Scientific Computing at Fermilab

Tammy Walton

54th Fermilab Annual (Virtual) Users Meeting

August 4, 2021

Outline

- Introduction
- Scientific Computing Services
- Scientific Computing Projects and Facilities
- Summary

Scientific Computing

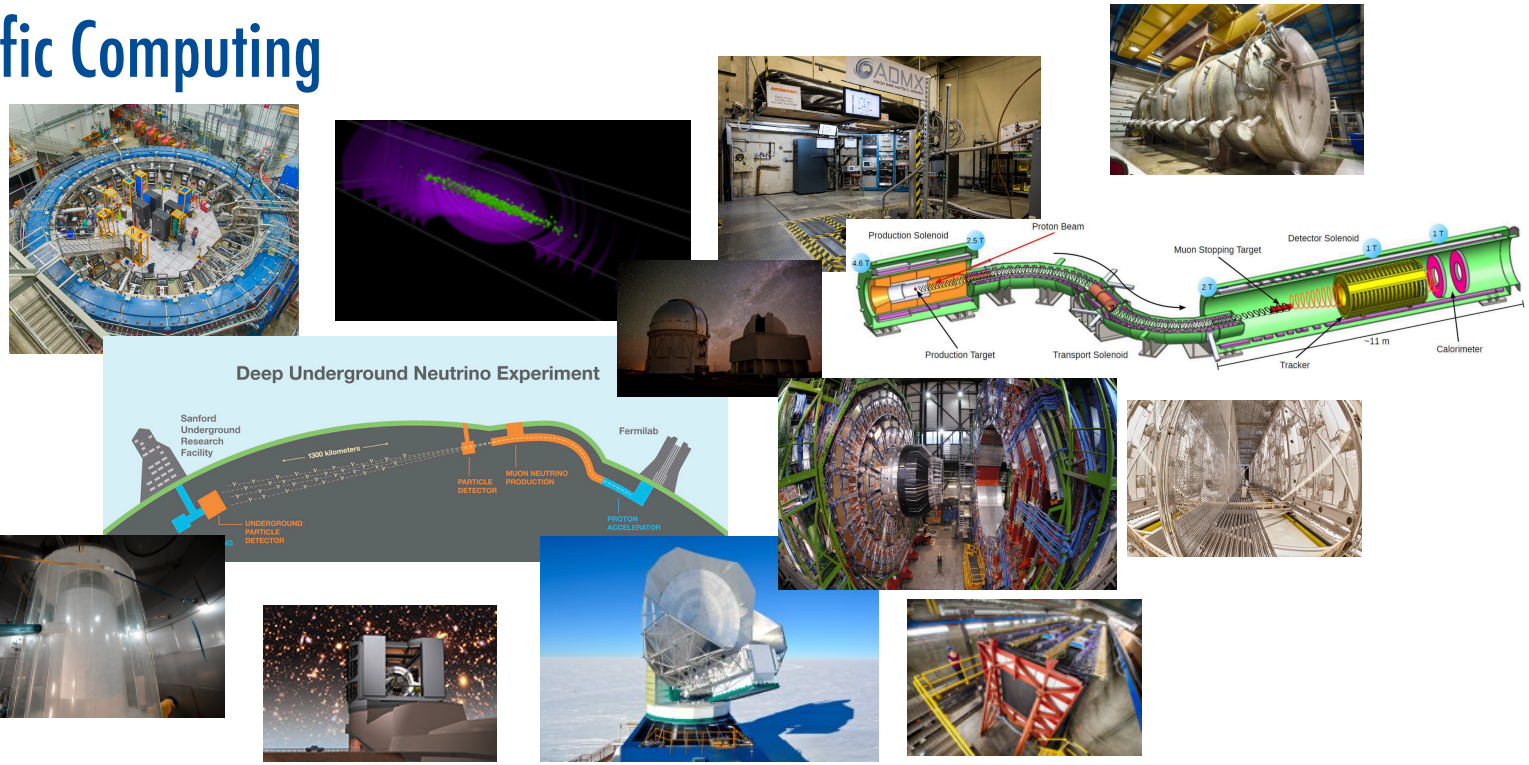

Fermilab hosts and participates in a variety of scientific physics programs

- Neutrino physics
- Muon physics
- CMS physics
- Dark matter and energy physics
- Accelerator physics

Fermilab

Scientific Computing

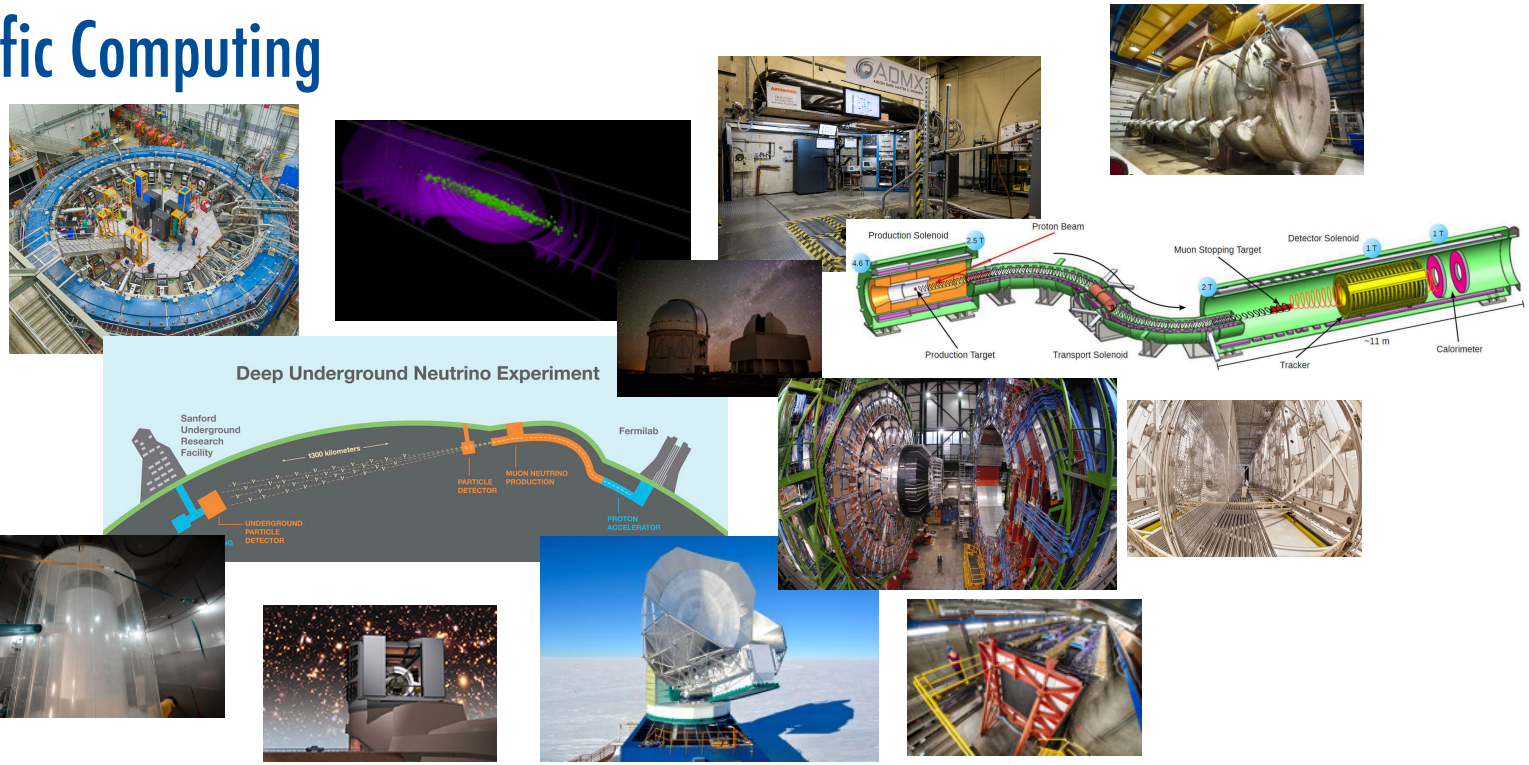

Fermilab hosts and participates in a variety of scientific physics programs

- Some experiments are projected to produce big data events.
  - DUNE supernova event : 185 terabytes
  - HL-LHC total data collection : ~20 exabytes = 20,000,000 terabytes

Scientific Computing

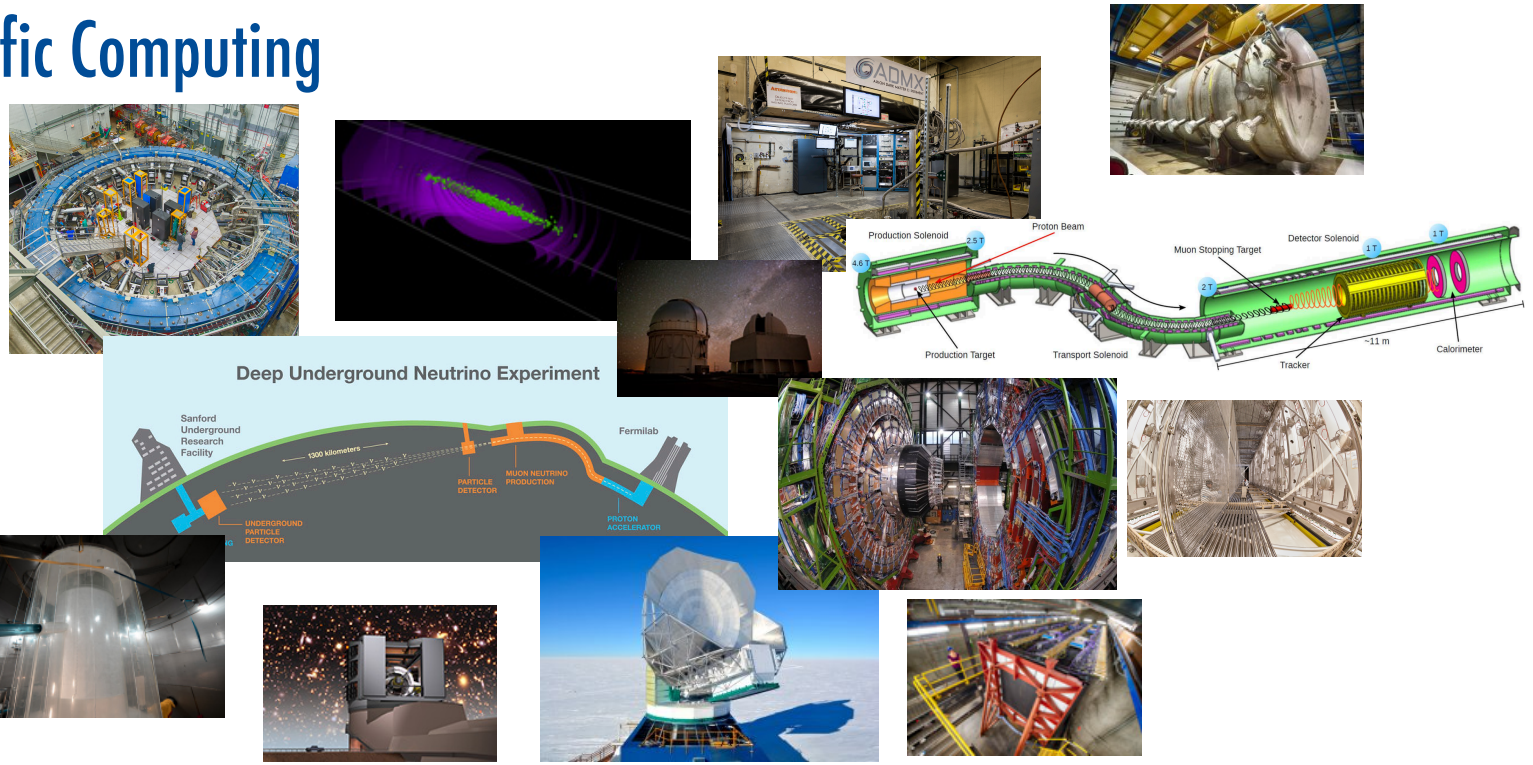

Fermilab hosts and participates in a variety of scientific physics programs

- Scientific computing supports and develops the leading-edge computational tools and facilities that experiments need for a successful physics program.

Fermilab

Scientific Computing

https://computing.fnal.gov/

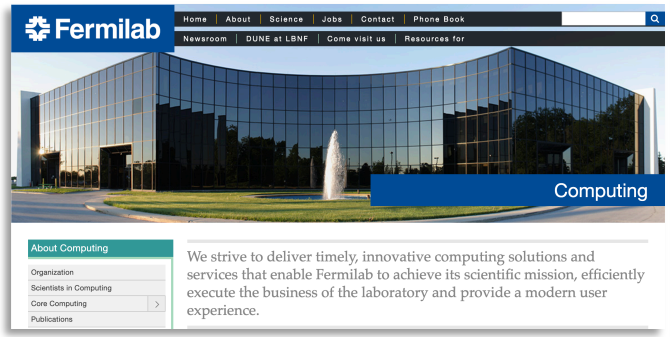

https://indico.fnal.gov/event/47936/

James Amundson
Head of Scientific Computing Division

Fermilab hosts and participates in a variety of scientific physics programs

- Scientific computing supports and develops the leading-edge computational tools and facilities that experiments need for a successful physics program.
- Scientific computing division consists of computing professionals, scientists and researchers; experimental, theoretical, and accelerator physicists; engineers, research associates

Scientific Computing

https://computing.fnal.gov/

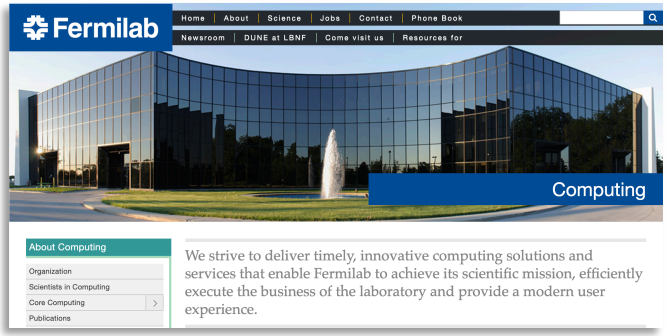

https://indico.fnal.gov/event/47936/

James Amundson
Head of Scientific Computing Division

Fermilab hosts and participates in a variety of scientific physics programs

- Scientific computing supports and develops the leading-edge computational tools and facilities that experiments need for a successful physics program.

This presentation provides an overview of contributions from some of the available services, projects, and facilities

Outline

- Introduction
- Scientific Computing Services
- Scientific Computing Projects and Facilities
- Summary

Scientific Computing Services

- Support and develop common computing tools for all experiments
  - Data management and submission
  - Software distribution and build systems
  - Source code version control systems and repositories
  - Access to Open Science Grid and High-Performance Computing centers
  - Interactive computing machines
  - Online and Offline software frameworks
  - Interactive Analysis Tools

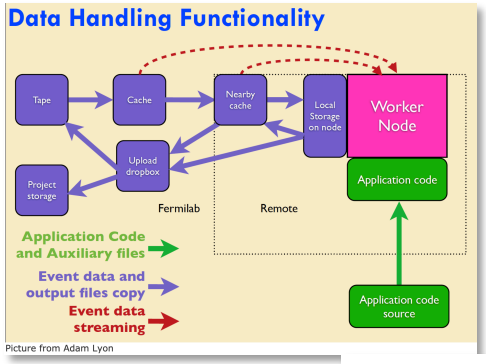

Picture from Adam Lyon

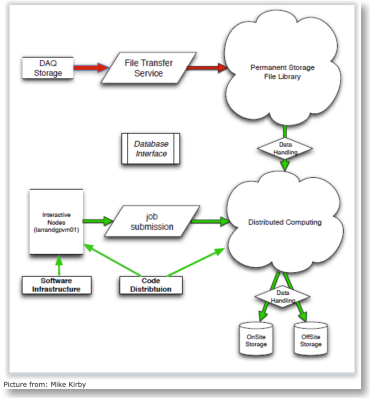

Picture from: Mike Kirby

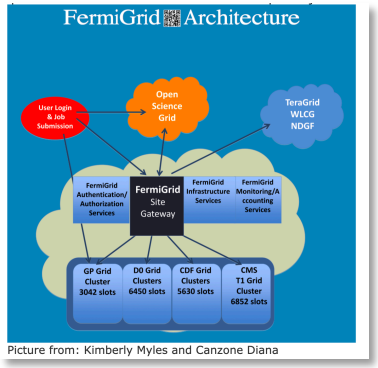

Picture from: Kimberly Myles and Canzone Diana

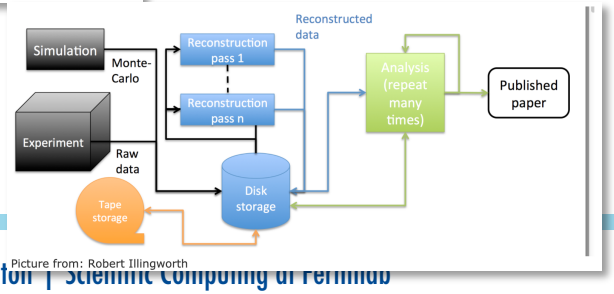

Picture from: Robert Illingworth

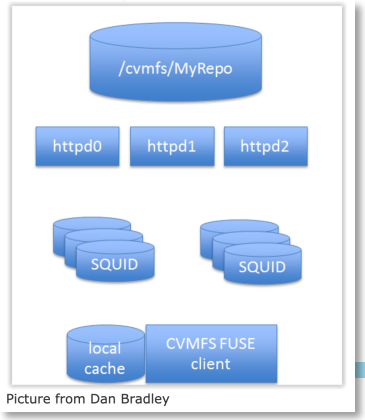

Picture from Dan Bradley

FABRIC FOR FRONTIER EXPERIMENTS

Scientific Computing Services

- Support and develop common computing tools for all experiments
- Provide services to answer questions from experiments about computing

https://fermi.servicenowservices.com

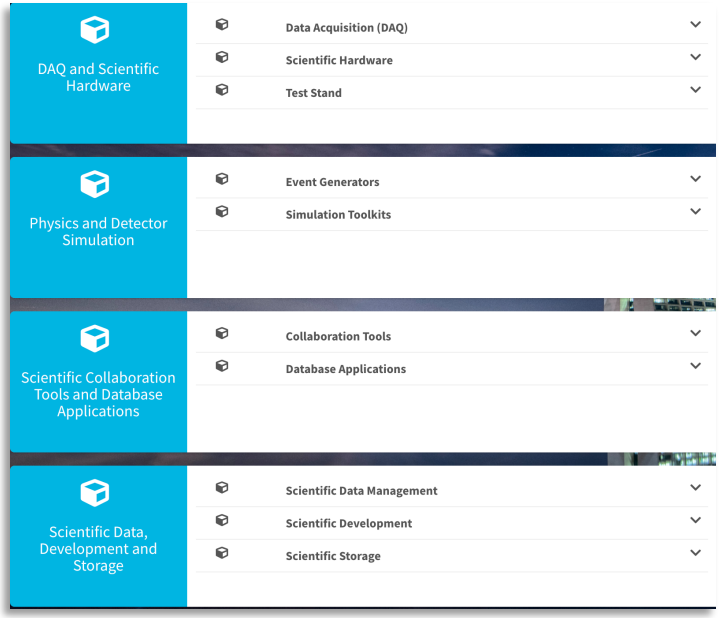

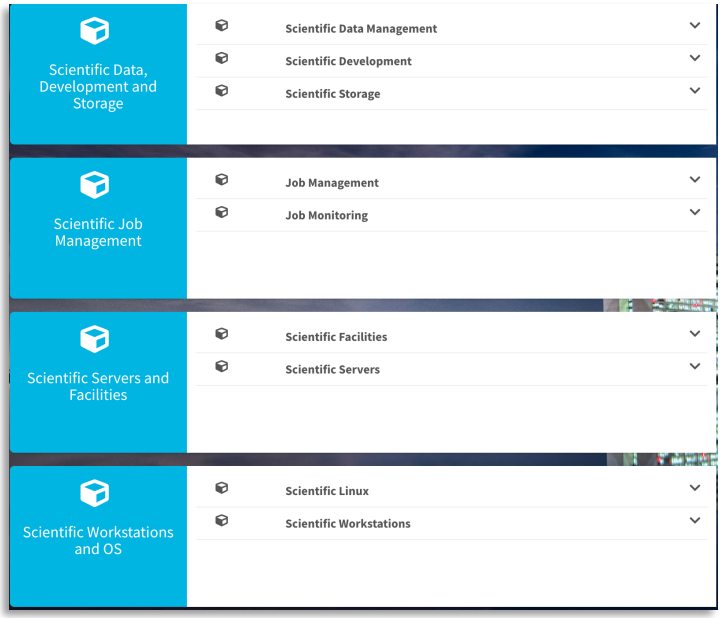

Fermilab

Scientific Computing Services

- Support and develop common computing tools for all experiments
- Provide services to answer questions from experiments about computing
- Provide tutorials and workshops to learn about the various computing tools

https://indico.fnal.gov/event/49414/

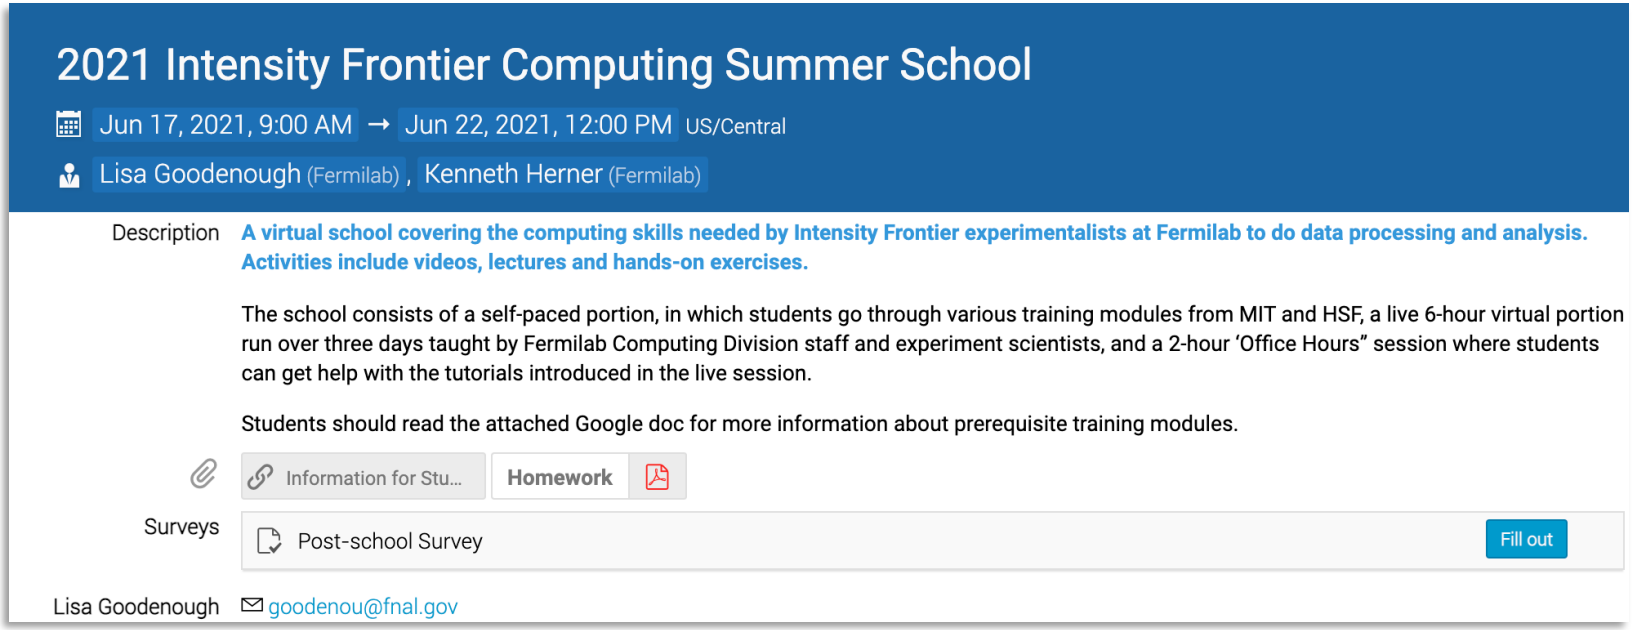

Scientific Computing Services

- Support and develop common computing tools for all experiments
- Provide services to answer questions from experiments about computing
- Provide tutorials and workshops to learn about the various computing tools
- Support and develop data monitoring tools for experiments

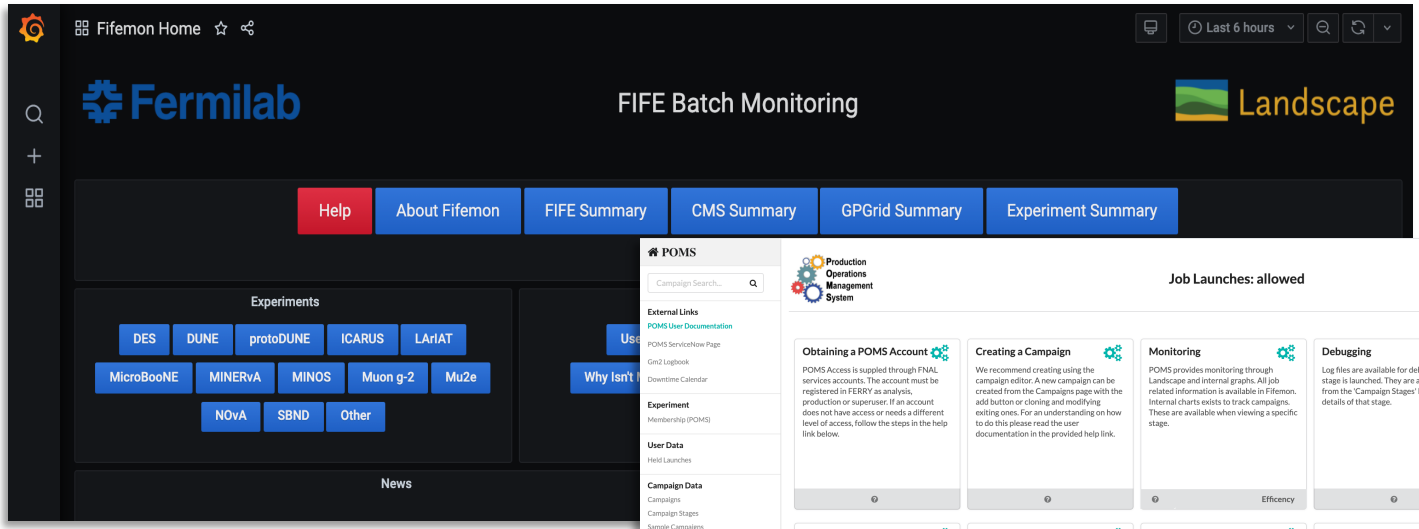

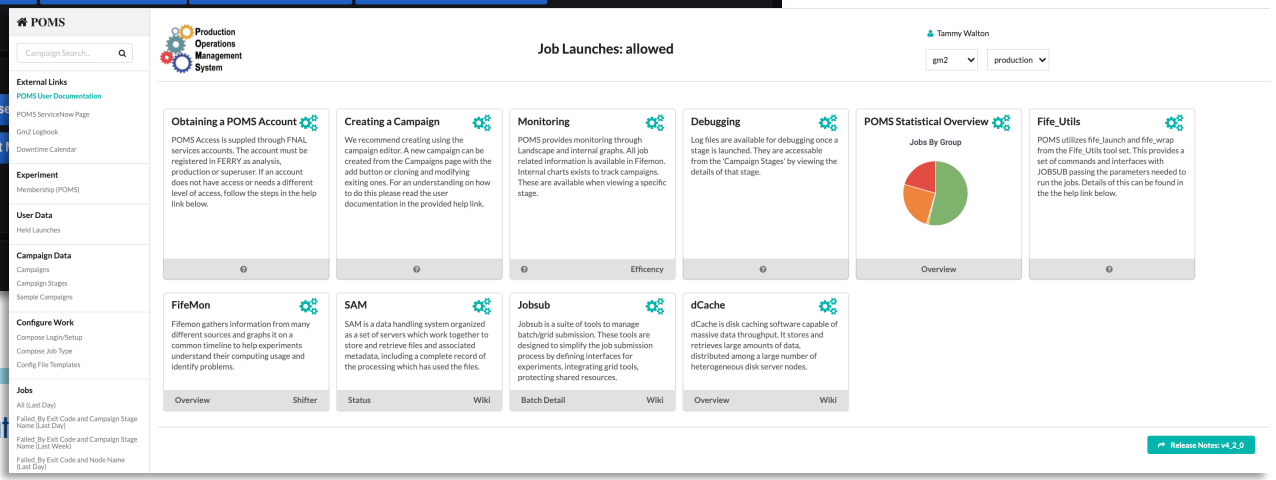

Scientific Computing Services

- Support and develop common computing tools for all experiments
- Provide services to answer questions from experiments about computing

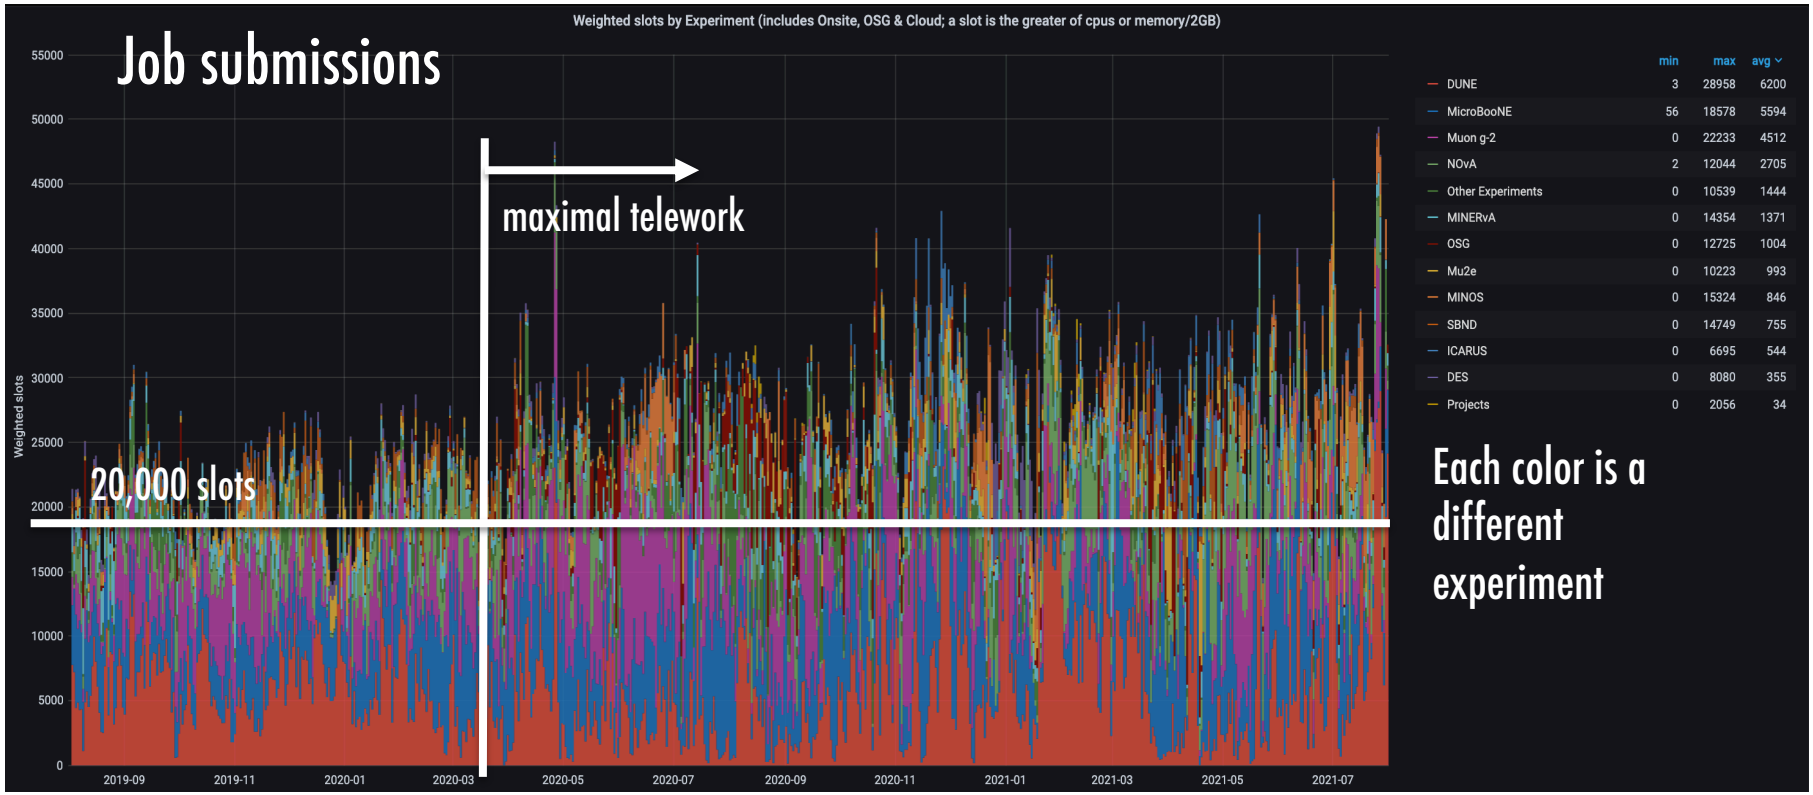

Fermilab

Outline

- Introduction
- Scientific Computing Services
- Scientific Computing Projects and Facilities
- Summary

SCIENTIFIC LINUX

art

G4

CernVM-FS
CernVM File System

Genie
UNIVERSAL NEUTRINO GENERATOR & GLOBAL FIT

FIFE
FABRIC FOR FRONTIER EXPERIMENTS

python™

SciDAC
Scientific Discovery through Advanced Computing

ROOT
Data Analysis Framework

LArSoft

dCache.org
Enstore

CMS
CMSSW

XRootD

HSF
HEP Software Foundation

Fermilab

Scientific Computing Projects and Facilities

- Computing resources
  - Spring 2019: rollout of HEPCloud (https://hepcloud.fnal.gov/)
    - Unified portal to access many computing resources
      - Local (FermiGrid) and Offsite (Open Science Grid) resources
      - High Performance Computing (HPC) sites (commercial cloud and scientific centers)
        - Supercomputers that use parallel processing for running advanced applications fast and efficient
        - Performance is measured in floating-point operations per second

Fermilab

Scientific Computing Projects and Facilities

- Computing resources
  - Spring 2019: rollout of HEPCloud (https://computing.fnal.gov/hep-cloud/)
    - Unified portal to access many computing resources
    - FIFE support is the interface between experiments and facilities

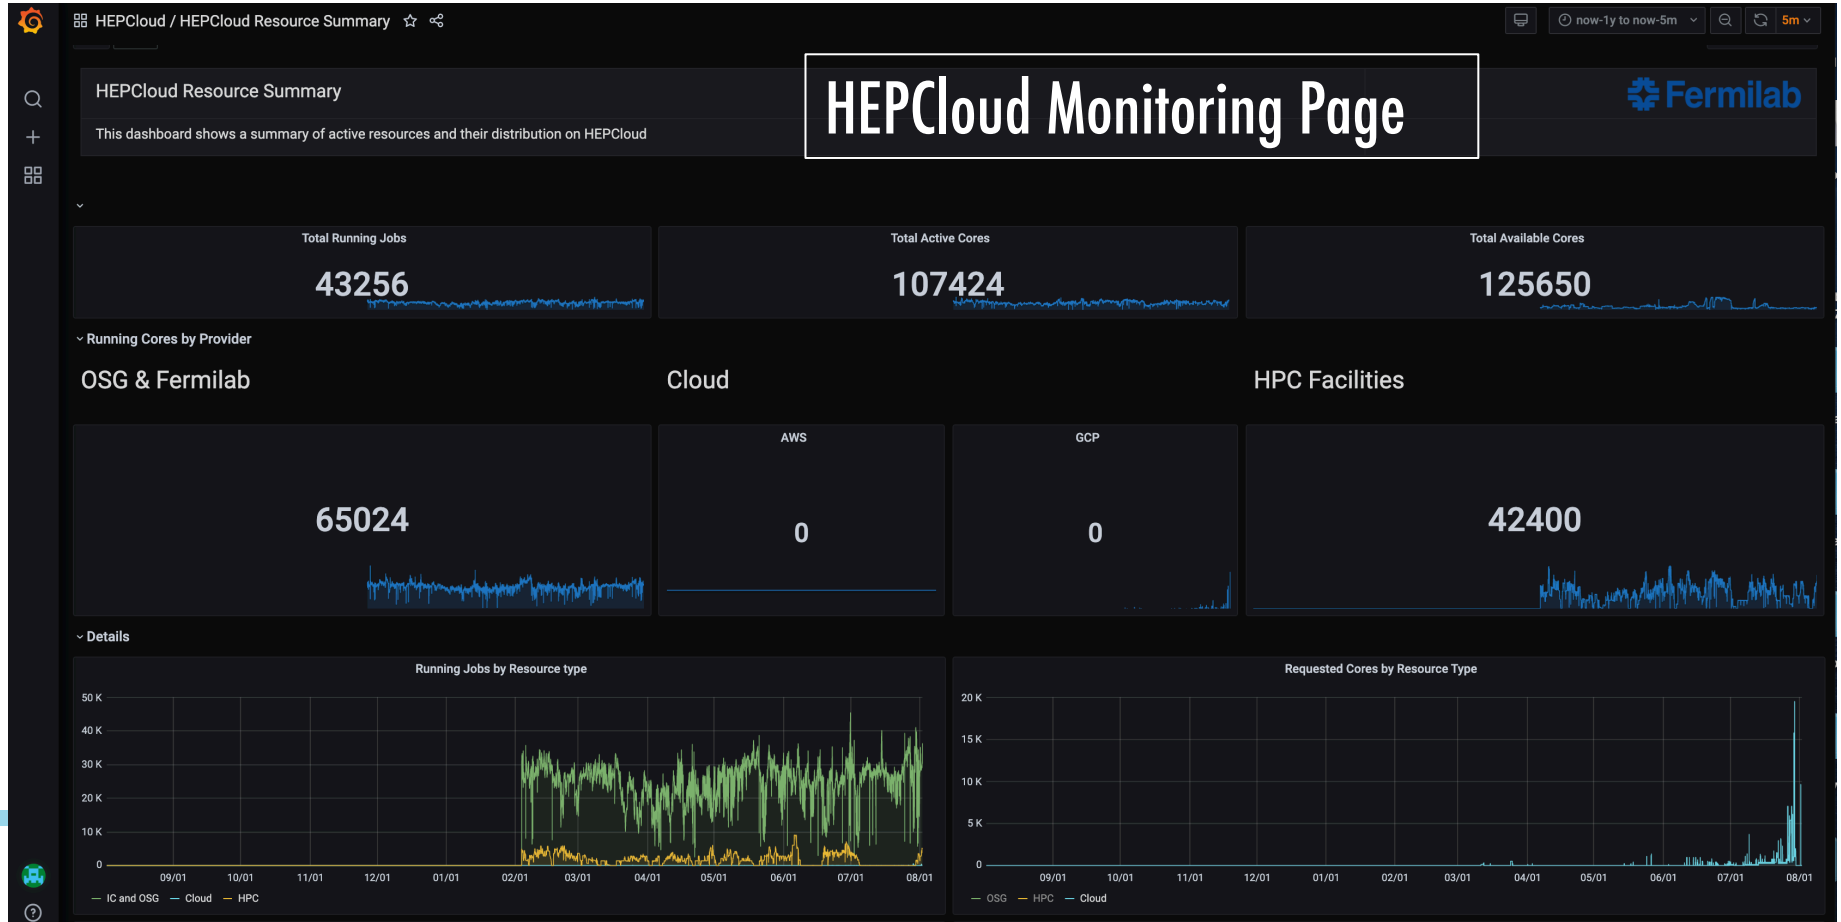

Scientific Computing Projects and Facilities

- Computing resources
  - Spring 2019: rollout of HEPCloud (https://computing.fnal.gov/hep-cloud/)
    - Unified portal to access many computing resources
    - FIFE support is the interface between experiments and facilities
    - Plans include to integrate facilities that will support exascale computing resources
      - Oak Ridge National Laboratory 200-petaflop supercomputer, Summit, is on the road to an exascale upgrade

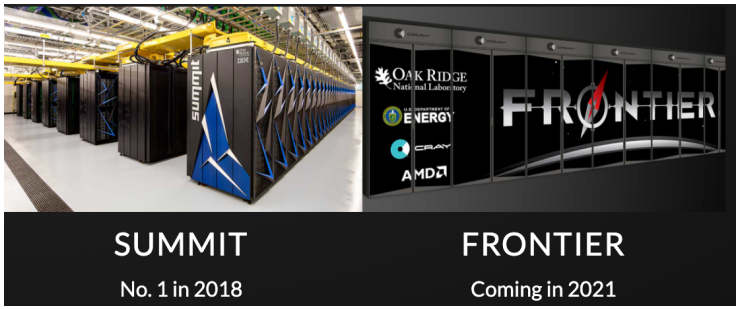

Scientific Computing Projects and Facilities

- Computing resources
- Advanced applications and algorithm techniques; large-scale data analytics
  - HPCs are required for success in HEP
  - Fermilab has hosted the Institutional Cluster for over a decade, and is used largely by the US Lattice QCD collaboration
  - Experiments (CMS, NOvA, DUNE, mu2e, g-2, etc) are actively running production jobs on HPC sites via HEPCloud

A very nice overview update on HPCs presented by Lisa Goodenough at 2020 Users Meeting

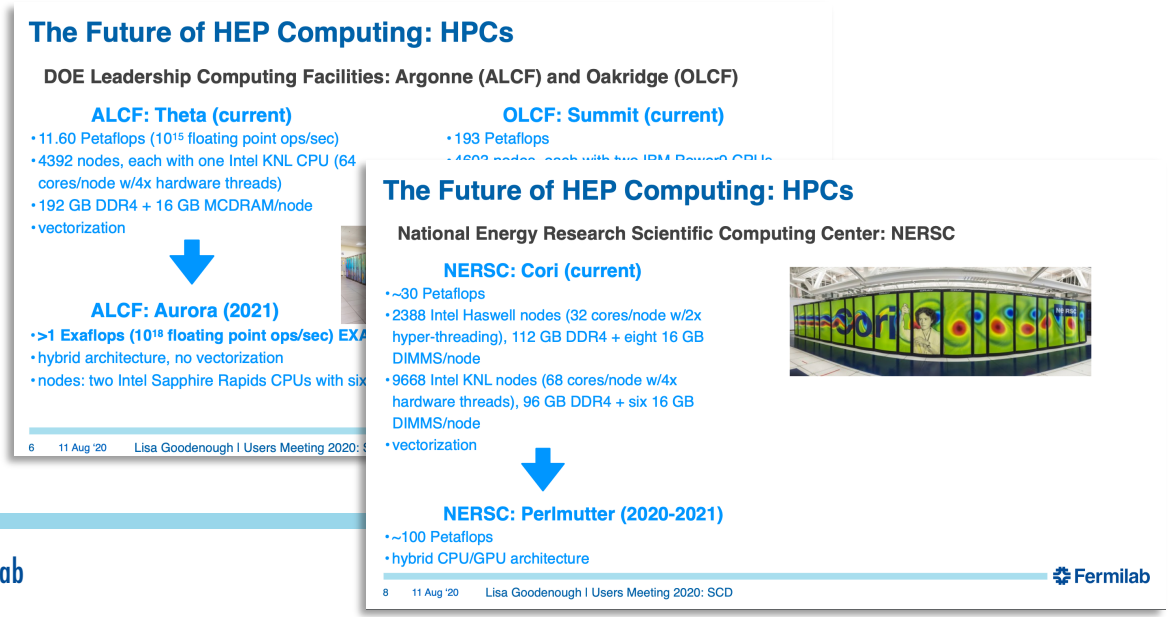

Scientific Computing Projects and Facilities

- Computing resources
- Advanced applications and algorithm techniques; large-scale data analytics
  - HPCs are required for success in HEP
  - Fermilab has hosted the Institutional Cluster for over a decade, and is used largely by the US Lattice QCD collaboration
  - Experiments (CMS, NOvA, DUNE, mu2e, g-2, etc) are actively running production jobs on HPC sites via HEPCloud
  - DOE is funding tools for the community to use HPCs

User Facilities

The Advanced Scientific Computer Research program supports the following national scienti

Argonne Leadership Computing Facility (ALCF)

The ALCF provides the computational science community with a world-class computing capa breakthrough science and engineering.

Energy Sciences Network (ESnet)

The ESnet is a high-speed network serving thousands of Department of Energy researchers worldwide.

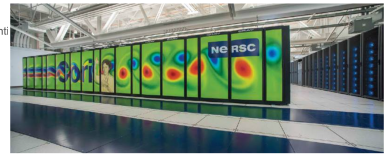

National Energy Research Scientific Computing Center (NERS

The NERSC is the mission high performance computing facility for Science, and is a world leader in accelerating scientific discovery t

Oak Ridge Leadership Computing Facility (OLCF)

The OLCF provides the computational science community with world-class computing capability dedicated to breakthrough science and engineering.

U.S. DEPARTMENT OF ENERGY | Office of Science

https://science.osti.gov/ascr

About

The mission of the Advanced Scientific Computing Research (ASCR) program is to discover, develop, and deploy computational and networking capabilities to analyze, model, simulate, and predict complex phenomena important to the Department of Energy (DOE). A particular challenge of this program is fulfilling the science potential of emerging computing systems and other novel computing architectures, which will require numerous significant modifications to today's tools and techniques to deliver on the promise of exascale science.

Scientific Computing Projects and Facilities

- Computing resources
- Advanced applications and algorithm techniques; large-scale data analytics
  - HPCs are required for success in HEP
  - Fermilab has hosted the Institutional Cluster for over a decade, and is used largely by the US Lattice QCD collaboration
  - Experiments (CMS, NOvA, DUNE, mu2e, g-2, etc) are actively running production jobs on HPC sites via HEPCloud
  - DOE is funding tools for the community to use HPCs
  - Develop and quantify performance on various HPC architects [CPU, GPU, and CPU/GPU architectures]

User Facilities

The Advanced Scientific Computer Research program supports the following national scienti

Argonne Leadership Computing Facility (ALCF)

The ALCF provides the computational science community with a world-class computing capa breakthrough science and engineering.

Energy Sciences Network (ESnet)

The ESnet is a high-speed network serving thousands of Department of Energy researchers worldwide.

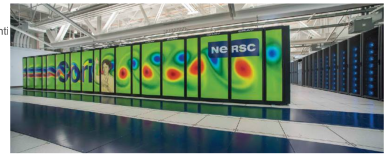

National Energy Research Scientific Computing Center (NERS

The NERSC is the mission high performance computing facility for Science, and is a world leader in accelerating scientific discovery t

Oak Ridge Leadership Computing Facility (OLCF)

The OLCF provides the computational science community with world-class computing capability dedicated to breakthrough science and engineering.

U.S. DEPARTMENT OF ENERGY | Office of Science

https://science.osti.gov/ascr

About

The mission of the Advanced Scientific Computing Research (ASCR) program is to discover, develop, and deploy computational and networking capabilities to analyze, model, simulate, and predict complex phenomena important to the Department of Energy (DOE). A particular challenge of this program is fulfilling the science potential of emerging computing systems and other novel computing architectures, which will require numerous significant modifications to today's tools and techniques to deliver on the promise of exascale science.

Scientific Computing Projects and Facilities

- Computing resources
- Advanced applications and algorithm techniques; large-scale data analytics
- Developed applications to run HPC platforms
  - HEP-Center for Computational Excellence (HEP-CCE)
    - Fermilab institutional lead, Liz Sexton-Kennedy
    - Focuses on HPC I/O systems and storage
      - Order of magnitude increase in data rates

Elizabeth Sexton-Kennedy
Chief Information Officer

Argonne NATIONAL LABORATORY · BROOKHAVEN NATIONAL LABORATORY · Fermilab · BERKELEY LAB Bringing Science Solutions to the World

https://www.anl.gov/hep-cce

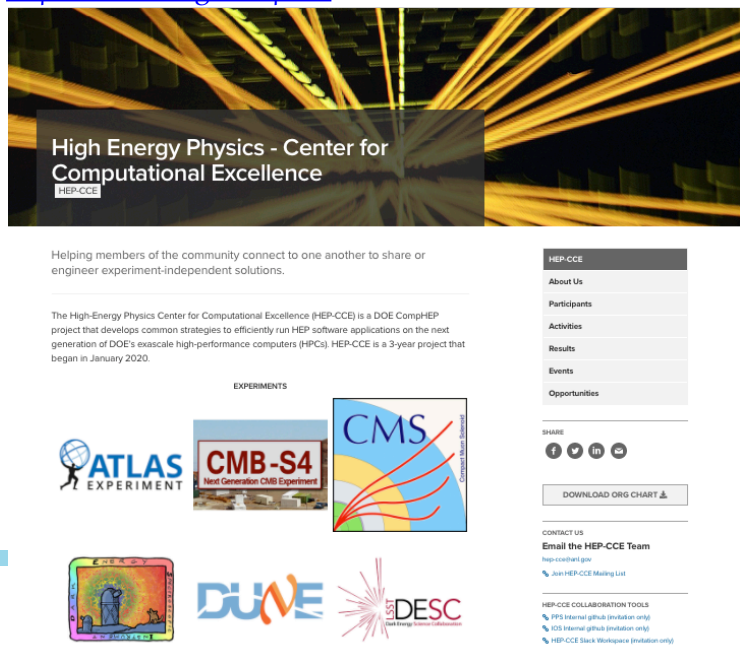

Scientific Computing Projects and Facilities

- Computing resources
- Advanced applications and algorithm techniques; large-scale data analytics
- Developed applications to run HPC platforms
  - HEP-Center for Computational Excellence (HEP-CCE)
    - Fermilab institutional lead, Liz Sexton-Kennedy
    - Focuses on HPC I/O systems and storage
    - Fermilab involvement
      - Measuring the performance of ROOT I/O on HPC (using the profiling tool Darshan)
        - Important because ROOT is the primary toolkit used in HEP
        - Must efficiently manage workflows for big and complex data

Scientific Computing Projects and Facilities

- Computing resources
- Advanced applications and algorithm techniques; large-scale data analytics
- Developed applications to run HPC platforms
  - HEP-Center for Computational Excellence (HEP-CCE)
    - Fermilab institutional lead, Liz Sexton-Kennedy
    - Focuses on HPC I/O systems and storage
    - Fermilab involvement
      - Measuring the performance of ROOT I/O on HPC (using the profiling tool Darshan)
      - Investigating Hierarchical Data Format (HDF5) as an intermediate event storage for processing data on HPC
        - Design to handle large amount of data
        - Organize the data in many different structures
        - Runs on a variety of computing platforms

Scientific Computing Projects and Facilities

- Computing resources
- Advanced applications and algorithm techniques; large-scale data analytics
- Developed applications to run HPC platforms
  - HEP-Center for Computational Excellence (HEP-CCE)
    - Fermilab institutional lead, Liz Sexton-Kennedy
    - Focuses on HPC I/O systems and storage
    - Fermilab involvement
      - Measuring the performance of ROOT I/O on HPC
      - Investigating Hierarchical Data Format (HDF5) as an intermediate event storage for processing data on HPC
      - Quantifying the disadvantages and advantages of running different data structures on GPUs

Scientific Computing Projects and Facilities

- Computing resources
- Advanced applications and algorithm techniques; large-scale data analytics
- Developed applications to run HPC platforms
- End user analysis toolkits beyond ROOT
  - Tools such as python, jupyter, julia, etc are becoming more popular
    - Well-maintained
    - Easy to use

ROOT Data Analysis Framework

https://root.cern.ch/

paradigm shift

CMS columnar-based python ecosystem

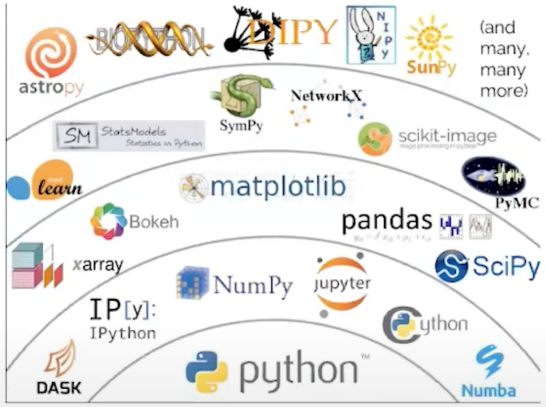

Scientific Computing Projects and Facilities

- Computing resources
- Advanced applications and algorithm techniques; large-scale data analytics
- Developed applications to run HPC platforms
- End user analysis toolkits beyond ROOT
  - CMS Elastic Analysis Facility (R&D project)
    - Column-wise analysis and supports row-wise analysis
    - Low-latency access to CPU resources
    - Fast access to large amount of disk
    - Containerized infrastructure (Kubernetes-based)
    - Jupyterhub deployment

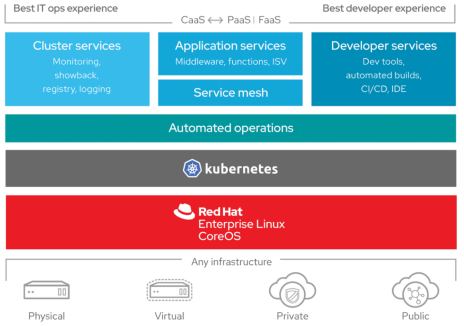

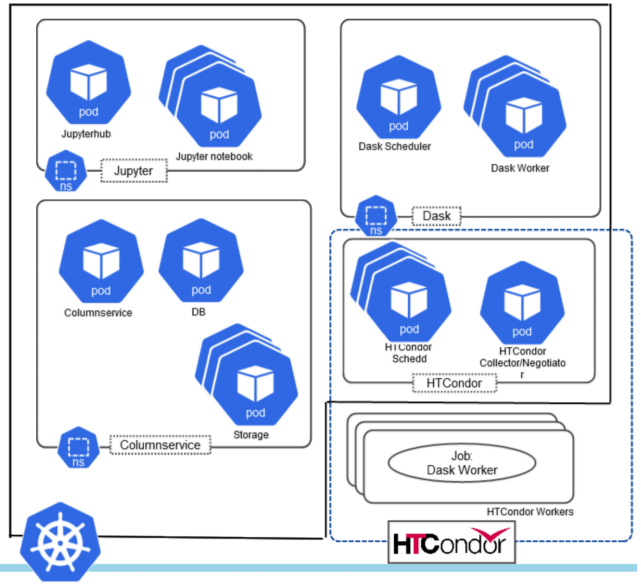

Outline

- Introduction
- Scientific Computing Services
- Scientific Computing Projects and Facilities
- Summary

"Computing is a catalyst for building bridges to other areas of science and society at large." – Snowmass Computation Frontier Conveners

Summary

- Fermilab offers a variety of services for experiments.
  - If you are unsure, please contact your Fermilab scientific computing liason for questions regarding SCD services.
- The next generation of HEP experiments are projected to produced big data.
  - High-Performance Computing will play a major role in data production and algorithm development.
- The community is adapting tools beyond ROOT.
  - Expect many analysis facilities to integrate tools (python, Julia,…) that are supported and developed by many fields within the scientific community.
- There are many ongoing computational R&D activities.
- HEP is not alone in facing the challenges of big data.

For more information on the future of scientific computing at Fermilab, please feel free to contact members within the SCD community

Thank you!

Special thanks to everyone who help with the contents of the talk

- Adam Lyon, Oliver Gutshe, Phil Demar, Andrey Bobyshev, Bo Jayatilaka, Burt Holzman, Lindsey Gray, Nick Smith, Lisa Goodenough, Kyle Knoepfel, Ken Herner

Back up Slides

Scientific Computing

https://computing.fnal.gov/

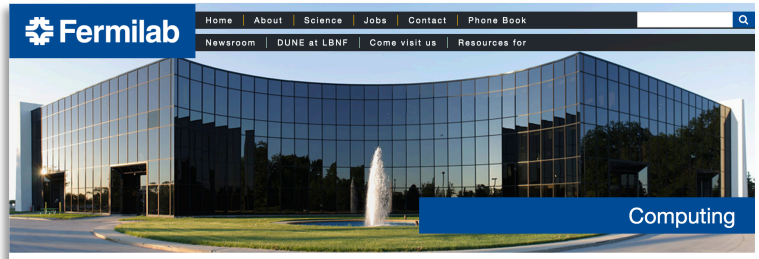

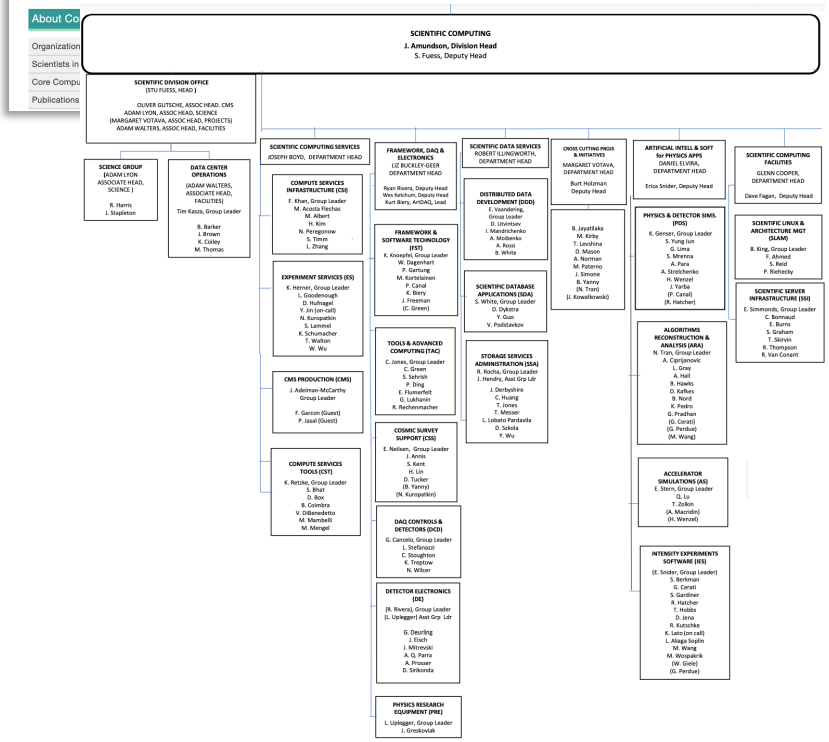

Scientific Computing Division consists of :

- computing scientists and researchers
- experimental, theoretical, and accelerator physicists
- engineers
- research associates

Scientific Computing Division contributes to:

- Accelerator simulations
- DAQ controls, detectors, and electronics
- Data storage, handling, distribution, and production
- Computing Infrastructure and architecture
- Software and frameworks
- Algorithms and reconstruction
- Physics Analysis
- And much more

Scientific Computing Projects and Facilities

- Computing resources
- Advanced applications and algorithm techniques; large data analytics
- Developed applications to run HPC platforms
  - HEP-Center for Computational Excellence (HEP-CCE)
    - Fermilab institutional lead, Liz Sexton-Kennedy
    - Focuses on HPC I/O systems and storage
    - Fermilab involvement

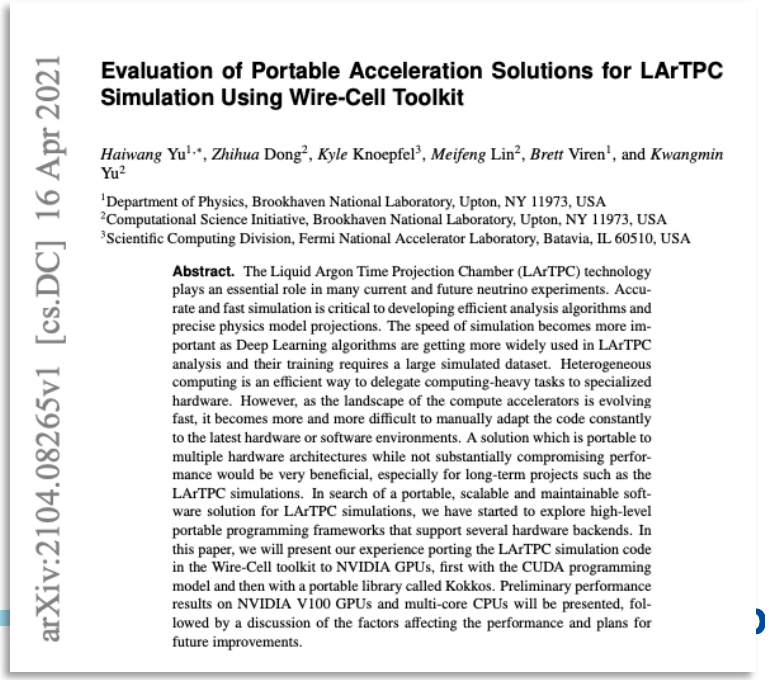

Evaluation of Portable Acceleration Solutions for LArTPC Simulation Using Wire-Cell Toolkit

Haiwang Yu[1,*], *Zhihua* Dong[2], *Kyle* Knoepfel[3], *Meifeng* Lin[2], *Brett* Viren[1], and *Kwangmin* Yu[2]

[1]Department of Physics, Brookhaven National Laboratory, Upton, NY 11973, USA
[2]Computational Science Initiative, Brookhaven National Laboratory, Upton, NY 11973, USA
[3]Scientific Computing Division, Fermi National Accelerator Laboratory, Batavia, IL 60510, USA

Abstract. The Liquid Argon Time Projection Chamber (LArTPC) technology plays an essential role in many current and future neutrino experiments. Accurate and fast simulation is critical to developing efficient analysis algorithms and precise physics model projections. The speed of simulation becomes more important as Deep Learning algorithms are getting more widely used in LArTPC analysis and their training requires a large simulated dataset. Heterogeneous computing is an efficient way to delegate computing-heavy tasks to specialized hardware. However, as the landscape of the compute accelerators is evolving fast, it becomes more and more difficult to manually adapt the code constantly to the latest hardware or software environments. A solution which is portable to multiple hardware architectures while not substantially compromising performance would be very beneficial, especially for long-term projects such as the LArTPC simulations. In search of a portable, scalable and maintainable software solution for LArTPC simulations, we have started to explore high-level portable programming frameworks that support several hardware backends. In this paper, we will present our experience porting the LArTPC simulation code in the Wire-Cell toolkit to NVIDIA GPUs, first with the CUDA programming model and then with a portable library called Kokkos. Preliminary performance results on NVIDIA V100 GPUs and multi-core CPUs will be presented, followed by a discussion of the factors affecting the performance and plans for future improvements.

Scientific Computing Projects and Facilities

- Computing resources
- Advanced applications and algorithm techniques; large data analytics
- Developed applications to run HPC platforms
  - HEP-Center for Computational Excellence (HEP-CCE)
    - Fermilab institutional lead, Liz Sexton-Kennedy
    - Focuses on HPC I/O systems and storage
      - Order of magnitude increase in data rates

Elizabeth Sexton-Kennedy
Chief Information Officer

https://www.anl.gov/hep-cce

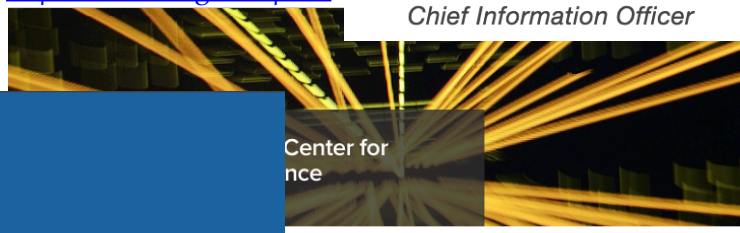

HEP-CCE All-Hands Meeting

Nov 4, 2020, 1:45 PM → Nov 6, 2020, 4:25 PM US/Central

Description Click Here to Join Meeting

The second HEP-CCE all-hands meeting will be dedicated to reviewing Year 1 progress for our four tasks, discussing Year 2 priorities and strategizing about possible new directions for our project. The idea is to have three two-hour sessions:

- Nov 4, 2-4 PM Central - Introduction, Task 2 (I/O Strategies) status and year 2 priorities
- Nov 5, 2-4 PM - Task 1 (Portable Parallezation) status and year 2 priorities
- Nov 6, 2-4 PM - Task 3 (Generators) and 4 (Workflows) year 2 priorities, CCE outreach. Year 2+ strategy discussion.

Scientific Computing Projects and Facilities

- Computing resources
- Advanced applications and algorithm techniques; large data analytics
- Developed applications to run HPC platforms
- End user analysis toolkits beyond ROOT
  - CMS Coffea python toolkit
    - Columnarbased analyses
    - Use by 10 teams and > 20 analyses
    - Backbone for many published analyses
      - http://cms-results.web.cern.ch/cms-results/public-results/publications/EXO-20-004/index.html
    - Analyses turn-around time from 1-2 weeks to 1-2 hours
    - Support a variety of computing clusters (condor, slurm,.)
    - And execution engines (dask, parsl,..)
    - Easy to implement data formats from others (ATLAS)
    - Tutorial videos are available
      - https://www.youtube.com/results?search_query=coffea+pyhep

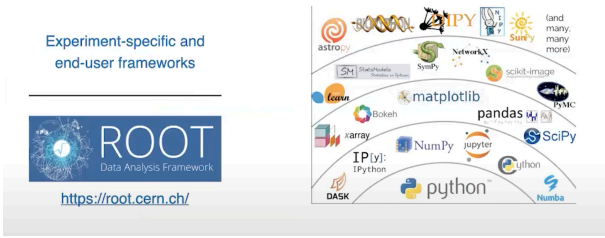

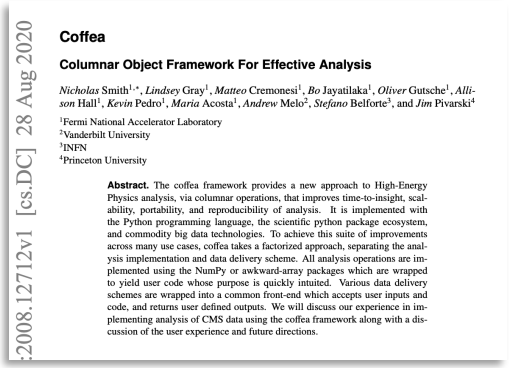

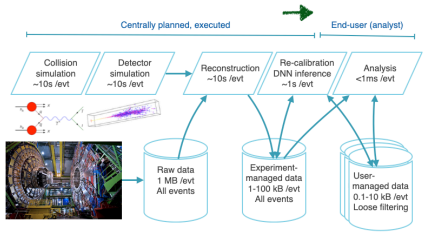

Scientific Computing Projects and Facilities

- Networking
  - Fast transportation of exascale data size around the world

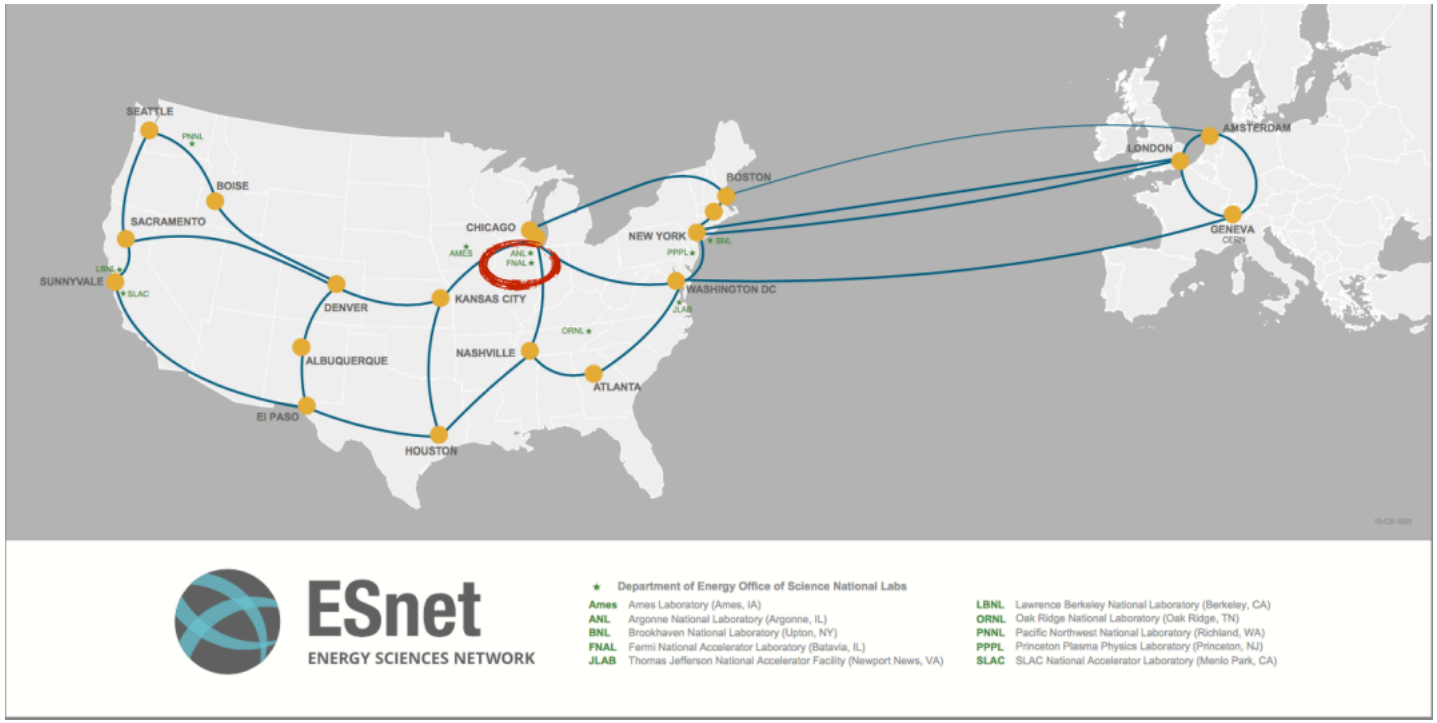

Scientific Computing Projects and Facilities

- Storage
  - R&D investment in scaling storage for upcoming scientific data
  - Current tape (managed by enstore) : 269 petabytes

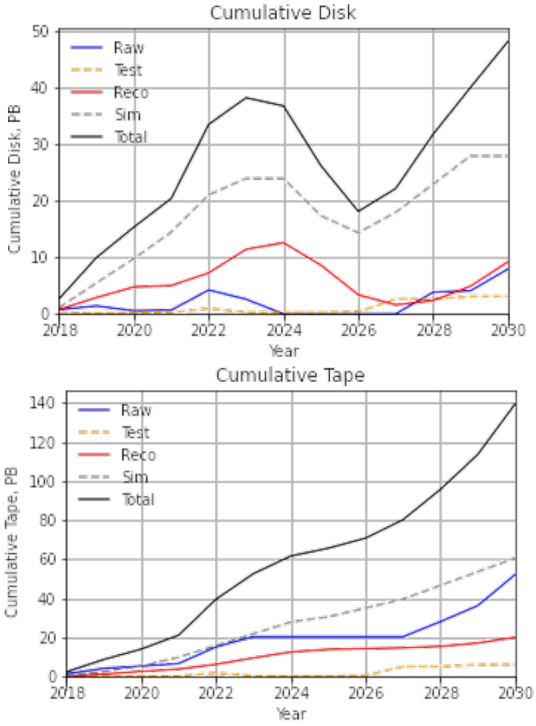

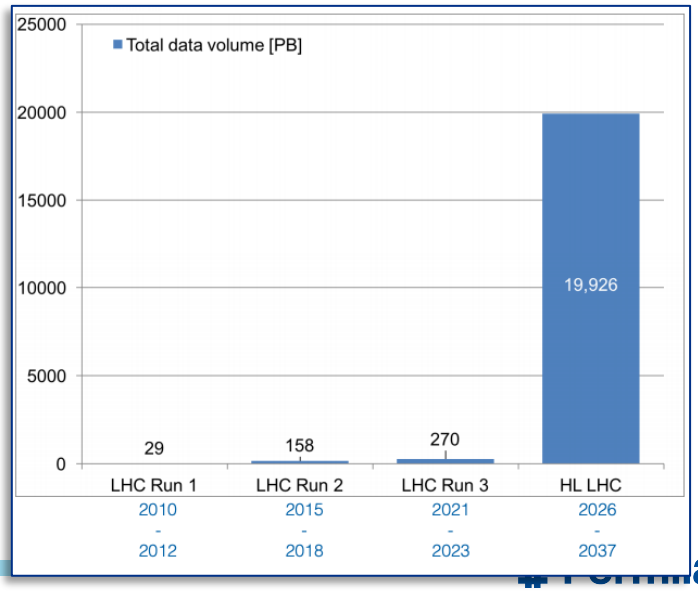

DEEP UNDERGROUND NEUTRINO EXPERIMENT

Scientific Computing Projects and Facilities

- Networking
  - Fast transportation of exascale data size around the world
    - HL LHC (~2027)
      - Current Projections are for US-CMS T1 Facility to require 1Tb/s – 1.5Tb/s of off-site bandwidth at start of HL-LHC:
      - Current FNAL off-site bandwidth is 3x100Gb/s for all off-site traffic
      - Soon to be 4x100Gb/s
      - Anticipating incremental ramp-up toward the terabit level in off-site bandwidth between now and 2027
      - Step-up from 100GE network technology to 400GE technology
      - US-CMS T1 internal network to be correspondingly scaled up:
      - 400GE network technology to become foundational basis of US-CMS T1 LAN
      - 100GE connectivity for server (& worker?) nodes
      - LHC Data Challenges to demonstrate ramp-up of CMS data movement capabilities
        - 2021 – 10%
        - 2023 – 30%
        - 2025 – 60%
        - 2027 – 100%

Scientific Computing Projects and Facilities

- Networking
  - Fast transportation of exascale data size around the world
    - DUNE (~2026-2027)
      - Support for remote operations of Far Detector (FD) is primary network challenge
      - Will follow same model as used to support NOvA and Soudan:
      - Remote LAN logically supported as being part of FNAL internal network
      - Lies within the FNAL network security perimeter
      - Fully redundant LAN at DUNE FD:
        - Includes redundant fiber paths between cavern (detectors…) and surface
        - Based on 400GE network technology
      - Geographically-diverse WAN circuits back to FNAL:
        - 100Gb/s guaranteed bandwidth on the primary path
        - 10Gb/s guaranteed bandwidth on the secondary path