

TPC Benchmark™ H
Full Disclosure Report
for
Lenovo System x®3850 X6
using
Microsoft® SQL Server® 2016
Enterprise Edition
and
Microsoft Windows Server® 2012
R2 Standard Edition

TPC-H™ Version 2.17.1

lenovo

First Edition
Submitted for Review
March 10, 2016

First Edition – March 2016

The information contained in this document is distributed on an AS IS basis without any warranty either expressed or implied. The use of this information or the implementation of any of these techniques is the customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item has been reviewed by Lenovo for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environment do so at their own risk.

In this document, any references made to Lenovo licensed program are not intended to state or imply that only Lenovo's licensed program may be used; any functionally equivalent program may be used.

This publication was produced in the United States. Lenovo may not offer the products, services, or features discussed in this document in other countries, and the information is subject to change without notice. Consult your local Lenovo representative for information on products and services available in your area.

© Copyright Lenovo Corporation 2016. All rights reserved.

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

Trademarks

Lenovo, the Lenovo logo and System x are trademarks or registered trademarks of Lenovo Corporation.

The following terms used in this publication are trademarks of other companies as follows: TPC Benchmark, TPC-H, QppH, QthH and QphH are trademarks of Transaction Processing Performance Council; Intel and Xeon are trademarks or registered trademarks of Intel Corporation; Microsoft, Windows and SQL Server are trademarks or registered trademarks of Microsoft Corporation. Other company, product, or service names, which may be denoted by two asterisks (**), may be trademarks or service marks of others.

Notes

¹ GHz and MHz only measures microprocessor internal clock speed, not application performance. Many factors affect application performance.

² When referring to hard disk capacity, GB, or gigabyte, means one thousand million bytes. Total user-accessible capacity may be less.



System x® 3850 X6 Microsoft® SQL Server® 2016

TPC-H Rev. 2.17.1
TPC-Pricing 2.0.0

Report Date 3/10/16

Description	Part Number	Price Source	Unit Price	Quantity	Extended Price	3-Yr. Maint. Price
Server Hardware						
Lenovo System x3850 X6 Configure-To-Order, includes:	6241AC1	1	270,101	1	270,101	
x3850 X6 4U Chassis	ASMH			1		
Midplane for 4U Chassis	A4A4			1		
x3850/x3950 X6 I/O Planar III	ASFB			1		
X6 DDR4 Compute Book E7-8890 v3 18C 2.5GHz 165W	AS8B, AS91			4		
X6 Storage Book	A4A1			1		
X6 Half-length I/O Book	A4A2			2		
4x 2.5" HDD Riser	A4A6			1		
ServeRAID M5210 SAS/SATA Controller	A3YZ			1		
300GB 10K 12Gbps SAS 2.5" G3HS HDD	AT89			2		
480GB Enterprise Entry SATA G3HS 2.5" SSD	AT8T			2		
Intel I350-T4 ML2 Quad Port GbE Adapter	A40R			1		
1400W HE Redundant Power Supply for altitudes >5000 meters	A54D			4		
Rail Kit	A4AA			1		
Power Cable	6311			4		
32GB TruDDR4 Memory PC4-17000 CL15 2133MHz LP RDIMM	A5UJ			96		
3200GB Enterprise Value io3 Flash Adapter	ARYM			6		
3m Blue Cat5e Cable	3803			1		
Preferred Pro Keyboard USB - US English 103P RoHS v2	00AM600	1	29	1	29	
2-Button Optical Mouse - Black - USB	40K9200	1	19	1	19	
ServicePac for 3-Year 24x7x4 Support (x3850 X6)	00GW689	1	1,500	1		1500
ThinkVision E1922 18.5-inch LED Backlit LCD Monitor	60B8AAR6US	1	100	1	100	
Subtotal					270,249	1,500
Server Software						
SQL Server 2016 Enterprise Edition	7JQ-00256	2	13,472.50	36	485,010	
Windows Server 2012 R2 Standard Edition	P73-06284	2	735	2	1,470	
Microsoft Problem Resolution Services	N/A	2	259	1		259
Subtotal					486,480	259
Infrastructure						
S2 42U Standard Rack	93074RX	1	1,565	1	1,565	
ServicePac for 3-Year 24x7x4 Support (Rack)	41L2760	1	315	1		315
Subtotal					1,565	315
Total					758,294	2,074
Dollar Volume Discount (See Note 1)					25.87%	1
						70,779

Pricing: 1 - Lenovo 1-877-782-7134; 2 - Microsoft

Note 1: Discount applies to all line items where Pricing=1; pricing is for these or similar quantities.

Discounts for similarly sized configurations will be similar to what is quoted here, but may vary based on the specific components priced.

Three-Year Cost of Ownership USD: \$689,589
QphH@3000GB: 969,504.20
SD/QphH@3000GB: \$0.72

Benchmark results and test methodology audited by Francois Raab for 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 section of the TPC benchmark specifications. If you find that stated prices are not available according to these terms, please inform the TPC at pricing@tpc.org. Thank you.



System x® 3850 X6
Microsoft® SQL Server® 2016

TPC-H Rev.
2.17.1
TPC-Pricing
2.0.0

Report Date
3/10/16

Measurement Results

Database Scale Factor	3000
Total Data Storage/Database Size	6.34
Memory/Database Size	102.4%
Start of Database Load	03/04/2016 21:43:17
End of Database Load	03/05/2016 01:01:35
Database Load Time	00d 02h 58m 00s
Query Streams for Throughput Test	8
TPC-H Power	1,346,018.5
TPC-H Throughput	698,310.1
TPC-H Composite Query-per-Hour (QphH@3000GB)	969,504.2
Total System Price over 3 Years	\$689,589 USD
TPC-H Price/Performance Metric (\$/QphH@3000GB)	\$ 0.72 USD

Measurement Interval

Measurement Interval in Throughput Test (Ts) = 2,722

Duration of Stream Execution:

Power Run	Seed	Query Start Time	Duration	RF1 Start Time	RF2 Start Time
		Query End Time	(sec)	RF1 End Time	RF2 End Time
	305010135	2016-03-05 09:43:45	254	2016-03-05 09:43:04	2016-03-05 09:47:59
		2016-03-05 09:47:59		2016-03-05 09:43:45	2016-03-05 09:49:01
Throughput Stream	Seed	Query Start Time	Duration	RF1 Start Time	RF2 Start Time
		Query End Time	(sec)	RF1 End Time	RF2 End Time
1	305010136	2016-03-05 09:49:00	1,782	2016-03-05 10:19:58	2016-03-05 10:20:40
		2016-03-05 10:18:42		2016-03-05 10:20:40	2016-03-05 10:21:59
2	305010137	2016-03-05 09:49:00	1,731	2016-03-05 10:21:59	2016-03-05 10:22:38
		2016-03-05 10:17:51		2016-03-05 10:22:38	2016-03-05 10:23:44
3	305010138	2016-03-05 09:49:00	1,814	2016-03-05 10:23:44	2016-03-05 10:24:24
		2016-03-05 10:19:14		2016-03-05 10:24:24	2016-03-05 10:25:41
4	305010139	2016-03-05 09:49:00	1,829	2016-03-05 10:25:42	2016-03-05 10:26:21
		2016-03-05 10:19:29		2016-03-05 10:26:21	2016-03-05 10:27:23
5	305010140	2016-03-05 09:49:00	1,859	2016-03-05 10:27:24	2016-03-05 10:28:04
		2016-03-05 10:19:59		2016-03-05 10:28:04	2016-03-05 10:29:10
6	305010141	2016-03-05 09:49:00	1,827	2016-03-05 10:29:10	2016-03-05 10:29:50
		2016-03-05 10:19:27		2016-03-05 10:29:50	2016-03-05 10:30:54
7	305010142	2016-03-05 09:49:00	1,835	2016-03-05 10:30:54	2016-03-05 10:31:33
		2016-03-05 10:19:35		2016-03-05 10:31:33	2016-03-05 10:32:38
8	305010143	2016-03-05 09:49:00	1,768	2016-03-05 10:32:39	2016-03-05 10:33:18
		2016-03-05 10:18:28		2016-03-05 10:33:18	2016-03-05 10:34:22

lenovo		System x® 3850 X6 Microsoft® SQL Server® 2016									TPC-H Rev. 2.17.1 TPC-Pricing 2.0.0	
											Report Date 3/10/16	
Stream ID	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11	Q12
0	5.8	1.8	10.1	9.7	11.2	0.4	6.9	9.5	29.2	5.7	18.1	3.5
1	15.0	16.8	42.8	17.7	140.6	18.9	54.9	88.8	92.6	98.2	50.8	72.4
2	15.9	15.5	57.6	27.0	55.4	1.8	70.2	128.8	86.0	105.0	49.5	119.6
3	93.7	21.8	25.5	27.2	21.6	8.2	182.2	92.4	177.8	131.7	42.3	77.4
4	8.3	15.5	34.5	26.8	92.9	3.7	33.3	125.8	152.7	51.8	38.5	85.5
5	16.2	16.3	25.4	38.5	138.1	20.2	46.6	60.5	33.3	67.0	46.9	150.3
6	85.8	14.5	25.7	13.0	44.5	37.9	87.2	106.9	119.8	118.8	29.8	160.1
7	52.7	15.3	20.9	76.6	81.1	2.2	83.6	109.8	133.5	115.4	31.6	48.5
8	10.8	17.6	57.1	81.1	37.2	15.8	46.5	98.3	222.6	41.8	28.5	97.8
Min	5.8	1.8	10.1	9.7	11.2	0.4	6.9	9.5	29.2	5.7	18.1	3.5
Avg	33.8	15.0	33.3	35.3	69.2	12.1	67.9	91.2	116.4	81.7	37.3	90.6
Max	93.7	21.8	57.6	81.1	140.6	37.9	182.2	128.8	222.6	131.7	50.8	160.1
Stream ID	Q13	Q14	Q15	Q16	Q17	Q18	Q19	Q20	Q21	Q22	RF1	RF2
0	42.3	2.6	4.3	7.2	3.9	41.0	1.7	4.9	28.2	5.0	40.7	61.5
1	174.0	9.3	30.8	31.5	12.5	250.5	132.8	41.5	357.5	30.9	40.4	78.8
2	170.1	10.1	27.2	23.0	23.4	329.6	55.9	45.8	236.8	76.1	39.1	65.9
3	75.5	13.3	19.2	28.4	106.7	89.1	135.5	71.5	295.4	76.3	39.3	77.3
4	52.2	8.7	31.4	30.4	33.3	498.5	163.0	33.4	187.0	120.4	39.3	62.2
5	138.5	7.8	28.9	31.2	16.7	445.0	108.0	47.9	323.3	51.4	40.2	65.9
6	80.4	9.7	8.8	20.6	58.6	347.4	147.2	25.8	220.4	63.3	39.4	64.1
7	112.3	9.3	9.4	20.6	116.2	318.7	149.0	27.7	246.5	52.6	38.5	65.1
8	192.9	21.8	17.8	47.0	15.5	262.4	113.8	37.0	214.9	88.4	39.6	63.2
Min	42.3	2.6	4.3	7.2	3.9	41.0	1.7	4.9	28.2	5.0	38.5	61.5
Avg	115.4	10.3	19.8	26.7	43.0	286.9	111.9	37.3	234.4	62.7	39.6	67.1
Max	192.9	21.8	31.4	47.0	116.2	498.5	163.0	71.5	357.5	120.4	40.7	78.8

Table of Contents

Preface	9
General Items	11
Benchmark Sponsor	11
Parameter Settings	11
Configuration Diagrams	11
Clause 1 – Logical Database Design Related Items	13
Database Table Definitions	13
Database Physical Organization	13
Horizontal/Vertical Partitioning	13
Replication	13
Clause 2 – Queries and Update Functions Related Items	14
Query Language	14
Random Number Generation	14
Substitution Parameters Generation	14
Query Text and Output Data from Database	14
Query Substitution Parameters and Seeds Used	14
Query Isolation Level	14
Refresh Function Implementation	15
Clause 3 – Database System Properties Related Items	16
Atomicity Requirements	16
Consistency Requirements	16
Isolation Requirements	18
Durability Requirements	19
Clause 4 – Scaling and Database Population Related Items	21
Initial Cardinality of Tables	21
Distribution of Tables and Logs	21
Database Partition / Replication Mapping	22
RAID Implementation	22
DBGEN Modifications	22
Database Load Time	22
Data Storage Ratio	22
Database Load Mechanism Details and Illustration	22
Qualification Database Configuration	23
Clause 5 – Performance Metrics and Execution Rules Related Items	24
System Activity between Load and Performance Tests	24
Steps in the Power Test	24
Timing Intervals for Each Query and Refresh Function	24
Number of Streams for the Throughput Test	24
Start and End Date/Times for Each Query Stream	24
Total Elapsed Time for the Measurement Interval	24
Refresh Function Start Date/Time and Finish Date/Time	24
Timing Intervals for Each Query and Each Refresh Function for Each Stream	25
Performance Metrics	25
Performance Metric and Numerical Quantities from Both Runs	25
System Activity between Tests	25
Clause 6 – SUT and Driver Implementation Related Items	26
Driver	26
Implementation-Specific Layer	26
Profile-Directed Optimization	26
Clause 7 – Pricing Related Items	27
Hardware and Software Components	27
Three-Year Cost of System Configuration	27

Availability Dates	27
Country-Specific Pricing	27
Clause 8 – Full Disclosure	28
8.1 Supporting Files Index Table	28
Clause 9 – Audit Related Items	29
Auditor	29
Attestation Letter	29
Clause 10 – Price Quotes.....	32

Preface

TPC Benchmark H Standard Specification was developed by the Transaction Processing Performance Council (TPC). It was released on February 26, 1999, and most recently revised (Revision 2.17.1). This is the full disclosure report for benchmark testing of the Lenovo System x3850 X6 according to the TPC Benchmark H Standard Specification.

The TPC Benchmark H is a decision support benchmark. It consists of a suite of business-oriented ad hoc queries and concurrent data modifications. The queries and the data populating the database have been chosen to have broad industrywide 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 set 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 (e.g., via a point-and-click GUI interface);
- 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 with specific population and scaling requirements;
- Are implemented with constraints derived from staying closely synchronized with an on-line production database.

The TPC-H operations are modeled as follows:

- The database is continuously available 24 hours a day, 7 days a week, for ad-hoc queries from multiple end users and data modifications against all tables, except possibly during infrequent (e.g., once a month) maintenance sessions.
- The TPC-H database tracks, possibly with some delay, the state of the OLTP database through ongoing refresh functions, which batch together a number of modifications impacting some part of the decision support database.
- Due to the worldwide nature of the business data stored in the TPC-H database, the queries and the refresh functions may be executed against the database at any time, especially in relation to each other. In addition, this mix of queries and refresh functions is subject to specific ACIDity requirements, since queries and refresh functions may execute concurrently.
- To achieve the optimal compromise between performance and operational requirements, the database administrator can set, once and for all, the locking levels and the concurrent scheduling rules for queries and refresh functions.

The minimum database required to run the benchmark holds business data from 10,000 suppliers. It contains almost 10 million rows representing a raw storage capacity of about 1 gigabyte. Compliant benchmark implementations may also use one of the larger permissible database populations (e.g., 100 gigabytes), as defined in Clause 4.1.3).

The performance metric reported by TPC-H is called the TPC-H Composite Query-per-Hour Performance Metric (QphH@Size), and reflects multiple aspects of the capability of the system to process queries. These aspects include the selected database size against which the queries are executed, the query processing power when queries are submitted by a single stream, and the query throughput when queries are submitted by multiple concurrent users. The TPC-H Price/Performance metric is expressed as \$/QphH@Size. To be compliant with the TPC-H standard, all references to TPC-H results for a given configuration must include all required reporting components (see Clause 5.4.6). The TPC believes that comparisons of TPC-H results measured against different database sizes are misleading and discourages such comparisons.

The TPC-H database must be implemented using a commercially available database management system (DBMS), and the queries executed via an interface using dynamic SQL. The specification provides for variants of SQL, as implementers are not required to have implemented a specific SQL standard in full.

Benchmarks results are highly dependent upon workload, specific application requirements, and systems design and implementation. Relative system performance will vary as a result of these and other factors. Therefore, TPC-H

should not be used as a substitute for specific customer application benchmarking when critical capacity planning and/or product evaluation decisions are contemplated.

General Items

Benchmark Sponsor

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

This benchmark was sponsored by Lenovo Corporation.

Parameter Settings

Settings must be provided for all customer-tunable parameters and options that 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.*

See the Supporting File, “Tunable Parameters,” which contains a list of all database parameters and operating system parameters.

Configuration Diagrams

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 and type of disk units (and controllers, if applicable)*
- *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 run-time execution location of software components (e.g., DBMS, query processing tools/languages, middleware components, software drivers, etc.).*

The configuration diagram for the tested and priced system is provided on the following page.

Measured Configuration



Lenovo System x3850 X6

Quantity	Lenovo System x3850 X6 Configure-To-Order, includes:
1	x3850 X6 4U Chassis + Midplane
4	X6 Compute Book with Intel Xeon Processor E7-8890 v3
96	32GB PC3L-12800 ECC DDR3 1600MHz LP LRDIMM
1	X6 Primary I/O Book + X6 Storage Book
2	Half Length PCI-e Expansion Card
2	Lenovo 4x 2.5" HS SAS/SATA/SSD HDD Backplane
1	ServerAID M5210 SAS/SATA Controller for Lenovo System x
4	Lenovo 1400W HE Redundant Power Supply
2	200GB 2.5" G3HS Enterprise SAS SSD
2	480GB SAS 2.5" 10K rpm HDD
6	3200GB io3 Enterprise Flash PCIe adapters

Clause 1 – Logical Database Design Related Items

Database Table Definitions

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)

See the Supporting Files for the scripts that were used to set up the TPC-H test and qualification databases.

Database Physical Organization

The physical organization of tables and indexes 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.

See the Supporting Files for the scripts that were used to create the indexes on the test and qualification databases

No column reordering is used.

Horizontal/Vertical Partitioning

Horizontal partitioning of tables and rows in the test and qualification databases must be disclosed (see Clause 1.5.4).

Horizontal partitioning on L_SHIPDATE and O_ORDERDATE is used and granularity is week.

Replication

Any replication of physical objects must be disclosed and must conform to the requirements of Clause 1.5.6).

Replication was not used.

Clause 2 – Queries and Update Functions Related Items

Query Language

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

SQL was the query language used.

Random Number Generation

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

The TPC-supplied DBGEN version 2.17.0 and QGEN version 2.17.0 were used to generate all database populations.

Substitution Parameters Generation

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

The supplied QGEN version 2.17.0 was used to generate the substitution parameters.

Query Text and Output Data from Database

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

See the Supporting Files for the query text and query output. The following modifications were used:

- In Q1, Q4, Q5, Q6, Q10, Q12, Q14, Q15 and Q20, the “dateadd” function is used to perform date arithmetic.
- In Q7, Q8 and Q9, the “datepart” function is used to extract part of a date (e.g., “YY”).
- In Q2, Q3, Q10, Q18 and Q21, the “top” function is used to restrict the number of output rows.
- In Q1, the “count_big” function is used in place of “count”.

Query Substitution Parameters and Seeds Used

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

See the Supporting Files for the seed and query substitution parameters used.

Query Isolation Level

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

The queries and transactions were run with repeatable read isolation level.

Refresh Function Implementation

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

See the Supporting Files for the source code for the refresh function

Clause 3 – Database System Properties Related Items

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. The results of the ACID tests must be disclosed, along with a description of how the ACID requirements were met. This includes disclosing the code written to implement the ACID Transaction and Query.

All ACID tests were conducted according to specifications. The Atomicity, Isolation, Consistency and Durability tests were performed on the Lenovo System x3850 X6 server. See the Supporting Files for the ACID transaction source code.

Atomicity of Completed Transactions

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

The following steps were performed to verify the Atomicity of completed transactions.

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

Atomicity of Aborted Transactions

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

The following steps were performed to verify the Atomicity of the aborted ACID transaction:

1. The total price from the ORDER table and the extended price from the LINEITEM table were retrieved for a randomly selected order key.
2. The ACID Transaction was performed using the order key from step 1. The transaction was stopped prior to the commit.
3. The ACID Transaction was ROLLED BACK. .
4. The total price from the ORDER table and the extended price from the LINEITEM table were retrieved for the same order key used in steps 1 and 2. It was verified that the appropriate rows had not been changed.

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:

$$O_TOTALPRICE = SUM(L_EXTENDEDPRICE * (1 - L_DISCOUNT) * (1 + L_TAX))$$

for each ORDER and LINEITEM defined by (O_ORDERKEY=L_ORDERKEY)

Consistency Tests

Verify that the ORDER and LINEITEM tables are initially consistent as defined in Clause 3.3.2.1, based on a random sample of at least 10 distinct values of O_ORDERKEY.

Consistency was tested as part of the durability tests.

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.

Isolation Test 1 - Read-Write Conflict with Commit

This test demonstrates isolation for the read-write conflict of a read-write transaction and a read-only transaction when the read-write transaction is committed.

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

1. An ACID Transaction was started for a randomly selected O_KEY, L_KEY and DELTA. The ACID Transaction was suspended prior to Commit.
2. An ACID query was started for the same O_KEY used in step 1. The ACID query blocked and did not see any uncommitted changes made by the ACID Transaction.
3. The ACID Transaction was resumed and committed.
4. The ACID query completed. It returned the data as committed by the ACID Transaction.

Isolation Test 2 - Read-Write Conflict with Rollback

This test demonstrates isolation for the read-write conflict of read-write transaction and read-only transaction when the read-write transaction is rolled back.

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

1. An ACID transaction was started for a randomly selected O_KEY, L_KEY and DELTA. The ACID Transaction was suspended prior to Rollback.
2. An ACID query was started for the same O_KEY used in step 1. The ACID query did not see any uncommitted changes made by the ACID Transaction.
3. The ACID Transaction was ROLLED BACK.
4. The ACID query completed.

Isolation Test 3 - Write-Write Conflict with Commit

This test demonstrates isolation for the write-write conflict of two update transactions when the first transaction is committed.

The following steps were performed to verify isolation of two update transactions:

1. An ACID Transaction T1 was started for a randomly selected O_KEY, L_KEY and DELTA. The ACID transaction T1 was suspended prior to Commit.
2. Another ACID Transaction T2 was started using the same O_KEY and L_KEY and a randomly selected DELTA.
3. T2 waited.
4. The ACID transaction T1 was allowed to Commit and T2 completed.
5. It was verified that:

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

Isolation Test 4 - Write-Write Conflict with Rollback

This test demonstrates isolation for write-write conflict of two update transactions when the first transaction is rolled back.

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

1. An ACID Transaction T1 was started for a randomly selected O_KEY, L_KEY and DELTA. The ACID Transaction T1 was suspended prior to Rollback.

2. Another ACID Transaction T2 was started using the same O_KEY and L_KEY used in step 1 and a randomly selected DELTA.
3. T2 waited.
4. T1 was allowed to ROLLBACK and T2 completed.
5. It was verified that T2.L_EXTENDEDPRICE = T1.L_EXTENDEDPRICE.

Isolation Test 5 - Concurrent Read and Write Transactions on Different Tables

This test demonstrates the ability of read and write transactions affecting different database tables to make progress concurrently.

The following steps were performed:

1. An ACID Transaction T1 for a randomly selected O_KEY, L_KEY and DELTA. The ACID Transaction T1 was suspended prior to Commit.
2. Another ACID Transaction T2 was started using random values for PS_PARTKEY and PS_SUPPKEY.
3. T2 completed.
4. T1 completed and the appropriate rows in the ORDER, LINEITEM and HISTORY tables were changed.

Isolation Test 6 - Update Transactions during Continuous Read-Only Query Stream

This test demonstrates that the continuous submission of arbitrary (read-only) queries against one or more tables of the database does not indefinitely delay update transactions affecting those tables from making progress.

The following steps were performed:

1. An ACID Transaction T1 was started, executing Q1 against the qualification database. The substitution parameter was chosen from the interval [0..2159] so that the query ran for a sufficient amount of time.
2. Before T1 completed, an ACID Transaction T2 was started using randomly selected values of O_KEY, L_KEY and DELTA.
3. T2 completed before T1 completed.
4. It was verified that the appropriate rows in the ORDER, LINEITEM and HISTORY tables were changed.

Durability Requirements

The SUT must guarantee durability: the ability to preserve the effects of committed transactions and ensure database consistency after recovery from any one of the failures listed in Clause 3.5.3.

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.

The OS was stored on a RAID-1 protected array of 2 physical drives. The database files were stored on 6 non-raided Enterprise io3 Flash drives. The log was stored on a 4-disk Raid10 array. The tests were conducted on the qualification database. The steps performed are shown below:

1. The database was backed up to the RAID-10 array.
2. The consistency of the ORDERS and LINEITEM tables were verified.
3. Nine streams of ACID transactions were started. Each stream executed a minimum of 100 transactions.
4. A checkpoint was issued.
5. While the test was running, one of the disks holding database table data was logically removed.
6. A checkpoint was issued to force a failure.
7. The 9 streams of ACID transactions failed and recorded their number of committed transaction in success files.
8. The database log was dumped to disk.
9. A new database drive was attached.
10. A database restore from back up was done.
11. A command was issued causing the database to run through its roll-forward recovery.
12. The success file and the HISTORY table counts were compared and were found to match.
13. The consistency of the ORDERS and LINEITEM tables were verified.

Loss of Log and System Crash Test

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.

1. The consistency of the ORDERS and LINEITEM tables were verified.
2. Nine streams of ACID transactions were started. Each stream executed a minimum of 100 transactions.
3. While the test was running, one of the disks from the database log RAID-10 was physically removed.
4. It was determined that the test would still run with the loss of a log disk, the system was powered off.
5. When the power was restored, the system booted and the log drive was rebuilt.
6. When the drive finished rebuilding, the database was restarted.
7. The database went through a recovery period.
8. The success file and the HISTORY table counts were compared and were found to match.
9. The consistency of the ORDERS and LINEITEM tables were verified.

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."

Clause 4 – Scaling and Database Population Related Items

Initial 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.

Table Name	Row Count
Orders	4,500,000,000
Lineitem	18,000,048,306
Customer	450,000,000
Part	600,000,000
Supplier	30,000,000
Partsupp	2,400,000,000
Nation	25
Region	5

Table 4-1. Initial Cardinality of Tables

Distribution of Tables and Logs

The distribution of tables and logs across all media must be explicitly described.

Database files were spread out on the 6 drives across the 6 Enterprise io3 Flash PCIe adapters. Database log was configured on a Raid-1 2-disk array of 480GB Enterprise value SATA 2.5” SSDs. Tempdb was spread out on the Enterprise io3 Flash drives also. The database and log distribution is shown in the table below.

Storage Distribution

Controller	Drives	RAID	Vol size GB	Format	Files	LUN
Slot 11	2x 200G	Raid-1	184	NTFS	OS, SQL	C:
ServeRaid M5210	2x 480G	Raid-10	446	NTFS	LOG	C:\mount\LOG
Slot 1	1x 3.2TB	Raid 0	2980	NTFS	Data, Tempdb	C:\mount\Fio1 C:\mount\Fio7
Slot 3	1x 3.2TB	Raid 0	2980	NTFS	Data, Tempdb	C:\mount\Fio2 C:\mount\Fio8
Slot 4	1x 3.2TB	Raid 0	2980	NTFS	Data, Tempdb	C:\mount\Fio3 C:\mount\Fio9
Slot 6	1x 3.2TB	Raid 0	2980	NTFS	Data, Tempdb	C:\mount\Fio4 C:\mount\Fio10
Slot 7	1x 3.2TB	Raid 0	2980	NTFS	Data, Tempdb	C:\mount\Fio5 C:\mount\Fio11
Slot 9	1x 3.2TB	Raid 0	2980	NTFS	Data, Tempdb	C:\mount\Fio6 C:\mount\Fio12
Totals	20560		19024			

Database Partition / Replication Mapping

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

The database was not replicated.

RAID Implementation

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.

RAID-10 was used for log disks. RAID-1 was used for the Operating System/Database install disk. The database disks and the temporary tablespace were placed on non-raided drives.

DBGEN Modifications

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 standard distribution DBGEN version 2.17.0 was used for database population. No modifications were made.

Database Load Time

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

The database load time was 02h 58m 00s.

Data Storage Ratio

The data storage ratio must be disclosed. It is computed as the ratio between the total amount of priced disk space and the chosen test database size as defined in Clause 4.1.3.

The calculation of the data storage ratio is shown in the following table.

Disk Type	Number of Disks	Space per Disk (GB)	Total Disk Space (GB)	Scale Factor	Storage Ratio
200GB 2.5" SAS SSD	2	186	372		
480GB 2.5" SATA SSD	2	446	892		
3200GB Enterprise io3 Flash SSD PCIe adapters	6	2960	17,880		
Total			19,024	3000 GB	6.34

The data storage ratio is 6.34, derived by dividing 19,024 GB by the database size of 3000 GB.

Database Load Mechanism Details and Illustration

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

Flat files for each of the tables were created using DBGEN. The tables were loaded as depicted in Figure 4.1.

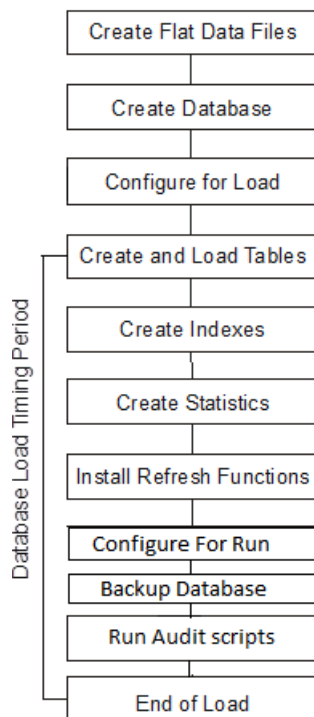


Figure 4-1. Database Load Procedure

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 and disk structure to create and load the data with adjustments for size difference.

Clause 5 – Performance Metrics and Execution Rules Related Items

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 was no activity between the load test and performance test.

Steps in the Power Test

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

The following steps were used to implement the power test:

1. Execute RF1 in refresh stream
2. Execute queries in query stream
3. Execute RF2 in refresh stream

Timing Intervals for Each Query and Refresh Function

The timing intervals for each query of the measured set and for both update functions must be reported for the power test.

See the Numerical Quantities Summary in the Executive Summary at the beginning of this report.

Number of Streams for the Throughput Test

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

Eight query streams and one refresh stream were used for the throughput test.

Start and End Date/Times for Each Query Stream

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

See the Numerical Quantities Summary in the Executive Summary at the beginning of this report.

Total Elapsed Time for the Measurement Interval

The total elapsed time for the measurement interval must be reported for the throughput test.

See the Numerical Quantities Summary in the Executive Summary at the beginning of this report.

Refresh Function Start Date/Time and Finish Date/Time

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

See the Numerical Quantities Summary in the Executive Summary at the beginning of this report.

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

The timing intervals for each query of each stream and for each update function must be reported for the throughput test.

See the Numerical Quantities Summary in the Executive Summary at the beginning of this report.

Performance Metrics

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

See the Numerical Quantities Summary in the Executive Summary at the beginning of this report.

Performance Metric and Numerical Quantities from Both Runs

The performance metric and numerical quantities from both runs must be disclosed.

Two consecutive runs of the TPC-H benchmark were performed. The following table contains the results for both runs.

	QppH @ 3000GB	QthH @ 3000GB	QphH @ 3000GB
Run1	1,322,573.3	716,471.9	973,440.6
Run2	1,346,018.5	698,310.1	969,504.2

System Activity between Tests

Any activity on the SUT that takes place between the conclusion of Run1 and the beginning of Run2 must be disclosed.

There was no activity on the system between Run1 and Run2.

Clause 6 – SUT and Driver Implementation Related Items

Driver

A detailed textual description of how the driver performs its functions, how its various components interact and any product functionality or environmental setting on which it relies must be provided. All related source code, scripts and configurations must be disclosed. The information provided should be sufficient for an independent reconstruction of the driver.

The TPC-H benchmark was implemented using a Microsoft tool called StepMaster. StepMaster is a general purpose test tool which can drive ODBC and shell commands. Within StepMaster, the user designs a workspace corresponding to the sequence of operations (or steps) to be executed. When the workspace is executed, StepMaster records information about the run into a database as well as a log file for later analysis.

StepMaster provides a mechanism for creating parallel streams of execution. This is used in the throughput tests to drive the query and refresh streams. Each step is timed using a millisecond resolution timer. A timestamp T1 is taken before beginning the operation and a timestamp T2 is taken after completing the operation. These times are recorded in a database as well as a log file for later analysis.

Two types of ODBC connections are supported. A dynamic connection is used to execute a single operation and is closed when the operation finishes. A static connection is held open until the run completes and may be used to execute more than one step. A connection (either static or dynamic) can only have one outstanding operation at any time.

In TPC-H, static connections are used for the query streams in the power and throughput tests. StepMaster reads an Access database to determine the sequence of steps to execute. These commands are represented as the Implementation Specific Layer. StepMaster records its execution history, including all timings, in the Access database. Additionally, StepMaster writes a textual log file of execution for each run.

The stream refresh functions were executed using multiple batch scripts. The initial script is invoked by StepMaster, subsequent scripts are called from within the scripts.

The source for StepMaster and the RF Scripts is disclosed in the supported file archive.

Implementation-Specific Layer

If an implementation-specific layer is used, then a detailed description of how it performs its functions must be supplied, including any related source code or scripts. This description should allow an independent reconstruction of the implementation-specific layer.

See Driver section for details.

Profile-Directed Optimization

Profile-directed optimization was not used.

Clause 7 – Pricing Related Items

Hardware and Software Components

A detailed list of the hardware and software used in the priced system must be reported. Each item must have a vendor part number, description and release/revision level, and 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) must also be reported.

A detailed list of all hardware and software, including the 3-year price, is provided in the Executive Summary at the front of this report. The price quotations are included in Appendix A.

Three-Year Cost of System Configuration

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

A detailed list of all hardware and software, including the 3-year price, is provided in the Executive Summary at the front of this report. The price quotations are included in Appendix A.

Availability Dates

The committed delivery date for general availability (availability date) of products used in the price calculations must be reported. When the priced system includes products with different availability dates, 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 (see Clause 7.3.1.3).

The Total System Availability Date is July 31, 2016.

Country-Specific Pricing

Additional Clause 7 related items may be included in the Full Disclosure Report for each country-specific priced configuration. Country-specific pricing is subject to Clause 7.1.7.

The configuration is priced for the United States of America.

Clause 8 – Full Disclosure

8.1 Supporting Files Index Table

An index for all files included in the supporting files archive as required by Clauses 8.3.2 must be provided in the report.

Clause	Description	Pathname
Clause 1	OS and DB settings	SupportingFilesArchive\Clause1
Clause 2	Qualification Queries and Output	SupportingFilesArchive\Clause2
Clause 3	ACID scripts and output	SupportingFilesArchive\Clause3
Clause 4	DB load scripts	SupportingFilesArchive\Clause4
Clause 5	Queries and output for measured runs	SupportingFilesArchive\Clause5
Clause 6	Implementation code for measured runs	SupportingFilesArchive\Clause6
Clause 8	RFs source and params	SupportingFilesArchive\Clause8

Clause 9 – Audit Related Items

Auditor

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, Inc. Further information regarding the audit process may be obtained from:

InfoSizing, Inc.

20 Kreg Ln.

Manitou Springs, CO 80829

Telephone: (719) 473-7555

Web address: www.sizing.com

For a copy of this disclosure, go to www.tpc.org.

Attestation Letter

The auditor's Attestation Letter is on the next two pages.

Benchmark sponsor: Vinay Kulkarni
Enterprise Business Group
Lenovo Corporation
3600 Carillon Point
Kirkland, WA 98033

March 9, 2016

I verified the TPC Benchmark H (TPC-H™ v2.17.1) performance of the following configuration:

Platform: Lenovo® System x®3850 X6
Operating System: Microsoft Windows Server 2012 R2 Standard Edition
Database Manager: Microsoft SQL Server 2016 Enterprise Edition
Other Software: n/a

The results were:

Performance Metric 969,504.2 QphH@3,000GB

TPC-H Power 1,346,018.5
TPC-H Throughput 698,310.1
Database Load Time 02h 58m 00s

Server Lenovo System x3950 X6

CPU	4 x Intel Xeon Processor E7-8890 v3 (2.5GHz, 45MB L3)		
Memory	3,072 GB		
Disks	Qty	Size	Type
	2	200 GB	SAS SSD
	2	480 GB	SATA 2.5" SSD
	6	3,200 GB	Enterprise Io3 Flash PCIe

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 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 and no query variant
- The execution of the queries against the SF1 database produced compliant answers
- A compliant implementation specific layer was used to drive the tests
- The throughput tests involved 8 query streams
- The ratio between the longest and the shortest query was such that no query timings were adjusted
- The execution times for queries and refresh functions were correctly measured and reported
- The repeatability of the measured results was verified
- The system pricing was verified for major components and maintenance
- The major pages from the FDR were verified for accuracy

Additional Audit Notes:

None.

Respectfully Yours,

A handwritten signature in black ink, appearing to read "François Raab", with a long horizontal flourish extending to the right.

François Raab, President

Microsoft Corporation
One Microsoft Way
Redmond, WA 98052-6399

Tel 425 882 8080
Fax 425 936 7329
<http://www.microsoft.com/>

Microsoft

March 9, 2016

Lenovo
Vinay Kulkarni
3600 Carillon Point
Kirkland, WA 98033

Here is the information you requested regarding pricing for several Microsoft products to be used in conjunction with your TPC-H benchmark testing.

All pricing shown is in US Dollars (\$).

Part Number	Description	Unit Price	Quantity	Price
Database Management System				
7JQ-00256	SQL Server 2016 Enterprise Edition <i>2 Core License Open Program - Level C</i>	\$13,472.50	36	\$485,010.00
Database Server Operating System				
P73-06284	Windows Server 2012 R2 Standard Edition <i>2 Processor License Open Program - Level C Unit Price reflects a 17% discount from the retail unit price of \$882.</i>	\$735.00	2	\$1,470.00
Support				
N/A	Microsoft Problem Resolution Services <i>Professional Support (1 Incident).</i>	\$259.00	1	\$259.00

Windows Server 2012 R2 Standard Edition are currently orderable and available through Microsoft's normal distribution channels. A list of Microsoft's resellers can be found in the Microsoft Product Information Center at <http://www.microsoft.com/products/info/render.aspx?view=22&type=how>

SQL Server 2016 Enterprise Edition will be orderable on March 10, 2016 and generally available by July 31, 2016.

Defect support is included in the purchase price. Additional support is available from Microsoft PSS on an incident by incident basis at \$259 call.

This quote is valid for the next 90 days.

Reference ID: TPCH_qhtplyIGYLKTVUKf77456jshr_2016_lvk.