

Keeping the TPC Relevant!

Raghunath Nambiar
Cisco Systems, Inc.
275 East Tasman Drive
San Jose, CA 95134, USA
RNambiar@cisco.com

Meikel Poess
Oracle Corporation
500 Oracle Pkwy
Redwood Shores, CA 94065, USA
Meikel.Poess@Oracle.com

ABSTRACT

The Transaction Processing Performance Council (TPC) is a nonprofit organization founded in 1988 to define transaction processing and database benchmarks. Since then, the TPC has played a crucial role in providing the industry with relevant standards for total system performance, price-performance, and energy-efficiency comparisons. TPC benchmarks are widely used by database researchers and academia. Historically known for database-centric standards, the TPC has developed a benchmark for virtualization and is currently developing a multisource data integration benchmark. The technology landscape is changing at a rapid pace, challenging industry experts and researchers to develop innovative techniques for evaluating, measuring, and characterizing the performance of modern application systems. The Technology Conference series on Performance Evaluation and Benchmarking (TPCTC), introduced in 2009, and the new TPC-Express initiatives are steps taken by the TPC to be relevant in the coming years and beyond.

1. INDUSTRY-STANDARD BENCHMARKS

Industry-standard benchmarks have played, and continue to play, a crucial role in the advancement of the computing industry. Without them we would regress to the situation of the mid-1980s, when the lack of standard database benchmarks led many system vendors to practice what is now referred to as “benchmarking”—performance claims based on self-designed, highly biased benchmarks. The goal of publishing results from such tailored benchmarks was to make marketing claims, regardless of the relevant and verifiable technical merits. In essence, these benchmarks were designed as forgone conclusions to fit a preestablished marketing message. Similarly, vendors would create configurations, referred to as “benchmark specials,” that were designed to maximize performance against a specific benchmark, with limited relevance to the real world.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Articles from this volume were invited to present their results at The 39th International Conference on Very Large Data Bases, August 26th - 30th 2013, Riva del Garda, Trento, Italy.
Proceedings of the VLDB Endowment, Vol. 6, No. 11
Copyright 2013 VLDB Endowment 2150-8097/13/09... \$ 10.00.

2. STATE OF TPC

Over the last quarter century, the TPC has developed an array of benchmark suites [1]: TPC-A and TPC-B are designed to measure the performance of transactional systems which later morphed into TPC-C. TPC also developed TPC-E, another transaction processing benchmark. TPC-D was the first decision support benchmark later split into TPC-H and TPC-R. TPC-DS is another decision support benchmark based on complex schema and queries representing modern decision support systems. TPC-VMS is the industry’s first benchmark for databases in virtualization environments. TPC-Energy and TPC-Pricing are common standards across all benchmark standards. These standards have had a significant impact on the computing industry’s use of industry-standard benchmarks. Vendors use TPC benchmarks to illustrate performance competitiveness for their existing products, and to improve and monitor the performance of their products under development. Many buyers use TPC benchmark results as points of comparison when purchasing new computing systems. TPC benchmarks are widely used by database research community; there have been hundreds of publications that referenced TPC benchmarks.

While TPC’s contributions have been impressive, they are now challenged by rapid changes in the industry landscape. Longer development cycles are not acceptable. The cloud and big data analytics have become important parts of the enterprise application portfolio.

TPC is also experiencing a diminishing number of benchmark publications. There are various reasons for this trend. Foremost, the traditional database and system vendors have consolidated into fewer companies, reducing TPC’s membership. Second, some benchmarks have been manipulated by publishing a very high score using unrealistic configurations. Third, TPC benchmarks have become extremely complicated to develop and run. While TPC’s first benchmarks consisted of a few pages of definition, today’s benchmarks consist of hundreds of pages. Contrary to other standard consortia’s benchmarks, the TPC provides benchmark specifications and not the implementation in terms of executables and tools.

Consequently, TPC benchmarks are technology agnostic but also difficult to run, as a benchmark sponsor needs to first develop and implement the benchmark using its own technologies and tools.

3. BENCHMARKING FOR 2013 AND BEYOND

TPC has been closely watching the industry and technology changes and continues to remain committed to serving the industry and research community by developing new standards and taking appropriate steps to reach these goals. TPC's Technology Conference on Performance Evaluation and Benchmarking (TPCTC) initiative, founded in 2009, [2][3][4][5] seeks to foster collaboration between industry experts and the research community to explore new ideas and methodologies in performance evaluation and benchmarking. TPC's virtualization initiative and data integration initiatives are the direct result of past TPCTC conferences.

To address the increased benchmark development cycle and development costs, the TPC is working on a new initiative called the "express benchmark" initiative, possibly using a meaningful subset of the complete workload while retaining the strengths of the existing "enterprise benchmarks" model.

Currently, there are two benchmark developments in progress, a new virtualization benchmark and data integration benchmark, both being developed in the existing enterprise model. As big data and analytics become mainstream in the enterprises, the TPC is expected to explore benchmark development to measure the effectiveness of hardware and software systems dealing with them, likely using the express model.

4. ACKNOWLEDGMENTS

Our thanks to past and current members of the TPC.

5. REFERENCES

- [1] TPC website: www.tpc.org.
- [2] Raghunath Nambiar, Meikel Poess (eds.): Selected Topics in Performance Evaluation and Benchmarking. 4th TPC Technology Conference, TPCTC 2012, Istanbul, Turkey, August 27, 2012, Revised Selected Papers. Lecture Notes in Computer Science 7755, Springer 2013. ISBN 978-3-642-36726-7.
- [3] Raghunath Nambiar, Meikel Poess (eds.): Topics in Performance Evaluation, Measurement and Characterization. Third TPC Technology Conference, TPCTC 2011, Seattle, WA, USA, August 29 to September 3, 2011, Revised Selected Papers. Lecture Notes in Computer Science 7144, Springer 2012. ISBN 978-3-642-32626-4.
- [4] Raghunath Nambiar, Meikel Poess (eds.): Performance Evaluation, Measurement and Characterization of Complex Systems. Second TPC Technology Conference, TPCTC 2010, Singapore, September 13 to 17, 2010, Revised Selected Papers. Lecture Notes in Computer Science 6417, Springer 2011. ISBN 978-3-642-18205-1.
- [5] Raghunath Nambiar, Meikel Poess (eds.): Performance Evaluation and Benchmarking. First TPC Technology Conference, TPCTC 2009, Lyon, France, August 24 to 28, 2009, Revised Selected Papers. Lecture Notes in Computer Science 5895, Springer 2009. ISBN 978-3-642-10423-7.