Contribution ID: **236** Type: **Poster**

Ceph S3 Object Storage for CMS data

Thursday, 27 October 2022 11:00 (30 minutes)

To support the needs of novel collider analyses such as long-lived particle searches, considerable computing resources are spent forward-copying data products from low-level data tiers like CMS AOD and MiniAOD to reduced data formats for end-user analysis tasks. In the HL-LHC era, it will be increasingly difficult to ensure online access to low-level data formats. In this talk, we present a novel online data storage mechanism that obviates the need for data tiers by storing individual data products in column objects using RadosGW, a Ceph object store technology. Benchmarks of the performance of storage and retrieval of the event data through the S3 protocol for a prototype of typical analysis workflows will be presented, and compared with traditional xrootd ROOT file access protocols.

Significance

The use of Ceph object stores and S3 protocol to access experiment data is novel within HEP. Our experience will help guide evaluation and possible adoption of these technologies.

References

https://indico.cern.ch/event/1125222/timetable/?view=standard#32-object-store-rd
https://uscms-software-and-computing.github.io/postdocs/nsmith-.html

Experiment context, if any

CMS

Primary authors: JAYATILAKA, Bo (Fermi National Accelerator Lab. (US)); MASON, David Alexander (Fermi National Accelerator Lab. (US)); SMITH, Nick (Fermi National Accelerator Lab. (US)); GUTSCHE, Oliver (Fermi National Accelerator Lab. (US))

Presenter: SMITH, Nick (Fermi National Accelerator Lab. (US))

Session Classification: Poster session with coffee break

Track Classification: Track 1: Computing Technology for Physics Research