

Report on the 12 GeV Software and Computing Review

Jefferson Laboratory

10-11 February 2015

<https://www.jlab.org/indico/conferenceDisplay.py?ovw=True&confId=93>

Committee members: Sergei Gerassimov (TU Muenchen/COMPASS), Martin Purschke (BNL/PHENIX), Elizabeth Sexton-Kennedy (FNAL/CMS), Torre Wenaus (BNL/ATLAS) (chair)

May 6 2015

Table of Contents

[Table of Contents](#)

[Executive Summary](#)

[Comments and Recommendations](#)

[General](#)

[Observations](#)

[Findings](#)

[Recommendations](#)

[Experimental Halls - General](#)

[Observations](#)

[Findings](#)

[Recommendations](#)

[Hall A and C](#)

[Observations](#)

[Findings](#)

[Recommendations](#)

[Hall B](#)

[Observations](#)

[Findings](#)

[Recommendations](#)

[Hall D](#)

[Observations](#)

[Findings](#)

[Recommendations](#)

[Management and Computing](#)

[Observations](#)
[Findings](#)
[Recommendations](#)
[Appendix: Agenda](#)
[Appendix: Charge](#)

Executive Summary

On February 10-11 2015 Jefferson Lab held its third Software and Computing Review for the 12 GeV program, following reviews in November 2013 and June 2012. The review was convened by Bob McKeown with a review committee consisting of 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 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
- Computing plans, including projections for cores, disk, and tape for the next 3 years
- Software and computing management

The committee was asked to address the questions posed in the charge (included as an Appendix), keeping in mind that the four halls are in various stages of preparation corresponding to their staggered commissioning and physics running schedules (i.e. more focus on Hall A and Hall D / GlueX).

Jefferson Lab has now embarked upon its exciting 12 GeV program with the first datataking having begun. Previous reviews of software and computing indicated that the Laboratory and the Hall experimental programs have given careful consideration and attention to preparing S&C for the needs of the physics program to enable expeditious analysis once data is in hand, minimizing the 'time to PRL' in the terminology we heard several times in this review. This review has reaffirmed this attention to S&C readiness and the state of preparation impressed the review committee. All the Halls are on track with S&C milestones and schedules consistent with the current commissioning schedule. The halls are generally making appropriate progress towards having their simulation, calibration and analysis software ready. New S&C developments are being appropriately deployed and hardened through exposure to real users and scale exercises in data challenges, and now in real datataking in some cases. User engagement and buy-in looks positive, with good attention to supporting users with tutorials. There has been good attention to leveraging and increasing commonality across the halls, and in leveraging common software and computing resources beyond the Lab. Staffing levels appear adequate, the JLab 12 GeV computing budget continues to be carefully managed and optimized to meet the evolving requirements and schedule of the program efficiently within the available means, indeed with economies relative to earlier planning. The resource sharing with LQCD discussed at the last review is now active practice. The recommendations of the last review were appropriately addressed.

The Committee appreciated the high quality of the materials and the informative, well focused presentations. Availability in advance of the presentations was much appreciated and very helpful.

Here follow the Committee's detailed comments and recommendations across the areas covered by the review.

Comments and Recommendations

The following sections present the comments and recommendations of the Committee across the areas covered by the review: general issues, common Experimental Hall issues, Hall-specific issues, and management and computing. Sections are organized as Observations, Findings and Recommendations. The Findings and Recommendations from the 2013 review are included in an Appendix for reference.

General

Observations

- The IT group provides a professional and centralized computing environment with about 4000 cores, a tape storage system, an advanced job and workflow management system (swift), and sophisticated accounting tools. The system manages node sharing between the ENC and HPC groups, with in-kind contributions at a later time balancing the shared resources.
- The underlying disk and file systems use inexpensive JBODs and the Lustre distributed file system, with ZFS as the disk format. An upgrade from Lustre from version 1.8 to 2.5 is planned for the near future.
- The sophisticated resource and performance tracking allows to identify bottlenecks, and allows to properly steer future procurements.
- The strategy of lab support for common community tools is a good one and needed as the tools are evolving and their developments must be followed.
- Engagement of and support for users appears to be uniformly good. Tutorials, workshops, documentation, expert assistance, JLab visits and so on are being used effectively. We did not hear evidence of the forking of software efforts because of dissatisfaction with mainstream efforts.
- Good attention continues to be given to identifying and cultivating common software efforts. We heard (though it wasn't presented) that the JLab-level committee responsible for common software continues its work.
- Management structures and processes continue to be effective at ensuring software and computing needs are met efficiently, and the analysis communities are served well by the software. The management is giving close and responsible attention to readiness for prompt physics publication. Processes to manage risk, provide contingency, and manage the still evolving 12 GeV program schedule appear to be operating well. Schedule and budget for S&C are consistent with the overall 12 GeV program.
- Careful tracking and provisioning of computing, storage and networking requirements continues. The Lab has a good plan for network provisioning and ensuring that fallback networking is upgraded. Cost management has been good, cost projections have actually decreased, for well documented reasons. The wider JLab program,

particularly LQCD, has been effectively leveraged to further control costs and maximize benefit.

Findings

- The Halls responded appropriately to the recommendations of the last review.
- While staffing is tight, and there are clearly areas where greater effort would bring worthwhile benefits (such as Hall A's long standing todo list), staffing appears adequate. Every effort should be made to engage young students and postdocs in software development as analysis ramps up, to find the next generation of software experts.
- It seems that none of the experiments have a good meta data catalog solution. For the larger experiments this will become important. A common solution should be considered.
- The halls are generally making appropriate progress towards having their simulation, calibration and analysis software ready. We have concerns in one area, Hall D simulation, which we discuss below.
- Data challenges and now the close and rapid study of the first 12 GeV era data are being used to ready the Halls for production scale operation and physics readiness. JLab has been an important facilitator of at-scale tests by providing the ability to spike usage for data challenges, and opportunistic resources have been effectively used to this purpose as well. The ~50% scale of the latest DCs is appropriate for the validation of production scale operation. The capability to spike processing resources will also be valuable for accelerating the completion of physics studies.
- We were not presented with data management and cataloging plans for Halls B and D where it will be important to have robust solutions. To the extent that this indicates plans and/or solutions are not in place, this will be a problem.
- We encourage the JLab-level committee responsible for common software to monitor and possibly engage with the recently established HEP Software Foundation (hepsoftwarefoundation.org) which is fostering common software in HEP and related fields such as NP, particularly on concurrency.

Recommendations

- None

Experimental Halls - General

Observations

- Although we did not hear CVS mentioned any longer, the code management tool in use is not yet fully standardized on git and github.

Findings

- Revised computing requirements from the experiments as well as other factors have reduced the cost of 12 GeV computing in the coming years. This makes it an opportune time to take another look at the cost/benefit of computing -- if spending

more in certain areas would have substantial benefit, this may be a moment when that could be built into the plan and budget. If for example it was only realized later that there would be great benefit to increasing the use of disk, the budget to do so might be harder to come by.

- JLab has organized a series of Geant4 tutorials in 2006, 2009, 2012 which have proven useful to the attendees. This is good and perhaps the time for another one is approaching.
- The groups should implement a way to perform bit-level checks on the output of identical jobs, which should come out completely identical. If not, this usually points to the use of uninitialized or inadvertently overwritten variables, which somehow pass the analysis tools.
- We heard that experiments generate and store online monitoring histograms, however, this is in most cases a manual process -- Hall D being an exception where it is done automatically, a ROOT histogram file is created during datataking and for each run is written out and archived to tape. The generation and storing of such monitoring output should be done in an automated fashion in general. Also, make a distinction between operator-level monitoring output (which is typically looked at by the shift crews) and expert-level output, which is usually only stored, but will be a valuable resource in case a problem is found later.

Recommendations

- It seems that some combination of code analysis tools such as cppcheck and valgrind are being used by all experiments. The applied tools should be unified to some extent to capture a larger phase space of potential problems, such as using clang's scan-build feature. It would be beneficial if a professional code analysis tool such as *coverity* would be licensed and made centrally available.
- Those groups that have not yet set up nightly rebuilds should do so, and flag the checked-in code that caused the rebuild to fail.

Hall A and C

Observations

- Hall A users express satisfaction with the software, and say it is easy to get started. There are not many questions from confused/frustrated users.
- Hall C is expecting software takeup to happen as detectors arrive. It is a bit early now to judge how the takeup is going. Users have successfully followed web documentation on their own to get started.
- Hall C seems most advanced in its use of github, with Hall A following suit.
- Hall C, having adopted the C++ framework more recently, is in a position to compare the results from the C++ and legacy (*engine*) frameworks, and they are doing so.

Findings

- We observe a commendable sharing and adoption of developments between Halls A and C (the C++ framework developed by Hall A, and many technologies such as the move to github and the implementation of nightly rebuilds by Hall C).
- Hall A has the most diverse user groups in the sense that different groups set up experiments with different requirements for tracking resolutions and data rates. We heard that each group employs their own run number system for data management, which has repercussions on the use of parameter and description files. We assume that a given experiment maintains unique run numbers for itself, so that the descriptions are internally consistent.
- The move to Geant4 is commended.
- The Hall A software is mature, current, popular with its users, well supported, and benefits from long experience in supporting the wide array of physics studies that will continue to characterize the Hall A program in the 12 GeV era. Hall A's readiness for 12 GeV era analysis appears to be good.
- The state of SBS software planning and development appears appropriate given the schedule. SBS construction is to be complete circa end 2017. Experiments in SBS will come about a year later, so the software timeline presented, with a design review circa summer 2015 and projecting readiness in 2018 is consistent.
- SBS software requirements differ across the experiments SBS will serve. There is a responsible plan to develop a design that promotes software commonality and reuse while also supporting well the diverse SBS physics community. An effort is being made to identify all requirements early and meet them in common components where possible. A summer 2015 design review will be timely and requires that significant effort be given to design work between now and then.

Recommendations

- Clarify for the users the role of timestamps and run numbers. Unless the condition is varying too rapidly, we recommend using run numbers as a primary key for constants. Treat the time as a secondary information to be stored with the collection of constants.

Hall B

Observations

- A set of common framework tools has been developed including a geometry package common to simulation, reconstruction and event display, I/O common library, ccdb access tools and other elements. A calibration framework is planned as a future development.
- The GEMC simulation is now based on the latest Geant4 version 10, covers most of the detectors and is progressing towards using multi-threading.
- Java, python and C++ are supported in CLARA. Java is the basis of all framework development. Python is used e.g. for steering scripts. There are no C++ modules at

present or planned; C++ is said to integrate less well. The user community continues to be happy with the framework.

- User experience was presented as good. The list of people and projects using the framework was presented, it is long, and has grown since the last review.

Findings

- We commend the development of common framework tools including common geometry across all geometry consumers.
- It's commendable that you have followed up on the Pythia6 to Pythia8 suggestion and have a good counter-proposal to contribute to HERWIG++.

Recommendations

- Explore the use of Analysis Trains in collaboration with GlueX (see Hall D recommendations), so the technology is in place once the data become available.

Hall D

Observations

- We highly commend the success of the first data taking and application of the software chain to real data in the Fall 2014 run. It was very quickly shown that the software works, with rho, pi0 and omega observations coming promptly. The team is making maximal use of the data now in hand to improve software and calibration, as they did with cosmics last spring.
- The offline software stack is broadly used, particularly for calibration at the moment. This will continue and expand as grad students and others move to analysis.
- There are plans to assign thesis topics by the summer, focusing initially on topics that do not demand full understanding of detector efficiency and normalization to publish.
- GlueX has a very modern and advanced reconstruction code base. However, its simulation engine is one of the oldest in use. Most of the field has moved from Geant3 to Geant4. The collaboration has taken the first important step in the migration with geometries that work both for Geant3 and Geant4. The remaining work is in wrapping/rewriting the current C/C++ routines that convert energy depositions to detector responses so they can be used by GEANT4.
- Data-taking rate was limited to 2 kHz in the fall run, for a number of reasons including packing only one event per block and a larger event size than nominal. The rate is expected to grow to ~20 kHz in the next run.
- Opportunistic (including OSG grid) resources were used effectively and at a substantial scale in a data challenge, much greater than JLab resources. The overall scale of the DC represented roughly half of a nominal data taking year. The software failure rate was ~0.1%, down by 1-2 orders of magnitude from previous grid runs. Of about 150 TB total data volume, 50 TB was replicated off-site (primarily to one institution) with gridftp.

- A software L3 trigger will be exercised for the first time in the spring, in pass-through mode, with tagging of events (e.g. pi0's) useful for calibration.
- A careful analysis and optimization of the event size has brought it down to 24 kB; the target size of 18 kB (up from the original 15 kB estimate) is said to be within reach.
- Bottlenecks with tape staging are beginning to be noticed. JLab's imminent tools for tape-aware workflow management will be timely.
- While the reconstruction gained a factor three speedup on the new hardware, the simulation ran at about the same speed. This is likely due to Geant3 Fortran not leveraging processor capability whereas the threaded reconstruction does.
- The digitization is very CPU heavy, taking about the same time as the simulation itself.
- We were not presented in the review with a data management/cataloging strategy. We were told of possible solutions including the Event Store and TagFS. Post-review it was clarified that the collaboration is in the process of implementing Event Store with the plan to use it on existing data. This product is both a data catalog as well as Event and run-based metadata. The expectation is that it will be running by the start of the spring 2015 run.

Findings

- The lack of simulation speedup on new architectures, in contrast to the multithreaded reconstruction, combined with the fact that most CPU resources go towards generating simulation samples, should raise the priority of investigating the simulation technical performance.
- The full implementation and validation of the geometry and simulation with Geant4 is a major effort, given for example the stricter geometrical requirements of Geant4. It is important that this effort continue with priority. This work is underway with a common geometry description in place that works with Geant3 and Geant4 and has been validated in Geant4.
- Unlike their colleagues in the other Halls, Hall D is not yet using git, and is perhaps seeing the effects in being faced with difficulty in merging the rapid development following the recent development (which was branched) with the trunk.
- It is commendable that EventStore (a product developed by CLEO III) was adopted for use in the collaboration instead of trying to re-implement your own.
- Experience with opportunistic resources points to their usefulness in delivering throughput, and also to the challenges they present. Opportunistic resources require short jobs, meaning small files -- this makes it a slow process to recover the files back to JLab. They are accessible at OSG sites also, at least up to space availability limits, but that implies cataloging where they are. This should motivate attention to data management/cataloging.
- Broad and early use of the software stack in analysis is vital to increase the likelihood that the first analyses and publications will proceed quickly. This seems to be happening, but a user presentation would probably have made it clearer. We encourage the assignment of the first round of PhD thesis topics soon.

- Understanding of detector and trigger performance through simulation is also essential to turning out physics results early. We encourage increased attention to simulation and digitization in general as physics running approaches.

Recommendations

- Establish milestones for the migration to Geant4, prioritized appropriately considering other activities and the needs of physics running, and identify more manpower to complete the milestones.
- Establish a strategy and timescale for meeting data management/cataloging needs, exploring whether common tools (e.g. with Hall B) can be part of the strategy.
- Raise the priority of investigating and tracking performance problems with profiling tools. The current choice of valgrind is heavy. Consider using a sampling profiler, and even better consult with the HPC staff to both borrow a licensed commercial tool and get help in understanding the results.

Management and Computing

Observations

- The resource sharing between LQCD and the experimental program described as an objective last time has been successfully implemented, and is having a positive impact on budgets and planning, with the experimental program currently banking CPU that will 'level the peaks' in coming years and provide a pool permitting large scale data challenges. There is appropriate accounting and tracking in place, taking account of hardware performance, with LQCD 'paying back what they gain' (though most of the payback is expected to be on the same architecture). Metrics are reported monthly and LQCD is reported to be happy with the arrangement.
- Computing resource requirements have been revisited in light of reduced experiment requirements, revised running times, stretched out schedules, etc. which yield a substantial overall reduction: about half the previous cost estimate in 2018.

Findings

- JLab will soon roll out their workflow management tools, which proactively include mechanisms that will allow optimizing the processing of tape-resident data (e.g. with a 'data carousel' approach). We expect this will prove to be timely and useful.
- We commend JLab for putting in place the promised resource sharing with LQCD. This will surely bring benefits for the science and opportunities for budget optimization.

Recommendations

- Explore, ideally in collaboration with Hall B, the use of *Analysis Trains* which have become the backbone of user data analysis at other facilities. Even if the current data sets are small enough to be kept disk-resident entirely, this is likely to change in the future. Trains are ideal to make the best use of scarce resources, such as tape

bandwidth. Assign a person to be responsible for the maintenance of train-managed data sets.

- As you move from the era of data challenges to that of data taking you should transition the people you have operating the challenges to a computing operations group that is responsible for both the reconstruction of collected data and the creation of monte carlo samples for analysis. If you decide that analysis trains are useful, the computing operations group would also insure that the coordination and services required are available.

Appendix: Agenda

Tuesday, February 10

- 8:00 – 8:40 Executive Session
- 8:40 – 9:00 Introduction and 12-GeV Overview Patrizia Rossi
(short update on where we are and what the detectors are)
- 9:00 – 9:30 Halls A and C Overview and Progress Ole Hansen / Mark Jones
(high-level overview and tracking versus recommendations/milestones,
include Computing Requirements and Budget (bottom up))
- 9:30 – 10:00 Hall B Overview and Progress Veronique Ziegler
(high-level overview and tracking versus recommendations/milestones,
include Computing Requirements and Budget (bottom up))
- 10:00 – 10:15 Break
- 10:15 – 10:45 Hall D Overview and Progress Curtis Meyer
(high-level overview and tracking versus recommendations/milestones,
include Computing Requirements and Budget (bottom up))
- 10:45 – 11:15 Computing Overview Graham Heyes / Sandy
Philpott
(roll-up of requirements and facility planning)
- 11:15 – 12:15 Halls A and C Deep Dive (< 50% presentation)
Software Maturity and associated labor, User Experiences, Discussion of how to go from
production beam to publication, examination of remaining risks
- Hall A deep dive (15) Ole Hansen
- C user experience and path to publication (15) Gabriel Niculescu
- 12:15 – 1:30 Working Lunch (Executive session)
- 1:30 – 3:00 Hall D Deep Dive (< 50% presentation)
- Online software performance and status (15) David Lawrence
- Offline software performance and status (30) Mark Ito
- 3:00 – 3:30 Break
- 3:30 – 5:00 Hall B Deep Dive (< 50% presentation)
Software maturity and path to publication Gagik Gavalian
- User experience and utilization Jerry Gilfoyle
- 5:00 – 6:00 Executive session
- 6:00 – 6:45 Questions

7:00 – 9:00 Reception and Dinner

Tuesday, November 26

8:00 – 9:00 Q&A

9:00 – 11:00 Executive session – prepare draft

11:00 – 11:30 Close Out

Appendix: Charge

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
- Computing plans, including projections for cores, disk, and tape for the next 3 years
- Software and computing management

The committee is asked to address the following questions, keeping in mind the different timelines for the different halls (i.e. more focus on Hall A and Hall D / GlueX):

1. Offline Software: Detector Simulation and Analysis
 - a. Are the halls making appropriate progress towards having their simulation, 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 and readiness for producing publication quality results?
 - c. Are the halls doing the right level of at-scale testing of each of simulation, event reconstructions, and physics analysis appropriate to the time before engineering and physics running?
 - d. 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?
 - e. 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? Are all of the assumptions clearly stated, and are all of the units clearly defined (e.g. "E2670v3 core" vs "core") ?
 - 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?