US ATLAS Software and Computing Review

May 18-20, 2009

Final Report

Table of Contents

Executive Summary

Overall, in Facilities and Software, the panel was presented with evidence of high quality work firmly committed to the success of ATLAS.

US ATLAS is doing at least its share of resource provision, operations and software development. Where US ATLAS is doing somewhat more than its share, this reflects that the US is "doing it right" and some others need to catch up.

The Tier 1 is ready for its planned role in ATLAS production and has a sound service-support model. The staffing level is not markedly different from that which Fermilab would require to perform the same functions. The individual Tier 2s were not presented in detail, but have also demonstrated their readiness for both simulation and data production. The Integration and Operations activities are effective and vigorous. The (non)existence of Tier 3 facilities at many sites and the need for central support for Tier 3s are important issues.

The fractional allocation of ATLAS facilities to analysis tasks appears much smaller than in recent large-scale experiments. Other major ATLAS-wide issues include the scalability and robustness of dcache, the scalability and robustness of the distributed data management system and the need for a policy-driven data-access strategy and its implementation.

Manpower needs, especially in software and operations, will not decrease in the first year of data taking, rather there will be increasing pressure on manpower when the data arrive, making it essential to have processes to prioritize work. Much of the priority work will require horizontal integration of WBS efforts.

US ATLAS Software contributions are highly regarded by, and clearly vital to, ATLAS. However, there are perceived, and in some cases real, issues relating to optimally aligning US efforts with ATLAS priorities and ensuring that the full potential benefit of US efforts is realized in ATLAS as a whole.

Summary of key recommendations:

1. Ensure that adequate management attention is focused on non-Tier-1 issues.
2. Flexible, nimble, and at times decisive management will be required to meet both the challenge of real data and the challenge of working within the context of international ATLAS
3. Ensure that the WBS structure does not become a barrier to efficient execution of activities requiring cross-WBS coordination.
4. Create and maintain crisp, clear, prioritized lists of software activities and tasks.
5. Consider ways to improve communication and collaboration with ATLAS central.

General

We thank the presenters for their excellent presentations and their open and productive interaction with the panel in response to questions.

Findings and Comments

1. Overall, in Facilities and Software we were presented with evidence of high quality work firmly committed to the success of ATLAS.
2. US ATLAS Software and Computing is doing its share (or slightly more than its share) of Resource Provision, Software development, and Operational support. This view is echoed by non-US members of ATLAS. In some areas, such as Operations, the level of US effort is well above 23% but this is likely to decline as other time zones follow the US lead in setting up this support.
3. Manpower needs will not decrease in the first year of data taking, rather there will be increasing pressure on manpower when the data arrive, making it essential to have processes to prioritize work.
4. Readiness for Analysis: Limited-scale tests have been performed successfully. More realistic tests are needed, but full realism is unattainable, both because the human resources required would be too high, and because history teaches us that the best estimates of how analysis will be performed are rapidly rendered invalid by experience with real data.

Recommendations

1. Flexible, nimble, and at times decisive management will be required to meet both the challenge of real data and the challenge of working within the context of international ATLAS.
2. Ensure that the WBS structure does not become a barrier to efficient execution of activities requiring cross-WBS coordination.

Computing Facilities

Findings and Comments

Ensemble of Facilities

1. Analysis capability: comparisons with Run II and BaBar indicate that the fraction of the computing facilities that will be devoted to analysis (10% to 25%) may have been underestimated. BaBar's approximate relative expenditures on Production, Simulation and (computer center supported) Analysis were 31%,26%,43%.

2. The organization of the data and of user access to the data remains unclear (this is not just a US ATLAS issue).

Tier 1 and Networking

1. The Tier 1 Facility is ready and equipped to carry out its role and will meet the pledged resources to the LCG, or come very close to meeting them.

2. We note that the Facility Manager is concerned about Storage Management and believes that
   "there is currently no solution implemented and deployed that is capable of handling the challenging data access requirements that are expected once hundreds of physicists start analyzing real data."
   However, there is little to no effort devoted to either supporting the dcache consortium to fix problems with dcache or investigating longer term solutions. This is a problem given the critical importance of storage management. There are contingency plans for the short term should workarounds be needed.

   A concern was raised that problems could become severe unless policies are enforced about movement of data by individuals (outside official data movement services) and unless the middleware can be "taught" about physics groups and roles of individuals.

3. The Tier 1 center at BNL is appropriately organized and managed as a set of services, with attention to service level agreements and supporting processes for running services efficiently, with sharing and leveraging across RHIC and ATLAS. The service provisioning model is impressive, but will certainly need to evolve to meet as yet unknown needs in the first years of data taking.

4. It is difficult, and likely error-prone, to compare the staffing levels and services with the Fermilab Tier 1 because of the following differences:

   i) Fermilab supports CDF, D0, CMS, several Neutrino experiments, several Astrophysics experiments, the Lattice QCD Facility, the Computational Cosmology Facility and a small HPC facility for Accelerator modeling all within a Scientific Facilities sector of the division (68 people). Shared services are therefore leveraged across many more programs.
   ii) At Fermilab Facility management and support costs are not included in the above, but come out of division overheads and lab overheads.
   iii) Administrative support at Fermilab comes out of division overheads
   iv) At Fermilab Basic OS support, cyber security program and site networking, as well as certain basic services like account management, help desk, Indico, document database, etc. are all provided out of lab overheads
   v) At Fermilab some of the work on WAN, cyber security for open science, grid services development, and storage development is paid for by competitively funded proposals to DOE and NSF, including Open Science Grid. At BNL OSG money does not seem to support the facility operations and grid services. Networking research funds are used at Fermilab to help support CMS WAN monitoring, diagnostics.

   Acknowledging the difficulty of comparing staffing levels between CMS Tier 1 and ATLAS T1 it seems that the staffing levels are not far off what they should be. The BNL T1 is missing some of

the leveraging opportunities and R&D/futures opportunities that the Fermilab T1 has. However, such funds only can be acquired after several years of investment by the lab – say from LDRD funds.

The BNL Tier 1 staffing levels are considerably above the corresponding levels at SLAC. We do not recommend using SLAC as a model at this time.

5. Service R&D is an integral responsibility of any computer center. Whatever the funding level, service R&D must continue at an appropriate level, even if a temporary slowdown due to time-limited budget discontinuities is acceptable.
6. There is a strong case for UPSs and a case for generators behind the UPSs. These cases need to be made more quantitative in terms of the long-term cost effectiveness for the type of science supported by the BNL Tier 1.
7. There are plans in place for BNL to provide the required physical infrastructure for the US-ATLAS T1 center. However, it is still unclear how outfitting of the space with power and cooling and electrical distribution is to be funded. We do not think that the immediate capital costs for these infrastructure upgrades should come out of the Tier-1 funding.
8. Networking appears to be in good shape for the immediate future. The cost-effectiveness of any possible use of project funds to provide additional funding for USLHCNET should be carefully examined.

Tier 2

1. The reviewers did not hear very much about the individual Tier 2 facilities other than the staffing levels at them and their accounted contributions to production. Their specific issues and concerns were not aired. The invisibility of the Tier 2s could be partly attributed to their successful integration into the Tier 1+2 complex.

Tier 3

1. The Tier 3 component of the US ATLAS computing system has received almost no attention (apart from the report mentioned below). The Tier 3s will likely be the only resource available to university physicists for many types of analysis, yet a major fraction of US ATLAS universities have no Tier 3 resources at all.
2. The Tier 3 report has raised expectations that US ATLAS will provide some support for Tier 3s, particularly for access to data. While such support for the university community is very important, it is not clear how it can be provided and what it will cost, since the ATLAS data management strategy is unclear, and the ability of Tier 2s to provide effort in support of Tier 3s is very questionable.

Integration and Operations

1. The Facilities integration program seems to be doing an excellent job and seems well organized and objectives-driven.

2. The Operations organization is doing an excellent job. However, it is not clear why 6.5 people funded by the project for Production and Operations Coordination is the right number. This is an area which straddles physics and facilities and where the staffing mixes physicists (doing shifts) and professionals providing support.
3. Overall, US ATLAS should strike an appropriate balance between monitoring shifts, developing automated monitoring, root-cause analysis of problems and engineering to overcome the problems. Consideration should be given to the most appropriate organizational structure for achieving this balance.

Recommendations

Ensemble of Facilities

1. Ensure that adequate management attention is focused on non-Tier-1 issues, for example:
   - how will university-based physicists have the resources and support to do analysis?
   - clarification of expectations on user-driven access to data and the priorities of the experiment with respect to use of resources.

Tier 1 and Networking

1. If there is a funding shortfall, consider cutting equipment rather than staff at the Tier 1 – don't (necessarily) meet the pledge to WLCG for resources. The disk and CPU requirements have huge error bars and only after data comes will the real needs be known and new pledges, or a different computing model, will be negotiated.
2. Plan to work with others (e.g. a small number of key Tier-1s that are major dcache users) to evaluate options and move forward with replacement for dcache in the longer term.

Software

Findings and Comments:

1. ATLAS must take care to not reduce manpower in this critical area for the near term. In the longer term, the trade off between reducing manpower for maintenance and support should be weighed against new developments for the upgrade, etc.
2. Priorities were difficult to divine from some presentations. This needs to be remedied to make effective review possible and productive.
3. The difficulty of fully integrating US-ATLAS organizational structures and work into ATLAS organizational structures and work was raised as an issue several times. The US ATLAS Software teams are regarded as very valuable, but slightly disconnected collaborators by ATLAS management. This must be addressed by a combination of better communication and a management-led effort to ensure that the US efforts are driven by the true priorities for the

success of ATLAS. It is also not clear that the US ATLAS organization structure is optimized to facilitate communication and integration.

Core services

1. Crucial Core personnel are addressing the Memory Crisis, which takes effort from planned upgrades and maintenance and user support. The memory crisis must be solved for production running as soon as possible, especially considering the expected unknowns with real data. This appears to be a mixture of technical and non-technical issues. The taskforce Paolo Calafiura co-leads with David Rousseau has the mandate to fix the problem and must do so soon so that other core priorities don't languish. A clear plan prioritizing work to fix production code and to down-select Algorithms and histograms from production running, alongside work to develop and deploy diagnostic tools must be defined and implemented and backed by ATLAS management. A clear plan prioritizing work to fix production code and to down-select Algorithms and histograms from production running, alongside work to develop and deploy diagnostic tools must be defined and implemented and backed by ATLAS management. A clear plan prioritizing work to fix production code and to down-select Algorithms and histograms from production running, alongside work to develop and deploy diagnostic tools must be defined and implemented and backed by ATLAS management.
2. The tools developed and demonstrated did not appear adequate to address the current and inevitable future optimization tasks. Tools for profiling CPU and memory usage need to be developed now so that any future optimization tasks can be dealt with in real time.

Data management

1. Data management has many stakeholder requests they must prioritize and address for the short and long term. The imminence of real data must focus attention on the most crucial issues of getting data to collaborators reliably and promptly. The highest priorities for the Data management team were difficult to identify. High priority issues like metadata handling and TAG DBs seemed to occupy similar mind-space with upper ATLAS management with low priority issues such as the Debug Stream. This may reflect collaborators' lack of clarity on priorities and trade-offs. The question of file size seemed to be solved, though it was unclear that this was effectively communicated to collaborators. A prioritized list of Data Management's 5 most urgent tasks for data startup should be made and used to drive effort.

Distributed software

1. PanDA has been officially accepted by global ATLAS and is in use in all ATLAS clouds. This is a good development, but we note some vestiges of other systems which need to be integrated or accommodated. Indeed, active use of competitive products continues. The current phase of PanDA should involve embracing other experts and concepts in distributed computing from within world-wide ATLAS. Seeking global ATLAS development effort will help in the short and long term with continued increasing acceptance of PanDA. Global ATLAS should explicitly find manpower for operations and for development in this area who can work effectively with the very productive BNL team.

Application software

1. We did not hear specific details about this WBS element. We encourage US ATLAS management to closely monitor this effort to ensure the proper usage of project *versus* base funds. This effort on project must be very tightly coupled with the core developers, even while addressing sub-detector-specific tasks. Efforts which properly belong on base funds should not be taken up by these developers.

Infrastructure support

1. We did not hear specific details about this WBS element. Certainly some US effort on Infrastructure support is necessary. To properly assess this effort we would need more information. The level seems of the right magnitude.

Analysis support

1. This WBS element is now a Level 2 position reflecting the increasing importance to the experiment of this effort as more physicists become active and as real data is about to flow. The amount of effort required is likely to increase, but should be considered a sub-responsibility of all US ATLAS software and computing groups. Specifically, developers responsible for critical components used by scientists, such as pyROOT and AthenaROOTAccess, will by necessity spend part of their effort doing this kind of work. The responsibility of the Analysis Support L2 will in large measure be one of coordination and prioritization of many efforts in US ATLAS.

Recommendations:

1. In general crisp, clear, prioritized lists of activities and tasks will facilitate communication of Software's efforts to both global and US ATLAS management.
2. Consider ways to improve communication and collaboration with ATLAS central – such as:
   - restructuring to better match roles in US-ATLAS with roles in ATLAS,
   - following negotiation with ATLAS management, make clear formal statements about what the US can and can't support and what the expectations are on both sides;
   - formalizing decision making in US-ATLAS and communicating those decisions and their rationale to International ATLAS,
   - publishing weekly US-ATLAS status bulletins?

Core services

1. Address the technical and non-technical aspects of the memory crisis in the context of the PMB. Priority must be on robustness of production and timely resolution so that core effort can be refocused on languishing core developments.
2. The PMB plan we were presented with should be diligently executed and communicated to ATLAS collaborators.
3. Develop and/or adopt more detailed tools needed for monitoring of performance in the short and long term. Use tools to keep ahead of such issues.

4. Develop a prioritized list of Core's highest priorities driven by physics priorities to define the coming year's work and make clear trade-offs when firefighting decisions are made.

Data management

1. The role of TAGS and metadata in ATLAS data management must be clarified and validated with respect to physics priorities.
2. At this late date, critical priorities for Data Management must be driven by the short term physics tasks and goals related to LHC startup. More explicit and trackable connection between physics tasks and data management features would help focus effort.

Distributed software

1. Explicitly request Global ATLAS provide operational and development support for PanDA.
2. Make sure that all commonalities between distributed production and distributed analysis are exploited to avoid unnecessary duplication of effort.

Application software

None.

Infrastructure support

None.

Analysis support

1. Because this effort is both important and difficult to scope, constant attention to ensure that all aspects are covered is required. For instance, core developers' support responsibilities for analysis must be a recognized, critical contribution and may grow.
2. Pay particular attention to this new effort as it may take some time to gel.

Appendix A: Review Charge

Preamble

This is a US-ATLAS-commissioned review of Software and Computing. The review report will be sent to the US ATLAS Management.

Charge to the Review Panel

Analyze current activities and future plans of US-ATLAS software and computing in terms of:

- Requirements arising from the agreed US role in ATLAS and the WLCG;
- The importance of each activity in delivering high quality physics results;
- The cost-effectiveness of each activity;
- The technical quality of each activity;
- The effectiveness of the management processes involved in each activity – is the organization in place optimal or should it be changed to make it more effective?

Make recommendations on improvements and optimizations, noting particularly the lowest priority activities that are in the current plan, and any higher priority activities that are not in the plan.

Examples of specific topics covered by this charge include:

- Is US ATLAS Software and Computing doing its share, more than its share, or less than its share of ATLAS software and computing?
- Is the Tier 1 appropriately staffed, in terms of staff numbers and skill sets?
- What are the staffing needs of the Tier 2s and Tier 3s in the operational phase, including staff needs to support physics analysis? How do these needs compare with current staffing?
- Is the Tier1/2/3 complex ready for physics analysis (in addition to production processing)? Is the complex flexible enough to deal with evolution of the computing model driven by experiences with real data?
- Are the analysis support activities appropriately resourced and are they poised to deliver efficiently the support required as data arrive?
- How should the US software effort make the transition from construction/development to operation? Which activities should shrink or move into maintenance mode? Are there software activities that need to grow? In the longer term, what are the likely needs for, and timing of, major software work, to prepare for the Phase 1 and even Phase 2 luminosity increases, or to re-engineer software components to meet long-term performance and maintainability requirements?
- Are the developments required to stay cost effective in the evolving hardware and software environment receiving adequate attention and resources?
- Comparisons with US CMS where possible (recognizing that the comparisons will require some effort from US CMS and that this effort cannot be insisted upon.)

Appendix B: Review Panel Membership

Peter Elmer / Princeton

Alexei Klimentov / BNL

Richard Mount / SLAC (Chair)

Ryszard Stroynowski / SMU

Craig Tull / LBNL

Vicky White / Fermilab