



Alibaba Cloud Computing Ltd.

TPC Benchmark™ DS

Full Disclosure Report

for

Alibaba Cloud E-MapReduce

(with 19 Alibaba Cloud Elastic Compute Service Servers)

using

E-MapReduce 3.21.2

and

CentOS Linux Release 7.4

First Edition

September 16, 2019

First Edition – September, 2019

Alibaba Cloud and the Alibaba Cloud Logo are trademarks of Alibaba Group and/or its affiliates in the U.S. and other countries.

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

Table of Contents

Abstract	5
Preface	11
TPC Benchmark™ DS Overview	11
General Items	12
0.1 Test Sponsor	12
0.2 Parameter Settings	12
0.3 Configuration Diagrams	12
Clause 2: Logical Database Design Related Items	15
2.1 Database Definition Statements	15
2.2 Physical Organization	15
2.3 Horizontal Partitioning	15
2.4 Replication	15
Clause 3: Scaling and Database Population	16
3.1 Initial Cardinality of Tables	16
3.2 Distribution of Tables and Logs Across Media	17
3.3 Mapping of Database Partitions/Replications	17
3.4 Implementation of RAID	18
3.5 DBGEN Modifications	18
3.6 Database Load time	18
3.7 Data Storage Ratio	18
3.8 Database Load Mechanism Details and Illustration	18
3.9 Qualification Database Configuration	19
Clause 4 and 5: Query and Data Maintenance Related Items	20
4.1 Query Language	20
4.2 Verifying Method of Random Number Generation	20
4.3 Generating Values for Substitution Parameters	20
4.4 Query Text and Output Data from Qualification Database	20
4.5 Query Substitution Parameters and Seeds Used	21
4.6 Refresh Setting	21
4.7 Source Code of Refresh Functions	21
4.8 Staging Area	21
Clause 6: Data Persistence Properties Related Items	22
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

	4
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
Auditor's Information and Attestation Letter	26
Supporting Files Index	28
Appendix A: Purchase Page of Creating Alibaba Cloud E-MapReduce Cluster with 3-Year Subscription	29
Appendix B: Third Party Price Quotes	30

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 2.11.0.



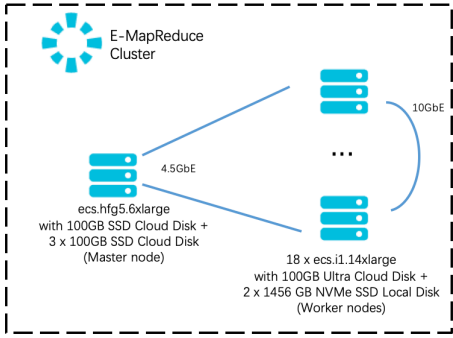
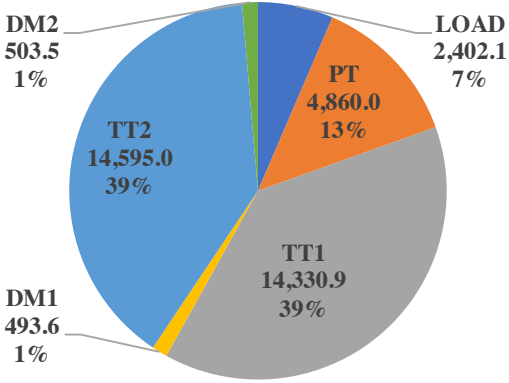
The test was conducted at a Scale Factor of 10000GB with 19 Alibaba Cloud Elastic Compute Service Servers running E-MapReduce 3.21.2 on CentOS Linux Release 7.4.


Measured Configuration


Company Name	Cluster Node	Database Software	Operation System
Alibaba Cloud Computing Ltd.	Alibaba Cloud Elastic Compute Service Server	Alibaba Cloud E-MapReduce 3.21.2	CentOS Linux Release 7.4

TPC Benchmark™ DS Metrics

Total System Cost (CNY)	TPC-DS Throughput (QphDS@10000GB)	Price/Performance (CNY / QphDS@10000GB)	Availability Date
¥ 2,742,339.36	5,261,414	¥ 0.53	As of Publication

 Alibaba Cloud	Alibaba Cloud E-MapReduce		TPC-DS: 2.11.0 TPC-Pricing: 2.4.0 Report Date: Sep. 16, 2019	
Total System Cost	TPC-DS Throughput	Price / Performance	System Availability Date	
¥2,742,339.36 CNY	5,261,414 QphDS@10000GB	¥0.53 CNY/QphDS@10000GB	As of Publication	
Dataset Size ¹	Database Manager	Operation System	Other Software	Cluster
10,000 GB	E-MapReduce 3.21.2	CentOS Linux Release 7.4	N/A	Yes
  Benchmarked Configuration		 Elapsed Time		
Load includes backup = No		RAID = RAID-1 for metadata; HDFS with 3-way replication for table data		
System Configuration:		Alibaba Cloud E-MapReduce Cluster		
Servers:		1 x ecs.hfg5.6xlarge + 18 x ecs.i1.14xlarge		
Total Processors/Cores/Threads:		19/516/1,032		
Total Memory:		4,128 GB		
Total Storage:		54,616 GB		
Storage Ratio:		5.46		
Server Configuration:		Per node (ecs.hfg5.6xlarge)		
Processors:		Intel(R)Xeon(R) Gold 6149 CPU @ 3.10GHz, 22 MB L3		
Memory:		96 GB		
Network:		Bandwidth: 4.5 Gbps, Packet forwarding rate: 2,000,000		
Storage Device:		3 x 100 GB SSD Cloud Disk (data disk) 1 x 100 GB SSD Cloud Disk (boot disk)		
Server Configuration:		Per node (ecs.i1.14xlarge)		
Processors:		Intel(R)Xeon(R) CPU E5-2682 v4 @ 2.50GHz, 40 MB L3		
Memory:		224 GB		
Network:		Bandwidth: 10.0 Gbps, Packet forwarding rate: 1,200,000		
Storage Device:		2 x 1456 GB NVMe SSD Local Disk (data disk) 1 x 100 GB Ultra Cloud Disk (boot disk)		
<p>1. Dataset Size includes only raw data (i.e., no temp, index, redundant storage space, etc.).</p> <p>2. Total Storage = (100 + 100 * 3) (Master node) + (100 + 1,456 * 2) * 18 (Worker nodes) = 54,616 GB</p> <p>3. Storage Ratio = Total Storage / SF = 54,616 GB / 10,000 GB</p>				

		Alibaba Cloud E-MapReduce			TPC-DS: 2.11.0 TPC-Pricing: 2.4.0 Report Date: Sep. 16, 2019	
Description	Part Number	Src	Unit Price (CNY)	Qty	Ext. Price (CNY)	3-Year Maint. (CNY)
Licensed Compute Services						
<u>Virtual cloud server</u>						
ECS Instance ecs.hfg5.6xlarge	ecs.hfg5.6xlarge (China North 2)	1	1,952.16	36	70,277.76	included
ECS System Disk (SSD Cloud Disk 100GB)	Option	1	50.00	36	1,800.00	included
ECS Data Disk (SSD Cloud Disk 100GB)	Option	1	50.00	108	5,400.00	included
<u>Virtual cloud server</u>						
ECS Instance ecs.i1.14xlarge	ecs.i1.14xlarge (China North 2)	1	4,088.70	648	2,649,477.60	included
- NVMe SSD Local Disk (2 x 1456 GB)	Included					
ECS System Disk (Ultra Cloud Disk 100GB)	Option	1	17.50	648	11,340.00	included
Licensed Compute Services Sub-Total					2,738,295.36	0.00
Licensed Software Services						
E-MapReduce for emr.hfg5.6xlarge	included	1		36		
E-MapReduce for emr.i1.14xlarge	included	1		648		
Licensed Software Services Sub-Total					0.00	0.00
Other Components						
Lenovo MIIX 210 Laptop (Includes spares)		2	1,348.00	3	4,044.00	
Other Components Sub-Total					4,044.00	0.00
1 = Alibaba Cloud, 2 = Tmall.com					3-Year Cost of Ownership: 2,742,339.36	
All Licensed Services prices are per month and based on 3-year pre-paid subscriptions.					QphDS@10000GB: 5,261,414	
Audited by Francois Raab, InfoSizing					¥/QphDS@10000GB: 0.53	
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 E-MapReduce		TPC-DS: 2.11.0 TPC-Pricing: 2.4.0 Report Date: Sep. 16, 2019	
Metrics Details:					
Name	Value	Unit			
Scale Factor (SF)	10,000	GB			
Streams	4	Stream			
Queries (Q)	396	Queries			
T_load	2,402.1	Second			
T_ld	0.0267	Hour			
T_power	4,860.0	Second			
T_pt	5.4000	Hour			
T_tt1	14,330.9	Second			
T_tt2	14,595.0	Second			
T_dm1	493.6	Second			
T_dm2	503.5	Second			
T_tt	8.0350	Hour			
T_dm	0.2770	Hour			
Load Step	Start	End	(sec.)	(hh:mm:ss)	
Build	08/22/19 10:37:57.08	08/22/19 11:17:59.18	2,402.09	0:40:02	
Audit	08/22/19 11:17:59.18	08/22/19 11:50:26.69	1,947.51	0:32:28	
Finish	08/22/19 11:50:26.69	08/22/19 11:50:26.69	0.00	0:00:00	
Reported	08/22/19 10:37:57.08	08/22/19 11:50:26.69	2,402.09	0:40:02	
Test	Start	End	(sec.)	(hh:mm:ss)	
Power	08/22/19 12:11:17.51	08/22/19 13:32:17.50	4,859.98	1:21:00	
Thruput-1	08/22/19 13:32:17.52	08/22/19 17:31:08.34	14,330.83	3:58:51	
Thruput-2	08/22/19 17:39:21.89	08/22/19 21:42:36.87	14,594.98	4:03:15	
DM-1	08/22/19 17:31:08.36	08/22/19 17:39:21.87	493.51	0:08:14	
DM-2	08/22/19 21:42:36.89	08/22/19 21:51:00.32	503.43	0:08:23	
Stream	Start	End	(sec.)	(hh:mm:ss)	
Pt - 0	08/22/19 12:11:17.51	08/22/19 13:32:17.50	4,859.98	1:21:00	
Tt1 - 1	08/22/19 13:32:17.52	08/22/19 17:23:17.69	13,860.17	3:51:00	
Tt1 - 2	08/22/19 13:32:17.52	08/22/19 17:27:28.24	14,110.73	3:55:11	
Tt1 - 3	08/22/19 13:32:17.52	08/22/19 17:29:13.15	14,215.63	3:56:56	
Tt1 - 4	08/22/19 13:32:17.52	08/22/19 17:31:08.34	14,330.83	3:58:51	
Tt2 - 5	08/22/19 17:39:21.89	08/22/19 21:40:08.53	14,446.64	4:00:47	
Tt2 - 6	08/22/19 17:39:21.89	08/22/19 21:41:24.20	14,522.31	4:02:02	
Tt2 - 7	08/22/19 17:39:21.89	08/22/19 21:38:18.52	14,336.63	3:58:57	
Tt2 - 8	08/22/19 17:39:21.89	08/22/19 21:42:36.87	14,594.98	4:03:15	
DMt1 - 1	08/22/19 17:31:08.36	08/22/19 17:35:32.87	264.51	0:04:25	
DMt1 - 2	08/22/19 17:35:32.88	08/22/19 17:39:21.87	228.99	0:03:49	
DMt2 - 3	08/22/19 21:42:36.89	08/22/19 21:46:38.52	241.63	0:04:02	
DMt2 - 4	08/22/19 21:46:38.52	08/22/19 21:51:00.32	261.79	0:04:22	

Timing Intervals for Each Query (In Seconds)

Query	Stream 0	Stream 1	Stream 2	Stream 3	Stream 4	Min	25%tile	Median	75%tile	Max	Stream 5	Stream 6	Stream 7	Stream 8	Min	25%tile	Median	75%tile	Max
1	14.3	72.1	34.2	51.7	38.9	34.2	37.7	45.3	56.8	72.1	61.2	65.0	135.2	61.3	61.2	61.3	63.2	82.6	135.2
2	36.2	356.0	76.3	91.5	157.9	76.3	87.7	124.7	207.4	356.0	122.7	100.1	81.0	62.5	62.5	76.4	90.6	105.8	122.7
3	4.0	8.7	23.2	30.8	27.7	8.7	19.6	25.5	28.5	30.8	5.6	8.0	8.7	36.4	5.6	7.4	8.4	15.6	36.4
4	331.2	385.0	404.7	462.2	371.4	371.4	381.6	394.9	419.1	462.2	430.6	382.3	412.3	364.4	364.4	377.8	397.3	416.9	430.6
5	46.1	175.6	119.5	89.1	96.6	89.1	94.7	108.1	133.5	175.6	435.3	128.4	439.6	125.2	125.2	127.6	281.9	436.4	439.6
6	9.6	105.1	80.8	362.5	35.6	35.6	69.5	93.0	169.5	362.5	93.3	21.7	132.2	52.2	21.7	44.6	72.8	103.0	132.2
7	18.1	23.3	54.4	31.7	18.4	18.4	22.1	27.5	37.4	54.4	16.4	42.7	19.6	78.4	16.4	18.8	31.2	51.6	78.4
8	11.9	146.5	127.5	67.5	54.4	54.4	64.2	97.5	132.3	146.5	17.2	61.4	165.3	51.7	17.2	43.1	56.6	87.4	165.3
9	23.4	49.7	32.0	42.7	24.7	24.7	30.2	37.4	44.5	49.7	36.6	52.9	51.9	57.8	36.6	48.1	52.4	54.1	57.8
10	16.5	25.9	40.0	83.9	35.4	25.9	33.0	37.7	51.0	83.9	47.9	44.8	47.1	104.9	44.8	46.5	47.5	62.2	104.9
11	68.8	97.0	159.2	112.9	125.2	97.0	108.9	119.1	133.7	159.2	231.7	93.2	122.5	137.6	93.2	115.2	130.1	161.1	231.7
12	4.0	128.0	77.4	10.0	34.2	10.0	28.2	55.8	90.1	128.0	183.6	19.6	26.4	39.0	19.6	24.7	32.7	75.2	183.6
13	17.4	74.3	55.9	39.2	102.9	39.2	51.7	65.1	81.5	102.9	197.8	30.0	71.5	44.4	30.0	40.8	58.0	103.1	197.8
14	433.4	1,099.1	1,066.1	1,222.6	985.1	985.1	1,045.9	1,082.6	1,130.0	1,222.6	892.7	1,124.9	869.5	1,095.0	869.5	886.9	993.9	1,102.5	1,124.9
15	13.1	22.0	72.9	40.0	49.0	22.0	35.5	44.5	55.0	72.9	101.4	148.5	67.9	13.9	13.9	54.4	84.7	113.2	148.5
16	42.1	84.3	225.1	161.6	132.2	84.3	120.2	146.9	177.5	225.1	157.3	62.5	68.6	98.8	62.5	67.1	83.7	113.4	157.3
17	22.3	185.7	174.3	252.5	78.1	78.1	150.3	180.0	202.4	252.5	86.4	63.5	159.4	32.5	32.5	55.8	75.0	104.7	159.4
18	29.8	69.4	60.5	70.2	40.1	40.1	55.4	65.0	69.6	70.2	301.3	96.0	130.9	65.7	65.7	88.4	113.5	173.5	301.3
19	10.6	46.8	62.3	50.5	42.6	42.6	45.8	48.7	53.5	62.3	55.8	35.9	84.2	53.6	35.9	49.2	54.7	62.9	84.2
20	4.0	6.0	11.3	8.9	112.7	6.0	8.2	10.1	36.7	112.7	17.1	6.6	27.5	376.2	6.6	14.5	22.3	114.7	376.2
21	2.6	119.7	34.6	38.2	6.2	6.2	27.5	36.4	58.6	119.7	13.7	18.3	83.9	20.7	13.7	17.2	19.5	36.5	83.9
22	16.2	42.7	106.6	42.3	58.0	42.3	42.6	50.4	70.2	106.6	60.4	50.0	75.8	41.6	41.6	47.9	55.2	64.3	75.8
23	663.5	1,152.7	1,441.4	1,470.8	1,295.7	1,152.7	1,260.0	1,368.6	1,448.8	1,470.8	1,167.6	1,333.5	1,205.1	1,282.2	1,167.6	1,195.7	1,243.7	1,295.0	1,333.5
24	228.5	920.8	528.5	658.8	1,198.8	528.5	626.2	789.8	990.3	1,198.8	794.5	779.9	690.3	633.3	633.3	676.1	735.1	783.6	794.5
25	29.4	195.3	55.6	70.5	36.7	36.7	50.9	63.1	101.7	195.3	28.1	43.8	52.7	76.7	28.1	39.9	48.3	58.7	76.7
26	7.5	155.2	117.5	57.4	30.5	30.5	50.7	87.5	126.9	155.2	13.1	22.1	24.6	161.1	13.1	19.9	23.4	58.7	161.1
27	9.0	110.1	67.5	246.9	36.4	36.4	59.7	88.8	144.3	246.9	28.5	55.0	19.4	13.5	13.5	17.9	24.0	35.1	55.0
28	54.7	115.8	130.9	94.3	92.9	92.9	94.0	105.1	119.6	130.9	87.8	73.2	250.7	301.5	73.2	84.2	169.3	263.4	301.5
29	61.7	103.2	70.5	77.9	101.9	70.5	76.1	89.9	102.2	103.2	75.1	150.8	150.1	427.6	75.1	131.4	150.5	220.0	427.6
30	21.7	52.7	72.2	182.4	133.1	52.7	67.3	102.7	145.4	182.4	74.5	54.0	89.8	66.6	54.0	63.5	70.6	78.3	89.8
31	39.0	88.0	88.9	256.3	362.7	88.0	88.7	172.6	282.9	362.7	48.9	99.8	98.7	39.2	39.2	46.5	73.8	99.0	362.7
32	10.0	26.6	110.8	13.6	44.6	13.6	23.4	35.6	61.2	110.8	78.4	33.0	20.0	84.4	20.0	29.8	55.7	79.9	84.4
33	13.3	38.9	122.9	32.5	31.8	31.8	32.3	35.7	59.9	122.9	125.1	100.0	29.8	7.1	7.1	24.1	64.9	106.3	125.1
34	14.0	64.5	270.4	103.4	28.3	28.3	55.5	84.0	145.2	270.4	168.6	40.1	48.1	95.6	40.1	46.1	71.9	113.9	168.6
35	46.1	182.0	192.2	115.3	277.0	115.3	165.3	187.1	213.4	277.0	434.2	114.4	127.5	109.4	109.4	113.2	121.0	204.2	434.2
36	8.9	34.4	301.0	98.5	70.6	34.4	61.6	84.6	149.1	301.0	21.7	24.4	19.6	138.4	19.6	21.2	23.1	52.9	138.4
37	28.3	281.0	69.1	83.3	305.8	69.1	79.8	182.2	287.2	305.8	54.8	56.3	80.8	129.5	54.8	55.9	68.6	93.0	129.5
38	116.4	383.7	253.7	311.4	395.0	253.7	297.0	347.6	386.5	395.0	168.4	347.3	256.6	219.1	168.4	206.4	237.9	279.3	347.3
39	24.8	122.6	106.5	106.9	85.7	85.7	101.3	106.7	110.8	122.6	56.9	78.4	183.7	46.2	46.2	54.2	67.7	104.7	183.7
40	19.1	29.6	27.2	381.0	100.6	27.2	29.0	65.1	170.7	381.0	23.6	44.3	199.2	55.1	23.6	39.1	49.7	91.1	199.2
41	1.3	31.2	40.4	8.8	19.1	8.8	16.5	25.2	33.5	40.4	18.7	4.9	1.9	4.1	1.9	3.6	4.5	8.4	18.7
42	2.4	9.4	43.4	11.2	14.4	9.4	10.8	12.8	21.7	43.4	133.1	56.7	8.2	3.9	3.9	7.1	32.5	75.8	133.1
43	7.7	25.2	14.7	75.2	52.1	14.7	22.6	38.7	57.9	75.2	9.8	31.4	12.8	28.2	9.8	12.1	20.5	29.0	31.4
44	6.9	55.2	13.2	27.8	234.1	13.2	24.2	41.5	99.9	234.1	52.6	359.7	220.3	105.9	52.6	92.6	163.1	255.2	359.7
45	10.0	36.2	70.9	12.7	33.2	12.7	28.1	34.7	44.9	70.9	36.1	39.4	28.1	77.9	28.1	34.1	37.8	49.0	77.9
46	16.6	22.0	126.1	147.2	123.5	22.0	98.1	124.8	131.4	147.2	42.6	135.8	242.0	66.4	42.6	60.5	101.1	162.4	242.0
47	26.2	80.6	76.2	37.5	38.8	37.5	38.5	57.5	77.3	80.6	113.0	161.6	44.0	410.0	44.0	95.8	137.3	223.7	410.0
48	18.6	54.6	23.0	28.0	234.7	23.0	26.8	41.3	99.6	234.7	62.5	79.6	75.4	25.3	25.3	53.2	69.0	76.5	79.6
49	18.4	45.2	47.2	21.2	37.0	21.2	33.1	41.1	45.7	47.2	100.2	139.9	58.1	103.3	58.1	89.7	101.8	112.5	139.9
50	152.3	297.5	265.6	262.1	182.3	182.3	242.2	263.9	273.6	297.5	280.6	250.6	217.5	229.1	217.5	226.2	239.9	258.1	280.6
51	31.3	263.7	133.9	224.5	66.5	66.5	117.1	179.2	234.3	263.7	57.0	187.6	48.6	169.2	48.6	54.9	113.1	173.8	187.6
52	2.5	3.0	5.2	8.5	6.0	3.0	4.7	5.6	6.6	8.5	251.3	3.8	14.3	5.3	3.8	4.9	9.8	73.6	251.3
53	5.0	12.8	26.0	143.3	5.2	5.2	10.9	19.4	55.3	143.3	25.7	54.4	15.4	110.0	15.4	23.1	40.1	68.3	110.0
54	15.0	98.9	172.6	40.8	43.9	40.8	43.1	71.4	117.3	172.6	105.8	276.3	77.2	249.9	77.2	98.7	177.9	256.5	276.3
55	2.5	4.4	5.5	27.6	91.7	4.4	5.2	16.6	43.6	91.7	5.7	4.0	174.5	6.3	4.0	5.3	6.0	48.4	174.5
56	14.1	25.6	31.5	63.3	26.5	26.5	26.3	29.0	39.5	63.3	10.2	164.1	12.7	44.0	10.2	12.1	28.4	74.0	164.1
57	16.2	31.2	24.4	18.8	38.8	18.8	23.0	27.8	33.1	38.8	70.1	60.7	333.2	34.5	34.5	54.2	65.4	135.9	333.2
58	25.6	70.2	167.4	72.9	177.5	70.2	72.2	120.2	169.9	177.5	48.3	93.0	65.8	43.7	43.7	47.2	57.1	72.6	93.0
59	112.4	149.7	65.3	138.1	233.6	65.3	119.9	143.9	170.7	233.6	46.9	55.3	43.6	337.9	43.6	46.1	51.1	126.0	337.9
60	13.7	140.1	76.5	152.4	60.7	60.7	72.6	108.3	143.2	152.4	150.9	17.1	52.0	21.5	17.1	20.4	36.8	76.7	150.9
61	17.7	40.4	387.6	45.7	28.0	28.0	37.3	43.1	131.2	387.6	39.9	47.0	47.3	94.5	39.9	45.2	47.2	59.1	94.5
62	21.4	17.8	15.3	88.7	156.9	15.3	17.2	53.3	105.8	156.9	32.0	42.2	152.3	362.4	32.0	122.2	257.4	377.4	422.4
63	4.8	39.7	144.1	10.7	59.5	10.7	32.5	49.6	80.7	144.1	197.4	25.0	16.9	5.4	5.4	14.0	21.0	68.1	197.4
64	204.3	674.9	478.8	594.9	515.9	478.8	506.6	555.4	614.9	674.9	409.7	574.9	338.4	352.8	338.4</				

95	70.8	90.1	125.7	110.2	83.6	83.6	88.5	100.2	114.1	125.7	136.1	204.0	128.0	222.5	128.0	134.1	170.1	208.6	222.5
96	36.0	22.3	35.0	95.2	17.8	17.8	21.2	28.7	50.1	95.2	43.6	89.4	13.2	19.4	13.2	17.9	31.5	55.1	89.4
97	62.7	298.8	207.7	295.6	106.2	106.2	182.3	251.7	296.4	298.8	105.8	210.8	225.7	72.1	72.1	97.4	158.3	214.5	225.7
98	5.9	27.5	18.4	13.5	20.4	13.5	17.2	19.4	22.2	27.5	9.1	51.3	53.4	9.9	9.1	9.7	30.6	51.8	53.4
99	25.1	42.1	98.4	43.9	87.4	42.1	43.5	65.7	90.2	98.4	447.9	50.8	55.9	353.9	50.8	54.6	204.9	377.4	447.9

Timing Intervals for Each Refresh Function (In Seconds)

DM Fx	R-Run 1	R-Run 2	R-Run 3	R-Run 4	Min	25%tile	Median	75%tile	Max
LF_CR	48.1	48.4	48.4	44.6	44.6	47.2	48.3	48.4	48.4
LF_CS	77.4	73.4	86.1	77.5	73.4	76.4	77.5	79.7	86.1
LF_I	31.6	29.6	31.6	31.8	29.6	31.1	31.6	31.7	31.8
LF_SR	44.8	45.1	44.5	48.0	44.5	44.7	45.0	45.8	48.0
LF_SS	75.9	76.4	84.0	71.5	71.5	74.8	76.2	78.3	84.0
LF_WR	39.1	38.0	38.8	38.7	38.0	38.5	38.8	38.9	39.1
LF_WS	65.0	68.1	73.7	64.5	64.5	64.9	66.6	69.5	73.7
DF_CS	109.8	104.6	90.6	111.2	90.6	101.1	107.2	110.2	111.2
DF_SS	136.5	102.6	104.5	137.1	102.6	104.0	120.5	136.7	137.1
DF_WS	107.9	94.3	93.2	73.0	73.0	88.2	93.8	97.7	107.9
DF_I	30.6	37.4	35.8	33.7	30.6	32.9	34.8	36.2	37.4

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 by 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*
- *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 (Clause 8) contains the Operating System and DBMS parameters used in this benchmark.

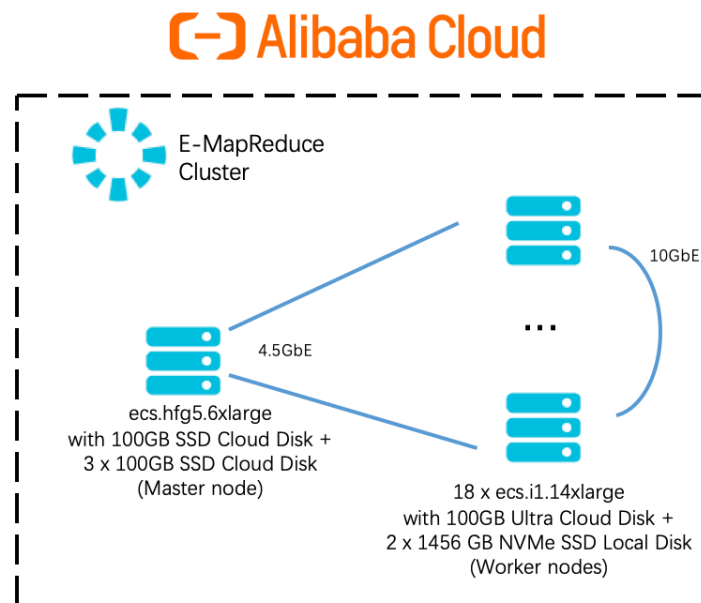
0.3 Configuration Diagrams

Diagrams of both measured and priced configurations must be provided, accompanied by a description of the differences. This includes, but is not limited to:

- *Number and type of processors*
- *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, including routers, workstations, terminals, etc., that were physically used in the test or are incorporated into the pricing structure.*
- *Type and the run-time execution location of software components (e.g., DBMS, query processing tools/languages, middle-ware components, software drivers, etc.).*

Measured Configuration

Figure 0.3: Measured Configuration



The measured configuration consisted of 19 Nodes:

Master node details (1 node):

- ECS Instance Type: ecs.hfg5.6xlarge
- Processors/Cores/Threads: 1/12/24
- Processor Model: Intel(R)Xeon(R) Gold 6149 CPU @ 3.10GHz, 22 MB L3
- Memory: 96 GB
- Storage:
 - 3 x 100 GB SSD Cloud Disk (data disk)
 - 1 x 100 GB SSD Cloud Disk (boot disk)
- Network:
 - Bandwidth (Gbit/s): 4.5
 - Packet forwarding rate (Thousand pps): 2,000
 - NIC queues: 6
 - ENIs: 8

Worker nodes details (18 nodes):

- ECS Instance Type: ecs.i1.14xlarge
- Processors/Cores/Threads: 1/28/56
- Processor Model: Intel(R)Xeon(R) CPU E5-2682 v4 @ 2.50GHz, 40 MB L3
- Memory: 224 GB
- Storage:
 - 2 x 1456 GB NVMe SSD Local Disk (data disk)
 - 1 x 100 GB Ultra Cloud Disk (boot disk)
- Network:
 - Bandwidth (Gbit/s): 10.0
 - Packet forwarding rate (Thousand pps): 1,200

- NIC queues: 4
- ENIs: 8

EMR System Components Configuration

	HDFS		YARN		Spark	
	NameNode	DataNode	Resource Manager	Node Manager	Thrift Server	Executor
Master	x		x		x	
Worker 1-18		x		x		x

Priced Configuration

There are no differences between the priced and measured configurations.

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.

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.

The store_sales, store_returns, catalog_sales, catalog_returns, web_sales, web_returns and inventory are partitioned. The partition columns for these tables respectively are ss_sold_date_sk, sr_returned_date_sk, cs_sold_date_sk, cr_returned_date_sk, ws_sold_date_sk, wr_returned_date_sk and inv_date_sk.

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.

Horizontal partitioning is used on store_sales, store_returns, catalog_sales, catalog_returns, web_sales, web_returns and inventory tables and the partitioning columns are ss_sold_date_sk, sr_returned_date_sk, cs_sold_date_sk, cr_returned_date_sk, ws_sold_date_sk, wr_returned_date_sk and inv_date_sk. The partition granularity is by day.

2.4 Replication

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

All the objects are replicated by HDFS in 3 replications.

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 as they existed upon completion of the build.

Table 3.1 Initial Number of Rows

Table Name	Row Count
call_center	54
catalog_page	40,000
catalog_returns	1,440,033,112
catalog_sales	14,399,964,710
customer	65,000,000
customer_address	32,500,000
customer_demographics	1,920,800
date_dim	73,049
household_demographics	7,200
income_band	20
inventory	1,311,525,000
item	402,000
promotion	2,000
reason	70
ship_mode	20
store	1,500
store_returns	2,879,544,876
store_sales	28,799,901,788
time_dim	86,400
warehouse	25
web_page	4,002
web_returns	720,020,485
web_sales	7,199,963,324
web_site	78

3.2 Distribution of Tables and Logs Across Media

The distribution of tables and logs across all media must be explicitly described using a format similar to that shown in the following example for both the tested and priced systems.

Table 3.2 Distribution of Tables and Logs

Server Node	Disk Type	Disk drive	Description of Content
emr-header-1	SSD Cloud Disk	/dev/vdb (/mnt/disk1)	logs
emr-header-1	SSD Cloud Disk	/dev/vd{c,d} (/mnt/disk2 RAID-1)	Hive metadata and HDFS metadata
emr-worker-{1 - 18}	Local SSD Disk	/dev/vd{b,c} (/mnt/disk[1-2])	logs, temp files, cache, replica of table data (See Section 3.4)
emr-header-1	SSD Cloud Disk	/dev/vda	Operating system, root directory, EMR software
emr-worker-{1 - 18}	Ultra Cloud Disk	/dev/vda	Operating system, root directory, EMR software

All the Table contents were on HDFS. Table size on HDFS:

175.9 K hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/call_center
 2.5 M hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/catalog_page
 115.2 G hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/catalog_returns
 1.1 T hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/catalog_sales
 3.8 G hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/customer
 835.9 M hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/customer_address
 20.7 M hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/customer_demographics
 2.1 M hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/date_dim
 108.1 K hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/household_demographics
 35.7 K hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/income_band
 11.3 G hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/inventory
 39.1 M hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/item
 234.4 K hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/promotion
 39.8 K hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/reason
 52.4 K hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/ship_mode
 349.5 K hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/store
 157.1 G hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/store_returns
 1.5 T hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/store_sales
 1.5 M hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/time_dim
 87.8 K hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/warehouse
 188.3 K hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/web_page
 56.9 G hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/web_returns
 519.8 G hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/web_sales
 161.6 K hdfs://emr-header-1:9000/user/hive/warehouse/tpcds_hdfs_parquet_10000.db/web_site

3.3 Mapping of Database Partitions/Replications

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

Neither database partitions nor replications are mapped to specific devices.

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

The database tables were on top of Hadoop Distributed Filesystem (HDFS). HDFS maintains 3 copies of table data.

For the database and file system metadata, they are stored on a RAID-1 device, which is built on top of 2 local drives of the master node.

3.5 DBGEN 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.

Dsdgen version 2.11.0 was used. Two minor changes are made to the dsdgen tool. To reduce the dsdgen execution time, the dsdgen code is wrapped as a Map/Reduce job. The wrapper does not change any of the TPC-provided code. Patches for dsdgen tool and the wrapper with source codes were included in the Supporting Files.

3.6 Database Load time

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

The database load time was 2,402.1 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.

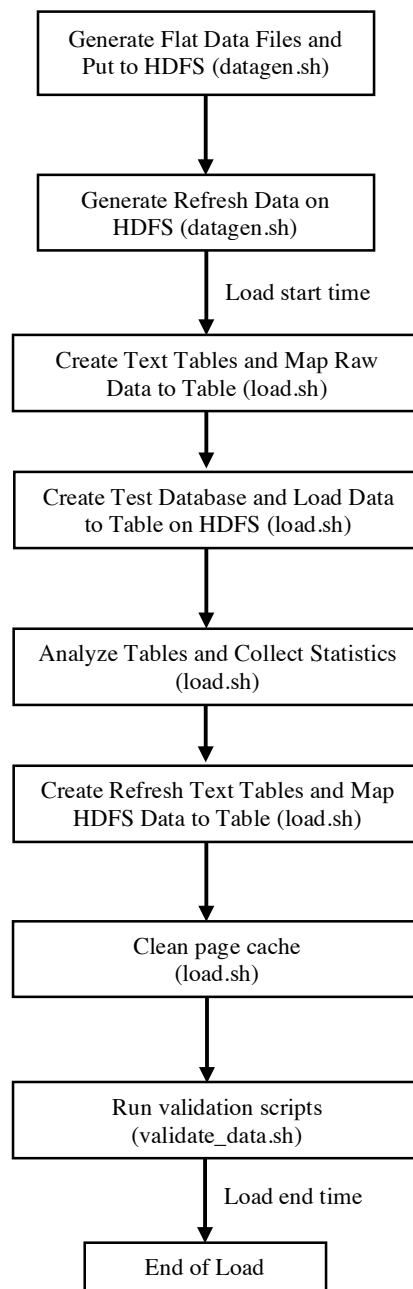
The data storage ratio is $54,616 / 10,000 = 5.46$.

Total Storage Capacity (Disk) = $(100 + 100 * 3)$ (Master node) + $(100 + 1,456 * 2) * 18$ (Worker nodes) = 54,616 GB

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.

The tables were loaded as shown in Figure 3.8. All of the related source code and scripts are included in the Supporting Files.

Figure 3.8: Block Diagram of Database Load Process

The final database load time is (load end time – load start time – duration of validation scripts).

3.9 Qualification Database Configuration

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

The qualification database is created using the same scripts as the test database with the following exceptions:

- The Scale factor is adjusted to 1GB
- The script create_qual_text_tables.sql is used instead of create_text_tables.sql to build the database on the local node.

All of the related source code and scripts are included in the Supporting Files.

Clause 4 and 5: Query and Data Maintenance Related Items

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 map/reduce wrapper based on TPC-supplied dsdgen version 2.11.0 and dsqgen version 2.11.0 were used.

4.3 Generating Values for Substitution Parameters

The method used to generate values for substitution parameters must be disclosed. The version number (i.e., the major revision number, the minor revision number, and third tier number) of dsqgen must be disclosed.

TPC supplied dsqgen version 2.11.0 was used to generate the substitution parameters:

```
./dsqgen -directory ../query_templates -input ../query_templates/templates.lst -scale 10000 -streams 9 -
output_dir ../queries -dialect sparksql -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 the execution of the query text against the qualification database. If minor modifications have been applied to any functional query definitions or approved variants in order to obtain executable query text, these modifications must be disclosed and justified. The justification for a particular minor query modification can apply collectively to all queries for which it has been used. The output data for the power and Throughput Tests must be made available electronically upon request.

Supporting Files Archive contains the actual query text and query output. Following are the modifications to the query.

The following MQM are used:

- Use vendor specific string concatenation operator. (MQM c.3)
 - Q5
 - Q66
 - Q80
 - Q84
- Use vendor-specific syntax of date expressions. (MQM f.1)
 - Q5
 - Q12
 - Q16
 - Q20
 - Q21
 - Q32
 - Q37
 - Q40
 - Q77

- Q80
- Q82
- Q94
- Q95
- Q98
- Use back quotes instead of double quotes to delimit column names. (MQM e.1)
 - Q16
 - Q32
 - Q50
 - Q62
 - Q94
 - Q95
 - Q99

Query results are inserted in a file (Clause 4.2.5) using an external table with column delimiter

- Q64 with an external table named q64_result_[s](stream[s])

The Supporting Files Archive contains the full set of executable query text template used in the test.

4.5 Query Substitution Parameters and Seeds Used

All the 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 (see Clause 5.1.8 for definition and usage restrictions) used to implement the data maintenance functions must be disclosed. Also, any disk storage used for the staging area must be priced, and any mapping or virtualization of disk storage must be disclosed.

No staging area was used.

Clause 6: Data Persistence Properties Related Items

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 test was performed by failing a disk drive on one worker node and failing one disk in the RAID-1 volume on the master node. These failures were included during the execution of the first data maintenance test.

The worker disk failure was simulated by removing and invalidating the corresponding data directory on the disk, and the master disk failure was simulated via the Linux utility mdadm. After the failures, the test continued to run until completion.

The Supporting Files Archive contains the logs of status before and after the disk failures.

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.

There only activity between the end of the load test and the beginning of the performance test was the generation of the executable query text.

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.

7.3 Timing Intervals for Each Query and Refresh Function

The timing intervals defined in Clause 7 must be disclosed.

See the Executive Summary at the beginning of this report.

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.

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.

7.7 Performance Metrics

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

QphDS@10000GB = 5,261,414

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.

beeline is the client of EMR Spark. It connects to the Spark Thrift Server by JDBC. The command is:

```
beeline -u jdbc:hive2://localhost:10001 -f sqlfile
```

The Spark Thrift Server accepts SQL queries from the beeline clients and processes the queries. The Thrift Server manages multiple executor nodes. All queries are compiled on the Thrift Server and then submitted to the Spark Executors as a job. When the job finishes, the Thrift Server takes the result from the Executors and sends it to beeline.

In the test, emr-header-1 is configured as the Spark Thrift Server, and all the EMR workers are configured as Spark Executors.

The Supporting Files Archive contains all the command, scripts and logs.

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 (<http://www.tpc.org>)

A detailed list of all licensed services, hardware and software, is provided in the Executive Summary of this report.

9.2 Availability Date

The System Availability Date (see Clause 7.6.5) 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 of the categories for which a pricing subtotal must be. All Availability Dates required to be reported must be disclosed to a precision of 1 day, but the precise format is left to the test sponsor.

The total system is available as of the date of this report.

9.3 Country-Specific Pricing

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

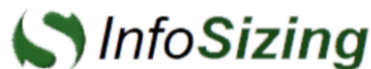
The configuration is priced for the Chinese market.

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: Francois Raab, of InfoSizing.



Benchmark sponsor: Wu Wei
Alibaba Cloud Computing Ltd.
969 West Wen Yi Road
Yu Hang District, Hangzhou,
Zhejiang
China

September 3, 2019

I verified the TPC Benchmark™ DS (TPC-DS™ v2.11.0) performance of the following configuration:

Platform: Alibaba Cloud E-MapReduce
Operating System: CentOS Linux Release 7.4
Database Manager: Alibaba Cloud E-MapReduce 3.21.2

The results were:

Performance Metric **5,261,414 QphDS@1000GB**
Database Load Time 40m 02s

Servers Alibaba Cloud Elastic Compute Service Server

1 Master Node, with:

CPU	1 x Intel Xeon Gold 6149 (3.10GHz, 22MB L3)		
Memory	96 GB		
Disks	Qty	Size	Type
	1	100 GB	SSD Cloud Disk (boot)
	3	100 GB	SSD Cloud Disk (data)

18 Worker Nodes, with:

CPU	1 x Intel Xeon E5-2682 v4 (2.50GHz, 40MB L3)		
Memory	224 GB		
Disks	Qty	Size	Type
	1	100 GB	SSD Ultra Cloud Disk (boot)
	2	1,456 GB	NVMe SSD Local Disk (data)

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

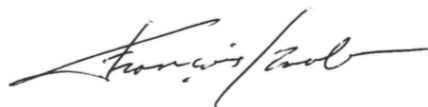
The following verification items were given special attention:

- The database records were defined with the proper layout and size
- The database population was generated using Dsdgen
- The database was properly scaled to 10,000GB and populated accordingly
- The database load time was correctly measured and reported
- The query templates were produced using approved minor query modifications and query variants
- The query input variables were generated by Dsqgen
- The execution of the queries against the qualification database produced compliant output
- The tests were driven and sequenced according to the requirements
- The throughput tests involved 4 query streams
- The execution times for queries and data maintenance functions were correctly measured and reported
- The data accessibility test was performed and verified
- 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,



François Raab, TPC Certified Auditor

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 Map/Reducer 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/mtsqs_[s]/
	Data file with delete dates	SupportingFiles/Clause_5/delete/
		SupportingFiles/Clause_5/inventory_delete/
Clause 6	Data accessibility test scripts and logs	SupportingFiles/Clause_6/
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 [q] in stream [s] (top 500)	SupportingFiles/Clause_7/stream_[s]_results/query_[q].out
Clause 8	EMR Configuration Inventory	SupportingFiles/Clause_8/

Appendix A: Purchase Page of Creating Alibaba Cloud E-MapReduce Cluster with 3-Year Subscription

E-MapReduce Quick Purchase Cluster Wizard Back

Software Settings Hardware Settings Basic Settings **4 Confirm**

Configuration List

Cluster Name: TPC-DS
 Region: China (Beijing)
 EMR Version: EMR-3.21.2
 Cluster Type: HADOOP
 Billing Method: Subscription
 Subscription Period: 3 Years
 Zone: China (Beijing) E
 Network Type: VPC
 Security Group Name: sg-tpc-ds

Master Machine Configuration: ecs.hfg5.6xlarge 24vCPU 96GB / 1 Nodes
 System Disk: SSD / 100G * 1
 Data Disk: SSD / 100G * 3

Core Machine Configuration: ecs.i1.4xlarge 56vCPU 224GB / 18 Nodes
 System Disk: Ultra Disk / 100G * 1
 Data Disk: Local Disk / 1456G * 2

Kerberos Mode: Standard
 High Availability: Standard
 Assign Public Network IP: No
 Remote Logon: No
 Metadatabase: Default

[E-MapReduce Service Terms](#)

	Total Cost	ECS	EMR
Total	¥ 2738295.36	¥ 2738295.36	¥ 0
Save	¥ 4233568.32	¥ 3329936.64	¥ 903631.68

Total **¥ 2738295.36** Save **¥ 4233568.32**

Previous: Basic Settings Create

Appendix B: Third Party Price Quotes

Lenovo MIIX 210 Laptop (Original Chinese version)

TMALL天猫 越禧数码专营店 描述 服务 物流 4.8↓ 4.7↓ 4.7↓ 手机逛

搜索 天猫 商品/品牌/店铺

越禧数码专营店
yuè xǐ shù mǎ zhuān yīng diàn
官方授权 | 正品保障 [❤️关注](#)

荣耀平板5 10.1英寸 麒麟八核芯片/人脸识别 [立即购买>>](#)

荣耀平板5 8英寸 八核/游戏加速/人脸识别 [立即购买>>](#)

所有商品 首页 荣耀平板电脑 联想平板电脑 华为平板电脑 荣耀笔记本电脑 政企客户采购

Lenovo/联想 MIIX 320/210四核平板电脑二合一笔记本10.1英寸Win10 学习办公娱乐pc轻薄便携笔记本电脑
三期分期免息&下单享暖心好礼&大量现货速发

价格 **¥1348.00**

本店活动 满599元减5元,包邮; 满999元减10元,包邮 [更多优惠](#)

运费 上海 至 杭州 快速: 0.00

月销量 **9** | 累计评价 **24** | 送天猫积分 **134**

颜色分类 **银色**

套餐类型 **MIIX 210 [HD/2G/32G]** MIIX 320 [HD/2G/32G]
MIIX 320 [FHD/4G/64G] MIIX 320 [FHD/4G/128G]
MIIX 320 [HD/2G/32G] 白色 MIIX 325 [HD/4G/64G] 黑色

数量 件 库存77件

服务 **延长保修一年 ¥0.00**

花呗分期 [该商品最高可享3期分期免息](#)

Lenovo MIIX 210 Laptop (Google Chrome translated English version)

TMALL天猫 越禧Digital franchise store description logistics Mobile 4.8↓ 4.7↓ 4.7↓ 5.0

Search Tmall merchandise/brand/store

越禧数码专营店
yuè xǐ shù mǎ zhuān yīng diàn
官方授权 | 正品保障 [❤️关注](#)

荣耀平板5 10.1英寸 麒麟八核芯片/人脸识别 [立即购买>>](#)

荣耀平板5 8英寸 八核/游戏加速/人脸识别 [立即购买>>](#)

All goods Home Glory tablet Lenovo tablet Huawei tablet Glory laptop Government and en

Lenovo/Lenovo MIIX 320/210 quad-core tablet two-in-one notebook 10.1 inch Win10 learning office entertainment pc thin portable laptop
Three-phase installment interest-free & order enjoy warmhearted gift & a lot of spot speed

price **¥1348.00**

Our event 599 yuan minus 5 yuan, 包邮; over 999 yuan minus 10 yuan [More offers](#)

Freight Shanghai to Hangzhou Express: 0.00

Monthly sales **9** | Cumulative evaluation **24** | Send Tmall Points **134**

Color Classification **Silver**

Package Type **MIIX 210 [HD/2G/32G]** MIIX 320 [HD/2G/32G]
MIIX 320 [FHD/4G/64G] MIIX 320 [FHD/4G/128G]
MIIX 320 [HD/2G/32G] white MIIX 325 [HD/4G/64G] black

Quantity parts inventory 77

service **Extended warranty one year ¥0.00**