Contribution ID: **204** Type: **Oral**

Using HEP experiment workflows for the benchmarking and accounting of computing resources

Monday 4 November 2019 11:30 (15 minutes)

The benchmarking and accounting of CPU resources in WLCG has been based on the HEP-SPEC06 (HS06) suite for over a decade. HS06 is stable, accurate and reproducible, but it is an old benchmark and it is becoming clear that its performance and that of typical HEP applications have started to diverge. After evaluating several alternatives for the replacement of HS06, the HEPIX benchmarking WG has chosen to focus on the development of a HEP-specific suite based on actual software workloads of the LHC experiments, rather than on a standard industrial benchmark like the new SPEC CPU 2017 suite.

This presentation will describe the motivation and implementation of this new benchmark suite, which is based on container technologies to ensure portability and reproducibility. This approach is designed to provide a better correlation between the new benchmark and the actual production workloads of the experiments. It also offers the possibility to separately explore and describe the independent architectural features of different computing resource types, which is expected to be increasingly important with the growing heterogeneity of the HEP computing landscape. In particular, an overview of the initial developments to address the benchmarking of non-traditional computing resources such as HPCs and GPUs will also be provided.

Consider for promotion

No

Primary authors: VALASSI, Andrea (CERN); GIORDANO, Domenico (CERN); ALEF, Manfred (Karlsruhe Institute of Technology (KIT)); MICHELOTTO, Michele (Università e INFN, Padova (IT))

Presenter: VALASSI, Andrea (CERN)

Session Classification: Track 7 –Facilities, Clouds and Containers

Track Classification: Track 7 –Facilities, Clouds and Containers