

TPC Express AI

TPCx-AI

Specification

Version 2.0.0

August 2024

Transaction Processing Performance Council (TPC)

www.tpc.org

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, AND 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 is a trademark of the Transaction Processing Performance Council.

Product names, logos, brands, and other trademarks featured or referred to within this Specification are the property of their respective trademark holders.

Acknowledgments

The TPC acknowledges the work and contributions of the TPCx-AI subcommittee member companies: Cisco, Dell, HPE, IBM, Intel, Lenovo, Microsoft, RedHat, TTA, VMware.

TPC Membership

Find out more about the current members of the TPC in
http://www.tpc.org/TPC_Documents_Current_Versions/pdf/TPC_Membership.pdf

Document Revision History

Date	Version	Description
Sept 7, 2021	1.0.0	Officially Approved Standard Specification
January 25, 2022	1.0.1	Change #1 - Include definition for High Availability System (HAS) Change #2 - clause 1.4.1 Change the word form to "formats" Change #3 - All figure captions to include '-' Change #4 - clause 7.3.7.1 to reflect link to clause 7.3.1 Change #5 - change in Tpst definition wording in clause 7.5.6.7 Change #6 - clause 7.5.7 change in description of sub-metrics to calculate primary metric Change #7 - clause 7.3.11.3 to reflect additional streams up to 200 streams and the streams.yaml file to reflect up to 200 streams. Change #8 - Fix in table 7-a (UC3 metric, lower is better) Change #9 - Fix in table 3-x (replace #Financial Customer to #Financial Account) Change #10- Fix table C in Appendix F. Tensorflow version change to 2.2.0 Change #11 – Change clause 2.2.9 ORDER table (add trip_type column) Change#12 – Appendix C.7, add num_round parameter to UC8 parameters that can be changed. Change#13 – Changes to Appendix F Change#14 – validation test clause 7.2
March 8, 2022	1.0.2	Change#15 – Addition of clause 10.3.5, CHECK_INTEGRITY requirement for audit. Change#16 – Refinement of clause 7.2.7 to validate accuracy with lower sample sizes. Change#17 – Correct clause 10.2.2 to reflect the correct number of members in the pre-publication board.
July 7, 2022	1.0.3	Change#18 – Change in Clause 6.4.3, reflects data redundancy requirements for metadata.
August 20, 2024	2.0.0	Change in third party pricing for single node configuration

Typographic Conventions

The following typographic conventions are used in this specification:

Convention	Description
Bold	Bold type is used to highlight terms that are defined in this document
<i>Italics</i>	Italics type is used to highlight a variable that indicates some quantity whose value can be assigned in one place and referenced in many other places.
UPPERCASE	Uppercase letters indicate database schema object names such as table and column names. In addition, most acronyms are in uppercase.

TABLE OF CONTENTS

0	PREAMBLE	6
0.1	INTRODUCTION	6
0.2	TPCx-AI KIT AND LICENSING	6
0.3	GENERAL IMPLEMENTATION GUIDELINES.....	6
0.4	GENERAL MEASUREMENT GUIDELINES	7
0.5	DEFINITIONS	8
1	BUSINESS AND BENCHMARK MODEL	15
1.1	OVERVIEW	15
1.2	BUSINESS MODEL	16
1.3	DATA MODEL	16
1.4	TPCx-AI DATA SCIENCE PIPELINE	16
2	LOGICAL DATASET DESIGN	18
2.1	DATASET SCHEMA OVERVIEW	18
2.2	DATASET & TABLE DESCRIPTIONS	18
3	DATA SCALING & POPULATION	22
3.1	SCALING MODEL.....	22
3.2	DATA POPULATION	22
4	USE CASE SUMMARY	24
4.1	SUMMARY OF THE LOGICAL AI DATA SCIENCE PIPELINE.....	24
4.2	USE CASE 1 - CUSTOMER SEGMENTATION.....	24
4.3	USE CASE 2 – CUSTOMER CONVERSATION TRANSCRIPTION.....	24
4.4	USE CASE 3 – SALES FORECASTING	25
4.5	USE CASE 4 – SPAM DETECTION	26
4.6	USE CASE 5 – PRICE PREDICTION	26
4.7	USE CASE 6 – HARDWARE FAILURE	27
4.8	USE CASE 7 – PRODUCT RATING	27
4.9	USE CASE 8 – CLASSIFICATION OF TRIPS.....	28
4.10	USE CASE 9 – FACIAL RECOGNITION	28
4.11	USE CASE 10 – FRAUD DETECTION	29
5	BENCHMARK KIT CONTENT AND DEVELOPMENT	30
5.1	BENCHMARK KIT	30
5.2	KIT USAGE FOR A COMPLIANT RESULT	30
5.3	KIT RUN REPORT.....	30
5.4	KIT PARAMETER SETTINGS.....	30

5.5	BENCHMARK KIT	31
5.6	BENCHMARK KIT MODIFICATIONS.....	31
5.7	KIT VALIDATION.....	34
6	SYSTEM UNDER TEST (SUT).....	35
6.1	LOGICAL BREAKDOWN OF SYSTEM UNDER TEST.....	35
6.2	COMMERCIALY AVAILABLE PRODUCTS	36
6.3	BENCHMARK DRIVER & COMPUTE SOFTWARE AND LIBRARIES.....	36
6.4	DATA REDUNDANCY REQUIREMENT.....	36
7	EXECUTION RULES AND PERFORMANCE METRICS.....	39
7.1	BENCHMARK EXECUTION	39
7.2	VALIDATION TEST	39
7.3	BENCHMARK RUN.....	40
7.4	CONFIGURATION AND TUNING	45
7.5	METRICS	45
8	PRICING	50
8.1	INTRODUCTION	50
8.2	PRICED CONFIGURATION.....	50
8.3	SPECIFICALLY EXCLUDED FROM THE PRICED CONFIGURATION CALCULATION ARE:	50
8.4	ADDITIONAL OPERATIONAL COMPONENTS	50
8.5	ALLOWABLE SUBSTITUTIONS.....	51
9	FULL DISCLOSURE	52
9.1	FULL DISCLOSURE REPORT REQUIREMENTS	52
9.2	FORMAT GUIDELINES.....	52
9.3	GENERAL ITEMS.....	52
9.4	SOFTWARE COMPONENTS AND DATASET DISTRIBUTION	54
9.5	WORKLOAD RELATED ITEMS	56
9.6	SUT RELATED ITEMS.....	56
9.7	METRICS AND SCALE FACTORS.....	57
9.8	AUDIT RELATED ITEMS.....	57
9.9	AVAILABILITY OF THE FULL DISCLOSURE REPORT	61
9.10	REVISIONS TO THE FULL DISCLOSURE REPORT	61
10	AUDIT.....	62
10.1	GENERAL RULES.....	62
10.2	INDEPENDENT AUDIT	62
10.3	AUDIT CHECKLIST	63
APPENDIX A:	PDGF USER GUIDE	65
APPENDIX B:	BENCHMARK KIT PARAMETERS.....	66
APPENDIX C:	USE CASE SPECIFIC PARAMETERS	67
APPENDIX D:	SAMPLE DEFAULT CONFIGURATION FILE	69
APPENDIX E:	THROUGHPUT TEST STREAM PLACEMENT.....	81
APPENDIX F:	MINIMUM SOFTWARE LIBRARY VERSIONS FOR THE SUT	83

0 PREAMBLE

0.1 Introduction

Artificial intelligence (AI) has become a key transformational technology of our times. Advances in neural networks and other **machine learning** techniques have made it possible to use AI on a variety of **use cases**. From the public sector to aerospace, defense and academia, new and improved ways to use AI techniques are changing the way we harness data and analytics. This along with advances in compute, interconnect and memory technologies have made possible to solve complicated challenges that will ultimately benefit customers in production datacenter and cloud environments.

Abundant volumes of rich data from text, images, audio and video are the essential starting point for creating a benchmark that would represent the myriad of **use cases** and customers. TPC Express Benchmark AI (TPCx-AI) is created in keeping with the TPC tradition of emulating real world AI scenarios and data science **use cases**. Unlike most other AI benchmarks, the TPCx-AI uses a diverse dataset and is able to scale across a wide range of scale factors. TPCx-AI may later expand with additional **use cases** and add additional flexibility for a greater variety of implementations.

The benchmark defines and provides a means to evaluate the **System Under Test (SUT)** performance as a general-purpose data science system that:

- Generates and processes large volumes of data.
- Trains preprocessed data to produce realistic **machine learning** models
- Conducts accurate insights for real-world customer scenarios based on the generated models
- Can scale to large scale distributed configurations
- Allows for flexibility in configuration changes to meet the demands of the dynamic AI landscape.

The benchmark models real-life examples of companies and public-sector organizations that use a range of analytics techniques, both AI and more traditional **machine learning** approaches, as well as the potential application of these techniques in situations like those in which they have already been successfully deployed. In addition, the benchmark measures end to end time to provide insights for individual **use cases**, as well as throughput metrics to simulate multiuser environments for a given hardware, operating system, and data processing system configuration under a controlled, complex, multi-user AI or **machine learning** data science workload.

Comment: While separated from the main text for readability, comments and appendices are a part of the standard and their provisions must be enforced.

0.2 TPCx-AI Kit and Licensing

The TPCx-AI kit is available from the TPC website (see www.tpc.org/tpcx-ai/ for more information). Users must sign-up and agree to the TPCx-AI End User Licensing Agreement (EULA) to download the kit. All related work (such as collaterals, papers, derivatives) must acknowledge the TPC and include the TPCx-AI copyright. The TPCx-AI kit includes: TPCx-AI Specification document (this document), TPCx-AI Users Guide (README.md) documentation, scripts to set up the benchmark environment, code to execute the benchmark workload, Data Generator, **use case** related files, and Benchmark Driver.

0.3 General Implementation Guidelines

The purpose of TPC benchmarks is to provide relevant, objective performance data to industry users. To achieve that purpose, TPC benchmark specifications require benchmark runs 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 (e.g., TPCx-AI models and represents complex, high data volume, decision support environments).
- Would plausibly be implemented by a significant number of users in the market segment modeled or represented by the benchmark.

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, which improve benchmark **Results** but not real-world performance or pricing, are prohibited.

Several characteristics shall be evaluated 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 this implementation is a **benchmark special**?"

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

- a) Is the implementation generally available, documented, and supported?
- b) Does the implementation have significant restrictions on its use or applicability that limits its use beyond TPC benchmarks?
- c) Is the implementation or part of the implementation poorly integrated into the larger product?
- d) Does the implementation take special advantage of the limited nature of TPC benchmarks (e.g., model selection, accuracy requirements, query templates etc.) in a manner that would not be generally applicable to the environment the benchmark represents?
- e) Is the use of the implementation discouraged by the vendor? (This includes failing to promote the implementation in a manner like other products and technologies.)
- f) Does the implementation require uncommon sophistication on the part of the end-user, programmer, or system administrator?
- g) Is the pricing unusual or non-customary for the vendor or unusual or non-customary compared to normal business practices? The following pricing practices are suspect:
 - Availability of a discount to a small subset of possible customers.
 - Discounts documented in an unusual or non-customary manner.
 - Discounts that exceed 25% on small quantities and 50% on large quantities.
 - Pricing featured as a close-out or one-time special.
 - Unusual or non-customary restrictions on transferability of product, warranty, or maintenance on discounted items.
- h) Is the implementation (including beta-release components) being purchased or used for applications in the market segment 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?

0.4 General Measurement Guidelines

TPCx-AI **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 implemented by the TPCx-AI Kit (Clause 5.1). 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 **Results**.
- The equipment used in measuring **Results** must conform to the requirements in Clause 6
- Fidelity and candor are maintained in reporting any anomalies in the **Results**, even if not specified in the benchmark requirements.

Comment: The use of new methodologies and approaches is encouraged so long as they meet the requirements above.

0.5 Definitions

A _____

- **Attestation Letter**

TPC-Certified **Auditor's** opinion regarding the compliance of a **Result** must be consigned in an **Attestation Letter** delivered directly to the Test Sponsor.

- **Availability Date**

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

B _____

- **Benchmark Special**

The **Benchmark Special** is defined as any aspect of the benchmark implementation with the primary purpose of the optimization of TPC Benchmark **Results** without any corresponding applicability to real-world applications and environments.

C _____

- **Commercially Available Product**

Commercially Available Product is defined in TPC Pricing Specification.

- **Compute Software**

Compute software runs on the **System Under Test (SUT)** hardware providing required software capabilities to successfully execute the benchmark. **Compute Software** unless otherwise stated will be part of the **System Under Test**.

D _____

- **Data Redundancy**

The ability to have no permanent data loss after the permanent irrecoverable failure of any single Durable Medium containing tables, input data, output data, or **metadata**.

- **Data Generation**

The process of using **PDGF** to create the data in a format suitable for presentation to the load facility.

- **Data Node**

Data Nodes store data in a Hadoop cluster and is the name of the daemon that manages the data. File data is replicated on multiple **Data Nodes** for reliability and so that localized computation can be executed near the data¹.

- **Deep Learning**

¹ https://docs.cloudera.com/documentation/enterprise/6/6.3/topics/cm_mc_dn.html

Deep learning (also known as deep structured learning) is part of a broader family of **machine learning** methods based on **artificial neural networks** with **representation learning**².

E _____

- **Executive Summary**

Defined by the TPC Policies, an **Executive Summary** is a two to four-page summary of the **Result**.

F _____

- **F-score**

The **F-score** or **F-measure** is a measure of a test's accuracy. It is calculated from the **precision** and **recall** of the test, where the precision is the number of correctly identified positive results divided by the number of all positive results, including those not identified correctly, and the recall is the number of correctly identified positive results divided by the number of all samples that should have been identified as positive³.

- **Full Disclosure Report (FDR)**

The **Full Disclosure Report** is a set of files that documents how a benchmark **Result** was implemented and executed in sufficient detail so that the **Result** can be reproduced given the appropriate hardware and software products.

- **Framework**

A Framework is a collection of software including API's, distributed computing engines, **Machine learning**, AI and libraries used to run TPCx-AI.

G _____

H _____

- **HDFS**

HDFS (Hadoop Distributed File System) is a file system that provides scalable and reliable data storage, and it was designed to span large clusters of commodity servers.

- **High Availability System (HAS)**

Computing environments configured to provide nearly full-time availability are known as **High Availability Systems**. Such systems typically have redundant hardware and software that makes the system available despite failures. Well-designed **high availability systems** avoid having single points-of-failure. Any hardware or software component that can fail has a redundant component of the same type.

I _____

J _____

² https://en.wikipedia.org/wiki/Deep_learning

³ <https://en.wikipedia.org/wiki/F-score>

- **JBOD**

JBOD (Just a Bunch of Disks) refers to a collection of hard disks that have not been configured to act as a redundant array of independent disks (RAID) array.

K _____

- **k-means clustering**

k-means clustering is a method of vector quantization, originally from signal processing, that aims to partition n observations into k clusters in which each observation belongs to the cluster with the nearest mean (cluster centers or cluster centroid), serving as a prototype of the cluster.⁴

L _____

- **LCS (Licensed Compute Services)**

Licensed Compute Service (LCS) is defined in TPC Pricing Specification: Publicly offered processing, storage, network, and software services that are hosted on remote computer servers accessed via a Wide Area Network (e.g. the Internet). A Customer pays a license fee to the **Licensed Compute Services** vendor for the use of the processing, storage, network, and software services. The **Licensed Compute Services** are not located or installed on a Customer's premises.

M _____

- **Machine Learning**

Machine learning (ML) is the study of computer algorithms that improve automatically through experience and by the use of data. It is seen as a part of artificial intelligence. **Machine learning** algorithms build a model based on sample data, known as "training data", in order to make predictions or decisions without being explicitly programmed to do so. **Machine learning** algorithms are used in a wide variety of applications, such as email filtering and computer vision, where it is difficult or unfeasible to develop conventional algorithms to perform the needed tasks.⁵

- **Measured Configuration**

See **System Under Test (SUT)**

- **Metastore/Metadata**

Descriptive information about the dataset including names and definitions of tables, indexes, and other schema objects. Various terms commonly used to refer collectively to the **Metadata** include Metastore, information schema, data dictionary, or system catalog. **Metadata** also includes additional information stored in the dataset to define, manage and use other objects, e.g. users, connections, etc.

- **Master Node**

Master Node(s) provide a variety of storage and processing coordination services for a cluster. These are conceptually distinct from **Data Nodes** but sometimes share physical hardware. Where necessary for correct execution of the benchmark, data in the **Master Node** services is considered **Metadata** unless it is data for a co-located **Data Node** service in which case it is considered table data. Some **Master Node** services can be configured with sufficient instances as part of a **High Availability System** while others require other approaches to protecting against loss of service or data

⁴ [k-means clustering - Wikipedia](#)

⁵ https://en.wikipedia.org/wiki/Machine_learning

loss. Examples of **Master Node** services include **Name Nodes**, Checkpoint Nodes, Journal Nodes and Resource Manager

- **Metadata**

Metadata is "data that provides information about other data". **Metadata** can be stored and managed in a database, often called a **Metadata registry** or **Metadata repository**.⁶

N _____

- **Name Node**

A **Name Node** is a particular class of **Master Node** service. ***Name Nodes** maintain the namespace tree for **HDFS** and a mapping of file blocks to **Data Nodes** where the data is stored. A simple **HDFS** cluster can have only one primary **Name Node**, supported by a secondary **Name Node***⁷

O _____

- **Operating System/OS**

The term **Operating System** refers to a commercially available program that, after being initially loaded into the computer by a boot program, manages all the other programs in a computer, or in a VM. The **Operating System** provides a software platform on top of which all other programs run. Without the **Operating System** and the core services that it provides no other programs can run and the computer would be non-functional. Other programs make use of the **Operating System** by making requests for services through a defined application program interface (API). All major computer platforms require an **Operating System**. The functions and services supplied by an **Operating System** include but are not limited to the following:

- manages a dedicated set of processor and memory resources
- maintains and manages a file system
- loads applications into memory
- ensures that the resources allocated to one application are not used by another application in an unauthorized manner

determines which applications should run in what order, and how much time should be allowed to run the application before giving another application a turn to use the systems resources

manages the sharing of internal memory among multiple applications

handles input and output to and from attached hardware devices such as hard disks, network interface cards, add-on cards and other hardware devices.

Some examples of **Operating Systems** are listed below:

- Windows
- Unix (Solaris, AIX)
- Linux (Red Hat, SUSE)
- Mac OS

P _____

⁶ <https://en.wikipedia.org/wiki/Metadata>

⁷ https://docs.cloudera.com/documentation/enterprise/6/6.3/topics/cm_mc_nn.html

- **PDGF**

The **PDGF** (Parallel Data Generator Framework) is part of TPCx-AI kit used to generate the **Test Dataset**.

- **Performance Metric**

The reported throughput as expressed in **AI use cases** per minute.

- **Performance Test**

The **Performance Test** is defined as the run following the validation test that consists of the 6 tests running in sequential order (Load test, power training test, power serving test I, power serving test II, scoring test and throughput test)

- **Priced Configuration**

The **Priced Configuration** consists of components defined in the TPCx-AI Benchmark Standard including all hardware, software and maintenance.

- **Price/Performance Metric**

The Price/Performance Metric is the total price of the **Priced Configuration** divided by the TPCx-AI Performance Metric.

Q _____

R _____

- **Report**

The **Report** is an Adobe Acrobat PDF file in the **FDR**. The contents of the **Report** are defined in Clause 5.3

- **Reported**

The term Reported an item that is part of the **FDR**.

- **Result**

[The following definition has been copied from the TPC Policies. Please refer to the latest revision of the TPC Policies for the most up-to-date text as well as definitions of defined terms.]

A **performance test** submitted to the **TPC** attested to meet the requirements of a **TPC Benchmark Standard** at the time of submission. A **Result** is documented by an **Executive Summary** and, if required, an **FDR**.

S _____

- **Scoring**

The phase where the accuracy metric of the generated model is calculated for a given **use case**

- **Serving**

The phase where the generated model is used to predict results for a given **use case**

- **Software Version**

A **Software Version** uniquely identifies a software product, its release level, update level, and/or patch level. It is typically a string of alphanumeric characters that allows the software manufacturer to uniquely identify the software.

- **Substitution**

Substitution is the use of components in the **Priced Configuration** which are different than those used in the measured configuration.

- **Supporting Files**

Supporting Files refers to the contents of the **Supporting Files** folder in the **FDR**. The contents of this folder, consisting of various source files, scripts, and listing files, are defined in Clause 9.1.1.

- **System Under Test (SUT)**

System Under Test (SUT) – is defined to be the sum of the components utilized in running a benchmark as specified in Clause 6.

T _____

- **Test Sponsor**

The **Test Sponsor** is the company officially submitting the **Result** with the **FDR** and will be charged the filing fee. Although multiple companies may sponsor a **Result** together, for the purposes of the TPC's processes the **Test Sponsor** must be a single company. A **Test Sponsor** need not be a TPC member. The **Test Sponsor** is responsible for maintaining the **FDR** with any necessary updates or corrections. The **Test Sponsor** is also the name used to identify the **Result**.

- **Test Dataset**

The **Test Dataset** is the dataset generated by **PDGF** for the defined scale factor used to execute the Load test, Power tests and Throughput test.

- **TPCx-AI approved Compute software**

Compute software and libraries that are not commercially available that cannot meet the pricing requirements that the TPCx-AI subcommittee has approved to be used as part of the benchmark kit configuration for a valid publication.

- **TPC-Certified Auditor (Auditor)**

The term **TPC-Certified Auditor** is used to indicate that the TPC has reviewed the qualification of the **Auditor** and has certified his/her ability to verify that benchmark **Results** are in compliance with a specification. (Additional details regarding the **Auditor** certification process and the audit process can be found in Section 9 of the TPC Policies document.)

- **Training**

The phase where the generated model is used to predict results for a given **use case**

U _____

- **Undo/Redo Log**

Undo/Redo Log: records all changes made in data files. The Undo/Redo Log makes it possible to replay all the actions executed by typical Big Data analytics systems as well as traditional Relational Databases Management Systems. If something happens to one of the data files, a backed up data file can be restored and the Undo/Redo Log that was written since the backup can be played and applied which brings the data file to the state it had before it became unavailable. Not all Big Data environments utilize an Undo/Redo Log to accommodate recovery.

- **Use Case**

We define a “**use case (s)**” as a targeted end to end pipeline of processes and technologies applied to a specific business challenge, with a measurable outcome. We recognize that **use cases** can be described at different levels of granularity. The **use cases** we chose for this benchmark correspond to descriptions of specific business challenges which industry experts acknowledged to the TPC as meaningful. For example, recommending the “next product to buy” or “detecting a fraudulent transaction” for e-commerce in the retail industry were considered examples of very important **use cases**. Higher level functions like “marketing” or “sales” were not sufficiently granular for the TPC to be considered **use cases**, even though they have relevance in the overall business development process.

A **Use Case** defines a single problem solved by the AI and **Machine learning** Data Science Pipeline. It is Framework and syntax agnostic and can be implemented in many ways. In the TPCx-AI kit, all **Use Cases** implemented include data generation, data management, training, scoring and serving phases.

V _____

W _____

X _____

Y _____

Z _____

1 BUSINESS AND BENCHMARK MODEL

1.1 Overview

TPC Express Benchmark AI (TPCx-AI) contains benchmark components that can be used to assess a broad range of system topologies and implementation methodologies in a technically rigorous and directly comparable, vendor-neutral manner. The benchmark has been mapped to typical retail businesses and on-premise, cloud and edge environments. This clause outlines the business modeling assumptions that were adopted during the development of the benchmark, and their impact on the benchmarking environments.

TPCx-AI models the end to end AI and **machine learning** data science system of a retail business datacenter. The supporting schema contains vital business information, such as customer, order, financial and product data. The benchmark models the two most important components of any mature data science system:

- Data aggregation and data management, which converts the data to relevant **Test datasets** for processing
- Insights, which transforms the relevant data sets into accurate business intelligence

The benchmark abstracts the diversity of operations in data science pipelines, while retaining essential performance characteristics.

In a retail datacenter scenario, it is necessary to execute a diverse number of AI or **machine learning** data science **use cases** based on various departments that may or may not be related within the datacenter. This poses an additional challenge of trying to benchmark scenarios and business intelligence pipelines that would benefit the end user. As a result TPCx-AI comprises of many independent AI and **machine learning use cases** that facilitate in helping any retail business datacenter address and manage any business analysis environment.

While TPCx-AI does not aspire to be a model of how to build actual AI or **machine learning** data science pipelines, the benchmark has been granted a realistic context. It imitates the activity of retail businesses and datacenters with customer information, department stores, sales and financial data, product catalog and reviews, emails, datacenter logs as well as facial images and audio conversations. In addition, it is quite possible that not all scenarios in TPCx-AI will be applicable to all users. To this end, in defining the **Performance Metric** for TPCx-AI, the TPC made sure that no one scenario dominates the **Performance Metric** so that all scenarios receive attention by the test sponsors.

The goal of selecting a retail business model is to assist the reader in relating intuitively to the components of the benchmark, without tracking that industry segment so tightly as to minimize the relevance of the benchmark. The TPCx-AI benchmark be used to characterize any industry that must transform operational and external data into business intelligence. The data represents a reasonable image of a business operation as they progress over time.

The TPCx-AI benchmark models the challenges of end to end artificial intelligence systems and pipelines where the power of **machine learning** and **deep learning** is used to: detect anomalies (fraud and failures), drive AI based logistics optimizations to reduce costs through real-time forecasts (classification, clustering, forecasting and prediction) and use **deep learning** AI techniques for customer service management and personalized marketing (face recognition and speech recognition).




Figure 1-a Benchmark Use Case pipeline flow

1.2 Business model

The TPCx-AI benchmark models the following **use cases**:

- Use Case 1 - Find customer segments based on their behavior. Clustering/segmentation of customers based on return behavior (return frequency, return/order ratio, ...) and buying behavior (frequency of purchases, recency of purchases, ...)
- Use Case 2 - Accurately transcript audio conversations of customers to text
- Use Case 3 - Forecast the weekly sales for each store department and each store of a retail chain with multiple stores based on a limited history of sales data.
- Use Case 4 - Find comments, reviews, or descriptions of items in a retail business with spam content. The problem to be solved is to identify those reviews that are spam
- Use Case 5 - Suggest a price of an item based on its brand, product name, and description on an online marketplace.
- Use Case 6- Predict imminent hardware failure, given existing logs of hardware events.
- Use Case 7 - Improve cross-selling by giving "next-product-to-buy" recommendations. Based on previously bought products recommend products that the customer might also be interested in. Those recommendations are found by comparing customers (by their products) and/ or products (by their customers)
- Use Case 8 - Classifying categories and trip types using data from existing customer shopping trips
- Use Case 9 - Accurately recognize customer facial images
- Use Case 10 - Detect if a given financial transaction is fraudulent or not.

1.3 Data model

The TPCx-AI benchmark models a unified dataset schema of a retail datacenter with a full **AI** and **machine learning data science pipeline**. This environment allows a single as well as concurrent data science pipelines to be run at any given point in time. The data rests in persistent storage so is always available. The TPCx-AI dataset represents a system that contains structured as well as unstructured data which is typical of modern data systems. Some examples of structured data are product, orders, customers, financial transactions tables. Customer images and audio conversations make up the unstructured dataset.

The Data generation phase of the benchmark scales each of the tables and unstructured datasets to simulate the growth of the datacenter data sizes. The hope of the benchmark is to mimic datasets of different company sizes that each would have the optimal data science pipelines (benchmark scaling or scale factor). Once generated, all the data is processed for subsequent stages of post-processing within the **machine learning** data science pipeline.

1.4 TPCx-AI Data Science Pipeline

1.4.1 Data Science Pipeline Assumptions

The **data management** stages, and data science pipeline modeled by the benchmark exhibit the following characteristics:

- a) They address complex business problems with the intent to answer specific questions
- b) The data is acquired from different sources and contains different formats
- c) The data management stages like cleansing, exploration and preprocessing mimic modern commercial pipelines used in current production environments
- d) They employ **training**, **serving** and **scoring** phases using production datasets available in the datacenter

To address the enormous range of **use cases** prevalent in today's modern data science pipelines, TPCx-AI utilizes a generalized data science pipeline model. This model allows the benchmark to capture important aspects of the pipeline where data scientists would train on preprocessed data, develop a production model, perform a scoring analysis to generate accuracy metrics and then use the built model in the serving phases with production quality data.

1.4.2 Data Generation

This phase of the benchmark models data acquisition wherein a built-in data generator (**PDGF**) provided by the TPC is used to generate ‘relevant’ data that will be used to answer specific business related questions the retail datacenter is trying to answer. The data is a combination of structured, and unstructured data formats.

1.4.3 Data Loading

Data loading takes existing raw data generated by the TPC provided **PDGF** data generator and loads unstructured data to persistent storage and structured data into the schema formats to persistent storage (e.g. **HDFS**, or other file systems). Some data science pipelines also call this stage Data Wrangling.

1.4.4 Cleansing and Transformation

Ensuring data is cleaned or cleansed can offer significant business value. Numerous surveys have found that many large enterprises do not use data effectively due to the complexity of the data and redundancies thereof. Cleaning or Cleansing the data can significantly help achieve a long list of benefits including accurate insights, increased productivity, and efficient pipeline execution. Examples of cleansing are included but not limited to removal of null records and duplicate records from the **Test dataset**.

1.4.5 Data preprocessing and labeling

This phase involves deep study of the data, columns or patterns and inter-relationships between the entities and variables. This analysis also helps in eliminating irrelevant data and narrowing down on the key parameters needed to build an effective model. Examples of preprocessing include normalizing and standardizing the cleansed data and labeling for getting the **Test dataset** ready for building the model.

1.4.6 Training and Model Generation

Training usually results in the generation of a model. This phase involves many stages like choosing one or many **machine learning** or AI methods to choose from, generating models and re-iterating the processing by changing hyper parameters to build the final model. As in real world scenarios, the TPCx-AI kit models the training phase using one representative **machine learning** or AI model that will meet key accuracy criteria set forth by the benchmark. The benchmark assumes that this model will be deployed in production by the datacenter to be used for subsequent serving phases. Each unique use case will represent its own unique model training phase in the benchmark.

1.4.7 Scoring

This stage may also be known as the model validation stage. Scoring in the benchmark involves taking the generated model and validating its accuracy with a labeled dataset (new dataset generated by the **PDGF** data generator tool) not used for training so that the accuracy of the model can be determined. Each **use case** pipeline in the benchmark models its unique scoring phase. The benchmark also defines accuracy metrics and criteria that need to be met to have a successful **run**.

1.4.8 Model Deployment and Serving

Real world production model deployment and **machine learning** inference is represented by the TPCx-AI benchmark by deploying only one production model per **use case**. The TPCx-AI benchmark uses the model generated in the training phases to use conduct the serving phase for each of the respective **use cases**. For each **use case**, a new dataset is generated especially for the serving phase to represent production data that needs quality business insights.

1.4.9 Business insights and Metrics

TPCx-AI metrics demonstrate the value obtained using the end to end data science pipeline execution. The results from the benchmark model the powerful use of insights that can be used to grow the retail business, predict failure rates or anomalies, accurately conduct classification or marketing personalization.

2 LOGICAL DATASET DESIGN

2.1 Dataset Schema Overview

The TPCx-AI dataset comprises of the retail datacenter for an organization that contains 14 structured tables and 2 unstructured datasets. Figure 2-a shows the structured and unstructured datasets, the high-level definitions for each table and its relationship to other tables in the form of a relationship diagram.

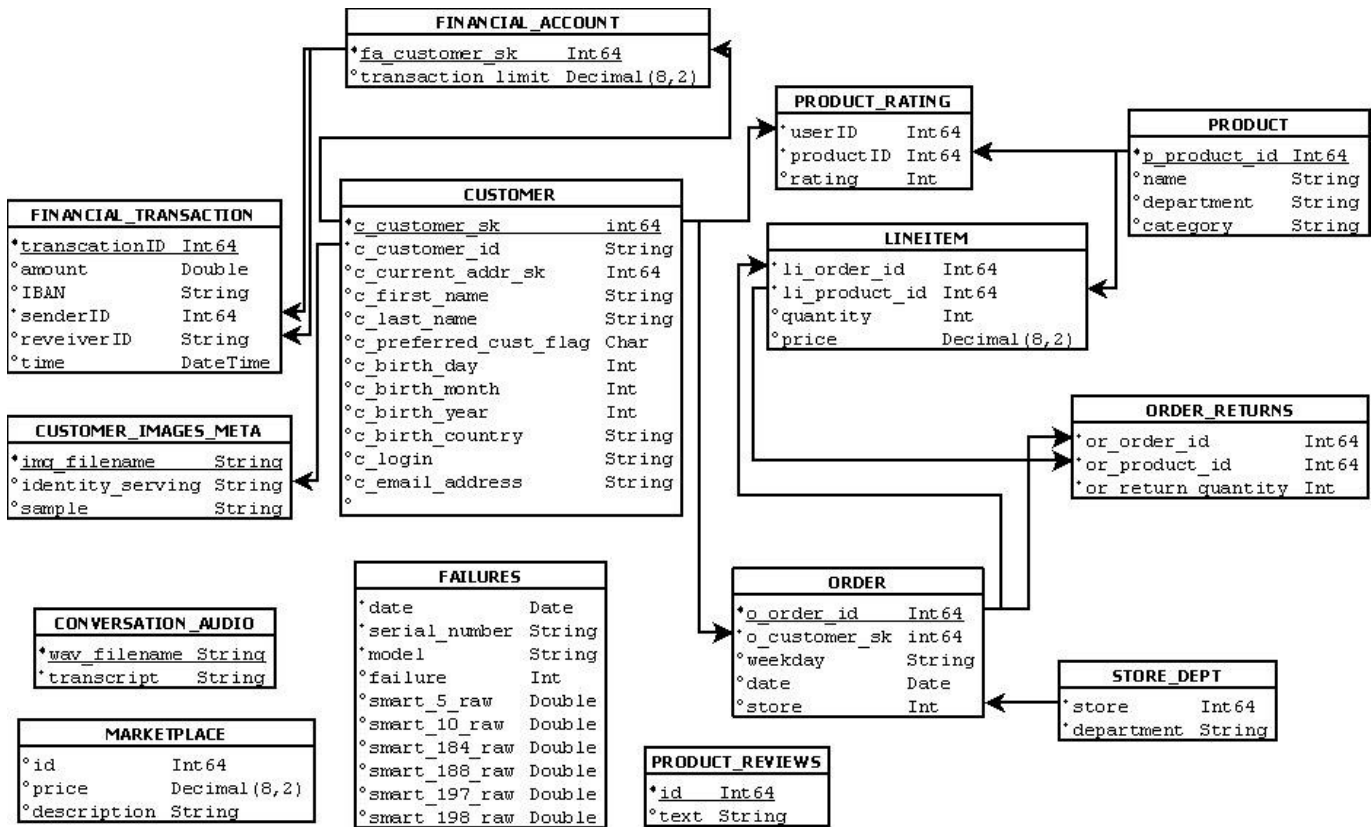


Figure 2-a – Training Dataset Schema

2.2 Dataset & Table Descriptions

2.2.1 Table description

2.2.1.1 Lists the high-level dataset or table generated for the power training test.

2.2.2 Column Definitions

2.2.2.1 Column Name

Each column is uniquely named. Columns that are part of the table's primary key are indicated in the column called Primary Key. If a table uses a composite primary key, then for convenience of reading the order of a given column in a table's primary key is listed in parentheses following the column name.

2.2.2.2 Datatype

Each column employs one of the following datatypes as generated by **PDGF**:

- a) 'Integer' means that the column shall be able to exactly represent INTEGER values (i.e., values in increments of 1) in the range of at least (-2^{n-1}) to $(2^{n-1} - 1)$, where n is 64.

- b) Decimal(d, f) means that the column shall be able to represent decimal values up to and including d digits, of which f shall occur to the right of the decimal place; the values can be either represented exactly or interpreted to be in this range.
- c) Char(N) means that the column shall be able to hold any string of characters of a fixed length of N.
- d) 'date' means that the column shall be able to express any calendar day between January 1, 1900 and December 31, 2199.
- e) 'string' is used to represent text rather than numbers. It is comprised of a set of characters that can also contain spaces and numbers.
- f) 'double' means the column shall be able to exactly represent a double-precision, 64-bit floating-point value.

The datatypes do not correspond to any specific SQL-standard datatype. The definitions are provided to highlight the properties that are required for a particular column.

2.2.2.3 Description

Description of each column as it relates to other columns within the training tables & training dataset.

2.2.3 CUSTOMER table

Each row in this table represents a unique customer.

Column	Datatype	Relations	Description
c_customer_sk	integer	PK	Primary key for customer
c_customer_id	string		Customer ID
c_current_addr_sk	integer		Foreign key to customer address
c_first_name	string		Customer first name
c_last_name	string		Customer last name
c_preferred_cust_flag	char(1)		Is this a preferred customer?
c_birth_day	integer		Day of the month 1...31
c_birth_month	integer		Month of year 1...12
c_birth_year	integer		Year 1930...2002
c_birth_country	string		Customer birth country
c_login	string		Login name
c_email_address	string		Well-formed email address

Table 2-a Customer Table Column Definitions

2.2.4 FAILURES table

Each row in this table represents the details of specific hardware failures within a datacenter

Column	Datatype	Relations	Description
date	date		Date of the event (yyyy-MM-dd)
serial_number	string		Disk drive serial number
model	string		Disk drive model number
failure	integer		Indicates a failure for a particular hard drive for a given date.
smart_5_raw	double		Parameter value for smart_5
smart_10_raw	double		Parameter value for smart_10
smart_184_raw	double		Parameter value for smart_184
smart_188_raw	double		Parameter value for smart_188
smart_197_raw	double		Parameter value for smart_197
smart_198_raw	double		Parameter value for smart_198

Table 2-b Table Failure Table Column Definitions

2.2.5 FINANCIAL_ACCOUNT table

Each row in this table represents a financial account customer with respective transaction limits

Column	Datatype	Relations	Description
fa_customer_sk	integer	PK FK(customer_)	Foreign key to customer
transaction limit	decimal(8,2)		Transaction limit set for the given customer

Table 2-c Financial Account Table Column Definitions

2.2.6 FINANCIAL_TRANSACTIONS table

Each row in this table represents financial transaction details for a given time.

Column	Datatype	Relations	Description
transactionID	integer	PK	Transaction ID
amount	double		Monetary amount that was sent
IBAN	string		Imitation of an IBAN
senderID	integer	FK(customer_)	Foreign key to financial account
receiverID	string	FK(customer_)	Foreign key to financial account or random number
time	datetime		yyyy-MM-dd T hh:mm

Table 2-a Financial Transactions Table Column Definitions

2.2.7 LINEITEM table

Each row in this table represents a single Lineitem for a particular order

Column	Datatype	Relations	Description
li_order_id	integer	FK(order_id)	Foreign key to order
li_product_id	integer	FK(product_id)	Foreign key to product
quantity	integer		Number of items
price	Decimal(8,2)		Price paid per unit

Table 2-b Lineitem Table Column Definitions

2.2.8 MARKETPLACE table

Each row in this table represents the description of each marketplace.

Column	Datatype	Relations	Description
Id	integer		Record ID
price	decimal (8,2)		Price of item
description	string		Product description (may include branch and product name)

Table 2-c Marketplace Table Column Definitions

2.2.9 ORDER table

Each row in this table represents a single Lineitem for details of an order

Column	Datatype	Relations	Description
o_order_id	integer	PK	Unique ID for each order
o_customer_sk	integer	FK(customer_)	Foreign key to customer
weekday	string		Day of week (Monday, Tuesday, etc.)
date	date		yyyy-MM-dd
store	integer		Foreign key to store_department table
trip_type	integer		Type of shopping trip to a store that a customer makes

Table 2-d Order Table Column Definitions

2.2.10 ORDER_RETURNS table

Each row in this table represents a single line item for an item returned

Column	Datatype	Relations	Description
or_order_id	integer	FK(order_id)	Foreign key to order
or_product_id	integer	FK(product_id)	Foreign key to product
or_return_quantity	integer		Number of units returned (at most the order quantity)

Table 2-e Order_returns Table Column Definitions

2.2.11 PRODUCT table

Each row in this table represents a unique product id and its location

Column	Datatype	Relations	Description
p_product_id	integer	PK	Unique ID for each product
name	string		Arbitrary product name
department	string		Department name

Table 2-f Product Table Column Definitions

2.2.12 PRODUCT_RATING table

Each row in this table represents a line item for the rating of a product

Column	Datatype	Relations	Description
userID	integer	FK(customer_)	Represents the customer id that rated the product
productID	integer	FK(product_)	Product ID for which the rating was given
Rating	integer		Rating for each product

Table 2-g Product Table Column Definitions

2.2.13 PRODUCT_REVIEWS table

Each row in this table represents a single line item for product review.

Column	Datatype	Relations	Description
Id	integer	PK	Represents line item of product
Text	string		Review text for corresponding product

Table 2-h Review Table Column Definitions

2.2.14 STORE_DEPT table

Each row in this table represents a unique store and department

Column	Datatype	Relations	Description
store	integer		Physical store id.
department	string		Name of the department

Table 2-i Store_department Table Column Definitions

2.2.15 CUSTOMER_IMAGES_META

Each row in this table represents the identity **Metadata** associated with each customer

Column	Datatype	Relations	Description
Img_filename	String	PK	Filename corresponding to the image file
identity_serving	string	FK(customer_id)	Unique ID for identity that does not leak name
sample	integer		Unique ID for image

Table 2-j Customer_images_meta Table Column Definitions

2.2.16 CONVERSATION_AUDIO

Each row in this table represents the identity **Metadata** associated with each audio conversation

Column	Datatype	Relations	Description
wav_filename	string	PK	Unique ID for identity that does not leak name
transcript	string		Unique ID for image

Table 2-k Conversation_audio Table Column Definitions

2.2.17 Customer images

The customer images dataset consists of varied sized .png files that contains a fixed set of identities of each customer.

2.2.18 Audio conversations

The customer audio conversations dataset consists of varied sized .wav files that contains conversation information of customers.

3 DATA SCALING & POPULATION

This clause defines the **Test dataset** population and how it scales.

3.1 Scaling Model

3.1.1 The TPCx-AI benchmark defines a set of discrete scaling points called “scale factors” based on the size of the raw data produced by **PDGF**. The actual byte count may vary depending on individual hardware and software configurations.

3.1.2 The set of scale factors defined for TPCx-AI is:

1GB, 3GB, 10GB, 30GB, 100GB, 300GB, 1000GB, 3000GB, 10000GB

Where GB stands for gigabyte, defined to be 2^{30} bytes.

Comment: The TPC recognizes that additional benchmark development work is necessary to allow TPCx-AI to scale beyond a certain limit of Scale factor size. Note that scale factors greater than 10000 have not been tested by the TPC. Please contact the TPC administrator for assistance on scale factors greater than 10000.

Scale Factor	SF
1GB	1
3GB	3
10GB	10
30GB	30
100GB	100
300GB	300
1000GB	1000
3000GB	3000
10000GB	10000

Table 3-a Scale Factor and SF

3.1.3 SF refers to a scale as shown in Table 3-a where it directly correlates to the Scale Factor (in GB). It is also used in the calculation of the **Performance Metric**. (clause 7.5.7)

3.1.4 Test sponsors may choose any scale factor from the defined series. No other scale factor may be used for a TPCx-AI **Result**.

3.1.5 **Results** at the different scale factors are not comparable, due to the substantially different computational challenges found at different data volumes.

3.1.6 The row size information provided is an estimate and may vary from one benchmark submission to another depending on the precise data base implementation that is selected. It is provided solely to assist benchmark sponsors in the sizing of benchmark configurations.

3.2 Data Population

The data generator used is based on an extension of the Parallel Data Generation Framework (**PDGF**). **PDGF** is a parallel data generator that can produce large amounts of data for an arbitrary schema. The existing **PDGF** can be used to generate the entire TPCx-AI **Test dataset**. **PDGF** handles the volume well since it can scale the size of the data based on a scale factor. It also runs efficiently for large scale factors since it runs in parallel and can leverage large systems dedicated for the benchmark.

The Dataset referred in Table 3-b is the approximate number of images and audio conversation files for the given scale factor that will be part of the Test Dataset for the Training Test.

SF	1	3	10	30	100	300	1,000	3,000	10,000
Number of Images	70	218	1,291	7,084	46,268	241,102	1,303,938	5,368,444	22,531,953
Number of Conversations	387	798	1,965	4,619	11,787	26,895	62,539	126,906	259,980

Table 3-b Table of number of images and audio conversations

Table	Sample Row count								
SF	1	3	10	30	100	300	1,000	3,000	10,000
# Customer	70,710	145,773	358,817	843,356	2,152,033	4,910,448	11,418,023	23,169,807	47,465,671
# Failures	49,490	211,265	1,284,504	7,109,019	46,311,040	241,100,640	1,303,707,240	5,368,210,962	22,529,547,040
# Financial Account	7,071	14,577	35,882	84,336	215,203	491,045	1,141,802	2,316,981	4,746,567
# Financial Transactions	7,353,840	15,160,080	55,975,920	131,564,160	447,622,240	1,276,717,000	3,562,422,240	7,228,980,720	17,277,503,880
# Lineitem	23,026,666	47,449,792	175,234,695	411,922,063	1,401,367,291	3,997,250,363	11,153,580,841	22,632,683,891	54,093,048,609
# Marketplace	70,710	145,773	358,817	843,356	2,152,033	4,910,448	11,418,023	23,169,807	47,465,671
# Order	3,676,955	7,580,237	27,987,765	65,781,798	223,811,527	638,358,245	1,781,211,593	3,614,490,039	8,638,752,268
# Order Returns	1,331,620	2,751,067	10,164,253	23,884,576	81,264,361	231,770,343	646,721,838	1,312,338,657	3,136,554,338
# Product	707	1,458	3,588	8,434	21,520	49,104	114,180	231,698	474,657
# Product Rating	120,694	317,665	851,651	2,089,626	5,775,780	14,168,182	34,109,998	73,936,803	156,322,863
# Review	134,349	306,117	825,286	2,024,064	5,595,278	13,749,260	33,112,258	71,826,411	151,890,144
#Store department	748	748	952	1,224	2,176	4,012	8,432	16,454	32,980

Table 3-c shows the approximate number of rows generated for each of the tables when the Parallel Data Generation Framework is given an input scale factor. The table files will be part of the Test Dataset for the Training Test.

Table 3-c Table of Rows

4 USE CASE SUMMARY

4.1 Summary of the Logical AI data science pipeline

This section summarizes a high level description of the data science pipeline for each use case after the dataset is generated and loaded on to persistent storage. Each use case lists the goal, design criteria and input tables. The figures show the logical stages that show the key preprocessing, training and serving stages for each use case. The actual implementation in the benchmark run and the order in which the training and serving tests run will differ.

4.2 Use case 1 - Customer Segmentation

The Customer segmentation (UC1) **use case** is designed to emulate the data science pipeline to find clusters of customers based on aggregate features where the customers are grouped based on their spending behavior. It involves creating subgroups of customers based on similar traits. The input in this **use case** consists of order and return transaction data from a retail business. The **use case** uses Tables Customer, Order, Lineitem and Order_returns. Refer to Figure 2-a

K-means clustering algorithm is used to derive the optimum number of clusters and understand the underlying customer segments based on the data provided. Clustering is an unsupervised **machine learning** technique, where there are no defined dependent and independent variables, i.e. the training samples are unlabeled. The pattern in the data is used to identify and group similar observations.

The output for **use case 1** comprises of 4 distinct cluster centers.




Figure 4-a Use case 1 *logical* data science pipeline

4.3 Use case 2 – Customer Conversation Transcription

The customer conversation transcription (UC2) **use case** is designed to emulate translating customer audio conversations to text. Usually and if not explicitly forbidden by the customer, customer calls are recorded. To index and classify the conversations it is necessary to transcribe them. Human transcription is costly; therefore an automatic system can be used to generate the transcriptions. This **use case** represents an automatic transcription of given customer calls. The input is a set of audio recordings of customer calls. Data acquisition consists of loading conversation data, transformation of data involves resampling audio to 16kHz, calculating Mel Frequency Cepstral Coefficients (MFCC) from raw audio, transforming labels to encoded sequences and then finally training a deep neural net model on the MFCC features.

DeepSpeech is an open source Speech-to-Text engine using a model trained by machine learning techniques based on the Deep Speech research paper by Baidu⁸. The core of the engine is a recurrent neural network (RNN) trained to ingest speech spectrograms and generate English text transcriptions.

⁸ [Deep-speech: Scaling up end-to-end speech recognition](#)

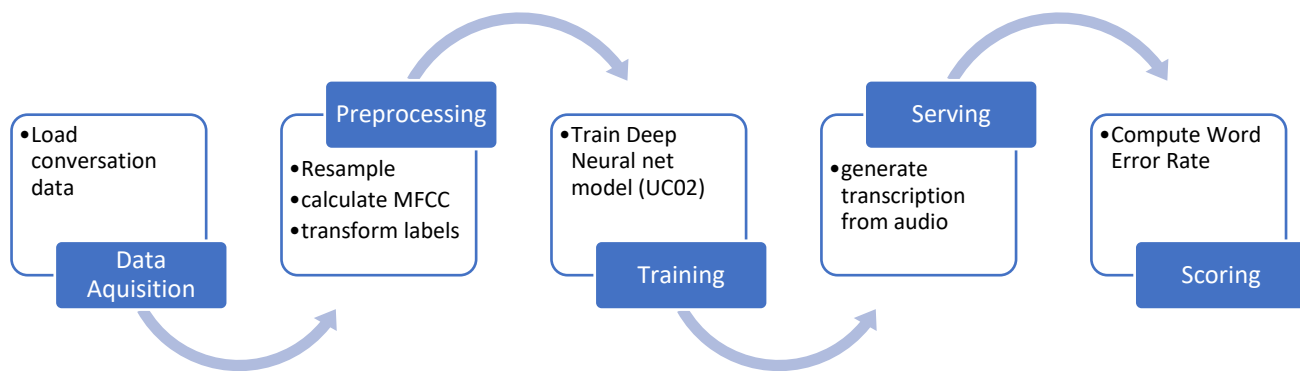


Figure 4-b Use case 2 logical data pipeline

4.4 Use Case 3 – Sales Forecasting

The sales forecasting (UC3) **use case** is designed to emulate weekly sales forecasting for up to 1 year for each department and each store given a limited history of sales data. The input tables are orders, product, Lineitem and store_department. Also included in the dataset are markdown events that might affect the sales of some departments. Each store of a retail chain with multiple stores has limited history of sales data. The retail chain has multiple stores, and each store has multiple departments, such as jewelry, toys, office supplies, kitchen accessories,

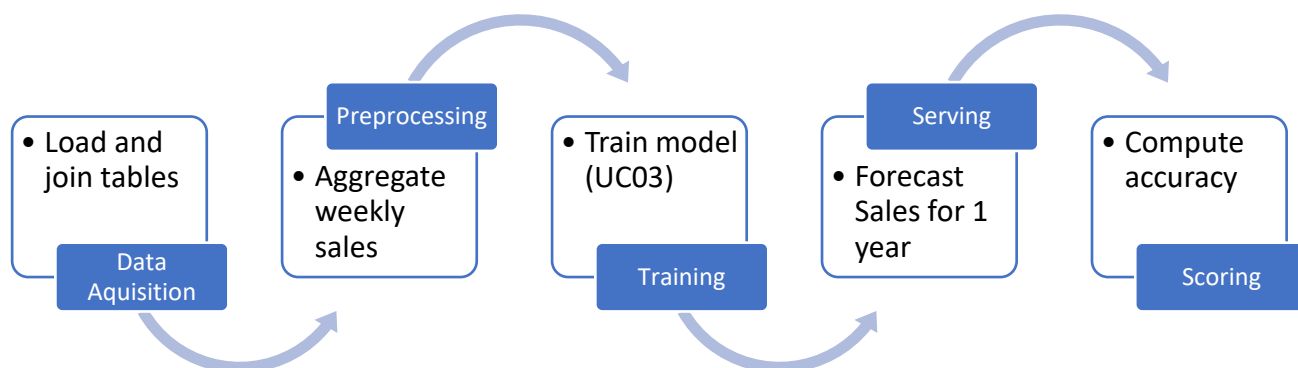


Figure 4-c Use case 3 data science pipeline

among others. The **use case** uses the tables Order, Lineitem, Product and Store_department.

The acronym ARIMA stands for Auto-Regressive Integrated Moving Average. The multiple regression model ARIMA forecasts the weekly sales using a linear combination of predictors like the past values of weekly sales. The term autoregression indicates that it is a regression of the weekly sales against itself. In the case of **use case** 03, there is an aspect of seasonality within the dataset and as such, training an SARIMA (seasonal ARIMA) model further helps accurately forecast the weekly sales.

4.7 Use Case 6 – Hardware failure

The hardware failure (UC6) **use case** is designed to emulate the data science pipeline that predicts and warns of impending failures of hardware in advance. This is based on data and logs collected for every drive. The input is logs of hardware events. Data is pre-aggregated per day and the log records contain events, component IDs, date and failures. The model will predict if component failure is imminent for each set of, previously unseen, log entries. **Use case 4** uses the Failures Table.

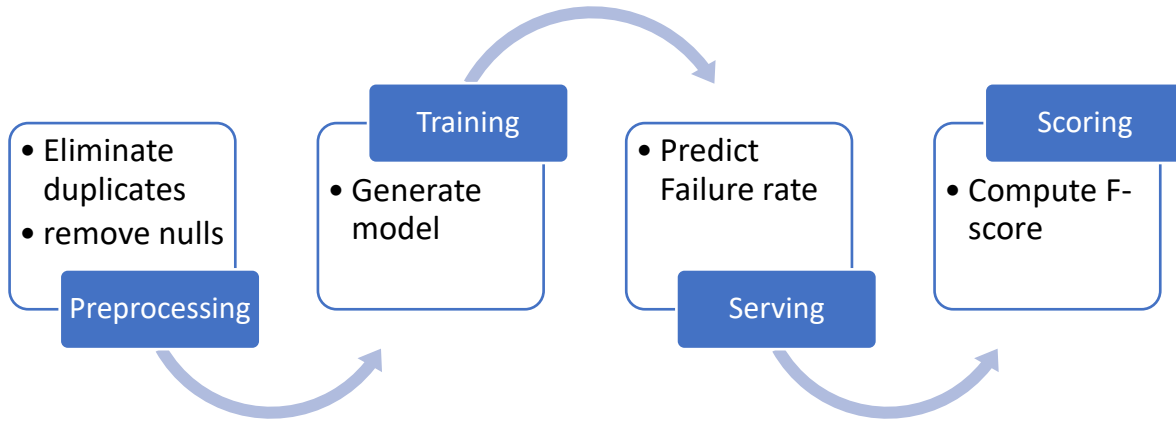


Figure 4-f Use case 6 Data Science Pipeline

Support-Vector machines (SVMs) are robust supervised learning models with associated learning algorithms that analyze data for classification and regression analysis. A support-vector machine constructs a hyperplane or set of hyperplanes in a high- or infinite-dimensional space, which can be used for classification and predictions.⁹

4.8 Use case 7 – Product Rating

One of the biggest challenges faced by the retail industry is generating accurate marketing based on available history of information. In other words, how to accurately inform the customer of products that are available. The product rating (UC7) is designed to emulate a data science pipeline for a recommender system that one would typically find in an online retail environment. The starting point is the shopping history of the customers and more important their rating of historically purchased products. This rating history will be used to train a recommender system that, upon being trained, is able to recommend products that the customer might also be interested in. **Use case 7** uses the Product_Rating table,

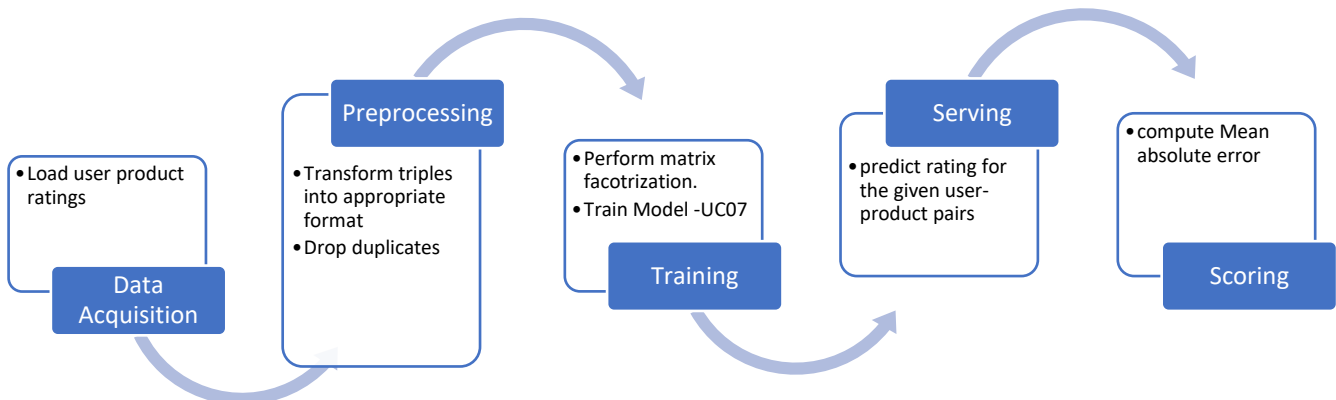


Figure 4-g Use case 7 Data Science Pipeline

The method of least squares is a standard approach in regression analysis to approximate the solution of overdetermined systems by minimizing the sum of the squares of the residuals made in the results of every single equation. Alternating Least Squares (ALS) used in **use case 7** minimizes two loss functions alternatively. It first holds the user matrix fixed

⁹ https://en.wikipedia.org/wiki/Support-vector_machine

and runs gradient descent (minimizing the sum of the squares) with the item matrix; then it holds the item matrix fixed and runs gradient descent with user matrix. Alternating between the two steps guarantees reduction of the cost function, until convergence.

4.9 Use case 8 – Classification of Trips

Customer shopping trips most often can be classified into different types of trips by the retail stores. The underlying assumption is that different trip types have different characteristics and patterns that make them unique. For instance, a weekly trip for grocery shopping would have different data associated with a holiday compared to a Christmas shopping trip. **Use Case 8** (Classification of trips) is designed to emulate training a classifier using a pre-labeled data set, i.e. a data set already containing said trip types, of shopping transactions and then predict the future trip. This prediction could be used in many ways, from tailoring promotional material to customers, marketing activities, placing products in key departments based on the trip type and historical transactions.

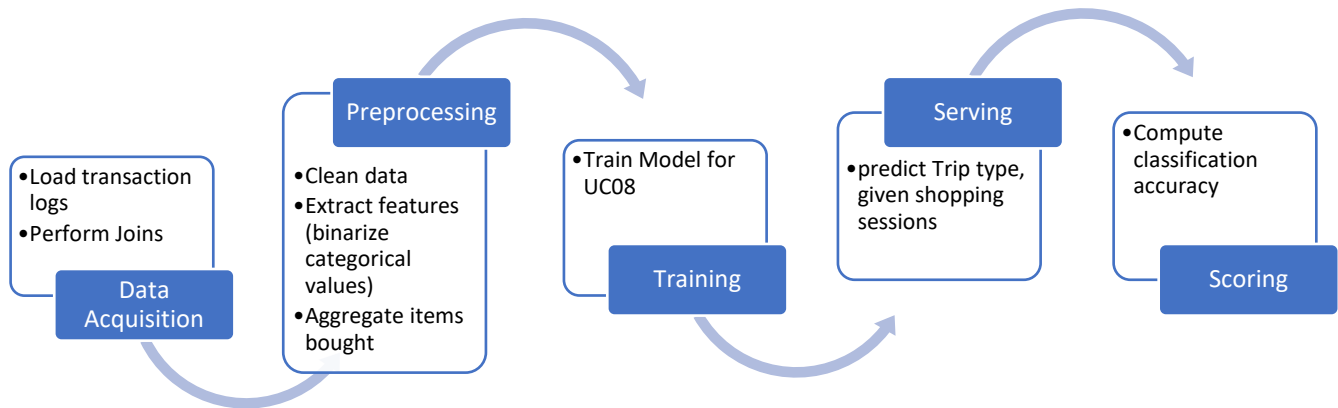


Figure 4-h Use case 8 Data Science Pipeline

4.10 Use case 9 – Facial recognition

The Facial recognition **use case** (UC9) is designed to emulate an end to end facial recognition data science pipeline. This **use case** identifies frequent customers or other persons of interest by training a classifier that recognizes previously seen faces. The data acquisition includes loading **Metadata** as well as importing images using paths from the **Metadata**. The images contain customer faces. The name of the customers is encoded and images aligned based on facial features. The image resolution is changed to fine-tune a pre-trained embedding.

A logistic regression model is trained on the embedding and the customer's name, and embedding is created for aligned faces. Finally, the logistic regression model is used to recognize the face and identify the name for the image.

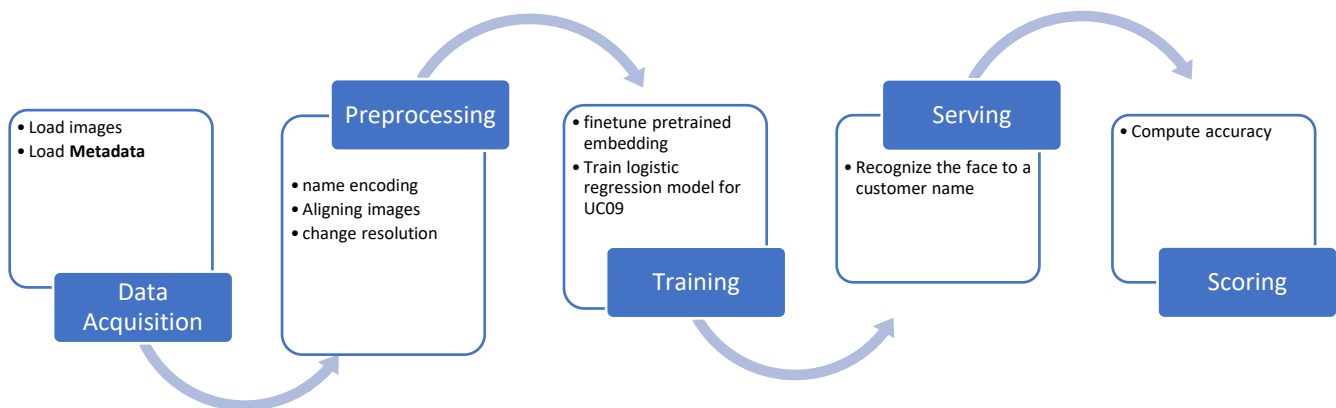


Figure 4-i Use case 9 Data Science Pipeline

4.11 Use case 10 – Fraud Detection

The Fraud detection **use case** (UC1) is designed to emulate whether a financial transaction is a normal transaction or a fraudulent transaction. The input in this **use case** consists of financial transactions in a retail business. The problem to be solved is to identify transactions that are fraudulent. Logistic regression is a linear model for classification that is trained for this **use case**. Logistic regression uses a logistic model that is used to model the probability of a certain class

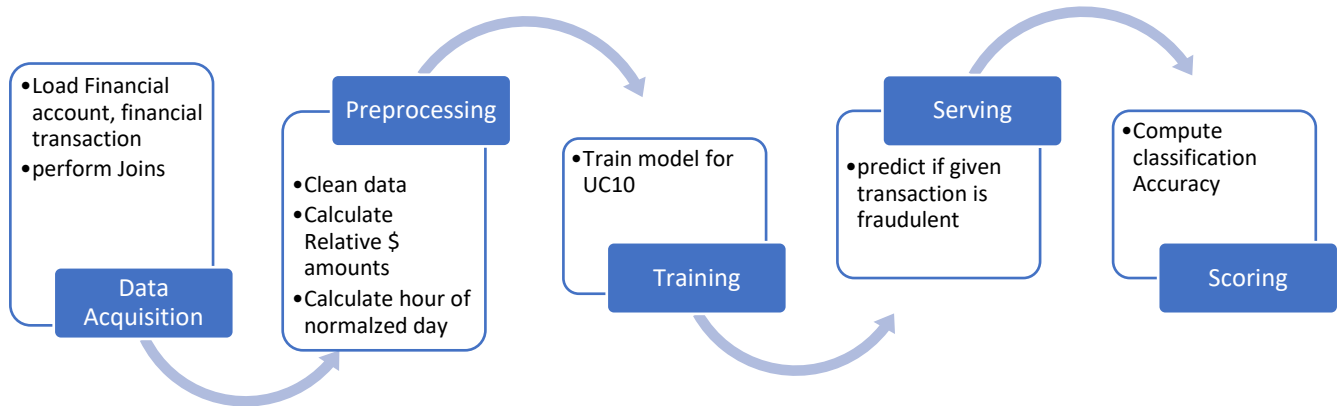


Figure 4-j Use case 10 Data Science Pipeline

or event existing such as a pass/fail or win/lose. This can be extended to model several classes of events. In the case of **use case** 10, the model tried to identify transactions that are fraudulent out of the input dataset. The output is an array of predictions (predicted class label per sample).

5 BENCHMARK KIT CONTENT AND DEVELOPMENT

5.1 Benchmark Kit

This clause defines TPCx-AI Kit contents, its workload execution process, allowed modification by the test sponsor, and contents of the run **report**. The TPCx-AI kit provides fully integrated benchmark and driver software to run on the SUT.

5.1.1 Kit Contents

The TPCx-AI kit contains the following:

- TPCx-AI Specification document.
- TPCx-AI Users Guide (README.md) documentation.
- Configuration files to adapt important parameters to the SUT.
- Scripts which control the benchmarking execution.
- A driver that implements the high-level run logic, time measurement and **Result** computation.
- A set of scripts which are called by the driver to perform benchmark and Use Case operations.
- A set of scripts to automate **Result** verification, checks on **Result** cardinality and **report** generation.
- A set of scripts defined by the TPC to build parts of the kit that are not prebuilt for the user.
- A set of scripts and notes to provide guidance and reference for the setup of execution environments

5.2 Kit Usage for a compliant Result

5.2.1 To submit a compliant TPCx-AI **Result**, the **Test Sponsor** is required to use the TPCx-AI kit as outlined in the TPCx-AI Users Guide (README.md) with the following two exceptions:

- The setting of Kit Parameter files specified in Clause 5.4.
- Test Sponsor Kit Modifications explicitly allowed by Clause 5.6.

5.2.2 If there is a conflict between the TPCx-AI Specification and the TPCx-AI kit, the TPCx-AI kit implementation prevails.

5.3 Kit Run report

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

- Version number of the TPCx-AI kit.
- The input parameters used to run the benchmark.
- The output of the successfully completed Validation test.
- The start, end and total elapsed times of the 6 tests (Clause 7.3) of each Benchmark Run (i.e. Load Test, Power Training Test, Power Serving Test I, Power Serving Test II, Scoring Test and Throughput Test).
- The computed TPCx-AI Secondary Metrics (Clause 7.5.2) for the Benchmark Run.
- The final performance metric for the benchmark run.
- Whether the run is considered valid (i.e. all phases were executed successfully, and all quality metric threshold conditions, as specified in clause 7.3.10.3, were met)

5.4 Kit Parameter settings

5.4.1 The following files and control the kit parameters that may be set by the **Test Sponsor**.

- Generic Benchmark parameters defined in Appendix B
- **Only the use case specific** parameters defined in Appendix C can be changed. All other parameters cannot be changed.

Comment: Global parameters are engine specific. The **Test Sponsor** can set their own parameters and must disclose as part of **FDR**. For example, when using the Apache Spark execution engine, the Test Sponsor can configure generic

parameters, such as the number of tasks used to run one or more **use cases**, the distribution of those tasks across the nodes of the cluster, the amount of compute resources allocated for each task, compression algorithms, and serializer interfaces, among others.

5.5 Benchmark Kit

- 5.5.1 Prior to donating any modifications to an existing TPCx-AI release kit, a CLA must be in place and provided by the sponsor.
- 5.5.2 The draft kit and **compute software** must be reviewed by the TPCx-AI subcommittee during the review process where the subcommittee decides the validity of the software stack that can be used for publications.
- 5.5.3 **Comment:** If the draft kit contains commercially available libraries that perform the same function as in clause 6.3.1, the libraries will need to meet the pricing criteria as listed in the pricing specification. The test sponsor will be responsible for support or can get support priced from a third party.

5.6 Benchmark Kit Modifications

- 5.6.1 For kit changes or modifications other than those allowed by Clause 5.4 any TPC Member, company or individual may bring forward proposed kit changes to the TPCx-AI Benchmark Subcommittee. There are two methods of bringing forward these proposed kit changes:
 - Direct Method – A TPC Member, company, or individual may propose kit changes directly to the TPCx-AI Subcommittee.
 - Indirect Method – If the TPC Member, company, or individual wishes to remain anonymous then a **TPC Certified Auditor** can be used as an intermediary to interact with the TPCx-AI Subcommittee.
- 5.6.2 Regardless of which method is used, the individual that will be interacting with the TPCx-AI Subcommittee becomes the Change Sponsor.
- 5.6.3 Test Sponsor modifications to the provided scripts and configuration files in the TPCx-AI kit to facilitate system, platform and **Framework** differences are allowed without TPC approval. The allowed Test Sponsor Modifications are as follows:
 - Script changes necessary for the kit scripts to execute on a particular Operating System, container platforms or compute service as long as the changes do not alter the execution logic of the script.
 - **Use case** specific optimization **Framework** parameters can be set for the configuration parameters Appendix B only
- 5.6.4 General guidelines used by the subcommittee to review Kit changes include but are not limited to:
 - perform the same function as in clause 6.3.1
 - The same input data must be used
 - To the extent that the new framework accepts similar parameters to existing frameworks (number of iterations, number of clusters, regularization parameters), the values for these parameters should be similar to those used for existing frameworks. If there is a need for the parameters to be different there must be sufficient technical justification provided.
 - If a different **compute software** is used, the **compute software** should be initialized using techniques that are comparable to the existing **compute software** (e.g. for clustering, a new **Framework** should use the same random initialization approach).
- 5.6.5 General guidelines used by the subcommittee to review and incorporate new or different versions of the **TPCx-AI list of Compute software** that are unsupported and not validated or approved by the TPC :
 - New or different versions of the **compute software** libraries and frameworks are intended to be used as part of the Kit (driver) can be used as long as it does not improve the performance of the SUT.
 - The new or different version should provide the same functionality as the previous approved **compute software versions** that are part of the existing benchmark kit
 - The new or different version of the **Compute Software** libraries and frameworks are not commercially available and do not meet the requirements set forth by the pricing specification.

- The new **compute software** should be capable of reporting the same accuracy/evaluation metrics (**F-score**, precision, etc.) as existing ML **Frameworks** and these metrics must meet or be better than the accuracy of the earlier **compute software** used in the comparison.

Comment: An example of a library that could be considered to be part of the approved list could be a scheduler that works with python to distribute execution of the data pipeline on multiple nodes.

Comment: An example of a software library that would not be considered to be part of the approved list could be replacing an older library (part of the approved list) with a newer one that would improve the performance of the kit such that it would not meet the performance requirements for minor kit modifications.

- 5.6.6 No modifications are allowed to the Bash scripts, Java code or the Python code provided in the TPCx-AI kit.
- 5.6.7 No JAR file optimizers are allowed to be used.
- 5.6.8 Any kit modifications not specified in Clause 5.4 or clause 5.6 must be brought forward to the Subcommittee as specified in Clause 5.6
- 5.6.9 Adding support for a new Framework or data science pipeline solution must be brought forward to the Subcommittee as specified in Clause 5.6.5
- 5.6.10 The **use case** implementations included in TPCx-AI use **machine learning** software to solve problems. **Machine learning** software applications and techniques often involve random number generation that might result in slight variations in their final answer. Therefore, doing an exact comparison with a known fixed answer set is not practical and some other criteria must be applied to determine whether modifications are yielding **Results** that should be considered comparable. There are two general categories of changes that could impact the **machine learning** pipeline in the **TPCx-AI** kit:
 - Changes to the version/implementation of the SUT's **machine learning** library (for example a new version of the Spark MLLib library) without any changes to kit itself. The concern in this case is that a new version of the **machine learning** library could make a different tradeoff in accuracy vs performance compared to earlier versions. The following criteria will be applied to evaluate whether **Results** using a new library version should be comparable to previous **Results**:
 - **Results** using the new library version must be generated without any changes to code or parameters in the kit
 - There can be no changes to the input data e.g. the number of clusters for **k-means**, algorithm initialization parameters including seeds for any random initialization, regularization parameters for classification algorithms, etc.
 - Introduction of new **machine learning Frameworks** (not just new versions of the previously supported framework) may require actual changes in the kit code or parameters. This case is more subjective, but the general guidelines for considering results from a new ML **Framework** to be comparable are:
 - The same input data must be used
 - To the extent that the new framework accepts similar parameters to existing frameworks (number of iterations, number of clusters, regularization parameters), the values for these parameters should be similar to those used for existing frameworks. If there is a need for the parameters to be different there must be sufficient technical justification provided.
 - The new **Framework** should be initialized using techniques that are comparable to the existing **Framework** (e.g. for clustering the new **Framework** should use the same random initialization approach).

5.6.11 Classification of Major, Minor and Third Tier Kit Modifications

It is necessary to ensure that the kit remains in sync with fast changing industry and technology landscape. The guidelines below illustrate the current structure of the Kit and help the Subcommittee to decide in a timely manner when evaluating a change proposal. These guidelines will help the Subcommittee do its due diligence and use its discretion to classify and process the change proposals. Modifications to the kit are divided into three types that follow the Revision classifications defined in the TPC Policies. They are:

- Major Kit Modifications
- Minor Kit Modifications
- Third Tier Kit Modifications

5.6.12 Major Kit Modifications:

Major Kit Modifications result in a significant change to the **Use Cases** or intent of the TPCx-AI Benchmark as to make **Results** from the new version non-comparable with the **Results** of the current TPCx-AI version.

These are a few examples of Major Kit Modifications:

- additions, deletions, and modifications to a **Use Case**
- changes to a Primary **Benchmark Metric**
- changes which may alter the reference **Result** set
- changes made to run rules and Benchmark execution process
- Maximum number of iterations
- Learning rate

5.6.13 Minor Kit Modifications:

Minor Kit Modifications do not significantly alter the reference **Result** set, the primary benchmark metrics, or the **Use Case**. **Results** are still comparable to the prior version. A few examples of Minor Kit changes:

- Addition of a new **Framework** support (e.g. Pytorch)
- bug fixes throughout the entire kit
- optimizations to the Framework specific code
- feature additions to Benchmark Driver
- modifications to configuration parameter files
- reference **Result** set changes due to bug fixes
- Framework feature support (or enhancements)
- updates to independent library files
- changes to the Data generator to support features and bugfixes that don't significantly change the characteristics of the datasets that are generated
- Additions or changes to TPCx-AI approved list of software or libraries as specified in 5.6.5

5.6.14 Third Tier Kit Modifications:

Third Tier Kit Modifications are those changes that clarify some confusing or ambiguous area of the kit, but do not alter the benchmark kit code or the **Use Cases**. **Results** are still comparable to the prior version. An example of a Third Tier change is:

- changes in documentation

5.6.15 Simple Review of Kit Modifications

5.6.15.1 For Third Tier (Clause 5.6.14) or Minor kit (Clause 5.6.13) modifications, the Change Sponsor shall present the proposed changes to the Subcommittee. The Subcommittee through its normal course of business will review the proposed changes, make the appropriate kit changes and bring forward the changes to the Council as a new revision of the TPCx-AI Benchmark.

5.6.15.2 If the proposed changes are significant, the Subcommittee may require that the Change Sponsor follow the Formal Review Process defined in Clause 5.6.17 and Clause 5.6.18.

5.6.16 Formal Review of Kit Modifications

For Major (Clause 5.6.12) kit Modifications, at the request to the Subcommittee or if the Change Sponsor so desires, the Change Sponsor shall adhere to the following Formal Review Process.

5.6.17 Formal Proposal of Kit Modifications

5.6.17.1 Step 1: The Change Sponsor must submit to the chair of the TPCx-AI Subcommittee the following information:

- The proposed code changes or new **Framework** code or model
- The reason for proposing the changes
- Result set from the proposed changes
- Complete source code access if the proposed change prototype is available

Comment: To facilitate decision making process change sponsor may provide hardware and software required to validate and review the proposed changes.

5.6.17.2 Step 2: The chair of the TPCx-AI Subcommittee will add a discussion of the proposed changes to the agenda of the next Subcommittee meeting that can be attended by the Change Sponsor.

5.6.17.3 Step 3: The Change Sponsor will present the proposed changes to the TPCx-AI Subcommittee.

5.6.17.4 Step 4: The TPCx-AI Subcommittee will vote on one of three courses of action for the proposed changes.

- Reject the proposed changes.
- Review the proposed changes as a Minor Kit Modification.
- Review the proposed changes as a Major Kit Modification.

5.6.17.5 If the proposed changes are rejected, no further action is necessary. Otherwise, the proposed changes immediately enter a Proposed Change Review period.

5.6.18 Formal Review of Proposed Major Kit Modifications – Approximately twelve week review period.

If the proposed changes were voted to be a Major Kit Modification, then the Subcommittee chair will select at least three members of the Subcommittee to act as primary reviewers of the proposed changes. The Subcommittee chair will also determine the length of the review period and communicate the due date to the primary reviewers and to the Subcommittee. The primary reviewers' job is to examine and test the proposed changes. The primary reviewers are to give their recommendation to the Subcommittee no later than the due date set by the Subcommittee chair which is approximately twelve weeks.

5.6.19 Formal Review of Proposed Minor Kit Modification – Approximately six week review period

If the proposed changes were voted to be a Minor Kit Modification, then the Subcommittee chair will select at least two members of the committee to act as primary reviewers of the proposed changes. The primary reviewers' job is to examine and test the proposed changes. The primary reviewers are to give their recommendation to the committee no more than six weeks later.

5.6.20 Formal Review by Subcommittee

Once the review period ends and the primary reviewers have given their recommendations, the subcommittee will vote on whether to accept the proposed changes into the TPCx-AI benchmark kit.

If the changes are accepted, then the changes will be added to the kit.

5.7 **Kit Validation**

Before any kit can be submitted for approval as a new revision of the TPCx-AI Benchmark Standard, all changes must pass the self-validation tests in the kit.

6 SYSTEM UNDER TEST (SUT)

6.1 Logical Breakdown of System Under Test

- 6.1.1 The tested and reported configurations are composed of the hardware and software components that are employed in the TPCx-AI benchmark run whose cost and performance are described by the benchmark metrics.
- 6.1.2 **System Under Test** consists of:
- 6.1.2.1 Hardware components that can be bare-metal, virtual machines or virtual instances. Examples include compute, storage and networking components or **licensed compute services (LCS)**
 - 6.1.2.2 **Compute Software:** **Compute software** runs on the hardware and/or **Licensed Compute Services** needed to communicate with user interface devices and providing required software capabilities to successfully execute the benchmark.
 - 6.1.2.3 **Data Storage Software:** Data Storage software runs on Data Storage hardware providing required software to create, store, and access input, output, intermediate, and temp data during the benchmark execution.
 - 6.1.2.4 Devices in addition to listed above used in the **SUT**, for example compute devices and/or data storage devices, for e.g. FPGA, Accelerator appliance, Accelerator cards, compression add-on cards, encryption add-on cards etc. and their supporting software stack, device driver software, plug-in software.
 - 6.1.2.5 Any hardware and software devices of all networks required to connect and support the **SUT** systems
 - 6.1.2.6 Device running benchmark driver hardware and software resides on a separate system facilitating the execution of the benchmark, without interfering and influencing the **SUT**. This device is not part of the **SUT** and contains necessary SW and configuration to interact with the **SUT** and can be in form of Desktop, Notebook, or a Server.
 - 6.1.2.7 The figures below show example models of the target SUT. Figure 6-a represents a single node SUT configuration.

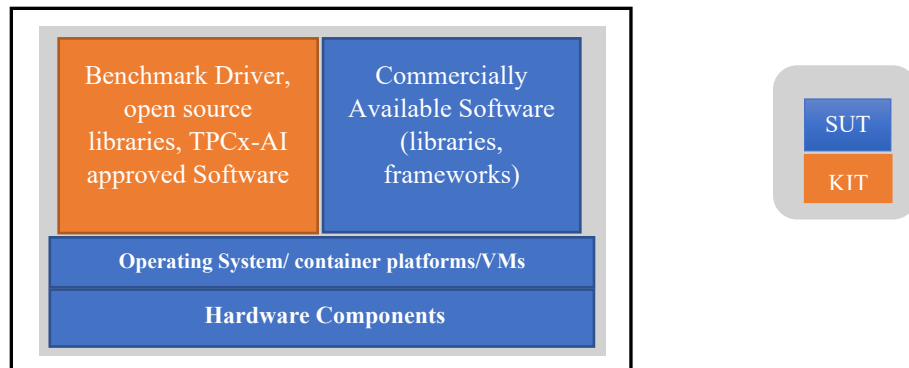


Figure 6-a Single SUT Configuration

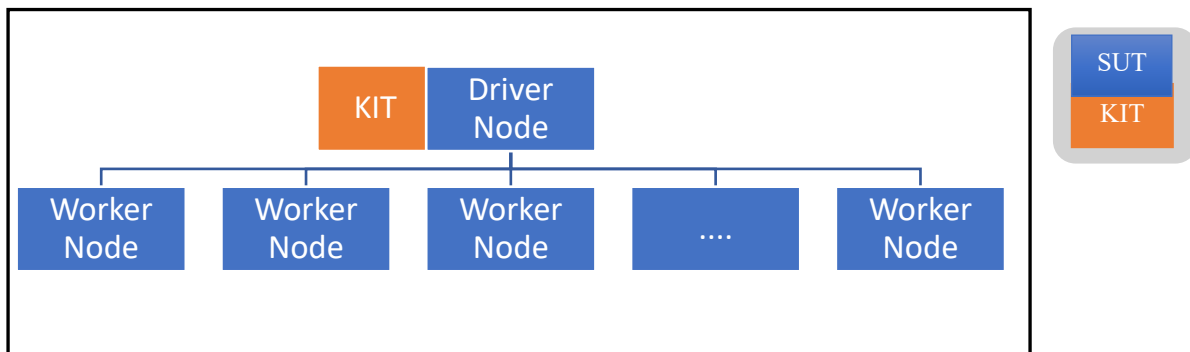


Figure 6-b Clustered SUT Configuration

6.1.2.8 Figure 6-b represents a typical clustered SUT configuration where the driver and worker nodes represent the hardware and compute software and/or licensed compute software needed to run the benchmark.

6.2 Commercially Available Products

Except for the TPCx-AI benchmark driver software and **TPCx-AI approved list of compute libraries** needed to run the benchmark, all **SUT** components must be **Commercially Available Products**. The source code of any non-commercially available products used to implement the **SUT** (including but not limited to scripts used to install, configure and tune the **SUT**) must be disclosed.

6.3 Benchmark driver & compute software and libraries

The following constitute the software that is part of the Benchmark Kit and **Compute Software** that is part of the **SUT**.

- 6.3.1 Benchmark kit or driver software needed to run the benchmark as shown in the figures 6-a and 6-b above will not be priced.
- 6.3.2 All software, libraries and frameworks used as part of the **compute software** (clause 6.1.2.2) are part of the **SUT**.
- 6.3.3 The TPCx-AI subcommittee will maintain a list of **TPCx-AI approved list of compute software (including version)** as allowable software to be used for the purposes of benchmark publications where the pricing specification may not apply. They may be considered as part of the TPCx-AI benchmark driver software and not the **SUT**.
- 6.3.4 The **TPCx-AI approved list of compute software** libraries listed by the TPCx-AI subcommittee will have been validated for a specific version of the release kit and the subcommittee shall vote in new items on that list (clause 5.6.5)

6.4 Data Redundancy Requirement

The following clauses describe required **Data Redundancy** characteristics of the **SUT**. The failures described are not induced during the benchmark Execution.

- 6.4.1 Durable Medium: A durable medium that is either:
 - An inherently non-volatile medium (e.g., magnetic disk, magnetic tape, optical disk, solid state disk, Phase Change Memory, or technology, similar to Phase Change Memory. etc.)or
 - A volatile medium with its own self-contained power supply that will retain and permit the transfer of data, before any data is lost, to an inherently non-volatile medium after the failure of external power.
- 6.4.2 The **System Under Test** must be configured to satisfy the requirements for **Data Redundancy** described in this clause. **Data Redundancy** is demonstrated by the **SUT** being able to maintain operations with full data access during and after the permanent irrecoverable failure of any single storage Medium containing tables, input, output, or **Metadata** (including **Master Node/Name Node metadata** [where present]).

Comment: A configured and priced Uninterruptible Power Supply (UPS) is not considered external power.

Comment: DRAM can be considered a durable storage medium if it can preserve data long enough to satisfy the requirement (b) above. For example, if memory is accompanied by an Uninterruptible Power Supply, and the contents of memory can be transferred to an inherently non-volatile medium during the failure, then the memory is considered durable. Note that no distinction is made between main memory and memory performing similar permanent or temporary data storage in other parts of the system (e.g., disk controller caches).

6.4.3 Data Redundancy Reporting Requirements

At the heart of this requirement is handling the failure of any single durable media for consistency with other TPC benchmarks. For distributed file systems, while **HDFS** 2-way replication would have satisfied the spirit of this requirement, **HDFS** deployments using 3-way replication are the norm (both for redundancy and for performance). The requirements below specify handling the failure of the TPCx-AI dataset, input Dataset, output files and the associated **Metadata**.

6.4.3.1 The test sponsor must guarantee that the test system will not lose the **Test dataset** used for all the benchmark runs due to a permanent irrecoverable failure of any single durable medium. **Use case** execution is not permitted to fail and no data can be lost in the presence of a single durable medium failure. This requirement applies to all Durable Medium containing TPCx-AI data (e.g., **Test dataset**, serving and scoring datasets including table data, **Metadata** (includes **Master Node/Name Node Metadata** [where present], **Undo/Redo Log** data [where present], and “temporary data” [where present])). This requirement also applies to any benchmark **use case** results (output data) stored on the SUT.

6.4.3.2 For **HDFS** file systems providing redundancy via replication, erasure coding, etc.:

The test sponsor must provide a **report** showing data resiliency. For example, with **HDFS** this can be done by running “hdfs fsck -blocks”. “hdfs dfs -du -s -h” and “hdfs ec -getPolicy -path /”. When TPCxAI_Benchmarkrun.sh is run, this **report** will automatically be generated at the end of each benchmark run.

- For 3-way replication of the TPCx-AI dataset, the “default replication factor” should be at least 3 and “under-replicated blocks” should be zero.
- For erasure coding, the Auditor must verify that the codec, node_count, and parity settings results in redundancy at least equivalent to 3-way replication.
- Metastore data in the Name Node/Master Node that hosts the TPCx-AI **Metadata** must be set to at least an equivalent of replication factor of two.

Comment: Typically this will be accomplished by verifying that “under-erasure-coded block groups”=0, num_parity blocks ≥ 3 and node_count \geq (num_data_blocks + num_parity_blocks) but the **Auditor** may need to perform an additional, implementation specific review.

6.4.3.3 For distributed file systems other than **HDFS**, like traditional JBOD, redundancy must be proved by the **Test sponsor**. They 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 at least equivalent to the data redundancy provided by traditional local-JBOD storage and **HDFS** replication factor of 3.

- If stored in a distributed filesystem, **Test Dataset**, Input Data and Output Data must be set to at least an equivalent of replication factor three for **HDFS** on JBOD.
- TPCx-AI **Metadata** stored must be set to at least an equivalent of replication factor of two.

6.4.3.4 For SUT components NOT using a distributed file system such as **HDFS**, but running on a single node SUT that provide redundancy via a **High Availability System**:

- a) The test sponsor needs to provide a **report** that explains the configuration in sufficient detail to satisfy the **Auditor**/PPB that outlines the use of distinct durable mediums for the individual service instances in the **HAS**.
- b) While encouraged, there is no requirement for triple redundancy for this class of data. This includes the TPCx-AI dataset, input Data, Output Data and associated **Metadata**.

Comment: A single durable medium failure could take down a service instance in the **HAS** but continued execution would be guaranteed by the existence of a secondary service instance using a distinct durable medium.

Comment: For consistency with the distributed file system model, no explicit test is necessarily required.

- c) The solution must guarantee uninterrupted access to the data on durable medium when a single Durable Media failure occurs.
- d) The test sponsor must provide a **report** from a system tool detailing the media redundancy hardware/software configuration to the satisfaction of the **Auditor**/PPB (e.g., a **report** showing that RAID-5 or RAID-10 is used).

Comment: Roll-forward recovery from an archive dataset copy (e.g., a copy taken prior to the run) using **Undo/Redo Log** data is not acceptable as the recovery mechanism in the case of durable medium failure. Note that “checkpoints”, “control points”, “consistency points”, etc. of the dataset taken during a run are not considered to be archives.

Comment: For consistency with the distributed file system model, no explicit test is necessarily required.

Comment: **Use Case execution** may not fail due to a permanent irrecoverable failure of any single durable medium containing TPCx-AI data. However, medium failures are not allowed during benchmark runs to be considered valid

(e.g., to avoid the possibility of higher performance when 3-way replication degrades into 2-way replication on medium failure).

7 EXECUTION RULES AND PERFORMANCE METRICS

7.1 Benchmark Execution

- 7.1.1 A Benchmark Execution is defined as a Validation test (Clause 7.2) followed by the **Benchmark Run** that consists of 6 tests.
- 7.1.2 The Test sponsor runs the following scripts, in the order as they are presented:
- TPCxAI_Validation.sh
 - TPCxAI_Benchmarkrun.sh
- 7.1.3 No part of the **SUT** may be restarted during the Benchmark Execution. If there is a non-recoverable error reported by any of the applications, operating system, or hardware in any of the tests (Clause 7.1.2) or between validation and benchmark runs, the run is considered invalid. If a recoverable error is detected in any of the tests 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 a recoverable error requires manual intervention to deal with or correct, then the run is considered invalid.
- 7.1.4 The sponsor may choose to use a different default user configuration file just for the validation test if they choose to do so.

7.2 Validation Test

- 7.2.1 The Validation test performs data generation, data load, power training test, power serving test, and scoring test with scale factor 1 to perform a **Result** validation against the reference **Result** set in the kit. Validation test ensures that the setup used by the Test Sponsor to produce the publication are comparable to the reference **Result** set generated.
- 7.2.2 The validation **Result** set for SF1 is the reference **Result** used to validate the **SUT** for result correctness.
- 7.2.3 The intent of **Result** validation is to validate all the **use cases** against SF1 and compare it against the reference **Result** set packaged with the benchmark kit. This is exercised against the **SUT** before the Benchmark Run as part of **SUT** Validation Test.
- 7.2.4 Populate the **SUT** with SF1 dataset and schema information.
- 7.2.5 Execute the **use cases** using same standard parameters as will be used by the benchmark run. Verify the **report** generated by the driver by comparing the output to the reference **Result** set.
- 7.2.6 **Result** Validation Process is defined in TPCx-AI_Validation.sh script and the generated **report** shall state that the output of scoring is the same or better than the reference accuracy metrics for all **use cases**.
- 7.2.7 The Validation steps are provided below:
- *ENGINE_VALIDATION_DATA_GENERATION*: This phase as defined in TPCx-AI_Validation.sh generates a dataset at a fixed scale factor of 1 (SF1).
 - *ENGINE_VALIDATION_LOAD_TEST*: During this phase, as defined in TPCx-AI_Validation.sh the data generated will be loaded into the final location from where they will be eventually accessed to execute each one of subsequent **Performance Tests** (Clause 7.3.1).
 - *ENGINE_VALIDATION_POWER_TRAINING_TEST*: During this phase as defined in TPCx-AI_Validation.sh, all **use cases**' training phase will be run in sequence and the results are stored in persistent storage.
 - *ENGINE_VALIDATION_POWER_SERVING_TEST*: During this phase as defined in TPCx-AI_Validation.sh, all **use cases**' serving phase will be run in sequence and the results are stored in persistent storage.
 - *ENGINE_VALIDATION_RESULT_VALIDATION*: During this automated phase as defined in TPCx-AI_Validation.sh, the benchmark driver compares the accuracy results from all **use cases** against a known reference result packaged with the kit. Due to the small training sample sizes, spark implementations may not be able to pass the accuracy test with use case 3. In that case, a failure to meet the accuracy metric threshold for use case 3 is acceptable and does not invalidate the validation test.

- 7.2.8 The elapsed time for Validation Test is not included as part of Benchmark Metric calculation.
- 7.2.9 The elapsed time for Validation Test is not counted as part of Benchmark Execution.
- 7.2.10 For all other scale factors, used in the Benchmark Run, the benchmark driver at the end of the benchmark performs output validation checking for the presence of output data from power test and throughput test in order to qualify successful benchmark execution.
- 7.2.11 Output data for Validation test to be verified:
- 7.2.11.1 The training phase of each **use case** has generated a model file successfully.
- 7.2.11.2 The model file for each **use case** is used to conduct the serving test for that **use case** successfully.
- 7.2.11.3 The model file for each **use case** is also used to conduct the scoring tests for that **use case**.
- 7.2.12 Software library versions for all the libraries used as part of the SUT must be listed.e.g. python, spark etc.
- 7.2.13 Software library versions for TPCx-AI approved compute libraries are defined in Appendix F and must be used in the **Result** validation.The audit will check for authenticity & versions of the libraries used.

7.3 Benchmark Run

- 7.3.1 All TPCx-AI tests are initiated by the TPCx-AI master scripts which can be executed from any of the nodes in the **SUT**. The tests are listed below:
- Load Test
 - Power Training Test
 - Power Serving Test I
 - Power Serving Test II
 - Scoring Test
 - Throughput Test
- 7.3.2 A valid run consists of 6 separate tests run sequentially. These tests may not overlap in their execution times. For example, the start of Power Serving Test I may not begin until Power Training Test is complete, the start of Power Serving Test II may not begin until Power Serving Test I is complete, etc.
- 7.3.3 The **Test Sponsor** sets the Benchmark Driver Parameters used during the tests are per Appendix B.
- 7.3.4 The elapsed time for each test in Clause 7.3.1 must be reported. However, the elapsed time of the Scoring test will not be considered for the computation of the final performance metric. Instead, the results of the Scoring test are used to determine whether the **Performance Test** was a successful one based on whether the thresholds of the quality metrics for each **use case** were met or not.
- 7.3.5 Parameters *BENCHMARK_START* and *BENCHMARK_STOP* in the Run **Report** determine the overall elapsed time for the **Performance Test**.
- 7.3.6 Data Generation
- 7.3.6.1 The process of using **PDGF** to create the data in a format suitable for presentation to the load facility. **PDGF** generates different types of data, including images, audio, and text-based flat files. **PDGF** generates data to local file systems in the SUT from where they can be loaded to a file system to be used as part of the **Use Cases** execution.
- 7.3.7 Load Test
- 7.3.7.1 The process of copying the input dataset files to the final location from where they will be eventually accessed to execute each one of the subsequent benchmark phases (Clause 7.3.1) is known as the Load Test. The Load Test may consist of an optional data relocation or data preparation phase.

7.3.7.2 The location of the **PDGF** output is different from the final Dataset Location and so the data must be copied into the final Dataset Location. This phase is timed and contributes to the load time. Note that this copy may be done as part of the optional format conversion in the Data Preparation phase, in which case the time is captured as part of the Data Preparation timing. If multiple data copies occur between the **PDGF** generation and the placement of the data in the final Dataset Location, only the final copy into the Dataset location is included in the load time. For example, if **PDGF** generates data initially to a location external to the **SUT**, the flat files are subsequently copied to a staging area on the **SUT**, and then the data is copied again from the staging area into the Dataset Location as part of the Data Preparation format conversion, only the final copy of the entire **Test Dataset** is included in the load time.

7.3.7.3 Optional Data preparation

7.3.7.3.1 Example of an optional data preparation phase includes all additional work, beyond the Generation and Relocation steps, required to prepare the data for the **use case** execution. This includes but is not limited to the following steps:

- Creation of **Metadata**.
- Computing statistics for the dataset.
- Conversion of the data into an alternative or optimized format. An example would be conversion from the row-oriented format in the flat files to a compressed and/or columnar format. Note this is an optional step – if the flat files have been placed in the final Dataset Location by earlier load steps, then it's permissible to run data preprocessing directly against the flat files in their original format in the Dataset Location.

7.3.7.3.2 Any format conversion or creation of auxiliary data structures must meet the following requirements:

- it must not lose information from the original **Test Dataset**.
- it cannot make use of any knowledge of the workloads in the benchmark.

7.3.7.3.3 For example, the conversion can't remove columns that aren't referenced by the benchmark **use case** execution, and creation of materialized views that pre-compute some or all the results needed for preprocessing, training, serving or scoring is not allowed.

7.3.7.3.4 All work done during Data Preparation is timed and included in the load time.

Any additional work required to prepare the data for execution (e.g. format conversion) is timed and included in the load time.

7.3.7.4 The **Load Test** must not include the execution of any of the queries in the **Power Test** or **Throughput Test** or any similar query.

7.3.8 Power Training Test

7.3.8.1 Power training test determines the maximum speed the **SUT** can process the Training phase of all 10 **use cases**. The Training pipelines of all **use cases** must run in sequential order. The result of the training test should be the generation of a training model files at the completion of the training stage of each of the **use cases**. All the model files will be saved in persistent storage where they will then be used for the Power Serving Test.

7.3.8.2 The Power training test is timed and included as part of the overall metric.

7.3.8.3 The output of the training test results is the generation of a model file that will be used for the remaining tests during the performance test. A test sponsor may however choose to make changes to the model file in order to optimize the performance of the remaining performance tests.

7.3.8.4 Model optimizations will be allowed on the generated original model as long as the original model has been trained on the test dataset. Examples of optimizations may include but are not limited to converting variables to constants, removing unused nodes or quantization.

7.3.8.5 The **Test Sponsor** can invoke post model optimizations after the training phase generates a model file.




Figure 7-a Power Training Test

- 7.3.8.6 All of the elapsed time spent in post model optimizations needs to be calculated as well as reported in the FDR. In addition, this time shall be added to the overall training time for the specific use case.
- 7.3.8.7 The audit process will verify the original model file generated from the output of the Power training test and also the optimized model file that can be used by the Power Serving Test I, Power Serving Test II, Scoring test and throughput test. In addition, the process followed by the test sponsor to generate the optimized model file will have to be reported.
- 7.3.9 Power Serving Test I and Power Serving Test II
- 7.3.9.1 Power Serving tests determine the maximum speed the **SUT** can process the Serving stages of all 10 **use cases**.
- 7.3.9.2 The Serving Tests follow the Power training Test after the model files for all the **use cases** have been created.
- 7.3.9.3 There shall be only one model file for each use case that will be used in the power serving test.
- 7.3.9.4 There are 2 Power Serving tests that run sequentially:
- Power Serving Test I
 - Power Serving Test II
- 7.3.9.5 The same model file for a use case shall be used for each of the two power serving tests.
- 7.3.9.6 There shall be no change in configurations or tuning between the two serving tests.

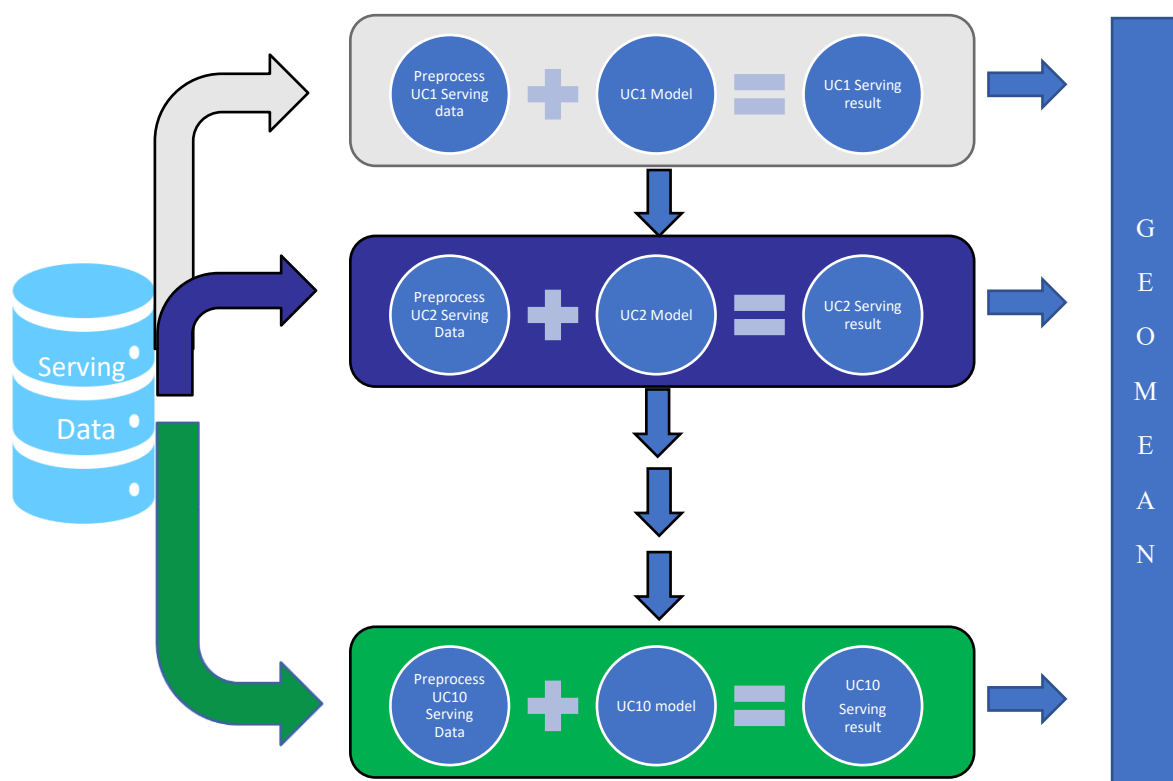


Figure 7-b Power Serving Test(s)

7.3.10 Scoring Test

7.3.10.1 The Scoring Test generates a small dataset with ground truth labels

7.3.10.2 During this scoring test, a separate serving phase is executed in sequence for all 10 **use cases** against a newly generated data (excluding the truth labels) and the resulting labels from separate serving phase are compared with the ground truth labels to determine the accuracy metric or error incurred by the **use case**.

7.3.10.3 The scoring results for each **use case** should meet or better the reference result set provided in the Kit as shown in Table 7-a.

Use Case	Quality metric name	Accuracy thresholds	Lower is Better
UC1	k-means clusters	N/A	N/A
UC2	Word error rate	0.500	Yes
UC3	Forecast accuracy	5.400	Yes
UC4	Matthews correlation coefficient	0.650	No
UC5	Prediction quality (RMSLE)	0.500	Yes
UC6	F-score	0.190	No
UC7	Mean Absolute Error	1.800	Yes
UC8	Classification accuracy	0.650	No
UC9	Face recognition accuracy	0.900	No
UC10	Classification accuracy	0.700	No

Table Benchmark Run Accuracy metrics

7.3.10.4 The scoring test is not part of the timed components of the benchmark. However, this step is crucial to determine whether the quality threshold defined for each **use case** is met as shown in figure 7-




Figure 7-c Scoring Test

7.3.11 Throughput Test

7.3.11.1 The Throughput Tests measure the ability of the system to process the most serving use cases in the least amount of time with multiple users

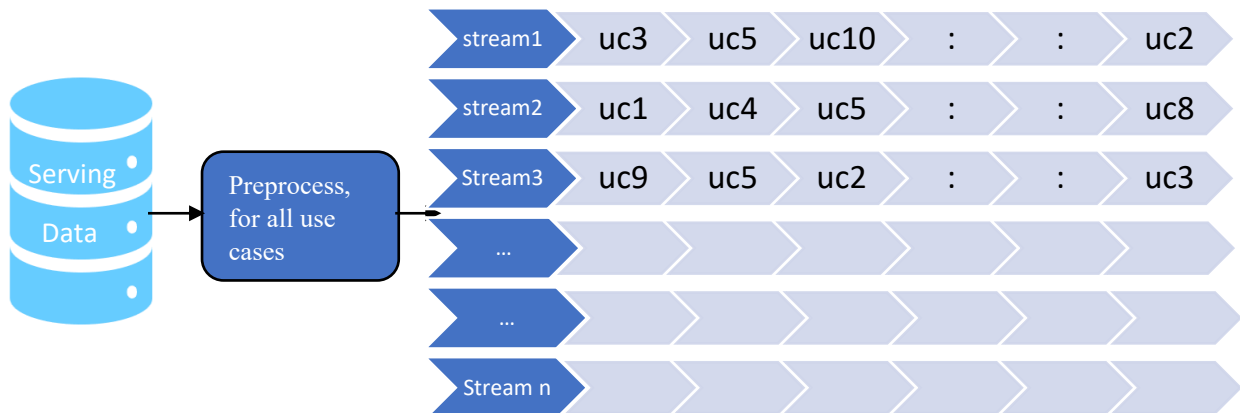


Figure 7-d Throughput test

7.3.11.2 The Throughput Test runs several streams concurrently where each stream is running the 10 **use cases** in a unique order. Each stream runs all **use cases** serving pipelines according to the placement order set in Appendix D. The Default streams for throughput test is set to 2, with the number of additional concurrent streams being configurable with no maximum limit.

- 7.3.11.3 Use Case placement in the serving throughput test is performed using the automatic shuffling of the streams. Use Case placement for the first 100 streams are shown in Appendix D. For a list of all the streams that can be used in the throughput test refer to the streams.yaml file included in the kit.
- 7.3.11.4 At the end of the throughput test, the total time taken to run the throughput test is used as one of the inputs to calculate the final performance metric of the benchmark.
- 7.3.12 The **Reported Performance Metric** is the TPCx-AI **Performance Metric** for the **Benchmark Run**. There must not be any interruption during the tests, and all tests should be run without intervention.

7.4 Configuration and Tuning

- 7.4.1 The **SUT** cannot be reconfigured, changed, or re-tuned by the **Test Sponsor** during or between any of the tests described in Clause 7.3.1, the Benchmark Run.
- 7.4.2 Any manual tunings or additional scripts and steps that need to be performed to the **SUT** must be performed before the beginning of the TPCxAI_Validation.sh script execution described in Clause 7.1.2, and must be fully disclosed. Clause (9.6.7)
- 7.4.3 Changes and tuning performed on the **SUT** between any of the tests are allowed as defined in Appendix C.
- 7.4.4 Any changes to default tunings or parameters of the applications, **Operating Systems**, or hardware of the **SUT** must be disclosed.
- 7.4.5 Any changes deemed with the characteristics of **Benchmark Special** in Clause 0.3 are not allowed.

If the **Test Sponsor** would like to perform model file optimizations after the Power Training Test, they can do so as long as the original model file, the new model file and all the relevant steps and files needed to perform the optimizations are disclosed as specified in Clause 9.6.8.

- The new model file must be derived from the original model file generated as a result of the Power Training Test.
- The new model file must be used for all the subsequent Power Serving Test I, Power Serving Test II, Scoring Test and Throughput Test.
- Both the original model files and the new model files need to be reported and shared for the audit as per Clause 9.6.8

7.5 Metrics

- 7.5.1 TPCx-AI defines three primary metrics:
- a) A Performance Metric, AIUCpm@SF, reflecting the TPCx-AI **use case** throughput (see Clause 7.5.7)
 - b) A Price-Performance metric, \$/AIUCpm@SF (see Clause 7.5.8)
 - c) System **availability date** (see Clause 7.5.10).
- 7.5.2 Secondary metrics are additional metrics defined below are provided as part of the **Report** :
- Computed Load Metric T_{LD} as defined in Clause 7.5.6
 - Computed Power Training Test Metric T_{PTT}
 - Computed Power Serving Test Metric T_{PST}
 - Computed Throughput Test Metric T_{TT}
 - When TPC_Energy option is chosen for reporting, the TPCx-AI energy metric reports the power per performance and is expressed as Watts/AIUCpm@SF. (see TPC-Energy specification for additional requirements).

- 7.5.3 Each secondary metric shall be referenced in conjunction with the scale factor at which it was achieved. For example, Load Time references shall take the form of Load Time @ SF, or “Load Time = 7900seconds @ 300GB”.
- 7.5.4 The start time for any test is when the timestamp is taken before the first use case of the test is submitted to the **SUT** by the driver.
- 7.5.5 The End time of any test is when the timestamp is taken after the last use case of the test completes and is acknowledged by the driver from the **SUT**.
- 7.5.6 The **Performance Metric** of the TPCx-AI benchmark, AIUCpm@SF, is computed by combining metric components representing the load, power, and throughput tests:

7.5.6.1 T_{LD} is the load factor computed as:

$$T_{LD} = T_{Load}$$

Where,

- T_{Load} is the elapsed time of the Load Test (Clause 7.3.7) in seconds.

7.5.6.2 T_{PTT} is the geometric mean of the elapsed time UT in seconds of each of the **Use case** Training times as measured during the Power Training Test (Clause 7.3.8):

$$T_{PTT} = \sqrt[N]{\prod_{i=1}^N UT(i)}$$

Where,

- $UT(i)$ is the elapsed time in seconds of the **Use Case** i during the Power Training Test and 10 is the number of **Use cases** in the Benchmark.
- N is the number of use cases in the benchmark.

7.5.6.3 The **Performance Test** includes the load test, the power training test, 2 power serving tests, throughput test and scoring test.

7.5.6.4 $TPST$, $TPST1$ and $TPST2$ are derived as follows:

7.5.6.5 $TPST1$ is the geometric mean of the elapsed time US in seconds of each of the **Use case** Serving times as measured during the Serving Power Test I (Clause 7.3.9):

$$T_{PST1} = \sqrt[N]{\prod_{i=1}^N US1(i)}$$

Where.

- $US1(i)$ is the elapsed time in seconds of the **Use Case** i during the Serving Power Test I and N is the number of **Use cases** in the Benchmark.
- N is the number of use cases in the benchmark.

7.5.6.6 $TPST2$ is the geometric mean of the elapsed time US in seconds of each of the **Use case** Serving times as measured during the Serving Power Test II (Clause 7.3.9):

$$T_{PST2} = \sqrt[N]{\prod_{i=1}^N US2(i)}$$

Where,

- $US2(i)$ is the elapsed time in seconds of the **Use Case** i during the Serving Power Test II and N is the number of **Use cases** in the Benchmark.
- N is the number of use cases in the benchmark.

7.5.6.7 $TPST$ is the higher value of the two Serving Power tests $TPST1$ and $TPST2$

7.5.6.8 T_{TT} is the throughput test metric computed as the total elapsed time of the throughput test divided by the number of streams and the number of use cases in the **Performance Test** as measured during the Throughput Test.

7.5.6.8.1 T_{Tput} is the elapsed time of all streams from the Throughput Test. This is the difference between the Throughput Test Start Time and Throughput test End Time.

- S is the number of streams in the Throughput Test.
- N is the number of use cases used in the Throughput Test. This will always be the total number of use cases used in the **Performance Test**.

$$T_{TT} = \frac{1}{N * S} T_{Tput}$$

7.5.7 The Performance Metric (AIUCpm@SF)

The primary performance metric of the benchmark is AIUCpm@SF, the effective use case throughput of the benchmarked configuration, defined as:

$$AIUCpm@SF = \frac{SF * N * 60}{\sqrt[4]{T_{LD} * T_{PTT} * T_{PST} * T_{TT}}}$$

Where:

- SF is defined in Clause 3.1.3, and is based on the scale factor used in the benchmark
- N is the total number of **use cases** executed in a Run
- T_{PTT} where T_{PTT} is the geomean of the training time of each of the Training Power Tests, as defined in Clause 7.3.86.2,
- T_{PST} is the higher time of the two Power serving test times (T_{PST1} and T_{PST2})
- T_{TT} is the total elapsed time of Throughput Test as defined in Clause 7.5.6.8.
- T_{LD} is the load factor computed as $T_{LD} = T_{Load}$, and T_{Load} is the time to finish the load, as defined in Clause 7.5.6.1
- T_{PTT} , T_{PST} , T_{TT} and T_{LD} quantities are in units of decimal seconds.

7.5.8 The Price Performance Metric $\$/AIUCpm@SF$ is defined as:

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

Where:

- P is the price of the Priced System as defined in Clause 8.2
- $AUCpm@SF$ is the reported performance metric as defined in Clause 7.5.7.

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

7.5.10 The System **Availability Date**, as defined in the TPC Pricing Specification must be disclosed in any references to either the performance or price-performance metric of the benchmark.

7.5.11 Fair Metric Comparison

7.5.11.1.1 **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.

7.5.12 If **results** measured against different dataset sizes (i.e., with different scale factors) appear in a printed or electronic communication, then each reference to a **result** or metric must clearly indicate the dataset size against which it was obtained. In particular, all textual references to TPCx-AI metrics (performance or price/performance) appearing must be expressed in the form that includes the size of the **Test Dataset** as an integral part of the metric's name, i.e. including the "@size" 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 **Test Dataset** size on which metric is based must be immediately discernible either by appropriate axis labeling or data point labeling.

7.5.13 In addition, the **Results** must be accompanied by a disclaimer stating:

- "The TPC believes that comparisons of TPCx-AI **Results** measured against different dataset sizes are misleading and discourages such comparisons".
- Any TPCx-AI **Result** is comparable to other TPCx-AI **Results** regardless of the number of query streams used during the test (as long as the scale factors chosen for their respective **Test Datasets** were the same).

7.5.14 Required Reporting Components

To be compliant with the TPCx-AI standard and the TPC's fair use policies, all public references to TPCx-AI **Results** for a given configuration must include the following components:

- The size of the test data expressed separately or as part of the metric's names (e.g., AIUCph@10GB).
- The TPCx-AI Performance Metric, AIUCpm@Size.
- The TPCx-AI Price/Performance metric, \$/AIUCpm@Size.
- The **Availability Date** of the complete configuration (see TPC Pricing Specification located on the TPC website (<http://www.tpc.org>)).

7.5.15 The following is an example of compliant reporting of TPCx-AI **Results**:

Example 1: At 10GB the RALF/3000 Server has a TPCx-AI **use case** per minute metric of 3010 when run against a 10GB dataset yielding a TPCX-AI Price/Performance of \$1,202 per **use case**-per-minute and will be available 26-Jan-21.

8 PRICING

8.1 Introduction

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

8.1.1 Pricing Methodology

- 8.1.1.1 A 1-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-AI benchmark.
- 8.1.1.2 The Pricing Model 1 – Default Pricing Model (as defined in the TPC Pricing Specification) is the only pricing model allowed in a TPCx-AI **Result**.

8.2 Priced Configuration

- 8.2.1 The system to be priced must 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.
- 8.2.2 Calculation of the **Priced Configuration** consists of the price of the SUT as tested and defined in Clause 6.1
- 8.2.3 Price of all software licenses for software used in the SUT
- 8.2.4 Price of a communication interface capable of supporting the required number of user interface devices
- 8.2.5 Price of additional products (software or hardware) required for customary operation, administration and maintenance of the SUT for a period of 1 year
- 8.2.6 Price of all products required to create, execute, administer, and maintain the executables or necessary to create and populate the test environment

Comment: Note that Clause 7.3.7.2 explicitly permits data generation to be external to the **SUT** in certain situations. In these situations, the products required for such external to **SUT** data generation would not be priced if the **Auditor** is satisfied that the solution meets the requirements of Clause 7.3.7.2

8.3 Specifically excluded from the Priced Configuration calculation are:

- end-user communication devices and related cables, connectors, and switches.

Comment: end-user communication device here means driver node used to start, stop and orchestrate the benchmark, however devices used to connect to the end-user device with its cable is part of pricing.

- equipment and tools used exclusively in the production of the **Full Disclosure Report**

8.4 Additional Operational Components

- 8.4.1 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**

8.5 Allowable Substitutions

- 8.5.1 **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. This also requires compliance with the TPC Pricing Specification.

Comments: 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.)

- 8.5.2 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 by the **Auditor**. The following hardware components may be substituted:

- durable medium (for example disk drives) and cables
- durable medium enclosure
- network interface card
- router
- bridge
- network switch
- repeater
- cables

Comment: If any hardware component is substituted, then the **Result** must be audited by an **Auditor**

9 FULL DISCLOSURE

9.1 Full Disclosure Report Requirements

9.1.1 A **Full Disclosure Report (FDR)** is required. This section specifies the requirements of the **FDR**.

The **FDR** is a zip file of a directory structure containing the following:

- An **Executive Summary Statement** in Adobe Acrobat PDF format
- A **Report** in Adobe Acrobat PDF format
- The **Supporting Files** consisting of any source files, configuration files, or scripts modified by the **Test Sponsor** and the output files generated by the TPCx-AI kit. Requirements for the **FDR** file directory structure are described below.

9.1.2 The **FDR** should be sufficient to allow an interested reader to evaluate and, if necessary, recreate an implementation of the TPCx-AI **Result** given the appropriate hardware and software products. If any sections in the **FDR** refer to another section of the **FDR**, the names of the referenced scripts/programs must be clearly labeled in each section. Unless explicitly stated otherwise, “disclosed” or “reported” refers to disclosing or reporting in the **FDR**.

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

9.2 Format Guidelines

9.2.1 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.

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

9.2.3 The **Executive Summary** must be included near the beginning of the Report.

9.2.4 The order and titles of sections in the **Report** and **Supporting Files** must correspond with the order and titles of sections from the TPCx-AI Specification (i.e., this document). The intent is to make it as easy as possible for readers to compare and contrast material in different Reports.

9.2.5 The directory structure of the **FDR** has three parts:

- **Executive Summary Statement** - contains **Executive Summary** statement
- **Report** - contains **Report**
- **TPCx-AI_Supporting_Files** Directory

9.3 General Items

9.3.1 The **FDR** must follow all reporting rules of the TPC Pricing Specification, located at www.tpc.org. For clarity and readability, the TPC Pricing Specification requirements may be repeated in the TPCx-AI Specification.

9.3.2 A statement identifying the benchmark **Test Sponsor** and other participating companies must be provided.

9.3.3 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 used by the **SUT**.
- Configuration parameters and options for the **Operating System** and file system components used by the **SUT**.
- Configuration parameters and options for any other software components (e.g. compiler optimization options) used by the **SUT**.

Comment: 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: 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.

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

9.3.5 Diagrams of both measured and **Priced Configurations** must be provided, accompanied by a description of the differences. For physical hardware, 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 disk units, controllers, and if applicable, **LCS** volumes
- 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 if applicable, other hardware components used in the test or are incorporated into the pricing structure
- type and the run-time execution location of software components

Depending on the implementation of the **SUT**, the configuration diagram shown in figure 9-a must include any key functional entities that were used during the execution of the benchmark. Examples of such entities include **Name Node**, Secondary **Name Node**, **Data Node**, Job/Task Tracker, Resource Manager/Node Manager, etc.:



Figure 9-a Sample Configuration Diagram

- $n \times$ Server Rack in scale out configuration.
- $n \times$ My Server Model B, 4/32/64 My CPU Model Z (2.7 GHz, 20MB cache, 130W), 128GB, My RAID Controller with 1GB BBWC
- $n \times$ My Storage Array Model A with 8 X 1TB 10K SAS HDD
- $n \times$ My Switch Model X 10GbE
- $n \times$ Top of the Rack switch.
- LCS results can show LCS instance configuration instead of physical hardware equipment.

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 **Test Sponsors** from using heterogeneous nodes if the system diagram reflects the correct system configuration.

9.4 Software Components and Dataset Distribution

The distribution of software components, roles and dataset across all media must be explicitly described using a format similar to that shown in the following example for the tested and **Priced Configuration**.

Server	Role(s)	Count	Virtual	Host Name(s)	HW/SW Configuration	Storage Setup
Worker	Yarn NM/ Spark Worker	50	N	TPCx-AI[100-150]	<ul style="list-style-type: none"> • Vendor Server Model Name. • HW/SW Config (Processor Model, socket count, Frequency, Core count). • DRAM capacity. • Storage x HDD Model. • Network and BW link speed. • OS Model and version. • Framework SW Model and version. • Graphics Adapters or Accelerators • Details of Additional HW/SW if any. 	OS: Model x GB SSD, Intermediate/Shuffle/T emp Data: x Model x GB SSD, Distributed FS: x Model 12x SAS/SATA Hard drive/
Distro Manger	Cloudera Manager	1	N	TPCx-AI-CM	<ul style="list-style-type: none"> • Vendor Server Model Name. • HW/SW Config (Processor Model, socket count, Frequency, Core count). • DRAM capacity. • Storage x HDD Model. • Network and BW link speed. • OS Model and version. • Framework SW Model and version. • Details of Additional HW/SW if any. 	OS: Model x GB SSD.
Gateway SUT Driver	YARN/SPARK Gateway	1	N	TPCx-AI-Driver 1	<ul style="list-style-type: none"> • Vendor Server Model Name. • HW/SW Config (Processor Model, socket count, Frequency, Core count). • DRAM capacity. • Storage x HDD Model. • Network and BW link speed. • OS Model and version. • Framework SW Model and version. • Details of Additional HW/SW if any. 	
Name Node/Resource Manager	YARN/NN/Zook eeper	1	N	TPCx-AI_NN1	<ul style="list-style-type: none"> • Vendor Server Model Name. • HW/SW Config (Processor Model, socket count, Frequency, Core count). • DRAM capacity. • Storage x HDD Model. • Network and BW link speed. • OS Model and version. • Framework SW Model and version. • Details of Additional HW/SW if any. 	

Table Example Layout Description

- 9.4.1 The distribution of various software components across the system must be explicitly described using a format similar to that shown in Table 9-a in Clause 9.4 for both the tested and **Priced Configuration**.
- Comment:** Software components might vary from across different implementations.
- 9.4.2 The file system used must be disclosed as well as any distributed file system used during the execution of the benchmark and its client API version. The **report** must clearly state in what file system the input dataset was generated and whether the data was moved to a different file system during the Load Test (Clause 7.3.7).
- 9.4.3 All **Frameworks** and tools used in the benchmark should be disclosed in the **report** (e.g. Python, **HDFS**, Spark, YARN, MPI, TensorFlow, Java).
- 9.4.4 If there were any additional vendor supported patches applied to the **SUT**, details of such patches should be disclosed.

9.5 Workload Related Items

- 9.5.1 Any script or text used to set any hardware and software tunable parameters must be included in the **Report**.
- 9.5.2 The version number of the TPCx-AI kit must be Included in the **FDR**.
- 9.5.3 Elapsed times of all **Use Cases** during the power and throughput **tests must be** reported from the **Performance Test**, grouped respectively as **Machine Learning (ML)** or **Deep Learning (DL)**.
- 9.5.4 Completion times for individual **Use Cases** run as part of the **Performance Test** should be included in the **Report**.
- 9.5.4.1 Output **report** from successful **SUT** Validation test must be included in the **Report** (Clause 5.3)
- 9.5.4.2 Global Benchmark **Parameter** files (Clause 5.4.1) must be included in the **Report**.
- 9.5.4.3 **Use case** specific configuration parameters (Clause 5.4.1) must be included in the **Report**.

9.6 SUT Related Items

- 9.6.1 Details of any Specialized Hardware/Software used in the **SUT** must be included in the **report**.
- 9.6.2 Relevant **Framework** configuration files from **SUT**, for the **Performance Test** must be included in the **FDR** e.g. Yarn-Site.xml, Hdfs-site.xml etc.
- 9.6.3 **General execution environment** information as well as any special environment configuration that is relevant to the benchmark must be included in the **report** in form of envinfo.log from a representative worker node from every role in the server.
- 9.6.4 The data storage ratio must be disclosed. It is computed by dividing the total physical data storage present in the **Priced Configuration** (expressed in GB) by the chosen Scale Factor as defined in Clause 3.1.3. 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 **SUT** configured with 96 disks of 1000GB capacity for a 100 Scale Factor has a data storage ratio of 960.
- Comment:** For consumption-based storage provisioning in **LCS**, the maximum storage provisioned during the entire benchmark run is considered to be the total physical data storage present.
- 9.6.5 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**. 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 1000GB of physical memory for a 10000GB Scale Factor has a memory ratio of 10.00.
- Comment:** For **LCS**, the maximum provisioned memory during the entire benchmark run is considered to be the total physical memory present.

- 9.6.6 All the files in the benchmark output directory and the TPCx-AI database file (TPCx-AI.db) must be included in the **Supporting Files** folder.
- 9.6.7 If the **Test sponsor** would like to use additional post-processing scripts and steps before starting a Benchmark run, the **Test Sponsor** would have to create a separate directory called 'additional_files' and create a special package at the end of the run for the audit that includes the 'additional_files' directory structure.
- 9.6.8 If the test sponsor performs model optimizations after the Power Training Test, then the additional files directory should contain the original model files (the output of the Power Training Test and relevant checksums), the output model file(s) as a result of the optimizations (along with the new relevant checksum) and the necessary files used to perform all the optimizations. The **Test Sponsor** should also disclose the tools used and any parameters used for the benchmark run.

9.7 Metrics and Scale Factors

- 9.7.1 The **Reported Performance Metric** (AIUCpm@SF for the Benchmark **Run**) must be disclosed in the **Report**.
- 9.7.2 The **Price/Performance Metric** (\$/AIUCpm@SF) for the Benchmark **Performance Test** must be disclosed in the **Report**.
- 9.7.3 The Scale Factor used for the **Result** must be disclosed in the **Report**.
- 9.7.4 The number of streams in the throughput run used for the **Result** must be disclosed in the **Report**.
- 9.7.5 The total elapsed time for the execution of the Benchmark **Performance Test** must be disclosed in the **Report**.
- 9.7.6 The start time, end time and total elapsed time for each of the six tests (Clause 7.3.1) must be disclosed for the Benchmark **Performance Test**.
- 9.7.7 The scoring metrics resulting from the Scoring test of the Benchmark **Performance Test** must be disclosed in the report. All timestamps must be expressed with a precision of at least 1/1000 of a second.

9.8 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 **Report**. A statement should be included specifying whom to contact to obtain further information regarding the audit process.

9.8.1 Executive Summary Statement

The **Executive Summary** is a high-level overview of a TPCx-AI 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.

9.8.2 The Executive Summary has three components:

- Implementation Overview
- Pricing Table
- Numerical Quantities

9.8.3 Page Layout

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.

9.8.4 Implementation Overview

The implementation overview page contains six 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.

9.8.4.1 The first section contains information about the 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-AI, TPC-Pricing and TPC-Energy (if reported)	16-20 pt. Bold
Report Date	16-20t. Bold

Table Sponsor and System Identification

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

Comment: The **report** date must be disclosed with a precision of 1 day. The precise format is left to the **Test Sponsor**.

9.8.4.2 The second section contains the Total System Cost, the TPCx-AI **Reported Performance Metric** and the **Price/Performance Metric**.

Title	Quantity	Precision	Font
Total System Cost	1 yr. Cost of ownership (Clause)	1	16-20 pt. Bold
Reported Performance Metric	AIUCpm@SF (Clause 7.5.7)	0.01	16-20 pt. Bold
Price/Performance	\$/ AIUCpm@SF (Clause 7.5.8)	0.01	16-20 pt. Bold

Table Benchmark Results

9.8.4.2.1 For results using currency conversions, refer to the current version of the TPC pricing specification for the process.

9.8.4.3 The third section contains detailed diagrams of the measured configuration (Clause 9.3.5) and the Software components distribution table (Clause 9.4)

Title	Quantity	Font
Framework /Engine 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 the TPC Pricing Specification	9-12 pt. Times
SF (Scale Factor)	SF in as defined in Clause 3.1.3	16-20pt. Bold
Streams	Concurrent Streams used in the Throughput Test (Clause 7.3.11)	16-20pt. Bold

Table System Configuration Information

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

9.8.4.4 The fourth section contains diagrams showing the Results for the Power Training Test, Power serving Tests and Throughput Test.

9.8.4.5 The fifth section contains the storage and memory ratios, see (Clause 9.6.4.)

Title	Precision	Font
Physical Storage/Scale Factor	0.01	9-12 pt. Times
Scale Factor/Physical Memory	0.01	9-12 pt. Times
Main Data Redundancy Model	n/a	9-12 pt. Times

Table Storage and Memory Ratios

9.8.4.6 The sixth section contains the components, 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, see (Clause 9.3.5)
- main and cache memory sizes
- network speed and topology (e.g. Top of the Rack switch, in-rack switch)
- storage type, quantity and configuration.

9.8.5 Pricing Table

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

- network(s)
- server(s) /node(s)
- storage
- software

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

9.8.6 Numerical Quantities Summary

9.8.6.1 The Numerical Quantities Summary page contains six sets of runtime data, presented in tabular form, detailing the Load Test, Power Training Test, Power Serving Test I, Power Serving Test II, Throughput Test, and Scoring Test for the **Benchmark Run**. Each set of data should be headed by its given title and clearly separated from the other tables.

9.8.6.2 The first section contains the Scale Factor, Number of streams, and **SUT** Validation test status for the reported benchmark execution **Result**.

9.8.6.3 The second section contains measurement results and metric from the **Performance Test**.

Performance Test	
Item Title	Precision
Overall Run Start Time	yyyy-mm-dd hh:mm:ss.sss
Overall Run End Time	yyyy-mm-dd hh:mm:ss.sss
Overall Run Total Elapsed Time	hh:mm:ss.sss
Start of Load Test	yyyy-mm-dd hh:mm:ss.sss
End of Load Test	yyyy-mm-dd hh:mm:ss.sss
Load Test Elapsed Time	hh:mm:ss.sss
Start of Power Training Test	yyyy-mm-dd hh:mm:ss.sss
End of Power Training Test	yyyy-mm-dd hh:mm:ss.sss
Power Training Test Elapsed Time	hh:mm:ss.sss
Start of Power Serving Test I	yyyy-mm-dd hh:mm:ss.sss
End of Power Serving Test I	yyyy-mm-dd hh:mm:ss.sss
Power Serving Test I Elapsed Time	hh:mm:ss.sss
Start of Power Serving Test II	yyyy-mm-dd hh:mm:ss.sss
End of Power Serving Test II	yyyy-mm-dd hh:mm:ss.sss
Power Serving Test II Elapsed Time	hh:mm:ss.sss
Start of Throughput Test	yyyy-mm-dd hh:mm:ss.sss
End of Throughput Test	yyyy-mm-dd hh:mm:ss.sss
Throughput Test Elapsed Time	hh:mm:ss.sss
Start of Scoring Test	yyyy-mm-dd hh:mm:ss.sss
End of Scoring Test	yyyy-mm-dd hh:mm:ss.sss
Scoring Test Elapsed Time	hh:mm:ss.sss

Performance Test	
Item Title	Precision
Performance Metric (AIUCpm@SF)	xx.xxxxx

Table Measurement Result for **Performance Test**

9.8.6.4 The Third section contains Individual loading, training and serving times and accuracy metrics for each **use case**.

Use Case	Load Time	Training Time	Serving Time1	Serving Time 2	Throughput Time	Accuracy Metric
UC01	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	xx.xxxxx
UC02	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	xx.xxxxx
UC03	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	xx.xxxxx
UC04	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	xx.xxxxx
UC05	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	xx.xxxxx
UC06	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	xx.xxxxx
UC07	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	xx.xxxxx
UC08	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	xx.xxxxx
UC09	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	xx.xxxxx
UC10	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	hh:mm:ss.sss	xx.xxxxx

Table Use case performance Metrics

9.8.7 TPCx-AI Run Report

The Run **Report** per Clause 5.3 from TPCx-AI must be included in the **Report** immediately after the **Executive Summary**.

9.8.8 “TPCx-AI_Supporting_Files” Directory

This directory must contain all the files and directories needed for completing the audit. This directory will contain all the contents of the ‘TPCx-AI_Benchmarkrun’ directory after the completion of the Benchmark Run.

9.9 Availability of the Full Disclosure Report

The **Full Disclosure Report** must be readily available to the public. The **Report** and **Supporting Files** must be made available when the **Result** is made public. In order to use the phrase “TPCx-AI Benchmark”, the **Full Disclosure Report** must have been previously submitted electronically to the TPC using the procedure described in the TPC Policies and Guidelines document.

9.9.1 The official **Full Disclosure Report** must be available in English but may be translated to additional languages.

9.10 Revisions to the Full Disclosure Report

The requirements for a revision to an **FDR** are specified in the TPC Pricing Specification.

10 AUDIT

This clause defines the audit requirements for TPCx-AI. The **Auditor** needs to ensure that the benchmark under audit complies with the TPCx-AI specification. 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.

10.1 General Rules

Before publication, a TPCx-AI **Result** must be certified to be compliant with the spirit and letter of the TPCx-AI Benchmark Standard by an Independent Certified TPC **Auditor** or a TPCx-AI Pre-Publication Board. The **Test Sponsor** can choose the certification be performed by either by a Certified TPC **Auditor** or a Pre-Publication Board.

10.2 Independent Audit

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 independent audit of the benchmark is described in TPC Policies on www.tpc.org. In addition, the auditing agency cannot have supplied any performance consulting under contract for the benchmark.

In addition, the following conditions must be met:

- a) 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.
- b) The auditing agency cannot be financially related to any one of the suppliers of the measured/**Priced Configuration**, e.g., the DBMS supplier, the disk supplier, etc.

10.2.1 The **Auditor**’s opinion regarding the compliance of a **Result** must be consigned in an **Attestation Letter** delivered directly to the **Test Sponsor**. To document that a **Result** has been audited, the **Attestation Letter** must be included in the **Report** and made readily available to the public. Upon request, and after approval from the **Test Sponsor**, a detailed audit **report** may be produced by the **Auditor**.

10.2.2 Pre-Publication Board

The term Pre-Publication Board as defined by the TPC Policies consists of three or more individuals that have been chosen by the TPCx-AI Benchmark Subcommittee to certify **Results** for publication. For TPCx-AI **Results** the Pre-Publication Board consists of 3 members from the TPCx-AI Benchmark Subcommittee. Each member serves a period of six months. The membership will be rotated through the TPCx-AI Benchmark Subcommittee membership. The submission is confidential to the Pre-Publication Board until the **Result** is published. The operating rules of the Pre-Publication Board are found in the current revision of the TPC Policies.

10.2.3 Results Based on Existing TPCx-AI Results

A **Test Sponsor** can demonstrate compliance of a new **Result** produced without running any **Performance Test** by referring to the certification of a **Result**, if the following conditions are all met:

- The referenced **Result** has already been published by the same or by another **Test Sponsor**.
- The new **Result** must have the same hardware and software architecture and configuration as the referenced **Result**. The only exceptions allowed are for elements not involved in the processing logic of the **SUT** (e.g., number of peripheral slots, power supply, cabinetry, fans, etc.)
- The **Test Sponsor** of the already published **Result** gives written approval for its use as referenced by the **Test Sponsor** of the new **Result**.
- The **Auditor** verifies that there are no significant functional differences between the priced components used for both **Results** (i.e., differences are limited to labeling, packaging and pricing.)
- The **Auditor** or Pre-Publication Board reviews the **FDR** of the new **Result** for compliance.

- If certification is performed by an independent **Auditor**, a new **Attestation Letter** must be included in the **Report** of the new **Result**.

Comment: The intent of this clause is to allow publication of benchmarks for systems with different packaging and model numbers that are considered to be identical using the same **Performance Test**. For example, a rack mountable system and a freestanding system with identical electronics can use the same **Performance Test** for publication, with, appropriate changes in pricing.

Comment: Although it should be apparent to a careful reader that the **FDR** for the two **Results** are based on the same set of **Performance Tests**, the **FDR** for the new **Result** is not required to explicitly state that it is based on the **Performance Tests** of another published **Result**.

Comment: When more than one **Result** is published based on the same set of **Performance Tests**, only one of the **Results** from this group can occupy a numbered slot in each of the benchmark **Result** “Top Ten” lists published by the TPC. The **Test Sponsors** of this group of **Results** must all agree on which **Result** from the group will occupy the single slot. In case of disagreement among the **Test Sponsors**, the decision will default to the **Test Sponsor** of the earliest publication from the group.

10.3 Audit Checklist

A generic audit checklist is provided as part of this specification. The **generic** audit checklist specifies the requirements that must be checked to ensure a **Result** is compliant with the TPCx-AI Specification. Not only should the TPCx-AI requirements be checked for accuracy but the **Auditor** or Pre-Publication Board must ensure that the **FDR** accurately reflects the **Result**.

Comment: An independent **Auditor** must be used for those audit checklist items that refer to pricing or energy.


- 10.3.1 Verify the version of the TPCx-AI provided kit used.
- 10.3.2 Verify that all 6 tests (Load, Power Training, Power Serving I, Power Serving II, Throughput, Scoring) (Clause 7.3) of the **Performance Test** completed with no error reported.
- 10.3.3 Verify Validation tests (Clause 7.2) is completed with no error reported.
- 10.3.4 Verify Benchmark Execution has been completed according to the sequence of scripts listed in Clause 7.1 and the sequence of tests in Clause 7.3
- 10.3.5 Verify that CHECK_INTEGRITY phase was included in the **Performance Test**.
- 10.3.6 Verify that all scripts and source code to implement the benchmark has been included in the **Report**.
- 10.3.7 Verify Kit run **report** contains all information mentioned in Clause 5.3
- 10.3.8 Verify Test Sponsor Kit Modifications Clause 5.6 have been followed to ensure the parameter settings were set as defined in the specification and required reports, files are provided as part of the **FDR**.
- 10.3.9 Verify Clause 5.6.6 and 5.6.7 is followed and no Java code files were modified and no JAR file optimizers were used.
- 10.3.10 Verify the test execution has produced the required output by:
 - checking the tpcx-ai.db logfiles to see if the **training test** has created all the model files and scoring test complete with accuracy metrics in the **Report**.

- 10.3.11 Verify that the **SUT** Pricing as reported by the Test Sponsor is in compliance with the TPC Pricing specification.
- Comment:** The **Auditor** should also review the **SUT** pricing details in Clause 6.2 to verify the **SUT** Pricing Report
- 10.3.12 Verify that all components of the **SUT** are commercially available as per the TPC Pricing Specification.
- 10.3.13 Verify that all components of the **SUT** are included in the pricing.
- 10.3.14 Verify no aspect of **SUT** including tuning parameters were changed between the validation test and the **Performance Test**.
- 10.3.15 Verify that the SF used for publication is valid according to Clause 3.1.3.
- 10.3.16 Verify that the metrics are **Reported** as per the requirements in Clause 7.5 and there are no errors in the **Report**
- 10.3.17 Verify that the Energy **report** is in compliance with the TPC Energy specification (if reported).
- 10.3.18 Verify that Full Disclosure **Report** and **Executive Summary** Reports are accurately reported and comply with the reporting requirements. This includes but not limited to.
- metric calculation
 - system availability
 - the diagrams of both measured and **Priced Configuration**
 - system pricing
 - the numerical quantity summary
 - Parameter files required as part of **FDR** are provided.

Appendix A: PDGF USER guide

A.1 **PDGF** is a TPC provided software package designed to facilitate the data generation of the TPCx-AI **Test Dataset**. This appendix provides information on how a Test Sponsor is to use the features and functionality of **PDGF**.

A.2 **PDGF** generates the **Test Dataset** in an output folder along with the dataset for serving and scoring.



A.3 With **PDGF** the **Test Sponsor** specifies the scale factor to be used and accordingly a unified dataset is generated where the size of the **Test dataset** will be very close to the scale factor specified.

Appendix B: Benchmark kit parameters

- B.1 The configuration files in the '/driver/dat' directory cannot be modified.
- B.2 The following benchmark parameters can be changed with the setenv.sh script.
- B.3 TPCxAI_SCALE_FACTOR: This parameter determines the size of the synthetic data set that will be generated and that will be the input to the use cases execution.
- B.4 TPCxAI_SERVING_THROUGHPUT_STREAMS: This parameter determines the number of parallel streams for the SERVING_THROUGHPUT test.
- B.5 TPCxAI_CONFIG_FILE_PATH: This parameter is used by the test sponsor to specify the configuration file that will be used to run the benchmark.
- B.6 TPCxAI_ENV_TOOLS_DIR: Location of the directory containing scripts to collect system configuration information. By convention this is a subdirectory of the tools directory.
- B.7 For the spark implementation in the kit:
- B.8 YARN_CONF_DIR: Location of the YARN configuration directory.
- B.9 PYSARK_PYTHON: Python executable used by Apache Spark. This must be accessible to all Apache Spark workers.
- B.10 PYSARK_DRIVER_PYTHON: Python executable used by Apache Spark. This must be accessible only by the driver node.

Appendix C: Use case specific parameters

- C.1 The configuration files in the '/driver/config' directory can be modified.
- C.2 Some use cases provide a set of settings that can be changed by the test sponsor prior to running the benchmark. Only the parameters listed below can be changed. If not stated otherwise the parameters apply to all reference implementations.
- C.3 The Scoring Test will use the same parameters as the Serving tests.
- C.4 Use case 2
- Epochs: Number of epochs for which the neural network will be trained. The use case will use a minimum of 5 epochs for training.
 - Batch: Batch size used during the Power Training Test and all the Serving tests. Different batch sizes can be used for training and for serving.
 - Learning_rate: Learning rate used during the Power Training Test has an impact on how fast the model can converge. This parameter needs to be adapted when training is done in parallel so that it meets the accuracy criteria set forth by the specification.
 - Task_cpus_horovod: The number of cores assigned to an Apache Spark task; this setting is used by Horovod during the Power Training Test. (Applies to Apache Spark only)
 - Executor_cores_horovod: The number of cores assigned to an Apache Spark executor; this setting is used by Horovod during the Power Training Test. (Applies to Apache Spark only)
- C.5 Use case 5
- Epochs: Number of epochs for which the neural network will be trained. The use case will use a minimum of 5 epochs for the Power Training Test.
 - Batch: Batch size used during the Power Training Test and all the Serving tests. Different batch sizes can be used for the Power Training Test and all serving tests
 - Learning_rate: Learning rate used during the Power Training Test has an impact on how fast the model can converge. This parameter needs to be adapted during the Power Training Test in parallel so that it meets the accuracy criteria set forth by the specification.
- C.6 Use case 7
- Num-blocks: The number of blocks used to parallelize computation, which has to be a positive integer. (Applies to Apache Spark only)
- C.7 Use case 8
- num-workers: Controls how many parallel workers can be used when running the Power Training Test for the XGBoost Classification Model.(Applies to Apache Spark only)num-threads: The number of threads used by each XGBoost worker. Spark requires that all of num-thread * num-Workers cores should be available before the start of the Power Training Test.(Applies to Apache Spark only)
 - num_rounds: The number of boosting rounds
- C.8 Use case 9
- Epochs: Number of epochs for which the neural network will be trained. The use case will use a minimum of 3 epochs for training.
 - Batch: Batch size used during the Power Training Test and all the Serving tests. Different batch sizes can be used for training and for serving.
 - Learning_rate: Learning rate used during the Power Training Test has an impact on how fast the model can converge. This parameter needs to be adapted when training is done in parallel so that it meets the accuracy criteria set forth by the specification.

- `Task_cpus_horovod`: The number of cores assigned to an Apache Spark task; this setting is used by Horovod during the Power Training Test. (Applies to Apache Spark only)
- `Executor_cores_horovod`: The number of cores assigned to an Apache Spark executor; this setting is used by Horovod during the Power Training Test. (Applies to Apache Spark only)

Appendix D: Sample Default Configuration file

D.1 Sample default configuration file for the python reference implementation.

Local filesystem on Windows, Linux, or MacOS

```
local_fs: &LOCAL_FS
name: "local_fs"
create: "tools/python/create.sh $destination"
load: "tools/python/load.sh $destination $source"
copy: "cp -f $source $destination"
delete: "rm -rf $destination"
delete_parallel: "pssh -t 0 -P -h nodes rm -rf $destination"
download: "cp $source $destination"
```

HDFS = Hadoop Distributed Filesystem

```
hdfs: &HDFS
name: "hdfs"
create: "tools/spark/create_hdfs.sh $destination"
load: "tools/spark/load_hdfs.sh $destination $source"
copy: "hdfs dfs -cp -f $source $destination"
delete: "hdfs dfs -rm -r -f -skipTrash $destination"

download: 'hdfs dfs -cat $source/* | awk "BEGIN{f=""}{if($0!=f){print $0}if(NR==1){f=$0}}" >
$destination/predictions.csv'
```

workload:

global definitions

```
engine_executable: &ENGINE "lib/python-venv/bin/python"
training_template: &TRAINING_TEMPLATE "$engine $name --stage training --workdir $output $input/$file"
serving_template: &SERVING_TEMPLATE "$engine $name --stage serving --workdir $model --output
$model/$phase $input/$file"
serving_throughput_template: &SERVING_THROUGHPUT_TEMPLATE "$engine $name --stage serving --
workdir $model --output $model/$stream $input/$file"
training_data_url: &TRAINING_DATA_URL "output/data/training"
serving_data_url: &SERVING_DATA_URL "output/data/serving"
scoring_data_url: &SCORING_DATA_URL "output/data/scoring"
datagen_datastore: *LOCAL_FS
include_datagen_in_tload: False
```

general/ benchmark-wide configuration parameters

```
pdgf_node_parallel: False
pdgf_home: "lib/pdgf"
raw_data_url: "output/raw_data"
temp_dir: '/tmp/tpcxai'
```

usecases:

```
1:
# general
name: "-m workload.UseCase01"
# engines
training_engine: *ENGINE
serving_engine: *ENGINE
# data stores
training_datastore: *LOCAL_FS # for storing the training data
model_datastore: *LOCAL_FS # for storing the trained models
```

```

serving_datastore: *LOCAL_FS # for storing the serving data
output_datastore: *LOCAL_FS # for storing the final output
# templates
datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"
training_template: "$engine $name --stage training --num_clusters 4 --workdir $output $input/order.csv
$input/lineitem.csv $input/order_returns.csv"
serving_template: "$engine $name --stage serving --workdir $model --output $model/$phase $input/order.csv
$input/lineitem.csv $input/order_returns.csv"
serving_throughput_template: "$engine $name --stage serving --workdir $model --output $model/$stream
$input/order.csv $input/lineitem.csv $input/order_returns.csv"
# URLs
training_data_url: *TRAINING_DATA_URL
serving_data_url: *SERVING_DATA_URL
scoring_data_url: *SCORING_DATA_URL
model_url: "output/model/uc01"
output_url: "output/output/uc01"
scoring_output_url: "output/scoring/uc01"

```

2:

```

# general
name: "-m workload.UseCase02"
# engines
training_engine: *ENGINE
serving_engine: *ENGINE
# data stores
training_datastore: *LOCAL_FS # for storing the training data
model_datastore: *LOCAL_FS # for storing the trained models
serving_datastore: *LOCAL_FS # for storing the serving data
output_datastore: *LOCAL_FS # for storing the final output
# templates
datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"
training_template: "$engine $name --stage training --epochs 25 --batch 32 --workdir $output $input/$file"
serving_template: "$engine $name --stage serving --batch 32 --workdir $model --output $model/$phase
$input/$file"
serving_throughput_template: "$engine $name --stage serving --batch 32 --workdir $model --output
$model/$stream $input/$file"
# URLs
training_data_url: *TRAINING_DATA_URL
serving_data_url: *SERVING_DATA_URL
scoring_data_url: *SCORING_DATA_URL
model_url: "output/model/uc02"
output_url: "output/output/uc02"
scoring_output_url: "output/scoring/uc02"

```

3:

```

# general
name: "-m workload.UseCase03"
# engines
training_engine: *ENGINE
serving_engine: *ENGINE
# data stores
training_datastore: *LOCAL_FS # for storing the training data
model_datastore: *LOCAL_FS # for storing the trained models
serving_datastore: *LOCAL_FS # for storing the serving data
output_datastore: *LOCAL_FS # for storing the final output
# templates
datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"

```

```

    training_template: "$engine $name --stage training --workdir $output $input/order.csv $input/lineitem.csv
$input/product.csv"
    serving_template: "$engine $name --stage serving --workdir $model --output $model/$phase
$input/store_dept.csv"
    serving_throughput_template: "$engine $name --stage serving --workdir $model --output $model/$stream
$input/store_dept.csv"
    # URLs
    training_data_url: *TRAINING_DATA_URL
    serving_data_url: *SERVING_DATA_URL
    scoring_data_url: *SCORING_DATA_URL
    model_url: "output/model/uc03"
    output_url: "output/output/uc03"
    scoring_output_url: "output/scoring/uc03"

```

4:

```

# general
name: "-m workload.UseCase04"
# engines
training_engine: *ENGINE
serving_engine: *ENGINE
# data stores
training_datastore: *LOCAL_FS # for storing the training data
model_datastore: *LOCAL_FS # for storing the trained models
serving_datastore: *LOCAL_FS # for storing the serving data
output_datastore: *LOCAL_FS # for storing the final output
# templates
datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"
training_template: *TRAINING_TEMPLATE
serving_template: *SERVING_TEMPLATE
serving_throughput_template: *SERVING_THROUGHPUT_TEMPLATE
# URLs
training_data_url: *TRAINING_DATA_URL
serving_data_url: *SERVING_DATA_URL
scoring_data_url: *SCORING_DATA_URL
model_url: "output/model/uc04"
output_url: "output/output/uc04"
scoring_output_url: "output/scoring/uc04"

```

5:

```

# general
name: "-m workload.UseCase05"
# engines
training_engine: *ENGINE
serving_engine: *ENGINE
# data stores
training_datastore: *LOCAL_FS # for storing the training data
model_datastore: *LOCAL_FS # for storing the trained models
serving_datastore: *LOCAL_FS # for storing the serving data
output_datastore: *LOCAL_FS # for storing the final output
# templates
datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"
training_template: "$engine $name --stage training --epochs 15 --batch 512 --workdir $output $input/$file"
serving_template: "$engine $name --stage serving --batch 512 --workdir $model --output $model/$phase
$input/$file"
serving_throughput_template: "$engine $name --stage serving --batch 512 --workdir $model --output
$model/$stream $input/$file"
# URLs
training_data_url: *TRAINING_DATA_URL

```

```
serving_data_url: *SERVING_DATA_URL
scoring_data_url: *SCORING_DATA_URL
model_url: "output/model/uc05"
output_url: "output/output/uc05"
scoring_output_url: "output/scoring/uc05"
```

6:

```
# general
name: "-m workload.UseCase06"
# engines
training_engine: *ENGINE
serving_engine: *ENGINE
# data stores
training_datastore: *LOCAL_FS # for storing the training data
model_datastore: *LOCAL_FS # for storing the trained models
serving_datastore: *LOCAL_FS # for storing the serving data
output_datastore: *LOCAL_FS # for storing the final output
# templates
datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"
training_template: *TRAINING_TEMPLATE
serving_template: *SERVING_TEMPLATE
serving_throughput_template: *SERVING_THROUGHPUT_TEMPLATE
# URLs
training_data_url: *TRAINING_DATA_URL
serving_data_url: *SERVING_DATA_URL
scoring_data_url: *SCORING_DATA_URL
model_url: "output/model/uc06"
output_url: "output/output/uc06"
scoring_output_url: "output/scoring/uc06"
```

7:

```
# general
name: "-m workload.UseCase07"
# engines
training_engine: *ENGINE
serving_engine: *ENGINE
# data stores
training_datastore: *LOCAL_FS # for storing the training data
model_datastore: *LOCAL_FS # for storing the trained models
serving_datastore: *LOCAL_FS # for storing the serving data
output_datastore: *LOCAL_FS # for storing the final output
# templates
datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"
training_template: *TRAINING_TEMPLATE
serving_template: *SERVING_TEMPLATE
serving_throughput_template: *SERVING_THROUGHPUT_TEMPLATE
# URLs
training_data_url: *TRAINING_DATA_URL
serving_data_url: *SERVING_DATA_URL
scoring_data_url: *SCORING_DATA_URL
model_url: "output/model/uc07"
output_url: "output/output/uc07"
scoring_output_url: "output/scoring/uc07"
```

8:

```
# general
name: "-m workload.UseCase08"
# engines
```



```

training_engine: *ENGINE
serving_engine: *ENGINE
# data stores
training_datastore: *LOCAL_FS # for storing the training data
model_datastore: *LOCAL_FS # for storing the trained models
serving_datastore: *LOCAL_FS # for storing the serving data
output_datastore: *LOCAL_FS # for storing the final output
# templates
datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"
training_template: "$engine $name --stage training --workdir $output $input/order.csv $input/lineitem.csv
$input/product.csv"
serving_template: "$engine $name --stage serving --workdir $model --output $model/$phase $input/order.csv
$input/lineitem.csv $input/product.csv"
serving_throughput_template: "$engine $name --stage serving --workdir $model --output $model/$stream
$input/order.csv $input/lineitem.csv $input/product.csv"
# URLs
training_data_url: *TRAINING_DATA_URL
serving_data_url: *SERVING_DATA_URL
scoring_data_url: *SCORING_DATA_URL
model_url: "output/model/uc08"
output_url: "output/output/uc08"
scoring_output_url: "output/scoring/uc08"

```

9:

```

# general
name: "-m workload.UseCase09"
# engines
training_engine: *ENGINE
serving_engine: *ENGINE
# data stores
training_datastore: *LOCAL_FS # for storing the training data
model_datastore: *LOCAL_FS # for storing the trained models
serving_datastore: *LOCAL_FS # for storing the serving data
output_datastore: *LOCAL_FS # for storing the final output
# templates
datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"
training_template: "$engine $name --stage training --epochs_embedding=15 --batch=64 --workdir $output
$input/CUSTOMER_IMAGES"
serving_template: "$engine $name --stage serving --batch=64 --workdir $model --output $model/$phase
$input/CUSTOMER_IMAGES"
serving_throughput_template: "$engine $name --stage serving --batch=64 --workdir $model --output
$model/$stream $input/CUSTOMER_IMAGES"
# URLs
training_data_url: *TRAINING_DATA_URL
serving_data_url: *SERVING_DATA_URL
scoring_data_url: *SCORING_DATA_URL
model_url: "output/model/uc09"
output_url: "output/output/uc09"
scoring_output_url: "output/scoring/uc09"

```

10:

```

# general
name: "-m workload.UseCase10"
# engines
training_engine: *ENGINE
serving_engine: *ENGINE
# data stores
training_datastore: *LOCAL_FS # for storing the training data

```

```

model_datastore: *LOCAL_FS    # for storing the trained models
serving_datastore: *LOCAL_FS  # for storing the serving data
output_datastore: *LOCAL_FS   # for storing the final output
# templates
datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"
training_template: "$engine $name --stage training --workdir $output $input/financial_account.csv
$input/financial_transactions.csv"
serving_template: "$engine $name --stage serving --workdir $model --output $model/$phase
$input/financial_account.csv $input/financial_transactions.csv"
serving_throughput_template: "$engine $name --stage serving --workdir $model --output $model/$stream
$input/financial_account.csv $input/financial_transactions.csv"
# URLs
training_data_url: *TRAINING_DATA_URL
serving_data_url: *SERVING_DATA_URL
scoring_data_url: *SCORING_DATA_URL
model_url: "output/model/uc10"
output_url: "output/output/uc10"
scoring_output_url: "output/scoring/uc10"

```

D.2 Sample default configuration file for the spark reference implementation.

```
# Local filesystem on Windows, Linux, or MacOS
local_fs: &LOCAL_FS
  name: "local_fs"
  create: "tools/python/create.sh $destination"
  load: "tools/python/load.sh $destination $source"
  copy: "cp -f $source $destination"
  delete: "rm -rf $destination"
  delete_parallel: "pssh -t 0 -P -h nodes rm -rf $destination"
  download: "cp $source $destination"

# HDFS = Hadoop Distributed Filesystem
hdfs: &HDFS_local
  name: "hdfs"
  create: "tools/spark/create_hdfs.sh $destination"
  load: "tools/spark/load_hdfs.sh $destination $source"
  copy: "hdfs dfs -cp -f $source $destination"
  delete: "hdfs dfs -rm -r -f -skipTrash $destination"
  download: 'hdfs dfs -cat $source/* | awk "BEGIN{f=\""} {if($0!=f){print $0}if(NR==1){f=$0}}" >
$destination/predictions.csv'

hdfs_parallel: &HDFS
  name: "hdfs"
  create: "tools/spark/create_hdfs.sh $destination"
  load: "tools/parallel-data-load.sh nodes 1 $destination $source"
  copy: "hdfs dfs -cp -f $source $destination"
  delete: "hdfs dfs -rm -r -f -skipTrash $destination"
  download: 'hdfs dfs -cat $source/* | awk "BEGIN{f=\""} {if($0!=f){print $0}if(NR==1){f=$0}}" >
$destination/predictions.csv'

workload:
  # global definitions
  engine_executable: &ENGINE "spark-submit --conf spark.executorEnv.NUMBA_CACHE_DIR=/tmp --jars
'Stpcxai_home/lib/*.jar' --executor-cores 5 --num-executors 1 --driver-memory 10g --executor-memory 40g --conf
spark.yarn.executor.memoryOverhead=4g --conf spark.kryoserializer.buffer.max=1g --conf
spark.rpc.message.maxSize=1024 --conf spark.executor.extraJavaOptions='-Xss128m' --driver-java-options '-
Xss128m' --master yarn --deploy-mode client "
  engine_executable_dl2: &ENGINE_DL2 "spark-submit --conf spark.executorEnv.NUMBA_CACHE_DIR=/tmp --
jars 'Stpcxai_home/lib/*.jar' --num-executors 1 --executor-cores 1 --driver-memory 10g --executor-memory 40g --conf
spark.yarn.executor.memoryOverhead=4g --conf spark.rpc.message.maxSize=1024 --conf
spark.kryoserializer.buffer.max=1g --conf spark.executor.extraJavaOptions='-Xss128m' --driver-java-options '-
Xss128m' --master yarn --deploy-mode client "
  engine_executable_dl5: &ENGINE_DL5 "spark-submit --conf spark.executorEnv.NUMBA_CACHE_DIR=/tmp --
jars 'Stpcxai_home/lib/*.jar' --num-executors 1 --executor-cores 1 --driver-memory 10g --executor-memory 40g --conf
spark.yarn.executor.memoryOverhead=4g --conf spark.rpc.message.maxSize=1024 --conf
spark.kryoserializer.buffer.max=1g --conf spark.executor.extraJavaOptions='-Xss128m' --driver-java-options '-
Xss128m' --master yarn --deploy-mode client "
  engine_executable_9: &SERVING_9 "spark-submit --conf spark.executorEnv.NUMBA_CACHE_DIR=/tmp --jars
'Stpcxai_home/lib/*.jar' --executor-cores 5 --conf spark.task.cpus=1 --num-executors 1 --driver-memory 10g --
executor-memory 40g --conf spark.yarn.executor.memoryOverhead=4g --conf spark.kryoserializer.buffer.max=1g --
conf spark.rpc.message.maxSize=1024 --conf spark.executor.extraJavaOptions='-Xss128m' --driver-java-options '-
Xss128m' --master yarn --deploy-mode client "

  datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"
  training_template: &TRAINING_TEMPLATE "$engine --class $name lib/workload-assembly-0.1.jar --stage training
--workdir $output $input/$file"
  serving_template: &SERVING_TEMPLATE "$engine --class $name lib/workload-assembly-0.1.jar --stage serving -
-workdir $model --output $model/$phase $input/$file"
```

```

serving_throughput_template: &SERVING_THROUGHPUT_TEMPLATE "$engine --class $name lib/workload-assembly-0.1.jar --stage serving --workdir $model --output $model/$stream $input/$file"
training_data_url: &TRAINING_DATA_URL "output/data/training"
serving_data_url: &SERVING_DATA_URL "output/data/serving"
scoring_data_url: &SCORING_DATA_URL "output/raw_data/scoring"
datagen_datastore: *LOCAL_FS

# general/ benchmark-wide configuration parameters
pdgf_node_parallel: True
pdgf_home: "lib/pdgf"
raw_data_url: "output/raw_data"
temp_dir: '/tmp/tpcxai'
usecases:
1:
  # general
  name: "org.tpc.tpcxai.UseCase01"
  # engines
  training_engine: *ENGINE
  serving_engine: *ENGINE
  # data stores
  training_datastore: *HDFS # for storing the training data
  model_datastore: *HDFS # for storing the trained models
  serving_datastore: *HDFS # for storing the serving data
  output_datastore: *HDFS # for storing the final output
  # templates
  datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"
  training_template: "$engine --class $name lib/workload-assembly-0.1.jar --stage training --num_clusters 4 --workdir $output $input/order.csv $input/lineitem.csv $input/order_returns.csv"
  serving_template: "$engine --class $name lib/workload-assembly-0.1.jar --stage serving --workdir $model --output $model/$phase $input/order.csv $input/lineitem.csv $input/order_returns.csv"
  serving_throughput_template: "$engine --class $name lib/workload-assembly-0.1.jar --stage serving --workdir $model --output $model/$stream $input/order.csv $input/lineitem.csv $input/order_returns.csv"
  # URLs
  training_data_url: *TRAINING_DATA_URL
  serving_data_url: *SERVING_DATA_URL
  scoring_data_url: *SCORING_DATA_URL
  model_url: "output/model/uc01"
  output_url: "output/output/uc01"
  scoring_output_url: "output/scoring/uc01"
2:
  # general
  name: "UseCase02.py"
  # engines
  training_engine: *ENGINE_DL2
  serving_engine: *ENGINE
  # data stores
  training_datastore: *HDFS # for storing the training data
  model_datastore: *HDFS # for storing the trained models
  serving_datastore: *HDFS # for storing the serving data
  output_datastore: *HDFS # for storing the final output
  # templates
  datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"
  training_template: "$engine $tpcxai_home/workload/spark/pyspark/workload-pyspark/$name --stage training --epochs 25 --batch 32 --executor_cores horovod 1 --task_cpus horovod 1 --workdir $output $input/$file $input/CONVERSATION_AUDIO.seq"
  serving_template: "$engine $tpcxai_home/workload/spark/pyspark/workload-pyspark/$name --stage serving --batch 32 --workdir $model --output $model/$phase $input/$file $input/CONVERSATION_AUDIO.seq"
  serving_throughput_template: "$engine $tpcxai_home/workload/spark/pyspark/workload-pyspark/$name --stage serving --batch 32 --workdir $model --output $model/$stream $input/$file $input/CONVERSATION_AUDIO.seq"
  # URLs

```

```

training_data_url: *TRAINING_DATA_URL
serving_data_url: *SERVING_DATA_URL
scoring_data_url: *SCORING_DATA_URL
model_url: "output/model/uc02"
output_url: "output/output/uc02"
scoring_output_url: "output/scoring/uc02"
working_dir: "/tmp"
3:
# general
name: "org.tpc.tpcxai.UseCase03"
# engines
training_engine: *ENGINE
serving_engine: *ENGINE
# data stores
training_datastore: *HDFS # for storing the training data
model_datastore: *HDFS # for storing the trained models
serving_datastore: *HDFS # for storing the serving data
output_datastore: *HDFS # for storing the final output
# templates
datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"
training_template: "$engine --class $name lib/workload-assembly-0.1.jar --stage training --workdir $output
$input/order.csv $input/lineitem.csv $input/product.csv"
serving_template: "$engine --class $name lib/workload-assembly-0.1.jar --stage serving --workdir $model --output
$model/$phase $input/store_dept.csv"
serving_throughput_template: "$engine --class $name lib/workload-assembly-0.1.jar --stage serving --workdir
$model --output $model/$stream $input/store_dept.csv"
# URLs
training_data_url: *TRAINING_DATA_URL
serving_data_url: *SERVING_DATA_URL
scoring_data_url: *SCORING_DATA_URL
model_url: "output/model/uc03"
output_url: "output/output/uc03"
scoring_output_url: "output/scoring/uc03"
4:
# general
name: "org.tpc.tpcxai.UseCase04"
# engines
training_engine: *ENGINE
serving_engine: *ENGINE
# data stores
training_datastore: *HDFS # for storing the training data
model_datastore: *HDFS # for storing the trained models
serving_datastore: *HDFS # for storing the serving data
output_datastore: *HDFS # for storing the final output
# templates
datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"
training_template: *TRAINING_TEMPLATE
serving_template: *SERVING_TEMPLATE
serving_throughput_template: *SERVING_THROUGHPUT_TEMPLATE
# URLs
training_data_url: *TRAINING_DATA_URL
serving_data_url: *SERVING_DATA_URL
scoring_data_url: *SCORING_DATA_URL
model_url: "output/model/uc04"
output_url: "output/output/uc04"
scoring_output_url: "output/scoring/uc04"
5:
# general
name: "UseCase05.py"
# engines

```

```

training_engine: *ENGINE_DL5
serving_engine: *ENGINE
# data stores
training_datastore: *HDFS # for storing the training data
model_datastore: *HDFS # for storing the trained models
serving_datastore: *HDFS # for storing the serving data
output_datastore: *HDFS # for storing the final output
# templates
datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"
# add namenode if necessary by specifying
# $engine [path]/$name --namenode [namenode.url:port]
training_template: &TRAINING_TEMPLATE_PY "$engine $tpcxai_home/workload/spark/pyspark/workload-
pyspark/$name --stage training --epochs 15 --batch 512 --workdir $output $input/$file"
serving_template: "$engine $tpcxai_home/workload/spark/pyspark/workload-pyspark/$name --stage serving --
batch 512 --workdir $model --output $model/$phase $input/$file"
serving_throughput_template: "$engine $tpcxai_home/workload/spark/pyspark/workload-pyspark/$name --stage
serving --batch 512 --workdir $model --output $model/$stream $input/$file"
# URLs
training_data_url: *TRAINING_DATA_URL
serving_data_url: *SERVING_DATA_URL
scoring_data_url: *SCORING_DATA_URL
model_url: "output/model/uc05"
output_url: "output/output/uc05"
scoring_output_url: "output/scoring/uc05"
working_dir: "/tmp"
6:
# general
name: "org.tpc.tpcxai.UseCase06"
# engines
training_engine: *ENGINE
serving_engine: *ENGINE
# data stores
training_datastore: *HDFS # for storing the training data
model_datastore: *HDFS # for storing the trained models
serving_datastore: *HDFS # for storing the serving data
output_datastore: *HDFS # for storing the final output
# templates
datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"
training_template: *TRAINING_TEMPLATE
serving_template: *SERVING_TEMPLATE
serving_throughput_template: *SERVING_THROUGHPUT_TEMPLATE
# URLs
training_data_url: *TRAINING_DATA_URL
serving_data_url: *SERVING_DATA_URL
scoring_data_url: *SCORING_DATA_URL
model_url: "output/model/uc06"
output_url: "output/output/uc06"
scoring_output_url: "output/scoring/uc06"
7:
# general
name: "org.tpc.tpcxai.UseCase07"
# engines
training_engine: *ENGINE
serving_engine: *ENGINE
# data stores
training_datastore: *HDFS # for storing the training data
model_datastore: *HDFS # for storing the trained models
serving_datastore: *HDFS # for storing the serving data
output_datastore: *HDFS # for storing the final output
# templates

```

```

    datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"
    training_template: "$engine --class $name lib/workload-assembly-0.1.jar --stage training --num-blocks 20 --workdir
$output $input/$file"
    serving_template: *SERVING_TEMPLATE
    serving_throughput_template: *SERVING_THROUGHPUT_TEMPLATE
    # URLs
    training_data_url: *TRAINING_DATA_URL
    serving_data_url: *SERVING_DATA_URL
    scoring_data_url: *SCORING_DATA_URL
    model_url: "output/model/uc07"
    output_url: "output/output/uc07"
    scoring_output_url: "output/scoring/uc07"
8:
    # general
    name: "org.tpc.tpcxai.UseCase08"
    # engines
    training_engine: *ENGINE
    serving_engine: *ENGINE
    # data stores
    training_datastore: *HDFS # for storing the training data
    model_datastore: *HDFS # for storing the trained models
    serving_datastore: *HDFS # for storing the serving data
    output_datastore: *HDFS # for storing the final output
    # templates
    datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"
    training_template: "$engine --class $name lib/workload-assembly-0.1.jar --stage training --num-workers 1 --num-
threads 1 --workdir $output $input/order.csv $input/lineitem.csv $input/product.csv"
    serving_template: "$engine --class $name lib/workload-assembly-0.1.jar --stage serving --num-workers 1 --num-
threads 1 --workdir $model --output $model/$phase $input/order.csv $input/lineitem.csv $input/product.csv"
    serving_throughput_template: "$engine --class $name lib/workload-assembly-0.1.jar --stage serving --num-workers
1 --num-threads 1 --workdir $model --output $model/$stream $input/order.csv $input/lineitem.csv $input/product.csv"
    # URLs
    training_data_url: *TRAINING_DATA_URL
    serving_data_url: *SERVING_DATA_URL
    scoring_data_url: *SCORING_DATA_URL
    model_url: "output/model/uc08"
    output_url: "output/output/uc08"
    scoring_output_url: "output/scoring/uc08"
9:
    # general
    name: "UseCase09.py"
    # engines
    training_engine: *ENGINE_DL2
    serving_engine: *SERVING_9
    # data stores
    training_datastore: *HDFS # for storing the training data
    model_datastore: *HDFS # for storing the trained models
    serving_datastore: *HDFS # for storing the serving data
    output_datastore: *HDFS # for storing the final output
    # templates
    datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"
    # add namenode if necessary by specifying
    # $engine [path]/$name --namenode [namenode.url:port]
    training_template: "$engine --files $tpcxai_home/workload/spark/pyspark/workload-
pyspark/resources/uc09/shape_predictor_5_face_landmarks.dat $tpcxai_home/workload/spark/pyspark/workload-
pyspark/$name --stage training --epochs_embedding=15 --batch=64 --executor_cores_horovod 1 --
task_cpus_horovod 1 --workdir $output $input/CUSTOMER_IMAGES_META.csv'
$input/CUSTOMER_IMAGES.seq"
    serving_template: "$engine --files $tpcxai_home/workload/spark/pyspark/workload-
pyspark/resources/uc09/shape_predictor_5_face_landmarks.dat $tpcxai_home/workload/spark/pyspark/workload-

```

```

pyspark/$name --stage serving --workdir $model --output $model/$phase
$input/CUSTOMER_IMAGES_META.csv' $input/CUSTOMER_IMAGES.seq'"
    serving_throughput_template: "$engine --files $tpcxai_home/workload/spark/pyspark/workload-
pyspark/resources/uc09/shape_predictor_5_face_landmarks.dat $tpcxai_home/workload/spark/pyspark/workload-
pyspark/$name --stage serving --workdir $model --output $model/$stream
$input/CUSTOMER_IMAGES_META.csv' $input/CUSTOMER_IMAGES.seq'"
    # URLs
    training_data_url: *TRAINING_DATA_URL
    serving_data_url: *SERVING_DATA_URL
    scoring_data_url: *SCORING_DATA_URL
    model_url: "output/model/uc09"
    output_url: "output/output/uc09"
    scoring_output_url: "output/scoring/uc09"
    working_dir: "/tmp"
10:
    # general
    name: "org.tpc.tpcxai.UseCase10"
    # engines
    training_engine: *ENGINE
    serving_engine: *ENGINE
    # data stores
    training_datastore: *HDFS # for storing the training data
    model_datastore: *HDFS # for storing the trained models
    serving_datastore: *HDFS # for storing the serving data
    output_datastore: *HDFS # for storing the final output
    # templates
    datagen_template: "java -jar $pdgf -ns -sf $scale_factor -s $table"
    training_template: "$engine --class $name lib/workload-assembly-0.1.jar --stage training --workdir $output
$input/financial_account.csv $input/financial_transactions.csv"
    serving_template: "$engine --class $name lib/workload-assembly-0.1.jar --stage serving --workdir $model --output
$model/$phase $input/financial_account.csv $input/financial_transactions.csv"
    serving_throughput_template: "$engine --class $name lib/workload-assembly-0.1.jar --stage serving --workdir
$model --output $model/$stream $input/financial_account.csv $input/financial_transactions.csv"
    # URLs
    training_data_url: *TRAINING_DATA_URL
    serving_data_url: *SERVING_DATA_URL
    scoring_data_url: *SCORING_DATA_URL
    model_url: "output/model/uc10"
    output_url: "output/output/uc10"
    scoring_output_url: "output/scoring/uc10"

```


Appendix E: Throughput Test Stream Placement

E.1 For convenience the following table displays the order of serving **use case** templates for the first 100 streams. The order is the same for all scale factors.

Stream	Use cases									
1	3	5	10	6	1	7	4	8	9	2
2	1	4	5	10	3	2	9	6	7	8
3	9	5	2	6	4	10	1	7	8	3
4	4	1	3	9	5	2	10	6	7	8
5	9	8	5	4	10	1	7	3	6	2
6	5	8	9	4	1	3	10	7	2	6
7	1	8	5	2	10	9	6	7	3	4
8	3	8	7	9	4	6	1	2	10	5
9	3	5	1	6	9	2	8	7	10	4
10	6	5	3	9	10	2	4	1	8	7
11	1	5	6	4	10	8	3	7	2	9
12	4	6	2	9	8	1	10	7	5	3
13	9	1	10	5	8	7	3	4	2	6
14	10	5	4	6	7	2	9	3	1	8
15	10	3	8	5	9	4	7	1	2	6
16	10	7	5	1	6	9	8	4	3	2
17	1	8	2	5	9	10	3	7	6	4
18	4	10	6	9	3	1	2	5	7	8
19	3	5	7	6	1	4	10	2	9	8
20	9	10	1	8	3	7	6	2	4	5
21	9	2	8	7	5	10	6	4	3	1
22	9	2	7	10	5	1	3	4	6	8
23	5	7	6	4	8	3	9	10	1	2
24	6	9	2	3	7	5	1	4	8	10
25	5	10	1	3	6	4	7	2	9	8
26	5	8	3	9	4	1	10	7	6	2
27	1	6	9	7	10	4	3	2	5	8
28	3	8	7	4	5	2	1	10	9	6
29	4	7	5	9	2	10	8	1	6	3
30	7	9	5	10	1	6	8	3	2	4
31	9	1	8	4	6	10	3	7	5	2
32	4	5	9	10	2	1	3	6	8	7
33	4	1	5	10	3	8	6	2	7	9
34	2	9	1	3	5	7	4	6	10	8
35	2	6	4	10	7	5	1	3	9	8
36	10	7	1	4	5	3	9	6	2	8
37	5	4	1	3	7	9	8	2	10	6
38	4	2	5	7	8	3	10	1	9	6
39	5	9	7	1	3	4	6	10	2	8
40	1	8	7	6	4	5	2	3	10	9
41	10	5	4	8	3	7	9	2	1	6
42	5	10	6	2	1	7	8	9	4	3
43	4	7	2	6	9	3	10	8	1	5
44	5	10	4	1	8	3	7	9	2	6
45	7	5	9	6	8	2	3	4	1	10
46	3	1	6	10	8	9	4	7	5	2
47	1	4	10	2	3	5	7	8	6	9
48	7	5	1	9	8	4	6	3	10	2
49	10	8	6	1	4	2	5	3	9	7
50	8	7	5	3	6	2	10	4	9	1

Stream	Use cases									
51	7	3	6	9	4	8	2	5	10	1
52	6	8	5	3	10	9	2	4	7	1
53	2	5	9	10	3	8	6	4	1	7
54	6	2	3	5	10	4	7	9	8	1
55	8	4	5	1	6	7	3	9	10	2
56	8	2	7	4	5	9	6	3	1	10
57	4	10	8	5	3	7	1	6	2	9
58	2	3	1	4	6	9	5	7	8	10
59	10	4	7	1	2	3	5	8	6	9
60	10	4	3	7	5	1	6	9	8	2
61	4	6	3	7	9	5	8	10	1	2
62	5	8	9	6	10	1	7	3	4	2
63	3	9	10	7	8	2	1	5	6	4
64	8	7	3	4	2	1	5	10	6	9
65	7	5	8	4	1	3	10	6	2	9
66	1	2	6	8	7	10	3	5	4	9
67	7	9	4	5	8	2	10	3	6	1
68	5	2	4	7	9	1	8	3	10	6
69	6	8	5	9	3	4	7	2	10	1
70	3	10	7	5	1	2	9	8	6	4
71	6	10	1	4	2	5	9	7	8	3
72	1	4	3	5	7	10	9	6	8	2
73	6	2	1	3	5	7	4	10	8	9
74	2	5	3	7	10	1	9	4	6	8
75	3	8	7	10	5	1	4	9	6	2
76	4	9	8	5	1	6	10	3	2	7
77	3	4	2	9	1	6	10	7	8	5
78	9	6	8	5	4	2	7	10	3	1
79	9	3	4	2	7	1	5	8	10	6
80	8	2	3	1	7	9	6	4	10	5
81	10	5	6	2	3	8	1	9	7	4
82	6	7	2	9	8	5	10	4	3	1
83	2	9	7	5	8	6	3	10	4	1
84	5	4	10	3	6	8	7	2	9	1
85	6	5	3	9	2	10	4	8	1	7
86	3	7	9	5	2	8	10	6	1	4
87	2	6	9	8	5	7	4	1	10	3
88	8	7	1	4	10	2	3	5	6	9
89	4	9	7	8	6	10	1	2	3	5
90	7	5	1	9	8	2	10	6	3	4
91	10	1	5	7	8	6	4	3	2	9
92	8	3	4	5	9	6	7	10	1	2
93	8	5	7	10	1	2	4	6	9	3
94	9	7	2	8	6	10	5	3	1	4
95	1	10	7	3	9	4	8	2	5	6
96	8	2	1	5	10	7	4	3	9	6
97	6	2	10	7	5	4	8	9	1	3
98	2	4	3	9	5	6	1	10	7	8
99	3	5	1	7	2	4	8	9	6	10
100	9	1	6	8	10	4	7	3	5	2

Appendix Table 10-a First 100 Streams

Appendix F: Minimum Software library versions for the SUT

- F.1 The list of TPC approved software libraries used by the SUT is listed below. The minimum version listed below is what the TPC has validated.
- F.2 For the scikit-learn implementation, the list of SUT software libraries approved by the TPC and may not be priced is below:

Library	Minimum Version	Software License
Scallop	3.1.3	MIT
Threeten	0.9	BSD 3
Xgboost4	1.0.0	Apache 2.0
Scikit-learn surprise	1.1	BSD 3
Librosa	0.8.1	ISC
Imbalanced-learn	0.9.0	MIT
Tensorflow-addons	0.15.0	Apache 2.0
OpenCV	4.5	Apache 2.0
Dlib	19.2	BSL 1.0
Tensorflow	2.1	Apache 2.0

Table 10-a TPCx-AI approved compute software list (scikit-learn)

- F.3 For the scikit-learn implementation, the list of SUT software that needs to be priced is below:

Library	Minimum Version	Software License
Python	3.7.12	PSFL
Setuptools	58	BSD 3
Pandas	1.2.4	BSD 3
Scikit-learn	1.0.2	BSD 3
Xgboost	1.5.0	Apache 2.0
Numpy	1.19.2	BSD 3
Nose	1.3.7	GPL 2+
Scipy	1.7.3	BSD 3
Statsmodels	0.12.2	BSD 3
Patsy	0.5.2	BSD 2

Tqdm	4.62	MPL V2.0
Keras	2.3.1	MIT
Joblib	1.1.0	BSD 3
Pyyaml	6	MIT
Matplotlib	3.5.0	BSD
Jinja2	3.0.2	BSD
Pycryptodome	3.12	BSD-2

Table 10-b: SUT Software to be priced

F.4 For the Spark implementation the list of SUT software libraries approved by the TPC and may not priced is below:

Library	Minimum Version	Software License
Scallop	3.1.3	MIT
Sparkts	0.4.1	Apache 2.0
Xgboost-spark	1.0.0	Apache 2.0
Xgboost4j	1.0.0	Apache 2.0
Horovod	0.19.1	Apache 2.0
Gcc linux64	9.3	GPL
Gxx_linux64	9.3	GPL
Openmpi mpicc	4	BSD 3
Numpy	1.19.2	BSD 3
Tensorflow	2.2	Apache 2.0
H5py	2.10.0	HDF5 Copyright
Tqdm	4.62	MPL V2.0
Joblib	1.1.0	BSD 3
Pyarrow	3	Apache 2.0
Librosa	0.8.1	ISC
Dlib	19.2	BSL 1.0
Petastorm	0.9.8	Apache 2.0
Tensorflow-addons	0.10.0	Apache 2.0
Pandas	1.2.3	BSD 3
Keras	2.4.3	MIT
Libsndfile-devel	1.0	LGPL
libsndfile	1.0	LGPL

Libglvnd-glx	1.0	LGPL2.1
cmake	3.12.1	BSD 3
pssh	2.3.1	GPL
Java	8	GPLV2

Table 10-c TPCx-AI approved compute software list (Spark)

F.5 For the spark implementation, the list of SUT software that needs to be priced is below:

Library	Minimum Version	Software License
Spark	2.4	Apache 2.0
Pyspark	2.4	Apache 2.0
CDP	7.1	ASLv2, AGPLv3

Table 10-d: SUT software to be priced (spark)