Dell Technologies

TPC Benchmark H Full Disclosure Report
For
Dell Technologies PowerEdge R940xa Server
While Using Microsoft SQL Server 2019 Enterprise Edition
and
Windows Server 2019 Standard Edition

First Edition: August 2020

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, August 2020

Abstract

Overview

This report documents the methodology and results of the TPC Benchmark H test conducted on Dell Technologies PowerEdge R940xa 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 Windows Server 2019 Standard Edition.

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

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

Standard and Executive Summary Statements

Pages iv−viii contain the Executive Summary and Numerical Quantities Summary of the benchmark results for the PowerEdge R940xa 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 Fancois Raab 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 R940xa Server**	TPC-H Rev. 3.0.0 TPC Pricing Rev. 2.6.0 Report Date: Aug 20, 2020
Total System Cost	Composite Query per Hour Rating	Price/Performance
$943,769 USD	**1,420,550.7** QphH@10,000 GB	**$ 664.37USD** $/kQphH@10,000GB

Database Size	Database Manager	Operating System	Other Software	Availability Date
10,000 GB	Microsoft® SQL Server 2019 Enterprise Edition	Windows Server 2019 Standard Edition	N/A	Aug 20, 2020

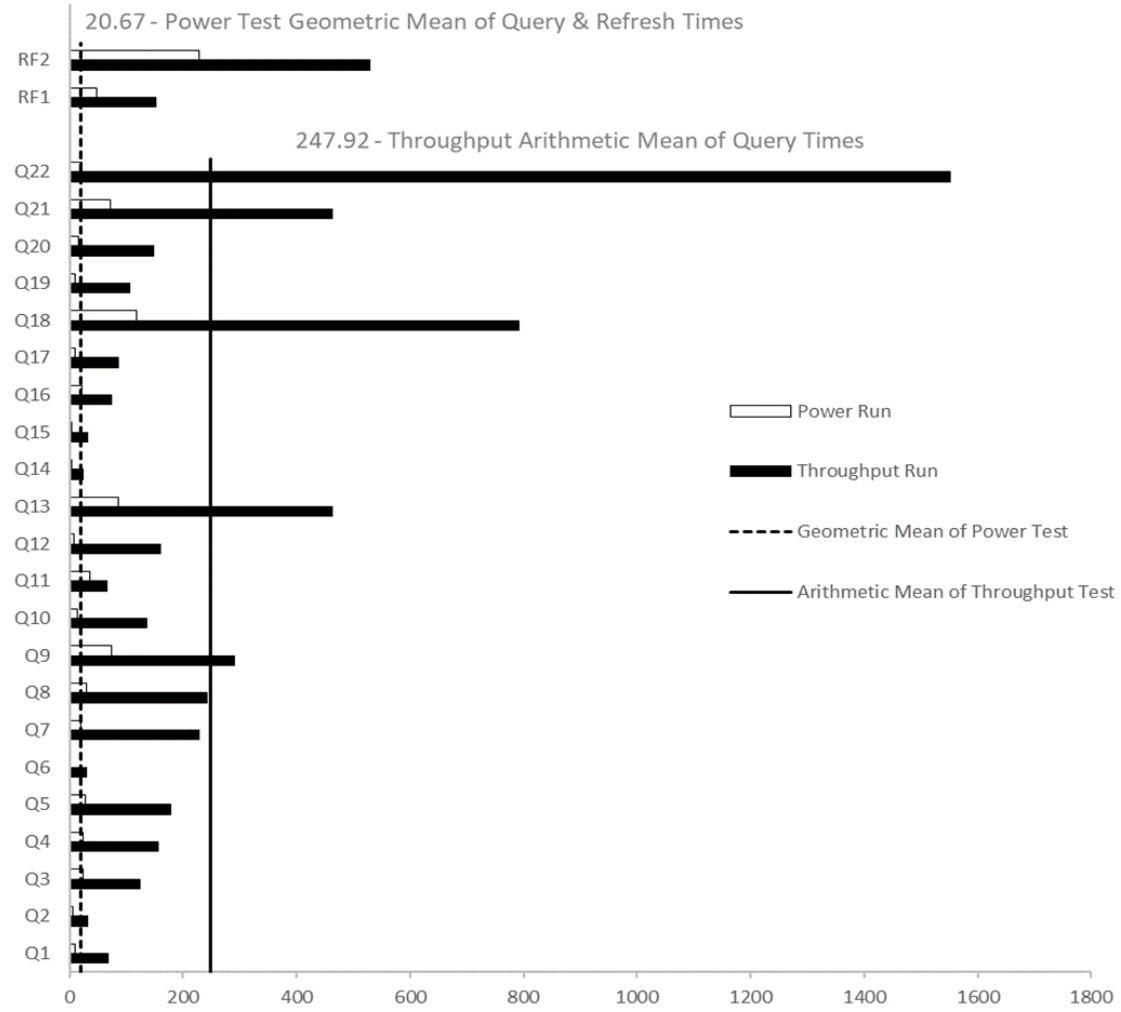

Database load Time = **00D 08H 08M 03S**	Storage Redundancy Level	
Load Includes Backup: **Y**	Base Tables and Auxiliary Data Structures	**No RAID**
Total Data Storage/Database Size = **3.78**	DBMS Temporary Space	**No RAID**
	Log	**RAID 10**
Percentage Memory/ Database size = **61.4** %	OS and DBMS Software	**RAID 1**
System Configuration:	Dell PowerEdge R940xa server	
Processors/Model:	**4x Intel Xeon Platinum 8280L 2.7G, 10.4GT/s, 38.5MCache, Turbo, HT, (205W)1.5TBDDR4-2666**	
Cores/Threads:	**112C/224T**	
Memory:	**6TB (48x 128GB)**	
Storage:	**10x 1.6TB SSD SAS Mix Use 12Gbps512e 2.5in Hot-plug AG Drive,3 DWPD, 8760 TBW**	
	4x Dell 6.4TB, NVMe, Mixed Use Express Flash, 2.5 SFF Drive, U.2, PM1725a with Carrier	
Total Storage:	**38745.32**	

PowerEdge R940xa Server

TPC-H Rev. 3.0.0
TPC Pricing Rev. 2.6.0
Report Date: Aug 20, 2020

Description	Part Number	Source	Unit Price	Qty	Ext. Price $	3 Yr Maint. Price
PowerEdge R940xa Server	210-AOQP	1	7,892.00	1	7,892.00	
PowerEdge R840/R940xa MLK Motherboard	329-BEOT	1	0.00	1	0.00	
No Trusted Platform Module	461-AADZ	1	0.00	1	0.00	
32 x 2.5" Chassis, 28 SAS/SATA + 4 NVME/SAS/SATA, Single-wide FPGA capable, 4 CPU Configuration	321-BDTC	1	2,999.00	1	2,999.00	
PowerEdge R940xa Shipping	340-CBCB	1	0.00	1	0.00	
PowerEdge R940xa Shipping Material, DAO	340-COPV	1	99.00	1	99.00	
Intel Xeon Platinum 8280L2.7G, 28C/56T, 10.4GT/s, 38.5M Cache, Turbo, HT (205W) DDR4-2933	338-BSHT	1	21,379.00	4	85,516.00	
4 CPU Heatsink	412-AANH	1	388.00	1	388.00	
2666MT/s LRDIMMs	370-ADNT	1	279.00	1	279.00	
Blank for 4CPU Configuration	370-AECN	1	0.00	1	0.00	
Performance Optimized	370-AAIP	1	0.00	1	0.00	
128GB LRDIMM, 2666MT/s, OctoRank	370-ADMX	1	4,749.00	48	227,952.00	
Unconfigured RAID	780-BCDS	1	0.00	1	0.00	
PERC H740P RAID Controller, 8GB NV Cache, Adapter, Full Height	405-AANR	1	999.00	1	999.00	
iDRAC9, Enterprise	385-BBKT	1	0.00	1	0.00	
iDRAC Group Manager, Disabled	379-BCQY	1	0.00	1	0.00	
iDRAC,Factory Generated Password	379-BCSF	1	0.00	1	0.00	
Riser Config4, 10 PCIe slots(10 x8 Full height slots), NVMe, 4 CPU	330-BBLB	1	569.00	1	569.00	
1.6TB SSD SAS Mix Use 12Gbps512e 2.5in Hot-plug AG Drive,3 DWPD, 8760 TBW	400-AZOY	1	2,199.00	10	21,990.00	
Dell 6.4TB, NVMe, Mixed Use Express Flash, 2.5 SFF Drive, U.2, P4610 with Carrier	400-BELR	1	7,918.00	4	31,672.00	
Broadcom 57412 Dual Port 10GbE SFP+ & 5720 Dual Port 1GbE BASE-T rNDC	540-BBUL	1	469.00	1	469.00	
Dell EMC PowerEdge SFP+ SR Optic 10GbE 850nm	407-BCBE	1	239.00	2	478.00	
No Internal Optical Drive	429-AAIQ	1	0.00	1	0.00	
1+1 Power Supply, Redundant Configuration, 2400W, 250 VoltPower Cord Required for	450-AHCJ	1	1,499.00	1	1,499.00	
No Power Cord	450-AAGG	1	0.00	1	0.00	
No Bezel	350-BBBW	1	0.00	1	0.00	
Dell EMC Luggage Tag	389-CHBB	1	0.00	1	0.00	
Quick Sync 2 (At-the-box mgmt)	350-BBOZ	1	149.00	1	149.00	
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	
No Rack Rails	770-BBBS	1	0.00	1	0.00	
No Systems Documentation, NoOpenManage DVD Kit	631-AACK	1	0.00	1	0.00	
Dell Hardware Limited Warranty Plus On-Site Service	818-7292	1	249.00	1	0.00	249.00
ProSupport Mission Critical: 7x24 HW/SW Technical Support & Assistance, 3 years	818-7402	1	5,668.00	1	0.00	5,668.00
ProSupport Mission Critical: 4-Hour 7x24 On-Site Service with Emergency Dispatch, 3 Ye	818-7394	1	3,031.00	1	0.00	3,031.00
Logitech MK120 Keyboard and Mouse	A6999510	1	15.99	1	15.99	
Dell Monitor	210-AIWG	1	169.99	1	169.99	
				Sub Total	383,135.98	8,948.00
Large Purchase Discount (58% for Hardware and 35% for Service)				Discounted Price	-222,218.87	-3,131.80
				Hardware Subtotal	160,917.11	5,816.20
SOFTWARE						
SQL SERVER 2019 Enterprice Edition for Windows (2 cores license, 28 Cores)	NA	2	13,748.00	56	769,888.00	
Windows Server 2019 Standard Edition(2core License)	NA	2	123.00	56	6,888.00	
Microsoft Problem Resolution Services	NA	2	259.00	1		259.00
				Software Subtotal	776,776.00	259.00
				Total	937,693.11	6,075.20

Source: 1: DELLEMC 2. MICROSOFT Audited by Francois Raab of InfoSizing (www.sizing.com)	Three Year cost Ownership	$943,769.00
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 @10000	1,420,550.70
	$/kQphH @10000	$664.37

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	: 3.78
Percentage Memory / Database Size	: 61.4%%
Start of Database Load	: 8/8/2020 11:03:57 PM
End of Database Load	: 8/9/2020 7:11:59 AM
Database Load Time h m s	: 00D 08H 08M 03S
Query Streams for Throughput Test	: 9
TPC-H Power	: 1,741,966.2
TPC-H Throughput	: 1,158,440.5
TPC-H Composite Query-per-Hour (QphH@10,000GB)	: 1,420,550.7
Total System Price over 3 Years ($ USD)	: USD 943,769.00
TPC-H Price/Performance Metric ($ USD/kQphH@10,000GB)	: USD 664.37

Measurement Interval

Measurement Interval in Throughput Test (Ts) : 6,153.1

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
	809071159	2020-08-09 10:09:42 / 2020-08-09 10:20:27	00:10:45	2020-08-09 10:08:52 / 2020-08-09 10:09:40	2020-08-09 10:20:29 / 2020-08-09 10:24:18
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	809071160	2020-08-09 10:24:17 / 2020-08-09 12:06:52	01:42:35	2020-08-09 10:24:17 / 2020-08-09 10:28:07	2020-08-09 10:28:07 / 2020-08-09 10:39:49
2	809071161	2020-08-09 10:24:17 / 2020-08-09 11:43:38	01:19:21	2020-08-09 10:39:50 / 2020-08-09 10:42:33	2020-08-09 10:42:33 / 2020-08-09 10:51:40
3	809071162	2020-08-09 10:24:17 / 2020-08-09 11:58:50	01:34:33	2020-08-09 10:51:40 / 2020-08-09 10:53:43	2020-08-09 10:53:43 / 2020-08-09 11:04:12
4	809071163	2020-08-09 10:24:18 / 2020-08-09 12:03:20	01:39:02	2020-08-09 11:04:12 / 2020-08-09 11:08:39	2020-08-09 11:08:39 / 2020-08-09 11:17:30
5	809071164	2020-08-09 10:24:18 / 2020-08-09 11:36:09	01:11:51	2020-08-09 11:17:30 / 2020-08-09 11:20:44	2020-08-09 11:20:44 / 2020-08-09 11:28:37
6	809071165	2020-08-09 10:24:18 / 2020-08-09 12:03:28	01:39:10	2020-08-09 11:28:37 / 2020-08-09 11:30:26	2020-08-09 11:30:26 / 2020-08-09 11:37:54
7	809071166	2020-08-09 10:24:18 / 2020-08-09 12:00:16	01:35:58	2020-08-09 11:37:54 / 2020-08-09 11:38:59	2020-08-09 11:38:59 / 2020-08-09 11:46:46
8	809071167	2020-08-09 10:24:18 / 2020-08-09 11:37:49	01:13:31	2020-08-09 11:46:47 / 2020-08-09 11:48:39	2020-08-09 11:48:39 / 2020-08-09 11:56:06
9	809071168	2020-08-09 10:24:18 / 2020-08-09 12:06:33	01:42:15	2020-08-09 11:56:06 / 2020-08-09 11:58:00	2020-08-09 11:58:00 / 2020-08-09 12:06:49

TPC-H Timing Intervals (Seconds)

Stream ID	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11	Q12
0	10.88	5.68	24.08	24.74	28.55	1.39	19.64	30.03	74.33	13.80	36.62	8.88
1	43.56	32.08	177.40	57.14	348.51	56.31	180.27	202.93	219.94	157.67	114.40	220.95
2	36.14	38.03	28.99	63.68	214.80	29.41	314.59	309.35	457.91	176.47	40.85	239.00
3	133.81	15.49	58.47	201.08	218.62	23.47	191.13	252.57	84.36	25.91	39.96	80.83
4	78.66	36.64	69.94	134.02	219.40	33.26	121.30	264.80	616.45	142.71	73.47	176.88
5	78.44	35.11	270.62	188.17	174.54	36.44	223.70	226.21	239.99	109.51	74.34	103.99
6	61.77	14.16	111.34	224.87	76.67	42.95	395.05	188.51	403.51	156.89	81.06	205.91
7	50.00	37.84	59.48	133.53	67.90	2.50	54.33	253.53	131.52	26.85	55.08	18.57
8	122.11	31.00	166.81	244.13	245.79	22.69	493.16	261.16	374.10	282.88	42.28	338.67
9	16.66	45.88	176.82	160.40	44.39	27.65	74.96	221.67	92.27	150.39	70.29	54.52
Ql Min	10.88	5.68	24.08	24.74	28.55	1.39	19.64	30.03	74.33	13.80	36.62	8.88
Ql Avg	63.20	29.19	114.40	143.18	163.92	27.61	206.81	221.08	269.44	124.31	62.84	144.82
Ql Max	133.81	45.88	270.62	244.13	348.51	56.31	493.16	309.35	616.45	282.88	114.40	338.67

Stream ID	Q13	Q14	Q15	Q16	Q17	Q18	Q19	Q20	Q21	Q22	RF1	RF2
0	85.16	4.53	4.78	22.50	9.57	118.36	9.71	16.33	72.81	21.81	47.27	229.13
1	520.85	12.94	24.40	61.43	73.30	877.41	131.78	153.85	702.69	1,783.31	228.26	702.39
2	508.79	45.10	40.86	126.37	126.85	1,070.14	105.74	555.32	138.31	93.02	163.80	546.50
3	190.96	15.09	5.99	60.21	117.44	772.40	15.70	59.60	305.44	2,803.20	122.84	628.39
4	463.60	50.59	38.77	111.46	134.26	352.76	129.32	47.93	605.57	2,039.90	266.68	530.89
5	380.25	32.55	46.81	49.77	119.93	749.13	250.69	102.11	699.67	118.00	194.18	472.49
6	725.46	11.13	47.92	58.58	18.29	388.64	17.27	29.67	184.77	2,505.16	108.64	448.17
7	197.75	9.35	12.09	32.03	34.02	1,274.31	77.09	123.11	833.35	2,272.97	65.00	466.91
8	471.45	28.75	30.77	122.61	65.04	458.56	93.05	146.71	213.20	155.44	111.98	446.84
9	711.22	5.27	31.69	43.44	93.95	1,187.57	130.13	112.59	479.68	2,202.79	113.77	528.90
Ql Min	85.16	4.53	4.78	22.50	9.57	118.36	9.71	16.33	72.81	21.81	47.27	229.13
Ql Avg	425.55	21.53	28.41	68.84	79.27	724.93	96.05	134.72	423.55	1,399.56	142.24	500.06
Ql Max	725.46	50.59	47.92	126.37	134.26	1,274.31	250.69	555.32	833.35	2,803.20	266.68	702.39

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 R940xa server depicted in Figure 0.1, consisted of:

- Dell Technologies PowerEdge R940xa server
  - Intel Xeon Platinum 8280L 2.7G, 28C/56T=4 no.
  - Memory size=48×128GB.

- Disk Drives (HDDs)
  - 1.6 TB SSD SAS Mix Use 12Gbps512e, 2.5in Hot-plug AG drive, 3 DWPD, 8,760 TBW=10 no.
  - Dell 6.4 TB, NVMe, Mixed Use Express Flash, 2.5 SFF drive, U.2, PM1725a with Carrier=4 no.

- PowerEdge RAID Controller (PERC)
  - PERC H740P RAID Controller, 8GB NV Cache, Adapter, Full Height=1 no.

Figure 0.1 Benchmark and priced configuration for PowerEdge R940xa server

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

- Dell Technologies PowerEdge R940xa server
  - Intel Xeon Platinum 8280L 2.7G, 28C/56T=4 no.
  - Memory size=48×128GB.

- Disk Drives (HDDs)
  - 1.6 TB SSD SAS Mix Use 12Gbps512e, 2.5in Hot-plug AG drive, 3 DWPD, 8,760 TBW=10 no.
  - Dell 6.4 TB, NVMe, Mixed Use Express Flash, 2.5 SFF drive, U.2, PM1725a with Carrier=4 no.

- PowerEdge RAID Controller (PERC)
  - PERC H740P RAID Controller, 8GB NV Cache, Adapter, Full Height=1 no.

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 2.18.0 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 2.18.0 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. At least 100 ACID Transactions were submitted.
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. T2 completed before T1 completed.
4. It was verified that the appropriate rows in the ORDER, LINEITEM and HISTORY tables were changed.

3.5 Durability Requirements

The tested system must guarantee durability: the ability to preserve the effects of committed transactions and ensure database consistency after recovery from any one of the failures listed in Clause 3.5.2.

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 OS disk followed immediately by the 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 and OS Disk Removal Test

1. The complete database was backed up.
2. The Consistency of the ORDERS and LINEITEM tables were verified.
3. 10 streams of ACID transactions were started. Each stream executed a minimum of 100 transactions.
4. While the test was running, one of the disks from the database/log R10 array 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. Immediately following the OS Disk Removal While the test was running, one of the RAID-10 configured Log

disk (slot 3) from the array was removed.
8. It was determined that the test kept on running. SQL Server did not error.
9. The pulled disk was replaced with a new disk. Log disk eventually completed its RAID rebuild process without any issues.

Data Disk Removal Test

1. The complete database was backed up.
2. The Consistency of the ORDERS and LINEITEM tables were verified.
3. 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 new volume with same letter and file structure
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. 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	# of Rows
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 R940xa server with the following storage configuration:
The storage system consisted of:

- 4X Dell 6.4 TB, NVMe, Mixed Use Express Flash, 2.5 SFF Drive, U.2, PM1725a with Carrier
- 1X PERC H740P RAID Controller, 8GB NV Cache, Adapter, Full-Height
    - 10X 1.6 TB SSD SAS Mix Use 12 Gbps512e 2.5in Hot-plug AG Drive, 3 DWPD, 8,760 TBW

The drives were distributed as follows:
- 1489.88 GB drives were used for the OS, Client Kit, and DB root file
- 5961.63 GB drives were used to hold Test DB table data and temporary database (TempDB).
- 1489.88 GB drives were used to hold DB Log and tempdb log.
- 1489.88 GB drives were used to hold backup.

A description of distribution of database file groups and log can be found in the Table below.

# of Disks	Drive Description	RAID Format	Size (GB)	Partition Format	Drive Letter	Content
2	1.6 TB 2.5-inch SSD SAS Mix Use 12 Gbps	RAID 1	1,489.88	ReFS	C	OS, SQL, Client Kit, & DB root
1	Dell 6.4 TB, NVMe, Mixed Use Express Flash	No RAID	5,961	ReFS	D	DB & Tempdb
1	Dell 6.4 TB, NVMe, Mixed Use Express Flash	No RAID	5,961	ReFS	E	DB & TempDB
1	Dell 6.4 TB, NVMe, Mixed Use Express Flash	No RAID	5,961	ReFS	F	DB & TempDB
1	Dell 6.4 TB, NVMe, Mixed Use Express Flash	No RAID	5,961	ReFS	G	DB & TempDB
4	1.6 TB 2.5-inch SSD SAS Mix Use 12 Gbps	RAID 10	2,979.75	ReFS	H	DB Log and TempDB Log
3	1.6 TB 2.5-inch SSD SAS Mix Use 12 Gbps	RAID 0	4,469.63	ReFS	B	Backup

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 4 spindles

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 four 6TB 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 four 1.6 TB SATA SSDs. OS on RAID1 array of two 1.6TB SSD SAS drives. The database backup was hosted on another RAID-0 array made of three 1.6 TB SATA SSDs.

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 2.18.0 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 08H 08M 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 (GB)	Total Disk Space (GB)	Total Storage Capacity (GB)	Scale factor	Data Storage Ratio
4 x Dell 6.4 TB, NVMe, Mixed Use Express Flash	5,961.63	23,846.52	38,745.32	10,000	3.87
10 x 1.6 TB 2.5-inch SSD SAS Mix Use 12Gbps	1,489.88	14,898.8			

Size of test database: 10,000 GiB Data Storage Ratio: 3.87

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.

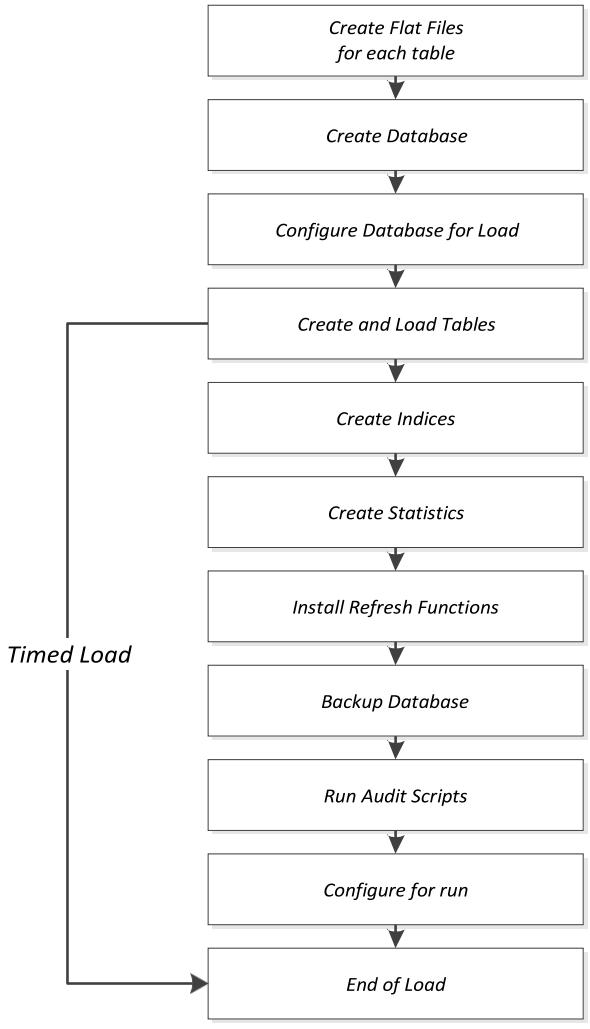

Figure 4.8: Block Diagram of Database Load Process

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=6,144 GB
- Scale Factor=10,000
- The memory to database size percentage=61.4%

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.

- There were only a couple minutes between the load and the run
- Restart the SQL Server database
-

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 by refresh stream
- RF2 Refresh Function by query stream

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	1,679,759.7	1,211,672	1,426,645.6
Run 2	1,741,966.2	1,158,440.5	1,420,550.7

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 58% discount from Dell Technologies, on all the hardware, and 35% on services.

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 AUG 26, 2020.

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 Francois Raab of InfoSizing, a certified TPC-H auditor. Further information regarding the audit process may be obtained from:

Francois Raab
InfoSizing (www.sizing.com)
20 Kreg Ln.
Manitou Springs, CO 80829
(719) 473-7555

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

Appendix A: Price Quotes

Microsoft Corporation
One Microsoft Way
Redmond, WA 98052-6399

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

Microsoft

July 22, 2020

Venkateswara Reddy Vatam
Dell
No. 65/2
Bagmane Tech Park
Bairasandra Main Rd
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	56	$769,888.00
Database Server Operating System			
Windows Server 2019 Standard 2 Core License Open Program – No Level - ERP	$123.00	56	$6,888.00
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 90 days.

Reference ID:

TPCH_djwyhdng4782604_2019