

Software Management News and Issues

Tom Junk

DUNE Software Architecture Meeting

June 13, 2024

DUNE Framework Workshop

Fermilab, June 27-28 2024

<https://indico.fnal.gov/event/65009/>

SL7 → AL9 Migration

- Done! All DUNE interactive computers are running AL 9.4
https://wiki.dunescience.org/wiki/DUNE_Computing/DUNE_Interactive_Computing_Resources#Hardware_Resources
- We have three build nodes. The old dunebuild01.fnal.gov was retired on the upgrade and replaced with a 64-CPU (not sure if it's real cores or just 32 with hyperthreading, but still...) with 125 GiB of RAM. Still called dunebuild01.fnal.gov
- Most things "just run" under AL 9 just fine
- UPS is not supported, however
- UPS was a major part of our software build, packaging, distribution and runtime setup environment
- Mostly extracted from the build environment with CETModules

SL7 Containers

- A requirement of the S&C infrastructure is we must be able to reproduce old calculations
- We keep many old releases of dune software in CVMFS for this purpose. But they are set up with UPS.
- To keep running existing and legacy workflows, SL7 containers are provided.
- Instructions for starting a SL7 container are available here: https://wiki.dunescience.org/wiki/SL7_to_Alma9_conversion
- Container image is in CVMFS and should be visible ~anywhere.
- I am impressed at how fast apptainer starts a new container (instantaneously!)
- You may have to customize the filesystem mount bindings if you run it on a machine that's not a dunegpvm. Such as a build node (no /pnfs)
- One minor annoyance – your .profile (or .shrc, .bashrc) are not sourced automatically. You have to do this yourself.
- Container task force meetings have taken place ~monthly since the beginning of the year, but they are winding down.

SL7 Containers

- We have been giving feedback to the container task force about packages that need to be present and have gotten good response from the team
 - editors (emacs, nedit...)
 - packages needed for development (dunesw compiles!)
- I haven't tried everything needed for a complete workflow, but some things are meant to be run on AL 9, like jobsub_lite, and I also sftp files around outside containers.
- Vito di Benedetto has a separate container with jobsub_lite features and has made it available to SBN.
- I found this one was needed to run forge_tools ddt. Problem with regular container was mounting a pseudo TTY. May just be in the invocation.
- https://sbnsoftware.github.io/GPVM_migration#sl7-development-container
- I helped a MINERvA 2x2 user who needed cvs in a container. Maybe that needs to be added to the standard one.

LArSoft v09_90_02

- We already have dunesw v09_90_01d00 released, and larsoft v09_90_02 came out June 5.
- https://github.com/LArSoft/larsoft/releases/tag/v09_90_02
- I've been a bit slow to get around to making the release
- Some PRs have been merged in dunesw v09_90_02d00, available June 19
 - Trigger decoder updates from Wes
 - ICEBERG fcls for DUNE-DAQ 4.4.1 (using DUNE-DAQ 4.4.0 code, assume to be compatible)
 - DUNEANA PR64 calibanatree
 - DUNEANA PR67 raw digits -> ROOT TTree
 - DUNERECO PR104 collection V in PDHD
 - DUNERECO PR102, 103 PDHD reco
 - DUNEPROTOTYPES PR47 DAPHNE Channel Map service
 - DUNEANA PR66 – tweak to AnaTree to save less G4 truth

duneanaobj and updated duneana

- The long-standing version shear between FD and ND caf files has been resolved
- Both ND and FD now use duneanaobj v03_05_00, though the old e20 build is maintained for ND use due to the need to run with the old HDF5 product with the C++ interface

Geometry refactoring

- Tom and Dom have been testing Kyle's PRs which are targeted at supporting more general geometries
- larsoft v10_00_00rc<n> (n=1 or 3) are the corresponding larsoft releases to test with it.
- Tom built with the PRs, asked Dom to test with a muon in the FDHD 1x2x6 workspace
 - Bugs found in rc1: Drift sign wrong in half of FDHD TPCs and wires that don't intersect were not flagged as such in the wire intersection calculator.
 - Kyle fixed these and rc3 was cut.
 - Tom tested the same muon workflow and it now looks good
 - But the event display is broken
<https://cdcv.s.fnal.gov/redmine/issues/28852>
 - I've been using v09.x to run the event display to test the rc.

Testing Generators

On May 15, Lynn sent this to the larsoft and larsoft-coordination lists:

SBND requested a new genie release: <https://cdcvns.fnal.gov/redmine/issues/28747>

Genie v3_04_02 is now available in larsoft v09_90_01_01 for testing. We understand that there are no physics changes, so we hope the update can be approved soon.

https://github.com/LArSoft/larsoft/releases/tag/v09_90_01_01

Testing Generators

The GENIE version separation project never finished. LArSoft version implies a GENIE version.

This is not enormously crazy as the API changed and various data structures were filled differently from GENIE 2 to GENIE 3.

But it does tie all the LArSoft-dependent experiments' GENIE versions together.

I sent the request to the Physcos and Mike Kirby before just signing off.

This one was an easy one – "no" physics changes.

Testing Generators

- Mike Kirby asked if we have an automated validation suite for the generators to test this sort of thing
- As far as I know, we do not
- Neutrino Interactions WG is a natural forum to ask for standardized metrics
- Would be great to have comparisons between generators
- Just Argon, standard flux would be good.
- I have a standalone workflow that runs gevgen, gntpc and a ROOT analysis script that makes lots of plots -- I'm sure lots of people have made things like this.
- Standard reference plots and overlays needed. Web interface? Or expert-level generator validation task?
- **Not limited to generators: GEANT4 versions too!**

Looking for Help

DUNE Software Coordinator

The DUNE Software Coordinator leads a team of software managers, developers and stakeholders. Holds regular meetings to discuss software architecture, algorithm integration, status of the DUNE software stacks, open pull requests, testing and documentation. Attends LArSoft coordination meetings, offline leads meetings, and steering group meetings, and meetings organized by CSAID relevant to software development, maintenance, and deployment. Sets project deliverables, schedules, and deadlines. The software manager works with developers and physics experts to identify reviewers for pull requests and ensures timely review, decisions and integration. Coordinates the production of coding guidelines. The software manager follows the main progress in the physics working groups in order to ensure that software solutions are targeted at physics deliverables. Presents status at DUNE Software and Computing meetings and at DUNE Collaboration meetings.

Estimated FTE: 0.5

Job Descriptions

DUNE Release Manager, LArSoft-based software stacks

The release manager produces the weekly releases of LArSoft-based DUNE software, which is used for far detector simulation and reconstruction, as well as the ProtoDUNEs, coldboxes and ICEBERG. The release manager tests approved pull requests using the CI system and interactively builds and tests before merging them into the main develop branches of the repositories. The release manager runs builds for all supported platforms and deploys the built software for collaboration use on distributed filesystems and web servers. The release manager writes release notes for the DUNE wiki. ND-GAr releases are also included in this position's role as the structure is similar to LArSoft-related software.

Estimated FTE: 0.2, though non-recurring work involving refactoring software and upgrading build systems may create large temporary needs for effort

Job Descriptions

DUNE Release Manager, Near Detector

The release manager produces software releases used for ND-LAr, TMS and SAND.

The release manager tests approved pull requests using the CI system and interactively builds and tests before merging them into the main develop branches of the repositories. The release manager runs builds for all supported platforms and deploys the built software for collaboration use on distributed filesystems and web servers. The release manager writes release notes for the DUNE wiki.

Estimated FTE: 0.2, though non-recurring work involving refactoring software and upgrading build systems may create large temporary needs for effort

Job Descriptions

FAIR Data and Software Distribution Manager

The distribution of data for use outside of the DUNE collaboration must be accompanied by the software and documentation necessary to make use of it. This position is much like that of a software release manager, creating persistent collections of software artifacts and testing them in an environment that can mimic one generally available outside the DUNE collaboration. This manager is responsible for responding to requests for assistance from users of the FAIR data distribution.

Estimated FTE: 0.1

Job Descriptions

Software Operations Manager – Integration and Testing

The Software Operations Manager will write tests for current software, maintain the reference test outputs on the continuous integration system, and alert the software coordinator and relevant developers and the Software Coordinator of test failures that require significant changes to address. This role also involves rearranging software that may be classified in the wrong package, or which may need its interfaces upgraded or to respond to external changes in the runtime environment.

Estimated FTE: 0.2

Job Descriptions

Software Reviewers, Far Detector and Near Detector

Physics experts need to check, sign off and approve pull requests for new software and modifications to old software. These are expected to be volunteers, and there is benefit from doing this work for the volunteers: they become more familiar with the software and are able to use the software better and help others.

So far, WG conveners and physics experts have stepped up and reviewed PRs.

Job Descriptions

Software Training Instructor

The Software Training Instructor prepares training materials for new and experienced users and leads training sessions scheduled near the DUNE collaboration meetings. The Software Training Instructor maintains and tests the materials to ensure that they continue to function as the computing environment evolves. The Software Training Instructor responds to collaboration feedback – questions on how the software works and best practices, questions about development workflows and style, and comments on the training materials.

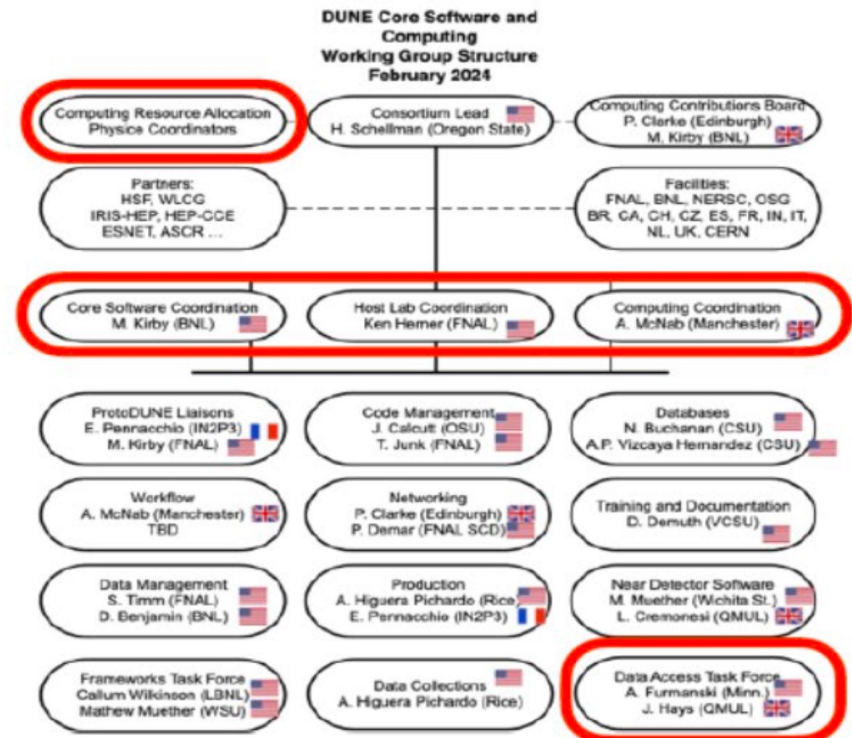
Estimated FTE: 0.1 per instructor. DUNE will need separate instructors for FD+ProtoDUNE software, ND software, and batch workflows.

From Mary's Talk @Collab Call April 19

DUNE Software and Computing reorg

Changes to the organization

- Additional groups:
 - Computing Resource Allocation Board (CRAB)
 - Data Access Task Force
- restructured Technical Roles
 - Computing Coordination
 - Host Lab Coordination
 - Core Software Coordination



April 19, 2024 DUNE Status, Collab Call April 2024



<https://indico.fnal.gov/event/64291/>

Mike Kirby's the new consortium lead, however. Still would like him to help with software coordination duties.

Spack Development and Tutorials

- <https://fifewiki.fnal.gov/wiki/Spack>
- Complete with wiki and video
- Transition not complete – unleashing it on DUNE will cause a lot of support need we do not know how to provide at the moment
- Kyle Knoepfel says he is working on a development environment that looks like mrb
- Patrick Gartung has submitted pull requests to each of the DUNE repositories and has created a dune_spack repository https://github.com/DUNE/dune_spack/
- I made small fixes in order to get the builds to work, and also merged the ones not in dune_suite (e.g. duneanaobj). garsoft still needs two fixes to de-upsify it.

Spack Progress and News

- Merging Patrick's cetmodules updates into dunesw v09_89_01d00 meant that we couldn't use larsoft v09_81_00 (Nov 2023) in spack – needed a new one.
- We asked the LArSoft team at one of the offline coordination meetings for updated spack distributions of LArSoft
- Patrick has heroically made a spack installation of larsoft v09_90_01

On an AL 9 machine:

```
source /cvmfs/larsoft.opensciencegrid.org/spack-v0.22.0-fermi/setup-env.sh
spack env activate larsoft-09-90-01-gcc-12-2-0-cxx17-prof-gcc-11-4-1
spack load larsoft
```

Spack Progress and News

- Met with Pengfei Ding after the collab meeting, and John Freeman the next week, to see how DUNE-DAQ works with spack (and has for the last three years or so – last UPS product update was June 2022).
- <https://dune-daq-sw.readthedocs.io/en/latest/packages/daq-buildtools/>
- <https://github.com/DUNE-DAQ/daq-buildtools>
- <https://github.com/DUNE-DAQ/daq-release>

Some takeaways from DUNE-DAQ Meetings

- development of individual repositories in private space on top of CVMFS works with DUNE-DAQ using daqbuildtools
- Pengfei tells me private builds like this still depend on LD_LIBRARY_PATH
- Separate procedure for cutting releases using github actions with containers
- containers are needed, not only to get an AL9 environment on github servers, but also to bind /cvmfs to local, writeable disk
- I prefer this to the procedure in our spack tutorial where a spack install is issued on the CVMFS publisher node.
- Would like a sandbox that looks like the release but not actually publish it before testing it.
- DUNE-DAQ scripts have DAQ-isms in them and would have to be adapted.
- I prefer using FNAL Jenkins nodes to GitHub's resources so we don't incur charges or run out of time – especially if we have to repeat builds until we get it right.

DUNE spack build recipes

- Many thanks to Patrick Gartung for making these!
- https://github.com/DUNE/dune_spack
- Need to be updated to build v09_90_01d00
- Some discussion in the DUNE spack channel about how much versioning information needs to go into these recipes
- Recipes should change as little as possible between versions, but:
 - dependencies change over time as code gets refactored and new repositories are added
 - APIs change and versions of software may depend on others even if the code builds and runs fine – get wrong answers if config parameters and code mismatch.
 - Can specify ranges of allowable versions
 - Or just tag this repo and each tag can only build that one version.