

HEPscore Workshop Summary

The workshop in numbers	1
Timeline for a recommendation to WLCG MB	1
Draft Executive Summary	2

NB: This draft is still under review within the TF

The workshop in numbers

1. Indico agenda: <https://indico.cern.ch/event/1170924/>
2. Participants
 - a. Registered: 51
 - b. Attended: ~20 from zoom, ~17 in the room 513/1-024
 - c. Among others the Computing Coordinators of Alice, Atlas, CMS, LHCb; sw coordinator of Juno
3. 5 sessions to address different, interconnected, topics: HEPscore, HEP Workloads, HEP Suite, Deployment, beyond x86
4. Reports from the Experiments and (some) sites representatives included answers to the guidelines' questions
 - a. Experiments: <http://cern.ch/go/zgb9>
 - b. Sites: <http://cern.ch/go/6xHC>
5. Notes via cernbox/codiMD ([link](#))

Timeline to get to a recommendation for the WLCG MB

1. GDB Oct. 12: Workshop report
2. HEPiX Workshop: Oct 31-Nov 3
3. WLCG Workshop Nov 7-11: Wider discussion. Draft of Recommendation document
4. Mid-Dec: submit to the MB

Draft Executive Summary

1. HEPscore:

- a. The data show that differences between an HEPscore_11¹ and HEPscore_6² are of few percents
 - i. **Consensus** in preferring a smaller workload set: shorter runtime, higher sensitivity to performance changes of servers
- b. Data show that differences between weighted workloads in HEPscore and unweighted is also minimal
 - i. **Consensus** in preferring the simplest approach of unweighted WLS
 - ii. Avoids the need to re-assess the relative weights in the future
- c. HEPscore composition:
 - i. It is important to include the **Alice Reco** workload which consists of the reconstruction of Pb-Pb events (other LHC experiments are based on proton-proton collisions). It implies very high combinatorial reco challenge
 1. Alice Reco should be prepared and tested before the end of the year, to enter in the HEPscore_7 configuration.
 2. HEPscore_7 is based on HEPscore_6 with the addition of the alice_reco workload
- d. The reference machine should have the following characteristics:
 - i. To be one of the top5 recent architectures.
 - ii. Located at CERN.
 - iii. Preserved for the years of validity of HEPscore (specify timescale)
 - iv. To be available for the Experiments for the exercise of resource estimation
- e. HEPscore absolute scale (also see section 4e). The current scale of HEPscore is based on the geometric mean of the benchmarks of the workloads normalized to the values obtained on the current reference machine.
 - i. Normalization value under discussion.

¹ HEPscore_11 is composed of the 11 experiment workloads: alice_gen_sim, belle2_gen_sim_reco, atlas_gen_sim, atlas_sim_mt, atlas_reco_mt, cms_reco, cms_digi, cms_gen_sim, juno_gen_sim, igwn_pw (gravity wave), and lhcb_gen_sim

² HEPscore_6 is a subset of workloads: atlas_gen_sherpa, atlas_reco, cms_gen_sim, cms_reco, belle2_gen_sim_reco and lhcb_gen_sim

Option a) Use same absolute scale for HEPscore and HS06 (on the reference machine)

1. *“Avoid the impression of obfuscation”*
2. Objection that this could cause confusion between HS06 and HEPscore. But the confusion will be easily avoided if sites benchmark using the HEP Benchmark Suite (see later)

Option b) use a normalization that makes the scale different w.r.t. HS06

1. Avoid confusion between HS06 and HEPscore values
2. Implies that the accounting portal should include this normalization number in order to rescale the plots to the same value (see 4e)

2. HEP Workloads

- a. Workloads may improve in the coming years. But there is no reason to wait for the newer ones to include in HEPscore
- b. The use of ARM and Power is increasing, but not all experiments have workloads for these architectures.
 - i. Therefore **HEPscore_7 won't benchmark ARM and Power CPUs**
- c. GPU adoption is also too early in offline computing activity. Few proof-of-concept workloads are available. Alice Reco natively supports GPUs.

3. HEP Benchmark Suite

- a. Run procedure (single script): positive feedback from site reports. Easier to run than HS06
- b. The **Benchmark DB** (ElasticSearch) at CERN is seen as a centralized solution for **long-term retention** and analysis of benchmark results
 - i. *“Published data is useful to sites : processor powers, memory population...that's a net improvement compared to former Hepix website publication”*
 - ii. NB: this Benchmark DB is not supposed to be used for Accounting purposes. It will report the performance of individual CPU models benchmarked at sites.
- c. Sites should run HEPscore (using the Suite) to benchmark new hardware. Using the Suite will upload the benchmark results and metadata to a DB instance of Elastic Search at CERN.
- d. Implications:

- i. ActiveMQ at CERN needs to improve the credentials validation, allowing ID tokens
- ii. Guarantee support and maintenance of the ELK infrastructure (ElasticSearch, Kibana): backup, data migration, etc

4. HEPscore deployment

- a. **Timeline** for the adoption of HEPscore is mainly driven by the cycles of submissions of resource requests: Sept YYYY for the YYYY+2
 - i. Therefore for HEPscore to be used (in parallel to or mixed with HS06) for requests in 2023 for the year 2025, the transition needs to be done before July 2023. This is the time needed for the request preparation.
- b. **Accounting tools** (APEL, UAO): evolution described. Timeline for a prototype: a few months, as detailed here (Adrian Coveney)
 - i. *September/October Repository development to support change in record format*
 - ii. *Prototype Repository ready for testing in October – check message flows, record loading and aggregation with test site(s)*
 - iii. *November/December make changes as necessary following testing and start on Repository-WAU interface if that is required*
- c. Accounting Strategy presented. Proposal from the Accounting TF:
 - i. Transition period with sites with double accounting reports, clients/portals to support both reports.
 - ii. Need to follow up the migration process. EGI and WLCG coordination. The whole procedure may take a few months.
- d. During the round table it emerges the consideration that this procedure could take more months than expected, as it was the case for previous migrations.
- e. The round table proposes a different transition scenario with
 - i. No re-benchmark of old hardware for accounting purposes. On the contrary the **very same value** reported in HS06 will be reported in HEPscore units
 - 1. for this the same absolute scale of HEPscore is needed (see 1e)
 - ii. Only new CPU models to be benchmarked with HEPscore
 - 1. Sites could still re-benchmark old hardware for their interest, but there is not a centralized campaign
 - iii. We remarked that this was also the approach followed for the transition from SPEC Int to HS06

5. Beyond x86 CPUs

- a. Power Consumption measurement: several, similar initiatives ongoing. See opportunities for collaborations and exchange of experience. The HEPiX Benchmarking WG is the forum where this exchange could happen.
- b. GPU Workloads
 - i. 3 GPU workloads have been introduced. They represent a playground to demonstrate how to benchmark CPU+GPU systems. Still a lot of development needs to be done on the experiments' side to consider these application for offline production

Draft