ghaff@
illuminata.com

Gordon R Haff
Illuminata, Inc.

TPC-C Passes Escape Velocity

Quick Note

Gordon Haff
2 December 2004

It was just earlier this year that Hewlett-Packard's Superdome broke the one million tpmC mark,[1] nudging past a transaction processing milestone that seemed impossibly distant just a few years ago. That result was especially notable because it surpassed one of those big round numbers to which people like to attach outsized significance. However, it was really just the latest in a game of leap-frog with IBM for the high water transaction processing mark—Sun having long since quit the race. HP's million-topper beat IBM's previous high mark by 31 percent. IBM would come back four months later, squeezing back into the lead by a mere 2 percent. This is the sort of back-and-forth seesaw battle that has characterized the TPC wars and benchmarking in general for years.

It's a costly battle for the vendors. Big labs with big staffs and millions of dollars of hardware go into running these benchmarks. To be sure, they provide some large configuration testing and tuning benefits that go beyond the benchmark itself. But a great deal of effort and money still goes toward the narrow task of producing the best benchmark number possible. And, although no meaningful system-level benchmark is cheap or simple to run, TPC-C has become the costliest of them all. By requiring the scaling of database size and query loads with system horsepower, it requires *oodles* of hardware, especially disk drives—so many disks, in fact, that storage costs dwarf the cost of the system itself under test.

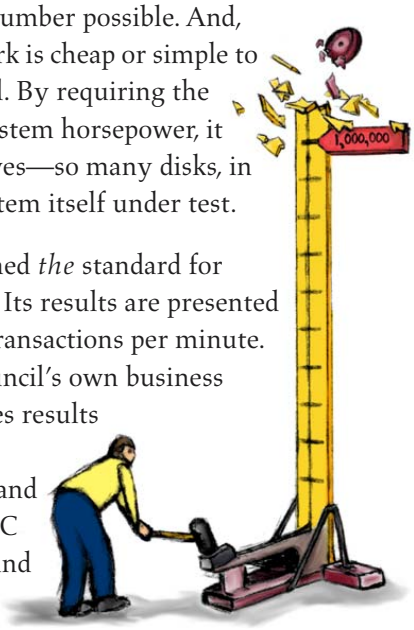

Yet, for several reasons, TPC-C has long remained *the* standard for measuring transaction processing performance. Its results are presented as a simple, relevant, understandable metric—transactions per minute. Contrast this to the Transaction Processing Council's own business intelligence benchmark, TPC-H, which expresses results non-intuitively as QppD. What in the heck is a QppD?! And unlike the also popular SAP Sales and Distribution (SD) application benchmark, TPC-C results are fully verified by an outside auditor and explicitly include both system and maintenance

1. The TPC-C discussion in this note refers specifically to "non-clustered" TPC-C results which benchmark a single system running a single OS image. Historically, clustered results have often been higher, but they also frequently depend on highly artificial database organization and optimizations, making them unrepresentative of "real world performance" available to most IT shops.

*E7DA3ACDF58200D8DC6

pricing, which is also independently verified. TPC-C also makes a reasonable effort to truly test system-level performance. For example, like many "real world" transactional applications, TPC-C penalizes systems with poor memory access latencies. Finally, TPC-C has historically been run by a wide range of companies on a wide range of system configurations—all intent on the game of benchmark leapfrog—thereby making comparisons between different vendors and different chip and system architectures possible. Of course, any single benchmark is—at best—an imperfect measure of system performance running real workloads and applications. That's why sensible customers evaluate systems based on a portfolio of benchmarks rather than just one—or even run custom benchmarks based on their own environments. But, that said, TPC-C has generally managed to maintain a reputation for modeling, as well as any synthetic benchmark can, real-world transaction processing performance.

All this now may have changed—at least at the high-end. For IBM has blown the doors off the benchmark. In a contest where a 3 percent advantage is generally considered a genuine win, and a 30 percent advantage considered a trouncing, IBM has turned in a result (over 3.2 million tpmC) that is more than 300 percent better—yes, *three times* the result of its closest competitor. That is not just a win, it is a brutal stomping.

Yet amidst the popping of champagne corks in IBM's Austin, TX benchmarking labs, they must also feel a bit of *fin de siecle* trepidation. Given the margin of the win, the staggering cost of running these benchmarks, the increased role of workload management and virtualization rather than single workload scale,[2] and the increasing distance between TPC-C benchmarks and real applications, it may *never* make sense for anyone to top the mark that IBM has now set.

Consider the size of the benchmark configuration. The system was an eServer p5 595 equipped with 64 1.9 GHz POWER5 processors and 2 TB of memory. Yet that takes up all of about half a cabinet (albeit a pricey half-cabinet, given that it represents 8 million dollars[3] of memory alone!). Adding a full boat of I/O Drawers brings us up to two loaded cabinets. But then there are the disks. It takes 6,400 disks to hold the database required for the benchmark.[4] That number a bit hard to physically picture? Let's put it this way: that's 45 cabinets full of disks. The total list price of this configuration was over 30 million dollars. Even the heavily-discounted Three-Year Cost of Ownership is almost 17 million dollars.[5]

Or consider the size of workload that this benchmark represents. Scaled to this performance level, the benchmark posits a company that has a transaction processing environment with over 2.5 million users servicing almost 8 billion customers from 2.5 million sales districts and 256,000 warehouses.[6] Now that's a company with *serious* market share! Of course, no such corporation or government agency actually exists—nor is one likely to in our lifetime. Even if we posit a real life company whose transactions are fewer but individually more intensive, thus generating a similar level of system load in ways not directly represented by the TPC-C benchmark, it's still evident that few companies or organizations could ever generate this sort of transactional load from a single application.

The conclusion is gratingly obvious: At the scale at which it is currently being run, TPC-C has lost its similarity to the business processing tasks that go on in real customer scenarios. Sure, someone—

2. See Illuminata report "IBMís p5 Mothership Whizzes into Orbit" (October 2004).

3. Unless otherwise noted, all pricing is list from the audited Full Disclosure Report (FDR) that is part of the benchmark result. See http://www.tpc.org/results/FDR/TPCC/IBM_595_64_20041118_FDR.pdf for IBM's FDR.

4. As is commonly done with TPC-C benchmarks, IBM used modest-sized 36 GB, 15K RPM drives so that the configuration would have more disk spindles—resulting in better I/O performance.

5. The discounted price is intended to reflect the actual price that a customer would pay for a one-time purchase.

6. See http://www.tpc.org/tpcc/detail.asp for a detailed description of the TPC-C benchmark.

whether HP or IBM or another—could spend the money to set a new top mark someday in their next product cycle or the one thereafter. But the system surpassing IBM's result would bear little relation to real customer systems and environments. The test configuration would probably cost something like 40 or 50 million dollars at list price, and would include a database modeling perhaps 12 billion customers. Perhaps such a result would be meaningful to the Intergalactic Acme Corporation—but to few, if any, earthbound companies today.

None of this diminishes the loftiness of IBM's accomplishment in the here and now. Indeed, it's hard to find anything at which to niggle. The next closest competitive result, an HP Superdome with 64 Itanium 2 processors, isn't particularly recent (November of 2003), but that system design hasn't changed and it seems unlikely that more recent Itanium 2 iterations would do more than nibble at IBM's three-fold advantage. Nor is IBM's price/performance anything to grouse about. At $5.19 per tpmC, it's second best in the top-10 tpmC results, trailing only another eServer, the 16-processor p5-570. Finally, although it's true that IBM ran this benchmark with its own DB2 UDB database—a database that competitors like to claim is better at benchmarking than running customer workloads—IBM has previously published 32-way p690 results that show Oracle and DB2 within a few percent of each other.[7] This strongly suggests that, even if DB2 does give IBM some incremental performance, it's not a difference on which the whole scale of the benchmark depends.

No, IBM has truly delivered a result on the venerable TPC-C metric that distances it from its competition far more decisively than is usually possible in this hyper-competitive industry. And if IBM's run the TPC-C benchmark up against its practical limits in the process? Beyond making it less feasible for others to follow in its footsteps, it decidedly highlights the industry's need to develop better benchmarks—ones that are more realistically aligned with what are now readily-achievable scale-points, as well as the increasing reality that today's Big Iron systems are far more about running multiple workloads in a highly virtualized environment than they are about hyper-scaling any single application. Perhaps TPC-C or a related successor can continue to have a place in such a future, but it will need to be run at more workaday sizes, and even side-by-side with benchmarks that measure other aspects of system prowess.

7. The 768,839 tpmC Oracle 10g result, published in February 2004, actually inched out the DB2 result from the previous November. One suspects that Oracle wouldn't have agreed to the publication (as they have the right to do) had it not come out on top.