12 GeV Software Review Report

Chip Watson
Jefferson Lab Scientific Computing
July 7, 2015

Outline:

- Software Review Process
- 2015 Charge
- Review Summary
- Impact

Software Review Process

Motivation:

Minimize time to first publication of excellent science!

Experimental physics software is not part of the 12 GeV scope, and so would not be reviewed for progress towards physics readiness as part of the (robust) 12 GeV Project reviews.

Reviews serve both to measure and aid progress, bringing in additional outside perspectives and suggestions for improvement. This is an important part of the peer review process that makes science successful.

History

Composite (multi-hall) computing plan evolved from 2007-2011

First formal ½ day review held in May 2011, with a broad scope (all IT)

First "Annual" (~18 month) 1½ day review held in June 2012, with subsequent reviews in Nov 2013 and Feb 2015

2015 Charge (high level)

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

Questions for the Review Committee

- Offline Software: Detector Simulation and Analysis
  - Are the halls making appropriate progress towards having simulation, calibration and analysis software ready? Are they meeting previously set milestones?
  - Have an adequate set of milestones been identified, and an appropriate set of tests been incorporated, to measure progress towards final production running?
  - 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?
  - 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?
  - Are appropriate efforts towards software commonality being made across the halls and/or with the wider HE/NP communities?

Questions (2)

- Management
  - Did the halls respond appropriately to the recommendations of the last review?
  - Are staffing levels for software development and documentation appropriate?
  - Are the current management structures and processes well-matched to the needs of the collaborations (including users)?
  - Are there appropriate contingency and risk-management processes in place? Have risks been appropriately identified?
  - Are reasonable change control processes being used to address scope and milestone changes?
  - Are there adequate plans for transitioning from a development phase into a deployment and operations phase? Are the timelines appropriate?
- Computing and Networking
  - 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. “E2670vs core” vs “core”) ?
  - Are computing and networking plans of the lab well matched to requirements? Are they cost effective, and are budgets appropriate for these plans?

2015 Review Summary

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

Recommendations

General

- 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.

Halls A & C

- 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.

Recommendations

Hall B / CLAS 12

- 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 / GlueX

- 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.

Recommendations

Management and Computing

- 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.

Laboratory View

Lots of progress!

Most things going very well!

Room for improvement: high quality software ready for Day 1 physics is now a requirement, not merely “nice”. The world has changed and expectations for software and computing are high.

Impacts of the 2015 Review on Computing

- Computing requirements thoroughly scrubbed just ahead of the review
  - CLAS requirements delayed one year
  - GlueX requirements reduced by ~2x, mostly due to reduced assumed weeks of running
  - Definition changed from vague "cores" to specific cores (note that Haswell cores are ~2x the performance of Ivy Bridge cores)
- Plans for Farm expansion in 2015 dropped (farm was already over provisioned and under utilized)
  - Next upgrade planned for Fall 2016
- Priority of larger disk capacity raised
  - Additional work needed, but currently being doubled

Next Review

The next review will probably be late Summer or early Fall 2016. Scrubbing computing requirements again will shape the build out of the farm and determine the capacity for the 2016 procurement.

Focus will include data management plans, and production calibration and analysis plans.

Questions?