Report of the 2006 U.S. ATLAS

Computing Advisory Panel Review

Held on January 5-6, 2006
Brookhaven National Laboratory

Peter Elmer, Ian Fisk, Tom Hemmick, Richard Mount, Chip Watson (editor)

Executive Summary 3
1 Introduction 4
2 Project Organization and Management 4
2.1 Change Control 5
2.2 Management Reserve (Contingency) 5
2.3 Project Tools 6
3 Software and Analysis Support 6
3.1 Data Challenge 2 6
3.2 PanDA 7
3.3 Computing System Commissioning 7
3.4 Performance and Architecture 7
3.5 Software Support 8
4 Grids and T1 / T2 Facilities 8
4.1 Tier 1 & Tier 2 8
4.2 Grids 10

Executive Summary

U.S. Atlas Computing is making very good progress towards detector commissioning and operation, now only a little over a year away. Significant progress has been achieved on all fronts, from facilities to software development. Software is maturing, facilities are growing; the network bandwidth issues have been significantly addressed, and plans for analysis support have made real progress.

During the data challenge and the run up to the Rome physics workshop, U.S. Atlas contributed more than its defined share towards production goals, with healthy production rates at both the tier 1 and the tier 2 centers, and with a higher quality of operation (percent success) than achieved by other regions.

The U.S. continues to provide strong software leadership, particularly in core software and data management. Overall good progress is being made, and no serious problems were identified. Last minute scope creep remains something to be guarded against.

PanDA (a new Production and Data Analysis tool) emerged from lessons learned in DC-2, and from significant concerns about Atlas prodsys scalability. This tool addresses those scalability concerns, and insulates Atlas from the immaturity of higher level grid components. The rapid pace of its development, and the probability of its uptake by International Atlas, are both good signs of the competence of the development team and of management for re-prioritizing effort.

Tight budgets remain (as always) an area of concern. The management reserve is sufficient to cover the shortfalls in software and computing for FY2006 and FY2007, but only under the assumption that nothing goes wrong (i.e. no other calls on the reserve). This is not true for subsequent years, when the additional 30% overhead rate at BNL (from which the project has been exempted to date as a construction project, and may or may not be exempted for the coming 2 years) will almost surely be in effect. Thus, the project is carrying no well defined management reserve or contingency capable of dealing with unexpected expenses. This should be addressed.

Project management should utilize change control more seriously to improve, among other things, the process of how this budget shortfall will be managed.

A modest number of performance challenges are actively being worked on, and will need to be watched closely. Among these are: ESD size, AOD size and performance, simulation time, dCache scaling, and job rate scaling.

Plans for analysis software support have made progress since last year, but more work is needed to flesh this out and put it into practice.

1 Introduction

The 2006 Computing Advisory Committee meeting was held at Brookhaven National Laboratory on January 5-6, 2006. The committee was asked to look at the following topics:

- Project Organization / Management
- Software and Analysis Support
- Grids and Facilities
- Participation in new Funding Initiatives

The short time for the review (one day plus responses to questions) did not allow the committee to address all subtopics in the charge. In particular, we were unable to address issues related to funding priority and allocation, especially the split between M&O and the Research Program. However, we were able to look at priorities and staffing within software and facilities, and have some comments in these areas.

The first day of the review was primarily presentations with question and answers, running from 9:00 a.m. until 6:15 p.m. A number of requests for additional information were given to the project at the end of the day, and responses were presented the next morning in a 1½ hour session, with additional discussion. A verbal closeout was held at 12:30 p.m. on the second day.

Three of the review committee had participated in previous reviews, two were new to the review committee. This level of continuity is essential to perform an adequate level of review in a two day meeting, and should be sustained into future reviews.

Background material was posted online for the committee only just before the review, and talks were not available before the meeting. The web sites were extremely helpful and convenient, however it would be most helpful to have this material available sooner.

The talks were well done, and the presenters exhibited great willingness to answer questions and help us in understanding the state of various aspects of the project (essential for such short reviews). We commend all the presenters and organizers for their hard work.

2 Project Organization and Management

Overall we feel that US Atlas has a good management team in place, and that the project is progressing at a healthy rate. We observed good interactions among the various levels of management and staff at the review, perceived no unusual problems in interactions with International Atlas, and thus anticipate continued success in the management of this project.

The greatest challenge to management remains the tight budget situation, and most of the comments in this section are intended to improve management's ability to deal with this challenge.

We were asked to comment on relative priorities with respect to the M&O portion of the project, but we did not receive sufficient information nor have sufficient time to address this split.

2.1 Change Control

Two years ago, the PCAP recommended that U.S. Atlas institute a formal change control mechanism. This was felt to be important to control scope creep in a project whose staff were already oversubscribed. By early 2005, a process was in place but was not being used, and the project was again urged to commit to a change control process. Already evidence of scope creep problems were apparent.

During the past 12 months, the change control process has been exercised once. However, the one change was in fact multiple changes rolled into one "omnibus" change, which is not a good use of the process.

Several other significant changes were executed in the past year, two of which should have triggered change control:

- The WBS underwent an important revision. Last year we were informed that the WBS needed updating, and the nature of the changes presented this year were quite reasonable. But, they should have triggered change control.
- A major new software task and WBS element were added: PanDA, Production and Data Analysis. We concur that in the end this tool will save (unbudgeted) FTE which would have otherwise been needed for production running, and so, in some sense, represents a cost savings. While the decision making process that led to this involved management at the highest level, it was a significant change and increase in scope for U.S. Atlas, and should have triggered change control.

While we did not examine all of the documents defining U.S. Atlas change control, we should point out that the table we were shown did not contain concrete levels for triggering the process. "Change of scope" did not have an associated level of effort (clearly insignificant changes should not be processed), and "Change in cost" did not have a dollar value included.

Recommendation: Refine the triggers for change control, and use the process!

2.2 Management Reserve (Contingency)

The budgets for Software and Computing presented to us showed a U.S. Atlas Research Program with a management reserve of scale $3M – 15% in FY 2006, falling to 10% in the out years. However, the project also showed budget shortfalls which would completely consume this reserve. Thus, the project has no real reserve, and no well defined mechanism for dealing with unexpected expenses or budget cuts other than a full blown de-scoping exercise.

Moreover, there is an anticipated 30% increase in the overhead rate at BNL as the project changes from a construction phase to an operations phase. This known and expected cost is not currently reflected in the budget planning for the tier 1, and represents an additional ~$1M expense which cannot be met even by consuming all of the management reserve.

We were told that the overhead rate is under discussion with management for the interim phase of the Research Project prior to LHC turn-on. It was quite clear, however, that everyone expected that once Atlas was fully operational, the overhead would return to the BNL standard overhead rate (30% higher than current).

Recommendations:

(1) Change the Tier 1 budget plan to correctly show the anticipated overhead rate for operating projects, at least for FY2007 and beyond, and also for FY2006 if negotiations don't produce an agreement for a continued exemption as a construction project.

(2) Produce a budget plan in which the management reserve is not used to balance a budget deficit. Clearly identify which tasks or procurements are not within the funded scope, and will be funded out of the management fund should those funds not be consumed by other unexpected expenses.

2.3 Project Tools

The management team is using MS Project and Access to hold project and WBS information. One minor observation is that the project is not resource loaded. Details of loading are held at a lower level in the project, as for example in the staffing plans for the Tier 1 center. Additionally, priorities were not clear, in that the budget planning did not indicate which items would be funded only if the management reserve was otherwise untouched.

3 Software and Analysis Support

The committee was pleased and impressed by the progress made since the last review in the areas of software and software support. Among the highlights of the efforts are the establishment of a concrete format of the AOD data files and the impressive implementation of the experiment geometry, leading to a functional event display. Despite the additional computing needs of the combined test beam and high priority developments on the conditions database, all major milestones were met.

3.1 Data Challenge 2

Among the most significant events in the past year was the Rome Productin, the last phase of what began as Data Challenge 2 (DC2). We were pleased to see that US Atlas accomplished much more than the 20% of the total computing that would be their charge. The large contribution to the grid production from the US-Atlas and tier-2's and smaller sites is a nice confirmation of the computing model.

In addition to (and arguably more valuable than) the accomplishments in mock data manipulation during DC2 are the lessons learned. Large scale computing on the grid during DC2 and the Rome Production highlighted those parts of the grid software that did not provide robust and reliable operation. Further, the single massive test was found to not provide a sufficient basis for smooth code development. We are pleased to see that US Atlas has responded positively and aggressively to the lessons learned, both by developing the PanDA software, and by restructuring the next series of software tests.

3.2 PanDA

DC2 and the Rome Production were of singular importance in the development of US Atlas software since they served to distinguish those parts of the grid software that were performing adequately for the needs of US Atlas from those that were not. This experience led to the design and (thus far) impressively rapid development of the PanDA distributed analysis system. PanDA internalizes those parts of the grid toolset that proved to be unreliable, while maintaining at the front side a standard executor interface to International Atlas and at the backside a use of a restricted set of well established and reliable grid tools.

The committee feels that the decision to develop PanDA was not only well justified, but also well timed, having started roughly 18 months prior to first beam. This timing is a good compromise between maximally leveraging/furthering the OSG grid development and leaving enough time to develop a robust tool for day one physics.

PanDA, in addition to providing a robust connection to the OSG grid, also integrates tightly with the Atlas Distributed Data Management (DDM) software by intrinsically implementing Don Quixote 2 (DQ2). PanDA has naturally become a key effort and achieving the present level of development has required several FTE. Although the enthusiasm of the authors in presenting the details of the PanDA system is understandable, they should be cautioned to avoid giving the impression that this is a deviation from grid technology or from the OSG effort.

3.3 Computing System Commissioning

The experience during DC2 and the Rome Production has also lead to a new paradigm for the final preparation of the software. According to the previous schedule, DC2 would be followed this year by Data Challenge 3 as the next major milestone in software development. Instead the model has been changed to replace DC3 with the Computing System Commissioning (CSC), consisting of a longer ordered set of smaller milestones to avoid the system integration issues which can arise from many sub-systems delivering just before a single “big bang” challenge. While this does represent a better strategy for software commissioning, US Atlas should remain keenly ware of the risks associated with delays of the full software system test.

Recommendation: Review the schedule of milestones, and ensure that sufficiently large scale testing is included to reveal any hidden scaling and/or integration problems.

3.4 Performance and Architecture

The committee is pleased to see that some initial work has been done on understanding and improving the I/O performance, especially for the AOD and ESD files. Further effort on optimizing this performance, including significant changes to simplify the data model if need be, is likely to be time well spent. As seen by many other experiments, poor read performance, in particular for the AOD, often leads to ad-hoc user and analysis group attempts to reimplement more performant versions, usually decoupled from the full eventstore data. The explicit exploration of Athena-aware ntuples as a means of managing derived physics data and preventing decoupling from the experiment framework also seems promising.

Regarding the possibility of decoupling the transient and persistent representations of the event data to simplify the schema evolution problem: while this would clearly have the stated benefits, care should be taken to insure that it doesn't put an ultimate cap on performance. The simplification of the data model and EDM classes themselves might itself lead to some significant gains, leaving the introduced transient/persistent split as a bottleneck at some point.

The Atlas wide reconsideration of where data streaming will happen in the data reduction chain, combined with reported discussions about the appropriate use of tag data make it somewhat unclear what the future role of the tag database (event level metadata) will be within the Atlas analysis model. Depending on where the Atlas-wide discussions wind up a re-evaluation of the value of the tag database and the level of effort placed there could be wise.

Effort must continue in defining a clear picture of the analysis model. This model should show explicitly and simply how an individual users physics analysis is accomplished, which tools they interact with directly, and what mechanisms are handled on their behalf behind the scenes.

3.5 Software Support

The committee was pleased to see that specific plans have been formulated for a physics analysis support organization and to establish Software Support Centers. The support center model is (appropriately) vague at this point, since no such facilities have been attempted before, but has more substance than the very initial ideas presented one year ago. The trial and error approach to achieving functionality and vitality of these centers is thus appropriate and must be started soon. The staff of the Software Support Groups should be assembled as quickly as possible. This staff should then be assigned a first task to begin designing how the system can work. Particular emphasis should be paid to the incentives required for attracting effort (volunteers) and defining/negotiating the level of effort appropriate to this task.

4 Grids and T1 / T2 Facilities

4.1 Tier 1 & Tier 2

There has been substantial progress in the facilities area both at the Tier-1 and Tier-2 centers during 2005 in both the scale of the facilities and in the services offered. The Brookhaven Tier-1 facility has nearly tripled in processing capacity during 2005 and has grown by a larger factor in disk based storage. The facility staff has grown to 11 FTE, which was the target for 2005. The first three US-ATLAS Tier-2 sites were chosen during the fall of 2004 and have made impressive progress during 2005 toward being effective computing resources for the ATLAS experiment.

The committee was pleased to see the successful large-scale deployment of dCache at the Tier-1, which was outlined as a goal during last year’s review. Operations experience with disk based mass storage is growing at the Tier-1 and Tier-2 centers. Disk storage is an important element of the US-ATLAS computing model. The committee was also

pleased to see the deployment of utilized grid interfaces at the Tier-1 and Tier-2 centers. These grid interfaces have been used to validate the PanDA prototype.

The news of the long awaited network upgrade for Brookhaven is exciting and eliminates a potentially serious problem for the ATLAS computing model. The network plans for the Tier-2 centers appear to be in good shape. The available network to the Southwest Regional Tier-2 center should be monitored, because it is not at the level of the other US-ATLAS Tier-2s and the Southwest center has a large number of computing resources available.

The Tier- and Tier-2 centers made an excellent contribution to the ATLAS data challenge (DC2). The US more than met their obligations to the experiment in terms of produced event count and several of the most productive sites were US Tier-2 sites.

The Facilities project appears to be following the defined schedule and achieving their goals. The only visible area of delay appeared to be the potential for a slip in the selection of the final Tier-2 centers due to uncertainties in the available funding. This should be monitored, but it not worrying currently.

The committee commends the hard work of the facility groups both at the Tier-1 and Tier-2 center. The final three years before the start of the physics run will be hard work. Over the next three years, there is a factor of ten increase anticipated in the processing resources and a factor of twenty increase in the disk storage resources.

Recommendation: While the deployment of the grid interfaces and a large-scale dCache is encouraging, the scale at which both of these components have been stressed remains relatively modest. The system has been demonstrated at approximately 10TB per day for write access and 5TB for read access. We recommend the facility teams at all centers perform stress tests of the storage system that allow them to demonstrate an achievable ramp of improving performance with milestones. The operations ramp should guide the facility from the current situation to the requirements of the ATLAS computing model at the start of operations.

Recommendation: The same recommendation for stress tests is made for the grid interfaces. The Tier-1 center has used 100 CPUs through the grid interfaces on average over the last year. The vast majority of the computing resources have been given to local users. Given the number of jobs expected to be submitted through executors like PanDA, it seems prudent to exercise the grid interfaces at a high level looking for scaling and reliability issues.

The US-ATLAS Tier-1 appears to be reaping many benefits from the co-location with the RHIC Computing Facility (RCF). This is a very positive situation and a concrete demonstration of the economy of scale when deploying computing at existing facilities. The Tier-1 center is enjoying mass storage support through the core program, but there is support in other areas as well. We encourage US-ATLAS to do a careful accounting of the resources available through the core program. The purpose of this is to demonstrate to the funding agencies the benefits enjoyed at BNL and to quantify the risk if the RCF suffered a significant cut.

In addition to effort and expertise from the core program, the Tier-1 center currently enjoys a very favorable overhead rate as a construction project. As the project moves to

operations this rate will transition to the normal BNL overhead rate and the effort cost will increase by approximately 30%. The current Tier-1 effort for operation is calculated based on the favorable rate and we recommend for the out years that the effort cost be recalculated with the expected overhead rate.

There is an issue to consider with regard to facility planning. Currently the ATLAS predictions for simulation and reconstruction are several factors above the nominal values outlined in the computing model document and the computing resource predictions.

Recommendation: US-ATLAS should do some contingency planning if the optimization of reconstruction and simulation do not yield sufficiently fast applications. The impact on the facility infrastructure and the and overall computing capacity should be considered.

The current US-ATLAS plan calls for deploying the computing resources needed in 2008 using the FY08 budget. Given the large increase in computing capacity between 2007 and 2008, roughly a factor of three in both CPU and disk resources, and the fact that sites rarely have fiscal authority to spend on the first day of the calendar year; there is the potential for slipping the deployment of properly commissioned resources early in 2008.

Recommendation: We recommend US-ATLAS develop a plan and schedule for the deployment in 2008 with a realistic timeline for procurement and commissioning.

4.2 Grids

Use of grid resources is planned to be an integral part of ATLAS computing. In Europe, most Tier-1 centers will serve several LHC experiments and will provide resources via the grid and more specifically the LCG version of the grid. In the U.S., computing resources will also be exploited via grid middleware. Although the majority of resources used by ATLAS today can be classified as ATLAS-dedicated, the greatest flexibility in access to non-Atlas resources in the U.S. will be achieved by using OSG-supported middleware, under the expectation that this will be used by a wide range of sciences including the majority of high-energy and nuclear physics experiments.

The connection to “external Grid Projects” was addressed by the presentations. The PanDA project started this year goes a long way towards de-coupling U.S.-ATLAS from the research aspects (and near term risks) of these projects, a situation that the committee finds appropriate. Nevertheless, these projects may produce future valuable enhancements, and it is appropriate for U.S. Atlas to remain engaged so as to maximize potential benefits. We observe that U.S. Atlas is appropriately represented in the LCG committee structure. The ATLAS Computing Management is committed to support and exploit the U.S. Open Science Grid (OSG), including a proposal to participate in the OSG executive team. In order to exploit the non-Atlas resources it will be necessary to ensure that the OSG software infrastructure be adequately mature at LHC turn on. Presentations stressed the need for ATLAS to see a net positive effect from its participation in OSG in a timeframe relevant to the start of LHC running (manpower concern).

The panel noted that there is still significant unease about ATLAS-OSG relations on the part of many ATLAS and OSG members.

The panel believes that the formal ATLAS participation in OSG management is on the right track and that the current level of ATLAS activity within OSG is appropriate.

Recommendation: ATLAS should pursue a success-oriented approach to participation in OSG management.

Were the panel to be invited to offer advice to OSG, it would recommend that OSG commit to achieving a net positive impact on the LHC experiments early in LHC running. This goal is of vital importance to ATLAS and to the success of OSG. As an integral part of OSG management, ATLAS should be well positioned to ensure the achievement of this goal.