



Dell Inc.

TPC Benchmark™ H Full Disclosure Report

**Dell PowerEdge R6525**

(with 18x Dell PowerEdge R6525 Servers)

*using*

**Exasol 7.1**

**First Edition  
June 1, 2021**

*Dell PowerEdge R6525  
using  
Exasol 7.1*

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

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

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

**Copyright © 2021 Dell Inc. All rights reserved.**

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 on the title page of each item reproduced.

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

PowerEdge R6525 is a trademark of Dell.

Exasol is trademark of Exasol AG.

Ubuntu and Canonical are registered trademarks of Canonical Ltd.

AMD and EPYC are trademarks of AMD Corporation.

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

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



## Dell PowerEdge R6525 using Exasol 7.1

TPC-H Rev. 3.0.0  
TPC-Pricing Rev. 2.7.0

Report Date  
June 1, 2021

Total System Cost

Composite Query per Hour Metric

Price / Performance

**\$ 1,925,242**

**22,664,825.1**  
QphH@30000GB

**\$ 84.95**  
\$ / kQphH@30000GB

Database Size

Database Manager

Operating System

Other Software

Availability Date


**30,000GB**

**Exasol 7.1**

**Ubuntu 20.04.2  
LTS**

**None**

**July 1, 2021**



Database Load Time = 00d 02h 48m 59s  
 Load Includes Backup: N  
 Total Data Storage / Database Size = 4.11  
 Memory to Database Size Percentage = 61.4%

**Storage Redundancy Levels:**

Base Tables	Level Zero
Auxiliary Data Structures	Level Zero
DBMS Temporary Space	Level Zero
OS and DBMS Software	Level One

**System Configuration:**

18 x Dell PowerEdge R6525 Server, each with:  
 1024 GB RAM  
 2 x AMD EPYC 7543 32-Core Processor (64 threads)  
 2 x 480 GB SSD SATA drives  
 2 x 3200 GB (3.2TB) NVMe drives

**Total Storage:** 123,382 GiB



## Dell PowerEdge R6525 using Exasol 7.1

TPC-H Rev. 3.0.0  
TPC-Pricing Rev. 2.7.0

Report Date  
June 1, 2021

Description	Part	Key	Unit Price	Qty	Ext. Price \$	3 Yr. Maint. Price \$
<b>SERVER COMPONENTS</b>						
<b>PowerEdge R6525 Server</b>		1	85,162.93	18	1,532,932.74	
2.5 Chassis	379-BDTF	1		18		
SAS/SATA/NVMe Capable Backplane	379-BDSW	1		18		
PowerEdge R6525 Server	210-ATCF	1		18		
Trusted Platform Module 2.0	461-AAGY	1		18		
Chassis with up to 10 x 2.5" Hot Plug Hard Drives, + 2 Rear NVMe Drives (2CPU & XGMI)	321-BFNL	1		18		
AMD 7543 2.8GHz,32C/64T,256M,225W,3200	338-BZRQ	1		18		
AMD 7543 2.8GHz,32C/64T,256M,225W,3200	338-BZRU	1		18		
Additional Processor Selected	379-BDCO	1		18		
Heatsink for 1st CPU, greater than or equal to 180W (R6525)	412-AAASH	1		18		
Heatsink for 2nd CPU, greater than or equal to 180W (R6525)	412-AASJ	1		18		
Performance Optimized	370-AAIP	1		18		
3200MT/s RDIMMs	370-AEVR	1		18		
No RAID	780-BCDI	1		18		
PERC H345 Controller, Front	405-AAUW	1		18		
Front PERC Mechanical Parts, rear load	750-ACFQ	1		18		
Performance BIOS Settings	384-BBBL	1		18		
UEFI BIOS Boot Mode with GPT Partition	800-BBDM	1		18		
Very High Performance Fan for 10 x 2.5" or 10 x 2.5" + 2 x 2.5" chassis	384-BCGH	1		18		
Dual, Hot-plug, Redundant Power Supply (1+1), 1400W	450-AIQZ	1		18		
Riser Config 5, 1 x 16 LP PCIe slot (CPU2)	330-BBNV	1		18		
R6525 MB.w2x1Gb LOM(BCM5720)MLK	384-BCUO	1		18		
iDRAC9,Enterprise 15G	385-BBOT	1		18		
Mellanox ConnectX-5 Dual Port 10/25GbE SFP28, OCP NIC 3.0	540-BCOF	1		18		
LCD Bezel	325-BCHG	1		18		
Dell EMC Luggage Tag (x8 or x10 chassis)	350-BBXP	1		18		
No Quick Sync	350-BBXM	1		18		
iDRAC,Legacy Password	379-BCSG	1		18		
iDRAC Group Manager, Enabled	379-BCQV	1		18		
ReadyRails Sliding Rails	770-BDMS	1		18		
BOSS Blank	403-BCID	1		18		
No Systems Documentation, No OpenManage DVD Kit	631-AACK	1		18		
PowerEdge R6525 Shipping NBIS	340-COGR	1		18		
PowerEdge R6525 Shipping Material for x4/10 Chassis	340-COXQ	1		18		
PowerEdge R6525, No CE or CCC Marking	389-DTIP	1		18		
US Order	332-1286	1		18		
Basic Next Business Day 36 Months	709-BBFM	1		18		
Prosupport Plus and 4Hr Mission Critical Initial, 36 Month(s)	865-BBNF	1		18		
On-Site Installation Declined	900-9997	1		18		
64GB RDIMM, 3200MT/s, Dual Rank	370-AEVP	1		288		
480GB SSD SATA Mix Use 6Gbps 512 2.5in Hot-plug AG Drive, 3 DWPD, 2628 TBW	400-AZUT	1		36		
3.2TB Enterprise NVMe Mixed Use AG Drive U.2 Gen4 Flex Bay	400-BKGB	1		36		
Power Cord - C13, 3M, 125V, 15A (North America, Guam, North Marianas, Philippines, Samoa, Vietnam)	450-AALV	1		36		
Mellanox ConnectX-5 EX Dual Port 40/100GbE QSFP28 Adapter, PCIe Low Profile	540-BCIX	1		18		
GCP Operations Management	929-8509	1		18		
Mellanox® Passive Copper cable, ETH 100GbE, 100Gb/s, QSFP28, 2m, Black, 30AWG, CA-N	AA319963	1	96.67	38	3,673.46	
APC NetShelter SX Deep Enclosure - Rack - cabinet - black - 24U - 19-inch	A7522217	1	1,129.99	1	1,129.99	
APC Basic Rack PDU AP9567 - 0U - 120V NEMA 5-15 Input / (14) NEMA 5-15 Output	A7541364	1	189.99	1	189.99	
ASUS VE248Q 24 Inch LED monitor - Widescreen Full HD Monitor	A6732050	1	136.99	3	410.97	
Logitech MK200 Media Keyboard and Mouse Combo	A6859396	1	24.99	3	74.97	

C2G 6ft Cat6 Snagless Unshielded (UTP) Ethernet Network Patch Cable - Black - patch cable - 6 ft - black	A7052140	1	5.99	20	119.80	
Dell C13 to C14, PDU Style, 250 V Power Cord, North America - 6.5 feet	450-ACHI	1	11.93	38	453.34	
32-port Mellanox Spectrum SN2700 - Switch - L3 - managed - 32 x 100 Gigabit QSFP28	A8628811	1	20,008.99	1	20,008.99	
Mellanox M-1 Global Support Gold Plus. 2hr committed TAC response (24x7). 3 years	A9878902	1	3,234.06	1		3,234.06
PowerSwitch S3048-ON		1	15,066.51	1	15,066.51	
Dell Networking S3048-ON, 48x 1GbE, 4x SFP+ 10GbE ports, Stacking, IO to PSU air, 1x AC PSU, DNOS 9	210-AEDM	1		1		
ProSupport Plus: Mission Critical 4-Hour 7x24 On-Site Service with Emergency Dispatch, Initial Year	802-7419	1		1		
ProSupport Plus: 7x24 HW/SW Tech Support and Assistance, 3 Year	802-7434	1		1		
ProSupport Plus: Mission Critical 4-Hour 7x24 On-Site Service with Emergency Dispatch, 2 Year Extended	802-7435	1		1		
<b>Hardware Subtotal</b>					1,574,060.76	3,234.06
<b>SOFTWARE</b>						
Exasol 30TB License (monthly subscription)		2	25,000.00	36m	900,000.00	
<b>Software Subtotal</b>					900,000.00	
<b>Total</b>					<b>2,474,060.76</b>	<b>3,234.06</b>
Large Purchase Discount (35%)*					-550,921.27	-1,131.92

**Three-Year Cost of Ownership: \$ 1,925,242**  
**QpH@30000GB: 22,664,825.1**  
**\$/kQpH@30000GB: \$ 84.95**

Price Key: 1 DELL Waseem Raja  
2 Exasol Gunther Schweer [sales@exasol.com](mailto:sales@exasol.com)

\* Discount applies to all line items where Key = 1. Discount based upon total system cost as purchased by a regular customer.

Audited by Doug Johnson of InfoSizing ([www.sizing.com](http://www.sizing.com))

*Prices used in TPC benchmarks reflect the actual prices a customer would pay for a one-time purchase of the stated components. Individually negotiated discounts are not permitted. Special prices based on assumptions about past or future purchases are not permitted. All discounts reflect standard pricing policies for the listed components. For complete details, see the pricing 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](mailto:pricing@tpc.org). Thank you.*



# Dell PowerEdge R6525 using Exasol 7.1

TPC-H Rev. 3.0.0  
TPC-Pricing Rev. 2.7.0

Report Date  
June 1, 2021

## Numerical Quantities

### Measurement Results

Database Scale Factor	30,000 GB
Total Data Storage / Database Size	4.11
Percentage Memory / Database Size	61.4%
Start of Database Load	2021-05-13 11:38:29
End of Database Load	2021-05-13 14:27:28
Database Load Time	00d 02h 48m 59s
Query Streams for Throughput Test	11
TPC-H Power	24,086,987.1
TPC-H Throughput	21,326,631.4
TPC-H Composite Query-per-Hour Metric (QphH@30000GB)	22,664,825.1
Total System Price Over 3 Years	\$1,925,242
TPC-H Price/ Performance Metric (\$/kQphH@30000GB)	\$84.95

### Measurement Interval

Measurement Interval in Throughput Test (Ts)	1,225.51 seconds
--	------------------

### Duration of Stream Execution

Power Run	Seed	Query Start Time	Total Time	RF1 Start Time	RF2 Start Time
		Query End Time	(hh:mm:ss)	RF1 End Time	RF2 End Time
	513142728	2021-05-13 14:27:43	00:02:30	2021-05-13 14:27:34	2021-05-13 14:30:13
		2021-05-13 14:30:13		2021-05-13 14:27:43	2021-05-13 14:30:14
Throughput Stream	Seed	Query Start Time	Total Time	RF1 Start Time	RF2 Start Time
		Query End Time	(hh:mm:ss)	RF1 End Time	RF2 End Time
1	513142729	2021-05-13 14:30:13	00:20:27	2021-05-13 14:30:13	2021-05-13 14:30:42
		2021-05-13 14:50:40		2021-05-13 14:30:42	2021-05-13 14:30:45
2	513142730	2021-05-13 14:30:13	00:19:47	2021-05-13 14:30:45	2021-05-13 14:31:13
		2021-05-13 14:50:00		2021-05-13 14:31:13	2021-05-13 14:31:16
3	513142731	2021-05-13 14:30:14	00:19:52	2021-05-13 14:31:16	2021-05-13 14:31:47
		2021-05-13 14:50:06		2021-05-13 14:31:47	2021-05-13 14:31:51
4	513142732	2021-05-13 14:30:14	00:19:56	2021-05-13 14:31:51	2021-05-13 14:32:25
		2021-05-13 14:50:10		2021-05-13 14:32:25	2021-05-13 14:32:27
5	513142733	2021-05-13 14:30:14	00:19:47	2021-05-13 14:32:27	2021-05-13 14:32:57
		2021-05-13 14:50:01		2021-05-13 14:32:57	2021-05-13 14:33:00
6	513142734	2021-05-13 14:30:14	00:19:52	2021-05-13 14:33:00	2021-05-13 14:33:35
		2021-05-13 14:50:06		2021-05-13 14:33:35	2021-05-13 14:33:40
7	513142735	2021-05-13 14:30:14	00:19:45	2021-05-13 14:33:40	2021-05-13 14:34:16
		2021-05-13 14:49:59		2021-05-13 14:34:16	2021-05-13 14:34:20
8	513142736	2021-05-13 14:30:14	00:19:57	2021-05-13 14:34:20	2021-05-13 14:35:00
		2021-05-13 14:50:11		2021-05-13 14:35:00	2021-05-13 14:35:05
9	513142737	2021-05-13 14:30:14	00:19:53	2021-05-13 14:35:05	2021-05-13 14:35:51
		2021-05-13 14:50:07		2021-05-13 14:35:51	2021-05-13 14:35:56
10	513142738	2021-05-13 14:30:14	00:19:52	2021-05-13 14:35:56	2021-05-13 14:36:38
		2021-05-13 14:50:06		2021-05-13 14:36:38	2021-05-13 14:36:43
11	513142739	2021-05-13 14:30:14	00:19:50	2021-05-13 14:36:43	2021-05-13 14:37:30
		2021-05-13 14:50:04		2021-05-13 14:37:30	2021-05-13 14:37:35



**Dell PowerEdge R6525**  
**using Exasol 7.1**

TPC-H Rev. 3.0.0  
TPC-Pricing Rev. 2.7.0

Report Date  
June 1, 2021

**TPC-H Timing Intervals (in seconds)**

Stream ID	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11	Q12
0	8.32	2.84	6.52	1.89	7.57	1.40	10.57	3.60	29.27	6.49	20.28	3.22
1	8.70	6.85	86.06	2.00	68.77	34.74	153.35	40.21	25.25	92.20	43.47	59.18
2	113.48	5.23	61.11	20.35	61.10	16.11	102.68	37.71	218.29	54.62	40.74	35.80
3	124.02	4.73	67.37	20.87	43.16	16.90	90.13	35.93	228.13	53.42	36.89	35.26
4	118.07	5.43	70.34	21.60	52.47	18.76	95.94	36.14	241.52	59.00	37.59	33.79
5	107.10	4.16	72.10	19.23	46.61	19.73	94.05	29.27	204.85	55.53	38.81	34.13
6	118.49	4.05	61.47	23.59	46.92	18.89	104.64	33.35	243.54	55.25	37.22	34.58
7	124.73	5.40	67.38	23.49	47.53	15.70	103.61	27.74	241.43	52.27	38.69	34.77
8	110.30	4.98	63.92	21.82	48.05	21.44	103.06	28.31	234.16	60.29	34.67	41.63
9	114.96	4.26	68.11	24.50	48.46	19.11	94.51	36.28	212.65	72.00	40.12	33.35
10	120.98	5.35	71.85	23.37	45.44	20.82	104.76	29.89	201.69	62.42	36.54	35.52
11	125.43	4.14	67.95	25.83	54.74	15.97	96.51	37.88	208.37	59.41	44.01	32.53
Qi Min	8.32	2.84	6.52	1.89	7.57	1.40	10.57	3.60	25.25	6.49	20.28	3.22
Qi Avg	99.55	4.79	63.68	19.05	47.57	18.30	96.15	31.36	190.76	56.91	37.42	34.48
Qi Max	125.43	6.85	86.06	25.83	68.77	34.74	153.35	40.21	243.54	92.20	44.01	59.18

Stream ID	Q13	Q14	Q15	Q16	Q17	Q18	Q19	Q20	Q21	Q22	RF1	RF2
0	8.29	4.96	6.16	6.04	0.99	9.01	1.92	4.92	4.22	0.84	8.85	0.74
1	114.57	58.71	67.58	88.34	9.80	141.26	16.00	51.77	55.76	0.90	28.24	2.63
2	79.07	31.69	50.29	55.06	6.06	89.43	11.27	32.47	53.39	9.69	28.37	2.60
3	81.94	38.43	36.51	59.57	5.16	99.82	12.24	34.47	55.33	11.03	31.34	3.08
4	80.03	27.37	46.71	66.14	5.47	82.40	12.44	25.82	48.58	9.70	34.01	2.74
5	97.05	36.93	40.67	65.33	3.92	98.22	15.86	35.04	55.77	12.03	29.27	3.70
6	82.14	27.47	46.50	63.94	4.76	84.54	13.66	30.74	47.28	8.62	34.92	4.34
7	79.60	29.04	44.78	63.59	4.23	77.93	12.07	29.62	50.38	9.31	36.25	3.78
8	82.02	34.94	44.26	59.49	6.17	88.19	19.22	29.20	49.89	9.94	40.73	4.29
9	73.53	26.48	46.38	60.21	4.62	101.13	14.63	30.53	55.84	10.60	46.29	4.64
10	102.25	33.08	44.57	57.71	4.42	75.78	12.11	33.11	58.10	10.98	42.59	4.67
11	85.13	27.74	47.75	70.30	7.22	73.67	11.71	31.76	50.75	9.73	46.86	5.07
Qi Min	8.29	4.96	6.16	6.04	0.99	9.01	1.92	4.92	4.22	0.84	8.85	0.74
Qi Avg	80.47	31.40	43.51	59.64	5.24	85.12	12.76	30.79	48.77	8.61	33.98	3.52
Qi Max	114.57	58.71	67.58	88.34	9.80	141.26	19.22	51.77	58.10	12.03	46.86	5.07

# Table of Contents

<b>TPC BENCHMARK H OVERVIEW</b> .....	<b>10</b>
<b>0 GENERAL ITEMS</b> .....	<b>11</b>
0.1 BENCHMARK SPONSOR.....	11
0.2 PARAMETER SETTINGS.....	11
0.3 CONFIGURATION DIAGRAM.....	11
<b>1 CLAUSE 1: LOGICAL DATABASE DESIGN RELATED ITEMS</b> .....	<b>13</b>
1.1 DATABASE DEFINITION STATEMENTS.....	13
1.2 PHYSICAL ORGANIZATION.....	13
1.3 HORIZONTAL PARTITIONING.....	13
1.4 REPLICATION.....	13
<b>2 CLAUSE 2: QUERIES AND REFRESH FUNCTIONS</b> .....	<b>14</b>
2.1 QUERY LANGUAGE.....	14
2.2 VERIFYING METHOD FOR RANDOM NUMBER GENERATION.....	14
2.3 GENERATING VALUES FOR SUBSTITUTION PARAMETERS.....	14
2.4 QUERY TEXT AND OUTPUT DATA FROM QUALIFICATION DATABASE.....	14
2.5 QUERY SUBSTITUTION PARAMETERS AND SEEDS USED.....	14
2.6 ISOLATION LEVEL.....	14
2.7 SOURCE CODE OF REFRESH FUNCTIONS.....	14
<b>3 CLAUSE 3: DATABASE SYSTEM PROPERTIES</b> .....	<b>15</b>
3.1 ACID PROPERTIES.....	15
3.2 ATOMICITY REQUIREMENTS.....	15
3.3 CONSISTENCY REQUIREMENTS.....	15
3.4 ISOLATION REQUIREMENTS.....	16
3.5 DURABILITY REQUIREMENTS.....	17
<b>4 CLAUSE 4: SCALING AND DATABASE POPULATION</b> .....	<b>19</b>
4.1 ENDING CARDINALITY OF TABLES.....	19
4.2 DISTRIBUTION OF TABLES AND LOGS ACROSS MEDIA.....	19
4.3 MAPPING OF DATABASE PARTITIONS/REPLICATION.....	19
4.4 IMPLEMENTATION OF RAID.....	19
4.5 DBGEN MODIFICATIONS.....	20
4.6 DATABASE LOAD TIME.....	20
4.7 DATA STORAGE RATIO.....	20
4.8 DATABASE LOAD MECHANISM DETAILS AND ILLUSTRATION.....	20
4.9 QUALIFICATION DATABASE CONFIGURATION.....	20
4.10 MEMORY TO DATABASE SIZE PERCENTAGE.....	21
<b>5 CLAUSE 5: PERFORMANCE METRICS AND EXECUTION RULES RELATED ITEMS</b> .....	<b>22</b>
5.1 SYSTEM ACTIVITY BETWEEN LOAD AND PERFORMANCE TESTS.....	22
5.2 STEPS IN THE POWER TEST.....	22
5.3 TIMING INTERVAL FOR EACH QUERY AND REFRESH FUNCTIONS.....	22
5.4 NUMBER OF STREAMS FOR THE THROUGHPUT TEST.....	22
5.5 START AND END DATE/TIME OF EACH QUERY STREAM.....	22
5.6 TOTAL ELAPSED TIME OF THE MEASUREMENT INTERVAL.....	22
5.7 REFRESH FUNCTION START DATE/TIME AND FINISH DATE/TIME.....	22
5.8 PERFORMANCE METRICS.....	22
5.9 THE PERFORMANCE METRIC AND NUMERICAL QUANTITIES FROM BOTH RUNS.....	22
5.10 SYSTEM ACTIVITY BETWEEN PERFORMANCE TESTS.....	23
5.11 DOCUMENTATION TO SATISFY CLAUSE 5.2.7.....	23
5.12 QUERY OUTPUT VALIDATION.....	23
<b>6 CLAUSE 6: SUT AND DRIVER IMPLEMENTATION RELATED ITEMS</b> .....	<b>24</b>
6.1 DRIVER.....	24
6.2 IMPLEMENTATION SPECIFIC LAYER (ISL).....	24
6.3 PROFILE-DIRECTED OPTIMIZATION.....	24
<b>7 CLAUSE 7: PRICING</b> .....	<b>25</b>
7.1 HARDWARE AND SOFTWARE USED IN THE PRICED SYSTEM.....	25
7.2 TOTAL THREE YEAR PRICE.....	25
7.3 AVAILABILITY DATE.....	25
<b>8 CLAUSE 8: FULL DISCLOSURE</b> .....	<b>26</b>
8.1 SUPPORTING FILES INDEX TABLE.....	26
<b>9 CLAUSE 9: AUDIT RELATED ITEMS</b> .....	<b>27</b>



9.1 AUDITOR'S REPORT.....	27
<b>APPENDIX A: PRICING INFORMATION .....</b>	<b>30</b>

## TPC Benchmark H Overview

The TPC Benchmark™ H (TPC-H) is a Decision Support benchmark. It is a suite of business-oriented adhoc queries and concurrent modifications. The queries and the data populating the database have been chosen to have broad industry-wide relevance while maintaining a sufficient degree of ease of implementation. This benchmark illustrates Decision Support systems that:

- Examine large volumes of data
- Execute queries with a high degree of complexity
- Give answers to critical business questions

TPC-H evaluates the performance of various Decision Support systems by the execution of sets of queries against a standard database under controlled conditions. The TPC-H queries:

- Give answers to real-world business questions
- Simulate generated ad-hoc queries
- Are far more complex than most OLTP transactions
- Include a rich breadth of operators and selectivity constraints
- Generate intensive activity on the part of the database server component of the system under test
- Are executed against a database complying to specific population and scaling requirements
- Are implemented with constraints derived from staying closely synchronized with an on-line production database

## 0 General Items

### 0.1 Benchmark Sponsor

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

This TPC-H benchmark is sponsored by Dell Inc. The benchmark implementation was developed and engineered by Exasol AG.

### 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 the database and OS parameters used in this benchmark.

### 0.3 Configuration Diagram

*Diagrams of both measured and priced configurations must be provided, accompanied by a description of the differences. This includes, but is not limited to:*

- *Number and type of processors.*
- *Size of allocated memory, and any specific mapping/partitioning of memory unique to the test.*
- *Number and type of disk units (and controllers, if applicable).*
- *Number of channels or bus connections to disk units, including their protocol type.*
- *Number of LAN (e.g. Ethernet) Connections, including routers, workstations, terminals, etc., that were physically used in the test or are incorporated into the pricing structure.*
- *Type and the run-time execution location of software components (e.g., DBMS, query processing tools/languages, middle-ware components, software drivers, etc.).*

The System Under Test (SUT), depicted in Figure 1.1, that was used to obtain the results in this benchmark consists of the following components:

<b>System components:</b>		
18 x	Dell PowerEdge R6525 servers, each with:	
	1024	GB RAM
	2 x	AMD EPYC 7543 32-Core Processor (64 threads)
	2 x	480 GB SSD SATA Mix Use 6Gbp
	2 x	3,200 GB Enterprise NVMe Mixed Use
	1 x	PERC H345 Controller
	1 x	Mellanox ConnectX-5 EX Dual Port 40/100GbE QSFP28
	1 x	BCM5720 Dual Port 1GbE (LOM)
1 x	32-port 100GbE Mellanox Spectrum SN2700 Switch	
1 x	48-port 1GbE PowerSwitch S3048-ON	

Network is 1x 100GbE and 1x 1GbE Ethernet. Each server has 2x 480 GB SSD configured as 1x mirrored pair (HW RAID 1) used for OS installation. In addition each server contains 2x 3200GB NVMe drives configured as raw

devices used for data storage. NUMA aware DBMS software configuration was used: 36x logical instances (n11..n46) running on 18 physical nodes (m1-m18). The CPU hyperthreading was enabled.

Priced configuration and benchmarked configuration are identical.

### DBMS-Cluster

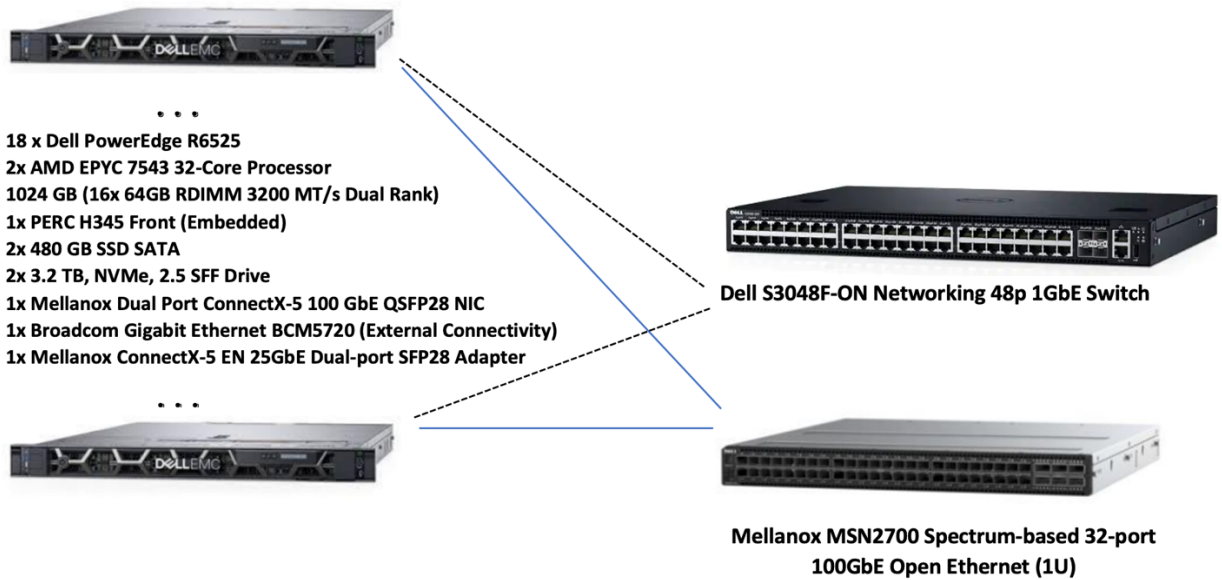


Figure 1.1: Benchmarked and priced system configuration

# 1 Clause 1: Logical Database Design Related Items

## 1.1 Database Definition Statements

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

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

## 1.2 Physical Organization

*The physical organization of tables and indices, within the test and qualification databases, must be disclosed. If the column ordering of any table is different from that specified in Clause 1.4, it must be noted.*

Physical organization requires no user input. All the database data is placed on the same partition.

## 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. The data is automatically distributed on the cluster nodes using a hash algorithm. The columns used for the hashing are controlled by DDL statements (see Supporting Files Archive).

## 1.4 Replication

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

No replication was used.

## 2 Clause 2: Queries and Refresh Functions

### 2.1 Query Language

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

SQL was the query language uniquely used throughout this benchmark.

### 2.2 Verifying Method for Random Number Generation

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

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

### 2.3 Generating Values for Substitution Parameters

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

QGEN version 3.0.0 was used to generate the substitution parameters.

### 2.4 Query Text and Output Data from Qualification Database

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

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

- Quoting of reserved keyword „value“ (Q11)
- LIMIT syntax used to restrict the number of output rows (Q2,Q3,Q10,Q18,Q21)
- Naming of columns of sub-select in Q13
- Used approved variant A of Q15 (Appendix B): ‘with clause’ instead of “create view/drop view”

### 2.5 Query Substitution Parameters and Seeds Used

*The query substitution parameters used for all performance tests must be disclosed in tabular format, along with the seeds used to generate these parameters.*

The Supporting Files Archive contains the seed and query substitution parameters.

### 2.6 Isolation Level

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

The queries and transactions were run with the isolation level 3.

### 2.7 Source Code of Refresh Functions

*The details of how the refresh functions were implemented must be disclosed (including source code of any non-commercial program used).*

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

## 3 Clause 3: Database System Properties

### 3.1 ACID Properties

*The ACID (Atomicity, Consistency, Isolation, and Durability) properties of transaction processing systems must be supported by the system under test during the timed portion of this benchmark. Since TPC-H is not a transaction processing benchmark, the ACID properties must be evaluated outside the timed portion of the test.*

All ACID tests were conducted according to specification. The Supporting Files Archive contains the source code of the ACID test scripts.

### 3.2 Atomicity Requirements

*The system under test must guarantee that transactions are atomic; the system will either perform all individual operations on the data, or will assure that no partially completed operations leave any effects on the data.*

#### 3.2.1 Atomicity of the Completed Transactions

*Perform the ACID Transaction for a randomly selected set of input data and verify that the appropriate rows have been changed in the ORDERS, LINEITEM, and HISTORY tables.*

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

1. The total price from the ORDERS table and the extended price from the LINEITEM table were retrieved for a randomly selected order key.
2. One ACID Transaction was performed using the order key from step 1.
3. The ACID Transaction was committed.
4. The total price from the ORDERS table and the extended price from the LINEITEM table were retrieved for the same order key.
5. It was verified that the appropriate rows had been changed.

#### 3.2.2 Atomicity of Aborted Transactions

*Perform the ACID Transaction for a randomly selected set of input data, substituting a ROLLBACK of the transaction for the COMMIT of the transaction. Verify that the appropriate rows have not been changed in the ORDERS, LINEITEM, and HISTORY tables.*

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

1. The total price from the ORDERS table and the extended price from the LINEITEM table were retrieved for a randomly selected order key.
2. One ACID Transaction was performed using the order key from step 1. The transaction was stopped prior to the commit.
3. The ACID Transaction was rolled back.
4. The total price from the ORDERS table and the extended price from the LINEITEM table were retrieved for the same order key.
5. It was verified that the appropriate rows had not been changed.

### 3.3 Consistency Requirements

*Consistency is the property of the application that requires any execution of transactions to take the database from one consistent state to another. A consistent state for the TPC-H database is defined to exist when:*

$$O\_TOTALPRICE = SUM(trunc(trunc(L\_EXTENDEDPRICE*(1-L\_DISCOUNT),2)*(1+L\_TAX),2))$$

*For each ORDER and LINEITEM defined by (O\_ORDERKEY = L\_ORDERKEY).*

#### 3.3.1 Consistency Test

*Verify that ORDERS and LINEITEM tables are initially consistent, submit the prescribed number of ACID Transactions with randomly selected input parameters, and re-verify the consistency of the ORDERS and LINEITEM.*

The following queries were executed before and after the durability tests to show that the database was always in a consistent state both initially and after submitting transactions:

```

SELECT *
FROM (
  SELECT o_orderkey, o_totalprice - sum(trunc(trunc(l_extendedprice *
    (1-l_discount),2)*(1+l_tax),2)) part_res
  FROM orders, lineitem
  WHERE o_orderkey=l_orderkey
  GROUP BY o_orderkey, o_totalprice
) WHERE not part_res=0;

```

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

1. The consistency of the ORDERS and LINEITEM tables was verified.
2. 100 transactions for each of the 16 execution streams were prepared.
3. The 100 ACID transactions per stream were executed from 16 execution streams.
4. The consistency of the ORDERS and LINEITEM tables was re-verified.

### 3.4 Isolation Requirements

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

The steps of the isolation tests were adapted to the Exasol isolation environment.

#### 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. Start a query and verify that the row was retrieved.
2. Start an update transaction, read and update the same row. Wait before commit.
3. Start the same query and verify that the row retrieved has not changed.
4. Commit the update transaction
5. Start the same query and verify that the new row is retrieved

#### 3.4.2 Isolation Test 2 – Read-Write Conflict with Rollback

*Demonstrate isolation for the read-write conflict of a read-write transaction and a read-only transaction when the read-write transaction is rolled back.*

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

1. Start a query and verify that the row was retrieved.
2. Start an update transaction, read and update the same row. Wait before commit.
3. Start the same query and verify that the row retrieved has not changed.
4. Rollback the update transaction
5. Start the same query and verify that the old row (step 1) is retrieved

#### 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. Start a query and verify that the row was retrieved.
2. Start an update transaction, read and update the same row. Wait before commit.
3. Start another update transaction, read and try to update the same row and verify that the transaction is forced to rollback.
4. commit the update transaction
5. Start the same query and verify that the new row is retrieved

#### 3.4.4 Isolation Test 4 – Write-Write Conflict with Rollback

*Demonstrate isolation for the write-write conflict of two update transactions when the first transaction is rolled back.*



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

1. Start a query and verify that the row was retrieved.
2. Start an update transaction, read and update the same row. Wait before commit.
3. Start another update transaction, read and try to update the same row and verify that the transaction is forced to rollback.
4. Rollback the update transaction
5. Start the same query and verify that the old row (step 1) is retrieved.

### 3.4.5 Isolation Test 5 – Concurrent Read and Write Transactions on Different Tables

*Demonstrate the ability of read and write transactions affecting different database tables to make progress concurrently.*

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

1. Start a query and verify that the row was retrieved.
2. Start an update transaction, read and update the same row. Wait before commit.
3. Start a second transaction that does the following:  
Select random values of PS\_PARTKEY and PS\_SUPPKEY. Return all columns of the PARTSUPP table for which PS\_PARTKEY and PS\_SUPPKEY are equal to the selected values.
4. Verify that the read transaction completes.
5. Commit the update transaction.
6. Start the same query and verify that the new row is retrieved.

### 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 query was used to ensure sufficient execution time to perform the test:

```
SELECT l1.l_quantity,
       SUM(l2.l_extendedprice),
       SUM(l3.l_extendedprice),
       SUM(l3.l_quantity)
FROM   lineitem l1, lineitem l2, lineitem l3, lineitem l4, lineitem l5
WHERE  l1.l_shipdate <= DATE '1998-12-01' -0
      AND l1.l_orderkey = l2.l_orderkey
      AND l1.l_linenumber = l2.l_linenumber
      AND l1.l_extendedprice = l3.l_extendedprice
      AND l3.l_quantity < 30
      AND l4.l_quantity = l1.l_quantity
      AND l4.l_orderkey < 150
      AND l5.l_receiptdate = l1.l_receiptdate
      AND l5.l_partkey < 140
GROUP BY l1.l_quantity;
COMMIT;
```

1. A Transaction, T1, which executed the above query against the qualification database, was started using a randomly selected DELTA.
2. An ACID Transaction, T2, was started for a randomly selected O\_KEY, L\_KEY and DELTA.
3. T2 completed and appropriate rows in the ORDERS, LINEITEM and HISTORY tables had been changed.
4. T1 was still executing.
5. Transaction T1 completed executing the query.

## 3.5 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 3.5.3.*

Exasol has serializable isolation level with table level lock concurrency control. The ACID Transaction of stream0 was expanded with 5 seconds delay after the update and before commit after it committed 100 transactions. Since only one update transaction can execute at any one time, the delay should guarantee that the active update transaction is "in-flight" at the time of the failure.

The following steps were performed for the durability test:

1. The consistency of the ORDERS and LINEITEM tables was verified.
2. 1000 transactions for each of the 16 execution streams were prepared.
3. After that at least 100 ACID transactions were submitted from each of the 16 execution streams.
4. A durability failure was induced (see details for each failure below).
5. After restoration of the system the consistency of the ORDERS and LINEITEM tables was re-verified.
6. The durability success files and the HISTORY table were compared.

All durability tests were performed on an 4-node cluster.

### **3.5.1 Permanent Unrecoverable Failure of Any Durable Medium**

*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.*

Disk, node and controller failure tests were combined as explained in section 3.5.5.

### **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.*

The system crash, memory failure and loss of external power tests were combined as explained in section 3.5.4.

### **3.5.3 Memory Failure**

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

The system crash, memory failure and loss of external power tests were combined as explained in section 3.5.4.

### **3.5.4 Loss of External Power**

*Loss of External Power: Guarantee the database and the effects of committed updates are preserved during the loss of all external power to the SUT for an indefinite time period.*

System crash was performed by turning off the power of 4 node cluster during the durability test. When power was restored, the system rebooted automatically and the database was restarted manually.

In addition single node crash was performed like described in section 3.5.5

### **3.5.5 Node or Controller Failure**

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

Cluster of 4 online nodes and 1 offline spare node were used for this test.

This test was executed in following order:

1. One of the online nodes was crashed by removing the power – the system stopped to work.
2. The offline spare node was turned on and manually configured as replacement of the failed node.
3. On the remaining 3 + 1 replacement nodes the database was restarted manually.

## 4 Clause 4: Scaling and Database Population

### 4.1 Ending Cardinality of Tables

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

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

Table	Rows
Lineitem	179,999,978,268
Order	45,000,000,000
Partsupp	24,000,000,000
Part	6,000,000,000
Customer	4,500,000,000
Supplier	300,000,000
Nation	25
Region	5

### 4.2 Distribution of Tables and Logs Across Media

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

Each server has 2x 480 GB SSD drives configured as 1x mirrored pair (HW RAID 1) used for OS installation. The resulting 1x 480 GB device (sda) is divided into 2 partitions as shown in the table below. In addition each server contains 2x 3200GB SSD drives configured as raw devices (nvme0n1, nvme1n1) used for data storage. All benchmark- and database-relevant data is stored on the raw devices.\*

Partition Name	Type	Partition Size** / Devices	Content
sda1	vfat	0.512 GiB of 447 GiB	/boot/efi
sda2	ext4	446.1 GiB of 447 GiB	OS installation
nvme0n1	raw	2,980 GiB of 2,980 GiB	DBMS data
nvme1n1	raw	2,980 GiB of 2,980 GiB	DBMS data

\* Disk manufacturer definition of 1 GB is  $10^9$  bytes

\*\* In this calculation 1 GiB is defined as  $2^{30}$  bytes

### 4.3 Mapping of Database Partitions/Replication

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

Horizontal partitioning is used. The data is automatically distributed on the cluster nodes using a hash algorithm. The columns used for the hashing are controlled by DDL statements (see Supporting Files Archive).

The data is mirrored across the cluster nodes to achieve redundancy for the purpose of recovery only.

### 4.4 Implementation of RAID

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

Please refer to chapter 4.2.

#### 4.5 DBGEN Modifications

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

The supplied DBGEN version 3.0.0 was used, no modifications were made.

#### 4.6 Database Load Time

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

See Numerical Quantities Summary in the Executive Summary.

#### 4.7 Data Storage Ratio

The data storage ratio must be disclosed. It is computed by dividing the total data storage of the priced configuration (expressed in GB) by the size chosen for the test database as defined in 4.1.3.1. The ratio must be reported to the nearest 1/100th, rounded up. For the reporting of configured disk capacity, gigabyte (GB) is defined to be  $2^{30}$  bytes.

Disk Type	GB per disk*	GiB per disk**	# of disks	Total (GiB)**
Internal	480 GB	447 GiB	2 * 18	16,093
Internal	3,200 GB	2,980 GiB	2 * 18	107,288
<b>Total Storage</b>				123,382
<b>Data Storage Ratio</b>				<b>4.11</b>

\* Disk manufacturer definition of 1 GB is  $10^9$  bytes


\*\* In this calculation 1 GiB is defined as  $2^{30}$  bytes

#### 4.8 Database Load Mechanism Details and Illustration

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

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

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



#### 4.9 Qualification Database Configuration

Any differences between the configuration of the qualification database and the test database must be disclosed.

The qualification database used identical scripts to create and load the data with changes to adjust for the database scale factor.

#### 4.10 Memory to Database Size Percentage

*The memory to database size percentage must be disclosed. It is computed by multiplying by 100 the total memory size priced on the SUT (see clause 6.2.1) and dividing this number by the size chosen for the test database as defined in Clause 4.1.3.1.*

<b>Nodes</b>	<b>RAM per Node</b>	<b>Total Memory</b>	<b>Scale Factor</b>	<b>Memory to Database Size Ratio</b>
18	1024 GB	18432 GB	30,000	61.4%

## **5 Clause 5: Performance Metrics and Execution Rules Related Items**

### **5.1 System Activity between Load and Performance Tests**

*Any system activity on the SUT that takes place between the conclusion of the load test and the beginning of the performance test must be fully disclosed.*

There is no activity on the SUT between the conclusion of the load test and the beginning of the performance test.

### **5.2 Steps in the Power Test**

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

The following steps were used to implement the power test:

1. RF1 refresh function from update stream
2. Stream 0 execution from query stream
3. RF2 refresh function from same update stream

### **5.3 Timing Interval for Each Query and Refresh Functions**

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

See Numerical Quantities Summary in the Executive Summary.

### **5.4 Number of Streams for the Throughput Test**

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

One stream was used for the refresh pairs. The number of query streams used is listed in the Numerical Quantities Summary in the Executive Summary.

### **5.5 Start and End Date/Time of Each Query Stream**

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

See Numerical Quantities Summary in the Executive Summary.

### **5.6 Total Elapsed Time of the Measurement Interval**

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

See Numerical Quantities Summary in the Executive Summary.

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

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

See Numerical Quantities Summary in the Executive Summary.

### **5.8 Performance Metrics**

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

See Numerical Quantities Summary in the Executive Summary.

### **5.9 The Performance Metric and Numerical Quantities from Both Runs**

*A description of the method used to determine the reproducibility of the measurement results must be reported. This must include the performance metrics ( $Q_{ppH}$  and  $Q_{thH}$ ) from reproducibility runs.*

	Run 1	Run 2
<b>QphH@30000GB</b>	22,664,825.1	22,794,751.9
<b>QppH@30000GB</b>	24,086,987.1	25,405,683.7
<b>QthH@30000GB</b>	21,326,631.4	20,452,144.5

### 5.10 System Activity between Performance Tests

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

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

### 5.11 Documentation to satisfy Clause 5.2.7

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

Exasol documentation is publicly available at [Exasol.com](https://www.exasol.com)

### 5.12 Query Output Validation

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

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

## 6 Clause 6: SUT and Driver Implementation Related Items

### 6.1 Driver

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

All stream executions are performed by a script. QGEN is used to produce query text.

For each power-test run:

1. A shell script is started, executes RF1 and then waits for the query stream to complete.
2. A shell script is started, executes the 22 queries in the required order for stream 0 and then signals to the shell script started in step 1.
3. The shell script started in step 1 is released and executes RF2.

For each throughput-test run:

1. The queries as generated by QGEN are submitted in the order defined by Clause 5.3.5.4 from the driver in several streams (the number of streams is listed in the Numerical Quantities).
2. In parallel with the queries, pairs of RF1/RF2 are executed sequentially in one update stream.

The source code of the used scripts are disclosed in the Supporting Files 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.*

The scripts used to implement the ISL are disclosed in the Supporting Files Archive.

### 6.3 Profile-Directed Optimization

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

Profile-directed optimization was not used.



## 7 Clause 7: Pricing

### 7.1 Hardware and Software Used in the Priced System

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

A detailed list of hardware and software used in the priced system is included in the pricing sheet in the executive summary. All prices are currently effective. Third-party price quotations are included in Appendix A.

### 7.2 Total Three Year Price

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

A detailed pricing sheet of all the hardware and software used in this configuration and the 3-year maintenance costs, demonstrating the computation of the total 3-year price of the configuration, is included in the executive summary at the beginning of this document.

Exasol uses raw data subscription pricing; for details see Appendix A.

### 7.3 Availability Date

*The committed delivery date for general availability of products used in the priced calculations must be reported. When the priced system includes products with different availability dates, the availability date reported on the executive summary must be the date by which all components are committed to being available. The full disclosure report must report availability dates individually for at least each of the categories for which a pricing subtotal must be provided.*

Component	Availability Date
Cluster Hardware	Now (date of publication)
Exasol 7.1	July 1, 2021

## 8 Clause 8: Full Disclosure

### 8.1 Supporting Files Index Table

An index for all files and/or directories included in the Supporting Files Archive as required by Clauses 8.3.2 through 8.3.8 must be provided in the report.

Clause	Description	Archive Files	Pathname
1	Parameter Settings	benchmark_scripts.zip	RUN/params.log
	OS Tunable Parameters		RUN/etc.zip
	DB Creation Scripts		KIT/sql/create_user.sql KIT/sql/create_schema.sql KIT/sql/create_indices.sql KIT/sql/analyse_database.sql
	System Verification		KIT/scripts/tools/hwinfo.sh
	Toolkit Common Scripts		KIT/scripts
2	Minor query modifications	benchmark_scripts.zip	KIT/tpch_archives/TPC-H_Tools_v3.0.0.zip.patch
3	ACID Test Scripts	benchmark_scripts.zip	KIT/ACID/
	ACID Test Results		ACID/
4	Database Load Scripts	benchmark_scripts.zip	KIT/scripts/load_init.sh
	Qualification Test Results		VLD/
5	Query Output Results	run1results.zip	RUN//run1/
		run2results.zip	RUN//run2/
6	Source Codes and Scripts of Driver	benchmark_scripts.zip	KIT/scripts/query_streams
7	There are no files to be included for Clause 7.	N/A	N/A
8	Query Parameters & Seeds	run1results.zip	RUN/run1/substitution_parameters.txt
	Executable Query Text		RUN/run1/stream*.sql
	RF function source code		KIT/scripts/tpc_h_run_full.sh

## 9 Clause 9: Audit Related Items

### 9.1 Auditor's Report

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

This implementation of the TPC Benchmark H was audited by Doug Johnson of InfoSizing, a certified TPC-H auditor. Further information regarding the audit process may be obtained from:

**Doug Johnson**

InfoSizing

20 Kreg Ln.

Manitou Springs, CO 80829

Phone: (719) 473-7555

Email: [doug@sizing.com](mailto:doug@sizing.com)

TPC Benchmark H Full Disclosure Report and other information can be downloaded from the Transaction Processing Performance Council website at [www.tpc.org](http://www.tpc.org)

Benchmark sponsor: Mr. Nicholas Wakou  
 Dell Inc.  
 701 E. Parmer Ln. Bldg. 2  
 Austin, TX 78753

May 27, 2021

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

Platform: Dell PowerEdge R6525 (18-node cluster)  
 Operating System: Ubuntu 20.04.2 LTS  
 Database Manager: Exasol 7.1  
 Other Software: n/a

The results were:

**Performance Metric 22,664,825.1 QpH@30,000GB**  
 TPC-H Power 24,086,987.1  
 TPC-H Throughput 21,326,631.4  
 Database Load Time 00d 02h 48m 59s

**Server 18x Dell PowerEdge R6525, each with:**

CPU	2x AMD EPYC 7543 32-Core Processor (64 threads)		
Memory	1,024 GiB		
Disks	<b>Qty</b>	<b>Size</b>	<b>Type</b>
	2	480 GB	SSD SATA
	2	3.2 TB	NVMe

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

The following verification items were given special attention:

- The database records were defined with the proper layout and size
- The database population was generated using DBGen
- The database was properly scaled to 30,000GB 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 approved query variant A of Q15
- 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 11 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 cursive script that reads "Doug Johnson". The signature is written in black ink and has a long, sweeping horizontal line extending to the right.

Doug Johnson, TPC Certified Auditor

63 Lourdes Dr. | Leominster, MA 01453 | 978-343-6562 | [www.sizing.com](http://www.sizing.com)

## Appendix A: Pricing Information

Exasol Inc  
1372 Peachtree Street  
Atlanta, GA, 30309  
United States



**Dell, Inc.**  
One Dell Way  
Round Rock, TX 78682  
United States

Product	Part Number	Unit Price \$/Mo	Qty	Ext. Price \$
Exasol 7.1	Exasol 30TB Subscription <sup>1)</sup>	\$25,000	36	\$900,000
<b>Total</b>				<b>\$900,000</b>

<sup>1)</sup> Includes software support 24x7

This quote is valid until December 15, 2021.

**Exasol Pricing Contact:**  
Gunther Schweer  
[sales@Exasol.com](mailto:sales@Exasol.com)