Fermilab U.S. DEPARTMENT OF ENERGY Office of Science

Updates from Computational Science and AI Directorate

Saba Sehrish

57th Annual Users Meeting (2024): Inspirations from P5

07/11/2024

Saba Sehrish – Introduction

- Computational Science Department Head (CSAID/DSSL Division)
- Grant coordinator for CSAID
- Level 2 manager for the Software area for the US-CMS Software and Computing.
- Co-lead of the Storage Optimization project under DOE HEP Center for Computational Excellence

Fermilab

Computational Science and AI Directorate (CSAID) Overview

Finance
Valena Sibley, Senior FFM

Administrative Support
Carla Lloyd, Lead Admin

Computational Science and AI Directorate

James Amundson, Associate Laboratory Director

(Ken Herner), DUNE Software and Computing Coordinator

(Oliver Gutsche), CMS Software and Computing Coordinator

(Robert Harris), Lead Scientist Mentor

(Saba Sehrish), Grant Coordinator

AI Office
(Nhan Tran), Head
(Burt Holzman), Deputy Head

Scientific Computing Systems and Services Division

Margaret Votava, Division Director
Bo Jayatilaka, Deputy Division Director

Data Science, Simulation, and Learning Division

Adam Lyon, Division Director
Erica Snider, Deputy Division Director

Real Time Processing Systems Division

Nhan Tran, Division Director
Ryan Rivera, Deputy Division Director

Elements of the Fermilab Computing Strategy

Data Centers

Mass Storage

Compute Resources

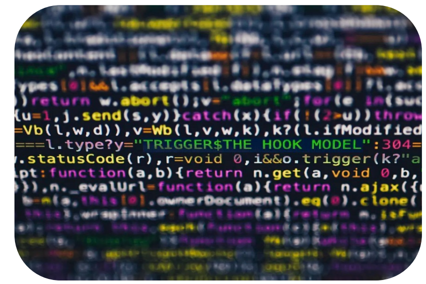

Software Development

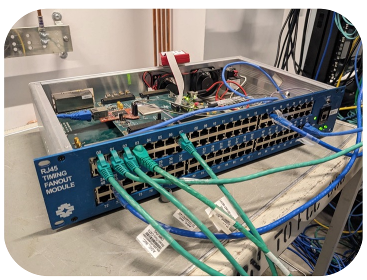

Engineering

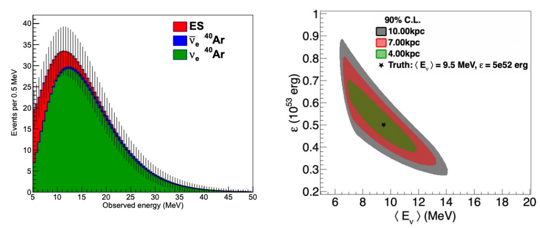

Data Analysis

Diversity, Equity, Inclusion, and Accessibility

Context: Life of HEP Data

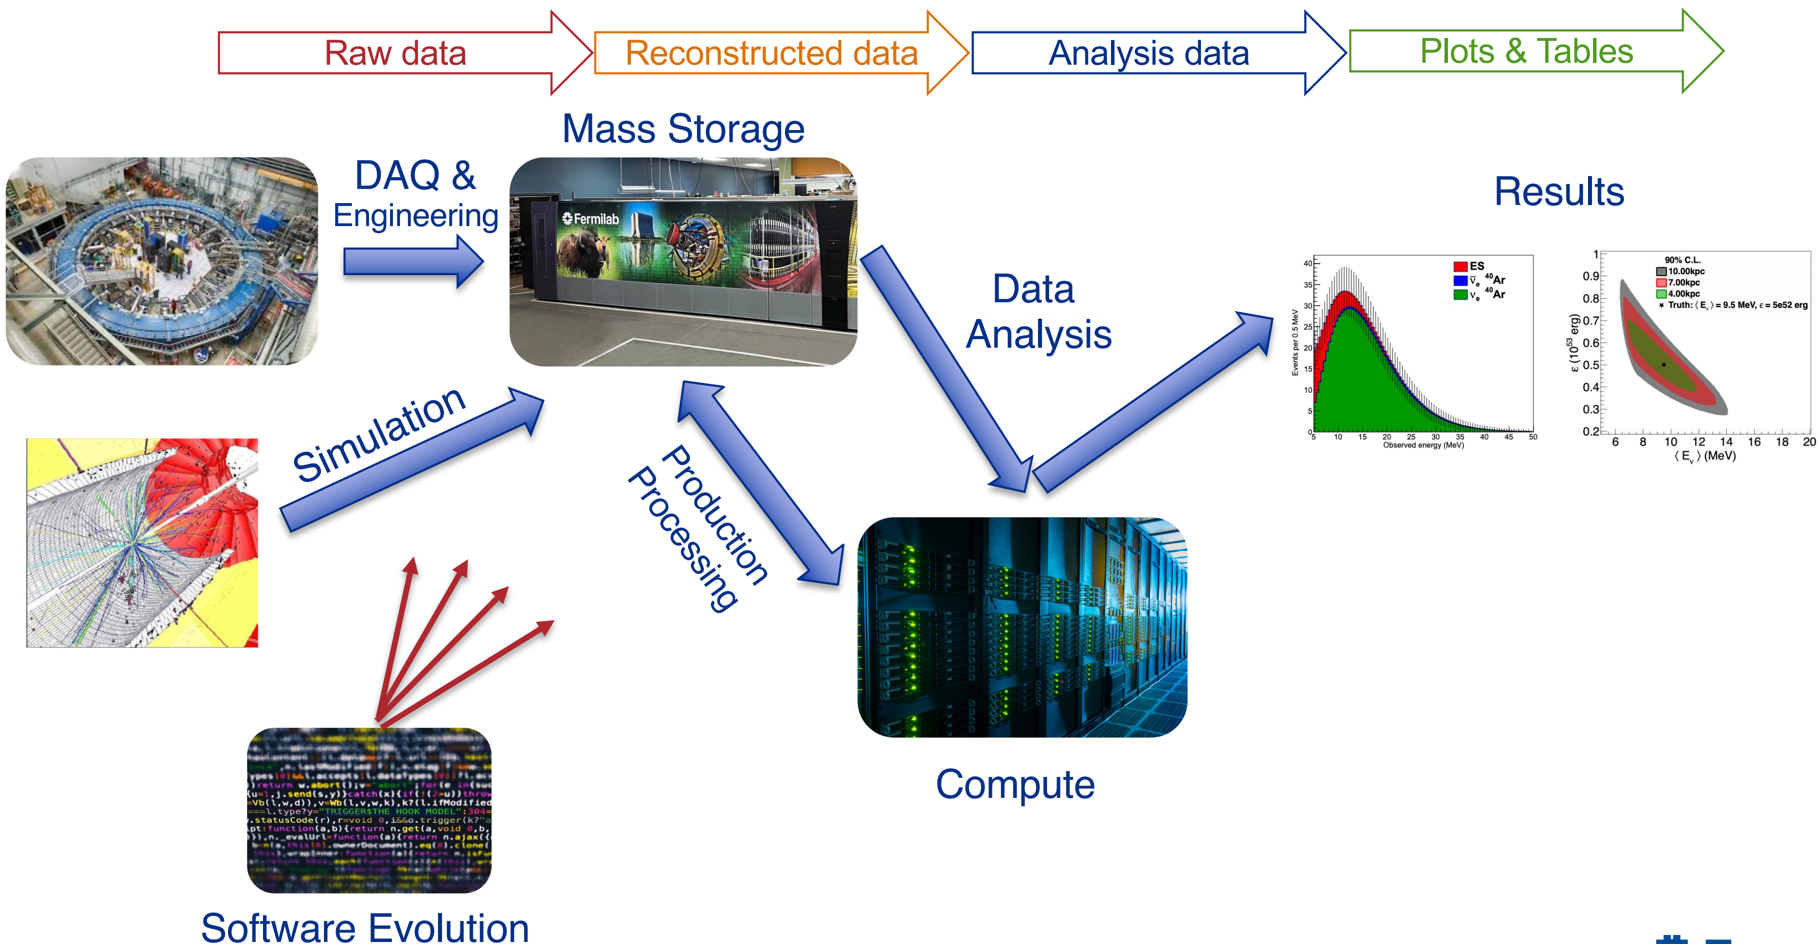

Context: Life of HEP Data

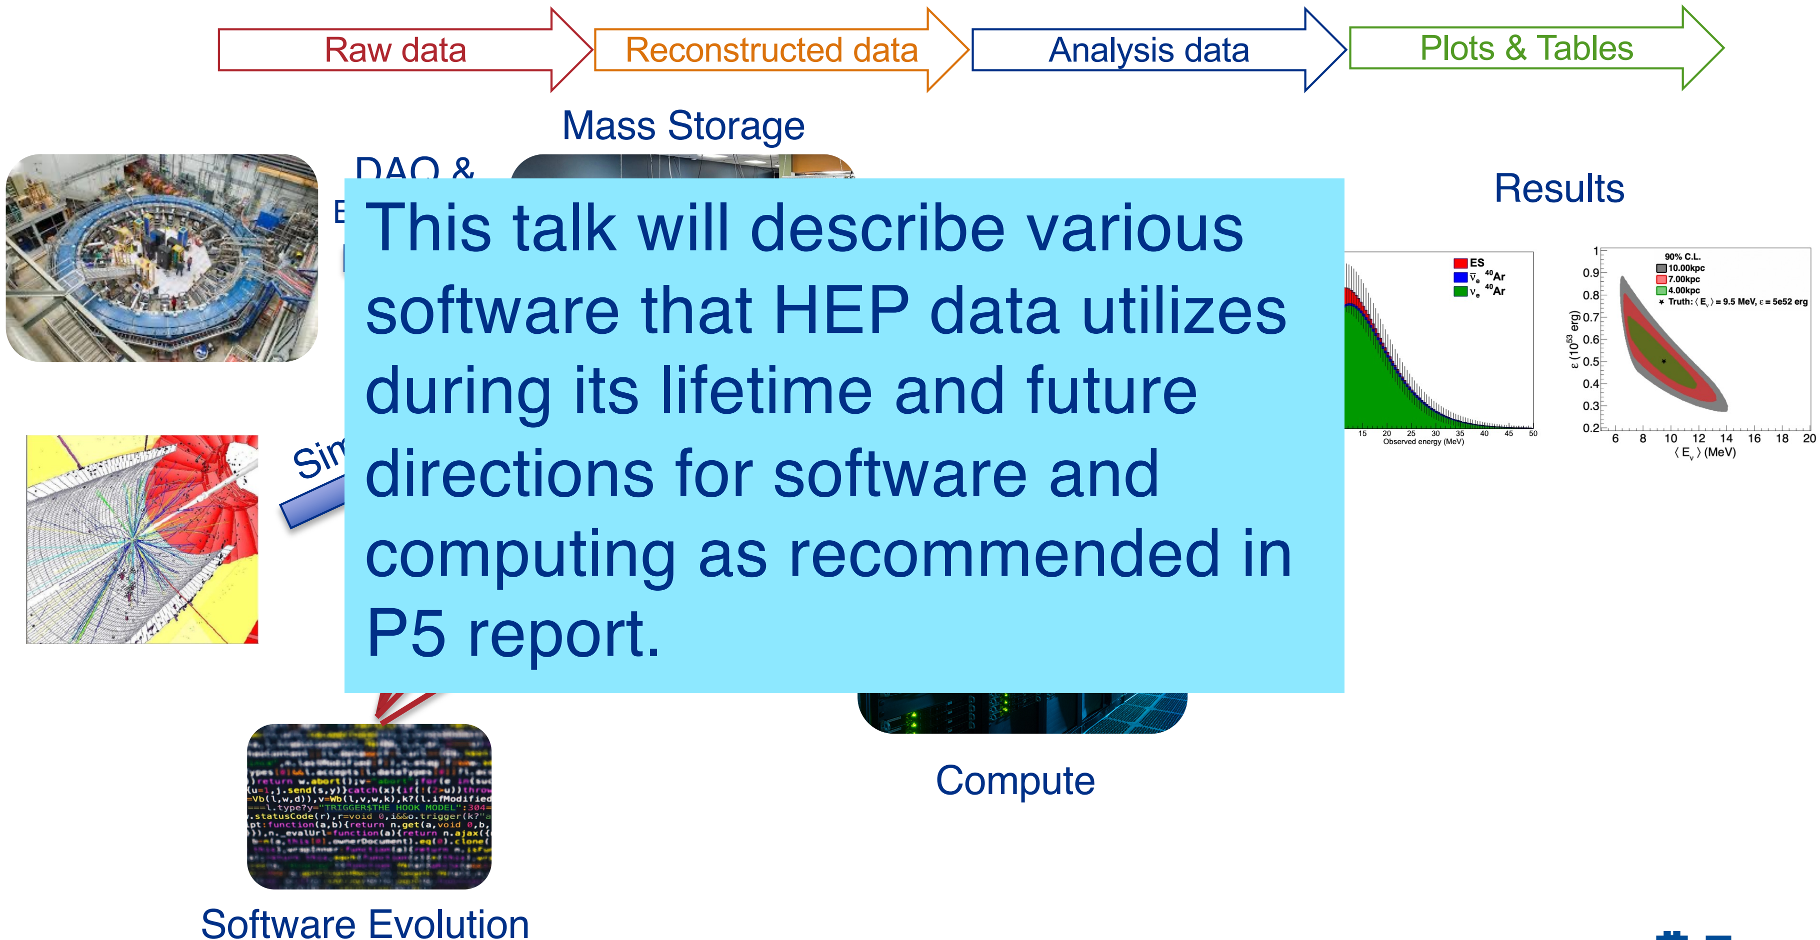

This talk will describe various software that HEP data utilizes during its lifetime and future directions for software and computing as recommended in P5 report.

What I learned about software and computing from P5 report

- Our strong and effective software and computing has played a vital role in the success of many HEP experiments, and we need to continue in the same direction
- Advances in software and computing including machine learning and AI will enable experiments to gather more data and detect rare events at a greater rate
  - Make use of the state-of-the-art hardware, heterogenous resources (GPUs, FPGAs) and modern technologies for effective and efficient software
- Support shared software tools and prioritize computing and novel data analysis techniques for maximizing science across the entire field

Some snippets from P5 report

Advances in **software and computing**, including **AI/ML**, will be key for solving the challenges associated with the data deluge and for enhancing the sensitivity of the experimental results.

Some snippets from P5 report

To promote robust R&D efforts across a range of enabling technologies, we recommend sustained investments in key areas essential to the future of particle physics: theory, an agile project program, detector instrumentation, particle accelerators, collider R&D, facilities and infrastructure, **software and computing**, and **data science.**

Fermilab

Some snippets from P5 report

We must ensure sustained development, maintenance, and user support for key cyberinfrastructure components, including **widely used software packages, simulation tools**, and information resources, such as the Particle Data Group and INSPIRE. Although most of these shared cyberinfrastructure components are not specifically tied to projects, nearly all scientists in the field rely on them.

Mass Storage is the highest priority

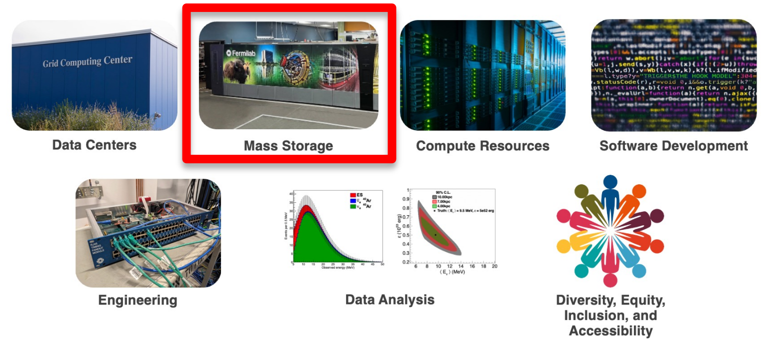

- HEP Experiments have unique requirements and needs for mass storage and data distribution/sharing
- We are atypical users of tape
  - Tape is not just for archiving; we read from tape extensively as part of workflows
    - Disk cache is not large enough to accommodate all popular datasets for IF experiments
  - ~50x more reads than a regular High-Performance Computing (HPC) center
  - Tape is still cost effective compared to other technology, but becoming more challenging
- **We have a long history and extensive experience in developing, operating and managing mass storage and data distribution systems**

Fermilab

New tape system coming

- Enstore the current tape system has served us well for over 20 years
  - As the needs of the experiments have grown, Enstore has reached its scalability limit
- Enstore is being replaced with the CERN Tape Archive (CTA) system
  - https://cta.web.cern.ch/cta/
  - CTA is designed to meet HL-LHC scale datasets needs
  - CTA is a community-based tape system and should grow with our needs
  - The change will be in the first half of 2025
- No changes for the end user. All existing Enstore data can be read with CTA
- dCache will still be used as a buffer system
- Testing of the transfer speed and file sizes is being done.

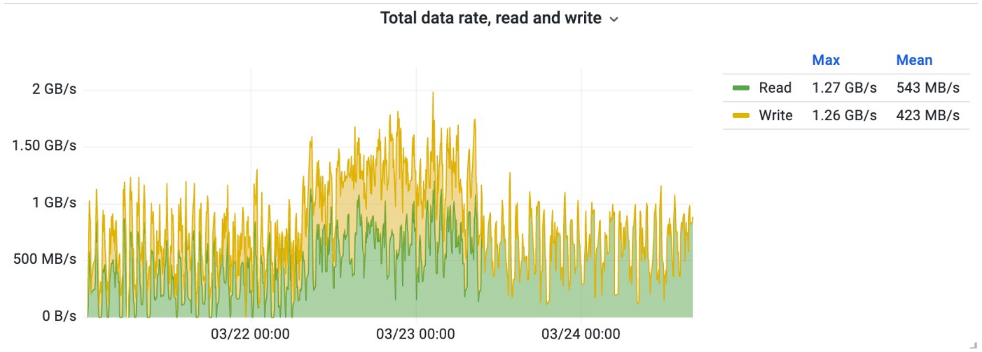

A test of 25TB/day reads and 25TB/day writes was done. This is showing the transfer rate over a three-day period.

New data management system

- Rucio is our exascale data management solution
- Designed with the scale and needs of the HL-LHC in mind, initiated by ATLAS
  - Adopted by CMS, DUNE, Mu2e, ICARUS, SBND, and Rubin observatory
- Provides a system to organize, transfer, archive, and restore large data
- Enables experiments to distribute data across sites taking full advantage of the specifics of those destinations (HPC sites!)
- https://rucio.cern.ch/

CMS data challenge DC 24

- All data managed by Rucio on top of normal CMS traffic (mostly green)
- Final target was 125 GB/s globally (25% of HL-LHC needs)

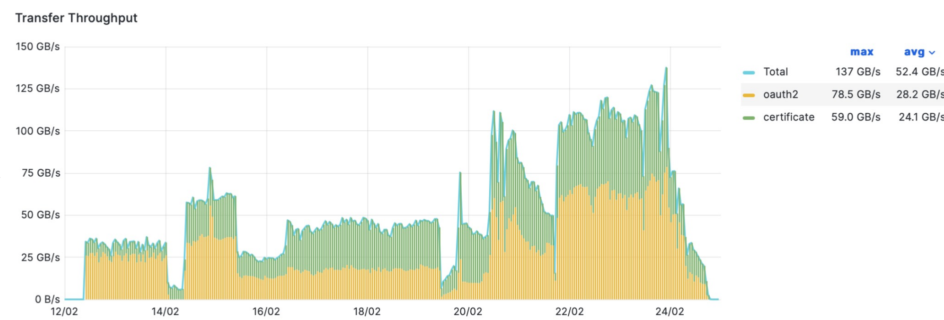

Software development

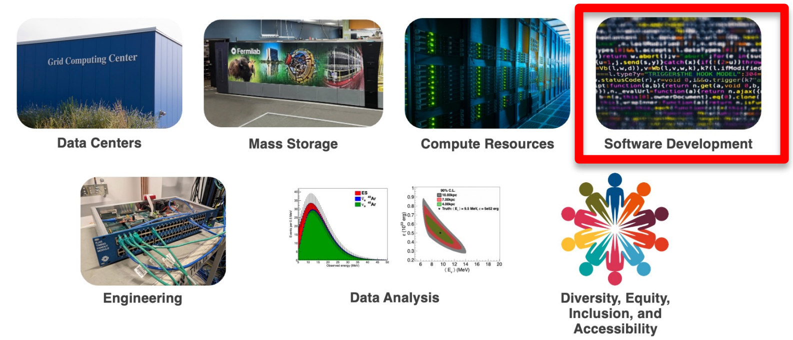

- Fermilab's CSAI directorate provides software framework and supporting infrastructure to many HEP experiments including CMS, DUNE, ICARUS, Mu2e, Muon g-2, NOvA, and SBND.
- Our program is aligned with P5 and DOE priorities, and HEP experiments needs
- Our strategy to provide software and infrastructure support is to
    - Use **community-developed** and **community-supported** software when suitable
    - Develop **shared software** and **common solutions** as needed
    - Keep our software and tools up-to-date with new hardware and language developments

Data processing software used by CMS -- CMSSW

- Collection of software used by the CMS experiment for data processing including software trigger, offline reconstruction, simulation and analysis
- CMS was the first LHC experiment to achieve multi-threaded framework
  - In production use since 2015
- Support for heterogenous resources (GPUs) was first done for run 3 trigger, and now expanding towards more general applications
  - Since 2022, NVIDIA specific (CUDA) infrastructure in production
  - In 2024, deploying portable infrastructure (Alpaka) on GPUs
- Modernize code base; support for C++20 enabled in summer 2024
- Thousands of users and hundreds of unique developers' contributions per month
- http://cms-sw.github.io

Data processing framework shared by IF experiments – *art*

- The application framework for creating runtime-configurable simulation and reconstruction programs
- Meeting experiments framework and data processing needs since last 15 years
- Provides the framework needs for most Fermilab-hosted experiments
  - Framework successfully used by 10+ experiments
  - Over 2000 users
  - Governed by stakeholder consensus and laboratory priorities
- Parallel processing of events supported since 2018
- Modernize the code base; the first HEP framework to completely support C++20
- https://art.fnal.gov

Future data processing framework – from R&D to development

- Framework requirements by DUNE could not be met by the existing framework, and we were awarded an LDRD (Meld) to address those requirements
  - Rigid collider-based event data model concept not workable for DUNE, require more flexibility
- Meld aims to directly support the frameworks needs of neutrino experiments like DUNE as well as the collider-based experiments
- Use state-of-the-art concurrency libraries and modern C++ to create a new framework with much more runtime configurable flexibility in both the organization of data and processing than existing frameworks
  - In collaboration with Intel to support HEP use cases
- A major task for next couple of years is working on the migration path from art to the new framework while supporting existing experiments

Shared toolkit for liquid argon experiments – LArSoft

- For over a decade LArSoft has successfully enabled code sharing among experiments and have helped small experiments to use sophisticated tools and algorithms
- Supporting multi-threading, and High-Performance Computing (HPC) for critical production workflows is the focus.
- Provide support for hardware acceleration where appropriate
    - Replace select algorithms with portable GPU-enabled algorithms
    - Enabling the use of GPUs in production workflows more generally
- Facilitate / simplify integration of machine learning workflows
- Expand adoption of community / industry supported tools
- https://larsoft.org

And more …

- Using a community supported system that could replace our homegrown package manager and all related tools
  - Spack is a package management tool designed to support multiple versions and configurations of software on a wide variety of platforms and environments.
  - It was designed for large supercomputing centers.
- Several experiments including DUNE, ICARUS, SBND, and Mu2e are successfully using Spack
- Contribute to advancing ROOT (which is used for both IO and data analysis)

Simulation software

- Working collaboratively with other institutions on development and support of simulation software, as well as on performing some of the related simulations
  - Geant4 - main HEP detector simulation toolkit
    - Celeritas - a project to enable detector simulation on heterogenous architectures using GPUs
  - GENIE, MARLEY - neutrino event generators
  - Pythia – multipurpose event generator
  - QUDA - library for lattice QCD on GPUs
  - Synergia - package for accelerator bunch simulations, extensively using portability approaches such as Kokkos

Data analysis

- Supporting interactive analyses with Jupyter notebooks and AI training with GPUs via Elastic Analysis Facility (EAF)
- Goal is for EAF to become the analysis hub for HEP experiments at Fermilab
- Future: “Elastically” expand into HPC and Cloud resources as needed. Add a flexible batch system.
- Try https://analytics-hub.fnal.gov

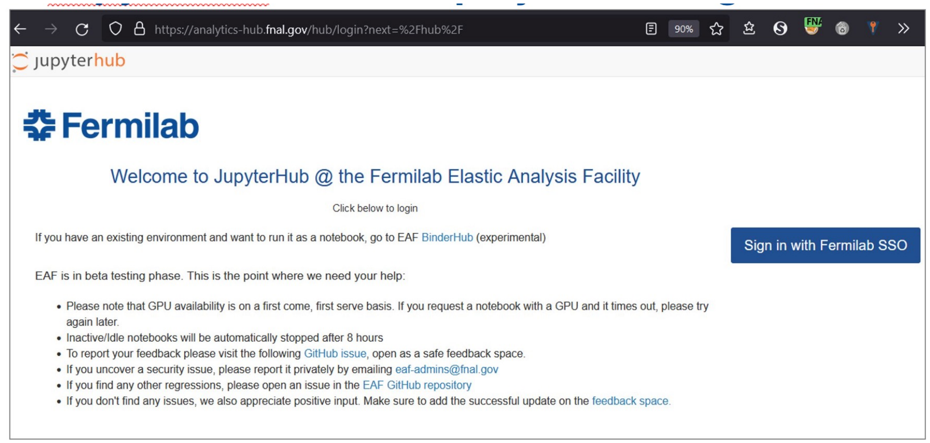

Fermilab

Engineering

- Examples of community tools and frameworks for experiment and accelerator data acquisition and control developed by specialized team of engineers, physicists, and computing professionals
  - DUNEDAQ, artdaq, and otsdaq for intensity frontier experiments
  - hls4ml for embedded system ML development
  - QICK for qubit read out and control

Fermilab

R&D

- Several R&D initiatives included 3 successfully completed SciDAC4 projects and several ASCR awards
  - SciDAC – Scientific Discovery through Advanced Computing
  - ASCR – Advanced Scientific Computing Research
- DOE HEP Center for Computational Excellence -- ASCR experts at DOE leadership facilities and other labs working with HEP experiments to effectively use exascale machines
- Active engagement and collaboration with ASCR researchers and programs

Summary

- We provide critical and effective software and computing for HEP data
- Our current and future software and computing plan of work is aligned with HEP experiment needs, DOE priorities and P5 recommendations
- We have several successful examples of using community-based and community-supported solutions
- We have invested in shared and common solutions successfully
- Our data processing software are making use of heterogenous resources (GPUs)
- Thanks to everyone who helped to put this material together.

Thank you!

Questions