

HEPScore23 status

D. Giordano (CERN/IT)

On behalf of HEPscore Deployment TF and HEPiX Benchmarking WG

GDB

12 April 2023

HEPScore23

We name HEP Score23 (HS23) the configuration of HEP Score adopted in production since April 1st

- ❑ 7 workloads included, as agreed at the HEP Score workshop
- ❑ All workloads have the most recent version of the experiments' software
 - Support x86_64 and aarch64
- ❑ Validation of the workloads finalized by end of Feb. '23

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

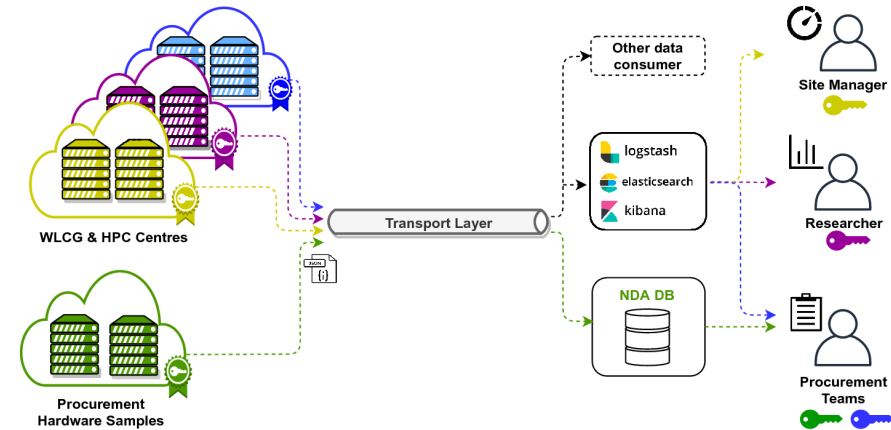
Distinguish configuration Vs code

- ❑ HEPScore v1.5 is the current beta version of the python code
 - Includes a single WL set (default – **HEPScore23**)
 - 2 configuration files available
 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
- Will be merged in a single configuration in the next HEPScore v2.0 release
- ❑ Code development of HEPScore is continuing to add minor features
 - No impact to the score computation

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

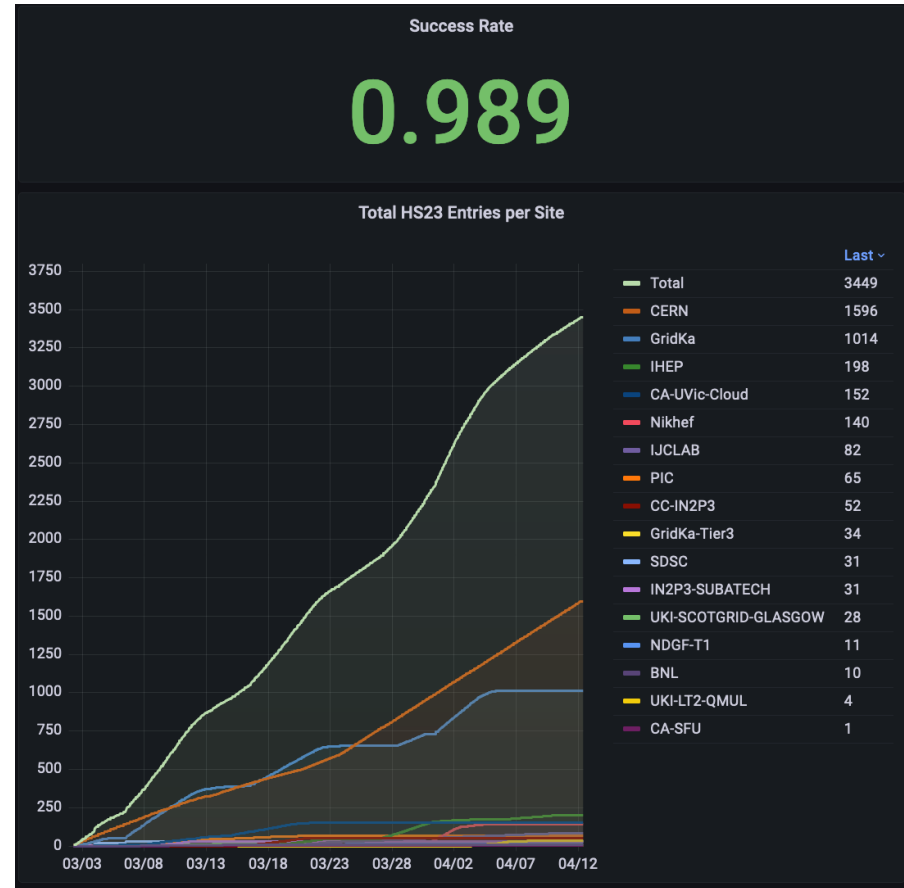
Running procedure

- ❑ Prepared documentation (see draft attached to the indico contribution)
 - Details the HW and SW requirements, the running procedure, the accounting process
- ❑ Preferred running approach
 - Use the Benchmark Suite to publish the results @CERN
 - Script made available to beta-tester sites (mainly from participant to previous measurements of the Task Force)



Data collected


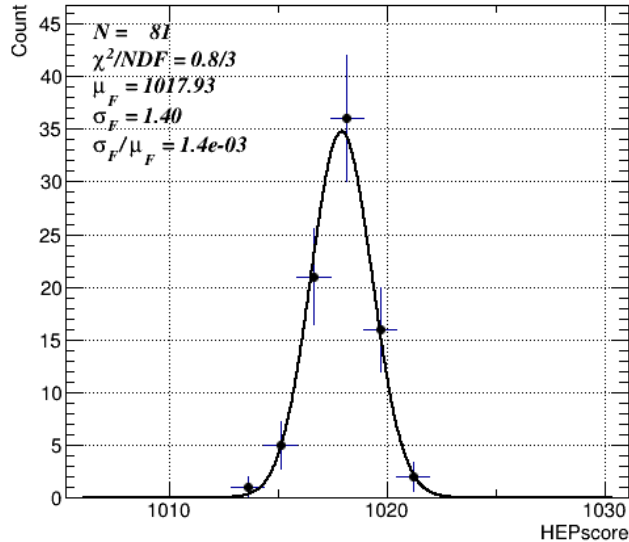
- ❑ 14 sites + CERN contributed so far
- ❑ ~2k measurements
(in addition to CERN)
- ❑ 46 distinct CPU models
(Intel, AMD, ARM)
- ❑ Success rate: >98%
 - Failures typically connected to cases not fulfilling the requirements



HS23 repeated measures

☐ Reproducible results


(spread $\lesssim 2\%$)



CPU	Online CPUs	Count	Score	spread
AMD EPYC 9654 96-Core Processor	0-383	26	6005	0.987%
AMD EPYC 9654 96-Core Processor	0-191	25	4957	0.454%
AMD EPYC 7773X 64-Core Processor	0-255	25	3388	0.461%
AMD EPYC 7773X 64-Core Processor	0-127	22	2974	0.992%
Neoverse-N1	0-159	12	2714	0.657%
AMD EPYC 7702 64-Core Processor	0-255	16	2691	1.46%
AMD EPYC 7702 64-Core Processor	0-255	20	2529	0.401%
AMD EPYC 7742 64-Core Processor	0-127	59	2500	1.50%

HS23 vs HEP-SPEC06


Clear trend: HS23 score is higher than HS06 for newer CPU models



Reference machine
(Gold 6326 @ 2.90GHz)
Fixed HS23==HS06

Performance as Events/kWh

The ability of HEP Score23 to run on x86 and ARM nodes has already enabled studies of energy consumption vs performance



On time with the schedule

As agreed at the MB of Dec22

Milestones

- ❑ 1st April 2023 HEPscore23 in production
- ❑ 1st March 2023 HEPscore23 configuration frozen
 - Allows for 1 month for tests
- ❑ 14th 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

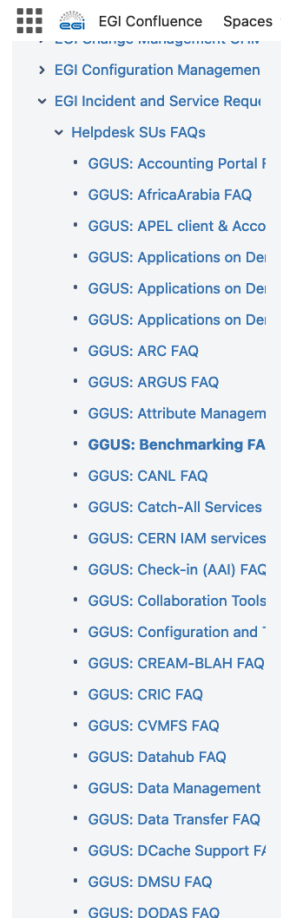


GGUS Support

For future support requests a GGUS support unit is available

(see [FAQ](#))

- Thanks to Maarten L. and A. Paolini (EGI)
- Redirects to CERN Service now



The screenshot shows a navigation menu for EGI Confluence. The menu is organized into a tree structure. The top level includes 'EGI Confluence' and 'Spaces'. Under 'EGI Confluence', there are several categories: 'EGI Configuration Management', 'EGI Incident and Service Request', and 'Helpdesk SUs FAQs'. The 'Helpdesk SUs FAQs' category is expanded, showing a list of links for various services and units, including 'GGUS: Accounting Portal', 'GGUS: Africa/Arabia FAQ', 'GGUS: APPEL client & Account', 'GGUS: Applications on Demand', 'GGUS: ARC FAQ', 'GGUS: ARGUS FAQ', 'GGUS: Attribute Management', 'GGUS: CANL FAQ', 'GGUS: Catch-All Services', 'GGUS: CERN IAM services', 'GGUS: Check-in (AAI) FAC', 'GGUS: Collaboration Tools', 'GGUS: Configuration and Management', 'GGUS: CREAM-BLAH FAQ', 'GGUS: CRIC FAQ', 'GGUS: CVMFS FAQ', 'GGUS: Datahub FAQ', 'GGUS: Data Management', 'GGUS: Data Transfer FAQ', 'GGUS: DCache Support FAQ', 'GGUS: DMSU FAQ', and 'GGUS: DODAS FAQ'. The 'GGUS: Benchmarking FAQ' link is highlighted in blue.

Pages / ... / Helpdesk SUs FAQs

GGUS: Benchmarking FAQ

Created by Alessandro Paolini, last modified by Maarten Litmaath 25 minutes ago

FAQ FOR Benchmarking SUPPORT UNIT

Responsible Unit: Benchmarking

Helpdesk: GGUS (Operations Helpdesk by TPM)

- What is the purpose of the Benchmarking Support?
- For which components does Benchmarking Support provide support?
- Which quality of service (QoS) will you provide?
- Who will assign tickets to Benchmarking Support?
- Are tickets typically solved in Benchmarking Support or reassigned elsewhere?
- Who is responsible for Benchmarking Support?
- What documentation is available on Benchmarking Support?
- What is the usual "Type of Issue" for the tickets of Benchmarking?
- Comments
- What if I have questions which are not dealt with by this FAQ?

What is the purpose of the Benchmarking Support?

To provide support to site administrators concerning the use of the new HEPscore23 benchmarking suite that will replace HEPspec06, as well as the implications for accounting.

For which components does Benchmarking Support provide support?

For the HEPscore23 benchmarking suite and related accounting configuration details.

Which quality of service (QoS) will you provide?


Base.

Who will assign tickets to Benchmarking Support?

Site admins.

Next steps

- ❑ Update the official documentation in the HEPiX [web page](#)
- ❑ Monitor the adoption by sites
 - Support the usage of the benchmark for future requests
 - Summer 2023 for 2025



Benchmarking Working Group

The Benchmarking WG is in charge of defining and maintaining a consistent and reproducible CPU benchmark to describe experiment requirements, lab commitments, existing compute resources, as well as procurements of new hardware.

HEP-SPEC06 (HS06)

HS06 is the HEP-wide benchmark for measuring CPU performance. It has been developed by the HEPiX Benchmarking Working Group in order to replace the outdated "kSI2k" metric.

The goal is to provide a consistent and reproducible CPU benchmark to describe experiment requirements, lab commitments, existing compute resources, as well as procurements of new hardware.

HS06 is based on the all_cpp benchmark subset (bset) of the widely used, industry standard SPEC@ CPU2006 benchmark suite. This bset matches the percentage of floating point operations which we have observed in batch jobs (~10%), and it scales perfectly with the experiment codes.

HS06 is the official CPU performance metric to be used by WLCG sites since 1 April 2009.

Although the HS06 benchmark was initially designed to meet the requirements of High Energy Physics (HEP) labs, it is by now widely used also by other communities.

Tables of HS06 results

Default system configurations

