

TPC Benchmark™ H
Full Disclosure Report
for
Lenovo® ThinkSystem™ SR665
using
Microsoft® SQL Server® 2019
Enterprise Edition
and
Red Hat® Enterprise Linux® 8.3

TPC-H™ Version 3.0.0



Second Edition
Submitted for Review
April 2, 2021

Second Edition – April 2021

THE INFORMATION CONTAINED IN THIS DOCUMENT IS DISTRIBUTED ON AN AS IS BASIS WITHOUT ANY WARRANTY EITHER EXPRESSED OR IMPLIED. The use of this information or the implementation of any of these techniques is the customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item has been reviewed by Lenovo for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environment do so at their own risk.

In this document, any references made to a Lenovo licensed program are not intended to state or imply that only Lenovo's licensed program may be used; any functionally equivalent program may be used.

This publication was produced in the United States. Lenovo may not offer the products, services, or features discussed in this document in other countries, and the information is subject to change without notice. Consult your local Lenovo representative for information on products and services available in your area.

© Copyright Lenovo 2021. All rights reserved.

Permission is hereby granted to reproduce this document in whole or in part, provided the copyright notice as printed above is set forth in full text on the title page of each item reproduced.

Trademarks

Lenovo, ServeRAID, ThinkSystem, and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available on the Web at <http://www3.lenovo.com/us/en/legal/copytrade/>.

The following terms used in this publication are trademarks of other companies as follows: TPC Benchmark, TPC-H, QppH, QthH and QphH are trademarks of Transaction Processing Performance Council; AMD, EPYC, and AMD EPYC are trademarks of Advanced Micro Devices, Inc; Intel and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries; Microsoft, Windows Server, and SQL Server are registered trademarks of Microsoft Corporation in the United States and/or other countries; Red Hat and Enterprise Linux are registered trademarks of Red Hat, Inc. in the United States and other countries. Other company, product, or service names, which may be denoted by two asterisks (**), may be trademarks or service marks of others.

Notes

¹ GHz and MHz only measures microprocessor internal clock speed, not application performance. Many factors affect application performance.

² When referring to hard disk capacity, GB, or gigabyte, means one thousand million bytes. Total user-accessible capacity may be less.

Abstract

Lenovo conducted the TPC Benchmark H (TPC-H) on the Lenovo ThinkSystem SR665. This report documents the full disclosure information required by the TPC Benchmark H Standard Specification, Revision 3.0.0, including the methodology used to achieve the reported results. All testing fully complied with this revision level.

The software used on the Lenovo ThinkSystem SR665 system included Red Hat Enterprise Linux 8.3 and Microsoft SQL Server 2019 Enterprise Edition.

Standard metrics, Composite Query-per-Hour (QphH™@size), price per QphH (\$/kQphH@size) and Availability Date, are reported as required by the TPC Benchmark H Standard Specification.



The benchmark results are summarized in the following table:

Hardware	Software	Total System Cost (\$USD)	QphH @3,000	\$ USD / kQphH @3,000	Total Solution Availability Date
Lenovo ThinkSystem SR665	Microsoft SQL Server 2019 Enterprise Edition Red Hat Enterprise Linux 8.3	\$149,906.00	624,778.0	\$239.94	April 28, 2021

The benchmark implementation and results were audited by Doug Johnson for InfoSizing (www.sizing.com). The auditor's attestation letter is contained in this report.

Executive Summary

The Executive Summary is included on the next several pages.

		Lenovo[®] ThinkSystem[™] SR665		TPC-H [™] 3.0.0 TPC Pricing 2.7.0	
				Report Date: 4/2/21 Revision Date: 4/2/21	
Total System Cost		TPC-H Composite Query per Hour Metric		Price/Performance	
\$149,906.00 USD		624,778.0 QphH [™] @3,000GB		\$239.94 \$ / kQphH@3,000GB	
Database Size	DBMS Manager	Operating System	Other Software	System Availability Date	
3,000GB*	Microsoft[®] SQL Server[®] 2019 Enterprise Edition	Red Hat[®] Enterprise Linux[®] 8.3	N/A	April 28, 2021	
					
Database Load Time 00d 04h 02m 09s	Load Includes Backup Y	Memory/Database Size Percentage 34.1%	Total Disk/Database Size 8.74		
Storage Redundancy Level Zero	Base Tables No RAID	Auxiliary Structures No RAID	DBMS Temp Space No RAID	OS and DBMS RAID-1	
<u>System Configuration</u>					
Processors/Cores/Threads	2/16/32	AMD EPYC[™] 72F3 8-Core Processor, 3.7GHz, 256MB L3 Cache			
Memory	16	64GB TruDDR4 3200 MHz RDIMM			
Total Memory		1TB			
Storage Controllers	1	ThinkSystem RAID 930-8i 2GB PCIe 12Gb Adapter			
Storage Devices	2	2.5" 5210 960GB Entry SATA 6Gb Hot Swap SSD			
	2	U.2 CM5-V 1.6TB Mainstream NVMe PCIe 3.0 x4 Hot Swap SSD			
	6	U.2 PM1733 3.84TB Entry NVMe PCIe 4.0 x4 Hot Swap SSD			
Total Storage		26,226.04GB			
*Database Size includes only raw data (e.g., no temp, index, redundant storage space, etc.).					



Lenovo ThinkSystem SR665

TPC-H 3.0.0
TPC Pricing 2.7.0

Report Date: 4/2/21
Revision Date: 4/2/21

Description	Part Number	Price Source	Unit Price	Quantity	Extended Price	3-Yr. Maint. Price
Server Hardware						
Server : ThinkSystem SR665-3yr Warranty	7D2VCTO1WW	1	60,999	1	60,999	
ThinkSystem 2U 2.5" Chassis with 8, 16 or 24 Bays	B8LZ			1		
ThinkSystem AMD EPYC 72F3 8C 180W 3.7GHz Processor	BF7H			2		
ThinkSystem 64GB TruDDR4 Performance+ 3200 MHz (2Rx4 1.2V) RDIMM-A	BCZZ			16		
ThinkSystem RAID 930-8i 2GB Flash PCIe 12Gb Adapter	AUNJ			1		
ThinkSystem 1610-8P NVMe Switch Adapter	B4PA			1		
ThinkSystem 2.5" 5210 960GB Entry SATA 6Gb Hot Swap QLC SSD	B9AC			2		
ThinkSystem U.2 Toshiba CM5-V 1.6TB Mainstream NVMe PCIe 3.0 x4 Hot Swap SSD	B21X			2		
ThinkSystem U.2 PM1733 3.84TB Entry NVMe PCIe 4.0 x4 Hot Swap SSD	BC4Z			6		
ThinkSystem 2U/4U 8x2.5" NVMe Backplane	BCQM			2		
ThinkSystem 2U 8x2.5" SAS/SATA Backplane	B8LU			1		
ThinkSystem Marvell QL41232 10/25GbE SFP28 2-Port OCP Ethernet Adapter	B5SW			1		
ThinkSystem 2U PCIe Gen4 x16/x16 Slot 1&2 Riser 1 or 2	B8LQ			1		
ThinkSystem 1100W (230V/115V) V2 Platinum Hot-Swap Power Supply	B8QC			2		
2.8m, 13A/100-250V, C13 to C14 Jumper Cord	6400			2		
ThinkSystem XClarity Controller Standard to Enterprise Upgrade	AUPW			1		
ThinkSystem 2U Performance Fan Option Kit	B8LY			6		
ThinkSystem Toolless Slide Rail Kit v2	B8LA			1		
ThinkSystem 2U EIA Latch Standard (Left)	B8L7			1		
Lenovo 1m Passive 25G SFP28 DAC Cable	AV1W			1		
ThinkSystem SR665 2U Refresh MB	BF6Z			1		
ThinkSystem 2U EIA Latch with FIO (right)	B8M9			1		
ThinkSystem 2U MS Air Duct Filler(For 2U Gap)	B8MP			2		
FBU345 SuperCap	AUNP			1		
MS 2FH Riser Filler	BC4X			1		
ThinkSystem 2U MS 3FH Riser Filler	B8MM			1		
ThinkSystem 2U MS 3FH Riser1 Cage	B8MN			1		
Premier Essential - 3Yr 24x7 4Hr Resp + YDYP SR665	5PS7A78552	1	2,786	1		2,786
				Subtotal	60,999	2,786
Server Software						
SQL Server 2019 Enterprise Edition (2 Core License)	N/A	2	13,748	8	109,984	
RHEL Server Physical or Virtual Node, 2 Skt Prem RH Sup 3Yr	2020	1	5,290	1	5,290	
Microsoft Problem Resolution Services (1 Incident)	N/A	2	259	1		259
				Subtotal	115,274	259
Infrastructure						
S2 42U Standard Rack	93074RX	1	1,565	1	1,565	
ThinkSystem Pref. Pro II USB Keyboard - US English	AXTL	1	29	1	29	
ThinkSystem Optical Wheel Mouse - USB	B0LN	1	39	1	39	
ThinkVision T1714p 17 Inch Square LED Backlit LCD Monitor	60FELAR1US	1	159	1	159	
Essential Service - 3Yr 24x7 4Hr Response (Rack)	41L2760	1	315	1		315
				Subtotal	1,792	315
				Total	178,065	3,360
Dollar Volume Discount (See Note 1)	44.28%	1			31,519	
Pricing: 1 - Lenovo 1-877-782-7134; 2 - Microsoft				Three-Year Cost of Ownership USD:		\$149,906.00
Note 1: Discount applies to all line items where Source=1; pricing is for these or similar quantities.				QphH@3000GB:		624,778.0
Discounts for similarly sized configurations will be similar to what is quoted here, but may vary based on the specific components priced.				\$ USD/kQphH@3000GB:		\$239.94

Benchmark results and test methodology audited by Doug Johnson for 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 assumptions about past or future purchases are not permitted. All discounts reflect standard pricing policies for the listed components. For complete details, see the pricing section of the TPC benchmark specifications. If you find that stated prices are not available according to these terms, please inform the TPC at pricing@tpc.org. Thank you.



Lenovo ThinkSystem SR665

TPC-H 3.0.0
 TPC-Pricing 2.7.0
 Report Date: 4/2/21
 Revision Date: 4/2/21

Measurement Results:

Database Scale Factor	3,000
Total Data Storage/Database Size	8.74
Memory/Database Size	34.1%
Start of Database Load	2021-03-14 12:58:01
End of Database Load	2021-03-14 17:00:10
Database Load Time	00d 04h 02m 09s
Query Streams for Throughput Test	8
TPC-H Power (QppH™@3,000)	853,014.6
TPC-H Throughput (QthH™@3,000)	457,609.5
TPC-H Composite Query-per-Hour Metric (QphH@3,000GB)	624,778.0
Total System Price over 3 Years	\$149,906.00 USD
TPC-H Price/Performance Metric (\$/kQphH@3,000GB)	\$ 239.94 USD

Measurement Intervals:

Measurement Interval in Throughput Test (Ts) **4153.76**

Duration of Stream Execution:

Power Run	Seed	Query Start Time	Total Time	RF1 Start Time	RF2 Start Time
		Query End Time	(hh:mm:ss)	RF1 End Time	RF2 End Time
	314170010	2021-03-14 22:52:54	00:07:51	2021-03-14 22:52:35	2021-03-14 23:00:47
		2021-03-14 23:00:45		2021-03-14 22:52:52	2021-03-14 23:01:40
Throughput Stream	Seed	Query Start Time	Total Time	RF1 Start Time	RF2 Start Time
		Query End Time	(hh:mm:ss)	RF1 End Time	RF2 End Time
1	314170011	2021-03-14 23:01:40	00:58:21	2021-03-15 00:01:03	2021-03-15 00:01:20
		2021-03-15 00:00:01		2021-03-15 00:01:20	2021-03-15 00:02:14
2	314170012	2021-03-14 23:01:40	00:58:06	2021-03-15 00:02:14	2021-03-15 00:02:32
		2021-03-14 23:59:46		2021-03-15 00:02:32	2021-03-15 00:03:28
3	314170013	2021-03-14 23:01:40	00:56:05	2021-03-15 00:03:29	2021-03-15 00:03:46
		2021-03-14 23:57:45		2021-03-15 00:03:46	2021-03-15 00:04:43
4	314170014	2021-03-14 23:01:40	00:59:21	2021-03-15 00:04:43	2021-03-15 00:05:00
		2021-03-15 00:01:01		2021-03-15 00:05:00	2021-03-15 00:05:56
5	314170015	2021-03-14 23:01:40	00:58:38	2021-03-15 00:05:56	2021-03-15 00:06:15
		2021-03-15 00:00:18		2021-03-15 00:06:15	2021-03-15 00:07:10
6	314170016	2021-03-14 23:01:40	00:58:54	2021-03-15 00:07:11	2021-03-15 00:07:28
		2021-03-15 00:00:34		2021-03-15 00:07:28	2021-03-15 00:08:26
7	314170017	2021-03-14 23:01:40	00:57:40	2021-03-15 00:08:26	2021-03-15 00:08:43
		2021-03-14 23:59:20		2021-03-15 00:08:43	2021-03-15 00:09:41
8	314170018	2021-03-14 23:01:40	00:54:46	2021-03-15 00:09:41	2021-03-15 00:09:58
		2021-03-14 23:56:26		2021-03-15 00:09:58	2021-03-15 00:10:54

Timing Intervals (in seconds):

Stream ID	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11	Q12
0	10.97	3.18	14.41	15.03	18.61	0.51	11.97	15.08	67.52	9.10	21.93	5.83
1	60.72	16.91	123.09	61.73	125.89	4.11	158.41	338.81	347.00	111.01	82.39	91.71
2	66.32	25.08	248.57	95.95	127.67	1.28	92.92	182.40	345.44	79.29	48.52	114.00
3	102.13	15.87	83.02	78.45	181.13	8.27	133.37	220.99	303.00	115.73	46.45	42.49
4	107.94	20.10	14.36	86.47	134.99	3.46	82.29	221.40	617.73	75.21	55.57	189.10
5	67.85	20.69	101.84	101.56	127.82	4.91	83.89	357.94	580.85	65.41	92.47	157.53
6	70.56	19.77	85.22	103.69	94.99	6.36	101.16	349.78	325.28	74.26	129.06	210.26
7	83.28	22.38	88.82	94.22	125.45	2.30	156.69	208.05	499.77	150.69	48.61	44.72
8	75.26	22.09	166.75	132.41	172.49	12.58	74.09	340.8	266.51	65.85	45.2	144.85
QI Min	10.97	3.18	14.36	15.03	18.61	0.51	11.97	15.08	67.52	9.10	21.93	5.83
QI Avg	71.67	18.45	102.90	85.50	123.23	4.86	99.42	248.36	372.57	82.95	63.36	111.17
QI Max	107.94	25.08	248.57	132.41	181.13	12.58	158.41	357.94	617.73	150.69	129.06	210.26

Stream ID	Q13	Q14	Q15	Q16	Q17	Q18	Q19	Q20	Q21	Q22	RF1	RF2
0	86.31	3.63	3.12	15.89	6.77	78.36	3.57	10.26	43.74	23.61	17.01	53.40
1	485.65	23.44	81.28	98.37	65.63	591.40	89.41	99.72	264.66	178.88	16.40	53.67
2	563.01	26.29	17.76	171.83	57.86	435.26	24.89	114.11	525.89	120.81	17.85	55.95
3	430.24	27.36	25.53	152.11	51.20	681.62	32.27	104.32	420.97	107.75	17.66	56.49
4	434.43	23.42	17.89	169.12	57.11	778.76	27.94	10.19	306.48	126.33	16.63	55.72
5	533.13	22.88	24.46	90.05	51.24	409.22	152.69	113.03	277.79	79.81	19.05	54.84
6	609.66	31.72	24.05	125.23	40.61	681.55	45.92	93.65	136.61	173.38	16.93	58.23
7	486.15	20.63	17.86	120.67	53.18	535.24	55.81	97.60	381.20	166.08	16.85	57.37
8	403.32	22.74	66.56	112.25	47.1	421.72	40.4	105.4	437.46	108.7	16.58	56.59
QI Min	86.31	3.63	3.12	15.89	6.77	78.36	3.57	10.19	43.74	23.61	16.40	53.40
QI Avg	447.99	22.46	30.95	117.28	47.86	512.57	52.54	83.14	310.53	120.59	17.22	55.81
QI Max	609.66	31.72	81.28	171.83	65.63	778.76	152.69	114.11	525.89	178.88	19.05	58.23

Table of Contents

Abstract	3
Executive Summary	3
Table of Contents	8
Preface	10
General Items	11
Benchmark Sponsor	11
Parameter Settings	11
Configuration Diagrams.....	11
Measured Configuration.....	11
Priced Configuration	12
Differences between the Priced and Measured Configurations	13
Substitution	13
Clause 1 – Logical Database Design Related Items.....	14
Database Table Definitions	14
Database Physical Organization.....	14
Horizontal/Vertical Partitioning.....	14
Replication	14
Clause 2 – Query and Refresh Function Related Items.....	15
Query Language.....	15
QGen.....	15
Query Text and Output Data from Database.....	15
Query Substitution Parameters and Seeds Used	15
Query Isolation Level	15
Refresh Function Implementation.....	15
Clause 3 – Database System Properties Related Items.....	16
Atomicity Requirements	16
Atomicity of Completed Transactions.....	16
Atomicity of Aborted Transactions	16
Consistency Requirements	16
Consistency Tests	17
Isolation Requirements	17
Isolation Test 1 - Read-Write Conflict with Commit	17
Isolation Test 2 - Read-Write Conflict with Rollback	17
Isolation Test 3 - Write-Write Conflict with Commit	17
Isolation Test 4 - Write-Write Conflict with Rollback	18
Isolation Test 5 - Concurrent Read and Write Transactions on Different Tables.....	18
Isolation Test 6 - Update Transactions during Continuous Read-Only Query Stream.....	18
Durability Requirements.....	18
Permanent Irrecoverable Failure of Any Single Durable Medium (Database Tables)	19
Permanent Irrecoverable Failure of Any Single Durable Medium (Database Log)	19
System Crash.....	19
Memory Failure.....	19
SUT Power Failure.....	19

Clause 4 – Scaling and Database Population Related Items	20
Initial Cardinality of Tables	20
Distribution of Tables and Logs.....	20
Database Partition / Replication Mapping	21
Storage Redundancy	21
DBGen	21
Database Load Time	21
Data Storage Ratio	22
Database Load Details	22
Qualification Database Configuration.....	22
Memory to Database Size Percentage.....	23
Clause 5 – Performance Metrics and Execution Rules Related Items	24
System Activity between Load and Performance Tests	24
Power Test Steps.....	24
Timing Intervals for Each Query and Refresh Function	24
Number of Query Streams Used for the Throughput Test	24
Start and End Date/Times for Each Query Stream	24
Total Elapsed Time for the Measurement Interval.....	24
Refresh Function Start Date/Time and Finish Date/Time.....	24
Performance Metrics.....	24
Performance Metric and Numerical Quantities from Both Runs	24
System Activity between Tests	25
Clause 6 – SUT and Driver Implementation Related Items.....	26
Driver.....	26
Implementation Specific Layer.....	26
Profile-Directed Optimization.....	26
Clause 7 – Pricing Related Items.....	27
Hardware and Software Components	27
Three-Year Cost of System Configuration	27
Availability Date.....	27
Country-Specific Pricing	27
Pricing Calculations	27
Clause 8 – Full Disclosure Report Related Items.....	28
Supporting Files Index	28
Clause 9 – Audit Related Items.....	29
Auditor	29
Attestation Letter	29
Appendix A: Price Quotes.....	32

Preface

TPC Benchmark H Standard Specification was developed by the Transaction Processing Performance Council (TPC). It was initially released on February 26, 1999. Revision 3.0.0 is the most recent version. This is the full disclosure report for benchmark testing of the Lenovo ThinkSystem SR665 according to the TPC Benchmark H Standard Specification.

The TPC Benchmark 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 industrywide 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 set 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 (e.g., via a point-and-click GUI interface);
- 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 with 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 (e.g., once a month) maintenance sessions.
- The TPC-H database tracks, possibly with some delay, the state of the OLTP database through ongoing refresh functions, which batch together a number of modifications impacting some part of the decision support database.
- Due to the worldwide 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 minimum database required to run the benchmark holds business data from 10,000 suppliers. It contains almost 10 million rows representing a raw storage capacity of about 1 gigabyte. Compliant benchmark implementations may also use one of the larger permissible database populations (e.g., 100 gigabytes), as defined in Clause 4.1.3.

The performance metrics 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 executed, 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 (see Clause 5.4.6). 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.

Benchmarks 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 specific customer application benchmarking when critical capacity planning and/or product evaluation decisions are contemplated.

General Items

Benchmark Sponsor

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

This benchmark was sponsored by Lenovo.

Parameter Settings

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


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

See the Supporting Files, “Tunable Parameters,” which contains a list of all database parameters and operating system parameters.

Configuration Diagrams

Diagrams of both measured and priced configurations must be provided, accompanied by a description of the differences.

Measured Configuration



Lenovo ThinkSystem SR665

Table 0-1. Measured Configuration Details

Quantity	Lenovo ThinkSystem SR665 Configure-To-Order, includes:
1	ThinkSystem 2U 2.5" Chassis with 8, 16 or 24 Bays
2	ThinkSystem AMD EPYC 72F3 8-Core 180W 3.7GHz Processor 2 / 16 / 32 Total Processors/Cores/Threads 1 x 256MB L3 cache per processor 8 x 512KB L2 cache per processor
16	ThinkSystem 64GB TruDDR4 Performance+ 3200 MHz RDIMM 1TB Total System Memory
1	ThinkSystem RAID 930-8i 2GB Flash PCIe 12Gb Adapter
1	ThinkSystem 1610-8P NVMe Switch Adapter
2	ThinkSystem 2.5" 5210 960GB Entry SATA 6Gb Hot Swap QLC SSD
2	ThinkSystem U.2 Toshiba CM5-V 1.6TB Mainstream NVMe PCIe 3.0 x4 Hot Swap SSD
4	ThinkSystem U.2 PM1733 3.84TB Entry NVMe PCIe 4.0 x4 Hot Swap SSD
2	ThinkSystem 2U/4U 8x2.5" NVMe Backplane
1	ThinkSystem 2U 8x2.5" SAS/SATA Backplane
1	ThinkSystem Marvell QL41232 10/25GbE SFP28 2-Port OCP Ethernet Adapter
1	ThinkSystem 2U PCIe Gen4 x16/x16 Slot 1&2 Riser 1 or 2
2	ThinkSystem 1100W (230V/115V) V2 Platinum Hot-Swap Power Supply
6	ThinkSystem 2U Performance Fan Option Kit
1	ThinkSystem SR665 2U Refresh MB
1	Premier Essential Support - 3Yr 24x7 4Hr Resp + YDYD SR665
	SQL Server 2019 Enterprise Edition
	Red Hat Enterprise Linux 8.3

Priced Configuration

Table 0-2. Priced Configuration Details

Quantity	Lenovo ThinkSystem SR665 Configure-To-Order, includes:
1	ThinkSystem 2U 2.5" Chassis with 8, 16 or 24 Bays
2	ThinkSystem AMD EPYC 72F3 8-Core 180W 3.7GHz Processor 2 / 16 / 32 Total Processors/Cores/Threads 1 x 256MB L3 cache per processor 8 x 512KB L2 cache per processor
16	ThinkSystem 64GB TruDDR4 Performance+ 3200 MHz RDIMM 1TB Total System Memory
1	ThinkSystem RAID 930-8i 2GB Flash PCIe 12Gb Adapter
1	ThinkSystem 1610-8P NVMe Switch Adapter
2	ThinkSystem 2.5" 5210 960GB Entry SATA 6Gb Hot Swap QLC SSD
2	ThinkSystem U.2 Toshiba CM5-V 1.6TB Mainstream NVMe PCIe 3.0 x4 Hot Swap SSD

6	ThinkSystem U.2 PM1733 3.84TB Entry NVMe PCIe 4.0 x4 Hot Swap SSD
2	ThinkSystem 2U/4U 8x2.5" NVMe Backplane
1	ThinkSystem 2U 8x2.5" SAS/SATA Backplane
1	ThinkSystem Marvell QL41232 10/25GbE SFP28 2-Port OCP Ethernet Adapter
1	ThinkSystem 2U PCIe Gen4 x16/x16 Slot 1&2 Riser 1 or 2
2	ThinkSystem 1100W (230V/115V) V2 Platinum Hot-Swap Power Supply
6	ThinkSystem 2U Performance Fan Option Kit
1	ThinkSystem SR665 2U Refresh MB
1	Premier Essential Support - 3Yr 24x7 4Hr Resp + YDYD SR665
	SQL Server 2019 Enterprise Edition
	Red Hat Enterprise Linux 8.3

Differences between the Priced and Measured Configurations

The priced configuration was the same as the measured configuration with the exception that the priced configuration included two additional ThinkSystem U.2 PM1733 3.84TB Entry NVMe PCIe 4.0 x4 Hot Swap SSDs, for a total of six. These additional drives were priced as space to hold the database backup.

Substitution

Some hardware components of the Priced Configuration may be substituted after the test sponsor has demonstrated to the auditor's satisfaction that the substituting components do not negatively impact the reported performance metric or numerical quantities. All substitutions must be reported in the FDR and noted in the auditor's attestation letter. Any information and/or measurement results used to prove the validity of a Component substitution must be included in the section of the FDR that describes the differences between the measured and Priced Configuration. Original and substituted Components must be clearly identified.

No components were substituted.

Clause 1 – Logical Database Design Related Items

Database Table Definitions

Listings must be provided for all table definition statements and all other statements used to set-up the test and qualification databases.

See the Supporting Files for the scripts that were used to set up the TPC-H test and qualification databases.

Database 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.

See the Supporting Files for the scripts that were used to create the indexes on the test and qualification databases.

No column reordering was used.

Horizontal/Vertical Partitioning

Horizontal partitioning of tables and rows in the test and qualification databases (see Clause 1.5.4) must be disclosed.

Horizontal partitioning on L_SHIPDATE and O_ORDERDATE was used and granularity was week.

Replication

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

Replication was not used.

Clause 2 – Query and Refresh Function Related Items

Query Language

The query language used to implement the queries must be identified (e.g., “RALF/SQL-Plus”).

SQL was the query language used.

QGen

The version number, release number, modification number, and patch level of QGen must be disclosed.

QGen version 2.18.0 was used to generate all database populations.

Query Text and Output Data from Database

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

See the Supporting Files for the query text and query output. The following modifications were used:

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

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.

See the Supporting Files for the seed and query substitution parameters used.

Query 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 repeatable read isolation level.

Refresh Function Implementation

The details of how the refresh functions were implemented must be reported in the supporting files archive (including source code of any non-commercial program used).

See the Supporting Files for the source code for the refresh function.

Clause 3 – Database System Properties Related Items

The results of the ACID tests must be disclosed along with a description of how the ACID requirements were met.

All ACID tests were conducted according to specifications. The Atomicity, Isolation, Consistency and Durability tests were performed on the Lenovo ThinkSystem SR665 server. See the Supporting Files for the ACID transaction source code.

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.

Atomicity of Completed Transactions

Perform the ACID Transaction (see Clause 3.1.5) 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 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.

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 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.

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 TPC-H database, when populated as defined in Clause 4.2, must meet the consistency condition defined in Clause 3.3.2.1:

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 ORDERS and LINEITEM defined by $(O_ORDERKEY = L_ORDERKEY)$*

If data is replicated, as permitted under Clause 1.5.7, each copy must meet the consistency condition defined in Clause 3.3.2.1.

Consistency Tests

To verify the consistency between the *ORDERS* and *LINEITEM* tables, perform the following steps:

1. Verify that the *ORDERS* and *LINEITEM* tables are initially consistent as defined in Clause 3.3.2.1, based on a random sample of at least 10 distinct values of *O_ORDERKEY*.
2. Submit at least 100 ACID Transactions from each of at least the number of execution streams (# query streams + 1 refresh stream) used in the reported throughput test (see Clause 5.3.4). Each transaction must use values of (*O_KEY*, *L_KEY*, *DELTA*) randomly generated within the ranges defined in Clause 3.1.6.2. Ensure that all the values of *O_ORDERKEY* chosen in Step 1 are used by some transaction in Step 2.
3. Re-verify the consistency of the *ORDERS*, and *LINEITEM* tables as defined in Clause 3.3.2.1 based on the same sample values of *O_ORDERKEY* selected in Step 1.

Consistency was tested as part of the durability tests.

Isolation Requirements

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

The six tests described here are designed to verify that the system under test is configured to support the required isolation levels, as defined in Clause 3.4.1. All Isolation Tests are performed using a randomly selected set of values (*P_KEY*, *S_KEY*, *O_KEY*, *L_KEY*, *DELTA*).

Isolation Test 1 - Read-Write Conflict with Commit

This test demonstrates 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.

Isolation Test 2 - Read-Write Conflict with Rollback

This test demonstrates 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.

Isolation Test 3 - Write-Write Conflict with Commit

This test demonstrates 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:

$$T2.L_EXTENDEDPRICE = T1.L_EXTENDEDPRICE + (DELTA1 * (T1.L_EXTENDEDPRICE / T1.L_QUANTITY))$$

Isolation Test 4 - Write-Write Conflict with Rollback

This test demonstrates 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 $T2.L_EXTENDEDPRICE = T1.L_EXTENDEDPRICE$.

Isolation Test 5 - Concurrent Read and Write Transactions on Different Tables

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

The following steps were performed:

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.

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

This test demonstrates that 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:

1. An ACID Transaction T1 was started, executing a modified version of Q1 against the qualification database.
2. Before T1 completed, another ACID Transaction T2 was started using randomly selected values of O_KEY, L_KEY and DELTA.
3. Transaction T3 was started like T1, executing a modified version of 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.

Durability Requirements

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

The test sponsor is required to guarantee that the test system will preserve the database and the effects of committed updates after recovery from any of the failures listed below. The intent of these tests is to demonstrate that all transactions whose output messages have been received by the driver have in fact been committed in spite of any single failure from the list in Clause 3.5.3 and that all consistency conditions are still met after the database is recovered.

Permanent Irrecoverable Failure of Any Single Durable Medium (Database Tables)

Permanent Irrecoverable Failure of Any Single Durable Medium (Database Log)

Guarantee the database and committed updates are preserved across a permanent irrecoverable failure of any single durable medium containing TPC-H database tables. The media to be failed is to be chosen at random by the auditor and cannot be specially prepared.

Guarantee the database and committed updates are preserved across a permanent irrecoverable failure of any single durable medium containing TPC-H recovery log data. The media to be failed is to be chosen at random by the auditor and cannot be specially prepared.

The OS was stored on two SATA SSDs in a RAID-1 array. The database tables were stored on four non-RAIDed NVMe SSDs. The database log was stored on a 2-disk NVMe RAID-1 array. The tests were conducted on the qualification database. The steps performed are shown below:

1. The consistency of the ORDERS and LINEITEM tables was verified.
2. Nine streams of ACID transactions were started. Each stream executed a minimum of 100 transactions.
3. While the test was running, one of the database log drives was removed. Since the database log array was RAID-1, the database continued processing.
4. While the test was running, one of the disks holding database table data was removed.
5. The 9 streams of ACID transactions failed and recorded their number of committed transactions in success files.
6. The database log was dumped to disk.
7. A checkpoint was issued to force a failure, terminating SQL Server.
8. The root database file was deleted, insuring SQL Server could not access the database when started.
9. The removed drives were replaced and the log array was rebuilt.
10. SQL Server was started and the database was restored from backup.
11. A command was issued directing the database to run through its roll-forward recovery.
12. The success file and the HISTORY table counts were compared and were found to match.
13. The consistency of the ORDERS and LINEITEM tables was verified.

System Crash

Memory Failure

SUT Power Failure

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

Guarantee the database and committed updates are preserved across the failure of all or part of memory (loss of contents) which may be caused by a loss of external power or the permanent failure of a memory board.

Guarantee the database and committed updates are preserved across the loss of all external power to the SUT for an indefinite time period.

These tests were all combined. The following steps were performed:

1. The consistency of the ORDERS and LINEITEM tables was verified.
2. Nine streams of ACID transactions were started. Each stream executed a minimum of 100 transactions.
3. The system was powered off.
4. Power was restored and the system booted.
5. The database was restarted.
6. The database went through a recovery period.
7. The success file and the HISTORY table counts were compared and were found to match.
8. The consistency of the ORDERS and LINEITEM tables was verified.

Clause 4 – Scaling and Database Population Related Items

Initial Cardinality of Tables

The cardinality (e.g., the number of rows) of each table of the test database, as it existed at the completion of the database load (see Clause 4.2.5), must be disclosed.

Table 4-1. Initial Cardinality of Tables

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

Distribution of Tables and Logs

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

The distribution of database tables and logs was the same for the measured and priced configurations except that the priced configuration had extra drives to hold database backups.

The database data files and tempdb were spread out across four ThinkSystem U.2 PM1733 3.84TB Entry NVMe PCIe 4.0 x4 SSDs. The database log was configured on a 2-disk RAID-1 array of ThinkSystem U.2 CM5-V 1.6TB Mainstream NVMe PCIe 3.0 x4 SSDs.

The database table and log distribution is shown in table 4-2.

Table 4-2. Physical Distribution of Tables and Logs

Controller	Drives	RAID	Volume Size (TB)	Format	Files	Mount Point
930-8i 2GB Flash PCIe 12Gb Adapter	2x 960GB SATA	RAID-1	0.872	XFS	OS, SQL Server	/
1610-8P NVMe Switch Adapter						
Slots 1 & 2	2x 1.6TB NVMe	RAID-1 (Software)	1.5	XFS	Log	/mnt/dblog
Slot 3	1x 3.84TB NVMe	No RAID	2.0	XFS	Data, tempdb	/mnt/data/data1
Slot 4	1x 3.84TB NVMe	No RAID	2.0	XFS	Data, tempdb	/mnt/data/data2
Slot 5	1x 3.84TB NVMe	No RAID	2.0	XFS	Data, tempdb	/mnt/data/data3
Slot 6	1x 3.84TB NVMe	No RAID	2.0	XFS	Data, tempdb	/mnt/data/data4
Slot 7	1x 3.84TB NVMe	No RAID	2.0	XFS	Backup (Priced Cfg)	/mnt/backup
Slot 8	1x 3.84TB NVMe	No RAID	2.0	XFS	Backup (Priced Cfg)	/mnt/backup2
Totals (TB)	28.16		14.37			

Database Partition / Replication Mapping

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

The database was not replicated.

Storage Redundancy

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

The Storage Redundancy Level of this result is zero (no redundancy).

The operating system and database software were placed on a RAID-1 array.

RAID-1 was used for the array that held the database log.

The database data and the temporary tablespace were placed on non-RAIDed drives.

Database backups were placed on non-RAIDed drives.

DBGen

The version number, release number, modification number, and patch level of DBGen must be disclosed.

The standard distribution DBGen version 2.18.0 was used for database population. No modifications were made.

Database Load Time

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

The database load time was 4h 2m 9s.

Data Storage Ratio

The data storage ratio must be disclosed. It is computed by dividing the total data storage of the priced configuration (expressed in GB) by the size chosen for the test database as defined in Clause 4.1.3.1.

The calculation of the data storage ratio is shown in table 4-3.

Table 4-3. Data Storage Ratio Calculations

Disk Type	Number of Disks	Space per Disk (GiB)	Total Disk Space (GiB)	Scale Factor	Storage Ratio
960GB SATA SSD	2	894.07	1,788.14		
1.6TB U.2 PCIe 3.0 x4 NVMe SSDs	2	1,490.12	2,980.23		
3.84TB U.2 PCIe 4.0 x4 NVMe SSDs	6	3,576.28	21,457.67		
Total			26,226.04	3,000	8.74

The data storage ratio is 8.74, derived by dividing 26,226.04 GiB by the database size of 3,000 GiB.

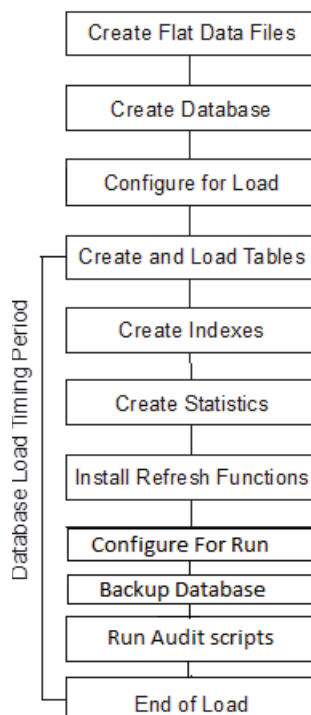
Database Load Details

The details of the database load must be disclosed including a block diagram illustrating the overall process.

Flat files for each of the tables were created using DBGen.

The tables were loaded as depicted in Figure 4.1.

Figure 4-1. Database Load Procedure



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 and disk structure to create and load the data with adjustments for size difference.

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.

The memory to database size percentage is 34.1%, derived by multiplying 1,024 GB * 100 and then dividing by the database size of 3,000 GB.

Clause 5 – Performance Metrics and Execution Rules Related Items

System Activity between Load and Performance Tests

Any system activity on the SUT that takes place between the conclusion of the load test and the beginning of the performance test must be fully reported in the supporting files archive including listings of scripts, command logs and system activity.

There was no activity between the load test and performance test.

Power Test Steps

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

The following steps were used to implement the power test:

1. RF1 Refresh function
2. Stream 00 Execution
3. RF2 Refresh function

Timing Intervals for Each Query and Refresh Function

The timing intervals (see Clause 5.3.7) for each query 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.

Number of Query Streams Used for the Throughput Test

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

Eight query streams and one refresh stream were used for the throughput test.

Start and End Date/Times for Each Query Stream

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

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

Total Elapsed Time for the Measurement Interval

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

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

Refresh Function Start Date/Time and Finish Date/Time

The start time and finish time for each refresh function in the refresh stream for the throughput test must be disclosed.

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

Performance Metrics

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

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

Performance Metric and Numerical Quantities from Both Runs

The performance metric (QphH@Size) and the numerical quantities (TPC-H Power@Size and TPC-H Throughput@Size) from both of the runs must be disclosed (see Clause 5.4).

Two consecutive runs of the TPC-H benchmark were performed. Table 5-1 contains the results for both runs.

Table 5-1. Performance Metrics from Both Runs

	QppH @ 3,000GB	QthH @ 3,000GB	QphH @ 3,000GB
Run 1	881,893.8	465,365.7	640,627.1
Run 2	853,014.6	457,609.5	624,778.0

System Activity between Tests

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

There was no activity between Run1 and Run2.

Clause 6 – SUT and Driver Implementation Related Items

Driver

A detailed textual description of how the driver performs its functions, how its various components interact and any product functionalities or environmental settings on which it relies and all related source code, scripts and configuration files must be reported in the supporting files archive. The information provided should be sufficient for an independent reconstruction of the driver.

The TPC-H benchmark was implemented using a Microsoft tool called StepMaster, which is a general purpose test tool that 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 in 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.

Implementation Specific Layer

If an implementation specific layer is used, then a detailed description of how it performs its functions, how its various components interact and any product functionalities or environmental setting on which it relies must be disclosed. All related source code, scripts and configuration files must be reported in the supporting files archive. The information provided should be sufficient for an independent reconstruction of the implementation specific layer.

See the Driver section above for details.

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.

Clause 7 – Pricing Related Items

Hardware and Software Components

A detailed list of hardware and software used in the Priced Configuration must be reported. The listing for each separately Orderable item must have vendor Part Number, description and applicable release/revision level, price source, unit price, quantity, extended price, applicable Discounted price and 3-year maintenance price. If package-pricing is used, the vendor Part Number of the package and a description uniquely identifying each of the Components of the package must be disclosed to a sufficient level of detail to meet the requirements of 1.4.1.1.

A detailed list of all hardware, software, and maintenance is provided in the Executive Summary at the front of this report. Price quotations are included in Appendix A.

Three-Year Cost of System Configuration

The total 3-year price of the entire Priced Configuration must be reported, including: hardware, software, and maintenance charges. The justification of any Discounts applied must be disclosed in the price sheet. Sufficient detail of what items are being discounted and by how much they are being discounted must be provided so that the Discount amount used in the computation of the total system cost can be independently reproduced.

A detailed list of all hardware, software, and maintenance, including the total 3-year price and discount information, is provided in the Executive Summary at the front of this report. Price quotations are included in Appendix A.

Availability Date

The committed Availability Date of Components used in the price calculations must be reported. The Availability Date must be reported on the first page of the Executive Summary and with a precision of one day. When the priced system includes products with different availability dates, the reported Availability Date for the priced system must be a date at which all Components are committed to be Generally Available. Each Component used in the Priced Configuration is considered to be Available on the Availability Date unless an earlier date is specified.

For each of the Line Items 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 Line Item 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 the order date arrives*
- *The method for verifying the price*

The total solution as priced will be generally available April 28, 2021.

All items are Orderable on the report date of the FDR.

Country-Specific Pricing

Pricing must be reported in the currency of the country where the system is priced.

The configuration is priced for the United States of America.

Pricing Calculations

A statement of the benchmark performance metric, as well as the respective calculations for 3-year pricing, price/performance, and the availability date must be included.

The performance metric, pricing calculations, price/performance, and availability dates are all included in the Executive Summary.

Clause 8 – Full Disclosure Report Related Items

Supporting Files Index

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.

Table 8-1. Supporting Files Index

Clause	Description	Pathname
Clause 1	OS and database settings	SupportingFilesArchive\Clause1
Clause 2	Qualification queries and output	SupportingFilesArchive\Clause2
Clause 3	ACID scripts and output	SupportingFilesArchive\Clause3
Clause 4	Database load scripts	SupportingFilesArchive\Clause4
Clause 5	Queries and output for measured runs	SupportingFilesArchive\Clause5
Clause 6	Implementation code for measured runs	SupportingFilesArchive\Clause6
Clause 7	There are no required files for Clause 7	
Clause 8	RFs source and parameters	SupportingFilesArchive\Clause8

Clause 9 – Audit Related Items

Auditor

The auditor's agency name, address, phone number, and attestation letter with a brief audit summary report indicating compliance must be included in the full disclosure report. A statement should be included specifying whom to contact in order to obtain further information regarding the audit process.

This implementation of the TPC Benchmark H was audited by Doug Johnson of Infosizing. Further information regarding the audit process may be obtained from:

InfoSizing
63 Lourdes Drive
Leominster, MA 01453
Telephone: (978) 343-6562
Web address: www.sizing.com

For a copy of this disclosure, go to www.tpc.org

Attestation Letter

The auditor's Attestation Letter is on the next two pages.

Benchmark sponsor: Lenovo
Michael Crutcher
Manager - DCG Performance Development
7001 Development Drive
Morrisville, NC 27560

March 25, 2021

I verified the TPC Benchmark H (TPC-H™ v2.18.0) performance of the following configuration:

Platform: Lenovo ThinkSystem SR665
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 624,778.0 QphH@3,000GB

TPC-H Power 853,014.6
TPC-H Throughput 457,609.5
Database Load Time 00d 04h 02m 09s

Server

Lenovo ThinkSystem SR665, with:

CPU	2x AMD EPYC 72F3 8-Core Processor, 3.7 GHz, 256 MB L3 Cache		
Memory	1,024 GiB		
Disks	Qty	Size	Type
	2	960 GB	SATA 6 Gb SSD
	2	1.6 TB	NVMe
	6	3.84 TB	NVMe

In my opinion, these performance results were produced in compliance with the TPC requirements for the benchmark.

The following verification items were given special attention:

- The database records were defined with the proper layout and size
- The database population was generated using DBGen
- The database was properly scaled to 3,000GB and populated accordingly
- The compliance of the database auxiliary data structures was verified
- The database load time was correctly measured and reported

- The required ACID properties were verified and met
- The query input variables were generated by QGen
- The query text was produced using minor modifications and no query variant
- The execution of the queries against the SF1 database produced compliant answers
- A compliant implementation specific layer was used to drive the tests
- The throughput tests involved 8 query streams
- The ratio between the longest and the shortest query was such that no query timings were adjusted
- The execution times for queries and refresh functions were correctly measured and reported
- The repeatability of the measured results was verified
- The system pricing was verified for major components and maintenance
- The major pages from the FDR were verified for accuracy

Additional Audit Notes:

None.

Respectfully Yours,

A handwritten signature in black ink that reads "Doug Johnson". The signature is written in a cursive style with a long horizontal flourish extending to the right.

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
February 24, 2021

Lenovo Corporation
Ray Engler
7001 Development Drive
Morrisville, NC 27560

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 <i>2 Core License Open Program - No Level - ERP</i>	\$13,748.00	8	\$109,984.00
Support			
Microsoft Problem Resolution Services <i>Professional Support (1 Incident).</i>	\$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: TPC_H_dhwuohj5494279_2019