Intensity Frontier Computing at Fermilab

Stephen Wolbers, Fermilab

DPF 2013, Santa Cruz

August 15, 2013

Outline

- Intensity Frontier Physics Program at Fermilab
- Intensity Frontier Computing – Requirements
- Computing for the Intensity Frontier
- Summary

Fermilab

INTENSITY FRONTIER PHYSICS PROGRAM

Intensity Frontier Program development

- Fermilab has a rich and growing Intensity Frontier (IF) experimental program in place consisting of
  - experiments taking data,
  - experiments soon to complete detector construction and begin data taking,
  - experiments about to ramp up to full construction,
  - proposed experiments
- In addition various activities such as liquid argon R&D are part of a larger strategy to develop additional capabilities and experiments.
- Other proposals and ideas such as polarized proton beams, a neutrino source, and an experimental program associated with Project X, are all being actively investigated.

Science and Facility Roadmap

Short- and medium-term efforts at the three frontiers fit together to support long-term strategy for science and facilities.

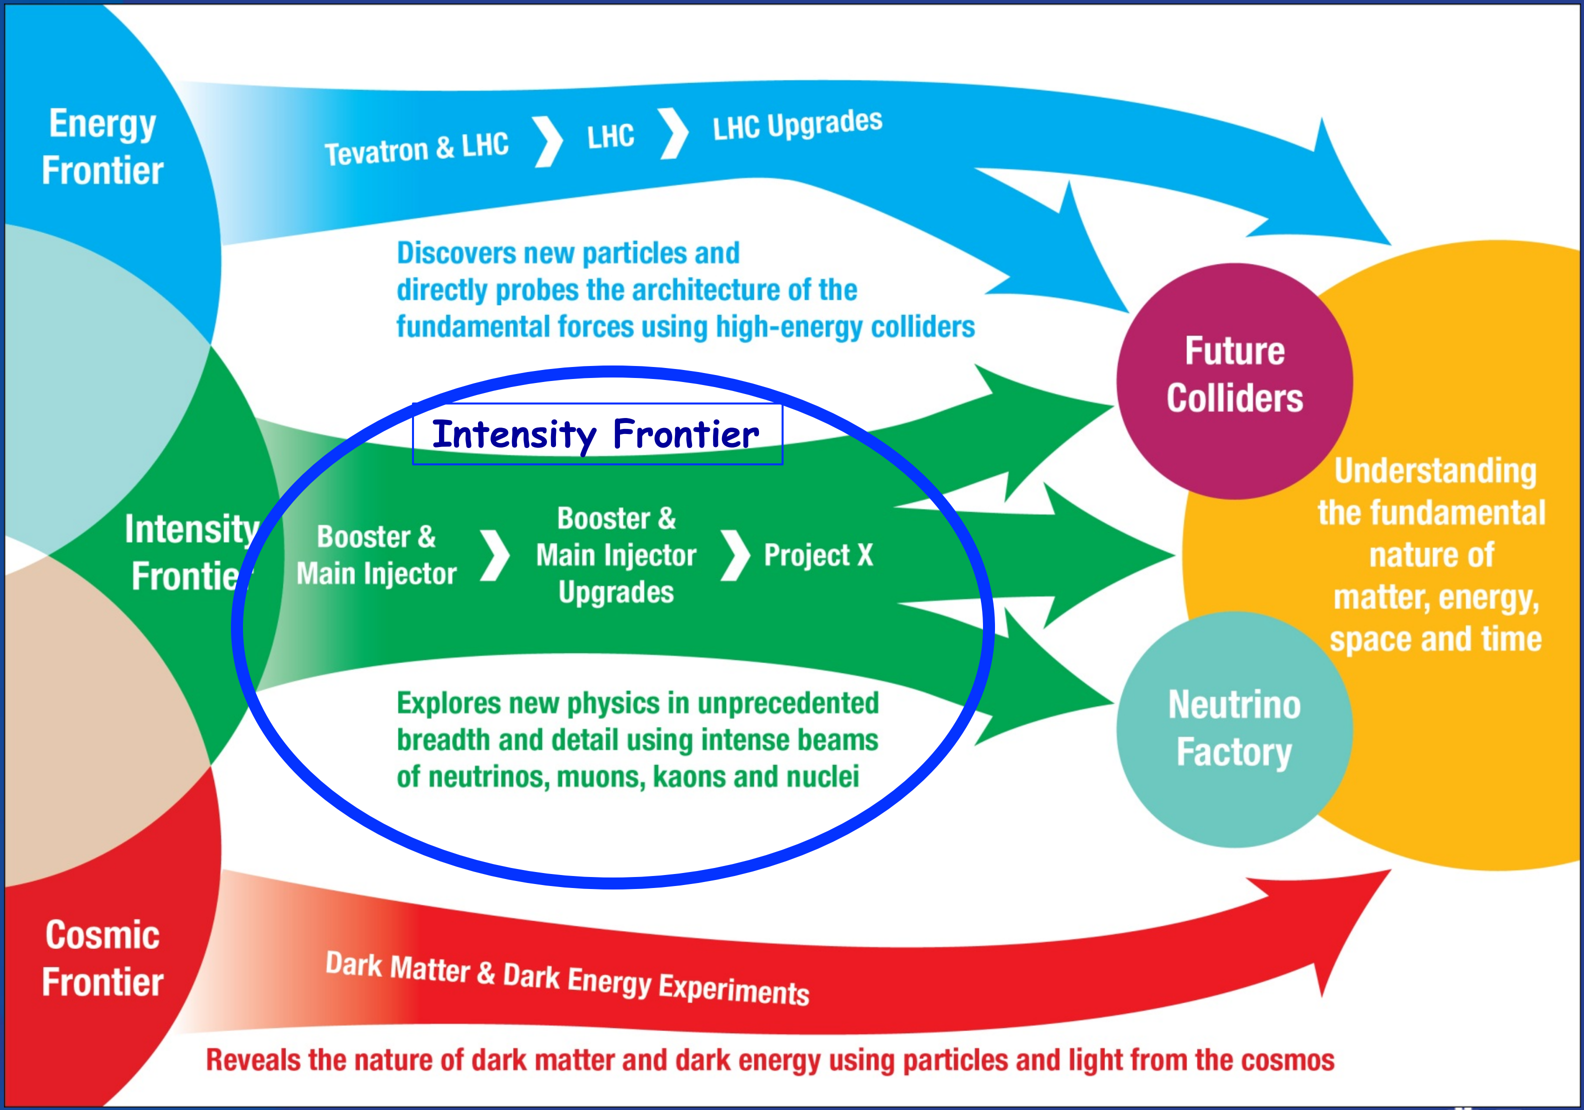

Fermilab Strategy for the Next Decade

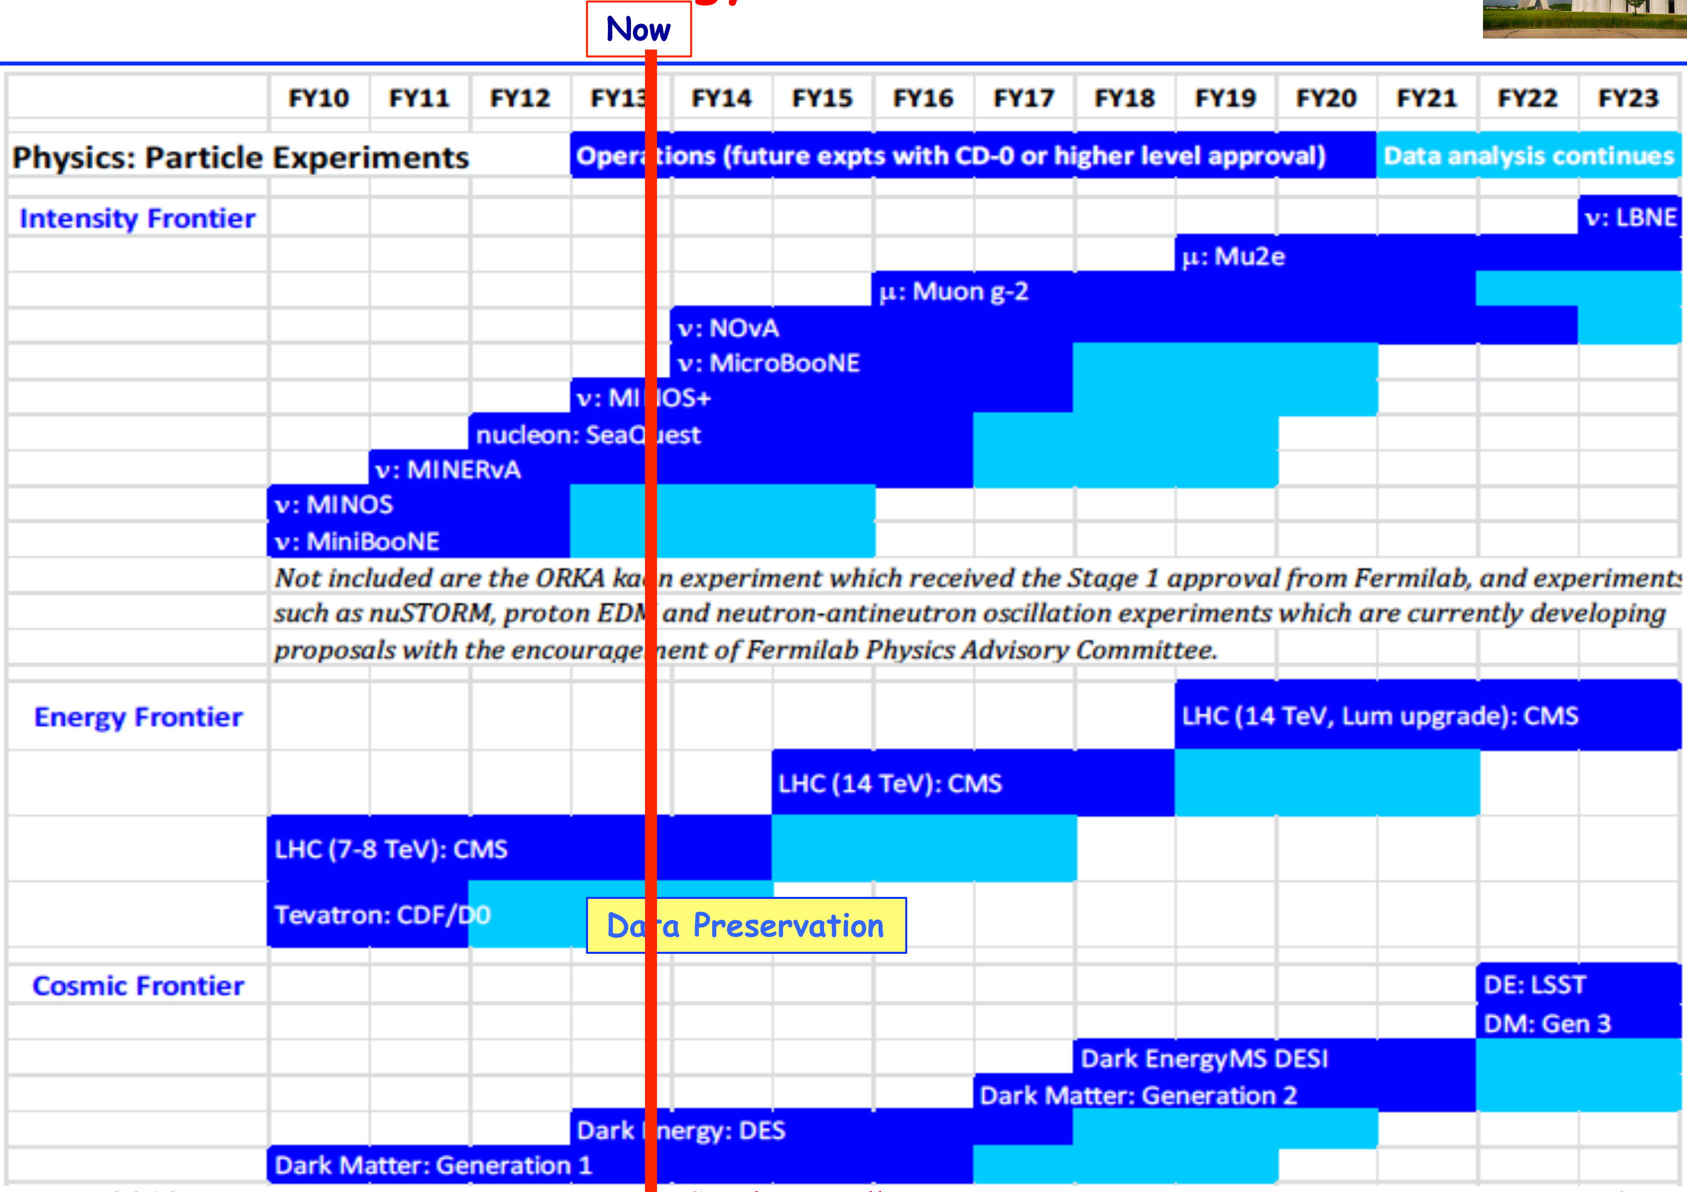

Fermilab experiments – next few years

Fermilab Program Planning
30---May---13

Fermilab Accelerator Experiments' Run Schedule

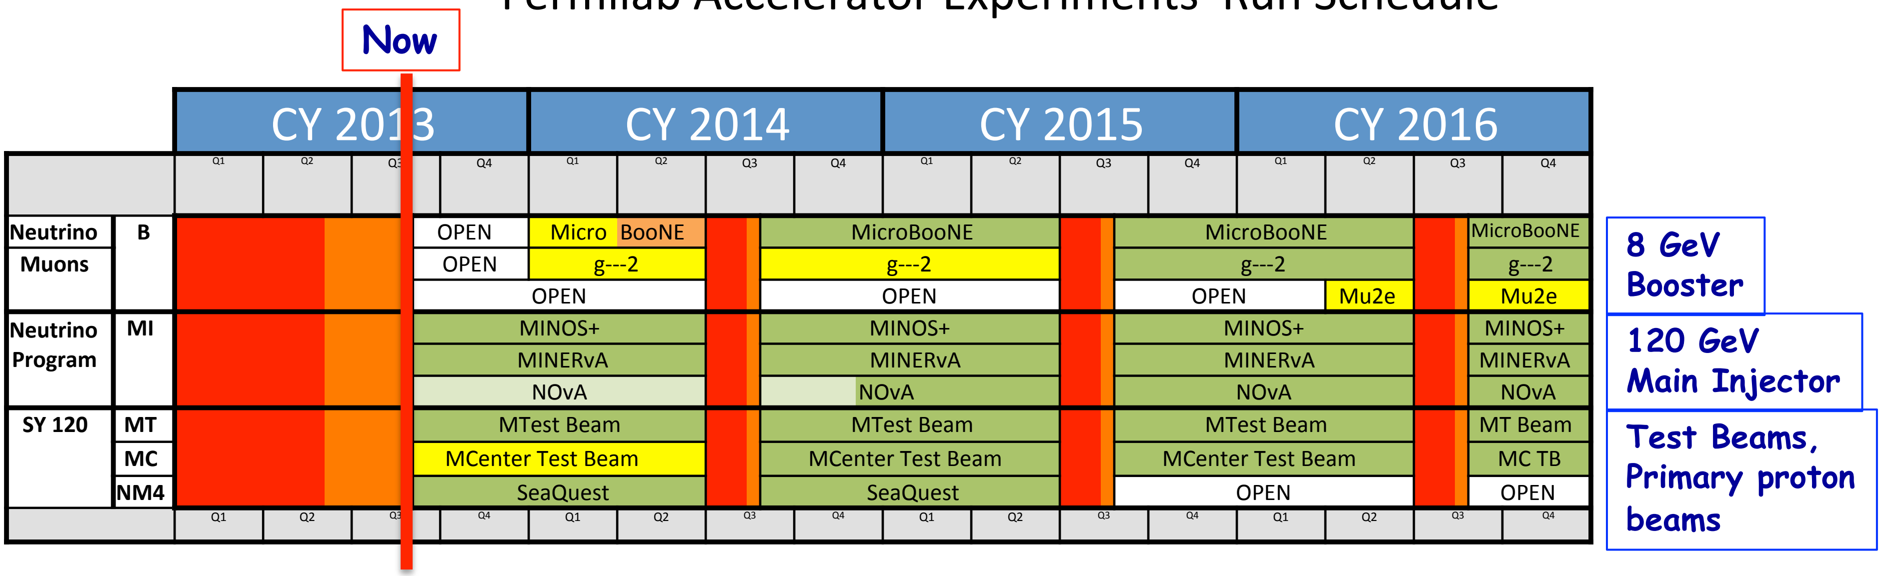

RUN/DATA STARTUP/
COMMISSIONING
INSTALLATION
M&D (SHUTDOWN)
INSTALLATION & RUNNING

Anticipate mid---year maintenance shutdowns of 6 week + 2 weeks commissioning

NOvA CD---4 end of Nov 2014, MicroBooNE first beam March/April 2014 MicroBooNE installation begins with cryostat arrival.

g---2 installation begins when building ready.

Mu2e istallation begins with benif. occup. mid---FY16. CD---4 Q3 of FY19.

CY13 MCenter Test Beam schedule not yet determined

Fermilab

Intensity Frontier Program (Diverse)

ArgoNeuT

Mu2e

SciBooNE

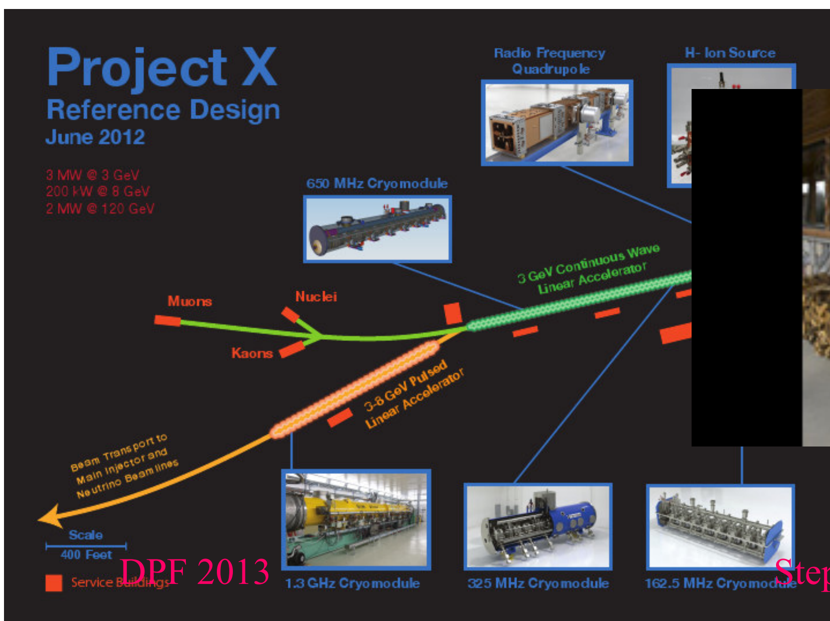

microboone

Muon g-2

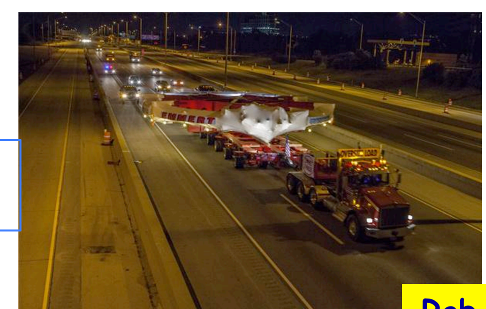

Rob Roser

Working toward LBNE

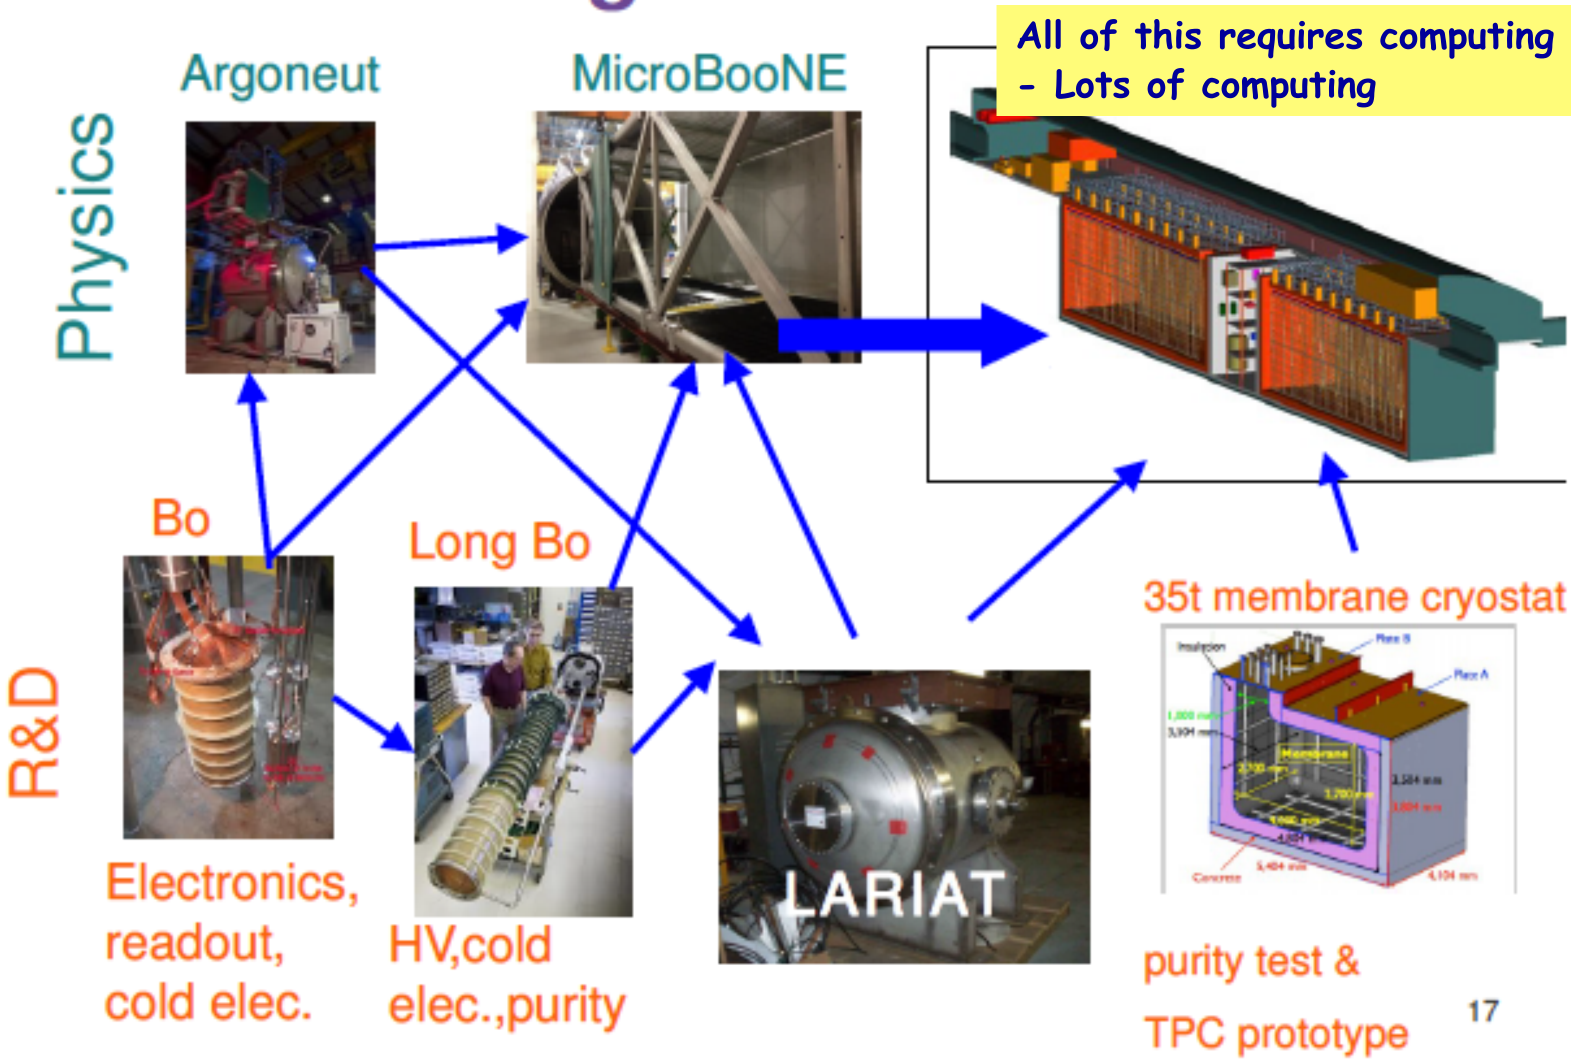

Slide from Mike Kordosky

INTENSITY FRONTIER COMPUTING REQUIREMENTS

Intensity Frontier Computing

- Computing is essential for a successful Intensity Frontier physics program.
- Computing is important for any physics program and this was discussed and investigated extensively in the recent Snowmass 2013 process.
  - It is important to properly recognize and plan the computing so that physics results are optimized.
- The transition from Run 2 to IF has led to a large shift in emphasis and resources (computing equipment and people – scientists, engineers, computer professionals) at Fermilab to the IF program.

IF Computing Requirements

- There are significant computing requirements for the Intensity Frontier program.
- The requirements in aggregate are significant, and approach the scale of the Run 2 experiments or a Tier 1 center for LHC (~5000 CPU slots at Fermilab, not including grid use off-site)
- The computing requirements are driven by high intensities and high precision for the following:
  - Beamline design and optimization,
  - Detector design,
  - DAQ and online systems,
  - Event simulation, including physics generators and detector response,
  - Reconstruction programs,
  - Trigger designs,
  - Large scale reconstruction,
  - Analysis.

Fermilab

The challenge

- Respond to the increasing demands of IF computing for new designs, new beamlines, new detectors, and analysis of data.
- Provide systems that allow the collaborations and others to do all of the above.
- Migrate experiment services to newer services that are more flexible, scalable, and aligned with Open Science Grid and Energy Frontier computing.
- Take advantage of and respond to major developments in computing technologies:
  - Parallel computing, multi-core, GPU's
  - Higher networking throughput
  - Multi-site computing
  - Virtualization
  - Other new developments

COMPUTING FOR THE INTENSITY FRONTIER

Strategy for IF computing

- Intensity Frontier experiments are typically 30-400 collaborators and are not each able to invent and support software, data handling and other systems required to collect and analyze data.
- Common approaches/solutions are desirable and essential to support this broad range of experiments with limited effort.
- The solutions are chosen with a view that the goal is to empower the physicists to concentrate on the physics code and not on the infrastructure around it.

Strategy (2)

- Effort will be made to integrate core software into robust solutions for each experiment.
- Data access and data handling is an important aspect of IF computing and allows use of worldwide resources.
- Technologies such as GPU's, multi-core processors, HPC computing, clouds all need to be considered and integrated.
- Training for physicists to allow developers, analyzers and others to contribute is being pursued.
- I'll talk about a few of the common projects and tools, with some emphasis on projects involving the Fermilab Scientific Computing Division.

Common projects, tools and facilities

- Framework (art)*
- Computing Architecture
- Generators, including GENIE and GEANT4
- Software Toolkits (LArSoft)*
- Tools and Utilities (svn, git, docDB, redmine,…)
- Data Handling (sam, FIFE project)*
- Computing facilities*
  - Interactive
  - Batch
  - GRID
  - Cloud
  - Parallel
- ROOT/analysis tools

* Will be discussed in this talk

Fermilab

art framework

- art is a C++ framework developed at Fermilab for IF (and other) experiments
- Essentially a follow-on of the CMS framework, also developed at Fermilab
- Used by g-2, NOvA, MicroBooNE, LBNE, Mu2e
- This allows shared development and support among the experiments and the developers.
- Integration into Fermilab's data-handling system, I/O, etc. is an important part of the art project.
- art-daq is variant of the framework used by Mu2e and others as part of the daq system
- Future directions include hooks for parallel processing at various levels of granularity.

Framework for the Intensity Frontier and others

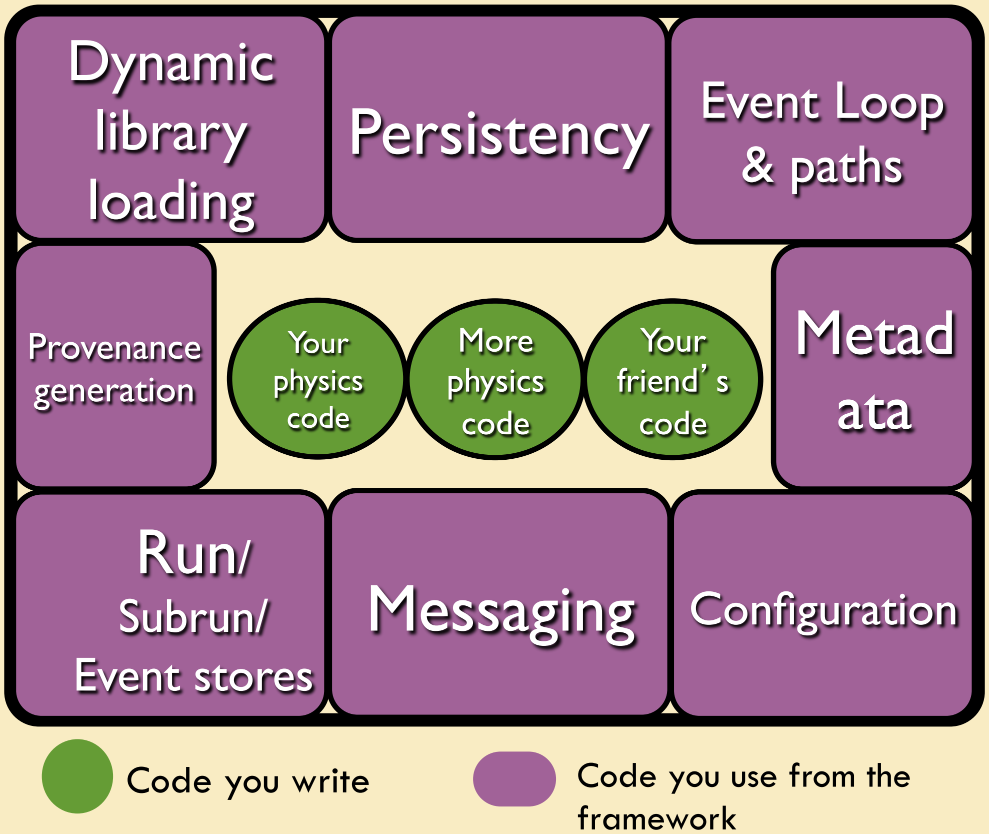

Next level of usability and collaboration

ART – A lite, forked version of CMSSW tailored for IF

Modularity makes it easy to collaborate

Physicists write physics code and algorithms, not infrastructure

Utilizes modern C++2011

Adopted by NOvA, Mu2e, MicroBoone, LBNE (LarSoft) Muon g-2, DarkSide50

Fermilab

LArSoft

- LArSoft is a common simulation, reconstruction and analysis toolkit for experiments using liquid argon time projection chamber (neutrino LArTPC's)
- The LArSoft package is managed by Fermilab's Scientific Computing Division (as of early 2013).
- The common joint effort allows multiple experiments to share ideas, algorithms, code, etc.
- Current experiments using LArSoft
  - ArgoNeut
  - MicroBooNE
  - LBNE
  - LArIAT
  - R&D activities

FIFE Project (FabrIc for Frontier Experiments)

integrated framework for offline analyses

7

- address all of the computing needs for experiments
- modular enough so that experiments can take what they need
- well enough designed so that while underlying solution may change, interface will be consistent
- allow for feedback from the experiments to incorporate their tools and experiments
- help experiments use computing beyond the walls of Fermilab
- allow new tools from outside communities to be incorporated as they develop

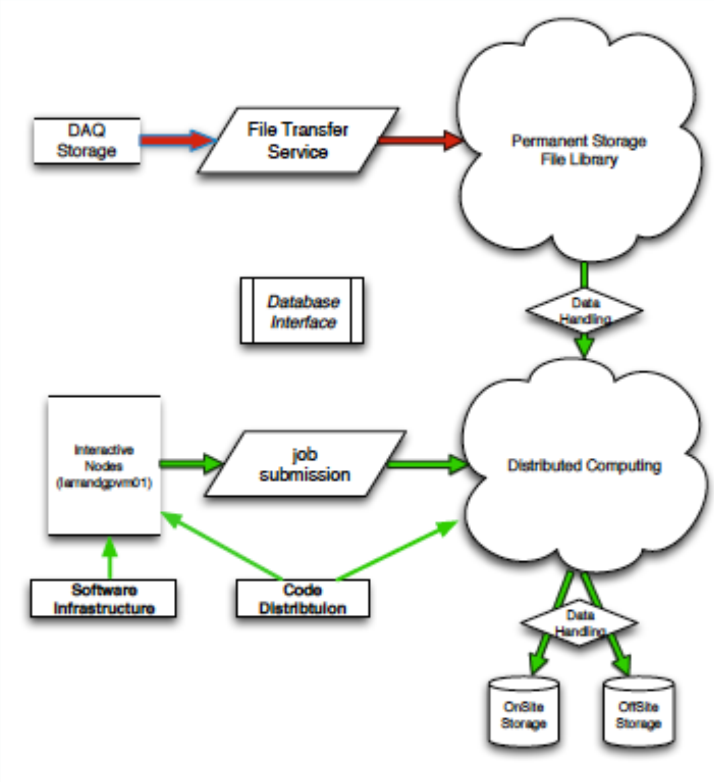

SCD FIFE Workshop 6/4/13

The FIFE Project for IF and others

A collection of projects that provide common computing services and interfaces needed to turn a physics task into results, enabling experiments to readily utilize offsite resources.

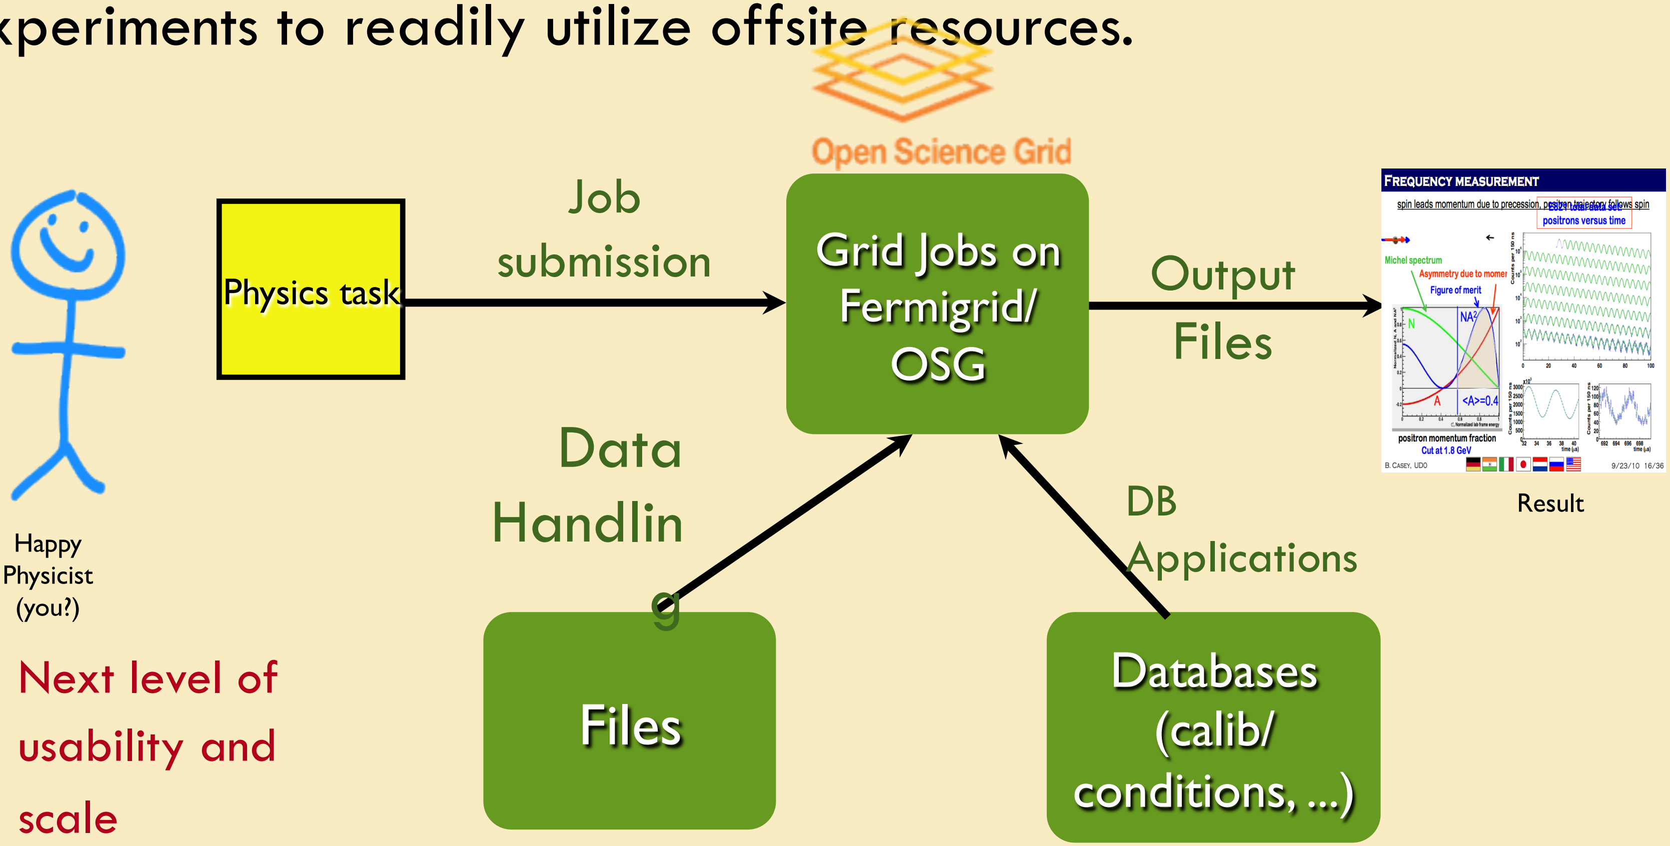

Next level of usability and scale

Tailored for opportunistic running;

CMS has similar system (workflow, glide-ins); adopting oppo running too

Computing Facilities at Fermilab

- Fermilab's computing facilities use a shared services model to ensure that all experiments, small and large, are able to make use of the facilities from "birth to grave"

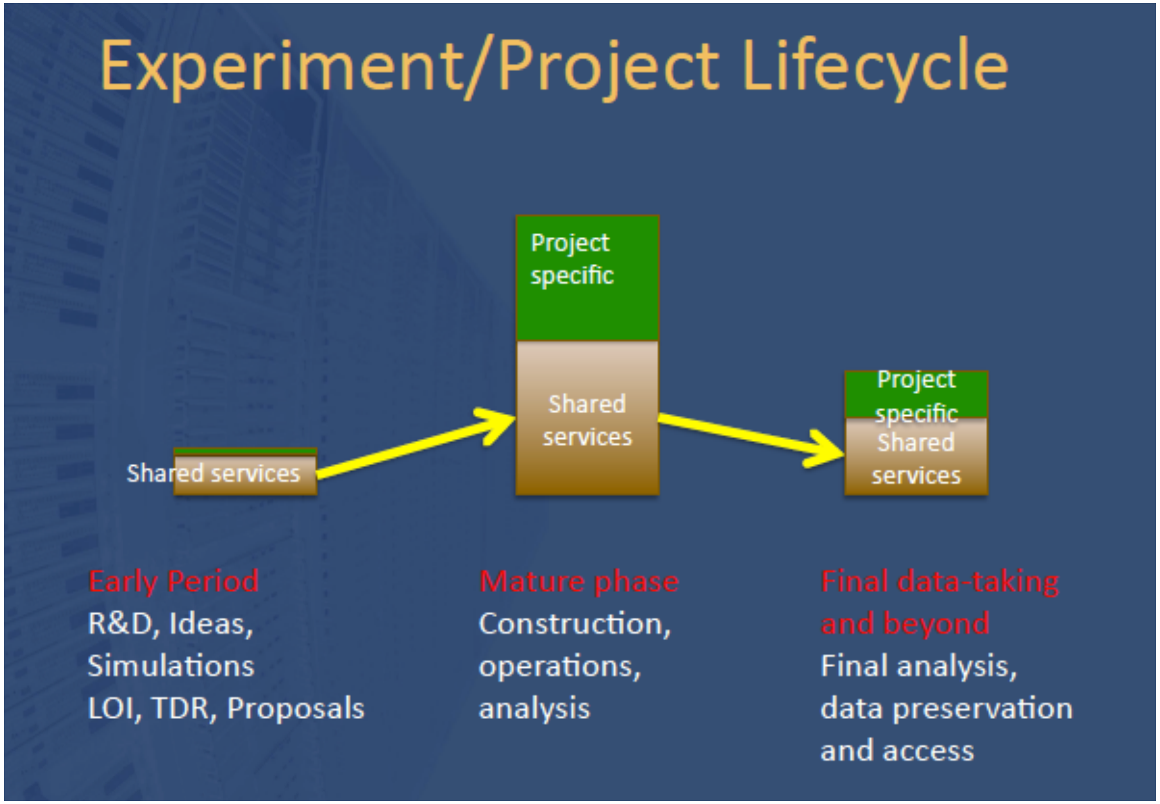

CPU: Recent growth of IF computing

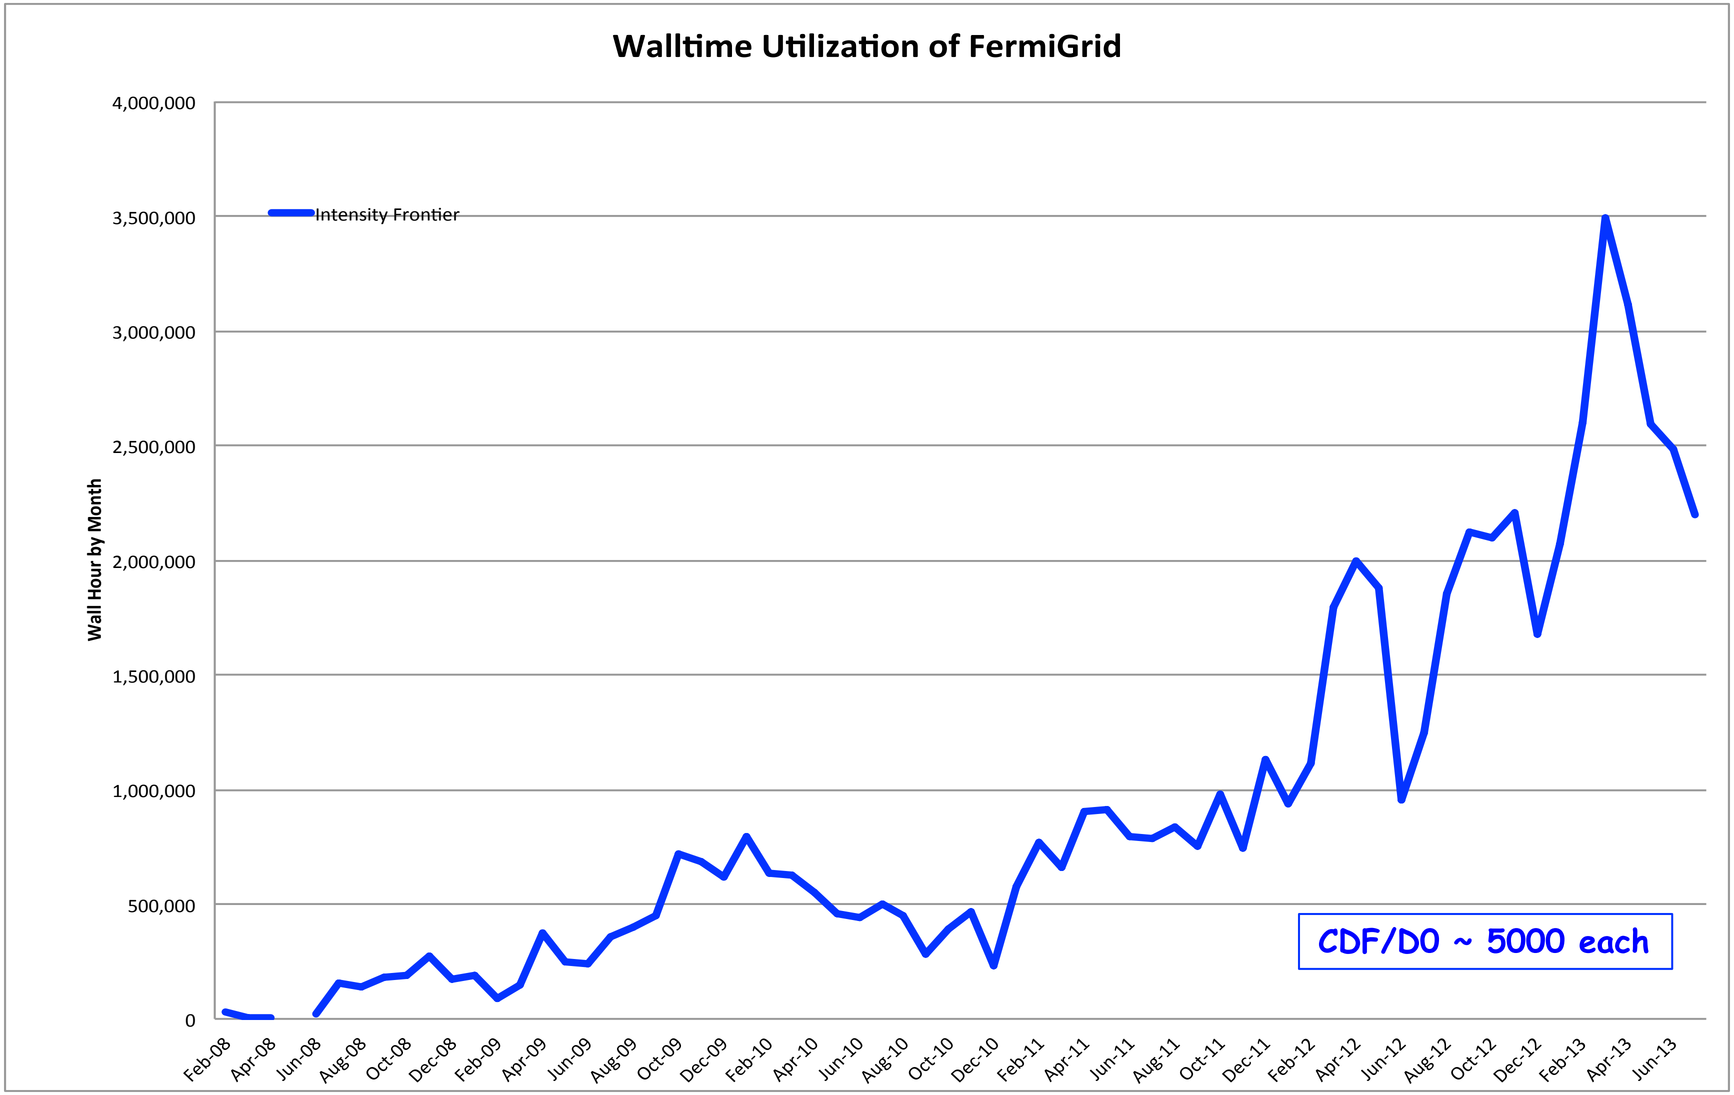

Cores

5500

2750

CPU: Individual experiments use >1000 cores

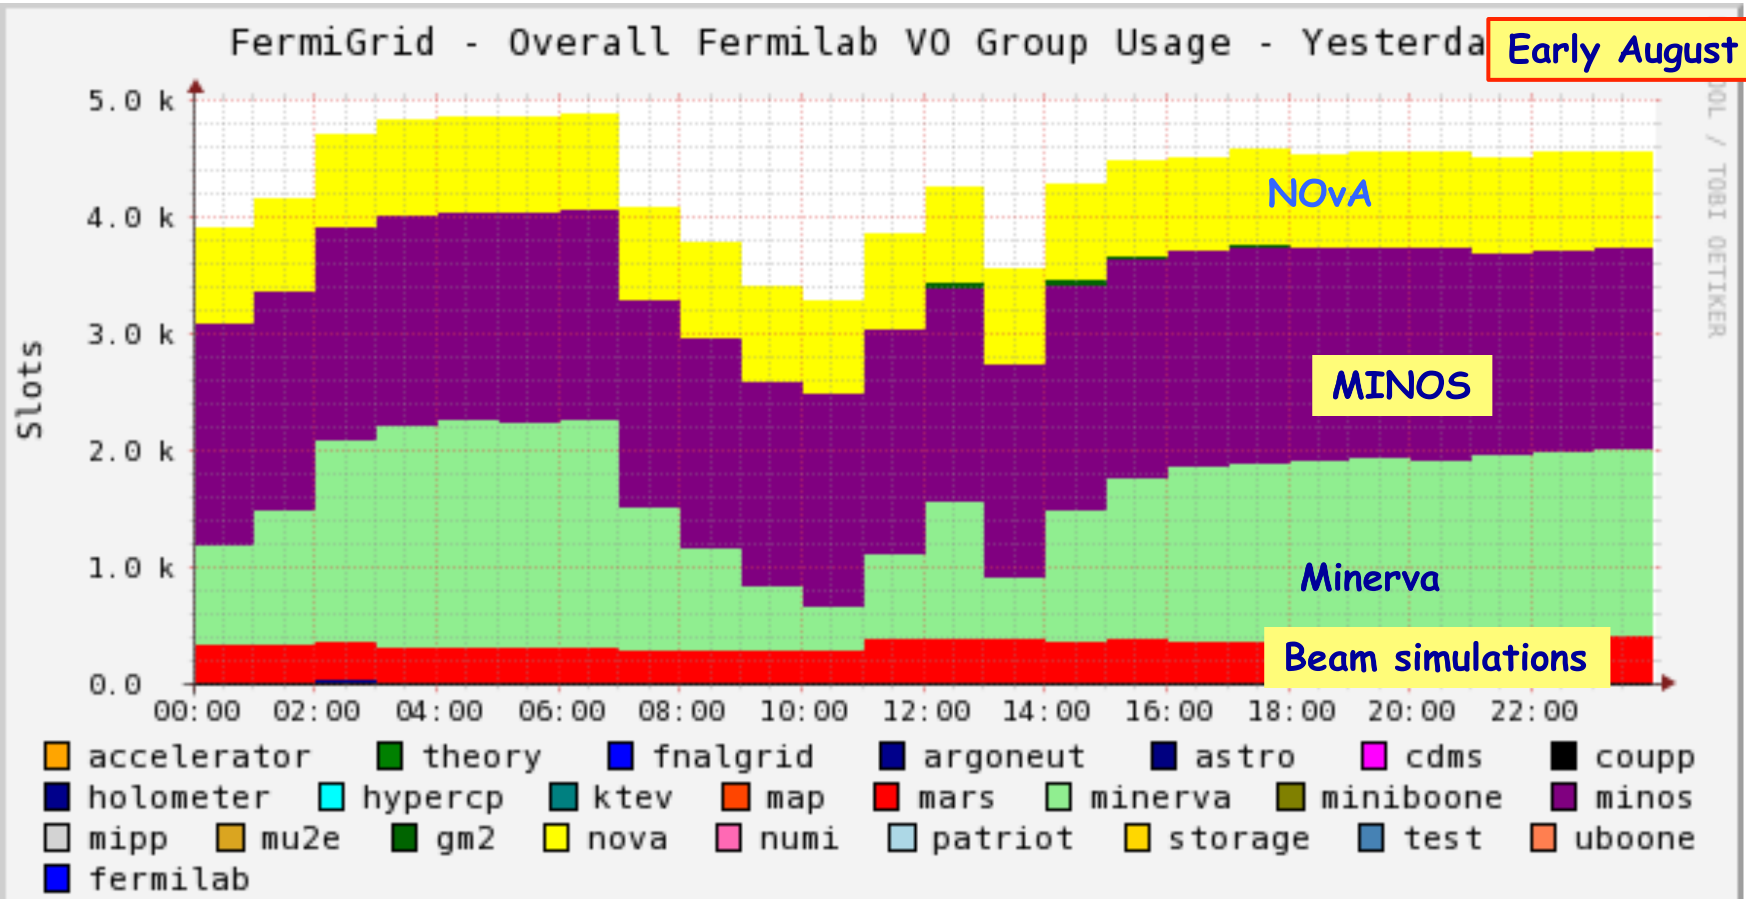

Fermilab

Projected growth of IF CPU needs

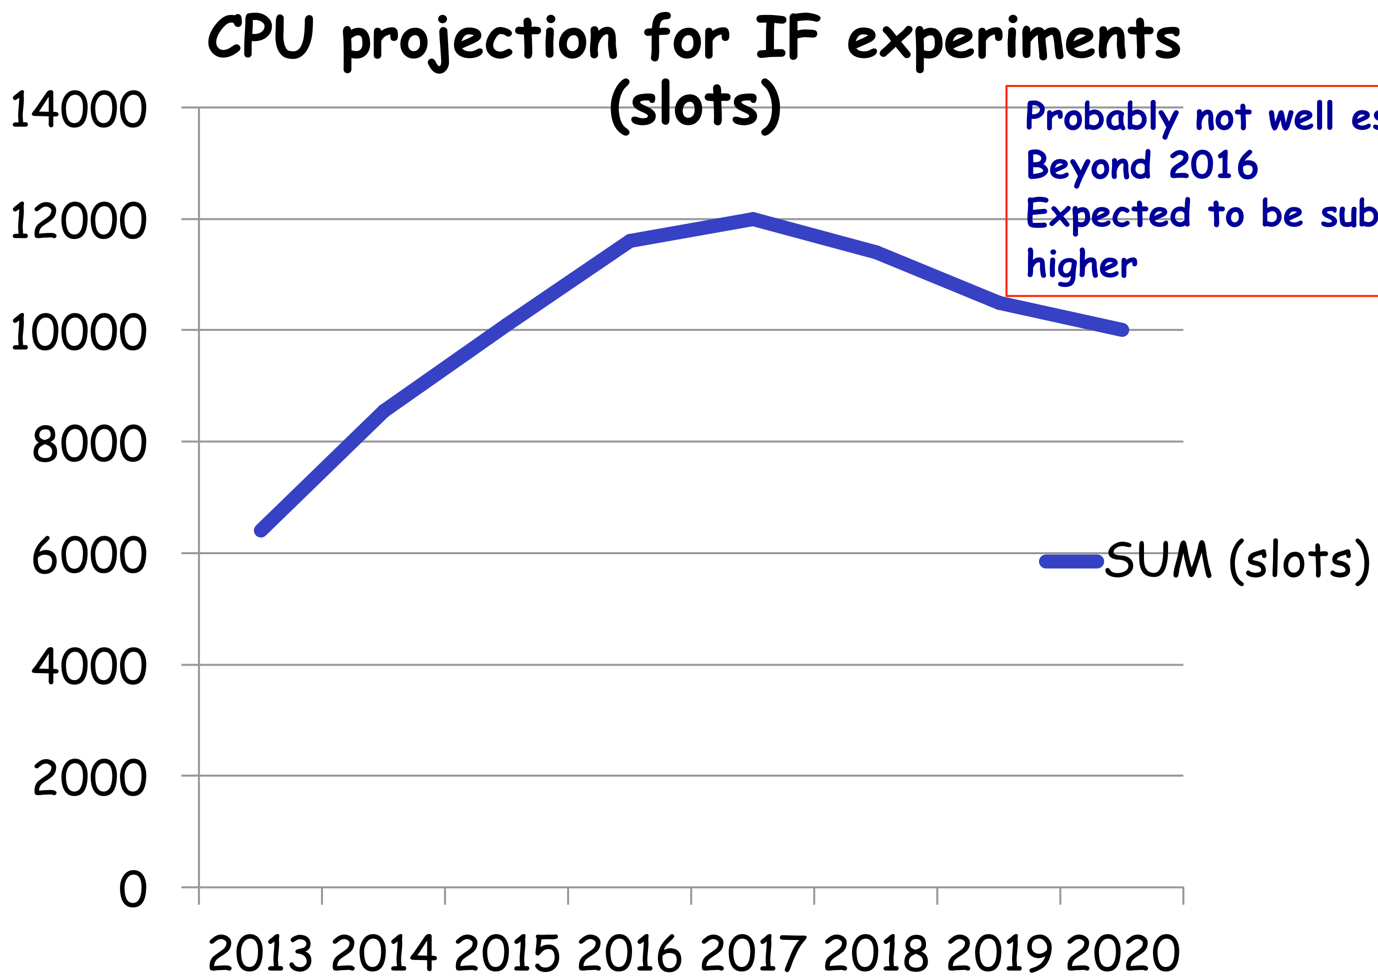

Probably not well estimated
Beyond 2016
Expected to be substantially higher

Taken from early Snowmass 2013 estimates. Will be updated.

Fermilab

Storage for IF experiments at FNAL

Fermilab Tape Storage
IF ~ 2 PB

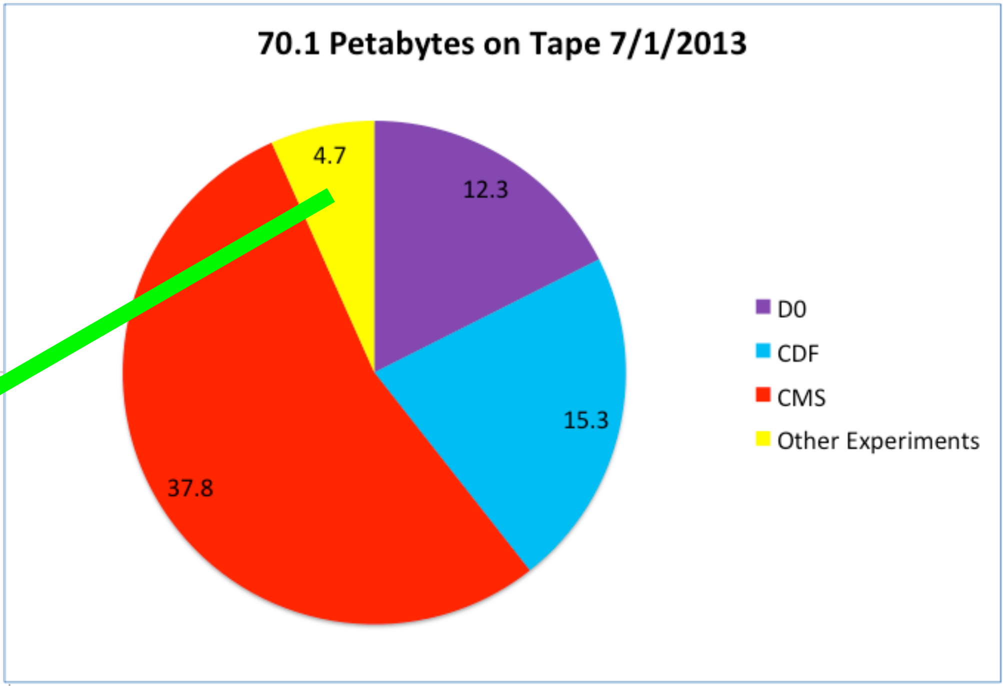

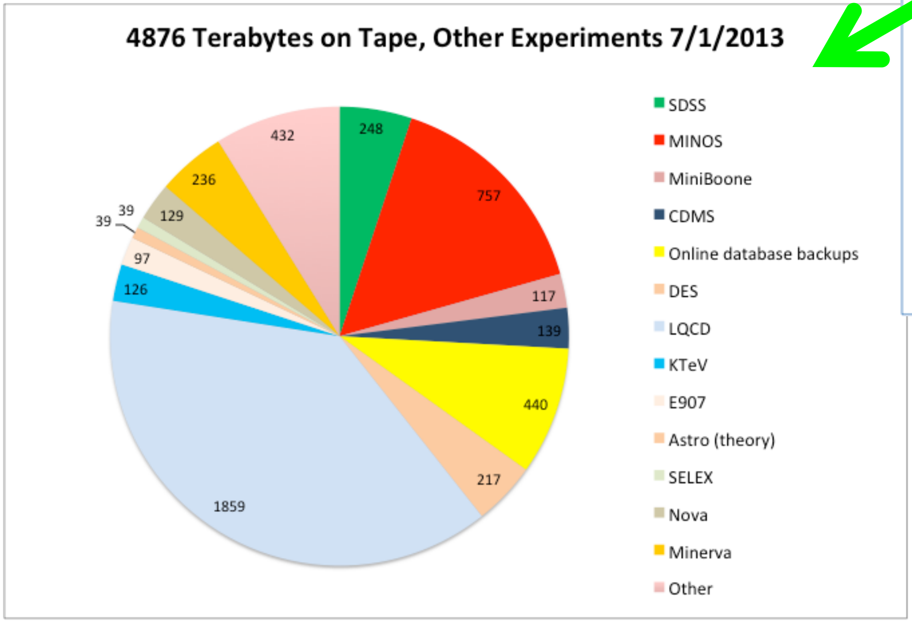

Fermilab

Projected growth in IF storage

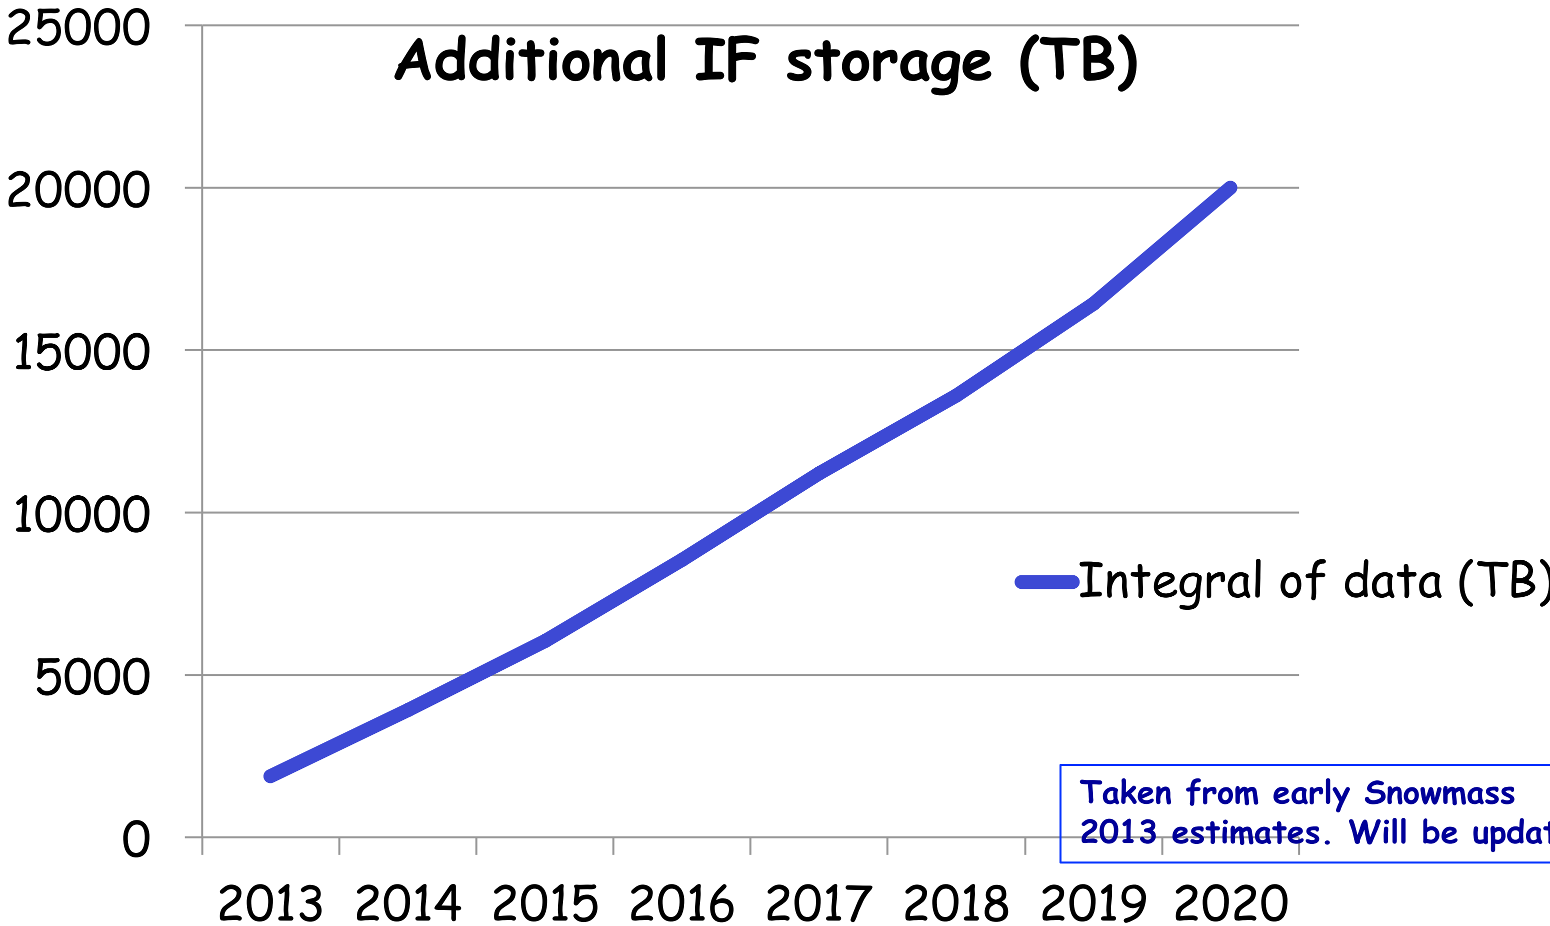

Taken from early Snowmass 2013 estimates. Will be updated.

Fermilab computing facilities

- Facilities include computer rooms, CPU, Storage, Networks

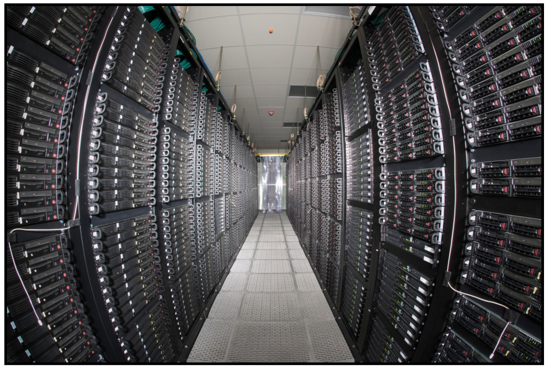

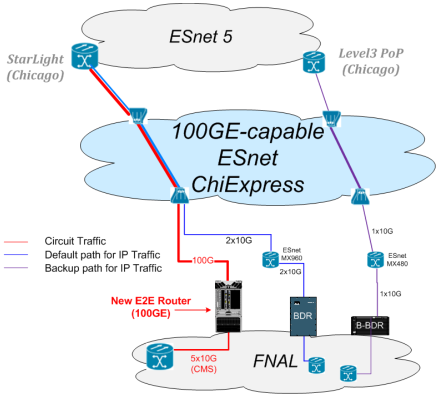

Upgrades in progress
Room for growth

Fermilab

Scientific Portfolio Review

- A yearly process has been initiated (starting in January 2013) to review the computing requirements and requests for the IF experiments.
- The result of the first review was used to
  - Prioritize effort,
  - Guide hardware purchases in FY13,
  - Study alternatives to computing model,
  - Change priorities for SCD from purchases to computing professional help
- The next review should occur early in FY14 to help inform and prioritize purchases and effort as the IF experiments take and analyze additional data.

Fermilab

Conclusion

- Fermilab has increased its emphasis on the Intensity Frontier program and continues to transition from Run 2 to Intensity Frontier experiments.
- Computing for the IF experiments is critical to ensuring that the experiments produce excellent science in a timely manner.
- Facilities, effort and projects aimed at improving the computing usability and capabilities are ramping up and this will continue as the IF program grows and matures.
  - Emphasis is on common and shared tools and solutions
  - Goal is to ensure smooth and effective physics results from the experiments as soon as possible

Extras

HEP Intensity Frontier Experiments

Outside US Taking data

List from DOE: There are MANY

Experiment	Location	Status	Description	#US Inst.	#US Coll.
Belle II	KEK, Tsukuba, Japan	Physics run 2016	Heavy flavor physics, CP asymmetries, new matter states	10 Univ, 1 Lab	55
BES III	IHEP, Beijing, China	Running	Precision measurements charm, charmonium, tau; search for and study new states of hadronic matter	6 Univ	26
CAPTAIN	Los Alamos, NM, USA	R&D; Test run 2015	Cryogenic apparatus for precision tests of argon interactions with neutrinos	5 Univ, 1 Lab	20
Daya Bay	Dapeng Penisula, China	Running	Precise determination of θ_{13}	13 Univ, 2 Lab	76
Heavy Photon Search	Jefferson Lab, Newport News, VA, USA	Physics run 2015	Search for massive vector gauge bosons which may be evidence of dark matter or explain g-2 anomaly	8 Univ, 2 Lab	47
KOTO	J-PARC, Tokai , Japan	Running	Discover and measure $K_L \to \pi^0 \nu\nu$ to search for CP violation	3 Univ	12
LArIAT	Fermilab, Batavia, IL	R&D; Phase I 2013	LArTPC in a testbeam; develop particle ID & reconstruction	11 Univ, 3 Lab	38
LBNE	Fermilab, Batavia, IL & Homestake Mine, SD, USA	CD1 Dec 2012; First data 2023	Discover and characterize CP violation in the neutrino sector; comprehensive program to measure neutrino oscillations	48 Univ, 6 Lab	336
MicroBooNE	Fermilab, Batavia, IL, USA	Physics run 2014	Address MiniBooNE low energy excess; measure neutrino cross sections in LArTPC	15 Univ, 2 Lab	101
MINERvA	Fermilab, Batavia, IL, USA	Med. Energy Run 2013	Precise measurements of neutrino-nuclear effects and cross sections at 2-20 GeV	13 Univ, 1 Lab	48
MINOS+	Fermilab, Batavia, IL & Soudain Mine, MN, USA	NuMI start-up 2013	Search for sterile neutrinos, non-standard interactions and exotic phenomena	15 Univ, 3 Lab	53
Mu2e	Fermilab, Batavia, IL, USA	First data 2019	Charged lepton flavor violation search for $\mu N \to eN$	15 Univ, 4 Lab	106
Muon g-2	Fermilab, Batavia, IL, USA	First data 2016	Definitively measure muon anomalous magnetic moment	13 Univ, 3 Lab, 1 SBIR	75
NOvA	Fermilab, Batavia, IL & Ash River, MN, USA	Physics run 2014	Measure ν_μ-ν_e and ν_μ-ν_μ oscillations; resolve the neutrino mass hierarchy; first information about value of δ_{cp} (with T2K)	18 Univ, 2 Lab	114
ORKA	Fermilab, Batavia, IL, USA	R&D; CD0 2017+	Precision measurement of $K^+ \to \pi^+ \nu\nu$ to search for new physics	6 Univ, 2 Lab	26
Super-K	Mozumi Mine, Gifu, Japan	Running	Long-baseline neutrino oscillation with T2K, nucleon decay, supernova neutrinos, atmospheric neutrinos	7 Univ	29
T2K	J-PARC, Tokai & Mozumi Mine, Gifu, Japan	Running; Linac upgrade 2014	Measure ν_μ-ν_e and ν_μ-ν_μ oscillations; resolve the neutrino mass hierarchy; first information about value of δ_{cp} (with NOvA)	10 Univ	70
US-NA61	CERN, Geneva, Switzerland	Target runs 2014-15	Measure hadron production cross sections crucial for neutrino beam flux estimations needed for NOvA, LBNE	4 Univ, 1 Lab	15
US Short-Baseline Reactor	Site(s) TBD	R&D; First data 2016	Short-baseline sterile neutrino oscillation search	6 Univ, 5 Lab	28

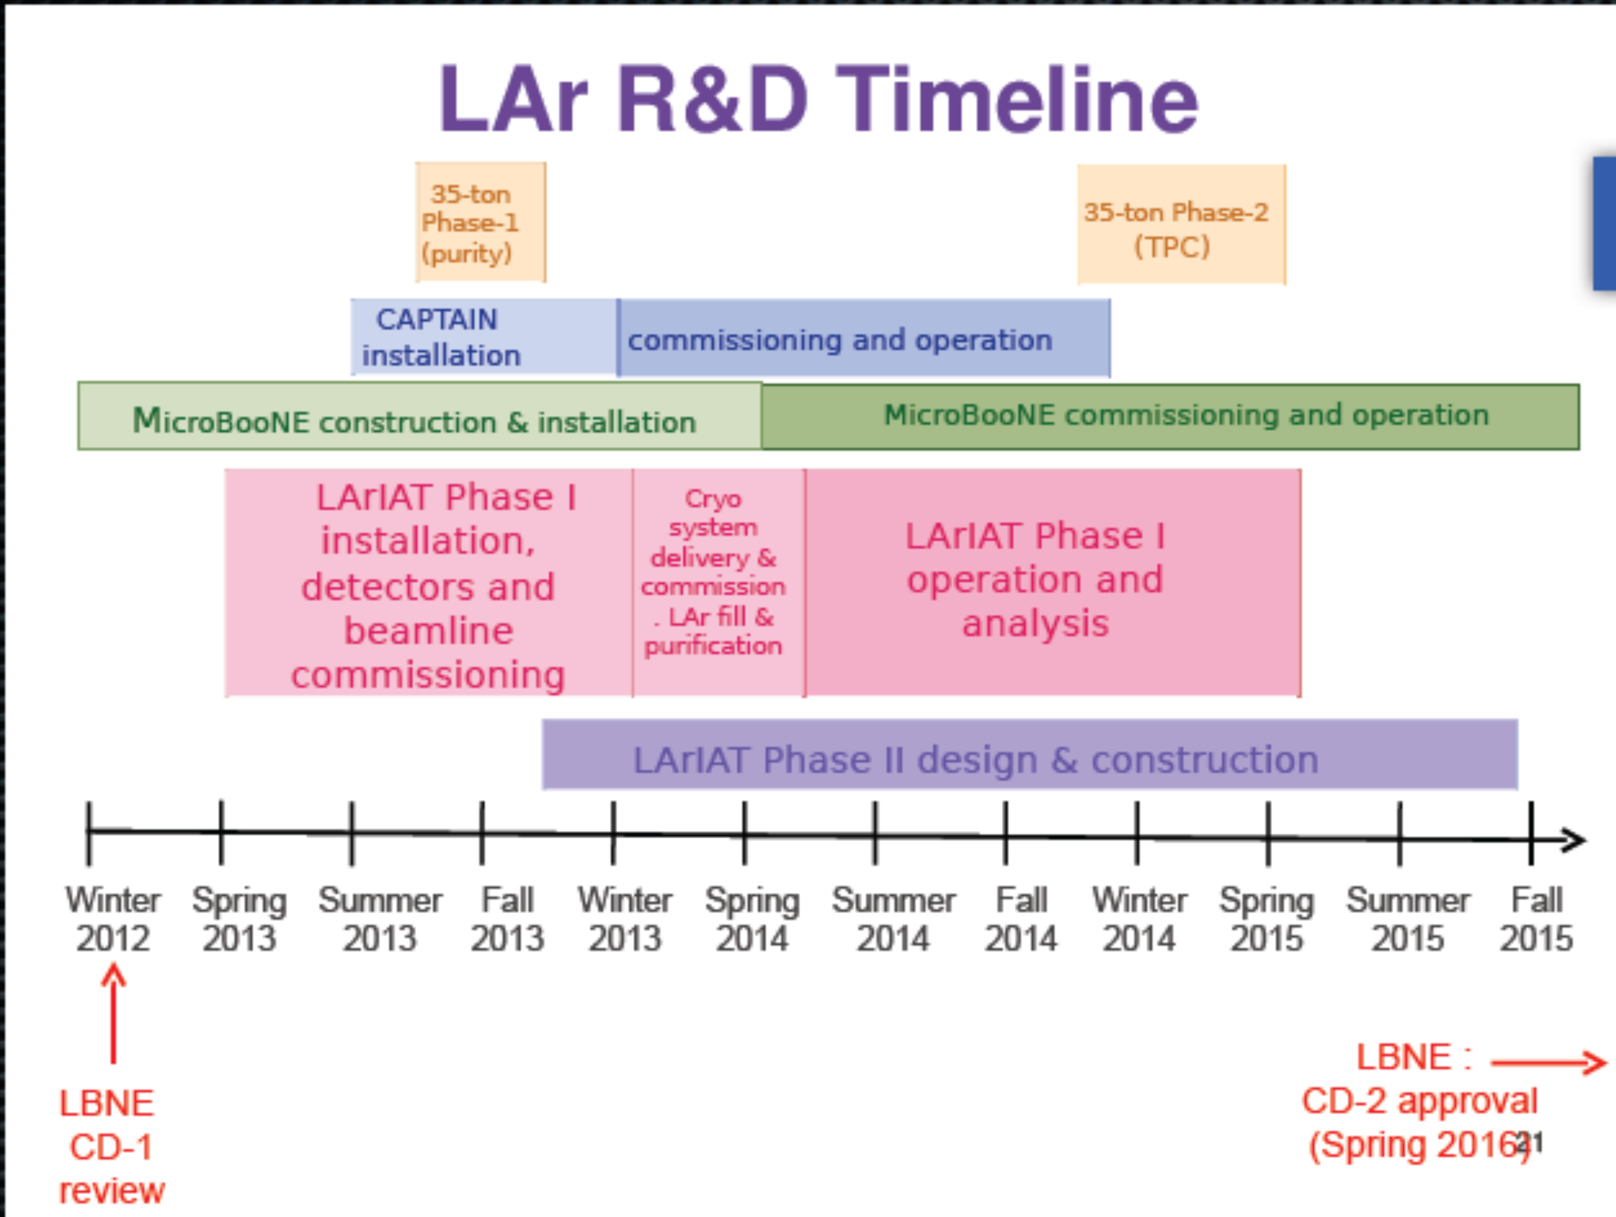

LAr1
2018

Slide from Mike Kordosky

Computing Frontier Charge

✦ Computing has become essential to advances in experimental and many areas of theoretical physics. Research requirements in these areas have led to advances in computational capabilities. The participants in the Computing Frontier will address these issues:

- What are the computational requirements for carrying out the experiments that will lead to advances in our physics understanding?
- What are the computational requirements for theoretical computations and simulations that will lead to advances in our physics understanding?
- What facility and software infrastructure must be in place in order to meet these requirements, and what research investments does it require in computing, storage, networking, application frameworks, algorithms, programming, etc. to provide that infrastructure?
- What are the training requirements to assure that personnel are available to meet the needs?

•Stephen Wolbers

From Mayly Sanchez, Snowmass 2013

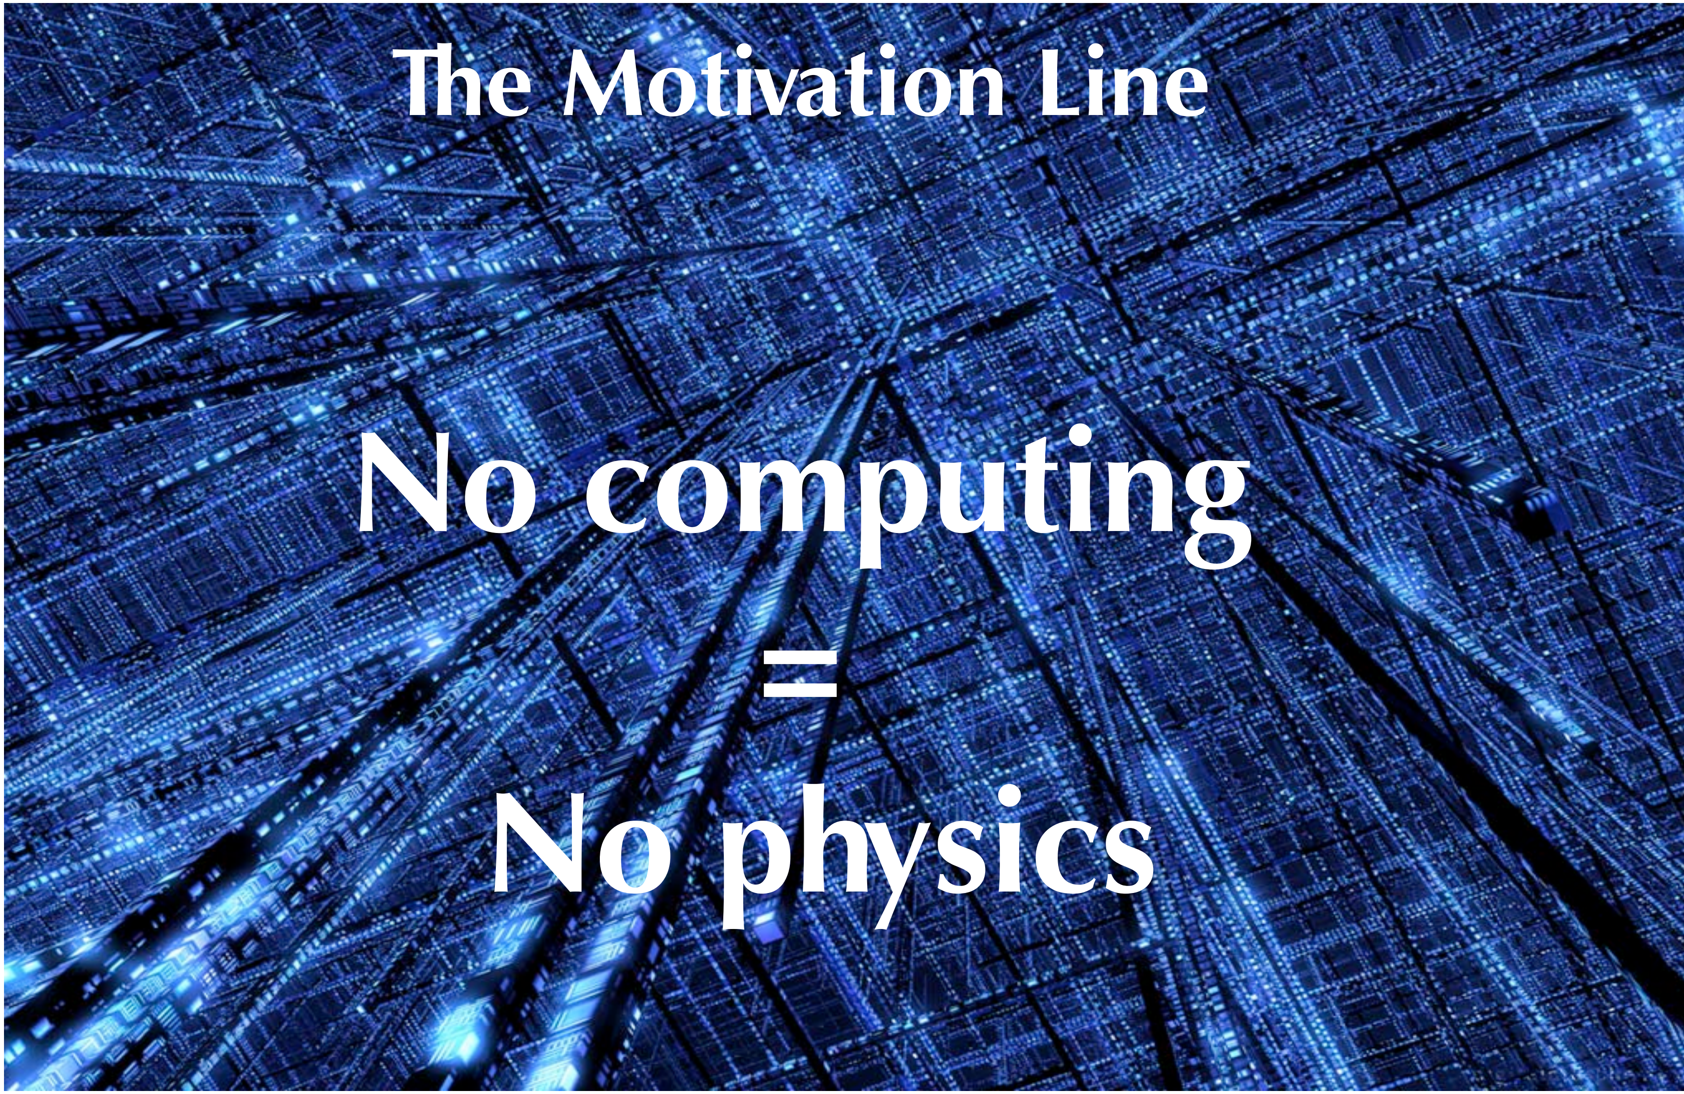

Scale of IF computing

- A survey of IF experiments was taken to identify the computing needs both in the scale of resources required and the software and system developments desired.
- The scale of each experiment request is in general not huge. However, the experiments are normally not large enough to provide the human and computing resources to handle the needs.
- There are many requests for tools and computing professional support.

Intensity Frontier Strategy

- Common approaches/solutions are essential to support this broad range of experiments with limited SCD staff. Examples include ArtDAQ, ART, SAM IF, LArSoft, Jobsub,…

- SCD has established a liaison between ourselves and experiments to insure communication and understand needs/ requirements

- Completing the process of establishing MOU's between SCD and experiment to clarify our roles/responsibilities

Rob Roser

Intensity Frontier Strategy - 2

- A shared analysis facility where we can quickly and flexibly allocate computing to experiments
- Continue to work to “grid enable” the simulation and processing software
  - Good success with MINOS, MINERvA and Mu2e
- All experiments use shared storage services – for data and local disk – so we can allocate resources when needed
- Perception that intensity frontier will not be computing intensive is wrong

Rob Roser

Fermilab

Fermilab

artdaq Introduction

artdaq is a toolkit for creating data acquisition systems to be run on commodity servers

- It is integrated with the *art* event reconstruction and analysis framework for event filtering and data compression.
- provides data transfer, event building, process management, system and process state behavior, control messaging, message logging, infrastructure for DAQ process and *art* module configuration, and writing of data to disk in ROOT format.
- The goal is to provide the common, reusable components of a DAQ system and allow experimenters to focus on the experiment-specific parts of the system. This software that reads out and configures the experiment-specific front-end hardware, the analysis modules that run inside of *art*, and the online data quality monitoring modules.
- As part of our work in building the DAQ software system for upcoming experiments, such as Mu2e and Darkside 50, we will be adding more features

art architecture

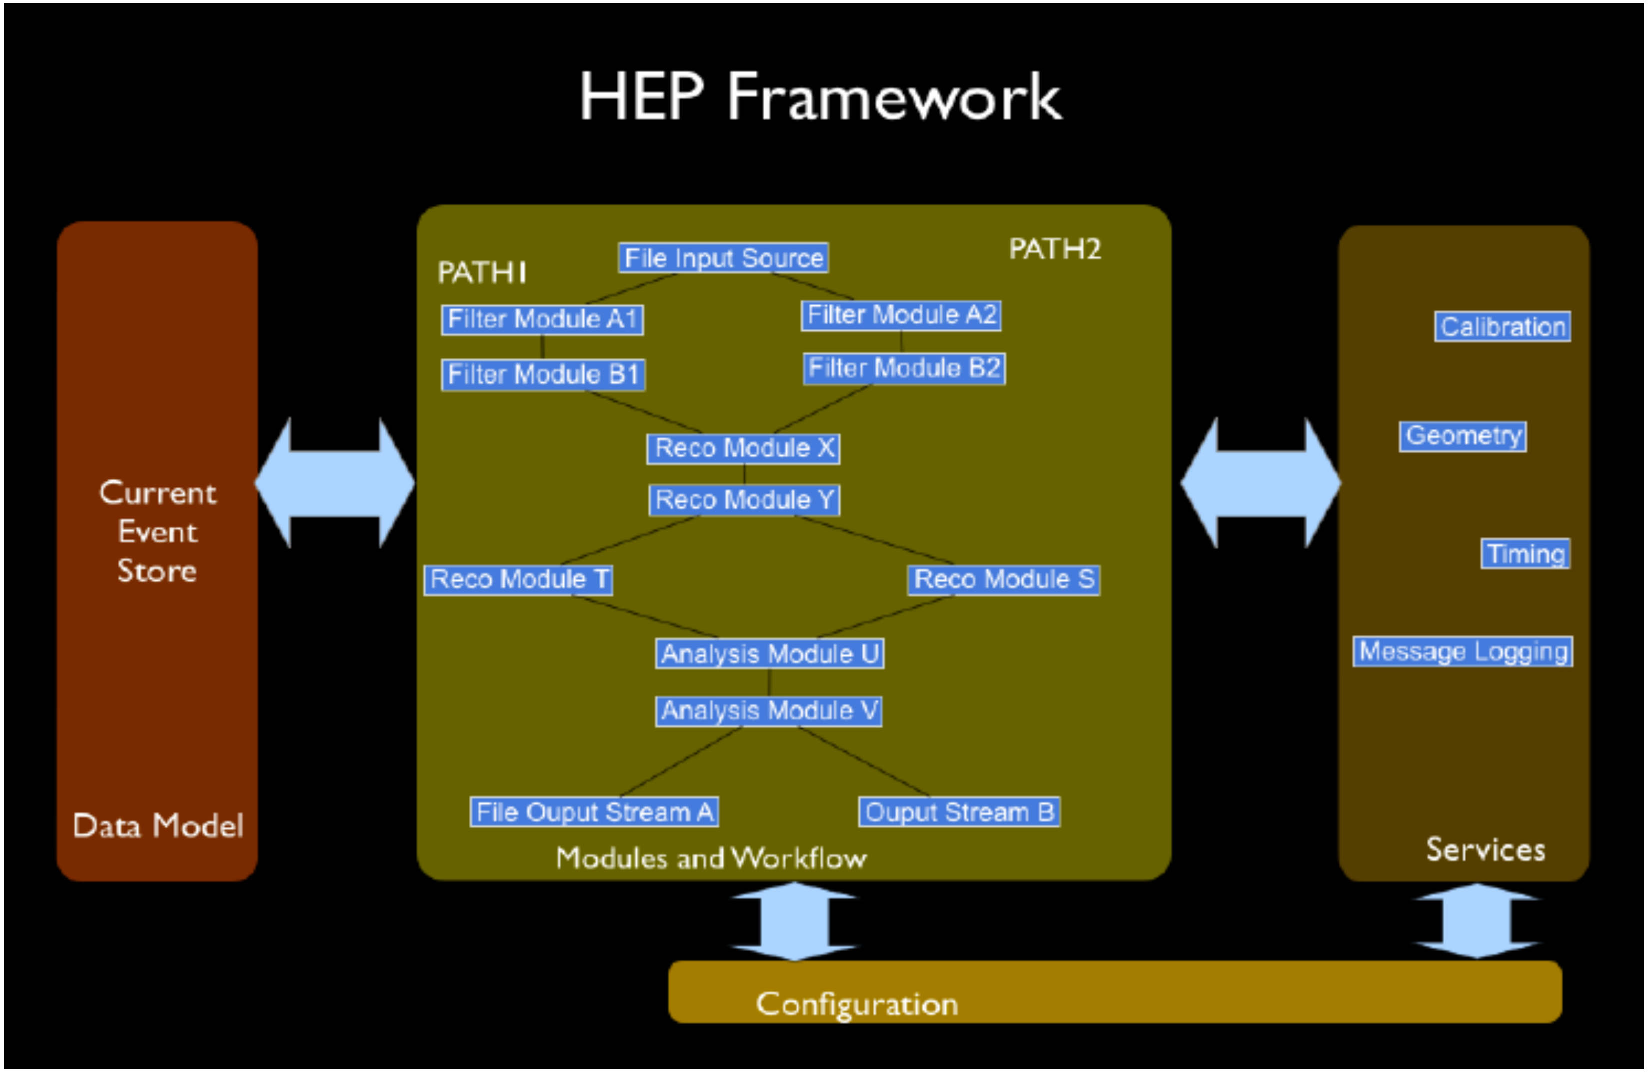

artdaq Introduction

We are currently working with the DarkSide-50 collaboration to develop and deploy their DAQ system using *artdaq*.

- The DS-50 DAQ reads out ~15 commercial VME modules into four front-end computers using commercial PCIe cards and transfers the data to five event builder and analysis computers over a QDR Infiniband network.
- The maximum data rate through the system will be 500 MB/s, and we have achieved a data compression factor of five.
- The DAQ system is being commissioned at LNGS, and it is being used to collect data and monitor the performance of the detector as it is being commissioned. (plots of phototube response?)

artdaq will be used for the Mu2e DAQ, and we are working toward a demonstration system which reads data from the candidate commercial PCIe cards, builds complete events, runs sample analysis modules, and writes the data to disk for later analysis.

- The Mu2e system will have 48 readout links from the detector into commercial PCIe cards, and the data rate into the PCIe cards will be ~30 GB/s. Event fragments will be sent to 48 commodity servers over a high-speed network, and the online filtering algorithms will be run in the commodity servers.
- We will be developing the experiment-specific artdaq components as part of creating the demonstration system, and this system will be used to validate the performance of the baseline design in preparation for the CD-review early next year.

Architecture task force

- After some years of experience it was realized that the computing architecture used by IF experiments (miniBooNE, MINOS, MINERvA) was not scalable.
- A task force was charged to propose a new architecture for the IF experiments.
- The output of this task force, along with the FIFE project and experience gained from US CMS and others, will be used to build a new architecture for the IF experiments that is expandable to scales well beyond the current scales and into the future.

IF computing use in 2010-2011

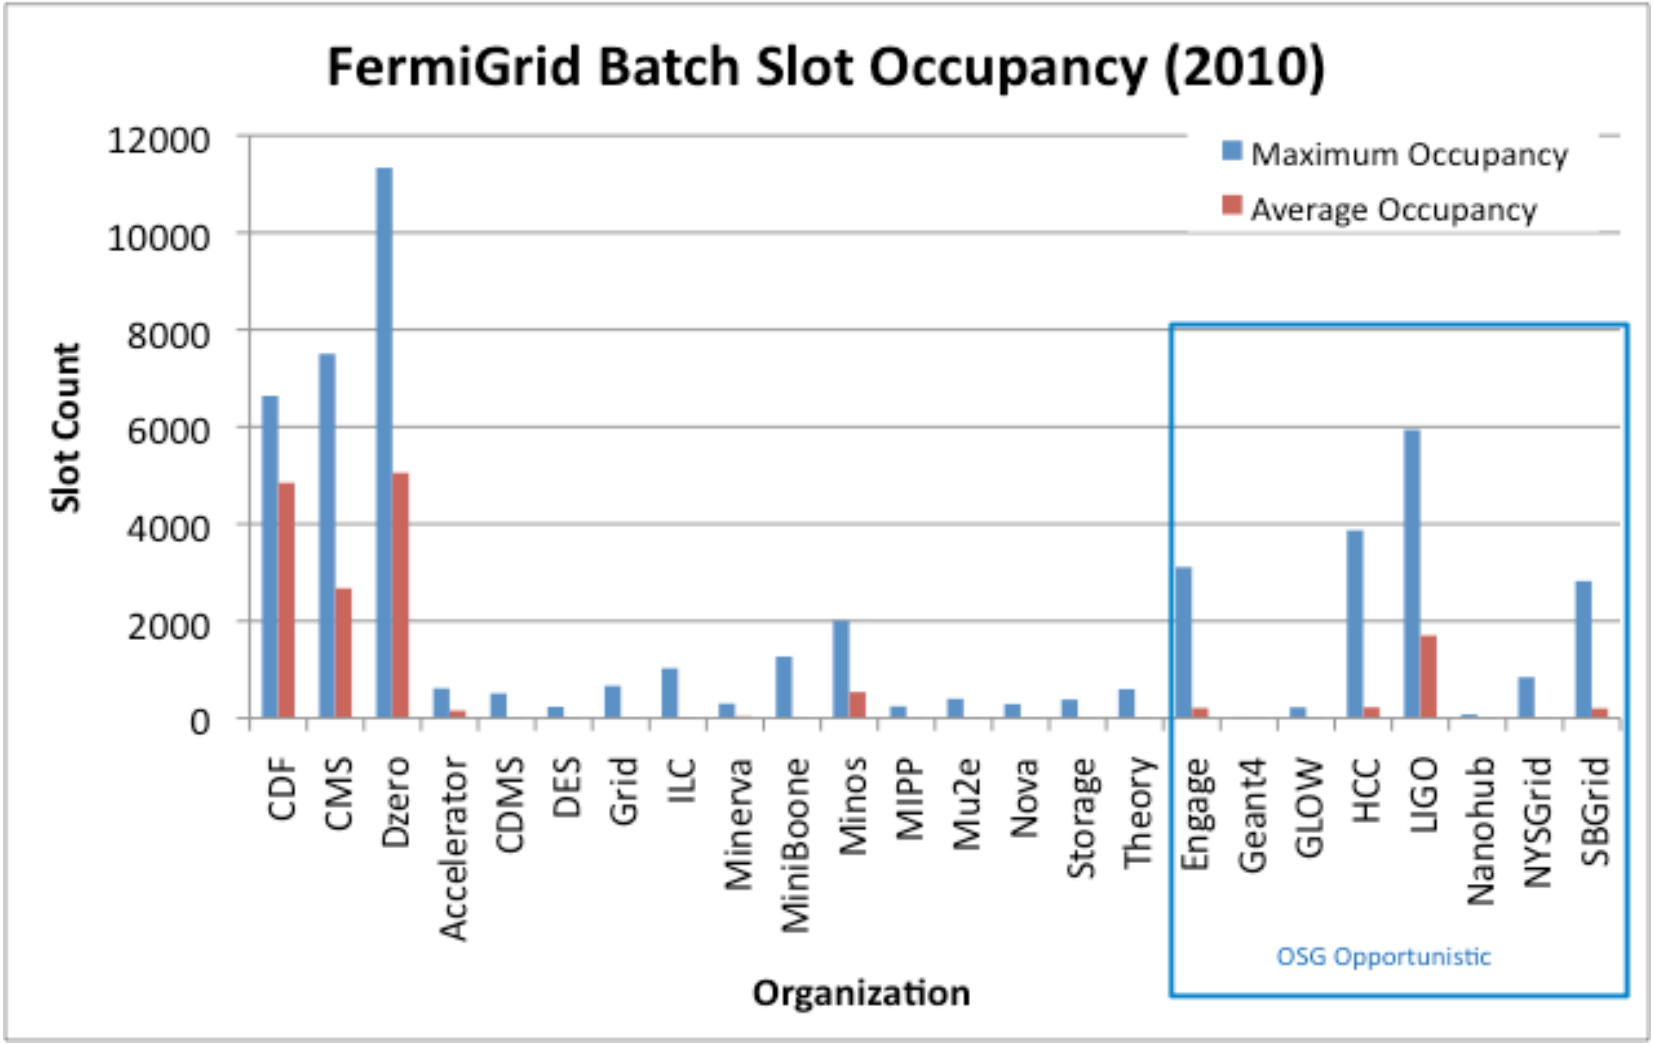

IF CPU usage for the past year

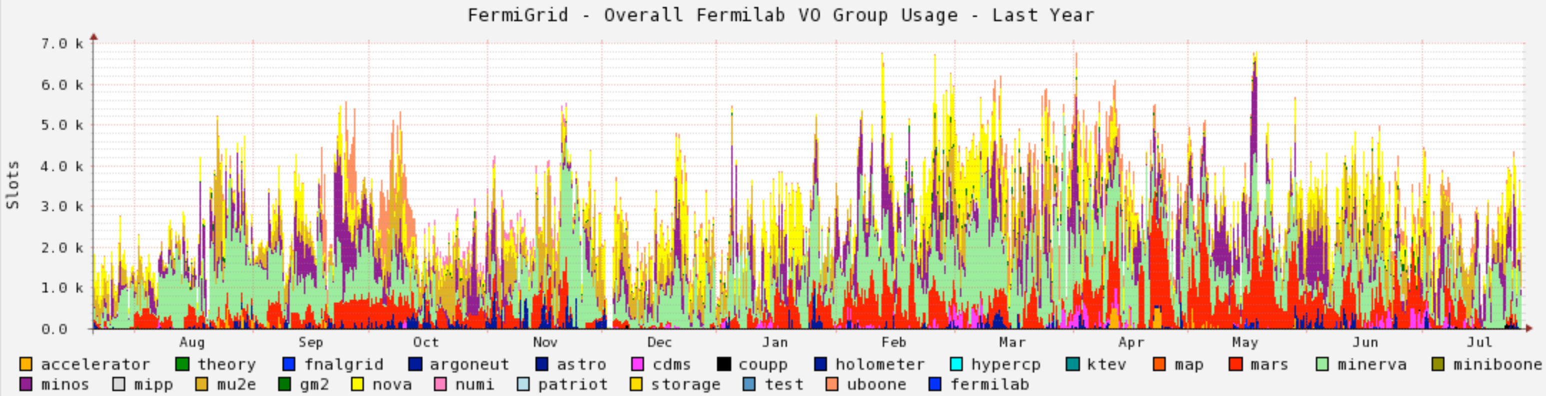

	Maximum	Average	Minimum	LastVal
total slots	25019	23130.02	12643	21950
fermilab	6796	2997.65	10	2862
accelerator	499	11.64	0	0
theory	29	0.04	0	0
fnalgrid	251	1.64	0	0
argoneut	971	55.02	0	0
astro	0	0.00	0	0
cdms	500	33.13	0	0
coupp	120	1.05	0	0
holometer	0	0.00	0	0
hypercp	0	0.00	0	0
ktev	0	0.00	0	0
map	178	0.61	0	0
mars	3228	465.98	0	2
minerva	4533	930.84	0	824
miniboone	0	0.00	0	0
minos	2899	466.69	1	130
mipp	200	5.72	0	55
mu2e	3456	443.59	0	849
gm2	233	17.66	0	0
nova	1783	438.15	0	1001
numi	206	15.83	0	0
patriot	5	0.14	0	0
storage	0	0.00	0	0
test	0	0.00	0	0
uboone	2065	117.83	0	1

Merged data from all FermiGrid clusters between 21-Jul-2012 and 28-Jul-2013
Plot generated at 28-Jul-2013 15:35:14 on monitor1.fnal.gov