



Status of the Scientific Computing Program at the Laboratory

Elizabeth Sexton-Kennedy

Fermilab PAC

13 Jan 2020

Outline

- Response to July 2019 PAC recommendations
- Advisory Committees and the flow of information
 - Plans for migrating HEP computing to high performance architecture(s)
 - support to current and future experiments' operations
- Things I want to personally advocate for and ask committee advice
 - Sustaining community and facility software within DOE
 - Open data and data lifetime cycle management

Response to July 2019 PAC Recommendations

Recommendation 1: Computing Advisory Committee Structure

- Fermilab Computing has 2 advisory boards:
 - International Computing Advisory Committee (ICAC) - addresses high level strategic, programatic and planning issues
 - Fermi - Computing Resource Scrutiny Group (FCRSG) - addresses local resource planning and prioritization issues
- The first has met twice in Mar. and Oct. and is well established
 - The Oct. meeting evaluated the progress made with respect to the recommendations of the Mar. review. See Indico for a posting of their report.
- The cadence of the second group is once a year so it will meet in the beginning of Mar.
 - This is required to prepare the experiments for the new documentation process
 - For Mu2e and DUNE their PEMP notables align with a Mar. meeting.

Recommendation 2: Enhance proactive interactions with experiments to clarify their computing needs

- The lab has been charged through DOE PEMP notables to create a operations plan for computing manpower and resource requests.
 - By June 2020, submit a strategic plan for CMS High-Luminosity LHC software & computing R&D activities (Objective 1.2) **Oliver Gutsche**
 - Develop a preliminary Operations plan for the Mu2e experiment, ...including software and computing, including resource estimates, suitable for external review, by February 2020. (Objective 2.3) **Rob Kutschke**
 - By February 2020, develop an initial pre-Operations plan for the DUNE ...including software and computing. Include a preliminary resource estimate based where possible on extrapolations from prior comparable experiments. (Objective 2.3) **Mike Kirby**
- I would like to use these plans as templates for the nearer term experiments SBN and g-2. It's not clear who should champion SBN computing.

Operations Plans Status for February - Mu2e

- The drafts of the subsections for the “Preliminary Experiment Operations Plan” are due to Greg Rakness on Jan 15. Rob is writing the Data Processing and Computing chapter and has given it to Greg.
- The document has to be delivered by 25-Feb.
- A preliminary version of the Computing WBS has been merged into the overall WBS.

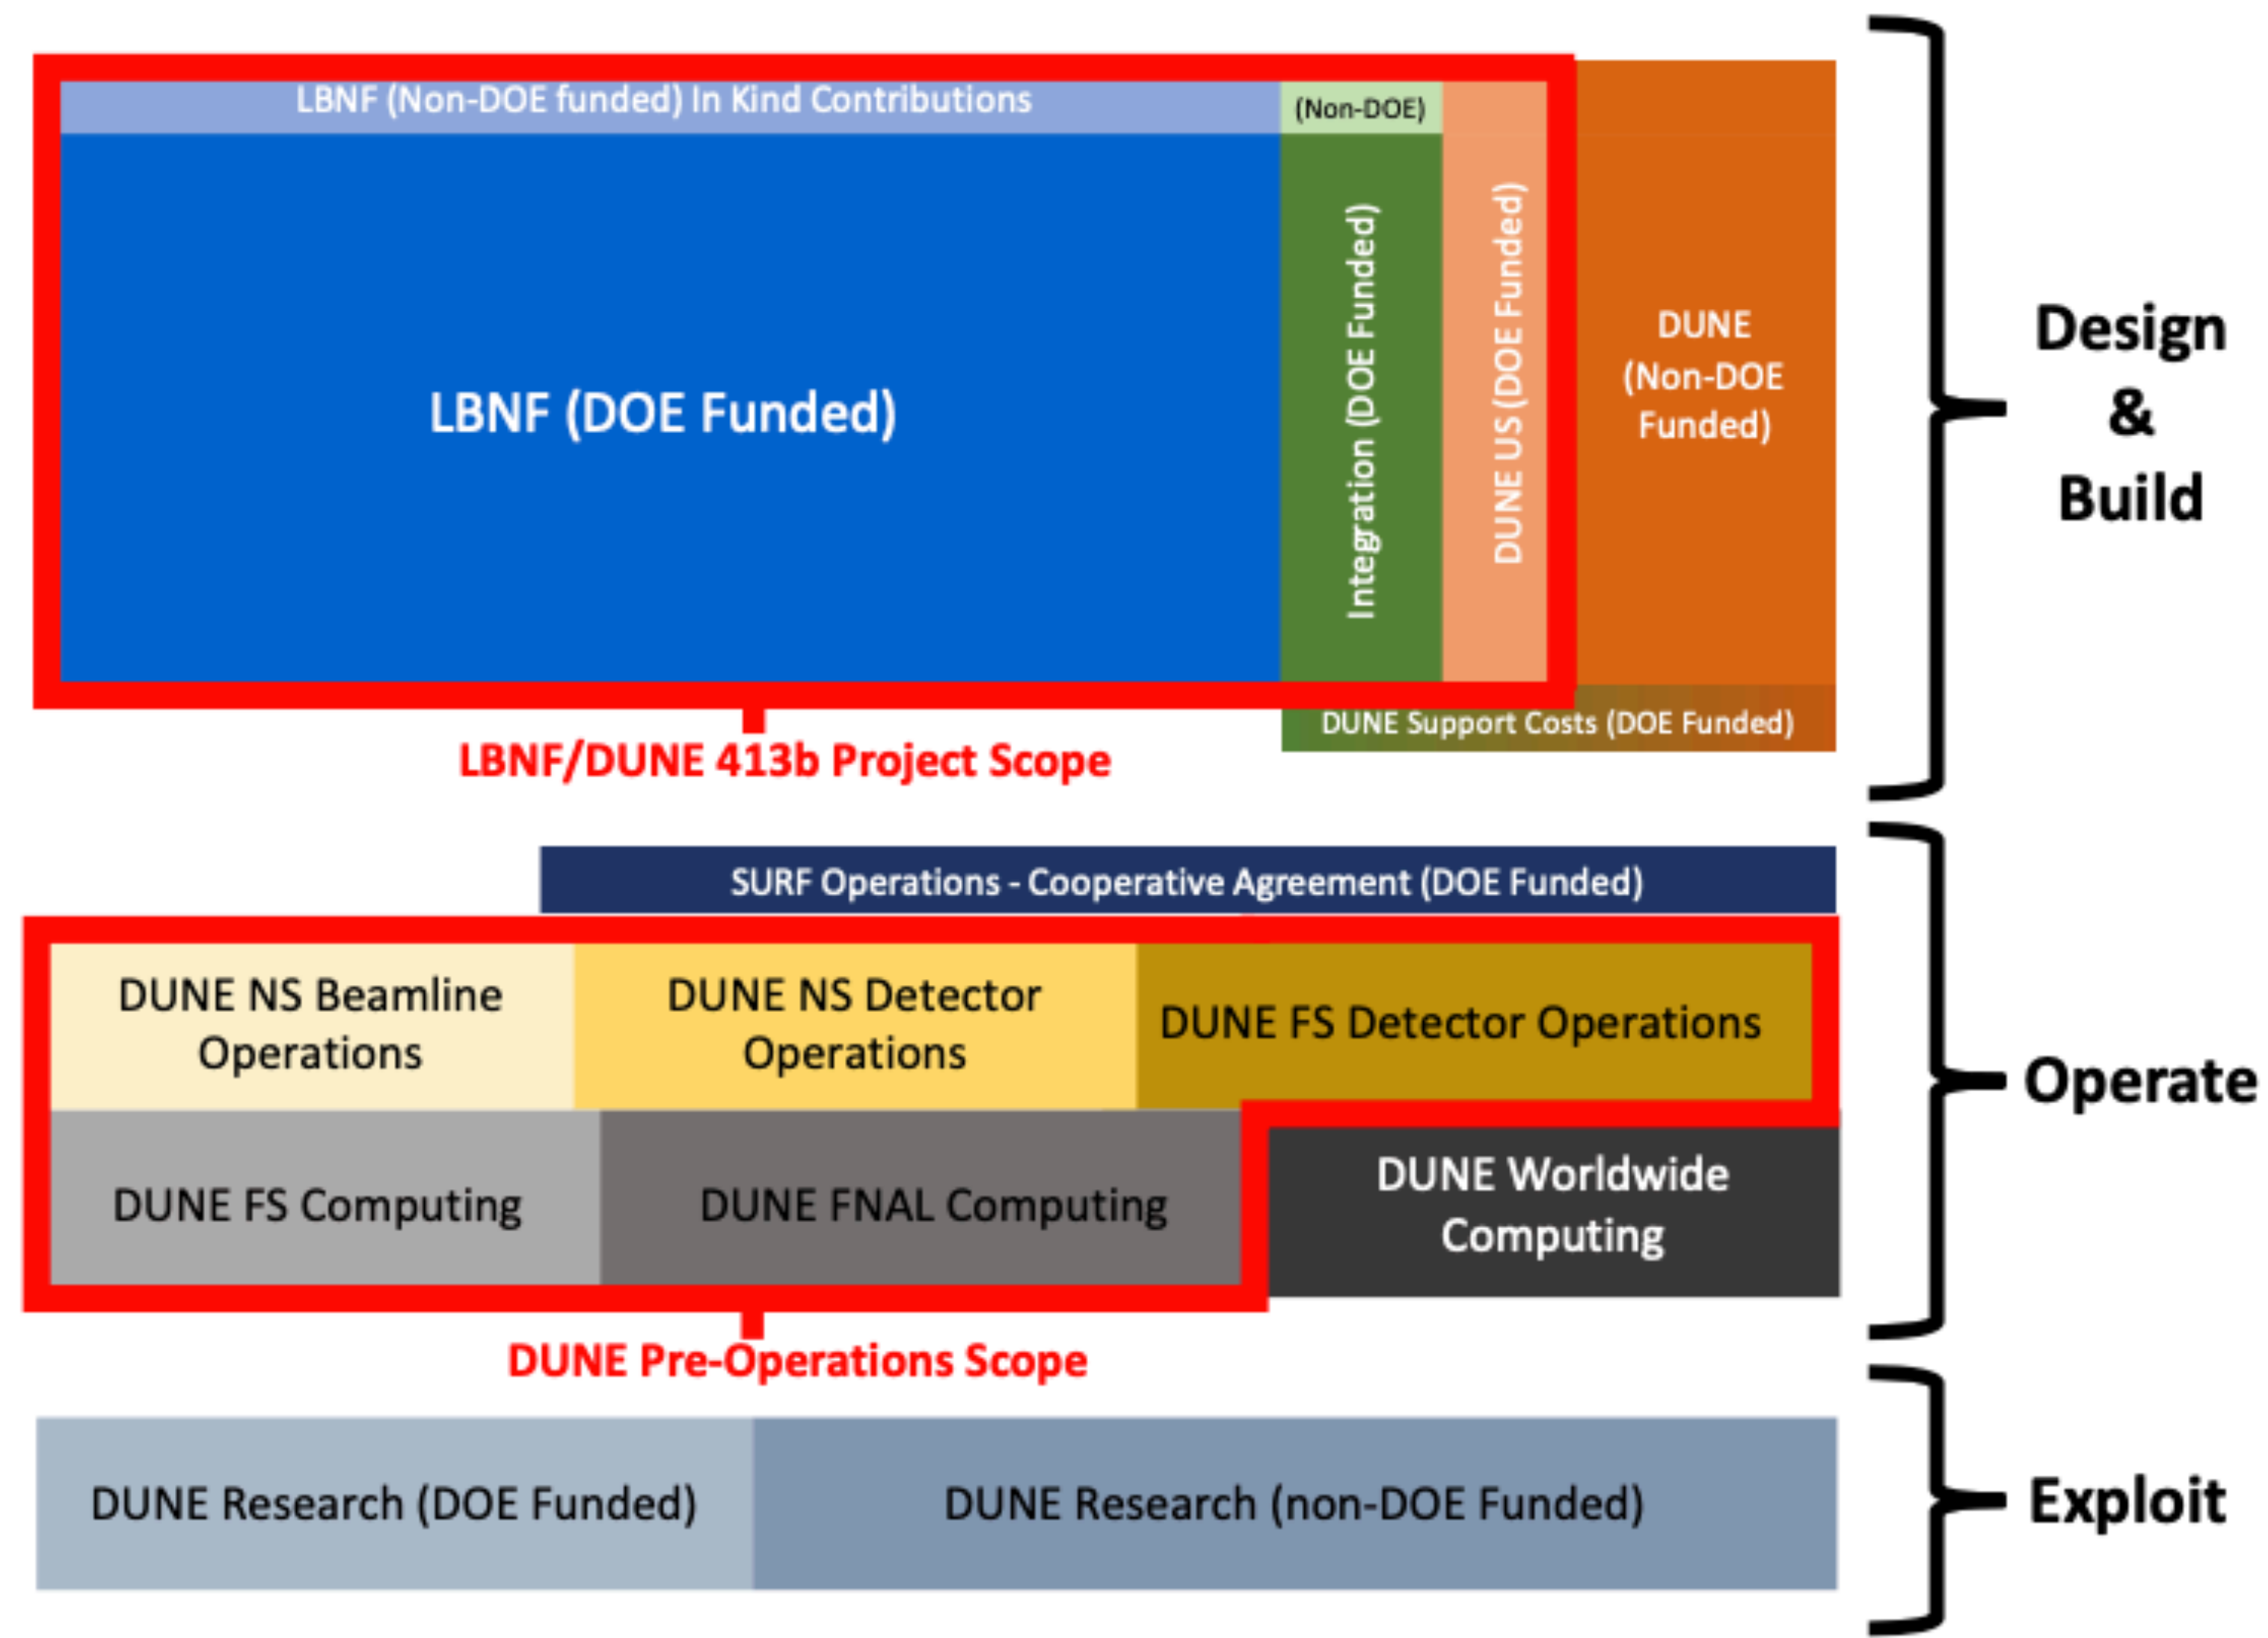
Table of Contents

1. Introduction	4
2. Mu2e Science	5
3. Mu2e Apparatus	5
4. Operations Roles and Responsibilities	5
4.1 Fermilab Organization	5
4.2 Collaboration Organization	7
4.3 Mu2e Operations Planning Team	9
5. Operations Scope of Work	10
6. Data Processing and Computing	10
7. Schedule Overview	11
8. Preliminary Budget and Resources	11
9. Operations Risks	11
10. Planning for Operations in FY2026 and Beyond	11
11. References	11

Operations Plans Status for February - DUNE

The DUNE Universe *(Brice)*

- Mike Kirby's charge: Develop a timeline of annual M&S and SWF for computing for each year from FY20 to FY30 and the story to go behind it



Fermilab

Recommendation 3: Inter-collaboration Information Transfer and Continued Education and Workforce Development


- In Sep. Fermilab hosted:
 - DUNE computing model workshop: <https://indico.fnal.gov/event/21231/>
 - WLCG SLATE security working group - Sept 10: <https://indico.fnal.gov/event/21485/>
 - WLCG pre-GDB: <https://indico.cern.ch/event/739896/>
 - WLCG Grid Deployment Board – Sept 11: <https://indico.cern.ch/event/739882/>
 - FIM4R – Sept 12: <https://indico.cern.ch/event/834658/>
 - IRIS-HEP blueprint workshop – Sept 12, Sept 13: <https://indico.cern.ch/event/840472/>
- Fermilab will host the next Rucio workshop in the second week of March 2020.
- Intend to make plans for a repeat of the successful C++ training course next summer.

C++ Training at Fermilab - Evaluations

C++ Course Evaluation Form

Q1 What part of the Fermilab program do you participate in?

Answered: 30 Skipped: 0




ANSWER CHOICES	RESPONSES	
Precision Muon	20.00%	6
Neutrino	26.67%	8
CMS	33.33%	10
Other (please specify)	20.00%	6
TOTAL		30

C++ Course Evaluation Form

Q9 How likely are you to recommend this workshop to a friend or colleague?

Answered: 30 Skipped: 0



ANSWER CHOICES	RESPONSES	
Very Likely	70.00%	21
Likely	30.00%	9
Neutral	0.00%	0
Unlikely	0.00%	0
Very Unlikely	0.00%	0
TOTAL		30

Recommendation 4: Prioritize Software R&D Efforts

- Software R&D efforts in FY20 are mostly funded competitively
 - Internal LDRDs
 - OHEP (Center for Computing Excellence-CCE)
 - OHEP & ASCR (SciDAC, Exa.Trx)
- Programmatic funding from CompHEP has been cut 80% but somewhat compensated by getting 30% back from CCE. Still this represents a 50% change in funding for R&D at Fermilab between FY19 and FY20.
- CMS contributions to R&D is heavy on development as is appropriate for an operations program. The national program sets its own priorities.
- Open calls like LDRDs and CCE are more helpful than highly targeted calls for proposals as favored by ASCR.

Recommendation 5: Engage ASCR More

- The CCE proposals are joint with ASCR and do address our most pressing R&D needs:
 - Portable Parallelization Strategies
 - Fine-Grained I/O and Storage
 - Event Generators
 - Complex Workflows (for Cosmic Frontier)
- These topics cover the lab's traditional strengths with the exception of simulation.
- We are trying to engage their help in creating a GPU enabled Geant application. Tom Evens of ORNL would be the PI and he is visiting this week to work out details.
 - All agree that Tom's approach is high risk and high reward
 - Continuing the less risky approach championed by LBNL could be a backup

Recommendation 6: Resource Allocation between HPC and Conventional Computing for the Near Future

- In the near term HPC resources are allocated in well defined programs not all of which match experimental HEP needs. Physics justifications for the allocations have to be specific and NOT programmatic for the bulk of the available cycles on HPCs.
- HEP gets 10% of the total as a program. Program managers (or detail-is when Tom was there) decide which experiment gets what within HEP.
- Fermilab experiments used all of the resources they were allocated last year.
- This represented 15-20% of the need, depending on experiment (CMS, Nova, ...)
- As exascale machines come online HPC resource constraints may disappear for those that can utilize GPUs... no one in HEP can at the moment.

Information Flow from SCD Review Committees

Recommendations of Resource Scrutiny Group

- Improve the SCPMT template and provide report in advance
 - this is a natural consequence of moving to the CRSG model.
- Improve efficiency of managing resources allocated to the experiments
- Facilitate on-boarding of the experiments and reduce the long-term direct support.
- Revamp storage resources and usage for improved sustainability.
- Continue efforts to develop and implement common tools across frontiers
 - Rucio is our one success in 2019, our efforts are funded by CompHEP
- SCD should identify 5% of its budget that can be used for R&D activities toward future hardware/software advances.
 - Undoable in 2019 due to 3% budget cuts in operations. The drastic reduction in CompHEP R&D funding in Fy20 has hurt. CompHEP no longer supports HepCloud or Geant
 - SCD headcount reduced by ~10; resignations and retirements were not replaced

International Computing Advisory Committee



ICAC Report Highlights

- The ICAC commented on the progress SCD has made on the 14 recommendations from the Mar. meeting.
 - I won't go through all of them due to time constraints however I'll highlight the recommendation evaluations that the PAC should be most interested in.
 - I've posted the full ICAC report if you'd like to see the entire response to their spring recommendations.
1. Resource Scrutiny Group
 2. DUNE Funding
 3. Funding lines for other non-LHC experiments
 4. Federated Identities|
 5. HPC Strategy
 6. DUNE Computing Management
 7. DUNE Computing Model
 8. Fermilab as Host Lab
 9. Storage Strategy
 10. Reducing Duplication
 11. Software R&D Strategy
 12. Facility Resources
 13. Separation of Environments
 14. Student Programs

ICAC Report Highlights - 1 Resource Scrutiny Group

- “The committee was pleased to see a concrete plan for setting up a resource scrutiny group.”
- It was clarified and agreed that the purpose of this group is to:
 - Receive “Resource Request Documents” from experiments. The RRDs should state the experiments’ usage over the last year, state the forward capacity requirements for the next year in detail, and the next n-years as preliminary requests. Resource requests should be based upon a sound computing model which should be described succinctly, but in enough detail to allow the panel to constructively scrutinise the requests.
 - Scrutinise the requests to ensure the model is sound in terms of data access and replication policy, CPU campaigns, etc., and that the capacity provided is used appropriately.
- Recommend two focci
 - Scrutiny of DUNE separately (if it receives a separate funding line) internationally
 - Scrutiny and prioritization between the smaller experiment needs

ICAC Report Highlights - 2 HPC Strategy

- “The work in progress around the use of HPC resources appears to be appropriate.”
- They acknowledge that re-engineering HEP codes to use GPUs is the primary goal of the CCE.
- They say, “This is a topic to be followed at the next ICAC meeting.” as there is some justifiable skepticism about the nature of this proposal.

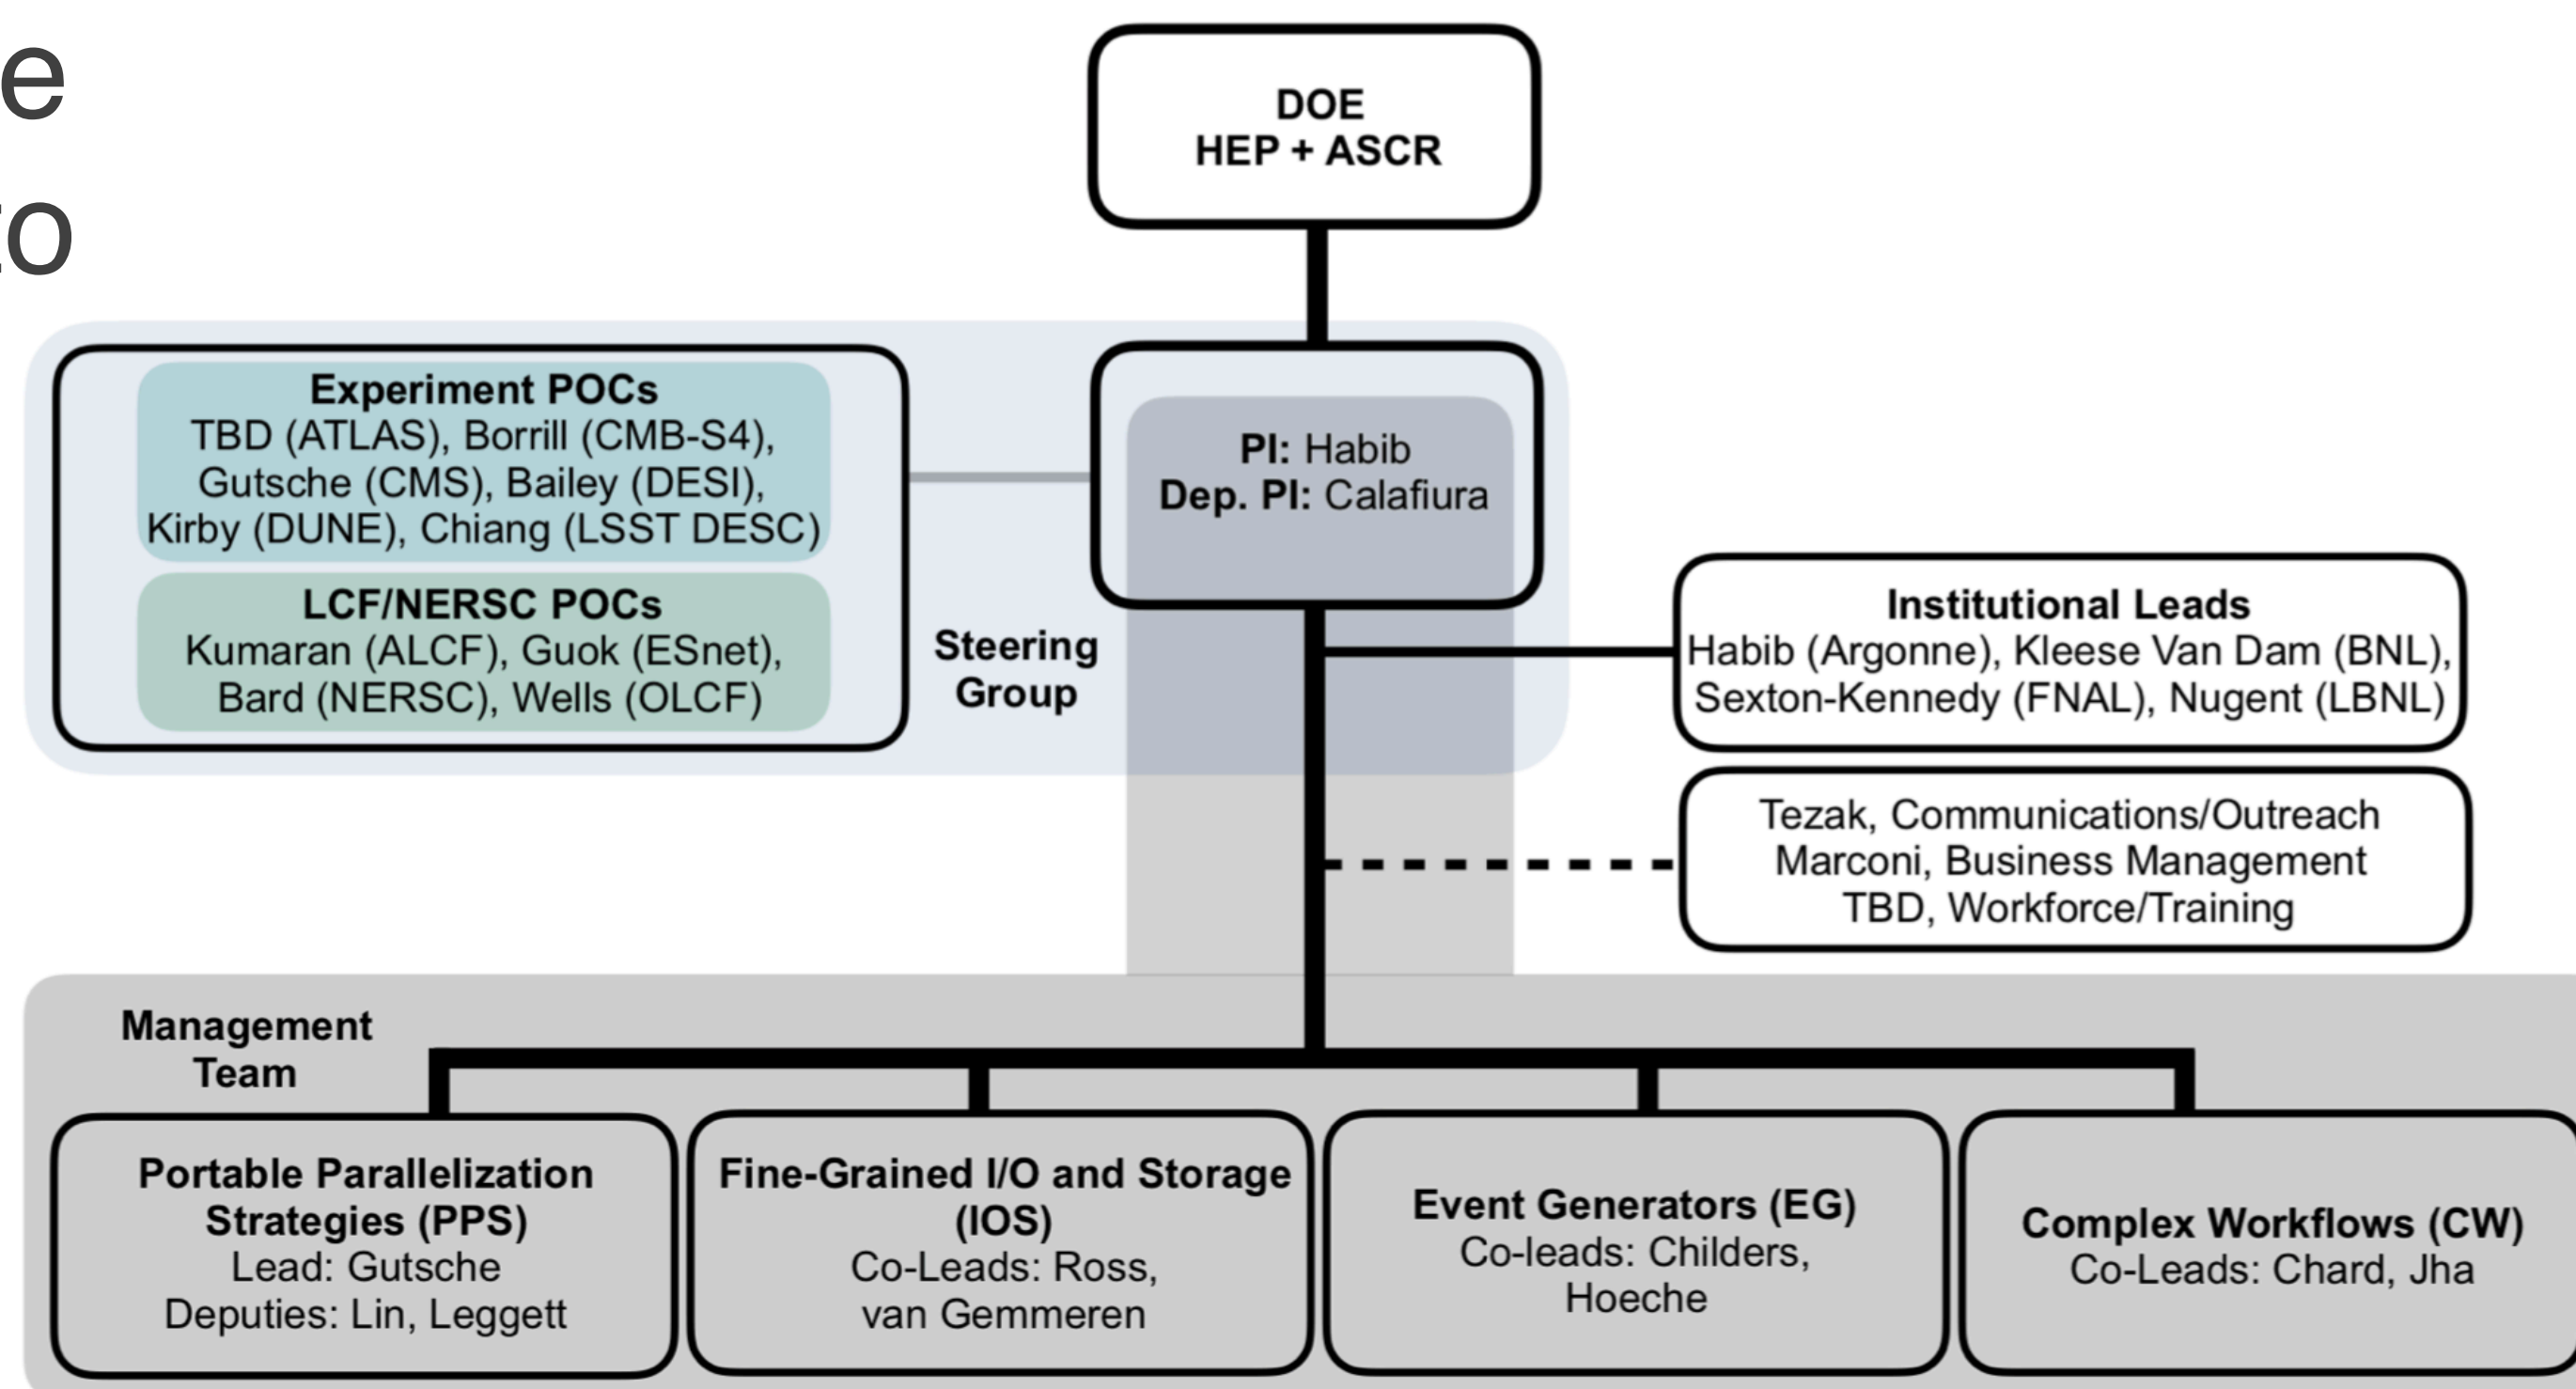


Figure 1. *Project organizational structure: The Management team is responsible for program execution, the Steering group provides guidance and feedback, and Institutional leads are responsible for technical and human resources at their institutions, as well as working together on cross-Institution interactions.*

ICAC Report Highlights - 3 DUNE Computing Model

- “We acknowledge the significant progress made by DUNE on its computing model definition.”
- Would like to see a more formal document at next meeting.
- Does not believe it is their job to review their CM but strongly advise that we put a team together to do this.

ICAC Report Highlights - 5 Storage Strategy

- “...actions taken so far, despite going in the right direction, tended to be opportunistic rather than driven by a long-term strategy ”
- “...concerns of aging hardware”
 - DOE is also concerned about this
 - In FY19 SCD spent 5.8M\$ on compute and storage hardware
 - DOE is right that this is ~10% of the facility and not adequate
- Driven by long term concerns about our tape facility we have been offered Ian’s help in initiating a collaboration with CERN around their new CERN Tape Archive.
- ICAC also encourages us to participate in WLCG Data Organization, Management and Access = DOMA working group.
 - We do this but already but could do more.

ICAC Report Highlights - 4 Software R&D Strategy

- Categorize R&D into Ops, sustaining capabilities (suscap) and long term R&D
- “Suscap R&D is a clear set of activities that are essential to keep operating as a facility, but there seems to be almost no funding for this, and it cannot easily be taken from the operations program without impact on ongoing operation.”
- “(suscap) is vital in order to keep Fermilab as a world-leading facility for the future.”
- “It is also necessary to identify how operation optimisation may allow to find these additional resources for suscap R&D.”

ICAC Report Highlights - 5 Facility Resources

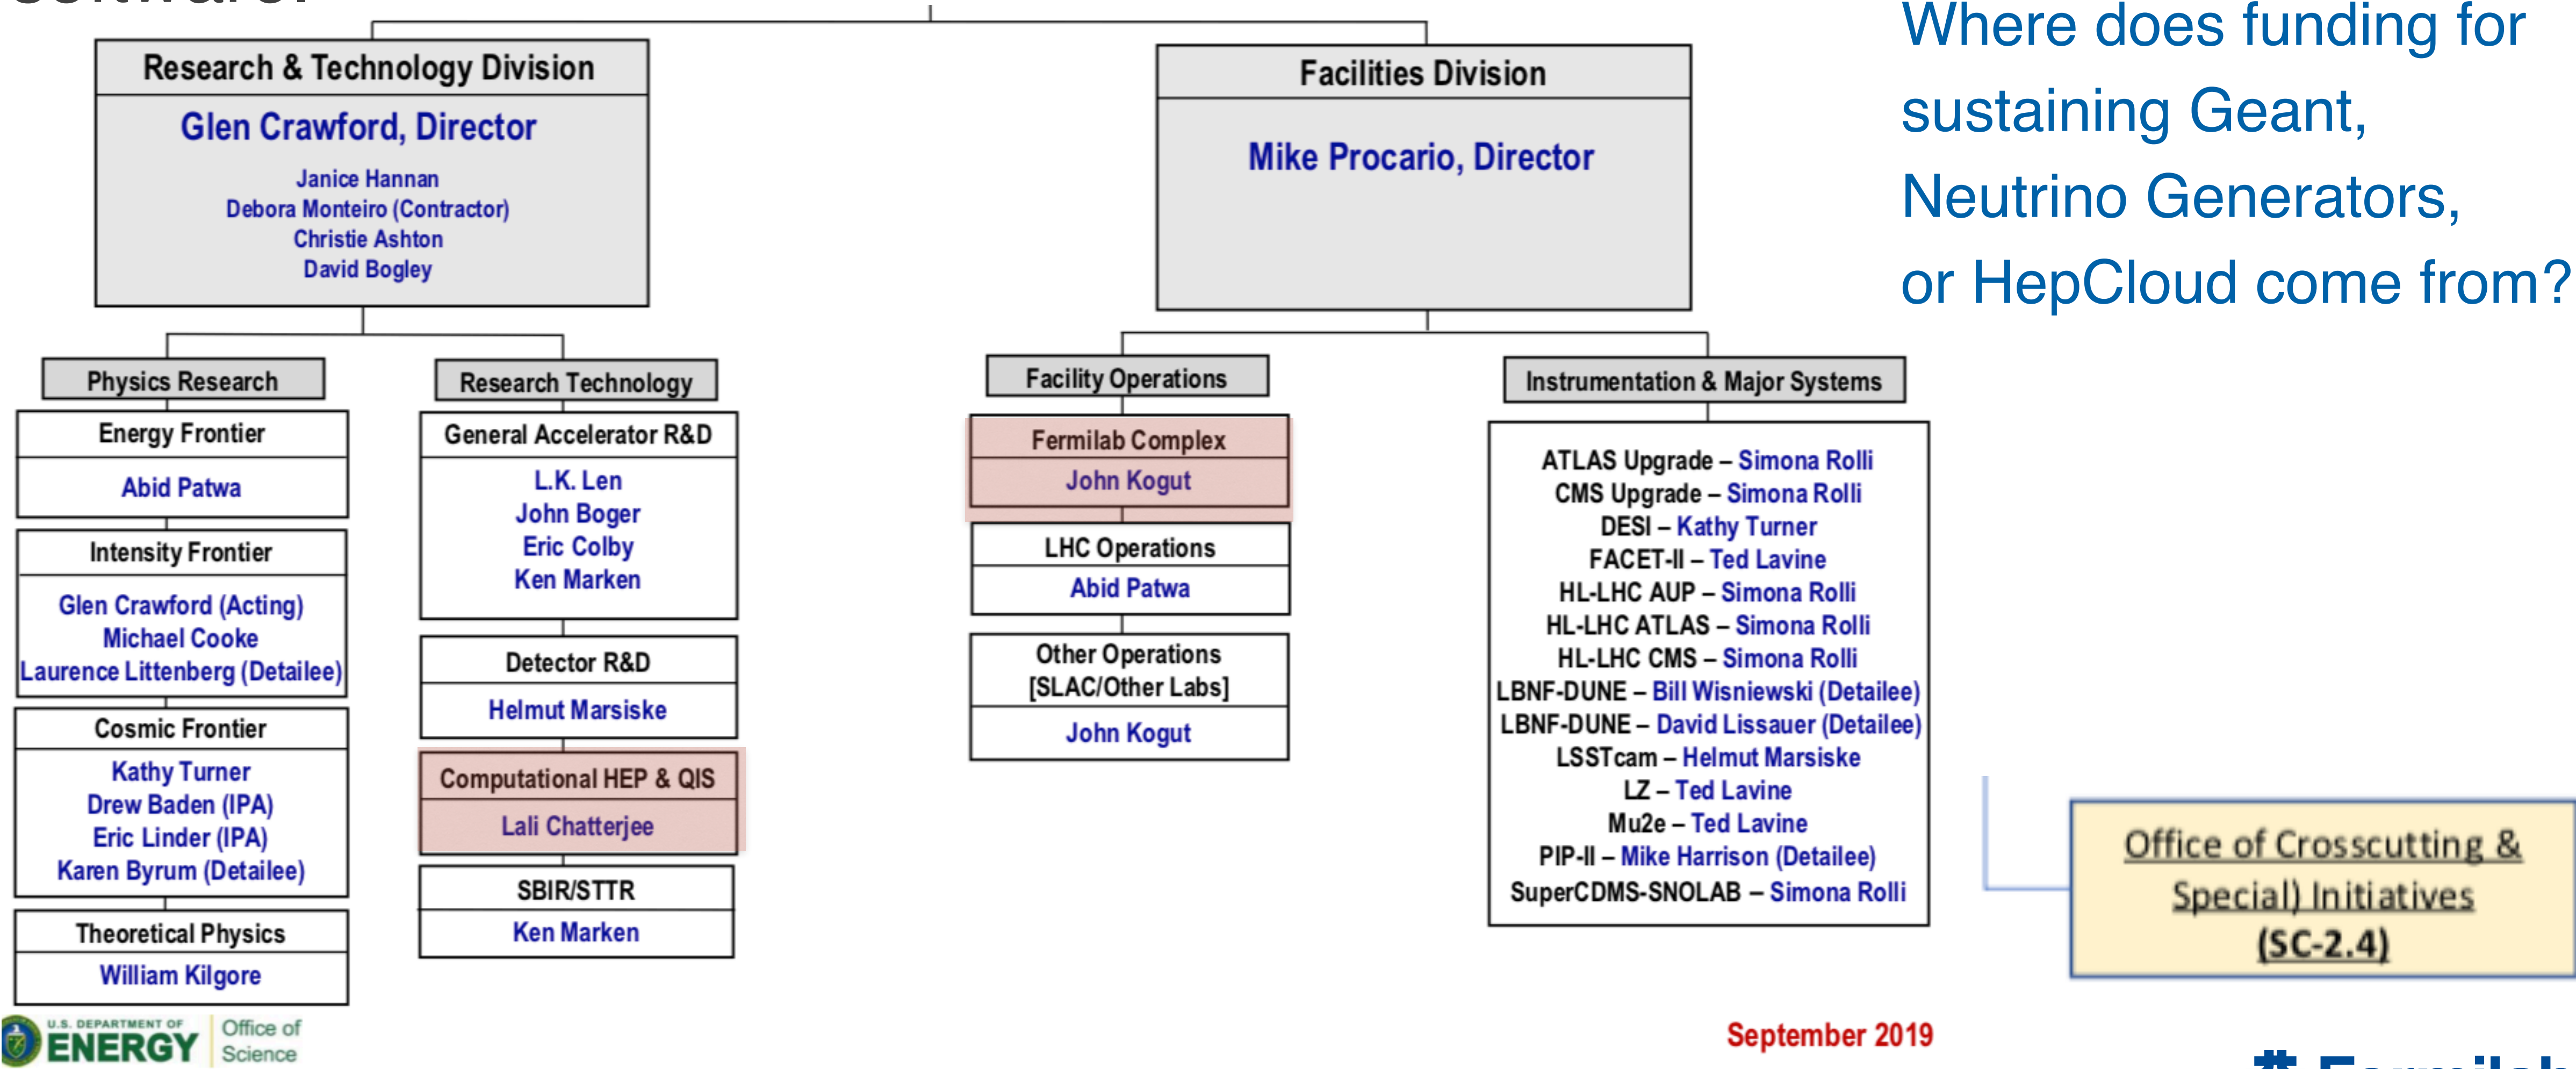
- “the ongoing reorganisation seems a good start that has much improved the potential for internal communication and integrated projects across the division”
- “The vision presented, is to develop the Institutional Cluster as a collection of resources – HTC, HPC, Storage, and networking, with users interacting via a scientific gateway (HEPCloud).”
- “In the experience of the committee it is essential to have a medium-term funding and resource planning outlook (5 years or so)”
- DOE is also worried about our facility funding...

CIO Advocacy

Sustaining community and facility software within DOE

- As discussed at our retreat with DOE there is no box sustaining field wide software.

Where does funding for sustaining Geant, Neutrino Generators, or HepCloud come from?



Open Data and Data Lifetime Cycle Management

- Fermilab computing would like to make it policy that all data stored must:
 - Be fully cataloged with appropriate meta-data to make it useful to someone other than it's creator.
 - Have a lifetime recorded with it even if that lifetime is infinite.
 - Not a problem for new experiments but we need your support with the old ones
- With the above making certain old datasets OPEN becomes a possibility
 - I view this as a health of the field issue. For instance microBoone data can be very interesting to those wanting to do LAr detector studies.
 - Funding agencies have advocated for clear data policies.
 - Discussion?

Summary

- SCD has made progress in many areas as confirmed by our review committees.
- SCD took a hard hit in funding between FY19 and FY20 causing some tasks to be very understaffed.
- We have moved staff to overhead funded tasks and operations in addition to not making new hires to replace retirements and resignations.
- Getting the planned funding for CCE will help with our goals and budget situation, but it is initially modest.

Back up

