

Newsletter

Issue 1, April 2023

Welcome to the first TPC Newsletter <p>The Quarterly Newsletter is a new approach by the Public Relations Committee to keep everyone in the TPC abreast of the most important activities of the TPC. It summarizes activities in each subcommittee and in key standing committees from the past few months, and gives detailed reports on key decisions that were made.</p> <p>As this is the first edition of this Newsletter, it spans the activities of the year 2022. Future issues of the newsletters will be released on a quarterly basis. The next newsletter is planned to be released in June, 2023.</p> <p>By the editors of the newsletter</p>	TPCx-AI Takes The World By Storm <p>TPCx-AI is a new benchmark (2021) that aims to measure the performance of hardware and software systems used to execute machine learning (ML) and artificial intelligence (AI) workloads. The benchmarks specification was officially published September 2021 with revisions in 2022. Since the specification was made available for use there have been audited publications by 5 vendors spanning 7 different scale factors – up to 3TB in size. The most recent publication was made in March 2023.</p>
TPC-OSS Reaches 242,000 Downloads from TPC’s GitHub Repository <p>For several years now the TPC has been boosting its capability to develop benchmark specifications and to enhance existing benchmarks, by pursuing an Open Source Initiative (TPC-OSS). This effort leverages community driven development and testing resources. It also promotes the TPC to a new generation of engineers and developers that is used to accessing source code, and gets the TPC brand out in the developer community that is more associated with easy-to-use database workloads. At the same time it becomes more approachable to Academia, which currently views the TPC as rigid and closed off.</p> <p>Full story by Rodrigo Escobar (Intel)</p>	TPCx-IoT: TPC Completes Integration of Cloud-Native Multi-Model Database, Lindorm <p>TPCx-IoT is the first IoT benchmark specifically designed to measure the performance of IoT Gateway Systems. The original version of TPCx-IoT, published in 2017, included support for two databases, Hbase and Couchbase-server. Due to high demand for TPCx-IoT results, in 2019 the TPC added support for Machbase, a time-series database that stores and analyzes high volumes of sensor data from various facilities in real time. In 2022 the TPC added Lindorm, a cloud-native multi-model database, supporting time series model, wide-column table model and queue model.</p>
TPCx-BB Version 1.6.0 Released <p>TPCx-BB is a benchmark created by the TPC to evaluate the performance of systems when executing data analytics workloads on massive datasets. It leverages widely used distributed data processing frameworks to execute 30 Online Analytics Processing (OLAP) queries and machine learning tasks with different runtime profiles. The varied characteristics of the queries enable users to form a more holistic picture of the system’s performance when running data analytics tasks. Users can execute TPCx-BB on several different environments, including in-Cloud and on-premises deployments, to determine the best hardware and software offerings for their data science workloads.</p>	TPC-Pricing Releases Version 2.8.0 <p>The first version of the TPC-Pricing specification was approved in February 2005. It assembles a consistent set of pricing procedures and methodologies for all TPC benchmarks executed on both physically acquired hardware and licensed compute services (Cloud). Prior to February 2005, each TPC benchmark defined its own pricing rules, which had let to inconsistent pricing across different benchmarks. The introduction of the pricing specification was a quantum leap towards consistent and credible pricing.</p> <p>TPC pricing is constantly adapting to new trends in the industry. In February 2022 the TPC published Version 2.8.0, which clarified that only one version of the TPC-Pricing specification can be used where two TPC-Pricing versions are active.</p>

<p>Technology Conference TPCTC</p> <p>TPCTC, held annually for the last 14 years, has been recognized as the international event for anyone interested in performance related topics in database technology, including Transaction Processing, Data Warehousing, Big Data Analytics, Internet of Things, Virtualization, and Artificial Intelligence. Last year’s conference, held in Sydney, Australia featured multiple talks about Artificial Intelligence benchmarks. It also included two panels of highly distinguished academic and industry leaders.</p> <p>Full story by Meikel Poess (Oracle)</p>	<p>Public Relationship Updates</p> <p>Like many standards bodies the TPC aims to maintain a level playing field for its members who play a dual role of collaborators in the creation of benchmark standards – and competitors in the marketplace. The TPC Public Relations committee promotes the efforts of all our members as well as fairness and innovation in benchmarking. Please engage with us in our efforts on Twitter @tpcbenchmarks and LinkedIn. We welcome your feedback to help make benchmarking interesting and useful.</p>
<p>Steering Committee Updates</p> <p>The Steering Committee completed both governance and management tasks over the course of the year. A new budget was approved, and committee chair evaluations were conducted. The SC also addressed several Fair Use violations, notably with Databricks, Snowflake, and Bodo. All resulted in compliance. The committee also approved minor revisions to TPCx-AI and TPCx-IoT.</p> <p>With the increasing number of members from China, the SC approved the creation of a TPC Mandarin site. This site was designed in February and implemented in March of 2023. Several security measures were also taken to protect the TPC online assets. The SC is now addressing how to move forward with PDGF after learning that Bankmarc, the owner of the copyright, is intending to dissolve. Finally, a minor revision of the Policies was brought forward and approved.</p>	<p>TPC-C and TPC-E</p> <p>The TPC-C and TPC-E are Online Transaction Processing (OLTP) benchmarks. While the TPC-C benchmark simulates an order-entry environment, the TPC-E benchmark simulates a brokerage firm. The performance metric for TPC-C is the number of transactions per minute (tpmC) and that for TPC-E is the number of transactions per second. (tpsE). These benchmarks are very stable and there were no revisions to these benchmarks in 2022. There were two TPC-C submissions in 2022 by TTA. The latest TPC-E publication is by Lenovo in March 2023.</p> <p>By Nirmala Sundararajan (Dell)</p>
<p>TPC-DS</p> <p>TPC-DS is a benchmark to evaluate the performance of decision support systems (DSS). It is a complex, multi-dimensional benchmark that simulates real-world decision support scenarios, including ad hoc queries, data mining, and online analytical processing (OLAP). TPC-DS includes a large data set, complex query workloads, and a set of rules for measuring and reporting performance metrics, such as response time and throughput. The benchmark is widely used in industry to evaluate the performance of DSS products and to compare different hardware and software configurations. As TPC-DS is in maintenance mode, new revisions are rare. No revisions, nor benchmark publications were made in 2022.</p> <p>By Meikel Poess (Oracle)</p>	<p>TPC-H</p> <p>TPC-H is a widely used decision support benchmark that measures the performance of data warehouses. It consists of a set of 22 ad-hoc queries that measure a system's ability to process large volumes of data and how complex queries perform in single- and multi-user modes. Results of TPC-H runs are used to compare the performance of different systems and to make informed decisions when purchasing and configuring data warehousing solutions. As TPC-H is in maintenance mode, new revisions are rare. However, Revision 3.0.1 was released on 4/28/ 2022, which contains minor changes, such as clarifications on change log history and comments to Clause 9.2.4.3 and Clause 2.4.19.5. In 2022 5 TPC-H benchmark results were published.</p> <p>By Meikel Poess (Oracle)</p>

TPC-OSS Reaches 242,000 Downloads from TPC's GitHub Repository

In recent years, many of TPC's benchmarks have been the foundation of open-source benchmarks. Database companies have been using these open-source benchmarks to prove their performance marketing claims. The extent to which these "TPC derived" benchmarks resemble their original version differs from a clone to running only a small subset of the original workload.

As long as TPC policies are not violated the TPC allows the use of TPC derived benchmarks. At the center of the TPC policies, which govern the use of TPC derived benchmarks, is the following, "It is permissible to create a non-TPC Benchmark Standard. Any non-TPC Benchmark Standard must adhere to the following requirements: ... The Use of TPC Material (Policies § 8.1) and Fair Use Policy (Policies § 8.2.1) must not be violated. ... All deviations from the TPC Benchmark Standard in question must be explicitly noted. ... Results based on a non-TPC Benchmark must be clearly identified as not being comparable to an official TPC Benchmark Result."

To partake in the development of open-source benchmarks, the TPC started its own open-source Initiative, TPC-OSS. It takes advantage of community driven development, testing and benchmark distribution. It further promotes TPC to a new generation of engineers and developers that expect access to source code. It also promotes the TPC brand in developer communities that are used to easy-to-use database workloads.

The first project of TPC-OSS was to host [HammerDB](#) in TPC's own [GitHub repository](#) starting in May 2019. HammerDB is a database benchmarking tool that is designed to stress test and evaluate the performance of various databases. It supports a wide range of databases, including Oracle, SQL Server, MySQL, and PostgreSQL, among others. HammerDB allows users to run multiple workloads concurrently, making it ideal for testing the scalability of a database under varying workloads. The tool is also highly customizable, with support for custom scripts and workloads.

HammerDB provides a user-friendly interface that simplifies the process of creating and running benchmark tests. The tool provides detailed reports that allow users to evaluate database performance and identify bottlenecks. These reports can be exported to various formats, including HTML and CSV, for further analysis. Additionally, HammerDB supports a range of performance metrics, including throughput, response time, and CPU utilization, among others.

Overall, HammerDB is a powerful and versatile tool that enables users to evaluate the performance of different databases under a range of workloads. Its user-friendly interface, customizability, and support for a wide range of performance metrics make it an ideal choice for anyone looking to benchmark and optimize their database performance.


The following list summarizes the accomplishment of TPC-OSS 2022.

1. HammerDB 4.4 was released March 3, 2022 with added support for Citus database

- a. HammerDB 4.6 was released December 1, 2022 with added GUI Icon and CLI command to drop the schema created by Build Schema and enhanced CLI to accept Python run scripts with .py extensions


The following graph shows the download activity of HammerDB over time and for all releases.

HammerDB Downloads Over Time



The following graph shows the download activity of HammerDB over time by release.

Downloads by Release over time



TPC Technology Conference - TPCTC



TPCTC is TPC's annual technology conference. Its mission is to bring together industry experts and researchers to explore new methodologies for measuring the performance of data-centric applications. Over the last 14 years TPCTC has been recognized as the international event for anyone interested in performance related topics in database technology, including Transaction Processing, Data Warehousing, Big Data Analytics, Internet of Things, Virtualization, and Artificial Intelligence.

Performance evaluation has and still is one of the main differentiators of computer systems. Constant hardware and software improvements require a fresh look at performance methodologies that allow performance evaluation in a technically sound, fair and meaningful ways. The TPC has always been at the forefront of these benchmark developments. While the TPC has focused historically on database centric benchmarks, recent developments include benchmarks for Artificial Intelligence, Internet-Of-Things, Hyper-Converged Infrastructure, Big Data and Virtualization. Many of these benchmarks were sparked by ideas that originated in papers presented at past TPCTC events. These papers were mostly academic papers, that inspired ideas for new benchmarks, identified deficiencies in existing benchmarks and motivated improvements.

At the same time the academic community and industry have benefited from TPCTC to define their own performance methodologies. The last ten years have seen the rise of many new DBMSs, some with unique approaches to traditional problems, like columnar databases and in-memory databases, others specializing on specific applications, such as graph databases, NoSQL databases, Timeseries databases etc. With them came a flurry of new performance methodologies resulting in customized benchmarks, many of which were based on methodologies originally developed in the TPC.

TPCTC has always served as a venue where both practitioners with real world performance expertise met with highly innovative academics to discuss performance methodologies for emerging technologies. The results, both in academia and TPC benchmarks, are a direct result of this knowledge exchange.

TPCTC is a peer reviewed conference. All papers were reviewed by at least 3 reviewers from the Program Committee, which consisted of: Dippy Aggarwal (Intel, USA), Daniel Bowers (Gartner, USA), Michael Brey (Oracle, USA), Ajay Dholakia (Lenovo, USA), Dhabaleswar Panda (The Ohio State University, USA), Tilmann Rabl (TU Berlin, Germany), Swapna Raj (Intel, USA), Anil Rajput (AMD, USA) and Reza Taheri (VMWare, USA). The following table shows the number of submitted papers versus the number of accepted papers in the last four years.

Year	Number of submitted papers	Number of accepted papers	Acceptance rate
2022	12	5	41%
2021	18	8	44%
2020	11	5	45%
2019	22	9	40%

TPCTC 2022 was held in Sydney, Australia in conjunction with the conference on Very Large Databases (VLDB), a top tier academic conference for databases. TPCTC 2022 marked TPCTC's 14th anniversary. Due to COVID, TPCTC was held in hybrid mode. That is, attendants could attend in person or over zoom.



Raghu Nambiar opened the conference by emphasizing how important TPC standards are for the industry and research and the contributions of the TPCTC in bringing new ideas and methodologies in performance evaluation and benchmarking. TPCTC was a great success featuring multiple talks about Artificial Intelligence benchmarks, “More the Merrier: Comparative evaluation of TPCx-AI and MLPerf Benchmarks for AI: by Ajay Dholakia, Miro Hodak, Yingrui Liu Olesiuk and David Ellison; “Preliminary Scaling Characterization with TPCx-AI” by Rodrigo D. Escobar Palacios, Sammy Nah, Kacper Ufa,

Amandeep Raina and Hamesh Patel; and “TPCx-AI: First Adopter's Experience Report” by Hyo-Sil Kim, Doohwan Kim, Byoungjun Seo and Sejin Hwang.

We offered two interesting panels with highly distinguished academic and industry leaders. The topic of the first panel, which was moderated by Shahram Ghandeharizadeh, was “Disaggregated Database Management Systems”. We invited industry and academic leaders who are working on Disaggregated Database Management Systems: Shahram Ghandeharizadeh (Moderator), Phil Bernstein, Dhruba Borthakur, Haoyu Huang, Jai Menon and Sumit Puri.



Shahram Ghandeharizadeh
(Moderator)
Associate Professor of Computer Science at University of Southern California (USC)



Phil Bernstein
Computer scientist specializing in database research in the Database Group of Microsoft Research



Dhruba Borthakur
Co-founder and CTO at Rockset



Haoyu Huang
Software Engineer in the database group at Google, before Meta (Facebook)



Jai Menon
Chief Scientist at Fungible, Inc.



Sumit Puri
CEO & Co-Founder at Liqid



Ajay Dholakia
(Moderator)
Principal Engineer, Sr Solution Architect, Chief Technologist for SW & Solution Dev at Lenovo

The second panel was entitled “Benchmarking considerations for Trustworthy and Responsible AI”. It was moderated by Ajay Dholakia. Participants were David Ellison, Miro Hodak, Debo Dutta, Raghu Nambiar, Rajesh Bordawekar

Below place find the full schedule for TPCTC2022.

September 5th, 2022		
All times are local times for Sydney, Australia		
Start Time	End Time	Paper Information
09:00 AM	09:30 AM	Opening Remarks. Industry Standards - A Look Back & A Look Ahead: Raghunath Nambiar.
09:30 AM	10:00 AM	Pick and Mix Isolation Levels: Mixed Serialization Graph Testing: Jack Waudby, Paul Ezhilchelvan and Jim Webber.
10:00 AM	10:30 AM	BODS: A Benchmark on Data Sortedness: Aneesh Raman, Kostas Karatsenidis, Subhadeep Sarkar, Matthaïos Olma and Manos Athanassoulis.
10:30 AM	11:00 AM	Coffee Break
11:00 AM	12:00 PM Noon	Disaggregated Database Management Systems (Panel): Shahram Ghandeharizadeh (Moderator), Phil Bernstein, Dhruva Borthakur, Haoyu Huang, Jai Menon, Sumit Puri.
12:00 PM Noon	12:30 PM	TPCx-AI on NVIDIA Jetson: Robert Bayer, Jon Voigt Tøttrup and Pinar Tozun.
12:30 PM	2:00 PM	Lunch Break
2:00 PM	2:30 PM	More the Merrier: Comparative evaluation of TPCx-AI and MLPerf Benchmarks for AI: Ajay Dholakia, Miro Hodak, Yingrui Liu Olesiuk, David Ellison.
2:30 PM	3:00 PM	Preliminary Scaling Characterization with TPCx-AI. Rodrigo D. Escobar Palacios, Sammy Nah, Kacper Ufa, Amandeep Raina, Hamesh Patel.
3:00 PM	3:30 PM	4mbench: Performance Benchmark of Manufacturing Business Database: Kazuo Goda, Yuto Hayamizu, Norifumi Nishikawa and Shinji Fujiwara.
3:30 PM	4:00 PM	Coffe Break
4:00 PM	5:00 PM	Benchmarking considerations for Trustworthy and Responsible AI (Panel). Ajay Dholakia (Moderator), David Ellison, Miro Hodak, Debo Dutta, Raghu Nambiar, Rajesh Bordawekar.
5:00 PM	5:30 PM	TPCx-AI: First Adopter's Experience Report: Hyo-Sil Kim, Doohwan Kim, Byoungjun Seo, Sejin Hwang.
5:30 PM	6:00 PM	New Initiatives in the TPC: Meikel Poess.
6:00 PM	6:30 PM	Closing Remarks Meikel Poess.

We are excited to announce that the VLDB committee has again accepted our proposal for TPCTC to be included in this year's VLDB conference. TPCTC will be held at the end of August 2023 in Toronto, Canada. Stay tuned for the exact date.

