

TPC Benchmark™ E
Full Disclosure Report
for
IBM® System x® 3850 X6
using
Microsoft® SQL Server® 2014
Enterprise Edition
and
Microsoft Windows Server® 2012
Standard Edition

TPC-E™ Version 1.12.0



First Edition
Submitted for Review
February 18, 2014

First Edition – February 2014

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

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

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

© Copyright International Business Machines Corporation 2014. All rights reserved.

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

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

Trademarks

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

The following terms used in this publication are trademarks of other companies as follows: TPC Benchmark, TPC-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

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

² When referring to hard disk capacity, GB, or gigabyte, means one thousand million bytes. Total user-accessible capacity may be less.

Abstract

IBM Corporation conducted the TPC Benchmark™ E on the IBM® System x®3850 X6 configured as a client/server system. This report documents the full disclosure information required by the TPC Benchmark E Standard Specification, Revision 1.12.0, including the methodology used to achieve the reported results. All testing fully complied with this revision level.



The software used on the IBM System x3850 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
IBM System x3850 X6	Microsoft SQL Server 2014 Enterprise Edition Microsoft Windows Server 2012 Standard Edition	\$1,052,173 USD	5576.27	\$188.69 USD	April 15, 2014

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

	IBM® System x® 3850 X6 Microsoft® SQL Server® 2014		TPC-E™ 1.12.0 TPC Pricing 1.7.0
			Report Date: February 18, 2014
			Revision Date: February 18, 2014
TPC-E Throughput 5576.27 tpsE	Price/Performance \$188.69 USD per tpsE™	Availability Date April 15, 2014	Total System Cost \$1,052,173 USD
Database Server Configuration			
Operating System Microsoft Windows Server® 2012 Standard Edition	Database Manager Microsoft SQL Server 2014 Enterprise Edition	Processors/Cores/Threads 4/60/120	Memory 2048GB
			
IBM System x3650 M4, with: - 2 x Intel® Xeon® Processor E5-2697 v2 2.70GHz (2 Procs/24 Cores/48 Threads) - 32GB Memory - 2 x 250GB SFF SATA (RAID-1) - 1 x IBM ServeRAID M5110e - Onboard Quad Gb Ethernet - Dual Port Gb Ethernet	IBM System x3850 X6, with: - 4 x Intel Xeon Processor E7-4890 v2 2.80GHz (4 Procs/60 Cores/120 Threads) - 2048GB Memory - 2 x 300GB 10K SAS (RAID-1) - 4 x 800GB SATA SSD (RAID-10) - 1 x IBM ServeRAID M5210 - 5 x IBM ServeRAID M5120 - 4 x Dual 10Gb-T Ethernet	9 x IBM System Storage® EXP2524 JBOD Storage Enclosures, with: - 208 x 200GB 2.5" SAS SSD (2 x 20-drive RAID-5 DB data) (7 x 24-drive RAID-5 DB data) - 2 x 600GB 2.5" SAS HDD (1 x 2-drive RAID-1 temp space) 210 Total External Drives	
Initial Database Size 23,393 GB	Redundancy Level: 1 RAID-10 Log RAID-5 Data	Storage 2 x 300GB 2.5" 10K SAS 4 x 800GB 2.5" SATA SSD 208 x 200GB 2.5" SAS SSD 2 x 600GB 2.5" 10K SAS	



IBM System x3850 X6 Microsoft SQL Server 2014

TPC-E 1.12.0

TPC Pricing 1.7.0

Report Date:
February 18, 2014

Revision Date:
February 18, 2014

Availability Date:
April 15, 2014

Description	Part Number	Price Source	Unit Price	Quantity	Extended Price	3-Yr. Maint. Price
Server Hardware						
IBM System x3850 X6 Configure-To-Order, includes:	3837AC1	1 *	144,293	1	144,293	
x3850 X6 4U Chassis + Midplane	44X3950, 44X4052			1		
X6 Compute Book with Intel Xeon Processor E7-4890 v2	44X3994, 44X3995			4		
X6 Primary I/O Book + X6 Storage Book	44X4042, 44X4046			1		
Half Length PCI-e Expansion Card	44X4048			2		
IBM 4x 2.5" HS SAS/SATA/SSD HDD Backplane	44X4103			2		
ServeRAID M5210 SAS/SATA Controller for IBM System x	46C9109			1		
Intel X540 ML2 Dual Port 10GbeseT Adapter for IBM System x	00D1993			1		
IBM 1400W HE Redundant Power Supply	44X4151			4		
S3700 800GB SATA 2.5" G3HS Enterprise SSD	00AJ168			4		
IBM 300GB 10K 6Gbps SAS 2.5" G3HS HDD	00AJ098			2		
Gen 3 Rail Kit	44X4113			1		
Power Cable	39Y5533			4		
32GB PC3L-12800 ECC DDR3 1600MHz LP LRDIMM	46W0675			64		
Intel X540-T2 Dual Port 10GbBase-T Adapter for IBM System x	49Y7973	1	749	3	2,247	
IBM ServeRAID M5120 SAS/SATA Controller	81Y4478	1	299	5	1,495	
ServeRAID M5100 Series 1GB Flash/RAID 5 Upgrade	81Y4559	1	499	5	2,495	
ServeRAID M5100 Series SSD Performance Key	90Y4273	1	399	1	399	
IBM Preferred Pro Keyboard USB - US English 103P RoHS	94Y6138	1	29	1	29	
IBM 2-Button Optical Mouse - Black - USB	40K9200	1	19	1	19	
ServicePac for 3-Year 24x7x4 Support (x3850 X6)	00GW689	1 *	1,500	1		1500
Hanns-G 16" Widescreen LED LCD Monitor (2 spares)	HL163ABB	3	100	3	300	
				Subtotal	151,277	1,500
Server Storage						
IBM S2 42U Standard Rack	93074RX	1	1,459	1	1,459	
IBM System Storage EXP2524 Storage Enclosure	174724X	1	3,999	9	35,991	
IBM 1M SAS cable	39R6529	1	119	9	1,071	
200GB 6Gb SAS 2.5" SSD	49Y6072	1 - S	3,079	208	640,432	
IBM 600GB 10K 6Gbps SAS 2.5" G2HS Hybrid HDD	00AD102	1	899	2	1,798	
ServicePac for 3-Year 24x7x4 Support (EXP2524)	91Y5785	1	1,200	9		10,800
ServicePac for 3-Year 24x7x4 Support (Rack)	41L2760	1	315	1		315
				Subtotal	680,751	11,115
Server Software						
SQL Server 2014 Enterprise Edition		2a *	13,472.50	30	404,175	
Windows Server 2012 Standard Edition	P73-05761	2	882	2	1,764	
Microsoft Problem Resolution Services	N/A	2a	259	1		259
				Subtotal	405,939	259
Client Hardware						
IBM System x3650 M4 Configure-To-Order, includes:	7915AC1	1	12,412	1	12,412	
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/FLL + 2 x8 FH/HLL Slots)	A1JT			1		
IBM System x Gen-III Slides Kit	A228			1		
IBM 250GB 7.2K 6Gbps NL SATA 2.5" SFF HS HDD	A1NX			2		
IBM 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	12,412	790
Client Software						
Windows Server 2012 Standard Edition	P73-05761	2	882	1	882	
				Subtotal	882	0
Infrastructure						
Ethernet Cables (2 spares)	RCW-720	3		4	24	
				Subtotal	24	0
				Total	1,251,285	13,664
Dollar Volume Discount (See Note 1)	24.76%	1			212,335	
Microsoft Open Program Discount Schedule	16.67%	2			441	
Pricing: 1 - IBM 1-866-883-8901; 2 - Microsoft; 3 - newegg.com				Three-Year Cost of Ownership USD: \$1,052,173		
Note 1: Discount applies to all line items where Pricing=1; pricing is for these or similar quantities.				TPC-E Throughput: 5,576.27		
Discounts for similarly sized configurations will be similar to what is quoted here, but may vary based on the specific components priced.				\$ USD/tps.e: \$188.69		
S: One or more components of the measured configuration have been substituted in the priced configuration. See the FDR for details.						
* These components are not immediately orderable. See the FDR for more information.						
Benchmark results and test methodology audited by Doug Johnson for InfoSizing, Inc. (www.sizing.com)						
Prices used in TPC benchmarks reflect the actual prices a customer would pay for a one-time purchase of the stated components. Individually negotiated discounts are not permitted. Special prices based on assumptions about past or future purchases are not permitted. All discounts reflect standard pricing policies for the listed components. For complete details, see the pricing section of the TPC benchmark specifications. If you find that stated prices are not available according to these terms, please inform the TPC at pricing@tpc.org. Thank you.						



IBM System x3850 X6 Microsoft SQL Server 2014

TPC-E 1.12.0 TPC Pricing 1.7.0

Report Date:
February 18, 2014

Revision Date:
February 18, 2014

Availability Date:
April 15, 2014

Numerical Quantities Summary				
Reported Throughput: 5576.27 tpsE		Configured Customers: 2,850,000		
Response Time (in seconds)	Minimum	Average	90 Th Percentile	Maximum
Broker-Volume	0.00	0.01	0.02	0.48
Customer-Position	0.00	0.01	0.01	0.49
Market-Feed	0.00	0.00	0.02	0.87
Market-Watch	0.00	0.01	0.01	0.33
Security-Detail	0.00	0.00	0.01	0.48
Trade-Lookup	0.00	0.06	0.10	0.46
Trade-Order	0.00	0.05	0.12	0.98
Trade-Result	0.00	0.01	0.02	1.45
Trade-Status	0.00	0.00	0.01	0.53
Trade-Update	0.00	0.07	0.11	0.42
Data-Maintenance	0.00	0.01	N/A	0.10
Transaction Mix		Transaction Count		Mix %
Broker-Volume		19,672,099		4.900%
Customer-Position		52,190,503		13.000%
Market-Feed		4,014,916		1.000%
Market-Watch		72,264,528		18.000%
Security-Detail		56,206,099		14.000%
Trade-Lookup		32,117,080		8.000%
Trade-Order		40,548,131		10.100%
Trade-Result		40,149,146		10.001%
Trade-Status		76,279,029		19.000%
Trade-Update		8,029,224		2.000%
Data-Maintenance		120		N/A
Test Duration and Timings				
Ramp-up Time (hh:mm:ss)			00:21:49	
Measurement Interval (hh:mm:ss)			02:00:00	
Business Recovery Time (hh:mm:ss)			00:19:28	
Total Number of Transactions Completed in Measurement Interval			401,470,755	

Table of Contents

Abstract.....	3
Table of Contents.....	7
Clause 0 – Preamble	9
Clause 1 – Introduction.....	11
Benchmark Sponsor.....	11
Configuration Diagrams	11
Measured and Priced Configurations.....	11
Figure 1-1. Measured Configuration.....	11
Table 1-1. Durable Media Substitution Information	12
Hardware and Software Configuration Steps	12
Clause 2- Database Design, Scaling, and Population.....	13
Database Creation and Table Definitions	13
Database Physical Organization	13
Horizontal/Vertical Partitioning	14
Replication.....	14
Table Attributes	14
Cardinality of Tables	14
Table 2-1. Initial Cardinality of Tables.....	15
Distribution of Tables and Logs	16
Table 2-2. Data Distribution for the Measured and Priced Configuration	16
Database Interface and Model Implemented	19
Database Load Methodology	19
Clause 3 – Transaction Related Items	20
Vendor-Supplied Code	20
Database Footprint of Transactions	20
Clause 4 – SUT, Driver, and Network	21
Network Configuration.....	21
Clause 5 – EGen.....	22
EGen Version	22
EGen Code and Modifications.....	22
EGen Files	22
Clause 6 – Performance Metrics and Response Time	23
EGen Instances	23
Measured Throughput.....	23
Throughput vs. Elapsed Time for Trade-Result Transaction.....	23
Figure 6-1. Test Run Graph	23
Steady State Methodology.....	24
Work Performed During Steady State	24
Transaction Statistics	24
Table 6-1. Transaction Statistics	25
Clause 7 – Transaction and System Properties.....	26
Atomicity Requirements.....	26
Consistency Requirements.....	26
Isolation Requirements	27
Durability Requirements.....	27
Durability Test for Data Accessibility	27
Figure 7-1. Data Accessibility Graph.....	28
Durability Test for Business Recovery	29
Figure 7-2. Business Recovery Time Graph	30

Clause 8 – Pricing.....	31
60-Day Space.....	31
Table 8-1. Disk Space Requirements	31
Availability Date.....	32
Table 8-2. Ordering and Pricing Information.....	32
Supporting Files Index.....	32
Auditor’s Attestation Letter	32
Appendix A – Price Quotes.....	35

Clause 0 – Preamble

Introduction

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

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

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

Goal of the TPC-E Benchmark

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

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

The benchmark defines:

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

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

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

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

Restrictions and Limitations

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

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

Clause 1 – Introduction

Benchmark Sponsor

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

This benchmark was sponsored by IBM Corporation.

Configuration Diagrams


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

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

Measured and Priced Configurations

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

Figure 1-1. Measured Configuration



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

IBM System x3850 X6, with:
- 4 x Intel Xeon Processor E7-4890 v2 2.80GHz (4 Procs/60 Cores/120 Threads)
- 2048GB Memory
- 2 x 300GB 10K SAS (RAID-1)
- 4 x 800GB SATA SSD (RAID-10)
- 1 x IBM ServeRAID M5210
- 5 x IBM ServeRAID M5120
- 4 x Dual 10Gb-T Ethernet

15 x IBM EXP2524 JBOD Enclosures, with:
- 208 x 200GB 2.5" SAS SSD (2 x 20-drive RAID-5 DB data)
- 7 x 24-drive RAID-5 DB data)
- 2 x 600GB 2.5" SAS HDD (1 x 2-drive RAID-1 temp space)
- 144 x 600GB 2.5" SAS HDD (6 x 24-drive RAID-10 backup)
354 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 nine 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. Two of these arrays used 20 SSDs each, which are the priced SSDs. The other seven of these arrays used 24 SSDs each. These arrays used older SSDs, but were priced one-for-one as the newer SSDs. This substitution was allowed based on performance measurements, taken while the configuration was running, that show that the 20-drive arrays using the newer SSDs were faster than the 24-drive arrays that used the older SSDs. Table 1-1 shows these measurements.

Table 1-1. Durable Media Substitution Information

Description	24-Drive Array	20-Drive Array
Drives Used	Older SSD	Newer SSD
Drive Capacity	200GB	200GB
Interface Type	SAS	SAS
Interface Speed	3 Gbps	6 Gbps
Drives Per Array	24	20
Avg Array Response Times (read / write, ms)	0.9 / 2.9	0.3 / 1.5
Avg Array Queue Depth	24.9	12.4
Priced As	24 New SSDs	20 New SSDs

Hardware and Software Configuration Steps

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

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

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

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

- Information specific to the Tier A client can be found in:
SupportingFiles\Introduction\TierA\TierA_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 in the Report. Any and all scripts or step by step GUI instructions are reported in the Supporting Files (see Clause 9.4.2). The description, scripts and GUI instructions must be sufficient such that a reader knowledgeable of database software environments and the TPC-E specification could recreate the database.

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

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

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

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

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

The files that were customized for this specific SUT hardware are included in the folder SupportingFiles\Clause2\2850000.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 in the Report.

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

- CASH_TRANSACTION
- SETTLEMENT
- TRADE
- TRADE_HISTORY
- TRADE_REQUEST
- HOLDING
- HOLDING_HISTORY
- HOLDING_SUMMARY

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

Horizontal/Vertical Partitioning

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

Partitioning was not used for this benchmark.

Replication

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

Replication was not used for this benchmark.

Table Attributes

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

No additional attributes were used for this benchmark.

Cardinality of Tables

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

The database was built with 2,850,000 customers. The cardinality is shown in Table 2-1.

Table 2-1. Initial Cardinality of Tables

Table Name	Rows
ACCOUNT_PERMISSION	20,235,711
ADDRESS	4,275,004
BROKER	28,500
CASH_TRANSACTION	45,308,170,225
CHARGE	15
COMMISSION_RATE	240
COMPANY	1,425,000
COMPANY_COMPETITOR	4,275,000
CUSTOMER	2,850,000
CUSTOMER_ACCOUNT	14,250,000
CUSTOMER_TAXRATE	5,700,000
DAILY_MARKET	2,547,686,250
EXCHANGE	4
FINANCIAL	28,500,000
HOLDING	2,521,516,990
HOLDING_HISTORY	66,000,661,220
HOLDING_SUMMARY	141,735,987
INDUSTRY	102
LAST_TRADE	1,952,250
NEWS_ITEM	2,850,000
NEWS_XREF	2,850,000
SECTOR	12
SECURITY	1,952,250
SETTLEMENT	49,248,000,000
STATUS_TYPE	5
TAXRATE	320
TRADE	49,248,000,000
TRADE_HISTORY	118,195,266,610
TRADE_REQUEST	0
TRADE_TYPE	5
WATCH_ITEM	285,007,996
WATCH_LIST	2,850,000
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 was stored on four 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 five IBM ServeRAID M5120 SAS/SATA controllers.

- One of these controllers was connected to two external EXP2524 enclosures, each holding 20 200GB SAS SSDs.
- Three of these controllers were each connected to two external EXP2524 enclosures, each holding 24 200GB SAS SSDs.
- One of these controllers was connected to one external EXP2524 enclosure, which was holding 24 200GB SAS SSDs.

In total, for database data, nine enclosures and 208 external SSDs were connected to the database server. Two data arrays were each configured as 20-drive RAID-5 and seven data arrays were each configured as 24-drive RAID-5. 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).

Additionally, one of those two enclosures with 20 SSDs also contained two 600GB 2.5” 10K SAS drives. These two drives were configured as a RAID-1 array and formatted as NTFS. This storage space was used to hold tempdb while the benchmark was running.

In addition to the priced configuration described above, the measured configuration included six additional external EXP2524 enclosures, each filled with twenty-four 600GB 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 six enclosures were attached to the previously mentioned ServeRAID M5120 SAS/SATA controllers, one directly and the other five via daisy-chaining. Six 24-drive RAID-10 arrays were created using this hardware and formatted as NTFS.

Adapter write caching was disabled for all controllers and arrays, except for the 600GB RAID-1 array.

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

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

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

Disk #	Controller	Drives Enclosure RAID Level (Pricing)	Partition (File System)	Size	Use
0	Internal M5210	2 x 300GB SAS HDD internal RAID-1	C: (NTFS)	277.95GB	OS
1	Internal M5210	4 x 800GB SATA SSD internal RAID-10	E: (RAW) F: (NTFS)	1269.53GB 218.59GB	tpce log MDF

Disk #	Controller	Drives Enclosure RAID Level (Pricing)	Partition (File System)	Size	Use
2	M5120 #1	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk6 (NTFS)	6694.21GB	backup & flatfiles
3	M5120 #1	24 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx1 (RAW) c:\mp\gw1 (RAW) c:\mp\xt1 (NTFS)	69.43GB 3304.20GB 888.90GB	fixed_fg growing_fg tempdb
4	M5120 #1	24 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx3 (RAW) c:\mp\gw3 (RAW) c:\mp\xt3 (NTFS)	69.43GB 3304.20GB 888.90GB	fixed_fg growing_fg tempdb
5	M5120 #2	2 x 600GB SAS HDD EXP2524 RAID-1	T: (NTFS)	557.73GB	tempdb
6	M5120 #2	20 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx4 (RAW) c:\mp\gw4 (RAW) c:\mp\xt4 (NTFS)	69.43GB 3304.20GB 147.57GB	fixed_fg growing_fg tempdb
7	M5120 #2	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk4 (NTFS)	6694.21GB	backup & flatfiles
8	M5120 #2	20 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx5 (RAW) c:\mp\gw5 (RAW) c:\mp\xt5 (NTFS)	69.43GB 3304.20GB 147.57GB	fixed_fg growing_fg tempdb
9	M5120 #3	24 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx6 (RAW) c:\mp\gw6 (RAW) c:\mp\xt6 (NTFS)	69.43GB 3304.20GB 888.90GB	fixed_fg growing_fg tempdb
10	M5120 #3	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk1 (NTFS)	6694.21GB	backup & flatfiles
11	M5120 #3	24 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx7 (RAW) c:\mp\gw7 (RAW) c:\mp\xt7 (NTFS)	69.43GB 3304.20GB 888.90GB	fixed_fg growing_fg tempdb
12	M5120 #4	24 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx2 (RAW) c:\mp\gw2 (RAW) c:\mp\xt2 (NTFS)	69.43GB 3304.20GB 888.90GB	fixed_fg growing_fg tempdb
13	M5120 #4	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk3 (NTFS)	6694.21GB	backup & flatfiles
14	M5120 #4	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk5 (NTFS)	6694.21GB	backup & flatfiles
15	M5120 #5	24 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx8 (RAW) c:\mp\gw8 (RAW) c:\mp\xt8 (NTFS)	69.43GB 3304.20GB 888.90GB	fixed_fg growing_fg tempdb
16	M5120 #5	24 x 200GB SAS SSD EXP2524 RAID-5	c:\mp\fx9 (RAW) c:\mp\gw9 (RAW) c:\mp\xt9 (NTFS)	69.43GB 3304.20GB 888.90GB	fixed_fg growing_fg tempdb

Disk #	Controller	Drives Enclosure RAID Level (Pricing)	Partition (File System)	Size	Use
17	M5120 #5	24 x 600GB SAS HDD EXP2524 RAID-10 (Measured)	c:\mp\bk2 (NTFS)	6694.21GB	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 in the Report.

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

Clause 3 – Transaction Related Items

Vendor-Supplied Code

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

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

The code to interface the stored procedures can be 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 cables. 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 cable was connected 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 v1.12.0 was used in the benchmark.

EGen Code and Modifications

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

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

EGen was not modified for use in this benchmark.

EGenLoader was not extended for this benchmark.

EGen Files

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

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

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

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

Clause 6 – Performance Metrics and Response Time

EGen Instances

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

There were 16 EGenDriverCEs with a total of 840 EGenDriverCE instances used in the benchmark.

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

Measured Throughput


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

The Measured Throughput was 5,576.27 tpsE.

Throughput vs. Elapsed Time for Trade-Result Transaction

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

Figure 6-1. Test Run Graph



Steady State Methodology

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

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

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

Work Performed During Steady State

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

Checkpoints 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	49.99%	48% to 52%
Market-Watch			
Securities chosen by	Watch List	60.00%	57% to 63%
	Account ID	35.00%	33% to 37%
	Industry	5.00%	4.5% to 5.5%
Security-Detail			
Access LOB	1	1.00%	0.9% to 1.1%
Trade-Lookup			
Frame to execute	1	30.00%	28.5% to 31.5%
	2	30.01%	28.5% to 31.5%
	3	30.00%	28.5% to 31.5%
	4	9.99%	9.5% to 10.5%
Trade-Order			
Transactions requested by a third party		10.01%	9.5% to 10.5%
By Company Name		40.01%	38% to 42%
Buy On Margin	1	7.99%	7.5% to 8.5%
Rollback	1	0.99%	0.94% to 1.04%
LIFO	1	35.00%	33% to 37%
Trade Quantity	100	25.01%	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	20.00%	19.8% to 20.2%
	Limit Sell	10.00%	9.9% to 10.1%
	Stop Loss	10.01%	9.9% to 10.1%
Trade-Update			
Frame to execute	1	33.00%	31% to 35%
	2	33.01%	31% to 35%
	3	33.99%	32% to 36%

Clause 7 – Transaction and System Properties

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

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

Atomicity Requirements

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

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

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

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

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

Consistency Requirements

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

Consistency condition 1:

Entries in the BROKER and TRADE tables must satisfy the relationship:

B_NUM_TRADES = count()*

For each broker defined by:

(B_ID = CA_B_ID) and (CA_ID = T_CA_ID) and (T_ST_ID = —CMPT').

Consistency condition 2:

Entries in the BROKER and TRADE tables must satisfy the relationship:

B_COMM_TOTAL = sum(T_COMM)

For each broker defined by:

(B_ID = CA_B_ID) and (CA_ID = T_CA_ID) and (T_ST_ID = —CMPT').

Consistency condition 3:

Entries in the HOLDING_SUMMARY and HOLDING tables must satisfy the relationship:

HS_QTY = sum(H_QTY)

For each holding summary defined by:

(HS_CA_ID = H_CA_ID) and (HS_S_SYMB = H_S_SYMB).

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

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

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

Isolation Requirements

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

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

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

Durability Requirements

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

Durability Test for Data Accessibility

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

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

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


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

1. 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 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 Durability Throughput Requirements were met again for at least 5 minutes.
6. Induced the second failure, which in this case was failing a drive in the database tempdb array by physically removing it from its enclosure. Since the database tempdb array is RAID protected, transaction processing continued.
7. Waited until the Durability 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 the data accessibility run. The timings of the induced failures as well as the recovery process are indicated.

Figure 7-1. Data Accessibility Graph



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

Durability Test for Business Recovery

The Test Sponsor must describe in the Report the test(s) used to demonstrate Business Recovery.

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

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

The tests for “Loss of Processing,” “Loss of Vulnerable Storage Component,” and “Loss of all External Power to the SUT” were combined.

The following steps were successfully performed to test Business Recovery:


1. 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.

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

The Database Recovery Time was 00:08:21. The Application Recovery Time was 00:11:07. The Business Recovery Time, which is the sum of the Database Recovery Time and the Application Recovery Time, was 00:19:28.

Figure 7-2. Business Recovery Time Graph



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

Clause 8 – Pricing

60-Day Space

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

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	2,850,000	Performance	5576.27 TpsE	Reported	5576.27 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	28,500	2,080	2,192	214	4,486	28,500	4,272	-	-	214
CASH_TRANSACTION	45,308,170,225	4,709,972,296	9,928,200	235,995,025	4,955,895,521	45,377,534,735	4,733,154,664	13,254,168	28,232,146	28,232,146
CHARGE	15	8	8	1	17	15	16	-	-	1
COMMISSION_RATE	240	16	16	2	34	240	32	-	-	2
SETTLEMENT	49,248,000,000	2,348,335,280	4,952,808	117,664,404	2,470,952,492	49,323,395,404	2,360,663,448	7,375,360	15,709,944	15,709,944
TRADE	49,248,000,000	5,879,471,152	3,274,523,016	457,699,708	9,611,693,876	49,324,384,682	9,187,875,904	33,881,736	72,170,060	72,170,060
TRADE_HISTORY	118,195,266,610	3,554,745,360	9,268,136	178,200,675	3,742,214,171	118,377,930,278	3,577,190,288	13,176,792	28,067,330	28,067,330
TRADE_REQUEST	-	-	-	-	-	331,519	863,488	863,488	1,839,280	1,839,280
TRADE_TYPE	5	8	1,032	52	1,092	5	1,040	-	-	52
ACCOUNT_PERMISSION	20,235,711	1,114,360	7,296	56,083	1,177,739	20,235,711	1,121,808	152	324	56,083
CUSTOMER	2,850,000	467,152	140,752	30,395	638,299	2,850,000	607,968	64	137	30,395
CUSTOMER_ACCOUNT	14,250,000	1,291,368	318,496	80,493	1,690,357	14,250,000	1,609,864	-	-	80,493
CUSTOMER_TAXRATE	5,700,000	118,936	2,008	6,047	126,991	5,700,000	121,128	184	392	6,047
HOLDING	2,521,516,990	168,626,864	115,282,792	14,195,483	298,105,139	2,523,434,407	288,422,312	4,512,656	9,612,219	9,612,219
HOLDING_HISTORY	66,000,661,220	2,400,024,512	1,603,310,200	200,166,736	4,203,501,448	66,102,324,130	4,018,894,632	15,559,920	33,143,531	33,143,531
HOLDING_SUMMARY	141,735,987	6,215,544	24,392	311,997	6,551,933	141,736,571	6,239,936	-	-	-
WATCH_ITEM	285,007,996	8,002,616	30,712	401,666	8,434,994	285,007,996	8,033,656	328	699	401,666
WATCH_LIST	2,850,000	71,080	66,464	6,877	144,421	2,850,000	137,544	-	-	6,877
COMPANY	1,425,000	304,136	93,280	19,871	417,287	1,425,000	397,440	24	52	19,871
COMPANY_COMPETITOR	4,275,000	114,848	106,352	11,060	232,260	4,275,000	221,200	-	-	11,060
DAILY_MARKET	2,547,686,250	119,489,104	350,096	5,991,960	125,831,160	2,547,686,250	119,840,912	1,712	3,647	5,991,960
EXCHANGE	4	8	8	1	17	4	16	-	-	1
FINANCIAL	28,500,000	3,211,768	10,440	161,110	3,383,318	28,500,000	3,222,608	400	853	161,110
INDUSTRY	102	8	24	2	34	102	32	-	-	2
LAST_TRADE	1,952,250	121,792	2,008	6,190	129,990	1,952,250	123,800	-	-	6,190
NEWS_ITEM	2,850,000	308,992,104	5,000	15,449,855	324,446,959	2,850,000	308,997,152	48	103	15,449,855
NEWS_XREF	2,850,000	71,056	2,008	3,653	76,717	2,850,000	73,064	-	-	3,653
SECTOR	12	8	24	2	34	12	32	-	-	2
SECURITY	1,952,250	270,736	77,080	17,391	365,207	1,952,250	347,840	24	52	17,391
STATUS_TYPE	5	8	8	1	17	5	16	-	-	1
ADDRESS	4,275,004	246,640	2,056	12,435	261,131	4,275,004	248,776	80	171	12,435
TAXRATE	320	24	16	2	42	320	56	16	35	35
ZIP_CODE	14,741	488	72	28	588	14,741	560	-	-	28
TOTALS (KB)		19,511,281,360	5,018,506,992	1,226,489,418	25,756,277,770		24,618,415,504	88,627,152	188,780,975	211,029,933
Initial Database Size (MB)		23,954,871	23,393 GB							
Database Filegroups	LUN Count	Partition Size (MB)	MB Allocated	MB Loaded	MB Required					
	0	-	-	-	-	OK				
growing_fg	9	3,325,100	29,925,900	23,520,196	23,704,546	OK				
	0	-	-	-	-	OK				
fixed_fg	9	69,800	628,200	434,676	456,409	OK				
Settlements	75,395,404									
Data Space Required (MB)	Data Space Configured (MB)					Log Space Required (MB)		Log Space Configured (MB)		
Initial Growing Space	23,520,196					Initial Log Size	7,414	Log LUNS	1	
Final Growing Space	23,606,743	Data LUNS	2	7	-	Final Log Size	517,497	Log Disks	4	
Delta	86,547	Disks per LUN	20	24	-	Log Growth	510,083	Disk Capacity	761,984	
Data Space per Trade	0.001147908	Disk Capacity	189,781	189,781	-	Log Growth/Trade	0.0006765436	RAID Overhead	50%	
1 Day Data Growth	184,350	RAID Overhead	95%	96%	0%	1 Day Log Space	1,093,920	Log Space	1,523,968	
60 Day Space	35,015,878	Total Space				37,766,417				
						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 April 15, 2014. The dates for ordering and availability are detailed in Table 8-2 for those components that are not immediately orderable.

Table 8-2. Ordering and Pricing Information

Description	Part Number	Order Date	Availability Date	Order Method	Price Verification
IBM System x3850 X6	3837-AC1	2014-02-19	2014-03-25	See note 1	See note 2
Microsoft SQL Server 2014 Enterprise Edition		2014-04-15	2014-04-15	See note 3	See note 3

Note 1: IBM - 1-866-883-8901

Note 2: These components are not immediately orderable. For price verification before the order date, call IBM at 1-866-883-8901.

Note 3: See the attached Microsoft price quote

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.

Joe Herman, Manager
System x Server Performance
IBM Systems and Technology Group
3039 Cornwallis Road
RTP, NC 27709

February 12, 2014

I verified the TPC Benchmark™ E v1.12.0 performance of the following configuration:

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

The results were:

Performance Metric 5576.27 tpsE

Trade-Result 90th %-tile 0.02 Seconds

Tier B (Server)

IBM System x3850 X6

CPU	4 x Intel Xeon Processor E7-4890 v2 (2.80 GHz, 15-core, 37.5 MB L3)		
Memory	2048 GB		
Storage	Qty	Size	Type
	2	300 GB	10K rpm SAS HDD
	4	800 GB	SATA SSD
	208	200 GB	SAS SSD
	2	600 GB	10K rpm SAS HDD

Tier A (Client)

IBM System x3650 M4

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

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

The following verification items were given special attention:

- All EGen components were verified to be v1.12.0
- The transaction were correctly implemented
- The database was properly scaled and populated for 2,850,000 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 00:19:28 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 (168) XceedIOPS SAS SSD disks that were substituted by (168) Optimus SAS SSD disks in the priced configuration. Based on the specifications of these disks and on I/O data collected during testing, it is my opinion that this substitution has no significant effect on performance.

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

February 4, 2014

IBM Corporation
Ray Engler
3039 Cornwallis Road
RTP, NC 27709

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

All pricing shown is in US Dollars (\$).

Part Number	Description	Unit Price	Quantity	Price
Database Management System				
*	SQL Server 2014 Enterprise Edition 2 Core License Open Program - Level C	\$13,472.50	30	\$404,175.00
Database Management System Operating System				
P73-05761	Windows Server 2012 Standard 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 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

Windows Server 2012 Standard is 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>

SQL Server 2014 Enterprise Edition will be orderable by April 15, 2014. The part number for SQL Server 2014 Enterprise Edition will be set and available by April 15, 2014.

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_qhtplylGYLKTUUKf95957fiii_2014_nggtp.

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



Shopping Cart

[Print](#)

Qty.	Product Description		Savings	Total Price
1		Hanns-G HL163ABB Black 16" Widescreen LED Backlight LCD Monitor Item #: N82E16824254115 Return Policy: Monitor Standard Return Policy	-\$30.00 Instant	\$99.99 \$69.99
1		Rosewill RCW-720 14ft. Network Cable Cat 6 (Crossover) Yellow Item #: N82E16812119156 Return Policy: Standard Return Policy <input type="checkbox"/> Send this item as a gift 📁		\$3.49
Subtotal:				\$73.48
<div>Calculate Shipping</div> <div>Zip Code: 27709</div> <div>Choose a delivery method 📁</div> <div><input checked="" type="radio"/> \$2.89 Super Eggsaver (4-7 Business Days)</div> <div><input type="radio"/> \$14.33 Newegg 3 Business Days</div> <div><input type="radio"/> \$26.89 Newegg 2 Business Days</div> <div><input type="radio"/> \$44.19 Newegg Next Business Day</div> <div>SAVE MONEY WITH SUPER EGGSAYER</div> <div>“Super EggSaver” is a low-cost 4-7 business day shipping method. Learn more.</div>			Shipping:	\$2.89
Redeem Newegg Gift Cards / Google Offer Code Card Number: <input type="text"/> Security Code: <input type="text"/>				
Redeem Eggpoints Please log in to redeem your EggPoints				
Apply Promo Code(s):			Promo Code:	\$0.00
Apply Promo Codes Promo codes have usage limitations that require you to enter your email address to enable Promo Codes. E-Mail Address: <input type="text"/> Enter We value your privacy. You won't get any emails from us unless you place an order or sign up for our newsletter.				

Grand Total:	\$76.37
--------------	---------

Print

[Policy & Agreement](#) | [Privacy Policy](#) © 2000-2014 Newegg Inc. All rights reserved.