Deployment HEPscore in production from the accounting perspective

Julia Andreeva (CERN)

15.03.2023

Migration to the new benchmark for the accounting workflow
What does it imply

- *Possibility to report job consumption using either HS06 or HEPscore benchmarks and to process it correctly in the APEL repository*
- *Show proper consumption in the EGI portal using HS06 units before the switch to HEPscore and in HEPscore units after the switch*
- *Show historical distribution in the EGI portal and WAU without disruption*
- *Possibility to follow up on the gradual increase of the resources benchmarked with HEPscore at the sites. This means a possibility to demonstrate which part of the consumed resources has been benchmarked with HEPscore vs HS06 in the EGI portal and/or WAU*
- *VO requirements and pledges in CRIC for 2023 will be in HEPscore and pledges and VO requirements for the previous years translated in HEPscore for the historical distributions*

Show proper consumption in the EGI portal using HS06 units before the switch to HEPscore and in HEPscore units after the switch

- Agreed with the EGI portal developers. Trivial change. Will be done shortly before the 1st of April

Show historical distribution in the EGI portal and WAU without disruption

- With normalisation of HEPSPEC to HEPscore 1 to 1 should work out of the box.

VO requirements and pledges in CRIC for 2023 will be in HEPscore and pledges and VO requirements for the previous years translated in HEPscore for the historical distributions

- CRIC team is already working on the changes.
- Will be deployed in production shortly before the 1st of April
- There are also some modifications required in the accounting report generation (to change units). This will be ready before the 1st of April, but deployed in production in May , when April accounting reports will be generated, since March reports still should be in HEPSPEC units

Possibility to report job consumption using either HS06 or HEPscore benchmarks and to process it correctly in the APEL repository

- It has two aspects : server side and client side

- Server side will be implemented in two steps:

  - Benchmark name tag is properly recorded in the DB (ready and deployed on the testing infrastructure).

  - Aggregation by a benchmark name and exposing aggregated data for export to the EGI portal repository (planned for the end of May)

- Client side

  - Depends on client used by the site

Testing infrastructure

- New specifications of the accounting records which include the name of the benchmark are ready

- Documentation for the sites joining testing and validation can be found on the [twiki page](#)

- The testing infrastructure capable of receiving accounting records in old and new formats and recording information in the DB with a proper benchmark name tag is ready

Client side, usecases (1)

- Sites which send normalized records: CERN, several T1 (IN2P3-CC, INFN-T1, NDGF), couple of european T2s and OSG

- New specification is published. All sites have been informed about the change and requested to join the tests. So far the successful tests have been performed by INFN-T1

- OSG is a special case. They do not need any changes at the sites. There should be some changes performed in the central repository. The required modifications have been discussed and agreed with GRACC developer.

Client side, usecases (2)

- Sites which send summaries or individual job records relying on the generic clients: APEL client (HTCondor European sites) or ARC client

- According to ARC developers there should not be any problem, but so far this has not been fully tested. We pinged ARC developers and some of the sites asking them to join the tests. First ARC site NCG-INGRID-PT started to send reports to the testing infrastructure today (to be still validated).

- New APEL client won't be ready before the 1st of April. However, KIT voluntered to patch the client and started to send data to the testing infrastructure today (to be validated). If it works, the patch will be documented and recommended for use for other sites in case they start to benchmark with HEPscore before new APEL client is released.

- The fact that APEL client won't be ready before the 1st of April is not a showstopper for the switch to a new benchmark. This just implies that probably for some fraction of sites which would start benchmark their resources with HEPscore, we won't be able to track it unless new APEL client is released and deployed, which should happen during April-May

Possibility to follow up on the gradual increase of the resources benchmarked with HEPscore at the sites. This means a possibility to demonstrate which part of the consumed resources has been benchmarked with HEPscore vs HS06 in the EGI portal and/or WAU

- This requires changes in APEL, EGI portal and WAU.
- The request for changes have been submitted to the EGI poartl developers, they started to look into it, but can not accomplish the task before APEL won't start reporting data considering the name of the benchmark used for a particular resource
- Similarly for WAU which exports data from the EGI portal
- The plan is to have everything in place by the end of May.
- Since data will be in APEL repository starting from the 1st of April, we would be able to follow gardual increase of the resources benchmarked with HEPscore starting from the 1st of April (apart of sites relying on APEL client)

Important note

- One of the frequently asked questions is how sites are supposed to introduce resources benchmarked with HEPscore

- The recommendation is to locate them behind a dedicated CE

- For the sites which prefer not to have a dedicated CE we have a following recommendation in the documentation:

  - In case CE provides access to a mixture of the resources benchmarked with two different benchmarks, the benchmarking factor should be calculated normalized by resource capacity of every resource set. In this case the accounting record should contain "HEPScore23" as a benchmark name, though technically resources are not homogeneous considering used benchmark.