



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

Composite Query per Hour Metric

Price / Performance

**\$4,222,459**

**11,612,395.5**

**\$0.37**

QphH@100000GB

\$/ QphH@100000GB

Database Size

Database Manager

Operating System

Other Software

Availability Date


**100,000GB**

**EXASolution 5.0**

**EXAClusterOS 5.0**

**None**

**September 22, 2014**



Database Load Time = 16:38:43  
Load Includes Backup: N  
Total Data Storage / Database Size = 7.95  
Memory to Database Size Percentage = 38.4%

### 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:

50 x Dell PowerEdge R720xd Server, each with:  
768 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  
14 x 1199 GB (10k rpm) internal SAS disks

**Total Storage:** 795,256 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.
<b>Hardware</b>						
Server Hardware	1		49,342.48	50	2,467,124.00	included
OEM R720XD Chassis up to 24 HDDs+2 Flex bays	1	210-ABNG		1		
Intel Xeon E5-2680v2, 2.8GHz, 10C	1	338-BCUH		2		
32GB LRDIMM , 1,866MT/s, Quad Rank	1	370-ABGL		24		
1.2TB SAS 6Gbit/s, 2.5", 10k, Hot Plug HDD	1	400-AAHL		14		
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 redundat 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				-740,137.20	
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	3	14,131.92	included
Rack TFT + Keyboard	1	A7485906	1,621.55	4	6,486.20	included
DELL DISCOUNT 30%	1				-25,869.41	
Dell Networking, Cable, SFP+ to SFP+, 10GbE, 3 meter	1	470-13555	186.48	100	18,648.00	
Dell Networking, Cable, 40GbE (QSFP+) to 4x10GbE SFP+	1	470-13548	748.61	2	1,497.22	
C2G Cat6 550MHz Snagless Patch Cable, 3 meter	1	A6928965	4.11	100	411.00	
DELL DISCOUNT 10%	1				-2,055.62	
APC Smart-UPS RT 12 kW 15000VA incl. Shutdown Software	1	A7067515	12,416.91	4	49,667.64	included
APC Basic Rack PDU Zero U, 120/208/230 V	1	A7067488	429.89	16	6,878.24	included
Dell NetShelter SX Deep Enclosure 42U	1	A7545497	1,486.42	4	5,945.68	included
APC IEC 320 EN 60320 C19 (W) - IEC 309 (M) – 2.4 m	1	A7155860	166.84	8	1,334.72	included
POWER CORD IEC 320 C19 TO IEC 320 C20 NS	1	A7294276	97.92	8	783.36	included
Reference offer Dell: 29862683	1					
<b>Hardware Subtotal</b>					<b>1,870,458.99</b>	
<b>Software</b>						
EXASolution 5.0 (incl. EXACluster OS 5.0)	2	35-40TB	2,100,000.00	1	2,100,000.00	252,000.00
<b>Software Subtotal</b>					<b>2,100,000.00</b>	<b>252,000.00</b>
<b>Total</b>					<b>\$3,970,458.99</b>	<b>\$252,000.00</b>

**Three-Year Cost of Ownership: \$4,222,459**  
**QphH@100000GB: 11,612,395.5**  
**\$/QphH@100000GB: \$0.37**

Price Key: 1 DELL Emanuel Lippmann [emanuel\\_lippmann@dell.com](mailto:emanuel_lippmann@dell.com)  
Harald Kalt [harald\\_kalt@dell.com](mailto:harald_kalt@dell.com)  
2 EXASOL Aaron Auld [sales@exasol.com](mailto:sales@exasol.com)

Audited by Francois Raab of InfoSizing, Inc. ([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 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	100,000 GB
Total Data Storage / Database Size	7.95
Percentage Memory / Database Size	38.4%
Start of Database Load	2014-08-27 23:43:04
End of Database Load	2014-08-28 16:21:47
Database Load Time	16:38:43
Query Streams for Throughput Test	11
TPC-H Power	11,754,288.9
TPC-H Throughput	11,472,214.9
TPC-H Composite Query-per-Hour Metric (QphH@100000GB)	11,612,395.5
Total System Price Over 3 Years	\$4,222,459
TPC-H Price/ Performance Metric (\$/QphH@100000GB)	\$0.37

### Measurement Interval

Measurement Interval in Throughput Test (Ts) 7,594 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
	828162147	2014-08-28 18:43:54	00 00:18:40	2014-08-28 18:43:00	2014-08-28 19:02:34
		2014-08-28 19:02:34		1,120	2014-08-28 18:43:54

Throughput Stream	Seed	Query Start Time	Duration (sec)	RF1 Start Time	RF2 Start Time
		Query End Time		RF1 End Time	RF2 End Time
1	828162148	2014-08-28 19:03:03	00 02:02:39	2014-08-28 19:03:03	2014-08-28 19:06:19
		2014-08-28 21:05:42		7,359	2014-08-28 19:06:19
2	828162149	2014-08-28 19:03:04	00 02:00:01	2014-08-28 19:06:50	2014-08-28 19:10:41
		2014-08-28 21:03:05		7,201	2014-08-28 19:10:41
3	828162150	2014-08-28 19:03:12	00 02:05:36	2014-08-28 19:11:06	2014-08-28 19:15:12
		2014-08-28 21:08:48		7,536	2014-08-28 19:15:12
4	828162151	2014-08-28 19:03:07	00 02:02:40	2014-08-28 19:15:37	2014-08-28 19:21:58
		2014-08-28 21:05:47		7,360	2014-08-28 19:21:58
5	828162152	2014-08-28 19:03:06	00 02:03:07	2014-08-28 19:22:17	2014-08-28 19:25:45
		2014-08-28 21:06:13		7,387	2014-08-28 19:25:45
6	828162153	2014-08-28 19:03:04	00 02:02:43	2014-08-28 19:26:18	2014-08-28 19:29:55
		2014-08-28 21:05:47		7,363	2014-08-28 19:29:55
7	828162154	2014-08-28 19:03:10	00 02:04:55	2014-08-28 19:30:43	2014-08-28 19:35:01
		2014-08-28 21:08:05		7,495	2014-08-28 19:35:01
8	828162155	2014-08-28 19:03:11	00 02:06:26	2014-08-28 19:35:40	2014-08-28 19:39:04
		2014-08-28 21:09:37		7,586	2014-08-28 19:39:04
9	828162156	2014-08-28 19:03:14	00 02:05:30	2014-08-28 19:39:46	2014-08-28 19:43:47
		2014-08-28 21:08:44		7,530	2014-08-28 19:43:47
10	828162157	2014-08-28 19:03:18	00 02:05:28	2014-08-28 19:44:34	2014-08-28 19:48:44
		2014-08-28 21:08:46		7,528	2014-08-28 19:48:44
11	828162158	2014-08-28 19:03:19	00 02:03:02	2014-08-28 19:49:17	2014-08-28 19:58:19
		2014-08-28 21:06:21		7,382	2014-08-28 19:58:19



**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	53.5	4.7	46.3	11.0	38.7	12.6	58.7	20.7	297.3	39.5	121.9	28.5
1	499.2	47.6	341.3	70.7	258.9	114.2	457.4	147.3	1,596.3	305.6	391.1	311.0
2	503.5	58.7	357.0	89.4	340.8	178.2	389.2	466.0	1,077.2	405.8	381.7	255.9
3	575.0	42.5	152.1	121.3	308.2	98.2	400.3	149.1	2,151.2	291.2	406.5	226.7
4	474.2	36.8	320.9	86.9	278.5	138.3	417.9	144.6	1,794.3	353.2	373.8	366.1
5	622.6	36.6	372.9	80.5	307.5	110.7	416.6	119.6	1,587.9	294.5	449.4	261.4
6	492.8	40.8	319.5	92.2	335.6	151.7	471.1	140.9	1,784.4	333.2	380.1	259.3
7	574.4	37.9	311.5	86.9	282.6	116.2	445.2	131.3	2,025.2	332.6	384.1	123.6
8	633.2	59.9	349.2	104.2	283.3	111.2	418.4	137.0	2,003.4	315.7	263.4	287.3
9	565.3	68.3	306.3	110.4	172.3	124.7	393.5	172.3	1,671.5	499.7	536.8	256.3
10	582.8	45.2	379.0	58.9	95.7	183.0	411.0	54.4	1,701.2	361.2	396.5	284.6
11	476.0	51.2	344.7	100.0	356.5	123.8	375.6	467.4	1,386.8	311.1	404.2	283.8
Min	53.5	4.7	46.3	11.0	38.7	12.6	58.7	20.7	297.3	39.5	121.9	28.5
Avg	504.4	44.2	300.1	84.4	254.9	121.9	387.9	179.2	1,589.7	320.3	374.1	245.4
Max	633.2	68.3	379.0	121.3	356.5	183.0	471.1	467.4	2,151.2	499.7	536.8	366.1

Stream ID	Q13	Q14	Q15	Q16	Q17	Q18	Q19	Q20	Q21	Q22	RF1	RF2
0	57.6	21.3	84.7	34.0	4.2	81.2	11.2	20.5	65.0	5.1	54.3	30.1
1	523.9	143.9	652.8	517.5	79.9	266.9	41.9	200.0	344.5	46.5	195.3	30.5
2	504.2	190.6	642.4	241.6	35.1	417.8	80.7	197.9	343.4	43.8	231.3	24.8
3	507.7	153.9	600.6	231.0	22.0	416.6	64.3	157.3	414.4	44.5	246.4	25.0
4	524.9	132.3	567.6	253.9	21.4	384.6	115.8	120.4	399.9	52.5	380.2	19.6
5	571.0	176.2	717.6	297.4	18.7	332.3	74.2	123.9	364.9	49.7	207.9	32.9
6	512.6	125.6	601.2	281.5	25.7	419.0	50.6	163.3	328.7	51.8	217.0	47.6
7	522.6	133.5	560.1	270.3	26.7	420.9	60.1	164.1	433.2	51.1	258.3	39.2
8	526.4	177.0	602.7	219.5	25.4	401.1	140.9	123.3	356.1	46.7	203.6	42.3
9	554.3	106.6	639.5	81.5	29.1	345.6	126.9	309.0	405.7	53.0	240.6	46.9
10	675.8	188.8	698.2	309.3	42.9	328.7	63.7	122.9	481.4	61.9	250.4	32.8
11	517.6	190.8	678.6	287.2	19.0	315.6	72.6	171.8	382.9	63.4	542.0	41.2
Min	57.6	21.3	84.7	34.0	4.2	81.2	11.2	20.5	65.0	5.1	54.3	19.6
Avg	499.9	145.0	587.2	252.1	29.2	344.2	75.2	156.2	360.0	47.5	252.3	34.4
Max	675.8	190.8	717.6	517.5	79.9	420.9	140.9	309.0	481.4	63.4	542.0	47.6

# Table of Contents

<b>TPC BENCHMARK H OVERVIEW</b> .....	<b>9</b>
<b>0 GENERAL ITEMS</b> .....	<b>10</b>
0.1 BENCHMARK SPONSOR .....	10
0.2 PARAMETER SETTINGS.....	10
0.3 CONFIGURATION DIAGRAM .....	10
<b>1 CLAUSE 1: LOGICAL DATABASE DESIGN RELATED ITEMS</b> .....	<b>12</b>
1.1 DATABASE DEFINITION STATEMENTS.....	12
1.2 PHYSICAL ORGANIZATION .....	12
1.3 HORIZONTAL PARTITIONING.....	12
1.4 REPLICATION .....	12
<b>2 CLAUSE 2: QUERIES AND REFRESH FUNCTIONS</b> .....	<b>13</b>
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
<b>3 CLAUSE 3: DATABASE SYSTEM PROPERTIES</b> .....	<b>14</b>
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
<b>4 CLAUSE 4: SCALING AND DATABASE POPULATION</b> .....	<b>18</b>
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
<b>5 CLAUSE 5: PERFORMANCE METRICS AND EXECUTION RULES RELATED ITEMS</b> .....	<b>21</b>
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
<b>6 CLAUSE 6: SUT AND DRIVER IMPLEMENTATION RELATED ITEMS</b> .....	<b>23</b>
6.1 DRIVER .....	23
6.2 IMPLEMENTATION SPECIFIC LAYER (ISL).....	23
6.3 PROFILE-DIRECTED OPTIMIZATION .....	23
<b>7 CLAUSE 7: PRICING</b> .....	<b>24</b>
7.1 HARDWARE AND SOFTWARE USED IN THE PRICED SYSTEM .....	24
7.2 TOTAL THREE YEAR PRICE.....	24
7.3 AVAILABILITY DATE .....	24
<b>8 CLAUSE 8: FULL DISCLOSURE</b> .....	<b>25</b>
8.1 SUPPORTING FILES INDEX TABLE .....	25
<b>9 CLAUSE 9: AUDIT RELATED ITEMS</b> .....	<b>26</b>

9.1 AUDITOR'S REPORT.....	26
<b>APPENDIX A: PRICING INFORMATION.....</b>	<b>29</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 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:**


- 50 x Dell PowerEdge R720xd servers, each with:
  - 768 GB RAM
  - 2 x CPU (Intel® Xeon E5-2680v2, 2.8GHz)
  - 2 x 146 GB SAS 15.000 rpm
- 14 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
- 3 x Dell Networking N2048 Switch – 48 x 1GbE Ports, 2 x 10GbE SPO

Network is 2 x 10Gb & 3 x 1Gb Ethernet. Each server contains 16 physical disks configured as 8 mirrored pairs at the controller level (HW RAID 1).

Priced configuration and benchmarked configuration are identical.

DBMS-Cluster


*Rack Servers*



*2 x 10 Gb Ethernet Switches*



*3 x 1 Gb Ethernet Switches*



*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	599,999,969,200
Order	150,000,000,000
Partsupp	80,000,000,000
Part	20,000,000,000
Customer	15,000,000,000
Supplier	1,000,000,000
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 16 physical disks configured as 8 mirrored pairs at the controller level (HW RAID 1). The resulting 8 devices (1 x 146GB and 7 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	SW RAID0 over 4 x 1199 GB	DBGEN RF1 & RF2 flat files
d04_storage	raw	3 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 * 50	13,597
Internal	1199 GB	14* 50	781,659
<b>Total Storage</b>			<b>795,256</b>
<b>Data Storage Ratio</b>			<b>7.95</b>

\* 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.

<b>Nodes</b>	<b>RAM per Node</b>	<b>Total Memory</b>	<b>Scale Factor</b>	<b>Memory to Database Size Ratio</b>
50	768 GB	38,400 GB	100,000	38.4%

## **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 5.3.6) 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
<b>QphH@100000GB</b>	11,899,598.7	11,612,395.5
<b>QppH@100000GB</b>	11,787,034.6	11,754,288.9
<b>QthH@100000GB</b>	12,013,237.7	11,472,214.9

## 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](http://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.

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](mailto: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](http://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, 50-node Cluster  
Operating System: EXAClusterOS 5.0  
Database Manager: EXASolution 5.0

The results were:

**Performance Metric 11,612,395.5 QphH@100000GB**  
TPC-H Power 11,754,288.9  
TPC-H Throughput 11,472,214.9  
Database Load Time 16:38:43

**Server**

**Dell PowerEdge R720xd, 50-node Cluster (each node with)**

CPU	2 x Intel Xeon E5-2680v2 10C 2.8 GHz (25 MB L3, 10 cores, 20 threads)		
Memory	768 GB		
Disks	<b>Qty</b>	<b>Size</b>	<b>Type</b>
	2	146 GB	15K rpm SAS HDD
	14	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 100,000GB 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)	35-40TB	\$2,100,000	1	\$2,100,000
Support 24x7 (4% per year)	EXA-SUP	\$84,000	3	\$252,000
<b>Total</b>				<b>\$2,352,000</b>

\* 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](mailto:sales@exasol.com)