Fast Benchmarking

Q/A (1)

- Whether your experiment would be interested that opportunistic resources are accounted in APEL?
  - ATLAS is interested, but it is not going to be easy.
- If yes, what are possible scenarios?
  - We are still discussing this, as there are different types of resources and some aren't benchmarked yet.
- Whether these opportunistic resources are already accounted in the experiment-specific systems?
  - Wallclock is recorded for every type of resource, the work depends on a benchmark value existing. The benchmark value exists only for owned/accessible resources.

Q/A (2)

- How/whether benchmarking of such resources performed?
  - HPC/cloud/boinc not benchmarked yet. Plan is to run the short benchmark (DB12) in the pilots but this doesn't work on all the resources. Cloud resources in Canada are using the CERN-benchmark suit already asynchronously. HPC and boinc are still under discussion. We don't have yet a way to store the results which is also still under discussion.
- How these resources are described regarding topology?
  - It depends on the resources. Over pledged resources installed in the grid way are declared in the usual way. Other resources like HPC/clouds and boinc aren't because they can come and go. This was discussed at length in the Information System TF. Many of these resources don't have explicitely declared end points and sometimes don't even have PandaQueues as they get dynamically added to existing pledged resources.

Q/A (3)

- Would it be possible to retrieve accounting data for the opportunistic resources from the experiment-specific systems via APIs?
  - At the moment the accounting is in the dashboard so it can be retrieved with a number of formats. I think kibana offers similar ways of getting the raw data of a selection (though not as many formats. We should check with the monitoring group).