Report on the 12 GeV Software and Computing Review

Jefferson Laboratory

November 25-26, 2013

David Nathan Brown (LBL/BaBar), Sergei Gerassimov (TU Muenchen/COMPASS), Martin Purschke (BNL/PHENIX), Elizabeth Sexton-Kennedy (FNAL/CMS), Torre Wenaus (BNL/ATLAS) (chair)

December 30, 2013

Executive Summary

On November 25-26 2013 Jefferson Lab held its second Software and Computing Review for the 12 GeV program, following a first review in June 2012. The review was convened by Bob McKeown with a review committee consisting of David Nathan Brown (LBL/BaBar), Sergei Gerassimov (TU Muenchen/COMPASS), Martin Purschke (BNL/PHENIX), Elizabeth Sexton-Kennedy (FNAL/CMS), and Torre Wenaus (BNL/ATLAS) (chair). The committee was asked to review the progress on software and computing for the 12 GeV program since the 2012 review, with particular emphasis on detector simulation, calibration, event analysis, and workflow tools for production and analysis. The review covered the four Halls and included planning for computing resources, and management. The full charge is included as an Appendix.

Jefferson Lab is well advanced on an exciting program of upgrading its accelerator and detector complex, the 12 GeV upgrade, that will enable new insights into the structure of the nucleon, the transition between the hadronic and quark/gluon descriptions of nuclei, and the nature of confinement. Doubling the energy of the JLab accelerator from 6 GeV to 12 GeV will enable three-dimensional imaging of the nucleon, revealing previously hidden aspects of the internal dynamics. The 12 GeV experimental program will be able to address a fundamental challenge for nuclear physics today, understanding the structure and interactions of nucleons and nuclei in terms of QCD. Completion of the 12 GeV Upgrade was recommended as the highest priority by the 2007 NSAC Long Range Plan. In 2013, the program has been modified in response to a changed funding profile to, among other adjustments, delay the completion date by 21 months ("re-baselining"). Commissioning is drawing near, with accelerator commissioning on target to begin in January 2014, Hall A in February, and Hall D in October 2014. Halls B and C will follow with commissioning starts in early 2016. A comprehensive program of 59 approved experiments across the four Halls, corresponding to over 7 years of operation, has been approved to date.

The review opened with an overview of the 12 GeV upgrade project and its status, followed by overview and status presentations from the four Halls, and JLab Computing. The four experimental Halls and their user communities with their distinct facilities and experimental programs each have distinct requirements for computing and software. For each of the Halls the Committee was presented with an overview of the facilities, the experimental program, software and computing, and progress and developments since the last review. Further more detailed presentations and discussions followed in two parallel breakout sessions, Halls A & C and Halls B & D. Finally, the whole Committee reconvened for more detailed presentations and discussion on computing requirements and IT implementation.

The Committee appreciated the high quality of the materials and the informative, well focused presentations. Early availability of materials, presentations and for Hall D detailed documentation on software and computing, was also much appreciated and very helpful. The Halls also provided access to their software wikis on request.

The review concluded with a tentative plan to hold a further review in about a year, by which time data taking with the upgraded accelerator and detectors in the first commissioned Halls will have begun.

A brief summary of the Halls, their 12 GeV programs and their current commissioning schedule:

Hall A: Study of nucleon form factors with the existing HRS spectrometer pair, the new Super BigBite (SBS) spectrometer, and future experiments not in this review's scope (Moller, SOLID). Hall A will be the first to come online in the 12 GeV era, with commissioning scheduled to begin in January 2014. SBS experiments will be coming online in 2017 at the earliest.

Hall B: The new CLAS12 multi-purpose large-acceptance detector with a polarized electron beam and a variety of polarized and unpolarized targets will enable a broad science program in understanding nuclear structure via generalized parton distributions. Hall B commissioning is scheduled to begin in Jan/Feb 2016.

Hall C: The existing HMS spectrometer with the addition of the new complementary SHMS spectrometer will perform precision studies of valence quark properties in nucleons and nuclei. Hall C commissioning is scheduled to begin in Jan/Feb 2016.

Hall D: The new GlueX hermetic detector installed in this new Hall, receiving a linearly polarized photon beam, will explore the origin of confinement by studying as-yet undiscovered gluonic excitation states, crucial to our understanding of QCD in the confinement regime. Hall D commissioning is scheduled to begin in October 2014.

Here follow summary comments on the Committee's assessment of issues identified in the Charge. More detailed and specific observations, findings and recommendations are found in the body of the report.

Software and analysis

The Halls are in general making good progress towards readiness for commissioning and datataking. Progress is measured against milestones consistent with the re-baselined schedule, with assessments made of the adequacy of manpower to complete the work. A program of progressively larger scale data challenges is underway to exercise the more intensive computing demands of the 12 GeV program. Particular attention is being given to tracking progress on critical deliverables such as calibration where substantial work remains to be done. Code evaluation tools are receiving attention but this could be increased.

User support and engagement appears good. Users are given opportunities for gaining real-world experience with new software. Tutorials and workshops have been organized and well attended. New software has been used in realistic studies such as Hall D's PAC-40 proposal, and is validated against legacy software. Due attention is given to documentation through channels like the wikis.

The Halls together with JLab Computing give good attention to encouraging and supporting the use of common software, both at the JLab level and the wider community. ROOT and Geant4 are widely used, with usage expanding since the last review. A new committee coordinating support for common tools is a welcome initiative that could also provide a context for engaging and collaborating with the external projects. The workflow tools initiative is addressing a shared need in a common way. There is further potential for common solutions; examples may be HTCondor, Hall B's own CLAS12 Data Management System (TagFS), and Pythia8.

Management

Management is doing an accomplished job. The state of preparedness for commissioning and datataking given the schedule is good, suggesting that staffing levels are adequate. Management structures and processes appear to be serving the Halls and users well. Changes since the last review show that management continues to evaluate and adjust structures and processes to best serve the program. Risks are assessed and mitigation steps taken both by JLab Computing and in the Halls. Extremely careful computing planning and budgeting helps mitigate risks in delivering the needed capacity and capability. The Halls, through measures such as data challenges and watchful committees covering areas on the critical path, are also working to assess and mitigate risk, and plan a smooth path from development to deployment and operations with minimal surprises. Adapting to scope and milestone changes has been well exercised with the re-baselining, and software and computing seems to have adapted well.

Computing and Networking

Processes in place to assess and deliver the needed computing requirements are impressive. Requirements, capacity planning and budgeting within overall Lab budgeting constraints are very carefully managed to ensure that computing, storage and networking requirements of the experimental program are scrutinized, translated into optimal plans to deliver the needed capacity on the basis of careful benchmarking, and integrated into budgets for just-in-time delivery to achieve the greatest cost-benefit. Technology tracking is actively pursued, and new technologies integrated where beneficial. The Lab could work with DOE Office of Science to minimize budget boundaries between programs, so as to maximize resource utilization.

Comments and Recommendations

The following sections present the comments and recommendations of the Committee across the areas covered by the review: general issues, data acquisition (only lightly touched in this review), common Experimental Hall issues, Hall-specific issues, and management and computing. Sections are organized as Observations, Findings and Recommendations. The recommendations from the 2012 review are included in an appendix for reference.

General comments and recommendations

Observations

- Progress has been excellent, the program seems to be well on track to being ready for data when it arrives at the different halls.
- The re-baselined program permits a smoother ramp-up of computing capacity as the Hall B schedule has been delayed, and the computing resource planning has been revised to take good advantage of this.
- The data challenge schedules dovetail well with the ramp towards commissioning and physics.
- The recommendations that ROOT be used as a primary data format have not been adopted but lots of use of ROOT translators is reported. Translated ROOT files are (re)generated by users as needed, they are not systematically produced and kept on disk.
- We heard about an effort to generate "self-documenting" output in the sense that a ROOT object in the datastream will capture a comprehensive parameter set that went into generating the file, such as macros, git software versions, input streams, and also database/calibration constants.
- The Halls have embraced the previous recommendation to use git as the main software management tool, and have selected the github service to host the master repository. Currently there are two persons with write access to the master who triage pull requests. In the future, this will be delegated to more "package owners" who will be responsible for managing the (currently about 8) different projects. There has been extensive training and Wikis to get collaborators on board with git, with tutorials how to manage and merge forks, and how to make use of the web-based management and visualization utilities which github provides.
- Standardization on SCons for build management is a good development.
- The legacy make system, which appears to have consisted of unmanaged Makefiles (e.g. no use of auto-tools), has been converted to SCons. At this time, the two make systems still coexist, but make will be phased out as soon as SCons is perceived to be a full replacement.
- All of the previous users of Geant3 reported that they are moving to Geant4, a C++ implementation, as their simulation engine.

Findings

- Reported issues such as memory leaks and non negligible rates of crashes and core dumps suggest that more attention to QA and the use of code evaluation tools could be beneficial.

Recommendations

- While we observe progress in the adoption of code evaluation tools in the development cycle we encourage increased attention to this, both static and dynamic analyzers. In the case of valgrind, investigate the use of a suppression file to suppress issues outside your own code.
- Consider code reviews, particularly in contexts where it could be most beneficial to get new developments started on the right footing in C++, such as new codes contributed by the physicist community, particularly if they are migrated legacy Fortran codes.
- An impressive program of software workshops and tutorials was described. Some were reported to be recorded; we recommend this as a valuable means of getting the greatest return on the workshops, making them an ongoing resource for new users.
- The move to Geant4 is very much encouraged. This however is a time of transition for the Geant collaboration as they try to adapt to the challenges of the multi-core era. We recommend that the support personnel within IT devoted to common tools make an effort to keep abreast of these changes by collaborating with the other labs like SLAC, FNAL, and CERN.

Data acquisition

Observations

- Data acquisition was not explicitly covered as it was last time, nor is it in the charge. We note that the data challenge programs include online data challenges as a means of ensuring DAQ/online readiness.

Experimental halls - general

Observations

- Halls were uniformly good in considering the recommendations of the last review.
- User presentations demonstrated the software is being actively and successfully used by the physics community. We appreciated their inclusion, three in the A/C breakout and two in the B/D breakout.
- Attention to documentation is evident in the webs/wikis for the halls.
- Computing budgets include resources sufficient to meet all projected simulation needs of the Hall physics programs. External (non-JLab) resources are not covered by MOUs at this time so they are not included in official resource assessments. They are regarded and used as opportunistic incremental resources.

Findings

- The committee asked whether data challenges include calibration activities. In particular, are there plans to test the extraction of critical calibration constants from simulated data.

Hall D reported they are doing this for CDC, it needs to be explored for calorimeter and beam-related calibrations. Hall B, not clear.

Recommendations

- We recommend the inclusion of calibration in the data challenges. Demonstration of a functional prompt calibration loop will be an important element of achieving rapid extraction of physics results.

Hall A and C

Observations

- The Committee was advised that with the SBS program scheduled for beyond 2017, the SBS software and GEM track reconstruction would be reviewed at a later date.
- The re-baselined schedule was taken as an opportunity to change the earlier plan to extend the support of the legacy Fortran to SHMS, instead migrating to C++ as recommended by the last review.
- Hall C is making use of Hall A's C++ infrastructure (PODD C++ library) in their Hall C specific C++ library. This consolidation is a very positive development.
- The Halls are implementing parallel processing schemes to make more efficient use of multi-core and hyper-threaded architectures. For Halls A and C it was stated that converting the existing reconstruction software to be fully thread-safe was not possible on a reasonable timescale. Rather, their plan is to parallelize the processing by forking subprocesses; each one working on one event at a time. The perceived benefit over the standard approach of running one fully independent, not multi-threaded analysis process per core is the re-use of calibration constants and other metadata for all forked subprocesses, rather than having potentially different metadata sets for each core. It was recognized that this approach requires a fair amount of inter-process communication, such as shared memory sections or sockets.
- The adoption of git and github by Halls A and C is a positive development. The issue tracking and branch tracking capabilities of github are being actively used to improve communication and reduce the barriers to development by users.
- Presentations by users showed support for the directions being followed by the Hall software supporters. In particular, the ability to easily generate ROOT diagnostic plots using text configuration seemed to be popular.
- Existing calibration procedures use text format data files. These files apparently can be produced by the PODD, but the stated goal is to convert the calibration procedures to run directly off ROOT structures.
- Comparisons between PODD and Engine low-level variables for Hall C detectors differ within acceptable tolerances, validating the migration of algorithms for tracking, calorimeter and particle ID reconstruction from Fortran to C++.
- Hall A is exploring adding a third tracking chamber to resolve ambiguities associated with high-rate running.

Findings

- We were presented with some conceptual plans for future software directed at SBS and SOLID. The requirements driving the conceptual designs were not clear.

Recommendations

- We recommend a design review of SBS software and GEM tracking, at an early enough date that plans can be adjusted before a substantial amount of effort has been invested.
- We recommend sharpening the case for event-parallel analysis. It should be viewed as a tool rather than an up-front goal. It is clear that a genuine thread-safe environment will have huge benefits and keep the number of open file handles, database connections, and the memory footprint to a minimum. The presented intermediate step to use forked processes has the potential to incur a complicated inter-process communications protocol (using shared memory or sockets), and will require inter-process synchronization if the master process discovers the need to adjust the meta-data.
- We recommend looking at TagFS as a potential tool for file metadata and data discovery.
- While basing the new Hall C C++ code on the well tested Fortran ENGINE code makes sense, we recommend reviewing the C++ code design to insure it conforms to standard OO practices.
- The goal of using the ability of Hall A and C “self documenting” output to include all the input parameters should be re-examined. While the capturing of some critical parameters and other input will help document the properties of a given output file, one should not expect to be able to exactly recreate the file.
- Provide a generic mechanism to capture the available monitoring histograms and other output at the end of a data acquisition run. This will be helpful if “forensics” are required in the case of a problem with the data, and can be used to improve the monitoring.
- Effort should be made to identify personnel capable of extending the Hall C software migration to full tracking, calibration, and physics calculations.

Hall B

Observations

- Software manpower is reported to be adequate, with the core team able to serve the needs of the experimental community. They report good collaboration and engagement with detector groups. Well attended software workshops are being used to engage and educate the community.
- The presentation on the CLAS12 Data Management system was very impressive and forward thinking. The goal of fully utilizing the IT resources of participating university collaborators and easily enabling that participation is very important. We note that as presented, the system apparently has features of a data preservation system, perhaps at

a further level of advancement than others.

- It was stated in Gerard Gilfoyle's talk that Hall B uses Pythia for event generation and that it is well supported by Lund and CERN, however this is true only for Pythia8, a C++ implementation, and it turns out that Hall B is using software that they inherited which is a forked version of Pythia6, a Fortran implementation.

Findings

- The last review questioned the value of multi language support relative to its cost, and recommended regular re-assessment. It was reported this time that support for the broadly popular python was added, and that the multi language capability is further used to allow early adopters to use C++11. The multi language support seems to be showing good value without visibly onerous costs.
- We commend the convergence of the simulation and reconstruction on a common geometry through gemc's use of the CLARA geometry service.
- We commend the establishment of a calibration and commissioning committee.

Recommendations

- We recommend that raw data be separated from reconstructed data such that (re)reconstruction outputs do not carry redundant raw data copies.
- We recommend exploring the extension of the Hall B data distribution and analysis framework for the needs of raw data processing, data calibration and physics analysis. This distributed computing capability will be an asset to the experiment and could help mitigate risks of insufficient computing funding.
- We encourage exploring whether the CLAS12 data management system has data preservation aspects that could be useful to a wider community.
- The authors of Pythia6 have stated very strongly (at least to CMS) that they no longer support it, so for Run 2 of the LHC CMS has already moved to Pythia8. We suggest that this may actually be an opportunity to collaborate with other labs on generator software. If Pythia8 could be made useful for Hall B, then both the fields of ENP and HEP would benefit. FNAL has one of the Pythia authors on staff and it would be good to have a face to face meeting with him to explore the idea.

Hall D

Observations

- Hall D software shows excellent planning, organization and documentation.
- The software is deployed and is being actively and successfully used in physics analyses with sophisticated event selection.
- As at the last review, Hall D reported on a Data Challenge exercise that included real physics analysis at a substantial scale. They reported on the 2012 DC which provided

the event samples for the PAC-40 strangeness program proposal, the analysis for which drove the development and use of the analysis software.

- Software workshops are being held, most recently a well attended 2 day workshop in July 2013, repeating an analysis reported last time (with substantial improvement in the analysis itself, 66%/25% purity/efficiency in 2012, now 95/22).
- The DC plan for 2014 is well defined with three DCs planned at progressively larger scales, reaching 5000 cores in Fall 2014. They include also scaling tests of production from the tape silo up to the same scale.

Findings

- Explicit inclusion of QA as a tracked activity with recognized manpower requirements is a positive step.
- Similarly data challenge management needs have been recognized and incorporated into planning.
- We commend the appointment of a calibration coordinator (since the summer). Manpower is reported to be in place for the work.

Recommendations

- Event tagging in the HLT is recommended as a mechanism for separating calibration data samples into streams for use in a prompt calibration loop.
- We recommend against consideration of SRM. The LHC grid community is moving away from it as a heavy and expendable layer.
- We recommend against consideration of LFC. It will soon have no LHC users and will be deprecated.
- We recommend TagFS be examined as a possible file metadata catalog solution.

Management and Computing

Observations

- Requirements, capacity planning and budgeting within overall Lab budgeting constraints are very carefully managed to ensure that requirements of the experimental program are scrutinized, translated into optimal plans to deliver the needed capacity on the basis of careful benchmarking, and integrated into budgets for just-in-time delivery to achieve the greatest cost-benefit. The requirements originally requested by the experiments were about twice the levels presented to us, suggesting that scrutiny and optimization is taking place.
- Plans to make LQCD resources available to the 12 GeV program when needed for peak usages such as data challenges have proceeded, not yet rolled out but will be soon as the LQCD capacity is needed to reach targeted scales for 2014 data challenges.
- Adjusting resource allocations between 12 GeV and LQCD is presently a manual process. In the long run, automation would be an asset. Other facilities are leveraging VM/cloud technology to this and other purposes.

- Regarding change control, the re-baselining entailed a broad revision of planning to assimilate the new schedule. Adaptation to the new schedule seems to have been comprehensively addressed across the Hall programs and JLab computing, leveraging the re-baselining as an opportunity for revising plans such a smoother capacity ramp and software planning adjustments including moving away from Fortran at Hall C. Taking this as a measure of change control in action, processes seem good.
- Regarding adequacy of contingency and risk management processes. The care taken in assessing computing requirements, how they can be technically delivered, and the associated budgeting lends confidence that risk exposure in terms of delivering the required computing capability is being minimized. The data challenge programs are also an important part of minimizing risks in meeting the schedule and functional/scaling requirements of commissioning and physics. In the individual Halls, the addition of a specific QA task to stress test the software is a positive development. Hall B did a risk mitigation assessment, concluding that calibration software readiness is the main risk, which is being addressed by seeking more manpower and regularly monitoring progress through their CalCom committee.
- In the talk given by Sandy Philpott it was stated that IT was able to support the ENP data challenges by being able to move nodes between the USQCD, and ENP clusters. They see that this need will not disappear in the future so they are developing a process to automate this procedure. It was also stated that strict accounting was done to make sure that one Program did not use more resources than it could afford to pay back within its computing budget. When asked what would happen if this condition was reached, the response was that the CPUs would have to sit idle.

Findings

- The Halls responded well to the recommendations of the last review. Many were followed, others were considered and investigated.
- Staffing levels for software development and documentation appear to be successfully delivering, and so seem adequate.
- As concluded at the last review also, management structures and processes appear well-matched to the needs of the collaborations and users, with improvements made in the interim, notably the common software support committee.
- The plans for transitioning from a development phase into a deployment and operations phase, centered around an ongoing program of data challenges of increasing scale, look appropriate. Timelines are consistent with commissioning.
- The establishment of a committee to coordinate support for common software (ROOT, Geant4, CERNLIB, EVIO, etc) is an excellent initiative, promoting and supporting software commonality among the experiments, and leveraging community software. We heard also about new community tools being used/evaluated, such as Globus Online. Further potential opportunities even within the JLab community itself came to light in the review, such as the TagFS. This suggests these reviews are useful for exposing opportunities within the community, but also that more day to day channels may be called

for.
- The workflow tools initiative as a common service to the Halls is well motivated.
- We were impressed by the extremely careful budgeting and planning based on detailed metrics and close consultation between physics and computing. Discussions take place ~monthly on priorities and current spending decisions.
- Requirements for computing, storage and networking were well stated and well justified.
- The computing and networking plans of the Laboratory are well matched to the 12 GeV program requirements, with effective processes in place to ensure this remains the case. Careful optimizations, metrics, and technology studies ensure they are cost effective. Budgets appear appropriate and carefully planned to optimize cost/benefit. The low 100Mb level of failover networking is striking, but the Laboratory finds this adequate until the upgrade anticipated for 2015.
- JLab Computing demonstrates a receptiveness to learning not only what they should do, but also what they don't have to do. This is not always seen in IT organizations. (For example, the experiments are using github rather than JLab creating a service supported in-house.)

Recommendations

- VM and cloud technology has been enthusiastically adopted at other facilities (CERN, BNL, FNAL, …) because of its benefits to operational flexibility and reducing operations load. It was reported that VMs are extensively used at JLab, particularly on the non-scientific side. On the scientific side, VMs are under study to preserve the ability to analyze data far in the future. Also being examined is running Linux containers on LQCD machines so that Linux jobs can run in parallel with LQCD running on the bare metal. We encourage such studies on VM exploitation on the scientific side, particularly to support flexible sharing of resources between LQCD and 12 GeV computing.
- HTCondor was reported to be under consideration for adoption as a batch service. We recommend it as a widely and successfully used community standard, used heavily for example among OSG sites.
- We were told that experiments are being encouraged to move to multi-threaded codes. While the known and potential drivers for this are many and were described to us (declining memory per core as core count grows; head thrashing on disks as the open file count grows with core count; cache utilization efficiency; quickening of the development cycle; latency-sensitive processing like L3 triggers), the investment effort to realize multi-threaded codes can be large. We recommend that care be taken that this guidance is offered in the context of case by case cost/benefit assessments that show multi-threading is worth the investment.
- The committee fully agrees that it is natural that demand for computing resources fluctuates due to accelerator up time and conference schedules. Moving towards a horizontal computing services layer within the lab is encouraged. Clearly the vertical funding nature of Projects within the DOE Office of Science is not something that JLab itself can do anything about. However since it is clearly sub-optimal to let CPUs go idle

because of funding boundaries, the lab may want to consider joining those encouraging the Office of Science to more broadly support computing services across Projects and Programs (see for example the recent Snowmass report).

Appendix: 12 GeV S&C Review Committee Charge

Nov 25-26, 2013

The committee is asked to review the state of software and computing developments for the 12 GeV program at Jefferson Lab, with particular emphasis upon

- Detector simulation, calibration, and event analysis
- Workflow tools for production analysis

The review will cover all 4 halls, including management and planning for computing resources. The committee is asked to address the following questions, keeping in mind the different timelines for the different halls:

1. Software and Analysis

a. Are the halls making appropriate progress towards having their simulation, data acquisition, calibration and analysis software ready? Are they meeting their previously set milestones?

b. Have an adequate set of milestones been identified, and an appropriate set of tests been incorporated into the milestones, to measure progress towards final production running?

c. Are the halls getting users engaged at an appropriate level to demonstrate usability and readiness from a user's perspective? Have the collaborations identified effective and appropriate mechanisms to support utilization of the software by the entire collaboration? Is the level of user documentation appropriate for this point in time?

d. Are appropriate efforts towards software commonality being made across the halls and/or with the wider HE/NP communities?

2. Management

a. Did the halls respond appropriately to the recommendations of the last review?

b. Are staffing levels for software development and documentation appropriate?

c. Are the current management structures and processes well-matched to the needs of the collaborations (including users)?

d. Are there appropriate contingency and risk-management processes in place? Have risks been appropriately identified?

e. Are reasonable change control processes being used to address scope and milestone changes?

f. Are there adequate plans for transitioning from a development phase into a deployment and operations phase? Are the timelines appropriate?

3. Computing and Networking

a. Are the requirements for computing, storage and networking well stated and well justified?

b. Are the computing and networking plans of the laboratory well matched to the requirements, are they cost effective, and are budgets appropriate for these plans?

2012 Recommendations (For reference)

2012 General Recommendations

1. Presentations in future reviews should cover end user utilization of and experience with the software in more detail. Talks from end users on usage experience with the software and analysis infrastructure would be beneficial.

2012 DAQ Recommendations

1. Once a modest all-way data path is established, plan a mock data challenge with fake data, in particular with nominal data rates from GlueX.

2012 Experimental Hall General Recommendations

1. Nightly builds are performed by some; we recommend them for all.
2. Evaluate standard code evaluation tools, such as valgrind, clang's scan-build, cppcheck, Gooda, ... for inclusion in the software development cycle. We suggest looking at an Insure++ license as well.
3. Run a code validation suite such as valgrind as part of the routine software release procedure.
4. Give full and early consideration to file management, cataloging and data discovery by physicists doing analysis. Report on this area in future reviews.

2012 Hall A Recommendations

1. Investigate the feasibility of event-based parallelization of C++ analysis in a multi- core batch environment.
2. Intensify efforts on the HRS tracking development, including calibration and alignment procedures. Define performance milestones which allow time to explore alternatives if problems arise.
3. Study the SBS track reconstruction algorithm efficiency under higher background conditions. It would be useful to know at what level of background the existing algorithm stops functioning.
4. Develop requirements for the SBS algorithm performance, along with a development timeline and a responsible contact. Requirements should include alignment and calibration.
5. Investigate if a move to git along with Hall C makes sense.

2012 Hall C Recommendations

1. With the somewhat aggressive schedule leading up to December 2013, make sure to engage a reasonable number of early adopters to stress test the new framework.
2. Re-use existing efforts from Hall A to decode CODA-formatted data in ROOT.

3. If resources are limited, the Fortran-based SHMS reconstruction should be a low priority.
4. While we encourage the move to git as a code management system, be sure not to underestimate the extent of the paradigm shift. Identify a workflow model for your use of git. Communicate clearly the new paradigm (easy branching, no central repository, etc) Set up (or link to) tutorials for users with a mapping of routine CVS tasks to their git equivalents (such as cvs diff, etc). Document or link to documentation for standard git tasks without obvious equivalent in CVS or SVN, such as git rebase, or bisect.

2012 Hall B Recommendations

1. A series of scaling tests ramping up using the LQCD farm should be planned and undertaken. Tests should begin soon; don't wait for completion of the software 18 months before startup.
2. Seriously consider using ROOT as the file format in order to make use of the steady advances in its I/O capabilities.
3. The costs and sustainability of supporting two languages, relative to the advantages, should be regularly assessed as the community of users grows, code development practices become clearer, the framework matures further, etc.

2012 Hall D Recommendations

1. A series of scale tests ramping up using JLab's LQCD farm should be planned and conducted.
2. The data volume and processing scale of GlueX is substantial but plans for data management and workload management systems supporting the operational scale were not made clear. They should be carefully developed.
3. Consider ROOT (with its schema evolution capabilities) as a possible alternative for the HDDM DST format.

2012 Management and Computing Recommendations

1. To ensure a smooth transition from development and deployment to operations, particularly for Halls B and D, an explicitly planned program of data challenges, directed both at exercising the performance of the full analysis chain and at exercising the scaling behavior and effectiveness of the computing model at scales progressively closer to operating scale, is recommended. We heard more explicit plans from Hall D than from Hall B in this respect. This data challenge program should be underway now, and should not await the full completion of the offline software.
2. In response to the question as to how the computing budget is scrubbed, the answer received was that scrubbing happens through this review. This review hasn't examined the requirements and associated budget sufficiently for this to be considered a scrubbing. Also it is not clear that an overall optimization of the computing models, associated resource requirements, and required budget levels has been done. A process should exist whereby this optimization takes place. For example are the relative roles of

disk and tape optimal for making analysis as effective as possible, within budgetary constraints.

3. The measures being planned to render LQCD resources usable by the 12 GeV community should have high priority.