

WBDB 2012
Abstract
John Galloway
jrg@toasterfish.com

Here are several issues I see as areas where either great improvement could be made compared to TPC benchmarks, or where such is likely to be required to make an effective benchmark suite for big data.

System performance characterization (i.e. not DBMS)

Model numbers or cluster resource level guarantees might be good for reproducing a result, but the appetite for doing so will decrease as the size expands. It may be more useful in the big data case to know the performance limits of the underlying hardware and OS layers. A set of secondary benchmarks could be utilized to measure CPU performance, memory bandwidth, I/O throughput and latency, interconnect throughput and latency etc.

Likewise knowing what percentage of the above performance limits are being exploited by a particular benchmark may be useful. This is typically the sort of thing that the vendors and sponsors would be doing when investigating system performance. I see no reason it also would not be of interest to potential customers if presented in a meaningful way (e.g. how much “headroom” is there?).

Both of the above could be accomplished by running a supplied set of monitoring tools while the system is being loaded either with the secondary performance tests or a primary benchmark.

Configuration Costs

In TPC benchmarks the difficulty of configuring the hardware and software (OS and DBMS) for a particular test (volume setup, table partitioning, etc) and the difficulty of porting the benchmark code is captured only in the copy of the scripts used to setup the system and run the benchmark. The more complex the hardware and software environments become the more there is a need for automated configuration and the more difficult it becomes to tailor DBMS application code to exploit a particular environment. Hence being able to quantify the configuration and porting or tailoring costs as a metric might be useful as a means to add pressure in this direction. However, I have no idea how to measure this.

Supplied Kit and Results Processing

The benchmark should record all relevant data including timings and metrics in a database. This implies little if any vendor written frameworks or drivers would be used which should ease audit requirements and lower the amount of work needed to be done by a vendor. The more that the code can be separated into a standard kit with a minimal vendor specific or porting layer, the easier it will be for everyone.

Result Presentation

Big data likely implies a greater importance for a capacity to grow with time and hence a key characteristic in big data benchmarks may be performance over a range of attributes (DBMS size, user load, system/cluster size, etc). A simple integer metric is unlikely to be sufficient to express this type of a family of results in a meaningful way especially when trying to compare multiple such results. Plenty of work is being done on ways to express and compare multi dimensional data sets and such methods should be exploited here as opposed to the traditional "Top 10" style of simple metrics.