DOE/NSF U.S. LHC PROGRAM OFFICE
JUNE 6, 2005

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

BNL, 1-4 MARCH 2005

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

TABLE OF CONTENTS

EXECUTIVE SUMMARY

ATLAS and CMS will be general-purpose detectors at the LHC, a high-luminosity proton-on-proton collider located at CERN in Geneva, Switzerland. The collider will operate at 7 TeV per beam and is expected to turn on in mid-2007. The U.S., recognizing the great physics opportunities at the energy frontier, has invested substantially in the construction of the LHC accelerator and its associated detectors. Software and computing will play a central role in handling the unprecedented data rates from the LHC. This activity is especially important for U.S. scientists who will be analyzing data collected at a facility overseas.

In order to fully exploit the U.S. investment, ATLAS and CMS Software and Computing (S&C) programs were developed as part of the U.S. LHC Research Program, with the aim of empowering U.S. physicists with the means and tools needed to access and analyze LHC data in an effective manner. The scope of the S&C projects include aspects such as the provision of U.S. computing facilities and of experiment-specific core software, a commitment of facilities and personnel to national and international grid projects, and responsibility for ensuring that U.S. researchers have access to LHC data and appropriate analysis tools.

A DOE/NSF review of the ATLAS and CMS S&C programs was held on March 1-4, 2005 at Brookhaven National Laboratory. The review covered the general areas of Management, Facilities and Grids and Core Software. The expert reviewers and the collaborations were instructed to address the progress in S&C to-date, and to assess the needs and plans for the period FY2007-FY2009. The reviewers asked questions during the presentations, and provided both oral and written comments to the collaborations. The panel members also discussed their observations and recommendations in executive sessions in the presence of representatives from the DOE and NSF funding agencies, and then presented their preliminary findings to the collaborations in a close-out session. This report captures the final observations and recommendations of the review committee on the S&C issues pertaining to the U.S. ATLAS and U.S. CMS research programs.

The reviewers made several observations that were common to both experiments. They noted that both experiments have sound S&C management structures that meet the needs of the projects. After delays resulting from past funding uncertainties, the experiments are now ramping up the deployment of their Tier-1 and Tier-2 facilities. Following the recent completion of their computing models, and the subsequent LHCC review, both U.S. collaborations made significant revisions to their estimates of required resources, and have now incorporated the increased needs into their planning.

The review committee expressed concern about the dependence of the success of U.S. Software and Computing on deliverables from external grid projects, feeling that both collaborations should work together to mitigate the risk of this dependence because of the limited duration of some of these projects.

The reviewers commended both experiments for their extensive leadership roles in their respective international collaborations. However, they expressed some concern that U.S. CMS might be over-extending its responsibilities within international CMS.

The reviewers also spoke of concerns specific to each collaboration, an important one being that U.S. CMS might be underestimating the personnel and time needed to complete the re-engineering of the CMS core-software framework and event-data model. This new task was taken over by U.S. CMS from the international collaboration after the recent Data Challenge (DC) indicated serious deficiencies in performance. The reviewers recommended that the collaboration should move quickly to develop plans for this new undertaking and identify any missing resources required to complete *all* U.S. CMS deliverables.

U.S. ATLAS and U.S. CMS presented very different approaches to the way they intend to support physics analysis. U.S. CMS is deploying its LHC Physics Center (LPC) on the 11th floor of Wilson Hall at Fermilab. The LPC provides a central gathering point for U.S. CMS collaborators, where physics meetings, training, and (possibly) remote shift operations can take place. The LPC is particularly advantageous for CMS scientists participating in Run II of the Tevatron. In general, the committee endorsed this approach for U.S. CMS.

U.S. ATLAS is considering a more distributed approach to analysis support for physicists. The prevailing model involves establishment of "virtual corridors" with provision of sufficient expertise to enable U.S. ATLAS scientists to carry out their analyses. U.S. ATLAS considers CERN to be one natural place for such a center. The reviewers, noting that what was presented appeared to comprise a significant departure from previous successful experience in data analysis at large experiments, were concerned with the lack of specificity in the ATLAS plans, and that the plans were "not sufficiently well-defined", thereby risking initial U.S. success in physics output. The reviewers feared that the model might also lead to a fragmentation of effort, adversely affecting the critical mass in personnel needed for a coherent program. The reviewers recommended that U.S. ATLAS clarify its approach by the time of the next review in August 2005.

In preparation for the review, the collaborations were asked to address three funding scenarios for the period FY2007-FY2009: (i) "minimal", where project scope is kept at a minimum level in support of U.S. physicists, (ii) "reduced", assuming a 10% reduction in funding, and (iii) "leadership", with a modest increase in funding to bring U.S. participation to a leadership level.[1] The expert reviewers felt that the funding scenarios were presented in a way that was difficult to interpret, and consequently asked the collaborations to clarify in writing their plans and priorities for these three scenarios.

[1] After the charge was issued, the collaborations were told to assume a 10% funding increase for a leadership scenario.

Under the leadership scenario, both collaborations would enhance the capabilities of their facilities by expanding the role of their Tier-1 centers and by adding an additional Tier-2 center. They would also expand support for physicists, and mitigate additional risk by protecting their respective management reserves.

In the minimal scenario, ATLAS reported that the capabilities of the Tier-1 center would be reduced to a level consistent with international LHC obligations. U.S. CMS would also be able to maintain its commitments to the international collaboration, but the rest of the program would then be supported at only the "bare bones" level, and with higher risk because of a diminished management reserve.

In the case of a 10% reduction in funding, the Tier-1 and Tier-2 facilities for both collaborations would be severely reduced in size and in ability to handle large data sets, thereby greatly reducing the analysis efficiency for U.S. physicists. Both collaborations would have to reduce support for physicists, with U.S. ATLAS redirecting some to perform support functions. In the case of U.S. ATLAS, the management reserve would be fully depleted just to provide basic functionality. Commenting on the impact on ATLAS, the committee observed that a 10% reduction "severely jeopardizes the ability of U.S. physicists to participate in early data analysis".

The following is a summary of the recommendations provided by the review committee.

U.S. ATLAS Management:

- The physics analysis model of U.S. ATLAS should be defined more clearly and in greater detail. Given that the physics centers must be established before turn-on, U.S. ATLAS should decide soon where to locate their centers and how to secure their funding.

- The committee recommends that U.S. ATLAS use the change/control procedure to set priorities and plan implementation of software features.

- The committee recommends that U.S. ATLAS take measures to ensure that facilities for U.S. physicists are available at CERN.

- The committee recommends that U.S. CMS and U.S. ATLAS work together to generate compelling proposals for funding common projects in grid operations and networking.

U.S. ATLAS Facilities and Grids:

- Personnel at the U.S. ATLAS Tier-1 center must continue to engage the attention of the DOE/ESnet program in order to ensure that BNL has sufficient bandwidth to support both U.S. ATLAS and its other operational commitments.

- The committee encourages U.S. ATLAS and U.S. CMS to work together in seeking a solution to the expected end of funding of grid-middleware projects currently supported by external U.S. Grid programs.

- U.S. CMS and U.S. ATLAS must continue to press for support of the Open Science Grid (OSG), and especially the middleware upon which the U.S. LHC grid is based. In parallel, the collaborations must continue to engage with the LCG and its EU partners to define stable interfaces that will ensure interoperability in the final software environment.

- U.S. ATLAS should strive to demonstrate as soon as possible that a 90% success rate for production jobs is acceptable, and that this rate can be sustained for full-time operations following the turn on of the LHC.

- The U.S. ATLAS facilities and grid teams must maintain participation in the effort on cybersecurity for the DOE/NSF grid. There is also a need for continued assessment of risk and pursuit of new measures and practices. The FTE requirements needed to support the cybersecurity policies promulgated by the funding agencies, in addition to any support provided by the host institutions, must be identified and fully assured.

U.S. ATLAS Core Software:

- U.S. ATLAS should press the international collaboration to keep the FTEs associated with infrastructure support consolidated into fewer individuals who could dedicate a larger fraction of their effort to their projects.

- Efforts to solicit personnel from universities to work on software issues should be expanded as much as possible.

- A more detailed prioritization of needs in personnel should be made for the out-years, and considered in light of the different funding scenarios given in the charge.

U.S. CMS Management:

- The committee recommends that U.S. CMS remain vigilant to any overextension of its responsibilities within CMS.

- The committee urges that U.S. CMS management ensure that the newly acquired responsibilities on framework software do not put current U.S. CMS deliverables at risk**.**

- The committee recommends that U.S. CMS conduct an internal review of the trade-offs of various funding scenarios. U.S. CMS should also engage the

international collaboration to further clarify the requirements and priorities for all Tier-1 centers.

- The committee recommends that U.S. CMS and U.S. ATLAS work together to generate compelling proposals for funding common projects in grid operations and networking.

- The committee encourages U.S. CMS to continue taking its alert and adaptive approach to project needs.

U.S. CMS Facilities and Grids:

- The committee recommends that U.S. CMS continue to pursue solutions to S&C issues that they consider critical, even when they lie beyond the direct control of U.S. CMS.

- Fermilab and U.S. CMS must continue their effort to find a long-term solution to the acquisition of a wide-area, transmission network of high bandwidth and high quality. The current arrangement using the FNAL-provided OC-192 link to Starlight (and then to CERN via LHCnet) is meeting the data-challenge needs, but is not adequate for the future. This must be resolved before the start of data taking in 2007.

- U.S. CMS must continue to work with International CMS in order to assure that the global collection of CMS Tier-1 centers has sufficient capacity and capability to meet CMS needs.

- The committee encourages U.S. ATLAS and U.S. CMS to work together in seeking a solution to the expected end of funding of grid-middleware projects currently supported by external U.S. Grid programs.

- U.S. CMS and U.S. ATLAS must continue to press for support of the Open Science Grid (OSG), and especially the middleware upon which the U.S. LHC grid is based. In parallel, the collaborations must continue to engage with the LCG and its EU partners to define stable interfaces that will ensure interoperability in the final software environment.

U.S. CMS Core Software:

- The committee recommends that the scope and work plan for the re-engineering of the CMS framework be clarified in detail. Comparisons should be drawn with contemporary experiments of similar complexity. The integration of applications, reconstruction and physics analysis software within the new framework has to be specified. Any shortfalls in personnel required for the re-engineering effort must be identified promptly. Should there be a shortfall, it is essential that U.S. CMS

management intervene to secure additional support from International CMS to ensure the success of the re-engineering effort.

1 INTRODUCTION

Software and Computing will play key roles as science enablers for U.S. physicists during the LHC era. The unprecedented data rates and the distributed nature of physics collaborations at the LHC pose logistic challenges of data access that can only be tackled with an efficient and ubiquitous software and computing infrastructure. This is particularly important for U.S. physicists, who must overcome the disadvantages of geographic separation from their experiments through good access to the data.

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 projects within the context of their Research Program, as defined in their Program Management Plans. The agencies, in order to monitor progress and establish support guidelines for the Research Program, regularly review the work of the collaborations. To this end, a joint DOE/NSF review of the U.S. LHC Software and Computing was held on March 1-4, 2005 at Brookhaven National Laboratory. This yearly comprehensive review was carried out at the request of the DOE/NSF Joint Oversight Group (JOG). It was organized with the aim of evaluating the progress and plans of the Software and Computing efforts of the U.S. CMS and U.S. ATLAS collaborations.

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 BNL. The panel was asked to evaluate progress and plans in three main areas: Management, Facilities and Grids, and Core Software. The collaborations were asked to provide enough information to facilitate the panel's evaluation. The members of the review panel were chosen for their expertise in the area of software and computing in large-scale experiments. The expert reviewers were chosen from outside the LHC community and came from U.S. universities, national laboratories, and Europe. The list of participating panel members is included in Appendix B.

The agenda of the review is included in Appendix C. The presentations given by the collaborations, a detailed agenda, and additional background material can be found at http://www.usatlas.bnl.gov/atlas_psc/review/rev_Mar_05/.

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 plenary presentations by the collaborations, on background material provided prior to the review, on detailed discussions during the parallel sessions, and on the answers to more in-depth questions posed by the panel members. The committee had extensive discussions during executive sessions in the presence of agency representatives. A draft of the committee's initial evaluations was presented to the collaborations at a close-out presentation.

Brookhaven National Laboratory, as host for this meeting, provided invaluable logistic support, which was essential to the success of the review.

2 MANAGEMENT

2.1 U.S. ATLAS

The U.S. ATLAS Management effort was summarized in presentations by J. Shank, M. Tuts, and J. Huth. A breakout session provided an opportunity for discussions with the ATLAS team.

2.1.1 FINDINGS

The organizational structure of the U.S. ATLAS Software and Computing Program was instituted in 2003, however, there have been some personnel changes in 2004, with Mike Tuts replacing William Willis as the manager of the U.S. ATLAS Research Program. R. Popescu has joined Bruce Gibbard as head of Facilities (WBS 2.3).

U.S. ATLAS remains well represented within the International ATLAS computing organization providing (11 of 20) project and subproject leaders. This organization has been quite stable.

U.S. ATLAS is in the process of defining the physics-analysis model, with a key requirement being the desire to leverage the distributed nature of the collaboration, to have a significant presence at CERN, and to let U.S. analysis support hubs emerge organically. They are exploring the concept of “virtual corridors”, utilizing collaborative tools to connect users and local experts, leading to distributed user training and user support. The resources to support these centers are envisioned to come from ”core” program funds. The analysis hubs are not expected to map directly onto Tier-1 and Tier-2 centers nor onto participating National Laboratories. However, the committee was told that one of the analysis centers will definitely be located at CERN. The distributed model is in use now, and U.S. ATLAS indicated that they have strong participation in physics working groups within international ATLAS.

The ATLAS computing model was revisited in 2004 in light of better estimates for the processing time, event sizes, and more realistic consideration of calibration and alignment needs. These revisions lead to a doubling of the total computing resource estimate for ATLAS. The requirements by 2008 on the BNL U.S. ATLAS Tier-1, including U.S. ATLAS controlled analysis resources, is 4.6 Petabyte of disk, 4 Petabyte of tape storage, and 8 MSI2K of CPU. By use of dCache to manage disks on worker nodes, the cost of the disk storage will become significantly smaller. Even so, meeting these requirements will require additional resources and funds in FY06 and beyond. These funds are anticipated to come from the management reserve. In FY08, the amount of funds required for this activity comprises almost the entire management reserve. Furthermore, there remains significant uncertainty in the amount of information that will be kept for the initial storage of events from the online system. The lack of availability of computing resources resulting from budgetary constraints could force a reduction of the trigger-accept rate.

There are significant calls on the management reserve to meet high-priority software-development (6 FTEs) and production-support (2 FTEs) priorities, as well as $600K to deploy hardware at the Tier-1 and Tier-2 facilities to support increased production and analysis needs, as well as to prepare for Data Challenge 3. In addition, management reserve will be used to compensate for missing functionality in grid middleware as well as to cover the end of PPDG funding. Finally, the management reserve is expected to cover the need for collaborative tools, which will form an essential component of the success of their distributed analysis model.

The U.S. ATLAS effort is continuing to provide roughly 40% of the FTEs working on Core Software issues for the ATLAS collaboration. The Core Software effort is about 35% understaffed, which is leading to a slippage of milestones for roughly 25% of the deliverables. The committee was told that U.S. ATLAS supplies 8.7 FTEs for infrastructure, spread among 22 individuals.

The U.S. ATLAS project managers are taking responsibility for overseeing any changes in scope or schedule, as well as adjusting the WBS to adapt to global ATLAS priorities. The combined test beam (CTB) and Data Challenge 2 proved to be a case in which conflicting priorities had to be resolved, in particular because international ATLAS management had expectations for the software and computing project that initially were not clearly explained or understood. The ramifications of the CTB on computing systems required new deliverables specifically to support the CTB and the reordering of established deliverables. Meeting the rigid deadlines of CTB required delaying some of the DC2 milestones, in particular demonstrating the analysis of reconstructed data.

For changes in control and scope within the core-software effort, the management team currently assesses schedules and requests and reassigns personnel on a case by case basis, without going through a formal procedure if this involved less than ½ FTE.

The Leadership scenario would enable full funding of all high-priority software-development tasks such that the detector software and infrastructure would receive the necessary attention to avoid conflicts with the analysis-support tasks. In this scenario, calls on management reserve to support the facilities would be reduced, enabling the full set of Tier-2 centers to be funded and the Tier-1 center to ramp up in a timely manner to enable full stress testing of all of components.

The committee was informed that, in the reduced (-10%) scenario, no hires for critical computing tasks could be made, and at least one position for core software would end. The Tier-1 center would be understaffed relative to current estimates, which could lead to a reduction of operating efficiency of computing equipment. One of the Tier-2 centers would be completely cut, reducing the computing resources for analysis and Monte Carlo generation.

2.1.2 OBSERVATIONS

The US-ATLAS Software and Computing management organization appears to be working well and meeting the needs of the project.

The effectiveness of a distributed physics analysis model is difficult to evaluate, as it comprises a significant departure from models used by most previous large experiments. Given that such a model is untested at this scale, it introduces a high risk into the analysis operation that potentially leaves U.S. ATLAS physicists unable to participate in early data analysis in an effective and timely fashion. The ideas presented on virtual corridors still appear preliminary, and not as yet fully developed or tested. In addition, CERN has not been supportive of providing properly equipped facilities for remote participation. The U.S. ATLAS analysis model is also vulnerable to ”core” funding cuts that may leave U.S. ATLAS institutions competing directly with each other for available funds.

It is not likely that 8.7 FTEs spread over 22 individuals can be fully effective as contributors to ATLAS infrastructure. U.S. ATLAS is not in a position to contribute more effort or to take on more scope, and the ATLAS computing and software managers have been diligent in bringing these concerns to international ATLAS. However, some personnel provided to global ATLAS are not trained computing professionals. The effort and lead time to train individuals can be an appreciable burden on management and developers. If effort is to be added to the team, it should be done soon so that new personnel have adequate time for integration into the operation.

The experience gained by management in 2004 with the Combined Test Beam and the Data Challenge will likely prove to be valuable. Understanding how to make trade-offs, assessing firmness of deadlines, and managing expectations will be key skills for setting priorities as ATLAS approaches operations. At that time, priority conflicts are likely to be the norm rather than the exception.

As ATLAS approaches turn-on, more requests for changes are likely, since more physicists will be getting involved and using the software. This will produce extensive pressure on the developers to respond to requests and add features, which could cause the developers to lose focus. Using established change-control procedure to set priorities and plan implementation of software features will help the developers stay focused on the project priorities.

In the proposed minimal scenario, shortfalls in developing software are anticipated to be covered by the management reserve. The 6 FTEs listed are anticipated to cover or augment 8 separate tasks. Priority would have to be given to supplying effort that directly enhances the ability of U.S. physicists to participate in analysis (such as a code librarian/web maintainer) at the expense of contributing more slowly to global functionality. An example of this type of trade-off would involve performing initial data analyses using sub-optimal calibration and alignment or tracking algorithms that have acceptable efficiency and purity at low instantaneous luminosity, but that need retuning or rewriting at higher instantaneous luminosity. Choices can also be made to use an initial

infrastructure that is cruder or less automated than the final version, usually at the cost of a steeper learning curve and more effort from the physicists doing analysis. Such a strategy maximizes the potential for quick physics returns early in the data collection cycle at the expense of taking much longer to understand and optimize the performance of the detector software for all the data. The risk is that sub-optimal detector software or infrastructure compromises the ability to get any physics out in a timely manner, or that temporary mitigations in the infrastructure add a maintenance burden that further stretches out delivery of the final infrastructure to the point that it is difficult to migrate users to a more efficient, leading to delays in publishing more precise measurements of higher luminosity.

Given the calls on the management reserve, in the minimal funding scenario, it is likely that one of the Tier-2 centers would not be funded at all, and it is also possible that the deployment of equipment at the Tier-1 center would proceed more slowly than is desirable to insure readiness by the time of first collisions. There are also issues of scaling of some of the software and grid middleware components that can only be tested thoroughly after all the hardware becomes deployed at the facilities. Thus, delaying procurement of equipment can lead to discovering difficulties too late to be addressed in an effective way prior to the start of data collection. Delaying staffing of the Tier-1 center can have the same effect of not thoroughly integrating all components at the scale needed for eventual analysis. Currently, no major risk has been identified with the assumed hardware technologies.

The leadership scenario reduces the risk that early physics results may be subject to compromises, and maximizes the ability to improve the precision of measurements as more integrated luminosity is added to the data samples. By reducing the calls on the management reserve, flexibility to deal with unforeseen problems is retained.

U.S. ATLAS had difficulties showing how their program would fit into the reduced scenario, with 10% funding cuts in FY2007, FY2008, and FY2009. After implementing severe cost-reduction measures in facilities and software, including elimination of one Tier-2 center and understaffing at the Tier-1, and reducing core-software personnel, a further $0.5M shortfall would need to be addressed. In this scenario, the management reserve is used, not as a reserve, but to provide basic levels of functionality, and any unanticipated problem could lead to long delays in producing physics as effort is redirected from critical tasks that are on track, to critical tasks that are off-track. At best, this compromises functionality in some number of critical tasks, and at worst it leads to a snowballing series of delays through the entire computing infrastructure. The committee feels that the reduced scenario severely jeopardizes the ability of U.S. physicists to participate in early data analysis.

2.1.3 RECOMMENDATIONS

The committee recommends that U.S. ATLAS use the change-control procedure to set priorities and plan implementation of software features.

The U.S. ATLAS model for physics analysis should be defined more clearly and in greater detail. Given that the physics centers will be needed prior to turn-on of the LHC, U.S. ATLAS must decide as soon as possible where these will be located and how they will be supported.

The committee also recommends that U.S. ATLAS take measures to ensure that facilities for U.S. physicists are available at CERN.

2.2 U.S. CMS

The U.S. CMS management team presented the status of the project during the open session.

2.2.1 FINDINGS

The U.S. CMS software and computing effort was reorganized in February 2004 and a new project execution team was put in place in response to suggestions made during last year's DOE/NSF review held at FNAL. The project is now organized in seven execution areas headed by area coordinators. The area coordinators together with the Level 1 manager and the project execution team leader form the newly-created project execution team (PET). The PET meets weekly to align project areas and assess effort issues. The split of the project into these areas separates deliverables that fall directly within the responsibility of the project from those that are merely tracked. The latter are collected into the SCC Liaison area. Four out of seven area coordinators have been identified. A search is currently ongoing for the missing three. A special effort is being made to strengthen university representation in these positions but this, the committee was told, turns out to be difficult. In reaction to lessons learned in DC04, there is also a proposal to reorganize the global CMS CPT. U.S. CMS plans to use this opportunity to further strengthen its leadership role within international CMS.

The committee was told that CERN support for visiting U.S. personnel is generally weak. This will become a critical issue at LHC startup. Therefore, U.S. CMS continues to pursue adequate support from CERN for visiting U.S. personnel. As an example, U.S. CMS has installed equipment at CERN and at the LPC to support remote operations and interaction.

The LHC Analysis Center (LPC) relies heavily on support from Fermilab staff and infrastructure. The aim of the LPC is to ensure that U.S. physicists will have an excellent opportunity for effective participation in the CMS physics program.

The committee heard that there is a strong U.S. CMS dependence on funding for grid middleware, the Open Science Grid (OSG), networking, and the newly proposed DISUN project. Individuals and key resources that need continued support have been identified.

The committee was explicitly charged to evaluate the effects of different funding scenarios on the project. The proposed trade offs protect (or enhance) effort at the expense of equipment at the Tier-1 center. In the reduced scenario, the number of Tier-2 centers will be reduced by one and the ramp up will be delayed. In addition, in order to meet schedules for software and computing deliverables, U.S. CMS anticipates reduced support for physics analysis during the crucial pre-data collection period. In the minimal scenario, U.S CMS would be able to maintain its commitments to international CMS, with some strain on the management reserve. The committee also heard that in the leadership scenario, the U.S. CMS Tier-1 center would either host the complete DST or provide 40% of all CMS Tier-1 capacity. In addition, CMS would reduce the risk of external grid deliverables by directly supporting OSG.

Fermilab is taking a strong and positive role in both U.S. CMS and in global CMS. Currently 6 of 8 L1 and L2 roles on the PET are filled by Fermilab personnel. The group intends to increase U.S. university participation in technical and managerial roles. The project benefits from the depth of technical expertise and support infrastructure inside the FNAL computing division.

2.2.2 OBSERVATIONS

U.S. CMS and international CMS continue to adapt and improve their structure based on evolving needs of the collaboration. They clearly demonstrate a strong commitment to the success of the global project. In response to deficiencies in the framework software, U.S. CMS reacted by proposing a complete redesign, and the development effort required to implement the redesigned framework is being taken up by U.S. CMS. U.S. CMS reacted quickly and enthusiastically, and the committee commends the group for this action. Global CMS will clearly benefit from this. The success of the framework is also critical so that U.S. CMS analyses can be effectively carried out away from CERN. However, it is also clear that the development of the framework software is not within the original scope of U.S. CMS and thus is an extension of the U.S. responsibilities. The committee is concerned that this additional commitment will put other U.S. CMS responsibilities at risk.

The LHC Analysis Center (LPC) provides considerable benefits to international and U.S. CMS. This model uses the knowledge gained in Run II in a coherent and productive way. It enables university and Fermilab physicists with commitments to Run II to contribute to CMS. Using the experience gained at the Tevatron, has already lead to improving readiness of CMS for turn-on. The LPC will lead to the formation of a critical mass of U.S. physicists able to contribute to physics analysis in a timely way. However, attention has to be given for support of activities as CERN from remote participation at the LPC.

In the reduced scenario, where the number of Tier-2 centers is reduced, there is a risk of having insufficient computing resources at startup. It will also reduce the ability to leverage effort at those sites. At the Tier-1 level, the reduced strategy introduces increased dependence on the worldwide Tier-1 centers and limits Monte Carlo and analysis capacity for the U.S. CMS group. However, evolving requirements and commitments from international CMS make it difficult to evaluate the impact of that strategy on U.S. CMS. Another impact of the reduced scenario is a reduction in support for physics analysis. However, the committee feels that experience in past experiments such as Run II indicates that scaling back core functionality to reduce the disproportionate impact on the physics analysis might form a better strategy.

The committee noted the dependence of U.S. CMS on externally-funded grid projects. It is not clear how these efforts will be funded in the future. This introduces a significant risk to the U.S. CMS project.

The prominent technical role that Fermilab is playing within CMS could limit the ability of FNAL scientists to participate in early physics analyses. The committee is concerned about the scarcity of university personnel in positions of technical and managerial responsibility.

2.2.3 RECOMMENDATIONS

The committee encourages U.S. CMS to continue its vigilance and adaptive approach to project needs.

The committee commends U.S. CMS for taking on additional leadership roles within international CMS. However, U.S. CMS must not to overextend its responsibilities.

The committee recommends that U.S. CMS ensure that the new responsibilities on framework software do not put U.S. CMS deliverables at risk**.**

The committee recommends that U.S. CMS conduct an internal review of the trade-offs of various funding scenarios. U.S. CMS should also engage the international collaboration to clarify more fully the Tier-1 requirements and priorities.

The committee recommends that U.S. CMS work together with U.S. ATLAS to make strong and compelling proposals for funding common projects, such as those needed for replacing the expiring funding related to U.S. grid issues.

3 FACILITIES AND GRIDS

3.1 U.S. ATLAS

Facilities and Grids were addressed in plenary talks by B. Gibbard (Tier-1/2 resources), R. Gardner (Grid R&D and deployment), and K. De (production status). Additional information was provided in breakout sessions with the committee members.

3.1.1 FINDINGS

The capacity of the U.S. ATLAS Tier-1 facility at Brookhaven was ramped up as planned this past year to a current level representing 2.5% of the capacity required in 2008. This fraction is relative to the revised 2008 capacity per the January 2005 ATLAS Computing Model, and corresponds to about 4% of last year's projected 2008 capacity. The revised ATLAS computing requirements represent an increase of 60% over previously assumed requirements for the Tier-1 Center.

Because of the increase in cost, the revised requirements have caused U.S. ATLAS to re-examine their model for disk-only Event Summary Data (ESD) data access. U.S. ATLAS is now pursuing a more modest 10% of ESD data availability through high performance RAID-based disk systems, with the remaining 90% available on “distributed disk” on each of the compute nodes served via dCache (system developed by FNAL and DESY). They will now deploy a storage-resource management system (SRM) interface to their cluster-based data stripping (as part of dCache), and the Berkeley-developed HRM/SRM as interface to the HPSS tape archive.

The plan to host 100% of the ESDs at the Tier-1 center for the U.S. ATLAS collaboration determines the Tier-1 storage and base computing requirements. They need the ability to process and skim the ESDs locally up to 30 times in a calendar year. In addition, U.S. ATLAS will allocate an extra 50% capacity at the Tier-1 center to support U.S. ATLAS physics analysis.

Over the past year, U.S. ATLAS was able to hire personnel according to their staffing plan for FY2004. They have also continued the deployment of grid-based infrastructure on the Tier-1 resources.

The Tier-2 selection process was defined, and a solicitation issued to collaborating institutions inviting proposals to site Tier 2 facilities for U.S. ATLAS. Three of the planned five Tier 2 centers for the collaboration were selected, and are situated strategically in the northeast (Harvard - Boston University), Midwest (Indiana University - University of Chicago) and the southwest (University of Texas at Arlington - Oklahoma University - University of New Mexico - Langston University).

The U.S. ATLAS Tier-1 center has defined its bandwidth requirements and provided these as inputs to the DOE/ESnet plan for major infrastructure upgrades. They expect to ramp up their WAN-connection bandwidth from OC48 in FY2005 to OC192 in FY2007 and to the level of four OC192 by the time of full LHC operations in FY2008-FY2010.

A major continuing concern is the future ability of U.S. ATLAS scientists to work at CERN. Office space and video-conference facilities remain an issue. This will be discussed at greater length in the Management section of this report, but the committee notes that the present plan is for U.S. ATLAS to primarily utilize the web and telephone communications technology in order to coordinate its U.S. and CERN activities.

In the area of grid deployment and usage, U.S. ATLAS reported continued progress since the last review. Members of the U.S. ATLAS team play significant roles in the U.S. grid projects, including GriPhyN, iVDGL and PPDG. In addition, the Grid3 effort that has grown out of iVDGL is organized and led by U.S. ATLAS and U.S. CMS scientists. Further, the nascent (but still embryonic) effort to organize U.S. Grid-project activities into a single unified framework, the Open Science Grid (OSG), is similarly being spearheaded by members of U.S. ATLAS and U.S. CMS.

U.S. ATLAS reported success in deploying its grid-work and data management systems (Capone) and demonstrating interoperability with equivalent EU derivatives within the common framework of the ATLAS grid-utilization model.

U.S. ATLAS facilities were able to support their reduced-scope DC2 effort, and were able to profit from the experience by obtaining useful grid performance metrics, such as job-failure rates and classification of failures according to impact on overall efficiency of grid utilization.

The software environment of U.S. ATLAS has been integrated with that of Grid3 at all major Grid3 sites. U.S. ATLAS has also implemented grid-based workload for event generation, G4 simulation, pileup and digitization software, and made extensive use of VDT components, such as the virtual data catalogue, Chimera and Pegasus for workflow management, and Condor-G for job execution and tracking.

Through such tools, U.S. ATLAS could monitor job-failure rates across its grid infrastructure, including the LCG and NorduGrid. On Grid3, they demonstrated a DC2 performance of about 64% success rate per submittal, leading to a reported overall efficiency of 77% after several retries of failed jobs. The reported failure rates on LCG and NorduGrid were greater than those achieved on Grid3. U.S. ATLAS was able to utilize 1.2 M CPU-hours in 2004, and completed more than 150K jobs across 20 sites. A significant outcome of this effort was the establishment of a set of grid performance specifications that will guide ATLAS into its “deployment endgame.”

U.S. ATLAS is contributing to and leading some areas of the effort for initial deployment of the OSG. This participation includes the major activities in deployment, development

of architecture, integration and validation, operations, interoperability, monitoring, policy/privilege, and governance of the OSG.

The upcoming initial deployment of OSG (Spring 2005) will take place along with several changes in Grid Service Infrastructure, such as the use of VDT 1.3.x, use of VOMS+GUMS for user authorization, SRM-dCache and SRM-drm deployment for managing cluster storage capacity and improvement of Core Monitoring Infrastructure to reduce loads on site gatekeepers. U.S. ATLAS is providing important input into the global ATLAS software effort, as lessons learned from DC2 and Common Test Beam (CTB) activities are factored into its distributed management infrastructure (Don Quixote and Windmill components).

U.S. ATLAS demonstrated successful LCG interoperability in Europe and with Teragrid in the U.S. As grid use grows, U.S. ATLAS is experiencing a dramatic increase in the need to support its many grid users.

The committee was reminded of the dependence of U.S. ATLAS on a growing body of grid tools and services, and the resulting perceived associated risk for U.S. ATLAS. The list of dependences includes SRM/dCache, SRM/DRM, grid-information and monitoring infrastructure, policy-defining activities within the OSG that may have impact on utilization of the grid. Furthermore, there is a pending need for continued funding of several U.S. Grid projects upon which U.S. ATLAS is very dependent for successful deployment of robust grid tools for the collaboration. The main Grid projects PPDG, GriPhyN, iVDGL are all coming to an end in one or two years – before the turn-on of LHC. Out of the total 6 FTEs working on grid-related activities within U.S. ATLAS, 3.75 come from non-LHC resources, including 0.6 from DOE core activities at the national laboratories and universities, and 3.15 from the various grid projects.

U.S. ATLAS provided information on their plans for dealing with the three budgetary scenarios presented by the DOE and NSF.

Under a continued minimal (baseline guidance) funding profile or a reduced (10% cut) funding profile for period FY2006 - FY2008, the plan is to maintain the Tier-1 center capabilities, if at all possible. Specifically, in the minimal funding scenario, U.S. ATLAS would be able to maintain the complement of five Tier-2 centers and their current baseline Tier-1 configuration.

In the event of a reduced funding scenario, the collaboration is prepared to reduce from five to four its complement of Tier 2 centers. Under the reduced funding scenario, the Tier-1 center would likely suffer one FTE reduction and a hardware reduction of about $300k. In addition, storage and Grid data management would be directly affected because the support for these activities comes from direct program funds. A lack of managed access to storage resources implies continued problems with production efficiency and robust operations. No effort could be dedicated to the integration of ATLAS distributed data management services with grid-based storage and site resources and Cataloging services. This shortfall would severely impact access and control of

ATLAS DC3 physics datasets and the subsequent Physics Readiness Run. There would necessarily be an increased load on the current development team, negatively affecting morale, with possible loss of expertise. Distributed software management and strengthening core infrastructure (base for reliable services) would also be impacted.

In the event of a "leadership" funding profile, corresponding to a 10% increase over current projections, U.S. ATLAS would take advantage of the opportunity to ramp up CPU and storage resources beyond the current ATLAS requirements. This would ensure that U.S. physicists could explore with greater ease the rich ATLAS data, and thereby assume additional leadership roles in physics analysis. The "leadership" scenario would provide a sixth Tier-2 center.

3.1.2 OBSERVATIONS

U.S. ATLAS reported considerable progress since the previous S&C review in the growth and development of their Tier-1 center, selection of their first Tier-2 centers, and continued deployment of grid technologies on collaboration resources. U.S. ATLAS continues to play a leading role in the U.S. grid efforts Grid3 and the planned follow-on, the Open Science Grid (OSG). Assuming that the funding agencies are able to provide the required level of support over the period 2006/2007, it appears that U.S. ATLAS will be able to have the required computing facilities in place for the turn-on of the LHC.

U.S. ATLAS has taken a reasonable approach toward accommodating the 60% increase in Tier-1 capacity requirements resulting from the recently released Computing Model. A decrease in ESD on RAID farms from 100% to 10% and concomitant introduction of file management within CPU farm disks using dCache, should yield adequate performance. The continued plan to host 100% of the ESDs at the U.S. ATLAS Tier-1 center should minimize the dependence of the U.S. collaboration to any reductions in Tier-1 center capacities of its international partners. Easy availability of all ESDs at a single site may result in U.S. ATLAS becoming a primary source of ESDs for the international collaboration. The committee was told that this eventuality can be accommodated within the baseline Tier-1 center. In the event that this becomes an issue of conflict between U.S. ATLAS and its international partners, collaboration management will have to keep an eye on the model for usage as LHC turns on.

U.S. ATLAS Tier 2 centers are following an aggregated model, whereby several regionally co-located universities pool resources to provide the functionality of a single Tier-2 center to the U.S. collaboration. The attractiveness of this situation arises from the added leveraging of resources from universities who wish to host Tier 2 facilities. There is a potential issue of critical mass in operating any single facility, and possible inefficiency if several subscale centers being aggregated into a single virtual Tier-2 center. The approach of strategically locating these centers across the U.S. is a good idea.

In the event of the 10% reduced funding scenario, given the negative impact on distributed data handling, U.S. ATLAS may want to reconsider the strategy of preserving

the funding profile of the Tier-1 facility, and instead see if some delay in Tier-1 capacity can allow the needed development work to proceed on schedule.

The committee was informed that DC3 has been de-scoped as a consequence of the experience with DC2. Discussion of these issues will continue at the upcoming Rome workshop in June 2005. DC3 will also be used as a vehicle for commissioning the Computing System. However, details for the role of Facilities were not presented.

The U.S. ATLAS Tier-1 center has a 2.5 Gbps WAN connection. In order to meet the LCG Service Challenge requirement, a 10 Gbps link is needed in late FY2005. The current Tier-1 bandwidth must be shared with RHIC, which poses a potential for conflict that has to be addressed by the DOE. While BNL Network and ESnet managers are coordinating progress on the needed upgrade, it appears that cost issues have not been resolved.

In addition, U.S. ATLAS is actively participating in the "Terapaths"- MPLS Project. This effort will investigate the integration and use of differentiated network services based on MPLS and LAN QoS in the ATLAS data network. This is an important activity that U.S. ATLAS/BNL and ESnet should be engaged in.

Cybersecurity was covered in a slide that was not presented in the Facilities talk. It did not appear that this area is an important component of their facilities and grid deployment and operations effort, and the perceived risk appears to be minimal and/or acceptable. As a follow-on to specific questions raised in the breakout session, U.S. ATLAS provided further information in this area. Specifically, the committee learned that the U.S. ATLAS Tier-1 is actively engaged in cybersecurity fora dealing with all international HEP Laboratories, and with the OSG and with LCG cybersecurity more generally. These fora are charged with developing policies and operations procedures. The OSG cybersecurity working group recently ran an exercise including both BNL and Fermilab to see if there would be appropriate notifications of a cybersecurity incident. Recently, there have been several incidents involving machines within the U.S. ATLAS Tier-1 center, and these have been based on passwords sniffing at external sites. None of the Tier-1 grid resources were affected. There are a few issues regarding differing cybersecurity policies across the U.S. ATLAS Grid that have affected ease of operation, but these are no showstoppers. The cybersecurity-related issue that currently exists with CERN is its requirement that all registered Grid users are in the CERN HR database. Since U.S. ATLAS currently shares this registration system, this implies that anyone wanting to have access to U.S. ATLAS grid resources through the ATLAS VO, even for a very brief time (e.g., a one month summer CS student), must be registered in the CERN HR database. Overall, the level of awareness of cybersecurity issues and the measures being taken seem appropriate, given that the whole issue of security, authorization and authentication on the Grid is an open and ongoing subject of discussion, definition and research.

U.S. ATLAS has done an excellent job of integrating its grid resources into the U.S. Grid3 and in the future follow-on OSG. The leadership role played by U.S. ATLAS is recognized and commendable. Its deployment of the U.S. ATLAS work-management

system Capone is a significant step. However, the fact that each of the international ATLAS Grid elements has developed a *separate* WMS environment (Dulcinea for NorduGrid and Lexor for LCG) raises some concern that the eventual operation of U.S.ATLAS resources may have inefficiencies and cause frustration for users due to the different underlying execution systems. However, at the higher level of integration, all major elements of the ATLAS grid have been shown to be interoperable under Windmill and Don Quixote WMS components.

The present failure rate of about 30% for jobs submitted to the grid is clearly too high, as acknowledged by the U.S. ATLAS team; however, it seems reasonable to expect a substantial reduction in this rate with the next generation of Grid Tools and Services (GTS) planned for availability by the time of OSG initial deployment in the Spring of 2005. The target of 10% required in production should be met by the time it is required. The level of understanding of classes of failures is commendable.

The issue of continued commitment to ensuring EU/U.S. Grid interoperability is perceived to be a risk for the U.S. effort. This risk is further increased if the support for U.S. ATLAS Grid projects is diminished under the reduced funding scenario. The evolution of EDG to LCG and appearance of the larger EU effort, EGEE, seem to make the EU grid definition a moving target. The U.S. can and does influence this process, but whether that is sufficient to ensure stability in definitions of interfacing is not clear. The implicit U.S. strategy of insulating its Grid and domestic production from the LCG / EGEE Grid development efforts appears to be a wise decision.

The committee was informed and concurs with the assessment that GTS deployment effort within U.S. ATLAS is critically short of effort for the scope of responsibility. The commitment to deploying the OSG, a presently unfunded effort, increases the workload for U.S. ATLAS Grid developers. OSG comes with few application-layer services (lack of advanced middleware services are a U.S. ATLAS VO responsibility). There are three significant areas in need of increased effort by a total of about 3 additional FTEs: grid workload management, grid data management, and tools for end-user distributed processing.

3.1.3 RECOMMENDATIONS

U.S. ATLAS is encouraged to keep up the good work to deploy its facilities in time for LHC turn-on. The proposed minimum funding scenario seems to protect Facilities and Grids from further cuts in support within the U.S. ATLAS S&C program.

At first glance, the U.S. ATLAS distributed physics analysis/expertise model that has been described appears at odds with the model of concentrating personnel/resources at the Tier-1 center. U.S. ATLAS should make an effort to better represent its thinking in these areas.

The U.S. ATLAS Tier-1 center personnel must continue to engage the DOE/ESnet program in order to ensure that BNL has sufficient bandwidth to support both the LHC

Service Challenge specifications *and* its other operational commitment to the RHIC program. The committee realizes that this problem transcends U.S. ATLAS; however, the voice of the U.S. ATLAS collaboration should remain strong and loud in this area. Otherwise, it is conceivable that the BNL operations may be required to provide an alternative to ESnet.

U.S. ATLAS has to continue to remain strongly and actively engaged on work involving interoperability with LCG and its EU partners. The interfaces must be stabilized as quickly as possible in order to assure full compatibility in the final software development/deployment endgame as the date of LHC turn-on approaches.

U.S. ATLAS must remain diligent in working on the problem of the looming end-of-funding for its Grid middleware support currently supported by other U.S. Grid projects.

U.S. ATLAS should strive to demonstrate as soon as possible that a 90% success rate for production jobs is acceptable, and that this rate can be sustained for operation in 24x7 conditions after the turn on of the LHC.

The U.S. ATLAS facilities and grid teams must remain diligent in their participation in the DOE/NSF grid Cybersecurity effort. There is a need for continued assessment of risk and in pursuing new measures and practices. The FTE requirements needed to support the cybersecurity policies of the funding agencies, beyond the umbrella of support provided by the host institutions, has to be identified and assured to be in compliance with all rules.

3.2 U.S. CMS

U.S. CMS Facilities and Grids were addressed in plenary talks by J. Bakken (Tier-1 resources), F. Wuerthwein (Tier-2C proposal), and R. Pordes (Grid Services & OSG). Additional information was provided in breakout sessions with committee members reviewing Facilities and Grids.

3.2.1 FINDINGS

The capacity of the U.S. CMS Tier-1 facility at Fermilab was ramped up as planned during the past year to a current level representing about 5% of the capacity required in 2008 at LHC turn-on. (This being calculated relative to the redefined capacity in 2008 per the December 2004 CMS Computing Model.) The revised CMS computing requirements, as presented in the recently released collaboration Computing Model, represent an increase of 100% over previous U.S. CMS assumptions for its Tier-1 Center. The staffing level for the Tier-1 facility is scheduled to be 9.5 FTE (12 individuals), which includes 1.25 FTE (3 individuals) still to be located. It was reported that the facilities functioned well during the data challenges of the past year. A commercial file system, IBRIX, has been chosen and deployed at the User Analysis Facility (UAF) to provide a scalable and robust file system for individual use.

The Tier-2 selection process was defined and a solicitation issued to collaborating institutions inviting proposals to host Tier-2 facilities for U.S. CMS. As a result of this process, four new Tier-2 sites were chosen, in addition to the three existing prototype Tier-2 centers. The new sites are MIT, Purdue, U. Nebraska and UW Madison and are budgeted in the current Research Program. A proposal (named DISUN) to augment the prototype centers is being prepared for submission to NSF by Caltech, UCSD and U. Florida, and UW Madison. The success of this proposal is assumed in the program plans of U.S. CMS S&C, and provides manpower essential for deploying the Tier1--Tier2 fabric.

The U.S. CMS Tier-1 center has defined its bandwidth requirements and provided these as inputs to the DOE/ESnet plan for major infrastructure upgrades. The present OC-12 connection to ESnet is insufficient for current U.S. CMS needs, and FNAL has established an optical connection (1 λ, OC-192) to Starlight in Chicago, which has a connection to CERN via LHCnet as well as the ESnet backbone. This connection is a research WAN, primarily intended for all of FNAL, and cannot provide the production-level traffic anticipated for LHC operations.

U.S. CMS provided information on their plans for dealing with the three budgetary scenarios presented by the DOE and NSF.

Under a continued minimal (baseline guidance) funding profile, there are two configurations being considered for the Tier-1 facility. This involves tradeoffs among the amount of data stored on disk vs. tape and the total CPU capacity. The differences correspond to a level of 20% capacity for disk and tape and 10% capacity for CPU.

In a reduced funding scenario, the amount of data stored on disk at the Tier-1 facility will be reduced, causing an increased access time to some original data (pulled from tape). It will also reduce the amount of original data stored at the Tier-1, so that access to some data will require retrieval from CMS Tier-1 facilities in other countries. The ramp-up profile for the Tier-2 centers is delayed in this scenario. This will reduce the effectiveness of U.S. scientists doing data analysis for CMS.

In the event of a “leadership” funding profile, corresponding to a 10% increase over current projections, U.S. CMS would develop a Tier-1 facility with a capacity in CPU and data storage of 40% of that of International CMS, considered by U.S. CMS to be a “fair share” for the U.S. This would provide full access to DST and additional CPU resources, and the most effective physics-analysis capabilities for U.S. scientists.

The CMS Computing Model, as well as that of ATLAS, has re-reconstruction of original data taking place at all the Tier-1 centers around the world, which creates a dependence of U.S. scientists on the capabilities of computing facilities in other countries. At the present time, CMS is seeking the needed resource commitments from the countries participating in CMS.

The LHC Physics Analysis Center (LPC) is funded at a level of $2M/year from U.S.CMS M&O. Arrangements have been made with Fermilab to host this center, which now occupies the 11th floor of Wilson Hall. It is planned that researchers visiting the LPC will have excellent connectivity and an integrated computing environment with the User Analysis Facility (UAF) component of the U.S. CMS computing facilities at Fermilab.

Members of the U.S. CMS team play significant roles in the U.S. Physics grid projects GriPhyN, iVDGL and PPDG. In addition, Grid3, a joint effort of the U.S. Physics grid projects, is organized and led by U.S. CMS and U.S. ATLAS scientists. Further, the nascent (but still embryonic) effort to organize U.S. Grid projects activities into a single unified framework, the Open Science Grid, is similarly being spearheaded jointly by members of U.S. CMS and U.S. ATLAS.

All 2004 milestones were met by the collaboration, but some incurred delays: production on the U.S. grid environment (Grid3, Spring), Requirements & Workplan for U.S. Grid environment (Grid3 to OSG plan, May to Sept.), U.S. grid interoperation with LCG (in November).

The job-failure rate on Grid3 in DC04 was at about 30%. The plan is to improve this in time for DC06, including an analysis of causes and improvements to the robustness.

The expected job-submission rate during 24x7 production running will be known only after DC06. Nonetheless, it is expected that a failure rate of 10% is achievable and tolerable. Inefficiencies associated with such failures have been factored into the CMS Computing Model.

The Data Challenge 2005 (DC05) was cancelled by CMS. It was decided that CMS Computing & Core Software (CCS) could not support the simultaneous preparation of the Physics TDR and a full data challenge. However, several activities planned in 2005 will replace some of the goals. DC06 is now defined as the “Physics Readiness Challenge”, to be fully integrated with Physics Reconstruction and Selection (PRS).

U.S. CMS is not directly dependent on the development and release schedule of the gLite middleware from EGEE. U.S. CMS has a commitment to interoperability with LCG and is confident that the LCG deployment of gLite will be consistent with interoperability with the U.S. LHC program.

The approach of U.S. CMS to cybersecurity is based on strong involvement with the FNAL Security Team, engagement of the grid community in various security forums, including the OSG Security Technical Group, GGF, and joint activities with LCG. Recent incidents have been limited to only a single one at one of its Tier-2 centers. The program plan includes a dedicated FTE in the area of U.S. CMS cybersecurity.

The looming end-of-funding for U.S. grid projects places the ability to deploy a grid-enabled framework for analysis at risk. The level of effort needed for CMS-specific work that is presently funded outside the U.S. CMS M&S budget is about 9 FTEs, distributed

among Tier-1 and Tier-2 centers across the collaboration. The NSF-funded GriPhyN project is currently scheduled to complete in 2005, and both the DOE-funded PPDG project and NSF-funded iVDGL project are scheduled to complete in 2006. These projects provide a critical and major share of the U.S. physics grid effort and are providing for the initial deployment of OSG.

CMS, led primarily by U.S. participants, has decided to redesign its Event Data Model and Framework. The committee heard that this was the outcome of lessons learned from DC04.

The CMS Computing Model requires an operational worldwide grid infrastructure and U.S. CMS is planning on OSG for the U.S. grid infrastructure. Funding for deployment and initial operation via the U.S. Physics grid projects is ending before FY07. U.S. CMS estimates that it will take $3.65M/year for continued development and deployment of grid infrastructure to support its program. Of this estimate, $0.5M/yr would support U.S. CMS FTEs.

3.2.2 OBSERVATIONS

The U.S. CMS facilities and grids team are to be commended for their clear presentations and the concise information they provided during the breakout session. The answers provided to the questions asked by the committee were extremely helpful.

U.S. CMS reported considerable progress since the last comprehensive review in the growth and development of their Tier-1 center, selection of four Tier-2 centers, and continued deployment of grid technologies. In addition, U.S. CMS continues to play a lead role in the U.S. grid efforts Grid3 and the planned Open Science Grid (OSG). Assuming that the funding agencies are able to provide the required level of support over the period 2006/2007, it appears that U.S. CMS will be able to have the computing facilities in place at LHC turn-on.

The success of the DISUN proposal is important in order to ramp up the existing prototype Tier-2 centers and to provide sufficient manpower to make most effective use of the Tier-1–Tier-2 grid fabric. We note that their proposal was turned down last year, and the success of this NSF proposal therefore cannot be guaranteed. If this funding ($1M/yr over 5 years) is not secured, the effectiveness and capacity of the Tier-2 centers in the U.S. will be significantly reduced. This will have a detrimental impact on the effectiveness of university-based physicists.

The LPC has gotten off to a good start and was noted to already have had an impact on the decision to redesign the event-data model and the analysis framework. The level of enthusiasm by U.S. CMS collaborators as well as the financial support from Fermilab appear to be sufficient to assure a positive impact on the effectiveness of U.S. physicists.

We note that the choice of the commercial software IBRIX for the user-data file system has some risk due to the dependence on vendor support for the chosen U.S. CMS

computing platform. There is also some risk that the vendor may go out of business, which would mean that the server-side hardware would have to be reconfigured to provide a different network file system (possibly NFS), but the hardware investment would be preserved. The performance of the ubiquitous NFS file system is well known to be degraded in data-intensive environments, so that the reason for seeking an alternative approach is understandable.

Fermilab is providing an OC192 research WAN connection to Starlight that is essential to satisfy current needs of U.S. CMS. This connection is for FNAL-wide use and is not intended for long-term production. This link has been adequate for the data challenge phase of CMS before LHC data-taking begins. It is not intended to be the production-level network required by U.S. CMS in 2007. ESnet has a plan but not the funding that would address this issue. The OC-192 link to Starlight, as well as the trans-Atlantic LHCnet link, are both critical to the success of the U.S. CMS physics program. Also, the CMS Computing Model has physics-quality event reconstruction portioned out to all the Tier-1 centers around the world, so that a quality high-bandwidth connection to all the CMS Tier-1 centers is essential for effective participation of U.S scientists in the CMS physics program.

U.S. CMS has been very effective at integrating its grid resources into the U.S. Grid3 and into the upcoming OSG. The excellent leadership role played by U.S. CMS principals is well recognized and commendable.

We note that the grid-based simulation production goals of DC04 were met using the Grid3 infrastructure running on 17 sites, utilizing 100 CPU years in a 3 month period, and achieved 50% greater CPU usage than that dedicated to U.S. CMS alone. This was achieved due because the U.S. CMS and U.S. ATLAS data challenges occurred in different time periods, with CMS winding down just as ATLAS was ramping up, illustrating one of the motivations for deploying shared grid resources.

The job failures rate of 30% deserves attention. The work at present to deploy storage management (SRM) on OSG should help reduce these failures. We note also the close and effective collaborative work with the Wisconsin Condor group has lead to significant improvements of middleware in the area of job management. These, along with a planned for 1 FTE starting this year in the area of workload accounting and auditing, should lead to a better understanding of failures, and bring the job-success rate to the 90% level deemed adequate for 24x7 operation. If this level is not achieved prior to DC06, it will be a cause for serious concern.

Although mentioned only briefly at this review, we note the significant effort in developing the Grid Analysis Environment lead by the Caltech group in U.S. CMS. This work will have an important role in the ability of scientists to carry out their data analysis activities. In the present schedule, the capabilities for grid-based analysis should be apparent by the time of the next review in early 2006, given the Jan 2006 milestone "Release of the DC06 Analysis Environment".

The U.S. CMS Grid effort is to be commended for its work in the area of cybersecurity and its effective integration with other projects and cybersecurity forums, including the past VOX project, the Privilege project, the OSG Security Technical Group and its participation in the LCG security group.

The committee requested U.S. CMS to state the level of funding that would be necessary to ensure a role for OSG in U.S. LHC computing. This is crucial because the current NSF/ITR and DOE/SciDAC grid projects are all scheduled for termination before FY07. We find the analysis provided to the committee to be very helpful. We note that there was no material or description presented concerning the acceptance or importance of the OSG beyond the physics and astronomy communities. It is essential to attract and to engage the broader science community with OSG for two reasons: it makes OSG more attractive to the funding agencies dealing with computer science, and it also makes development and support of the middleware for OSG more attractive to the computer scientists involved. It is a commendable effort for HEP to champion the OSG model in contrast to maintaining a long-held reputation of "doing their own thing", but it will be a challenge to attain broad participation in the OSG in time with the onset of LHC data taking. Maintaining this effort will require strong commitment from the U.S. LHC program.

3.2.3 RECOMMENDATIONS

The committee felt that the U.S. CMS facilities and grids team was already well-aware of the problems they face in the coming two years before LHC turn-on. Therefore, our recommendations are that U.S. CMS continue to pursue solutions to those critical issues that are beyond the direct control of U.S. CMS.

Fermilab and U.S. CMS must continue their effort to find a long-term solution to the problem of developing quality wide-area network bandwidth. The current arrangement with the FNAL-provided OC-192 link to Starlight (and CERN via LHCnet) is meeting the data challenge needs but the situation should be resolved in time for LHC data-taking in 2007.

U.S. CMS must continue to work with International CMS in order to assure that the global collection of CMS Tier-1 centers has sufficient capacity and capability to meet CMS needs.

U.S. CMS, together with U.S. ATLAS, must continue to press the case for support of OSG, and especially the middleware upon which the U.S. LHC grid is based, even if OSG does not remain in its current form.

4 CORE SOFTWARE

4.1 U.S. ATLAS

The U.S. ATLAS effort on core software was summarized in presentations by J. Shank and S. Rajagopalan, and presented in detail by P. Calafiura and D. Malon. A breakout session provided an opportunity for discussions with the ATLAS team.

4.1.1 FINDINGS

First collisions at the LHC are expected in July 2007. In a very real sense, U.S. ATLAS core software is entering a stage in which the software, tools and organization being developed and deployed today will likely be in use for processing and analyzing those first collisions. This makes it all the more important to ensure that the priorities underlying the decisions currently being taken are well considered. The charge to the committee and to U.S. ATLAS included a directive to weigh various directions and priorities in three different funding scenarios: reduced, minimal and leadership.

U.S. ATLAS provided comprehensive material regarding core software. They should be particularly commended for their quick research and cogent responses to questions regarding that material, which significantly clarified the issues on which the committee was directed to comment.

The responses of U.S. ATLAS core software to the recommendations in the 2004 Comprehensive Review of U.S. LHC Computing were quite reasonable. There has been a notable effort to implement a review of the code – at least its certain key components. For instance, in the data management effort, there is a "two-eyes" rule requiring that a second programmer examine revisions before they are committed. There is also an intention to hold more formal reviews for particularly critical packages.

An increased priority has been placed on support for users of U.S. ATLAS software. There is notable activity focused on improving the usability of the software. The addition of personnel from several university groups is a step in this direction.

There has also been progress on strengthening the interaction between the HLT/DAQ group and the core-software effort. An example of this can be found in the response to the "database coordination problem" identified in the previous review. Additional personnel, familiar with the Oracle DBMS, has been added, and there is a technical plan to achieve replication of data at the main Oracle database and at secondary MySQL databases.

The U.S. ATLAS effort in core software accounts for roughly 20% of the overall ATLAS effort in this area. This is a broad average, and there are some areas of development where U.S. ATLAS personnel have particular expertise and in which U.S. ATLAS has therefore chosen to concentrate its effort, and as a consequence carries a higher fraction of the effort. For instance, about three-quarters of the ATHENA framework development

effort comes from U.S. ATLAS.

The ramp-up of U.S. ATLAS computing effort has been reasonably successful, with several positions having already been filled. As a result, the previously identified personnel shortage with respect to international ATLAS has improved. In addition, there is anticipation of further increases in personnel from universities. The addition of SMU and Arizona is contingent upon funding at least at the level envisioned in the "minimal" scenario.

U.S. ATLAS scientists serve in several key positions in the management of computing for international ATLAS.

U.S. ATLAS core software met its schedule for most important deliverables for the combined test beam (CTB). However, 20% of the overall FY'04 milestones were missed.

Memoranda of understanding are being discussed and the computing model has been reviewed by the LHCC. The addition of core-software effort in Europe has helped with the U.S. ATLAS workload. Support for GEANT 3 is being dropped. Nonetheless, there has been some enlargement of scope in the U.S. ATLAS computing portfolio. One concrete example of this is the adoption of the ATHENA framework by the HLT/DAQ, which has led to new milestones for U.S. ATLAS computing.

4.1.2 OBSERVATIONS

The specific example of 8.7 FTEs based on the effort of 22 individuals in international ATLAS working on support of infrastructure was rather worrying. It seems very unlikely that such a fragmented effort will be very effective.

The personnel estimates presented to the committee are approximately constant in the out-years, this is somewhat surprising in that the menu of activities is expected to evolve from one primarily focused on development to one slewed toward maintenance and support. Based on the experience of CDF and D0, the plan and the justification that were presented do not seem *prima facie* unreasonable, but a stronger case for the desired support may be needed in light of the pressure on funding.

There was very good progress made in the ATHENA framework since the previous review, which was clearly validated during DC2. There was also good progress on data management. One concern, however, is that the incomplete DC2 phase II Tier-0 exercise leaves the subsequent phase III exercise untested and therefore at risk. The fact that the large scale of the CTB was only appreciated at a rather late stage seems to indicate inadequate communication between U.S. ATLAS computing and ATLAS management.

4.1.3 RECOMMENDATIONS

U.S. ATLAS should press the international collaboration to keep the FTEs associated with infrastructure support consolidated into fewer individuals dedicating a larger fraction of their effort. Efforts to solicit personnel from universities should be continued. Finally, a more detailed prioritization of the personnel levels needed in the out-years should be made. That prioritized list of needs should be weighed against the different funding scenarios defined in the charge.

4.2 U.S. CMS

The U.S. CMS effort on core software was summarized in presentations by L. Bauerdick, D. Green and I. Fisk, and presented in detail by D. Stickland, R. Clare and L. Sexton-Kennedy. A breakout session provided an opportunity for additional discussions with the CMS team.

4.2.1 FINDINGS

U.S. groups are continuing to make vital contributions in the development of CMS core software. This is in the form of individuals filling management roles within international CMS, and direct contributions to software development through two subprojects within the U.S. CMS research program.

As in previous years, the U.S. is well represented in management roles, with D. Stickland (project manager) and L. Taylor (deputy project manager) in leadership positions in the "Computing and Core Software" (CCS) Level-2 branch of the international CMS management structure. The committee heard that the CCS management structure was currently undergoing reorganization, but that individuals from U.S. groups were expected to fill several roles in the new structure.

The Core Applications Software (CAS) subproject of the U.S. CMS research program is a strong ongoing activity. Led by Clare, CAS employs 9 FTEs who contribute to a wide range of software projects as determined through close interaction with CCS management. This tight coupling ensures that U.S. resources are deployed efficiently to achieve maximum impact within CMS.

The U.S. CMS "Software and Support" (S&S) subproject was established within the past year to conduct major development projects. This effort is led by Sexton-Kennedy, and consists of both on-project and "base"-supported individuals. It is centered at Fermilab, benefiting from a synergistic relationship with the LPC, as well as from expertise and resources within the Fermilab Computing Division.

Overall, the U.S. CMS effort in core software accounts for about 40% (about 20 out of 47 FTEs) of the global effort on the CMS CCS project.

The U.S. core software activities have been effective, and good progress has been made in 2004. The successful production and reconstruction of Monte Carlo samples in the DC04 data challenge is an indicator of this. The full production chain was exercised, providing an opportunity for evaluation of the core and reconstruction software under realistic conditions. This pointed to areas (e.g., in the DST and Event Data Model) where desired functionality was lacking.

Also in the past year, CAS group members authored or contributed significantly to a number of packages and projects: McRunJob, PhEDEx, and publishing catalogs (for data production, moving and publishing); DAR (for binary distribution); and IgProf (performance monitoring tool). In addition, they have an ongoing involvement with conditions databases, Physh (shell environment), data management and flow, visualization, QA tools, calorimetry framework, OSCAR (CMS Monte Carlo package), production management, and the POOL file catalog.

The development of the CMS Computing Model and its recent review and approval by LHCC also reflects the important input from the U.S. core-software effort.

The committee was told that the DC04 experience pointed to weaknesses in major elements of CMS core software. In particular, the failure to produce a useful DST was a serious impediment to the analysis stage of the data challenge.

In fact, this element (DST and Event Data Model) of core software was recognized as being at risk at the time of the 2004 DOE/NSF review. In part, it is explained as a consequence of the migration from Objectivity to POOL for persistency management and ROOT for the data format. The symptoms observed in DC04 included poor performance of data management (for example, data from a single event being split across 12 files), poor metadata organization, and incomplete functionality (several key uses not being supported).

In part from the lessons learned in DC04, the international CMS CCS management structure is currently being reorganized. The committee was told that a new structure has been defined, and that candidates for the management positions have been identified, pending ratification by CMS within several weeks following this review. It was noted by the U.S. CMS group that the reorganization would provide an opportunity for additional U.S. participation in CMS CCS leadership roles.

Independent of this reorganization, a decision was taken by CMS to re-engineer their framework software. The re-evaluation of the existing framework that led to this decision was initiated by U.S. CMS collaborators, bringing experience from the FNAL collider experiments and from the LPC.

The U.S. CMS S&S project led by Sexton-Kennedy has taken on the challenge of carrying out this re-engineering effort. The committee was told that the scope of the framework redesign was limited to event-processing services, and includes at least four packages: (1) the EDM base classes, including collection classes and associations, (2) I/O

modules and service classes, (3) core-framework services, such as scheduling, event loop, provenance persistence, and (4) parameter-set configuration and management. Already, this group has made progress in redesigning the Event Data Model (EDM).

The committee was told that this re-engineering program involves an essentially complete re-write of the framework code. No code from the corresponding packages in the existing COBRA framework is expected to be re-used. The goal is to migrate forward in advance of the Magnet Test in November 2005. However, the upcoming physics TDR will be based on data analysis carried out within the existing COBRA framework.

The personnel committed to the framework re-engineering include on-project personnel, the effort of some of whom has been redirected from other CMS software-development work. The committee was told that by virtue of being based at Fermilab and involving members of the Collider experiments, the effort will benefit from previous experience with frameworks for experiments of comparable complexity. Furthermore, the development of the new EDM has benefited significantly from the connection with detector and event reconstruction groups based in the U.S., facilitated by the LPC. Finally, the recent addition of new U.S. institutions to CMS (such as Cornell) has made available personnel with experience in framework-software design and implementation, which will hopefully provide additional support for this effort.

4.2.2 OBSERVATIONS

The committee commends the U.S. CMS software and computing group on its successful and many contributions to core software for CMS. The group represents a significant fraction of the overall CMS core-software effort. Furthermore, the U.S. group is dominated by individuals who devote more than 50% of their time to CMS software, suggesting that effective use is being made of U.S. funded personnel.

The contributions from the CAS group are deemed to be very important for the successful integration of CMS software components. The CAS group also serves an essential service to the U.S. user community in its role as a source of expertise on many facets of the CMS core-software environment. Bolstering the effort in this area during the critical years ahead would be most valuable; at least, the level of effort must remain steady if the U.S. is to continue its successful contributions to the overall CMS software effort.

The committee applauds the initiative of the S&S group in taking on the challenge of re-engineering the CMS framework. This is an activity in which the U.S. group is poised to make a very large and visible impact on the CMS experiment. The committee believes every effort should be made by U.S. CMS Research-Program management and the funding agencies to provide enough support and flexibility to ensure that this effort has the resources needed for success.

However, the committee notes there is significant risk, given the very aggressive time scale necessary for readiness for physics in 2007. The size of the development team, as

presented to the committee, is much smaller than seen in comparable experiments. For example, the framework for the BaBar Experiment consists of 93 packages, with 127 classes comprising the EDM, while the ATLAS framework represents tens of FTE-years of development effort. (The committee was told that the new EDM for CMS was anticipated to consist of about 50 classes).

The committee is also concerned about the impact of the rapid development cycle that is needed to meet the goal for readiness for the November 2005 Magnet Test. In addition, the adoption of a new framework by individuals writing reconstruction and physics analysis code can be a lengthy process. Finally, it was not clear to the committee how much support for the existing framework would be required before it is retired, and to what extent this might represent a burden for the U.S. groups.

4.2.3 RECOMMENDATIONS

In light of the significant change in focus that the re-engineering of CMS framework software represents, the committee is directing its recommendations toward this aspect of the core-software effort.

The scope and work plan for the re-engineering of the CMS framework should be clarified in detail. Comparisons with contemporary experiments of similar complexity should be drawn. The plan for integrating applications, reconstruction and physics analysis software within the new framework has to be specified. Any shortfalls in personnel available for the re-engineering effort must be identified promptly. Should there be a shortfall, it is essential that U.S. CMS management intervene to secure additional effort from International CMS to ensure success of the re-engineering.

5 APPENDIX A

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

U.S. LHC Joint Oversight Group

To: D. Agarwal (LBNL)
A. Boehnlein (FNAL)
S. Gowdy (SLAC)
T. Haas (DESY)
B. Jacak (Stony Brook)
A. Lazzarini (Caltech),
P. McBride (FNAL)
D. Morrison (BNL)
J. Urheim (Indiana)

Subject: Charge for the March 2005 Review of the Software and Computing (S&C) Plans for the U.S. LHC Research Program

The Joint Oversight Group (JOG) for the U.S. Large Hadron Collider (LHC) Program of the Department of Energy and the National Science Foundation (DOE/NSF) greatly appreciates your willingness to participate in a review of the Software and Computing (S&C) plans for the U.S. LHC Research Program, which will take place at Brookhaven National Laboratory on March 1-4, 2005.

The review is intended to evaluate the progress and plans of the U.S. ATLAS and U.S. CMS S&C activities in order to assess the effectiveness of the management structures, and to learn whether the S&C activities are sufficiently strong and focused to facilitate research by U.S. collaborators at the LHC. The review will focus on the scope, cost, and schedule of the S&C plans for the period leading up to LHC turn-on, and scrutinize the needs of U.S. ATLAS and CMS for the initial period of LHC running. To this end, the collaborations will provide sufficient financial information to evaluate the following funding scenarios for U.S. LHC participation in Software and Computing for the period FY2005-FY2009, with emphasis on FY2007 and beyond:

- Minimal: The U.S. maintains most of its prior commitments and provides a minimum level of service that does not place U.S. collaborators at a disadvantage.
- Leadership: The U.S. plays a substantive role in selected areas, and is able to influence major decisions related to the architecture and deployment of hardware and software components. This implies an active development program in S&C, with the goal of addressing any future upgrades of the LHC machine and the associated ATLAS and CMS detectors.
- Reduced: Discuss the impact and priorities for the U.S. S&C effort in FY2007 assuming a 10% reduction in funding relative to current FY2007 guidelines.

The review will also consider the appropriateness of the U.S. S&C scopes, and how well tasks taken up by the U.S. are conceived, executed, and monitored. Is risk evaluated in a reasonable manner, and is scope well-matched to funding guidance? The review will also examine the adequacy of communication among the U.S. collaborations, their international counterparts and other relevant entities such as the LCG and EGEE.

The collaborations are expected to present an analysis of the critical paths for their major S&C milestones, including the impact of any deferments or de-scoping in the international program. In particular, the review should attempt to assess whether the U.S. collaborations are minimizing the possibility of a disproportionately harmful effect on the U.S. program from the possibility of any unmet international milestones.

As a guide, we point to the following issues:

MANAGEMENT AND EXTERNAL INTERACTIONS

- The overall scope of the U.S. S&C efforts and their connections to the international S&C efforts and the LCG project.
- The "portfolio balance" of U.S. involvement – does it give U.S. researchers a realistic chance for effective participation in LHC science?
- The risk of international S&C manpower shortfalls affecting U.S. milestones – are the U.S. collaborations taking appropriate steps to minimize this risk?
- Is the level of U.S. participation in S&C issues commensurate with the overall participation in the LHC program?
- Flexibility in Management: Does management have a well-conceived algorithm (including management reserve) to react and adapt to budgetary uncertainties? Do the organizational structures provide enough flexibility?
- Is the methodology of determining the funding split between M&O and S&C well conceived? Does it optimize the returns on U.S. investment in the LHC?
- Is there sufficient coordination with the international managers of the two LHC collaborations, and is U.S. management effective in communicating its unique needs and perspectives? Is there sufficient communication between U.S. ATLAS and U.S. CMS?

FACILITIES, GRIDS, AND PHYSICS-ANALYSIS MODELS

- The function and scope of the national U.S. LHC computing facilities (Tier-1 centers), their relationship to CERN (Tier-0 center) and to the regional facilities (Tier-2 centers), and whether present plans (hardware, Grid software, and networking) are adequate for satisfying the needs as outlined in the experiments' documentation of computing models. Are there additional facilities accessible to U.S. LHC that could be leveraged?
- The effectiveness of the physics-analysis models and whether they take into account the unique U.S. perspective and interactions with the whole international collaboration. What are the plans (if any) for regional analysis centers?
- Do the results of the latest round of data and service challenges lend support to the computing models proposed by U.S. scientists? Are U.S. scientists ("users") providing sufficient feedback on problems specific to U.S. involvement?

CORE SOFTWARE

- Has the personnel shortage in the international efforts improved since the previous comprehensive review of 2004, and is the U.S. carrying a fair burden of the effort in core software, including leadership responsibilities?

- Are the U.S. collaborations sufficiently vigilant in controlling "scope creep"? Is there a well-defined strategy for defining the scope of U.S. participation and for the transition from development to production software?
- How does the progress in core software stack up against the milestones shown at the previous comprehensive review? Are U.S. milestones on track and realistic? Is there any critical dependence on international milestones that brings substantial risk to U.S. deliverables?

The review will be chaired by the U.S. LHC Research Program Manager, Tom Ferbel. Jim Whitmore will serve as the NSF representative, and Saul Gonzalez as the Technical Coordinator and Secretary. You will receive all available documentation at least one week prior to the start of the review. We will appreciate closeout statements following the reviews of both ATLAS and CMS, and more formal written reports within one month of the completion of your evaluation. This will provide valuable input to the Agencies and to the experiments prior to the meeting of the Resource Review Boards in the Spring of 2005 at CERN. Your reports will also be made available to other DOE/NSF committees that review the U.S. ATLAS and U.S. CMS projects.

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

Sincerely,

_____ _____

John R. O'Fallon
Co-Chair
U.S. LHC Joint Oversight Group
Department of Energy

John Lightbody, Jr.
Co-Chair
U.S. LHC Joint Oversight Group
National Science Foundation

cc: Tom Ferbel, SC-20
Aesook Byon-Wagner, SC-20
Jim Whitmore, NSF
Glen Crawford, SC-20
Saul Gonzalez, SC-20
Craig Tull, SC-20
Dan Green, Fermilab
Bob Cousins, UCLA
Jim Shank, Boston, University
Lothar Bauerdick, Fermilab
Michael Tuts, Columbia University
Howard Gordon, BNL

5 APPENDIX B

Review Panel Membership:

D. Agarwal (LBNL)
A. Boehnlein (FNAL)
S. Gowdy (SLAC)
T. Haas (DESY)
A. Lazzarini (Caltech),
D. Olson (LBL)
D. Morrison (BNL)
J. Urheim (Indiana)

Agency Participation:

Tom Ferbel (DOE, Chair)
Glen Crawford (DOE)
Saul Gonzalez (DOE)
Craig Tull (DOE)
Jim Whitmore (NSF)

5 APPENDIX C

U.S. ATLAS SESSION: March 1-2, 2005

	March 1, 2005	
8:30	Executive Session	
9:00	Plenary Session I	
9:00	**U.S. ATLAS Program Status**	J. Shank
9:30	**U.S. ATLAS Physics Analysis Model**	M. Tuts
10:00	**ATLAS Software & Computing Status**	D. Barberis/D. Quarrie
10:30	**Coffee**	
10:45	Plenary Session II	
10:45	**U.S. ATLAS Software : Management**	S. Rajagopalan
11:15	**U.S. ATLAS Software : Core Services**	P. Calafiura
11:35	**U.S. ATLAS Software : Data Management**	D. Malon
12:00	**U.S. ATLAS Facilities**	B. Gibbard
12:30	**Lunch**	
13:15	Plenary Session III	
13:15	**U.S. ATLAS Grid Tools and Services**	R. Gardner
13:40	**U.S. ATLAS Production Status**	K. De
14:05	**U.S. ATLAS : Relation to External Groups**	J. Huth
14:45	Parallel Session	
16:30	Executive Session	
	March 2, 2005	
9:00	ATLAS answers to questions	
10:30	Executive Session	
15:00	ATLAS Closeout Session	

U.S. CMS SESSION: March 3-4, 2005

	March 3, 2005	
8:30	Plenary Session I	
8:30	**U.S. CMS S&C Overview**	LATBauerdick
9:15	**CMS Status and U.S. CMS Research Program**	Dan Green
9:40	**The LHC Physics Center**	Bob Cousins
10:05	**CMS Computing Model and Computing TDR**	David Stickland
10:30	**Coffee**	
10:45	Plenary Session II	
10:45	**U.S. CMS S&C Workplan, schedule and plans**	Ian Fisk
11:10	**CMS Computing and Core Software**	Bob Clare
11:35	**Tier-1 Facilities and Ramping to Start of Physics**	Jon Bakken
12:00	**Lunch**	
12:45	Plenary Session III	
12:45	**U.S. CMS Software Development and Support**	Liz Sexton-Kennedy
13:10	**Grid Services and Interactions with the OSG**	Ruth Pordes
13:35	**Tier-2C Proposal and Physics Analysis at Tier-2 Centers**	Frank Wuerthwein
14:00	**Coffee**	
14:15	Parallel Session	
16:00	Executive Session	
	March 4, 2005	
8:30	CMS Response to questions	
10:00	Executive Session	
12:00	CMS Closeout Session	