



Hewlett Packard
Enterprise

Hewlett-Packard Enterprise

TPC Benchmark™ H
Full Disclosure Report
for
HPE ProLiant DL325 Gen10
using
Microsoft® SQL Server 2017 Enterprise Edition
and
Red Hat® Enterprise Linux® 8

First Edition
August 2019

Hewlett Packard Enterprise (HPE), 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.

Copyright © 2018 Hewlett Packard Enterprise.

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

Printed in the United States, August 2019

HPE and HPE StorageWorks are registered trademarks of Hewlett Packard Enterprise.

Microsoft and SQL Server 2017 Enterprise Edition are registered trademarks of Microsoft Corporation.

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

All other brand or product names mentioned herein must be considered trademarks or registered trademarks of their respective owners.

Abstract

Overview

This report documents the methodology and results of the TPC Benchmark™ H test conducted on the HPE ProLiant DL325 Gen10 using Microsoft SQL Server 2017 Enterprise Edition in conformance with the requirements of the TPC Benchmark™ H Standard Specification, Revision 2.18.0. The operating system used for the benchmark was Red Hat® Enterprise Linux 8.

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 HPE ProLiant DL325 Gen10.

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 Francois 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."

 Hewlett Packard Enterprise		HPE ProLiant DL325 Gen10		TPC-H Rev. 3.0.0 TPC Pricing Rev. 2.7.0	
Total System Cost		Composite Query per Hour Rating		Report Date: August 7, 2019	
\$252,286 USD		743,750.6 QpH @ 1,000GB		\$339.21 USD \$/kQpH@ 1,000GB	
Database Size	Database Manager	Operating System	Other Software	Availability Date	
1000GB	Microsoft® SQL Server 2017 Enterprise Edition	Red Hat® Enterprise Linux® 8	None	August 26, 2019	



Database Load Time = 00d 6h 56m 29s	Storage Redundancy Level	
Load included backup: Y	Base Tables and Auxiliary Data Structures	No RAID
Total Data Storage / Database Size = 12.33	DBMS Temporary Space	No RAID
Memory/Database Size = 51.2%	OS and DBMS Software	RAID 1

System Components	System Total	Per Node
Nodes:	1	n.a.
AMD® EPYC ROME 7502P @ 2.50GHz:	1	1
Cores/Threads:	32/64	32/64
Memory:	512GB	512GB
Network Interfaces (on-board GigE):	1	1
(E208i-a embedded)	1	1
1.6TB NVMe Drives	8	n.a.
Disk Drives (240GB 6G SATA SSD):	2	n.a.
Total Storage:	16,393.8 GB	n.a.

 Hewlett Packard Enterprise	HPE ProLiant DL325 Gen10		TPC-H Rev. 3.0.0 TPC Pricing Rev. 2.7.0				
			Report Date: August 7, 2019				
Description	Price Key	Part Number	Unit Price	Qty	Extended Price	3 Yr Maint Price	
Server Hardware							
HPE DL325 Gen10 NVMe SFF CTO Server	1	P04662-B21	\$1,600	1	\$1,600		
HPE DL325 Gen10 7502P AMD Kit	1	P16639-L21	\$3,459	1	\$3,459		
HPE 64GB 2Rx4 DDR4-2933 R Smart Memory Kit	1	P19045-B21	\$2,129	8	\$17,032		
HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	1	865414-B21	\$379	1	\$379		
HPE USB US Keyboard/Mouse Kit	1	631341-B21	\$29	1	\$29		
Add to Rack 48U 600mmx1075mm G2 Advanced Pallet	1	P9K19A	\$1,499	1	\$1,499		
HPE C13-JIS C8303 JP 100V 12Amp 2.0m Power Cord	1	AF572A	\$49	1	\$49		
HPE 3Y Foundation Care 24x7 SVC	1	H7J34A3	\$2,448	1		\$2,448	
					Subtotal	\$24,047	\$2,448
Storage							
HPE 1.6TB NVMe x4 Lanes Mixed Use SFF (2.5in) SCN 3yr Wty Digitally Signed Firmware SSD	1	P10222-B21	\$3,049	8	\$24,392		
HPE Smart Array E208i-a SR Gen10 (8 Internal Lanes/No Cache) 12G SAS Modular LH Controller	1	869079-B21	\$299	1	\$299		
HPE 240GB SATA 6G Read Intensive SFF (2.5in) SC 3yr Wty Digitally Signed Firmware SSD	1	P05924-B21	\$409	2	\$818		
HPE 1U Gen10 2SFF SAS/SATA FIO Enablement Kit without Cables	1	881179-B21	\$59	1	\$59		
HPE DL325 Gen10 2SFF Smart Array SAS FIO Cable Kit	1	P05036-B21	\$30	1	\$30		
HPE DL325G10 Standard Riser	1	StandardDL325G10Riser		1			
					Subtotal	\$25,598	\$0
Other							
Other HP 23 Inch Monitor 3 HP	3	HP 23f	\$199	1	\$199		
					Subtotal	\$199	\$0
Server Software							
SQL Server 2017 Enterprise Edition for Linux (2 core license, 32 cores)	2	N/A	\$13,472.00	16	\$215,560	Included	
Microsoft Problem Resolution Services	2	N/A	\$259.00	1		\$259.00	
RHEL Svr 2 Sckt/2 Gst 3yr 24x7 E-LTU	1	G3J30AAE	\$3,702.00	1	\$3,702		
					Subtotal	\$219,262	\$259
					Total Extended Price	\$269,106	\$2,707
					Total Discounts	\$18,671	\$856
Price Key: 1 - HPE , 2 – Microsoft, 3 – Hewlett Packard Inc. Audited by Francois Raab of InfoSizing (www.sizing.com)			Grand Total		\$250,435	\$1,851	
All discounts are based on US list prices and for similar quantities and configurations. A 35% discount was based on the overall specific components pricing from vendor 1 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.			3 year cost of ownership USD:		\$252,286		
			QpH @ 1000GB:		743,750.6		
			\$ USD/kQpH @ 1000GB		\$339.21		
Sales contact: HPE WW Headquarters, 3000 Hanover St., Palo Alto, CA 94304-1185 (855)472-5233 / (800)786-7967 / (888)269-4073 (US) 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 at pricing@tpc.org. Thank you.							



Numerical Quantities

Measurement Results

Database Scale Factor	1,000
Total Data Storage / Database Size	12.33
Percentage Memory / Database Size	51.2%
Start of Database Load	2019-07-26 18:00:57
End of Database Load	2019-07-27 04:42:38
Database Load Time	00d 6h 56m 29s
Query Streams for Throughput Test	7
TPC-H Power	805,960.6
TPC-H Throughput	686,342.5
TPC-H Composite Query-per-Hour (QphH@1000GB)	743,750.6
Total System Price over 3 Years (\$ USD)	\$252,286
TPC-H Price/Performance Metric (\$ USD / kQphH@1000GB)	\$339.21

Measurement Interval

Measurement Interval in Throughput Test **807.76 seconds**

Duration of Stream Execution

Power Run	Seed	Query Start Time	Total Time (hh:mm:ss)	RF1 Start Time	RF2 Start Time
		Query End Time		RF1 End Time	RF2 End Time
	727044238	2019-07-27 04:49:09	00:02:21	2019-07-27 04:48:26	2019-07-27 04:51:30
		2019-07-27 04:51:30		2019-07-27 04:49:08	2019-07-27 04:52:01
Throughput Stream	Seed	Query Start Time	Total Time (hh:mm:ss)	RF1 Start Time	RF2 Start Time
		Query End Time		RF1 End Time	RF2 End Time
1	727044239	2019-07-27 04:52:01	00:13:16	2019-07-27 04:52:01	2019-07-27 04:53:01
		2019-07-27 05:05:17		2019-07-27 04:53:01	2019-07-27 04:53:45
2	727044240	2019-07-27 04:52:01	00:12:18	2019-07-27 04:53:45	2019-07-27 04:54:44
		2019-07-27 05:04:19		2019-07-27 04:54:44	2019-07-27 04:55:24
3	727044241	2019-07-27 04:52:01	00:12:44	2019-07-27 04:55:24	2019-07-27 04:56:25
		2019-07-27 05:04:45		2019-07-27 04:56:25	2019-07-27 04:57:09
4	727044242	2019-07-27 04:52:01	00:13:29	2019-07-27 04:57:10	2019-07-27 04:58:09
		2019-07-27 05:05:30		2019-07-27 04:58:09	2019-07-27 04:58:54
5	727044243	2019-07-27 04:52:01	00:13:18	2019-07-27 04:58:55	2019-07-27 04:59:51
		2019-07-27 05:05:19		2019-07-27 04:59:51	2019-07-27 05:00:33
6	727044244	2019-07-27 04:52:01	00:12:52	2019-07-27 05:00:34	2019-07-27 05:01:35
		2019-07-27 05:04:53		2019-07-27 05:01:35	2019-07-27 05:02:20
7	727044245	2019-07-27 04:52:01	00:12:54	2019-07-27 05:02:21	2019-07-27 05:03:19
		2019-07-27 05:04:55		2019-07-27 05:03:19	2019-07-27 05:04:08

TPC-H Timing Intervals (in seconds):

Query	Stream 0	Stream 1	Stream 2	Stream 3	Stream 4	Stream 5	Stream 6	Stream 7	Min Qi	Max Qi	Avg. Qi
Q1	3.05	6.14	22.18	17.14	5.88	30.30	24.38	7.64	3.05	30.30	14.59
Q2	1.55	5.17	7.44	13.97	7.38	8.24	10.02	7.19	1.55	13.97	7.62
Q3	4.09	18.72	39.43	14.85	4.16	67.06	12.75	66.25	4.09	67.06	28.41
Q4	4.35	12.21	43.60	6.52	30.22	19.21	51.23	31.38	4.35	51.23	24.84
Q5	5.16	34.87	11.46	14.88	34.54	81.48	46.10	37.10	5.16	81.48	33.20
Q6	0.22	6.72	0.38	4.92	6.36	0.78	5.50	10.52	0.22	10.52	4.43
Q7	3.83	13.42	15.35	41.57	16.03	40.66	18.39	11.55	3.83	41.57	20.10
Q8	3.97	19.14	63.75	41.87	17.02	21.08	62.64	67.73	3.97	67.73	37.15
Q9	21.00	27.58	51.06	71.79	44.81	20.24	59.76	54.68	20.24	71.79	43.87
Q10	4.08	43.26	45.63	42.56	65.57	3.89	45.01	37.90	3.89	65.57	35.99
Q11	6.27	23.88	12.69	11.78	12.41	10.80	23.88	12.08	6.27	23.88	14.22
Q12	2.05	78.42	69.18	85.77	25.36	61.96	19.57	35.74	2.05	85.77	47.26
Q13	26.11	120.41	77.01	81.55	49.62	29.80	46.62	79.42	26.11	120.41	63.82
Q14	1.28	1.95	2.69	2.78	4.75	5.20	13.83	5.42	1.28	13.83	4.74
Q15	1.38	6.88	9.21	6.91	6.35	9.56	5.06	7.17	1.38	9.56	6.57
Q16	2.73	12.88	9.55	7.20	17.75	21.82	15.16	7.83	2.73	21.82	11.87
Q17	1.91	31.93	35.63	28.22	53.46	17.21	5.24	13.49	1.91	53.46	23.39
Q18	21.07	114.96	72.25	41.24	210.43	123.07	134.16	109.47	21.07	210.43	103.33
Q19	0.86	36.66	15.46	87.56	54.09	36.77	42.52	45.06	0.86	87.56	39.87
Q20	3.74	13.55	27.63	23.21	3.58	11.92	15.75	15.91	3.58	27.63	14.41
Q21	12.77	138.01	73.76	58.03	92.91	119.05	86.98	75.14	12.77	138.01	82.08
Q22	8.89	28.32	31.58	58.39	44.95	56.62	26.26	33.85	8.89	58.39	36.11
RF1	42.71	59.12	58.91	60.53	58.62	55.85	61.57	58.29	42.71	61.57	56.95
RF2	31.49	43.51	39.39	44.44	45.29	42.13	45.15	48.23	31.49	48.23	42.45

Abstract.....	iii
Overview.....	iii
Standard and Executive Summary Statements	iii
Auditor	iii
0 General Items	1
0.1 Test Sponsor.....	1
0.2 Parameter Settings	1
0.3 Configuration Items	1
1.0 Clause 1: Logical Database Design	3
1.1 Table Definitions	3
1.2 Physical Organization of Database	3
1.3 Horizontal Partitioning.....	3
1.4 Replication	3
2.0 Clause 2: Queries and Refresh Functions - Related Items.....	4
2.1 Query Language.....	4
2.2 Random Number Generation	4
2.3 Substitution Parameters Generation.....	4
2.4 Query Text and Output Data from Database	4
2.5 Query Substitution Parameters and Seeds Used	4
2.6 Isolation Level	4
2.7 Refresh Functions	4
3.0 Clause 3: Database System Properties.....	5
3.1 Atomicity Requirements	5
3.1.1 Atomicity of the Completed Transactions	5
3.1.2 Atomicity of Aborted Transactions	5
3.2 Consistency Requirements.....	5
3.3 Isolation Requirements	6
3.4 Durability Requirements.....	8
4.0 Clause 4: Scaling and Database Population.....	9
4.1 Initial Cardinality of Tables	9
4.2 Distribution of Tables and Logs Across Media	10
4.3 Mapping of Database Partitions/Replications.....	11
4.4 Implementation of RAID	11
4.5 DBGEN Modifications	11
4.6 Database Load time.....	11
4.7 Data Storage Ratio	11
4.8 Database Load Mechanism Details and Illustration	12
4.9 Qualification Database Configuration	13
4.10 Memory to Database Size Percentage.....	13
5.0 Clause 5: Performance Metrics and Execution Rules Related Items.....	14
5.1 Steps after the Load Test.....	14
5.2 Steps in the Power Test.....	14
5.3 Timing Intervals for Each Query and Refresh Function.....	14
5.4 Number of Streams for The Throughput Test.....	14
5.5 Start and End Date/Times for Each Query Stream	14
5.6 Total Elapsed Time for the Measurement Interval	14
5.7 Refresh Function Start Date/Time and Finish Date/Time	14
5.8 Timing Intervals for Each Query and Each Refresh Function for Each Stream.....	15
5.9 Performance Metrics.....	15
5.10 The Performance Metric and Numerical Quantities from Both Runs	15

5.11 System Activity Between Tests	15
5.12 Documentation to satisfy Clause 5.2.7	15
5.13 Query Validation Output.....	15
6.0 Clause 6: SUT and Driver Implementation Related Items	16
6.1 Driver	16
6.2 Implementation Specific Layer (ISL)	16
6.3 Profile-Directed Optimization	16
7.0 Clause 7: Pricing Related Items	17
7.1 Hardware and Software Used	17
7.2 Three-Year Cost of System Configuration	17
7.3 Availability Dates	17
8.0 Clause 8: Support Files Index Table.....	18
8.1 Supporting Files Index Table.....	18
9.0 Clause 9: Audit Related Items	19
9.1 Auditors' Report	19
Appendix A: Price Quotes	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 Hewlett Packard Enterprise. The benchmark was developed and engineered by Hewlett Packard Enterprise. Testing took place at HPE facilities in Houston, TX.

0.2 Parameter Settings

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

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

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

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

0.3 Configuration 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 (e.g., DBMS, query processing tools /languages, middle-ware components, software drivers, etc.).*

The System Under Test (SUT), an HPE ProLiant DL325 Gen10, depicted in Figure 0.1, consisted of :

HPE ProLiant DL325 Gen10

- 1 AMD® EPYC ROME 7502P (2.50GHz/32-core)
- 8 HPE 64GB RDIMM (DDR4)

Storage adapters


- 1 E208i-a SR on-board storage adapter (boot drives and DB Root only)

Spindles

- 2 240GB 6G Read Intensive SSD
- 8 1.6 TB Mixed Use NVMe

Note: This system is the same system used for both the measured and priced configurations.

HP ProLiant DL325 Gen 10



HP ProLiant DL325 Gen 10
EPYC ROME 7502P (2.50GHz/32-core)
512 GB MEM: 8x64GB Dimms

8 x 1.6TB Mixed Use NVMe drives	Data files and tempdb data files Under software RAID1 for transaction log and DB backup
1 x E208i-a on board controller 2x240GB SATA 6G SSD	OS Boot and DB Root (RAID 1)

Figure 0.1 Benchmark and priced configuration for HPE ProLiant DL325 Gen10

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.

The Supporting Files Archive contains the DDL for the index definitions.

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 Random Number Generation

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

DBGEN version 2.17.3 and QGEN version 2.17.3 were used to generate all database populations. Version 2.18.0 of DBGEN and QGEN is identical to version 2.17.3, except for the version number.

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 2.17.3 was used to generate the substitution parameters. Version 2.18.0 of DBGEN and QGEN is identical to version 2.17.3, except for the version number.

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 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 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.

All ACID tests were conducted according to specification. The steps performed are outlined below.

3.1.1 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.

3.1.2 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.

3.2 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)

3.2.1 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.

Consistency was tested as part of the durability tests.

3.3 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.

3.3.1 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.

3.3.2 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.

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

3.3.4 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 $T2.L_EXTENDEDPRICE = T1.L_EXTENDEDPRICE$.

3.3.5 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:

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.

3.3.6 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:

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.

3.4 Durability Requirements

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

3.4.1 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

Two tests were completed in this section.

1st test was the removal of a log disk followed immediately by the removal of a data disk

2nd test was the 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 Disk Removal Test.

1. The complete database was backed up.
2. Eight streams of ACID transactions were started. Each stream executed a minimum of 100 transactions.
3. While the test was running, one of the disks from the database/log R10 array was removed.
4. It was determined that the test kept on running. SQL Server did not error.
5. The pulled disk was replaced with a new disk. Log disk eventually completed it's RAID rebuild process without any issues.

Data Disk Removal Test

1. Immediately following the Log Disk Removal, we initiated the data disk removal.
2. Benchmark Client was still running and transaction count was within range, allowing for the immediate testing of data disk removal.
3. While the test was running, one of the data disks (Non RAID) was removed.
4. A checkpoint was issued to force a failure.
5. Database errorlog recorded the checkpoint.
6. The eight streams of ACID transactions failed and recorded their number of committed transactions in success files.
7. The Database log was backed up and the Database was dropped.
8. The pulled disk was replaced with a new one.
9. The database was restored.
10. When database restore completed, issued a command to apply the backed up log file.
11. 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. Eight streams of ACID transactions were started. Each stream executed a minimum of 100 transactions.
2. While the streams of ACID transactions were running the System was powered off by pulling power plugs.
3. When power was restored the system booted and the Database engine was restarted.
4. The database went through a recovery period.
5. Rolled forward, Rolled backward transactions captured by DB ERRORLOG file.
6. Recovery complete.
7. The counts in the history table and success files were compared and verified, and the consistency of the ORDERS and LINEITEM tables was verified.

3.4.2 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 3.4.1

3.4.3 Memory Failure

Guarantee the database and committed updates are preserved across failure of all or part of memory (loss of contents).

See section 3.4.1

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	5,999,989,709
Orders	1,500,000,000
Partsupp	800,000,000
Part	200,000,000
Customer	150,000,000
Supplier	10,000,000
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 an HPE ProLiant DL325 Gen10 with the following configuration:

NVMe Storage Configuration:					
Physical Device	OS Device	Mount	Physical Size In TB	Available Size in TiB	Usage Comments
Box 1 Bay 2	/dev/md0	/var/opt/mssql/data/dblog	1.6	1.45	1 x transaction log 1 x tempdb transaction log
Box 1 Bay 1			1.6	1.45	
Box 1 Bay 7	Not Static	/mnt/data/data1	1.6	1.45	128 x data files 64 x tempdb files
Box 1 Bay 8	Not Static	/mnt/data/data2	1.6	1.45	
Box 1 Bay 4	Not Static	/mnt/data/data3	1.6	1.45	
Box 1 Bay 3	Not Static	/mnt/data/data4	1.6	1.45	
Box 1 Bay 6	Not Static	/mnt/data/data5	1.6	1.45	
Box 1 Bay 5	Not Static	/mnt/data/data6	1.6	1.45	

Traditional SATA Disk Configuration						
Controller	Device Location	Logical Drive	RAID Level	Mount Point	Spec Size(GB)	Available Size(GiB)
E208i-a Internal Controller	Port 1I Box 2 Bay 1	Logical Drive 01	RAID1	OS Mounts	240	223.52
	Port 1I Box 2 Bay 2				240	223.52

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 partition granularity is by week.

The database partitions are evenly distributed across 12 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.

RAID 0 was used for database file groups and TempDB. RAID 10 for the database recovery logs and tempDB recovery logs. RAID10 for the Backup drives.

4.5 DBGEN Modifications

The version number, release number, modification number, and patch level of DBGEN must be disclosed. Any modifications to the DBGEN (see Clause 4.2.1) source code must be disclosed. In the event that a program other than DBGEN was used to populate the database, it must be disclosed in its entirety.

DBGEN version 2.17.3 was used, no modifications were made.

Version 2.18.0 of DBGEN and QGEN is identical to version 2.17.3, except for the version number.

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 00d 6h 47m 21s

Load	
Load Start	07/26/2019 18:00:57
Load End	07/27/2019 04:42:38
Load Delay	00d 03h 45m 13s
Load Time	00d 6h 56m 29s

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.

Disk Type	# of Disks	RAW SIZE	Total Size
1.6TB NVMe	8	1.45 TiB	12.04 TiB 12325.44 GiB
240GB SAS SSD	2	223.52 GiB	

Size of test database: 1,000GiB

Data Storage Ratio: 12.33

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.

The memory to database size percentage is 51.2%

Memory/Database Size		
Physical Mem GB	Scale Factor	Percentage
512	1000	51.20%

5.0 Clause 5: Performance Metrics and Execution Rules Related Items

5.1 Steps after the Load Test

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

The queries were generated using QGen with the seed collected at the end of the load test.
There was only a couple minutes between the load and the run.
The system sat idle, without any interruption during this time frame.

5.2 Steps in the Power Test

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

The following steps were used to implement the power test:

1. Execute RF1 from the update stream
2. Execute Power queries from the query stream
3. Execute RF2 from the update 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.

The timing intervals for each query and both refresh functions are given in the Numerical Quantities Summary earlier in this document on page vii.

5.4 Number of Streams for The Throughput Test

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

7 query streams were used for the Throughput Test. Each stream running all 22 queries.
1 stream was used for RF.

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.

The Numerical Quantities Summary on page vii contains the start and stop times for the query execution streams run on the system reported.

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.

The Numerical Quantities Summary on page vii contains the timing intervals for the throughput test run on the system reported.

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.

The Numerical Quantities Summary on page vii contains the start and finish times for the refresh functions of each stream.

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.

The timing intervals for each query and each update function are given in the Numerical Quantities Summary earlier in this document on page vii.

5.9 Performance Metrics

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

The Numerical Quantities Summary contains the performance metrics, related numerical quantities, and the price/performance metric for the system reported.

5.10 The Performance Metric and Numerical Quantities from Both Runs

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

Run ID	QppH@1000G	QthH@1000G	QphH@1000G
Run 1	805,960.6	686,342.5	743,750.6
Run 2	852,885.7	659,638.8	750,064.3

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 only a couple minutes between Run1 and Run2. The Database log confirmed that no database activity took place during that time.

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 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.

The pricing summary sheet is given on page v in the Executive Summary at the front of this report. The source for all prices is indicated.

Server and all storage components are available at publication date.

The pricing and availability of the Microsoft software used is given in a quote from Microsoft, which is included in this report 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.

The pricing summary sheet on page v in the front of this report contains all details.

7.3 Availability Dates

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 (see Clause 7.3.1.4). All availability dates, whether for individual components or for the SUT as a whole, must be disclosed to a precision of 1 day, but the precise format is left to the test sponsor.

Category	Available
Server Hardware	August 26, 2019
Storage	Now (date of publication)
Server Software	Now (date of publication)
SQL Server 2017	Now (date of publication)

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	Pathname
Clause 1	DB Creation/Load Scripts– QUAL	n/a	SupportingFiles/Clause1/DB_Creation_QualDB
	DB Creation/Load Scripts – TEST	n/a	SupportingFiles/Clause1/DB_Creation_TestDB
	OS_DB_Settings	n/a	SupportingFiles/Clause1/OS_DB_Settings
Clause 2	Query Validation Output	n/a	SupportingFiles/Clause2/QueryValidationOutput
Clause 3	Atomicity	n/a	SupportingFiles/Clause3/Atomicity
	Consistency	n/a	SupportingFiles/Clause3/Consistency
	Isolation	n/a	SupportingFiles/Clause3/Isolation
	Durability	benchmark_scripts.zip	SupportingFiles/Clause3/Durability
Clause 4	DB Load Scripts	benchmark_scripts.zip	SupportingFiles/Clause4/DB_Load_Scripts
Clause 5	Query Output Run 1	Run1.zip	SupportingFiles/Clause5/Run1
	Query Output Run 2	Run2.zip	SupportingFiles/Clause5/Run2
	Queries	Queries.zip	SupportingFiles/Clause5/Queries
Clause 6	Implementation Specific layer source code and Driver	StepMaster_FDR.zip	SupportingFiles/Clause6/ClientReadme.txt
Clause 7	There are no files required to be included for Clause 7.	n/a	n/a
Clause 8	There are no files required to be included for Clause 8.	n/a	n/a

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.

Benchmark sponsor: Ankit Chouksey
Hewlett-Packard Enterprise
11445 Compaq Center Dr. West
Houston, TX 77070

July 31, 2019

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

Platform: HPE ProLiant DL325 Gen10
Operating System: Red Hat Enterprise Linux 8
Database Manager: Microsoft SQL Server 2017 Enterprise Edition
Other Software: n/a

The results were:

Performance Metric 743,750.6 QphH@1,000GB

TPC-H Power 805,960.6
TPC-H Throughput 686,342.5
Database Load Time 06h 56m 29s

Server HPE ProLiant DL325 Gen10, with:

CPU	1 x AMD EPYC ROME 7502P 2.50GHz processor (32 cores, 64 threads)		
Memory	512 GB		
Disks	Qty	Size	Type
	8	1.6 TB	SFF NVMe
	2	240 GB	SATA 6G SSD

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 1,000GB 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 7 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, appearing to read "François Raab". The signature is fluid and cursive, with a long horizontal stroke extending to the right.

François Raab, 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

July 29, 2019

Hewlett-Packard Enterprise
Ankit Chouksey
Survey 192, Whitefield Road
Mahadevapura Post
Bangalore, India 560048

Here is the information you requested regarding pricing for 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 2017 Enterprise Edition <i>2 Core License</i> <i>Open Program – No Level</i> Unit Price reflects a 5% discount from the retail unit price of \$14,256.00	\$13,472.50	16	\$215,560.00
Support			
Microsoft Problem Resolution Services <i>Professional Support</i> <i>(1 Incident).</i>	\$259.00	1	\$259.00

All Microsoft 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: TPC-H_VhmT2pjjqZjEJUWA5DeBRw7122__hac.doc