



# Fermilab Computing Sector (CS)

Elizabeth Sexton-Kennedy, Jon Bakken, Panagiotis Spentzorous

Computing Sector All Hands

5 June 2018

# Outline

- Nigel's top 5 and what it means for Computing
- Some changes in the ORG chart
- FY18 and 19 Budget update
- News from HR and Communications
- Lab Modernization for Computing
- Building for the Future of Software and Computing
- R&D to Bridge the Gap

# The Top 5

# Top 5 messages from director's April 27 all-hands meeting: 1

- Safety is paramount at Fermilab and the new **South Dakota Operations Division**.
  - Fermilab is now operating two sites: Batavia and South Dakota.
  - Although SURF is not a DOE lab, we will work together for the next few decades as partners in a common goal: great science.
  - There have been a number of incidents at SURF.
  - It is critical that we do not have a serious safety incident, either in Batavia or South Dakota.
- Hosting the long- baseline neutrino facility and deep underground neutrino project (LBNF/DUNE) is a **lab-wide** effort.
  - Fermilab will integrate work from across the lab.
  - Key areas include the accelerator complex and PIP-II, the Integrated Engineering Research Center, short-baseline neutrino program, **advanced computing capabilities**, and the necessary support systems such as safety, procurement and finance.

# Top 5 messages from director's April 27 all-hands meeting: 2

- Broad support for LBNF/DUNE and PIP-II
  - Outstanding LBNF/DUNE FY18 budget support allows restart for beamline design here at Fermilab.
  - DUNE continues to secure **international support** (1,078 collaborators from 175 institutions in 32 countries).
- Integrated Engineering Research Center
  - The Center will provide state-of-the-art laboratories, and engineering and technical spaces within a collaborative environment.
  - Will host scientists and engineers from around the world at Fermilab in support of the international DUNE experiment, LBNF and PIP-II.
- LHC/CMS
  - CMS physics program makes historical step in the understanding of the Higgs boson: first observation of the associated production of Higgs bosons and top quarks, ttH.
  - Fermilab leads the High-Luminosity LHC Accelerator Upgrade Project that will provide components for the upgraded accelerator at CERN.

# How is SCD supporting LBNF/DUNE?

- The merging of computing for CMS and the IF experiments years ago has given us the experience of supporting the computing of an international collaboration.
  - protoDUNE will take data in Sept. and Oct., and we need to be ready to help them store, process and analyze that data at CERN and FNAL.
  - Their second data challenge in April was promising. The rates achieved will meet their baseline goals. However, there are stretch goals....
  - DAQ rates are still uncertain. We may need to store tens of PB of data.
  - We need to help them develop winning strategies for writing their technical design report.
  - DUNE uses more scientific computing services than any other Fermilab experiment.
  - Many in SCD are members of the scientific collaboration in leadership positions.

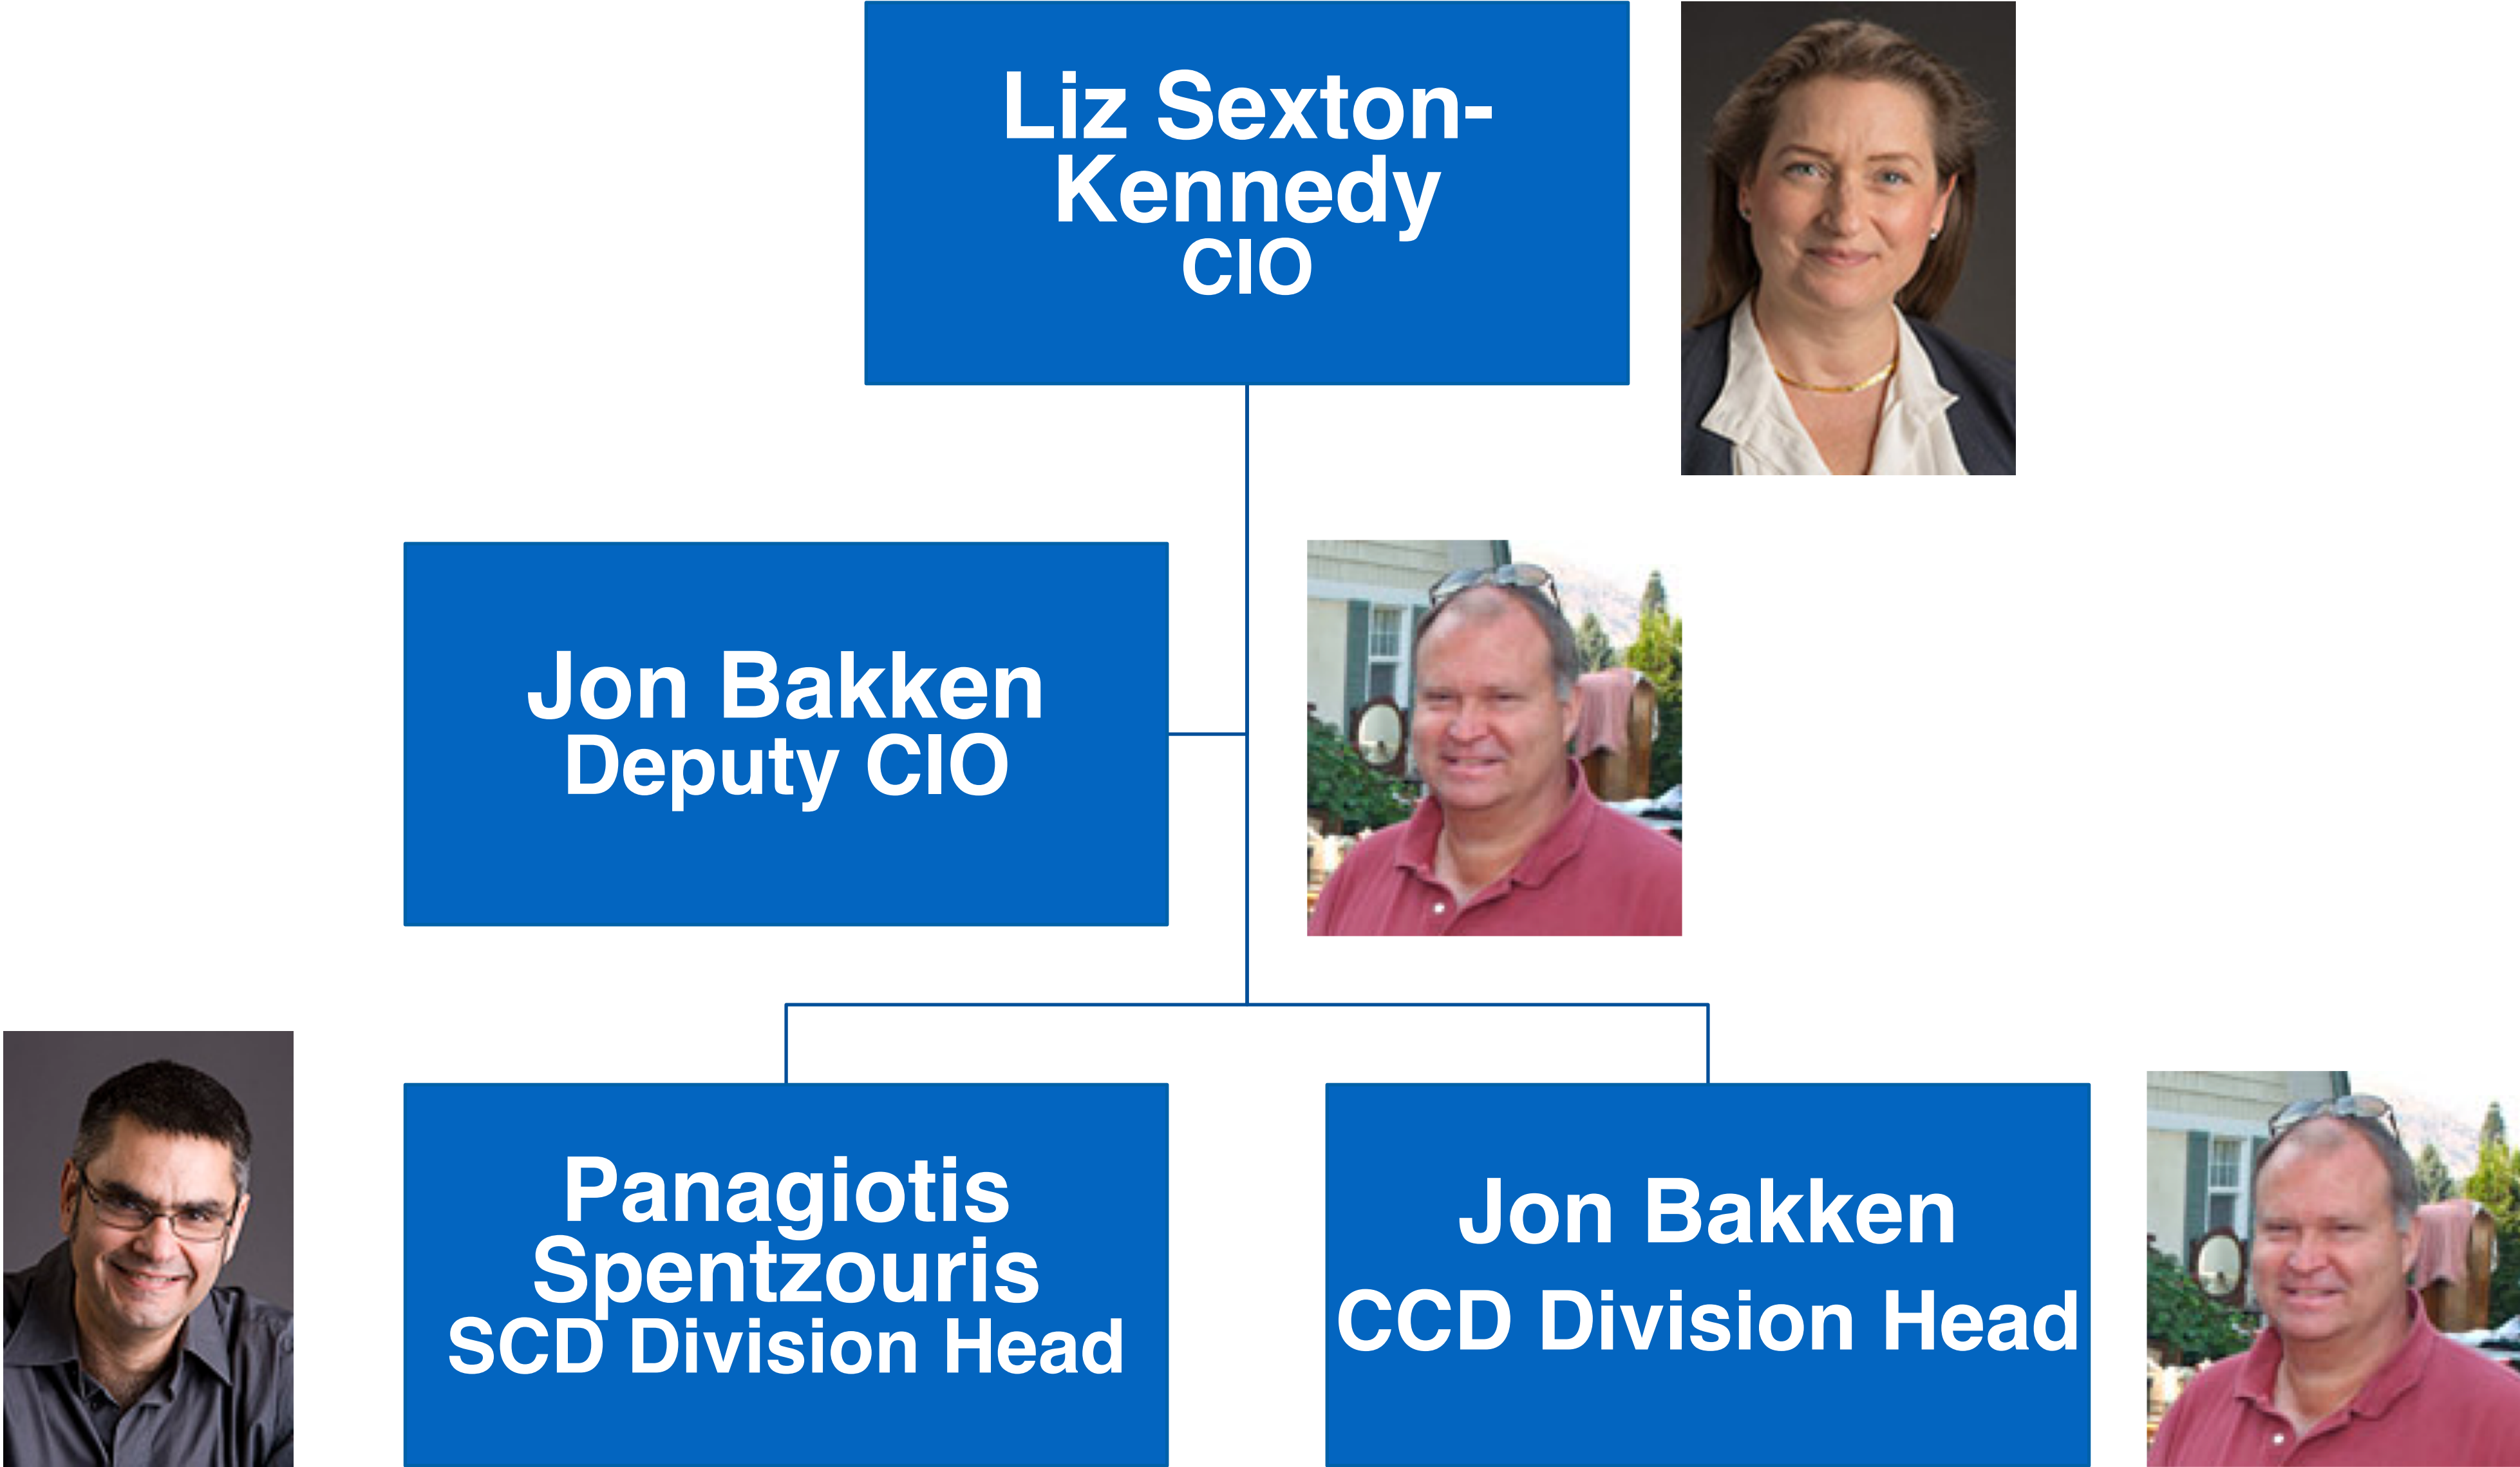
# How is CCD supporting LBNF/DUNE?

- We are helping to bring the Fermilab way of doing things to SURF.
  - Now that we have a division in South Dakota, it needs an IT management system for Fermilab office activities.
  - **Bob Sieloff** will be our IT liaison with the backup of the rest of his department led by Mike Rosier.
- Networking to SURF
  - For the time being, we'll use the VPN to secure access to Fermilab business systems.
  - Eventually this will become a limitation.
  - Bandwidth to the site is an inadequate 1Gb/sec connection.
  - Our networking group is starting the process of applying to ESNET to boost networking for both enterprise computing and the science DMZ.


# Organizational Changes



# Organizational Chart



# Organization Chart: 2



# New Hires since Dec 2017

**Sudha Balakrishnan** - CCD/Business Infrastructure

**Gabriela Garcia** – OCIO/Financial Management

**Josephine Fazio** – OCIO/Project Management

**Joshua Kenward** - CCD/Enterprise Services Operations/Unix Server Services

**Andy Li** – SCD/Systems for Scientific Applications/Scientific Software Infrastructure

# HR Updates and Reminders

**Stacey Vassallo**

**Computing HR Partner**

**[masters@fnal.gov](mailto:masters@fnal.gov), x4363**

**FCC Office: FCC 1 West (near coffee station)**



## **General Timeline for Performance Review**

- **Employee self-evaluation is June 5 – July 3**
- **Manager evaluation completed and approved – July 27**
- **Ratings finalized/salary review completed – Aug./Sept.**
- **Review/salary increase discussions – late Sept/early Oct.**
- **2018/2019 goals anytime – due Oct. 31**
- **Pay Increase effective Oct. 1**

# Reminder about Interactions with Reporters

- Media coverage builds support for Fermilab and for our science. The Office of Communication is actively seeking to grow media coverage about the lab.
  - Make sure to let us know at least 2 weeks in advance when you have a publication coming up, or another milestone worthy of news coverage.
  - You are encouraged to answer questions from reporters about published results, and science/research topics where you are an expert. You must notify the Office of Communication about the exchange.
- DOE and Fermilab rules regarding how we work with reporters have changed over the last decade. See the lab's Communication Policy for more.
  - You must get advance approval from the Office of Communication to answer questions from reporters about lab or DOE policy, management or budget; 413.3B construction projects; emergencies/crises; or topics not connected with your work or the Fermilab research program.
  - All visits by reporters to the Fermilab site must be approved in advance by the Office of Communication (media@fnal.gov, x3351)

# FY18 and 19 Budget update

- Overall HEP budget up from 825M\$ (FY17) to 908M\$ (FY18)
- Expect the final FY18 research budgets down ~10% from FY17 as the new funding is going to Projects as forecasted.
- Every June the lab submits Field Work Proposals (FWPs) to OHEP to support Research and Operations for the next fiscal year. These describe tasks to be performed and estimated costs with budget guidance from the Directorate. Panagiotis has been preparing these for SCD.
- Lab budget guidance for Research in FY19 is expected to be ~5% lower than FY17, but ~5% higher than final FY18.

# Lab Modernization Plan

- Purpose
  - Develop a plan that defines what we need (expertise and facilities) to achieve our goals and satisfy commitments between now and 2030.
- Steering Committee
  - A cross-function team drawn from across the lab, chaired by Mike Lindgren, Accelerator Division Head.
  - Jon, Panagiotis and I are working with Erik Gottschalk and Bill Boroski to update and refine the Computing lab capabilities.
- This summer, all employees need to update their skills in FermiWorks to ensure we account for all of our existing talent and abilities.
  - The skills will be used to establish a baseline for expertise we have at the lab today and what we will need for the future.
  - This will help us figure out our hiring and training needs. Our people are our greatest asset.

# Building for the Future of Software and Computing



# 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.
- BNL and Fermilab have a big head start as BigData facilities successfully delivering for their LHC users. We need to expand the consumer base for these services.
- ANL and LBNL have HPC centers that have significant DOE investment, which HEP must learn to more fully leverage. HEPCloud should provide a low cost of entry to these. It is our “science gateway”.
- We need to develop a national cyber-infrastructure to serve the needs of the scientific community and have dynamic sharing of this resource.

# DUNE challenges

- If you multiply the front-end data taking rates by the number of channels in the full DUNE detector, the full stream collection rates would be 145ExaB/year-impossibly large.
- Much of the detector research will go into reducing that to reasonable levels. Heidi assures me that Trigger and DAQ projects have committed to constrain their designs to 100Gbit/s and 30PB/year.
- Suppression of  $^{39}\text{Ar}$  decay, cold electronics noise, space charge effects, argon purities all play a role.
- This means that protoDUNE already has challenging data needs from now through 2020.

# Projecting Needs for 10 Years - CMS

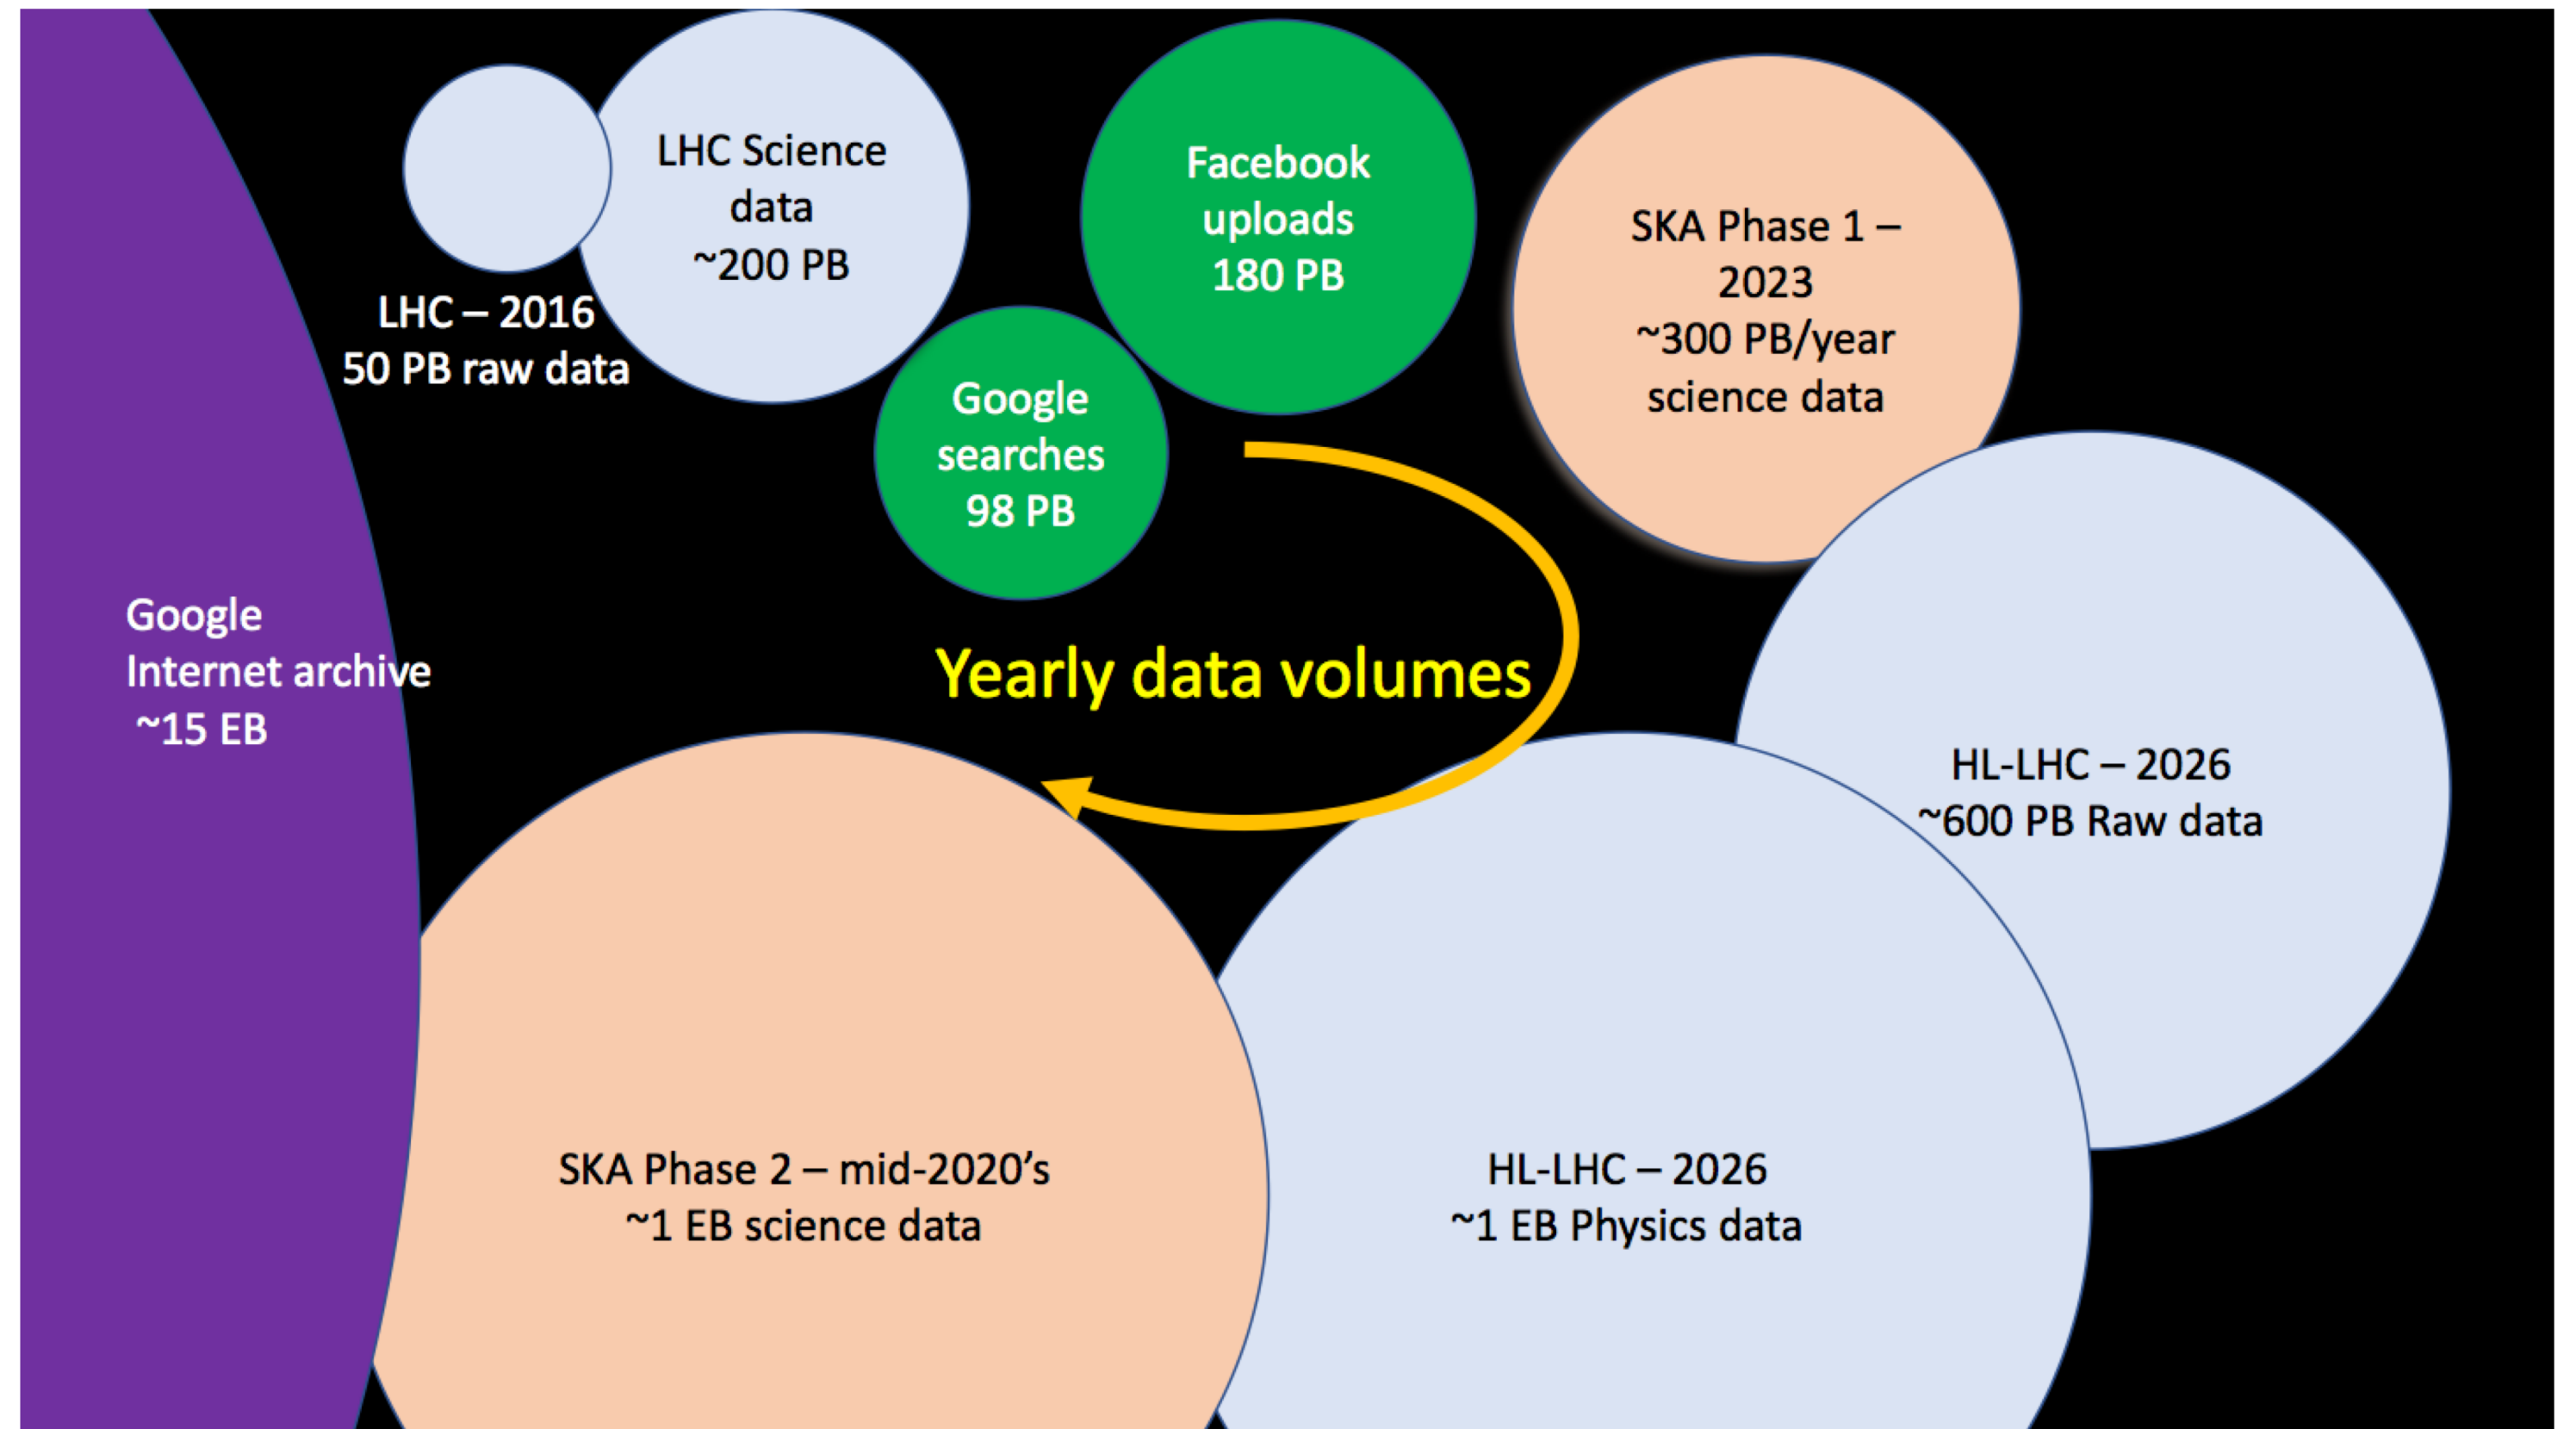
- The biggest challenge is not knowing how much and where you can economize without hurting the science goals.
- Last year CMS took its current computing models, made a naïve extrapolation and found it was unaffordable. This is true for others as well.
- This kicked off the flurry of activity to prove we could use smaller formats.
  - CMS can now satisfy 95% of its users with the miniAOD and is investigating an even smaller format.
  - The cost is that these derived formats will need to be calculated more frequently. It explicitly trades processing for storage, organized for unorganized.
- These strategies are a big help. However, we still face the question, “How do we provide even organized access to hundreds of PB of data?”

# Evolving the Facility

- We need to move more broadly to an institutional cluster (IC) model for HEP computing.
- The U.S. is the only country that has its T1 LHC virtual organizations hosted at different sites.
  - Efficient sharing infrastructures, like a global batch system, have enabled much broader sharing of resources than could be envisioned 20 years ago.
  - It is easier now to guarantee that science customers get what they paid for.
- Going forward there are too many customers to allocate one site to each.
- 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
- FermiGrid used by IF is also an IC with a long successful history.

# Yearly International Data Needs

- We do this today with a worldwide computing grid. It will need to grow.
- Reliable and performant networking is key to our federated data model.
- Usage of this infrastructure will have to expand to support other HEP domains as well.
  - LSST plans to collect 50PB/year.



# Meeting the Needs of All: LQCD, Accelerator, Cosmic

- Many of our users have multiple needs:
- **LQCD** requires InfiniBand-connected high-performance machines including processing accelerators.
  - They also need data movement and storage facilities for the results of their calculations.
  - Turning the data into science
- **Accelerator modeling**
  - The easy one, since they have been using HPC for years
- **Cosmic**
  - Largely a high-throughput application requiring nodes with high memory capabilities
- **HepCloud**
  - This project is our entry/offering for providing a science gateway

# R&D to Bridge the Gap

# Community White Paper <sup>1</sup>

- Inspired by the P5 process and guided by its goals
- The Global Community White Paper provides a roadmap to extend commonality to a broader set of software.
  - The HEP Software Foundation (HSF) who shepherded this process has little funds of its own.
  - It is a grass roots organization of volunteers.
  - NSF funded travel but not the effort
  - In order to continue, it must attract funding partners.
  - DOE and NSF met at Catholic University in Dec. 2018 to work out **focus areas for each agency.**
  - This seems like an ideal role for the Center for Computing Excellence (CCE).
    - Encouragement from this funding source should encourage broader use of common solutions.

[1] <https://arxiv.org/pdf/1712.06982.pdf>



# R&D Focus for the Labs

- The labs have been participating in R&D in several areas and lead the field in the areas of:
  - **Simulation**; most notably in the pushing forward of first, the multi-threaded Geant toolkit and later a vectorized version.
  - **Frameworks**; DOE labs are developers of 3 of the 4 HEP frameworks.
  - **Physics Generators**; actively developing Genie, Sherpa and a generic phase space integration package for other generators to use.
  - **Workload Management**; both Panda and WMAgent are developed at the labs.
  - **Networking and Federated Storage**; we have a long collaboration with ESNNet.
- At the Dec. workshop at Catholic University, the university PIs represented there agreed that the above areas should remain a focus for the labs.
- There are a number of areas where there is common interest, notably reconstruction and there is plenty to do.
- **It will take lab manpower to develop and deploy the results of this research.**

# Strategy of Common Solutions

We need to continue to push for common solutions across all frontiers.

- Envisioning a goal of merging solutions across Energy Frontier and Intensity Frontier
- Fermilab Scientific Computing Division is a great place to foster this vision.
  - Some facilities are already shared like configuration services (infrastructure as a service)
  - Software can be leveraged by multiple experiments (software as a service)
- Currently have separate instances or solutions for:
  - Compute
  - Storage
  - Software
  - Data Management
  - Software repository
  - ...
  - LPC CAF
  - EOS
  - CMSSW
  - PhEdEx
  - GitHub
  - ...
  - FermiGrid
  - NAS - dCache
  - art - LArsoft
  - SAM ← Plans for fixing in 2018
  - Redmine
  - ...

# Summary and Plans

- Continue the successful model of providing services to a broad range of scientific endeavors with a lean staff.
- Continue the adoption of common solutions across frontiers by negotiating with our partners.
- Unprecedented scale of computing will be needed in the next decade.
  - The ECP will have to be a part of the strategy for meeting the computing challenge.
  - The SciDAC and other R&D efforts should help us learn how to get to the exascale.
  - HEPCloud is instrumental in bridging our user base into this exceptional resource.
- Data movement and storage is an even larger challenge.
  - The scale of investment required to meet this need is international.
  - Common solutions as outlined in the CWP must be found, and Fermilab must participate.