Contribution ID: **182** Type: **Short Talk**

Exploring Object Stores for High-Energy Physics Data Storage

Wednesday, May 19, 2021 6:45 PM (13 minutes)

Over the last two decades, ROOT TTree has been used for storing over one exabyte of High-Energy Physics (HEP) events. The TTree columnar on-disk layout has been proved to be ideal for analyses of HEP data that typically require access to many events, but only a subset of the information stored for each of them. Future accelerators, and particularly HL-LHC, will bring an increase of at least one order of magnitude in the volume of generated data. To this end, RNTuple has been designed to overcome TTree's limitations, providing improved efficiency and taking advantage of modern storage systems, e.g. low-latency high-bandwidth NVMe devices and object stores. In this paper, we extend RNTuple with a backend that leverages Intel DAOS as the underlying storage, proving that RNTuple's architecture can accommodate such changes. From the RNTuple user's perspective, this data can be accessed with minimal changes to the user code, i.e. replacing a filesystem path by a DAOS URI. Our performance evaluation shows that the contributed backend can be used for realistic analyses, while outperforming the compatibility solution provided by the DAOS project.

Primary authors: LOPEZ GOMEZ, Javier (CERN); BLOMER, Jakob (CERN)

Presenter: LOPEZ GOMEZ, Javier (CERN)

Session Classification: Storage

Track Classification: Distributed Computing, Data Management and Facilities