



---

Managed by Fermi Research Alliance, LLC for the U.S. Department of Energy Office of Science

---

## FIFE Overview

Ken Herner  
OSG All-Hands Meeting  
15 March 2016



*This picture is in the public domain*

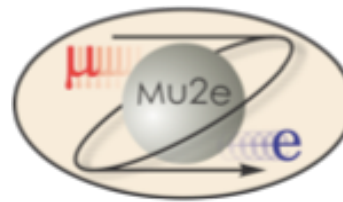
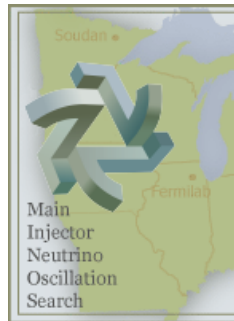
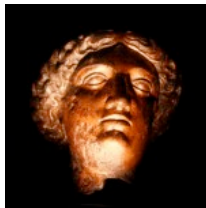
# Introduction to FIFE

---

- The **F**abric for **F**rontier **E**xperiments aims to
- Lead the development of the computing model for non-LHC experiments
- Provide a robust, common, modular set of tools for experiments, including
  - Job submission, monitoring, and management software
  - Data management and transfer tools
  - Database and conditions monitoring
  - Collaboration tools such as electronic logbooks, shift schedulers
- Work closely with experiment contacts during all phases of development and testing
- <https://web.fnal.gov/project/FIFE/SitePages/Home.aspx>

# A Wide Variety of Stakeholders

- At least one experiment in energy, intensity, and cosmic frontiers, studying all physics drivers from the P5 report, uses some or all of the FIFE tools
- Experiments range from those built in 1980s to fresh proposals



LArIAT



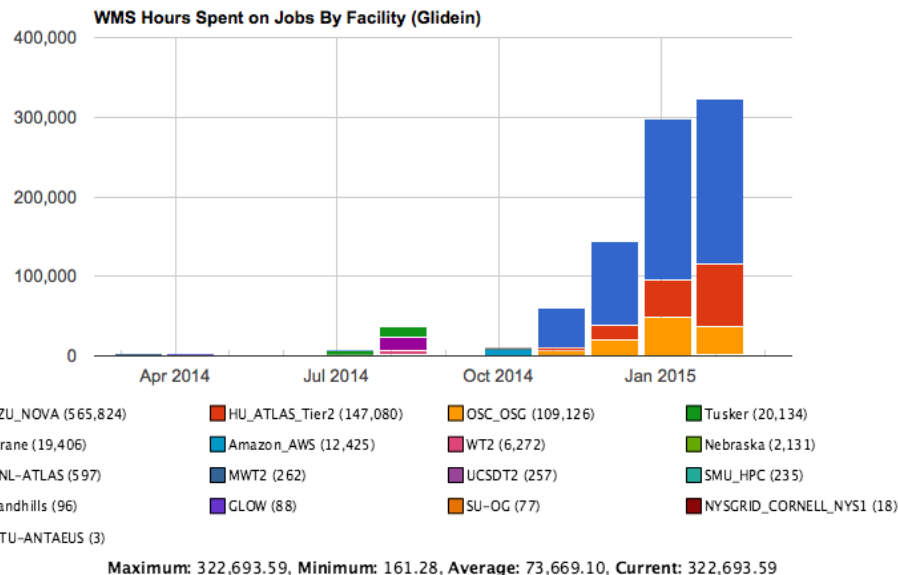
# Common problems, common solutions

---

- FIFE experiments on average are 1-2 orders of magnitude smaller than LHC experiments; often lack sufficient expertise or time to tackle all problems, e.g. software frameworks or job submission tools
  - Also much more common to be on multiple experiments in the neutrino world
- By bringing experiments under a common umbrella, can leverage each other's expertise and lessons learned
  - Greatly simplifies life for those on multiple experiments
- Common software frameworks are also available (ART, based on CMSSW) for most experiments

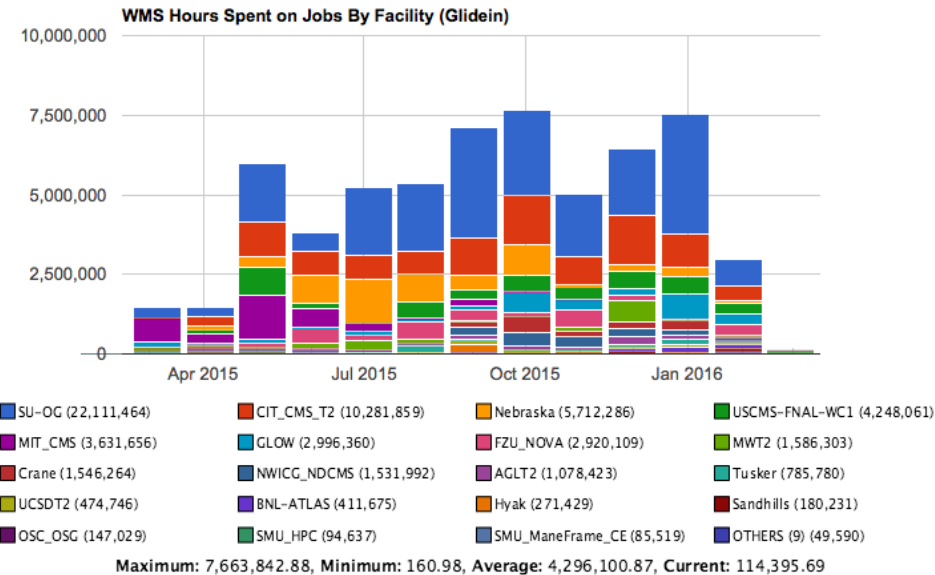
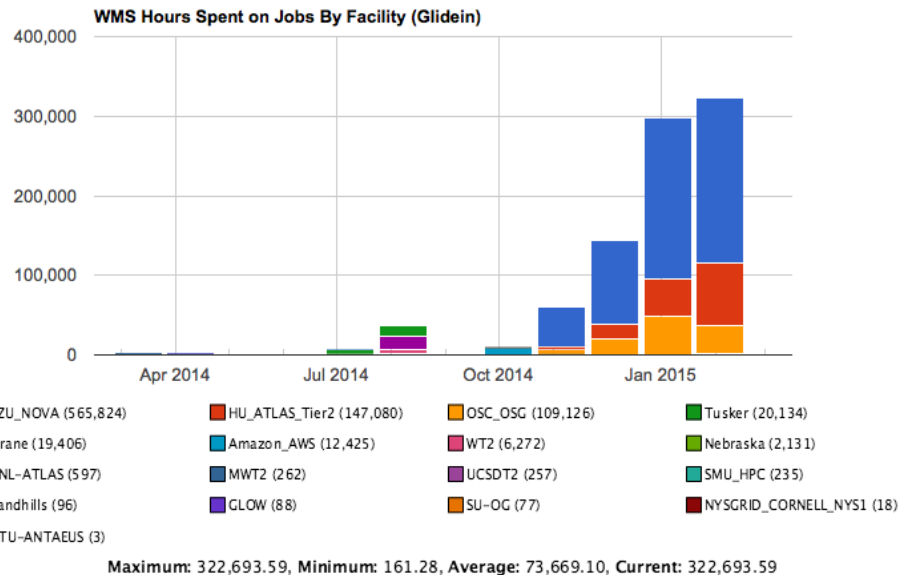
# A long way since the last AHM

- One year ago were doing hardly anything (about 2.5M hours on OSG in 2014)
- Large auxiliary files (the “flux files” for the GENIE simulation) were creating a lot of pressure on data transfer systems and/or CVMFS Stratum 1s



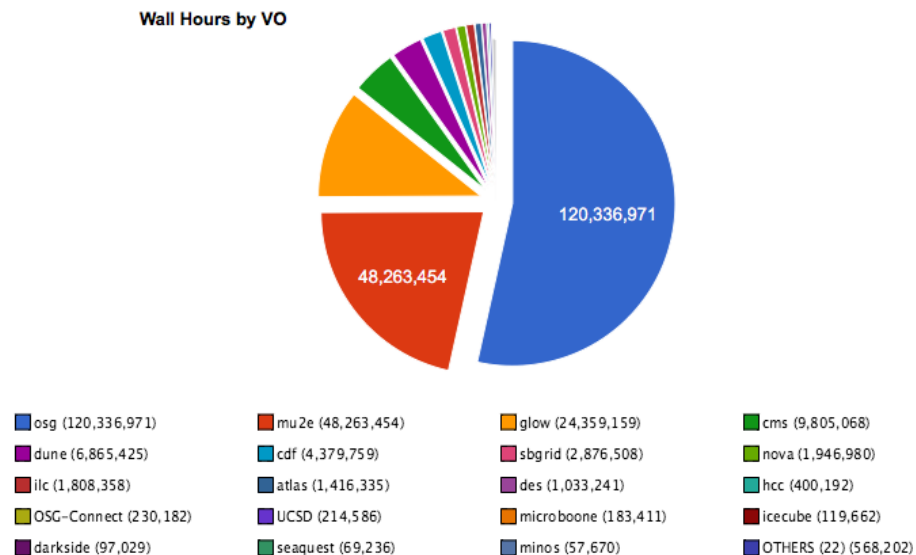
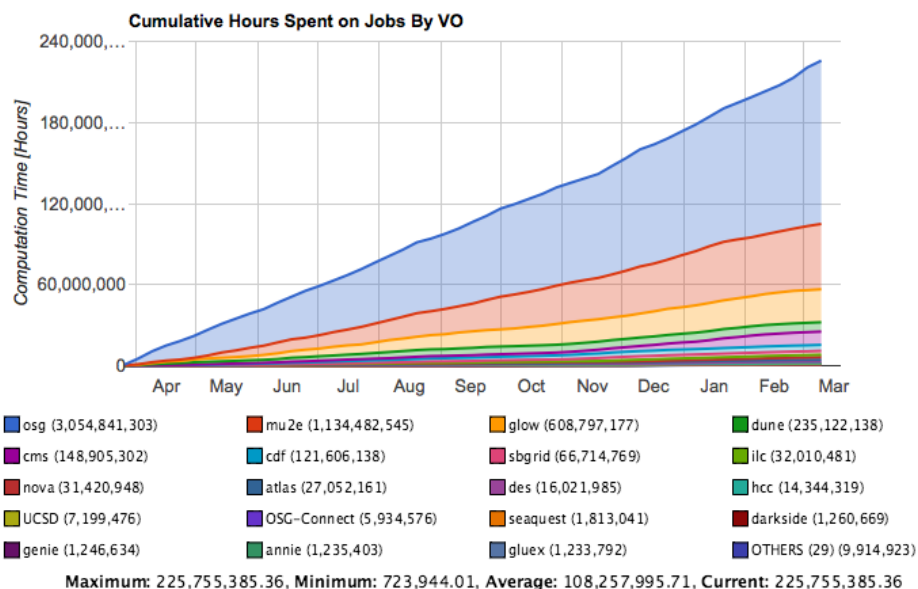
# A long way since the last AHM

- Average hours per week has increased by a factor of 60
  - Seeing upwards of 40% of all FIFE jobs on non-FNAL OSG sites
  - Success rates now typically  $\approx 99\%$
- Playing a role in testing new technologies like StashCache



# Numbers since March 2015 (Last AHM)

- FIFE experiments have used about 58M opportunistic hours since 1 April 2015-- second only to OSG VO
  - About 25% of all opportunistic hours (defined has hours run on facility not owned by the VO)



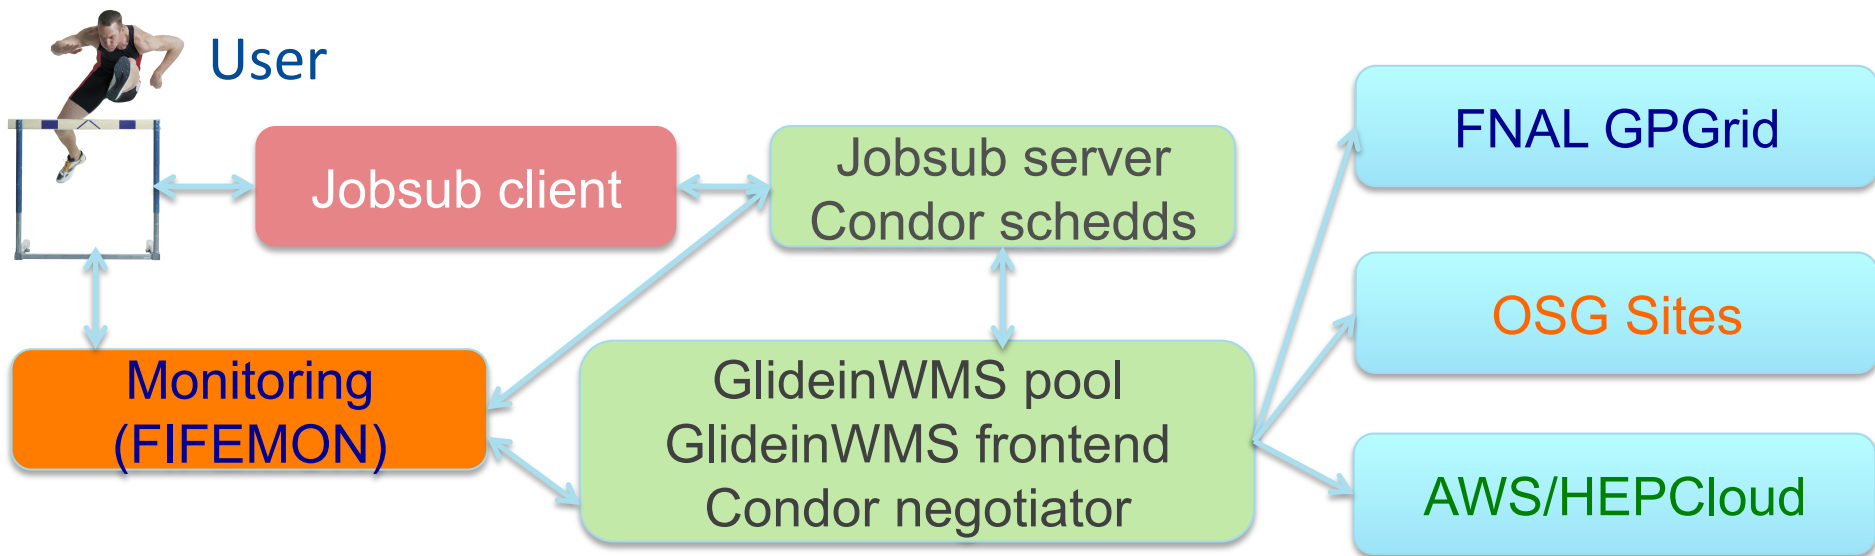
# OSG's place in the FIFE model

---

- As more experiments come online soon (Microboone, Lariat Run II, DUNE prototypes to name a few) computing demand rapidly increasing
  - FIFE experiments expect to use about 150M CPU hours this FY, 190M next year, **220M** after that (plus add g-2 datataking)
- **Fermilab resources, while vast, cannot meet full demand**
  - Either scale back scope of work, or find more resources
- OSG is **critical** to compute strategy in coming years

# Job Submission and management architecture

- Common infrastructure is the **fifebatch** system: one GlideInWMS pool, 2 schedds, frontend, collectors, etc.
- Users interface with system via “jobsub”: middleware that provides a common tool across all experiments; shields user from intricacies of Condor
  - Simple matter of a command-line option to steer jobs to different sites
- Common monitoring provided by FIFEMON tools (recently updated; see Kevin Retzke’s talk)
  - Now also helps users to understand why jobs aren’t running



## Somewhat Typical Usage Model Promoted by FIFE

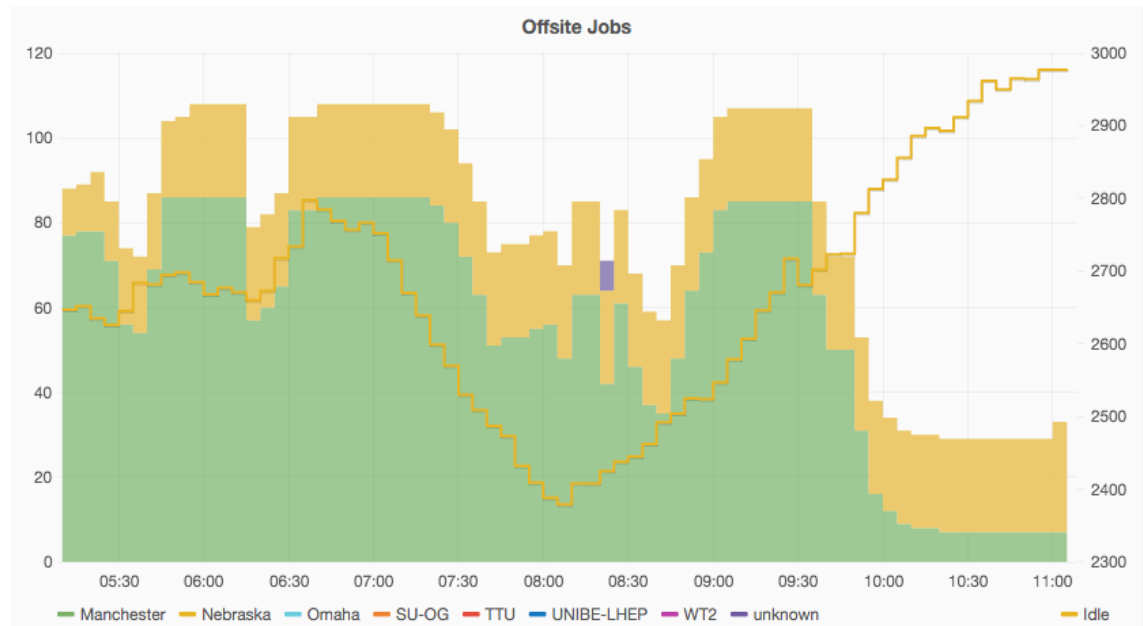
---

- Most experiments prefer to run simulation on OSG and data reconstruction at FNAL: lower payloads and less disruption if preempted
- Push for workflows to be site-agnostic whenever possible
- Code distribution via CVMFS; file I/O through Fermilab dCache or dedicated site SEs (mostly for NOvA)
- For newer experiments, build in expectation to run on OSG from the beginning.
- The Fermilab cluster has additional NFS mounts (going away soon) that experiments have been (overly) relied upon. Difficult to wean the older experiments, but some have been able to make it.

## New International Sites

- Previously had allocation for NOvA at FZU in Prague
- Have since added Manchester and Bern for Microboone (only) in recent weeks
  - Alessandra Forti very helpful at Manchester; Gianfranco Sciacca at Bern
- Setup in both cases was about one week in both cases

Made so smooth by GWMS and **OSG's ongoing work** to make variety of different sites compatible- you have “flattened the globe”



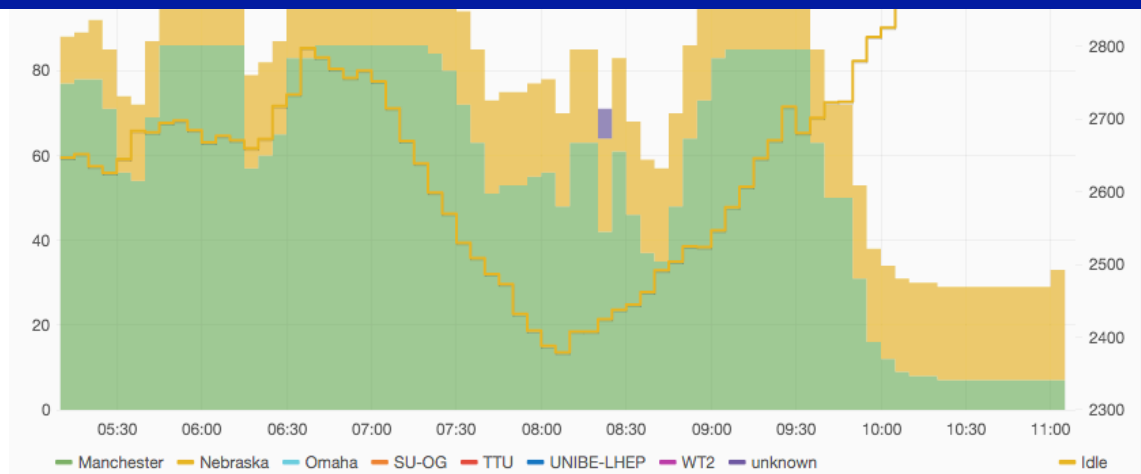
## New International Sites

- Previously had allocation for NOvA at FZU in Prague
- Have since added Manchester and Bern for Microboone (only) in recent weeks
  - Alessandra Forti very helpful at Manchester; Gianfranco Sciacca at Bern

How “smooth” was it really?

The very first test job at Manchester worked

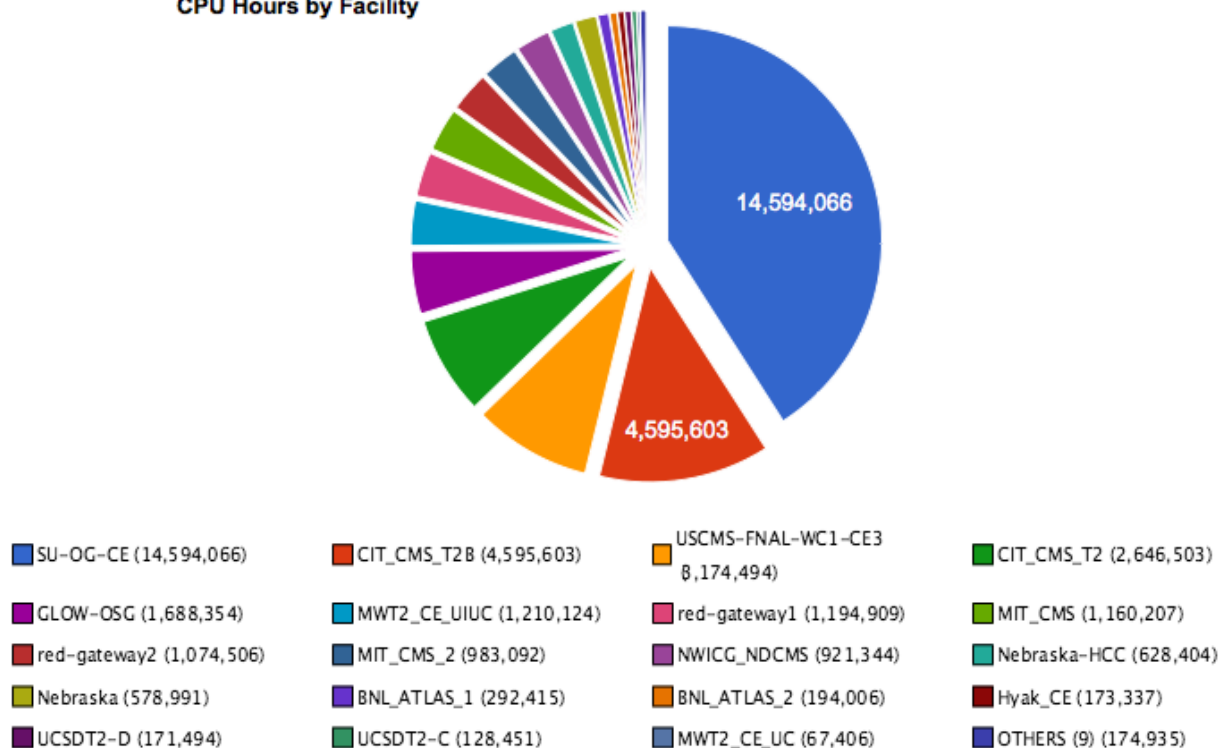
Ma  
GWMS and **OSG’s ongoing work** to make variety of different sites compatible- you have “flattened the globe”



# The Pacesetter: Mu2e

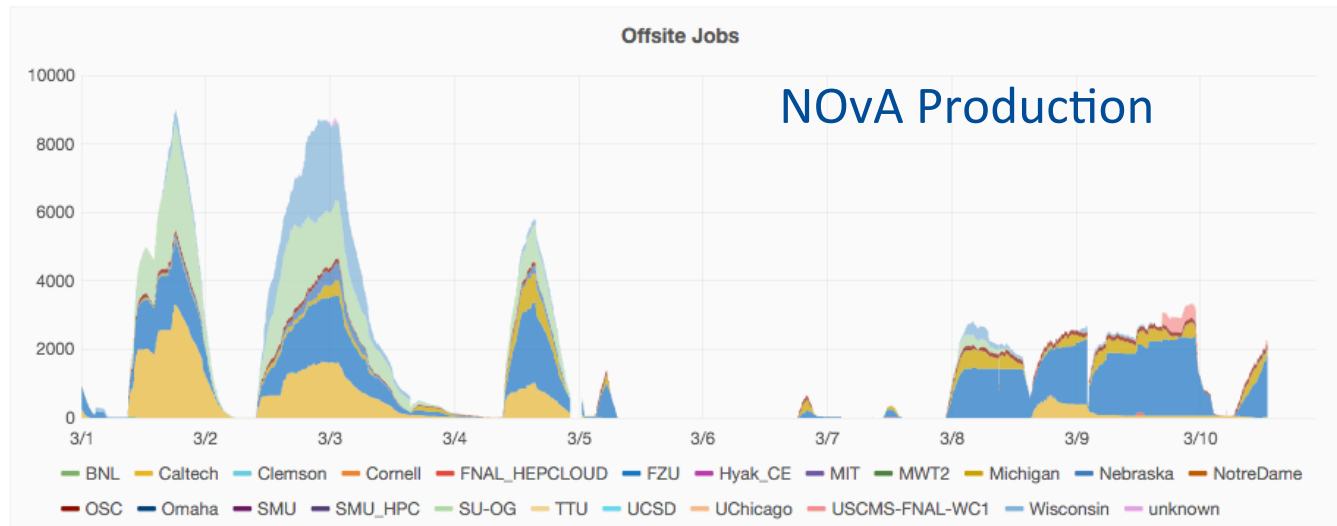
- Mu2e: new experiment set to look for lepton flavor violation
  - Ray Culbertson will go into more detail tomorrow
- Has become the standard for other FIFE Experiments
- Over 60M hours in past year, regularly a 99.9% success rate

CPU Hours by Facility



# NOvA: Current Flagship Neutrino Experiment

- NOvA was one of the first experiments to go offsite; more complex workflows due to being a running experiment
- Making increasing use of OSG resources in recent weeks; necessary for analysis campaign aimed at Neutrino 2016
- Have been somewhat slower than Mu2e going to OSG due to memory requirements on jobs (main FW executable often needs 2.5 - 3 GB depending on plugins used) and library dependencies
  - Made good progress on dependency issue recently



# Recent challenges for FIFE Experiments

---

- Code distribution via CVMFS generally works very well
  - Differences in installed software on worker nodes causes occasional problems (mostly X11 libs, i.e. things users assume are always installed)
  - Difficult to get non-power users to take a sparser environment than on FNAL machines into account
- Memory requirements
  - Younger experiments (particularly LAr TPC expts.) have workflows requiring  $> 2$  GB memory per job. Somewhat limited resources available going above 2 GB/1 core.
    - My opinion: this tends to scare some smaller experiments away. Not sure a lot can be done in the near term though...
- Large auxiliary files
  - StashCache looking promising
- Understanding preemption policies; getting clear signals of preemption into jobs

# Summary and Future Directions

---

- FIFE Experiments have made significant in the past year
  - OSG was completely foreign to almost all experiments
  - Making up a significant portion of OSG work
- Still much work to do
  - Work on building more robust software frameworks, bring along needed libraries
  - Reduce overall memory footprint (multithreading?)
  - Continue to expand site list
- Many thanks to effort of OSG staff in setting up sites and responding quickly to (my numerous) GOC tickets
- Looking forward to another productive year

# Backup

---

