

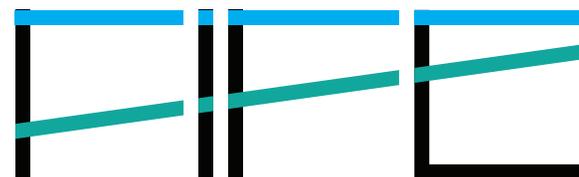


Fabric for Frontier Experiments - FIFE

Michael Kirby

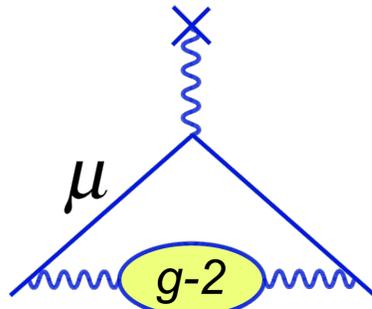
Fermilab Operations Review

16th-18th May 2016

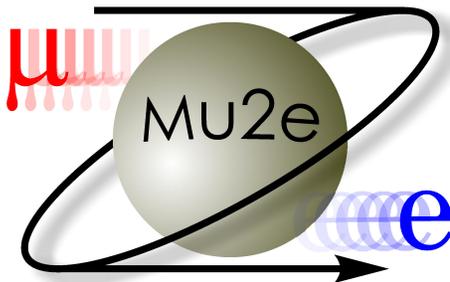


FABRIC FOR FRONTIER EXPERIMENTS

Need for FIFE Support Group



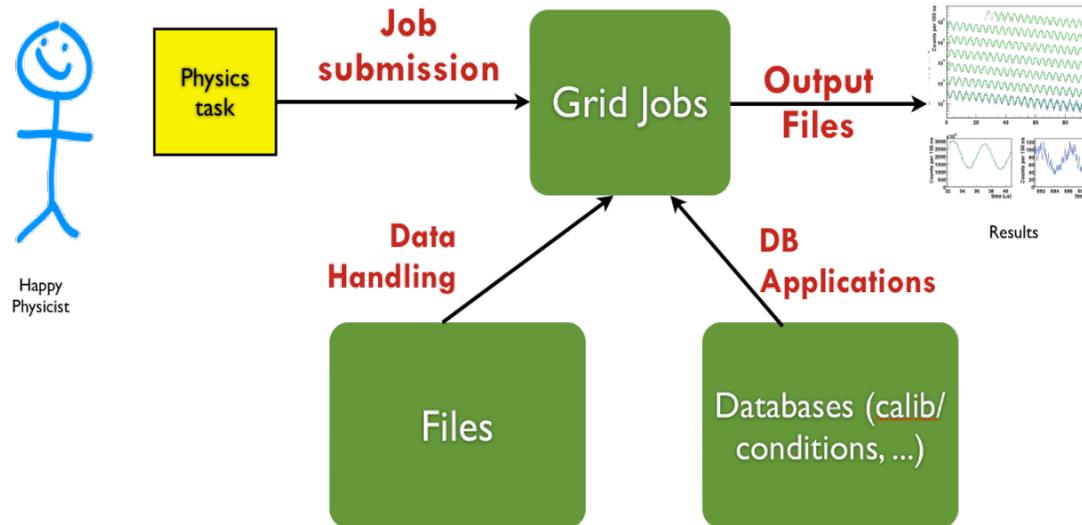
DARK ENERGY SURVEY



MINOS

- Proliferation of smaller, diversified experiments at Fermilab
- Computing needs critical mass of effort for design and operations
- FIFE team brings to the table full set of services and matches them to each experiment
- Educate the experiments about central services, ensure consistent security controls, and aid in the implementation and customization of services they request
- Serve as conduit for communicating ideas
 - Improving services as based upon new requirements and experimental ideas
 - Transmitting ideas and tools across experiment boundaries
 - Centralized information for service improvements
- Customer support for SCD services

FIFE Group Goals



- Provide collaborative scientific-data processing solutions
- Based on common toolsets wherever possible
 - Increase flexibility
 - Improve efficient evolution
 - Reduce maintenance load for experiments and service providers
- Include services over a broad spectrum of tasks
 - Grid submission and monitoring of jobs
 - Data management and data handling with integration to analysis framework
 - Database applications: beam, hardware, and conditions
 - Security controls and authentication mechanism – allows tracability
 - Collaborations with experiments to build integrated solutions

Computing requirements for IF experiments @ Fermilab

- 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
- greater detector granularity requires higher bandwidth DAQ
- event generation and detector response simulation
- reconstruction and analysis algorithms

Fermilab Scientific Computing Review

Mil CPU Hours	FY15 Actual	FY15 Req	FY16 Req
MINERvA	9.6	22	34
MINOS	17.2	18	22
NOvA	18.4	16	16
Mu2e	14.1	14	20
DUNE	8.4	11	11.4
LArIAT	0.30	8.8	2.5
Muon g-2	0.58	5.2	7
MARS	6.3	6.3	4
Seaquest	2.6	1.8	1.5
MicroBooNE	2.7	0.5	10.6
Total	80	104	129

Centralized Services through FIFE Support

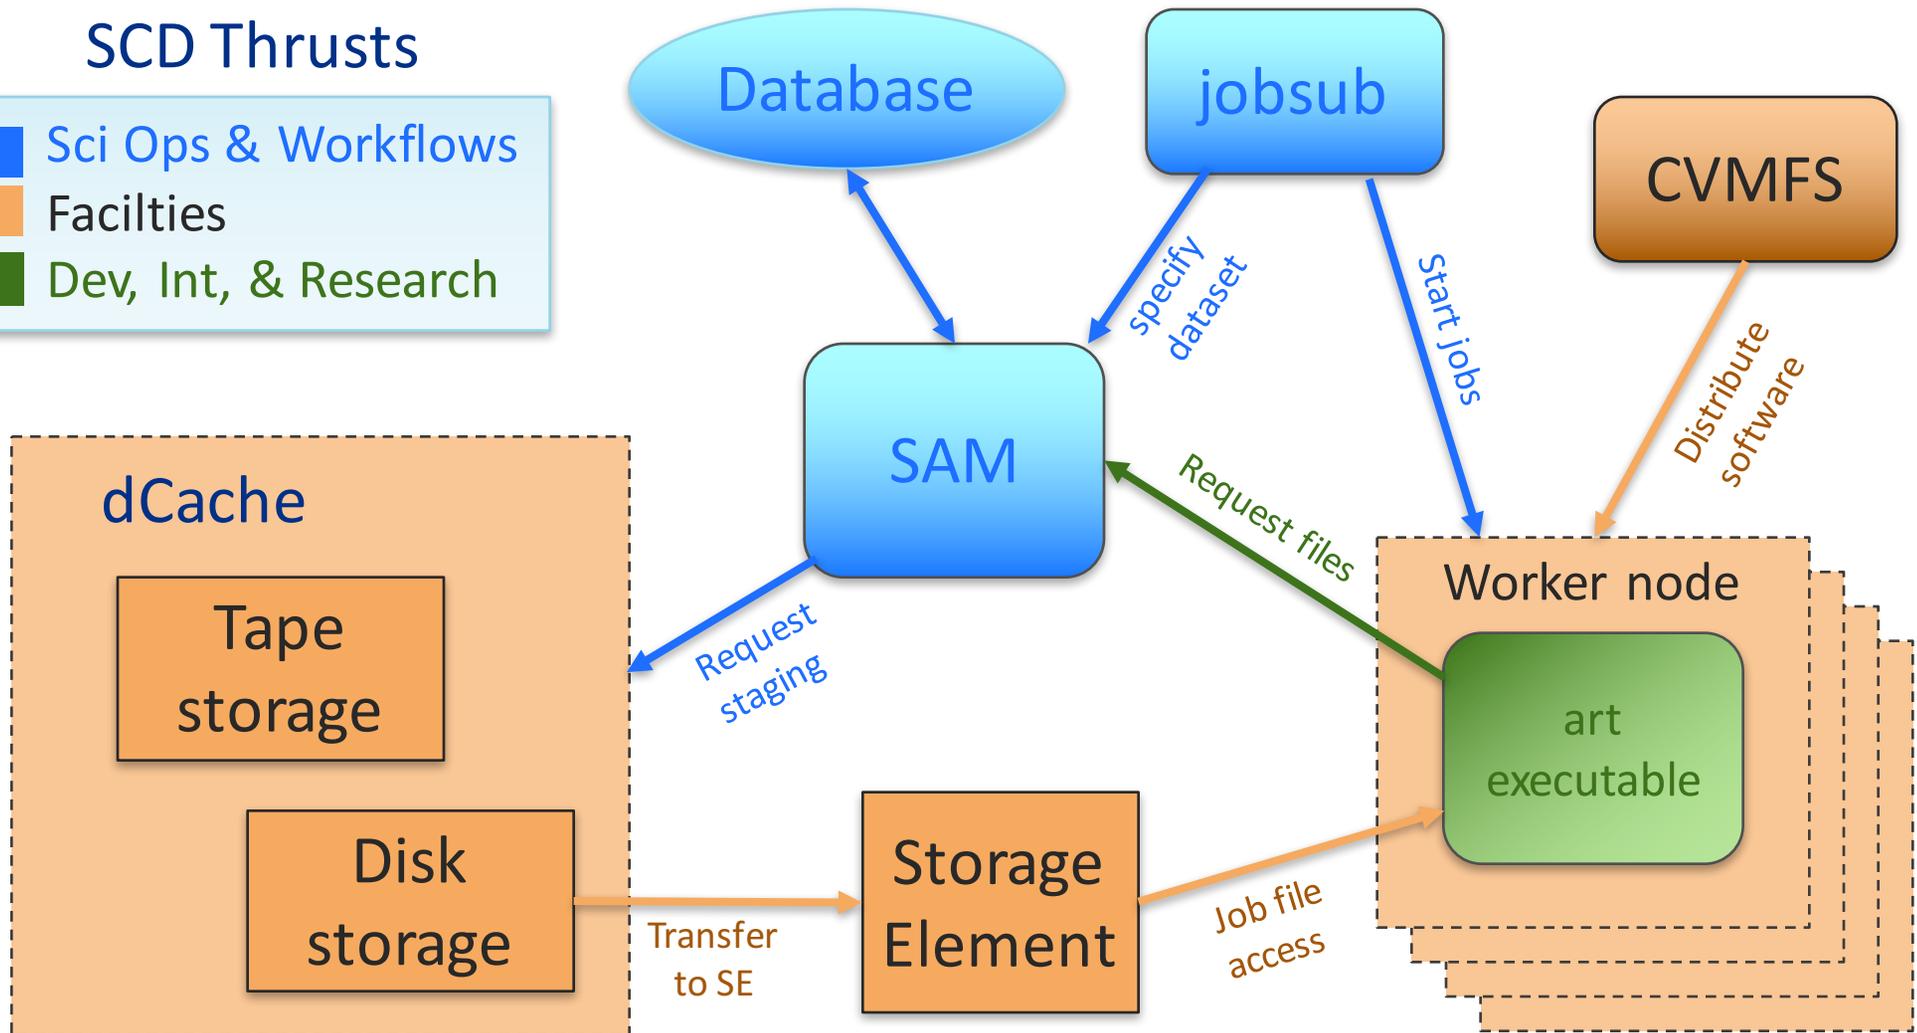


- Submission to distributed computing – JobSub
- Processing Monitors, Alarms, and Automated Submission
- Data Handling and Distribution
 - Sequential Access Via Metadata (SAM)
 - dCache/Enstore
 - File Transfer Service
 - Intensity Frontier Data Handling Client
- Software stack distribution – CERN Virtual Machine File System (CVMFS)
- User Authentication, Proxy generation, and security
- Electronic Logbooks, Databases, and Beam information
- Integration with future projects, e.g. HEPCloud

Example of centralized FIFE Services

SCD Thrusts

- Sci Ops & Workflows
- Facilities
- Dev, Int, & Research



How the FIFE Group Achieves Goals



- Experiment onboarding
 - Initial consultation with the experiment – get them off the ground
 - Provide guidance defining the requirements for different stages of experiment
 - Instantiate services for experiment as needed
 - Assist experiments to expand their computing capabilities

Area	FIFE	DAQ and Controls						Data Management				
Services												
Experiments	FIFE	DAQ	Fermigrid	Job Sub	FIFEMON	OSG Enabled	Amazon	IFDH	SAM	SAM Web	FTS	dcache/enstore
Lariat	Yes	artdaq	Deployed	Deployed	Deployed	In Progress		Deployed	Deployed	Deployed	Planning	Planning
DUNE	Yes	artdaq	Deployed	Deployed	Deployed	Deployed		Deployed	Deployed	Deployed	Planning	Planning
LBNE 35T	No	Deployed	Deployed	Deployed	Deployed	Deployed		Planning	Planning	Planning	Planning	Planning
Mars	Yes		Deployed	Deployed	Deployed	In Progress		Deployed	Deployed	Deployed		
uBoone	Yes	Custom	Deployed	Deployed	Deployed	Deployed	Planning	Deployed	Deployed	Not Required	Deployed	Deployed
Minerva	Yes	Custom	Deployed	Deployed	Deployed			Deployed	Deployed	Planning	Deployed	Deployed
MiniBoone	No	Custom	Deployed					Not Required	Not Required	Not Required		
Minos+	Yes	artdaq	Deployed	Deployed	Deployed	In Progress		Deployed	Deployed	Planning		
Mu2e	Yes	artdaq	Deployed	Deployed	Deployed	Planning		In Progress	In Progress	In Progress	Planning	Planning
Muon g-2	Yes	midas	Deployed	Deployed	Deployed	Deployed		Deployed	Deployed	Deployed	Planning	Planning
NoVA	Yes	novadaq	Deployed	Deployed	Deployed	Deployed	Deployed	Deployed	Deployed	Deployed	Deployed	Deployed

How the FIFE Group Achieves Goals

- FIFE Notes Newsletter every two months

FIFE Notes – April 2016

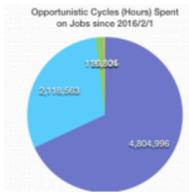
- 2016 Open Science Grid all-hands meeting
- HEP Cloud: How to add thousands of computers to your data center in a day
- DCAFI moving forward
- AFS transition
- Experiment with the most opportunistic hours Feb. – March 2016
- Most efficient big non-production users Feb. – March 2016
- Most efficient experiments Feb. – March 2016

Past FIFE Notes

FIFE Notes – April 2016

Best in Class

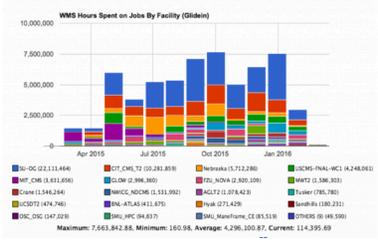
Experiment with the most opportunistic hours Feb. – March 2016



The experiment with the most opportunistic hours on OSG between Feb. 1, 2016 and March 31, 2016 was Mu2e with 4,804,996 hours.

Feature Articles

2016 Open Science Grid all-hands meeting



Every spring, the entire Open Science Grid (OSG) community--consisting of resource owners and operators, users, and staff--gathers at the annual OSG all-hands meeting. The 2016 OSG all-hands meeting was held between Monday, March 14 and

- Hold a FIFE workshop every year
 - Announce new services, service enhancements
 - Consultation sessions with all service providers present
 - Tutorials and discussion of best practices
 - Attended by computing leaders and analyzers from all experiments and SCD experts from all services – 57 attendees in 2015

8 5/13/16 Michael Kirby | Fermilab Operations Review



NOvA - full integration of FIFE Services



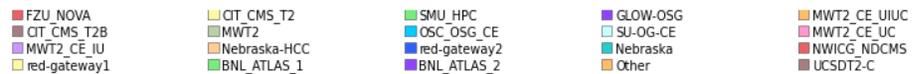
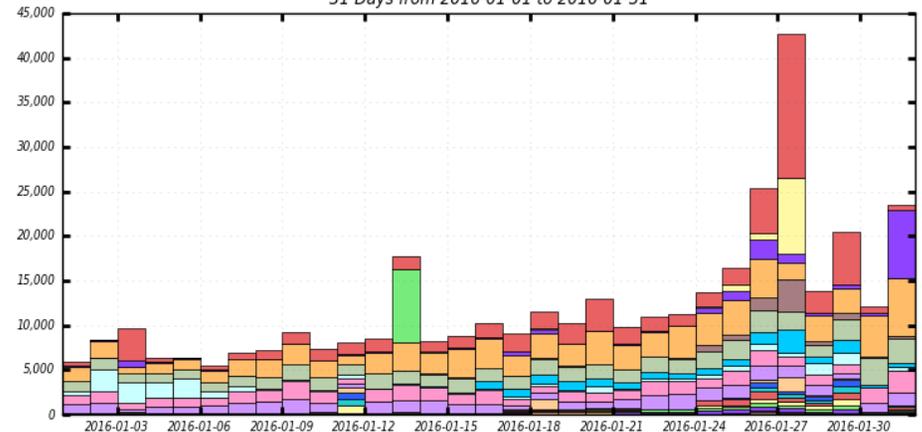
- Utilized early prototype of artdaq for DAQ
- Lean heavily on SAM and jobsub for running data and MC processing pipeline
- File Transfer Service utilized for catalog and transfer of Near and Far Detector files
- With dedicated resources on HPC sites, motivated to have remote site capabilities
- NOvA institutional knowledge and rapport with FIFE make it a natural test case for utilizing Amazon Web Service for on-demand process – step towards HEPCloud

Area	FIFE	DAQ and Controls						Data Management				
Services	FIFE	DAQ	Fermigrd	Job Sub	FIFEMON	OSG Enabled	Amazon	IFDH	SAM	SAM Web	FTS	dcache/enstore
Lariat	Yes	artdaq	Deployed	Deployed	Deployed	In Progress		Deployed	Deployed	Deployed	Planning	Planning
DUNE	Yes	artdaq	Deployed	Deployed	Deployed	Deployed		Deployed	Deployed	Deployed	Planning	Planning
LBNE 35T	No	Deployed	Deployed	Deployed	Deployed	Deployed		Planning	Planning	Planning	Planning	Planning
Mars	Yes		Deployed	Deployed	Deployed	In Progress		Deployed	Deployed	Deployed		
uBoone	Yes	Custom	Deployed	Deployed	Deployed	Deployed	Planning	Deployed	Deployed	Not Required	Deployed	Deployed
Minerva	Yes	Custom	Deployed	Deployed	Deployed			Deployed	Deployed	Planning	Deployed	Deployed
MiniBoone	No	Custom	Deployed					Not Required	Not Required	Not Required		
Minos+	Yes	artdaq	Deployed	Deployed	Deployed	In Progress		Deployed	Deployed	Planning		
Mu2e	Yes	artdaq	Deployed	Deployed	Deployed	Planning		In Progress	In Progress	In Progress	Planning	Planning
Muon g-2	Yes	midas	Deployed	Deployed	Deployed	Deployed		Deployed	Deployed	Deployed	Planning	Planning
NoVA	Yes	novadaq	Deployed	Deployed	Deployed	Deployed	Deployed	Deployed	Deployed	Deployed	Deployed	Deployed

NOvA – full integration of FIFE Services

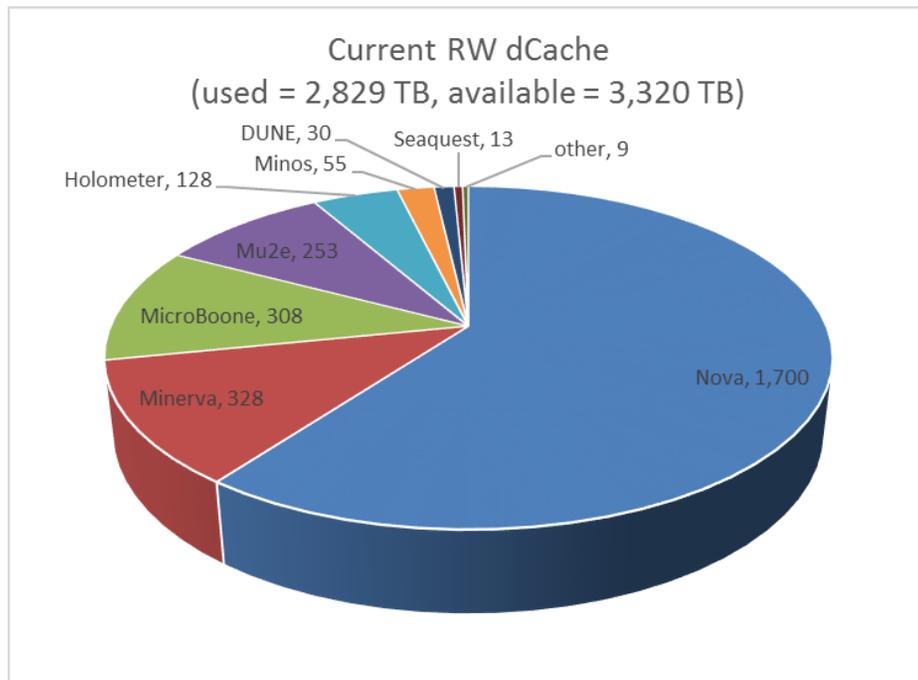
- File Transfer Service stored 1.7 PB of NOvA data in dCache and Enstore
- SAM Catalog contains more than 41 million files
- Helped develop SAM4Users as lightweight catalog

Hours Spent on Jobs By Facility
31 Days from 2016-01-01 to 2016-01-31



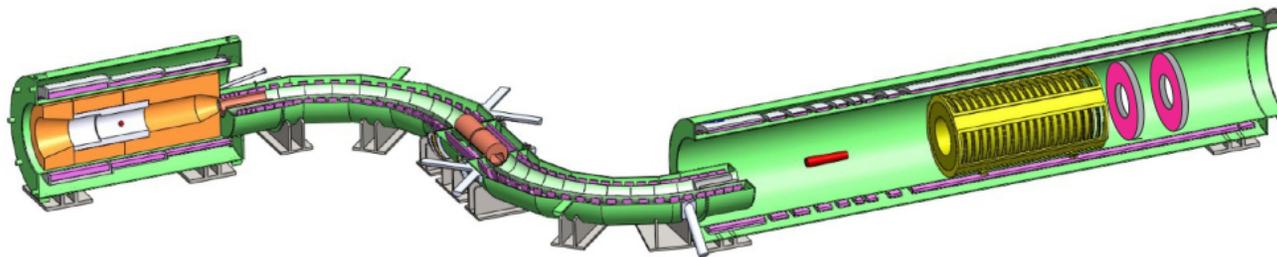
Maximum: 42,744 , Minimum: 5,535 , Average: 12,222 , Current: 23,484

- 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



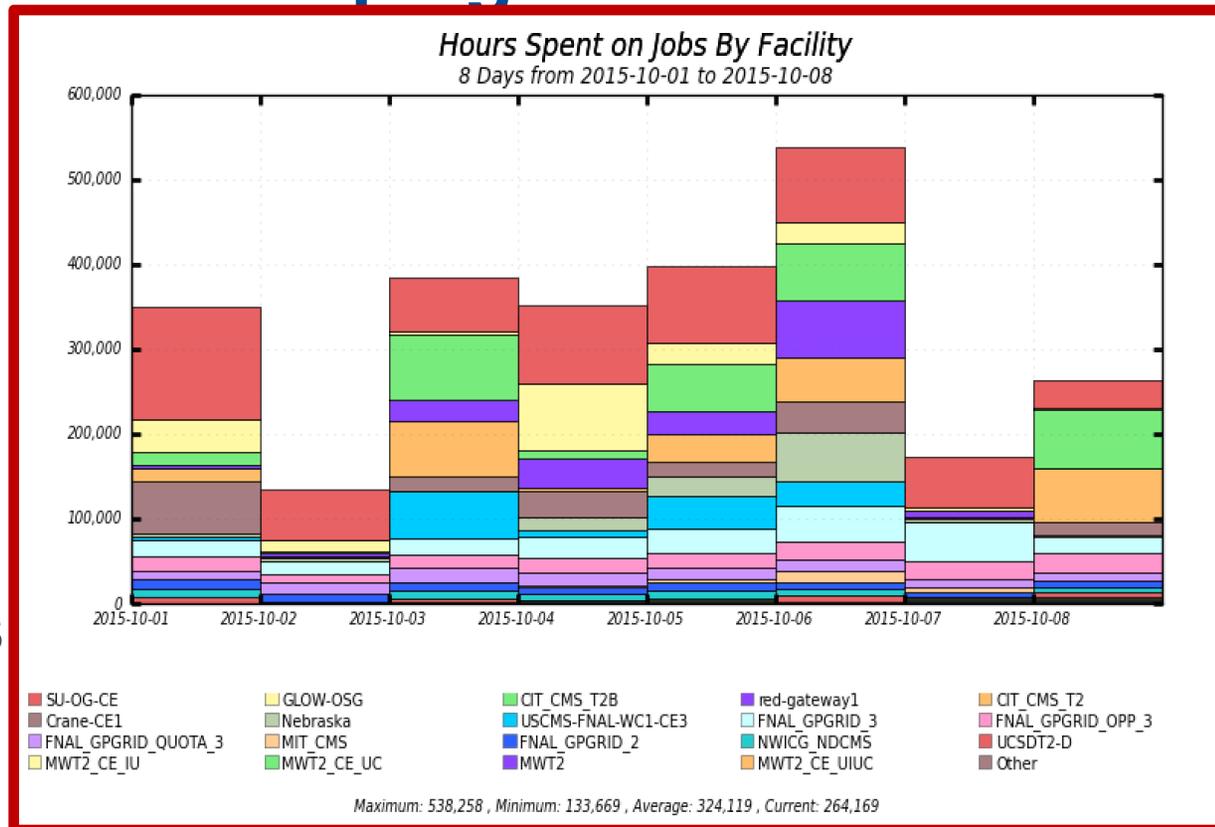
Mu2e Beam Simulations Campaign

- Mu2e is preparing to go through the CD-3a review – review design of beam transport, magnets, detectors, and radiation
- Combination of beam intensity and magnet complexity necessitates significant simulation studies
- 12 Million CPU hours in 6 months estimate for required precision
- Well beyond the available resources at Fermilab allocated to Mu2e
- FIFE support group helped deploy MARS beam simulation software stack through CVMFS to remote sites
- Helped probe additional remote resources and integrate into job submission – ideally without user knowledge



Mu2e Beam Simulations Campaign

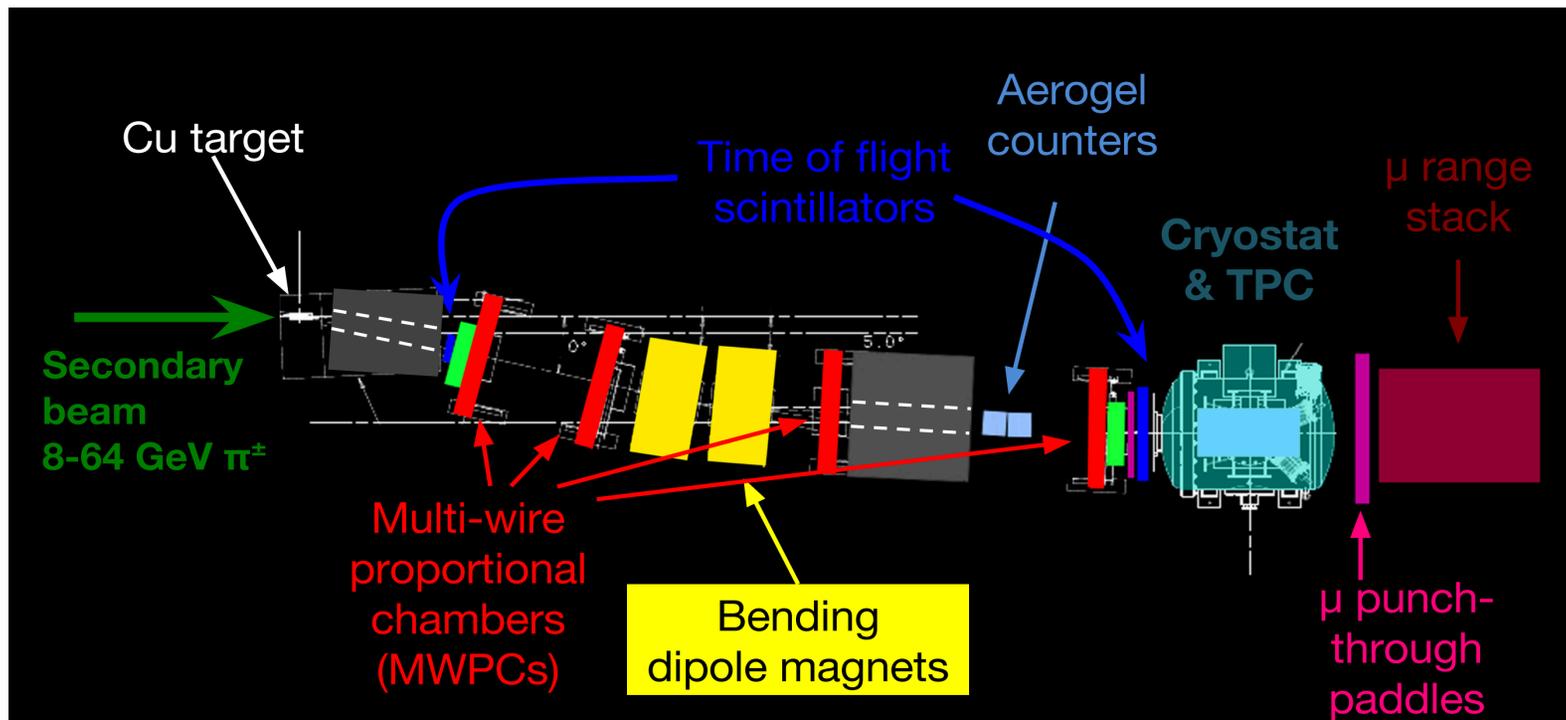
- Almost no input files
- Heavy CPU usage
- < 100 MB output
- Ran > 20M CPU-hours in under 5 months
- Avg 8000 simultaneous jobs across > 15 remote sites



- Usage as high as 20,000 simultaneous jobs and 500,000 CPU hours in one day – peaked usage 1st wk Oct 2015
- Achieved stretch goal for processing 24 times live-time data for 3 most important backgrounds

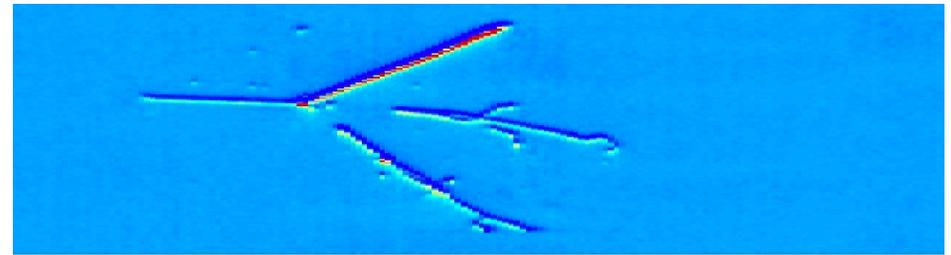
Enhancement of LArIAT SAM File catalog

- Liquid Argon In A Testbeam - exploring the cross-sections on LAr for final state particles
- Important for understanding the response in future detectors
- Incident beam can change every day, but DAQ not coupled to bending magnets – incorporate beam db into file catalog



Enhancement of LArIAT SAM File Catalog

- Extended the capability of SAM to be able to interface with external databases
- Allows for LArIAT to select data based upon criteria from the beam condition database
- DAQ and Offline processing are independent of beam database so that this is not a blocking situation
- FIFE Support team helped to instantiate and configure this beam db integration with LArIAT SAM Catalog
- Analyzers focused on physics instead of computing
- LArIAT presented first cross-sections at W&C April 8, 2016



LArIAT

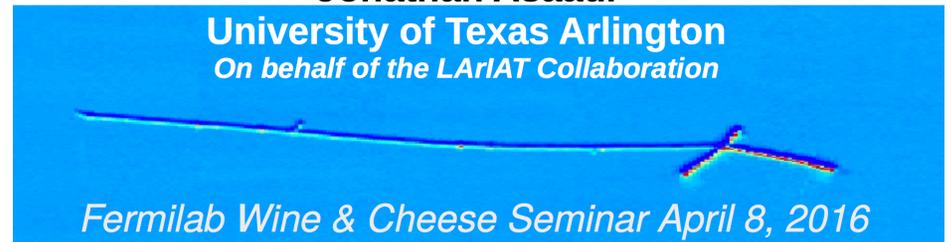
Liquid Argon TPC In A Testbeam

First Total π -Ar Cross Section Measurement

Jonathan Asaadi

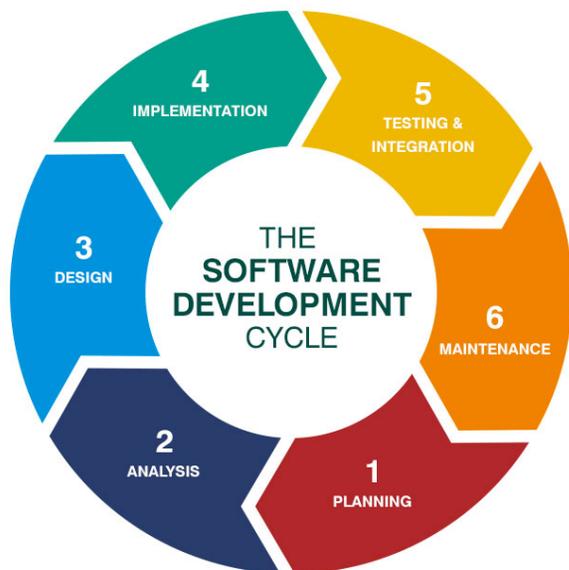
University of Texas Arlington

On behalf of the LArIAT Collaboration

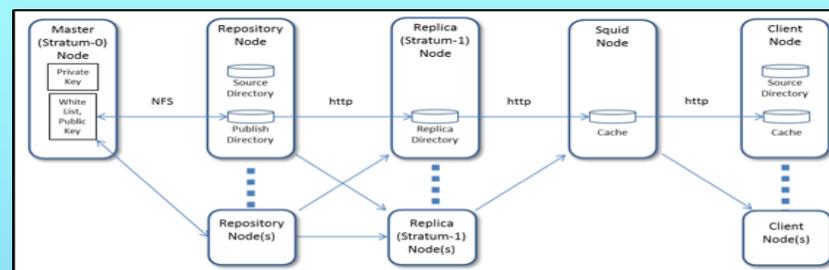


Fermilab Wine & Cheese Seminar April 8, 2016

Contributing back to software stack



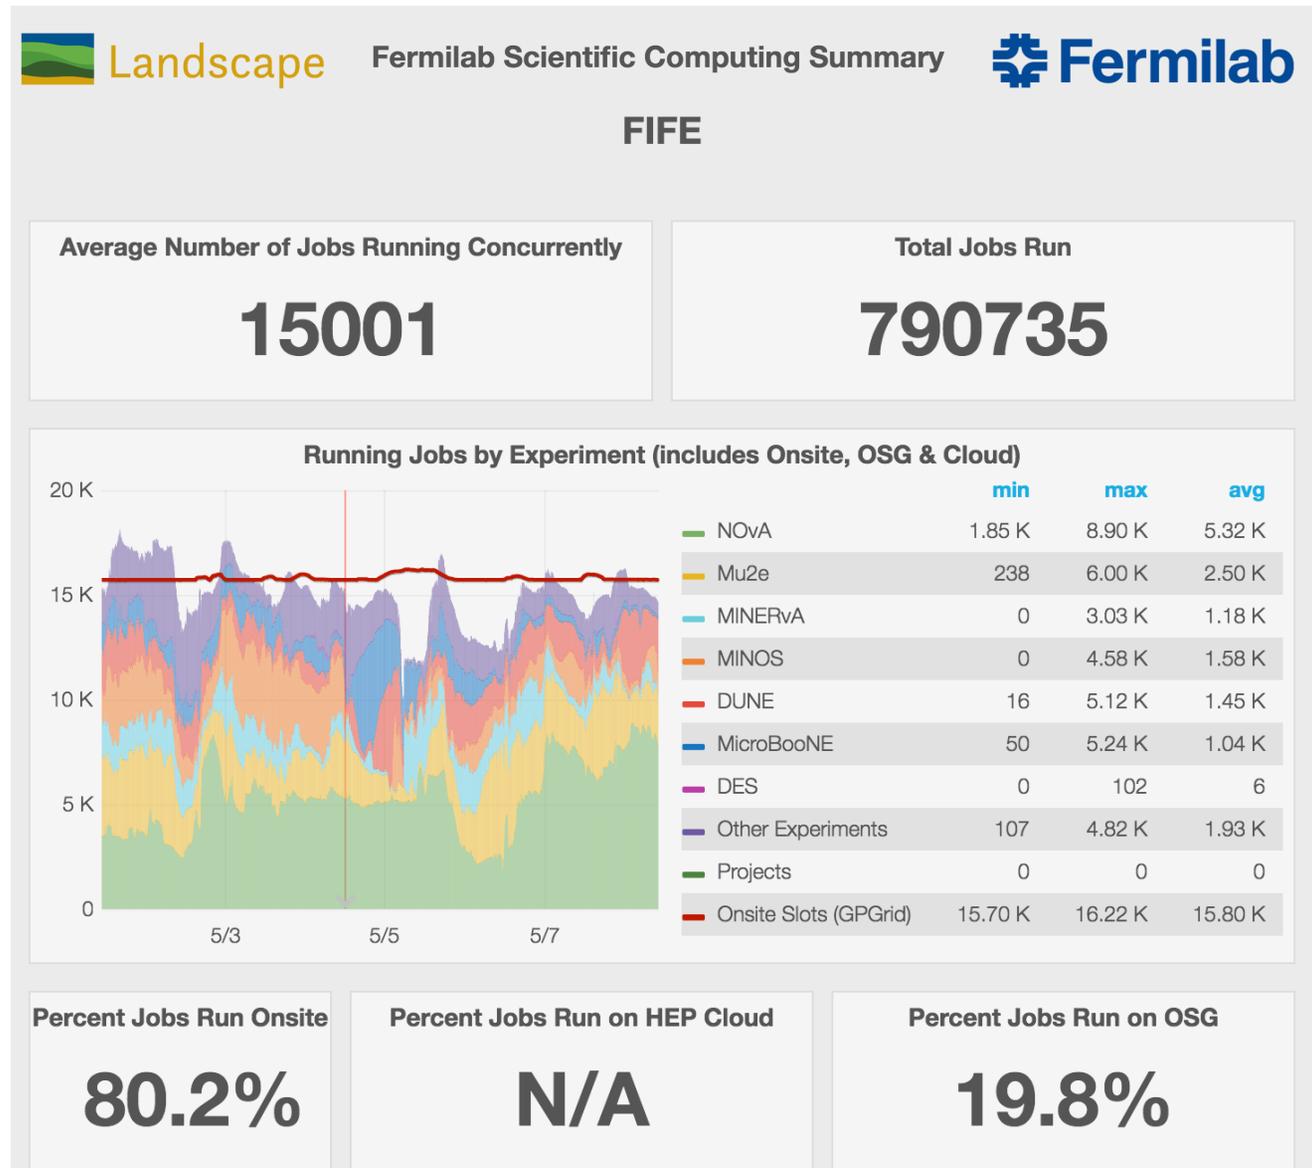
Software distribution to Grid worker nodes: CVMFS infrastructure



- increase of Fermilab experiments utilizing OASIS CVMFS caused conflicts updating and syncing software on OASIS
- To relieve conflicts Fermilab worked with CERN to update CVMFS and OASIS to integrate remote CVMFS repositories
- CVMFS repositories located at sites (Fermilab, other labs)
- distribution of large files for simulation tasks -> development of StashCache
- FIFE served the role of collating and communicating requirements, and contributing to design, testing, and implementation to include monitoring and tracking usage

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



Overview of Experiment Computing Operations

Select Experiment:

ANNIE

CDF

CDMS

D0

DUNE

LArIAT

MINERvA

MINOS

MicroBooNE

Mu2e

NOvA

SBND

SeaQuest

g-2



MicroBooNE Computing Summary



Average Jobs Running Concurrently

1042

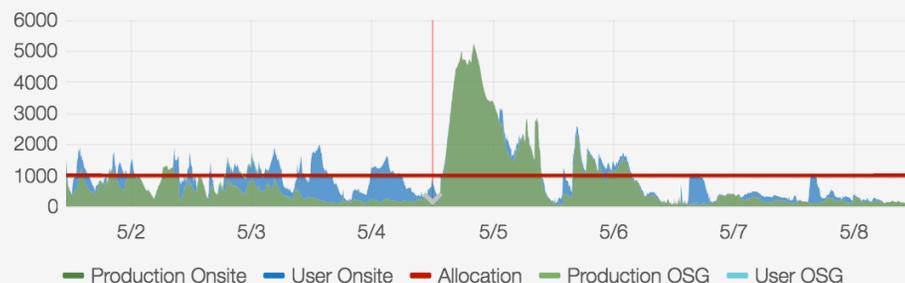
Total Jobs Run

168855

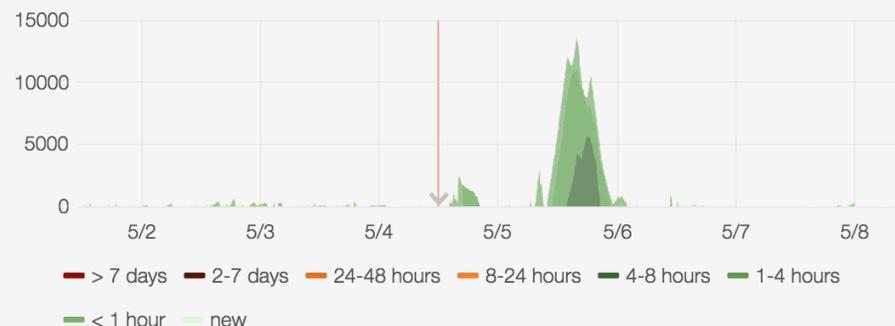
Average Time Spent Waiting in Queue (Production)

25.6 min

Running Batch Jobs



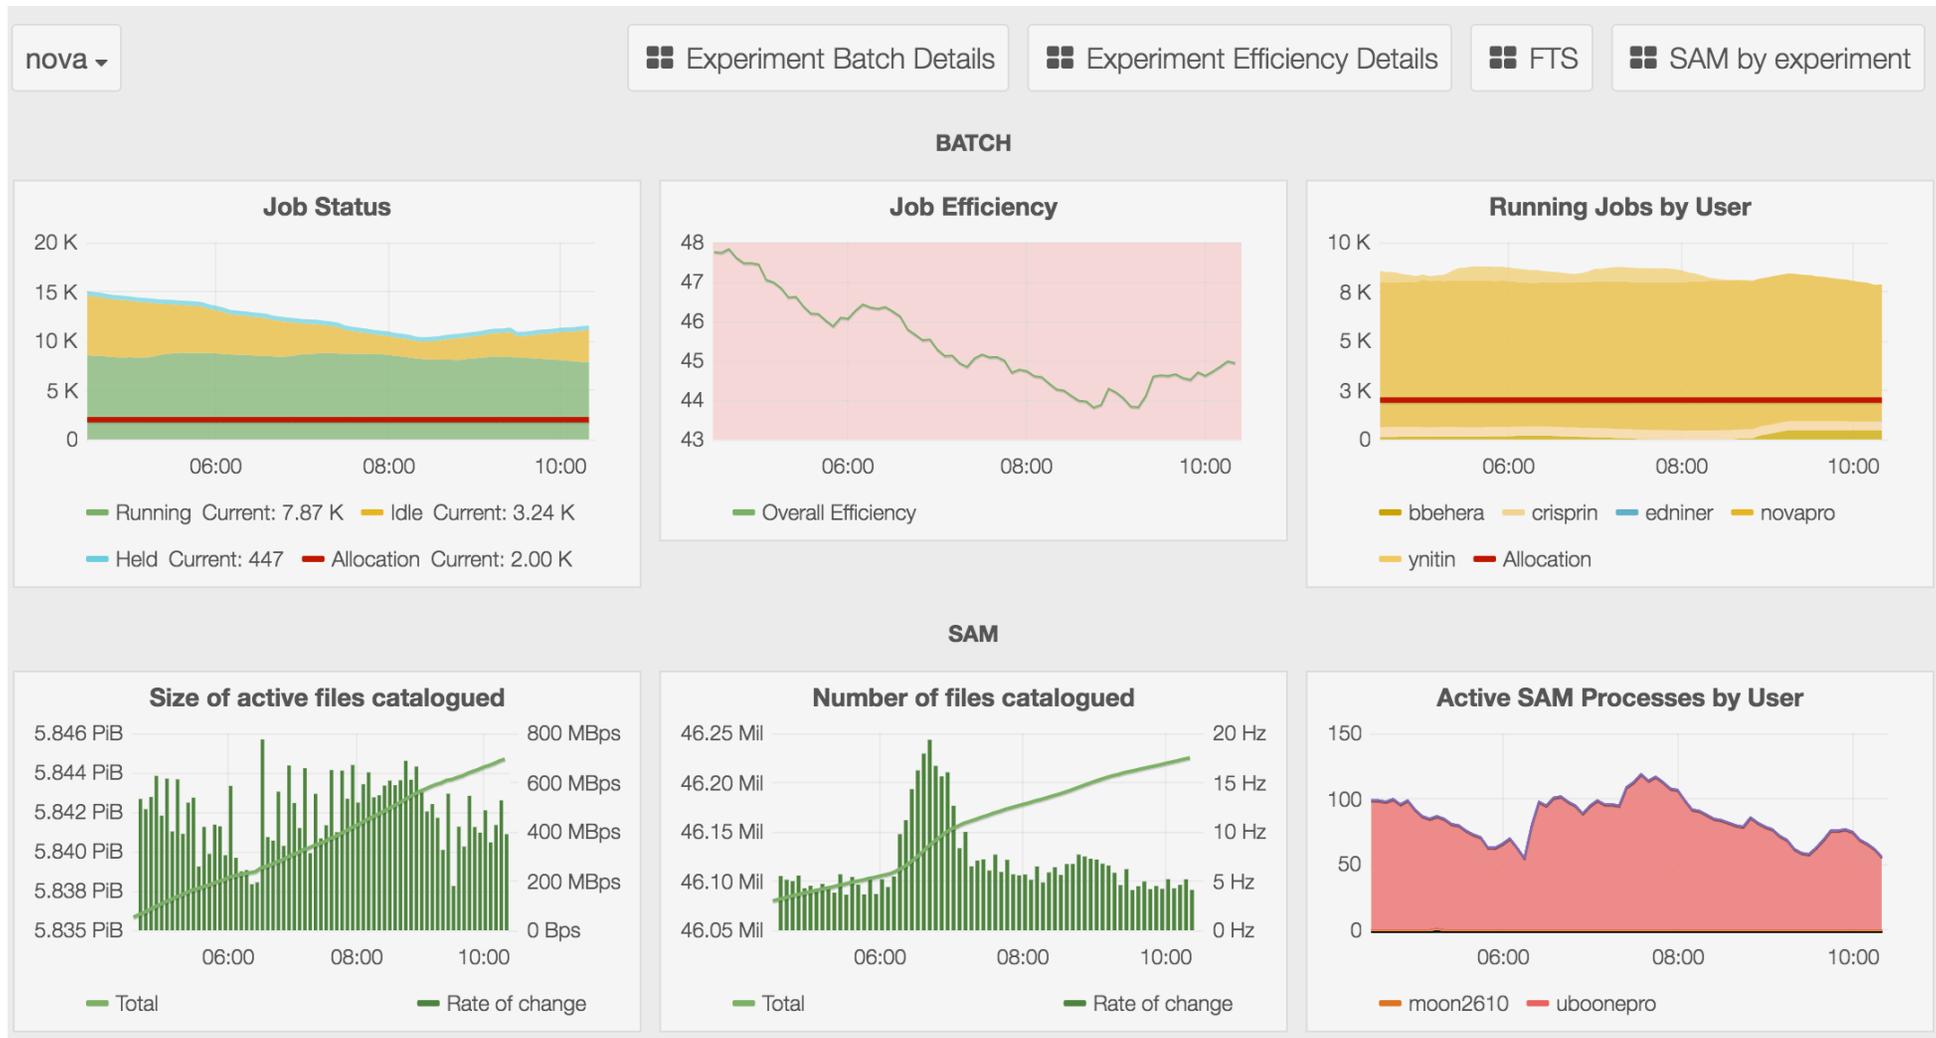
Queued Production Jobs by Wait Time



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

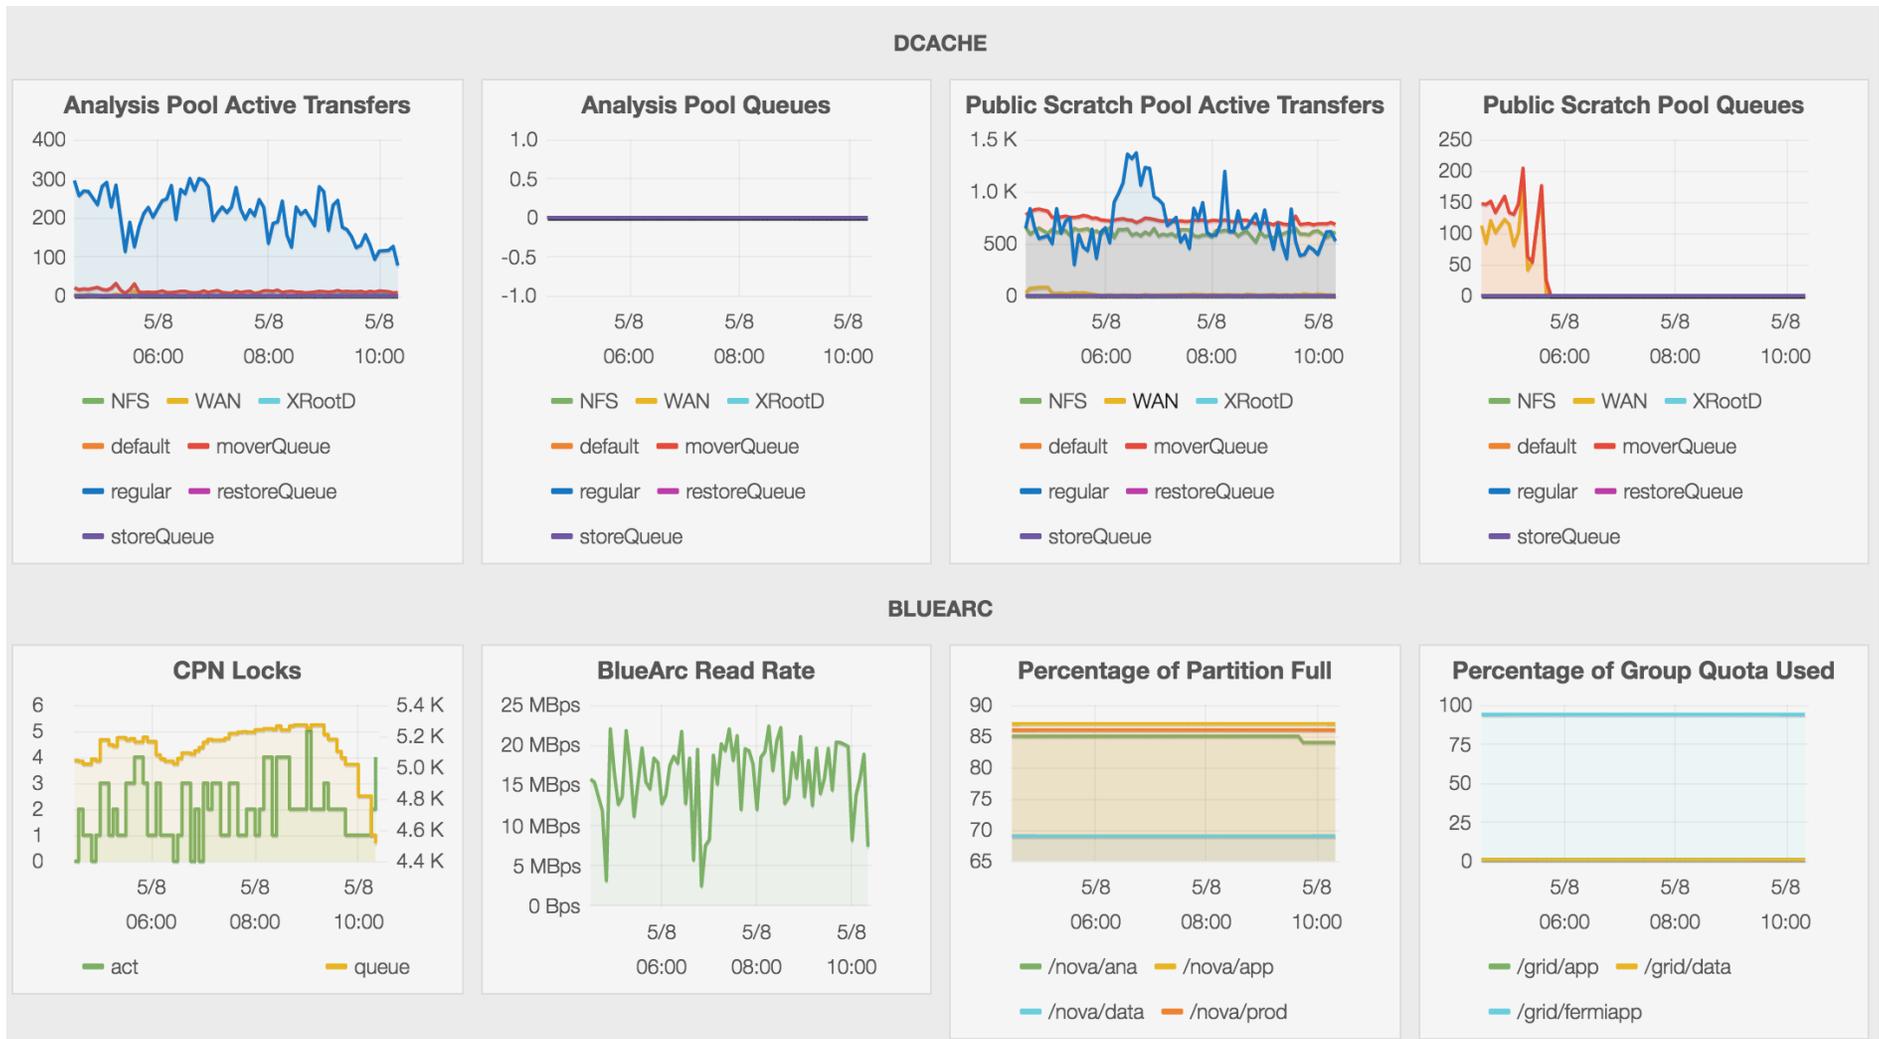


Detailed profiling of experiment operations



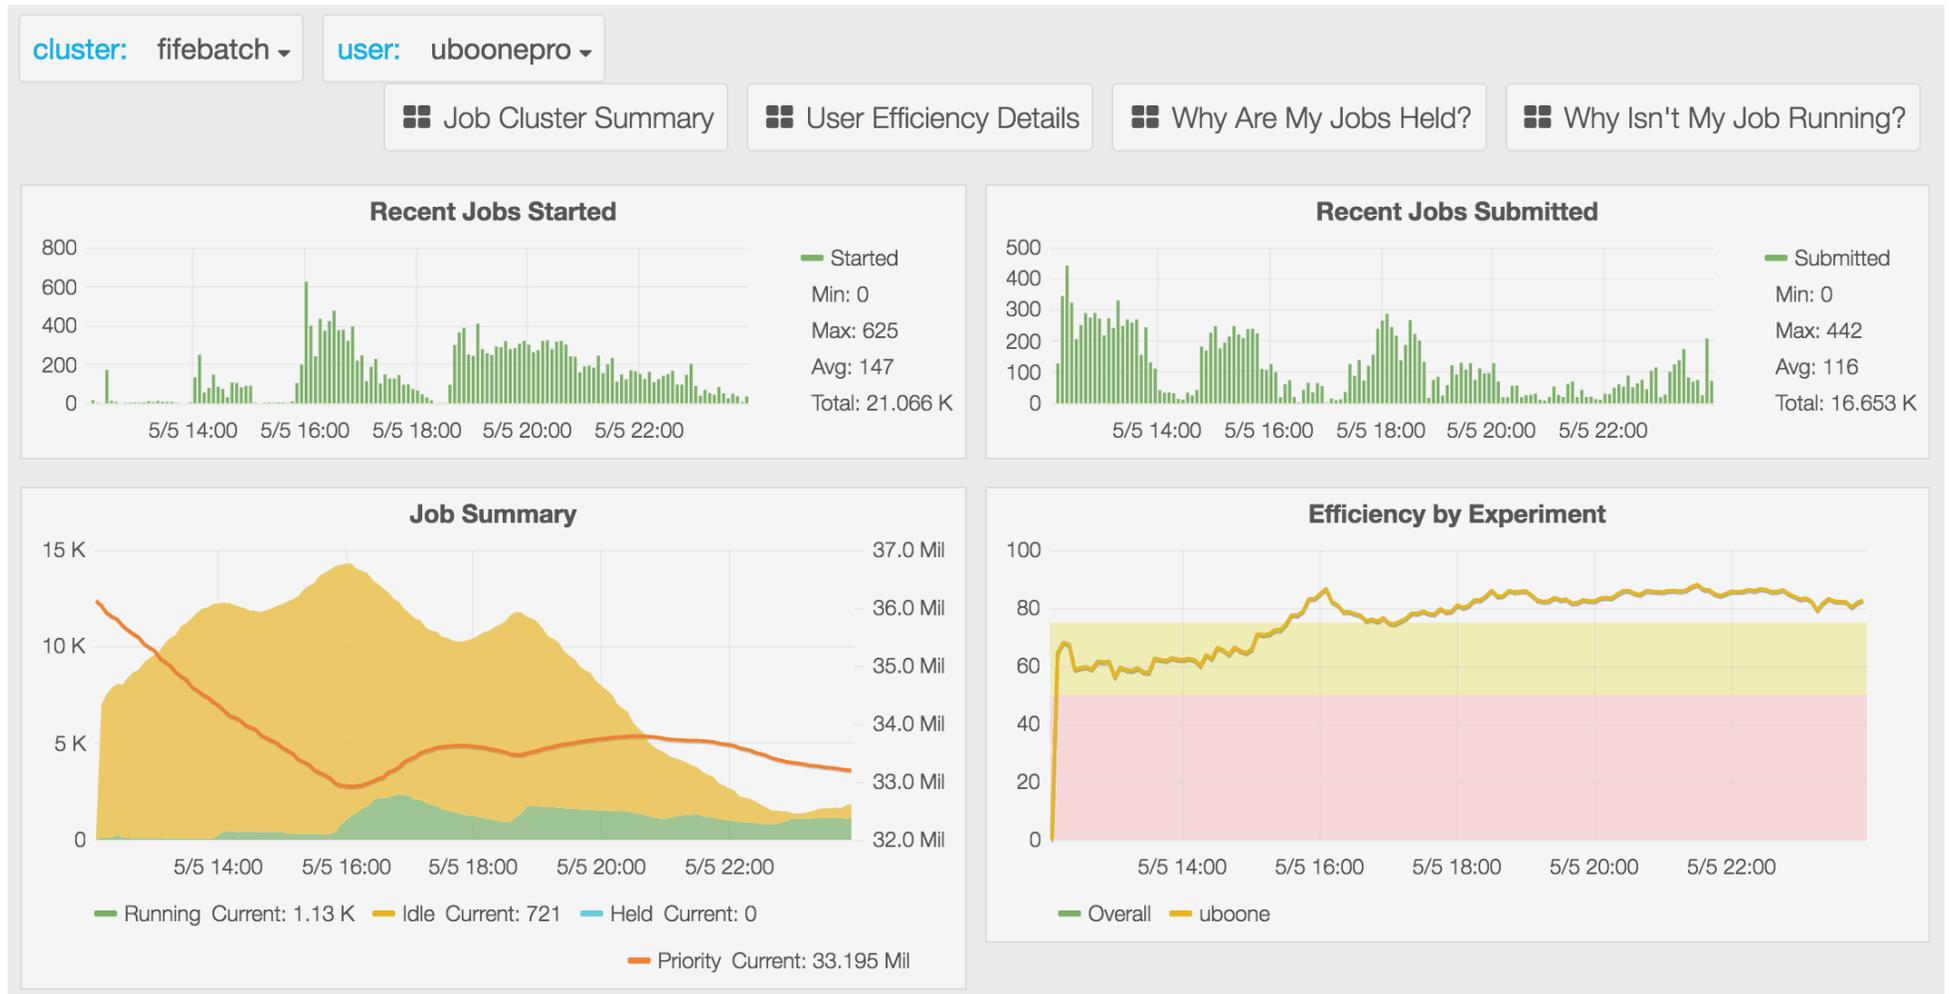
Allows identification for inefficiencies, potential slow downs, or blocking conditions in workflows

Detailed profiling of experiment operations



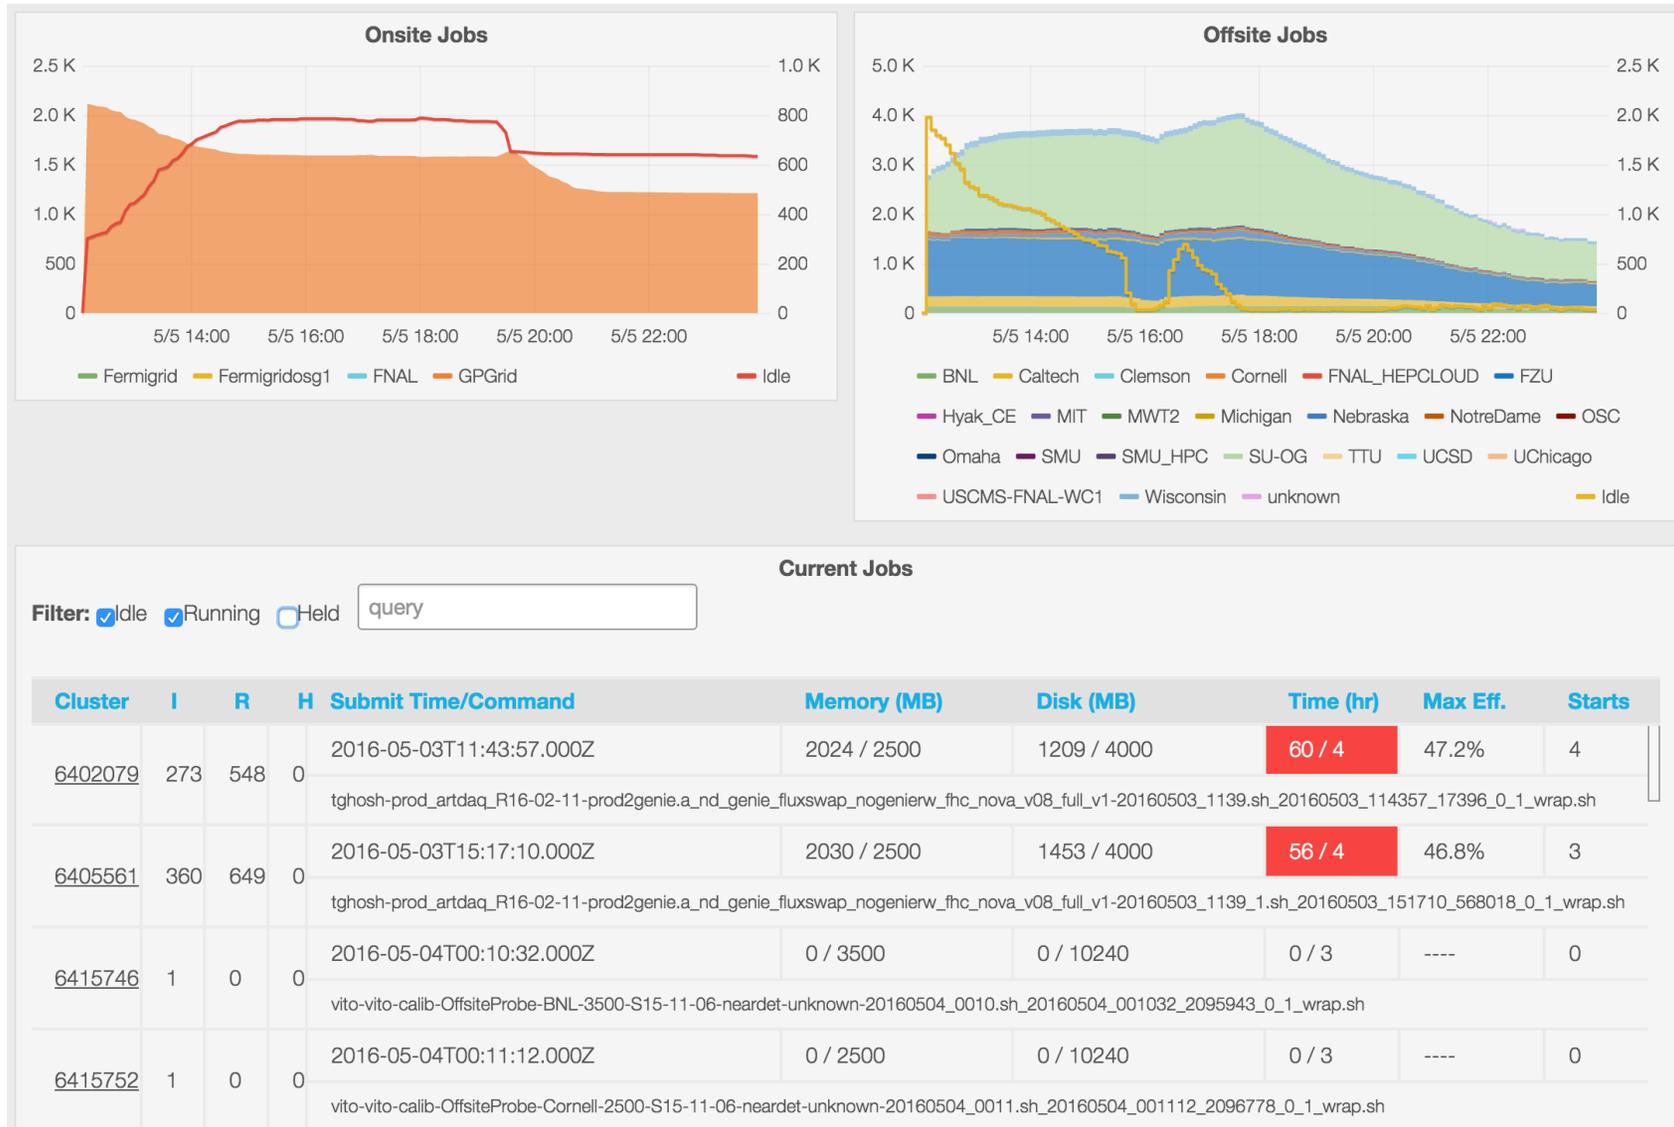
Monitor usage of slow moving resources so that projections can be made for projecting future need and limitations

Monitoring of jobs and experimental dashboards



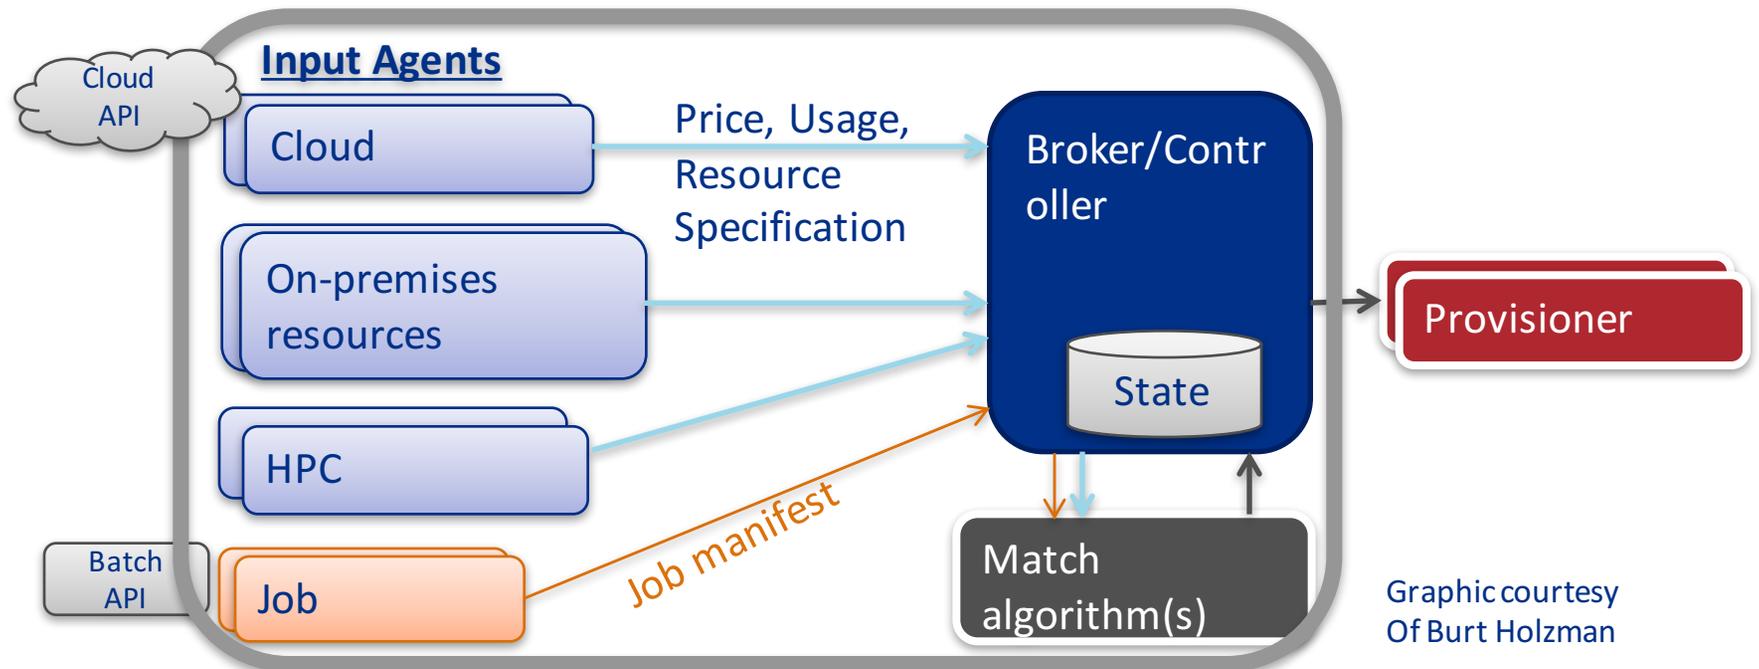
Monitoring for individual users to track their distributed computing workflows and understand their resource allocation and needs

Monitoring of jobs and experiment dashboards



Connection to HEPCloud and future services

- using the FIFE integrated experiments to test and build out the HEPCloud infrastructure
- gathering requirements and feedback from experiments for design of HEPCloud user facing tools
- NOvA running on AWS was very successful first step
- helping experiments be location agnostic aids in allowing for quick transition to HEPCloud resources and tools when available
- building monitoring tools with an eye towards HEPCloud and metrics needed as input to decision engine



FIFE Operations Summary

- FIFE has been successful helping experiments incorporate centralized services into computing operations
- Help to educate experiments and analyzers about best practices and security and how to utilize services
- Improve experiment efficiency by providing enhanced monitoring and alert them to potentially problematic workflows
- Fermilab General Purpose Grid being successfully utilized at near capacity by more than a dozen experiments
- FIFE experiment usage of offsite resources has increased drastically in the last two years
- Preparing to integrate and transition experiments to HEPCloud operation
- Continue to provide conduit for improved services, feature development, and connection with other laboratories

