Mu2e Offline & Computing Review March 2015

Charge

The Mu2e experiment recently completed a successful CD-2/3b review (Oct. 2014). The collaboration is now working towards the next step in the approval process – a CD-3c review in late 2015. This requires significant simulation and analysis campaigns. In addition, the experiment needs to plan for successful physics analysis once data taking begins. The Mu2e spokespersons and SCD management would like the committee to review and comment on:

1. The current offline computing infrastructure and tools, including build and release tools, simulation tools, framework, database, workflow, workflow management, data management, and operations. Are the tools, infrastructure, and established processes appropriate for this stage in the experiment? Are the tools, infrastructure, and established processes sufficient to engage non-expert resources from the collaboration? Are best practices being employed as appropriate? Is the experiment appropriately leveraging tools and expertise provided by SCD?
2. The requirements for the simulation campaign leading to the late 2015 CD review and the plan for meeting these requirements. Is the plan reasonable and achievable? Are the current infrastructure and tools capable of meeting these requirements in a timely manner? If not, what additional infrastructure and tools are required? Have adequate personnel resources, both from the experiment and SCD, been identified? If not, where are the personnel shortfalls?
3. The requirements for the simulation campaign leading to the late 2015 CD review and the plan for meeting these requirements. Is the plan reasonable and achievable? Are the current infrastructure and tools capable of meeting these requirements in a timely manner? If not, what additional infrastructure and tools are required? Have adequate personnel resources, both from the experiment and SCD, been identified? If not, where are the personnel shortfalls?

Charge Question 1

The current offline computing infrastructure and tools, including build and release tools, simulation tools, framework, database, workflow, workflow management, data management, and operations. Are the tools, infrastructure, and established processes appropriate for this stage in the experiment? Are the tools, infrastructure, and established processes sufficient to engage non-expert resources from the collaboration? Are best practices being employed as appropriate? Is the experiment appropriately leveraging tools and expertise provided by SCD?

Overview

The current setup of the Mu2e infrastructure and tools are appropriate for the current construction stage of the experiment - preparing for the next steps in the DOE project management process, with data-taking a few years down the road. There are 30 people working part-time on the Mu2e software and computing effort, totaling 14.5 FTE, which also is appropriate for this stage. The experiment has done an excellent job in incorporating the tools provided by Scientific Computing Division. They have integrated art at a high level and the collaboration is using it in a productive fashion.They are planning to migrate most of their infrastructure to FIFE tools in the near term; to that end, there is good bi-directional communication with the experiment and the FIFE toolkit maintainers.

The experiment has asked for additional SCD support in various areas; the committee recognizes that this support is more likely to be realized if specific needs are expressed by other experiments.

The committee suggests to SCD that specific time be allocated in the CS Liaison process to facilitate finding commonalities between experiments, and to leverage the SC Portfolio Management Team to keep senior SCD management informed and involved.

Event Display

Mu2e requested that SCD provide a graphics/user interface expert to work along with experiment scientists to work on event displays.

Findings

The experiment has multiple event displays that they have developed, and demonstrated them for the review committee. The displays that were presented are fully featured. They contain detailed knowledge of the detector and data from the reconstruction process. They seem to match the current needs of the experiment. When asked for further justification for the effort, the experiment noted that professionally-developed event displays improve experimental performance, that the right event display would be a powerful tool to investigate tracking algorithms as well as reconstruction in the calorimeter and cosmic ray veto system, and that an expert would be able to maximize performance and improve the user interface. The experiment also mentioned that the people who work on the event displays have physics responsibilities which limits their ability to maintain and extend them.

Comments

The committee was unconvinced that general SCD expertise in the area of event displays would meet the experimental needs. From experience, members of the review

committee know that event display requirements and deliverables tend to be highly specialized and require experiment-specific knowledge and expertise. Effort for a graphics/UI expert from within SCD seems unlikely, unless there is sharing with other experiments involved.

Recommendations

To Mu2e:

The committee recommends that Mu2e find commonalities with other experiments in event display and development needs. Experiments may make specific requests for assistance to SCPMT and those requests will be prioritized among the other division tasks.

G4Beamline

The experiment is actively using the G4Beamline software product. This product is currently supported by Muons, Inc.

Findings

The experiment is using G4Beamline, a framework independent from the Mu2e-offline framework. It is being used as a script wrapper to define and manipulate Geant4 geometries, which then allows users to run simulations with very limited knowledge of the underlying C++ language.

The current experts are very familiar with the product, and it would take significant effort to migrate to Mu2e-offline; if quick shielding studies are presently needed, G4beamline would be used. With the ramp-down of Muons, Inc., the support is unclear going forward. The experiment requested that SCD step in to support this software.

Comments

The G4Beamline tool is only needed for a limited time and scope. When asked, the experiment indicated that G4Beamline will not be needed after the shielding design is finalized; however, this may take a few years to complete the optimization project.

Recommendations

To Mu2e:

The committee encourages Mu2e to continue the efforts to migrate away from G4Beamline, recognizing that Mu2e uses it to only providing a simpler interface to the underlying Geant4 software package.

To SCD:

The committee recommends to SCD to monitor the general use of G4Beamline, as it is in use by other collaborations (e.g. g-2), and that other experiments are using additional features that may require more effort to replace. The accelerator division also uses G4BeamLine heavily and perhaps they could be a resource for assistance.

Code development and software distribution

Findings

There are roughly 10 active developers contributing code to the Mu2e experiment. The experiment expects more, but no growth estimate is available and none was given. A large git repository in the FNAL redmine system holds the experiment code. They are using a simple git-based development workflow. The developers clone the main repo, commit code only in their local repository, and then privately “rebase and merge” code into the master branch; at that time, they push their changes into a central repository. There is one remote repository (from BaBar) overlaid on top of the central one for a complete build. There is no clear process for integrating contributions into official analysis code -- the experiment relies on the coordination between 2-3 core contributors to do the work and assess its validity. They do maintain stable tags in the repository.

The experiment uses relocatable UPS tarballs available from the standard scisoft.fnal.gov distribution server for non-experiment code, but does not provide binary installable package for the Mu2e software (no usable Mu2e UPS product). They have identified this as a limiting factor in the long run. For experiment software, developers and analyzers download the source code and build locally, which can take up to 30 minutes for the first build (and generally a few minutes for subsequent builds). Mu2e said that compilation time is not the limiting factor. Rather, the space for developer builds is becoming a limitation on FNAL GPVMs. The experiment is using SCONS as the build tool, and they believe this tool adequately meets their needs. They reported regular builds off-site at LBNL. The SCD continuous integration system is used to ensure that the code in the master branch can compile and link successfully. Gdb and Valgrind are used to aid in debugging code.

The next planned steps in development that will soon be underway are code/build cleanups (including targeted code reviews), and additional validation and testing. Plans were shown for splitting the large repository into smaller pieces to minimally cover areas of reconstruction and analysis. The experiment has already prototyped a base-release / test environment that links a development build (for an individual) and a pre-built managed release area.

Comments

While this software development model is satisfactory given the construction phase of the experiment and for the near-term computing campaign for CD-3c, there is concern that this workflow will not scale to more users and developers. The expectation is that there will be on the order of 200 collaborators, with 30-40 developers by FY20. Pushing changes to a single master branch ensures that code conflicts will become part of the normal development

process, and has the potential to destabilize the experimental software. There is also currently no release management, which would declare which software tags are production, describe their features to the users, document them appropriately, and so on.

The ongoing work to split the repository into “common”, “offline”, and “online” repositories and the development of a better incremental / test release build environment is very important for productivity and will make it easier for new collaborators to contribute to the software projects. The experiment developers made it clear that they are aware that problems can come about due to multiple build areas (test / base releases) and parallel build tools. Other experiments and projects are gaining experience with tying together several separately-versioned (and installed) packages across repositories. Tools such as MRB are starting to mature and their design may help in defining the build / release as the development structure becomes more complex and additional users are added.

Recommendations

To Mu2e:

The committee recommends that the experiment prepare itself for the expansion in its community, mainly after CD-3c is reached. In particular, a more scalable and robust software development model should be investigated, including official code librarians, comprehensive code reviews, and a more advanced git workflow (e.g. feature branches and the “git flow” set of tools). The experiment should adopt a binary distribution model, leveraging CVMFS and a clear release management process that includes declaration and documentation of both base and test releases. The experiment should be aware that SCD does not support the SCONS tool, and that Mac OS X is not supported in a production environment.

Engaging non-expert users

Findings

The engagement of non-experts was directly called out as a priority during the roadmap summary presentation. Mu2e has already completed some unique work in this area. They produced (via Rob Kutschke) the “art workbook” to introduce new users to development using art. This is proving to be an excellent contribution to the entire art community. The Mu2e software suite has been developed using C++ and art framework best practices wherever possible. This was made possible through a small group of dedicated and experienced physicist-developers within Mu2e. The current offline software package requires intermediate to advanced software development skills. This is not only due to the complexities of using a large full-featured framework, but also because of the quickly evolving nature of the software in this early phase of the experiment from a handful of experts. Other contributing factors have already been discussed in this report (build/release procedures and some missing end-user reconstruction products). A few noticeable effects of skill level can already be seen even in these somewhat early stages of the experiment. An example is the CDF “stdntuple” package that has been brought forward and attached directly to the repository as an aid in

some user's analyses. The experiment explicitly stated that a standard ntuple format will be designed and chosen that is acceptable by the entire collaboration and that there is a desire to not have one appear from individual wants and desires. The Mu2e team has sketched out a plan that alleviates many of the current and upcoming difficulties. By December 2015, the base/satellite builds are expected to be in place, along with a decision for how to move ahead with a standard analysis n-tuple. The execution of the plans for engaging non-experts also begins in December 2015.

Comments

The presented plans include many "usability" tasks that will help non-experts in the use of the Mu2e-offline software. Many of the important usability tasks appear to be well underway or complete by the July 2016 milestone. By this time, the restructuring of the repository, standard packaging of Mu2e software, and builds using the base/satellite system ought to be complete and in common use. The December 2015 milestone states that Mu2e will "start to execute plan for engagement of non-experts." This plan should include starting work on the ntuple format and the production of it, the code housekeeping and cleaning tasks, moving towards better code management practices, and providing the missing persistent and summary track data product. To quickly ramp up new users, Mu2e wants to consider providing them with standard ntuples. This has some risk, as the training of users in the basics of the Mu2e-offline framework could be forgotten. This timeline for advancing the Mu2e-offline framework is fairly aggressive given the workload of the collaboration software experts.

The art team has been planning a series of usability features that may also help Mu2e engage non-experts.

Recommendations

To SCD:

Continue to support the art workbook as a tool to help Mu2e and the community quickly learn to use central processing frameworks and the C++ programming language. Support workshops and tutorials that will help non-experts come up to speed quickly.

Support feature development within art and the surrounding tools that make applications and tools more usable for non-expert experiment collaborators.

Ensure that a SCD framework developer is available for regular consultation with Mu2e, to help make sure that the underlying infrastructure software is well-aligned with the needs of the experiment.

To Mu2e:

We recommend that the experiment follow the proposed plans and keep the engagement of non-experts as a priority.

Participate in any upcoming art workshops or courses that will help new users learn about processing framework and quickly become productive.

Provide input and track progress on art usability features that will help engage non-experts.

Three geometries

Findings

Mu2e-offline, G4Beamline and MARS are all used by the Mu2e collaboration for different purposes. Each of the packages have their own geometry implementation, which has to be maintained and changes have to be synchronized.

G4Beamline's geometry parsing code is not used in Mu2e-offline, because some of the geometry description features that Mu2e-offline needs are missing in G4Beamline's geometry parsing code.

The geometry as a whole is not expected to change significantly before the start of data taking. However, optimizations and tweaks will be necessary -- for example to optimize the detector shielding.

Comments

The synchronization of changes to the geometry is time-intensive but necessary. Mu2e-offline cannot easily stop using G4Beamline, for example (see comment in G4Beamline section above). Mu2e mentioned that a tool to automatically synchronize the three geometries would need 6 FTE-months of an expert to write.

Charge question 2

The requirements for the simulation campaign leading to the late 2015 CD review and the plan for meeting these requirements. Is the plan reasonable and achievable? Are the current infrastructure and tools capable of meeting these requirements in a timely manner? If not, what additional infrastructure and tools are required? Have adequate personnel resources, both from the experiment and SCD, been identified? If not, where are the personnel shortfalls?

Findings

Mu2e presented the plan for the simulation campaign leading up to the CD-3c review. The review will take place March 2016. Mu2e's plan foresees that all computing activities including analysis of the produced simulation samples are completed by December 2015. Mu2e estimates to need 14M CPU-hours, and needs a very significant amount of opportunistic resources both inside and outside of Fermilab with the current Mu2e allocation. To keep the schedule, the campaign needs to start production on April 1st 2015. The schedule and the CD-3c review is at risk if less than 70% of the planned work can be completed in the presented time frame.

In general, the experiment believes that the offline code is ready for the challenge. They have made available 3 FTE until April 1st to develop some specific simulation code. The highest priority are geometry updates and changes to data products. Many of the other tasks are reconstruction or analysis phase projects and can be deferred until simulation production has started.

On the computing infrastructure side, the Mu2egrid script suite needs to be updated for SAM and dCache usage. The commissioning of CVMFS and OSG running needs to be finished, which includes load testing. These will start immediately after the review.

The analysis component of the CD-3c production campaign needs to be completed by early December 2015 and consists of about 15 topics that will be investigated by 15 Mu2e collaborators. The analysis jobs will read 0-40 TB of pre-mixed datasets about 200 times, with single step jobs producing small output.

Mu2e describes the plan as ambitious and presented FTE estimates. For the weeks after the review until April 1st, Mu2e estimates that 2.65 FTE weeks from Andrei Gaponenko and Ray Culbertson will be sufficient to accomplish all tasks. About 9 FTE weeks are planned from other collaborators. Mu2e estimates the need of 1.35 FTE weeks of effort from SCD for consulting and help. During the campaign, Mu2e estimates the need of 0.2 FTE from collaborators, 0.5 FTE from the FNAL operations team and 0.05 FTE from SCD for consulting and help.

Comments

In the committee's opinion, the presented plan taking into account the presented effort estimates is extremely optimistic, based on experience from other experiments with the same scale of operations and problems. There is a lot of work to do in a short amount of time. Directory trees and text file lists are being used as metadata for experimental bookkeeping -- the experiment needs to integrate SAM for this purpose. The use of the POSIX-compliant BlueArc networked filesystem must be replaced with the use of dCache. If on-site resources are insufficient for the planned opportunistic use, the experiment must be able to exploit OSG cycles. This new workflow needs to be debugged, integrated, hardened, and load-tested.

It might be advisable to have a more detailed exchange with the storage group to optimize file families, data tiers and to improve file grouping on SFA. One suggestion is to use IFDH_STAGE_VIA for storing log files on Bluearc.

Judging from the magnitude of tasks to accomplish, the committee expects that the time schedule for completion will slip into the April 1st production window.

Recommendations

To Mu2e:

We recommend that Mu2e develop and adopt a phased approach for the migration of the existing workflow tools to the planned new implementations and for the transition to large

scale operations. We recommend a similar short-term milestone-based approach as was presented for the roadmap to operations on a longer time scale.

We recommend that Mu2e modify the request for operations support to SCD from FTE-based to task-based.

Because of the similarities to NOvA's production tasks, we recommend that Mu2e closely collaborate with NOvA and reuses as much as possible of their production model.

We recommend the selection of a single person responsible for the CD-3c production campaign deliverables to keep an overview and guarantee coherence of the effort.

To SCD:

We recommend the identification of SCD personnel to participate in close and regular interaction with experiment people (i.e. a "code sprint") to finalize the implementation of the production script setup including dCache/PNFS and SAM integration.

We recommend increasing the OSG integration and operation support immediately for Mu2e, following the example of NOvA, with very close and regular interaction of experiment and SCD personnel (i.e. an "OSG sprint"). We would first concentrate on provisioning resources from collaborating institutes of Mu2e, and then branch out.

We recommend that SCD include the needs of Mu2e for operations group support in their planning for the year 2015.

We recommend defining a dedicated contact in SC Facilities to follow the code and OSG sprints to ensure optimal use of the FNAL facilities, especially for storage interactions.

Charge question 3

The plan for developing the offline computing infrastructure, tools, and processes (as defined in (1) above), leading to Mu2e data taking and beyond. Is the plan reasonable and achievable? Does the plan appropriately leverage tools and services provided by SCD? Are best practices employed? Do the required personnel resources, both from the experiment and SCD, seem reasonable? Will the plan result in tools, infrastructure, and processes that are capable of producing "analysis ready" data in a timely manner and allow for significant engagement of non-expert resources from the collaboration?

Overall plan

Findings

Mu2e presented their milestone plan leading up to the start of data taking of the experiment. The CD-3c review milestone in March 2016 is preceded by the completion of the simulation production campaign and subsequent analysis or produced samples by December 2015. The overall schedule then foresees yearly milestones starting in 2017. Q4 of FY20 will mark the start of data taking with a cosmic ray test with the complete detector outside of the magnet (in the so-called "garage" position).

Mu2e continues to fully utilize the SCD toolkit and anticipates no notable future conflicts. Mu2e is planning to follow best practices in consultation with SCD and will implement internal review and planning processes to accommodate a growing collaboration. Especially the analysis workflows and user interaction are a priority and will be discussed and planned by the whole collaboration. Mu2e acknowledges the need for a standard ntuple data format for analysis. They have a plan to proactively address it with involvement of the whole collaboration.

Mu2e is well positioned for the coming years in terms of simulation production setup. Although complex, Mu2e designed a workflow with many intermediate output stages that facilitate rapid development and efficient reuse of previous steps.

Mu2e presented a comprehensive list of housekeeping tasks and plans to work on them over time. The tasks are roughly prioritized.

Mu2e presented a first plan how to evolve the software and computing management structure which will be further discussed and refined over time in the collaboration.

Comments

Overall, the presented plans seem reasonable and appropriate for the different stages of the experiment. The committee thinks that Mu2e is on a well-organized and solid path towards start of data taking. The collaboration is to be commended for its excellent integration and utilization of the SCD FIFE toolkit.

In more detail, the plan for code and release management foresees a significant amount of effort which seems reasonable for the later stages of the experiment. Concerns were raised that this level of effort is too high for the experiment to commit in the earlier stages of development.

Recommendations

To Mu2e:

Due to the many necessary and dependent steps in their production workflows, we recommend that Mu2e be actively involved in the design of the production workload/workflow management system that is currently planned as part of the FIFE toolkit.

To SCD:

We recommend that SCD continues to work on a plan to provide a workload/workflow management solution for the FIFE toolkit, and to engage with the Mu2e experiment to incorporate their requirements.

Questions to Mu2e (submitted on day 1 at 16:50):

- Mu2e requested event display development support from SCD. Please provide more details and concrete justification.
- Give an indication of how many people and how many FTE are behind the management boxes. Where are SCD people involved?
- CD3 MC production is planned to start on April 1st; many of the ingredients are new compared to the former TDR production. Please present:
    - timeline and scope for development tasks that still have to be done
    - timeline for integration tests and scale tests that still have to be performed
    - timeline and plan for operating the infrastructure during the 5 months of the campaign
    - These should include manpower estimates from Mu2e and SCD for the items listed above.
- Please include manpower estimates from Mu2e and SCD for the commissioning and operation of production on the Open Science Grid.
- Explain in detail the expectations and timeline for analyses based on CD3 MC production, including workflows and number of users.
- What is the impact on the CD3 deliverables if Mu2e only achieves half the CPU hours that are required? What is the contingency plan in case of a deficit?