

WLCG HEP-SCORE Deployment Task Force

Meeting on 28 Sep 2022 at 14:00 h UTC (teleconference)

Notes

Indico event page: <https://indico.cern.ch/event/1204945/>

Welcome, note-taking, notes from the previous meeting, matters arising

The minutes from the previous meeting are approved. Thanks to Michel Jouvin for preparing them.

After workshop debriefing (Domenico Giordano)

Round table

Domenico proposes a round table of our impressions from the HEPscore Workshop, and he shared a [summary document](#) to discuss it.

Helge congratulates Domenico and Randy on the agenda. He found it very useful being able to discuss the next steps of the TF and other aspects like Vincenzo's contributions. He highlighted the importance of having knowledge about what is actually going on at a low level which leads to optimization.

There was good consensus on the benchmarks proposed for HEPscore 22, and we should continue pushing for this year. We should make sure Alice reconstruction is ready to close everything as soon as possible.

Helge later congratulated Gonzalo on all the work done around the benchmarking infrastructure and the HEPscore workshop itself, which Gonzalo appreciated.

Andrew commented briefly on how happy he was about the workshop. He pointed out the importance of meeting the next resource request and later provisioning not to wait for another cycle.

Bernd Panzer found the workshop very useful as well.

Matthias thought the workshop helpful, highlighting the importance of meeting in person and remarked on the importance of going further than x86, especially the part considering GPUs, both alone and together with CPUs.

Oxana thought the workshop was a good summary of what was achieved and managed to set clearer steps for the future. She added that more follow-up on the accounting implementation is needed.

Stefano agreed on the usefulness of the workshop, as well as the discussion between computing coordinators. He agrees on HEPscore 22 so next spring we can take it into account for the machines that will arrive in 2025. Regarding Alice, they already have a new tag and are ready to start the testing of the reco part alone.

Randy asked whether we would like to run the new reco workload only at CERN or on as many sites as possible. Domenico answered that we will certainly do it at CERN and will also request the sites to run it as well, even though not as many results as in the past are expected due to some machines perhaps not being available any longer. Still, the amount of data should be sufficient to validate it.

Tony Wong, thanked Domenico, Randy, and Gonzalo for all the effort we have put into this project. He stated that he had a very keen interest in this benchmark, and in it being available soon to cover the discrepancies that HS06 creates at this point, which seems to be a matter of interest in the US. Domenico concurred about the importance of pushing for this.

Xiaofei thanked us for running the workshop and was keen on being able to compare HEPscore results to know how good a result actually is. Domenico replied that he could do that following Randy and Tristan's presentation if he was referring to the score as such, whereas if he meant it in terms of efficiency it was Vincenzo's presentation that should be followed instead.

HEPscore

At the GBD on 12 Oct the summary of the HEPscore workshop will be presented (slides). The recommendation document should be ready on time for the WLCG workshop (Nov 7-11), so by mid-December a decision can be taken by the Management Board.

Helge suggested presenting the status also at the HEPiX workshop as well, which was effectively carried out. Regarding his impressions of the workloads, he preferred a HEPscore of 6 workloads rather than 11 since it was proven the difference made was little. The same thing applies to workload weights.

Alice reco should be finished ASAP so as to be included and be part of HEPscore 7 (i.e 7 workloads):

- Atlas
 - gen_sherpa
 - reco
- CMS
 - gen_sim
 - reco
- Belle2
 - gen_sim_reco
- LHCb
 - sim
- Alice
 - reco

Domenico moved on to say that we should pick the reference machine from the top 5 scoring machines that should be available at CERN and for the experiments. Randy pointed out that we're going to need the reference machine very soon, in case the one picked is not present already at CERN.

With regards to the scaling between HS06 and HEPscore, he proposed using the same absolute scale for HEPscore and HS06 on the reference machine, especially to make the migration of the accounting easier. Helge has concerns about this and proposed to discuss it after this meeting further, as more people like him had similar concerns during the meeting.

HEP Workloads

The workloads will improve with time, but there is no need to wait longer. ARM and POWER are gaining strength, but they won't be addressed by HEPscore 7. Despite some workloads already support ARM, none of the Experiments claimed a need to have HEPscore on ARM. A similar comment was made on GPUs.

HEP Benchmark Suite

The ease of use of the suite was highlighted by the sites, which considered it easier to run than HS06. Moreover, the central DB where the data is stored was seen as a centralized solution for long-term retention and analysis of benchmarking results. This vision is the purpose the DB wanted to fulfill so this is great feedback. It is hence reasonable to ask WLCG to run future benchmarks using the suite.

HEPscore Deployment

The provisioning deadline already mentioned in the meeting was brought up. In order to make it to the next batch, the adoption deadline is July 2023.

As for the accounting, the evolution of the tools was addressed, with a prototype expected in a few months. A transition period where both HS06 and HEPscore will live together is expected. Domenico remarked how accounting was perceived as a critical point for the HEPscore adoption.

His view on the outcome from the round table was that we should not re-benchmark old machines. However, a more controversial point was whether the reference machine should report the same value for HS06 and for HEPscore, by multiplying the latter by a conversion factor equal to its HS06 value. This would make accounting plots look the same, but make the values reported by HS06 and HEPscore similar, making them hard to identify without the units.

Helge agreed that we should leave old capacity alone and there should be a conversion factor for the only purpose that we can compare the machines benchmarked before the transition. For new machines, the real HEPscore value should be reported instead.

Beyond x86 CPUs

Not covered since it was covered in the initial round table and the meeting was running late.

Any other business

Domenico suggested meeting now every week rather than every two weeks, given the amount of work and upcoming deadlines, which was approved.

Next meeting

05 October 2022

Annex: Attendance

Present:

Domenico Giordano (CERN)
Helge Meinhard (CERN, chair)
Andrew Melo (Vanderbilt U)
Gonzalo Menendez Borge (CERN; notes)
Gonzalo Merino (PIC)
Bernd Panzer-Steindel (CERN)
Stefano Piano (INFN Trieste)
Matthias Schnepf (KIT)
Oxana Smirnova (U Lund)
Randall Sobie (U Victoria)
Tony Wong (BNL)

Apologies: