The journey towards HEPScore, the HEP-specific CPU benchmark for WLCG

Domenico Giordano

CERN, Geneva, Switzerland
on behalf of
the HEPiX Benchmarking working group and
the HEPScore deployment task force

E-mail: domenico.giordano@cern.ch

Abstract. HEPScore is a CPU benchmark, based on HEP applications, that the HEPiX Benchmarking working group is proposing as a replacement of the currently used HEP-SPEC06 benchmark, adopted in WLCG for procurement, computing resource pledges and performance studies. In 2019, we presented at ACAT the motivations for building a benchmark for the HEP community based on HEP applications. The process from the conception to the implementation and validation of this objective has been inspiring and challenging. In the spirit of the HEP community, it has involved many contributions from software developers, data analysts, experts of the experiments, representatives of several WLCG computing centres, as well as the WLCG HEPScore Deployment Task Force. In this contribution, we review this long journey, the technological solutions selected, the readiness of HEPScore, and the deployment plans for 2023.

1. Introduction

The Worldwide LHC Computing Grid (WLCG) [1] collaboration is constantly looking for ways to improve and increase the efficiency of the computing resources used by the High Energy Physics (HEP) community. To achieve this goal, the community needs benchmarking tools that accurately reflect the computing needs of HEP experiments. The HEP community has been using for many years the HEP-SPEC06 (HS06) [2] benchmark to evaluate the performance of computing resources. However, as the computing landscape in HEP has changed over the past decade, the community has recognized the limitations of HS06 and the need for a more comprehensive and up-to-date benchmark. To this end, the HEPiX Benchmarking working group [3] has proposed HEPScore as a replacement for HS06.

2. Phases of the HEP-Benchmark project

HEPScore is a CPU benchmark based on HEP applications and has been conceived to innovate the approach to resource profiling for HEP. The journey to HEPscore has been a long process involving many members of the high energy physics community. It began in 2017, when the potential for building such a benchmark was presented at the WLCG workshop [4]. In the following years, the HEPiX Benchmarking working group has been at the forefront of the HEPScore development. First of all, the team confirmed the conceptual basis of the project, investigated its technical feasibility, and presented evidence of the differences between HS06 and

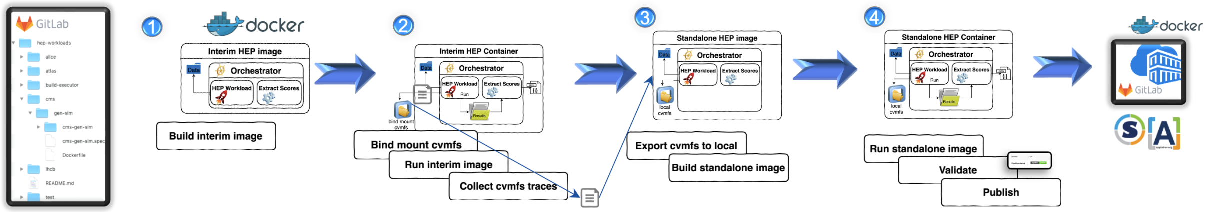

Figure 1: The Gitlab CI/CD infrastructure builds HEP Workload container images in multiple phases. The process starts with preparing the experiment-specific plugin that sets the runtime environment and configures the application, assuming the access of CVMFS is available. When the pipeline is triggered, a base image is built with the workload base code, the plugin code, and software packages needed by the plugin and absent in CVMFS (phase 1). In the second phase, CVMFS endpoints are automatically mounted to the container and the experiment's application is executed. During the execution, the accessed binaries are traced and then exported to the final standalone container that is built in the third phase. In the last phase, the standalone container is validated by executing it in dedicated GitLab runners, before publication in the GitLab repository. Docker [8] and Apptainer/Singularity [9] formats are made available.

HEP applications, using hardware performance monitoring tools that provide low-level insight into the CPU activity [5].

The second phase of the journey has been the definition of a software project, the HEP Benchmarks project [6], to develop all the components needed to achieve the prefixed goal. HEPScore is one of the pillars of the project, together with the fully containerized HEP applications, named HEP Workloads, and the tool developed for the submission and collection of the benchmark runs, named HEP Benchmark Suite. All components are released under the GPLv3 license.

The project has been developed with a focus on meeting the common requirements of benchmarking tools, including reproducibility, portability, and intuitiveness. The HEP Workloads are implemented as standalone container images to distribute software and data, addressing the intricacies posed by the large HEP applications: extensive codebases, necessary input data, configurations, and dedicated software environments. This approach offers a controlled environment for benchmark runs, a simplified deployment on diverse computing resources, and ultimately a smooth experience for the end users.

The HEP Benchmarks project infrastructure addresses also the challenge of collecting, maintaining, and extending workloads from multiple experiments by eliminating the need for *ad-hoc* recipes for each workload. The build infrastructure leverages GitLab CI/CD for a fully automated build of the containers' images on multiple architectures including x86, aarch64, GPUs, with potential support for IBM-Power (Fig. 1) The infrastructure has been designed to minimize the contribution of experiments' experts, that become responsible only for developing an experiment specific plugin. The plugin role is (i) to setup the experiment environment within the container relying on the accessibility of the experiment's software from CVMFS [7], (ii) to configure the application that processes the given input data, and (iii) to implement the report the performance metric, that generally is the event throughput. The execution monitoring, error logging and reporting components are included in the core-software package. To ensure a uniform approach across workloads, the project has established strict requirements for a common command line interface and a standardized report structure.

To date, multiple production workloads from 7 experiments have been successfully prepared through a collaboration between experts from the experiments and the HEPiX Benchmark working group.

3. HEPScore23

In 2020/21, a comprehensive validation of the HEPScore benchmark concept was conducted using the available workloads from the Run2 software version of various experiments [10]. The workloads covered the four typical phases of the processing pipeline in HEP: Monte Carlo generation of physics events, simulation of the particle propagation throughout a detector and electric signals production, digitization of the signals, and reconstruction. A demonstrator benchmark, HEPScore$_\beta$, was released and tested on 15 different Intel and AMD CPU models based on the x86 architecture. The results indicated that HEPScore had the potential to replace HS06 as a benchmark for CPUs.

Building upon the results of this study, the WLCG collaboration upgraded the project from an evaluation phase to a viable replacement for HS06. At the end of 2020, the WLCG HEPScore Deployment task force [11] was established to further support this effort with the objective of creating the migration plan from HS06 to HEPScore, determining the final composition of HEPScore, and onboarding a larger set of WLCG sites for validation purposes. In this respect, the role of the task force complemented that of the HEPiX Benchmarking working group.

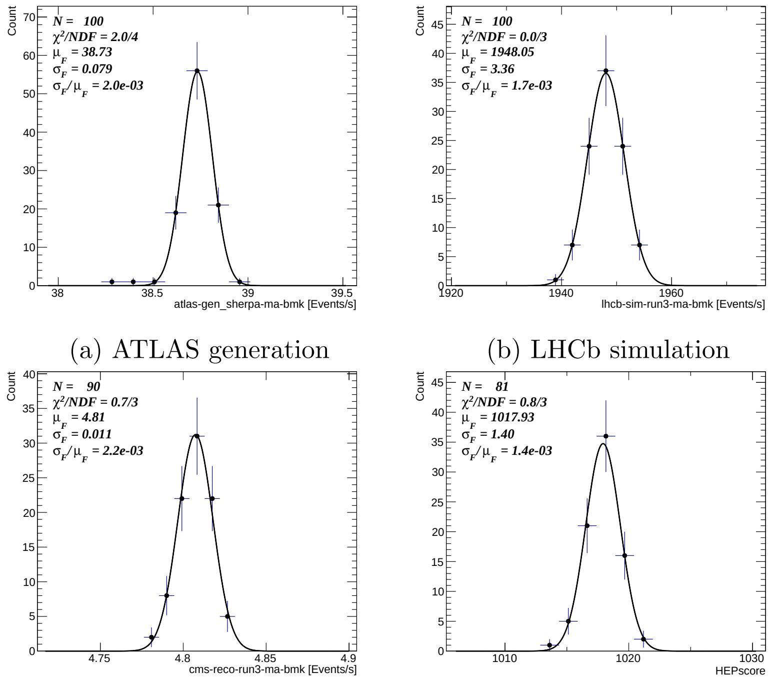

(a) ATLAS generation (b) LHCb simulation

(c) CMS reconstruction (d) HEPscore23

Figure 2: Histograms of the measurements of the events processed per second for three of the seven workloads composing HEPScore23 (a, b, c). The measurements of HEPscore23 are also reported (d). The measurements have been performed on the reference server Intel Gold 6326 CPU @ 2.90GHz with 256 GB of RAM and hyper-threading enabled. Note that each entry is the median of three sequential measurements of the workload.

In 2022, a significant step was taken towards the final composition of HEPScore. Eleven workloads among the most recent ones from the LHC experiments, Belle2 [12], Juno [13], and IGWN [14] were containerized. These workloads were then executed repeatedly on approximately 40 different CPU models from 15 WLCG sites, to assess the robustness of the workloads and their resolution on repeated measurements, proven to be at the level of per mill. A summary of the studied workloads is reported in Tab. 1 and Fig. 2. Additionally, the execution of the workloads via the HEP Benchmark Suite provided valuable operational experience and helped to expose the tool to a larger number of site managers and drive the benchmarking campaign and data collection. The collected data was also utilized to explore various combinations of HEPScore, such as including all workloads, excluding the long-running ones, and weighting each workload either equally or based on the Grid fraction of jobs. The analysis showed little difference in the ranking score for more than 30 CPU models, with only a few percent variation between the candidates (Fig. 3). Therefore, other requirements were considered in the selection process of the HEPScore composition, such as a small workload set, shorter runtime, unweighted scores.

HEP-workload	Threads	Runtime [min]	wl-score [events/s]	σ/μ $[10^{-3}]$
ALICE DIGI-RECO∗	4	16	0.76	7.5
ATLAS GEN-SHERPA∗	1	6	39	2.0
ATLAS SIM	4	13	0.34	2.3
ATLAS RECO∗	4	15	9.1	1.5
Belle2 GEN-SIM-RECO∗	1	5	15	0.9
CMS GEN-SIM∗	4	9	2.7	4.0
CMS DIGI	4	4	11	2.0
CMS RECO∗	4	12	4.8	2.2
IGWN PE	4	33	2310	4.7
JUNO GEN-SIM-RECO	1	14	3.9	9.1
LHCb SIM∗	1	10	1948	1.7

Table 1: Measurements of the event throughput and their resolution on repeated measurements for the workloads studied in 2022. The workloads composing HEPScore23 are flagged with a ∗. The measurements have been performed on the reference server Intel Gold 6326 CPU @ 2.90GHz with 256 GB of RAM and hyper-threading enabled. The event throughput of each given workload refers to what delivered by the server when running a number of parallel copies of the workload that would utilize all available server's cores.

The final composition of HEPScore was discussed during a 2-days workshop held at CERN [15], which brought together a diverse group of stakeholders, including representatives from experiments, sites, and the WLCG board. The agreed composition consists of seven workloads that mimic the processing pipelines of production jobs running on the WLCG infrastructure: a digitization and reconstruction workload from ALICE [16], a ATLAS [17] generation workload, a ATLAS reconstruction workload, a CMS [18] generation and simulation workload, a CMS reconstruction workload, a LHCb [19] simulation workload, and a Belle2 generation, simulation and reconstruction workload. This composition of HEPScore is identified as HEPScore23, being expected to enter in production during the year 2023. All seven workloads are based on the latest available versions of the experiments' software, and are available for both x86 and aarch64 architectures. This expansion creates new opportunities for studies that combine performance benchmarks and energy consumption benchmarks not only on x86-based CPU models, but also on aarch64-based CPU models, making HEPScore a valuable tool for the WLCG community [20].

The strategy for transitioning from HS06 to HEPScore23 was drafted during the HEPScore workshop and later consolidated in subsequent collaboration meetings such as the Grid Deployment Board [21], WLCG workshop, and HEPiX workshop [22] all in Q4 2022. This transition was designed to be as simple and straightforward as possible for WLCG site

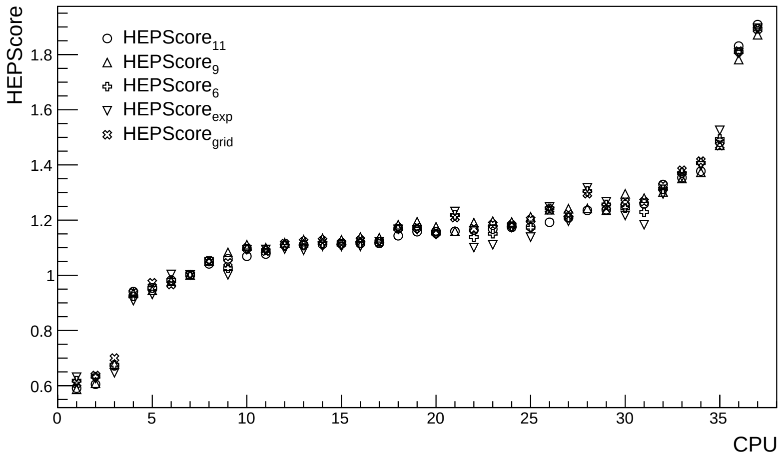

Figure 3: Benchmark score value measured on 37 different CPU models evaluated with five different candidate compositions of HEPScore: three unweighted averages of eleven, nine, six workloads respectively, and 2 weighted averages of six workloads (exp and grid). The value of the weights is determined equally weighting each experiment contribution (exp case) or weighting by the percentage of WLCG resources used to process a specific workload (Grid case).

administrators, minimizing changes to the accounting framework and reducing operational efforts for site managers. The transition from HS06 to HEPScore23 aims to be gradual, requiring that only new hardware acquired from 2023 onward is benchmarked with HEPScore23; the old hardware still deployed at sites is accounted in terms of HS06. This mixture of two benchmark values in the accounting reports is made possible because the same scale factor of HS06 is applied to HEPScore, fixed on a reference server [1], to ensure consistency in the scoring system. As the relation between HS06 and HEPScore is within 10% for the newest model available in the WLCG infrastructure, at the time of writing, the discrepancy is expected to be gradually absorbed with the rollout of new hardware at sites.

4. Conclusions

HEPScore is a valuable addition to the HEP community's tools for evaluating the performance of computing resources. The ability to execute a range of HEP applications make HEPScore a representative benchmark. The future inclusion of workloads running also on GPUs, including physics analysis and machine learning algorithms will make it an even more useful tool. The deployment of the configuration HEPScore23 in the WLCG infrastructure from 2023 onward will progressively substitute HS06.

References

[1] WLCG, The Worldwide LHC Computing Grid, `http://wlcg.web.cern.ch` Accessed 20 March 2023.
[2] Michelotto M et. al. (2010), A comparison of HEP code with SPEC1 benchmarks on multi-core worker nodes. J. Phys.: Conf. Ser. 219 052009. `https://doi.org/10.1088/1742-6596/219/5/052009`.
[3] HEPiX Benchmarking working group, `https://w3.hepix.org/benchmarking.html`. Accessed 20 March 2023.
[4] WLCG workshop 2017, `https://indico.cern.ch/event/609911/`. Accessed 20 March 2023.
[5] Giordano D and Santorinaiou E, Next Generation of HEP CPU benchmarks. J. Phys. Conf. Ser. 1525 012073 (2020) doi: https://doi.org/10.1088/1742-6596/1525/1/012073
[6] HEP Benchmarks Repository, `https://gitlab.cern.ch/hep-benchmarks` Accessed 20 March 2023.

[1] The reference server is equipped with a dual socket Intel Gold 6326 CPU @ 2.90GHz with 256 GB of RAM

[7] Blomer J et al (2017), New directions in the CernVM file system. J Phys Conf Ser 898:062031. https://doi.org/10.1088/1742-6596/898/6/062031
[8] Docker containers, https://www.docker.com/resources/what-container Accessed: 20 March 2023
[9] Singularity, https://sylabs.io/singularity/ Accessed: 20 March 2023
[10] Giordano, D., Alef, M., Atzori, L. et al., HEPiX Benchmarking Solution for WLCG Computing Resources. *Comput Softw Big Sci* **5**, 28 (2021). doi: https://doi.org/10.1007/s41781-021-00074-y
[11] HEPScore Deployment task force, https://indico.cern.ch/category/12951/ Accessed 20 March 2023.
[12] Kou E et. al. (2019), The Belle II Physics Book. Progress of Theoretical and Experimental Physics 12 123C01 https://doi.org/10.1093/ptep/ptz106
[13] Adam T et. al. (2015), JUNO Conceptual Design Report https://arxiv.org/abs/1508.07166
[14] Brady, P., Losurdo, G., Shinkai, H. (2021), LIGO, VIRGO, and KAGRA as the International Gravitational Wave Network. Handbook of Gravitational Wave Astronomy. Springer, Singapore. https://doi.org/10.1007/978-981-15-4702-7_51-1
[15] HEPScore Workshop, Sept. 19-20 2022 https://indico.cern.ch/event/1170924/ Accessed 20 March 2023.
[16] Aamodt, K. and others (2008), The ALICE experiment at the CERN LHC JINST 3 S08002 10.1088/1748-0221/3/08/S08002
[17] Aad G et. al. (2008), The ATLAS Experiment at the CERN Large Hadron Collider. JINST 3 S08003 https://doi.org/10.1088/1748-0221/3/08/S08003
[18] The CMS Collaboration (2008), The CMS experiment at the CERN LHC. JINST 3 S08004 https://doi.org/10.1088/1748-0221/3/08/S08004
[19] Alves AA et. al. (2008), The LHCb Detector at the LHC. JINST 3 S08005 https://doi.org/10.1088/1748-0221/3/08/S08005
[20] Simili E, Stewart G, Skipsey S, Spiteri D, Britton D (2023), Power Efficiency in HEP (x86 vs. arm) *to appear in this same issue*
[21] Grid Deployment Board, https://indico.cern.ch/category/6890/ Accessed 20 March 2023.
[22] HEPiX Autumn 2022 Workshop, https://indico.cern.ch/event/1200682/ Accessed 20 March 2023.