Report on the

12 GeV Software and Computing Review

Jefferson Laboratory

June 7-8, 2012

Executive Summary

On June 7-8 2012 Jefferson Lab held a Software and Computing Review for the 12 GeV program, convened by Bob McKeown and with a review committee consisting of David Nathan Brown (LBL/BaBar), Sergei Gerassimov (TU Muenchen/COMPASS), Chris Jones (FNAL/CMS), Martin Purschke (BNL/PHENIX), and Torre Wenaus (BNL/ATLAS) (chair). The committee was asked to review the state of software and computing developments for the 12 GeV program with particular emphasis on detector simulation, calibration, event analysis, and workflow tools for production and analysis. The review covered all four Halls and included data acquisition, planning for computing resources, and management. The full charge is included as an Appendix.

Jefferson Lab is engaged in 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 its present 6 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 last (2007) NSAC Long Range Plan, and is on track for commissioning in 2014 and first beam for early experiments in 2015, with the schedule and cost performance indices at the time of the review both at 96%. A comprehensive program of 48 experiments across the four Halls, corresponding to over 5 years of operation, has been approved to date.

The review opened (see Appendix 2 for the timetable of presentations) with overview presentations on the 12GeV project as a whole, and the detectors, software and physics. The four experimental Halls and their user communities with their distinct facilities and experimental programs each have distinct requirements for 12 GeV program computing and software. For each of the four Halls the Committee was presented with background information on the facility and physics program to date (for the three existing Halls), an overview of the new facilities and experimental programs under development for the 12 GeV program, and details of their software and computing requirements, status, and planning. Further more detailed discussions took place in two parallel sessions, Halls A & C and Halls B & D. Other presentations covered DAQ and common online tools, offline computing, and networking.

A brief summary of the Halls and their 12GeV programs:

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 begin its experimental program with early experiments in 2015, with SBS experiments coming online in late 2016.

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. The Hall B experimental program will begin in 2015.

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. The Hall C experimental program will begin in 2015.

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. The Hall D program will begin with commissioning in Run I, ramping up to full scale GlueX running in Run III.

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 frameworks, simulation, analysis and calibration software were generally found to be in an impressive state. Software development appropriately leverages the long legacies of Hall software while moving to modern languages and approaches. Development timelines are consistent with the 12 GeV schedule, with required manpower levels presented that appear credible, achievable and are generally rather carefully estimated. Community tools such as ROOT are being well leveraged, migration to Geant4 is everywhere at least foreseen and in some cases complete, there is commonality across Halls on calibration and visualization software. Potential exists for further commonality (e.g. in code management). Frameworks and analysis software are being developed with new computing technologies (MT, MP, co-processors) in mind, to an impressive degree. On a cautionary note, C++ development (in particular) should make use of code evaluation and optimization tools to ensure code efficiency and to discover and correct deleterious coding practices early. Consider lessons learned from HENP community experience in OO frameworks: code for performance.
- Regarding software usability and readiness from users' perspective, usability and readiness seemed to have an appropriately high priority in the development efforts, expressed for example in assuring the continuity of trusted usable software in the transition from 6 GeV to 12 GeV codes, but the Committee would have welcomed more explicit attention to usability (talks from analysis users) in the presentations.
- Plans delineating specific testing programs and milestones to measure progress towards at-scale production running were mixed. Hall efforts included well developed plans for successive Data Challenges progressively scaling up testing of the software and computing systems; we recommend such plans be made general practice, and make full use of JLab's available computing resources for realistic scaling tests. Such plans will be key not only to software and computing readiness but to a smooth transition from development to operations. When datataking begins, computing operations at realistic datataking levels should not be a new experience.
- The Committee heard little on data management plans despite the fact that this will be an important part of the infrastructure. Also plans for workload management were unevenly developed among the Halls. These are particularly important for the more data and processing intensive programs of Halls B and D. Plans in these areas should be carefully developed, and are good candidates for common solutions.
- Software and computing management and support seemed appropriate and functional, reflecting long successful histories in the existing Halls. JLab physics and computing management seems well attuned to giving computing and software the importance necessary in order for the 12 GeV program to deliver physics results in a timely way.

- Budgets and schedules appear adequate. The resource requirements have been developed through close consultation between physics, software and computing, and so should be accurate, but (as is recognized) should be regularly revisited in an ongoing process. The resource requirements do not strain the ability of JLab and the 12 GeV experimental community to provision sufficient computing resources assuming foreseen budgets can be relied upon.
- The early and close attention that is being given to software and computing development and readiness, together with the fact that the 12 GeV program does not demand 'bleeding edge' or exceptionally large scale computing, lends confidence that risk mitigation and the ability to react to emerging contingencies is in good shape. The committee did not identify areas in which risks or contingencies are being seriously underestimated.
- The Laboratory appears comfortably able and ready to meet its responsibilities in computing and networking for the JLab program, with a close consultative process established to ensure plans and expectations continue to coincide. Equipment provisioning and budget planning are consonant with long standing practice at the Lab and can be expected to be cost effective. The plans to leverage QCD computing resources at the lab for spikes in computing activity in the 12 GeV program (Data Challenges, periods of intense processing demand) is excellent and efficient.

Overall, the Committee was very impressed with the current state of software and computing preparations and the plans leading up to 12 GeV datataking. The Committee was struck by the extremely high quality of the presentations and excellent preparations made for the review. We wish to thank the presenters and all involved in the preparations and logistics for a very successful and, for the reviewers at least, enjoyable and enlightening review.

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, common Experimental Hall issues, Hall-specific issues, and management and computing. Sections are organized as Observations, Findings and Recommendations.

General Comments and Recommendations

Observations

- Presentations were comprehensive and very well prepared, thank you to all.
- The documentation received was much appreciated, special thanks to Hall D for early and comprehensive documentation. Hall B materials were also appreciated, the Committee would have welcomed them somewhat earlier.
- The level of preparation of software at this stage prior to 12GeV startup is impressive. Plans have been developed, communicated within the experimental communities, major components such as offline frameworks implemented and in use, testing and scale-up programs to ready offline systems for startup are underway.

- The extent to which specific plans and milestones were given for measuring the functionality, performance, scalability and readiness of the software in advance of startup varied across the Halls as will be mentioned in the sections following.
- The Committee did not hear a great deal on mechanisms to support the utilization of the software by the full collaborations. We were shown examples of the software in real use for physics and performance studies.

Findings

- There is evidence of broad examination and substantial use of common components, from both the wider HENP and open source communities, and the other Halls.
- Collaborations are undertaking the migration from legacy codes, Geant3 and other Fortran codes, in a managed way, validating new codes against the well-understood legacy codes.
- Software planning in general appears comprehensive and responsibly managed. The Committee did not observe areas appearing strongly at risk due to inattention or absence of planning.
- Manpower for software development appears under control and well managed, without serious deficiencies or inefficient use of the limited existing manpower. Quantitative estimations of FTE-years spent and still required for completion were given, and indicated available manpower levels should be sufficient to meet readiness milestones.

Recommendations

- **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.**

Data Acquisition

Observations

- The data volume is dominated by Hall D (3GByte/s raw, 300MB/s after L3 trigger)
- A new readout scheme is required to cope with 200KHz+ trigger rates: pipelined readout of blocks of data (multiple events) which get built into full events downstream.
- An agent-based experimental control system is under development, with the goal to speed up the start of a run.
- The network will be a mix of Infiniband and 1Gb Ethernet (10Gb Ethernet does not seem cost-effective compared to Infiniband). The front-ends are typically confined to Gb Ethernet. Infiniband interface cards are being recycled from older machines as they are decommissioned.

Findings

- DAQ is supported by an experienced dedicated team, which looks appropriately sized for the current intensive effort to ready DAQ and online systems for the 12GeV program.

- Support is organized through an appropriate hierarchy of experts: General support in Physics and IT, and Hall-based online groups with enough authority and privileges to tend to problems confined to a Hall.
- Kudos on CODA as an established, Lab-supported common online framework and raw data format across the Halls.

Recommendations

- **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.**

Experimental Halls - General

Observations

- Several commendable examples of commonality in software across the Halls:
  - Halls B and D share a common conditions database system CCDB.
  - Hall D uses parts of the Hall B event display system.
  - Hall C is adopting the ROOT based C++ analysis framework of Hall A.
- File management and data discovery by physicists doing analysis was not much addressed. Particularly important for Halls B and D.

Findings

- There is good attention to multi-threading/multi-processing support to accommodate new computing architectures. Event-level, but not subevent-level (or at least not before considering event level) parallelism is being pursued, consistent with trends in HENP.
- Software profiling and performance analysis early and often pays off many-fold. It is important to identify and change bad (C++) coding habits before they become too ingrained. All Halls are giving attention to this.

Recommendations

- **Nightly builds are performed by some; we recommend them for all.**
- **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.**
- **Run a code validation suite such as valgrind as part of the routine software release procedure.**
- **Give full and early consideration to file management, cataloging and data discovery by physicists doing analysis. Report on this area in future reviews.**

Hall A

Observations

The Hall A 12 GeV physics program will use much of the same apparatus as was used in the 6 wGeV era. The notable exception is the introduction of the SBS, which includes GEM trackers that are substantially more complex devices than have been used in this hall previously. SBS won't start operation before 2016.

A C++ based framework was deployed for Hall A analysis in 2003. The framework is based on ROOT, and uses plug-in modules. Modules for reconstructing the spectrometer instruments are fully functional. An SDK exists, and support is provided to users developing reconstruction for custom devices. Online monitoring and calibration can be performed in the analysis framework. Several independent simulation frameworks are supported, ranging from standalone matrix-based codes to full Geant4 simulations. Raw data is in CODA.

Findings

- The existing Hall A analysis framework is well adapted for use in the 12 GeV program. Synergies between the experiments in Hall A and Hall C suggest that a common framework and common code management system may be practical, and may improve the efficiency of users who work on experiments in different halls.
- The C++ analysis framework is single-threaded, but can potentially be parallelized by forking.
- A tracking algorithm for the HRS VDC has been implemented, and has been shown to provide acceptable resolution, but it is at an early stage in optimization and validation. This device is needed for the first physics programs.
- The SBS detector has been simulated, and a prototype track reconstruction algorithm has shown efficiency > 90% under nominal background conditions. Institutional responsibilities for all the software packages are defined, but individual contacts are not yet formalized.

Recommendations

- **Investigate the feasibility of event-based parallelization of C++ analysis in a multi-core batch environment.**
- **Intensify efforts on the HRS tracking development, including calibration and alignment procedures. Define performance milestones which allow time to explore alternatives if problems arise.**
- **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.**
- **Develop requirements for the SBS algorithm performance, along with a development timeline and a responsible contact. Requirements should include alignment and calibration.**
- **Investigate if a move to git along with Hall C makes sense.**

Hall B

Observations

The CLAS12 collaboration is a continuation of the CLAS experiment from the 6 GeV period, with a new detector designed for the 12 GeV era. GEMC is their very flexible Geant4 based data simulation package in use since 2007. GEMC derives its geometry information from a database, thereby making it easy to share one geometry description between the simulation and reconstruction. They have taken the lessons they learned from maintaining the CLAS software and developed a new event processing framework to improve future maintainability. Although the framework was developed in-house, the ideas it uses are based on the existing LHCb Gaudi framework. Unlike Gaudi, the framework is specifically designed to do thread based event level parallelism which will make good use of new multicore machines. A novel aspect of the processing framework is that it is implemented in both Java and C++. This decision was made both to ‘future proof’ the code by not tying it to only one language but also because Java is now a dominant language used in university courses.

CLAS12 is considering whether to adopt nightly builds of their software (as CLAS does)..

In the CLAS era, calibrating was a lengthy process. For CLAS12 the collaboration is actively working on being able to do initial calibration while taking data.

CLAS12 is reusing the calibration database developed by GlueX.

Findings

- The offline software framework is innovative and a departure from common practice. It requires close involvement from and buy-in by the analysis community. Exposure of the plans and progress at collaboration meetings, tutorials etc. over a period of years, as well as an internal review in February of this year was described, which is encouraging.
- A performance penalty of Java and the dual-language framework was roughly estimated during discussion at ~15%. A quantitative measure would be useful.
- The team is only beginning to outline how processing can be done on a batch farm.
- In the CLAS era the collaboration routinely did regression tests using a standard data sample but this has not yet been done for CLAS12.

Recommendations

- **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.**
- **Seriously consider using ROOT as the file format in order to make use of the steady advances in its I/O capabilities.**
- **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.**

Hall C

Observations

Hall C has a large established code base written in Fortran (compiled with gfortran). It is recognized that the framework has to be converted to modern standards, C++ and ROOT I/O. The existing code base is considered a reference, which has to be matched in the new framework. In this process, the current plan is to implement the reconstruction of the new SHMS spectrometer in Fortran to extend the reference/baseline to the SHMS. The assumption is that the analyses of both the HMS and SHMS spectrometers are very similar. In addition, the new detector/readout modules for the SHMS will be implemented in Fortran initially. This is an effort to use the existing software, which has been refined over 15 years, as a stepping stone to the new software framework, and provide (or bring up to date) extensive documentation of the existing software. Using the results from the Fortran-based framework as the reference, the new C++/ROOT based framework and C++ modules will be evaluated. The new framework will be based on the existing (and regarded as more advanced) Hall A efforts.

Hall C has made the decision to use git as the code management system (the only Hall as we can tell, all others use the more classic CVS or SVN).

Findings

- There is self-admitted limited expertise in C++/ROOT by Hall C staff.
- The conversion to C++ requires a large manpower effort.
- New C++ framework will likely not simply map on the existing framework, but will make some re-design efforts necessary.
- A tentative work breakdown has been presented, ending with a C++ implementation of the HMS analysis in the new framework in December 2013. The intermediate steps (selecting reference 6GeV data, DAQ data decoding and raw data level analysis in October) seem aggressive but achievable.
- The expected synergy between Hall A and Hall C software development efforts, with the stated goal of making a transition between the Halls painless, will avoid duplication of effort and code and will streamline the development and debugging efforts.
- The use of standard code evaluation tools does not seem to be an integral part of the development cycles.
- The envisioned online documentation efforts (Wiki, ROOT documentation, and potentially doxygen) seems reasonable.

Recommendations

- **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.**
- **Re-use existing efforts from Hall A to decode CODA-formatted data in ROOT.**
- **If resources are limited, the Fortran-based SHMS reconstruction should be a low priority.**
- **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.**

Hall D

Observations

Hall D is a new hall specifically built for the 12 GeV era. The hall houses the GlueX detector which is new to Jefferson lab. The collaboration has demonstrated the ability to do MC generation, reconstruction and end user partial wave analysis. The full end-to-end test is very encouraging and has shown its utility in uncovering reconstruction software problems, subsequently corrected.

The framework is written in C++ and is decomposed into a series of components. The event processing framework does event level parallelism, which is the appropriate scope for their problem domain.

They do a nightly build of the software, which signals experts in the case of problems. In addition, they have twice weekly regression tests using known MC samples in order to find performance problems in new code.

Calibration and alignment software are in a rather advanced state. Nonetheless completing the development of the calibration software is estimated to be the largest remaining offline software effort to complete in terms of FTE-years.

They have implemented a run based calibration system, CCDB, which can use ASCII or MySQL back ends. The system keeps a full history and makes it easy to create a new calibration 'era' by cloning an existing era and modifying only the calibrations one is interested in. It also includes an easy to use python shell interface.

HDDS as the single source of detector geometry description for MC and reconstruction looks very promising. This guarantees that reconstruction and simulation are in synchronization. In addition it should aid in the migration from Geant3 to Geant4.

Innovative use is made of GPUs for Partial Wave Analysis.

They are evaluating the use of the visualization library built by CLAS12 in order to do event visualization.

Findings

- Workflow tools for handling bulk processing on the batch farm are only in conceptual stage.
- The detector simulation is based on Geant3.
- The collaboration has explored using grid resources; this is worthwhile so long as required manpower levels are low, to improve flexibility and capability in analysis.
- The JANA framework is very stable. They report they haven't changed the code for a year.
- Their JANA framework is not specific to GlueX, it could be adopted by others. It hasn't been thus far.

Recommendations

- **A series of scale tests ramping up using JLab's LQCD farm should be planned and conducted.**
- **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.**
- **Consider ROOT (with it's schema evolution capabilities) as a possible alternative for the HDDM DST format.**

Management and Computing

Observations

The 12GeV software and computing effort does not include explicit management of software and computing activities across the Halls. Each Hall has an Offline Software Coordinator with responsibility for managing the provisioning of the offline software needed by the Hall's experimental program, and for communicating and negotiating the computing requirements of the Hall with JLab scientific computing and physics management. The Offline Software Coordinators are JLab staff and so are fully integrated into both 12GeV program management and line management at JLab. We observed that the JLab physics and scientific computing management pays close attention to the computing and software programs and needs of the Halls.

Findings

- The software and computing management structure seems to have been successful in the past and we see no reason it should not be successful in the future.
- Offline software is not formally a part of the 12GeV project, but is treated as such by the Hall communities. This is appropriate.
- The internal review of one year ago, together with the present review, and the expressed intent to hold similar reviews in the future provide a useful and appropriate level of oversight and external consultation on software and computing. The reviews are also highly valuable as an incentive and driver for software and computing planning and advancement within the Halls.
- The Committee did not have concerns over contingency and risk management. Plans for delivering computing capacity and capability at the level required for the 12 GeV program and with the available (largely static) manpower levels seem well in hand.
- The process for resource estimation, planning and budgeting was described, and a capacity ramp with associated budget plan to reach required capacity levels was presented. The process looks to have appropriate and effective consultation between the Halls, 12GeV program management and JLab Scientific Computing management. Requested resource levels, budget planning, and ability to provision the resources on schedule all look reasonable.
- Planning for CPU provisioning is based on sustaining high (>~90%) utilization levels. This is cost-effective and appropriate but leaves open the question of how the natural peaks in usage demands will be handled. The prospect of using a fraction of the substantial JLab lattice QCD resources for peak handling on a temporary basis -- and for large scale data challenges in advance of startup -- is an excellent one.
- JLab is ready to procure GPUs for the 12GeV program as soon as it is useful to do so.

- On networking, the planned upgrade of the backup link to 100Mbit is timely. The Committee was glad to see plans for a redundant 10Gbit link in the future, which hopefully will become firm.

Recommendations

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

Appendix 1 – Review Charge

12 GeV Software and Computing Review Committee Charge

April 26, 2012

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 data acquisition, planning for computing resources, and management.

The committee is asked to address the following questions:

1) Software and Analysis
   a. What is the state of simulation, data acquisition, calibration and analysis software, including usability and readiness from a user's perspective? Are the software plans complete, and is the scope appropriate?
   b. Is there adequate progress in software maturity, and is there a defined set of goals leading towards full readiness ahead of production running?
   c. To what extent will software tools and components common across the halls and/or with the wider HE/NP communities be utilized? Are efforts towards commonality being made?
   d. Have milestones been identified, and an appropriate set of tests been incorporated into the milestones, to measure progress towards final production running?
   e. Have the collaborations identified effective and appropriate mechanisms to support utilization of the software by the entire collaboration ?
2) Management
   a. Are the current management structures and processes well-matched to the needs of the collaborations (including users)?
   b. Are there appropriate contingency and risk-management processes in place? Have risks been appropriately identified?
   c. Are there adequate plans for transitioning from a development phase into a deployment and operations phase?
   d. Have the required resources been correctly assessed. Are the assumptions of resource requirements well justified? Have the resources been identified? Is the proposed schedule for implementation reasonable?
3) Computing and Networking

a. Are the requirements for computing 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?

Appendix 2 – Presentations

For the full timetable see
https://www.jlab.org/indico/conferenceTimeTable.py?confId=4#all.detailed

12GeV Project Overview	*ENT, ROLF*	*L102-104, Jefferson Lab*	08:45 - 09:15
Detector/Software/Physics Overview	*YOUNG, GLENN*	*L102-104, Jefferson Lab*	09:15 - 09:45
Hall A	*HANSEN, OLE*	*L102-104, Jefferson Lab*	09:45 - 10:15
Hall C	*STEVE WOOD/MARK JONES*	*L102-104, Jefferson Lab*	10:15 - 10:45
Break		*L102-104, Jefferson Lab*	10:45 - 11:00
Hall B	*LATIFA ELOUADHRIRI/DENNIS WEYGAND*	*L102-104, Jefferson Lab*	11:00 - 11:40
Hall D	*CURTIS MEYER(CMU)/MARK ITO*	*L102-104, Jefferson Lab*	11:40 - 12:20
Working Lunch		*L102-104, Jefferson Lab*	12:20 - 13:30
DAQ and Common Online Tools	*HEYES, GRAHAM*	*L102-104, Jefferson Lab*	13:30 - 14:00
Offline Computing (roll up)	*PHILPOTT, SANDY*	*L102-104, Jefferson Lab*	14:00 - 14:30
Networking	*KOWALSKI, ANDY*	*L102-104, Jefferson Lab*	14:30 - 14:50
Break		*L102-104, Jefferson Lab*	14:50 - 15:15
Hall A & C		*B207, Jefferson Lab*	15:15 - 17:15
Halls B & D	*LAWRENCE, David et al.*	*L102-104, Jefferson Lab*	15:15 - 17:15