Fermilab | U.S. DEPARTMENT OF ENERGY | Office of Science

Scientific Computing status and vision

(with focus on neutrino program support)

Panagiotis Spentzouris & Wesley Ketchum

Fermilab PAC

January 20, 2016

The charge

We ask the committee to comment on the SCD status, plans and vision, and their consistency with programmatic priorities. In particular, are the proposed activities in support of the neutrino program likely to be adequate for the success of the experiments within the program

The organization

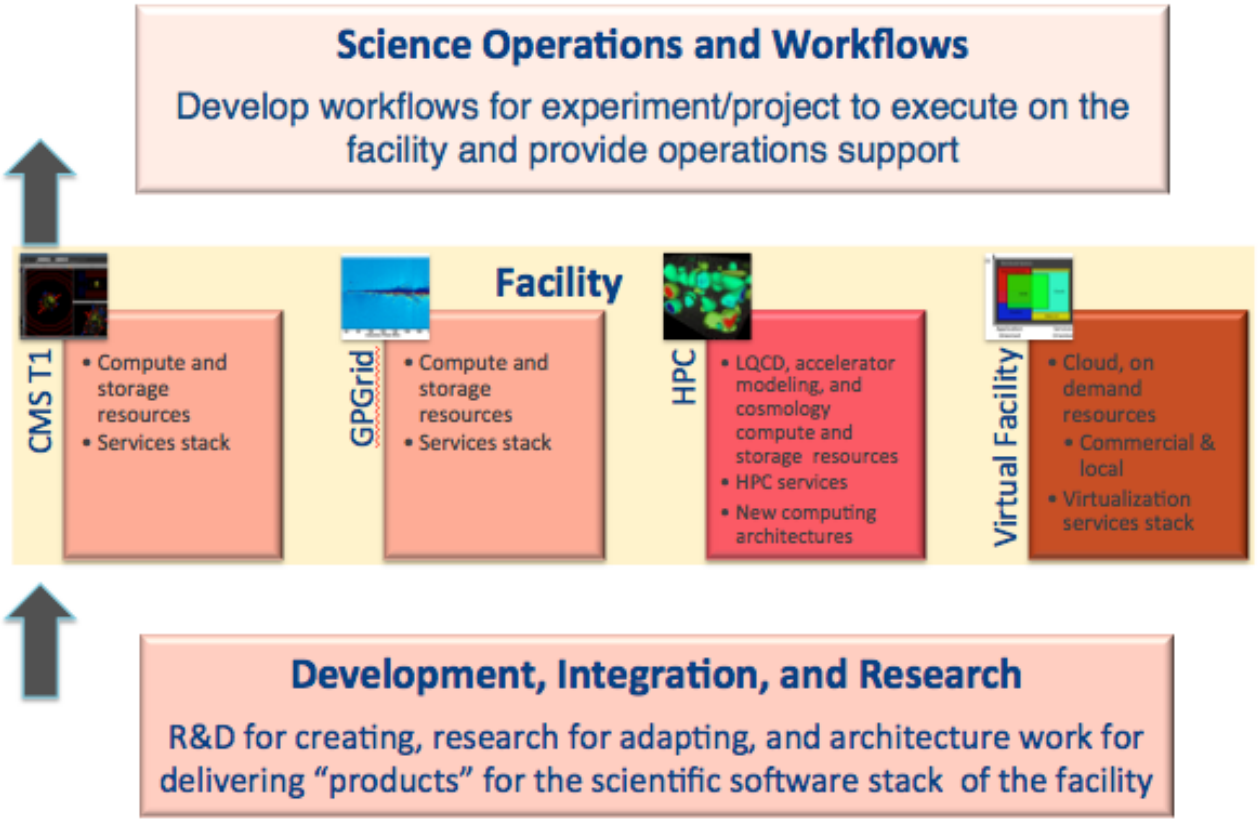

Staff ~equally distributed in the three activity areas

Headcount of 143, including 27 Scientists, 10 Application Physicists, 52 PhDs (physics and computer science) in technical jobs

The challenges (at least some of them…)

- Scientific results from all programs depend critically on complex software and computing infrastructure
- Infrastructure and application development and their *support* requires significant investment
- Most projects/experiments don't include programmatic funding for computing
- Long term *support* necessary but no clearly defined funding model
- Application development and computing infrastructure *support* requires specialized expertise
  - Especially as we move to new techniques and technologies

The Strategy

- **Develop and maintain core expertise, tools and infrastructure,** aiming to support the entire lifecycle of scientific programs
  - Focus on **areas of general applicability** (common to all/most programs) **with long term support requirements**
    - **Continuity:** Well matched to lab environment
    - **Effectiveness through Collaboration:** Work in partnership with individual programs/experiments
    - **Applying Research Opportunities:** Enabling and taking advantage of innovation
  - Participate in **collaborative projects to develop scientific computational infrastructure** (both within and outside HEP)
- **Incorporate expertise and best-of-class tools through partnerships** with individual projects and make them available to the whole program
  - Benefits both new and mature (diminishing resources) experiments

The intended benefits

- Programmatically, gain in cost effectiveness and efficiency (leveraging, sharing)
  - Application deployment, operations of existing capabilities
  - R&D for evolving/new capabilities
- For the user community, provides a *de facto* support model for the software stack
  - availability, maintenance, consultation, porting to new platforms…
- For new projects or upgrades, cost effectives
  - benefits of leveraging R&D between programs which might have not been able to afford individually
- Foster community involvement, by shared ownership
  - provide (elements) of necessary training on computing for new generation of HEP scientists

The Status: Scientific Computing Portfolio Drivers (1/2)

- Support the **CMS science program,** by
  - hosting and operating the **CMS Tier-1** facility and the LHC Physics Center (**LPC**),
  - developing and supporting the core software framework and key computing tools.
- Support the diverse **neutrino and muon programs,** in all aspects of their computing needs, by providing
  - **Facility** with **Tier-0** performance and capabilities,
  - common tools, services, and operations to enable science.
- Support selected **Cosmic Frontier experiments** per P5 and Fermilab priorities
  - Focus on DES operations, software frameworks and workflows

Portfolio Drivers (2/2)

- Provide **Real-time systems** solutions for the entire program
  - Emphasis on neutrino and muon program DAQ and test-beam
- Support the **LQCD** program by hosting a High Performance Computing (HPC) center
- Study and optimize **current and future FNAL accelerators**
  - Utilizing HPC modeling capabilities
- **Perform R&D for new tools and services**: evolution of computing architectures and technologies calls for major re-engineering to maintain capabilities
  - multicore, co-processors, reduced memory/core footprint
  - emergence of clouds as a resource
  - **Focus on selected high impact/relevance areas: facility evolution, software frameworks, workflow management, Geant4, accelerator modeling**

Planning and Resource Allocation: Scientific Computing Project Portfolio Management Process

- Programmatic resource allocation is based on lab-wide scientific needs (hardware and effort for services)
  - Process is Science driven – ask experiments to present annually their goals for the coming two years
  - Utilize external committee for scrutiny and recommendations
- Continue to monitor and communicate through frequent meetings
  - adjusting as priorities/needs change
- Many other points of contact
  - Computing liaisons provide bi-directional status and information
  - Stakeholder meetings for major computing projects

➢ **We support operations of 24 service areas across 32 scientific collaborations and projects (23 experiments)**

High Impact Common Tool Solutions

- **art** is a software framework, for HEP experiments
  - Allows **shared development and support among experiments**
  - Used by Mu2e, g-2, NOvA, DS50, LArSoft
- **LArSoft** is a common simulation, reconstruction and analysis toolkit for LArTPC experiments, utilizing art
  - managed by Fermilab, **contributions from all experiments**
- **FIFE:** provide common computing services and interfaces needed to turn a physics task into results, enabling experiments to seamlessly utilize onsite and offsite resources.
  - Enables use of grid and cloud resources

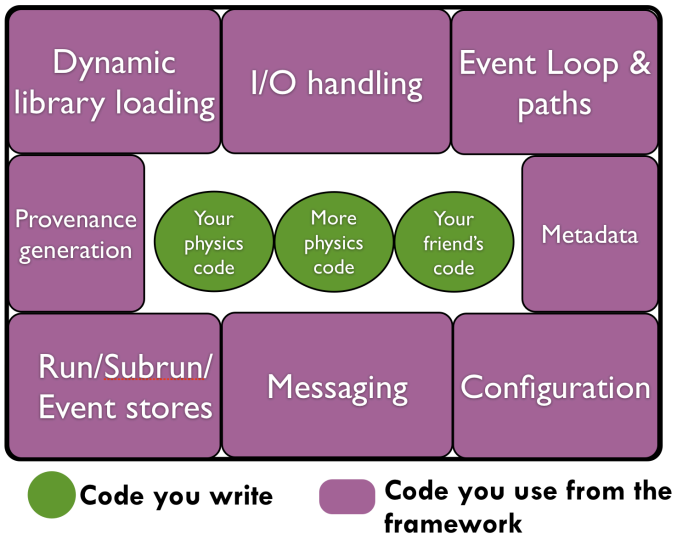

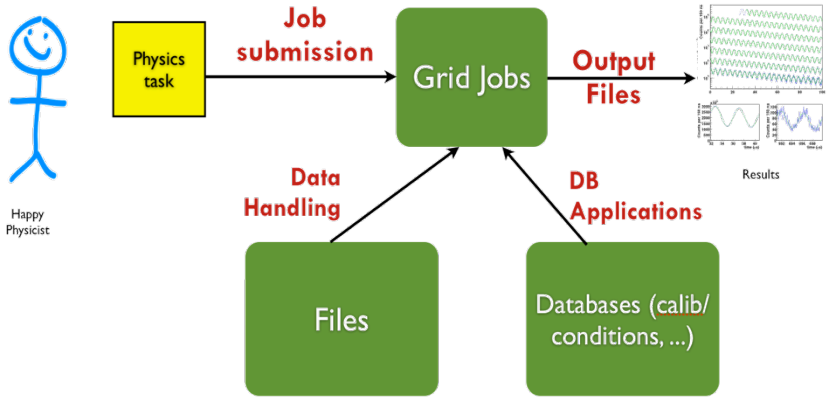

FIFE: FabrIc for Frontier Experiments

Fermilab

Utilization of FIFE

Experiments	jobsub grid submission	fifemon/gratia	SAM	FTS	IFDH	Production Operations	
ArgoNeuT	Testing	Testing					Intensity Frontier
CHIPS	In Production	In Production				Planning	
DUNE (35T)	In Production	In Production	In Production	In Production	In Production		
DUNE	In Production	In Production	In Production			In Production	
LArIAT	In Production	In Production	In Production	In Production	In Production	Planning	
MARS	In Production	In Production	In Production		In Production		
MicroBooNE	In Production	In Production	In Production	In Production	In Production	Planning	
MINERvA	In Production	In Production	In Production	In Production	In Production	In Production	
MiniBooNE						No intention	
MINOS+	In Production	In Production	In Production	In Production	In Production	In Production	
Mu2e	In Production	In Production	In Production	In Production	In Production	Planning	
Muon g-2	In Production	In Production	In Production	In Production	In Production	Planning	
NOvA	In Production	In Production	In Production	In Production	In Production	In Production	
NumiX	Testing	Testing	Planning				
SBND	In Production	In Production		In Production	In Production		
SeaQuest	In Production	In Production	In Progress				
CDMS	Testing	Testing			In Production		Cosmic Frontier
COUPP	In Production	In Production	In Production		In Production		
DarkSide	In Production	In Production	In Production	In Production	In Production		
DES	Testing	Testing			In Production		
DES-GW	In Production	In Production			In Production		
Holometer	Upgrading	Upgrading	In Production				
LSST	In Production	In Production					
SDSS							

High Impact Common Tools Solutions

- **artdaq** is a real-time software system for data acquisition, utilizing art (for monitoring, filtering,…)
  - Conceptualizes common DAQ tasks
  - Allows experiment to focus on design/ configuration of system
  - Used by Mu2e, Darkside50, ICARUS test system, uBooNE cosmic ray tagger, SBND
    - Note that we provide support for NOvA (FNAL pre-artdaq) and uBOONE DAQ
- **Geant4,** collaboration member: provide validation, development, and expertise on physics configuration and user application development
  - Relevant to the whole program
- **GENIE** (neutrino generator), collaboration member: modernize infrastructure for incorporation of new data and physics validation, provide consultation

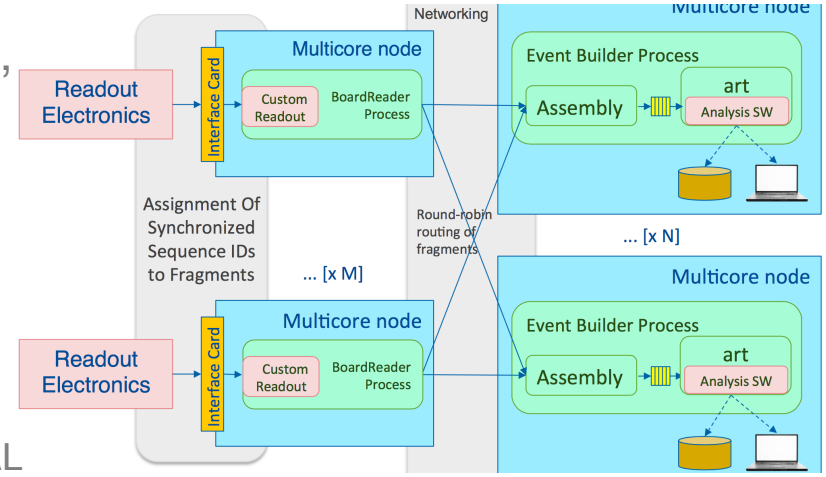

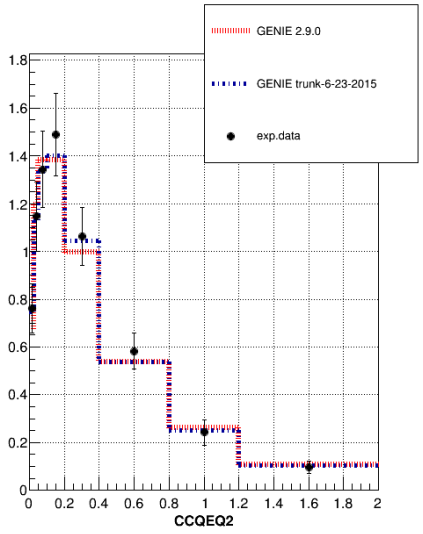

Scientific Computing services and operations in high demand

- Data ingress at record rates
  - LHC restarted taking data at 13 TeV
  - MicroBooNE started data taking at high volumes
- CPU resources are in high demand
  - Mu2e used Fermilab and opportunistic resources on the Open Science Grid (OSG) to produce simulations for the CD-3c review
- Delivering the software, services and operations for storing, distributing, and processing the data
  - Workflow management and distributed data tools, operations both for the facility and experiment workflows
    - From running workflows (MINOS, MINERvA, NOvA, DUNE simulation...) to monitoring and troubleshooting jobs (for all), to providing tools and expertise to experiments for utilization of remote resources (OSG)

CPU utilization on Fermilab resources, CY2015 (Reference)

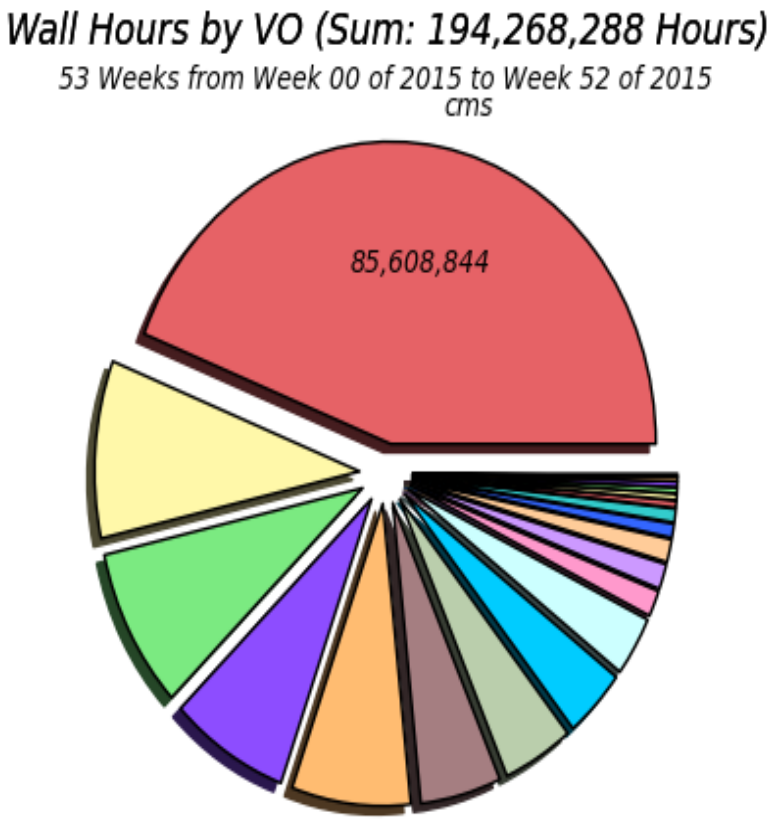

CPU utilization on all OSG resources, CY2015 (CMS excluded) (Reference)

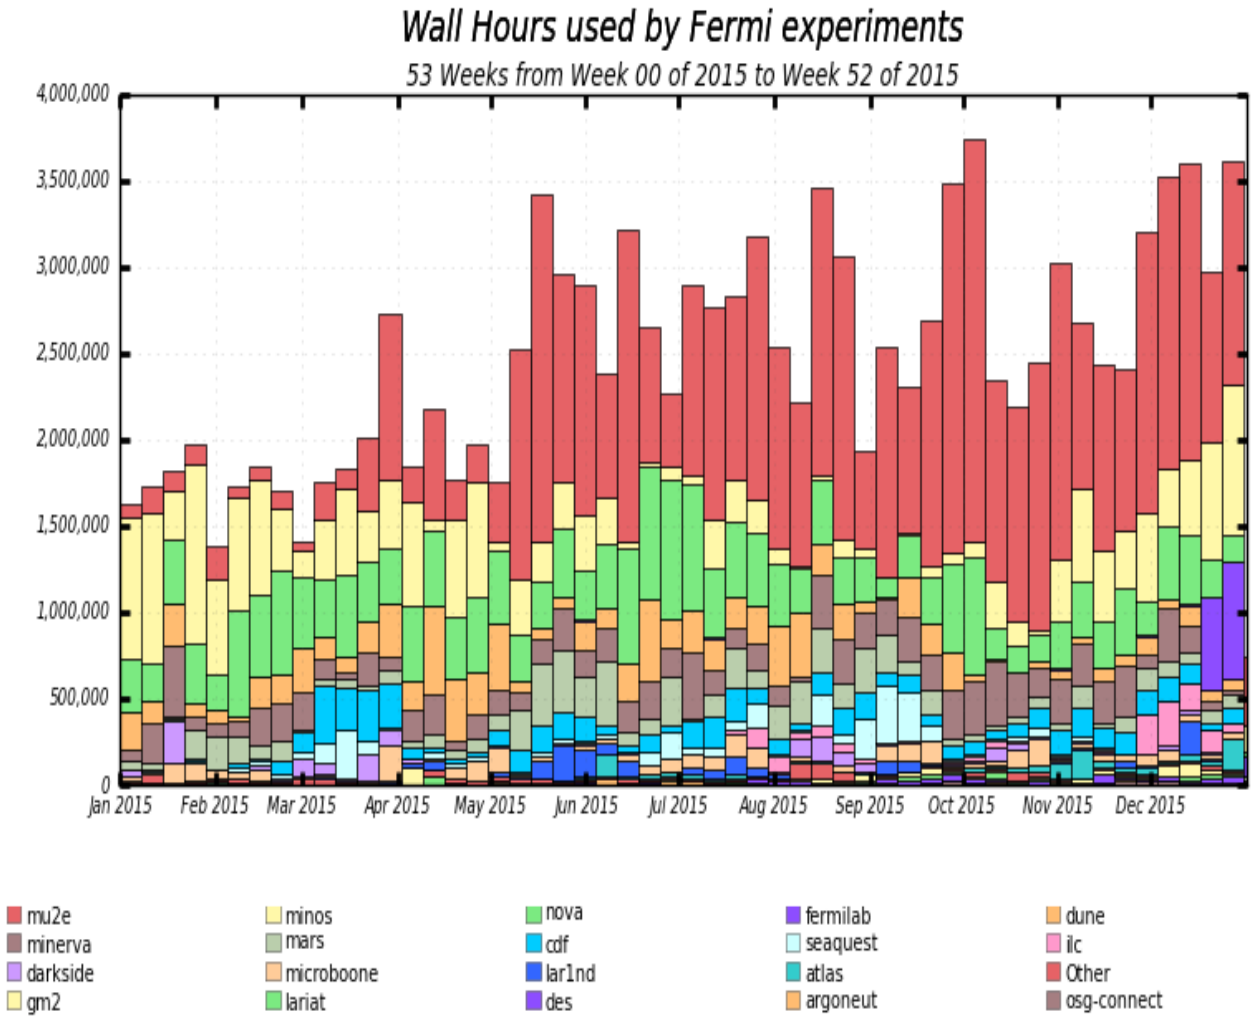

Storage

- Disk and tape utilization in the Fermilab Active Archive Facility
  - "active": catalogs, tools to access and distribute
    - CMS excluded
- NOvA dataset already ~ size of CDF Run II !

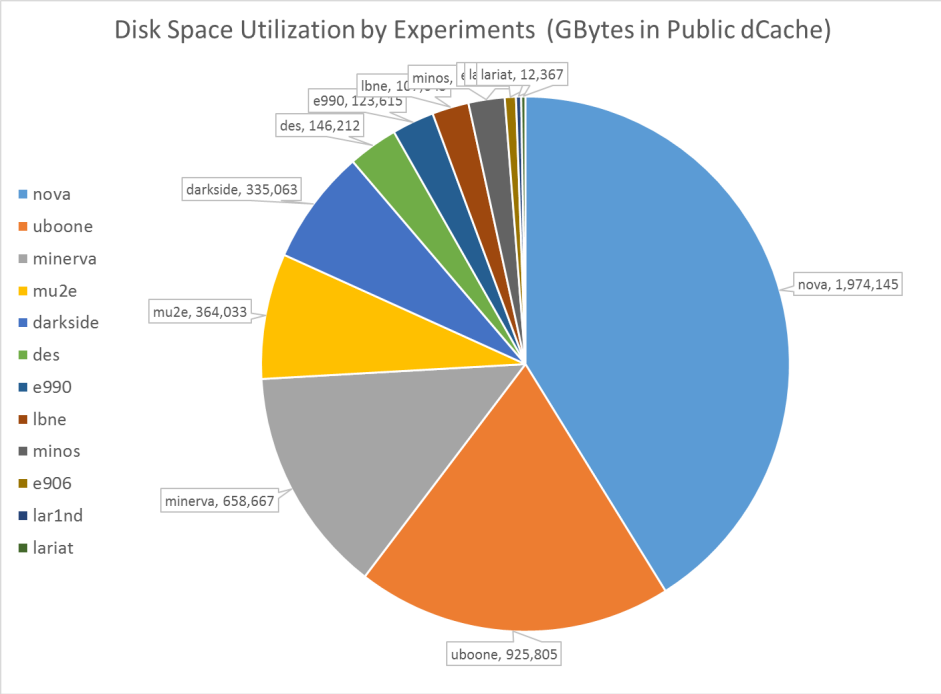

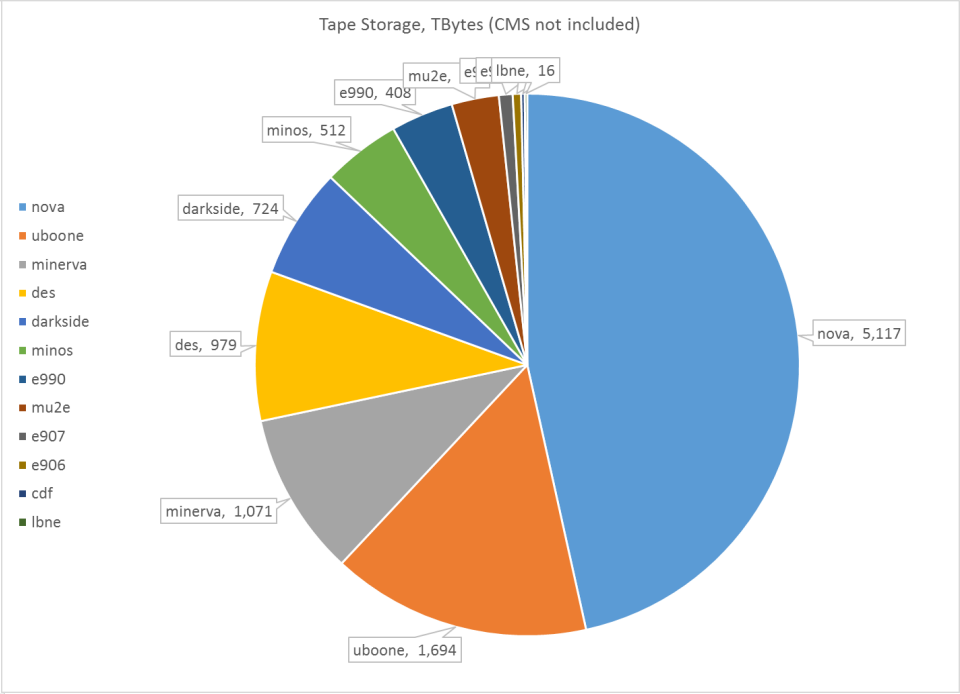

Model works well: for example, NOvA, where SCD contributes to all aspects of software and computing

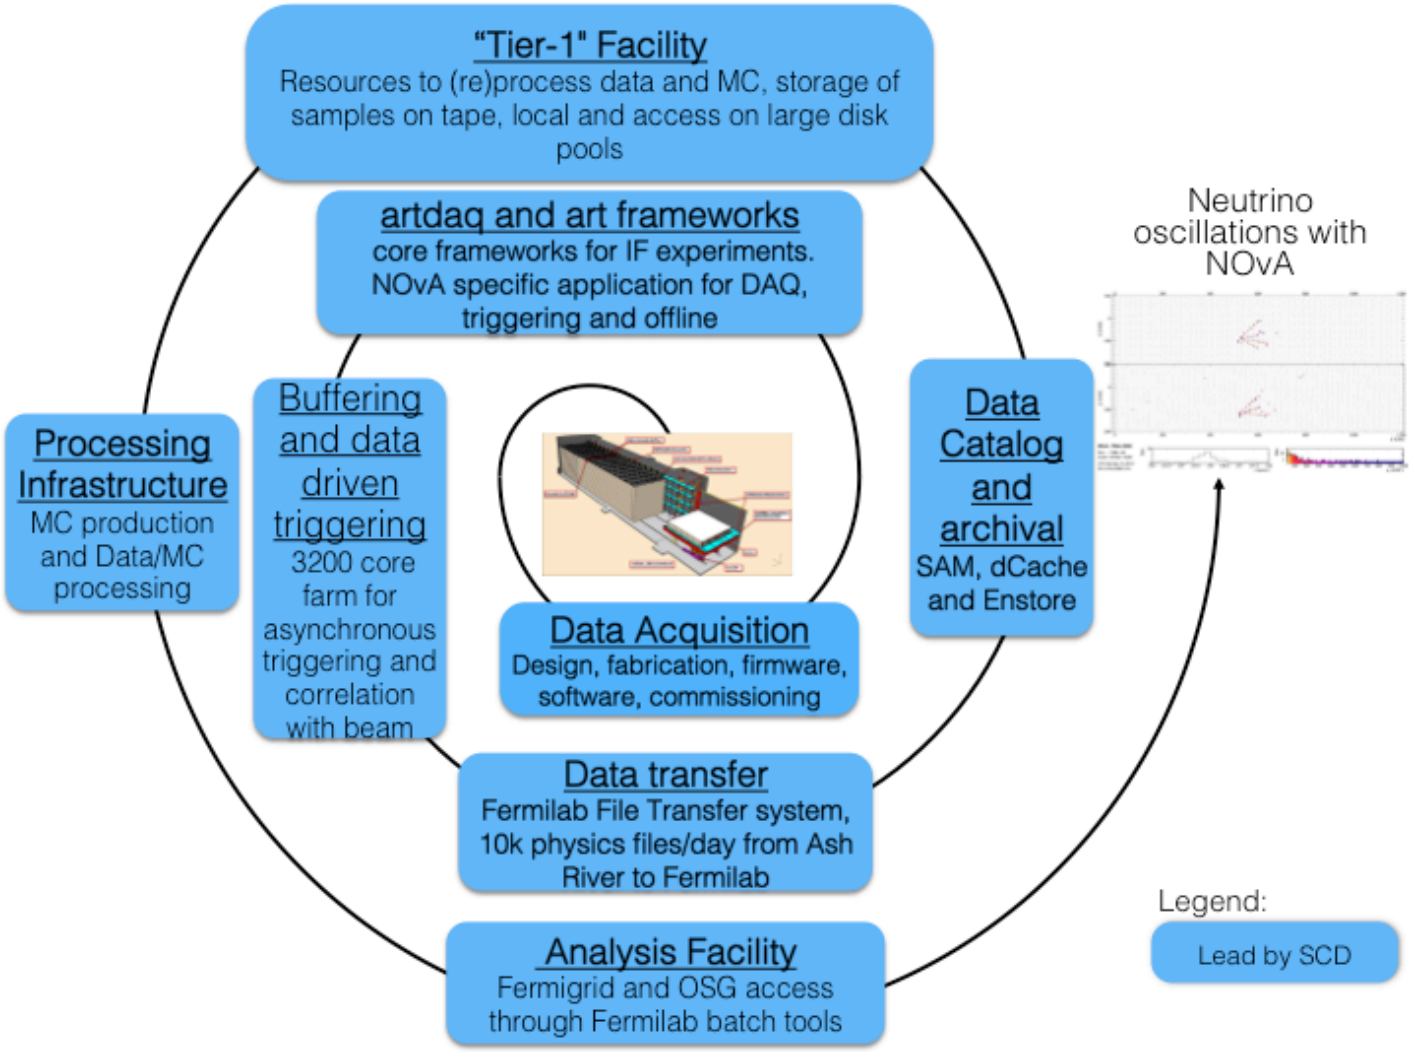

High priority to provide support for LArTPC based Neutrino Program:

SBN Far Detector

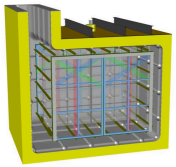

SBN Near Detector

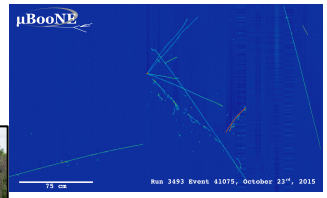

MicroBooNE

2015 | 2016 | 2017 | 2018 | 2019 | 2020 and beyond

DUNE 35-ton

Proto-DUNE(s)

DUNE

Computing Challenges include:

- High data acquisition rates and large data volume.
- Detector resolution demands powerful and robust reconstruction tools.
- Sophisticated simulations for particle interactions.
- Computing resources for both small and large collaborations
- Must address both immediate and long-timescale needs smoothly

Neutrino program support

- Neutrino experiments use, are supported in, or are currently adopting “offerings” in the following Service Areas

Service / Experiment	artdaq	art Framework	Software Processing (larsoft)	Tape Storage and Disk Caches	SAM Data Management	Genie Geant4	Beam Information Databases	OPOS	Scientific Servers	Control Room System Management	Code Repository Hosting	FIFE	OSG Enabled
DUNE (35T)	In Progress	In Production	In Production	In Production	In Progress	In Progress	In Progress	In Production	In Production	In Progress	In Production	In Production	In Progress
DUNE	Planning	In Production	In Production	In Production	Planning	In Progress	In Progress	In Production	In Production	Planning	In Production	In Production	In Progress
protoDUNE	Planning	In Production	In Progress	In Production	Planning	In Progress	In Progress	In Production	In Production	Planning	In Production	In Production	In Progress
LArIAT	In Production	In Production	In Production	In Production	In Production	Planning		Planning	In Production	In Production	In Production	In Production	In Progress
MicroBooNE		In Production	In Production	In Production	In Production	In Production	In Production	Planning	In Production	In Production	In Production	In Production	In Progress
MINERvA	No intention	Custom		In Production	In Production	In Production	In Production	In Production	In Production	In Production	In Production	In Production	Planning
MINOS+		Custom		In Production	In Production	In Production	In Production	In Production	In Production	In Production	In Production	In Production	In Production
NOvA		In Production		In Production	In Production	In Production	In Production	In Production	In Production	In Production	In Production	In Production	In Production
SBND	Planning	In Progress	In Progress	In Production	In Production				In Production		In Production	In Production	Planning
ArgoNeuT		In Production	In Production	In Production	No intention				In Production		In Production	Testing	

Includes:

- Experiment specific work such as running production operations (OPOS) for Minerva, MINOS, DUNE simulations, NOvA, upcoming for MicroBoone
- Cross-experiment common services e.g. support for Tape Storage and Disk Caching, use of distributed resources.

Neutrino Program support

Projects for further development of our software stack and new services. These are directly driven by experiment/stakeholder needs & computing/software evolution.

- Key ones include DAQ (artDAQ), access to commercial clouds and HPC systems (HEPCloud), Frameworks (art), physics reconstruction and analysis toolkits (LArSoft, ROOT), simulation (Geant4, GENIE, accelerator modeling)

R&D for the future:

- Examples include use of Big Data technologies to reduce time to analysis results - NOvA evaluation; multi-threading frameworks and infrastructure - art-HPC; discussions of Deep Learning (advanced neural networks) - with DUNE s&c leads.

Moving Forward

- LArTPC offline software infrastructure
  - Ensure adequacy for experiment **requirements**
    - **workshop** October 2015, draft report available, **implementation review** mid 2016
  - Important needs: **automated reconstruction** and **assisted reconstruction** capabilities; readiness for ProtoDUNEs (**dual phase**); **interfaces** with external packages; **algorithms**
    - Increased **LArSoft resources**, created **reconstruction group** with SCD experts from all programs tasked to provide expertise, consultation, and assistance
  - Evolution challenge: **adiabatic vs step** change (especially relevant for DUNE).
    - Incorporating new techniques/technologies essential for future efforts
    - Choice driven by experiment needs.
      - For SCD, to be able to provide support (if it is expected), essential to participate in architecture design and implementation

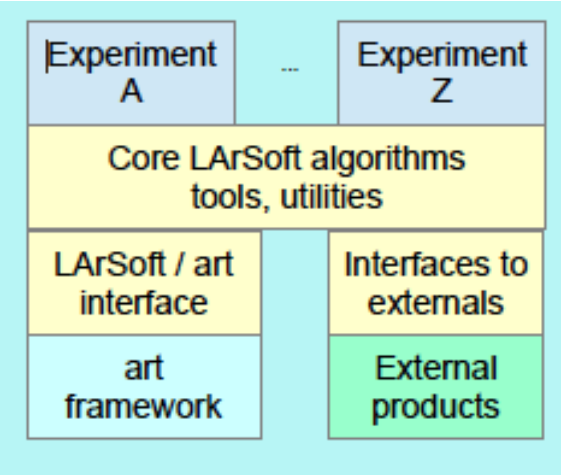

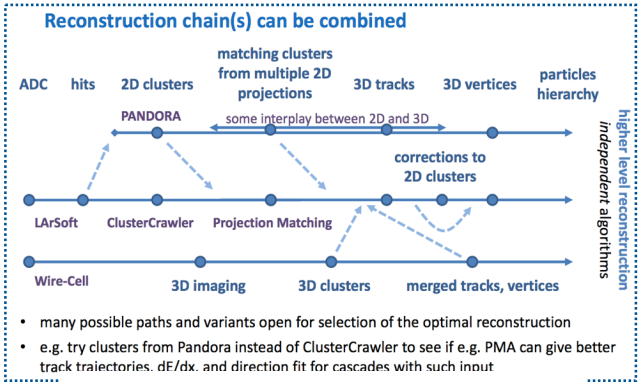

Courtesy Robert Sulej (DUNE)

Moving Forward

- Essential that SBN experiments **share tools & infrastructure** (more than efficiency, it simplifies combined analysis)
  - uBooNE, SBND on-board, engaging ICARUS (including presence in Europe)
- Providing support to the DUNE program (35ton, ProtoDUNEs, DUNE)
  - 35ton utilizes artdaq (a success story) and LArSoft
  - Recently established artdaq as a “community” project (like LArSoft) to encourage collaboration with experiment members
    - Any collaborator can have “ownership” of components, as long as they are contributed to the toolkit, thus SCD can develop “know-how” and provide support
  - Will respond to ProtoDUNE(s) needs, working with the collaboration.
    - given constraints will have to adjust (re-prioritize) efforts internally

Moving Forward

- The DUNE program (including prototypes), being an international effort, requires an appropriate *framework* for the software and computing (S&C) program
  - SCD has started discussions with international partners (CERN neutrino platform) and work with other international collaborators, will contribute to establishing this S&C *framework* and will contribute services and operations in its context
  - An important issue for the short term is the computing model (driven by ProtoDUNE timeline)
    - Establishing roles and responsibilities for various facilities, tools
    - SCD is planning to provide Tier0 facility services for DUNE,
    - will work with the collaboration to implement any missing functionality in our "**Neutrino Grid Tools**" (FIFE, glideinwms, ifdh, jobsub, GPGrid)

Moving Forward

- As we move toward the era of HL-LHC, the demands for computing (~60 times more computing, exabytes of data), deploying our own resources for peak computing becomes unsustainable.
- **HEPCloud:** Delivering a new paradigm for HEP computing facilities through a single, management portal to a dynamic, heterogeneous set of computing and storage resources. Provide cost effective and efficient "elastic" deployment of resources.

Project aiming to integrate "rental" resources into the current Fermilab facility in a manner transparent to the user.

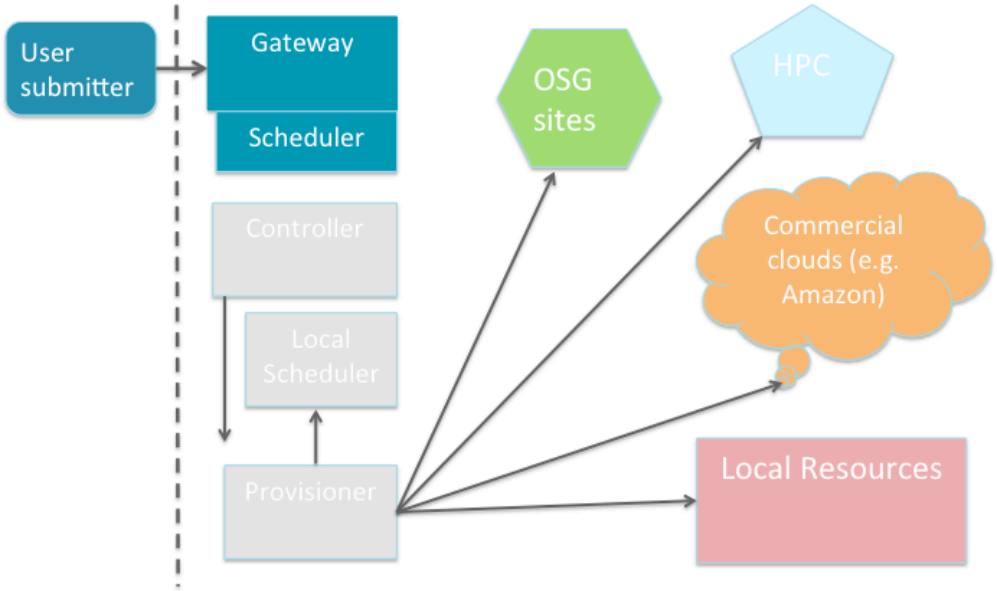

Instead of a summary, the vision…

- Provide scientific software and facility solutions and support their applications, within the context of international HEP and the international HEP computing ecosystem, to enable Fermilab's scientific program.
  - Thus, "provide" should be understood as develop, contribute, adopt, participate, …, with emphasis on maintaining the "know-how" to develop and support applications

So, enable science through working with the community and supporting users (and improve the mood of spokespeople and lab directors…)

From Ken Bloom, DPF2015

Backups

- More detailed material for various SCD tools and approaches

art: the event-processing software framework

- *art* is an event-processing framework used for online data monitoring, data calibration, reconstruction, simulation, and analysis
  - Provides a *common software infrastructure* for…
    - defining experimental data types
    - defining and applying a configurable algorithm workflow to data
    - providing ancillary detector information (*e.g.* geometry, calibration)
    - recording data provenance
  - Used by Mu2e, g-2, NOvA, and currently is the LArSoft engine
- We maintain infrastructure, development toolkit, documentation → experiments focus on writing algorithms
  - *art Workbook*: a user-focused guide to getting started
  - "Stakeholders" meetings to ensure communication with user community

LArSoft: common tools for LArTPC experiments

- LArSoft is a general software toolkit for LArTPC simulation, reconstruction, and analysis
  - Utilizing the *art* event processing framework
  - Incorporates common algorithms and data formats, and experiment-specific geometry, signal calibration, and configurations
  - See table for utilization
- SCD works to ensure quality design and performance for successful cross-experiment collaboration
  - Recently held “requirements workshop” to establish and document experiment needs
    - Participation from all SBN experiments, DUNE and prototypes

artdaq: common tools for DAQ systems

- artdaq is a real-time software system for data acquisition
  - Conceptualizes common DAQ tasks: data packet transit, event-building and writing, monitoring data flow, …
  - Allows experiment to focus on design/configuration of system, data extraction from hardware, and validating collected data
    - Online monitoring/processing of events using *art* framework
- SCD works with experiments to design and commision DAQ systems, and extend functionality of *artdaq* to meet needs
  - Heavily involved in *artdaq* system for DUNE 35-ton prototype
  - Designing and building *artdaq* systems for SBN program
    - SBND, MicroBooNE cosmic-ray-tagger system, ICARUS test system

Simulation Common Tools: Geant4

- High-fidelity simulation tools essential for translating experimental data to physics results
- Fermilab is a member of the Geant4 Collaboration
  - Leads computing performance working group
  - Develops and operates **physics validation infrastructure**
  - Works with user community to **support major applications**.
    - Very successful CMS involvement, developing partnerships with muon and neutrino programs (physics lists, new physics, application development)
- Established R&D program to **evolve G4 to take advantage of HPC and modern architectures**
  - Partnered with DOE/ASCR institutes to **improve G4 performance**.
  - Partnered with CERN, UNESP (Brazil) to re-engineer G4 to **run on modern computing architectures** – Geant Vector Prototype (GeantV)

GENIE

- GENIE (Generates Events for Neutrino Interaction Experiments) is a Monte Carlo event generator package used by nearly every accelerator-based neutrino experiment.
- High quality generators are key in a discovery experiment like DUNE.
- We are transformomg GENIE operations and science to operate at the scale of DUNE and other neutrino-based discovery experiments.
  - Run community meetings to train experimenters.
  - redesigned the physics performance framework and built an automated validation framework to enable faster physics development and faster release cycle.
- The current goal at the laboratory is to leverage the validation automation to build machinery for producing new global physics tunes that will enable physics measurements with the smallest possible interaction model uncertainties.

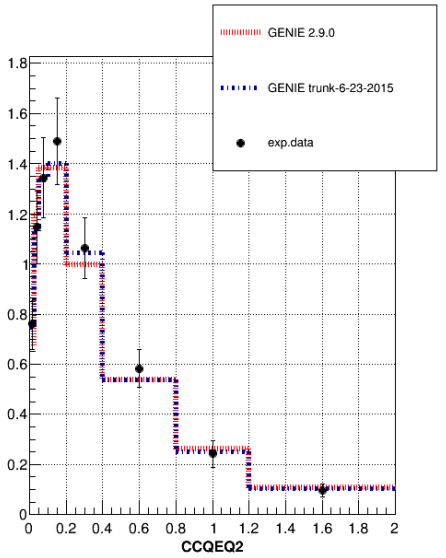

Validation plot from the new framework comparing MINERvA cross section data (PRL 111, 022501 and 022502) to two recent versions of GENIE. The validation framework is capable of fully-utilizing reported error covariance matrices and can compare multiple publications against variations in the same internal physics model. Furthermore, it provides a convenient API for future developers to "plug in" new data with very little effort.

Simulations: Accelerator Modeling

- Reaching the goals of PIP-II requires detailed understanding of collective effects to control/avoid beam instabilities and minimize losses
- Utilize Synergia a 3D Particle-in-Cell HPC accelerator simulation framework, developed at Fermilab under the ComPASS SciDAC project
  - Wakefield and space-charge induced instabilities in Booster
  - Space-charge in Recycler slip-stacking & collimator design

Partner with AD scientists to run the applications

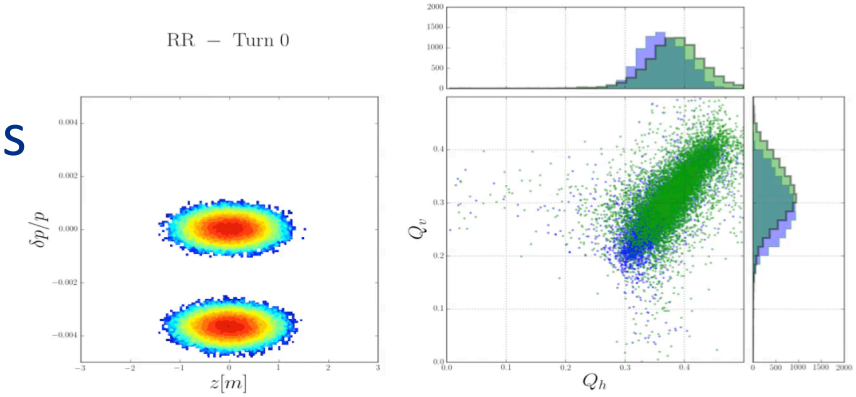

Facilities for data storage and retrieval

- Neutrino experiments aren't "small"
  - MicroBooNE has written 1.7PB total on tape so far
  - Raw, reconstructed, and simulated data needs for SBN program, ProtoDUNE, and DUNE will be multi-PB-scale too
- SCD maintains common disk and tape storage systems and develops the tools for efficiently using them
  - dCache and Enstore manage the storage and access of data on disk and tape systems, respectively
  - *File transfer service* (FTS) and *intensity frontier data handling* (ifdh) enable experiments automate movement of files
  - SAM (Sequential Access via Metadata) automates the retrieval of data based on experiment-defined properties (run number, types of file, etc.)

Facilities for data processing

- Lots of data → need for lots of processing power
  - MC production alone can take ~1E6 CPU hours for single experiments
    - Which need to be reproduced and re-reconstructed several times in life of experiment
  - Leverage all available resources to be time and cost efficient
    - Distributed computing model already becoming a need
- SCD manages and maintains local resources and develops the infrastructure for using shared resources worldwide
  - Common grid nodes (GPGrid) for neutrino experiments at FNAL
  - Leader in the Open Science Grid (OSG), providing software and services for opportunistic usage on external resources
  - HEPCloud as path to resources in the HPC and commercial world

FIFE: bridging software and facilities

- FabrIc for Frontier Experiments (FIFE): catalog of tools and services to help experiments make the most use of available resources
  - Streamlined interfaces for job submission, data access, and software distribution
  - Common utilities for database and dataset applications and collaborative tools (e.g. electronic logbooks)
  - Experiments pick and choose elements they find most useful
    - And help improve/add to the existing portfolio
- SCD leads coordination of resources and technical work
  - Helps experiments establish workflows and build integrated solutions
  - Let users focus on *what* to do with data, not *how* and *where* to do it

FIFE: bridging software and facilities

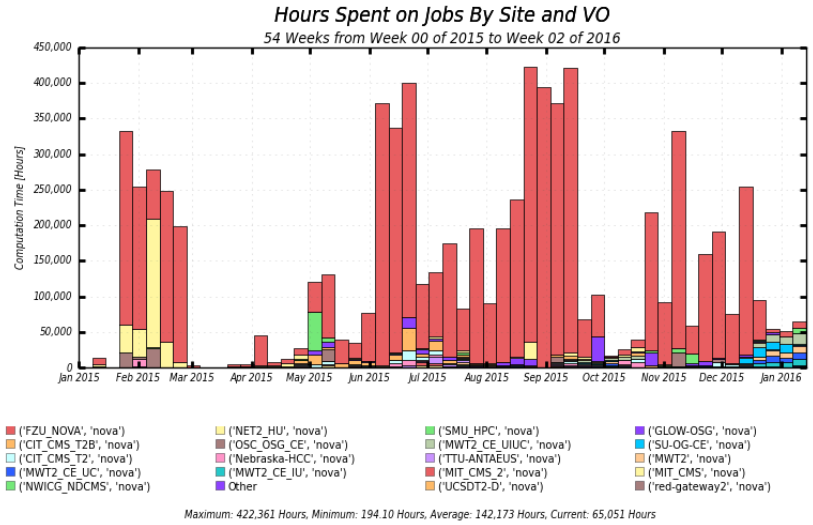

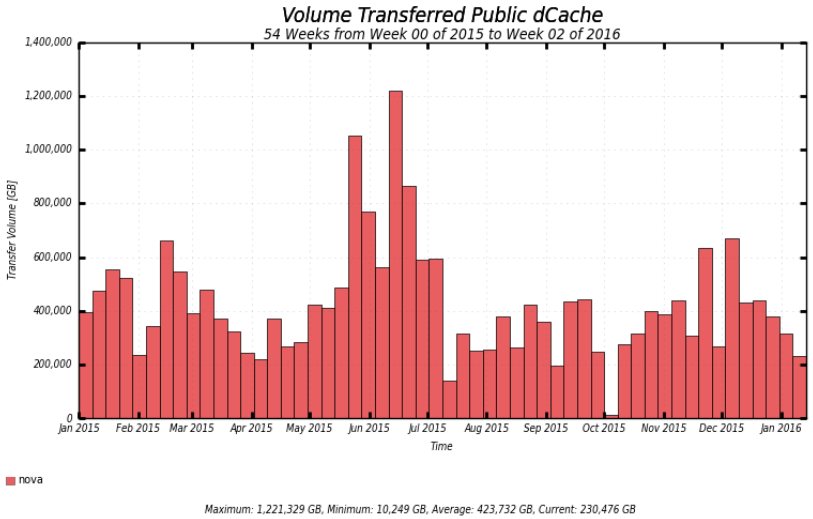

- Example: NOvA uses FIFE tools to achieve physics goals
  - *jobsub* utilities allow users to submit jobs to FermiGrid and OSG with same interface
    - Over last year: average of 142,000 CPU hours per week across > 20 sites
  - *ifdh* tools allow users to easily access data properly and efficiently from variety of storage architectures
    - Over last year: >22 PB data transferred in and out of dCache (1.2 PB in one week!)

CPU utilization on Fermilab resources, CY2015 (CMS excluded)

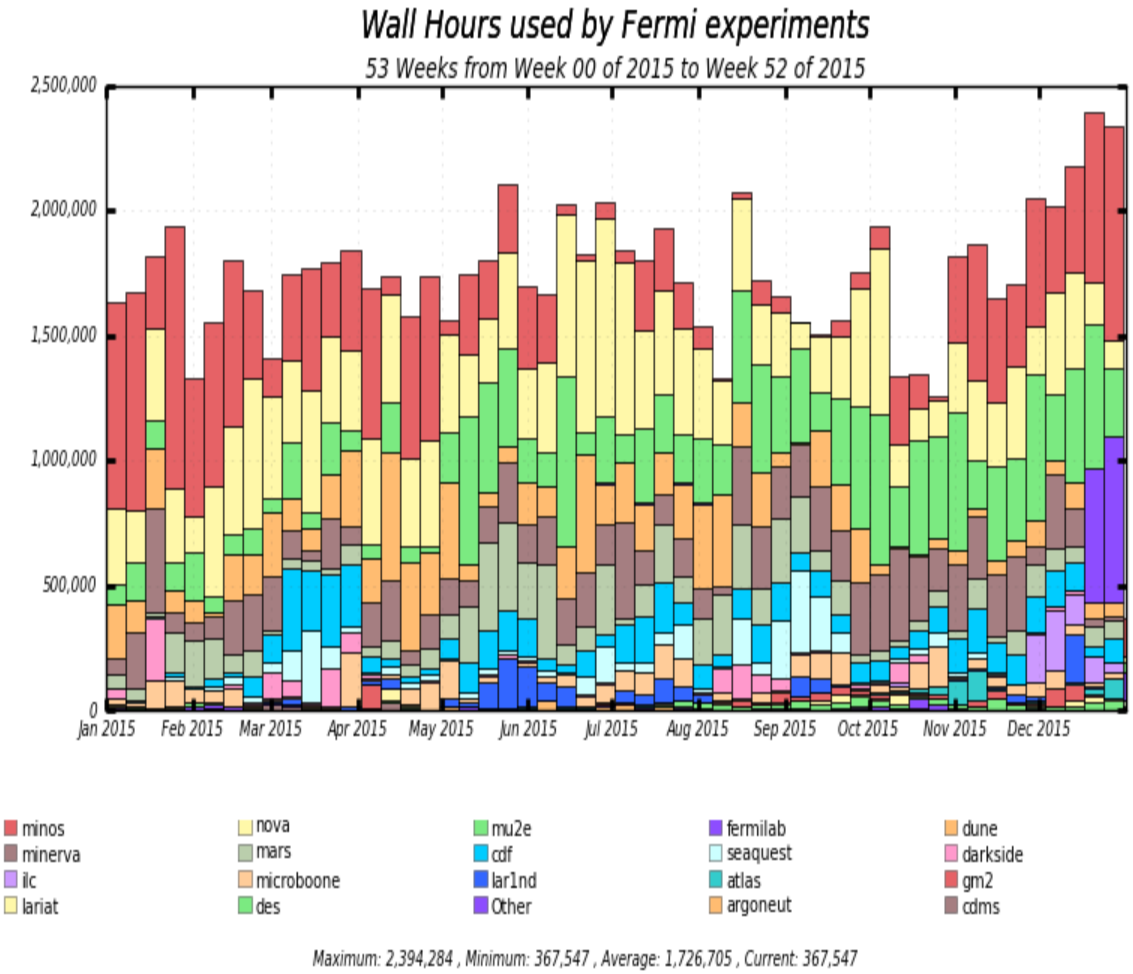

LArTPC software requirements workshop

- Goal was to capture requirements for software infrastrucute that will support the analysis needs of LArTPC experiments over the next ~decade
- ➢ Requirements document now in draft, https://cdcvs.fnal.gov/redmine/projects/lartpc-requirements/repository/revisions/master/entry/new-document/requirements.pdf
  - Examples of major areas: **i)** physics algorithm performance, **ii)** ability to use multiple physics algorithms in end-to-end analysis of data, **iii)** increased functionality of event visualizations, **iv)** enable effective use of multi-core and new computer hardware technologies, **v)** ease of use and distribution for international collaborations, **vi)** inclusion of new external software components such as event generators and hadronic simulation codes

Requirements workshop: bringing the community together

- Participation of ~40 people from all experiments, including 3 ICARUS physics software developers, in virtual (face-to-face and remote) rooms arranged to help experiments interact
  - detailing needs and ideas for LArTPC reconstruction and development
- Informal as well as formal (notes/requirements captured in many breakout sessions) articulation and discuss of longer term future (…as well as near term requirements, methods and technologies)
- Many one-on-one physics/technical conversations
  - on common topics and thrusts of interest across multiple experiments.
- Captured automated reconstruction needs
- In depth demonstration of the ICARUS QSCAN (interactive visualization and hand-directed analysis).
- Broad buy-in and agreement on principles and requirements to guide the development of the software
- Increased appreciation of benefits from sharing infrastructure and algorithms/algorithmic implementations
  - resulted in more codes being brought to LArSoft Coordination forum for inclusion in core releases!

Evolution of FNAL experimental program

2015-2020 (large, mid-size programs)

MI-LB neutrinos

- MINOS+, MINERvA, NOvA

Booster-SB neutrinos

- uBooNE, SBND, ICARUS

FNAL Recycler-muons

- g-2

LHC beams: Run 2

- CMS

2020-2025

…

Recycler-muons

- Mu2e

LHC Run 3 (phase 1 upgrade)

…

2025-…

Long Baseline Neutrino Facility (LBNF)

- **DUNE**

LHC Run4 (HL-LHC)

…