TPC EXPRESS BENCHMARK™ IoT

(TPCx-IoT)

Standard Specification

Draft Version 1.0.0

May 1, 2017

Transaction Processing Performance Council (TPC)

www.tpc.org

mailto:info@tpc.org

Legal Notice

No Warranty

Trademarks

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-IoT subcommittee member companies in developing the TPCx-IoT specification.

The TPCx-IoT subcommittee would like to acknowledge the contributions made by the many members during 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 to this version includes <TBD>

Document Revision History

Table 1: Document Revision History

Date	Version	Description
02/08/2017	1.0.0	Initial Draft

TPC Membership

TPC membership as of February 2017.

Full Members <TBD>

Actian	Cisco	Cloudera	Dell	Fujitsu
Hewlett Packard Enterprise	Hitachi	Huawei	IBM	Inspur
Intel	Lenovo	MapR	Microsoft	
Oracle	Pivotal	Redhat	SAP	Teradata
Unisys	VMware			

Associate Members

IDEAS International	ITOM International Co	TTA	

Clause 1 Introduction

1.1 Preamble

Devices and sensors connected to the internet are expected to number in the billions within a few years. This Internet of Things (IoT) is being adopted across almost every industry triggering a massive influx of data that has to be analyzed in real-time for insights. Millions upon millions of connected things will generate massive amounts of data all of which needs to be processed by the systems reliably and with high-performance.

TPC Express Benchmark™ IoT (TPCx-IoT) was developed to provide an objective measure of hardware, operating system, data storage and data management systems to provide the industry with verifiable performance, price-performance and availability metrics for systems which are meant to ingest and persist massive amounts of data from large number of devices, and provide real-time insights typical in IoT gateway systems running commercially available NoSQL database systems.

The TPCx-IoT benchmark models a continuous system available 24 hours a day, 7 days a week. The TPCx-IoT can be used to assess a broad range of system topologies and implementation methodologies in a technically rigorous, directly comparable, vendor-neutral manner.

1.2 TPCx-IoT Kit and Licensing

A full benchmark kit is provided by the TPC. Vendors are required to use this kit for benchmark publications. The TPCx-IoT kit includes a set of scripts to generate data simulating IoT sensors. The data is ingested and persisted into the System Under Test (SUT) and continuously queried to simulate simple edge analytics use cases, generate metrics and create a full run report. The SUT represents an IoT edge gateway system consisting of commercially available servers, switches and storage systems running a NoSQL database (See Clause 3).

The TPCx-IoT kit is available from the TPC (see www.tpc.org/TPCx-IoT for more information). Users must sign-up and agree to the TPCx-IoT User Licensing Agreement (ULA) to download the kit. Re-distribution of the kit is prohibited. All related work (e.g., collateral, papers, derivatives, etc.) must acknowledge the TPC and include the TPCx-IoT copyright. The TPCx-IoT Kit includes: TPCx-IoT Specification document (this document), TPCx-IoT User Guide,

Scripts to set up the benchmark environment, run the benchmark, run data validation and capture system inventory, Java code to execute the benchmark load against the SUT. See Clause 2.

1.3 General Implementation Guidelines

The purpose of TPC benchmarks are to provide relevant, objective, and verifiable 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-IoT models representing IoT gateway systems that can be used for intermediate/final data storage);
- Would plausibly be implemented by a significant number of users in the market segment the benchmark models or represents.

The fault-recovery attributes of the system under test (SUT) 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, data loading/updating 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 running 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 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” implementation. 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's primary purpose is performance optimization of TPC benchmark results without any corresponding applicability to real-world applications and environments?"

The following characteristics shall be used to make this judgment:

- 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 TPCx-IoT 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-IoT 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.

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

The following sections provides the contents of the benchmark kit and usage guidelines.

2.1.1 Kit Contents

The TPCx-IoT kit contains the following:

- TPCx-IoT Specification (this document)
- TPCx-IoT User Guide
- Driver Program
- Scripts to setup the benchmark environment, capture system inventory, run the benchmark, and validate the run
- Java code to execute the benchmark load

2.1.2 Kit Usage

To submit a compliant TPCx-IoT benchmark result, the test sponsor is required to use the TPCx-IoT kit as provided except for modifications explicitly listed in Clause 2.1.3
The kit must be used as outlined in the TPCx-IoT User Guide. The output of the TPCx-IoT kit is called the run report which includes the following:

- Version number of TPCx-IoT kit
- Checksum for the TPCx-IoT programs
- Validation for compliance (number of records ingested, data replication factor)
- Verification of data

If there is a conflict between the TPCx-IoT 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-IoT kit to facilitate operating system differences or the storage that is being used are allowed without TPC approval.
The following changes are considered minor modifications:

- Shell script changes necessary for the 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

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-IoT 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.

2.1.3.2 Java Code Modifications

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

2.1.4 Future Kit Releases

The TPC will release future TPCx-IoT 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-IoT 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-IoT benchmark kit.

2.2 Benchmark Workload

The TPC Benchmark™ IoT (TPCx-IoT) benchmark workload is designed based on Yahoo Cloud

Serving Benchmark (YCSB)[1] with data ingestion and concurrent queries simulating workloads on typical IoT Gateway systems. The dataset represents data from a power transmission unit sensors with 50 attributes as listed below:

Angle_Diff_Chester_Madison
Angle_Diff_Madison_Fairview
Angle_Diff_Hudson_Fairview
Angle_Diff_Chester_Hudson
Angle_Diff_Fairview_Hudson
Angle_Diff_Chester_Madison
Chester_Curr_MVar
Chester_Curr_MW
Chester_Curr_Angle
Chester_Curr_Mag
Chester_Volt_Angle
Chester_Volt_Mag
Chester_Speed_Mag
Chester_Throttle_Mag
Chester_Temp_MW
Chester_Wheel_Speed_MW
Madison_Freq
Madison_Curr_MVar
Madison_Curr_MW
Madison_Curr_Angle
Madison_Curr_Mag
Madison_Volt_Angle
Madison_Volt_Mag
Madison_Speed_Mag
Madison_Throttle_Mag
Madison_Temp_MW
Madison_Wheel_Speed_MW
Fairview_Freq
Fairview_Curr_MVar
Fairview_Curr_MW
Fairview_Curr_Angle
Fairview_Curr_Mag
Fairview_Volt_Angle
Fairview_Volt_Mag
Fairview_Speed_Mag
Fairview_Throttle_Mag
Fairview_Temp_MW
Fairview_Wheel_Speed_MW
Hudson_Freq

[1] Yahoo! Cloud Serving Benchmark (YCSB) References (i) Wikipedia: https://en.wikipedia.org/wiki/YCSB (ii) Benchmarking Cloud Serving Systems with YCSB: https://www.cs.duke.edu/courses/fall13/cps296.4/838-CloudPapers/ycsb.pdf (iii) YCSB+T: Benchmarking web-scale transactional databases: https://www.computer.org/csdl/proceedings/icdew/2014/3481/00/06818330.pdf

Hudson_Curr_MVar
Hudson_Curr_MW
Hudson_Curr_Angle
Hudson_Curr_Mag
Hudson_Volt_Angle
Hudson_Volt_Mag
Hudson_Speed_Mag
Hudson_Throttle_Mag
Hudson_Temp_MW
Hudson_Wheel_Speed_MW

A key is generated based on the sensor attribute, time stamp, and driver(s) identification. The driver system(s) identification is used to ensure no duplicate data is inserted. The data management platform must be a commercially available NoSQL system and data must be replicated in a minimum of two ways, and persisted in a non-volatile durable media.

Each record is a 1000-byte record. Each run (see Clause 2) consists of one billion record inserts. The workload has the following parameters with an insert proportion of 100% and a query proportion of 0.1 %. For each 'n' records inserted, records inserted in the last 120 seconds are queried, and an average of each attribute value is calculated and logged in the report. At the end of the run (included in the run time) the latency times at 99 and 95 percentiles are also calculated and logged in the report.

2.3 Benchmark Execution

Data ingesting and querying will be performed against the NoSQL database by the driver program. A run consists of ingesting one billion records in to the SUT and querying every 120 seconds from SUT using driver system(s). In a multi driver environment the keys are distributed and partitioned among the driver systems(s). The maximum time taken by a driver system(s) is considered as the elapsed time for that run.

The benchmark test consists of two runs, Run1 and Run2. Each run consists of a Warmup Run and Measured Run. No activities are allowed between Warmup Run and Measured Run. 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 Performance Metric calculation. The Performance Run is defined as the Measured Run with the lower Performance Metric. The repeatability run is defined as the Measured Run with the higher Performance Metric. The Reported Performance Metric is the Performance Metric for the Performance Run.

No configuration or tuning changes are allowed between the runs. The benchmark execution phases are shown in the Figure below:

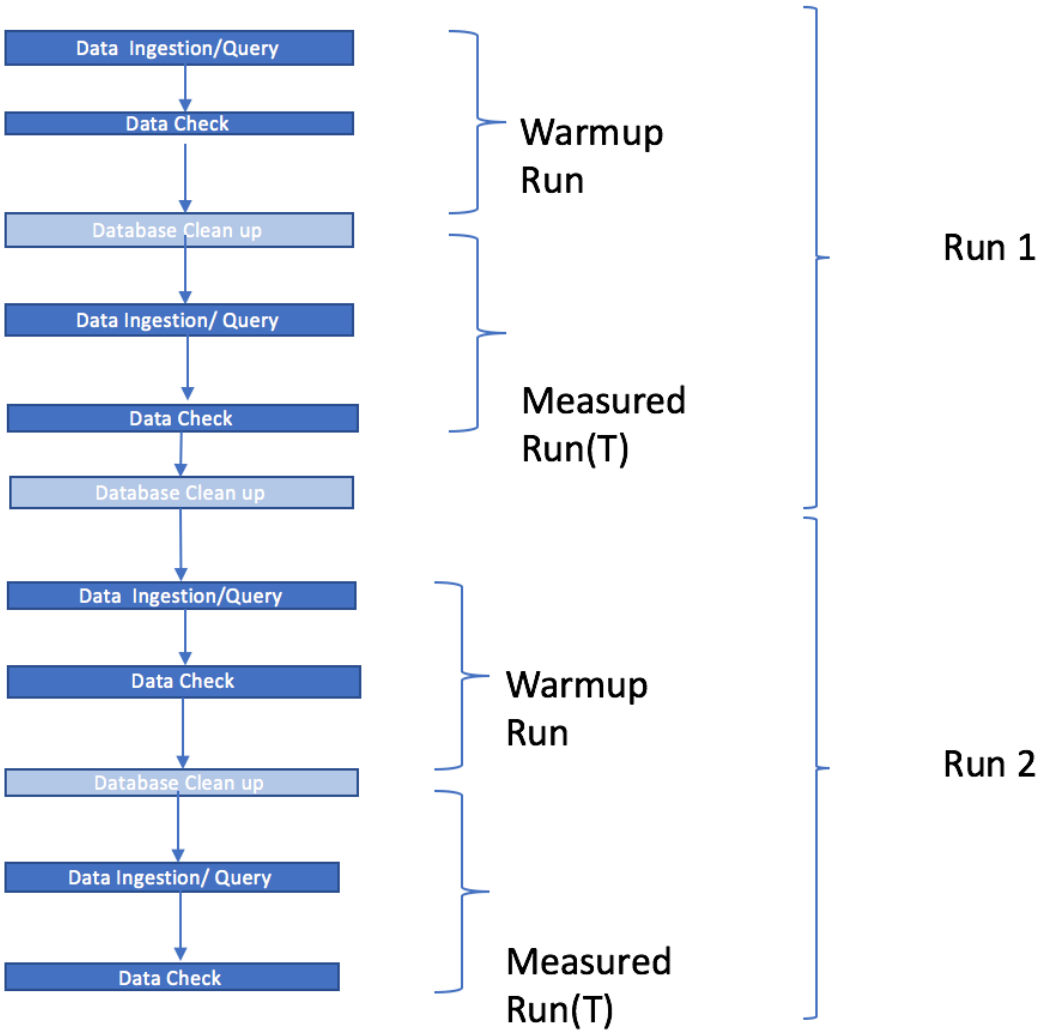

Figure 1: Benchmark Execution Phases

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.

2.4 Configuration and Tuning

The SUT cannot be reconfigured, changed, or re-tuned by the test sponsor during or between any of the 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.

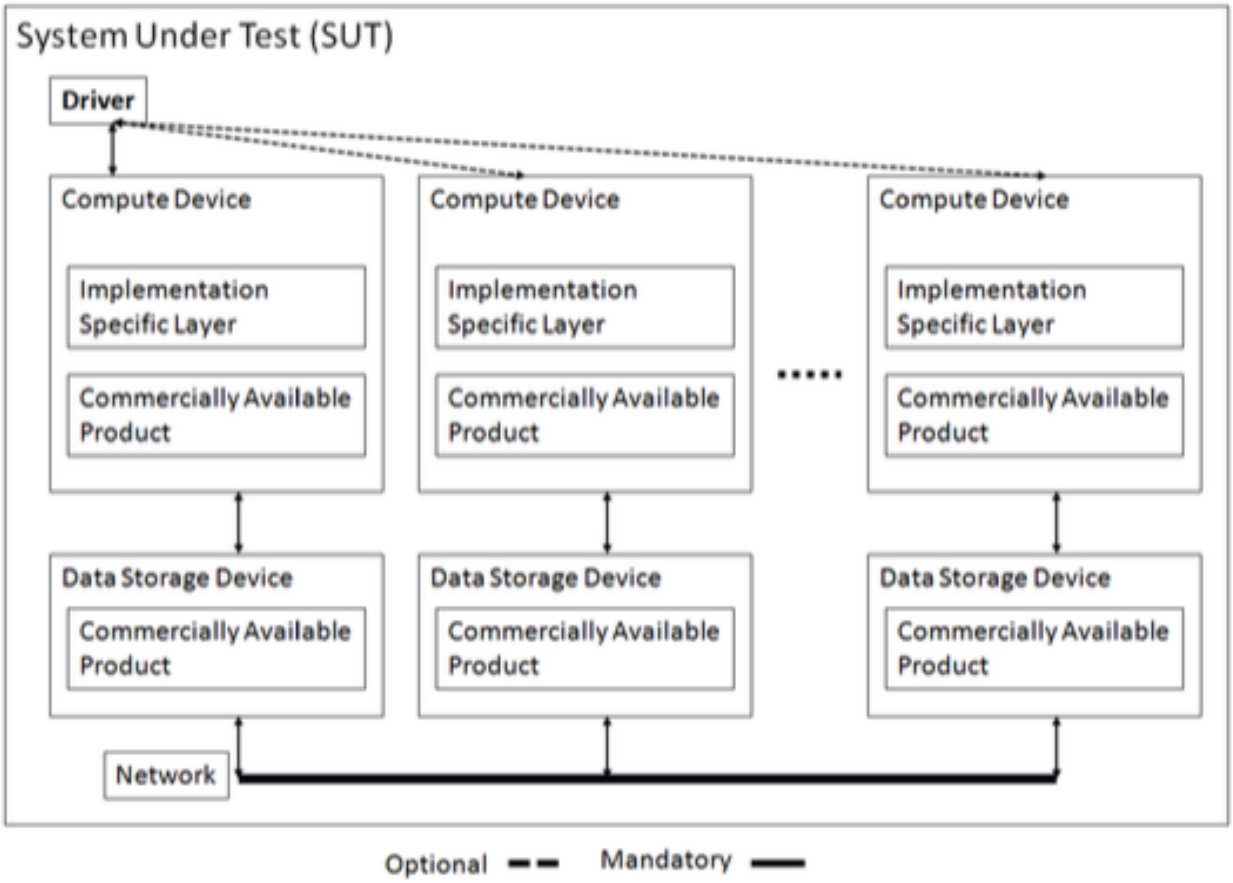

Figure 2 System Under Test (SUT)

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(s) present the workload to the SUT.

Clause 4: Scale Factor and Metrics

This clause defines Scale Factor and Metrics.

4.1 Scale Factor

The current version of the TPCx-IoT kit follows a fixed scale factor model which is one billion records.

Comment: The TPC will continuously evaluate increasing the scale factor based on industry trends.

4.2 Metrics

TPCx-IoT defines the following primary metrics:

1. IoTph, the Performance Metric, reflecting the TPCx-IoT throughput
2. $/IoTph , the Price-Performance metric
3. System availability date

4.3 Performance Metric

The **performance metric** of the benchmark is IoTph, the effective time for inserting on billion records. The metric represents data ingestion capability of the SUT.

IoTph = 1,000,000,000/(T)

T is the time elapsed in hours

4.4 Price Performance Metric

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

$/IoTph = P/ IoTph

P is the total cost of ownership of the SUT.

4.5 System Availability Date

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

4.6 Metric Comparison

A TPCx-IoT Result is only comparable with other TPCx-IoT Results. (see Clause 4).

- 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.
- In addition, the results must be accompanied by a disclaimer stating: "The TPC believes that comparisons of TPCx-IoT results measured against different Scale Factors are misleading and discourages such comparisons."

4.7 Required Reporting Components

To be compliant with the TPCx-IoT standard and TPC Polices, the URL to the benchmark result and Availability Date of the complete configuration must be included for all public references (See Clause 7: 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 System and associated software and maintenance are included in the TPC Pricing Specification located at www.tpc.org.

5.1 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 system consists of:

- Price of the SUT as tested and defined in Clause 3
- Price of a communication interface capable of supporting the required number of user interface devices defined in Clause 5.2
- 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 system 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

5.2 Additional Operational Components

Additional products included on a customer installed configuration are also to be included in the priced system if explicitly required for the operation, administration or maintenance of the priced system. 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 5.1 Priced System)

5.3 Software

All software licenses must be priced.

5.4 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 (see Clause 7: Audit).

Clause 6: Full Disclosure

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

6.1 Reporting Requirements

6.1.1 A Full Disclosure Report (FDR) in pdf format and Executive Summary are required.

6.1.2 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.

6.2 Format Guidelines

6.2.1 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.

6.2.2 All sections of the report, including appendices, must be printed using font sizes of a minimum of 8 points.

6.2.3 The Executive Summary must be included near the beginning of the full disclosure report.

6.3 Full Disclosure Report

The FDR should be sufficient to allow an interested reader to evaluate and, if necessary, recreate an implementation of TPCx-IoT. 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-IoT standard specification (i.e., this document).

6.4 General Items

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

6.4.2 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 components incorporated into the pricing structure
- Configuration parameters and options for operating system and file system components incorporated into the pricing structure
- Configuration parameters and options for any other software components 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.

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

6.4.4 Diagrams of both measured and priced configurations must be provided, accompanied by a description of the differences.

This includes, but is not limited to:

- Total number 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 disk units (and controllers, if applicable)
- Number of channels or bus connections to disk units, including their protocol type

- Number of LAN (for example, Ethernet) connections and speed for switches and other hardware components physically used in the test 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-IoT benchmark measurement.

Depending on the implementation of the SUT, the components for the storage system being used, the head node, the worker nodes etc. or the functional equivalents must be specified in the diagram

Figure 3: Sample Configuration Diagram

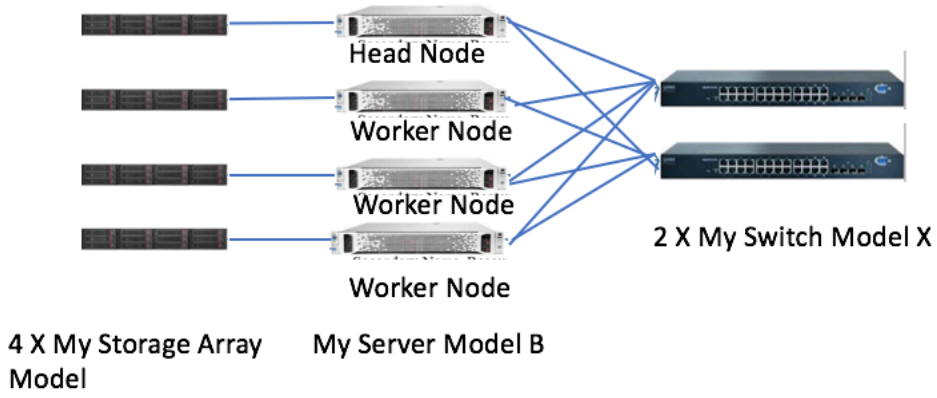

- 4 x My Server Model B, 4/32/64 My CPU Model Z (2.7 GHz, 20MB cache, 130W), 128GB, My RAID Controller with 1GB BBWC
- 4 x My Storage Array Model A with 8 X 1TB 10K SAS HDD
- 2x My Switch Model X 10GbE

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.

6.4.5 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 1: Sample Layout Description

Server	Controller	Disk Drive	Description of Content
1	40A	0	Operating system, root, swap, NoSQL Master
		1-12	Master Server File system Metadata
2	40A	0	Operating system, root, swap
		1-12	NoSQL worker data nodes
3	40A	0	Operating system, root, swap
		1-12	NoSQL worker data nodes
4	40A	0	Operating system, root, swap
		1-24	NoSQL worker data nodes

Number of nodes on which the TPCx-IoT driver is running is typically equal to or less than the number of nodes in the SUT. The driver nodes are usually comparable to SUT in compute, memory and network, but with no additional storage requirements.

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

6.4.7 Storage System used (for example, Cassandra/HBase) and corresponding version information must be disclosed.

6.5 Workload Related Items

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

6.5.2 Version number of TPCx-IoT kit and checksum for the kit components must be reported.

6.5.3 The run report generated by TPCx-IoT benchmark kit must be reported.

6.6 SUT Related Items

6.6.1 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 dataset size. In the current release the dataset size is fixed.

Let r be the ratio. The reported value for r must be rounded to the nearest 0.01. That is, reported value=round(r,2).

Comment: For the reporting of configured disk capacity, terabyte (TB) is defined to be 10^12 bytes.

6.6.2 The dataset size to memory ratio must be disclosed. It is computed by dividing the dataset size by the total physical memory present in the priced configuration (see 6.6.1).

> Let r be this ratio. The reported ratio must be rounded to the nearest 0.01. That is, reported value=round(r,2).

6.7 Metrics

Run and query time for the various workloads must be disclosed. The performance metric (IoTph) must be disclosed for Run1 and Run2. Price-performance metric ($/IoTph) must be disclosed for the performance run. See section 2.3 Benchmark Execution and Clause 4.

6.8 Clause 7 – Audit Related Items

6.8.1 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.

6.9 Executive Summary

6.9.1 The executive summary is meant to be a high-level overview of a TPCx-IoT 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.

6.9.2 The executive summary has three components:

- Implementation Overview
- Pricing Spreadsheet
- Numerical Quantities

6.10 Page Layout

6.10.1 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.

6.11 Implementation Overview

6.11.1 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.

6.11.2 The first section contains information about the sponsor and system identification.

Table 2: 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-IoT, 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 one day. The precise format is left to the test sponsor.

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

Table 3: Test Results

Title	Quantity	Precision	Font
Total System Cost	3 yr. Cost of Ownership (see Clause 5)	1	16-20 pt. Bold
TPCx-IoT Performance Metric	IoTph(see Clause 4: Scale Factor and Metrics)	0.01	16-20 pt. Bold
Price/Performance	$/IoTph(see Clause 4: Scale Factor 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.

6.11.4 The third section contains detailed the system configuration.

Table 4: System Configuration Information

Title	Quantity	Font
Storage System Software	Product Name and Product Version	9-12 pt. Times
Operating System	Product Name, Software Version for 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)

6.11.4 The fourth section contains the storage and memory ratios.

Table 5: Storage and Memory Ratios

Title	Precision	Font
Physical Storage/Data Set size	0.01	9-12 pt. Times
Data Set size /Physical Memory	0.01	9-12 pt. Times

6.11.5 The fifth section contains the components, including:

- Total number 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, see Clause 7: Audit
- Main and cache memory sizes
- Network and I/O connectivity
- Disk quantity and geometry
- 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.

Ex: 8 Servers (2RU each) and 2 Switches (1RU each) mounted in a 42RU rack. The Rack Units occupied by the SUT is 8x 2 + 2x 1 = 18RU

6.12 Pricing Spreadsheet

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

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

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

6.13 Numerical Quantities Summary

6.13.1 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.

6.13.2 The first section contains measurement results from the benchmark execution.
Table 6: Results for Measured Run

Item Title	Precision
Scale	1
Run Start Time	yyyy-mm-dd hh:mm:ss
Run End Time	yyyy-mm-dd hh:mm:ss

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

6.14 TPCx-IoT Run Report

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

6.15 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 IoT," the full disclosure report must be 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.

6.16 Revisions to the Full Disclosure Report

6.16.1 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 IoTph rating of the system to decrease by more than two percent from the initially reported IoTph, 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 7: 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-IoT result must be audited by a TPC-Energy certified auditor.

7.1 General Rules

An independent audit or peer audit of the benchmark result is required before publication. The vendor may choose an independent audit or peer audit.

7.1.1 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-IoT. 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.

7.1.2 The term peer review is defined as the process of reviewing benchmark results for compliance prior to publication by a committee named the peer review committee. The peer review committee consists of 3 members from the TPCx-IoT committee. Each member serves a period of three months. The membership will be rotated through the TPCx-IoT 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.

7.1.3 TPCx-IoT results can be used as the basis for new TPCx-IoT 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 re- seller's name a TPCx-IoT result already published by the supplier.

7.2 Audit Check List

7.2.1 Clause 2: Workload and Execution Related Items

7.2.1.1 Verify that the TPC provide kit is used

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

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

7.2.2 Clause 3: System Under Test and Driver Related Items

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

7.2.2.2 Verify that all components of the SUT is included in the pricing

7.2.3 Clause 4: Scale Factors and Metrics Related Items

7.2.3.1 Verify that the system is scaled as per the specification

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

7.2.4 Clause 5: Pricing Related Items

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

7.2.5 Clause 7: Full Disclosure Related Items

7.2.6.1 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 8: Sample Executive Summary

My Company Logo		My Server Model B	TPCx-IoT Rev. 1.1.0 TPC-Pricing Rev. 2.0.1 Report Date: May1, 2017	
Total System Cost		Performance Metric	Price / Performance	
$99,996.13 USD		390.99 IoTph	$255.76 USD $ / IoTph	
Scale Factor	Apache Hadoop Compatible software	Operating System	Other Software	Availability Date
100000000	My NoSQL Software 1.0	My OS V2.0	None	May 1, 2017

System Configuration

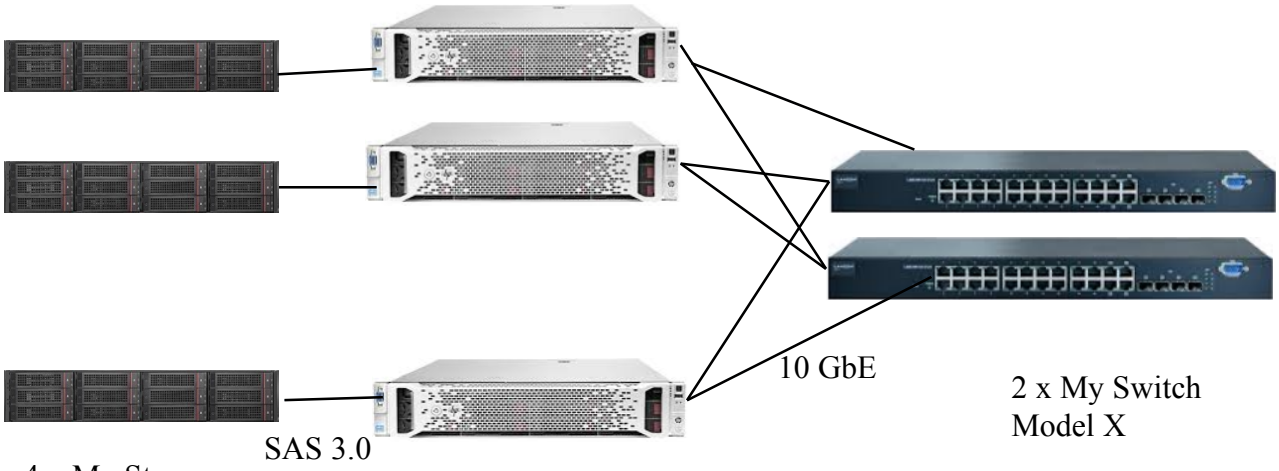

Physical Storage /DataSet Size	DataSet Size/Physical Memory
Servers	4 x My Server Model B
Processors/Cores/Threads/Model	4/32/64 My CPU Model Z (2.7 GHz, 20MB cache, 130W)
Memory	128GB
Storage	2 x 600GB 10K SFF SAS (internal) 1 x My Storage Array Model A with 8 X 1TB 7.2K SAS LFF HDD
Network:	2x My Switch Model X 10GbE

My Company Logo	My Server Model B	TPCx-IoT Rev. 1.0.0 TPC-Pricing Rev. 2.0.1 Report Date: May 1, 2017
Total System Cost	Performance Metric	Price / Performance
$99,996.13 USD	390.99 IoTph	$255.76 USD $ / IoTph

Scale Factor	Apache Hadoop Compatible software	Operating System	Other Software	Availability Date
100000000	My NoSQL Software 1.0	My OS V2.0	None	May 1, 2017

My Company Logo	My Server Model B			TPCx-IoT Rev. 1.1.0 TPC-Pricing Rev. 2.0.1		
				Report Date:		1-May-2017
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 NoSQL Software	MY-NoSQL	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.	IoTphH:	1,100.1
	$ / IoTphH:	$90.9

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.