

Wednesday, October 14, 2009

IBM's reaction to Sun/Oracle TPC-C announcement

I expected a reaction from IBM in regard of Sun/Oracle. But not such a lame one as provided by Mrs. Stahl of IBM. Let's forget for a minute that TPC-C is a horrible benchmark. What is a benchmark? A benchmark is a measure how fast a system can do things. Measuring how fast a system can count to one million is as well a benchmark as doing the general ledger for a company. When it's your job to count to one million you would use the fastest system in that benchmark. If it's your job to do a general ledger, you would use a benchmark testing that task.

TPC-C is now a benchmark testing the capabilities of the sytem to do OLTP loads. Of course this benchmarks reached numbers beyond any reasoning in the real world. But for a long time IBM saw in this benchmark a good representation of OLTP loads.

So let's just assume you have a load that you think is reasonably near to the TPC-C load. What are the interesting numbers for you: Well of course performance of the complete architecture. But there are several important other metrics: Performance per \$ for example (as your budget has limits), or performance per square meter (as your datacenter space has limits). Given that you have a load like the one specified by TPC-C it's irrelevant how you reached the performance, just that you reach the performance.

The Sun benchmarking team has demonstrated, that their results were better as well in performance per dollar (Sun \$2.34 vs. IBM \$2.81 per tpmC) as well as in performance per rack (Sun 9 Racks vs. IBM 76 Racks).

IBM tries now to tell the public, that this benchmark shows that Power is still superior, because they have more tpmC per Core. 4.7 times more tpmC per core - to be exact. I could calculate another derivate out of this - Price per Core in the configuration used for the benchmark: \$18,051,719 / 384 cores = \$47,009/core at Sun and \$17,111,788 / 64 cores = \$267,371/core for IBM. In the TCO calculation the single core at IBM is 5.6 times more expensive. This is the context you have to know, when Mrs. Stahl writes that a single core delivers more performance per core.

But both is totally irrelevant for this benchmark. What counts performancewise is the number of TPC-C transactions per minute. Period. When you want to proof the impact of a better single-core performance, use a different benchmark. The business of a customer is doing transactions in his architecture on his complete conglomerate of servers, storage and networking, not doing transactions on a single core. And the customer buy a certain number of cores to do the job, not a single core. For the customer the way to yield a certain level of performance is irrelevant. Cluster or single system? Irrelevant! Fewer fast cores or more slower cores? Irrelevant! Of course the customer must be sure, that he has chosen the right benchmark ... that TPC-C is representative for his or her workload, but that's the real problem of TPC-C.

IBM has three options now: Admitting that Sun/Oracle has beaten IBM in TPC-C or admitting that TPC-C isn't a good benchmark. Both wont happen from my perspective: The first option would be really bad PR-wise. The second one would be bad, as TPC-C is the spine of many relative performance numbers, that IBM is using in competitive sizings as far as i know. Even the independent RPE benchmark from IDEAS contained TPC-C. As IBM mastered the art of optimizing for TPC-C and the TPC-Values for Sun were estimated, those relative numbers were somethat skewed. However, both outcomes would be advantageous for Sun.

But i assume, they will just opt for the third option and will spindoctoring the results like in this article of Mrs. Stahl at developerworks. Must be bad, when the student is beaten by a horde of toddlers in squirrel costumes.

PS: I've read the FDR for the Sun result the second time now ... i have somehow the inkling, that this system use by Sun fought with one arm bound to the back. But i think i have to read it a third time to substantiate my thought.

Disclosure Statement: TPC Benchmark C, tpmC, and TPC-C are trademarks of the Transaction Performance Processing Council (TPC). 12-node Sun SPARC Enterprise T5440 Cluster (1.6GHz UltraSPARC T2 Plus, 4 processor) with Oracle 11g Enterprise Edition with Real Application Clusters and Partitioning, 7,717,510.6 tpmC, \$2.34/tpmC. Available 12/14/09. IBM Power 595 (5GHz Power6, 32 chips, 64 cores, 128 threads) with IBM DB2 9.5, 6,085,166 tpmC, \$2.81/tpmC, available 12/10/08. Source: <http://www.tpc.org>, results as of 10/11/09.

Posted by Joerg Moellenkamp in English, Oracle at 16:17

Blog Export: c0t0d0s0.org, http://www.c0t0d0s0.org/

also compare response times of the top three results for power/sparc/intanic - sun is way in front here
Anonymous on Oct 14 2009, 16:57

I'm writing an article about some key takeaways from this benchmark. The response time is part of it.
Anonymous on Oct 14 2009, 17:03

Hmmm... if Mrs. Stahl of IBM claims that the cores are 4.7 times more powerful, perhaps Oracle needs to raise the licensing fee per core on the IBM so it's 4.7x more expensive than the Sun?

I mean, even IBM is saying it's 4.7x more powerful... therefore to equalize the cost of licensing for the same work done.
Anonymous on Oct 14 2009, 17:31

Well, the results are 'In review', not yet accepted.

Salient points (imho) in the Executive Summary :

a. Why does the champion of Linux (IBM) use as 'other soft' Microsoft Visual C++ and COM?
b. So twice as much for rebate for IBM (-20,273,753) as for Oracle (\$2,006,628)? That seems to me that IBM is giving their hardware(+maintenance) away? (cfr. subtotal page2)

Questions :

1. POWER 7 comes next year : does that mean that all the DDR2 memory is for the dustbin? (\$2,602,598)
2. Last they told me that IBM still counts \$1 = 1€ : does Sun/Oracle do so too? If not what would the prices in Europe/rest of the world then be?

Elephant in the room : do I see anywhere the cost of the computer room? Which if in Tokio would dwarf the rest of the costs? (advantage Sun, no contest)

Congratulations Sun & Oracle!

P.S. Which entity has six or seven million users for one box? That's nearly 20 times all the employees of IBM, isn't it? Just THINKing for myself
Anonymous on Oct 14 2009, 17:43

Funny how fast IBM yell foul when they discovered someone else beat them using a clustered results. They themselves run a clustered results for TPC-H using 32 POWER6 p570 when they could have done it with a fully configured p595 (http://www.tpc.org/results/individual_results/IBM/IBM_570_10000GB_20071015_ES.pdf) , not to mentioned running it on AIX 5.3 instead of the so-called latest and greatest AIX 6.1.
Anonymous on Oct 14 2009, 17:50

Errata :

So twice as much for rebate for IBM (-20,273,753) as for Oracle (\$2,006,628)?
Ten times as much, decimally speaking of course. Doh!

P.S. six or seven : 5184000 of 6480000 users, I seem to have seriously rounded up the numbers here (people in meeting?), but the point is still valid. (Can't we knife this benchmark just for this reason alone : common sense?).
Anonymous on Oct 14 2009, 19:04

I would like to kill it ... perhaps this get easier, now we lead it. But you don't believe how many people consider this as a good benchmark ...
Anonymous on Oct 14 2009, 19:07

One arm tied to the back? Isn't the Sun config fault tolerant, whereas the IBM config isn't? At least, that was mentioned by Larry during his keynote on Sunday...
Anonymous on Oct 14 2009, 19:26

IBM:

Winter is coming, do you feel the cold? The Sun is still shining....
Anonymous on Oct 14 2009, 20:13

They will probably just add "non clustered" to the marketing material with a small font now, and use that from now on, just like she is doing in her post:
"...with IBM DB2 9.5 is the best overall non-clustered TPC-C system..."
Anonymous on Oct 14 2009, 20:28

Keep in mind:

- IBM used 11,000 disks in their benchmark for the server and storage
(8 x 146GB SCSI disks and 10,992 73.4GB FC disks)

calculate:

- 5W per disks = 55,000 W for all disks

- then you have to add some power for the controllers assume 100W per controller - that's 68 * 100W = 6.8 KW for the disk controllers)

- next you have to power your Power server (will be 20,000 W)

- adding this: $55 + 6.8 + 20 = 81.8\text{KW}$

Now we know that the IT has also to be cooled - they didn't mention any water cooling in the FDR, so it will be conventional cooling technology - the rule of thumb is "1W cooling per 1W IT" - so you need another 81.8KW for cooling. In total that's 163.6KW for the IT and cooling. Assuming a price of 5 cent per KW/hour (very cheap) you have to pay 8€ (rounding down) for each hour you're using this equipment. That's 192€ per day that will be 70,080€ per year and that will leave you with an energy bill of 210,000€ in three years.

Anonymous on Oct 15 2009, 08:54

Joerg,

All this time, you whine about how TPC-C is a "terrible" benchmark. Now, Sunacle replaces HDD with Flash, which results in a whopping 26% increase over an IBM result with HDD (big surprise), and all of a sudden TPC-C is important.

Who are you kidding?

Anonymous on Oct 16 2009, 03:58

Sorry, but I wrote in every article about this TPC-C article, that I still think that TPC-C that tpc-c is horrible. Just read the blog entries. The day the TPC is retiring TPC-C is a day worth of celebration with wine and singing all night long.

From my perspective this result is important out of a single reason: Many customers still think that that TPC-C is a reasonable benchmark. In the years before this, we had to tell them "It's a nonsense benchmark." and the customers thought "They say this because they are bad in it". And when they saw, that we didn't report a result they thought "Oh, they are really bad in it".

Now we can say "It's a nonsense benchmark." and the customer thinks "Oops, they lead the benchmark, and still say it's a nonsense benchmark". And then we can talk about reasonable methods to compare the performance of architectures. That's the real importance of this TPC-C benchmark result.

Besides of it, it gives some interesting insight about the load stability of some features in Solaris (like the FC target in COMSTAR)

If marketing has some opportunities to fire salvos in direction of IBM ... I'm fine with it, but in my perspective it's not the important outcome.

Anonymous on Oct 16 2009, 07:15

I mind, what the opinion of Mr. Bluetothebones may be.

Anonymous on Oct 17 2009, 12:15

Well my point is, did Oracle/Sun did the apple to apple comparison.?

This new benchmark from Oracle/Sun is the same TPC-C benchmark that Oracle promised to reveal on October 14, which led to a reprimand by the Transaction Processing Council (TPC) for having made an unsubstantiated superior performance claim. (link to press release) In order to demonstrate "the world's fastest database performance with a TPC-C benchmark" Oracle's used an Oracle RAC configuration (database for clustered servers) and a 12-Node Sun SPARC Enterprise T5440 cluster. It achieved 7,717,510 tpmC, with \$2.34/per tpmC. Oracle compared this result to IBM's single system 64-core IBM Power 595 with IBM DB2 9.5, based on a TPC benchmark submitted on June 10, 2008.

Here's a deeper look at Oracle's claim and the comparison to IBM's system:

- IBM system is a single server while the Sun system is a configurations with a 12 node cluster.
- The Sun cluster had a total of 384 processor cores and 3,072 threads compared to 64 cores and 128 threads in the IBM system.
- With 512GB of memory per node the Sun cluster had a total of 6TB of memory, compared to 4TB in the IBM Power 595 system.
- While the Sun tpmC per core is 20,097; the IBM tpmC per core is 95,080. The IBM result had 4.7 times higher performance.
- The Sun system will not available until 12/14/09. The IBM system has been available since December 10, 2008.
- The flash technology used in the Sun result was integral to the performance and price/performance. Flash technology wasn't even being used in benchmarks when the TPC-C published the IBM system result on June 10, 2008. IBM has leadership results using SSDs in benchmarks .
- For clustered systems, see the outstanding systems performance of IBM DB2 pureScale on Power Systems.

Anonymous on Nov 13 2009, 14:27

Did you read the article I wrote, or do you just copy over the IBM website here ? (<http://www-03.ibm.com/systems/migratetoibm/getthefacts/powersystem.html>).

Of course it's an apple to apple comparison, because it's exactly the same workload. When you have a workload that is represented by TPC-C then the Sun system is clearly faster and cheaper than the IBM system. It does it differently, but that's irrelevant. I've explained that in my article ...

Of course you can argue about TPC-C in its entirety, but that's another question.

To your copied text I stated already in the article you copied: "IBM tries now to tell the public, that this benchmark shows that Power is still superior, because they have more tpmC per Core. 4.7 times more tpmC per core - to be exact. I could calculate another derivative out of this - Price per Core in the configuration used for the benchmark: $\$18,051,719 / 384 \text{ cores} = \$47,009/\text{core}$ at Sun and $\$17,111,788 / 64 \text{ cores} = \$267,371/\text{core}$ for IBM. In the TCO calculation the single core at IBM is 5.6 times more expensive. This is the context you have to know, when Mrs. Stahl writes that a single core delivers more performance per core.

But both is totally irrelevant for this benchmark. What counts performance-wise is the number of TPC-C transactions per minute. Period. When you want to prove the impact of a better single-core performance, use a different benchmark. The business of a customer is doing transactions in his architecture on his complete conglomerate of servers, storage and networking, not doing

Blog Export: c0t0d0s0.org, http://www.c0t0d0s0.org/

transactions on a single core. And the customer buy a certain number of cores to do the job, not a single core. For the customer the way to yield a certain level of performance is irrelevant. Cluster or single system? Irrelevant! Fewer fast cores or more slower cores? Irrelevant! Of course the customer must be sure, that he has chosen the right benchmark ... that TPC-C is representative for his or her workload, but that's the real problem of TPC-C."

I appreciate any discussion but please stop to copy IBM marketing in the comment field ...
Anonymous on Nov 13 2009, 14:43

Yes I copied the from the same side cause you cannot measure 3000CC car with 3500CC car, for measuring you have to be on same level that's the important point, I think you didn't read the copied article that's why I said "ITS NOT APPLE TO APPLE COMPAIRISON" and secondly your comparison and measurements are from Oracle side or if you want me to copy the side ...
Anonymous on Nov 20 2009, 14:10

Yes I copied the from the same side cause you cannot measure 3000CC car with 5000CC car, for measuring you have to be on same level that's the important point, I think you didn't read the copied article that's why I said "ITS NOT APPLE TO APPLE COMPAIRISON" and secondly your comparison and measurements are from Oracle side or if you want me to copy the side 
... Anonymous on Nov 21 2009, 20:52

Who cares how fast one core is? The entire CPU is the important thing. If my old Volvo car has a faster piston in the motor, does that make my old Volvo faster than a Porsche? No. You can not study one core, you must stude the entire CPU. No one buys one core, you buy an entire CPU. Therefore, compare cpu vs cpu. Not core vs core. Or, I might as well as compare ALU vs ALU, or register speed vs register speed, etc.
Anonymous on Nov 25 2009, 11:47

At first: I saw your comment this weekend, but as you were so eager to repost it, i've decided your article could need some maturing in the queue of articles that await moderation ;--> i'm not looking into that queue every day ... thus it can take a while until i'm approving comments on articles older than 30 days. This time limit helps me to reduce the level of spam ... as spammers often choose older articles.

At second: You didn't understand my point. I can do this comparision because it's an apple to apple comparision. It's an application benchmark. This benchmark tests architectures on their speed for solving the problem defined in the TPC-C benchmark rules. Of course i've read the article because it was the reason why i wrote the article you've commented at first. I would like to ask you, to think about this point from my article: " But both is totally irrelevant for this benchmark. What counts performancewise is the number of TPC-C transactions per minute. Period. When you want to proof the impact of a better single-core performance, use a different benchmark. The business of a customer is doing transactions in his architecture on his complete conglomerate of servers, storage and networking, not doing transactions on a single core. And the customer buy a certain number of cores to do the job, not a single core. For the customer the way to yield a certain level of performance is irrelevant. Cluster or single system? Irrelevant! Fewer fast cores or more slower cores? Irrelevant! Of course the customer must be sure, that he has chosen the right benchmark ... that TPC-C is representative for his or her workload, but that's the real problem of TPC-C."

TPC-C measures performance of an architectures on a certain workload, it doesn't test the CPUs. Of course IBM could just use multiple p595 but there is a problem with it and this is the only good thing about TPC-C, it includes a performance/price comparison.

At the end this discussion is useless, as TPC-C has left the realm of realistic workloads long ago. From my field perspective this benchmark just helps me to kill of TPC-C in the minds of customers. This is easier when you in lead the benchmark. Marketing may thing otherwise, but that the basic essence of my view to it.
Anonymous on Nov 25 2009, 14:14