



TPC Benchmark™ H Full Disclosure Report

Dell PowerEdge R720xd *using* **EXASolution 5.0**

First Edition
September 22, 2014

*Dell PowerEdge R720xd
using
EXASolution 5.0*

Dell 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. We assume 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 © 2014 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.

PowerEdge R720xd is a trademark of Dell.

EXASOL, EXASolution and EXACluster OS are trademarks of EXASOL AG.

Intel and Xeon are trademarks of Intel 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 R720xd using EXASolution 5.0

TPC-H Rev. 2.17.0
TPC-Pricing Rev. 1.7.0

Report Date
September 22, 2014

Total System Cost

\$2,566,209

Composite Query per Hour Metric

11,223,614.4
QphH@30000GB

Price / Performance

\$0.23
\$ / QphH@30000GB

Database Size

30,000GB

Database Manager

EXASolution 5.0

Operating System

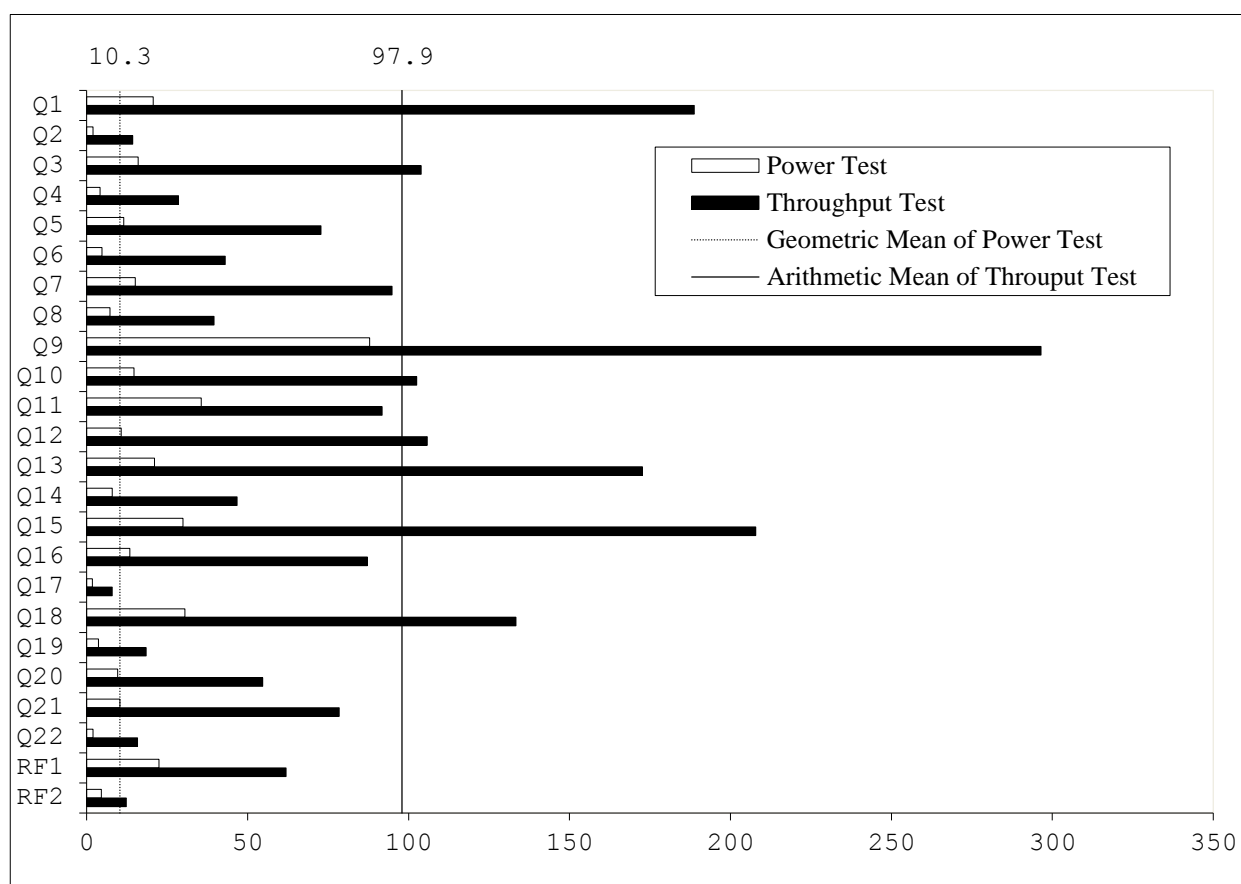
EXAClusterOS 5.0

Other
Software

None

Availability Date

September 22, 2014



Database Load Time = 4:56:36

Load Includes Backup: N

Total Data Storage / Database Size = 6.32

Memory to Database Size Percentage = 42.7%

Storage Redundancy Levels:

Base Tables Level One

Auxiliary Data Structures Level One

DBMS Temporary Space Level One

OS and DBMS Software Level One

System Configuration:

40 x Dell PowerEdge R720xd Server, each with:

320 GB RAM

2 x Intel Xeon E5-2680v2 10C 2.8 GHz processors (each 10 cores, 20 threads)

2 x 146 GB (15k rpm) internal SAS disks

4 x 1199 GB (10k rpm) internal SAS disks

Total Storage: 189,543 GB



Dell PowerEdge R720xd using EXASolution 5.0

TPC-H Rev. 2.17.0
TPC-Pricing Rev. 1.7.0

Report Date
September 22, 2014

Description	Key	Part Number	Unit Price \$	Qty	Ext. Price \$	3 Yr. Maint.
Hardware						
Server Hardware	1		24,102.90	40	964,116.00	included
OEM R720XD Chassis up to 24 HDDs+2 Flex bays	1	210-ABNG		1		
Intel Xeon E5-2680v2, 2.8Hz, 10C	1	338-BCUH		2		
16GB RDIMM, 1,866MT/s	1	370-AAWL		16		
8GB RDIMM, 1,866MT/s	1	370-AAUK		8		
1.2TB SAS 6Gbit/s, 2.5", 10k, Hot Plug HDD	1	400-AAHL		4		
146GB, SAS 6Gbit/s, 2.5", 15k Hit Plug HDD (Flex Bay)	1	408-10128		2		
PERC H710p integrated RAID-Controller 1GB Cache	1	405-12071		1		
Hot-Plug redundant Power Supply (1+1) 1100W	1	450-17883		1		
Intel Ethernet i350 QP 1Gbit/s Network Daughtercard	1	540-11050		1		
Intel X520 DP 10Gbit/s DA/SFP+ Serveradapter	1	540-11063		2		
iDRAC7 Enterprise including 8GB SD Card	1	528-10006		1		
Sliding Ready Rack Rails with cable management arm	1	770-11607		1		
39M ProSupport for OEM and Next Business Day On-Site Service	1	710-33226		1		
Exasol Bezel Branding	1	CFI BGMK01		1		
DELL DISCOUNT 30%	1				-289,234.80	
S4810P, 48x10GbE SFP+ Ports, 4xQSFP 40GbE Ports	1	210-38654	32,806.62	2	65,613.24	included
Dell Networking N2048, L2, 48x 1GbE + 2x10GbE SFO	1	210-ABNX	4,710.64	2	9,421.28	included
Rack TFT + Keyboard	1	A7485906	1,621.55	3	4,864.65	included
DELL DISCOUNT 30%	1				-23,969.75	
Dell Networking, Cable, SFP+ to SFP+, 10GbE, 3 meter	1	470-13555	186.48	80	14,918.40	
C2G Cat6 550MHz Snagless Patch Cable, 3 meter	1	A6928965	4.11	80	328.80	
DELL DISCOUNT 10%	1				-1,524.72	
APC Smart UPS RT 8000VA RM incl. Shutdown Software	1	A1176411	5,337.64	3	16,012.92	included
APC Smart UPS RT 192V 5kVA and 6kVA RM Battery Pack	1	A7773129	1,568.27	3	4,704.81	included
APC Basic Rack PDU Zero U, 120/208/230 V	1	A7067488	429.89	8	3,439.12	included
Dell NetShelter SX Deep Enclosure 42U	1	A7545497	1,486.42	3	4,459.26	included
APC IEC 320 EN 60320 C19 (W) - IEC 309 (M) - 2.4 m	1	A7155860	166.84	4	667.36	included
POWER CORD IEC 320 C19 TO IEC 320 C20 NS	1	A7294276	97.92	4	391.68	included
Reference offer Dell: 29862105	1					
Hardware Subtotal					774,208.25	
Software						
EXASolution 5.0 (incl. EXACluster OS 5.0)	2	10-15TB	1,600,000.00	1	1,600,000.00	192,000.00
Software Subtotal					1,600,000.00	192,000.00
Total					\$2,374,208.25	\$192,000.00

Three-Year Cost of Ownership: **\$2,566,209**
QphH@30000GB: **11,223,614.4**
\$/QphH@30000GB: **\$0.23**

Price Key: 1 DELL Emanuel Lippmann emanuel_lippmann@dell.com
2 EXASOL Harald Kalt harald_kalt@dell.com
Aaron Auld sales@exasol.com

Audited by Francois Raab of InfoSizing, Inc. (www.sizing.com)

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



Dell PowerEdge R720xd using EXASolution 5.0

TPC-H Rev. 2.17.0
TPC-Pricing Rev. 1.7.0

Report Date
September 22, 2014

Numerical Quantities

Measurement Results

Database Scale Factor	30,000 GB
Total Data Storage / Database Size	6.32
Percentage Memory / Database Size	42.7%
Start of Database Load	2014-08-27 07:41:23
End of Database Load	2014-08-27 12:37:59
Database Load Time	4:56:36
Query Streams for Throughput Test	11
TPC-H Power	10,478,181.0
TPC-H Throughput	12,022,079.1
TPC-H Composite Query-per-Hour Metric (QphH@30000GB)	11,223,614.4
Total System Price Over 3 Years	\$2,566,209
TPC-H Price/ Performance Metric (\$/QphH@30000GB)	\$0.23

Measurement Interval

Measurement Interval in Throughput Test (Ts)	2,174 seconds
--	---------------

Duration of Stream Execution

Power Run	Seed	Query Start Time	Duration (sec)	RF1 Start Time	RF2 Start Time
		Query End Time		RF1 End Time	RF2 End Time
	827123759	2014-08-27 13:21:07	00 00:06:01	2014-08-27 13:20:44	2014-08-27 13:27:08
		2014-08-27 13:27:08	361	2014-08-27 13:21:07	2014-08-27 13:27:13

Throughput Stream	Seed	Query Start Time	Duration (sec)	RF1 Start Time	RF2 Start Time
		Query End Time		RF1 End Time	RF2 End Time
1	827123760	2014-08-27 13:27:14	00 00:35:33	2014-08-27 13:27:12	2014-08-27 13:28:26
		2014-08-27 14:02:47	2,133	2014-08-27 13:28:26	2014-08-27 13:28:35
2	827123761	2014-08-27 13:27:13	00 00:36:13	2014-08-27 13:28:35	2014-08-27 13:29:36
		2014-08-27 14:03:26	2,173	2014-08-27 13:29:36	2014-08-27 13:29:48
3	827123762	2014-08-27 13:27:16	00 00:36:03	2014-08-27 13:29:48	2014-08-27 13:30:39
		2014-08-27 14:03:19	2,163	2014-08-27 13:30:39	2014-08-27 13:30:48
4	827123763	2014-08-27 13:27:15	00 00:36:01	2014-08-27 13:30:48	2014-08-27 13:31:48
		2014-08-27 14:03:16	2,161	2014-08-27 13:31:48	2014-08-27 13:31:58
5	827123764	2014-08-27 13:27:12	00 00:36:02	2014-08-27 13:31:58	2014-08-27 13:32:45
		2014-08-27 14:03:14	2,162	2014-08-27 13:32:45	2014-08-27 13:32:54
6	827123765	2014-08-27 13:27:18	00 00:36:01	2014-08-27 13:32:54	2014-08-27 13:33:52
		2014-08-27 14:03:19	2,161	2014-08-27 13:33:52	2014-08-27 13:34:01
7	827123766	2014-08-27 13:27:24	00 00:36:02	2014-08-27 13:34:01	2014-08-27 13:34:51
		2014-08-27 14:03:26	2,162	2014-08-27 13:34:51	2014-08-27 13:35:07
8	827123767	2014-08-27 13:27:20	00 00:35:59	2014-08-27 13:35:07	2014-08-27 13:36:40
		2014-08-27 14:03:19	2,159	2014-08-27 13:36:39	2014-08-27 13:36:55
9	827123768	2014-08-27 13:27:22	00 00:35:40	2014-08-27 13:36:56	2014-08-27 13:38:11
		2014-08-27 14:03:02	2,140	2014-08-27 13:38:11	2014-08-27 13:38:23
10	827123769	2014-08-27 13:27:26	00 00:36:00	2014-08-27 13:38:23	2014-08-27 13:39:24
		2014-08-27 14:03:26	2,160	2014-08-27 13:39:24	2014-08-27 13:39:44
11	827123770	2014-08-27 13:27:28	00 00:35:39	2014-08-27 13:39:44	2014-08-27 13:41:14
		2014-08-27 14:03:07	2,139	2014-08-27 13:41:14	2014-08-27 13:41:36



Dell PowerEdge R720xd using EXASolution 5.0

TPC-H Rev. 2.17.0
TPC-Pricing Rev. 1.7.0

Report Date
September 22, 2014

TPC-H Timing Intervals (in seconds)

Stream ID	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11	Q12
0	20.7	2.0	16.0	4.1	11.5	4.7	15.1	7.2	87.9	14.7	35.6	10.7
1	180.9	14.5	112.3	32.7	69.5	45.9	106.8	42.8	342.3	103.0	86.4	103.5
2	221.2	13.5	118.6	38.3	90.5	51.9	107.8	50.1	266.5	100.0	91.9	108.2
3	203.2	16.9	106.3	36.2	79.1	43.9	81.3	44.9	334.2	103.5	95.3	105.0
4	218.8	14.4	107.4	33.5	72.0	46.1	104.6	43.8	317.2	100.2	95.1	128.5
5	184.9	14.2	120.0	28.3	85.6	37.4	88.6	46.4	295.7	126.0	111.9	127.9
6	221.4	17.1	103.7	30.2	91.3	51.2	95.6	36.0	327.8	116.9	106.7	119.5
7	216.7	14.4	117.4	24.9	81.6	55.0	107.7	35.8	330.6	115.3	92.5	78.1
8	196.1	14.9	109.1	29.7	84.9	49.1	105.2	36.3	292.8	105.0	89.7	131.8
9	178.6	16.4	98.8	28.5	69.4	34.4	106.1	51.3	346.7	124.5	98.8	125.3
10	205.4	13.4	129.3	26.9	55.4	61.4	116.9	40.2	308.2	103.1	94.6	128.1
11	217.5	19.3	108.6	28.6	82.2	35.4	102.1	39.9	308.6	118.1	102.8	102.8
Min	20.7	2.0	16.0	4.1	11.5	4.7	15.1	7.2	87.9	14.7	35.6	10.7
Avg	188.8	14.3	104.0	28.5	72.8	43.0	94.8	39.6	296.5	102.5	91.8	105.8
Max	221.4	19.3	129.3	38.3	91.3	61.4	116.9	51.3	346.7	126.0	111.9	131.8

Stream ID	Q13	Q14	Q15	Q16	Q17	Q18	Q19	Q20	Q21	Q22	RF1	RF2
0	21.1	7.9	29.9	13.4	1.8	30.5	3.6	9.6	10.3	2	22.4	4.5
1	187.1	55.2	234.6	109.4	11.6	105.2	27.9	62.2	79.4	18.8	73	9.7
2	176.6	61.7	234	79.7	15.9	179.4	13.8	65.8	67.2	20.2	60.7	12
3	208.6	47	213.9	95.3	6.6	150.8	17.3	63.4	92.9	16.2	51.3	8.2
4	207.7	46	207.2	83.7	5.8	167	11.1	50.3	86	13.6	60.8	9.7
5	185.6	51.5	237.3	94.5	9.9	135.4	18.9	63.5	76	21.8	47.4	8.9
6	180.8	42.6	204.1	98.6	7.4	150.5	14.8	52.3	78	13.9	57.9	8.8
7	173.7	58.9	218.4	108.8	6.6	151.5	17.9	55.2	85.5	14.7	49.7	16.2
8	192.3	48.3	215.8	90.2	12.4	147.3	41.5	58	90.2	17.4	92.4	16
9	162.6	48	230.3	86.7	6.2	141.6	21.4	58.6	90	15	75.7	11.6
10	185.8	44.5	230.2	104.4	5.6	120.5	17.6	62.2	85.4	20	61.1	20.2
11	189.9	49.1	238.5	82	5	120.3	15.4	55.4	101.1	15.3	90.3	21.6
Min	21.1	7.9	29.9	13.4	1.8	30.5	3.6	9.6	10.3	2.0	22.4	4.5
Avg	172.7	46.7	207.9	87.2	7.9	133.3	18.4	54.7	78.5	15.7	61.9	12.3
Max	208.6	61.7	238.5	109.4	15.9	179.4	41.5	65.8	101.1	21.8	92.4	21.6

Table of Contents

TPC BENCHMARK H OVERVIEW	9
0 GENERAL ITEMS.....	10
0.1 BENCHMARK SPONSOR	10
0.2 PARAMETER SETTINGS.....	10
0.3 CONFIGURATION DIAGRAM	10
1 CLAUSE 1: LOGICAL DATABASE DESIGN RELATED ITEMS	12
1.1 DATABASE DEFINITION STATEMENTS.....	12
1.2 PHYSICAL ORGANIZATION	12
1.3 HORIZONTAL PARTITIONING.....	12
1.4 REPLICATION	12
2 CLAUSE 2: QUERIES AND REFRESH FUNCTIONS	13
2.1 QUERY LANGUAGE	13
2.2 VERIFYING METHOD FOR RANDOM NUMBER GENERATION.....	13
2.3 GENERATING VALUES FOR SUBSTITUTION PARAMETERS	13
2.4 QUERY TEXT AND OUTPUT DATA FROM QUALIFICATION DATABASE.....	13
2.5 QUERY SUBSTITUTION PARAMETERS AND SEEDS USED	13
2.6 ISOLATION LEVEL	13
2.7 SOURCE CODE OF REFRESH FUNCTIONS	13
3 CLAUSE 3: DATABASE SYSTEM PROPERTIES.....	14
3.1 ACID PROPERTIES.....	14
3.2 ATOMICITY REQUIREMENTS	14
3.3 CONSISTENCY REQUIREMENTS	14
3.4 ISOLATION REQUIREMENTS	15
3.5 DURABILITY REQUIREMENTS.....	16
4 CLAUSE 4: SCALING AND DATABASE POPULATION	18
4.1 ENDING CARDINALITY OF TABLES.....	18
4.2 DISTRIBUTION OF TABLES AND LOGS ACROSS MEDIA.....	18
4.3 MAPPING OF DATABASE PARTITIONS/REPLICATION	18
4.4 IMPLEMENTATION OF RAID	18
4.5 DBGEN MODIFICATIONS	18
4.6 DATA BASE LOAD TIME	19
4.7 DATA STORAGE RATIO	19
4.8 DATABASE LOAD MECHANISM DETAILS AND ILLUSTRATION.....	19
4.9 QUALIFICATION DATABASE CONFIGURATION	19
4.10 MEMORY TO DATABASE SIZE PERCENTAGE	20
5 CLAUSE 5: PERFORMANCE METRICS AND EXECUTION RULES RELATED ITEMS	21
5.1 SYSTEM ACTIVITY BETWEEN LOAD AND PERFORMANCE TESTS	21
5.2 STEPS IN THE POWER TEST.....	21
5.3 TIMING INTERVAL FOR EACH QUERY AND REFRESH FUNCTIONS	21
5.4 NUMBER OF STREAMS FOR THE THROUGHPUT TEST.....	21
5.5 START AND END DATE/TIME OF EACH QUERY STREAM	21
5.6 TOTAL ELAPSED TIME OF THE MEASUREMENT INTERVAL.....	21
5.7 REFRESH FUNCTION START DATE/TIME AND FINISH DATE/TIME	21
5.8 PERFORMANCE METRICS	21
5.9 THE PERFORMANCE METRIC AND NUMERICAL QUANTITIES FROM BOTH RUNS	21
5.10 SYSTEM ACTIVITY BETWEEN PERFORMANCE TESTS.....	22
5.11 DOCUMENTATION TO SATISFY CLAUSE 5.2.7.....	22
5.12 QUERY OUTPUT VALIDATION	22
6 CLAUSE 6: SUT AND DRIVER IMPLEMENTATION RELATED ITEMS	23
6.1 DRIVER	23
6.2 IMPLEMENTATION SPECIFIC LAYER (ISL).....	23
6.3 PROFILE-DIRECTED OPTIMIZATION	23
7 CLAUSE 7: PRICING.....	24
7.1 HARDWARE AND SOFTWARE USED IN THE PRICED SYSTEM	24
7.2 TOTAL THREE YEAR PRICE.....	24
7.3 AVAILABILITY DATE	24
8 CLAUSE 8: FULL DISCLOSURE	25
8.1 SUPPORTING FILES INDEX TABLE	25
9 CLAUSE 9: AUDIT RELATED ITEMS	26

9.1 AUDITOR'S REPORT.....	26
APPENDIX A: PRICING INFORMATION.....	29

TPC Benchmark H Overview

The TPC BenchmarkTM 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 parameters used in this benchmark. Default EXACluster OS system parameters were used.

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:

System components:

- 40 x Dell PowerEdge R720xd servers, each with:
 - 320 GB RAM
 - 2 x CPU (Intel® Xeon E5-2680v2, 2.8GHz)
 - 2 x 146 GB SAS 15.000 rpm
 - 4 x 1199 GB SAS 10.000 rpm
 - 1 x PERC H710p RAID controller 1GB NV cache
 - 1 x Ethernet Intel I350 QP 1Gb
 - 2 x Ethernet Intel X520 DP 10Gb
- 2 x Dell Force 10 S4810P Switch – 48 x 10GbE SFP+, 4 x QSFP 40GbE
- 2 x Dell Networking N2048 Switch – 48 x 1GbE Ports, 2 x 10GbE SPO

Network is 2 x 10Gb & 2 x 1Gb Ethernet. Each server contains 6 physical disks configured as 3 mirrored pairs at the controller level (HW RAID 1).

Priced configuration and benchmarked configuration are identical.




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 (8.1.2.1).

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 1.5.6.

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 2.17.0 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 2.17.0 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 12 execution streams were prepared.
3. The 100 ACID transactions per stream were executed from 12 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 3.5.2.

EXASolution 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. 400 transactions for each of the 12 execution streams were prepared.
3. After that at least 100 ACID transactions were submitted from each of the 12 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 6-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 failure test was performed as step in *Node or Controller Failure* test 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 and memory failure tests were combined as explained in section 3.5.3.

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.

System crash was performed by turning off the power of 6 node cluster during the durability test. When power was restored, the system rebooted automatically and the database was restarted manually.

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.

Each cluster node is configured with redundant power supplies. The first one is protected by UPS (one for the entire cluster) and the other is connected directly to the power grid.

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.

This test was executed in following order:

1. Disk of one node was removed (disk failure) during in-flight transaction – nothing happened, the system has continued to run.
2. Another node was crashed by removing the power (node failure) – the DBMS has stopped working.
3. Additional controller cache failure was injected:
 - a. All disk of that node were removed
 - b. The node was powered up
 - c. The controller cache was discarded using the raid configuration utility
 - d. The node was powered down
 - e. All disk were reinserted
4. The node was powered up and the system booted automatically
5. 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 contains 6 physical disks configured as 3 mirrored pairs at the controller level (HW RAID 1). The resulting 3 devices (1 x 146GB and 2 x 1199GB) are divided into 5 partitions as shown in the table below. All benchmark- and database-relevant data is stored in partition *d04_storage* (Tables, Indexes, Temp).

Partition Name	Type	Partition Size / Devices	Content
d00_os	ext4	50 GB of 146 GB	OS
d01_swap	swap	4 GB of 146 GB	Swap
d02_data	ext4	70 GB of 146 GB	Event logs for OS & DBMS
d03_data	ext4	1 x 1199 GB	DBGEN RF1 & RF2 flat files
d04_storage	raw	1 x 1199 GB	DBMS Data Devices

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 2.17.0 was used, no modifications were made.

4.6 Data Base 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.

Disk Type	GB per disk*	# of disks	Total (GB)**
Internal	146 GB	2 * 40	10,878
Internal	1199 GB	4 * 40	178,665
Total Storage			189,543
Data Storage Ratio			6.32

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


** In this calculation 1 GB 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.

Nodes	RAM per Node	Total Memory	Scale Factor	Memory to Database Size Ratio
40	320 GB	12,800 GB	30,000	42.7%

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 5.3.6) 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 5.3.5) 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
QphH@30000GB	11,382,610.2	11,223,614.4
QppH@30000GB	10,762,283.6	10,478,181.0
QthH@30000GB	12,038,691.8	12,022,079.1

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.

EXASolution documentation is publicly available at exasol.com

5.12 Query Output Validation

The output of the Query Output Validation Test must be 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.

EXASolution is licensed by the amount of main memory allocated to the database software (DBRAM size). This is independent of the physical RAM per node and the number of nodes. The database data doesn't need to fit into the licensed memory, although best performance can be reached in that case. Due to compression, this can be achieved with much less DBRAM size than raw data size.

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)
EXASolution 5.0	Now (date of publication)

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
	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/tpch_2_17_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 Francois Raab of InfoSizing, a certified TPC-H auditor. Further information regarding the audit process may be obtained from:

Francois Raab
InfoSizing, Inc.
531 Crystal Hills Blvd.
Manitou Springs, CO 80829
Phone: (719) 473-7555
Email: francois@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

Brian Bassett
Primary TPC Representative
Dell
Round Rock, Texas
United States

September 11, 2014

I verified the TPC Benchmark™ H v2.17.0 performance of the following configuration:

Platform: Dell PowerEdge R720xd, 40-node Cluster
Operating System: EXAClusterOS 5.0
Database Manager: EXASolution 5.0

The results were:

Performance Metric **11,223,614.4 QphH@30000GB**

TPC-H Power 10,478,181.0

TPC-H Throughput 12,022,079.1

Database Load Time 04:56:36

Server

Dell PowerEdge R720xd, 40-node Cluster (each node with)

CPUs 2 x Intel Xeon E5-2680v2 10C 2.8 GHz (25 MB L3, 10 cores, 20 threads)

Memory 320 GB

	Qty	Size	Type
Disks	2	146 GB	15K rpm SAS HDD
	4	1199 GB	10K rpm SAS HDD

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:

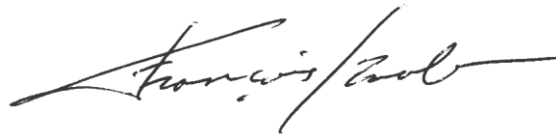
- The database records were defined with the proper layout and size
- The database population was generated using DBGen
- The database was properly scaled to 30,000GB and populated accordingly
- The compliance of the database auxiliary data structures was verified
- The database load time was correctly 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 one query variant
- The execution of the queries against the SF1 database produced compliant answers
- The implementation used Redundancy Level 1
- 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 the 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,



François Raab, President

Appendix A: Pricing Information

EXASOL Inc
625 2nd Street
San Francisco, CA 94107
United States



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

Product	Part Number	Unit Price \$	Qty	Ext. Price \$
EXASolution 5.0 (incl. EXACluster OS 5.0)	10-15TB	\$1,600,000	1	\$1,600,000
Support 24x7 (4% per year)	EXA-SUP	\$64,000	3	\$192,000
Total				\$1,792,000

* License size doesn't depend on the physical main memory, but the amount of memory allocated to EXASolution.

This quote is valid until December 15, 2014.

EXASOL Pricing Contact:

Aaron Auld
sales@exasol.com