

TPC Benchmark™ E
Full Disclosure Report
for
Lenovo® ThinkSystem™ SR950
using
Microsoft® SQL Server® 2017
Enterprise Edition
and
Microsoft Windows Server® 2016
Standard Edition

TPC-E™ Version 1.14.0



First Edition
Submitted for Review
November 6, 2017

First Edition – November 2017

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

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

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

© Copyright Lenovo 2017. All rights reserved.

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

Trademarks

Lenovo, System x, ServeRAID, ThinkSystem, the Lenovo logo, and For Those Who Do are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available on the Web at <http://www3.lenovo.com/us/en/legal/copytrade/>.

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

Lenovo conducted the TPC Benchmark E (TPC-E) on the Lenovo ThinkSystem SR950. This report documents the full disclosure information required by the TPC Benchmark E Standard Specification, Revision 1.14.0, including the methodology used to achieve the reported results. All testing fully complied with this revision level.

The software used on the Lenovo ThinkSystem SR950 system included Microsoft Windows Server 2016 Standard Edition and Microsoft SQL Server 2017 Enterprise Edition.

Standard metrics, transactions per second-E (tpsETM), 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 (\$USD)	tpsE	\$ USD / tpsE	Total Solution Availability Date
Lenovo ThinkSystem SR950	Microsoft SQL Server 2017 Enterprise Edition Microsoft Windows Server 2016 Standard Edition	\$1,122,410	11,357.28	\$98.83	November 6, 2017

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.

Executive Summary

The Executive Summary is included on the next several pages.

	Lenovo[®] ThinkSystem[™] SR950		TPC-E[™] 1.14.0 TPC Pricing 2.1.1
			Report Date: Nov 6, 2017 Revision Date: Nov 6, 2017
TPC-E Throughput 11,357.28 tpsE [™]	Price/Performance \$98.83 USD per tpsE	Availability Date November 6, 2017	Total System Cost \$1,122,410 USD
Database Server Configuration			
Operating System Microsoft[®] Windows Server[®] 2016 Standard Edition	Database Manager Microsoft SQL Server[®] 2017 Enterprise Edition	Processors/Cores/Threads 4/112/224	Memory 3072GB
			
4 x 10Gb Ethernet		12Gb SAS x12	
		6 x 	
TIER A (Client)		TIER B (Database Server and Storage)	
1 x Lenovo System x3650 M5, with: <ul style="list-style-type: none"> - 2 x Intel[®] Xeon[®] Processor E5-2699 v4 2.20GHz (2 Procs/44 Cores/88 Threads) - 64GB Memory - 2 x 600GB SFF SAS (RAID-1) - 1 x ServeRAID M5210 - 1 x Onboard Quad Gb Ethernet - 3 x Intel x540 Dual 10Gb-T Ethernet 		1 x Lenovo ThinkSystem SR950, with: <ul style="list-style-type: none"> - 4 x Intel Xeon Platinum 8180 Processor 2.50GHz (4 Procs/112 Cores/224 Threads) - 3072GB Memory - 2 x 800GB SAS SSD (RAID-1) - 8 x 800GB SAS SSD (RAID-10) - 1 x ThinkSystem RAID 930-16i - 6 x ThinkSystem RAID 930-8e - 1 x Onboard Dual Gb Ethernet - 1 x Broadcom NX-E Dual 10Gb-T - 3 x Intel x550 Dual 10Gb-T Ethernet 	
		6 x Lenovo Storage D1224 DAS Enclosures, with: <ul style="list-style-type: none"> - 106 x 800GB 2.5" SAS SSD (6 x 17-drive RAID-5) (1 x 4-drive RAID-10) - 3 x 1200GB 2.5" 10K SAS (1 x 3-drive RAID-5) 	
		109 Total External Drives	
Initial Database Size 49,276 GB	Redundancy Level: 1 RAID-10 Log, tempdb RAID-5 Data, 60-Day Space	Storage 3 x 1200GB 2.5" 10K SAS 116 x 800GB 2.5" SAS SSD	



Lenovo ThinkSystem SR950

TPC-E 1.14.0 TPC Pricing 2.1.1

Report Date: Nov 6, 2017
Revision Date: Nov 6, 2017

Availability Date:
November 6, 2017

Description	Part Number	Price Source	Unit Price	Quantity	Extended Price	3-Yr. Maint. Price
Server Hardware						
ThinkSystem SR950 Configure-To-Order, includes:	7X12CT01WWW	1	220,142	1	220,142	
SR950 2S to 4S Base + Main I/O Planar	AUNY, AJUMT			1		
Intel Xeon Platinum 8180 28C 205W 2.5GHz Processor	AX6G			4		
ThinkSystem 64GB TruDDR4 2666 MHz (4Rx4 1.2V) LRDIMM	AUNE			48		
ThinkSystem 2.5" SAS/SATA 2x2 Bay Backplane Kit	AUN6			3		
ThinkSystem RAID 930-16i 4GB Flash PCIe 12Gb Adapter	AUNK			1		
ThinkSystem 2.5" 800GB Mainstream SAS 12Gb Hot Swap SSD	AUMD			2		
ThinkSystem 2.5" 800GB Performance SAS 12Gb Hot Swap SSD	ALMH			8		
ThinkSystem SR950 (4) x8 PCIe Riser	AUN2			2		
ThinkSystem 1Gb 2-port RJ45 LOM	AUKG			1		
ThinkSystem Broadcom NX-E ML2 10Gb 2-Port Base-T Ethernet Adapter	AUKQ			1		
Intel X550-T2 Dual Port 10GBase-T Adapter	ATPX			3		
ThinkSystem RAID 930-8e 4GB Flash PCIe 12Gb Adapter	AUNQ			6		
High Efficiency 1600W Power Supply for SR950 (200-240VAC only)	AUPJ			4		
Power Cable	6311			4		
FBU345 SuperCap	AUNP			7		
ThinkSystem SR950 2-CPU, 24 DIMM Compute System Board	AUN0			2		
Front operator panel with LCD display	AUMW			1		
Rail Kit + SR950 Lift Handles	A4AA, AUMY			1		
Essential Service - 3Yr 24x7 4Hr Response + YourDrive YourData (SR950)	5PS7A01828	1	1,999	1		1,999
				Subtotal	220,142	1,999
Server Storage						
Lenovo Storage D1224 SFF Chassis, Dual ESM	4587A31	1	2,999	6	17,994	
External MiniSAS HD 8644/MiniSAS HD 8644 1M Cable	00YL848	1	59	12	708	
Lenovo Storage 800GB 10 DWD 2.5" SAS SSD	01DC452	1 S	2,999	106	317,894	
Lenovo Storage 1.2TB 10K 2.5" SAS HDD	01DC407	1	989	3	2,967	
Essential Service - 3Yr 24x7 4Hr Response (D1224)	01JR577	1	1,169	6		7,014
				Subtotal	339,563	7,014
Server Software						
SQL Server 2017 Enterprise Edition (2 Core License)	N/A	2	13,472.50	56	754,460	
Windows Server 2016 Standard Edition (2 Core License)	N/A	2	92	56	5,152	
Microsoft Problem Resolution Services (1 Incident)	N/A	2	259	1		259
				Subtotal	759,612	259
Client Hardware						
System x3650 M5 Configure-To-Order, includes:	8871AC1	1	21,033	1	21,033	
System x3650 M5 Base + Planar + Power Paddle Card	ATDY, ATE4, A5FT			1		
System x 750W High Efficiency Platinum AC Power Supply	A5EU			2		
Intel Xeon Processor E5-2699 v4 22C 2.2GHz 55MB Cache 2400MHz 145W	ATF2, ATFT			2		
System x3650 M5 PCIe Risers 1 & 2 + Thermal Kit	A5FQ, A5R6, ASQD			1		
System x Gen-II Universal Slides Kit	A5FW			1		
600GB 10K 12Gbps SAS 2.5" G3HS HDD	AT8A			2		
x3650 M5 8x 2.5" HS HDD Assembly Kit (Single RAID)	A5G6			1		
8GB TruDDR4 Memory (2Rx8, 1.2V) PC4-19200 CL17 2400MHz LP RDIMM	ATC9			8		
Power Cable	6311			2		
ServeRAID M5210 SAS/SATA Controller	A3YZ			1		
Intel X540-T2 Dual Port 10GBaseT Adapter	A2ED			2		
Intel X540 ML2 Dual Port 10GbaseT Adapter	A40P			1		
System x3650 M5 Riser Bracket	A5G5			2		
System x3650 M5 EIA L- Blank, Right EIA, & EIA Plate	ATEA, A5V5, A5G1			1		
System x3650 M5 Single Rotor Fan	ATEC			4		
Essential Service - 3Yr 24x7 4Hr Response + YourDrive YourData (x3650 M5)	01GX547	1	1,114	1		1,114
				Subtotal	21,033	1,114
Client Software						
Windows Server 2016 Standard Edition (2 Core License)	N/A	2	92	22	2,024	
				Subtotal	2,024	0
Infrastructure						
S2 42U Standard Rack	93074RX	1	1,565	1	1,565	
0U 36 C13/6 C19 24A/200-240V 1 Phase PDU with line cord	00YJ776	1	479	1	479	
Preferred Pro Keyboard USB - US English 103P RoHS v2	00AM600	1	29	1	29	
2 Button Optical Wheel Mouse - Black - USB	40K9200	1	19	1	19	
ThinkVision E2054 19.5-inch LED Backlit LCD Monitor	60DFAAR1US	1	119	1	119	
1.0m CAT6 Green Cable	00WE127	1	19	4	76	
Essential Service - 3Yr 24x7 4Hr Response (Rack)	41L2760	1	315	1		315
				Subtotal	2,287	315
				Total	1,344,661	10,701
Dollar Volume Discount (See Note 1)		39.25%	1		232,952	

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.

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 Line Items. 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 Line Items. For complete details, see the pricing section of the TPC Benchmark Standard. If you find that the stated prices are not available according to these terms, please inform the TPC at pricing@tpc.org. Thank you.

Three-Year Cost of Ownership USD: \$1,122,410
TPC-E Throughput: 11,357.28
\$ USD/tpsE: \$98.83



Lenovo ThinkSystem SR950

TPC-E 1.14.0 TPC Pricing 2.1.1

Report Date: Nov 6, 2017
Revision Date: Nov 6, 2017

Availability Date:
November 6, 2017

Numerical Quantities Summary

Reported Throughput:

11,357.28 tpsE

Configured Customers:

6,000,000

Response Time (in seconds)	Minimum	Average	90 th Percentile	Maximum
Broker-Volume	0.01	0.01	0.02	0.44
Customer-Position	0.01	0.01	0.02	2.24
Market-Feed	0.01	0.01	0.02	2.23
Market-Watch	0.01	0.01	0.02	0.43
Security-Detail	0.01	0.01	0.01	0.43
Trade-Lookup	0.01	0.04	0.06	0.60
Trade-Order	0.01	0.03	0.07	0.63
Trade-Result	0.01	0.02	0.04	0.59
Trade-Status	0.01	0.01	0.01	0.43
Trade-Update	0.01	0.05	0.07	0.52
Data-Maintenance	0.01	0.01	N/A	0.05

Transaction Mix	Transaction Count	Mix %
Broker-Volume	40,068,102	4.900%
Customer-Position	106,302,916	13.000%
Market-Feed	8,177,257	1.000%
Market-Watch	147,188,055	18.000%
Security-Detail	114,479,771	14.000%
Trade-Lookup	65,416,789	8.000%
Trade-Order	82,588,501	10.100%
Trade-Result	81,772,482	10.000%
Trade-Status	155,365,918	19.000%
Trade-Update	16,354,137	2.000%
Data-Maintenance	120	N/A

Test Duration and Timings

Ramp-up Time (hh:mm:ss)	00:43:46
Measurement Interval (hh:mm:ss)	02:00:00
Business Recovery Time (hh:mm:ss)	00:35:57
Total Number of Transactions Completed in Measurement Interval	817,713,928

Table of Contents

Abstract	3
Executive Summary	3
Table of Contents	7
Clause 0 – Preamble	9
Introduction.....	9
Goal of the TPC-E Benchmark	9
Restrictions and Limitations	9
Clause 1 – Introduction	11
Benchmark Sponsor	11
Configuration Diagrams.....	11
Measured Configuration.....	11
Figure 1-1. Measured Configuration	11
Priced Configuration	11
Differences between the Priced and Measured Configurations.....	12
Substitution	12
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 Configurations.....	16
Database Interface and Model Implemented.....	18
Database Load Methodology	18
Clause 3 – Transaction Related Items.....	19
Vendor-Supplied Code.....	19
Database Footprint of Transactions	19
Clause 4 – SUT, Driver, and Network.....	20
Network Configuration	20
Clause 5 – EGen	21
EGen Version.....	21
EGen Code and Modifications	21
EGen Files	21
Clause 6 – Performance Metrics and Response Time.....	22

EGen Instances	22
Reported Throughput.....	22
Throughput vs. Elapsed Time for Trade-Result Transaction.....	22
Figure 6-1. Test Run Graph.....	22
Steady State Methodology	23
Work Performed During Steady State	23
Transaction Statistics	23
Table 6-1. Transaction Statistics	24
Clause 7 – Transaction and System Properties	25
Atomicity Requirements	25
Consistency Requirements	25
Isolation Requirements	26
Durability Requirements.....	26
Data Accessibility	26
Table 7-1. Combinations of Durable Media Technologies Tested for Data Accessibility	26
Figure 7-1. Data Accessibility Graph.....	28
Business Recovery	28
Figure 7-2. Business Recovery Time Graph	30
Clause 8 – Pricing	31
60-Day Space.....	31
Table 8-1. Disk Space Requirements	31
Hardware and Software Components	32
Three-Year Cost of System Configuration	32
Availability Date.....	32
Country-Specific Pricing	32
Pricing Calculations	32
Supporting Files Index.....	32
Auditor’s Attestation Letter	33
Appendix A – Price Quotes	36

Clause 0 – Preamble

Introduction

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

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

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

This benchmark was sponsored by Lenovo.


Configuration Diagrams

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

Measured Configuration

The measured configuration is shown in Figure 1-1.

Figure 1-1. Measured Configuration



TIER A (Client)

1 x Lenovo System x3650 M5, with:

- 2 x Intel Xeon Processor E5-2699 v4 2.20GHz (2 Procs/44 Cores/88 Threads)
- 64GB Memory
- 2 x 600GB SFF SAS (RAID-1)
- 1 x ServeRAID M5210
- 1 x Onboard Quad Gb Ethernet
- 3 x Intel x540 Dual 10Gb-T Ethernet

TIER B (Database Server and Storage)

1 x Lenovo ThinkSystem SR950, with:

- 4 x Intel Xeon Platinum 8180 Processor 2.50GHz (4 Procs/112 Cores/224 Threads)
- 3072GB Memory
- 2 x 800GB SAS SSD (RAID-1)
- 8 x 800GB SAS SSD (RAID-10)
- 1 x ThinkSystem RAID 930-16i
- 6 x ThinkSystem RAID 930-8e
- 1 x Onboard Dual Gb Ethernet
- 1 x Broadcom NX-E Dual 10Gb-T
- 3 x Intel x550 Dual 10Gb-T Ethernet

18 x Lenovo Storage D1224

DAS Enclosures, with:

- 202 x 800GB 2.5" SAS SSD (6 x 17-drive RAID-5) (1 x 4-drive RAID-10) (6 x 16-drive RAID-5)
- 99 x 1200GB 2.5" 10K SAS (4 x 24-drive RAID-5) (1 x 3-drive RAID-5)
- 48 x 900GB 2.5" 10K SAS (2 x 24-drive RAID-5)

349 Total External Drives

Priced Configuration

The priced configuration is shown above in the Executive Summary.

Differences between the Priced and Measured Configurations

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.

Substitution

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 section of the FDR that describes the differences between the measured and Priced Configuration. Original and substituted Components must be clearly identified.

Some of the database data SSDs were substituted. Five of the six database data arrays used the priced SSDs. The remaining database data array used older SSDs. All of the database data arrays were 17-drive RAID-5 and did the same amount and type of work. See Table 1-1 for details.

Table 1-1. Durable Media Substitution Information

Description	Database Data Array with Priced SSDs	Database Data Array with Older SSDs
Drives Used	Priced SSD	Older SSD
Drive Capacity	800GB	800GB
Drive Interface Type	SAS	SAS
Drive Interface Speed	12 Gbps	12 Gbps
Drives Per Array	17-drive RAID-5	17-drive RAID-5
Array Utilization	1.018 GB/s	1.017 GB/s
Avg. Array Response Times (read / write, ms)	0.16 / 0.29	0.19 / 0.34
Avg. Array Queue Depth	21.0	24.4
Priced As	17 x 800GB SSDs	17 x 800GB SSDs

Hardware and Software Configuration Steps

A description of the steps taken to configure all of the hardware and 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 and 9.4.1.2). The description, scripts and GUI instructions must be sufficient such that a reader knowledgeable of computer systems and the TPC-E specification could recreate the hardware and software environments.

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

- Information specific to the Tier A client can be found in:
SupportingFiles\Introduction\TierA\TierA_x3650M5_Setup.pdf
- Information specific to the Tier B database server and storage can be found in:
SupportingFiles\Introduction\TierB\TierB_SR950_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\6000000.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 in the Report.

Partitioning was not used for this benchmark.

Replication

Replication of tables, if used, must be reported in the Report (see Clause 2.3.4).

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 database load (see Clause 2.6), must be reported in the Report.

The database was built with 6,000,000 customers. The cardinality is shown in Table 2-1.

Table 2-1. Initial Cardinality of Tables

Table Name	Rows
ACCOUNT_PERMISSION	42,597,820
ADDRESS	9,000,004
BROKER	60,000
CASH_TRANSACTION	95,385,505,113
CHARGE	15
COMMISSION_RATE	240
COMPANY	3,000,000
COMPANY_COMPETITOR	9,000,000
CUSTOMER	6,000,000
CUSTOMER_ACCOUNT	30,000,000
CUSTOMER_TAXRATE	12,000,000
DAILY_MARKET	5,363,550,000
EXCHANGE	4
FINANCIAL	60,000,000
HOLDING	5,308,256,057
HOLDING_HISTORY	138,948,556,646
HOLDING_SUMMARY	298,390,703
INDUSTRY	102
LAST_TRADE	4,110,000
NEWS_ITEM	6,000,000
NEWS_XREF	6,000,000
SECTOR	12
SECURITY	4,110,000
SETTLEMENT	103,680,000,000
STATUS_TYPE	5
TAXRATE	320
TRADE	103,680,000,000
TRADE_HISTORY	248,832,006,404
TRADE_REQUEST	0
TRADE_TYPE	5
WATCH_ITEM	600,034,034
WATCH_LIST	6,000,000
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.

The OS was stored on one RAID-1 array created on two 800GB 2.5” SAS SSDs in the server accessed by the internal ThinkSystem RAID 930-16i SAS/SATA controller. The OS partition was formatted NTFS.

The database log was stored on one RAID-10 array created on eight 800GB 2.5” SAS SSDs in the server accessed by the internal ThinkSystem RAID 930-16i SAS/SATA controller. The database log partitions created on this array were not formatted (RAW).

The database data was stored on six RAID-5 arrays, 17 drives each, created on 102 800GB 2.5” SAS SSDs in six Lenovo Storage D1224 enclosures accessed by six ThinkSystem RAID 930-8e SAS/SATA controllers. Each data array was broken into three partitions: one for fixed_fg (RAW), one for growing_fg (RAW), and one for extra space (NTFS).

The run-time tempdb database and log were stored on one RAID-10 array created on four 800GB 2.5” SAS SSDs in one of the Lenovo Storage D1224 enclosures mentioned above. The run-time tempdb partition was formatted NTFS.

An extra array was added to meet the benchmark’s 60-Day Space requirements. This RAID-5 array was created on three 1200GB 2.5” 10K SAS HDDs in one of the Lenovo Storage D1224 enclosures mentioned above. This space was partitioned and formatted NTFS.

In addition to the priced configuration described above, the measured configuration included additional hardware that was used during the database load process to hold load-time tempdb, flat files, and database backups. This hardware performed no function during benchmark runs:

- Six Lenovo Storage D1224 enclosures each holding 16 800GB 2.5” SAS SSDs, creating six 16-drive RAID-5 arrays, each partitioned and formatted NTFS
- Four Lenovo Storage D1224 enclosures each holding 24 1200GB 2.5” 10K SAS HDDs, creating four 24-drive RAID-5 arrays, each partitioned in half and formatted NTFS
- Two Lenovo Storage D1224 enclosures each holding 24 900GB 2.5” 10K SAS HDDs, creating two 24-drive RAID-5 arrays, each partitioned in half and formatted 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	930-16i	2 x 800GB SAS SSD internal RAID-1	C: (NTFS)	743.57GB	OS
1	930-16i	8 x 800GB SAS SSD internal RAID-10	E: (RAW) F: (NTFS) G: (RAW)	1953.22GB 534.77GB 488.38GB	tpce log MDF tpce log
2	930-8e #1	16 x 800GB SAS SSD Lenovo D1224 RAID-5 (Measured)	c:\mp\le5 (NTFS)	11,161.76GB	flatfiles & tempdb
3	930-8e #1	17 x 800GB SAS SSD Lenovo D1224 RAID-5	c:\mp\fx5 (RAW) c:\mp\gw5 (RAW) c:\mp\xt5 (NTFS)	215.04GB 10,254.10GB 1436.75GB	fixed_fg growing_fg tempdb
4	930-8e #1	24 x 900GB SAS HDD Lenovo D1224 RAID-5 (Measured)	c:\mp\bk9 (NTFS) c:\mp\bk10 (NTFS)	9,209.78GB 9,209.78GB	backup

Disk #	Controller	Drives Enclosure RAID Level (Pricing)	Partition (File System)	Size	Use
5	930-8e #2	16 x 800GB SAS SSD Lenovo D1224 RAID-5 (Measured)	c:\mp\le6 (NTFS)	11,161.76GB	flatfiles & tempdb
6	930-8e #2	17 x 800GB SAS SSD Lenovo D1224 RAID-5	c:\mp\fx6 (RAW) c:\mp\gw6 (RAW) c:\mp\xt6 (NTFS)	215.04GB 10,254.10GB 1436.75GB	fixed_fg growing_fg tempdb
7	930-8e #2	24 x 900GB SAS HDD Lenovo D1224 RAID-5 (Measured)	c:\mp\bk11 (NTFS) c:\mp\bk12 (NTFS)	9,209.78GB 9,209.78GB	backup
8	930-8e #3	17 x 800GB SAS SSD Lenovo D1224 RAID-5	c:\mp\fx4 (RAW) c:\mp\gw4 (RAW) c:\mp\xt4 (NTFS)	215.04GB 10,254.10GB 1436.75GB	fixed_fg growing_fg tempdb
9	930-8e #3	24 x 1200GB SAS HDD Lenovo D1224 RAID-5 (Measured)	c:\mp\bk7 (NTFS) c:\mp\bk8 (NTFS)	12,283.14GB 12,283.14GB	backup
10	930-8e #3	16 x 800GB SAS SSD Lenovo D1224 RAID-5 (Measured)	c:\mp\le3 (NTFS)	11,161.76GB	flatfiles & tempdb
11	930-8e #3	24 x 1200GB SAS HDD Lenovo D1224 RAID-5	S: (NTFS)	2233.18GB	60-Day
12	930-8e #4	17 x 800GB SAS SSD Lenovo D1224 RAID-5	c:\mp\fx3 (RAW) c:\mp\gw3 (RAW) c:\mp\xt3 (NTFS)	215.04GB 10,254.10GB 1436.75GB	fixed_fg growing_fg tempdb
13	930-8e #4	24 x 1200GB SAS HDD Lenovo D1224 RAID-5 (Measured)	c:\mp\bk5 (NTFS) c:\mp\bk6 (NTFS)	12,283.14GB 12,283.14GB	backup
14	930-8e #4	16 x 800GB SAS SSD Lenovo D1224 RAID-5 (Measured)	c:\mp\le2 (NTFS)	11,161.76GB	flatfiles & tempdb
15	930-8e #5	17 x 800GB SAS SSD Lenovo D1224 RAID-5	c:\mp\fx1 (RAW) c:\mp\gw1 (RAW) c:\mp\xt1 (NTFS)	215.04GB 10,254.10GB 1436.75GB	fixed_fg growing_fg tempdb
16	930-8e #5	24 x 1200GB SAS HDD Lenovo D1224 RAID-5 (Measured)	c:\mp\bk1 (NTFS) c:\mp\bk2 (NTFS)	12,283.14GB 12,283.14GB	backup
17	930-8e #5	16 x 800GB SAS SSD Lenovo D1224 RAID-5 (Measured)	c:\mp\le4 (NTFS)	11,161.76GB	flatfiles & tempdb
18	930-8e #5	4 x 800GB SAS SSD Lenovo D1224 RAID-10	T: (NTFS)	1488.13GB	run-time tempdb
19	930-8e #6	17 x 800GB SAS SSD Lenovo D1224 RAID-5	c:\mp\fx2 (RAW) c:\mp\gw2 (RAW) c:\mp\xt2 (NTFS)	215.04GB 10,254.10GB 1436.75GB	fixed_fg growing_fg tempdb
20	930-8e #6	24 x 1200GB SAS HDD Lenovo D1224 RAID-5 (Measured)	c:\mp\bk3 (NTFS) c:\mp\bk4 (NTFS)	12,283.14GB 12,283.14GB	backup

Disk #	Controller	Drives Enclosure RAID Level (Pricing)	Partition (File System)	Size	Use
21	930-8e #6	16 x 800GB SAS SSD Lenovo D1224 RAID-5 (Measured)	c:\mp\le1 (NTFS)	11,161.76GB	flatfiles & tempdb

Database Interface and Model Implemented

A statement must be provided in the Report that describes:

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

Microsoft SQL Server 2017 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 in the Report.

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 in the Report.

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 in the Report. 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 ten Ethernet ports. Four of these were provided by the onboard Gb Ethernet solution and the remaining six were provided by three dual-port 10Gb Ethernet adapters.

The Tier B database server had ten Ethernet ports. Two of these were provided by the onboard Gb Ethernet solution and the remaining eight were provided by four dual-port 10Gb Ethernet adapters.

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

Two crossover connections were setup between the Tier A client and the Driver. On the client, these cables were both plugged into one dual-port 10Gb adapter. These networks, which fulfill the mandatory network between the Driver and Tier A, were used by the client to report its results to the Driver as a benchmark run was underway.

Another network connected the Driver, the Tier B database server, the Tier A client, and a time server. This network, which was connected via a Gb Ethernet switch, used one of the onboard Gb Ethernet ports on the Tier A client and on the Tier B 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 in the Report (see Clause 5.3.1).

EGen v1.14.0 was used in the benchmark.

EGen Code and Modifications

A statement that all required TPC-provided EGen code was used in the benchmark must be reported in the Report.

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 in the Report (see Clause 5.3.7.1). If any of the changes to EGen do not have a formal waiver that must also be reported in the Report.

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 in the Report (see Clause 5.7.4).

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

EGen was not modified for 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 in the Report (see Clause 6.2.5).

There were 16 EGenDriverCEs with a total of 1312 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 in the Report (see Clause 6.7.1.2).

The Reported Throughput was 11,357.28 tpsE.

Throughput vs. Elapsed Time for Trade-Result Transaction

A Test Run Graph of throughput versus elapsed wall clock time must be reported in the Report 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 in the Report.

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 in the Report (for example 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 in the Report.

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	60.00%	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	30.00%	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		40.00%	38% to 42%
Buy On Margin	1	8.00%	7.5% to 8.5%
Rollback	1	0.99%	0.94% to 1.04%
LIFO	1	35.01%	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	20.00%	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	32.99%	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 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 0. 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 1. 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 four tests described in Clause 7.4.2. These isolation tests are designed to verify that the configuration and implementation of the System Under Test provides the Transactions with the required isolation levels defined in Clause 7.4.1.3.

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.

Data Accessibility

The System Under Test must be configured to satisfy the requirements for Data Accessibility. Data Accessibility is demonstrated by the SUT being able to maintain database operations with full data access after the permanent irrecoverable failures of any single Durable Medium containing database tables, recovery log data, or Database Metadata. Data Accessibility tests are conducted by inducing failures of Durable Media within the SUT. The failures of Clause 7.6.3 test the ability of the SUT to maintain access to the data. The specific failures addressed in Clause 7.6.3 are defined sufficiently significant to justify demonstration of Data Accessibility across such failures. However, the limited nature of the tests listed must not be interpreted to allow other unrecoverable single points of failure.

The Test Sponsor must report in the 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 in the Report.

A Data Accessibility Graph for each run demonstrating a Redundancy Level must be reported in the Report (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	930-8e
Database Log	SSD	SAS	RAID-10	930-16i
Database tempdb	SSD	SAS	RAID-10	930-8e

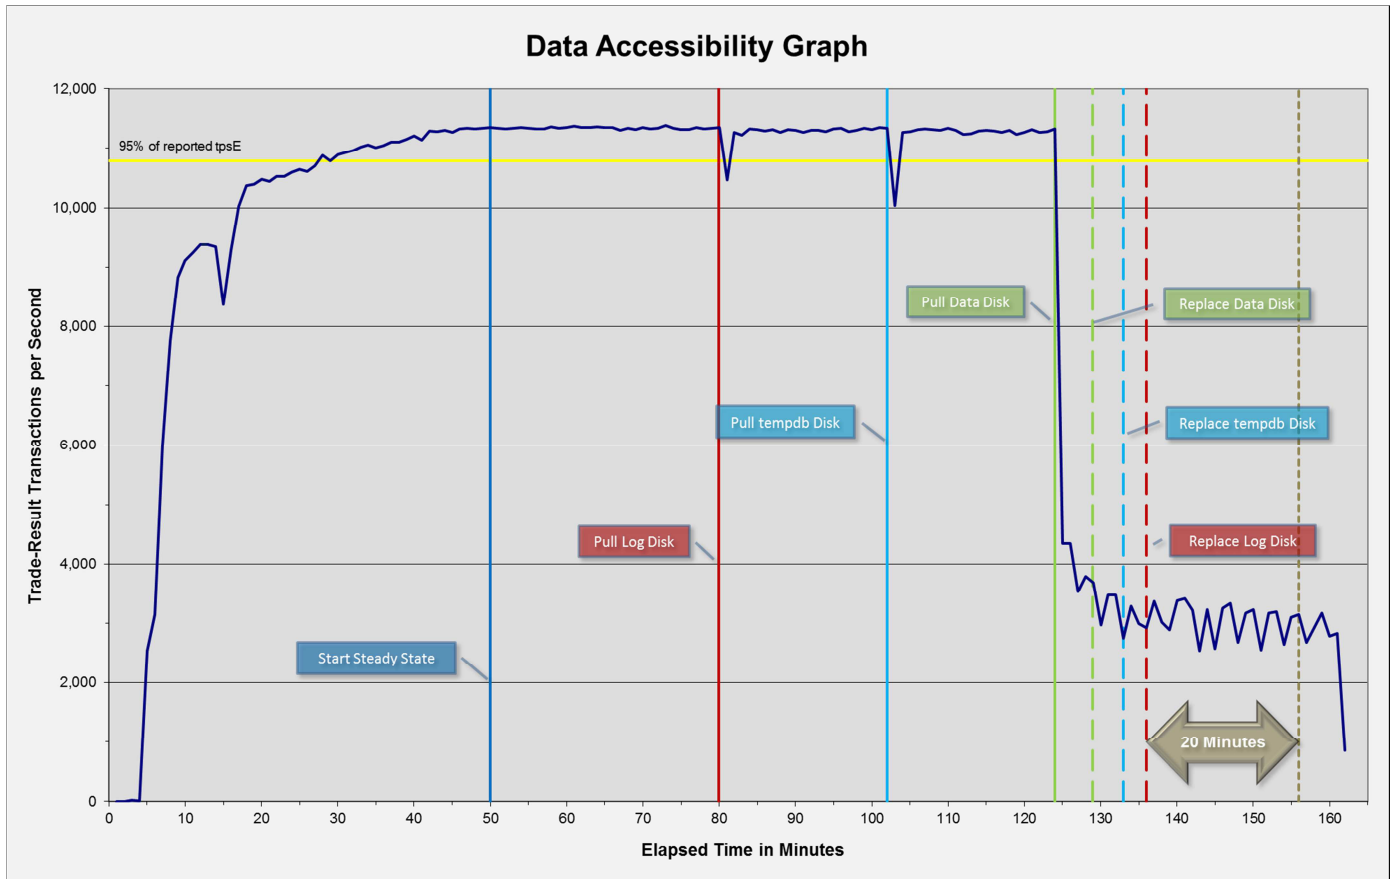
To prove Redundancy Level 1, the following steps were successfully performed:

1. Restored the database to its freshly-loaded, proven-consistent state.
2. Determined the current number of completed trades in the database, *count1*.
3. Started a run, using the profile from the measured run, with checkpoints, and met the Data Accessibility Throughput Requirements for at least 5 minutes.

4. Induced the first failure, which in this case was failing a drive in the database log array by physically removing it from its enclosure. Since the database log array was RAID protected, transaction processing continued.
5. Waited until the Data Accessibility Throughput Requirements were met again for at least 5 minutes.
6. Induced the second failure, which in this case was failing a drive in the tempdb array by physically removing it from its enclosure. Since the tempdb array was RAID protected, transaction processing continued.
7. Waited until the Data Accessibility Throughput Requirements were met again for at least 5 minutes.
8. Induced the third failure, which in this case was failing a drive in a database data array by physically removing it from its enclosure. Since the database data arrays were RAID protected, transaction processing continued.
9. After a few minutes passed, a new drive was inserted into the data enclosure to replace the failed data drive. The data array rebuilding process was started.
10. After a few minutes passed, a new drive was inserted into the tempdb enclosure to replace the failed tempdb drive. The tempdb array rebuilding process was started.
11. After a few minutes passed, a new drive was inserted into the log enclosure to replace the failed log drive. The log array rebuilding process was started.
12. Continued running the benchmark for at least 20 minutes.
13. Terminated the run gracefully.
14. Retrieved the new number of completed trades in the database by running *select count(*) as count2 from SETTLEMENT*.
15. Verified that $(count2 - count1)$, which is the number of actual completed Trade-Result Transactions done during the run, equaled the number of successful Trade-Result transactions reported by the Driver.
16. Allowed the recovery process to complete.

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

Figure 7-1. Data Accessibility Graph



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

Business Recovery

Business Recovery is the process of recovering from a Single Point of Failure and reaching a point where the business meets certain operational criteria.

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 7.5.3.1, 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 in the Report for all Business Recovery tests.

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

The following steps were successfully performed to test Business Recovery:

1. Restored the database to its freshly-loaded, proven-consistent state.
2. Determined the current number of completed trades in the database, *count1*.
3. Started a run, using the profile from the measured run, with checkpoints, and met the Durability Throughput Requirements for at least 20 minutes.
4. Pulled the power cords from the database server, 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 all of the data and log files 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 “Recovery 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 00:10:51. The Application Recovery Time was 00:25:06. The Business Recovery Time, which is the sum of the Database Recovery Time and the Application Recovery Time, was 00:35:57.

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 in the Report.

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

Table 8-1. Disk Space Requirements

Customers	6,000,000		Measured Throughput			11,357.28	Trade-Results/s	Reported Throughput			11,357.28	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	60,000	4,616	6,392	550	11,558	60,000	11,008	-	-	550		
CASH_TRANSACTION	95,385,505,113	9,936,198,384	20,941,368	497,856,988	10,454,996,740	95,496,922,969	9,978,103,912	20,964,160	56,619,889	56,619,889		
CHARGE	15	8	8	1	17	15	16	-	-	1		
COMMISSION_RATE	240	16	56	4	76	240	72	-	-	4		
SETTLEMENT	103,680,000,000	4,943,862,328	10,431,112	247,714,672	5,202,008,112	103,801,108,682	4,965,956,328	11,662,888	31,499,064	31,499,064		
TRADE	103,680,000,000	12,379,010,608	6,899,421,648	963,921,613	20,242,353,869	103,801,789,452	19,342,971,512	64,539,256	174,307,269	174,307,269		
TRADE_HISTORY	248,832,006,404	7,483,670,152	19,508,400	375,158,928	7,878,337,480	249,122,939,638	7,523,956,528	20,777,976	56,117,044	56,117,044		
TRADE_REQUEST	-	-	-	-	-	680,691	2,003,728	2,003,728	5,411,658	5,411,658		
TRADE_TYPE	5	8	1,032	52	1,092	5	1,040	-	-	52		
ACCOUNT_PERMISSION	42,597,820	2,345,536	12,272	117,890	2,475,698	42,597,820	2,357,944	136	368	117,890		
CUSTOMER	6,000,000	983,120	285,648	63,438	1,332,206	6,000,000	1,268,824	56	152	63,438		
CUSTOMER_ACCOUNT	30,000,000	2,718,336	663,216	169,078	3,550,630	30,000,000	3,381,552	-	-	169,078		
CUSTOMER_TAXRATE	12,000,000	250,144	1,584	12,586	264,314	12,000,000	251,864	136	368	12,586		
HOLDING	5,308,256,057	355,453,280	243,000,248	29,922,676	628,376,204	5,311,346,423	605,597,888	7,144,360	19,295,449	19,295,449		
HOLDING_HISTORY	138,948,556,646	5,052,675,528	3,375,378,456	421,402,699	8,849,456,683	139,111,861,021	8,452,071,296	24,017,312	64,865,825	64,865,825		
HOLDING_SUMMARY	298,390,703	13,105,992	48,472	657,723	13,812,187	298,390,647	13,154,464	-	-	-		
WATCH_ITEM	600,034,034	16,879,384	61,952	847,067	17,788,403	600,034,034	16,941,592	256	692	847,067		
WATCH_LIST	6,000,000	149,440	134,336	14,189	297,965	6,000,000	283,776	-	-	14,189		
COMPANY	3,000,000	640,744	188,808	41,478	871,030	3,000,000	829,552	-	-	41,478		
COMPANY_COMPETITOR	9,000,000	241,544	218,232	22,989	482,765	9,000,000	459,776	-	-	22,989		
DAILY_MARKET	5,363,550,000	251,844,968	735,888	12,629,043	265,209,899	5,363,550,000	252,582,096	1,240	3,349	12,629,043		
EXCHANGE	4	8	8	1	17	4	16	-	-	1		
FINANCIAL	60,000,000	6,761,352	18,848	339,010	7,119,210	60,000,000	6,780,496	296	800	339,010		
INDUSTRY	102	8	24	2	34	102	32	-	-	2		
LAST_TRADE	4,110,000	256,136	1,352	12,874	270,362	4,110,000	257,488	-	-	12,874		
NEWS_ITEM	6,000,000	650,509,408	6,720	32,525,806	683,041,934	6,000,000	650,516,144	16	44	32,525,806		
NEWS_XREF	6,000,000	149,376	1,576	7,548	158,500	6,000,000	150,952	-	-	7,548		
SECTOR	12	8	24	2	34	12	32	-	-	2		
SECURITY	4,110,000	570,536	156,832	36,368	763,736	4,110,000	727,384	16	44	36,368		
STATUS_TYPE	5	8	8	1	17	5	16	-	-	1		
ADDRESS	9,000,004	519,040	1,648	26,034	546,722	9,000,004	520,752	64	173	26,034		
TAXRATE	320	24	56	4	84	320	96	16	44	44		
ZIP_CODE	14,741	488	56	27	571	14,741	544	-	-	27		
TOTALS (KB)		41,098,800,528	10,571,226,280	2,583,501,340	54,253,528,148		51,821,138,720	151,111,912	408,122,232	454,982,280		
Initial Database Size (MB)		50,459,011	49,276 GB									
Database Filegroups	LUN Count	Partition Size (MB)	MB Allocated	MB Loaded	MB Required							
	0	-	-	-	-					OK		
growing_fg	6	10,500,100	63,000,600	49,543,658	49,942,209					OK		
	0	-	-	-	-					OK		
fixed_fg	6	220,100	1,320,600	915,352	961,120					OK		
Settlements	121,108,682											
Data Space Required (MB)	Data Space Configured (MB)	Log Space Required (MB)	Log Space Configured (MB)									
Initial Growing Space	49,543,658	Data LUNS	6	1	-	Initial Log Size	39,489	Log LUNS	1			
Final Growing Space	49,691,226	Delta	17	3	-	Final Log Size	853,611	Log Disks	8			
Delta	147,568	Disk Capacity	761,984	1,141,899	-	Log Growth	814,122	Disk Capacity	761,984			
Data Space per Trade	0.001218476	RAID Overhead	94.12%	67%	0%	Log Growth/Trade	0.006722246	RAID Overhead	50%			
1 Day Data Growth	398,551	Total Space				1 Day Log Space	2,238,266	Log Space	3,047,936			
60 Day Space	74,372,069											

OK

OK

Hardware and Software Components

A detailed list of hardware and software used in the Priced Configuration must be reported. The listing for each separately Orderable item must have vendor Part Number, description and applicable release/revision level, price source, unit price, quantity, extended price, applicable Discounted price and 3-year maintenance price. If package-pricing is used, the vendor Part Number of the package and a description uniquely identifying each of the Components of the package must be disclosed to a sufficient level of detail to meet the requirements of 1.4.1.1.

A detailed list of all hardware, software, and maintenance is provided in the Executive Summary at the front of this report. Price quotations are included in Appendix A.

Three-Year Cost of System Configuration

The total 3-year price of the entire Priced Configuration must be reported, including: hardware, software, and maintenance charges. The justification of any Discounts applied must be disclosed in the price sheet. Sufficient detail of what items are being discounted and by how much they are being discounted must be provided so that the Discount amount used in the computation of the total system cost can be independently reproduced.

A detailed list of all hardware, software, and maintenance, including the total 3-year price and discount information, is provided in the Executive Summary at the front of this report. Price quotations are included in Appendix A.

Availability Date

The committed Availability Date of Line Items used in the price calculations must be reported. The Availability Date must be reported on the first page of the Executive Summary and with a precision of one day. When the priced system includes products and/or Licensed Compute Services with different Availability Dates, the reported Availability Date for the priced system must be a date at which all Line Items are committed to be Generally Available. Each Line Item used in the Priced Configuration is considered to be Available on the Availability Date unless an earlier date is specified.

For each of the Line Items 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 Line Item can be ordered (on or before the Availability Date)*
- *The method to be used to order the Component (at or below the quoted price) when the order date arrives*
- *The method for verifying the price*

The total solution as priced will be generally available November 6, 2017.

Country-Specific Pricing

Pricing must be reported in the currency of the country where the system is priced.

The configuration is priced for the United States of America.

Pricing Calculations

A statement of the benchmark performance metric, as well as the respective calculations for 3-year pricing, price/performance, and the availability date must be included.

The performance metric, pricing calculations, price/performance, and availability dates are all included in the Executive Summary.

Supporting Files Index

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

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.

Dilip Pendyal
DCG Server Performance
Lenovo Data Center Group
8001 Development Drive
Morrisville, NC 27560

October 17, 2017

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

Platform: Lenovo ThinkSystem SR950
Operating System: Microsoft Windows Server 2016 Standard Edition
Database Manager: Microsoft SQL Server 2017 Enterprise Edition

The results were:

Performance Metric **11,357.28 tpsE**
Trade-Result 90th %-tile 0.04 Seconds

Tier B (Server)

Lenovo ThinkSystem SR950

CPU	4 x Intel Xeon Platinum 8180 Processor (2.50 GHz, 28-core, 38.5 MB L3)		
Memory	3,072 GB		
Storage	Qty	Size	Type
	116	800 GB	SAS SSD
	3	1,200 GB	10K rpm SAS HDD

Tier A (Client)

Lenovo System x3650 M5

CPU	2 x Intel Xeon Processor E5-2699 v4 (2.20 GHz, 22-core, 55 MB L3)
Memory	64 GB
Storage	2 x 600 GB 10K rpm SAS HDD

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

The following verification items were given special attention:

- All EGen components were verified to be v1.14.0
- The transactions were correctly implemented
- The database was properly scaled and populated for 6,000,000 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 00:35:57 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 (17) older Lenovo Storage 800 GB SAS SSD disks that were substituted by (17) newer Lenovo Storage 800 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.

Respectfully Yours,

A handwritten signature in black ink that reads "Doug Johnson". The signature is written in a cursive style with a long horizontal flourish extending to the right.

Doug Johnson, Certified TPC Auditor

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
October 10, 2017

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 (\$).

Description	Unit Price	Quantity	Price
Database Management System			
SQL Server 2017 Enterprise Edition 2 Core License Open Program - Level C Unit Price reflects a 6% discount from the retail unit price of \$14,256	\$13,472.50	56	\$754,460.00
Database Server Operating System			
Windows Server 2016 Standard Edition 2 Core License Open Program - Level C Unit Price reflects a 17% discount from the retail unit price of \$110.25.	\$92.00	56	\$5,152.00
Tier-A Operating System(s)			
Windows Server 2016 Standard Edition 2 Core License Open Program - Level C Unit Price reflects a 17% discount from the retail unit price of \$110.25.	\$92.00	22	\$2,024.00
Support			
Microsoft Problem Resolution Services Professional Support (1 Incident).	\$259.00	1	\$259.00

All software components are currently orderable and available. 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_qhtplylGYLKTUVUK876923wwbo_2017_Ire.