

Report from the 2018 Review of Jefferson Lab Scientific Software and Computing

David Brown (LBNL), Frank Golf (UNL), Adam Lyon (FNAL), and Arjun Shankar (ORNL)

[Introduction](#)

[Executive Summary](#)

[Observations](#)

[Findings](#)

[Recommendations](#)

[Responses to the Review Charge Questions](#)

[Appendix A: Charge Questions](#)

[Appendix B: Review Agenda](#)

[Appendix C: Questions for Presenters and Organizers](#)

Introduction

The fifth bi-annual review of the Thomas Jefferson National Accelerator Facility (JLAB) Scientific Software and Computing (S&C) for the experimental program was held November 27 and 28, 2018. The committee consisted of David Brown (LBNL) (chair), Frank Golf (UNL), Adam Lyon (FNAL), and Arjun Shankar (ORNL). Previous review chair Torre Wenaus (BNL) was invited to the review but was unable to attend due to sudden illness.

The committee received advance materials on November 7 consisting of the previous S&C review report, the review charge and draft agenda, and the JLAB response to the 2017 S&T review. Most presentation materials became available to the committee through the [indico site](#) on November 20. The organization of the review was very clear, and the organizers did an excellent job providing for the needs of the committee.

Presentations on Scientific Software and Computing were given November 27. The presentations were of excellent quality and covered all the issues relevant to the charge. The organizers and speakers did an excellent job keeping to the schedule.

Based on the presented material the committee prepared written questions for presenters and the JLAB S&C coordinators at the end of the day November 27. Responses to these questions were presented by JLAB staff November 28. These responses were appropriate and answered the questions to the committee's satisfaction. The review committee presented a draft report during the review closeout on November 28. The general conclusion of the committee is that the Software and Computing are very effective in supporting the current JLAB physics program, and that the development plans are appropriate to continue support for the upcoming physics program.

Below we list observations, findings, and recommendations. Observations are statements of fact the committee chooses to highlight. Items in bold deserve special focus. Findings are opinions and suggestions from the committee. Recommendations are items that require a specific response.

Observations

Hall A + C

- Onsite resources suffice now and in the mid-term future computing projections for the needs of Halls A and C
- Podd software is managed using standard tools (Github, Redmine issue tracking, Travis validation)
- Podd processing is a single pass from raw data to physics observable plots/trees
- SWIF is used to manage workflows (also in Halls B+D)
- **Hall A tritium cross-section ratio analysis is near publication**
- **Hall C F2 deuteron/proton ratio measurement has a preliminary result**
- Podd framework will depend on root 6 through the SBS era
- Both Halls project that their simulation and reconstruction needs will be met by their 5% (each) allocation of local computing resources
- SBS has tested reconstruction performance at up to 200% expected background rates in simulation
- SBS software is supported by a large but coordinated pool with a core of dedicated postdocs
- SBS GEM simulation has been tuned to prototype response measurements using cosmic rays and sources
- **Virtual experiments are planned in the near future to test SBS readiness for data**

Hall B (CLAS12)

- The CLAS12 physics program has begun since the last review.
- The CLAS12 program is divided into 42 experiments grouped into 12 “Run Groups”.
- The collaboration successfully addressed recommendations from a May 2018 (internal?) software review. Improvements include DAQ/trigger efficiency and purity, monitoring, event size, and reconstruction time.
- The JLAB community DAQ system is used, but at one version behind. There are CLAS12 customizations in their version.
- **Run Group A was able to process (calibrate, reconstruct, produce DSTs) within three months of collecting data. Significant physics results are planned for fall 2019. First results on 10% of the data were shown at DNP2018 (Congratulations!)**
- Processing estimates were shown. The total processing time for five years of data is two years.
- Current processing efficiency is 50%. This is reasonable for a first attempt at processing a significant fraction of the data. Although resources are sufficient to support this efficiency moving forward, the collaboration is actively working to improve this and expects to do so.
- The collaboration structure was shown and computing looks to be well integrated.
- The CLARA scientific framework is used with the COATJAVA toolkit (libraries). These codes are available on Github. Pull request model is used as well as the Travis CI system.
- Designing the reconstruction tools in a composable framework is a good design approach. However, the general creation of custom lab-dependent software (CLARA, COAT, HIPO) is a risk to manage. (Appreciate the provided weblinks to the tools, but some documentation has empty links - e.g., Tutorial python bindings into CLARA).
- CLARA code review, profiling, and reconstruction improvements is a valuable step forward.
- Enabling fault isolation and reuse of the analyzer modules will help the deployment of software tools.
- Analysis trains allow for efficient use of computing resources.
- Bit packing is a clever way to decrease the event size.
- Inefficient use of tapes for the Spring 2018 running is corrected for the future, but accessing that data will require some amount of extra resources (either to re-write tapes or to stage the entire dataset)
- Plan to run simulation offsite with OSG and NERSC. 30M NERSC allocation has been requested to cover 50% of simulation workload. Docker and CVMFS will be used like GlueX.
- Using software coordinators across DAQ and Offline software will be important to ensure the data, tools, and capabilities are aligned.
- We value the analysis on simulation time required per event and the analysis of cycles available and needed. Based on previous lessons, it would be useful to track these requirements and usage over time.

- **Overall, CLAS12 is on track to produce timely and important science and is leveraging expertise and processes from other experiments and the lab itself.**

Hall D (GlueX)

- The physics program is well organized, with a plan for timely publication of physics results. This is facilitated by an elected physics analysis coordinator.
- **Early beam asymmetry measurements based on 12 GeV data have already been published.** New measurements (e.g. J/psi photoproduction, SDME, additional beam asymmetry measurements) are expected to be submitted for publication imminently. Kudos to the collaboration for the timely production of these results.
- **Steps have been taken to reduce data processing burdens on analyzers through simple APIs/interfaces. This is strongly commended.**
- Monte Carlo is generated on-the-fly, also reducing burdens on users. There may be some duplication, but given the modest MC needs balanced against the savings in user effort, this is justified and commended.
- There is a standard run conditions database used for bookkeeping of data, allowing for reproducible run lists used throughout the collaboration.
- Studies of potential benefits of trading some CPU for data compression are still under study. Data compression may also be important as more processing moves off-site.
- The collaboration has begun to explore high intensity operation. Pilot runs already demonstrating >80% livetime are very encouraging. Work to address limitations from disk write speed is in progress.
- Processing of the full 2018 dataset is expected to be ready for analysis by late summer 2019.
- NERSC is now heavily relied on for reconstruction. Running reconstruction (instead of simulation) matches NERSC preferences. OSG resources are used to meet most simulation needs.
- The collaboration now has two years experience with physics running. Estimates for future needs are informed by this experience.
- **Overall, GlueX plans and actual developments are excellent and appear to match what is needed to produce timely and important science.**

Experimental Computing Overview

- **We commend the quantitative response to the previous S&T recommendations to make a plan for resourcing over the coming years.** The “Response to 2017 S&T recommendation for Experimental Computing” (Boehnlein, Hayes, Watson) was a logical presentation of the plan and approaches being investigated.
- The estimates and the planning steps (resources use by Hall, going off premises, using SSD, etc.) are sensible.
- Computing resource requirements have already grown beyond what can be satisfied with local resources. This has been driven by the growing needs of halls B and D. **The**

lab has worked with the experimental halls to make use of off-site resources such as NERSC and OSG and incorporate these into the out year computing plan.

- Investigating increased use of SSDs as a pass-through and to enhance performance of I/O intensive workflows.
- It is important that the plan is regularly reviewed to ensure this continues in the future.

Scientific Computing and Scope

- Monitoring of the infrastructure for file system issues is a valuable step.
- The web tools for monitoring and scientific computing is valuable to the operators and users.
- Pivot away from X86 to Phi (now being phased out by Intel) and GPU is logical for LQCD. It will be worth being attuned to any unintended consequences of having a progressively different footprint between the QCD and the farm clusters. Discussion on the efficiencies of the jobs running on the other cluster was well received.
- The continued use of older hardware and the arguments for extracting value from it is commendable.
- **The decision process for the investments (tactical and strategic) were useful and appeared well thought through.**
- **We commend the increased emphasis/investment on disk availability to make data access easier, and to enable jobs to run offsite.**
- Plans to use test nodes as infrastructure as a service in AWS will help JLab and potentially be experiences to share with other laboratories.

Scientific Computing Implementation

- The move to SLURM is logical and the testbed approach is a welcome step forward.
- Flash data buffer is an appropriate implementation for improved DAQ and reliability.
- Use of the Lustre data management and movement back and forth from tape is well appreciated. A robust offering of the Jasmine tool could have broader appeal to other laboratories.
- Planned near term and longer-term upgrades are appropriate and will improve the operational reliability. This will be directionally correct to anticipate the growth of data in the 12 GeV era and the data analysis needs.
- **We strongly commend and are impressed by the overall operational posture and support given to JLab's programs with a lean operations team in scientific computing.**
- Uptime for the ENS cluster is impressive.
- An OSG head node is now run by JLab with central support and help from the OSG community.
- SLURM provides a mechanism for offering interactive nodes with fast access to the data.
- Off-site computing is limited by the bandwidth available to read data from the lab. Plans to add another 10 Gb/s link in 2019, followed by a 100 Gb/s link in 2020, are important to a) support growing demand for off-site computing resources and b) have flexibility in the mixture of job types that can be run off-site.

Outlook and Gaps Talks

- The scientific computing interest in machine learning/AI is timely. In addition to data analysis for the Halls, its use in detector monitoring and accelerator monitoring and tuning could have a high impact.
- We agree that new classes of data intensive explorations will require more interactive computing and resources, tools, and consulting expertise. That you are planning for such activities is a welcome development. Staffing up to meet this need is an appropriate response.
- Keeping an eye on mitigating one-deep risks is important.
- Potential future experiments (Moeler, SOLID) will require additional computing resources
- JLab is an EIC software consortium member
- **A grand challenge to better integrate the data taking-calibration-analysis cycle is in preparation**
- Opportunities exist for improving uptime and reliability through redundant tape and disk resources
- A central L3 processing farm might be an effective future way of pooling experiment resources
- A paradigm shift to interactive computing would require new infrastructure and support.
- A global analysis strategy for NP physics parameters requires new resources and support.

Findings

- Subject-matter-expert/staff turnover is a risk for future support of custom JLab-specific software tools.
- NERSC allocations are now an important resource to support ENP computing needs. This is a positive development. It is important to continue investigating additional offsite resources as part of future planning.
- Argonne is starting to open their KNL machine (THETA) to running with containers via Singularity, which could be another remote HPC resource along with NERSC. Similarly, OLCF discretionary resources (leading to proposals like ALCC or INCITE) is a possibility.
- More efficient data transfer mechanisms for OSG (e.g XrootD) would allow for running reconstruction at these sites.
- Additional operations and data processing associated with a 34 week run schedule may stress existing computing and software support resources.
- At this time, we do not see a compelling benefit of 24/7 support that outweighs the cost. Experiments seem to have proper expectations of best-effort off-hours support. This

situation, however, should be closely monitored and revisited as the load both onsite and offsite increases.

- Need for a central data cataloging system should be gauged. If there is demand, tools and expertise from other laboratories should be considered.
- An alternative way to fill unused cycles on the local ENS cluster would be to open those to outside OSG users, which may provide JLab leverage in access to OSG.

Recommendations

- Prepare to support increasing interest in machine learning and modern data science tools, possibly in collaboration with other labs to leverage existing solutions.
- Consider increasing the central support for offsite resource access, especially for OSG and data transfers, leveraging work already done by GlueX and CLAS12 and at other laboratories.

Responses to the Review Charge Questions

1. Offline Software: Detector Simulation and Analysis
 - a. Are the event reconstructions and physics analysis appropriate for publication quality physics results for the foreseen experimental program over the next five years? Are the halls producing appropriate levels of simulated events?
 - Yes
 - b. Has each hall developed multi-year estimates consistent with CEBAF operations of 34 weeks/year for offline computing resources that appropriately support physics analysis and timely publication of results? Are those estimates consistent with the experience gained to date for data rates, processing time and need for simulation to support analysis?
 - Yes
 - c. Are there any identifiable gaps in computing resources that would impede the timely production of publishable results? Are there any bottlenecks to doing user level analysis? If so, what are the plans towards addressing them?
 - No, but support for offsite resource usage may become an issue in future
 - d. What opportunities exist for common tools and approaches with the broader HEP/NP community that can be utilized by Jefferson Lab? How can Jefferson Lab collaborations increase the utilization of new approaches and tools, such as machine learning?
 - Opportunities exist for collaborating with outside groups working on topics such as vectorized tracking code being developed by LHC collaborations.
 - JLab can leverage the work done at other DOE labs in deploying scalable machine learning and deep learning frameworks in native and containerized forms.
 - e. Are there functionality gaps with respect to the deployed tools?
 - We see no major functionality gaps
 - f. What, if any, steps have been taken to support distributed computing, either for processing or for analysis?
 - Reconstruction and Simulation processing is being done on NERSC and OSG, respectively, and the tools (e.g., the SWIF modifications) support distributed analysis.

2. Management
 - a. Did the halls respond appropriately to the recommendations of the last review?
 - Yes
 - b. What lessons have been learned in transitioning from a development phase into physics?

- We noted that lessons were learned from the development phase: improved data transfer patterns from DAQ, that onsite resources need to be augmented with offsite, estimates of resource requirements needed to be updated, and operational use is improved with system monitoring and feedback to the users.
- c. Please comment on the organizational model and expertise considering the need to simultaneously support operations, algorithmic and infrastructure enhancements and to deliver timely physics results.
 - We notice an informal but collegial relationship to gather requirements which appears to be effective for the JLab processes. Specific coordination meetings were noted which will clearly help operations and science outcomes.

3. Computing and Networking

- a. Are the computing and networking plans of the laboratory well matched to the requirements, including CEBAF operations of 34 weeks/year? Are they cost effective, and are budgets appropriate for these plans?
 - Yes. The link upgrade and ESNet plans will help.
- b. Is the implemented scientific computing architecture at JLab appropriate for the needs of the physics program? (Balance of tape vs. disk, networking to the halls, proposed online and trigger farms, onsite computing farms). Are there plans for achieving sufficient robustness and reliability for online and offline systems?
 - Yes. The balance of disk vs. compute investment was done with some thought. Reliability improvement plans for Lustre and processes for the management of the data across disk and tape are sound.
- c. Is the role of distributed computing appropriate and achievable? What is the experience with distributed computing to date? Are there barriers to scaling out the use of distributed computing?
 - Experiments have demonstrated use of remote computing resources at NERSC, OSG, and other sites. The lab plans to address WAN bandwidth limits. Improved data transfer mechanisms for OSG jobs should be explored (see finding above).
- d. Over all, is the plan for provision computing using local and distributed resources consistent with a target of publication within a year for key measurements?
 - Yes, some key measurements appear to be publishable within 1 year

Appendix A: Charge Questions

12 GeV Experimental Computing Review Committee 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

- ∅ Computing plans, including projections for cores, disk, and tape for the next five years
- ∅ Performance of the Scientific Computing Systems and planning for future needs
- ∅ Software and Computing considerations for Physics Readiness

1. Offline Software: Detector Simulation and Analysis
 - a. Are the event reconstructions and physics analysis appropriate for publication quality physics results for the foreseen experimental program over the next five years? Are the halls producing appropriate levels of simulated events?
 - b. Has each hall developed multi-year estimates consistent with CEBAF operations of 34 weeks/year for offline computing resources that appropriately support physics analysis and timely publication of results? Are those estimates consistent with the experience gained to date for data rates, processing time and need for simulation to support analysis?
 - c. Are there any identifiable gaps in computing resources that would impede the timely production of publishable results? Are there any bottlenecks to doing user level analysis? If so, what are the plans towards addressing them?
 - d. What opportunities exist for common tools and approaches with the broader HEP/NP community that can be utilized by Jefferson Lab? How can Jefferson Lab collaborations increase the utilization of new approaches and tools, such as machine learning?
 - e. Are there functionality gaps with respect to the deployed tools?
 - f. What, if any, steps have been taken to support distributed computing, either for processing or for analysis?
2. Management
 - a. Did the halls respond appropriately to the recommendations of the last review?
 - b. What lessons have been learned in transitioning from a development phase into physics?

- c. Please comment on the organizational model and expertise considering the need to simultaneously support operations, algorithmic and infrastructure enhancements and to deliver timely physics results.
3. Computing and Networking
 - a. Are the computing and networking plans of the laboratory well matched to the requirements, including CEBAF operations of 34 weeks/year? Are they cost effective, and are budgets appropriate for these plans?
 - b. Is the implemented scientific computing architecture at JLab appropriate for the needs of the physics program? (Balance of tape vs. disk, networking to the halls, proposed online and trigger farms, onsite computing farms). Are there plans for achieving sufficient robustness and reliability for online and offline systems?
 - c. Is the role of distributed computing appropriate and achievable? What is the experience with distributed computing to date? Are there barriers to scaling out the use of distributed computing?
 - d. Over all, is the plan for provision computing using local and distributed resources consistent with a target of publication within a year for key measurements?

Appendix B: Review Agenda

Jefferson Lab 12-GeV Experimental Computing Review

November 27-28, 2018

L102/104 CEBAF Center

Agenda

Tuesday, November 27

8:00 – 8:45 Executive Session – also includes few slides on where we are with 12-GeV

8:45 – 9:00 Introduction – General JLab Computing and Experimental Software Structure
(to introduce reviewers to mode of operations at JLab)

Amber Boehnlein

9:00 – 9:30 Halls A and C Overview and Progress Ole Hansen / Mark Jones
(include what the detectors are, the schedule for the next five years, software experience from first 12-GeV experiments, path to publication, high-level overview and tracking versus recommendations/milestones, estimates of Computing and Storage Requirements for 5 years)

- 9:30 – 10:00 SBS Software and Computing Seamus Riordan (ANL)?
(include status of GEM tracking/reconstruction software & validation)
- 10:00 – 10:15 Break
- 10:15 – 10:45 GlueX/Hall D Overview and Progress: A Blueprint for Physics
Curtis Meyer / David Lawrence
(include what the detectors are, the schedule for the next five years, computing/software experience so far, lessons learned, prospects for high-intensity run, tracking versus recommendations/milestones, efforts and planning for off-site computing, estimates of on-site Computing and Storage Requirements for 5 years)
- 10:45 – 11:15 CLAS12/Hall B Overview and Progress Stepan Stepanyan?
(include what the detectors are, the schedule for the next five years, software experience from first 12-GeV experiments, path to publication, high-level overview and tracking versus recommendations/milestones, efforts and planning for off-site computing, estimates of on-site Computing and Storage Requirements for 5 years)
- 11:15 – 12:15 CLAS12 Software and Computing Deep Dive Raffaella De Vita / G. Heyes
(maybe split in 2 talks? Talk about CLARA software status, gaps, items to complete, changes made in computing requirements, software, organization since last review, recent successes, etc.)
- 12:15 – 12:45 Experimental Computing Overview Graham Heyes
(roll-up of requirements and facility planning, roll-up of off-site planning assumptions, DOE/NP report)
- 12:45 – 2:00 Working Lunch (Executive session)
- 2:00 – 2:30 Scientific Computing and Scope at Jefferson Lab Chip Watson
(overview of current resources, and the process by which we arrived at this point including optimizing for science (compute, storage and bandwidth, memory, networking), ongoing evolution of the computing environment to support swings in demand and to integrate available ("free") external resources, near-term (FY2019) evolution (cloud, increasing offsite resources), including expected changes in computer architecture (component) and how we choose likely future winners)
- 2:30 – 3:00 Scientific Computing Implementation Sandy Philpott

(Current capacity (compared to requirements) and FY2019 likely purchases to respond to requirements, Operations (everything that serves the users: queues, allocations, fair share, trouble tickets), overview of work flow tools: Auger, Jasmine, SWIF (present and FY2019 evolution), management (decision making) and software support (Physics div. "volunteers"))

- 3:00 – 4:30 Outlook and Gaps for Future Nuclear Physics Program
(3 talks, 30 minutes each, <50% presentation)
- 1) Gaps for Possible Future Experimental Program Steve Wood?
(include GEANT4 needs, TDIS/streaming, MOLLER, SoLID, etc.)
 - 2) Gaps for Possible Future Scientific Computing Chip Watson
(include planning for 24/7 computing support, plans/hurdles for advanced scientific computing in out-years)
 - 3) Gaps for Possible Future Phenomenology Nobuo Sato/Jianwei Qiu?
(include state-of-the-art global fit/phenomenology, further software development needed for radiative corrections & 3D global fits, needs)
- 4:30 – 4:45 Break
- 4:45 – 6:30 Executive session
- 6:30 – 6:45 Questions
- 7:00 – 9:00 Reception

Wednesday, November 28

- 8:00 – 10:00 Q&A
- 10:00 – 12:00 Executive session – prepare draft
- 1:00 – 1:30 Close Out

Appendix C: Questions for Presenters and Organizers

- Graham/Chip: What criteria are used to decide which workflows run locally vs remotely?

- Chip: Is there a list of new feature requests from the Halls and individual users? Is there a process for prioritizing computing/data/software work to address these feature requests?
- Sandy/Chip/Graham: Are there metrics (e.g., user satisfaction, ticket turn around times, etc.) that you can share from user feedback?
- Sandy/Chip: Please give us performance information on the current computing facility. (E.g., what were significant downtimes of the farm and Lustre for the past year? How does Lustre perform under high load?) Can you give some examples of significant improvements that could be made if additional resources were available?
- All Halls: Roughly, how many Phd theses are currently underway that use 12GeV data?
- GlueX/CLAS12: What is the impact of NERSC allocations/scheduling on your publication timeline?
- Amber: Do you have a count by year of publications that have relied on JLab resources?
- Chip: What duration (number of years) are you obligated to hold the Halls and LQCD data?
- Gluex: How does the NERSC preference for non-simulation submissions (Graham's last talk, slide 13) affect your plans to run simulation on NERSC?
- CLAS12: What is the source of the 50% inefficiency on slide 13 of Graham's talk?