Gonzalo Menéndez Borge (CERN)

WLCG HEP-SCORE Deployment Task Force

Meeting on 18 Jan 2023 at 14:00 h UTC (teleconference)
Notes

Indico event page: https://indico.cern.ch/event/1209420/

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

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

Domenico mentions that the abstract submitted for CHEP23 has been approved as oral contribution and that candidatures for the speaker can be sent to Domenico and Randy. Michel enquired who is going to be at CHEP and therefore possibly give the talk "HEPscore: a new benchmark for WLCG compute resources". Randy will be present. Other TF members will be present too.

Summary of the Dec WLCG MB report (Domenico Giordano)

Domenico summarizes the outcome of the HEPscore deployment plan presented at the December WLCG MB meeting. For details see contribution in the agenda.
Alastair asked if sites that are procuring next-generation CPUs are meant to benchmark them with HS23. He wanted to know at which point HS23 would definitely be the benchmark to use instead of HS06 and others.
Domenico replied that for procurement actions that happen before April 1st, HS06 can be still used, given that HS23 is still not released. Year 2023 is a transition year in which sites will gradually move to HS23 after 1 April, to benchmark new capacity.

HEPscore test configurations (Domenico Giordano)

For details see contribution in agenda.

- Reference server will use HT ON for calibration since it makes the discrepancy between HS06 32-bits and HS23 smaller.

- The inclusion of Alice digi-reco introduces higher load and larger spread, but still within acceptable thresholds. Still, validation is needed, specially on servers with a very large # of cores.

No comments on Domenico's contribution.

HEP-Workloads status (Andrea Valassi, Randall Sobie, Stefano Piano)

LHCb workload (Andrea Valassi)

- New LHCB sim run3 workload more representative of Run 3 than the previous one, though similar. Released for x86, but could be replaced soon by another considering the following activities:

    - Port to ARM ongoing, still requires validation.

    - Suggestionts from V. Innocente to speed up the workload are being implemented. WIP.

Domenico wanted Andrea to clarify whether the part requiring still physics validation was only the ARM one, not x86, which Andrea confirmed. They also agreed on the fact that the physics validation does not need to be exhaustive for now, so long as it's reasonably accurate in terms of what it reconstructs.

Belle2 workload (Randall Sobie)

Work done by Tristan Sullivan, in collaboration with the Belle2 software group. Software has been ported successfully on ARM using the VM at CERN. The Belle2 software group finds the results good enough for benchmark. There are no plans to run physics validation. Effort is now on building the container that should be finished next week.

Domenico pointed out how things are looking bright about HEPscore 23 covering ARM.

Alice workload (Stefano Piano)

- Port to ARM started in December; it run into some issues, but they were fixed. A stable multiarch setup is available since last week.
- New apptainer images for both archs are available. They have been tested and proven to run fine. The load has been reduced to a maximum of 4 for 4 cores, to meet HEPscore's running conditions.
- Digi-reco will be chosen over reco alone, since it's been proven to be more representative.
- As future work, a workload that runs both on cpu and gpu is in the to-do list, but will still take some time.

Domenico remarked that targetting GPUs is indeed a very interesting approach, even though it won't be included in HS23. He also concurred on using digi-reco since he has also seen empirically how the results were more accurate and had less spread.

Lastly, he called for help again from those who have servers with a high number of cores available (> 64, ideally 256). Andrew volunteered to pledge his servers for that purpose.

Matthias mentioned he had run it on some 256-core machines and it failed. It seems to work fine with a lower number of copies, but starts to struggle with a higher number of them. E.g. 3 per core is fine, 10 breaks it.

Any other business

The fact that the next meeting is the last one before the deadline for the workloads was pointed out by Domenico, urging the experiments to wrap them up and report on them during the next meeting.

Next meeting

01 February 2023

Annex: Attendance

Present:

Domenico Giordano (CERN, chair)
Alastair Dewhurst (STFC UKRI)
Andrew Melo (Vanderbilt U)
Andrea Valassi (CERN)
Giuseppe Andronico (INFN)
Gonzalo Menendez Borge (CERN; notes)
Matthias Schnepf (KIT)
Michel Jouvin (IJCLab)
Randall Sobie (U Victoria)
Stefano Piano (INFN Trieste)
Tony Wong (BNL)

Apologies:

Helge Meinhard (CERN)