


Mu2e-doc-4269-v2



Mu2e and Muon g-2: FIFE Plans and Issues


Rob Kutschke and Adam Lyon
FIFE Workshop
June 15, 2013



Mu2e and Muon $g-2$




- I am speaking for Mu2e
- Presenting some slides on behalf of Adam Lyon for Muon $g-2$.



Most Important Slide




- Thanks to the FIFE team and related service providers for all of their hard work.
 - We have used lots of resources.
 - We have received lots of help and advice, some of it on a rush basis.
 - In most cases we have very good communication.
- We very much appreciate all of it.
- There is a lot more work to do ...



Mu2e Timeline




- We are in the P5 plan in all scenarios!
 - Some stretch-out possible
- CD-2/3 Director's Readiness Review: July 8-11
 - Internal deadline this week!
- CD-2/3 DOE Review: August 19-21
 - No disruptive changes until after the DOE review
 - CD 3 only for selected subsystems
 - Other subsystems: CD-3 mini review over ~1 year
 - Continue heavy simulations through this period
- Projected construction start this fall
- Projected engineering runs ~2019



Grid Usage - 1




- GPGrid and CDFgrid
 - Average ~1000 slots active at any time
 - Peaks to > 4000 slots
 - Likely to continue for at least a year (but see next page ...)
- Two major classes of jobs
 - CPU heavy, little IO
 - IO dominated
 - Dominated by “wait for lock” to access bluearc
 - Stage in O(10 GB) and compute for only 1 or 2 hours
 - Often hundreds of processes waiting to stage in files
 - **Expect that dcache will help a lot – testing now**
- Expect fraction of IO dominated jobs to increase



Grid Usage - 2




- A new need has arisen.
- Known weak points in cosmic ray veto coverage
 - Over the river; off the expressway; through the window; off the scoreboard; nothin' but net.
- Need to simulate at least 10 x the live time of the experiment with targeted simulations
 - ~4,000,000 cpu hours over 6 months+
 - (= 1000 cores for ~5.6 months)
- Learning to use offsite OSG will be important
 - We will initiate this in the next weeks



Grid Issues




- Intermittent problem: jobs that die for reasons we do not understand – $O(1\%)$.
 - Hundreds of processes in one cluster will complete successfully and a few will die seconds after starting.
 - Current ticket: [INC000000418734](#)
- Mu2e processes sit idle when there are thousands of free slots on both CDF Grid and GP Grid
- Wait-for-lock issue from page 5 – hopefully dcache and/or xrootd is enough mitigation.
- Have not yet tried client-server jobsub.



FermiCloud




- One important use case: making pre-mixed background overlays
 - IO limited 1000 x [O(8 GB) in O(8 GB) out]
 - Enough additional degrees of freedom to safely oversample many times.
 - **Needs a machine with about 8 GB peak virtual size**
 - Because of IO limitations we can only use a few such machines at one time.
 - Fermicloud has been a great solution
- This will be an intermittent but ongoing need.



Disk - 1




- Bluearc: 82 TB data, 1.4 TB app - full
- Recently started to use dcache scratch
 - SCD response was fast
 - Some glitches but we are working through them
 - Some feature requests in the ifdh queue
- We are gambling that we can use dcache scratch as a safe place to stash files for a few more weeks until we have our SAM instance configured.
- Migrating our high IO jobs to stage-in large files from dcache, not from bluearc.
 - Will also test drive at xrootd.



Disk – 2




- We renew our request for individual bluearc quotas
- For now, only a group quota
- We would like to be able to control individual quotas ourselves, ie without a service desk ticket
- Under normal operating conditions we have a few spare TB of bluearc disk space
 - Not all should be written to tape – in fact most of it should not.
 - It's easy for someone who is not paying attention to use it all and cause everyone else's work to fail.
 - Expect this to get worse as more people join
 - dcache can help but it is not really the complete answer



SAM




- We have had the initial meetings to set up SAM
- On our plate:
 - Decide on file families, meta-data and policy for what gets written to tape.
 - Ray Culbertson has prepared a draft plan
 - Expect to make decisions $O(\leq 1 \text{ week})$
- I understand that we can be up and running $O(1 \text{ week})$ after that.
- Decision was made a few weeks ago that new files will be written to T10000 tapes.
- Then we will need to adapt scripts $O(1 \text{ week})$ including debugging in the field.



Software Environment




- Relocatable UPS
 - Our own instance plus products/common
 - Everything except our own code.
- Code management
 - cvs on cdcvs
 - plan to change to git on redmine
- Build system: scons (no plans to change soon)
- Still use expwww and are ramping up redmine wiki.
- Use our own G4 interface (do not use artG4)
- SLF5/6 only.
 - Want to test-drive Mac OS when we have the people-cycles



art




- Overall very happy with art.
- About bug-fix and feature requests.
 - Things that are both important and easy to do are usually dealt with quickly.
 - In recent months NOvA has had priority so other work has been delayed (understandably so)
 - There is a large and growing backlog of requests, including many Mu2e requests.
 - The client base is growing
 - The art team needs to grow or it will not be able to keep up.



jobsub




- We have our own submit scripts that live on top of jobsub
 - mu2eart
 - mu2eg4bl
 - mu2emars
- So far only use GP and CDF grids
- Good response on jobsub feature requests
- Have not yet exercised client server jobsub. Will do so as soon as well can post TDR.



ifdh




- We use it and in general works well.
- Periodic issues and good response.
- We have an issue open now regarding stage-in from mixed sources (both bluearc and dcache).



Conditions Databases




- Three needs:
 - Fall 2014: DBs for travellers for construction project
 - DB's will be $O(\text{a few GB?})$
 - ~2015-2016?: conditions DB for prototyping the conditions system and learning to do alignment
 - Guess that scale is $O(100 \text{ GB?})$
 - ??: Test beam conditions data
 - Guess that the scale is $< O(1 \text{ GB})$
- GUI's for data entry for the traveller DBs.
 - Historically done by postdocs/grad students (aka free labor) on project but Mu2e is too far from data to have many of these
 - **Is this available within FIFE? Elsewhere in SCD?**



Other




- OASIS/CVMFS
 - We know that we need it to enable offsite running.
 - On our list so deploy soon.
- ELOG
 - We have an instance but it is not heavily used.
 - Expect use to grow as testbeam starts
 - Not clear yet if construction projects will use it



Not Sure If this is the Right Venue ...




- The Mu2e collaboration asked for a discussion forum technology ~2 years ago.
 - I understand that a prototype is being designed in Sharepoint 2013 but I have not yet heard when it might be available to test.
- Our outside collaborators have urged me to talk about this in the strongest possible terms – they are already telling their colleagues that FNAL CS just does not care.



FIFE & Muon $g-2$


Adam Lyon for the Muon $g-2$
collaboration



What we use now




- We're a big user, supporter, and cheerleader for the art framework
 - Simulation and offline code is under development in art
 - Our ArtG4 is now supported by the SCD Simulation group
- Our development environment has become “mrb”, in use by LArSoft and others
- We've had new users go through the art workbook and they like it



What we use now




- We rely on Redmine and git (writing the conceptual and technical design reports were managed by git – even the spokespersons learned git [though not happily])
- We use CMake for builds (same as art itself)
- We deliver applications/libraries with CVMFS (including to our mac laptops; Cornell U and U Washington groups use OASIS to get releases on their interactive nodes)



A glaring omission



- Several collaborators want to use their Windows Laptops (including a spokesperson)
 - The lab basically does not support kerberos on non-lab owned windows laptops. There is no good free kerberos Windows solution
 - Installing and using kerberos via Cygwin is painful
 - Some people don't appreciate being told to buy a Mac (but of course that's the right thing to do 😊)
- Our solution:
 - Windows user installs VirtualBox (free virtualization system - easy)
 - Download and install our specially prepared SLF6 Vbox Image file
 - This gives the user a virtual SLF6 machine that comes with Kerberos, ssh, git, LaTeX, X-windows, and CVMFS Oasis automounted – all installed and instantly ready to go. Files can be shared with the Windows side.
 - This solution has worked nicely, but wish we didn't have to do this




Plans for this summer



- Adopt SAM/IFDH/dCache
 - Right now we're still using Bluearc via jobsub's nice copyback mechanism (uses IFDH)
 - Plan to put important simulations on tape and use the scratch cache
 - SAM is already set up and is waiting for us
- Try new JobSub
- Try Build Facility


Going slowly because people are busy with work on TDR/reviews.



Plans




- By fall we'll be a fully on-boarded OSG experiment using SAM, dCache, etc
- Thanks for everyone's hard work and help!



Backup Slides






Computing Projections



- Continue current CPU and IO load for about another year to finalize the design
 - Background rates as we modify details of shielding.
 - Shielding for system A can increase rates in system B!
 - Weak points in cosmic ray veto system.
 - Improve robustness of reco algorithms in higher background scenarios.
 - Calibrate-out some some effects that are modeled in the simulation but not yet in reco



Mu2e Background Model



- 1 Event = 3.1×10^7 protons on target
- ~3,000 hits in the tracker during the live gate
- Simulate one proton at a time
 - Done 6 stages: 4 simulation + premix + reco/analysis
 - The last 2 are IO heavy
 - Premixing requires 8 GB peak virtual size
 - Some stages resample the output of previous stage many times
 - Lots of knobs and dials to increase the effective statistics so that we can have effective statistics of many millions of events.
 - Form samples of pre-mixed background to overlay on signal events