



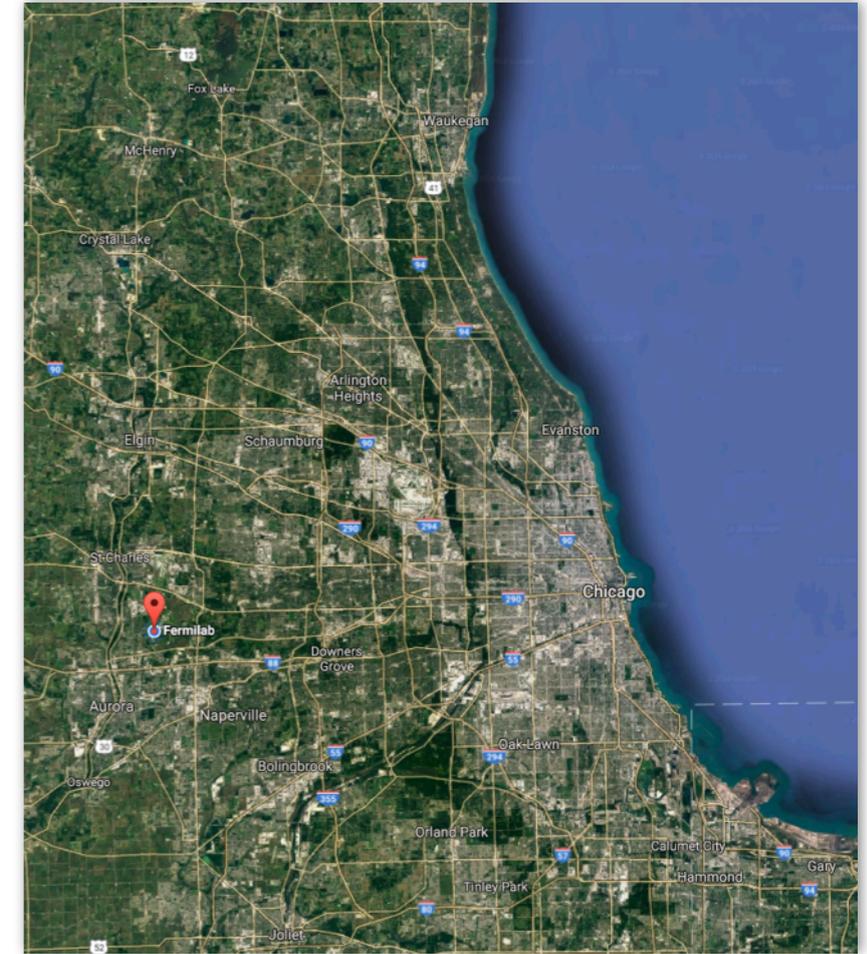
Monitoring Large-Scale Scientific Computing with Grafana

Kevin Retzke
GrafanaCon
November 30 2016

Disclaimer

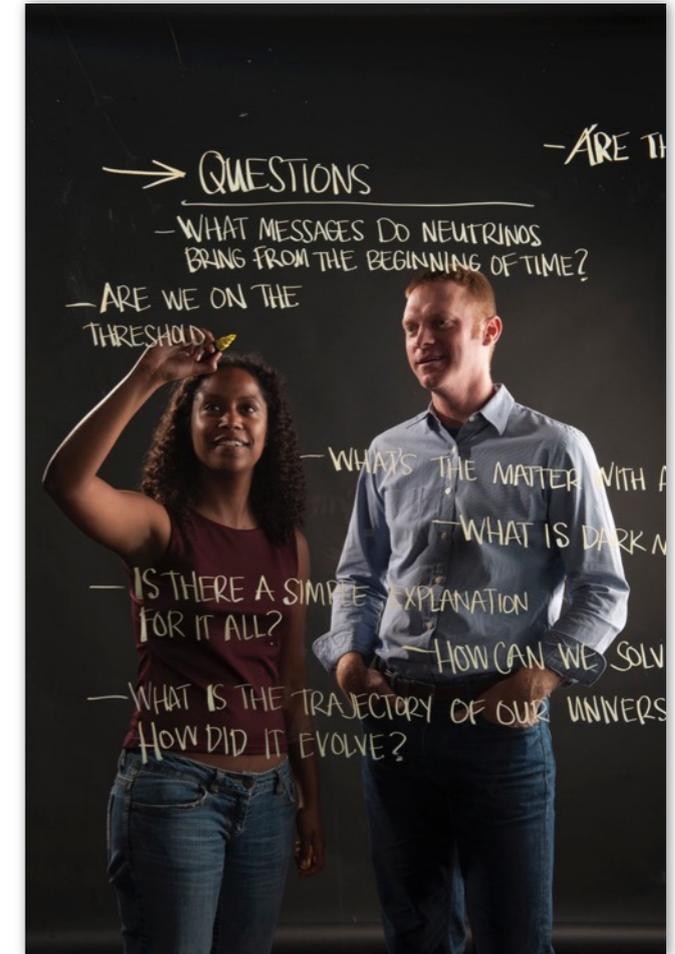
Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof.

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



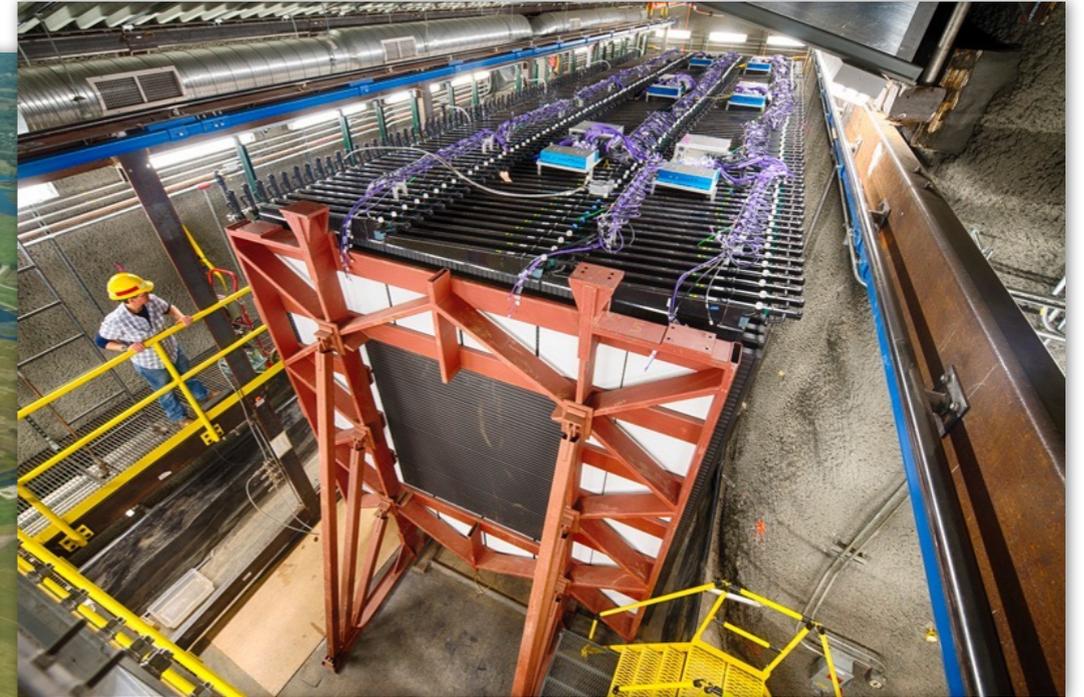
More than 3,500 scientists worldwide use Fermilab and its accelerators, detectors, and computers for their research. About 2,600 researchers from 44 countries collaborate on experiments at Fermilab.

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



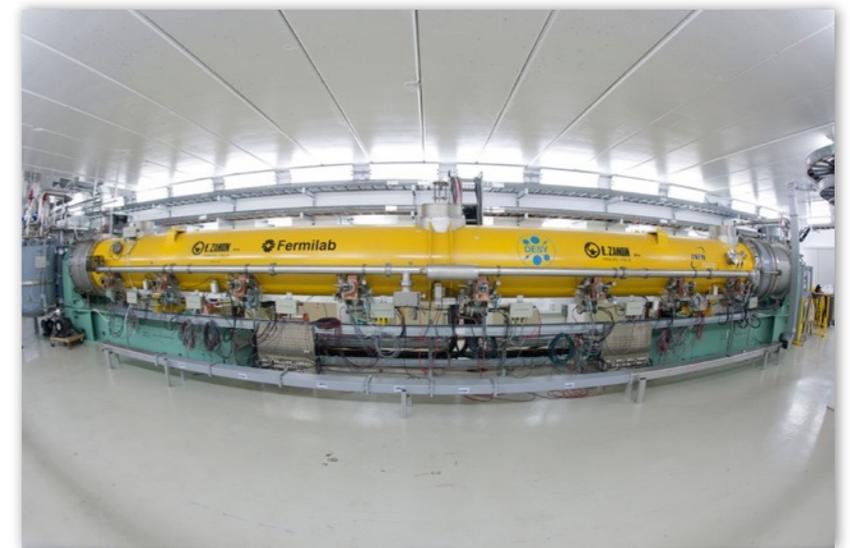
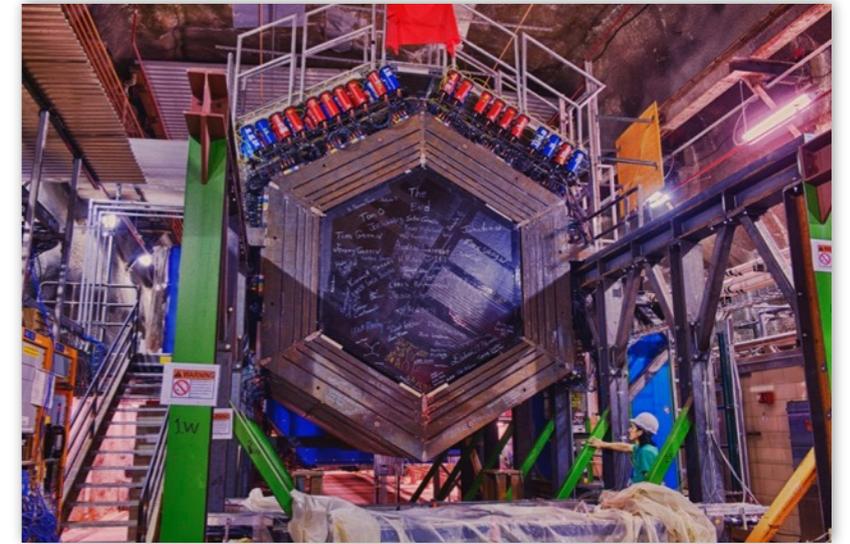
The Dark Energy Camera, designed and built at Fermilab, is one of the world's most powerful cameras (570-megapixel), and now takes images on a telescope in Chile. The heart of the Dark Energy Survey, it advances the quest to understand the nature of the dark energy that pushes the universe apart.

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 - Science!



DUNE - The Next Mega-Science Project

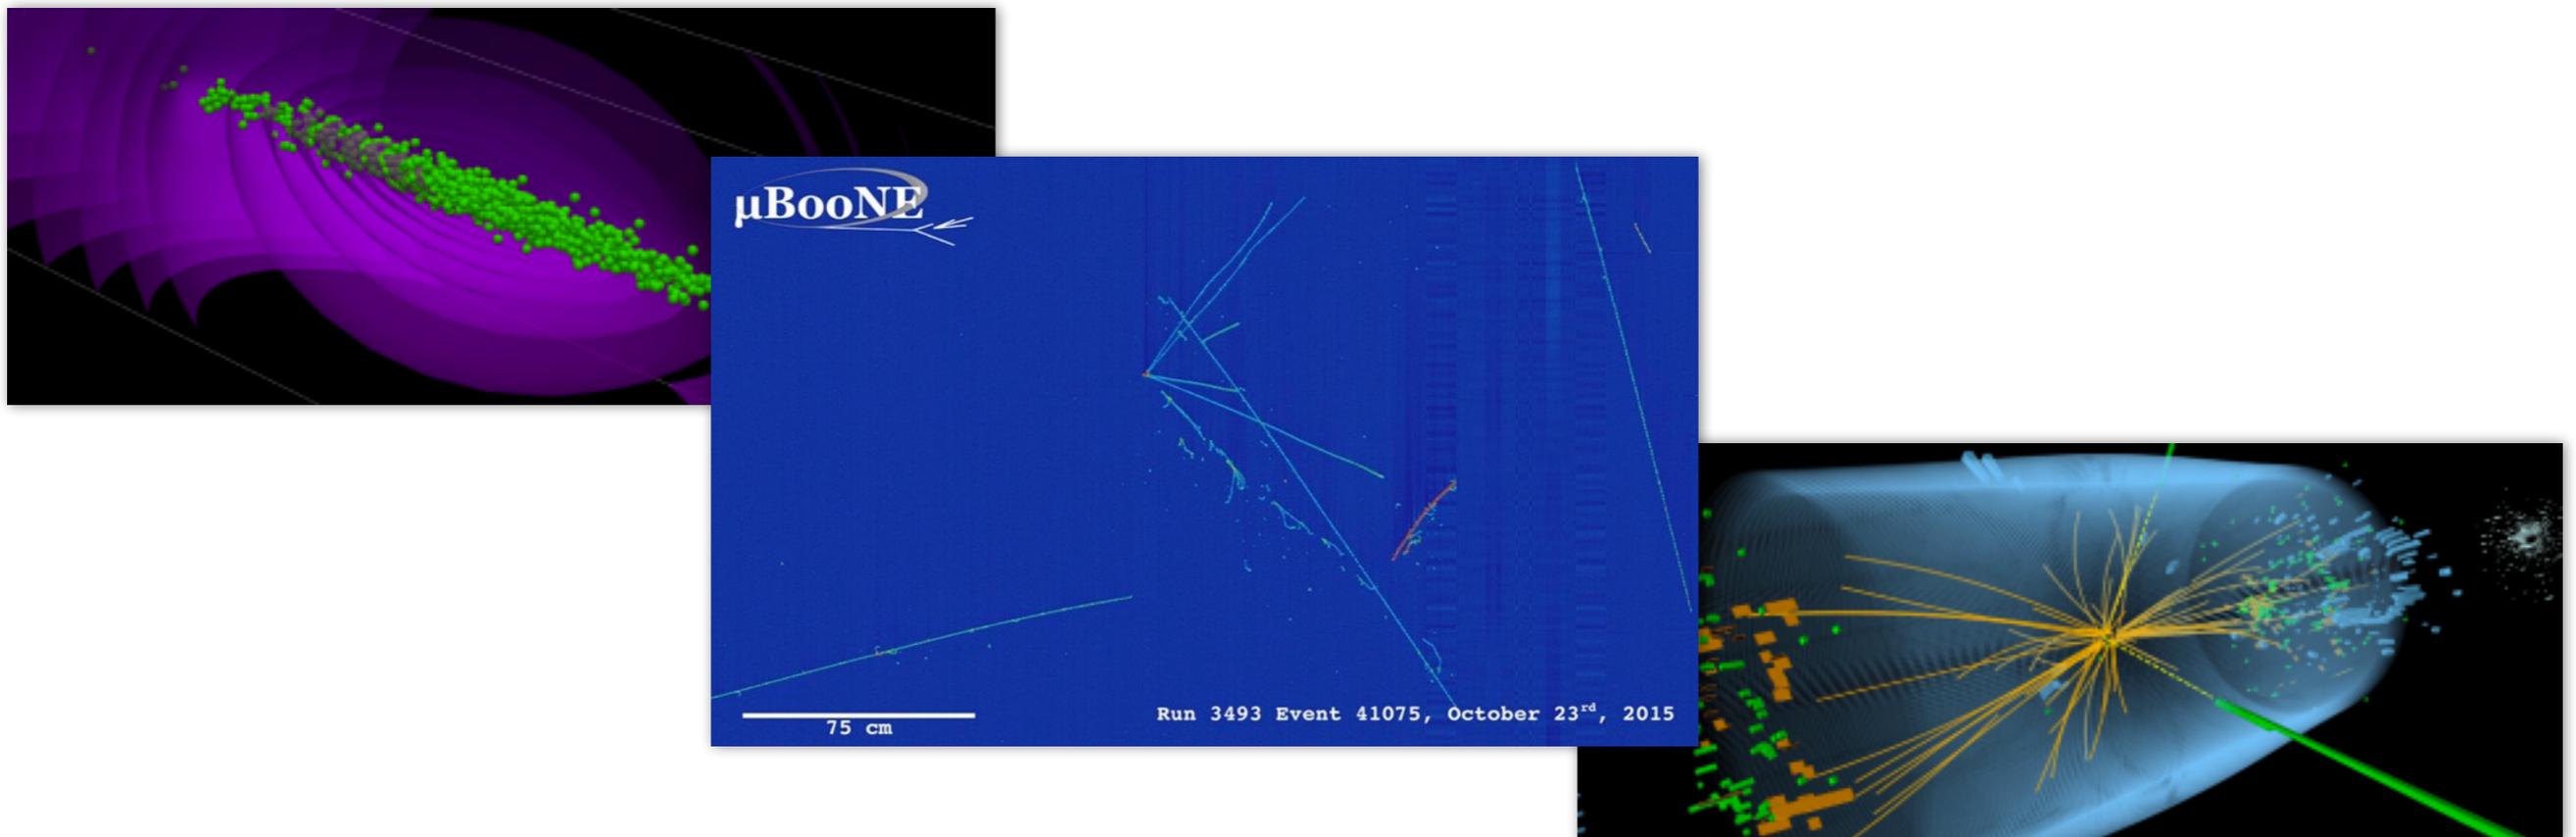
https://youtu.be/AYtKcZMJ_4c

Scientific Computing

Scientific Computing @ Fermilab

High energy physics requires a **huge** amount of computing

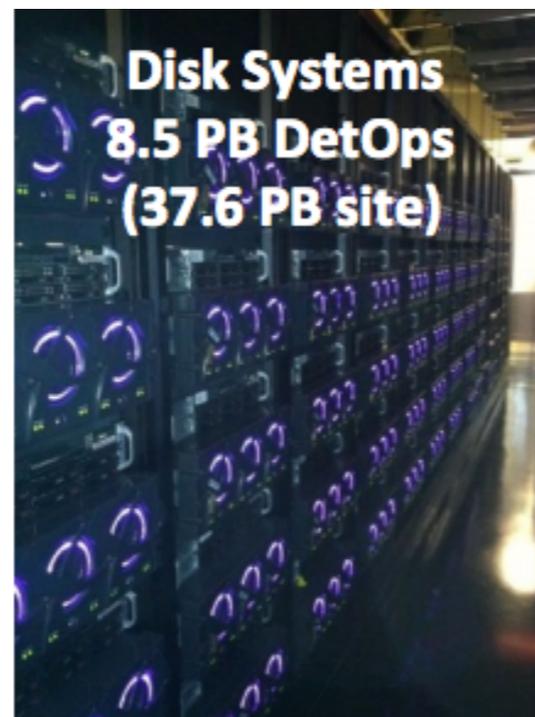
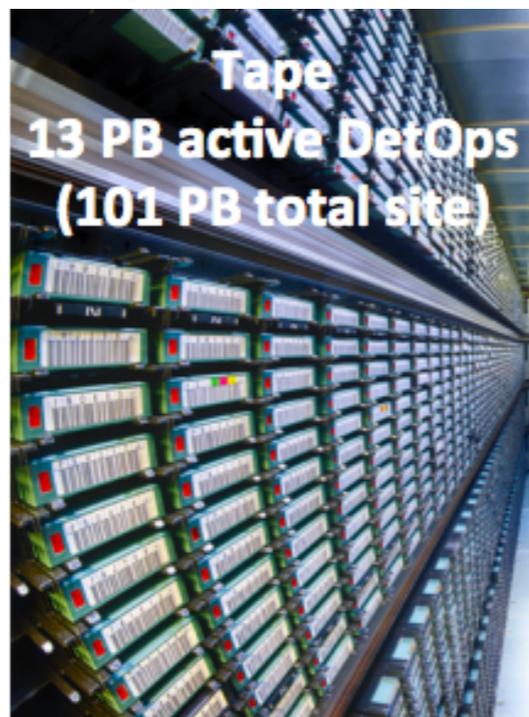
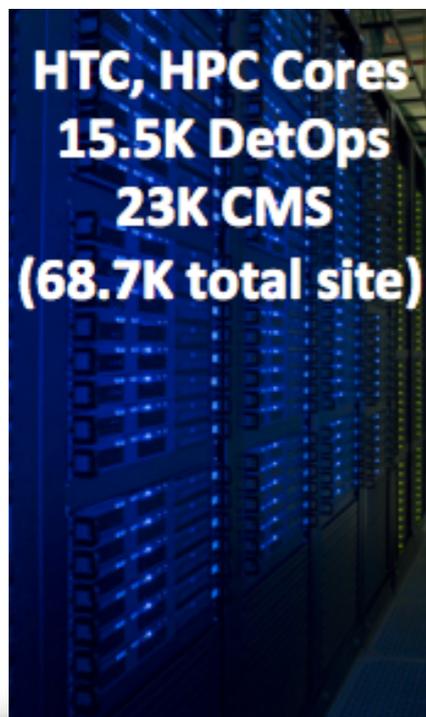
- Detector and accelerator design and simulation
- Data collection and reconstruction
- Event analysis



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



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

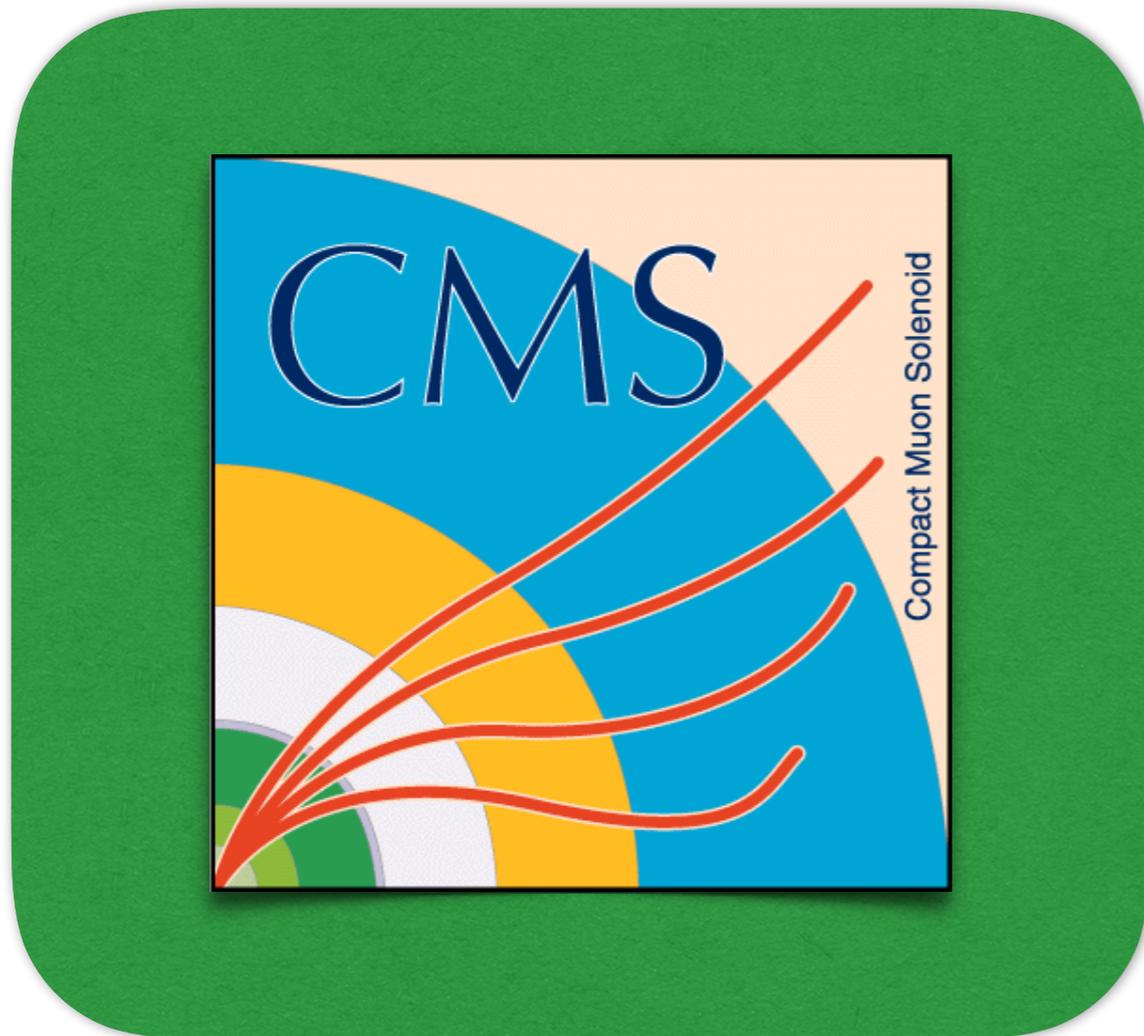
HPC/GPU systems, primarily dedicated to Lattice QCD calculations

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

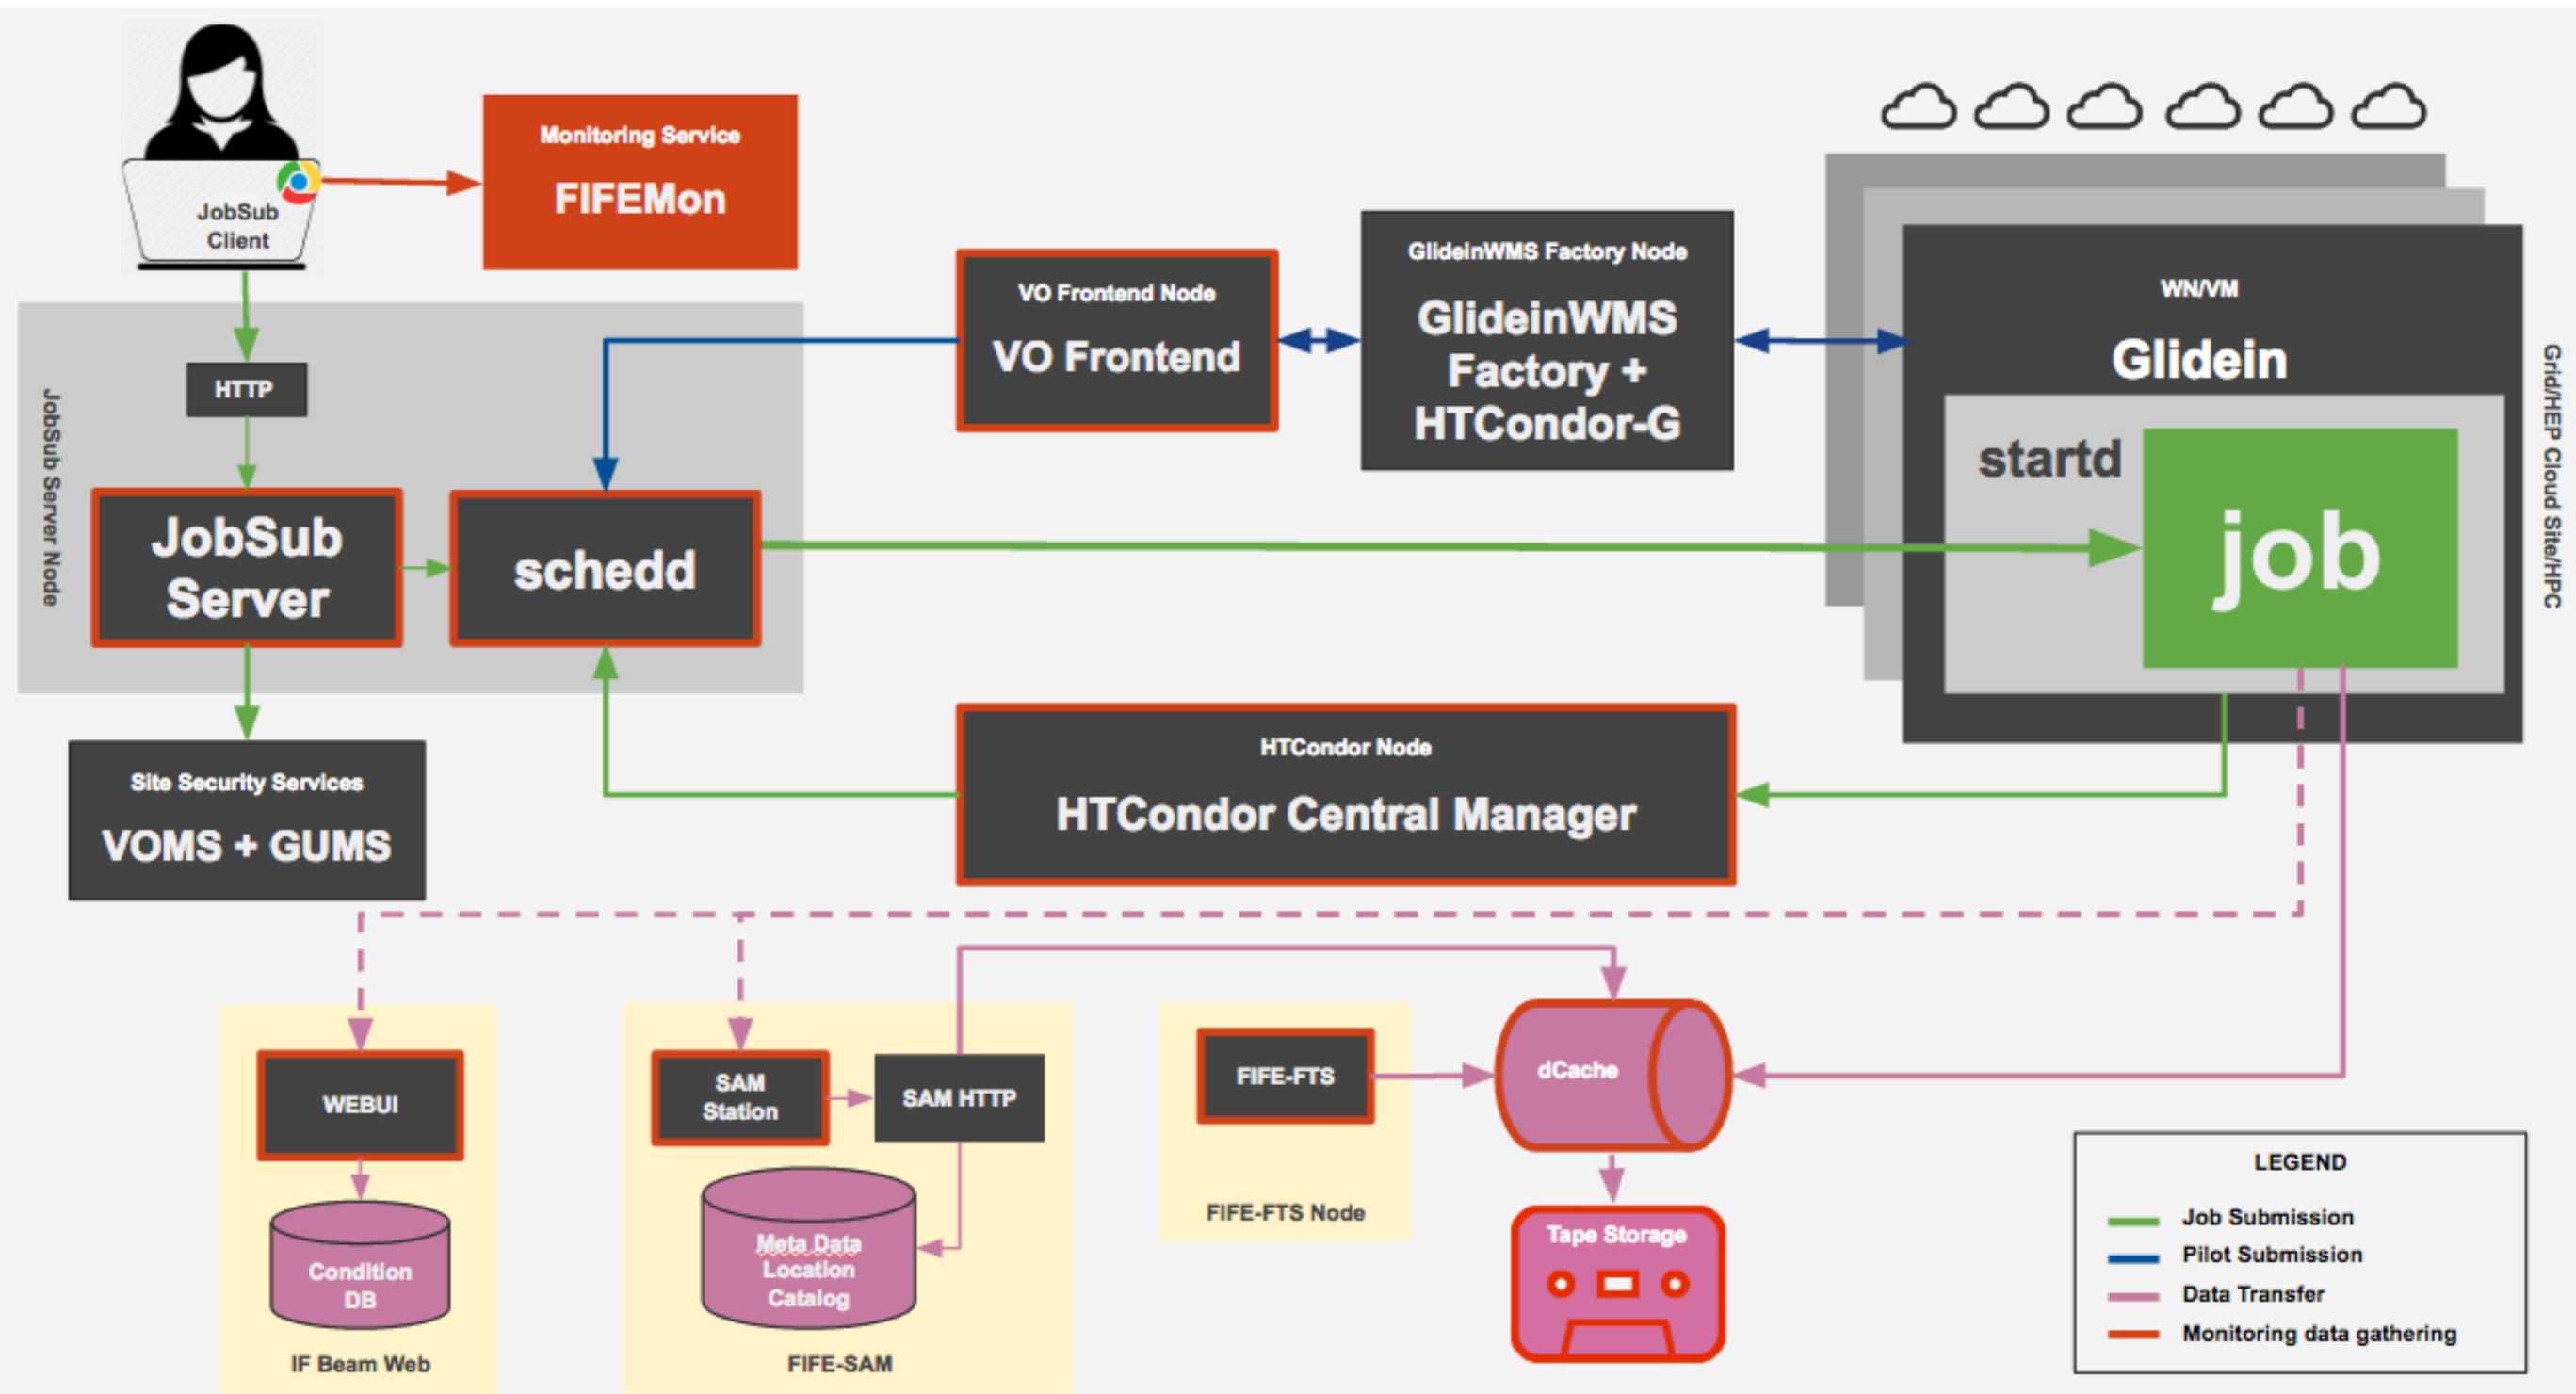
Batch/Grid Computing - FIFE and CMS

Two main HTC systems:

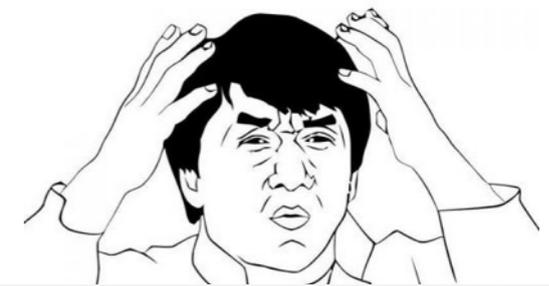
- CMS Tier 1 - dedicated computing for CMS experiment
- FIFE (Fabric for Frontier Experiments) - everyone else



FIFE Batch System Architecture



Monitoring as it was...



Store/Restore to/from Enstore by Storage Groups. Last 6 Months

Month	2014-01		2014-02		2014-03		2014-04		2014-05		2014-06	
Storage Group	Store	Restore										
HCC	13	0	2,419	0	2,420	1	0	0	2,706	0	0	0
astro	944	0	0	372	0	0	0	0	0	0	0	0
cdms	4,260	0	2,363	0	2,366	0	1,976	270	0	0	0	0

FTS status for novadaq-ctrl-datadisk-02.fnal.gov

Generated at 2014-06-13 15:54:08 CDT (refresh)

Summary

FTS: OK FSS: OK Stager: OK

Completed files: 147395
Failed transfers: 2276

Project Id	1666
Status	running
Owner	satish
Start time	2014-06-11 08:03:29
Dataset definition	prod_federat_artdaq_cosmic_15400_16000_draining_513-03-24
Files in snapshot	3543
Files seen	3543
Processes	179
Busy processes	0
Finished processes	178
Waiting processes	0
Error processes	1
Mean wait time (per file)	1s
Mean busy time (per file)	21min 45s
Last activity	process ended at 2014-06-11 16:33:36

Job ID	Comp	Subst time	Price	Waiting	Queue	Hold	Completed	Remarks	Cluster	Command
7741	Jun 09 08:29	Jan	10	0	0	10	0	fifebatch1		

Active Transfers

Dest	Domain	Seq	Prot	Owner	Proc	Preld	Pool	Host	Status	Size	S	Trans. (KB)	Speed (KB/s)
DCap11-Data-00213	dcap11-Data-00213	3	http	1666	682			WaitingForDataTransfer	Running	101140	2070280	10413629	201120

Processes

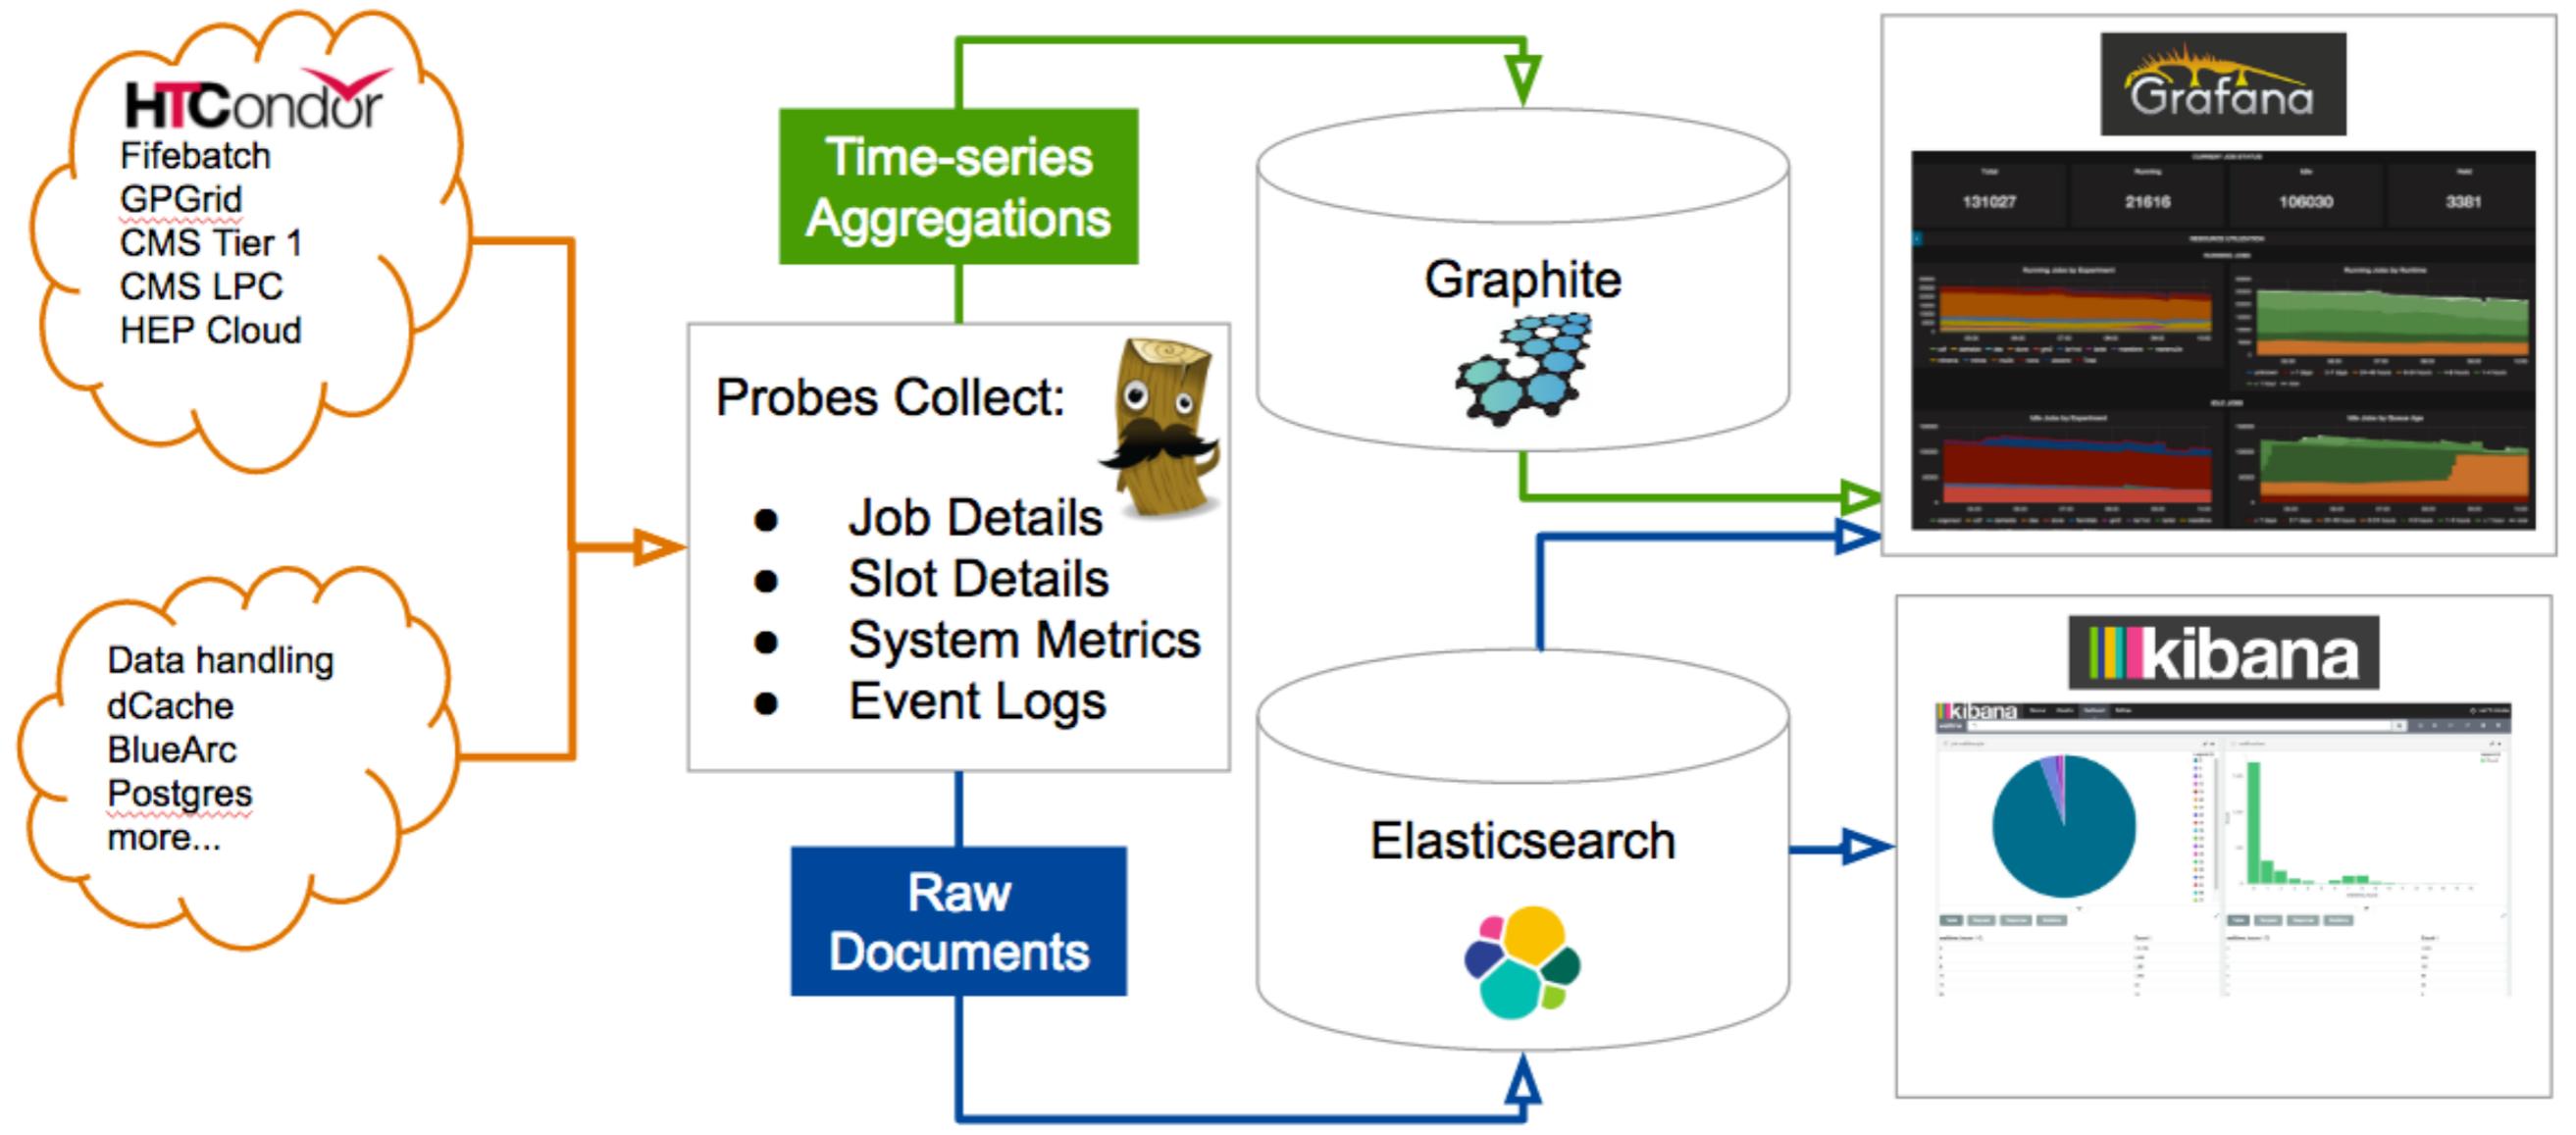
Process Id	Node name	Status	Description	Files seen	Last change	Waiting for	Mean wait time (per file)	Mean busy time (per file)
732712	fnpc3002.fnal.gov	error	16625758.8	1	2014-06-11 08:21:29 (process ended - error)	-	7min 40s	-
732713	fnpc5015.fnal.gov	completed	16625758.0	20	2014-06-11 14:32:36 (process ended - completed)	-	0s	19min 24s
732714	fnpc4005.fnal.gov	completed	16625758.4	20	2014-06-11 15:17:43 (process ended - completed)	-	0s	21min 39s

Fifemon & Landscape

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



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



Fermilab Scientific Computing Summary
FIFE



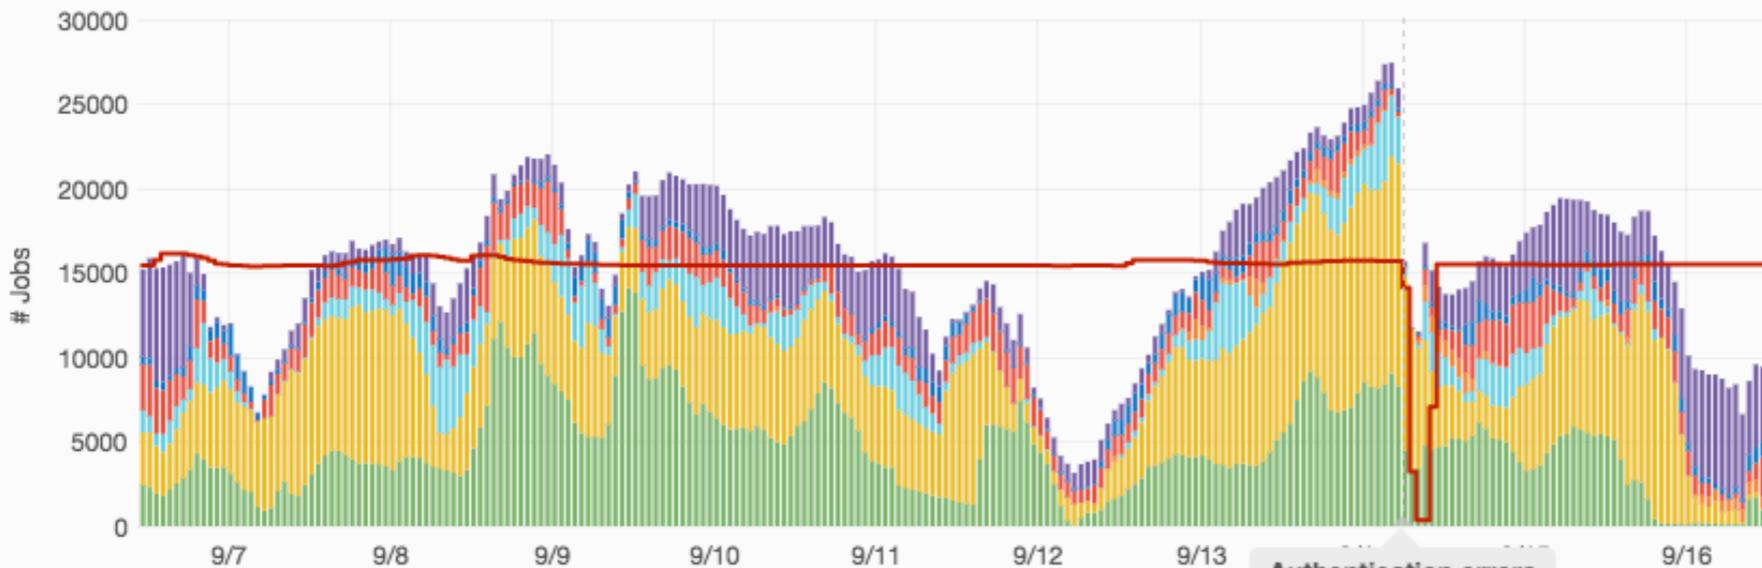
Average Number of Jobs Running Concurrently

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 (GPGGrid)	347	16145	15349

Percent Jobs Run Onsite

57.5%

Percent Jobs Run

N/A

Percent Jobs Run on OSG

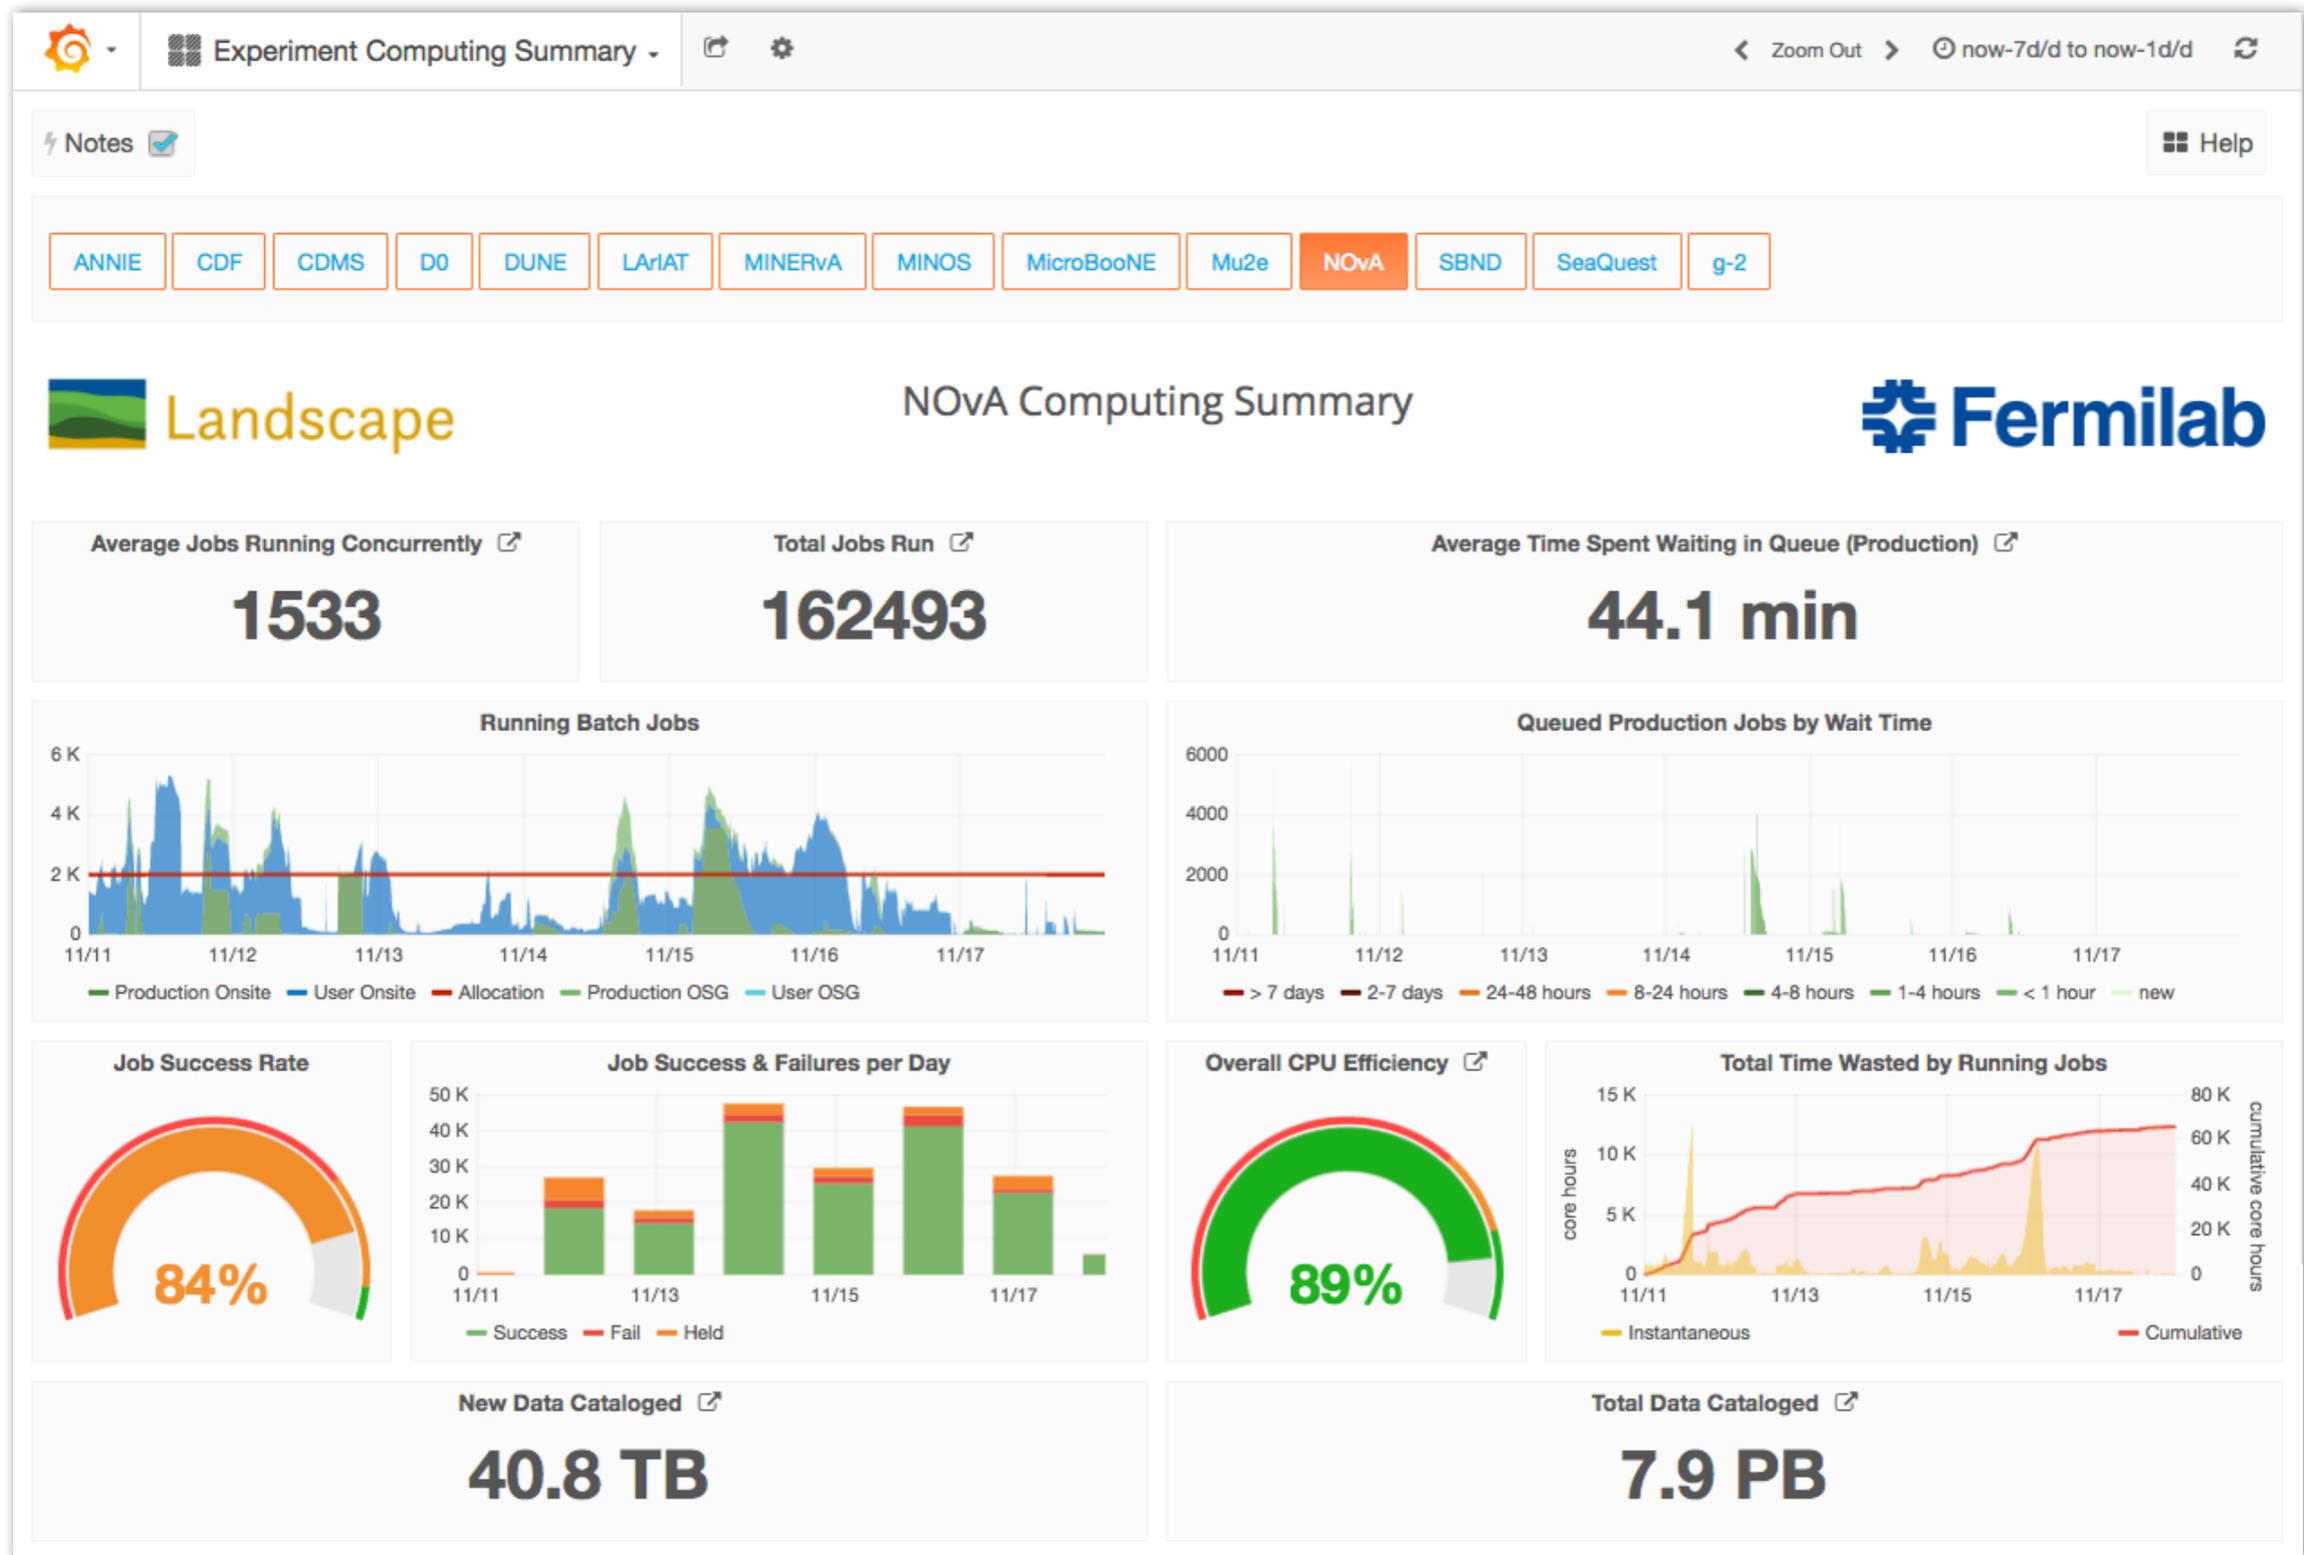
42.8%

NOTES

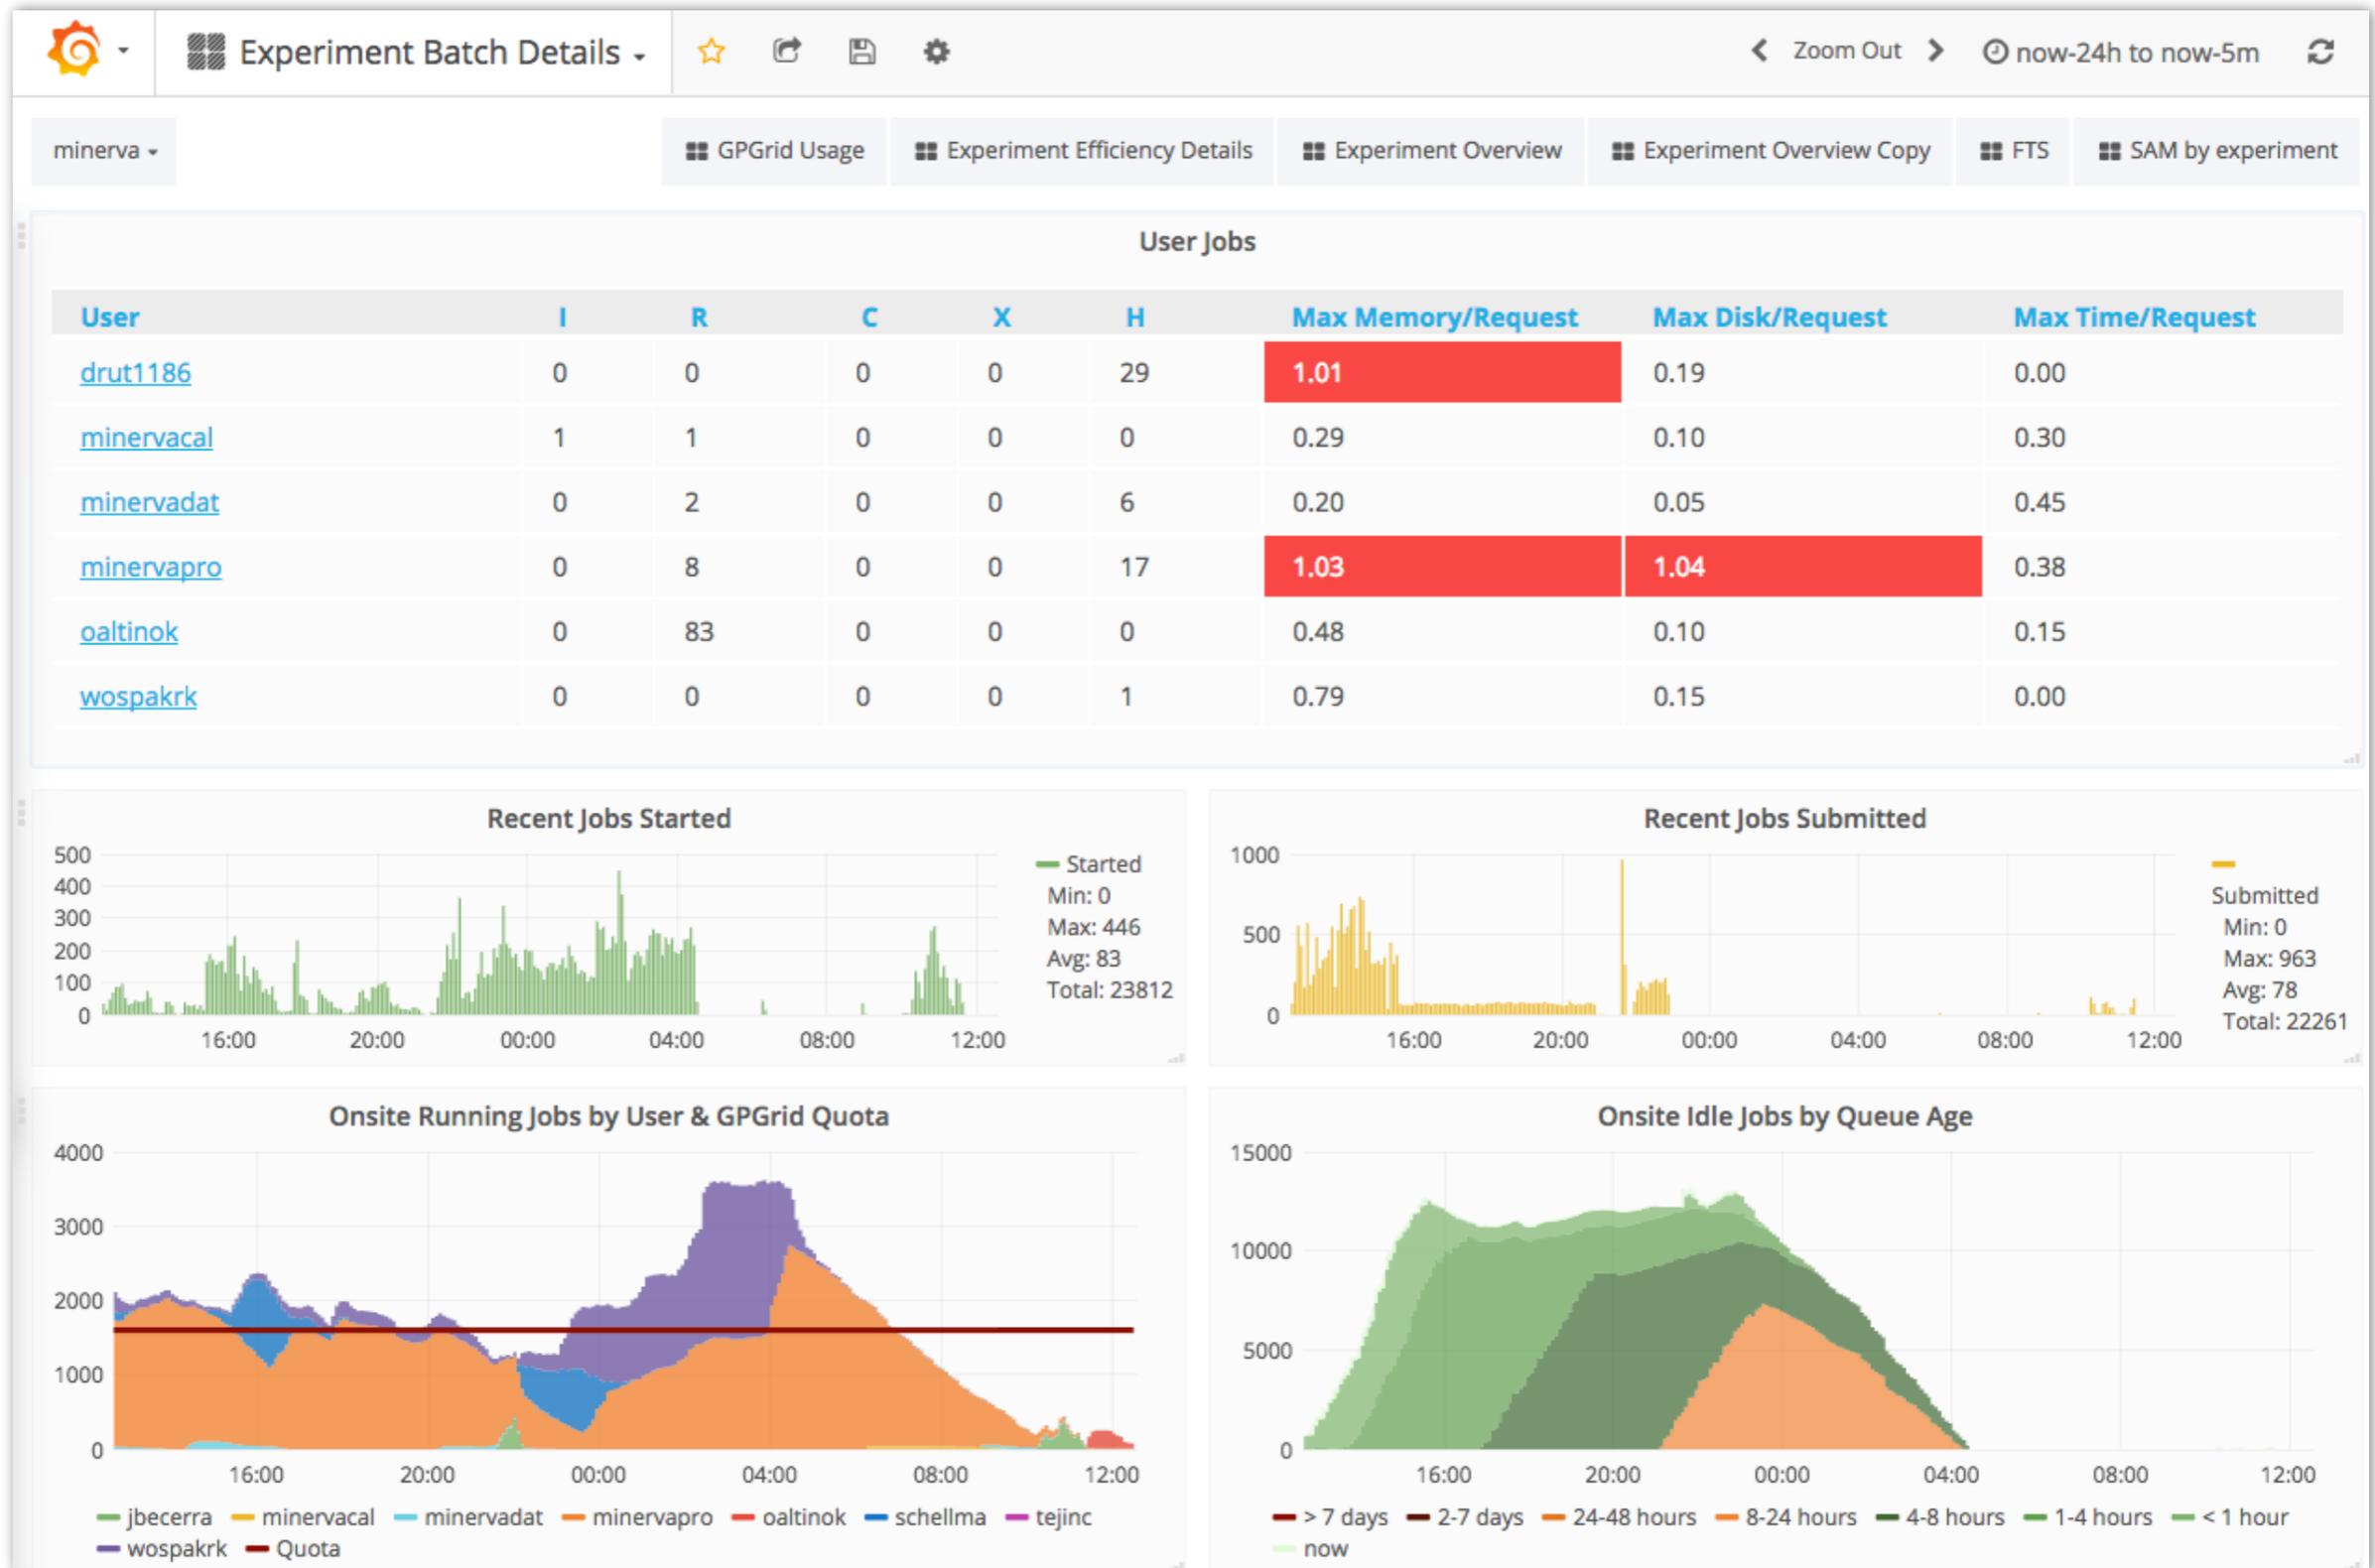
Time	Description	Details
2016-09-14 06:00:00	Authentication errors	GUMS authentication was failing



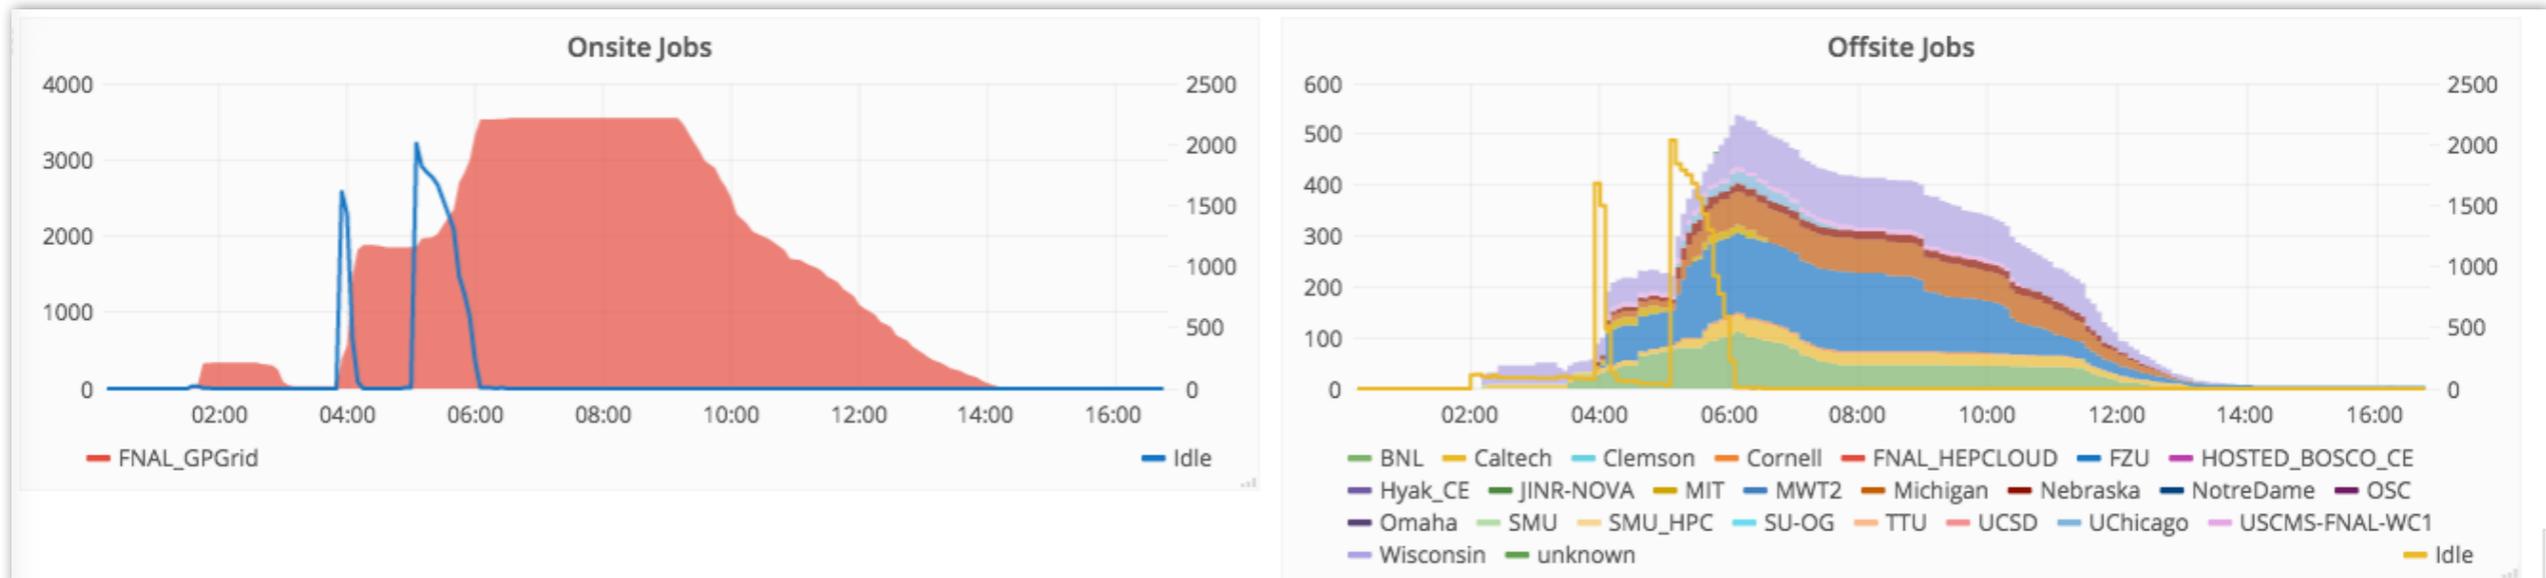
Fifemon - Experiment Summary



Fifemon - Experiment Batch Details



Fifemon - User Batch Details



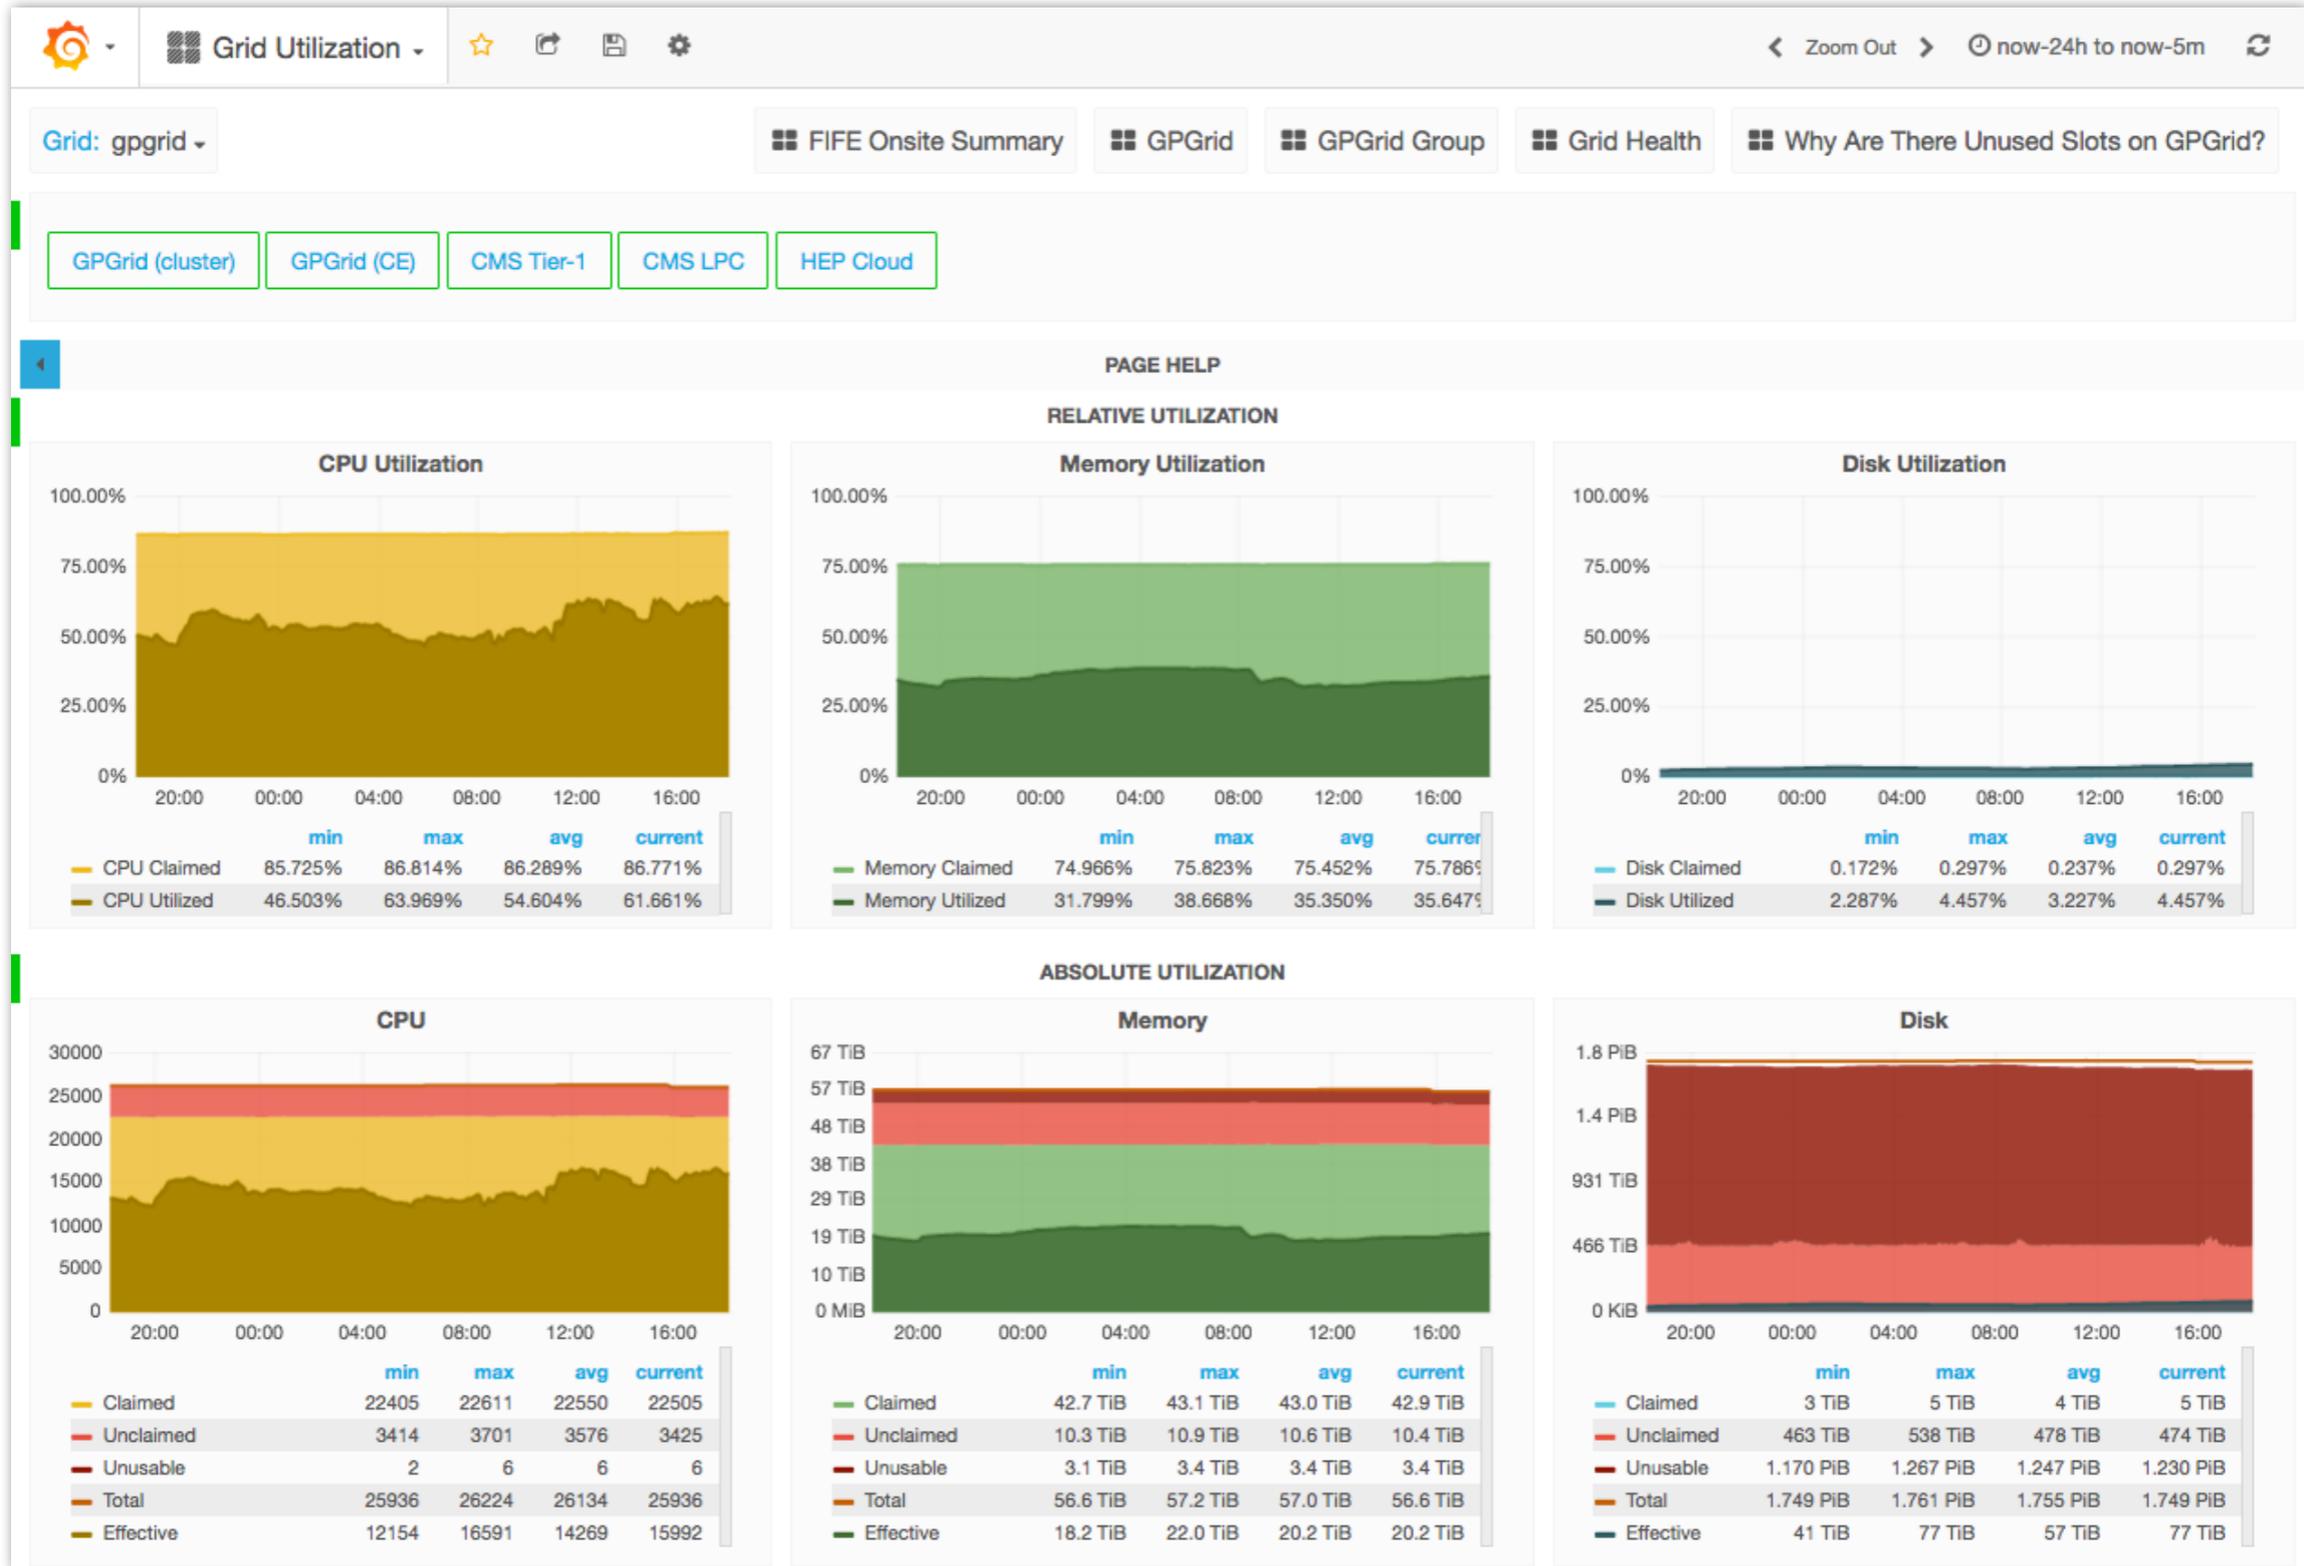
Current Job Clusters [↗](#)

2799 jobs in 21 clusters.

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 tghosh-prod_calibration_S15-03-11_fd_cosmic_preshutdown_draining-20161114_1256.sh_20161114_130049_2693892_0_1_wrap.sh	689 / 1950	2219 / 10000	0 / 0	----	1
11868298@fifebatch1.fnal.gov	0	0	3	Wed Nov 16 17:58 -0600 offsite_test_JINR.sh_20161116_175833_1168516_0_1_wrap.sh	0 / 500	0 / 1024	0 / 1	----	0
11875126@fifebatch1.fnal.gov	0	0	3	Wed Nov 16 22:09 -0600 offsite_test_JINR.sh_20161116_220946_2202517_0_1_wrap.sh	0 / 500	0 / 1024	0 / 1	----	0
14927498@fifebatch2.fnal.gov	0	0	6	Thu Nov 17 18:54 -0600 rijeeshk-horn_off_data_reco-20161117_1853.sh_20161117_185450_1430195_0_1_wrap.sh	7069 / 7000	915 / 5000	0 / 0	----	1
11899167@fifebatch1.fnal.gov	0	0	7	Fri Nov 18 02:10 -0600	2114 / 2000	172 / 10240	0 / 15	----	3

Fifemon - Grid/Cluster Utilization



Fifemon - Troubleshooting Guides

Are there FIFE jobs requesting onsite resources?

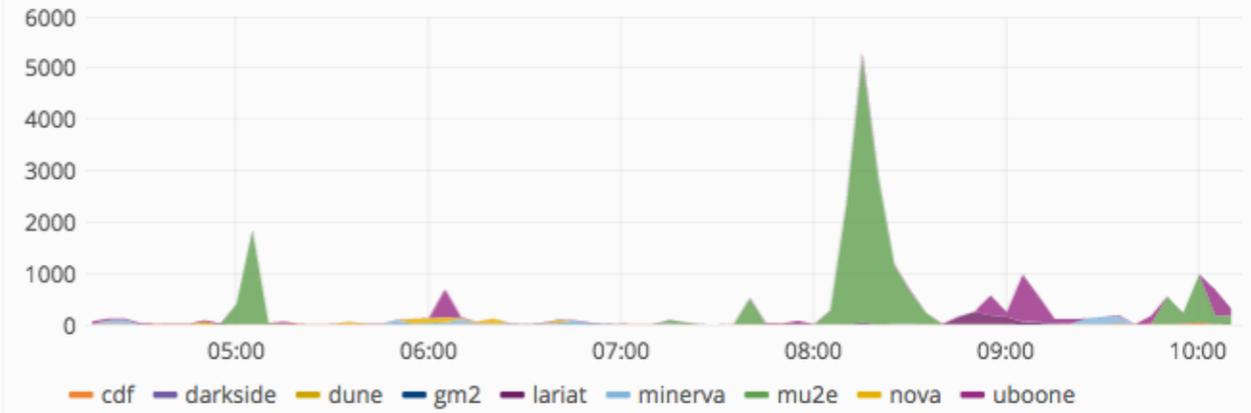
If jobs are requesting only OFFSITE, they will not run on GPGrid, unless they come back through the OSG opportunistic gatekeeper.

See also:

[FIFE Onsite Summary](#)

[Fifebatch - Onsite](#)

Idle Fifebatch Jobs Requesting Dedicated or Opportunistic (Onsite)



Are the remaining resources in the Glideins "unusable"?

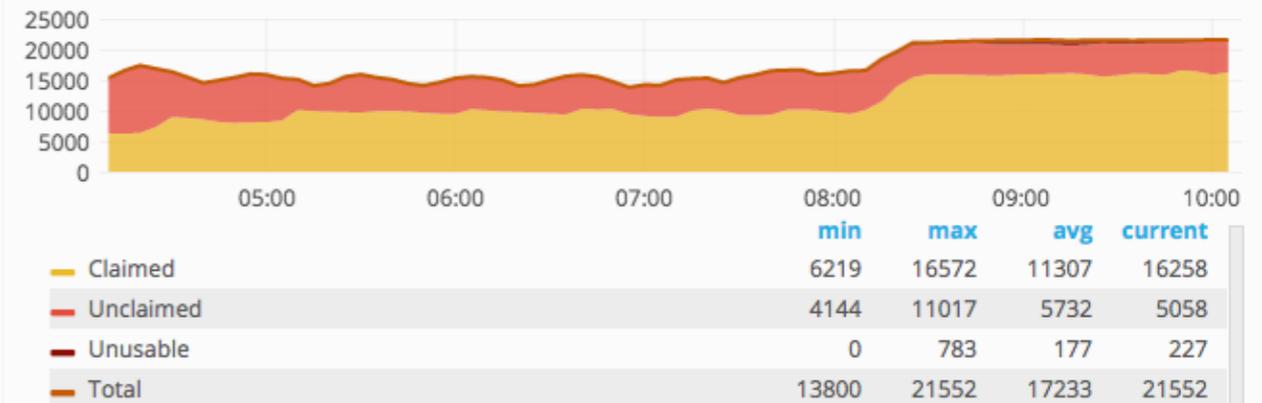
If there are lots of multicore or high-memory (>2 GB) jobs running there will be unusable resources left in the glideins.

See also:

[Grid Utilization \(GPGrid CE\)](#)

[Fifebatch Slots \(GPGrid\)](#)

GPGrid Glidein CPU Usage



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:

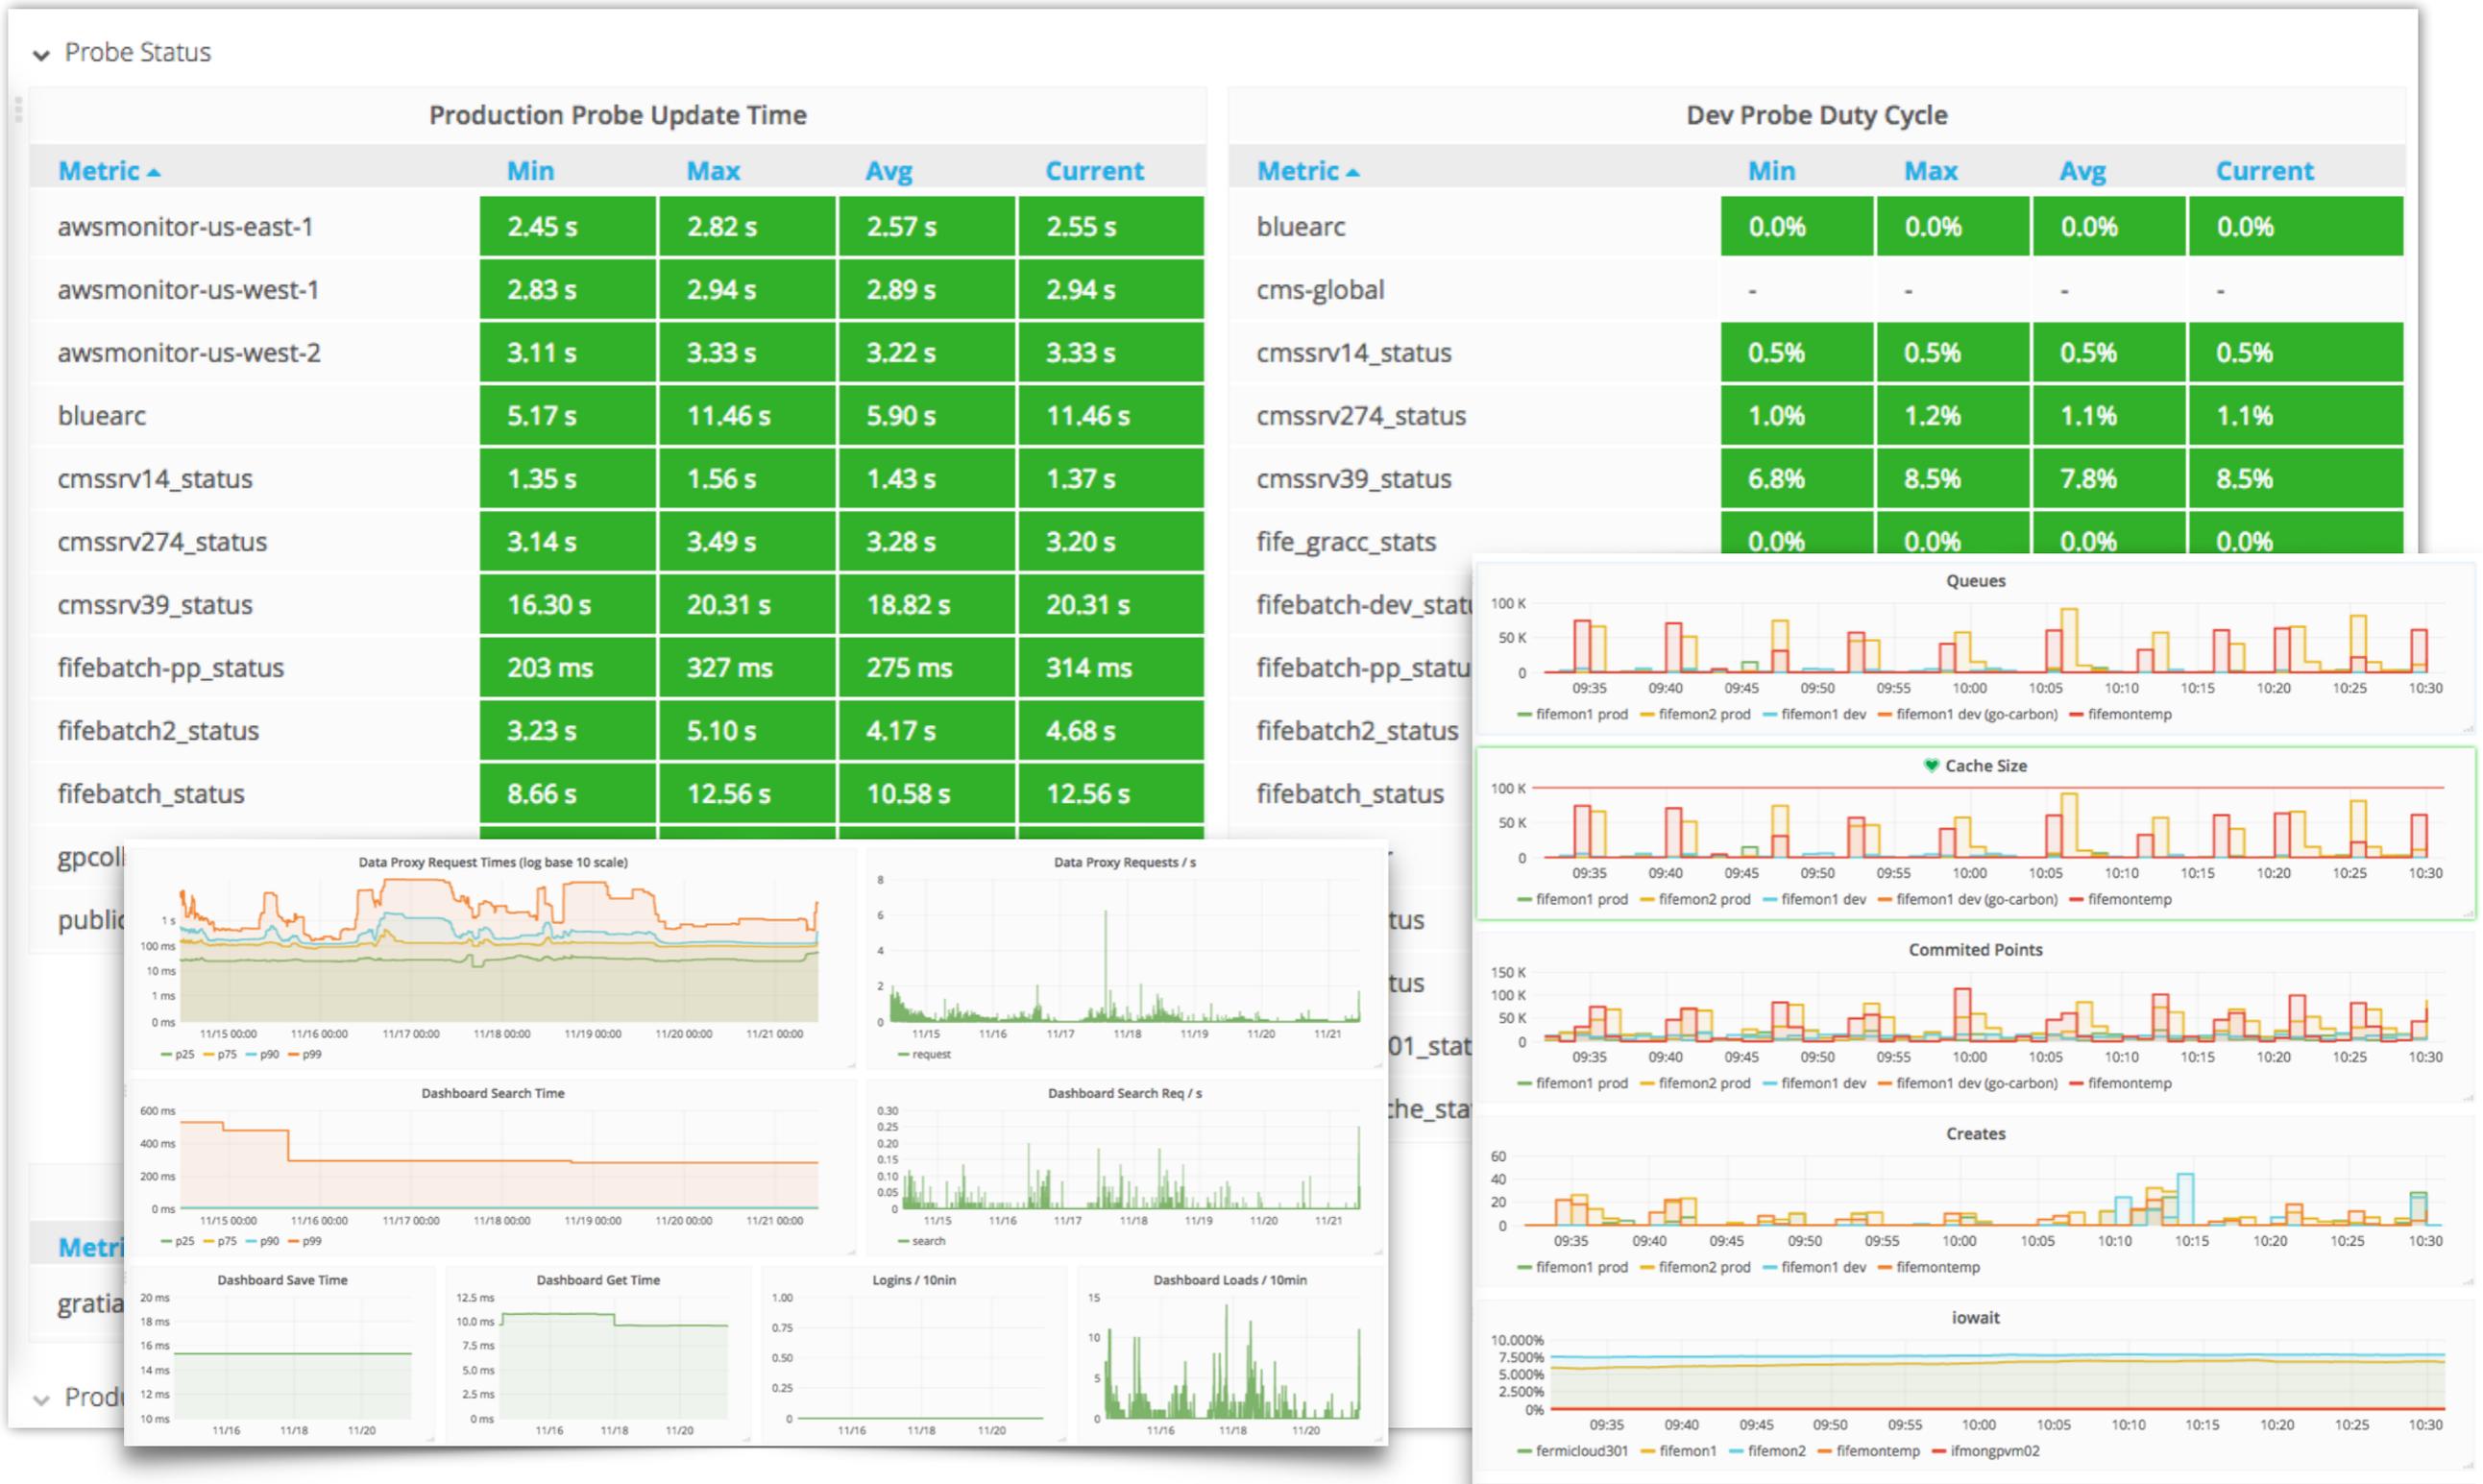
[Fifebatch Slots \(GPGrid\)](#)

[Fifebatch Slots Unclaimed \(GPGrid\)](#)

Slots with remaining resources exceeding JobSub defaults

96

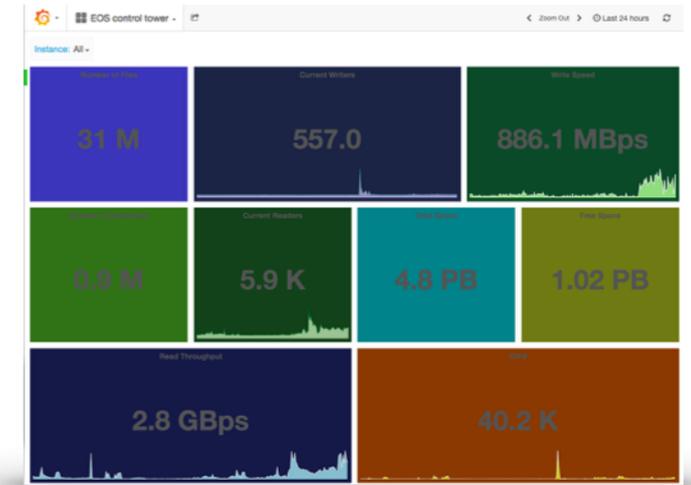
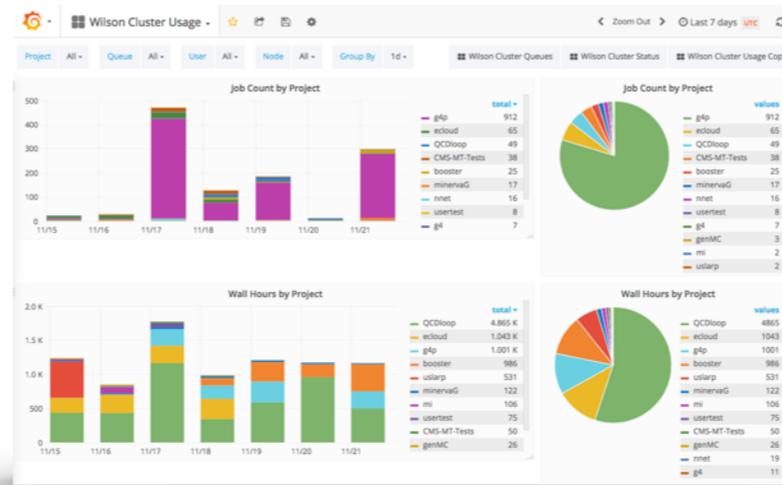
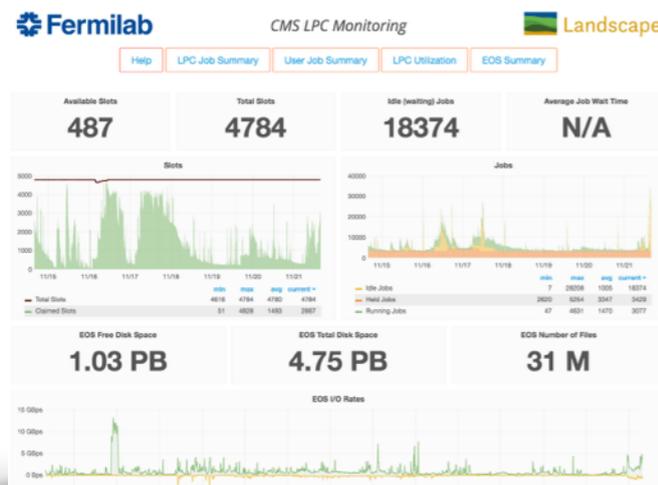
Fifemon - System Monitoring



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



Landscape



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

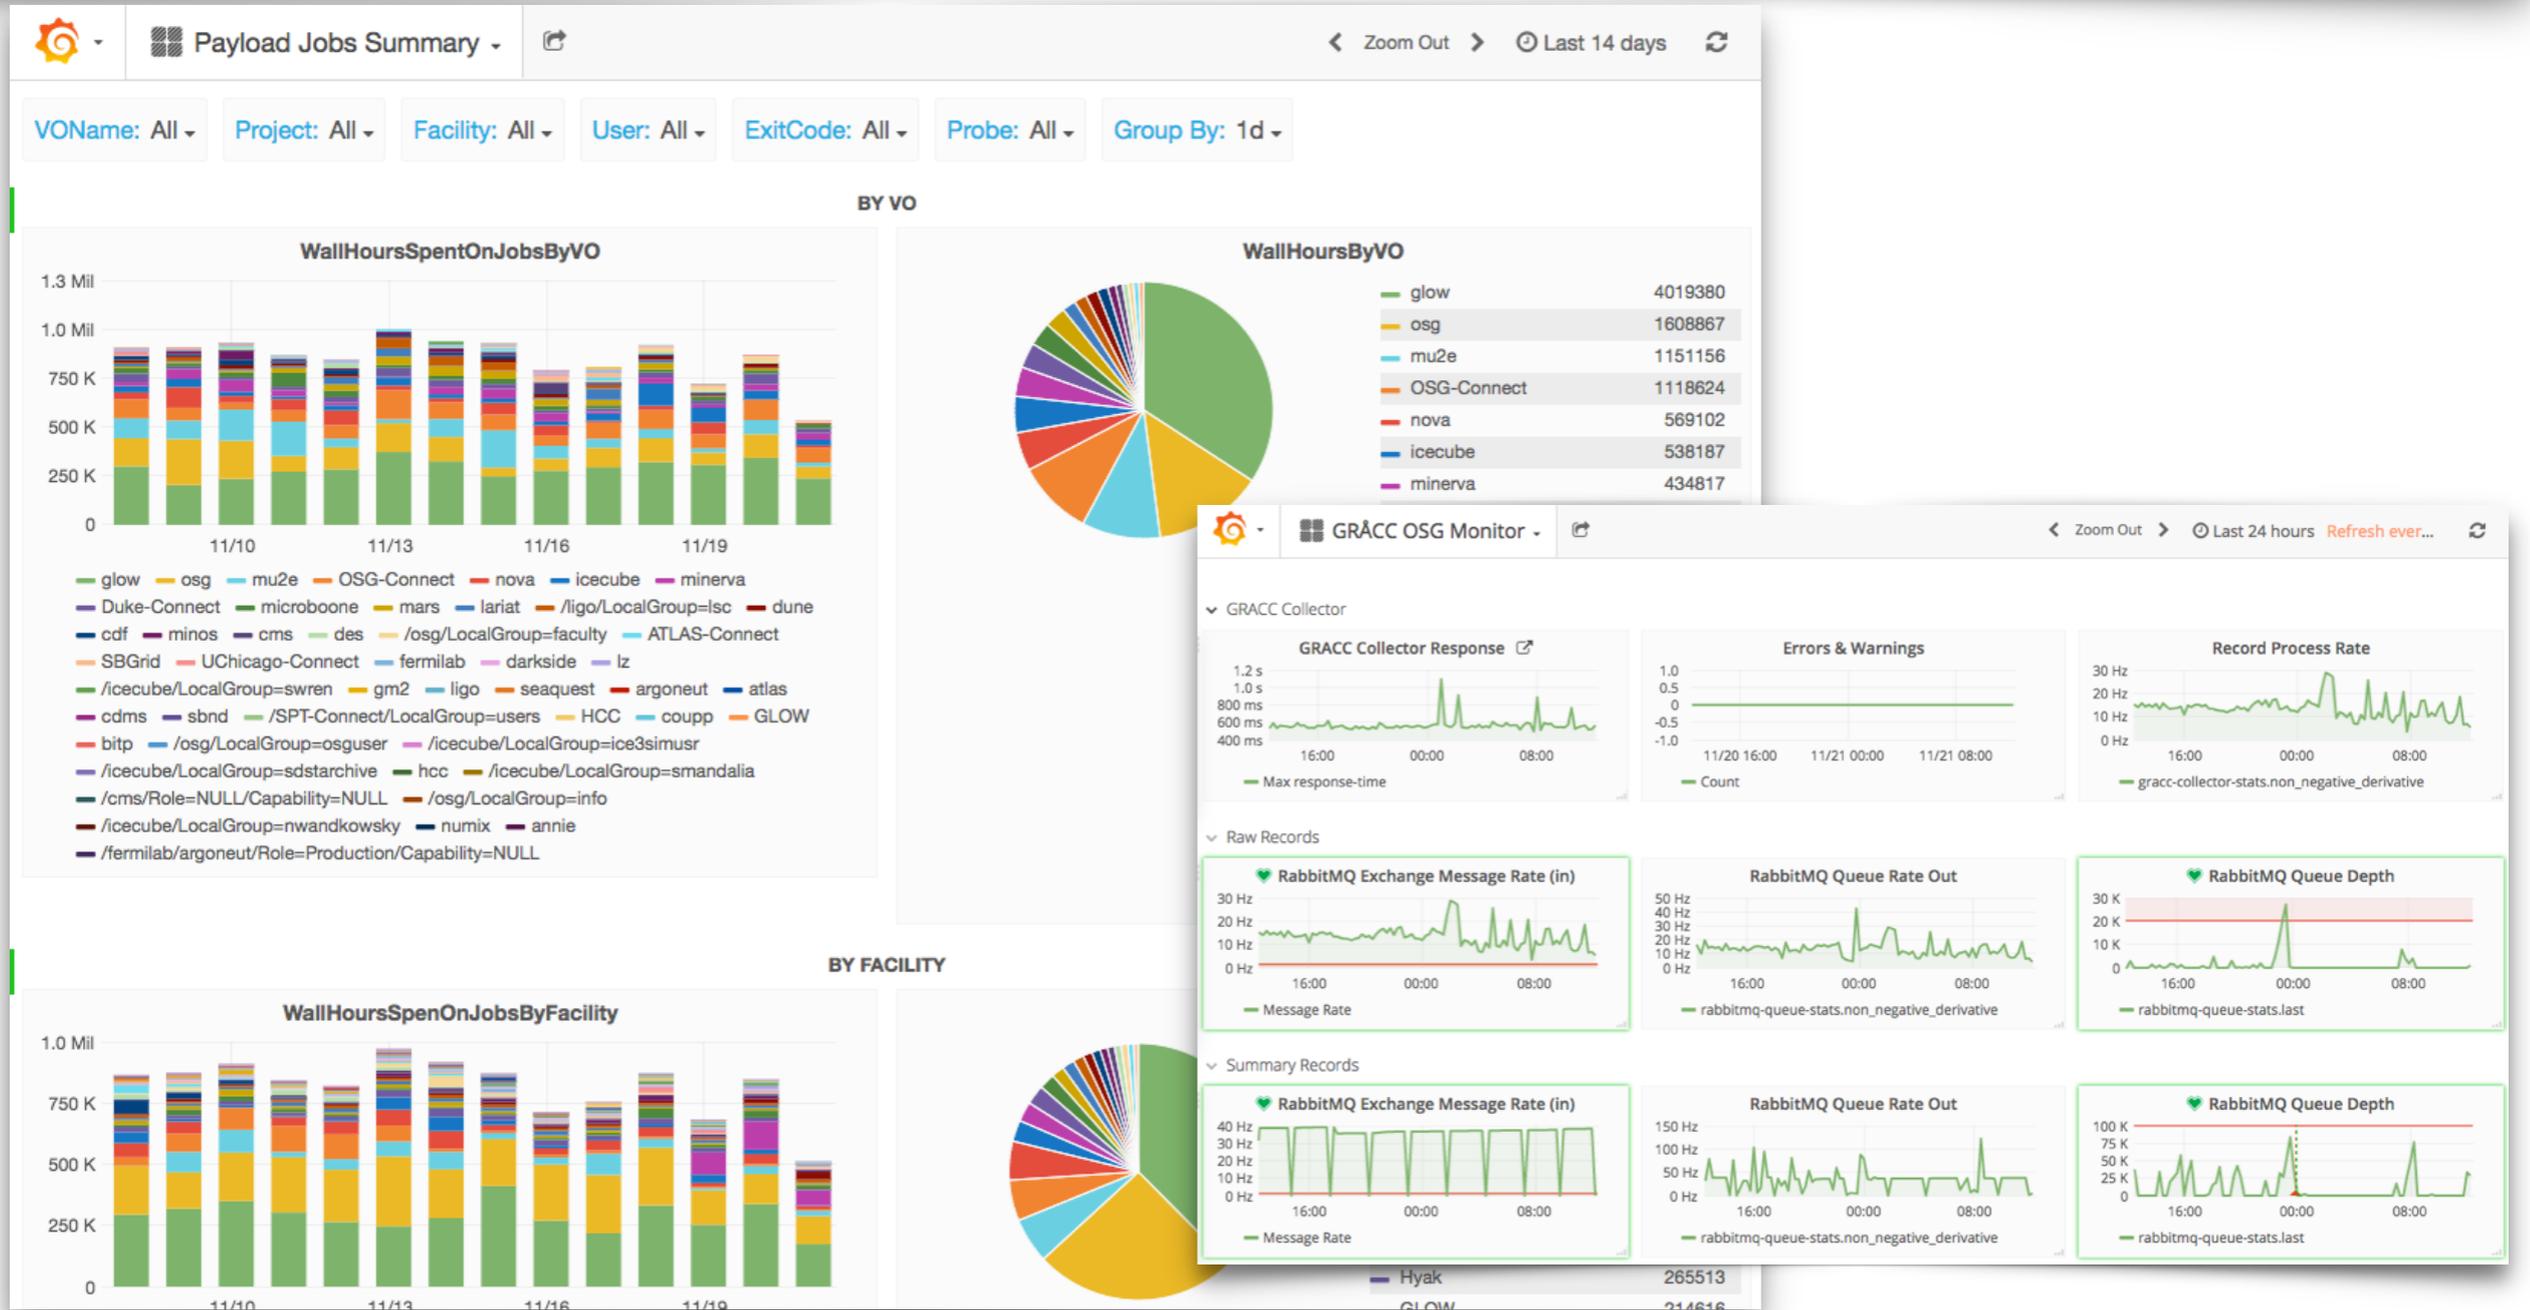


Open Science Grid



GRACC Grid Accounting

Status
OK



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.

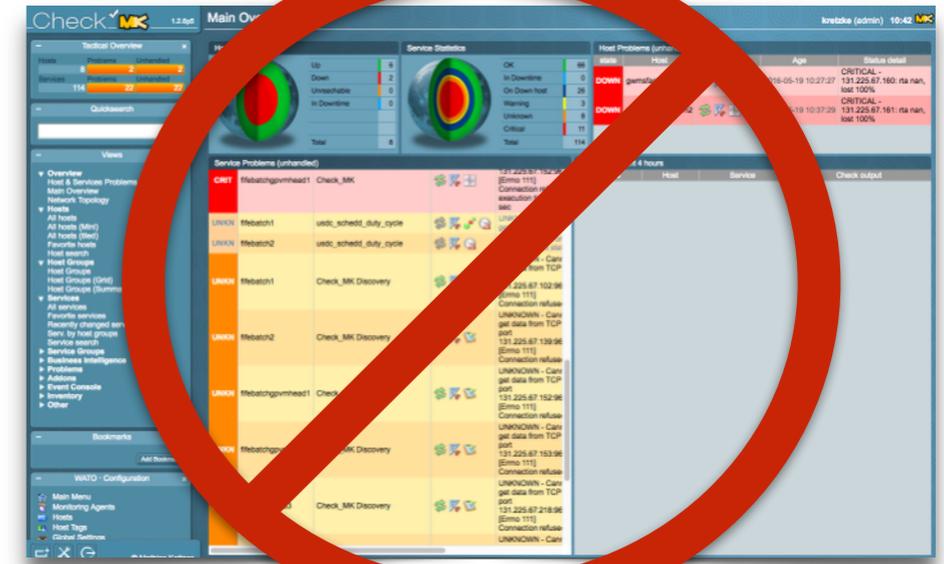
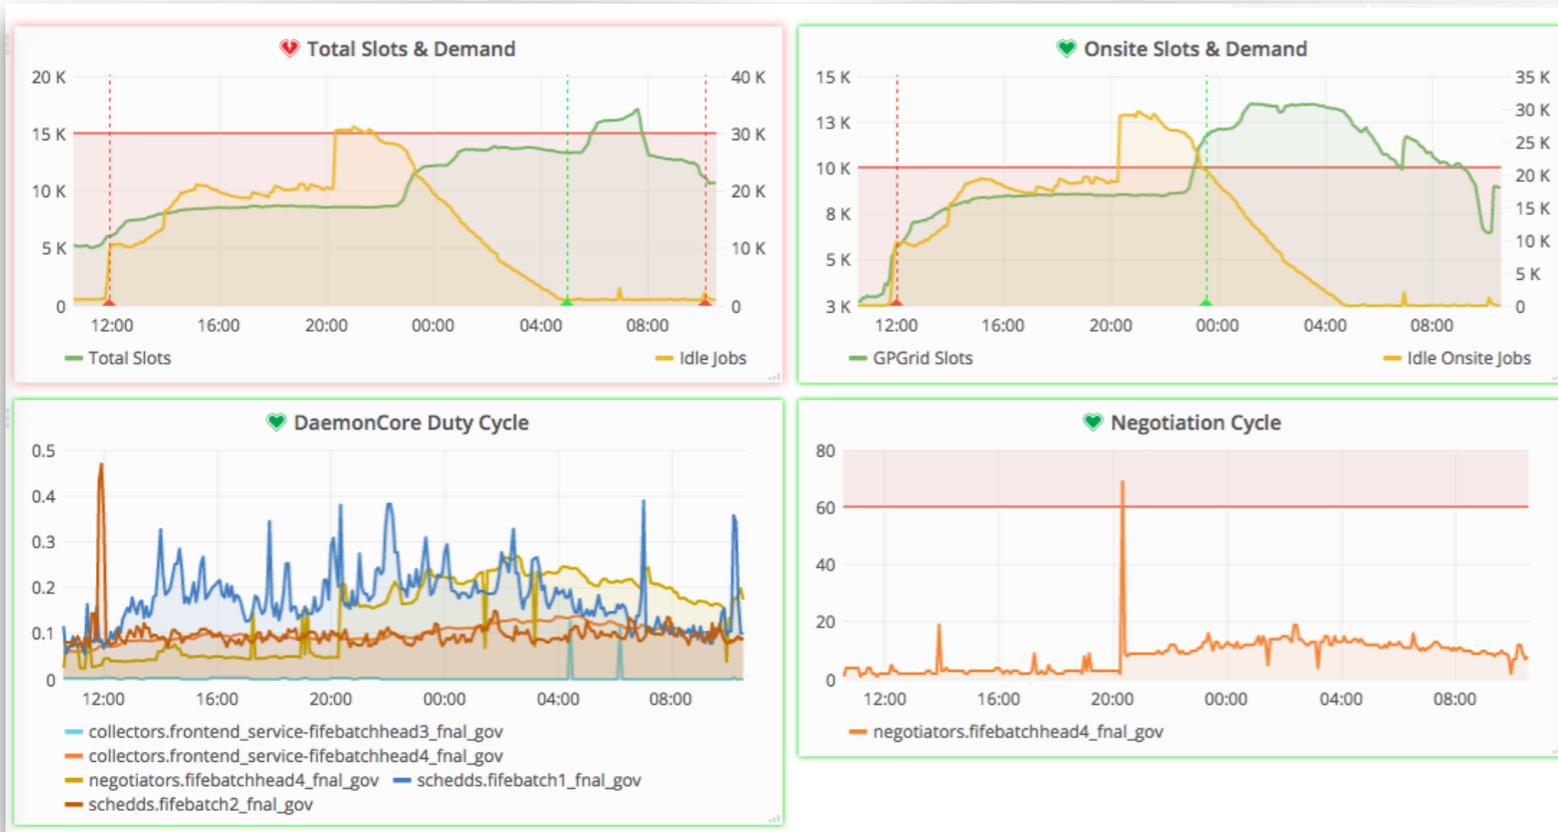
Built-in Auth

- 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!



Carbon Cache Size	OK for 15 hours	⚙
Data Proxy Request Times	OK for 2 days	⚙
EOS file offline	OK for 4 days	⚙
Fifebatch DaemonCore Duty Cycle	OK for 6 days	⚙
Fifebatch Negotiaton Cycle	OK for a day	⚙
Fifebatch Onsite Slots	OK for 11 hours	⚙
Fifebatch Total Slots	ALERTING for 26 minutes	⚙
Gratia MySQL Replication Lag	OK for 2 days	⚙
Pool Read Rates < 10 MBps		⚙

Grafana BOT 5:00 AM
[OK] Fifebatch Total Slots
 Total Slots
 13345.578
 Grafana v4.0.0-beta1 | Today at 5:00 AM

Grafana BOT 6:27 AM
[Alerting] Pool Write Rates < 5 MBps
 scratch
 4615314.5163569
 Grafana v4.0.0-beta1 | Today at 6:27 AM

Grafana BOT 6:33 AM
[Alerting] Pool Read Rates < 10 MBps
 testing 123
 scratch
 9464946.9738174
 Grafana v4.0.0-beta1 | Today at 6:33 AM

Grafana BOT 6:55 AM
[OK] Pool Read Rates < 10 MBps
 Grafana v4.0.0-beta1 | Today at 6:55 AM

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

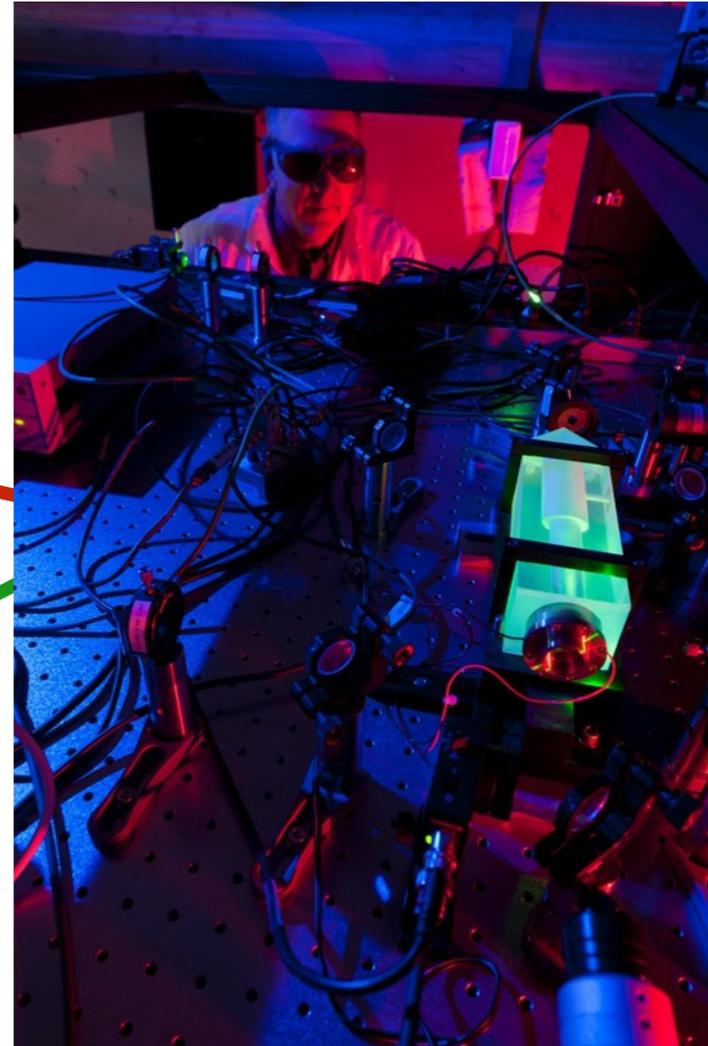
landscape.fnal.gov

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.