



Tencent Cloud Computing Ltd.

TPC Benchmark™ DS

Full Disclosure Report

for

Tencent Cloud TDSQL For PostgreSQL

(Dedicated cluster with 8 nodes)

using

TDSQL For PostgreSQL 5.21.18

and

TencentOS Server 3.2

First Edition

September, 2024

First Edition – September, 2024

Copyright © 2024 Tencent Cloud Computing (Beijing) Co., Ltd. and/or its affiliates. All rights reserved.

Tencent Cloud Computing (Beijing) Co., Ltd., as the Sponsor of this benchmark test, believes that the pricing information in this document is accurate as of the publication date, and is subject to change without notice. The Sponsor assumes no responsibility for any errors that may appear in this document. The pricing information is believed to accurately reflect the current prices as of the publication date. However, the Sponsor provides no warranty of the pricing information in this document.

The performance information in this document is for guidance only. Benchmark results are highly dependent on many factors including workload, hardware, operating environments, specific application requirements, and system design and implementation. Relative system performance may vary significantly as a result of these and other factors. The Sponsor does not warrant or represent that a user can or will achieve the same or similar performance. No warranty on system performance is either expressed or implied.

TDSQL is a registered trademark of Tencent and/or its affiliates.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation.

TPC Benchmark and TPC-DS are trademarks of the Transaction Processing Performance Council.

All other products mentioned herein are trademarks or registered trademarks of their respective owners.

©Copyright 2024 Tencent Cloud Computing (Beijing) Co., Ltd.

All rights reserved. This product and related documentation are protected by copyright and distributed under licenses restricting its use, copying, distribution, and de-compilation. No part of this product or related documentation may be reproduced in any form by any means without the prior written authorization of Tencent Cloud Computing (Beijing) Co., Ltd. and its licensors, if any.

THIS PUBLICATION IS PROVIDED WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS PUBLICATION COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. TENCENT CLOUD COMPUTING (BEIJING) CO., LTD. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS DOCUMENT.

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

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 Tencent 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 3.2.0.


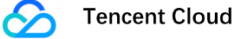

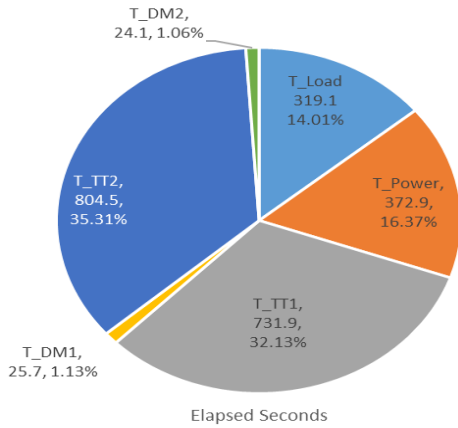
The test was conducted at a Scale Factor of 10000GB with 8 TDSQL for PostgreSQL Cluster Node T0-CI81X-100GS running TDSQL for PostgreSQL 5.21.18 on TencentOS Server 3.2.


Measured Configuration

Company Name	Cluster Node	Database Software	Operation System
Tencent Cloud Computing Ltd.	TDSQL for PostgreSQL Cluster Node T0-CI81X-100GS	Tencent Cloud TDSQL for PostgreSQL 5.21.18	TencentOS Server 3.2

TPC Benchmark™ DS Metrics

Total System Cost (CNY)	TPC-DS Throughput (QphDS@10000GB)	Price/Performance (CNY / kQphDS@10000GB)	Availability Date
¥2,723,728	72,603,042	¥37.52	12/12/2024

		Tencent Cloud TDSQL for PostgreSQL		TPC-DS: 3.2.0 TPC-Pricing: 2.9.0 Report Date: 2024-10-08	
Total System Cost		TPC-DS Throughput	Price / Performance		System Availability Date
¥2,723,728 CNY		72,603,042 QphDS@10000GB	¥37.52 CNY/kQphDS@10000GB		12/12/2024
Dataset Size¹		Database Manager	Operation System		Other Software
10,000 GB		TDSQL for PostgreSQL 5.21.18	TencentOS Server 3.2		N/A
 TDSQL for PostgreSQL Instance  Cluster with 2 coordinator nodes 128 data nodes 8 x T0-CI81X-100GS Intel Emerald Rapids(-/3.0GHz) 112 vCPU/ 5600 64G*16/ NVMe 7.68T*16/ SATA 960G*2/ Silver Fir 100G(2P)*1			 Elapsed Seconds		
Benchmarked Configuration			Load includes backup = No		
			RAID = RAID-10 for all database objects (table data, logs, EADs, metadata)		
System Configuration:			TDSQL for PostgreSQL Cluster		
Servers:			8 x T0-CI81X-100GS		
Total Processors/Cores/Threads:			896 vCPU		
Total Memory:			8,192 GB		
Total Storage²:			929,832.46 GB		
Storage Ratio³:			92.99		
Server Configuration:			Per node		
Processors:			Intel Emerald Rapids(-/3.0GHz)		
Memory:			1,024 GB		
Network:			Silver Fir Intelligent NIC Speed 100 Gbps		
Storage Device:			16 x 7.68 TB NVMe SSD Local Disk (data disk) 2 x 960 GB SATA Local Disk (system disk)		
¹ . Dataset Size includes only raw data (i.e., no temp, index, redundant storage space, etc.). ² . Total Storage = 929,832.46 GB (2x 960 GB and 16x 7.68TB per node should be converted to base-2) ³ . Storage Ratio = Total Storage / SF = 92.99					

		Tencent Cloud TDSQL for PostgreSQL			TPC-DS: 3.2.0 TPC-Pricing: 2.9.0 Report Date: 2024-10-08			
Description	Part Number	Src	Unit Price	Qty(CNY)	Ext. Price (CNY)	3-Year Maint. (CNY)		
TDSQL for PostgreSQL Cluster								
- T0-CI81X-100GS		1	included	8	2,685,363.95	included		
- System Local Disk (2 x 960GB)			included	8				
- NVMe Local Disk (16 x 7.68TB)			included	8				
- TencentOS Server 3.2			included	8				
- TDSQL for PostgreSQL 5.21.18			included	8				
- Silver Fir Intelligent NIC			included	8				
Licensed Software Services Sub-Total							2,685,363.95	0.00
Other Components								
MacBook Pro 16 inch Apple M1 Pro (includes spares)		2	12,788.00	3	38,364.00			
Other Components Sub-Total					38,364.00	0.00		
1 = Tencent Cloud, 2 = jd.com.cn All Licensed Services prices are per month and based on 3-year pre-paid subscriptions. Audited by Doug Johnson, InfoSizing				3-Year Cost of Ownership: QphDS@10000GB: ¥/kQphDS@10000GB:		2,723,728 72,603,042 37.52		
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.								

Metric Details

Name	Value	Description / Unit
SF	10,000	Scale Factor 10TB
S	8	Total Throughput Streams
Sq	4	Streams / Throughput Test
Q	396	Queries / Throughput Test
T_LD	0.0036	hours@10000
T_PT	0.4144	hours@10000
T_TT	0.4268	hours@10000
T_DM	0.0139	hours@10000

Secondary Metrics

Name	Value	Unit
T_Load	319.1	seconds@10000
T_Power	372.9	seconds@10000
T_TT1	731.9	seconds@10000
T_TT2	804.5	seconds@10000
T_DM1	25.7	seconds@10000
T_DM2	24.1	seconds@10000

Test Timeline

Test	Start	End	Seconds	hh:mm:ss
Load	2024-09-11 03:39:56.381	2024-09-11 03:45:15.449	319.068	0:05:19
Audit/Admin	2024-09-11 03:45:15.457	2024-09-11 03:58:37.289	801.832	0:13:22
Power	2024-09-11 03:58:37.313	2024-09-11 04:04:50.150	372.837	0:06:13
TT-1	2024-09-11 04:04:50.203	2024-09-11 04:17:02.061	731.858	0:12:12
DM-1	2024-09-11 04:17:02.101	2024-09-11 04:17:27.781	25.680	0:00:26
TT-2	2024-09-11 04:17:27.832	2024-09-11 04:30:52.264	804.432	0:13:24
DM-2	2024-09-11 04:30:52.303	2024-09-11 04:31:16.345	24.042	0:00:24

Stream	Start Time	End Time	Seconds	hh:mm:ss
Power - 0	2024-09-11 03:58:37.313	2024-09-11 04:04:50.150	372.837	0:06:13
TT-1 - 1	2024-09-11 04:04:50.203	2024-09-11 04:16:54.686	724.483	0:12:04
TT-1 - 2	2024-09-11 04:04:50.203	2024-09-11 04:16:23.443	693.240	0:11:33
TT-1 - 3	2024-09-11 04:04:50.203	2024-09-11 04:16:54.026	723.823	0:12:04
TT-1 - 4	2024-09-11 04:04:50.204	2024-09-11 04:17:02.061	731.857	0:12:12
DM-1 - 1	2024-09-11 04:17:02.101	2024-09-11 04:17:15.577	13.476	0:00:13
DM-1 - 2	2024-09-11 04:17:15.599	2024-09-11 04:17:27.781	12.182	0:00:12
TT-2 - 5	2024-09-11 04:17:27.832	2024-09-11 04:30:37.395	789.563	0:13:10
TT-2 - 6	2024-09-11 04:17:27.833	2024-09-11 04:30:52.264	804.431	0:13:24
TT-2 - 7	2024-09-11 04:17:27.833	2024-09-11 04:30:42.498	794.665	0:13:15
TT-2 - 8	2024-09-11 04:17:27.833	2024-09-11 04:30:48.568	800.735	0:13:21
DM-2 - 3	2024-09-11 04:30:52.303	2024-09-11 04:31:03.989	11.686	0:00:12
DM-2 - 4	2024-09-11 04:31:04.010	2024-09-11 04:31:16.345	12.335	0:00:12

Timing Intervals for Refresh Functions (in Seconds)

Function	DM-1		DM-2		<i>Min</i>	<i>25th</i>	<i>Median</i>	<i>75th</i>	<i>Max</i>
	Run 1	Run 2	Run 3	Run 4					
ID	Run 1	Run 2	Run 3	Run 4	<i>Min</i>	<i>25th</i>	<i>Median</i>	<i>75th</i>	<i>Max</i>
DF_CS	3.0	3.0	3.1	3.2	3.00	3.00	3.05	3.13	3.20
DF_I	0.4	0.5	0.4	0.4	0.40	0.40	0.40	0.43	0.50
DF_SS	8.3	7.3	7.2	7.9	7.20	7.28	7.60	8.00	8.30
DF_WS	2.9	2.5	2.6	2.7	2.50	2.58	2.65	2.75	2.90
LF_CR	3.8	3.6	3.0	2.9	2.90	2.98	3.30	3.65	3.80
LF_CS	5.3	5.0	4.5	4.5	4.50	4.50	4.75	5.08	5.30
LF_I	0.7	0.7	1.0	0.9	0.70	0.70	0.80	0.93	1.00
LF_SR	5.1	2.8	1.9	4.4	1.90	2.58	3.60	4.58	5.10
LF_SS	5.1	5.0	4.4	4.4	4.40	4.40	4.70	5.03	5.10
LF_WR	3.1	2.9	2.3	2.3	2.30	2.30	2.60	2.95	3.10
LF_WS	4.9	4.5	4.0	4.1	4.00	4.08	4.30	4.60	4.90

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 Tencent 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.

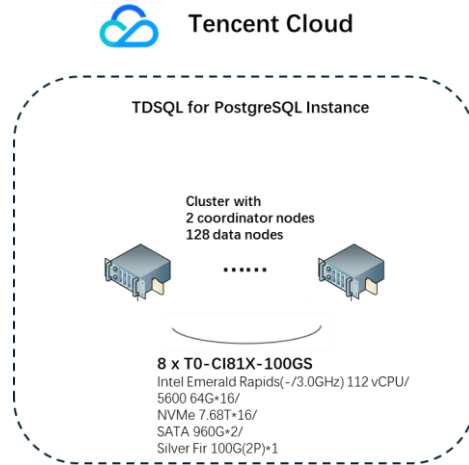
0.3 Configuration Diagrams

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 8 Nodes:

Node details (8 nodes):

- T0-CI81X-100GS
- Processors/Cores/Threads: 112 vCPU
- Processor Model: Intel Emerald Rapids(-/3.0GHz)
- Memory: 1,024 GB
- Storage:
 - 2 x 960 GB SATA Disk (system disk)
 - 16 x 7.68 TB NVMe SSD Disk (data disk)
- Network: Silver Fir Intelligent NIC Speed: 100 Gbps

System Components Configuration

	1	2	3	4	5	6	7	8
Coordinator	cn001	cn002						
Datanode	dn001- dn016	dn017- dn032	dn033- dn048	dn049- dn064	dn065- dn080	dn081- dn096	dn097- dn112	dn113- dn128
GTM		gtm						

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.

Horizontal partitioning is used as described.

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.

All tables are partitioned. The partition columns for the tables are:

customer: c_customer_sk
customer_address: ca_address_sk
customer_demographics: cd_demo_sk
catalog_sales: cs_item_sk, cs_sold_date_sk
catalog_returns: cr_item_sk, cr_returned_date_sk
inventory: inv_item_sk, inv_date_sk
store_returns: sr_item_sk, sr_returned_date_sk
store_sales: ss_item_sk, ss_sold_date_sk
web_returns: wr_item_sk, wr_returned_date_sk
web_sales: ws_item_sk, ws_sold_date_sk
call_center: cc_call_center_sk
catalog_page: cp_catalog_page_sk
date_dim: d_date_sk
household_demographics: hd_demo_sk
income_band: ib_income_band_sk
item: i_item_sk
promotion: p_promo_sk
reason: r_reason_sk
ship_mode: sm_ship_mode_sk
store: s_store_sk
time_dim: t_time_sk
warehouse: w_warehouse_sk
web_page: wp_web_page_sk
web_site: web_site_sk

2.4 Replication

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

No physical object was replicated.

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	2,879,611,062
store_sales	28,799,864,841
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
1-8	Local SATA SSD	/dev/sda /dev/sdb	Operating system, root
	Local NVMe	/dev/md127 (/dev/nvme[0-15]n1 (Soft raid 10)	TDSQL for PostgreSQL, logs, tables, temp files, swap

Table size on local storage:

catalog_sales: 1696 GB

catalog_returns: 129 GB

store_sales: 2388 GB

store_returns: 199 GB

web_sales: 878 GB

web_returns: 64 GB

inventory: 8060 MB

date_dim: 189 MB

time_dim: 88 MB

item: 206 MB

customer: 4730 MB

customer_demographics: 80 MB

household_demographics: 44 MB

customer_address: 1649 MB

store: 92 MB

promotion: 58 MB

warehouse: 2608 kB

ship_mode: 2352 kB

reason: 3544 kB

income_band: 2352 kB

call_center: 5288 kB

web_page: 97 MB

catalog_page: 74 MB

web_site: 6104 kB

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

For each node, a soft RAID10 volume is created on 16 PCIe NVMe drives using the mdadm command and all TDSQL for PostgreSQL data is stored in this volume.

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 3.2.0 was used. No changes were made to the dsdgen tool.

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 319.068 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.

Total Storage Capacity (8 Local nodes) = 929,832.46 GB (2x 960 GB and 16x 7.68TB per node should be converted to base-2)

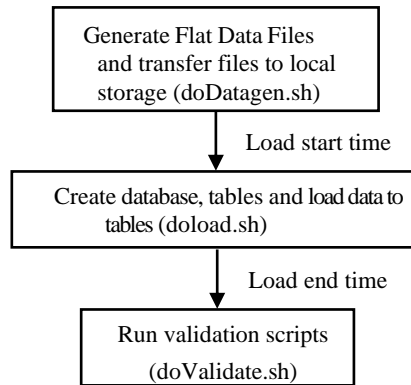
The data storage ratio is $929,832.46 \text{ GB} / 10,000 \text{ GB} = 92.99$

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).

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

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.

TPC-supplied dsdgen version 3.2.0 and dsqgen version 3.2.0 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 3.2.0 was used to generate the substitution parameters:

```
./dsqgen -streams $stream -input ../query_templates/templates.lst -directory ../query_templates -dialect tbase -scale $TPCDS_SCALE -rngseed $SEED -verbose y -output_dir $STREAM_TMP
```

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 syntax of date expressions. (MQM f.1)
 - Q5
 - Q12
 - Q16
 - Q20
 - Q21
 - Q32
 - Q37
 - Q40
 - Q77

- Q80
- Q82
- Q92
- Q94
- Q95
- Q98
- Use column references expression in ORDER BY clause (MQM e.2)
 - Q36
 - Q70
 - Q86
- Query results are inserted in a file (Clause 4.2.5) using an external table with column delimiter
 - Q34 with readable/writable external table named q34_result_[s]_r_ext/ q34_result_[s]_w_ext (stream[s])
 - Q39 with readable/writable external table named q39_result_[s]_r_ext/ q39_result_[s]_w_ext (stream[s])
 - Q64 with readable/writable external table named q64_result_[s]_r_ext/ q64_result_[s]_w_ext (stream[s])
 - Q71 with readable/writable external table named q71_result_[s]_r_ext/ q71_result_[s]_w_ext (stream[s])
 - Q98 with readable/writable external table named q98_result_[s]_r_ext/ q98_result_[s]_w_ext (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 the local storage of one node. This failure was induced during the execution of the first data maintenance test. The storage on each nodes made of 16 PCIe NVMe. The storage failure was simulated by removing access to 1 of the PCIe NVMe. The Supporting Files Archive contains the logs of status before and after the storage 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.

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 and the execution of audit scripts.

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 = 72,603,042

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.

Psql was used to submit the queries. It connects to the TDSQL for PostgreSQL instance.

The command run a SQL file

```
psql -h ${host} -p ${port} -d tpcds -f SQL.sql
```

The command run a SQL command:

```
psql -h ${host} -p ${port} -d tpcds -c "SQL command"
```

The TDSQL for PostgreSQL instance accepts SQL queries from the psql clients and processes the queries. All queries are compiled on the Coordinator node and then dispatched to the data nodes as distributed tasks. When the tasks finish, their results are collected by the Coordinator node which send the query output to the psql client.

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: Doug Johnson, of InfoSizing.

Anqun Pan
Director
TEG Database R&D Department
No. 16 Gaoxin South Avenue 10
Shenzhen, China

October 7, 2024

I verified the TPC Benchmark™ DS v3.2.0 performance of the following configuration:

Platform: Tencent Cloud TDSQL For PostgreSQL
(Dedicated cluster with 8 nodes)
Operating System: TencentOS Server 3.2
Database Manager: TDSQL for PostgreSQL 5.21.18

The results were:

Performance Metric 72,603,042 QphDS@10000

Secondary Metrics (seconds@10000)

T _{Load}	319.1	T _{TT1}	731.9	T _{DM1}	25.7
T _{Power}	372.9	T _{TT2}	804.5	T _{DM2}	24.1

System Under Test 8x T0-CI81X-100GS each with:

CPU	112x vCPU		
Memory	1,024 GiB		
Storage	Qty	Size	Type
	2	960 GB	SATA
	16	7.68 TB	NVMe SSD

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 v3.2.0
- The database was properly scaled to 10,000 GB and populated accordingly.
- The primary and secondary metrics were correctly measured and reported.

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

- The query templates were produced using approved minor query.
- The query substitution parameters were generated using DSQGen v3.2.0.
- 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 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,

A handwritten signature in black ink, reading "Doug Johnson", with a long horizontal flourish extending to the right.

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	Tools for data generation, data split (including base table and table for maintenance)	SupportingFiles/Clause_3/doDatagen.sh SupportingFiles/Clause_3/split_tpcds.sh SupportingFiles/Clause_3/split_tpcds_dm.sh
	Scripts, SQL and logs for load base tables (Specifically, tfdw_import.sh is for creating base/foreign tables, loading data to base tables and analyzing base tables, all tasks are scheduled in parallel)	SupportingFiles/Clause_3/doLoad.sh SupportingFiles/Clause_3/sqls/ddl.sql SupportingFiles/Clause_3/tfdw_import.sh SupportingFiles/Clause_3/tfdw_import_sql/*.sql SupportingFiles/Clause_3/create_tfdw_server.sql SupportingFiles/Clause_3/sqls/create_tpcds_ddl.schedule SupportingFiles/Clause_3/sqls/analyze_tpcds_job.schedule SupportingFiles/Clause_3/sqls/analyze_tpcds_job.schedule SupportingFiles/Clause_3/logs/create_tpcds_tables.out SupportingFiles/Clause_3/logs/load_tpcds_time.log SupportingFiles/Clause_3/logs/load_tpcds.out SupportingFiles/Clause_3/logs/split_tpcds_log/* SupportingFiles/Clause_3/logs/split_tpcds_dm_log/* SupportingFiles/Clause_3/logs/tfdw_import_log/*
	Scripts and SQL for validation and log files	SupportingFiles/Clause_3/doValidate.sh SupportingFiles/Clause_3/sqls/count_tables.sql SupportingFiles/Clause_3/sqls/desc_tables.sql SupportingFiles/Clause_3/sqls/Validate_Data.sql SupportingFiles/Clause_3/sqls/Check_Insert.sql SupportingFiles/Clause_3/sqls/Check_RI.sql SupportingFiles/Clause_3/sqls/create_tpcds_vld.sql SupportingFiles/Clause_3/sqls/load_vld.sql SupportingFiles/Clause_3/logs/row_count.out SupportingFiles/Clause_3/logs/table_schema.out SupportingFiles/Clause_3/logs/validate_data.log SupportingFiles/Clause_3/logs/Check_Insert.log SupportingFiles/Clause_3/logs/Check_RI.log
	pre-generated SQL to data maintenance and output	SupportingFiles/Clause_3/sqls/mtsqls/* SupportingFiles/Clause_3/logs/fetch*.out
Clause 4	The script to execute qualification test and log file	SupportingFiles/Clause_4/doQualification_test.sh SupportingFiles/Clause_4/logs/qualification.log SupportingFiles/Clause_4/logs/qual_*.out
	SQL for qualification queries	SupportingFiles/Clause_4/queries/
	Output from executing qualification queries	SupportingFiles/Clause_4/output/

	Query templates modify	SupportingFiles/Clause_4/query_templates_modify/ query_templates.patch
Clause 5	Data maintenance execution scripts and logs files for each stream [s]	SupportingFiles/Clause_5/doRefresh.sh SupportingFiles/Clause_5/refresh.sh SupportingFiles/Clause_5/logs/mt_[r]_time.log SupportingFiles/Clause_5/logs/refresh_[s]_timing.log
	SQL scripts for DM functions for stream [s]	SupportingFiles/Clause_5/mtloadsq/* SupportingFiles/Clause_5/mtsqs_[s]/LF*.sql SupportingFiles/Clause_5/mtsqs_[s]/DF*.sql SupportingFiles/Clause_5/mtsqs_[s]/ddl.sql
	Output from executing DM functions	SupportingFiles/Clause_5/output/*.out
	Raw data files for maintenance	SupportingFiles/Clause_5/data/delete_[s].dat SupportingFiles/Clause_5/data/inventory_delete_[s].dat
	MT function and data verification sqs and outputs	SupportingFiles/Clause_5/doVerify_mt.sh SupportingFiles/Clause_5/logs/run_verify_mt_[s].log SupportingFiles/Clause_5/mtsqs_[s]/fetch*.sql SupportingFiles/Clause_5/mtsqs_[s]/verify*.sql SupportingFiles/Clause_5/mtsqs_[s]/count_mt_tables.sql SupportingFiles/Clause_5/output/mt_verify/*.out
Clause 6	Data accessibility test scripts, logs and output files	SupportingFiles/Clause_6/data_access.sh SupportingFiles/Clause_6/data_access_test.log SupportingFiles/Clause_6/data_disk_remove.out SupportingFiles/Clause_6/data_disk_status_fail.out SupportingFiles/Clause_6/data_disk_status_good.out
Clause 7	Performance test scripts and logs	SupportingFiles/Clause_7/doQueryGen.sh SupportingFiles/Clause_7/generate_queries.log SupportingFiles/Clause_7/doPower.sh SupportingFiles/Clause_7/doTT.sh SupportingFiles/Clause_7/doStream.sh SupportingFiles/Clause_7/logs/pt_time.log SupportingFiles/Clause_7/logs/tt_[r]_time.log SupportingFiles/Clause_7/logs/stream_[s]_time.log
	Query text for query [q] in stream [s]	SupportingFiles/Clause_7/stream_[s]_queries/query_[q].sql
	Output of query [q] in stream [s]	SupportingFiles/Clause_7/stream_[s]_results/query_[q].out
Clause 8	System config	SupportingFiles/Clause_8/system_profiles/ SupportingFiles/Clause_8/collect_system_profiles.sh SupportingFiles/Clause_8/collect_system_profiles.log

	Database config	SupportingFiles/Clause_8/tbase_config/ SupportingFiles/Clause_8/collect_tbase_config.sh SupportingFiles/Clause_8/collect_tbase_config.log SupportingFiles/Clause_8/collect_tbase_dir.sh SupportingFiles/Clause_8/collect_tbase_dir.log
General	Compile tpchs	SupportingFiles/General/compile_tpchs.sh SupportingFiles/General/DSGen-software-code-3.2.0rc1 SupportingFiles/General/D78DE658-6BBB-4970-B72E-B3E295448072-TPC-DS-Tool.zip
	Run all test	SupportingFiles/General/common.sh SupportingFiles/General/config.sh SupportingFiles/General/hosts SupportingFiles/General/runall.sh SupportingFiles/General/run.log
	Scripts for get q34 / q39 / q64 / q71 / q98 results, sort stream outputs and get QphH result	SupportingFiles/General/get_q34_result.sh SupportingFiles/General/get_q39_result.sh SupportingFiles/General/get_q64_result.sh SupportingFiles/General/get_q71_result.sh SupportingFiles/General/get_q98_result.sh SupportingFiles/General/result_report.sh

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

数据库配置

版本选择: PostgreSQL 兼容版 | Oracle 兼容版 | **融合版**

实例架构: **分布式** | 集中式

内核版本: 5.21.18

机型

机型	处理器型号	vCPU	内存	磁盘	数量
<input checked="" type="checkbox"/> TD-C181X-100GS	Intel Emerald Rapids(-3.00Hz)	112核	1024 GB	124800 GB	- 8 +

数据复制方式: **强同步 (可退化)** | 强同步 | 异步

用户名: dbadmin

密码:

确认密码:

高级配置

安全组: 请选择安全组, 可输入名称或完整ID | 新建安全组, 可输入名称或完整ID进行搜索

[使用说明](#)

如您有业务需要放通其他端口, 您可以 [自定义安全组](#)

标签

标签键	标签值	删除
<input type="text"/>	<input type="text"/>	<input type="button" value="删除"/>

[添加](#)

键值对板

您可以去控制台[新建标签](#)

购买时长: 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 1年 | 2年 | **3年** | 4年 | 5年 | [收起](#)

密码长度不能少于8位; 请输入确认密码; 请选择安全组; 请阅读并同意《云数据库服务条款》;

配置费用 **2685363.95元** 4319997.60元 | 备份费用 **0.05元/GB** (按GB收费, 超出赠送空间收费) | 流量费用 **0.00** | [立即购买](#)

English version (Edge translated)

Database configuration

Version selection: PostgreSQL Compatible Edition | Oracle Compatible Edition | **Fusion Edition**

Instance architecture: **Distributed** | Centralized

Kernel version: 5.21.18

Models

Models	Processor model	vCPU	memory	disk	quantity
<input checked="" type="checkbox"/> TD-C181X-100GS	Intel Emerald Rapids(-3.00Hz)	112 cores	1024 GB	124800 GB	- 8 +

Data replication mode: **Strong synchronization (degradable)** | Strong synchronization | asynchronous

Username: dbadmin

password:

Confirm your password:

Advanced configuration

Security group: Select a security group and search | Select a security group and search by name or full ID

[Directions for use](#)

If you need to allow access to other ports, you can [Customize the security group](#)

label

Label key	Tag value	删除
<input type="text"/>	<input type="text"/>	<input type="button" value="删除"/>

[Add to](#)

Key-value pairboard

You can go to the Tags console [Create a new label](#)

Duration of purchase: 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 1 year | 2 years | **3 years** | 4 years | 5 years | [Collapse](#)

The length of the password must not be less than 8 characters; Please enter the confirmation password; Please select a security group; Please read and agree to the TencentDB Terms of Service.

Configuration fees **2,685,363.95 yuan** 4319997.60 yuan | Backup charges **0.05 yuan/GB** (Free at the beginning, beyond the charge for free space) | Data transfer charges **0.00** | [Buy now](#)

Appendix B: Third Party Price Quotes

MacBook Pro 16 inch Apple M1 Pro (Original Chinese version)



MacBook Pro 16 inch Apple M1 Pro (Google Chrome translated English version)

