



# The FIFE Project: Computing for Experiments

Ken Herner for the FIFE Project  
CHEP 2016  
11 October 2016




## 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; standing meetings w/developers
- <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




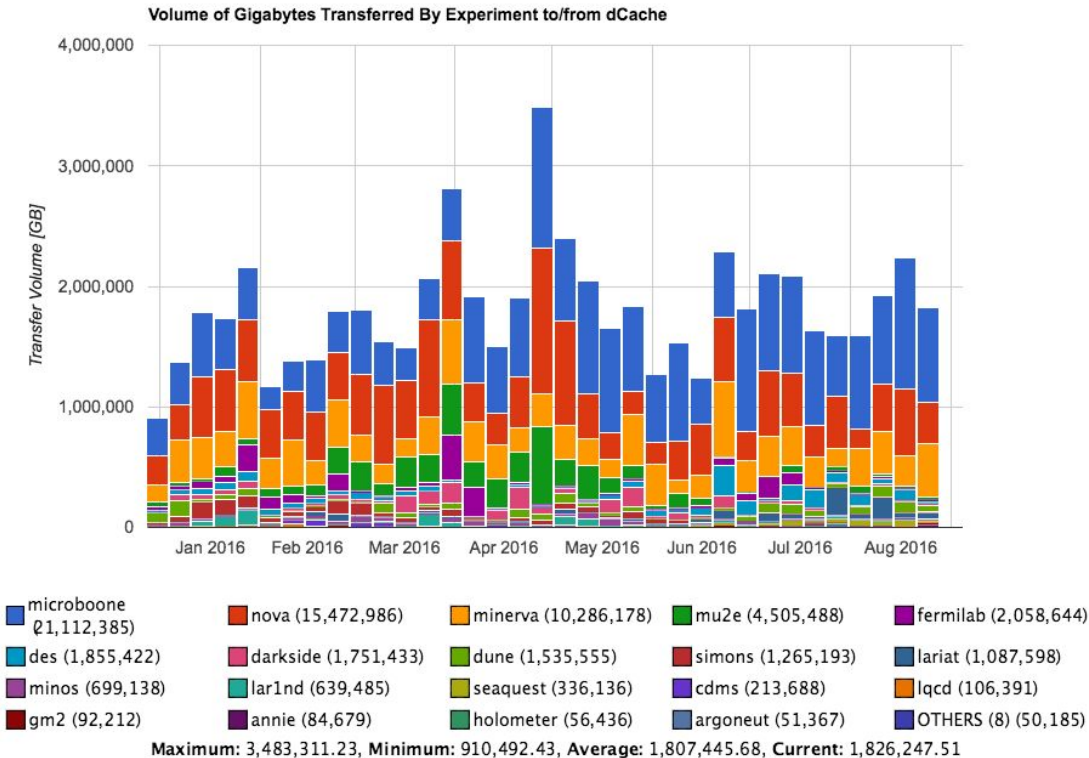
## Common problems, common solutions

- 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 modular software framework is also available (ART, based on CMSSW) for most experiments
- Common problem of large auxiliary files needed by many jobs; trying out StashCache (see B. Bockelman's talk) as a solution. Testing and providing feedback to Open Science Grid and developers



# Data and Job volumes

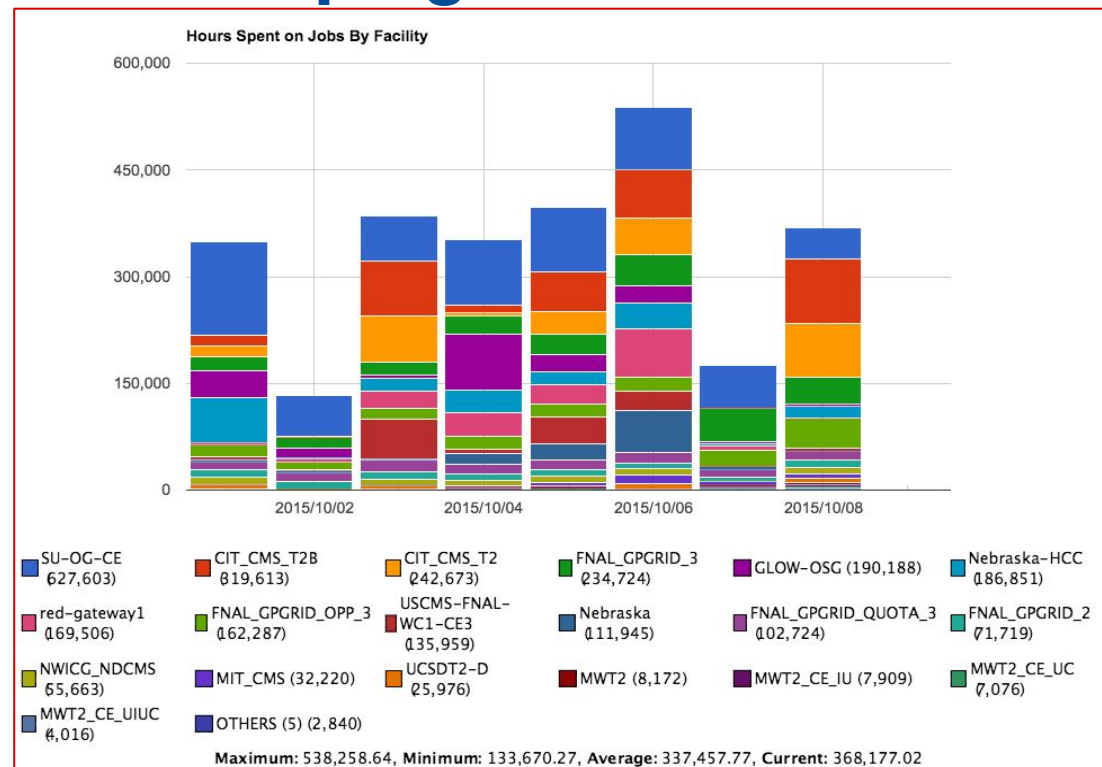
- Nearly 5 PB new data catalogued over past 6 months
- Average throughput of 1.8 PB/wk through FNAL dCache
- Typically 15K simultaneous jobs running; peak over 30K
- **Combined numbers approaching scale of LHC experiments**



# Mu2e Beam Simulations Campaign

- Almost no input files
- Heavy CPU usage
- <100 MB output per job
- Ran > 20M CPU-hours in under 5 months

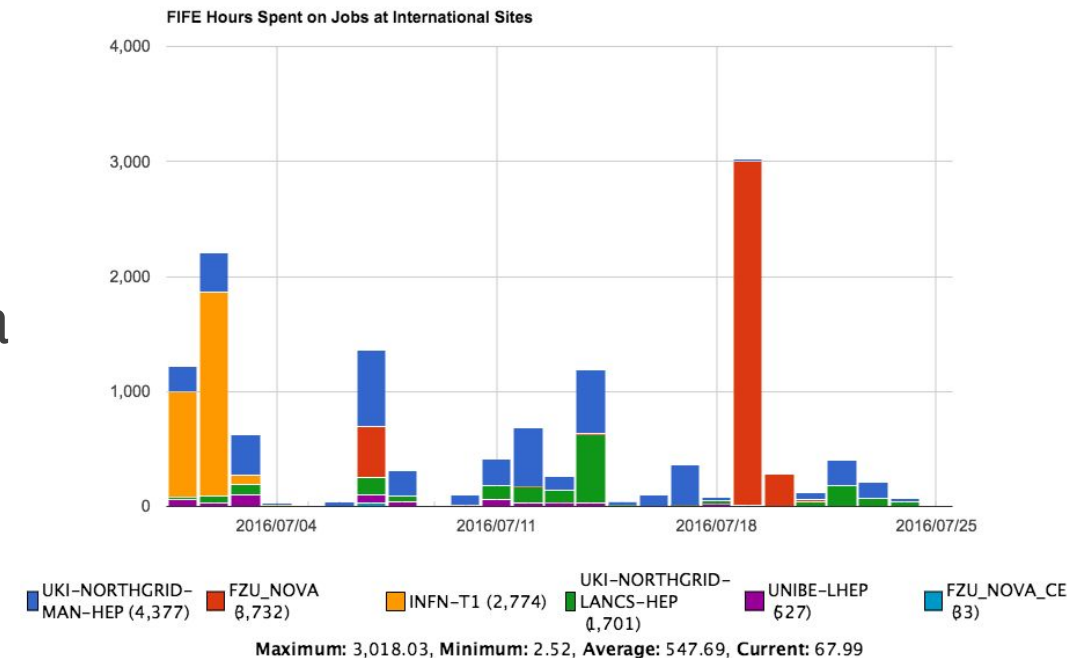
- Avg 8000 simultaneous jobs across > 15 remote sites



- Usage as high as 20,000 simultaneous jobs and 500,000 CPU hours in one day – peaked usage 1<sup>st</sup> wk Oct 2015
- *Achieved stretch goal* for processing 24 times live-time data for 3 most important backgrounds
- **Total cost to Mu2e for these resources: \$0**

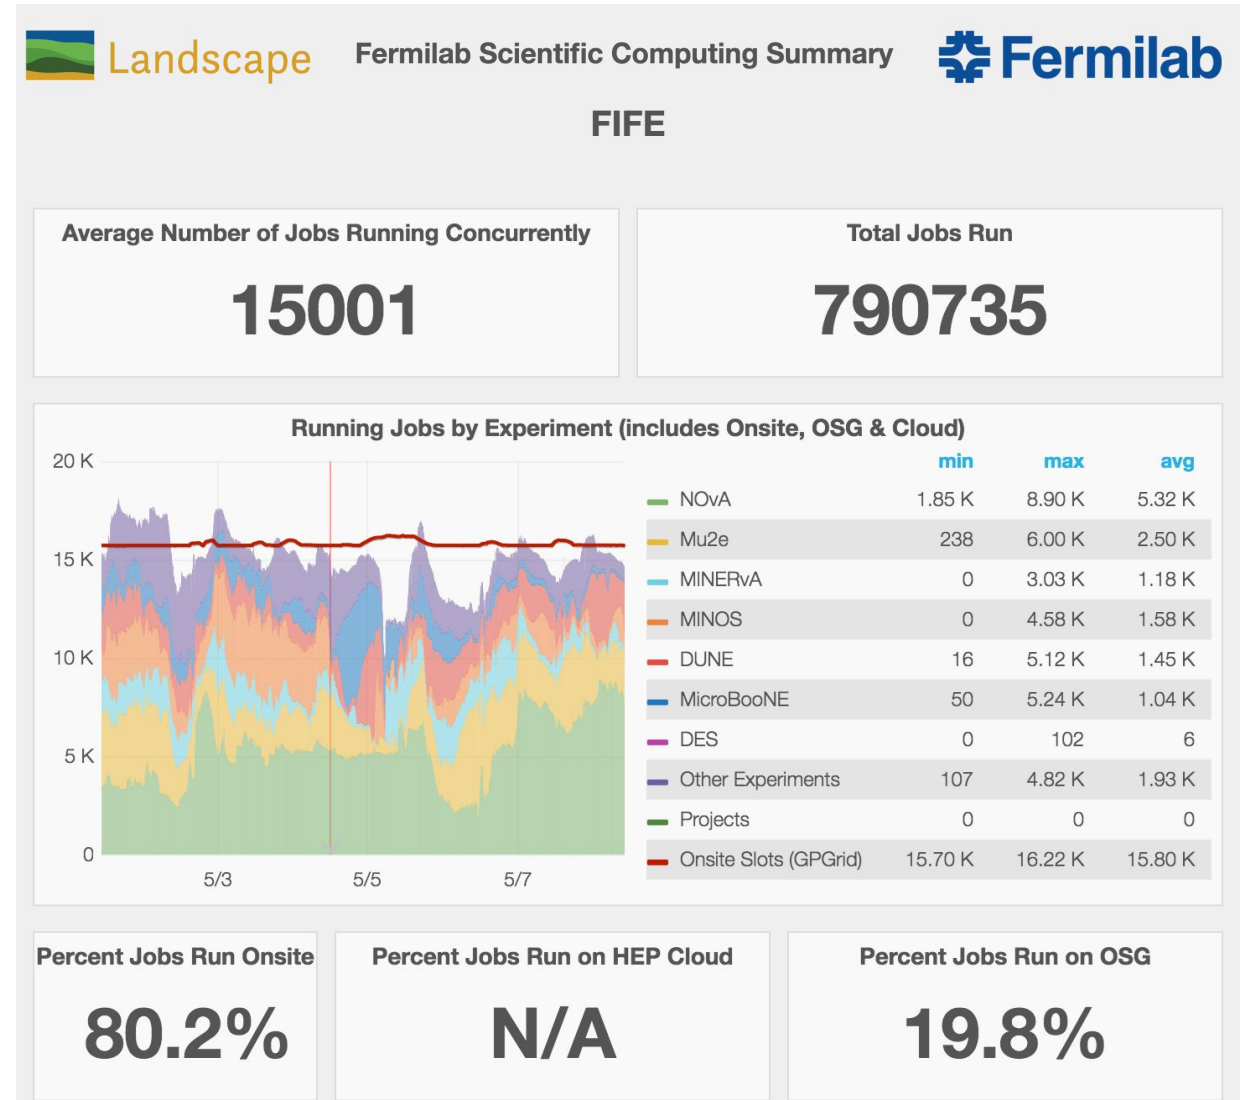
## New International Sites for running jobs

- International collaborators can often bring additional computing resources to bear, but may have specific configurations. users want to be able to seamless run at all sites with unified submission command
  - First International location was for NOvA at FZU in Prague; have added Manchester, Lancaster, and Bern for Microboone
- Following OSG prescription makes it easy to have sites around the globe communicate with a common interface, with a variety of job management systems underneath (ARC, CREAM, HTCondor, PBS,...)
- Integration times as short as 1-2 weeks



# FIFE Monitoring of resource utilization


- Extremely important to understand performance of system
- Critical for responding to downtimes and identifying inefficiencies
- Focused on improving the real time monitoring of distributed jobs, services, and user experience
- Completely new project built on open source tools (ELK stack, Grafana for visualization)
  - Access to historical information using same toolset





# From Gratia to GRACC

- Next-generation OSG accounting service, based on open-source technology (same used for new FIFEMON tools)
- Provides access to historical information that includes pilot and payload jobs data, and transfers since 2004.



# Full workflow management

- Now combining job submission, data management, databases, and monitoring tools into complete workflow management system
  - Production Operations Management Service (POMS)
- Can specify user-designed “campaigns” via GUI describing job dependencies, automatic resubmission of failed jobs, complete monitoring and progress tracking in DB
  - Visible in standard job monitoring tools
- Usable for production-level running and user analysis
- REST API for data I/O
- Command line tools for needed operations
- Supports POMS launching jobs, or experimenters launching jobs and using POMS only for tracking

**POMS**

- Dashboard
- Calendar
- Requests
- DB Admin
  - Experiments
  - Users
  - Raw Tables
- Campaigns
  - Active
  - InActive
  - Compose Campaign
  - Compose Definition
  - Compose Launch Template
- Jobs
  - All (Last Day)
  - Failed\_By Exit Code and Campaign Name (Last Day)
  - Failed\_By Exit Code and Campaign Name (Last Week)
  - Failed\_By Exit Code and Node Name (Last Day)
  - Failed\_By Exit Code and CPU Type (Last Day)

Home / Jobs by exit code

## Jobs by user\_exe\_exit\_code,name,experiment ?

2016-09-27 17:53 to 2016-09-28 17:53

<previous 1 days | next 1 days >



user\_exe\_exit\_code x name x experiment x Submit Query

user_exe_exit_code	name	experiment	count
1	cry_generation	nova	4
134	cry_generation	nova	2

Kenneth Herner

Home / Campaigns

## Campaigns ?



### Active Campaigns

<previous 1 days | next 1 days > 2016-09-27 17:54 to 2016-09-28 17:54

Campaign ?		Active Jobs			Jobs in 1.000000 day ending 2016-09-28 17:54		
ExperimentName		Idle ?	Running ?	Held ?	Completed ?	Located ?	Removed ?
lariat	Lariat Raw2Digits	0	0	0	0	0	0
nova	Nova raw2root keepup ND	0	11	0	0	0	0
nova	Nova raw2root keepup FD	0	162	0	0	0	0
nova	NOvA Reco Keepup FD	0	0	0	0	0	0
nova	NOvA Reco Keepup ND	0	7	0	0	0	0
nova	prod_daq_R16-Q2-11-geniepreview.a_nd_genie_fhc_nonswap_DefaultPlusMECWithNC	0	0	0	0	0	0

# Improving Productivity with Continuous Integration

- Have built up a Jenkins-based Continuous Integration system designed for both common software infrastructure (e.g. Art) and experiment-specific software, full web UI
- In addition to software builds, can also perform physics validation tests of new code (run specific datasets as grid jobs and compare to reference plots)
- Supporting SL6/7, working on OSX and Ubuntu support, experiments free to choose

Multiplatform continuous integration for Art

Select builds:	Build	Start Time	Platform	Build Type	checkout	pullProducts	build	unit_test	install	Progress Legend
From build: <input type="text" value="Buildname"/> <input type="button" value="GO"/>										Running
# of builds: <input type="text" value="number"/>										Pending
Select platforms:										Succeeded
<input type="checkbox"/> Darwin 13.4.0										Failed
<input type="checkbox"/> Darwin 14.5.0										Skipped
<input type="checkbox"/> Linux 2.6.32-573.26.1.el6.x86_64										
<input type="checkbox"/> Linux 3.10.0-327.18.2.el7.x86_64										
Select build types:										
<input type="checkbox"/> d13-s35-e10:debug										
<input type="checkbox"/> d13-s35-e10:nu:debug										
<input type="checkbox"/> d13-s35-e10:nu:prof										
<input type="checkbox"/> d13-s35-e10:prof										
<input type="checkbox"/> d14-s35-e10:debug										

Build	Start Time	Platform	Build Type	checkout	pullProducts	build	unit_test	install
art_ci/195 (Art)	2016-08-24 18:20:30.084410	Darwin 14.5.0	d14-s35-e10:nu:debug	✓	⚠	⚠	⚠	⚠
	2016-08-24 18:19:53.538640	Darwin 14.5.0	d14-s35-e10:nu:prof	✓	⚠	⚠	⚠	⚠
	2016-08-24 18:19:04.143585	Linux 2.6.32-573.26.1.el6.x86_64	slf6-s35-e10:nu:prof	✓	✓	⚠	⚠	⚠
	2016-08-24 18:19:04.601685	Linux 2.6.32-573.26.1.el6.x86_64	slf6-s35-e10:nu:debug	✓	✓	⚠	⚠	⚠
art_ci/194 (Art)	2016-08-24 18:09:00.921544	Linux 2.6.32-573.26.1.el6.x86_64	slf6-s35-e10:debug	✓	✓	✓	✓	✓
	2016-08-24 18:08:49.644149	Linux 2.6.32-573.26.1.el6.x86_64	slf6-s35-e10:prof	✓	✓	✓	✓	✓
	2016-08-24 18:08:59.198921	Darwin 14.5.0	d14-s35-e10:prof	✓	✓	⚠	⚠	⚠
	2016-08-24 18:08:59.098807	Darwin 14.5.0	d14-s35-e10:debug	✓	✓	⚠	⚠	⚠
	2016-08-24 18:08:45.978269	Linux 3.10.0-327.18.2.el7.x86_64	slf7-s35-e10:prof	✓	✓	✓	⚠	⚠
	2016-08-24 18:08:45.881568	Linux 3.10.0-327.18.2.el7.x86_64	slf7-s35-e10:debug	✓	✓	✓	✓	⚠

NOvA experiment's CI tests

2 tests with Warning:  
They are successful  
BUT the Data Product are different  
from the reference files

## Phase: ci\_tests

 [ci\\_tests](#)  
Started 2016-08-24 18:04:34.519615

- [ci\\_calib fd regression test novasoft](#)
- [ci\\_calib nd regression test novasoft](#)
- [ci\\_raw2root fd t00 regression test novasoft](#)
- [ci\\_raw2root fd t02 regression test novasoft](#)
- [ci\\_raw2root nd t00 regression test novasoft](#)
- [ci\\_raw2root nd t02 regression test novasoft](#)
- [ci\\_reco fd regression test novasoft](#)
- [ci\\_reco nd regression test novasoft](#)

Finished 2016-08-24 18:08:58.841689  
exit code: 3

# FIFE Plans for the future

- Increase production teams' productivity with workflow management tools
- Help define the overall computing model of the future
  - Seamlessly integrating dedicated, opportunistic, and commercial computing resources via HEPCloud
  - Increase access to HPC resources for job submission
    - Already doing this by enabling access to allocation-based resources through existing GlideinWMS system. **MINOS+ now able to run jobs on Stampede at TACC via XSEDE allocation**
  - Usher in easy access to GPU resources for those experiments interested
- **Lower barriers to accessing computing elements around the world in multiple architectures**
  - DCAFI Project (see D. Dykstra's poster) is a key piece of this effort
  - Help to connect experimenters and computing professionals to drive experiment SW to increased multithreading and smaller memory per core footprints
- Augment data management tools (SAM) to also allow a "jobs to the data" model
- Scale up and improve UI to existing services




# Summary

- FIFE providing access to world class computing to help accomplish world class science
  - FIFE Project aims to provide common, modular tools useful for the full range of HEP computing tasks
  - Stakeholders in all areas of HEP, wide range of maturity in experiments
  - Experiments, datasets, and tools are not limited to Fermilab
- **Overall scale now approaching LHC experiments**; plan to heavily leverage opportunistic resources
- Plan to provide full WMS functionality not limited to Fermilab resources
- Work hand-in-hand with experiments and service providers to move into new computing models via HEPCloud (see talk by B. Holzman)

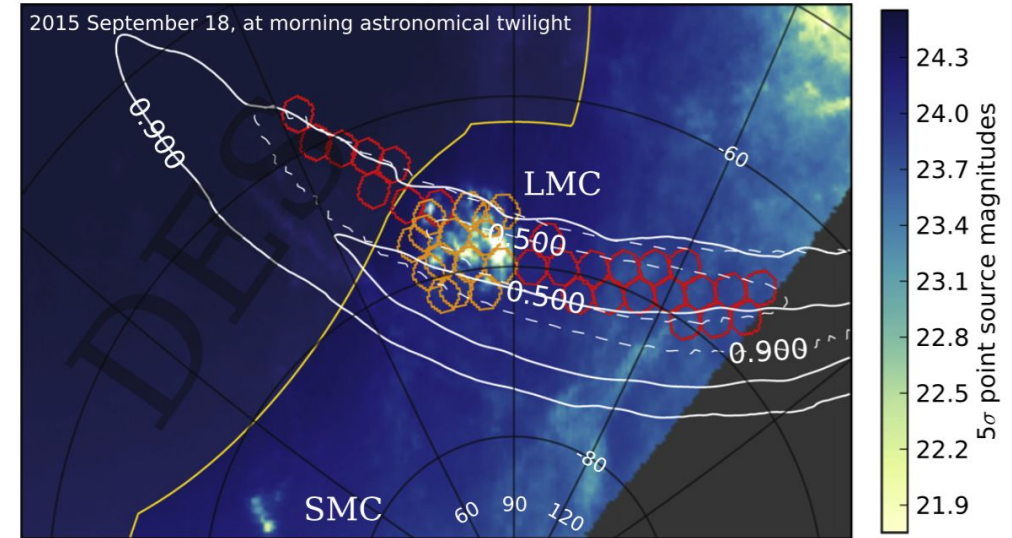


# Backup


# Selected results using the FIFE Tools



NOvA: excludes maximal mixing at 2.5 s.d.




Dark Energy Survey: Optical follow-up of gravitational wave triggers





Microboone: first results

MINOS+: limits on LEDs



# NOvA – full integration of FIFE Services


- File Transfer Service stored over 6.5 PB of NOvA data in dCache and Enstore
- SAM Catalog contains more than 41 million files
- Helped develop SAM4Users as lightweight catalog



- Jan 2016 - NOvA published first papers on oscillation measurements
- avg 12K CPU hours/day on remote resources
- > 500 CPU cores opportunistic
- FIFE group enabled access to remote resources and helped configure software stack to operate on remote sites
- Identified inefficient workflows and helped analyzers optimize




# Overview of Experiment Computing Operations



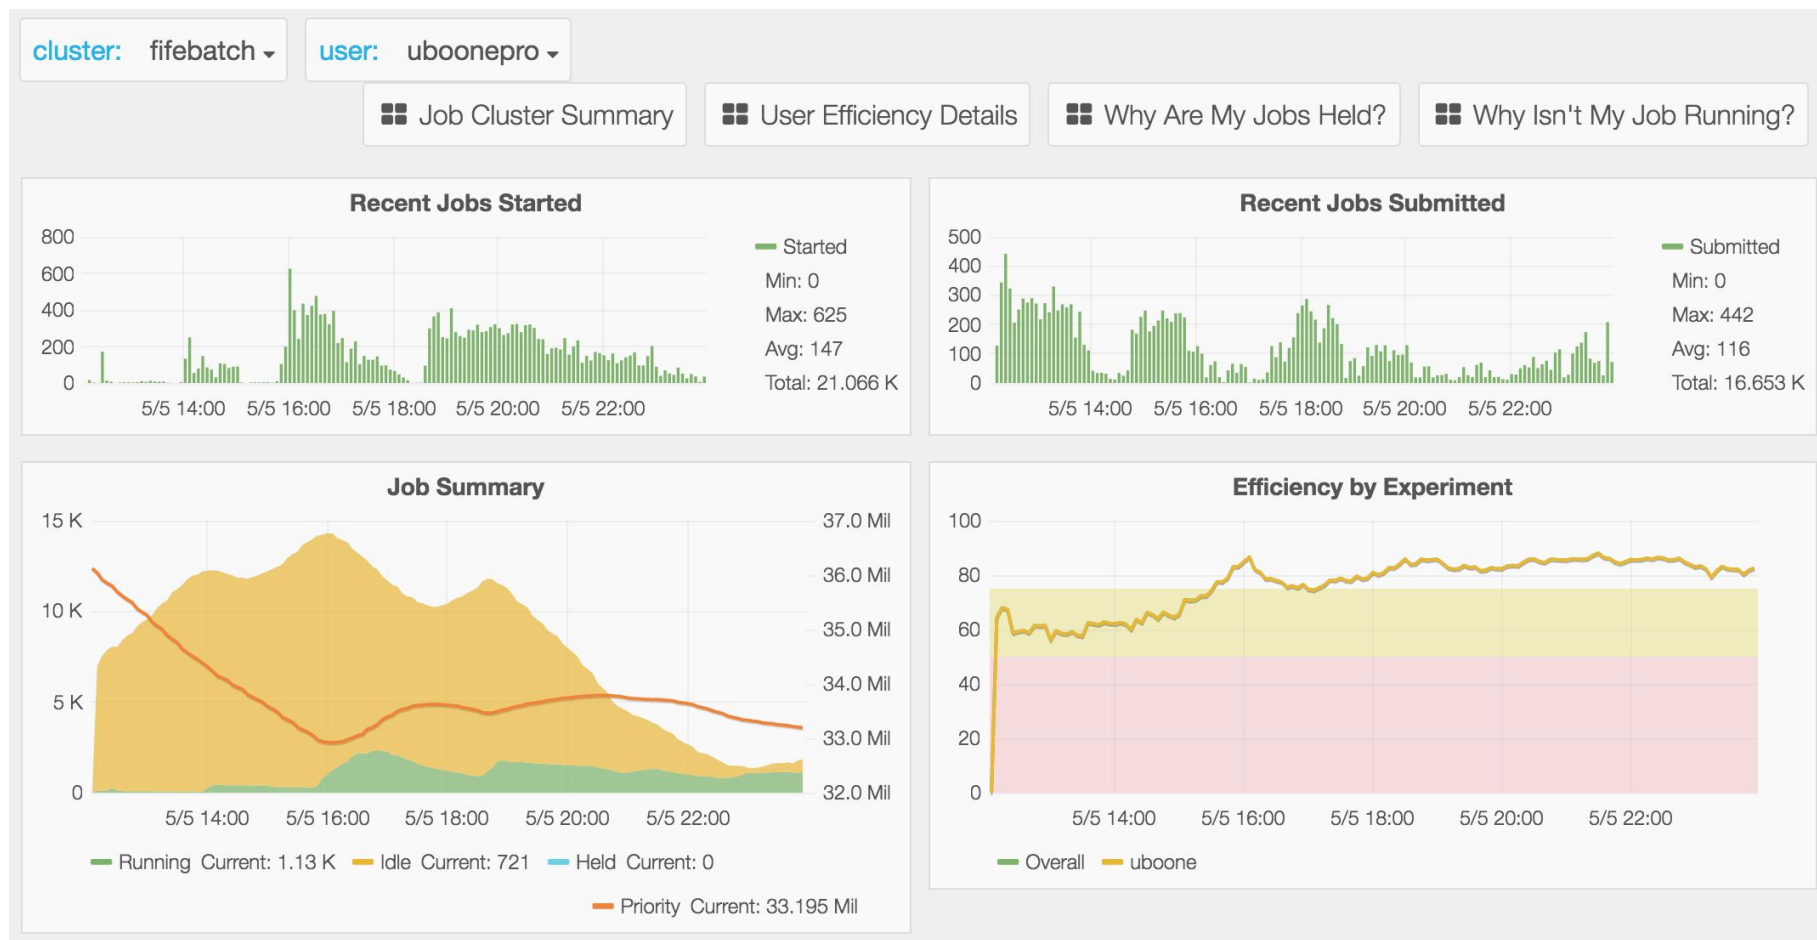
quickly understand the usage pattern for the last week of each experiment and collectively get a picture of distributed computing operations for the FIFE experiments

# Detailed profiling of experiment operations



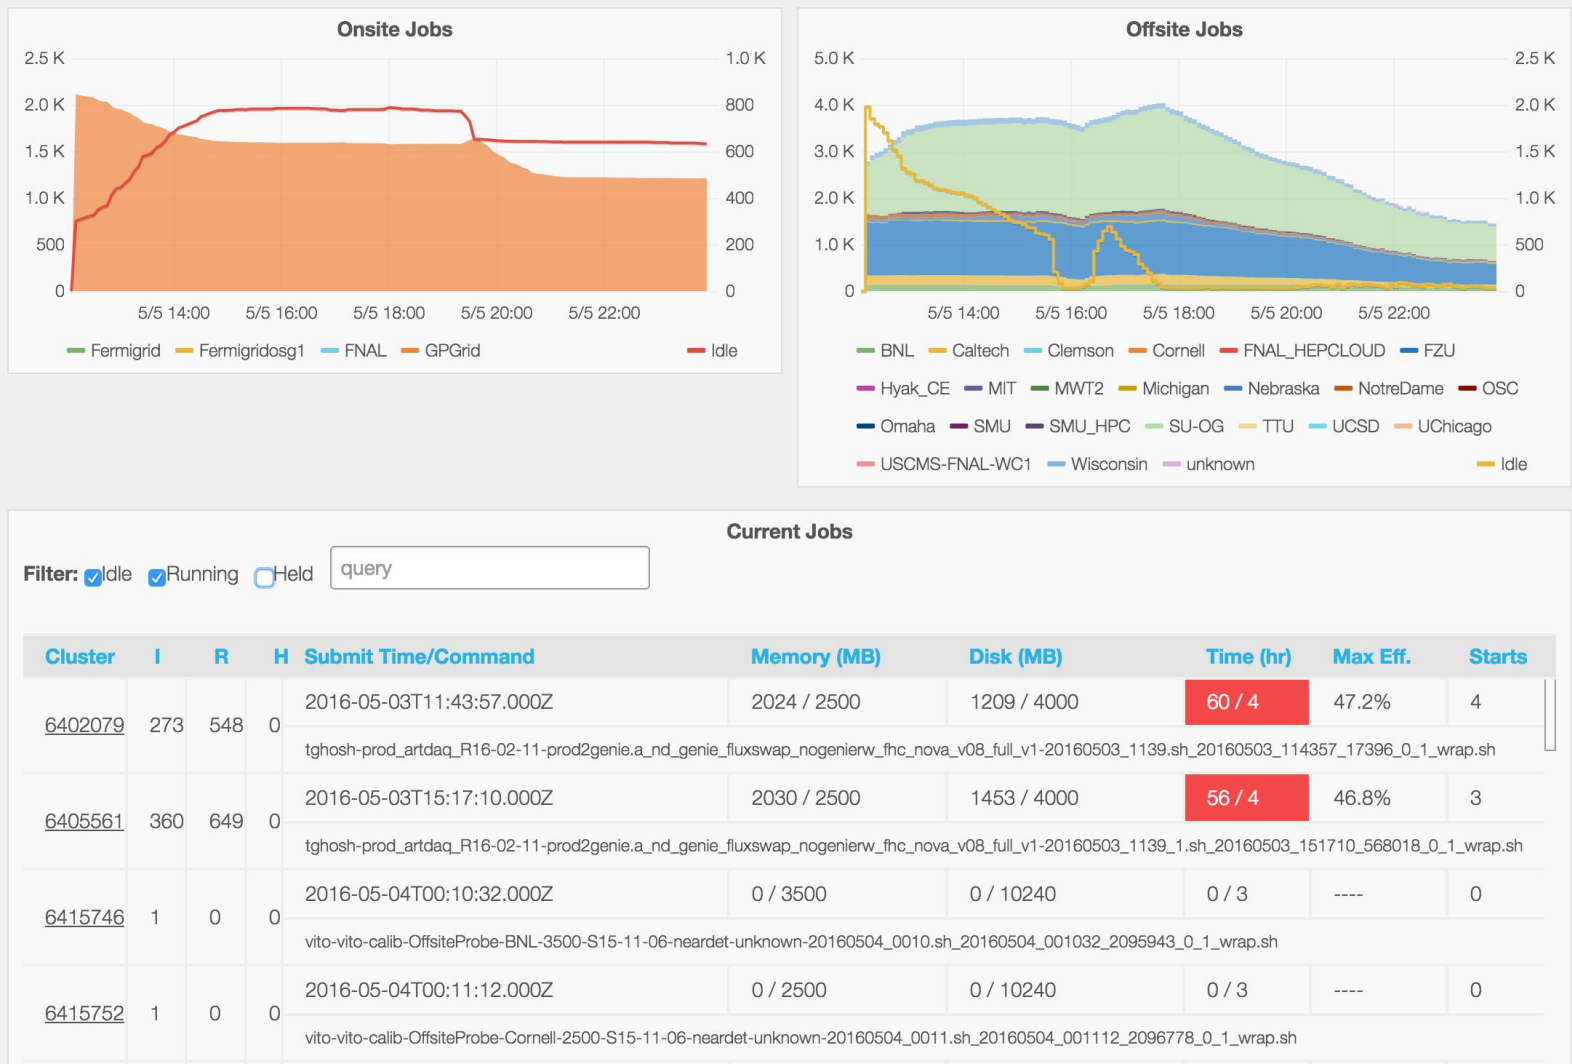
Monitor usage of slow moving resources so that projections can be made for projecting future need and limitations

# Monitoring of jobs and experimental dashboards



Monitoring for individual users to track their distributed computing workflows and understand their resource allocation and needs

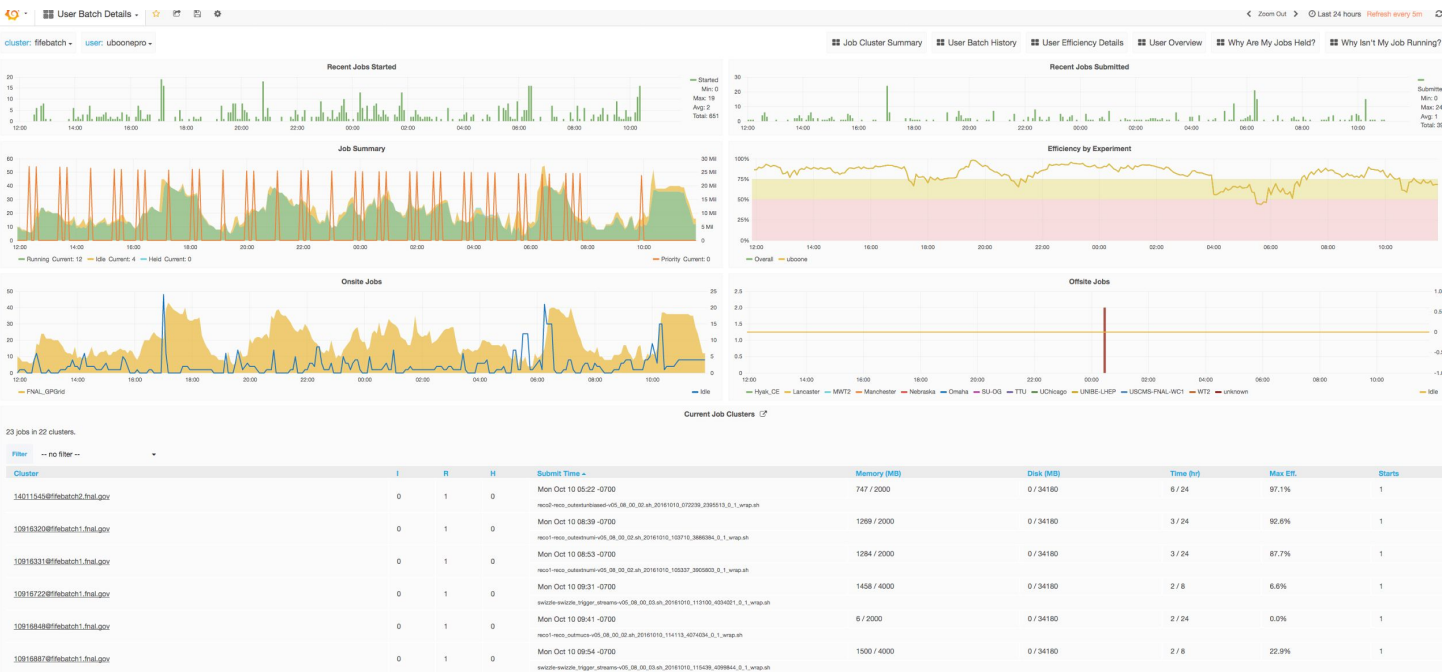
# Monitoring of jobs and experiment dashboards





# Monitoring at user level

Users have access to their own page, including special page with details of held jobs



### General Tips

#### What is the hold reason?

You can see this on your [User Batch Details](#) page, in the table below (select your username from the dropdown above), or by running:

```
jobsq --hold --user=<your username>
```

- SYSTEM\_PERIODIC\_HOLD  
This means your job exceeded requested resources.

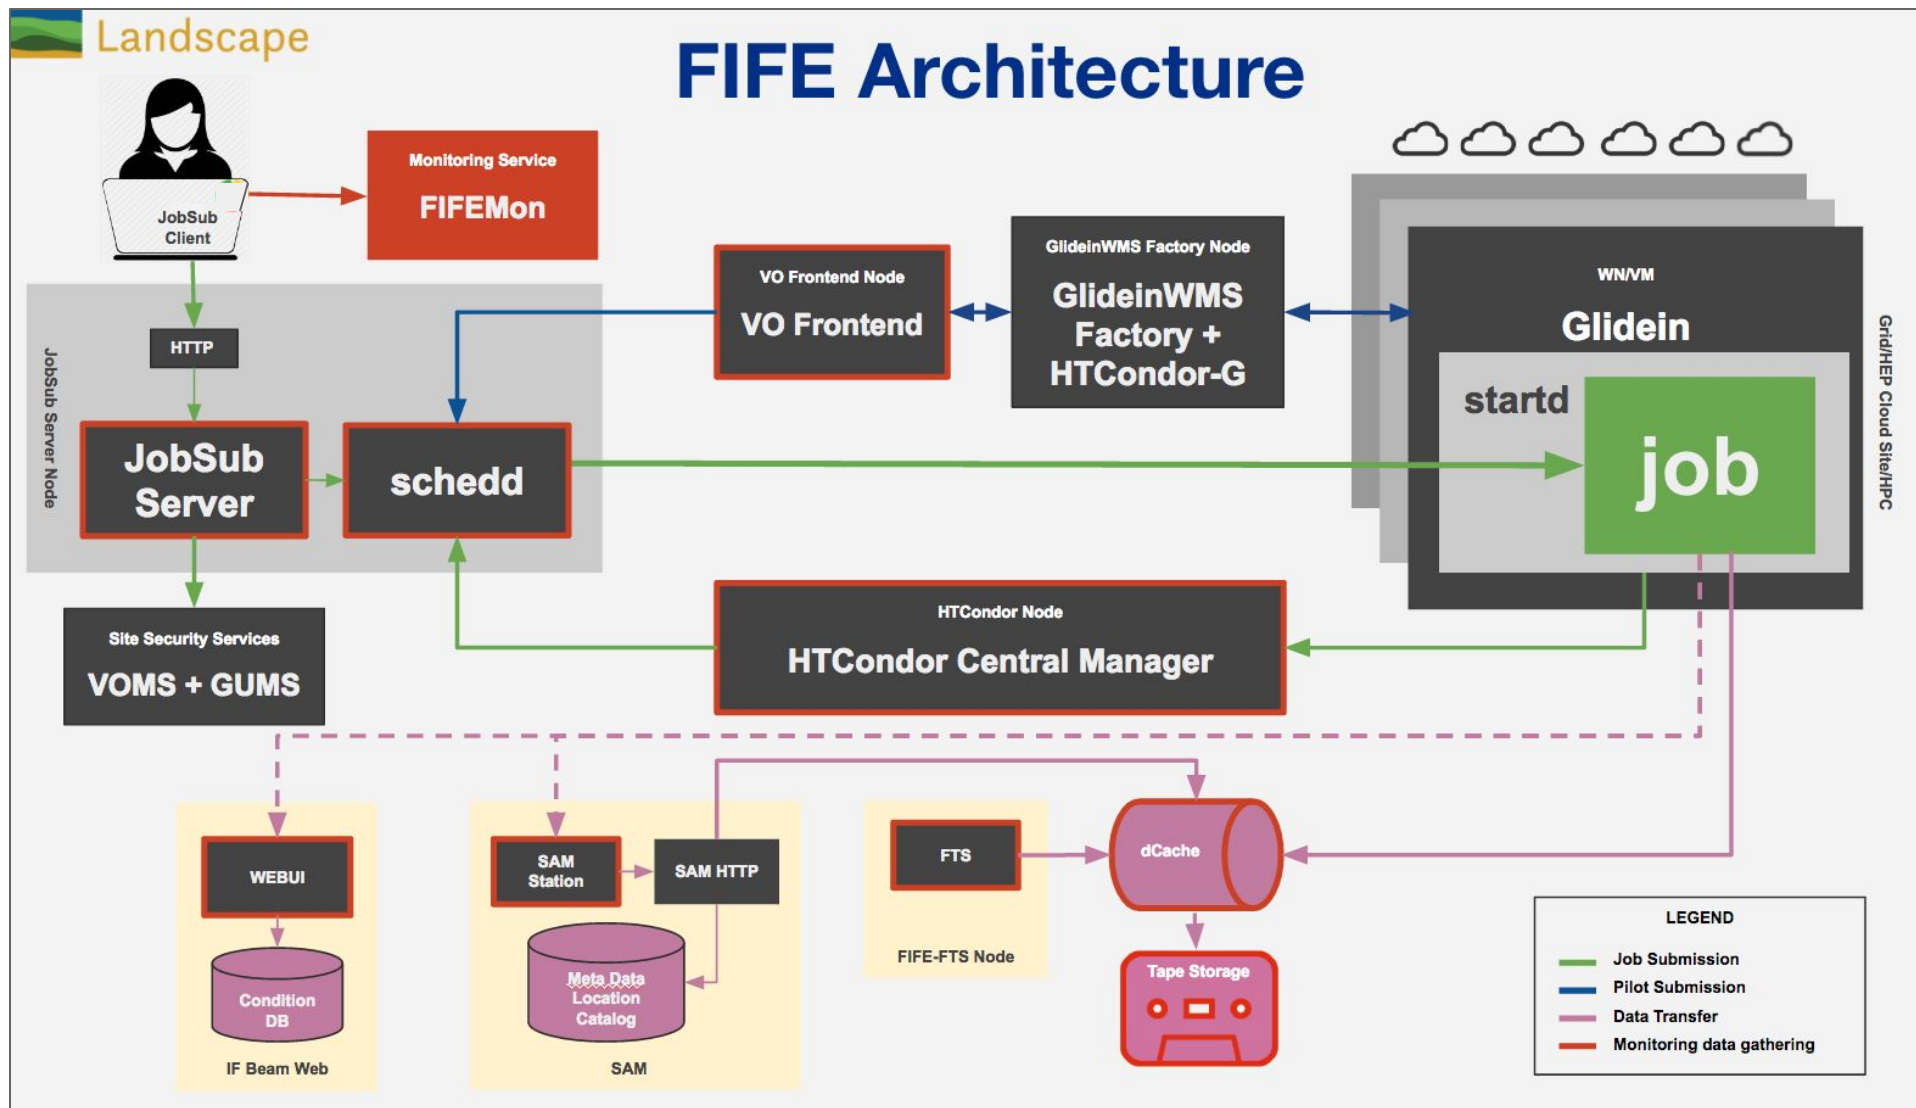
### HELD JOBS

Held Jobs				
jobid	hold_date	HoldReasonCode	HoldReasonSubcode	HoldReason
1397.0@fife-jobs-dev01.fnal.gov	2016-10-06 10:00:43	26	8	SYSTEM_PERIODIC_HOLD Run Time/limit 3607/3600
1394.0@fife-jobs-dev01.fnal.gov	2016-10-06 10:00:43	26	8	SYSTEM_PERIODIC_HOLD Run Time/limit 3608/3600
1396.0@fife-jobs-dev01.fnal.gov	2016-10-06 10:00:42	26	8	SYSTEM_PERIODIC_HOLD Run Time/limit 3606/3600
1395.0@fife-jobs-dev01.fnal.gov	2016-10-06 10:00:42	26	8	SYSTEM_PERIODIC_HOLD Run Time/limit 3606/3600
1392.0@fife-jobs-dev01.fnal.gov	2016-10-06 10:00:42	26	8	SYSTEM_PERIODIC_HOLD Run Time/limit 3607/3600

# Processing Data with SAM Projects and jobs

When processing data with SAM, one:

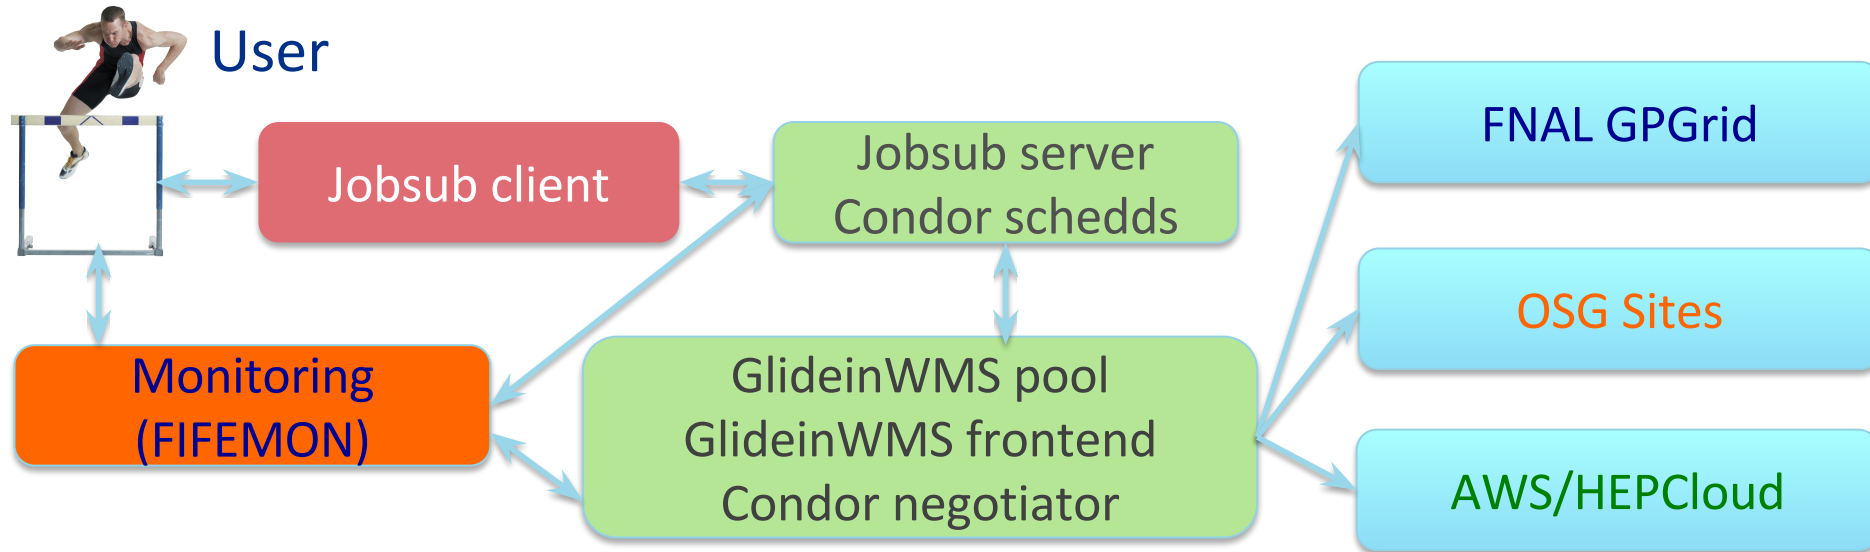
- Defines a dataset containing the files you want to process
- Start a SAM “Project” to hand them out
- Start one or more jobs which register as “Consumers” of the Project, including their location.
- Consumer Jobs then request files from the project, process them, and request another file, etc.
- Projects can prestage data while handing out data already on disk, and refer consumers to the “nearest” replica.
- Generally output is copied to an FFTS dropbox for production work, or to a user’s personal disk area.
- Thus the data is sent to the job, not the other way around
- However projects have limits; only so much at one submission.



Provide a modular architecture: experiments do not need to take all services. Can insert experiment-specific services as well (e.g. dedicated local SEs or local lab/university clusters)

# Job Submission and management architecture

- Common infrastructure is the **lifebatch** 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
  - Now also helps users to understand why jobs aren’t running
- **Automatic enforcement of memory, disk, and run time requests (jobs held if they exceed their request)**





# Simplifying I/O with IFDH

- File I/O is a complex problem (Best place to read? What protocol? Best place to send output?)
- **I**ntensity **F**rontier **D**ata **H**andling client developed as common wrapper around standard data movement tools; shield user from site-specific requirements and choosing transfer protocols
- Nearly a drop-in replacement for cp, rm, etc., but also extensive features to interface with SAM (can fetch files directly from SAM project, etc.)
- Supports a wide variety of protocols (including xrootd); automatically chooses best protocol depending on host machine, source location, and destination (can override if desired)
  - Backend behavior can be changed or new protocols added in completely transparent ways

## Centralized Services available from FIFE

- Submission to distributed computing: JobSub
  - GlideinWMS frontend
- Workflow monitors, alarms, and automated job submission
- Data handling and distribution
  - Sequential Access Via Metadata (SAM)
  - dCache/Enstore
  - File Transfer Service
  - Intensity Frontier Data Handling Client (data transfer)
- Software stack distribution via CVMFS
- User authentication, proxy generation, and security
- Electronic logbooks, databases, and beam information
- Integration with future projects, e.g. HEPCloud

# Data management: SAM and FTS

SAM originally developed for CDF and D0; many FNAL experiments now using it

- A File metadata/provenance catalog
- A File replica catalog ([data need not be at Fermilab](#))
- Allows metadata query-based “dataset” creation
- An optimized file delivery system (command-line, C++, Python APIs available)
- Originally a Oracle backend; now PostgreSQL
- Communication via CORBA for CDF/D0; now via http for everyone
  - Eliminates need to worry about opening ports for communication with server in nearly all cases


# Data management: SAM and FTS (2)

## Fermilab File Transfer Service

- Watches one or more dropboxes for new files
- Can extract metadata from files and declare to SAM, or handle files already declared
- Copies files to one or more destinations based on file metadata and/or dropbox used, register locations w/SAM
- Can automatically clean dropboxes, usually N days after files are on tape
- **Does *not* have to run at Fermilab, nor do source or destination have to be at Fermilab**

# CI Existing Plans

- Fermilab has already applied the Continuous Integration practice to the LArSoft-based experiments. Experiments on-boarded in Lar CI are: MicroBooNE, DUNE, LArIAT and ArgoNeuT.
- Because of the given justification, the CI project **plan** is to apply the Continuous Integration development practice to all IF experiments at Fermilab:
  - Extend Lar-CI practice to other no-LArSoft based experiments
  - Add additional features to the existing LAr-CI
  - Improve performance like: speed the response time of the DB/ schema changes (it requires some code and dataflow analysis to optimize the queries, it may need some DB model changes ... suspect scalability issue), create dynamic plots ....
  - Provide documentation to “facilitate” the use of the CI practice among the experiments.
- See CI redmine:  
<https://cdcv.sfnal.gov/redmine/projects/ci>
- Apply the The Plan Do Check Act (PDCA) cycle: work together with the experiments to define needs and priorities and receive feedback.






# Monitoring in the CI system - NOvA

ci\_calib\_regression\_test\_novasoftware on EL6  
Build nova\_ci\_beta/121  
Trigger: trigger script by  
vito@novagpvm12.fnal.gov  
Failed  
Success

Test ci\_calib\_regression\_test\_novasoftware Usage History



Statistics from the NOvA CI Calibration Test

## Test: ci\_calib\_regression\_test\_novasoftware

[stdout](#)  
[stderr](#)

Registered	2016-06-10 16:20:27.469562
Started	2016-06-10 16:20:27.832907
exitcode	0
rusage_user_cpu	10.110000
rusage_scaled_user_cpu	46.506607
rusage_system_cpu	0.610000
rusage_scaled_system_cpu	2.806037
rusage_elapsed	28.090000
rusage_%cpu	38.000000
rusage_avgtext	0.000000
rusage_avgdata	0.000000
rusage_maxrss	262256.000000
rusage_inputs	65368.000000
rusage_outputs	1424.000000
rusage_major_faults	436.000000
rusage_minor_faults	50025.000000
rusage_swaps	0.000000
valerrs	0
success	True
Finished	2016-06-10 16:20:56.768274
exit code:	0.0

build for EL6  
Build nova\_ci\_beta/121, Trigger:  
trigger script by  
vito@novagpvm12.fnal.gov  
on EL6  
status:  
Success

Stage build Runtime History



## Phase: build

[build](#)  
Started 2016-06-10 16:04:33.750498  
  
Finished 2016-06-10 16:19:51.991690  
exit code: 0

- Found an issue in the reco processing stage and in a commit of the NOvA code from a user (contacted and solved)

# Monitoring in the CI system - MicroBooNE

- Memory usage history plot: uboonecode geant4 stage as an example.



- Using CORSIKA as cosmic shower generator, memory usage goes from ~2Gb to ~3.5Gb.
- After the intervention of a memory profiling “task force” the memory usage went down to ~1.2Gb.

# POMS: Example Campaign Info

POMS

Dashboard

Calendar

Requests

DB Admin

Experiments

Users

Raw Tables

Campaigns

Active

InActive

Compose Campaign

Compose Definition

Compose Launch Template

Jobs

All (Last Day)

Failed\_By Exit Code and Campaign Name (Last Day)

Failed\_By Exit Code and Campaign Name (Last Week)

Failed\_By Exit Code and Node Name (Last Day)

Failed\_By Exit Code and CPU Type (Last Day)

Search...

Campaign uBooNE Electron Lifetime

Campaign

**Name:** uBooNE Electron Lifetime

**Experiment:** uboone

**Dataset:** none

**Software Version:** v05\_08\_00\_03

**Created:** 2016-08-05 17:42:27.461118-05:00

**Creator:** vito@fnal.gov

**VO Role:** Production

**Param Overrides:** [{"--configfile":"","ConfigFiles/Config\_ElectronLifetime\_test.cfg"}]

**cs\_split\_type:** None

**cs\_split\_dimensions:** None

**cs\_last\_split:** None

**Active:** True

Actions

[Job Efficiency Histogram](#)

[Day by Day Spreadsheet](#)

[Submission Time Bars](#)

[Campaign Submission Files](#)

[Launch Campaign Jobs Now](#)

[Kill Jobs for Campaign](#)

[Schedule Future Job Launches](#)

Campaign Definition

**Name:** uBooNE Electron Lifetime

**Creator:** vito@fnal.gov

**Created:** 2016-08-05 17:49:08.688918-05:00

**Launch Script:** /uboone/app/home/uboonepro/KeepUp/ProductionKeepUp\_uBooNE.sh

**Definition Parameters:** [{"--configfile":"","ConfigFiles/Config\_ElectronLifetime.cfg"}]

**Input Files Per Job:** 0

**Output Files Per Job:** 0

**Output File Patterns:** SwizRecoLifetime\_hist\_%

Tags

+

Recent Launch Outputs

- 20160824\_161120
- 20160824\_141115
- 20160824\_121133
- 20160824\_101132
- 20160824\_081132

Launch Template

**Name:** uboone template

**Launch Host:** uboonegvm07.fnal.gov

**Launch Account:** uboonepro

**Launch Setup:** echo Launch Template

# POMS: Example of Troubleshooting

## Jobs by user\_exe\_exit\_code,node\_name,experiment ?

🕒 2016-08-23 19:01 to 2016-08-24 19:01

<previous 1 days | next 1 days >



node\_name ✕

user\_exe\_exit\_code ✕

experiment ✕



Submit

user_exe_exit_code	node_name	experiment	count
65	fnpc8001.fnal.gov	uboone	4
65	fnpc7015.fnal.gov	uboone	1
65	fnpc8002.fnal.gov	uboone	1
65	fnpc7002.fnal.gov	uboone	1
65	fnpc3274.fnal.gov	uboone	1
250	acas1396.usatlas.bnl.gov	nova	1
65	fnpc3284.fnal.gov	uboone	1
65	fnpc2066.fnal.gov	uboone	1
65	fnpc4217.fnal.gov	uboone	1
65	fnpc2126.fnal.gov	uboone	1