

TPC Benchmark™ H Full Disclosure Report

Alibaba Cloud Hologres
(with 69 Compute Containers)

using

Hologres V4.1

and

Pangu File System V2.0

Alibaba Cloud Linux 3.2104 U10 (OpenAnolis Edition)

First Edition

March 4, 2026



First Edition – March 4, 2026

Hologres V4.1

Pangu File System V2.0

Hologres and Hologres Logo are trademarks of Alibaba Group and/or its affiliates in China and other countries.

The Hologres 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 Hologres 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 Hologres Cloud products will vary depending on individual customer configurations and conditions

Abstract

Overview

This report documents the methodology and results of the TPC Benchmark™ H test conducted on Alibaba Cloud Server F91 using Hologres in conformance with the requirements of the TPC Benchmark™ H Standard Specification, Revision 3.0.1. The operating system used for the benchmark was Alibaba Cloud Linux 3.2104 U10 (OpenAnolis Edition).

The TPC Benchmark™ H was developed by the Transaction Processing Performance Council (TPC). The TPC was founded to define transaction processing benchmarks and to disseminate objective, verifiable performance data to the industry.

TPC Benchmark™ H Full Disclosure Report and other information can be downloaded from the Transaction Processing Performance Council web site at www.tpc.org.


Executive Summary Section

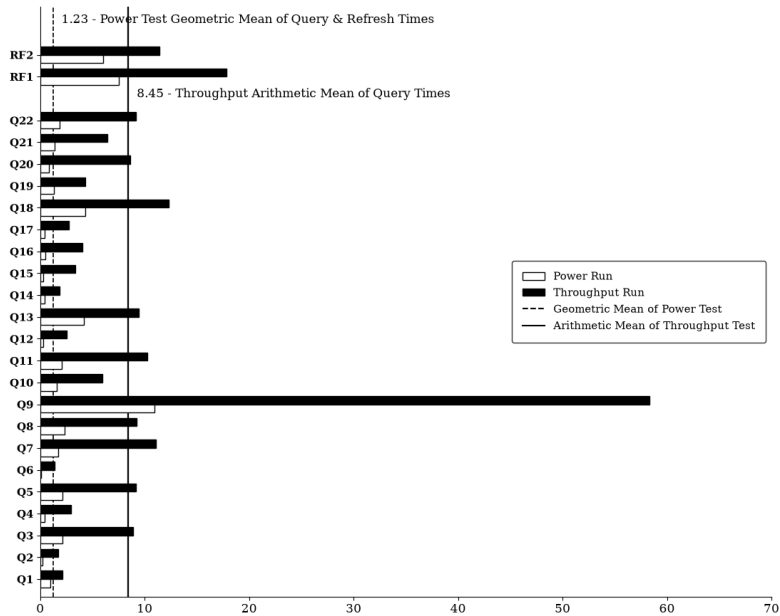
The first section of this report contains the Executive Summary and Numerical Quantities Summary of the benchmark results.

Auditor

The benchmark configuration, environment and methodology used to produce and validate the test results, and the pricing model used to calculate the cost per kQpH were audited by Doug Johnson of InfoSizing to verify compliance with the relevant TPC specifications.

The auditor's information is available in Section 9.1.

	Alibaba Cloud Hologres		TPC-H Rev.3.0.1 TPC Pricing Rev.2.9.0	
			Report Date: 4-March-26	
Total System Cost	Composite Query per Hour Metric		Price/Performance	
3,296,977 CNY	8,443,627.5 QphH@3000GB		390.47 CNY Price/kQphH@3000GB	
Database Size	Database Manager	Operating System	Other Software	Availability Date
3,000 GB	Hologres 4.1	Alibaba Cloud Linux 3.2104 U10 (OpenAnolis Edition)	Pangu File System V2.0	03/04/2026



Database Load Time: 00:20:31
 Load Includes Backup: N
 Total Data Storage/Database Size: 56.03
 Memory to Database Size Percentage: 138.7%

Storage Redundancy Levels:
 Base Tables: Level Two
 Auxiliary Data Structures: Level Two
 DBMS Temporary Space: Level Two
 OS and DBMS Software: Level Two

System Configuration

Compute Resources:

Role	Per Container		Per Role		
	vCPUs	Memory GB	# Containers	vCPUs	Memory GB
Gateway	2	8	2	4	16
JobMaster	4	16	1	4	16
MetaStore	4	16	1	4	16
HoloMaster	4	16	1	4	16
HoloWorker	16	64	64	1,024	4,096
Instance Total:			69	1,040	4,160


Storage Resources:


Instance GB	
Local	174,240
Pangu FS	6,200
MaxCompute	50

Local Details GB		
Node*	7,920	(240 + (2*3,840))
Cluster	174,240	(22 Nodes)

*Node local storage shared by all containers on the node.

Total Storage GiB: 168,094.41

	Alibaba Cloud Hologres			TPC-H Rev.3.0.1 TPC Pricing Rev.2.9.0		
				Report Date: 4-March-26		
Description	Part Number	Source	Unit Price (CNY)	Qty	Extended Price (CNY)	3 yr Maint. Price (CNY)
License Computer and Software Service						
Hologres Computing Resources						
1,024 vCPUs, 4,096GB memory, 3 years prepay						
1 meta-store, each with 4 vCPUs, 16 GB RAM						
1 job-master, each with 4 vCPUs, 16 GB RAM	China (Beijing)	1	6,266,880	1	6,266,880	Included
1 holo-master, each with 4 vCPUs, 16 GB RAM						
64 holo-workers, each with 16 vCPUs, 64 GB RAM						
2 Gateways, each with 2 vCPUs, 8 GB RAM	China (Beijing)	1	24,480	1	24,480	Included
Hologres Data Storage Resources						
6,200 GB Storage Capacity, 3 years prepay						
Discount (50%, Hologres line items above)					3,290,760	
MaxCompute Data Storage Resources 50 GB Storage Capacity, 3 years prepay	China (Beijing)	1	220	1	220	Included
License Computer and Software Service Sub-Total					3,290,980	0.00
Other Components						
Ningmei 15.6-inch Slim Notebook		2	1999.00	3	5997.00	
Other Components Sub-Total					5997.00	
1 = Alibaba Cloud 2 = Tmall.com				3-year Cost of Ownership:		3,296,977.00
All prices are based on 3-year pre-paid subscriptions.				QphH@3000GB:		8,443,627.5
				¥/kQphH@3000GB		390.47
Audited by: Doug Johnson, InfoSizing						
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 section of the TPC benchmark specifications. If you find that the stated prices are not available according to these terms, please inform the TPC at pricing@tpc.org. Thank you.						

	Alibaba Cloud Hologres	TPC-H Rev.3.0.1 TPC Pricing Rev.2.9.0
		Report Date: 4-March-26

Measurement Results

Database Scale Factor	3,000
Total Data Storage / Database Size	56.03
Percentage Memory / Database Size	138.7%
Start of Database Load Time	2026-02-25 12:55:57
End of Database Load Time	2026-02-25 13:16:29
Database Load Time	00:20:31
Query Streams for Throughput Test (S)	8
TPC-H Power	8,789,205.2
TPC-H Throughput	8,111,637.4
TPC-H Composite	8,443,627.5
Total System Price Over 3 Years	3,296,977 (CNY)
TPC-H Price/Performance Metric (¥/kQpH@3000GB)	390.47 (CNY)

Measurement Interval

Measurement Interval in Throughput Test (Ts)	234.33
--	--------

Duration of stream execution:

Power Run	Seed	Query Start Time	Total Time	RF1 Start Time	RF2 Start Time
		Query End Time	(hh:mm:ss)	RF1 End Time	RF2 End Time
	225131629	2026-02-25 13:21:27	00:00:42	2026-02-25 13:21:20	2026-02-25 13:22:09
		2026-02-25 13:22:09		2026-02-25 13:21:27	2026-02-25 13:22:15

Throughput Stream	Seed	Query Start Time	Total Time	RF1 Start Time	RF2 Start Time
		Query End Time	(hh:mm:ss)	RF1 End Time	RF2 End Time
1	225131630	2026-02-25 13:22:14	00:03:00	2026-02-25 13:22:14	2026-02-25 13:22:40
		2026-02-25 13:25:14		2026-02-25 13:22:40	2026-02-25 13:22:58
2	225131631	2026-02-25 13:22:14	00:03:00	2026-02-25 13:22:58	2026-02-25 13:23:32
		2026-02-25 13:25:14		2026-02-25 13:23:32	2026-02-25 13:23:49
3	225131632	2026-02-25 13:22:14	00:03:16	2026-02-25 13:23:49	2026-02-25 13:24:04
		2026-02-25 13:25:30		2026-02-25 13:24:04	2026-02-25 13:24:15
4	225131633	2026-02-25 13:22:14	00:03:13	2026-02-25 13:24:15	2026-02-25 13:24:45
		2026-02-25 13:25:27		2026-02-25 13:24:45	2026-02-25 13:25:13
5	225131634	2026-02-25 13:22:14	00:03:12	2026-02-25 13:25:13	2026-02-25 13:25:28
		2026-02-25 13:25:26		2026-02-25 13:25:28	2026-02-25 13:25:35
6	225131635	2026-02-25 13:22:14	00:03:15	2026-02-25 13:25:35	2026-02-25 13:25:44
		2026-02-25 13:25:29		2026-02-25 13:25:44	2026-02-25 13:25:47
7	225131636	2026-02-25 13:22:14	00:03:13	2026-02-25 13:25:47	2026-02-25 13:25:54
		2026-02-25 13:25:27		2026-02-25 13:25:55	2026-02-25 13:25:58
8	225131637	2026-02-25 13:22:14	00:02:46	2026-02-25 13:25:58	2026-02-25 13:26:06
		2026-02-25 13:25:00		2026-02-25 13:26:06	2026-02-25 13:26:09



Alibaba Cloud Hologres

TPC-H Rev.3.0.1
TPC Pricing Rev.2.9.0

Report Date:
4-March-26

TPC-H Timing Intervals (in seconds)

Duration of query execution:

Stream ID	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11	Q12
0	0.96	0.25	2.14	0.41	2.11	0.11	1.72	2.32	10.96	1.62	2.10	0.26
1	1.95	1.15	9.58	1.73	12.40	0.53	14.20	9.34	36.55	5.37	25.09	1.46
2	1.74	1.17	12.75	2.73	9.85	0.35	9.38	6.70	51.37	4.03	5.12	1.05
3	1.76	1.93	2.28	1.86	8.48	6.05	12.83	9.66	76.00	8.74	6.75	1.24
4	2.34	2.17	5.82	2.79	9.81	1.39	10.25	9.78	72.92	6.32	6.20	3.64
5	2.03	1.61	7.11	1.13	9.99	0.85	15.93	9.96	45.88	7.56	13.99	2.71
6	1.24	0.48	8.45	5.85	7.94	1.15	9.38	8.60	73.56	6.19	7.89	3.77
7	1.79	4.65	17.40	3.68	6.44	0.24	4.42	6.41	72.18	2.90	6.11	1.35
8	4.18	0.47	7.76	3.74	8.23	0.46	12.10	13.40	38.36	6.69	11.05	5.22
Minimum	0.96	0.25	2.14	0.41	2.11	0.11	1.72	2.32	10.96	1.62	2.10	0.26
Average	2.00	1.54	8.14	2.66	8.36	1.24	10.02	8.46	53.09	5.49	9.37	2.30
Maximum	4.18	4.65	17.40	5.85	12.40	6.05	15.93	13.40	76.00	8.74	25.09	5.22

Stream ID	Q13	Q14	Q15	Q16	Q17	Q18	Q19	Q20	Q21	Q22	RF1	RF2
0	4.20	0.46	0.31	0.52	0.42	4.33	1.33	0.82	1.40	1.89	7.56	6.01
1	9.32	1.94	1.67	4.09	2.28	9.63	3.72	12.00	8.45	6.25	24.78	18.18
2	8.49	2.45	6.56	4.07	5.85	11.66	1.90	19.99	5.99	5.98	34.45	16.82
3	5.33	1.75	6.51	1.71	2.17	17.27	4.56	5.26	3.61	9.11	14.50	11.94
4	11.27	2.55	3.81	3.93	5.02	11.75	8.39	3.20	5.42	3.61	29.37	28.24
5	10.39	2.74	2.05	9.47	2.27	17.15	4.06	6.72	9.39	8.07	14.93	6.47
6	16.81	0.56	2.52	2.79	0.89	8.50	3.37	6.71	1.90	15.24	9.06	3.45
7	6.91	0.78	1.15	1.90	2.08	8.40	6.12	10.66	8.94	18.04	7.47	3.18
8	7.08	2.02	2.56	4.33	1.63	13.97	2.67	4.39	7.89	7.27	8.41	3.03
Minimum	4.20	0.46	0.31	0.52	0.42	4.33	1.33	0.82	1.40	1.89	7.47	3.03
Average	8.87	1.69	3.02	3.65	2.51	11.41	4.01	7.75	5.89	8.38	16.73	10.81
Maximum	16.81	2.74	6.56	9.47	5.85	17.27	8.39	19.99	9.39	18.04	34.45	28.24

TPC BENCHMARK H OVERVIEW	10
0. GENERAL ITEMS	11
0.1 BENCHMARK SPONSOR	11
0.2 PARAMETER SETTINGS	11
0.3 CONFIGURATION DIAGRAM	11
1. CLAUSE 1: LOGICAL DATABASE DESIGN	14
1.1 DATABASE DEFINITION STATEMENTS	14
1.2 PHYSICAL ORGANIZATION	14
1.3 HORIZONTAL PARTITIONING	14
1.4 REPLICATION	14
2. CLAUSE 2: QUERIES AND REFRESH FUNCTIONS RELATED ITEMS	15
2.1 QUERY LANGUAGE	15
2.2 VERIFYING METHOD OF RANDOM NUMBER GENERATION	15
2.3 QUERY TEXT AND OUTPUT DATA FROM QUALIFICATION DATABASE	15
2.4 QUERY SUBSTITUTION PARAMETERS AND SEEDS USED	15
2.5 ISOLATION LEVEL	15
2.6 SOURCE CODE OF REFRESH FUNCTIONS	15
3. CLAUSE 3: DATABASE SYSTEM PROPERTIES	16
3.1 ACID PROPERTIES	16
3.2 ATOMICITY REQUIREMENTS	16
3.2.1 Atomicity of Completed Transactions	16
3.2.2 Atomicity of Aborted Transactions	16
3.3 CONSISTENCY REQUIREMENTS	17
3.3.1 Consistency Test	17
3.4 ISOLATION REQUIREMENTS	17
3.4.1 Isolation Test 1 – Read-Write Conflict with Commit	17
3.4.2 Isolation Test 2 – Read-Write Conflict with Rollback	17
3.4.3 Isolation Test 3 – Write-Write Conflict with Commit	18
3.4.4 Isolation Test 4 – Write-Write Conflict with Rollback	18
3.4.5 Isolation Test 5 – Concurrent Read and Write Transactions on Different Tables	18
3.4.6 Isolation Test 6 – Update Transactions during Continuous Read-Only Query Stream	18
3.5 DURABILITY REQUIREMENTS	19
3.5.1 Permanent Unrecoverable Failure of Any Durable Medium	20
3.5.2 System Crash	20
3.5.3 Memory Failure	20
3.5.4 Loss of External Power	20
3.5.5 Node or Controller Failure	20
4. CLAUSE 4: SCALING AND DATABASE POPULATION	22

4.2	INITIAL CARDINALITY OF TABLES.....	22
4.3	DISTRIBUTION OF TABLES AND LOGS ACROSS MEDIA.....	22
4.4	MAPPING OF DATABASE PARTITIONS/REPLICATION	22
4.5	IMPLEMENTATION OF DATA REDUNDANCY	23
4.6	DBGEN MODIFICATIONS.....	23
4.7	DATABASE LOAD TIME	23
4.8	DATA STORAGE RATIO	23
4.9	DATABASE LOAD MECHANISM DETAILS AND ILLUSTRATION.....	23
4.10	QUALIFICATION DATABASE CONFIGURATION.....	24
4.11	MEMORY TO DATABASE SIZE PERCENTAGE.....	24
5	CLAUSE 5: PERFORMANCE METRICS AND EXECUTION RULES RELATED ITEMS	25
5.2	SYSTEM ACTIVITY BETWEEN LOAD AND PERFORMANCE TESTS	25
5.3	STEPS IN THE POWER TEST	25
5.4	TIMING INTERVAL FOR EACH QUERY AND REFRESH FUNCTIONS	25
5.5	NUMBER OF STREAMS FOR THE THROUGHPUT TEST	25
5.6	START AND END DATE/TIME OF EACH QUERY STREAM	25
5.7	TOTAL ELAPSED TIME OF THE MEASUREMENT INTERVAL.....	25
5.8	REFRESH FUNCTIONS START DATE/TIME AND FINISH DATE/TIME.....	25
5.9	TIMING METRIC PRECISION	26
5.10	PERFORMANCE METRICS	26
5.11	THE PERFORMANCE METRIC AND NUMERICAL QUANTITIES FROM BOTH RUNS.....	26
5.12	SYSTEM ACTIVITY BETWEEN PERFORMANCE TESTS.....	26
5.13	DOCUMENTATION TO SATISFY CLAUSE 5.2.7	26
5.14	QUERY OUTPUT VALIDATION	26
6	CLAUSE 6: SUT AND DRIVER IMPLEMENTATION RELATED ITEMS.....	27
6.2	DRIVER.....	27
6.3	IMPLEMENTATION SPECIFIC LAYER (ISL).....	27
6.4	PROFILE-DIRECTED OPTIMIZATION	27
7	CLAUSE 7: PRICING	28
7.2	PRICED CONFIGURATION	28
7.3	AVAILABILITY DATE AND ORDERABILITY DATE	28
7.4	COUNTRY-SPECIFIC PRICING	28
8	CLAUSE 8: FULL DISCLOSURE.....	29
8.2	SUPPORTING FILES INDEX TABLE.....	29
9	CLAUSE 9: AUDIT RELATED ITEMS	30
9.2	AUDITOR'S REPORT	30
	APPENDIX A: PRICE QUOTES	31

TPC BENCHMARK H OVERVIEW

The TPC Benchmark™ H (TPC-H) is a Decision Support benchmark. It is a suite of business-oriented adhoc queries and concurrent modifications. The queries and the data populating the database have been chosen to have broad industry-wide relevance while maintaining a sufficient degree of ease of implementation. This benchmark illustrates Decision Support systems that:

- Examine large volumes of data
- Execute queries with a high degree of complexity
- Give answers to critical business questions

TPC-H evaluates the performance of various Decision Support systems by the execution of sets of queries against a standard database under controlled conditions. The TPC-H queries:

- Give answers to real-world business questions
- Simulate generated ad-hoc queries
- Are far more complex than most OLTP transactions
- Include a rich breadth of operators and selectivity constraints
- Generate intensive activity on the part of the database server component of the system under test
- Are executed against a database complying to specific population and scaling requirements
- Are implemented with constraints derived from staying closely synchronized with an on-line production database

Further information is available at www.tpc.org.

0. General Items

0.1 Benchmark Sponsor

A statement identifying the benchmark sponsor(s) and other participating companies must be provided.

This benchmark is sponsored by Alibaba Cloud Computing Ltd.

0.2 Parameter Settings

Settings must be provided for all customer-tunable parameters and options that have been changed from the defaults found in actual products, including but not limited to:

- *Database tuning options;*
- *Optimizer/Query execution options;*
- *Query processing tool/language configuration parameters;*
- *Recovery/commit options;*
- *Consistency/locking options;*
- *Operating system and configuration parameters;*
- *Configuration parameters and options for any other software component incorporated into the pricing structure;*
- *Compiler optimization options.*

In the event that some parameters and options are set multiple times, it must be easily discernible by an interested reader when the parameter or option was modified and what new value it received each time.

This requirement can be satisfied by providing a full list of all parameters and options, as long as all those that have been modified from their default values have been clearly identified and these parameters and options are only set once.

The Supporting File Archive contains the Operating System and DBMS parameters used in this benchmark.

0.3 Configuration Diagram

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.).*

The system diagram of the measured system is depicted in Figure 1. and the configuration of the measured systems is shown in Table 1. The system is composed of 69 compute containers, a Pangu File System service with storage capacity of 6,200 GB, and an Alibaba Cloud MaxCompute storage of 50 GB used to stage the refresh data set flat files. There is no difference between the priced and measured configurations.

Configuration	Gateway
#Gateway	2
Per Gateway Config	2 vCPUs
	8 GB RAM
	7,920 GB Local Disk (Shared)
Network	100 Gbits/s

Configuration	JobMaster
#JobMaster	1
Per JobMaster Config	4 vCPUs
	16 GB RAM
	7,920 GB Local Disk (Shared)
Network	100 Gbits/s

Configuration	MetaStore
#MetaStore	1
Per MetaStore Config	4 vCPUs
	16 GB RAM
	7,920 GB Local Disk (Shared)
Network	100 Gbits/s

Configuration	HoloMaster
#HoloMaster	1
Per HoloMaster Config	4 vCPUs
	16 GB RAM
	7,920 GB Local Disk (Shared)
Network	100 Gbits/s

Configuration	HoloWorker
#HoloWorker	64
Per HoloWorker Config	16 vCPUs
	64 GB RAM
	7,920 GB Local Disk (Shared)
Network	100 Gbits/s

Configuration	PanguStorage
Storage Capacity	6,200 GB

Configuration	MaxCompute Storage
Storage Capacity	50 GB

Table 1. Configuration of the measured system.

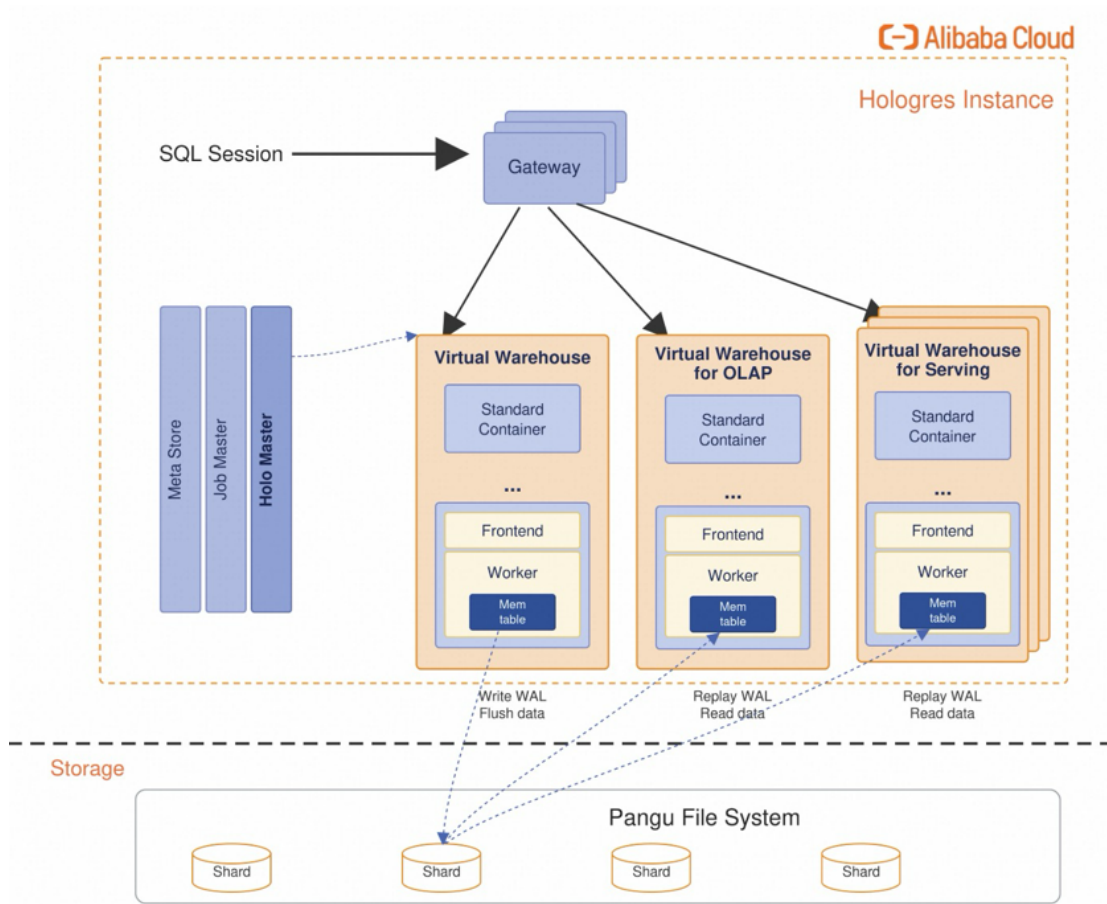


Fig. 1. System diagram of the measured configuration.

1. Clause 1: Logical Database Design

1.1 Database Definition Statements

Listings must be provided for all table definition statements and all other statements used to set-up the test and qualification databases. All listings must be reported in the supporting files archive.

The Supporting File Archive contains the table definitions and all other statements to set up the tables used in this benchmarking.

1.2 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 1.4, it must be noted. The physical organization of tables must be reported in the supporting files archive.

The concept of physical organization includes, but is not limited to: record clustering (i.e., rows from different logical tables are co-located on the same physical data page), index clustering (i.e., rows and leaf nodes of an index to these rows are co-located on the same physical data page), and partial fill-factors (i.e., physical data pages are left partially empty even though additional rows are available to fill them).

Physical organization requires no user input. No column reordering nor other physical alteration to the database is used.

1.3 Horizontal Partitioning

Horizontal partitioning of tables and rows in the test and qualification databases (see Clause 1.5.4) must be disclosed. Scripts to perform horizontal partitioning must be reported in the supporting files archive.

The tested database consists of 320 data shards. The mapping between data shards and compute instances (holo-workers) is managed by the DBMS automatically. Tables are horizontally distributed on data shards using a hash algorithm. The columns used for the hashing (e.g., distribution key) are controlled by DDL statements (see Supporting Files Archive, keyword: `distribution_key`).

1.4 Replication

Any replication of physical objects must be disclosed and must conform to the requirements of Clause 1.5.7. Scripts to perform any replication must be reported in the supporting files archive.

No replication was used.

2. Clause 2: Queries and Refresh Functions Related Items

2.1 Query Language

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

SQL was the query language used to implement the queries.

2.2 Verifying Method of Random Number Generation

The version number, release number, modification number, and patch level of QGen must be disclosed. Any modifications to the QGen source code must be reported in the supporting files archive.

TPC-supplied DBGen version 3.0.1 and QGen version 3.0.1 were used.

2.3 Query Text and Output Data from Qualification Database

The executable query text used for query validation must be reported in the supporting files archive along with the corresponding output data generated during the execution of the query text against the qualification database. If minor modifications (see Clause 2.2.3) 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 actual query text and query output are included in the Supporting Files Archive. The standard queries were used throughout with the following modifications:

- LIMIT syntax used to restrict the number of output rows (Q2, Q3, Q10, Q18, Q21).

2.4 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 seed and query substitution parameters.

2.5 Isolation Level

The isolation level used to run the queries must be disclosed. If the isolation level does not map closely to one of the isolation levels defined in Clause 3.4, additional descriptive detail must be provided.

The queries and transactions were run with isolation level 3.

2.6 Source Code of Refresh Functions

The details of how the refresh functions were implemented must be disclosed (including source code of any non-commercial program used).

Supporting Files Archive contains the Source Code of the refresh functions.

3. Clause 3: Database System Properties

The results of the ACID tests must be disclosed along with a description of how the ACID requirements were met. All code (including queries, stored procedures etc.) used to test the ACID requirements and their entire output must be reported in the supporting files archive.

The results of the ACID tests are disclosed as required.

3.1 ACID Properties

The ACID (Atomicity, Consistency, Isolation, and Durability) properties of transaction processing systems must be supported by the system under test during the timed portion of this benchmark. Since TPC-H is not a transaction processing benchmark, the ACID properties must be evaluated outside the timed portion of the test.

All ACID tests were conducted according to the specifications. The Supporting Files Archive contains the source code of the ACID test scripts.

3.2 Atomicity Requirements

The system under test must guarantee that transactions are atomic; the system will either perform all individual operations on the data, or will assure that no partially completed operations leave any effects on the data.

3.2.1 Atomicity of Completed Transactions

Perform the ACID Transaction for a randomly selected set of input data and verify that the appropriate rows have been changed in the ORDERS, LINEITEM, and HISTORY tables.

The following operations were performed to verify the atomicity of the completed transactions:

1. Randomly selected one order key, and then retrieved the total price from table ORDERS and the extended price from table LINEITEM.
2. One transaction was performed using the selected order key.
3. The transaction was committed.
4. Retrieved the total price from the table ORDERS and the extended price from the table LINEITEM for the order key.
5. It was verified that the rows had been changed.

3.2.2 Atomicity of Aborted Transactions

Perform the ACID Transaction for a randomly selected set of input data, substituting a ROLLBACK of the transaction for the COMMIT of the transaction. Verify that the appropriate rows have not been changed in the ORDERS, LINEITEM, and HISTORY tables.

The following operations were performed to verify the atomicity of the aborted transactions:

1. Randomly selected one order key, and then retrieved the total price from table ORDERS and the extended price from table LINEITEM.
2. One transaction was performed using the selected order key and the transaction was stopped before commit.
3. The transaction was rolled back.
4. Retrieved the total price from the table ORDERS and the extended price from the table LINEITEM for the order key.
5. It was verified that the rows had not been changed.

3.3 Consistency Requirements

Consistency is the property of the application that requires any execution of transactions to take the database from one consistent state to another. A consistent state for the TPC-H database is defined to exist when:

$O_TOTALPRICE = SUM(trunc(trunc(L_EXTENDEDPRICE(1 - L_DISCOUNT), 2) * (1 + L_TAX), 2))$ For each ORDER and LINEITEM defined by ($O_ORDERKEY = L_ORDERKEY$).*

3.3.1 Consistency Test

Verify that ORDERS and LINEITEM tables are initially consistent, submit the prescribed number of ACID Transactions with randomly selected input parameters, and re-verify the consistency of the ORDERS and LINEITEM.

The following query was executed before and after the consistency tests to demonstrate that the database kept staying in a consistent state both initially and after submitting transactions:

```
select * from (  
select o_orderkey, (o_totalprice - sum( truncate((truncate((l_extendedprice * (1-l_discount)), 2)*(1+l_tax),2))) part_res  
from orders, lineitem  
where o_orderkey=l_orderkey group by o_orderkey, o_totalprice  
) temp where not part_res=0;
```

The following operations were performed to verify the consistency of ACID transactions:

1. The consistency of the ORDERS and LINEITEM tables was verified.
2. For each of the 11 execution streams, 100 transactions were prepared.
3. For all 11 execution streams, the prepared ACID transactions were executed.
4. The consistency of the ORDERS and LINEITEM tables was verified again.

3.4 Isolation Requirements

Operations of concurrent transactions must yield results, which are indistinguishable from the results, which would be obtained by forcing each transaction to be serially executed to completion in some order.

The steps of the isolation tests were adapted to the Hologres isolation environment.

3.4.1 Isolation Test 1 – Read-Write Conflict with Commit

Demonstrate isolation for the read-write conflict of a read-write transaction and a read-only transaction when the read-write transaction is committed.

The following operations were performed to verify the isolation for a read-only and a read-write committed transaction:

1. Started a query and verified that the row was retrieved.
2. Started an update transaction, read and updated the same row. Stalled before commit.
3. Started the same query and verified that the row has not changed.
4. Committed the update transaction.
5. Started a query and verified that the new row was retrieved.

3.4.2 Isolation Test 2 – Read-Write Conflict with Rollback

Demonstrate isolation for the read-write conflict of a read-write transaction and a read-only transaction when the read-write transaction is rolled back.

The following operations were performed to verify the isolation for a read-only and a rolled back read-write transaction:

1. Started a query and verified that the row was retrieved.
2. Started an update transaction, read and updated the same row. Stalled before commit.
3. Started the same query and verified that the row has not changed.
4. Rolled back the update transaction.
5. Started a query and verified that the old row was retrieved.

3.4.3 Isolation Test 3 – Write-Write Conflict with Commit

Demonstrate isolation for the write-write conflict of two update transactions when the first transaction is committed.

The following operations were performed to verify the isolation of two update transactions:

1. Started an update transaction T1, and stopped T1 immediately prior to COMMIT.
2. Started an update transaction T2. T2 was blocked.
3. Committed T1. T2 should now complete and commit too.
4. Compared extendedprice after update in T1 and extendedprice before update in T2. They should be same.

3.4.4 Isolation Test 4 – Write-Write Conflict with Rollback

Demonstrate isolation for the write-write conflict of two update transactions when the first transaction is rolled back.

The following operations were performed to verify isolation of two update transactions after the first one is rolled back:

1. Started an update transaction T1, and stopped T1 immediately prior to COMMIT.
2. Started an update transaction T2. T2 was blocked.
3. Rolled back T1. T2 should now complete and commit, too.
4. Compared extendedprice before update in T1 and extendedprice before update in T2. They should be same.

3.4.5 Isolation Test 5 – Concurrent Read and Write Transactions on Different Tables

Demonstrate the ability of read and write transactions affecting different database tables to make progress concurrently.

The following operations were performed to demonstrate the ability of read and write transactions involving different tables to make progress concurrently:

1. Started a query and verified that the row was retrieved.
2. Started an update transaction, read and updated the same row. Stalled before commit.
3. Started another transaction that performed the following operation: Select random values of PS_PARTKEY AND PS_SUPPKEY. Return all columns of the PARTSUPP table for which PS_PARTKEY and PS_SUPPKEY are equal to the selected values.
4. Verified that the read transaction completed.
5. Committed the update transaction.
6. Started the same query and verified that the new row was retrieved.

3.4.6 Isolation Test 6 – Update Transactions during Continuous Read-Only Query Stream

Demonstrate the continuous submission of arbitrary (read-only) queries against one or more tables of the database does not indefinitely delay update transactions affecting those tables from making progress.

The following query was used to ensure sufficient execution to perform the test:

```
SELECT
  11.1_orderkey,
  count(*) AS cnt,
  sum(11.1_extendedprice * (1 - 11.1_discount) * (1 + 11.1_tax)) AS revenue
FROM
  lineitem 11,
  lineitem 12,
  lineitem 13,
  lineitem 14,
  lineitem 15
WHERE
  11.1_shipdate <= date '1998-12-01'
  AND 11.1_orderkey = 12.1_orderkey
  AND 11.1_linenumber = 12.1_linenumber
  AND 11.1_extendedprice = 13.1_extendedprice
  AND 13.1_quantity < 30
  AND 14.1_quantity = 11.1_quantity
  AND 14.1_orderkey < 20000
  AND 15.1_receiptdate = 11.1_receiptdate
  AND 15.1_partkey < 2000
GROUP BY
  11.1_orderkey
ORDER BY
  revenue DESC,
  11.1_orderkey;
```

1. A transaction, T1, which executed the above query on the qualification database, was started using a random DELTA.
2. A transaction, T2, was started for a randomly selected O KEY, L KEY, and DELTA.
3. T2 completed and appropriate rows in the tables ORDERS, LINEITEM, and HISTORY had been changed.
4. T1 was still executing.
5. Transaction T1 completed executing the query.

3.5 Durability Requirements

The tested system must guarantee durability: the ability to preserve the effects of committed transactions and ensure database consistency after recovery from any one of the failures listed in Clause 3.5.3.

The following steps were performed for the durability test:

1. The consistency of the ORDERS and LINEITEM tables was verified.
2. 2000 transactions for each of the 9 execution streams were prepared.
3. After at least 100 ACID transactions were completed by each of the 9 execution streams.

4. A durability failure was induced (see details for each failure shown below).
5. The consistency of the ORDERS and LINEITEM tables was verified again.
6. The durability success files were compared with the HISTORY table.

All durability tests were performed on the cluster shown in Section 0.3.

3.5.1 Permanent Unrecoverable Failure of Any Durable Medium

Guarantee the database and committed updates are preserved across a permanent irrecoverable failure of any single durable medium containing TPC-H database tables or recovery log tables.

The system crash, memory failure and node or controller failure tests were performed together as shown in Section 3.5.5.

3.5.2 System Crash

Guarantee the database and committed updates are preserved across an instantaneous interruption (system crash/system hang) in processing which requires the system to reboot to recover.

The system crash, memory failure and node or controller failure tests were performed together as shown in Section 3.5.5.

3.5.3 Memory Failure

Guarantee the database and committed updates are preserved across failure of all or part of memory (loss of contents). See the previous section.

The system crash, memory failure and node or controller failure tests were performed together as shown in Section 3.5.5.

3.5.4 Loss of External Power

Loss of External Power: Guarantee the database and the effects of committed updates are preserved during the loss of all external power to the SUT for an indefinite time period.

Each Hologres compute node, PanguFS node and PanguFS Meta node are deployed on Alibaba Cloud IDC and Alibaba Cloud IDC is configured with redundant power supplies and Uninterruptible Power Supply (UPS). The UPS capabilities are sufficient to keep the entire SUT running for a period of at least 30 minutes in the event of a total loss of all external power.

In addition, various roles of single node crash were performed as described in section 3.5.5.

3.5.5 Node or Controller Failure

Guarantee the database and committed updates are preserved across failure of the controller or the whole node.

The tested cluster consists of 2 gateways, 1 job-master, 1 meta-store, 1 holo-master and 64 holo-worker instances and 6,200GB PanguFS storage. Five tests were performed to test the durability:

1. Shutdown the host machine of one of the computing Holo-Worker instances. The full cluster was restored automatically after the Holo-Worker container was started automatically.
2. Shutdown the host machine of the Gateway container. The full cluster was restored automatically after the Gateway container was started automatically.

3. Shutdown the host machine of the Job-Master container. The full cluster was restored automatically after the Gateway container was started automatically.
4. Shutdown the host machine of the Meta-Store container. The full cluster was restored automatically after the Meta-Store container was started automatically.
5. Shutdown the host machine of the Holo-Master container. The full cluster was restored automatically after the Holo-Master container was started automatically.

4 Clause 4: Scaling and Database Population

4.1 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 4.2.5), must be disclosed.

Table 2 lists the TPC Benchmark™ H defined tables and the row count for each table as they existed upon completion of the test database build.

Table	#Rows
Lineitem	18,000,048,306
Order	4,500,000,000
Partsupp	2,400,000,000
Part	600,000,000
Customer	450,000,000
Supplier	30,000,000
Nation	25
Region	5

Table 2. Table Cardinalities.

4.2 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 measured and priced configurations.

Server Node	Disk Type	Disk Drive	Content
Gateway	Local NVMe SSD	/dev/nvme0n1, /dev/nvme1n1	Event logs, temp files
MetaStore			
HoloMaster			
HoloWorker			
JobMaster			

Table 3. Distribution of tables and logs

All the base tables and transaction logs (WAL) were stored on the PanguFS storage, the sizes of tables are shown in Table 4.

Table Name	Lineitem	Orders	Partsupp	Part	Customer	Supplier	Nation	Region
Table Size	1,168 GB	340 GB	234 GB	91 GB	62 GB	10 GB	87 KB	19 KB

Table 4. Table Size on Pangu Storage

4.3 Mapping of Database Partitions/Replication

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

The intent is to provide sufficient detail about partitioning and replication to allow independent reconstruction of the test database.

Tables are horizontally sharded using a hash algorithm. The columns used for the hashing (e.g., distribution key) are controlled by DDL statements (see Supporting Files Archive, keyword: distributed_key). In this benchmarking, the data volume is partitioned into 320 shards. Each shard corresponds to one physical data directory. Hologres

automatically manages the physical mapping between shards and holo-workers. Shards are evenly distributed over the holo-workers.

4.4 Implementation of Data Redundancy

Implementations may use data redundancy mechanism(s). The type of data redundancy mechanism(s) and any configuration parameters (e.g., RAID level used must be disclosed for each device). If data redundancy mechanism(s) are used in an implementation, the logical intent of their use must be disclosed.

RAID is not used. The Data Redundancy is internally managed by Pangu FS, a distributed file system, which uses multiple replicas and erasure code to ensure high availability for data access service.

4.5 DBGen Modifications

The version number, release number, modification number, and patch level of DBGen must be disclosed. Any modifications to the DBGen (see Clause 4.2.1) source code (see Appendix D) must be reported in the supporting files archive.

The supplied DBGen version 3.0.1 was used, no modifications were made.

4.6 Database Load Time

The database load time for the test database (see Clause 4.3) must be disclosed.

See Numerical Quantities Summary in the Executive Summary.

4.7 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 the size chosen for the test database as defined in Clause 4.1.3.1. Let r be the ratio. The reported value for r must be rounded to the nearest 0.01. That is, reported value= $\text{round}(r,2)$. 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.

For the reporting of configured disk capacity, gigabyte (GB) is defined to be 2^{30} bytes.

Server Node	Total Storage Capacity (GB)
Pangu FS	5,774.20
Hologres Compute Node	162,273.64
MaxCompute	46.57

Table 5. Data Storage Capacity

Total disk capacity: 168,094.41 GB Scale factor: 3,000.

The database storage ratio is 56.03.

4.8 Database Load Mechanism Details and Illustration

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

Figure 2 shows the database build procedure. The raw data flat files were created using DBGen and stored on Alibaba Cloud MaxCompute. The configuration for loading data from Alibaba Cloud MaxCompute is disclosed in the Supporting Files Archive. The refresh data set flat files generated by DBGen are also staged on Alibaba Cloud MaxCompute. The MaxCompute storage holding the refresh data set flat files is included in the priced configuration.

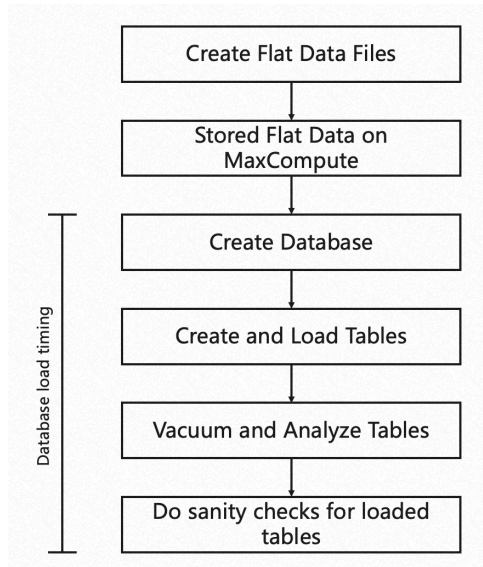


Fig. 2. Block Diagram of Database Loading Procedure

4.9 Qualification Database Configuration

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

The qualification database used identical scripts to create and load the data with changes to adjust for the database scale factor.

4.10 Memory to Database Size Percentage

The memory to database size percentage must be disclosed. It is computed by multiplying by 100 the total memory size priced on the SUT (see clause 6.2.1) and dividing this number by the size chosen for the test database as defined in Clause 4.1.3.1. Let r be this ratio. The reported ratio must be rounded to the nearest 0.1. That is, reported value= $\text{round}(r,1)$. For example, a system configured with 256GB of memory for a 1000GB test database has a memory/database size percentage of 25.6.

Available Memory: 4,160 GB Scale Factor: 3,000.

The memory to database size percentage is 138.7%.

5 Clause 5: Performance Metrics and Execution Rules Related Items

5.1 System Activity between Load and Performance Tests

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 reported in the supporting files archive including listings of scripts, command logs and system activity.

There is no activity on the SUT between the conclusion of the load test and the beginning of the performance test.

5.2 Steps in the Power Test

The details of the steps followed to implement the power test (e.g., system boot, database restart, etc.) must be reported in the supporting files archive.

The following steps were used to implement the power test:

- RF1 Refresh Function from the refresh stream.
- Query Execution from query stream 0.
- RF2 Refresh Function from the refresh stream.

5.3 Timing Interval for Each Query and Refresh Functions

The timing intervals (see Clause 5.3.7) for each query and for both refresh functions must be reported for the power test. The output for each query and for both refresh functions must be reported in the supporting files archive.

See the Numerical Quantities Summary in the Executive Summary at the beginning of this report.

5.4 Number of Streams for the Throughput Test

The number of query streams used for the throughput test must be disclosed.

8 query streams were used for the throughput test.

5.5 Start and End Date/Time of Each Query Stream

The start time and finish time for each query stream for the throughput test must be disclosed. The output for each query stream for the throughput test must be reported in the supporting files archive.

See the Numerical Quantities Summary in the Executive Summary at the beginning of this report.

5.6 Total Elapsed Time of the Measurement Interval

The total elapsed time of the measurement interval (see Clause 5.3.6) must be disclosed for the throughput test.

See the Numerical Quantities Summary in the Executive Summary at the beginning of this report.

5.7 Refresh Functions Start Date/Time and Finish Date/Time

The start time and, finish time for each refresh function in the refresh stream for the throughput test must be disclosed. The output of each refresh function in the refresh stream for the throughput test must be reported in the supporting files archive.

See the Numerical Quantities Summary in the Executive Summary at the beginning of this report.

5.8 Timing Metric Precision

The start time and finish time for each query and refresh stream shall be reported to the hundredth of a second. If times are measured with the precision greater than one hundredth of a second, the reported times shall be truncated to the hundredth of a second.

See the Numerical Quantities Summary in the Executive Summary at the beginning of this report.

5.9 Performance Metrics

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

See the Numerical Quantities Summary in the Executive Summary at the beginning of this report.

5.10 The Performance Metric and Numerical Quantities from Both Runs

The performance metric (QphH@Size) and the numerical quantities (TPC-H Power@Size and TPC-H Throughput@Size) from both of the runs must be disclosed.

Run ID	QphH@3,000GB	QppH@3,000GB	QthH@3,000GB
Run 1	8,710,708.5	8,977,993.8	8,451,380.6
Run 2	8,443,627.5	8,789,205.2	8,111,637.4

Table 6. Performance Metric

5.11 System Activity between Performance Tests

Any activity on the SUT that takes place between the conclusion of Run1 and the beginning of Run2 must be fully disclosed including system activity, listings of scripts or command logs along with any system reboots or database restarts.

There was no activity between Run 1 and Run 2.

5.12 Documentation to satisfy Clause 5.2.7

All documentation necessary to satisfy Clause 5.2.7 must be made available upon request.

Hologres document can be obtained from <https://www.alibabacloud.com/>.

5.13 Query Output Validation

The output of the Query Output Validation Test must be reported in the supporting files archive.

The output of the validation test is available in the Supporting Files Archive.

6 Clause 6: SUT and Driver Implementation Related Items

6.1 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 and all related source code, scripts and configuration files must be reported in the supporting files archive. The information provided should be sufficient for an independent reconstruction of the driver.

All stream executions are performed by a script. QGEN is used to produce query text. For each power-test run:

1. A shell script is started, executes RF1 and then waits for the query stream to complete.
2. A shell script is started, executes the 22 queries in the required order for stream 0 and then signals to the shell script started in step 1.
3. The shell script started in step 1 is released and executes RF2.

For each throughput-test run:

1. The queries as generated by QGEN are submitted in the order defined by Clause 5.3.5.4 from the driver in several streams (the number of streams is listed in the Numerical Quantities) and runs concurrently.
2. Then throughput RF function pairs are scheduled to run sequentially in one update stream.

The source code of the used scripts are disclosed in the Supporting Files Archive.

6.2 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 disclosed. All related source code, scripts and configuration files must be reported in the supporting files archive. The information provided should be sufficient for an independent reconstruction of the implementation specific layer.

The scripts used to implement the ISL are available in the Supporting Files Archive.

6.3 Profile-Directed Optimization

If profile-directed optimization as described in Clause 5.2.9 is used, such use must be disclosed. In particular, the procedure and any scripts used to perform the optimization must be reported in the supporting files archive.

Profile-directed optimization was not used.

7 Clause 7: Pricing

7.1 Priced Configuration

The pricing methodology used for pricing the Priced Configuration is the "Default 3-Year Pricing Methodology", as defined in the current revision of the TPC Pricing specification.

The system to be priced shall include the hardware, Licensed Compute Services and software components present in the System Under Test (SUT), a communication interface that can support user interface devices, additional operational components configured on the test system, and maintenance on all of the above:

- *System Under Test.*
- *User Interface Devices and Communications.*
- *Database Storage and Recovery Log*
- *Additional Operational Components.*
- *Software.*

A 3-Year Pricing Methodology is used for this benchmark. A detailed list of hardware and software used in the priced system is included in the pricing sheet in the Executive Summary at the beginning of this report. The price quotations are included in Appendix A.

7.2 Availability Date and Orderability Date

The committed delivery date for general availability of products used in the price calculations must be reported. When the priced system includes products with different availability dates, the availability date reported on the executive summary must be the date by which all components are committed to being available. The full disclosure report must report availability dates individually for at least each of the categories for which a pricing subtotal must be provided.

All components of the priced configuration will be available for purchase on 03/04/2026.

For each of the components that are not orderable on the report date of the FDR, the following information must be included in the FDR:

- *Name and part number of the item that is not orderable*
- *The date when the component can be ordered (on or before the Availability Date)*
- *The method to be used to order the component (at or below the quoted price) when that date arrives*
- *The method for verifying the price*

All priced components will be orderable on 03/04/2026.

7.3 Country-Specific Pricing

The Priced Locale and Priced Currency of the Priced Configuration must be disclosed.

The configuration is priced in CNY for the China market.

8 Clause 8: Full Disclosure

8.1 Supporting Files Index Table

An index for all files and/or directories included in the Supporting Files Archive as required by Clauses 8.3.2 through 8.3.8 must be provided in the report.

Clause	Description	Archive Files	Path Name
1	Parameter Settings	benchmark_scripts.zip	RUN/params.log
	DB Creation Scripts		KIT/sql/create_user.sql KIT/sql/create_schema.sql KIT/sql/analyze_database.sql
	System Verification		KIT/scripts/tools/hwinfo.sh KIT/scripts/tools/podinfo.sh
	Toolkit Common Scripts		KIT/scripts/
2	Minor query modifications	benchmark_scripts.zip	KIT/tpch_archives/TPC-H_Tools_v3.0.0.zip.patch
3	ACID Test Scripts	benchmark_scripts.zip	KIT/scripts/ACID/
	ACID Test Results		ACID/
4	Database Load Scripts	benchmark_scripts.zip	KIT/scripts/load_init.sh
	Qualification Test Results		VLD/
5	Query Output Results	run1results.zip	run1/
		run2results.zip	run2/
6	Source Code and Scripts of Driver	benchmark_scripts.zip	KIT/scripts/query_streams/
7	There are no files to be included for Clause 7.	N/A	N/A
8	Query Parameters	run1results.zip	run1/substitution_parameters.txt
	Executable Query Text		run1/stream*.sql
	RF function source code	benchmark_scripts.zip	KIT/scripts/tpc_h_full.sh

Table 7. Supporting File Index Table

9 Clause 9: Audit Related Items

9.1 Auditor's Report

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 who to contact in order to obtain further information regarding the audit process.

This implementation of the TPC Benchmark™ H was audited by Doug Johnson of InfoSizing, a certified TPC- H auditor. Further information regarding the audit process may be obtained from:

Doug Johnson
InfoSizing (www.sizing.com)

63 Lourdes Dr
Leominster, MA 01453
(978) 343-6562

TPC Benchmark™ H Full Disclosure Report and other information can be downloaded from the Transaction Processing Performance Council web site at www.tpc.org.

Benchmark sponsor: Jiahui Zhang
Alibaba Group, Alibaba Cloud
Yungu Park, No. 1008 Dengcai Street
Xihu District, Hangzhou, Zhejiang, China

June 24, 2026

I verified the TPC Benchmark H (TPC-H™ v3.0.1) performance of the following configuration:

Platform: Alibaba Cloud Hologres (with 69 Compute Containers)
Operating System: Alibaba Cloud Linux 3.2104 U10 (OpenAnolis Edition)
Database Manager: Hologres V4.1
Other Software: Pangu File System V2.0

The results were:

Performance Metric	8,443,627.5 QphH@3,000GB
TPC-H Power	8,789,205.2
TPC-H Throughput	8,111,637.4
Database Load Time	00d 00h 20m 31s

Server	<u>Alibaba Cloud Hologres (69 Compute Containers) with:</u>
vCPUs	1,040
Memory	4,160 GiB
Storage	174,240 GB Node Local Storage 6,200 GB Pangu File System V2.0 50 GB MaxCompute

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 DBGen
- The database was properly scaled to 3,000GB and populated accordingly
- The compliance of the database auxiliary data structures was verified
- The database load time was correctly measured and reported
- The required ACID properties were verified and met

- The query input variables were generated by QGen
- The query text was produced using minor modifications and no query variant
- The execution of the queries against the SF1 database produced compliant answers
- A compliant implementation specific layer was used to drive the tests
- The throughput tests involved 8 query streams
- The ratio between the longest and the shortest query was such that no query timings were adjusted
- The execution times for queries and refresh functions were correctly measured and reported
- The repeatability of the measured results was 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 cursive script that reads "Doug Johnson". The signature is written in black ink and has a long, horizontal flourish extending to the right.

Doug Johnson, TPC Certified Auditor

Appendix A: Price Quotes

The screenshot displays the configuration interface for provisioning Alibaba Cloud Hologres. Key settings include:

- Region:** China (Beijing)
- Specifications:** Virtual Warehouse Type: General-purpose
- Zone:** Zone L
- Gateway Nodes:** 2
- Reserved Computing Resources of Virtual Warehouse:** 1924 CU
- Storage Redundancy Type:** ZRS
- Standard:** 6200 GB
- Infrequent Access:** 0 GB
- Backup Set Storage:** Resources for storing periodic backup data.
- VPC:** vpc-2ze6d9f1gkhtoyvnnz / huadong2_vpc
- vSwitch:** vsw-2ae0289e20e3m3y8lch / beijing_LL_vsw

Expense Details:

List Price (¥6,581,520.00)	
Discounted Price	¥3,290,760.00
Gateway Nodes	¥12,240.00
Reserved Computing Resources of Virtual Warehouse	¥3,133,440.00
Standard	¥145,080.00

Amount Payable: **¥3,290,760.00**

Fig 3. Purchase Page for provisioning the tested Alibaba Cloud Hologres with 3-Year Subscription

The screenshot shows the pricing page for MaxCompute storage. The main table details the following:

存储类型	计费公式	价格	说明
单 AZ 存储-标准存储	日存储费用=日存储量平均值×单价	<ul style="list-style-type: none"> 一般情况下，计费单价如下： <ul style="list-style-type: none"> 公共云：0.004 元/GB/天。 金融云：0.0076 元/GB/天。 当项目的实际数据存储量大于 0 MB 小于等于 512 MB 时，计费单价如下： <ul style="list-style-type: none"> 公共云：0.01 元/天。 金融云：0.019 元/天。 	日存储量平均值=采集到的存储量之和÷24
单 AZ 存储-低频存储	<ul style="list-style-type: none"> 日存储费用=日存储量平均值×单价 数据访问费=数据访问量×单价 	<ul style="list-style-type: none"> 低频存储计费单价如下： <ul style="list-style-type: none"> 公共云：0.00266 元/GB/天。 金融云：0.00505 元/GB/天。 低频存储数据访问费单价如下： <ul style="list-style-type: none"> 公共云：0.04 元/GB。 金融云：0.076 元/GB。 	<ul style="list-style-type: none"> 日存储量平均值=采集到的存储量之和÷24 当低频存储的表或者分区转成标准存储或者长期存储时，会收取当前整张表或者分区的低频存储访问费。
单 AZ 存储-长期存储	日存储费用=日存储量平均值×单价	<ul style="list-style-type: none"> 长期存储计费单价如下： <ul style="list-style-type: none"> 公共云：0.0011 元/GB/天。 金融云：0.0021 元/GB/天。 长期存储数据访问费单价如下： 	<ul style="list-style-type: none"> 当长期存储的表或者分区转成标准存储或者低频存储时，会收取当前整张表或者分区的长期存储访问费。


Fig 4. Alibaba Cloud MaxCompute storage pricing page (standard storage, 0.004 CNY/GB/day).

天猫 便携轻薄笔记本电脑 Port 搜索 Search our store

淘宝网首页 Taobao homepage 已买到的宝贝 Bought baby 我的淘宝 My Taobao 购物车 Shopping cart 收藏夹 Favorites 免费开店 Free store opening 千牛卖家中心 QianNiu Seller Center 帮助中心 Help Center

宁美官方旗舰店 Ningmei official flagship store

店铺评分 4.8 88VIP好评率90% 88VIP praise rate: 90%



视频 video 图集 Gallery 参数 Parameters

用户评价 User testimonials 参数信息 Parameter information 图文详情 Graphic details 本店推荐 Our store reco

用户评价 · 76 User reviews · 76

客服很耐心 11 Customer service is very patient 11 用起来很方便 7 Easy to use? 方便携带 6 Easy to carry6

质量很好 6 Good quality6 性价比很高 5 Great value for money5 必须好评 4 Must praise 4

清潭畔666 Ye-Ye 666

2026-01-19 已购: D4内存大屏笔记本N150 / 1TB / 16GB

2026-01-19 Purchased: D4 memory large screen notebook N150 / 1TB / 16GB

机器不错,是我喜欢的类型,性能够用,运行流畅 The machine is good, it is the type I like, the sex can be used, and the running process is good

选购优选 宁美国度新款笔记本电脑15.6英寸轻薄本酷睿i5 1240P/英特尔N150处理器DDR5内存便携家用办公游戏本炒股正品...

已售 300+ 300+ sold 可开发票 Invoiceable 多人评价 客服很耐心 Many people commented in

券后 After the ticket ¥ **1999** 优惠券 Voucher 立减立减 The whole audience stms

原价 ¥ 2399

超级立减400元 Super instant discount of 400 yuan

48小时内发货,预计后天送达 Shipped within 48 hours, expected to arrive the day after tomorrow

快递: 免运费 Express Shipping: Free shipping

湖北省 武汉市 余杭区

湖北鄂州 至 杭州市 余杭区

Hubei Ezhou to Hangzhou Yuhang District

7天无理由 退换货 88VIP送运费险 赠一年四 极速退款 7天无理由退换

7 days of price guarantee return treasure 88VIP return package freight fake one

信用卡支付 对公支付

颜色分类 Color classification 切换大图模式 Toggle the large picture mode

D5内存轻薄本N150 D5 memory thin and light notebook N150

D4内存大屏笔记本N150 D4 memory large screen notebook N150

D4内存大屏笔记本N5095 D4 memory large screen notebook N5095

D5内存大屏笔记本i5 1240P D5 memory large screen notebook i5 1240P

硬盘容量 Hard disk capacity

256GB 512GB 1TB

内存容量 Memory capacity

16GB

数量 quantity

- 1 + 即将售罄 Soon to be sold out

送装服务 选购更多 Shop more

加入购物车 Add to cart 领券购买 Get a coupon to buy

Fig. 5. Ningmei 15.6-inch Slim Notebook tablet purchase page (Google translated English version).