

WLCG HEP-SCORE Deployment Task Force

Meeting on 24 Aug 2022 at 15:00 h UTC (teleconference)

Notes

Indico event page: <https://indico.cern.ch/event/1067909/>

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

The minutes from the previous meeting are approved. Thanks to Tony Wong for preparing them.

Progress on workloads (Domenico Giordano)

Scheduled on Sept 19th and 20th, with TF members + personnel from experiments (Comp Coord + experts) + WLCG people.

The room is not enormous, hence please state asap (upon registration) whether you will attend in person or virtually.

The agenda is somehow atypical, starting in the morning and finishing early. This is mostly in order to avoid other meetings / workshops / weeks already scheduled.

On Monday, 3 sessions: TF review, analysis of workloads and finally the suite.

On Tuesday, 2 sessions: strategy to deploy hepscore and developments beyond x86.

In order to have uniform reports from experiments, templates will be provided to speakers. Same for site reports, a template is provided.

Questions:

Andrea Valassi: there are clashes with scrutiny group meetings + CMS Week + other stuff.

Yan Xiaofei: will attend the WS remotely.

Helge Meinhard: would like to congratulate Domenico and Randy for the very good agenda, and its preparedness.

Tony Wang: is there a registration page? YES – we need to make sure people see this asap → <https://indico.cern.ch/event/1170924/>

Tommaso Boccali: about starting to use HepScore, should we propose a date or wait for the bottom-up discussion?

Domenico: we should start having an idea during the first day + the survey results; not to be prepared NOW. We do NOT need to have a definite answer + date by the end of the workshop; that is the task of the task force, to be endorsed by the WLCG MB eventually.

Randy Sobie: the final solution could end up being more political than technical. Also having an accounting system compatible with HepScore is needed before any transition.

Status of the analysis (Randall Sobie)

Randy shows an updated picture of the benchmark analysis. Today is mostly on the finalization of the measurements.

A total of 2540 measurements was done, on 34 systems. There is a good coverage, and a lot of statistics.

Each workload was compared with 4 HS/SP benchmarks. CMS for example seems to scale with HS06 at most with a 5% difference, while LHCb closer to 40%.

Other studies show the effect of hyperthreading being on/off.

The proposed mixture of workloads deviates by less than 10% from HS06.

Having a technical proposal for a workload mixture by the WS is realistic, given that the relative weight of workloads is not too critical.

Andrea Valassi: we should thank Randy for this deep analysis. The main outcome of a LHCb meeting was interesting: while CMS and ATLAS scale very similarly, LHCb does not and they think this is due to the RICH which has a particular geometry with mirrors, which might lead to the difference. Bottom line: LHCb is 20% off the others for a good physics reason, and its appearing in HepScore is thus good since it stresses a different physics regime.

Tommaso Boccali: is there evidence that newer CPUs are deviating more wrt HS06 – since seeing deviations was one of the reason to abandon HS06? Randy: we can produce the plot, and in the past we did something similar. Helge concurs it would be a nice plot to see (deviation as a function of the CPU introduction date). Domenico thinks we can see this at the next task force meeting (Sept 7th).

Any other business

Next meeting

Next meeting will be Sept 7th. We can cancel the meeting after the WS; nobody seems to disagree.

Annex: Attendance

Present:

Tommaso Boccali (INFN Pisa)
Domenico Giordano (CERN; notes)
Helge Meinhard (CERN, chair)
Gonzalo Menendez Borge (CERN)
Stefano Piano (INFN Trieste)
Matthias Schnepf (KIT)
Randall Sobie (U Victoria)
Andrea Valassi (CERN)
Tony Wong (BNL)

Apologies:

Andrew Melo (Vanderbilt U)