

TPC Benchmark™ H Full Disclosure Report

Alibaba Cloud Hologres
(with 459 Compute Containers)
using
Hologres V1.3
and
Pangu File System V2.0
Alibaba Group Enterprise Linux Server 7.2 (Paladin)

First Edition
Oct 31, 2022



First Edition – Oct 31, 2022

Hologres V1.3

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 F53 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 Group Enterprise Linux Server release 7.2 (Paladin).

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 kQphH 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.2.8.0 | |
| | | | Report Date: 31-Oct-22 | |
| Total System Cost | Composite Query per Hour Metric | | Price/Performance | |
| 44,358,597 CNY | 27,865,885.7 QphH@30000GB | | 1,591.87 CNY Price/kQphH@30000GB | |
| Database Size | Database Manager | Operating System | Other Software | Availability Date |
| 30,000 GB | Hologres V1.3 | Alibaba Group Enterprise Linux Server 7.2 (Paladin) | Pangu File System V2.0 | 10/31/2022 |

5.03 - Power Test Geometric Mean of Query & Refresh Times

24.84 - Throughput Arithmetic Mean of Query Times





Legend:

- Power Run (light gray bar)
- Throughput Run (black bar)
- Geometric Mean of Power Test (dashed line at 5.03)
- Arithmetic Mean of Throughput Test (solid line at 24.84)

| | |
|--|--|
| Database Load Time: 00:58:14 Load Includes Backup: N Total Data Storage/Database Size: 53.70 Memory to Database Size Percentage: 293.3% | 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 | | | | | | |
|----------------------|------------------|-----------|------------|------------|--------------|-----------|
| | HoloFrontend | MetaStore | HoloMaster | HoloWorker | System Total | Pangu FS |
| #Container | 1 | 1 | 1 | 456 | 459 | N/A |
| vCPUs | 24 | 18 | 14 | 24 | 11,000 | N/A |
| Memory | 192 GB | 144 GB | 112 GB | 192 GB | 88,000 GB | N/A |
| Storage | 3,673 GB | 3,673 GB | 3,673 GB | 3,673 GB | 1,685,907 GB | 44,000 GB |
| Total | | | | | | |
| Storage: | 1,611,101.45 GiB | | | | | |

|  | Alibaba Cloud Hologres | | | TPC-H Rev.3.0.1 TPC Pricing Rev.2.8.0 | | |
|---|---------------------------|--------|------------------|--|----------------------|-------------------------|
| | | | | Report Date: 31-Oct-22 | | |
| Description | Part Number | Source | Unit Price (CNY) | Qty | Extended Price (CNY) | 3 yr Maint. Price (CNY) |
| License Computer and Software Service | | | | | | |
| Hologres Computing Resources | | | | | | |
| 11,000 vCPUs, 88,000GB memory, 3 years prepay | | | | | | |
| 1 holo-frontend, each with | | | | | | |
| 24 vCPUs, | | | | | | |
| 192 GB RAM, | | | | | | |
| 1 meta-store, each with | | | | | | |
| 18 vCPUs, | | | | | | |
| 144 GB RAM | | | | | | |
| 1 holo-master, each with | | | | | | |
| 14 vCPUs, | | | | | | |
| 112 GB RAM | | | | | | |
| 456 holo-workers, each with | | | | | | |
| 24 vCPUs, | | | | | | |
| 192 GB RAM | | | | | | |
| Data Storage Resources | | | | | | |
| (44,000 GB Storage Capacity, 3 years prepay) | | | | | | |
| | China (Shanghai) | 1 | 43,560,000 | 1 | 43,560,000 | Included |
| License Computer and Software Service Sub-Total | | | | | 44,352,000 | 0.00 |
| Other Components | | | | | | |
| Lenovo E41-55 laptop | | | | | | |
| | | 2 | 2199.00 | 3 | 6597.00 | |
| Other Components Sub-Total | | | | | 6597.00 | |
| 1 = Alibaba Cloud 2 = Tmall.com | | | | 3-year Cost of Ownership: | | 44,358,597.00 |
| All prices are based on 3-year pre-paid subscriptions. | | | | QphH@30000GB: | | 27,865,885.7 |
| | | | | ¥/kQphH@30000GB | | 1,591.87 |
| 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.2.8.0 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|---|---------------------------|---------------------|--|---------------------|---------------------|------|------------------|------------|----------------|----------------|----------------|------------|--------------|--------------|---------------------|----------|---------------------|---------------------|---------------------|---------------------|---------------------|-------------------|------|------------------|------------|----------------|----------------|----------------|------------|--------------|--------------|---------------------|----------|---------------------|---------------------|---------------------|---------------------|---------------------|---|------------|---------------------|----------|---------------------|---------------------|---------------------|---------------------|---------------------|---|------------|---------------------|----------|---------------------|---------------------|---------------------|---------------------|---------------------|---|------------|---------------------|----------|---------------------|---------------------|---------------------|---------------------|---------------------|---|------------|---------------------|----------|---------------------|---------------------|---------------------|---------------------|---------------------|---|------------|---------------------|----------|---------------------|---------------------|---------------------|---------------------|---------------------|---|------------|---------------------|----------|---------------------|---------------------|---------------------|---------------------|---------------------|---|------------|---------------------|----------|---------------------|---------------------|---------------------|---------------------|---------------------|---|------------|---------------------|----------|---------------------|---------------------|---------------------|---------------------|---------------------|----|------------|---------------------|----------|---------------------|---------------------|---------------------|---------------------|---------------------|
| | | | Report Date: 31-Oct-22 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| <div>Measurement Results</div> <div><div>Database Scale Factor</div><div>30,000</div></div> <div><div>Total Data Storage / Database Size</div><div>53.70</div></div> <div><div>Percentage Memory / Database Size</div><div>293.3%</div></div> <div><div>Start of Database Load Time</div><div>2022-10-20 18:56:58</div></div> <div><div>End of Database Load Time</div><div>2022-10-20 19:55:12</div></div> <div><div>Database Load Time</div><div>00:58:14</div></div> <div><div>Query Streams for Throughput Test (S)</div><div>10</div></div> <div><div>TPC-H Power</div><div>21,468,670.0</div></div> <div><div>TPC-H Throughput</div><div>36,169,338.3</div></div> <div><div>TPC-H Composite</div><div>27,865,885.7</div></div> <div><div>Total System Price Over 3 Years</div><div>44,358,597 (CNY)</div></div> <div><div>TPC-H Price/Performance Metric (¥/kQphH@30000GB)</div><div>1,591.87 (CNY)</div></div> <div>Measurement Interval</div> <div><div>Measurement Interval in Throughput Test (Ts)</div><div>656.91</div></div> | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| <div>Duration of stream execution:</div> <table><tr><th rowspan="3">Power Run</th><th rowspan="3">Seed</th><th>Query Start Time</th><th>Total Time</th><th>RF1 Start Time</th><th>RF2 Start Time</th></tr><tr><th>Query End Time</th><th>(hh:mm:ss)</th><th>RF1 End Time</th><th>RF2 End Time</th></tr><tr><td>2022-10-20 19:56:23</td><td rowspan="2">00:02:29</td><td>2022-10-20 19:56:05</td><td>2022-10-20 19:58:52</td></tr><tr><td>2022-10-20 19:58:52</td><td>2022-10-20 19:56:22</td><td>2022-10-20 19:59:02</td></tr></table> <table><tr><th rowspan="3">Throughput Stream</th><th rowspan="3">Seed</th><th>Query Start Time</th><th>Total Time</th><th>RF1 Start Time</th><th>RF2 Start Time</th></tr><tr><th>Query End Time</th><th>(hh:mm:ss)</th><th>RF1 End Time</th><th>RF2 End Time</th></tr><tr><td>2022-10-20 19:59:01</td><td rowspan="2">00:09:16</td><td>2022-10-20 19:59:01</td><td>2022-10-20 20:00:04</td></tr><tr><td>2022-10-20 20:08:17</td><td>2022-10-20 20:00:04</td><td>2022-10-20 20:00:55</td></tr><tr><td>2</td><td>1020195514</td><td>2022-10-20 19:59:01</td><td rowspan="2">00:08:58</td><td>2022-10-20 20:00:55</td><td>2022-10-20 20:02:03</td></tr><tr><td>2022-10-20 20:07:59</td><td>2022-10-20 20:02:03</td><td>2022-10-20 20:02:52</td></tr><tr><td>3</td><td>1020195515</td><td>2022-10-20 19:59:01</td><td rowspan="2">00:09:08</td><td>2022-10-20 20:02:52</td><td>2022-10-20 20:03:48</td></tr><tr><td>2022-10-20 20:08:09</td><td>2022-10-20 20:03:48</td><td>2022-10-20 20:04:51</td></tr><tr><td>4</td><td>1020195516</td><td>2022-10-20 19:59:01</td><td rowspan="2">00:09:23</td><td>2022-10-20 20:04:51</td><td>2022-10-20 20:05:44</td></tr><tr><td>2022-10-20 20:08:24</td><td>2022-10-20 20:05:44</td><td>2022-10-20 20:06:15</td></tr><tr><td>5</td><td>1020195517</td><td>2022-10-20 19:59:01</td><td rowspan="2">00:09:12</td><td>2022-10-20 20:06:15</td><td>2022-10-20 20:07:13</td></tr><tr><td>2022-10-20 20:08:13</td><td>2022-10-20 20:07:13</td><td>2022-10-20 20:07:52</td></tr><tr><td>6</td><td>1020195518</td><td>2022-10-20 19:59:01</td><td rowspan="2">00:09:03</td><td>2022-10-20 20:07:52</td><td>2022-10-20 20:08:21</td></tr><tr><td>2022-10-20 20:08:04</td><td>2022-10-20 20:08:21</td><td>2022-10-20 20:08:28</td></tr><tr><td>7</td><td>1020195519</td><td>2022-10-20 19:59:01</td><td rowspan="2">00:09:12</td><td>2022-10-20 20:08:28</td><td>2022-10-20 20:08:45</td></tr><tr><td>2022-10-20 20:08:13</td><td>2022-10-20 20:08:45</td><td>2022-10-20 20:08:51</td></tr><tr><td>8</td><td>1020195520</td><td>2022-10-20 19:59:01</td><td rowspan="2">00:08:44</td><td>2022-10-20 20:08:51</td><td>2022-10-20 20:09:09</td></tr><tr><td>2022-10-20 20:07:45</td><td>2022-10-20 20:09:09</td><td>2022-10-20 20:09:16</td></tr><tr><td>9</td><td>1020195521</td><td>2022-10-20 19:59:01</td><td rowspan="2">00:09:08</td><td>2022-10-20 20:09:16</td><td>2022-10-20 20:09:31</td></tr><tr><td>2022-10-20 20:08:09</td><td>2022-10-20 20:09:31</td><td>2022-10-20 20:09:37</td></tr><tr><td>10</td><td>1020195522</td><td>2022-10-20 19:59:01</td><td rowspan="2">00:09:13</td><td>2022-10-20 20:09:37</td><td>2022-10-20 20:09:53</td></tr><tr><td>2022-10-20 20:08:14</td><td>2022-10-20 20:09:53</td><td>2022-10-20 20:09:59</td></tr></table> | | | | | 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 | 2022-10-20 19:56:23 | 00:02:29 | 2022-10-20 19:56:05 | 2022-10-20 19:58:52 | 2022-10-20 19:58:52 | 2022-10-20 19:56:22 | 2022-10-20 19:59:02 | 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 | 2022-10-20 19:59:01 | 00:09:16 | 2022-10-20 19:59:01 | 2022-10-20 20:00:04 | 2022-10-20 20:08:17 | 2022-10-20 20:00:04 | 2022-10-20 20:00:55 | 2 | 1020195514 | 2022-10-20 19:59:01 | 00:08:58 | 2022-10-20 20:00:55 | 2022-10-20 20:02:03 | 2022-10-20 20:07:59 | 2022-10-20 20:02:03 | 2022-10-20 20:02:52 | 3 | 1020195515 | 2022-10-20 19:59:01 | 00:09:08 | 2022-10-20 20:02:52 | 2022-10-20 20:03:48 | 2022-10-20 20:08:09 | 2022-10-20 20:03:48 | 2022-10-20 20:04:51 | 4 | 1020195516 | 2022-10-20 19:59:01 | 00:09:23 | 2022-10-20 20:04:51 | 2022-10-20 20:05:44 | 2022-10-20 20:08:24 | 2022-10-20 20:05:44 | 2022-10-20 20:06:15 | 5 | 1020195517 | 2022-10-20 19:59:01 | 00:09:12 | 2022-10-20 20:06:15 | 2022-10-20 20:07:13 | 2022-10-20 20:08:13 | 2022-10-20 20:07:13 | 2022-10-20 20:07:52 | 6 | 1020195518 | 2022-10-20 19:59:01 | 00:09:03 | 2022-10-20 20:07:52 | 2022-10-20 20:08:21 | 2022-10-20 20:08:04 | 2022-10-20 20:08:21 | 2022-10-20 20:08:28 | 7 | 1020195519 | 2022-10-20 19:59:01 | 00:09:12 | 2022-10-20 20:08:28 | 2022-10-20 20:08:45 | 2022-10-20 20:08:13 | 2022-10-20 20:08:45 | 2022-10-20 20:08:51 | 8 | 1020195520 | 2022-10-20 19:59:01 | 00:08:44 | 2022-10-20 20:08:51 | 2022-10-20 20:09:09 | 2022-10-20 20:07:45 | 2022-10-20 20:09:09 | 2022-10-20 20:09:16 | 9 | 1020195521 | 2022-10-20 19:59:01 | 00:09:08 | 2022-10-20 20:09:16 | 2022-10-20 20:09:31 | 2022-10-20 20:08:09 | 2022-10-20 20:09:31 | 2022-10-20 20:09:37 | 10 | 1020195522 | 2022-10-20 19:59:01 | 00:09:13 | 2022-10-20 20:09:37 | 2022-10-20 20:09:53 | 2022-10-20 20:08:14 | 2022-10-20 20:09:53 | 2022-10-20 20:09:59 |
| 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 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 2022-10-20 19:56:23 | 00:02:29 | 2022-10-20 19:56:05 | 2022-10-20 19:58:52 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 2022-10-20 19:58:52 | 2022-10-20 19:56:22 | 2022-10-20 19:59:02 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 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 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | 2022-10-20 19:59:01 | 00:09:16 | 2022-10-20 19:59:01 | 2022-10-20 20:00:04 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 2022-10-20 20:08:17 | 2022-10-20 20:00:04 | 2022-10-20 20:00:55 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 2 | 1020195514 | 2022-10-20 19:59:01 | 00:08:58 | 2022-10-20 20:00:55 | 2022-10-20 20:02:03 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 2022-10-20 20:07:59 | 2022-10-20 20:02:03 | 2022-10-20 20:02:52 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 3 | 1020195515 | 2022-10-20 19:59:01 | 00:09:08 | 2022-10-20 20:02:52 | 2022-10-20 20:03:48 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 2022-10-20 20:08:09 | 2022-10-20 20:03:48 | 2022-10-20 20:04:51 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 4 | 1020195516 | 2022-10-20 19:59:01 | 00:09:23 | 2022-10-20 20:04:51 | 2022-10-20 20:05:44 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 2022-10-20 20:08:24 | 2022-10-20 20:05:44 | 2022-10-20 20:06:15 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 5 | 1020195517 | 2022-10-20 19:59:01 | 00:09:12 | 2022-10-20 20:06:15 | 2022-10-20 20:07:13 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 2022-10-20 20:08:13 | 2022-10-20 20:07:13 | 2022-10-20 20:07:52 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 6 | 1020195518 | 2022-10-20 19:59:01 | 00:09:03 | 2022-10-20 20:07:52 | 2022-10-20 20:08:21 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 2022-10-20 20:08:04 | 2022-10-20 20:08:21 | 2022-10-20 20:08:28 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 7 | 1020195519 | 2022-10-20 19:59:01 | 00:09:12 | 2022-10-20 20:08:28 | 2022-10-20 20:08:45 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 2022-10-20 20:08:13 | 2022-10-20 20:08:45 | 2022-10-20 20:08:51 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 8 | 1020195520 | 2022-10-20 19:59:01 | 00:08:44 | 2022-10-20 20:08:51 | 2022-10-20 20:09:09 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 2022-10-20 20:07:45 | 2022-10-20 20:09:09 | 2022-10-20 20:09:16 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 9 | 1020195521 | 2022-10-20 19:59:01 | 00:09:08 | 2022-10-20 20:09:16 | 2022-10-20 20:09:31 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 2022-10-20 20:08:09 | 2022-10-20 20:09:31 | 2022-10-20 20:09:37 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 10 | 1020195522 | 2022-10-20 19:59:01 | 00:09:13 | 2022-10-20 20:09:37 | 2022-10-20 20:09:53 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 2022-10-20 20:08:14 | 2022-10-20 20:09:53 | 2022-10-20 20:09:59 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

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 | 2.73 | 2.65 | 4.06 | 2.20 | 4.88 | 0.48 | 5.36 | 6.26 | 31.52 | 6.08 | 18.44 | 2.64 |
| 1 | 10.72 | 10.84 | 16.00 | 3.16 | 24.12 | 9.96 | 23.02 | 32.76 | 88.72 | 13.99 | 59.55 | 11.68 |
| 2 | 19.18 | 24.28 | 19.76 | 10.38 | 29.82 | 0.60 | 40.39 | 32.32 | 86.81 | 7.92 | 42.96 | 9.51 |
| 3 | 18.24 | 9.00 | 25.77 | 17.39 | 22.24 | 5.66 | 21.34 | 22.48 | 98.24 | 33.05 | 33.00 | 11.24 |
| 4 | 22.62 | 13.96 | 9.08 | 14.61 | 14.87 | 11.84 | 23.93 | 36.99 | 82.61 | 15.77 | 36.15 | 41.17 |
| 5 | 13.93 | 41.86 | 16.76 | 10.24 | 23.86 | 3.78 | 19.93 | 27.80 | 109.95 | 15.50 | 32.95 | 5.17 |
| 6 | 23.75 | 9.57 | 27.19 | 10.70 | 30.35 | 4.83 | 24.47 | 35.48 | 73.79 | 10.21 | 39.43 | 21.21 |
| 7 | 20.52 | 14.81 | 17.47 | 45.97 | 21.84 | 6.38 | 30.53 | 34.21 | 82.66 | 33.26 | 46.66 | 4.22 |
| 8 | 30.21 | 9.26 | 15.35 | 18.24 | 18.45 | 3.76 | 33.62 | 25.64 | 69.99 | 13.58 | 37.22 | 15.05 |
| 9 | 14.32 | 10.80 | 24.93 | 32.80 | 31.08 | 4.26 | 26.14 | 22.60 | 76.25 | 14.27 | 43.86 | 9.29 |
| 10 | 26.68 | 20.31 | 15.44 | 33.57 | 7.87 | 0.73 | 37.00 | 47.27 | 86.80 | 17.34 | 37.74 | 9.77 |
| Minimum | 2.73 | 2.65 | 4.06 | 2.20 | 4.88 | 0.48 | 5.36 | 6.26 | 31.52 | 6.08 | 18.44 | 2.64 |
| Average | 18.45 | 15.21 | 17.44 | 18.11 | 20.85 | 4.75 | 25.98 | 29.44 | 80.67 | 16.45 | 38.91 | 12.81 |
| Maximum | 30.21 | 41.86 | 27.19 | 45.97 | 31.08 | 11.84 | 40.39 | 47.27 | 109.95 | 33.26 | 59.55 | 41.17 |

| Stream ID | Q13 | Q14 | Q15 | Q16 | Q17 | Q18 | Q19 | Q20 | Q21 | Q22 | RF1 | RF2 |
|-----------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|
| 0 | 8.87 | 3.34 | 1.77 | 2.04 | 3.69 | 11.93 | 4.27 | 8.69 | 12.62 | 3.57 | 17.32 | 10.68 |
| 1 | 22.66 | 10.47 | 15.66 | 31.67 | 22.86 | 46.82 | 7.47 | 27.65 | 44.10 | 21.19 | 61.50 | 50.70 |
| 2 | 22.63 | 18.76 | 7.77 | 9.17 | 15.77 | 49.80 | 13.05 | 16.37 | 41.36 | 18.20 | 68.83 | 48.82 |
| 3 | 32.44 | 8.90 | 7.19 | 13.71 | 19.43 | 56.36 | 23.49 | 22.88 | 31.13 | 13.10 | 55.68 | 63.24 |
| 4 | 32.79 | 7.14 | 15.96 | 21.02 | 22.55 | 37.28 | 8.26 | 11.61 | 59.96 | 21.66 | 52.42 | 31.04 |
| 5 | 28.59 | 14.68 | 9.84 | 11.68 | 18.26 | 43.76 | 20.57 | 20.45 | 44.99 | 16.26 | 58.75 | 38.76 |
| 6 | 18.94 | 5.11 | 7.89 | 10.48 | 29.59 | 40.43 | 14.39 | 28.02 | 58.03 | 17.39 | 29.29 | 6.77 |
| 7 | 12.48 | 13.21 | 18.66 | 14.88 | 18.22 | 28.37 | 21.72 | 21.61 | 28.14 | 14.96 | 16.52 | 6.33 |
| 8 | 21.63 | 11.08 | 6.50 | 43.42 | 18.68 | 51.07 | 9.76 | 17.41 | 39.26 | 13.71 | 18.09 | 6.74 |
| 9 | 21.79 | 20.20 | 18.66 | 19.32 | 22.96 | 47.05 | 12.05 | 14.38 | 46.75 | 13.33 | 15.12 | 6.01 |
| 10 | 16.97 | 7.75 | 7.82 | 12.73 | 18.22 | 59.49 | 12.50 | 17.39 | 42.78 | 15.65 | 16.25 | 5.94 |
| Minimum | 8.87 | 3.34 | 1.77 | 2.04 | 3.69 | 11.93 | 4.27 | 8.69 | 12.62 | 3.57 | 15.12 | 5.94 |
| Average | 21.80 | 10.97 | 10.70 | 17.28 | 19.11 | 42.94 | 13.41 | 18.77 | 40.83 | 15.37 | 37.25 | 25.00 |
| Maximum | 32.79 | 20.20 | 18.66 | 43.42 | 29.59 | 59.49 | 23.49 | 28.02 | 59.96 | 21.66 | 68.83 | 63.24 |

| | |
|--|-----------|
| 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 |
| 3.1.1 Atomicity of Completed Transactions | 16 |
| 3.1.2 Atomicity of Aborted Transactions..... | 16 |
| 3.2 Consistency Requirements | 17 |
| 3.2.1 Consistency Test..... | 17 |
| 3.3 Isolation Requirements | 17 |
| 3.3.1 Isolation Test 1 – Read-Write Conflict with Commit | 17 |
| 3.3.2 Isolation Test 2 – Read-Write Conflict with Rollback | 17 |
| 3.3.3 Isolation Test 3 – Write-Write Conflict with Commit | 18 |
| 3.3.4 Isolation Test 4 – Write-Write Conflict with Rollback..... | 18 |
| 3.3.5 Isolation Test 5 – Concurrent Read and Write Transactions on Different Tables..... | 18 |
| 3.3.6 Isolation Test 6 – Update Transactions during Continuous Read-Only Query Stream..... | 18 |
| 3.4 Durability Requirements | 19 |
| 3.4.1 Permanent Unrecoverable Failure of Any Durable Medium | 19 |
| 3.4.2 System Crash..... | 20 |
| 3.4.3 Memory Failure..... | 20 |
| 3.4.4 Loss of External Power | 20 |
| 3.4.5 Node or Controller Failure | 20 |
| 4 Clause 4: Scaling and Database Population..... | 21 |
| 4.1 Initial Cardinality of Tables..... | 21 |
| 4.2 Distribution of Tables and Logs Across Media..... | 21 |
| 4.3 Mapping of Database Partitions/Replication | 21 |

| | | |
|------|--|----|
| 4.4 | Implementation of Data Redundancy | 22 |
| 4.5 | DBGen Modifications..... | 22 |
| 4.6 | Database Load Time..... | 22 |
| 4.7 | Data Storage Ratio | 22 |
| 4.8 | Database Load Mechanism Details and Illustration..... | 22 |
| 4.9 | Qualification Database Configuration..... | 23 |
| 4.10 | Memory to Database Size Percentage | 23 |
| 5 | Clause 5: Performance Metrics and Execution Rules Related Items..... | 24 |
| 5.1 | System Activity between Load and Performance Tests | 24 |
| 5.2 | Steps in the Power Test..... | 24 |
| 5.3 | Timing Interval for Each Query and Refresh Functions..... | 24 |
| 5.4 | Number of Streams for the Throughput Test..... | 24 |
| 5.5 | Start and End Date/Time of Each Query Stream..... | 24 |
| 5.6 | Total Elapsed Time of the Measurement Interval..... | 24 |
| 5.7 | Refresh Functions Start Date/Time and Finish Date/Time | 24 |
| 5.8 | Timing Metric Precision | 25 |
| 5.9 | Performance Metrics..... | 25 |
| 5.10 | The Performance Metric and Numerical Quantities from Both Runs..... | 25 |
| 5.11 | System Activity between Performance Tests..... | 25 |
| 5.12 | Documentation to satisfy Clause 5.2.7 | 25 |
| 5.13 | Query Output Validation | 25 |
| 6 | Clause 6: SUT and Driver Implementation Related Items | 26 |
| 6.1 | Driver | 26 |
| 6.2 | Implementation Specific Layer (ISL) | 26 |
| 6.3 | Profile-Directed Optimization..... | 26 |
| 7 | Clause 7: Pricing..... | 27 |
| 7.1 | Priced Configuration..... | 27 |
| 7.2 | Availability Date and Orderability Date | 27 |
| 7.3 | Country-Specific Pricing | 27 |
| 8 | Clause 8: Full Disclosure..... | 28 |
| 8.1 | Supporting Files Index Table | 28 |
| 9 | Clause 9: Audit Related Items..... | 28 |
| 9.1 | Auditor' s Report..... | 29 |
| | Appendix A: Price Quotes | 30 |

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 459 compute containers and an Pangu File System service with storage capacity of 44,000 GB. There is no difference between the priced and measured configurations.

| Configuration | HoloFrontend |
|-------------------------|----------------------------------|
| #HoloFrontend | 1 |
| Per HoloFrontend Config | 24 vCPUs |
| | 192 GB RAM |
| | 3,673 GB Local SSD Disk (Shared) |
| Network | 25 Gbits/s |

| Configuration | MetaStore |
|----------------------|------------------------------|
| #MetaStore | 1 |
| Per MetaStore Config | 18 vCPUs |
| | 144 GB RAM |
| | 3,673 GB Local Disk (Shared) |
| Network | 25 Gbits/s |

| Configuration | HoloMaster |
|----------------------|----------------------------------|
| #MetaStore | 1 |
| Per MetaStore Config | 14 vCPUs |
| | 112 GB RAM |
| | 3,673 GB Local SSD Disk (Shared) |
| Network | 25 Gbits/s |

| Configuration | HoloWorker |
|-----------------------|----------------------------------|
| #HoloWorker | 456 |
| Per HoloWorker Config | 24 vCPUs |
| | 192 GB RAM |
| | 3,673 GB Local SSD Disk (Shared) |
| Network | 25 Gbits/s |

| Configuration | PanguStorage |
|------------------|--------------|
| Storage Capacity | 44,000 GB |

Table 1. Configuration of the measured system.

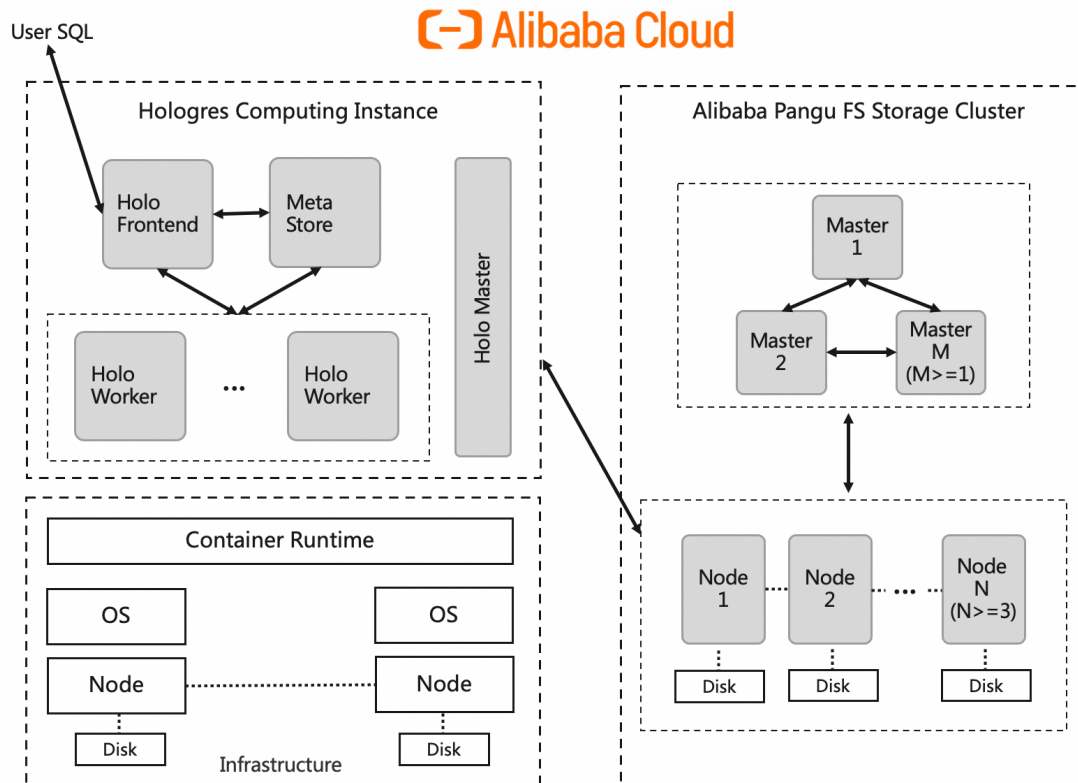


Fig. 1. System diagram of the measured configuration.

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 3456 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.

3.2.01 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.

3.2.02 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 ($O_ORDERKEY = L_ORDERKEY$).*

3.3.01 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.

3.4.01 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.

3.4.02 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.

3.4.03 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.

3.4.04 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.

3.4.05 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.

3.4.06 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  l1.l_orderkey
FROM    lineitem l1, lineitem l2, lineitem l3, lineitem l4, lineitem l5
WHERE   l1.l_shipdate <= date '1998-12-01'
        AND    l1.l_orderkey = l2.l_orderkey
        AND    l1.l_linenumber = l2.l_linenumber
        AND    l1.l_extendedprice = l3.l_extendedprice
        AND    l3.l_quantity < 30
        AND    l4.l_quantity = l1.l_quantity
        AND    l4.l_orderkey < 20000
        AND    l5.l_receiptdate = l1.l_receiptdate
        AND    l5.l_partkey < 140
GROUP BY l1.l_orderkey
ORDER BY l1.l_orderkey
LIMIT 2;
```

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. 400 transactions for each of the 11 executions streams were prepared.
3. After at least 100 ACID transactions were completed by each of the 11 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.

3.5.01 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.

Disk, node, and controller failure tests were performed together as shown in Section 3.5.5.

3.5.02 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.

3.5.03 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.

3.5.04 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.

3.5.05 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 1 holo-frontend, 1 meta-store, 1 holo-master and 456 holo-worker instances and 44,000 PanguFS storage. Four 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 Holo-Frontend container. The full cluster was restored automatically after the Holo-Frontend container was started automatically.
3. Shutdown the host machine of the Meta-Store container. The full cluster was restored automatically after the Meta-Store container was started automatically.
4. Shutdown the host machine of the Holo-Master container. The full cluster was restored automatically after the Meta-Store 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 | 179,999,978,268 |
| Order | 45,000,000,000 |
| Partsupp | 24,000,000,000 |
| Part | 6,000,000,000 |
| Customer | 4,500,000,000 |
| Supplier | 300,000,000 |
| 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 |
|--------------|----------------|----------------|------------------------|
| HoloFrontend | Local SSD DISK | /dev/nvme0n1p2 | Event logs, temp files |
| MetaStore | | | |
| HoloMaster | | | |
| HoloWorker | | | |

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 | 11,225 GB | 3,276 GB | 2,414 GB | 419 GB | 1,105 GB | 98 GB | 82 KB | 16 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 3456 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 | 40,978.19 |
| Hologres Compute Container | 3,420.75 GB x 459 = 1,570,123.26 |

Table 5. Data Storage Capacity

Total disk capacity: 1,611,101.45 GB Scale factor: 30,000.

The database storage ratio is 53.70.

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.

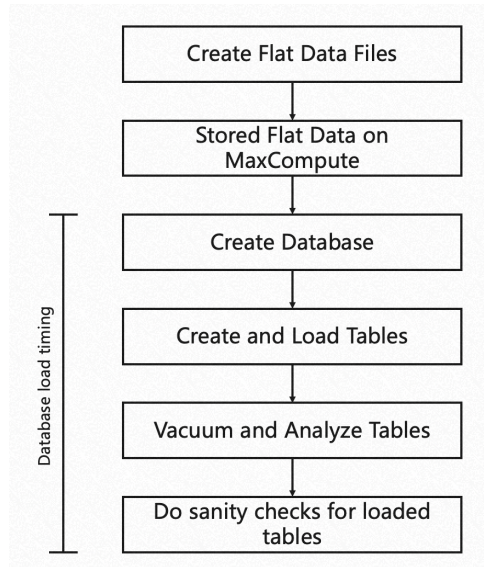


Fig. 2. Block Diagram of Database Loading Procedure

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: 88,000 GB Scale Factor: 30,000.

The memory to database size percentage is 293.3%.

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.

10 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@30,000GB | QppH@30,000GB | QthH@30,000GB |
|--------|---------------|---------------|---------------|
| Run 1 | 27,865,885.7 | 21,468,670.0 | 36,169,338.3 |
| Run 2 | 29,824,664.0 | 26,098,300.4 | 34,083,084.7 |

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. F*
- *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 10/31/2022.

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 10/31/2022.

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/ACID/ |
| | ACID Test Results | | ACID/ |
| 4 | Database Load Scripts | benchmark_scripts.zip | KIT/script/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.

Appendix A: Price Quotes

Hologres_combo

[Note] You can also pay for the storage resources of subscription exclusive instances in pay-as-you-go mode. Bills are generated on an hourly basis.

商品类型

Exclusive Instance (Subscription)Exclusive Instance (Pay-as-you-go)Shared Cluster (for Accelerated BI Analysis in MaxCompute)

Region

China

亚太

Europe and America

China (Hangzhou)

China (Shanghai)

China (Beijing)

China (Zhangjiakou)

China (Shenzhen)

China (Hong Kong)

Specifications

General-purpose (Recommended)Read-only Secondary InstanceMemory-optimized

Computing Resources (Memory-optimized)

11000 Cores, 88000 GB

To purchase an instance with more than 1,024 CUs, [submit a ticket](#).

Hot Data Storage Resources

44000

+

GB

Excess resources are charged in pay-as-you-go mode.

Instance Name

hologres_tpc_h

The content must be 2 to 64 characters in length.

Resource Group

Default Resource Group

If you want to create a resource group, click [here](#) to create.

Duration

1 Month2 Months3 Months4 Months5 Months6 Months3 Years

Computing Resources (Memory-optimized)

¥87,120,000.00

Saved ¥43,560,000.00

Hot Data Storage Resources

¥1,584,000.00

Saved ¥792,000.00

Total Configuration Cost

¥44,352,000.00

50% off for 3 years of paying Saved ¥44,352,000.00

Buy Now

Add to Cart

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

lenovo jingxu store

4.7

Seller Services

4.8

Logistics Services

4.9

Contact Customer Service

enter a shopping cart

Lenovo E41-55 14-inch commercial notebook AMD Athlon dual-core processor 3050U 4G 256GSSD WIN10 FHD Full HD

Monthly sales 0

¥2199

AssureparameterQualification

Shipping: Nanjing, Jiangsu to Yuhang District, Hangzhou

Express: Free Shipping

Color classification: black

Hard disk capacity: 256G SSD

Memory capacity: 4GB

Package Type: Official standard

Guarantee service: Comprehensive Warranty (Accident + Performance Failure Repair)

Two years (10% off for a limited time on Double 11) ¥89.1

Three years (20% off for a limited time on Double 11) ¥99.2

Extended Warranty (Performance Failure Extended Repair)

one year ¥58two years ¥91three years ¥103

quantity: 1 Inventory is tight

Buy nowadd to Shopping Cart

Favorites

Lenovo E41-55 14-inch commercial notebook AMD Athlon dual-core processor 3050U 4G 256GSSD WIN10 FHD Full HD

Monthly sales 0

¥2199

AssureparameterQualification

Shipping: Nanjing, Jiangsu to Yuhang District, Hangzhou

Express: Free Shipping

Color classification: black

Hard disk capacity: 256G SSD

Memory capacity: 4GB

Package Type: Official standard

Guarantee service: Comprehensive Warranty (Accident + Performance Failure Repair)

Two years (10% off for a limited time on Double 11) ¥89.1

Three years (20% off for a limited time on Double 11) ¥99.2

Extended Warranty (Performance Failure Extended Repair)

one year ¥58two years ¥91three years ¥103

quantity: 1 Inventory is tight

Buy nowadd to Shopping Cart

Favorites

Fig. 4. Lenovo E41-55 tablet purchase page (Google translated English version).

30