

WLCG HEP-SCORE Deployment Task Force

Meeting on 16 June 2021 at 14:00 h UTC (teleconference)

Notes

Indico event page: <https://indico.cern.ch/event/1030676/>

Welcome, note-taking, notes from previous meeting

The minutes from the [previous meeting](#) are approved, with a comment by Tommaso Boccali: “it is not exactly true CMS uses -O2; in mission critical parts of code, we use -Ofast (tracking is one example)”.

Thanks to Simone Campana for preparing the minutes.

Workload updates

Meeting today is the point in time when we should declare the workflows frozen.

Round table of workloads:

LHCb (Andrea Valassi): waiting for input from LHCb; no real time scale but by the end of the months it **could** be ready.

ALICE (Stefano Piano): a standalone container with all the software needed has been built. These are typical O2 workflows, based on physics use cases. The second workflow is optimized for 8 threads running (2 GB / thread) via an internal scheduling mechanism. A todo list for Run3 is presented, for example to move to the 4 threads configuration needed for the benchmark. ALICE could even add an analysis workflow. How to access the calibration DB is still a problem, and that needed if we want to be completely standalone and not use network connections (it is accessed via HTTP). **Domenico:** the need for a local calibration DB is an important aspect; for a first phase in which we test on WLCG sites, it is not a showstopper. But it is on the longer run when we will want to use HEP-SCORE on other platforms. **Helge:** are measurements per step interesting and collected? **Stefano:** yes, but we will have a final overall score. **Domenico:** is the graph pipeliner (the internal scheduler) running Simulation and Digitization (for example) at the same time? **Stefano:** yes. **Domenico:** it is a little atypical and different from the standard way to run HEP-SCORE workloads, so from our point of view it will be single workflow doing all the steps. **Helge:** timeline for the analysis workflow? **Stefano:** 1-2 months. **Helge:** given the timeline, we should not count on this analysis workflow; it is possibly fully IO bound and checks other aspects of a single machine than the CPU (which is not bad per se, but “different”).

GW (Josh Willis): a containerized benchmark is available, 4 threads. Need to prepare the parsing logic needed to fill the JSON.

ATLAS (Walter Lampl): some progresses have been made using Geant4MT; could be ready “pretty soon”. Reconstruction MT, if used now, would be just a snapshot of the development process. Main problem is to create the snapshot of the DB needed to run unconnected. If we need to use the external connection to DB, we should at the very least avoid considering the first event where most conditions are loaded. **Domenico:** can you still use the DB from local sqlite file? **Walter:** yes, and it is on CVMFS, so it is for current workflows already on the node. **Helge:** can we close by the end of the month? **Walter:** probably.

CMS (Tommaso Boccali): all done, waiting for integration; Domenico is in contact with the person doing it. Andrea Valassi will also help centrally.

BELLEII (Randall Sobie): looking into some new data collected on the CERN testbed, and we see some deviations on new Intel machines. About containerization, it is done (Domenico confirms, it is in the HepScore20 proof of concept).

IHEP (Yan Xiaofei): now building the container, results seem to be good; timescale is ~ mid July.

Non HEP workflows (Domenico Giordano): not much about them (HS06, SpecInt etc). **Jeff:** not known if the solution by Domenico has been tried or not. In one machine classes (EPYC) there have been stability problems (BIOS settings??). **Domenico:** with these tests, results are accessible via Kibana and Elastic Search. **Michel:** on publishing the results on the web at CERN; can instructions on how to publish be circulated since they seem to involve e-groups and certificates? **Domenico:** if you have a valid X509 certificate, it should be ok.

Any other business

Helge: a proposal to change the Zoom room, and use one via Indico, has been received. Helge will take care.

Domenico: will be on holidays from the next meeting.

Yan: 2 workloads will be ready by the beginning of July.

Next meeting

Already scheduled for 07 July

Annex: Attendance

Present:

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

Apologies: