



Deployment HEPscore from the accounting perspective

Julia Andreeva (CERN)
Adrian Coveney (STFC UKRI)

GDB 14.12.2022

Migration to the new benchmark for the accounting workflow

Main conclusions from the WS in Lancaster

Emerging consensus from the Benchmarking Task Force (and the experiments/sites) is that site should not re-benchmark existing hardware with HEPScore

This means that we will have sets of resources

benchmarked with two different benchmarks for a long while

Outcome of the discussion at the WLCG WS in Lancaster is that HEPscore will be normalised to HS06 with factor 1

This implies minimum changes in the accounting workflow, though some modifications are still required

The switch to HEPscore in the official accounting reports is planned for the 1st of April when we switch to the new yearly pledges

This means that some of the required modifications should be performed and tested before the 1st of April 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*

APEL client side. Existing workflow

- Most sites include benchmarks as part of the APEL accounting records (job or summary)
- Some sites only provide final usage × benchmark values in “normalised summaries”
- Benchmark name is defined in the APEL local configuration. For the time being APEL client expects certain values as a benchmark name and it does not include HEPscore
- Sites apply a scaling factor between execute nodes , so that they can use a single benchmark per cluster

Required changes. APEL and CEs

- **APEL client side.**
 - *For all APEL reports the name of the benchmark should be defined along with the actual benchmarking value. Such possibility already exists in APEL, though since only a single benchmark was used at a time , benchmark name was not considered for APEL processing*
 - *APEL repository should be able to deal with the reports corresponding to two different benchmarks. This implies that the schema and APEL aggregation procedures have to be changed to take into account one more meta attribute (benchmark name)*
- **ARC and HTCondor CEs**
 - *ARC and HTCondor CEs should be enabled to deal with two benchmarks for APEL reporting (one at a time)*
- **OSG accounting**
 - *OSG accounting GRACC sends normalised records to APEL, changes on their side should be performed in order to include correct name of the benchmark in the record*
- **Sites (like CERN) which send normalised records**
 - *Such sites would need to perform necessary modifications and report the name of the benchmark*

EGI portal and WAU

- EGI portal UI should change units at the point of the switch
- Collector from the APEL repository , DB schema, aggregation procedures and UI should be modified to properly take into account the name of the benchmark
- Since HEPscore is normalised to HS06 with factor 1 no recalculation of the historical data is required. All mentioned points are also relevant to WAU
- Good news, CESGA has subcontracted Quanta Labs company. Starting from November to the end of June there will be effort to work on the EGI portal

CRIC and accounting reports

- Since HEPscore is normalised to HS06 with factor 1 no recalculation of history is required. Therefore, changes are minimal and mostly cosmetic, like change of the units in the accounting reports.

When changes should be implemented

- **Mandatory changes for the 1st of April**
 - APEL should be modified in order to ensure that the name of the benchmark could be properly reported with the accounting records from the site to APEL repository, and benchmark name should be properly recorded in the repository. More details will be provided by Adrian in the second part of presentation
 - CRIC should be modified for the switch for the 1st of April to show proper units in the accounting reports
- **At the point when the site would like to benchmark resources with HEPscore it would need to perform the following changes**
 - If it uses individual or summary records, it would need to upgrade APEL client in case APEL client is used to generate accounting records. We are checking with CE developers whether any changes would be required on the ARC or HTCondor side
 - If the site uses normalised records (this also relates to the OSG GRACC accounting), the site needs to perform changes in order to properly report benchmark name in the normalised record. Slightly modified specification will be provided by Adrian
- **Changes which are not mandatory though desirable by the 1st of April**
 - APEL aggregation procedures should be changed to enable aggregation by the benchmark name.
 - EGI and WAU should be modified to enable distributions showing which part of the resources is benchmarked with a given benchmark

Reporting consumption using several benchmarks

- Described scenarios are minimalistic ones
- In a longer term (after 1st of April 2023), we might need more flexibility in the accounting workflow, which allows to report and process consumption in several benchmarks in parallel.
- The specification for usage record has been developed. See backup slides
- There is no consensus whether we need to enable this feature and would be good to understand whether this work is required

Backup slide. Draft of the specification to report consumption using two benchmarks (individual record).

This draft was created when we thought that there is no common conversion factor which will be used across the infrastructure and therefore two benchmarks had to be reported.

```
APEL-individual-job-message: v0.4
Site: SOME-SITE
SubmitHost: host.ac.uk/cluster
LocalJobId: 9aef372d-e26f-42ce-7acb-5e1c479dc47f
LocalUserId: bob
GlobalUserName: /DC=ac/DC=uni/DC=/DC=vac
FQAN: /host.org/Role=NULL/Capability=NULL
WallDuration: 47248
CpuDuration: 46871
Processors: 1
InfrastructureDescription: APEL-CREAM-HTCONDOR
InfrastructureType: grid
StartTime: 1531869580
EndTime: 1623693622
ServiceLevel: {hepspec: 11.4, HEPscore22: 0.153}
```

Backup slide. Draft of the specification to report consumption using two benchmarks (summary record)

```
APEL-summary-job-message: v0.4
Site: SOME-SITE
SubmitHost: host.ac.uk/cluster
Month: 9
Year: 2022
GlobalUserName: /DC=ac/DC=uni/DC=/DC=vac
WallDuration: 47248
CpuDuration: 46871
Processors: 1
NumberofJobs: 3
InfrastructureType: grid
EarliestStartTime: 1531869580
LatestEndTime: 1623693622
ServiceLevel: {hepspec: 11.4, HEPscore22: 0.153}
```