

Inspur Electronic Information Industry Co., Ltd

TPC Benchmark™ H
Full Disclosure Report
for
Inspur NF8380M5
using
Microsoft SQL Server2017 Enterprise Edition
and
Windows Server 2016 Standard Edition

**First Edition
October 2018**

Inspur Electronic Information Industry Co.,Ltd (Inspur), the Sponsor of this benchmark test, believes that the information in this document is accurate as of the publication date. The information in this document is subject to change without notice. The Sponsor assumes no responsibility for any errors that may appear in this document.

The pricing information in this document is believed to accurately reflect the current prices as of the publication date. However, the Sponsor provides no warranty of the pricing information in this document.

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

All performance data contained in this report was obtained in a rigorously controlled environment. Results obtained in other operating environments may vary significantly. No warranty of system performance or price/performance is expressed or implied in this report.

Copyright © 2018 Inspur.

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

Printed in the United States, October 2018

Inspur is registered trademarks of Inspur Group.

Microsoft and SQL Server 2017 Enterprise Edition are registered trademarks of Microsoft Corporation.

TPC Benchmark™, TPC-H, QppH, QthH and QphH are registered certification marks of the Transaction Processing Performance Council.

All other brand or product names mentioned herein must be considered trademarks or registered trademarks of their respective owners.

Abstract

Overview

This report documents the methodology and results of the TPC Benchmark™ H test conducted on the Inspur NF8380M5 using Microsoft SQL Server 2017 Enterprise Edition in conformance with the requirements of the TPC Benchmark™ H Standard Specification, Revision 2.17.3. The operating system used for the benchmark was Microsoft Windows Server 2016 Standard Edition for x64.

The TPC Benchmark™ H was developed by the Transaction Processing Performance Council (TPC). The TPC was founded to define transaction processing benchmarks and to disseminate objective, verifiable performance data to the industry.

TPC Benchmark™ H Full Disclosure Report and other information can be downloaded from the Transaction Processing Performance Council web site at www.tpc.org.

Standard and Executive Summary Statements

Pages 4 - 7 contain the Executive Summary and Numerical Quantities Summary of the benchmark results for the Inspur NF8380M5.

Auditor

The benchmark configuration, environment and methodology used to produce and validate the test results, and the pricing model used to calculate the cost per QppH and QthH were audited by Francois Raab of InfoSizing to verify compliance with the relevant TPC specifications.

The auditor's letter of attestation is attached in Section 9.1 "Auditors' Report."



Inspur NF8380M5

TPC-H Rev. 2.17.3
 TPC Pricing Rev. 2.3.0

Report Date:
 October 1, 2018

Total System Cost

Composite Query per Hour Rating

Price/Performance

\$825,114

1,100,113.5
 QpH @ 10,000GB

\$0.76 USD
 \$ / QpH @10,000GB

Database Size

Database Manager

Operating System

Other Software

Availability Date

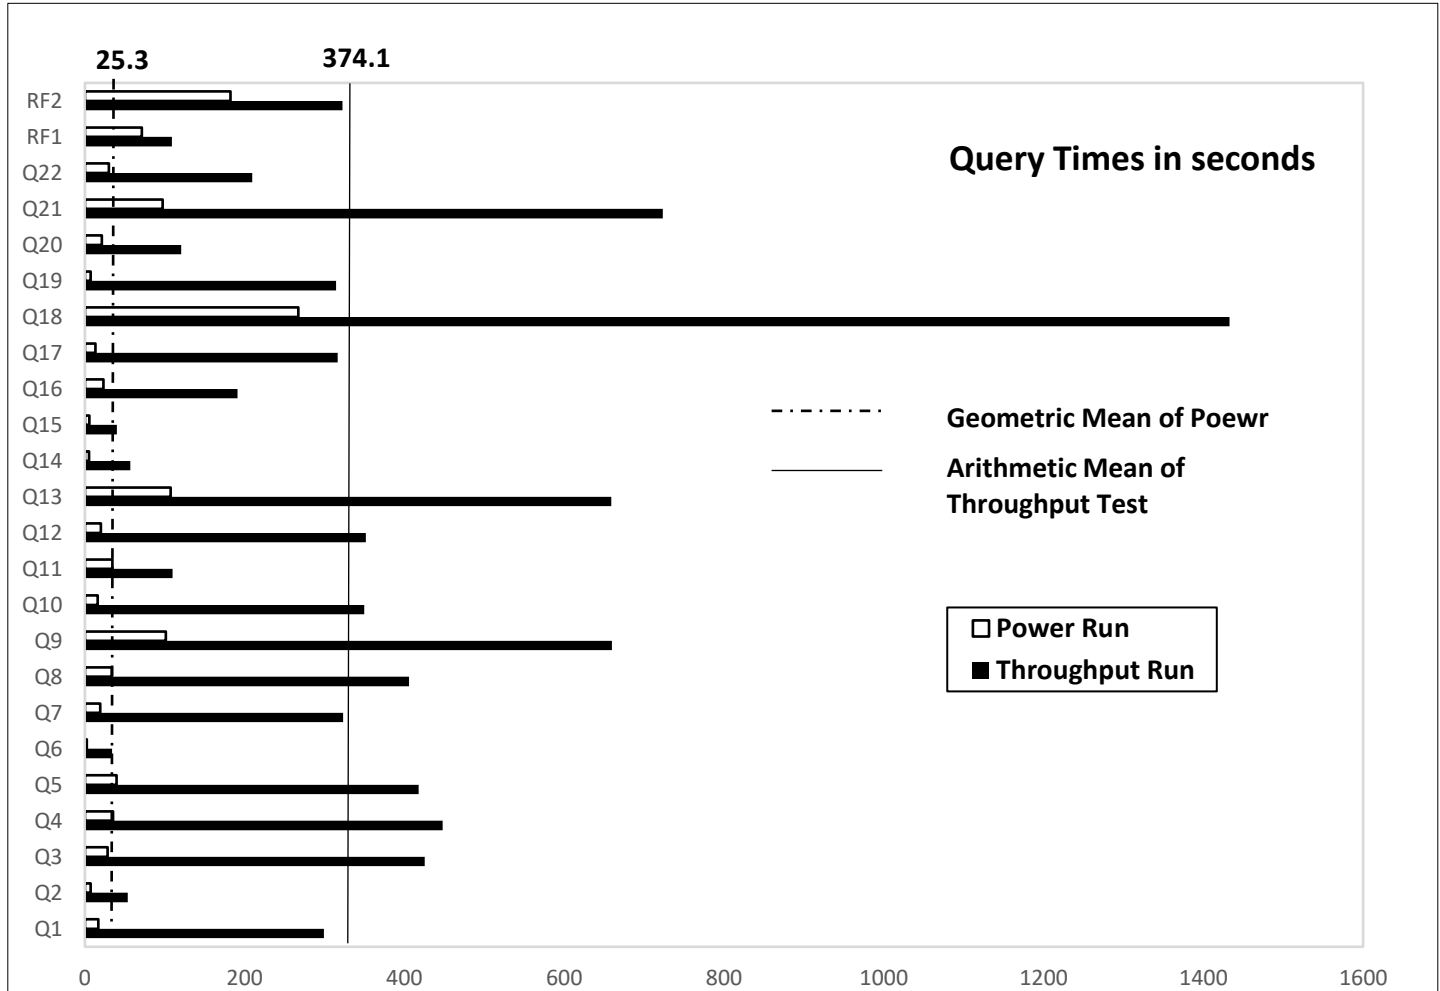
10000GB

Microsoft®
 SQL Server 2017
 Enterprise Edition

Microsoft®
 Windows Server 2016
 Standard Edition

None

December 2018



Database Load Time = 00d 05h 58m 13s

Storage Redundancy Level

Load included backup: **Y**

Base Tables and Auxiliary Data Structures

No RAID

Total Data Storage / Database Size = **3.71**

DBMS Temporary Space

No RAID

Memory/Database Size = **30.7%**

OS and DBMS Software

No RAID

System Components

System Total

Per Node

Nodes:	1	n.a.
Intel(R) Xeon(R) Platinum 8168 @ 2.70GHz:	4	4
Cores/Threads:	96/192	96/192
Memory:	3,072GB	3,072GB
Network Interfaces (40G NIC):	1	1
Storage PCIe4 NVMe 960GB M.2 SSD:	39	39
Total Storage	37,073 GB	n.a.



Inspur NF8380M5

TPC-H Rev. 2.17.3
 TPC Pricing Rev. 2.3.0

Report Date:
 October 1, 2018

Description	Price Key	Part Number	Unit Price	Qty	Extended Price	3 Yr Maint Price
Server Hardware						
Inspur NF8380M5 Server (Customer to Order), Includes:	1	N/A	\$248,227.00	1		
Intel® Xeon® Platinum 8168 (2.70GHz/24-core)	1	V001033000000000	Included	4		
SK Hynix 64G DDR4-2666LR	1	V00402L000000000	Included	48		
960G PCIex4 NVMe M.2 SSD (For OS)	1	V0230SM000000000	Included	1		
Inspur Standard 3 Year Warranty	1	N/A	Included	1		
Monitor, Keyboard And Mouse Bundle	1	N/A	Included	1		
42U Standard 19 inch Rack	1	N/A	Included	1		
Server Storage						
Inspur HHHL M.2 riser card	1	YZRI-01201-101	Included	9		
960G PCIex4 NVMe M.2 SSD	1	V0230SM000000000	Included	38		
Subtotal					\$248,227.00	
Server Software						
SQL Server 2017 Enterprise Edition for Windows (2 core license, 24 cores)	2	N/A	\$13,472.50	48	\$646,680.00	Included
Windows Server 2016 Standard Edition	2	N/A	\$92.00	48	\$4,416.00	Included
Microsoft Problem Resolution Services	2	N/A	\$259.00	1		\$259.00
Subtotal					\$651,096.00	
Total Extended Price					\$899,582.00	
Total Discounts					\$74,468.10 Included	
Price Key: 1 - Inspur , 2 - Microsoft Audited by Francois Raab of InfoSizing (www.sizing.com)		Grand Total			\$825,113.90	
All discounts are based on US list prices and for similar quantities and configurations. A 30% discount was based on the overall specific components pricing from vendor 1 in this single quotation. Discounts for similarly sized configurations will be similar to those quoted here but may vary based on the components in the configuration.		3 year cost of ownership USD:				\$825,114.00
		QpH @ 10000GB:				1,100,113.5
		\$ USD/QpH @ 10000GB				\$0.76
Sales contact: Inspur Systems Inc, 47451 Fremont Blvd, Fremont, CA 94538 (800) 697-5893 Microsoft Corporation, One Microsoft Way, Redmond, WA 98052-6399 (425) 882 8080						
Prices used in TPC benchmarks reflect the actual prices a customer would pay for a one-time purchase of the stated components. Individually negotiated discounts are not permitted. Special prices based on assumptions about past or future purchases are not permitted. All discounts reflect standard pricing policies for the listed components. For complete details, see the pricing sections of the TPC benchmark specifications. If you find that the stated prices are not available according to these terms, please inform at pricing@tpc.org. Thank you.						



Inspur NF8380M5

TPC-H Rev. 2.17.3
TPC Pricing Rev. 2.3.0

Report Date:
October 1, 2018

Numerical Quantities

Measurement Results

Database Scale Factor	10,000
Total Data Storage / Database Size	3.71
Percentage Memory / Database Size	30.7%
Start of Database Load	2018-9-26 14:04:55
End of Database Load	2018-9-26 20:03:08
Database Load Time	00d 05h 58m 13s
Query Streams for Throughput Test	9
TPC-H Power	1,423,843
TPC-H Throughput	849,988.1
TPC-H Composite Query-per-Hour (QphH@10000GB)	1,100,113.5
Total System Price over 3 Years (\$ USD)	\$825,114
TPC-H Price/Performance Metric (\$ USD / QphH@10000GB)	\$0.76

Measurement Interval

Measurement Interval in Throughput Test	8,386 seconds
---	---------------

Duration of Stream Execution

Power Run	Seed	Query Start Time	Duration (sec)	RF1 Start Time	RF2 Start Time
		Query End Time		RF1 End Time	RF2 End Time
	926200308	2018-09-26 23:13:37	00 00:16:41	2018-09-26 23:12:28	2018-09-26 23:30:20
		2018-09-26 23:30:18	1,001	2018-09-26 23:13:35	2018-09-26 23:33:38
Throughput Stream	Seed	Query Start Time	Duration (sec)	RF1 Start Time	RF2 Start Time
		Query End Time		RF1 End Time	RF2 End Time
1	926200309	2018-09-26 23:33:37	00 02:12:02	2018-09-26 23:33:37	2018-09-26 23:35:16
		2018-09-27 01:45:39	7,922	2018-09-26 23:35:16	2018-09-26 23:40:45
2	926200310	2018-09-26 23:33:37	00 02:19:40	2018-09-26 23:40:46	2018-09-26 23:42:31
		2018-09-27 01:53:17	8,380	2018-09-26 23:42:31	2018-09-26 23:46:32
3	926200311	2018-09-26 23:33:37	00 02:17:59	2018-09-26 23:46:33	2018-09-26 23:49:22
		2018-09-27 01:51:36	8,279	2018-09-26 23:49:22	2018-09-26 23:53:36
4	926200312	2018-09-26 23:33:37	00 02:19:46	2018-09-26 23:53:36	2018-09-26 23:56:12
		2018-09-27 01:53:23	8,386	2018-09-26 23:56:12	2018-09-27 00:00:57
5	926200313	2018-09-26 23:33:37	00 02:16:55	2018-09-27 00:00:57	2018-09-27 00:02:39
		2018-09-27 01:50:32	8,215	2018-09-27 00:02:39	2018-09-27 00:07:30
6	926200314	2018-09-26 23:33:37	00 02:14:26	2018-09-27 00:07:31	2018-09-27 00:10:48
		2018-09-27 01:48:03	8,066	2018-09-27 00:10:48	2018-09-27 00:15:42
7	926200315	2018-09-26 23:33:37	00 02:19:02	2018-09-27 00:15:42	2018-09-27 00:17:27
		2018-09-27 01:52:39	8,342	2018-09-27 00:17:27	2018-09-27 00:24:27
8	926200316	2018-09-26 23:33:37	00 02:19:04	2018-09-27 00:24:28	2018-09-27 00:26:14
		2018-09-27 01:52:41	8,344	2018-09-27 00:26:14	2018-09-27 00:32:33
9	926200317	2018-09-26 23:33:37	00 02:15:53	2018-09-27 00:32:33	2018-09-27 00:34:01
		2018-09-27 01:49:30	8,153	2018-09-27 00:34:01	2018-09-27 00:39:02



Inspur NF8380M5

TPC-H Rev. 2.17.3
TPC Pricing Rev. 2.3.0

Report Date:
October 1, 2018

TPC-H Timing Intervals (in seconds):

Query	Stream 0	Stream 1	Stream 2	Stream 3	Stream 4	Stream 5	Stream 6	Stream 7	Stream 8	Stream 9	Min Qi	Max Qi	Avg. Qi
Q1	15.8	324.5	88.3	453.9	273.2	177.0	156.2	93.1	221.2	343.7	15.8	214.7	453.9
Q2	11.3	44.8	55.4	47.2	59.8	42.9	70.4	49.3	67.1	58.8	11.3	50.7	70.4
Q3	27.7	519.0	625.3	431.9	401.2	668.8	592.1	586.9	188.0	232.7	27.7	427.4	668.8
Q4	34.1	232.3	453.0	530.0	547.8	529.6	359.8	478.6	739.7	491.7	34.1	439.7	739.7
Q5	37.1	488.7	500.5	547.9	240.1	638.7	512.4	404.6	387.1	262.4	37.1	402.0	638.7
Q6	1.0	54.9	5.9	53.7	38.9	42.9	23.1	48.6	41.8	30.1	1.0	34.1	54.9
Q7	17.8	258.0	130.3	380.1	248.1	380.8	329.6	522.5	424.5	471.5	17.8	316.3	522.5
Q8	29.9	549.0	401.4	207.4	216.9	419.5	321.0	351.3	551.5	436.7	29.9	348.5	551.5
Q9	208.0	633.4	556.1	677.1	870.1	479.7	568.2	988.1	654.3	744.3	208.0	637.9	988.1
Q10	16.7	249.3	431.5	499.6	682.0	351.2	109.3	246.5	556.1	666.0	16.7	380.8	682.0
Q11	34.1	161.5	86.8	96.0	149.5	192.5	186.3	97.0	68.3	129.9	34.1	120.2	192.5
Q12	11.5	222.5	807.4	224.8	426.9	362.9	329.2	132.8	469.5	408.2	11.5	339.6	807.4
Q13	104.4	734.8	677.1	806.0	667.0	450.5	586.4	718.9	544.8	478.1	104.4	576.8	806.0
Q14	6.0	23.2	18.3	24.9	103.0	28.1	68.2	19.7	68.6	98.6	6.0	45.9	103.0
Q15	5.9	41.2	57.1	48.6	60.8	41.7	44.1	28.8	13.0	28.6	5.9	37.0	60.8
Q16	23.4	166.1	370.5	355.8	193.7	279.9	235.6	153.9	159.5	163.3	23.4	210.2	370.5
Q17	10.2	540.8	235.8	138.4	412.9	399.4	427.8	706.9	491.9	356.5	10.2	372.1	706.9
Q18	251.0	1,412.4	1,511.9	1,624.0	1,285.5	1,493.8	1,754.7	1,548.0	1,539.0	1,375.2	251.0	1,379.6	1,754.7
Q19	5.5	323.9	343.0	140.5	206.8	461.2	327.8	199.8	11.7	113.0	5.5	213.3	461.2
Q20	21.0	109.5	150.5	107.9	36.7	80.7	132.1	124.6	186.3	146.5	21.0	109.6	186.3
Q21	95.6	555.9	575.1	613.6	1,101.1	555.6	683.2	660.7	775.9	919.0	95.6	653.6	1,101.1
Q22	31.6	275.7	298.0	269.2	163.6	136.3	248.0	180.8	183.3	197.2	31.6	198.4	298.0
RF1	67.3	98.7	105.8	169.7	155.6	102.2	197.2	104.3	106.2	87.8	67.3	119.5	197.2
RF2	197.5	328.9	240.9	253.3	285.0	291.1	294.2	420.6	378.8	301.5	197.5	299.2	420.6

Abstract.....	3
Overview.....	3
Standard and Executive Summary Statements.....	3
Auditor.....	3
0 General Items.....	10
0.1 Test Sponsor.....	10
0.2 Parameter Settings.....	10
0.3 Configuration Items.....	10
1.0 Clause 1: Logical Database Design.....	12
1.1 Table Definitions.....	12
1.2 Physical Organization of Database.....	12
1.3 Horizontal Partitioning.....	12
1.4 Replication.....	12
2.0 Clause 2: Queries and Refresh Functions - Related Items.....	13
2.1 Query Language.....	13
2.2 Random Number Generation.....	13
2.3 Substitution Parameters Generation.....	13
2.4 Query Text and Output Data from Database.....	13
2.5 Query Substitution Parameters and Seeds Used.....	13
2.6 Isolation Level.....	13
2.7 Refresh Functions.....	14
3.0 Clause 3: Database System Properties.....	15
3.1 Atomicity Requirements.....	15
3.1.1 Atomicity of the Completed Transactions.....	15
3.1.2 Atomicity of Aborted Transactions.....	15
3.2 Consistency Requirements.....	15
3.2.1 Consistency Tests.....	15
3.3 Isolation Requirements.....	17
3.4 Durability Requirements.....	19
4.0 Clause 4: Scaling and Database Population.....	20
4.1 Initial Cardinality of Tables.....	20
4.2 Distribution of Tables and Logs Across Media.....	21
4.3 Mapping of Database Partitions/Replications.....	22
4.4 Implementation of RAID.....	22
4.5 DBGEN Modifications.....	22
4.6 Database Load time.....	22
4.7 Data Storage Ratio.....	23
4.8 Database Load Mechanism Details and Illustration.....	23
4.9 Qualification Database Configuration.....	24
4.10 Memory to Database Size Percentage.....	24
5.0 Clause 5: Performance Metrics and Execution Rules Related Items.....	25
5.1 Steps after the Load Test.....	25
5.2 Steps in the Power Test.....	25
5.3 Timing Intervals for Each Query and Refresh Function.....	25
5.4 Number of Streams for The Throughput Test.....	25
5.5 Start and End Date/Times for Each Query Stream.....	25
5.6 Total Elapsed Time for the Measurement Interval.....	25
5.7 Refresh Function Start Date/Time and Finish Date/Time.....	26
5.8 Timing Intervals for Each Query and Each Refresh Function for Each Stream.....	26
5.9 Performance Metrics.....	26
5.10 The Performance Metric and Numerical Quantities from Both Runs.....	26
5.11 System Activity Between Tests.....	26

5.12 Documentation to satisfy Clause 5.2.7	27
5.13 Query Validation Output.....	27
6.0 Clause 6: SUT and Driver Implementation Related Items.....	28
6.1 Driver	28
6.2 Implementation Specific Layer (ISL).....	28
6.3 Profile-Directed Optimization	28
7.0 Clause 7: Pricing Related Items.....	29
7.1 Hardware and Software Used	29
7.2 Three-Year Cost of System Configuration	29
7.3 Availability Dates.....	29
8.0 Clause 8: Support Files Index Table.....	30
8.1 Supporting Files Index Table	30
9.0 Clause 9: Audit Related Items	31
9.1 Auditors' Report.....	31
Appendix A: Price Quotes	34

0 General Items

0.1 Test Sponsor

A statement identifying the benchmark sponsor(s) and other participating companies must be provided.

This benchmark was sponsored by Inspur. The benchmark was developed and engineered by Inspur. Testing took place at Inspur facilities in Jinan, China.

0.2 Parameter Settings

Settings must be provided for all customer-tunable parameters and options which have been changed from the defaults found in actual products, including but not limited to:

- *Database Tuning Options*
- *Optimizer/Query execution options*
- *Query processing tool/language configuration parameters*
- *Recovery/commit options*
- *Consistency/locking options*
- *Operating system and configuration parameters*
- *Configuration parameters and options for any other software component incorporated into the pricing structure*
- *Compiler optimization options*

This requirement can be satisfied by providing a full list of all parameters and options, as long as all those which have been modified from their default values have been clearly identified and these parameters and options are only set once.

The supporting files archive contains a list of all database parameters and operating system parameters.

0.3 Configuration Items

Diagrams of both measured and priced configurations must be provided, accompanied by a description of the differences. This includes, but is not limited to:

- *Number and type of processors*
- *Size of allocated memory, and any specific mapping/partitioning of memory unique to the test.*
- *Number and type of disk units (and controllers, if applicable).*
- *Number of channels or bus connections to disk units, including their protocol type.*
- *Number of LAN (e.g. Ethernet) Connections, including routers, workstations, terminals, etc., that were physically used in the test or are incorporated into the pricing structure.*
- *Type and the run-time execution location of software components (e.g., DBMS, query processing tools /languages, middle-ware components, software drivers, etc.).*

The System Under Test (SUT), an Inspur NF8380M5, depicted in Figure 0.1, consisted of:

Inspur NF8380M5

4 Intel(R) Xeon(R) Platinum 8168 (2.70GHz/24-core)

3,072 GB RAM (64GBx48) (LRDIMM, DDR4)

Storage

39 960GB PCIe4 NVMe M.2 SSD (internal)

Note: This system is the same system used for both the measured and priced configurations.




Figure 0.1 Benchmark and priced configuration for Inspur NF8380M5

1.0 Clause 1: Logical Database Design

1.1 Table Definitions

Listings must be provided for all table definition statements and all other statements used to set up the test and qualification databases.

The Supporting Files Archive contains the table definitions and the program used to load the database.

1.2 Physical Organization of Database

The physical organization of tables and indices, within the test and qualification databases, must be disclosed. If the column ordering of any table is different from that specified in Clause 1.4, it must be noted.

No column reordering was used.

The Supporting Files Archive contains the DDL for the index definitions.

1.3 Horizontal Partitioning

Horizontal partitioning of tables and rows in the test and qualification databases (see Clause 1.5.4) must be disclosed.

Horizontal partitioning is used on LINEITEM and ORDERS tables and the partitioning columns are L_SHIPDATE and O_ORDERDATE. The partition granularity is by week.

1.4 Replication

Any replication of physical objects must be disclosed and must conform to the requirements of Clause 1.5.6.

No replication was used.

2.0 Clause 2: Queries and Refresh Functions - Related Items

2.1 Query Language

The query language used to implement the queries must be identified.

T-SQL was the query language used.

2.2 Random Number Generation

The method of verification for the random number generation must be described unless the supplied DBGEN and QGEN were used.

DBGEN version 2.17.3 and QGEN version 2.17.3 were used to generate all database populations.

2.3 Substitution Parameters Generation

The method used to generate values for substitution parameters must be disclosed. If QGEN is not used for this purpose, then the source code of any non-commercial tool used must be disclosed. If QGEN is used, the version number, release number, modification number and patch level of QGEN must be disclosed.

The TPC source based QGEN version 2.17.3 was used to generate the substitution parameters

2.4 Query Text and Output Data from Database

The executable query text used for query validation must be disclosed along with the corresponding output data generated during the execution of the query text against the qualification database. If minor modifications (see Clause 2.2.3) have been applied to any functional query definitions or approved variants in order to obtain executable query text, these modifications must be disclosed and justified. The justification for a particular minor query modification can apply collectively to all queries for which it has been used. The output data for the power and throughput tests must be made available electronically upon request..

The Supporting Files Archive contains the query text and query output. The following modifications were used:

- The “dateadd” function is used to perform date arithmetic in Q1, Q4, Q5, Q6, Q10, Q12, Q14 ,Q15 and Q20.
- The “datepart” function is used to extract part of a date (“YY”) in Q7, Q8 and Q9.
- The “top” function is used to restrict the number of output rows in Q2, Q3, Q10, Q18 and Q21
- The “count_big” function is used in place of “count” in Q1

2.5 Query Substitution Parameters and Seeds Used

All the query substitution parameters used during the performance test must be disclosed in tabular format, along with the seeds used to generate these parameters.

The Supporting Files Archive contains the seed and query substitution parameters used.

2.6 Isolation Level

The isolation level used to run the queries must be disclosed. If the isolation level does not map closely to one of the isolation levels defined in Clause 3.4, additional descriptive detail must be provided.

The queries and transactions were run with isolation level Read Committed.

2.7 Refresh Functions

The details of how the refresh functions were implemented must be disclosed

The Supporting Files Archive contains the source code for the refresh functions.

3.0 Clause 3: Database System Properties

3.1 Atomicity Requirements

The results of the ACID tests must be disclosed along with a description of how the ACID requirements were met. This includes disclosing the code written to implement the ACID Transaction and Query.

All ACID tests were conducted according to specification. The steps performed are outlined below.

3.1.1 Atomicity of the Completed Transactions

Perform the ACID Transaction for a randomly selected set of input data and verify that the appropriate rows have been changed in the ORDER, LINEITEM, and HISTORY tables.

The following steps were performed to verify the Atomicity of completed transactions:

1. The total price from the ORDER table and the extended price from the LINEITEM table were retrieved for a randomly selected order key.
2. The ACID Transaction was performed using the order key from step 1.
3. The ACID Transaction committed.
4. The total price from the ORDER table and the extended price from the LINEITEM table were retrieved for the same order key. It was verified that the appropriate rows had been changed.

3.1.2 Atomicity of Aborted Transactions

Perform the ACID transaction for a randomly selected set of input data, submitting a ROLLBACK of the transaction for the COMMIT of the transaction. Verify that the appropriate rows have not been changed in the ORDER, LINEITEM, and HISTORY tables.

The following steps were performed to verify the Atomicity of the aborted ACID transaction:

1. The total price from the ORDER table and the extended price from the LINEITEM table were retrieved for a randomly selected order key.
2. The ACID Transaction was performed using the order key from step 1. The transaction was stopped prior to the commit.
3. The ACID Transaction was ROLLED BACK. .
4. The total price from the ORDER table and the extended price from the LINEITEM table were retrieved for the same order key used in steps 1 and 2. It was verified that the appropriate rows had not been changed.

3.2 Consistency Requirements

Consistency is the property of the application that requires any execution of transactions to take the database from one consistent state to another. A consistent state for the TPC-H database is defined to exist when:

$$O_TOTALPRICE = SUM(trunc(trunc((L_EXTENDEDPRICE - L_DISCOUNT) * (1 + L_TAX)))$$

for each ORDER and LINEITEM defined by (O_ORDERKEY = L_ORDERKEY)

3.2.1 Consistency Tests

Verify that ORDER and LINEITEM tables are initially consistent as defined in Clause 3.3.2.1, based upon a random sample of at least 10 distinct values of O_ORDERKEY.

Consistency was tested as part of the durability tests.

3.3 Isolation Requirements

Operations of concurrent transactions must yield results which are indistinguishable from the results which would be obtained by forcing each transaction to be serially executed to completion in some order.

3.3.1 Isolation Test 1 - Read-Write Conflict with Commit

Demonstrate isolation for the read-write conflict of a read-write transaction and a read-only transaction when the read-write transaction is committed)

The following steps were performed to satisfy the test of isolation for a read-only and a read-write committed transaction:

1. An ACID Transaction was started for a randomly selected O_KEY, L_KEY and DELTA. The ACID Transaction was suspended prior to Commit.
2. An ACID query was started for the same O_KEY used in step 1. The ACID query blocked and did not see any uncommitted changes made by the ACID Transaction.
3. The ACID Transaction was resumed and committed. The ACID query completed. It returned the data as committed by the ACID Transaction.

3.3.2 Isolation Test 2 - Read-Write Conflict with Rollback

Demonstrate isolation for the read-write conflict of a read-write transaction and a read-only transaction when the read-write transaction is rolled back.

The following steps were performed to satisfy the test of isolation for read-only and a rolled back read-write transaction:

1. An ACID transaction was started for a randomly selected O_KEY, L_KEY and DELTA. The ACID Transaction was suspended prior to Rollback.
2. An ACID query was started for the same O_KEY used in step 1. The ACID query did not see any uncommitted changes made by the ACID Transaction.
3. The ACID Transaction was ROLLED BACK.
4. The ACID query completed.

3.3.3 Isolation Test 3 - Write-Write Conflict with Commit

Demonstrate isolation for the write-write conflict of two update transactions when the first transaction is committed.

The following steps were performed to verify isolation of two update transactions:

1. An ACID Transaction T1 was started for a randomly selected O_KEY, L_KEY and DELTA. The ACID transaction T1 was suspended prior to Commit.
2. Another ACID Transaction T2 was started using the same O_KEY and L_KEY and a randomly selected DELTA.
3. T2 waited.
4. The ACID transaction T1 was allowed to Commit and T2 completed.
5. It was verified that: $T2.L_EXTENDEDPRICE = T1.L_EXTENDEDPRICE + (DELTA1 * (T1.L_EXTENDEDPRICE / T1.L_QUANTITY))$

3.3.4 Isolation Test 4 - Write-Write Conflict with Rollback

Demonstrate isolation for the write-write conflict of two update transactions when the first transaction is rolled back.

The following steps were performed to verify the isolation of two update transactions after the first one is rolled back:

1. An ACID Transaction T1 was started for a randomly selected O_KEY, L_KEY and DELTA. The ACID Transaction T1 was suspended prior to Rollback.
2. Another ACID Transaction T2 was started using the same O_KEY and L_KEY used in step 1 and a randomly selected DELTA.
3. T2 waited.

4. T1 was allowed to ROLLBACK and T2 completed.
5. It was verified that T2.L_EXTENDEDPRICE = T1.L_EXTENDEDPRICE

3.3.5 Isolation Test 5 – Concurrent Read and Write Transactions on Different Tables

Demonstrate the ability of read and write transactions affecting different database tables to make progress concurrently.

The following steps were performed:

1. An ACID Transaction T1 for a randomly selected O_KEY, L_KEY and DELTA. The ACID Transaction T1 was suspended prior to Commit.
2. Another ACID Transaction T2 was started using random values for PS_PARTKEY and PS_SUPPKEY.
3. T2 completed.
4. T1 completed and the appropriate rows in the ORDER, LINEITEM and HISTORY tables were changed.

3.3.6 Isolation Test 6 – Update Transactions During Continuous Read-Only Query Stream

Demonstrate the continuous submission of arbitrary (read-only) queries against one or more tables of the database does not indefinitely delay update transactions affecting those tables from making progress.

The following steps were performed:

1. An ACID Transaction T1 was started, executing a modified Q1 against the qualification database. The substitution parameter was chosen from the interval [0..2159] so that the query ran for a sufficient amount of time.
2. Before T1 completed, an ACID Transaction T2 was started using randomly selected values of O_KEY, L_KEY and DELTA.
3. T2 completed before T1 completed.
4. It was verified that the appropriate rows in the ORDER, LINEITEM and HISTORY tables were changed.

3.4 Durability Requirements

The tested system must guarantee durability: the ability to preserve the effects of committed transactions and insure database consistency after recovery from any one of the failures listed in Clause 3.5.2.

3.4.1 Permanent Unrecoverable Failure of Any Durable Medium and Loss of System Power

Guarantee the database and committed updates are preserved across a permanent irrecoverable failure of any single durable medium containing TPC-H database tables or recovery log tables

Three tests were completed in this section.

1st test was the removal of a log disk

2nd test was the removal of a data disk

3rd test was the Power Loss test.

Each of these tests were performed against the qualification database. The qualification database is identical to the test database in virtually every regard except size.

Log Disk Removal Test.

1. The complete database was backed up.
2. Ten streams of ACID transactions were started. Each stream executed a minimum of 100 transactions.
3. While the test was running, one of the disks from the database log RAID 5 volume was removed.
4. A system bluescreen occurred, followed by an automatic reboot. The system was then manually shutdown.
5. The pulled disk was replaced with a new disk. Log disk eventually completed its RAID rebuild process without any issue.
6. The counts in the history table and success files were compared and verified, and the consistency of the ORDERS and LINEITEM tables was verified

Data Disk Removal Test

1. The complete database was backed up.
2. Ten streams of ACID transactions were started. Each stream executed a minimum of 100 transactions.
3. While the test was running, one of the disks from the data volume was removed.
4. A system bluescreen occurred, followed by an automatic reboot
5. The Database log was backed up and the Database was dropped.
6. The system was shutdown manually. The pulled disk was replaced with a new one and the system was restarted.
7. The database was restored.
8. When database restore completed, issued a command to apply the backed up log file.
9. The counts in the history table and success files were compared and verified, and the consistency of the ORDERS and LINEITEM tables was verified.

Power Loss test.

1. Ten streams of ACID transactions were started. Each stream executed a minimum of 100 transactions.
2. While the streams of ACID transactions were running the System was powered off by pulling power plugs.
3. When power was restored the system booted and the Database engine was restarted.
4. The database went through a recovery period.
5. Rolled forward and backward transactions were captured by DB ERRORLOG file.
6. The database recovery completed.
7. The counts in the history table and success files were compared and verified, and the consistency of the ORDERS and LINEITEM tables was verified.

3.4.2 System Crash

Guarantee the database and committed updates are preserved across an instantaneous interruption (system crash/system hang) in processing which requires the system to reboot to recover.

See section 3.4.1

3.4.3 Memory Failure

Guarantee the database and committed updates are preserved across failure of all or part of memory (loss of contents).

See section 3.4.1

4.0 Clause 4: Scaling and Database Population

4.1 Initial Cardinality of Tables

The cardinality (i.e., the number of rows) of each table of the test database, as it existed at the completion of the database load (see clause 4.2.5) must be disclosed.

Table 4.1 lists the TPC-H Benchmark defined tables and the row count for each table as they existed upon completion of the build.

Table	# of Rows
Lineitem	59,999,994,267
Orders	15,000,000,000
Partsupp	8,000,000,000
Part	2,000,000,000
Customer	1,500,000,000
Supplier	100,000,000
Nation	25
Region	5

Table 4.1 Initial Number of Rows

4.2 Distribution of Tables and Logs Across Media

The distribution of tables and logs across all media must be explicitly described for the tested and priced systems.

Microsoft SQL Server was configured on an Inspur NF8380M5 with the following storage configuration:

39 x 960GB PCIe4 NVMe M.2 SSD

The disks were distributed as follows:

- 1 960GB NVMe SSDs was used for the OS.
- 6 960GB NVMe SSDs were used for database LOG
- 32 960GB NVMe SSDs were used for DATA and TempDB

A description of distribution of database file groups and log can be found in the Table below.

Slot Controller/Channel	# of Disks	Array Fault Tolerance	Size in GB	Partition Format	Content
Onboard NVMe SSD	1	Non RAID	960	NTFS	OS, SQL Server
Onboard&PCIe Slot 9 NVMe SSD	6	RAID 5	4470.63	NTFS	DB Log
PCIe Slot 1 NVMe SSD	4	Non RAID	894.13	NTFS	DB & TempDB
PCIe Slot 2 NVMe SSD	4	Non RAID	894.13	NTFS	DB & TempDB
PCIe Slot 3 NVMe SSD	4	Non RAID	894.13	NTFS	DB & TempDB
PCIe Slot 4 NVMe SSD	4	Non RAID	894.13	NTFS	DB & TempDB
PCIe Slot 5 NVMe SSD	4	Non RAID	894.13	NTFS	DB & TempDB
PCIe Slot 6 NVMe SSD	4	Non RAID	894.13	NTFS	DB & TempDB
PCIe Slot 7 NVMe SSD	4	Non RAID	894.13	NTFS	DB & TempDB
PCIe Slot 8 NVMe SSD	4	Non RAID	894.13	NTFS	DB & TempDB
			Used		
Total Space	39		34,042.8 GB		

4.3 Mapping of Database Partitions/Replications

The mapping of database partitions/replications must be explicitly described.

Horizontal partitioning is used on LINEITEM and ORDERS tables and the partitioning columns are L_SHIPDATE and O_ORDERDATE. The partition granularity is by week.

The database partitions are evenly distributed across 32 NVMe SSDs.

4.4 Implementation of RAID

Implementations may use some form of RAID to ensure high availability. If used for data, auxiliary storage (e.g. indexes) or temporary space, the level of RAID used must be disclosed for each device.

No RAID was used for the database file groups and the TempDB files. RAID 5 was used for the database recovery logs and tempDB recovery logs. No RAID was used for the Backup drives..

4.5 DBGEN Modifications

The version number, release number, modification number, and patch level of DBGEN must be disclosed. Any modifications to the DBGEN (see Clause 4.2.1) source code must be disclosed. In the event that a program other than DBGEN was used to populate the database, it must be disclosed in its entirety.

DBGEN version 2.17.3 was used, no modifications were made.

4.6 Database Load time

The database load time for the test database (see clause 4.3) must be disclosed.

The database load time was 00d 05h 58m 13s

Database Load Time	
Load Start	09/26/2018 14:04:55
Load End	09/26/2018 20:03:08
Load Time	00d 05h 58m 13s

4.7 Data Storage Ratio

The data storage ratio must be disclosed. It is computed by dividing the total data storage of the priced configuration (expressed in GB) by the size chosen for the test database as defined in 4.1.3.1. The ratio must be reported to the nearest 1/100th, rounded up.

Disk Type	# of Disks	RAW SIZE (GB)	Total (GB)
960G NVMe SSD	39	950.58G	37,073

Size of test database: 10,000GB

Data Storage Ratio: 3.71

4.8 Database Load Mechanism Details and Illustration

The details of the database load must be disclosed, including a block diagram illustrating the overall process. Disclosure of the load procedure includes all steps, scripts, input and configuration files required to completely reproduce the test and qualification databases.

Flat files for each of the tables were created using DBGEN. The tables were loaded as depicted in Figure 4.8. All steps, scripts and configuration files are included in the Supporting Files.

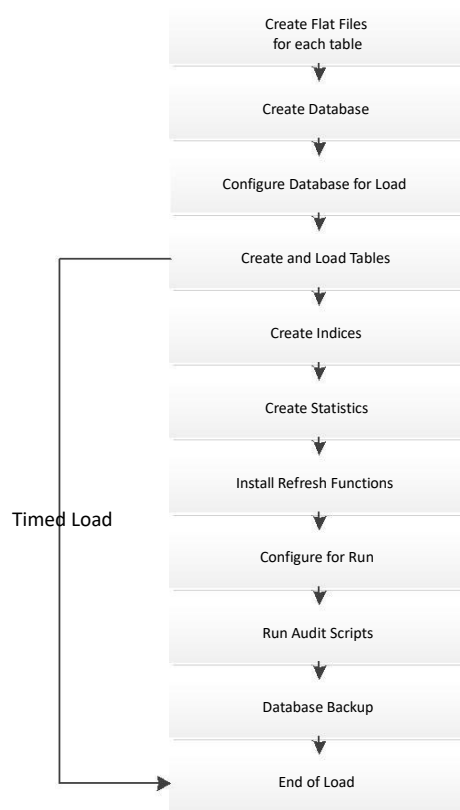


Figure 4.8: Block Diagram of Database Load Process

4.9 Qualification Database Configuration

The details of the database load must be disclosed, including a block diagram illustrating the overall process. Disclosure of the load procedure includes all steps, scripts, input and configuration files required to completely reproduce the test and qualification databases.

The qualification database used identical scripts to create and load the data with changes to adjust for the database scale factor.

4.10 Memory to Database Size Percentage

The memory to database size percentage, as defined in clause 8.3.6.10, must be disclosed.

The memory to database size percentage is 30.72%

Memory/Database Size		
Physical Mem GB	Scale Factor	Percentage
3072	10000	30.72%

5.0 Clause 5: Performance Metrics and Execution Rules Related Items

5.1 Steps after the Load Test

Any system activity on the SUT that takes place between the conclusion of the load test and the beginning of the performance test must be fully disclosed including listings of scripts or command logs.

The queries were generated using QGen with the seed collected at the end of the load test. The system sat idle, without any interruption during this time frame.

5.2 Steps in the Power Test

The details of the steps followed to implement the power test (e.g., system boot, database restart, etc.) must be disclosed.

The following steps were used to implement the power test:

1. Execute RF1 from the update stream
2. Execute Power queries from the query stream
3. Execute RF2 from the update stream

5.3 Timing Intervals for Each Query and Refresh Function

The timing intervals (see Clause 5.3.6) for each query of the measured set and for both refresh functions must be reported for the power test.

The timing intervals for each query and both refresh functions are given in the Numerical Quantities Summary earlier in this document on page vii.

5.4 Number of Streams for The Throughput Test

The number of execution streams used for the throughput test must be disclosed.

9 query streams were used for the Throughput Test. Each stream running all 22 queries. 1 stream was used for RF.

5.5 Start and End Date/Times for Each Query Stream

The start time and finish time for each query execution stream must be reported for the throughput test.

The Numerical Quantities Summary on page vii contains the start and stop times for the query execution streams run on the system reported.

5.6 Total Elapsed Time for the Measurement Interval

The total elapsed time of the measurement interval(see Clause 5.3.5) must be reported for the throughput test.

The Numerical Quantities Summary on page vii contains the timing intervals for the throughput test run on the system reported.

5.7 Refresh Function Start Date/Time and Finish Date/Time

Start and finish time for each update function in the update stream must be reported for the throughput test.

The Numerical Quantities Summary on page vii contains the start and finish times for the refresh functions of each stream.

5.8 Timing Intervals for Each Query and Each Refresh Function for Each Stream

The timing intervals (see Clause 5.3.6) for each query of each stream and for each update function must be reported for the throughput test.

The timing intervals for each query and each update function are given in the Numerical Quantities Summary earlier in this document on page vii.

5.9 Performance Metrics

The computed performance metrics, related numerical quantities and the price performance metric must be reported.

The Numerical Quantities Summary contains the performance metrics, related numerical quantities, and the price/performance metric for the system reported.

5.10 The Performance Metric and Numerical Quantities from Both Runs

The performance metric (QphH@Size) and the numerical quantities (TPC-H Power@Size and TPC-H Throughput@ Size) from both of the runs must be disclosed

Run ID	QppH@10000GB	QthH@10000GB	QphH@10000GB
Run 1	1,484,851.6	909,647.8	1,162,192.8
Run 2	1,423,843	849,988.1	1,100,113.5

5.11 System Activity Between Tests

Any activity on the SUT that takes place between the conclusion of Run1 and the beginning of Run2 must be disclosed.

There was no activity between Run1 and Run2.

5.12 Documentation to satisfy Clause 5.2.7

All documentation necessary to satisfy Clause 5.2.7 must be made available upon request

The supporting files archive contains the documentation

5.13 Query Validation Output

The output of the Query Output Validation Test must reported in the supporting files archive

The supporting files archive contains the documentation

6.0 Clause 6: SUT and Driver Implementation Related Items

6.1 Driver

A detailed description of how the driver performs its functions must be supplied, including any related source code or scripts. This description should allow an independent reconstruction of the driver.

The TPC-H benchmark was implemented using a Microsoft tool called StepMaster. StepMaster is a general purpose test tool which can drive ODBC and shell commands. Within StepMaster, the user designs a workspace corresponding to the sequence of operations (or steps) to be executed. When the workspace is executed, StepMaster records information about the run into a database as well as a log file for later analysis.

StepMaster provides a mechanism for creating parallel streams of execution. This is used in the throughput tests to drive the query and refresh streams. Each step is timed using a millisecond resolution timer. A timestamp T1 is taken before beginning the operation and a timestamp T2 is taken after completing the operation. These times are recorded in a database as well as a log file for later analysis.

Two types of ODBC connections are supported. A dynamic connection is used to execute a single operation and is closed when the operation finishes. A static connection is held open until the run completes and may be used to execute more than one step. A connection (either static or dynamic) can only have one outstanding operation at any time.

In TPC-H, static connections are used for the query streams in the power and throughput tests. StepMaster reads an Access database to determine the sequence of steps to execute. These commands are represented as the Implementation Specific Layer. StepMaster records its execution history, including all timings, in the Access database. Additionally, StepMaster writes a textual log file of execution for each run.

The stream refresh functions were executed using multiple batch scripts. The initial script is invoked by StepMaster, subsequent scripts are called from within the scripts.

The source for StepMaster and the RF Scripts is disclosed in the supported file archive.

6.2 Implementation Specific Layer (ISL)

If an implementation-specific layer is used, then a detailed description of how it performs its functions must be supplied, including any related source code or scripts. This description should allow an independent reconstruction of the implementation-specific layer.

See Section 6.1 for details.

6.3 Profile-Directed Optimization

If profile-directed optimization as described in Clause 5.2.9 is used, such use must be disclosed.

Profile-directed optimization was not used.

7.0 Clause 7: Pricing Related Items

7.1 Hardware and Software Used

A detailed list of hardware and software used in the priced system must be reported. Each item must have a vendor part number, description, and release/revision level, and indicate General Availability status or committed delivery date. If package pricing is used, contents of the package must be disclosed. Pricing source(s) and effective date(s) of price(s) must also be reported.

The pricing summary sheet is given on page v in the Executive Summary at the front of this report. The source for all prices is indicated.

Server and all storage components are available at publication date.

The pricing and availability of the Microsoft software used is given in a quote from Microsoft, which is included in this report in Appendix A.

7.2 Three-Year Cost of System Configuration

The total 3-year price of the entire configuration must be reported, including: hardware, software, and maintenance charges. Separate component pricing is required.

The pricing summary sheet on page v in the front of this report contains all details.

7.3 Availability Dates

The committed delivery date for general availability (availability date) of products used in the priced calculations must be reported. When the priced system includes products with different availability dates, the single availability date reported on the first page of the executive summary must be the date by which all components are committed to being available. The full disclosure report must report availability dates individually for at least each of the categories for which a pricing subtotal must be provided (see Clause 7.3.1.4). All availability dates, whether for individual components or for the SUT as a whole, must be disclosed to a precision of 1 day, but the precise format is left to the test sponsor.

Category	Available
Server Hardware	Dec, 2018
Storage	Dec, 2018
Server Software	Now (date of publication)
SQL Server 2017	Now (date of publication)

8.0 Clause 8: Support Files Index Table

8.1 Supporting Files Index Table

An index for all files included in the supporting files archive as required by Clauses 8.3.2 must be provided in the report.

Clause	Description	Archive File	Pathname
Clause 1	OS_DB_BIOS_Settings	benchmark_scripts.zip	SupportingFiles/Clause1
Clause 2	Queries Qualifcation and Output	benchmark_scripts.zip	SupportingFiles/Clause2
Clause 3	ACID Test scripts and Result	benchmark_scripts.zip	SupportingFiles/Clause3
Clause 4	DB Load Scripts	benchmark_scripts.zip	SupportingFiles/Clause4
Clause 5	Queries and output for measured runs	Run1Run2Results.zip	SupportingFiles/Clause5
Clause 6	Implementation Specific layer source code and Driver	benchmark_scripts.zip	SupportingFiles/Clause6
Clause 7	There are no files required to be included for Clause 7.	n/a	n/a
Clause 8	There are no files required to be included for Clause 8.	n/a	n/a

9.0 Clause 9: Audit Related Items

9.1 Auditors' Report

The auditor's agency name, address, phone number, and Attestation letter with a brief audit summary report indicating compliance must be included in the full disclosure report. A statement should be included specifying who to contact in order to obtain further information regarding the audit process.

This implementation of the TPC Benchmark™ H was audited by Francois Raab of InfoSizing, a certified TPC-H auditor. Further information regarding the audit process may be obtained from:

Francois Raab
InfoSizing (www.sizing.com)
20 Kreg Ln.
Manitou Springs, CO 80829
(719) 473-7555

TPC Benchmark™ H Full Disclosure Report and other information can be downloaded from the Transaction Processing Performance Council web site at www.tpc.org.

Benchmark sponsor: Xin Qiao
Inspur Electronic Information Industry Co., Ltd.
2F, Tower C, No.2 Xixi Rd.
Shangdi Haidian District, Beijing 100085
People's Republic of China

September 29, 2018

I verified the TPC Benchmark H (TPC-H™ v2.17.3) performance of the following configuration:

Platform: Inspur NF8380M5
Operating System: Microsoft® Windows Server 2016 Standard Edition
Database Manager: Microsoft SQL Server 2017 Enterprise Edition
Other Software: n/a

The results were:

Performance Metric 1,100,113.5 QphH@10,000GB

TPC-H Power 1,423,843.0
TPC-H Throughput 849,988.1
Database Load Time 5h 58m 13s

Server

Inspur NF8380M5

CPU	4 x Intel® Xeon® Platinum 8168 (2.70GHz, 24-core)		
Memory	3,072 GB		
Storage	Qty	Size	Type
	39	960 GB	PCIEx4 NVMe M.2 SSD

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:

- The database records were defined with the proper layout and size
- The database population was generated using DBGen
- The database was properly scaled to 10,000GB and populated accordingly
- The compliance of the database auxiliary data structures was verified
- The database load time was correctly measured and reported
- The required ACID properties were verified and met

- The query input variables were generated by QGen
- The query text was produced using minor modifications and no query variant
- The execution of the queries against the SF1 database produced compliant answers
- A compliant implementation specific layer was used to drive the tests
- The throughput tests involved 9 query streams
- The ratio between the longest and the shortest query was such that no query timings were adjusted
- The execution times for queries and refresh functions were correctly measured and reported
- The repeatability of the measured results was verified
- The system pricing was verified for major components and maintenance
- The major pages from the FDR were verified for accuracy

Additional Audit Notes:

None.

Respectfully Yours,

A handwritten signature in black ink, appearing to read "François Raab", with a long horizontal flourish extending to the right.

François Raab, TPC Certified 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
September 28, 2018

Inspur Systems, Inc.
Yinan Liu
2002 156th Ave NE #150
Bellevue, WA 980007

Here is the information you requested regarding pricing for Microsoft products to be used in conjunction with your TPC-H 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</i> <i>Open Program - Level C</i> Unit Price reflects a 6% discount from the retail unit price of \$14,256	\$13,472.50	48	\$646,680.00
Database Server Operating System			
Windows Server 2016 Standard Edition <i>2 Core License</i> <i>Open Program - Level C</i> Unit Price reflects a 17% discount from the retail unit price of \$110.25	\$92.00	48	\$4,416.00
Support			
Microsoft Problem Resolution Services <i>Professional Support</i> <i>(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: Tpch_X14u8lmpst2nsClb0CXVQGG2426_2018_iyl.doc