

TPC Benchmark™ E
Full Disclosure Report
for
Lenovo® ThinkSystem™ SR650
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
April 2, 2019

First Edition – April 2019

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 2019. 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 SR650. 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 SR650 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 SR650	Microsoft SQL Server 2017 Enterprise Edition Microsoft Windows Server 2016 Standard Edition	\$638,052	7,012.53	\$90.99	April 17, 2019

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™ SR650		TPC-E™ 1.14.0 TPC Pricing 2.4.0
			Report Date: April 2, 2019 Revision Date: April 2, 2019
TPC-E Throughput 7,012.53 tpsE™	Price/Performance \$90.99 USD per tpsE	Availability Date April 17, 2019	Total System Cost \$638,052 USD
Database Server Configuration			
Operating System Microsoft® Windows Server® 2016 Standard Edition	Database Manager Microsoft SQL Server® 2017 Enterprise Edition	Processors/Cores/Threads 2/56/112	Memory 1536GB
<div>  </div> <div> <div>  <div> 4 x 10Gb Ethernet </div> </div> <div>  <div> 4 x 12Gb SAS 8 </div> </div> </div> <div> <div> TIER A (Client) 1 x Lenovo ThinkSystem SR650, with: <ul style="list-style-type: none"> - 2 x Intel® Xeon® Platinum 8168 Processor 2.70GHz (2 Procs/48 Cores/96 Threads) - 96GB Memory - 2 x 300GB SFF SAS (RAID-1) - 1 x ThinkSystem RAID 930-8i - 1 x Onboard Dual Gb Ethernet - 1 x Intel x550 Dual 10Gb-T Ethernet - 2 x Broadcom NX-E Dual 10Gb-T Enet </div> <div> TIER B (Database Server and Storage) 1 x Lenovo ThinkSystem SR650, with: <ul style="list-style-type: none"> - 2 x Intel Xeon Platinum 8280 Processor 2.70GHz (2 Procs/56 Cores/112 Threads) - 1536GB Memory - 2 x 800GB SAS SSD (RAID-1) - 6 x 800GB SAS SSD (RAID-10) - 1 x ThinkSystem RAID 930-8i - 4 x ThinkSystem RAID 930-8e - 1 x Onboard Dual Gb Ethernet - 2 x Intel x550 Dual 10Gb-T Ethernet </div> <div> 4 x Lenovo Storage D1224 DAS Enclosures, with: <ul style="list-style-type: none"> - 74 x 800GB 2.5" SAS SSD (4 x 17-drive RAID-5) (1 x 6-drive RAID-10) </div> </div> <div> 74 Total External Drives </div>			
Initial Database Size 29,145 GB	Redundancy Level: 1 RAID-10 Log, tempdb RAID-5 Data	Storage 82 x 800GB 2.5" SAS SSD	



Lenovo ThinkSystem SR650

TPC-E 1.14.0 TPC Pricing 2.4.0

Report Date: April 2, 2019
Revision Date: April 2, 2019

Availability Date:
April 17, 2019

Description	Part Number	Price Source	Unit Price	Quantity	Extended Price	3-Yr. Maint. Price
Server Hardware						
ThinkSystem SR650 Configure-To-Order, includes:	7X06CT01WW	1	141,024	1	141,024	
ThinkSystem SR650 2.5" Chassis with 8, 16 or 24 bays	AUVV			1		
Intel Xeon Platinum 8280 28C 205W 2.7GHz Processor	B4H7			2		
ThinkSystem 64GB TruDDR4 Performance+ 2933MHz (2Rx4 1.2V) RDIMM	B5N8			24		
ThinkSystem 2U/Twr 2.5" SATA/SAS 8-Bay Backplane	AURA			1		
ThinkSystem RAID 930-8i 2GB Flash PCIe 12Gb Adapter	AUNJ			1		
ThinkSystem 2.5" PM1645 800GB Mainstream SAS 12Gb Hot Swap SSD	B4A0			2		
ThinkSystem 2.5" HUSMM32 800GB Performance SAS 12Gb Hot Swap SSD	AUMH			6		
ThinkSystem 2U x8/x8/x8 PCIe FH Riser 1	AUR4			1		
ThinkSystem SR550/SR590/SR650 (x16/x8)/(x16/x16) PCIe FH Riser 2 Kit	AURC			1		
ThinkSystem 1Gb 2-port RJ45 LOM	AUKG			1		
Intel X550-T2 Dual Port 10GBase-T Adapter	ATPX			2		
ThinkSystem RAID 930-8e 4GB Flash PCIe 12Gb Adapter	AUNQ			4		
ThinkSystem Toolless Slide Rail	AXCA			1		
ThinkSystem 2U left EIA Latch Standard	AURD			1		
ThinkSystem 1100W (230V/115V) Platinum Hot-Swap Power Supply	AWWF			2		
2.8m, 13A/100-250V, C13 to C14 Jumper Cord	6400			2		
Lenovo ThinkSystem 2U 3FH Riser Bracket	AURQ			1		
Lenovo ThinkSystem 2U 2FH Riser Bracket	AURP			1		
FBU345 SuperCap	AUNP			5		
ThinkSystem SR650 Refresh MB	B4NL			1		
Essential Service - 3Yr 24x7 4Hr Response + YourDrive YourData (SR650)	5PS7A01558	1	2,079	1		2079
Subtotal					141,024	2,079
Server Storage						
Lenovo Storage D1224 SFF Chassis, Dual ESM	4587A31	1	2,999	4	11,996	
External MiniSAS HD 8644/MiniSAS HD 8644 1M Cable	00YL848	1	59	8	472	
Lenovo Storage 800GB 10 DWD 2.5" SAS SSD	01DC452	1	2,999	74	221,926	
Essential Service - 3Yr 24x7 4Hr Response (D1224)	01JR577	1	1,169	4		4,676
Subtotal					234,394	4,676
Server Software						
SQL Server 2017 Enterprise Edition (2 Core License)	N/A	2	13,472.50	28	377,230	
Windows Server 2016 Standard Edition (2 Core License)	N/A	2	92	28	2,576	
Microsoft Problem Resolution Services (1 Incident)	N/A	2	259	1		259
Subtotal					379,806	259
Client Hardware						
ThinkSystem SR650 Configure-To-Order, includes:	7X06CT01WW	1	32,215	1	32,215	
ThinkSystem SR650 2.5" Chassis with 8, 16 or 24 bays	AUVV			1		
Intel Xeon Platinum 8168 24C 205W 2.7GHz Processor	AWDJ			2		
ThinkSystem 8GB TruDDR4 2666 MHz (1Rx8 1.2V) RDIMM	AUU1			12		
ThinkSystem 2U/Twr 2.5" SATA/SAS 8-Bay Backplane	AURA			1		
ThinkSystem RAID 930-8i 2GB Flash PCIe 12Gb Adapter	AUNJ			1		
ThinkSystem 2.5" 300GB 10K SAS 12Gb Hot Swap 512n HDD	AULY			2		
ThinkSystem 2U x8/x8/x8 PCIe FH Riser 1	AUR4			1		
ThinkSystem SR550/SR590/SR650 (x16/x8)/(x16/x16) PCIe FH Riser 2 Kit	AURC			1		
ThinkSystem 1Gb 2-port RJ45 LOM	AUKG			1		
ThinkSystem Broadcom 57416 10GBASE-T 2-Port PCIe Ethernet Adapter	AUKP			2		
Intel X550-T2 Dual Port 10GBase-T Adapter	ATPX			1		
ThinkSystem 1100W (230V/115V) Platinum Hot-Swap Power Supply	AWWF			2		
2.8m, 13A/100-250V, C13 to C14 Jumper Cord	6400			2		
ThinkSystem Toolless Friction Rail	AXFM			1		
ThinkSystem 2U left EIA Latch Standard	AURD			1		
Lenovo ThinkSystem Mainstream MB - 2U	AUQB			1		
Lenovo ThinkSystem 2U 3FH Riser Bracket	AURQ			1		
Lenovo ThinkSystem 2U 2FH Riser Bracket	AURP			1		
FBU345 SuperCap	AUNP			1		
Essential Service - 3Yr 24x7 4Hr Response + YourDrive YourData (SR650)	5PS7A01558	1	2,079	1		2,079
Subtotal					32,215	2,079
Client Software						
Windows Server 2019 Standard Edition (2 Core License)	N/A	2	123	24	2,952	
Subtotal					2,952	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 NEMA L6-30P line cord	00YJ776	1	479	1	479	
ThinkSystem Keyboard w/ Int. Pointing Device USB - US Eng 103P RoHS v2	7ZB7A05470	1	99	1	99	
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					2,338	315
Total					792,729	9,408
Dollar Volume Discount (See Note 1)					164,085	
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.						
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:						\$638,052
TPC-E Throughput:						7,012.53
\$ USD/tpsE:						\$90.99



Lenovo ThinkSystem SR650

TPC-E 1.14.0
TPC Pricing 2.4.0

Report Date: April 2, 2019
Revision Date: April 2, 2019

Availability Date:
April 17, 2019

Numerical Quantities Summary

Reported Throughput:

7,012.53 tpsE

Configured Customers:

3,550,000

Response Time (in seconds)	Minimum	Average	90 Th Percentile	Maximum
Broker-Volume	0.01	0.01	0.02	0.31
Customer-Position	0.01	0.01	0.02	3.62
Market-Feed	0.01	0.02	0.04	3.62
Market-Watch	0.01	0.01	0.01	0.28
Security-Detail	0.01	0.01	0.01	0.27
Trade-Lookup	0.01	0.04	0.06	0.34
Trade-Order	0.01	0.03	0.05	0.32
Trade-Result	0.01	0.03	0.05	0.31
Trade-Status	0.01	0.01	0.01	0.31
Trade-Update	0.01	0.05	0.07	0.35
Data-Maintenance	0.01	0.02	N/A	0.04

Transaction Mix	Transaction Count	Mix %
Broker-Volume	24,739,811	4.900%
Customer-Position	65,635,811	13.000%
Market-Feed	5,049,028	1.000%
Market-Watch	90,880,788	18.000%
Security-Detail	70,684,757	14.000%
Trade-Lookup	40,391,307	8.000%
Trade-Order	50,994,302	10.100%
Trade-Result	50,490,226	10.000%
Trade-Status	95,929,354	19.000%
Trade-Update	10,097,957	2.000%
Data-Maintenance	120	N/A

Test Duration and Timings

Ramp-up Time (hh:mm:ss)	03:05:33
Measurement Interval (hh:mm:ss)	02:00:00
Business Recovery Time (hh:mm:ss)	00:25:39
Total Number of Transactions Completed in Measurement Interval	504,893,341

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
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.....	17
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)	TIER B (Database Server and Storage)	
1 x Lenovo ThinkSystem SR650, with: <ul style="list-style-type: none"> - 2 x Intel Xeon Platinum 8168 Processor 2.70GHz (2 Procs/48 Cores/96 Threads) - 96GB Memory - 2 x 300GB SFF SAS (RAID-1) - 1 x ThinkSystem RAID 930-8i - 1 x Onboard Dual Gb Ethernet - 1 x Intel x550 Dual 10Gb-T Ethernet - 2 x Broadcom NX-E Dual 10Gb-T Enet 	1 x Lenovo ThinkSystem SR650, with: <ul style="list-style-type: none"> - 2 x Intel Xeon Platinum 8280 Processor 2.70GHz (2 Procs/56 Cores/112 Threads) - 1536GB Memory - 2 x 800GB SAS SSD (RAID-1) - 6 x 800GB SAS SSD (RAID-10) - 1 x ThinkSystem RAID 930-8i - 4 x ThinkSystem RAID 930-8e - 1 x Onboard Dual Gb Ethernet - 2 x Intel x550 Dual 10Gb-T Ethernet 	12 x Lenovo Storage D1224 DAS Enclosures, with: <ul style="list-style-type: none"> - 140 x 800GB 2.5" SAS SSD (6 x 17-drive RAID-5) - (1 x 6-drive RAID-10) - (2 x 16-drive RAID-5) - 99 x 1200GB 2.5" 10K SAS (4 x 24-drive RAID-5) - (1 x 3-drive RAID-5)
		239 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.

No components were substituted.

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_SR650_Setup.pdf
- Information specific to the Tier B database server and storage can be found in:
SupportingFiles\Introduction\TierB\TierB_SR650_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\3550000.Cust\Database:

- Tempdb_load.sql specifies the tempdb files to use when loading the database
- Tempdb_run.sql specifies the tempdb files to use when running the database
- Shrinktempdb.sql removes tempdb files
- Backupdev.sql creates SQL Server database backup devices
- 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 defines the locations of the flatfiles during the database 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 3,550,000 customers. The cardinality is shown in Table 2-1.

Table 2-1. Initial Cardinality of Tables

Table Name	Rows
ACCOUNT_PERMISSION	25,206,358
ADDRESS	5,325,004
BROKER	35,500
CASH_TRANSACTION	56,436,535,937
CHARGE	15
COMMISSION_RATE	240
COMPANY	1,775,000
COMPANY_COMPETITOR	5,325,000
CUSTOMER	3,550,000
CUSTOMER_ACCOUNT	17,750,000
CUSTOMER_TAXRATE	7,100,000
DAILY_MARKET	3,173,433,750
EXCHANGE	4
FINANCIAL	35,500,000
HOLDING	3,140,759,821
HOLDING_HISTORY	82,211,064,212
HOLDING_SUMMARY	176,556,595
INDUSTRY	102
LAST_TRADE	2,431,750
NEWS_ITEM	3,550,000
NEWS_XREF	3,550,000
SECTOR	12
SECURITY	2,431,750
SETTLEMENT	61,344,000,000
STATUS_TYPE	5
TAXRATE	320
TRADE	61,344,000,000
TRADE_HISTORY	147,225,714,782
TRADE_REQUEST	0
TRADE_TYPE	5
WATCH_ITEM	355,072,194
WATCH_LIST	3,550,000
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 930-8i SAS/SATA controller. The OS partition was formatted NTFS.

The database log was stored on one RAID-10 array created on six 800GB 2.5” SAS SSDs in the server accessed by the internal ThinkSystem RAID 930-8i SAS/SATA controller. This array was broken into two partitions, one RAW and one NTFS.

The database data was stored on four RAID-5 arrays, 17 drives each, created on 68 800GB 2.5” SAS SSDs in four Lenovo Storage D1224 enclosures accessed by four 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 six 800GB 2.5” SAS SSDs in one of the Lenovo Storage D1224 enclosures mentioned above. The run-time tempdb partition was 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:

- Two Lenovo Storage D1224 enclosures each holding 16 800GB 2.5” SAS SSDs, creating two 16-drive RAID-5 arrays, each partitioned and formatted NTFS
- Two Lenovo Storage D1224 enclosures each holding 17 800GB 2.5” SAS SSDs, creating two 17-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
- One 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.

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	930-8e #1	6 x 800GB SAS SSD Lenovo D1224 RAID-10	T: (NTFS)	2,232.38GB	run-time tempdb
1	930-8e #1	17 x 800GB SAS SSD Lenovo D1224 RAID-5	c:\mp\fx1 (RAW) c:\mp\gw1 (RAW) c:\mp\xt1 (NTFS)	191.41GB 9,277.34GB 2,437.26GB	fixed_fg growing_fg tempdb
2	930-8e #1	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
3	930-8e #1	16 x 800GB SAS SSD Lenovo D1224 RAID-5 (Measured)	c:\mp\le1 (NTFS)	11,161.89GB	flatfiles & tempdb
4	930-8e #2	17 x 800GB SAS SSD Lenovo D1224 RAID-5	c:\mp\fx2 (RAW) c:\mp\gw2 (RAW) c:\mp\xt2 (NTFS)	191.41GB 9,277.34GB 2,437.26GB	fixed_fg growing_fg tempdb

Disk #	Controller	Drives Enclosure RAID Level (Pricing)	Partition (File System)	Size	Use
5	930-8e #2	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
6	930-8e #2	16 x 800GB SAS SSD Lenovo D1224 RAID-5 (Measured)	c:\mp\le2 (NTFS)	11,161.89GB	flatfiles & tempdb
7	930-8i	6 x 800GB SAS SSD internal RAID-10	E: (RAW) F: (NTFS)	1,562.50GB 669.75GB	tpce log MDF
8	930-8i	2 x 800GB SAS SSD internal RAID-1	C: (NTFS)	743.57GB	OS
9	930-8e #3	17 x 800GB SAS SSD Lenovo D1224 RAID-5	c:\mp\fx3 (RAW) c:\mp\gw3 (RAW) c:\mp\xt3 (NTFS)	191.41GB 9,277.34GB 2,437.26GB	fixed_fg growing_fg tempdb
10	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
11	930-8e #3	17 x 800GB SAS SSD Lenovo D1224 RAID-5 (Measured)	c:\mp\le3 (NTFS)	11,905.89GB	flatfiles & tempdb
12	930-8e #3	3 x 1200GB SAS HDD Lenovo D1224 RAID-5 (Measured)	S: (NTFS)	2,233.18GB	Unused
13	930-8e #4	17 x 800GB SAS SSD Lenovo D1224 RAID-5	c:\mp\fx4 (RAW) c:\mp\gw4 (RAW) c:\mp\xt4 (NTFS)	191.41GB 9,277.34GB 2,437.26GB	fixed_fg growing_fg tempdb
14	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
15	930-8e #4	17 x 800GB SAS SSD Lenovo D1224 RAID-5 (Measured)	c:\mp\le4 (NTFS)	11,905.89GB	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 eight Ethernet ports. Two 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 six Ethernet ports. Two of these were provided by the onboard Gb Ethernet solution and the remaining four were provided by two 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 both ports of each 10Gb adapter 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 800 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 7,012.53 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	50.01%	48% to 52%
Get History	1	50.01%	48% to 52%
Market-Watch			
Securities chosen by	Watch List	60.01%	57% to 63%
	Account ID	34.99%	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	10.01%	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	8.00%	7.5% to 8.5%
Rollback	1	0.99%	0.94% to 1.04%
LIFO	1	34.99%	33% to 37%
Trade Quantity	100	25.00%	24% to 26%
	200	25.00%	24% to 26%
	400	25.00%	24% to 26%
	800	25.00%	24% to 26%
Trade Type	Market Buy	30.01%	29.7% to 30.3%
	Market Sell	30.00%	29.7% to 30.3%
	Limit Buy	19.99%	19.8% to 20.2%
	Limit Sell	10.00%	9.9% to 10.1%
	Stop Loss	10.00%	9.9% to 10.1%
Trade-Update			
Frame to execute	1	33.02%	31% to 35%
	2	33.01%	31% to 35%
	3	33.97%	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	930-8i
Database tempdb	SSD	SAS	RAID-10	930-8e


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

1. Performed Trade-Cleanup to remove remnants of previous benchmark runs from the database.
2. Determined the current number of completed trades in the database, *count1*.
3. Started a run, using the profile from the measured run, with checkpoints, and met the Data Accessibility Throughput Requirements for at least 5 minutes.

4. Induced the first failure, which in this case was failing a drive in the database log array by physically removing it from its enclosure. Since the database log array 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 00:15:33. The Application Recovery Time was 00:10:06. The Business Recovery Time, which is the sum of the Database Recovery Time and the Application Recovery Time, was 00:25:39.

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\tpce_space.xls.

Table 8-1. Disk Space Requirements

Customers	3,550,000	Measured Throughput			7,012.53	Trade-Results/s		Reported Throughput		7,012.53	tpsE		
Table	Initial Rows	Data Size (KB)	Index Size (KB)	Extra 5% (KB)	Total + 5% (KB)	Rows After	After Run (KB)	Growth (KB)	Bus. Day Growth (KB)	Req. Add. (KB)			
BROKER	35,500	2,592	2,552	257	5,401	35,500	5,144	-	-	257			
CASH_TRANSACTION	56,436,535,937	5,871,330,136	12,374,280	294,185,221	6,177,889,637	56,575,555,030	5,909,883,392	26,178,976	34,988,739	34,988,739			
CHARGE	15	8	8	1	17	15	16	-	-	1			
COMMISSION_RATE	240	16	40	3	59	240	56	-	-	3			
SETTLEMENT	61,344,000,000	2,925,119,888	6,169,928	146,564,491	3,077,854,307	61,495,109,438	2,945,864,264	14,574,448	19,479,049	19,479,049			
TRADE	61,344,000,000	7,323,828,952	4,080,060,576	570,194,476	11,974,084,004	61,495,939,193	11,461,296,504	57,406,976	76,725,602	76,725,602			
TRADE_HISTORY	147,225,714,782	4,427,841,504	11,544,848	221,969,318	4,661,355,670	147,589,526,046	4,465,385,544	25,999,192	34,748,454	34,748,454			
TRADE_REQUEST	-	-	-	-	-	420,916	2,108,840	2,108,840	2,818,508	2,818,508			
TRADE_TYPE	5	8	1,032	52	1,092	5	1,040	-	-	52			
ACCOUNT_PERMISSION	25,206,358	1,387,944	7,176	69,756	1,464,876	25,206,358	1,395,240	120	161	69,756			
CUSTOMER	3,550,000	581,704	169,392	37,555	788,651	3,550,000	751,152	56	75	37,555			
CUSTOMER_ACCOUNT	17,750,000	1,608,320	391,768	100,004	2,100,092	17,750,000	2,000,088	-	-	100,004			
CUSTOMER_TAXRATE	7,100,000	148,024	928	7,448	156,400	7,100,000	149,208	256	343	7,448			
HOLDING	3,140,759,821	210,139,448	143,661,128	17,690,029	371,490,605	3,144,611,191	362,985,464	9,184,888	12,275,792	12,275,792			
HOLDING_HISTORY	82,211,064,212	2,989,493,704	1,997,094,984	249,329,434	5,235,918,122	82,414,824,247	5,018,631,536	32,042,848	42,825,924	42,825,924			
HOLDING_SUMMARY	176,556,595	7,747,072	28,496	388,778	8,164,346	176,557,764	7,776,464	896	1,198	1,198			
WATCH_ITEM	355,072,194	9,976,712	36,520	500,662	10,513,894	355,072,194	10,013,744	512	685	500,662			
WATCH_LIST	3,550,000	88,384	79,440	8,391	176,215	3,550,000	167,824	-	-	8,391			
COMPANY	1,775,000	378,904	111,984	24,544	515,432	1,775,000	490,912	24	33	24,544			
COMPANY_COMPETITOR	5,325,000	142,896	128,744	13,582	285,222	5,325,000	271,640	-	-	13,582			
DAILY_MARKET	3,173,433,750	148,901,424	434,512	7,466,797	156,802,733	3,173,433,750	149,338,336	2,400	3,208	7,466,797			
EXCHANGE	4	8	8	1	17	4	16	-	-	1			
FINANCIAL	35,500,000	4,000,456	10,992	200,572	4,212,020	35,500,000	4,012,000	552	738	200,572			
INDUSTRY	102	8	24	2	34	102	32	-	-	2			
LAST_TRADE	2,431,750	151,584	944	7,626	160,154	2,431,750	152,528	-	-	7,626			
NEWS_ITEM	3,550,000	384,884,672	3,584	19,244,413	404,132,669	3,550,000	384,888,328	72	97	19,244,413			
NEWS_XREF	3,550,000	88,328	912	4,462	93,702	3,550,000	89,240	-	-	4,462			
SECTOR	12	8	24	2	34	12	32	-	-	2			
SECURITY	2,431,750	337,312	92,856	21,508	451,676	2,431,750	430,232	64	86	21,508			
STATUS_TYPE	5	8	8	1	17	5	16	-	-	1			
ADDRESS	5,325,004	307,064	1,040	15,405	323,509	5,325,004	308,232	128	172	15,405			
TAXRATE	320	24	40	3	67	320	80	16	22	22			
ZIP_CODE	14,741	488	40	26	554	14,741	528	-	-	26			
TOTALS (KB)		24,308,487,600	6,252,408,808	1,528,044,820	32,088,941,228		30,728,397,672	167,501,264	223,868,886	251,586,358			
Initial Database Size (MB)		29,844,625	29,145 GB										
Database Filegroups	LUN Count	Partition Size (MB)	MB Allocated	MB Loaded	MB Required								
	0	-	-	-	-	OK							
growing_fg	4	9,318,800	37,275,200	29,303,159	29,521,776	OK							
	0	-	-	-	-	OK							
fixed_fg	4	195,300	781,200	541,466	568,540	OK							
Settlements		151,109,438											
Data Space Required (MB)		Data Space Configured (MB)					Log Space Required (MB)		Log Space Configured (MB)				
Initial Growing Space	29,303,159						Initial Log Size	50,085	Log LUNS	1			
Final Growing Space	29,466,730	Data LUNS	4	-	-	-	Final Log Size	1,059,223	Log Disks	6			
Delta	163,571	Disks per LUN	17	-	-	-	Log Growth	1,009,138	Disk Capacity	761,989			
Data Space per Trade	0.001082469	Disk Capacity	761,989	-	-	-	Log Growth/Trade	0.006678194	RAID Overhead	50%			
1 Day Data Growth	218,616	RAID Overhead	94.12%	0%	0%	0%	Log Growth/Trade	0.006678194	RAID Overhead	50%			
60 Day Space	42,961,613	Total Space				48,767,304	1 Day Log Space	1,398,819	Log Space	2,285,967			
						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 April 17, 2019.

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
7001 Development Drive
Morrisville, NC 27560

March 27, 2019

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

Platform: Lenovo ThinkSystem SR650
Operating System: Microsoft Windows Server 2016 Standard Edition
Database Manager: Microsoft SQL Server 2017 Enterprise Edition

The results were:

Performance Metric **7,012.53 tpsE**
Trade-Result 90th %-tile 0.05 Seconds

Tier B (Server)

Lenovo ThinkSystem SR650

CPU	2 x Intel Xeon Platinum 8280 Processor (2.70 GHz, 28-core, 38.5 MB L3)		
Memory	1,536 GB		
Storage	Qty	Size	Type
	82	800 GB	SAS SSD

Tier A (Client)

Lenovo ThinkSystem SR650

CPU	2 x Intel Xeon Platinum 8168 Processor (2.70 GHz, 24-core, 33 MB L3)
Memory	96 GB
Storage	2 x 300 GB 10K rpm SAS HDD

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

The following verification items were given special attention:

- All EGen components were verified to be v1.14.0
- The transactions were correctly implemented
- The database was properly scaled and populated for 3,550,000 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 00:25:39 was correctly measured
- The 60-day storage requirement was correctly computed
- The system pricing was verified for major components and maintenance

Additional Audit Notes:

None.

Respectfully Yours,

A handwritten signature in black ink, reading "Doug Johnson". The signature is fluid and cursive, 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
March 26, 2019

Lenovo Corporation
Ray Engler
7001 Development Drive
Morrisville, NC 27560

Here is the information you requested regarding pricing for 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 <i>2 Core License Open Program - Level C Unit Price reflects a 5% discount from the retail unit price of \$14,256.00</i>	\$13,472.50	28	\$377,230.00
Database Server Operating System			
Windows Server 2016 Standard <i>2 Core License Open Program - Level C Unit Price Reflects a 17% discount from the retail unit price of \$110.25</i>	\$92.00	28	\$2,576.00
Tier-A Operating System(s)			
Windows Server 2019 Standard <i>2 Core License</i>	\$123.00	24	\$2,952.00
Support			
Microsoft Problem Resolution Services <i>Professional Support (1 Incident).</i>	\$259.00	1	\$259.00

All Microsoft 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: TPC-E_jQs+jifUh4zcv1mjwRBRAg@@5081__lre.doc