



Deployment HEPscore from the accounting perspective

Julia Andreeva (CERN)

19.10.2022

Migration to the new benchmark. Some considerations (1)

- We need to support transition from HEP-SPEC 06 to HEPscore
- Transition to HEP-SPEC 06 was more straightforward since a single conversion factor could be applied for all sites. Therefore from the accounting point of view the transition could be done in one go

Initial scenario

- Sites gradually benchmark their resources (existing at the site and purchased during transition period) in HEPscore, as soon as it is done they report consumption in both benchmarks
- WLCG ops follows on the progress. As soon as we see that the vast majority of sites do enable HEPscore consumption in their reports we switch official accounting resources from HS06 to HEPscore.
- In order to preserve accounting history, we re-calculate HS06 consumption into HEPscore applying to HS06 historical data conversion factor calculated per site.

What came out at the WS (1)

- Apparently sites **do not intend to re-benchmark** old resources with HEPscore and do not plan to benchmark new ones with HS06.
- In difference to earlier plan, there **will be a common conversion factor** applied for translation between HS06 and HEPscore benchmarks

What is currently reported to APEL

- Main metrics used for CPU accounting:
 - wallclock and CPU **time**
 - wallclock and CPU **work** (HS06 consumption , represents time multiplied by HS06 benchmarking factor)
- APEL accounting record, contains wall clock and CPU time and a benchmarking factor which APEL uses to translate time to work
- We need to understand how we deal with reporting of the benchmarking factors when two benchmarks coexist

Possible solution (1). On the side of the site

- **We can not sum up oranges with apples!**
- Keep in mind that during the transition period HS06 will stay the official benchmark used for the accounting reports. So we need to be able to express overall site work both in HS06 and HEPscore
- Old and new resources at the site **should be organised in separate clusters.**
- For the old clusters HS06 is a basic benchmark
- For the new ones HEPscore is a basic benchmark
- For old resources there won't be re-benchmark and there will be a common conversion factor which can be applied to calculate HEPscore consumption out of HS06 one.
- Similarly, for the new resources in order to keep HEPscore consumption in the same scale as HS06 one, a common conversion factor will be introduced for the whole infrastructure.
- This implies, that we can go with a single benchmark reported for a given cluster, and the second benchmark can be derived by the accounting system from the one coming with the report.

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}
```

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}
```


What does it imply in practice

- Site has old cluster which is benchmarked with HS06 with value 11.4. The report related to this cluster should contain:
 - `ServiceLevel : {hepspec: 11.4}`.
- Probably APEL client can complement reported info using common conversion factor, let's say 0.01 and the accounting record sent to message queue will contain:
 - `ServiceLevel : {hepspec: 11.4, HEPscore22: 0.114 }`
- Site got new resources and benchmarked them with HEPscore, benchmarking factor is 0.175.
 - `ServiceLevel : {HEPscore22: 0.175 }`
- Probably APEL client can complement reported info using common conversion factor, let's say 100, and the accounting record sent to message queue will contain:
 - `ServiceLevel : {hepspec: 17.5, HEPscore22: 0.175 }`
- This is just one possible scenario which should be confirmed with APEL developers. Conversion can happen later in the chain , for example in the repository.

Possible solution(2). On the side of accounting tools

- During the migration period HS06 stays an official benchmark and we generate accounting reports using it.
- We still collect second set of metrics based on HEPscore.
- Parallel use of two benchmarks might require update of APEL local components at the site.
- We monitor the situation and at the point when vast majority of sites publish accounting information in the agreed form, we start generating two sets of accounting reports using two benchmarks. This lasts for 3 months, people have an opportunity to check both reports and conclude whether they make sense. These 3 months are considered to be a validation period before the final switch. HS06 still stays official.
- After 3 months of validation, accounting reports based on HEPscore become official (point X). We can still generate reports based on HS06 just for crosschecking for few more months.
- We recalculate old HS06 monthly statistics, using common conversion factor till point X, then starting from point X we complement this new metrics with HEPscore consumption. This new metric becomes the official one used both for historical distributions and for the updates.

Questions?