

TPC EXPRESS BENCHMARK TM HS
(TPCx-HS)

Standard Specification

Version 2.0.3

Mar 5, 2018

Transaction Processing Performance Council (TPC)

www.tpc.org

info@tpc.org

© 2018 Transaction Processing Performance Council

All Rights Reserved

Legal Notice

The TPC reserves all right, title, and interest to this document and associated source code as provided under U.S. and international laws, including without limitation all patent and trademark rights therein.

Permission to copy without fee all or part of this document is granted provided that the TPC copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the Transaction Processing Performance Council. To copy otherwise requires specific permission.

No Warranty

TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, THE INFORMATION CONTAINED HEREIN IS PROVIDED "AS IS" AND WITH ALL FAULTS, AND THE AUTHORS AND DEVELOPERS OF THE WORK HEREBY DISCLAIM ALL OTHER WARRANTIES AND CONDITIONS, EITHER EXPRESS, IMPLIED OR STATUTORY, INCLUDING, BUT NOT LIMITED TO, ANY (IF ANY) IMPLIED WARRANTIES, DUTIES OR CONDITIONS OF MERCHANTABILITY, OF FITNESS FOR A PARTICULAR PURPOSE, OF ACCURACY OR COMPLETENESS OF RESPONSES, OF RESULTS, OF WORKMANLIKE EFFORT, OF LACK OF VIRUSES, AND OF LACK OF NEGLIGENCE. ALSO, THERE IS NO WARRANTY OR CONDITION OF TITLE, QUIET ENJOYMENT, QUIET POSSESSION, CORRESPONDENCE TO DESCRIPTION OR NON-INFRINGEMENT WITH REGARD TO THE WORK.

IN NO EVENT WILL ANY AUTHOR OR DEVELOPER OF THE WORK BE LIABLE TO ANY OTHER PARTY FOR ANY DAMAGES, INCLUDING BUT NOT LIMITED TO THE COST OF PROCURING SUBSTITUTE GOODS OR SERVICES, LOST PROFITS, LOSS OF USE, LOSS OF DATA, OR ANY INCIDENTAL, CONSEQUENTIAL, DIRECT, INDIRECT, OR SPECIAL DAMAGES WHETHER UNDER CONTRACT, TORT, WARRANTY, OR OTHERWISE, ARISING IN ANY WAY OUT OF THIS OR ANY OTHER AGREEMENT RELATING TO THE WORK, WHETHER OR NOT SUCH AUTHOR OR DEVELOPER HAD ADVANCE NOTICE OF THE POSSIBILITY OF SUCH DAMAGES.

Trademarks

TPC Benchmark and TPC Express are trademarks of the Transaction Processing Performance Council.

Acknowledgments

Developing a TPC benchmark for a new environment requires a huge effort to conceptualize research, specify, review, prototype, and verify the benchmark. The TPC acknowledges the work and contributions of the TPCx-HS subcommittee member companies in developing the TPCx-HS specification.

The TPCx-HS subcommittee would like to acknowledge the contributions made by the many members to the development of the benchmark specification. It has taken the dedicated efforts of people across many companies, often in addition to their regular duties. The list of contributors includes Andrew Bond, Andrew Masland, Avik Dey, Brian Caufield, Chaitanya Baru, Da Qi Ren, Dileep Kumar, Jamie Reding, John Fowler, John Poelman, Karthik Kulkarni, Matthew Emmerton, Meikel Poess, Mike Brey, Mike Crocker, Paul Cao, Raghunath Nambiar, Reza Taheri, Simon Harris, Tariq Magdon-Ismael, Wayne Smith, Yanpei Chen, David Grimes, Chinmayi Narasimhadevara and Dave Jaffe.

Document Revision History

Table 1: lists the document revision history.

Table 1: Document Revision History

Date	Version	Description
06/17/2014	1.1.0	Incorporated new comments
09/06/2014	1.1.2	Added file Benchmark_Parameters.sh to provide number of Mappers and Reducers. Updated ReadMe File for the above file. Updated Script to print out "Performance Metric (HSph@SF) Report". Updated the source code for 3-way replication. The source code also includes files for HSSort with MapReduce (MR2) for experimental future use.
11/13/2014	1.2.0	Updated requirements for independent audit, requirements for auditing energy audit, substitution. Removed the performance requirements for new kit revisions.
02/19/2015	1.3.0	Updated Full Disclosure requirement. Sample Executive Summary is included. Fixed cross references and alignments
03/24/2016	1.3.1	Updated TPC membership.
04/19/2016	1.4.0	Updated Executive Summary to include the Total Rack Units (RU) occupied by the SUT.
07/12/2016	1.4.1	Document Revision History Table formatting changes.
11/03/2016	1.4.2	Updated to align with Pricing v2 and allow for LCS publications. Added support for non-JBOD storage.
04/20/2017	2.0.0	Updated to include Spark. Updated Acknowledgments.
06/07/2017	2.0.1	Updated version number to match kit.
09/20/2017	2.0.2	Updated Section 8.1 to include the Pre-Publication Board.
03/05/2018	2.0.3	Include source files in kit.

TPC Membership

TPC membership as of March 2018.

Full Members

Associate Members

		
---	---	--

TABLE OF CONTENTS

ACKNOWLEDGMENTS	3
DOCUMENT REVISION HISTORY	3
TPC MEMBERSHIP	4
FULL MEMBERS	4
ASSOCIATE MEMBERS	4
CLAUSE 1: INTRODUCTION	7
1.1 PREAMBLE	7
1.2 TPCx-HS KIT AND LICENSING	7
1.3 GENERAL IMPLEMENTATION GUIDELINES	7
1.4 GENERAL MEASUREMENT GUIDELINES	8
CLAUSE 2: WORKLOAD AND EXECUTION	9
2.1 BENCHMARK KIT	9
2.1.1 <i>Kit Contents</i>	9
2.1.2 <i>Kit Usage</i>	9
2.1.3 <i>Kit Modification</i>	9
2.1.3.1 <i>Minor Shell Script Modifications</i>	9
2.1.3.2 <i>Major Shell Script Modifications</i>	9
2.1.3.3 <i>Java Code Modifications</i>	10
2.1.4 <i>Future Kit Releases</i>	10
2.2 BENCHMARK WORKLOAD	10
2.3 BENCHMARK EXECUTION	10
2.3.1 DATA REPLICATION	12
2.4 CONFIGURATION AND TUNING	12
CLAUSE 3: SYSTEM UNDER TEST AND DRIVER	13
3.1 SYSTEM UNDER TEST	13
CLAUSE 4: SCALE FACTORS AND METRICS	15
4.1 SCALE FACTOR	15
4.2 METRICS	15
4.3 PERFORMANCE METRIC (HSPH@SF)	15
4.4 PRICE PERFORMANCE METRIC (\$/HSPH@SF)	15
4.5 SYSTEM AVAILABILITY DATE	16
4.6 METRIC COMPARISON	16
4.7 REQUIRED REPORTING COMPONENTS	16
CLAUSE 5: PRICING	17
5.1 PRICING METHODOLOGY AND MODEL	17
5.2 PRICED SYSTEM	17
5.3 ADDITIONAL OPERATIONAL COMPONENTS	17
5.4 SOFTWARE	18
5.5 ALLOWABLE SUBSTITUTIONS	18
CLAUSE 6: ENERGY	19

6.1 GENERAL MEASUREMENT GUIDELINES	19
6.3 TPCx-HS ENERGY METRIC	21
6.4 REPORTING TPC-ENERGY OPTIONAL METRICS	21
CLAUSE 7: FULL DISCLOSURE	22
7.1 REPORTING REQUIREMENTS	22
7.2 FORMAT GUIDELINES	22
7.3 FULL DISCLOSURE REPORT	22
7.4 GENERAL ITEMS	22
7.5 WORKLOAD RELATED ITEMS	25
7.6 SUT RELATED ITEMS	26
7.7 SCALE FACTORS AND METRICS	26
7.8 CLAUSE 8 – AUDIT RELATED ITEMS	26
7.9 EXECUTIVE SUMMARY	27
7.10 PAGE LAYOUT	27
7.11 IMPLEMENTATION OVERVIEW	27
7.12 PRICING SPREADSHEET	29
7.13 NUMERICAL QUANTITIES SUMMARY	29
7.14 TPCx-HS RUN REPORT	30
7.15 AVAILABILITY OF THE FULL DISCLOSURE REPORT	30
7.16 REVISIONS TO THE FULL DISCLOSURE REPORT	31
CLAUSE 8: AUDIT	32
8.1 GENERAL RULES	32
8.2 AUDIT CHECK LIST	33
8.2.1 <i>CLAUSE 2 Related Items</i>	33
8.2.2 <i>CLAUSE 3 Related Items</i>	33
8.2.3 <i>CLAUSE 4 Related Items</i>	33
8.2.4 <i>CLAUSE 5 Related Items</i>	33
8.2.5 <i>CLAUSE 6 Related Items</i>	33
CLAUSE 9: SAMPLE EXECUTIVE SUMMARY	34

CLAUSE 1: INTRODUCTION

1.1 Preamble

Big Data technologies like Hadoop and Spark have become an important part of the enterprise IT ecosystem. The TPC Express Benchmark™HS (TPCx-HS) was developed to provide an objective measure of hardware, operating system and commercial Apache Hadoop File System API compatible software distributions, and to provide the industry with verifiable performance, price-performance and availability metrics. The benchmark models a continuous system availability of 24 hours a day, 7 days a week.

Even though the modeled application is simple, the results are highly relevant to hardware and software dealing with Big Data systems in general. **TPCx-HS** stresses both the hardware and software stack including the execution engine (MapReduce or Spark) and Hadoop Filesystem API compatible layers. This workload can be used to assess a broad range of system topologies and implementation of Hadoop/Spark clusters. The TPCx-HS benchmark can be used to assess a broad range of system topologies and implementation methodologies in a technically rigorous and directly comparable, in a vendor-neutral manner.

1.2 TPCx-HS Kit and Licensing

The **TPCx-HS** kit is available from the TPC (See www.tpc.org/tpcx-hs for more information). User must sign-up and agree to the **TPCx-HS** User Licensing Agreement (ULA) to download the kit. Re-distribution of the kit is prohibited. All related work (such as collaterals, papers, derivatives) must acknowledge the TPC and include TPCx-HS copyright. The TPCx-H Kit includes: TPCx-HS Specification document (this document), TPCx-HS Users Guide documentation, Shell scripts to set up the benchmark environment, Java code to execute the benchmark load. See CLAUSE 2: WORKLOAD AND EXECUTION

1.3 General Implementation Guidelines

The purpose of TPC benchmarks are to provide relevant, objective performance data to industry users. To achieve that purpose, TPC benchmark specifications require that benchmark tests be implemented with systems, products, technologies and pricing that:

- Are generally available to users;
- Are relevant to the market segment that the individual TPC benchmark models or represents (for example, TPCx-HS models and represents Hadoop MapReduce/Spark execution engines and Hadoop Filesystem API compatible systems);
- Would plausibly be implemented by a significant number of users in the market segment the benchmark models or represents.

The system under test (SUT)'s fault-recovery attributes must transparently and automatically allow recovery from partial failures. Partial failure is defined as a failure that is limited to a given component (for example, machines or hardware nodes) or phase of processing (for example, map or reduce tasks). In case of a partial failure, the implementation will only allow restarting of the impacted portion of the work in progress. The work that has already successfully completed in earlier phases of processing in other non-failed components of the SUT cannot be restarted. For example, if a node fails, only the tasks that were scheduled and ran on that node can be restarted.

The use of new systems, products, technologies (hardware or software) and pricing is encouraged so long as they meet the requirements above. Specifically prohibited are benchmark systems, products, technologies or pricing (hereafter referred to as "implementations") whose primary purpose is performance optimization of the TPC benchmark results without any corresponding applicability to real-world applications and environments. In other words, all "benchmark special" implementations that improve benchmark results but not real-world performance or pricing, are prohibited.

The following characteristics shall be used as a guide to judge whether a particular implementation is a benchmark special. It is not required that each point below be met, but that the cumulative weight of the evidence be considered to identify an unacceptable implementation. Absolute certainty or certainty beyond a reasonable doubt is not required to make a judgment on this complex issue. The question that must be answered is: "Based on the available evidence, does the clear preponderance (the greater share or weight) of evidence indicate that this implementation is a benchmark special?"

The following characteristics shall be used to judge whether a particular implementation is a benchmark special:

- Is the implementation generally available, externally documented, and supported?
- Does the implementation have significant restrictions on its use or applicability that limits its use beyond the TPCx-HS benchmark?
- Is the implementation or part of the implementation poorly integrated into the larger product?
- Does the implementation take special advantage of the limited nature of the TPCx-HS benchmark in a manner that would not be generally applicable to the environment the benchmark represents?
- Is the use of the implementation discouraged by the vendor? (This includes failing to promote the implementation in a manner similar to other products and technologies.)
- Does the implementation require uncommon sophistication on the part of the end-user, programmer, or system administrator?
- Is the implementation (including beta) being purchased or used for applications in the market area the benchmark represents? How many sites implemented it? How many end-users benefit from it? If the implementation is not currently being purchased or used, is there any evidence to indicate that it will be purchased or used by a significant number of end-user sites?

The rules for pricing are included in the TPC Pricing Specification located at www.tpc.org.

The rules for energy measurement are included in the TPC Energy Specification located at www.tpc.org.

1.4 General Measurement Guidelines

TPC benchmark results are expected to be accurate representations of system performance. Therefore, there are certain guidelines that are expected to be followed when measuring those results. The approach or methodology to be used in the measurements are either explicitly described in the specification or left to the discretion of the test sponsor. When not described in the specification, the methodologies and approaches used must meet the following requirements:

- The approach is an accepted engineering practice or standard;
- The approach does not enhance the result;
- Equipment used in measuring the results is calibrated according to established quality standards;
- Fidelity and candor is maintained in reporting any anomalies in the results, even if not specified in the TPC benchmark requirements.

CLAUSE 2: WORKLOAD AND EXECUTION

This clause defines workload and execution.

2.1 Benchmark Kit

2.1.1 Kit Contents

The TPCx-HS kit contains the following:

- TPCx-HS Specification document
- TPCx-HS Users Guide documentation
- Shell scripts to set up the benchmark environment
- Java code to execute the benchmark load

2.1.2 Kit Usage

To submit a compliant TPCx-HS benchmark result, the test sponsor is required to use the TPCx-HS kit as provided except for modifications explicitly listed in Clause 2.1.3

The kit must be used as outlined in the TPCx-HS Users Guide.

The output of the TPCx-HS kit is called the **run report** which includes the following:

- Version number of TPCx-HS kit
- Checksum for HSGen, HSDataCheck, HSSort and HSValidate Programs
- Output from HSGen, HSDataCheck (presort and post sort), HSSort and HSValidate Programs
- Cardinality of the test dataset presort and post sort
- Verification of data replication

If there is a conflict between the TPCx-HS specification and the TPC provided code, the TPC provided code prevails.

2.1.3 Kit Modification

2.1.3.1 Minor Shell Script Modifications

Minor modifications to the provided shell scripts in the TPCx-HS kit to facilitate operating system differences are allowed without TPC approval.

The following changes are considered minor modifications:

- Shell script changes necessary for the kit scripts to execute on a particular operating system as long as the changes do not alter the execution logic of the script

2.1.3.2 Major Shell Script Modifications

All changes to the scripts included in the kit other than those allowed in Clause

2.1.3 Kit Modification

2.1.3.1 Minor Shell Script Modifications would be considered major modifications. Major modifications must be approved by the TPC prior to being used in a benchmark submission.

It will be the judgment of the TPC members reviewing the submission or the TPCx-HS certified auditor (if being used) as to whether scripting changes are considered minor or major. If the test sponsor has any doubts they are encouraged to have the changes approved by the TPC prior to being used in a submission.

2.1.3.3 Java Code Modifications

No modifications are allowed to the java code provided in the TPCx-HS kit.

2.1.4 Future Kit Releases

The TPC will release future TPCx-HS benchmark kits at its discretion to fix bugs or add features. When a new kit version is released the TPC will release a timetable regarding the last date a benchmark submission can be made using the previous kit version. After this date only submissions using the new kit version will be considered, and submissions using the previous kit version will immediately be found non-compliant.

If the test sponsor would like new scripts or existing script changes to be included in a future release of the TPCx-HS benchmark kit, then the test sponsor can donate the scripts or script code changes to the TPC and work with the TPC to get them included in the next release.

If a test sponsor would like to see changes made to the java code of the kit, then the changes should be provided to the TPC for potential inclusion in the next release of the TPCx-HS benchmark kit.

2.2 Benchmark Workload

The TPC Benchmark™HS (TPCx-HS) benchmark workload consists of the following modules:

- **HSGen** is a program to generate the data at a particular **Scale Factor**. (See CLAUSE 4: SCALE FACTORS AND METRICS). HSGen is based on TeraGen
- **HSDDataCheck** is a program to check the compliance of the dataset and replication.
- **HSSort** is a program to sort the data into a total order. HSSort is based on TeraSort [1]
- **HSValidate** is a program that validates the output is sorted. HSValidate is based on TeraValidate [1]

[1] http://hadoop.apache.org/docs/current/api/org/apache/hadoop/examples/terasort/package-summary.html#package_description

2.3 Benchmark Execution

A **valid run** will consist of five separate phases run sequentially. These phases may not overlap in their execution times. For example, the start of Phase 2 may not begin until Phase 1 is complete and the start of Phase 3 may not begin until Phase 2 is complete etc. All phases are initiated by the <TPCx-HS-master> script, which can be executed from any of the nodes in the SUT. The phases are listed below:

1. Generation of input data via HSGen. The data generated must be written on a Durable Medium.
2. Dataset (See CLAUSE 4: SCALE FACTORS AND METRICS) verification via HSDDataCheck. The program is to verify the cardinality, size and replication factor of the generated data. If the HSDDataCheck program reports that the data replication factor is less than 3, then the test sponsor must provide additional documentation (see Clause 2.3.1); otherwise, the run is considered invalid.
3. Running the sort using HSSort on the input data. This phase samples the input data and sorts the data. The sorted data must be written on a Durable Medium.

4. Dataset (See CLAUSE 4: SCALE FACTORS AND METRICS) verification via HSDataCheck. The program is to verify the cardinality, size and replication factor of the sorted data. If the HSDataCheck program reports that the data replication factor is less than 3, then the test sponsor must provide additional documentation (see Clause 2.3.1); otherwise, the run is considered invalid.
5. Validating the sorted output data via HSValidate. HSValidate validates the sorted data. If the HSValidate program reports that the HSSort did not generate the correct sort order, then the run is considered invalid.


The elapsed time for each phase must be reported.

The benchmark test consists of two runs, Run1 and Run2, which must follow the Run Phases above. No activities except file system cleanup are allowed between Run 1 and Run2. The total elapsed time for the run in seconds (T) is used for the TPCx-HS Performance Metric calculation. The **performance run** is defined as the run with the lower TPCx-HS Performance Metric. The **repeatability run** is defined as the run with the higher TPCx-HS Performance Metric. The **reported performance metric** is the TPCx-HS Performance Metric for the **performance run**.

No configuration or tuning changes are allowed between the two runs.

The benchmark execution phases are listed below and illustrated in Figure 1:

Figure 1 TPCx-HS Execution Phases



Comment: No part of the SUT may be rebooted or restarted during or between the runs. If there is a non-recoverable error reported by any of the applications, operating system, or hardware in any of the five phases or between Run 1 and Run 2, the run is considered invalid. If a recoverable error is detected in any

of the phases, and is automatically dealt with or corrected by the applications, operating system, or hardware then the run is considered valid provided the run meets all other requirements. However, manual intervention by the test sponsor is not allowed. If the recoverable error requires manual intervention to deal with or correct then the run is considered invalid.

2.3.1 Data Replication

Either a minimum of three-way data replication must be maintained for the datasets generated by HSGen and HSSort throughout the run, or the test sponsor must provide a description of the data redundancy approach describing both hardware and software used to achieve the data redundancy and explain why it is equivalent to or better than the data redundancy provided by traditional local-JBOD storage and replication factor of three. For Licensed Compute Services where visibility into the hardware and software used to implement the durable medium is limited, the test sponsor must provide a description of the Licensed Compute Service Service Level Agreement and explain why it is equivalent to or better than the data redundancy provided by traditional local-JBOD storage and replication factor of three.

Comment: The intent is to allow test sponsors to publish TPCx-HS results on alternative storage media types besides traditional local-JBOD. The alternative hardware or service must prevent data loss in the event of a permanent irrecoverable failure of any single durable medium containing dataset data generated by HSGen or HSSort. The distributed file system must be HDFS or an implementation of Hadoop File System compatible APIs (see Clause 7.4.7).

2.4 Configuration and Tuning

The SUT cannot be reconfigured, changed, or re-tuned by the test sponsor during or between any of the five phases or between Run 1 and Run 2. Any manual tunings to the SUT must be performed before the beginning of Phase 1 of Run 1, and must be fully disclosed. Automated changes and tuning performed between any of the phases are allowed. Any changes to default tunings or parameters of the applications, operating systems, or hardware of the SUT must be disclosed.

CLAUSE 3: SYSTEM UNDER TEST AND DRIVER

This clause defines the System Under Test (SUT) and the benchmark driver.

3.1 System Under Test

The tested and reported configuration is composed of those hardware and software components that are employed in the performance test and whose cost and performance are described by the benchmark metrics. Specifically, the SUT consists of (See Figure 2):

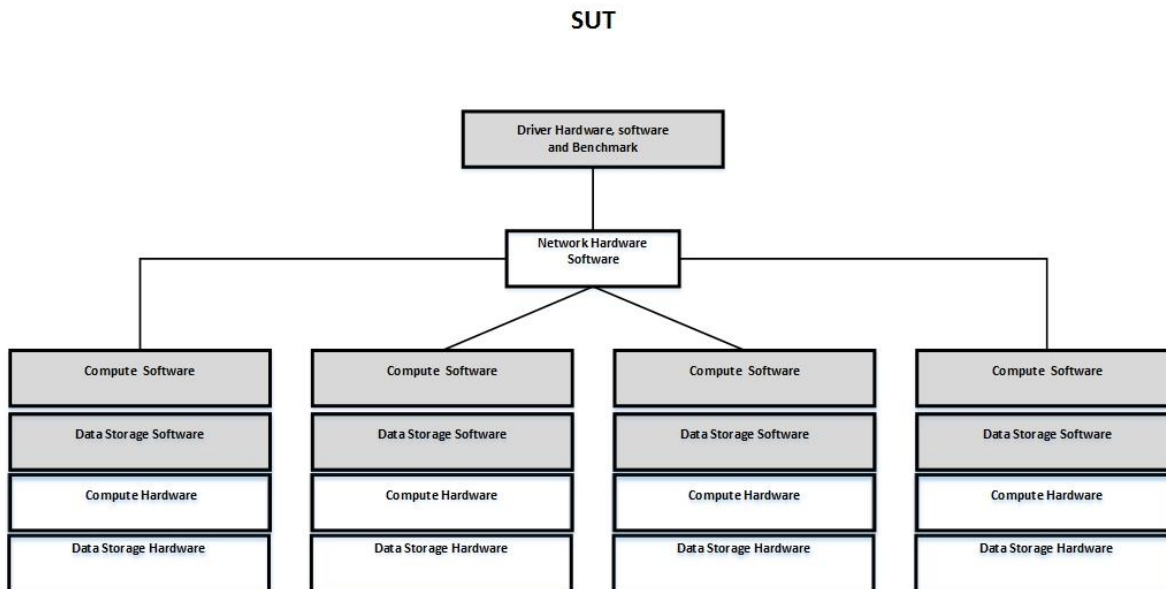
- Devices, for example, compute devices and/or data storage devices, including hardware and software components,
- Any hardware and software devices of all networks required to connect and support the SUT systems,
- Each compute device includes a benchmark specific software layer, the benchmark implementation, and other commercially available software products, and
- One benchmark specific driver, the benchmark driver. The driver may reside on one of the compute devices or on a separate system. In case the driver resides on a separate compute device, this device is part of the SUT.

Comment: Except for the benchmark implementation and the benchmark driver, all SUT components must be commercially available software or hardware products.

Comment: Any hardware or software component may be part of a **Licensed Compute Service** as defined in the TPC Pricing Specification.

Comment: A hardware component may be a physical or virtual device.

Figure 2 System Under Test (SUT)



Comment: The source code of any non-commercially available components used to implement the SUT (such as scripts to configure the system, set tunables, etc.) must be disclosed.

Comment: The driver presents the workload to the SUT. The driver is a logical entity that can be implemented using one or more programs, processes, or systems.

CLAUSE 4: SCALE FACTORS AND METRICS

This clause defines Scale Factor and Metrics.

4.1 Scale Factor

The TPCx-HS follows a stepped size model. **Scale factor (SF)** used for the test dataset must be chosen from the set of fixed Scale Factors defined as follows:

1TB, 3TB, 10TB, 30TB, 100TB, 300TB, 1000TB, 3000TB, 10000TB.

The corresponding number of records are as follows:

10B, 30B, 100B, 300B, 1000B, 3000B, 10000B, 30000B, 100000B, where each record is 100 bytes generated by HSGen.

Comment: The TPC will continuously evaluate adding larger Scale Factors and retiring smaller Scale Factors based on industry trends.

4.2 Metrics

4.2.1 TPCx-HS defines the following primary **metrics**:

1. HSph@SF, the Performance Metric, reflecting the TPCx-HS throughput; where **SF** is the Scale Factor (see Clause 4.1 Scale Factor);
2. \$/HSph@SF, the Price-Performance metric;
3. System availability date (see CLAUSE 5: PRICING);
4. When TPC-Energy option is chosen for reporting, the TPCx-HS energy metric reports the power per performance and is expressed as Watts/HSph@SF (see Clause 6.3 TPCx-HS Energy Metric for additional requirements).

4.3 Performance Metric (HSph@SF)

The **performance metric** of the benchmark is HSph@SF, the effective sort throughput of the benchmarked configuration:

$$HSph @ SF = \frac{SF}{T / 3600}$$

Where:

- SF is the Scale Factor (see Clause 4.1 Scale Factor),
- T is the total elapsed time for the run in seconds

Comment: The **performance run** is defined as the run with the lower TPCx-HS Performance Metric (see Clause 2.3 Benchmark Execution).

4.4 Price Performance Metric (\$/HSph@SF)

The price-performance metric for the benchmark is defined as:

$$\$/HSph @ SF = \frac{P}{HSph @ SF}$$

Where:

P is the total cost of ownership of the SUT.

If a benchmark configuration is priced in a currency other than US dollars, the units of the price-performance metrics must be adjusted to employ the appropriate currency.

4.5 System Availability Date

The System Availability Date is defined in the TPC Pricing Specification.

4.6 Metric Comparison

A TPCx-HS Result is only comparable with other TPCx-HS Results of the same **Scale Factor** (see CLAUSE 4: SCALE FACTORS AND METRICS).

- Results at the different Scale Factors are not comparable, due to the substantially different computational challenges found at different data volumes. Similarly, the system price/performance may not scale down linearly with a decrease in dataset size due to configuration changes required by changes in dataset size.
- If results measured against different Scale Factors appear in a printed or electronic communication, then each reference to a result or metric must clearly indicate the Scale Factors against which it was obtained. In particular, all textual references to TPCx-HS metrics (performance or price/performance) appearing must be expressed in the form that includes the Scale Factor as an integral part of the metric's name; i.e. including the "@SF" suffix. This applies to metrics quoted in text or tables as well as those used to annotate charts or graphs. If metrics are presented in graphical form, then the Scale Factor on which metric is based must be immediately discernible either by appropriate axis labeling or data point labeling.
- In addition, the results must be accompanied by a disclaimer stating: "The TPC believes that comparisons of TPCx-HS results measured against different Scale Factors are misleading and discourages such comparisons".

4.7 Required Reporting Components

To be compliant with the TPCx-HS standard and TPC Policies, the URL to the benchmark result and Availability Date of the complete configuration must be included for all public references (See CLAUSE 8: AUDIT).

CLAUSE 5: PRICING

This section defines the components, functional requirements of what is priced, and what substitutions are allowed. Rules for pricing the **Priced Configuration** and associated software and maintenance are included in the TPC Pricing Specification located at www.tpc.org.

5.1 Pricing Methodology and Model

- 5.1.1 The **Default 3-Year Pricing Methodology** (as defined in the TPC Pricing Specification) must be used to calculate the price and the price/performance result of the TPCx-HS benchmark.
- 5.1.2 The **Pricing Model 1 – Default Pricing Model** (as defined in the TPC Pricing Specification) is the only pricing model allowed in a TPCx-HS result.

5.2 Priced System

The system to be priced shall include the hardware and software components present in the System Under Test (SUT), a communication interface that can support user interface devices, additional operational components configured on the test system, and maintenance on all of the above

Calculation of the priced configuration consists of:

- Price of the SUT as tested and defined in CLAUSE 3: SYSTEM UNDER TEST AND DRIVER
- Price of a communication interface capable of supporting the required number of user interface devices defined in Clause 5.3
- Price of additional products (software or hardware) required for customary operation, administration and maintenance of the SUT for a period of 3 years
- Price of all products required to create, execute, administer, and maintain the executables or necessary to create and populate the test environment.

Specifically excluded from the priced configuration calculation are:

- End-user communication devices and related cables, connectors, and switches;
- Equipment and tools used exclusively in the production of the full disclosure report;

5.3 Additional Operational Components

Additional products included on a customer installed configuration are also to be included in the priced configuration if explicitly required for the operation, administration, or maintenance, of the priced configuration. Examples of such products are:

- Operator console
- User interface terminal
- CD drive
- Software if required for initial load or maintenance updates
- All cables used to connect components of the SUT (except as noted in section 5.2 Priced System)

5.4 Software

All software licenses must be priced.

5.5 Allowable Substitutions

Substitution is defined as a deliberate act to replace components of the Priced Configuration by the test sponsor as a result of failing the availability requirements of the TPC Pricing Specification or when the part number for a component changes.

Comment: Corrections or "fixes" to components of the Priced Configuration are often required during the life of products. These changes are not considered Substitutions so long as the part number of the priced component does not change. Suppliers of hardware and software may update the components of the Priced Configuration, but these updates must not negatively impact the **reported performance metric** or numerical quantities more than two percent.

The following are not considered substitutions:

- software patches to resolve a security vulnerability
- silicon revision to correct errors
- new supplier of functionally equivalent components (for example, memory chips, disk drives etc.)

Some hardware components of the Priced Configuration may be substituted after the test sponsor has demonstrated to the auditor's satisfaction that the substituting components do not negatively impact the reported performance metric or numerical quantities. All substitutions must be reported in the FDR and noted in the auditor's attestation letter. The following hardware components may be substituted:

- Durable Medium (for example disk drives) and Cables

Comment: Durable Medium is defined as a data storage medium that is inherently non-volatile such as a magnetic disk or tape.

Comment: If any hardware component is substituted then the result must be audited by a TPC Certified Auditor or the TPCx-HS Pre-Publication Board (see CLAUSE 8: AUDIT).

CLAUSE 6: ENERGY

This section contains the rules and methodologies for measuring and reporting energy metric in TPCx-HS benchmarks. This includes the energy consumption of system components associated with typical business information technology environments, which are characterized by:


- Energy consumption of servers
- Energy consumption of applications
- Energy consumption of other items that consume power and are required by the benchmark specification as components of the System Under Test (SUT).
- Rules for the Energy configuration and associated software and maintenance are included in the TPC-Energy Specification located at www.tpc.org.

6.1 General Measurement Guidelines

The TPCx-HS energy results are expected to be accurate representations of system performance and energy consumption. Therefore, there are certain requirements which must be followed. The approach and methodology are explicitly detailed in this specification and the TPC Benchmark Standards, as defined in TPC-Energy.

When TPC-Energy option is chosen for reporting, power consumption is measured for Run 1 and Run 2. The energy metric reported is calculated for the **reported run**.

Figure 3 Power Measurement Interval



TPCx-HS Energy metric reports the power per performance and is expressed as Watts/HSph@SF.

6.2.1 Energy Calculation

The computation of the total energy consumption for the **performance run** must be disclosed. If the energy of the entire Priced Configuration is not derived from direct measurements, the methods for deriving the energy for components which were not measured must be disclosed. The average power consumption for the **performance run** must be disclosed. The TPC-Energy Primary Metric must be disclosed, including the calculation that is used to derive it.

If the optional TPC-Energy Secondary Metrics are reported, the components which are included in each subsystem must be identified. This can be achieved with separate lists to be included in the FDR or with a specific designation in the price spreadsheet (see Clause 7.12). Every component that consumes energy must be included in exactly one subsystem. For each defined subsystem, the calculations defined for the TPC-Energy Secondary Metrics must be reported, using the **reported performance metric** of the entire SUT and the energy consumption for each subsystem under report.

$$P = \sum_{1 \leq i \leq m} p_i$$

$$E = \int_0^T P(t) dt$$

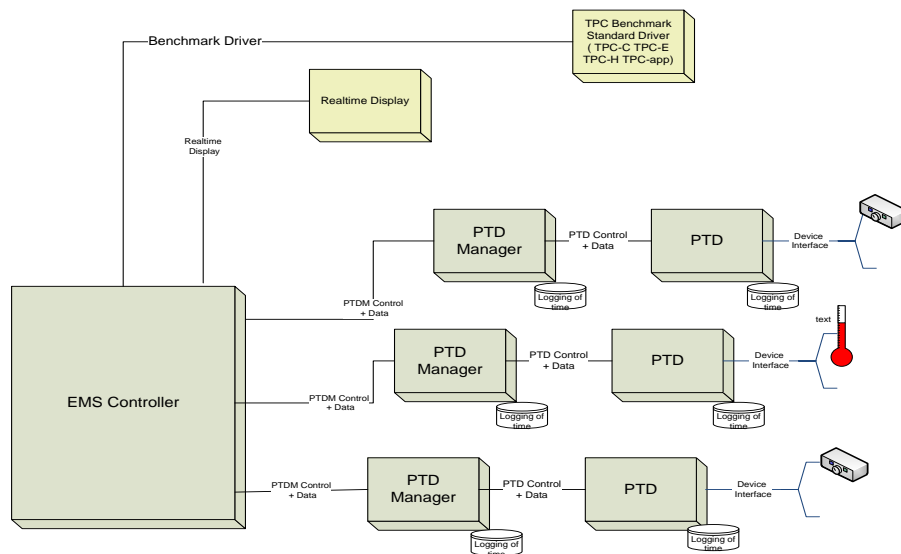
Where p_i is power measurement of each subsystem i during the run.

For example, $p_i = p_{\text{storage}}$ is the power consumption of the storage subsystem for the run.

T is the time elapsed time for the performance run.

$P(t)$ power consumption measured at time t

Figure 4 Power Measurement for Subsystems



Comment: For idle Power reporting please refer the TPC-Energy Specification.

6.3 TPCx-HS Energy Metric

The energy measurement from the **reported run** is used to compute the TPCx -HS power metric at the chosen Scale Factor. It must be computed as:

$$E / (T * HSph@SF)$$

- Where:
 - E is the energy consumption for the **reported run**
 - T is the elapsed time in seconds for the **reported run**
 - *HSph@SF* is the reported performance metric

Comment: The units of Energy Metric are reported to one digit after the decimal point, rounded to the nearest 0.1.

6.4 Reporting TPC-Energy Optional Metrics

When the TPC-Energy optional reporting is selected by the test sponsor, the additional requirements and format of TPC-Energy related items in the executive summary and FDR are included in the TPC Energy Specification, located at www.tpc.org.

CLAUSE 7: FULL DISCLOSURE

Rules for reporting **Pricing** information are included in the TPC Pricing Specification located at www.tpc.org.

7.1 Reporting Requirements

- 7.1.1 A Full Disclosure Report (FDR) in pdf format and Executive Summary are required.
- 7.1.2 The intent of this disclosure is to simplify comparison between results and for a customer to be able to replicate the results of this benchmark given appropriate documentation and products.

7.2 Format Guidelines

- 7.2.1 While established practice or practical limitations may cause a particular benchmark disclosure to differ from the examples provided in various small ways, every effort should be made to conform to the format guidelines. The intent is to make it as easy as possible for a reviewer to read, compare and evaluate material in different benchmark disclosures.
- 7.2.2 All sections of the report, including appendices, must be printed using font sizes of a minimum of 8 points.
- 7.2.3 The Executive Summary must be included near the beginning of the full disclosure report.

7.3 Full Disclosure Report

The FDR should be sufficient to allow an interested reader to evaluate and, if necessary, recreate an implementation of TPCx-HS. If any sections in the FDR refer to another section of the report, the names of the referenced scripts/programs must be clearly labeled in each section. Unless explicitly stated otherwise “disclosed” refers to disclosed in the FDR.

Comment: Since the building test environment may consist of a set of scripts and corresponding input files, it is important to disclose and clearly identify, by name, scripts and input files in the FDR.

The order and titles of sections in the test sponsor's full disclosure report must correspond with the order and titles of sections from the TPCx-HS standard specification (i.e., this document).

7.4 General Items

- 7.4.1 A statement identifying the benchmark sponsor(s) and other participating companies must be provided.
- 7.4.2 Settings must be provided for all customer-tunable parameters and options that have been changed from the defaults found in actual products, including but not limited to:
 - Configuration parameters and options for server, storage, network and other hardware component incorporated into the pricing structure;
 - Configuration parameters and options for operating system and file system component incorporated into the pricing structure;

- Configuration parameters and options for any other software component incorporated into the pricing structure.
- Compiler optimization options.

Comment 1: In the event that some parameters and options are set multiple times, it must be easily discernible by an interested reader when the parameter or option was modified and what new value it received each time.

Comment 2: This requirement can be satisfied by providing a full list of all parameters and options, as long as all those that have been modified from their default values have been clearly identified and these parameters and options are only set once.

7.4.3 Explicit response to individual disclosure requirements specified in the body of earlier sections of this document must be provided.


7.4.4 Diagrams of both measured and priced configurations (as defined in the TPC Pricing Specification) must be provided, accompanied by a description of the differences. This includes, but is not limited to:

- Total number and type of nodes used;
- Total number and type of processors used/total number of cores used/total number of threads used (including sizes of L2 and L3 caches);
- Size of allocated memory, and any specific mapping/partitioning of memory unique to the test;
- Number and type of data storage units (for example, disks or Licensed Compute Service volumes);
- Number and type of storage controllers, if applicable;
- Number of channels or bus connections to disk units, including their protocol type (if applicable);
- Number of LAN (for example, Ethernet) connections and speed for switches and other hardware components used in the test (if applicable) or are incorporated into the pricing structure;
- Type and the run-time execution location of software components.

The following sample diagram illustrates a measured benchmark configuration using Ethernet, an external driver, and four processors each with two cores and four threads per node in the SUT. Note that this diagram does not depict or imply any optimal configuration for the TPCx-HS benchmark measurement.

Depending on the implementation of the SUT, the Name Node, Secondary Name Node, Data Node, Job/Task Tracker, Resource Manager/Node Manager, etc. or the functional equivalents must be specified in the diagram.

Figure 5 Sample Configuration Diagram



- 4 x My Server Model B, 4/32/64 My CPU Model Z (2.7 GHz, 20MB cache, 130W), 128GB, My RAID Controller with 1GB BBWC
- 4 x My Storage Array Model A with 8 X 1TB 10K SAS HDD
- 2x My Switch Model X 10GbE

Comment: Detailed diagrams for system configurations and architectures can vary widely, and it is impossible to provide exact guidelines suitable for all implementations. The intent here is to describe the system components and connections in sufficient detail to allow independent reconstruction of the measurement environment. This example diagram shows homogeneous nodes. This does not preclude tests sponsors from using heterogeneous nodes as long as the system diagram reflects the correct system configuration.

7.4.5 The distribution of dataset across all media must be explicitly described using a format similar to that shown in the following example for the tested system.

Table 2: Sample Layout Description

Server	Controller	Disk Drive	Description of Content
1	40A	0	Operating system, root, swap, Hadoop Master
		1-12	Distributed file system Metadata

2	40A	0	Operating system, root, swap
		1-12	25% of HDFS data, 25% of temporary data
3	40A	0	Operating system, root, swap
		1-12	25% of HDFS data, 25% of temporary data
4	40A	0	Operating system, root, swap
		1-2424	50% of HDFS data, 50% of temporary data

7.4.6 The distribution of various software components across the system must be explicitly described using a format similar to that shown in the following example for both the tested and priced configurations.

Table 3: Distribution of Software Components

Server	Software Component
1	Name Node Job Tracker Benchmark driver
2	Data Node, Task Tracker
3	Data Node
4	Data Node, Task Tracker

Comment: The software components might vary from implementation to implementation.

7.4.7 Distributed file system implementation (for example Apache HDFS, Red Hat Storage, IBM GPFS, EMC Isilon OneFS) and corresponding Hadoop File System API version must be disclosed.

7.4.8 MapReduce or Spark implementation (for example, Apache MapReduce, IBM Platform Symphony) and corresponding version must be disclosed.

7.5 Workload Related Items

7.5.1 Script or text used to set for all hardware and software tunable parameters must be reported

- 7.5.2 Version number of TPCx-HS kit and checksum for HSGen, HSSort and HSValidate Programs must be reported
- 7.5.3 The run report generated by TPCx-HS benchmark kit must be reported.

7.6 SUT Related Items

- 7.6.1 The data storage ratio must be disclosed. It is computed by dividing the total physical data storage present in the priced configuration (expressed in TB) by the chosen Scale Factor as defined in Clause 4.1. Let r be the ratio. The reported value for r must be rounded to the nearest 0.01. That is, reported value= $\text{round}(r,2)$. For example, a system configured with 96 disks of 1TB capacity for a 1TB Scale Factor has a data storage ratio of 96.

Comment: For the reporting of data storage capacity, terabyte (TB) is defined to be 10^{12} bytes.

Comment: For consumption based storage provisioning in Licensed Compute Services, the maximum storage provisioned during the entire benchmark test is considered to be the total physical data storage present.

- 7.6.2 The Scale Factor to memory ratio must be disclosed. It is computed by dividing the Scale Factor by the total physical memory present in the priced configuration (see CLAUSE 5.2). Let r be this ratio. The reported ratio must be rounded to the nearest 0.01. That is, reported value= $\text{round}(r,2)$. For example, a system configured with 1TB of physical memory for a 10TB Scale Factor has a memory ratio of 10.00.

Comment: For Licensed Computing Services, the maximum provisioned memory during the entire benchmark test is considered to be the total physical memory present.

7.7 Scale Factors and Metrics

- 7.7.1 The HSGen time must be disclosed for Run1 and Run2
- 7.7.2 The HSSort time must be disclosed for Run1 and Run2.
- 7.7.3 The HSValidate time must be disclosed for Run1 and Run2.
- 7.7.4 Both HSDataCheck times must be disclosed for Run1 and Run2.
- 7.7.5 The performance metric (HSph@SF) must be disclosed for Run1 and Run2. Price-performance metric ($\$/\text{HSph@SF}$) must be disclosed for the **performance run**. See section 2.3 Benchmark Execution and CLAUSE 4: SCALE FACTORS AND METRICS.

7.8 Clause 8 – Audit Related Items

- 7.8.1 If the benchmark is audited by an Independent Auditor, 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 whom to contact in order to obtain further information regarding the audit process.

7.9 Executive Summary

7.9.1 The executive summary is meant to be a high level overview of a TPCx-HS implementation. It should provide the salient characteristics of a benchmark execution (metrics, configuration, pricing, etc.) without the exhaustive detail found in the FDR. When the TPC-Energy optional reporting is selected by the test sponsor, the additional requirements and format of TPC-Energy related items in the executive summary are included in the TPC Energy Specification, located at www.tpc.org.

7.9.2 The executive summary has three components:

- Implementation Overview
- Pricing Spreadsheet
- Numerical Quantities

7.10 Page Layout

7.10.1 Each component of the executive summary should appear on a page by itself. Each page should use a standard header and format, including

- 1/2 inch margins, top and bottom
- 3/4 inch left margin, 1/2 inch right margin
- 2 pt. frame around the body of the page. All interior lines should be 1 pt.

7.11 Implementation Overview

7.11.1 The implementation overview page contains five sets of data, each laid out across the page as a sequence of boxes using 1 pt. rule, with a title above the required quantity. Both titles and quantities should use a 9-12 pt. Times font unless otherwise noted.

7.11.2 The first section contains information about the sponsor and system identification

Table 4: Sponsor and System Identification

Title	Font
Sponsor Name or Logo	16-20 pt. Bold (for Name)
System Identification	16-20 pt. Bold
Version Numbers for TPCx-HS, TPC-Pricing and TPC-Energy (if reported)	16-20 pt. Bold
Report Date	16-20 pt. Bold

Comment 1: It is permissible to use or include company logos when identifying the sponsor.

Comment 2: The report date must be disclosed with a precision of 1 day. The precise format is left to the test sponsor.

7.11.3 The second section contains the Total System Cost; and, TPCx-HS Performance Metric and Price/Performance for the **performance run**.

Table 5: Test Results

Title	Quantity	Precision	Font
Total System Cost	3 yr. Cost of ownership (see CLAUSE 5)	1	16-20 pt. Bold
TPCx-HS Performance Metric	HSph (see CLAUSE 4: SCALE FACTORS AND METRICS Error! Reference source not found.)	0.01	16-20 pt. Bold
Price/Performance	\$/ HSph (see CLAUSE 4: SCALE FACTORS AND METRICS)	0.01	16-20 pt. Bold

Depending on the currency used for publication this sign has to be exchanged with the ISO currency symbol.

7.11.4 The third section contains detailed the system configuration.

Table 6: System Configuration Information

Title	Quantity	Font
Apache Hadoop/Spark Compatible Software	Product name and Product Version	9-12 pt. Times
Operating System	Product name, Software Version of OS, File System Type and Version	9-12 pt. Times
Other Software	Product name and Software Version of other software components (example Java)	9-12 pt. Times
System Availability Date	The Availability Date of the system, defined in Clause 0 of the TPC Pricing Specification.	9-12 pt. Times

Comment: The Software Version must uniquely identify the orderable software product referenced in the Priced Configuration (for example, RALF/2000 4.2.1)

7.11.4 The fourth section contains the storage and memory ratios. See section **Error! Reference source not found.**

Table 7: Storage and Memory Ratios

Title	Precision	Font
Physical Storage /Scale Factor	0.01	9-12 pt. Times
Scale Factor/Physical Memory	0.01	9-12 pt. Times

7.11.5 The fifth section contains the components (see CLAUSE 7.4.4), including:

- Total number and type of nodes used;
- Total number of processors used with their types and speeds in GHz/ total number of cores used/total number of threads used;
- Main and cache memory sizes;
- Network and I/O connectivity;
- Number and type of data storage units (for example, disks or Licensed Compute Service volumes);
- Total Rack Units (RU) occupied by the SUT

Comment: Rack Units (RU) occupied by the SUT include servers, storage, connectivity devices and any additional rack space required to be kept empty by the devices. Free space available on the rack, space for vertical mount PDU etc., are not to be included.

E.g.: 8 Servers (2RU each) and 2 Switches (1RU each) mounted in a 42RU rack. The Rack Units occupied by the SUT is $8 \times 2 + 2 \times 1 = 18\text{RU}$

Comment: If the SUT as priced does not occupy any Rack Units (for example in a fully virtual Licensed Compute Service configuration) then the Total Rack Units must be reported as “NA”.

7.12 Pricing Spreadsheet

7.12.1 The major categories in the Price Spreadsheet, as appropriate, are:

- Network(s)
- Server(s) /Node(s)
- Storage
- Software

7.12.2 Discounts (may optionally be included with above major category subtotal calculations).

7.13 Numerical Quantities Summary

7.13.1 The Numerical Quantities Summary page contains two sets of data, presented in tabular form, detailing the execution timings for the reported execution of the performance test. Each set of data should be headed by its given title and clearly separated from the other tables.

7.13.1 The first section contains measurement results from the benchmark execution.

Table 8: Measurement Results for Performance Run

Item Title	Precision
Scale Factor	1
Run Start Time	yyyy-mm-dd hh:mm:ss
Run End Time	yyyy-mm-dd hh:mm:ss
Run Elapsed Time	ss.sss
Start of HSGen	yyyy-mm-dd hh:mm:ss
End of HSGen	yyyy-mm-dd hh:mm:ss
HSGen Time	ss.sss
Start of HSSort	yyyy-mm-dd hh:mm:ss
End of HSSort	yyyy-mm-dd hh:mm:ss
HSSort Time	ss.sss
Start of HSValidate	yyyy-mm-dd hh:mm:ss
End of HSValidate	yyyy-mm-dd hh:mm:ss
HSValidate Time	ss.sss

7.13.2 Second section contains the measurement result for the **repeatability run**. See Table 8: for contents and precision.

7.14 TPCx-HS Run Report

7.14.1 The run report from TPCx-HS must be included in page 4 of the Executive Summary

7.15 Availability of the Full Disclosure Report

- The full disclosure report must be readily available to the public. The report must be made available when results are made public. In order to use the phrase “TPC Benchmark HS”, the full disclosure report submitted electronically to the TPC using the procedure described in the TPC Policies and Guidelines document.
- The official full disclosure report must be available in English but may be translated to additional languages.

7.16 Revisions to the Full Disclosure Report

7.16.1 Revisions to the full disclosure documentation shall be handled as follows:

- Substitutions will be open to challenge for a 60 day period. No other portion of the FDR and supporting files archive are challengeable.
- During the normal product life cycle, problems will be uncovered that require changes, sometimes referred to as ECOs, FCOs, patches, updates, etc. When the cumulative result of applied changes causes the HSph rating of the system to decrease by more than two percent from the initially reported HSph, then the test sponsor is required to re-validate the benchmark results. The complete revision history is maintained following the timing interval section showing the revision date and description.
- Full disclosure report and supporting files archive revisions may be required for other reasons according to TPC policies (see Transaction Processing Performance Council (TPC)).

CLAUSE 8: AUDIT

Rules for auditing **Pricing** information are included in the TPC Pricing Specification located at www.tpc.org.

When the TPC-Energy optional reporting is selected by the test sponsor, the rules for auditing of TPC-Energy related items are included in the TPC Energy Specification located at www.tpc.org. If TPC-Energy metrics are reported the TPCx-HS result must be audited by a TPC-Energy certified auditor.

8.1 General Rules

An audit of the benchmark result is required before publication.

The vendor may choose either a TPC Certified, independent, Auditor or the TPCx-HS Pre-Publication Board (peer review committee) to conduct the audit.

8.1.1 The term independent is defined as “the outcome of the benchmark carries no financial benefit to the auditing agency other than fees earned directly related to the audit.” The auditing agency cannot have supplied any performance consulting under contract for the benchmark. The independent auditor must be certified by the TPC to audit TPCx-HS. The independent auditor must meet the following::

- The auditor holds an active TPC certification for a TPC enterprise benchmark,
- The auditing agency cannot be financially related to the sponsor. For example, the auditing agency is financially related if it is a dependent division of the sponsor, the majority of its stock is owned by the sponsor, etc.
- The auditing agency cannot be financially related to any one of the suppliers of the measured/priced configuration.
- The auditor's attestation letter is to be made readily available to the public as part of the full disclosure report. A detailed report from the auditor is not required.

8.1.2 The term peer review is defined as the process of reviewing benchmark results for compliance prior to publication by the TPCx-HS Pre-Publication Board. The Pre-Publication Board consists of 3 members from the TPCx-HS committee. Each member serves a period of six months. The membership will be rotated through the TPCx-HS membership. The submission is confidential to the peer review committee until the result is published. The peer review committee must complete the review in 10 business days. If no issues are raised in 10 days the result is considered valid.

8.1.3 TPCx-HS results can be used as the basis for new TPCx-HS results if and only if:

- The auditor or peer review committee ensure that the hardware and software products are the same as those used in the prior result;
- The auditor or peer review committee reviews the FDR of the new results and ensures that they match what is contained in the original sponsor's FDR;
- The auditor or peer review committee can attest to the validity of the pricing used in the new FDR.

Comment 1: The intent of this clause is to allow a reseller of equipment from a given supplier to publish under the reseller's name a TPCx-HS result already published by the supplier.

8.2 Audit Check List

8.2.1 CLAUSE 2: WORKLOAD AND EXECUTION Related Items

8.2.1.1 Verify that the TPC provide kit is used

8.2.1.2 Verify that all phases are complete with no error in Run1 and Run2

8.2.1.3 Verify that all scripts and source code to implement the benchmark is included.

8.2.2 CLAUSE 3: SYSTEM UNDER TEST AND DRIVER Related Items

8.2.2.1 Verify that all components of the SUT are commercially available as per TPC Pricing Specification

8.2.2.2 Verify that all components of the SUT are included in the pricing

8.2.3 CLAUSE 4: SCALE FACTORS AND METRICS Related Items

8.2.3.1 Verify that the system is scaled as per the specification

8.2.3.2 Verify that the metrics are reported as per the precision requirements in clause 7

8.2.4 CLAUSE 5: PRICING Related Items

8.2.4.1 Verify that the benchmark is in compliance with the TPC Pricing specification

8.2.5 CLAUSE 6: ENERGY Related Items


8.2.5.1 Verify that the benchmark is in compliance with the TPC Energy specification (if reported)

8.2.6 CLAUSE 7: FULL DISCLOSURE Related Items

8.2.6.1 Verify that full disclosure report and executive summary report are accurate and comply with the reporting requirements. This includes:

- Metric calculation;
- System availability;
- The diagrams of both measured and priced configuration.
- System pricing;
- The numerical quantity summary

CLAUSE 9: SAMPLE EXECUTIVE SUMMARY

My Company Logo		My Server Model B	TPCx-HS Rev. 1.1.0	
			TPC-Pricing Rev. 2.0.1	
Total System Cost		Performance Metric	Price / Performance	
\$99,996.13 USD		390.99 HSph@3TB	\$255.76 USD \$ / HSph @3TB	
Scale Factor	Apache Hadoop/Spark Compatible Software	Operating System	Other Software	Availability Date
3TB	My HDFS Software 1.0	My OS V2.0	None	December 15, 2014
System Configuration				
 <p>The diagram illustrates the system configuration. On the left, four storage arrays (labeled '4 x My Storage Array Model A') are connected via SAS 3.0 to four servers (labeled '4 x My Server Model B'). Each server is connected to two network switches (labeled '2 x My Switch Model X') via 10 GbE connections. The switches are stacked vertically.</p>				
Physical Storage /Scale Factor		Scale Factor/Physical Memory		
Servers		4 x My Server Model B		
Processors/Cores/Threads/Model		4/32/64 My CPU Model Z (2.7 GHz, 20MB cache, 130W)		
Memory		128GB		
Storage		2 x 600GB 10K SFF SAS (internal) 1 x My Storage Array Model A with 8 X 1TB 7.2K SAS LFF HDD		
Network:		2x My Switch Model X 10GbE		
Rack Space:		18 RU		

My Company Logo	My Server Model B			TPCx-HS Rev. 1.1.0		
				TPC-Pricing Rev. 2.0.1		
				Report Date:	15-Dec-2014	
Description	Part Number	Source	Unit Price	Qty	Extended Price	3 Year Maint. Price
My Server Model B, 4 My CPU Model Z, 128GB, 2 x 600GB 10K SFF SAS	MY-S-001	1	12,100.77	4	\$48,403	\$100
My Storage Array Model A	MY-SE-002	1	1,988.00	4	\$7,952	\$200
My HDD Model xyz 1TB SATA 7.2K LFF	MY-HDD-011	1	800.47	40	\$32,019	
My OS	MY-OS	1	485.24	4	\$1,941	
My HDFS Software	MY-Hadoop	1	2,700.00	4	\$10,800	
My Switch Model X	My-Switch	1	1,922.12	2	\$3,844	
				Subtotal	\$104,959	\$300
Large Purchase Discount	5.0%	1			-\$5,248	-\$15
Pricing: 1=My Company			Three-Year Cost of Ownership: \$99,996.1			
Audited by My Auditor						
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.						
			HSph@3TB: 390.99			
			\$ / HSph @3TB: \$255.76			
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.						

My Company Logo	My Server Model B	TPCx-HS Rev. 1.1.0
		TPC-Pricing Rev. 2.0.1
		December 15, 2014
Measurement Results		
Scale Factor	3000	
Start of Run	10/02/2014 02:01:09	
End of Run	10/02/2014 08:11:31	
Run Elapsed Time	6:10:22	
Start of HSDataCheck	10/02/2014 02:01:09	
End of HSDataCheck	10/02/2014 02:01:16	
Start of HSGen	10/02/2014 02:01:16	
End of HSGen	10/02/2014 03:08:26	
Start of HSSort	10/02/2014 03:08:27	
End of HSSort	10/02/2014 05:24:26	
Start of HSValidate	10/02/2014 05:24:27	
End of HSValidate	10/02/2014 08:11:26	
Start of HSDataCheck	10/02/2014 08:11:27	
End of HSDataCheck	10/02/2014 08:11:31	
Performance Metric (HSph@3TB)	390.99@ HSph 3TB	
Price/Performance (\$/HSph@3TB)	\$255.76/HSph@3TB	