

**TPC Benchmark™ E**  
**Full Disclosure Report**  
**for**  
**System x® 3950 X6**  
**using**  
**Microsoft® SQL Server® 2014**  
**Enterprise Edition**  
**and**  
**Microsoft Windows Server® 2012 R2**  
**Standard Edition**

**TPC-E™ Version 1.14.0**



**First Edition**  
**Submitted for Review**  
**December 17, 2015**

## **First Edition – December 2015**

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 2015. 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; 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**

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

<sup>2</sup> 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 1.14.0, including the methodology used to achieve the reported results. All testing fully complied with this revision level.


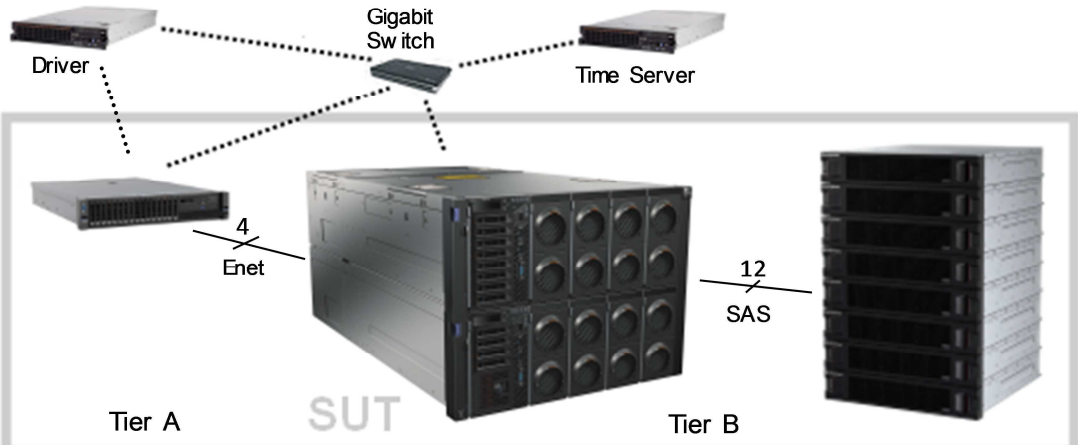
The software used on the System x3950 X6 system included Microsoft® Windows Server® 2012 R2 Standard Edition and Microsoft SQL Server® 2014 Enterprise Edition.

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:

<b>Hardware</b>	<b>Software</b>	<b>Total System Cost</b>	<b>tpsE</b>	<b>\$ USD /tpsE</b>	<b>Total Solution Availability Date</b>
System x3950 X6	Microsoft SQL Server 2014 Enterprise Edition Microsoft Windows Server 2012 R2 Standard Edition	\$1,591,485 USD	11,058.99	\$143.91 USD	December 17, 2015

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

	<b>System x<sup>®</sup> 3950 X6</b> <b>Microsoft<sup>®</sup> SQL Server<sup>®</sup> 2014</b>		<b>TPC-E<sup>™</sup> 1.14.0</b> <b>TPC Pricing 2.0.0</b>
			Report Date: December 17, 2015
			Revision Date: December 17, 2015
TPC-E Throughput <b>11,058.99 tpsE</b>	Price/Performance <b>\$143.91</b> <b>USD per tpsE<sup>™</sup></b>	Availability Date <b>December 17, 2015</b>	Total System Cost <b>\$1,591,485 USD</b>
<b>Database Server Configuration</b>			
Operating System <b>Microsoft Windows Server<sup>®</sup> 2012 R2 Standard Edition</b>	Database Manager <b>Microsoft SQL Server 2014 Enterprise Edition</b>	Processors/Cores/Threads <b>8/144/288</b>	Memory <b>4096GB</b>
			
System x3650 M5, with: - 2 x Intel <sup>®</sup> Xeon <sup>®</sup> Processor E5-2699 v3 2.30GHz (2 Procs/36 Cores/72 Threads) - 64GB Memory - 2 x 300GB SFF SAS (RAID-1) - 1 x ServeRAID M5210 - Onboard Quad Gb Ethernet - 2 x Dual 10Gb-T Ethernet	System x3950 X6, with: - 8 x Intel Xeon Processor E7-8890 v3 2.50GHz - 4096GB Memory - 2 x 900GB 10K SAS (RAID-1) - 6 x 800GB SAS SSD (RAID-10) - 1 x ServeRAID M5210 - 8 x ServeRAID M5225 - 4 x Dual 10Gb-T Ethernet	12 x Lenovo E1024 JBOD Enclosures, with: - 104 x 800GB 2.5" SAS SSD (6 x 17-drive RAID-5) (1 x 2-drive RAID-1)	104 Total External Drives
Initial Database Size <b>47,222 GB</b>	Redundancy Level: <b>1</b> <b>RAID-10 Log</b> <b>RAID-5 Data</b> <b>RAID-1 tempdb</b>	Storage <b>2 x 900GB 2.5" 10K SAS</b> <b>6 x 800GB 2.5" SAS SSD</b> <b>104 x 800GB 2.5" SAS SSD</b>	



# System x3950 X6 Microsoft SQL Server 2014

**TPC-E 1.14.0**  
**TPC Pricing 2.0.0**

Report Date:  
December 17, 2015

Revision Date:  
December 17, 2015

Availability Date:  
December 17, 2015

Description	Part Number	Price Source	Unit Price	Quantity	Extended Price	3-Yr. Maint. Price
<b>Server Hardware</b>						
System x3950 X6 Configure-To-Order, includes:	6241AC2	1	276,329	1	276,329	
x3950 X6 8U Chassis + Midplane	ASMJ, A4A5			1		
X6 DDR4 Compute Book Intel Xeon E7-8890 v3	AS8B, AS91			8		
X6 Primary I/O Book + X6 Storage Book	ASFB, A4A1			2		
X6 Half-length I/O Book	A4A2			4		
4x 2.5" HDD Riser	A4A6			2		
ServeRAID M5210 SAS/SATA Controller	A3YZ			2		
Intel X540 ML2 Dual Port 10GbBaseT Adapter	A40P			1		
Intel X540-T2 Dual Port 10GbBaseT Adapter	A2ED			3		
1400W HE Redundant Power Supply	A54E			8		
800GB 12G SAS 2.5" MLC G3HS Enterprise SSD	AS7G			6		
900GB 10K 6Gbps SAS 2.5" G3HS HDD	A4TN			2		
System x Rail Kit	A4AA			1		
Power Cable	6311			8		
32GB TruDDR4 Memory PC4-17000 CL15 2133MHz LP RDIMM	A5UJ			128		
ServeRAID M5225-2GB SAS/SATA Controller	00AE938	1	899	8	7,192	
ServeRAID M5200 Series Performance Accelerator- FoD	47C8710	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	100	1	100	
ServicePac for 3-Year 24x7x4 Support (x3950 X6)	67568BU	1	1,500	1		1500
				<b>Subtotal</b>	<b>283,718</b>	<b>1,500</b>
<b>Server Storage</b>						
S2 42U Standard Rack	93074RX	1	1,565	1	1,565	
Lenovo Storage E1024	64111B4	1	2,999	12	35,988	
1.5m SAS Cable (mSAS HD to mSAS)	00MU163	1	129	12	1,548	
Lenovo Storage 2.5" 800GB SSD (SAS)	00MM715	1	4,999	104	519,896	
ServicePac for 3-Year 24x7x4 Support (E1024)	00WX880	1	1,146	12		13,752
ServicePac for 3-Year 24x7x4 Support (Rack)	41L2760	1	315	1		315
				<b>Subtotal</b>	<b>558,997</b>	<b>14,067</b>
<b>Server Software</b>						
SQL Server 2014 Enterprise Edition	7JQ-00750	2a	13,472.50	72	970,020	
Windows Server 2012 R2 Standard Edition	P73-06284	2b	882	4	3,528	
Microsoft Problem Resolution Services	N/A	2a	259	1		259
				<b>Subtotal</b>	<b>973,548</b>	<b>259</b>
<b>Client Hardware</b>						
System x3650 M5 Configure-To-Order, includes:	5462AC1	1	20,101	1	20,101	
x3650 M5 Base + Planar + Power Paddle Card	A5FD, A5EA, A5FT			1		
750W High Efficiency Platinum AC Power Supply	A5EU			2		
Intel Xeon Processor E5-2699 v3 18C 2.3GHz 45MB 145W	ARYJ, ARYT			2		
x3650 M5 PCIe Risers 1 & 2 + Thermal Kit	A5FQ, A5R6, ASQD			1		
System x Gen-II Universal Slides Kit	A5FW			1		
300GB 10K 6Gbps SAS 2.5" G3HS HDD	A4TL			2		
x3650 M5 8x 2.5" HS HDD Assembly Kit (Single RAID)	A5G6			1		
8GB TruDDR4 Memory PC4-17000 CL15 2133MHz LP RDIMM	A5B8			8		
Power Cable	6311			2		
ServeRAID M5210 SAS/SATA Controller	A3YZ			1		
Intel X540-T2 Dual Port 10GbBaseT Adapter	A2ED			1		
Intel X540 ML2 Dual Port 10GbBaseT Adapter	A40P			1		
x3650 M5 Riser Bracket	A5G5			2		
x3650 M5 Right EIA & EIA Plate	A5V5, A5G1			1		
ServicePac for 3-Year 24x7x4 Support (x3650 M4)	67567XR	1	790	1		790
				<b>Subtotal</b>	<b>20,101</b>	<b>790</b>
<b>Client Software</b>						
Windows Server 2012 R2 Standard Edition	P73-06284	2b	882	1	882	
				<b>Subtotal</b>	<b>882</b>	<b>0</b>
<b>Infrastructure</b>						
10 ft. Cat 6A Network Cable (2 spares)	CAT6A-10S	3	7	6	42	
				<b>Subtotal</b>	<b>42</b>	<b>0</b>
				<b>Total</b>	<b>1,837,288</b>	<b>16,616</b>
Dollar Volume Discount (See Note 1)	29.76%	1			261,684	
Microsoft Open Program Discount Schedule	16.67%	2b			735	
Pricing: 1 - Lenovo 1-877-782-7134; 2 - Microsoft; 3 - Newegg					<b>Three-Year Cost of Ownership USD:</b> \$1,591,485	
Note 1: Discount applies to all line items where Pricing=1; pricing is for these or similar quantities.					<b>TPC-E Throughput:</b> 11,058.99	
Discounts for similarly sized configurations will be similar to what is quoted here, but may vary based on the specific components priced.					<b>\$ USD/tpsE:</b> \$143.91	

Benchmark results and test methodology audited by Doug Johnson for InfoSizing, Inc. ([www.sizing.com](http://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](mailto:pricing@tpc.org). Thank you.



**System x3950 X6  
Microsoft SQL Server 2014**

**TPC-E 1.14.0  
TPC Pricing 2.0.0**

Report Date:  
December 17, 2015

Revision Date:  
December 17, 2015

Availability Date:  
December 17, 2015

Numerical Quantities Summary				
<b>Reported Throughput: 11,058.99 tpsE</b>		<b>Configured Customers: 5,750,000</b>		
Response Time (in seconds)	Minimum	Average	90 <sup>Th</sup> Percentile	Maximum
Broker-Volume	0.01	0.01	0.03	1.02
Customer-Position	0.01	0.01	0.02	2.96
Market-Feed	0.01	0.01	0.02	2.93
Market-Watch	0.01	0.01	0.02	1.00
Security-Detail	0.01	0.01	0.01	0.82
Trade-Lookup	0.01	0.05	0.08	1.03
Trade-Order	0.01	0.07	0.17	1.26
Trade-Result	0.01	0.01	0.04	0.97
Trade-Status	0.01	0.01	0.01	0.98
Trade-Update	0.01	0.05	0.08	1.05
Data-Maintenance	0.01	0.02	N/A	0.26
Transaction Mix		Transaction Count	Mix %	
Broker-Volume		39,013,149	4.900%	
Customer-Position		103,504,876	13.000%	
Market-Feed		7,962,484	1.000%	
Market-Watch		143,314,195	18.000%	
Security-Detail		111,466,412	14.000%	
Trade-Lookup		63,694,909	8.000%	
Trade-Order		80,414,718	10.100%	
Trade-Result		79,624,781	10.001%	
Trade-Status		151,275,822	19.000%	
Trade-Update		15,923,804	2.000%	
Data-Maintenance		120	N/A	
Test Duration and Timings				
Ramp-up Time (hh:mm:ss)				01:03:58
Measurement Interval (hh:mm:ss)				02:00:00
Business Recovery Time (hh:mm:ss)				00:41:16
Total Number of Transactions Completed in Measurement Interval				796,195,150

# **Table of Contents**

<b>Abstract</b> .....	<b>3</b>
<b>Table of Contents</b> .....	<b>7</b>
<b>Clause 0 – Preamble</b> .....	<b>9</b>
<b>Clause 1 – Introduction</b> .....	<b>11</b>
Benchmark Sponsor.....	11
Configuration Diagrams .....	11
Measured and Priced Configurations.....	11
Figure 1-1. Measured Configuration.....	12
Hardware and Software Configuration Steps .....	12
<b>Clause 2- Database Design, Scaling, and Population</b> .....	<b>13</b>
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 Configurations.....	16
Database Interface and Model Implemented .....	19
Database Load Methodology.....	19
<b>Clause 3 – Transaction Related Items</b> .....	<b>20</b>
Vendor-Supplied Code .....	20
Database Footprint of Transactions .....	20
<b>Clause 4 – SUT, Driver, and Network</b> .....	<b>21</b>
Network Configuration.....	21
<b>Clause 5 – EGen</b> .....	<b>22</b>
EGen Version .....	22
EGen Code and Modifications.....	22
EGen Files .....	22
<b>Clause 6 – Performance Metrics and Response Time</b> .....	<b>23</b>
EGen Instances .....	23
Reported Throughput.....	23
Throughput vs. Elapsed Time for Trade-Result Transaction.....	23
Figure 6-1. Test Run Graph .....	23
Steady State Methodology.....	23
Work Performed During Steady State .....	24
Transaction Statistics.....	24
Table 6-1. Transaction Statistics.....	25
<b>Clause 7 – Transaction and System Properties</b> .....	<b>26</b>
Atomicity Requirements.....	26
Consistency Requirements.....	26
Isolation Requirements .....	27
Durability Requirements.....	27
Durability Test for Data Accessibility .....	27
Table 7-1. Combinations of Durable Media Technologies Tested for Data Accessibility.....	27
Figure 7-1. Data Accessibility Graph.....	29
Durability Test for Business Recovery .....	29
Figure 7-2. Business Recovery Time Graph .....	31

<b>Clause 8 – Pricing</b> .....	<b>32</b>
60-Day Space.....	32
Table 8-1. Disk Space Requirements .....	32
Availability Date.....	33
Supporting Files Index.....	33
Auditor’s Attestation Letter.....	33
<b>Appendix A – Price Quotes</b> .....	<b>36</b>



# **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 M5, with:  
 - 2 x Intel® Xeon® Processor E5-2699 v3 2.30GHz (2 Procs/36 Cores/72 Threads)  
 - 64GB Memory  
 - 2 x 300GB SFF SAS (RAID-1)  
 - 1 x ServeRAID M5210  
 - Onboard Quad Gb Ethernet  
 - 2 x Dual 10Gb-T Ethernet

System x3950 X6, with:  
 - 8 x Intel Xeon Processor E7-8890 v3 2.50GHz (8 Procs/144 Cores/288 Threads)  
 - 4096GB Memory  
 - 2 x 900GB 10K SAS (RAID-1)  
 - 6 x 800GB SAS SSD (RAID-10)  
 - 1 x ServeRAID M5210  
 - 8 x ServeRAID M5225  
 - 4 x Dual 10Gb-T Ethernet

12 x E1024 JBOD with:  
 - 104 x 800GB 2.5" SAS SSD (6 x 17-drive RAID-5)  
 (1 x 2-drive RAID-1)  
 14 x EXP2524 JBOD with:  
 - 144 x 600GB 2.5" SAS HDD (6 x 24-drive RAID-10)  
 - 96 x 1200GB 2.5" SAS HDD (4 x 24-drive RAID-10)  
 - 48 x 200GB 2.5" SAS SSD (2 x 24-drive RAID-5)  
 - 28 x 400GB 2.5" SAS SSD (2 x 14-drive RAID-5)  
 420 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.

## 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\_x3650M5\_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\5750000.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 5,750,000 customers. The cardinality is shown in Table 2-1.

**Table 2-1. Initial Cardinality of Tables**

Table Name	Rows
ACCOUNT_PERMISSION	40,823,550
ADDRESS	8,625,004
BROKER	57,500
CASH_TRANSACTION	91,411,180,729
CHARGE	15
COMMISSION_RATE	240
COMPANY	2,875,000
COMPANY_COMPETITOR	8,625,000
CUSTOMER	5,750,000
CUSTOMER_ACCOUNT	28,750,000
CUSTOMER_TAXRATE	11,500,000
DAILY_MARKET	5,140,068,750
EXCHANGE	4
FINANCIAL	57,500,000
HOLDING	5,086,647,986
HOLDING_HISTORY	133,159,313,267
HOLDING_SUMMARY	285,955,273
INDUSTRY	102
LAST_TRADE	3,938,750
NEWS_ITEM	5,750,000
NEWS_XREF	5,750,000
SECTOR	12
SECURITY	3,938,750
SETTLEMENT	99,360,000,000
STATUS_TYPE	5
TAXRATE	320
TRADE	99,360,000,000
TRADE_HISTORY	238,463,746,183
TRADE_REQUEST	0
TRADE_TYPE	5
WATCH_ITEM	574,990,134
WATCH_LIST	5,750,000
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 900GB 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 was stored on six 800GB 2.5” SAS 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 six ServeRAID M5225 SAS/SATA controllers. Each of these controllers was connected to two Lenovo E1024 enclosures, one holding eight 800GB 2.5” SAS SSDs and the other holding nine 800GB 2.5” SAS SSDs. These 17 total SSDs per RAID adapter were used to create a RAID-5 array. Additionally, for one of these RAID adapters, each of the two attached enclosures held one additional 800GB 2.5” SAS SSD. These two SSDs were used to create a RAID-1 array to hold the run-time tempdb database.

In total, for database data, twelve Lenovo E1024 enclosures and 102 external SSDs were connected to the database server and were used to create six RAID-5 data arrays. Each data array was broken into seven partitions: three for fixed\_fg (RAW), three for growing\_fg (RAW), and one for extra space (NTFS). One partition was created on the run-time tempdb array and formatted NTFS.

In addition to the priced configuration described above, the measured configuration included fourteen additional external EXP2524 enclosures, all attached to two ServeRAID M5225 SAS/SATA controllers via daisy-chaining:

- Six with twenty-four 600GB SAS HDDs each, used for six 24-drive RAID-10 arrays (NTFS)
- Four with twenty-four 1200GB SAS HDDs each, used for four 24-drive RAID-10 arrays (NTFS)
- Two with twenty-four 200GB SAS SSDs each, used for two 24-drive RAID-5 arrays (NTFS)
- Two with fourteen 400GB SAS SSDs each, used for two 14-drive RAID-5 arrays (NTFS)

All of this additional space was used during the database load process to hold load-time tempdb, flat files, and database backups. This hardware- all fourteen of these enclosures and the drives contained in them- performed no function during benchmark runs.

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	Internal M5210	2 x 900GB SAS HDD internal RAID-1	C: (NTFS)	836.74GB	OS
1	Internal M5210	6 x 800GB SAS SSD internal RAID-10	E: (RAW) F: (NTFS)	1954.10GB 278.15GB	tpce log MDF



Disk #	Controller	Drives Enclosure RAID Level (Pricing)	Partition (File System)	Size	Use
2	M5225 #1	17 x 800GB SAS SSD E1024 RAID-5	c:\mplfx1 (RAW) c:\mplfx2 (RAW) c:\mplfx3 (RAW) c:\mplgw1 (RAW) c:\mplgw2 (RAW) c:\mplgw3 (RAW) c:\mplxt1 (NTFS)	68.85GB 68.85GB 68.85GB 3275.78GB 3275.78GB 3275.78GB 1872.00GB	fixed_fg fixed_fg fixed_fg growing_fg growing_fg growing_fg tempdb
3	M5225 #2	17 x 800GB SAS SSD E1024 RAID-5	c:\mplfx4 (RAW) c:\mplfx5 (RAW) c:\mplfx6 (RAW) c:\mplgw4 (RAW) c:\mplgw5 (RAW) c:\mplgw6 (RAW) c:\mplxt2 (NTFS)	68.85GB 68.85GB 68.85GB 3275.78GB 3275.78GB 3275.78GB 1872.00GB	fixed_fg fixed_fg fixed_fg growing_fg growing_fg growing_fg
4	M5225 #3	17 x 800GB SAS SSD E1024 RAID-5	c:\mplfx7 (RAW) c:\mplfx8 (RAW) c:\mplfx9 (RAW) c:\mplgw7 (RAW) c:\mplgw8 (RAW) c:\mplgw9 (RAW) c:\mplxt3 (NTFS)	68.85GB 68.85GB 68.85GB 3275.78GB 3275.78GB 3275.78GB 1872.00GB	fixed_fg fixed_fg fixed_fg growing_fg growing_fg growing_fg
5	M5225 #4	14 x 400GB SAS SSD EXP2524 RAID-5 (Measured)	c:\mplt1 (NTFS)	4830.64GB	tempdb
6	M5225 #4	24 x 1200GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mplbk7 (NTFS) c:\mplbk8 (NTFS)	6699.87GB 6699.87GB	backup & flatfiles
7	M5225 #4	24 x 1200GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mplbk13 (NTFS) c:\mplbk14 (NTFS)	6699.87GB 6699.87GB	backup & flatfiles
8	M5225 #4	24 x 1200GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mplbk11 (NTFS) c:\mplbk12 (NTFS)	6699.87GB 6699.87GB	backup & flatfiles
9	M5225 #4	14 x 400GB SAS SSD EXP2524 RAID-5 (Measured)	c:\mplt2 (NTFS)	4830.64GB	tempdb
10	M5225 #4	24 x 1200GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mplbk9 (NTFS) c:\mplbk10 (NTFS)	6699.87GB 6699.87GB	backup & flatfiles
11	M5225 #4	24 x 200GB SAS SSD EXP2524 RAID-5 (Measured)	c:\mplt3 (NTFS)	4262.53GB	tempdb
12	M5225 #4	24 x 200GB SAS SSD EXP2524 RAID-5 (Measured)	c:\mplt4 (NTFS)	4262.53GB	tempdb

Disk #	Controller	Drives Enclosure RAID Level (Pricing)	Partition (File System)	Size	Use
13	M5225 #5	17 x 800GB SAS SSD E1024 RAID-5	c:\mp\fx10 (RAW) c:\mp\fx11 (RAW) c:\mp\fx12 (RAW) c:\mp\gw10 (RAW) c:\mp\gw11 (RAW) c:\mp\gw12 (RAW) c:\mp\xt4 (NTFS)	68.85GB 68.85GB 68.85GB 3275.78GB 3275.78GB 3275.78GB 1872.00GB	fixed_fg fixed_fg fixed_fg growing_fg growing_fg growing_fg
14	M5225 #6	17 x 800GB SAS SSD E1024 RAID-5	c:\mp\fx13 (RAW) c:\mp\fx14 (RAW) c:\mp\fx15 (RAW) c:\mp\gw13 (RAW) c:\mp\gw14 (RAW) c:\mp\gw15 (RAW) c:\mp\xt5 (NTFS)	68.85GB 68.85GB 68.85GB 3275.78GB 3275.78GB 3275.78GB 1872.00GB	fixed_fg fixed_fg fixed_fg growing_fg growing_fg growing_fg
15	M5225 #7	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk6 (NTFS)	6694.21GB	Backup & flatfiles
16	M5225 #7	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk3 (NTFS)	6694.21GB	Backup & flatfiles
17	M5225 #7	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk1 (NTFS)	6694.21GB	Backup & flatfiles
18	M5225 #7	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk4 (NTFS)	6694.21GB	Backup & flatfiles
19	M5225 #7	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk2 (NTFS)	6694.21GB	Backup & flatfiles
20	M5225 #7	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk5 (NTFS)	6694.21GB	Backup & flatfiles
21	M5225 #8	17 x 800GB SAS SSD E1024 RAID-5	c:\mp\fx16 (RAW) c:\mp\fx17 (RAW) c:\mp\fx18 (RAW) c:\mp\gw16 (RAW) c:\mp\gw17 (RAW) c:\mp\gw18 (RAW) c:\mp\xt6 (NTFS)	68.85GB 68.85GB 68.85GB 3275.78GB 3275.78GB 3275.78GB 1872.00GB	fixed_fg fixed_fg fixed_fg growing_fg growing_fg growing_fg
22	M5225 #8	2 x 800GB SAS SSD E1024 RAID-1	T: (NTFS)	744.00GB	run-time 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 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 eight Gb Ethernet ports. Four of these are provided by the onboard Ethernet chip and the other four are provided by two dual-port PCI-e 10Gb Ethernet adapters.

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 all four of the 10Gb Ethernet adapter ports. These crossover networks, all running at 10Gb, 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 v1.14.0 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.

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 11,058.99 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
<b>Customer-Position</b>			
By Tax ID	1	50.00%	48% to 52%
Get History	1	50.00%	48% to 52%
<b>Market-Watch</b>			
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%
<b>Security-Detail</b>			
Access LOB	1	1.00%	0.9% to 1.1%
<b>Trade-Lookup</b>			
Frame to execute	1	30.00%	28.5% to 31.5%
	2	30.00%	28.5% to 31.5%
	3	30.00%	28.5% to 31.5%
	4	10.00%	9.5% to 10.5%
<b>Trade-Order</b>			
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	8.00%	7.5% to 8.5%
Rollback	1	0.99%	0.94% to 1.04%
LIFO	1	35.00%	33% to 37%
Trade Quantity	100	24.99%	24% to 26%
	200	25.00%	24% to 26%
	400	25.00%	24% to 26%
	800	25.00%	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	20.00%	19.8% to 20.2%
	Limit Sell	10.00%	9.9% to 10.1%
	Stop Loss	10.01%	9.9% to 10.1%
<b>Trade-Update</b>			
Frame to execute	1	33.01%	31% to 35%
	2	32.98%	31% to 35%
	3	34.00%	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 M5225
Database Log	SSD	SAS	RAID-10	ServeRAID M5210
Database tempdb	SSD	SAS	RAID-1	ServeRAID M5225


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 array by physically removing it from its enclosure. Since the database log 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 the tempdb array by physically removing it from its enclosure. Since the tempdb array is RAID protected, transaction processing continued.
7. Waited until the Data Accessibility Throughput Requirements were met again for at least 5 minutes.
8. Induced the third 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.
9. 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.
10. After a few minutes passed, a new drive was inserted into the tempdb enclosure to replace the failed tempdb drive. The tempdb array rebuilding process was started.
11. 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.
12. Continued running the benchmark for at least 20 minutes.
13. Terminated the run gracefully.
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 run, equaled the number of successful Trade-Result transactions reported by the Driver.
16. 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 00:15:10. The Application Recovery Time was 00:26:06. The Business Recovery Time, which is the sum of the Database Recovery Time and the Application Recovery Time, was 00:41:16.

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				
	5,750,000	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	57,500	4,616	6,392	550	11,558	57,500	11,008	-	-	550
CASH_TRANSACTION	91,411,180,729	9,521,348,952	20,081,192	477,071,507	10,018,501,651	91,531,806,747	9,564,285,704	22,855,560	55,517,936	55,517,936
CHARGE	15	8	8	1	17	15	16	-	-	1
COMMISSION_RATE	240	16	16	2	34	240	32	-	-	2
SETTLEMENT	99,360,000,000	4,737,869,208	10,009,240	237,393,922	4,985,272,370	99,491,119,266	4,760,584,040	12,705,592	30,862,873	30,862,873
TRADE	99,360,000,000	11,863,181,832	6,611,730,184	923,745,601	19,398,657,617	99,491,780,768	18,539,212,088	64,300,072	156,189,885	156,189,885
TRADE_HISTORY	238,463,746,183	7,171,842,784	18,702,896	359,527,284	7,550,072,964	238,778,695,340	7,213,219,120	22,673,440	55,075,553	55,075,553
TRADE_REQUEST	-	8	8	1	17	661,459	2,086,872	2,086,856	5,069,136	5,069,136
TRADE_TYPE	5	8	1,032	52	1,092	5	1,040	-	-	52
ACCOUNT_PERMISSION	40,823,550	2,247,848	11,952	112,990	2,372,790	40,823,550	2,259,920	120	292	112,990
CUSTOMER	5,750,000	942,256	274,768	60,851	1,277,875	5,750,000	1,217,072	48	117	60,851
CUSTOMER_ACCOUNT	28,750,000	2,605,136	635,872	162,050	3,403,058	28,750,000	3,240,992	-	-	162,050
CUSTOMER_TAXRATE	11,500,000	239,648	1,688	12,067	115,000,000	11,500,000	241,496	160	389	12,067
HOLDING	5,086,647,986	340,596,552	232,843,992	28,672,027	602,112,571	5,089,984,684	581,205,096	7,764,552	18,860,702	18,860,702
HOLDING_HISTORY	133,159,313,267	4,842,157,568	3,234,755,848	403,845,671	8,480,759,087	133,336,117,987	8,103,386,384	26,472,968	64,304,902	64,304,902
HOLDING_SUMMARY	285,955,273	12,559,104	46,608	630,286	13,235,998	285,955,946	12,605,712	-	-	-
WATCH_ITEM	574,990,134	16,173,864	59,464	811,666	17,044,994	574,990,134	16,233,632	304	739	811,666
WATCH_LIST	5,750,000	143,208	129,064	13,614	285,886	5,750,000	272,272	-	-	13,614
COMPANY	2,875,000	613,984	181,376	39,768	835,128	2,875,000	795,384	24	59	39,768
COMPANY_COMPETITOR	8,625,000	231,536	209,240	22,039	462,815	8,625,000	440,776	-	-	22,039
DAILY_MARKET	5,140,068,750	241,340,504	705,304	12,102,290	254,148,098	5,140,068,750	242,407,384	1,576	3,829	12,102,290
EXCHANGE	4	8	8	1	17	4	16	-	-	1
FINANCIAL	57,500,000	6,479,680	18,232	324,896	6,822,808	57,500,000	6,498,232	320	778	324,896
INDUSTRY	102	8	24	2	34	102	32	-	-	2
LAST_TRADE	3,938,750	245,584	1,472	12,353	259,409	3,938,750	247,056	-	-	12,353
NEWS_ITEM	5,750,000	623,405,000	6,704	31,170,585	654,582,289	5,750,000	623,411,736	32	78	31,170,585
NEWS_XREF	5,750,000	143,160	1,672	7,242	152,074	5,750,000	144,832	-	-	7,242
SECTOR	12	8	24	2	34	12	32	-	-	2
SECURITY	3,938,750	546,776	150,456	34,862	732,094	3,938,750	697,264	32	78	34,862
STATUS_TYPE	5	8	8	1	17	5	16	-	-	1
ADDRESS	8,625,004	497,432	1,784	24,961	524,177	8,625,004	499,264	48	117	24,961
TAXRATE	320	24	16	2	42	320	56	16	39	39
ZIP_CODE	14,741	488	104	30	622	14,741	592	-	-	30
<b>TOTALS (KB)</b>		<b>39,385,416,816</b>	<b>10,130,566,648</b>	<b>2,475,799,173</b>	<b>51,991,782,637</b>		<b>49,674,845,168</b>	<b>158,861,720</b>	<b>385,887,502</b>	<b>430,793,898</b>
Initial Database Size (MB)		48,355,453	47,222 GB							
Database Filegroups	LUN Count	Partition Size (MB)	MB Allocated	MB Loaded	MB Required					
growing_fg	18	3,354,300	60,377,400	47,478,248	47,855,085					OK
fixed_fg	18	70,400	1,267,200	877,205	921,065					OK
Settlements	131,119,266									
Data Space Required (MB)	Data Space Configured (MB)	Log Space Required (MB)	Log Space Configured (MB)							
Initial Growing Space	47,478,248	Initial Log Size	92,214	Log LUNS	1					
Final Growing Space	47,633,384	Data LUNS	6	Log Disks	6					
Delta	155,136	Disks per LUN	17	Disk Capacity	761,985					
Data Space per Trade	0.001183165	Disk Capacity	761,985	RAID Overhead	50%					
1 Day Data Growth	376,837	RAID Overhead	94%	0%	0%	0%	0.006739760	2,238,820	2,285,955	
60 Day Space	70,965,667	Total Space	73,150,562	1 Day Log Space	2,238,820					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 December 17, 2015.

## 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

December 14, 2015

I verified the TPC Benchmark™ E v1.14.0 performance of the following configuration:

Platform: System x3950 X6  
Operating System: Microsoft Windows Server 2012 R2 Standard Edition  
Database Manager: Microsoft SQL Server 2014 Enterprise Edition

The results were:

**Performance Metric** **11,058.99 tpsE**  
Trade-Result 90<sup>th</sup> %-tile 0.04 Seconds

**Tier B (Server)**

**System x3950 X6**

CPU	8 x Intel Xeon Processor E7-8890 v3 (2.50 GHz, 18-core, 45 MB L3)		
Memory	4096 GB		
Storage	<b>Qty</b>	<b>Size</b>	<b>Type</b>
	2	900 GB	10K rpm SAS HDD
	6	800 GB	SAS SSD
	104	800 GB	SAS SSD

**Tier A (Client)**

**System x3650 M5**

CPU	2 x Intel Xeon Processor E5-2699 v3 (2.30 GHz, 18-core, 45 MB L3)		
Memory	64 GB		
Storage	2 x 300 GB 10K rpm SAS 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 v1.14.0
- The transactions were correctly implemented
- The database was properly scaled and populated for 5,750,000 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 00:41:16 was correctly measured
- The 60-day storage requirement was correctly computed
- The system pricing was verified for major components and maintenance

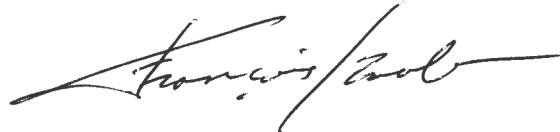
Additional Audit Notes:

None.

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**  
December 7, 2015

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
<b>Database Management System</b>				
7JQ-00750	<b>SQL Server 2014 Enterprise Edition</b> 2 Core License Open Program - Level C	\$13,472.50	72	\$970,020.00
<b>Database Server Operating System</b>				
P73-06284	<b>Windows Server 2012 R2 Standard Edition</b> 2 Processor License Open Program - Level C Unit Price reflects a 17% discount from the retail unit price of \$882.	\$735.00	4	\$2,940.00
<b>Tier-A Operating System(s)</b>				
P73-06284	<b>Windows Server 2012 R2 Standard Edition</b> 2 Processor License Open Program - Level C Unit Price reflects a 17% discount from the retail unit price of \$882.	\$735.00	1	\$735.00
<b>Support</b>				
N/A	<b>Microsoft Problem Resolution Services</b> Professional Support (1 Incident).	\$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\_qhtplylGYLKTVUKf44789fftr\_2015\_Ire.

TRY **PREMIER** ▶ Free 30-day trial when you sign up for a 12-month membership OR Download mobile app & get a free 3-month membersh...

Log in or Register 🛒 \$6.49 (1 Items) ♥ Wish List 🔍 Customer Service ▼

  **newegg Business**

ALL PRODUCTS ▼ DEALS & SERVICES ▼ FEATURED SELLERS ▼

Home > *My Shopping Cart***SHOPPING SATISFACTION**

Nearly all of Newegg.com's offerings are covered by a full 30-day refund policy.


**PRIVACY PROTECTED**

Our Privacy Policy states that we will only share your personal information with third parties working on our behalf to complete your order such as UPS and FedEx.

**SECURITY SAFEGUARDS**

Security is a top priority at Newegg. When you submit sensitive information via the website, your information is protected both online and offline.

ADVERTISEMENT **1 Item(s) in your shopping cart: \$6.49** Select All**Newegg Standard Shipping Service** [Important Shipping Information](#)

<input type="checkbox"/>	 <b>Kaybles 10ft CAT6A-10S 10 ft. Cat 6A Blue Color Shielded Stranded STP Network Cable Blue Color 10 feet - OEM</b> <a href="#">Extended Holiday Return Policy</a>	1 IN STOCK	<b>\$6.49</b>
<input type="checkbox"/> Send this item as a gift			

**Subtotal:** \$6.49**Shipping Options**Zip/Postal Code  

Choose a delivery method

- APO/FPO-Military ONLY
- Super Eggsaver (4-7 Business Days)**
- Newegg 3 Business Days
- Newegg 2 Business Days
- Newegg Next Business Day

**Will Call Service****Shipping:**  
\$0.00

Your products can be picked up at our warehouse in City Of Industry, CA.

You will be notified by E-mail once your products are ready for pick-up.

**Apply Promo Codes**

Promo codes have usage limitations that require you to enter your email address to enable Promo Codes.

E-Mail Address 

We [value your privacy](#). You won't get any emails from us unless you place an order or [sign up for our newsletter](#).

**Redeem Newegg Gift Cards / Google Offer Code**Card Number  Security Code **Redeem Eggpoints****6 OR 12 MONTHS SPECIAL FINANCING\***

- 6 months offer applies to purchases of **\$249-\$498.99**.
- 12 months offer applies to purchases of **\$499** or more.

\*Subject to credit approval. Minimum monthly payments required. See site for details.

[Learn More](#) ▶**Grand Total**

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

**Having problems with your cart?**

[Check out the FAQ](#) for help or try [emptying your cart](#) to start over.