



Deployment HEPscore in production from the accounting perspective

Julia Andreeva (CERN)

GDB 12.04.2023

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

- DONE**

Accounting Portal High-Throughput Compute Cloud Compute Disciplines Research Infrastructures Reports Restricted Views Help

The EGI Accounting Portal is an EGI service provided by . This work is co-funded by the EOSC-hub project (Horizon 2020) under Grant number 777536.

Tier 1 Grid Accounting

Metric: Metric Unit:

Start Time: End Time:

Row Variable: Column Variable:

VO Filtering: LHC TOP 10 ALL EGI Official Custom VO Selection

Job submission type: Global Jobs Only All Jobs Local Jobs Only

This view shows the accounting data from all Grid Sites that are classified as WLCG Tier1 . Only production certified Sites are included. The metric shown is Total number of jobs, grouped by Tier 1 Node and Month shown. No local jobs are shown

Tier1 — Total number of jobs by Tier 1 Node and Month (LHC VO)


Tier 1 Node	Apr 2022	May 2022	Jun 2022	Jul 2022	Aug 2022	Sep 2022	Oct 2022	Nov 2022
CA-TRIUMF	540,051	871,412	823,182	692,855	482,752	355,490	559,025	614,387
CH-CERN	5,468,837	5,262,579	5,579,382	5,535,857	5,225,278	5,718,585	5,706,438	4,775,188
DE-KIT	1,578,093	1,654,618	1,392,788	1,303,411	1,367,109	1,315,441	1,420,838	1,329,129
ES-PIC	259,743	243,211	212,163	240,287	298,539	304,022	288,313	253,037
FR-CCHN2P3	1,033,742	1,108,485	1,149,977	1,064,949	1,108,860	1,022,334	1,203,701	1,157,916
IT-INFN-CNAF	724,220	729,581	1,106,260	871,421	881,104	763,346	841,036	783,287
KR-KISTI-GSDC	103,630	90,224	212,871	227,574	246,489	228,066	267,027	253,249
NDGF	472,554	725,646	775,862	1,063,908	1,116,225	1,023,085	1,021,218	822,172
NL-T1	581,503	715,392	673,003	661,377	697,854	579,349	690,671	816,150

Show historical distribution in the EGI portal and WAU without disruption

- DONE
- With normalisation of HEPSPec to HEPscore 1 to 1 was trivial

[Download JSUN Data](#) / [Download GSV Data](#)

The information in the previous table is also shown in the following graph.



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

- DONE
- Has been deployed in production shortly before the 1st of April
- There are also some modifications required in the accounting report generation (to change units). Code is ready but will be deployed in production in May after generation of April accounting reports, since March reports still will 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 by the end of March.
 - 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 kindly asked 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
- ARC client has been tested by couple of pioneer sites, some issues have been detected. The patched version is ready, will be released in production this month.
- New APEL client was not ready before the 1st of April. However, KIT has patched APEL client and joined the tests.
- New APEL client should be ready by the end of this month, tests still might take some time.
- The fact that APEL client was not ready before the 1st of April is not a showstopper for the switch to a new benchmark. This just implies that 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, 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 portal developers, they started to look into it, but can not accomplish the task before APEL starts 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.

Important note

- One of the frequently asked questions is how sites are supposed to introduce resources benchmarked with HEPscore
- In case, when the new resources are located behind a dedicated CE, the solution is simple, since benchmark info (benchmark name and benchmark factor) is defined per 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. Example for calculation is included in the [draft of documentation](#)

Practical instructions for the sites

- The sites which will start to benchmark part of their resources with Hepscore23 first are supposed to test the flow using APEL test infrastructure. Instructions can be found on the [twiki page](#)
- In parallel we prepared a [document](#) which contains detailed instructions for sites (how to run benchmark and what should be done for accounting). The document is being improved following site feedback.
- GGUS/SNOW support has been enabled to report issues or ask questions
- We would like to thank INFN-T1, FZK, PIC and NCG-INGRID-PT teams for active participation in testing and debugging and encourage other sites to join the tests ASAP.

QUESTIONS?