

Filesystem-less High Performance HEP Data Analysis

D. Piparo for the **ROOT** Team and J. Cervantes



- ▶ Key aspect of HEP data analysis: go through lots of data, fast.
- ▶ Huge success of LHC experiments data processing model: read event data from files organised in filesystems.
- ▶ Future Physics programmes: recorded datasets' sizes explosion
 - Analysis - not only real but also simulation datasets, I/O is a bottleneck when it comes to parallel data analysis/processing
 - Calibration constants
 - Quality monitoring (10^{10} histos stored in a CMS DQM server)
- ▶ Is the paradigm of "files" stored on "disks" the only one viable?

Can we identify alternatives from which key use cases can benefit?

- ▶ Investigate exploitation of **object stores** - no blocks, no files: get blobs of data by an identifier
 - Minimal interface: PUT, GET, DELETE + little else
 - Widely adopted by industry
 - Scalable, distributed solution, resilience
- ▶ **Two scenarios: columnar and row-wise storage**
 - columnar: blobs made of some rows (but also blobs of some columns for some rows to allow partial reading)
 - row-wise: individual objects written, e.g. histograms (the data quality monitoring use-case)



Challenges, Potential Milestones, Bonuses

- ▶ Tentative intermediate milestones can be already identified, e.g.
 - 1) Row-wise case study based on experience of existing DQM solutions
 - 2) Transpose present columnar dataset chunking to new backend and exploration of optimisations
- ▶ New backend implies improved separation of concerns in existing I/O sw
 - Benefits for current, traditional solution too
- ▶ New compression/sizes compromises to be explored
- ▶ Built-in error handling