Welcome and Charge

James Amundson
Margaret Votava
2023 Fermilab Computing Scrutiny Group
Feb 15th, 2023

Updates since 2022

- Response to FY22 recommendations
- FCRSG
  - Kept basic structure from last summer
  - No updates on computing models
  - Asked experiments for 5-year view
  - Only a few months since last year – returning to February cycle
- Had storage performance issues late summer
  - File sizes, organization of file families, tape library location all factors
  - Asking for campaign roadmap to better understand potential I/O contention issues.
  - Detailed review of existing monitoring with recommendations for improvement
- Lab organization
  - SCD promoted into a directorate level with three underlying divisions
- Mu2e computing review in the next few months

Fermilab Organization

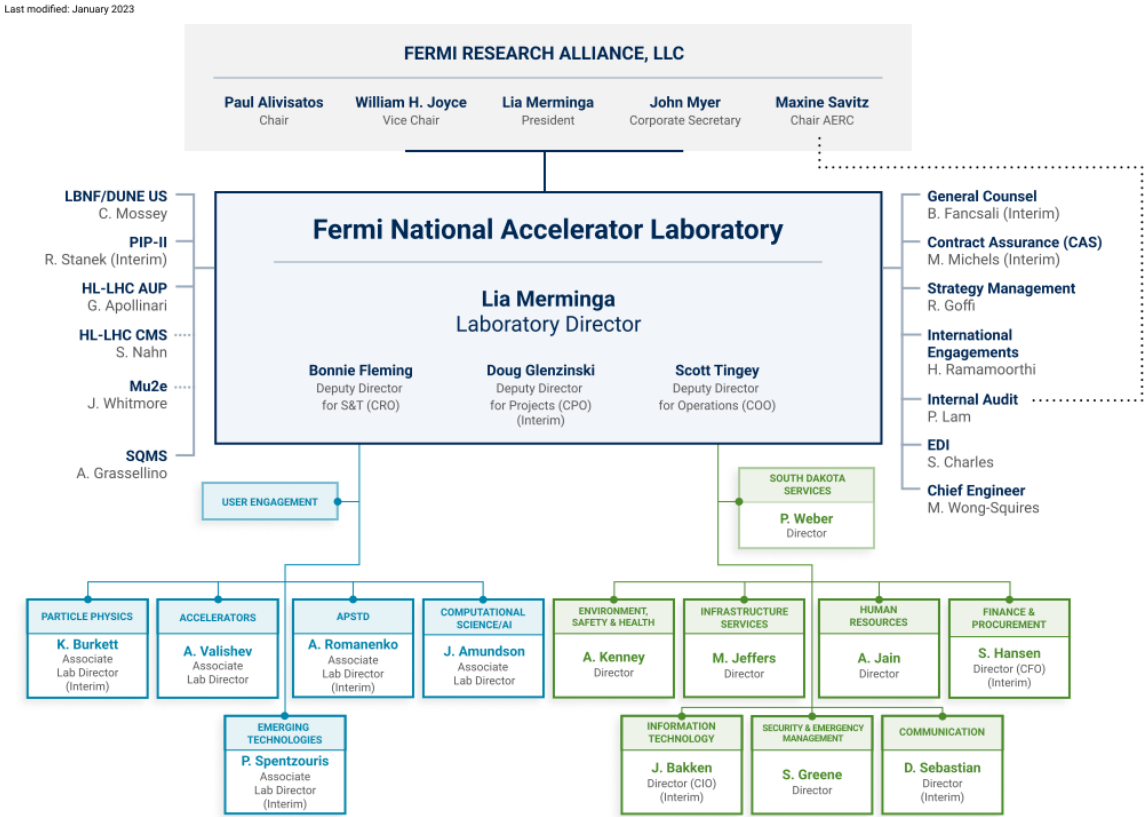

Computational Science and Artificial Intelligence Directorate

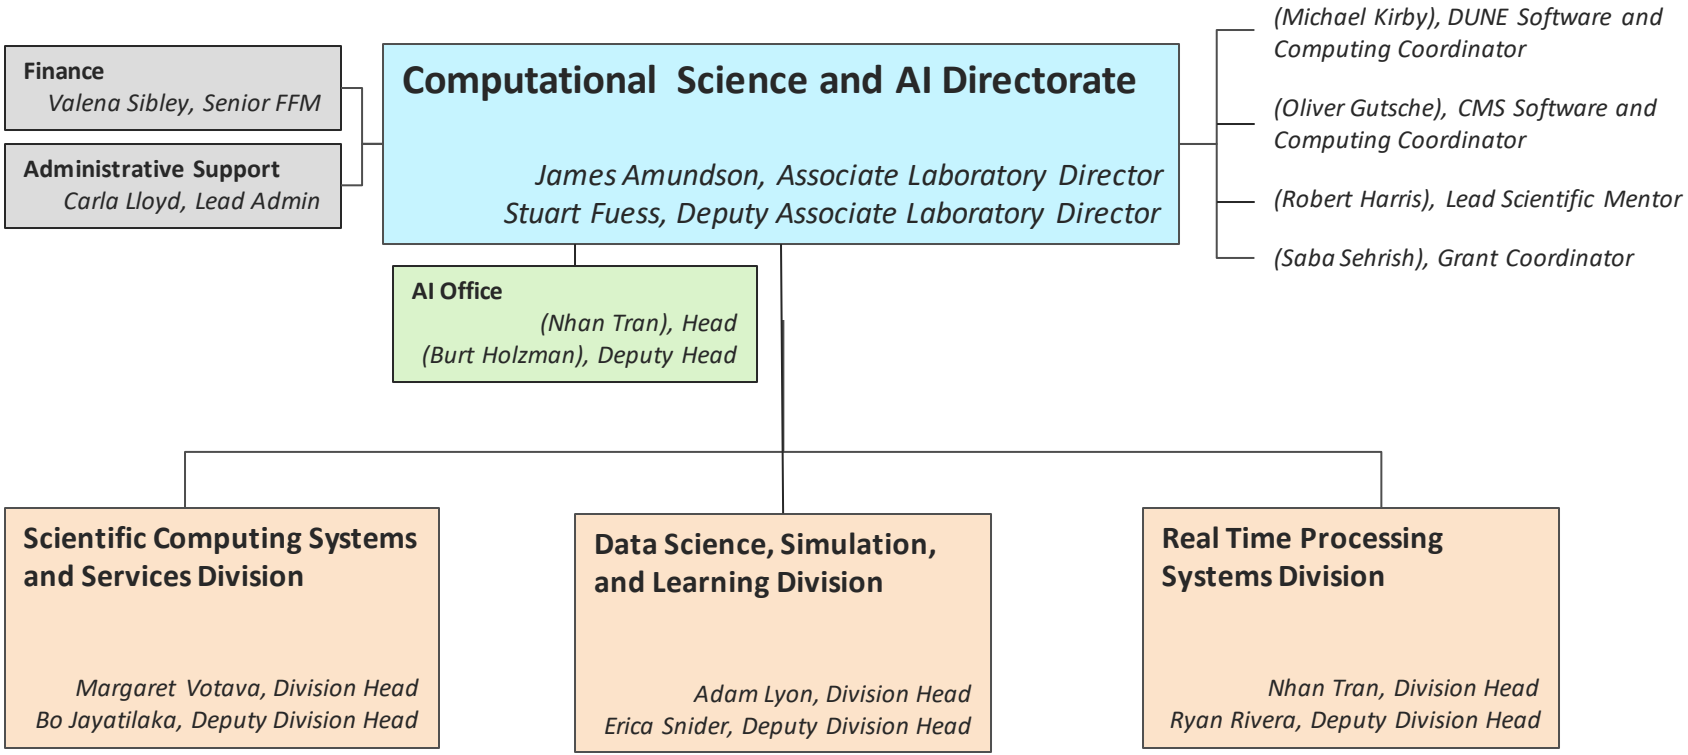

Fermilab Computing Strategy (hasn't changed since last review)

Detailed View

- Compute
  - Work towards getting a substantial fraction of US HEP computing from the Exascale/HPC machines
  - High-throughput computing at Fermilab will remain important
  - Utilize cloud resources when cost effective
    - Increase peak capacity
    - Access to non-standard hardware
- Storage
  - Mass storage will remain the foundation of Fermilab's computing capability
    - Tape libraries
    - Disk systems
    - Full-stack storage software
      - Including support for data lifetime management
- Analysis
  - Build an elastic analysis facility taking advantage of industry tools and Fermilab storage
- AI/ML
  - Provide GPU resources
  - Develop AI-enhanced operations
- Software
  - Pursue community-wide solutions
  - Engage ASCR partners
  - Collaborate with CERN
  - Leverage industry-standard tools where available

Fermilab

Facility modernization

- **C**omputing **R**esources **E**volution **ST**rategy (CREST)
  - ***Provide access to exascale-size resources for our experiments either on-site or off-site for the 2030s***
  - Document guiding principles and strategy
  - Chaired by Oli Gutsche
  - 1st draft document by August
  - Data Center big cost driver
- Elastic Analysis Facility (EAF) has moved [partially] into production in January
  - New Facility added to our service catalog
  - Some pieces of infrastructure still in R&D phase
  - Experiments are gaining experience. Will expect demand to grow and will understand more in 2023
  - Can expect to see resource requests in FCRSG 2024.
  - $500K GPU purchase in FY23. Planning will start soon

Facility modernization

- Storage
  - CTA
    - Making progress on conversion from Enstore to CTA
    - CMS will transition first, eta
    - Public will follow -- still need to sort out some small file handling
  - Rucio
    - Beta version of metadata catalog released -- adopted by DUNE
  - NAS –> CEPH
    - NAS falling off warranty
    - Small scale test of CEPH successful
    - Will start transition when new servers are production ready this spring

Facility modernization

- New DOE cybersecurity requirements
  - IPv6
  - Strong preference by site office for fedramp versions of all cloud services
- Transition to tokens
- In 2023, develop plan for deprecating older services
  - Redmine, ups

Funding Outlook

- Operations slightly better in FY23
  - No need to reduce existing staffing level, as would be required by a flat-flat budget
  - $0.5M for production GPUs
- New approach to carryover at Fermilab
  - Good news: will no longer require last-minute hardware purchases
  - Bad news: overhead needs to be used strategically across Computing and Detector Operations
- Recent (2022) push for enhanced operations funding from Office of Science
  - Starting in FY24
  - Billed as new normal
  - Requested DUNE analysis and operations support
  - Requested hardware and personnel for facility modernization
  - Roughly $4M (20%) additional funds requested
  - *Not yet demonstrated to be real*
- Expect DOE to create dedicated DUNE Operations funding this year
  - Includes computing
  - Work in progress – expect DUNE computing review this summer

Charge for this Review

- Goal: evaluate usage requests, plan for future of computing at Fermilab
  - Not limited to M&S – SWF is the largest portion of our budget
  - Not limited to Computing and Detector Operations
    - Note, however, that CMS computing is separately reviewed
  - Not including Scientific Software development
- Time frame: next year + experiment-specific horizon
- Scrutiny should focus on incremental costs (SWF + M&S)
  - Including custom solutions in software, etc.

Finally

- Thank you for your help
  - These reviews are very valuable to the division!

- Questions?