

TPC BenchmarkTM H
Full Disclosure Report
for
IBM[®] System x[®]3850 X5
using
Microsoft[®] SQL Server 2008
R2 Enterprise x64 Edition
and
Microsoft Windows[®] Server 2008
R2 Enterprise x64 Edition

TPC-H Version 2.13.0

Submitted for Review
April 5, 2011

IBM Corporation

First Edition – April 2011

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 IBM 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 an IBM licensed program are not intended to state or imply that only IBM's licensed program may be used; any functionally equivalent program may be used.

This publication was produced in the United States. IBM 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 IBM representative for information on products and services available in your area.

© Copyright International Business Machines Corporation 2011. 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.

U.S. Government Users - Documentation related to restricted rights: Use, duplication, or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Trademarks

IBM, the IBM logo, System x and System Storage are trademarks or registered trademarks of International Business Machines Corporation.

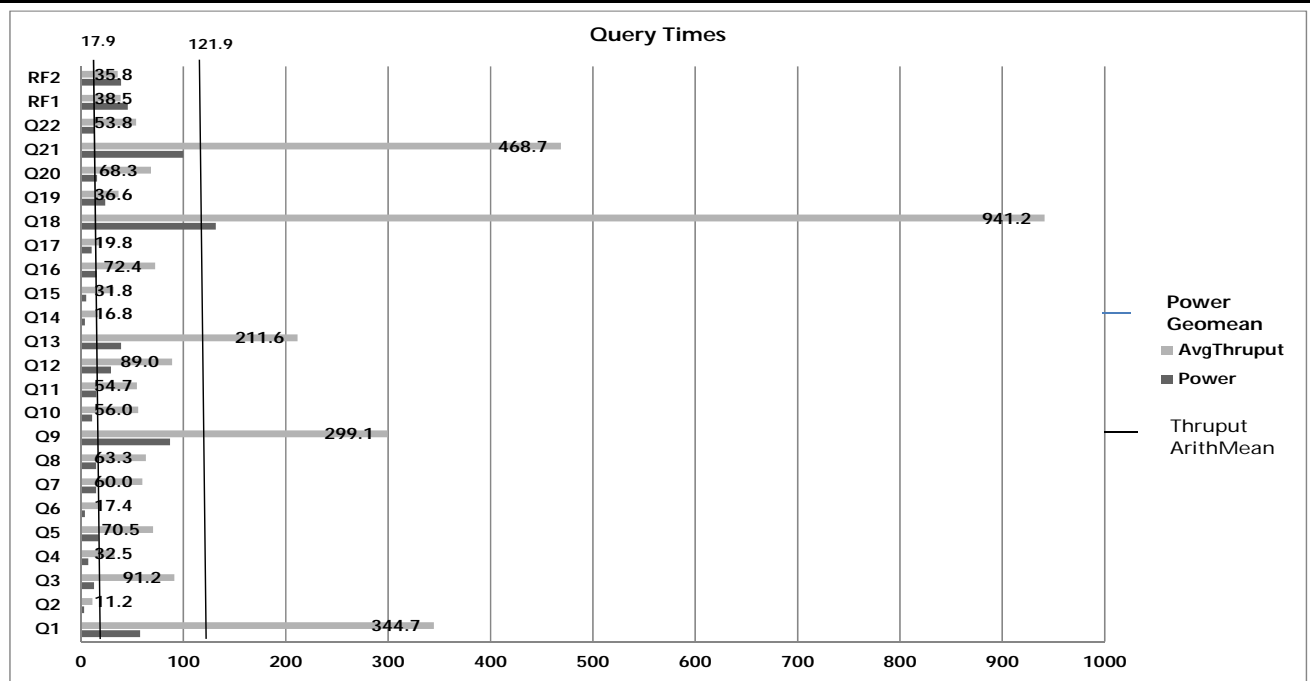
The following terms used in this publication are trademarks of other companies as follows: TPC Benchmark, TPC-H, QppH QthH and QphH are trademarks of Transaction Processing Performance Council; Intel and Xeon are trademarks or registered trademarks of Intel Corporation; Microsoft and Windows are trademarks 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.

IBM Corporation	IBM® System x®3850 X5 8P Microsoft® SQL Server 2008 R2			TPC-H Rev. 2.13.0
				Report Date 4/5/11
Total System Cost	Composite Query-per-Hour Metric			Price/Performance
\$238,375 USD	173,961.8 QphH @ 1000GB			\$1.37 per QphH @ 1000GB
Database Size	Database Manager	Operating System	Other software	Availability
1000GB	Microsoft SQL Server 2008 R2 Enterprise x64 Edition	Microsoft Windows Server R2 Enterprise x64 Edition	n.a.	May 20, 2011



Database Load Time: 14:44:52		Total Data Storage / Database Size: 7.0	Load Included Backup: Y	Memory Ratio=200%
RAID (Base Tables Only): N		RAID (Base Tables and Auxiliary Data Structures): N	RAID (All): N	
Configuration				
Processors/Cores/Threads	8/80/80	Intel Xeon Processor E7-8870, 2.4GHz, 30MB L3 Cache		
Memory	128	16GB (1x16GB) 4R PC3-8500 1066MHz LP RDIMM		
Disk Controllers	2	IBM ServeRAID-M1015 SAS/SATA Controller		
Disk Drives	2	146GB 15K 6Gbps SAS 2.5" SFF Slim-HS HDD		
	14	300GB 10K 6Gbps SAS 2.5" SFF Slim-HS HDD		
	1	640GB High IOPS MLC Duo Adapter for IBM System x		
	6	320GB High IOPS MS Class SSD PCIe Adapter		
Total Disk Storage		7052GB		

**IBM
Corporation**

**IBM® System x®3850 X5 8P
Microsoft® SQL Server 2008 R2**

**TPC-H
Rev. 2.13.0**

Report Date 4/5/11

IBM Corporation	IBM System x3850 X5				TPC-H Revision 2.13.0			
	Microsoft SQL Server 2008 R2				Report Date: April 5, 2011			
Description	Part Number	Brand	Price Source	Unit Price	Quantity	Extended Price	3-Yr. Maint. Price	
Server Hardware								
x3850 X5 (2 x Intel Xeon Processor E7-8850, 4x4GB Memory, 2 Memory Cards)	7143-C1Y	IBM	1a*	\$16,519				
Remove 2 Intel Xeon E7-8850 processors				(8,790)				
Subtotal								
Add 2 Intel Xeon E7-8870 processors				14,550				
Total				22,279	2	44,558		
Intel Xeon Processor E7-8870 (2.4GHz / 30MB L3)	69Y1899	IBM	1a*	7,275	4	29,100		
IBM x3850 X5 QPI Scalability Kit	46M0072	IBM	1a	1,795	1	1,795		
16GB (1x16GB) QuadRank PC3-8500 1066MHz LP RDIMM	46C7483	IBM	1	999	128	127,872		
IBM x3850 X5 Memory Expansion Card	46M0071	IBM	1	299	12	3,588		
IBM ServeRAID-M1015 SAS/SATA Controller	46M0831	IBM	1	299	2	598		
IBM 146GB 15K 6Gbps SAS 2.5" SFF Slim-HS HDD	42D0677	IBM	1	399	2	798		
IBM 300GB 10K 6Gbps SAS 2.5" SFF Slim-HS HDD	42D0637	IBM	1	449	14	6,286		
640GB High IOPS MLC Duo Adapter for IBM System x	81Y4519	IBM	1	14,999	1	14,999		
IBM 320GB High IOPS MS Class SSD PCIe Adapter	46M0898	IBM	1	8,099	6	48,594		
IBM 3-Button Optical Mouse - Black - USB	40K9201	IBM	1a	19	1	19		
IBM Preferred Pro USB Keyboard	40K9584	IBM	1	29	1	29		
ServicePac for 3-Year 24x7x4 Support (x3850 X5)	84Y2233	IBM	1	1,250	2		2,500	
Acer V173Djb Black 17" 5ms LCD Monitor (2 spares)	V173Djb		3	97	3	291		
					Subtotal	278,527	2,500	
Server Software								
Microsoft SQL Server 2008 R2 Enterprise Edition w/ 25 CALs	810-08553	Microsoft	2	\$8,318	1	8,318		
Microsoft SQL Server 2008 R2 Client Access License	359-05354	Microsoft	2	\$114	45	5,130		
Microsoft Windows Server 2008 R2 Enterprise Edition	P72-04217	Microsoft	2	2,280	1	2,280		
Microsoft Problem Resolution Services	N/A	Microsoft	2	259	1		259	
					Subtotal	15,728	259	
					Total	294,255	2,759	
Dollar Volume Discount (See Note 1.)	28.57%		1			58,639		
Pricing: 1 - IBM - 1-800-656-0833, ext. 35334; 2 - Microsoft; 3 - newegg.com					Three-Year Cost of Ownership USD		\$238,375	
Note 1: Discount applies to all line items where Pricing=1; pricing is for these or similar quantities.					QpH@1000GB:		173,961.80	
Discounts for similarly sized configurations will be similar to what is quoted here, but may vary based on the components in the price quotation					\$ USD/QpH@1000GE		\$1.37	
* These components are not immediately orderable. See the FDR for more information.								
Benchmark results and methodology audited by Francois Raab, 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 stated prices are not available according to these terms, please inform the TPC at pricing@tpc.org. Thank you.								

IBM Corporation	IBM® System x®3850 X5 8P Microsoft® SQL Server 2008 R2	TPC-H Rev. 2.13.0
		Report Date 4/5/11

Measurement Results

Database Scale Factor	1000
Total Data Storage/Database Size	7.0
Start of Database Load	03/21/11 00:46:47
End of Database Load	03/21/11 15:31:39
Database Load Time	14:44:52
Query Streams for Throughput Test	7
TPC-H Power	200,899.9
TPC-H Throughput	150,635.8
TPC-H Composite Query-per-Hour (QphH@1000GB)	173,961.8
Total System Price over 3 Years	\$238,375 USD
TPC-H Price/Performance Metric (\$/QphH@1000GB)	\$ 1.37 USD

Measurement Interval

Measurement Interval in Throughput Test (Ts) = 3680.8

Duration of Stream Execution:

	Seed	Query Start (D/T) Query End (D/T)	RF1 Start (D/T) RF1 Stop (D/T)	RF2 Start (D/T) RF2 Stop (D/T)	Duration
Stream 0	321153138	3/21/11 15:35:50 3/21/11 15:46:20	3/21/11 15:35:04 3/21/11 15:35:50	3/21/11 15:46:20 3/21/11 15:46:59	00:10:30
Stream 1	321153139	3/21/11 15:47:00 3/21/11 16:39:31	3/21/11 16:39:34 3/21/11 16:40:12	3/21/11 16:40:12 3/21/11 16:40:50	00:52:30
Stream 2	321153140	3/21/11 15:47:00 3/21/11 16:39:33	3/21/11 16:40:50 3/21/11 16:41:29	3/21/11 16:41:29 3/21/11 16:42:06	00:52:33
Stream 3	321153141	3/21/11 15:47:00 3/21/11 16:39:34	3/21/11 16:42:06 3/21/11 16:42:45	3/21/11 16:42:45 3/21/11 16:43:21	00:52:33
Stream 4	321153142	3/21/11 15:47:00 3/21/11 16:38:33	3/21/11 16:43:22 3/21/11 16:44:01	3/21/11 16:44:01 3/21/11 16:44:36	00:51:32
Stream 5	321153143	3/21/11 15:47:00 3/21/11 16:39:26	3/21/11 16:44:37 3/21/11 16:45:15	3/21/11 16:45:15 3/21/11 16:45:52	00:52:26
Stream 6	321153144	3/21/11 15:47:00 3/21/11 16:38:24	3/21/11 16:45:52 3/21/11 16:46:31	3/21/11 16:46:31 3/21/11 16:47:06	00:51:23
Stream 7	321153145	3/21/11 15:47:00 3/21/11 16:35:30	3/21/11 16:47:07 3/21/11 16:47:46	3/21/11 16:47:46 3/21/11 16:48:20	00:49:53

IBM Corporation	IBM® System x®3850 X5 4P Microsoft® SQL Server 2008 R2										TPC-H Rev. 2.13.0	
	Report Date 4/5/11											
TPC-H Timing Intervals (in seconds)												
Query	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11	Q12
Stream 00	57.8	2.9	12.8	7.1	17.1	3.9	14.7	14.7	86.9	10.8	15.0	29.2
Stream 01	351.9	12.4	85.9	24.2	60.1	16.1	52.3	61.1	313.9	47.0	54.6	113.1
Stream 02	347.4	9.8	102.9	34.6	69.2	18.4	48.4	61.1	347.4	58.8	83.0	115.8
Stream 03	369.3	11.1	63.1	34.8	77.8	18.2	70.3	72.9	270.7	55.6	64.0	117.6
Stream 04	365.2	10.5	101.9	35.8	75.6	16.5	67.2	50.9	290.3	52.8	39.5	63.7
Stream 05	366.4	11.8	88.5	35.8	69.2	16.2	60.3	68.2	290.1	57.4	51.3	61.5
Stream 06	327.0	13.4	107.0	33.5	65.8	19.3	58.2	72.5	304.8	67.7	38.3	67.2
Stream 07	285.8	9.4	89.1	28.5	75.5	17.1	63.0	56.4	276.8	52.5	52.1	84.2
Minimum	57.8	2.9	12.8	7.1	17.1	3.9	14.7	14.7	86.9	10.8	15.0	29.2
Maximum	369.3	13.4	107.0	35.8	77.8	19.3	70.3	72.9	347.4	67.7	83.0	117.6
Average	308.9	10.2	81.4	29.3	63.8	15.7	54.3	57.2	272.6	50.3	49.7	81.5
Query	Q13	Q14	Q15	Q16	Q17	Q18	Q19	Q20	Q21	Q22	RF1	RF2
Stream 00	39.1	3.9	5.0	15.1	10.4	131.7	23.5	15.8	99.9	12.8	45.6	39.0
Stream 01	189.7	18.2	28.2	68.5	14.7	947.8	36.5	67.2	527.8	59.6	37.7	37.2
Stream 02	208.3	19.6	29.8	79.0	22.6	931.6	37.6	77.2	401.3	49.7	38.4	36.9
Stream 03	218.3	12.8	28.7	76.1	17.7	972.3	42.5	70.5	449.2	40.0	38.6	36.1
Stream 04	232.7	16.9	33.0	64.8	22.1	886.9	34.8	63.5	499.7	68.4	38.9	35.2
Stream 05	201.8	16.8	35.2	71.1	21.6	933.9	33.1	74.7	523.4	57.6	38.5	36.3
Stream 06	228.9	17.4	35.4	67.5	20.7	894.0	37.0	72.4	489.5	46.3	38.6	35.1
Stream 07	201.2	16.2	32.2	80.0	19.3	1022.1	34.5	52.6	390.1	55.2	38.9	34.1
Minimum	39.1	3.9	5.0	15.1	10.4	131.7	23.5	15.8	99.9	12.8	37.7	34.1
Maximum	232.7	19.6	35.4	80.0	22.6	1022.1	42.5	77.2	527.8	68.4	45.6	39.0
Average	190.0	15.2	28.4	65.3	18.6	840.0	34.9	61.7	422.6	48.7	39.4	36.2

Table of Contents

Preface	9
General Items	10
Benchmark Sponsor.....	10
Parameter Settings	10
Configuration Diagrams	10
Clause 1 – Logical Database Design Related Items	12
Database Table Definitions.....	12
Database Physical Organization	12
Horizontal/Vertical Partitioning	12
Replication.....	12
Clause 2 – Queries and Update Functions Related Items	13
Query Language	13
Random Number Generation.....	13
Substitution Parameters Generation.....	13
Query Text and Output Data from Database	13
Query Substitution Parameters and Seeds Used	13
Query Isolation Level	13
Refresh Function Implementation	14
Clause 3 – Database System Properties Related Items	15
Atomicity Requirements.....	15
Consistency Requirements	15
Isolation Requirements	17
Durability Requirements.....	18
Clause 4 – Scaling and Database Population Related Items	20
Initial Cardinality of Tables.....	20
Distribution of Tables and Logs	20
Database Partition / Replication Mapping	20
RAID Implementation	20
DBGEN Modifications	20
Database Load Time	21
Data Storage Ratio.....	21
Database Load Mechanism Details and Illustration	21
Qualification Database Configuration	22
Clause 5 – Performance Metrics and Execution Rules Related Items	23
System Activity between Load and Performance Tests.....	23
Steps in the Power Test	23
Timing Intervals for Each Query and Refresh Function.....	23
Number of Streams for the Throughput Test.....	23
Start and End Date/Times for Each Query Stream	23
Total Elapsed Time for the Measurement Interval	23
Refresh Function Start Date/Time and Finish Date/Time	23
Timing Intervals for Each Query and Each Refresh Function for Each Stream	24
Performance Metrics.....	24
Performance Metric and Numerical Quantities from Both Runs.....	24
System Activity between Tests.....	24
Clause 6 – SUT and Driver Implementation Related Items	25
Driver.....	25
Implementation-Specific Layer	25
Profile-Directed Optimization	25
Clause 7 – Pricing Related Items	26
Hardware and Software Components	26
Three-Year Cost of System Configuration	26
Availability Dates	26

Country-Specific Pricing	26
Clause 8 – Full Disclosure.....	27
8.1 Supporting Files Index Table.....	27
Clause 9 – Audit Related Items	28
Auditor.....	28
Attestation Letter	28
Appendix A: 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 2.7.0). This is the full disclosure report for benchmark testing of the IBM System x3850 X5 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:

- v Examine large volumes of data;
- v Execute queries with a high degree of complexity;
- v 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:

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

The TPC-H operations are modeled as follows:

- v 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.
- v 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.
- v 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.
- v 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 metrics 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.

Benchmark 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 IBM 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 Files, “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



IBM System 8S x3850 X5 with QPI scaling

The measured and priced configurations are the same. For the priced configuration, see the Executive Summary.

The priced configuration for the x3850 X5 contained:

- Eight Intel Xeon Processor E7-8870 (2.40GHz, 2MB L2 cache and 30MB L3 cache)
- One hundred and twenty-eight 16GB (1x16GB) QuadRank PC3-8500 1066MHz LP RDIMM
- One embedded dual-port 10/100/1000 Gigabit Ethernet interface
- One IBM ServeRAID-M1015 SAS/SATA Controller
- One 640GB High IOPS MLC Duo Adapter for IBM System x
- Six IBM 320GB High IOPS MS Class SSD PCIe Adapters
- Two IBM 146GB 15K 6Gbps SAS 2.5" SFF Slim-HS HDDs
- Fourteen IBM 300GB 10K 6Gbps SAS 2.5" SFF Slim-HS HDDs

The measured configuration and the priced configuration were identical.

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.

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 was not used.

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 2.13.0 and QGEN version 2.13.0 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 2.13.0 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:

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

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 isolation level 1.

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 IBM System x3850 X5 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.

The following steps were performed to verify consistency:

1. The consistency of the ORDER and LINEITEM tables was verified based on a sample of O_ORDERKEYS.
2. One hundred ACID Transactions were submitted from each of eight execution streams.
3. The consistency of the ORDER and LINEITEM tables was reverified.

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 14 non-raided IBM high IOPs drives. The log was stored on a 6-disk Raid1 array. A backup of the database was taken and stored on the Raid1 array of 6 physical drives. The tests were conducted on the qualification database. The steps performed are shown below:

1. The complete database was backed up to the RAID-1 array.
2. The consistency of the ORDERS and LINEITEM tables were verified.
3. Eight 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 from the database was removed.
6. A checkpoint was issued to force a failure.
7. The 8 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. Eight 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 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
Order	1,500,000,000
Lineitem	5,999,989,709
Customer	150,000,000
Part	200,000,000
Supplier	10,000,000
Partsupp	800,000,000
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 Partition / Replication Mapping

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

Database files were spread out on the 14 drives split across the 7 IBM high IOPS PCIe adapters. Database log was configured on a Raid-10 8-disk arrays of 300GB SAS 2.5” drives. Tempdb was spread out on the IBM high IOPs drives also.

The database was not replicated or partitioned.

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 2.13.0 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 14 hours 44 minutes and 52 seconds.

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	Total Disk Space	Scale Factor	Storage Ratio
146GB 15K 2.5 inch SAS SFF H/S Drive	2	146GB	292GB		
300GB 10K 2.5 inch SAS H/S Drive	14	300GB	4.2TB		
320GB IBM High IOPs MS class SSD PCIe adapters	6	320GB	1.92TB		
640GB IBM high IOPs MLC DUO adapter for IBM System x	1	640GB	0.64TB		
Total			7.0TB	1000GB	7.0

The data storage ratio is 7.0, derived by dividing 7.0TB by the database size of 1000GB.

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.

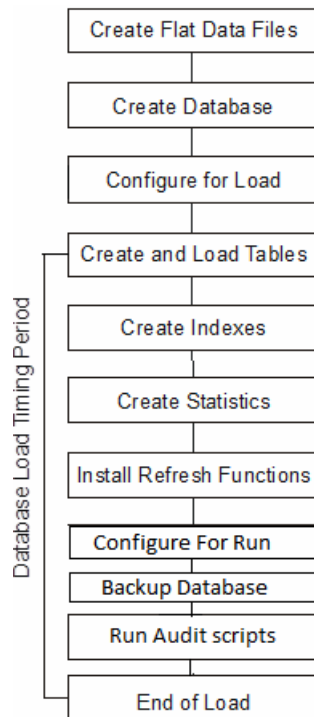


Figure 4-1. Database Load Procedure

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. See Section 5.2 for details.

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.

The auditor requested queries were run against the database to verify the correctness of the database load.

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. RF1 Refresh Transaction
2. Stream 00 Execution
3. RF2 Refresh Transaction

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.

Seven streams 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 @ 1000GB	QthH @ 1000GB	QphH @ 1000GB
Run1	200,899.9	150,635.8	173,961.8
Run2	213,075.8	149,959.4	178,753.2

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.

Two scripts were used. The first one was used to create and load the database; the second was used to run the Power and Throughput tests. These scripts are listed in the Supporting Files. The C program, semaphore.c, was used for coordination of parallel processes.

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.

A command script was used to control and track the execution of queries. The scripts are contained in the Supporting Files. Qgen was used to generate the query streams, along with the appropriate substitution values.

The following steps are performed, to accomplish the Power and Throughput Runs:

Power Run

1. Execute 48 concurrent RF1 processes, each of which will apply a segment of an update set generated by dbgen.
2. Each process submits multiple transactions, where a transaction spans a set of orders and their associated line items. The first stream waits on a semaphore prior to beginning its execution. RF streams wait on a semaphore before beginning execution of RF2.
3. Execute the Stream0 queries, in the prescribed order. Signal the waiting RF2 semaphore to continue with RF stream execution and run RF2.
4. Execute 48 concurrent RF2 processes, each of which will apply a segment of an update set generated by dbgen. Each thread submits multiple transactions, where a transaction spans a set of orders and their associated line items.

Throughput Run

1. Execute seven concurrent query streams. Each stream executes queries in the prescribed order for the appropriate Stream Id (1-5). Upon completion of each stream, a semaphore is set to indicate completion.
2. Execute five consecutive RF1/RF2 transactions, against ascending Update sets produced by dbgen. The first RF1 waits on a semaphore prior to beginning its insert operations.

Each step is timed by the script. The timing information is stored in the database for later analysis. The inputs and outputs of steps are stored in text files for later analysis.

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.

The IBM hardware used in this configuration qualified for a 28.57% discount based on the dollar-volume.

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 May 20, 2011.

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

InfoSizing

125 West Monroe Street

Colorado Springs, CO 80907

Telephone: (719) 473-7555

Fax: (719) 473-7554

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
 IBM System x Performance Engineer
 3600 Carillon Point
 Kirkland, WA 98033

April 4, 2011

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

Platform: **IBM System x3850 X5 8-Socket Server**
 Database Manager: **Microsoft SQL Server 2008 R2 Enterprise x64 Edition**
 Operating System: **Microsoft Windows Server 2008 R2 Enterprise x64 Edition**

The results were:

CPU (Speed)	Memory	Storage	QphH@1000GB
IBM System x3850 X5 8-Socket Server			
8 x Intel E7-8870 (2.4GHz)	2TB	14 x 300GB 10Krpm 2 x 146GB 15Krpm 6 x 320GB SSD 1 x 640GB SSD	173,961.8

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 1,000GB 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 7 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

Appendix A: Price Quotes

Microsoft Corporation
One Microsoft Way
Redmond, WA 98052-6399

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

Microsoft

March 29, 2011

IBM Corporation
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
810-08553	SQL Server 2008 R2 Enterprise Edition <i>Server License with 25 CALs</i> <i>Open Program - Level C</i> <i>Unit Price reflects a 23% discount from the retail unit price of \$3,999.</i>	\$8,318	1	\$8,318
359-05354	SQL Server Client Access License <i>Open Program - Level C</i> <i>Unit Price reflects a 30% discount from the retail unit price of \$114.</i>	\$114	45	\$5,130
P72-04217	Windows Server 2008 R2 Enterprise Edition <i>Server License with 25 CALs</i> <i>Open Program - Level C</i> <i>Unit Price reflects a 43% discount from the retail unit price of \$3,999.</i>	\$2,280	1	\$2,280
N/A	Microsoft Problem Resolution Services <i>Professional Support</i> <i>(1 Incident).</i>	\$259	1	\$259

SQL Server 2008 R2 Enterprise Edition and Windows Server 2008 R2 Enterprise 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>

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: TPC_H_qhtplyIGYLKTVUKfhkjQjhilimkJipf85757.

Skip to: [Content](#) | [Footer](#)

Shopping Cart

Print

Qty.	Product Description	Savings	Total Price
1	 Acer V173DJb Black 17" 5ms LCD Monitor Item #: N82E16824009230 Return Policy: Monitor Standard Return Policy	-\$43.00 Instant	\$139.99 \$96.99
Subtotal:			\$96.99
Calculate Shipping Zip Code: <input type="text" value="UPS Guaranteed 3 Day Service"/>		Shipping:	\$0.00
Redeem Newegg Gift Cards Card Number: <input type="text"/> Security Code: <input type="text"/>			
Apply Promo Code(s):		Promo Code:	\$0.00
New Newegg Preferred Account Customers: Get \$15 off \$100! Plus, No Interest if paid in full in up to 12 Months. Minimum purchase required. Subject to credit approval. Learn More			
Grand Total:*			\$96.99

* Above total does not include shipping or taxes. Please input zip code to calculate your grand total.

Print

[Policy & Agreement](#) | [Privacy Policy](#) © 2000-2011 Newegg Inc. All rights reserved.