

Information System

– BDII current system not needed and OTT but there are still some dependencies. Most people agree WLCG should become independent from it, some sites are worried about smaller VOs but EGI will still support the BDII and they can simplify their requirements too (they have already started discussions about this). WLCG should work on moving the last parameters somewhere else. The point of this is to reduce the number of parameters that requires publishing and give clear definition for those. Work has already started to use only GOCDB for experiments targeted service discovery for example.

– WLCG IS: do we need it? Each experiments has already their IS with specific parameters. ATLAS and WLCG ops may agree on AGIS, CMS showed some interest, LHCb only if proved reliable. Everybody agrees it is paramount to have reliable information, so the new info system has to guarantee that otherwise not worth the effort to move.

– Glue2.x maintained for compatibility reasons. Can be used in a simplified form as a dictionary more than a fully blown schema with constrictions.

– Benchmark publishing currently done over the BDII which imposes constraints due to the multiple use cases it has to satisfy. Clearly this needs to be separated. In particular for experiments that want to use it for payload brokering (currently only LHCb but ATLAS keeping an eye on things) MJF is a more precise solution which removes the uncertainty due to the averages. Numbers used for accounting currently still in the BDII could be moved somewhere else GOCDB/OIM good candidates.

Benchmark

– HS06 v fast benchmark. We agreed we still need HS06 (or replacement when available) for procurement of hardware and accounting of standard sites. Fast benchmark can be used to do pilot and job brokering and for procurements where there is no control over the hardware (i.e. commercial clouds). The two different benchmarks respond to different use cases.

– Concerning HS06 some of the discrepancies with experiments software maybe due to compiler flags which don't activate boosting features on new hardware. It should be investigated further.

– Concerning the short benchmarks, there are already several. We should avoid a proliferation and aim at adopting one solution (maybe a bundle if doesn't become too big) for all the experiments that need it. Sites should be able to run it and replicate the tests. Fast benchmarks are still a less precise alternative than HS06 on a per worker node basis (also because HS06 is calculated in the worst case scenario).

– In cooperation with IS and MJF TFs agree on definitions and publishing instructions for sites. This another step to reduce uncertainty on the numbers. Ongoing.

Accounting

– Cloud accounting infrastructure already there for EGI and will continue to be developed and improved. It is not used by WLCG though, because WLCG so far have used only virtualised WNs. Not clear anything else is required ATM.

– Wallclock vs cputime there was no time to discuss, but we agreed we will produce a review document for the discussion at the MB.

– Commercial clouds and HPC resources accounting not clear how and what we want to publish ATM that would mean publishing fast benchmarks for these resources (ATLAS is working in this direction). Some sites in the US (FNAL) have started to do some work to extract information from real jobs run on commercial clouds which is an alternative. To be discussed further.