

## Database Performance and Benchmarking

“My one’s faster than your one!”. “Oh no it’s not, mine is much faster than yours!”. No, not a playground conversation between 6 year olds but part of a discussion between two database administrators, each favouring their chosen product. Of course we all like to be associated with good products and like to have this decision vindicated by independent third parties. One of the most trusted organisations that produce objective reviews of database performance is the Transaction Processing Council (known as the TPC, Website [www.tpc.org](http://www.tpc.org)) and the TPC benchmarks are often quoted in product reviews.

But what does the TPC actually do and what is the relevance to database people working in the real world?

### History of the TPC

In the early 1980’s the world of computing was exploding with new innovations and ideas that made computing power available to many organisations that had previously fought shy of IT. As new innovations came about it was impossible to objectively measure the performance of one vendor’s software against another. A benchmark test, called TP1, was leaked from its creator IBM and entered the public domain for the first time. This test simulated a network running ATM machines but did not factor in any user interaction (i.e. delays whilst people think) or network performance limitations. As you would expect this benchmark was wholly abused by a lot of vendors in an attempt to prove the performance of their systems, which lead to a huge damage to the notions of benchmarks and a discrediting of this test in particular.

In 1985 25 computer experts, including Jim Gray (now of Microsoft), designed another benchmark which went by the name DebitCredit. This expanded the idea of the TP1 benchmark and more importantly encompassed real-life issues such as network traffic delays and user think time. Despite this innovation the benchmark tests were still being abused by vendors so dramatic action was required. In 1988 a frustrated database expert Omri Serlin managed to persuade 8 companies to sign up to the newly formed Transaction Processing Council, with the objective of making sense of the benchmarking that was going on. By 1989 the first TPC benchmark, called TPC-A, was released. This built on the success of the DebitCredit benchmark but had a number of additional features including support for local area networks and better use of transactions that complied with the well defined ACID (atomic, consistent, isolated and durable) requirements.

The future of the Transaction Processing Council had been set.

### TPC and Benchmarks Today

Today the TPC is a non-profit corporation whose mission is to define transaction processing benchmarks and ensure that objective, verifiable performance data is available to the industry. The TPC-C benchmark is the most commonly referred to when talking about database performance and we

will cover that in a bit more detail, but there are a number of other benchmarks that the Transaction Processing Council complete, including:

- TPC-H used to measure decision support systems. It represents a decision support system that examines large amounts of data and executes some fairly complex database queries.
- TPC-R which is similar to TPC-H but allows some precompiled queries using knowledge of the data structure
- TPC-W is a transactional Web benchmark simulating an Internet commerce environment. It incorporates activities such as multiple browser sessions, on-line transactions and management of update data contention

Probably the most commonly referred to TPC benchmark is TPC-C.

### **TPC-C**

This benchmark is designed to simulate a business with a number of terminal operators executing one of 5 predefined transactions, such as entering orders and recording payments, against a database. These tasks are fairly generic and the TPC-C test is designed to be applicable across any industry that sells or distributes a product or service.

The business model behind the benchmark is an expansive one. The benchmark comprises a number of warehouses that supply ten sales areas and each sales area has 3,000 customers. The warehouse must try and maintain stock for 100,000 products. The benchmark forces a scale out of the system as new warehouses are added and the impact of this on transactional throughput is measured.

A new order transaction comprises of an average of 10 items which are filled from the 100,000 hopefully in stock. If there is an out of stock (TPC-C forces 10% of these transactions to be out of stock) then the items must be filled from another warehouse. In parallel with this transaction are another four used for stock item queries and payment processing.

The final metric reported by the benchmark is the number of orders that can be filled per minute and is called tpm-C.

The problem with cramming through as many orders as possible is that it can most easily be achieved by throwing as much hardware as possible at the benchmark, with no expense spared. This has resulted in excellent transactional throughput but astronomical hardware costs outside the scope of most average organisations. To overcome this objection the TPC benchmarks now incorporate a cost-per-transaction. The total system cost includes all of the hardware and software along with enough disk capacity to store the benchmark throughput for 180 days, plus the cost of the ongoing hardware/software maintenance. The cost-per-transaction tends to be a favourite measure by many organisations in the “real” world.

For a vendor to undertake a TPC benchmark it requires both time and money. Typically a benchmark test will take 6 months to complete and can cost many hundreds of thousands of dollars.

I suppose at this point I must stress a need to tarnish the silver bullet of official benchmark numbers. Vendors assign some of their top database engineers, often with intimate knowledge of the internal code of the database product and years of experience, to make these benchmarks as good as they can. They work in an ideal lab environment (which means they don't get dragged into pointless meetings every five minutes) and they have access to the engineers that designed and built the hardware. With all this in their favour they will squeeze every last atom of performance out of the system. The chances are that you work in a real job with less of an ideal environment, so the TPC numbers need to be considered as an excellent guide but not an absolute metric that you will obtain. I think of it as the mpg figures given by car manufacturers. This is always calculated with a minimum all up weight of the car driven by an expert test driver. Take your car on the school run and you are unlikely to hit such good fuel economy.

So getting the benchmark right is very important, but sometimes fur does fly and companies disagree.

### **Oracle and the TPC**

In April 1993 a US based consulting firm called the Standish Group alleged that Oracle had added a special feature to their product called discrete transactions. The objective of this new feature was to enable the product to achieve extra high throughput in the TPC-A benchmark. Although Oracle followed the letter of the TPC law it was believed that they had violated the spirit of the Transaction Processing Council. As expected this gained front page news in the IT press and did the credibility of the TPC no good at all, at a time when "benchmarketeering" threatened the whole notion of standard benchmarks.

Luckily this activity spurred the TPC onwards and they produced a prohibition of special code in databases solely for the purposes of forcing higher benchmark results. Interestingly Oracle refused to resubmit the test without the discrete transaction code and withdrew all of their TPC results shortly afterwards although they still remain on the council.

### **And the scores on the doors are...**

On the 19<sup>th</sup> September last year Microsoft submitted a benchmark test for review which was approved by the TPC. Microsoft SQL Server 2000 running on a HP ProLiant DL760-900-256P achieved a whopping 709,220 tpm-C at a cost of \$14.96 per transaction. This is still number one in the TPC-C benchmarks and in fact SQL Server 2000 occupies places 1,2,3,6 and 9 in the TPC-C top ten, running on different manufacturers hardware.

Is my one faster than yours? You bet it is!