



# Fabric for Frontier Experiments at Fermilab

Gabriele Garzoglio

Grid and Cloud Services Department, Scientific Computing Division, Fermilab

ISGC – Thu, March 27, 2014

# Overview

- Frontier Physics at Fermilab: Computing Challenges
- Strategy: common infrastructure on distributed resources
- Early successes with multiple communities

# Diverse program of experiments





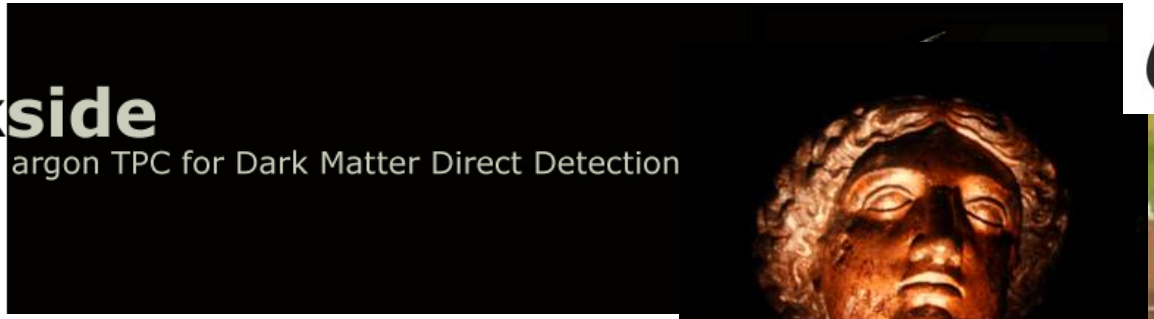







Lariat

# Fermilab Experiment Schedule

- Measurements at all frontiers – Electroweak physics, neutrino oscillations, muon g-2, dark energy, dark matter
- 8 major experiments in 3 frontiers running simultaneously in 2016
- Sharing both beam and computing resources
- Impressive breadth of experiments at FNAL

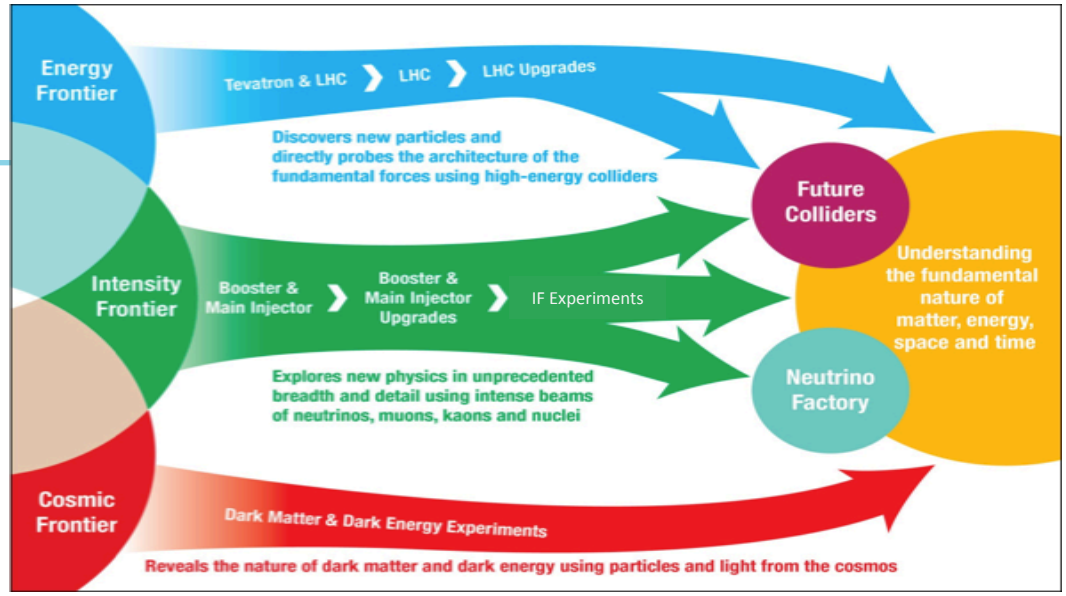


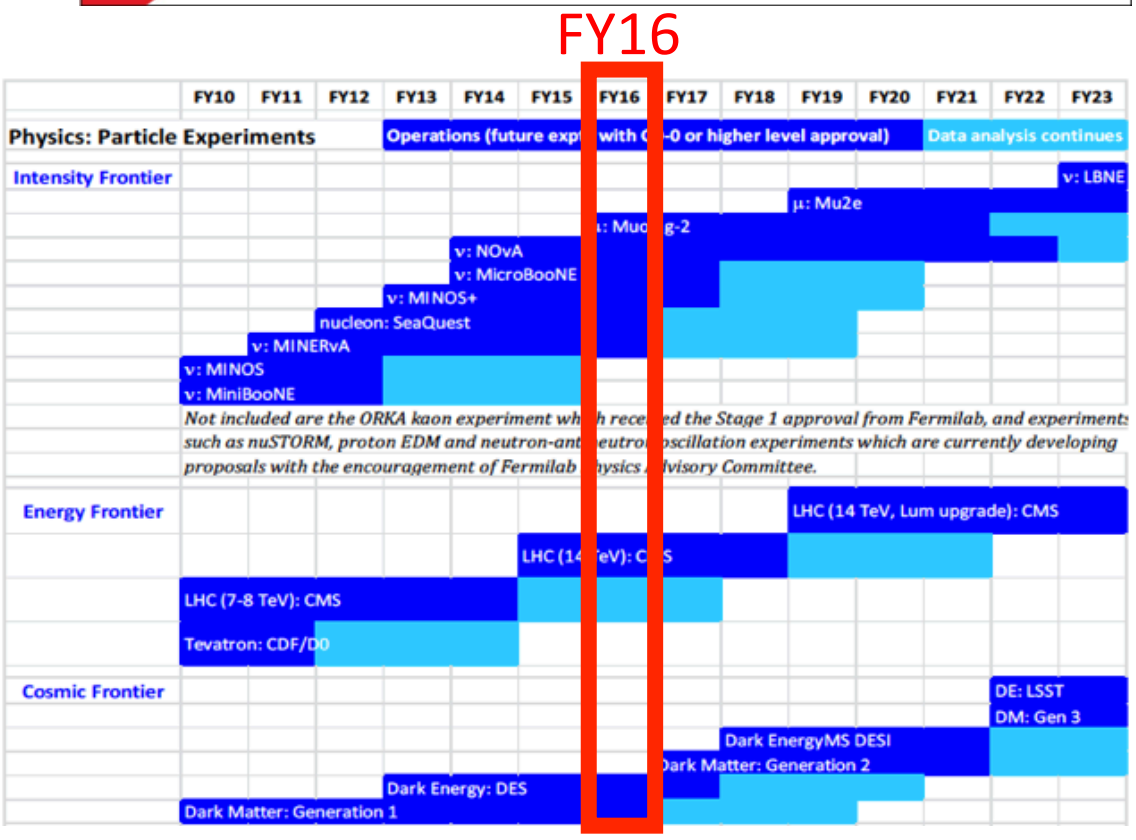

# Computing requirements for experiments

- higher intensity and higher precision measurements are driving request for more computing resources than previous “small” experiments
- beam simulations to optimize experiments - make every particle count
- detector design studies - cost effectiveness and sensitivity projections
- higher bandwidth DAQ and greater detector granularity
- event generation and detector response simulation
- reconstruction and analysis algorithms

Fermilab Scientific Computing Review*

| Experiment | Allocation (Slots) | Average Utilization (last quarter) Average Slots | Average Utilization (last quarter) Peak Slots | FY14 Request (Slots) |
|---|---|---|---|---|
| ArgoNeuT | 200 | 75 | 1712 | 0 |
| Muon g-2 | 200 | 9 | 195 | 0 |
| LBNE | 500 | 30 | 1202 | 200 |
| MARS | 1225 | 212 | 1444 | 100 |
| MicroBooNE | 500 | 39 | 597 | 100 |
| Minerva | 1600 | 667 | 3580 | 500 |
| Minos | 1200 | 294 | 2427 | 0 |
| Mu2e | 500 | 597 | 2491 | 500 |
| Nova | 1300 | 410 | 1799 | 500 |
| Total Requested | 7225 | 2333 | 15447 | 1900 |

* Number of slots have been updated periodically since

# Available onsite resources

- allowed experiments to develop architectures based on local resources
- now reintegrating them into a distributed computing model
- 30,000 slots on-site available through Open Science Grid interfaces
  - ~7,000 slots for smaller experiments, but aggregate peak need is twice as much
  - off-site resources are needed to meet experiments needs
  - OSG provides a framework for universities and laboratories to contribute resources to experiments remotely

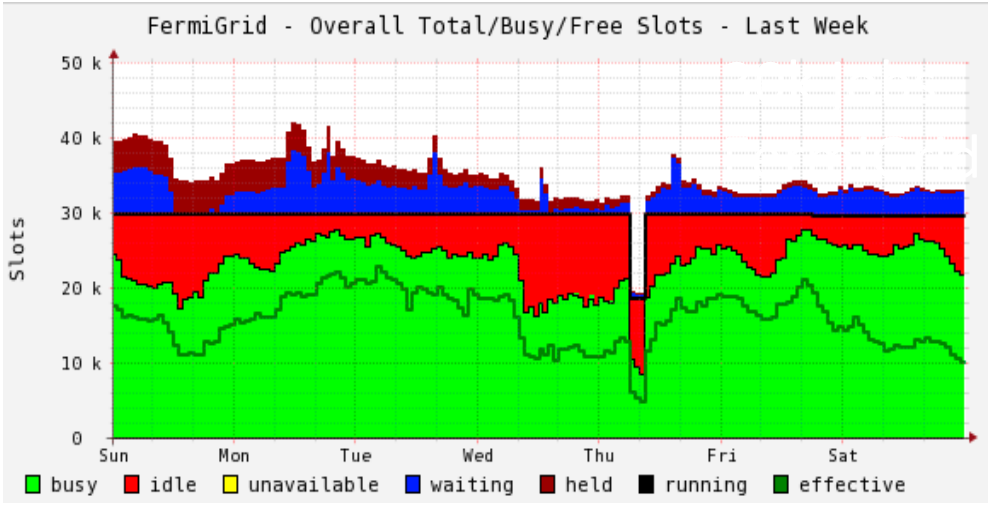


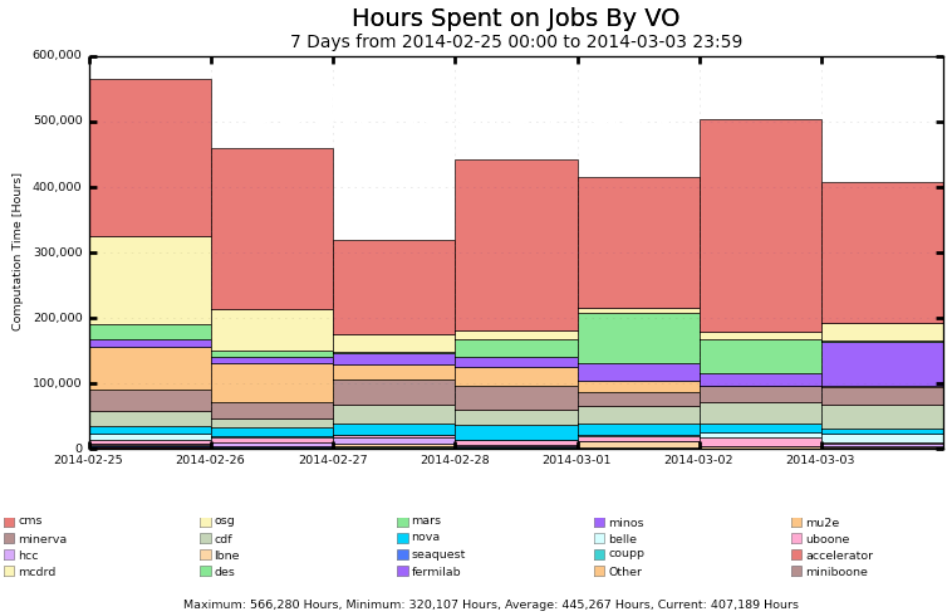

# Transition to OSG important for smaller experiments

- Most experiments are relatively small (~100 researchers)
  - Unlike Tevatron or LHC experiments, no resources to develop large scale computing infrastructure
- Implementing a unique computing infrastructure is too hard
- Fermilab SCD goal: integrate tools to meet the needs of smaller experiments, focusing on local resources, OSG, and Clouds
- Need to inform experiments and gather requirements from them
  - what resources are available?
  - how do you get a job onto the OSG? And onto Amazon?
  - what storage elements are available?
  - how does it all fit together?
- FIFE represents an integrated solution from a new architecture

# FabrIc for Frontier Experiments (FIFE)

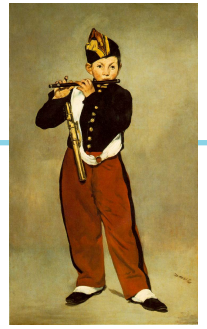

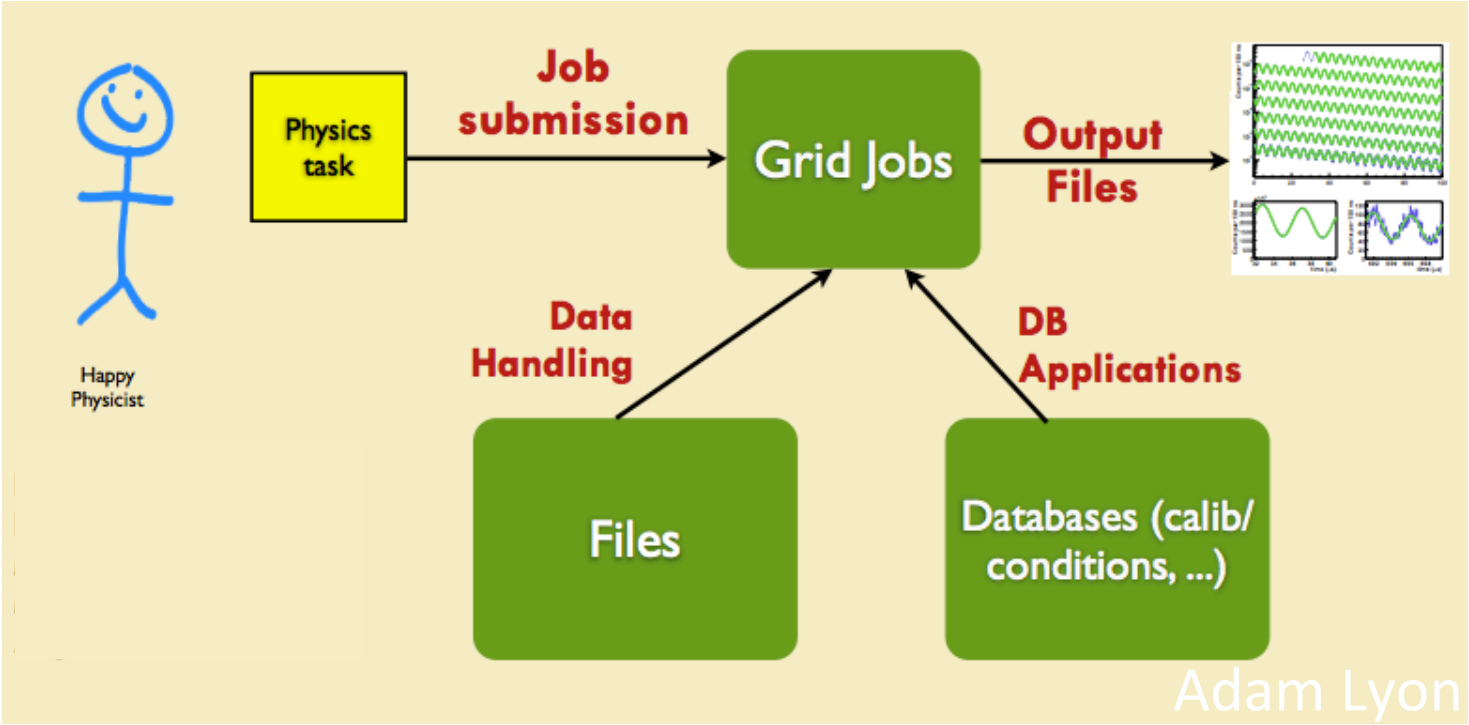


- provide resources while experiments focus on science
- integrate existing solutions into a consistent model

# FIFE Strategy

- support the experiments in their computing needs
- modular: experiments can take only what they need
- designed so that while underlying solution may change, interface will be consistent
- provide mechanism for experiments to incorporate their tools and solutions
- help experiments utilize computing beyond the Fermilab campus
- integrate new tools and resources from outside Fermilab and other communities as they develop

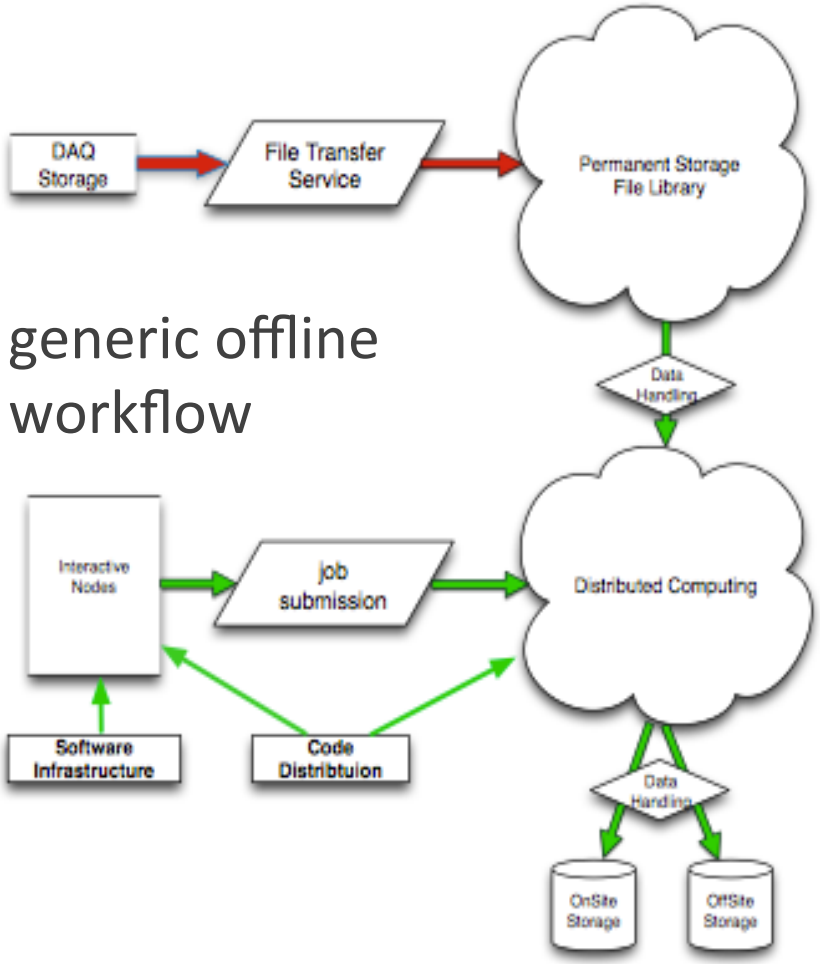

# Current FIFE implementations

- Software Framework – ART supports integrated services (Data Handling, Geant4, ROOT)
- build environment and distribution
  - git, svn, gmake, cmake and distribution with CERN Virtual Machine File System (CVMFS)
  - build process and machine in development
- data handling - access to file catalog, tape storage, cache disk, and file transfer
- job submission infrastructure – based on GlideinWMS pilot system
- database infrastructure and access
- shared software (LArSoft for Liquid Argon TPC reconstruction)
- additional infrastructure - authentication, electronic control room log book, new user accounts

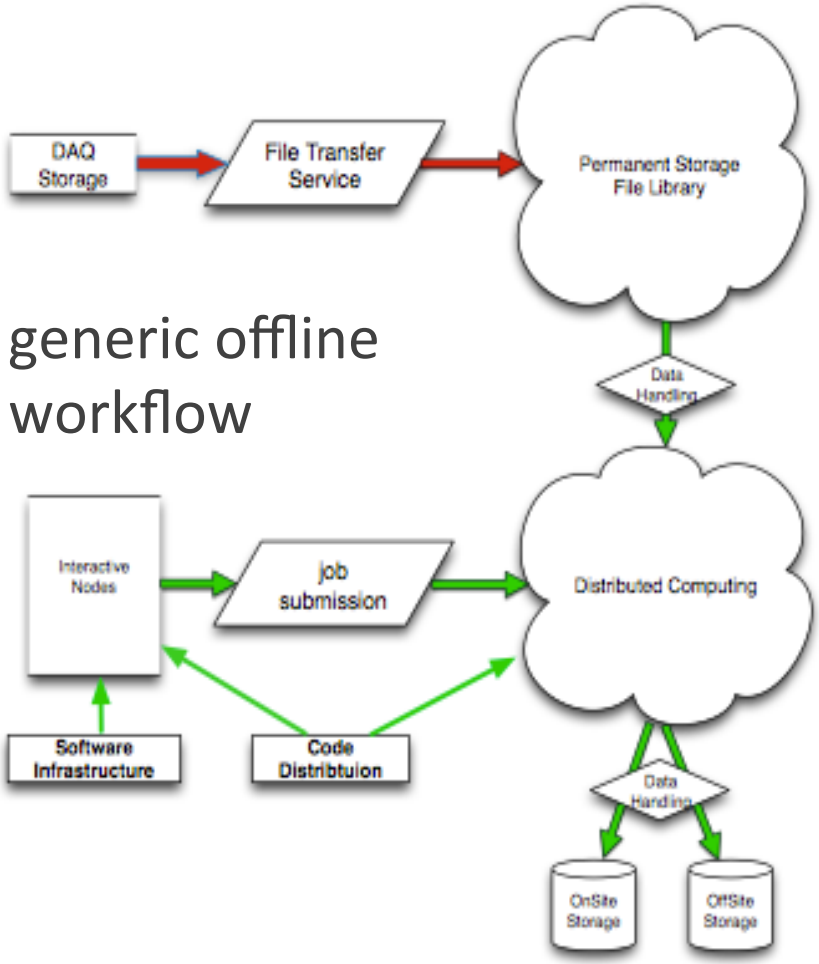

# Common Software Layers across the experiments

| | Minos | Minerva | NoVA | uBoone | LBNE | Gm2 | Mu2e | Lariat | 35T prototype |
|---|---|---|---|---|---|---|---|---|---|
| **Software Framework** | root/SRT | gaudi | art | art | art | art | art | art | art |
| **software processing** | | | | larsoft | larsoft | | | larsoft | larsoft |
| **Generators** | Genie | Genie | Genie, Flugg, Cry and Coriska | Genie | Genie | | | Genie, Cry | |
| **Data Management** | SAM | SAM | SAM | SAM | | SAM | None | SAM | |
| **Job Submission** | Jobsub/ glideinwms | Jobsub/ glideinwms | Jobsub/ glideinwms | Jobsub/ glideinwms | glideinwms | Jobsub/ glideinwms | Jobsub/ glideinwms | Jobsub/ glideinwms | |
| **Onsite Grid Enabled** | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | |
| **CVMFS** | Yes | | Yes | In Progress | In Progress | Yes | Planned | In Design | |
| **OSG Enabled** | In Progress | | Partially | In Progress | Planned | Planned | Planned | In Design | |
| **e-control room logbook** | Yes | Yes | Yes | Yes | | | Planned | Yes | |
| **Beam Condition DB** | Yes | Yes | Yes | Yes | | | | Yes | |
| **Scientific Document Management (docdb)** | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| **tape storage** | enstore | enstore | enstore | enstore | enstore | enstore | enstore | enstore | enstore |
| **data cache** | dcache | dcache | dcache | dcache | dcache | dcache | dcache | dcache | dcache |

# §1 – Integration success: early NOvA MonteCarlo on OSG

- Initial contained computing campaign with a single well-understood application
- 1,000,000 ev generated with ~10k jobs for 88,535 CPU h in 2 weeks of ops and 2 TB of data
- Run on OSG at SMU (dedicated), UNL, Uwisc, UC, UCSD and O(100) jobs at FermiCloud
- Operations consisted of 1 submission + 2 recoveries
- Spent about 10% more resources than expected due to preemption

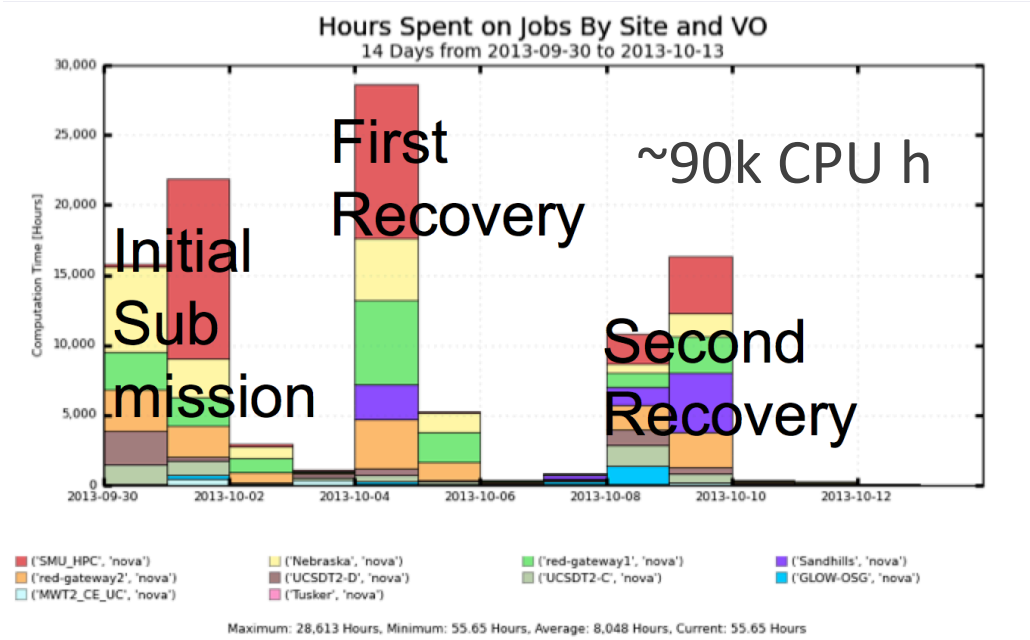


- using same submission command from user perspective, launched jobs "on-demand" at FermiCloud to provide proof-of-principle for Cloud bursting grid jobs

# §2 – Integration success: Ongoing NOvA computation on OSG …

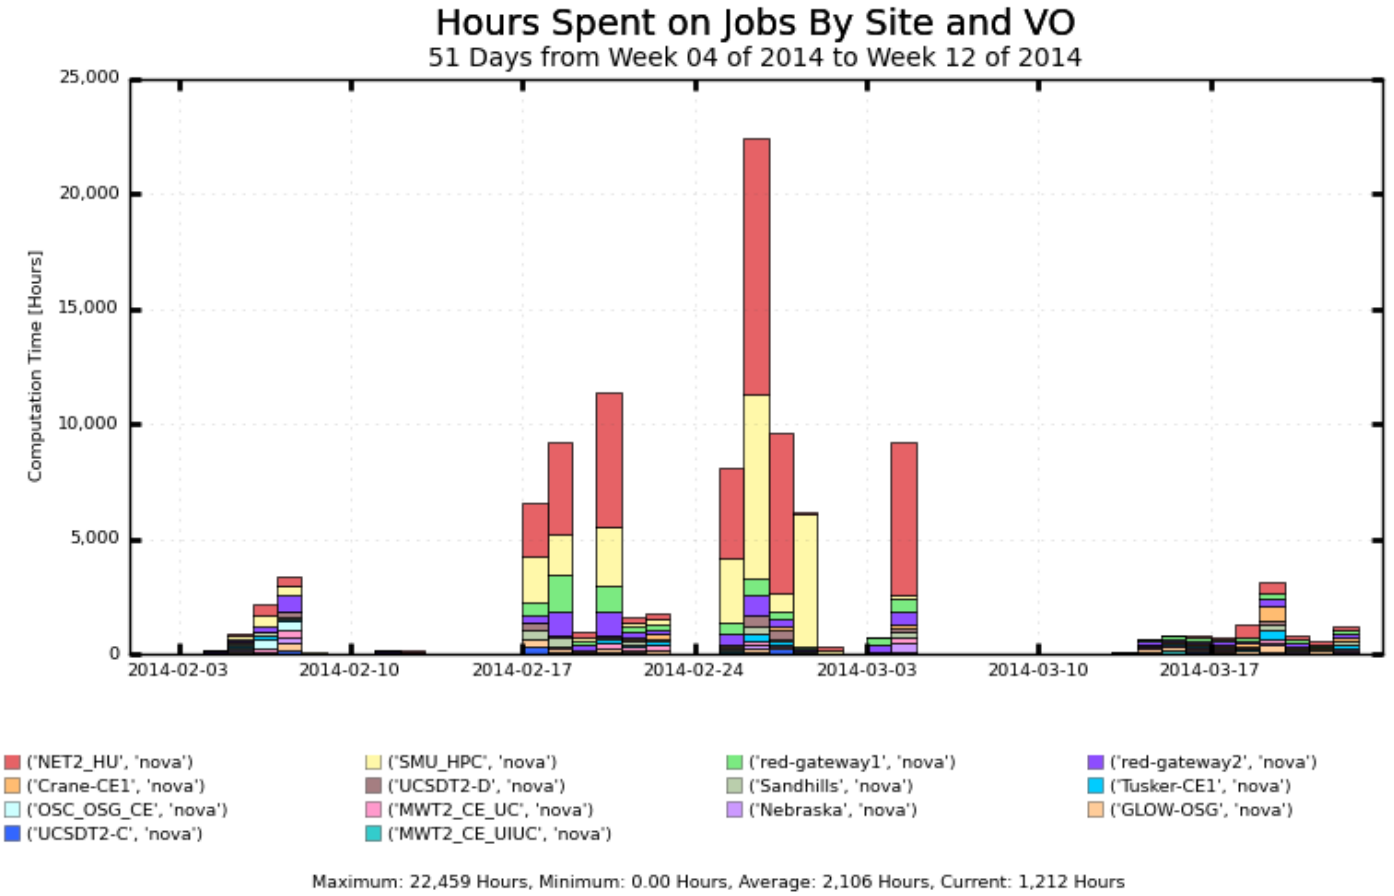


## …and Amazon Web Services

"On-demand Services for the Scientific Program at Fermilab"

Thu 3/27 16:25 – Rm1

# §3 – Integration success: early experience of MicroBoone on OSG

- LArSoft: package common to multiple Intensity Frontier experiments – Liquid Argon Simulation of Time-Projection Chambers
- Code with dependencies were made fully portable and deployed through CVMFS.
- Demonstrated the ability to run interactively at OSG sites (Fermilab and SLAC)
- Preparing for batch job submission to OSG sites – Nebraska, UCSD, SLAC

# Plans for improvements

- FIFE architecture has recently undergone a re-evaluation and re-architecture process
- job submission infrastructure modify to client-server model
- local storage element making transition to shared dCache pools
- data handling project continues to integrate new resources without any change in user interface
- starting new integration push for several experiments in both the Intensity Frontier and Cosmic Frontier

# Conclusions

- Fermilab physics program is extremely active at all three frontiers
- large scale computing needs for new “small” experiments encourages shared services designed to ease integration to distributed resource – OSG, Commercial, and Community Clouds
- FIFE project has developed a strategy and architecture to achieve this
- currently integrated OSG processing into NOvA and MicroBoone. Next, targeting DarkSide50, mu2e, and Muon g-2
- looking forward to bringing more resources and techniques to enable increased scientific output from frontier experiments