

Industry Standards for the Analytics Era: TPC Roadmap

Raghunath Nambiar^{1(✉)} and Meikel Poess²

¹ Cisco Systems, Inc., 275 East Tasman Drive, San Jose, CA 95134, USA
rnambiar@cisco.com

² Oracle Corporation, 500 Oracle Parkway, Redwood Shores, CA 94065, USA
meikel.poess@oracle.com

Abstract. The Transaction Processing Performance Council (TPC) is a non-profit organization focused on developing data-centric benchmark standards and disseminating objective, verifiable performance data to industry. This paper provides a high-level summary of TPC benchmark standards, technology conference initiative, and new development activities in progress.

Keywords: Industry standards · Database benchmarks

1 TPC Benchmark Timelines

Founded in 1988, the Transaction Processing Performance Council (TPC) is a non-profit corporation dedicated to creating and maintaining benchmarks which measure database performance in a standardized, objective and verifiable manner. As of November 2017, 21 full members and three associate members comprise the TPC.

To date the TPC has approved a total of sixteen different benchmarks. Of these benchmarks, twelve are currently active. TPC currently defines two benchmark classes: Enterprise and Express. See Fig. 1 for the benchmark timelines.

- Enterprise benchmarks are technology agnostic. They are specification based, typically complex, and have long development cycles. Their specifications are provided by the TPC, but their implementation is up to the vendor. The vendor may choose any commercially available combination of software and hardware products to implement benchmarks. Examples of enterprise benchmarks are: TPC-C, TPC-E, TPC-H, TPC-DS, TPC-DI, TPC-VMS
- Express benchmarks are kit based, typically based on exiting workloads have shorter development cycles. It is required to use TPC provided kits for the publication of express benchmarks. Examples of express benchmarks: TPCx-HS, TPCx-BB, TPCx-V, TPCx-IoT

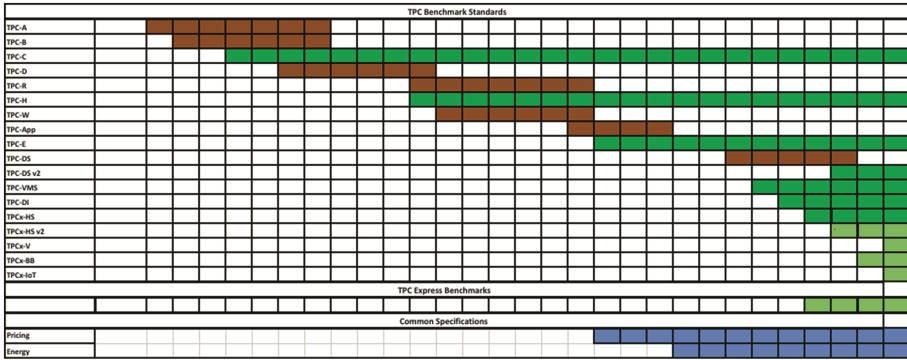


Fig. 1. TPC benchmark timelines

A high-level summary of current active standards are listed below:

1.1 Transaction Processing

TPC-C: Approved in July of 1992, TPC Benchmark C is an on-line transaction processing (OLTP) benchmark. TPC-C is more complex than previous OLTP benchmarks such as TPC-A because of its multiple transaction types, more complex database and overall execution structure. TPC-C involves a mix of five concurrent transactions of different types and complexity either executed on-line or queued for deferred execution. The database is comprised of nine types of tables with a wide range of record and population sizes. TPC-C is measured in transactions per minute (tpmC). While the benchmark portrays the activity of a wholesale supplier, TPC-C is not limited to the activity of any particular business segment, but, rather represents any industry that must manage, sell, or distribute a product or service.

TPC-E: Approved in February of 2007, TPC Benchmark E is an on-line transaction processing (OLTP) benchmark. TPC-E is more complex than previous OLTP benchmarks such as TPC-C because of its diverse transaction types, more complex database and overall execution structure. TPC-E involves a mix of twelve concurrent transactions of different types and complexity, either executed on-line or triggered by price or time criteria. The database is comprised of thirty-three tables with a wide range of columns, cardinality, and scaling properties. TPC-E is measured in transactions per second (tpsE). While the benchmark portrays the activity of a stock brokerage firm, TPC-E is not limited to the activity of any particular business segment, but rather represents any industry that must report upon and execute transactions of a financial nature.

1.2 Decision Support

TPC-H: The TPC BenchmarkTMH (TPC-H) is a decision support benchmark. It consists of a suite of business oriented ad-hoc queries and concurrent data modifications. The queries and the data populating the database have been chosen to have broad

industry-wide relevance. This benchmark illustrates decision support systems that examine large volumes of data, execute queries with a high degree of complexity, and give answers to critical business questions. The performance metric reported by TPC-H is called the TPC-H Composite Query-per-Hour Performance Metric (QphH@Size), and reflects multiple aspects of the capability of the system to process queries. These aspects include the selected database size against which the queries are executed, the query processing power when queries are submitted by a single stream, and the query throughput when queries are submitted by multiple concurrent users. The TPC-H Price/Performance metric is expressed as $\$/\text{QphH@Size}$.

TPC-DS: The TPC Benchmark DS (TPC-DS) is a decision support benchmark that models several generally applicable aspects of a decision support system, including queries and data maintenance. The benchmark provides a representative evaluation of performance as a general purpose decision support system. A benchmark result measures query response time in single user mode, query throughput in multi user mode and data maintenance performance for a given hardware, operating system, and data processing system configuration under a controlled, complex, multi-user decision support workload. The purpose of TPC benchmarks is to provide relevant, objective performance data to industry users. TPC-DS Version 2 enables emerging technologies, such as Big Data systems, to execute the benchmark [3, 4].

TPC-DI: Historically, the process of synchronizing a decision support system with data from operational systems has been referred to as Extract, Transform, Load (ETL) and the tools supporting such process have been referred to as ETL tools. Recently, ETL was replaced by the more comprehensive acronym, data integration (DI). DI describes the process of extracting and combining data from a variety of data source formats, transforming that data into a unified data model representation and loading it into a data store. The TPC-DI benchmark combines and transforms data extracted from an On-Line Transaction Processing (OLTP) system along with other sources of data, and loads it into a data warehouse. The source and destination data models, data transformations and implementation rules have been designed to be broadly representative of modern data integration requirements [5].

1.3 Big Data and Analytics

TPCx-HS v1: Big Data technologies like Hadoop has become an important part of the enterprise IT ecosystem. Introduced in 2014, the TPC Express Benchmark HS (TPCx-HS) Version 1 is industry's first ever standard for benchmarking big data systems. It was developed to provide an objective measure of hardware, operating system and commercial Apache Hadoop File System API compatible software distributions, and to provide the industry with verifiable performance, price-performance and availability metrics. Even though the modeled application is simple, the results are highly relevant to hardware and software dealing with Big Data systems in general. TPCx-HS stresses both the hardware and software stacks including the execution engine (MapReduce or Spark) and Hadoop Filesystem API compatible layers. This workload can be used to

assess a broad range of system topologies and implementation of Hadoop clusters. The TPCx-HS benchmark can be used to assess a broad range of system topologies and implementation methodologies in a technically rigorous and directly comparable, in a vendor-neutral manner [6].

TPCx-HS v2: The Hadoop ecosystem is moving fast beyond batch processing with MapReduce. Introduced in 2016 TPCx-HS V2 is based on TPCx-HS V1 with support for Apache Spark - a popular platform for in-memory data processing that enables real-time analytics on Apache Hadoop. TPCx-HS V2 also supports MapReduce (MR2) and supports publications on traditional on premise deployments and clouds. More information about TPCx-HS v1 can be found at <http://www.tpc.org/tpcx-hs/default.asp?version=1>. The TPCx-HS v2 benchmark can be used to assess a broad range of system topologies and implementation methodologies in a technically rigorous and directly comparable, in a vendor-neutral manner.

TPCx-BB: TPCx-BB Express Benchmark BB (TPCx-BB) measures the performance of Hadoop-based Big Data systems. It measures the performance of both hardware and software components by executing 30 frequently performed analytical queries in the context of retailers with physical and online store presence. The queries are expressed in SQL for structured data and in machine learning algorithms for semi-structured and unstructured data. The SQL queries can use Hive or Spark, while the machine learning algorithms use machine learning libraries, user defined functions, and procedural programs [7].

1.4 Virtualization

TPC-VMS: Introduced in 2012, the TPC Virtual Measurement Single System Specification (TPC-VMS) leverages the TPC-C, TPC-E, TPC-H and TPC-DS Benchmarks by adding the methodology and requirements for running and reporting performance metrics for virtualized databases. The intent of TPC-VMS is to represent a Virtualization Environment where three database workloads are consolidated onto one server. Test sponsors choose one of the four benchmark workloads (TPC-C, TPC-E, TPC-H, or TPC-DS) and runs one instance of that benchmark workload in each of the 3 virtual machines (VMs) on the system under test. The 3 virtualized databases must have the same attributes, e.g. the same number of TPC-C warehouses, the same number of TPC-E Load Units, or the same TPC-DS or TPC-H scale factors. The TPC-VMS Primary Performance Metric is the minimum value of the three TPC Benchmark Primary metrics for the TPC Benchmarks run in the Virtualization Environment [8].

TPCx-V: The TPC Express Benchmark V (TPCx-V) benchmark measures the performance of a virtualized server platform under a demanding database workload. It stresses CPU and memory hardware, storage, networking, hypervisor, and the guest operating system. TPCx-V workload is database-centric and models many properties of cloud services, such as multiple VMs running at different load demand levels, and large fluctuations in the load level of each VM. Unlike previous TPC benchmarks, TPCx-V

has a publicly-available, end-to-end benchmarking kit, which was developed specifically for this benchmark. It loads the databases, runs the benchmark, validates the results, and even performs many of the routine audit steps. Another unique characteristic of TPCx-V is an elastic workload that varies the load delivered to each of the VMs by as much as 16x, while maintaining a constant load at the host level [8].

1.5 Internet of Things (IoT)

TPCx-IoT: TPCx-IoT is the industry's first benchmark which enables direct comparison of different software and hardware solutions for IoT gateways. Positioned between edge architecture and the back-end data center, gateway systems perform functions such as data aggregation, real-time analytics and persistent storage. TPCx-IoT was specifically designed to provide verifiable performance, price-performance and availability metrics for commercially available systems that typically ingest massive amounts of data from large numbers of devices, while running real-time analytic queries. The workload is representative of activities typical in IoT gateway systems, running on commercially available hardware and software platforms. The TPCx-IoT can be used to assess a broad range of system topologies and implementation methodologies in a technically rigorous and directly comparable, in a vendor-neutral manner.

2 TPCTC Conference Series

To keep pace with rapid changes in technology, in 2009, the TPC initiated a conference series on performance analysis and benchmarking. The TPCTC has been challenging industry experts and researchers to develop innovative techniques for performance evaluation, measurement, and characterization of hardware and software systems. Over the years it has emerged as a leading forum to present and debate the latest and greatest in the world of benchmarking. The topics of interest included:

- Big data and analytics
- Complex event processing
- Database Optimizations
- Data Integration
- Disaster tolerance and recovery
- Emerging storage technologies (NVMe, 3D XPoint Memory etc.)
- Hybrid workloads
- Energy and space efficiency
- In-memory databases
- Internet of Things
- Virtualization
- Enhancements to TPC workloads
- Lessons learned in practice using TPC workloads
- Collection and interpretation of performance data in public cloud environments

2.1 Summary of the TPCTC Conferences Are Listed Below

The first TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2009), held in conjunction with the 35th International Conference on Very Large Data Bases (VLDB 2009) in Lyon, France from August 24th to August 28th, 2009 [9].

The second TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2010) was held in conjunction with the 36th International Conference on Very Large Data Bases (VLDB 2010) in Singapore from September 13th to September 17th, 2010 [10].

The third TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2011), held in conjunction with the 37th International Conference on Very Large Data Bases (VLDB 2011) in Seattle, Washington from August 29th to September 3rd, 2011 [11].

The fourth TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2012), held in conjunction with the 38th International Conference on Very Large Data Bases (VLDB 2012) in Istanbul, Turkey from August 27th to August 31st, 2012 [12].

The fifth TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2013), held in conjunction with the 39th International Conference on Very Large Data Bases (VLDB 2013) in Riva del Garda, Trento, Italy from August 26th to August 30st, 2013 [13].

The sixth TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2014), held in conjunction with the 40th International Conference on Very Large Data Bases (VLDB 2014) in Hangzhou, China, from September 1st to September 5th, 2014 [14].

The seventh TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2015), held in conjunction with the 41st International Conference on Very Large Data Bases (VLDB 2015) in Kohala Coast, USA, from August 31st to September 4th, 2015 [15].

The eighth TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2016), held in conjunction with the 42nd International Conference on Very Large Data Bases (VLDB 2016) in New Delhi, India, from September 5th to September 9th, 2016.

The ninth TPC Technology Conference on Performance Evaluation and Benchmarking (TPCTC 2017), held in conjunction with the 43nd International Conference on Very Large Data Bases (VLDB 2017) in Munich, India, from August 28th to September 1th, 2017.

TPCTC has had a significant positive impact on the TPC. TPC is able to attract new members from industry and academia to join the TPC. The formation of working groups on Big Data, Virtualization, Hyper-convergence, Internet of Things (IoT) and Artificial Intelligence were a direct result of TPCTC conferences.

3 Outlook

TPC remains committed to develop relevant standards in collaboration with industry and research communities and continue to enable fair comparison of technologies and products in terms of performance, cost of ownership.

Foreseeing the industry transition to digital transformation the TPC has created a working group to develop set of standards for hardware and software pertaining to Artificial Intelligence. Companies, research and government institutions who are interested in influencing the development of such benchmarks are encouraged to join the TPC [2].

Acknowledgements. Developing benchmark standards require a huge effort to conceptualize, research, specify, review, prototype, and verify the benchmark. The authors acknowledge the work and contributions of past and present members of the TPC.

References

1. Nambiar, R., Poess, M.: Reinventing the TPC: from traditional to big data to Internet of Things. In: Nambiar, R., Poess, M. (eds.) TPCTC 2015. LNCS, vol. 9508, pp. 1–7. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-31409-9_1
2. Nambiar, R., Poess, M.: Keeping the TPC relevant! PVLDB 6(11), 1186–1187 (2013)
3. Nambiar, R., Wakou, N., Masland, A., Thawley, P., Lanken, M., Carman, F., Majdalany, M.: Shaping the landscape of industry standard benchmarks: contributions of the Transaction Processing Performance Council (TPC). In: Nambiar, R., Poess, M. (eds.) TPCTC 2011. LNCS, vol. 7144, pp. 1–9. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32627-1_1
4. Nambiar, R., Poess, M.: The making of TPC-DS. In: VLDB 2006, pp. 1049–1058
5. Poess, M., Nambiar, R., Walrath, D.: Why you should run TPC-DS: a workload analysis. In: VLDB 2007, pp. 1138–1149
6. Poess, M., Rabl, T., Caufield, B.: TPC-DI: the first industry benchmark for data integration. PVLDB 7(13), 1367–1378 (2014)
7. Nambiar, R., Poess, M., Cao, P., Magdon-Ismail, T., Ren, D.Q., Bond, A.: Introducing TPCx-HS: the first industry standard for benchmarking big data systems. In: Nambiar, R., Poess, M. (eds.) TPCTC 2014. LNCS, vol. 8904, pp. 1–12. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-15350-6_1
8. Baru, C., et al.: Discussion of BigBench: a proposed industry standard performance benchmark for big data. In: Nambiar, R., Poess, M. (eds.) TPCTC 2014. LNCS, vol. 8904, pp. 44–63. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-15350-6_4
9. Bond, A., Johnson, D., Kopczyński, G., Taheri, H.R.: Profiling the performance of virtualized databases with the TPCx-V benchmark. In: Nambiar, R., Poess, M. (eds.) TPCTC 2015. LNCS, vol. 9508, pp. 156–172. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-31409-9_10
10. Nambiar, R., Poess, M. (eds.): Performance Evaluation and Benchmarking. LNCS, vol. 5895. Springer, Heidelberg (2009). <https://doi.org/10.1007/978-3-642-10424-4>
11. Nambiar, R., Poess, M. (eds.): Performance Evaluation, Measurement and Characterization of Complex Systems. LNCS, vol. 6417. Springer, Heidelberg (2011). <https://doi.org/10.1007/978-3-642-18206-8>
12. Nambiar, R., Poess, M. (eds.): Topics in Performance Evaluation, Measurement and Characterization. LNCS, vol. 7144. Springer, Heidelberg (2012). <https://doi.org/10.1007/978-3-642-32627-1>

13. Nambiar, R., Poess, M. (eds.): Selected Topics in Performance Evaluation and Benchmarking. LNCS, vol. 7755. Springer, Heidelberg (2013). <https://doi.org/10.1007/978-3-642-36727-4>
14. Nambiar, R., Poess, M. (eds.): Performance Characterization and Benchmarking. LNCS, vol. 8391. Springer, Cham (2014). <https://doi.org/10.1007/978-3-319-04936-6>
15. Nambiar, R., Poess, M. (eds.): Performance Characterization and Benchmarking. Traditional to Big Data. LNCS, vol. 8904. Springer, Cham (2015). <https://doi.org/10.1007/978-3-319-15350-6>
16. Nambiar, R., Poess, M. (eds.): Performance Evaluation and Benchmarking: Traditional to Big Data to Internet of Things. LNCS, vol. 9508. Springer, Cham (2016). <https://doi.org/10.1007/978-3-319-31409-9>