

# DUNE Software and Computing News and Announcements

Tom Junk

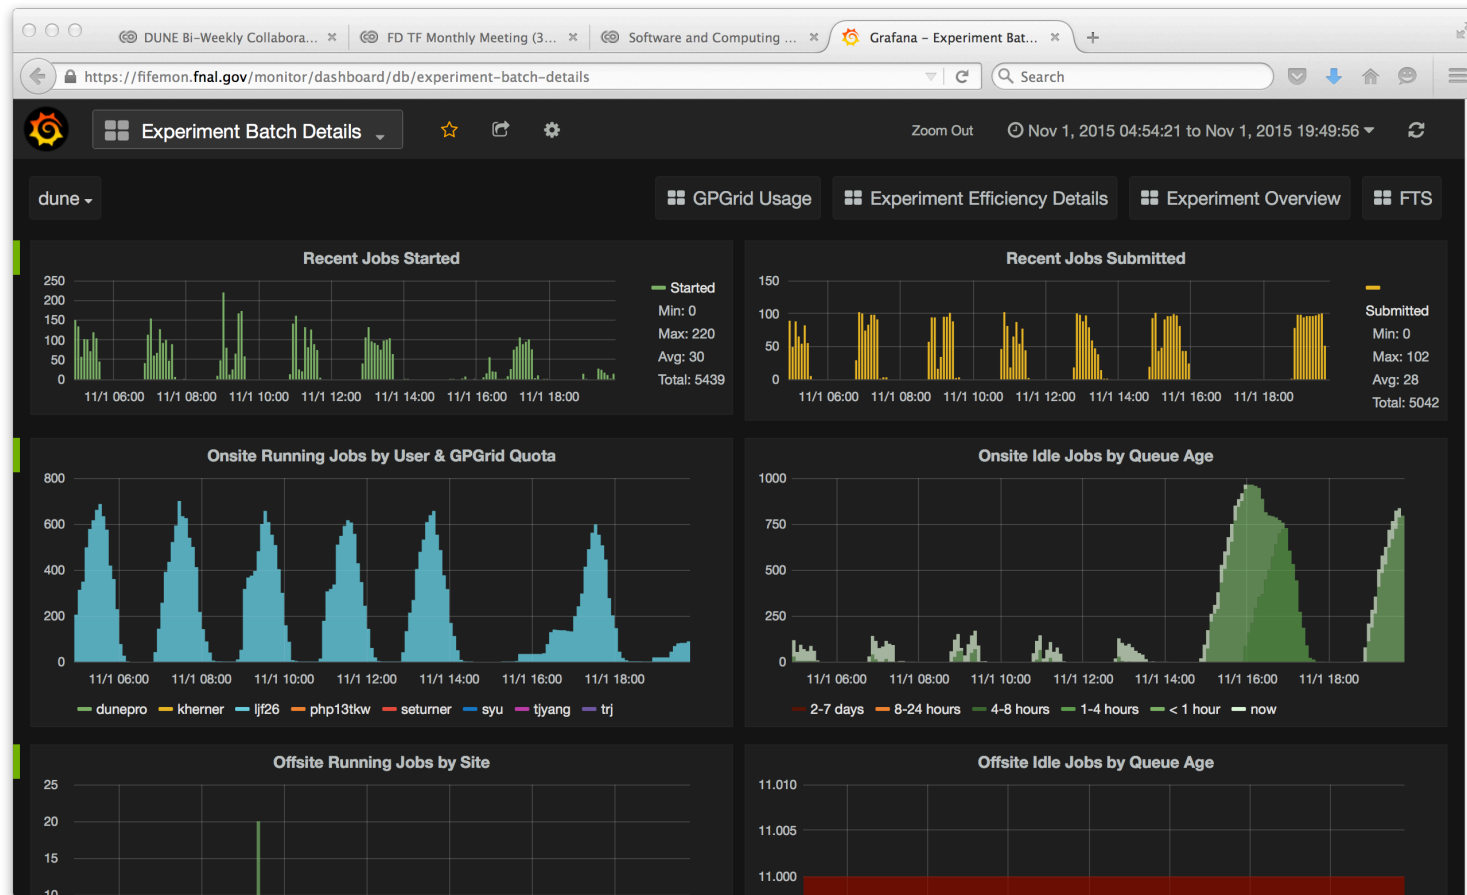
DUNE Software and Computing General Meeting

November 3, 2015

# New FIFEMON

Sign in with your Fermilab Services credentials

<https://fifemon.fnal.gov/monitor/>



## Fermigrid Batch Operations News

- GUMS transition Sep. 29
  - jobs used to run as user lbneana (for the LBNE VO) and as duneana (in the DUNE VO).
  - Now they run as the user (yay!). Files transferred back now owned by the user by default.
  - This helps esp. for users in many VO's (LBNE and DUNE for example). Users had gotten used to making their output directories on BlueArc group writeable so that lbneana could write to them. But duneana is in a different group from lbneana's group (!). World-write permission wouldn't be great.
  - Also I was not fond of the lbneana account filling up its quota on the BlueArc data areas /lbne/data and /lbne/data2
  - hard to tell who's responsible and what files to clean up

## Fermigrid Batch Operations News

- Job limits being phased in
  - Jobs satisfying the following constraints will be put on Hold
    - if they use twice the requested memory (they're being nice!)
    - use twice the requested disk space
    - get restarted 10 times
- jobsub v1\_1\_7 went in – no one noticed, though perhaps stronger checking of xml files when submitting dags.

## Migration away from AFS at Fermilab

- We use it for
  - user home areas – migrate to NFS
  - shared web areas (but these have largely been migrated to NFS already. Added benefit of migrating away from old Solaris web servers.
  - Connecting to non-FNAL AFS cells (like /afs/cern.ch)

[Andy Romero's talk Oct 14](#)

## dCache News

- New Disk installed: 2.4 PBytes
- We have 140 TBytes of persistent space in /pnfs/lbne/persistent which is shared with /pnfs/dune/persistent but which has a separate namespace.
- Not backed up!
- Do we need to ask for more? We've used 11% already. Maybe ask next cycle, not this one.

[dCache presentation to CS Liaisons](#) Oct 14, 2015

A monitor page of who's using what on the persistent dCache disks (Thanks, Stu!!)

[http://fndca3a.fnal.gov/cgi-bin/du\\_cgi.py](http://fndca3a.fnal.gov/cgi-bin/du_cgi.py)

## Speaking of Disk Space ...

- /lbne/app (= /dune/app) is now 90% full, of a total of 2 TBytes.
- Some users had been storing data files.
- Other users have left DUNE
- But there is a lot of useful work in these files too.
- Some of it can be archived (data disk, dCache, tape..)
- In a pinch, there is an admin machine that has permission over all lbne files

## KCA News

- Kerberos Certificate Authority
- certificates used for job submission and operation, and web authentication. Some ifdh commands need a KCA certificate obtained with KX509.
- Certificates have a lifetime of 7 days
- Moving to SHA2 (SHA1 is being deprecated, concerns over the integrity of the hash).
- each experiment and each application needs to test the new certs (jobsub and ifdh have teams for development/testing)

## [KCA presentation to CS Liaisons](#)



## Fermigrid Job Efficiency Web Pages

Some jobs wait long times for i/o, occupying a slot that could be used for computing.

Important to chase down worst workflows and see how to improve the fraction of the time the job is doing useful work.

Daily, weekly, and monthly reports of low-efficiency jobs:

- [http://web1.fnal.gov/scoreboard/daily\\_reports/fife-efficiency.daily.latest](http://web1.fnal.gov/scoreboard/daily_reports/fife-efficiency.daily.latest)
- [http://web1.fnal.gov/scoreboard/weekly\\_reports/fife-efficiency.weekly.latest](http://web1.fnal.gov/scoreboard/weekly_reports/fife-efficiency.weekly.latest)
- [http://web1.fnal.gov/scoreboard/monthly\\_reports/fife-efficiency.monthly.latest](http://web1.fnal.gov/scoreboard/monthly_reports/fife-efficiency.monthly.latest)

Can contact users and try to help them. Frequently they already know (“Why isn’t my job finishing faster?”)

## Upcoming SCPMT Requests

- Will get a template and guidance for how to format the request
  - last year – CPU hours (not “slots”), though SCD was interested in the integrals of spikes: big demands in advance of conferences, etc.

### Last Year's request

#### Summary: FY15 Requests: Batch and Storage

| Group               | CPU Hours | Disk (TBytes)<br>Permanent | Tape (TBytes) | Peak Slots |
|---------------------|-----------|----------------------------|---------------|------------|
| Beam Simulations    | 2M        | 20                         |               | 2000       |
| MARSLBNE            | 3.6M      | 8                          |               | 1000       |
| 35-ton, FD Sim/Reco | 4M        | 10                         | 400           | 2000       |
| ND Sim/Reco         | 0.6M      | 4                          |               | 500        |
| Cosmogenics         | 0.5M      | 10                         | 50            | 200        |
| FastMC/Physics WG's | 3M        | 7                          |               | 2000       |
| DAQ                 | modest    | incl in 35t                |               |            |
| Water Cherenkov     | 1M        | 0.5                        |               | 120        |
| Total               | 14.7M     | 59.5                       | 450           | 6000       |

Need to make new requests for each group and justify them.

# Computing Model Document

- Just got started:
  - Maxim Potekhin – main author
  - Amir Farbin
  - Tom Junk

<https://github.com/DUNE/dune-computing-model>
- Timeline:
  - CD-3a (Early Dec.) Progress (LBNC’s wording). We interpret this to mean that a rough draft should be ready.
  - End of 2015: polished draft.
  - We will need reviewers!
- Major Divisions:
  - Design and prototyping era (2015-2020)
  - DUNE experiment construction, commissioning, and operations

## *art* News

- *art* release 1.16.02, on which larsoft v04\_27\_00 (use `-q e9:prof` or `-q debug:e9` now) is built, has a bug in that *art*-formatted files written with earlier versions cannot be read. Reason: there's a new ResultsTree produced in the output file, and this version of *art* insists on its presence and chokes if it is absent.

The ResultsTree is now added so that trees, histograms, etc. can be stored in the *art*-formatted output file and not just in the TFileService-supported output file. This way you get provenance info for your ntuples and histograms, not just events.

Fixed in *art* release 1.17.02, on which larsoft v04\_28\_00 is built

## *art* News

- v1.18 is a “Technology Preview Release” built on ROOT6

[https://cdcv.s.fnal.gov/redmine/projects/art/wiki/Release\\_Notes\\_11800](https://cdcv.s.fnal.gov/redmine/projects/art/wiki/Release_Notes_11800)

Specifically, ROOT v5\_34\_32 -> v6\_04\_06

separate cpp0x package is now retired. *art* now assumes that the compiler used will be fully compliant with the standards required by that version. (We distribute appropriate compilers along with the larsoft/art bundle).

From the release notes page: there is a significant memory increase relative to [1.17.03](#) due to ROOT 6's autoparse facility. The exact amount will vary with experiment's configurations and data product use, but could be over 200MiB. We are confident we will be able to address this significantly in a future *art* release by taking steps to avoid triggering the autoparse behavior.

## 35t Computing News

- Data Handling
  - Following [instructions](#) from Qizhong.
  - Tom has written simple scripts that obtain Kerberos authentication, scp files, and check the error codes and checksums of the copied files.
  - Tom has also received an example online database query python script from Jonathan Paley that takes as input the run number and returns the run mode, the configuration name, the component list, and the start and end times. Tom reformatted the string output so it fits in the JSON format from Qizhong but needs to integrate this script with Qizhong's to get complete JSON files (an afternoon's work)
  - To do – automatically upload metadata to SAM – Qizhong's scripts provide examples. Tested on lbne35t-gateway02.
  - We have all the pieces – maybe a day's worth of work to get it all put together.

## 35t Computing News

- Online database filled by run control and replicated offline. Tom has verified that the metadata extraction works using queries to this database.
- OPOS (Offline Production Operations Service)  
Tingjun, Karl, and Tom met with the OPOS group  
[https://cdcvs.fnal.gov/redmine/projects/offline\\_production\\_operations\\_service](https://cdcvs.fnal.gov/redmine/projects/offline_production_operations_service)
- Their job is to shepherd jobs through the batch system and check for proper completion and resubmit failed ones.
- Tingjun and Karl demonstrated the project.py workflow to the OPOS team and they were pleased with how automated it is. Thanks to all of Herb Greenlee's hard work!
- MCC5 to be started in about 1 week. Single particles and neutrinos

<http://indico.fnal.gov/getFile.py/access?contribId=1&resId=0&materialId=slides&confId=10676>

## SAM Tools for Analysis Users

[Link to Andrew Norman's Presentation](#)

Simplified tools for creating dataset definitions and archiving data to tape.

Probably won't use it for production work – project.py is already interfaced to SAM via samweb commands. We already make metadata JSON files for storing the raw data.

But ... users may be interested in a more lightweight dataset definition tool for managing their large sets of analysis ntuples. See Andrew's NOvA examples.



## Cleaning up Old LArSoft Releases in /grid/fermiapp

- e-mail from Erica Snider, Nov. 2
- In order to free up space in the LArSoft products area on GPCF, /grid/fermiapp/products/larsoft, **we plan to remove all releases prior to the v04 series — including production releases — on Nov 9 (next Monday!!)**. Note that this change will also be reflected in the releases available via cvmfs.
- Per the LArSoft release retention policy, we need approval for removing the production releases, so please let us know if you are using any production release prior to v04.02.00.
- Should the need arise, any production release that is removed can be restored upon request by any experiment

## Simulations for Neutrinos Meeting Summary

<https://indico.fnal.gov/conferenceDisplay.py?confId=10677>

GENIE news G. Perdue

GEANT4 news K. L. Genser

Version 4.10.1p02 is the latest

Distributed as a scisoft.fnal.gov bundle

GEANT4 parameter sensitivity study – J. Yarba

GEANT4 validation data summary database H. Wenzel

NOvA simulation tuning real-world example: hadronic recoil energy measurement is not simulated well and it is “recalibrated” in the data – P. Vahle