# WLCG HEP-SCORE Deployment Task Force

Meeting on 02 June 2021 at 14:00 h UTC (teleconference)
Notes

Indico event page: https://indico.cern.ch/event/1030673/

## Welcome, note-taking, notes from previous meeting

The minutes from the previous meeting are approved. Thanks to Manfred Alef for preparing them.

## Vectorisation in benchmarking workloads (Andrea Valassi)

Only few HEP software workload exploit vectorisation today, so little discussion in the benchmarking WG so far. While the hardware might support SIMD, the workflow needs to support it as well (compute the same function on multiple data at the same time) and the code needs to exploit it properly. Event generation looks like the most promising workflow, as the most CPU consuming part can be vectorised in a very relatively simple way (~x3 speedup in AVX2). Note that the compiler flags for vectorization are architecture-specific: an option is to have separate builds for different architectures (used by most experiments), otherwise one can have fat binaries supporting multiple SIMD levels. HEP workloads today do not use SIMD builds. In future we could have problem of definition in HEP-score due to possibly different degrees of freedom (same as for GPUs) depending on the level of vectorisation of different workflows. But today in WLCG, HEP software is largely not vectorised. The proposal is to use the workflows that are there now for the definition of a single HEP-score benchmark (so little/no effect of SIMD). ATLAS commented that there is little use of vectorisation in their software (confirming A. Valassi's statements). The situation is similar for CMS. General agreement to include a workflow heavily using vectorisation (potentially not from HEP) in the benchmark matrix (while not using it to define HEP-score now) to demonstrate the potential gains.

## Compilers and flags in benchmarking workloads (Helge Meinhard)

Inspired by a presentation in the GDB ~18 months ago. Workloads in HEP still compiled with -O2 to be conservative (facilitate physics validation) and inclusive, but at the cost of performance as does not leverage new features. SIMD is one aspect but there are other compiler features that can provide a gain. The proposal made at the GDB was to have a systematic (and recurrent) activity investigating the advantages of using new compiler features at the cost of not being able to continue running on old resources that do not support them. Questions to address include which workflows to focus on, what is the impact on physics validation and what is the impact on pledges and resource requests. There were positive statements after the GDB discussion, but the activity was not kickstarted because of lack of time and effort. Can we merge this effort now with the work in the HEP-SCORE task force (also in light of the previous presentation on SIMD)? One opinion is that the experiments are already considering the "best" compiler options for a given application, taking into account different effects including reproducibility. At the same time, there is an interest in exploring potential gains of features currently not in use, again, not entering in the HEP-SCORE definition at least today.

*Comment by T Boccali received after the meeting: It is not exactly true that CMS uses -O2; in mission critical parts of code, we use -Ofast (tracking is one example).*

Concluding the previous two discussions: we should use in HEP-SCORE the workloads with the set of options as are in use by the experiments "today". At the same time we should continue investigating various options for more performance to give the right incentives and include them in the suite (while not in the SCORE)

There was a discussion about the C-RSG process and how the performance (HS06/event) is monitored and scrutinised as incentive to the experiments. The process was explained, with some of the caveats.

## Updates on workloads and server configurations

Invitation to upload updated information to the agenda of the meeting.

## Any other business

The next meeting date (June 16) will be taken as the freeze for workflows

### Next meeting

Already scheduled for 16 June.

# Annex: Attendance

## Present:

Manfred Alef (KIT)
Simone Campana (CERN, notes)
Helge Meinhard (CERN, chair)
Josh Willis (Caltech)
Xiaofei Yan (IHEP)
Oxana Smirnova (U Lund)
Randall Sobie (U Victoria)
Domenico Giordano (CERN)
Jeff Templon (Nikhef)
Stefano Piano (INFN Trieste)
Walter Lampl (U Arizona)
Tony Wong (BNL)
Andrea Valassi (CERN)
Andrew Melo (Vanderbilt U)
Gonzalo Merino (PIC)
Michel Jouvin (IJCLab)
Ian Collier (STFC-RAL)

## Apologies:

Bernd Panzer-Steindel (CERN)
Tommaso Boccali (INFN Pisa)