Monitoring Large-Scale Scientific Computing with Grafana

Kevin Retzke

GrafanaCon

November 30 2016

Disclaimer

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof.

Fermilab

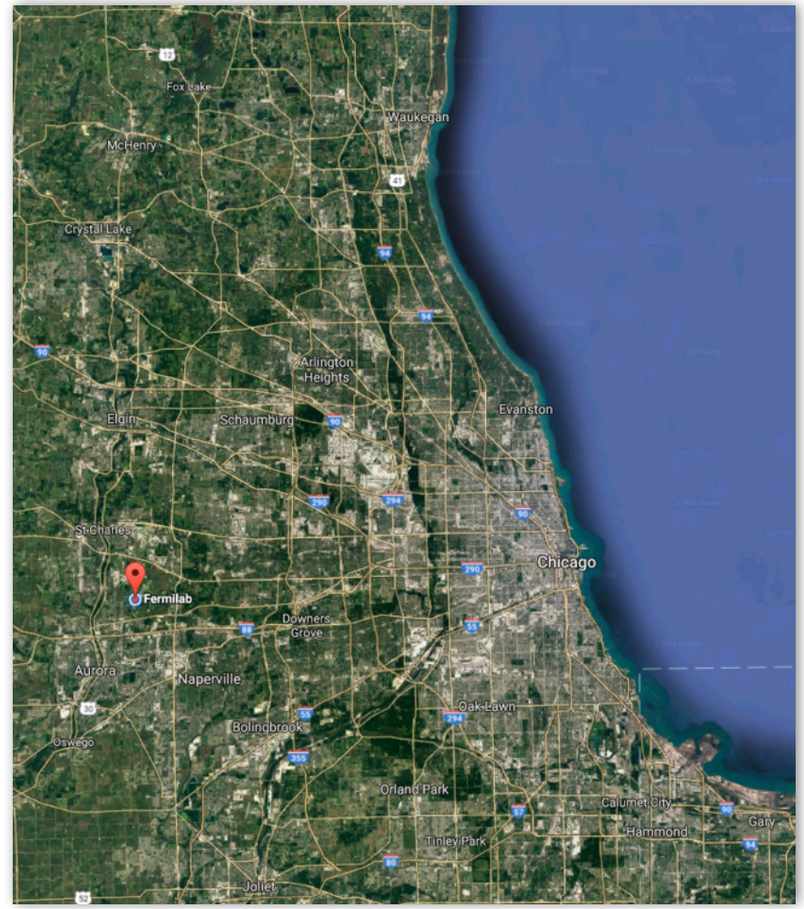

Fermi National Accelerator Laboratory was founded 50 years ago (Dec 16!) on a 6,800-acre site 40 miles west of Chicago in Batavia, Illinois.

Its mission is to solve the mysteries of matter, energy, space, and time for the benefit of all.

Fermilab - Open Science

More than 3,500 scientists worldwide use Fermilab and its accelerators, detectors, and computers for their research. About 2,600 researchers from 44 countries collaborate on experiments at Fermilab.

Fermilab - Accelerators

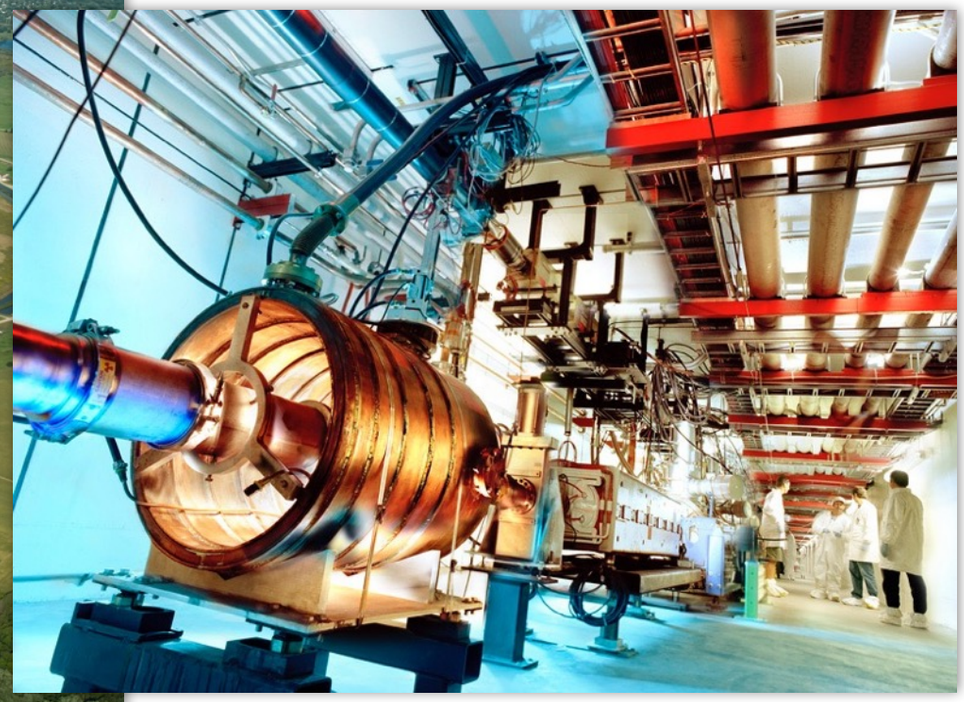

Fermilab produces the world's most intense beam of high-energy neutrinos, particles that may hold the key to understanding why the universe is made of matter.

Fermilab - NOvA

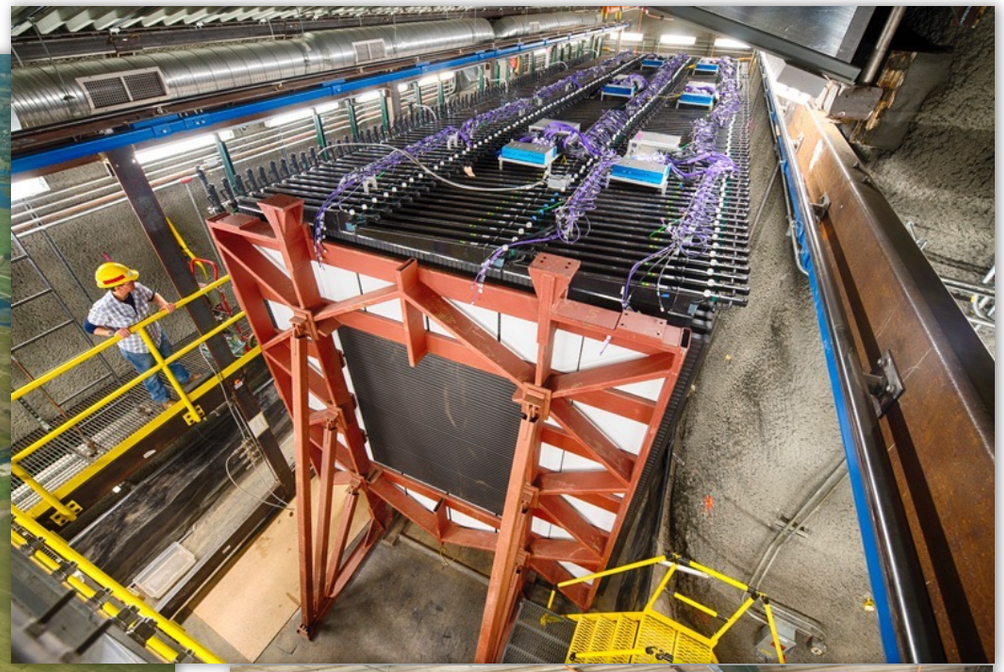

The NOvA experiment will help answer some of the most important scientific questions about neutrino masses, neutrino oscillations, and the role neutrinos may have played in the evolution of the universe.

Fermilab - CMS

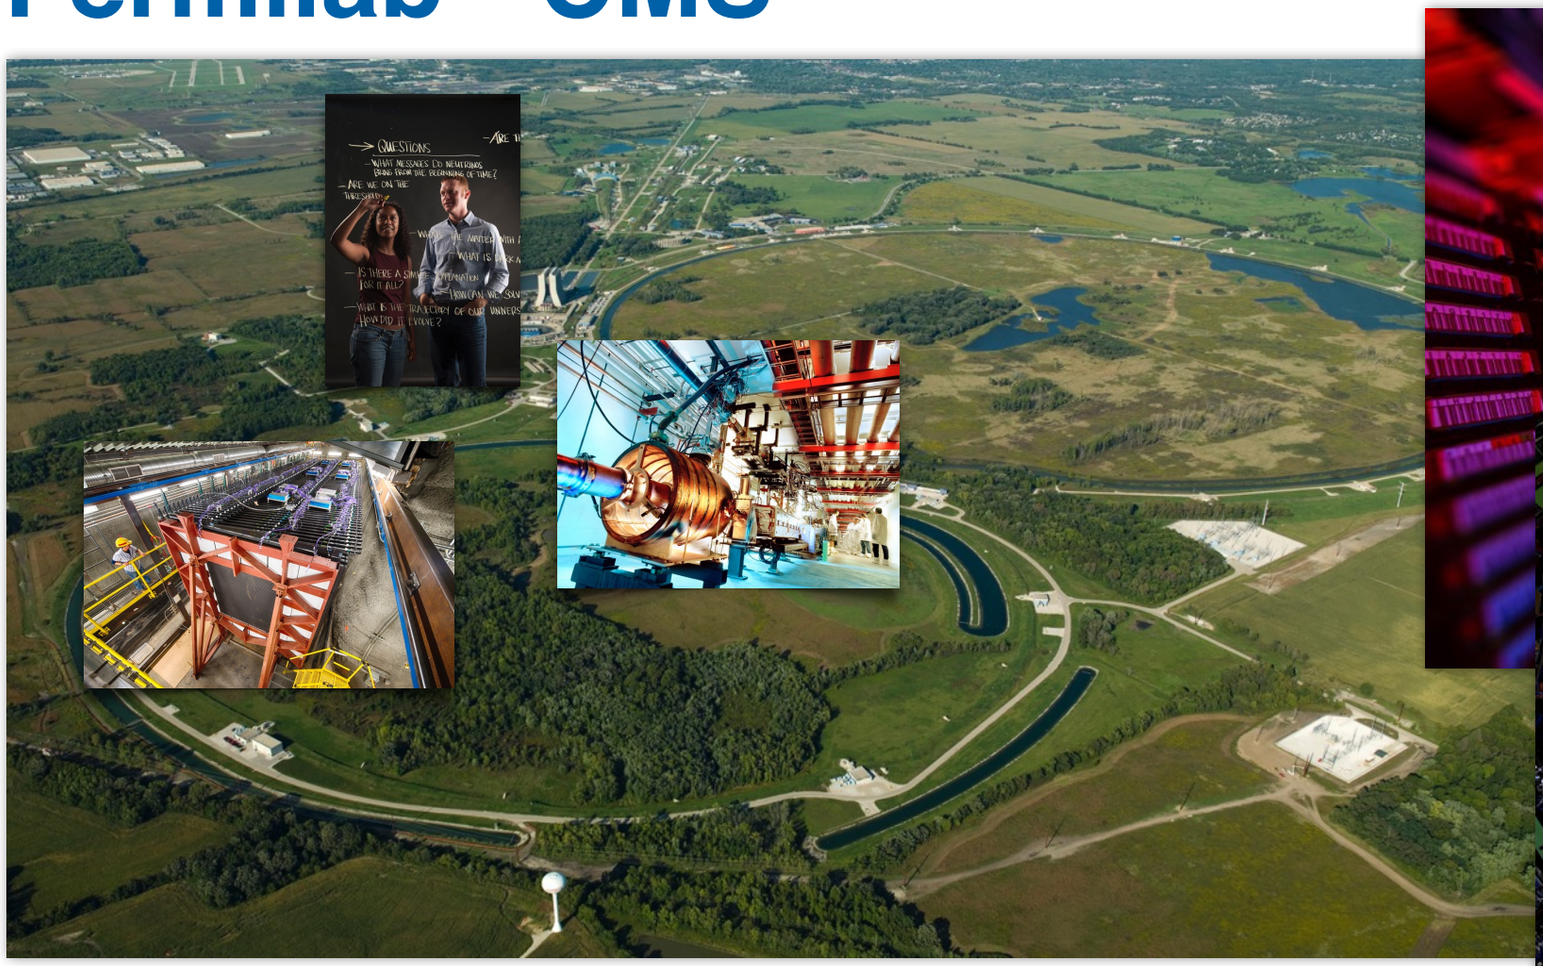

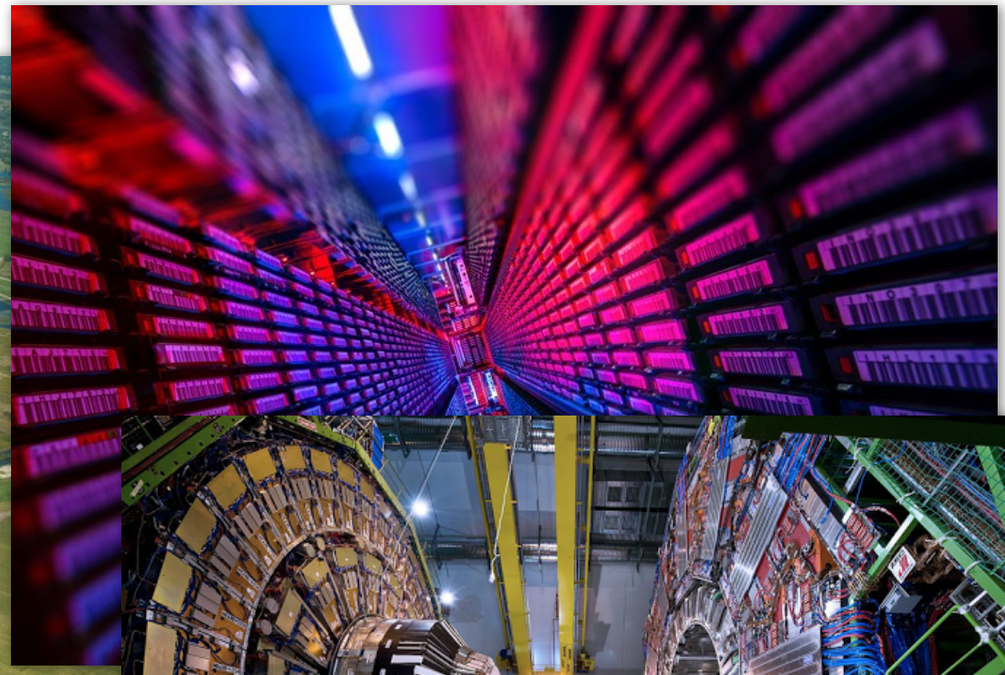

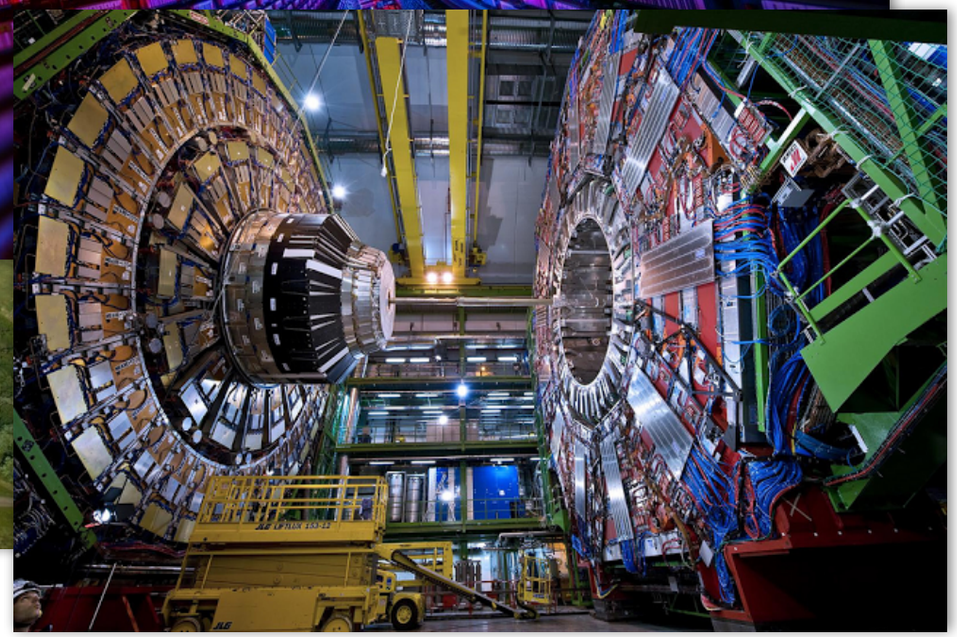

The LHC at CERN is the world's highest-energy particle collider and enabled the discovery of the Higgs particle in 2012. Fermilab houses an LHC Remote Operations Center, provides a quarter of the computing power for the CMS experiment and designs and builds components for upgrades to the LHC and CMS.

Fermilab - DES

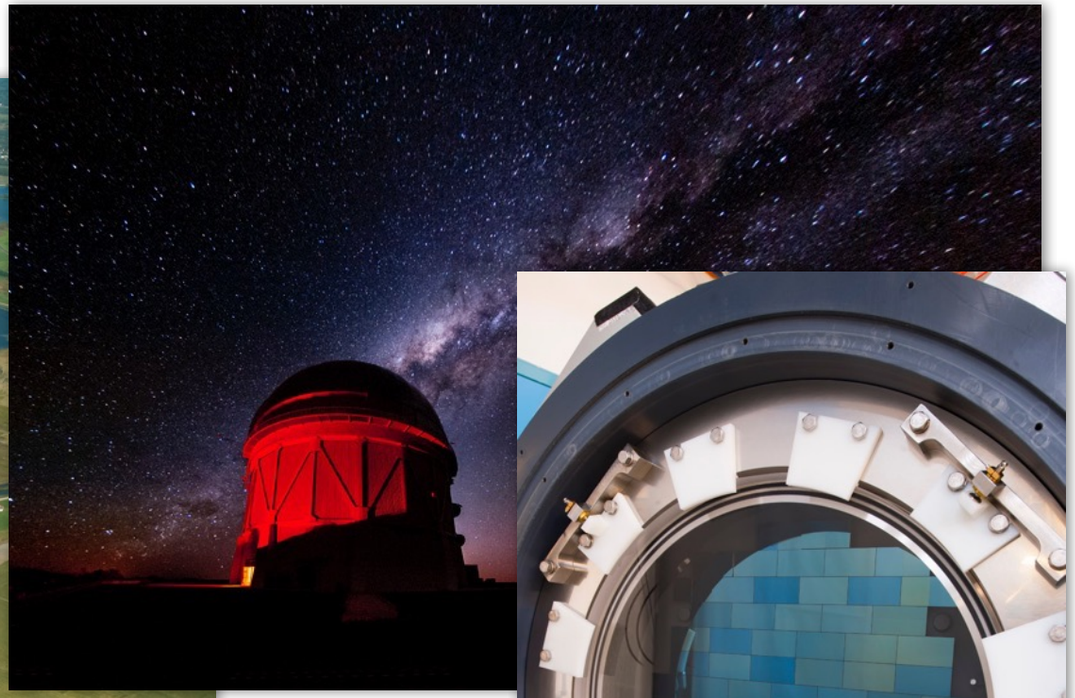

The Dark Energy Camera, designed and built at Fermilab, is one of the world's most powerful cameras (570-megapixel), and now takes images on a telescope in Chile. The heart of the Dark Energy Survey, it advances the quest to understand the nature of the dark energy that pushes the universe apart.

Fermilab - D’awww

Fermilab’s herd of bison is part of an effort to restore and protect the native Illinois prairie. And they’re SOOO CUTE! (and tasty)

Fermilab

Fermilab - Science!

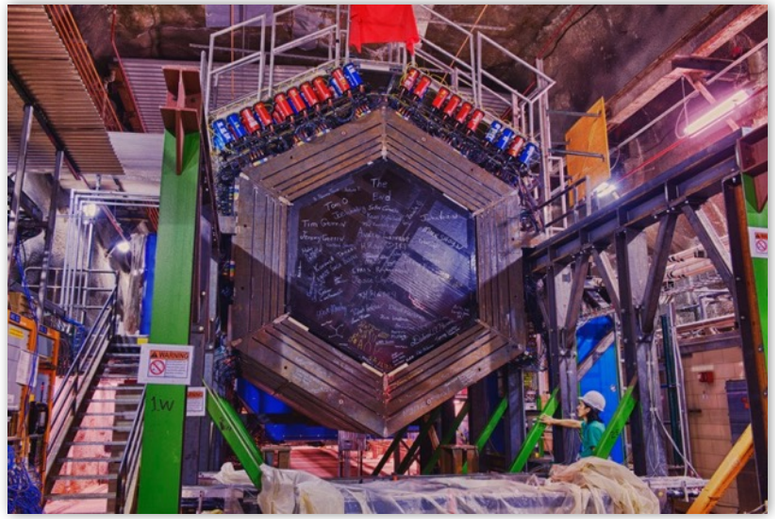

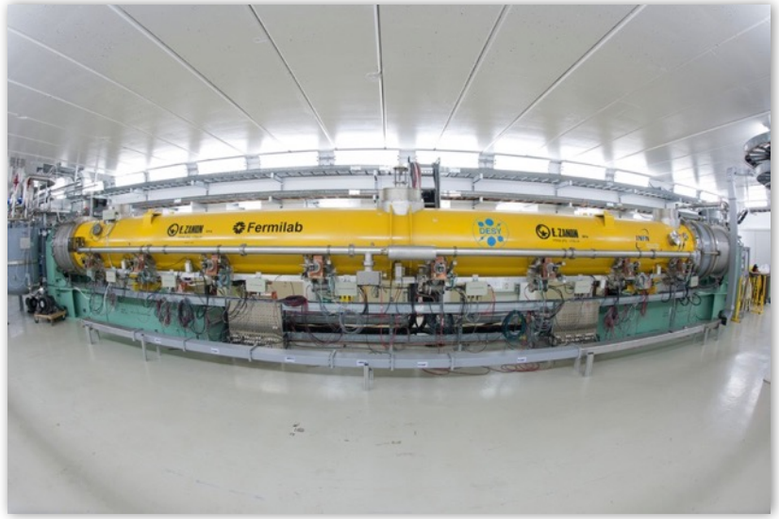

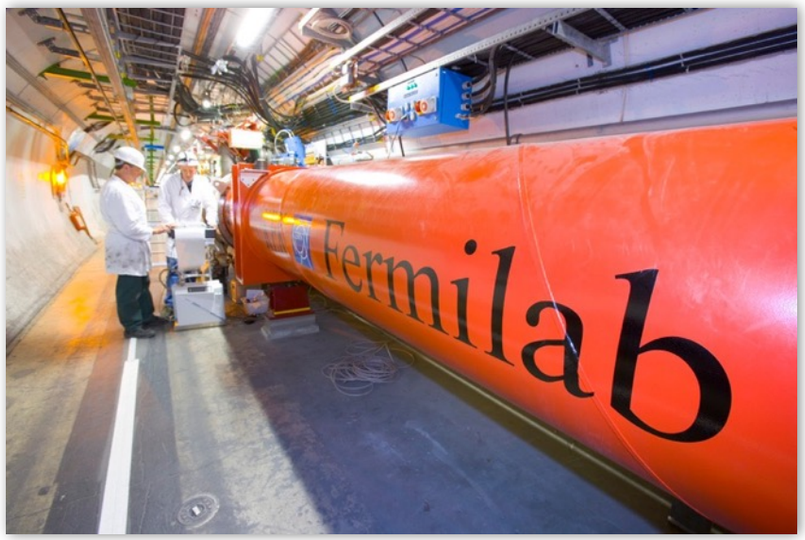

Fermilab

DUNE - The Next Mega-Science Project

https://youtu.be/AYtKcZMJ_4c

Scientific Computing

Scientific Computing @ Fermilab

High energy physics requires a huge amount of computing

- Detector and accelerator design and simulation
- Data collection and reconstruction
- Event analysis

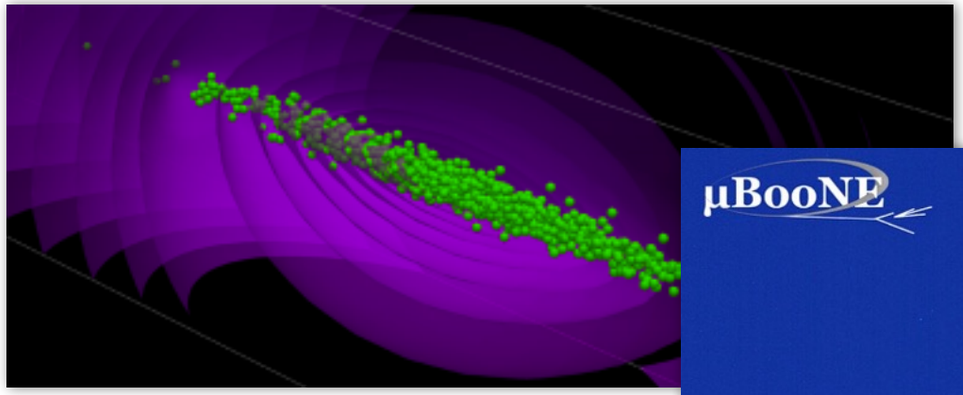

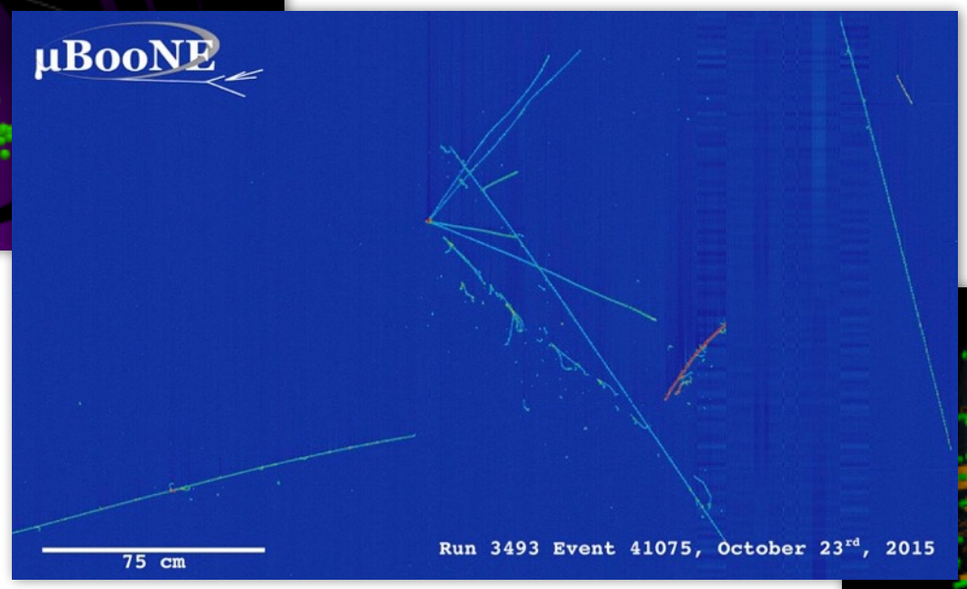

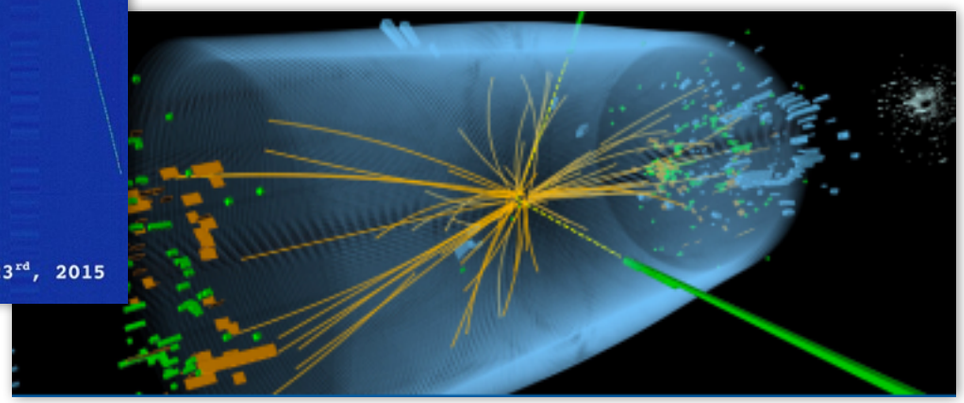

Fermilab

Scientific Computing @ Fermilab

Fermilab Computing Sector provides scientists with the computing resources and support they need

- Three on-site data centers
- Access to dozens of computing sites world-wide
  - connected by 100 Gbps dedicated network (ESNet)
- Scientific toolkits, frameworks, and services

Fermilab

Batch/Grid Computing - HTC and HPC

General calculations are handled by our high-throughput (HTC) batch systems

- HTCondor cluster on bare metal (commodity x86), plus GlideinWMS-driven virtual clusters (also HTCondor)
- Running millions of batch jobs per week, processing petabytes of data

HPC/GPU systems, primarily dedicated to Lattice QCD calculations

- Little siblings to leadership-class supercomputers, e.g. Mira BlueGene/Q at Argonne (#9 Top500, Nov 2016)

Batch/Grid Computing - FIFE and CMS

Two main HTC systems:

- CMS Tier 1 - dedicated computing for CMS experiment
- FIFE (Fabrlc for Frontier Experiments) - everyone else

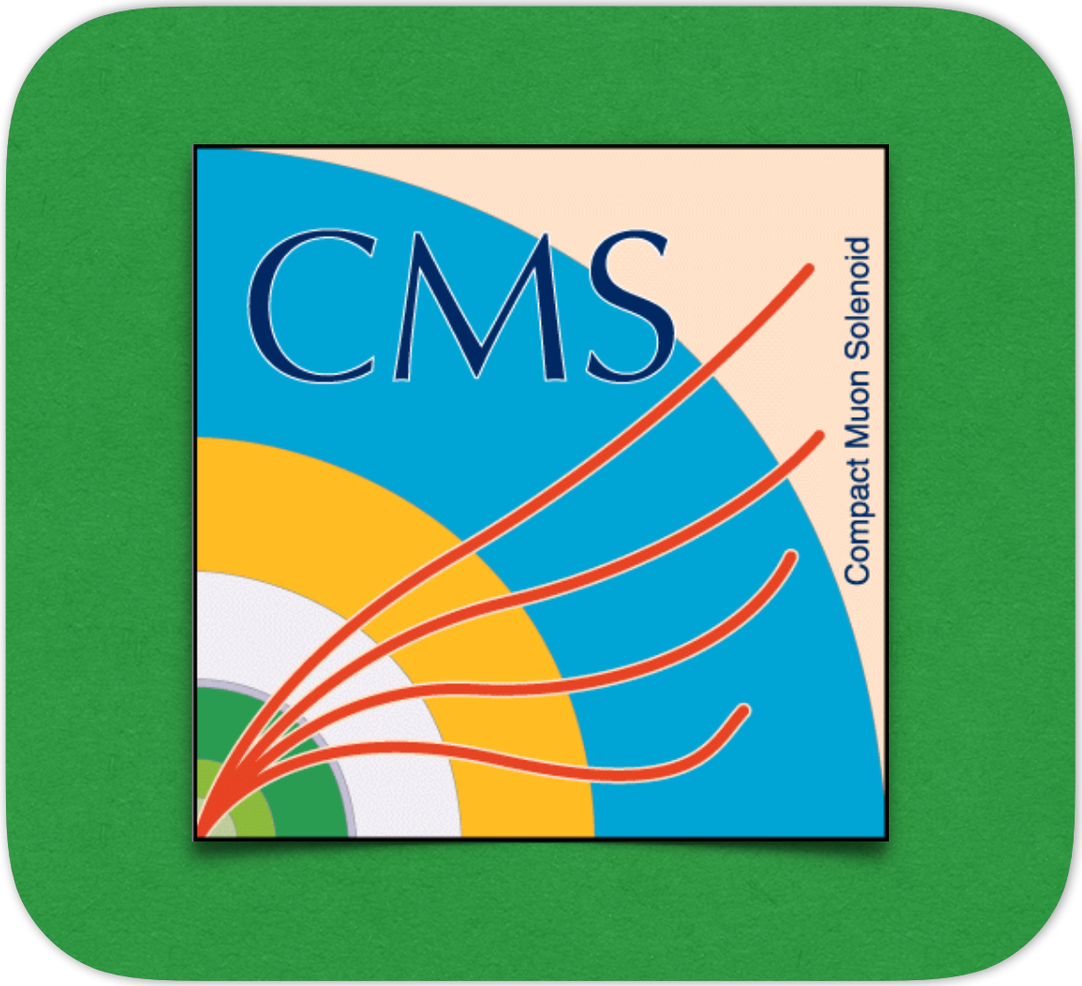

FIFE Batch System Architecture

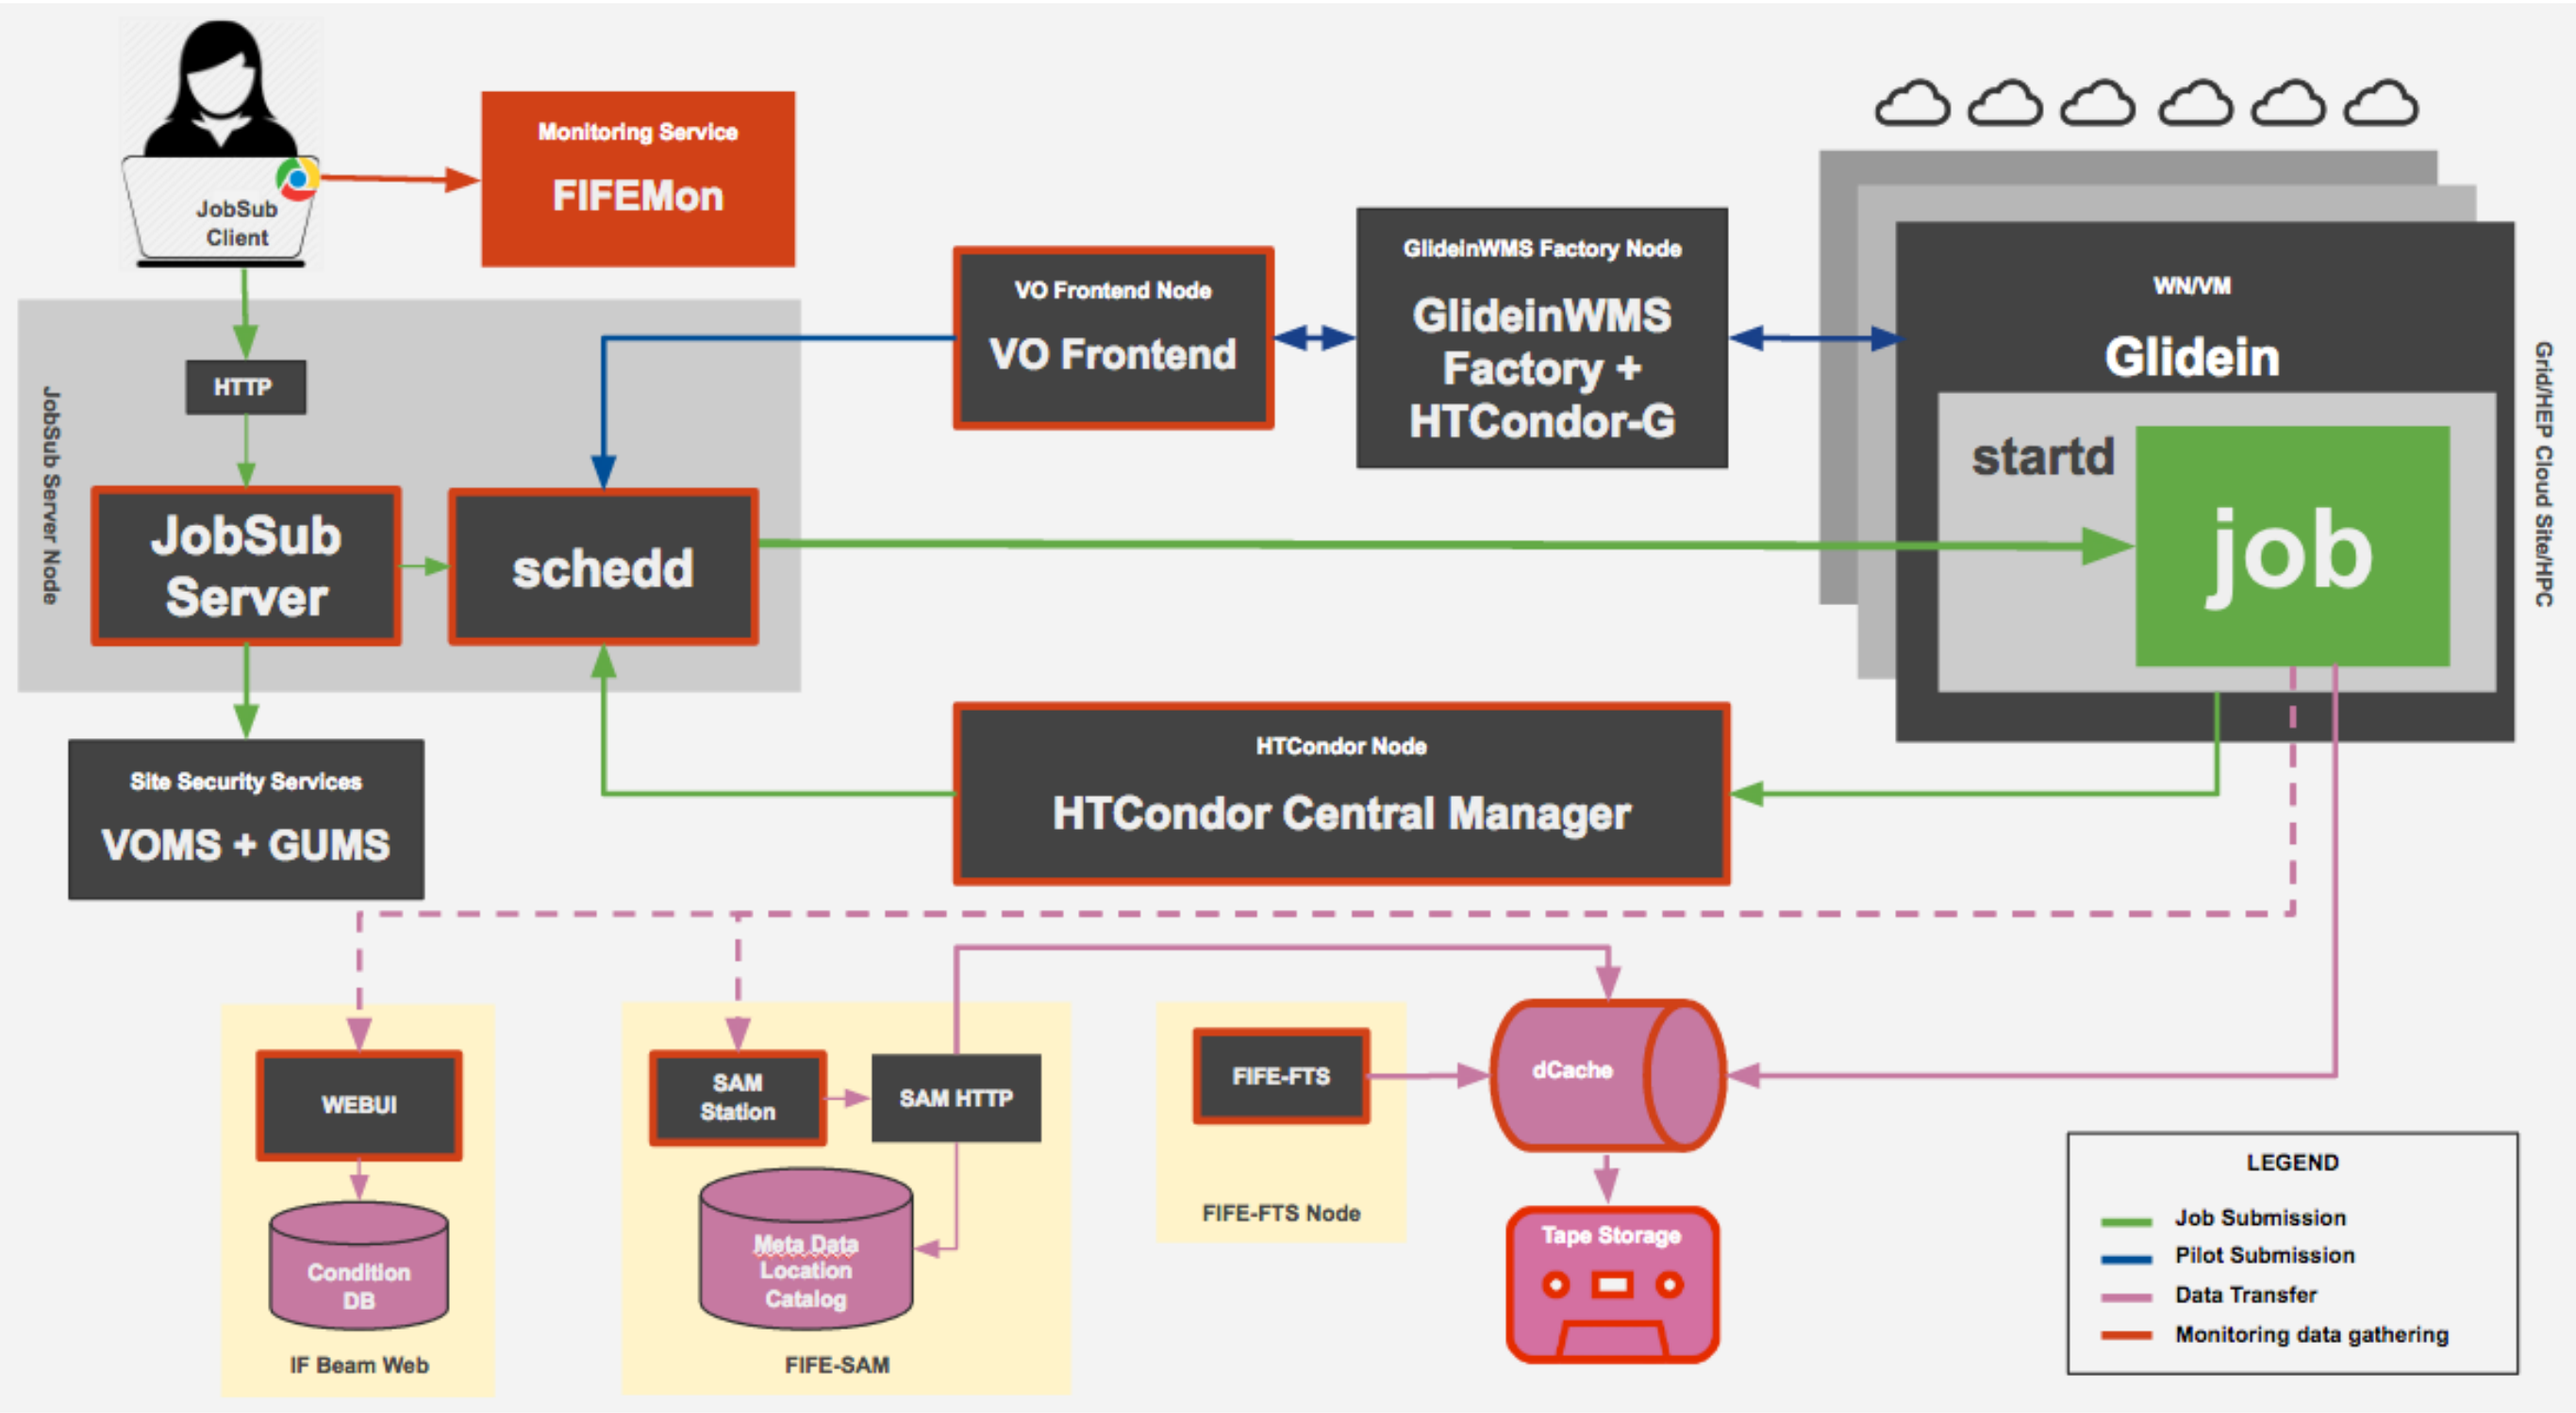

Fermilab

Monitoring as it was…

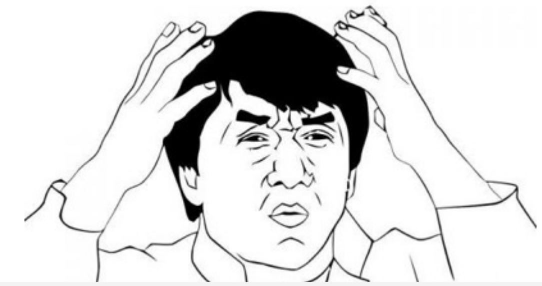

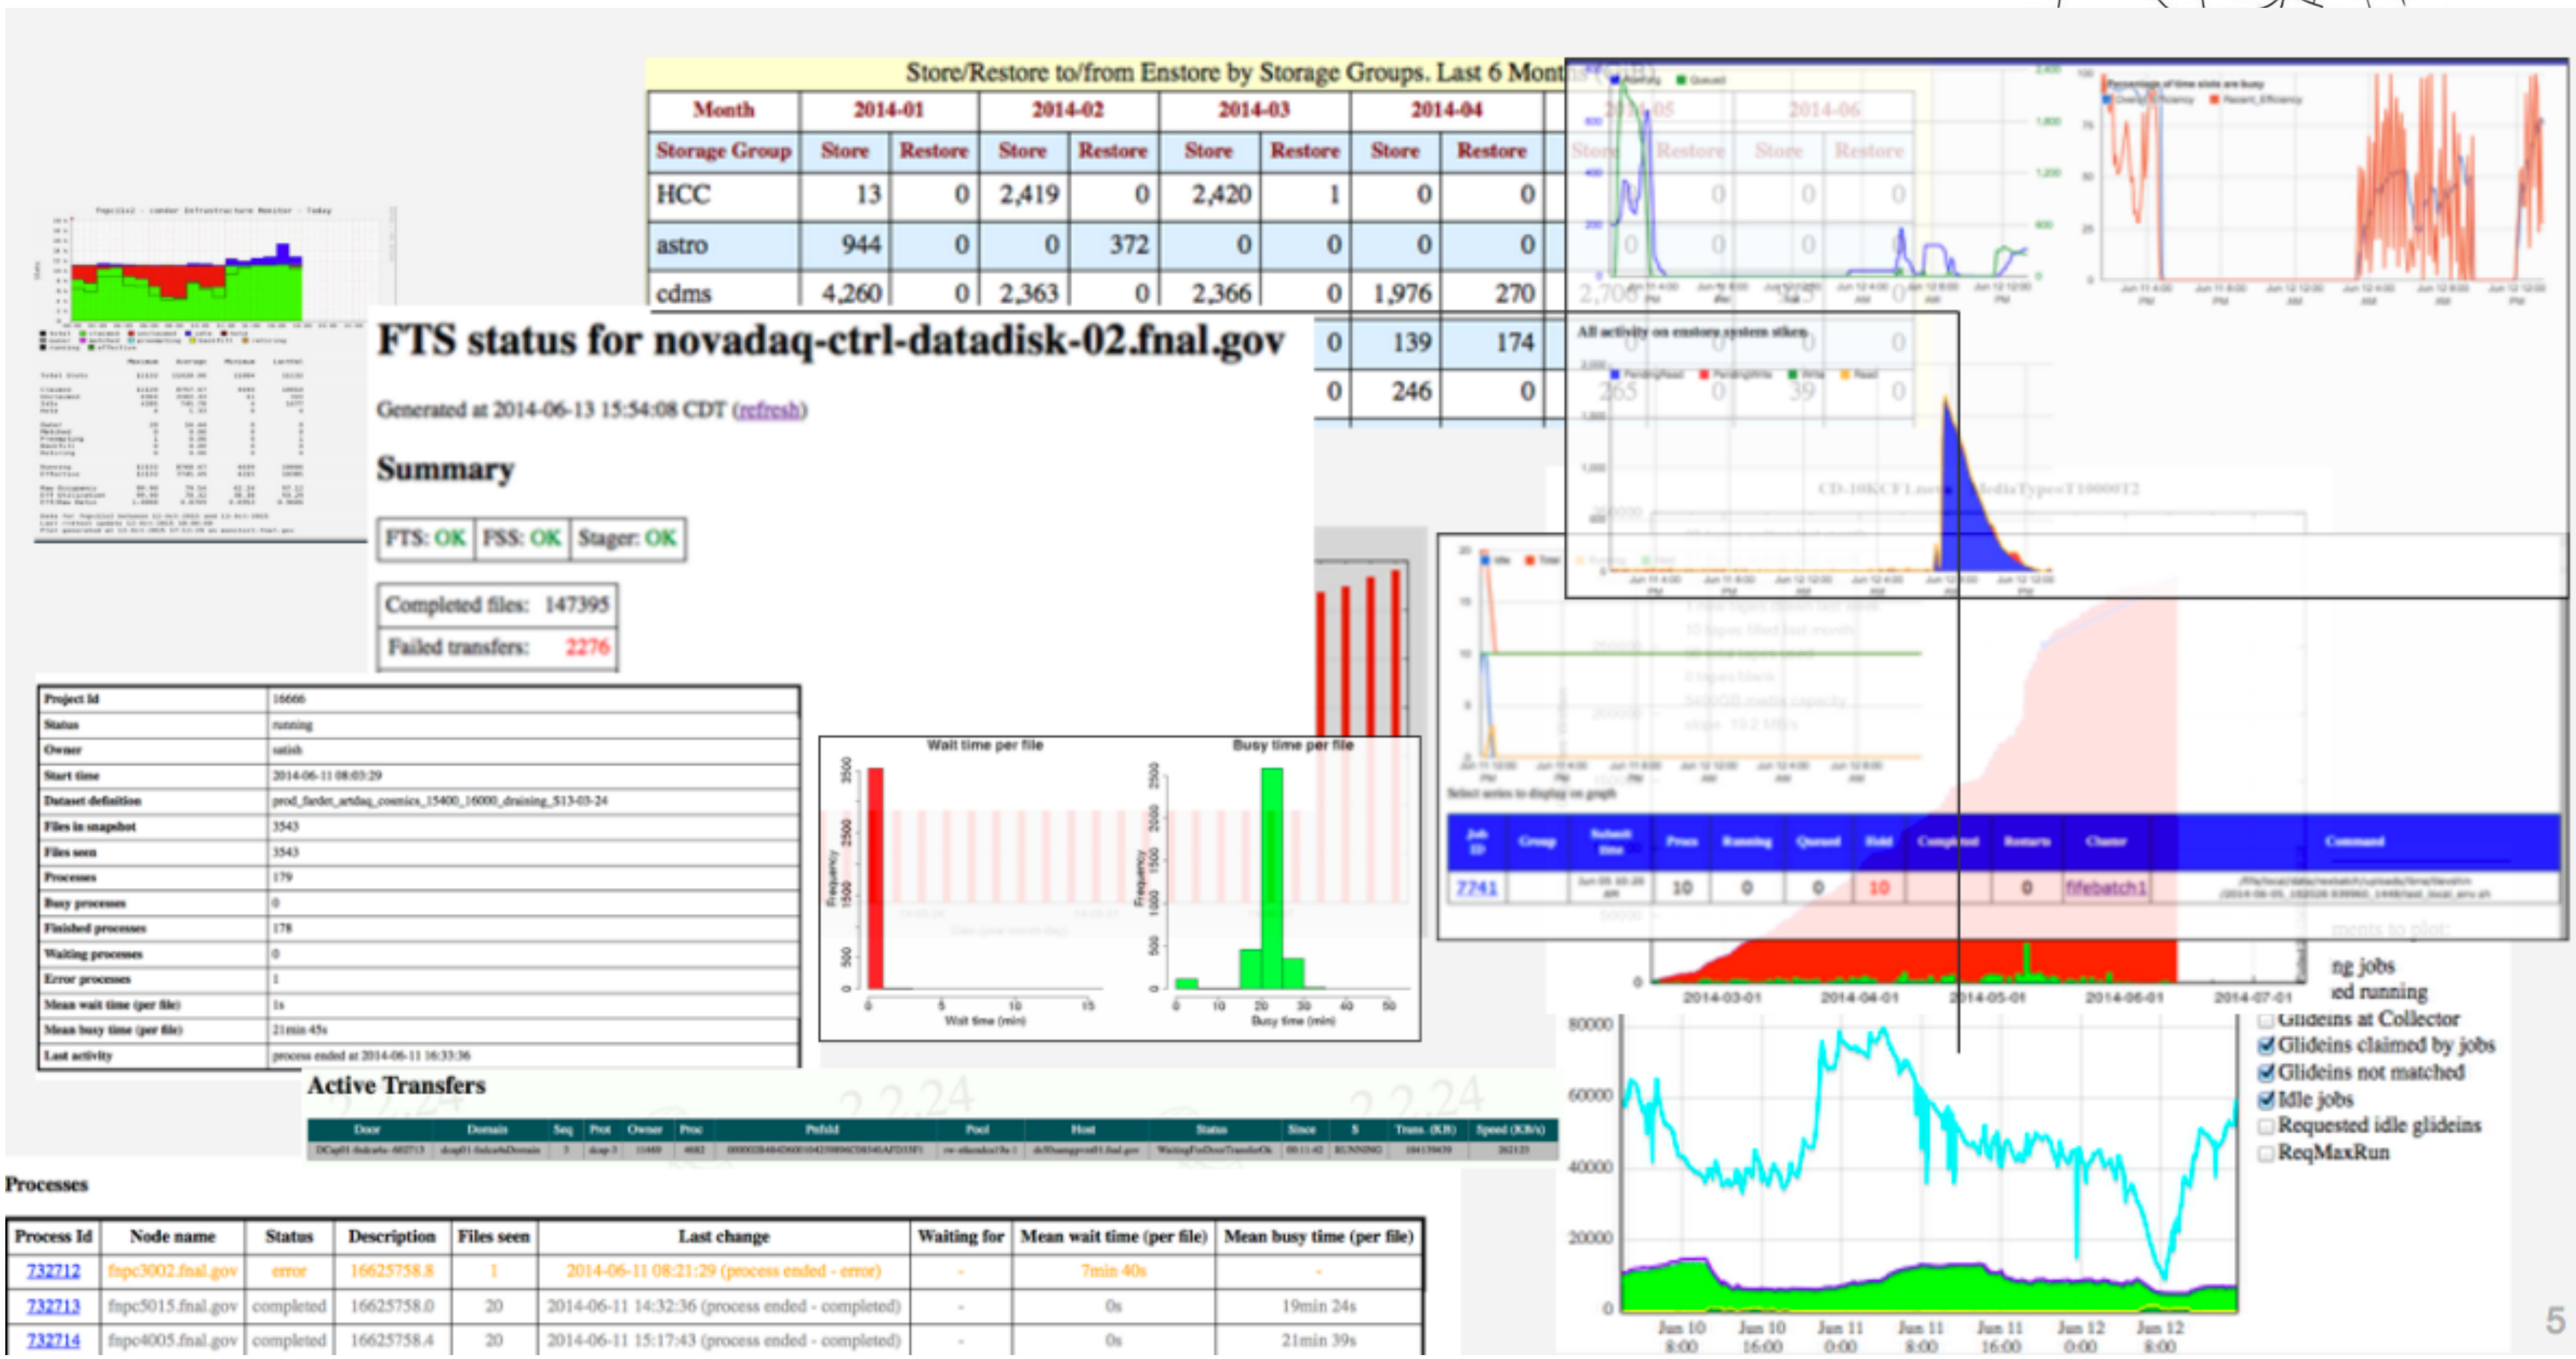

Fifemon & Landscape

Fifemon is a comprehensive monitoring platform for all FIFE experiments, services, and stakeholders.

Fermilab

Fifemon Architecture & Stats

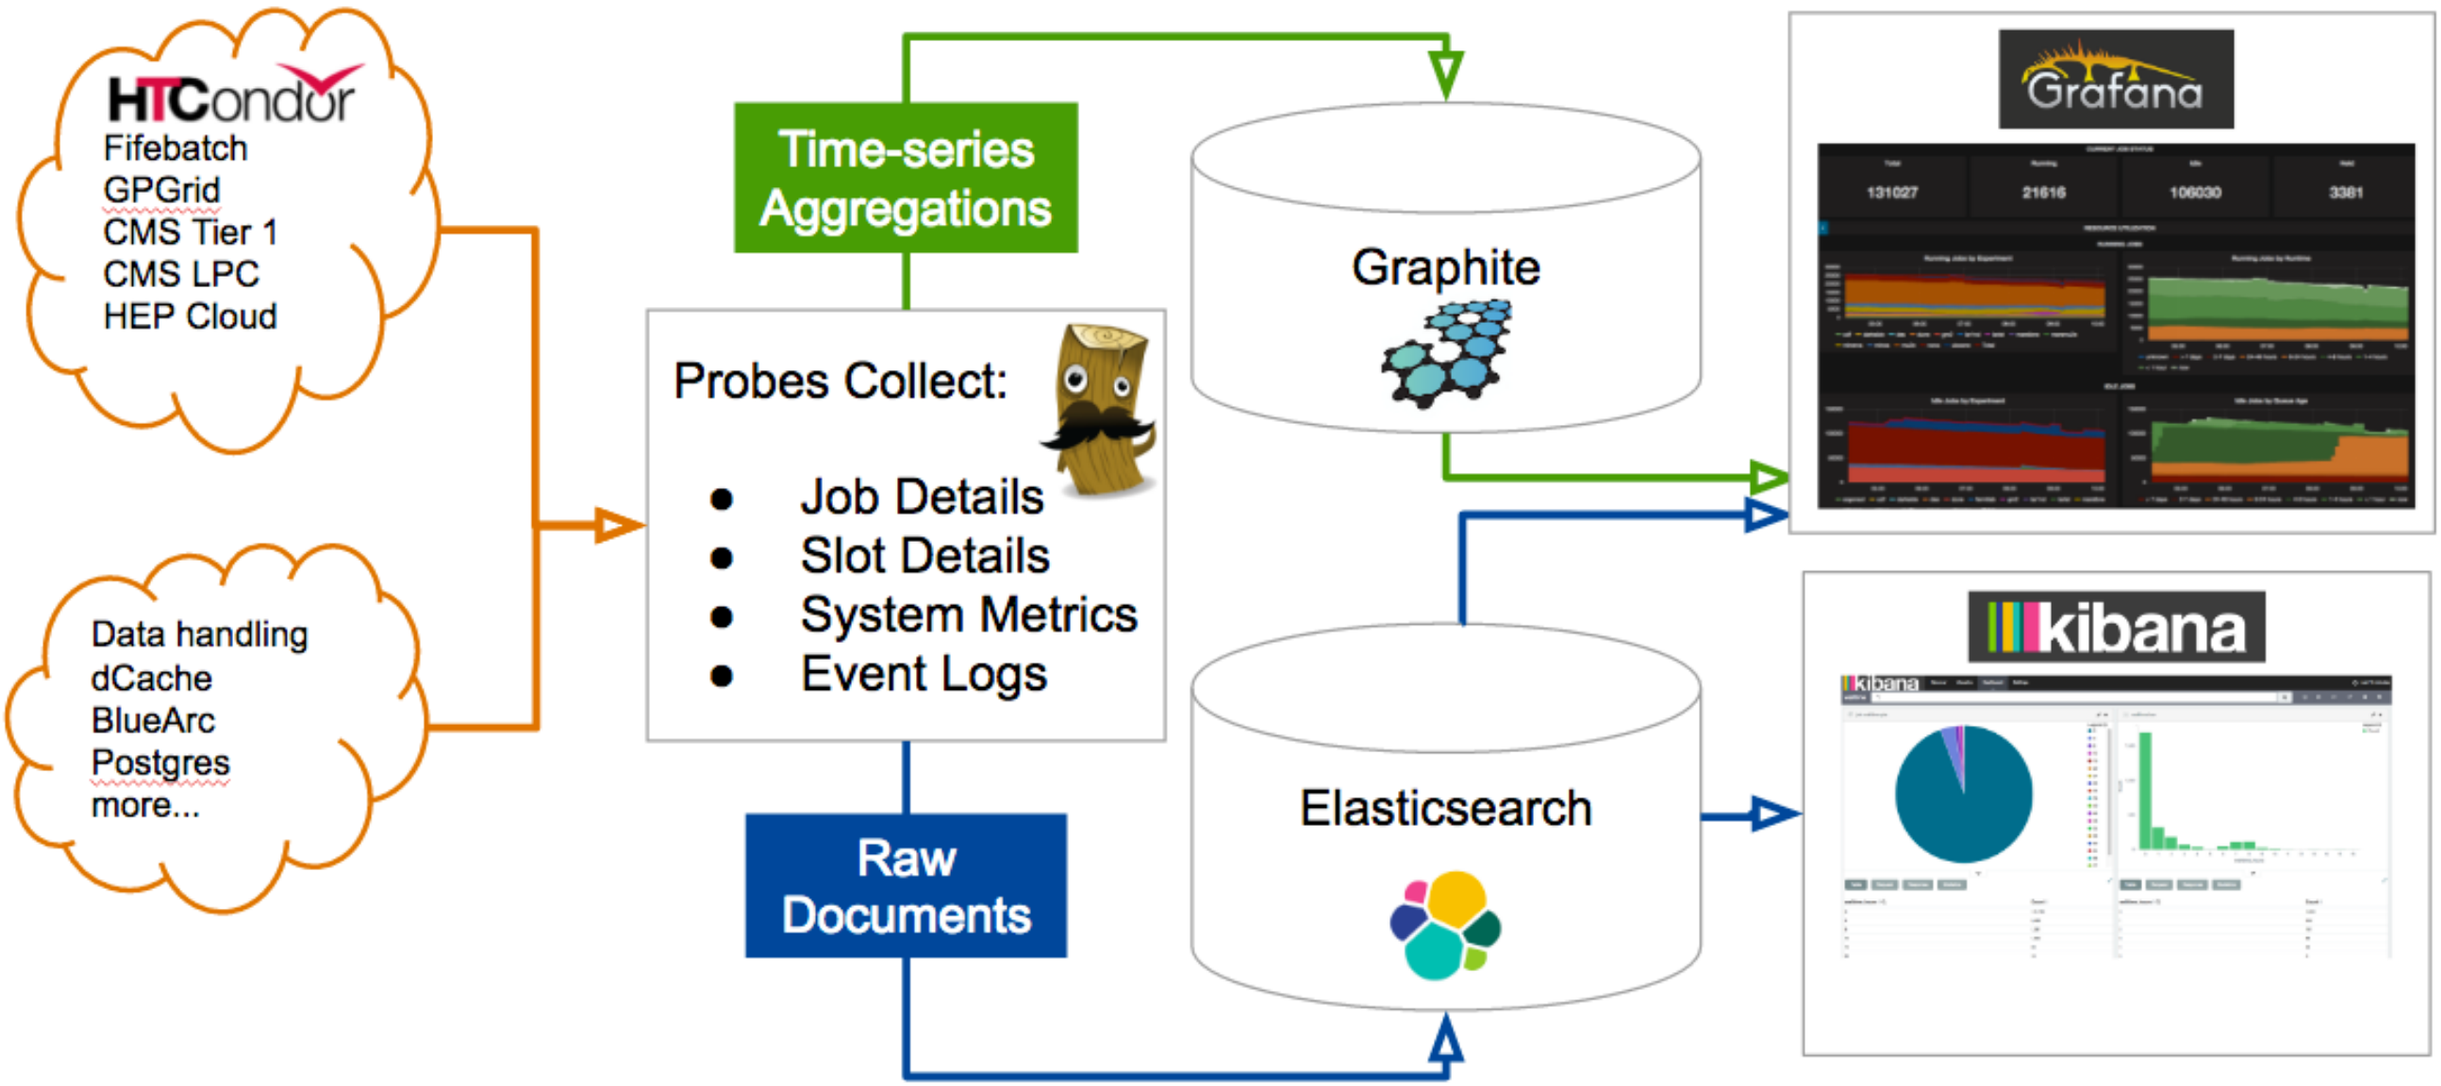

Graphite
1 node + replica
~500K total metrics
~1.5M points per hour

Elasticsearch
6 VM nodes
650 indices, 1.3Bil docs
20 TB Ceph storage

Grafana
400 users
50 dashboards
~1 dashboard load/min

Fifemon - FIFE Summary

Landscape

Fermilab Scientific Computing Summary
FIFE

Fermilab

Average Number of Jobs Running Concurrently

15511

Total Jobs Run

1199091

Running Jobs by Experiment (includes Onsite, OSG & Cloud)

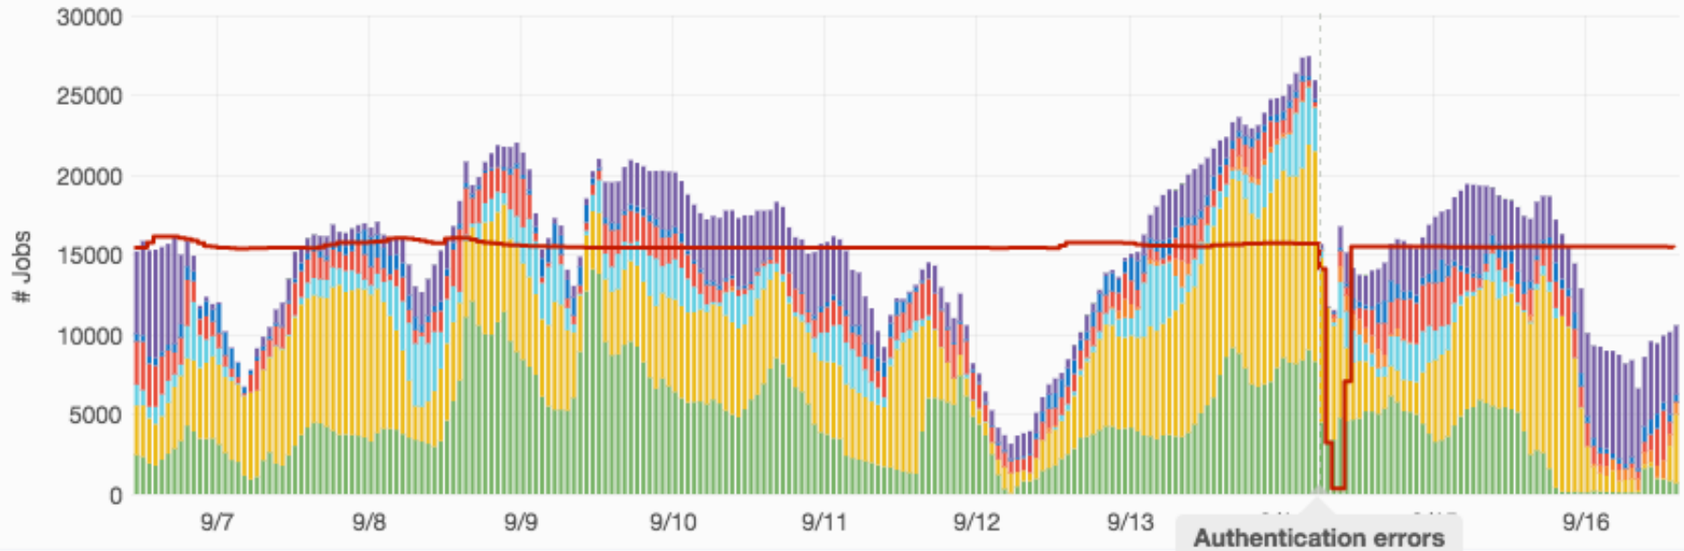

	min	max	avg
NOvA	56	14083	4561
Mu2e	64	13176	5514
MINERvA	0	4415	1215
MINOS	0	1668	133
DUNE	13	3934	1329
MicroBooNE	26	1884	544
DES	0	29	2
Other Experiments	29	6783	2200
Projects	0	42	3
Onsite Slots (GPGrid)	347	16145	15349

Percent Jobs Run Onsite

57.5%

Percent Jobs Ru

N/A

Percent Jobs Run on OSG

42.8%

NOTES

Time	Description	Details
2016-09-14 06:00:00	Authentication errors	GUMS authentiation was failing

Fifemon - Experiment Summary

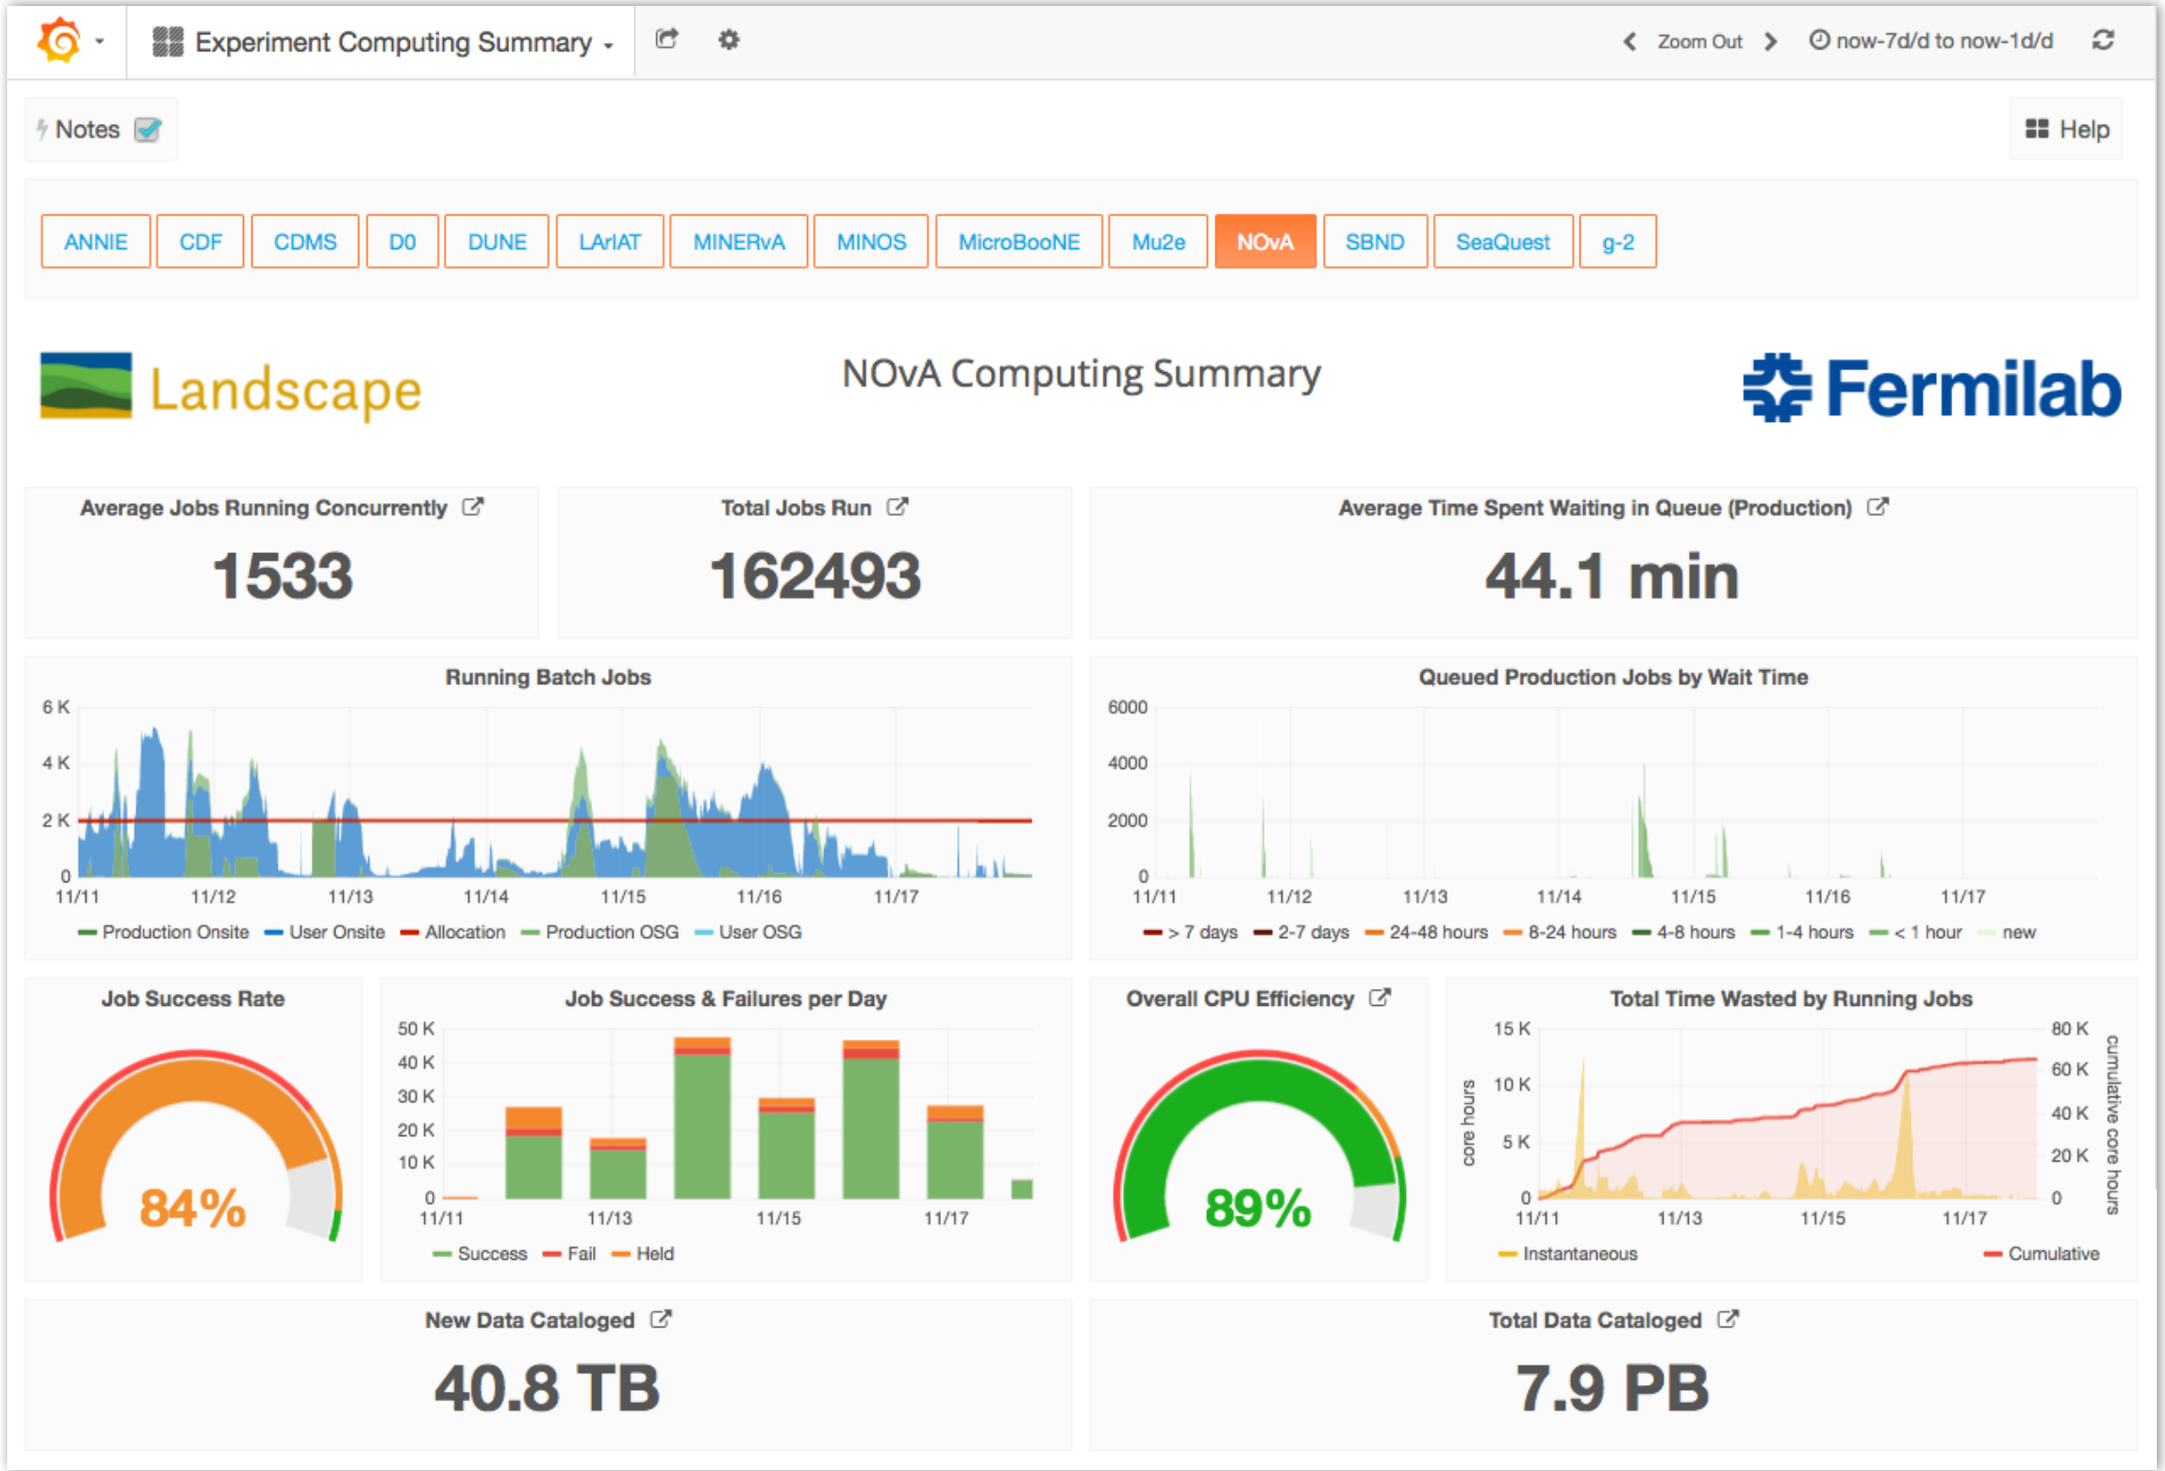

Fifemon - Experiment Batch Details

Experiment Batch Details | Zoom Out | now-24h to now-5m

minerva | GPGrid Usage | Experiment Efficiency Details | Experiment Overview | Experiment Overview Copy | FTS | SAM by experiment

User Jobs

User	I	R	C	X	H	Max Memory/Request	Max Disk/Request	Max Time/Request
drut1186	0	0	0	0	29	1.01	0.19	0.00
minervacal	1	1	0	0	0	0.29	0.10	0.30
minervadat	0	2	0	0	6	0.20	0.05	0.45
minervapro	0	8	0	0	17	1.03	1.04	0.38
oaltinok	0	83	0	0	0	0.48	0.10	0.15
wospakrk	0	0	0	0	1	0.79	0.15	0.00

Recent Jobs Started

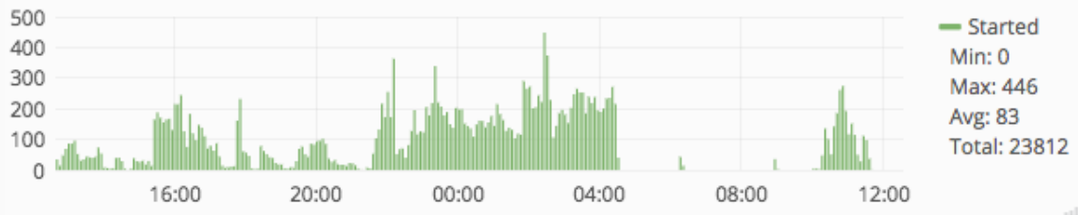

Recent Jobs Submitted

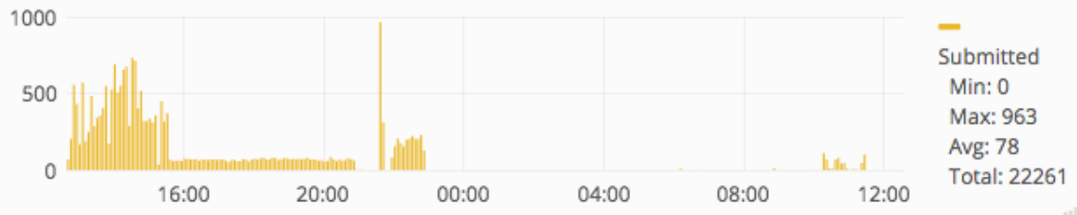

Onsite Running Jobs by User & GPGrid Quota

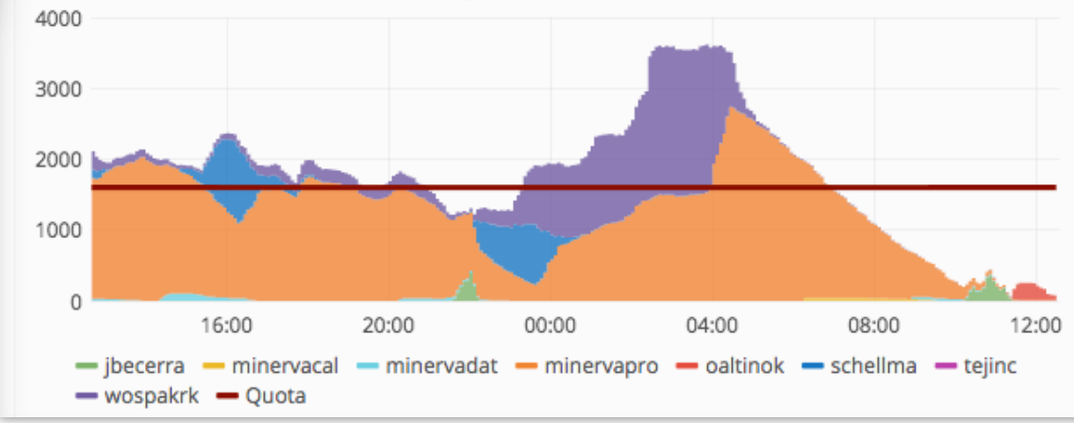

Onsite Idle Jobs by Queue Age

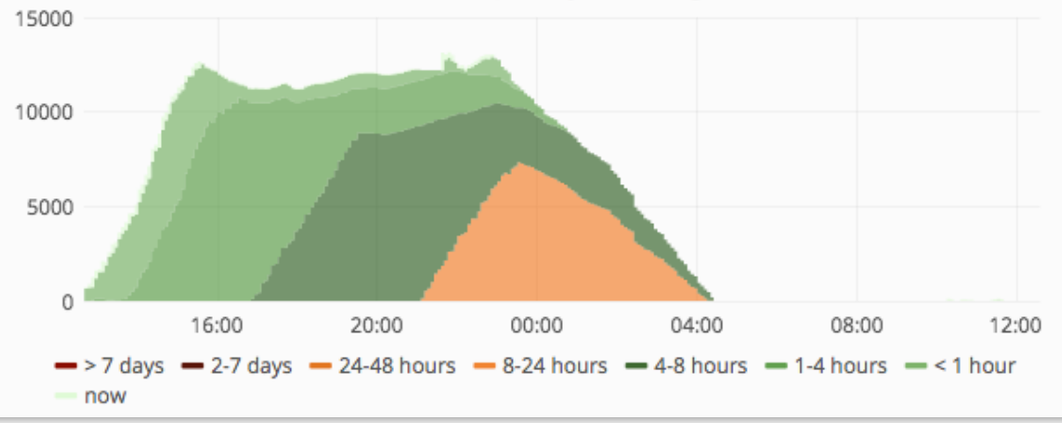

Fifemon - User Batch Details

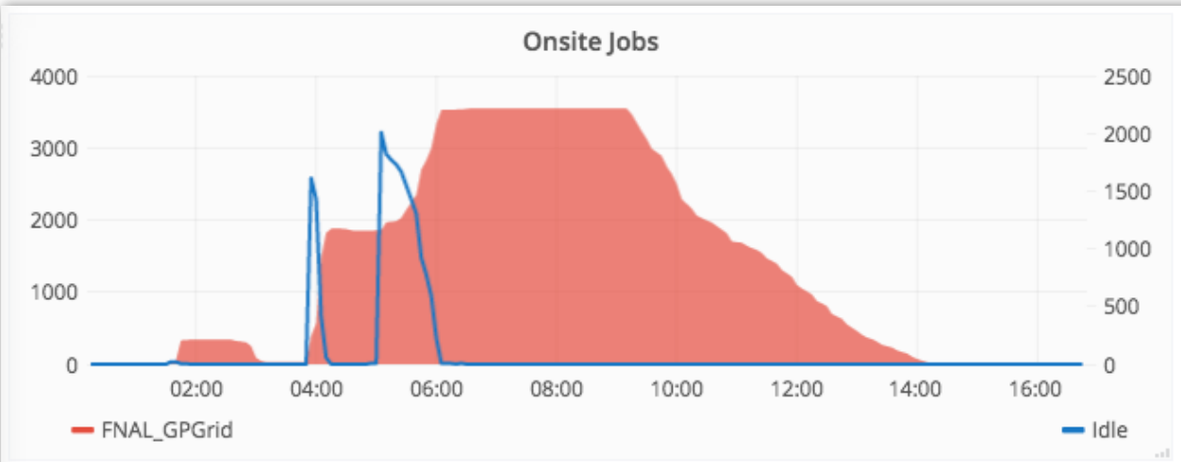

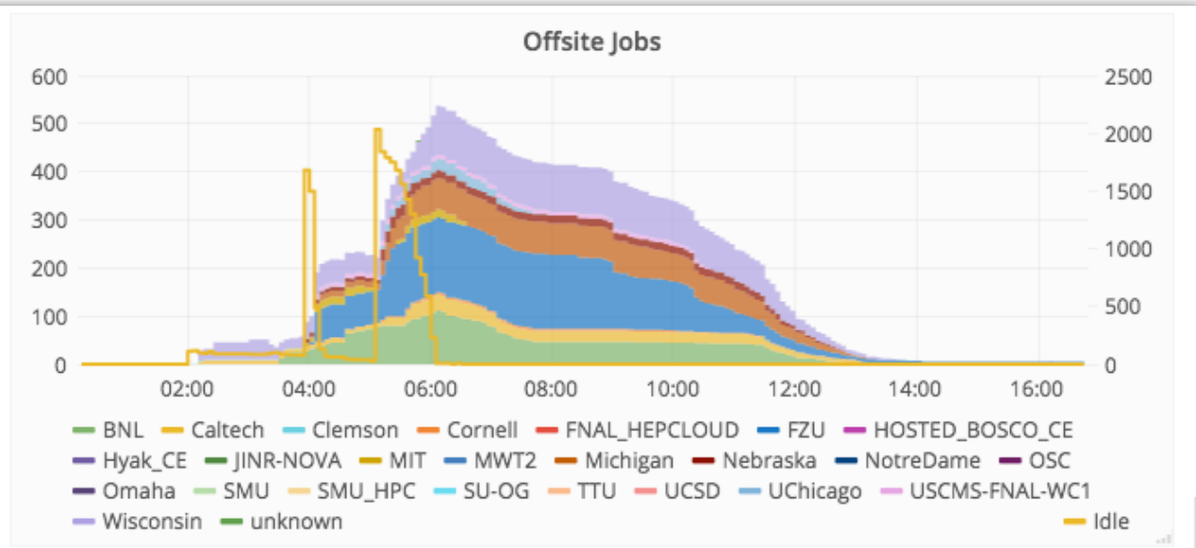

Current Job Clusters

2799 jobs in 21 clusters.

Filter -- no filter --

Cluster	I	R	H	Submit Time	Memory (MB)	Disk (MB)	Time (hr)	Max Eff.	Starts
14806744@fifebatch2.fnal.gov	0	0	2404	Mon Nov 14 13:00 -0600	689 / 1950	2219 / 10000	0 / 0	----	1
tghosh-prod_calibration_S15-03-11_fd_cosmic_preshutdown_draining-20161114_1256.sh_20161114_130049_2693892_0_1_wrap.sh									
11868298@fifebatch1.fnal.gov	0	0	3	Wed Nov 16 17:58 -0600	0 / 500	0 / 1024	0 / 1	----	0
offsite_test_JINR.sh_20161116_175833_1168516_0_1_wrap.sh									
11875126@fifebatch1.fnal.gov	0	0	3	Wed Nov 16 22:09 -0600	0 / 500	0 / 1024	0 / 1	----	0
offsite_test_JINR.sh_20161116_220946_2202517_0_1_wrap.sh									
14927498@fifebatch2.fnal.gov	0	0	6	Thu Nov 17 18:54 -0600	7069 / 7000	915 / 5000	0 / 0	----	1
rijeeshk-horn_off_data_reco-20161117_1853.sh_20161117_185450_1430195_0_1_wrap.sh									
11899167@fifebatch1.fnal.gov	0	0	7	Fri Nov 18 02:10 -0600	2114 / 2000	172 / 10240	0 / 15	----	3

Fifemon - Grid/Cluster Utilization

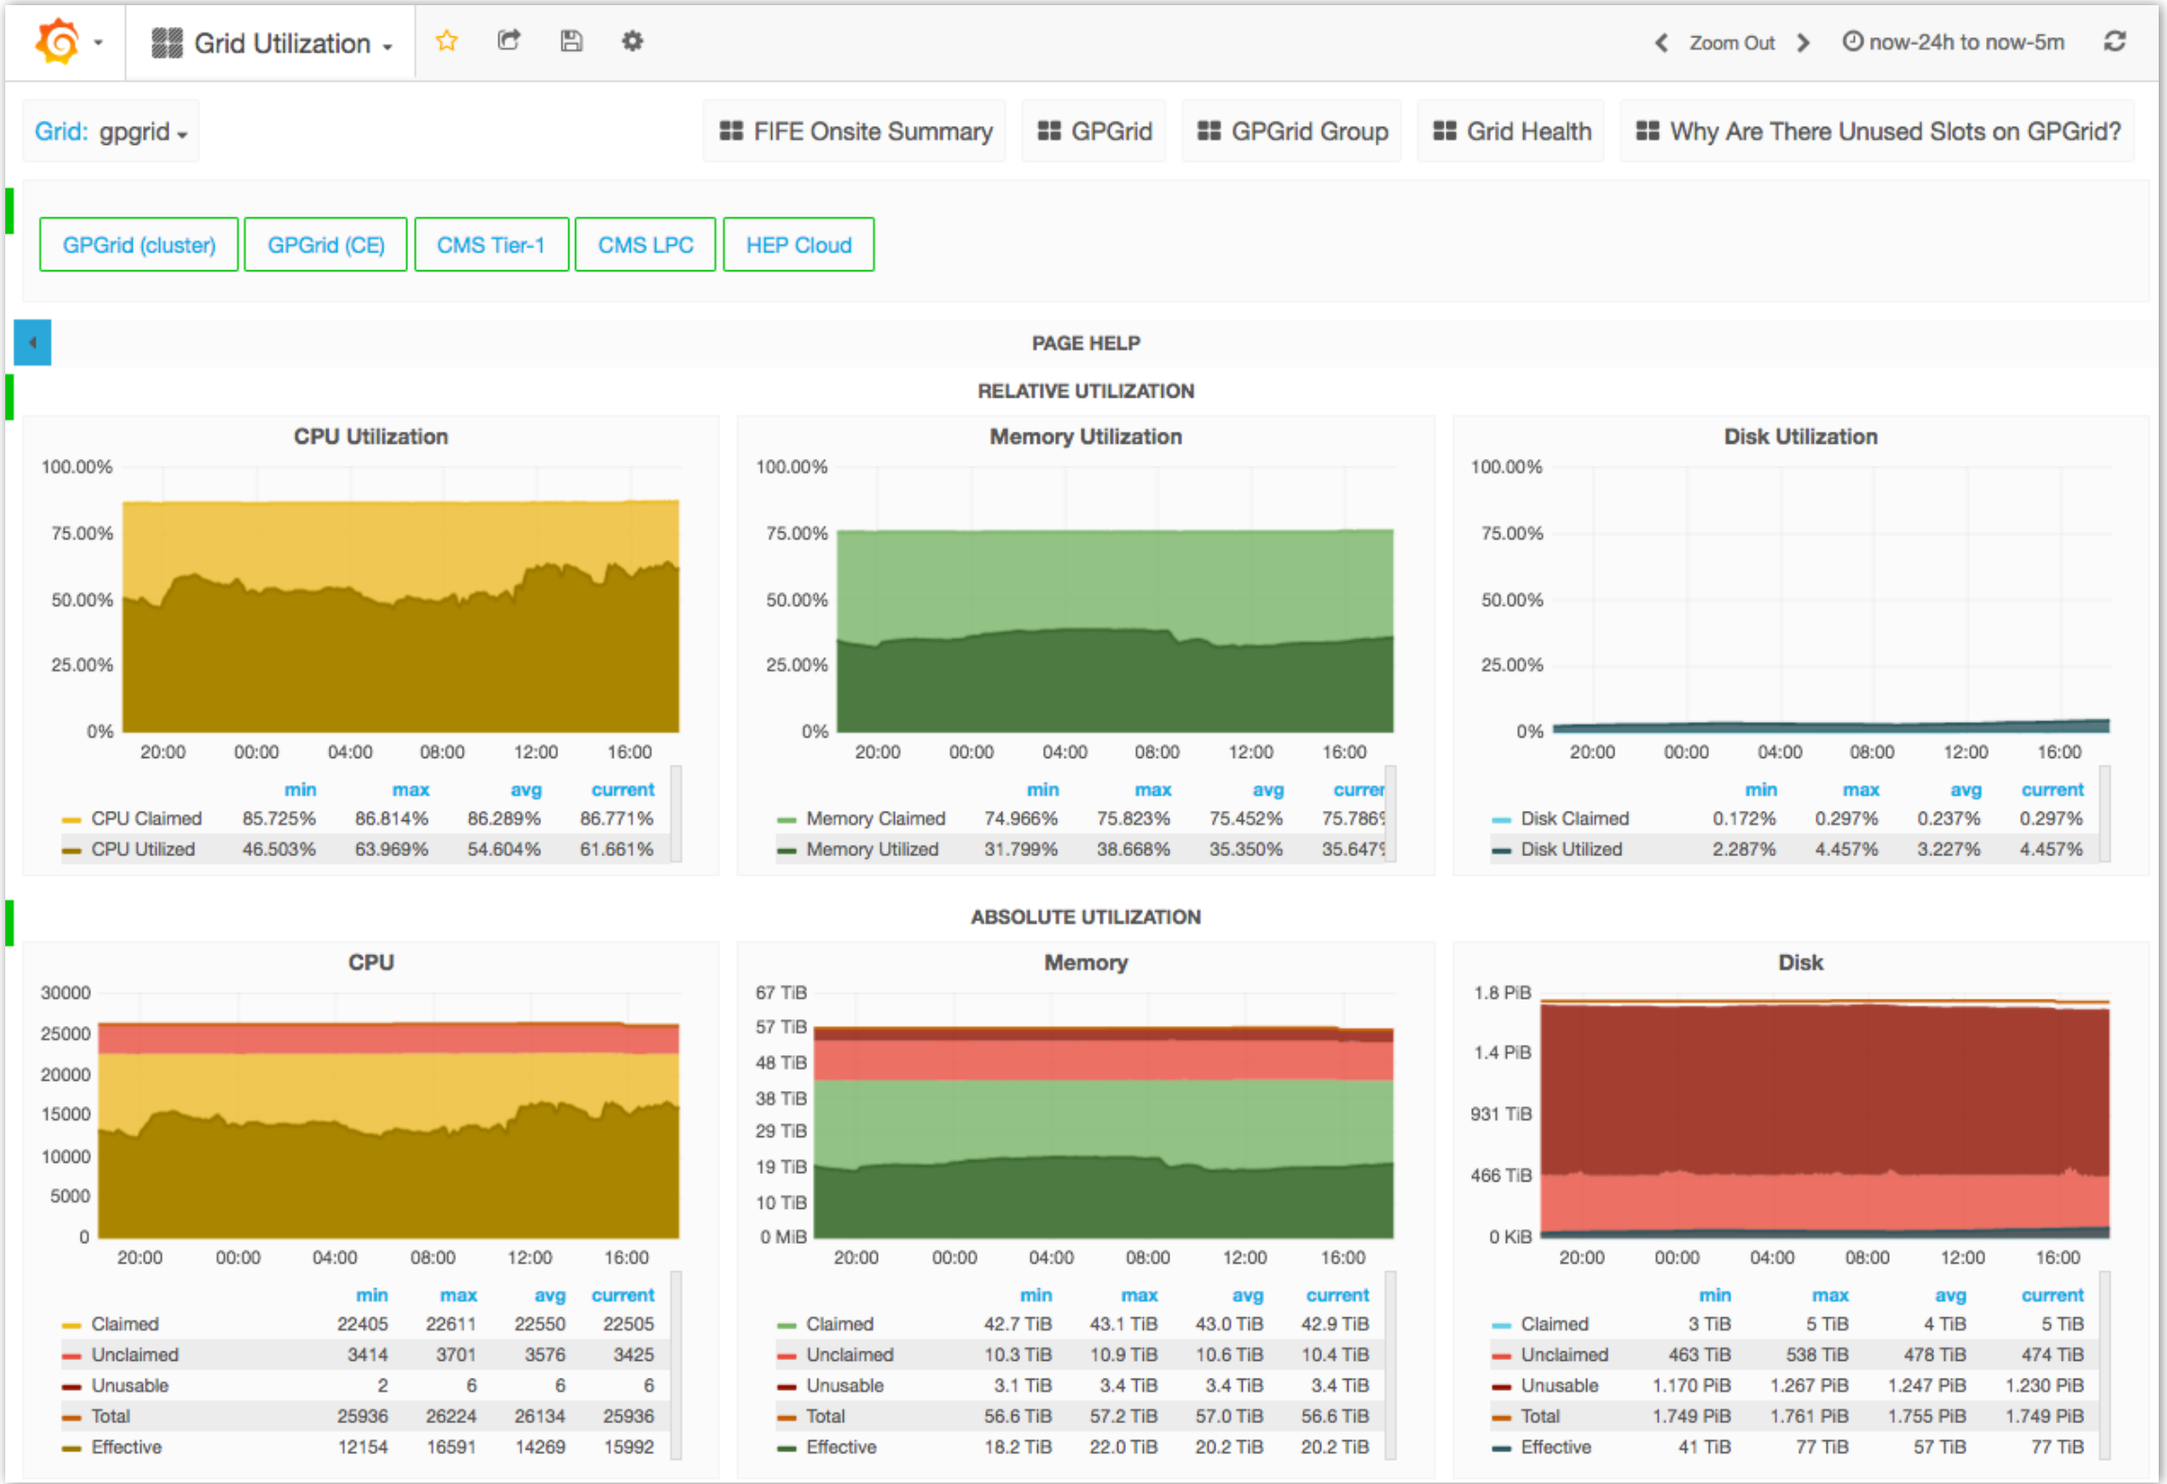

Fermilab

Fifemon - Troubleshooting Guides

Are there FIFE jobs requesting onsite resources?

If jobs are requesting only OFFSITE, they will not run on GPGrid, unless they come back through the OSG opportunistic gatekeeper.

See also:

FIFE Onsite Summary | Fifebatch - Onsite

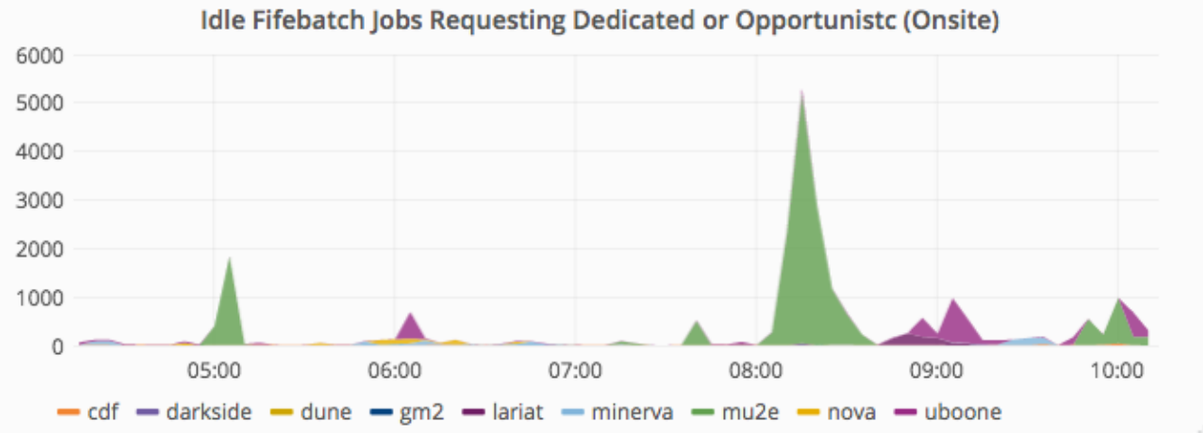

Are the remaining resources in the Glideins "unusable"?

If there are lots of mulitcore or high-memory (>2 GB) jobs running there will be unusable resources left in the glideins.

See also:

Grid Utilization (GPGrid CE) | Fifebatch Slots (GPGrid)

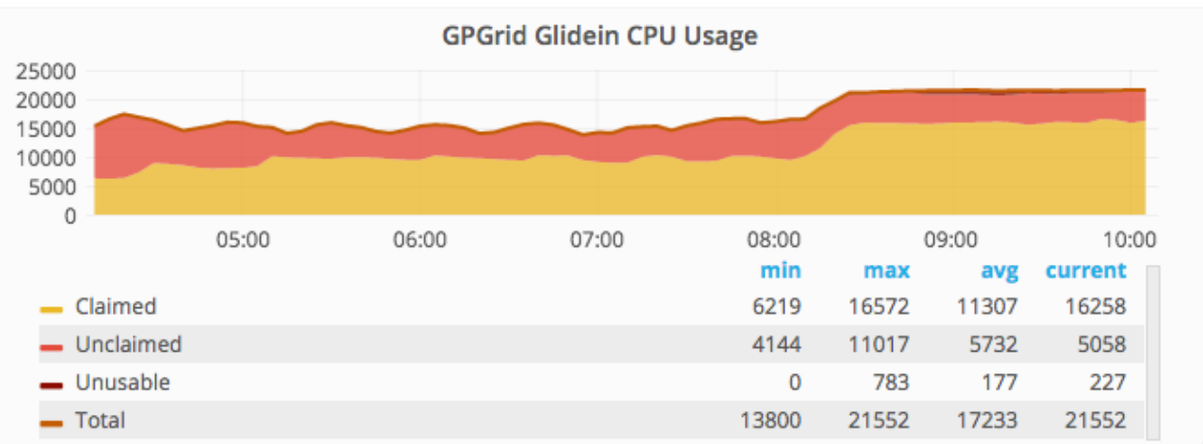

Are the remaining Glidein resources unusable due to job resource requests?

"Unusable" above means that no job could possibly run; if there's no jobs requesting resources within the limits of what's remaining in the Glideins then they too are effectively "unusable."

See also:

Fifebatch Slots (GPGrid) | Fifebatch Slots Unclaimed (GPGrid)

Fifemon - System Monitoring

Probe Status

Production Probe Update Time

Metric	Min	Max	Avg	Current
awsmonitor-us-east-1	2.45 s	2.82 s	2.57 s	2.55 s
awsmonitor-us-west-1	2.83 s	2.94 s	2.89 s	2.94 s
awsmonitor-us-west-2	3.11 s	3.33 s	3.22 s	3.33 s
bluearc	5.17 s	11.46 s	5.90 s	11.46 s
cmssrv14_status	1.35 s	1.56 s	1.43 s	1.37 s
cmssrv274_status	3.14 s	3.49 s	3.28 s	3.20 s
cmssrv39_status	16.30 s	20.31 s	18.82 s	20.31 s
fifebatch-pp_status	203 ms	327 ms	275 ms	314 ms
fifebatch2_status	3.23 s	5.10 s	4.17 s	4.68 s
fifebatch_status	8.66 s	12.56 s	10.58 s	12.56 s

Dev Probe Duty Cycle

Metric	Min	Max	Avg	Current
bluearc	0.0%	0.0%	0.0%	0.0%
cms-global	-	-	-	-
cmssrv14_status	0.5%	0.5%	0.5%	0.5%
cmssrv274_status	1.0%	1.2%	1.1%	1.1%
cmssrv39_status	6.8%	8.5%	7.8%	8.5%
fife_gracc_stats	0.0%	0.0%	0.0%	0.0%
fifebatch-dev_statu				
fifebatch-pp_statu				
fifebatch2_status				
fifebatch_status				

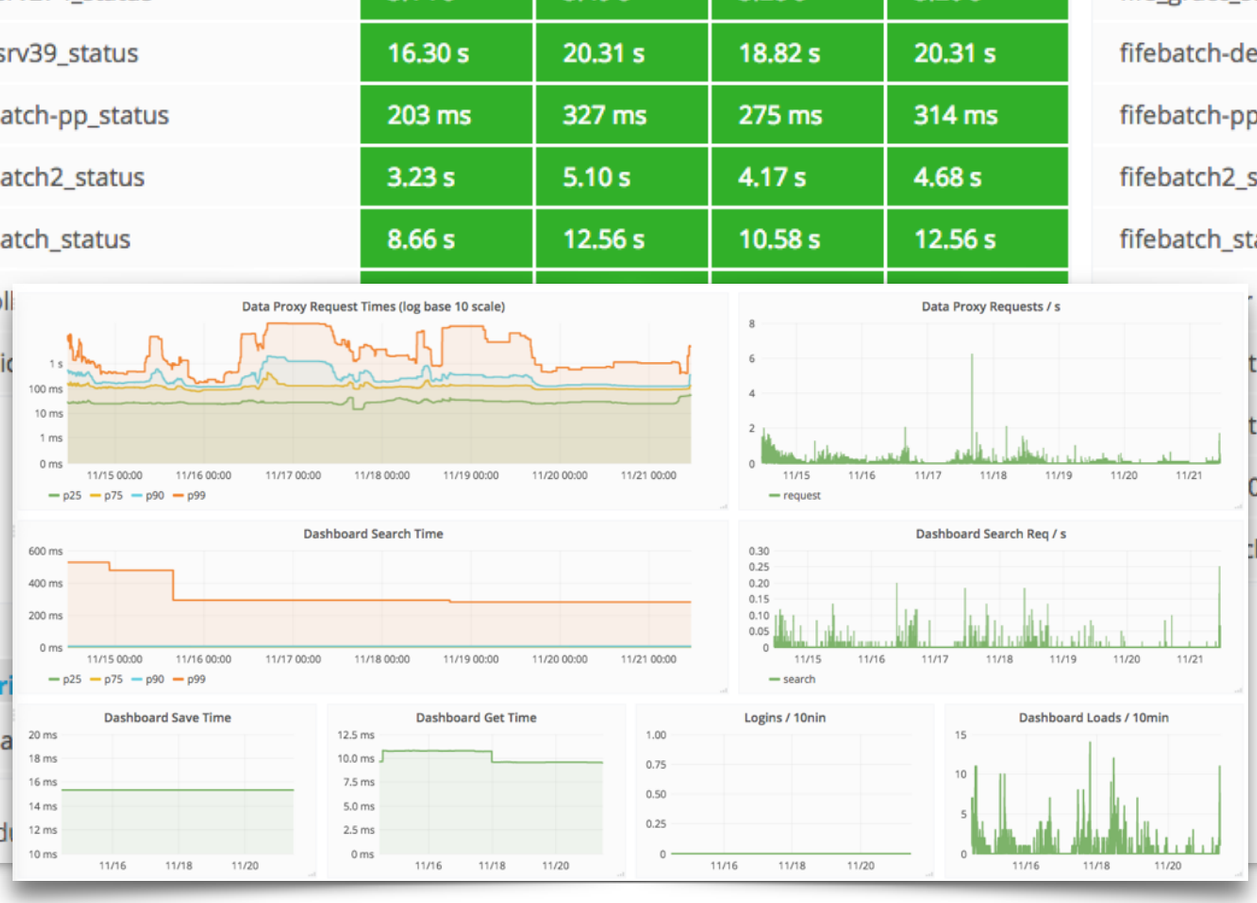

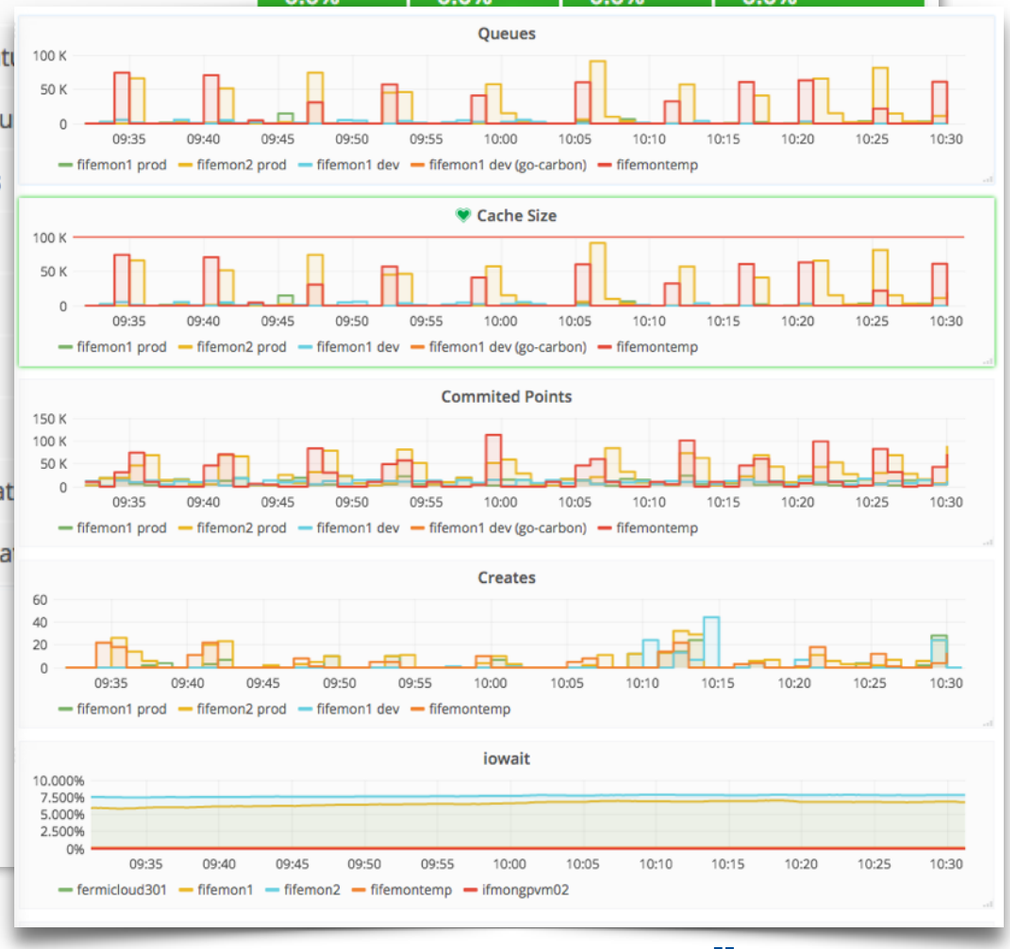

Fermilab

The success of Fifemon led to the Landscape program: common monitoring tools and support for Scientific Computing.

Landscape

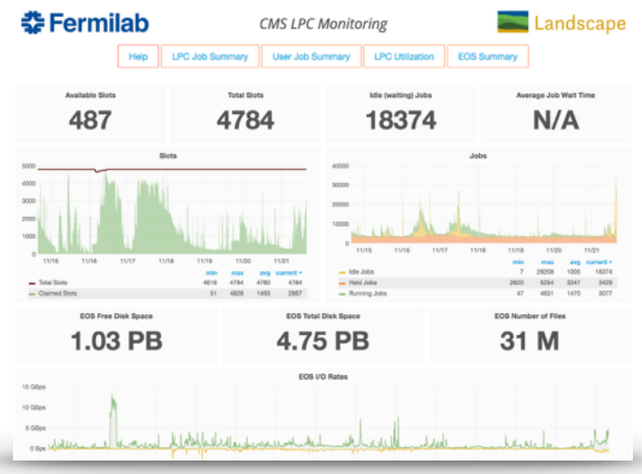

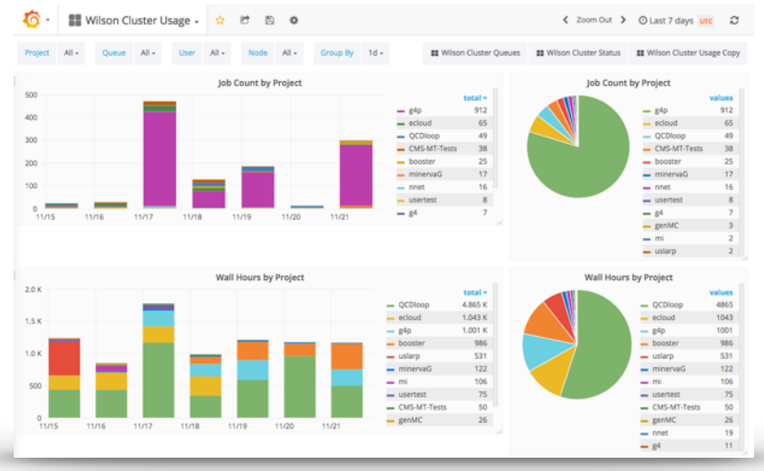

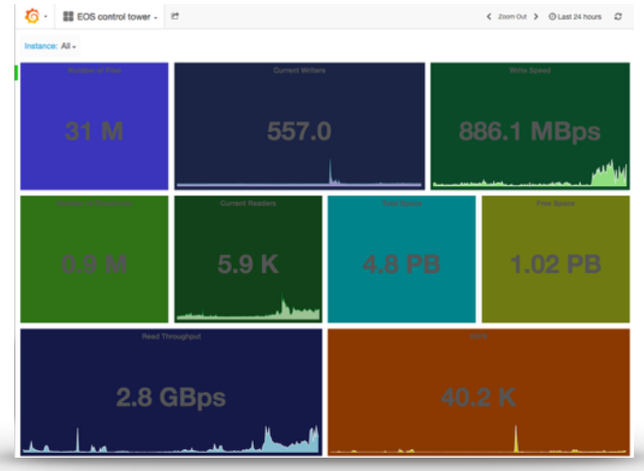

Fermilab

HEP Cloud / GCP Demo Live @ SC16

Google Cloud
@googlecloud

Following

160k cores double @CMSexperiment capacity via @fermilab's #HEPCloud. Powering particle physics w/ Preemptible VMs: goo.gl/CaQ9v7

Fermilab

Open Science Grid

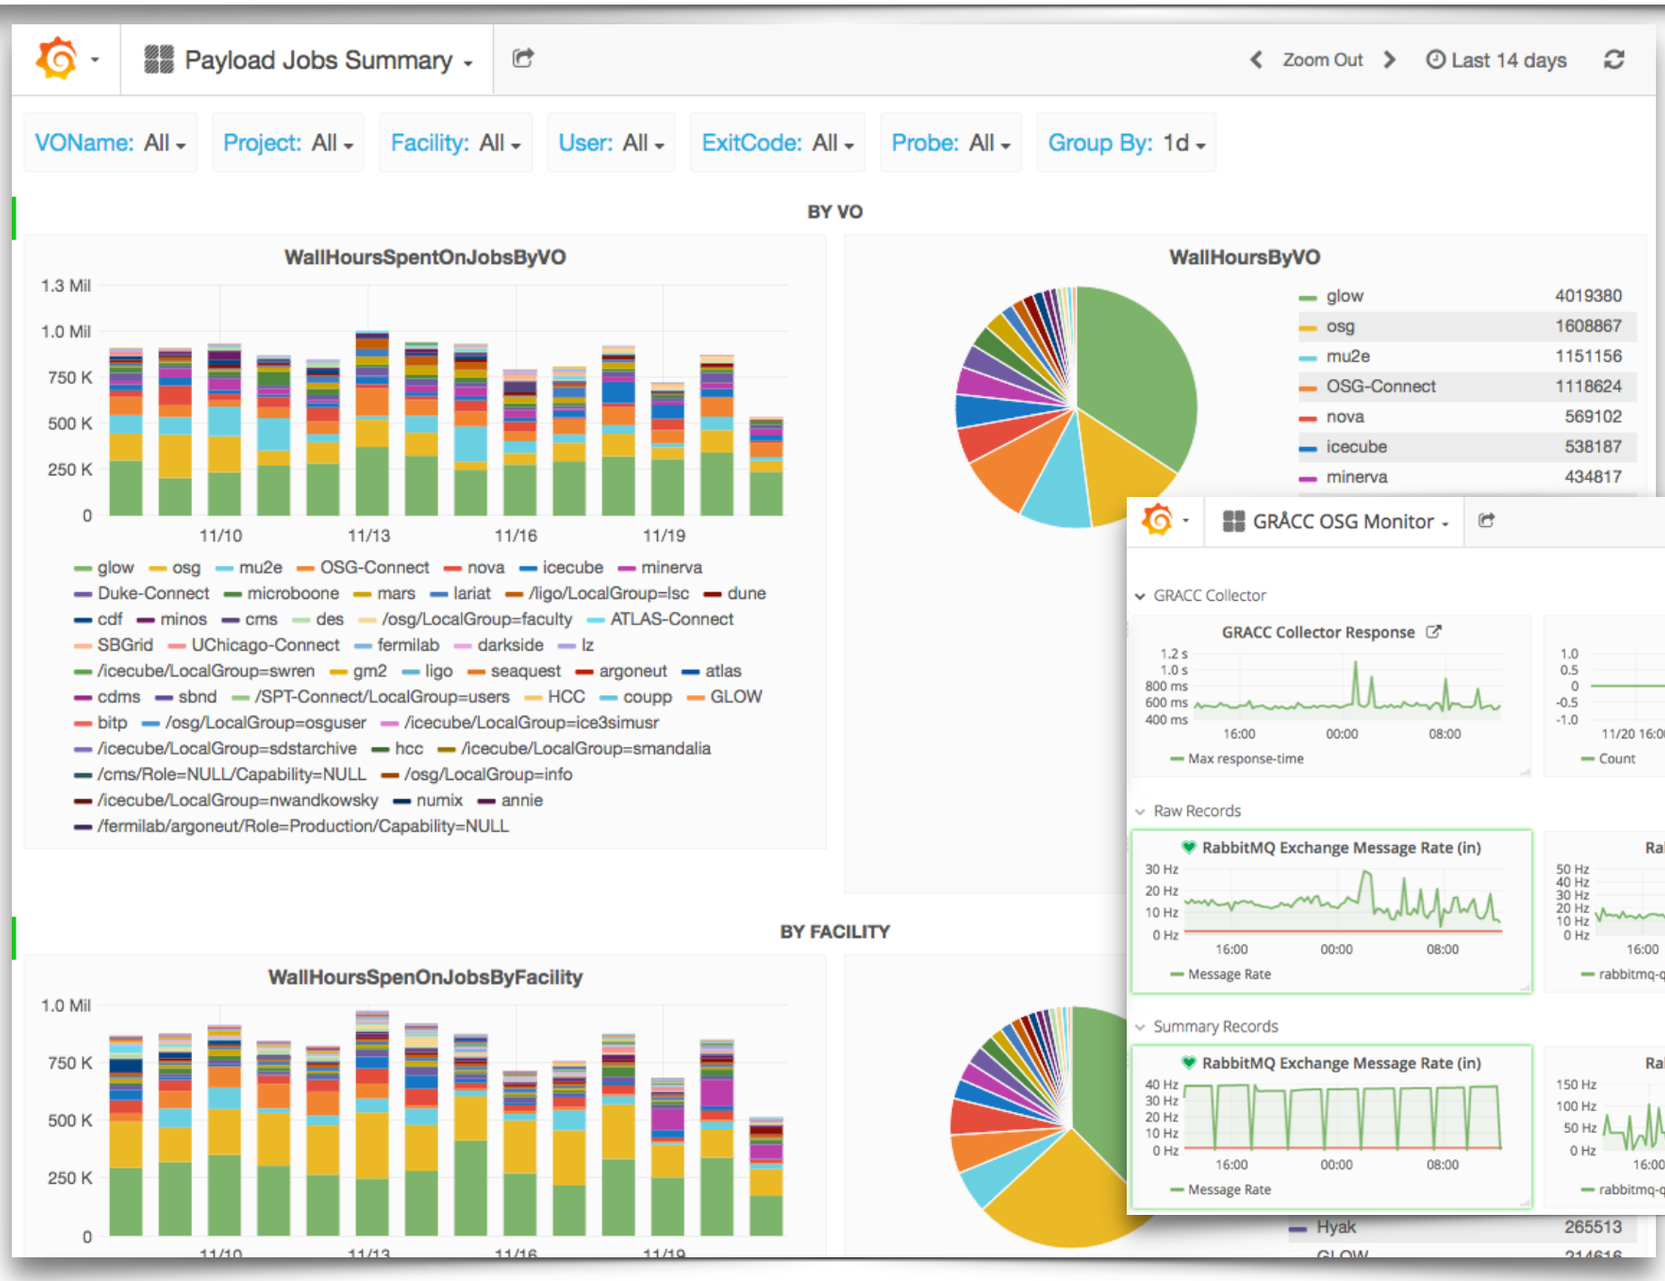

Fermilab

Fermilab ❤ Grafana

(for the most part)

Things we love

Understanding system & service data is ~~hard~~ easy

- Trivial to add graphs and dashboards
- What used to take an expert to figure out (if anyone could), is now easily understood by (almost) anyone.

Built-in Auth

- LDAP authN & authZ
- SSO (proxy) authN

It’s beautiful!

- Great web UX
- Plugin system makes it easy to customize with native look & feel

… and Alerting!

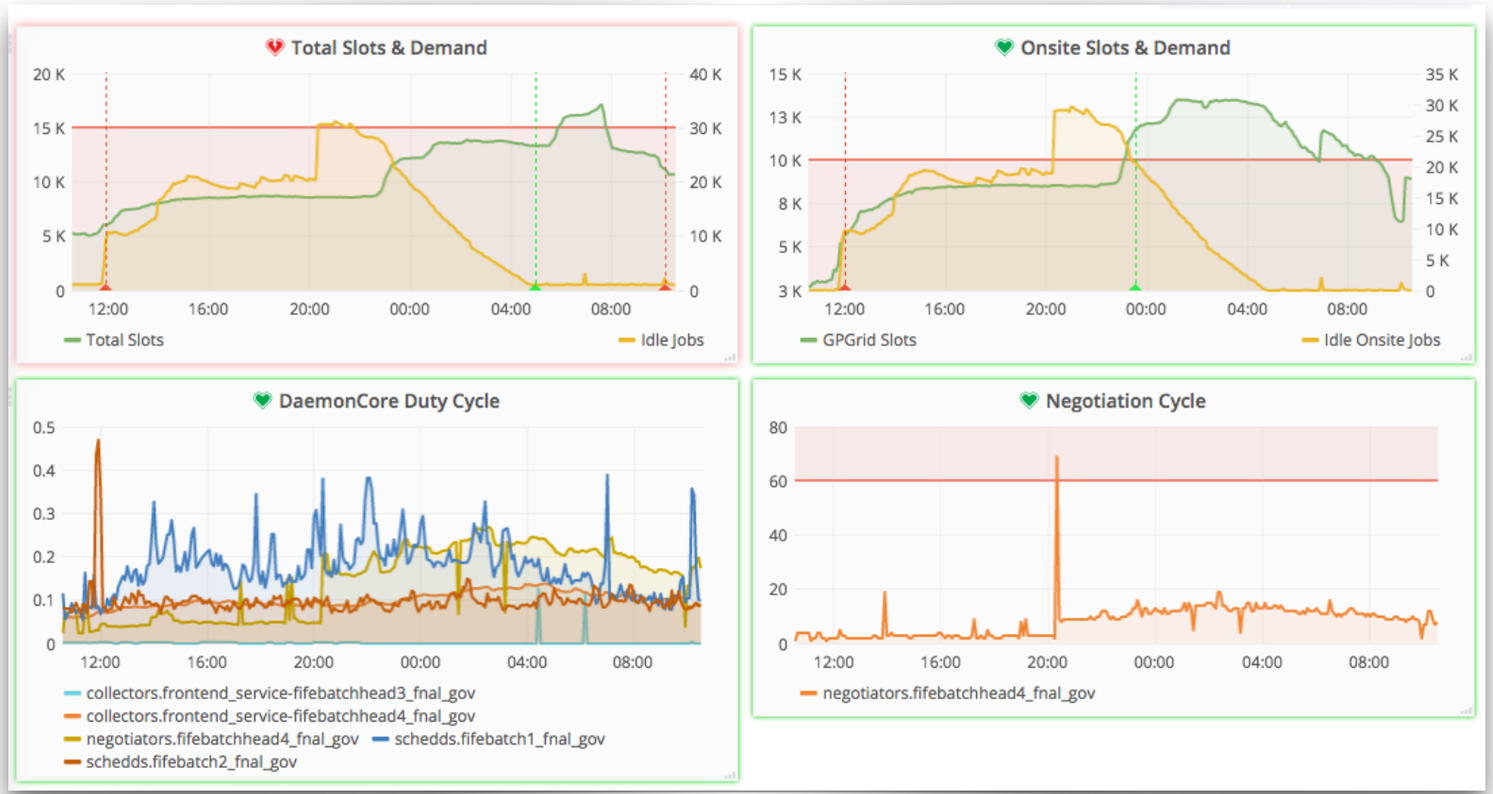

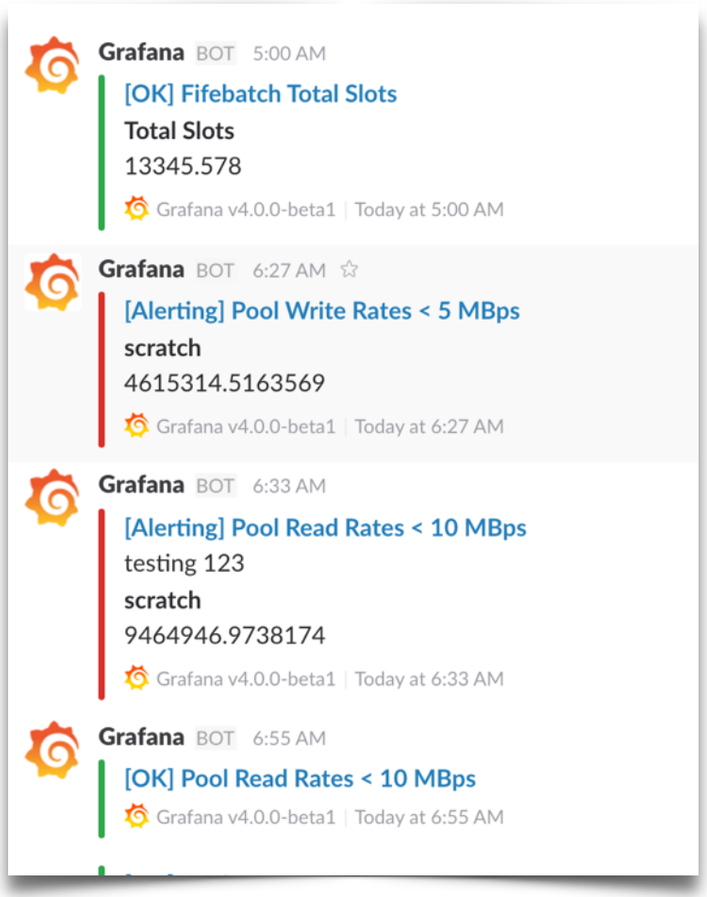

Fermilab

Things we'd love to see

User & group sandboxes

- Organizations don't work
  - different data sources, dashboards & authZ
  - broken URLs

Dashboard revision control & management

- Who changed what, when, and why
- Dashboards are code
- CLI? github.com/retzkek/grafanactl

Better Elasticsearch & non-time-series support

More drill-down links

More configurable home page (e.g. template vars)

Links

fifemon.github.io

landscape.fnal.gov

gracc.opensciencegrid.org

Contact

kretzke@fnal.gov

@kevin_retzke

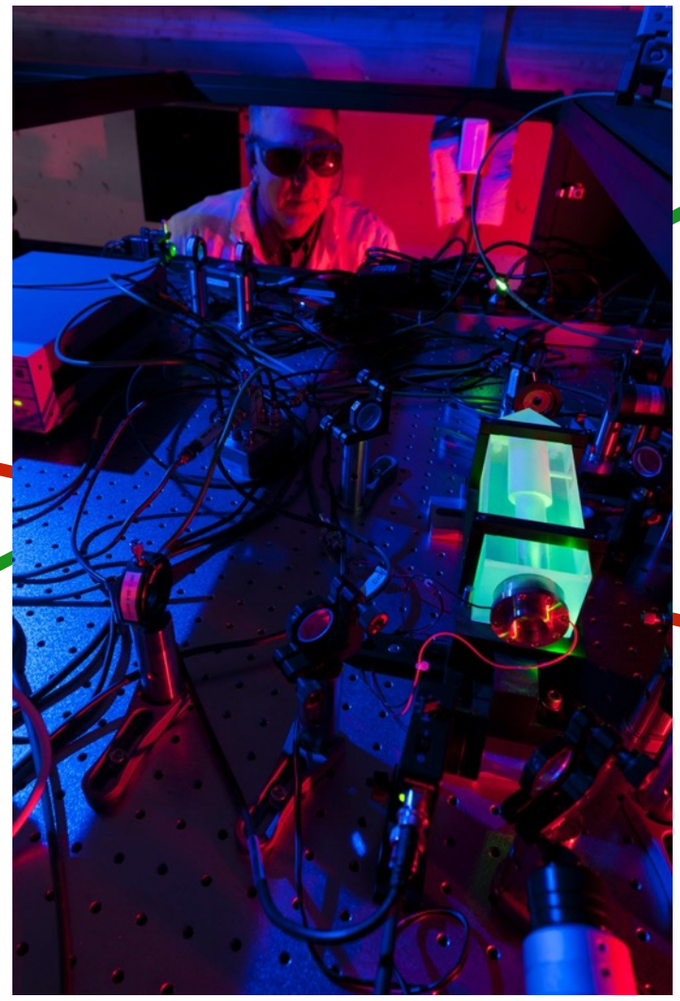

Lasers!

Fermilab is Operated by the Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the United States Department of Energy.