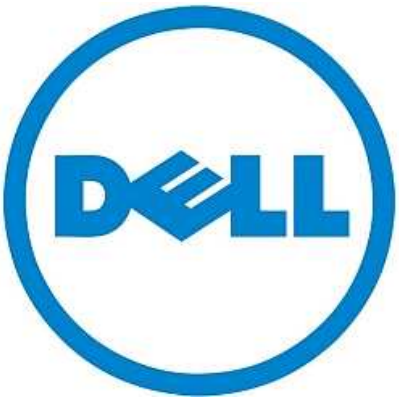

TPC Benchmark™ H Full Disclosure Report

Dell PowerEdge R710
using
EXASolution 4.0

First Edition
April 5, 2011

Dell PowerEdge R710
using
EXASolution 4.0

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.

	Dell PowerEdge R710 using EXASolution 4.0		TPC-H Rev. 2.14.0 Report Date **April 5, 2011**	
Total System Cost	Composite Query per Hour Metric		Price / Performance	
$ 1,736,850	**5,556,993.3** QphH@3000GB		**$ 0.32** $ / QphH@3000GB	
Database Size	Database Manager	Operating System	Other Software	Availability Date
3000 GB	**EXASolution 4.0**	**EXACluster OS 4.0**	**C++ Compiler**	**October 1, 2011**

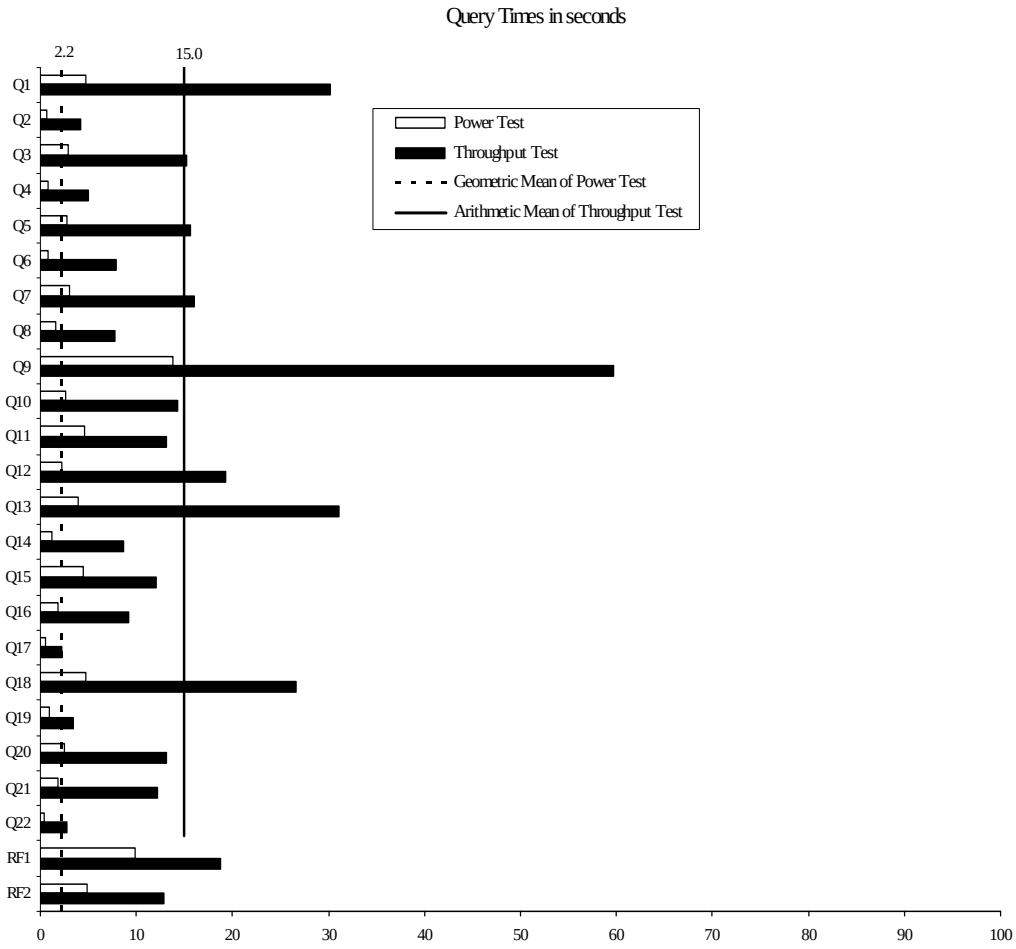

	Storage Redundancy Levels:	
Database Load Time = 0:54:15	Base Tables	2
Load Includes Backup: N	Auxiliary Data Structures	2
Total Data Storage / Database Size = 14.90	DBMS Temporary Space	1
Memory to Database Size Percentage = 64.0	OS and DBMS Software	1

System Configuration:

40 x Dell PowerEdge R710 Server, each with:
2 Intel Xeon X5690 QC 3.46 GHz processors (each is 1 chip, 6 cores, 12 threads)
48 GB RAM
2x 600 GB (10k rpm) internal SAS disks

Total Storage: 44,703 GB (GB = 1024*1024*1024 bytes)

Dell PowerEdge R710 using EXASolution 4.0

TPC-H Rev. 2.14.0

Report Date
April 5, 2011

Description	Part Number	Pricing	Unit Price	Qty	Extended Price	3 yr. Maint. Price
Server Hardware						
R710 with Chassis for up to 8, 2.5-in. Hard Drives	210-31964	1	$12,542.70	40	$501,708.00	included
Intel Xeon X5690,3.46GHz,12M,XN,6.4GT/s,6C	213-12953	1	$0.00	40	$0.00	
Intel Xeon X5690,3.46GHz,12M,XN,6.4GT/s,6C	374-13955	1	$0.00	40	$0.00	
HO Pwr Sply,Redundant,870W UPG	450-12451	1	$0.00	40	$0.00	
Rizer with 2 PCIe x8 +2 PCIe x4 Slots	330-10106	1	$0.00	40	$0.00	
48GB,12x4G,1333M,2R,RDIMM	370-19192	1	$0.00	40	$0.00	
iDRAC6 Enterprise	565-10114	1	$0.00	40	$0.00	
PERC H700 INTG RAID CTL 512MB,x8	405-11392	1	$0.00	40	$0.00	
HDD, 600GB 10K SAS,6G,2.5,Hot Plug	400-20817	1	$0.00	80	$0.00	
DVD+/-RW, SATA, INTERNAL	429-14713	1	$0.00	40	$0.00	
Broadcom GbE 4P LOM with TOE	541-10013	1	$0.00	40	$0.00	
Qlogic Infinband Adapter QLE7340-CK	690-86219	1	$0.00	40	$0.00	
Sliding Ready Rack Rails with Cable Management	770-10763	1	$0.00	40	$0.00	
3Yr ProSupport and Next Business Day On-Site Service	710-16539	1	$0.00	40	$0.00	
PE 4220 42U Rack with Doors and Side Panels	210-26840	1	$2,705.21	3	$8,115.63	
DELL PowerConnect 5548 Switch 48 GbE Ports	210-34476	1	$1,903.24	2	$3,806.48	
RPS720 External Red. PSU for PC 5524 Switch	450-16399	1	$770.87	1	$770.87	
Dell Discount (30%)		1			-$154,320.29	
PATCH CABLE CAT5E 5M GREY STP	A2900570	1	$4.39	80	$351.20	
Mice : Dell Optical USB	570-11147	1	$9.78	1	$9.78	
Keyboard Dell Standard Quietkey USB	580-14755	1	$15.53	1	$15.53	
Qlogic 12300-040, 96-port Infiniband QDR Switch	A4897761	1	$59,250.00	1	$59,250.00	
Qlogic 5.0M Copper 26 AWG QSFP to QSFP Passive	A4897774	1	$160.33	40	$6,413.20	
DELL 17" TFT Display E1709W	855-10385	1	$141.22	3	$423.66	
Dell Discount (10%)		1			-$6,646.34	
				Subtotal	**$419,897.72**	**$0.00**
Storage						
No external storage required						
Software						
EXASolution 4.0 licence (1GB DB RAM*)	EXA-1G	2	$1,560.00	1,260	$1,965,600.00	
EXASOL Support 24x7 (4% per year)	EXA-SUP	2				$235,872.00
Discount (45%)		2			-$884,520.00	
				Subtotal	**$1,081,080.00**	**$235,872.00**
				Total	**$1,500,977.72**	**$235,872.00**

3-Year Cost of Ownership $1,736,850

QphH Rating: 5,556,993.3

$/QphH@3000GB: $0.32

Price Key: 1-Dell: Pricing may be verified by calling 1-800-BUY-DELL
and referencing quote # 24898481.2
contact: Ronny Lenz, ronny_lenz@dell.com
2-EXASOL, contact: Steffen Weissbarth, sales@exasol.com

* Memory size allocated to EXASolution (see also section 7.2)

All discounts are based on list prices and for similar quantities and configurations.

Results independently 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.

Numerical Quantities

Measurement Results

Database Scale Factor	3000 GB
Total Data Storage / Database Size	14.90
Start of Database Load	13:13:39
End of Database Load	14:07:54
Database Load Time	0h 54m 15 s
Query Streams for Throughput Test	9
TPC-H Power	4,866,544.6
TPC-H Throughput	6,345,400.6
TPC-H Composite Query-per-Hour Metric (QphH@3000GB)	5,556,993.3
Total System Price Over 3 Years	$1,736,850
TPC-H Price/ Performance Metric ($/QphH@3000GB)	$0.32

Measurement Interval

Measurement Interval in Throughput Test (Ts)	337 seconds

Duration of Stream Execution

Stream ID	Seed	Start Date	Start Time	End Date	End Time	Duration
Stream 0	330140754	2011-03-30	14:14:54	2011-03-30	14:16:11	1:17 min
Stream 1	330140755	2011-03-30	14:16:11	2011-03-30	14:21:40	5:29 min
Stream 2	330140756	2011-03-30	14:16:12	2011-03-30	14:21:40	5:28 min
Stream 3	330140757	2011-03-30	14:16:11	2011-03-30	14:21:38	5:27 min
Stream 4	330140758	2011-03-30	14:16:12	2011-03-30	14:21:39	5:27 min
Stream 5	330140759	2011-03-30	14:16:11	2011-03-30	14:21:44	5:33 min
Stream 6	330140760	2011-03-30	14:16:11	2011-03-30	14:21:43	5:32 min
Stream 7	330140761	2011-03-30	14:16:11	2011-03-30	14:21:48	5:37 min
Stream 8	330140762	2011-03-30	14:16:11	2011-03-30	14:21:41	5:30 min
Stream 9	330140763	2011-03-30	14:16:11	2011-03-30	14:21:42	5:31 min
Refresh		2011-03-30	14:16:11	2011-03-30	14:20:57	4:46 min

TPC-H Timing Intervals (in seconds)

Query	Power	Stream 1	Stream 2	Stream 3	Stream 4	Stream 5	Stream 6	Stream 7	Stream 8	Stream 9	Min Q_i	Avg Q_i	Max Q_i
1	4.7	27.2	32.0	28.8	28.7	31.5	31.6	30.8	30.4	30.1	27.2	30.1	32.0
2	0.7	2.9	4.1	4.0	6.2	4.6	4.1	3.4	4.5	4.4	2.9	4.3	6.2
3	2.9	16.5	14.4	13.4	14.3	15.8	16.6	14.0	15.5	16.4	13.4	15.2	16.6
4	0.7	3.9	5.9	3.7	5.3	3.9	4.9	5.9	6.8	5.1	3.7	5.1	6.8
5	2.7	15.7	15.2	18.1	16.6	12.9	16.2	17.0	15.9	13.4	12.9	15.7	18.1
6	0.8	8.8	8.3	7.0	8.2	8.5	8.4	7.0	8.2	6.1	6.1	7.8	8.8
7	3.0	16.6	15.8	15.7	16.1	15.4	17.3	16.0	17.6	14.0	14.0	16.1	17.6
8	1.6	6.1	8.6	9.0	7.1	6.8	8.0	7.0	7.6	9.0	6.1	7.7	9.0
9	13.8	59.9	54.8	53.9	55.1	47.0	60.3	91.1	56.1	59.3	47.0	59.7	91.1
10	2.6	12.3	13.3	14.5	14.8	16.2	15.5	13.6	13.0	16.0	12.3	14.4	16.2
11	4.5	11.2	12.3	11.2	14.3	13.8	13.5	14.6	11.2	15.8	11.2	13.1	15.8
12	2.3	22.5	21.3	20.1	21.6	24.7	20.0	2.2	20.8	20.9	2.2	19.3	24.7
13	4.0	31.4	29.5	32.5	32.0	33.4	31.0	30.5	31.8	28.2	28.2	31.2	33.4
14	1.2	8.3	12.4	9.7	7.2	9.6	7.6	5.9	8.7	8.6	5.9	8.7	12.4
15	4.5	14.0	13.6	11.2	12.1	13.8	12.7	4.8	12.6	13.9	4.8	12.1	14.0
16	1.9	9.4	8.0	10.2	11.1	10.6	10.2	8.2	8.5	6.6	6.6	9.2	11.1
17	0.5	2.1	2.2	2.9	2.2	1.9	1.9	1.9	2.3	2.2	1.9	2.2	2.9
18	4.7	24.0	24.1	30.9	24.8	27.8	24.5	30.8	25.7	27.5	24.0	26.7	30.9
19	0.9	4.0	2.8	3.1	3.2	3.3	3.0	2.7	5.3	3.8	2.7	3.5	5.3
20	2.5	14.7	15.7	11.9	11.7	12.7	12.3	14.4	11.4	13.5	11.4	13.2	15.7
21	1.9	13.7	11.3	11.9	12.9	14.4	8.1	12.8	13.0	12.4	8.1	12.3	14.4
22	0.4	2.9	2.7	2.4	2.2	3.4	3.2	2.2	3.0	2.9	2.2	2.8	3.4
RF1	9.8	15.9	16.4	19.2	21.4	19.7	21.9	17.0	18.5	18.9	15.9	18.8	21.9
RF2	4.9	13.3	12.0	14.9	11.2	11.5	14.1	10.1	13.8	15.0	10.1	12.9	15.0

Table of Contents

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 system and database parameters used in this benchmark.

0.3 Configuration Diagram

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

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

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

	System components
40	Dell PowerEdge R710 servers, each with:
2	CPU (Intel® Xeon X5690, 3.46GHz)
48	GB RAM
2	600 GB SAS 10.000 rpm
1	RAID controller onboard (BBU included)
4	Ethernet Ports
1	Infiniband network adapter
2	Dell PowerConnect 5548 (1Gb Ethernet Switch, 48 ports)
1	QLogic 12300-BS40 (40Gb Infiniband Switch, 96 ports)

Network is 1Gb Ethernet and 40Gb Infiniband. Each server contains 2 physical disks that are configured as a mirrored pair at the controller level (HW RAID 1).

Priced configuration and benchmarked configuration are identical.

DBMS-Cluster

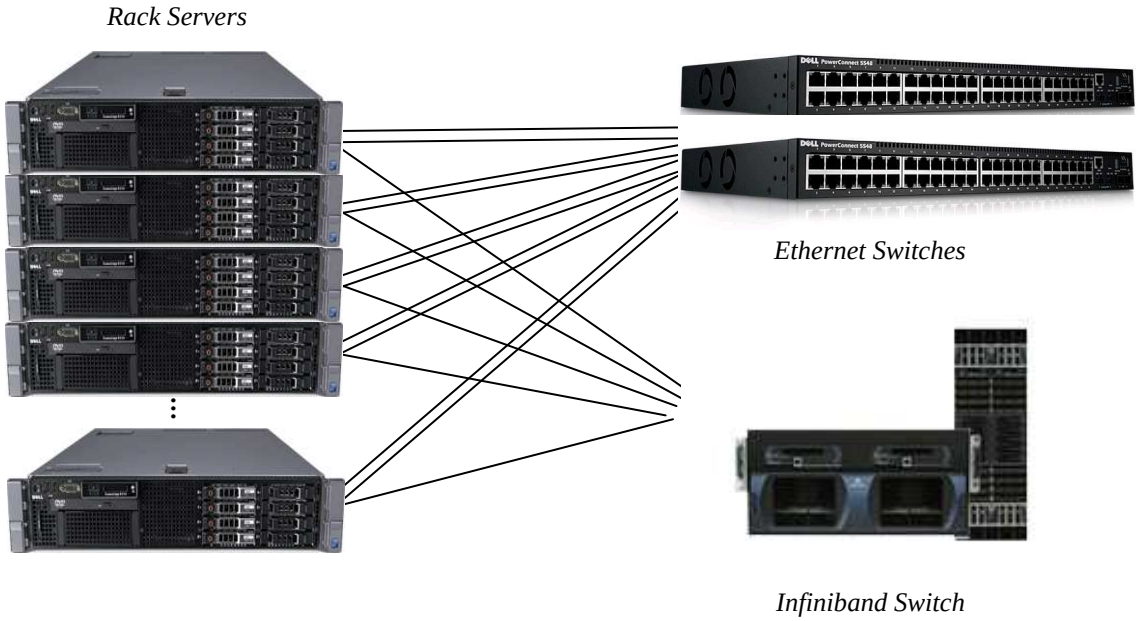

Figure 1.1: Benchmarked and priced system configuration

1 Clause 1: Logical Database Design Related Items

1.1 Database Definition Statements

Listings must be provided for all table definition statements and all other statements used to set up the test and qualification databases (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 2.14.0 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 2.14.0 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)
- 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 15 execution streams were prepared.
3. The 100 ACID transactions per stream were executed from 15 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 adopted 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:

5. Start a query and verify that the row was retrieved.
6. Start an update transaction, read and update the same row. Wait before commit.
7. Start the same query and verify that the row retrieved has not changed.
8. Commit the update transaction
9. 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 15 executions streams were prepared.
3. After that at least 100 ACID transactions were submitted from each of the 15 execution streams.
4. A durability failure was inducted (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 8-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.

During the durability test, one disk was removed. The test continued uninterrupted, because of the RAID protection (see also section 1.3).

Upon replacement/re-insert of the drive, the data files were rebuild to a consistent state by the RAID controller.

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. Two system crashs were performed by turning off the power during the durability test: one test for only one node and the other test for the whole cluster.

When power was restored, in both cases the system rebooted automatically and the database was restarted manually.

3.5.3 Memory Failure

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

The system crash and memory failure tests were combined as explained in section 3.5.2.

3.5.4 Node or Controller Failure

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

During that durability test, one server was turned off and both disks were replaced by empty disks. The node was started again and 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	18,000,048,306
Order	4,500,000,000
Partsupp	2,400,000,000
Part	600,000,000
Customer	450,000,000
Supplier	30,000,000
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 2 physical disks that are configured as a mirrored pair at the controller level (HW RAID 1).

The resulting device is divided into 3 partitions. All benchmark- and database-relevant data is stored in partition `Data` (Tables, Indexes, Temp, Logs, Flat Files).

Partition	Size
System	49 GB
Swap	2 GB
Data	499 GB

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 files are 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 5.2.

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 2.14.0 was modified (changes made to a header file) to generate the database population for this benchmark. This header file change is included in the supporting files archive.

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	600 GB	80	44,703

Scale Factor	Data Storage Ratio
3000	14.90

** 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.

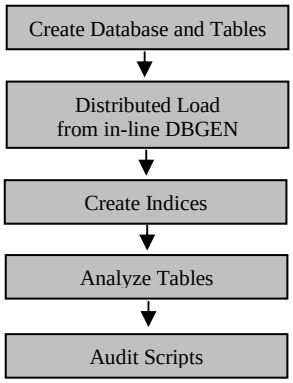

4.9 Qualification Database Configuration

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 percentage
40	48 GB	1920 GB	3,000	64.0%

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.

See 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 Timing Intervals for Each Query and Each Refresh Function for Each Stream

The timing intervals (see Clause 5.3.6) for each query of each stream and for each refresh function must be reported for the throughput test.

See Numerical Quantities Summary in the Executive Summary.

5.9 Performance Metrics

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

See Numerical Quantities Summary in the Executive Summary.

5.10 The Performance Metric and Numerical Quantities from Both Runs

A description of the method used to determine the reproducibility of the measurement results must be reported. This must include the performance metrics (QppH and QthH) from reproducibility runs.

Run ID	QppH@3000GB	QptH@3000GB	QphH@3000GB
Run 1	4,906,506.3	6,460,423.0	5,630,107.1
Run 2	4,866,544.6	6,345,400.6	5,556,993.3
% Difference	0.82 %	1.81 %	1.32 %

5.11 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.

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 (DB RAM 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 DB RAM 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 4.0	October 01, 2011

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	Pathname
1	OS and DB Parameters	`SupportingFiles/TPCH/settings.txt`

	DB Creation scripts	`SupportingFiles/TPCH/sql/create_user.sql`
		`SupportingFiles/TPCH/sql/create_schema.sql`
		`SupportingFiles/TPCH/sql/create_indices.sql`

2	Query validation text	`SupportingFiles/TPCH/scripts/tpc_h_run_validation.sh` `SupportingFiles/TPCH/sql/validation.sql`
	Refresh function details	`SupportingFiles/TPCH/scripts/tpc_h_run_full.sh`
3	ACID Test Scripts	`SupportingFiles/TPCH/ACID/`
	ACID Test Results	`SupportingFiles/TPCH/ACID_output/`
4	DBGEN Modifications	`SupportingFiles/TPCH/tpch_archives/tpch*.patch`
	Database Load Details	`SupportingFiles/TPCH/scripts/load_init.sh`
5	Execution log	`SupportingFiles/TPCH/result/run1/` `SupportingFiles/TPCH/result/run2/`
	Query validation output	`SupportingFiles/TPCH/result/validation/`
6	Implementation scripts	`SupportingFiles/TPCH/scripts/`

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.
125 West Monroe Street
Colorado Springs, CO 80907
Phone: (719) 473-7555
Fax: (719) 473-7554
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

INFO SIZING

TPC TRANSACTION PROCESSING PERFORMANCE COUNCIL CERTIFIED AUDITOR

Benchmark Sponsor: David J. Morse
Primary TPC Representative
Dell
Round Rock, Texas
United States

April 02, 2011

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

Platform: **Dell PowerEdge R710, 40-node Cluster**

Database Manager: **EXASolution 4.0**

Operating System: **EXACluster OS 4.0**

The results were:

CPU (Speed)	Memory	Disks	**QphH@3000GB**
Dell PowerEdge R710, 40-node cluster (each node with)			
2 x Intel Xeon X5690 Quad (3.46GHz)	48GB Main	2 x 600GB SAS	**5,556,993.3**

In my opinion, this performance result was produced in compliance with the TPC's 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 3,000GB 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
- The execution of the queries against the SF1 database produced compliant answers
- A compliant implementation specific layer was used to drive the tests
- The throughput tests involved 9 query streams
- The ratio between the longest and the shortest query was such that no query timing was adjusted
- The execution times for queries and refresh functions were correctly measured and reported
- The repeatability of the measured results was verified
- The system pricing was verified for major components and maintenance
- The major pages from the FDR were verified for accuracy

Additional Audit Notes:

DBGEN version 2.14.0 was modified (changes made to a header file) to add an EXASolution option. This change complies with the requirements for the benchmark. A description of this change is included in the supporting files provided with the Full Disclosure Report.

Respectfully Yours,

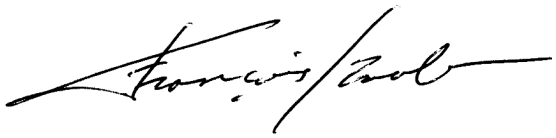

François Raab
President

Appendix A: Pricing Information

EXASOL AG
Neumeyerstrasse 48
D-90411 Nürnberg
T +49 911 23991 0
F +49 911 23991 5241

March 21, 2011

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

Product	Part Number	Price per GB	Quantity	Extended Price
EXASolution 4.0 (1GB DB RAM*) (EXACluster OS 4.0 included)	EXA-1G	$ 1,560	1260	$1,965,600
Discount (45%)				-$884,520
Support 24x7 (4% per year)	EXA-SUP			$235,872
Total				**$1,316,952**

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

This quote is valid for 60 days.

EXASOL Pricing Contact:
Steffen Weissbarth
sales@exasol.com