

TPC EXPRESS BENCHMARK™ IoT (TPCx-IoT) Standard Specification Version 2.1.1

June, 2024

Transaction Processing Performance Council (TPC)

www.tpc.org

<mailto:info@tpc.org>

© 2024 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 like the Internet of Things (IoT) required a huge effort to work and contributions of the TPCx-IoT subcommittee member conceptualize research, specify, review, prototype, and verify the benchmark. The TPC acknowledges the companies in developing the TPCx-IoT Specification. The list of contributors to this version includes Andy Bond, Bhaskar Gouda, Karthik Kulkarni, Chaitanya Kundety, Chinmayi Narasimhadevara, Da Qi Ren, David Grimes, Meikel Poess, Nicholas Wakou, Jamie Reding, John Poelman, Ken Rule, Hamesh Patel, Mike Brey, Matthew Emmerton, Paul Cao, Reza Taheri, and Tariq Magdon-Ismail.

Document Revision History

Table 1: Document Revision History

Date	Version	Description
06/07/2017	1.0.0	<ul style="list-style-type: none">• Draft proposed for GC approval with all changes since formal review
09/21/2017	1.0.1	<ul style="list-style-type: none">• Editorial fixes. Workload is modified to include analytics query over randomly selected interval
12/6/2017	1.0.2	<ul style="list-style-type: none">• Precision upto 3 decimals for all metrics
1/29/18	1.0.3	<ul style="list-style-type: none">• Add list of supported NoSQL Databases
10/3/19	1.0.4	<ul style="list-style-type: none">• Add Machbase as supported NoSQL Database in Clause 2.5• Substitute member table under “TPC Membership” with link to TPC website on Page 5• Add TPCx-IoT version number to foot note
2/6/2020	1.0.5	<ul style="list-style-type: none">• Changed datatypes from integer to long in client code to allow for more than 2.1 billion records to be inserted by a single client.• Commented out not-used random number calls in function “acknowledge”• Switched to use java.util.concurrent for random number generation• Updated version number
02/10/2021	2.0.0	<ul style="list-style-type: none">• Changes related to price performance metric: affected clauses are 4.2, 4.4, 6.8, 8• Change timestamp interval for analytic query from 30 seconds to 5 seconds to make the specification consistent with the KIT: Affected Clause 2.2• General code cleanup• Change insert timestamp interval to 100ms in TPCx-IoT 1.0.4 from current 10s

		<ul style="list-style-type: none"> • Add reporting of the number of queries with 0 rows returned in the final report in the next release of TPCx-IoT
3/3/2022	2.0.1	<ul style="list-style-type: none"> • Issue 50: md5sum error for core-0.13.0-SNAPSHOT.jar component • Issue 51: Remote testing instances execution problem • Integration of Lindorm into the KIT
11/3/2022	2.1.0	<ul style="list-style-type: none"> • Support for two IP addresses in Machbase
06/26/2024	2.1.1	<ul style="list-style-type: none"> • Add TimechoDB as supported NoSQL Database in Clause 2.5

TPC Membership

A list of the current TPC member companies can be found at
http://www.tpc.org/tpc_documents_current_versions/pdf/tpcmembers.pdf

Table of Contents

Clause 1 Introduction	8
1.1 Preamble	8
1.2 TPCx-IoT Kit and Licensing	8
1.3 General Implementation Guidelines	8
1.4 General Measurement Guidelines	10
Clause 2: Workload and Execution	11
2.1 TPCx-IoT Kit	11
2.1.1 Kit Contents	11
2.1.2 TPCx-IoT Kit Usage	11
2.1.3 Kit Modification	11
2.1.3.1 Minor Shell Script Modifications	11
2.1.3.2 Major Shell Script Modifications	12
2.1.3.2 Java Code Modifications	12
2.1.4 Future Kit Releases	12
2.2 Benchmark Workload	12
2.3 Benchmark Execution	13
2.4 Configuration and Tuning	14
2.5 NoSQL Databases supported by the Benchmark	15
Clause 3: System Under Test and Benchmark Driver	16
3.1 System Under Test	16
Clause 4: Scale Factor and Metrics	18
4.1 Scale Factor	18
4.2 Metric	18
4.3 Performance Metric	18
4.4 Price Performance Metric	18
4.5 Availability Date	18
4.6 Metric Comparison	18
4.7 Required Reporting Components	18
Clause 5: Pricing	20
5.1 Priced System	20
5.2 Allowable Substitutions	20
Clause 6: Full Disclosure Report and Executive Summary	22
6.1 Reporting Requirements	22
6.2 Format Guidelines	22
6.3 Full Disclosure Report	22
6.4 General Items	22
6.5 Workload Related Items	25
6.6 Audit Related Items	25
6.7 Executive Summary	25
6.8 Implementation Overview	26
6.9 Pricing Spreadsheet	28
6.10 Numerical Quantities Summary	28
6.11 TPCx-IoT Run Report	28

The run report from TPCx-IoT must be included in the Executive Summary.....	28
6.12 Availability of the Full Disclosure Report.....	28
6.13 Revisions to the Full Disclosure Report	29
Clause 7: Audit	30
7.1 General Rules	30
<i>The Pre-Publication Board consists of three members of the TPCx-IoT committee.</i>	<i>30</i>
7.2 Audit Check List	31
7.2.1 Clause 2: Workload and Execution Related Items.....	31
7.2.2 Clause 3: System Under Test and Driver Related Items.....	31
7.2.3 Clause 4: Scale Factors and Metrics Related Items	31
7.2.4 Clause 5: Pricing Related Items	31
7.2.5 Clause 7: Full Disclosure Related Items	31
Clause 8: Sample Executive Summary	32

Clause 1 Introduction

1.1 Preamble

Internet of Everything (IoT) represents a global market transition driven by a surge in connections among people, processes and things. IoT is being adopted across almost every industry triggering a massive influx of data that has to be analyzed for insights. Typical IoT topology consists of three tiers: edge devices, gateway systems and backend data center. While there exist workloads for backend data center, there are no realistic and proven measures to compare different software and hardware solutions for gateway systems. To address this, TPC has developed TPC Express Benchmark™ IoT (TPCx-IoT).

TPCx-IoT provides 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 software and hardware.

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

TPCx-IoT is a TPC Express benchmark and a full kit (TPCx-IoT Kit) is provided by the TPC. Vendors are required to use this Kit for benchmark publications. The Kit includes a set of scripts to generate data simulating IoT sensors, data inject, analytics, calculate the metrics and validation.

The data generated is ingested and persisted into the System Under Test (SUT) and continuously queried to simulate simple analytics use cases. The System Under Test (SUT) represents an IoT gateway system consisting of commercially available servers and storage systems running a commercially available NoSQL data management system.

The Kit is available at [TPC Downloads page](#). Users must sign-up and agree to the TPCx-IoT User Licensing Agreement (ULA) to download the Kit.

To add support for a new database, follow the instructions in the 'How to Add a New Database' document included in the Kit .

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, the TPC Benchmark Specifications require that benchmark tests be implemented with systems, products, technologies and pricing that:

- Are commercially available;
- Are generally available to all users;
- Are relevant to the market segment that the individual TPC benchmark models;
- Would plausibly be implemented by a significant number of users in the market segment the benchmark models.

The use of new systems, products, technologies (software or hardware) 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 the [TPC Documentation](#) webpage.

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 TPCx-IoT Kit

The following sections provides the contents of the benchmark TPCx-IoT 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
- A document with instructions on how to add a new database
- 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 TPCx-IoT Kit Usage

To submit a compliant TPCx-IoT benchmark result, the test sponsor is required to use the TPCx-IoT kKt 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 Kit is called the run report which includes the following:

- Version number of 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 Kit 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 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 Kit.

2.2 Benchmark Workload

The TPC Benchmark™ IoT (TPCx-IoT) benchmark workload is designed based on Yahoo Cloud Serving Benchmark (YCSB)¹. It is not comparable to YCSB due to significant changes. The TPCx-

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>

IoT workloads consists of data ingestion and concurrent queries simulating workloads on typical IoT Gateway systems. The dataset represents data from sensors from electric power station(s). The data ingestion and query workloads are detailed in the following section.

Each record generated consists of driver system id, sensor name, time stamp, sensor reading and padding to a 1 Kbyte size. The driver system id represents a power station. The dataset represents data from 200 different types of sensors. The SUT must run a data management platform that is commercially available and data must be persisted in a non-volatile durable media with a minimum of two-way replication. The workload represents data inject in to the SUT with analytics queries in the background. The analytic queries retrieve the readings of a randomly selected sensor for two 5 second time intervals, TI_1 and TI_2 . The first time interval TI_1 is defined between the timestamp the query was started T_s and the timestamp 5 seconds prior to T_s , i.e. $TI_1=[T_s-5, T_s]$. The second time interval is a randomly selected 5 seconds time interval TI_2 within the 1800 seconds time interval prior to the start of the first query, T_s-5 . If $T_s \leq 1805$, then the time interval is 1705 seconds prior to the start of the first query, T_s-5 .

2.3 Benchmark Execution

Data ingestion and query are performed against the SUT by the driver program included in the TPCx-IoT Kit.

The benchmark test consists of two runs, Run1 and Run2. Each run consists of a Warmup Run and Measured Run. No activities other than database cleanup triggered by the control scripts are allowed between Warmup Run and Measured Run. No activities are allowed between Run 1 and Run2. The total elapsed time for the Performance 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 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 1.

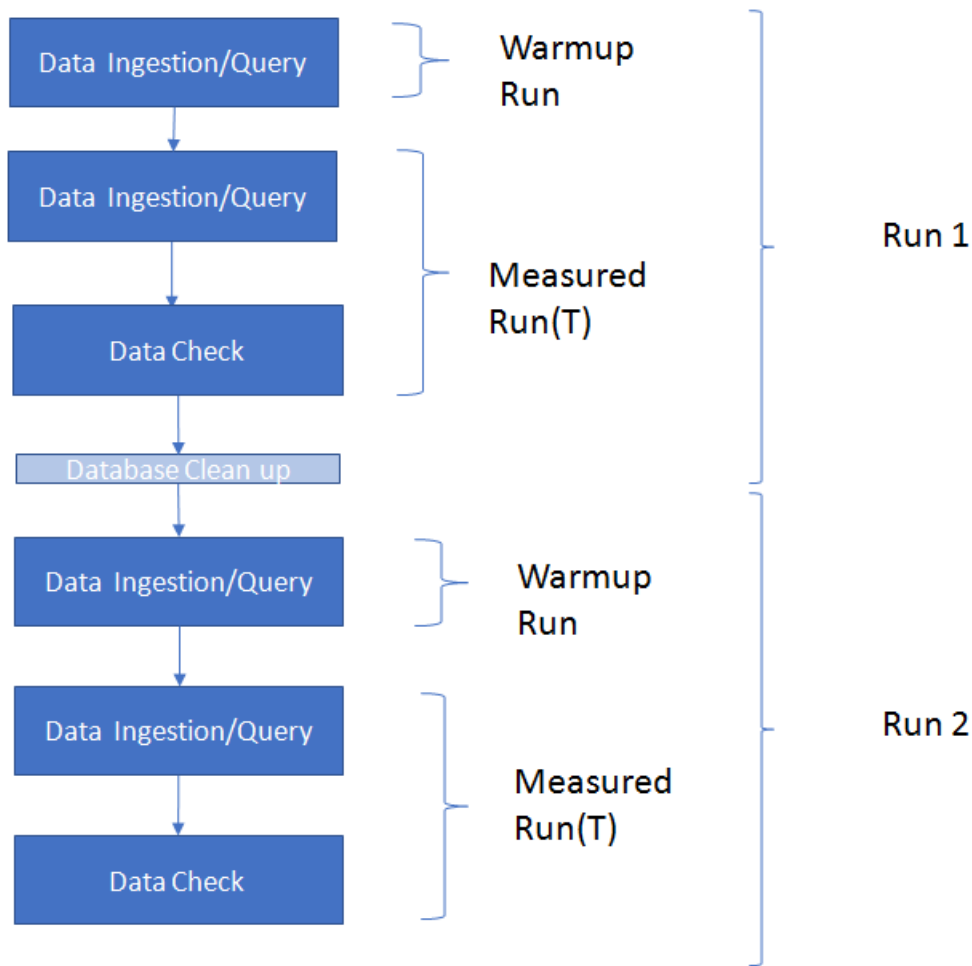


Figure 1: Benchmark Execution Phases

Comment: No part of the SUT and driver(s) 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 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.

2.5 NoSQL Databases supported by the Benchmark

The benchmark currently supports the following NoSQL databases.

- Hbase – 1.2.1
- Couchbase-server 5.0.0
- Machbase
- Lindorm
- TimechoDB

Additional support for new databases can be added by following the instructions in the user guide provided as part of the benchmark kit.

Clause 3: System Under Test and Benchmark Driver

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

3.1 System Under Test

The SUT is composed of those software and hardware components that are employed in the performance test and whose performance and cost are described by the benchmark metrics.

See Figure 2. Specifically, the SUT consists of:

- 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,

The benchmark driver(s) 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 not considered as 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.

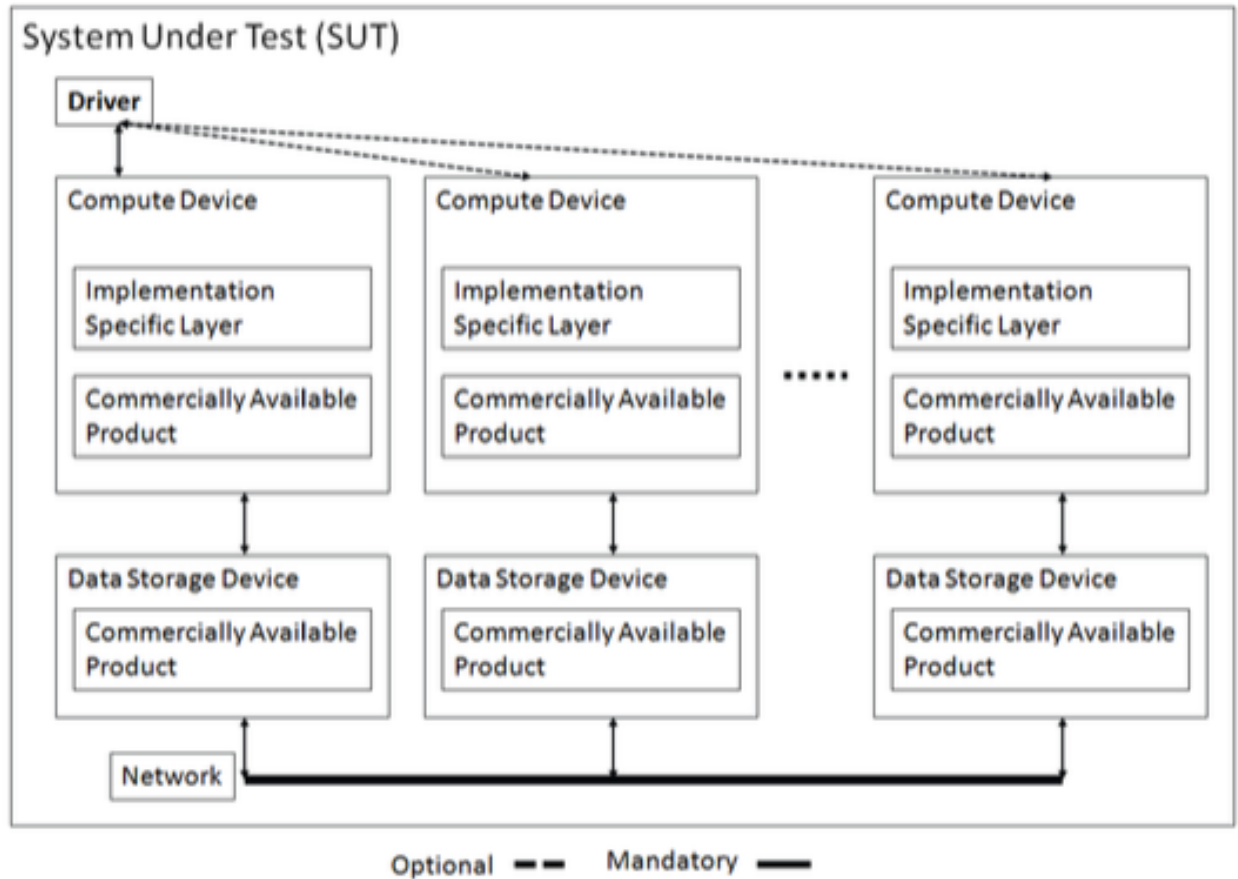


Figure 2: System Under Test (SUT)

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 included in the Full Disclosure Report. See Clause 6.

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 continuous Scale Factor. Scale Factor is defined as the number of records to be ingested. Benchmark sponsor can pick any Scale Factor.

4.2 Metric

TPCx-IoT defines the following primary metrics:

1. IoTps, the Performance Metric
2. \$/kIoTps , the Price-Performance metric
3. System availability date

4.3 Performance Metric

The **performance metric** represents the effective throughput capability of the SUT

$$\text{IoTps} = \text{SF}/(\text{T})$$

Where

SF is the Scale Factor. See Clause 4.1.

T is the time elapsed in seconds

4.4 Price Performance Metric

The **price-performance metric** is defined as:

$$\$/\text{kIoTps} = \frac{1000 * P}{\text{IoTps}}$$

P is the total cost of ownership of the SUT.

4.5 Availability Date

The Availability Date is defined in the TPC Pricing Specification located at the [TPC Documentation webpage](#).

4.6 Metric Comparison

A TPCx-IoT Result is only comparable with other TPCx-IoT Results.

4.7 Required Reporting Components

To be compliant with the TPCx-IoT Specification and TPC Policies, the URL to the benchmark result and Availability Date of the complete configuration must be included for all public

references. TPC Policies can be found at the [TPC Documentation webpage](#).

Clause 5: Pricing

Rules for pricing the Priced System and associated software and maintenance are included in the TPC Pricing Specification located at the [TPC Documentation webpage](#).

5.1 Priced System

The system to be priced shall include the software and hardware components present in the System Under Test (SUT) and maintenance.

Calculation of the priced system consists of:

- Price of the SUT as in Clause 3;
- Price of additional products (software or hardware) present in the system;
- 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 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 Full Disclosure Report 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 Pre-Certification Board. See Clause 7.

Clause 6: Full Disclosure Report and Executive Summary

6.1 Reporting Requirements

A Full Disclosure Report and Executive Summary in pdf are required as part of the benchmark submission. 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

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.

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

The Executive Summary must be included near the beginning of the Full Disclosure Report.

6.3 Full Disclosure Report

The Full Disclosure Report should be sufficient to allow an interested reader to evaluate and, if necessary, recreate an implementation of result. If any sections in the Full Disclosure Report 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 Full Disclosure Report.

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 Full Disclosure Report. 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 Specification (i.e., this document).

6.4 General Items

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

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.

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

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.

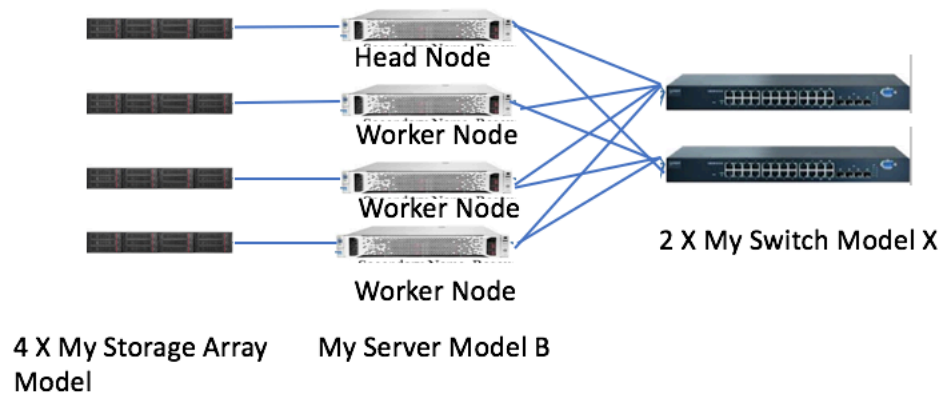


Figure 3: Sample Configuration Diagram

- 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.

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

6.5 Workload Related Items

Script or text used to set all hardware and software tunable parameters must be reported. The run report generated by the TPCx-IoT Kit for Performance Run and Repeatability Run must be reported.

6.6 Audit Related Items

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.

6.7 Executive Summary

The Executive Summary is meant to be a high-level overview of the implementation. It should provide the salient characteristics of a benchmark execution (metrics, configuration, pricing, etc.) without the exhaustive detail found in the Full Disclosure Report.

The executive summary has three components:

- Implementation Overview
- Pricing Spreadsheet

- Numerical Quantities

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.

6.8 Implementation Overview

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.

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	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.

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 (1	16-20 pt. Bold
TPCx-IoT Performance Metric	IoTps	0.01	16-20 pt. Bold
Price/Performance	\$/kIoTps	0.01	16-20 pt. Bold

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

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)

The fourth 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.
- 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 $8 \times 2 + 2 \times 1 = 18\text{RU}$

6.9 Pricing Spreadsheet

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

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

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

6.10 Numerical Quantities Summary

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.

The first section contains measurement results from the benchmark execution.

Table 5: Results for Measured Run

Item Title	Precision
Run Start Time	yyyy-mm-dd hh:mm:ss
Run End Time	yyyy-mm-dd hh:mm:ss
IoTps	0.01

Second section contains the measurement result for the repeatability run. See Table 5: for contents and precision.

6.11 TPCx-IoT Run Report

The run report from TPCx-IoT must be included in the Executive Summary.

6.12 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 Express Benchmark

IoT”, the Full Disclosure Report must be submitted electronically to the TPC using the procedure described in the TPC Policies located at the [TPC Documentation webpage](#). The Full Disclosure Report must be available in English but may be translated to additional languages.

6.13 Revisions to the Full Disclosure Report

Revisions to the Full Disclosure Report shall be handled as follows:

- Substitutions will be open to challenge for a 60-day period. No other portion of the Full Disclosure Report 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 IoTps rating of the system to decrease by more than two percent from the initially reported IoTps, 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 located at the [TPC Documentation webpage](#).

Clause 7: Audit

Rules for auditing Pricing information are included in the TPC Pricing Specification located at the [TPC Documentation webpage](#).

7.1 General Rules

The benchmark result must be certified by a TPC Certified Auditor or the TPCx-IoT Pre-Publication Board. See Section 10 of TPC Policies located at the [TPC Documentation webpage](#).

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 meet the following:

- The auditor holds an active TPC certification for a TPC enterprise benchmark or an express 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.

The Pre-Publication Board consists of three members of the TPCx-IoT committee. Each member serves a period of six 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.

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

- The Auditor or [Pre-Publication Board](#) ensure that the hardware and software products are the same as those used in the prior result;
- The Auditor or [Pre-Publication Board](#) review the F of the new results and ensures that they match what is contained in the original sponsor's Full Disclosure Report;
- The Auditor or [Pre-Publication Board](#) can attest to the validity of the pricing used in the new Full Disclosure Report.

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

Verify that the TPCx-IoT kit is used

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

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

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

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

7.2.3 Clause 4: Scale Factors and Metrics Related Items

Verify that the system is scaled as per the Specification

Verify that the metrics are reported as per the precision requirements

7.2.4 Clause 5: Pricing Related Items









Verify that the benchmark is in compliance with the TPC Pricing Specification

7.2.5 Clause 7: Full Disclosure Related Items

Verify that Full Disclosure Report and executive summary report are accurate and comply with the reporting requirements. This includes:

- 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. 2.0.0 TPC-Pricing Rev. 2.6.0 Report Date: May1, 2017	
Total System Cost		Performance Metric		Price / Performance	
\$99,996.13 USD		390,000.99 IoTps		\$256.40 USD \$ / kIoTps	
Database Management System		Operating System		Other Software	Availability Date
My NoSQL Software 1.0		My OS V2.0		None	May 1, 2017
System Configuration					
<div><div><div> 4 x My Storage Array Model A</div><div> 4 x My Server Model B</div><div> SAS 3.0</div><div></div><div> Model A</div><div> Model B</div></div><div> 10 GbE</div><div> 2 x My Switch Model X</div></div>					
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 Rack Unit		2x My Switch Model X 10GbE 10RU			

My Company Logo		My Server Model B			TPCx-IoT Rev. 2.0.0 TPC-Pricing Rev. 2.6.0 Report Date: May 1, 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:			\$99,996.13
Audited by My Auditor or Pre-Publication Board						
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.			IoTps:			390,000.99
			\$ / kIoTps:			\$256.40
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.						