Fermilab | U.S. DEPARTMENT OF ENERGY | Office of Science

Object store research

for the exascale era and beyond

M. Paterno, J. Kowalkowski, and S. Sehrish

20 April 2018

What is our objective?

- Current workflows and frameworks are based around *batch processing* of *files*.
- This is not appropriate for exacale-era machines, without much disk access and with many layers of “storage” (memory).
- Our alternative is to produce a view of all relevant data, accessible through a sufficiently general API, that abstracts away the dealing with the layers of memory.
- We are working with collaborators at ANL that have an *object store* technology they are extending to meet this need.

Deliverables of the work

This work is a part of the SciDAC4 project “HEP Analytics on HPC”. The deliverables are:

- Object store demo at ALCF with HEP data model in KNL system (2018)
- Conclude object store demo and begin storage tier integration for data representation, handling and placement (2019)
- Begin exploring additions to `art` framework (new i/o modules) needed to use the new storage tiers (2019)
- Conclude storage tier integration for data representation, handling and placement (2020)
- Continue exploring upgrades to `art` framework to use the new storage tiers (2020)