Contribution ID: **290** Type: **Poster**

Performance evaluation of distributed file systems for the phase-II upgrade of the ATLAS experiment at CERN

Over the next few years, the LHC will prepare for the upcoming High-Luminosity upgrade in which it is expected to deliver ten times more p-p collisions. This will create a harsher radiation environment and higher detector occupancy. In this context, the ATLAS experiment, one of the general purpose experiments at the LHC, plans substantial upgrades to the detectors and to the trigger system in order to efficiently select events. Similarly, the Data Acquisition System (DAQ) will have to redesign the data-flow architecture to accommodate for the large increase in event and data rates.
The Phase-II DAQ design involves a large distributed storage system that buffers data read out from the detector, while a computing farm (Event Filter) analyzes and selects the most interesting events. This system will have to handle 5.2 TB/s of input data for an event rate of 1 MHz and provide access to 3 TB/s of these data to the filtering farm. A possible implementation for such a design is based on distributed file systems (DFS) which are becoming unavoidable among the big data industry. Features of DFS such as replication strategies and smart placement policies match the distributed nature and the requirements of the new data-flow system.
This paper presents an up-to-date performance evaluation of some of the DFS currently available: GlusterFS, HadoopFS and CephFS. After characterization of the future data-flow system's workload, we report on small-scale raw performance and scalability studies. Finally, we conclude on the suitability of such systems to the tight constraints expected for the ATLAS experiment in phase-II and, in general, what the HEP community can profit from these storage technologies.

Primary authors: AVOLIO, Giuseppe (CERN); ABED ABUD, Adam (Universita and INFN (IT)); LE GOFF, Fabrice (CERN)

Presenter: AVOLIO, Giuseppe (CERN)

Session Classification: Poster Session

Track Classification: Track 1: Computing Technology for Physics Research