U.S. Department of Energy
and the
National Science Foundation

DOE/NSF U.S. LHC PROGRAM OFFICE

REPORT ON THE JOINT DOE/NSF REVIEW OF U.S. LHC SOFTWARE AND COMPUTING

HELD AT THE UNIVERSITY OF TEXAS, ARLINGTON
17-19 JANUARY 2007

REPORT TO THE U.S. DEPARTMENT OF ENERGY AND THE NATIONAL SCIENCE FOUNDATION

TABLE OF CONTENTS

EXECUTIVE SUMMARY 3

1. INTRODUCTION 5

2. U.S. CMS 6

3. U.S. ATLAS 12

4. APPENDICES

A. Charge to Committee 20

B. Committee membership 23

C. Agenda for the meeting 24

D. Answers by U.S. ATLAS

Questions Prior to the Review 27

Questions During the Review 58

E. Answers by U.S. CMS

Questions Prior to the Review 72

Questions During the Review 89

EXECUTIVE SUMMARY

The external reviewers of the U.S. CMS and U.S. ATLAS Software and Computing (S&C) programs commended the U.S. teams on their progress since the previous comprehensive review of March 2006, as well as for the clear way they addressed their charge and presented the status of their programs.

Both U.S. S&C teams have sound management structures, and they are well represented in leadership positions in their respective collaborations. They are addressing readiness for startup through a series of well-planned data and service challenges. Nevertheless, it is important that managements establish specific milestones for tracking this year's activities, with special focus on system and integration tests. The grid-production and analysis capabilities demonstrated by U.S. ATLAS and U.S. CMS during the review at UT-Arlington indicated that, although the software is still under development, it already has the essential functional elements for usability during data taking.

On the matter of the impact of 10% reductions in funding guidelines, the reviewers concluded that this would significantly diminish the flexibility of the collaborations to respond to unanticipated problems before and during LHC startup. A serious concern is that, in order to retain flexibility, both groups would have to de-scope their computing facilities and reduce user support, impacting adversely the physics-analysis productivity of the U.S. The collaborations were urged to maintain flexibility and avoid expanding scope so as not to strain the available resources. The reviewers recommended that U.S. scope be clearly defined in Memoranda of Agreement (MoA) within their collaborations.

A concern was voiced about U.S. CMS, with the reviewers stressing that any commitments to scope must be reconciled with available resources and through inclusion of software tasks in the MoA with CMS. The latter, by defining and crediting physicist participation in S&C operations tasks could provide additional personnel from other elements of the collaboration, freeing expert software personnel for maintenance and development.

In addition to the LHC Physics Analysis Center (LPC), U.S. CMS has established an effective organization for supporting user needs in software. Nevertheless, the reviewers noted that one area needing improvement is dCache, whose performance should be optimized for the CMS challenge "CSA07." It was also recommended that U.S. CMS develop and deploy the appropriate QA tools in order to optimize overall software performance. Such tools could also help monitor performance at the computing facilities.

The overall plan for U.S. CMS Tier-1 and Tier-2 facilities is reasonable. However, there is concern that assumptions about the trans-Atlantic transfer of data from Tier-2 to Tier-1 centers may be incompatible with plans for the bandwidth expected for LHCNet. U.S. CMS was therefore urged to increase the scale of testing of its Tier-2 analysis system in order to validate its computing model.

U.S. ATLAS demonstrated good progress in the ability of physicists to make use of grid services with the PanDA production tool. U.S. ATLAS should continue to develop PanDA, and pursue plans to engage the rest of the ATLAS collaboration in support of operations. The reviewers had concern about the ATLAS Distributed Data Management (DDM) system, and urged that a plan for DDM be developed and executed with urgency. U.S. ATLAS should also work closely with ATLAS to reconcile the projected event size with the assumptions of the ATLAS computing model, and evaluate the potential impact and cost of a possible increase in the size of an event. U.S. ATLAS should develop and implement proper monitoring tools so that it can do a better job of reporting the rates and nature of job failures with PanDA.

U.S. ATLAS is to be commended for investing considerable resources in the support of analysis activities of its users. However, U.S. ATLAS should assess the workload of support functions on key developers in order to limit the impact on essential software development.

Good progress was reported on production activities at the U.S. ATLAS Tier-1 and Tier-2 facilities. The ATLAS model of locating the complete Event Summary Data (ESD) at the BNL Tier-1 center will facilitate analysis, but at the risk of having a single point of failure. U.S. ATLAS should evaluate this risk and develop a contingency plan with BNL that would include mitigation strategies for possible cyber-security incidents.

Regarding infrastructure at facilities, the network connectivity plan for U.S. ATLAS sites is adequate down to the Tier-2 level. In addition, although resources are available to satisfy near-term computing needs, U.S. ATLAS must develop a plan to address the predicted shortfall in Tier-2 resources in 2010 and 2011.

1 INTRODUCTION

Software and Computing will serve as a key element for enabling the science of U.S. physicists at the LHC. The large data rates and the distributed nature of collaborative analyses pose logistic challenges of access to data that can only be tackled with effective and broad software and computing infrastructure. A viable system is particularly important for U.S. physicists, who must overcome the disadvantages of geographic separation from their experiments.

The U.S. Department of Energy and the National Science Foundation, partner agencies supporting U.S. participation in the LHC, have recognized the importance of Software and Computing (S&C) to the success of the U.S. investment in the LHC. The U.S. collaborations have established S&C programs within the context of their Research Program Management Plans, and, in order to monitor progress and establish guidelines for support, the agencies regularly review the work of the collaborations.

At the request of the DOE/NSF Joint Oversight Group (JOG), a review of the ATLAS and CMS S&C programs was held on January 17-19, 2007 at the University of Texas at Arlington (UTA). The review covered the general areas of Management and External Interactions, Facilities, Grids, and Infrastructure, and the Support of Software and Analysis. A charge to the review committee (Appendix A) was prepared by the JOG, and distributed to the panel members and the collaborations prior to the meeting at UTA. The members of the review panel were chosen for their expertise in software and computing, and came from U.S. universities, U.S. laboratories and European institutes, all from outside the LHC community. The list of participating panel members is given in Appendix B. The agenda of the review is reproduced in Appendix C. The presentations by the collaborations and additional supporting material can be found at the review web site http://www.uta.edu/physics/meetings/hep/doe/01172006/index.html.

The reviewers and the collaborations were instructed to address the progress in S&C to-date, and to assess the cost, scope, and plans for the period FY2007–FY2011. The reviewers were asked to pay particular attention to the collaborations' preparedness to support physics analysis in the U.S., including the usability of developed software analysis tools. This report, prepared by the external reviewers and edited by the agencies, summarizes the findings, observations, and recommendations of the panel members. The observations and recommendations are based on the presentations, on background material provided prior to the review, on detailed discussions during the parallel sessions, and on answers to questions posed by the panel (Appendices D and E). The panel members also deliberated on their observations and recommendations in executive sessions in the presence of representatives from the DOE and NSF, and then presented their preliminary findings in close-out sessions to the collaborations.

UTA, as host for this meeting, provided invaluable logistical support, which contributed greatly to the success of the review.

2 U.S. CMS

2.1 MANAGEMENT

The U.S. CMS management team for Software and Computing, represented by Lothar Bauerdick et al, gave a comprehensive overview of the U.S. CMS effort along with written responses to pre-selected questions from the agencies.

2.1.1 KEY FINDINGS AND OBSERVATIONS

The U.S. CMS team had well-prepared presentations and was responsive to questions from the committee. There is now a new international CMS management structure in place that is targeted towards moving the experiment into an operational mode. This process appears to be converging and the new structure is appropriate for the task at hand. The organization appears sensible, although many appointments remain open, which could be a concern, given the need of maintaining a stable and effective organization all the time. We observe that the U.S. is well represented in the new international management structure, with an appropriate mix of university and FNAL personnel in positions of leadership. A goal of the CMS management structure is to provide well defined U.S. CMS S&C deliverables, and we commend U.S. CMS for their commitment to the success of the experiment.

The major issues for experiment readiness are being addressed with a series of well planned and incrementally more demanding challenges and tests.

We commend the completion of the Physics TDR, which included many contributions from U.S. physicists, and encourage a continued and appropriate participation by U.S. CMS physicists, including FNAL scientific staff, in such involvements. The physics TDR (Vols. 1 and 2) utilized about 200M generated and analyzed events in the CSA06 challenge. The latest iteration of software and infrastructure was tested successfully in CSA06, which met all its major milestones, with the Tier-1 and all 7 U.S. Tier-2 centers participating. The test also demonstrated interoperability between the U.S. OSG and the European EGEE grids.

U.S. CMS S&C management has proved to be very effective in allocating its resources and in addressing technical issues as they arise. The team is currently focusing on achieving a sustainable operations model with an expectation that physicists can be integrated into operations tasks. There is a plan to allocate between one-half to one week per month to global data taking from now on, to further exercise new code releases.

There is an active effort to assure that the LPC remains an effective organization, and that the LPC and the Fermilab CAF provide the resources needed by U.S. CMS physicists.

While U.S. CMS has analyzed the impact of a possible 10% reduction in funding, and has developed a mitigation plan, the committee finds that any funding reduction would remove needed flexibility, leading to added risk for addressing unanticipated problems.

2.1.3 RECOMMENDATIONS

The committee makes the following recommendations:

- We recommend that U.S. CMS S&C work with CMS to define and credit physicist participation in software and computing operation tasks.

- We recommend a timely process for including software tasks in a Memorandum of Agreement (MoA) between U.S. CMS and international CMS.

- We recommend that U.S. CMS S&C specifies its scope and reconciles it with its available resources. This should be expressed within the MoA.

- We recommend retaining flexibility within U.S. CMS S&C to address the urgent problems that will inevitably arise during the year before data-taking commences.

- In order to facilitate integration, we recommend that U.S. CMS develop and publish milestones that reflect the system functionality required to support the experiment and data collection as a function of time.

- U.S. CMS should continue to work with international CMS management to fill the positions in the new CMS Computing organization in a timely way.

2.2 FACILITIES AND GRIDS

The material relevant to the part of the charge pertaining to U.S. CMS facilities and grids was presentated by Lothar Bauerdick, Patricia McBride, Ian Fisk, Darin Acosta, Ruth Pordes, Sara Eno, John Bakken, Ken Bloom, Frank Würthwein, Don Petravick and Oliver Gutsche. There was also some additional material made available prior to and during the review.

2.2.1 KEY FINDINGS

U.S. CMS demonstrated adequate progress on all recommendations made at the March 2006 review. During that review, it was noted with concern that the total Tier-1 resources pledged for International CMS were significantly below the requirements for 2008. In the CMS computing model, an aggregate shortage of resources would also impair the capabilities of U.S. scientists to analyze data. The committee was pleased to

note that the aggregate resources now pledged for 2008 meet the requirements of CMS, which is partially due to the delay in LHC startup.

U.S. CMS has made very good progress in the past year in preparing its facilities and grids capabilities, and has met or exceeded all the goals of CSA06 in achieving 25% of the requirements for 2008. The current U.S. CMS resources for Tier-1 and Tier-2 centers are on target for computing capacities, network, CPU, and disk space needed for LHC start of operations.

2.2.2 OBSERVATIONS

The activities in CSA06 demonstrated that the U.S. Tier-2 sites can be used for CMS data analysis, as planned in the CMS Computing Model. The use of Analysis Object Data (AOD) in data analysis has attained sufficiently high levels of performance to provide users with an effective analysis procedure.

The CMS Computing Model calls for Tier-1 centers to serve data with equal priority to all CMS Tier-2 centers around the world. LHCNet is provisioned to provide links from Tier-0 (CERN) to Tier-1 centers in the U.S. However, there is risk that the incumbent networks might not be able to handle full capacity of international Tier-2 centers. In addition, if any Tier-1 center goes offline for an extended period of time, there would be significant impact on overall CMS computing operations. However, U.S. CMS has demonstrated a good start on provisioning of Tier-2 links in their network topology to both U.S. and non-U.S. facilities. The U.S. distributed network has successfully handled Tier-1 to Tier-2 transfers, with a factor of two safety margin relative to requirements for 2008.

The grid is a key element of the current package of software tools and, as demonstrated during the review, has made good progress in terms of usability and readiness. The ProdAgent production system is able to fully utilize the OSG and EGEE grids. A substantial portion of the data analysis is running on the grid using CRAB. However, scalability and reliability of tracking and correction procedures for failures in grid-submitted jobs was not explicitly demonstrated during the review.

The procedures and policies for grid-level cybersecurity are in place and have been exercised for a few minor “alarms”. No intrusion was detected in any of these cases. U.S. CMS stated during the review that responsibilities for cybersecurity at its Tier-1 and Tier-2 centers and for OSG have now been established.

2.2.3 RECOMMENDATIONS

The committee makes the following recommendations:

- In order to validate its computing model, U.S. CMS should follow its stated plan to scale up the testing of analyses at Tier-2 sites to meet required targets.

- CMS should adopt checksum techniques for all data transfers to ensure integrity of data.
- CMS should develop and deploy monitoring and diagnostic tools that allow operators to manage the predicted scale of job submissions for data, and demonstrate the success of these tools in CSA07.
- U.S. CMS should develop a plan to address end-to-end data transfers between U.S. Tier-2 sites and non-U.S. Tier-1 sites, and work with the appropriate parties to ensure that the CMS computing model matches the bandwidths available on trans-Atlantic networks. This plan should be presented at the next agency review.

2.3 CORE SOFTWARE

Material on recent accomplishments and planning of software support for user analyses was presented during the plenary sessions by Lothar Bauerdick, Ian Fisk, Patty McBride, Darin Acosta and Sarah Eno. Also, a one hour demonstration of the use of CMS software in the grid environment was given by Oliver Gutsche. In addition, three members of the review committee met with the U.S. CMS representatives for a 2 hour discussion.

2.3.1 KEY FINDINGS

During 2006, the CMS software group worked mostly on CMS integration and commissioning, with the goal of developing a usable and scalable system. The most important accomplishment in that context was the complete integration of the CMS software with the new framework and data model. The U.S. CMS S&C group carries direct responsibility for the two latter components. The new suite of software is called "CMSSW". As part of this activity, tools for software release integration, configuration, and distribution had to be developed and put in place. The success of this activity was demonstrated convincingly during the data challenge CSA06, which exercised the software by processing 50 million simulated events. This corresponds to 25% of the capacity that will be needed in 2008. The exercise established a 10^{-6} failure rate from generic software errors. The major components tested in CSA06 corresponded to (i) the use of CMS software (CMSSW) for simulation and reconstruction with rates of up to 250 Hz, (ii) the use of the software calibration system, (iii) generation of both the detector raw data and reconstructed data (FEVT), (iv) Analysis Object Data (AOD) formats, (v) skims of calibration data, and (vi) splitting of a sample with a high-level trigger ("HLT-tagged sample") into multiple streams.

The CMS S&C activities have evolved successfully from a software project into an offline organization, and are moving from a development into a maintenance and operations phase. The emphasis from here on will be on those tasks that are crucial for the startup of preparations and for commissioning, including establishing software validation procedures and establishing strong connections with the detector groups.

As the transition to operations develops, the analysis support group is being strengthened, and the analysis-support model validated. This effort is tightly integrated into the global CMS support organization. In fact, the global nature of this integrated effort provides around-the-clock response to user questions. Examples of useful tools are HyperNews for responding to user questions and Savanna for tracing bugs. The plan is to extend this effort with the help of U.S. CMS physicists, who will be contributing as part of their common service-task responsibilities. Introductory software tutorials are being organized regularly. The LHC physics center at Fermilab (LPC) continues to expand, currently occupying the 10th and the 11th floors of Wilson Hall. The LPC also provides meeting rooms, video conferencing facilities, large scale computing, and space for informal discussions. Currently, there are between 40 and 60 visitors at the LPC on average.

The emphasis in 2007 is to establish a sustainable operations model. The CMS activities will be driven by three major milestones: (i) A computing challenge, "CSA07," planned for late summer 2007, with the goal of reproducing analyses based on previously published CMS physics analyses, but now using the new software package; this requires significant effort by the software team, including U.S. personnel at CERN, on software releases needed for production; (ii) another magnet test and cosmic challenge (MTCC3) will take place in October 2007, which will include the detector, the data acquisition system and both online and offline computing; and (iii) the December pilot run will provide a small amount of data, but will be very useful for testing the detector calibration system.

All the above will comprise "end-to-end" activities, involving every component in the software and computing system, starting from the High Level Trigger and down to the Tier-2 centers.

2.3.2 OBSERVATIONS

The review committee commends the U.S. CMS S&C group on the achievements during 2006 and for its careful planning for 2007. Also, the committee commends U.S. CMS for their clear and concise presentations. The group responded clearly to the questions from the review committee. In addition, the software demonstration was well done and very useful in making the committee appreciate the state of the system. It is the general conclusion of the committee that the group is well prepared to handle the upcoming challenges.

The committee observes that an effective software-support organization for users is in place, which will be extended further through leveraging the effort of other CMS physicists. This effort is well integrated into the global CMS organization, making good use of available resources and thereby strengthening the whole collaboration.

The rate for generic software failures during CSA06 was established at a level of one in a million. However, the methodology for obtaining this failure rate was unfortunately not presented at the review. In fact, judging from the demonstration, it was evident that

measuring generic software failure rates is difficult in the absence of effective quality assurance tools. It is also important to note that this was done using simulated events. With data from detectors, the failure rate is likely to grow. Nevertheless, the committee observes that the established rate of failures is reasonable at this stage of the experiment.

U.S. CMS showed that data in the present AOD format can be analyzed rapidly at a rate exceeding 1 kHz using simple analysis algorithms. A significant overhead due to the dCache protocol was evident during the demonstration; however, the use of dCache was not optimized for the exercise. The committee believes that the analysis rate shown for very simple analysis tasks is appropriate, but that the dCache overhead is currently too large.

During CSA06, the data management procedures were exercised successfully between the Tier-0 and the Tier-1 sites. The upcoming studies will extend this to the high level trigger (HLT) and Tier-0 interfaces. Initially, the exercises will rely on manual operation by an expert team. Tools are also being developed to monitor the data-access patterns more effectively. The committee notes that this is a sound strategy at the initial stage of the experiment.

The committee observes that, while the production aspects of the computing system have been tested and validated, the validation of the analysis aspects of the computing model, and in particular the performance of the software under significant load from user analysis, is not yet tested and therefore still presents certain risks.

2.3.3 RECOMMENDATIONS

The committee makes the following recommendations:

- The performance of dCache should to be optimized. We recommend this be done for CSA07.
- The appropriate quality assurance and monitoring tools must be delivered in order to enable optimization of software performance.
- CMS should try to decrease the rate of software failures
- Analysis rates should be optimized, in particular using more realistic algorithms.
- CSA07 should be used to gain as much experience as possible on Tier-2 access patterns, and as a means for validating the functionality of monitoring tools

3 U.S. ATLAS

3.1 MANAGEMENT

The U.S. ATLAS management team for Software and Computing of J. Shank et al. presented the status of their effort during the plenary and parallel sessions.

3.1.1 KEY FINDINGS

The U.S. ATLAS S&C project is making a transition from development to operations in preparation for data collection. This has led to the creation of a manager for operations and a new board to guide the allocation of resources. U.S. ATLAS is making good use of its change-control board, and has addressed five important issues in the past year. The U.S. collaboration continues to grow slowly, and still plans to add a few new collaborators. However, the incremental costs for such growth are not well understood. U.S. ATLAS management takes responsibility for ensuring that new collaborators can contribute effectively to ATLAS, and encourages smaller groups to join in affiliation with larger established U.S. ATLAS institutions. U.S. collaborators are well integrated within the global collaboration and are well represented in leadership positions throughout the organization.

The deployment of the Tier-1 and Tier-2 sites is progressing well, and is contributing 28% of the computing to the global ATLAS Computing System Commissioning (CSC) test. The multi-site Tier-2 model continues to deliver computing resources to the overall system, with the Southwest Tier-2 contributing 30% of the total U.S. ATLAS computing resources (compared to 40% at the Tier-1).

PanDA, developed by U.S. ATLAS to provide a distributed analysis and production system, has proven to be highly efficient, and is gaining international acceptance. Performance and scaling issues continue to be addressed. Full acceptance of PanDA by the collaboration would enable a sharing of the operations support within the collaboration, freeing up personnel to further develop PanDA. This will enable the U.S. to retain a leadership role in ATLAS computing and the ability to adapt to changing grid technologies.

Issues with Distributed Data Management (DDM) remain a concern and a serious risk for the experiment. The DDM was the subject of a November 2006 ATLAS review, and a forthcoming report is expected to contain extensive recommendations.

The U.S. ATLAS user analysis support model relies on help from Analysis Support Groups (teams of on-call experts), the Analysis Support Centers (at ANL, BNL and LBNL) and running of a variety of Analysis Forums. In order to train the user community, a series of “Jamborees” have been held at the three Analysis Support Centers—targeting new users of ATLAS software. Approximately 100 physicists have participated. The jamborees have been successful with very positive feedback received

from collaborators. The feedback is used to improve the jamborees and to develop a more focused curriculum. Readiness for analysis has been reviewed internally by U.S. ATLAS, and provided several surprises, among which was that the on-call experts in the Analysis Support Groups were underutilized. To facilitate U.S. analysis, the plan is to locate a complete copy of all Event Summary Data (ESD) from ATLAS at the BNL U.S. Tier-1 center.

The ATLAS Computing Model has recently been refined to provide a better accounting of resource requirements. The assumptions for the model have budgetary impact in the out-years, and remain under discussion and evaluation. The charge for the review requested an examination of the implication of a 10% funding shortfall. The committee observed that this size of shortfall relative to the agency-provided guidelines would require staff reductions at the Tier-1, reduced staff for software and operations tasks and a reduction of hardware at the Tier-2 centers, leading to highly detrimental consequences on U.S. analysis at a critical time of need.

3.1.2 OBSERVATIONS

The committee commends the U.S. ATLAS team on significant progress in the past year, including the success of PanDA, participation in global commissioning tests and in improvements in analysis support. The committee appreciates the responsiveness and candor of U.S. ATLAS management, and notes that, generally, the recommendations from last year's review have been addressed in commendable ways

The management structure of U.S. ATLAS is effective and appropriate, with problems being identified and addressed in a timely fashion. Sensible processes for change control and for resource allocation are in place and working well. U.S. ATLAS has established a plan for transition from development to operations, but it is difficult at this point to assess the planning and its flexibility in addressing unanticipated problems. We observed that many aspects of ATLAS and U.S. ATLAS were under internal review in the past year. In general, this is a commendable approach as such committees provide an excellent way to achieve consensus within an international collaboration. However, this can be a very slow and labor intensive process, which does not always converge in a timely manner.

We note the lack of reports on the recent end-to-end cosmic ray test. Such tests, relying on current software and computing are crucial for exposing problems in overall system integration, and we therefore encourage full participation in testing with cosmic rays and in the proposed full dress rehearsal (FDR) of the S&C in mid-2007, even if it distracts somewhat from testing of separate U.S. deliverables.

Locating the complete copy of the ESD at the U.S. Tier-1 center is likely to make analysis easier, but at the risk of an increased single-point of failure. We note that ATLAS has considered mitigations of this risk, for example, by partnering with other Tier-1 sites and by using more flexible software. The metrics presented for the

effectiveness of support of software and analysis exhibited a significant divergence from the computing model in terms of the initial expectation for the size of the data set, as well as a changing data model. This could lead to an increase in storage requirements at a time of convergence and consolidation of the model within an already tight budget.

U.S. ATLAS is taking reasonable steps to educate its user community and to build a consensus that will ensure that U.S. collaborators are ready and able to do their physics analyses in an effective manner.

We note that any funding shortfall would limit flexibility, leading to extreme risk for meeting essential functionality. Even with full funding (including the proposed call on management reserve), US ATLAS S&C has limited flexibility to reassign resources for addressing pressing issues that will inevitably arise in the early periods of data collection.

3.1.3 RECOMMENDATIONS

The committee makes the following recommendations:

- U.S. ATLAS should work with ATLAS to get timely delivery of reports such that risks can be identified and addressed expeditiously.
- U.S. ATLAS should strive to maintain sufficient flexibility to be able to address urgent problems that will inevitably arise during the year before data-taking.
- U.S. ATLAS, working with ATLAS, is urged to focus on a full system test with well defined milestones. There should be public and published milestones that reflect the full system-functionality required to support the experiment during data collection, and include all activities such as the full dress rehearsal, cosmic-ray tests and commissioning of the entire computing system.
- We urge U.S. ATLAS to work closely with the ATLAS collaboration on clarifying and converging on its data model and projected event size (ESD, EOD and Derived Physics Data format or DPD), and to evaluate cost and impact of a potential growth in the size of events.
- We recommend that U.S. ATLAS work urgently with ATLAS to execute the Distributed Data Management plan. This requires a plan to be put in place to solve both short-term needs and long-term issues in DDM, and to make sure that sufficient resources are dedicated to guarantee its success. This plan should include concrete milestones that will lead to a working DDM as soon as feasible, but certainly functional within 2007.
- We recommend maintaining the development effort of PanDA in the U.S. and pursuing plans to increase international support for operations.
- We recommend continued use of the change-control procedures and for focusing on existing U.S. ATLAS responsibilities; avoiding an expansion of scope should help optimize the use of available resources

3.2 FACILITIES AND GRIDS

The material relevant to the charge for U.S. ATLAS facilities and grids was available in the presentations of Jim Shank, Bruce Gibbard, Torre Wenaus, Kaushik De and Stephane Willocq, as well as in additional responses prepared before and during the review.

3.2.1 KEY FINDINGS

U.S. ATLAS facilities continued to expand in 2006 to match the projected performance and capacity requirements of the U.S. ATLAS computing model and the additional needs to support its U.S. community. The new standalone mass storage system was deployed at the BNL Tier-1 center with the expected capacity (600 TB) and data transfer rates (400 MB/s). The WAN links were upgraded from OC48 (2.5 Gb/sec) to 2 lambda (20 Gb/s) circuits, one dedicated to general ESnet traffic and one to the LHC Optical Private Network. As foreseen, the U.S. ATLAS staff grew to 15 FTEs.

The selection of U.S. ATLAS Tier-2 sites was completed with the addition of the SLAC Tier-2 and the "Great Lakes" (University of Michigan and MSU) Tier-2. The Tier-1 and Tier-2 centers participated in the Computing System Commissioning effort in 2006. The fraction of total jobs run on US facilities (28%) exceeded their “fair share” value based on PhD author-count (20%). The U.S. ATLAS OSG resources were used efficiently and extensively in production runs under the PanDA system, which will be extended for opportunistically exploiting more resources through the OSG and for supporting LCG for ATLAS-wide analysis and production.

Functionality tests of the full Tier-0/Tier-1/Tier-2 chain were carried out successfully. Particularly noticeable has been the effort to operate and optimize the SRM/dCache storage system at the Tier-1 center.

Recently, the capacity and funding profiles of Tier-1 and Tier-2 facilities were revised to take into account the change in LHC startup schedule, the inclusion of efficiency factors for utilization of resources, the fact that CERN does not provide any computing resources aside from the Tier-0 and the CAF, and the separation of contributions related to running of heavy ions from the HEP requirements of U.S. ATLAS. The net result is that the currently estimated CPU needs at the Tier-1 were significantly reduced beyond 2008, while the opposite is true for disk space. The yearly projected total costs show a maximum increase of 18% in 2010. The main impact on the Tier-2 is a large imbalance between target capacities and currently planned resources that peaks at -36% in 2010 and 2011.

3.2.2 OBSERVATIONS

Progress has been demonstrated on all recommendations made at the '06 S&C review. There remains a concern with dCache performance for chaotic data analysis, which requires further investigation.

Good progress has been demonstrated in utilizing the Tier-1 as well as the Tier-2 facilities, where, in the latter case, there has been considerable success in starting event simulation in production mode.

Although most of the analysis activities continue to be carried out at the BNL Tier-1, the U.S. collaboration has a plan to test data analyses at the Tier-2 sites in advance of data taking. We deem important that such a plan is carried out in a timely manner in order to guarantee that the Tier-2 centers are ready to accept the analysis load, once the bulk of resources available at the Tier-1 are taken over for intended production tasks.

The full copy of ESD at Tier-1 and a full copy of AOD in each of the U.S. Tier-2 sites will surely benefit the analysis capabilities of U.S. groups. However, if ESD and AOD sizes remain larger than originally foreseen in the Computing TDR, this approach may not be possible.

Since the BNL Tier-1 is the only Tier-1 serving the full ESD data set, it may become an attractive choice of access for international ATLAS, potentially impacting U.S. analyses. Tier-0 to Tier-1 integration has been demonstrated, but because of problems encountered in using DDM at Tier-2 sites, Tier-1 to Tier-2 connections have not been tested sufficiently. The plan to integrate Tier-3 centers via the OSG is reasonable, but defining desktops or PCs as Tier-3 sites is probably not appropriate.

The Service Challenges and CSC-related activities have been very important to demonstrate that the infrastructure at facilities is capable of supporting major production tasks. The autonomy of the Tier-2 centers was demonstrated when they ran successfully during an 8 hour stand-down of BNL, which took their Tier-1 facility offline.

There has been good progress by U.S. physicists in the wide use of the grid via PanDA. U.S. ATLAS has demonstrated that the appropriate management links exist between OSG and WLCG. U.S. ATLAS has also shown that PanDA can be used to submit jobs to both EGEE and OSG.

U.S. ATLAS now has a cybersecurity officer in place. Both U.S. ATLAS and BNL have an understanding of the cybersecurity responsibilities at the Tier-1 facility. The U.S. ATLAS group is also leveraging OSG expertise to address cybersecurity issues for the grid.

The currently budgeted plan to add a fully redundant path on top of the existing 20 Gb/s links from BNL to ESNet seem adequate for the initial running of the LHC. The U.S. ATLAS model, which foresees storage of the full ESD sample at BNL, requires good connectivity between each Tier-2 center and the Tier-1 center at BNL. The connectivity between the Tier-0 and BNL Tier-1 center, which consists of a set of existing or planned 10 Gb/s links to be completed by 2008, also looks to be on track to fulfill the bandwidth requirements..

3.2.3 RECOMMENDATIONS

The committee makes the following recommendations:

- To minimize the impact of large event sizes on the U.S. ATLAS storage model, U.S. ATLAS should ask the ATLAS task force on ESD and AOD event formats to report its conclusions by early summer 2007 to U.S. ATLAS management. In parallel, U.S. ATLAS should develop a contingency plan to address any detrimental impact from a significant excess in event size.
- ATLAS should adopt checksum techniques for all data transfers so as to ensure the integrity of data.
- By the time of the next agency review, U.S. ATLAS should develop a plan for addressing the Tier-2 shortfall expected in 2010 and 2011 .
- U.S. ATLAS should work with the BNL to put in place a plan aimed at mitigating the effects of cybersecurity incidents that might affect the availability of the BNL Tier-1 facility.

3.3 CORE SOFTWARE

Material on the accomplishments and on the planning within the software and support efforts was presented during the plenary session in talks given by Jim Shank, Bruce Gibbard and Srini Rajagopalan. A one hour demonstration of how to use the ATLAS software in the grid environment was given by Kaushik De and Torre Wenaus, and a software usability demonstration (pAthena) was provided by Stephane Willocq. In addition, the review committee met with the U.S. ATLAS representatives for a 2-hour discussion.

3.3.1 KEY FINDINGS

The software demonstration was very satisfactory, showing that the software is currently functional. The committee considered this an important milestone. The committee also found it commendable that PanDA worked very well in last year's data production. It was shown that PanDA produced more events than any other grid submission tool in ATLAS. Some scaling issues were identified and these are being addressed. The committee also saw it as a positive step that Stephen Gowdy will be taking over the integration of Monte Carlo production into PanDA.

The committee notes that 2007 is a year of intense activity for ATLAS S&C. The important tests and milestones in 2007 include the Computing System Commissioning (CSC) activities, the Calibration Data Challenge (CDC), the integrated cosmic-ray run, the TDAQ large-scale test and the 900 GeV engineering run. It was stated that work is ongoing towards all of these milestones, and that the plan calls for their completion by summer 2007. The plan also calls for a "full dress rehearsal" of S&C to take place this year. The committee feels that the dress rehearsal is an important milestone for demonstrating the end-to-end functionality of the software and computing chain. It was stated that the strategy for the full dress rehearsal was about to be completed.

The current ATLAS-wide event production rate is about 2 million events per week. At this rate, about 10 million events needed for the full dress rehearsal can be generated in 5-6 weeks. An increase by another factor of two in generation rate was anticipated by summer 2007, which would reduce the total time to 2-3 weeks.

U.S. ATLAS has invested a significant amount of effort in Analysis Support Centers and Analysis Support activities for US ATLAS users, including tutorials and jamborees.

3.3.2 OBSERVATIONS

The review committee commends the U.S. ATLAS S&C SW group on the achievements during 2006 and for the careful planning for 2007. The committee notes that the software demonstrations were well done and useful for developing an understanding of the state of the system. U.S. ATLAS responded clearly to the questions from the review committee, and provided clear and concise presentations. It is the general conclusion of the committee that the group is well prepared to handle the upcoming challenges.

The committee feels that the computing resources are adequate for completing the full dress rehearsal in 2007.

The committee commends U.S. ATLAS for investing considerable resources in the Analysis Support Centers and Analysis Support activities of its user community, including organizing tutorials and jamborees. Calendar year 2007 will produce a large burden on U.S. ATLAS core software, and there is a chance that U.S. ATLAS support may suffer while the needs of global ATLAS are being addressed. Potentially, U.S. support groups are exposed to high levels of requests from across the collaboration. While this is a sign of a job well done in terms of being recognized experts and as leaders in the development of core software (framework and event store) and physics analysis tools, U.S. ATLAS developers will likely be asked to field many of the support requests. We commend the strategy of developing documentation for users, which will mitigate the anticipated load on the U.S. experts. We note, however, that there was no response this year to last year's recommendation for regular assessments of the level of load needed for satisfying support services.

ATLAS recognized the need for a user-defined ntuple (DPD) for analysis. It is commendable that this need was recognized early on and made part of the plan for developing software and analysis codes. U.S. ATLAS is developing a common framework to generate and write analysis ntuples, allowing physics groups and individuals to create these directly from ATHENA. This is to be commended.

At the request of the ATLAS technical coordinator, U.S. ATLAS plans to devote one new FTE to develop an online event display tool called vATLAS,. This is seen as an opportunity to leverage graphics expertise in the U.S. and provide a collaboration-wide

tool. While we feel that this is a sensible strategy, we believe that more effort than 1 FTE will be needed for the development of vATLAS.

We note that validation of releases appears to be quite inadequate. Procedures and implementation of system-wide testing were not discussed. We observe a lack of Quality Assurance tools, and believe this deficiency should be corrected so that it does not impact the success of the project.

It was shown graphically that the software failure rate spikes with each new release. We find that the failure rate is unacceptably high. In addition, there was no measure offered for the quaity of the algorithms in use.

In summary, we commend U.S. ATLAS for their assumed level of responsibilities in international ATLAS, and for the key projects they develop and support. While noting that significant progress had been made in the area of core software, the committee would like to state that a stable software release with a low failure rate and acceptable algorithm performance is a crucial milestone for the full dress rehearsal in 2007.

3.3.3 RECOMMENDATIONS

The committee makes the following recommendations to U.S. ATLAS:

- U.S. ATLAS should press ATLAS for an integration schedule with milestones to define the year's activities and priorities.
- U.S. ATLAS should demonstrate during the next agency review the outcome on the DPD effort, including to what extent it was adopted by the collaboration
- The additional resources invested by U.S. ATLAS in developing vATLAS should be matched with comparable support within international ATLAS. As in the case of PanDA, a strategy should be developed for ATLAS-wide adoption of vATLAS.
- U.S. ATLAS should assess the potential workload generated by software support and develop a mitigation strategy to limit exposure of U.S. ATLAS developers. For example, the U.S. should collaborate with other ATLAS partners who are already developing support solutions.
- U.S. ATLAS should press ATLAS to develop a culture of code QA and testing, and to develop system tests this year. This should also lead to a reduction in software failures.
- U.S. ATLAS should report job failure rates, including a clear breakdown by point of failure and per event probabilities.

4 APPENDIX A

U.S. Department of Energy
and the
National Science Foundation

U.S. LHC Joint Oversight Group

To: J. K. Blackburn (Caltech), A. Boehnlein (FNAL), R. Dubois (SLAC), T. Haas (DESY), A. Kotwal (Duke Univ.), J. Lauret (BNL), M. Morandin (INFN, Padova), D. Olson (LBNL), T. Schalk (UC, Santa Cruz)

Subject: Charge for the January 2007 review of the software and computing (S&C) efforts of the U.S. LHC Research Program

The Joint Oversight Group (JOG) of the Department of Energy and the National Science Foundation (DOE/NSF) greatly appreciates your willingness to participate in the review of the software and computing (S&C) programs of U.S. ATLAS and U.S. CMS. This review will take place at the University of Texas at Arlington on January 17-19, 2007.

The goal is to evaluate recent progress and the effectiveness of current and planned U.S. S&C activities in enabling the research of U.S. physicists at the LHC. You are asked to examine the scope, cost and schedule of the S&C plans for the period of FY2007 through FY2011, with special emphasis on readiness for initial LHC running. You are asked to judge whether the U.S. S&C portfolio is balanced so as to give U.S. researchers a sound basis for effective participation in the analysis of LHC data, and to examine the realism of the plans in the context of the funding guidelines provided by the agencies.

In addition, we ask you to evaluate the progress made by each collaboration in implementing the recommendations of previous reviews (February 2006 and August 2006). As a guide, we point to the following issues:

1) Management

- Are the current management structures and techniques well-matched to the needs of the U.S. collaboration?
- Are their internal contingency and risk-management mechanisms appropriate?
- Are there adequate plans for transitioning from a development phase to a deployment and operations phase? Are the assumptions for resource requirements well justified?
- Are the priorities of the S&C program conducive to effective participation in data analysis by U.S. physicists?
- Does management have adequate S&C plans to accommodate new collaborators? Have they developed a reasonable model for the corresponding incremental costs?

- What would be the impact of a 10% S&C funding shortfall on current U.S. deliverables and on productivity in physics analysis?
- Do the U.S. projects interact sufficiently with the international S&C efforts?
- Does the U.S. play a role in the international S&C leadership that is commensurate with its overall participation in the experiment?

2) Facilities, Grids, Networking, and Infrastructure

- Are the current computing models of the experiments appropriate for U.S. needs?
- Have infrastructure and operating costs of the Tier-1 and Tier-2 facilities been fully considered in their deployment? Are there any high-risk assumptions? Are the estimated personnel requirements, equipment and infrastructure costs valid and well-justified?
- Has there been adequate progress made in deploying the U.S. Tier-1 and Tier-2 centers and in their integration with the CERN Tier-0 center?
- From a user's perspective, is the usability and readiness of grid-based production software in good shape? (The collaborations should provide sufficient information to help the committee evaluate typical user experience with grid-based tools.)
- Are cybersecurity issues given adequate priority by management? On matters of security, are the lines of authority clearly spelled out? Has the collaboration assessed the impact of a cybersecurity incident on user access to data and to computing cycles? Is there a mitigation plan in place?
- Have network bandwidth and connectivity requirements been appropriately identified by the U.S. collaborations? Are these requirements consistent with their latest computing models? Is there a roadmap to achieve the required T0-T1-T2 connectivity?
- Do the U.S. S&C programs have adequate links to the Worldwide LHC Computing Grid (WLCG) and the Open Science Grid (OSG)?

3) Core Software and Analysis Support

- Are the current models for the support of data analysis well thought out, and is the support structure responsive to the needs of the U.S. community, both in the U.S. and at CERN? Are there adequate metrics to monitor progress in this area? Have all the required resources been identified by the collaborations? Will there be adequate support during all phases of the experiment?
- From a user's perspective, comment on the usability and readiness of the analysis software. (The collaborations should provide sufficient information to help the committee evaluate typical user experience with analysis tools.)
- Are the personnel requirements for the maintenance and operation phase of production software well understood, well justified and available? On what basis are commitments made to the international collaboration? Are these commitments realistic and consistent with U.S. interests?
- Is the role of Tier-3 centers well-defined and integrated into the S&C facilities plan? Are the plans for Tier-3 facilities sufficiently developed to guarantee capability for data analysis at interested U.S. institutions by November 2007?
- Has progress in core software relative to the milestones presented at the February 2006 comprehensive DOE/NSF review of the U.S. program been adequate? Are the forthcoming U.S. milestones on track and realistic? Is there any critical dependence on international milestones that could put U.S. deliverables at risk?
- Is the U.S. core software portfolio sufficiently balanced to offer U.S. researchers a good chance to participate effectively in the initial science of the LHC?

The review will be chaired by the U.S. LHC Associate Program Manager for Computing, Saul Gonzalez, with other program staff members from the DOE and the NSF in attendance. The proponents will post their presentations on the web a week prior to the meeting, and you will have access to additional supporting documentation at least two weeks prior to the start of the review. We would appreciate close-out statements following the reviews of both U.S. ATLAS and U.S. CMS, and more formal written reports within two weeks of the completion of your evaluation. This will provide valuable and timely input to the agencies and to the experiments. Your reports will also be made available to other DOE and NSF committees that review U.S. ATLAS and U.S. CMS projects.

Again, we wish to express our deepest appreciation for your willingness to participate in this important activity.

Sincerely,

John R. O'Fallon	John Lightbody, Jr.
Co-Chair	Co-Chair
U.S. LHC Joint Oversight Group	U.S. LHC Joint Oversight Group
Department of Energy	National Science Foundation

cc: Tom Ferbel, SC-25
Aesook Byon-Wagner, SC-25
Moishe Pripstein, NSF/MPS
Jim Whitmore, NSF/MPS
Miriam Heller, NSF/OCI
Glen Crawford, SC-25
Saul Gonzalez, SC-25
Craig Tull, SC-25
Dan Green, Fermilab
Joel Butler, Fermilab
Bob Cousins, UCLA
Jim Shank, Boston University
Lothar Bauerdick, Fermilab
Michael Tuts, Columbia University
Howard Gordon, BNL
Hugh Montgomery, Fermilab
Peter Bond, BNL

5 APPENDIX B

Review Panel Membership:

J.K. Blackburn (Caltech)	T. Haas (DESY)	M. Morandin (INFN-Padova)
A. Boehnlein (FNAL)	A. Kotwal (Duke U.)	D. Olson (LBL)
R. Dubois (SLAC)	L. Lauret (BNL)	T. Schalk (UC, Santa Cruz)

Agency Participation:

Glen Crawford (DOE/HEP)
Tom Ferbel (DOE/HEP)
Saul Gonzalez (DOE/HEP, Chair)
Kevin Thompson (NSF/OCI)
Moishe Pripstein (NSF/EPP)
Craig Tull (DOE/HEP)

4 APPENDIX C

U.S. LHC Software and Computing Review
January 17-19, 2007
University of Texas, Arlington

AGENDA

Room Assignments: Plenary and Executive Sessions: Room 303
Management Session: Room 303
Core Software and Analysis: TBA
Facilities & Grids: TBA

January 17 (Wednesday)

8:30 Executive Session (30')

ATLAS

9:00 Plenary Session (1:40)
US ATLAS S&C Project – J. Shank (40'+10')
Facilities – B. Gibbard (40'+10')

10:40 COFFEE (0:15)

10:55 Plenary Session (0:50)
Software and Analysis Support – S. Rajagopalan (40'+10')

11:45 Executive Session (0:30)

12:15 LUNCH (1:00)

13:15 Parallel Discussion Session (2:00)

Management	Facilities and Grids	Software / Analysis
Analysis Support Model – M. Tuts (45')	(common session on usability) Usability of PanDA and production running – De/Wenaus (30') Usability of pathena and user analysis – Willocq (30')	
Discussion	Discussion	Discussion

15:15 COFFEE (0:15)

15:30 Executive Session (2:30)

18:00 Adjourn

January 18 (Thursday)

ATLAS

8:00 Executive session (2:30)
(Further discussions with ATLAS, preparations for closeout)

10:30 COFFEE (0:15)

CMS

10:45 Plenary Session (1:30)

Introduction – LAT Bauerdick (20'+5')
CMS Offline and Computing – P. McBride (15'+5')
US CMS S&C Project – I. Fisk (35'+10')

12:15 LUNCH (1:00)

13:15 Plenary Session (1:00)
CSA06 – D. Acosta (15'+5')
Readiness and Operation of Grid Infrastructure for US CMS – R. Pordes (15'+5')
Physics Readiness – S. Eno (15'+5')

14:15 Executive Session (0:30)

14:45 COFFEE (0:15)

15:00 Parallel Discussion Session (2:00)

Management	Facilities and Grids	Software / Analysis
Bauerdick, Butler et al.	(common session on usability) Usability of production tools (30') Usability of software and analysis tools (30')	
Discussion	T2s, DISUN, WAN – J. Bakken et al. Discussion	L. Sexton-Kennedy et al. Discussion

17:00 Executive Session (1:30)

18:30 Adjourn

January 19 (Friday)

CMS

8:00 Executive Session (1:00)
(Including further discussions with CMS)

9:00 Executive Session (1:30)

10:30 COFFEE (0:15)

10:45 Executive Session (1:00)
(Preparations for closeout)

11:45 LUNCH (1:00)

12:45 ATLAS Closeout (0:30)

13:15 CMS Closeout (0:30)

14:00 Adjourn

5 APPENDIX D

Material in *italics* to be provided prior to S&C review in support of review charge:

MANAGEMENT

1. **Are the current management structures and techniques well-matched to the needs of the U.S. collaboration?**
   - ***Please provide an organization chart of US ATLAS / US CMS***

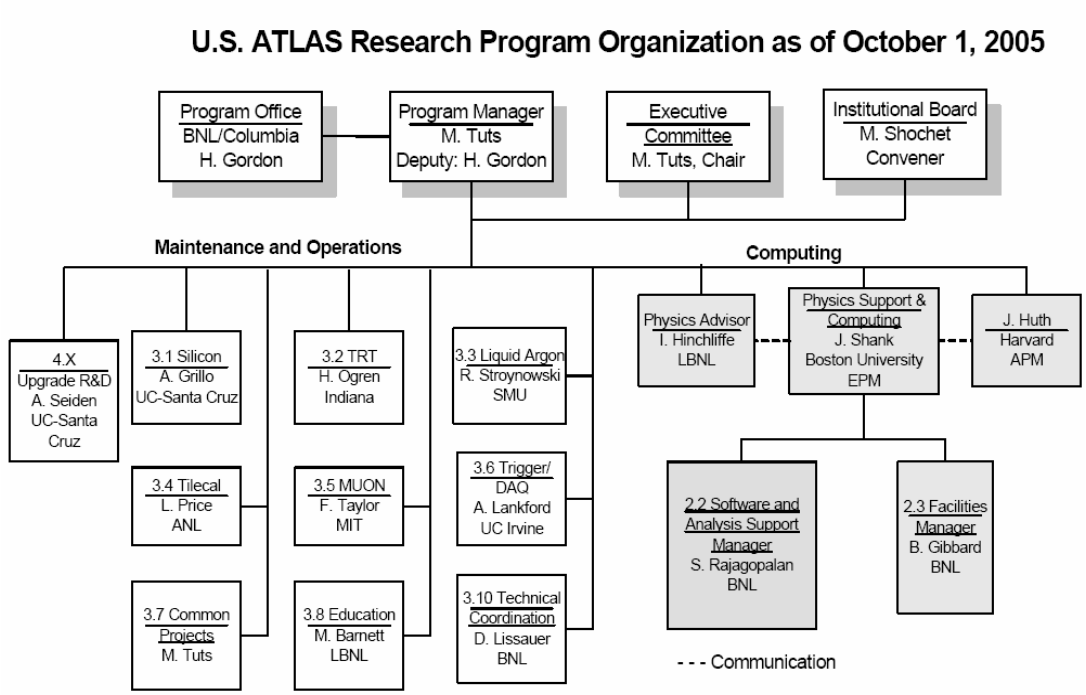

2. **Are their internal contingency and risk-management mechanisms appropriate?**
   - ***Please provide a table with S&C funding history FY2004-FY2006. If applicable, provide initial calls on M.R. and granted M.R per year.***

S&C Funding History: Fiscal Years 2004 - 2006

	FY04	FY05	FY06
Initial Approved Amount	4,633	7,795	9,816
MR Allocation	10	617	900
Final FY Budget	4,643	8,412	10,716

3. **Are there adequate plans for transitioning from a development phase to a deployment and operations phase? Are the assumptions for resource requirements well justified?**
    - ***Provide a table with personnel requirements for the period FY06-FY11; estimate distribution of FTEs: Universities, Labs, US-based, CERN-based.***

RP Funded FTE levels

		FY06	FY07	FY08	FY09	FY10	FY11
University		**15**	**20.5**	**21.5**	**20.5**	**20.5**	**20.5**
	CERN	4.5	7.5	7.5	6.5	6.5	6.5
	Home inst.	10.5	13	14	14	14	14
Lab		**34.75**	**40.5**	**41.5**	**41.5**	**41.5**	**41.5**
	CERN	5.5	5.5	6	6	6	6
	Home inst.	29.25	35	35.5	35.5	35.5	35.5
Total		**49.75**	**61**	**63**	**62**	**62**	**62**

4. **Are the priorities of the S&C program conducive to effective participation in data analysis by U.S. physicists?**
    - ***Please list your priorities and metrics in this area.***

In the software effort, from the onset, the research program funds have been targeted at projects that form a concrete foundation that is conducive for U.S. physicists to effectively participate in physics analysis. The core of our involvement lies with the reconstruction and analysis framework, the Event Store and related data management effort, and the ability to process and analyze data in a distributed manner. The three key components are essential components in providing a framework for physics analysis for U.S. physicists. During the years, we have striven to establish the required personnel in these three primary areas to ensure success to our commitments and to establish the local expertise that will help U.S. physicists. During the past years, we have fine tuned our program further to assist U.S. physicists directly. We have done this by using our research program funds to establish expertise in sub-system (tracking, calorimeter, muons) reconstruction and infrastructure support. Understanding the performance of the detector sub-systems is critical for a successful physics analysis. We are now establishing support for the Analysis Support Centers and areas such as Graphics and ROOT - again these are projects that directly serve the needs of a U.S. physicist.

In the facilities area, we are ramping up our Tier 1 and 2 sites in accordance with the ATLAS Computing model, but we also have always planned on having extra capacity at the T1 to facilitate U.S. physics analysis.

We have an established Analysis Support infrastructure that was designed with broad input from the U.S. collaboration and this has been recently reviewed, again by a panel broadly representing the U.S. ATLAS collaboration. Our support model evolves based on input from reviews such as this.

5. **Does management have adequate S&C plans to accommodate new collaborators? Have they developed a reasonable model for the corresponding incremental costs?**
    - ***Summarize your model for incremental cost of new collaborators***

Currently the U.S. has 286 PhD authors on ATLAS out of a total of 1486, so 19.25% of total. Currently the cost for Common Fund per additional collaborator is about $8k
The additional support load for new user at T1/T2 is difficult to assess since there is an economy of scale, but we estimate this to be $10k/yr/author.

U.S. ATLAS management, together with CERN ATLAS management do screen new collaborators to make sure they can really make a significant contribution to ATLAS. Small groups are encouraged to join as part of larger, established ATLAS inst. such as has recently happened with: Iowa, UT Dallas, Louisiana Tech. U, South Carolina U.

6. **What would be the impact of a 10% S&C funding shortfall on current U.S. deliverables and on productivity in physics analysis?**
    - ***Provide specific impacts assuming "reasonable calls on MR/contingency" are granted.***

A 10% reduction in FY08, would require us to cut $1.6M

- Facilities T1/T2 must be ramped up.
    1. We will have to look closely at cost and ramp-up schedule:
        - FTE's at T1➔ delay the ramp to 20 FTE
        - Hardware: ramp up late
        - Reduce hardware at T2.
            - Already short of predicted need for FY10
- Production and Distributed Software must be maintained
    1. Physicists will do analysis on OSG sites ➔ only in US
- Software must take a hit
    1. No new planned hires in software. We would have to fire existing FTE's.
        - ATLAS would suffer. Our strategy would be to redirect current effort to other ATLAS countries
            - Not a very effective way to go: we are using our expertise in areas that are crucial to ATLAS
            - No one outside of the US is going to support our facilities nor our grid.
    2. At least 3 (one at each Lab.)
- This would be very painful for the whole collaboration

1. User support would suffer and contention for CPU time would rise
    - Hard to quantify, but a typical physicists analysis may take 30% longer.
- Will seriously hurt the U.S. ability to participate in physics
    1. 30% longer analysis ➔ others get there first.

7. **Do the U.S. projects interact sufficiently with the international S&C efforts?**
8. **Does the U.S. play a role in the international S&C leadership that is commensurate with its overall participation in the experiment?**
    - ***Provide International ATLAS/CMS organizational chart.***

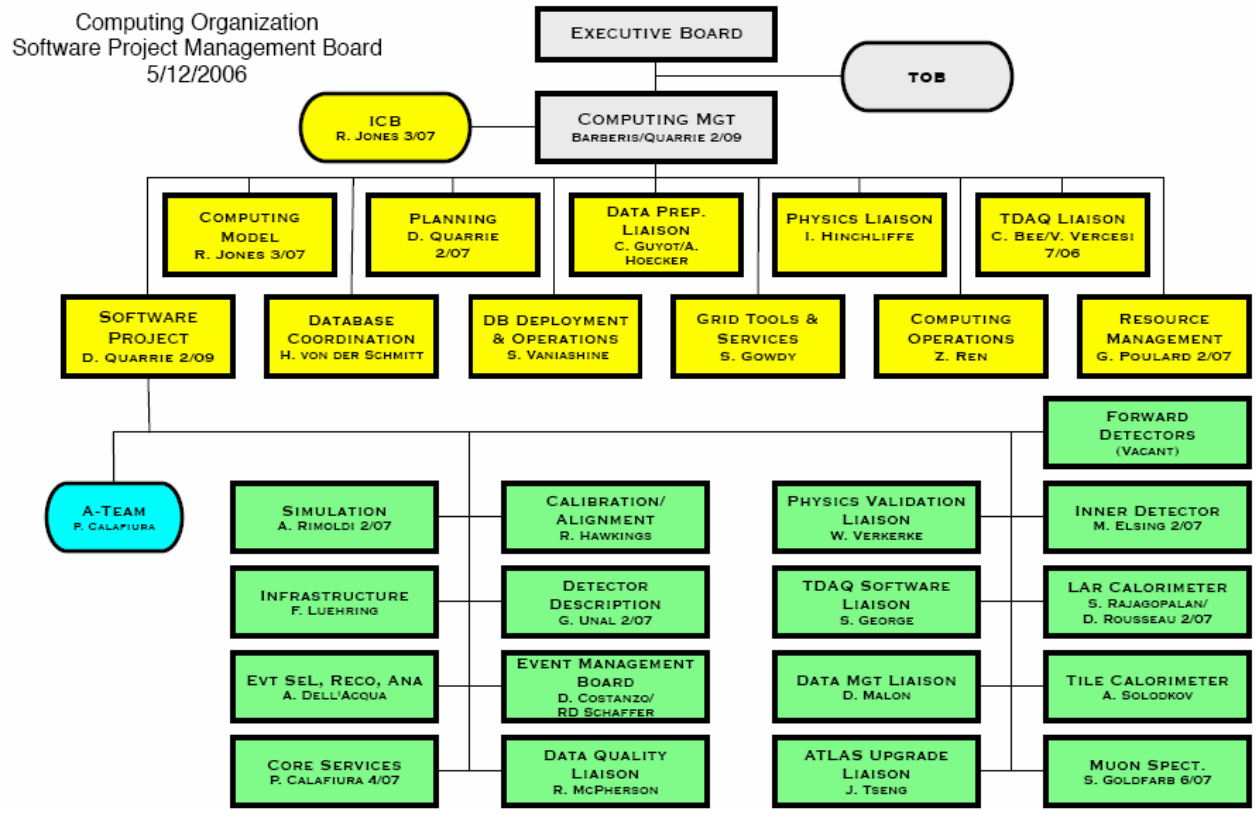

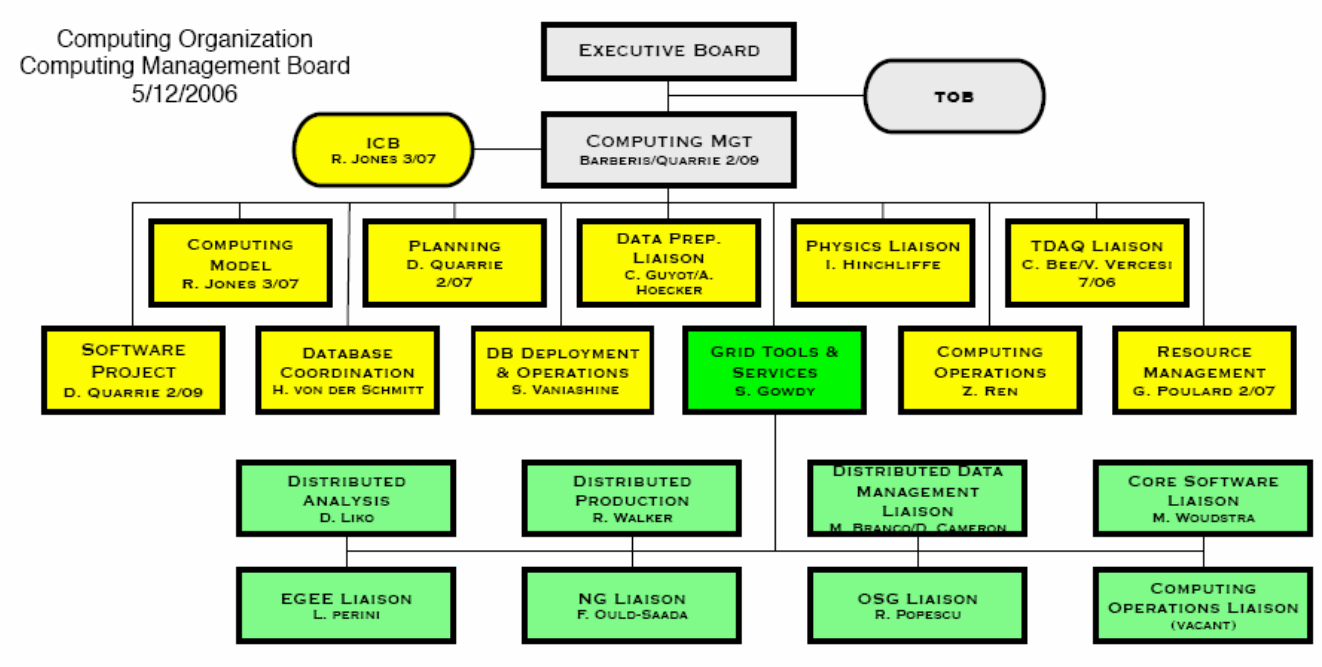

FACILITIES, GRIDS, NETWORKING, AND INFRASTRUCTURE

9. **Are the current computing models of the experiments appropriate for U.S. needs?**
   - ***Provide links or supporting material describing the computing model.***

The ATLAS computing model is formally described as part of the ATLAS Computing Technical Design Report at http://atlas-proj-computing-tdr.web.cern.ch/atlas-proj-computing-tdr/Html/Computing-TDR-4.htm. While many of the numbers in this TDR have been revised over the past two years the basic model itself as described in that document remains intact. The U.S. ATLAS computing model places some additional requirements on top of the ATLAS ones, for example, we have a copy of the full Event Summary Data (ESD) at the T1. In addition, we anticipate having analysis capability at the T1.

10. **Have infrastructure and operating costs of the Tier-1 and Tier-2 facilities been fully considered in their deployment? Are there any high-risk assumptions? Are the estimated personnel requirements, equipment and infrastructure costs valid and well-justified?**
    - ***Provide a table outlining all facilities costs (T1, T2), namely, personnel requirements and infrastructure costs for the period FY06-FY11.***

All costs; equipment, personnel, and infrastructure have been considered. Those costs which are on project are indicated in the table below. Significant infrastructure costs are born by the various institutions and significant personnel costs are also covered by some of the Tier 2's.

Projected Total Facilities Cost

($ Items below include overheads)	2007	2008	2009	2010	2011
On Program Staff Level (FTE's beyond '07)	*20*	*20*	*20*	*20*	*20*
Labor (Fully loaded salaries)	*2,892*	*3,855*	*4,048*	*4,250*	*4,463*
MST (travel, maint, licen, etc)	*927*	*1,220*	*1,394*	*1,833*	*1,552*
Facility Space & Power	*248*	*356*	*469*	*562*	*598*
Capital Equipment	*2,228*	*3,762*	*3,902*	*4,989*	*3,124*
Tier 1 Total	6,295	9,193	9,813	11,634	9,736
Tier 2 Total	3,000	3,000	3,000	3,000	3,000
Total	9,295	12,193	12,813	14,634	12,736

11. **Has there been adequate progress made in deploying the U.S. Tier-1 and Tier-2 centers and in their integration with the CERN Tier-0 center?**
    - ***If possible, provide progress in the context of deployment in other countries.***

As indicated in the slides of various presentations, the US ATLAS facilities are performing at or above the expected level relative to typical ATLAS facilities in both the Computer System Commissioning and Data Challenge exercises. This indicated that US deployment is abreast or ahead of deployment in other countries.

12. **From a user's perspective, is the usability and readiness of grid-based production software in good shape? (The collaborations should provide sufficient information to help the committee evaluate typical user experience with grid-based tools.)**
    - ***Any metrics in this area?***

Pathena has been available to users for the past 6 months. During this time, there has been numerous well attended tutorials (I can provide some web links?), and a steady increase in the number of users using pathena. Currently, there are 77 users who have run 50k jobs using pathena. There is a Savannah bug reporting system for Panda - with only 50 items submitted during the past 12 months for Panda and Pathena.

There have been many ATLAS-wide talks by users on good experience with Pathena. (Should I make a list)? The feedback from jamborees have been positive. We have been asked by European groups to install pathena outside the U.S., so that they can use it with non U.S. resources.

We expect the number of users and jobs to increase significantly in 2007, once the data from the current Release 12 production becomes fully available.

13. **Are cybersecurity issues given adequate priority by management? On matters of security, are the lines of authority clearly spelled out? Has the collaboration assessed the impact of a cybersecurity incident on user access to data and to computing cycles? Is there a mitigation plan in place?**
    - ***Describe the cybersecurity lines of authority and mitigation plans.***

Within US ATLAS, Bob Cowles has been designated as the coordinator of cyber security. The OSG and WLCG have cyber security teams which have developed protocols for responding to cyber security incidences across these Grids. However, actual cyber security authority resides in the cyber security groups at each of the institutions (Lab or University) that make up the Grid or virtual organization. Planning factors which mitigated the impact of cyber security incidents for US ATLAS users include 1) the maintenance of a complete ESD date set at the BNL Tier 1 so that if access to 1 or more other ATLAS Tier 1's is lost the complete data set is nonetheless still accessible and 2) if there is a cyber security incident at BNL external to the Tier 1 itself, the firewall protected enclave within which the Tier 1 is located at BNL affords the same protection against such a problem as it does against the external world in general and so should allow it to continue to operate and service its user community.

14. **Have network bandwidth and connectivity requirements been appropriately identified by the U.S. collaborations? Are these requirements consistent with their latest computing models? Is there a roadmap to achieve the required T0-T1- T2 connectivity?**
    - ***Briefly describe the status of connectivity down to Tier-2 in the context of your latest computing model.***

As described in the presentations the Wide Area Network requirement for ATLAS (and US ATLAS) is derived from a hierarchical model across the various Tiers. The current connectivity is adequate for 2008 operations in most cases and is expected to be adequate in all cases by 2008. See table below:

US ATLAS Site Network Connectivity

Type of Site	Site	Connectivity Current	Connectivity in 2008
Tier-1	BNL	2 x 10GE	2 x 10GE
Tier-2	GL (UM/MSU)	10GE	2 x 10GE
	MW (IU/UC)	10GE	10GE
	NE (BU/Harvard)	1GE	10GE
	SW (UTA/OU/Langstor	1GE-10GE	1GE-10GE
	W (SLAC)	4GE	2 x 10GE

15. **Do the U.S. S&C programs have adequate links to the Worldwide LHC Computing Grid (WLCG) and to the Open Science Grid (OSG)?**

- ***Provide relevant organizational chart (if applicable)***

In the wLCG, US ATLAS is currently represented by
John Huth on the Overview Board
Bruce Gibbard on the WLCG Management Board
Bruce Gibbard on the Grid Deployment Board
In addition, Jim Shank is on the ATLAS wLCG Team—a body that advises Dario Barberis and Peter Jenni on wLCG/ATLAS interactions.
This representation is augmented by ad hoc bilateral meetings between WLCG management and US LHC computing management (some face to face but mostly via phone on a biweekly schedule when needed).

Jim Shank and Howard Gordon are on the OSG Council.
ATLAS is invested in and contributing to the success of the OSG as an integrator and provider of distributed computing capability, both in terms of facilities and software. ATLAS is represented well in OSG management. Rob Gardner is Deputy to the Executive Director and Integration Coordinator, responsible for integration and validation of middleware for OSG software releases. Torre Wenaus is Applications Co-Coordinator, responsible with Frank Wuerthwein (CMS) for science-driven middleware extensions projects, and user support and liaison. The BNL Tier 1 and US ATLAS Tier 2s are operating as OSG facilities, open to the OSG beyond ATLAS, and reporting usage to the OSG. The Panda production/analysis system is integrated with OSG accounting, and is being extended as part of an OSG Extensions collaboration with Condor and CMS to develop a just-in-time workload management service and toolkit for OSG. The OSG is funding 3 FTEs at BNL (hiring in progress) for work on this program, applications management deputy and other extensions program activities. The OSG is a positive force as seen by US ATLAS and we understand the OSG sees the converse as true also.

CORE SOFTWARE AND ANALYSIS SUPPORT

16. **Are the current models for the support of data analysis well thought out, and is the support structure responsive to the needs of the U.S. community, both in the U.S. and at CERN? Are there adequate metrics to monitor progress in this area? Have all the required resources been identified by the collaborations? Will there be adequate support during all phases of the experiment?**
    - ***Please provide specific metrics your collaboration is using to monitor this area. Also provide planned resources for the period FY06-FY11.***

The US ATLAS Analysis Support Model has been defined after multiple interactions with the collaboration and consists of three main components designed to address the needs of US ATLAS physicists:
1) The Analysis Support Group, which consists of experts in core and analysis software, detector performance and trigger, provides support for technical issues related to all aspects of performing a physics analysis;

2) The Analysis Support Centers, which consist of 3 regional centers at ANL, BNL and LBNL, designed to provide a focus for analysis jamborees, workshops and longer-term visitors;
3) The Analysis Forums, which provide a way for US ATLAS physicists to meet, discuss and develop detector performance studies and physics analyses.

During 2006, the Analysis Support Centers hosted a series of physics workshops, tutorials and analysis jamborees designed to introduce US ATLAS physicists to the overall ATLAS physics picture, train them on how to use ATLAS software and jump starting them on running and developing their own analysis code.

At each of these tutorials and analysis jamborees, participants filled out "student evaluations" with both quantitative scores and feedback questionnaires. These provide metrics to gauge the success of these tutorials and analysis jamborees. The feedback has been very positive and has also been valuable to adapt our analysis support to better serve US ATLAS physicists. We also note increased participation in subsequent jamborees by U.S. physicists and more than a 100 U.S. physicists have received significant support.

The fraction of physics working group conveners and CSC notes in ATLAS led by US ATLAS physicists is consistent or better than the fraction of US ATLAS physicists within the ATLAS Collaboration. The number of talks given by U.S. physicists at major ATLAS meetings is good, with a number of younger U.S. people who are able to make more matured presentations due to the support they have received.

As we approach first data taking and are beginning to analyze cosmic ray data, we are planning for an increase in analysis support. This is the main motivation for allocating additional FTEs at the Analysis Support Centers: 0 FTE in FY06, 2 FTEs in FY07, 3 FTEs in FY08-11.

17. **From a user's perspective, comment on the usability and readiness of the analysis software. (The collaborations should provide sufficient information to help the committee evaluate typical user experience with analysis tools.)**
    - ***Any metrics in this area?***

Usability of ATLAS software in general has been the main focal point for the US ATLAS core software developers during the past year. After the ATLAS physics workshop in Rome (July 2004), and together with the experience gathered at the Combined Test Beam monitoring and data analysis (2005), it became clear that there were two main obstacles for users of ATLAS software
i) The difficulty to understand the configuration of an Athena job and to adapt it to their specific needs.
ii) The complexity and size of the AOD format and the ability to use it in subsequent analysis with ease.

i) Job Configuration

After preliminary meetings at CERN and U of Arizona, Tucson in Spring 2005, the Usability Task Force composed of Athena developers and representatives from the users and developers communities, designed a new job configuration scheme with an emphasis on:
a) automatic generation of default module configuration and of its documentation
b) easy and safe modification of module configuration
c) restructuring of high-level job configuration in terms of tasks producing data as requested by the user. https://twiki.cern.ch/twiki/bin/view/Atlas/UsabilityTaskForce

All the infrastructure for a) and b) is available in developers' releases since months and should be put in production with release 13, currently scheduled for the week after the review. The migration of high-level job configuration has started and is a deliverable for release 14.

In parallel with the Usability Task Force, ATLAS Software Management initiated in 2005 a complete review of ATLAS software from the user's perspective: https://twiki.cern.ch/twiki/bin/view/Atlas/ReviewsPlans2005

(ii). Analysis Data Model
Over the course of last year, lots of effort from the ATLAS Reconstruction and Physics Analysis Tools groups went into unifying the AOD transient data model with the ESD one (thus allowing to develop physics code which works transparently on both). At the same time the PAT and Database groups redesigned the persistency mechanism of AOD and ESD with the goal of improving its performance . The results have been impressive, with an average order-of-magnitude improvent in data access speed, but still not quite sufficient to use AOD as main format for interactive data analysis of large data samples. To this end, ATLAS computing model introduced a new analysis format, DPD (Derived Physics Data) that has been implemented (mostly by US developers) as a Structured, Athena-aware Ntuple. SAN allows to analyze events in ROOT at ~1KHz and should soon be accessible in Athena with similar speeds. SAN data model is inspired by, and soon will be almost identical to, the AOD data model. By mid 2007 ATLAS should have achieve a grand-unification of its analysis data model allowing to use the same Athena Algorithms to analyze data in ESD, AOD and DPD format, while at the same time allowing to analyze AOD and DPD files directly from ROOT.

iii). User-level documentation
An essential component of software usability is user-level documentation. The community-updated, UK-mantained workbook (inspired by the Babar one) https://twiki.cern.ch/twiki/bin/view/Atlas/WorkBook is the main resource for Athena users. Univ of Indiana is developing a new Physics Analysis Workbook with the goal to give a new user enough information that they can start doing a physics analysis https://twiki.cern.ch/twiki/bin/view/Atlas/PhysicsAnalysisWorkBook. Finally ATLAS has always put a significant effort in users training and we have a comprehensive array of tutorials, many of which prepared by US ATLAS developers

http://atlas-computing.web.cern.ch/atlascomputing/documentation/tutorialRef/currentTutorials.php

iv) We also note that most of the software components in ATLAS have gone through a "usability" review, namely do they satisfy the user requirements and do they provide adequate documentation and understandable interfaces. The recommendations of those reviews were subsequently adopted by the developers to provide a more usable software.

v) The validation of the software and access to data still needs to be improved. Recent analysis exercises and the validation of the software have exposed difficulties. A robust validation infrastructure capable of automated validation of software as they are released needs more work to becomes the source of quickly locating and solving problems. This has led to end physicists finding many of the problems, a process that can be slow and painful. Access to data over the grid has also caused difficulties for end physicists. These problems are now actively being worked on.

18. Are the personnel requirements for the maintenance and operation phase of production software well understood, well justified and available? On what basis are commitments made to the international collaboration? Are these commitments realistic and consistent with U.S. interests?

- ***Similar to 3rd bullet in Management Section.***

The software profile for our commitments to Framework and Database effort through the operations phase is well understood. We have a minimal software effort that focuses on our key interests in these two areas. The [Athena] Framework and associated tools are primary U.S. deliverables and long term support for this has been reduced to a maintenance and support mode focusing on optimization and usability issues into the LHC turn on. The database effort focuses on providing development and support for Event base data and the profile has been kept stable well through turn-on based on past experience from other experiments.

In the production and distributed software, we believe personnel requirements are understood because distributed software has been in operations mode for some time so we understand our operations and maintenance support loads. We have a plan that shows essentially flat level of personnel for Panda development and operations. Software has been developed to minimize the operations load as we scale (automation, monitoring), this has been successful so far -- operation requirements of one shifter is low, so is well justified -- so we believe we know how to scale manageably so as to fulfill US ATLAS obligations to US and international ATLAS for production and analysis. Needed manpower is available, once the vacated shift captain slot is filled again. As we extend Panda to LCG, we will not scale US-supported operations support for LCG needs, we expect LCG participation to cover any additional ops requirements (i.e. extending to LCG

is contingent on LCG participating in ops, which is certainly doable, since Panda ops load is low and LCG resources are extensive).

Overall, our involvement in Framework, [Event] Database and distributed software form the foundation of our involvement in ATLAS core software effort and is directly related to the success of LHC physics and hence U.S. involvement in LHC physics. We have proposed and endorsed additional hires in key related areas that are beneficial to U.S. physicists. These include developing the expertise in support of ROOT and its integration into the ATLAS analysis activities and support for 3-D event displays, a critical tool for physics analysis.

19. **Is the role of Tier-3 centers well-defined and integrated into the S&C facilities plan? Are the plans for Tier-3 facilities sufficiently developed to guarantee capability for data analysis at interested U.S. institutions by November 2007?**
    - ***Provide short statement on Tier-3 plans and coordination.***

The U.S. ATLAS Management formed a panel to write a white paper on the role of Tier 3 centers in U.S. ATLAS. The complete report is on the web page of background material for this review (http://www.usatlas.bnl.gov/computing/meet/0701AgencyReview/)

The summary of that report is included here:

- Some local compute resources, beyond Tier-1 and Tier-2, are required to do physics analysis in ATLAS.
- These resources are termed Tier-3 and could be as small as a modern desktop computer on each physicist's desk, or as large as Linux farm, perhaps operated as part of a shared facility from an institution's own resources.
- Resources outside of the U.S. ATLAS Research Program are sometimes available for Tier-3 centers. A small amount of HEP Core Program money can sometimes leverage a large amount of other funding for Tier-3 centers. Decisions on when it is useful to spend Core money in this way will have to be considered on a case by case basis.
- Support for Tier-3 centers can be accommodated in the U.S. Research Program provided the Tier-3 centers are part of the Open Science Grid and that they provide access those resources with appropriate priority settings to US ATLAS via the VO authentication, authorization and accounting infrastructure.

20. **Has progress in core software relative to the milestones presented at the February 2006 comprehensive DOE/NSF review of the U.S. program been adequate? Are the forthcoming U.S. milestones on track and realistic? Is there any critical dependence on international milestones that could put U.S. deliverables at risk?**
    - ***Provide list of milestones for FY06-FY08 and comment on status.***

The milestones and status are updated in every quarterly report. Here are the software milestones form the latest complete report:

Milestone	Baseline	Previous	Forecast	Status
Release 12.0.0 for commissioning	15-Mar-06	--	30-May-06	Completed (See #1)

Milestone	Baseline	Previous	Forecast	Status
Release 12.0.x validated	30-Aug-06	--	30-Nov-06	Delayed (See #1)
Release 13.0.0	15-Feb-07	[New]	15-Feb-07	On Schedule

Note #1 Release 12.0.x continues to be validated, expected to be complete by end of 2006 and used for CSC simulation production

2.2.2 Core Services

2.2.2.1 Framework

Milestone	Baseline	Previous	Forecast	Status
port of Gaudi and Control to gcc 3.4.4	1-Feb-06	--	1-Feb-06	Completed
Full chain and Interactive Tutorials	8-Feb-06	--	8-Feb-06	Completed

Milestone	Baseline	Previous	Forecast	Status
Review ATLAS software documentation/workbook (Indiana)	1-Mar-06	--	1-Mar-06	Completed

Milestone	Baseline	Previous	Forecast	Status
Integration of Seal plug-in Mechanism	14-Nov-05	14-Oct-06	11-Dec-06	Delayed (See #1)
python-accessible Property Repository	1-Feb-06	--	1-Sep-06	Completed
use GaudiPython to provide interactive access to (selected) framework functionality	8-Feb-06	14-Sep-06	31-Jan-07	Delayed (See #2)
High-level job configuration design and tools	14-Feb-06	14-Nov-06	11-Dec-06	Delayed (See #3)
History & Property Mech Integ	14-Feb-06	--	31-Jan-07	Delayed (See #4)
flexible job reinitialization	1-Jun-06	14-Sep-06	14-Jun-07	Delayed (See #5)
port Gaudi and Control to 64-bit architectures	1-Jul-06	1-Jul-06	1-Sep-06	Completed (See #6)

Review ATLAS Python Scripts for usability (Indiana)	1-Sep-06	--	1-Sep-06	Completed
Generic event dump (needed for validation tests)	14-Sep-06	14-Sep-06	31-Jan-07	Delayed (See #7)
Enhance AtRndmGenSvc	11-Dec-06	[New]	11-Dec-06	On Schedule (See #8)
common online/offline error reporting/handling	31-Dec-06	--	31-Dec-06	On Schedule

Note #1 we implemented a prototype plugin mgr in Gaudi, still based on Gaudi DLL mechanism and the new Configurables. This will be put in production in release 13. It may be replaced in the future by a reflex-based scheme.

Note #2 started work on StoreGate access.

Note #3 Two CS summer students at LBL delivered a framework to access the PropertyRepository as their project work. Focus now is on

new-style job configuration: Job-level properties, ToolHandle, ServiceHandle (possibly DataObjectHandle), migrate all core examples to configurables, document migration procedure (basic documentation exists).

Note #4 awaits history persistency.

Note #5 some progress with framework tool but still no real-life test of the functionality. This may come with TDAQ "Large scale test" scheduled for next spring

Note #6 port completed in time. Unfortunately while Gaudi and Control compile and run fine in 64-bit machines, athena reconstruction jobs are much too big to run on lxplus. We should try to address this problem after release 13

Note #7 prototype in c++ available. During prototype development it appeared obvious that a python implementation is probably easier to write and maintain. This will be investigated till December and a final implementation will go in release 13

Note #8 automatically seed generator from the run/evt# stream ID, evaluate new random engine, possibly allow to choose random engine

2.2.2.2 EDM Infrastructure

Milestone	Baseline	Previous	Forecast	Status
Integration with POOL-Cache Manager	31-Dec-05	6-Sep-06	6-Sep-07	Delayed (See #1)
Support for History Objects	14-Feb-06	14-Jun-06	14-Jun-07	Delayed (See #2)
evaluate the need of object aliases and versioning in StoreGate	14-Sep-06	--	14-Sep-06	Completed (See #3)
Integrate CLID Database Generation	14-Sep-06	14-Sep-06	11-Feb-07	Delayed (See #4)

provide common base class for Element/DataLinks	14-Sep-06	14-Sep-06	11-Feb-07	Delayed (See #5)
support DataLinks across different stores	23-Sep-06	23-Sep-06	11-Feb-07	Delayed (See #6)
Prototype Support for Integer Keys	30-Sep-06	30-Sep-06	11-Feb-07	Delayed (See #7)
update DataList, evaluate DataMap and association objects	1-Dec-06	1-Dec-06	1-Sep-06	Completed
DataVector safe symlink mechanism	11-Dec-06	[New]	11-Dec-06	On Schedule (See #8)
New Style INavigableFourMomentum	11-Dec-06	[New]	11-Dec-06	On Schedule

Note #1 low priority

Note #2 prototyped persistency using pickle. POOL needs Transient/Persistent separation. Low priority

Note #3 still need to test alias persistency, but a scheme is in place

Note #4 Not critical for release 13

Note #5 prototype under development. Not critical for release 13

Note #6-7 Not critical for release 13

Note #8 depends on ROOT schedule

2.2.2.3 Detector Description

Milestone	Baseline	Previous	Forecast	Status
All Identifiers in DD database	1-Oct-05	1-Jun-06	1-Jan-06	Completed (See #1)
DC3 Detector Description Quality Control Tests Passed	1-Jan-06	--	1-Oct-06	Completed
All numbers in DB. All file-based info removed (Pitt)	1-Mar-06	--	1-Sep-06	Completed

Milestone	Baseline	Previous	Forecast	Status
CPU/Memory Optimization Effort Final Report (Pitt)	1-May-06	1-Jul-06	1-Jan-07	Delayed (See #1)
Initialization time performance study	1-May-06	1-Jul-06	1-Jan-07	Delayed (See #2)
Database preemption system delivery date (Pitt)	1-Sep-06	--	1-Jan-07	Delayed (See #3)
Geometry Database overriders	1-Sep-06	--	1-Jan-07	Delayed (See #4)
LAr Readout Geometry Revision	1-Dec-06	[New]	1-Dec-06	On Schedule

Readout geometry working for LAr (all clients)	1-Jan-07	[New]	1-Jan-07	On Schedule

Note #1-2 Tied to the completion of the LAr Readout Geometry Revision.

Note #3-4 Delayed till the completion of the LAr Readout Geometry Revision.

2.2.2.4 Graphics

Milestone	Baseline	Previous	Forecast	Status
Configuration Presets available in v-atlas	1-Sep-06	1-Sep-06	1-Jan-07	Delayed (See #1)

Note #1 Joe will do this over the holiday break.

2.2.2.5 Analysis Tools

Milestone	Baseline	Previous	Forecast	Status
access c++ classes from python	8-Feb-06	--	8-Feb-06	Completed

Milestone	Baseline	Previous	Forecast	Status
access Atlas Data from python	8-Feb-06	--	31-May-06	Completed
support event thinning for AOD data	14-Sep-06	14-Sep-06	11-Dec-06	Delayed (See #1)
runtime recompilation/reloading	15-Sep-06	15-Sep-06	15-Jun-07	Delayed (See #2)
automatic generation of transformations	31-Dec-06	--	31-Dec-06	On Schedule

Note #1 prototype exists. Will go in release 13

Note #2 low priority. Awaits experience with interactive athena.

2.2.2.6 Grid Integration

Milestone	Baseline	Previous	Forecast	Status
Job Transformations in Python	14-Feb-06	--	14-Nov-06	Delayed (See #1)
Prototype Implementation for Grid Monitoring Architecture	30-Sep-06	--	30-Sep-06	Completed
extract job metadata from athena	15-Dec-06	--	15-Dec-06	On Schedule
Improve message formatting and filtering	15-Dec-06	--	15-Dec-06	On Schedule

Note #1 infrastructure deployed in release 12 preproduction. Will need continuous maintenance, further development at least till release 13

2.2.2.7 Core Service Usability

Milestone	Baseline	Previous	Forecast	Status
add and use help string to Gaudi Property	18-Feb-06	--	18-Nov-06	Delayed (See #1)
Configurable-based job configuration	1-May-06	29-Aug-06	31-Jan-07	Delayed (See #2)
provide job-level properties a la SimFlags	1-May-06	29-Aug-06	11-Dec-06	Delayed (See #3)
Integrated athena configuration and development environment	31-Dec-06	--	31-Dec-06	On Schedule
Job Configuration editor/browser/debugger	31-Dec-06	--	31-Dec-06	On Schedule

Note #1 Help string handling added in gaudi. First tool using it is Job Options Inspector.

Note #2 Infrastructure almost finished. Further delays in release schedule have led to target the migration for release 13.

Note #3 prototype available. Further delays in release schedule have led to target the migration for release 13.

2.2.3 Data Management

2.2.3.1 Database Services and Servers

Milestone	Baseline	Previous	Forecast	Status
User-level documentation for access to database services	27-Mar-06	--	27-Mar-06	Completed

Milestone	Baseline	Previous	Forecast	Status
Database release model decoupled from software release	30-Jun-06	--	30-Jun-06	Completed

Milestone	Baseline	Previous	Forecast	Status
Completion of FroNTier (data caching technology) testing	1-Sep-06	1-Sep-06	31-Dec-06	Delayed (See #1)
LCG 3D service deployment at U.S. ATLAS Tier 1	1-Nov-06	--	1-Nov-06	On Schedule
Database Deployment and	1-Dec-06	1-Dec-06	31-Mar-07	Delayed (See #2)

Operations infrastructure ready for calibration data challenge				
Completion of database deployment tests via DDM DQ2	1-Feb-07	[New]	1-Feb-07	On Schedule
DDO infrastructure ready for Full Dress Rehearsal	30-Jun-07	[New]	30-Jun-07	On Schedule

Note #1 September 2006: This milestone is now a collaborative milestone including non-U.S. contributors, subject to international ATLAS schedules.

Note #2 September 2006: delayed because of delays to the international ATLAS calibration data challenge itself

2.2.3.2 Common Data Mgmt Software

Milestone	Baseline	Previous	Forecast	Status
Technology evaluation and selection for DDM catalogs	27-Feb-06	--	27-Feb-06	Completed
Transient/persistent separation of DataHeader	3-Apr-06	--	3-Apr-06	Completed
ROOT5/Reflex integration	24-Apr-06	--	24-Apr-06	Completed
Transient/persistent separation of EventInfo and TriggerInfo	24-Apr-06	--	24-Apr-06	Completed

Milestone	Baseline	Previous	Forecast	Status
Support for Placement Control	7-Nov-05	7-Aug-06	31-Jan-07	Delayed (See #1)
Support for Multiple Transaction Contexts	14-Nov-05	9-Oct-06	17-Apr-07	Delayed (See #2)
User-level documentation for Athena I/O infrastructure	27-Mar-06	17-Jul-06	28-Aug-06	Completed (See #3)
Transient/persistent separation infrastructure for heterogeneous containers	21-Aug-06	--	21-Aug-06	Completed
Extended ROOT/POOL options setting (e.g., 64-bit values)	4-Sep-06	--	4-Sep-06	Completed
Support for files larger than 2 GB	4-Sep-06	--	4-Sep-06	Completed
Refactorization of persistence packages and services	11-Sep-06	--	11-Sep-06	Completed
Address cascade ("chain") effect in transient/persistent separation	6-Nov-06	[New]	6-Nov-06	On Schedule
Complete move to Configurables, and new Property model	31-Jan-07	[New]	31-Jan-07	On Schedule
I/O performance optimization and	10-Apr-07	[New]	10-Apr-07	On Schedule

tuning				

Note #1 December 2005: Agreement with ATLAS software management is to postpone this milestone until it is specifically requested or required by physics.

June 2006: This work will be deferred until after the ATLAS streaming study group makes its recommendations.

September 2006: International ATLAS streaming study group has not yet made recommendations. In the interim, POOL placement control capabilities have been extended.

Note #2 December 2005: A limited version will be introduced into early 2006 releases, but a more general strategy awaits the outcome of a spring 2006 event store developers workshop.

September 2006: Transaction context control will require changes to the LCG POOL code base. These changes are not expected before early 2007.

Note #3 31 March 2006: Partially complete; completion pending documentation of Release 12 infrastructure, which was delayed by international ATLAS.

2.2.3.3 Event Store

Milestone	Baseline	Previous	Forecast	Status
Schema evolution infrastructure prototype based upon transient/persistent separation	27-Jan-06	--	27-Jan-06	Completed
Athena interface to set/get technology-specific attributes (POOL)	27-Feb-06	--	27-Feb-06	Completed
Transient/persistent separation of StoreGate support classes for EDM	1-May-06	--	1-May-06	Completed
Event selector with scope-based retrieval support	22-May-06	--	22-May-06	Completed
File/dataset association machinery	22-May-06	--	22-May-06	Completed

Milestone	Baseline	Previous	Forecast	Status
Job-level history object and configuration persistence prototype	2-Oct-05	18-Dec-06	10-Apr-07	Delayed (See #1)
Unique EDO Identification Infrastructure	31-Jan-06	25-Sep-06	10-Apr-07	Delayed (See #2)
Infrastructure to support stream- and run-level metadata	27-Mar-06	25-Sep-06	31-Jan-07	Delayed (See #3)
Strategy for run-level, stream-level, file-level metadata	22-May-06	25-Sep-06	18-Dec-06	Delayed (See #4)
Back navigation support for bytestream data	7-Aug-06	7-Aug-06	31-Jan-07	Delayed (See #5)

EventHeader/EventInfo extensions to support luminosity blocks, etc.	11-Sep-06	11-Sep-06	31-Jan-07	Delayed (See #6)
EventSelector improvements	11-Sep-06	--	11-Sep-06	Completed
Trigger decision representation in EventInfo/TriggerInfo	11-Sep-06	--	11-Sep-06	Completed
Event store components for streaming studies	25-Sep-06	--	25-Sep-06	Completed
Prototype alternative event streaming strategies	11-Dec-06	--	11-Dec-06	On Schedule (See #7)
AthenaPOOL/bytestream event selector unification	31-Jan-07	[New]	31-Jan-07	On Schedule
DataHeader evolution to support POOL tokens more efficiently, and to support bytestream references	31-Jan-07	[New]	31-Jan-07	On Schedule
User-level event store documentation improvements	12-Feb-07	[New]	12-Feb-07	On Schedule
Utilities and optimizations to support an efficient skimming service	10-Apr-07	[New]	10-Apr-07	On Schedule

Note #1 30 September 2005: Actual delivery may be later in the quarter because the prototype requires new versions of LCG dictionary software that will not be delivered to us until later in the quarter.

December 2005: This milestone must be delayed until some time after the corresponding control framework history infrastructure is in place.

September 2006: The corresponding control framework support has been descoped from Release 13, so event store support cannot be provided until a subsequent release.

Note #2 31 March 2006: Delayed pending outcome of May 2006 Physics Analysis Tools workshop

September 2006: Descoped from core and event data model plans for Release 13, and hence from event store work as well.

Note #3 31 March 2006: Delayed pending proposed ATLAS metadata task force

September 2006: International ATLAS metadata task force has begun work, but has provided no guidance or recommendations yet.

Note #4 31 March 2006: Delayed pending proposed ATLAS metadata task force

September 2006: International ATLAS metadata task force has begun work, but has provided no guidance or recommendations yet.

Note #5 September 2006: Implementation requires changes to bytestream I/O services provided by TDAQ. These will not be available until after the next TDAQ software release (currently proposed for November 2006).

Note #6 September 2006: Extensions sufficient to support luminosity block identification in streams tests have been added, but "final" work awaits TDAQ work on evolution of the ATLAS event header.

Note #7 September 2006: While event-store-specific components are ready, delays in preparation of data samples may cause some reduction in the scale of this prototype.

2.2.3.4 Non-Event Data Management

Milestone	Baseline	Previous	Forecast	Status
DDM support for tile commissioning data	13-Mar-06	--	13-Mar-06	Completed

Milestone	Baseline	Previous	Forecast	Status
Maintain NOVA for use by LAr calorimeter until COOL transition is complete	31-Jan-06	1-Jul-06	31-Jan-07	On Schedule (See #1)
Ensure event store integration with non-event data needed for cross section calculation	3-Jul-07	[New]	3-Jul-07	On Schedule

Note #1 September 2007: NOVA may need support for a bit longer, while COOL transition (not a U.S. responsibility) is incomplete. Forecast is based upon international ATLAS Release 13 schedule.

2.2.3.5 Collections, Catalogs, Metadata

Milestone	Baseline	Previous	Forecast	Status
Tag database infrastructure for Tier0 scaling tests	16-Jan-06	--	16-Jan-06	Completed
Tag test integration into Service Challenge 4	19-Jun-06	--	19-Jun-06	Completed

Milestone	Baseline	Previous	Forecast	Status
Athena Interface/Read/Write Access to Collection-Level Metadata	19-Dec-05	24-Jul-06	31-Jan-07	Delayed (See #1)
Integration of Collection Support & Bookkeeping	19-Dec-05	24-Jul-06	31-Jan-07	Delayed (See #2)
Evaluation of strategies for support	27-Feb-06	25-Sep-06	31-Jan-07	Delayed (See #3)

of variable-length structures in tags				
Transition POOL collections to CORAL AttributeLists and component library model	17-Apr-06	--	7-Aug-06	Completed
Support extensible collections, and collection reblocking	22-May-06	7-Aug-06	31-Jan-07	Delayed (See #4)
Integration and support of event-level metadata in commissioning tests	25-Sep-06	--	25-Sep-06	Completed
Demonstrate tag database building and selection at a Tier 2 center	11-Dec-06	--	11-Dec-06	On Schedule
Demonstrate tag-based selection for alternative streaming models	11-Dec-06	--	11-Dec-06	On Schedule
Address primary key infrastructure issues	31-Jan-07	[New]	31-Jan-07	On Schedule
Consistent reference handling, and primary reference naming	31-Jan-07	[New]	31-Jan-07	On Schedule
Propose and implement appropriate primary key for ATLAS event tags	31-Jan-07	[New]	31-Jan-07	On Schedule
Prototype tag query integration with detector status and quality information	26-Mar-07	--	26-Mar-07	On Schedule
Prototype tag query integration with luminosity block information	26-Mar-07	--	26-Mar-07	On Schedule
Implement relational collection model in MySQL	10-Apr-07	[New]	10-Apr-07	On Schedule
Prototype indexing strategies for tag database	3-Jul-07	[New]	3-Jul-07	On Schedule
Migration of POOL collections infrastructure to "component model"	4-Dec-07	[New]	4-Dec-07	On Schedule

Note #1 September 2006: An underlying component (AMI, from Grenoble) is behind schedule. New effort from Glasgow has been recruited.

The revised schedule calls for delivery of this component in ATLAS Release 13, which has itself been delayed until January 2007.

Note #2 September 2006: An underlying component (AMI, from Grenoble) is behind schedule. New effort from Glasgow has been recruited.

Note #3 September 2006: Delayed pending LCG major POOL collections infrastructure changes.

Note #4 September 2006: The revised LCG schedule calls for delivery of this functionality, along with major POOL collections infrastructure changes, in time for ATLAS Release 13 in January 2007.

2.2.3.6 Distributed Data Management

Milestone	Baseline	Previous	Forecast	Status
Bulk transfer tests based upon realistic T0->T1 model	31-Mar-06	--	31-Mar-06	Completed (See #1)
Deployment of SC4 version of DDM infrastructure on Tier 1s	31-Mar-06	--	31-Mar-06	Completed (See #2)

Milestone	Baseline	Previous	Forecast	Status
Initial tests of DDM in distributed analysis chain	17-Apr-06	--	17-Apr-06	Completed
DDM 0.2 on all Tier 1s	10-Jun-06	--	10-Jun-06	Completed
DDM support for LAr commissioning data	15-Jun-06	--	15-Jun-06	Completed

Milestone	Baseline	Previous	Forecast	Status
Set up Savannah portal for DDM operations	10-Jul-06	--	10-Jul-06	Completed
Submit note on ATLAS DDM operations to ATLAS computing management	25-Jul-06	--	25-Jul-06	Completed (See #1)
Automate input dataset delivery for production jobs	25-Aug-06	--	25-Aug-06	Completed
Data integrity checking for Monte Carlo production	25-Aug-06	--	25-Aug-06	Completed
U.S. ATLAS Data Management and Production Workshop at BNL	30-Sep-06	--	30-Sep-06	Completed
Data transfer functional test (est. 9 Tier 1s, 50 Tier 2s)	3-Nov-06	[New]	3-Nov-06	On Schedule
Implement ATLAS SW projects table in MC production	15-Nov-06	[New]	15-Nov-06	On Schedule
Data integrity procedure for DDM catalogs	30-Nov-06	[New]	30-Nov-06	On Schedule
Set up MySQL local file catalog at CERN (backup to ATLAS LFC)	21-Dec-06	[New]	21-Dec-06	On Schedule

Note #1 ATLAS note ATL-COM-SOFT-2006-008 (approved Aug 2006)

2.2.3.7 Data Access Support

Milestone	Baseline	Previous	Forecast	Status

Database representation on U.S. ATLAS Analysis Support Group	30-Jan-06	--	30-Jan-06	Completed (See #1)

Milestone	Baseline	Previous	Forecast	Status
Provide database expertise via U.S. ATLAS Analysis Support Group	2-Jul-07	--	2-Jul-07	On Schedule

2.2.4 Distributed Software

2.2.4.1 Distributed Analysis

Milestone	Baseline	Previous	Forecast	Status
Physics data access tools	9-Jan-06	[New]	9-Jan-06	Completed
Distributed analysis benchmarks for DIAL and PanDA	15-Feb-06	[New]	15-Feb-06	Completed
PanDA/distributed analysis review	8-Mar-06	[New]	8-Mar-06	Completed
Decide relative roles of PanDA and DIAL for DA	15-Mar-06	[New]	15-Mar-06	Completed

Milestone	Baseline	Previous	Forecast	Status
Physics data access tools	9-Jan-06	--	9-Jan-06	Completed
Distributed analysis benchmarks for DIAL and PanDA	15-Feb-06	--	15-Feb-06	Completed
PanDA/distributed analysis review	8-Mar-06	--	8-Mar-06	Completed
Decide relative roles of PanDA and DIAL for DA	15-Mar-06	--	15-Mar-06	Completed
Pilot delivery via dedicated analysis queues deployed	15-Apr-06	[New]	15-Apr-06	Completed
Local analysis pilot submission deployed at BNL, UTA	20-May-06	[New]	20-May-06	Completed
Personal PanDA analysis interface deployed	1-Jun-06	[New]	1-Jun-06	Completed
Stable analysis pilot delivery achieved	1-Jun-06	[New]	1-Jun-06	Completed
PanDA analysis in support of Physics Analysis Jamboree	5-Jun-06	[New]	5-Jun-06	Completed

Milestone	Baseline	Previous	Forecast	Status

Support for long user jobs via long queues	15-Jul-06	[New]	15-Jul-06	Completed
Support event collection analysis with pathena	20-Jul-06	[New]	20-Jul-06	Completed
PanDA analysis in support of Physics Analysis Jamboree	20-Aug-06	5-Jun-06	20-Aug-06	Completed
Multitasking pilot deployment for analysis	31-Aug-06	[New]	31-Mar-07	Delayed (See #1)
Initiate extension of scheduler/pilot operation to LCG for pathena	10-Sep-06	[New]	10-Sep-06	Completed

Note #1 Multitasking pilot itself was deployed in production, but the analysis multitasking feature itself remains under test and is not yet activated for general use.

2.2.4.2 Production System

Milestone	Baseline	Previous	Forecast	Status
Achieve <10% production error rate due to PanDA	31-Jan-06	--	31-Jan-06	Completed
PanDA migration to ATLAS production DQ2	10-Feb-06	--	10-Feb-06	Completed
PanDA production scalability test	20-Feb-06	--	20-Feb-06	Completed
Active disk space management	28-Feb-06	--	28-Feb-06	Completed
PanDA deployment at Indiana U	15-Mar-06	--	15-Mar-06	Completed
Brokerage based on dynamic site resource information	10-Apr-06	--	10-Apr-06	Completed
Site usage monitoring and reporting to OSG	15-Apr-06	--	15-Apr-06	Completed
Teraport site integrated	15-Apr-06	--	15-Apr-06	Completed
CERN instance of PanDA server deployed	15-May-06	--	15-May-06	Completed
PanDA monitor extended to multi-instance, multi-site deployment	20-May-06	--	20-May-06	Completed
User quota system integrated	10-Jun-06	--	10-Jun-06	Completed
New DQ2 major release integrated	15-Jun-06	--	15-Jun-06	Completed
Support grid data access from pilot jobs based on uberftp	15-Jun-06	--	15-Jun-06	Completed
SLAC site integrated	20-Jun-06	--	20-Jun-06	Completed
Support direct posix access to SE-resident data from PanDA jobs	15-Jul-06	[New]	15-Jul-06	Completed
Deploy grid data access from pilot jobs, for opportunistic site support	31-Jul-06	[New]	31-Jul-06	Completed
User quota system deployed	31-Aug-06	[New]	31-Aug-06	Completed

Define initial program in generic PanDA (OSG extensions)	10-Sep-06	[New]	10-Sep-06	Completed
Define initial program in PanDA/Condor OSG collaboration (OSG extensions)	10-Sep-06	[New]	10-Sep-06	Completed
Establish BNL Condor testbed for PanDA/Condor program	30-Sep-06	[New]	30-Sep-06	Completed

2.2.4.3 Production Support

2.2.5 Application Software

2.2.5.1 Generator Support

Milestone	Baseline	Previous	Forecast	Status
Migrate to Herwig 6.510	11-Apr-06	11-Apr-06	31-Mar-06	Completed
Migrate to Pythia 6.326	11-Apr-06	11-Apr-06	31-Mar-06	Completed

Milestone	Baseline	Previous	Forecast	Status
First version of EvtGen interface for inclusive decays with Pythia and Herwig	11-Apr-06	--	11-Apr-06	Completed
Migration of production scripts to python	30-Jun-06	--	30-Jun-06	Completed
Evaluation of migration of Sherpa into atlas releases	5-Jul-06	5-Jul-06	30-Jun-06	Completed

Milestone	Baseline	Previous	Forecast	Status
Customized persistency for HepMC	5-Jul-06	--	5-Jul-06	Completed
First version of enforced consistency of masses among all generators	5-Jul-06	--	5-Jul-06	Completed
Herwig++ available and integrated	5-Jul-06	--	5-Jul-06	Completed
Validated version of EvtGen interface for inclusive decays with Pythia and Herwig	5-Jul-06	--	5-Jul-06	Completed
Sherpa in production.	1-Dec-06	[New]	1-Dec-06	On Schedule
Evaluation of Herwig++ for inclusion in release 13.	10-Jan-07	[New]	10-Jan-07	On Schedule
Evaluation of full migration of sherpa into releases (no external files).	15-Jan-07	[New]	15-Jan-07	On Schedule

Evaluation of Pythia 8 for inclusion in release 14.	1-Jun-07	[New]	1-Jun-07	On Schedule
Complete validation of all packages to be used for first data.	1-Aug-07	[New]	1-Aug-07	On Schedule

2.2.5.2 Tracking Infrastructure

Milestone	Baseline	Previous	Forecast	Status
Understand existing b-tagging software.	1-Aug-06	1-Oct-06	1-Dec-06	Delayed (See #1)
Modify CBNT for increased functionality in monitoring b-tagging	15-Oct-06	[New]	15-Oct-06	On Schedule
Port modifications of CBNT to Athena Aware Ntuple.	1-Dec-06	[New]	1-Dec-06	On Schedule
Optimize b-tagging algorithm for release 13.0.0	26-Jan-07	[New]	26-Jan-07	On Schedule

Note #1 Best's replacement (Jain) has started on August 1 but will take some time to learn the b-tagging code.

2.2.5.3 Calorimeter Infrastructure

Milestone	Baseline	Previous	Forecast	Status
move cluster corrections into database, provide software framework for fill/retrieval	11-Apr-06	11-Apr-06	31-Mar-06	Completed
conversion/transport of mySQL etc. testbeam conditions database to COOL	1-Jan-07	1-Jan-07	30-Jun-06	Completed

Milestone	Baseline	Previous	Forecast	Status
clean-up of CaloCluster class & new fast CaloTower implementation	11-Apr-06	31-Oct-06	27-Jan-07	Delayed (See #1)
move hadronic calibration parameters into database	3-May-06	31-Oct-06	27-Jan-07	Delayed (See #2)
improvement of Navigation package wrt performance and storage requirements	5-Jul-06	31-Oct-06	27-Jan-07	Delayed (See #3)
investigate data compression for CaloCluster	5-Jul-06	31-Oct-06	27-Jan-07	Delayed (See #4)
fast CaloTower implementation	1-Aug-06	31-Oct-06	27-Jan-07	Delayed (See #5)
implementation of new EDM for navigable four-momentum typed	31-Oct-06	--	31-Oct-06	On Schedule

objects				
testbeam converter support	31-Oct-06	31-Oct-06	27-Jan-07	Delayed (See #6)
use of conditioning data in reconstruction and simulations, esp. validation of "realistic" calo.	6-Dec-06	6-Dec-06	25-Jun-07	Delayed (See #7)
optimizations for commissioning/ATLAS data	1-Jan-07	--	1-Jan-07	On Schedule (See #8)
conversion/transport of mySQL etc. conditions database to COOL.	27-Jan-07	[New]	27-Jan-07	On Schedule (See #9)
improve CaloCluster features to provide uncalibrated and calibrated signals for hadronic clusters	27-Jan-07	[New]	27-Jan-07	On Schedule (See #10)
transient/persistent object separation for CaloCluster done	27-Jan-07	[New]	27-Jan-07	On Schedule (See #11)

Note #1 Delayed to release 13.0.0. framework problems

Note #2 postponed after first implementation ideas circulated - will be picked up again mid-October 2006. Expected for 13.0.0.

Note #3 Release 13.0.0 is delayed to 27-Jan-07, ongoing,10% done

Note #4 Release 13.0.0 is delayed and so is this milestone, started in context of new navigation implementation, 75% done.

Note #5 100% done, final test postponed due to framework pbs, now expected for 13.0.0

Note #6 Release 13.0.0 is delayed, ongoing, 95% done

Note #7 Release 14 is delayed to June 2007 starting, 25% done

Note #8 discussions only at this time, no demand

Note #9 100% done, software tested and ready to go, actual conversion is on demand from testbeam analysis: presently providing "on demand conversion service" for calorimeter testbeam communities

Note #10 New activity, 90% done.

Note #11 new activity, working prototype viable, awaiting final approval from calorimeter group, 90% done.

2.2.5.4 Muon Infrastructure

Milestone	Baseline	Previous	Forecast	Status
Implement and Test new EDM for	28-Apr-06	--	28-Apr-06	Completed

Muon reconstruction for 12.0.0				

Milestone	Baseline	Previous	Forecast	Status
Implement and Test new EDM for Muon reconstruction for 13.0.0	12-Jul-06	26-Sep-06	26-Jan-07	Delayed (See #1)
Implement and Test Truth EDM for 13.0.0	12-Jul-06	31-Oct-06	26-Jan-07	Delayed (See #2)
Improve documentation of EDM packages	31-Jul-06	--	31-Jul-06	Completed
Finalize EDM design, and work on transient / persistent separation for 13.0.0	26-Sep-06	31-Oct-06	26-Jan-07	Delayed (See #3)
Integration of new muon reconstruction packages for 13.0.0	26-Sep-06	31-Oct-06	26-Jan-07	Delayed (See #4)

Note #1-2 This item is delayed because release 13.0.0 is delayed until the end of January 2007.

Note #3-4 Release 13.0.0 has been delayed to January 26, 2007.

2.2.5.5 Monitoring Infrastructure

Milestone	Baseline	Previous	Forecast	Status
first functioning Gatherer	1-Oct-05	--	1-Oct-05	Completed
first successful running of LAr partition with online display	31-Jan-06	[New]	31-Jan-06	Completed
initial discussions on data quality monitoring as separate from the Gatherer	31-Jan-06	[New]	31-Jan-06	Completed
irst successful running of complete online environment through full chain to display	5-Feb-06	[New]	5-Feb-06	Completed
first version of Gatherer for tdaq-01-06	6-Mar-06	[New]	6-Mar-06	Completed
first draft of DQM requirements document to ATLAS MWG	31-Mar-06	[New]	31-Mar-06	Completed

Milestone	Baseline	Previous	Forecast	Status
Commit of Gatherer version incorporating further Run Control-related mods. in prep. for pre-release.	31-Mar-06	--	30-Jun-06	Completed

Milestone	Baseline	Previous	Forecast	Status
Gatherer in tdaq-01-06	31-Jul-06	[New]	31-Jul-06	Completed
DQMF Design Document	31-Aug-06	[New]	31-Aug-06	Completed
working tdaq-01-06 installation for LAr	30-Sep-06	[New]	30-Sep-06	Completed

2.2.6 Infrastructure Support

Milestone	Baseline	Previous	Forecast	Status
Install and support ATLAS releases and associated external software at BNL	31-Dec-06	--	31-Dec-06	On Schedule
Support ATLAS software nightly builds at BNL and CERN	31-Dec-06	--	31-Dec-06	On Schedule

2.2.6.2 Librarian

Milestone	Baseline	Previous	Forecast	Status
Release NICOS version 1.1 with the better multiplatform support	30-Jun-06	--	30-Jun-06	Completed
Setup mirror of ATLAS Pacman cache at BNL	30-Jun-06	--	30-Jun-06	Completed

Milestone	Baseline	Previous	Forecast	Status
Install and support ATLAS releases and associated external software at BNL	31-Dec-06	--	31-Dec-06	On Schedule
Setup builds of the nightly releases distribution kits	31-Dec-06	--	31-Dec-06	On Schedule
Support ATLAS software nightly builds at BNL and CERN	31-Dec-06	--	31-Dec-06	On Schedule

2.2.7 Analysis Support

Milestone	Baseline	Previous	Forecast	Status
Formation of Analysis Support Group	15-Feb-06	--	1-Feb-06	Completed
Preparation of Analysis Support Centers Role and Implementation Document	15-Feb-06	--	15-Feb-06	Completed
Selection of Analysis Forum	15-Feb-06	--	15-Feb-06	Completed

conveners				

Milestone	Baseline	Previous	Forecast	Status
Signing of Analysis Support Center MOU	1-Apr-06	--	1-Apr-06	Completed

Milestone	Baseline	Previous	Forecast	Status
Signing of Analysis Support Center MOU	1-Apr-06	--	1-Apr-06	Completed
Full operation of Analysis Forums	1-Sep-06	--	1-Feb-07	Delayed (See #1)
First analysis jamboree at each ASC	1-Oct-06	--	1-Oct-06	On Schedule
Follow-up with analysis jamboree participants	1-Oct-06	--	1-Oct-06	On Schedule
US ATLAS Analysis Support HyperNews in working order	31-Dec-06	[New]	31-Dec-06	On Schedule
CSC Notes Jamborees	31-Jan-07	[New]	31-Jan-07	On Schedule (See #2)

Note #1 Physics Analysis Forums awaiting dedicated CSC Notes Jamborees planned for Dec 2006 and Jan 2007 at BNL and ANL, respectively.

Note #2 Analysis jamborees focused on the physics analysis CSC notes will take place in Dec 06 and Jan 07 to stimulate analysis activities.

21. Is the U.S. core software portfolio sufficiently balanced to offer U.S. researchers a good chance to participate effectively in the initial science of the LHC?

- ***Please list your priorities and metrics in this area.***

The software effort portfolio is well balanced, emphasizing core software, data access and I/O infrastructure, and distributed data management, while covering the spectrum of software needs. This effort allocation is part of a deliberate strategy that has been in place for several years, one that will pay dividends for U.S. ATLAS physics in early data taking and analysis. We plan to retain these foci in the coming years, with some shift of emphasis within these areas from development into operations and analysis support as indicated in the answer to Question 18.

ATLAS ANSWERS TO OVERNIGHT QUESTIONS
Management sub-group

1. What are the plans to test scalability of the DDM?

- Jan : ATLAS AODs replication
    - All MC Production AODs files will be replicated to all Tier-1s sites. For US ATLAS cloud we will replicate data not only to BNL, but also to all Tier-2s
- Files Catalogues stress tests
    - Check performance and robustness of File Replicas Catalogs used by the experiment (the tests are in progress and they are conducted in collaboration with SW developers)
- Feb : Tests of DDM central datasets catalogs (ORACLE implementation and new schema).
    - Set up DDM test-bed at CERN (already exists for US ATLAS)
- Feb-May : Regular DDM functional and performance tests
    - All ATLAS sites. Data distribution functional tests to check new version of DDM catalogs.

2. What is the plan for continuing to provide for analysis in the event of US Tier 1 unavailability?
    - Up to 8 hours: no serious problems
    - From many days to 1 week or more: we overflow buffers at CERN for real data and at T2's for simul. Data.

Depending on the period of unavailability there are various concerns and differing appropriate courses of action.

If the Tier 1 becomes unavailable for a period of order 8 hours as was the case during the cyber security stand-down at BNL back at the beginning of November or even long say 1-2 days, the impact is modest. The Tier 2's have adequate local disk storage to hold the Monte Carlo they are generating until the Tier 1 is again available. They also have local copies of those AOD, TAG, PDP, and ESD sets of most current interest and so can continue most analyses. The Tier 0 at CERN also has adequate disk storage to buffer the raw and ESD data destined for BNL until it is again available.

If the Tier 1 becomes unavailable for a period of many days, a week or more the problem becomes very serious on many fronts. The Tier 2's can of course send the Monte Carlo data the are producing to other Tier 1's for storage and it is also within the capabilities of the distributed data management system to allow the US Tier 2's to subscribe to receive the data they need for analysis, AOD, TAG, etc. from other Tier 1's. More of a problem is the fact that the US Tier 1 is expected to be storing and then reconstructing 25% of the ATLAS RAW data. CERN only has adequate disk buffering to hold the data destine for the BNL Tier 1 for a couple of days. All contingency planning within WLCG for Tier 1 unavailability has assumed down periods of order 24 hours not a week or more. Given its size, if the US Tier 1 were to go offline for a period of weeks, it would produce a major disruption to the whole of ATLAS computing, conceivably even data taking.

At the BNL Tier 1 great effort is being made to assure its cyber security, even in the eventuality of incidents elsewhere at BNL. To data this effort has been very effective. The real answer to the above it that there is need for adequate push back against punitive cyber security actions to the effect that the services of the US Tier 1's are scientifically very important and an integral part of international agreements and so can not be disrupted except in the case of actual technical justification. This push back needs to come from the scientific community, its program management and its agency oversight and be directed at over reaching cyber security establishments at the institutional or agency levels.

3. What are the roles and expectations from the tier 3s? How do you support the overhead to make Tier 3 resources effective?

Tier 3 centers provide local users with facilities to develop analysis algorithms plots for publication.
One of the definitions of a T3 in US ATLAS is that it is NOT funded from the RP. We expect some effort form RP funded people to help setup T3's. As mentioned this is easiest if the form of that help is to install the OSG software stack. We think this will be a minor additional workload. Most support effort for tier 3's will come from the local group it serves.

4. What would it take to complete the CDC tasks by August, 2007? What are the three critical dates/milestones for the CDC?

CDC has several components and stages. The first stage is to test the basic conditions database infrastructure, both in terms of deployment and for access by reconstruction algorithms to both static and time varying conditions and geometry information. Later stages are to exercise this in the context of the 24-hour latency that's part of the computing model before the Tier-0 processing of raw data begins.

The first stage has already been performed for all detector subsystems, although some will not be able to handle time varying information until release 13 in March. However, all have already demonstrated the ability to reconstruct using a different statically misaligned geometry from that which was used to simulate the data.

The major deliverables for release 13 are:

o Complete migration to COOL for conditions DB
o Completion of support for time-varying rather than static misalignments

The Tier-0 infrastructure already implements the 24-hour delay and handling of the calibration processing, albeit today with dummy algorithms. It is intended that the Tier-0 test scheduled for May will exercise the real calibration processing infrastructure.

5. What issues have been uncovered in any end to end tests, such as the proposed cosmic test from 2006?

The primary feedback from various standalone or combined cosmics tests that took place during 2006 and are ongoing in 2007 have been:

- Exercising the geometry database access with real detector hardware
- Exercising the conditions (calibration, cabling maps and alignment) database

infrastructure with real detector-derived information
- Exercising the event data model from bytestream raw data through to ESD/AOD
- Exercising the detector monitoring and reconstruction algorithms
- Performing combined monitoring and reconstruction
- All experience has been fed back into the primary software development.

6. Can you please clarify the milestones and impact of the full dress rehearsal on other deliverables?

Because the full dress rehearsal is really a demonstration of full-chain readiness for data-taking, it imposes few requirements that would not need to be met for 2008 running in any case. The main risk is that it could, because of its summer 2007 time frame, push development, deployment, and testing of functionality not needed for 900 GeV running ahead of functionality that _is_ needed for 900 GeV running, in a period of already-intense development pressure. With careful scoping of the dress rehearsal, this risk can be mitigated. The current strategy of expanding the ongoing Data Streaming Test to adiabatically meet the goals of the final dress rehearsal is designed to address these concerns while still meeting the FDR goals.

The culminating milestone of the dress rehearsal is successful completion by late summer 2007 of the exercise described in Srini's talk. Intermediate milestones are the subject of a planning meeting that will take place on 24 January.

7. What are the mechanisms and plans for getting long-term ATLAS support for PANDA?

There currently are 3-5 people working on support for the other 2 grid flavors. As the production system moves over to all PanDA worldwide, we expect to have these people available to support the ATLAS production system.

8. Are the major issues for experiment readiness being sufficiently addressed by ATLAS? What mechanisms are used to reassess requirements for experiment readiness? Are they effective and efficient?

    - Yes. The SPMB(with broad rep.) and the CMB effectively accomplish this.

The major issues for readiness are discussed at either the Software Project Management Board (SPMB) or Computing Management Board Meetings (CMB) which are held in alternating weekly meetings. The SPMB, which has representatives from every major body in ATLAS - such as core software, databases, sub-systems, trigger, etc., focuses on software deliverables for major production releases which are driven by major milestones (e.g. combined cosmics tests, final dress rehearsal, calibration data challenge, ATLAS low energy running). Each group prepares a list of milestones for the next production release which are followed to ensure progress and therefore readiness for the data taking. The CMB focuses on the global computing model and deployment and operations aspects of computing.

Facilities & Grids sub-group

1. Quantitatively clarify the effect of the LHC startup delay on the T1/T2 facility ramp up. What are the changes in the T1/T2 plans?

The first two tables below describe roughly the incremental cost for the BNL Tier 1 in adapting to meet the new ATLAS estimate of requirements first in absolute dollars and then in per cent change. The second set of two tables attempt to remove most of the effects aside from the actually running schedule change. You will see that with all effects taken into account there are small savings in 2007 and 2008 and significant additional costs in the out years. If only the schedule change is considered the cost savings are quite substantial in all years except 2010.

Cost Increment to New Model from Old (@Yr $k)

All effects

	2007	2008	2009	2010	2011
Labor (Fully loaded salaries)	0	0	0	0	0
MST (travel, maint, licen, etc)	-122	-190	-154	342	16
Facility Space & Power	0	0	0	0	0
Capital Equipment	8	-56	723	1236	512
Total	-114	-246	569	1578	528

% Cost Increment

All effects

	2007	2008	2009	2010	2011
Labor (Fully loaded salaries)	0%	0%	0%	0%	0%
MST (travel, maint, licen, etc)	-12%	-13%	-10%	23%	1%
Facility Space & Power	0%	0%	0%	0%	0%
Capital Equipment	0%	-1%	23%	33%	20%
Total	-2%	-3%	6%	16%	6%

Cost Increment to New Model from Old (@Yr $k)

~ Running schedule change only

	2007	2008	2009	2010	2011
Labor (Fully loaded salaries)	0	0	0	0	0
MST (travel, maint, licen, etc)	-295	-225	-238	219	-66
Facility Space & Power	0	0	0	0	0
Capital Equipment	-1488	-661	-257	57	-236
Total	-1783	-885	-494	276	-301

% Cost Increment

~ Running schedule change only

	2007	2008	2009	2010	2011
Labor (Fully loaded salaries)	0%	0%	0%	0%	0%
MST (travel, maint, licen, etc)	-28%	-16%	-15%	15%	-4%
Facility Space & Power	0%	0%	0%	0%	0%
Capital Equipment	-67%	-17%	-8%	2%	-9%
Total	-28%	-9%	-5%	3%	-3%

2. The T2 profile is over-target in 2007 and under-target after 2009. What funding profile is needed to meet projected resource targets (Bruce's slide 10)? What are the expectations that leveraged resources will meet this target?

Assuming that of the $3,000k per year going to Tier 2's approximately $600k of it is going into personnel costs (6 FTE's) leaving the remaining $2,400k for equipment, it would require ~doubling this equipment level to produce the target capacities over the three years 2009, 2010, 2011. So the total Tier 2 cost would go in 2009 from its current $3.0M/Yr to $5.4M/Yr to achieve the target capacities. This assumes the leveraged dedicated capacity would scale with equipment expenditure. It does not take account of non-dedicated capacities which might be significant but whose actual level can probably not be accurately estimated until the first couple of years of actual high demand operation.

3. What validation process is planned to verify analysis capabilities at T2s in advance of data taking?

Validation will be through deployment and operation of Panda/pathena based analysis at T2s well in advance of data taking. Principal issue is deployment of AODs to the Tier 2s, which is starting now with UTA
and will be extended to all Tier 2s over the next ~2 months; once AODs are at each site, pathena will be activated there.

4. What is the fraction of pathena (grid-submitted) analysis vs. local analysis?

This information is not immediately available. The following can be said. At the Tier 1 in November, 90,000 PanDA production jobs were run, 27,000 PanDA analysis jobs were run, and 70,000 other ATLAS jobs were run. It is not possible to determine from the current accounting summary what fraction of these "other" ATLAS jobs were analysis as opposed to simulation. Note that this is accounting by job not CPU or wall clock. Such accounting seems appropriate for typically short running analysis jobs. It appears that all that can be said is that ~1/3 and perhaps more likely 1/2 of the analysis jobs being run are using pathena.

5. How certain is the expectation that event size (ESD & AOD) can be reduced sufficiently? What is the plan if event size (ESD & AOD) stays much larger than computing model predicted sizes? When will the decision be taken to synchronize the model and reality? What are the current sizes?

ESD and AOD have expanded in size and in content because physicists have argued vigorously that they need this information to set up the experiment. The role of software experts is to help physicists increase the data content/size ratio (as it has been done with the transient/persistent separation), to provide tools that will help them making the necessary hard choices by measuring and monitoring data sizes down to the data member level, and to limit the impact of this hard choices on the algorithmic by promoting wherever possible the use of a common interface between ESD and AOD data objects.
Current ESD size is ~1.5MB (target steady-state size 0.5MB) and current AOD size is ~250KB of which 40% is MC truth(target size 100KB).

6. What are the current comparisons between empirical CPU usage and computing model predictions?

The following two plots are taken from the WLCG accounting summary as of the end of November. This information is in the context of Tier 1's only. There is currently no such comprehensive accounting for Tier 2's being done by the WLCG. They are first a summary for BNL and then a summary for the entire WLCG (there is no ATLAS only plot). The blue line shows the expected (MOU agreed) level of wall-clock time, the green line shows the installed level and the bars should the actual usage level. Usage levels can exceed installed capacity because of the possible use of opportunistic resources and/or because of operation at an efficiency greater than the assumed efficiency (85%). Usage significantly below the installed levels can be a result of lack of demand from the experiments or inefficient use of computing resources by the production management system. You will see that the BNL Tier 1 has installed its MOU committed capacity and in the Oct & Nov time frame utilized approximately that full capacity. The overall WLCG does not quite but is near meeting its MOU committed installed capacity but has to date only been utilizing about 2/3's of that capacity.

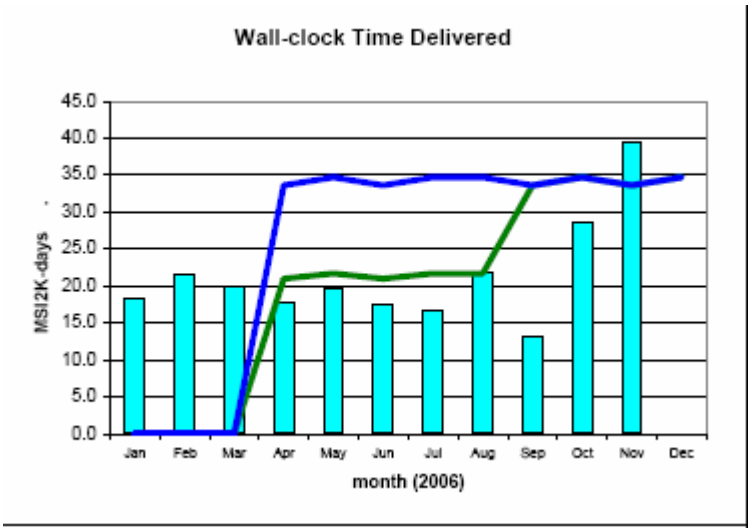

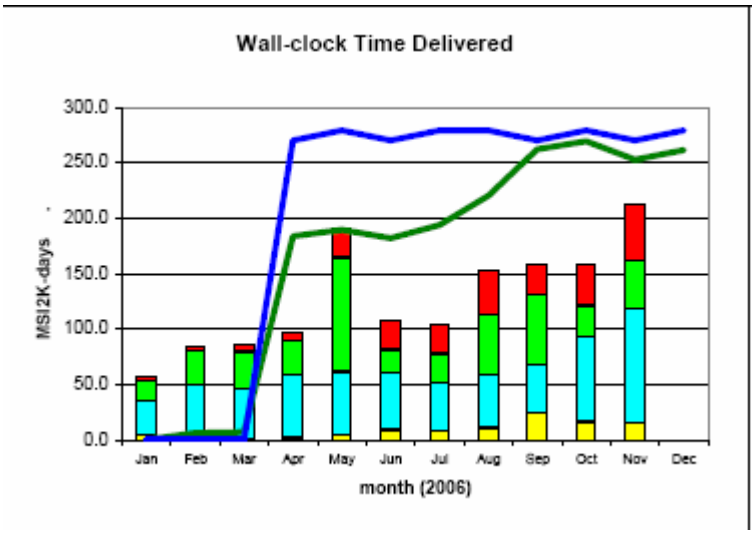

7. What motivates keeping both ATLAS' AOD and DPD? Is this a change in computing model? How does this comport with previous experience of other experiments?

AOD and DPD have been part of the computing model since this was written down in the Software and Computing TDR. At that time it was thought the main difference among the two was access speed and direct browsability of DPD from ROOT. While the gap in access speeds and ROOT functionality has narrowed, the DPD/AOD difference in scope has grown: we now expect that physicists will use AOD mainly for high statistics systematics and reconstruction tuning work and for the more complex analysis. DPD will be the choice of physicists doing simple particle-level analysis work. Another important difference is that one or few AOD will be centrally maintained and produced, while most physics groups and possibly many individuals will define and produce their own DPD.

8. There is evidence that some ATLAS users still eschew use of the Grid, and that some ATLAS T2s can be underutilized

```
without apparent explanation. How does this compare with
your presented demo? How do users submit trouble tickets
and/or convey these experiences?
```

ATLAS has gone through a series of data production exercises – DC1, DC2 CSC11, CSC12. With each exercise, it has become obvious that the required scale of data processing in ATLAS can only be achieved through the grid. Grid systems contributed ~20% of DC1 production, 100% of DC2 and 100% of CSC exercises. Only a small fraction of data processing in ATLAS is done outside these exercises, using non-grid resources. A few people still eschew the grid – but they are a very small minority of users, with special use cases, or with access to large local resources for private use. For the vast majority of ATLAS users, the grid is the only choice for doing physics.

The Tier2 utilization rate has been very high with Panda. We do not have accounting of idle CPU's from all sites, but the rate of production is well matched to available capacity. Exceptions occur due to scheduled downtime for upgrades, site problems, data movement problems, or availability of jobs. Unused CPU is rarely due to Panda problems – only recent exception being a few days due to Panda Database problems.

Consider for example, the following plot from SLAC T2. The red dots show average number of running jobs, which fluctuate considerably because of shared resources. The green dots show the number of Panda pilots waiting in queue. Note it is mostly constant at 50, the limit we set. As long as there is any green dot greater than zero in this plot, all available CPU at the site is being used. There were some exceptions during the holidays, when we had DQ2 problems:

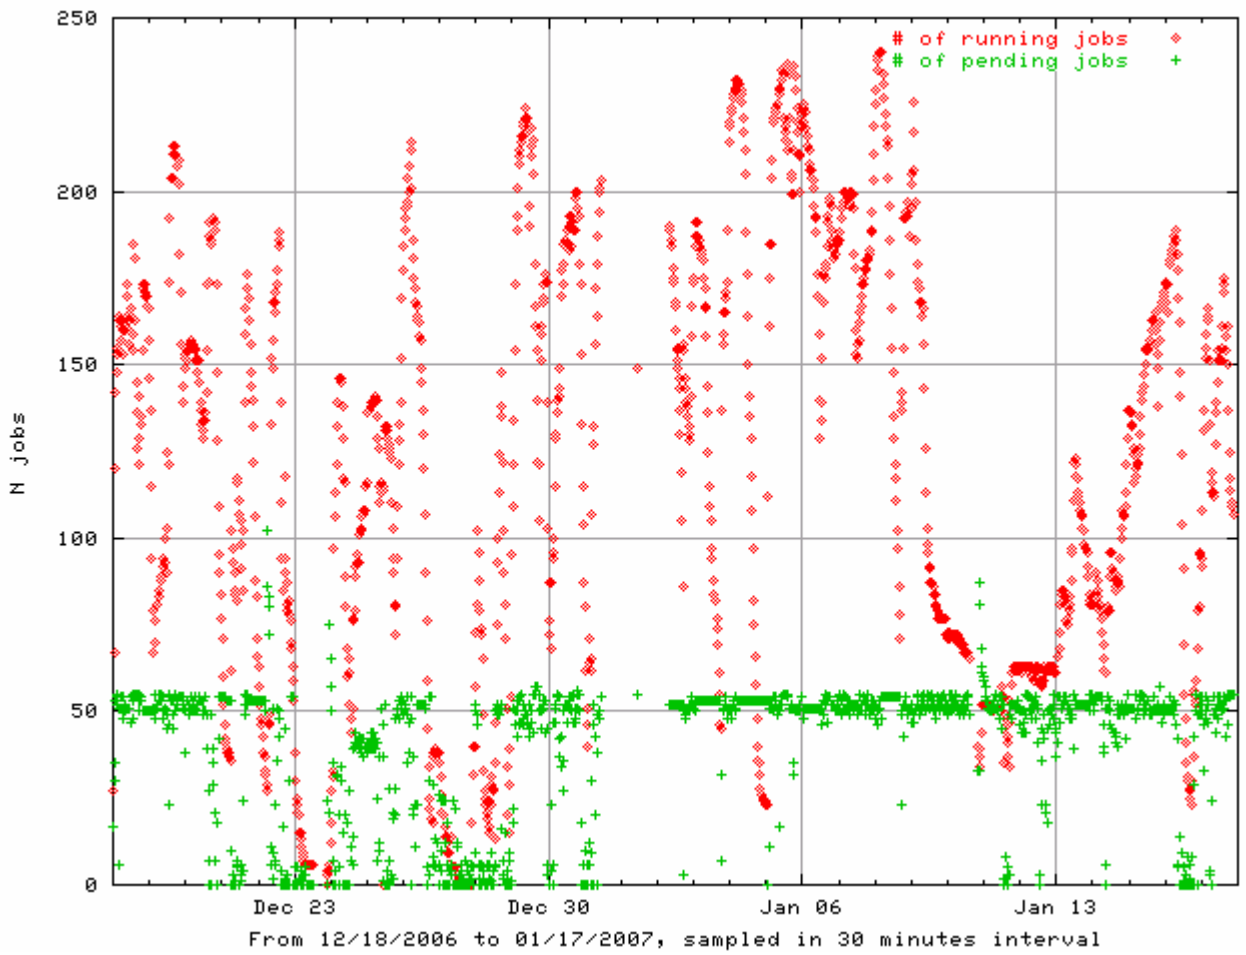

Panda is still a relatively new system for users. Few people (~75 analysis users compared to ~1800 collaborators) have tried Panda. User feedback has been mostly positive. We have set up wiki pages, a Savannah bug reporting system, a very active listserve, and a RT trouble ticket system for Panda.

9. Show Tier2 utilization for the last several months.

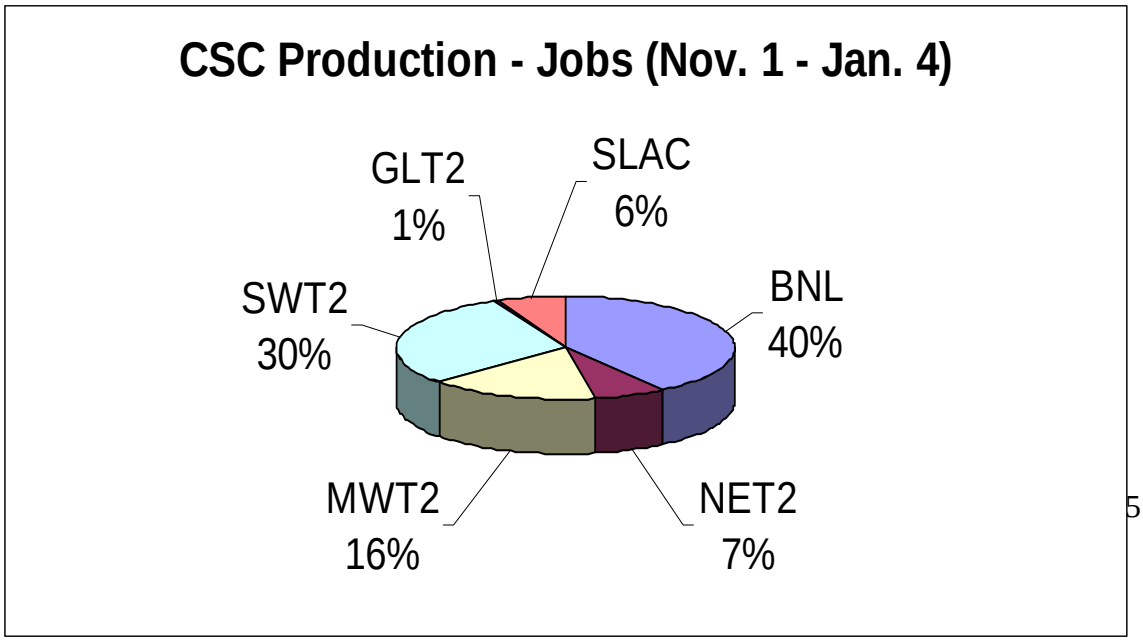

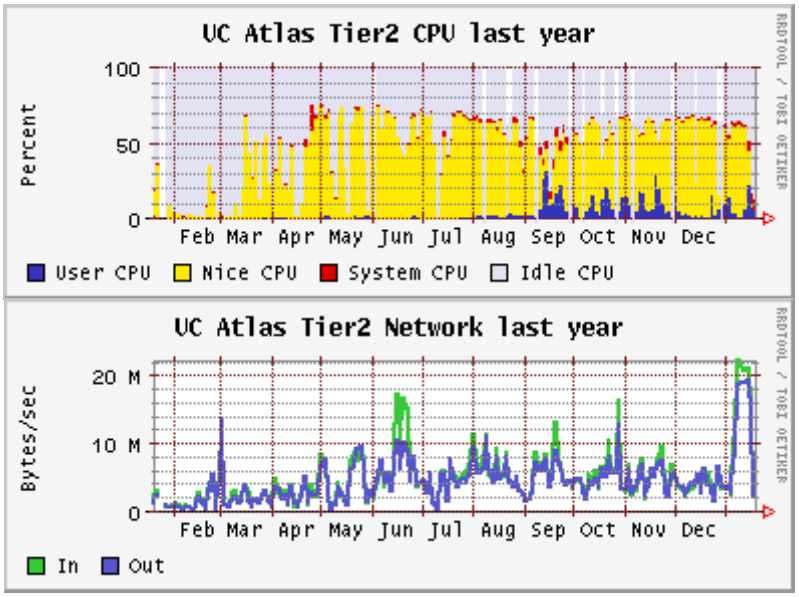

Software sub-group

1. The Full Dress Rehearsal seems unobtainable with currently allocated resources. Is this true of other important tests/milestones in 2007/2008? What are the minimum acceptable required tests for success:
   - with first beams in 2007?
   - with the first physics run in 2008?
   - IE. In triage-mode, what would be deferred on the current list of responsibilities?

Sufficient resources exist for the completion of other tests (listed below) and the strategy for FDR s being re-worked to not overburden these resources. These exercises must be completed by Summer 2007, do not foresee any major exercises after the engineering run.

The other important tests/milestones in 2007 include:
- The Computing System Commissioning
- The Calibration Data Challenge
- The integrated cosmic ray run
- The TDAQ Large Scale Tests
- 900 GeV engineering run.

All of them are ongoing and planned for completion by Summer 2007. Many of the software infrastructure components (framework, conditions database, schema evolution, I/O optimization, data replication) are common issues to all of them. These are being addressed by dedicated group of people (U.S. plays a major role in framework, Event Database and production software) with all these exercises in mind. Their work on this is not five-fold or even doubled because of several tests. The actual tests are decouple in many cases - decouple implying that different groups are addressing different tests. For example the cosmic ray run is mostly handled by dedicated people within sub-systems, the TDAQ LST by the trigger community, etc. Hence we believe that these tests have reasonably equipped resources to achieve successful completion.

As noted in the response to item 6 in the Management section, the strategy for the final dress rehearsal is to adiabatically extend the existing Data Streaming tests and to couple it to other tests such as the Tier-0 scaling and distribution test scheduled for May 2007. We believe this will minimize the disruption to the ongoing ATLAS commissioning activities, while maximizing the usefulness of FDR.

The time between the 900 GeV run and the 14 TeV run is a critical phase that will essentially be devoted to code optimizations and to rectify many of the problems we will encounter during the 900 GeV run - from pathological detector readouts to algorithm optimizations. It is not possible to foresee dedicated large scale exercises to be carried out during this critical phase. Smaller component tests could be carried out to demonstrate resolution of problems that we would have encountered during the 900 GeV running.

2. Please clarify the effort required to establish the DPD and outline the boundaries of your responsibility. How do you leverage the experience gained in existing experiments?

Structured Athena-Aware Ntuples can be (and are) written out using the current production release. We expect them to be used by physicists for their physics analysis work on CSC Monte Carlo sample. The work which remains to be done between now and this spring (release 14) is to complete the Transient/Persistent separation of AOD data objects and then to merge the classes used for the persistent representation of AOD with the corresponding SAN ones.

Extensive input from other experiments was obtained by the Analysis Model working group as can be seen by the agenda here: http://indico.cern.ch/conferenceDisplay.py?confId=6866 where talks were given by experts from D0, CDF, BaBar, H1, CMS .

3. Did you perform a requirements analysis for the development of vATLAS? If yes, what are the requirements? Did your resource estimate result from these? Demonstrate that 1 FTE is sufficient.

The requirements were drafted at a joint meeting between USATLAS and ATLAS management and the outcome of that workshop was to establish 1 FTE to provide the requested integrated framework view. Sub-system non-experts will be able to provide “plug-ins” for their individual sub-system views.

A meeting in Brookhaven in May 2006 included the Technical Coordination Leader and representatives of USATLAS. The meeting identified specific areas in which the existing 3D graphics were insufficient for the purpose of a monitoring event display. The chief difficulties were the simultaneous availability of several standardized detector views, and the requirement of a high degree of coordination between two dimensional and three dimensional views. A plan was developed to provide the missing functionality. It entails additional work at the level of the framework and the GUI (both operating within ATHENA). A three month period was envisioned for the design phase (Boudreau and Kittelman). About one month after that (4 months into the project) specific displays for one or two detector subsystems are foreseen to be available. One year into the project, good displays of all the major systems are

working within the monitoring framework. Two years into the project, the subsystem-specific displays have been improved after exposure to a year of actual detector operations.

The price of this extra functionality is the additional manpower that we requested, which is tailored to the requirements requested by ATLAS technical coordination. The system is foreseen to be open so that contributions from nonexperts could be easy to sollicit and incorporate, once the new framework was mature enough. However the framework itself and a reasonable level of sophistication in the display can be assured by a combination of Boudreau, Kittelman and Tsulaia.

4. Will International ATLAS provide sufficient user support so that US ATLAS support will not be unfairly burdened? What is the quality of the user workbook? Is it sufficient to cover standard user questions?

Various UK groups have official responsibilities in this area (workbook, tutorials) and they are putting quite some effort into user support especially for what concerns beginners. The workbook is indeed a valuable tool for beginners and together with the tutorials, goes a long way into addressing typical beginners questions. US ATLAS leads the work in core software (framework and event store), and physics analysis tools, hence US ATLAS developers tend to field many of the support requests from advanced users or from other developers.

5. What is the status of the calibrations database? Is it part of the US responsibilities?

The calibrations database is well advanced, and is being stress tested in the calibration and alignment data challenge. The underlying temporal database infrastructure comes from an LCG common project (COOL). Replicas of the content of the calibration database will be deployed in the U.S., but the database itself is not a U.S. deliverable.

6. Quantify job & software failure rates. How do they compare to the established targets? What are those targets?

   The requirements coming from physics (measurement of integrated luminosity) is that untraceable event losses due to software failure not exceed few percent (based on Tevatron experience). We are not there yet after every release – however this rate is achieved after an extensive software validation process.

   Every software release undergoes validation, over many months. During this period, various bugs are fixed. The release is considered to be ready for production when the error rate reaches less than a few percent due to software errors. An intensive robustness and performance period is planned for mid 2007, aimed at reducing software-induced job failure rates, improving the cpu & memory performance, and reducing e.g. memory leaks. Although these are all ongoing activities, the intention is to make these the highest priority for a preiod of several months, even at the expense of improved functionality. Examples of decreasing failure rates after each validation release (for the Release 12 series) can be seen in the plot below (showing successful vs failed walltime):

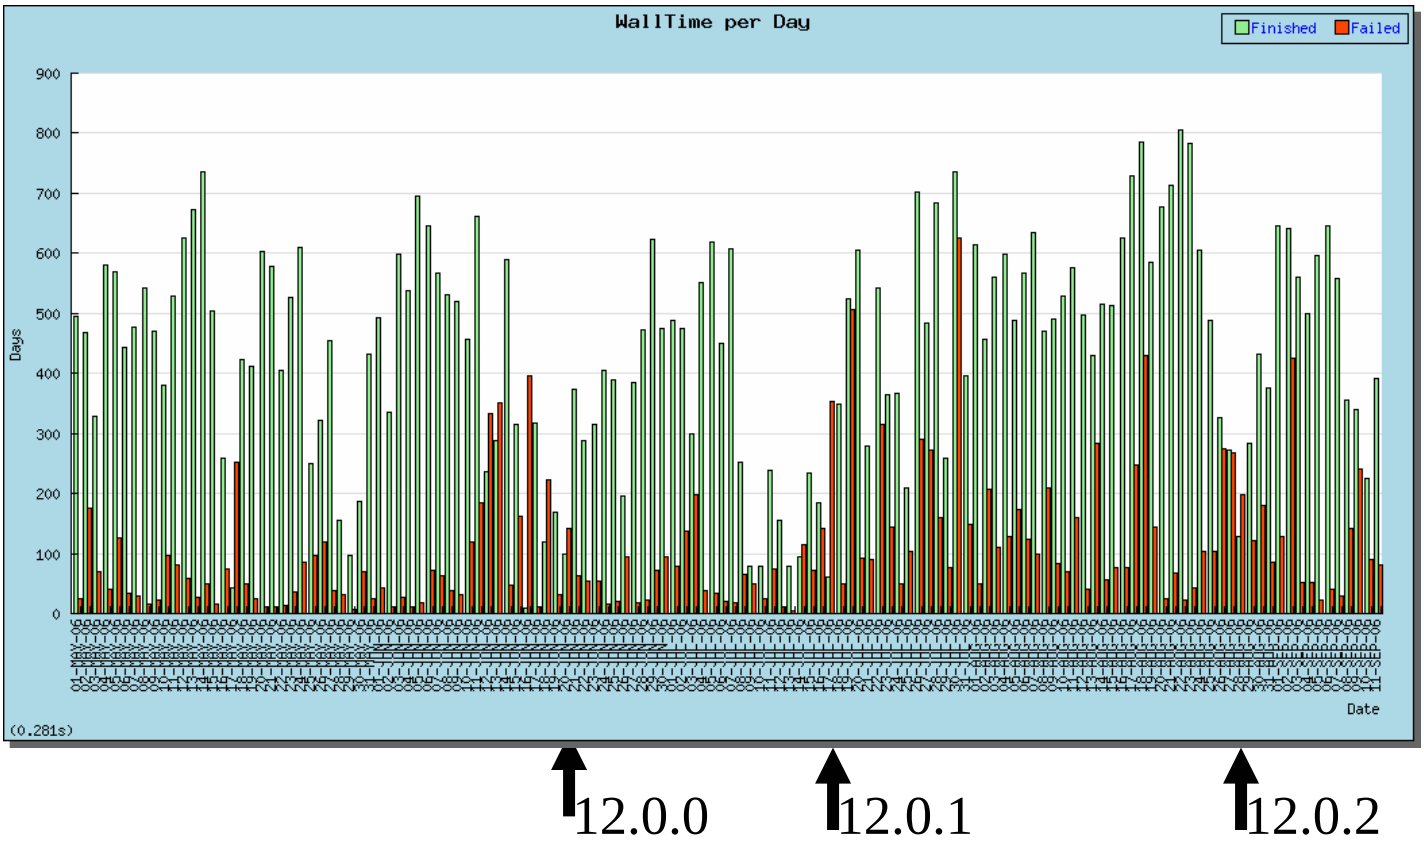

7. How long would it take to generate 10M dress-rehearsal events with the current system?

Our current ATLAS-wide production rate is ~2M events per week. This includes a factor of two increase in rate achieved this past winter. At this rate, we can generate the dress-rehearsal sample in 5-6 weeks. However, we expect to achieve another increase of factor of two by summer 2007. If achieved, the dress-rehearsal sample can produced in 2-3 weeks.

Q: "What are US ATLAS' dependencies on OSG and how is OSG bringing added value to your Software and Computing program?"

US ATLAS depends crucially on the OSG infrastructure. All our facilities have the OSG software stack as the base upon which we install the ATLAS software system. In moving to the PanDA system over one year ago, we eliminated our dependency on SOME OSG software components, limiting our dependency to only those components that had been shown to be robust and scalable. PanDA was designed to be modular enough to be able to plug in software components from middleware projects, like the OSG, as these components pass scalability tests.

We've sought to minimize grid middleware dependencies in general for well known reasons. Where we still have middleware dependencies, we often still have problems. The OSG has been helpful in mitigating or minimizing the problems. Examples:

- GRAM dependency in CondorG submission of pilots. Not sufficiently scalable to support analysis in Panda. OSG program includes a 'pilot factory' effort to work around this by doing site-local pilot submission without every pilot seeing the gatekeeper and GRAM.
- glexec. This has been validated and deployed for production by FNAL/CMS. It will be added to the OSG software stack so we can pick it up easily for integration ourselves via the OSG stack.
- OSG-standard site configuration. Provides a 'known' environment on OSG WNs for more homogeneity in execution environment. Still a bit rough but will be more useful as it improves. This lessens the application-level work of establishing this homogeneity.
- Tools for resource discovery. The tools we have are provided by OSG, and while they need improvement, we'd be in much worse shape without them; it would be very difficult to gather the info on resource availability, health, and access rights that is required to run a distributed system on all the resources available.

A new OSG component that we are particularly interested in is the Workflow Management System which we (T. Wenaus) are working on now with the OSG. This is an important piece of PanDA and we have a temporary solution now which will work when data first starts to flow from the LHC. Thus our dependency on this component is not critical, but we would really benefit from a common OSG developed solution. We are working to make this happen and will deploy it as soon as it is ready.
Another area in OSG we are very interested in is the attempt being made in the OSG in the extensions area and elsewhere to have OSG 'add value' in the area of storage management. For example, by taking a close look at dCache, code and performance, to look for ways to improve it; and by developing a monitoring system providing uniform grid-wide monitoring of dCache systems. At BNL we (ATLAS and STAR) are also interested in getting xrootd into the mix. The OSG may not have much impact on SRM 2.2 deployment timeline and quality, but even if it doesn't, there are other ways in which we're getting value out of OSG.
We will also benefit from the new OSG-wide accounting services being introduced in the next release of OSG, as well as the information services that will provide OSG site information in both Condor ClassAd and ldiff (for LCG interoperability) formats.
We benefit and rely on the VDT and OSG packaging, installation, and configuration processes that lead to a well-documented and easily deployable OSG software stack, and the OSG integration and validation processes that precede incorporation of new services into the VDT.
We also benefit from and rely on the infrastructure maintenance aspects of the OSG such as the GOC and others that keep the OSG facility as a whole operational. We have PanDA running on ALL OSG sites (not just ATLAS) are poised to take advantage of opportunistic resources on the OSG to do further Monte Carlo simulation for ATLAS.

We should point out that the OSG 'dependencies' could be a bit misleading because the dependencies don't come from OSG, they come from the grid middleware and facilities underneath, and the OSG is a layer in between where the experiments work together and with the grid projects to mitigate the problems and the risks coming from the layer below. The signs so far are that that is working.

5 APPENDIX E

CMS ANSWERS TO QUESTIONS PRIOR TO THE REVIEW

Department of Energy/National Science Foundation
U.S. LHC Research Program Review, January 2007
University of Texas at Arlington

Responses to Questions pertaining to the charge of the review

QUESTIONS TO MANAGEMENT

Question M1) Are the current management structures and techniques well-matched to the needs of the U.S. collaboration?
Please provide an organization chart of US ATLAS / US CMS

Answer M1) The organizational chart for U.S. CMS Software and Computing is shown here:

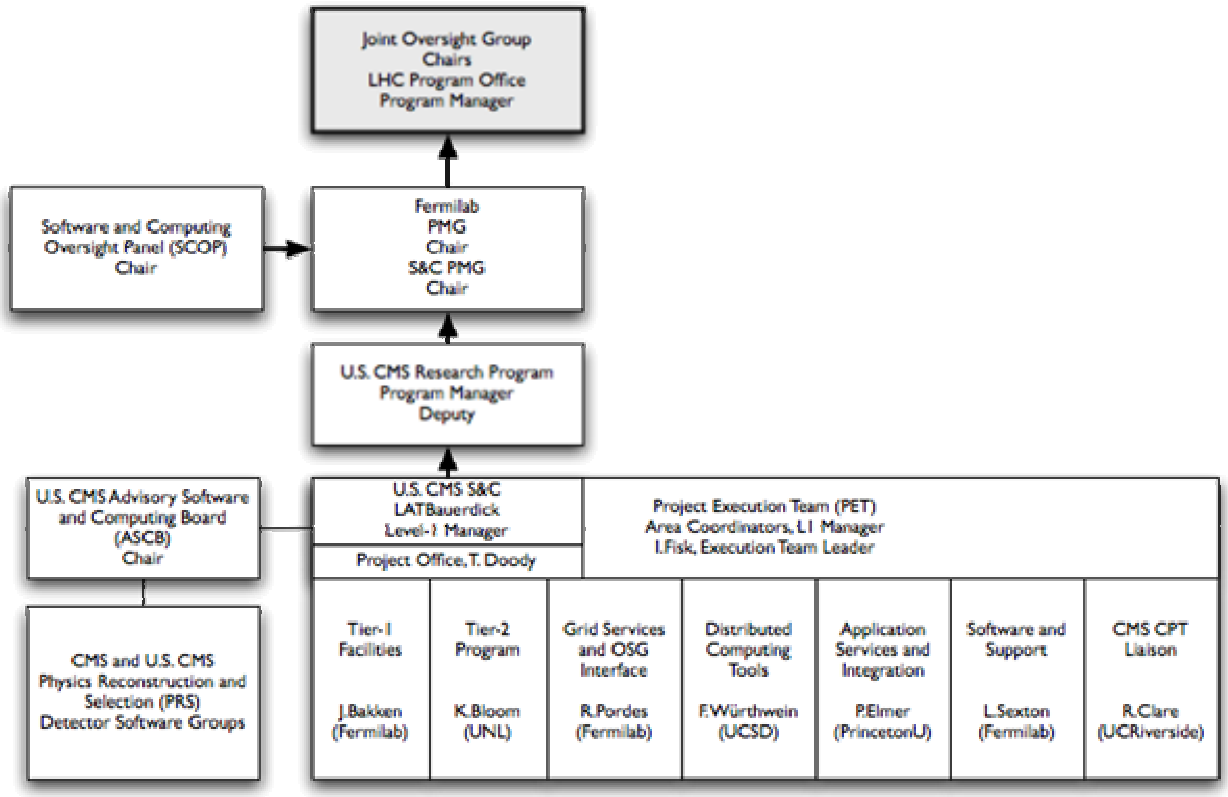

Question M2) Are their internal contingency and risk-management mechanisms appropriate?
Please provide a table with S&C funding history FY2004-FY2006. If applicable, provide initial calls on M.R. and granted M.R per year.

Answer M2) Table of S&C Funding History FY04-FY06:

DOE Funding Available	**$5,609k**	**$7,546k**	**$11,861k**
DOE FinPlan	$4,060k	$7,000k	$10,500k
Carry-Over from previous year	$1,549k	$546k	$1,361k
Systems & Operations (UF personnel)	$2,011k	$2,856k	$2,956k
Software & Support (CAS personnel)	$1,216k	$1,983k	$1,756k
Tier-1 Equipment	$1,508k	$1,144k	$3,012k
Project Office	$327k	$201k	$464k
DOE Funds Total Spent	**$5,062k**	**$6,184k**	**$8,188k**
NSF Funding (RP, DISUN, iVDGL)	$1,295k	$4,500k	$4,690k
Tier-2 Labor	$720k	$2,063k	$2,149k
Software & Support	$480k	$370k	$670k
Tier-2 Equipment	$95k	$1,490k	$1,501k
Project Office, Reserve		$577k	$370k
NSF Funds Total	**$1,295k**	**$4,500k**	**$4,690k**

In U.S. CMS the management reserve is part of the yearly allocation of funds for the S&C (and M&O) sub-projects. Project management applies management reserve through the change request mechanism, a list of change requests is given below. Over the last three years most of the management reserve had been applied for increasing the equipment funding for the Tier-1 center, given the change of the estimates of CMS computing resource needs for the Tier-1 processing and analysis resources (LPC-CAF). This is shown in the following table:

Table with list of change requests

Date	Institution	Amount	Reason
7//22/04	Fermilab	$383,000	Additional T1 equipment, Robotic tape storage, User disks, Network reliability upgrades
2/17/05	UCSD	$70,000	Additional funds to upgrade local network infrastructure at UCSD Tier 2 center
4/6/05	Princeton	$95,400	To fund an additional .5 FTE for FY05 (P.Elmer) on Application Services WBS 1.5
4/20/05	Iowa	$75,000	To fund an additional staff on workload management
5/1/05	Virginia Tech	$8,000	To fund additional/extended effort Virginia Tech on priviledge project
6/24/05	Fermilab	$63,000	To fund additional equipment for Tier 1 center WBS 1.1
8/10/05	Cornell	$60,986	Cornell work on CMS core software WBS 1.6
8/10/05	U. Chicago	$50,000	University of Chicago work on Edge Services
8/29/05	Fermilab	$13,000	To fund consulting (databases) work by S. Iqbal at CERN
9/21/05	Fermilab	$38,000	Push forward the STK tape purchase to FY05
11/30/05	Fermilab	$1,600,000	To cover increased costs of computing equipment over a 3 year period.
3/10/06	Fermilab	$1,013,000	To cover computing & networking equipment for the LPC-CAF
3/10/06	Cornell	$300,000	To fund 2 FTEs working on WBS 1.4 Application Services for 1 year at Cornell
3/30/06	U. Chicago	$50,000	For continued work on Edge Services WBS 1.5

Question M3) Are there adequate plans for transitioning from a development phase to a deployment and operations phase? Are the assumptions for resource requirements well justified?

Provide a table with personnel requirements for the period FY06-FY11; estimate distribution of FTEs: Universities, Labs, US-based, CERN-based.

Answer M3)

The table of U.S. CMS Software and Computing personnel is given below. The categories are "Systems & Operations", which includes system administrators, system integrators and trouble shooters, and the computing, storage, infrastructure (security, accounting etc), and grid related development and integration efforts, at Fermilab and at Tier-2 /Tier-2C centers (2 FTE operations and system support for each of the Tier-2 sites). The other category is "Software & Support" which includes the application software developments and maintenance, down to the Application services (WBS 1.4), the physics support etc.

Most of the university personnel is actually based at CERN or/and is traveling regularly to Fermilab, both places have critical mass of developers, project coordinators and physicists.

While we are still ramping the facilities efforts in FY07, from then on we plan for a flat overall profile in the effort. Within this envelope the maintenance, support and trouble shooting functions will increase when the development efforts goes down.

In particular the CMS computing integration and operations tasks in CMS will require manpower support from S&C that will have to come out of such a transition of function and people.

U.S. CMS S&C Personnel	FY07	FY08	FY09	FY10	FY11
Total Effort	**62**	**67**	**67**	**63**	**63**
Systems & Operations	**39**	**42**	**42**	**38**	**38**
Fermilab based	18.25	21.25	21.25	21.25	21.25
Tier-2 based	18	18	18	14	14
University based	2.75	2.75	2.75	2.75	2.75
Software & Support	**23**	**25**	**25**	**25**	**25**
Fermilab based	11.5	11.5	11.5	11.5	11.5
University based	11.5	13.5	13.5	13.5	13.5
of these, at based at CERN	10	10	10	10	10

Question M4) Are the priorities of the S&C program conducive to effective participation in data analysis by U.S. physicists?
Please list your priorities and metrics in this area.

Answer M4) The main focus of U.S. CMS Software and Computing is on delivering a useable, fully functional computing and software environments to U.S. physicists, with increasing levels of functionality and usability, and in supporting CMS members to use it for physics. This means a focus on integrating the systems we build into functional environments, while components get further developed to increase functionality. This approach is apparent in the choice of milestones, the focus on data challenges and production and the choice of the specific U.S. contributions to the overall CMS software and computing efforts (like the framework, data management, storage systems etc). This focus was helped by several U.S. S&C leaders taking up roles in the CMS software and computing organization, and by shaping it such that software releases and computing integration was focused on delivering an environment that CMS physicists can already use for data analysis.

The main priority for the software release plan and the computing integration plan thus was getting ready for the data challenges:

* MTCC "Magnet Test and Cosmic Challenge" during the summer of 2006, that involved reading out a slice of the detector and fully reconstruct data with the new CMS framework, and the transferring data to remote operations center for data monitoring.
* CSA06 Computing, Software and Analysis challenge that exercised a large part of the CMS data and workflows: from Tier-0 prompt reconstruction, calibration, to data transfers to Tier-1 and down to each Tier-2 center, with individual users running analysis on data they skimmed from Tier-1 centers and downloaded to Tier-2 sites.

All these challenges, and the milestones leading up to them, including the software release plan and the pre-challenges had well-defined metrics to establish that effective remote participation was possible:

- for MTCC, data was analyzed at several US universities and at Fermilab, including use the Remote Operations Center (ROC) for quasi-online data quality monitoring,
- for CSA06 U.S. physicists participated in the analysis challenge and all U.S. Tier-2 centers hosted analysis datasets.

For the coming year this approach will continue with an “always on” functional system and a series of challenges that lead into real data taking later this year. The main “deliverables” to ensure effective participation of U.S. physicists are

1. well-functioning Tier-2 centers and LPC-CAF (that will replace the already existing User Analysis Facility at Fermilab) supporting some 100 active users
2. scalable and well performing job submission and monitoring systems, supporting some 100k jobs per day in the system
3. reliable data transfers to make sure data will be accessible for physics users skimming specific dataset from large FEVT/RECO or AOD datasets hosted at Tier-1 centers into analysis dataset that then get downloaded to Tier-2 centers, and to push down the data from the Tier-0 to the Tier-1 (e.g. MTCC) and from Tier-2 to Tier-1 (MC samples)
4. usability of the framework, documentation (workbook), training of individual users, monitoring the usage of the new CMS software and the user feedback
5. coordination of the placement of datasets and other elements in the physics infrastructure

Question M5) Does management have adequate S&C plans to accommodate new collaborators? Have they developed a reasonable model for the corresponding incremental costs?
Summarize your model for incremental cost of new collaborators

Answer M5) The incremental computing costs for new collaborators have been studied, and in our model the hypothetical incremental cost per additional active physicist would be about $20k/year, using the latest cost estimates. However, this number is just a pro-rata of the analysis facility costs (above a large fixed cost for infrastructure and for data processing needs) and not necessarily a realistic estimate of incremental cost for additional individuals: adding a member to an already existing group will not necessarily increase the number of skimmed datasets needed, etc because of a lot of sharing of resources amongst group members with common interest etc. Also, the model of resource needs is not at all precise enough to estimate within 10%, and we do not expect fluctuations in the number of collaborators even at the 10% level.

For these reasons, and also because we fundamentally build the facilities “to cost” anyway, we do not foresee having to adapt the size of the resource to eventual additional collaborators. In our plans the facilities are sized for about 500 active physicists, corresponding to the numbers of U.S. CMS collaborators that we expected to have at the start of data, and we have no indication yet that we will deviate too much from that rough number.

Question M6) What would be the impact of a 10% S&C funding shortfall on current U.S. deliverables and on productivity in physics analysis?
Provide specific impacts assuming “reasonable calls on MR/contingency” are granted.

Answer M6) The total allocation for S&C from the U.S. CMS Research Program in FY2007 is $14.13M, and there is also $2M in the DISUN project. The general impact of a 10% cut in FY07 has been studied last year, and the contingency plan foresees to delay purchases of Tier-1 equipment, delaying the ramp-up of the facility, and possibly cuts in the Tier-2 program. We analyzed this again and came to the following conclusions:

A 10% cut in the RP funding corresponds to $1.4M. With the current cost estimates and facility deployment plan we keep a $1.1M management reserve in S&C, and such a cut would eat up the management reserve. We are also planning to carry over $3.5M into next year, to fund the large procurements for the 2008 facility. This allows some flexibility to respond to eventual shortfalls during FY07.

On the Tier-2 side, with seven Tier-2 centers and the DISUN commitments our funding is stretched too thin to allow significant cost savings short of closing one of the centers, or delaying all FY07 hardware procurements. Both would have drastic impact on the program, in particular as Tier-2 MC resources are committed to CMS in the WLCG MoU and the effective cut would mean significant impact on the data analysis opportunities for U.S. physicists. Thus, for a short-term cut we would not touch the Tier-2 funding but would need to identify savings in other areas.

To recover some costs we would save funds through delaying personnel ramp-up to FY08, and/or delaying to refill positions, thus avoiding personnel layoffs as much as possible. We would try to recover additional reserve by down-sizing the Fermilab facility ramp and trying to catch up by delaying procurements. Because of the central importance of the Tier-1 processing and data serving capabilities for the whole CMS system, we would have to cut the LPC-CAF first, with proportional impact on analysis capabilities of U.S. physicists.

The impacts of such cuts have been studied in some detail by the facility group, and several courses of action have been considered:

- Reducing the LPC-CAF disk by 20% would gain $150k. Using cheaper disks for the LPC-CAF instead was also considered, but leads to no significant savings due to the additional operations and maintenance manpower required (expected to be 1 additional FTE!).
- Reducing the LPC-CAF CPU resources by 20% would free about $280k.
- Delaying procurements of 125 Tier-1 worker nodes into FY08 may yield $110k.
- Delaying 300TB of Tier-1 disk procurements into the next fiscal year would free $100k, however with a large risk of not being able to study and remedy expected scaling behavior of the storage systems (last year's review had given a recommendation to not delay these disk procurements).

These cuts would add risks to the ramp-up of the facility and would only be done in case of urgent need for management reserve, which otherwise would be used to mitigate most of the hypothetical cut.

Question M7) Do the U.S. projects interact sufficiently with the international S&C efforts? Does the U.S. play a role in the international S&C leadership that is commensurate with its overall participation in the experiment?
Provide International ATLAS/CMS organizational chart.

Answer M7) The CMS Offline and Computing organization chart is shown below, U.S. people are highlighted in red (in the on-screen version...).

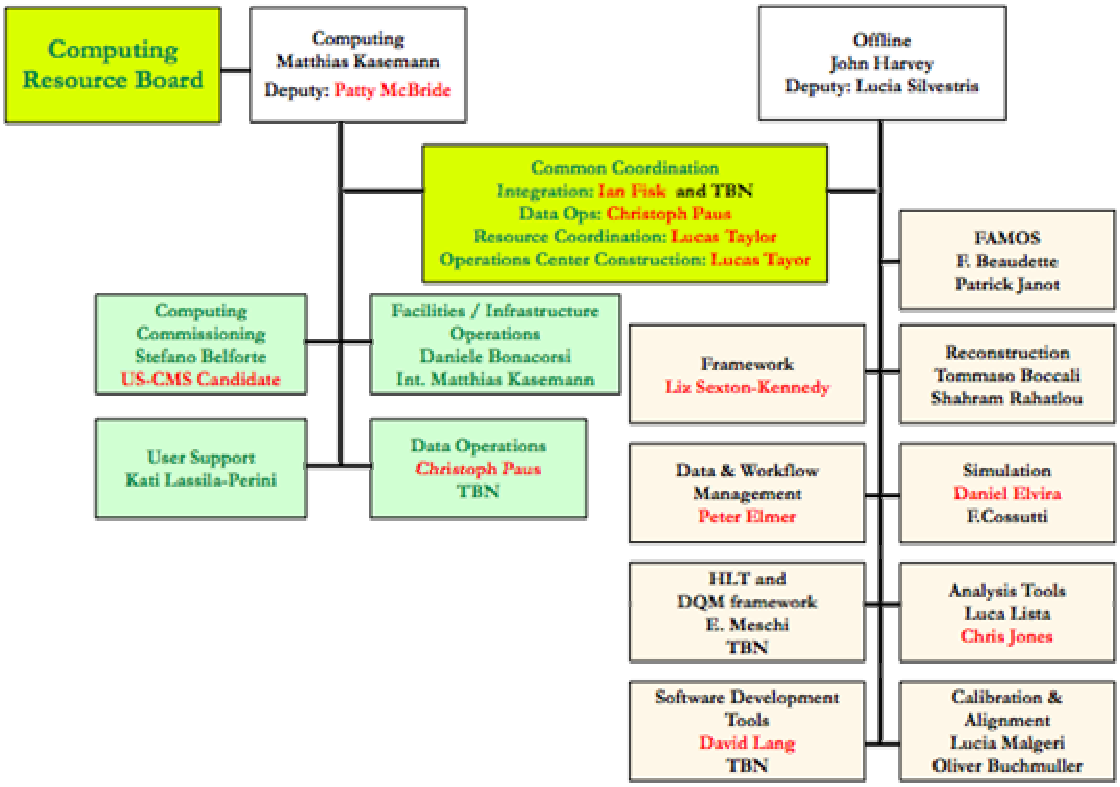

QUESTIONS TO FACILITIES, GRIDS, NETWORKING, AND INFRASTRUCTURE

Question F1) Are the current computing models of the experiments appropriate for U.S. needs?
Provide links or supporting material describing the computing model.

Answer F1) The main references for the computing model are the Computing Model CMS-Note CERN-LHCC-2004-35 (Dec 2004) and the CMS Computing TDR CERN-LHCC-2005-23 (June 2005). Both documents are linked to the documentation section of the review webpage.

Question F2) Have infrastructure and operating costs of the Tier-1 and Tier-2 facilities been fully considered in their deployment? Are there any high-risk assumptions? Are the estimated personnel requirements, equipment and infrastructure costs valid and well-justified?

Provide a table outlining all facilities costs (T1, T2), namely, personnel requirements and infrastructure costs for the period FY06-FY11.

Answer F2) The following table shows the facilities costs and personnel requirements for the Fermilab facility (Tier-1 + LPC-CAF) and the set of Tier-2 centers.

U.S. CMS Computing Facilities	FY06	FY07	FY08	FY09	FY10	FY11
Fermilab Facilities Cost	$3,160k	$3,778k	$5,789k	$5,000k	$5,150k	$5,305k
Tier-2 Facilities Cost	$1,569k	$1,740k	$1,740k	$1,740k	$1,750k	$1,750k
Fermilab Systems & Operations Effort	16.25	18.25	21.25	21.25	21.25	21.25
Tier-2 Systems & Operations Effort	18	18	18	18	14	14

The costing model details for the Fermilab facility will be described in the facilities parallel session.

The cost developments and performance increases of CPU and disk have been significantly less favorable than anticipated in the Jan 2006 cost estimates. With the delay of the LHC machine startup and the start of the initial physics run in spring 2008 the Fermilab facility procurement plan now foresees to significantly delay procurements for the “2008 facility” (sized to support the 2008 run) such a large part will be procured with FY08 funds. Despite the machine schedule now allowing to delay significant procurements of CPU and disk well into CY2008, the total accumulated (multi-FY) estimated cost for the 2008 facility has increased by ~10% to 12.4M (was 11.2M). Reasons are the slow introduction of quad-core CPUs while CPU speeds almost stagnate, and the long time it takes for the 1TB disks to get to market.

These unforeseen cost increases are only partially offset by the procurement delays. However, in terms of cash flow, with the procurements crossing the FY boundary into FY08 we expect to have enough funds to build the 2008 facility.

The budget model for replacements/upgrades of the Tier-1 and LPC-CAF beyond the ramp-up to the initial “2008 facility” is a flat budget of $5M/year of LHC operation (including operations costs like media and tape library maintenance).

Although the precise machine schedule for these years of running are not yet defined we expect start of data taking each year to be “early”, so that some equipment procurements need to happen in the previous year to commission the upgraded computing and storage capacity. The budget foresees about 1/3 of the costs to occur in the previous FY.

We studied in particular for the Tier-1 facility, using updated cost and technology projections, and with an operations-driven replacement cycle (typically replacing equipment after 4 years), if indeed the facility would scale up according to the estimated increases of CMS resource needs every year.

Similar to the Fermilab facility, in the out years FY2009 to FY2011, the budget model for Tier-2 centers is to maintain a budget of $500K/center/year for the upcoming five years, covering both Tier-2 operations costs and equipment initial procurements and replacements.

With the experiences of a slow-down of CPU and in particular disk developments over the last years, these cost developments are less favorable than originally anticipated. However, we estimate that until about FY2011 the decrease of CPU and disk costs per performance unit will allow staying on top of the CMS resource requests.

However, as we have seen in previous years, there are large uncertainties on these projections.

Question F3) Has there been adequate progress made in deploying the U.S. Tier-1 and Tier-2 centers and in their integration with the CERN Tier-0 center?
If possible, provide progress in the context of deployment in other countries.

Answer F3) The recent data challenge CSA06 provides opportunity to compare the progress of U.S. Tier-1 and Tier-2/Tier-2C centers with the international efforts, in terms of concrete metrics of available resources, throughputs, productivity for MC production and analysis job execution etc. In most of these metrics the U.S. both the individual sites and in particular the ensemble of U.S. Tier-1 and Tier-2 sites stands head and shoulders above other countries. As an extreme example, the U.S. Tier-2 centers in terms of computing resources are similar in scale to the UK Tier-1.

U.S. Tier-2 sites now provide some 3 MSI2k CPU power and 0.7 PB total storage, dedicated to U.S. CMS use. U.S. sites have produced 21M of 50M events produced in preparation of the CSA06 challenge. Daily averaged throughputs for data transfers totaled to up to 230 MB/s over the 7 sites, and in some cases sustained transfer rates of above 300MB/s (2.5Gbps) were achieved from Fermilab. Also, data transfers from all non-U.S. Tier-1 sites to all U.S. Tier-2 sites have been achieved, so access to the total CMS data sample (which will be distributed over all Tier-1 sites) has been demonstrated. In total, CMS transferred more than a PB of data during CSA06.

Regarding the Tier-1 center, the facility is now at 2MSI2k with 1800 batch slots, and has 0.7PB dCache-based grid connected storage. All goal metric for CSA06 have been achieved, and in most cases exceeded. Transfers from the Tier-0 happen routinely now, with throughputs above what is needed for 2008 sustained for tens of days. Regarding analysis running, very high throughputs have been achieved for the Fermilab storage systems, above the throughputs of the Run-II experiments, demonstrating the ability of the Fermilab facilityto support analysis running at the LPC-CAF.

Question F4) From a user's perspective, is the usability and readiness of grid-based production software in good shape? (The collaborations should provide sufficient information to help the committee evaluate typical user experience with grid-based tools.)
Any metrics in this area?

Answer F4) CMS' Grid based production software consists of organizing the processing activities with a system called ProdAgent. This includes event selection (skimming), event re-processing, and simulated event production. The ProdAgent system is a US deliverable. In the CMS computing model organized processing is normally operated by select groups of representatives from the experiment and analysis groups. This model was successfully tested in CSA06 with 4 simulated event production teams generating the events needed for CSA06 during July and August of 2006. The entire production was completed on time using the available distributed resources. Event selection and re-reconstruction was also performed with ProdAgent during CSA06 and the processing goals were achieved. The CMS Computing Integration program has tested and monitored a set of metrics, in particular the success rate of distributed jobs, which in CSA06 generally above 90%. We also monitor the success of the end-to-end workflows using a site availability metric.

We will demonstrate the use and usability of the production systems in the parallel session during the review.

Question F5) Are cybersecurity issues given adequate priority by management? On matters of security, are the lines of authority clearly spelled out? Has the collaboration assessed the impact of a cybersecurity incident on user access to data and to computing cycles? Is there a mitigation plan in place?
Describe the cybersecurity lines of authority and mitigation plans.

Answer F5) Cybersecurity is integrated through the US CMS organization.

Security for the US CMS Tier-1 Facility is integrated into the Fermilab Computing Division Security infrastructure and program. The CD Security program implements an automated scanning and host based intrusion detection processes, analyses all network flows passing through the Fermilab border router (see http://security.fnal.gov/).

The Tier-1 facility audits and monitors all services for security and includes such matters in the weekly Tier-1 operations meetings. Joe Kaiser, a system administrator in the Tier-1 facility is a member of the Computing Division Security Working Group.

The storage systems at the Fermilab Tier-1 (Enstore, dCache) provide for data and meta-data integrity through:

- File protection procedural policies
- File removed by user are not deleted
- Physical protection against modification of filled tapes
- "cloning" of problem or tapes with large mount counts
- File protection software policies
- Extensive CRC checking throughout all transactions
- Automated periodic CRC checking of randomly selected files on tapes
- Automated CRC checking of files in the disk caches.
- Automated disabling of access to tapes or tape drives that exhibit read or write problems.
- Redundancy in metadata database information (file, volume, pnfs) and automated database checks,
- Metadata backups and journaling – 1-4 hr cyclic and archived to tape.

Data transport using GridFTP to/from dCache has selectable CRC checking.

The WLCG has a set of security policies and agreements agreed to by CMS as part of the collaboration.
http://proj-lcg-security.web.cern.ch/proj-lcg-security/security_policy.html
https://edms.cern.ch/file/428035/LAST_RELEASED/Incident_Response_Guide.pdf

Open Science Grid has a risk assessment and security plan for operational, management and technical controls for the core OSG assets (common services, software caches, software stack) and guidance, awareness training, utilities and operational templates for the Site and VO administrators and management.
http://osg-docdb.opensciencegrid.org/0003/000389/013/OSGSecurityPlan-2-3-1.pdf
https://osg-docdb.opensciencegrid.org:440/cgi-bin/RetrieveFile?docid=440&version=1&filename=OSGRiskAssessment.pdf (protected)

OSG incident response is stimulated by 24x7 coverage, initiated from the Grid Operations Center. The Security Officer, Don Petravick, leads the determination of the nature of and appropriate response to security events. This determination and response involves the OSG core staff and OSG site and VO Support Center security contacts. OSG collaborates with the EGEE and WLCG with joint or agreed upon policies.
http://osg-docdb.opensciencegrid.org/0000/000019/002/OSG_incident_handling_v1.0.pdf

Response and mitigation to an incident involving a user access is removal of access privileges. This is accomplished through removal of the user from the CMS Virtual Organization.

OSG has experienced 3 alarms for security events in the past six months. One affected middleware, a second was an alert from the WLCG of a potentially misused user account, and the third a potential local account compromise on a machine which hosted grid certificates. In no case was an intrusion detected as having occured. In each case we:

- ensured broad notification of the OSG participants and partners, including the EGEE and WLCG. A meeting of the OSG Security team was convened within a few hours of notification.
- discussed, researched, and then responded appropriately to the risk and the vulnerability identified.
  - In the middleware case patches were developed and thoroughly. In practice testing and code review determined that the initial patch was not complete and another round of patch development was done. Then the patch was made available to the OSG sites and the Operations Center monitored for installation.
- held a root cause and debriefing session to determine future best practice.
  - for the account compromise we identified some needed scripts/utilities that would make an intrusion more easily and automatically detected, and added these to the program of work.

Question F6) Have network bandwidth and connectivity requirements been appropriately identified by the U.S. collaborations? Are these requirements consistent with their latest computing models? Is there a roadmap to achieve the required T0-T1- T2 connectivity? Briefly describe the status of connectivity down to Tier-2 in the context of your latest computing model.

Answer F6) Here is a table of Tier-2 site connectivity

US CMS T2 site	Local Connectivity	Connectivity to FNAL
Caltech	10 Gbps	CENIC – PacificWave – ESNet
Florida	10 Gbps	Florida Lambda Rail – ESNet
MIT	1 Gbps, upgr. to 10 Gbps	NOX – ESNet
Nebraska	0.622 Gbps, upgr. to 10 Gbps	Greatplains – Abilene – ESNet
Purdue	10 Gbps	Starlight – Fermi Light Path
UCSD	10 Gbps	ESNet/sdsc – ESNet/SDN/chi
Wisconsin	10 Gbps	Starlight – ESNet/SDN/Chi

Regarding the networking provisioning CMS has the following priorities:

1) The ingest of data from the T0 to the T1 through the Optical Private Network (OPN)
2) All other production in the formal data management system -- traffic between T[12]<->T[12].
3) Any other traffic.

The OPN has two functions, with T0 -> T1 transfers being the main task, while its function for inter-Tier-1 connectivity is secondary. Because of that, non-OPN networks may be used for T1 to T1 transfers on a routine (as opposed to a backup) basis.

The OPN is an overlay network. In the US, there are two segments involved in its provisioning. US LHCNET is used for transatlantic provisioning. It is present at MANLAN and STARLIGHT. BNL and ATLAS reach US LHCNET via ESNET MAN's.

US Tier-1 to Offshore Tier-2s; Offshore Tier-1s to US Tier-2s
This is an important CMS use case because the data is distributed over **all** Tier-1 centers. Only complete AOD samples will be cached at local Tier-1 sites, all FEVT/RECO and MC data will be available only once across the ensemble of CMS Tier-1 centers (most likely at Fermilab due to the relative size of the US Tier-1).
There is a set of concerns related to these:

- These flows are, and always have been served by the ESNET/GEANT/I2. We expect to saturate the available Trans Atlanic connectivity used by these networks. While the OHEP funds Trans-Atlanic connectivy, That connectivity is only to CERN. (The LHCOPN provides some generalization of that, but that generalization is not available for these data flows.)
- There needs to be some way to provide resources for these transfers. The path being investigated is to contribute bandwidth from the USLHCNET Manhattan - Amsterdam link.
- Any diminishment of USLHCNET resouces must not impact transfers from CERN to US T1's.

Networking for access to data hosted in Europe:
Since the LHC OPN mission is T0<->T0 and some of T1<->T1, the networking provisioning for T2's needs its own consideration. There are two cases:
1) T2 connectivity to CERN.
2) T2 connectivity elsewhere in Europe.

US T2 Connectvity to CERN
US LHCnet is availble for this mission. It is known that at least several T2's are connected to USLHCNet at starlight and ManLan. How this connectivity figures into the computing model needs to be understood. However, it is important that this connectivity not impact the ingest of data by the US T1's.

US T2 Connectivity elsewhere in Europe
The CMS Computing model calls for the data sets under study to be distributed among its Tier-1 centers. A Tier-2 supporting analysis needs to download a data set from the unique Tier-1 which happens to host it. The Tier-1 center at FNAL, though large, only holds a proportional share of these data sets; it does not hold a comprehensive collection of all CMS data of interest to the US Tier-2s.

What does this mean for transatlantic networking? When estimating transatlantic T1 to T2 bandwidth , one can begin with the approximation that 1/2 of the T1 capacity is in Europe, 1/2 in North America, while 2/3 of the T2 capacity is in Europe, and 1/3 in the US.

In such a crude model one expects
1/6 of these transfers to be within North America.
1/6 of these transfers to be from Europe to North America.
1/3 of these transfers to be from North America to Europe.
1/3 of these transfer to occur with Europe.
While one might quibble about the merits of this simple model, it is clear that transatlanic transfers are of concern to the US project. As US LHCNet has no means to transit these transfers, the transiting has been done by the incumbent networks.

The current situation is as follows: eastbound transatlantic transfers originate from Fermilab. Because Fermilab is a DOE lab, ESNET has responsibility. Transfers are routed from Fermilab to Manhattan, then handed off to GEANT for transit to national research networks. Westbound transfers terminate in US universities; therefore Internet 2 has responsibility. GEANT routes transfers from the European T1s to Manhattan, and then hands off to I2. These transfers use the same 10 gig TA link as ESNet. Moreover, the link is shared with general purpose IP transfers.

In production, it is our estimate that demand will be of the order 7-20 Gigabits. With the exception of some of the Russian sites, all these transfers happen to flow across the Atlantic on a GEANT-furnished link from Manhattan (32 AoA) to Amsterdam. The link is general purpose and not dedicated to HEP.

In October 2006 that link was not saturated; however, commissioning of all anticipated T1 to T2 data flows had not begun. As this ramps up, we expect we will congest the link. In CSA06 we observed better than 1 Gbit/sec sustained for over an hour to DESY, and have achieved 0.6 Gbit/sec sustained for over an hour to several European T2's after modest, ad hoc tuning.

Fermilab has set a goal of being concerned if data service from Fermilab to a given European T2s achieves less than 80% of rate obtained from reasonably distant inter-European T1s. Fermilab networking experts will help commission rates from overseas T1's to the US T2's, to make sure the CMS computing model can successfully be employed.

Given this plan, the incumbent networks will see saturated capacity and have a bona-fide need for resources to support LHC Science. With the strictures of not spending additional HEP funds on transatlantic networking, it was agreed in October, in a LHC networking meeting at Fermilab, to investigate to offer some or all of the new USLHCNet Manhattan - Amsterdam link up for reciprocity. This link could then be exploited by I2/ESNET/GEANT, and be seen as a US contribution. We understand that negotiations are underway, and are not hopeless although there are issues. For example, the simplest implementation has the HEP-funded reciprocal link carrying general (i.e non-HEP) IP traffic. After considerations of how IP routing works, the LHC networking experts prefer the model of USLHCNET giving bandwidth to I2/GEANT/ESNET in this way, to an alternative of USHLCNET offering a competing routed IP service.

Lastly, the cost of networking is not just the cost of bandwidth. At this point we are concerned that we do not yet understand the support costs involved and who will bear them.

Question F7) Do the U.S. S&C programs have adequate links to the Worldwide LHC Computing Grid (WLCG) and to the Open Science Grid (OSG)?
Provide relevant organizational chart (if applicable)

Answer F7) There are no changes in the WLCG organization since last year. The U.S. is represented in all executive and oversight bodies of the WLCG, in particular in the Management Board (Ian Fisk for U.S. CMS S&C, Ruth Pordes for OSG), in the Oversight Board (John Huth and Irwin Gaines for the U.S.) and in the Collaboration Board (Institute PIs from each Tier-1 and Tier-2 center). Also, the U.S. facilities are represented in the Grid Deployment Board (Ruth Pordes and Bruce Gibbard). Also U.S. people participate in the different working groups of the WLCG etc.

The Open Science Grid has a new organizational structure in which U.S. CMS is well represented and linked into. Here is an organization chart showing the communication lines to U.S. CMS (shown as "OSG VOs").

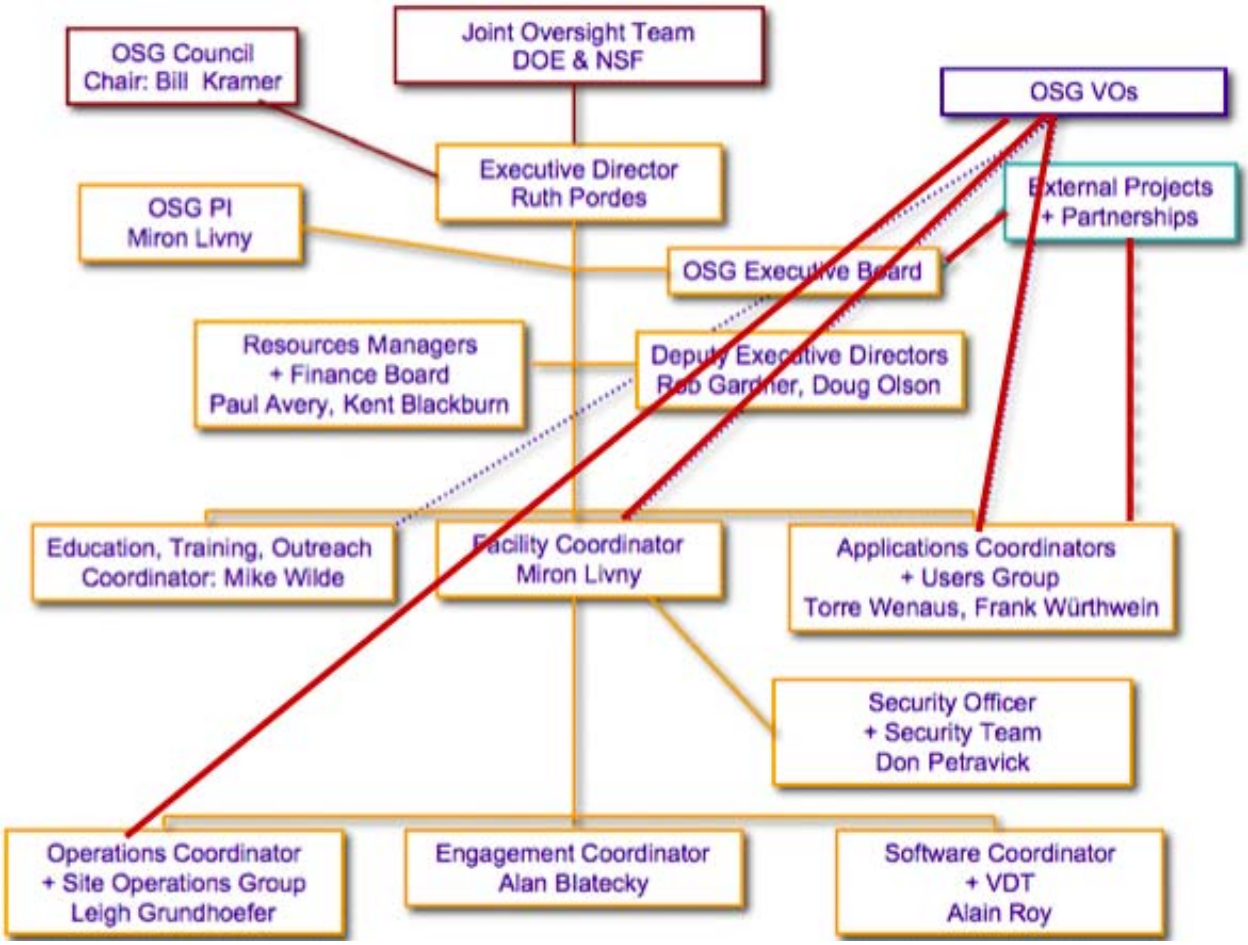

QUESTIONS TO CORE SOFTWARE AND ANALYSIS SUPPORT

Question S1) Are the current models for the support of data analysis well thought out, and is the support structure responsive to the needs of the U.S. community, both in the U.S. and at CERN? Are there adequate metrics to monitor progress in this area? Have all the required resources been identified by the collaborations? Will there be adequate support during all phases of the experiment?

Please provide specific metrics your collaboration is using to monitor this area. Also provide planned resources for the period FY06-FY11.

Answer S1) User support is now a L2 task in the CMS computing organization. It is lead by Kati Lassila-Perini. Support starts with detailed documentation, see https://twiki.cern.ch/twiki/bin/view/CMS/WorkBook. It is close to being completed. U.S. CMS has made major contributions to this effort by providing a technical writer at Fermilab. For in-person user support CMS also provides help desk assistance at the CERN (trough Kati) and at FNAL (through Patrick Gartung). If a physics user cannot find an answer this way, an expert will need to be contacted. For this purpose CMS uses a community based support model. A set of HyperNews fora has been setup for both computing and software technical questions and problems, that users utilize to contact experts, and that experts use to channel support.

Question S2) From a user's perspective, comment on the usability and readiness of the analysis software. (The collaborations should provide sufficient information to help the committee evaluate typical user experience with analysis tools.)
Any metrics in this area?

Answer S2) Unlike the MC event production grid tools, the user analysis tools are intended to be used by a broad community of users. CMS has developed the CMS Remote Analysis Batch system (CRAB), which is being used by a large number of physicists since Summer 2005. Much of the data analysis for the Physics Technical Design Reviews (PTDR) was done with analysis jobs submitted to the Grid through CRAB.

One important aspect is the reliability of running jobs in the Grid computing environment, and for the Grid the ability to run jobs for all CMS users at realistic scales.

The number of jobs submitted by users peaked at 100k jobs per month during the last month or the Physics TDR Volume 2. During CSA06, 50k jobs per day were submitted during the final week of the challenge (about 80% were submitted with automatic job load generators). The success rate of analysis jobs was approximately 90%. Over the last 12 months approximately 300 individuals have submitted a total of more than 1 million analysis processing requests through CRAB.

Tutorials on the use of CRAB have been conducted at CERN and FNAL and the feedback from the user community is generally positive on the system as a whole.

We are also monitoring the use of CMSSW by regular physicists a the UAF.

Question S3) Are the personnel requirements for the maintenance and operation phase of production software well understood, well justified and available? On what basis are commitments made to the international collaboration? Are these commitments realistic and consistent with U.S. interests?
Similar to 3rd bullet in Management Section.

Answer S3) The personnel planning for software and support is shown in Answer M3). However, the concrete assessment of these requirements are starting now by the CMS Offline and Computing coordinators, but at this point the maintenance operations needs are not yet well understood. The U.S. Software and Support coordinator (Liz Sexton-Kennedy) in CMS also has

the role of the Framework software coordinator, and she is putting together manpower estimates for framework M&O right now. Likewise the CMS Data and Workload Management coordinator is also the U.S. CMS Applications Services coordinator. Thus the U.S. is well informed and integrated into the process, and know about the existing and planned resources for the U.S. Software and Support areas.

Later this year, with a better understanding of the scope of this support needs the experiment will negotiate Memoranda of Agreement with the resource providers (currently started for the detector maintenance efforts), and commitments will be negotiated in this process.

Question S4) Is the role of Tier-3 centers well-defined and integrated into the S&C facilities plan? Are the plans for Tier-3 facilities sufficiently developed to guarantee capability for data analysis at interested U.S. institutions by November 2007?
Provide short statement on Tier-3 plans and coordination.

Answer S4) In the U.S. CMS model Tier-3 centers are all computing sites that are not Tier-1/2 centers (including departmental servers, work group computing for University groups etc). This infrastructure is very important for individual groups to prepare jobs that eventually get run at Tier-2 centers, or for keeping locally the final analysis samples for histogramming etc, and they are an integral part of the computing model. There are already several Tier-3 sites active, and more U.S. CMS university groups have announced they will setup Tier-3 computing sites soon.

The U.S. CMS project is not funded nor staffed to coordinate and manage all these sites (some 50 CMS University groups!), but instead we have a successful strategy to include these sites into the CMS computing environment through the OSG. The OSG was founded a couple of years ago and is now a funded project that is committed to help individual sites to become part of the OSG, and to provide the "facilities services" that these groups need in order to succeed on the Grid.

By deploying the Grid interfaces, each Tier-3 site will become part of the CMS system, and will be able to submit jobs, transfer data, access the calibration/alignment databases etc. Also, CMS has a system to push down software releases to sites, and interested Tier-3 sites are being included into this.

Beyond that, CMS and U.S. CMS are providing framework for exchange of knowledge, for participating in the challenges etc in the CMS Computing Integration Program and in the regular meetings that are open to interested Tier-3 site people. During the OSG meeting in March 2007 U.S. CMS is organizing a first Tier-3 workshop to disseminate information, give training to site administrators and get new sites started running CMS software.

Question S5) Has progress in core software relative to the milestones presented at the February 2006 comprehensive DOE/NSF review of the U.S. program been adequate? Are the forthcoming U.S. milestones on track and realistic? Is there any critical dependence on international milestones that could put U.S. deliverables at risk?
Provide list of milestones for FY06-FY08 and comment on status.

Answer S5)

FY06 Milestones:
Simulated Event Production for CSA06 (April 2006)

The CMSSW software based on the new framework was deliver on time to begin validation runs in the spring of 2006. Large scale simulation began in summer and 60M events were produced in time for the CSA06 data challenge. The new framework ran with a predictable memory footprint and low failure rate allowing the simulated event production to proceed at the required pace.

Participation in Service Challenge 4 (June 2006)

Service Challenge 4 served as a dress rehearsal for CSA06. Sites were commissioned and services were deployed.

Participating in the Computing Software and Analysis Challenge 2006 (Sep 2006)

The challenge began on Oct 1st 2006. The new CMSSW software performed well. The reconstruction software was delivered on time and was within expectations for memory usage and stability. Problems were discovered in specific elements in the reconstruction and calibration software which resulted in 4 new versions over the course of the 6 week challenge. The build and deployment infrastructure performed well allowing new releases to be created quickly.

Transition to new framework (October 2006)

After the release of the Physics TDR, CMS physics analysis was proactively moved over to CMSSW and old datasets were removed, while only a select subset. The milestone itself was related to deletion of all old data and was shifted to the end of 2006 to allow several graduate students to complete their thesis work based on the old data and software. Otherwise the CMS data analysis has successfully transitioned to the new framework, and during December of 2006 95% of the analysis activities were based on CMSSW. At the beginning of 2007, several samples of old framework data were identified as valuable for comparison and validation purposes.

Achieve 50% Facility (December 2006)

The facility milestone primarily relates to facility capacity and scaling of facility services, but the strong performance and reliability of the framework has allowed large scale validation of the computing systems.

FY07 Milestones:

- Participation in the start of global data taking (March 2007)
- Simulated data available for Physics Preparation activities (June 2007)
- Computing Facilities ready to receive data (July 2007)
- Participating in CSA07 (August 2007)
- Magnet Test and Cosmic Challenge 3 (October 2007)
- Pilot Run (December 2007)

Question S6) Is the U.S. core software portfolio sufficiently balanced to offer U.S. researchers a good chance to participate effectively in the initial science of the LHC?

Please list your priorities and metrics in this area.

Answer S6) The US has had a reasonably diversified core software portfolio with contributions to a wide variety of activities, while maintaining a critical mass of effort in areas where the US has particular expertise. The leveraging of core framework development expertise at FNAL and elsewhere in the US has ensured the successful deployment of the new software framework and improved the preparation of CMS as a whole for the start of data taking. The US has also played a leading role in the grid distributed production tool, data distribution, calibration information distribution, and data bookkeeping tools. All of these will help to facilitate the participation of remote communities more effectively. In areas where US-CMS has not led we have made strong contributions to ensure the needs of the US communities are being addressed. The U.S. collaboration with the Italian CRAB development team to adopt it to be able to submit to Open Science Grid sites is a good example for this successful approach.

CMS ANSWERS TO OVERNIGHT QUESTIONS

Responses of U.S. CMS to committee questions --- January 19, 2007

Software sub-group:

1. *What is the understanding between CMS and US CMS on expected response times to user questions?*

There is no separation between CMS and US CMS with regard to user support. The experiment as a whole has an integrated user support organization that receives effort from US CMS, as well as non-US collaborators. There is a very active community in CMS paying attention the HyperNews for answers. The global distribution of the CMS collaboration has a very positive impact on round the clock response to user questions. Bugs are being tracked in Savanna. The quality of the support is monitored by the software managers.

2. *Quantify job & software failure rates. How does the CMSSW contribute to these failure rates and how is this measured?*

During CSA06 the dominant failure modes were grid submission problems and data access problems. The job failure rate was approximately 10%. The failures due to software problems were less than 1 per 1 million events. There was one road-search tracking algorithm which had memory consumption problems for pathological events and was thus switched off.

3. *How fast will a CMS physicist be able to produce a Z->ee mass plot from 1 Million AOD events? Please specify wall-clock time which reflects software performance.*

Using the EDAnalyzer from our demonstration in the review, the time to analyze 1,000,000 AOD events to produce a Z->ee mass plot is about 10 minutes (wall clock time) reading files from a local file system and about 40 minutes reading files from the dCache mass storage system. It has to be noted that no possible optimizations for the dCache system have been applied which can increase the performance by about a factor of 10 (based on CSA06 investigations).

4. *You stated in the parallel session that data management/placement in response to analysis patterns will be a major challenge at the beginning. How will you determine the optimal data locations and how will you then execute the necessary policies in a timely fashion? At what point in time do you plan to demonstrate this?*

In the parallel session it was stated that data management issues in general were likely to be those most critical at the beginning of data taking *relative* to core framework/Software issues. This has been seen in all of the currently running experiments. In addition it was stated that the analysis patterns were likely to be particularly challenging given their dynamic nature.

The data handling in the HLT to Tier-0 to Tier-1 data and work flow for real data is clearly the first challenge. CSA06 only tested Tier-0 activities and data transfers to Tier-1 sites and onward, without a "full chain" exercise including the HLT. The HLT—Tier-0 exercise will begin this spring with a "full chain" functional exercise using the global DAQ.

Given the CMS computing model, the most dynamic analysis patterns are concentrated in the Tier-2 sites (and the LPC-CAF, which however is special as it has full access to all data samples stored at Fermilab) and (unlike in previous experiments) Tier-2 sites will likely be important very early on. In terms of the S&C deliverables, the main missing ingredient is the full "site-view" of all data access, production and transfer activities at the site and direct control by the site of the data available at the site. CMS has all of the elements of the necessary monitoring, but in disparate monitoring systems. The administrative control of transfers to all sites currently takes place through the central operations team. Work is underway to combine these things to provide a "full view" for a "data manager" at each Tier-2 site, as well as direct control of transfers to the site. This will be deployed this spring and will be exercised during the large MC production, follow-on physics studies and CSA07.

Management sub-group:

1 *Describe the ways in which additional physicists can contribute to US CMS S&C. Can data operations be used to satisfy the experiment detector service requirements? How many physicists could be productively integrated and on what time scale? What sort of recruiting would be effective?*

CMS has recently restructured its organization to focus more on operations. Physicists could easily contribute to the new computing tasks in the areas of Data Operations and User Support. In fact, it is expected that physicists will supply much of the effort for data operations. The tasks are described in more detail below.

The Data Operations Group will be responsible for operational aspects of real and MC data processing and executing production workflows. Physicists in this group are charged to ensure that detector calibration and alignment corrections are applied at prompt processing and data are properly distributed. The data operations tasks include:

* Prompt calibration, alignment at Tier-0/CAF, processing of environmental data, first pass reconstruction at Tier-0, operations for obtaining refined calibration and alignment at Tier-0, CAF, Tier-1s and Tier-2s.

* Distributed production and storage of MC events; coordinate MC queues, bookkeeping of MC data

* Distribution of first-pass FEVT and AOD to Tier-1s; data management and bookkeeping

* Organized large-scale analysis activities; reprocessing at Tier-1s, RECO and AOD distribution

* Application of Data Quality Monitoring algorithms and procedures to the data; Luminosity determination; organizing offline shifts for DQM; establishing good-run lists.

CMS expects to have two L2 managers and five L3 managers in this area. It is likely that these positions will be filled by physicists. One L2 manager (a US-CMS physicist) and one L3 manager have been identified. The other leaders are expected to be identified over the next 6 months.

The data operations group will grow over the coming year in anticipation of the pilot run. The Run II experiments estimate that the collaborations provide 20-25 FTE's of scientific staff in

offline computing and data processing areas. This is a reasonable estimate for CMS at the beginning of data taking if one excludes development work.

MC production for CSA06 was done by 5 physicists and 1 software professional working nearly full time in 4 teams. It is expected that this level of effort will be needed until the production system is fully automated when it will become a production facilities operations task.

Physicists will be needed for integration and commissioning tasks. In 2006, 3-4 physicists worked on grid integration tasks for CMS. The level of effort in 2007 could expand to 5-10 physicist FTEs to take on the full the integration of databases and calibration and alignment tasks.

There are also user support tasks that require physicists. One expects that about 3-5 FTE could be integrated to help with these tasks over the next year particularly in the tutorials and in providing documentation for CMS software. The US-CMS tutorials in 2006 conducted primarily by physicist effort.

The US has the LPC as a platform to make physicists effective in contributing to these tasks.

We expect to be able to recruit the required physicists for these tasks from inside the collaboration over the coming 6-9 months. It should be possible to integrate them into the effort if enough attention is paid to documentation. Tutorials on the software and computing are already run at a regular basis.

CMS has not yet come to an agreement on an operations model and has not finalized the definition of service requirements. We believe that most compelling incentive for recruiting physicist to data operations tasks would be if CMS would credit Software and Computing activities in the same way that detector operations shifts for the authorship requirements. An addition incentive for US physicist would be if offline and data quality monitoring shifts could be taken at the FNAL ROC.

2 *Please provide the "lessons learned" document from CS06, and please summarize the major elements of a program of work to address the needs for stable operations, item by item and the FTE allocation needed. Please list the missing components identified in CS06. How are these missing components being addressed?*

Lessons learned is the final chapter of the CSA06 wrap up document attached in the review supporting documentation. The chapter has been carved off as a separate document linked to

http://www.uscms.org/SoftwareComputing/ProjectManagement/DOE_NSFReviews/2007-01/docs/lessons.pdf

The specific technical elements with learned lessons include:

* Improving the site control for Tier-2 data management. Provide sites with tools to control subscriptions, delete data, and recover from data loss. This is the primary development effort for the PhEDEx team during the first quarter of 2007 and estimated at 2FTE of effort for initial development followed by 0.5FTE of continued support.

* CMS needs to complete the implementation of the organized processing request system. This includes simulation requests, event selection and re-reconstruction requests. The request

system is in the work plan for work flow development and is estimated to require 1FTE of effort for 6 months.

* The DBS was discovered to have scaling limitations while performing the merging of small files into larger files. This is being addressed in DBS2, which will have an initial release. The total DBS effort over 2007 is 5FTE, which includes integration into offline conditions databases.

* A variety of issues were found with the ability to transfer large files through slow networks using the file transfer service. Some of this was related to site configuration, time-out parameters in FTS. Improvements in the Tier-1 to Tier-2 transfer quality is an effort of the CMS Computing Commissioning activity. The site effort from the integration program is estimated at 10FTE.

The general lesson learned from CSA06 is that scale and commissioning testing is very valuable and the operations model needs work to arrive at a sustainable load. The first operations development activity is being addressed with the transition to periodic global data taking periods of increasing frequency starting in March of 2007. The scale will be stress tested with the next Computing Software and Analysis Challenge (CSA07)

The largest missing components that were outside the goal functionality of CSA06 were the components required to integrate the online and DAQ and the conditions data. The integration with conditions data will be largely addressed in the DBS effort. The integration of on-line and off-line computing is an international CMS effort.

3 *What are the explicit US CMS S&C deliverables for 2007?*

Regarding USCMS facilities, the US, as part of the WLCG, has signed an MoU with CERN which commits to deliver the Tier-1 and Tier-2 facilities facilities at the pledged scales, and to provide the required grid services. Also, the USCMS collaboration expects us to deliver the corresponding analysis facilities (Tier-2s and LPC-CAF) at the baselined scale. These deliverables are covered by the WBS areas 1.1 (Facilities), 1.2 (Tier-2s), 1.3 (Grid Services) and 1.5 Distributed Computing Tools.

USCMS also delivers contributions to the CMS-wide core software and systems (Data management, workload management etc), which are covered by WBS areas 1.4 (Application Services), 1.6 (Software and Support) and 2.0 (CAS). The responsibility for delivering these systems is with the CMS offline group, in particular the CMS Framework and the CMS Data & Workflow Management tasks, which are led by the corresponding US area coordinators (WBS 1.4, P.Elmer and WBS 1.6, L.Sexton-Kennedy). US deliverables (in which US people serve as project leads and that are done predominantly with US manpower) are: Storage Manager, DBS, Frontier.

The US contribution to the CMS Framework are delivered according to the Software release plan (WBS 1.6). The US is also contributing to Phedex, ProdRequest, ProdManager, ProdAgent (WBS 1.4, 2.0).

In addition there are deliverables related to support (including engineering support for US detector software), integration, and operations. All these efforts are at the “baselined” level of US FTEs.

4 *Are there mechanisms and plans for getting long term global CMS operational support for US deliverables? What are the mechanisms for transitioning from expert or developer support to operations for deliverables such as the framework/edm?*

No explicit mechanisms in place at this point. This poses a problem in particular because most of the detector related commissioning and operations (including computing operations, but not including offline) are to be covered in Memoranda of Agreement while software and offline is not included (and the scope of Computing Operations is not really defined yet).

CMS will probably need to agree on a process to include the software operations and maintenance efforts so they can be accounted for in the overall needs of the detector for maintenance and operations.

5 *Provide the compilation of the milestones for 2007, in particular, targeted to achieving stable and sustainable production operations and for the transition to data collection.*

The milestones intended to achieve stable and sustainable operations include

∗ The start of global data taking (March of 2007)

∗ CSA07 in September (September 2007)

∗ Beginning of Magnet Test and Cosmic Challenge 3 (MTCC3) (October 2007)

The start of global data taking has integration aspects to connect the detector, the on-line, and the off-line computing. From the operational side it is intended to function as a periodic opportunity to test the end-to-end system. The milestone includes DAQ and Tier-0 elements as well as integration with data management components.

CSA07 is an opportunity to test improvements in the operations model as well as demonstrate the CMS computing system in a higher scale test. The CSA07 test will run sustained for a 30-45 day period.

MTCC3 is an opportunity for the experiment to take data through the full chain with cosmics. CMS will be in a continuos operations mode at a reduced event size.

6 *What are the formal processes by which the software deliverables are evaluated by the CMS collaboration?*

The software deliverables are evaluated in a schedule sense against the global CMS schedule and in an annual internal review. The process by which the global CMS schedule is formed is directly shaped by the participation of USCMS managers as L2 managers in the CMS computing and offline organization.

If the question is intended to mean quality assurance and quality control, there are three aspects:

a) A set of standard, basic MC samples used to verify that each new release is functional (no crashes) for event generation, simulation, reconstruction, etc. These are run both during pre-release testing (at and immediately after a release is announced (with somewhat larger statistics)

b) A set of "physics validation" samples and a testing suite (SVSuite and "Validation" test packages) doing statistical tests with histograms to evaluate physics performance from release to release. A "Physics Validation" task force has been working since Nov2006 validating CMSSW against the old SW/EDM using these tools.

c) We currently do not systematically measure and document CPU/memory performance from release to release, although this is done in an ad-hoc fashion. The Framework infrastructure exists to measure this and a dedicated "Code Performance Task Force" has been recently formed (Jan2007) in CMS whose objectives include adding the functionality to properly document this in an automated fashion.

7 *What is the plan for continuing to provide for analysis in the event of any Tier 1 unavailability?*

In the CMS model there are two copies of all real event data, one tape archive copy at CERN and one active copy distributed across the Tier-1's. For MC events, there will be typically only one copy of individual samples at Tier-1 center and CERN, except for optimizations.

There are not enough resources at CERN to serve the CERN archive copy of real even data to Tier-2s for analysis. Since each Tier-1 is responsible for a fraction of the FEVT, this fraction of the FEVT data will then be unavailable while the Tier-1 is down. Each Tier-1 has a copy of the full AOD data, so physicists will be able to continue their AOD analysis by getting their AOD sub-sample from another Tier-1.

In the event of an extended outage of a Tier-1 the experiment would fail-over and cope by making a one-time copy to another Tier-1 to serve it. The model foresees good connectivity between all Tier-2s to all Tier-1s no matter where the copy is placed, so the experiment should see, to first order, normal quality service. The PhEDEx architecture fully supports this already today. It has the ability to automatically serve the data from the most available location if the data exists in more than one location.

8 *What are the changes in the FNAL organization with respect to US CMS and what was the charge to re-organize the LPC?*

Fermilab has created the "CMS center" which puts the Fermilab CMS efforts onto the lab org chart in the Research Sector at the same level as the division. The role of this center is to create a focus for all the CMS efforts at Fermilab. Pier Oddone has appointed Lothar Bauerdick as Director of the CMS Center.

The charge of the LPC Governance Committee is not to re-organize the LPC, but to provide advice to the host lab about a governance framework for the LPC, including proposing a process to select the next LPC leads. To quote from the charge to the LPC governance committee:

"The Fermilab Director would like to request that the LPC Governance Committee provide advice on the appropriate structure and leadership for the LPC. The LPC Governance Committee, convened by Lothar Bauerdick, comprises:

Lothar Bauerdick/Fermilab -- Chair
Joel Butler/Fermilab -- U.S. CMS Research Program Manager
Maxwell Chertok/UC Davis -- LPC Advisory Board Chair
Sarah Eno/U.Maryland -- current LPC co-lead

Dan Green/Fermilab -- Member of the CMS Executive Board
Joe Incandela/UC Santa Barbara -- CMS Physics Coordinator Deputy
Boaz Klima/Fermilab -- LPC working group leader, Fermilab contact
Harvey Newman/Caltech -- as chair of the US CMS Collaboration Board
Paris Sphicas/CERN and U.Athens -- CMS Physics Coordinator
Avi Yagil/Fermilab -- current LPC co-lead

To quote from an email by Harvey Newman, chair of the US CMS Collaboration Board, to the US CMS collaboration:

The advice of the Committee is meant "to provide transparency, ensure buy-in and ownership of all collaborating institutions, enable full participation of US University groups in the LPC", and above all to help ensure that the LPC will be able to fulfill its primary goal of "helping U.S. CMS scientists at Universities and labs to contribute to CMS and to participate efficiently in CMS physics analysis."

9 *Are the International CMS management structures making decisions and schedules in a timely manner such that you can plan the US CMS S&C commissioning? Are the major issues for the experiment readiness being addressed?*

Yes, although the transition to the new management happened only recently and thus the planning process has not yet concluded. And, yes, the major issues for experiment readyness are being addressed with a series of challenges and tests as presented in the plenary and parallel sessions.

The major milestones for computing and offline activities for the commissioning phase are known, which allowed US-CMS to do coarse planning and set priorities for development, integration and commissioning activities. The change of management combined with the change of thrust from integration into operations has delayed some of the detailed planning, but this has not negatively impacted CMS planning particularly.

Grid & Facilities sub-group:

1. *We were shown plots of efficiency/failure rates of PhEDex data transfers over Oct. 2006. (Commendable to show the plots). Integrated over the month, do the failures only impact bandwidth utilization or does it indicate the amount of data intended to be sent that was never transfered? What fraction of the total data was not transferred during Oct?*

No. All intended data was transferred.

The denominator is # of attempts to transfer files. All files that were attempted to be transferred were actually transferred. PhEDEx architecture includes a “download agent” that keeps track of errors, and automatically retries the transfers.

Fundamentally, the important metric is transfer rate, not failure rate, as failure rate only reduces effective transfer rate. We pay attention to the failure rate as an indication of health of both end points.

2. *What would be the impact if the AOD size doubled, and what would you do to mitigate the problem?*

It would effectively double the bandwidth required for data transfer from T1 to T2 as fewer events can be cached at the T2. The T1 has adequate disk and network bandwidth to deal with a doubling of the AOD, especially as the AOD volume total is small compared to the total data volume at the T1. The size of the AOD in the end is a parameter in the model of "usefulness" of the AOD (so people don't have to go to the FEVT) vs. space and throughput required at the Tier-1 and Tier-2 centers.

3. *Do you know the event rate for simple analysis of AOD (KHz/processor)? Is it satisfactory for physicists doing data analysis?*

This is an issue that CMS will need to pay close attention to within the upcoming year as the analysis algorithms for start of data taking take a more complete form. This has not been a priority to measure thus far, as we are not sure that numbers measured today are representative of the start of data taking. The CMS framework supports making these measurements today.

One data point was given in the answer to software question 2, which corresponds to 1.7kHz.

4. *What validation process is planned to verify analysis capabilities at T2s in advance of data taking? Including data access from the non-U.S. centers by the U.S. Tier2s.*

(a) In CSA06, we have established the ability to sink data from any T1 to any US CMS T2 . We are presently starting to exercise this with a goal of understanding the bottlenecks that keep us from reaching 1Gbps for any link between a US CMS T2 and a global CMS T1. We expect this to be a significant integration problem with a variety of different problems to be identified and overcome.

(b) We have established CMS data analysis capability at all US CMS T2 centers. This was demonstrated in CSA06 where the US CMS sites dominated the user analysis activities during the challenge. We have an ongoing program of analysis exercises and challenges in 2007 to scale up our capabilities. This includes the focus on Monte Carlo based physics analysis preparing global CMS for the first 1fb–1 of data, as well as the work leading up to, and including CSA07.