

# HEPscore23 status

D. Giordano (CERN/IT)

WLCG Operation Coordination

2 March 2023

# HEP Benchmarks project

*HEPscore* has been proposed by the HEPiX Benchmarking WG

- Uses the workloads of the HEP experiments
- Combine them in a single benchmark score

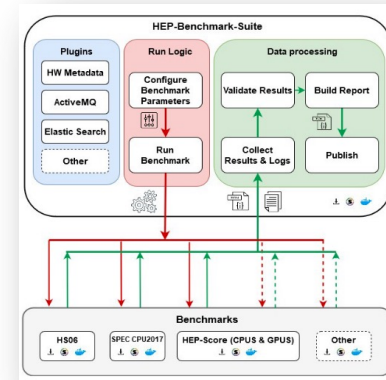
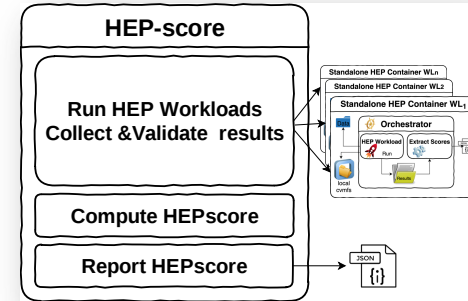
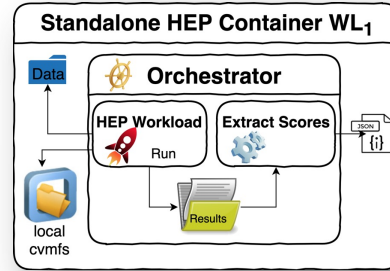
*HEPscore* relies on *HEP Workloads*

- Individual **reference** HEP applications

In addition, *HEP Benchmark Suite*

- Orchestrator of multiple benchmark (HEPscore, HS06, SPEC CPU2017)
- Central collection of benchmark results

All released under GPLv3 license





# Progress reports

In several forums since the beginning of the project. Constructive feedback received

Only in the Last months

- HEPscore Deployment Task force (2 meetings/month since 2020)
- HEPscore Workshop (19-20 Sept 2022)
- HEPiX Workshop (31 Oct - 03 Nov 2022)
- WLCG Workshop in Lancaster (7-11 Nov 2022)
- GDB (14 Dec 2022)

In the recent future

- HEPiX Workshop / GDB (March)
- CHEP23 (May)
- EGI (June)

# HEPscore23

We name HEPscore23 (HS23) the configuration of HEPscore to be adopted in production

- ❑ 7 workloads included, as agreed at the HEPscore workshop
- ❑ All workloads have the most recent version of the experiments' software
  - Support x86\_64 and aarch64
- ❑ Validation finalized

Exp	WL	x86_64 / aarch64
ALICE	digi-reco	✓
ATLAS	gen_sherpa	✓
	reco_mt	✓
Belle2	gen-sim-reco	✓
CMS	gen-sim	✓
	reco	✓
LHCb	sim	✓

# HEPscore23 configuration

## ❑ HEPscore v1.5 is the current beta version

- Includes a single WL set (default)
- 2 configuration files:
  1. Access SIF images from registry
  2. Access SIF unpacked images from CVMFS. Useful for runs in job slots or sites with cvmfs unpacked.cern.ch

## ❑ Custom configurations are still possible for other studies

- Could be included in the configuration folder if desired

Name
..
➔ hepscore-cvmfs.yaml
➔ hepscore-default.yaml
{-} hepscore23-cvmfs.yaml
{-} hepscore23.yaml

# Validation of HS23

- ❑ On a CERN server testbed
- ❑ Measurement reproducible (spread  $\lesssim 2\%$ )

CPU	Online CPUs	# reps	Count	Score ↓	spread
Neoverse-N1	0-159	3	12	2713	0.837%
Intel(R) Xeon(R) Gold 6326 CPU @ 2.90GHz	0-63	3	26	1019	0.445%
AMD EPYC 7302 16-Core Processor	0-63	3	27	982	0.887%
Intel(R) Xeon(R) Silver 4216 CPU @ 2.10GHz	0-63	3	17	715	2.11%
Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz	0-63	3	6	708	0.465%
Intel(R) Xeon(R) Gold 6130 CPU @ 2.10GHz	0-63	3	7	690	0.316%
Intel(R) Xeon(R) CPU E5-2680 v4 @ 2.40GHz	0-55	3	16	631	2.04%
Intel(R) Xeon(R) CPU E5-2650 v4 @ 2.20GHz	0-47	3	16	482	1.56%
Intel(R) Xeon(R) CPU E5-2630 v3 @ 2.40GHz	0-31	3	16	319	2.27%

# Plan to release

- ❑ Presented and approved at the MB of Dec22
- ❑ Defines the main M/D for the benchmark release the 1<sup>st</sup> of April
  - Starting date of the resource scrutiny year

## Milestones

- ❑ 1<sup>st</sup> April 2023 HEPscore23 in production
- ❑ 1<sup>st</sup> March 2023 HEPscore23 configuration frozen
  - Allows for 1 month for tests
- ❑ 14<sup>th</sup> Feb. 2023 HEP Workloads frozen
  - Allows for 2 weeks of tests/fixes
  - Latest date to have HEPscore23 for x86 and ARM
    - Otherwise ARM support will be added in a next version HEPscore2x, with x>3
  - In case a new workload does not pass the validation:
    - (a) the corresponding current one can be used; (b) exclude from HEPscore23





# Also at the Dec 22 MB

## Accounting perspective

- ❑ Migration procedure detailed by Julia in the last GDB ([talk](#))
- ❑ The 1:1 normalization of HS06 and HS23 simplifies the transition
  - Less changes to the code
- ❑ Sites will be expected to only benchmark new hardware with HEPscore. Old hardware does not need to be re-benchmarked
  - However, sites wishing to re-benchmark old hardware may do it (*outcome of WLCG workshop and the last GDB*)



# Next steps

Discussed yesterday at the  
HEPscore TF meeting

- ❑ The HEPscore TF (and HEPiX WG) are on track with the planned schedule
- ❑ 1 month to test and validate the configuration
  - Confirm the stability and usability of the benchmark
  - Add minor features to the HEPscore python code before release v2.0
- ❑ Resume the measurement procedure as done in the past
  - Reach an increasing number of WLCG sites, starting with the ones already involved in the past
  - Need to access as many CPU models as possible
    - Including ARM, and large core workloads

# Extend plan after 1<sup>st</sup> of April

- ❑ Document the procedure
  - Clarify operational questions
- ❑ Monitor the adoption by sites
- ❑ Support computing coordinators in the usage of the benchmark for future requests
  - Summer 2023 for 2025

Contribution needed from Accounting WG and Op Coordination

