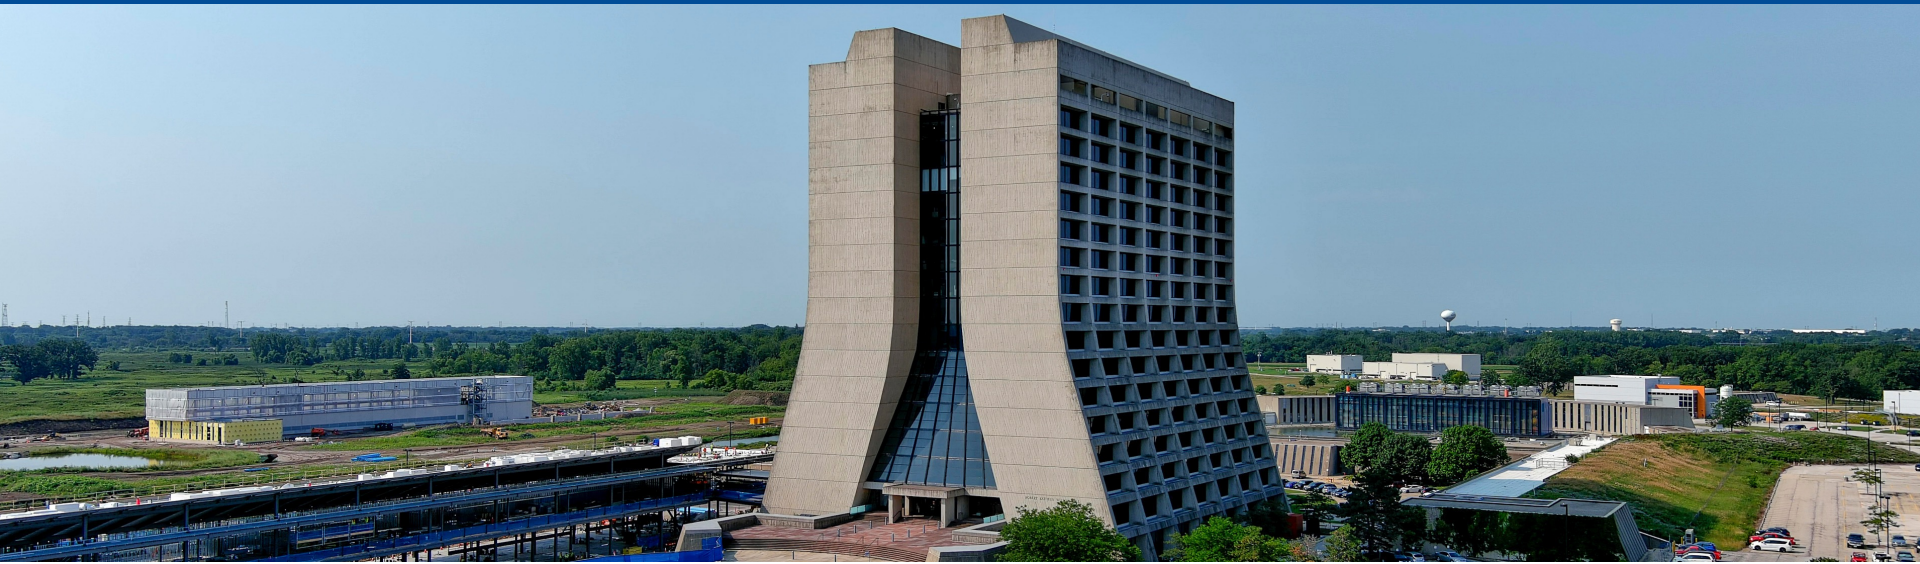

Strategic Planning for Software and Computing at the Laboratory

James Amundson

PAC Meeting June 5-9, 2023

June 8, 2023

Charge

Following up on the recommendation from the last PAC, we recommend that the Laboratory develops a strategic computing plan based on the experiments' needs. The plan should project computing and storage resources that will be made available to the experiments, as well as central services like power and cooling. The plan should include a timeline and milestones for critical technology decision points e.g., the integration of heterogeneous computing resources provided by US HPCs and the role of commercial clouds.

Outline

- Lab Strategy Pillars
- Elements of the Computing Strategy
- Background
- Strategy
  - Data Centers
  - Mass Storage
  - Compute Resources
  - Software Development
  - Data Analysis
  - Diversity, Equity, Inclusion, and Accessibility

Strategic Thrusts: Pillars of our vision for Fermilab

Deliver groundbreaking science and technology innovation

Building for Discovery: Project Execution

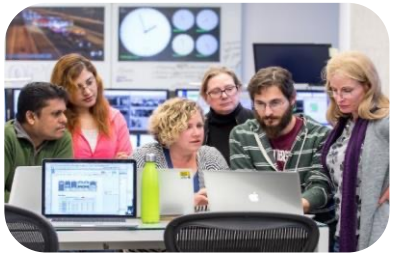

Diversify and empower our workforce

Transform business operations, infrastructure management and campus sustainability

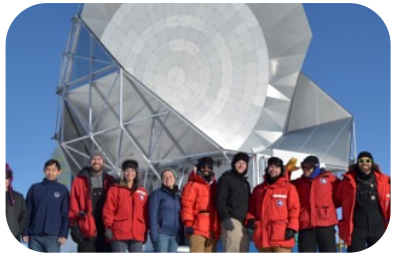

Forge strong alliances with regional, national & international institutions

Position Fermilab for a vibrant and successful future

Strategic Thrusts: Pillars of our vision for Fermilab

Deliver groundbreaking science and technology innovation

Building for Discovery: Project Execution

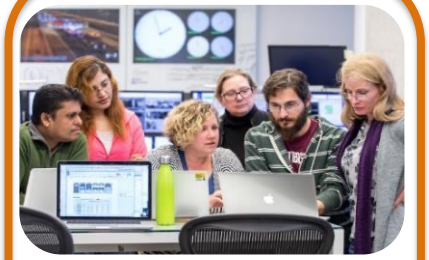

Diversify and empower our workforce

Transform business operations, infrastructure management and campus sustainability

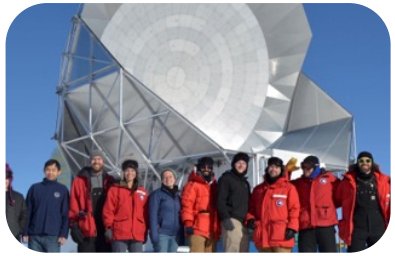

Forge strong alliances with regional, national & international institutions

Position Fermilab for a vibrant and successful future

Elements of the Fermilab Computing Strategy

Data Centers

Mass Storage

Compute Resources

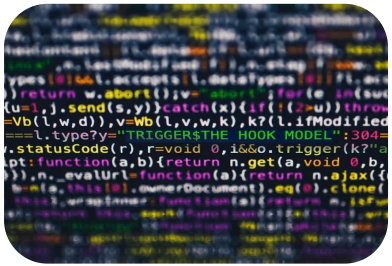

Software Development

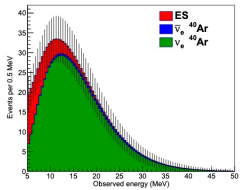

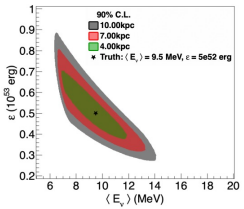

Data Analysis

Diversity, Equity, Inclusion, and Accessibility

CSAID Mission

The mission of the Computational Science and Artificial Intelligence Directorate is to advance the state of scientific computing for high energy physics and to provide computational services for the Fermilab scientific program.

Background

Understanding the problem: the experimental landscape

- Supporting HEP experiments is the largest portion of our scientific mission
- Schedules and data volumes are the most important inputs to planning
  - Computational requirements are (very) roughly proportional to data volumes

Planning through the end of the decade

Beam delivery plan

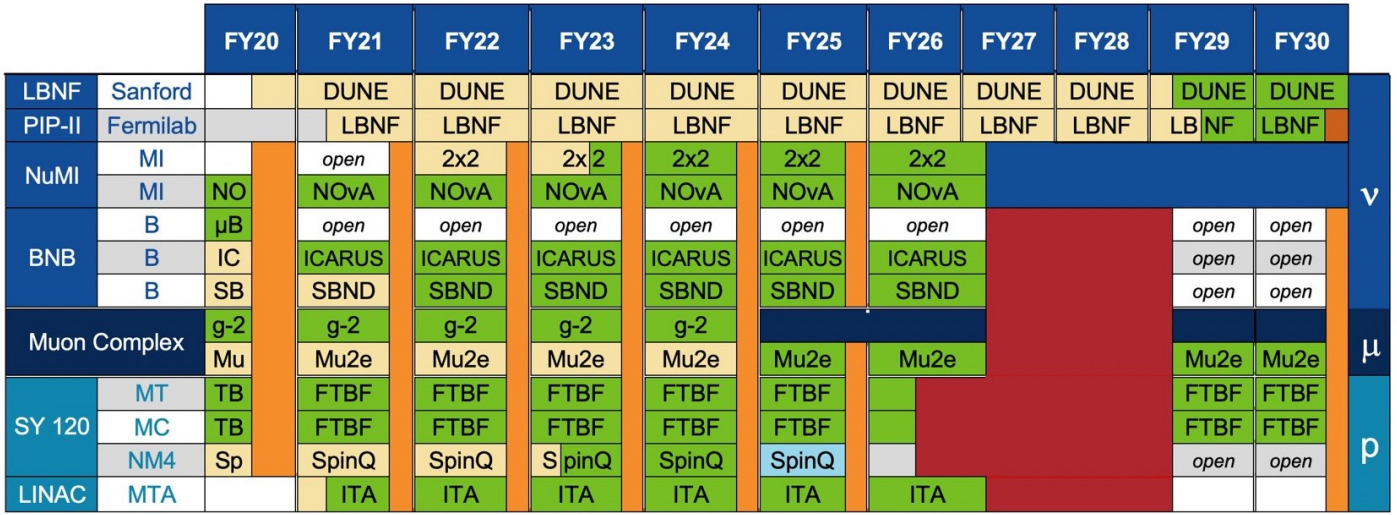

Fermilab

Including CMS

Fermilab

USCMS and DUNE Storage Projections

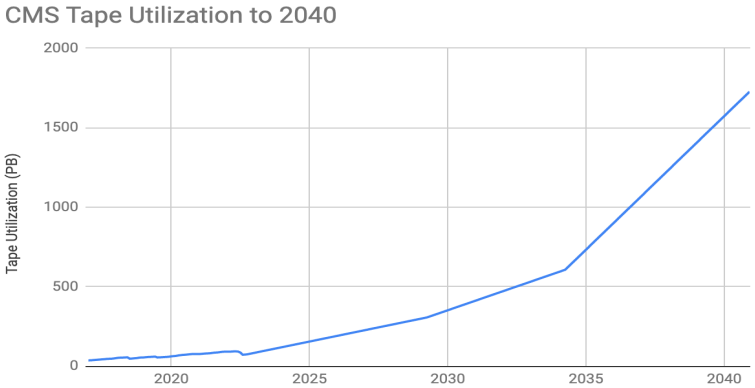

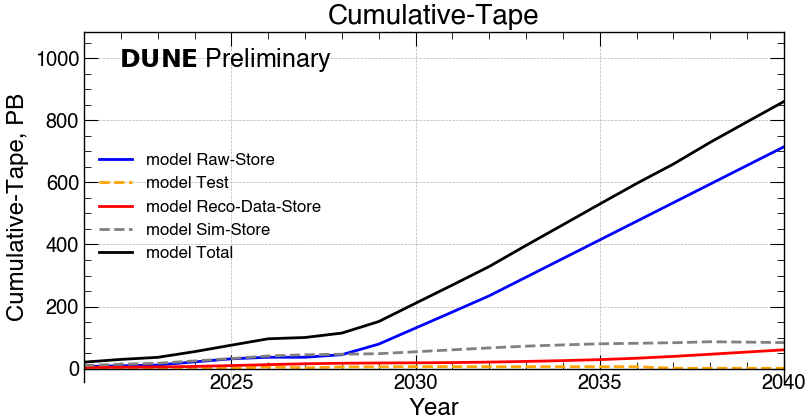

Fermilab USCMS tape storage forecast
Roughly 700 PB at Fermilab in 2035
(Global is 4x Fermilab)
Uncertainties are still in the 2x range

Global DUNE tape storage forecast
Roughly 250 PB at Fermilab in 2035
(Fermilab is 1/2 Global)
Uncertainties are still large, especially for the Near Detector

Compute usage much more difficult to predict, but roughly assume tracking storage
DUNE compute/storage overall lower than CMS compute/storage

Other Storage

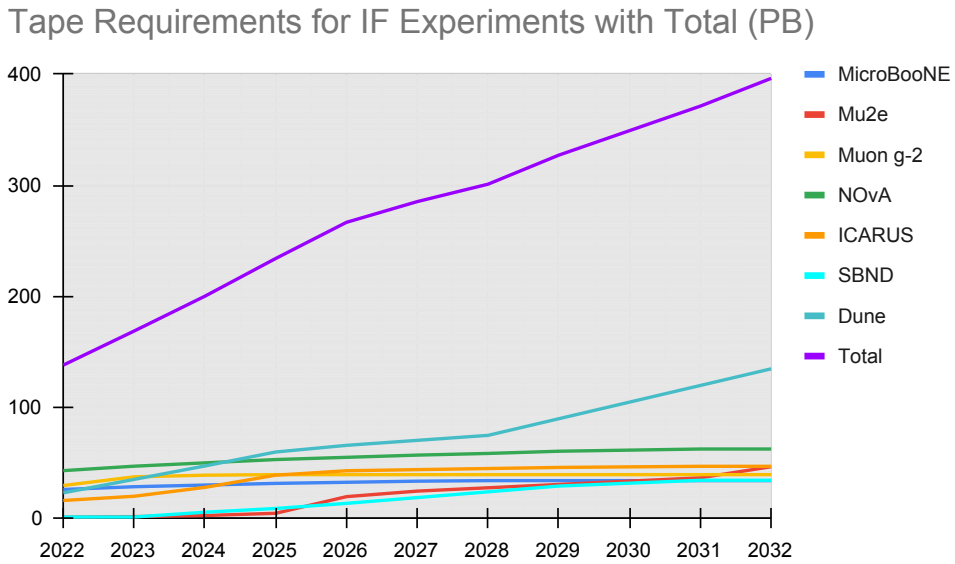

- DUNE is not yet the largest component of Fermilab non-CMS storage, but will be soon
- Other experiments generate significant volume
  - Will eventually be dominated by DUNE

Strategy

Data Center Strategy

Ensure that Fermilab has the data center capacity to host the necessary storage and compute resources necessary to meet the experimental needs for the next decade, with an outlook the the following decade. The data centers must be cost-effective, sustainable, and provide the redundancy necessary to fulfill DUNE host lab responsibilities.

- Current facilities
  - Grid Computing Center (GCC)
  - Feynman Computing Center (FCC) 2nd and 3rd floor data centers

Data Center Plans

- Objective: Consider off-site hosting options
  - Economics are unlikely to favor off-site hosting
  - Security and logistics also argue against off-site solutions
  - Continue to monitor
- Objective: Study future facility needs
  - Task force (CREST) currently making detailed study of future facility needs
  - **Milestone**: Initial draft of Computing Facility Evolution Plan scheduled to be released before the end of August 2023 – milestone from the Lab Agenda.
- Objective: Maintain existing infrastructure
  - HVAC infrastructure in existing computer rooms reaching end of life soon
  - Upgrade project in lab GPP list
  - **Milestone**: GPP for HVAC upgrade approval in FY2024 or FY2025
- Objective: Water cooling for compute resources to be evaluated
  - **Milestone**: Determine need based on Computing Facility Evolution Plan. Early FY2024
  - **Milestone**: Engineering estimates. FY2024
  - **Milestone**: GPP for cooling infrastructure. FY2025 or FY2026

Mass Storage Strategy

High Energy Physics has uniquely large data storage and processing requirements. Fermilab will provide systems capable of scaling to meet the needs of USCMS, DUNE and other Fermilab experiments through the next decade with a plan for the following decade. DUNE mass storage should be brought into parity with USCMS mass storage.

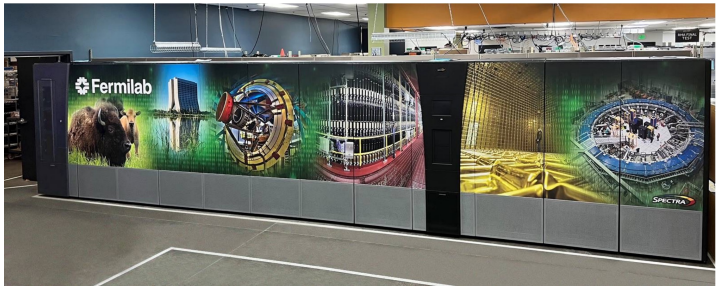

- Fermilab currently operates two segregated mass storage systems
  - CMS
  - “Public”
- The tape portion relies on Enstore, which was developed for Tevatron Run-II
  - The limitations of Enstore are rapidly emerging
- Isolation between experiments in the Public system is poor

Mass Storage Plans

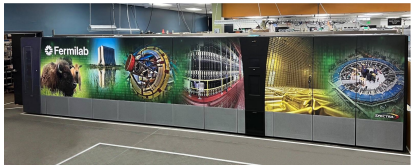

- Objective: Evaluate cloud storage options
  - Cloud storage options are available
  - Not currently economically viable
    - Data access fees would rapidly dominate
    - Concerns about long-term ownership
  - **Milestone**: Revaluate biennially
- Objective: Migrate Fermilab tape storage to the CERN Tape Archive (CTA) software
  - Working jointly with CERN on CTA development
  - **Milestone**: Testing of CTA software for USCMS system at Fermilab to begin this summer – milestone from the Lab Agenda
- Objective: Establish parity between USCMS storage and DUNE storage
  - **Milestone**: Evaluate isolation possibilities with CTA-based storage System. Early FY2025
  - **Milestone**: Decide optimal course between unified storage system and three-component system: USCMS, DUNE, and other. Late FY2025
  - **Milestone**: Implement chosen system. Completion in FY2027

Compute Resource Strategy

Provide experiments access to the computational resources needed to process their data efficiently. Minimize the effort of experiments to access resources, even heterogenous ones.

- Most complex portion of the computing strategy
  - Plan for mixture of local resources, HPC, grid and cloud
  - Commercial market is rapidly evolving
  - Compute resource strategy is closely tied to software development strategy

~~Computing~~ ~~CPU~~ Compute Resource Strategy

Provide experiments access to the computational resources needed to process their data efficiently. Minimize the effort of experiments to access resources, even heterogenous ones.

- Most complex portion of the computing strategy
  - Plan for mixture of local resources, HPC, grid and cloud
  - Commercial market is rapidly evolving
  - Compute resource strategy is closely tied to software development strategy

Compute Resource Plans

- Objective: Provide uniform interface to compute resources
  - HEPCloud will remain our approach to providing access to heterogeneous resources
- Objective: Provide local resources
  - Annual resource allocation
    - Determined by the Fermilab Computing Resource Scrutiny Group (FCRSG) process
    - **Milestone**: Resource allocation determined by FCRSG process. Annual (9/2023 for FY2023) – milestone from the Lab Agenda
  - Annual refreshes of CPUs and GPUs
    - **Milestone(s)**: CPU refreshes.
      - Annual for DUNE and other Intensity Frontier experiments, driven by FCRSG process
      - USCMS CPU driven by USCMS Software and Computing.
  - Strategic GPU purchases
    - **Milestone:** Annual GPU purchases at the level of $500k/year until experiment demand grows to be part of mainstream request

Compute Resource Plans

- Objective: Provide access to cloud resources
  - Cloud resources are currently economically viable for peak periods or special pricing opportunities and,
    possibly, access to exotic resources
    - HEPCloud for uniform access
    - **Milestone**: Annual evaluation of cloud economics
- Objective: Provide access to the Open Science Grid
  - Continue to participate in Open Science Grid Consortium

Compute Resource Plans

- Objective: Provide access to HPC
  - We anticipate utilizing HPC resources to
    - Allow large-scale analyses not otherwise possible, e.g., NOvA's Feldman-Cousins analysis
    - Reduce the need for local resources
  - We will continue to provide HEPCloud for uniform access
  - GPU development is required for newest HPC resources
    - See HEP CCE
  - Other developments necessary for HPC to come from HEP CCE
    - **Milestone**: New HEP CCE proposal. August 2023
  - Fermilab plans do not rely on a specific level of HPC utilization

Software Development Strategy

Fermilab, leveraging its unique position as a national lab, will provide specialized expertise to experiments. Our support will encompass aiding experiments to adapt to evolving computing technologies and delivering core software products. Fermilab will supply the core software development necessary to ensure the success of DUNE.

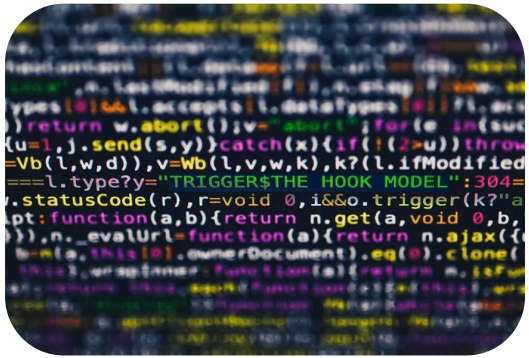

- Historically, lab contributions to experiment software included infrastructure and framework, but nothing beyond
- Better support for software throughout the entire experiment software stack is necessary to deal with technology evolution

Software Development Plans

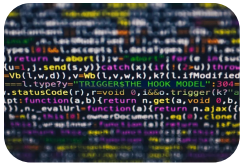

- Objective: Provide DUNE with a software framework
  - Existing frameworks are based on a collider model with well-defined beam crossings. Some DUNE requirements (e.g., responding to supernova events) do not map well into this model
  - Research new approaches to framework software for DUNE
    - **Milestone**: Completion of framework research LDRD. End of FY2023
  - Establish experiment needs
    - **Milestone**: Host framework workshop. Occurred this week– milestone from the Lab Agenda
  - Develop framework
    - **Milestone**: Framework design document. Early FY2024
    - **Milestone**: Deliver working framework. Date to be determined after requirements are set
- Objective: Support DUNE reconstruction software
  - Fermilab computing must become an active participant in DUNE reconstruction development
    - **Milestone**: Hire DUNE scientist in Data Science Simulation and Learning Division. FY2024
      - Funding will be a challenge

Software Development Plans, continued

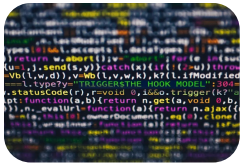

- Objective: Provide software development support to utilize GPUs
  - CPUs are the new Fortran
  - GPU-enabled software necessary to utilize emerging generation of HPC
  - GPU usage is, in principle, more energy efficient
  - **Milestone**: Hire two GPU software experts. FY24.
- Objective: Engage in software development to enable HPC utilization
  - HEP CCE is a multi-lab DOE-funded project to enable HPC utilization in HEP
  - Supports experiment adoption of GPUs
    - Not HPC-specific
  - Supports data handling and workflow efforts
  - **Milestone**: New multi-year HEP CCE proposal. August 2023

Software Development Plans, continued[2]

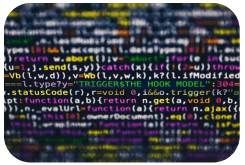

- Objective: Engage with broader scientific computing committee
  - Bring in a deputy ALD for CSAID from the ASCR community
    - **Milestone**: Job ad to be posted in June
- Objective: Support AI-based developments in experiment software
    - For details, see AI presentation
    - **Milestone**: Host AI Infrastructure Workshops. First one was April 6, 2023 – milestone from the Lab Agenda
- Objective: Support event simulation software
  - Neutrino generator support critical for DUNE
    - **Milestone**: Hire postdoc for recently promoted Steven Gardiner, Associate Scientist working on neutrino generators. FY2023
    - Building working relationship with neutrino generator efforts in Theory Division

Data Analysis Strategy

Fermilab will provide resources that empower scientists to examine their data efficiently, thereby reducing the duration to reach significant discoveries.

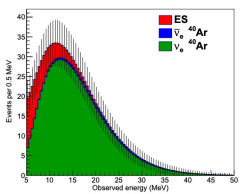

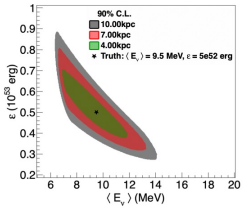

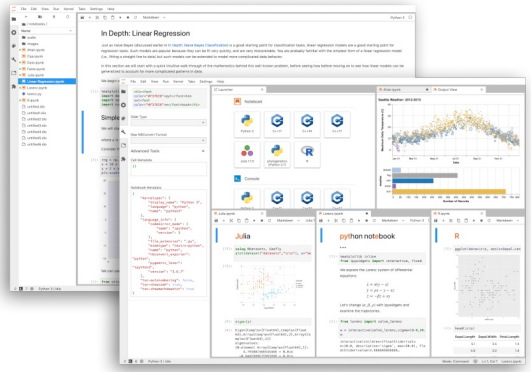

- Historically, analysis systems have been an ad hoc mixture of batch computing and local resources
- Reducing time-to-insight will improve scientific output

Fermilab

Data Analysis Plans

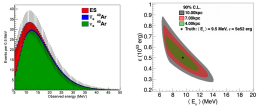

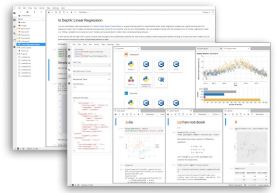

- Objective: Deploy dedicated Elastic Analysis Facility (EAF) for all experiments
  - Provides modern interface based on Jupyter notebooks
  - Includes CPU and GPU resources
  - **Milestone**: EAF in limited release. Done
  - **Milestone**: EAF in full production mode. FY2024
  - **Milestone**: Reevaluate fraction of local computing resources dedicated to EAF. Annual
- Objective: Support analysis at HPC centers
  - HPC centers enable types of analysis as scale that would not be available otherwise
  - Continue to support through HEPCloud

Diversity, Equity, Inclusion, and Accessibility Strategy

Fermilab is committed to cultivating a workforce that embraces diversity and fosters an inclusive environment.

- Work with the lab EDIA initiatives
- Emphasize strategies that have had demonstrated successes
- Fill the pipeline

Diversity, Equity, Inclusion, and Accessibility Plans

- Objective: Work with Lab EDIA efforts
  - Senior leaders all participate in Senior Leadership EDI Council
  - Individual EDIA contributions throughout the organization
- Objective: Recruit through professional organizations
  - Past participation in the Grace Hopper Conference has succeeded in recruiting more women into computing
  - **Milestone**: Participate in the 2023 Grace Hopper Conference. September 2023
  - **Milestone**: Find similar conferences for URM professional organizations and participate. FY2024
- Objective: Enhance summer student/intern pipeline
  - Assisted by new Fermilab Alumni Database
  - **Milestone**: Hire permanent employees who have come in as interns. By FY2025
- Objective: Use recruiting firms to identify more candidates from underrepresented groups.
  - Two of our most prominent women in computing came in through recruiting firms
  - Organizational issues need to be overcome

Summary: Fermilab Scientific Computing Strategy

Data Centers

- Ensure that Fermilab has the data center capacity to host the necessary storage and compute resources necessary to meet the experimental needs for the next decade, with an outlook the the following decade. The data centers must be cost-effective and sustainable and provide the redundancy necessary to fulfill DUNE host lab responsibilities.

Mass Storage

- High Energy Physics has uniquely large data storage and processing requirements. Fermilab will provide systems capable of scaling to meet the needs of USCMS, DUNE and other Fermilab experiments through the next decade with a plan for the following decade. DUNE mass storage should be brought into parity with USCMS mass storage.

Compute Resources

- Provide experiments access to the computational resources needed to process their data efficiently. Minimize the effort of experiments to access resources, even heterogenous ones.

Software Development

- Fermilab, leveraging its unique position as a national lab, will provide specialized expertise to experiments. Our support will encompass aiding experiments to adapt to evolving computing technologies and delivering core software products. Fermilab will supply the core software development necessary to ensure the success of DUNE.

Data Analysis

- Fermilab will provide resources that empower scientists to examine their data efficiently, thereby reducing the duration to reach significant discoveries.

Equity, Diversity, Inclusion, and Accessibility

- Fermilab is committed to cultivating a workforce that embraces diversity and fosters an inclusive environment.

Summary: Fermilab Scientific Computing Strategy

Data Centers

- Ensure that Fermilab has the data center capacity to host the necessary storage and compute resources necessary to meet the experimental needs for the next decade, with an outlook the the following decade. The data centers must be cost-effective and sustainable and provide the redundancy necessary to fulfill DUNE host lab responsibilities.

Mass Storage

- High Energy Physics has uniquely large data storage and processing requirements. Fermilab will provide systems capable of scaling to meet the needs of USCMS, DUNE and other Fermilab experiments through the next decade with a plan for the following decade. DUNE mass storage should be brought into parity with USCMS mass storage.

Compute Resources

- Provide experiments access to the computational resources needed to process their data efficiently. Minimize the effort of experiments to access resources, even heterogenous ones.

Software Development

- Fermilab, leveraging its unique position as a national lab, will provide specialized expertise to experiments. Our support will encompass aiding experiments to adapt to evolving computing technologies and delivering core software products. Fermilab will supply the core software development necessary to ensure the success of DUNE.

Data Analysis

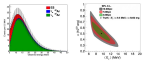

- Fermilab will provide resources that empower scientists to examine their data efficiently, thereby reducing the duration to reach significant discoveries.

Equity, Diversity, Inclusion, and Accessibility

- Fermilab is committed to cultivating a workforce that embraces diversity and fosters an inclusive environment.

Fermilab

Compute Resource Plans for HPC (backup)

- Objective: Provide access to HPC resources
  - Many barriers to be overcome
    - Allocations
      - LCF Allocation mechanisms are not compatible with HEP computing
      - Political problem, not a technical problem
      - Plan: Work with ASCR on long-term solution
    - Job submission/authentication, etc.
      - Plan: HEPCloud
    - Workflow management
      - Plan: HEPCloud
    - Data access
      - Not yet limiting
      - Addressed by CCE – see software development strategy
    - GPU Utilization
      - Addressed by CCE – see software development strategy
  - Plan: Work on all barriers. If they cannot be overcome, either obtain more operations funds for computing or limit experiment computing
  - Theoretical outcome
    - If enough resources were to come from HPC, we could avoid updating *some* of the cooling infrastructure in the GCC
      - Not enough information will be available in time
      - Would substantially limit lab capabilities
      - Leadership class HPC machines change on timescales shorter than data centers
        - Possibility of ending up in an impossible situation in out years
    - Not a viable strategy