Dell Technologies

TPC Benchmark H Full Disclosure Report
For
Dell Technologies PowerEdge R7525 Server
While Using Microsoft SQL Server 2019 Enterprise Edition
and
Red Hat® Enterprise Linux® 8.3

First Edition: Dec 2021

Dell Technologies, 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, the Sponsor provides no warranty of the pricing information in this document.

Benchmark results are highly dependent upon workload, specific application requirements, and system design and implementation. Relative system performance will vary as a result of these and other factors. Therefore, the 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 was obtained in a rigorously controlled environment. Results obtained in other operating environments may vary significantly. No warranty of system performance or price/performance is expressed or implied in this report.

Printed in the United States, Dec 2021

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.

Dell and DELL EMC PowerEdge are registered trademarks of Dell Technologies.

Microsoft SQL Server 2019 Enterprise Edition are registered trademarks of Microsoft Corporation.

TPC Benchmark™, TPC-H, QppH, QthH, and QphH are registered certification marks of the Transaction Processing Performance Council.

All other brand or product names mentioned herein must be 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 R7525 server that was using Microsoft SQL Server 2019 Enterprise Edition in conformance with the requirements of the TPC Benchmark H Standard Specification, Revision 3.0.0. The operating system used for the benchmark was Red Hat® Enterprise Linux® 8.3

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**.

Standard and Executive Summary Statements

Pages iv−vii contain the Executive Summary and Numerical Quantities Summary of the benchmark results for the PowerEdge R7525 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	**PowerEdge R7525 Server**	TPC-H Rev. 3.0.0 TPC Pricing Rev. 2.7.0 Report Date: Dec 15, 2021
Total System Cost	Composite Query per Hour Rating	Price/Performance
$379,133 USD	**960,382.0** QphH@10000 GB	**$394.78** $/kQphH@10000GB

Database Size	Database Manager	Operating System	Other Software	Availability Date
10,000 GB	Microsoft® SQL Server 2019 Enterprise Edition	Red Hat® Enterprise Linux® 8.3	N/A	Dec 15, 2021

Database load Time = **00D 09H 57M 03S**	**Storage Redundancy Level**	
Load Includes Backup: **Y**	Base Tables and Auxiliary Data Structures	**No RAID**
Total Data Storage/Database Size = 6.65	DBMS Temporary Space	**No RAID**
Percentage Memory/ Database size = 41.0 %	OS and DBMS Software	**RAID 1**
	Log	**RAID 10**
System Configuration:	**Dell PowerEdge R7525 Server (2S)**	
Processors/Model:	**2 * AMD EPYC 73F3 3.5GHz, 16C/32T, 256M Cache (240W) DDR4-3200**	
Cores/Threads:	**32C/64T**	
Memory:	**4096 GiB**	
Storage:	**BOSS controller card + with 2, 480GB ,** **10 * 1.92TB SSD SAS Mix Use 12Gbps ,** **8 * 6.4TB Enterprise NVMe Mixed Use AG Drive U.2 Gen4**	
Total Storage:	**66,459 GiB**	

DELL Technologies

PowerEdge R7525 Server

TPC-H Rev. 3.0.0
TPC Pricing Rev. 2.7.0
Report Date: Dec 15, 2021

Description	Part Number	Source	Unit Price	Qty	Ext. Price $	3 Yr Maint. Price
PowerEdge R7525 Server	210-AUVQ	1	$200.00	1	$200.00	
SAS/SATA/NVMe Capable Backplane	379-BDSW	1	$0.00	1	$0.00	
Trusted Platform Module 2.0 V3	461-AAIG	1	$99.00	1	$99.00	
16X 2.5 SAS/SATA + 8X 2.5 NVME without XGMI	321-BFEE	1	$610.00	1	$610.00	
AMD EPYC 73F3 3.5GHz, 16C/32T, 256M Cache (240W) DDR4-3200	338-BZZJ	1	$6,099.00	2	$12,198.00	
Additional Processor Selected	379-BDCO	1	$0.00	1	$0.00	
Heatsink for 2 CPU + GPU/FPGA/Full Length card configs configuration	412-AASZ	1	$199.00	1	$199.00	
Performance Optimized	370-AAIP	1	$0.00	1	$0.00	
3200MT/s LRDIMMs	370-AFNV	1	$0.00	1	$0.00	
128GB LRDIMM, 3200MT/s, QuadRank	370-AGEW	1	$6,249.00	32	$199,968.00	
Unconfigured RAID	780-BCDS	1	$0.00	1	$0.00	
PERC H755 SAS Front	405-AAZB	1	$1,449.00	1	$1,449.00	
Front PERC Mechanical Parts,rear load	750-ACFQ	1	$0.00	1	$0.00	
1.92TB SSD SAS Mix Use 12Gbps 512e 2.5in Hot-plug AG Drive,3 DWPD	400-AZOY	1	$3,449.00	10	$34,490.00	
6.4TB Enterprise NVMe Mixed Use AG Drive U.2 Gen4 with carrier	400-BKGI	1	$10,529.00	8	$84,232.00	
Performance BIOS Settings	384-BBBL	1	$0.00	1	$0.00	
UEFI BIOS Boot Mode with GPTPartition	800-BBDM	1	$0.00	1	$0.00	
Very High Performance Fan x6	750-ACSM	1	$289.00	1	$289.00	
Dual, Hot-Plug, Power Supply 2400W Redundant, D Mixed Mode	450-AJEV	1	$2,099.00	1	$2,099.00	
Jumper Cord - C13/C14, 0.6M,250V, 13A (North American, Guam, North Marianas, Philippines, Samoa)	492-BBDH	1	$20.00	2	$40.00	
Riser Config 3, Half Length,5 x16 slots	330-BBPI	1	$469.00	1	$469.00	
R7525 Motherboard with LOM,MLK	384-BCUP	1	$0.00	1	$0.00	
iDRAC9,Enterprise 15G	385-BBOT	1	$489.00	1	$489.00	
Broadcom 57414 Dual Port 10/25GbE SFP28, OCP NIC 3.0	540-BCOC	1	$709.00	1	$709.00	
PowerEdge 2U Standard Bezel	325-BCHU	1	$49.00	1	$49.00	
BOSS controller card + with 2 M.2 Sticks 480GB (RAID 1),LP	403-BCHH	1	$1,499.00	1	$1,499.00	
Quick Sync 2 (At-the-box mgmt)	350-BBYW	1	$89.00	1	$89.00	
iDRAC,Legacy Password	379-BCSG	1	$0.00	1	$0.00	
iDRAC Group Manager, Enabled	379-BCQV	1	$0.00	1	$0.00	
No Operating System	611-BBBF	1	$0.00	1	$0.00	
No Media Required	605-BBFN	1	$0.00	1	$0.00	
ReadyRails Sliding Rails	770-BBBQ	1	$149.00	1	$149.00	
Assembly BOSS Blank	329-BERC	1	$0.00	1	$0.00	
Fan Foam, HDD 2U	750-ACOM	1	$0.00	1	$0.00	
No Systems Documentation, NoOpenManage DVD Kit	631-AACK	1	$0.00	1	$0.00	
PowerEdge R7525 Shipping	340-COFR	1	$0.00	1	$0.00	
R7525 Shipping Material	481-BBFE	1	$0.00	1	$0.00	
PowerEdge R7525 no CE or with CCC Marking	389-DUZG	1	$0	1	$0.00	
Prosupport and 4 Hrs Mission Critical (36 Months)	865-BBNB	1	$13,116	1	$0.00	$13,115.52
Basic Next Business Day 36 Months	709-BBFM	1	$200.01	1	$0.00	$200.01
				Sub Total	$339,326.00	$13,315.53
Large Purchase Discount (56%)				Discounted Price	-190,022.56	-7,456.70
				Server Subtotal	149,303.44	5,858.83
Logitech MK120 USB Keyboard and Mouse Combo (1+2 SPARE)	A6999510	1	$19.99	3	59.97	
Dell 22 Monitor - E2220H, 54.7cm (21.5")	210-AUNF	1	$235.49	1	$235.49	
Dell Limited Hardware Warranty	814-9340	1	$4.50	1	$4.50	
Advanced Exchange Service, 3 Years	814-9341	1	$0.00	1	$0.00	
				Peripharal Subtotal	$299.96	
				Discounted Price	-71.99	
				Peripharal Subtotal	227.97	
				Hardware Subtotal	149,531.41	
SOFTWARE						
SQL SERVER 2019 Enterprice Edition for LINUX (2 cores license, 32 Cores)	NA	2	$13,748.00	16	219,968.00	
Red Hat Enterprise Linux Server, Premium (Physical or Virtual Nodes), 3 Year x 1	RH00003F3	1	$3,515.20	1	3,515.20	
Microsoft Problem Resolution Services	NA	2	$259.00	1		259.00
				Software Subtotal	223,483.20	259.00
				Total	373,014.61	6,117.83

Source: 1: DELLEMC 2. MICROSOFT		
Audited by Doug Johnson of InfoSizing (www.sizing.com)	Three Year cost Ownership	$379,133.00
All discounts are based on US list prices and for similar quantities and configurations. The discounts was based on the overall specific components pricing from respective vendors in this single quotation. Discounts for similarly sized configurations will be similar to those quoted here, but may vary based on the components in the configuration.	QphH @10,000	960,382.00
	$/kQphH @10,000	$394.78

Dell EMC Sales Contact: Microsoft.Assist@dell.com

Windows and SQL Server Sales Contact: Microsoft Corporation, One Microsoft Way, Redmond, WA 98052 6399 (425) 882 8080

Prices used in TPC benchmarks reflect the actual prices a customer would pay for a one-time purchase of the stated components. Individually negotiated discounts are not permitted. Special prices based on assumptions about past or future purchases are not permitted. All discounts reflect standard pricing policies for the listed components. For complete details, see the pricing sections of the TPC benchmark specifications. If you find that the stated prices are not available according to these terms please inform the TPC at pricing@tpc.org. Thank you.

Numerical Quantities

Measurement Results

Database Scale Factor	10,000
Total Data Storage / Database Size	6.65
Percentage Memory / Database Size	41.0 %
Start of Database Load	12-10-2021 2:40:35 AM
End of Database Load	12-10-2021 12:37:38 PM
Database Load Time h m s	00d 09h 57m 03s
Query Streams for Throughput Test	9
TPC-H Power	1,213,443.2
TPC-H Throughput	760,096.2
TPC-H Composite Query-per-Hour (QphH@10,000 GB)	960,382.0
Total System Price over 3 Years ($ USD)	$379,133.0
TPC-H Price/Performance Metric ($ USD / kQphH@10,000 GB)	$394.78

Measurement Interval

Measurement Interval in Throughput Test (Ts) 9,377.76 Seconds

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
	1210123738	2021-12-10 16:26:47	00:16:32	2021-12-10 16:23:25	2021-12-10 16:43:21
		2021-12-10 16:43:19		2021-12-10 16:26:45	2021-12-10 16:48:55
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	1210123739	2021-12-10 16:48:54	02:28:36	2021-12-10 16:48:54	2021-12-10 16:52:30
		2021-12-10 19:17:30		2021-12-10 16:52:30	2021-12-10 17:00:47
2	1210123740	2021-12-10 16:48:54	02:35:24	2021-12-10 17:00:47	2021-12-10 17:04:10
		2021-12-10 19:24:18		2021-12-10 17:04:09	2021-12-10 17:12:36
3	1210123741	2021-12-10 16:48:54	02:27:14	2021-12-10 17:12:36	2021-12-10 17:16:05
		2021-12-10 19:16:08		2021-12-10 17:16:05	2021-12-10 17:24:44
4	1210123742	2021-12-10 16:48:54	02:24:31	2021-12-10 17:24:44	2021-12-10 17:28:13
		2021-12-10 19:13:25		2021-12-10 17:28:13	2021-12-10 17:35:44
5	1210123743	2021-12-10 16:48:54	02:35:21	2021-12-10 17:35:44	2021-12-10 17:39:24
		2021-12-10 19:24:15		2021-12-10 17:39:24	2021-12-10 17:48:20
6	1210123744	2021-12-10 16:48:54	02:36:01	2021-12-10 17:48:20	2021-12-10 17:52:08
		2021-12-10 19:24:55		2021-12-10 17:52:08	2021-12-10 17:59:59
7	1210123745	2021-12-10 16:48:54	02:33:34	2021-12-10 17:59:59	2021-12-10 18:03:25
		2021-12-10 19:22:28		2021-12-10 18:03:25	2021-12-10 18:11:09
8	1210123746	2021-12-10 16:48:54	02:25:32	2021-12-10 18:11:09	2021-12-10 18:14:32
		2021-12-10 19:14:26		2021-12-10 18:14:32	2021-12-10 18:22:40
9	1210123747	2021-12-10 16:48:54	02:36:18	2021-12-10 18:22:40	2021-12-10 18:26:11
		2021-12-10 19:25:12		2021-12-10 18:26:11	2021-12-10 18:34:11

TPC-H Timing Intervals (Seconds)

Query	Stream 0	Stream 1	Stream 2	Stream 3	Stream 4	Stream 5	Stream 6	Stream 7	Stream 8	Stream 9	Ql Min	Ql Avg	Ql Max
Q1	21.34	175.03	201.01	282.41	481.52	448.12	392.95	155.19	178.30	408.34	21.34	274.42	481.52
Q2	5.37	89.99	81.84	114.52	81.24	97.71	75.12	124.30	76.77	91.61	5.37	83.85	124.30
Q3	36.73	564.51	445.29	245.25	465.21	762.71	297.35	413.59	387.41	464.61	36.73	408.27	762.71
Q4	37.32	202.62	295.03	577.82	219.01	463.62	304.50	708.52	352.90	487.32	37.32	364.87	708.52
Q5	43.65	841.87	611.11	314.31	373.68	661.85	334.06	378.03	710.99	163.72	43.65	443.33	841.87
Q6	0.84	17.88	2.61	42.06	48.37	19.08	26.47	3.84	23.53	58.57	0.84	24.33	58.57
Q7	30.32	347.54	297.63	315.56	282.94	310.18	422.38	465.74	876.90	355.41	30.32	370.46	876.90
Q8	45.24	703.24	598.53	379.18	606.70	307.43	383.77	368.20	425.78	384.62	45.24	420.27	703.24
Q9	122.70	777.55	1,046.84	640.91	756.53	667.00	1,487.93	1,107.09	789.20	991.90	122.70	838.77	1,487.93
Q10	17.03	412.44	595.74	465.85	172.29	148.05	173.81	115.18	596.69	472.66	17.03	316.97	596.69
Q11	36.60	134.68	78.45	96.65	86.03	95.12	76.64	77.39	83.60	123.07	36.60	88.82	134.68
Q12	12.64	302.81	284.01	117.97	337.76	378.74	558.53	73.37	72.00	463.97	12.64	260.18	558.53
Q13	184.12	1,173.13	1,419.40	1,545.35	1,325.13	880.43	1,593.60	1,499.36	960.51	1,015.57	184.12	1,159.66	1,593.60
Q14	7.43	32.92	24.47	51.10	50.48	45.36	41.81	35.43	29.80	37.42	7.43	35.62	51.10
Q15	6.06	42.45	46.74	46.27	50.93	60.39	51.17	33.12	42.31	23.47	6.06	40.29	60.39
Q16	29.28	143.25	154.15	231.24	186.11	242.76	186.80	157.62	104.39	43.69	29.28	147.93	242.76
Q17	16.54	174.94	294.80	335.59	281.06	283.40	188.25	324.59	289.64	222.37	16.54	241.12	335.59
Q18	168.88	1,148.28	1,448.83	1,624.70	1,284.49	1,745.88	1,605.42	1,482.28	1,561.58	1,332.62	168.88	1,340.30	1,745.88
Q19	8.09	90.04	113.27	168.66	156.42	206.39	222.10	368.44	98.99	496.71	8.09	192.91	496.71
Q20	21.09	239.88	153.56	170.02	244.11	197.65	244.00	169.48	217.04	254.62	21.09	191.15	254.62
Q21	103.89	1,015.18	790.58	882.08	947.50	1,015.01	509.26	935.37	690.62	1,285.33	103.89	817.48	1,285.33
Q22	35.27	285.17	339.29	185.74	232.82	282.96	184.73	217.20	161.54	199.67	35.27	212.44	339.29
RF1	200.06	215.66	202.63	208.40	209.02	219.94	228.55	205.59	202.82	210.53	200.06	210.32	228.55
RF2	333.77	496.30	506.26	519.12	451.33	535.22	470.73	463.75	488.23	479.65	333.77	474.44	535.22

Contents

Preface
TPC Benchmark H Overview

The TPC Benchmark H (TPC-H) is a decision support benchmark. It consists of a suite of business-oriented ad-hoc queries and concurrent data modifications. The queries and the data populating the database have been chosen to have broad industry-wide relevance while maintaining an enough 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—for example, by using a point-and-click Graphical User Interface (GUI)
- 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 modeled 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 July 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 July 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 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 are run, 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 an 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**.

0.0 General Items

0.1 Test Sponsor

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

This benchmark was sponsored by Dell Technologies. The benchmark was developed and engineered by Dell Technologies. Testing took place at the Dell Technologies, Durham 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, if 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 Items

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 (for example, DBMS, query processing tools /languages, middle-ware components, and software drivers).

The System Under Test (SUT), a PowerEdge R7525 server depicted in Figure 0.1, consisted of:

- Dell Technologies PowerEdge R7525 server
  - 2 x AMD EPYC 73F3 3.5GHz, 16C/32T, 256M Cache (240W) DDR4-3200
  - 32 x 128GiB memory.

- Disk Drives
  - 10 x 1.92TB SSD SAS Mix Use 12Gbps 512e 2.5in Hot-plug AG Drive,3 DWPD
  - 8 x 6.4TB Enterprise NVMe Mixed Use AG Drive U.2 Gen4 with carrier

- Controllers
  - PERC H755 SAS Front RAID Controller, 8Gb NV Cache
  - BOSS controller card + with 2 M.2 Sticks 480GB (RAID 1), LP

Fig0.1: R7525 Server

Figure 0.3 Benchmark and priced configuration for PowerEdge R7525 server

Note—There were no differences between the tested and priced configurations.

- Dell Technologies PowerEdge R7525 server
  - 2 x AMD EPYC 73F3 3.5GHz, 16C/32T, 256M Cache (240W) DDR4-3200
  - 32 x 128GiB memory.

- Disk Drives
  - 10 x 1.92TB SSD SAS Mix Use 12Gbps 512e 2.5in Hot-plug AG Drive,3 DWPD
  - 8 x 6.4TB Enterprise NVMe Mixed Use AG Drive U.2 Gen4 with carrier

- Controllers
  - PERC H755 SAS Front RAID Controller, 8Gb NV Cache
  - BOSS controller card + with 2 M.2 Sticks 480GB (RAID 1), LP

1.0 Clause 1: Logical Database Design

1.1 Table Definitions

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 table definitions and the program used to load the database.

1.2 Physical Organization of Database

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 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 1.5.6.

No replication was used.

2.0 Clause 2: Queries and Refresh Functions - Related Items

2.1 Query Language

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

T-SQL was the query language used.

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.

The TPC source based DBGEN version 3.0.0 and QGEN was used to generate all database populations.

2.3 Substitution Parameters Generation

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.

The TPC source based QGEN version 3.0.0 was used to generate the substitution parameters.

2.4 Query Text and Output Data from 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 definitions or approved variants in order to obtain executable query text, these modifications must be disclosed and justified. The justification for a particular minor query modification can apply collectively to all queries for which it has been used. The output data for the power and throughput tests must be made available electronically upon request.

The Supporting Files Archive contains the query text and query output. The following modifications were used:

- 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

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 used.

2.6 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 as Read-Committed.

2.7 Refresh Functions

The details of how the refresh functions were implemented must be disclosed

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

3.0 Clause 3: Database System Properties

3.1 ACID Properties

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 results of the ACID tests must be disclosed along with a description of how the ACID requirements were met. This includes disclosing the code written to implement the ACID Transaction and Query.

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 ORDER, LINEITEM, and HISTORY tables.

The following steps were performed to verify the Atomicity of completed transactions:

1. The total price from the ORDER table and the extended price from the LINEITEM table were retrieved for a randomly selected order key.
2. The ACID Transaction was performed using the order key from step 1.
3. The ACID Transaction committed.
4. The total price from the ORDER 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, submitting a ROLLBACK of the transaction for the COMMIT of the transaction. Verify that the appropriate rows have not been changed in the ORDER, LINEITEM, and HISTORY tables.

The following steps were performed to verify the Atomicity of the aborted ACID transaction:

1. The total price from the ORDER table and the extended price from the LINEITEM table were retrieved for a randomly selected order key.
2. The 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 ORDER table and the extended price from the LINEITEM table were retrieved for the same order key used in steps 1 and 2. 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 – L_DISCOUNT) * (1 + L_TAX)))

for each ORDER and LINEITEM defined by (O_ORDERKEY = L_ORDERKEY)

3.3.1 Consistency Tests

Verify that ORDER and LINEITEM tables are initially consistent as defined in Clause 3.3.2.1, based upon a random sample of at least 10 distinct values of O_ORDERKEY.

The following steps were performed to verify consistency:
1. The consistency of the ORDER and LINEITEM tables was verified based on a sample of O_ORDERKEYs.
2. Each stream executed a minimum of 100 transactions.
3. The consistency of the ORDER and LINEITEM tables was re-verified.

The Consistency test was performed as part of the Durability test explained in section 3.5.

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. 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 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 the 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 verify isolation of concurrent read and write transactions on different tables:

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. An ACID Transaction T1 was started, executing a modified Q1 against the qualification database. The substitution parameter was chosen from the interval [0...2159] so that the query ran for a sufficient amount of time.
2. Before T1 completed, an ACID Transaction T2 was started using randomly selected values of O_KEY, L_KEY and DELTA.
3. Transaction T3 has started like T1, executing a modified Q1 against the qualification database.
4. T1 completed while T2 and T3 were running.
5. T2 completed while T3 was running.
6. 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 3.5.2.

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

Three tests were completed in this section:

1. Removal of a Log disk.
2. Removal of Data disk
3. Power Loss test.

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.

Log Removal Test

1. The complete database was backed up.
2. The Consistency of the ORDERS and LINEITEM tables were verified.
3. The ten streams of ACID transactions were started. Each stream executed a minimum of 100 transactions.
4. While the test was running, one of the RAID-10 configured log disk was removed.
5. It was determined that the test kept on running. SQL Server did not error.
6. The pulled disk was replaced with a new disk. Log disk eventually completed its RAID rebuild process without

any issues.
7. The consistency of the database was reconfirmed at the end of the test.

Data Disk Removal Test
1. The complete database was backed up.
2. The Consistency of the ORDERS and LINEITEM tables were verified.
3. The ten streams of ACID transactions were started. Each stream executed a minimum of 100 transactions.
4. While the test was running, one of the data disks (Non-RAID) was removed.
5. The ten streams of ACID transactions failed and recorded their number of committed transactions in success files.
6. Stop the SQL server and start the SQL server with -f option.
7. The Database log was backed up and the Database was dropped.
8. The pulled disk was replaced with a new one.
9. Stopped the SQL Server and removed -f option.
10. Formatted the drive and created XFS file system and added entries into fstab.
11. Started the SQL Server, the database was restored.
12. When database restore completed, issued a command to apply the backed-up log file.
13. The counts in the history table and success files were compared and verified, and the consistency of the ORDERS and LINEITEM tables was verified.

System Crash test
1. The Consistency of the ORDERS and LINEITEM tables were verified.
2. The ten streams of ACID transactions were started. Each stream executed a minimum of 100 transactions.
3. While the streams of ACID transactions were running the System was powered off by pulling power plugs.
4. When power was restored the system booted and the Database engine was restarted.
5. The database went through a recovery period.
6. Rolled-forward, Rolled-backward transactions captured by the DB ERRORLOG file.
7. Recovery complete.
8. The counts in the history table and success files were compared and verified, and the consistency of the ORDERS and LINEITEM tables was verified.

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.

See section 3.5.1.

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 section 3.5.1

4.0 Clause 4: Scaling and Database Population

4.1 Initial Cardinality of Tables

The cardinality (i.e., 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 4.1 lists the TPC-H Benchmark defined tables and the row count for each table as they existed upon completion of the build.

Table Name	Total Row Count
LINEITEM	59,999,994,267
ORDERS	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

Table 4.1 Initial Number of Rows

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 R7525 server with the following storage configuration:

NVMe and SSD Storage Configuration					
Drive type	**RAID Format**	**Mount**	**Physical Size**	**Available Size**	**Logical Allocation**
BOSS controller card + with 2 M.2 Sticks 480GB	RAID 1	/sdc	480 GB	447 GiB	OS,SQL Binaries
6 * 1.92TB SSD SAS Mix Use 12Gbps	RAID 10	/opt/db/log	1.92 TB	5.3 TiB	LOG
4 * 1.92TB SSD SAS Mix Use 12Gbps	RAID 0	/opt/db/backup	1.92 TB	7.0 TiB	Backup
Dell 6.4 TB, NVMe, Mixed Use Express Flash	No RAID	/opt/db/data1	6.4 TB	5.9 TiB	SQL DB DATA Files and TEMP DB Files
Dell 6.4 TB, NVMe, Mixed Use Express Flash	No RAID	/opt/db/data2	6.4 TB	5.9 TiB	SQL DB DATA Files and TEMP DB Files
Dell 6.4 TB, NVMe, Mixed Use Express Flash	No RAID	/opt/db/data3	6.4 TB	5.9 TiB	SQL DB DATA Files and TEMP DB Files
Dell 6.4 TB, NVMe, Mixed Use Express Flash	No RAID	/opt/db/data4	6.4 TB	5.9 TiB	SQL DB DATA Files and TEMP DB Files
Dell 6.4 TB, NVMe, Mixed Use Express Flash	No RAID	/opt/db/data5	6.4 TB	5.9 TiB	SQL DB DATA Files and TEMP DB Files
Dell 6.4 TB, NVMe, Mixed Use Express Flash	No RAID	/opt/db/data6	6.4 TB	5.9 TiB	SQL DB DATA Files and TEMP DB Files
Dell 6.4 TB, NVMe, Mixed Use Express Flash	No RAID	/opt/db/data7	6.4 TB	5.9 TiB	SQL DB DATA Files and TEMP DB Files
Dell 6.4 TB, NVMe, Mixed Use Express Flash	No RAID	/opt/db/data8	6.4 TB	5.9 TiB	SQL DB DATA Files and TEMP DB Files

4.3 Mapping of Database Partitions/Replications

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 8 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 used must be disclosed for each device.

The database tables were hosted on eight 6.4 TB Dell NVMe drives. The temporary files were hosted on the same drives as the database tables. The database log files resided on a RAID-10 array of six 1.92 TB SAS SSDs. OS on BOSS controller card + with2 M.2 Sticks 480G (RAID 1). The database backup was hosted on another RAID-0 array made of four 1.92 TB SAS SSD.

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

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.

The database load time was 09h 57m and 03s.

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.

Storage Devices	Space per Disk (GiB)	Total Disk Space (GiB)	Total Storage Capacity (GiB)	Scale factor	Data Storage Ratio
8 x Dell 6.4 TB, NVMe	5961.63	47,693.04	66,459.18	10,000	6.65
10 x 1.92 TB SSD SAS	1787.88	17,878.8			
2 x 480GB SSD (BOSS)	447.13	894.26			

Size of test database: 10,000 GiB Data Storage Ratio: 6.65

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.

Flat files for each of the tables were created using DBGEN. The tables were loaded as depicted in Figure 4.8.

All steps, scripts and configuration files are included in the Supporting Files.

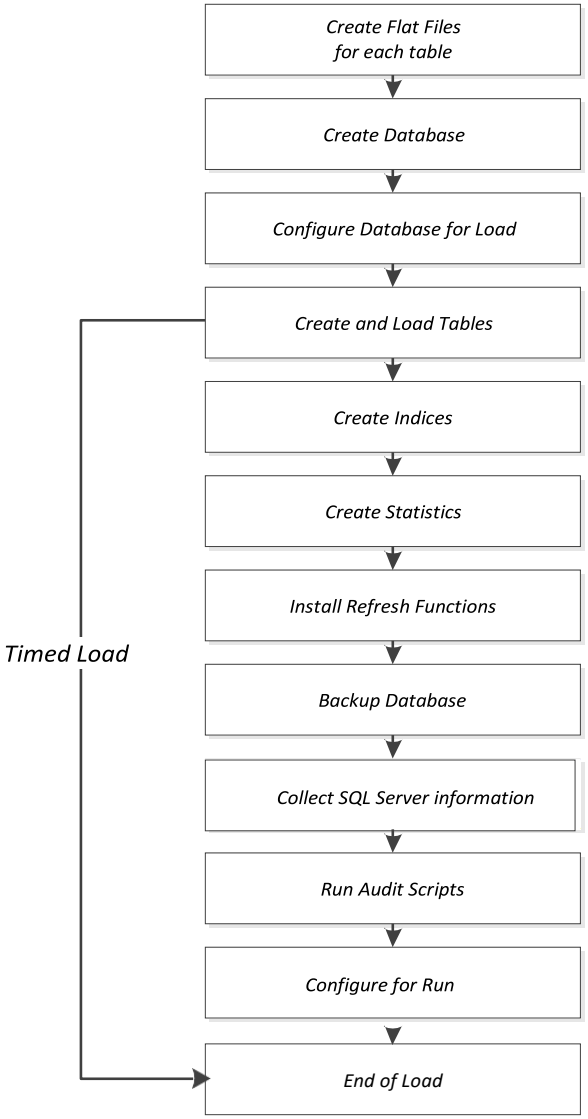

Figure 4.8: Block Diagram of Database Load Process

4.9 Qualification Database Configuration

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 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, as defined in clause 8.3.6.10, must be disclosed.

- Available Memory=4096 GiB
- Scale Factor=10,000
- The memory to database size percentage= 41.0%

5.0 Clause 5: Performance Metrics and Execution Rules Related Items

5.1 Steps after the Load Test

Any system activity on the SUT that takes place between the conclusion of the load test and the beginning of the performance test must be fully disclosed including listings of scripts or command logs.

The queries were generated using QGen at the end of the load test.
There were 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 (for example, system boot, and database restart) must be disclosed.

The following steps were used to implement the power test:

- RF1 Refresh Function
- Stream 00 Execution
- RF2 Refresh Function

5.3 Timing Intervals for Each Query and Refresh Function

The timing intervals (see Clause 5.3.6) for each query of the measured set and for both refresh functions must be reported for the power test.

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 execution streams used for the throughput test must be disclosed.

Nine query streams and one refresh stream were used for the Throughput test.

5.5 Start and End Date/Times for Each Query Stream

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

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

5.6 Total Elapsed Time for the Measurement Interval

The total elapsed time of the measurement interval (see Clause 5.3.5) must be reported for the throughput test.

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

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

5.8 Timing Intervals for Each Query and Each Refresh Function for Each Stream

The timing intervals (see Clause 5.3.6) for each query of each stream and for each update function must be reported for the throughput test.

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

5.9 Performance Metrics

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

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

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 the reproducibility runs.

Performance results from the first two executions of the TPC-H benchmark indicated the following difference for the metric points:

Run ID	QppH@10,000GB	QthH@10,000GB	QphH@10,000GB
Run 1	1,284,827.6	775,915.3	998,457.5
Run 2	1,213,443.2	760,096.2	960,382.0

5.11 System Activity Between Tests

Any activity on the SUT that takes place between the conclusion of Run1 and the beginning of Run2 must be disclosed.

There was no activity between Run1 and Run2.

5.12 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.

5.13 Query Validation Output

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

The supporting files archive contains the documentation.

6.0 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.

7.0 Clause 7: Pricing Related Items

7.1 Hardware and Software Used

A detailed list of hardware and software used in the priced system must be reported. Each item must have a vendor part number, description, and release/revision level, and indicate 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. The 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 required.

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 56% discount on Server and 24% discount on peripherals from Dell Technologies.

7.3 Availability Date

The committed delivery date for general availability (availability date) of products used in the priced calculations must be reported. When the priced system includes products with different availability dates, the single availability date reported on the first page of 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 Dec 15, 2021.

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.

8.0 Clause 8: Support Files Index Table

8.1 Supporting Files Index Table

An index for all files included in the supporting files archive as required by Clauses 8.3.2 must be provided in the report.

Clause	*Description*	*Archive File*
Clause 1	OS and DB parameter settings	*SupportingFilesArchive\Clause1*
Clause 2	DB creation scripts	*SupportingFilesArchive\Clause2*
Clause 3	ACID scripts, ACID output	*SupportingFilesArchive\Clause3*
Clause 4	DB Load scripts, Qualification output	*SupportingFilesArchive\Clause4*
Clause 5	Query output results	*SupportingFilesArchive\Clause5*
Clause 6	Implementation Specific layer source code	*SupportingFilesArchive\Clause6*
Clause 7	There are no files required to be included for Clause 7	*SupportingFilesArchive\Clause7*
Clause 8	Query substitution parameters, RF function source	*SupportingFilesArchive\Clause8*

9.0 Clause 9: Audit Related Items

9.1 Auditors' 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 of Doug Johnson 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.

InfoSizing

TPC Certified Auditor

Benchmark sponsor:

Venkateswara Reddy Vatam
Senior Manager, Software Engineering
Dell Technologies
No. 65/2
Bagmane Tech Park
Bairasandra Main Rd
CV Raman Nagar
Bengaluru, Karnataka,
India 560093

December 17, 2021

I verified the TPC Benchmark H (TPC-H™ v3.0.0) performance of the following configuration:

Platform:	Dell PowerEdge R7525 Server
Operating System:	Red Hat® Enterprise Linux® 8.3
Database Manager:	Microsoft® SQL Server 2019 Enterprise Edition
Other Software:	n/a

The results were:

Performance Metric	**960,382.0 QphH@10,000GB**
TPC-H Power	1,213,443.2
TPC-H Throughput	760,096.2
Database Load Time	00d 09h 57m 03s

Server	**Dell PowerEdge R7525 Server, with:**		
CPUs	2x AMD EPYC 73F3 3.5 GHz, 16C/32T, 256 MB Cache		
Memory	4,096 GiB		
Disks	***Qty***	***Size***	***Type***
	1	480 GB	BOSS controller card + with 2 M.2 Sticks (RAID 1)
	10	1.92 TB	SSD SAS Mix Use 12 Gbps
	8	6.4 TB	Enterprise 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 10,000GB and populated accordingly
- The compliance of the database auxiliary data structures was verified
- The database load time was correctly measured and reported
- The required ACID properties were verified and met
- The query input variables were generated by QGen
- The query text was produced using minor modifications and no query variant
- The execution of the queries against the SF1 database produced compliant answers
- A compliant implementation specific layer was used to drive the tests
- The throughput tests involved 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:

None.

Respectfully Yours,

Doug Johnson, TPC Certified Auditor

Appendix A: Price Quotes

Microsoft Corporation
One Microsoft Way
Redmond, WA 98052-6399

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

Microsoft

December 15, 2021

Venkateswara Reddy Vatam
Dell
No. 65/2
Bagmane Tech Park
Bairasandra Main Rd
Byrasandra
C V Raman Nagar
Bengaluru, Karnataka, India 560093

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 2019 Enterprise Edition *2 Core License* *Open Program – No Level - ERP*	*$13,748.00*	*16*	*$219,968.00*
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 120 days.

Reference ID: TPCH_swfvvejh4485273_2019