TPC Benchmark™ E
Full Disclosure Report
for
IBM® System x® 3850 X5
using
Microsoft® SQL Server 2008 R2
Enterprise Edition
and
Microsoft Windows® Server 2008 R2
Enterprise Edition

TPC-E Version 1.9.0

Second Edition
Submitted for Review
June 7, 2010

IBM Corporation

Second Edition – June 2010

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-E, tpsE, and $/tpsE trademark 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.

Abstract

IBM Corporation conducted the TPC Benchmark™ E on the IBM® System x®3850 X5 configured as a client/server system. This report documents the full disclosure information required by the TPC Benchmark E Standard Specification, Revision 1.9.0, including the methodology used to achieve the reported results. All testing fully complied with this revision level.

The software used on the IBM System x3850 X5 system included Microsoft® Windows® Server 2008 R2 Enterprise Edition and Microsoft SQL Server 2008 R2 Enterprise Edition.

Standard metrics, transactions per second-E (tpsE), price per tpsE ($/tpsE) and Availability Date, are reported as required by the TPC Benchmark E Standard Specification.

The benchmark results are summarized in the following table:

Hardware	Software	Total System Cost	tpsE	$ USD /tpsE	Total Solution Availability Date
IBM System x3850 X5	Microsoft SQL Server 2008 R2 Enterprise Edition Microsoft Windows Server 2008 R2 Enterprise Edition	$718,065 USD	2022.64	$355.02 USD	July 30, 2010

The benchmark implementation and results were audited by Doug Johnson for InfoSizing (www.sizing.com). The auditor's attestation letter is contained in this report.

	IBM® System x® 3850 X5 **Microsoft® SQL Server 2008 R2**	**TPC-E 1.9.0** **TPC Pricing 1.5.0** Report Date: March 30, 2010 Revision Date: June 7, 2010

TPC-E Throughput	Price/Performance	Availability Date	Total System Cost
2022.64 tpsE	**$355.02 USD per tpsE**	**July 30, 2010**	**$718,065 USD**

Database Server Configuration

Operating System	Database Manager	Processors/Cores/ Threads	Memory
Microsoft Windows Server 2008 R2 Enterprise Edition	**Microsoft SQL Server 2008 R2 Enterprise Edition**	**4/32/64**	**1024GB**

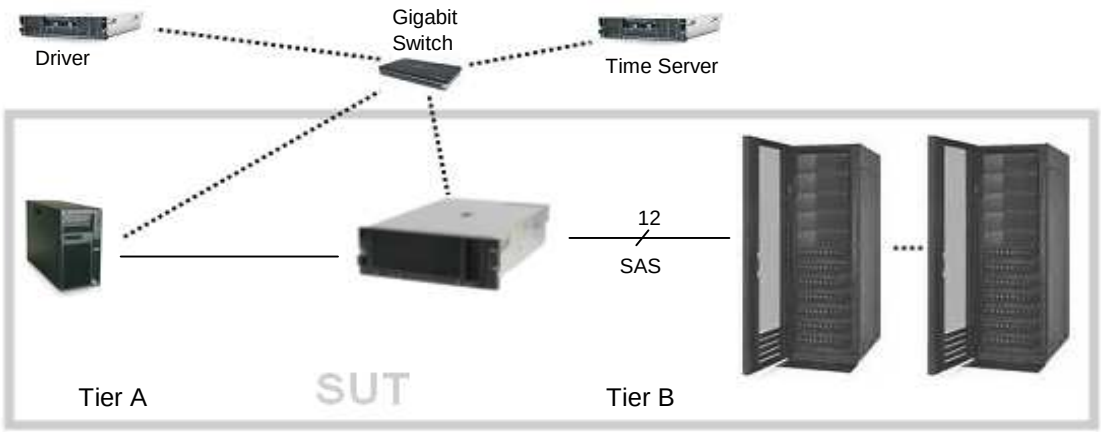

IBM x3500 M2
- 2 x Intel Xeon Processor X5570 2.93GHz (2 Procs/8 Cores/16 Threads)
- 8GB Memory
- 2 x 160GB SFF SATA (RAID-1)
- 1 x IBM ServeRAID-BR10i
- Onboard Gigabit Ethernet
- 1 x 10Gb Ethernet

IBM System x3850 X5
- 4 x Intel Xeon Processor X7560 2.26GHz (4 Procs/32 Cores/64 Threads)
- 1TB Memory
- 2 x 160GB SFF SATA (RAID-1)
- 6 x 300GB SFF SAS (RAID-10)
- 1 x IBM ServeRAID-M5015
- 6 x IBM ServeRAID-M5025
- Onboard Gigabit Ethernet
- 1 x 10Gb Ethernet

84 x IBM System Storage EXP3000 Enclosure
Four contain:
- 12 x 300GB 3.5" 15K SAS

80 contain:
- 12 x 146GB 3.5" 15K SAS

= 42 x 24-drive RAID-10
1008 Total External Drives

Initial Database Size	Redundancy Level: **1**	Storage
8,512 GB	**RAID-10 Log & Data**	**2 x 160 GB** **960 x 146GB** **54 x 300 GB**

IBM System x3850 X5
Microsoft SQL Server 2008 R2

TPC-E 1.9.0

TPC Pricing 1.5.0

Report Date:
March 30, 2010

Revision Date:
June 7, 2010

Availability Date:
July 30, 2010

Description	Part Number	Price Source	Unit Price	Quantity	Extended Price	3-Yr. Maint. Price
Server Hardware						
x3850 X5 with 2 x Intel Xeon Processor X7560 (2.26GHz / 2MB L2 Cache / 24MB L3 Cache) 4 x 4GB Memory, 2 memory cards	71455RU	1	21,705	1	21,705	
Intel Xeon Processor X7560 (2.26GHz/2MB L2/24MB L3)	49Y4300	1	6,455	2	12,910	
16GB (1x16GB) QuadRank PC3-8500 1066MHz LP RDIMM	46C7483	1	1,549	64	99,136	
IBM x3850X5 Memory Expansion Card	46M0071	1	299	6	1,794	
IBM ServeRAID-M5015 SAS/SATA Controller	46M0829	1	649	1	649	
IBM 160GB 7200 NL SATA 2.5" SFF Slim-HS HDD	42D0747	1	259	2	518	
IBM 300GB 10K 6Gbps SAS 2.5" SFF Slim-HS HDD	42D0637	1	559	6	3,354	
IBM ServeRAID-M5025 SAS/SATA Controller	46M0830	1*	649	6	3,894	
IBM Preferred Pro USB Keyboard	40K9584	1	29	1	29	
IBM 3-Button Optical Mouse - Black - USB	40K9201	1	19	1	19	
Single-Port SFP+ 10GbE PCI-E x8 Adapter	46M1809	1	1,499	1	1,499	
ServicePac for 3-Year 24x7x4 Support (x3850 X5)	10N3059	1	1,695	1		1,695
Acer V173 Bb Black 17" 5ms LCD Monitor (2 spares)	V173Bb	3	120	3	360	
				Subtotal	145,867	1,695
Server Storage						
IBM S2 42U Standard Rack	93074RX	1	1,609	5	8,045	
IBM System Storage EXP3000 Enclosure	1727-01X	1	3,199	84	268,716	
IBM 1M SAS cable	39R6529	1	119	72	8,568	
IBM 3M SAS cable	39R6531	1	135	12	1,620	
IBM 300GB 15K 6Gbps SAS 3.5" Hot-Swap HDD	44W2234	1	599	48	28,752	
146GB 15K 3.5" Hot-Swap SAS	40K1044	1 - S	359	960	344,640	
ServicePac for 3-Year 24x7x4 Support (EXP3000)	41L2768	1	760	84		63,840
ServicePac for 3-Year 24x7x4 Support (Rack)	41L2760	1	300	5		1,500
				Subtotal	660,341	65,340
Server Software						
Microsoft SQL Server 2008 R2 Enterprise x64 Edition		2	28,749	4	114,996	
Microsoft Windows Server 2008 R2 Enterprise x64 Edition	P73-04217	2	3,999	1	3,999	
Microsoft Problem Resolution Services	N/A	2a	259	1		259
				Subtotal	118,995	259
Client Hardware						
x3500 M2 with 1 x Intel Xeon Processor X5570 (2.93GHz / 1MB L2 Cache / 8MB L3 Cache) 2x 1GB Memory	783982U	1	3,845	1	3,845	
Intel Xeon Processor X5570 (2.93GHz/1MB L2/8MB L3 Cache)	46D1357	1	2,135	1	2,135	
2GB (1x2GB) DualRank PC3-10600 DDR3-1333 LP RDIMM	44T1481	1	125	4	500	
ServeRAID-BR10i SAS/SATA Controller	44E8689	1	229	1	229	
IBM 160GB 7200 NL SATA 2.5" SFF Slim-HS HDD	42D0747	1	259	2	518	
Single-Port SFP+ 10GbE PCI-E x8 Adapter	46M1809	1	1,499	1		1,499
ServicePac for 3-Year 24x7x4 Support (x3500 M2)	21P2078	1	600	1		600
				Subtotal	7,227	2,099
Client Software						
Microsoft Windows Server 2008 R2 Standard x64 Edition	P73-04980	2	1,029	1	1,029	
				Subtotal	1,029	0
Infrastructure						
1M Fibre Optic Cable LC-LC	39M5696	1	79	1	79	
				Subtotal	79	0
				Total	933,538	69,393
Dollar Volume Discount (See Note 1)	27.72%	1			244,585	
Microsoft Open Program Discount Schedule	33.49%	2			40,281	

Pricing: 1 - IBM - 1-800-656-0833, x35330; 2 - Microsoft; 3 - newegg.com

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 components in the price quotation

S: One or more components of the measured configuration have been substituted in the priced configuration. See the FDR for details.

* These components are not immediately orderable. See the FDR for more information.

Three-Year Cost of Ownership USD:	$718,065
TPC-E Throughput:	2,022.64
$ USD/tpsE:	$355.02

Benchmark results and test methodology audited by Doug Johnson 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.

IBM System x3850 X5
Microsoft SQL Server 2008 R2

TPC-E 1.9.0
TPC Pricing 1.5.0

Report Date:
March 30, 2010

Revision Date:
June 7, 2010

Availability Date:
July 30, 2010

Numerical Quantities Summary				
Reported Throughput: 2022.64 tpsE	**Configured Customers: 1,050,000**			
Response Time (in seconds)	**Minimum**	**Average**	**90Th Percentile**	**Maximum**
Broker-Volume	0.01	0.05	0.09	1.72
Customer-Position	0.01	0.03	0.06	2.19
Market-Feed	0.01	0.03	0.05	17.06
Market-Watch	0.01	0.03	0.07	2.46
Security-Detail	0.01	0.02	0.03	2.16
Trade-Lookup	0.01	0.40	0.60	3.34
Trade-Order	0.01	0.08	0.12	2.77
Trade-Result	0.01	0.09	0.15	14.72
Trade-Status	0.01	0.02	0.04	2.45
Trade-Update	0.01	0.45	0.61	10.32
Data-Maintenance	0.01	0.07	N/A	0.32

Transaction Mix	Transaction Count	Mix %
Broker-Volume	7,135,411	4.900
Customer-Position	18,930,528	13.000
Market-Feed	1,456,309	1.000
Market-Watch	26,211,229	18.000
Security-Detail	20,386,162	14.000
Trade-Lookup	11,649,210	8.000
Trade-Order	14,707,405	10.100
Trade-Result	14,563,015	10.001
Trade-Status	27,667,569	19.000
Trade-Update	2,912,297	2.000
Data-Maintenance	120	N/A

Test Duration and Timings	
Ramp-up Time (hh:mm:ss)	00:35:53
Measurement Interval (hh:mm:ss)	02:00:00
Business Recovery Time (hh:mm:ss)	00:39:39
Total Number of Transactions Completed in Measurement Interval	145,619,135

Table of Contents

Abstract **3**

Table of Contents **7**

Clause 0 – Preamble **9**

Clause 1 – Introduction **11**
Benchmark Sponsor 11
Configuration Diagrams 11
Measured and Priced Configurations 11
Figure 1-1. Measured Configuration 11
Table 1-1. Durable Media Substitution Information 12
Hardware and Software Configuration Steps 12

Clause 2- Database Design, Scaling, and Population **13**
Database Creation and Table Definitions 13
Database Physical Organization 13
Horizontal/Vertical Partitioning 14
Replication 14
Table Attributes 14
Cardinality of Tables 14
Table 2-1. Initial Cardinality of Tables 15
Distribution of Tables and Logs 16
Table 2-2. Data Distribution for the Measured and Priced Configuration 16
Database Interface and Model Implemented 20
Database Load Methodology 20

Clause 3 – Transaction Related Items **21**
Vendor-Supplied Code 21
Database Footprint of Transactions 21

Clause 4 – SUT, Driver, and Network **22**
Network Configuration 22

Clause 5 – EGen **23**
EGen Version 23
EGen Code and Modifications 23
EGen Files 23

Clause 6 – Performance Metrics and Response Time **24**
EGen Instances 24
Measured Throughput 24
Throughput vs. Elapsed Time for Trade-Result Transaction 24
Figure 6-1. Test Run Graph 24
Steady State Methodology 25
Work Performed During Steady State 25
Transaction Statistics 25
Table 6-1. Transaction Statistics 26

Clause 7 – Transaction and System Properties **27**
Atomicity Requirements 27
Consistency Requirements 27
Isolation Requirements 28
Durability Requirements 28
Durability Test for Data Accessibility 28
Figure 7-1. Data Accessibility Graph 29
Durability Test for Business Recovery 29
Figure 7-2. Business Recovery Time Graph 31

Clause 8 – Pricing **32**
60-Day Space 32
Table 8-1. Disk Space Requirements 32
Availability Date 33
Table 8-2. Ordering and Pricing Information 33
Supporting Files Index 33
Auditor's Attestation Letter 33
Appendix A – Price Quotes **36**

Clause 0 – Preamble

Introduction

TPC Benchmark™ E (TPC-E) is an On-Line Transaction Processing (OLTP) workload. It is a mixture of read-only and update intensive transactions that simulate the activities found in complex OLTP application environments. The database schema, data population, transactions, and implementation rules have been designed to be broadly representative of modern OLTP systems. The benchmark exercises a breadth of system components associated with such environments, which are characterized by:

- The simultaneous execution of multiple transaction types that span a breadth of complexity
- Moderate system and application execution time
- A balanced mixture of disk input/output and processor usage
- Transaction integrity (ACID properties)
- A mixture of uniform and non-uniform data access through primary and secondary keys
- Databases consisting of many tables with a wide variety of sizes, attributes, and relationships with realistic content
- Contention on data access and update

The TPC-E operations are modeled as follows: The database is continuously available 24 hours a day, 7 days a week, for data processing from multiple sessions and data modifications against all tables, except possibly during infrequent (e.g., once a month) maintenance sessions. Due to the worldwide nature of the application modeled by the TPC-E benchmark, any of the transactions may be executed against the database at anytime, especially in relation to each other.

Goal of the TPC-E Benchmark

The TPC-E benchmark simulates the OLTP workload of a brokerage firm. The focus of the benchmark is the central database that executes transactions related to the firm's customer accounts. In keeping with the goal of measuring the performance characteristics of the database system, the benchmark does not attempt to measure the complex flow of data between multiple application systems that would exist in a real environment.

The mixture and variety of transactions being executed on the benchmark system is designed to capture the characteristic components of a complex system. Different transaction types are defined to simulate the interactions of the firm with its customers as well as its business partners. Different transaction types have varying run-time requirements.

The benchmark defines:

- Two types of transactions to simulate Consumer-to-Business as well as Business-to-Business activities
- Several transactions for each transaction type
- Different execution profiles for each transaction type
- A specific run-time mix for all defined transactions

For example, the database will simultaneously execute transactions generated by systems that interact with customers along with transactions that are generated by systems that interact with financial markets as well as administrative systems. The benchmark system will interact with a set of driver systems that simulate the various sources of transactions without requiring the benchmark to implement the complex environment.

The performance metric reported by TPC-E is a "business throughput" measure of the number of completed Trade-Result transactions processed per second. Multiple transactions are used to simulate the business activity of processing a trade, and each transaction is subject to a response time constraint. The performance metric for the benchmark is expressed in transactions-per-second-E (tpsE). To be compliant with the TPC-E standard, all references to tpsE results must include the tpsE rate, the associated price-per-tpsE, and the availability date of the priced configuration.

TPC-E uses terminology and metrics that are similar to other benchmarks, originated by the TPC and others. Such similarity in terminology does not imply that TPC-E results are comparable to other benchmarks. The only benchmark results comparable to TPC-E are other TPC-E results that conform to a comparable version of the TPC-E specification.

Restrictions and Limitations

Despite the fact that this benchmark offers a rich environment that represents many OLTP applications, this benchmark does not reflect the entire range of OLTP requirements. In addition, the extent to which a customer can achieve the results reported by a vendor is highly dependent on how closely TPC-E approximates the customer application. The relative performance of systems derived from this benchmark does not necessarily hold for other workloads or environments. Extrapolations to any other environment are not recommended.

Benchmark results are highly dependent upon workload, specific application requirements, and systems design and implementation. Relative system performance will vary because of these and other factors. Therefore, TPC-E should not be used as a substitute for specific customer application benchmarking when critical capacity planning and/or product evaluation decisions are contemplated.

Clause 1 – Introduction

Benchmark Sponsor

A statement identifying the benchmark Sponsor(s) and other participating companies must be reported.

This benchmark was sponsored by IBM Corporation.

Configuration Diagrams

Diagrams of both the Measured and Priced Configurations must be reported, accompanied by a description of the differences.

Any information and/or measurement results used to prove the validity of a Component substitution must be included in the FDR. Original and substituted Components must be clearly identified. Proof of comparable performance for substitution without a full benchmark run must be cited in the FDR.

Measured and Priced Configurations

The measured configuration is shown in Figure 1-1. The priced configuration is shown above in the executive summary.

Figure 1-1. Measured Configuration

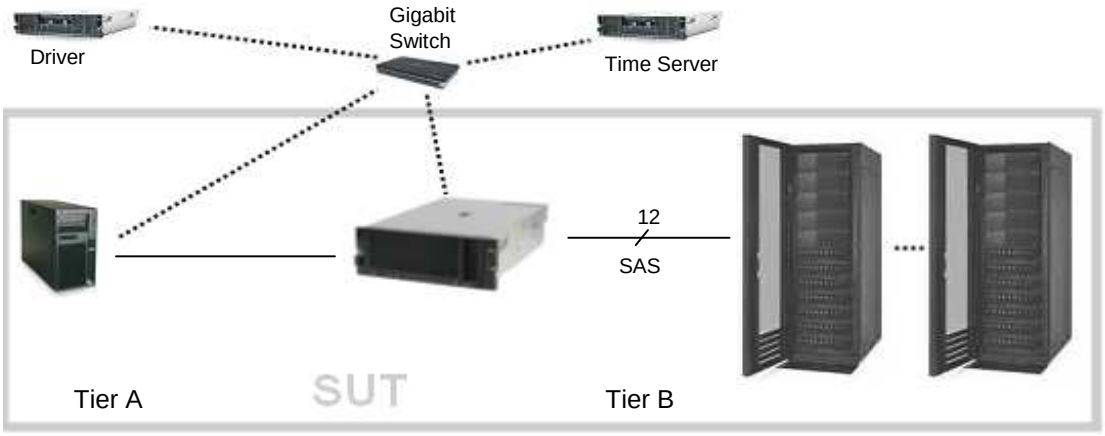

IBM x3500 M2
- 2 x Intel Xeon Processor X5570 2.93GHz (2 Procs/8 Cores/16 Threads)
- 8GB Memory
- 2 x 160GB SFF SATA (RAID-1)
- 1 x IBM ServeRAID-BR10i
- Onboard Gigabit Ethernet
- 1 x 10Gb Ethernet

IBM System x3850 X5
- 4 x Intel Xeon Processor X7560 2.26GHz (4 Procs/32 Cores/64 Threads)
- 1TB Memory
- 2 x 160GB SFF SATA (RAID-1)
- 6 x 300GB SFF SAS (RAID-10)
- 1 x IBM ServeRAID-M5015
- 6 x IBM ServeRAID-M5025
- Onboard Gigabit Ethernet
- 1 x 10Gb Ethernet

84 x IBM System Storage EXP3000 Enclosure
Four contain:
- 12 x 300GB 3.5" 15K SAS

80 contain:
- 12 x 73GB 3.5" 15K SAS

= 42 x 24-drive RAID-10
1008 Total External Drives

The measured and priced configurations differed in only the durable media used for the database data:

- The priced configuration used 960 146GB 15K 3.5" SAS drives, in addition to 48 300GB 15K 3.5" SAS drives.
- The measured configuration used 960 73GB 15K 3.5" SAS drives, in addition to the 48 300GB 15K 3.5" SAS drives.

This substitution was allowed based on the following information:

Table 1-1. Durable Media Substitution Information

Description	73GB (Measured)	146GB (Priced)
Capacity	73.4 GB	146.8 GB
Interface Type	SAS	SAS
Track-to-Track Seek (R/W)	0.2 / 0.4	0.2 / 0.4
Average Seek (R/W)	3.5 / 4.0	3.5 / 4.0
Interface Speed	3 Gb/s	3 Gb/s
Buffer Size	16 MB	16 MB
Rotational Speed	15,000 RPM	15,000 RPM
Media Density	110 Gbits/inch^2	110 Gbits/inch^2

Hardware and Software Configuration Steps

A description of the steps taken to configure all the hardware must be reported in the Report.
A description of the steps taken to configure all the software must be reported in the Report.
Any and all configuration scripts or step by step GUI instructions are reported in the Supporting Files (see Clauses 9.4.1.1 and 9.4.1.2). The description, scripts and GUI instructions must be sufficient such that a reader knowledgeable of computer systems and the TPC-E specification could recreate the hardware and software environments.

Detailed instructions for installing and configuring the SUT, hardware and software, are included in the supporting files:

- Information specific to the Tier A client can be found in: SupportingFiles\Introduction\TierA\TierA_x3500M2_Setup.pdf
- Information specific to the Tier B database server and storage can be found in: SupportingFiles\Introduction\TierB\TierB_x3850X5_Setup.pdf

Clause 2- Database Design, Scaling, and Population

Database Creation and Table Definitions

A description of the steps taken to create the database for the Reported Throughput must be reported in the Report. Any and all scripts or step by step GUI instructions are reported in the Supporting Files (see Clause 9.4.2). The description, scripts and GUI instructions must be sufficient such that a reader knowledgeable of database software environments and the TPC-E specification could recreate the database.

The database was created and populated using the Microsoft TPC-E benchmark kit. Instructions for doing so are included in the supporting files. See SupportingFiles\Clause2\MSTPCE Database Setup Reference.pdf.

Changes and customizations were made to some of the kit files. First, the filegroups the database was loaded onto were changed in number from three filegroups to two. Second, several scripts were modified to customize the load to the specific hardware configuration of this SUT.

The default kit files create the database on three filegroups: fixed_fg, scaling_fg, and growing_fg. That was changed so that only two filegroups were used, fixed_fg and growing_fg. All of the items that would have been loaded onto scaling_fg were loaded instead onto fixed_fg.

The modified files are included as part of SupportingFiles\Clause2:

- Utility\Create_TID_Ranges_Table.sql
- DDL\ Create_Indexes_Scaling_Tables.sql
- DDL\ Create_Tables_Scaling.sql
- DDL\ Create_Tables_Scaling_Flat.sql
- DDL\ Unified_Create_Indexes.sql

The files that were customized for this specific SUT hardware are included in the folder SupportingFiles\Clause2\1050000.Cust\Database:

- Tempdb.sql creates a larger temporary database for SQL Server
- Shrinktempdb.sql shrinks it back down
- Backupdev.sql creates devices for SQL Server to backup the database to
- Dropbackupdev.sql removes those devices
- Backup_Database.sql backs up the tpce database to the specified device names
- Restore_Database.sql restores the tpce database from the specified device names
- Create_Database.sql maps the database filegroups and log to physical storage
- Flatfile.txt tells the database loader where to store the database flatfiles during the load
- Remove_Database.sql drops the current tpce database

Database Physical Organization

The physical organization of tables and User-Defined Objects, within the database, must be reported in the Report.

The following tables and related indexes were on the growing_fg filegroup:

- CASH_TRANSACTION
- SETTLEMENT
- TRADE
- TRADE_HISTORY
- TRADE_REQUEST
- HOLDING

- HOLDING_HISTORY
- HOLDING_SUMMARY

The remaining tables and their related indexes were all on the fixed_fg filegroup.

Horizontal/Vertical Partitioning

While few restrictions are placed upon horizontal or vertical partitioning of tables and rows in the TPC-E benchmark (see Clause 2.3.3), any such partitioning must be reported.

Partitioning was not used for this benchmark.

Replication

Replication of tables, if used, must be reported in the Report.

Replication was not used for this benchmark.

Table Attributes

Additional and/or duplicated columns in any table must be reported in the Report along with a statement on the impact on performance (see Clause 2.3.5).

No additional attributes were used for this benchmark.

Cardinality of Tables

The cardinality (e.g., the number of rows) of each table, as it existed after the database load (see Clause 2.6), must be reported in the Report.

The database was built with 1,050,000 customers. The cardinality is shown in Table 2-1.

Table 2-1. Initial Cardinality of Tables

Table Name	Rows
ACCOUNT_PERMISSION	7,454,828
ADDRESS	1,575,004
BROKER	10,500
CASH_TRANSACTION	16,692,481,567
CHARGE	15
COMMISSION_RATE	240
COMPANY	525,000
COMPANY_COMPETITOR	1,575,000
CUSTOMER	1,050,000
CUSTOMER_ACCOUNT	5,250,000
CUSTOMER_TAXRATE	2,100,000
DAILY_MARKET	938,621,250
EXCHANGE	4
FINANCIAL	10,500,000
HOLDING	928,855,578
HOLDING_HISTORY	24,316,103,321
HOLDING_SUMMARY	52,216,631
INDUSTRY	102
LAST_TRADE	719,250
NEWS_ITEM	1,050,000
NEWS_XREF	1,050,000
SECTOR	12
SECURITY	719,250
SETTLEMENT	18,144,000,000
STATUS_TYPE	5
TAXRATE	320
TRADE	18,144,000,000
TRADE_HISTORY	43,545,559,454
TRADE_REQUEST	0
TRADE_TYPE	5
WATCH_ITEM	104,941,231
WATCH_LIST	1,050,000
ZIP_CODE	14,741

Distribution of Tables and Logs

The distribution of tables, partitions and logs across all media must be explicitly depicted for the Measured and Priced Configurations.

There were two 160GB 2.5" SFF 7.2K SATA drives and six 300GB 2.5" SFF 10K SAS drives in the database server, all accessed by an internal ServeRAID-M5015 SAS/SATA controller. The OS was loaded onto a RAID-1 array located on the two 160GB drives. The database log was stored on a RAID-10 array located on the six 300GB drives.

The database data was stored on external SAS storage. This storage was accessed by six IBM ServeRAID-M5025 SAS/SATA controllers, filling six of the seven PCI-E slots in the database server. Each of these controllers was connected to fourteen IBM System Storage EXP3000 enclosures, which held twelve 146GB or 300GB 3.5" 15K SAS drives each. In total, eighty-four EXP3000 enclosures and 1,008 external drives were connected to the server. Forty-two data arrays were each configured as 24-drive RAID-10. Each data array was broken into three partitions: one for fixed_fg (RAW), one for growing_fg (RAW), and one for backup, tempdb, and flatfiles (NTFS).

The measured configuration was the same as the priced configuration, except that the measured configuration used 73GB 3.5" 15K SAS drives in eighty of the eighty-four EXP3000 enclosures and forty of the forty-two data arrays.

Adapter write caching was disabled for all controllers and arrays.

Further details on the storage configuration are available in the supporting files. See the files in the directory SupportingFiles\Introduction\TierB.

Table 2-2 depicts the database configuration of the measured and priced systems to meet the 8-hour steady state requirement.

Table 2-2. Data Distribution for the Measured and Priced Configuration

Disk #	Controller	Drives Enclosure RAID level	Partition (File System)	Size	Use
0	Internal M5015	2x160GB SATA internal RAID-1	C: (NTFS)	147.86GB	OS
1	Internal M5015	6x300GB SAS internal RAID-10	E: (RAW) F: (NTFS)	488.28GB 346.99GB	Log MDF & tempDB
2	M5025 #1	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx1 (RAW) c:\mp\gw1 (RAW) c:\mp\bk1 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
3	M5025 #1	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx2 (RAW) c:\mp\gw2 (RAW) c:\mp\bk2 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
4	M5025 #1	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx3 (RAW) c:\mp\gw3 (RAW) c:\mp\bk3 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
5	M5025 #1	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx4 (RAW) c:\mp\gw4 (RAW) c:\mp\bk4 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB

Disk #	Controller	Drives Enclosure RAID level	Partition (File System)	Size	Use
6	M5025 #1	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx5 (RAW) c:\mp\gw5 (RAW) c:\mp\bk5 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
7	M5025 #1	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx6 (RAW) c:\mp\gw6 (RAW) c:\mp\bk6 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
8	M5025 #1	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx7 (RAW) c:\mp\gw7 (RAW) c:\mp\bk7 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
9	M5025 #2	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx8 (RAW) c:\mp\gw8 (RAW) c:\mp\bk8 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
10	M5025 #2	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx9(RAW) c:\mp\gw9 (RAW) c:\mp\bk9 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
11	M5025 #2	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx10 (RAW) c:\mp\gw10 (RAW) c:\mp\bk10 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
12	M5025 #2	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx11 (RAW) c:\mp\gw11 (RAW) c:\mp\bk11 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
13	M5025 #2	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx12 (RAW) c:\mp\gw12 (RAW) c:\mp\bk12 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
14	M5025 #2	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx13 (RAW) c:\mp\gw13 (RAW) c:\mp\bk13 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
15	M5025 #2	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx14 (RAW) c:\mp\gw14 (RAW) c:\mp\bk14 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
16	M5025 #3	24x300GB SAS (Priced & Measured) EXP3000 RAID-10	c:\mp\fx15 (RAW) c:\mp\gw15 (RAW) c:\mp\bk15 (NTFS)	5.57GB 256.54GB 3079.34GB	fixed_fg growing_fg backup & tempDB
17	M5025 #3	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx16 (RAW) c:\mp\gw16 (RAW) c:\mp\bk16 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
18	M5025 #3	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx17 (RAW) c:\mp\gw17 (RAW) c:\mp\bk17 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB

Disk #	Controller	Drives Enclosure RAID level	Partition (File System)	Size	Use
19	M5025 #3	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx18 (RAW) c:\mp\gw18 (RAW) c:\mp\bk18 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
20	M5025 #3	24x300GB SAS (Priced & Measured) EXP3000 RAID-10	c:\mp\fx19(RAW) c:\mp\gw19 (RAW) c:\mp\bk19 (NTFS)	5.57GB 256.54GB 3079.34GB	fixed_fg growing_fg backup & tempDB
21	M5025 #3	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx20 (RAW) c:\mp\gw20 (RAW) c:\mp\bk20 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
22	M5025 #3	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx21 (RAW) c:\mp\gw21 (RAW) c:\mp\bk21 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
23	M5025 #4	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx22 (RAW) c:\mp\gw22 (RAW) c:\mp\bk22 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
24	M5025 #4	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx23 (RAW) c:\mp\gw23 (RAW) c:\mp\bk23 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
25	M5025 #4	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx24 (RAW) c:\mp\gw24 (RAW) c:\mp\bk24 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
26	M5025 #4	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx25 (RAW) c:\mp\gw25 (RAW) c:\mp\bk25 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
27	M5025 #4	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx26 (RAW) c:\mp\gw26 (RAW) c:\mp\bk26 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
28	M5025 #4	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx27 (RAW) c:\mp\gw27 (RAW) c:\mp\bk27 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
29	M5025 #4	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx28 (RAW) c:\mp\gw28 (RAW) c:\mp\bk28 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
30	M5025 #5	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx29(RAW) c:\mp\gw29 (RAW) c:\mp\bk29 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
31	M5025 #5	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx30 (RAW) c:\mp\gw30 (RAW) c:\mp\bk30 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB

Disk #	Controller	Drives Enclosure RAID level	Partition (File System)	Size	Use
32	M5025 #5	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx31 (RAW) c:\mp\gw31 (RAW) c:\mp\bk31 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
33	M5025 #5	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx32 (RAW) c:\mp\gw32 (RAW) c:\mp\bk32 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
34	M5025 #5	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx33 (RAW) c:\mp\gw33 (RAW) c:\mp\bk33 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
35	M5025 #5	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx34 (RAW) c:\mp\gw34 (RAW) c:\mp\bk34 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
36	M5025 #5	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx35 (RAW) c:\mp\gw35 (RAW) c:\mp\bk35 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
37	M5025 #6	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx36 (RAW) c:\mp\gw36 (RAW) c:\mp\bk36 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
38	M5025 #6	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx37 (RAW) c:\mp\gw37 (RAW) c:\mp\bk37 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
39	M5025 #6	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx38 (RAW) c:\mp\gw38 (RAW) c:\mp\bk38 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
40	M5025 #6	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx39(RAW) c:\mp\gw39 (RAW) c:\mp\bk39 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
41	M5025 #6	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx40 (RAW) c:\mp\gw40 (RAW) c:\mp\bk40 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
42	M5025 #6	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx41 (RAW) c:\mp\gw41 (RAW) c:\mp\bk41 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB
43	M5025 #6	24x73.4GB SAS (Meas) 24x146GB SAS (Priced) EXP3000 RAID-10	c:\mp\fx42 (RAW) c:\mp\gw42 (RAW) c:\mp\bk42 (NTFS)	5.57GB 256.54GB 542.42GB	fixed_fg growing_fg backup & tempDB

Database Interface and Model Implemented

A statement must be provided in the Report that describes:

- *The Database Interface (e.g., embedded, call level) and access language (e.g., SQL, COBOL read/write) used to implement the TPC-E Transactions. If more than one interface / access language is used to implement TPC-E, each interface / access language must be described and a list of which interface /access language is used with which Transaction type must be reported.*
- *The data model implemented by the DBMS (e.g., relational, network, hierarchical).*

Microsoft SQL Server 2008 R2 Enterprise Edition is a relational database. The interface used was Microsoft SQL Server stored procedures accessed with Remote Procedure Calls embedded in C++ code using the Microsoft ODBC interface.

Database Load Methodology

The methodology used to load the database must be reported in the Report.

The database was loaded using the flat files option on the EGenLoader command line. This will generate flat files first, then bulk insert the data into the tables. A further description is provided in SupportingFiles\Clause2\MSTPCE Database Setup Reference.pdf.

Clause 3 – Transaction Related Items

Vendor-Supplied Code

A statement that vendor-supplied code is functionally equivalent to Pseudo-code in the specification (see Clause 3.2.1.6) must be reported.

The stored procedure code for the transactions was functionally equivalent to the pseudo-code. The stored procedures can be seen in SupportingFiles\Clause3\StoredProcedures.

The code to interface the stored procedures can be seen in:

- SupportingFiles\Clause3\BaseServer
- SupportingFiles\Clause3\TransactionsSP
- SupportingFiles\Clause3\TxnHarness

Database Footprint of Transactions

A statement that the database footprint requirements (as described in Clause 3.3) were met must be reported.

The database footprint requirements were met.

Clause 4 – SUT, Driver, and Network

Network Configuration

The Network configurations of both the Measured and Priced Configurations must be described and reported. This includes the mandatory Network between the Driver and Tier A (see Clause 4.2.2) and any optional Database Server interface networks (see Clause 4.1.3.12).

The network configurations of the measured and priced configurations were the same. Refer to Figure 1-1 for a diagram of the network connections.

The Tier A client and Tier B database server were each configured with one 10Gb Ethernet adapter. These two adapters were connected by a FC cable, creating a 10Gb network between the two systems, which handled all of the network traffic between Tier A and Tier B while a measurement was underway.

Another network connected the driver, the database server, the client, and a time server. This network, which was connected via a gigabit Ethernet switch, used one of the onboard Ethernet ports on the client and database server. This network fulfills the mandatory network between the driver and Tier A. It also allows the driver, client, and database server to synchronize and verify their times with the time server.

Clause 5 – EGen

EGen Version

The version of EGen used in the benchmark must be reported (see Clause 5.3.1).

EGen v1.9.0 was used in the benchmark.

EGen Code and Modifications

A statement that all required TPC-provided EGen code was used in the benchmark must be reported. If the Test Sponsor modified EGen, a statement EGen has been modified must be reported in the Report. All formal waivers from the TPC documenting the allowed changes to EGen must also be reported (see Clause 5.3.7.1). If any of the changes to EGen do not have a formal waiver, that must also be reported. If the Test Sponsor extended EGenLoader, the use of the extended EGenLoader and the audit of the extension code by an Auditor must be reported (see Clause 5.7.4).

All required TPC-provided EGen code was used in the benchmark.

EGen was not modified for use in this benchmark.

EGenLoader was not extended for this benchmark.

EGen Files

The make/project files used to compile/link EGenLoader and EGenValidate must be reported in the Supporting Files. The compiler/linker options and flags used to compile/link EGen objects for the SUT must be reported in the Supporting Files.

See the supporting files directory SupportingFiles\Clause3\prj for the files related to EGenLoader and EGenValidate.

See the supporting files directory SupportingFiles\Clause3\SUT_CE_Server for the files related to the SUT_CE_Server.

See the supporting files directory SupportingFiles\Clause3\SUT_MEE_Server for the files related to the SUT_MEE_Server.

Clause 6 – Performance Metrics and Response Time

EGen Instances

The number of EGenDriverMEE and EGenDriverCE instances used in the benchmark must be reported (see Clause 6.2.5).

There were eight EGenDriverCEs with a total of 1365 EGenDriverCE instances used in the benchmark.

There were eight EGenDriverMEEs with a dynamic number of instances used in the benchmark.

Measured Throughput

The Measured Throughput must be reported (see Clause 6.7.1.2).

The Measured Throughput was 2022.64 tpsE.

Throughput vs. Elapsed Time for Trade-Result Transaction

A Test Run Graph of throughput versus elapsed wall clock time must be reported for the Trade-Result Transaction (see Clause 6.7.2).

Figure 6-1. Test Run Graph

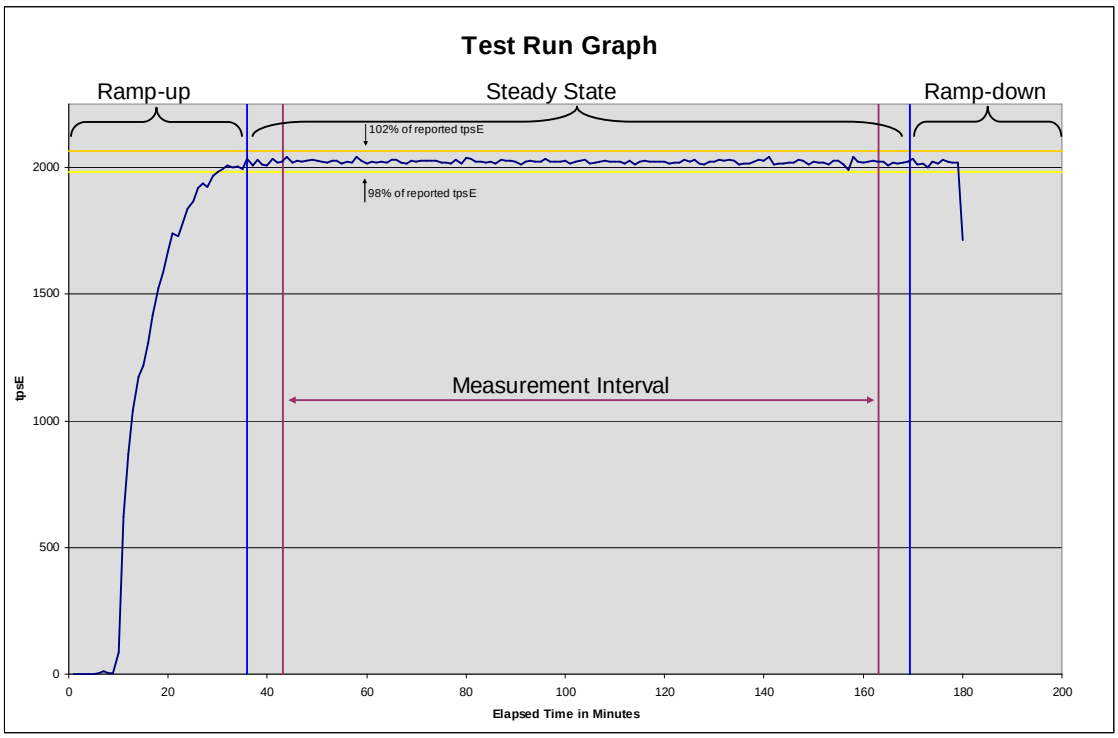

Steady State Methodology

The method used to determine that the SUT had reached a Steady State prior to commencing the Measurement Interval must be reported.

During the run, observation of the tpsE as the benchmark ran was used to determine Steady State. After the run, Steady State was confirmed by:

1. Looking at the Test Run Graph and verifying that tpsE was steady prior to commencing the Measurement Interval.
2. Calculating the average tpsE over 60-minute windows during Steady State, with the start of each window 10 minutes apart. Then it was confirmed that the minimum 60-minute average tpsE was not less than 98% of the Reported Throughput, and that the maximum 60-minute average tpsE was not greater than 102% of the Reported Throughput.
3. Calculating the average tpsE over 10-minute windows during Steady State, with the start of each window 1 minute apart. Then it was confirmed that the minimum 10-minute average tpsE was not less than 80% of the Reported Throughput, and the maximum 10-minute average tpsE was not greater than 120% of the Reported Throughput.

Work Performed During Steady State

A description of how the work normally performed during a Test Run, actually occurred during the Measurement Interval must be reported (e.g., checkpointing, writing Undo/Redo Log records).

Checkpoints were run continuously every 7½ minutes during the entire run.

Data-Maintenance was run every 60 seconds.

Transaction Statistics

The recorded averages over the Measurement Interval for each of the Transaction input parameters specified by clause 6.4.1 must be reported.

Table 6-1 contains the transaction statistics.

Table 6-1. Transaction Statistics

Input Parameter	Value	Actual Percentage	Required Range
Customer-Position			
By Tax ID	1	50.00%	48% to 52%
Get History	1	50.00%	48% to 52%
Market-Watch			
Securities chosen by	Watch List	60.00%	57% to 63%
	Account ID	35.00%	33% to 37%
	Industry	5.00%	4.5% to 5.5%
Security-Detail			
Access LOB	1	1.00%	0.9% to 1.1%
Trade-Lookup			
Frame to execute	1	29.99%	28.5% to 31.5%
	2	30.01%	28.5% to 31.5%
	3	30.00%	28.5% to 31.5%
	4	10.00%	9.5% to 10.5%
Trade-Order			
Transactions requested by a third party		9.99%	9.5% to 10.5%
By Company Name		39.99%	38% to 42%
Buy On Margin	1	8.01%	7.5% to 8.5%
Rollback	1	0.99%	0.94% to 1.04%
LIFO	1	35.00%	33% to 37%
Trade Quantity	100	25.01%	24% to 26%
	200	25.00%	24% to 26%
	400	25.00%	24% to 26%
	800	24.99%	24% to 26%
Trade Type	Market Buy	29.99%	29.7% to 30.3%
	Market Sell	30.00%	29.7% to 30.3%
	Limit Buy	20.01%	19.8% to 20.2%
	Limit Sell	10.00%	9.9% to 10.1%
	Stop Loss	10.01%	9.9% to 10.1%
Trade-Update			
Frame to execute	1	32.99%	31% to 35%
	2	33.00%	31% to 35%
	3	34.01%	32% to 36%

Clause 7 – Transaction and System Properties

The ACID (Atomicity, Consistency, Isolation, and Durability) properties of transaction processing systems must be supported by the System Under Test during the running of this benchmark. It is the intent of this section to define the ACID properties informally and to specify a series of tests that must be performed to demonstrate that these properties are met.

The results of the ACID tests must be reported in the Report along with a description of how the ACID requirements were met, and how the ACID tests were run.

Atomicity Requirements

The System Under Test must guarantee that Database Transactions are atomic; the system will either perform all individual operations on the data, or will ensure that no partially completed operations leave any effects on the data..

All ACID tests were conducted according to specification. The following steps were performed to verify the Atomicity of the Trade-Order transactions:

- *Perform a market Trade-Order Transaction with the roll_it_back flag set to false. Verify that the appropriate rows have been inserted in the TRADE and TRADE_HISTORY tables.*
- *Perform a market Trade-Order Transaction with the roll_it_back flag set to true. Verify that no rows associated with the rolled back Trade-Order have been added to the TRADE and TRADE_HISTORY tables.*

The procedure for running the atomicity tests is documented in the file SupportingFiles\Clause7\MSTPCE ACID Procedures.pdf

The atomicity scripts and outputs are located in the directory SupportingFiles\Clause7\Atomicity

Consistency Requirements

Consistency is the property of the Application that requires any execution of a Database Transaction to take the database from one consistent state to another. A TPC-E database when first populated by EGenLoader must meet these consistency conditions. The three consistency conditions must be tested after initial database population and after any Business Recovery tests.

Consistency condition 1:
Entries in the BROKER and TRADE tables must satisfy the relationship:
B_NUM_TRADES = count(*)
For each broker defined by:
(B_ID = CA_B_ID) and (CA_ID = T_CA_ID) and (T_ST_ID = ―CMPT‘).

Consistency condition 2:
Entries in the BROKER and TRADE tables must satisfy the relationship:
B_COMM_TOTAL = sum(T_COMM)
For each broker defined by:
(B_ID = CA_B_ID) and (CA_ID = T_CA_ID) and (T_ST_ID = ―CMPT‘).

Consistency condition 3:
Entries in the HOLDING_SUMMARY and HOLDING tables must satisfy the relationship:
HS_QTY = sum(H_QTY)
For each holding summary defined by:
(HS_CA_ID = H_CA_ID) and (HS_S_SYMB = H_S_SYMB).

Consistency conditions 1, 2, and 3 were tested using a batch file to issue queries to the database after the database was loaded and after the Business Recovery Test. The results of the queries demonstrated that the database was consistent for all three tests.

The procedure for running the consistency tests is documented in the file SupportingFiles\Clause7\MSTPCE ACID Procedures.pdf

The consistency scripts and outputs are located in the directory SupportingFiles\Clause7\Consistency

Isolation Requirements

The isolation property of a Transaction is the level to which it is isolated from the actions of other concurrently executing Transactions. Systems that implement Transaction isolation using a locking and/or versioning scheme must demonstrate compliance with the isolation requirements by executing the tests described in Clause 7.4.2.

Isolation tests 1 through 4 were successfully done following the procedure documented in the file SupportingFiles\Clause7\MSTPCE ACID Procedures.pdf

The isolation scripts and outputs are located in the directory SupportingFiles\Clause7\Isolation

Durability Requirements

The SUT must provide Durability. In general, state that persists across failures is said to be Durable and an implementation that ensures state persists across failures is said to provide Durability. In the context of the benchmark, Durability is more tightly defined as the SUT's ability to ensure all Committed data persist across a Single Point of Failure.

Durability Test for Data Accessibility

The Test Sponsor must report in the Report the Redundancy Level (see Clause 7.6.7.1) and describe the Data Accessibility test(s) used to demonstrate compliance.

A Data Accessibility Graph for each run demonstrating a Redundancy Level must be reported in the Report (see Clause 7.6.7.2).

This benchmark result used Redundancy Level 1. The test for Redundancy Level 1 is the test for permanent irrecoverable failure of any single Durable Medium.

To prove Redundancy Level 1, the following steps were successfully performed:

1. Restored the database to its freshly-loaded, proven-consistent state.
2. Determined the current number of completed trades in the database, *count1*.
3. Started a run, using the profile from the measured run, with checkpoints, and met the Durability Throughput Requirements for at least 5 minutes.
4. Induced the first failure, which in this case was failing a drive in a database data array by physically removing it from its enclosure. Since the database data arrays are RAID protected, transaction processing continued.
5. Waited until the Durability Throughput Requirements were met again for at least 5 minutes.
6. Induced the second failure, which in this case was failing a drive in the database log array by physically removing it from its enclosure. Since the database log array is RAID protected, transaction processing continued.
7. After a few minutes passed, a new drive was inserted into the log enclosure to replace the failed log drive. The log array rebuilding process was started.

8. After a few minutes passed, a new drive was inserted into the data enclosure to replace the failed data drive. The data array rebuilding process was started.
9. Continued running the benchmark for at least 20 minutes.
10. Terminated the run gracefully.
11. Retrieved the new number of completed trades in the database by running *select count(*) as count2 from SETTLEMENT.*
12. Verified that (*count2 – count1*), which is the number of actual completed Trade-Result Transactions done during the run, equaled the number of successful Trade-Result transactions reported by the Driver.
13. Allowed the recovery process to complete.

Figure 7-1 is a graph of the measured throughput versus elapsed time for the data accessibility run. The timings of the induced failures as well as the recovery process are indicated.

Figure 7-1. Data Accessibility Graph

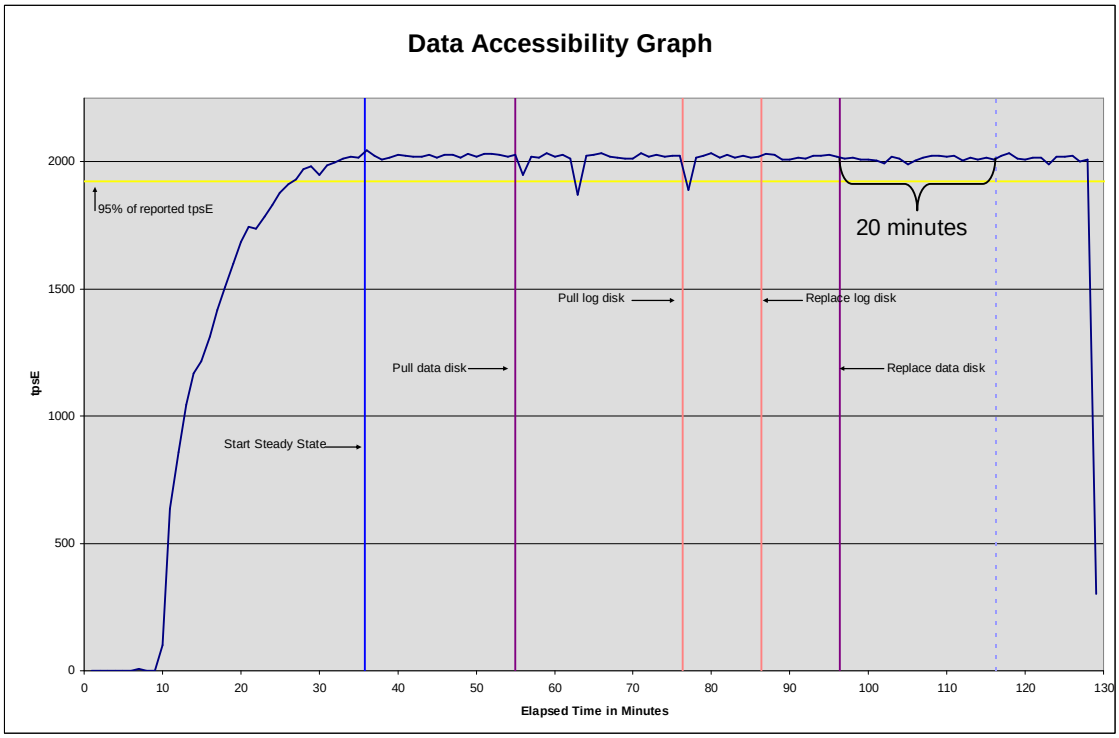

The files related to this data accessibility test are located in SupportingFiles\Clause7\Durability\DataAccessibility

Durability Test for Business Recovery

The Test Sponsor must describe in the Report the test(s) used to demonstrate Business Recovery
.
The Business Recovery Time must be reported on the Executive Summary Statement and in the Report. If the failures described in Clauses 7.6.2.2, 7.6.2.3 and 7.6.2.4 were not combined into one Durability test (usually powering off the Database Server during the run), then the Business Recovery Time for the failure described for instantaneous interruption is the Business Recovery Time that must be reported in the Executive Summary Statement. All the Business Recovery Times for each test requiring Business Recovery must be reported in the Report.

The Business Recovery Time Graph (see Clause 7.6.7.4) must be reported in the Report for all Business Recovery tests.

The tests for "Loss of processing," "Loss of Vulnerable Storage Component," and "Loss of all external power to the SUT" were combined.

The following steps were successfully performed to test Business Recovery:

1. Restored the database to its freshly-loaded, proven-consistent state.
2. Determined the current number of completed trades in the database, *count1*.
3. Started a run, using the profile from the measured run, with checkpoints, and met the Durability Throughput Requirements for at least 20 minutes.
4. Pulled the power cords from the database server.
5. Stopped submitting Transactions.
6. Plugged in and restarted the database server.
7. Started SQL Server on the database server. It automatically began recovery of the tpce database. The timestamp in the SQL Server ERRORLOG of the first message related to database tpce is considered the start of Database Recovery.
8. Waited for SQL Server to finish recovering the database. The timestamp in the SQL Server ERRORLOG of the message indicating that the recovery of database tpce is complete is considered the end of Database Recovery.
9. Since there was a time gap between the end of Database Recovery and the start of Application Recovery, and the Drivers and Transactions needed to be started again (not just continued), the Trade-Cleanup Transaction was executed during this time gap.
10. Started a run, using the profile from the measured run, with checkpoints. The time when the first transaction is submitted to the database is considered the start of Application Recovery.
11. Let the run proceed until a 20 minute window existed such that the first minute of the window and the entire window both scored at least 95% of the Reported Throughput. The time of the beginning of that 20-minute window is considered the end of Application Recovery.
12. Terminated the run gracefully.
13. Verified that no errors were reported during steps 8 through 12.
14. Retrieved the new number of completed trades in the database by running *select count(*) as count2 from SETTLEMENT.*
15. Verified that (*count2 – count1*), which is the number of actual completed Trade-Result Transactions done during the two runs, was greater than or equal to the combined number of successful Trade-Result Transactions reported by the Driver for both runs. In the case of an inequality, verified that the difference was less than or equal to the maximum number of transactions that could be simultaneously in-flight from the Driver to the SUT.
16. Verified database consistency.

Figure 7-2 is a graph of the measured throughput versus elapsed time for Business Recovery.

The Database Recovery Time was 00:21:59. The Application Recovery Time was 00:17:40. The Business Recovery Time, which is the sum of the Database Recovery Time and the Application Recovery Time, was 00:39:39.

Figure 7-2. Business Recovery Time Graph

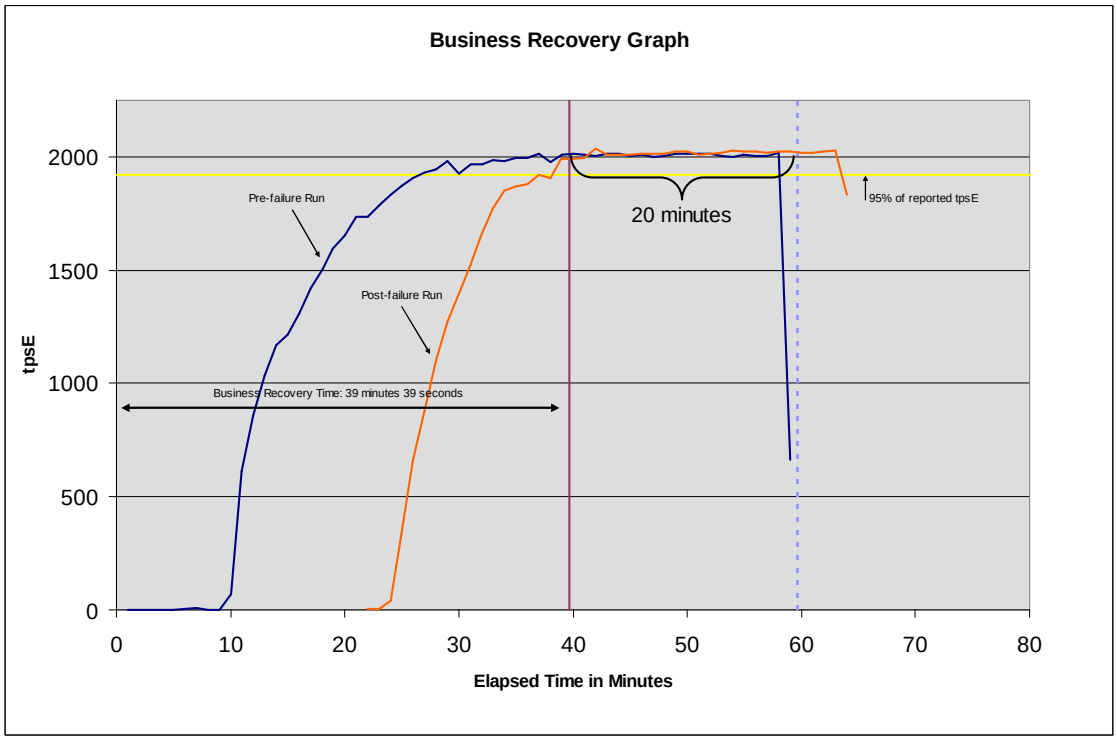

The files related to this business recovery test are located in SupportingFiles\Clause7\Durability\BusinessRecovery

Clause 8 – Pricing

60-Day Space

Details of the 60-Day Space computations (see Clause 6.6.6.6) along with proof that the database is configured to sustain a Business Day of growth (see Clause 6.6.6.1) must be reported in the Report.

Table 8-1. Disk Space Requirements

TPC-E Disk Space Requirements

Customers 1,050,000 **Performance** 2022.64 **TpsE** **Reported** 2022.64 **TpsE**

Table	Initial Rows	Data Size (KB)	Index Size (KB)	Extra 5% (KB)	Total + 5% (KB)	Rows After	After Run (KB)	Growth (KB)	Bus. Day Growth (KB)	Req. Add. (KB)
BROKER	10,500	768	1,008	89	1,865	10,500	1,776	-	-	89
CASH_TRANSACTION	16,692,481,567	1,723,585,096	3,636,760	86,361,093	1,813,582,949	16,710,873,795	1,730,925,312	3,703,456	10,790,165	10,790,165
CHARGE	15	8	8	1	17	15	16	-	-	1
COMMISSION_RATE	240	16	16	2	34	240	32	-	-	2
SETTLEMENT	18,144,000,000	961,819,224	2,030,776	48,192,500	1,012,042,500	18,163,993,563	966,030,728	2,180,728	6,353,637	6,353,637
TRADE	18,144,000,000	2,164,943,072	1,089,211,392	162,707,723	3,416,862,187	18,164,114,357	3,264,968,944	10,814,480	31,508,413	31,508,413
TRADE_HISTORY	43,545,559,454	1,309,641,144	3,416,480	65,652,881	1,378,710,505	43,593,595,968	1,316,754,680	3,697,056	10,771,519	10,771,519
TRADE_REQUEST	-	-	-	-	-	120,794	49,080	49,080	142,997	142,997
TRADE_TYPE	5	8	1,032	52	1,092	5	1,040	-	-	52
ACCOUNT_PERMISSION	7,454,828	634,704	4,552	31,963	671,219	7,454,828	639,464	208	607	31,963
CUSTOMER	1,050,000	177,976	47,288	11,263	236,527	1,050,000	225,304	40	117	11,263
CUSTOMER_ACCOUNT	5,250,000	475,848	102,784	28,932	607,564	5,250,000	578,632	-	-	28,932
CUSTOMER_TAXRATE	2,100,000	43,896	1,360	2,263	47,519	2,100,000	45,360	104	304	2,263
HOLDING	928,855,578	61,933,392	39,226,664	5,058,003	106,218,059	929,367,953	102,378,256	1,218,200	3,549,274	3,549,274
HOLDING_HISTORY	24,316,103,321	884,222,232	511,080,328	69,765,128	1,465,067,688	24,343,058,137	1,399,172,488	3,869,928	11,275,188	11,275,188
HOLDING_SUMMARY	52,216,631	2,277,984	9,584	114,378	2,401,946	52,216,100	2,287,568	-	-	-
WATCH_ITEM	104,941,231	2,928,528	11,872	147,020	3,087,420	104,941,231	2,940,608	208	607	147,020
WATCH_LIST	1,050,000	26,248	22,680	2,446	51,374	1,050,000	48,928	-	-	2,446
COMPANY	525,000	114,416	33,152	7,378	154,946	525,000	147,568	-	-	7,378
COMPANY_COMPETITOR	1,575,000	42,408	35,504	3,896	81,808	1,575,000	77,912	-	-	3,896
DAILY_MARKET	938,621,250	48,618,208	172,264	2,439,524	51,229,996	938,621,250	48,791,536	1,064	3,101	2,439,524
EXCHANGE	4	8	8	1	17	4	16	-	-	1
FINANCIAL	10,500,000	1,235,520	4,664	62,009	1,302,193	10,500,000	1,240,408	224	653	62,009
INDUSTRY	102	8	24	2	34	102	32	-	-	2
LAST_TRADE	719,250	44,792	1,360	2,308	48,460	719,250	46,152	-	-	2,308
NEWS_ITEM	1,050,000	113,839,688	2,504	5,692,110	119,534,302	1,050,000	113,842,208	16	47	5,692,110
NEWS_XREF	1,050,000	26,280	1,352	1,382	29,014	1,050,000	27,632	-	-	1,382
SECTOR	12	8	24	2	34	12	32	-	-	2
SECURITY	719,250	113,656	27,440	7,055	148,151	719,250	141,120	24	70	7,055
STATUS_TYPE	5	8	8	1	17	5	16	-	-	1
ADDRESS	1,575,004	90,944	1,360	4,615	96,919	1,575,004	92,352	48	140	4,615
TAXRATE	320	24	16	2	42	320	56	16	47	47
ZIP_CODE	14,741	488	176	33	697	14,741	664	-	-	33
TOTALS (KB)		**7,276,836,600**	**1,649,084,440**	**446,296,052**	**9,372,217,092**		**8,951,455,920**	**25,534,880**	**74,396,886**	**82,835,584**

Initial Database Size (MB) **8,716,720** **8,512 GB**

Database Filegroups	LUN Count	Partition Size (MB)	MB Allocated	MB Loaded	MB Required	
	0	-	-	-	-	OK
growing_fg	42	262,700	11,033,400	8,551,791	8,624,439	OK
	0	-	-	-	-	OK
fixed_fg	42	5,700	239,400	164,929	173,175	OK

Settlements **19,993,563**

Data Space Required (MB)		Data Space Configured (MB)					Log Space Required (MB)		Log Space Configured (MB)	
Initial Growing Space	8,551,791									
Final Growing Space	8,576,726	Data LUNS	40	2	-	-	Initial Log Size	10,980	Log LUNS	1
Delta	24,935	Disks per LUN	24	24	-	-	Final Log Size	141,745	Log Disks	6
Data Space per Trade	0.001247126	Disk Capacity	139,532	285,147	-	-	Log Growth	130,765	Disk Capacity	285,147
1 Day Data Growth	72,648	RAID Overhead	50%	50%	0%	0%	Log Growth/Trade	0.006540345	RAID Overhead	50%
60 Day Space	**13,075,579**	Total Space				**73,819,030**	1 Day Log Space	**391,969**	Log Space	**855,441**
						OK				OK

The 60-day space calculations are included in SupportingFiles\Clause8\ tpce_space.xls

Availability Date

The committed Availability Date of Components used in the price calculations must be reported with a precision of one day. All hardware, software and support used in the calculations must be Orderable by Any Customer on the Availability Date. For each of the Components that are not Orderable on the report date of the FDR, the following information must be included in the FDR:

- *Name and Part Number of the item that is not Orderable*
- *The date when the Component can be ordered (on or before the Availability Date)*
- *The method to be used to order the Component (at or below the quoted price) when the order date arrives*
- *The method for verifying the price*

The total solution as priced will be generally available July 30, 2010. The dates for ordering and availability are detailed in Table 8-2 for those components that are not immediately orderable.

Table 8-2. Ordering and Pricing Information

Description	Part Number	Order Date	Availability Date	Order Method	Price Verification
IBM ServeRAID-M5025 SAS/SATA Controller	46M0830	7-30-10	7-30-10	See note 1	See note 2

Note 1: IBM - 1-800-656-0833, x35330

Note 2: These components are not immediately orderable. For price verification before the order date, call IBM - 1-800-656-0833, x35330.

Supporting Files Index

An index for all files required by Clause 9.4 Supporting Files must be provided.

An index of the files contained in the supporting files is here: SupportingFiles\SupportingFilesIndex.pdf

Auditor's Attestation Letter

The Auditor's Attestation Letter, which indicates compliance, must be included in the Report.

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

INFO SIZING

TPC TRANSACTION PROCESSING PERFORMANCE COUNCIL CERTIFIED AUDITOR

Celia Schreiber, Manager
IBM System x and BladeCenter Performance Analysis and Benchmarking
IBM Systems and Technology Group
3039 Cornwallis Road
RTP, NC 27709

June 6, 2010

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

Platform: IBM System x3850 X5
Operating System: Microsoft Windows Server 2008 R2 Enterprise Edition
Database Manager: Microsoft SQL Server 2008 R2 Enterprise Edition

The results were:

CPU's Speed	Memory	Disks	Trade-Result 90% Response Time	tpsE
Tier B, Server: IBM System x3850 X5				
4 x Intel Xeon X7560 (2.26GHz)	1 TB (4 x 2 MB L2) (4 x 24 MB L3)	2 x 160GB 7.2K SATA 960 x 146GB 15K SAS 48 x 300 GB 15K SAS 6 x 300 GB 10K SAS	0.15 Seconds	2022.64
Tier A, OneClient: IBM System x3500 M2				
2 x Intel Xeon X5570 (2.93 GHz)	8 GB (2 x 1MB L2) (2 x 8 MB L3)	2 x 160 GB 7.2K SATA	n/a	n/a

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:

- All EGen components were verified to be v1.9.0.
- The transactions were correctly implemented.
- The database was properly scaled and populated for 1,050,000 customers.
- The mandatory network between the driver and the SUT was configured.
- The ACID properties were met.

- Input data was generated according to the specified percentages.
- The reported response times were correctly measured.
- All 90% response times were under the specified maximums.
- The measurement interval was representative of steady state conditions.
- The reported measurement interval was 120 minutes.
- The implementation used Redundancy Level 1.
- The Business Recovery Time of 00:39:39 was correctly measured.
- The 60 day storage requirement was correctly computed.
- The system pricing was verified for major components and maintenance.

Additional Audit Notes:

This implementation was originally audited on March 25, 2010. This Letter of Attestation was issued after reviewing materials in support of publishing a revised FDR.

The measured system included (960) 73GB 15K 3.5" SAS drives that were substituted by (960) 146GB 15K 3.5" SAS drives in the priced configuration. Based on the specifications of these drives, it is my opinion that this substitution has no significant effect on performance.

Respectfully Yours,

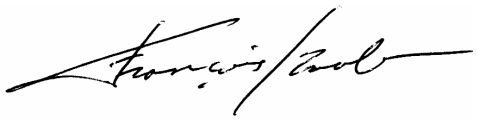

Doug Johnson, Auditor

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 16, 2010

IBM Corporation
Chris King
3079 Cornwallis Road
Durham, NC 27709

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

All pricing shown is in US Dollars ($).

Part Number	Description	Unit Price	Quantity	Price
*	**SQL Server 2008 R2 Enterprise Edition** *Per Processor License* *Open Program - Level C* *Unit Price reflects a 33% discount from the retail unit price of $28,749.*	$19,188	4	$76,752
P73-04217	**Windows Server 2008 R2 Enterprise Edition** *Server License with 25 CALs* *Open Program - Level C* *Unit Price reflects a 43% discount from the retail unit price of $3,999.*	$2,280	1	$2,280
P73-04980	**Windows Server 2008 R2 Standard Edition** *Server License with 5 CALs* *Open Program - Level C* *Unit Price reflects a 31% discount from the retail unit price of $1,029.*	$711	1	$711
N/A	**Microsoft Problem Resolution Services** *Professional Support* *(1 Incident).*	$259	1	$259

Windows Server 2008 R2 Enterprise Edition and Windows Server 2008 R2 Standard Edition is currently orderable and available through Microsoft's normal distribution channels. A list of Microsoft's resellers can be found at the Microsoft Product Information Center at
http://www.microsoft.com/products/info/render.aspx?view=22&type=how

SQL Server 2008 R2 Enterprise Edition will be orderable and available by May 6, 2010.

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

This quote is valid for the next 90 days.

Reference ID: TPCE_g3wOpiq6ZAu+JROPXruxPoytABrZB6hm_V1.0.0.

Shopping Cart

Print

Qty.	Product Description	Savings	Total Price
1	 Acer V173Bb Black 17" 5ms LCD Monitor Item #: N82E16824009185 Return Policy: Monitor Replacement Only Return Policy	-$5.00 Instant	~~$119.99~~ $114.99
		Subtotal:	$114.99
Calculate Shipping Zip Code: 27513 UPS Guaranteed 3 Day Service -- $8.99		Shipping:	$8.99
Redeem Newegg Gift Cards Card Number: Security Code:			
Apply Promo Code(s):		Promo Code:	$0.00
New Newegg Preferred Account Customers: Get $20 off $100. Subject to credit approval. Details			
		Grand Total:	$123.98

Print