TPC Benchmark™ E
Full Disclosure Report
for
System x® 3850 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
May 5, 2015

First Edition – May 2015

Trademarks

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®3850 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 x3850 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 x3850 X6	Microsoft SQL Server 2014 Enterprise Edition Microsoft Windows Server 2012 Standard Edition	$1,713,135 USD	6964.75	$245.98 USD	July 31, 2015

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.

lenovo. FOR THOSE WHO DO.	**System x® 3850 X6 Microsoft® SQL Server® 2014**		**TPC-E™ 1.13.0 TPC Pricing 2.0.0** Report Date: May 5, 2015 Revision Date: May 6, 2015
TPC-E Throughput **6964.75 tpsE**	Price/Performance **$245.98 USD per tpsE™**	Availability Date **July 31, 2015**	Total System Cost **$1,713,135 USD**
Database Server Configuration			
Operating System **Microsoft Windows Server® 2012 Standard Edition**	Database Manager **Microsoft SQL Server 2014 Enterprise Edition**	Processors/Cores/ Threads **4/72/144**	Memory **4096GB**

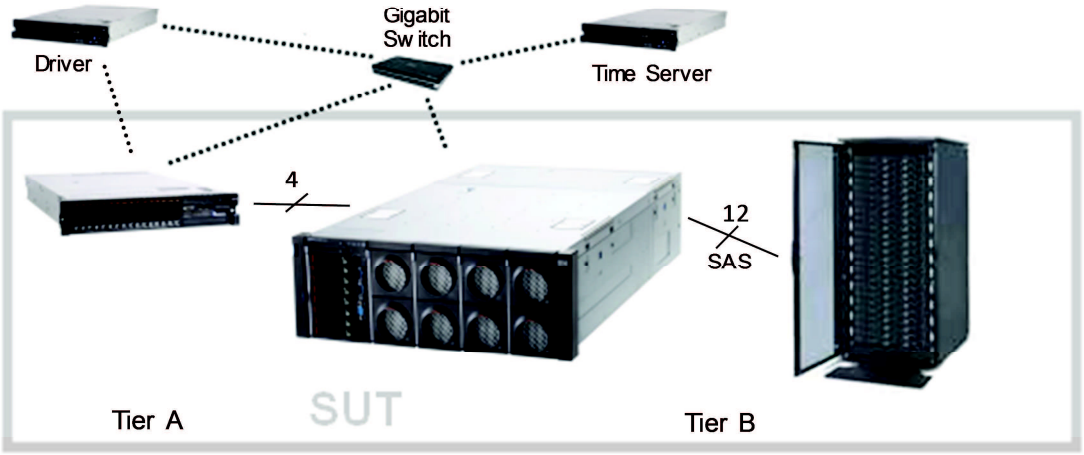

System x3650 M4, with:
- 2 x Intel® Xeon® Processor E5-2697 v2 2.70GHz (2 Procs/24 Cores/48 Threads)
- 32GB Memory
- 2 x 300GB SFF SAS (RAID-1)
- 1 x ServeRAID M5110e
- Onboard Quad Gb Ethernet
- Dual Port Gb Ethernet

System x3850 X6, with:
- 4 x Intel Xeon Processor E7-8890 v3 2.50GHz (4 Procs/72 Cores/144 Threads)
- 4096GB Memory
- 2 x 300GB 10K SAS (RAID-1)
- 6 x 800GB SATA SSD (RAID-10)
- 1 x ServeRAID M5210
- 6 x ServeRAID M5225
- 4 x Dual 10Gb-T Ethernet

12 x EXP2524 JBOD Enclosures, with:
- 223 x 400GB 2.5" SAS SSD
(3 x 19-drive RAID-5 DB data)
(5 x 22-drive RAID-5 DB data)
(4 x 14-drive RAID-5 DB data)

223 Total External Drives

Initial Database Size **28,734 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** **223 x 400GB 2.5" SAS SSD**

System x3850 X6
Microsoft SQL Server 2014

TPC-E 1.13.0

TPC Pricing 2.0.0

Report Date:
May 5, 2015

Revision Date:
May 6, 2015

Availability Date:
July 31, 2015

Description	Part Number	Price Source	Unit Price	Quantity	Extended Price	3-Yr. Maint. Price
Server Hardware						
System x3850 X6 Configure-To-Order, includes:	6241AC1	1	532,162	1	532,162	
x3850 X6 4U Chassis + Midplane	ASMH, A4A4			1		
X6 DDR4 Compute Book Intel Xeon E7-8890 v3	AS91			4		
X6 Primary I/O Book + X6 Storage Book	ASFB, A4A1			1		
X6 Half-length I/O Book	A4A2			2		
4x 2.5" HS SAS/SATA/SSD HDD Backplane	A4A6			2		
ServeRAID M5210 SAS/SATA Controller for System x	A3YZ			1		
Intel X540 ML2 Dual Port 10GbaseT Adapter for System x	A40P			1		
Intel X540-T2 Dual Port 10GBase-T Adapter for System x	A2ED			3		
1400W HE Redundant Power Supply	A54E			4		
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			4		
64GB TruDDR4 PC4-17000 CL15 2133MHz LP LRDIMM	A5UK			64		
ServeRAID M5225-2GB SAS/SATA Controller for System x	00AE938	1	899	6	5,394	
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	110	1	110	
ServicePac for 3-Year 24x7x4 Support (x3850 X6)	67568BU	1	1,500	1		1500
				Subtotal	537,763	1,500
Server Storage						
S2 42U Standard Rack	93074RX	1	1,565	1	1,565	
EXP2524 Storage Enclosure	610024X	1	3,999	12	47,988	
1.5m SAS Cable (mSAS HD to mSAS)	00MJ163	1	129	12	1,548	
400GB 2.5 Inch Flash Drive	00NC575	1 - S	6,279	223	1,400,217	
ServicePac for 3-Year 24x7x4 Support (EXP2524)	91Y5785	1	1,200	12		14,400
ServicePac for 3-Year 24x7x4 Support (Rack)	41L2760	1	315	1		315
				Subtotal	1,451,318	14,715
Server Software						
SQL Server 2014 Enterprise Edition	7JQ-00750	2a	13,472.50	36	485,010	
Windows Server 2012 Standard Edition	P73-05761	2	882	2	1,764	
Microsoft Problem Resolution Services	N/A	2a	259	1		259
				Subtotal	486,774	259
Client Hardware						
System x3650 M4 Configure-To-Order, includes:	7915AC1	1	12,436	1	12,436	
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		
300GB 10K 6Gbps SAS 2.5" SFF G2HS HDD	A2XC			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	12,436	790
Client Software						
Windows Server 2012 Standard Edition	P73-05761	2	882	1	882	
				Subtotal	882	0
Infrastructure						
Ethernet Cables	78004256	1	6	4	24	
				Subtotal	24	0
				Total	2,489,197	17,264
Dollar Volume Discount (See Note 1)	39.28%	1			792,885	
Microsoft Open Program Discount Schedule	16.67%	2			441	

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:	$1,713,135
TPC-E Throughput:	6,964.75
$ USD/tpsE:	$245.98

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.

Numerical Quantities Summary				
Reported Throughput: 6964.75 tpsE	**Configured Customers: 3,500,000**			
Response Time (in seconds)	**Minimum**	**Average**	**90Th Percentile**	**Maximum**
Broker-Volume	0.01	0.01	0.02	0.29
Customer-Position	0.01	0.01	0.01	3.69
Market-Feed	0.01	0.01	0.02	3.68
Market-Watch	0.01	0.01	0.01	0.26
Security-Detail	0.01	0.01	0.01	0.54
Trade-Lookup	0.01	0.05	0.08	0.59
Trade-Order	0.01	0.06	0.12	0.48
Trade-Result	0.01	0.01	0.02	0.48
Trade-Status	0.01	0.01	0.01	0.55
Trade-Update	0.01	0.05	0.09	0.51
Data-Maintenance	0.01	0.01	N/A	0.17

Transaction Mix	**Transaction Count**	**Mix %**
Broker-Volume	24,571,013	4.900%
Customer-Position	65,188,338	13.000%
Market-Feed	5,014,628	1.000%
Market-Watch	90,260,609	18.000%
Security-Detail	70,202,478	14.000%
Trade-Lookup	40,115,332	8.000%
Trade-Order	50,645,868	10.100%
Trade-Result	50,146,204	10.001%
Trade-Status	95,275,516	19.000%
Trade-Update	10,028,918	2.000%
Data-Maintenance	120	N/A

Test Duration and Timings	
Ramp-up Time (hh:mm:ss)	00:21:50
Measurement Interval (hh:mm:ss)	02:00:00
Business Recovery Time (hh:mm:ss)	00:23:18
Total Number of Transactions Completed in Measurement Interval	501,448,904

Table of Contents

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

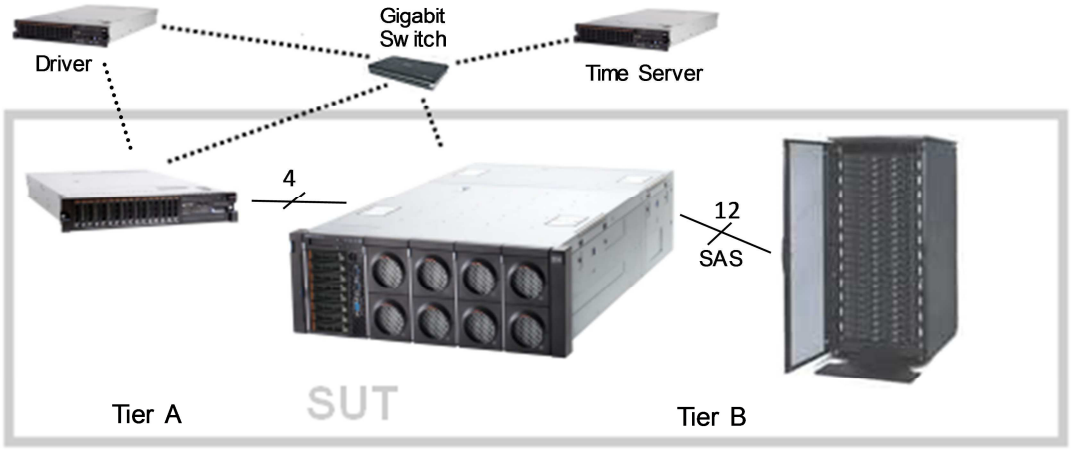

System x3650 M4, with:
- 2 x Intel® Xeon® Processor E5-2697 v2 2.70GHz (2 Procs/24 Cores/48 Threads)
- 32GB Memory
- 2 x 300GB SFF SAS (RAID-1)
- 1 x ServeRAID M5110e
- Onboard Quad Gb Ethernet
- Dual Port Gb Ethernet

System x3950 X6, with:
- 4 x Intel Xeon Processor E7-8890 v3 2.50GHz (4 Procs/72 Cores/144 Threads)
- 4096GB Memory
- 2 x 300GB 10K SAS (RAID-1)
- 6 x 800GB SATA SSD (RAID-10)
- 1 x ServeRAID M5210
- 6 x ServeRAID M5225
- 4 x Dual 10Gb-T Ethernet

22 x EXP2524 JBOD Enclosures, with:
- 56 x 400GB 2.5" SAS SSD (4 x 14-drive RAID-5 DB data)
- 167 x 200GB 2.5" SAS SSD (3 x 19-drive RAID-5 DB data) (5 x 22-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)

463 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 twelve 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, and five of these arrays used 22 older 200GB SSDs each. All these 200GB SSDs were priced one-for-one as 400GB SSDs. This substitution was allowed based on performance measurements, taken while the configuration was running, that show that the arrays using the 400GB SSDs were faster than the arrays using either of the 200GB SSDs. Table 1-1 shows these measurements.

Table 1-1. Durable Media Substitution Information

Description	14-Drive Array	22-Drive Array	19-Drive Array
Drives Used	SSD	Older SSD	SSD
Drive Capacity	400GB	200GB	200GB
Interface Type	SAS	SAS	SAS
Interface Speed	6 Gbps	3 Gbps	6 Gbps
Drives Per Array	14	22	19
Avg Array Response Times (read / write, ms)	0.21 / 0.52	0.50 / 1.20	0.22 / 0.55
Avg Array Queue Depth	8.2	21.5	9.4
Priced As	14 400GB SSDs	22 400GB SSDs	19 400GB 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_x3850X6_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\3500000.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 3,500,000 customers. The cardinality is shown in Table 2-1.

Table 2-1. Initial Cardinality of Tables

Table Name	Rows
ACCOUNT_PERMISSION	24,851,082
ADDRESS	5,250,004
BROKER	35,000
CASH_TRANSACTION	55,641,580,672
CHARGE	15
COMMISSION_RATE	240
COMPANY	1,750,000
COMPANY_COMPETITOR	5,250,000
CUSTOMER	3,500,000
CUSTOMER_ACCOUNT	17,500,000
CUSTOMER_TAXRATE	7,000,000
DAILY_MARKET	3,128,737,500
EXCHANGE	4
FINANCIAL	35,000,000
HOLDING	3,096,806,241
HOLDING_HISTORY	81,053,233,799
HOLDING_SUMMARY	174,068,725
INDUSTRY	102
LAST_TRADE	2,397,500
NEWS_ITEM	3,500,000
NEWS_XREF	3,500,000
SECTOR	12
SECURITY	2,397,500
SETTLEMENT	60,480,000,000
STATUS_TYPE	5
TAXRATE	320
TRADE	60,480,000,000
TRADE_HISTORY	145,151,921,464
TRADE_REQUEST	0
TRADE_TYPE	5
WATCH_ITEM	350,010,375
WATCH_LIST	3,500,000
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 six ServeRAID M5225 SAS/SATA controllers. Each of these controllers was connected to two EXP2524 enclosures for database data:

- Five of these enclosures held 22 200GB SAS SSDs each → 5 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, twelve enclosures and 223 external SSDs were connected to the database server and were used to create twelve 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 M5225 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	M5225 #1	19 x 200GB SAS SSD EXP2524 RAID-5 (Measured)	c:\mp\fx9 c:\mp\gw9 c:\mp\xt9	65.33GB 3108.40GB 162.14GB	fixed_fg growing_fg tempdb
1	M5225 #1	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk3 (NTFS)	6694.21GB	backup & flatfiles
2	M5225 #1	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk1 (NTFS)	6694.21GB	backup & flatfiles

Disk #	Controller	Drives Enclosure RAID Level (Pricing)	Partition (File System)	Size	Use
3	M5225 #1	19 x 200GB SAS SSD EXP2524 RAID-5 (Measured)	c:\mp\fx10 c:\mp\gw10 c:\mp\xt10	65.33GB 3108.40GB 162.14GB	fixed_fg growing_fg tempdb
4	Internal M5210	2 x 300GB SAS HDD internal RAID-1	C: (NTFS)	277.95GB	OS
5	Internal M5210	6 x 800GB SATA SSD internal RAID-10	E: (RAW) F: (NTFS)	1954.10GB 278.15GB	tpce log MDF tempdb
6	M5225 #2	22 x 200GB SAS SSD EXP2524 RAID-5 (Measured)	c:\mp\fx11 c:\mp\gw11 c:\mp\xt11	65.33GB 3108.40GB 718.13GB	fixed_fg growing_fg tempdb
7	M5225 #2	22 x 200GB SAS SSD EXP2524 RAID-5 (Measured)	c:\mp\fx12 c:\mp\gw12 c:\mp\xt12	65.33GB 3108.40GB 718.13GB	fixed_fg growing_fg tempdb
8	M5225 #2	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk5 (NTFS)	6694.21GB	backup & flatfiles
9	M5225 #3	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk6 (NTFS)	6694.21GB	backup & flatfiles
10	M5225 #3	22 x 200GB SAS SSD EXP2524 RAID-5 (Measured)	c:\mp\fx5 c:\mp\gw5 c:\mp\xt5	65.33GB 3108.40GB 718.13GB	fixed_fg growing_fg tempdb
11	M5225 #3	22 x 200GB SAS SSD EXP2524 RAID-5 (Measured)	c:\mp\fx6 c:\mp\gw6 c:\mp\xt6	65.33GB 3108.40GB 718.13GB	fixed_fg growing_fg tempdb
12	M5225 #4	19 x 200GB SAS SSD EXP2524 RAID-5 (Measured)	c:\mp\fx4 c:\mp\gw4 c:\mp\xt4	65.33GB 3108.40GB 162.14GB	fixed_fg growing_fg tempdb
13	M5225 #4	22 x 200GB SAS SSD EXP2524 RAID-5 (Measured)	c:\mp\fx3 c:\mp\gw3 c:\mp\xt3	65.33GB 3108.40GB 718.13GB	fixed_fg growing_fg tempdb
14	M5225 #4	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk4 (NTFS)	6694.21GB	backup & flatfiles
15	M5225 #4	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk2 (NTFS)	6694.21GB	backup & flatfiles
16	M5225 #5	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

Disk #	Controller	Drives Enclosure RAID Level (Pricing)	Partition (File System)	Size	Use
17	M5225 #5	24 x 1200GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk7 (NTFS) c:\mp\bk8 (NTFS)	6699.87GB 6699.87GB	backup & flatfiles
18	M5225 #5	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
19	M5225 #5	24 x 1200GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk11 (NTFS) c:\mp\bk12 (NTFS)	6699.87GB 6699.87GB	backup & flatfiles
20	M5225 #6	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
21	M5225 #6	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
22	M5225 #6	24 x 1200GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk13 (NTFS) c:\mp\bk14 (NTFS)	6699.87GB 6699.87GB	backup & flatfiles
23	M5225 #6	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 960 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 6964.75 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

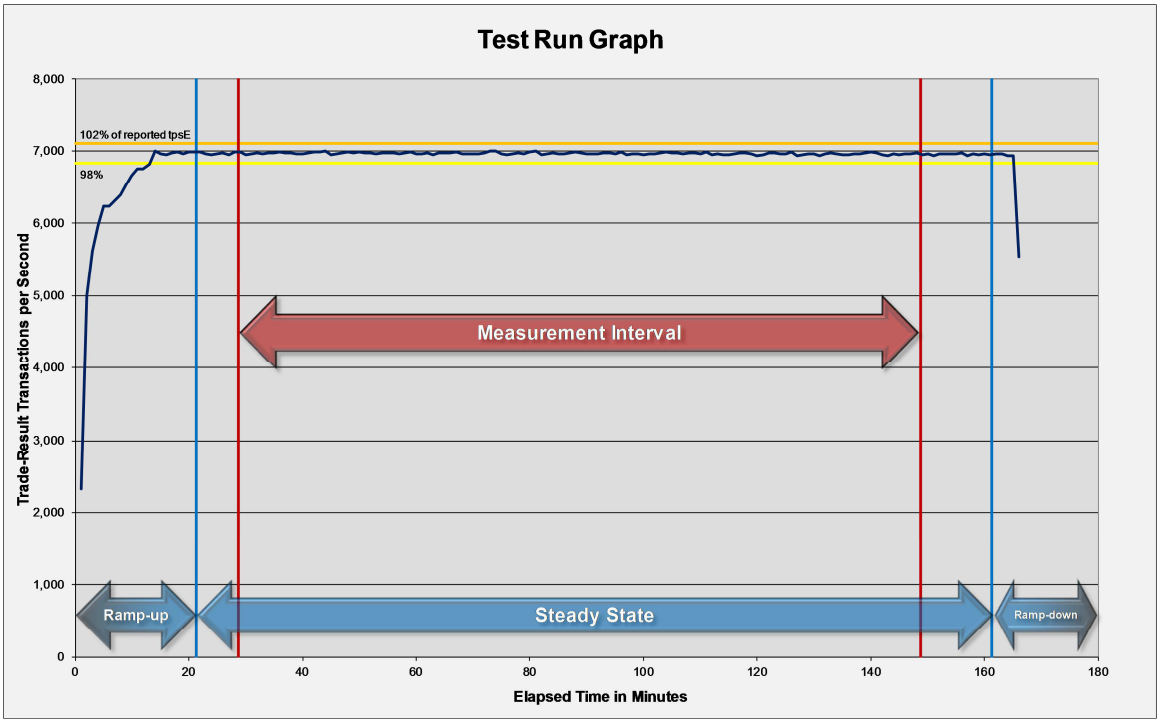

Steady State Methodology

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	49.99%	48% to 52%
Get History	1	50.00%	48% to 52%
Market-Watch			
Securities chosen by	Watch List	60.01%	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	30.00%	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%
Trade-Order			
Transactions requested by a third party		10.00%	9.5% to 10.5%
By Company Name		39.99%	38% to 42%
Buy On Margin	1	8.01%	7.5% to 8.5%
Rollback	1	0.99%	0.94% to 1.04%
LIFO	1	34.99%	33% to 37%
Trade Quantity	100	25.00%	24% to 26%
	200	25.00%	24% to 26%
	400	25.00%	24% to 26%
	800	25.00%	24% to 26%
Trade Type	Market Buy	30.00%	29.7% to 30.3%
	Market Sell	30.00%	29.7% to 30.3%
	Limit Buy	20.01%	19.8% to 20.2%
	Limit Sell	10.00%	9.9% to 10.1%
	Stop Loss	10.00%	9.9% to 10.1%
Trade-Update			
Frame to execute	1	33.00%	31% to 35%
	2	33.01%	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 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

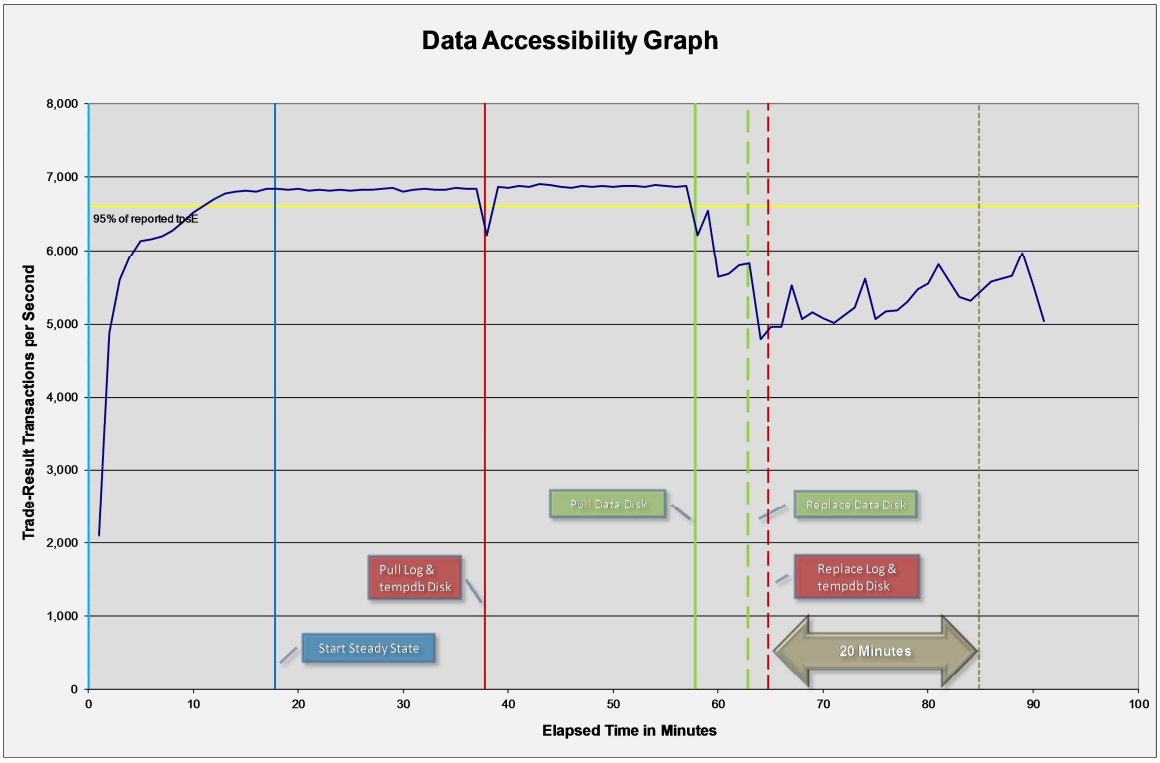

The files related to this data accessibility test are located in SupportingFiles\Clause7\Durability\DataAccessibility.

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 00:11:11. The Application Recovery Time was 00:12:07. The Business Recovery Time, which is the sum of the Database Recovery Time and the Application Recovery Time, was 00:23:18.

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

Figure 7-2. Business Recovery Time Graph

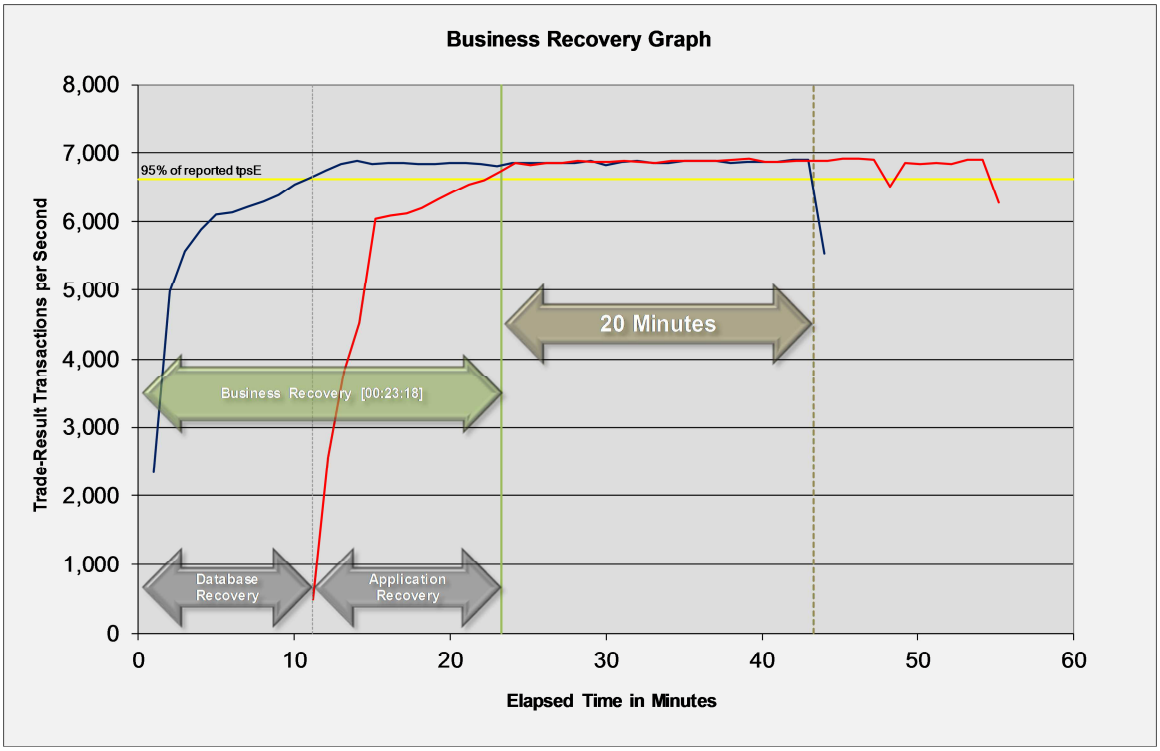

The files related to this business recovery test are located in SupportingFiles\Clause7\Durability\BusinessRecovery.

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

Customers	3,500,000	Measured Throughput	6964.75	Trade-Results/s	Reported Throughput	6964.75	tpsE

Table	Initial Rows	Data Size (KB)	Index Size (KB)	Extra 5% (KB)	Total + 5% (KB)	Rows After	After Run (KB)	Growth (KB)	Bus. Day Growth (KB)	Req. Add. (KB)
BROKER	35,000	2,552	2,712	263	5,527	35,000	5,264	-	-	263
CASH_TRANSACTION	55,641,580,672	5,788,375,816	12,200,304	290,028,806	6,090,604,926	55,727,895,643	5,816,986,496	16,410,376	35,082,982	35,082,982
CHARGE	15	8	8	1	17	15	16	-	-	1
COMMISSION_RATE	240	16	16	2	34	240	32	-	-	2
SETTLEMENT	60,480,000,000	2,883,920,064	6,083,008	144,500,154	3,034,503,226	60,573,825,321	2,899,143,368	9,140,296	19,540,615	19,540,615
TRADE	60,480,000,000	7,220,657,888	4,022,529,560	562,159,372	11,805,346,820	60,575,059,816	11,285,062,696	41,875,248	89,523,150	89,523,150
TRADE_HISTORY	145,151,921,464	4,365,471,792	11,382,176	218,842,698	4,595,696,666	145,379,228,704	4,393,165,872	16,311,904	34,872,463	34,872,463
TRADE_REQUEST	-	-	-	-	-	418,253	1,109,752	1,109,752	2,372,488	2,372,488
TRADE_TYPE	5	8	1,032	52	1,092	5	1,040	-	-	52
ACCOUNT_PERMISSION	24,851,082	1,368,392	8,464	68,843	1,445,699	24,851,082	1,376,928	72	154	68,843
CUSTOMER	3,500,000	573,528	170,728	37,213	781,469	3,500,000	744,264	8	18	37,213
CUSTOMER_ACCOUNT	17,500,000	1,585,808	389,960	98,788	2,074,556	17,500,000	1,975,768	-	-	98,788
CUSTOMER_TAXRATE	7,000,000	145,944	2,032	7,399	155,375	7,000,000	148,152	176	377	7,399
HOLDING	3,096,806,241	207,192,704	141,649,664	17,442,118	366,284,486	3,099,192,931	354,439,832	5,597,464	11,966,559	11,966,559
HOLDING_HISTORY	81,053,233,799	2,947,390,728	1,968,971,048	245,818,089	5,162,179,865	81,179,751,865	4,935,715,192	19,353,416	41,374,771	41,374,771
HOLDING_SUMMARY	174,068,725	7,637,672	29,376	383,352	8,050,400	174,069,258	7,667,048	-	-	-
WATCH_ITEM	350,010,375	9,834,224	37,240	493,573	10,365,037	350,010,375	9,871,776	312	668	493,573
WATCH_LIST	3,500,000	87,200	80,632	8,392	176,224	3,500,000	167,832	-	-	8,392
COMPANY	1,750,000	373,568	113,624	24,360	511,552	1,750,000	487,224	32	69	24,360
COMPANY_COMPETITOR	5,250,000	140,944	129,480	13,521	283,945	5,250,000	270,424	-	-	13,521
DAILY_MARKET	3,128,737,500	146,800,640	430,024	7,361,533	154,592,197	3,128,737,500	147,232,336	1,672	3,575	7,361,533
EXCHANGE	4	8	8	1	17	4	16	-	-	1
FINANCIAL	35,000,000	3,944,160	12,280	197,822	4,154,262	35,000,000	3,956,832	392	839	197,822
INDUSTRY	102	8	24	2	34	102	32	-	-	2
LAST_TRADE	2,397,500	149,496	2,048	7,577	159,121	2,397,500	151,544	-	-	7,577
NEWS_ITEM	3,500,000	379,463,848	5,584	18,973,472	398,442,904	3,500,000	379,469,488	56	120	18,973,472
NEWS_XREF	3,500,000	87,200	2,032	4,462	93,694	3,500,000	89,232	-	-	4,462
SECTOR	12	8	24	2	34	12	32	-	-	2
SECURITY	2,397,500	332,600	94,160	21,338	448,098	2,397,500	426,808	48	103	21,338
STATUS_TYPE	5	8	8	1	17	5	16	-	-	1
ADDRESS	5,250,004	302,808	2,192	15,250	320,250	5,250,004	305,072	72	154	15,250
TAXRATE	320	24	16	2	42	320	56	16	35	35
ZIP_CODE	14,741	488	96	29	613	14,741	584	-	-	29
TOTALS (KB)		**23,965,840,152**	**6,164,329,560**	**1,506,508,486**	**31,636,678,198**		**30,239,971,024**	**109,801,312**	**234,739,140**	**262,066,957**

Initial Database Size (MB) **29,423,994** **28,734 GB**

Database Filegroups	LUN Count	Partition Size (MB)	MB Allocated	MB Loaded	MB Required	
	0	-	-	-	-	OK
growing_fg	12	3,062,600	36,751,200	28,890,129	29,119,360	OK
	0	-	-	-	-	OK
fixed_fg	12	64,300	771,600	533,865	560,558	OK

Settlements **93,825,321**

Data Space Required (MB)		Data Space Configured (MB)					Log Space Required (MB)		Log Space Configured (MB)	
Initial Growing Space	28,890,129									
Final Growing Space	28,997,354	Data LUNS	3	5	4	-	Initial Log Size	94,084	Log LUNS	1
Delta	107,225	Disks per LUN	19	22	14	-	Final Log Size	726,244	Log Disks	6
Data Space per Trade	0.001142816	Disk Capacity	189,781	189,781	380,516	-	Log Growth	632,160	Disk Capacity	761,985
1 Day Data Growth	229,231	RAID Overhead	95%	95%	93%	0%	Log Growth/Trade	0.006737630	RAID Overhead	50%
60 Day Space	**43,177,882**	Total Space				**49,962,028**	1 Day Log Space	**1,445,550**	Log Space	**2,285,955**
						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 July 31, 2015. The dates for ordering and availability are detailed in Table 8-2 for those components that are not immediately orderable.

Table 8-2. Ordering and Pricing Information

Description	Part Number	Order Date	Availability Date	Order Method	Price Verification
X6 DDR4 Compute Book Intel Xeon Processor E7-8890 v3 18C 2.5GHz 165W	00ML982 (AS91)	7-31-15	7-31-15	See note 1	See note 2
x3850/x3950 X6 I/O Planar	00MY848 (ASFB)	7-31-15	7-31-15	See note 1	See note 2

Note 1: Lenovo 1-877-782-7134

Note 2: These components are not immediately orderable. For price verification before the order date, call Lenovo 1-877-782-7134.

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.

InfoSizing
The Right Metric For Sizing IT

TPC Transaction Processing Performance Council
Certified Auditors

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

April 19, 2015

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

Platform:	System x3850 X6
Operating System:	Microsoft Windows Server 2012 Standard Edition
Database Manager:	Microsoft SQL Server 2014 Enterprise Edition

The results were:

Performance Metric	**6964.75 tpsE**
Trade-Result 90th %-tile	0.02 Seconds

Tier B (Server)	**System x3850 X6**		
CPUs	4 x Intel Xeon Processor E7-8890 v3 (2.50 GHz, 18-core, 45 MB L3)		
Memory	4096 GB		
Storage	***Qty***	***Size***	***Type***
	2	300 GB	10K rpm SAS HDD
	6	800 GB	SATA SSD
	223	400 GB	SAS SSD

Tier A (Client)	**System x3650 M4**
CPUs	2 x Intel Xeon Processor E5-2697 v2 (2.70 GHz, 12-core, 30 MB L3)
Memory	32 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 v1.13.0
- The transactions were correctly implemented
- The database was properly scaled and populated for 3,500,000 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 00:23:18 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 (110) XceedIOPS 200 GB SAS SSD disks and (57) Optimus 200 GB SAS SSD disks that were substituted by (167) Optimus 400 GB 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

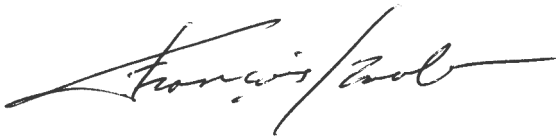

François Raab, President

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

April 13, 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
Database Management System				
7JQ-00750	**SQL Server 2014 Enterprise Edition** *2 Core License* *Open Program - Level C*	$13,472.50	36	$485,010.00
Database Server Operating System				
P73-05761	**Windows Server 2012 Standard Edition** *2 Processor License* *Open Program - Level C* *Unit Price reflects a 17% discount from the retail unit price of $882.*	$735.00	2	$1,470.00
Tier-A Operating System(s)				
P73-05761	**Windows Server 2012 Standard Edition** *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
Support				
N/A	**Microsoft Problem Resolution Services** *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_qhtplylGYLKTVUKf28479lsje_2015_lre.