Contribution ID: **87** Type: **Short Talk**

Design of a Resilient, High-Throughput, Persistent Storage System for the ATLAS Phase-II DAQ System

Tuesday, 18 May 2021 11:03 (13 minutes)

The ATLAS experiment will undergo a major upgrade to take advantage of the new conditions provided by the upgraded High-Luminosity LHC. The Trigger and Data Acquisition system (TDAQ) will record data at unprecedented rates: the detectors will be read out at 1 MHz generating around 5 TB/s of data. The Dataflow system (DF), component of TDAQ, introduces a novel design: readout data are buffered on persistent storage while the event filtering system analyses them to select 10000 events per second for a total recorded throughput of around 60 GB/s. This approach allows for decoupling the detector activity from the event selection process. New challenges then arise for DF: design and implement a distributed, reliable, persistent storage system supporting several TB/s of aggregated throughput while providing tens of PB of capacity. In this paper we first describe some of the challenges that DF is facing: data safety with persistent storage limitations, indexing of data at high-granularity in a highly-distributed system, and high-performance management of storage capacity. Then the ongoing R&D to address each of them is presented and the performance achieved with a working prototype is shown.

Primary authors: ABED ABUD, Adam (University of Liverpool (GB) and CERN); BONAVENTURA, Matias Alejandro (CERN); FARINA, Edoardo Maria (CERN); LE GOFF, Fabrice (CERN)

Presenter: BONAVENTURA, Matias Alejandro (CERN)

Session Classification: Storage

Track Classification: Online Computing