Benchmarking Changes and Accounting

John Gordon, STFC-RAL

WLCG Accounting Task Force, 9th March 2017

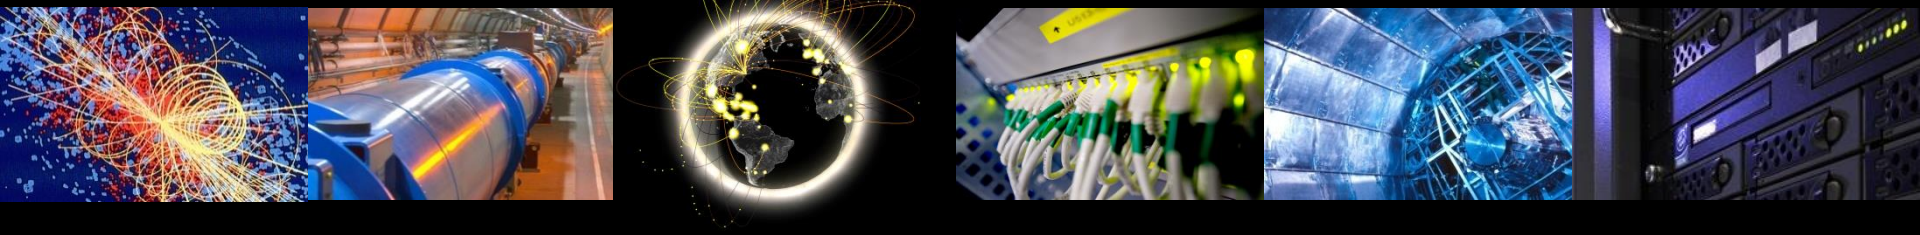

Outline

- Issues concerning any change of benchmark
  - Fast or slow
- Fast Benchmarking inside jobs
- Summary Publishing

Change of Benchmark

- Apel supports multiple benchmarks but currently only one per job record.
- **ServiceLevelDescription**: the flavour of benchmark (text string)
  ServiceLevel: numerical value of benchmark for the core the job ran on (float)
- Apel relies on getting the appropriate benchmark value from somewhere in order to add it to the usage record.
  - For the Apel client this is currently the BDII but could be some other place (CRIC, GOCDB, other database)
- During a migration from one benchmark to another APEL needs to make an assumption about the relative values of the two benchmarks (i.e have a conversion factor) so that it can integrate data using old and new benchmarks into a single view in order to compare sites/pledges etc.
  - E.g. Tier2 Report
  - Typically one might view the data in old units until sufficient sites had migrated and then use the new units with conversion from old for sites not yet migrated.
- In addition the portal could display data normalised to both benchmarks as the old portal did for SI2K and HS06 but for each *site*date* there would only actually be one measurement . The other value would be from conversion.

Fast Benchmarking

- Apel requires access to benchmark information to create Usage Records.
  - Not typically available in batch logs
  - It does this after the job has finished so it doesn't have any access to any data inside the job unless it is stored somewhere persistently.
- A fast benchmark run inside each job would need the benchmark measurements to be stored in a form that it could be retrieved later by the apel client and identified with the job or the WN with a relevant timestamp so that it could be matched to the conditions at the time the job ran.
- Matching to the job might be feasible.
  Matching to the WN, probably not as with multiple jobs running on a WN there might be a variety of benchmark results over different time windows.
- If fast benchmarks run in each job were averaged and collected in a database somewhere then APEL could query this provided it has sufficient information to identify what to ask for.
  - e.g a site average would be simple, for different node types APEL would need to know about different node flavours.

Summary Publishing

- OSG, CERN, and some other sites publish normalised summary data to APEL.
- There is nothing in the summary record to identify the benchmark used to normalize. It is currently HS06 by convention only.
- Managing a migration would be difficult.
- Extending the SUR to include the benchmark type would allow APEL to do conversion as described previously

Summary

- An alternative benchmark evaluated as at present is feasible.
- Needs an agreed conversion factor during migration
- Benchmarking in every job is more challenging.
- Normalised Summaries will need a change to label them.

Unpledged Resources

- APEL is a flat database of sites.
- Other metadata like WLCG Tier, EGI NGI, country, is applied at the portal.