

Transition to a new CPU benchmarking unit for the WLCG

Date: 27/01/2009

Contributors/Authors:

- Chairman: Gonzalo Merino
- From the HEPIX Benchmarking WG: Manfred Alef (GridKa), Helge Meinhard (CERN), Michele Michelotto (INFN), Alex Iribarren (CERN).
- From the experiments: Peter Hristov (ALICE), Alessandro De Salvo and Franco Brasolin (ATLAS), Marie-Christine Sawley and Gabriele Benelli (CMS), Hubert Degaudenzi (LHCb).

Introduction

In the WLCG MoU both the experiment requirements and the site pledges use thousands of SPECint2000 (kSI2K) as the units for CPU capacity. These were also the units used in the elaboration of the Computing TDRs of the experiments, which date from 2005. At that time the SI2K benchmark results published at the www.spec.org site for different processors showed a good linear scaling with the experiment applications.

In June 2007, after noting that the published SI2K results for the then new processors did not scale linearly any more with the HEP applications, the WLCG project agreed on an ad hoc recipe to run the SI2K benchmark that corrected the problem (see Ref. [1]). This was adopted as an interim solution and WLCG sites were officially requested to compute their pledged CPU capacity following this recipe from September 2007.

The SPECint2000 benchmark was retired in February 2007, and as a consequence no new results were published at the www.spec.org site nor was support offered any longer afterwards. The HEPIX Benchmarking Working Group was therefore asked to help in the definition of a new CPU benchmark unit. In July 2008 the HEPIX Benchmarking WG presented a proposal for a new CPU benchmark to the WLCG Management Board (MB), based on the reasonable correlations seen with the four experiment applications. The MB formally accepted the proposal.

In October 2008 the MB agreed to create a Working Group to propose a plan for converting the existing site installations and the experiment requirements to the new CPU benchmarking unit. The mandate for this group, which is addressed with this report, was set as:

- Agree and publish the recipe and the conditions under which the benchmark is run.
- Agree on conversion factors to convert experiment requirements and site pledges tables from SI2K units to the new units.

A detailed description of this new benchmark, proposed to be codenamed HEP-SPEC, has been written by the HEPIX Benchmarking WG specifying the conditions under which it is to be run by the WLCG sites to compute the power of their CPU installed capacity and is available at Ref. [2].

Background

Within the work of the HEPIX Benchmarking WG, a cluster (LXBENCH¹) with eight machines was setup at CERN to be used as reference platforms for running benchmarking codes.

The old (SI2K) and the new proposed (HEP-SPEC) benchmarks have been run on every machine. The results obtained are summarized in Table 1. During 2008, the four experiments also ran their applications in this reference cluster and they all see good correlation between their relevant codes and the new proposed benchmark². For the detailed results of these experiment results see Ref. [3].

¹ <https://twiki.cern.ch/twiki/bin/view/FIOgroup/TsiLxbench>

² For CMS, a number of runs using a production version of the CMSSW suite were executed in January 09 on the most representative architectures of the LXBENCH cluster. The results of the cases that could be run in the short time period showed no significant differences between the behaviour of the specific production version and the early CMSSW pre-prod used in spring 08 for the more complete analysis.

Hostname	Processor type	HEP-SPEC (new)	KSI2K (old)	Ratio new/old
lxbench01	Intel Xeon 2.8 GHz	10,24	2,25	4,55
lxbench02	Intel Xeon 2.8 GHz	9,63	2,24	4,29
lxbench03	AMD Opteron 275	28,03	6,20	4,52
lxbench04	Intel Xeon 5150	35,58	8,51	4,18
lxbench05	Intel Xeon 5160	38,21	9,27	4,12
lxbench06	AMD Opteron 2218	31,67	6,85	4,62
lxbench07	Intel Xeon E5345	57,52	14,19	4,05
lxbench08	Intel Xeon E5410	60,76	15,83	3,84

Table 1 – Results for the old and new proposed CPU benchmarks on the eight reference machines of the LXBENCH cluster at CERN.

An additional exercise has been carried out which is to compute the CPU performance delivered by the whole GridKa farm with both the old and the new benchmarks. This provides a measurement with a realistic mixture of processors in a large WLCG site. The obtained result for the ratio new/old benchmarks is 4,19.

Proposal

To find a conversion factor between the old and new CPU benchmarks this group proposes to focus on the results obtained in the most modern reference machines. In particular, those with Intel Quad-core processors. An additional argument to support this choice is the fact that, as of today, more than 70% of the CPU cores provided by the Tier-0 and Tier-1 centres are of these type of processors. From the results in Table 1 we can see that the new/old units ratio for the machines with this type of processor (lxbench07 and lxbench08) fluctuates around 4,00 within a 5% band, as also does the measurement for the whole GridKa site previously described.

The members of this group believe that uncertainties of the order of 5% are inherent to this kind of CPU power measurements when realistic scenarios are considered. On the other hand, the requirements that these measurements have to be compared with, also suffer from uncertainties of the same order.

Given this situation, rather than arguing about different values within this 5% resolution, this group proposes a simple conversion rule which eases the whole transition process. Therefore the proposal is to adopt 4,00 as conversion factor representing the benchmarks ratio HEP-SPEC/kSI2K.

Once a fixed value has been agreed for the conversion factor, this group proposes the following scenario for the transition period, which should conclude before the April 2009 CRRB:

1. The current tables for requirements and pledges are translated into the new units as they are, by just applying the agreed factor.
2. Sites buy the SPECCPU2006 benchmark and calibrate their farms to be able to report their current CPU power in the new unit. The appropriate scripts and configuration files are provided by the HEPIX Benchmarking WG at Ref. [2].
3. A technical group is appointed to propose a detailed plan on how to make the migration to the new units regarding the CPU power published by sites through the Information System and stored in the Accounting system. Ideally, the plan should allow a coexistence period for the two units.
4. For the April CRRB, sites publish their updated pledges tables in the new units.
5. For the April CRRB, experiments re-compute their requirements tables given the new LHC schedule. New numbers should be computed already in the new unit.

Finally, this group also proposes to set up a web site for WLCG sites to publish the HEP-SPEC results obtained for their machines, so that other sites and the experiments can query these results. Publication of results can be automated at the level of the script to run the benchmark by enabling the reporting flag.

References

- [1] CPU capacity metric for WLCG, May 2007:
http://lcg.web.cern.ch/LCG/documents/Processor%20performance%20metrics_26may07.pdf
- [2] Running HEP-SPEC, Jan 2009:
<https://twiki.cern.ch/twiki/bin/view/FIOgroup/TsiBenchHEPSPEC>
- [3] http://hepix.caspur.it/processors/michele_combined_fits_spec-cpp.xls