

Transaction Processing Performance Council (TPC): Twenty Years Later – A Look Back, a Look Ahead

Raghunath Othayoth Nambiar¹, Matthew Lanken², Nicholas Wakou³,
Forrest Carman⁴, and Michael Majdalany⁵

¹ Hewlett-Packard Company, 11445 Compaq Center Dr. W, Houston, TX-77070, USA
Raghu.Nambiar@hp.com

² Oracle Corporation, 500 Oracle Parkway, Redwood Shores, CA-94065, USA
Matthew.Lanken@oracle.com

³ Dell Inc., One Dell Way, Round Rock, TX-78682, USA
Nicholas_Wakou@dell.com

⁴ Owen Media, 3130 E. Madison St., Suite 206, Seattle, WA-98112, USA
forrestc@owenmedia.com

⁵ LoBue & Majdalany Magt Group, 572B Ruger St. San Francisco, CA- 94129, USA
majdalany@lm-mgmt.com

Abstract. The Transaction Processing Performance Council (TPC) [1] is a non-profit corporation founded to define transaction processing and database benchmarks and to disseminate objective, verifiable TPC performance data to the industry. Established in August 1988, the TPC has been integral in shaping the landscape of modern transaction processing and database benchmarks over the past twenty years. Today the TPC is developing an energy efficiency metric and a new ETL benchmark, as well as investigating new areas for benchmark development in 2010 and beyond.

Keywords: Industry Standard Benchmark.

1 Industry Standard Benchmarks

Historically, robust and meaningful benchmarks have been crucial to the advancement of the computing industry. Without them, assessing relative performance between disparate vendor architectures is virtually impossible. Demands for audited and verifiable benchmarks have existed since buyers were first confronted with a choice between purchasing one piece of hardware over another, and have been driven by their desire to compare price and performance on an apples-to-apples basis.

Over the years, benchmarks have proven useful to both systems/software vendors and purchasers. Vendors use benchmarks to demonstrate performance competitiveness for their existing products and to improve/monitor performance of products-under-development; in addition, many buyers reference benchmark results when considering new equipment. Finally, benchmarks help vendors improve their products through competition.

The two most prominent industry standard benchmark organizations to emerge from the 1980's are the Transaction Processing Performance Council (TPC) and

Systems Performance Evaluation Corporation (SPEC) [2]. The TPC's primary focus is total system performance under a database workload, including: hardware, operating system, and I/O system. All results have a price-performance metric audited by an independent TPC certified auditor. Like the TPC, SPEC "develops suites of benchmarks intended to measure computer performance. These suites are packaged with source code and tools and are extensively tested for portability before release."⁷ Unlike the TPC, SPEC results are peer audited.

2 Transaction Processing Performance Council

The TPC is a non-profit corporation founded to define vendor-neutral transaction processing benchmarks and to disseminate objective, verifiable performance data to the industry. Omri Serlin and Tom Sawyer founded the TPC in 1988 as a response to the growing problem of "benchmarking," the inappropriate use of questionable benchmark results in marketing promotions and competitive comparisons. At the time, two frequently referenced benchmarks were the TP1 benchmark, originally developed by IBM, and the debit-credit benchmark, which appeared in a 1985 Tandem Computers Inc. technical article that was published by a team led by Jim Gray. With no standard body oversight, vendors took such liberties with these benchmarks that muddied the waters even further [3].

The TPC's first benchmark was TPC-A, which was a formalization of the TP1 and Debit/Credit benchmarks. However, while a formal and accepted benchmark for system performance now existed, there continued to be many complaints of benchmarking. In response, the TPC initiated a review process wherein each benchmark test had to be extensively documented and then carefully vetted by an independent auditor before it could be published as a formal TPC benchmark result. Today, all published results of TPC benchmarks have been audited and verified by TPC certified auditors.

2.1 User and Vendor Benefits

Over the past two decades, the TPC has had a significant impact on the industry and expectations around benchmarks. TPC benchmarks have permanently raised the bar; vendors and end users rely on TPC benchmarks to provide real-world data that is backed by a stringent and independent review process. The main user and vendor benefits of TPC benchmarks are listed below.

- **A trusted and respected auditing process.** TPC-certified independent auditors verify all results as a requirement for publishing a benchmark result. Additionally, after a benchmark result is published the TPC allows for a peer review process –for 60 days every company in the TPC has the right to challenge any published result based on technical correctness.
- **An objective means of comparing price and price/performance.** The TPC has been the most successful benchmarking group in developing a standard means of comparing the price and price/performance of different systems. All TPC testing requires vendors to detail their hardware and software components, along with the associated costs and three years of maintenance

fees, in order to provide the industry's most accurate price and price/performance metrics.

- **Standard benchmarks for corporate and governmental acquisitions.** Typically, as corporations and governments develop their request for proposals(RFPs) to purchase new systems, they scramble to find an objective means of evaluating the performance of different vendor architectures, technologies and products. Prior to the TPC, these users would spend enormous time and resources trying to define a custom benchmark and convincing the vendor community that it was technically sound and fair to all parties. Using a TPC benchmark already accepted by the user and vendor communities eliminates much of this wasted time and resources.
- **Complete system evaluation vs. subsystem or processor evaluation.** The TPC benchmarking model has been the most successful in modeling and benchmarking a complete end-to-end business computing environment. This has helped TPC benchmarks gain recognition as credible, realistic workloads. Most past and many current benchmarks only measure the hardware performance (processor and memory subsystem). TPC benchmarks have led the way in developing a benchmark model that most fully incorporates robust software testing.
- **Objective engineering tools which spur real hardware and software improvements.** TPC benchmarks, especially TPC-C and TPC-H, are well-understood, stable workloads that engineers use on a continuous basis to eliminate hardware and software bottlenecks that lead to real-world performance improvements for users.

2.2 What Makes the TPC Unique?

Since its formation, the TPC has had a marked impact on the server performance industry. It boasts a membership that is who-is-who in the computer industry. These companies can be considered to be business rivals and yet they willingly work together to showcase and compare their products using the TPC organization. Given the marketing stakes, why do companies trust the TPC to fairly validate their claims? What makes the TPC unique?

First of all the TPC provides cross-platform performance and technology comparisons. The organization provides, manages, and maintains the benchmark specification. The results' sponsors independently decide on the best test configuration to showcase their system as long as it complies with the benchmark specification. This has enabled fair and verifiable cross-platform and technology comparisons.

The TPC is the only benchmark organization that requires price/performance scores across all of its benchmarks. This is a realization of the fact that performance has its costs and those costs have to be declared. The price/performance metric has encouraged test sponsors to be more realistic in their test configurations because cost affects purchasing decisions. The availability date of all components of the test configurations has to be declared. These components have to be orderable at the time the result is published. This has limited test configurations to commercially available and viable products.

Furthermore, all results have to be verified and certified by an independent Auditor prior to publication. Next, a 60-day peer review period begins where the result can be challenged.

All tests require full documentation of the components, applications under test and benchmark procedures to enable test replication by any interested party. This full disclosure makes it possible to question and challenge a result and ensures that all published results are credible and verifiable.

The TPC is the foremost database performance verifying organization. Most of its benchmarks are database-centric and database agnostic and a wide variety of databases have been showcased and compared.

Furthermore, all TPC benchmarks demand that a system support ACID (atomicity, consistency, isolation, and durability) requirements for the database system to demonstrate they can meet the reliability and security features required for real systems. Each benchmark defines a series of tests that are administered by the auditor before each benchmark result is published.

In 2006 TPC introduced a common pricing specification across all its benchmarks that is designed to allow a fair and honest comparison for customers to review. It was created in order to guide the customer, the vendors implementing a benchmark, and the auditors on what is acceptable pricing for the purposes of publication.

The TPC pricing specification does not attempt to dictate or exclude business practices in the marketplace. There may be some restrictions on pricing for publication (such as excluding sales and closeouts) that are different from some business transactions that actually take place in the marketplace, but those restrictions are intended to make publication both tractable and comparable during the lifetime of the publication for the majority of customers and vendors.

Several factors make TPC's pricing specification unique and creditable. Pricing must be based upon some pricing model that the sponsoring company actually employs with customers. Furthermore, the published price must be a price that any customer would pay for the priced configuration, and the methodology used must generate a similar price for a similar configuration for any customer.

Pricing must also be verifiable. In a competitive environment, aggressive discounting may occur in certain situations, so the pricing model employed for TPC Benchmark publications might not represent the best or lowest price some customer would pay. It must, however, represent the pricing that could be obtained by any customer in a request for bid to a single vendor.

Situations that occur when requests for bids go out to multiple vendors, and then those bids are used in multiple negotiations to get a better price, are not represented.

2.3 TPC Organization

At the helm of the TPC organization is the General Council, which is composed of all member companies. To expedite the work of the TPC, the General Council has created two types of subcommittees: standing and technical subcommittees. Standing subcommittees are permanent committees that supervise and manage administrative, public relations, and documentation issues for the TPC. The technical subcommittees are formed to develop a benchmark proposal, maintain the benchmark, and evolve the benchmark after development work is complete.

- **General Council:** All major decisions are made by the General Council. Each member company of the TPC has one vote, and a two-thirds vote is required to pass any motion.
- **Steering Committee:** Consists of five representatives, elected annually, from member companies. The Steering Committee is responsible for overseeing TPC administration, supporting activities, and providing overall direction and recommendations to the General Council. The General Council, however, has the final decision in all substantive TPC matters.
- **Technical Advisory Board:** This subcommittee is tasked with maintaining document and change control over the complex benchmark proposals and methodologies. In addition, the TAB studies issues involving interpretation/compliance of TPC specifications and makes recommendations to the Council.
- **Public Relations Committee:** This subcommittee is tasked with promoting the TPC and establishing the TPC benchmarks as industry standards.
- **Technical Subcommittees**
 - **Benchmark Development Subcommittees:** A development subcommittee is the working forum within the TPC for development of a Specification.
 - **Benchmark Maintenance Subcommittees:** A maintenance subcommittee is the working forum within the TPC for developing and recommending changes to an approved TPC Benchmark Standard.

The TPC organization adheres strictly to the TPC Bylaws [4] and the TPC Policies [5].

2.4 TPC Membership and Benefits

The TPC has two classes of members, full members and associate members. While only select organizations are eligible for associate membership, all members enjoy the rights and benefits detailed in the TPC bylaws and policies. The following types of organizations are eligible for associate membership: non-profit organizations, educational institutions, market researchers, publishers, consultants, governments, and organizations and businesses that do not create, market or sell computer products or services [6].

Benefits of membership include:

- **Influence in the TPC benchmarking development process.** Given the TPC's tremendous influence on the competitive arena, members have an opportunity to help the TPC decide which benchmarks to consider for future development, and how the current set of benchmarks should evolve. TPC member participation also provides advance, detailed knowledge of upcoming changes to the TPC benchmarking process. These changes can have a profound impact on how the market perceives products and product performance. Members gain insight into how products will be measured in advance of actual benchmark publication.

- **Timely access to ongoing proceedings.** The TPC's membership roster is virtually a "Who's Who" list of commercial computing, featuring many prominent computer system vendors, software database vendors, market research firms and systems integrators. TPC membership provides access to the TPC's internal Web site, which contains all day-to-day rulings, along with ongoing discussions and member proposals.
- **Product Improvement.** Some organizations participate in the TPC's benchmarking process for the same reasons car companies participate in Formula 1 racing. In the world of racing, many of the new components and technologies applied to producing the fastest Formula 1 cars are incorporated into consumer car design. Similarly, the TPC testing process enables members to produce more robust, higher performing retail products. TPC benchmarks are designed to put systems under maximum stress. Since TPC benchmarks model the basic types of operations that a typical transaction processing system might use, they have a wide-range of applicability to customers' environments.

Since the inception of the TPC, the organization's membership roster has grown substantially. In 1988, the TPC had eight members, while in 1994 the organization had 45 members. After many acquisitions and mergers, the TPC now has 25 full members and three associate members.

3 TPC Benchmark Standards: Past and Present

Over the past two decades, the TPC has had a significant impact on the industry and expectations around benchmarks. Prior to the late 1980's, the industry lacked objective, verifiable benchmarks that were relevant to real world computing. TPC benchmarks quickly raised the bar for what the industry came to expect in terms of benchmarks themselves.

Vendors and end users now rely on TPC benchmarks to provide real-world data that is backed by a stringent and independent review process. TPC benchmarks have encouraged the comparison of price and price/performance on an apples-to-apples basis.

3.1 TPC Benchmark Standards Overview

The TPC's primary focus has been on online transaction processing (OLTP) benchmarks. Over the years, the original TPC benchmarks evolved into new benchmarks or were ultimately retired. The chart below, created by TPC associate member IDEAS International, illustrates the progression of TPC's benchmarks over the past 20 years.

TPC Evolving Benchmarks

- TPC-A evolved into TPC-B, and was ultimately replaced by TPC-C.
- TPC-E is a new OLTP benchmark, which currently coexists with TPC-C.

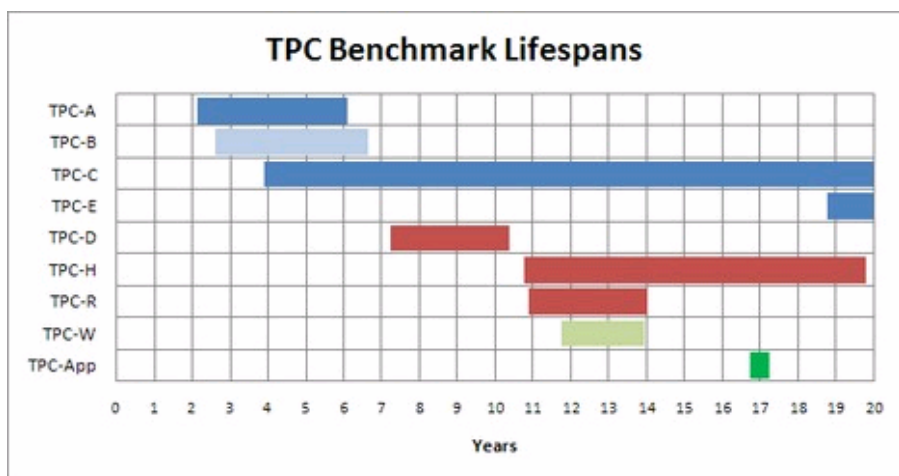


Fig. 1. TPC Benchmark Lifespans [7]

- TPC-D was the first Decision Support benchmark, which as evolved into TPC-H and TPC-R.
- The TPC also developed a web server benchmark, TPC-W, which was later replaced by TPC-App.

TPC Retired Benchmarks

- TPC-R and TPC-W were retired due to lack of industry acceptance.
- TPC-S, TPC-Enterprise and TPC-CS benchmark development efforts were aborted due to lack of support.

TPC “Undecided” Benchmarks

- TPC-DS is the next generation Decision Support benchmark designed to overcome some of the limitations of TPC-H but TPC is unable to reach consensus on TPC-DS.

TPC Current Benchmarks

- TPC-C and TPC-E for OLTP workloads.
- TPC-H for decision support workloads.
- TPC-App for Application Server and Web services.

The TPC-C and TPC-H benchmarks continue to be a popular yardstick for comparing OLTP performance and Decision Support performance respectively. The longevity of these benchmarks means that hundreds of results are publicly available over a wide variety of systems. TPC-E is also gaining momentum.

3.2 TPC Current Developments

Since 1988, the TPC has developed nine benchmarks, each addressing requirements of industry demands. Currently, the TPC’s metrics-under-development include the TPC Energy Specification and the TPC-ETL benchmark [8].

3.2.1 The TPC Energy Specification

Performance and price/performance metrics are key criteria in data center purchasing decisions, but the demands of today's corporate IT environment also include energy consumption as one of the most important considerations. Energy efficiency has become one of the most significant factors in evaluating computing hardware. To address this shift of IT purchasers' priorities, the TPC is developing a new Energy Specification to enhance its widely used benchmark standards. The addition of energy consumption metrics to the TPC's current arsenal of price/performance and performance benchmarks will help buyers identify the energy efficiency of computing systems to meet their computational and budgetary requirements.

The TPC Energy metrics will provide an additional dimension to computing systems' performance and price. As with the TPC's price/performance metrics, which rank computing systems according to their cost-per-performance, the TPC Energy metrics will rank systems according to their energy-consumption-per-performance rates, and will take the form of watts/performance. The ranking of the Top Ten energy/performance systems will be available on the TPC website.

Buyers now demand an objective method of comparing all three factors to select equipment that best meets their changing requirements, and the TPC's Energy Specification is being carefully designed to address this need. Like the TPC Pricing Specification, the TPC Energy Specification is a supplement to existing TPC benchmark standards, rather than a standalone measurement framework. This means that it is intended to be compatible with TPC benchmark standards currently in use, including TPC-App, TPC-C, TPC-E and TPC-H. The result will be metrics that enable comparison of systems on all three axes: price, performance and energy consumption.

3.2.2 The TPC-ETL Benchmark

The TPC-ETL (extract, transform and load) benchmark committee was formed in November 2008. The TPC's intent is to develop an ETL environment workload, which manipulates a defined volume of data, preparing the data for use in a traditional data warehouse (DW). In particular, data extracted from an online transaction processing (OLTP) system is transformed along with data from ancillary data sources (including database and non-database sources), and loaded into a data warehouse. The source and destination schemas, data transformations and implementation rules have been designed to be broadly representative of modern ETL systems. The TPC is encouraging companies interested in participating in the development of this benchmark to join the TPC.

4 Benchmark Development Process

The TPC encourages industry experts and the research community to submit draft standard specifications in a format similar to other TPC benchmark standards. The proposal can be a new benchmark in a new domain (e.g. TPC-ETL), a new benchmark in an existing domain (e.g. TPC-E in OLTP domain where TPC-C is a

standard) or a refinement of an existing benchmark (e.g. TPC-Energy initiative, adding and energy metric to all TPC benchmarks).

After the TPC Steering Committee reviewed all submitted benchmark proposals for contents, applicability and viability it presents recommendations to the TPC General Council, identifying advantages/disadvantages for each submission and a proposed course of action. The TPC General Council may then vote to formally accept the proposal for future work.

If approved, the TPC General Council establishes and empowers a subcommittee to develop a formal benchmark specification. To speed up the benchmark development cycle, the subcommittee is empowered to brief non-members on their benchmark in order to obtain timely feedback.

At each General Meeting, the subcommittee provides a status update on its work, including a working draft of the specification. The TPC General Council may provide direction and feedback to the subcommittee to further their work.

The TPC General Council may also authorize the release of a draft specification to the public. Principal goals include encouraging companies to implement the draft specification, gathering more experimental data, and speeding up specification approval.

Once the specification is of sufficient quality, the subcommittee will submit it to the General Council for formal review and approval. During this phase, the specification will be made available to all members and the public. All comments and proposed changes will be posted to the TPC's private Web site and considered by the subcommittee for resolution.

The subcommittee will propose resolution of comments as an updated specification, which is then reviewed by the General Council. The General Council approves the updated specification by voting to send the specification out for mail ballot. To become a benchmark standard, the specification must be approved by a mail ballot in accordance with policies.

After the specification has been approved by mail ballot, the General Council will establish a corresponding maintenance subcommittee, which will automatically supersede the development subcommittee. Results on different versions of a TPC benchmark standard are considered comparable unless the General Council stipulates restrictions for publicly comparing older version results with newer version results.

Such complex development and approval processes are necessary to the benchmark development process. Creating carefully designed, robust metrics takes a considerable amount of time and resources, and the TPC's extensive review process is designed to minimize potential future revisions.

5 A Look Ahead

The world is in the midst of an extraordinary information explosion, and the need for industry standard benchmarks remains crucial to the computing industry. To meet industry demands, the TPC is exploring potential areas for benchmark development in 2010 and beyond.

The following areas are considered high-priority for future benchmark exploration:

- Appliance
- Business Intelligence
- Cloud computing
- Complex event processing
- Database performance optimizations
- Green computing
- Data compression
- Disaster tolerance and recovery
- Energy and space efficiency
- Hardware innovations
- High speed data generation
- Hybrid workloads or operational data warehousing
- Unstructured data management
- Software management and maintenance
- Virtualization
- Very large memory systems

In the 1980's, the TPC established the basic framework for price and price/performance metrics based on industry need, and ultimately many of the TPC's original benchmarks have evolved and are still in widespread use today. Now, with its first ever Technology Conference on Performance Evaluation and Benchmarking, the TPC is again laying a basic framework for the development of future benchmarks. This conference is meant to solicit new ideas and to provide a discussion forum for the research community and industry experts. The TPC's Technology Conference on Performance Evaluation and Benchmarking is a concrete step towards identifying and fostering the development of the benchmarks of the future. Tomorrow's metrics will likely have as profound an impact on systems/software vendors and purchasers, as the TPC's price and price/performance benchmarks have today.

Acknowledgements

The authors would like to thank the past and present members of the TPC, especially Jerry Buggert, Andreas Hotea, Charles Levine, Mike Molloy and Kim Shanley for their contributions to documents referenced in this paper. Additionally, the authors would like to thank Meikel Poes for his comments and feedback.

References

1. Transaction Performance Council website (TPC), <http://www.tpc.org>
2. Standard Performance Evaluation Corporation website (SPEC), <http://www.spec.org>
3. Shanley, K.: History and overview of the TPC, <http://www.tpc.org/information/about/history.asp>
4. TPC Bylaws, <http://www.tpc.org/information/about/documentation/bylaws.asp>
5. TPC Policies, http://www.tpc.org/information/about/documentation/TPC_Policies_v5.17.htm
6. TPC Invitation to Join, <http://www.tpc.org/information/about/join.asp>
7. IDEAS International, <http://www.ideasint.blogs.com/ideasinsights/2008/10/twenty-years-of.html>
8. TPC Benchmark Status Report, <http://www.tpc.org/reports/status/default.asp>