



Cisco Systems, Inc.

TPC Benchmark™ H
Full Disclosure Report
for
Cisco UCS C480 M5 Rack-Mount Server
using
Microsoft SQL Server 2019 Enterprise Edition
And
Red Hat Enterprise Linux 8.0

First Edition

Nov 4, 2019

First Edition – Nov 4, 2019

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company.

The Cisco products, services or features identified in this document may not yet be available or may not be available in all areas and may be subject to change without notice. Consult your local Cisco business contact for information on the products or services available in your area. You can find additional information via Cisco's World Wide Web server at <http://www.cisco.com>. Actual performance and environmental costs of Cisco products will vary depending on individual customer configurations and conditions.

Table of Contents

ABSTRACT	6
TPC-H REV. 2.18.0.....	8
PREFACE.....	12
TPC BENCHMARK™ H OVERVIEW	12
GENERAL ITEMS.....	14
0.1 TEST SPONSOR.....	14
0.2 PARAMETER SETTINGS	14
0.3 CONFIGURATION DIAGRAMS	14
CLAUSE 1: LOGICAL DATABASE DESIGN	16
1.1 DATABASE DEFINITION STATEMENTS.....	16
1.2 PHYSICAL ORGANIZATION.....	16
1.3 HORIZONTAL PARTITIONING	16
1.4 REPLICATION	16
CLAUSE 2: QUERIES AND REFRESH FUNCTIONS RELATED ITEMS	17
2.1 QUERY LANGUAGE	17
2.2 VERIFYING METHOD OF RANDOM NUMBER GENERATION	17
2.3 GENERATING VALUES FOR SUBSTITUTION PARAMETERS.....	17
2.4 QUERY TEXT AND OUTPUT DATA FROM QUALIFICATION DATABASE.....	17
2.5 QUERY SUBSTITUTION PARAMETERS AND SEEDS USED	17
2.6 ISOLATION LEVEL	18
2.7 SOURCE CODE OF REFRESH FUNCTIONS.....	18
CLAUSE 3: DATABASE SYSTEM PROPERTIES	19
3.1 ACID PROPERTIES.....	19
3.2 ATOMICITY REQUIREMENTS.....	19
3.2.1 ATOMICITY OF THE COMPLETED TRANSACTIONS	19
3.2.2 ATOMICITY OF ABORTED TRANSACTIONS.....	19
3.3 CONSISTENCY REQUIREMENTS.....	19
3.3.1 CONSISTENCY TEST	20
3.4 ISOLATION REQUIREMENTS	20
3.4.1 ISOLATION TEST 1 - READ-WRITE CONFLICT WITH COMMIT	20
3.4.2 ISOLATION TEST 2 - READ-WRITE CONFLICT WITH ROLLBACK.....	20

3.4.3 ISOLATION TEST 3 - WRITE-WRITE CONFLICT WITH COMMIT	20
3.4.4 ISOLATION TEST 4 - WRITE-WRITE CONFLICT WITH ROLLBACK	21
3.4.5 ISOLATION TEST 5 – CONCURRENT READ AND WRITE TRANSACTIONS ON DIFFERENT TABLES.....	21
3.4.6 ISOLATION TEST 6 – UPDATE TRANSACTIONS DURING CONTINUOUS READ-ONLY QUERY STREAM	21
3.5 DURABILITY REQUIREMENTS	21
3.5.1 PERMANENT UNRECOVERABLE FAILURE OF ANY DURABLE MEDIUM	21
3.5.2 LOSS OF LOG AND SYSTEM CRASH TEST	22
3.5.3 MEMORY FAILURE	22
CLAUSE 4: SCALING AND DATABASE POPULATION	23
4.1 INITIAL CARDINALITY OF TABLES.....	23
4.2 DISTRIBUTION OF TABLES AND LOGS ACROSS MEDIA	23
4.3 MAPPING OF DATABASE PARTITIONS/REPLICATIONS	25
4.4 IMPLEMENTATION OF RAID	25
4.5 DBGEN MODIFICATIONS	25
4.6 DATABASE LOAD TIME	25
4.7 DATA STORAGE RATIO	25
4.8 DATABASE LOAD MECHANISM DETAILS AND ILLUSTRATION	26
4.9 QUALIFICATION DATABASE CONFIGURATION	28
4.10 MEMORY TO DATABASE SIZE PERCENTAGE	28
CLAUSE 5: PERFORMANCE METRICS AND EXECUTION RULES RELATED ITEMS.....	29
5.1 STEPS IN THE POWER TEST	29
5.2 TIMING INTERVALS FOR EACH QUERY AND REFRESH FUNCTION.....	29
5.3 NUMBER OF STREAMS FOR THE THROUGHPUT TEST	29
5.4 START AND END DATE/TIMES FOR EACH QUERY STREAM.....	29
5.5 TOTAL ELAPSED TIME FOR THE MEASUREMENT INTERVAL	29
5.6 REFRESH FUNCTION START DATE/TIME AND FINISH DATE/TIME	29
5.7 TIMING INTERVALS FOR EACH QUERY AND EACH REFRESH FUNCTION FOR EACH STREAM.....	29
5.8 PERFORMANCE METRICS	30
5.9 THE PERFORMANCE METRIC AND NUMERICAL QUANTITIES FROM BOTH RUNS.....	30
5.10 SYSTEM ACTIVITY BETWEEN TESTS	30
CLAUSE 6: SUT AND DRIVER IMPLEMENTATION RELATED ITEMS	31
6.1 DRIVER	31
6.2 IMPLEMENTATION SPECIFIC LAYER (ISL)	31

6.3 PROFILE-DIRECTED OPTIMIZATION.....	31
CLAUSE 7: PRICING RELATED ITEMS.....	32
7.1 HARDWARE AND SOFTWARE USED	32
7.2 TOTAL 3 YEAR PRICE	32
7.3 AVAILABILITY DATE	32
7.4 ORDERABILITY DATE	32
7.5 COUNTRY-SPECIFIC PRICING.....	32
7.6 TESTED AND PRICED CONFIGURATIONS	32
CLAUSE 8: FULL DISCLOSURE	34
8.1 SUPPORTING FILE INDEX	34
CLAUSE 9: AUDIT RELATED ITEMS	35
AUDITORS' INFORMATION AND ATTESTATION LETTER.....	35
APPENDIX A: PRICE QUOTES.....	38

Abstract


This document contains the methodology and results of the TPC Benchmark™ H (TPC-H) test conducted on the Cisco UCS C480 M5 Rack-Mount Server, in conformance with the requirements of the TPC-H Standard Specification, Revision 2.18.0. The operating system used for the benchmark was Red Hat Enterprise Linux 8.0 and database software used for the benchmark was Microsoft SQL Server 2019 Enterprise Edition.

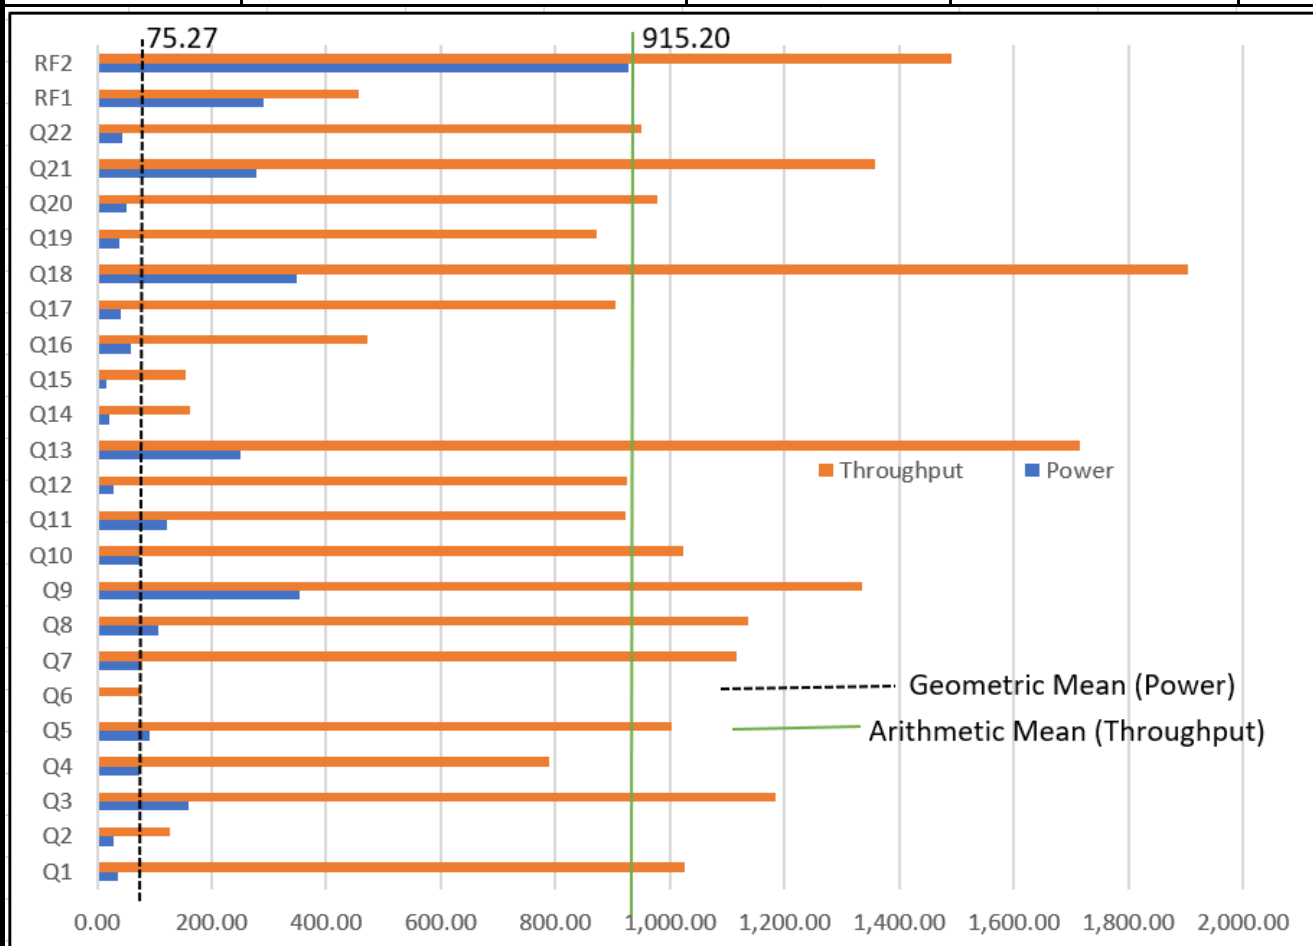
Cisco UCS C480 M5 Server

Company Name	System Name	Database Software	Operating System
Cisco Systems, Inc	Cisco UCS C480 M5 Server	Microsoft SQL Server 2019 Enterprise Edition	Red Hat Enterprise Linux 8.0

TPC Benchmark© H Metrics

Total System Cost	TPC-H Throughput	Price/Performance	Availability Date
1,197,396.28 USD	1,278,277.8 QphH@30,000GB	936.73 USD / kQphH@30,000GB	Nov 4, 2019

		Cisco UCS C480 M5 Server		TPC-H Rev. 3.0.0 TPC-Pricing Rev. 2.7.0	
				Report Date: 4-Nov-2019	
Total System Cost		Composite Query per Hour Metric		Price / Performance	
1,197,396.28 USD		1,278,277.8 QphH@30000GB		936.73 USD \$/kQphH@30000GB	
Database Size	Database Manager		Operating System	Other Software	Availability Date
30000GB	Microsoft SQL Server 2019 Enterprise Edition for Linux		Red Hat Enterprise Linux 8.0		4-Nov-2019



Database Load Time = 18h 1m 9s		Storage Redundancy Level	
Load Includes Backup: Y		Base Tables and Auxiliary Data Structures	0
Total Data Storage / Database Size = 2.81		DBMS Temporary Space	0
Percentage Memory / Database Size = 20.48%		OS and DBMS Software	1

System Configuration:	Cisco UCS C480 M5 Server
Processors/Cores/Threads/Model:	4/112/224 Intel 2 nd Gen Xeon Scalable 8280M (2.7GHz, 38.5MB cache)
Memory:	6 TB (48x 128G)
Storage:	10 x 1.9TB 2.5-inch Enterprise Value 12G SAS SSDs, 4x 3.8TB 2.5-inch Enterprise Value 6G SATA SSD, 4 x Cisco 2.5-inch U.2 7.6TB HGST SN200 NVMe High Perf. Value Endurance, 8 x Cisco HHHL AIC 3.2T HGST SN260 NVMe Extreme Perf High Endurance
Table Storage:	82.44 TB



Cisco UCS C480 M5 Server

TPC-H Rev. 3.0.0
TPC-Pricing Rev. 2.7.0

Report Date: 4-Nov-2019

Description	Part Number	Source	Qty	Unit Net Price	Extended Net Price	3 Year Maint. Price
UCS C480 M5 Std base chassis w/o CPU, mem, HDD, PCIe, PSU	UCSC-C480-M5	1	1	18,237.00	18,237.00	
UCS C480 M5 CPU Module w/o CPU, mem	UCSC-C480-CM	1	2	3,500.00	7,000.00	
Cisco UCS 1600W AC Power Supply for Rack Server	UCSC-PSU1-1600W	1	4	929.00	3,716.00	
Power Cord Jumper, C13-C14 Connectors, 2 Meter Length	CAB-C13-C14-2M	1	4	0.00	0.00	
Cable Management Arm for UCS C480 M4	UCSC-CMA-4U-M5	1	1	150.00	150.00	
IMC SW (Recommended) latest release for C-Series Servers.	CIMC-LATEST	1	1	0.00	0.00	
Other Infrastructure	UCS-SID-INFR-OI	1	1	0.00	0.00	
Rail Kit for UCS C480 M5	UCSC-RAIL-4U-M5	1	1	0.00	0.00	
Super Cap for UCSC-RAID-M5, UCSC-MRAID1GB-KIT	UCSC-SCAP-M5	1	1	0.00	0.00	
Super Cap cable for UCSC-RAID-M5HD	CBL-SC-MR12GM5P	1	1	0.00	0.00	
CPU Heat Sink for UCS C480 M5 Rack Server	UCSC-HS-02-EX	1	4	0.00	0.00	
Cisco 12G Modular RAID controller with 4GB cache	UCSC-RAID-M5HD	1	1	2,900.00	2,900.00	
Cisco HHHL AIC 3.2T HGST SN260 NVMe Extreme Perf High Endrnc	UCSC-NVME-H32003	1	8	10,391.00	83,128.00	
UCS C480 M5 Drive Module for 8x HDD	UCSC-C480-8HDD	1	1	800.00	800.00	
UCS C480 M5 Drive Module for 8x HDD	UCSC-C480-8HDD	1	1	800.00	800.00	
128GB DDR4-2933-MHz LRDIMM/4Rx4/1.2v	UCS-ML-128G4RT-H	1	48	11,600.00	556,800.00	
Intel 8280M 2.7GHz/205W 28C/38.50MB DCP DDR4 2TB 2933 MHz	UCS-CPU-I8280M	1	4	43,000.00	172,000.00	
1.9TB 2.5 inch Enterprise Value 12G SAS SSD	UCS-SD19TB121X-EV	1	6	8,064.00	48,384.00	
3.8TB 2.5 inch Enterprise Value 6G SATA SSD	UCS-SD38T61X-EV	1	4	9,888.00	39,552.00	
1.9TB 2.5 inch Enterprise Value 12G SAS SSD	UCS-SD19TB121X-EV	1	4	8,064.00	32,256.00	
UCS C480 M5 Drive Module for 8x HDD	UCSC-C480-8HDD	1	1	800.00	800.00	
UCS C-Series M5 SFF drive blanking panel	UCSC-BBLKD-S2	1	6	0.00	0.00	
Cisco VIC 1385 Dual Port 40Gb QSFP+ CNA w/RDMA	UCSC-PCIE-C40Q-03	1	1	2,428.00	2,428.00	
7.6TB 2.5in U.2 HGST SN200 NVMe High Perf. Value Endurance	UCSC-NVMEHW-H7680	1	2	22,484.00	44,968.00	
7.6TB 2.5in U.2 HGST SN200 NVMe High Perf. Value Endurance	UCSC-NVMEHW-H7680	1	2	22,484.00	44,968.00	
Cisco R42612 standard rack, w/side panels	RACK2-UCS2	1	1	6,241.00	6,241.00	
3YR SNTC 24X7X4OS UCS C480 M5 Std base chassis w/o CPU, mem,	CON-3OSP-SCC480M5	1	1	7,189.50	0.00	7,189.50
				Subtotal	1,065,128.00	7,189.50
Large Purchase discount: 61% for hardware and 35% for service			1	Discounted Price	-649,728.08	-2,516.33
IOGEAR GKM513Spill Proof Keyboard & Mouse Combo	GKM513	3	1	19.95	19.95	
ASUS 19.5" VS207D-P Widescreen LED 1600x900 VGA	VS207D-P	3	1	74.61	74.61	
				Hardware Subtotal	415,494.48	4,673.17
Software						
SQL Server2019 Enterprise Edition for Linux	NA	2	56	13,748.00	769,888.00	
Red Hat Enterprise Linux Server, 2 Sockets w/ Premium 24x7 support, with 23% discount on 4,598.46	NA	1	2	3,540.81	7,081.63	
Microsoft Problem Resolution Services	NA	2	1	259		259.00
				Software Subtotal	776,969.63	259.00
				Total	1,192,464.11	4,932.17
Source: 1-Cisco, 2-Microsft, 3-Provantage		Three-Year Cost of Ownership			1,197,396.28	USD
Audited by Francois Raab from InfoSizing.(sizing.com)						
All discounts are based on US list prices and for similar quantities and configurations. The discounts are based on the overall specific components pricing from respective vendors 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.		QpH			1,278,277.80	
		\$/kQpH			936.73	
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 pricing specifications. If you find that the stated prices are not available according to these terms, please inform the TPC at pricing@tpc.org. Thank you						



Cisco UCS C480 M5 Server

TPC-H Rev. 3.0.0
TPC-Pricing Rev.
2.7.0

Report Date: 4-Nov-2019

Measurement Results

Database Scaling (SF/Size)	30,000
Total Data Storage/Database Size	2.81
Percentage Memory/Database Size	20.48
Start of Database Load Time	10/22/2019 08:03:30
End of Database Load Time	10/23/2019 02:04:48
Database Load Time	00d 18h 01m 19s
Query Streams for Throughput Test (S)	10
TPC-H Power	1,434,669.1
TPC-H Throughput	1,138,910.7
TPC-H Composite	1,278,277.8
Total System Price Over 3 Years	\$1,197,396.28
TPC-H Price/Performance Metric (\$/kQpH@30000GB)	\$936.73

Measurement Interval

Measurement Interval in Throughput Test (Ts)	20,862.04
--	-----------

Duration of stream execution:

Power Run	Seed	Query Start Time	Total Time (hh:mm:ss)	RF1 Start Time	RF2 Start Time
		Query End Time		RF1 End Time	RF2 End Time
	1023020448	2019-10-23 10:44:43	00:38:23	2019-10-23 10:39:51	2019-10-23 11:23:09
		2019-10-23 11:23:06		2019-10-23 10:44:41	2019-10-23 11:38:37

Throughput Stream	Seed	Query Start Time	Duration (sec)	RF1 Start Time	RF2 Start Time
		Query End Time		RF1 End Time	RF2 End Time
1	1023020449	2019-10-23 11:38:36	05:47:42	2019-10-23 11:38:36	2019-10-23 11:47:14
		2019-10-23 17:26:18		2019-10-23 11:47:14	2019-10-23 12:13:16
2	1023020450	2019-10-23 11:38:36	05:47:43	2019-10-23 12:13:16	2019-10-23 12:21:05
		2019-10-23 17:26:19		2019-10-23 12:21:05	2019-10-23 12:43:42
3	1023020451	2019-10-23 11:38:36	05:37:39	2019-10-23 12:43:43	2019-10-23 12:51:33
		2019-10-23 17:16:15		2019-10-23 12:51:33	2019-10-23 13:16:18
4	1023020452	2019-10-23 11:38:36	05:34:15	2019-10-23 13:16:18	2019-10-23 13:22:31
		2019-10-23 17:12:51		2019-10-23 13:22:31	2019-10-23 13:47:30
5	1023020453	2019-10-23 11:38:36	05:28:39	2019-10-23 13:47:30	2019-10-23 13:55:41
		2019-10-23 17:07:15		2019-10-23 13:55:41	2019-10-23 14:17:39
6	1023020454	2019-10-23 11:38:36	05:05:46	2019-10-23 14:17:39	2019-10-23 14:24:15
		2019-10-23 16:44:22		2019-10-23 14:24:15	2019-10-23 14:48:01
7	1023020455	2019-10-23 11:38:36	05:39:02	2019-10-23 14:48:01	2019-10-23 14:55:47
		2019-10-23 17:17:38		2019-10-23 14:55:47	2019-10-23 15:19:10
8	1023020456	2019-10-23 11:38:36	05:40:58	2019-10-23 15:19:10	2019-10-23 15:26:08
		2019-10-23 17:19:34		2019-10-23 15:26:08	2019-10-23 15:50:02
9	1023020457	2019-10-23 11:38:36	05:30:30	2019-10-23 15:50:02	2019-10-23 15:58:24
		2019-10-23 17:09:06		2019-10-23 15:58:24	2019-10-23 16:22:54
10	1023020458	2019-10-23 11:38:36	05:43:38	2019-10-23 16:22:54	2019-10-23 16:30:38
		2019-10-23 17:22:14		2019-10-23 16:30:38	2019-10-23 17:03:18



Cisco UCS C480 M5 Server

TPC-H Rev. 3.0.0
TPC-Pricing Rev. 2.7.0
Report Date: 4-Nov-2019

TPC-H Timing Intervals (in seconds)

Stream ID	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11	Q12
0	35.91	29.25	160.65	74.10	91.15	2.86	75.56	108.00	353.80	73.34	123.20	28.25
1	97.18	129.61	1,260.04	151.47	1,203.41	74.43	1,219.28	837.67	1,609.52	670.24	1,064.28	1,143.40
2	1,167.32	162.57	1,272.53	789.71	918.45	42.02	1,482.01	1,581.47	1,586.41	1,417.75	1,022.59	520.72
3	1,200.89	133.24	1,045.60	1,146.01	1,665.51	92.06	865.05	1,091.16	948.69	1,005.01	927.11	868.10
4	890.80	65.16	1,223.41	745.89	326.84	178.27	816.42	596.41	1,177.22	1,157.95	634.38	1,360.65
5	593.05	123.42	1,121.44	1,030.92	927.00	71.99	1,136.23	1,398.15	942.58	1,231.36	1,856.49	613.61
6	939.90	125.64	1,498.34	980.13	783.68	41.63	769.51	848.16	1,379.55	314.07	962.08	1,106.60
7	1,526.46	131.77	798.67	1,131.77	1,921.63	83.96	797.35	1,272.89	1,932.25	1,226.64	779.66	976.99
8	1,489.57	101.81	1,096.30	197.49	1,270.87	95.97	1,822.44	1,460.57	1,362.16	1,017.82	439.12	1,140.70
9	1,273.49	191.85	1,179.25	793.01	631.94	81.40	1,099.53	1,090.33	1,429.25	1,206.20	774.40	567.68
10	1,067.44	96.51	1,347.13	927.17	375.92	14.25	1,152.12	1,180.25	995.01	995.65	751.24	939.54
Qi Min	35.91	29.25	160.65	74.10	91.15	2.86	75.56	108.00	353.80	73.34	123.20	28.25
Qi Avg	934.73	117.35	1,091.21	724.33	919.67	70.80	1,021.41	1,042.28	1,246.95	937.82	848.60	842.39
Qi Max	1,526.46	191.85	1,498.34	1,146.01	1,921.63	178.27	1,822.44	1,581.47	1,932.25	1,417.75	1,856.49	1,360.65

Stream ID	Q13	Q14	Q15	Q16	Q17	Q18	Q19	Q20	Q21	Q22	RF1	RF2
0	249.88	22.13	15.39	59.95	40.49	347.66	37.65	52.68	277.13	42.99	289.79	927.93
1	1,940.85	131.02	128.99	445.70	1,084.97	2,393.00	1,093.02	2,991.89	1,050.90	140.61	517.72	1,561.14
2	1,329.83	403.24	155.58	348.86	666.34	1,902.87	1,048.65	659.88	780.07	1,603.15	469.19	1,357.33
3	1,589.18	55.27	163.63	435.31	1,164.57	2,065.13	1,023.99	716.28	1,428.87	627.83	470.02	1,485.22
4	1,827.19	207.11	151.97	1,027.32	564.42	2,067.10	1,156.06	1,186.55	1,886.12	806.70	372.47	1,499.20
5	1,579.46	105.20	259.17	232.09	549.48	2,160.25	997.33	556.54	1,400.04	832.56	490.33	1,317.94
6	1,310.91	103.11	120.41	305.59	727.03	1,037.56	1,166.26	848.00	1,274.79	1,701.94	396.19	1,425.75
7	1,032.87	166.10	114.05	223.12	1,161.06	1,666.59	735.39	742.61	891.65	1,028.03	465.86	1,402.80
8	2,365.17	111.17	210.75	216.65	822.20	1,876.24	161.91	961.50	1,214.22	1,022.52	418.02	1,433.25
9	2,060.36	161.51	143.86	442.68	888.85	1,682.69	1,015.74	424.81	1,958.79	731.52	502.03	1,469.74
10	2,116.46	175.99	105.53	1,046.82	1,422.37	2,201.32	320.84	693.37	1,680.72	1,011.02	464.15	1,960.01
Qi Min	249.88	22.13	15.39	59.95	40.49	347.66	37.65	52.68	277.13	42.99	289.79	927.93
Qi Avg	1,582.01	149.26	142.67	434.92	826.53	1,763.67	796.08	894.01	1,258.48	868.08	441.43	1,440.03
Qi Max	2,365.17	403.24	259.17	1,046.82	1,422.37	2,393.00	1,166.26	2,991.89	1,958.79	1,701.94	517.72	1,960.01

Preface

TPC Benchmark™ H Overview

The TPC Benchmark™ H (TPC-H) is a decision support benchmark. It consists of a suite of business oriented ad-hoc queries and concurrent data modifications. The queries and the data populating the database have been chosen to have broad industry-wide relevance while maintaining a sufficient degree of ease of implementation. This benchmark illustrates decision support systems that

- Examine large volumes of data;
- Execute queries with a high degree of complexity;
- Give answers to critical business questions.

TPC-H evaluates the performance of various decision support systems by the execution of sets of queries against a standard database under controlled conditions. The TPC-H queries:

- Give answers to real-world business questions;
- Simulate generated ad-hoc queries (e.g., via a point and click GUI interface);
- Are far more complex than most OLTP transactions;
- Include a rich breadth of operators and selectivity constraints;
- Generate intensive activity on the part of the database server component of the system under test;
- Are executed against a database complying to specific population and scaling requirements;
- Are implemented with constraints derived from staying closely synchronized with an on-line production database.

The TPC-H operations are modeled as follows:

- The database is continuously available 24 hours a day, 7 days a week, for ad-hoc queries from multiple end users and data modifications against all tables, except possibly during infrequent (e.g., once a month) maintenance sessions;
- The TPC-H database tracks, possibly with some delay, the state of the OLTP database through on-going refresh functions which batch together a number of modifications impacting some part of the decision support database;
- Due to the world-wide nature of the business data stored in the TPC-H database, the queries and the refresh functions July be executed against the database at any time, especially in relation to each other. In addition, this mix of queries and refresh functions is subject to specific ACIDity requirements, since queries and refresh functions July execute concurrently;
- To achieve the optimal compromise between performance and operational requirements, the database administrator can set, once and for all, the locking levels and the concurrent scheduling rules for queries and refresh functions.

The performance metric reported by TPC-H is called the TPC-H Composite Query-per-Hour Performance Metric (QphH@Size), and reflects multiple aspects of the capability of the system to process queries. These aspects include the selected database size against which the queries are executed, the query processing power when queries are submitted by a single stream and the query throughput when queries are submitted by multiple concurrent users. The TPC-H Price/Performance metric is expressed as \$/QphH@Size. To be compliant with the TPC-H standard, all references to TPC-H results for a given configuration must include all required reporting components. The TPC believes that comparisons of TPC-H results measured against different database sizes are misleading and discourages such comparisons.

The TPC-H database must be implemented using a commercially available database management system (DBMS) and the queries executed via an interface using dynamic SQL. The specification provides for variants of SQL, as implementers are not required to have implemented a specific SQL standard in full.

TPC-H uses terminology and metrics that are similar to other benchmarks, originated by the TPC and others. Such similarity in terminology does not in any way imply that TPC-H results are comparable to other benchmarks. The only benchmark results comparable to TPC-H are other TPC-H results compliant with the same revision.

Despite the fact that this benchmark offers a rich environment representative of many decision support systems, this benchmark does not reflect the entire range of decision support requirements. In addition, the extent to which a customer can achieve the results reported by a vendor is highly dependent on how closely TPC-H 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 as a result of these and other factors. Therefore, TPC-H should not be used as a substitute for a specific customer application benchmarking when critical capacity planning and/or product evaluation decisions are contemplated.

Further information is available at www.tpc.org

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 Cisco Systems, Inc.

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 by 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 File Archive contains the Operating System and DBMS parameters used in this benchmark.

0.3 Configuration Diagrams

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 Cisco UCS C480 M5 server features:

- Up to four Intel® Xeon® Scalable Processors (up to 28 cores per socket) – Skylake or Cascade Lake family
- 2933-MHz DDR4 memory, 48 DDR4 DIMM slots: 16, 32, 64, and 128 GB; up to 2933 MHz
- 12 PCIe 3.0 slots plus 1 dedicated 12-Gbps RAID controller slot
- RAID controllers
 - Cisco 12-Gbps Modular RAID Controller (PCIe 3.0) with 4-GB Flash-Backed Write Cache (FBWC), providing enterprise-class data protection for up to 24 SAS and SATA HDDs and SSDs

- 12-Gbps 9460-8i RAID controller with 2-GB FBWC provides support for up to 8 SAS and SATA HDDs and SSDs in the auxiliary drive modules
- PCIe NVMe switch card for up to 8 PCIe NVMe drives in the auxiliary drive module
- Internal Storage
 - Support for up to 32 hot-swappable 2.5-inch Small Form Factor (SFF) drives
 - Up to 24 front loading 2.5-inch SAS/SATA HDDs and SSDs and PCIe NVMe drives
 - Up to 8 top loading 2.5-inch SAS/SATA/PCIe HDDs and SSDs or NVMe drives in the C480 M5 auxiliary drive module
 - DVD drive option
- Internal Secure Digital (SD) or M.2 boot options
- Dual 10GBASE-T Intel x550 Ethernet ports



The measured configuration consists of a Cisco UCS C480 M5 Rack-Mount Server with:

- 4 x Intel 2nd Gen Xeon 2nd Gen Scalable 8280M Processors (2.7 GHz, 38.5MB L1 cache, 205W)
- 6 TB of memory (48x 128GB DDR4 2933MHz LRDIMM)
- 8 x Cisco HHL AIC 3.2TB HGST SN260 NVMe Extreme Performance High Endurance
- 4 x Cisco 2.5in U.2 4.0TB Intel P4500 NVMe High Perf. Value Endurance
- 1 x Cisco 12-Gbps modular RAID controller with 4-GB cache module
 - 10 x 1.9TB 2.5-inch Enterprise Value 12G SAS SSD
 - 4 x 3.8TB 2.5-inch Enterprise Value 6G SATA SSD

In the priced configuration the four 4.0TB Intel P4500 NVMe were substituted by four equivalent 7.6TB 2.5in U.2 HGST SN200 NVMe. This substitution was based on the documented specifications of these NVMe devices. According to these specifications, all aspects of the priced devices that affected these benchmark results were equal or better than the tested devices.

Clause 1: Logical Database Design

1.1 Database Definition Statements

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

The Supporting File Archive contains the table definitions and all other statements used to set up the test and qualification databases.

1.2 Physical Organization

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.

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.

Clause 2: Queries and Refresh Functions Related Items

2.1 Query Language

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

SQL was the query language used to implement the queries.

2.2 Verifying Method of Random Number Generation

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

TPC-supplied DBGEN version 2.18.0 and QGEN version 2.18.0 were used.

2.3 Generating Values for Substitution Parameters

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.

TPC supplied QGEN version 2.18.0 was used to generate the substitution parameters.

2.4 Query Text and Output Data from Qualification 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.

Supporting Files Archive contains the actual query text and query output. Following are the modifications to the query.

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

Supporting Files Archive contains the query substitution parameters and seed 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 “Read committed” isolation level.

2.7 Source Code of Refresh Functions

The details of how the refresh functions were implemented must be disclosed (including source code of any non-commercial program used).

Supporting Files Archive contains the Source Code of refresh functions.

Clause 3: Database System Properties

3.1 ACID Properties

The ACID (Atomicity, Consistency, Isolation, and Durability) properties of transaction processing systems must be supported by the system under test during the timed portion of this benchmark. Since TPC-H is not a transaction processing benchmark, the ACID properties must be evaluated outside the timed portion of the test.

All ACID tests were conducted according to specification. The Supporting Files Archive contains the source code of the ACID test scripts.

3.2 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.

3.2.1 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.

3.2.2 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.

3.3 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:

$$O_TOTALPRICE = SUM(L_EXTENDEDPRICE - L_DISCOUNT) * (1 + L_TAX)$$
For each ORDER and LINEITEM defined by (O_ORDERKEY = L_ORDERKEY)

3.3.1 Consistency Test

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.

The following steps were performed to verify consistency:

1. The consistency of the ORDER and LINEITEM tables was verified based on a sample of O_ORDERKEYs.
2. At least 100 ACID Transactions were submitted.
3. The consistency of the ORDER and LINEITEM tables was re-verified.

The Consistency test was performed as part of the Durability test explained in section 3.5.

3.4 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.

3.4.1 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.
4. The ACID query completed. It returned the data as committed by the ACID Transaction.

3.4.2 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.

3.4.3 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:

$$T2.L_EXTENDEDPRICE = T1.L_EXTENDEDPRICE \\ + (DELTA1 * (T1.L_EXTENDEDPRICE / T1.L_QUANTITY))$$

3.4.4 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 $T2.L_EXTENDEDPRICE = T1.L_EXTENDEDPRICE$.

3.4.5 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 to verify isolation of concurrent read and write transactions on different tables:

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.

3.4.6 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 to verify isolation of update transaction during continuous read-only query:

1. An ACID Transaction T1 was started, executing 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.

3.5 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.

3.5.1 Permanent Unrecoverable Failure of Any Durable Medium

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.

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.

A backup of the database was taken. The tests were conducted on the qualification database.

The steps performed to demonstrate that committed updates are preserved across a permanent irrecoverable failure of disk drive containing data tables:

1. The database was backed up.
2. The consistency of the ORDERS and LINEITEM tables were verified.
3. Eleven streams of ACID transactions were started. Each stream executed a minimum of 100 transactions.
4. While the test was running, one of the 3200GB HGST SN260 NVMe was detached (making it logically unavailable).
5. A checkpoint was issued to force a failure.
6. Database error log recorded the failure.
7. The running ACID transactions were stopped.
8. The Database log was backed up.
9. The disk drive was reattached.
10. The database was dropped and restored.
11. When database restore completed, issued a command to apply the backed up log file.
12. The counts in the history table and success files were compared and verified, and the consistency of the ORDERS and LINEITEM tables was verified.

Testing the permanent irrecoverable failure of disk drive containing database log file was tested as part of the system crash test (see section 3.5.2).

3.5.2 Loss of Log and System Crash Test

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.

1. The consistency of the ORDERS and LINEITEM tables were verified.
2. Eleven 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-10 array was physically removed.
4. The database log RAID-10 volume went to a degraded state.
5. The tests were still running without any problem even after the log disk was in a degraded state.
6. While the streams of ACID transactions were still running, the system was powered off.
7. When power was restored, the system booted and the database was restarted.
8. The database went through a recovery period.
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.

3.5.3 Memory Failure

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

See section 3.5.2

Clause 4: Scaling and Database Population

4.1 Initial Cardinality of Tables

The cardinality (e.g., 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 Benchmark H defined tables and the row count for each table as they existed upon completion of the build.

Table 4. 1: Initial Number of Rows

Table Name	Row Count
Region	5
Nation	25
Supplier	300,000,000
Customer	4,500,000,000
Part	6,000,000,000
Partsupp	24,000,000,000
Orders	45,000,000,000
Lineitem	179,999,978,268

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.

The storage system of the tested configuration consisted of:

- 8 x Cisco HHHL AIC 3.2T HGST SN260 NVMe Extreme Perf High Endurance
- 4 x Cisco 2.5in U.2 4.0TB Intel P4500 NVMe High Perf. Value End
- 1 x Cisco 12-Gbps modular RAID controller with 4-GB cache module
 - 10 x 1.9TB 2.5-inch Enterprise Value 12G SAS SSD
 - 4 x 3.8TB 2.5-inch Enterprise Value 6G SATA SSD

The database tables were hosted across eight Cisco HHHL AIC 3.2T HGST SN260 NVMe Extreme Perf High Endurance cards. The tempdb data files were stored across four 4.0TB Intel P4500 NVMe High Perf. Value End SSD drives. The database log and tempdb log files resided on a RAID-10 array of ten 1.9 TB 2.5-inch Enterprise Value 12G SAS SSD drives. The database backup was hosted on RAID-0 array made of four 3.8TB 2.5-inch Enterprise Value 6G SATA SSD drives. A detailed description of distribution of database filegroups and log can be found in Table 4.2.

Table 4.2: Disk Array to Logical Drive Mapping

Logical Allocation	Drive Description	Usable Drive Size (TB)	RAID Format	Disk Group Spindles	Total Space (TB)	Drive Letter/Mount Point
OS, SQL Binaries	1.9TB 2.5-inch Enterprise Value 12G SAS SSD	1.7	10	10	0.5	/sda/ - XFS Partition
Swap					1.5	/sdb/ - XFS Partition [SWAP]
SQL DB LOG					6.5	/sdd/ - XFS Partition Mount Point: /sqllog
SQL DB DATA Files #1	Cisco HHL AIC 3.2T HGST SN260 NVMe Extreme Perf High Endurance	2.98	No RAID	1	2.98	/nvme0n1/- XFS Partition; Mount Point: /CPU1_NVMe0_DATA1
SQL DB DATA Files #2	Cisco HHL AIC 3.2T HGST SN260 NVMe Extreme Perf High Endurance	2.98	No RAID	1	2.98	/nvme1n1/- XFS Partition; Mount Point: /CPU1_NVMe1_DATA2
SQL DB DATA Files #3	Cisco HHL AIC 3.2T HGST SN260 NVMe Extreme Perf High Endurance	2.98	No RAID	1	2.98	/nvme4n1/- XFS Partition; Mount Point: /CPU2_NVMe4_DATA3
SQL DB DATA Files #4	Cisco HHL AIC 3.2T HGST SN260 NVMe Extreme Perf High Endurance	2.98	No RAID	1	2.98	/nvme5n1/- XFS Partition; Mount Point: /CPU2_NVMe5_DATA4
SQL DB DATA Files #5	Cisco HHL AIC 3.2T HGST SN260 NVMe Extreme Perf High Endurance	2.98	No RAID	1	2.98	/nvme6n1/- XFS Partition; Mount Point: /CPU3_NVMe6_DATA5
SQL DB DATA Files #6	Cisco HHL AIC 3.2T HGST SN260 NVMe Extreme Perf High Endurance	2.98	No RAID	1	2.98	/nvme7n1/- XFS Partition; Mount Point: /CPU3_NVMe7_DATA6
SQL DB DATA Files #7	Cisco HHL AIC 3.2T HGST SN260 NVMe Extreme Perf High Endurance	2.98	No RAID	1	2.98	/nvme9n1/- XFS Partition; Mount Point: /CPU4_NVMe9_DATA7
SQL DB DATA Files #8	Cisco HHL AIC 3.2T HGST SN260 NVMe Extreme Perf High Endurance	2.98	No RAID	1	2.98	/nvme10n1/- XFS Partition; Mount Point: /CPU4_NVMe10_DATA8
TempDB Drive #1	Cisco 2.5in U.2 4.0TB Intel P4500 NVMe High Perf. Value End	3.7	No RAID	1	3.7	/nvme2n1 - XFS Partition; Mount Point: /CPU1_NVMe2_TempdbDATA1

TempDB Drive #2	Cisco 2.5in U.2 4.0TB Intel P4500 NVMe High Perf. Value End	3.7	No RAID	1	3.7	/nvme3n1 - XFS Partition; Mount Point: /CPU2_NVMe3_TempdbDATA2
TempDB Drive #3	Cisco 2.5in U.2 4.0TB Intel P4500 NVMe High Perf. Value End	3.7	No RAID	1	3.7	/nvme8n1 - XFS Partition; Mount Point: /CPU3_NVMe8_TempdbDATA3
TempDB Drive #4	Cisco 2.5in U.2 4.0TB Intel P4500 NVMe High Perf. Value End	3.7	No RAID	1	3.7	/nvme11n1 - XFS Partition; Mount Point: /CPU4_NVMe11_TempdbDATA4
Backup	3.8TB 2.5-inch Enterprise Value 6G SATA SSD	3.5	0	4	14	/sdc - XFS Partition; Mount Point: /sqlbcp

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.

4.4 Implementation of RAID

Implementations July 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.

The database log files resided on a RAID-10 array of ten 1.9 TB 2.5-inch Enterprise Value 12G SAS SSD drives. The database backup was hosted on RAID-0 array made of four 3.8 TB 2.5 inch Enterprise Value 6G SATA SSD 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 2.18.0 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 18 hours 01 minutes and 19 seconds.

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.

The database storage ratio can be found in Table 4.7

Table 4.7: Data Storage Ratio

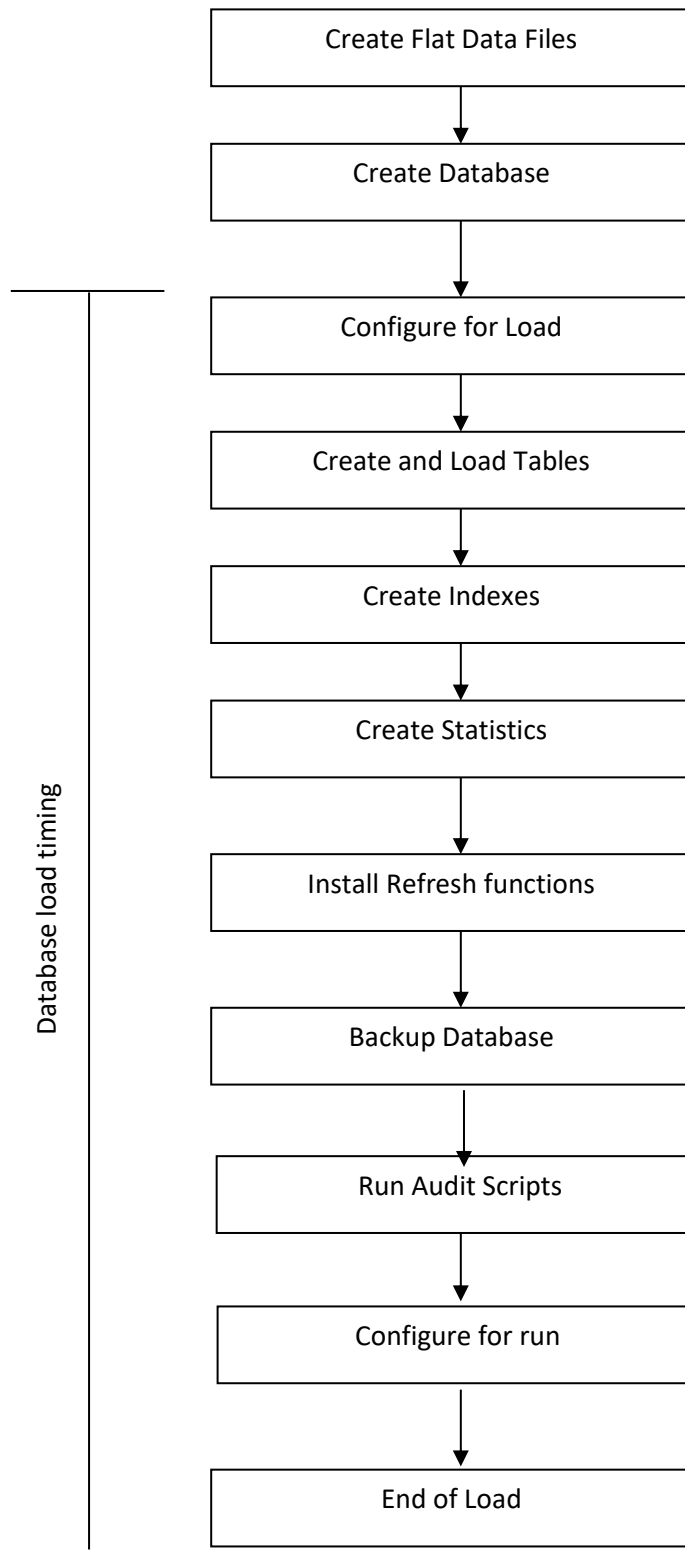
Storage Devices	Space per Disk(GB)	Total Disk Space(GB)	Total Storage Capacity(GB)	Scale factor	Data Storage Ratio
10 x 1.9 TB 2.5-inch Enterprise Value 12G SAS SSD in RAID 10	1,740.8	17,408	84,419	30,000	2.81
4 x 3.8 TB 2.5-inch Enterprise Value 6G SATA SSD in RAID 0	3,584	14,336			
8 x Cisco HHL AIC 3.2T HGST SN260 NVMe Extreme Perf High Endurance	3051.52	24,412.16			
4 x Cisco 2.5in U.2 7.6TB HGST SN200 NVMe High Perf. Value Endurance	7065.6	28,262.4			

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 were created using DBGEN. The tables were loaded as shown in Figure 4.8.

Figure 4.8: Block Diagram of Database Load Process



4.9 Qualification Database Configuration

Any differences between the configuration of the qualification database and the test database must be disclosed.

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 must be disclosed.

Available Memory: 6144GB

Scale Factor:30000

The memory to database size percentage is 20.48%.

Clause 5: Performance Metrics and Execution Rules Related Items

5.1 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. RF1 Refresh Function
2. Stream 00 Execution
3. RF2 Refresh Function

5.2 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.

See the Numerical Quantities Summary in the Executive Summary at the beginning of this report.

5.3 Number of Streams for The Throughput Test

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

Ten query streams were used for throughput test. Each stream running all twenty-two queries. One stream was used for RF.

5.4 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.

See the Numerical Quantities Summary in the Executive Summary at the beginning of this report.

5.5 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.

See the Numerical Quantities Summary in the Executive Summary at the beginning of this report.

5.6 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.

See the Numerical Quantities Summary in the Executive Summary at the beginning of this report.

5.7 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.

See the Numerical Quantities Summary in the Executive Summary at the beginning of this report.

5.8 Performance Metrics

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

See the Numerical Quantities Summary in the Executive Summary at the beginning of this report.

5.9 The Performance Metric and Numerical Quantities from Both Runs

A description of the method used to determine the reproducibility of the measurement results must be reported. This must include the performance metrics (QppH and QthH) from the reproducibility runs.

Performance results from the first two executions of the TPC-H benchmark indicated the following difference for the metric points:

Run	QppH @ 30,000GB	QthH @ 30,000GB	QphH @ 30,000GB
Run 1	1,634,382.0	1,198,495.8	1,399,571.3
Run 2	1,434,699.1	1,138,910.7	1,278,277.8

5.10 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.

SQL Server was restarted between Run1 and Run2.

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 July 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. Step Master 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 and subsequent scripts are called from within the scripts.

The source for Step Master and the RF scripts is disclosed in the Supporting Files 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 Driver section for details.

6.3 Profile-Directed Optimization

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

Profile-directed optimization was not used.

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 vendor part number, description, and release/revision level, and either 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.

A detailed list of all hardware and software, including the 3-year support, is provided in the Executive Summary in the Abstract section of this report. The price quotations are included in Appendix A.

7.2 Total 3 Year Price

The total 3-year price of the entire configuration must be reported including: hardware, software, and maintenance charges. Separate component pricing is recommended. The basis of all discounts used must be disclosed.

A detailed list of all hardware and software, including the 3-year support, is provided in the Executive Summary in the Abstract section of this report. The price quotations are included in Appendix A. This purchase qualifies for a 61% discount from Cisco Systems, Inc. on all the hardware and 35% on services.

7.3 Availability Date

The committed delivery date for general availability of products used in the price calculations must be reported. When the priced system includes products with different availability dates, the availability date reported on 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.

The total system availability date is Nov 4, 2019.

7.4 Orderability Date

For each of the components that are not orderable on the report date of the FDR, the following information must be included in the FDR:

- Name and part number of the item that is not orderable*
- The date when the component can be ordered (on or before the Availability Date)*
- The method to be used to order the component (at or below the quoted price) when that date arrives*
- The method for verifying the price*

All components are orderable at the time of publication date.

7.5 Country-Specific Pricing

Additional Clause 7 related items must be included in the Full Disclosure Report for each country-specific priced configuration. Country-specific pricing is subject to Clause 7.1.7.

The configuration is priced for the United States of America.

7.6 Tested and Priced configurations

Additional Clause 5.7.3.3 of the Pricing specification related items must be included in the Full Disclosure Report.

If the following criteria are completely satisfied, an allowed storage device substitution can be done without additional measurement.

- 1. The formatted capacity of the substitute device must be equal or greater than the substituted device.*
- 2. The substitute device must have the same interface type as the substituted device.*
- 3. Characteristics of the substitute devices, such as those listed below must be the same or better than the substituted devices.*
 - read/write latency, read/write throughput, form factor, NAND type, cache size, controller sizing, memory population model, write endurance, amount of NAND on the board, PCIe lanes*

Although four 4.0TB 2.5inch U.2 Intel P4500 NVMe High Perf. Value End SSDs have been used for storing the Tempdb data files, the system was priced with 4 x 7.6TB 2.5in U.2 HGST SN200 NVMe High Perf. Value End SSDs. The below table describes the tested and Priced configurations.

Tested Configuration with Intel SSDs	Priced Configuration with HGST SSDs
Used 4 x 4.0TB 2.5inch U.2 Intel P4500 NVMe High Perf. Value End SSDs for storing TempDB data files	Used 4 x 7.6TB 2.5in U.2 HGST SN200 NVMe High Perf. Value End for storing TempDB data files

Clause 8: Full Disclosure

8.1 Supporting File Index

An index for all files included in the supporting files archive as required by Clause 8.3.2 through 8.3.8 must be provided in the report.

Clause	Description	Archive File Pathname
Clause 1	OS and DB parameter settings	SupportingFilesArchive\Clause1
Clause 2	DB creation scripts	SupportingFilesArchive\Clause2
Clause 3	ACID scripts, ACID output	SupportingFilesArchive\Clause3
Clause 4	DB Load scripts, Qualification output	SupportingFilesArchive\Clause4
Clause 5	Query output results	SupportingFilesArchive\Clause5
Clause 6	Implementation Specific layer source code	SupportingFilesArchive\Clause6
Clause 8	Query substitution parameters, RF function source	SupportingFilesArchive\Clause8

Clause 9: Audit Related Items

Auditors' Information and Attestation Letter

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 benchmark was audited by:

Francois Raab,
Infosizing,
20 Kreg Lane.
Manitou Springs, CO 80829.
Phone Number: 719-473-7555.

The auditor's letter is included in the following section.

Benchmark sponsor: Siva Sivakumar
Sr. Director, UCS Solutions,
Computing Systems Product Group,
Cisco Systems
3800 Zanker Road
San Jose, CA 95134

November 1, 2019

I verified the TPC Benchmark H (TPC-H™ v2.18.0) performance of the following configuration:

Platform: Cisco UCS C480 M5 Rack-Mount Server
Operating System: Red Hat Enterprise Linux 8.0
Database Manager: Microsoft SQL Server 2019 Enterprise Edition for Linux
Other Software: n/a

The results were:

Performance Metric 1,278,277.8 QphH@30,000GB

TPC-H Power 1,434,699.1

TPC-H Throughput 1,138,910.7

Database Load Time 18h 01m 19s

Server

Cisco UCS C480 M5 Rack-Mount Server

CPU 4 x Intel 2nd Gen Xeon Scalable 8280M (2.7GHz, 38.5MB L3)

Memory 6 TB

Disks

Qty	Size	Type
10	1.9 TB	2.5.inch 12G SAS SSD
4	3.8 TB	2.5.inch 12G SAS SSD
4	7.6 TB	2.5.inch HGST SN200 NVMe
8	3.2 TB	2.5.inch HGST SN260 NVMe

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 30,000GB 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 10 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:

The measured system included four (4) 4.0TB Intel P4500 NVMe that were substituted by four (4) equivalent 7.6TB 2.5in U.2 HGST SN200 NVMe, in the priced configuration. Based on the specifications of these devices, it is my opinion that this substitution has no significant effect on performance.

Respectfully Yours,



François Raab, TPC ~~Certified~~ Auditor

Appendix A: Price Quotes

ASUS 19.5" VS207D-P Widescreen LED 1600x900 VGA HDCP Splendid Video Int



Mfr List: \$89.00

Only **\$74.61**

Save \$14.39 (16%)

[ADD TO CART](#)

[209 IN STOCK](#)

[ADD TO WISH LIST](#)

Click image to enlarge or see more views.



See all products

Model VS207D-P

UPC Code: 886227347377

Provantage Code: ASUM01C

Condition: Factory New

- ▶ Product Type: LCD Monitor
- ▶ Aspect Ratio: 16:9
- ▶ Screen Mode: HD+
- ▶ Maximum Resolution: 1600 x 900
- ▶ Viewable Screen Size: 19.5"



[Computing](#) > [Monitors and Displays](#) > [Monitors](#)

Specifications

Accessories

Related/Similar

ASUS VS207D-P Specifications

19.5" VS207D-P Widescreen LED 1600x900 VGA HDCP Splendid Video Int

Overview

Superior Image Quality Meets Classic Elegant Design





Spill-Resistant Keyboard and Mouse Combo

MPN: **GKM513** UPC: 881317508560

MSRP: \$19.95



OVERVIEW FEATURES REQUIREMENTS PACKAGE SPECS SUPPORT

Distributor	Part #
Ingram	GF9752
DH	GKM513
Synnex	2959476
SKU	ECCN
GKM513	EAR99
Function	GKM513
Housing	
Case	ABS plastic
Operating System Support	Windows Vista®, Windows® 7, Windows® 8, Windows® 8.1, Mac OS X 10.2+

Microsoft Corporation
One Microsoft Way
Redmond, WA 98052-6399

Tel 425 882 8080
Fax 425 936 7329
<http://www.microsoft.com/>

Microsoft

October 28, 2019

Babu Mahadevan V
Cisco Systems
3800 Zanker Road
San Jose, CA 95134

Here is the information you requested regarding pricing for several 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 2019 Enterprise Edition <i>2 Core License/ Open Program – No Level - ERP</i>	\$13,748.00	56	\$769,888.00
Support			
Microsoft Problem Resolution Services <i>Professional Support (1 Incident).</i>	\$259.00	1	\$259.00

All software components are currently orderable and available. A list of Microsoft's resellers can be found in the Microsoft Product Information Center at <http://www.microsoft.com/products/info/render.aspx?view=22&type=how>

Defect support is included in the purchase price. Additional support is available from Microsoft PSS on an incident by incident basis at \$259 call.

This quote is valid for the next 90 days.

Reference ID: TPCH_zdwhgbl4272499_2019_cbm
