

DUNE Software and Computing News and Announcements

Tom Junk
DUNE S&C General Meeting
December 15, 2015

Fermilab Computing Resource News

- New interactive node: `dunegpvm11.fnal.gov`, configured much like the `lbnegpvm01` through `lbnegpvm10`, though with these differences:
 - Runs SLF 6.7 (others run 6.6)
 - Separate set of users from the `lbnegpvm` machines
 - Users on `dunegpvm11` are members of the `dune` (`gid=9010`) group and the `lbne` (`gid=9960`) group. The `9960` group has been re-named “`lbnf`” but that change has not yet propagated everywhere. Users are able to read and write their old files, but new files get created in the `dune` group. A `dune` file may not be writeable by someone on an `lbnegpvm` machine.
 - Users on `lbnegpvm*` machines are only members of the `lbne` group
 - Users who have not been active in the last few years or never joined DUNE were not automatically given accounts on `dunegpvm11`. If you were missed, apply for a new DUNE account at <http://servicedesk.fnal.gov>

Fermilab Computing Resource News

- New interactive node: `dunegpvm11.fnal.gov`, configured much like the `lbnegpvm01` through `lbnegpvm10`, though with these differences:
 - The BlueArc mounts `/lbne/data`, `/lbne/data2`, and `/lbne/app` are not available on `dunegpvm11`. The same directories are mounted as `/dune/data`, `/dune/data2` and `/dune/app` on both the `lbnegpvm` and `dunegpvm11` machines to allow users to migrate their scripts.
 - The `$HOSTNAME` variable says `dunegpvm11`. The `lbnegpvm`'s have `dunegpvm*` convenience names but `$HOSTNAME` still says `LBNE`.

We'd like to migrate our full suite of `lbnegpvm`'s to look like `dunegpvm11`.

Propose: `lbnegpvm06` through `lbnegpvm10` to `dunegpvm`'s today, Dec. 15. Follow on with the remainder on Dec. 21. Do we keep one spare longer? Probably just force people to change.

Moving away from the LBNE VO

- Well, at least for submitting jobs at Fermilab.
- The gid=9960 name change to lbnf broke jobsub client job submission to gpgrid
- cpn lock could not be acquired to transfer job wrapper scripts
- probably cpn would fail for user file transfer also
- Many users have expired AUP's in the LBNE VO anyway
- Still some users running LBNE jobs as of Dec. 11 – Beam sims.

- Users should use DUNE VO for job submission

```
jobsub_submit -G dune ...
```

```
jobsub_q -G dune
```

The “-G <group>” name is a “required” command-line option (according to the jobsub_client wiki), but the JOBSUB_GROUP environment variable can supply its value. Users may not know to look there.

- DUNE VO membership AUP is automatically refreshed as users must re-sign the Fermilab Computing Policy to renew their accounts.

DUNE Production Certificate

- I got a new production certificate for DUNE on Neha Sharma's request, to make its name consistent with other certificate names.
- Installed on fifebatch machines
- Used to declare metadata from the 35t file transfer (samweb decaire-file needs it)
- The reason I bring this up is because the issuer line for OSG Certificates is now /CN=DigiCert Grid CA-1 G2 (The G2 is new! Wasn't in the SAM list of known CA's).
- New OSG certificates will be issued by a new provider starting February 2016. Old certificates are still good until they expire.

To Do item: Data Catalog Web Pages

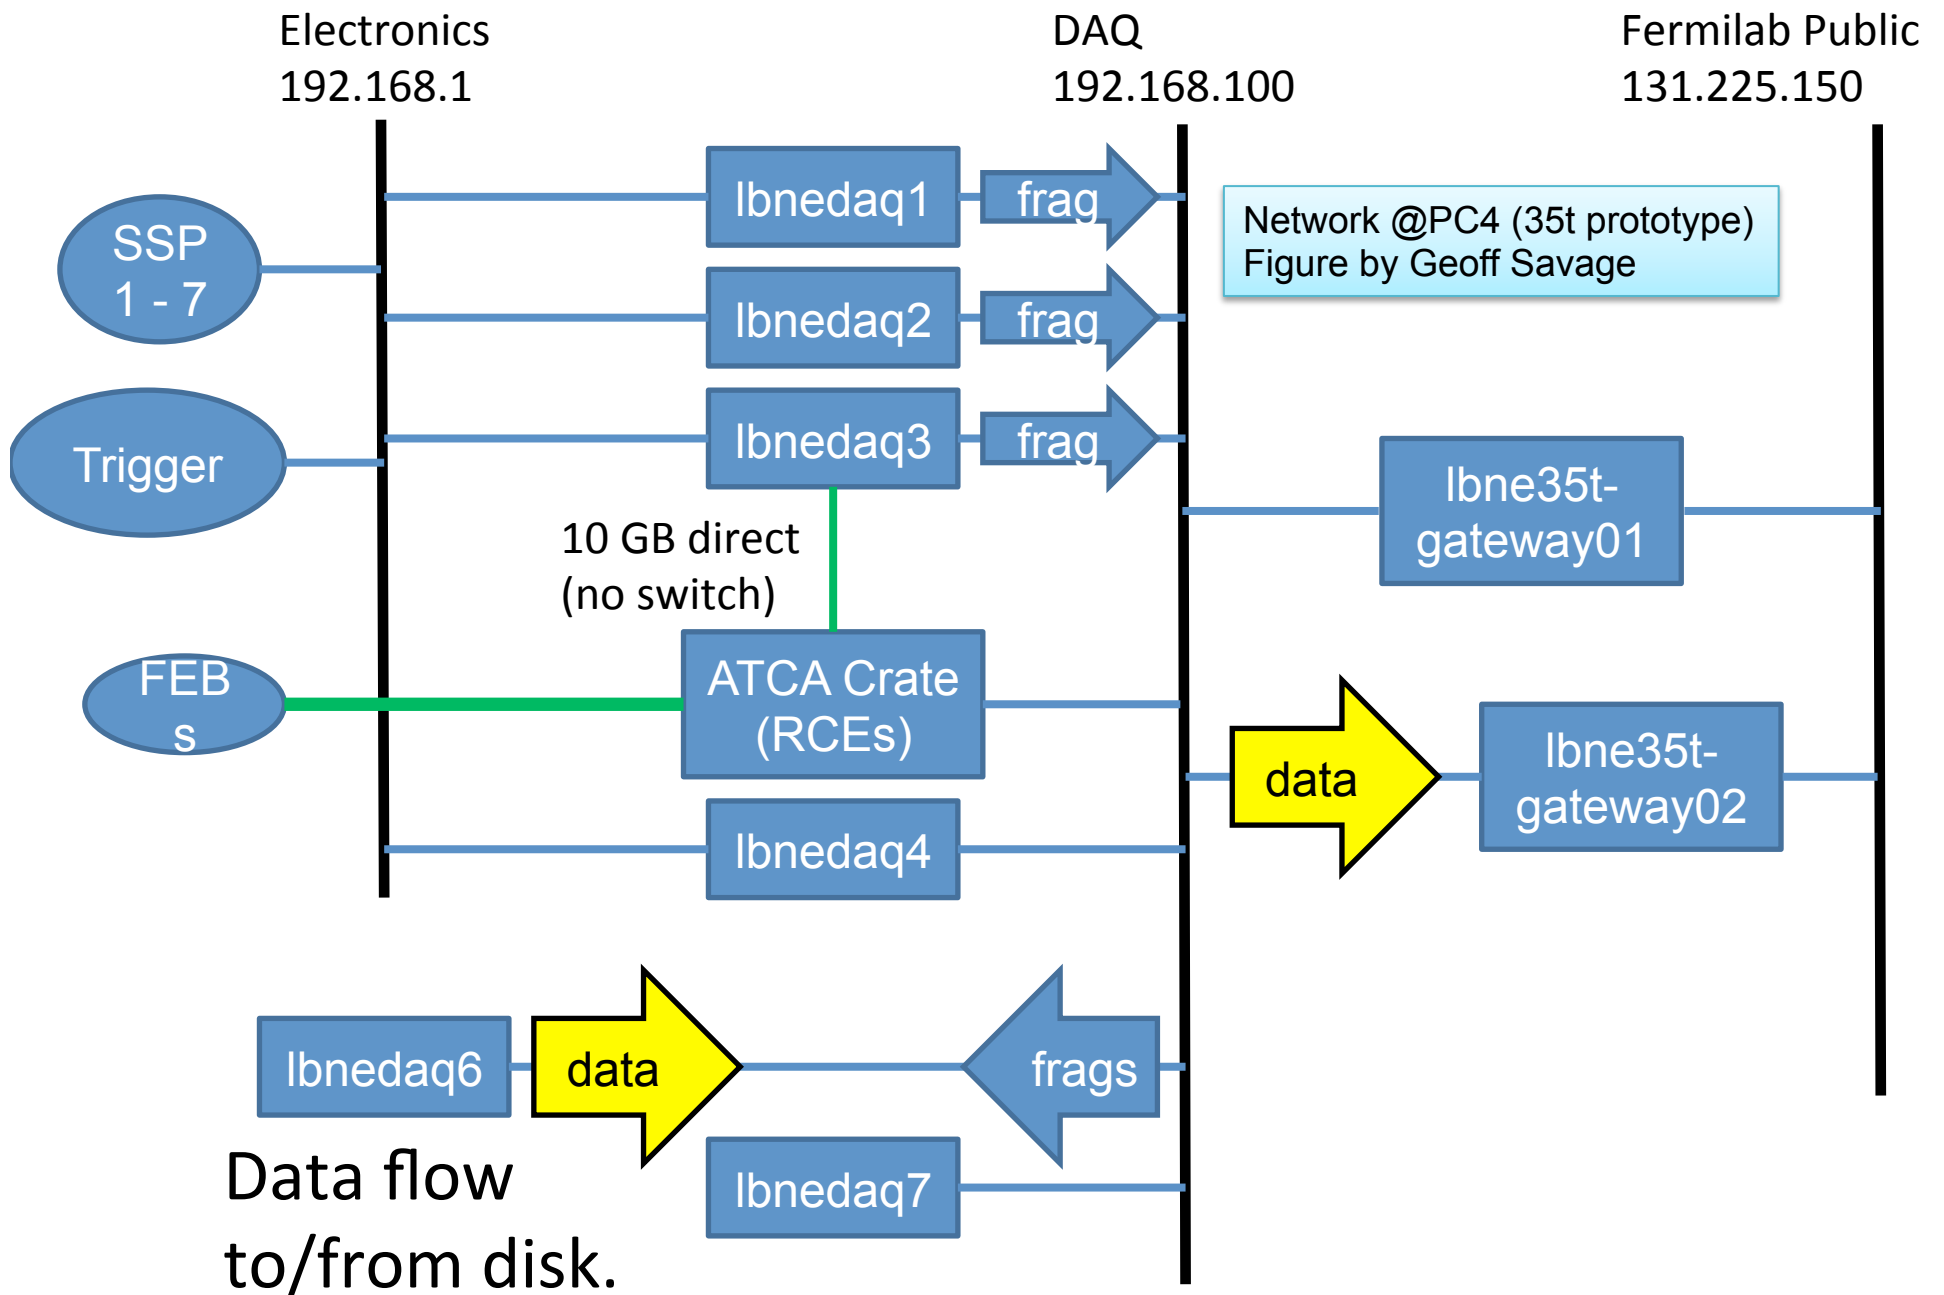
- I requested a centrally-managed web site for data catalog use. I haven't heard back from the Service Desk yet.
- Currently online and offline we're listing locations of data and MC files in Redmine Wiki pages – this doesn't scale at all and it is impossible to script.
- The 35-ton data transfer script makes a file listing all the transferred data files and their sizes and event counts.
- SAM can retrieve other metadata, though the metadata files are kept on disk on lbne35t-gateway02 for now (they're small).

Collecting Requests for Fermilab Computing Resources

- Tom is scheduled to give a talk on Feb. 4, 2016, presenting DUNE's computing requests.
- Last year's talk and spreadsheet are in DUNE DocDB 862 – covered years 2015, 2016.
- Template spreadsheet for the upcoming request is in DUNE DocDB 863
- Similar format to last year's request
 - CPU hours, Fermigrid and OSG. Now broken into 2 GB VM's, 4 GB, and 4+ (or multi-threaded) request. CPU gets accounted for how many 2 GB slots are occupied even if only one core is doing any work. Some experiments have some very large memory needs.
 - peak CPU demand (how many peaks, how many slots for how long)
 - Storage: Disk, tape
 - Services we need from SCD

Jobsub News from the Dec. 9 Liaisons' Meeting

- Jobsub 1.1.9
 - expected production deployment 12/17/15
 - specify 'expected lifetime' of the job - short (6 hrs), medium (12 hrs), long (24 hrs). Default - short
 - used to match against resources advertising that they have REMAINING_LIFETIME seconds left
 - sandbox_readable_by_group
 - anyone in group can fetch logs for all jobs submitted under that group
 - **liaisons -- please let us know if you would like this enabled for your group**



35-ton DAQ and DAQ interface work


- Working to identify bottlenecks
 - Data → disk on lbndaq6 (or lbndaq7)
 - Raw disk write speeds in excess of 300 MBytes/sec
 - DAQ file write speeds of 40-60 MBytes/sec
 - Network out of PC4: 1 GBit/sec = 125 MBytes/sec
 - We don't always fill that network bandwidth even if we try, but we'd like that to be the bottleneck, and not other limitations within the DAQ network/artdaq processing/file copying and checking in PC4.

Transferring 35-ton data files

- Currently rather awkward – copy files from lbnedaq6 to lbne35t-gateway02 with a cron job. Checksum validated on this copy on both sides (slow!).
- Then make metadata JSON file
 - Get run number from filename
 - Get file size from ls
 - Get event count by running root on the file and looking at the Events tree entry count
 - Get component list, start and stop times from online database
- Declare metadata to SAM
- Copy file to dropbox in dCache which FTS monitors.

File transfer service monitoring on lbnesamgpvm01

<http://lbnesamgpvm01.fnal.gov:8787/fts/status>



Most errors before Nov 30 due to no end time recorded in online DB for the run. Time was written as “None”. Now the metadata set the end time to the begin time if it is “None”.

Failed transfers: One 595 GByte file never got copied to tape (it made it to the dropbox in dCache). Another mystery file didn’t transfer in Nov. One failed copy in Dec – different file size on gateway and on dCache last Sunday. Files kept on gateway until verified that they are on tape (scripted, but not a cron job).

S&C & Physics Tools Related Sessions at the Collab Meeting

Parallel sessions

- S&C + Computing Model (1 session)
- Physics Needs for simulation/reconstruction (1 session)
(Physics? Sim/Reco?)
- 35t (1 session)
- FD Sim/Reco (2 sessions)
- ND Sim/Reco (1 Session)
- Analysis Tools (1 Session)

Plenary Talks

- 35t: Michelle Stancari
- DUNE Computing: Amir Farbin or Tom Junk
- DUNE Reco: Tingjun Yang