


FIFE Roadmap Workshop 2021

Agenda: <https://indico.fnal.gov/event/50150/overview>

Slides:  FIFE Roadmap Summer 2021

Management vision (Liz Sexton-Kennedy)

Q: What comes after GPUs? Shared memory architectures for example. GPUs will evolve but still be called GPUs. Still efforts going on to make hardware optimized for AI/linear algebra.

Neuromorphic chips? Several years before they are competitive though.

Q: What about RISC-V? No buzz in HPC circles right now.

For next-generation HPC machines, we need to optimize our software for these highly parallel architectures in order to get access.

Storage challenges for LHC and DUNE; capacity isn't so much the issue as the bandwidth is for the other neutrino and muon experiments.

Q: What's the decision for EL8 support? Final decision expected in Q4, joint with CERN. Not final, but will likely be a choice of one of the various repackagers.

Intro (Lisa)

Offsite running included in jobsub defaults; Singularity support also required offsite by default (expert option to override available).

Storage (Rafael)

New library coming

PNFS mounts on gpvms behave better in SL7 than they did in SL6, but that is not a green light to use at will. It is still better to access files in pnfs via other mechanisms such as xrootd or WebDAV whenever possible.

Q: Are you open to a fair-use tape algorithm that will give priority to an expt that hasn't used tape drives recently? Being considered, though it's much easier said than done.

Data management (Brandon, Lisa as speaker)

samweb is https-only now

Important note: when you do samweb prestige-dataset, it does NOT reset the LRU clock on files that are already on disk. See the slides for a way to reset the clock on files by reading only the first few bytes of each file via webDAV. However, don't use it as a backdoor way of keeping files in cache in perpetuity. It's more for the cases where you need to stage a large dataset, know you need the files soon, but there may be a bit of a delay between the beginning of the prestige part and actually running the jobs.

MetaCat (Igor)

Q: If user metadata is arbitrary and contained in json instead of native database objects, is it efficient to query on it? MetaCat uses a native json representation internally.

RCDS (Shreyas)

Usage has increased since it was made the default. Tarballs are generally available to workers within 5 minutes of submission.

Jobsub/Jobsub-lite (Kevin)

Candidate expected to be available for testing in the fall; initial target is DUNE global pool

Landscape(Kevin)

POMS (Marc)

Project-py (Vito, Lisa as speaker)

Several informative links in the slides for more information and tutorials.

END of the morning session. The next talk will start promptly at 12:30 CT.

Security

Note: There was a mixup about what slides were supposed to be in the main deck; we will get the correct slides inserted as soon as we can.

The capability sets requested will be configured in vault. The client just sees issuer and optionally role, for example “-i dune -r production”.

Authenticate one time only or once per month/year? That was the old problem (for instance on certs for DocDB) of having to re-authorize but so infrequently that it was problematic. A: Re-authenticate only if you have no qualifying activity for one month.

Q: Can we get a token in a cron job, using a kcron principle? kcron can be used to talk to Vault, yes.

Q: Can service providers create an initial list of existing structures and access privileges within dCache?

Some concern about making the initial authorization/release step too complicated. If it is simply a case of htgettext spitting out a URL that people can cut and paste into the browser of their choice, it's probably OK, but if it calls xdg_open and tries to launch the browser on a gpvm, for example, they might get a second browser asking them to do the same thing that they already did, leading to confusion. Should figure out what services should disable that part of things. Also, some detailed documentation about how to go through the initial authorization (or re-authorization when coming back after inactivity) would be very important for the infrequent/less savvy user.

For now, for all vo users it is read everything from /vo-name for analysis users to write /vo-name/scratch/users/user-name, for production anywhere write under /vo-name

Q: are the transition timelines (1+ year) sufficiently aligned with OSG/WLCG? They seem to be moving a little faster.

The host side of things (host/service certs) will not change.

Elastic Analysis Facility (Burt, Lisa as speaker)

Spack (Marc)

Just listing myself (Robert Hatcher) as super invested in the eventuality of NULL flavor spack packages - I have 3 similar use cases in UPS that use this ability: `genie_xsec`, `genie_phyopt`, `genie_fluxopt`, as a means of distributing configuration files. The first is big enough that one really doesn't want the duplication. None of these are oriented at the user building the package but rather distributing pre-built versions and setting up the environment for the user.

Continuous Integration (Vito, Lisa as speaker)

Q: Is the chi2 test by itself making the experiments happy, as opposed to e.g. a KS test. CMS allows tests to pick their own

Wilson Cluster

User Activity Mitigation Policy (Ken H.)

Q: How many times have you emailed people and the result has not been a ban? A: O(dozens). It probably averages about one per week. It does seem to help a little bit in that most people who go through the process and improve their workflows usually don't show up again.

Wrap-up (Lisa)