPUs	32 vCPUs (2,048 total) AMD EPYC 9T95 192-Core Processor		
Memory	128 GB (8,192 GB total)		
Storage	***Qty***	***Size***	***Type***
	1	60 GB	Cloud Disk (ESSD)
	1	1,229 GB	Cloud Disk (ESSD)
	8	256 GB	Cloud Disk (ESSD)
MaxCompute Storage	1	175.44 TB	Shared by all pods

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

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

The following verification items were given special attention:

- The database records were defined with the proper layout and size.
- The database population was generated using DSDGen v4.0.0
- The database was properly scaled to 100,000 GB and populated accordingly.
- The primary and secondary metrics were correctly measured and reported.
- The query templates were produced using approved minor query modifications.
- The query substitution parameters were generated using DSQGen v4.0.0.
- The execution of the queries against the qualification database produced compliant output.
- The tests were driven and sequenced according to the requirements.
- Each throughput test comprised 4 query streams.
- The execution times for queries and data maintenance functions were correctly measured and reported.
- The data accessibility test was satisfied via documentation.
- 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,

Doug Johnson, Certified TPC Auditor

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

Supporting Files Index

Clause	Description	Archive File Pathname
Clause 3	Database create and load scripts, SQL scripts for table creation and validation	SupportingFiles/Clause_3/
	The code for the Spark wrapper of dsdgen	SupportingFiles/Clause_3/datagen
	Patches for data generation tools	SupportingFiles/Clause_3/patches/tools/
Clause 4	The script to execute qualification test	SupportingFiles/Clause_4/
	Patches for query templates	SupportingFiles/Clause_4/patches/query_templates/
	SQL for qualification queries	SupportingFiles/Clause_4/queries/
	Output from executing qualification queries	SupportingFiles/Clause_4/output/
Clause 5	Data maintenance execution scripts and logs files	SupportingFiles/Clause_5/
	SQL scripts for DM functions for stream [s]	SupportingFiles/Clause_5/mtsqls_[s]/
	Data file with delete dates	SupportingFiles/Clause_5/delete/
		SupportingFiles/Clause_5/inventory_delete/
Clause 6	Data accessibility documentation	SupportingFiles/Clause_6/DA_Documentation.docx
Clause 7	Performance test scripts and logs	SupportingFiles/Clause_7/
	Query text for query [q] in stream [s]	SupportingFiles/Clause_7/stream_[s]_queries/query_[q].sql
	Output of query in stream [s], [exec #] is the execution order number. [Q #] is the query template number.	SupportingFiles/Clause_7/stream_[s]_results/query_[exec #]_[Q #].out
Clause 8	MaxCompute Query Accelerator Configuration Inventory	SupportingFiles/Clause_8/

Appendix A: Purchase Page of Creating Alibaba Cloud MaxCompute with 1-Year Subscription

Appendix B: Aliyun MaxCompute Resource Utilization

Appendix C: Third Party Price Quotes