

DUNE FD/Prototypes Software Status

Tom Junk

DUNE Software Architecture and Management

November 22, 2024

LArSoft v10 – Breaking Geometry Changes

- Laura Paulucci reported a segfault in the VD detsim stage.
- She had asked about photon detector methods which have been moved around a bit in v10
- ProtoDUNE-VD CI tests pass
- Reaching out to Laura about the segfault.

dunetrigger

- Wes Ketchum gave me a heads-up about a DUNE trigger simulation product he's been working on:

<https://github.com/wesketchum/dunetrigger>

- It depends on dunecore, though some of the example fcls require other things (duneprototypes, dunesw)
- It would need updating to the new geometry API – just a couple of calls need to be adjusted.
- Wes would like to include it in the dunesw stack.
- Would need a dune repo (and protected develop branch), a Spack recipe, a home in the SciSoft web server, and updates to some scripts that loop over dune products.

dunetrigger Progress

- We now have a centrally-managed repo in the DUNE org for dunetrigger.
- Managed by dune-software-managers, license added to README
- I touched up makefiles and a few bits of code to make the develop branch compile.
- There's a PR from James Shen, Simranjit Chhibra and David Drobner that introduces a Git submodule for DUNE-DAQ's trigger algorithms.
- <https://indico.fnal.gov/event/66867/>
- There are scripts as part of the PR that do some touchups to the submodule code so it can work offline.

dunetrigger Status

- Tom submitted tickets for mrb gitcheckout (and add to dune_suite) and also scisoft web site directories.
- Need to add to appropriate ups/product_deps to get dependency in there.
- Need to do when it's integrated: CI test config files, release version scripts and CVMFS publishing scripts. release wiki.
- Questions of versioning – what happens if trigger algorithms evolve in the DAQ? We may need to simulate old and new algs offline, on the same events. And compare new data taken with new triggers with old data taken with old triggers.
- Prefer managing our own copy in case we need to change the code or keep several versions. Build-time code tweakers make understanding what's in GitHub harder.

Spack Status

- Kyle has been helping Tom with trying out Spack MPD.
- A local environment needed to be built first before Kyle's examples could run. So far, so good.
- Tom ran into an error where no compilers were defined by default, when running the Spack MPD unit tests
- Error message was rather unhelpful, but investigation narrowed it down to having zero compilers defined. Kyle gave instructions for moving forwards
- Tom chatted with Marc Mengel – who wants to refresh the January tutorial. Marc says a new version of Spack has been released and is incorporated in Fermilab's stack, and testing is underway
- V says they are doing NOvA's spack releases.

PDSP use of IFDH

- `spdp::DetectorPropertiesProtoDUNEsp::UpdateReadoutWindowSize` calls `ifdh_ns::ifdh::getMetadata`
- Hangs on my desktop in `dunesw v09_91_04d00` which sets up `ifdhc v2_7`
- Setting up `dunesw v10_00_02d00` which sets up `ifdhc v2_7_2` lets that work.
- setting up `dunesw v09_91_04d00` and unsetting up `ifdhc`, setup `ifdhc v2_7_2` gives a sqlite error when running a program.

Updated November 2024

- Solved with addition of
`IFDH_BASE_URI="https://samdune.fnal.gov:8483/sam/dune/api"` and
`SAM_WEB_HOST="samdune.fnal.gov"`
in DUNE setup script in CVMFS

The DUNE setup script

- `/cvmfs/dune.opensciencegrid.org/products/dune/setup_dune.sh`
- It is very UPS-centered.
- Need another one for Spack
- Recent additions since the last meeting:
 - setting `UPS_OVERRIDE` to `"-H Linux64bit+3.10-2.17"` (needed for some operations in SL7 containers)
 - setting `IFDH_BASE_URI` and `SAM_WEB_HOST` as mentioned on previous page
 - setting `ALLINEA_USE_PSEUDO_TTY="no"` (needed to make arm forge ddt to work in a SL7 container)

What the DUNE Setup Script Does

- sets up UPS
- Checks your `.bash_profile`, `.bash_login`, `.profile`, `.shrc`, `.bashrc` for UPS setups of "login" or "shrc" and gives user a message (some of these products are veeery old and may not always exist).
- I found "setup login" can help with making manpages accessible, but I execute that by hand if a manpage is missing.
- Executes LArSoft setup script
- Adds the current working directory "." to `FW_SEARCH_PATH`
- sets `JOBSUB_GROUP=dune`, `EXPERIMENT=dune`, `SAM_EXPERIMENT=dune`

What the DUNE Setup Script Does

- defines some old art workbook variables (delete?)
- defines some 35t database variables (delete?)
- calls UPS setup for mrb, gdb, ninja, gitflow

- Some of these things need to exist in a Spack setup
- I have this alias:
`alias larspack2='source /cvmfs/larsoft.opensciencegrid.org/spack-v0.22.0-fermi/setup-env.sh'`
- This sets up Spack, but I still have to `spack env activate` and `spack load larsoft` to get stuff.
- Environment variables needed by a stack can be put in the recipes.
- Some may be of more general use (e.g. mpd)

Procedures for Updating setup script

- version-controlled copy is in duneutil/scripts
- DUNE software managers (Tingjun years ago, now me) get a request to change it, or we know of something needed
- I test a modified script in my own interactive session using a copy in my own disk space
- I put the tested version in duneutil/scripts and push
- I copy to /grid/fermiapp/products/dune/setup_dune.sh
- I publish on CVMFS and notify people

- Issue identified by Chris Green: Bash scripts on disk are the cache for running scripts – replacing a script that is currently being executed can cause inconsistency in results.

The DUNE setup script

- For the most part, it exits very quickly (though sometimes CVMFS access can slow it down). CVMFS race conditions.
- Updates are very infrequent (months or years pass between updates)
- But grid jobs run it. Small failure rates could show up in grid jobs.
- UPS's setup command makes an auxiliary script in /tmp and sources it.
- One could imagine the script copying itself to /tmp and sourcing itself. Assumes /tmp is defined.
- Possibly overkill for this application, but may be only a few extra lines at the beginning of the script.

CI Status

- CI tests were integral to updating to LArSoft v10
- Test coverage is not 100%. We could include more tests.
- Some integration tests are as easy as listing a fcl file and an input file in the appropriate test config
- Reference histograms were last updated Feb. 5, 2024
- Andrew and Tom tried updating them
- Vito helped – x509 proxy authentication is not working on AL9. Answer is to move to tokens
- Service Desk ticket submitted RITM2291872
- But there's an older one RITM1627792 where token roles are assigned – duneci is one of them. Conferring with Vito.

CI done and to-do

- Vito added Tom as a manager on the Redmine lar_ci page so he could add and subtract users.
- We still have some old DUNE software managers on there who have since moved on. And some new ones!
- We ought to disable most integration tests on Clang due to missing TensorFlow
- Add more test coverage.

Style Guide

- https://docs.google.com/document/d/1rRgGwGD0h5Z9XV0_R46M1RrHRimWPVjEeDMTp7KSJMk/edit?tab=t.0
- But it's functional rather than style oriented
- Links to other style guides in the appendix
- Running larsoft format-code script – probably a good idea!
- Might break some pending PR's.
- The DUNE-DAQ one is quite prescriptive. e.g.:

For DUNE DAQ software, if a variable name begins with a single letter followed by an underscore, that's meant to convey something about the variable's type (details below). For that reason, don't use this convention for any other purpose (e.g., use num_widgets instead of n_widgets).

Code Management and Release Management

- UPS and SL7 -> Spack and AL9
- SL7 Containers
- Establish Reviewers for Pull Requests
 - improve automated testing of PR (CMSbot)
- Continuous Integration development and maintainance
 - make the tests more comprehensive
 - integrate more code stacks and tools
- Write a style guide!
- algorithm integration to take advantage of new tools
- greater coordination on dev vs integration vs production releases with Physics Groups

Multiplatform Continuous Integration for LarCI

Build	Start Time	Build Type	setup_environment_cvmfs	checkout_github	build	unit_test	install	ci_tests
dune_ci/18144 (LArSoft DUNE)	2024-11-06 02:40:57	slf7 c14:prof	✓	✓	✓	✓	✓	✗
dune_ci/18143 (LArSoft DUNE)	2024-11-06 02:20:52	slf7 c14:debug	✓	✓	✓	✓	✓	✗
dune_ci/18142 (LArSoft DUNE)	2024-11-06 02:00:47	slf7 e26:prof	✓	✓	✓	✓	✓	⚠
dune_ci/18141 (LArSoft DUNE)	2024-11-06 01:40:57	slf7 e26:debug	✓	✓	✓	✓	✓	⚠

Progress Legend

- Pending
- Unknown
- Succeeded
- Warning
- Failed
- Skipped

"And the last thing is greater coordination of release management in terms of a development release vs an integration release vs a production release with the physics groups so that we know we will be marching forward so that when we can declare something a production release it will be stable. It might have patch fixes, bug fixes, but it's going to be a stable release for a given period of time so that everyone can work against that production release. Or be working in a dev release, but if you're doing physics, we want you in a production release. So, there is greater effort and coordination between those two that is needed."

dunesw release scheduling

- What we've been doing:
 - Follow the LArSoft weekly release schedule
 - New features in LArSoft need to be made available to DUNE collaborators
 - someone somewhere will ask for the latest LArSoft features
 - Often the driving factors are contributing products interfacing with the environment, such as ifdhc, jobsub, or the compiler
- Extra releases between the LArSoft releases are made on request.
 - "Hey Tom, can you merge dunereco PR#xyz and cut a release for us?" is typical.
 - That's what the d<nn> in the release tag (e.g. v10_01_00d00) is for.
- Patch releases are made on request. This is a bit rare, but it has happened
- We do not have a large, semi-uniform data sample that many collaborators are simultaneously analyzing. So the premium is on new features and not patching old releases.
- Old releases are always available. But operability with the environment isn't always tested. If someone needs a patch, let us know.

LArSoft Releases

- Lynn maintains this list

https://larsoft.github.io/LArSoftWiki/releases/LArSoft_release_list.html

- Production releases (used by any experiment for production) are highlighted in bold type
- Any currently available release is fine for production samples from a software coordination viewpoint. Physics groups must agree on content.
- Aaron Higuera has promised better communication about releases used in DUNE production.
- Older non-production releases sometimes get deleted from CVMFS.
- Build artifacts are available on the SciSoft web site for deleted releases.
- Andrew has a tool for making release notes automatically

Pros and Cons of Weekly Releases

- Pro:
 - Rapid deployment of new features. "We'll fix it in the next release" means you don't have to wait long.
 - External shocks are addressed quickly. New ifdhc? No problem! Include it in the next release.
 - Keeps the wheels greased and people's skills fresh
 - Keeps the backlog of updates down
- Con:
 - It's a lot of work
 - develop quickly evolves past users' test releases. Often, it's just the version tag. Check out and set up are different!