

Chapter 17

A Review of System Benchmark Standards and a Look Ahead Towards an Industry Standard for Benchmarking Big Data Workloads

Raghunath Nambiar
Cisco Systems, Inc., USA

Meikel Poess
Oracle Corp., USA

ABSTRACT

Industry standard benchmarks have played, and continue to play, a crucial role in the advancement of the computing industry. Demands for them have existed since buyers were first confronted with the choice between purchasing one system over another. Over the years, industry standard benchmarks have proven critical to both buyers and vendors: buyers use benchmark results when evaluating new systems in terms of performance, price/performance, and energy efficiency; while vendors use benchmarks to demonstrate competitiveness of their products and to monitor release-to-release progress of their products under development. Historically, we have seen that industry standard benchmarks enable healthy competition that results in product improvements and the evolution of brand new technologies. Over the past quarter-century, industry standard bodies like the Transaction Processing Performance Council (TPC) and the Standard Performance Evaluation Corporation (SPEC) have developed several industry standards for performance benchmarking, which have been a significant driving force behind the development of faster, less expensive, and/or more energy efficient system configurations. The world has been in the midst of an extraordinary information explosion over the past decade, punctuated by

DOI: 10.4018/978-1-4666-4699-5.ch017

rapid growth in the use of the Internet and the number of connected devices worldwide. Today, we're seeing a rate of change faster than at any point throughout history, and both enterprise application data and machine generated data, known as Big Data, continue to grow exponentially, challenging industry experts and researchers to develop new innovative techniques to evaluate and benchmark hardware and software technologies and products. This chapter looks into techniques to measure the effectiveness of hardware and software platforms dealing with big data.

1. INTRODUCTION TO SYSTEM BENCHMARKS

System benchmarks have played, and continue to play, a crucial role in the advancement of the computing industry. Existing system benchmarks are critical to both buyers and vendors. Buyers use benchmark results when evaluating new systems in terms of performance, price/performance, and energy efficiency, while vendors use benchmarks to demonstrate the competitiveness of their products and to monitor release-to-release progress of their products under development. With no standard system benchmarks available for Big Data systems, today's situation is similar to that of the middle 1980s, when the lack of standard database benchmarks led many system vendors to practice what is now referred to as "benchmarking," a practice in which organizations make performance claims based on self-designed, highly biased benchmarks. The goal of publishing results from such tailored benchmarks was to state marketing claims, regardless of the absence of relevant and verifiable technical merit. In essence, these benchmarks were designed as forgone conclusions to fit a pre-established marketing message. Similarly, vendors would create configurations, referred to as "benchmark specials," that were specifically designed to maximize performance against a specific benchmark with limited benefit to real-world applications.

As a direct consequence of the benchmarking era, two benchmarking consortia emerged: the Transaction Processing Performance Council (TPC) and the Standard Performance Evaluation Corporation (SPEC). The TPC, founded in

1988, defines transaction processing and database benchmarks and disseminates objective, verifiable TPC performance data to the industry. While TPC benchmarks involve the measurement and evaluation of computer transactions, the TPC regards a transaction as it is commonly understood in the business world — as a commercial exchange of goods, services, or money. The TPC offers currently two benchmarks to measure On Line Transaction Processing (OLTP) systems, TPC-C and TPC-E, two others to measure decision support performance (TPC-H, TPC-DS), and one to measure virtualized databases. SPEC has been known mostly for their component-based benchmarks, such as SPEC CPU. It is a non-profit corporation formed to establish, maintain, and endorse a standardized set of relevant benchmarks that can be applied to the newest generation of high-performance computers, including processor-intensive benchmarks, benchmarks to measure graphics and workstation performance, high performance computing benchmarks, Java client/server benchmarks, mail server benchmarks, network file system benchmarks, and SPECpower_ssj2008, a benchmark focused on the relationship of power and performance.

While both consortia follow different methodologies in terms of benchmark development, benchmark dissemination, and benchmark compositions, they follow the same primary goal, namely to provide the industry and academia with realistic, verifiable, and fair means to compare performance.

Since the early days, other more specialized consortia arose, such as the Storage Performance Council (SPC), which is a non-profit corporation

16 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:

www.igi-global.com/chapter/a-review-of-system-benchmark-standards-and-a-look-ahead-towards-an-industry-standard-for-benchmarking-big-data-workloads/85466

Related Content

Multi-Objective Optimization-Based Networked Multi-Label Active Learning

Lei Li, Yuqi Chu, Guanfeng Liu and Xindong Wu (2019). *Journal of Database Management* (pp. 1-26).
www.irma-international.org/article/multi-objective-optimization-based-networked-multi-label-active-learning/232719

A Survey of SQL Language

Hock Chuan Chan, Hong Jun Lu and Kwok Kee Wei (1993). *Journal of Database Management* (pp. 4-16).
www.irma-international.org/article/survey-sql-language/51124

Using Harel's Statecharts to Model Business Workflows

Wai Y. Mok and David Paper (2002). *Journal of Database Management* (pp. 17-34).
www.irma-international.org/article/using-harel-statecharts-model-business/3281

INDUSTRY AND PRACTICE: The Database Designer is a Person Too!

John M. Artz (1994). *Journal of Database Management* (pp. 31-35).
www.irma-international.org/article/industry-practice-database-designer-person/51137

INDUSTRY AND PRACTICE: The Database Designer is a Person Too!

John M. Artz (1994). *Journal of Database Management* (pp. 31-35).
www.irma-international.org/article/industry-practice-database-designer-person/51137