Comprehensive Grid and Job Monitoring with Fifemon

Kevin Retzke
User Support for Distributed Computing @ Fermilab

HTCondor Week 2016

Landscape

Fermilab

FIFE Project

FabrIc for Frontier Experiments:

Common computing for “not CMS” experiments at Fermilab

- O(10) experiments
- O(100) users
- O(10 000) simultaneous jobs
- O(1 000 000) jobs per week
- O(1 PB) data collected per month
- One global HTCondor pool (via GlideinWMS)
  - ~ ⅔ jobs run on dedicated local cluster
  - ~ ⅓ opportunistic through Open Science Grid

Why Do We Need Monitoring?

Grid admins want to know:

- Overall health of the batch system
- Worker node status and availability
- Efficiency in matching jobs to resources
- Identify and fix problems quickly (before users and stakeholders notice... and open tickets)

Users want to know:

- State of their jobs
- Availability of resources
- WHY ISN'T MY JOB RUNNING?

Stakeholders want to know:

- Each group is getting the resources it needs
- Resources are being used effectively

Fermigrid Monitor (ca. 2004)

Monitoring for local HTCondor cluster (GPGrid).

- Aggregate metrics for grid and VOs.
- No offsite information, no user job information.
- Difficult to alter or expand.

OK for grid admins, good for stakeholders, bad for users.

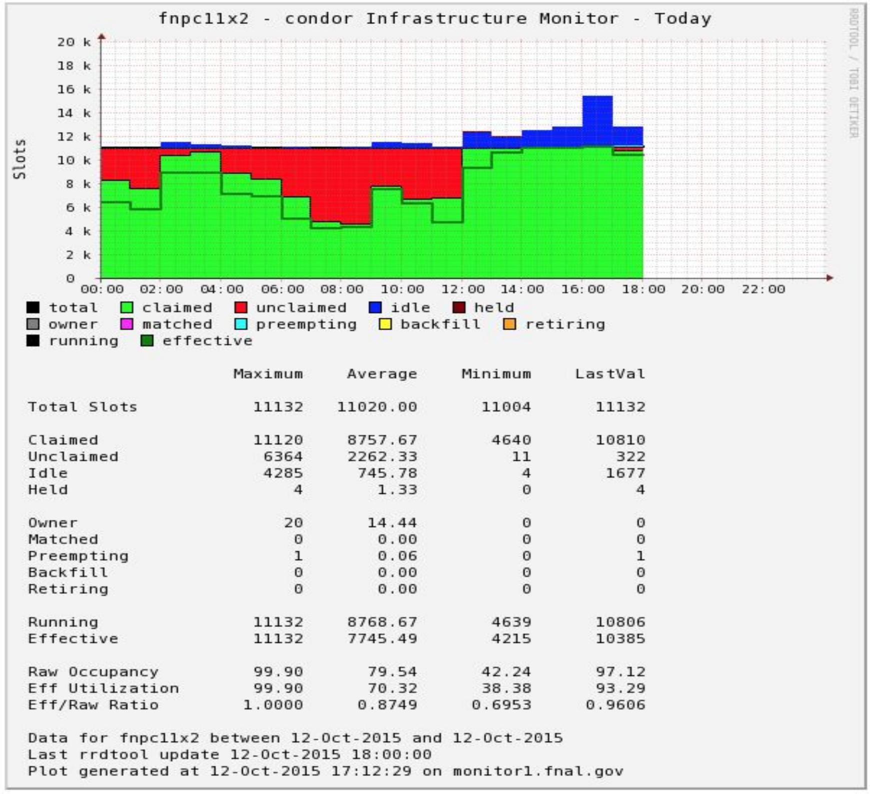

```
                        Maximum    Average    Minimum    LastVal

Total Slots               11132   11020.00      11004      11132

Claimed                   11120    8757.67       4640      10810
Unclaimed                  6364    2262.33         11        322
Idle                       4285     745.78          4       1677
Held                          4       1.33          0          4

Owner                        20      14.44          0          0
Matched                       0       0.00          0          0
Preempting                    1       0.06          0          1
Backfill                      0       0.00          0          0
Retiring                      0       0.00          0          0

Running                   11132    8768.67       4639      10806
Effective                 11132    7745.49       4215      10385

Raw Occupancy             99.90      79.54      42.24      97.12
Eff Utilization           99.90      70.32      38.38      93.29
Eff/Raw Ratio            1.0000     0.8749     0.6953     0.9606

Data for fnpc11x2 between 12-Oct-2015 and 12-Oct-2015
Last rrdtool update 12-Oct-2015 18:00:00
Plot generated at 12-Oct-2015 17:12:29 on monitor1.fnal.gov
```

Fifemon v1 (ca. 2014)

Growing usage of offsite resources through OSG; needed new monitoring.

- Aggregate metrics for users and VOs.
- No cluster information.
- Cumbersome to maintain and expand.

OK for grid admins, bad for stakeholders, good for users.

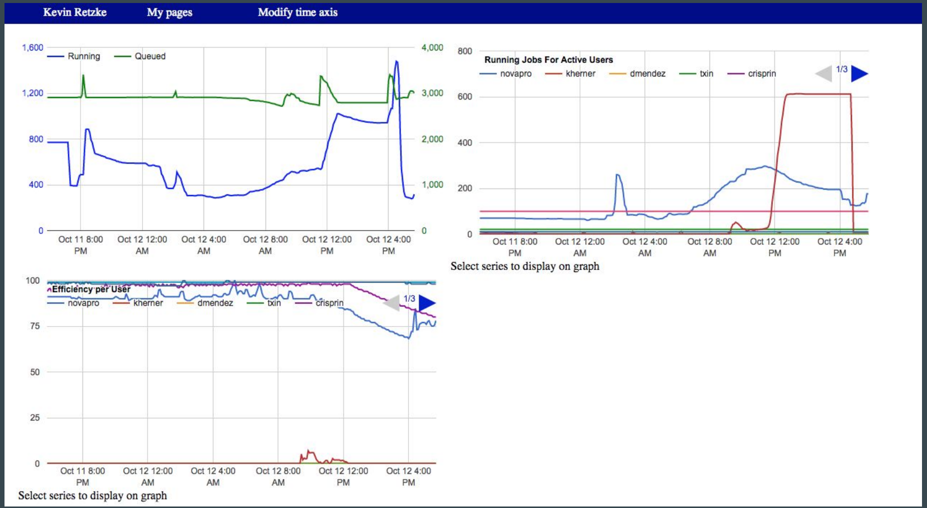

Fifemon v2+ (ca. 2015)

Landscape Program: develop comprehensive monitoring for FIFE, HEP Cloud, and beyond.

- Leverage open-source monitoring technology
- Focus on incorporating new data sources and new dashboards
- Rapid development and iteration of tailored views for each target audience.

Good for grid admins, stakeholders, and users alike.

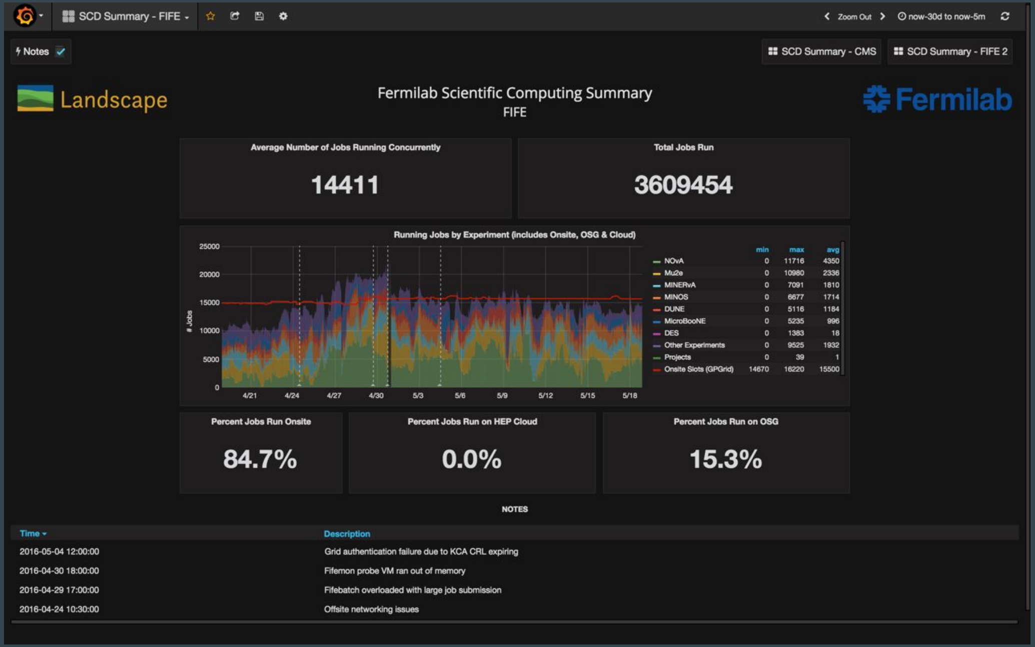

Fermilab Scientific Computing Summary
FIFE

Average Number of Jobs Running Concurrently
14411

Total Jobs Run
3609454

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

	min	max	avg
NOvA	0	11716	4350
Mu2e	0	10980	2336
MINERvA	0	7091	1810
MINOS	0	6677	1714
DUNE	0	5116	1184
MicroBooNE	0	5235	996
DES	0	1383	18
Other Experiments	0	9525	1932
Projects	0	39	1
Onsite Slots (GPGrid)	14670	16220	15500

Percent Jobs Run Onsite
84.7%

Percent Jobs Run on HEP Cloud
0.0%

Percent Jobs Run on OSG
15.3%

NOTES

Time	Description
2016-05-04 12:00:00	Grid authentication failure due to KCA CRL expiring
2016-04-30 18:00:00	Fifemon probe VM ran out of memory
2016-04-29 17:00:00	Fifebatch overloaded with large job submission
2016-04-24 10:30:00	Offsite networking issues

Fifemon Backend

Data collection:

- Generic HTCondor probe collecting daemon, machine, and job status
- Logstash collecting live HTCondor Events
- Several other centrally-run probes querying other resources
- Some services directly reporting to Graphite

- Most probes report stats every five minutes
- Graphite:
  - 250K individual metrics
  - ~80GB
  - 10 year history
- Elasticsearch: ~8GB per day

Graphite:

- Time-series database, stores data in files similar to RRD with caching layer.
- Simple line protocol
- Powerful query manipulations and aggregations

Elasticsearch:

- “NoSQL” document database, powered by Apache Lucene.
- Store full details on current jobs, batch slots, and logs.

Fifemon Frontend

Grafana:

- Time-series (primarily) visualization dashboard platform.
- Supports numerous data sources (Graphite, InfluxDB, Elasticsearch, etc).
- Several auth methods (LDAP, OAuth, proxy).
- Rich user interface for graphing metrics and composing dashboards.
- Scripted and templated dashboards and raw HTML panels allow extensive customization.
- V3 (released last week) introduces new plugin system to support custom datasources and panels.

Kibana:

- Elasticsearch data only
    - Current jobs and machine status
    - Event logs
- Explore data, create ad-hoc visualizations, combine into dashboards
- Used for analytics and troubleshooting
- Access limited to grid admins and power users

Fifemon Architecture

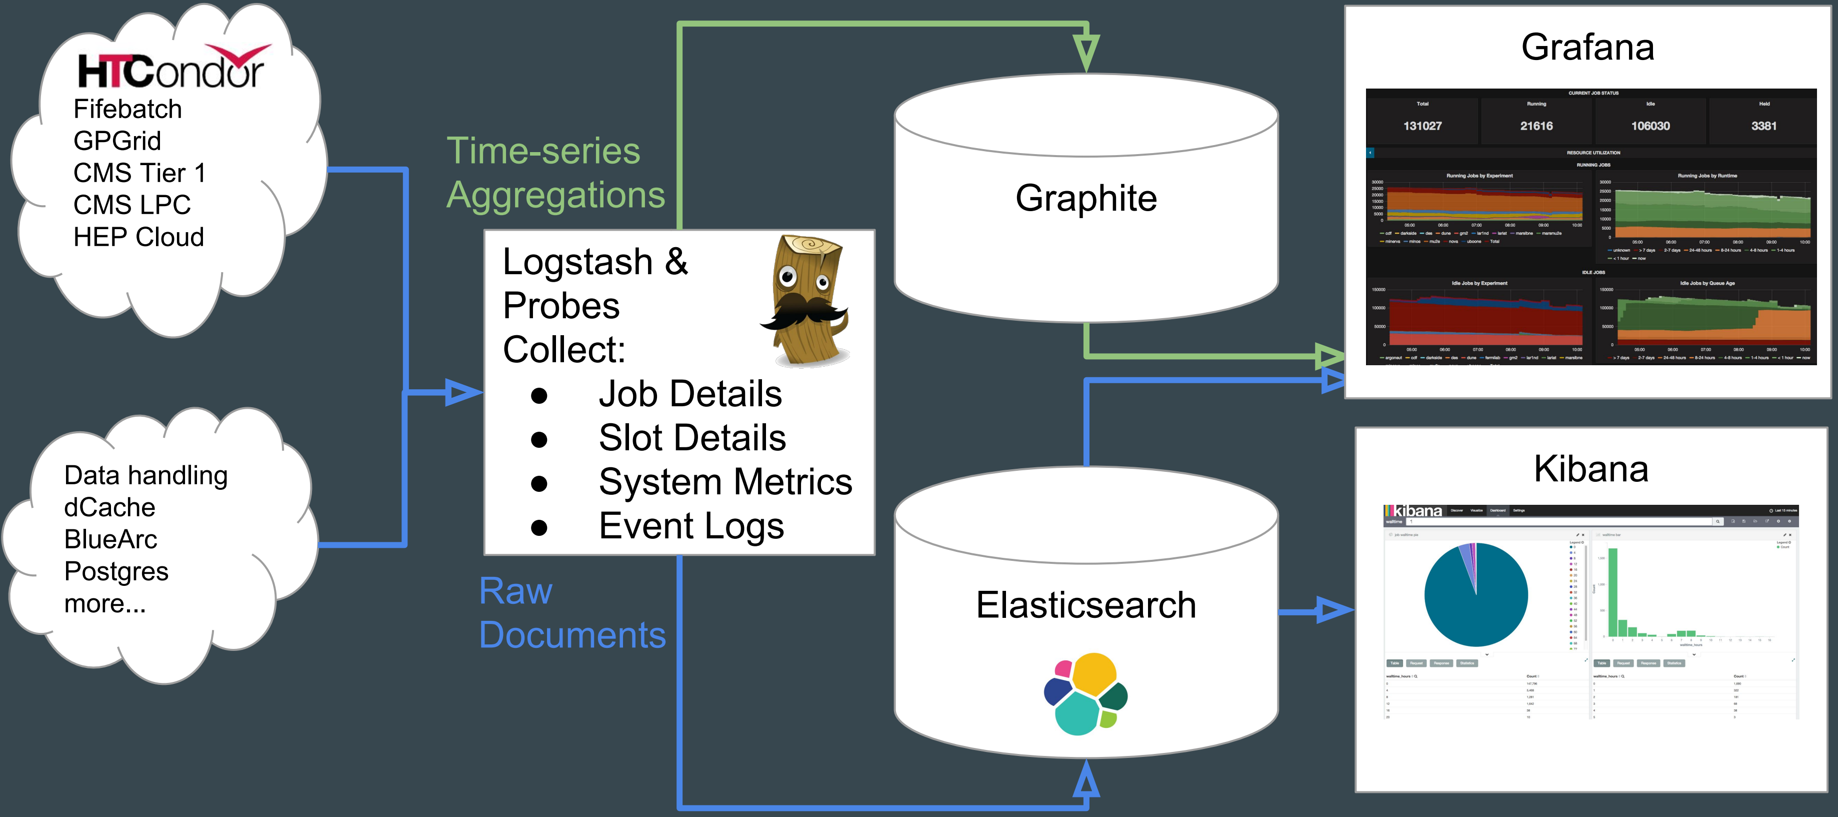

Next Steps

Fifemon is constantly evolving:

- Adding new data sources and metrics
- New dashboards:
  - Tailored views based on user request
  - Discovering new ways of looking at the data
- New Grafana panels
- Further leverage HTCondor event logs & gangliad/metricsd for true real-time monitoring

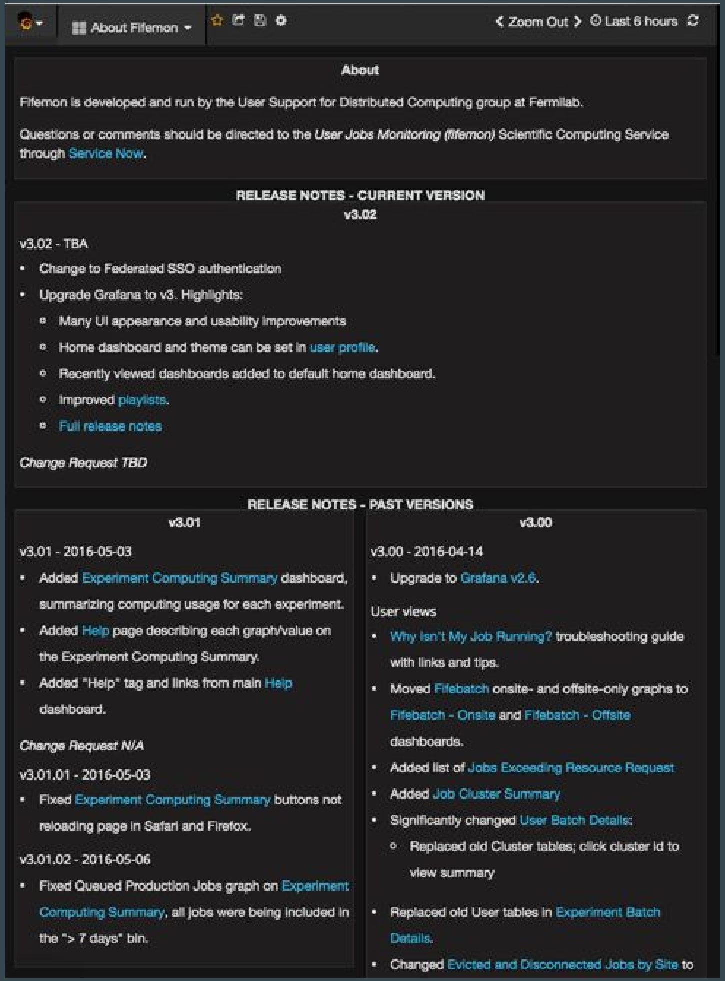

About

Fifemon is developed and run by the User Support for Distributed Computing group at Fermilab.

Questions or comments should be directed to the *User Jobs Monitoring (fifemon)* Scientific Computing Service through Service Now.

RELEASE NOTES - CURRENT VERSION

v3.02

v3.02 - TBA

- Change to Federated SSO authentication
- Upgrade Grafana to v3. Highlights:
  - Many UI appearance and usability improvements
  - Home dashboard and theme can be set in user profile.
  - Recently viewed dashboards added to default home dashboard.
  - Improved playlists.
  - Full release notes

Change Request TBD

RELEASE NOTES - PAST VERSIONS

v3.01

v3.01 - 2016-05-03

- Added Experiment Computing Summary dashboard, summarizing computing usage for each experiment.
- Added Help page describing each graph/value on the Experiment Computing Summary.
- Added "Help" tag and links from main Help dashboard.

Change Request N/A

v3.01.01 - 2016-05-03

- Fixed Experiment Computing Summary buttons not reloading page in Safari and Firefox.

v3.01.02 - 2016-05-06

- Fixed Queued Production Jobs graph on Experiment Computing Summary, all jobs were being included in the "> 7 days" bin.

v3.00

v3.00 - 2016-04-14

- Upgrade to Grafana v2.6.

User views

- Why isn't My Job Running? troubleshooting guide with links and tips.
- Moved Fifebatch onsite- and offsite-only graphs to Fifebatch - Onsite and Fifebatch - Offsite dashboards.
- Added list of Jobs Exceeding Resource Request
- Added Job Cluster Summary
- Significantly changed User Batch Details:
  - Replaced old Cluster tables; click cluster id to view summary
- Replaced old User tables in Experiment Batch Details.
- Changed Evicted and Disconnected Jobs by Site to

Case Studies

“There’s a dashboard for that…”

Case Study: Grid Admin

“Is the batch system healthy?”

Cinnamon

Photo:Public Domain

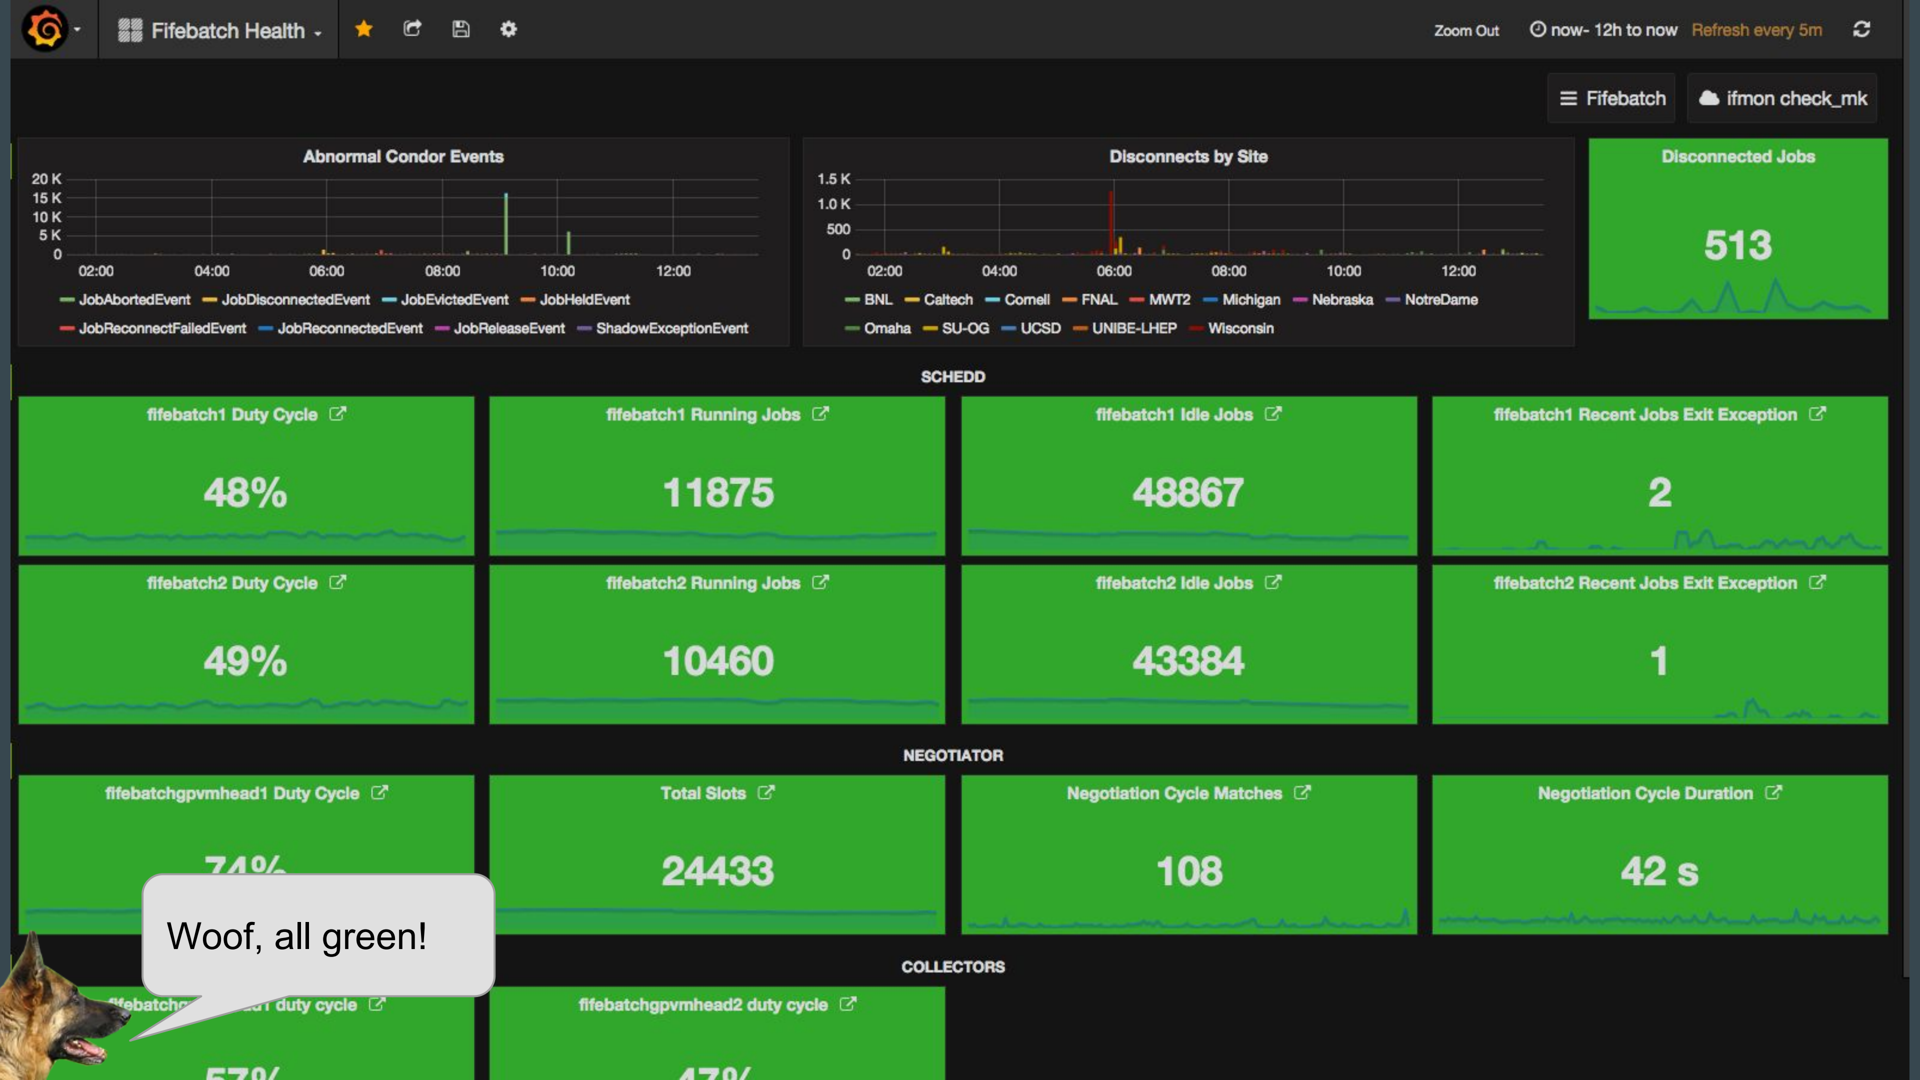

Disconnected Jobs: 513

SCHEDD

fifebatch1 Duty Cycle: 48%

fifebatch1 Running Jobs: 11875

fifebatch1 Idle Jobs: 48867

fifebatch1 Recent Jobs Exit Exception: 2

fifebatch2 Duty Cycle: 49%

fifebatch2 Running Jobs: 10460

fifebatch2 Idle Jobs: 43384

fifebatch2 Recent Jobs Exit Exception: 1

NEGOTIATOR

fifebatchgpvmhead1 Duty Cycle: 74%

Total Slots: 24433

Negotiation Cycle Matches: 108

Negotiation Cycle Duration: 42 s

COLLECTORS

fifebatchg... 1 duty cycle: 57%

fifebatchgpvmhead2 duty cycle: 47%

Woof, all green!

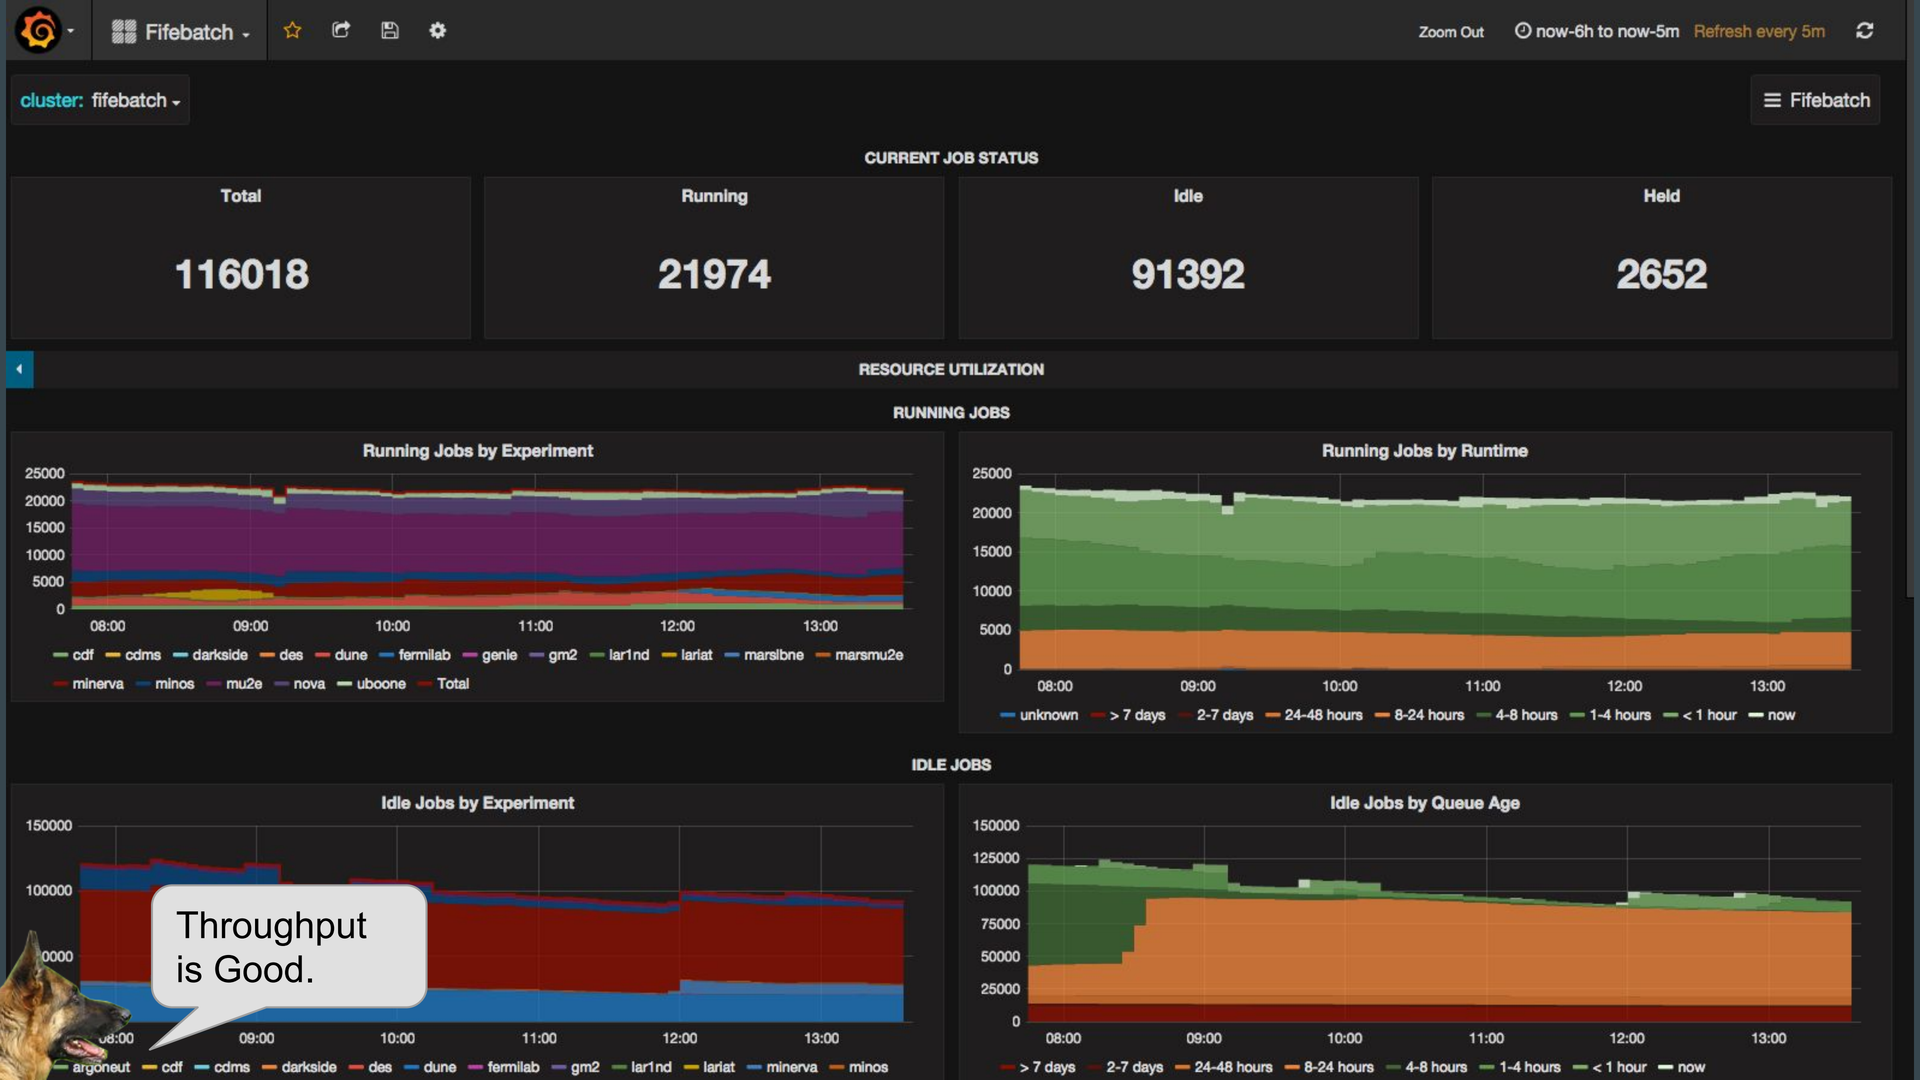

CURRENT JOB STATUS

Total	Running	Idle	Held
116018	21974	91392	2652

RESOURCE UTILIZATION

RUNNING JOBS

Running Jobs by Experiment

Running Jobs by Runtime

IDLE JOBS

Idle Jobs by Experiment

Idle Jobs by Queue Age

Throughput is Good.

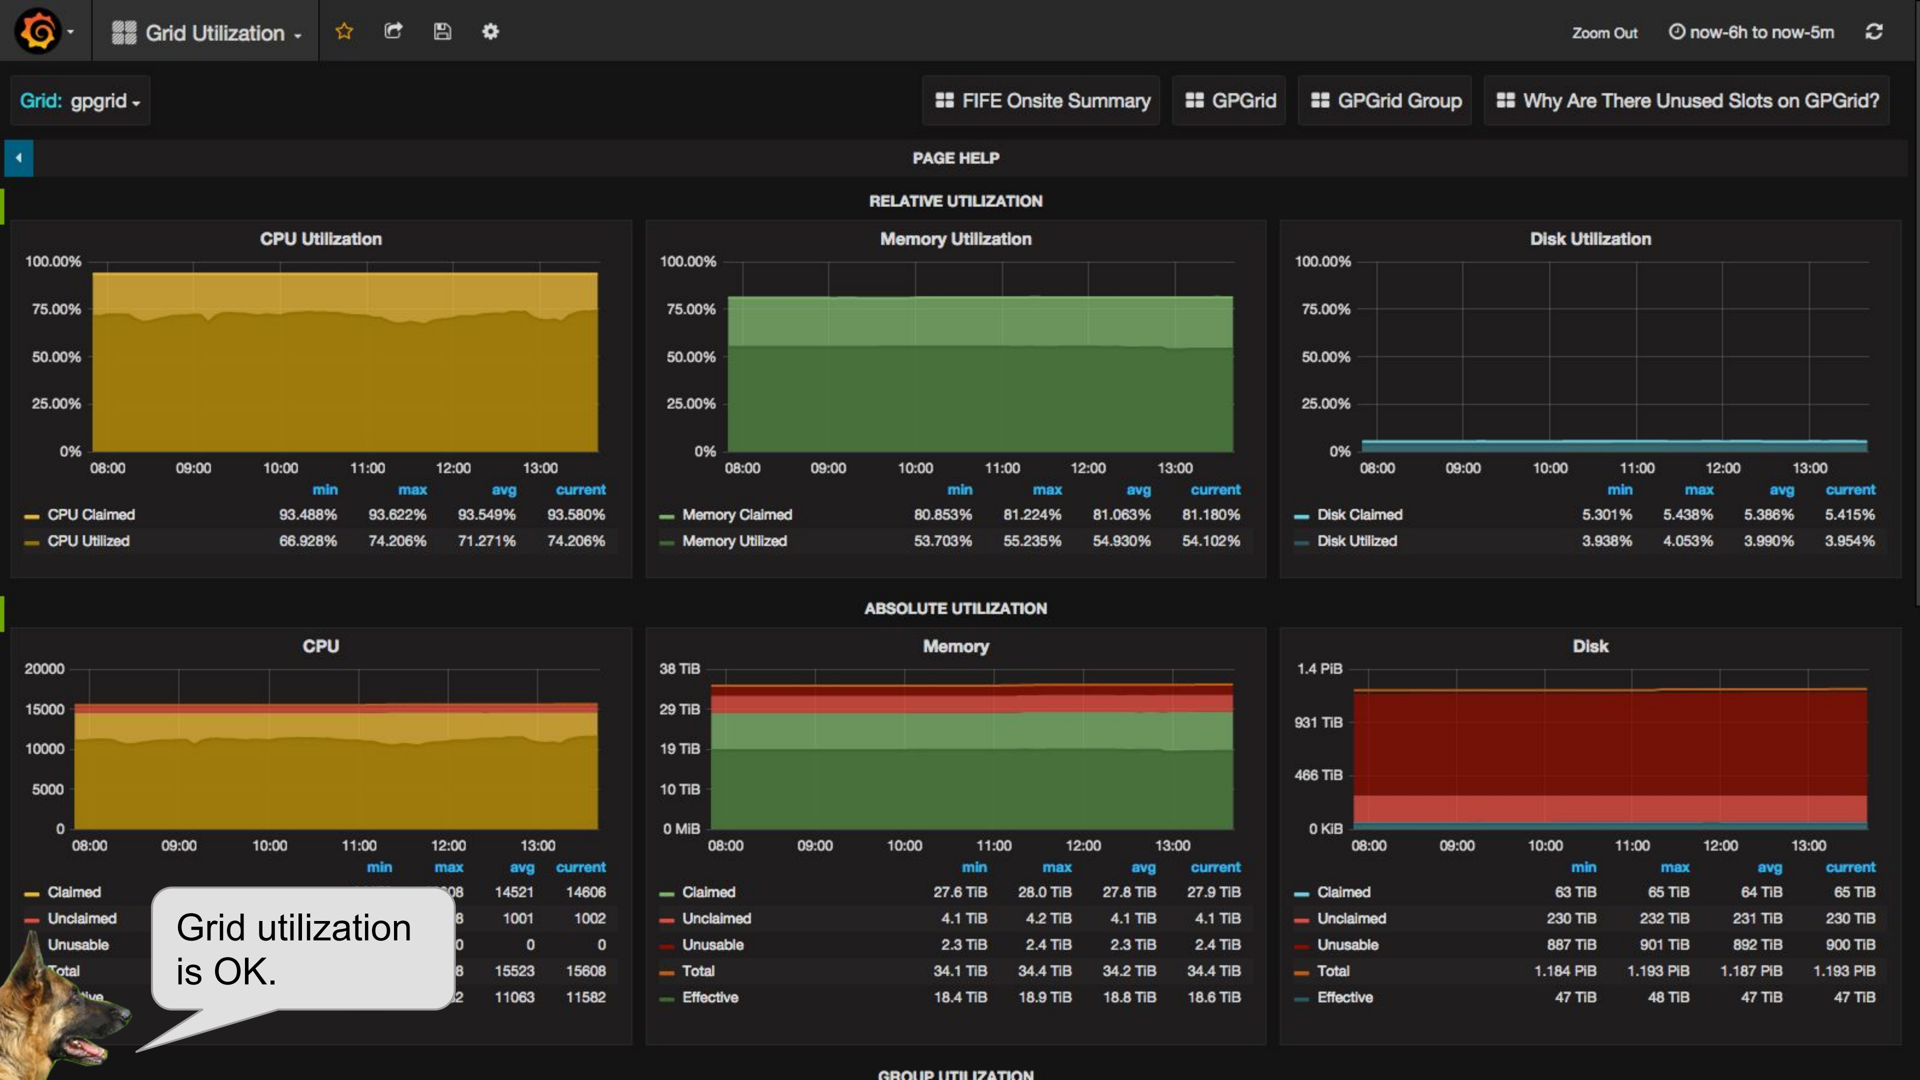

RELATIVE UTILIZATION

CPU Utilization

	min	max	avg	current
CPU Claimed	93.488%	93.622%	93.549%	93.580%
CPU Utilized	66.928%	74.206%	71.271%	74.206%

Memory Utilization

	min	max	avg	current
Memory Claimed	80.853%	81.224%	81.063%	81.180%
Memory Utilized	53.703%	55.235%	54.930%	54.102%

Disk Utilization

	min	max	avg	current
Disk Claimed	5.301%	5.438%	5.386%	5.415%
Disk Utilized	3.938%	4.053%	3.990%	3.954%

ABSOLUTE UTILIZATION

CPU

	min	max	avg	current
Claimed	[illegible]	[illegible]	14521	14606
Unclaimed	[illegible]	[illegible]	1001	1002
Unusable	[illegible]	0	0	0
Total	[illegible]	[illegible]	15523	15608
Effective	[illegible]	[illegible]	11063	11582

Memory

	min	max	avg	current
Claimed	27.6 TiB	28.0 TiB	27.8 TiB	27.9 TiB
Unclaimed	4.1 TiB	4.2 TiB	4.1 TiB	4.1 TiB
Unusable	2.3 TiB	2.4 TiB	2.3 TiB	2.4 TiB
Total	34.1 TiB	34.4 TiB	34.2 TiB	34.4 TiB
Effective	18.4 TiB	18.9 TiB	18.8 TiB	18.6 TiB

Disk

	min	max	avg	current
Claimed	63 TiB	65 TiB	64 TiB	65 TiB
Unclaimed	230 TiB	232 TiB	231 TiB	230 TiB
Unusable	887 TiB	901 TiB	892 TiB	900 TiB
Total	1.184 PiB	1.193 PiB	1.187 PiB	1.193 PiB
Effective	47 TiB	48 TiB	47 TiB	47 TiB

Grid utilization is OK.

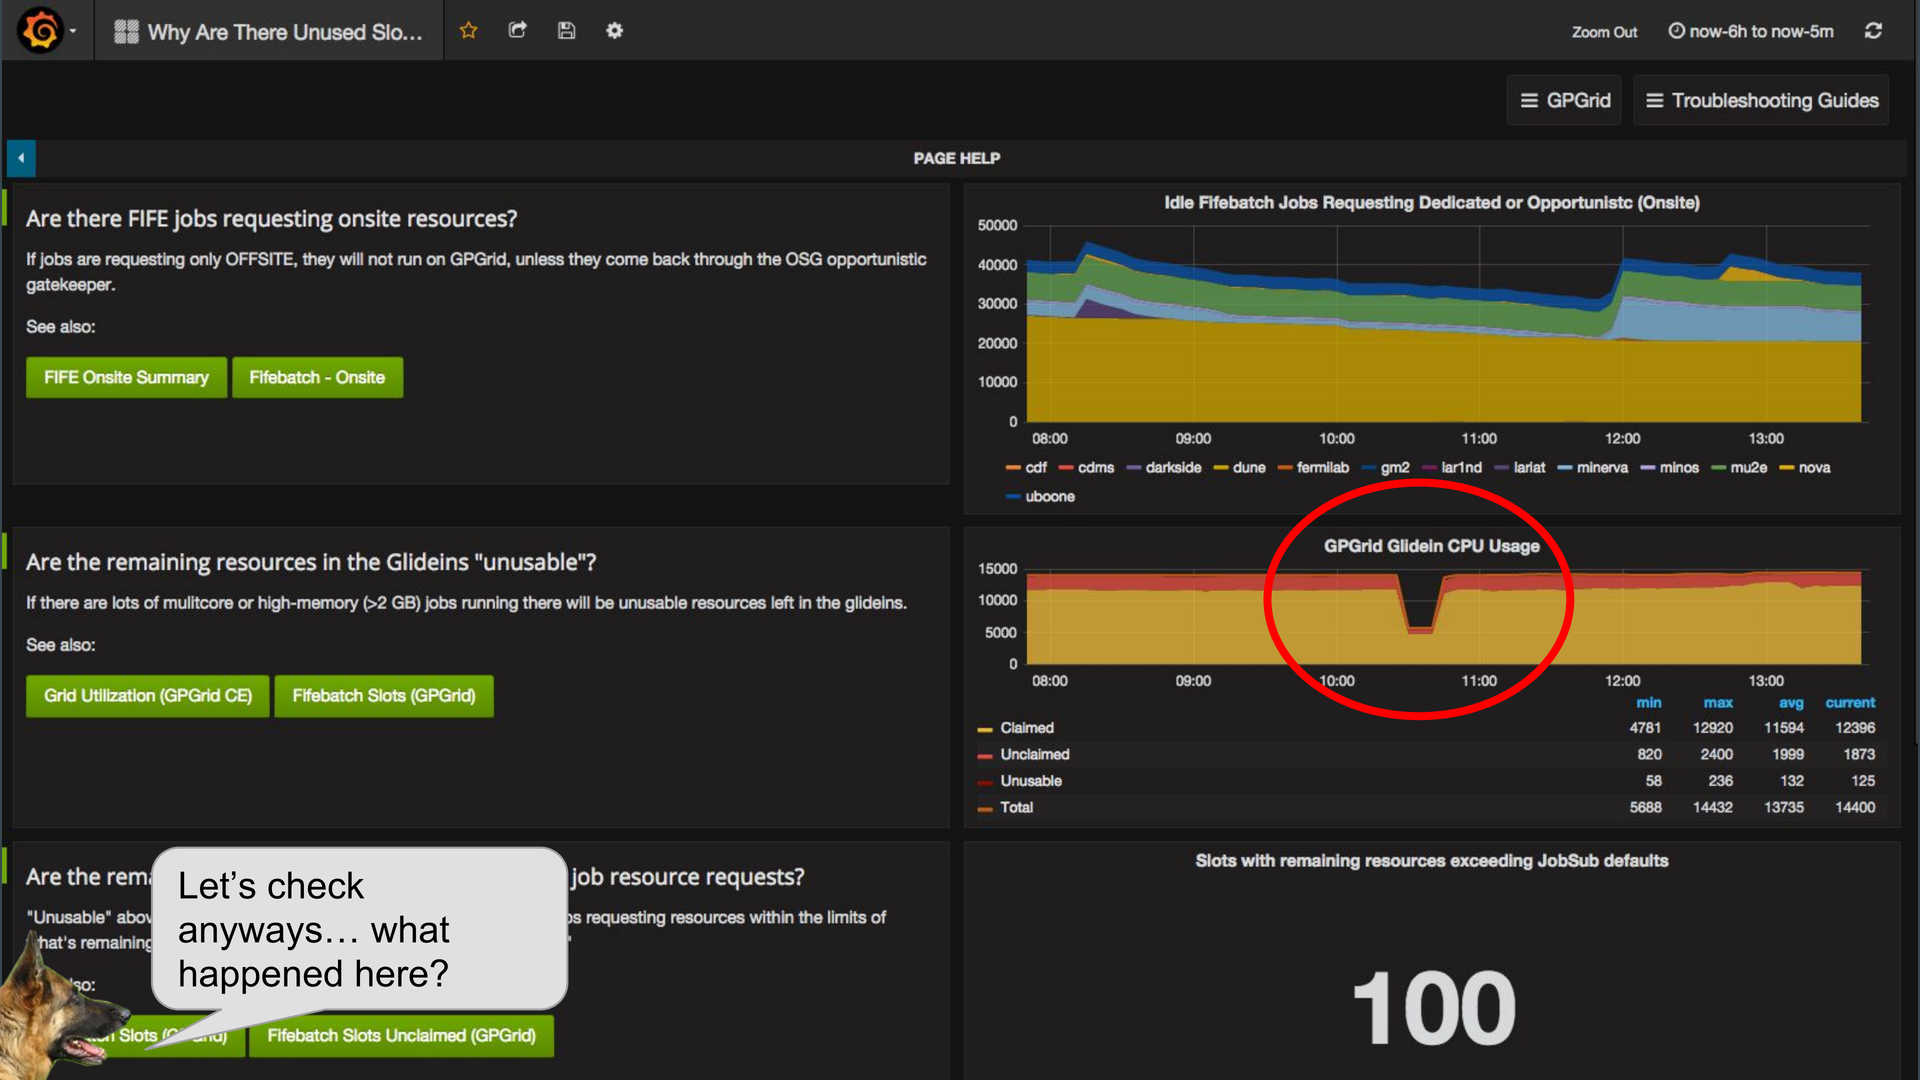

Are there FIFE jobs requesting onsite resources?

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)

	min	max	avg	current
Claimed	4781	12920	11594	12396
Unclaimed	820	2400	1999	1873
Unusable	58	236	132	125
Total	5688	14432	13735	14400

Are the rema job resource requests?

"Unusable" abov s requesting resources within the limits of hat's remaining

lso:

Slots (G rid) | Fifebatch Slots Unclaimed (GPGrid)

Let’s check anyways… what happened here?

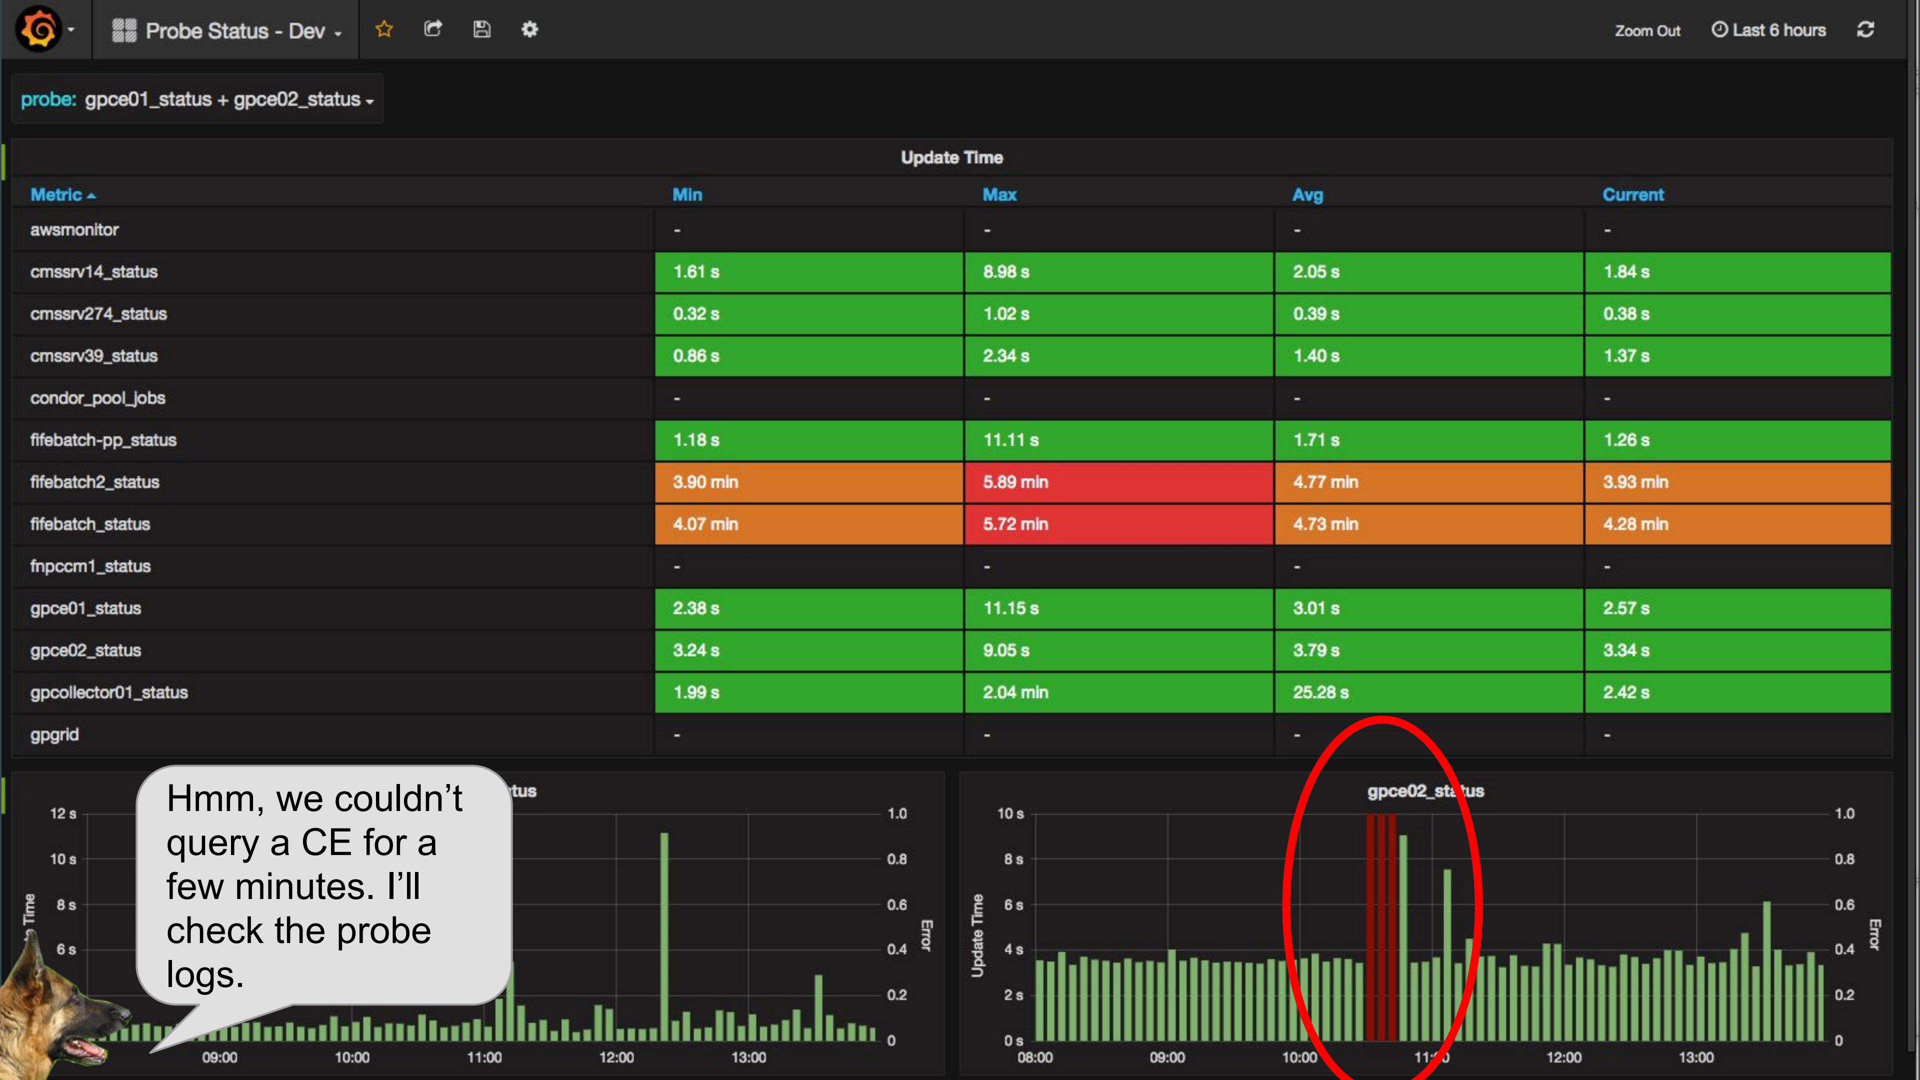

Update Time

Metric ▲	Min	Max	Avg	Current
awsmonitor	-	-	-	-
cmssrv14_status	1.61 s	8.98 s	2.05 s	1.84 s
cmssrv274_status	0.32 s	1.02 s	0.39 s	0.38 s
cmssrv39_status	0.86 s	2.34 s	1.40 s	1.37 s
condor_pool_jobs	-	-	-	-
fifebatch-pp_status	1.18 s	11.11 s	1.71 s	1.26 s
fifebatch2_status	3.90 min	5.89 min	4.77 min	3.93 min
fifebatch_status	4.07 min	5.72 min	4.73 min	4.28 min
fnpccm1_status	-	-	-	-
gpce01_status	2.38 s	11.15 s	3.01 s	2.57 s
gpce02_status	3.24 s	9.05 s	3.79 s	3.34 s
gpcollector01_status	1.99 s	2.04 min	25.28 s	2.42 s
gpgrid	-	-	-	-

Hmm, we couldn’t query a CE for a few minutes. I’ll check the probe logs.

Case Study: Stakeholder

“Is my experiment getting the resources it needs and using them effectively?”

Hazel

Photo: Claudio Gennari (CC-BY-2.0)

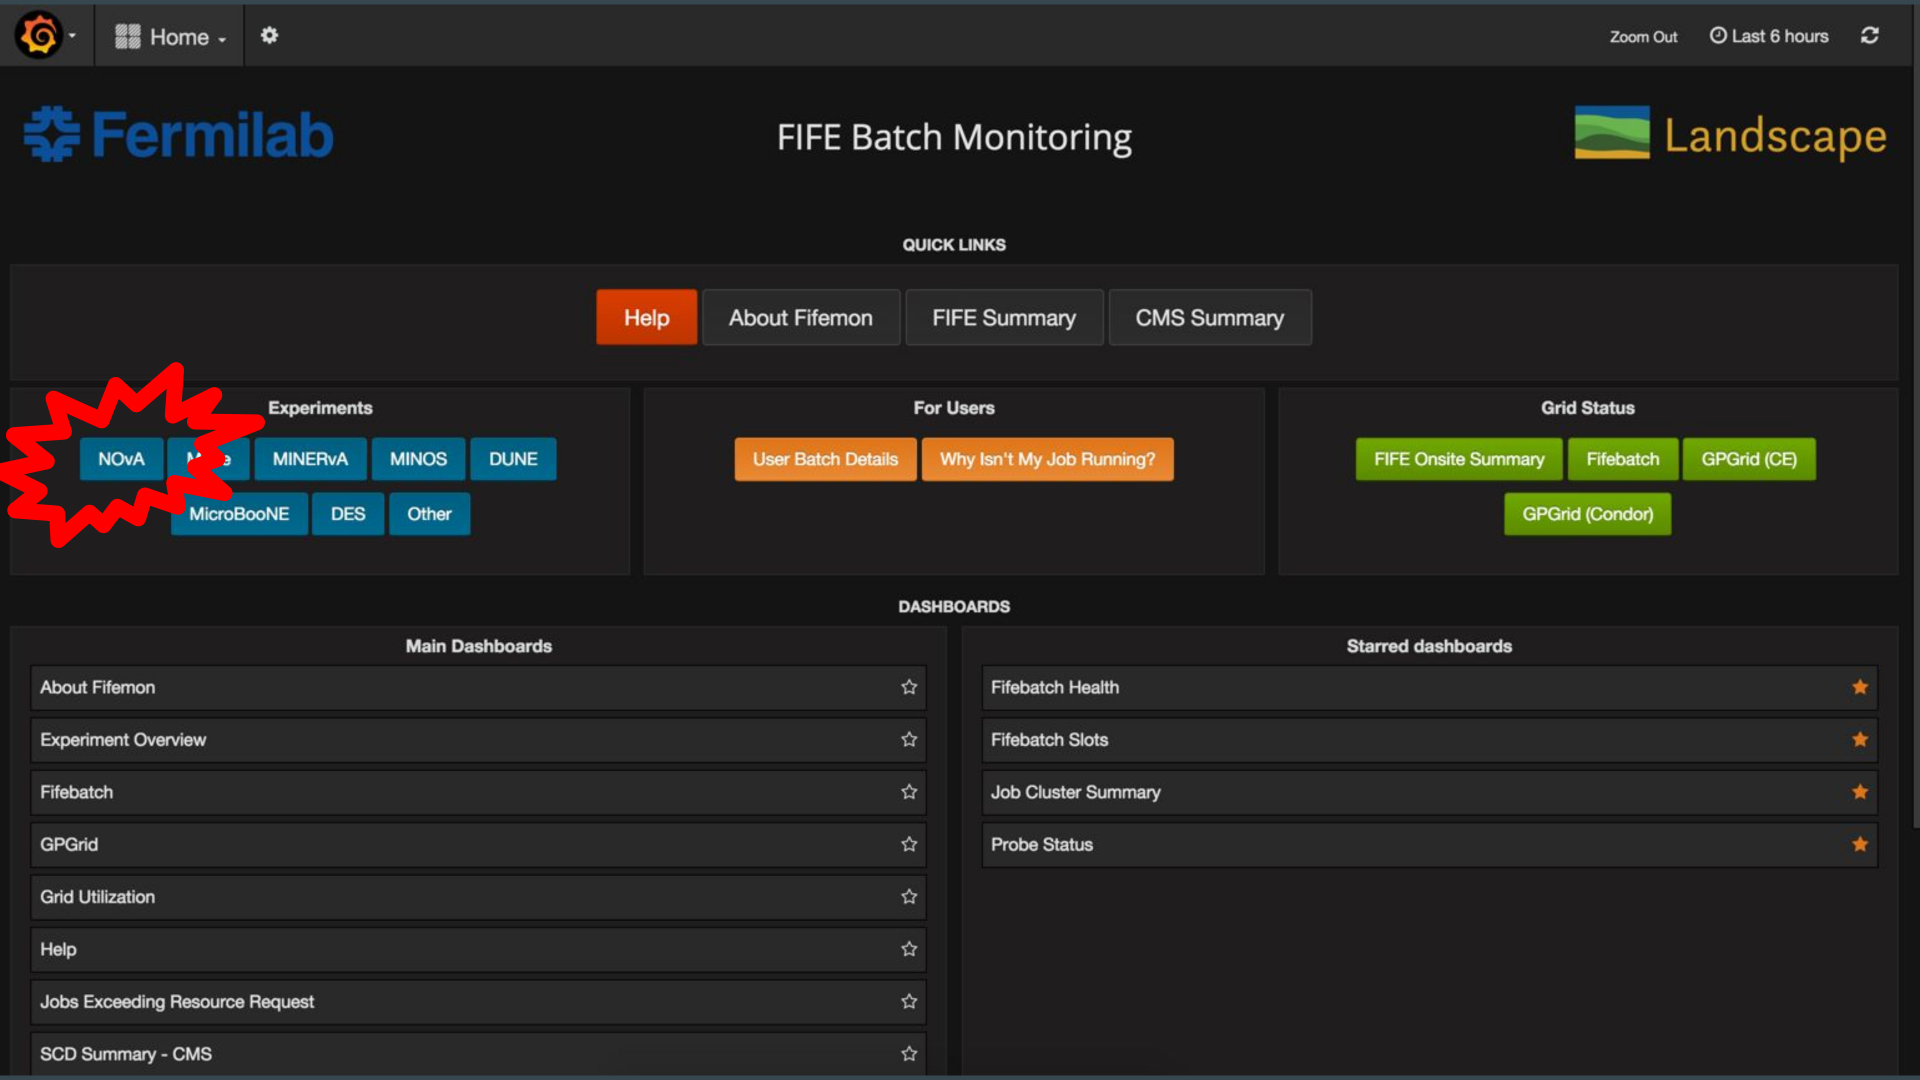

Home | Zoom Out | Last 6 hours

Fermilab

FIFE Batch Monitoring

Landscape

QUICK LINKS

Help | About Fifemon | FIFE Summary | CMS Summary

Experiments

NOvA | Mu2e | MINERvA | MINOS | DUNE

MicroBooNE | DES | Other

For Users

User Batch Details | Why Isn't My Job Running?

Grid Status

FIFE Onsite Summary | Fifebatch | GPGrid (CE)

GPGrid (Condor)

DASHBOARDS

Main Dashboards

- About Fifemon ☆
- Experiment Overview ☆
- Fifebatch ☆
- GPGrid ☆
- Grid Utilization ☆
- Help ☆
- Jobs Exceeding Resource Request ☆
- SCD Summary - CMS ☆

Starred dashboards

- Fifebatch Health ★
- Fifebatch Slots ★
- Job Cluster Summary ★
- Probe Status ★

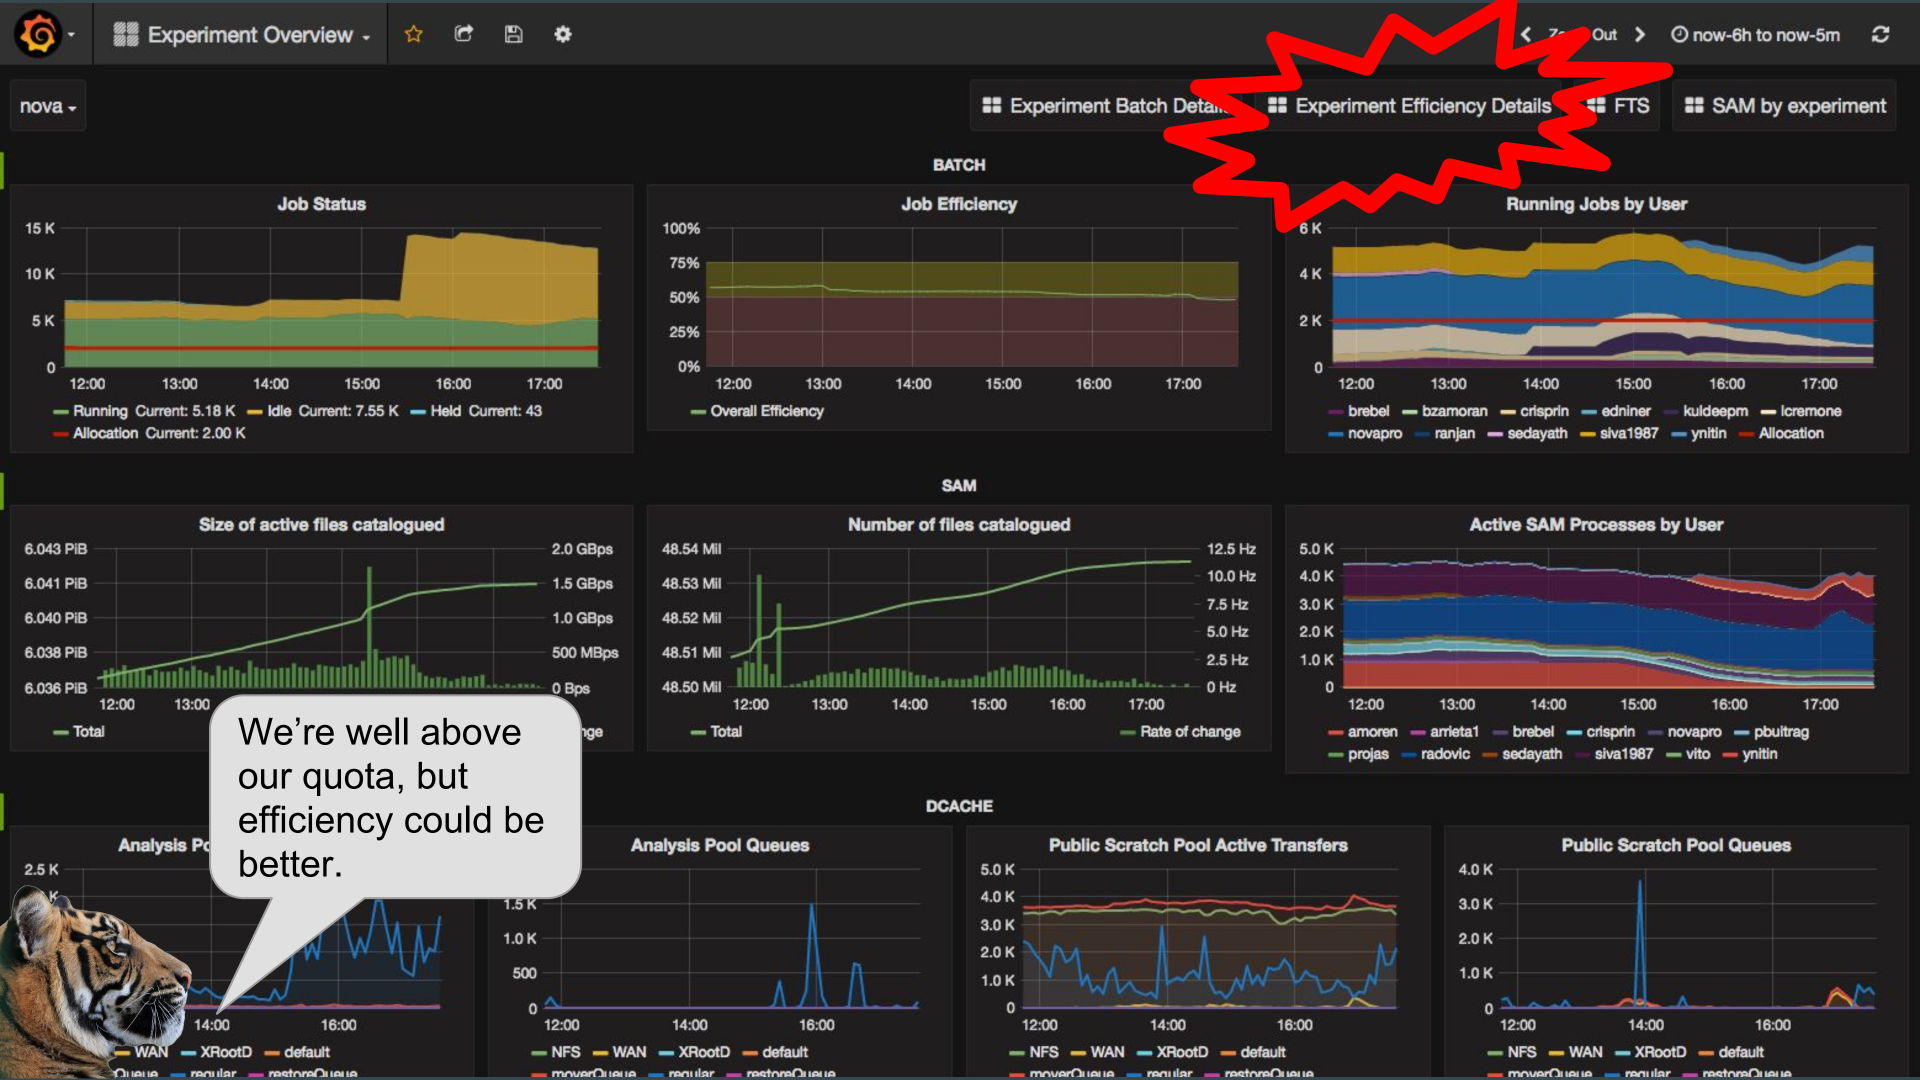

We’re well above our quota, but efficiency could be better.

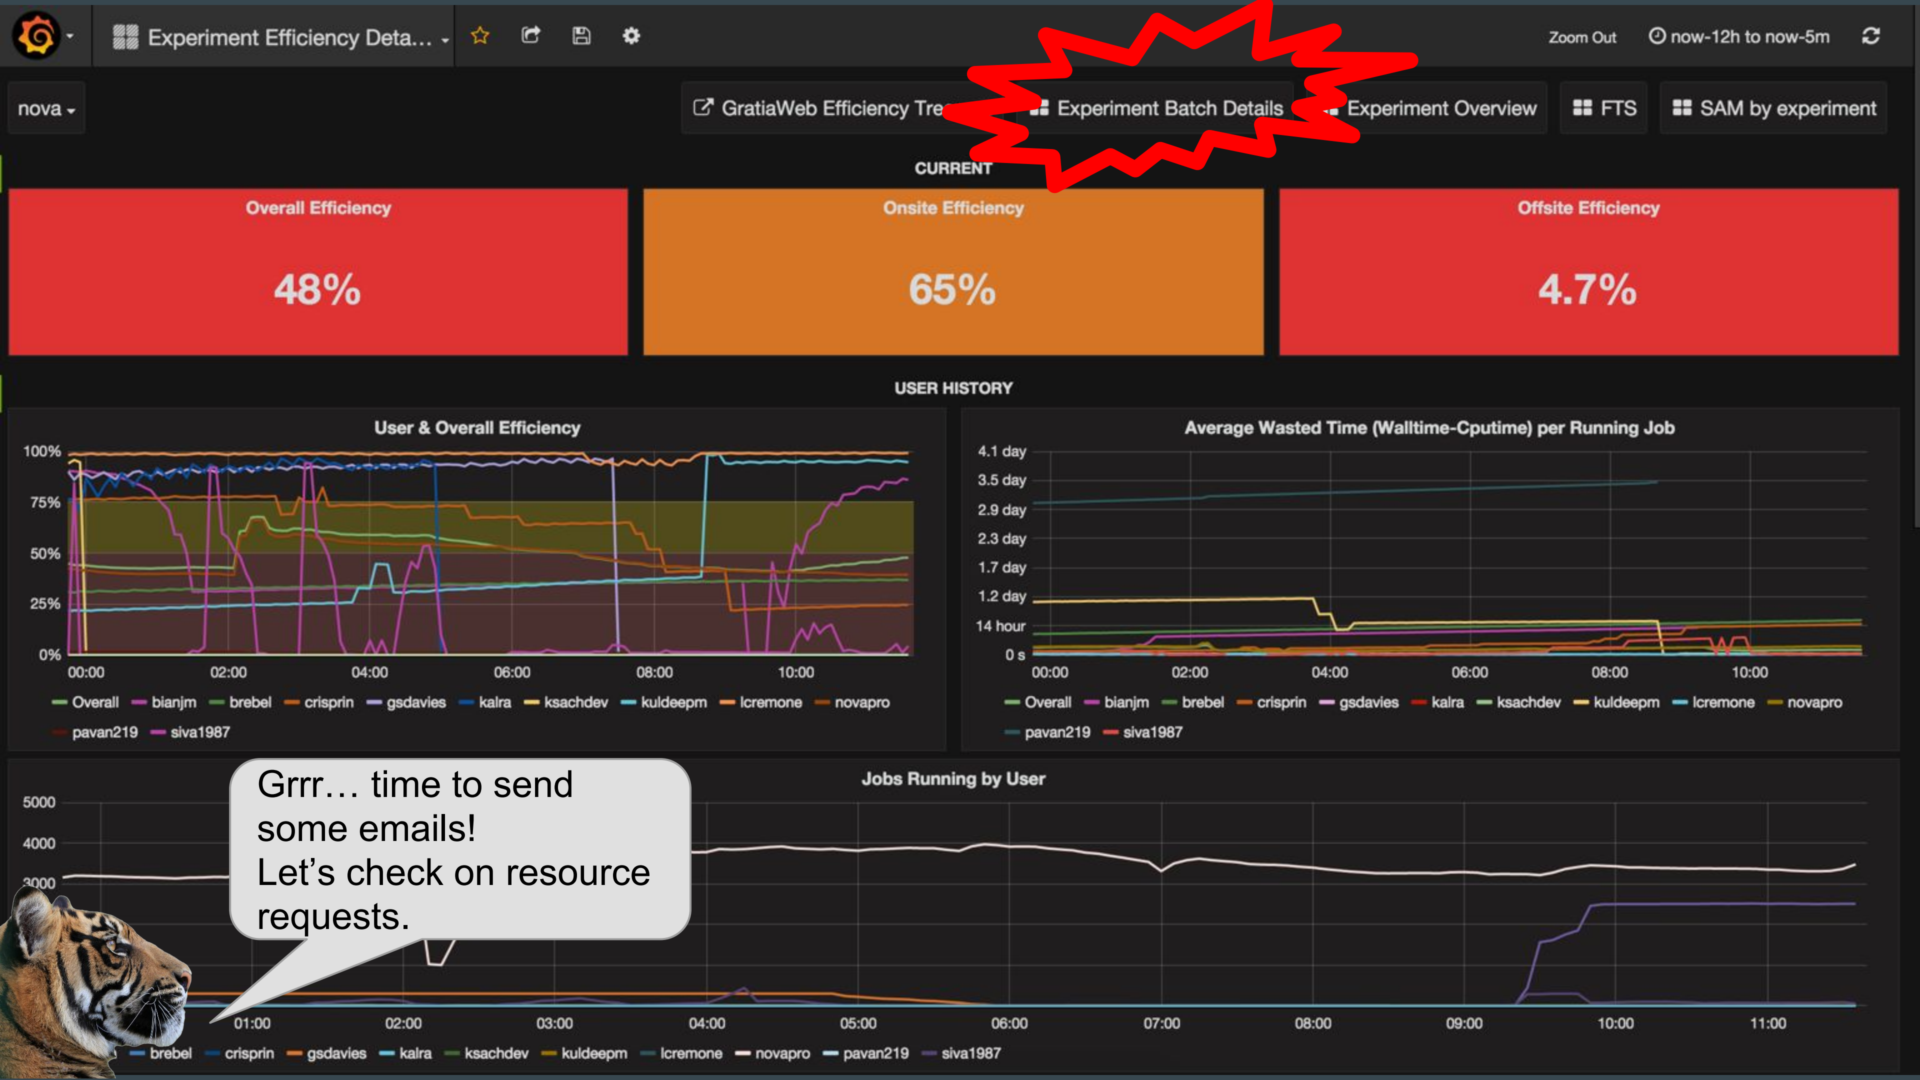

CURRENT

Overall Efficiency	Onsite Efficiency	Offsite Efficiency
48%	65%	4.7%

USER HISTORY

Grrr… time to send some emails!
Let's check on resource requests.

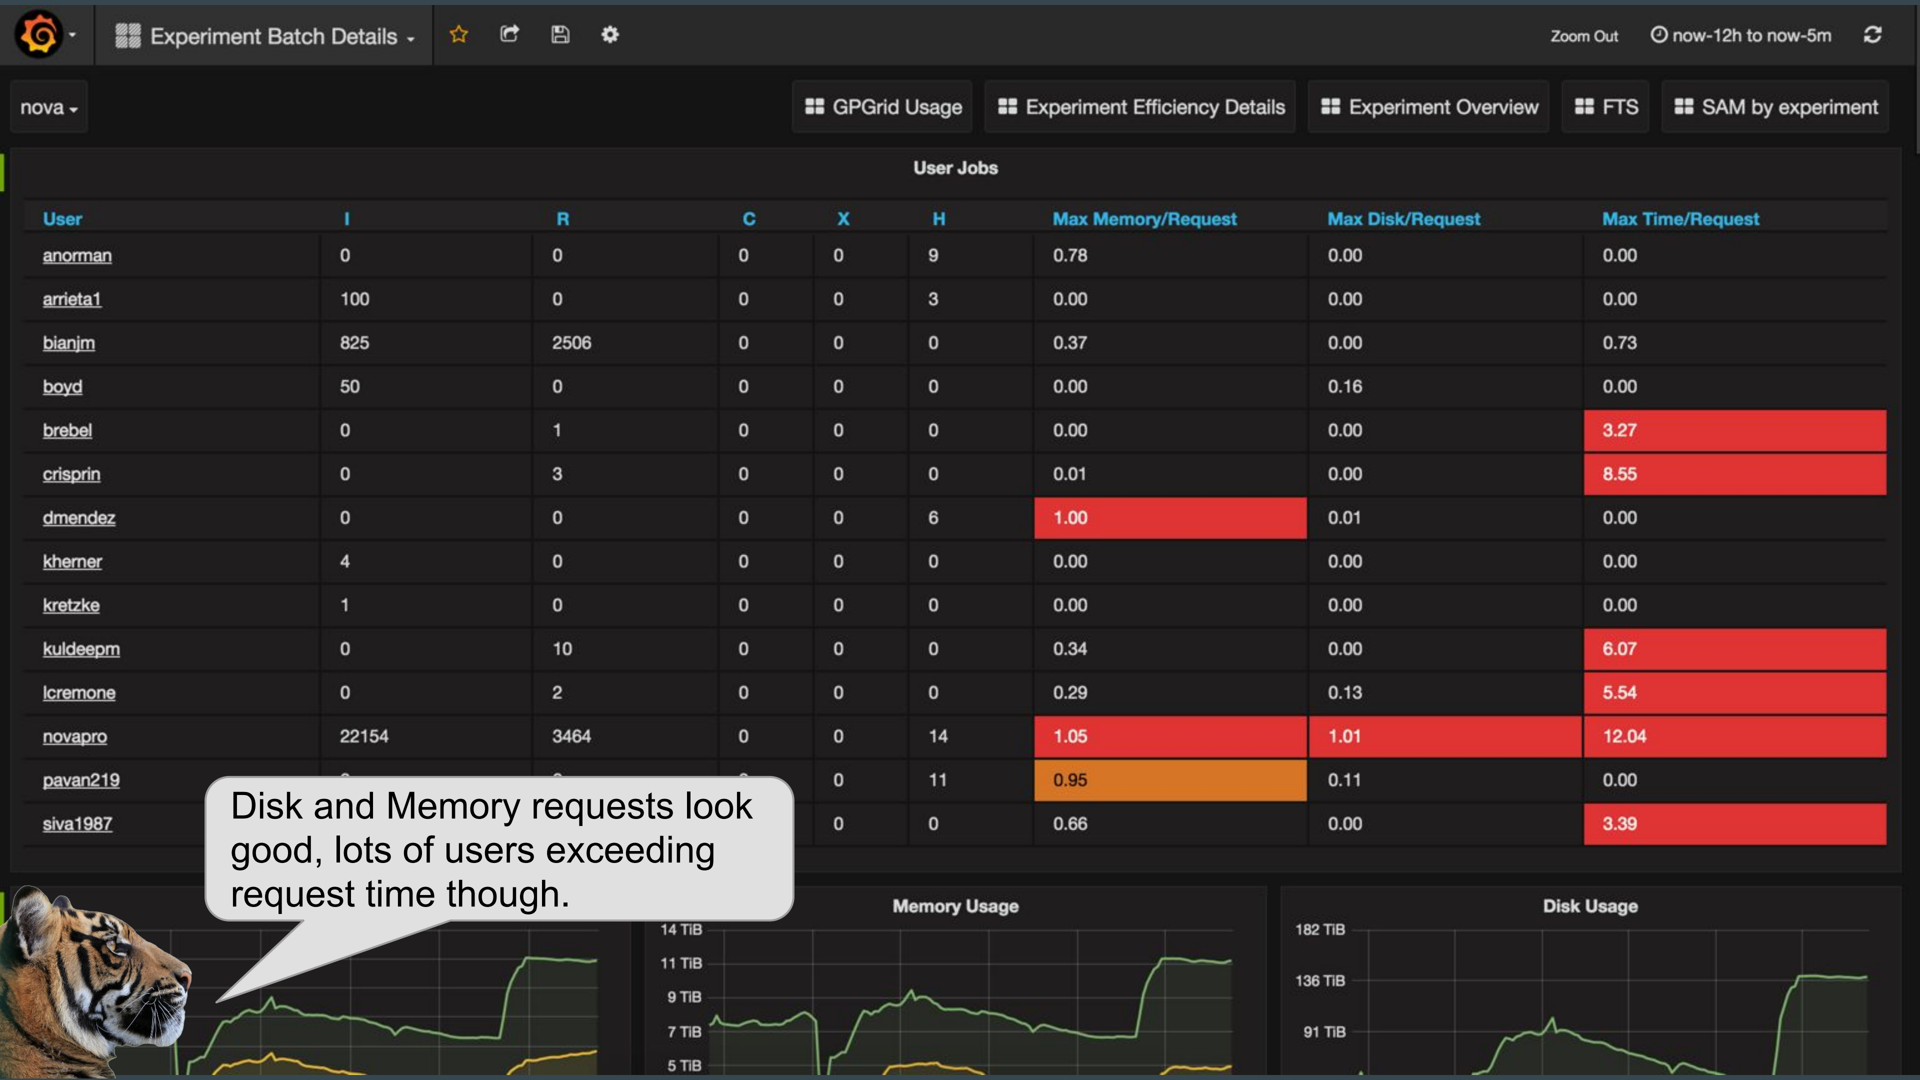

User Jobs

User	I	R	C	X	H	Max Memory/Request	Max Disk/Request	Max Time/Request
anorman	0	0	0	0	9	0.78	0.00	0.00
arrieta1	100	0	0	0	3	0.00	0.00	0.00
bianjm	825	2506	0	0	0	0.37	0.00	0.73
boyd	50	0	0	0	0	0.00	0.16	0.00
brebel	0	1	0	0	0	0.00	0.00	3.27
crisprin	0	3	0	0	0	0.01	0.00	8.55
dmendez	0	0	0	0	6	1.00	0.01	0.00
kherner	4	0	0	0	0	0.00	0.00	0.00
kretzke	1	0	0	0	0	0.00	0.00	0.00
kuldeepm	0	10	0	0	0	0.34	0.00	6.07
lcremone	0	2	0	0	0	0.29	0.13	5.54
novapro	22154	3464	0	0	14	1.05	1.01	12.04
pavan219	[illegible]	[illegible]	[illegible]	0	11	0.95	0.11	0.00
siva1987	[illegible]	[illegible]	[illegible]	0	0	0.66	0.00	3.39

Disk and Memory requests look good, lots of users exceeding request time though.

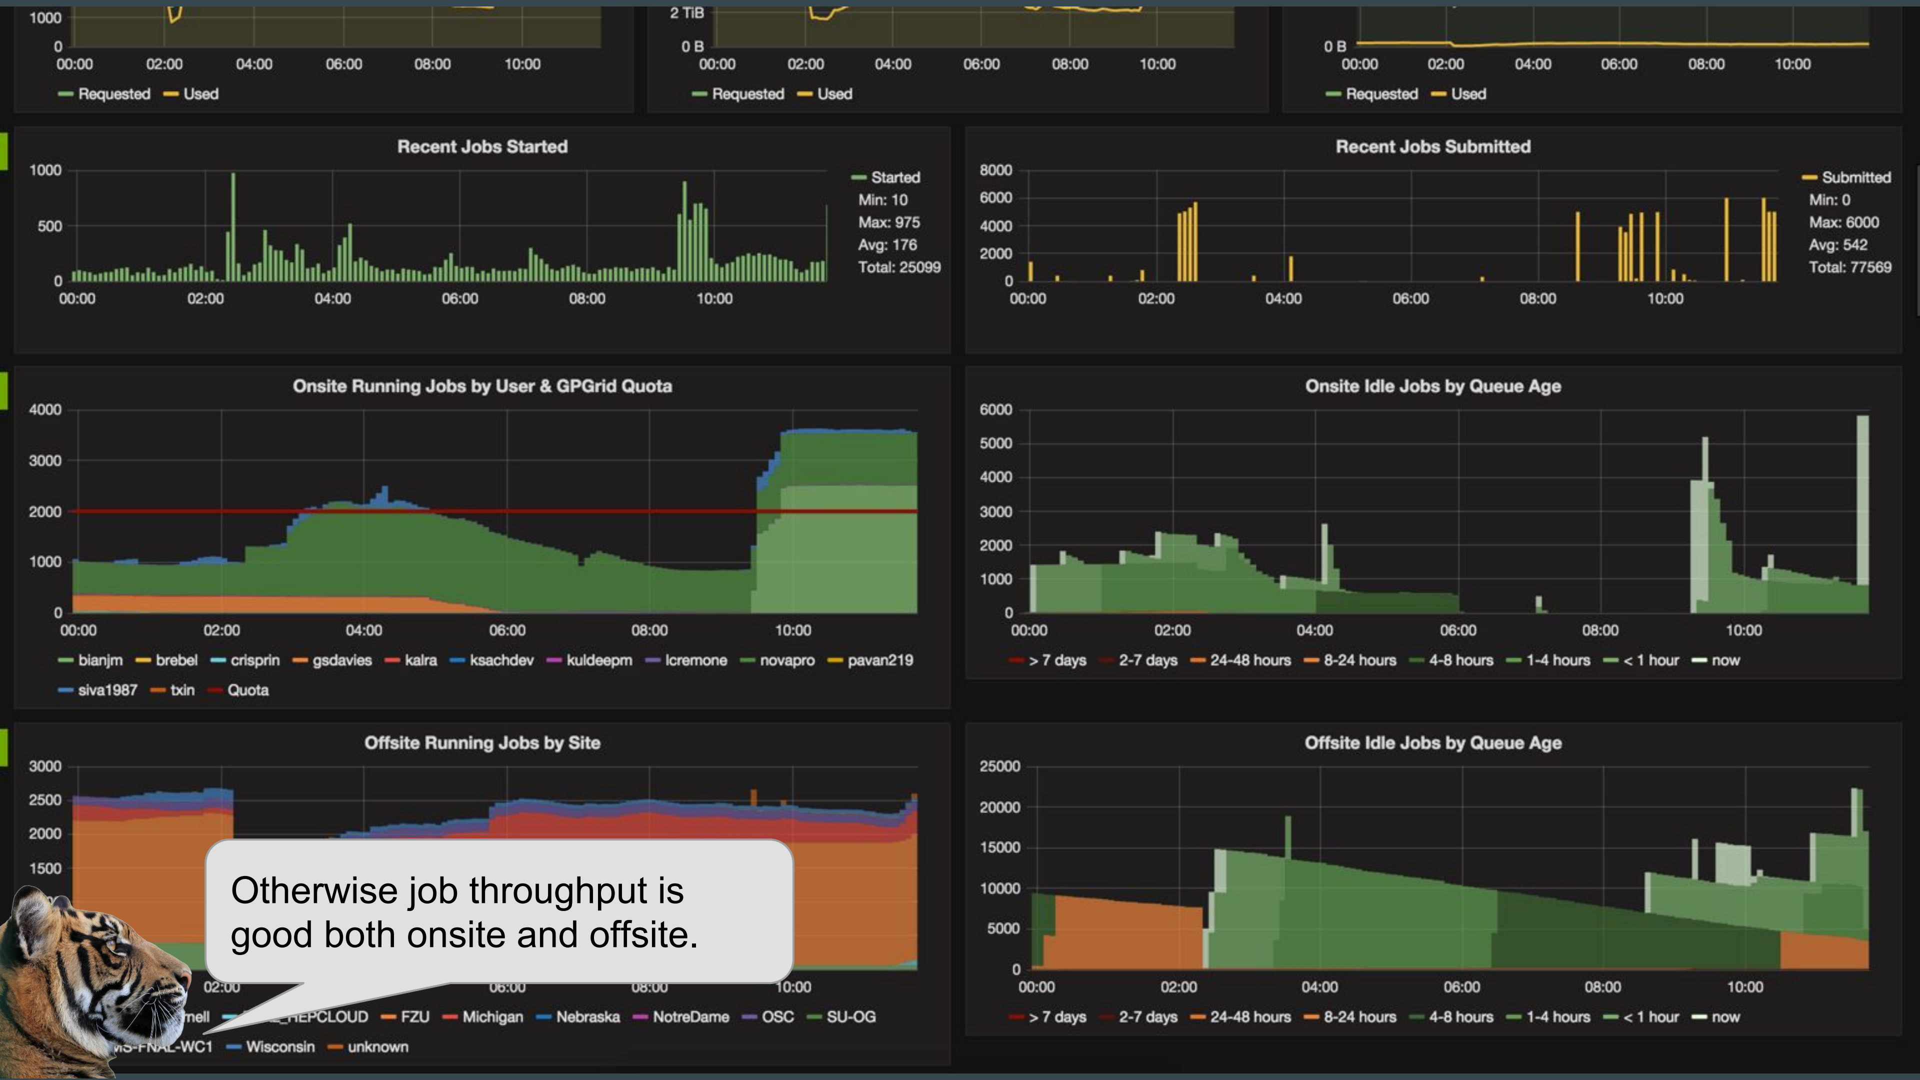

Otherwise job throughput is good both onsite and offsite.

Case Study: User

“What’s the status of my jobs?”

Cocoa

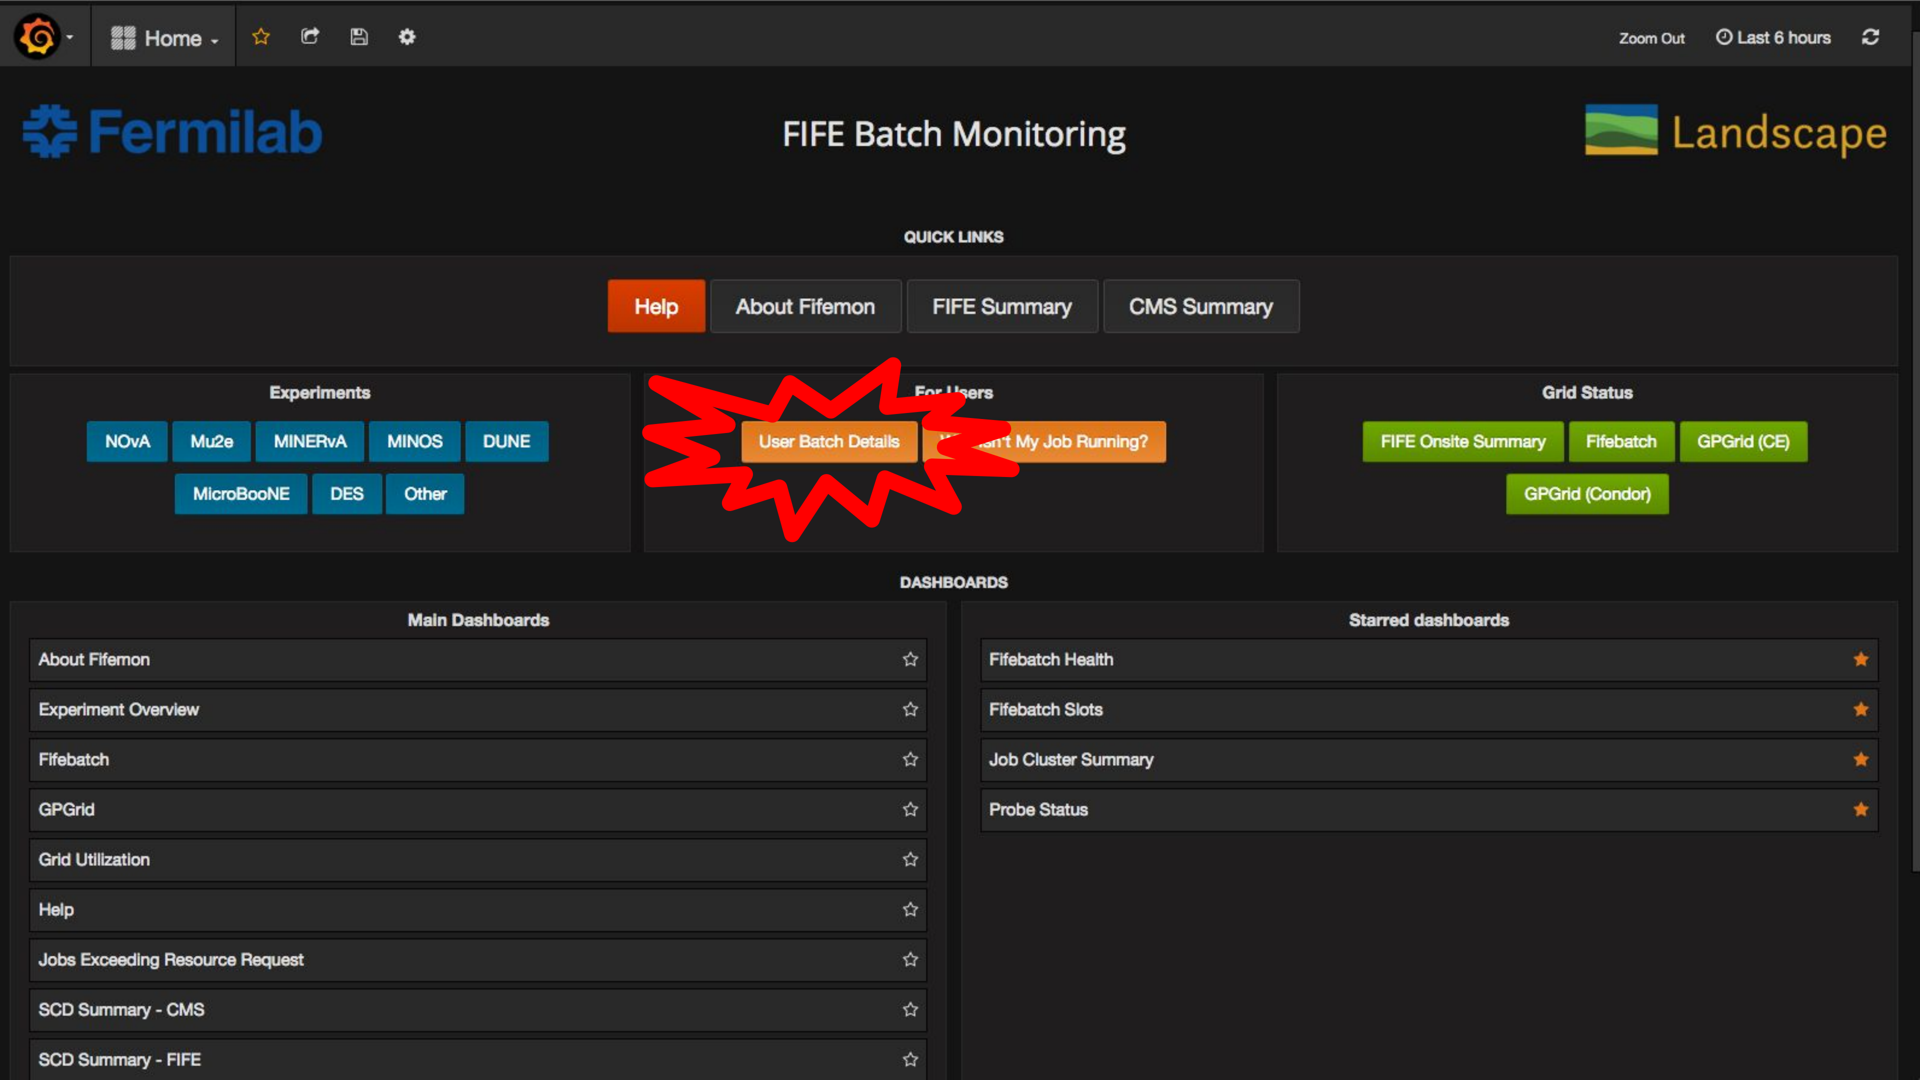
Home
Zoom Out
Last 6 hours
Fermilab
FIFE Batch Monitoring
Landscape
QUICK LINKS
Help
About Fifemon
FIFE Summary
CMS Summary
Experiments
NOvA
Mu2e
MINERvA
MINOS
DUNE
MicroBooNE
DES
Other
For Users
User Batch Details
Why Isn't My Job Running?
Grid Status
FIFE Onsite Summary
Fifebatch
GPGrid (CE)
GPGrid (Condor)
DASHBOARDS
Main Dashboards
About Fifemon
Experiment Overview
Fifebatch
GPGrid
Grid Utilization
Help
Jobs Exceeding Resource Request
SCD Summary - CMS
SCD Summary - FIFE
Starred dashboards
Fifebatch Health
Fifebatch Slots
Job Cluster Summary
Probe Status

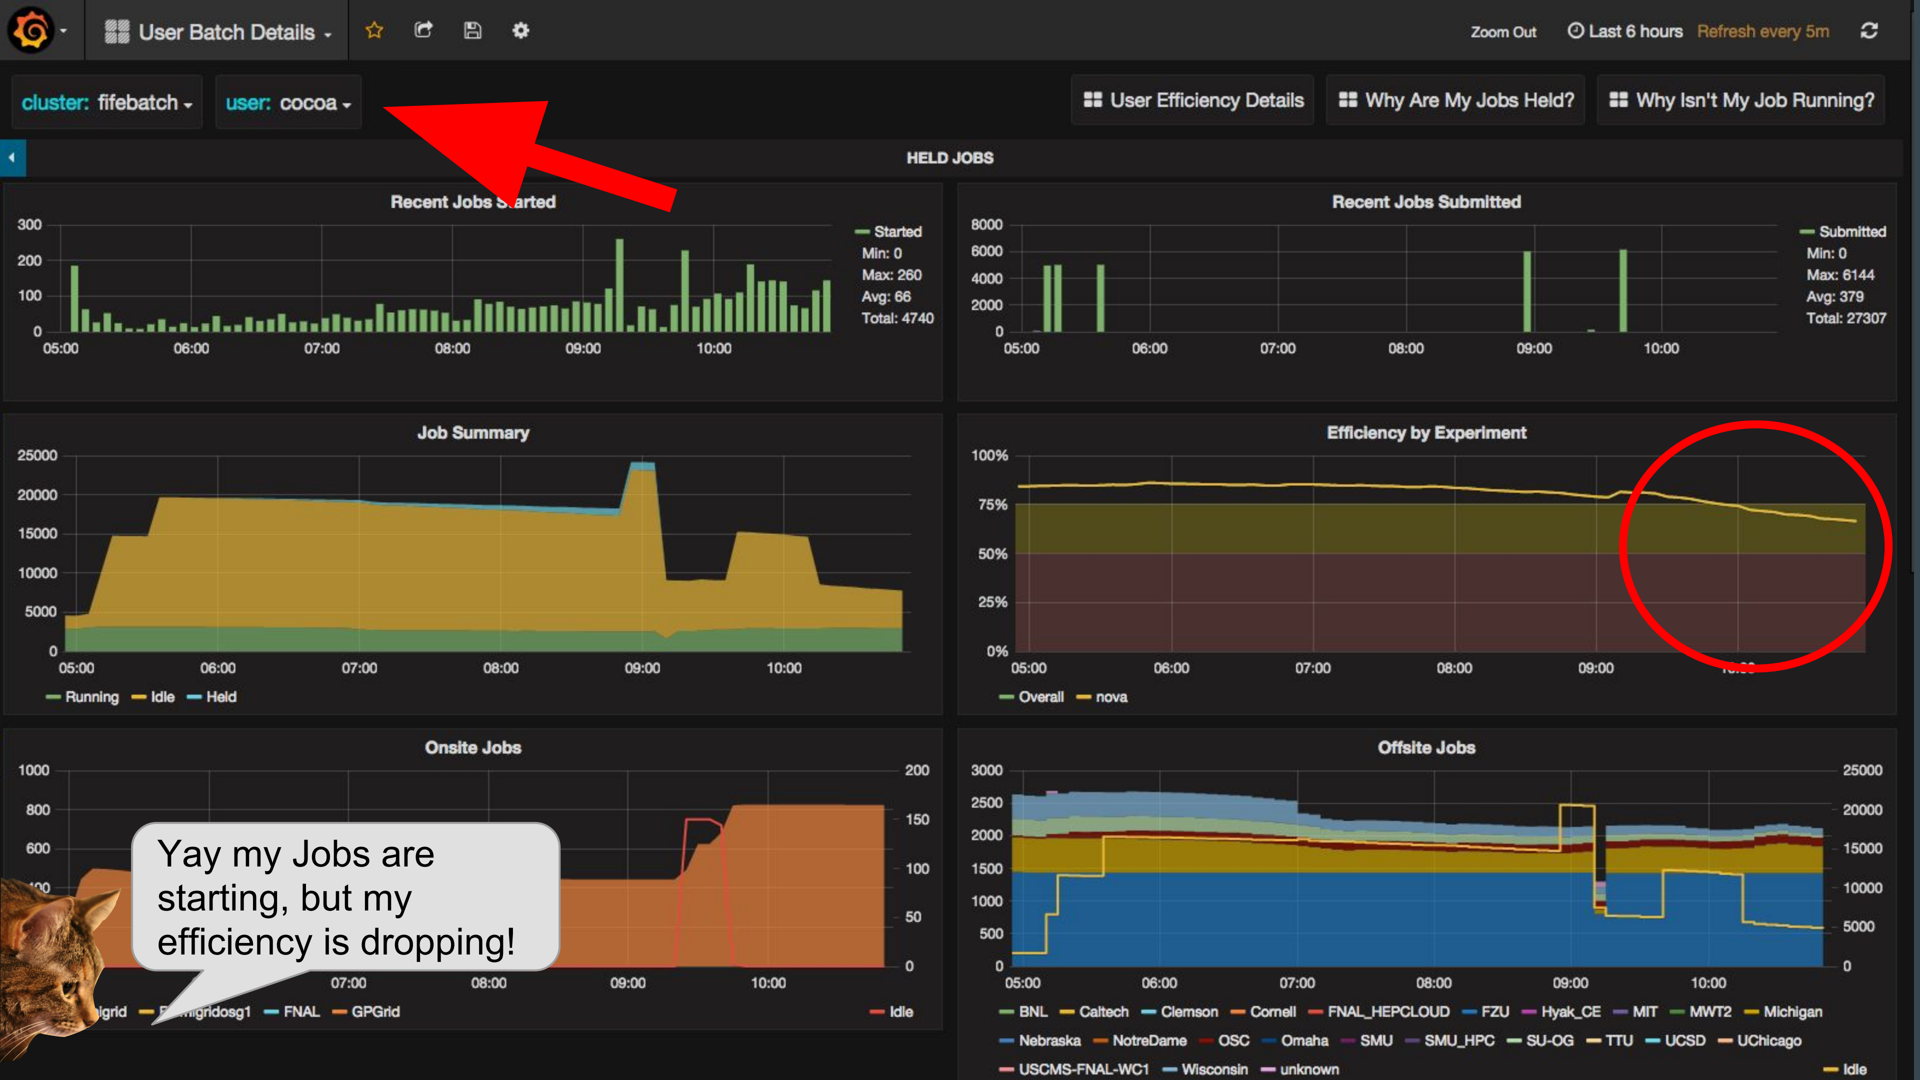

cluster: fifebatch

user: cocoa

Yay my Jobs are starting, but my efficiency is dropping!

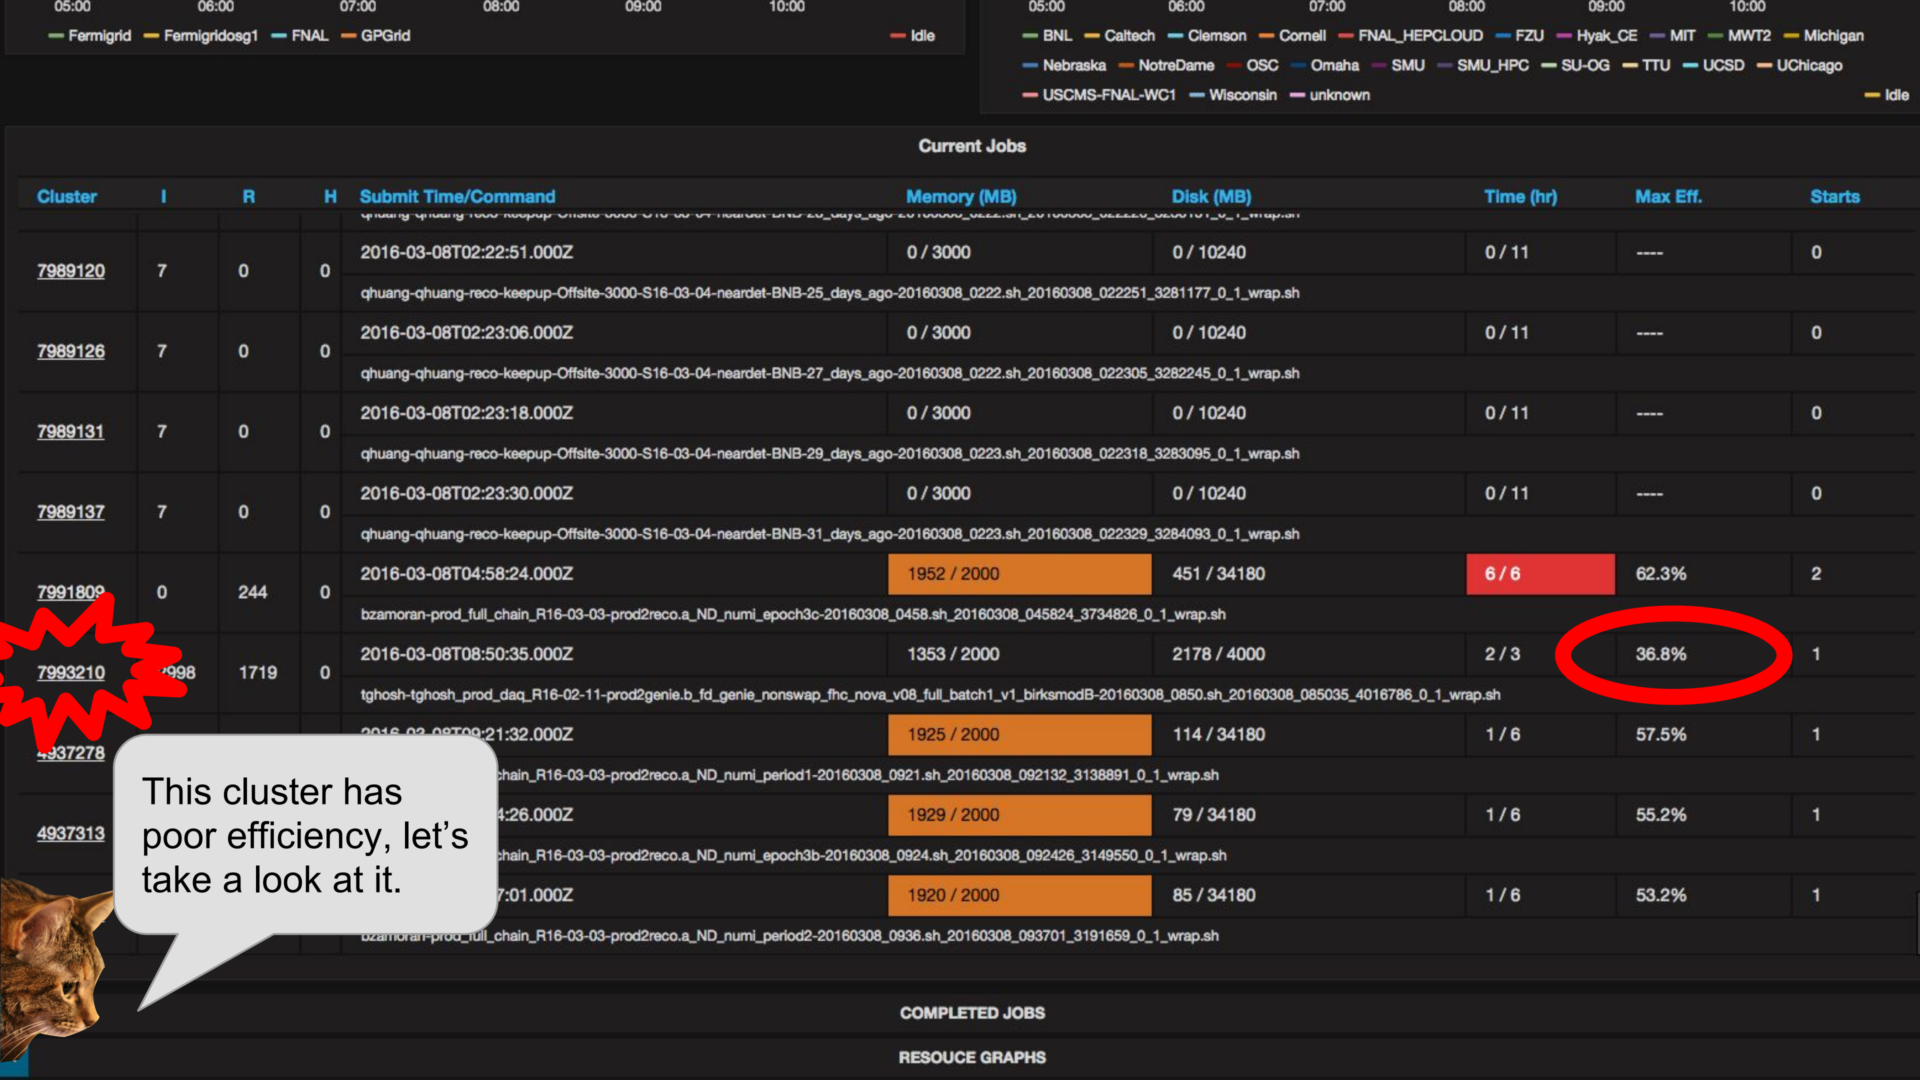

Current Jobs

Cluster	I	R	H	Submit Time/Command	Memory (MB)	Disk (MB)	Time (hr)	Max Eff.	Starts
7989120	7	0	0	2016-03-08T02:22:51.000Z qhuang-qhuang-reco-keepup-Offsite-3000-S16-03-04-neardet-BNB-25_days_ago-20160308_0222.sh_20160308_022251_3281177_0_1_wrap.sh	0 / 3000	0 / 10240	0 / 11	----	0
7989126	7	0	0	2016-03-08T02:23:06.000Z qhuang-qhuang-reco-keepup-Offsite-3000-S16-03-04-neardet-BNB-27_days_ago-20160308_0222.sh_20160308_022305_3282245_0_1_wrap.sh	0 / 3000	0 / 10240	0 / 11	----	0
7989131	7	0	0	2016-03-08T02:23:18.000Z qhuang-qhuang-reco-keepup-Offsite-3000-S16-03-04-neardet-BNB-29_days_ago-20160308_0223.sh_20160308_022318_3283095_0_1_wrap.sh	0 / 3000	0 / 10240	0 / 11	----	0
7989137	7	0	0	2016-03-08T02:23:30.000Z qhuang-qhuang-reco-keepup-Offsite-3000-S16-03-04-neardet-BNB-31_days_ago-20160308_0223.sh_20160308_022329_3284093_0_1_wrap.sh	0 / 3000	0 / 10240	0 / 11	----	0
7991809	0	244	0	2016-03-08T04:58:24.000Z bzamoran-prod_full_chain_R16-03-03-prod2reco.a_ND_numi_epoch3c-20160308_0458.sh_20160308_045824_3734826_0_1_wrap.sh	1952 / 2000	451 / 34180	6 / 6	62.3%	2
7993210	[illegible]998	1719	0	2016-03-08T08:50:35.000Z tghosh-tghosh_prod_daq_R16-02-11-prod2genie.b_fd_genie_nonswap_fhc_nova_v08_full_batch1_v1_birksmodB-20160308_0850.sh_20160308_085035_4016786_0_1_wrap.sh	1353 / 2000	2178 / 4000	2 / 3	36.8%	1
4937278	[illegible]	[illegible]	[illegible]	2016-03-08T09:21:32.000Z chain_R16-03-03-prod2reco.a_ND_numi_period1-20160308_0921.sh_20160308_092132_3138891_0_1_wrap.sh	1925 / 2000	114 / 34180	1 / 6	57.5%	1
4937313	[illegible]	[illegible]	[illegible]	4:26.000Z chain_R16-03-03-prod2reco.a_ND_numi_epoch3b-20160308_0924.sh_20160308_092426_3149550_0_1_wrap.sh	1929 / 2000	79 / 34180	1 / 6	55.2%	1
[illegible]	[illegible]	[illegible]	[illegible]	7:01.000Z bzamoran-prod_full_chain_R16-03-03-prod2reco.a_ND_numi_period2-20160308_0936.sh_20160308_093701_3191659_0_1_wrap.sh	1920 / 2000	85 / 34180	1 / 6	53.2%	1

COMPLETED JOBS

RESOUCE GRAPHS

This cluster has poor efficiency, let's take a look at it.

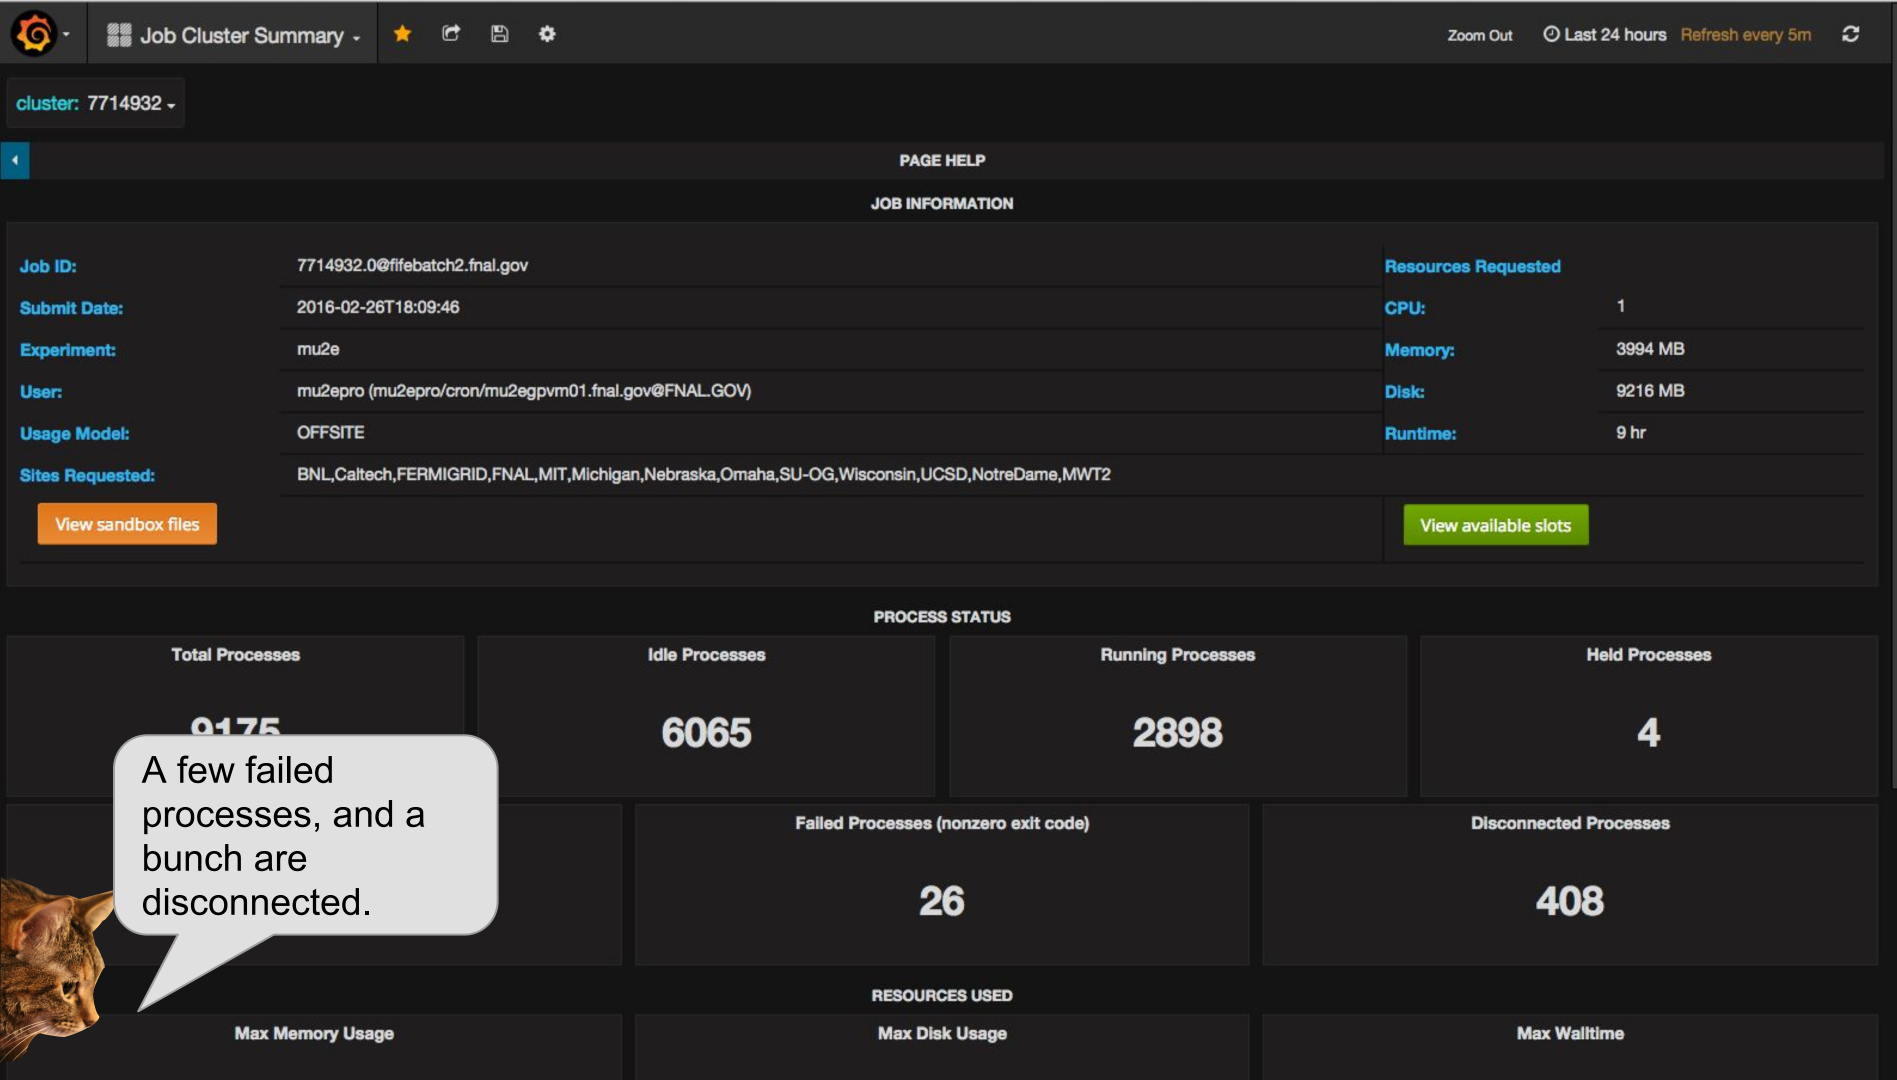

JOB INFORMATION

Job ID:	7714932.0@fifebatch2.fnal.gov
Submit Date:	2016-02-26T18:09:46
Experiment:	mu2e
User:	mu2epro (mu2epro/cron/mu2egpvm01.fnal.gov@FNAL.GOV)
Usage Model:	OFFSITE
Sites Requested:	BNL,Caltech,FERMIGRID,FNAL,MIT,Michigan,Nebraska,Omaha,SU-OG,Wisconsin,UCSD,NotreDame,MWT2

Resources Requested	
CPU:	1
Memory:	3994 MB
Disk:	9216 MB
Runtime:	9 hr

PROCESS STATUS

Total Processes	Idle Processes	Running Processes	Held Processes
9175	6065	2898	4

Failed Processes (nonzero exit code)	Disconnected Processes
26	408

RESOURCES USED

A few failed processes, and a bunch are disconnected.

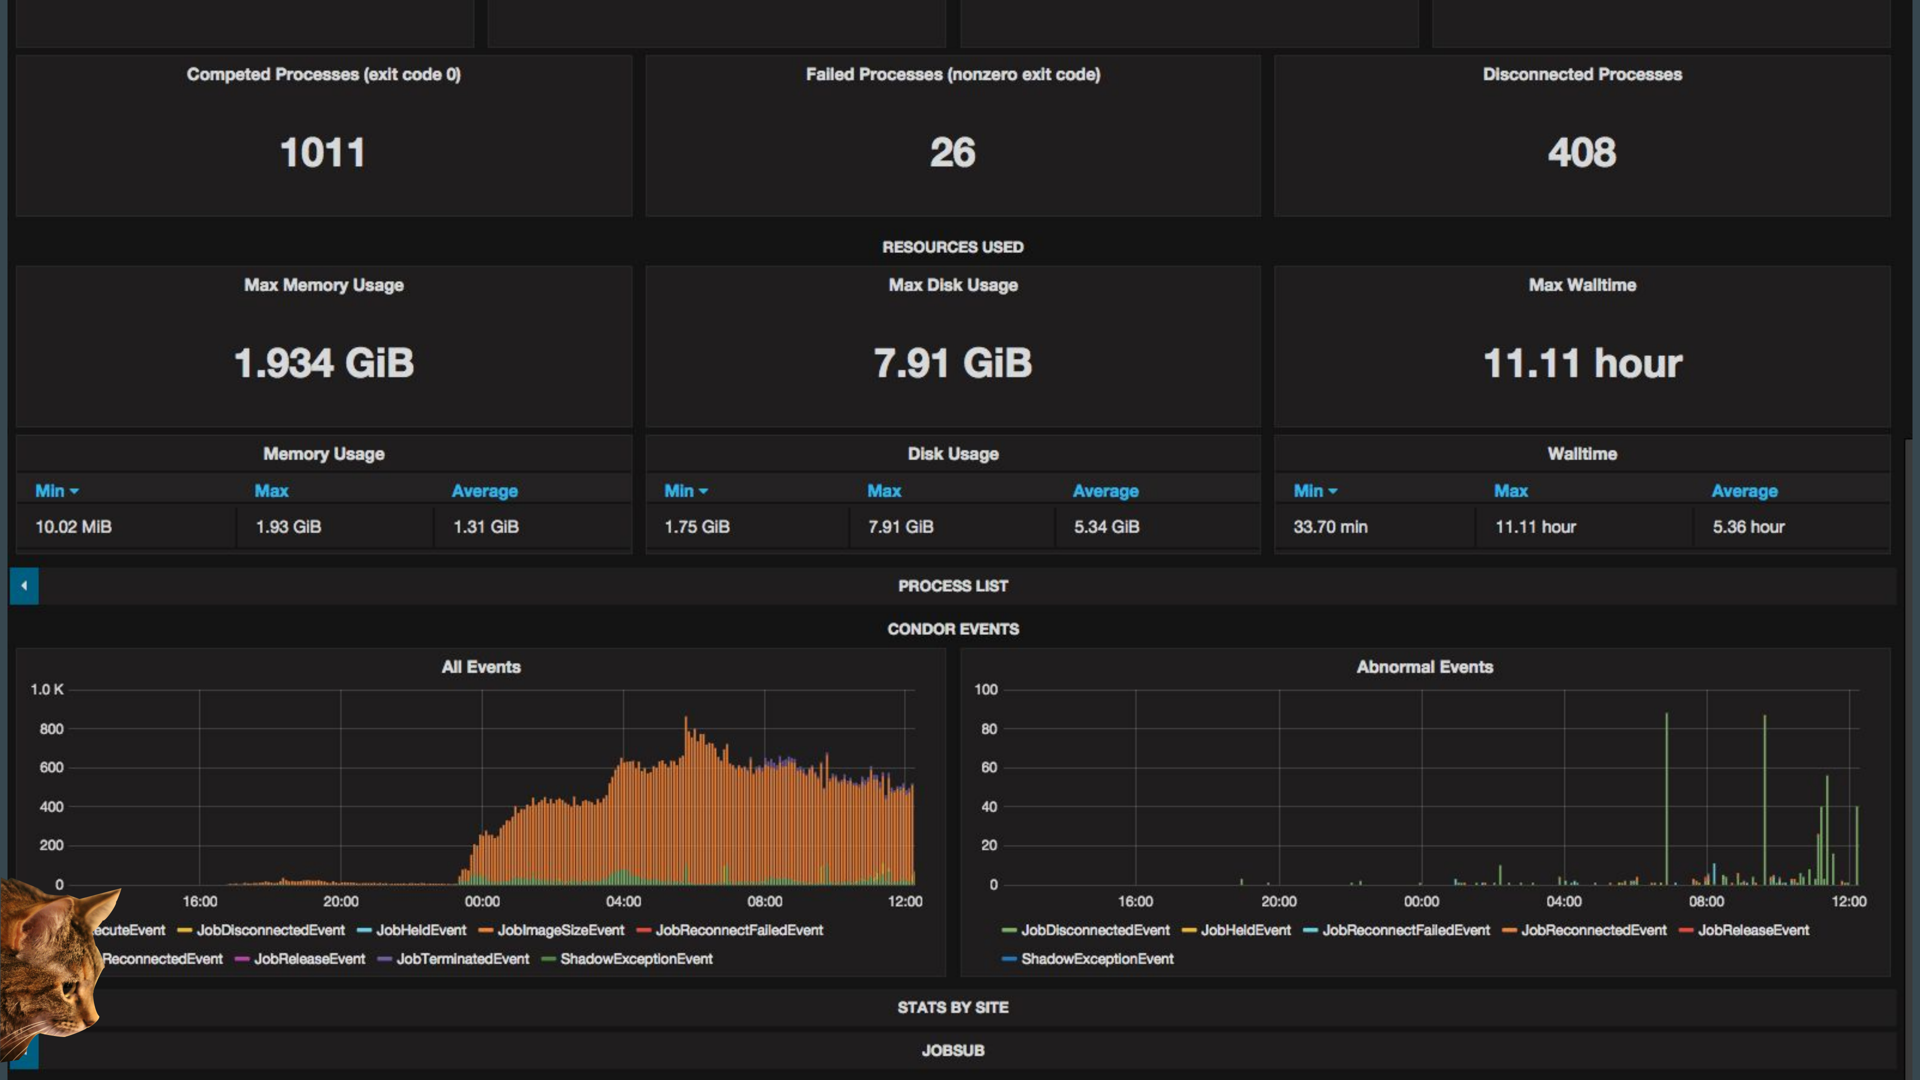

Competed Processes (exit code 0)

1011

Failed Processes (nonzero exit code)

26

Disconnected Processes

408

RESOURCES USED

Max Memory Usage

1.934 GiB

Max Disk Usage

7.91 GiB

Max Walltime

11.11 hour

Memory Usage

Min ▾	Max	Average
10.02 MiB	1.93 GiB	1.31 GiB

Disk Usage

Min ▾	Max	Average
1.75 GiB	7.91 GiB	5.34 GiB

Walltime

Min ▾	Max	Average
33.70 min	11.11 hour	5.36 hour

PROCESS LIST

CONDOR EVENTS

All Events

Abnormal Events

STATS BY SITE

JOBSUB

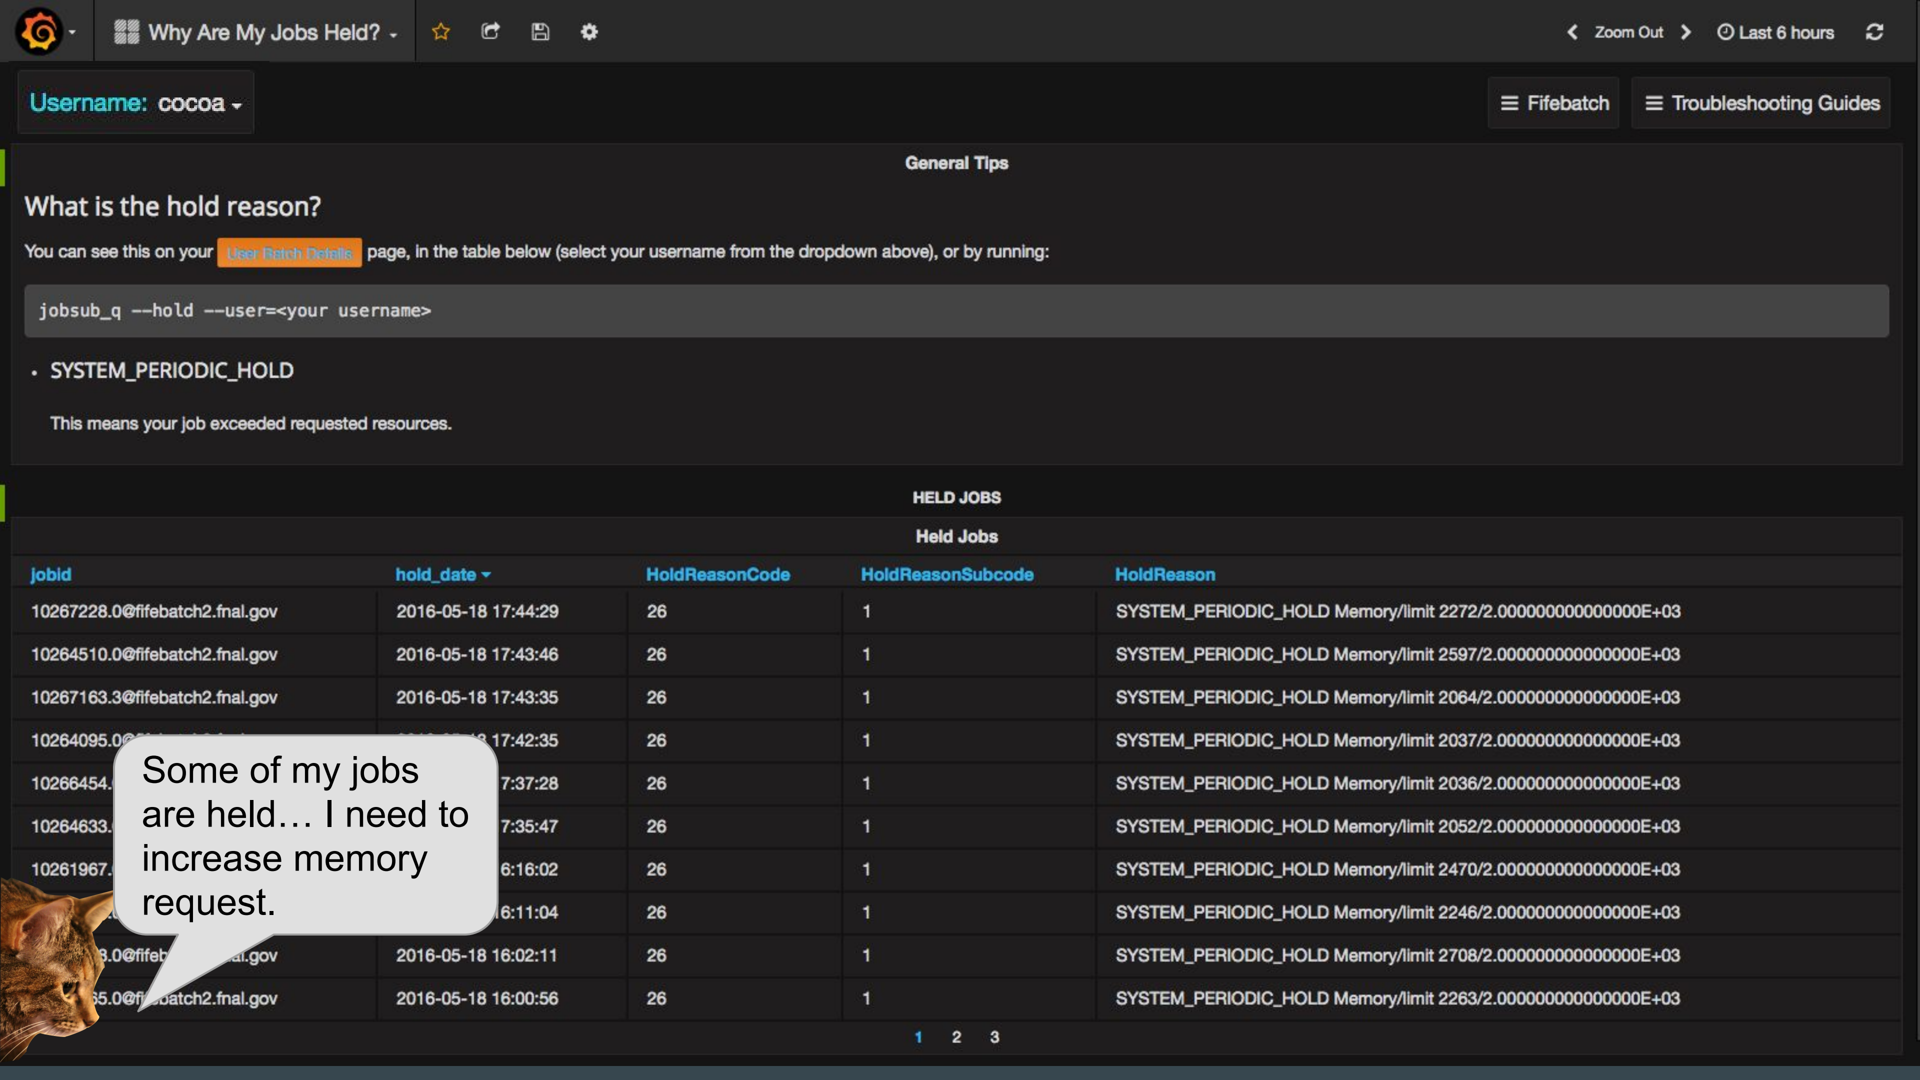

Username: cocoa

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:

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

- SYSTEM_PERIODIC_HOLD

This means your job exceeded requested resources.

HELD JOBS

Held Jobs

jobid	hold_date	HoldReasonCode	HoldReasonSubcode	HoldReason
10267228.0@fifebatch2.fnal.gov	2016-05-18 17:44:29	26	1	SYSTEM_PERIODIC_HOLD Memory/limit 2272/2.000000000000000E+03
10264510.0@fifebatch2.fnal.gov	2016-05-18 17:43:46	26	1	SYSTEM_PERIODIC_HOLD Memory/limit 2597/2.000000000000000E+03
10267163.3@fifebatch2.fnal.gov	2016-05-18 17:43:35	26	1	SYSTEM_PERIODIC_HOLD Memory/limit 2064/2.000000000000000E+03
10264095.0	17:42:35	26	1	SYSTEM_PERIODIC_HOLD Memory/limit 2037/2.000000000000000E+03
10266454.	7:37:28	26	1	SYSTEM_PERIODIC_HOLD Memory/limit 2036/2.000000000000000E+03
10264633.	7:35:47	26	1	SYSTEM_PERIODIC_HOLD Memory/limit 2052/2.000000000000000E+03
10261967.	6:16:02	26	1	SYSTEM_PERIODIC_HOLD Memory/limit 2470/2.000000000000000E+03
[illegible]	6:11:04	26	1	SYSTEM_PERIODIC_HOLD Memory/limit 2246/2.000000000000000E+03
8.0@fifeb al.gov	2016-05-18 16:02:11	26	1	SYSTEM_PERIODIC_HOLD Memory/limit 2708/2.000000000000000E+03
65.0@fi batch2.fnal.gov	2016-05-18 16:00:56	26	1	SYSTEM_PERIODIC_HOLD Memory/limit 2263/2.000000000000000E+03

1 2 3
Some of my jobs are held… I need to increase memory request.

Case Study: Upper Management

“What does the computing division do again?”

Sage (and minions)

Photo: Public Domain

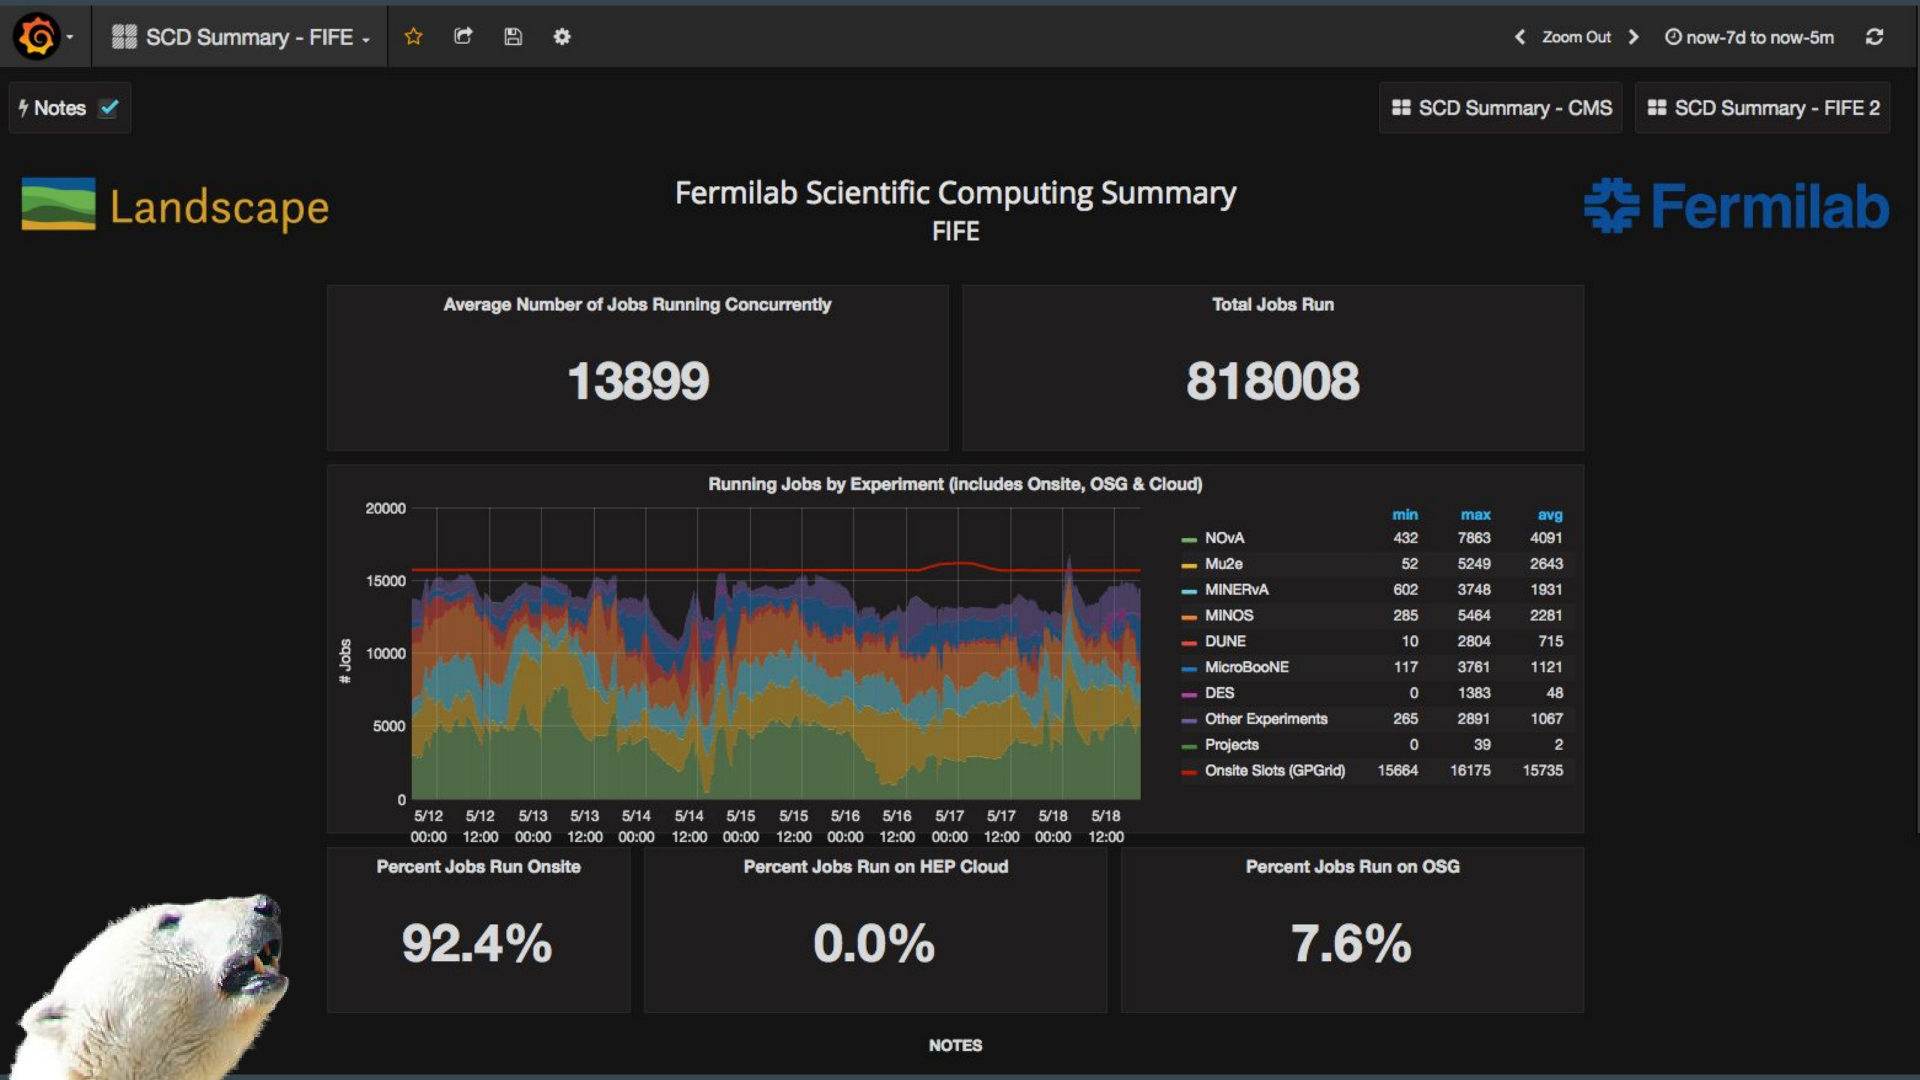

Fermilab Scientific Computing Summary

FIFE

Average Number of Jobs Running Concurrently

13899

Total Jobs Run

818008

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

	min	max	avg
🟩 NOvA	432	7863	4091
🟨 Mu2e	52	5249	2643
🟦 MINERvA	602	3748	1931
🟧 MINOS	285	5464	2281
🟥 DUNE	10	2804	715
🟦 MicroBooNE	117	3761	1121
🟪 DES	0	1383	48
🟪 Other Experiments	265	2891	1067
🟩 Projects	0	39	2
🟥 Onsite Slots (GPGrid)	15664	16175	15735

Percent Jobs Run Onsite

92.4%

Percent Jobs Run on HEP Cloud

0.0%

Percent Jobs Run on OSG

7.6%

NOTES

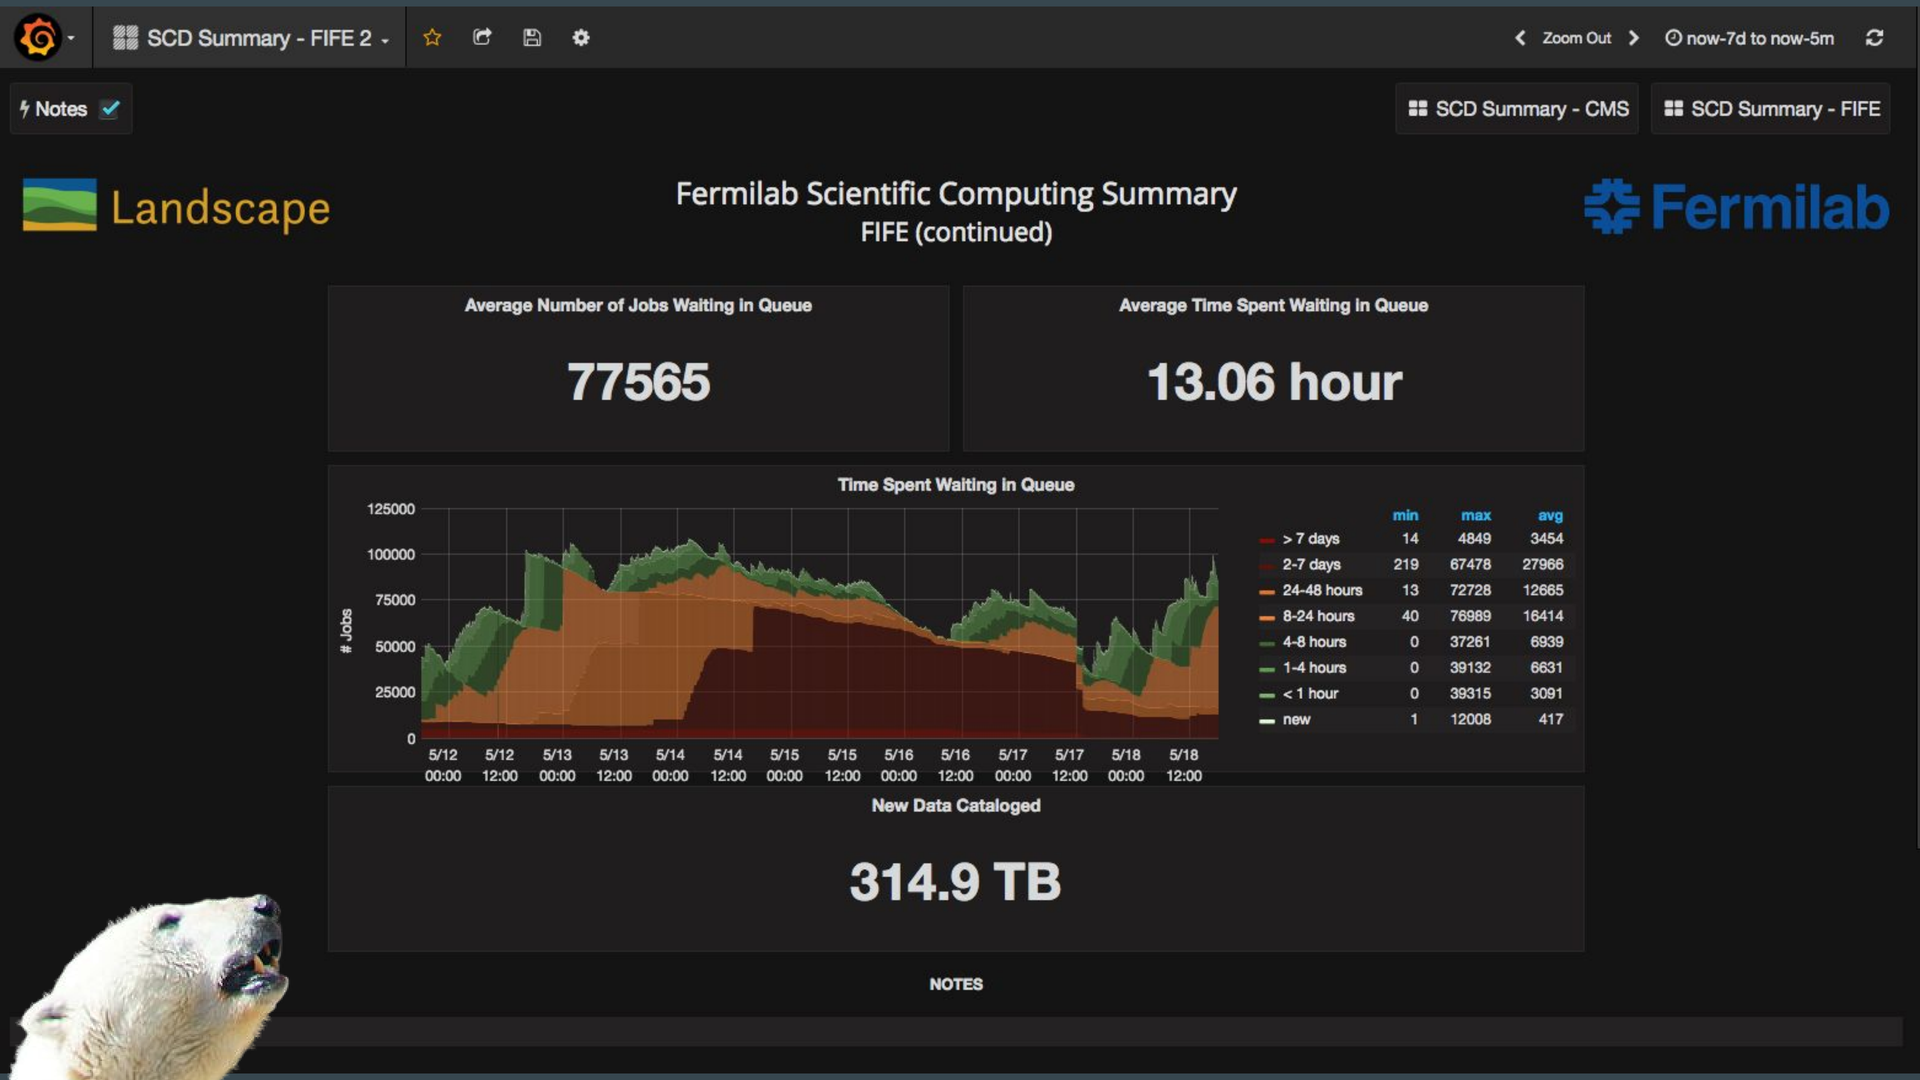

Fermilab Scientific Computing Summary

FIFE (continued)

Average Number of Jobs Waiting in Queue

77565

Average Time Spent Waiting in Queue

13.06 hour

Time Spent Waiting in Queue

	min	max	avg
🟥 > 7 days	14	4849	3454
🟥 2-7 days	219	67478	27966
🟧 24-48 hours	13	72728	12665
🟧 8-24 hours	40	76989	16414
🟩 4-8 hours	0	37261	6939
🟩 1-4 hours	0	39132	6631
🟩 < 1 hour	0	39315	3091
⬜ new	1	12008	417

New Data Cataloged

314.9 TB

NOTES

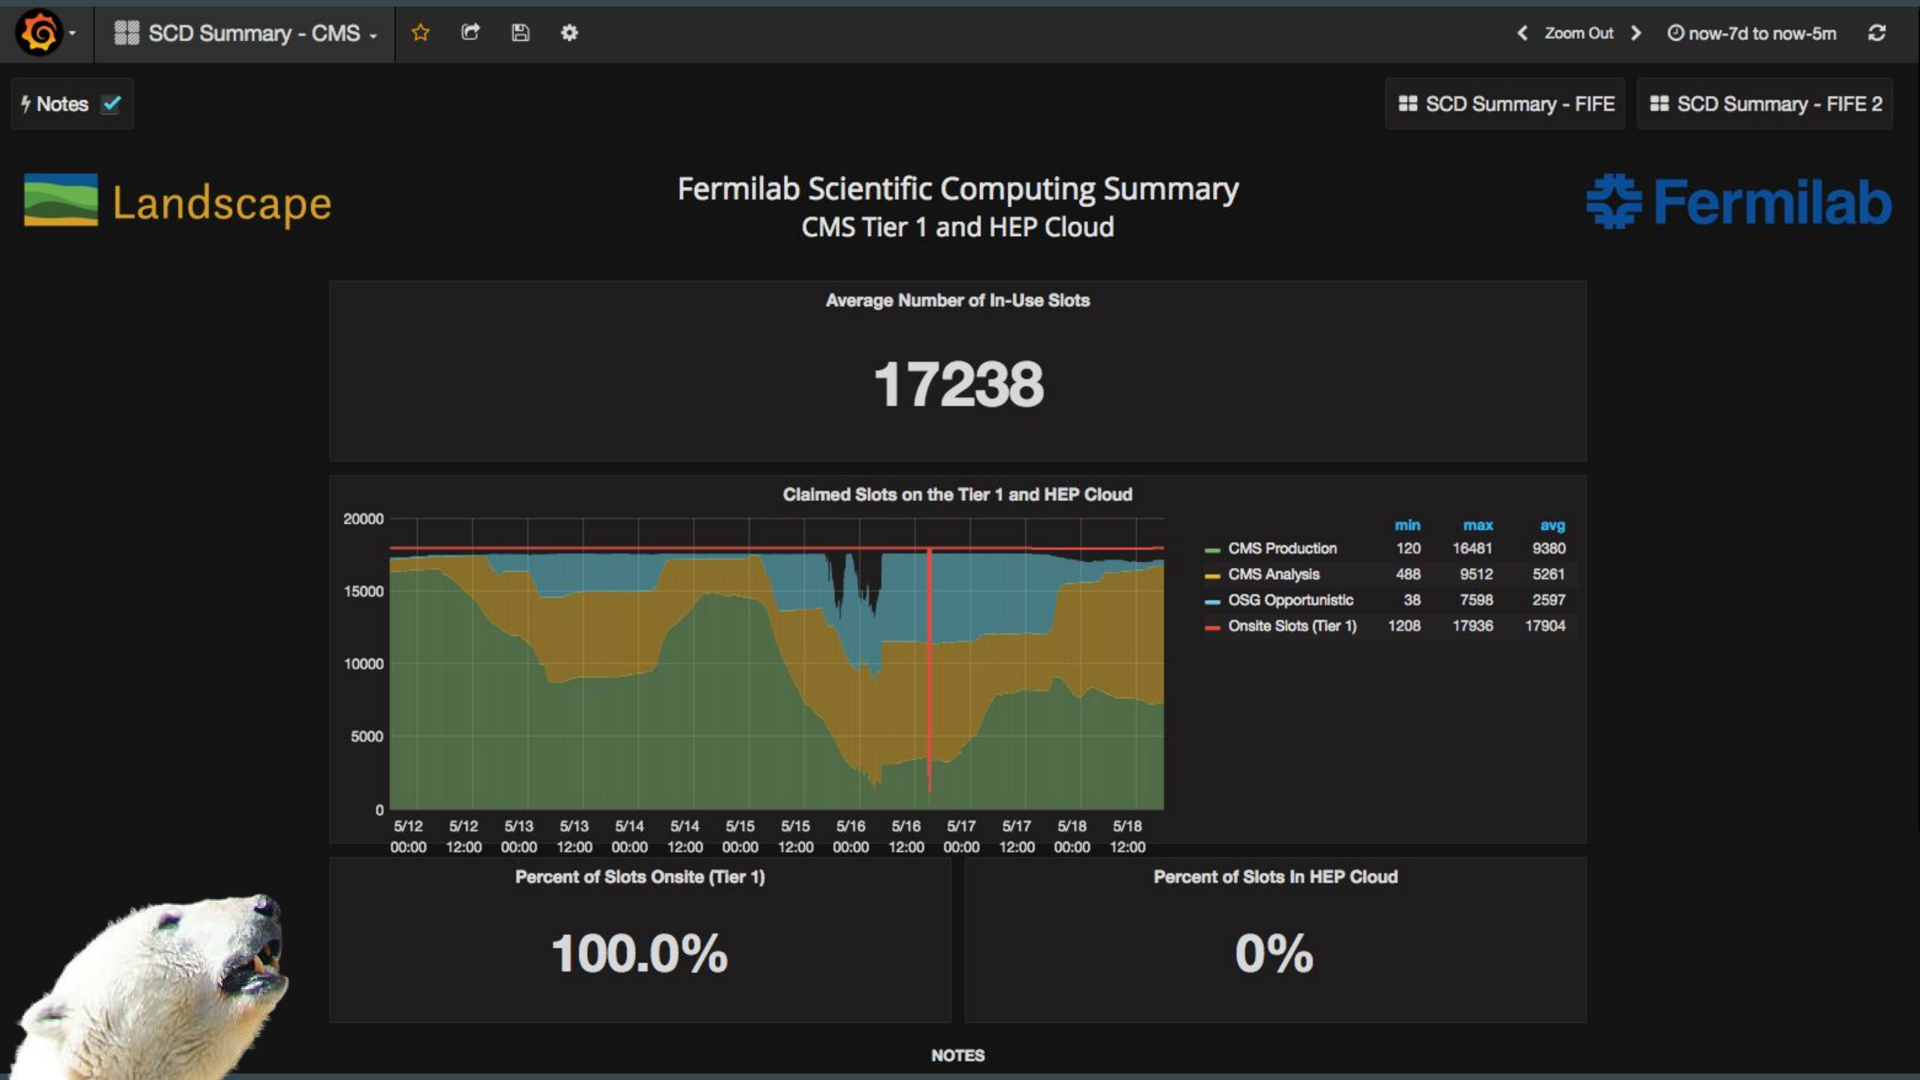

Fermilab Scientific Computing Summary

CMS Tier 1 and HEP Cloud

Average Number of In-Use Slots

17238

Claimed Slots on the Tier 1 and HEP Cloud

	min	max	avg
🟩 CMS Production	120	16481	9380
🟨 CMS Analysis	488	9512	5261
🟦 OSG Opportunistic	38	7598	2597
🟥 Onsite Slots (Tier 1)	1208	17936	17904

Percent of Slots Onsite (Tier 1)

100.0%

Percent of Slots In HEP Cloud

0%

NOTES

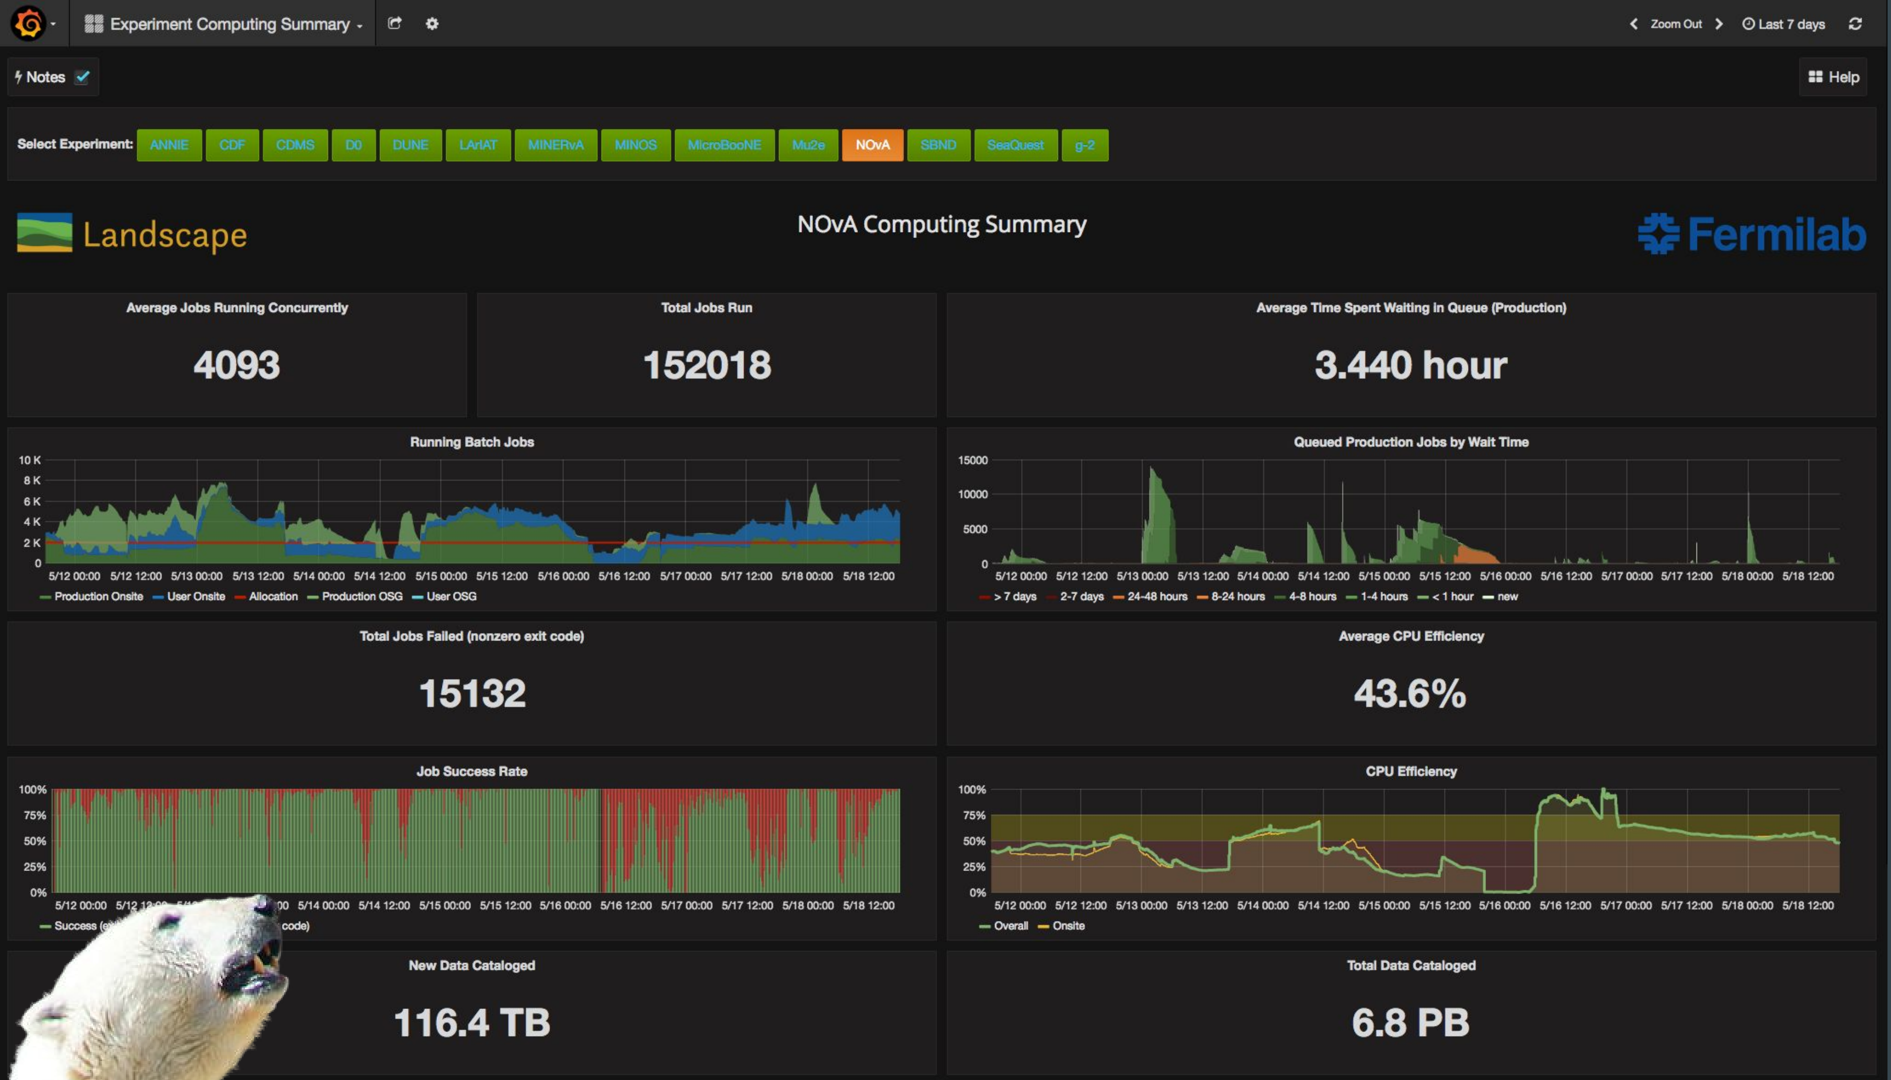

Experiment Computing Summary

Zoom Out

Last 7 days

Notes

Help

Select Experiment: ANNIE CDF CDMS D0 DUNE LArIAT MINERvA MINOS MicroBooNE Mu2e NOvA SBND SeaQuest g-2

Landscape

NOvA Computing Summary

Fermilab

Average Jobs Running Concurrently

4093

Total Jobs Run

152018

Average Time Spent Waiting in Queue (Production)

3.440 hour

Running Batch Jobs

Queued Production Jobs by Wait Time

Total Jobs Failed (nonzero exit code)

15132

Average CPU Efficiency

43.6%

Job Success Rate

CPU Efficiency

New Data Cataloged

116.4 TB

Total Data Cataloged

6.8 PB

Comprehensive grid monitoring with Fifemon has improved resource utilization, job throughput, and computing visibility at Fermilab.

Probes, dashboards, and docs at:

https://github.com/fifemon