Cisco Systems, Inc.

TPC Benchmark™ H Full Disclosure Report

Cisco UCS C420 M3 Rack-Mount Server

using

Sybase IQ Version 16

and

Red Hat Enterprise Linux Server Release 6.4

First Edition

October, 2013

Table of Contents

Abstract

This report documents the methodology and results of the TPC Benchmark© H test conducted on the Cisco UCS C420 M3 Server using Red Hat Enterprise Linux Server Release 6.4 and database Sybase IQ Version 16.

Cisco UCS C420 M3 Server

Company Name	System Name	Database Software	Operating System
Cisco Systems, Inc	Cisco UCS C420 M3 Server	Sybase IQ Version 16.0 SP02	Red Hat Enterprise Linux Server Release 6.4

TPC Benchmark© H Metrics

Total System Cost	TPC-H Throughput	Price/Performance	Availability Date
$295,599 USD	230.119.9 QphH@3000GB	$1.29 USD $/QphH@3000GB	December 30, 2013

	Cisco UCS C420 M3 Server	TPC-H Rev. 2.16.0 TPC-Pricing Rev. 1.7.0 Report Date: October 31, 2013
Total System Cost	Composite Query per Hour Metric	Price / Performance
$295,599 USD	**230,119.9** QphH@3000GB	**$1.29 USD** $ / QphH@3000GB

Database Size	Database Manager	Operating System	Other Software	Availability Date
3000GB	**Sybase IQ Version 16.0 SP02**	**Red Hat Enterprise Linux Server Release 6.4**		**December 30, 2013**

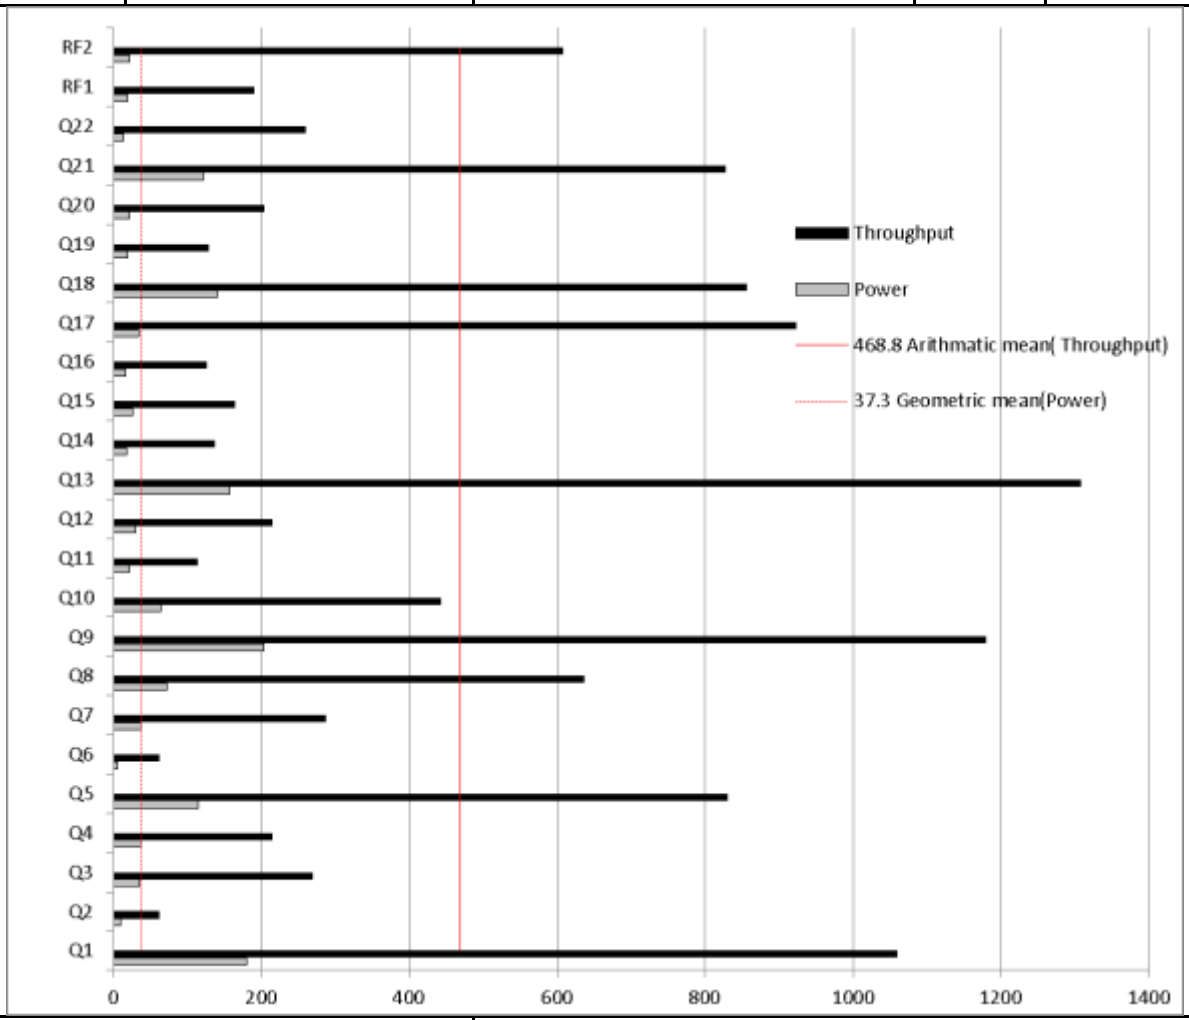

Database Load Time = 2h 25m 34s	**Storage Redundancy Level**	
Load Includes Backup: N	Base Tables and Auxiliary Data Structures	1
Total Data Storage / Database Size = 4.44	DBMS Temporary Space	0
Percentage Memory / Database Size = 34.1%	OS and DBMS Software	1

System Configuration:	**Cisco UCS C420 M3 Server**
Processors/Cores/Threads/Model:	4/32/64 Intel Xeon E5-4650 Processor (2.7 GHz, 20MB cache, 130W)
Memory:	1024 GB
Storage:	8 x 600GB 6Gb SAS 10K RPM SFF HDD 8 x Cisco UCS 200 GB SATA 2.5" Enterprise Performance SSD 6 x Cisco UCS Fusion ioDrive2 1205M Adapter
Table Storage:	13.6 TB

Cisco	Cisco UCS C430 M3 Server	TPC-H Rev. 2.16.0 TPC-Pricing Rev. 1.7.0 Report Date: 31-Oct-2013

Description	Part Number	Brand	Source	Unit Price	Qty	Extended Price	3 Year Maint. Price
Server Hardware and OS							
UCS C420 M3 w/o CPU mem HDD PCIe PSU rail kit	UCSC-C420-M3	Cisco	1	7,500.00	1	$7,500	
SMARTNET Onsite 24X7X4 UCS C420 M3 Server, 36 Month(s)	CON-OSP-C420	Cisco	1	2,631.00	1		$2,631
2.70 GHz E5-4650 130W 8C/20MB Cache/DDR3 1600MHz	UCS-CPU-E5-4650	Cisco	1	10,255.00	4	$41,020	
32GB DDR3-1600-MHz LR DIMM/PC3-12800/quad rank/x4/1.35v	UCS-ML-1X324RY-A	Cisco	1	2,639.00	32	$84,448	
600GB 6Gb SAS 10K RPM SFF HDD/hot plug/drive sled mounted	A03-D600GA2	Cisco	1	1,086.00	8	$8,688	
200GB 2.5 SATA Enterprise performance SSD	UCS-SD200G0KS2-EP	Cisco	1	3,598.00	8	$28,784	
Cisco UCS 1205GB MLC Fusion ioDrive2	UCSC-F-FIO-1205M	Cisco	1	37,500.00	6	$225,000	
1200W 2u Power Supply For UCS	UCSC-PSU2-1200	Cisco	1	652.00	2	$1,304	
MegaRAID 9271CV with 8 internal SAS/SATA ports with Supercap	UCS-RAID9271CV-8I	Cisco	1	1,686.00	1	$1,686	
8 Drive Backplane W/ Expander For C-Series	UCSC-DBKP-08E	Cisco	1	528.00	2	$1,056	
SAS/SATA cable (long) for C420 M3	UCSC-CABLE-L	Cisco	1	Included	2		
Heatsink for the C420 M3	UCSC-HS-C420M3	Cisco	1	Included	4		
Full height PCIe filler for C-Series	UCSC-PCIF-01F	Cisco	1	Included	2		
Half height PCIe filler for UCS	UCSC-PCIF-01H	Cisco	1	Included	3		
2U Rail Kit for UCS C-Series servers	UCSC-RAIL-2U	Cisco	1	Included	1		
RHEL/4 Socket/1 Guest/3Yr Svcs Required	RHEL-4S-1G-3A	Cisco	1	Included	1		
AC Power Cord - 250V 10A - PRC	CAB-250V-10A-CN	Cisco	1	Included	2		
ISV 24X7 Rhel/4 Socket/1 Guest 36 Month(s)	CON-ISV1-RH4S1G3A	Cisco	1	9,588.00	1		$9,588
IOGEAR GKM502 Compact Wired Keyboard and Mouse Combo	GKM502	Provantage	3	$11	3	$32	
Acer 17" V173 DJOb LCD Monitor 1280X1024	UM.BV3AA.D01	Provantage	3	$112	3	$336	
Cisco R42610 expansion rack	RACK-UCS	Cisco	1	$2,857	1	$2,857	
					Subtota	$402,711	$12,219
Large Purchase Discount [i] 57%		Cisco	1			-$229,545	-$6,964
				Hardware Subtotal		$173,166	$5,255
Software							
Sybase IQ Single App Svr, per cpu core	13545	SAP	2	$2,595	32	$83,040	
3 yr support Single App Svr, per cpu core	98480	SAP	2	$1,713	32		$54,816
Sybase IQ Discount (15% of Net)		SAP	2			-$12,456	-$8,222
					Softwar	$70,584	$46,594
					Total	$243,750	$51,849

Pricing: 1=Cisco 2=SAP 3=Provantage;

Audited by Francois Raab from InfoSizing, Inc. (sizing.com)

[i]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.

Three-Year Cost of Ownership: $295,599 USD

QphH: 230,119.9

$ / QphH: $1.29 USD

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.

Cisco	Cisco UCS C420 M3 Server	TPC-H Rev. 2.16.0 TPC-Pricing Rev. 1.7.0 October 31, 2013

Measurement Results

Database Scaling (SF/Size)	3,000
Total Data Storage/Database Size	4.44
Percentage Memory/Database Size	34.1
Start of Database Load Time	10/22/2013 21:02:40
End of Database Load Time	10/22/2013 23:28:14
Database Load Time	02:25:34
Query Streams for Throughput Test (S)	8
TPC-H Power	289,320.0
TPC-H Throughput	183,033.2
TPC-H Composite	230,119.9
Total System Price Over 3 Years	$295,599
TPC-H Price/Performance Metric ($/QphH@3000GB)	$1.29

Measurement Interval

Measurement Interval in Throughput Test (Ts)	10,385

Duration of throughput stream execution:

Power Run	Seed	Query Start Time Query End Time	Duration (sec)	RF1 Start Time RF1 End Time	RF2 Start Time RF2 End Time
	1022232814	10/22//2013 23:28:45 10/22//2013 23:51:45	1,434	10/22/2013 23:28:15 10/22/2013 23:28:36	10/22/2013 23:51:47 10/22/2013 23:52:09

Stream	Seed	Query Start Time Query End Time	Duration (sec)	RF1 Start Time RF1 End Time	RF2 Start Time RF2 End Time
1	1022232815	10/22/2013 23:52:09 10/23/2013 02:45:07	10,378	10/22/2013 23:53:10 10/22/2013 23:58:56	10/22/2013 23:58:57 10/23/2013 00:04:06
2	1022232816	10/22/2013 23:52:09 10/23/2013 02:45:14	10,385	10/23/2013 00:04:06 10/22/2013 00:06:45	10/23/2013 00:06:45 10/23/2013 00:16:26
3	1022232817	10/22/2013 23:52:09 10/23/2013 02:44:55	10,366	10/23/2013 00:16:27 10/23/2013 00:19:18	10/23/2013 00:19:18 10/23/2013 00:32:26
4	1022232818	10/22/2013 23:52:09 10/23/2013 02:44:32	10,343	10/23/2013 00:32:26 10/23/2013 00:36:22	10/23/2013 00:36:22 10/23/2013 00:38:46
5	1022232819	10/22/2013 23:52:09 10/23/2013 02:44:37	10,348	10/23/2013 00:38:46 10/23/2013 00:40:48	10/23/2013 00:40:49 10/23/2013 00:51:56
6	1022232820	10/22/2013 23:52:09 10/23/2013 02:44:41	10,352	10/23/2013 00:51:57 10/23/2013 00:54:41	10/23/2013 00:54:41 10/23/2013 01:15:37
7	1022232821	10/22/2013 23:52:09 10/23/2013 02:44:33	10,344	10/23/2013 01:15:37 10/23/2013 01:18:16	10/23/2013 01:18:16 10/23/2013 01:20:07
8	1022232822	10/22/2013 23:52:09 10/23/2013 02:38:48	9,999	10/23/2013 01:20:08 10/23/2013 01:22:57	10/23/2013 01:22:57 10/23/2013 01:39:49

TPC-H Timing Intervals (in seconds)

Duration of stream execution:

Stream ID	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11	Q12
0	179.9	9.5	35.5	37	114.4	5.3	37.5	72.7	203.6	64	21.5	29.9
1	1110.7	52.3	314.1	108.8	874.4	42.3	236.6	640.6	1067.5	422.5	104.7	217.3
2	1138.4	31.8	229.4	240.2	876.2	72.6	342.2	771	1406.7	424.2	145.7	218
3	1125.4	64.3	207.4	221.5	885.3	46.7	315.2	479.8	1116.1	335	110.5	220.2
4	1096.5	60.9	272.9	167	764.6	65.2	318.6	620.3	1239	492.8	102.4	225.7
5	989.0	80.1	304.5	233.4	758.1	52.7	307.5	534	1137.3	408.1	121.3	260.1
6	1067.4	52.4	360.1	218	805	61.1	249.6	673.2	1237.3	646.2	96.9	228.5
7	1001.4	66.2	259.9	205.6	748.6	60.6	265.4	678.1	1061.6	445.3	101.7	166.7
8	947.4	84.3	210.8	329	930.1	90.1	265.5	699.4	1176.4	368.7	123.7	188.4
Minimum	947.4	31.8	207.4	108.8	748.6	42.3	236.6	479.8	1061.6	335	96.9	166.7
Maximum	1138.4	84.3	360.1	329	930.1	90.1	342.2	771	1406.7	646.2	145.7	260.1
Average	1059.5	61.5	269.9	215.4	830.3	61.4	287.6	637.1	1180.2	442.9	113.4	215.6

Stream ID	Q13	Q14	Q15	Q16	Q17	Q18	Q19	Q20	Q21	Q22	RF1	RF2
0	157.1	17.5	26.3	16.8	34.4	140.8	19.3	21.2	121.1	14.1	19.6	21.4
1	1268.2	128.1	123.5	116.1	994	733.6	132.3	235.4	1141.1	312.9	345.7	308.5
2	1337.5	219.3	133.8	147.2	775.7	791.6	93	219.3	497.5	272	159	581.3
3	1303.1	147.7	168.2	120.2	1038.2	1013.6	118.2	301.6	827.5	199.3	170.9	787.8
4	1329.9	138.3	206.1	104.6	965.9	837.2	145.5	140.3	701.3	347.1	235.3	143.6
5	1376.8	132	167.5	168.6	897.3	780.1	151.7	134	1060.9	291.6	122.2	667.8
6	1333.6	106.1	250.6	148.5	811.8	896.9	99.2	185.2	629.8	193.4	164.5	1255.6
7	1323.3	133.1	154.8	134	1071.2	1007.1	74.6	196.4	1005.3	182.3	158.9	111.2
8	1193.4	92.6	114.1	73.6	833.7	794.1	215.4	219.2	762	286	169.7	1011.3
Minimum	1193.4	92.6	114.1	73.6	775.7	733.6	74.6	134	497.5	182.3	122.2	111.2
Maximum	1376.8	219.3	250.6	168.6	1071.2	1013.6	215.4	301.6	1141.1	347.1	345.7	1255.6
Average	1308.2	137.2	164.8	126.6	923.5	856.8	128.7	203.9	828.2	260.6	190.8	608.4

Preface

The Processing Performance Council (TPC) is a non-profit corporation founded to define transaction processing and database benchmarks and to disseminate objective, verifiable TPC performance data to the industry. The TPC Benchmark© H (TPC-H) is a decision support benchmark.

TPC Benchmark© H Overview

The TPC Benchmark© H (TPC-H) 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 may 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 may 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. and developed and engineered in partnership with Sybase an SAP company.

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 C420 M3 server features:

- Intel Xeon processor E5-4600 series processors
- 4-rack unit (RU) rack-mount chassis
- 48 slots for registered DIMMs (RDIMM) and load-reduced DIMMs (LRDDIMMs)
- Drives are installed in configurable (1 or 2) drive bay modules that provide hot-pluggable front-panel access
- Each drive bay module can hold up to eight 2.5 x 0.55 in. (63.5 x 14 mm) SAS or SATA hard disk drives (HDDs) or solid state drives (SSDs), for a total of 16 drives
- 7 PCIe expansion slots
- Integrated quad-port Gigabit Ethernet
- Baseboard management controller (BMC)

Both the measured and priced configurations are same and consist of a Cisco UCS C420 M3 Rack-Mount Server with:

- 4 x Intel Xeon E5-4650 Processor (2.7GHz, 20 MB Cache, 130W)
- 1024 GB of memory
- 1 x MegaRAID 9271CV with 8 internal SAS/SATA ports with Supercap
    - 8 x 600GB 6Gb SAS 10K RPM SFF HDD
    - 8 x Cisco UCS 200 GB SATA 2.5" Enterprise Performance SSD
    - 2 x 8 Drive Backplane W/ Expander For C-Series
- 6 x Cisco UCS Fusion ioDrive2 1205M Adapter (1.2TB Usable)

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 was used for the lineitem and orders tables. The Supporting Files Archive contains the partitioning definitions used.

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

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 and QGEN were used (see sections 2.3 and 4.5)

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.

Due to the lack of availability of QGen 2.16.0, QGen 2.15.0 was used with modifications approved by the TPC for release 2.16.0. These approved modifications were detailed in the FDR of the TPC-H result published by Oracle on June 7, 2013 for the SPARC T5-4 Server.

In addition, the reference dataset for the query substitution parameters could not be verified since the 2.16.0 dataset is not available

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 Q2, Q3, Q10, Q18 and Q21, the “top” function is used to restrict the number of output rows
- In Q7, Q8 and Q9, the “year” function is used to extract part of a date
- The semicolon ';' is used as a command delimiter.

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 isolation level 3.

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 orderkey, o_key, and l_linenumber, l_key. The number of records in the HISTORY table with selected o_key, l_key was also retrieved.

2. The ACID transaction T1 was executed for the Orderkey used in Step 1.

3. The ACID transaction committed.

4. The total price and the extended price were retrieved for the same orderkey used in step 1 and step 2. It was verified that: T1.EXTENDEDPRICE = OLD.EXTENDEDPRICE + ((T1.DELTA) * (OLD.EXTENDEDPRICE/OLD.QUANTITY)), T1.TOTALPRICE = OLD.TOTALPRICE + ((T1.EXTENDEDPRICE-OLD.EXTENDEDPRICE)*(1- DISCOUNT)*(1+TAX)), and that the number of records in the history table with o_key, l_key used in step 1 had increased by 1.

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 random Orderkey. The number of records in the HISTORY table was also retrieved.

2. The ACID transaction was executed for the Orderkey used in step 1.

3. The transaction was rolled back.

4. The total price and the extended price were retrieved for the same orderkey used in step 1 and step 2. It was verified that the extended price and the total price were the same as in step 1. The number of records in the HISTORY table was retrieved again and verified to be the same as in step1.

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(trunc(trunc(L_EXTENDEDPRICE(1-L_DISCOUNT),2)*(1+L_TAX),2))*
for each ORDERS 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. One hundred ACID Transactions were submitted from each of eight execution streams.
3. The consistency of the ORDER and LINEITEM tables was re-verified.

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.

1. 1st session: Started an ACID transaction with a randomly selected O_KEY, L_KEY and DELTA. The transaction was delayed for 10 seconds just prior to the Commit.
2. 2nd session: Started an ACID query for the same O_KEY as in the ACID transaction. The query completed without blocking and did not see any of the uncommitted changes made by the ACID transaction.
3. 1st session: the ACID transaction resumed and successfully completed the Commit.

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. 1st session: Performed the ACID transaction for a random O_KEY, L_KEY and DELTA. The transaction was delayed for 10 seconds just prior to the Rollback.

2. 2nd session: Started an ACID query for the same O_KEY as in the ACID transaction. The query completed without blocking and did not see any of the uncommitted changes made by the ACID transaction.

3. 1st session: the ACID transaction resumed and successfully completed the Rollback.

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.

The database files are distributed across 6 x Cisco UCS Fusion ioDrive2 1205M Adapters (1.2TB Usable) configured as RAID10 volume at the Operating System level

3.5.1 Power Failure

Guarantee the database and committed updates are preserved across the loss of all external power to the SUT for an indefinite time period.

This test was conducted on the qualification database. The following steps were performed:

1. The consistency condition described in section 3.3 was verified.
2. The current count of the total number of records in the HISTORY table was determined giving hist1.
3. A test to run ACID transactions on each of 10 execution streams was started such that each stream executes a different set of transactions.
4. At least 100 ACID transactions were completed from each of the execution streams.
5. While ACID transactions continued to be executed by all execution streams, the system was powered down by removing the power source. Once powered down, the system's enclosure was opened and one of the FusionIO Drive2 Flash Cards was removed.
6. The system was powered back on and rebooted and the database was restarted.

7. Step 2 was repeated giving hist2. It was verified that hist2 – hist1 was greater than or equal to the number of records in the success file. The content of the history file was compared against the content of the execution streams' success files.
8. The consistency condition described in section 3.3 was verified.

3.5.2 Instantaneous Interruption & Memory Failure

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 or across the failure of all or part of memory (loss of contents).

The test in section 3.5.1 was used for the Instantaneous Interruption and the Memory Loss tests.

3.5.3 Loss of Durable Media

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

The test in section 3.5.1 was used for the Loss of Durable Media, which stores the IQ database tables.

Following steps were performed for the Loss of Durable Media, which stores the IQ catalogue data and log file.

1. The consistency condition described in section 3.3 was verified.
2. The current count of the total number of records in the HISTORY table was determined giving hist1.

3. A test to run ACID transactions on each of 10 execution streams was started such that each stream executes a different set of transactions.
4. At least 100 ACID transactions were completed from each of the execution streams.
5. While ACID transactions continued to be executed by all execution streams, one of the disks which stores the IQ catalogue data and log file was removed from the enclosure.
6. It was observed that the test continued without any failure after the disk was removed.
7. All ACID transaction streams were stopped.
8. Step 2 was repeated giving hist2. It was verified that hist2 – hist1 was greater than or equal to the number of records in the success file. The content of the history file was compared against the content of the execution streams' success files.
9. The consistency condition described in section 3.3 was verified

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	30,000,000
Customer	450,000,000
Part	600,000,000
Partsupp	2,400,000,000
Orders	4,500,000,000
Lineitem	18,000,048,306

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 consisted of:

- 1 x MegaRAID 9271CV with 8 internal SAS/SATA ports with Supercap
  - 8 x 600GB 6Gb SAS 10K RPM SFF disk drives
  - 8x Cisco UCS 200 GB SATA 2.5" Enterprise Performance SSD
  - 2 x 8 Drive Backplane W/ Expander For C-Series
- 6 x Cisco UCS Fusion ioDrive2 1205M Adapter

The IQ database files are distributed across six Cisco UCS Fusion ioDrive2 1205M Adapters configured as RAID10 volume at the Operating System. The IQ Temp space is distributed across the Cisco UCS 200 GB SATA 2.5" Enterprise Performance SSDs. The IQ Catalog database and Log files are placed on RAID 1 Volume on which the OS and Sybase IQ installations are stored.

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

	GB	Type	RAID Level	File System	Capacity (GB)	Device	Data
1	600	SAS HDD	RAID1	Ext4	600	/dev/sda	OS, Sybase IQ installation, .IQ Catalog data and log files
2	600	SAS HDD					
3	600	SAS HDD	RAID0	Ext4	600	/dev/sdb	DBGen Flat files
4	600	SAS HDD	RAID0	Ext4	600	/dev/sdd	
5	600	SAS HDD	RAID0	Ext4	600	/dev/sde	
6	600	SAS HDD	RAID0	Ext4	600	/dev/sdf	
7	600	SAS HDD	RAID0	Ext4	600	/dev/sdg	
8	600	SAS HDD	RAID0	Ext4	600	/dev/sdh	
9	200	SATA SSD	RAID0	RAW	1.6 TB	/dev/sdc	IQ Temp space and DBGen Flat Files
10	200	SATA SSD					
11	200	SATA SSD					
12	200	SATA SSD					
13	200	SATA SSD					
14	200	SATA SSD					
15	200	SATA SSD					
16	200	SATA SSD					
17	1200		Software RAID10 using mdadm /dev/md127	RAW	1200	/dev/fioa	3.6 TB for IQ Database tables
18	1200	PCI Flash		RAW	1200	/dev/fiob	
19	1200	PCI Flash		RAW	1200	/dev/fioc	
20	1200	PCI Flash		RAW	1200	/dev/fiod	
21	1200	PCI Flash		RAW	1200	/dev/fioe	
22	1200	PCI Flash		RAW	1200	/dev/fiof	

4.3 Mapping of Database Partitions/Replications

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

The lineitem and orders tables were partitioned by year. The database was not replicated..

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.

The database files are distributed across six Cisco UCS Fusion ioDrive2 1205M Adapters configured as RAID10 volume at the Operating System.

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 2.15.0 was used with no modifications.

Note: DBGEN version 2.15.0 was used due to the lack of availability of DBGen 2.16.0. TPC did not make any modifications between the two versions. Aside from the release number, the two versions are identical.

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 2 hours 25 minutes and 34 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: Database Storage Ratio

Storage Devices	Storage Capacity	Total Storage Capacity	Scale factor	Data Storage Ratio
8 x 600 GB (SAS HDDs)	4,800 GB	13,310 GB	3000	4.44
8 x 200 GB (SATA SSDs)	1,600 GB			
6 x 1.2 TB (PCI Flash)	7,200 GB			

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

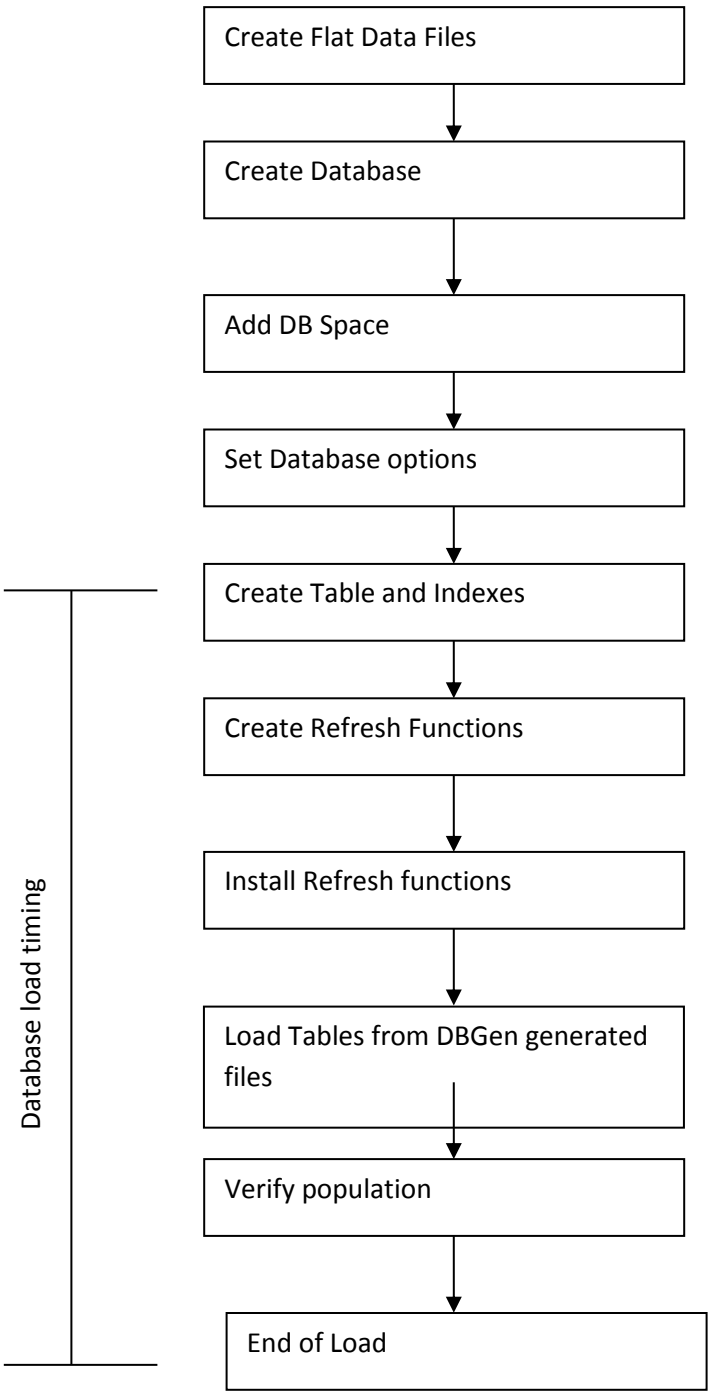

4.9 Qualification Database Configuration

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.

The memory to database size percentage is 34.1%

Clause 5: Performance Metrics and Execution Rules

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 Transaction
2. Stream 00 Execution
3. RF2 Refresh Transaction.

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.

The timing intervals for each query and both refresh functions are given in the Numerical Quantities Summary earlier in the executive summary.

5.3 Number of Streams for The Throughput Test

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

Eight streams were used for the Throughput Test.

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.

The Numerical Quantities Summary section in the Executive Summary of this report contains the start and end times for the query execution streams.

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.

The Numerical Quantities Summary section in the Executive Summary of this report contains the timing intervals for the throughput test (Ts).

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.

The Numerical Quantities Summary section in the Executive Summary of this report contains the start and end times for each refresh function.

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.

The Numerical Quantities Summary section in the Executive Summary of this report contains the the timing intervals for each query and each refresh function for each stream.

5.8 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.

5.9 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.

Two consecutive runs were executed. Following are the performance metrics for both runs:

Table 5.9: Performance Metric

Run	QppH @ 3000GB	QthH @ 3000GB	QphH @ 3000GB
Run 1	289,320.0	183,033.2	230,119.9
Run 2	297,616.9	183,351.0	233,598.7

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.

There was no activity performed between Run1 and Run2.

Clause 6: SUT and Driver Implementation

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.

Two scripts were used. The first one was used to create and load the database, while the second was used to run the Power and Throughput tests. These scripts are in Supporting files.. A C program, semaphore.c, was used for coordination of parallel processes.

6.2 Implementation Specific Layer

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.

The performance tests are executed using the dbisqlc and iqisql utilities. The dbisqlc and iqisql utilities are Sybase-provided utilities which allow SQL statements to be executed against a Sybase IQ database. The dbisqlc and iqisql utilities are invoked from the command line on the SUT. They read input from files containing SQL statements and sends results to stdout. The dbisqlc utility uses information in the.odbc.ini file to connect to the database while iqisql uses information in the interfaces file to connect to the database.
The ACID tests are performed using the dbtest utility. The dbtest utility is a Sybase-provided utility, similar to dbisqlc, but providing additional scripting capabilities. It is invoked from the command-line on the SUT. It reads input from files that determine how and what to execute. It uses information in the.odbc.ini file to connect to the database**.**

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

7.1 Hardware and Software Pricing

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 Three 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.

The pricing details for this disclosure are contained in the executive summary pages.

7.3 Availability Dates

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.

Hardware is available today except the 200GB 2.5 SATA Enterprise performance SSD(UCS-SD200G0KS2-EP). 200GB 2.5 SATA Enterprise performance SSD is available by December 30, 2013. Sybase IQ SP02 is available December 30, 2013.

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.

Table 8.0: Supporting File Index

Clause	Description	Archive File
Clause 1	Operating System and Database settings	Supporting_file/Clause1
Clause 2	Qualification Queries and Output	Supporting_file/Clause2
Clause 3	ACID scripts and output	Supporting_file/Clause3
Clause 4	Database load scripts	Supporting_file/Clause4
Clause 5	Queries and output	Supporting_file/Clause5
Clause 6	Implementation code for measured runs	Supporting_file/Clause6
Clause 8	RFs source and parameters	Supporting_file/Clause8

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.

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

This benchmark was audited by:

Francois Raab
InfoSizing, Inc.
531 Crystal Hills Blvd
Manitou Springs, CO 80829
Phone: 719-473-7555.

InfoSizing
The Right Metric For Sizing IT

TPC Transaction Processing Performance Council
Certified Auditors

Raghunath Nambiar
Cisco Systems Inc.
3800 Zanker Road
San Jose, CA 95134

October 31, 2013

I verified the TPC Benchmark H (TPC-H™ v2.16.0) performance of the following configuration:

Platform:	Cisco UCS C420 M3
Operating System:	Red Hat Enterprise Linux Server Release 6.4
Database Manager:	Sybase IQ Version Service Pack 2
Other Software:	n/a

The results were:

Performance Metric	**230,119.9 QphH@3000GB**
TPC-H Power	289,320.0
TPC-H Throughput	183,033.2
Database Load Time	02:25:34

Server	**Cisco UCS C420 M3**		
CPUs	4 x Intel Xeon E5-4650 Processor (2.7 GHz, 20MB Cache)		
Memory	1024 GB		
Disks	***Qty***	***Size***	***Type***
	8	600 GB	SAS 10K rpm SFF HDD
	8	200 GB	SATA 2.5" SSD
	6	1205 GB	PCI Flash

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 3,000GB and populated accordingly
- The compliance of the database auxiliary data structures was verified
- The database load time was correctly measured and reported

521 Crystal Hills Blvd • Manitou Springs, CO 80829 • 719 473 7555 • www.sizing.com

- 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 8 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:

Version 2.16.0 of the DBGen package was not available at the time of testing. DBGen version 2.15.0 was used instead. The TPC did not make any modifications between the two versions. Aside from the release number, the two versions are identical. QGen 2.15.0 was used with the modifications approved by the TPC for release 2.16.0. These approved modifications were detailed in the FDR of the TPC-H result published by Oracle on June 7, 2013 for the SPARC T5-4 Server. In addition, the reference dataset for the query substitution parameters could not be verified since the 2.16.0 dataset is not available.

Respectfully Yours,

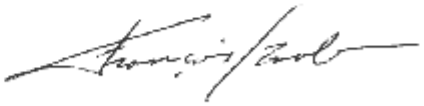

François Raab, President

521 Cresta Hills Blvd • Manitou Springs, CO 80829 • 719 473 7555 • www.sut.com

Appendix A

For:

Quotation for Software and Support

Company Cisco Systems Inc.
Contact Raghunath Nambiar
Phone (408) 527-3052
Fax rnambiar@cisco.com
Address 275 East Tasman Drive, San Jose, CA 95134

SYBASE Sales Rep: Ronald Keusch
Phone: 212-596-1170
Fax:

Sybase, an SAP Company, 1114 Avenue of the Americas, New York, NY 10036

	Catalogue Number	Product Description	License Type	Machine	P/S	List Price Per Unit	Quantity	Price	Discount	Extended Price
1	13545	Sybase IQ Single App Svr, per cpu core	CP	Linux x86-64	P	2,595	32	83,040	12,456	70,584
3	98480	3 yr support Single App Svr, per cpu core				1,713	32	54,816	8,222	46,594
4		Discounts:								
5		15.00%								
6										
7										
8										
9										
10										
11										
12										

Quote Date: ***10/23/13***
Valid thru: ***12/31/13***

Total (USD) 117,178.00

Licence + 3 year support

Payment terms : Net 30 Days