Vision and Strategy for Computing at Fermilab

Elizabeth Sexton-Kennedy
Fermilab PAC
18 Jul 2019

Introduction

- What is the strategic direction and high level goals for Fermilab Computing
  - HPC migration strategy
  - Mid scale computing
  - Fermilab as a cross-cutting hub for data movement and storage
- Fermilab support for experiment operations
  - CMS
  - DUNE and the rest of the intensity frontier program
  - LQCD and other theory, Accelerator modeling, and Cosmic
- Fermilab scientific computing divisions's ambitions in R&D
- Advisory Committees and the flow of information
- How can the PAC support us

Fermilab Computing Vision

"Moore's Law" – the good old days

https://www.karlrupp.net/2018/02/42-years-of-microprocessor-trend-data/

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2017 by K. Rupp

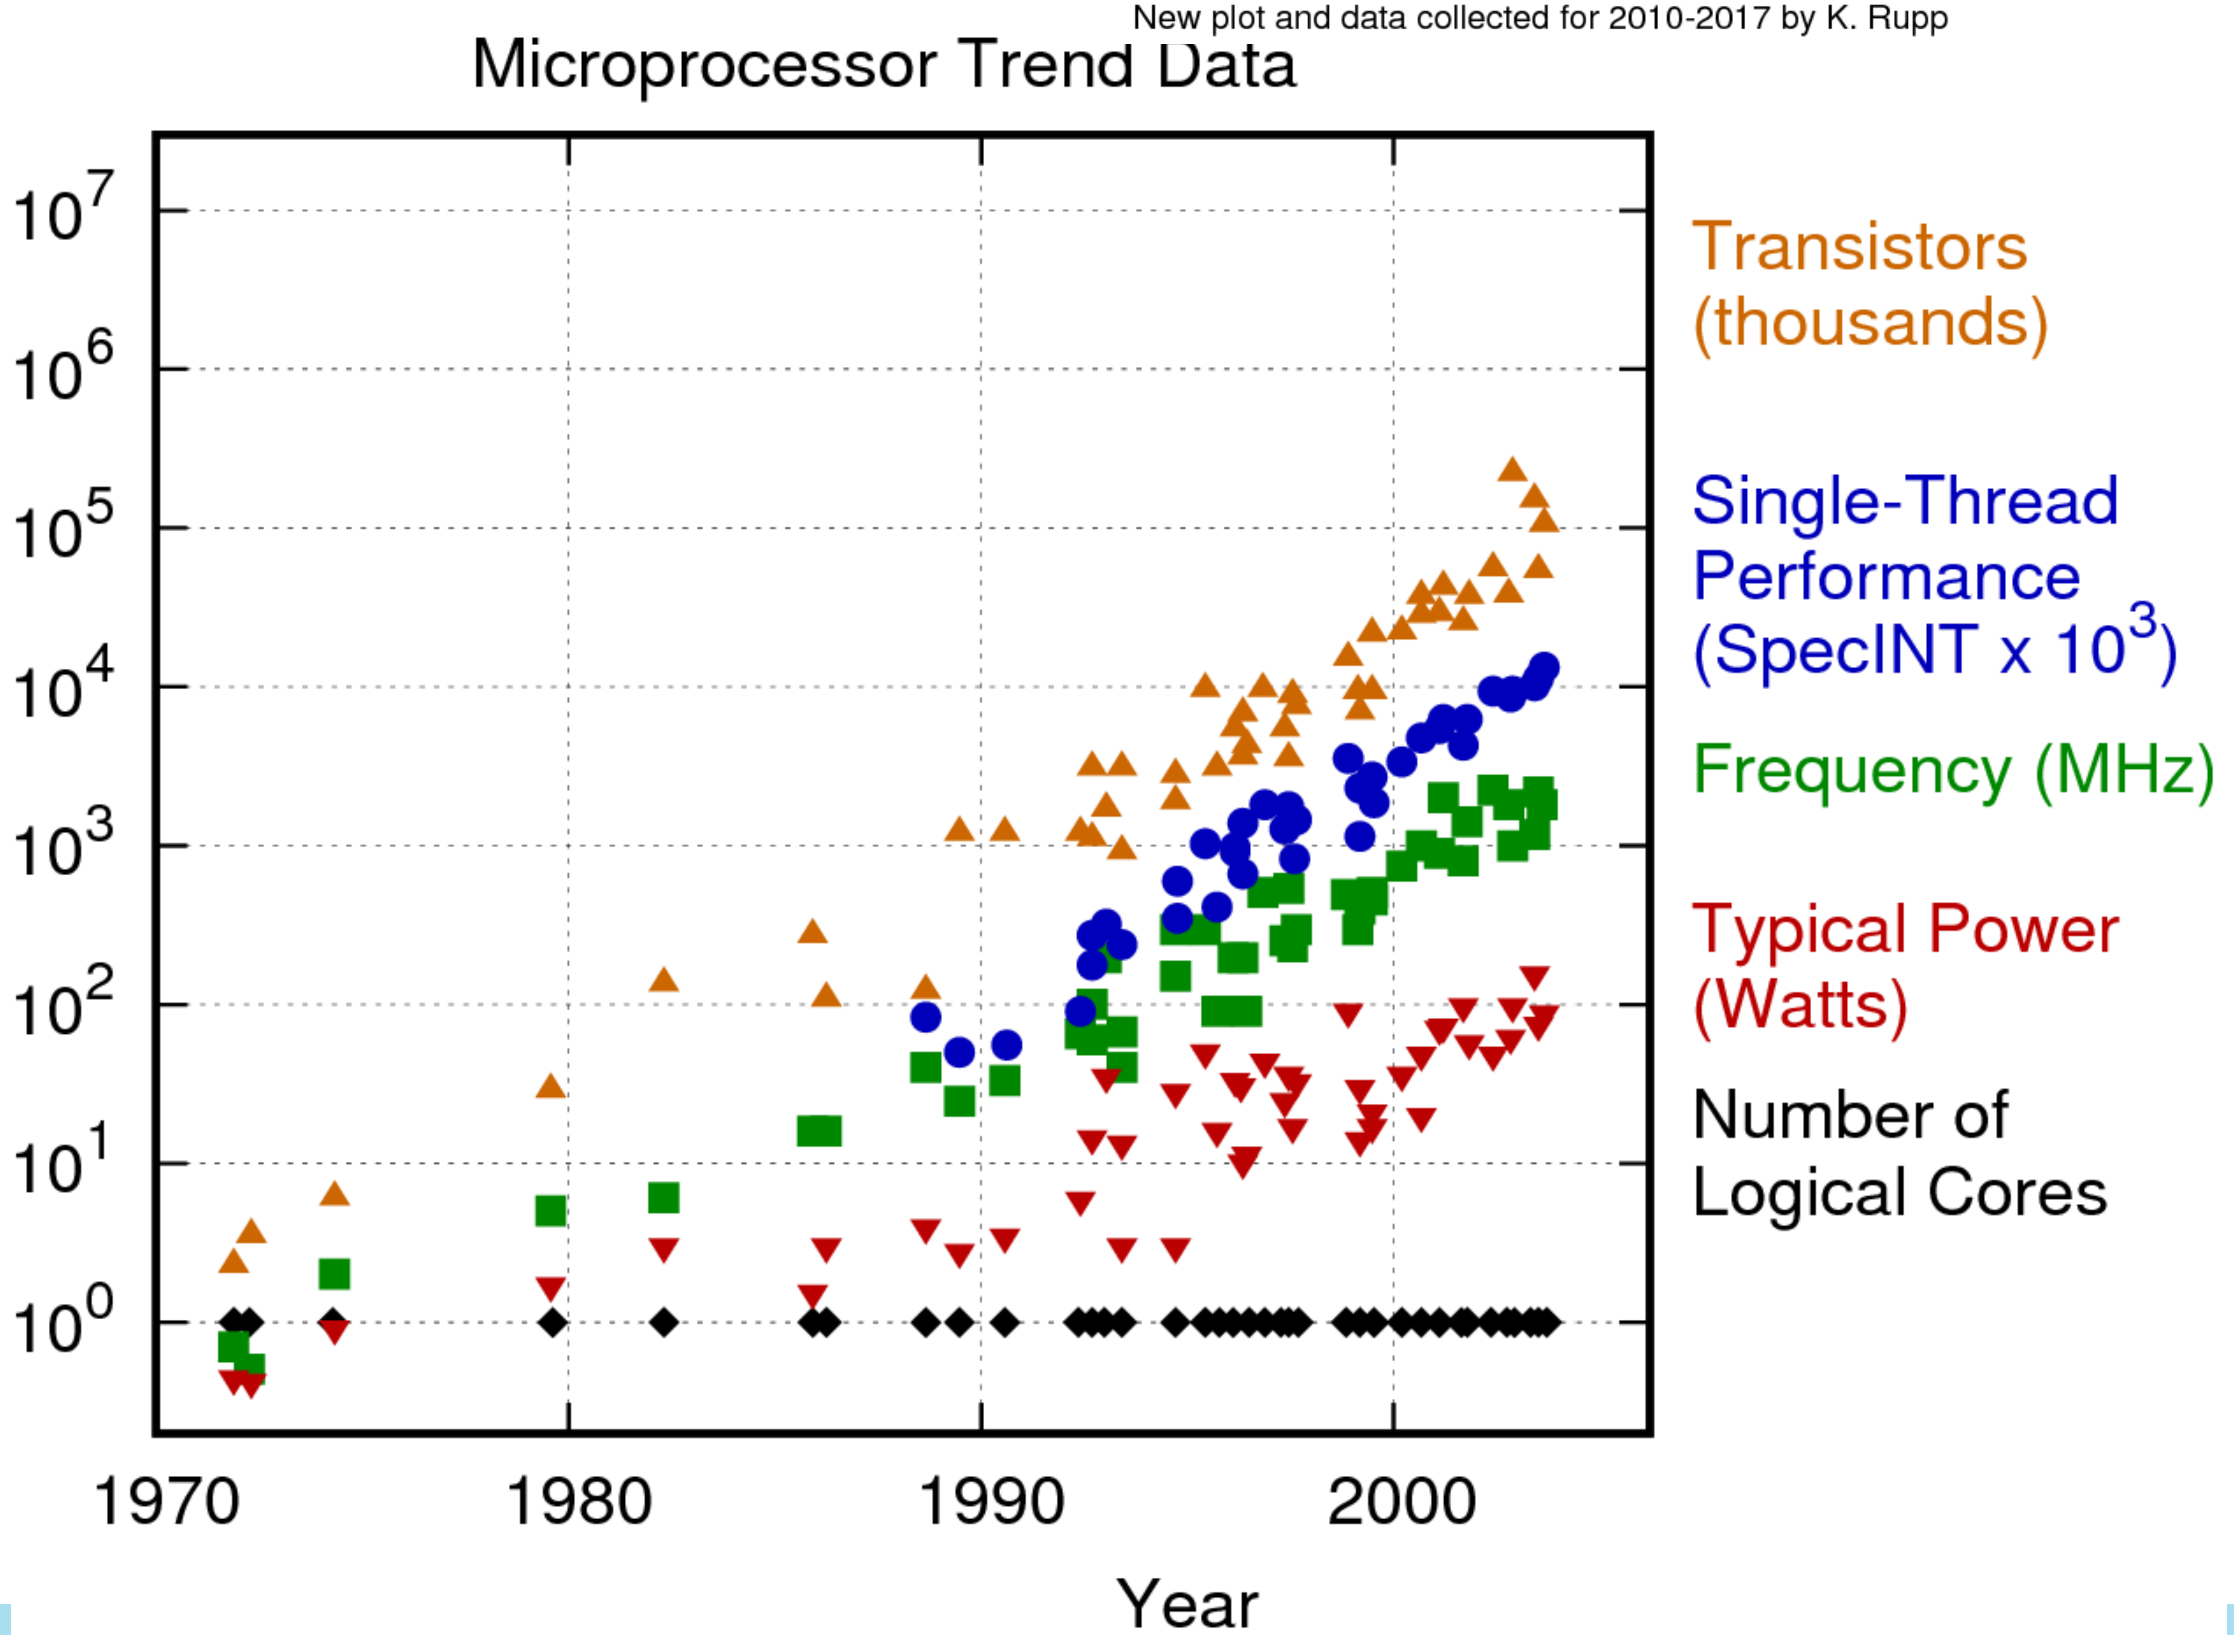

The world computing grid was built during these years and the policies still in place today where shaped by this reality.

The software work could be de-prioritized because applications improved by themselves.

“Moore’s Law” – recent times

Trends have changed

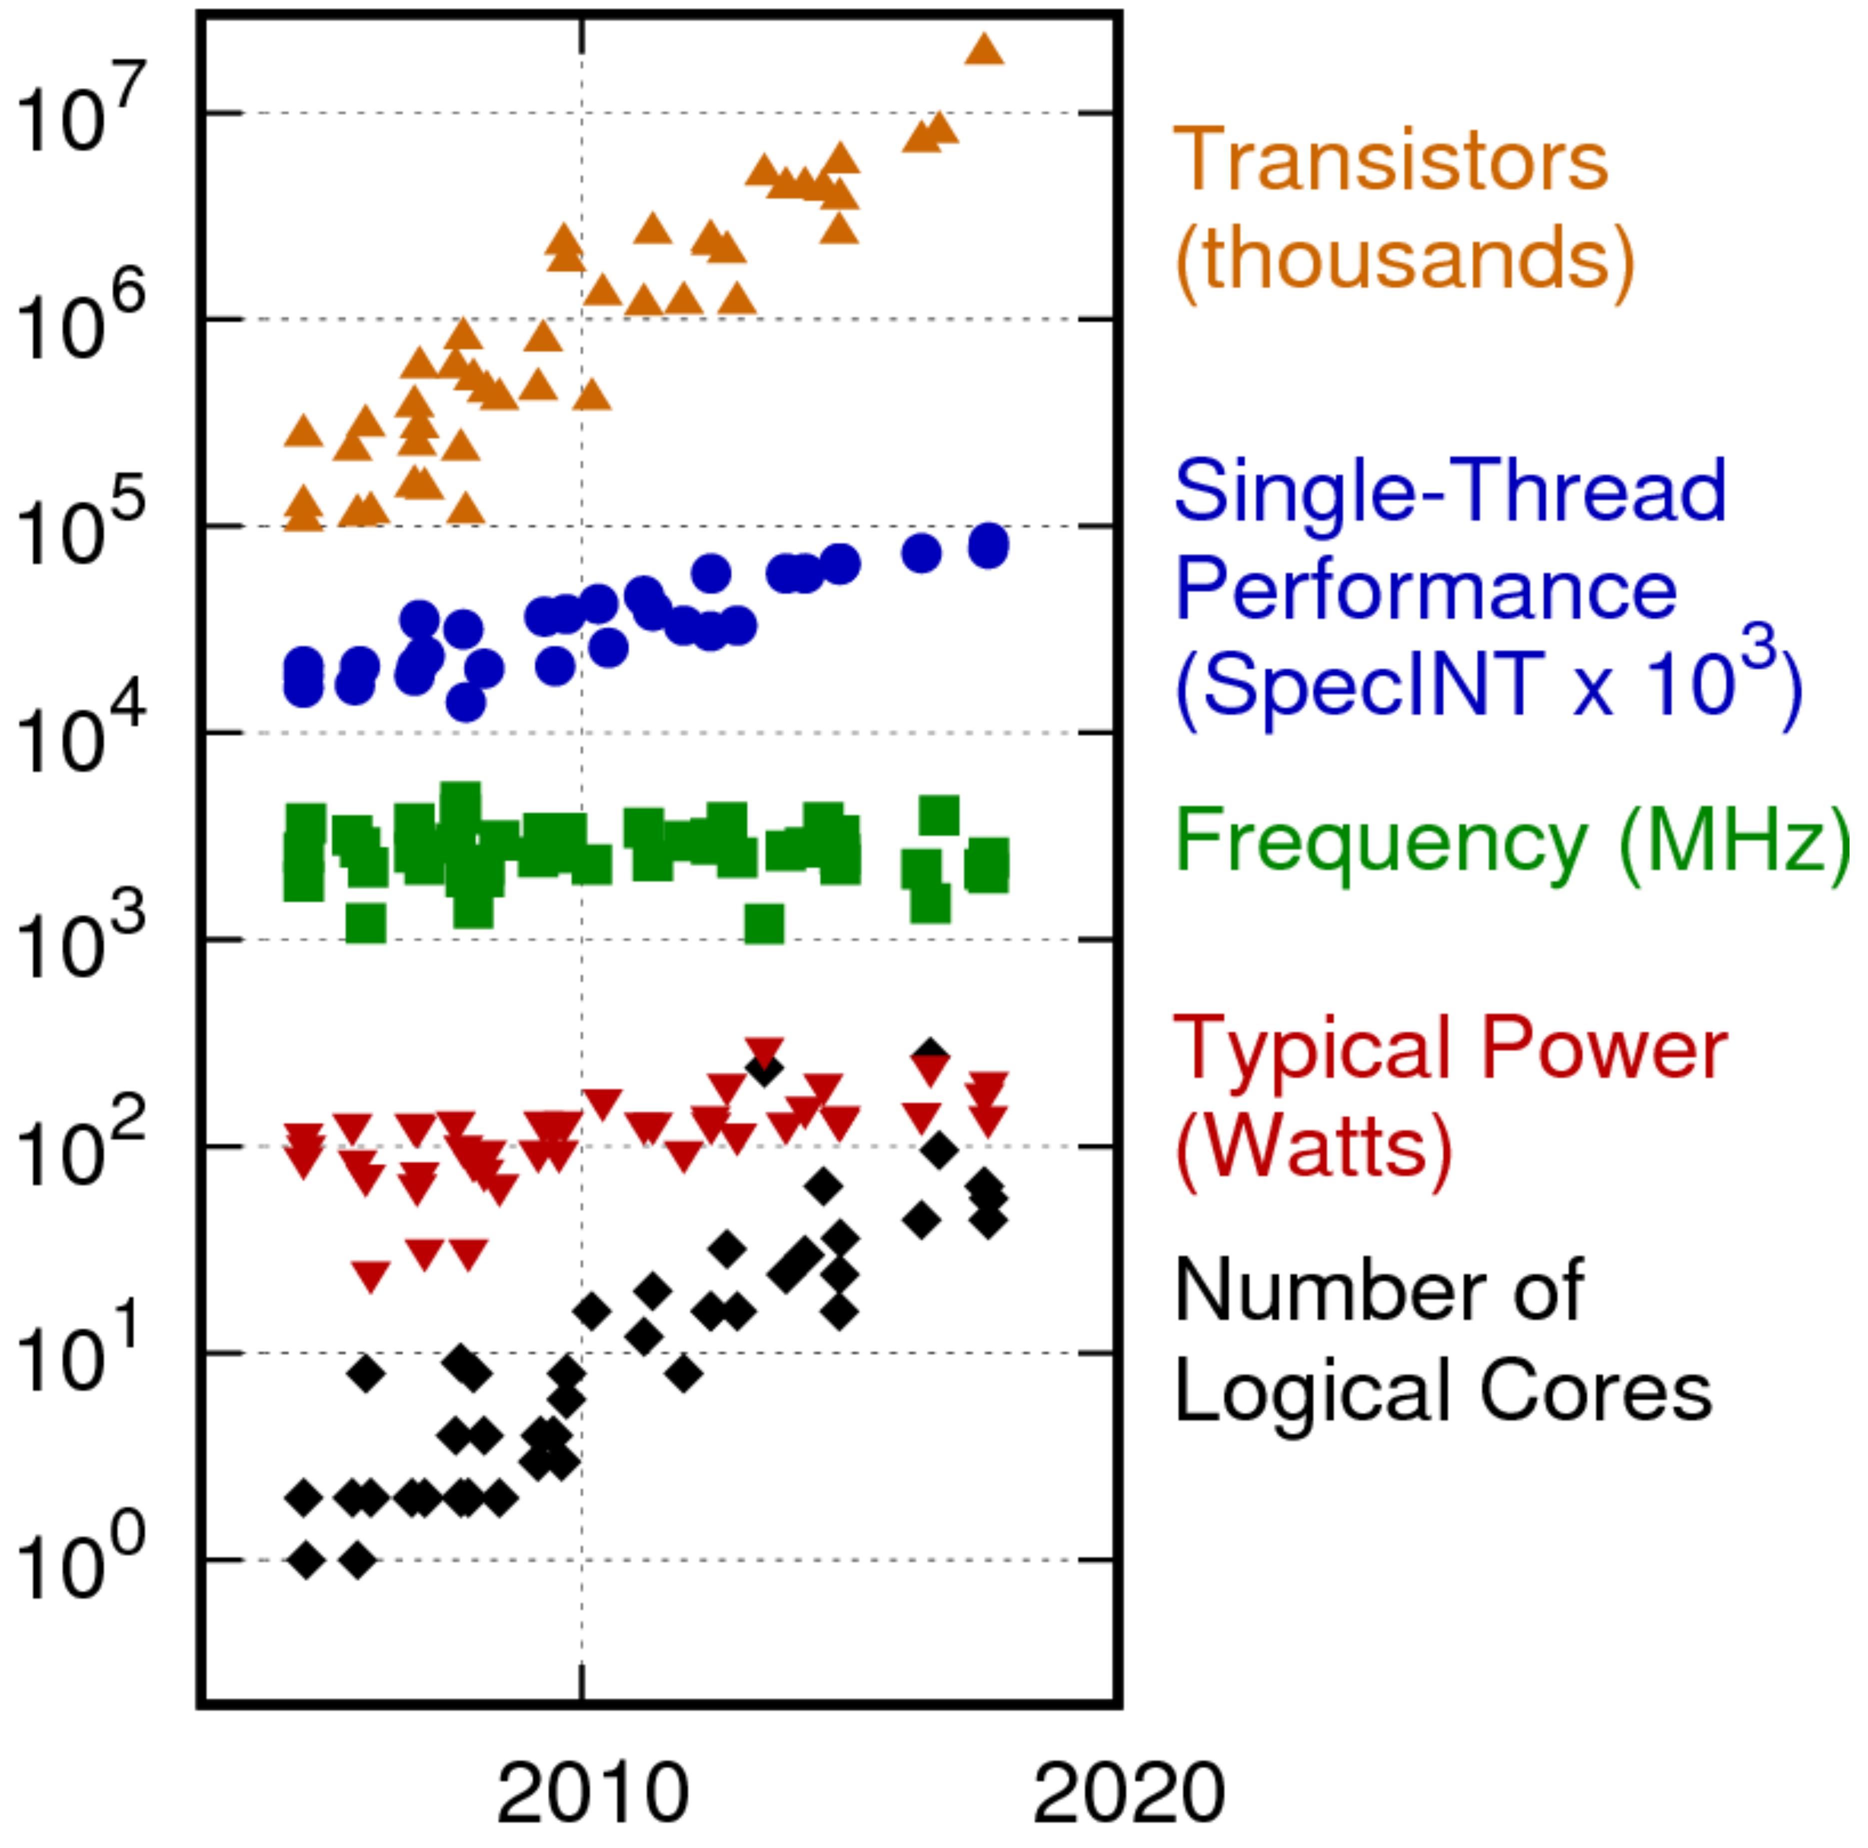

- Architectures are changing
  - Driven by solid state physics of CPUs
    - Multi-core
    - Limited power/core
    - Limited memory/core
    - *Memory bandwidth increasingly limiting*
- High Performance Computing (HPC, aka Supercomputers) are becoming increasingly important for HEP
  - 2000s: HPC meant Linux boxes + low-latency networking
    - No advantage for experimental HEP
  - Now: HPC means power efficiency
    - Rapidly becoming important for HEP, everyone else
- New technologies will change our workflows even on traditional resources

BIG DATA AND EXTREME-SCALE COMPUTING:

- Fermilab should be a major player in reconciling the split between traditional HPC and HTC ecosystems, discussed by an international gr

"Combining HPC and HTC applications and methods in large-scale workflows that orchestrate simulations or incorporate them into the stages of large-scale analysis pipelines for data generated by simulations, experiments, or observations"

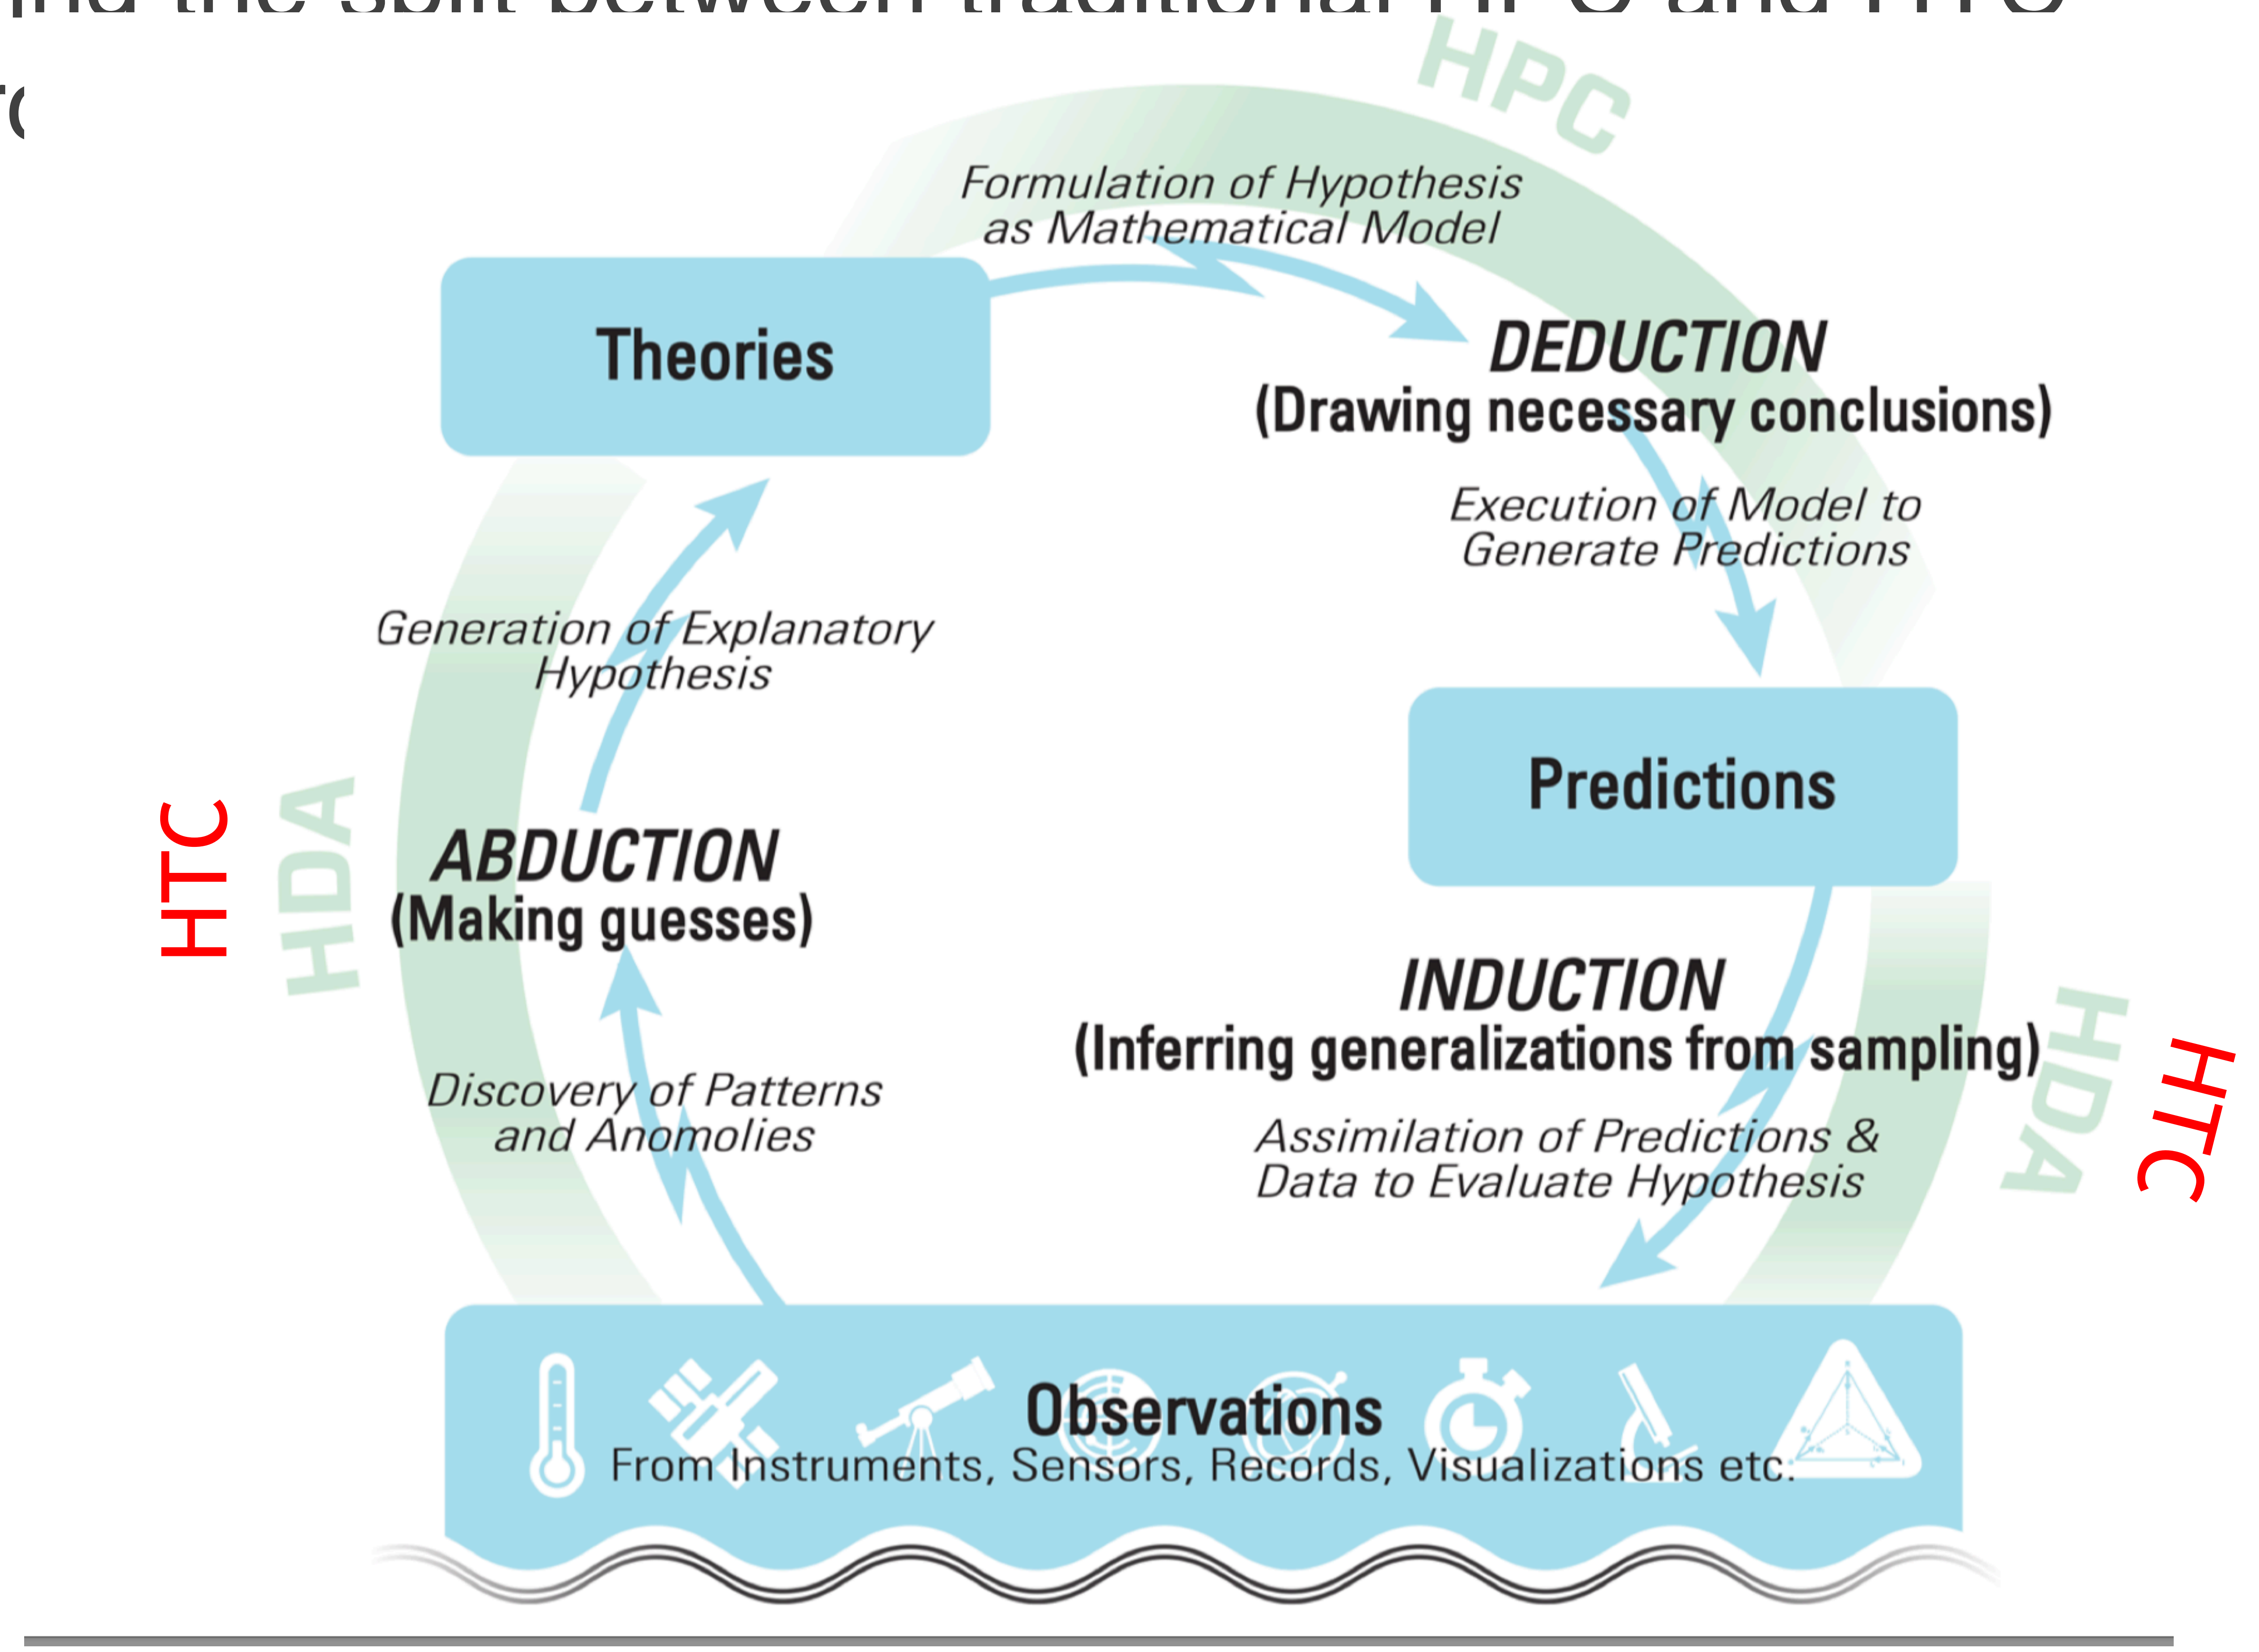

[1] http://www.exascale.org/bdec/sites/www.exascale.org.bdec/files/whitepapers/bdec2017pathways.pdf

Next Generation HPC

- Architectures for Exascale machines have been announced
  - x86_64 + GPUs (!)
    - Not NVIDIA GPUs like Summit

  Need a portable programming model
- ALCF (Argonne)
  - Aurora
    - 2021
    - > 1 Exaflop
    - Intel
- OLCF (Oak Ridge)
  - Frontier
    - 2021
    - > 1.5 Exaflop
    - AMD

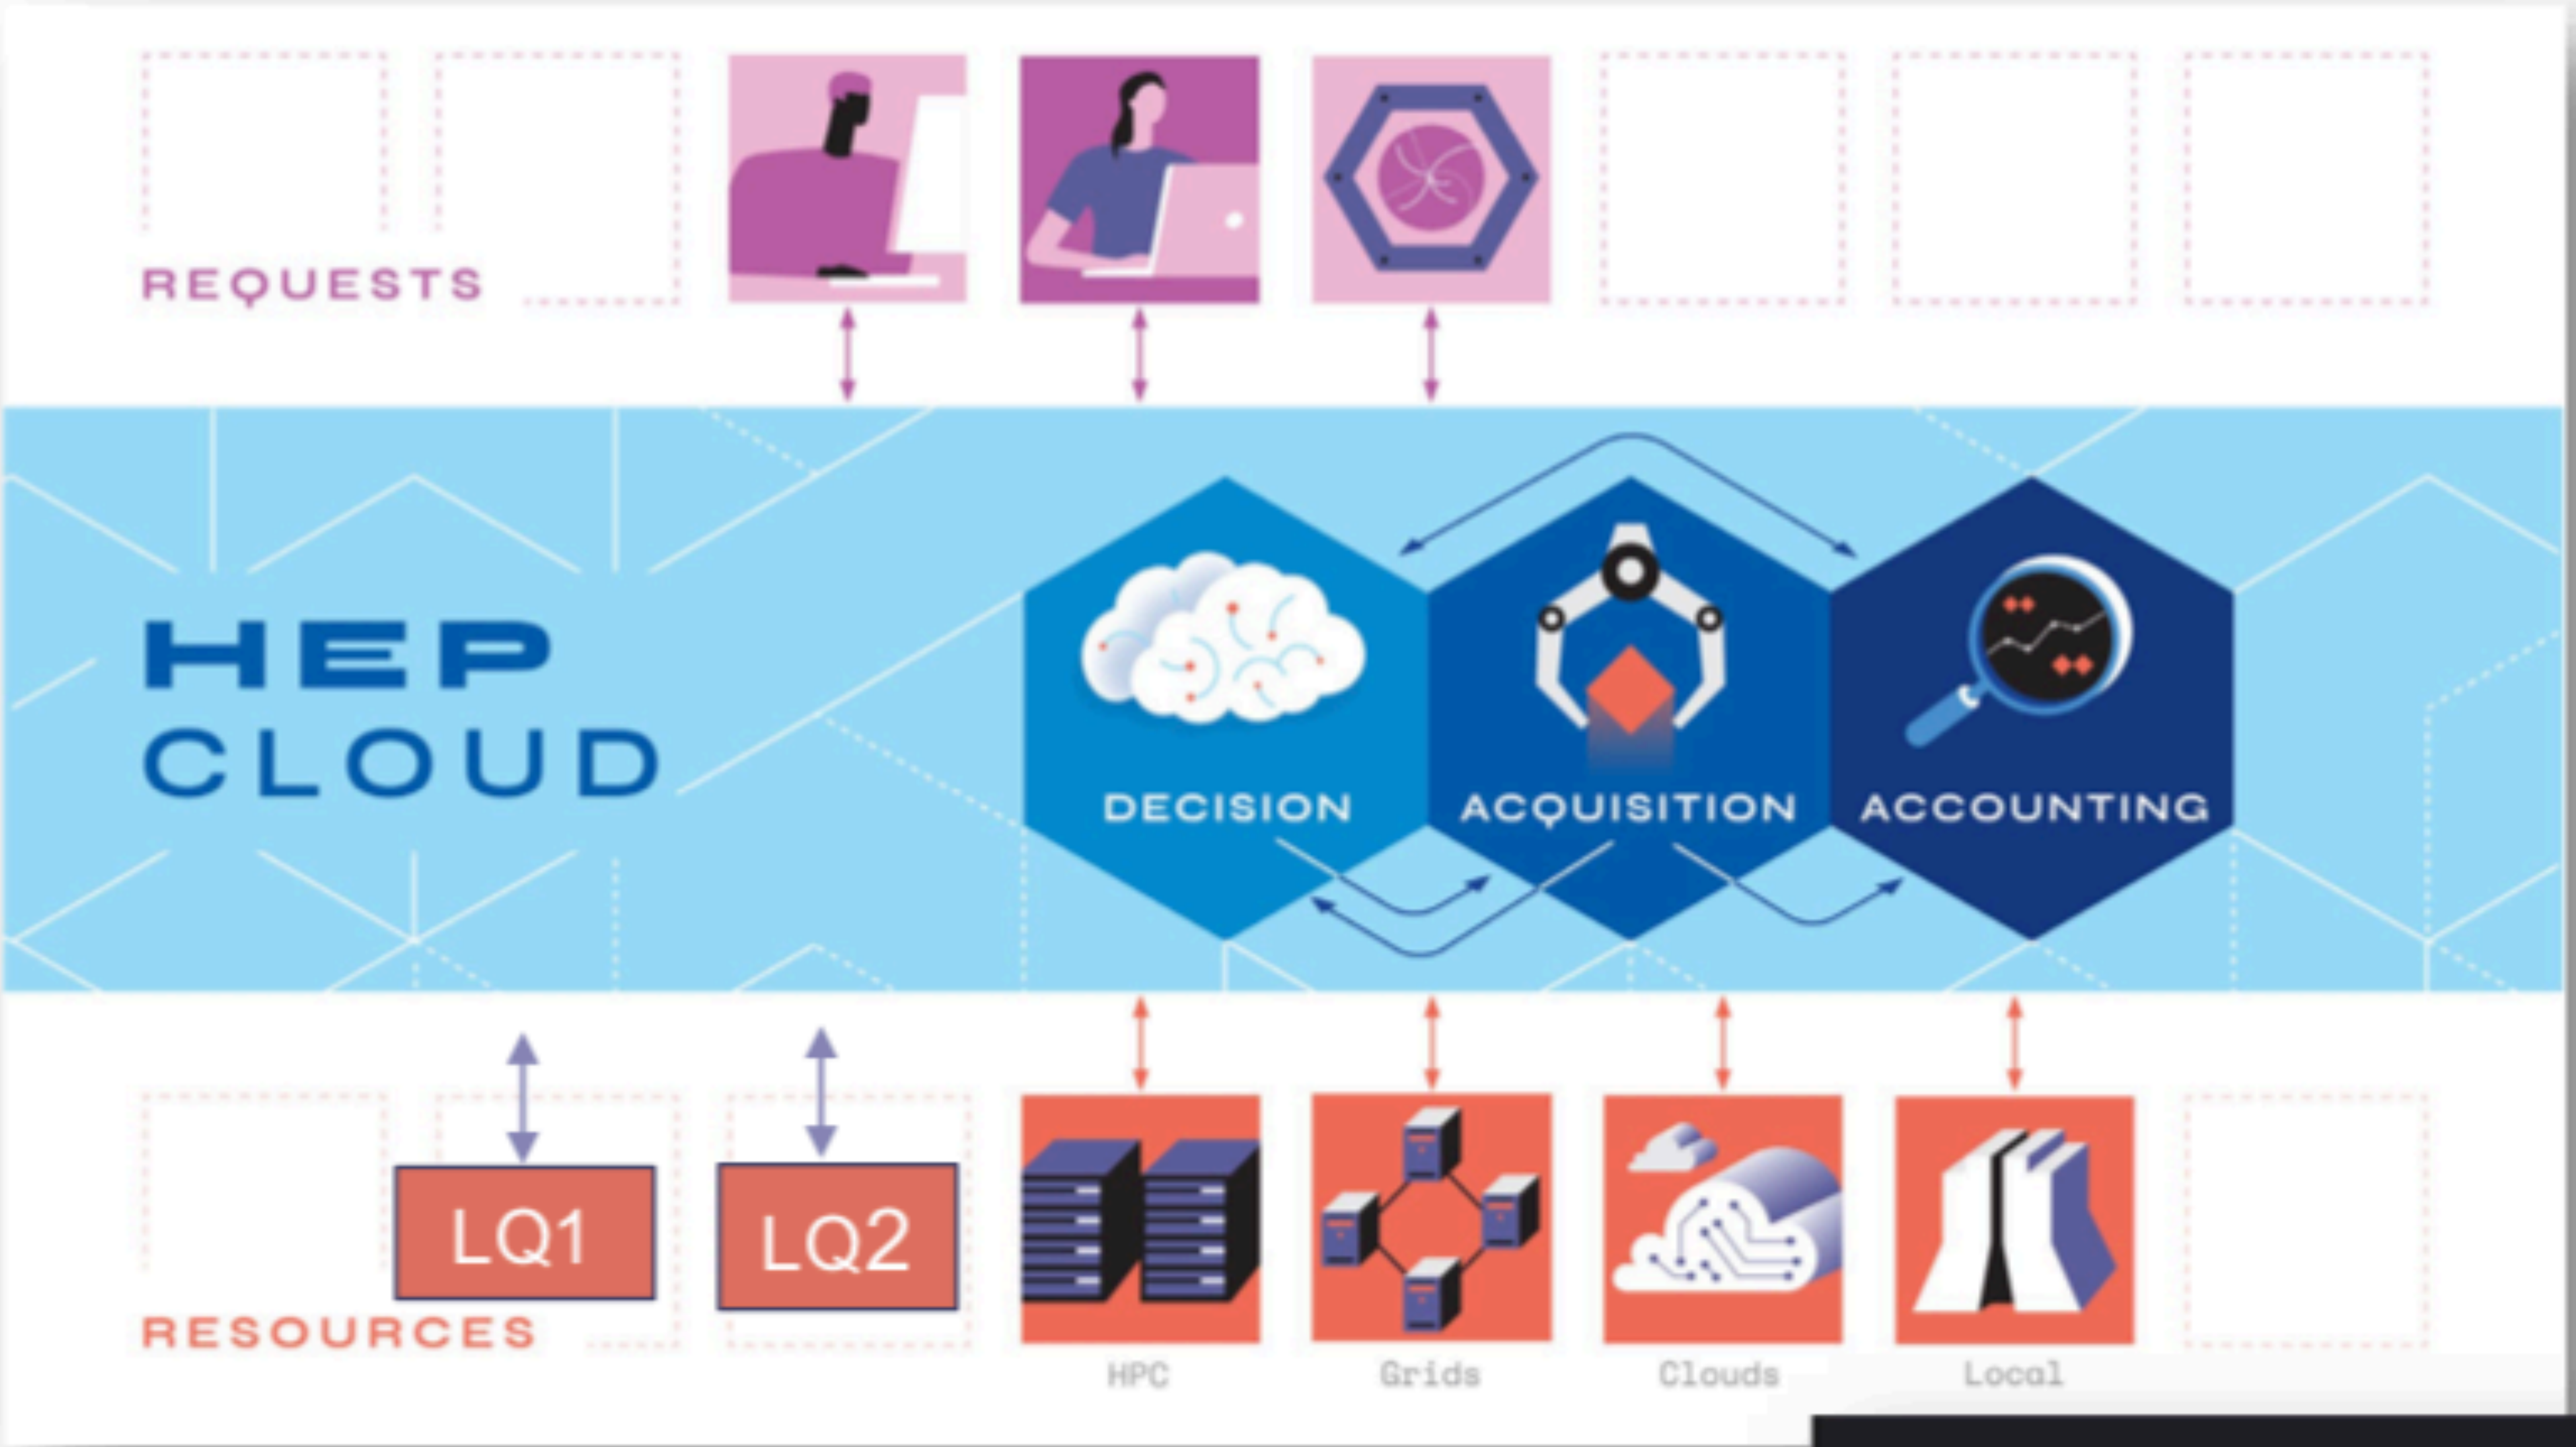

NOvA using HEPCloud to claim ~1M cores at NERSC to perform a large-scale analysis over a short timeframe.

Resource demand spikes that cannot be met by on-premise resources can be satisfied by routing jobs to commercial clouds, OSG resources or LCF's.

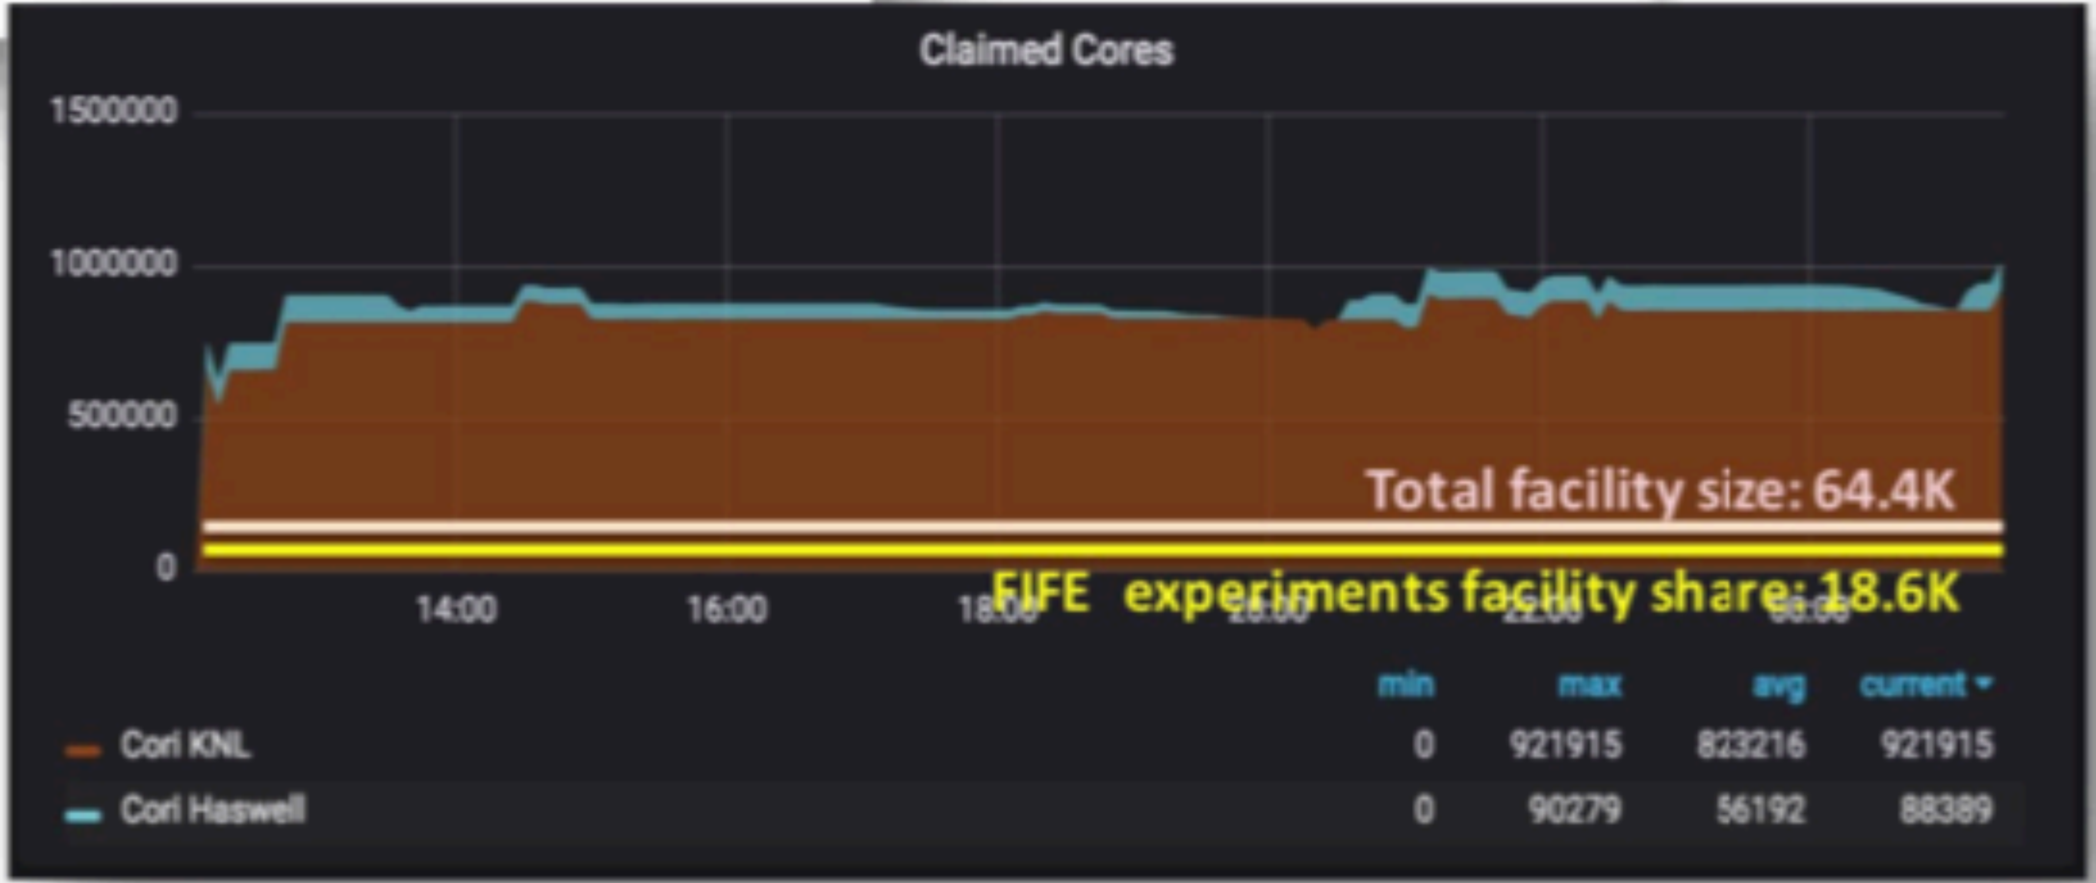

Laboratory Complex Program for Computing

- The computing challenges of the next decade are large. We need a new era of laboratory complex cooperation to create the data facilities so necessary for scientific insights we aim for.
- HPCs at 3 of the labs, data facilities at 2 (FNAL,BNL).
- We need to develop a national cyber-infrastructure to serve the needs of the scientific community, and have dynamic sharing of our resources.
  - Provide a smooth onramp to exascale computing
  - Provide mid-scale facilities that can be used to test work-flows and codes
  - Provide custodial storage for our experimental and theory communities
  - Provide networking and the expertise to run them in a cyber safe way
    - Fermilab networking engineers worked with ESNET to put a proposal for the far site

Fermilab Computing support for experiment operations

Are the Experiments Ready for Exascale DOE Facilities?

- No, CMS may be ahead of DUNE but it also needs them sooner (2021)

- Strategy:

1. Bring DUNE to the level of CMS - Establish host lab responsibilities
2. Help them both with doing projects to move into the Exascale era

- Provide support and manpower to put together funding proposals to engage ASCR:

1. Have already succeeded with SiDAC
2. Putting together a CCE proposal together with other labs
3. Cooperating with IRIS-HEP

Software & Computing Research and Development

Why - causes:

A. Requirements from experiments based on upcoming needs
B. Forward thinking to keep up with evolving computing landscape
C. Useful technologies that scientists adopt and needs support
D. Fruitful collaborations

What - drivers:

A. CMS in the HL-LHC era and DUNE
B. New computing architectures/accelerators and the Exascale High Performance Computing Era
C. Machine Intelligence's impact on HEP reconstruction and analysis
D. Specific funding calls (e.g. SciDAC from DOE-ASCR)

These guide the HOW:

- Software and Computing requirements from CMS and DUNE
- Community White Papers (HEP Software Foundation and IRIS-HEP)
- Goals of SciDAC and ECP
- Strive for common tools where possible and common principles for moving forward
- Domain and computer scientists working in cooperation

High Priority Technologies (Unordered)

Strategy

- Be the leader in data management and storage
- Be the leader in access to heterogeneous computing
- Be the center of core software development
- Be the center of scientific software R&D
- Be the leader in HEP AI/ML R&D
- Be the leader in DAQ integration
- Provide the home for physics analysis

Community data management system (Rucio)

R&D into storage technologies *such as*

Wide area network storage (data lakes)

Object stores

Root i/o & serialization

Monitoring technologies

High Priority Technologies (Unordered)

Strategy

- Be the leader in data management and storage
- Be the leader in access to heterogeneous computing
- Be the center of core software development
- Be the center of scientific software R&D
- Be the leader in HEP AI/ML R&D
- Be the leader in DAQ integration
- Provide the home for physics analysis

HEPCloud (HEP portal to computing resources)

R&D in efficient use of accelerators (GPUs, TPUs, FPGAs, QPUs)

Institutional Cluster (local access to heterogeneous computing technologies to aid scaling up to HPC)

R&D in Containerization

Monitoring technologies

High Priority Technologies (Unordered)

Strategy

- Be the leader in data management and storage
- Be the leader in access to heterogeneous computing
- Be the center of core software development
- Be the center of scientific software R&D
- Be the leader in HEP AI/ML R&D
- Be the leader in DAQ integration
- Provide the home for physics analysis

Further a common scientific data processing framework

R&D in containerization for deployment

Leadership in community efforts for software development (software management [e.g. Github], build [e.g. spack/spackdev] & CI systems)

High Priority Technologies (Unordered)

Strategy

- Be the leader in data management and storage
- Be the leader in access to heterogeneous computing
- Be the center of core software development
- Be the center of scientific software R&D
- Be the leader in HEP AI/ML R&D
- Be the leader in DAQ integration
- Provide the home for physics analysis

Further a common scientific data processing framework

Scientific Toolkit Development (e.g. LArSoft)

Modernization for new computing architectures (e.g. in simulation [Geant] & reconstruction)

Exploit open source software (e.g. concurrency libraries, Machine learning libraries)

Root development for future

High Priority Technologies (Unordered)

<u>*Strategy*</u>

- Be the leader in data management and storage
- Be the leader in access to heterogeneous computing
- Be the center of core software development
- Be the center of scientific software R&D
- Be the leader in HEP AI/ML R&D
- Be the leader in DAQ integration
- Provide the home for physics analysis

Exploit open source Machine learning software - provide expertise in turning your challenge into a ML application

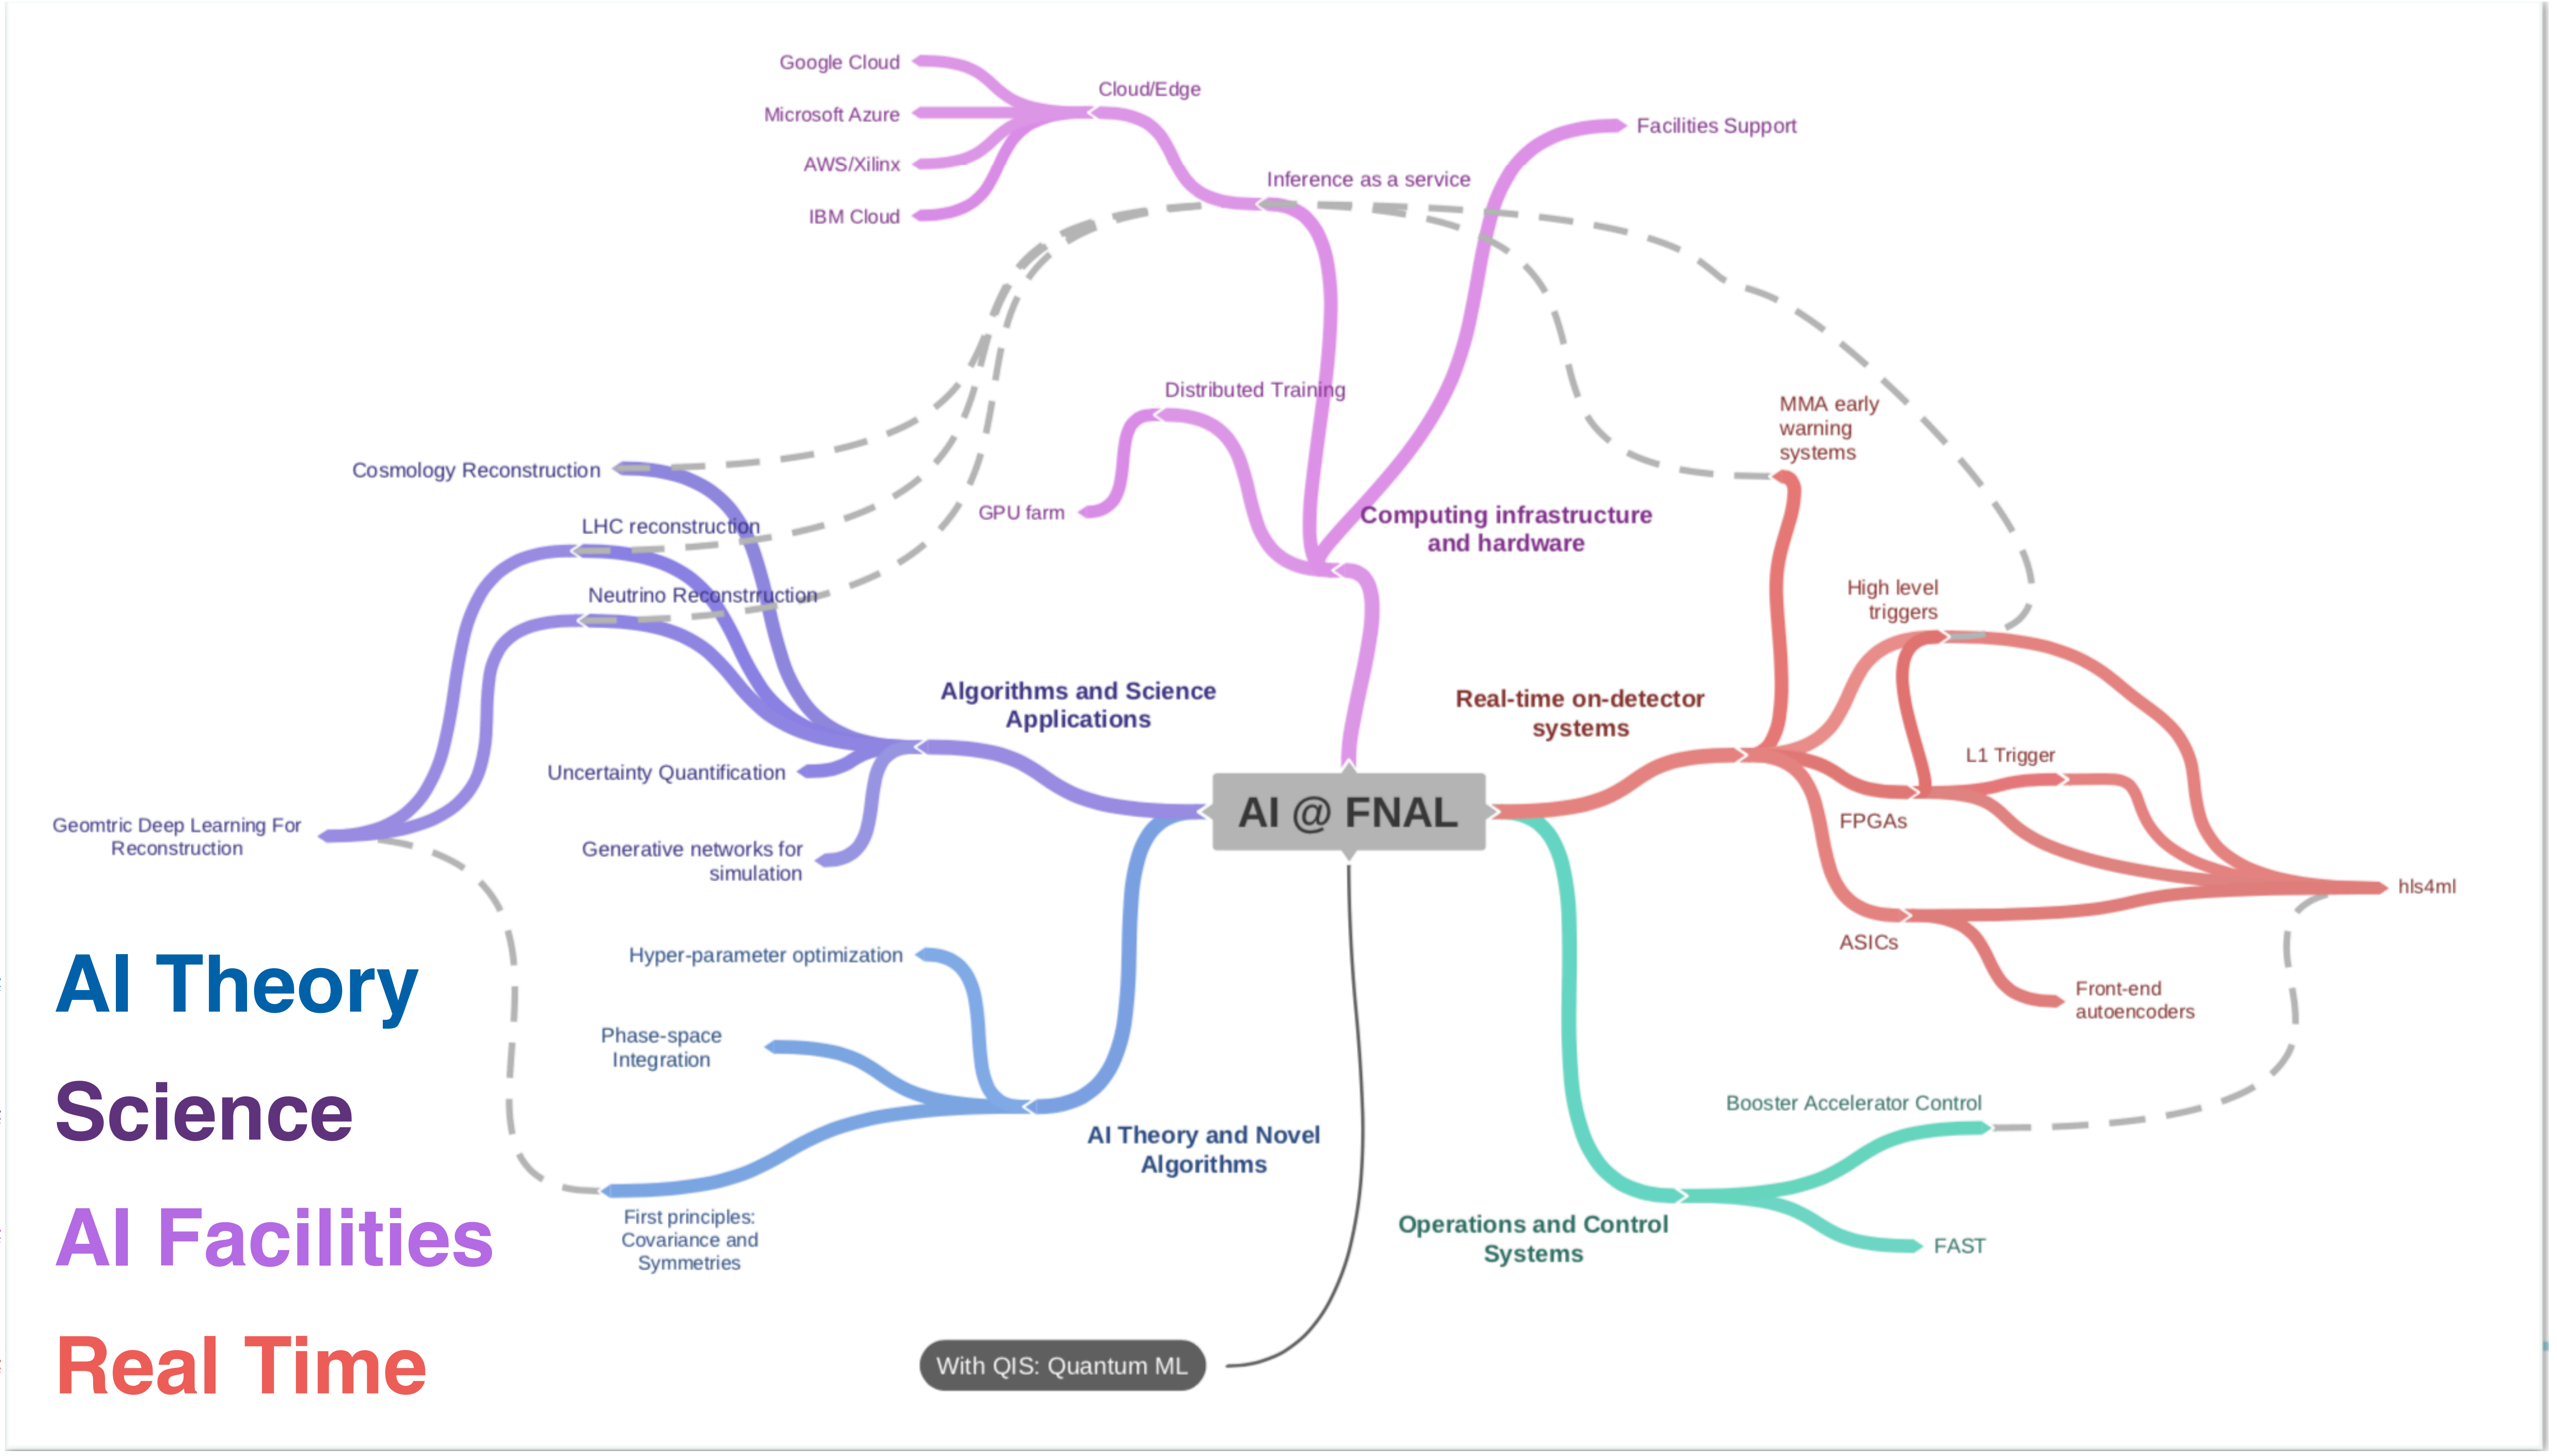

AI Theory

Science

AI Facilities

Real Time

Control & Ops

Fermilab

High Priority Technologies (Unordered)

Strategy

- Be the leader in data management and storage
- Be the leader in access to heterogeneous computing
- Be the center of core software development
- Be the center of scientific software R&D
- Be the leader in HEP AI/ML R&D
- Be the leader in DAQ integration
- Provide the home for physics analysis

Continued R&D in DAQ toolkits and off-the-shelf systems

R&D in efficient use of accelerators (GPUs, TPUs, FPGAs, QPUs)

High Priority Technologies (Unordered)

Strategy

- Be the leader in data management and storage
- Be the leader in access to heterogeneous computing
- Be the center of core software development
- Be the center of scientific software R&D
- Be the leader in HEP AI/ML R&D
- Be the leader in DAQ integration
- Provide the home for physics analysis

R&D on exploiting big data toolkits for analysis

R&D on object stores

Root development for future

Fermilab S&C R&D Afternoon

Thursday, June 13, 2019	
12:00 - 12:10	Introduction, Goals for the Meeting *10'* Speakers: Prof. Maria Spiropulu (Caltech), Dr. Lothar Bauerdick (Fermilab)
12:10 - 12:35	Innovation and Engineering for HL-LHC Computing Infrastructure *25'* *Computing HL-LHC challenges and the envisioned computation ecosystem to address them* Speakers: Dr. Oliver Gutsche (Fermi National Accelerator Laboratory), Dr. Lothar Bauerdick (Fermilab)
12:35 - 13:45	Innovation & Engineering (I) *These talks will convey the strategy of novel computation methods applicability towards solutions for HL-LHC.* 12:35 **"on-chip" computation (FPGA-based framework)** *20'* Speakers: Michail Bachtis, Javier Duarte (Fermilab) 12:55 **Advanced Methods for Data Processing and Reconstruction** *30'* Speakers: Lindsey Gray (Fermilab), Allison Reinsvold Hall, Nhan Tran (FNAL) 13:25 **Advanced Data Analysis and Simulation Methods** *20'* Speakers: Nicholas Smith (Fermilab), Joosep Pata (Caltech), Kevin Pedro (Fermilab)
13:45 - 14:00	Break, Discussions
14:00 - 15:00	Innovation & Engineering (II) *These talks will convey the strategy of novel computation frameworks and systems towards solutions for HL-LHC* 14:00 **Data Organization, Management, and Access** *15'* Speaker: Brian Bockelman 14:15 **Software Framework and HPC** *15'* Speaker: Matti Kortelainen 14:30 **Machine Learning Ecosystem** *15'* Speaker: Dr. Jean-Roch Vlimant (California Institute of Technology) 14:45 **Mapping Computing Innovation onto HL-LHC Needs** *15'* Speaker: Markus Klute (MIT)
15:00 - 15:20	Strategy towards HL-LHC Computing, Next Steps *20'* Speakers: Dr. Lothar Bauerdick (Fermilab), Dr. Oliver Gutsche (Fermi National Accelerator Laboratory), Prof. Maria Spiropulu (Caltech)
15:20 - 15:50	Executive Session/Adjourn *30'*

- Successful showing in DC
- Jim Siegrist interested, spent the whole afternoon with us
- He also spent the afternoon discussing our CCE proposals
- The SciDAC team is reviewing well this week - gives ASCR confidence that we can deliver

ICAC

- Our new **I**nternational **C**omputing **A**dvisory **C**ommittee
- Ian Bird (CERN, chair), Peter Clarke (University of Edinburgh), Suchandra Dutta (Saha Institute of Nuclear Physics), Peter Elmer (Princeton), Eric Lancon (Brookhaven National Laboratory), Michel Jouvin (LAL, Universite Paris-Sud and CNRS/IN2P3), Margaret Votava (FNAL, secretary)

Fermilab

ICAC Review March 15-16

- Inaugural Meeting of the International Computing Advisory Committee
  - https://indico.fnal.gov/event/20100

- Charter: The ICAC
  - Reviews and Advises the laboratory on:
    - computing operations,
    - cyber security,
    - upgrade plans, and
    - software and computing R&D aimed towards
      - the development and exploitation of future facilities
      - as well as advancing scientific tools and methods in general

  - Monitors progress with respect to the established laboratory objectives, currently encompassing:
    - Software and Computing for the Intensity Frontier Experiments;
    - Fermilab’s involvement in the HL-LHC Software and Computing Upgrades;
    - Progress toward common solutions for the above domains;
    - National and International cooperation and collaboration with partner institutions;
  - The ICAC is expected to address high-level strategic, programmatic, and planning issues, rather than specific implementation details.

ICAC Presentations

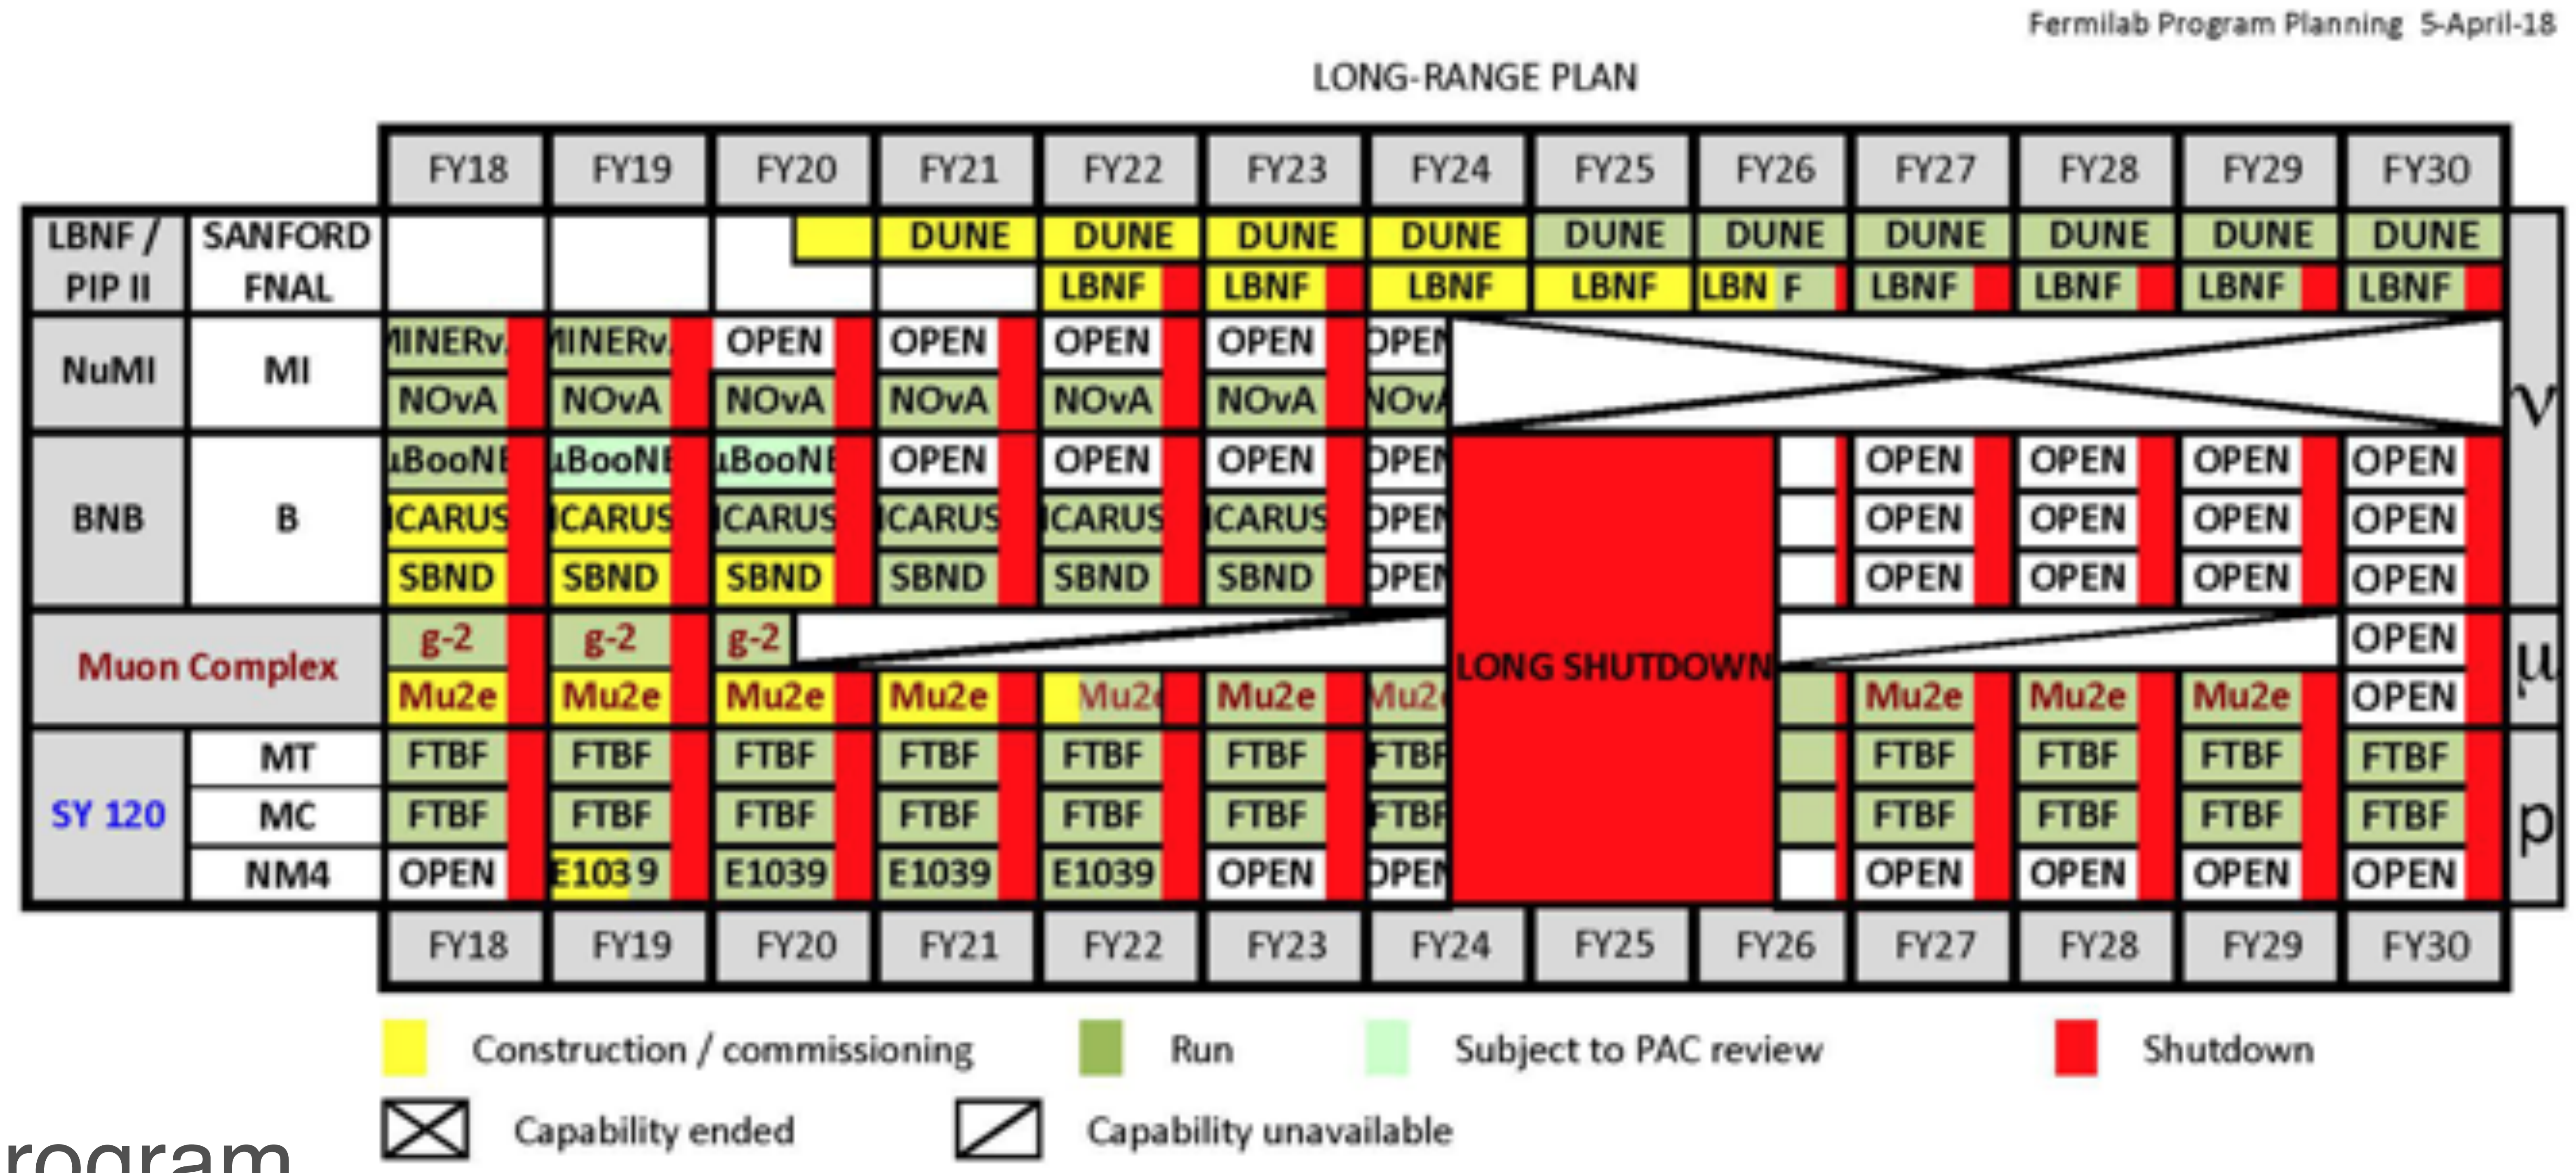

- Introductions
  - Computing Sector
  - Scientific Computing
- Strategy
  - HPC Strategy & US Exascale Program
  - International Cooperation Strategy
  - Cyber Security and Other DOE Mandates
  - Future Facility Plans
  - Local Operations Review (SCPMT) <——Their Recommendations are input to ICAC
  - Software and Computing R&D

SCPMT Charge

- Scope: Computing and Detector Operations funded activities
  - not Cosmic, not CMS, not SciDAC, etc.
- Priorities
  - We ask the committee for comments on priorities of support
    1. Are the lab / P5 priorities satisfied?
    2. Are the needs of the major experiments met?
    3. Are there low priority efforts that should be discontinued?
    4. We have expressed the effects of our plan in terms of risks; are the risk mitigations appropriate?
  - In an era where funding is diminishing at the same time needs are growing, we need to have a clear set of priorities
    - We ask for the committee's guidance on computing support for
      - The current experimental program
      - The future experimental program
      - … and the balance between the two

2019 SCPMT Review February 25-26

- Fermilab Computing Resource Scrutiny Group
- Committee: Lothar Bauerdick, Pushpa Bhat (Chair), Brian Bockelman, Taylor Childers, Ian Fisk, Kate Scholberg
- https://indico.fnal.gov/event/18685/
- Division Presentations:
  - Conventional Resources and Requests
  - HPC Resources
  - Service Requests
- Liquid Argon Experiments: DUNE, MicroBooNE, ICARUS, SBND
- Other Neutrino and Muon: NOvA, Muon g-2, mu2e, “everybody else”
- Externally funded experiments: CMS, DES, LSST

History and Future of Computing Resources

Data collected from the experiments for the portfolio review

This demonstrates the needs are increasing

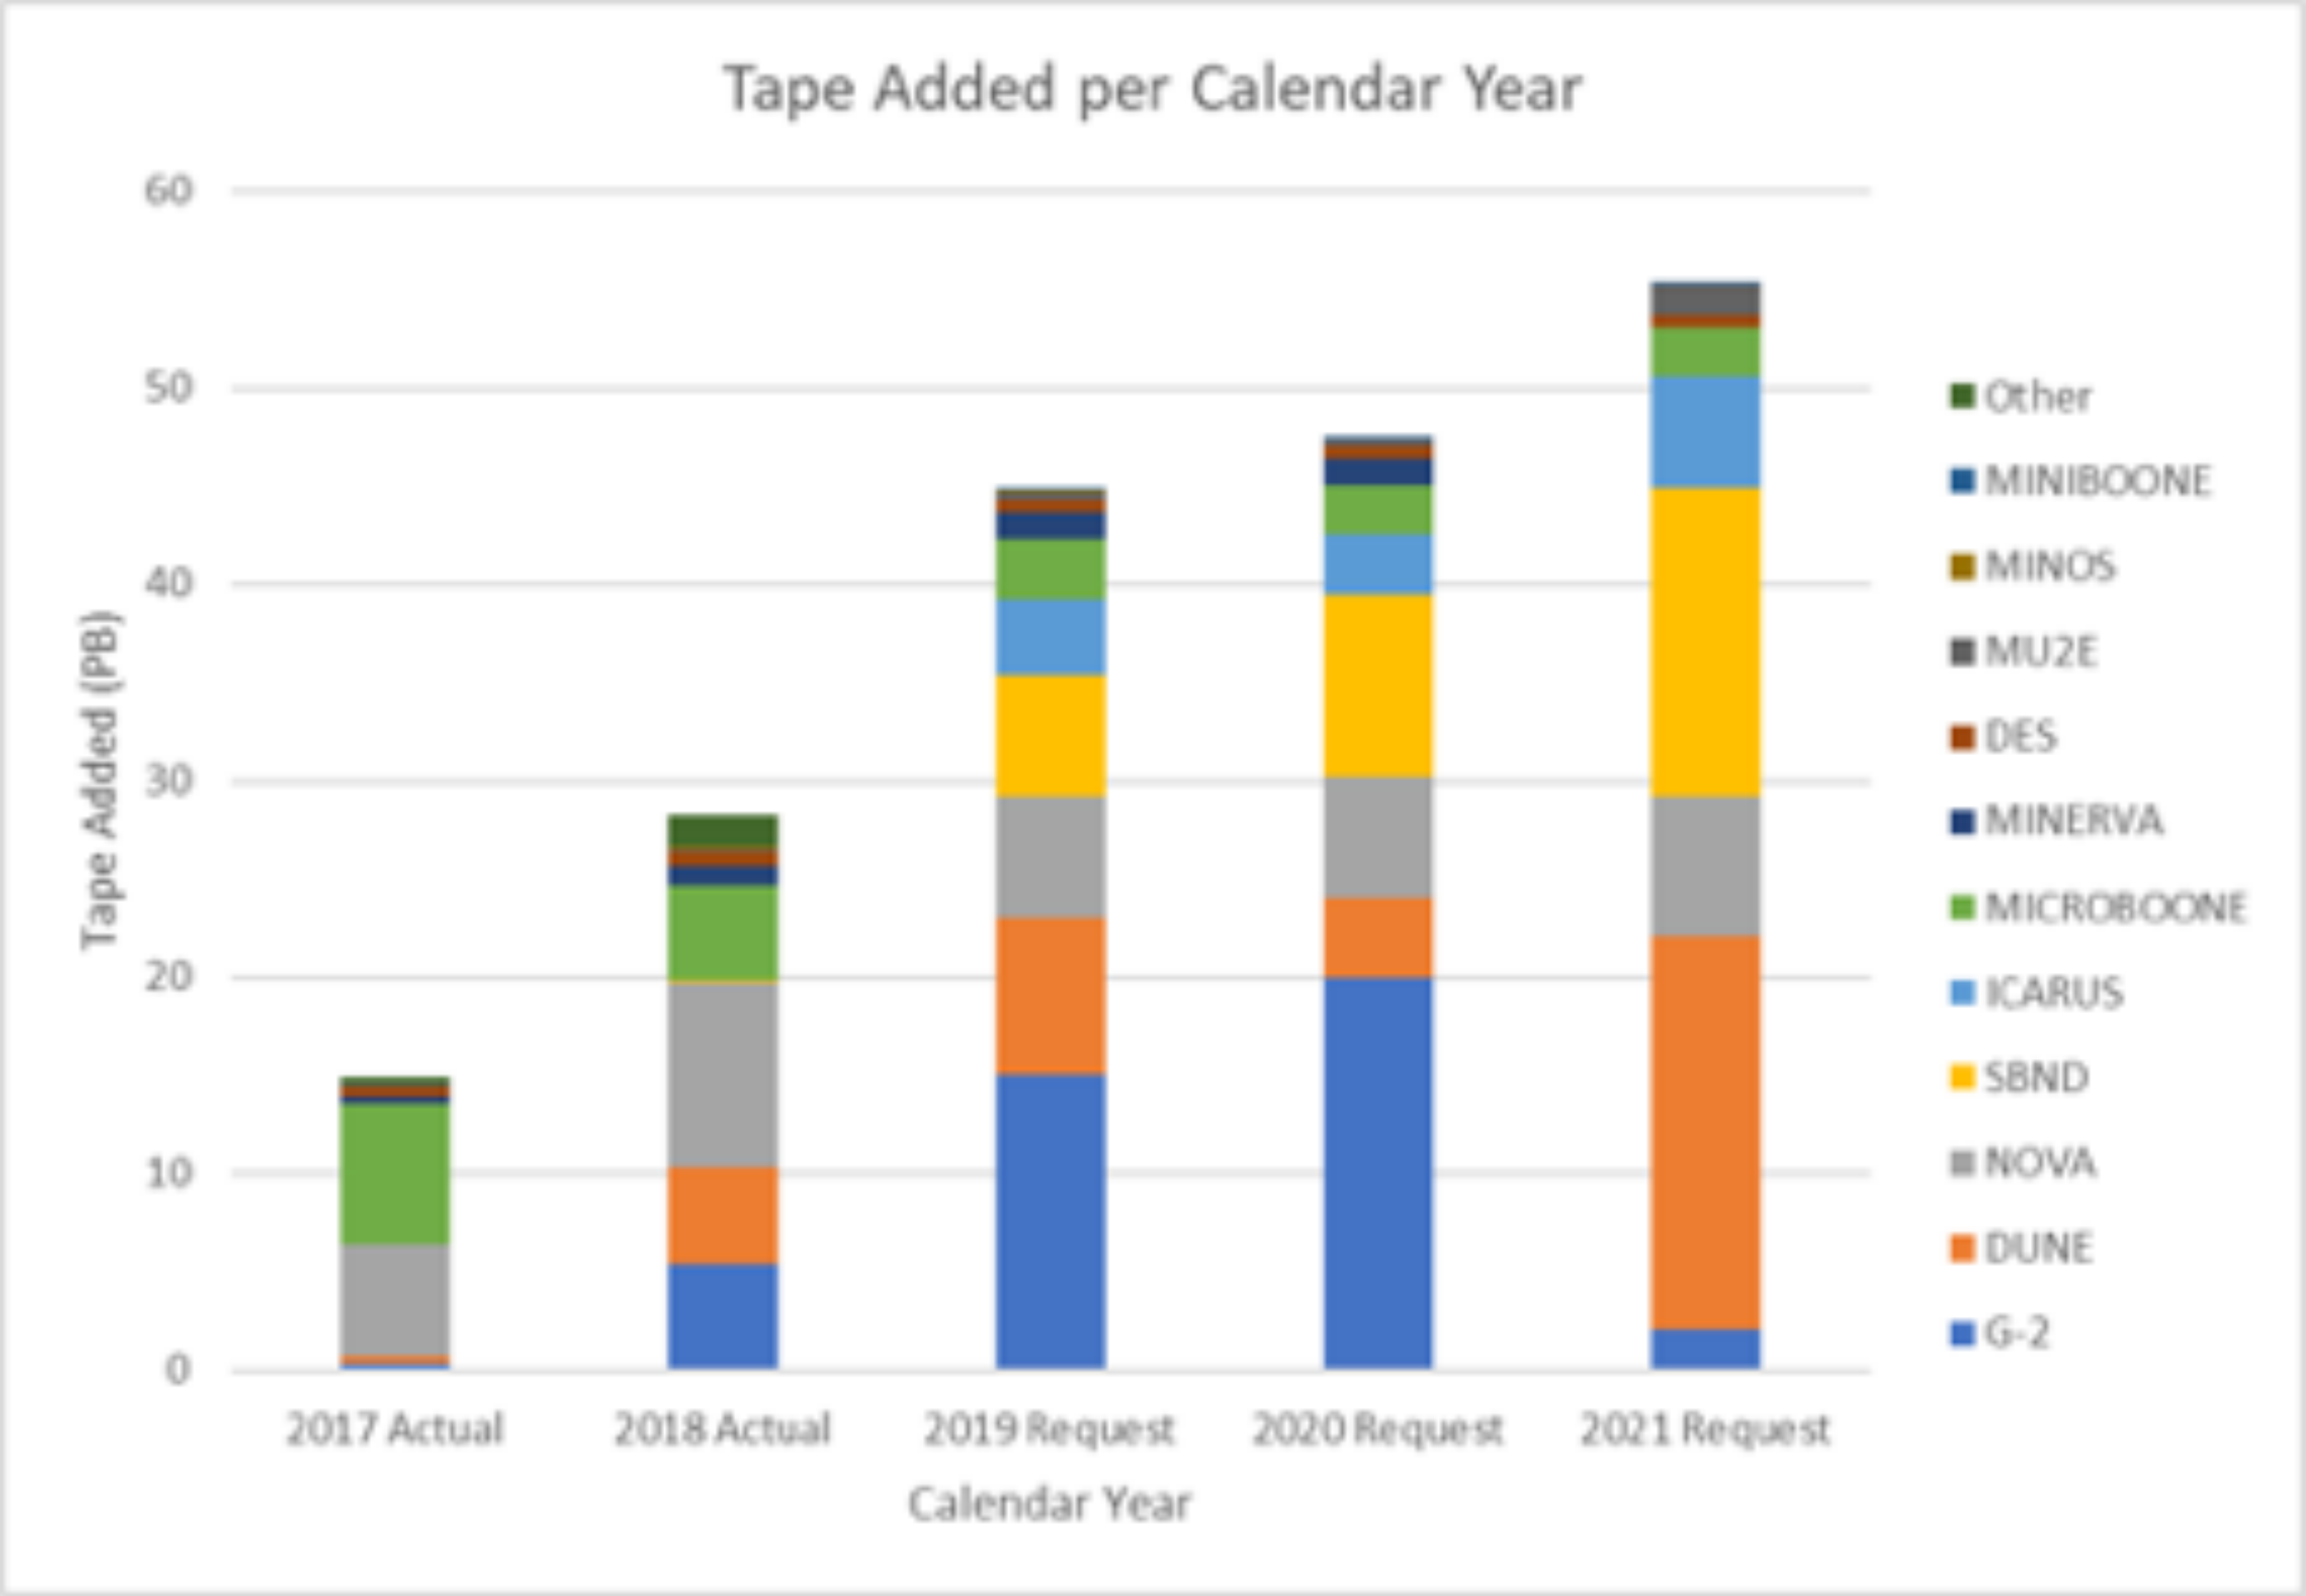

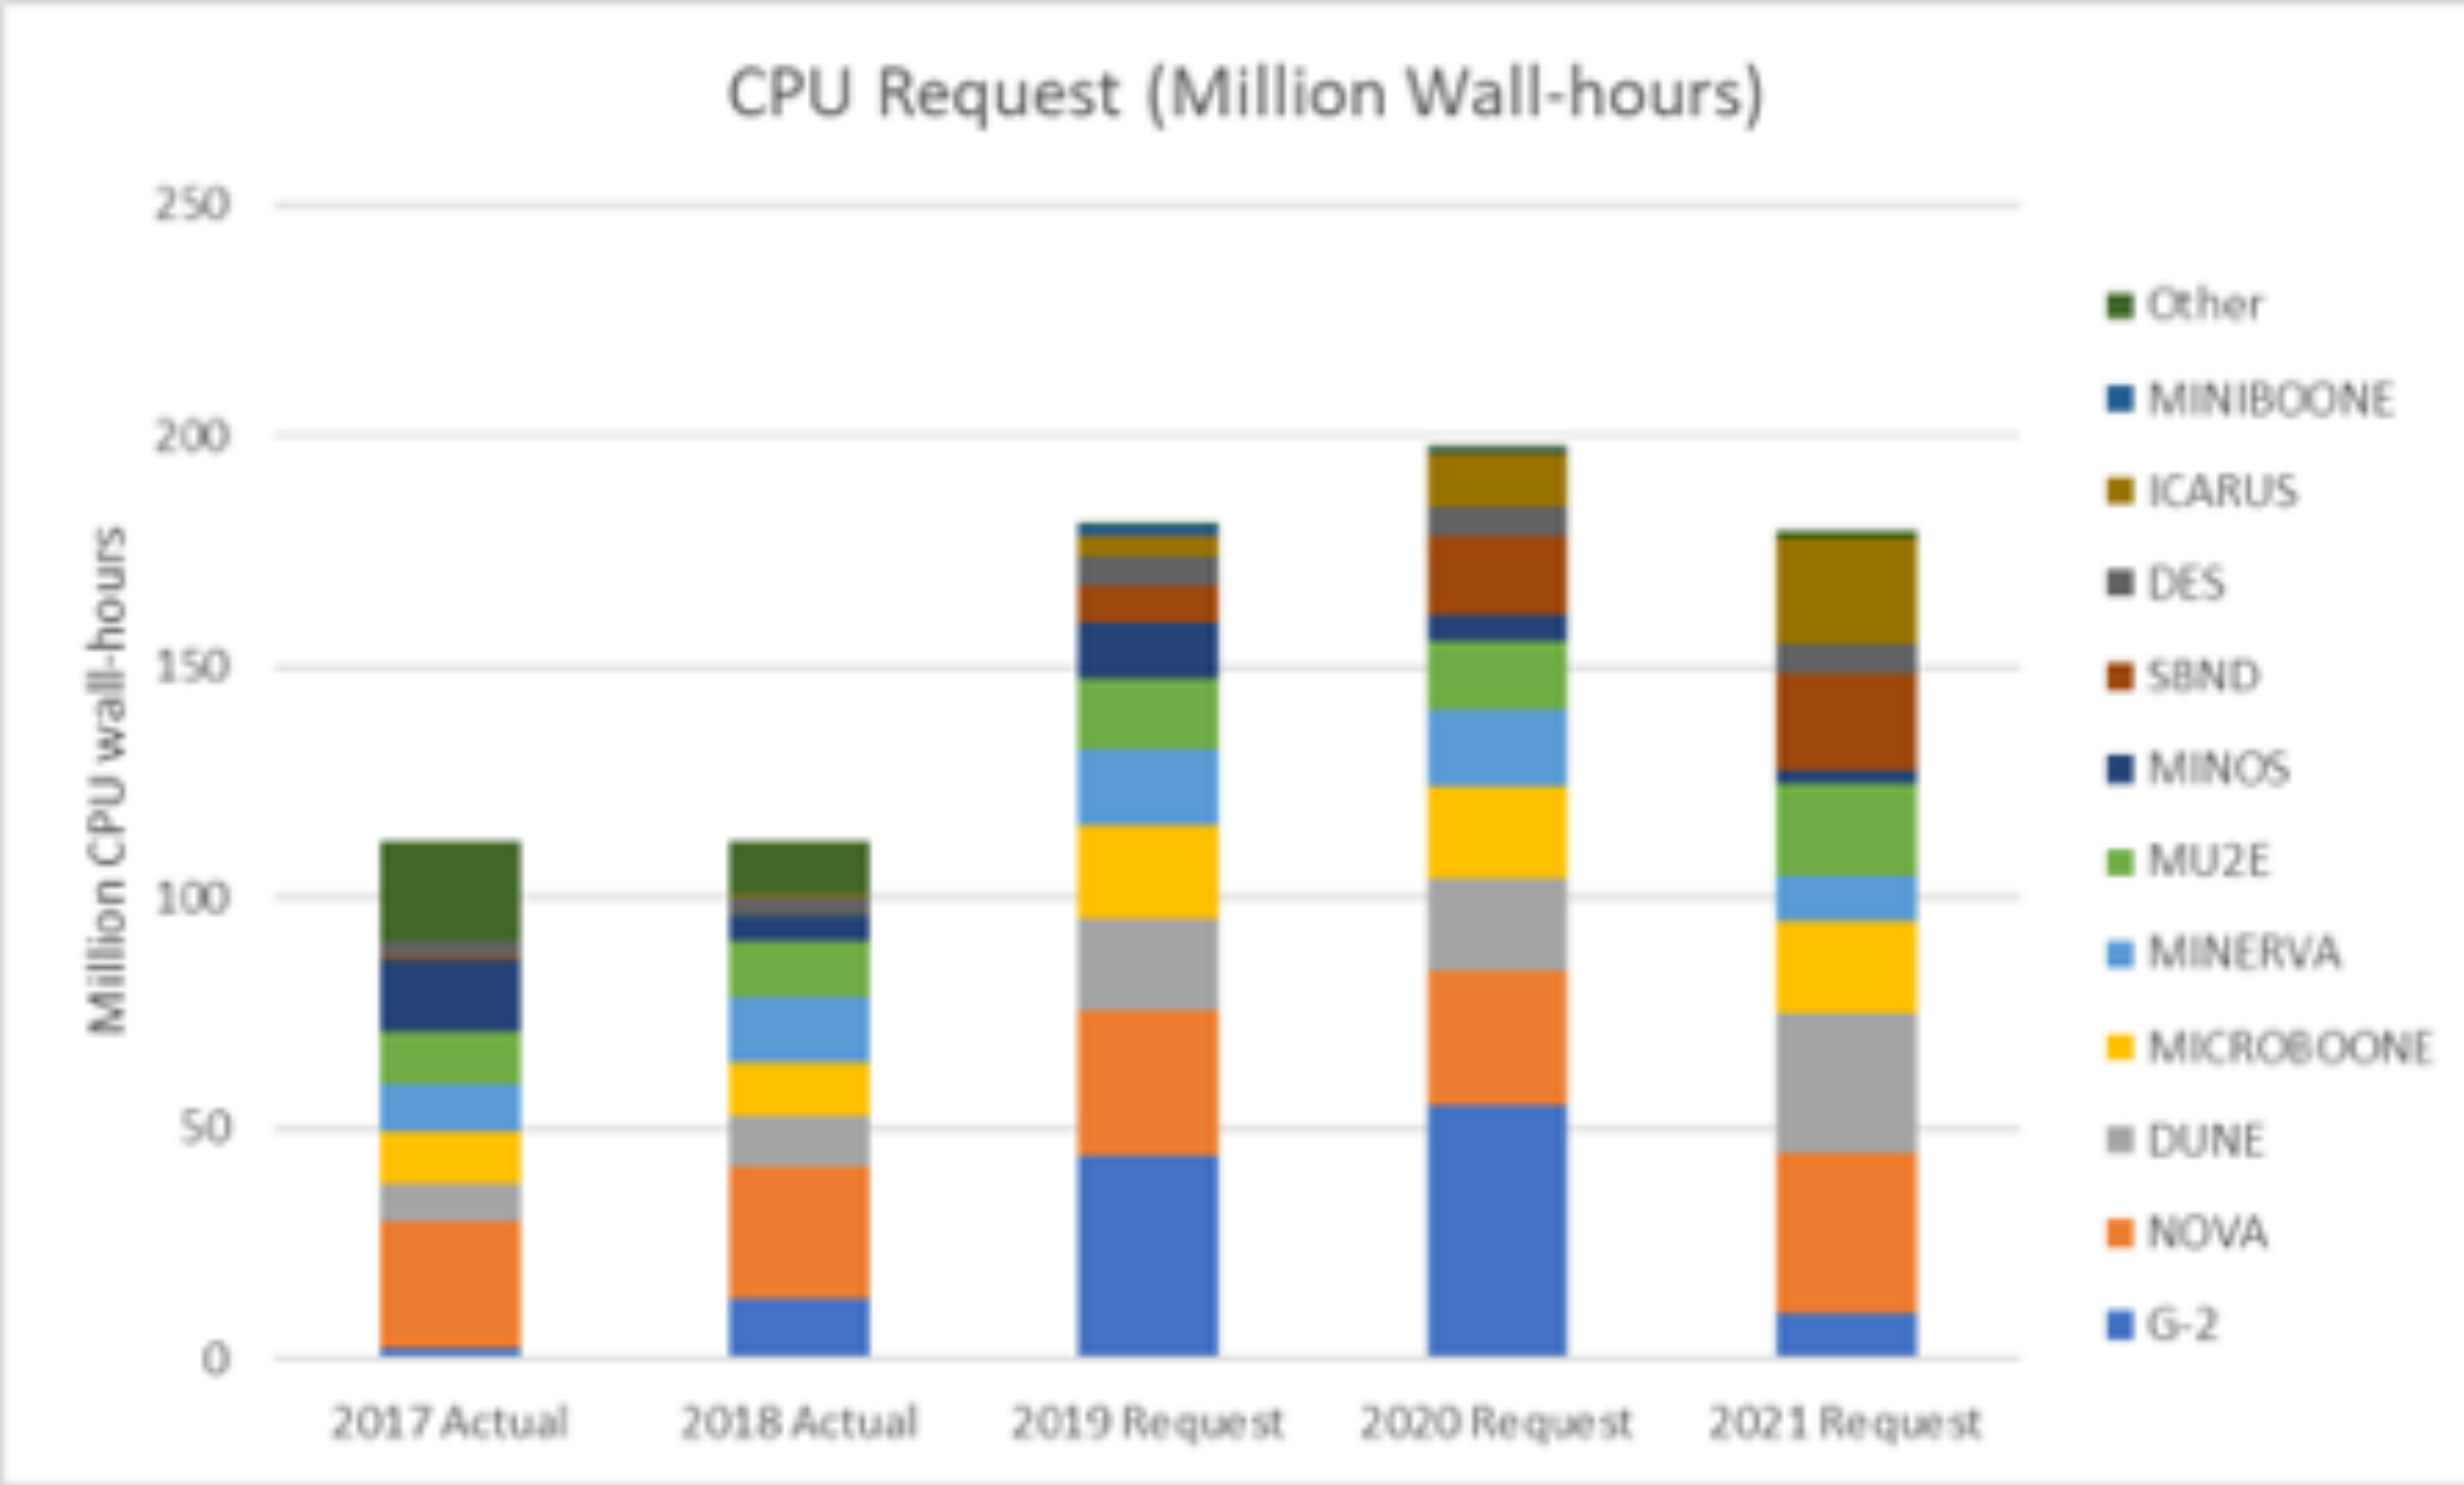

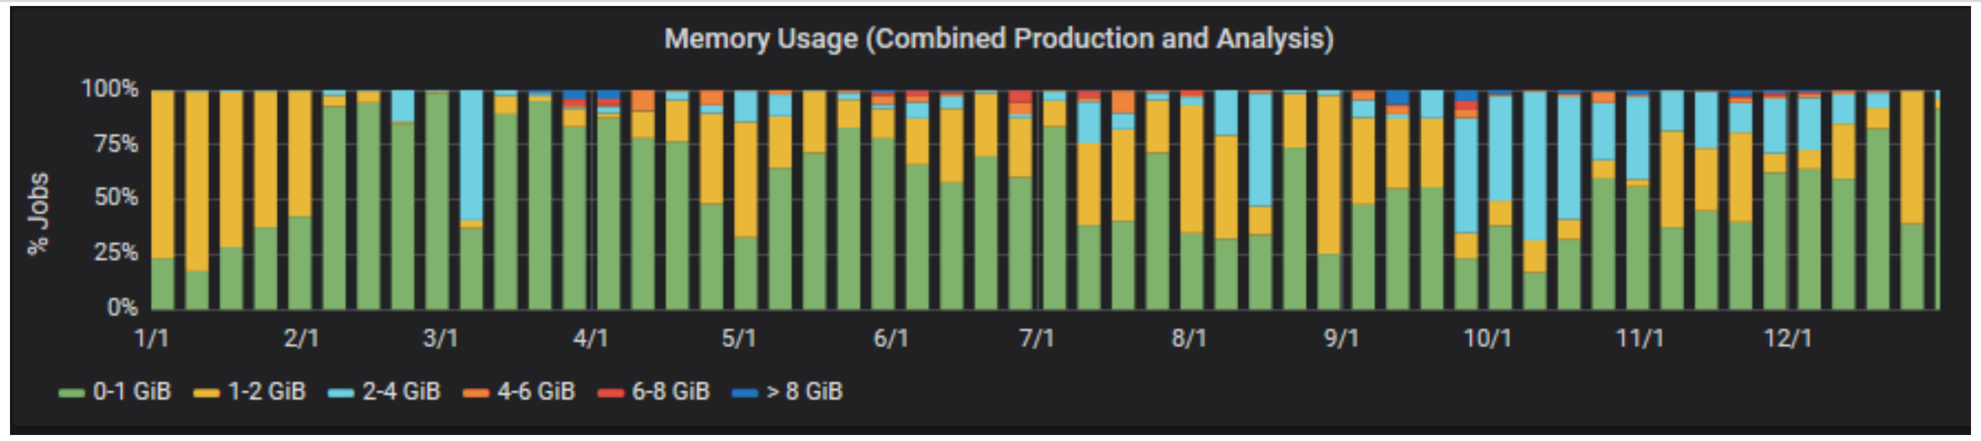

Fermilab

Effects of Postponing Equipment Purchasing

End of Warranty Reliability Period for Disk Servers

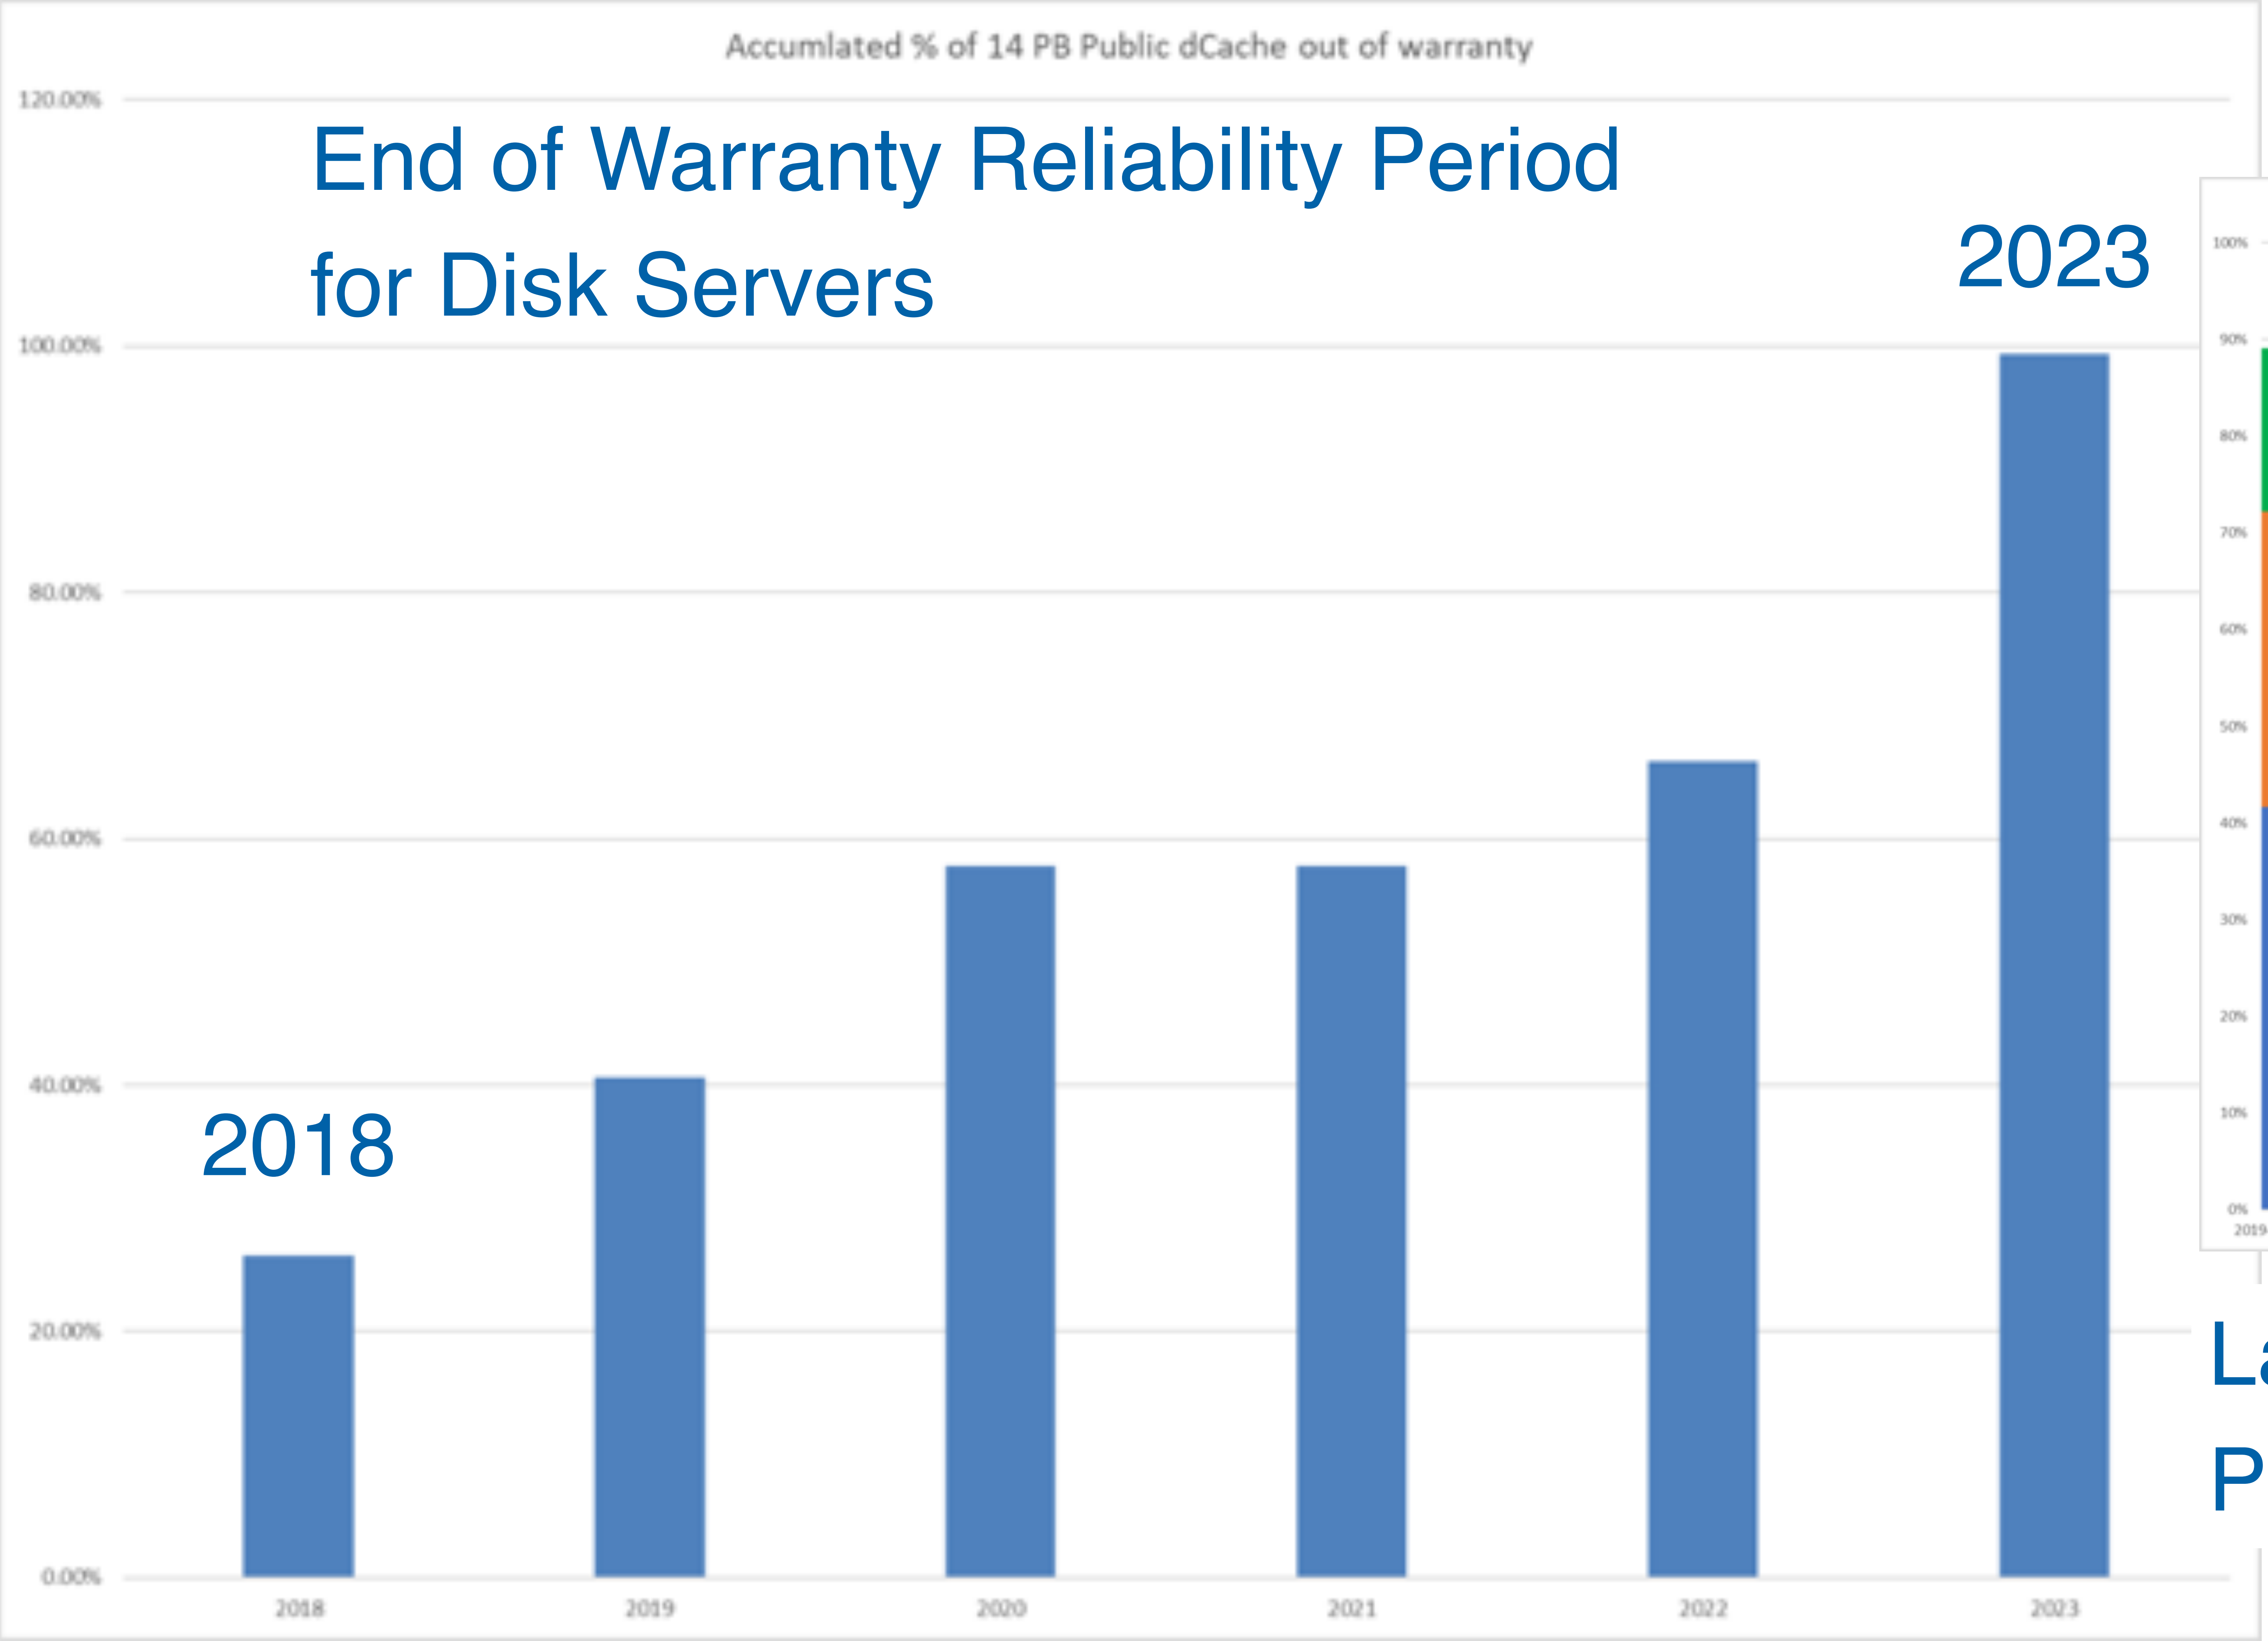

Reduction of FermiGrid Capacity 2019-2023

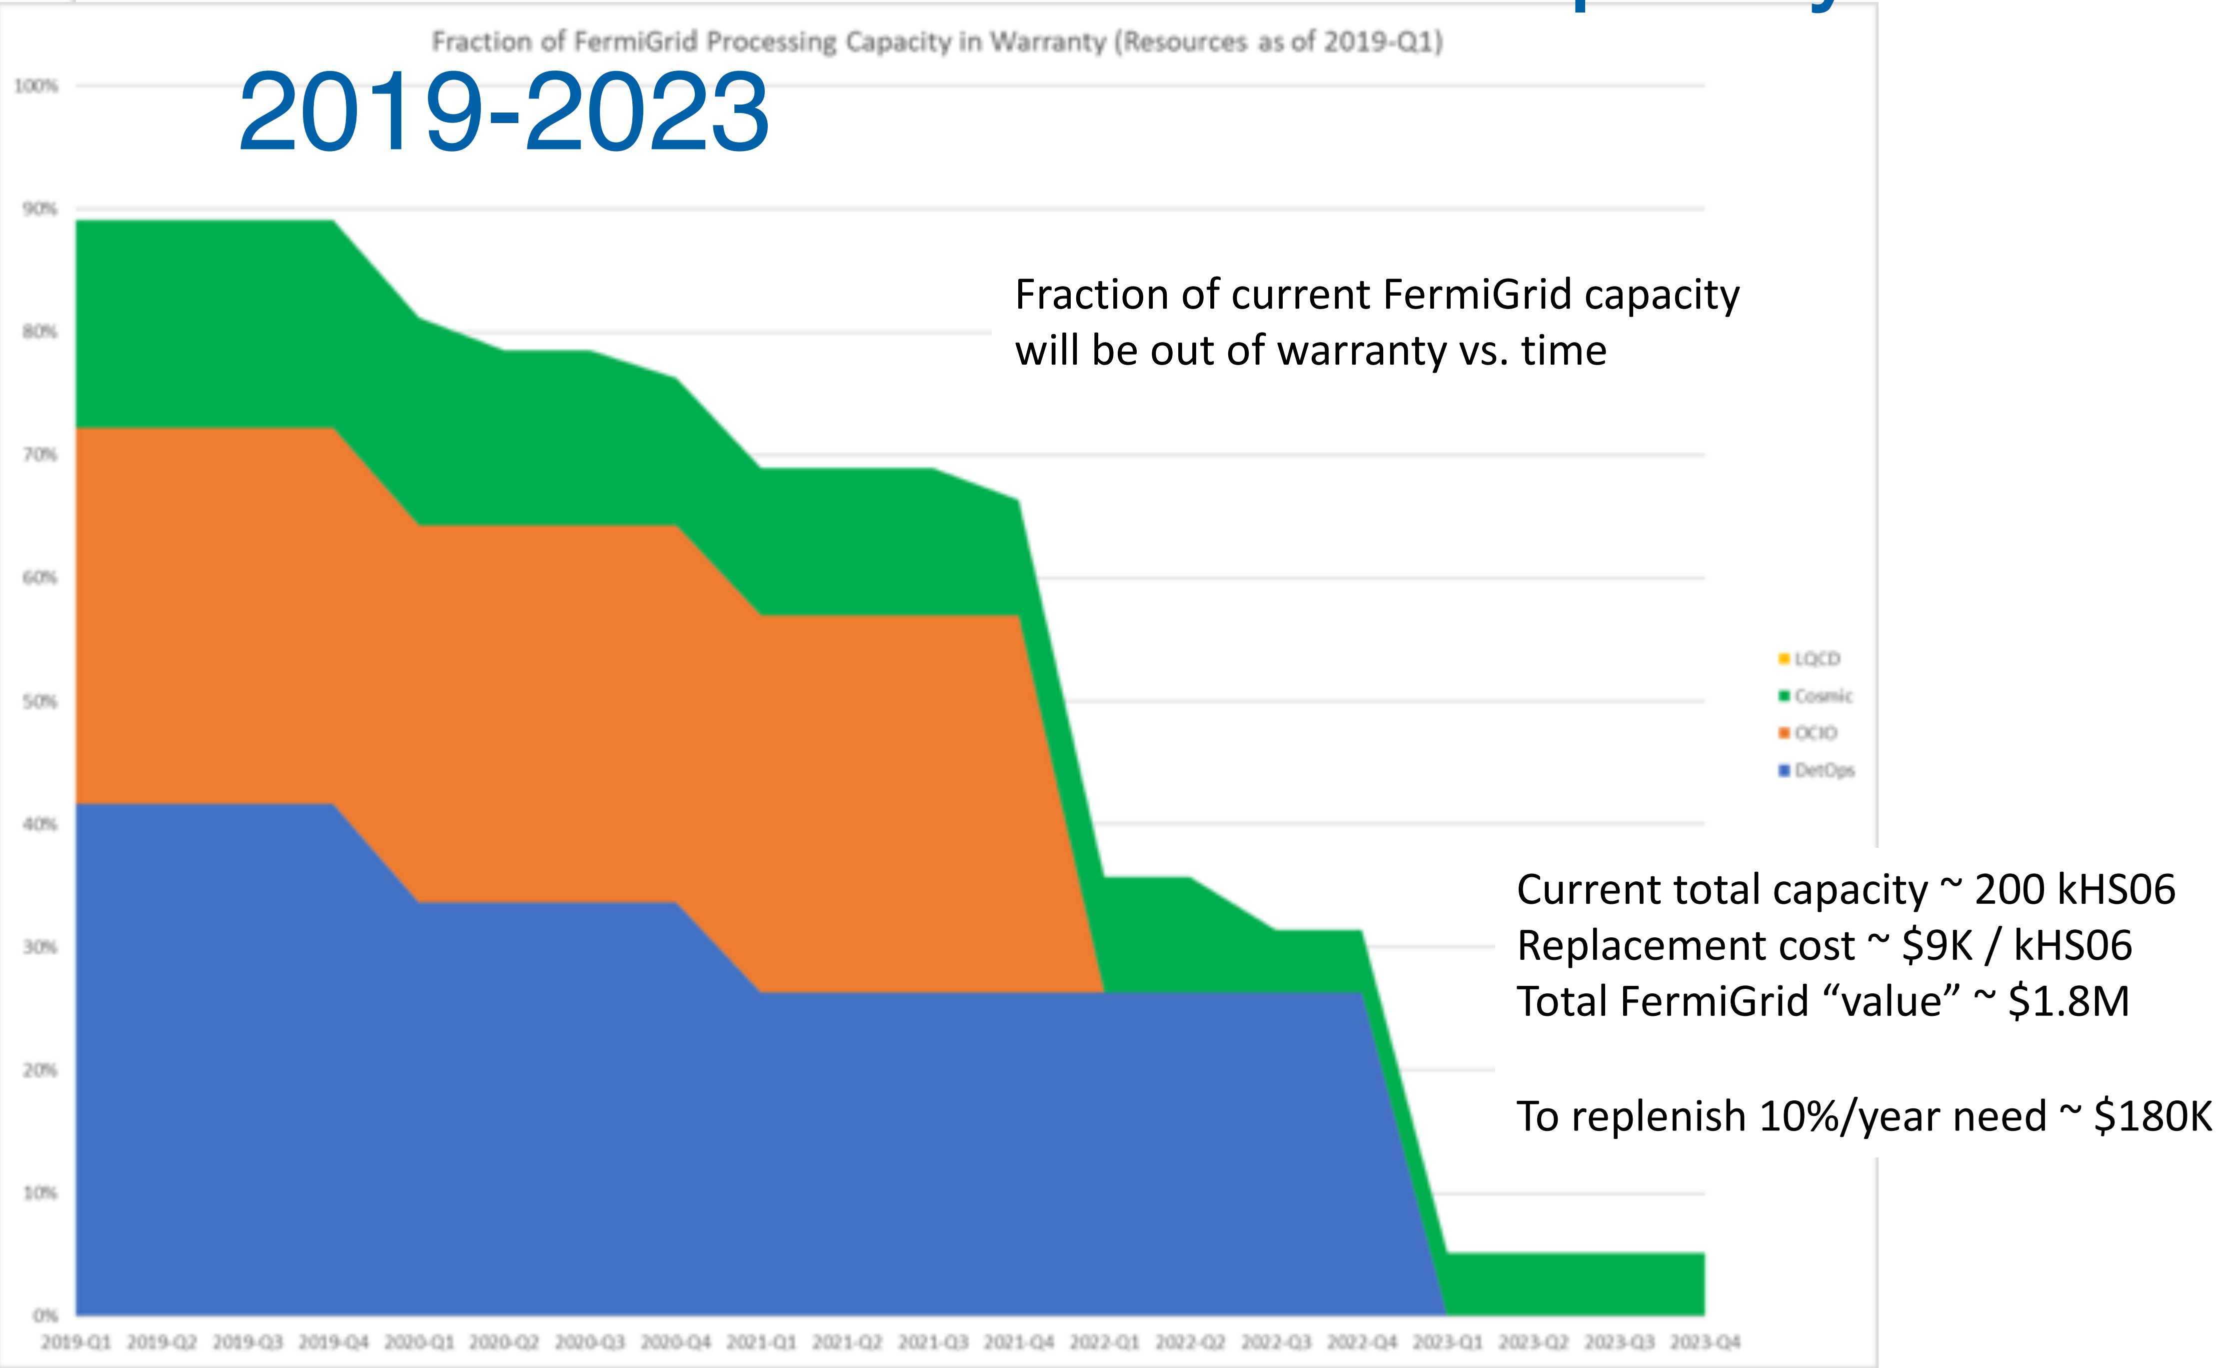

Last purchase of servers was in 2017

Population of Institutional Cluster is critical

Fermilab

SCPMT Recommendations

1. **Improve the SCPMT template by reexamining the technical metrics.** Make the responses available in advance to provide more time for discussions with experiments and of SCD's action plan. Have larger projects outline their computing models and methods used to estimate the requested resources.
2. **Improve efficiency of managing resources allocated to the experiments by developing well-defined policies for CPU performance and storage. Enforce policies via automated quotas and allocations. Develop tools to incentivize users who follow the policies.**
3. **Facilitate onboarding of the experiments and reduce the long-term direct support.**
4. **Storage resources and usage need a sustainable philosophy.** An example would be the NAS, which, as implemented, has led to dependence on expensive and old technology. The absence of high performance solutions has forced the experiments to use expensive storage systems in an inefficient way.
5. **Continue efforts to develop and implement common tools across frontiers.**
6. In light of constrained budgets, no flexibility remains for identifying and updating current services and infrastructure. To be a sustainable enterprise, **SCD should identify 5% of its budget that can be used for R&D activities toward future hardware/software advances.**

ICAC Recommendations

(*Excerpts*)

- **Look at ways to speed up adoption of federated identity** use as a building block of collaborative services, particularly needed for DUNE.
- **DUNE should be encouraged to draft a computing model**, in order that Fermilab (and other sites) can plan their facilities.
- **Fermilab should have a plan for how it becomes an international laboratory for DUNE**, what collaborative tools will be provided, etc. The plan should clarify the responsibilities of Fermilab as a host lab, and as part of the computing model.
- **The future storage strategy requires particular attention.** In particular, a vision and a roadmap is needed to address the needs in the Public cluster and a plan should be elaborated to address concerns over the sustainability of Enstore, possibly by adopting a solution with greater support in the community.
- **Within SCD we recommend that CMS and other projects should be less stovepiped.** This is a source of duplication of effort and inefficiency. This must be avoided for DUNE. Facilities and services should be as far as possible common across supported experiments, focusing on function rather than specific requested solutions. We encourage the computing management to continue to re-evaluate the organisational structures in the light of constrained resources and with an eye to the evolving needs of the lab and the experiments.

How can the PAC Support us?

Workforce Development

- In general the field lacks the training it needs to accomplish our goals.
- Encourage intellectual participation from University community in the computing challenges we are facing.
  - Would a Guest and Visitor's program for Software and Computing be possible?
  - Past collaborations between SW engineers and professors have been very fruitful.
- Agree that Education is important
  - We already do CMSDAS, FIFE Workshops, LArSoft workshop, Experiment led training
  - Doing more … Recent week long C++ class with invited Prof. Glenn Downing from UTA

	STRONGLY AGREE	AGREE	SLIGHTLY AGREE	SLIGHTLY DISAGREE	DISAGREE	STRONGLY DISAGREE	TOTAL
This training was a worthwhile investment of my time.	73.33% 22	23.33% 7	3.33% 1	0.00% 0	0.00% 0	0.00% 0	30

Reorganizations

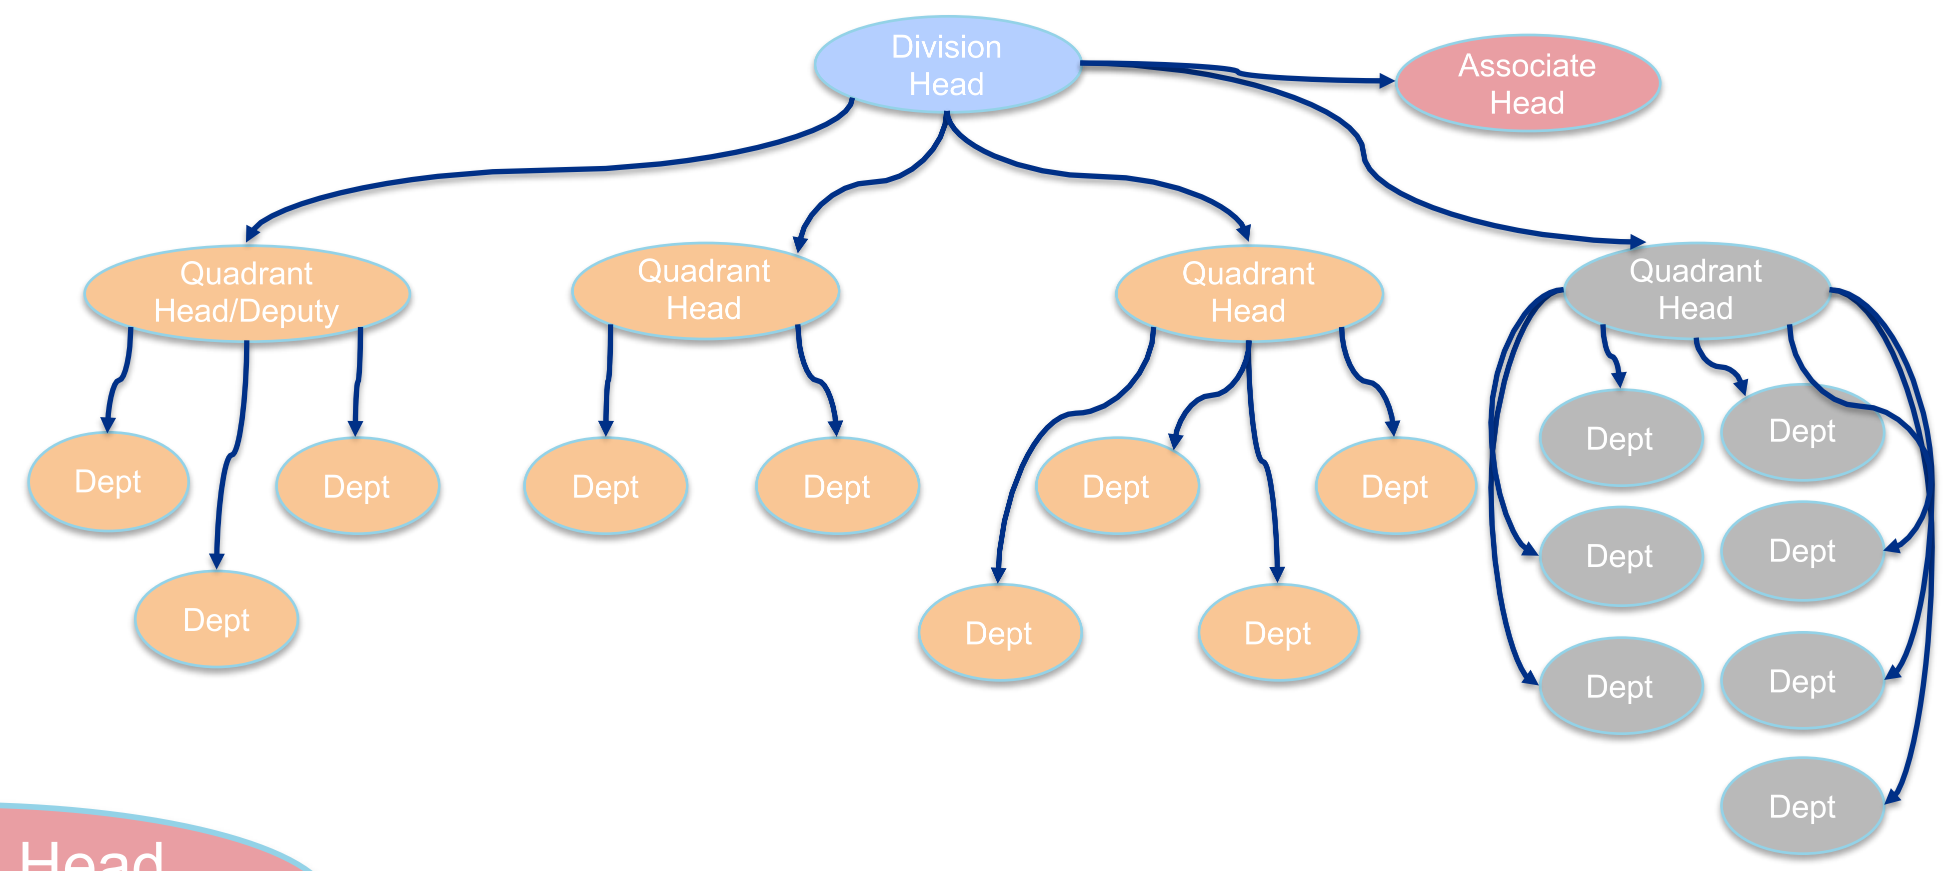

- Reorganized OCIO to eliminate non-essential groups making room for needed skills
- SCD reorganization
- USCMS S&C reorganization

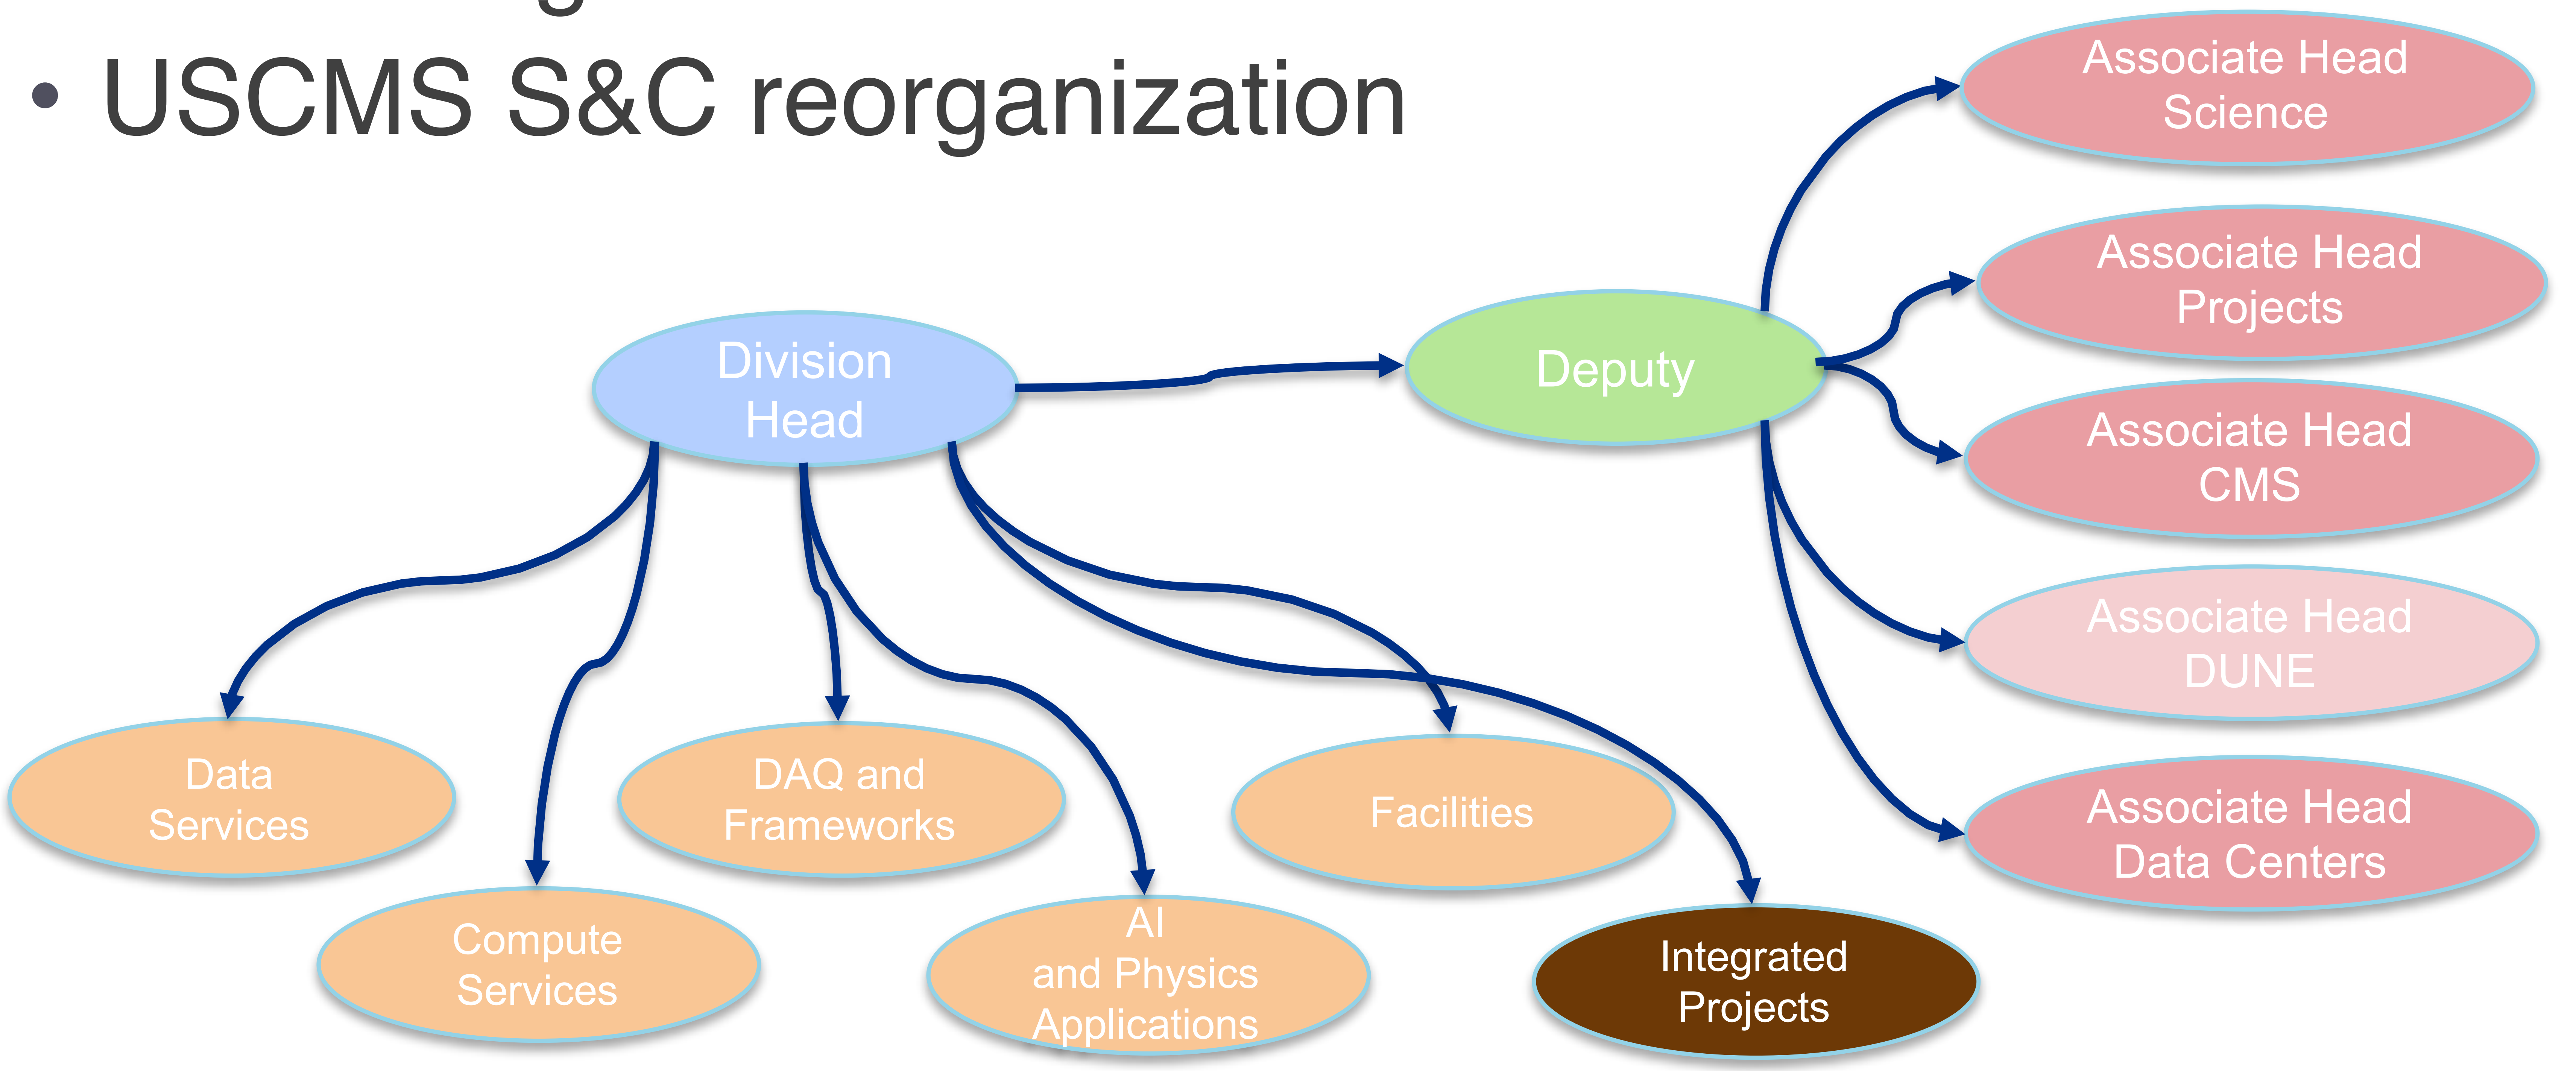

Departments labeled by description, not final name

- Improve communication by removing a layer of management; also making it lighter.
- Elevating cross cut Projects to emphasize technical leadership.

Workforce Training

The Office for Professional & Organization Development has scheduled

C++/STL

May 20-24, 2019
IARC Lecture Hall
with

Professor Glenn Downing
Assistant Professor of Instruction
Department of Computer Science
University of Texas at Austin

Professor Downing will present a five-day class on the syntax and semantics of C++ and the Standard Template Library (STL)

This course is designed for software developers and technical managers, post docs and graduate students.

To enroll, visit The Office for Professional & Organization Development homepage at:

http://www-esh.fnal.gov/pls/cert/schedule.show_course_details?this_course_code=FN000632&this_instr_type=CR&this_fermi_id=

Software Carpentry Workshop - Fermilab - (1-2 April 2019)

iris hep
Institute for Research and Innovation in Software for High Energy Physics

Software Carpentry Workshop

1-2 April 2019
Fermilab Feynman Computing Center (Room FCC2A)
US/Central timezone

- Overview
- Scientific Programme
- Timetable
- Contribution List
- Registration
  - Registration Form
- Participant List
- Getting a Fermilab Computing Account
- Map and directions (Fermilab)
- Main roads around Fermilab
- Fermilab site map
- Carpentry Link
- Supplemental materials
- Pre-Workshop Survey
- Post-workshop Survey

Support

We are very excited to announce the Software Carpentry Workshop at Fermilab. Our experiment-specific and advanced software trainings has shown that participants' kno software skills can be quite variable, depending on their particular background. Some basic skills from university courses or self-training, but holes are very common. This i ability to profit from the advanced trainings being offered. In offering the workshop t establish and provide a uniform set of basic skills for all HEP graduate students and p broadening participation from institutions lacking such courses.

The topics will cover *python, python plottting, access physics data in Python with Py as well as manipulating irregular data as jagged arrays.*

NOTE: The regsitration is strictly limited to 25 on first come first serve basis. There will be a waitlist of 5 in case a spot opens up. To be waitlisted, send email to Sudhir Malik (malik@fnal.gov).

NOTE: Coffee/cookies will be served. Lunch is on your own.

Tutors:

David Yakobovitch - Enterprise Data Scientist at Galvanize, AI Instructor

Will Trimble - bioinformatician, based at ANL

Jim Pivarski - Physicist, Princeton University

Organisers:

Sudhir Malik (University of Puerto Rico Mayaguez)

Peter Elmer (Princeton University)

- 25 Participants + tutors
- mostly neutrino community

https://www-esh.fnal.gov/pls/cert/schedule.show_course_details?cid=11499

Summary

- It's not possible to do science without computing
- The nature of computers is changing -> heterogeneous hardware
- The way computers are used is changing -> new algorithms and ML
- High Velocity Exascale Data is still a core capability at Fermilab
- The challenge is so great that we need computing and domain scientists to work together.

Back up

Evolving the Fermilab Facility

- We need to move more broadly to an institutional cluster (IC), model for HEP computing. The idea is that different programs by a fraction of the cluster that they get priority access to, and share when not in use.
  - Efficient sharing infrastructures have enabled a much broader sharing of resources then could be envisioned years ago.
  - It is easier now to guarantee that science customers get what they paid for.
- Recently LQCD (the project behind the USQCD collaboration) bought in to the BNL institutional cluster , and are very happy with the arrangement.
  - Volume discounts on original purchases
  - Shared support ongoing operations
- Fermilab is transitioning to an IC as FermiGrid ramps down.