Full Disclosure Report

Dell Technologies

TPC Benchmark H Full Disclosure Report
For
Dell Technologies PowerEdge R7625 Server
While Using Microsoft SQL Server 2022 Enterprise Edition
and
Windows Server 2022 Standard Edition

First Edition: March 2024

Dell Technologies (Dell), the Sponsor of this benchmark test, believes that the information in this document is accurate as of the publication date. The information in this document is subject to change without notice. The Sponsor assumes no responsibility for any errors that may appear in this document. The pricing information in this document is believed to accurately reflect the current prices as of the publication date. However, we provide no warranty of the pricing information in this document.

Benchmark results are highly dependent upon workload, specific application requirements, system design and implementation. Relative system performance will vary as a result of these and other factors. Therefore, TPC Benchmark™ H should not be used as a substitute for a specific customer application benchmark when critical capacity planning and/or product evaluation decisions are contemplated.

All performance data contained in this report were obtained in a rigorously controlled environment. Results obtained in other operating environments may vary significantly. We do not warrant or represent that a user can or will achieve similar performance. No warranty of system performance or price/performance is expressed or implied in this report.

Dell and the Dell EMC Logo are trademarks of Dell Technologies and/or its affiliates in the U.S. and other countries. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Dell and any other company.

PowerEdge R7625 is a trademark of Dell Technologies.

Microsoft SQL Server 2022 is a registered trademark of Microsoft Corporation.

AMD and EPYC are trademarks of AMD Corporation.

TPC Benchmark, TPC-H, QppH, QthH, and QphH are trademarks or registered trademarks of the Transaction Processing Performance Council (TPC).

All other brand or product names mentioned are considered trademarks or registered trademarks of their respective owners.

Abstract

Overview

This report documents the methodology and results of the TPC Benchmark H test conducted on Dell Technologies PowerEdge R7625 server that was using Microsoft SQL Server 2022 Enterprise Edition in conformance with the requirements of the TPC Benchmark H Standard Specification, Revision 3.0.1. The operating system used for the benchmark was Microsoft Windows Server 2022.

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 website at www.tpc.org.

Standard and Executive Summary Statements

Pages iv-vii contain the Executive Summary and Numerical Quantities Summary of the benchmark results for the PowerEdge R7625 server.

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 QppH and QthH were audited by Doug Johnson of InfoSizing to verify compliance with the relevant TPC specifications.

The auditor's letter of attestation is attached in Section 9.1 "Auditors' Report."

DELL Technologies	**Dell PowerEdge R7625** **Using Microsoft SQL Server 2022**		TPC-H Rev. 3.0.1 TPC-Pricing Rev. 2.9.0 Report Date March 20, 2024	
Total System Cost	Composite Query per Hour Metric		Price / Performance	
$1,332,342.00	**2,720,098.1** QphH@10000GB		**$ 489.82** Price/ kQphH@10000GB	
Database Size	Database Manager	Operating System	Other Software	Availability Date
10,000GB	**SQL Server 2022**	**Windows Server 2022**	**N/A**	**March 20, 2024**

Database Load Time = 00d 004h 15m 38s	Storage Redundancy Levels:
Load includes backup = Y	Base Tables and Auxiliary Data Structures: Level 0
Total Data Storage / Database Size = 5.48	DBMS Temporary Space: Level 0
Memory to Database Size Percentage = 61.4%	OS and DBMS Software: Level 1 DBMS Log: Level 1

System Configuration:	Dell PowerEdge R7625 Server
Processors/Model:	2x AMD EPYC 9554 3.1GHz/3.75GHz 64c/128t, 245MB Cache
Cores/Thread	128C/256T
Memory	6144 (24x 256 GB DDR5-4800MT/s)
Storage	PERC 755 Controller w/ 4x 1.92TB SAS SSDs 8x Dell 6.4 TB NVMe, Mixed Use
Total Storage	54836.27

DELL Technologies	Dell PowerEdge R7625	TPC-H Revision 3.0.1 TPC Pricing Revision 2.9.0 Report Date: March-20-2024

Description	Part Number	Key	Unit Price	Qty	Extended Price	3 yr. Maint Price
HARDWARE COMPONENTS						
PowerEdge R7625 Server	210-BFMO	1	$907,633.84	1	$907,633.84	
2.5 Chassis	379-BDTF	1	$0.00	1	$0.00	
SAS/SATA/NVMe Capable Backplane	379-BDSW	1	$0.00	1	$0.00	
No Rear Storage	379-BDTE	1	$0.00	1	$0.00	
Trusted Platform Module 2.0 V3	461-AAIG	1	$0.00	1	$0.00	
2.5" Chassis with up to 24 SAS4/SATA Drives including 8 Universal Slots (NVMe Direct), Front PERC 11	321-BICT	1	$0.00	1	$0.00	
AMD EPYC 9554 3.10GHz, 64C/128T, 256M Cache (360W) DDR5-4800	338-CGXC	1	$0.00	2	$0.00	
Additional Processor Selected	379-BDCO	1	$0.00	1	$0.00	
High Performance Heatsink for 2 CPU Configuration	412-ABCZ	1	$0.00	1	$0.00	
Performance Optimized	370-AHLL	1	$0.00	1	$0.00	
4800MT/s RDIMMs	370-AHCL	1	$0.00	1	$0.00	
256GB RDIMM, 4800MT/s Octo Rank	370-AHCO	1	$0.00	24	$0.00	
RAID 1	780-BCDN	1	$0.00	1	$0.00	
PERC H755 SAS Front	405-AAZB	1	$0.00	1	$0.00	
Front PERC Mechanical Parts, for 2.5" x24 SAS/SATA Chassis	750-ADUB	1	$0.00	1	$0.00	
1.92TB SSD SAS ISD, Mixed Use up to 24Gbps 512e 2.5in Hot-Plug ,AG Drive	345-BELI	1 S	$0.00	4	$0.00	
6.4TB Enterprise NVMe Mixed Use AG Drive U.2 Gen4 with carrier	400-BKGI	1	$0.00	8	$0.00	
Performance BIOS Settings	384-BBBL	1	$0.00	1	$0.00	
UEFI BIOS Boot Mode with GPT Partition	800-BBDM	1	$0.00	1	$0.00	
No Energy Star	387-BBEY	1	$0.00	1	$0.00	
Fan Foam, HDD 2U	750-ACOM	1	$0.00	1	$0.00	
Very High Performance Fan x6	750-ADGJ	1	$0.00	1	$0.00	
Dual, Hot-Plug, Power Supply 2400W Redundant, D Mixed Mode	450-AJEV	1	$0.00	1	$0.00	
C19 to C20, PDU Style, 16 AMP, Power Cord	450-AEJJ	1	$0.00	2	$0.00	
Riser Config 2, 6 x8 FH + 2 x16 LP	330-BCBH	1	$0.00	1	$0.00	
PowerEdge R7625 Motherboard	329-BHNU	1	$0.00	1	$0.00	
No OCP 3.0 mezzanine NIC card	412-AASK	1	$0.00	1	$0.00	
Broadcom 5720 Dual Port 1GbE LOM	540-BDKD	1	$0.00	1	$0.00	
Broadcom 57414 Dual Port 10/25GbE SFP28 Adapter, PCIe Full Height, V2	540-BDHF	1	$0.00	1	$0.00	
No Cables Required	470-AEYU	1	$0.00	1	$0.00	
Dell Luggage Tag	321-BHXZ	1	$0.00	1	$0.00	
No Bezel	350-BBBW	1	$0.00	1	$0.00	
Assembly BOSS Blank	329-BERC	1	$0.00	1	$0.00	
iDRAC9, Enterprise 16G	528-CTIC	1	$0.00	1	$0.00	
Secured Component Verification	528-COYT	1	$0.00	1	$0.00	
No Quick Sync	350-BBYX	1	$0.00	1	$0.00	
iDRAC,Factory Generated Password	379-BCSF	1	$0.00	1	$0.00	
iDRAC Service Module (ISM), NOT Installed	379-BCQX	1	$0.00	1	$0.00	
iDRAC Group Manager, Disabled	379-BCQY	1	$0.00	1	$0.00	
No Rack Rails	770-BBBS	1	$0.00	1	$0.00	
No Systems Documentation, No OpenManage DVD Kit	631-AACK	1	$0.00	1	$0.00	
PowerEdge R7625 Shipping	340-DCWB	1	$0.00	1	$0.00	
PowerEdge R7625 Shipping Material	340-DCXY	1	$0.00	1	$0.00	
PowerEdge R7625 No CE or CCC Marking	389-EFMJ	1	$0.00	1	$0.00	
Basic Next Business Day 36 Months	709-BBFM	1	$249.00	1	$0.00	$249.00
ProSupport and 4Hr Mission Critical Initial, 36 Month(s)	865-BBNB	1	$110,037.32	1	$0.00	$110,037.32
Dell 24 Monitor – S2421HN	210-AXHJ	1	158.49	1	158.49	
Dell Wireless Keyboard and Mouse - KM3322W	580-AKCW	1	29.99	1	29.99	
HARDWARE COMPONENTS				**Subtotal**	$907,822.32	$110,286.32
SOFTWARE COMPONENTS						
Windows Server 2022 Standard,16CORE License		2	984.00	8	7,872.00	
Microsft Problem Resolution		2	259.00	1	259.00	
Microsoft SQL Server 2022 Enterprise Edition, 2 core licenses		2	15,123.00	64	967,872.00	
SOFTWARE COMPONENTS				**Subtotal**	$976,003.00	$0.00
Total					**$1,883,825.32**	110,286.32
Large Purchase Discount (65%)*					-590,084.51	-71,686.11

Pricing: 1 - Dell Mohan Rokkam; 2- Microsoft Corporation	**Three-Year Cost of Ownership:**	**$1,332,342.00**
* Discount based upon total system cost as purchased by a regular customer.	QphH@SF10,000GB:	**2,720,098.10**
S- One or more components of the Measured Configuration have been substituted in the Priced Configuration. See the FDR for details.	**$ /kQphH@SF10,000GB:**	**489.82**
Audited by Doug Johnson, InfoSizing Inc. (www.sizing.com)		

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

Numerical Quantities

Measurement Results

Database Scale Factor	10,000 GB
Total Data Storage / Database Size	5.48
Percentage Memory / Database Size	61.4%
Start of Database Load	2024-03-14 19:57:11
End of Database Load	2024-03-15 00:12:49
Database Load Time	00d 04h 15m 38s
Query Streams for Throughput Test	9
TPC-H Power	3,025,456.0
TPC-H Throughput	2,445,559.9
TPC-H Composite Query-per-Hour Metric (QphH@10000GB)	2,720,098.1
Total System Price Over 3 Years	USD 1,332,342.00
TPC-H Price/ Performance Metric ($/QphH@10000GB)	USD 489.82

Measurement Interval

Measurement Interval in Throughput Test (Ts)	2,914.67

Duration of Stream Execution

Power Run	Seed	Query Start Time Query End Time	Total Time (hh:mm:ss)	RF1 Start Time RF1 End Time	RF2 Start Time RF2 End Time
	315001249	2024-03-15 01:17:18 2024-03-15 01:23:03	00:05:45	2024-03-15 01:16:03 2024-03-15 01:17:16	2024-03-15 01:23:05 2024-03-15 01:25:15
Throughput Stream	**Seed**	**Query Start Time Query End Time**	**Total Time (hh:mm:ss)**	**RF1 Start Time RF1 End Time**	**RF2 Start Time RF2 End Time**
1	315001250	2024-03-15 01:25:14 2024-03-15 02:06:23	00:41:09	2024-03-15 01:25:14 2024-03-15 01:26:41	2024-03-15 01:26:41 2024-03-15 01:30:17
2	315001251	2024-03-15 01:25:14 2024-03-15 02:09:52	00:44:38	2024-03-15 01:30:17 2024-03-15 01:31:50	2024-03-15 01:31:50 2024-03-15 01:36:09
3	315001252	2024-03-15 01:25:14 2024-03-15 02:06:14	00:41:00	2024-03-15 01:36:09 2024-03-15 01:37:39	2024-03-15 01:37:39 2024-03-15 01:41:54
4	315001253	2024-03-15 01:25:14 2024-03-15 02:10:00	00:44:46	2024-03-15 01:41:55 2024-03-15 01:43:29	2024-03-15 01:43:29 2024-03-15 01:47:38
5	315001254	2024-03-15 01:25:14 2024-03-15 02:08:20	00:43:06	2024-03-15 01:47:38 2024-03-15 01:49:07	2024-03-15 01:49:07 2024-03-15 01:53:28
6	315001255	2024-03-15 01:25:14 2024-03-15 02:06:10	00:40:56	2024-03-15 01:53:28 2024-03-15 01:54:56	2024-03-15 01:54:56 2024-03-15 01:58:56
7	315001256	2024-03-15 01:25:14 2024-03-15 02:07:10	00:41:56	2024-03-15 01:58:56 2024-03-15 02:00:30	2024-03-15 02:00:30 2024-03-15 02:04:22
8	315001257	2024-03-15 01:25:14 2024-03-15 02:05:26	00:40:12	2024-03-15 02:04:23 2024-03-15 02:05:56	2024-03-15 02:05:56 2024-03-15 02:10:13
9	315001258	2024-03-15 01:25:14 2024-03-15 02:07:49	00:42:35	2024-03-15 02:10:13 2024-03-15 02:11:29	2024-03-15 02:11:29 2024-03-15 02:13:49

TPC-H Timing Intervals (in seconds)

Stream ID	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11	Q12
0	5.62	1.92	11.44	15.13	20.95	0.52	11.66	15.58	34.13	7.57	25.72	4.98
1	95.78	13.90	92.75	70.40	103.76	14.88	183.22	101.96	136.21	148.67	51.30	244.40
2	153.14	26.22	41.96	22.85	75.91	6.19	124.07	206.63	180.56	154.26	45.14	292.38
3	93.27	20.77	118.03	115.46	39.21	33.28	89.55	70.62	72.03	166.33	40.22	90.39
4	58.91	16.33	24.60	66.00	77.76	21.28	161.29	91.52	121.79	153.91	81.64	225.04
5	155.16	17.17	57.01	72.28	149.18	16.15	155.20	50.07	67.96	103.71	85.03	207.59
6	121.27	8.73	90.74	41.82	102.01	8.69	56.53	217.04	134.19	140.31	44.95	170.34
7	104.73	17.48	54.34	85.17	115.19	5.77	100.82	146.59	122.05	94.26	99.33	202.70
8	147.34	13.30	30.13	70.12	73.85	28.24	82.28	240.73	96.82	108.14	57.05	137.81
9	89.33	11.69	31.13	41.53	33.03	25.81	49.36	105.70	126.06	180.20	66.20	79.78
Qi Min	5.62	1.92	11.44	15.13	20.95	0.52	11.66	15.58	34.13	7.57	25.72	4.98
Qi Avg	102.46	14.75	55.21	60.08	79.09	16.08	101.40	124.64	109.18	125.74	59.66	165.54
Qi Max	155.16	26.22	118.03	115.46	149.18	33.28	183.22	240.73	180.56	180.20	99.33	292.38

Stream ID	Q13	Q14	Q15	Q16	Q17	Q18	Q19	Q20	Q21	Q22	RF1	RF2
0	42.85	2.51	4.72	11.39	5.59	53.14	9.09	14.25	38.24	6.62	72.69	129.61
1	67.32	7.05	13.64	16.10	117.15	477.01	43.12	43.17	347.45	78.76	86.21	216.00
2	94.29	10.16	25.30	27.06	162.02	745.21	116.79	27.11	74.13	65.50	92.91	258.50
3	59.87	4.85	8.12	33.22	165.36	654.27	38.93	51.40	352.15	142.20	90.55	255.03
4	68.52	18.75	24.42	33.09	142.58	647.21	169.77	67.61	391.52	21.23	94.32	249.11
5	113.49	5.70	20.63	33.56	118.27	515.90	62.16	48.23	491.20	39.41	88.20	260.94
6	54.45	5.20	8.50	24.21	148.89	556.59	161.34	36.98	282.30	39.57	88.36	239.49
7	68.14	11.55	30.77	44.53	64.95	514.94	166.00	71.66	374.39	19.16	93.71	232.60
8	57.36	5.57	17.19	37.80	131.97	663.55	21.19	21.86	313.92	55.00	93.34	256.92
9	58.07	8.22	12.36	37.99	157.73	538.33	209.87	54.09	459.71	178.04	75.30	140.76
Qi Min	42.85	2.51	4.72	11.39	5.59	53.14	9.09	14.25	38.24	6.62	72.69	129.61
Qi Avg	68.44	7.96	16.57	29.90	121.45	536.62	99.83	43.64	312.50	64.55	87.56	223.90
Qi Max	113.49	18.75	30.77	44.53	165.36	745.21	209.87	71.66	491.20	178.04	94.32	260.94

Contents

Preface

TPC Benchmark H Overview

The TPC Benchmark™ H (TPC-H) is a Decision Support benchmark. It is a suite of business-oriented ad hoc 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.

The TPC-H operations are modelled as follows:

- The database is continuously available 24 hours a day, 7 days a week for ad-hoc queries from multiple end users and data modifications against all tables, except possibly during infrequent (for example, once a month) maintenance sessions.
- The TPC-H database tracks, possibly with some delay, the state of the OLTP database through on-going refresh functions which batch together several modifications impacting some part of the decision support database.
- Due to the world-wide nature of the business data stored in the TPC-H database, the queries and the refresh functions may be executed against the database at any time, especially in relation to each other. In addition, this mix of queries and refresh functions is subject to specific ACIDity requirements, since queries and refresh functions may execute concurrently.
- To achieve the optimal compromise between performance and operational requirements, the database administrator can set, once and for all, the locking levels and the concurrent scheduling rules for queries and refresh functions.

The performance metric reported by the TPC-H is called the TPC-H Composite Query-Per-Hour Performance Metric (QphH@Size) and reflects multiple aspects of the capability of the system to process queries. These aspects include the selected database size against which the queries rerun, the query processing power when queries are submitted by a single stream and the query throughput when queries are submitted by multiple concurrent users. The TPC-H Price/Performance metric is expressed as $/kQphH@Size. To be compliant with the TPC-H standard, all references to TPC-H results for a given configuration must include all required reporting components. The TPC believes that comparisons of TPC-H results measured against different database sizes are misleading and discourages such comparisons.

The TPC-H database must be implemented using a commercially available database management system (DBMS) and the queries executed via and interface using dynamic SQL. The specification provides for variants of SQL, as implementers are not required to have implemented a specific SQL standard in full.

TPC-H uses terminology and metrics that are like other benchmarks, originated by the TPC and others. Such similarity in terminology does not in any way imply that TPC-H results are comparable to other benchmarks. The only benchmark results comparable to TPC-H are other TPC-H results compliant with the same revision.

Even though this benchmark offers a rich environment representative of many decision support systems, this benchmark does not reflect the entire range of decision support requirements. In addition, the extent to which a customer can achieve the results reported by a vendor is highly dependent on how closely TPC-H approximates the customer application. The relative performance of systems derived from this benchmark does not necessarily hold for other workloads or environments. Extrapolations to any other environment are not recommended.

Benchmark results are highly dependent upon workload, specific application requirements, and systems design and implementation. Relative system performance will vary as a result of these and other factors. Therefore, TPC-H should

not be used as a substitute for a specific customer application benchmarking when critical capacity planning and/or product evaluation decisions are contemplated.

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 TPC-H benchmark is sponsored by Dell Technologies. The benchmark implementation was developed and engineered by Dell Technologies. Testing took place at the Dell Technologies lab.

0.2 Parameter Settings

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

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

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

The supporting files archive contains a list of all database parameters and operating system parameters.

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 Under Test (SUT), a PowerEdge R7625 server depicted in Figure 0.1, that was used to obtain the results in this benchmark consists of the following components:

- CPU: 2x AMD EPYC 9554
- Memory: 24x 256 GB DDR5 4800MHz
- Storage:
  - Controller
    - Perc H755 RAID Controller
  - Disks
    - 4x Dell 960GB SAS SSD.
    - 8x Dell 6.5 TB NVMe.
- Network:
  - Embedded Broadcom 5720 Dual Port 1GbE LOM
  - Broadcom 57414 Dual Port 10/25 GbE SFP28 Adapter

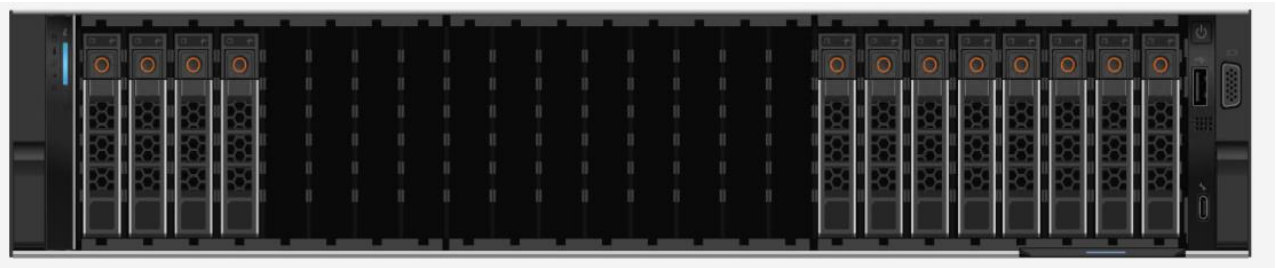

Figure 0.1: PowerEdge R7625 Server

The Priced configuration substituted Dell 1.92TB SAS SSD for each of the 960GB SAS SSD in the tested configuration.

- CPU: 2x AMD EPYC 9554
- Memory: 24x 256 GB DDR5 4800MHz
- Storage:
  - Controller
    - Perc H755 RAID Controller
  - Disks
    - 4x Dell 1.92 TB SAS SSD.
    - 8x Dell 6.4 TB NVMe.
- Network:
  - Embedded Broadcom 5720 Dual Port 1GbE LOM
  - Broadcom 57414 Dual Port 10/25 GbE SFP28 Adapter

1 Clause 1: Logical Database Design Related Items

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.

The Supporting Files Archive contains the build scripts that define the tables and indices for the TPC-H database.

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.

No column reordering was 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.

Horizontal partitioning is used on the LINEITEM and ORDERS tables and the partitioning columns are L_SHIPDATE and O_ORDERDATE. The partition granularity is by week.

1.4 Replication

Any replication of physical objects must be disclosed and must conform to the requirements of Clause 1.5.7.

No replication was used.

2 Clause 2: Queries and Refresh Functions

2.1 Query Language

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

T-SQL was the query language uniquely used throughout this benchmark.

2.2 Verifying Method for Random Number Generation

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

TPC supplied version 3.0.0 of DBGEN and QGEN were used in this benchmark.

2.3 Generating Values for Substitution Parameters

The method used to generate values for substitution parameters must be disclosed. If QGEN is not used for this purpose, then the source code of any non-commercial tool used must be disclosed. If QGEN is used, the version number, release number, modification number, and patch level of QGEN must be disclosed.

QGEN version 3.0.0 was used to generate the substitution parameters.

2.4 Query Text and Output Data from Qualification Database

The executable query text used for query validation must be disclosed along with the corresponding output data generated during the execution of the query text against the qualification database. If minor modifications (see Clause 2.2.3) have been applied to any functional query definition or approved variants in order to obtain executable query text, these modifications must be disclosed and justified. The justification for a particular minor query modification can apply collectively to all queries for which it has been used. The output data for the power and throughput tests must be made available electronically upon request.

The Supporting Files Archive contains the qualification query text and query output. The standard queries were used throughout with the following modifications:

- The “dateadd” function is used to perform date arithmetic in Q1, Q4, Q5, Q6, Q10, Q12, Q14, Q15, and Q20.
- The “datepart” function is used to extract part of a date (“YY”) in Q7, Q8, and Q9.
- The “top” function is used to restrict the number of output rows in Q2, Q3, Q10, Q18, and Q21.
- The “count_big” function is used in place of “count” in Q1.

2.5 Query Substitution Parameters and Seeds Used

The query substitution parameters used for all performance tests 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.

2.6 Isolation Level

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

The queries and transactions were run with the isolation level as Read-Committed.

2.7 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).

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

3 Clause 3: Database System Properties

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 specification. 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 the 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 steps were performed to verify the atomicity of the completed transactions:

1) The total price from the ORDERS table and the extended price from the LINEITEM table were retrieved for a randomly selected order key.
2) One ACID Transaction was performed using the order key from step 1.
3) The ACID Transaction was committed.
4) The total price from the ORDERS table and the extended price from the LINEITEM table were retrieved for the same order key. It was verified that the appropriate 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 steps were performed to verify the atomicity of the aborted ACID transactions:

1) The total price from the ORDERS table and the extended price from the LINEITEM table were retrieved for a randomly selected order key.
2) One ACID Transaction was performed using the order key from step 1. The transaction was stopped prior to the commit.
3) The ACID Transaction was ROLLED BACK.
4) The total price from the ORDERS table and the extended price from the LINEITEM table were retrieved for the same order key. It was verified that the appropriate 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 queries were executed before and after the durability tests to show that the database was always in a consistent state both initially and after submitting transactions:

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

1) The consistency of the ORDERS and LINEITEM tables was verified.
2) The 100 ACID transactions per stream were executed from 10 execution streams.
3) The consistency of the ORDERS and LINEITEM tables was re-verified.

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.

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 steps were performed to satisfy the test of isolation for a read-only and a read-write committed transaction:

1. An ACID Transaction was started for a randomly selected O_KEY, L_KEY and DELTA. The ACID Transaction was suspended prior to Commit.
2. An ACID query was started for the same O_KEY used in step 1. The ACID query blocked and did not see any uncommitted changes made by the ACID Transaction.
3. The ACID Transaction was resumed and committed.
4. The ACID query completed. It returned the data as committed by the ACID Transaction.

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 steps were performed to satisfy the test of isolation for a read-only and a rolled back read-write transaction:

1) An ACID transaction was started for a randomly selected O_KEY, L_KEY and DELTA. The ACID Transaction was suspended prior to Rollback.
2) An ACID query was started for the same O_KEY used in step 1. The ACID query did not see any uncommitted changes made by the ACID Transaction.
3) The ACID Transaction was ROLLED BACK.
4) The ACID query completed.

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 steps were performed to verify isolation of two update transactions:

1) An ACID Transaction T1 was started for a randomly selected O_KEY, L_KEY and DELTA. The ACID transaction T1 was suspended prior to Commit.
2) Another ACID Transaction T2 was started using the same O_KEY and L_KEY and a randomly selected DELTA.
3) T2 waited.
4) The ACID transaction T1 was allowed to Commit and T2 completed.
5) It was verified that: T2.L_EXTENDEDPRICE=T1.L_EXTENDEDPRICE
   +(DELTA1*(T1.L_EXTENDEDPRICE/T1.L_QUANTITY))

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

1) An ACID Transaction T1 was started for a randomly selected O_KEY, L_KEY and DELTA. The ACID Transaction T1 was suspended prior to Rollback.

2) Another ACID Transaction T2 was started using the same O_KEY and L_KEY used in step 1 and a randomly selected DELTA.
3) T2 waited.
4) T1 was allowed to ROLLBACK and T2 completed.
5) It was verified that T2.L_EXTENDEDPRICE = T1.L_EXTENDEDPRICE.

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 steps were performed to demonstrate the ability of read and write transactions affecting different tables to make progress concurrently:

1. An ACID Transaction T1 for a randomly selected O_KEY, L_KEY and DELTA. The ACID Transaction T1 was suspended prior to Commit.
2. Another ACID Transaction T2 was started using random values for PS_PARTKEY and PS_SUPPKEY.
3. T2 completed.
4. T1 completed and the appropriate rows in the ORDER, LINEITEM and HISTORY tables were changed.

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 steps were performed to verify isolation of update transaction during continuous read-only query:

1) A series of 6 overlapping transactions, T1.[1-6], each executed the above query against the qualification database, was started using a randomly selected DELTA such that the query runs longer. In this series T1.2 starts before T1.1 finishes. T1.3 starts before T1.2 finishes, etc.
   An ACID Transaction, T2, was started for a randomly selected O_KEY, L_KEY and DELTA.
2) T2 completed before the final query in the T1 series, and appropriate rows in the ORDERS, LINEITEM and HISTORY tables had been changed.
3) Transaction T1 completed executing the series of queries.
4) It was verified that the appropriate rows in the ORDER, LINEITEM and HISTORY tables were changed.

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.

3.5.1 Permanent Unrecoverable Failure of Any Durable Medium and Loss of System Power

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.

Determined that two tests were needed to complete in this section:
1) Removal of Log disk combined with removal of data disk.
2) System crash, memory failure and system power loss tests are combined.

Each of these tests were performed against the qualification database. The qualification database is identical to the test database in virtually every regard except size.

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.

System Crash Test:

3) The consistency of the ORDERS and LINEITEM tables was verified.
4) 200 transactions for each of the 10 executions streams were prepared.
5) After that at least 100 ACID transactions were submitted from each of the 10 execution streams.
6) System was powered off by pulling the plugs.

7) The 10 streams failed and recorded their number of committed transactions in success files.
8) Restored power to the system.
9) The database went through a recovery period.
10) Rolled forward, rolled-backward transactions captured in the DB ERRORLOG.
11) Recovery complete message captured in the DB ERRORLOG.
12) The consistency of the ORDERS and LINEITEM tables was re-verified.
13) The durability success files and the HISTORY table were compared.

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 loss of external power tests were combined as explained in section 3.5.2

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.

The system crash, memory failure and loss of external power tests were combined as explained in section 3.5.2

3.5.5 Log Disk and Data Disk Failure Test

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

Log and Data disk failure test was performed by the following steps:

1. The Complete database was backed up.
2. The Consistency of ORDERS and LINEITEM tables was verified.
3. 10 Streams of ACID transactions were started. Each stream executed a minimum of 100 transactions.
4. While the test was running, one of the RAID-1 configured log disks was removed.
5. It was determined that the test continued to run. SQL Server did not error.
6. All 10 streams continued to execute ACID transactions.
7. One of the data disks (6.4TB NVMe drives) was removed.
8. The 10 Streams of ACID transactions failed and recorded their number of committed transactions in SUCCSS files.
9. Stop the SQL server and start the SQL Server with -f option.
10. The Database log was backed up and the Database was dropped.
11. The pulled disk was replaced with a new one.
12. Stopped SQL Server and removed -f option.
13. Formatted the drive and created new volume with same letter and file structure.
14. Started the SQL Server, the database was restored.
15. When database restore completed, issued a command to apply the backed-up log file.
16. The counts in the history table and success files were compared and verified.
17. The Consistency of the ORDERS and LINEITEM tables was verified.

4 Clause 4: Scaling and Database Population

4.1 Ending 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.

The following table lists the TPC Benchmark H defined tables and the row count for each table as they existed upon completion of the build.

Table	Rows
Lineitem	59,999,994,267
Order	15,000,000,000
Partsupp	8,000,000,000
Part	2,000,000,000
Customer	1,500,000,000
Supplier	100,000,000
Nation	25
Region	5

4.2 Distribution of Tables and Logs Across Media

The distribution of tables and logs across all media must be explicitly described for the tested and priced systems.

Microsoft SQL Server was configured on PowerEdge R7625 Server with the following storage configuration:

- 8x 6.4 TB NVMe Drives.
- PERC H755 Controller
    - 2x 980GB SAS SSD disks Mirrored.
    - 2x 980GB SAS SSD disks Mirrored.

Physical Device	RAID	Drive Letter	Physical Size	Available Size	Content
PERC H755 controller 1st Virtual Disk	RAID-1	C:	1.92TB		OS/SQL Server 2022 binaries
PERC H755 controller 2nd Victual Disk	RAID-1	L:	1.92TB		DB and tempdb Logs
Dell 6.4 TB NVMe Drive	No Raid	D:	6.4 TB		Flatfiles, DB datafiles, tempdb files, database backups
Dell 6.4 TB NVMe Drive	No Raid	E:	6.4 TB		Flatfiles, DB datafiles, tempdb files, database backups
Dell 6.4 TB NVMe Drive	No Raid	F:	6.4 TB		Flatfiles, DB datafiles, tempdb files, database backups
Dell 6.4 TB NVMe Drive	No Raid	G:	6.4 TB		Flatfiles, DB datafiles, tempdb files, database backups
Dell 6.4 TB NVMe Drive	No Raid	H:	6.4 TB		Flatfiles, DB datafiles, tempdb files, database backups
Dell 6.4 TB NVMe Drive	No Raid	I:	6.4 TB		Flatfiles, DB datafiles, tempdb files, database backups
Dell 6.4 TB NVMe Drive	No Raid	J:	6.4 TB		Flatfiles, DB datafiles, tempdb files, database backups
Dell 6.4 TB NVMe Drive	No Raid	K:	6.4 TB		Flatfiles, DB datafiles, tempdb files, database backups

4.3 Mapping of Database Partitions/Replication

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

Horizontal partitioning is used on LINEITEM and ORDERS tables and the partitioning columns are L_SHIPDATE and O_ORDERDATE. The database partitions are evenly distributed across the eight NVMe drives.

4.4 Implementation of RAID

Implementations may use some form of RAID to ensure high availability. If used for data, auxiliary storage (e.g. indexes) or temporary space, the level of RAID must be disclosed for each device.

The database tables are hosted on eight 6.4TB Dell NVMe drives. The temporary files were hosted on the same drives as the database tables. The database log files resided on a RAID-1 array of two 9800GB SAS SSD's on the Perc H755 controller. Database backup was stored on the same NVMe drives as the data using mirrored backup sets in case of drive loss.

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 must be disclosed. In the event that a program other than DBGEN was used to populate the database, it must be disclosed in its entirety.

The supplied DBGEN version 3.0.0 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 4.1.3.1. The ratio must be reported to the nearest 1/100th, rounded up. For the reporting of configured disk capacity, gigabyte (GB) is defined to be 2^30 bytes.

Disk Type	GB per disk*	GiB per disk**	# of disks	Total (GiB)**
Internal	1920GB	1788.139 GiB	4	**7152.56**
Internal	6400GB	5960.464 GiB	8	**47683.72**
Total Storage				**54836.27**
	Data Storage Ratio			**5.48**

** Disk manufacturer definition of 1 GB is 10^9 bytes*
*** In this calculation 1 GiB is defined as 2^{30} bytes*

4.8 Database Load Mechanism Details and Illustration

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

The database was loaded using data generation stored on the flat files all on the tested and priced configuration. DBGEN was used to create the flat files.

The following block diagram describes the process used to load the database.

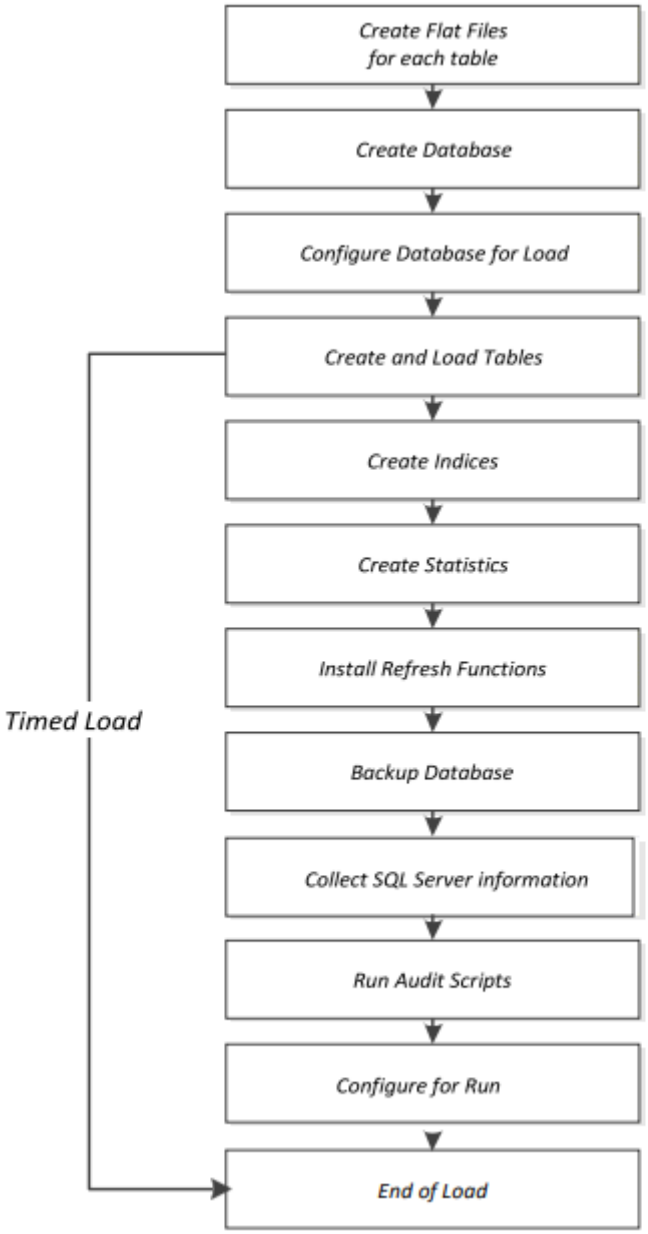

4.9 Qualification Database Configuration

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.

Total Memory	Scale Factor	Memory to Database Size Ratio
6144 GB	**10,000**	**61.4**%

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 disclosed.

The queries were generated using QGen at the end of the load test.
There were a few minutes between the load and the run.

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 disclosed.

The following steps were used to implement the power test:

1) RF1 refresh function.
2) Stream 0 execution.
3) RF2 refresh function.

5.3 Timing Interval for Each Query and Refresh Functions

The timing intervals (see Clause 5.3.7) for each query of the measured set for both refresh functions must be reported for the power test.

See Numerical Quantities Summary in the Executive Summary.

5.4 Number of Streams for the Throughput Test

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

Nine query streams and one refresh stream were used in the throughput test.

5.5 Start and End Date/Time of Each Query Stream

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

See Numerical Quantities Summary in the Executive Summary.

5.6 Total Elapsed Time of the Measurement Interval

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

See Numerical Quantities Summary in the Executive Summary.

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

Start and finish time for each update function in the update stream must be reported for the throughput test.

See Numerical Quantities Summary in the Executive Summary.

5.8 Performance Metrics

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

See Numerical Quantities Summary in the Executive Summary.

5.9 The Performance Metric and Numerical Quantities from Both Runs

A description of the method used to determine the reproducibility of the measurement results must be reported. This must include the performance metrics (QppH and QthH) from reproducibility runs.

Run ID	QppH@10000GB	QthH@10000GB	QphH@10000GB
Run 1	3,041,283.9	2,564,324.0	2,792,639.8
Run 2	3,025,456.0	2,445,559.9	2,720,098.1

5.10 System Activity between Performance Tests

Any activity on the SUT that takes place between the conclusion of Run 1 and the beginning of Run 2 must be disclosed.

There was no system activity between Run 1 and Run 2.

5.11 Documentation to satisfy Clause 5.2.7

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

The Supporting Files Archive contains the documentation.

5.12 Query Output Validation

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

The Supporting Files Archive contains the output of the validation test.

6 Clause 6: SUT and Driver Implementation Related Items

6.1 Driver

A detailed description of how the driver performs its functions must be supplied, including any related source code or scripts. This description should allow an independent reconstruction of the driver.

The TPC-H benchmark was implemented using a Microsoft tool called StepMaster. StepMaster is a general-purpose test tool which can drive ODBC and shell commands. Within StepMaster, the user designs a workspace corresponding to the sequence of operations (or steps) to be executed. When the workspace is executed, StepMaster records information about the run into a database as well as a log file for later analysis.

StepMaster provides a mechanism for creating parallel streams of execution. This is used in the throughput tests to drive the query and refresh streams. Each step is timed using a millisecond resolution timer. A timestamp T1 is taken before beginning the operation and a timestamp T2 is taken after completing the operation. These times are recorded in a database as well as a log file for later analysis.

Two types of ODBC connections are supported. A dynamic connection is used to execute a single operation and is closed when the operation finishes. A static connection is held open until the run completes and may be used to execute more than one step. A connection (either static or dynamic) can only have one outstanding operation at any time.

In TPC-H, static connections are used for the query streams in the power and throughput tests. StepMaster reads an Access database to determine the sequence of steps to execute. These commands are represented as the Implementation Specific Layer. StepMaster records its execution history, including all timings, in the Access database. Additionally, StepMaster writes a textual log file of execution for each run.

The stream refresh functions were executed using multiple batch scripts. The initial script is invoked by StepMaster, subsequent scripts are called from within the scripts.

The source for StepMaster and the RF Scripts is disclosed in the supported file archive.

6.2 Implementation Specific Layer (ISL)

If an implementation specific layer is used, then a detailed description of how it performs its functions must be supplied, including any related source code or scripts. This description should allow an independent reconstruction of the implementation-specific layer.

See section 6.1 for details.

6.3 Profile-Directed Optimization

If profile-directed optimization as described in Clause 5.2.9 is used, such use must be disclosed.

Profile-directed optimization was not used.

7 Clause 7: Pricing

7.1 Hardware and Software Used in the Priced System

A detailed list of hardware and software used in the priced system must be reported. Each item must have vendor part number, description, and release/revision level, and either general availability status or committed delivery date. If package-pricing is used, contents of the package must be disclosed. Pricing source(s) and effective date(s) of price(s) must also be reported.

A detailed list of all hardware and software, including the 3-year support, is provided in the Executive Summary in the Abstract section of this report. Third-party price quotations are included in Appendix A.

7.2 Three-Year Cost of System Configuration

The total 3-year price of the entire configuration must be reported including: hardware, software, and maintenance charges. Separate component pricing is recommended. The basis of all discounts used must be disclosed.

A detailed list of all hardware and software, including the 3-year support, is provided in the Executive Summary in the Abstract section of this report. The price quotations are included in Appendix A. This purchase qualifies for a 65% discount on all products from Dell Technologies.

7.3 Availability Date

The committed delivery date for general availability of products used in the priced 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.

The total system availability date is March 20, 2024.

7.4 Orderability Date

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 components are orderable at the time of publication date.

7.5 Country-Specific Pricing

Additional Clause 7 related items July be included in the Full Disclosure Report for each country-specific priced configuration. Country-specific pricing is subject to Clause 7.1.7.

The configuration is priced for the United States of America.

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 File Folder
Clause 1	OS and DB Parameter settings	SupportingFIlesArchive\Clause1
Clause 2	DB creation and Load scripts	SupportingFIlesArchive\Clause2
Clause 3	ACID Scripts, ACID output	SupportingFIlesArchive\Clause3
Clause 4	Qualification Query output	SupportingFIlesArchive\Clause4
Clause 5	Query output results	SupportingFIlesArchive\Clause5
Clause 6	Implementation Specific Layer source code	SupportingFIlesArchive\Clause6
Clause 7	No files required	SupportingFIlesArchive\Clause7
Clause 8	Query substitution parameters, RF Function source	SupportingFIlesArchive\Clause8

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.

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

InfoSizing

TPC Certified Auditor

Benchmark sponsor: Nicholas Wakou
Dell Technologies
701 E. Parmer Ln. Bld. 2
Austin, TX 78753

March 20, 2024

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

Platform:	Dell Technologies PowerEdge R7625 Server
Operating System:	Windows Server 2022 Standard Edition
Database Manager:	Microsoft SQL Server 2022 Enterprise Edition
Other Software:	n/a

The results were:

Performance Metric	**2,720,098.1 QphH@10000GB**
TPC-H Power	3,025,456.0
TPC-H Throughput	2,445,559.9
Database Load Time	0d 4h 15m 38s

Server	**Dell Technologies PowerEdge R7625 Server, with:**		
CPUs	2x AMD EPYC 9554 (3.1 GHz, 64-core, 245 MB Cache)		
Memory	6,144 GiB		
Disks	***Qty***	***Size***	***Type***
	4	1.92 TB	SAS SSD
	8	6.4 TB	NVMe, Mixed Use

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 10,000 GB 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

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

- 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 9 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:

The measured configuration included (4) 960 GB SAS SSDs drives that were substituted by (4) 1.92 TB SAS SSDs in the priced configuration. Based on the specifications of these drives, it is my opinion that this substitution has no significant impact on performance.

Respectfully Yours,

Doug Johnson, TPC Certified Auditor

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

Appendix A: Pricing Information

Microsoft Corporation
One Microsoft Way
Redmond, WA 98052-6399

Tel 425 882 8080
Fax 425 936 7329
http://www.microsoft.com/

Microsoft

March 14, 2024

Nicholas Wakou
Dell Technologies
701 E. Parmer Ln. Bld. 2
Austin, TX 78753

Here is the information you requested regarding pricing for several Microsoft products to be used in conjunction with your TPC-H benchmark testing.
All pricing shown is in US Dollars ($).

Description	Unit Price	Quantity	Price
Database Management System			
SQL Server 2022 Enterprise Edition *2 Core License* *Open Program – No Level - ERP*	$15,123.00	64	$967,872.00
Database Server Operating System			
Windows Server 2022 Standard *16 Core License* *Open Program – No Level – ERP*	$984.00	8	$7,872.00
Support			
Microsoft Problem Resolution Services *Professional Support* *(1 Incident).*	$259.00	1	$259.00

All software components are currently orderable and available. A list of Microsoft's resellers can be found in the Microsoft Product Information Center at **http://www.microsoft.com/products/info/render.aspx?view=22&type=how**
Defect support is included in the purchase price. Additional support is available from Microsoft PSS on an incident-by-incident basis at $259 call.
This quote is valid for the next 160 days.

Reference ID: TPCH_orclgvio491110458_2024