



Deployment HEPscore from the accounting perspective

Julia Andreeva (CERN)
Adrian Coveney (STFC UKRI)

WLCG WS, 09.11.2022

Migration to the new benchmark for the accounting workflow

What we need to consider

From the presentation of Randall:

Emerging consensus from the 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

Should HEPScore be normalised to HS06? (No consensus in the task force)

Answer to this question has a lot of implications on the accounting. Would be great if we could conclude at the WS.

Initial scenario (before benchmarking WS in September 2022)

- Sites gradually benchmark their resources (existing at the site and purchased during transition period) in HEPscore and 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 reports from HS06 to HEPscore.
- For the accounting history, we re-calculate HS06 consumption into HEPscore applying conversion factor calculated per site.
- **This scenario is not considered any more. Main change is that existing resources won't be re-benchmarked with HEPscore.**

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*

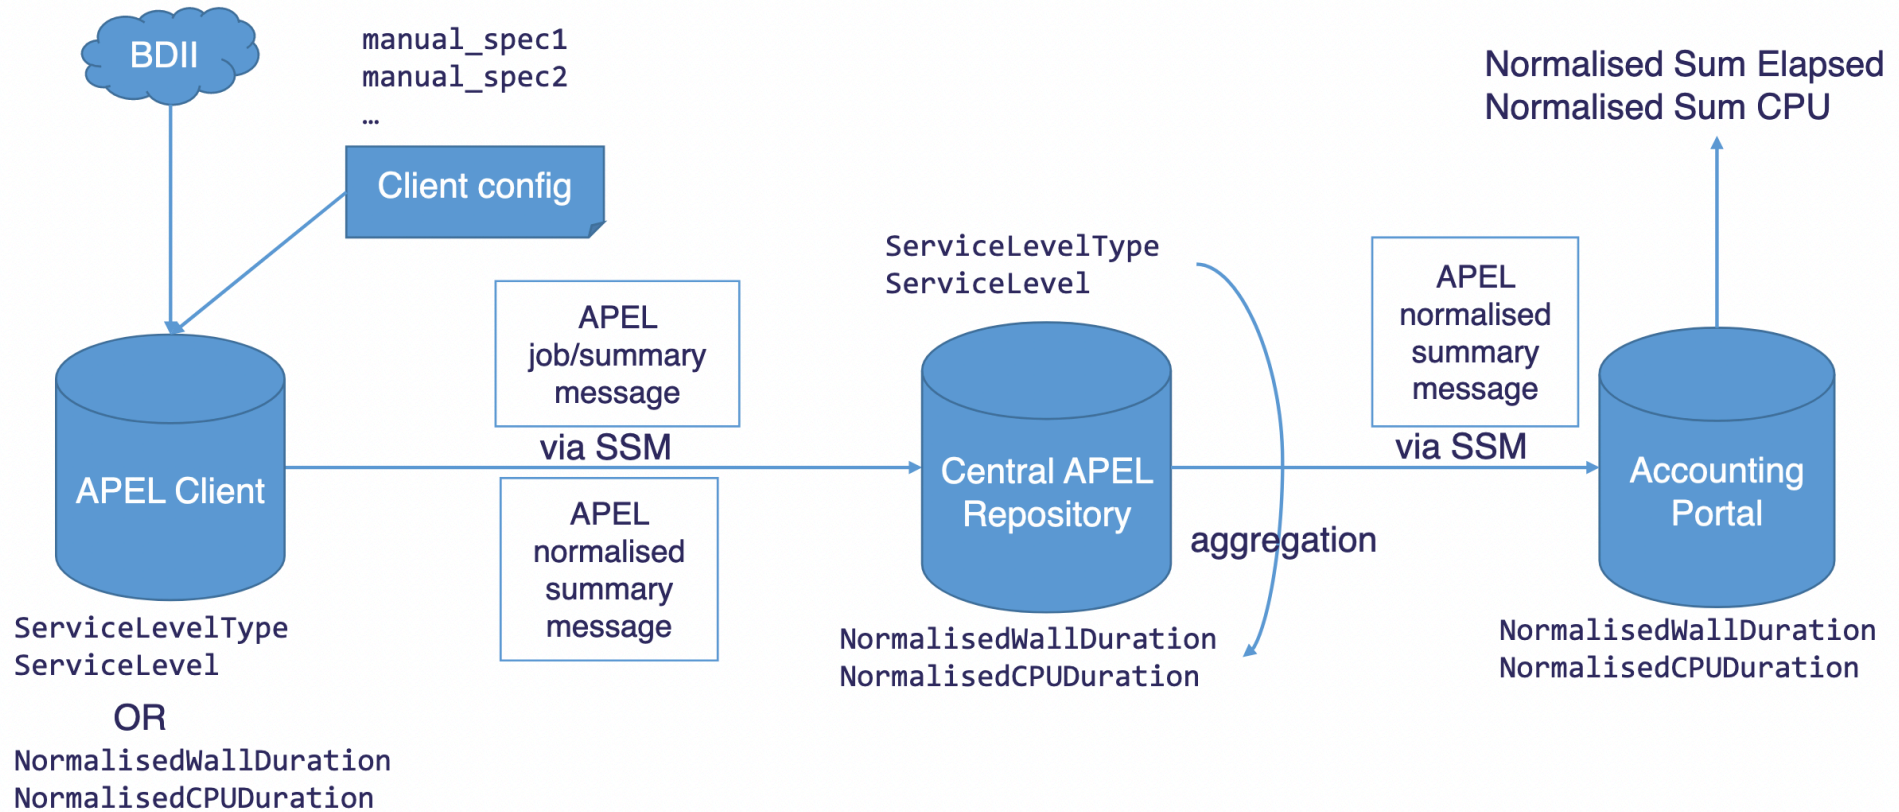
Scenario when existing resources are not re-benchmarked and HEPscore is normalized to HS06 would minimise changes in the accounting workflow, though modifications will still be required.

APEL client side. Existing workflow

- Most sites include benchmarks as part of the APEL accounting records (job or summary)
- Many sites only provide final usage \times benchmark values in “normalised summaries”
- Benchmarks can be pulled in from BDIs or set manually locally to the accounting client
- Sites apply a scaling factor between execute nodes , so that they can use a single benchmark per cluster

APEL client side. Existing workflow

GLUE_CE_CAPABILITY
GLUE_HOST_BENCHMARK



Required changes. APEL and CEs

If we want to follow which fraction of the resources has been benchmarked with HEPscore or HEPscore is not normalised to HS06

- **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)*

If we do not care which fraction of the resources has been benchmarked with HEPscore and HEPscore normalised to HS06 these changes are not required

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 be able show which fraction of the consumed resources have been benchmarked with HEPscore (if this feature is desired)
- If HEPscore is normalised to HS06 no recalculation of the historical data is required. Otherwise such recalculation should be performed using conversion factor between HS06 and HEPscore.
- 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

- If HEPscore is not normalised to HS06 recalculation of history is required, this includes consumption, pledges and VO requirements.
- Otherwise changes are minimal and mostly cosmetic

So how we go

If we want to follow which fraction of the WLCG resources has been benchmarked with HEPscore or HEPscore is not normalised to HS06

- **Changes described above should be implemented in APEL, EGI portal, WAU and CRIC**
- **It should be ready by the 1st of APRIL if this is the time of the switch**
- **Sites which report ‘normalised summaries’ should include benchmark in their report. This change is mandatory, when these summaries contain resources benchmarked with HEPscore, otherwise HS06 is a default.**

OSG sends to APEL summaries and does not use APEL infrastructure as such. However, depending on the selected scenario changes might be also needed for OSG accounting.

Reporting consumption using several benchmarks

- Described scenarios are minimalistic ones
- In a longer term (after 1st of April 2023), we certainly need more flexibility in the accounting workflow, which allows to report and process consumption in several benchmarks in parallel.
- We intend to continue work in this direction, though the time line for implementation and deployment is longer than 1st of April next year
- The specification for usage record has been developed. See backup slides

Discussion

- Do we decide to normalise HEPscore to HS06 (factor 1)?
- Do we need to follow which fraction of our resources is benchmarked with HEPscore?
- In a medium term, do we put effort in making APEL accounting flow more flexible, supporting reporting and processing of consumption in several benchmarks in parallel?

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}
```