HEP-CCE: Fine-grained I/O and Storage (IOS) Activities Year 1

Peter van Gemmeren, Rob Ross

People Involved
(sorry if we missed you!)

- **High Energy Physics**
  - Doug Benjamin (ANL)
  - Paolo Calafiura (LBL)
  - Philippe Canal (FNAL)
  - Oliver Gutsche (FNAL)
  - Salman Habib (ANL)
  - Kenneth Herner (FNAL)
  - Patrick Gartung (FNAL)
  - Lisa Goodenough (FNAL)
  - Christopher Jones (FNAL)
  - Liz Sexton Kennedy (FNAL)
  - Kyle Knoepfel (FNAL)
  - Peter Van Gemmeren (ANL)
- **(More) High Energy Physics**
  - Tammy Walton (FNAL)
  - Torre Wenaus (BNL)
- **Computer Science**
  - Suren Byna (LBL)
  - Matthieu Dorier (ANL)
  - Rob Latham (ANL)
  - Rob Ross (ANL)
  - Saba Sehrish (FNAL)
  - Shane Snyder (ANL)
  - John Wu (LBL)

Statement of the Challenge

- Next-generation HEP experiments will have an order of magnitude increase in data rates and will face increasing data complexity and data volume.
  - New capabilities and methods for data processing, simulations, and data management are needed .
- ASCR HPC resources are expected to be key in addressing several of these issues.
- HEP software and workflows need to be developed to efficiently use these resources.

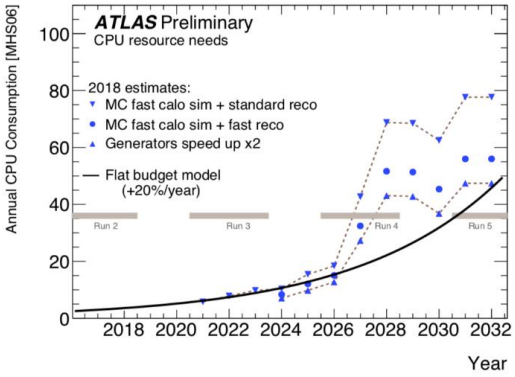

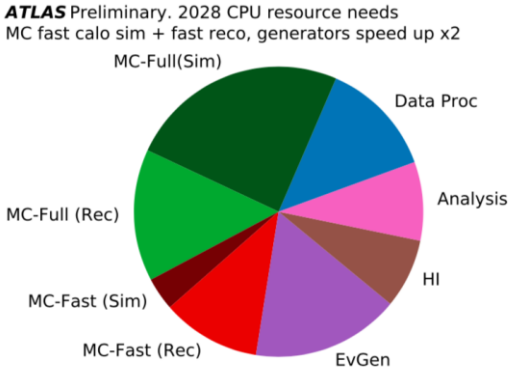

Goals of IOS

- The project is concentrating effort on:
  - Efficient serialization/de-serialization of data representations under parallel I/O models
    - both single node and multi-node access patterns
  - Development of persistable data representations that can be tuned for access on HPC storage systems and optimized for HEP I/O patterns.
    - may involve accelerator technologies
    - can benefit from Write-Once/Read-Many access models
  - Optimization of reads of partial, partitioned or sub-event data blocks from storage which are matched to specific algorithm consumption requirement
  - Optimization of runtime memory mapping of data to exploit batched, vectorized, and data parallel operations and transforms on columnar data.

Year 1 Plans and Accomplishments

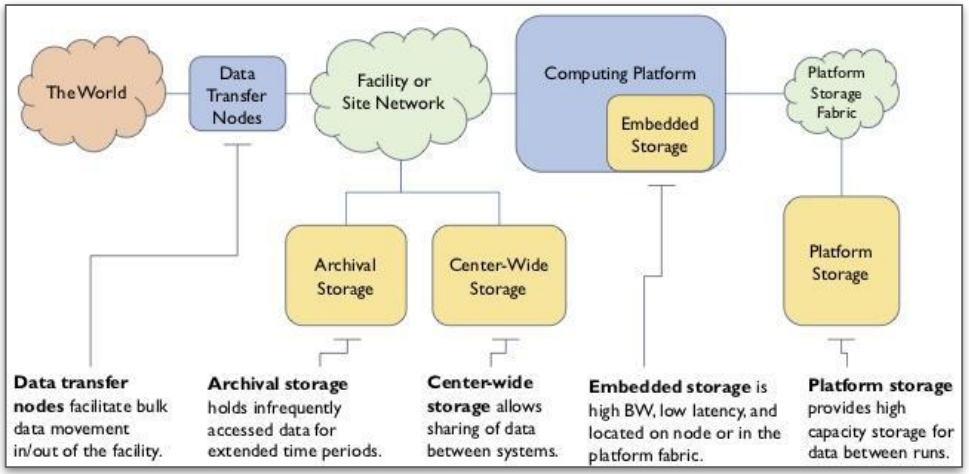

Notional diagram of HPC system (the "platform") and surrounding storage environment.

- **1st quarter**: Document I/O patterns and event data models (EDMs)
  - Short presentations on background topics with Q&A
- **2nd quarter**: Performance of HEP experiment benchmarks on Grid resources
  - ROOT: Optimizable for HPC, xCache, Instrument ROOT I/O patterns.
  - ATLAS: EventService Simulation (fined-grained (event-wise) processing).
- **3rd quarter**: Produce benchmarks either by packaging experiments workflows or by building synthetic benchmarks
- **4th quarter**: Decide on optimization targets for memory infrastructures for phase 2.
  - For example explicit synchronous/asynchronous CPU-GPU data transfers vs unified GPU/CPU memory architectures

Ongoing Activities: Three Deeper Dives

- **Darshan for ROOT I/O in HEP workflows on HPC**
  - ROOT I/O is central to all HEP experiments. Measurements of its performance on HPC using tools like Darshan, could give valuable insights for possible improvements.
  - Shane Snyder presenting.
- **Investigate HDF5 as intermediate event storage for HPC processing**
  - In some workflows, such as the ATLAS EventService, temporary data is written to ROOT files. Moving this data to a parallel file format such as HDF5 could be beneficial.
  - Saba Sehrish presenting.
- **Testing framework for understanding scalability and performance of HEP output methods**
  - An ability to simulate HEP output of specific data products (e.g., RECO, AOD, miniAOD) in different scenarios prepares us for deeper analysis of intermediate data storage options.
  - Chris Jones presenting.