

Final Report - First Draft

Muon g-2 Computing Review

1. Introduction

At the request of the spokespeople of the Muon g-2 experiment and Scientific Computing Division Head Panagiotis Spentzouris, members of the experiment and members of SCD met recently to review the readiness of the collaboration for commissioning the experiment in the areas of online and offline computing. The charge is provided in Appendix A, and the membership of the committee is given in Appendix B. The review was conducted November 7-8, 2016 and the materials presented are available on the meeting agenda at this URL, <https://indico.fnal.gov/conferenceDisplay.py?confid=13276>. After the review, the committee met to discuss the review and presented the experiment with requests for additional information. This additional material was presented to the committee and incorporated into the review process.

This report presents the findings, comments, and recommendations of the committee based upon the committee's understanding of the material presented at the review, additional material provided immediately following the review, and discussions between the committee members. A draft of the report will be presented to the collaboration for review and comment, and any additional material or changes will be incorporated. The report summary and recommendations present the consensus opinion of the entire committee.

Each section below addresses one of the three questions from the charge and lists the findings, comments, and recommendations of the committee. All areas of the charge were reviewed by all members of the committee and contributions were collected from all.

2. Charge question #1

The current offline computing infrastructure and tools, including build and release tools, simulation tools, framework, analysis tools and approaches, database, workflow, workflow management, data management, and operations. Is the experiment efficiently leveraging tools

and expertise offered by SCD? Does the experiment have sufficient resources to implement, deploy and operate the infrastructure?

2.1 Findings

- Yes. In general, the experiment is efficiently leveraging or planning to leverage tools and expertise offered by SCD. The areas where this is lacking are called out below.
- No. The committee finds that the experiment is stretched critically thin for effort to implement, deploy, and operate all of the necessary offline infrastructure. There are specific recommendations below to mitigate this situation to assure the success experiment.
- The g-2 collaboration as an organization has defined a management structure that has identified management roles for their computing and analysis tasks/groups.
- The omega_a and B-field analysis paths are almost entirely decoupled in terms of infrastructure and available effort. This separation is by design to ensure blindness of data.
- The build-and-release tools, simulation tools, framework, analysis tools and approaches are in place for the omega_a analysis path(s); the corresponding tools and approaches are less developed for the B-field analysis path.
- A tight schedule is defined to bring the the B-field analysis up to speed and ready to perform a full measurement by the end of March 2017.
- *See response to charge question no. 2 regarding databases.*
- A data-management system is in place, but it has not yet been exercised at a scale comparable to that expected during data-taking.
- There is not currently a TSW in place with SCD.
- The g-2 simulation team recognizes its need for a release manager and is taking steps to address that need.
- The g-2 simulation team is cognizant of the benefits of unit-testing algorithms. Such recognition is uncommon among particle-physics experiments and is to be commended.
- Large portions of the storage ring are modeled using CADMesh. The simulation team is to be commended for pursuing an innovative approach handling detailed geometric structures, which can be clumsy to manage using standard G4 C++ interface.
- The review committee is uncertain that the experiment's resources are adequate to complete the Offline implementation and workflow and database infrastructure given (a) the commissioning needs of the online DAQ and detector systems and (b) the (potential) overlap of students that contribute to both software and hardware.
- The Offline simulation group would like to improve their software management practices by switching to github instead of redmine, following the pull-request and sparse-checkout model that CMSSW uses.
- The Offline simulation group regularly interacts with the *art* development team, requesting advice in terms of implementation and design.

- The collaboration acknowledges their tracking developments are behind schedule, and that key expertise in the area of track fitting is missing. They advocate for external help quantifiable in 3 FTE months from a tracking expert to focus on track fitting.
- The development of the verification package for the simulation physics validation is well advanced, and appears to be appropriate to serve the collaboration needs.
- The beam simulation studies demonstrate that storage ring entrance parameters match the desired criteria both in terms of polarization (>95%), in terms of momentum (within 0.5% of magic momentum), and in terms of robustness against magnet displacement.
- The Delivery Ring introduces spin momentum correlations which intensify with the number of turns. They estimate that it contributes to ~10-20 ppb error but this needs to be verified with more simulations inside the storage ring
- The beam simulation software runs at NERSC (Edison), and it is not clear how many nodes/threads are needed and not clear how easy it is for the collaboration to get access to those resources. It is also unclear whether the code requires on highly parallel hardware, or if HTC facilities would be able to run the simulation.
- The beamline simulation is used to optimize and tune beam parameters in order to achieve the optimal muon yield. The beamline simulation is also used to study spin-momentum correlations, which form a potentially important systematic uncertainty. It is not clear how many full production runs of the beamline simulation are required.
- Multiple beamline simulation codes have been compared to each other for predictions, but relatively little comparison has been done to relevant external data for validation.
- The beamline simulation does not currently include magnetic fringe fields.
- Automated workflows are not yet in place. A workflow coordinator will be appointed in early December. Some of the individual components (JobSub, SAM, IFDH) have been tested at small scale, using simple scripts. SAM services are not yet integrated into the current workflow. POMS is being investigated.
- An Offline production team is not in place. Production team members will be solicited from the collaboration at the next g-2 collaboration meeting in December.
- Reconstruction for calorimeters is in good shape, with clustering performance and energy resolution (3%) well within expectations for first data taking.
- One full copy of the raw, compressed data is estimated to require 10 petabytes of space on disk. The collaboration currently plans on storing one copy of the raw data.
- The collaboration has a prioritized request for a senior scientist to work with graduate students and postdocs on simulation.

2.2 Comments

- Writing a Technical Statement of Work (TSW) with SCD will help the collaboration to focus its computing efforts.
- The Offline software efforts should continue to interact with the *art* team, informing them of their desires and needs.

- When the make-study tool is available for testing, a member of the g-2 collaboration developing algorithms should test the facility and provide feedback to the *art* team.
- Due to the urgency of other computing needs, the committee suggests that low-impact or deferrable projects (e.g. migrating to a github model *a la* CMSSW, expanding and supporting use of Docker, etc...) be deferred until more critical needs have been met.
- The CADMesh work g-2 has done is potentially a major contribution to detector simulation for HEP experiments. As such, we would encourage that the collaboration pursue the publication of a paper that detailed the procedures and important lessons learned.
- Although innovative, the g-2 experiment realizes that there exists a small risk associated with using CADMesh, which is an external product. A statement has been made that, if necessary, the software team is comfortable with maintaining its (reportedly simple) implementation.
- The committee is concerned that the independence of analysis-path infrastructure has led, and will lead, to duplication of effort. To the extent possible, the collaboration should define the boundary where (a) common infrastructure should be used to avoid multiple points of maintenance and duplicated effort, and where (b) independent infrastructure is necessary to maintain blindness of the data and independence of the analysis methods.
- While it is good that a workflow manager is being appointed, the committee is concerned that full end-to-end production workflows are not yet in place, and that no mock data challenge has been performed.
- The complexity of implementing Kalman tracking in a high-gradient, non-uniform magnetic field is a difficult project--it may be challenging to get the algorithm correct, and may be computationally expensive. The committee does not believe that there is expertise within the collaboration to implement such an algorithm within the current 3 month schedule. Nor does the committee believe that external effort would be able to ramp up sufficiently fast enough to meet that timeline.
- We believe that the tracking algorithm is needed for the start of physics production data taking in FY18, giving a time window giving of almost 12 months from the date of this review.
- An external tracking expert for a limited period of time may not be sufficient to ensure longer term maintenance of the tracking algorithms.
- The current tracking working group could benefit from presenting their current status and needs to the Reconstruction group (led by Rob Kutschke) in the Software Scientific Infrastructure department of SCD.
- We recognize g-2's prioritized request for senior scientific effort to engage in the scientific mentoring of students and postdocs as valuable for the scientific progress of the experiment. However, we see this request as not within the scope of this computing review. We would encourage the collaboration to address this need internally or to pursue funding opportunities (such as the Intensity Frontier Fellowships) that will enable collaborating faculty to participate at increased levels in the experiment.
- The collaboration should consider a literature survey to determine if any important external data sets exist for validating the physics of their beamline simulation. If any such

sets are identified, the collaboration should perform the most important of these comparisons.

- The collaboration should make use of the Continuous Integration Service as part of the release management, build infrastructure, and simple integration tests similar to NOvA production.
- The verification package should be expanded to include monitoring of reconstruction performance plots and key physics observables.

2.3 Recommendations

- The collaboration should establish direct communication with the head of SCD to ensure that necessary resources, consultation, and effort are available to ensure the on-going success of the experiment.
- The collaboration should plan to store two full copies of the raw data in an archival manner and should formalize their data retention and data management plan.
- Availability of resources at NERSC or at other centers is critical for the beam simulation group. The collaboration must establish beam simulation computing needs, especially as required during the commissioning phase. The collaboration should establish whether the beamline simulation is portable from NERSC to Fermilab and Open Science Grid facilities. The collaboration should estimate the required resources to achieve beamline simulation goals and execute a strategy to procure the needed allocation.
- The committee recommends that by the start of production physics running in FY18 a mature Kalman tracking/fitting algorithm be established and thoroughly tested. We see two possible effort profiles that will match this goal and recommend that the collaboration, in consultation with the laboratory, take steps to implement one of them:
 - reprogram 1-2 FTEs of collaboration effort to expand the tracking group, expand the timescale for completion from 3 months to 6-9 months, which is consistent with the schedule for return of beam from the '17 shutdown; or
 - seek specific track-reconstruction expertise from outside of the collaboration to join the collaboration and implement Kalman tracking algorithms, while at the same time building the expertise to maintain the algorithm in the longer term.
- The committee recommends that an active production team is promptly built around the newly appointed workflow manager. Then:
 - A detailed production plan for the first publication should be defined, including samples needed and what resources should be allocated to match the publication goal.
 - In order to ensure that a production workflow (including simulation, calibration, reconstruction and validation steps) is ready in time for data analysis, the collaboration should define specific milestones required to meet the goal. Resources from the laboratory and established practices from other experiments should be used as much as possible.

- The collaboration should conduct a mock data campaign that fully exercises data handling and storage tools.

3. Charge question #2

The current online computing infrastructure and tools, including DAQ systems for the precession frequency measurement and the magnetic field, online monitoring, nearline analysis for data quality checks, database and slow control monitoring. Does the experiment have sufficient resources to implement, deploy and operate the infrastructure? Is there enough familiarity with the experiment's approaches and solutions in SCD to be able to provide expert consultation if necessary?

3.1 Findings

- Yes. The committee finds that the experiment has enough effort to implement, deploy, and operate all of the necessary online infrastructure, but that effort is stretched thin to the point where they may not be able to cope with unforeseen difficulties.
- No. The utilization of MIDAS and other customized DAQ solutions means that SCD does not have the expertise to provide expert consultation. The experiment recognizes this risk and accepts the responsibility.
- The average rate of muon fills is 12 Hz. In the calorimeter DAQ, the front-end data rate is 20 GB/s, and this rate is reduced by factor of 100 in GPUs with a lossless compression algorithm. The DAQ computers, networking, and disk have been largely purchased and installed in MC-1. SCD personnel provided guidance for the purchases and will be providing OS support.
- The custom calorimeter microTCA crates and custom digitizer electronics are the responsibility of the experiment. The timing system for the experiment was not discussed in any detail, and it is also expected to be the responsibility of the experiment.
- g-2 uses the MIDAS software suite for their DAQ, and the experiment has developed several MIDAS front-ends to interface to their electronics. The version of the MIDAS software that is used is a modern [current] version that is produced by PSI/Triumf. There is no statement of support from PSI/Triumf for MIDAS, and the experiment will rely on internal expertise for any MIDAS issues.
- The development of the GPU algorithms is done in CUDA. The number of CUDA developers on the experiment is four, all of whom have other responsibilities. The CUDA code is not currently versioned.
- Data is stored in MIDAS files and converted to art/ROOT format offline using a custom *art* Input Source that is based on the one from NOvA.
- Testing of the calorimeter DAQ has been done in a testbeam run at SLAC and in test setups using software emulation of calorimeter electronics. The data throughput at various points in the system has been measured. The necessary throughput can be achieved, and the various components provide 33% or more headroom. Simultaneous writing and reading from the DAQ disk storage has not been tested in a systematic way.

- A laser calibration system will be used to monitor and correct for variations in the SiPM gain. At the moment, the storage of these correction constants is done in FHiCL files, but this is planned to be moved to a database.
- A separate MIDAS instance will be used for the magnetic field DAQs for fixed-probes, the trolley, and the supplemental devices. The fixed-probes system will produce 250 TB of data over 2 years of running. A custom DAQ has been developed for the plunging probes. The DAQ systems for these various field subsystems are in various states of readiness; estimates for remaining work and target deadlines were presented. Despite that, it was not clear if the remaining work can be completed on the expected timescale. The plunging probe subsystem appears to have only one person responsible.
- General plans for databases were presented, as were estimates of the size of the data. Several database types are currently in use (e.g. MySQL, MIDAS ODB, Postgres), and the plan is to move to Postgres. Members of the experiment have started working with members of SCD to learn what is available and design the eventual system. It was not clear that the full set of databases have been clearly identified and the relationship between online and offline databases understood.
- The data quality monitoring for most of the subsystems appears to be in good shape. And, there are plans for lots of slow controls monitoring. DAQ monitoring (monitoring the performance of the DAQ system itself) does not appear to have received much attention so far.
- A nearline processing system has been developed and uses TBB to achieve the necessary performance (ability to process each primary-DAQ file immediately after it is written). The writing of the derived data to ROOT files is observed to be slower than expected.
- Slow controls infrastructure has been developed and tested with sample sensors. This system will not be responsible for any safety-related interlocks. The online slow controls database has not yet been created.
- Files and “databases” for initial hardware and tests are scattered across various locations and stored in different manners.
- The formation of a DAQ operations team is estimated to take 1 FTE day. A training workshop is estimated to take 3 FTE days.
- Particular ASUS motherboards were identified as important for the calorimeter DAQ system. Approximately one dozen are required and the experiment has one spare (perhaps two - the response was slightly ambiguous).

3.2 Comments

- The g-2 experiment should be commended on the progress being made on all aspects of the data acquisition system. The DAQ team in particular should be congratulated on their progress towards meeting a very aggressive schedule. The calorimeter testbeam run at SLAC, the tracker testbeam run, and the cosmic ray teststand, have been used

extensively for DAQ testing and are significant milestones towards a fully integrated and operational vertical slice of the readout systems.

- The experiment should be aware that the SCD lacks the required knowledge and expertise to provide either consultation or support for the MIDAS-based DAQ components.
- The DAQ effort is spread rather thin and no subsystem has the depth of expertise to ensure completion of the work towards commissioning and early operations, if key personnel are lost or unforeseen challenges arise. Based on extensive experience with the commissioning, support and operations of other similarly sized DAQ systems, we believe that the level of support for the core DAQ systems (2 individuals) is unsustainable. The DAQ group is planning to appeal to the collaboration for additional effort at the upcoming collaboration meeting and provide training for interested individuals. It is essential that the collaboration respond to this call in such a manner that there is sufficient expert support for the DAQ. Beyond this call, the committee believes strongly that there may be a need to re-focus and re-direct effort from low priority and low impact offline and online tasks into the DAQ areas that are most time-critical to prepare for data-taking.
- For the magnetic field DAQ, a robust work plan must be re-examined and re-developed that is fully effort-loaded. The effort estimates that were presented tap into a very small pool of individuals, and the committee worries that the total available effort may be incorrect or impossible to sequence within the constraints of the timeline and the individuals involved.
- To mitigate the risks associated with the magnetic field DAQ effort profiles, the experiment's management is advised to seriously assess the priorities of all ongoing and planned tasks to determine their priorities and the impacts of deferring the lowest-priority projects till after the completion and transition into operation of the DAQ. In particular they should consider redirecting this effort into the high priority DAQ and commissioning groups, like the magnetic field DAQ, to ensure its timely completion.
- The plunger DAQ was presented to be in an operational and tested state, however the plans to migrate, through a complete re-write, to a MIDAS based system would consume a relatively large fraction of an FTE. While it would be nice to move this to the MIDAS framework, it is not essential for beam running, and given the effort, we would advise against doing this.
- Compared to the rest of the project, the plan for the online databases is not well-defined or enumerated. In particular, there is not a well-developed plan for recording critical accelerator beam information.
- The IFBeam service provided by SCD, in conjunction with AD, appears to meet the needs of g-2.
- The collection of notes, files, spreadsheets and other "databases" that have been created in relation to hardware testing and construction are problematic from an organizational and operational standpoint. The committee would advise the collaboration to combine and centralize this information into a single hardware database

that is hosted and supported in a standard manner. The experiment would be encouraged to adopt existing solutions either from SCD or from other experiments.

- The monitoring of the DAQ system itself should be added to the DQM and slow controls monitoring that is planned.
- The compression settings that are provided by the RootOutput module should be investigated as a possible source of the slow disk-writing that is observed in the nearline processing.
- The slow controls database is still in the planning state. We strongly advise working tightly with database experts in SCD over the next 2 months to implement and deploy a solution that meets the monitoring and control needs.
- The committee believes, based on extensive experience in DAQ operations and training, that the effort estimates for the creation and training of expert DAQ personnel are severely underestimated.

3.3 Recommendations

- We recommend that the experiment use the IFBeam service provided by SCD and AD for the recording and monitoring of accelerator data.
 - The development work associated with this is estimated to require three months of SCD-provided effort. The experiment should complete their plans for the beamline data and be fully engaged with the SCD and AD IFBeam experts no later than January 3, 2017.
- We recommend that the same data movement tools be used for magnetic field data as are used for the calorimeter and tracker data (these tools are ones provided by SCD), and the storage requirements for this data be communicated to SCD via the existing SPPM process for the purposes of program and budget planning.
- The collaboration should carefully detail and justify the number of required spares to keep the electronics operational for the full planned data-taking run and procure additional spares (e.g., ASUS motherboards for the calorimeter DAQ system) if needed.
- We recommend the experiment develop DAQ monitoring tools that will ensure prompt notification of system problems and minimize downtime.

4. Charge question #3

Are the tools, infrastructure, and established processes sufficient to engage non-expert resources from the collaboration? Are best practices employed in these processes?

4.1 Findings

- The committee was unable to ascertain if the tools, infrastructure, and established processes sufficient to engage non-expert resources because the tools are not yet at a mature enough stage.
- Without further details, it is difficult to know if best practices are being exercised in each case.

- The necessity of providing adequate documentation has been recognized by most if not all of the major computing efforts in the g-2 experiment.
- Current documentation sources include: redmine wiki pages, *art*/LArSoft course material from Aug. 2015, a dedicated g-2 pdf manual, etc.
- The g-2 computing team would like to consolidate documentation as much as possible; it is considering using Doxygen and/or README text files resident to the directories of the components that need documenting.
- A physics analysis user's guide was developed for the 2016 SLAC test-beam and an offline physics analysis user's guide was developed for the omega_a analysis by Adam Lyon.
- A simplified interface for accessing data products, using a bare ROOT interface, was developed for analyzing 2016 SLAC test-beam data.
- The verification package for the simulation physics validation can be easily and efficiently used by non-expert collaborators.
- Computing infrastructure exists for each analysis path, but it is not clear how these infrastructures are tailored to engaging non-experts.

4.2 Comments

- The Offline software team should continue to pursue methods of improving documentation, but not at the expense of developing critical infrastructure. For user-friendly descriptions of *art* modules and services, the description and validation facilities provided by *art* should be strongly considered.
- If a common Offline computing website or other information portal does not exist for the g-2 experiment, the committee strongly advocates that one be created. The website would direct users to the relevant computing information regarding a given aspect of the Offline computing infrastructure.
- The physics analysis guide for the 2016 SLAC test-beam and omega_a offline analysis could serve as a pattern for a mainline analysis guide and a unified analysis users guide should be collated from these.
- Creating a simplified interface for non-experts (*a la* the 2016 SLAC ROOT/C++ functionality) can be very helpful. However, care should be taken to ensure that such functionality is not duplicated elsewhere.

4.3 Recommendations

- None

5. Summary

The Muon g-2 Experiment computing was reviewed in preparation for the start of commissioning data-taking in late Spring of 2017 and commencement of physics data-taking in Fall of 2017. The Collaboration presented material covering all aspects of the experiment's computing during a 1.5 day meeting. The committee is grateful for the care and thoroughness of the presentations and additional material provided as part of the review process.

The Muon g-2 Collaboration should be commended for readily incorporating SCD tools and services within their offline computing infrastructure. The use of the standard tools has placed the experiment in a good position to have a successful and efficient computing model. The committee finds that the DAQ and online systems are on schedule for first data-taking and should be able to achieve the requirements stated for operations. The choice of MIDAS as the DAQ software for the experiment means that there is very limited support within SCD for the g-2 online systems, but the experiment is aware of the risks.

The committee recognized several areas where resources appeared to be stretched critically thin and is concerned by the need to reallocate effort or seek additional external effort. Due to the limited depth of DAQ expertise combined with the start of data-taking rapidly approaching, the experiment is strongly encouraged to broaden the knowledge base of the DAQ systems. Additionally, the development of tracking algorithms needed for studying and monitoring the muon beam position will require considerable additional effort from some combination of reassignment within the experiment and additional consultation external to the experiment.

Appendix A - Charge

Muon g-2 Computing Review November 7-8, 2016

The g-2 collaboration is preparing to begin data taking in late Spring of 2017 to commission beam and detector systems, followed by physics data taking over the next few years. The experiment spokes and SCD management would like the committee to review and evaluate the online and offline software and computing infrastructure readiness of the experiment for successfully carrying out the planned commissioning and physics analysis tasks. In particular, the reviewers should comment on:

1. The current offline computing infrastructure and tools, including build and release tools, simulation tools, framework, analysis tools and approaches, database, workflow, workflow management, data management, and operations. Is the experiment efficiently leveraging tools and expertise offered by SCD? Does the experiment have sufficient resources to implement, deploy and operate the infrastructure?
2. The current online computing infrastructure and tools, including DAQ systems for the precession frequency measurement and the magnetic field, online monitoring, nearline analysis for data quality checks, database and slow control monitoring. Does the experiment have sufficient resources to implement, deploy and operate the infrastructure? Is there enough familiarity with the experiment's approaches and solutions in SCD to be able to provide expert consultation if necessary?
3. Are the tools, infrastructure, and established processes sufficient to engage non-expert resources from the collaboration? Are best practices employed in these processes?

The committee is charged with producing a written report by the end of November 2016 that addresses these questions and makes recommendations for correcting any problems or issues identified.

Appendix B - Committee Members

Kurt Biery (Real-Time Systems Engineering)

Giuseppe Cerati (Scientific Software Infrastructure - Reconstruction)

Burt Holzman (Coordination of Technical Activities - Associate Head/Facilities)

Mike Kirby - Chair (Scientific Distributed Computing Solutions - Associate Deputy Head)

Kyle Knoepfel (Scientific Software Infrastructure - Framework and Software Technology)

Jim Kowalkowski (Coordination of Technical Activities - Associate Head/R&D and Architecture)

Andrew Norman (Coordination of Technical Activities - Associate Head/Sci Workflows and Ops)

Gabe Perdue (Scientific Computing Simulations - Physics and Detector Simulations)

Stephen White (Scientific Data Processing Solutions - Scientific Database Applications)