

TPC Benchmark™ E
Full Disclosure Report
for
System x® 3950 X6
using
Microsoft® SQL Server® 2014
Enterprise Edition
and
Microsoft Windows Server® 2012
Standard Edition

TPC-E™ Version 1.13.0



First Edition
Submitted for Review
November 25, 2014

First Edition – November 2014

THE INFORMATION CONTAINED IN THIS DOCUMENT IS DISTRIBUTED ON AN AS IS BASIS WITHOUT ANY WARRANTY EITHER EXPRESSED OR IMPLIED. The use of this information or the implementation of any of these techniques is the customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item has been reviewed by Lenovo for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environment do so at their own risk.

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

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

© Copyright Lenovo Corporation 2014. All rights reserved.

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

Trademarks

Lenovo, System x, the Lenovo logo, and For Those Who Do are trademarks or registered trademarks of Lenovo Corporation.

The following terms used in this publication are trademarks of other companies as follows: TPC Benchmark, TPC-E, and tpsE are trademarks of the Transaction Processing Performance Council; IBM, the IBM logo, and System Storage are trademarks or registered trademarks of International Business Machines Corporation; Intel and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries; Microsoft, Windows Server, and SQL Server are registered trademarks of Microsoft Corporation in the United States and/or other countries. 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

Lenovo® Corporation conducted the TPC Benchmark™ E on the System x®3950 X6 configured as a client/server system. This report documents the full disclosure information required by the TPC Benchmark E Standard Specification, Revision 1.13.0, including the methodology used to achieve the reported results. All testing fully complied with this revision level.


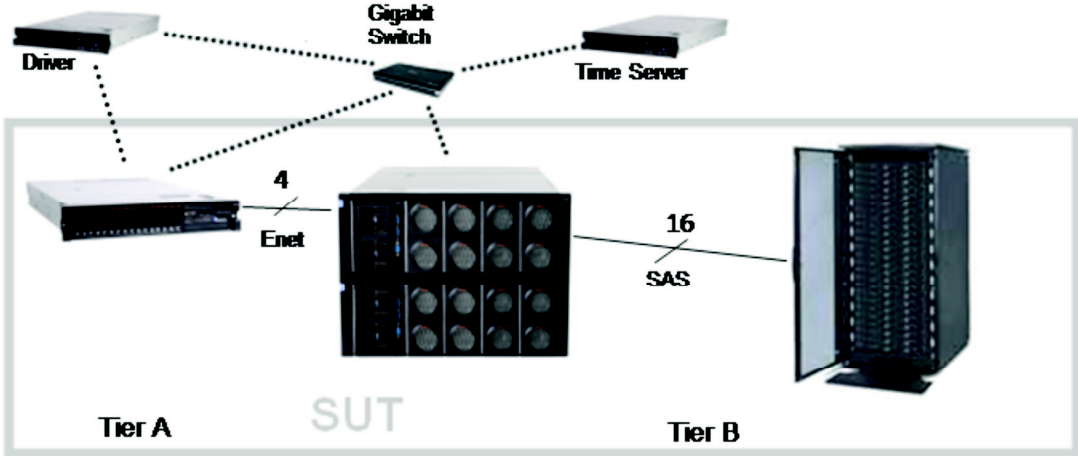
The software used on the System x3950 X6 system included Microsoft® Windows Server® 2012 Standard Edition and Microsoft SQL Server® 2014 Enterprise Edition.

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
System x3950 X6	Microsoft SQL Server 2014 Enterprise Edition Microsoft Windows Server 2012 Standard Edition	\$1,759,232 USD	9145.01	\$192.38 USD	November 25, 2014

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.

	System x[®] 3950 X6 Microsoft[®] SQL Server[®] 2014		TPC-E[™] 1.13.0 TPC Pricing 2.0.0
			Report Date: November 25, 2014
			Revision Date: November 25, 2014
TPC-E Throughput 9145.01 tpsE	Price/Performance \$192.38 USD per tpsE[™]	Availability Date November 25, 2014	Total System Cost \$1,759,232 USD
Database Server Configuration			
Operating System Microsoft Windows Server[®] 2012 Standard Edition	Database Manager Microsoft SQL Server 2014 Enterprise Edition	Processors/Cores/Threads 8/120/240	Memory 4096GB
			
System x3650 M4, with: <ul style="list-style-type: none"> - 2 x Intel[®] Xeon[®] Processor E5-2697 v2 2.70GHz (2 Procs/24 Cores/48 Threads) - 32GB Memory - 2 x 250GB SFF SATA (RAID-1) - 1 x ServeRAID M5110e - Onboard Quad Gb Ethernet - Dual Port Gb Ethernet 	System x3950 X6, with: <ul style="list-style-type: none"> - 8 x Intel Xeon Processor E7-8890 v2 2.80GHz (8 Procs/120 Cores/240 Threads) - 4096GB Memory - 2 x 300GB 10K SAS (RAID-1) - 6 x 800GB SATA SSD (RAID-10) - 2 x ServeRAID M5210 - 8 x ServeRAID M5120 - 4 x Dual 10Gb-T Ethernet 	16 x EXP2524 JBOD Enclosures, with: <ul style="list-style-type: none"> - 255 x 200GB 2.5" SAS SSD (3 x 19-drive RAID-5 DB data) - 9 x 22-drive RAID-5 DB data) - 56 x 400GB 2.5" SAS SSD (4 x 14-drive RAID-5 DB data) 311 Total External Drives	
Initial Database Size 38,595 GB	Redundancy Level: 1 RAID-10 Log RAID-5 Data	Storage 2 x 300GB 2.5" 10K SAS 6 x 800GB 2.5" SATA SSD 255 x 200GB 2.5" SAS SSD 56 x 400GB 2.5" SAS SSD	



System x3950 X6 Microsoft SQL Server 2014

TPC-E 1.13.0

TPC Pricing 2.0.0

Report Date:
November 25, 2014

Revision Date:
November 25, 2014

Availability Date:
November 25, 2014

Description	Part Number	Price Source	Unit Price	Quantity	Extended Price	3-Yr. Maint. Price
Server Hardware						
System x3950 X6 Configure-To-Order, includes:	3837AC2		281,892	1	281,892	
x3950 X6 8U Chassis + Midplane	A4BM, A4A5			1		
X6 Compute Book with Intel Xeon Processor E7-8890 v2	A4AZ, A4BF			8		
X6 Primary I/O Book + X6 Storage Book	A4A0, A4A1			2		
X6 Half-length I/O Book	A4A2			4		
4x 2.5" HS SAS/SATA/SSD HDD Backplane	A4A6			3		
ServeRAID M5210 SAS/SATA Controller for System x	A3YZ			2		
Intel X540 ML2 Dual Port 10GbBase-T Adapter for System x	A40P			1		
Intel X540-T2 Dual Port 10GbBase-T Adapter for System x	A2ED			3		
1400W HE Redundant Power Supply	A54E			8		
S3700 800GB SATA 2.5" G3HS Enterprise SSD	A4U5			6		
300GB 10K 6Gbps SAS 2.5" G3HS HDD	A4TL			2		
System x Rail Kit	A4AA			1		
Power Cable	6311			8		
32GB PC3L-12800 ECC DDR3 1600MHz LP LRDIMM	A3SR			128		
ServeRAID M5120 SAS/SATA Controller	81Y4478	1	259	8	2,072	
ServeRAID M5100 Series 1GB Flash/RAID 5 Upgrade	81Y4559	1	475	8	3,800	
ServeRAID M5100 Series SSD Performance Key	90Y4273	1	49	1	49	
Preferred Pro Keyboard USB - US English 103P RoHS v2	00AM600	1	29	1	29	
2-Button Optical Mouse - Black - USB	40K9200	1	19	1	19	
ThinkVision E1922 18.5-inch LED Backlit LCD Monitor	60B8AAR6US	1	110	1	110	
ServicePac for 3-Year 24x7x4 Support (x3950 X6)	67568BU	1	1,500	1		1500
	Subtotal				287,971	1,500
Server Storage						
S2 42U Standard Rack	93074RX	1	1,565	1	1,565	
EXP2524 Storage Enclosure	610024X	1	3,999	16	63,984	
1M SAS cable	39R6529	1	119	16	1,904	
200GB 2.5 Inch Flash Drive	00NC573	1 - S	3,079	255	785,145	
400GB 2.5 Inch Flash Drive	00NC575	1	6,279	56	351,624	
ServicePac for 3-Year 24x7x4 Support (EXP2524)	91Y5785	1	1,200	16		19,200
ServicePac for 3-Year 24x7x4 Support (Rack)	41L2760	1	315	1		315
	Subtotal				1,204,222	19,515
Server Software						
SQL Server 2014 Enterprise Edition	7JQ-00750	2a	13,472.50	60	808,350	
Windows Server 2012 Standard Edition	P73-05761	2	882	4	3,528	
Microsoft Problem Resolution Services	N/A	2a	259	1		259
	Subtotal				811,878	259
Client Hardware						
System x3650 M4 Configure-To-Order, includes:	7915AC1	1	12,328	1	12,328	
x3650 M4 Base + Planar	A1KF, A3V6			1		
750W High Efficiency Platinum AC Power Supply	A1H5			2		
Intel Xeon Processor E5-2697 v2 12C 2.7GHz 30MB 130W	A3VM, A3W5			2		
NetXtreme II 1000 Express Dual Port Ethernet Adapter	2995			1		
x3650 M4 PCIe Riser Card 1 (1 x8 FH/FL + 2 x8 FH/HL Slots)	A1JT			1		
System x Gen-III Slides Kit	A228			1		
250GB 7.2K 6Gbps NL SATA 2.5" SFF HS HDD	A1NX			2		
System x Lightpath Kit	A1LF			1		
x3650 M4 8x 2.5" HS HDD Assembly Kit	A1JX			1		
4GB PC3L-12800 CL11 ECC DDR3 1600MHz LP RDIMM	A3QE			8		
Power Cable	6263			2		
ServeRAID M5100 Series 512MB Cache/RAID 5 Upgrade	A1J3			1		
ServicePac for 3-Year 24x7x4 Support (x3650 M4)	67567XR	1	790	1		790
	Subtotal				12,328	790
Client Software						
Windows Server 2012 Standard Edition	P73-05761	2	882	1	882	
	Subtotal				882	0
Infrastructure						
Ethernet Cables	78004256	1		6	24	
	Subtotal				24	0
	Total				2,317,305	22,064
Dollar Volume Discount (See Note 1)	37.96%	1			579,402	
Microsoft Open Program Discount Schedule	16.67%	2			735	

Pricing: 1 - Lenovo 1-877-782-7134; 2 - Microsoft
 Note 1: Discount applies to all line items where Pricing=1; pricing is for these or similar quantities.
 Discounts for similarly sized configurations will be similar to what is quoted here, but may vary based on the specific components priced.
 S: One or more components of the measured configuration have been substituted in the priced configuration. See the FDR for details.

Three-Year Cost of Ownership USD: \$1,759,232
TPC-E Throughput: 9,145.01
\$ USD/tpsE: \$192.38

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.



**System x3950 X6
Microsoft SQL Server 2014**

**TPC-E 1.13.0
TPC Pricing 2.0.0**

Report Date:
November 25, 2014

Revision Date:
November 25, 2014

Availability Date:
November 25, 2014

Numerical Quantities Summary				
Reported Throughput: 9145.01 tpsE		Configured Customers: 4,700,000		
Response Time (in seconds)	Minimum	Average	90 th Percentile	Maximum
Broker-Volume	0.00	0.02	0.04	3.12
Customer-Position	0.00	0.01	0.03	1.78
Market-Feed	0.00	0.01	0.04	2.50
Market-Watch	0.00	0.01	0.03	3.28
Security-Detail	0.00	0.01	0.02	1.52
Trade-Lookup	0.00	0.07	0.11	2.46
Trade-Order	0.00	0.04	0.08	1.74
Trade-Result	0.00	0.03	0.05	2.40
Trade-Status	0.00	0.01	0.02	2.04
Trade-Update	0.01	0.08	0.12	2.13
Data-Maintenance	0.00	0.02	N/A	0.25
Transaction Mix		Transaction Count	Mix %	
Broker-Volume		32,261,772	4.900%	
Customer-Position		85,592,352	13.000%	
Market-Feed		6,584,432	1.000%	
Market-Watch		118,512,357	18.000%	
Security-Detail		92,175,983	14.000%	
Trade-Lookup		52,672,331	8.000%	
Trade-Order		66,498,305	10.100%	
Trade-Result		65,844,115	10.001%	
Trade-Status		125,097,126	19.000%	
Trade-Update		13,167,894	2.000%	
Data-Maintenance		120	N/A	
Test Duration and Timings				
Ramp-up Time (hh:mm:ss)			00:43:46	
Measurement Interval (hh:mm:ss)			02:00:00	
Business Recovery Time (hh:mm:ss)			00:38:27	
Total Number of Transactions Completed in Measurement Interval			658,406,667	

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.....	12
Table 1-1. Durable Media Substitution Information	13
Hardware and Software Configuration Steps	13
Clause 2- Database Design, Scaling, and Population	14
Database Creation and Table Definitions	14
Database Physical Organization	14
Horizontal/Vertical Partitioning	15
Replication.....	15
Table Attributes	15
Cardinality of Tables	15
Table 2-1. Initial Cardinality of Tables.....	16
Distribution of Tables and Logs	17
Table 2-2. Data Distribution for the Measured and Priced Configurations.....	17
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
Reported Throughput.....	24
Throughput vs. Elapsed Time for Trade-Result Transaction.....	24
Figure 6-1. Test Run Graph	24
Steady State Methodology.....	24
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
Table 7-1. Combinations of Durable Media Technologies Tested for Data Accessibility.....	28
Figure 7-1. Data Accessibility Graph.....	30
Durability Test for Business Recovery	30

Figure 7-2. Business Recovery Time Graph	32
Clause 8 – Pricing	33
60-Day Space.....	33
Table 8-1. Disk Space Requirements	33
Availability Date.....	34
Supporting Files Index.....	34
Auditor’s Attestation Letter.....	34
Appendix A – Price Quotes.....	37

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 Lenovo Corporation.

Configuration Diagrams


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

Some hardware components of the Priced Configuration may be substituted after the Test Sponsor has demonstrated to the Auditor's satisfaction that the substituting components do not negatively impact the Reported Throughput. All Substitutions must be reported in the Report and noted in the Auditor's Attestation Letter. Any information and/or measurement results used to prove the validity of a Component substitution must be included in the Report. Original and substituted Components must be clearly identified.

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



System x3650 M4, with:
 - 2 x Intel® Xeon® Processor E5-2697 v2 2.70GHz (2 Procs/24 Cores/48 Threads)
 - 32GB Memory
 - 2 x 250GB SFF SATA (RAID-1)
 - 1 x ServeRAID M5110e
 - Onboard Quad Gb Ethernet
 - Dual Port Gb Ethernet

System x3950 X6, with:
 - 8 x Intel Xeon Processor E7-8890 v2 2.80GHz (8 Procs/120 Cores/240 Threads)
 - 4096GB Memory
 - 2 x 300GB 10K SAS (RAID-1)
 - 6 x 800GB SATA SSD (RAID-10)
 - 1 x ServeRAID M5210
 - 8 x ServeRAID M5120
 - 4 x Dual 10Gb-T Ethernet

26 x EXP2524 JBOD Enclosures, with:
 - 255 x 200GB 2.5" SAS SSD (3 x 19-drive RAID-5 DB data) (9 x 22-drive RAID-5 DB data)
 - 56 x 400GB 2.5" SAS SSD (4 x 14-drive RAID-5 DB data)
 - 144 x 600GB 2.5" SAS HDD (6 x 24-drive RAID-10 backup)
 - 96 x 1200GB 2.5" SAS HDD (4 x 24-drive RAID-10 backup)
 551 Total External Drives

Compared to the priced configuration, the measured configuration contained extra external enclosures and drives used strictly for database backup files and flat file space used during the benchmark database load process. These extra enclosures and drives were not used at all during the benchmark runs.

As shown above, there were sixteen RAID-5 arrays containing database data. Each of these arrays held the same type and amount of database data, and did the same amount of work. Three of these arrays used 19 200GB SSDs each, which are the priced 200GB SSDs. Nine of these arrays used 22 200GB SSDs each. These arrays used older SSDs, but were priced one-for-one as the newer SSDs. This substitution was allowed based on performance measurements, taken while the configuration was running, that show that the 19-drive arrays using the newer SSDs were faster than the 22-drive arrays that used the older SSDs. Table 1-1 shows these measurements.

Table 1-1. Durable Media Substitution Information

Description	22-Drive Array	19-Drive Array
Drives Used	Older SSD	Newer SSD
Drive Capacity	200GB	200GB
Interface Type	SAS	SAS
Interface Speed	3 Gbps	6 Gbps
Drives Per Array	22	19
Avg Array Response Times (read / write, ms)	0.62 / 1.62	0.24 / 0.63
Avg Array Queue Depth	23.4	9.1
Priced As	22 New 200GB SSDs	19 New 200GB SSDs

Hardware and Software Configuration Steps

A description of the steps taken to configure all the hardware must be reported.

A description of the steps taken to configure all the software must be reported.

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_x3650M4_Setup.pdf
- Information specific to the Tier B database server and storage can be found in:
SupportingFiles\Introduction\TierB\TierB_x3950X6_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. 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

The files that were customized for this specific SUT hardware are included in the folder SupportingFiles\Clause2\4700000.Cust\Database:

- Tempdb_load.sql specifies temporary database files to use when loading the database
- Tempdb_run.sql specifies temporary database files to use when running the database
- Shrinktempdb.sql removes extra tempdb files
- Backupdev.sql creates devices for SQL Server to back up 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.

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.

Replication was not used for this benchmark.

Table Attributes

Additional and/or duplicated columns in any table must be reported 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.

The database was built with 4,700,000 customers. The cardinality is shown in Table 2-1.

Table 2-1. Initial Cardinality of Tables

Table Name	Rows
ACCOUNT_PERMISSION	33,370,932
ADDRESS	7,050,004
BROKER	47,000
CASH_TRANSACTION	74,718,714,407
CHARGE	15
COMMISSION_RATE	240
COMPANY	2,350,000
COMPANY_COMPETITOR	7,050,000
CUSTOMER	4,700,000
CUSTOMER_ACCOUNT	23,500,000
CUSTOMER_TAXRATE	9,400,000
DAILY_MARKET	4,201,447,500
EXCHANGE	4
FINANCIAL	47,000,000
HOLDING	4,158,080,490
HOLDING_HISTORY	108,843,343,118
HOLDING_SUMMARY	233,740,702
INDUSTRY	102
LAST_TRADE	3,219,500
NEWS_ITEM	4,700,000
NEWS_XREF	4,700,000
SECTOR	12
SECURITY	3,219,500
SETTLEMENT	81,216,000,000
STATUS_TYPE	5
TAXRATE	320
TRADE	81,216,000,000
TRADE_HISTORY	194,918,461,958
TRADE_REQUEST	0
TRADE_TYPE	5
WATCH_ITEM	470,061,785
WATCH_LIST	4,700,000
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 300GB 2.5” 10K SAS drives in the server accessed by the internal ServeRAID M5210 SAS/SATA controller. The OS was loaded onto a RAID-1 array located on these two drives.

The database log and run-time tempdb were stored on six 800GB 2.5” SATA SSDs in the server accessed by the internal ServeRAID M5210 SAS/SATA controller. These drives were used to create a RAID-10 array.

The database data was stored on external SAS SSD storage. This storage was accessed by eight ServeRAID M5120 SAS/SATA controllers. Each of these controllers was connected to two EXP2524 enclosures for database data:

- Nine of these enclosures held 22 200GB SAS SSDs each → 9 x 22-drive RAID-5
- Three of these enclosures held 19 200GB SAS SSDs each → 3 x 19-drive RAID-5
- Four of these enclosures held 14 400GB SAS SSDs each → 4 x 14-drive RAID-5

In total, for database data, sixteen enclosures and 311 external SSDs were connected to the database server and were used to create sixteen RAID-5 data arrays. Each data array was broken into three partitions: one for fixed_fg (RAW), one for growing_fg (RAW), and one for load-time tempdb (NTFS).

In addition to the priced configuration described above, the measured configuration included ten additional external EXP2524 enclosures. Six of these were each filled with twenty-four 600GB SAS HDDs; the other four of these were each filled with twenty-four 1200GB SAS HDDs. This space was used to generate and load the TPC-E benchmark database, and during database backup and restore operations. This hardware performed no function during benchmark runs. These additional ten enclosures were attached to the previously mentioned ServeRAID M5120 SAS/SATA controllers via daisy-chaining. Ten 24-drive RAID-10 arrays were created using this hardware and formatted as NTFS.

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 Configurations

Disk #	Controller	Drives Enclosure RAID Level (Pricing)	Partition (File System)	Size	Use
0	M5120 #1	19 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx4 c:\mp\gw4 c:\mp\xt4	65.33GB 3108.40GB 162.14GB	fixed_fg growing_fg tempdb
1	M5120 #1	22 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx3 c:\mp\gw3 c:\mp\xt3	65.33GB 3108.40GB 718.13GB	fixed_fg growing_fg tempdb
2	M5120 #1	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk2 (NTFS)	6694.21GB	backup & flatfiles
3	Internal M5210	2 x 300GB SAS HDD internal RAID-1	C: (NTFS)	277.95GB	OS

Disk #	Controller	Drives Enclosure RAID Level (Pricing)	Partition (File System)	Size	Use
4	Internal M5210	6 x 800GB SATA SSD internal RAID-10	E: (RAW) F: (NTFS)	1954.10GB 278.15GB	tpce log MDF tempdb
5	M5120 #2	19 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx9 c:\mp\gw9 c:\mp\xt9	65.33GB 3108.40GB 162.14GB	fixed_fg growing_fg tempdb
6	M5120 #2	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk1 (NTFS)	6694.21GB	backup & flatfiles
7	M5120 #2	19 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx10 c:\mp\gw10 c:\mp\xt10	65.33GB 3108.40GB 162.14GB	fixed_fg growing_fg tempdb
8	M5120 #3	22 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx5 c:\mp\gw5 c:\mp\xt5	65.33GB 3108.40GB 718.13GB	fixed_fg growing_fg tempdb
9	M5120 #3	22 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx6 c:\mp\gw6 c:\mp\xt6	65.33GB 3108.40GB 718.13GB	fixed_fg growing_fg tempdb
10	M5120 #3	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk4 (NTFS)	6694.21GB	backup & flatfiles
11	M5120 #4	14 x 400GB SAS SSD EXP2524 RAID-5	c:\mp\fx7 (RAW) c:\mp\gw7 (RAW) c:\mp\xt7 (NTFS)	65.33GB 3108.40GB 1656.91GB	fixed_fg growing_fg tempdb
12	M5120 #4	24 x 1200GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk7 (NTFS) c:\mp\bk8 (NTFS)	6699.87GB 6699.87GB	backup & flatfiles
13	M5120 #4	14 x 400GB SAS SSD EXP2524 RAID-5	c:\mp\fx8 (RAW) c:\mp\gw8 (RAW) c:\mp\xt8 (NTFS)	65.33GB 3108.40GB 1656.91GB	fixed_fg growing_fg tempdb
14	M5120 #4	24 x 1200GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk11 (NTFS) c:\mp\bk12 (NTFS)	6699.87GB 6699.87GB	backup & flatfiles
15	M5120 #5	22 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx15 c:\mp\gw15 c:\mp\xt15	65.33GB 3108.40GB 718.13GB	fixed_fg growing_fg tempdb
16	M5120 #5	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk3 (NTFS)	6694.21GB	backup & flatfiles
17	M5120 #5	22 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx16 c:\mp\gw16 c:\mp\xt16	65.33GB 3108.40GB 718.13GB	fixed_fg growing_fg tempdb
18	M5120 #6	22 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx11 c:\mp\gw11 c:\mp\xt11	65.33GB 3108.40GB 718.13GB	fixed_fg growing_fg tempdb

Disk #	Controller	Drives Enclosure RAID Level (Pricing)	Partition (File System)	Size	Use
19	M5120 #6	22 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx12 c:\mp\gw12 c:\mp\xt12	65.33GB 3108.40GB 718.13GB	fixed_fg growing_fg tempdb
20	M5120 #6	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk5 (NTFS)	6694.21GB	backup & flatfiles
21	M5120 #7	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk6 (NTFS)	6694.21GB	backup & flatfiles
22	M5120 #7	22 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx13 c:\mp\gw13 c:\mp\xt13	65.33GB 3108.40GB 718.13GB	fixed_fg growing_fg tempdb
23	M5120 #7	22 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx14 c:\mp\gw14 c:\mp\xt14	65.33GB 3108.40GB 718.13GB	fixed_fg growing_fg tempdb
24	M5120 #8	14 x 400GB SAS SSD EXP2524 RAID-5	c:\mp\fx1 (RAW) c:\mp\gw1 (RAW) c:\mp\xt1 (NTFS)	65.33GB 3108.40GB 1656.91GB	fixed_fg growing_fg tempdb
24	M5120 #8	14 x 400GB SAS SSD EXP2524 RAID-5	c:\mp\fx2 (RAW) c:\mp\gw2 (RAW) c:\mp\xt2 (NTFS)	65.33GB 3108.40GB 1656.91GB	fixed_fg growing_fg tempdb
26	M5120 #8	24 x 1200GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk13 (NTFS) c:\mp\bk14 (NTFS)	6699.87GB 6699.87GB	backup & flatfiles
27	M5120 #8	24 x 1200GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk9 (NTFS) c:\mp\bk10 (NTFS)	6699.87GB 6699.87GB	backup & flatfiles

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 2014 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.

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 found 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 had six Gb Ethernet ports. Four of these are provided by the onboard Ethernet chip and the other two are provided by a dual-port PCI-e Gb Ethernet adapter.

The Tier B database server had eight 10Gb Ethernet ports. These were provided by four dual-port 10Gb Ethernet adapters.

The Tier A client and Tier B database server were connected by four Ethernet crossover connections. These cables were plugged into one of the two ports of each 10Gb adapter in the database server. On the client, these cables plugged into two of the onboard Gb Ethernet ports and both of the Gb Ethernet adapter ports. These crossover networks, all running at 1Gb, handled all of the network traffic between Tier A and Tier B while a measurement was underway.

An additional crossover connection was setup between the Tier A client and the driver. This network, which fulfills the mandatory network between the driver and Tier A, was used by the client to report its results to the driver as a benchmark run was underway.

Another network connected the driver, the database server, the client, and a time server. This network, which was connected via a Gb Ethernet switch, used one of the onboard Ethernet ports on the client and a free 10Gb Ethernet port on the database server. It was used for miscellaneous file sharing and time syncing. It was not used during a benchmark run.

Clause 5 – EGen

EGen Version

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

EGen v1.13.0 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. 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 (as described in Appendix A.6), 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 v1.13.0 introduces non-trivial constructors for certain classes defined in TxnHarnessStructs.h. As a consequence it is a compile-time error to use any of these classes as a member of a union. The TPC-E subcommittee has been informed of this situation. This change in EGen compile-time behavior is unintentional so the TPC-E subcommittee has classified this as a logic error (per TPC Policies v6.2 Clause 5.4.4) and will address it in a future release of EGen. In the interim, the TPC-E subcommittee recommends that affected test sponsors wishing to publish a result proceed according to TPC-E v1.13.0 Clause 5.3.6. Accordingly, EGen was modified for this publication by removing the constructors in question. The TPC-E subcommittee has discussed this solution and found no compliance issues with it. The file TxnHarnessStructs.h can be found in Supporting Files Clause5.

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 16 EGenDriverCEs with a total of 1760 EGenDriverCE instances used in the benchmark.

There were 16 EGenDriverMEEs with a dynamic number of instances used in the benchmark.

Reported Throughput


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

The Reported Throughput was 9,145.01 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, Steady State was determined by observation of the Trade-Result transactions per second. After the run, Steady State was confirmed by:

1. Looking at the Test Run Graph and verifying that the Trade-Result transactions per second was steady prior to commencing the Measurement Interval.
2. Calculating the average Trade-Result transactions per second 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 Trade-Result transactions per second was not less than 98% of the Reported Throughput, and that the maximum 60-minute average Trade-Result transactions per second was not greater than 102% of the Reported Throughput.
3. Calculating the average Trade-Result transactions per second 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 Trade-Result transactions per second was not less than 80% of the Reported Throughput, and the maximum 10-minute average Trade-Result transactions per second 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, etc.).

Checkpoints had a duration of 430 seconds and were scheduled to run every 447 seconds.

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	59.99%	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	30.01%	28.5% to 31.5%
	2	29.99%	28.5% to 31.5%
	3	30.01%	28.5% to 31.5%
	4	10.00%	9.5% to 10.5%
Trade-Order			
Transactions requested by a third party		10.00%	9.5% to 10.5%
By Company Name		40.00%	38% to 42%
Buy On Margin	1	7.99%	7.5% to 8.5%
Rollback	1	0.99%	0.94% to 1.04%
LIFO	1	35.00%	33% to 37%
Trade Quantity	100	25.00%	24% to 26%
	200	25.00%	24% to 26%
	400	25.00%	24% to 26%
	800	25.01%	24% to 26%
Trade Type	Market Buy	30.00%	29.7% to 30.3%
	Market Sell	29.99%	29.7% to 30.3%
	Limit Buy	20.00%	19.8% to 20.2%
	Limit Sell	10.01%	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	32.99%	31% to 35%
	3	34.02%	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 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 zero. 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 one. 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. These 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 any Single Point of Failure.

Durability Test for Data Accessibility

The Test Sponsor must report the Redundancy Level (see Clause 7.6.3.4) and describe the Data Accessibility test(s) used to demonstrate compliance. A list of all combinations of Durable Media technologies tested in Clause 7.6.3.5 must be reported.

A Data Accessibility Graph for each run demonstrating a Redundancy Level must be reported (see Clause 7.6.4.2).

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.

The combinations of Durable Media technologies that were tested are shown in table 7-1. All unique combinations that contained database data, the database log, and/or the tempdb database were tested.

Table 7-1. Combinations of Durable Media Technologies Tested for Data Accessibility


Contents	Durable Media Type	Bus Type	Array Redundancy	Controller
Database Data	SSD	SAS	RAID-5	ServeRAID M5120
Database Log and tempdb	SSD	SATA	RAID-10	ServeRAID M5210

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

1. Performed Trade-Cleanup to remove remnants of previous benchmark runs from the database.
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 Data Accessibility Throughput Requirements for at least 5 minutes.
4. Induced the first failure, which in this case was failing a drive in the database log & tempdb array by physically removing it from its enclosure. Since the database log & tempdb array is RAID protected, transaction processing continued.
5. Waited until the Data Accessibility Throughput Requirements were met again for at least 5 minutes.
6. Induced the second 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.
7. 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.
8. After a few minutes passed, a new drive was inserted into the log & tempdb enclosure to replace the failed log drive. The log 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 Data Accessibility. 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 the test(s) used to demonstrate Business Recovery.

The Business Recovery Time must be reported. If the failures described in Clauses 7.5.3.1, 7.5.3.2 and 7.5.3.3 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 7.5.8.2) must be reported 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. Performed Trade-Cleanup to remove remnants of previous benchmark runs from the database.
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, causing it to immediately cease functioning. All the contents of the server’s main memory and caches were lost. All the disk controllers were inside the server, and none of their batteries were present, so all disk controller cache contents were lost.
5. Stopped submitting Transactions.
6. Plugged in and restarted the database server. It booted a fresh copy of the OS from the OS array.

7. Deleted the data file and log file for tempdb.
8. 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.
9. 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.
10. 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.
11. 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.
12. 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.
13. Terminated the run gracefully.
14. Verified that no errors were reported during steps 8 through 13.
15. Retrieved the new number of completed trades in the database by running *select count(*) as count2 from SETTLEMENT*.
16. 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.
17. Verified database consistency.

The Database Recovery Time was 00:19:21. The Application Recovery Time was 00:19:06. The Business Recovery Time, which is the sum of the Database Recovery Time and the Application Recovery Time, was 00:38:27.

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

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.

The 60-day space calculations shown in Table 8-1 are included in SupportingFiles\Clause8\tpce_space.xls.

Table 8-1. Disk Space Requirements

Table	Customers	Measured Throughput				Reported Throughput				
	4,700,000	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	47,000	4,624	10,728	768	16,120	47,000	15,352	-	-	768
CASH_TRANSACTION	74,718,714,407	7,779,281,224	16,398,248	389,783,974	8,185,463,446	74,829,551,643	7,816,976,528	21,297,056	46,559,720	46,559,720
CHARGE	15	8	8	1	17	15	16	-	-	1
COMMISSION_RATE	240	16	16	2	34	240	32	-	-	2
SETTLEMENT	81,216,000,000	3,872,693,552	8,178,336	194,043,594	4,074,915,482	81,336,471,937	3,892,694,664	11,822,776	25,847,006	25,847,006
TRADE	81,216,000,000	9,696,667,512	5,403,433,632	755,005,057	15,855,106,201	81,337,563,207	15,155,191,880	55,090,736	120,439,614	120,439,614
TRADE_HISTORY	194,918,461,958	5,862,210,168	15,288,656	293,874,941	6,171,373,765	195,210,219,234	5,898,649,536	21,150,712	46,239,782	46,239,782
TRADE_REQUEST	-	-	-	-	-	-	-	-	-	-
TRADE_TYPE	5	8	1,032	52	1,092	5	1,040	-	-	52
ACCOUNT_PERMISSION	33,370,932	1,837,640	11,936	92,479	1,942,055	33,370,932	1,849,656	80	175	92,479
CUSTOMER	4,700,000	770,296	231,368	50,083	1,051,747	4,700,000	1,001,696	32	70	50,083
CUSTOMER_ACCOUNT	23,500,000	2,129,496	524,928	132,721	2,787,145	23,500,000	2,654,424	-	-	132,721
CUSTOMER_TAXRATE	9,400,000	196,112	3,216	9,966	209,294	9,400,000	199,512	184	403	9,966
HOLDING	4,158,080,490	278,342,800	190,290,704	23,431,675	492,065,179	4,161,112,204	475,850,192	7,216,688	15,777,156	15,777,156
HOLDING_HISTORY	108,843,343,118	3,957,940,528	2,644,058,936	330,099,973	6,932,099,437	109,005,807,042	6,626,815,960	24,816,496	54,253,935	54,253,935
HOLDING_SUMMARY	233,740,702	10,262,568	40,192	515,138	10,817,898	233,740,724	10,302,760	-	-	-
WATCH_ITEM	470,061,785	13,217,088	50,648	663,387	13,931,123	470,061,785	13,268,088	352	770	663,387
WATCH_LIST	4,700,000	117,192	109,408	11,330	237,930	4,700,000	226,600	-	-	11,330
COMPANY	2,350,000	501,944	152,576	32,726	687,246	2,350,000	654,576	56	123	32,726
COMPANY_COMPETITOR	7,050,000	189,392	173,864	18,163	381,419	7,050,000	363,256	-	-	18,163
DAILY_MARKET	4,201,447,500	197,222,056	578,856	9,890,046	207,690,958	4,201,447,500	197,802,584	1,672	3,656	9,890,046
EXCHANGE	4	8	8	1	17	4	16	-	-	1
FINANCIAL	47,000,000	5,296,464	17,080	265,677	5,579,221	47,000,000	5,313,912	368	805	265,677
INDUSTRY	102	8	24	2	34	102	32	-	-	2
LAST_TRADE	3,219,500	200,816	2,832	10,182	213,830	3,219,500	203,648	-	-	10,182
NEWS_ITEM	4,700,000	509,565,896	8,112	25,478,700	535,052,708	4,700,000	509,574,080	72	158	25,478,700
NEWS_XREF	4,700,000	117,176	3,216	6,020	126,412	4,700,000	120,392	-	-	6,020
SECTOR	12	8	24	2	34	12	32	-	-	2
SECURITY	3,219,500	446,928	126,584	28,676	602,188	3,219,500	573,536	24	53	28,676
STATUS_TYPE	5	8	8	1	17	5	16	-	-	1
ADDRESS	7,050,004	406,760	3,232	20,500	430,492	7,050,004	410,080	88	193	20,500
TAXRATE	320	24	16	2	42	320	56	16	35	35
ZIP_CODE	14,741	488	104	30	622	14,741	592	-	-	30
TOTALS (KB)		32,189,618,808	8,279,698,528	2,023,465,867	42,492,783,203		40,610,714,744	141,397,408	309,123,654	345,828,760
Initial Database Size (MB)		39,520,818	38,595 GB							
Database Filegroups	LUN Count	Partition Size (MB)	MB Allocated	MB Loaded	MB Required					
	0	-	-	-	-					OK
growing_fg	16	3,084,500	49,352,000	38,803,796	39,105,668					OK
	0	-	-	-	-					OK
fixed_fg	16	64,700	1,035,200	717,022	752,873					OK
Settlements	120,471,937									
Data Space Required (MB)	Data Space Configured (MB)				Log Space Required (MB)		Log Space Configured (MB)			
Initial Growing Space	38,803,796	Data LUNS	3	9	4	Initial Log Size	47,277	Log LUNS	1	
Final Growing Space	38,941,876	Disks per LUN	19	22	14	Final Log Size	858,255	Log Disks	6	
Delta	138,081	Disk Capacity	189,781	189,781	380,516	Log Growth	810,978	Disk Capacity	761,985	
Data Space per Trade	0.001146163	RAID Overhead	95%	95%	93%	Log Growth/Trade	0.006731676	RAID Overhead	50%	
1 Day Data Growth	301,872	Total Space				1 Day Log Space	1,820,241	Log Space	2,285,955	
60 Day Space	57,633,154									

OK

OK

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 November 25, 2014.

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.

Marc Baker, Manager
System x Server Performance
Lenovo Enterprise Business Group
8001 Development Drive
Morrisville, NC 27560

November 20, 2014

I verified the TPC Benchmark™ E v1.13.0 performance of the following configuration:

Platform: System x3950 X6
Operating System: Microsoft Windows Server 2012 Standard Edition
Database Manager: Microsoft SQL Server 2014 Enterprise Edition

The results were:

Performance Metric 9145.01 tpsE
Trade-Result 90th %-tile 0.05 Seconds

Tier B (Server)

System x3950 X6

CPU	8 x Intel Xeon Processor E7-8890 v2 (2.80 GHz, 15-core, 37.5 MB L3)		
Memory	4096 GB		
Storage	Qty	Size	Type
	2	300 GB	10K rpm SAS HDD
	6	800 GB	SATA SSD
	255	200 GB	SAS SSD
	56	400 GB	SAS SSD

Tier A (Client)

System x3650 M4

CPU	2 x Intel Xeon Processor E5-2697 v2 (2.70 GHz, 12-core, 30 MB L3)
Memory	32 GB
Storage	2 x 250 GB 7.2K rpm SATA HDD

In my opinion, these performance results were produced in compliance with the TPC requirements for the benchmark.

The following verification items were given special attention:

- All EGen components were verified to be v1.13.0
- The transactions were correctly implemented
- The database was properly scaled and populated for 4,700,000 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 120 minutes
- The implementation used Redundancy Level 1
- The Business Recovery Time of 00:38:27 was correctly measured
- The 60-day storage requirement was correctly computed
- The system pricing was verified for major components and maintenance

Additional Audit Notes:

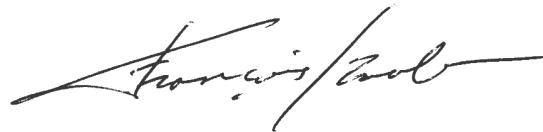
The measured system included (198) XceedIOPS SAS SSD disks that were substituted by (198) Optimus SAS SSD disks in the priced configuration. Based on the specifications of these disks and on I/O data collected during testing, it is my opinion that this substitution has no significant effect on performance.

EGen v1.13.0 introduces non-trivial constructors for certain classes defined in TxnHarnessStructs.h. As a consequence it is a compile-time error to use any of these classes as a member of a union. The TPC-E subcommittee has been informed of this situation. This change in EGen compile-time behavior is unintentional so the TPC-E subcommittee has classified this as a logic error (per TPC Policies v6.2 Clause 5.4.4) and will address it in a future release of EGen. In the interim, the TPC-E subcommittee recommends that affected test sponsors wishing to publish a result proceed according to TPC-E v1.13.0 Clause 5.3.6. Accordingly, EGen was modified for this publication by removing the constructors in question. The TPC-E subcommittee has discussed this solution and found no compliance issues with it.

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

November 7, 2014

Lenovo
Ray Engler
8001 Development Drive
Morrisville, NC 27560

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
Database Management System				
7JQ-00750	SQL Server 2014 Enterprise Edition <i>2 Core License Open Program - Level C</i>	\$13,472.50	60	\$808,350.00
Database Server Operating System				
P73-05761	Windows Server 2012 Standard Edition <i>2 Processor License Open Program - Level C Unit Price reflects a 17% discount from the retail unit price of \$882.</i>	\$735.00	4	\$2,940.00
Tier-A Operating System(s)				
P73-05761	Windows Server 2012 Standard Edition <i>2 Processor License Open Program - Level C Unit Price reflects a 17% discount from the retail unit price of \$882.</i>	\$735.00	1	\$735.00
Support				
N/A	Microsoft Problem Resolution Services <i>Professional Support (1 Incident).</i>	\$259.00	1	\$259.00

SQL Server 2014 Enterprise Edition and Windows Server 2012 Standard Edition are currently orderable and available through Microsoft's normal distribution channels. A list of Microsoft's resellers can be found in the Microsoft Product Information Center at

<http://www.microsoft.com/products/info/render.aspx?view=22&type=how>

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: TPCE_qhtplyIGYLKTVUKf28479jydi_2014_rewhep.