DOE/NSF U.S. LHC Program Office

Report on the Joint DOE/NSF Review of U.S. LHC Computing

Fermi National Accelerator Laboratory
November 27-30, 2001

Executive Summary

CMS and ATLAS will be large, general-purpose detectors for observation of very high-energy proton-proton collisions at the Large Hadron Collider (LHC). This facility is now under construction at CERN, the European Laboratory for Particle Physics near Geneva, Switzerland. In order to reap the scientific benefits of over $0.5 billion of U.S. investment in the LHC, the LHC software and computing projects must be successful in enabling physics analysis. The US Computing Projects must not only contribute an appropriate share of computing resources and software effort, but they must also ensure that U.S. Physicists will be able to fully and immediately contribute to analysis of the data and to the physics groups' research work.

A peer review of the U.S. LHC software and computing efforts was held on November 27-30, 2001, at Fermi National Accelerator Laboratory in Batavia, Illinois. The primary purpose of the review was to assess the scope, cost and schedule baselines for the near-term (defined as FY2002-2003). Included in the charge to the reviewers (Appendix A) was an assessment of the risk to the LHC Software and Computing Projects schedule or scope, given current funding profiles and overall LHC project schedules. The expert reviewers provided comments during the review, directly to the U.S. LHC collaboration members and to the DOE and the National Science Foundation (NSF) representatives. These comments, and those provided in writing from the reviewers, form the basis of this report.

The committee found that US-ATLAS had assembled a strong core software team, some of whom are making broad contributions to the international computing project in roles such as chief architect and database manager. They were also commended for their willingness to adopt and adapt other's software. Nevertheless the committee found the scope and schedule of US ATLAS core software effort, and indeed the whole international software effort, to be at risk, due to a difficult overall management environment and management indecision. There is also concern about the upcoming Data Challenge 1, which is behind schedule, and threatens to create a great deal of extra load for the US team.

To address these issues there were recommendations to international ATLAS management to enforce decisions about choices of software and to provide a concrete staffing plan for Data Challenge 1, plus a recommendation to the core-software group to become less willing to take up additional work (although to still be prepared to fire-fight when really needed, for the good of ATLAS).

The current state of core software in CMS and the plans for the future met with approval and the fact that the software is actually in use by a substantial number of people (200) five years ahead of experiment turn-on was described as a marvelous achievement. Some of the software is behind schedule due to delays in hiring. The uncertainty coming from expected dependence on “GRIDs” was noted and the reviewers were dismayed to find almost no software projects in common with other experiments, despite presumably similar requirements.

The committee found that the US-CMS software engineers were visible and effective and endorsed the request for an additional developer to work in the area of architecture. They also endorsed the plan to re-evaluate the persistency model and hoped that a simple persistent data model might assist in a rather late move to a unified GRID approach. They wished to see better coordination of software development between LHC experiments.

In the area of User Facilities both US-ATLAS and US-CMS have established working Tier 1 prototypes, in keeping with the overall LHC computing model that envisages 3 tiers of computing facilities. US-CMS has established three Tier 2 prototypes – a combined facility split between Caltech and SDSC and one at FSU and these (100 CPUs + 3 TB disk) plus the Tier 1 center (80 CPUs + 2.75 TB disk) are reportedly serving the community well as production platforms for simulation and reconstruction of monte carlo events for jet physics. US-ATLAS hardware acquisitions are delayed relative to CMS (giving some Moore’s law advantage in final capacity) and they are just now, with funding from iVDGL, establishing two Tier 2 prototypes at IU and BU. The US-ATLAS Tier 1 facility (128 CPUs + 0.75 TB disk) has reduced its dependence on the RHIC Computing Facility and has been contributing to the monte carlo production effort of ATLAS. The compute capability of the US-ATLAS Tier 1 center has not changed in the past year, due to funding constraints, but disk capacity has been doubled and independent (from RCF) HPSS servers have been installed.

US-ATLAS is considering a change of scope in their plans for the Tier 1 facility; they plan to move away from tape-based primary storage to fully disk-based primary storage, with 100% of the event summary data available on disk in the US Tier 1 facility, at all times. The committee applauded the goal of this change and hoped it could be realized within budget constraints. They also cautioned that the ramifications of this change in terms of user access patterns, network bandwidth and a revised balance of load and components between Tier 1 and Tier 2 centers, as well as implications on cost and schedule, should be studied. US-ATLAS were also encouraged to test the system with a number of boxes closer to that of the final system, and more mature software, by holding Data Challenge3 no later than early 2005. A 20% complexity test should be considered; this would require doubling the capacity of the planned 2004 Tier 1 facility and fully integrating the Tier 2 sites.

The committee found the staffing at the Tier 1 centers of both experiments to be very lean. The plans for both experiments call for more than doubling that number in FY03. The lack of staffing has caused some R&D delays and a lack of attention on interactive data analysis facilities. They felt that an additional FTE for US-CMS before FY03 was called for. They found the proposed ramp-up to a staff of 25 FTEs at the Tier 1 for US-ATLAS to be reasonable and suggested that the figure of 32 FTEs for

US-CMS might be excessive and should be re-evaluated. In fact, they wanted to see both experiments re-evaluate their long range staffing plans for their Tier 1 centers, perhaps comparing against "best-in-class" commercial operations. Similarly hardware cost estimates for the two centers will need to be re-evaluated; the two experiments are using different Moore's law-based assumptions and different procurement strategies that if modified slightly could result either in reduced costs, or increased capacity at physics turn-on.

US-ATLAS and US-CMS were found to have taken different approaches to networking needs and staffing and to the cost of such. Both experiments need to carefully document their network requirements for both the Tier 1 and Tier 2 centers and convey these to the appropriate planning bodies and funding agencies, as well as work together, and in a more broadly based effort, to address technical networking issues.

In examining the overall organization and management of the two experiments Software and Computing Projects the committee made many positive statements about the progress made in defining the project structure and filling key management positions with experienced personnel.

US-ATLAS was found to have their key US managers overloaded with multiple tasks. Some concerns were expressed about how well the interaction with the US effort and the international effort is going and a lack of ATLAS-wide responsibility assignments for software and computing. The committee found that some of the key milestones are not well defined and that project completion by FY07 may be at risk. US-ATLAS were urged to move to define their projects as well as possible to avoid mission creep; to push the international organization for clear decisions on technical issues; to monitor the productivity of the staff and ensure that is commensurate with the (relative to CMS) high costs. The committee recommended that US-ATLAS ensure that they are properly represented in the decision-making process at the CERN LHC Computing Grid project level and are able to clearly and accurately state the impact of any major changes. They will also need to be prepared to revise schedule, milestones and budget profile as the LHC schedule becomes better defined.

US-CMS contribution to CMS in the area of core application software is defined by a level of effort, with the US CMS team contributing 25% of the effort. They will rely on delivery of robust grid tools produced by the Grid projects. The budgetary profile is imposed by the agencies, therefore scope contingency is the one management tool available to the project to manage risk. However, the lack of well-defined scope may make it difficult for the US CMS SCP management to assess progress and take appropriate management–level actions in a timely fashion; especially for the Grid software packages. The multiplicity of independent Grid projects, both in the US and internationally, poses concern and perhaps a risk that some part of the US-CMS SCP grid-supported effort will eventually have to be replaced.

The committee recommended that US-CMS transition to project-imposed configuration, schedule, and change control procedure now.

DOE and NSF were urged to coordinate their resources at agency levels in order to minimize impact on project progress. There were several recommendations that spoke to the complex issue of integrating and relying on Grid project products – in the short and long term. This included a recommendation to US CMS to establish and maintain close interactions and relationships with other projects and assure that it is adequately represented on the various technical and oversight boards being constituted by CERN for LHC Computing. Also a recommendation to agencies to consider how to ensure continuing support for grid software.

Table of Contents

1 Introduction

This report is the product of the DOE/NSF review of U.S. LHC Computing Projects held at Fermilab on November 27-30, 2001. The review was charged with examining the technical scope, cost, and schedule baselines and the project management of these efforts, focusing on the near-term (through FY2003) plans of both collaborations in developing software and user facilities for the LHC experiments, ATLAS and CMS, that are now scheduled to begin taking physics data in 2007.

Eight outside experts reviewed presentations made by both collaborations on their individual projects. Their evaluations are contained in this report. At the review, they provided many recommendations to the collaborations and the agencies.

Observers from the funding agencies were also present and participated in the open discussions and executive sessions. The review was chaired by Glen Crawford from DOE with assistance from: Pepin Carolan, Mike Procario, Vicky White, Jim Yeck (DOE); and Alex Firestone (NSF). Marilyn Smith (Fermilab) provided invaluable local support.

The charge to the reviewers is shown in Appendix A. The review committee was composed of experts in computing for high-energy physics and related fields, and the committee membership is detailed in Appendix B. The agenda for the review is given in Appendix C. Separate presentations were made for the U.S. ATLAS and U.S. CMS computing efforts, totaling 1.5 days for each experiment. In addition a half-day was devoted to talks on International and Common projects, with presentations about the newly formed LHC Computing Grid project and its Software and Computing Steering Committee, as well as presentations on Grid projects and SciDAC and Network initiatives. Cost tables for both efforts for FY2002-2003 are given in Appendix D, and milestones and work schedules in Appendix E. Organization charts are collected in Appendix F.

This report, including its recommendations, represents the views of committee members on issues raised during the review, but it does not attempt to portray the personal opinions of every reviewer nor provide a comprehensive summary of all issues related to LHC computing efforts. It is intended as a compendium of expert advice to the funding agencies, and to the U.S. and international collaborators on the ATLAS and CMS experiments, on how best to achieve the goals of the software and computing projects.

2 Program Overview

CMS and ATLAS will be large, general-purpose detectors for observation of very high-energy proton-proton collisions at the Large Hadron Collider (LHC), now under construction at CERN, the European Laboratory for Particle Physics near Geneva, Switzerland. The LHC will be the highest energy accelerator in the world for many years following its completion in 2007. It will provide two proton beams, circulating in opposite directions, at an energy of 7 trillion electron volts (TeV) each, almost an order of magnitude more energy than presently achieved at the Tevatron (1 TeV per beam), at Fermi National Accelerator Laboratory (Fermilab) outside Chicago.

The two large detectors will measure and record the results of the more interesting collisions. They will be among the largest and most complex devices for experimental research ever built, and the events that they see are expected to point to exciting, even revolutionary, advances in our understanding of matter and energy. The large increase in energy over that presently available may well lead to an understanding of the origin of mass and the discovery of new families of subatomic particles.

The U.S. scientific community strongly and repeatedly endorsed U.S. involvement in the LHC program. Numerous groups of U.S. scientists at universities and national laboratories, historically supported by both the Department of Energy (DOE) and the National Science Foundation (NSF), expressed great interest in the potential physics of the LHC and in 1994 they tentatively joined the international collaborations designing the CMS and ATLAS detectors. In 1996, DOE and NSF formed the Joint Oversight Group to coordinate and manage these efforts and to negotiate an appropriate U.S. role in the LHC program.

In December 1997, the heads of DOE, NSF and CERN signed an agreement on U.S. participation in the LHC program. This was further detailed by the Experiments and Accelerator Protocols signed later that month, committing the U.S. to spend a total of $531 million on LHC construction projects, with $200 million for aspects of the accelerator and the remainder supporting the efforts of the U.S. high energy physics (HEP) community in the construction of the two large detectors. The U.S. efforts on the detectors were formalized into construction projects with baselines established in 1998.

U.S. physicists are participating in many aspects of the detectors, including important management roles. With approximately 300 physicists from 30 U.S. universities and 3 national laboratories working on each of the two large detectors, the U.S. groups comprise roughly 20% of the full collaborations and the U.S. groups plan to provide a comparable portion of each detector.

As with past large detector projects, the LHC research program, including the computers and software needed for the physics data analysis, was not made part of the detector construction projects. However, the U.S. LHC research program must be successful if the U.S. HEP community is to reap the scientific benefits of the U.S. investment in the LHC. In addition, the international scientific community is depending on the U.S. to hold up its share of the collaborative effort.

With the construction projects for both of the large general purpose detectors and the accelerator well underway, the Joint Oversight Group, in November 2000, held the first “baseline” review and assessment of the formal organization of the U.S. LHC Research Program, including the software and computing projects that will be required to generate physics results over the life of the experiments.

The U.S. LHC Research Program is a joint effort of DOE and NSF, utilizing the oversight structures established for the U.S. LHC Construction Project, as detailed in the DOE/NSF Memorandum of Understanding concerning U.S. participation in the LHC Program. In particular, this report is the result of the second formal “baseline” review of the Software and Computing Projects of both U.S. ATLAS and U.S. CMS. This review was conducted in a manner analogous to the DOE/NSF reviews of the U.S. ATLAS and U.S. CMS Detector Construction Projects. However, due to the dynamic nature of the software and computing fields, we do not expect that complete long-term(5-year) baselines could or should be set at this time. The reviewers were therefore asked to evaluate the detailed technical, cost, schedule and management plans for the near-term project efforts, up through Fiscal Year 2003

3 Core Software

US ATLAS – Findings and Evaluation

The US-ATLAS core software team brings together a strong group of physicists and senior and junior software engineers. This includes the important leadership positions of chief software architect (David Quarrie) and database manager (David Malon). They have a record of addressing important issues in computing, devising appropriate solutions, and providing valuable guidance on technical issues to the larger ATLAS international computing project. They have shared code with other projects in a variety of ways.

Nonetheless the committee found the scope and schedule of US ATLAS core software effort and the whole international software effort to be at risk, not as the result of mismanagement on the part of the US ATLAS core software effort, but due to the difficult international management environment within which they operate. The chief problem is the inability of the international software effort to take and to enforce decisions on coding standards and products. The proliferation of alternative solutions (sometimes contrary even to signed agreements), within what ought to be a streamlined architected project, is costing the US ATLAS core software effort money and manpower and is putting mission-critical exercises, such as the data challenges, at risk. Indeed we suspect that we are now seeing actual damage as the result.

We detail, in the following sections, our findings on the US ATLAS software core effort, the reasons that the impact of their important work is often diminished within their collaboration, and some suggestions for remedy.

The US ATLAS core group should be commended for their willingness to adopt HEP tools, rather than always inventing new ones. Specific examples are Gaudi and CMT. Gaudi was originally written by the LHCb experiment and is now an open source project, primarily developed by LHCb and ATLAS, with ATLAS extending it to create it's ATHENA framework. CMT is a code management tool developed at Orsay, also in prior use by LHCb, Virgo and GLAST.

They have also created a HepMC package and submitted it to CLHEP for wider use in the community.

Under the Chief Architect, the US group was instrumental in getting ATHENA approved by the collaboration and deployed. ATHENA provides the long-term insulation from the choice of persistency model. The US group also prototyped an interface allowing the use of GEANT4 as an ATHENA service. This latter work will eliminate the need for any significant standalone GEANT4 initiatives in ATLAS.

The US group is strong and qualified. Having the Chief Architect and the Database Manager on board will ensure a central and vital role for the group.

The Database group has grown from 1 to 4.5 FTEs in 2001. This has allowed them to take on significant tasks. A critical effort will be the evaluation of the persistency model, with the serious look at what may well become the LHC-standard hybrid ROOT scheme. The group's ties to the STAR experiment, which already uses this, will facilitate the effort.

The loss of Ed Frank at Chicago is unfortunate. He was instrumental in leading the effort on the design of the event model, gathering up input from the collaboration. The group will try to complete this task with existing resources, but it will be much more difficult.

The upcoming data challenge, DC1, is looked upon with some trepidation by the database group. While management assures them that CERN is organizing sufficient manpower to run the data challenge, the DB group worries that it will still get tagged for this task. This would be a serious diversion of effort for them.

Upper management appears to have troubles making decisions. This has affected the US efforts adversely in 2001: the alternative GOOFY framework for GEANT4 has lived too long, causing considerable angst and skewing of priorities. It was felt necessary to prototype the ATHENA/G4 interface early in order to deflect the GOOFY impact. Another example is the continuing support of Root/Objy persistency in the DC0 tools (eg ATLFAST). While necessary in the process of validating ATHENA, support was supposed to be assumed by the ATLFAST developers. This has not happened and continues to be a drain on the DB group. These are just examples of the US group's exposure to upper management's indecision.

The preparations for DC1 are behind schedule. This is no surprise to the US group. Managing the transitions during DC1 with a partial transition to ATHENA and bringing in GEANT4 as a test will stress the available manpower.

US ATLAS -- Recommendations

1. With the chief software architect, the US-ATLAS group commands an important position in the international collaboration. This will lead inevitably to the diversion of manpower for solving short-term problems in the future, as it has in the past. **The committee recommends to the international collaboration that it provide the chief architect with resources and authority to fulfill that role. Until that issue is resolved we recommend to US-ATLAS to continue this kind of fire-fighting for the common good of ATLAS.**

2. An additional drain of manpower is the continuous support of software that has not been approved by the collaboration. **The committee recommends to international ATLAS management to enforce decisions about choices of software in the collaboration.**

3. **The committee recommends to the US-ATLAS software group to be less willing to take up additional workload** purely driven by part of the collaborations non-compliance with already taken software decisions.

4. **ATLAS International is strongly encouraged to provide a concrete staffing plan for DC1.**

US CMS – Findings and Evaluation

The committee was on the whole very pleased with the both the current state of the core software in CMS and the plans for future development. The experiment now already regularly conducts large productions of simulated data. It has been reported that a user community of approximately 200 regularly uses the software and the samples of simulated data that have been already generated, and that the level of satisfaction is high. This is obviously a marvelous achievement five years before the experiment is to be launched, and will likely result in a well-tested software environment and a well-trained user community.

The production scripts are operational for distributed production and the original bash scripts are being converted to Python for improved functionality and maintenance.

The IGUANA project has produced a working interactive display of the detector geometry and nearly all of the physics objects currently defined. It includes powerful features such as geometry clash detection. The display is now used to debug the geometrical description of the detector while transitions are occurring in the both the unified geometrical description of the detector and in the particle simulation engine. It is also used to debug reconstruction algorithms such as track reconstruction. The same project has provided an analysis framework to the experiment.

The IGUANA project is behind schedule due to delay in hiring in 2001. The hire is now imminent. Effort was also diverted to attend to an emergency situation with SCRAM. The aim is to allow visualization of the DDD geometry by June of '02.

There has been significant progress in defining the grid requirements. However, there is significant uncertainty coming from expected dependence on GRIDs. These affect the core group in terms of effort needed to "GRID-enable" CMS tools. This is made more difficult by the multitude of GRID competitors. There was insufficient information presented for us to evaluate the GRID situation. We could not assess the effect of GRID uncertainties on CMS.

We see signs that the software is moving towards more maturity in the area of testing and support. 25% of the manpower is devoted to support activities. Automatic rebuilds of the project will commence soon, followed soon by automatic validation procedures. We think it would be beneficial to CMS to consider establishing in due course a bug-tracking system in place of the informal exchange of email messages.

The core software is poised now to support three big efforts of critical importance to the experiment: a triggering and DAQ TDR next spring, a computing TDR in 2003, and a Physics TDR in 2004. We have confidence that the core software will succeed in supporting these efforts.

US-CMS sees the need to play an important role in the re-evaluation of the persistency model. They plan to devote 1 FTE to this starting in '02 and then have that engineer move on to improve the persistency interface layer, removing Objectivity.

While CMS is not in principle different from the rest of the major LHC experiments, there is little to no common software projects shared between them. Examples of one-off projects shown to us during the review were SCRAM, DDD and Ignominy.

Half of the group's effort is being applied to the common CMS projects. These engineers are well integrated into CMS software and are an investment in the collaboration. We deem this effort well placed and that the common projects the US engineers are applied to are significant.

Though we looked carefully at the activities of mostly software engineers hired through US CMS, we also perceive that the US CMS Software leadership is both highly effective and highly visible. We support both the notion that well-placed software engineers can significantly enhance the productivity of physicists in the experiment, and the notion that some part at least of the contingent of manpower work together on joint software initiatives with other CMS and/or LHC efforts that have a bearing on the success of the CMS software and experiment, as well as specifically U.S. software initiatives and contributions.

US CMS -- Recommendations

1. **The committee endorses the request for “one additional developer to work in architecture** “to participate in the persistency reevaluation leading to the final database choice at the end of the year” (Ian Fisk). Although an explicit statement was carefully avoided in the presentations by US-CMS, the committee did get the impression that a probable outcome of such an evaluation would be the decision to change the underlying persistency mechanism. The current software is. to a large extent, directly dependent on Objectivity, the currently used Object Oriented Database. The additional manpower could be used to disentangle the software from the underlying database package, e.g. by an abstract layer, and thus facilitate a possible change of said package.

2. Since it is at present unclear what the GRID standard (if any) will be, **an intelligent, i.e. sufficiently simple, persistent data model might assist in a rather late move to a unified GRID approach.**

3. **The committee would like to find evidence of successful coordination of software needs between the LHC experiments; they could find none.** Examples named were SCRAM and DDD, packages which surely do not depend strongly on the particular experiment. The worry was expressed that the sum of additional manpower requested to develop independent software even in areas where a more coordinated approach would be feasible binds a large amount of funding which could be more efficiently used.

4 User Facilities

US ATLAS -- Summary

In keeping with the overall LHC computing plan for 3 tiers of computing facilities, ATLAS has established a working prototype Tier 1 center at BNL and, with funding from iVDGL, is establishing two tier 2 prototypes at IU and BU. The Tier 1 facility has reduced its dependence on the RHIC Computing Facility (RCF). The Tier 1 facility has been contributing to the Monte Carlo production effort of the ATLAS experiment.

The US-ATLAS group presented a change in scope of the Tier 1 facility away from a tertiary-store/disk cache model for data storage to a completely disk-resident model. The committee supports the goal of enhanced physics analysis capabilities of this change. Further consideration of the impact upon architecture scalability and Tier 2 versus Tier 1 task balance is suggested.

The US-ATLAS group is commended for their strategy of backloading their component acquisition to take advantage of Moore's law to meet their capacity requirements. They are further commended for plans to test integration of the Tier 2 facility in their Data Challenge (DC) 2, exceeding the scope aimed for by international ATLAS. The US facilities, however, will not meet the level of complexity targeted for DC2. The committee feels that a further DC meeting a 20% complexity test by early 2005 be planned. This may require a change in the component procurement plan, though it should remain compatible with a highly-backloaded plan.

The US-ATLAS group has made a commendable effort towards exposure of potential network capacity issues for the LHC program as a whole.

US ATLAS -- Findings and Evaluation

Two members of the review panel plus DOE representatives met with representatives of ATLAS for about two hours to focus upon the User Facilities sub-project. ATLAS personnel included Bruce Gibbard and Rich Baker (Tier 1 project management), Robert Gardner (IU Tier 2 project manager) plus a few others. While some prepared material was presented, much of the time was spent in friendly, informal discussions.

While the primary focus of this review is the near term (the prototype phase for the user facilities), part of the time during the facilities session was also spent on the longer range Tier 1 and Tier 2 equipment cost and staffing projections, and upon long term networking requirements

Current Status of the Facilities

The tier 1 facility is located at Brookhaven National Laboratory, adjacent to the RHIC computing facility (and sharing some infrastructure with that facility. It includes 62 dual Intel boxes, ¾ TB disk, and servers for HPSS, AFS, NFS and web and interactive access.This represents no change in the compute cluster and a tripling of the disk capacity during the past year. Compute capacity was constrained by funding priorities during this period.

The tier 1 facility operating at BNL remains a CPU-dominated facility, mainly serving Monte Carlo production for the collaboration. The facility has recently added a dedicated HPSS server, relying on the adjacent RHIC Computing Facility (RCF) for the "core" service into the RHIC storage-tek vault. Tape use is minimal at this time, and minimal activity in this area, given lean staffing, is appropriate. A new ATLAS AFS cell was recently commissioned, furthering independence from the RCF.

Tier 1 Architecture Change

US ATLAS is considering a change of scope in their plans for the tier 1 facility. The facility plan moves away from tape-based primary storage for event summary data (ESD) with 25% (cycled monthly) of the ESD available on disk, to disk-based primary storage with 100% of the ESD available on disk at the US tier 1 facility at all times. An intermediate option has 1/3 of the ESD disk resident at each of 3 ATLAS tier 1 sites. We applaud the goal of enhanced physics capability of such a plan if it can be realized within the budget guidance.

Ramifications of the change in user access patterns given the enhanced ESD access have been considered, with increases in Tier 2 capacity of roughly 100% for CPU and 200% for disk over the previous plan. However, re-evaluation of network load, task balancing between tier1 and tier 2 sites and architecture scalability remain.

The cost of this change is partially hidden by a one year increase in deployment time, taking advantage of a year slip in the CERN schedule. Whereas the plan presented a year ago projected completion in 2006 the new plan projects completion in 2007. Assuming a $7M per year operating budget in 2007 for the old plan, there is no change in total expenses through 2007 (that is, new construction costs equal old construction costs plus one year of operating costs). A comparison of capability should then be made between new plan and old plan plus one year of upgrades (not done).

One possible weakness of the new plan is the impact of future luminosity upgrades. The old plan allows the increased ESD size to be entirely tape resident. If the ESD size becomes larger than a disk pool sustainable under the operating budget, then the Tier 1 facility will be forced to fall back to the old architecture (less than 100% disk resident).

Because of this possibility, Atlas software development should not assume a 100% disk-resident ESD.

Staffing levels are lean for the tier 1 facility, with funding available for only 2.7 of the 5 FTE's projected for 2001. Two recent hires will bring the 2002 levels to 4.7. User support has benefitted from the increased support.

Current Activities

With approval of the iVDGL project, two tier 2 test facilities are being established at Indiana University and Boston University at the level of 0.25% of final capacity. The data generated in the "5%" DC2 in 2003 will be used to test grid integration.

FY02-FY03 Plans

The tier 2 sites will be integrated into production for that data challenge within the US, even though the ATLAS collaboration has targeted testing of only the tier 0 and tier 1 sites at that time.

The hardware acquisition plans are heavily backloaded for both the tier 1 and tier 2 sites. We agree with the strategy to utilize technology improvements (Moore's law) to obtain the final required capacity under the current funding profile. However, at the time of the"5%" capacity test in 2003, only about 1% of the final tier1 capacity will be installed, with correspondingly much less complexity: the US portion will be smaller by a factor of 2 from the 5% (2003) and 10% (2004) specified in the international plan. The ability to examine the scalability of their proposed system before data acquisition begins is limited under the current acquisition plan.

Staffing levels remain lean, but appear adequate over 2002 and 2003. We applaude the strategy not to build up the staff here at the expense of the software development that is much more critical during this period. The levels are too lean for this group to move out of their scope to manage the production for the data challenges, and every effort should be made to prevent such a situation. The plan to double the staffing levels between 2003 and 2004, when deployment and integration ramps up, appears reasonable. All effort should be made to meet these staffing levels. Levels then ramp smoothly to meet the estimated level of 25 needed to support production by 2007.

Networking

US-ATLAS and RHIC dominate the network bandwidth requirements of the BNL networking infrastructureNetworking requirements for the Tier 1 center are OC-192 (10 Gb/s). Funding for this capacity has not been identified. Atlas must continue to bring this problem to the attention of BNL management and the funding agencies, and should work collaboratively to pursue appropriate network upgrades within the U.S. and between the U.S. and CERN.
Atlas has also identified a need for expertise and tools for the exploitation of these next generation networks. To address this need Atlas has recommend the creation of a "SWAT team" to resolve end-to-end networking issues. The staffing suggested for this team is considerably greater than the network staff suggested for Atlas.
We believe this is a science-wide problem, and should be addressed in a way which meets the needs of Atlas, CMS, and other projects within DOE and NSF. Atlas must assist in ensuring that this need is addressed.

Long Range Cost Estimates

Experience with the RHIC computing facility serves as the basis for cost estimates. That facility primarily utilizes high-end components, and use of such components (particularly for disk) for the tier 1 and tier 2 sites has been assumed for ATLAS. Moore's law doubling used times for CPU (20 months) and disk capacity (1.4 years to 20 months) are more conservative by about 20% than those used by CMS.

US ATLAS -- Recommendations

1. **To test the system under a higher level of complexity ("number of boxes") closer to that of the final system and with more mature software, a DC3 should be attempted no later than early 2005.** A 20% complexity test should be considered, which would require doubling the capacity of the planned 2004 tier 1 facility. Furthermore, the tier 2 sites should be fully integrated into the tests at this time, with tier 1 and tier 2 playing their expected final roles with the majority of MC being generated at tier 2 and user analysis being supported at tier 1.

2. The level of 25 FTE's to support the Tier 1 facility during production appears reasonable. Nevertheless, **benchmarking against against best-in-class operations such as Celera Genomics is suggested.**

3. The data challenge 2 set should be used as effectively as possible to examine the analysis usage patterns and determine what change in scope for the Tier 1 center appears most reasonable. **ATLAS should coordinate with CMS (as they have with the disk technology studies) in technology evaluation of effective disk caching strategies as an alternative to the proposed scope change.**

4. With the base plan still including tape storage for the ESD, as well as ability to retrieve ESD from archival at the tier 0, **balanced use of commodity components at both the tier 1 and tier 2 sites should be seriously evaluated before procurement begins.**

5. **Attention must be paid to the need for increased network bandwidth and an appropriate support team beyond what is currently projected for ESnet.** Atlas should carefully document their requirements for both the Tier 1 and Tier 2 centers and convey these requirements to the appropriate planning bodies and funding agencies, and should carefully track the evolution of planning by those entities.

US CMS – Summary

In keeping with the overall LHC computing plan for 3 tiers of computing facilities, CMS has established a working prototype tier 1 center at Fermilab and three tier 2 prototypes at Cal Tech, SDSC, and FSU. Collectively, these 4 facilities have made good progress in the past year, and are well supporting US CMS activities. They have become valuable production platforms for simulation and reconstruction of monte carlo events for jet physics. Users of the facilities are reportedly very pleased with the level of support achieved in this early stage of the project.

In addition to these production and support activities, the tier 1 center has begun a series of necessary R&D activities, the most significant one for last year being an investigation of disk and RAID sub-systems. This investigation will yield both insight into the current state of this technology, plus a set of tools for benchmarking and qualifying future systems. A second project to develop an interactive data analysis cluster is now ramping up, with deployment and first experience evaluation scheduled for the next two quarters. Grid R&D activities are likewise underway at the tier 2 centers.

Overall, plans for the R&D phase (FY02-04) and the deployment phase (05-07) are well thought out, and appropriate to well support CMS goals. Personnel and hardware cost estimates are conservative, increasingly so in the later years.

US CMS – Findings and Evaluation

Two members of the review panel plus DOE representatives met with representatives of CMS for about two hours to focus upon the User Facilities sub-project. CMS personnel included Vivian O'Dell (Tier 1 project manager), Paul Avery (FSU Tier 2 project manager) plus a half a dozen others. While some prepared material was presented, much of the time was spent in friendly, informal discussions.

These discussions supplemented the excellent materials provided in advance for this review.

While the primary focus of this review is the near term (the prototype phase for the user facilities), part of the time during the facilities session was also spent on the longer range Tier 1 and Tier 2 equipment cost and staffing projections.

Good progress has been made in the last year for the user facilities on a number of fronts. To start with, a new tier 2 prototype at FSU was successfully brought on-line.

The distributed facilities are now serving both production needs and R&D activities. It is particularly noteworthy that these four sites are being used as a coherent resource for detector simulations, marking a first modest step towards the seamless computational and data grid planned for CMS. Interactions among the sites are healthy, with tier 1 staff on several occassions providing on-site assistance to the tier 2 sites.

Current Status of the Facilities

The current Tier 1 prototype consists of a cluster of 40 dual processor Linux nodes plus a small number of file servers with disk space totaling around 2 ¾ TB. The three Tier 2 prototypes (CalTech, UCSB, FSU) have an aggregate of 100 dual processor cluster nodes and 3 TB of disk.

Current Activities

Over the past year, these distributed facilities, augmented by general use computers at Fermilab, have generated over 2 million events in a significant distributed production exercise. Multiple strategies for serving pile-up events from large minimum bias samples were investigated.

In addition to this direct support of the CMS project, the user facilities have undertaken a number of on-going R&D activities. The first of these is an evaluation of RAID disk sub-systems. This evaluation has included a variety of solutions (IDE, SCSI, hardware and software RAID), using hardware deployed at the various tier 1 and tier 2 centers. As part of these evaluations, CMS has assembled and developed a number of useful benchmarking tools. The performance results, the tools, and the operational experience with the different solutions will be invaluable as the ramp up of the facilities continues. In particular, the results from the next 2 quarters will be used to guide acquisitions in FY03.

A second activity is the selection of hardware and software for a commodity user analysis cluster. Options for operating system, batch system, and other components are being evaluated, with a planned deployment at the end of the next quarter, to be followed by an evaluation period and a report on lessons learned.

The tier 2 centers, in addition to production running and collaborating in disk evaluations, have also begun the development and evaluation of necessary grid software components.

FY02-FY03 Plans

In each of the next 2 years, the tier 1 center will add an additional 40 computers (approx.) and 10-15 TB of disk, plus an appropriate amount of networking and other infrastructure, leading up to the 5% (complexity) data challenge test in mid 2003. During this same period, the 3 tier 2 centers will roughly double the number of compute nodes and make smaller increases in disk capacity.

Current staffing at the tier 1 center is very lean, roughly 5 ½ FTE, with plans to more than double that in FY03 to 12 ½. The current tight staffing has caused some R&D delays relative to initial plans and have had some impact on operations and support. Nevertheless, users are reportedly happy, and there appears to be no related critical path issues.

Staffing at the 3 tier 2 prototypes currently is 3.5 FTE, with plans to increase to 5.5 FTE by 2004, appropriate for the increase in size and activity.

Networking

The informal discussion confirmed the understanding that the Tier 1 center would be embedded within and leverage the existing and anticipated FNAL computing infrastructure. This has implications primarily for scope and staffing plans, particularly in the area of networking. Fermilab currently provides all off-site networking, as well as cyber security. This seems like an appropriate model, although as the Tier 1 center becomes fully operational it will no longer present a small perturbation on the host institution. The cost of the necessary additional high bandwidth (either ESnet or other research network connections), is currently part of the tier 1 plan. (Atlas makes the opposite assumption that bandwidth provision is off-project.) An additional concern is the expertice to effectively exploit the next generation of high speed networks.

Staffing plans for the Tier 1 center includes 2 FTE in FY03 (growing to 3 FTE in 05). This level of staffing is believed by the reviewers to be excessive under the shared understanding that the host institution is providing off-site networking. (Under this assumption, 1 FTE would be more than adequate during the proto-typing phase, possibly ramping up to 1.5 FTE during production.) Nevertheless, the panel agrees that there will need to be some involvement by LHC and HENP in the exploitation of these networks, including feeding requirements to network providers, understanding new hardware and software, and assisting in network optimization. It is our opinion that this effort should not be a CMS problem (beyond the necessary contributing / interface role), but should be a wider activity. (Similar comments can be nade about bandwidth provision.)

CMS management should insure that requirements for both network capacity and performance as well as this type of manpower be well specified and addressed, with appropriate interaction with host institutions and funding agencies.

Long Range Cost Estimates

As in last year's review, detailed cost extrapolations for the various components (disks, CPU's, etc.) were presented in the supplemental material, including data for both commodity and high end disk, and commodity and high end computing nodes. The procurement schedule has shifted relative to last year to better match funding guidance and to match the slip in the CERN schedule.

While not a focus of this review, a number of aspects of the tier 1 and tier 2 costs are worth commenting on.

First of all, several of the assumptions made in the costing of the tier 1 and tier 2 centers are probably unduly pessimistic. Disks are assumed to be used at only a 50% efficiency. Assuming instead 75% would yield a savings of 1/3 the total disk cost, or 10% of the Tier 1 hardware cost. Additionally, more expensive data center disks were used for the extrapolation (CMS is still evaluating mid-range options) and high end SMP nodes were selected for a number of data serving functions, where future quad processor commodity systems will likely be adequate (opinion).

A second point relates to the procurement strategy. Funding for both tier 1 and the tier 2 centers is flat for FY05-07, roughly $5M / year. Skewing this distribution to be more back loaded could take better advantage of Moore's Law effects and result in either reduced costs or increased capacity at physics turn-on, easily a 10% effect.

A perhaps more significant point is the large anticipated staff to run the tier 1 center, 32 FTE. This can be criticized on both a bottoms-up basis as well as overall. The cost estimates detail a large number of tasks, with a minimum assigned effort of 0.125 or 0.25 FTE, which when summed leads to an overestimate of needed manpower. On some details (user support, networking, security, data import/export) the sums seem particularly high. By way of comparison, Atlas estimates 25 FTE based upon equally valid experiences in running the RHIC computing facility. The additional 7 FTE has an associated cost equal to 25% of the annual hardware cost (non-negligible). As an even more extreme example, Celera Genomics announced at SuperComputing 2001 that they operate their very large computing center (which supports a diverse mix of jobs and a very large disk pool) with only 9 FTE.

US CMS -- Recommendations

The User Facilities are in good shape, doing good work, and there are no urgent matters which need to be addressed. The following items are minor recommendations:

1. **Staffing of the Tier 1 center at FNAL is tight, and could benefit from an additional 1 FTE before FY03 to help with a backlog of system administration and R&D activities.**

2. Network staffing proposed for FY03 is high for a tier 1 center. **A more widely based effort to address network end-to-end performance and related issues should be supported, including CMS, Atlas, BaBar, BNL/RHIC and Jefferson Lab and their associated major university partners.**

3. **Long range staffing plans for the tier 1 center should be re-evaluated, and perhaps benchmarked against best-in-class operations such as Celera Genomics.**

4. **Hardware procurements for the centers should be re-evaluated to be more back loaded.**

5. **(Future, as more experience is gained.) Hardware cost estimates for the production centers should be re-evaluated.**

5 Project Management

US ATLAS – Findings and Evaluation

Organization

The organization of the US ATLAS Physics and Computing Project is described in the US ATLAS Computing Project Management Plan (PMP), which has been recently updated. The Physics and Computing project is organized under the US ATLAS construction project and employs project management tools and procedures developed for it. This seems to be working. The Physics and Computing project itself consists of three subtasks: the core software subtask, the facilities subtask, and the physics subtask. All key management positions are filled with experienced personnel and the organization is functioning.

The relationship and interfaces to International ATLAS are well-defined and the US project is in principle well-aligned with it. US computing personnel and physicists have a prominent role in the international organization, which is to be expected given the size and strength of the US effort in ATLAS. However, we have some concerns about how well the interaction between the US effort and the international effort is really going. There does not yet seem to be an ATLAS wide set of software and computing responsibility assignments and certain decisions that could affect the US effort are not getting made.

We also note that some of the key US managers are overloaded with multiple tasks and responsibilities. A particular example of this and a cause for concern is the WBS 2.2 Level 2 manager, who has become the planning officer for the whole ATLAS computing project. The Chief Architect and the Database Coordinator for ATLAS are also from the US. While this enhances the US role and facilitates coordination between the US and International effort, it holds the danger of stressing the resources that are available to US-specific tasks and is closely related to the problem of mission creep discussed below.

Scope:

Core Software

The scope is now better defined than Nov '00. In Core Software, the US has major responsibilities for

6. The control framework

7. A portion of the data management system

8. The Event Store

- Various Collaboratory tools

These are intended to be formalized by Software Agreements between the International ATLAS organization and the US ATLAS Physics and Computing Project. These agreements will include the task description, the deliverables, the estimated level of effort, the requirements, and some process for modifying the agreement. One such agreement has now been completed for the "Development and Maintenance of a Common ATLAS Software Control Framework, and of a Data Dictionary for ATLAS Software." The software agreements would be included in an overall "International Memorandum of Understanding" between ATLAS and CERN, should that be decided to be necessary. (This is described in the PMP section 3.4. We are not clear from the presentation, that this is the final plan.)

Institutional commitments to the project are defined in Institutional Memoranda of Understanding made once in the lifetime of the project. These are augmented by Annual Statements of Work with each institution. Reporting is discussed below.

The incorporation of projects supported by multiple funding sources and some not within the control of the ATLAS collaboration create many new challenges. These include GriPhyN, PPDG, EU, iVDGL, the Grid Telemetry Project, etc. Good progress has been made in defining the scope of grid-related projects and incorporating them into the same tracking system used for tasks directly provided through the US ATLAS Computing Project. An important step has been the appointment of a single point of contact to each of the external projects to act as an ATLAS liaison, each with an identified supervisor within the US project. Issues connected to networking have been identified as an area where effort is needed but is not yet within the defined project scope. A liaison has been designated for HEP networking.

The issues connected with International ATLAS can also result in "mission creep" that can increase the project scope if not carefully controlled. The development of the software agreements is viewed as a safeguard against this and should be pursued for the remainder of the projects described in the US ATLAS Physics and Computing project WBS.

We are concerned that decisions made by the new CERN LHC Grid project, which will emphasize commonality among LHC experiments may introduce new burdens on the ATLAS and US ATLAS. ATLAS believes that it is writing the software to be resilient against the changes most likely to be required by the LHC Grid Project. US ATLAS needs to make clear at each level of the decision process, the costs of any such changes and the likely re-scoping that would be needed to accommodate it.

US ATLAS has sorting out to do with International ATLAS and within its own part of the software on definitions of the scope of many of these projects as soon as possible. There is a lot of duplication, including dagman vs magda, GRAPPA vs MOP and Athena vs Japanese Simulation Framework. This level of duplication of effort is too large and constitutes a waste of scarce resources. Firm decisions must be made by ATLAS if the full project and the US part of it are to proceed to achieve their goals in an efficient manner.

Facilities Subproject

The Facilities Subproject is based on a hierarchical model of distributed analysis facilities now accepted for the LHC. The sizing of these facilities, which include a Tier 1 regional computing center at BNL and of order 5 Tier 2 centers at US universities, has been understood for some time. The design and costing has been based on a particular model, which included reliance of tape-based robotic mass storage systems. The Facilities Subproject presented some new proposals for an all-disk based system to facilitate access to the ESDs. We agree that this will make a more flexible facility and probably improve the physics output. We are not clear that the cost of this approach has really been understood. This is essentially a new model of access, which encourages more reading of the ESDs and may have a large impact on CPU requirements and networking needs. We would like to see this very interesting idea more fully developed and would like to see a new cost estimate, schedule, and set of milestones.

ATLAS has recently established two prototype Tier2 centers using a well-defined selection process defined in an appendix to their PMP. They have a plan, also outlined in the PMP, for the selection of the permanent Tier2 centers with definite dates for completing the selection rules and moving forward with the process.

Physics Subproject

This task involves maintaining a variety of event generation packages and providing a uniform interface between each of them and the simulation package and providing support for the various Data Challenges as they occur. A search is being started to get the support person that is required to keep this on track. It is a small, well-defined, and very important effort.

Cost

It is hard to tell whether total budget request and obligation profile for FY02-FY06 is well matched with the guidance given by funding agencies but seems not too far off.

However if we assume flat funding beyond FY07, US ATLAS will be $5M short on completing full scope of the system.

Schedule

Overall, there is more breathing room due to LHC overall slip.
However, if one defines the projection completion date as having full scope of the planned system in place, the project will not be completed even beyond FY07. (We cannot estimate project completion date due to lack of agency funding guidance beyond FY07)

Some of key milestones are still not well defined .We have heard that the data challenges, especially DC2, might slip. However, we also heard that there would be more regular and “continuous” data challenge, which is probably a better plan, and that will mitigate the effect of the slippage..

Manpower

A small but persistent shortage was presented. The staffing plan for software developers is flat. Due to uncertainties in scope and lack of "standard" (ala International ATLAS), there is high probability of needing more manpower than planned (i.e.. cost increase and possible schedule delay) while scope and “standards” issues get sorted out. There is also a likelihood that more staff will be needed as data-taking approaches.

Some of the “apparent” shortfall is coming from the higher labor rate which effectively reduces the allocation for facility hardware. When compared to US-CMS, US-ATLAS in FY02-07 has

- labor cost is greater by $6.5M
- facility budget is less by $6.5M
- labor rate per FTE is 1.37 times higher

We comment that the higher labor rate is justified if the productivity of the people supported by the project is correspondingly higher. However, US ATLAS should fold that into its FTE requirements and should be monitoring productivity to make sure that the higher labor costs can be justified.

Tier1 Centers

We expect CERN to be particularly active in defining interfaces between the Tier0 and Tier1 center. It is important that the US be represented in these discussions and the US ATLAS and US CMS support each other rather than mutually annihilate.

Management and Control systems

As mentioned above, the Physics and Computing Project uses the same tools that are used to manage the detector construction project and is well-integrated with the International computing and construction projects.

Reporting and Reviews

Review Committees:

The US ATLAS project manager now appoints a Physics and Computing Advisory Panel (PCAP) to advise on the status and progress of the project. BNL also has a Project Advisory Committee, PAP, which includes computing expertise and receives the report from the PCAP.

Reporting :

The US ATLAS Physics and Computing Project now produces written quarterly reports that being with reports form each participating institution and are rolled up and collated at higher levels to give a good picture of progress towards milestones, performed vs budgeted work, etc.

Quality Assurance/Quality Control

ATLAS has some rudimentary QA/QC procedures in place and appears to have plans for extending and improving this important area. The collaboration has a QC expert. There will soon be a Software agreement on QA/QC activities.

Change Control

There is a well-defined change control process in place.

Procurements

There is a well-established plan for procurements with a clear line of signature authority for requests for proposal, purchase orders, and awards.
Summary and Recommendations:

Comments:

1. The organization of the Physics and Computing project, as defined in the revised Project Management Plan, has been improved over the last year and should adequate to the accomplishment of the project.

2. The staff is in place, is well-qualified, and is functioning well.

3. The practices and procedures that will help keep the project on-track, such as quarterly reports and software agreements, are all beginning to take place.

4. Significant progress has been made in integrating the “external projects” into the management structure and scope of the project

US ATLAS -- Recommendations

1. **US ATLAS PCP should move to define its projects as well as possible so that mission creep can be avoided**
2. **US ATLAS PCP should watch for and prevent or mitigate overload on its personnel from accepting too many responsibilities at the international level if this could compromise its ability to deliver its commitments.**
3. **The US project should push the International Organization for clear decisions on technical issues and ATLAS standards so as to avoid duplication and wasted efforts and must work to do the same within the US part of the project.**
4. **US ATLAS should monitor the productivity of its staff and make sure that it is commensurate to its costs**
5. It is important to make sure that the scope and deliverables of the project are not severely impacted by decisions made at the CERN/LHC level. US ATLAS must make sure that it is properly represented in **the decision-making process and must be prepared to clearly and accurately state the impact of any major changes to its ability to deliver**
6. As the LHC schedule becomes better defined over the next 6 months, **US ATLAS, working with International ATLAS and US funding agencies must be prepared to revise its schedule, milestones, budget profile accordingly**
7. **US ATLAS should present at the next meeting a detailed cost estimate, schedule, and milestones for its proposed modification of the architecture of the Tier1 Center at BNL** to use a disk-based system for ESD storage. They should include an attempt to describe how this new capability will change the envisioned analysis model and what additional resources will be required by those changes.

US CMS – Summary

The US CMS Software and Computing Project (SCP) is responsible for providing a number of software components to the international CMS experiment collaboration as well as software and facilities for the US CMS part of the experiment. Successful delivery of SCP products is critical to enabling the US CMS team to participate fully in the international experiment at LHC and to exploit the scientific potential within the United States.

The US CMS SC Project has responsibility for 25% of the overall CMS software development. This fraction is measured in the number of FTEs being requested for the US CMS part of the SC effort. In addition, the project is responsible for deploying and supporting a U.S. (regional) Tier 1 center at FNAL and 5 Tier 2 centers at universities across the U.S.

US CMS SC activities are well coordinated with their corresponding international CMS counterparts. The US CMS SC project management has been set up to execute and to oversee the tasks that have been identified as the responsibility of the US collaboration.

The committee was pleased to see that the project management organization that was put in place over the past year has demonstrated the leadership and management that is required to successfully complete this project. Funding constraints imposed by the agencies require a close coordination between both US LHC projects and the US grid computing R&D and implementation projects (GriPhyN, iVDGL, PPDG). This coordination is evident in the progress reported over the past year.

The US CMS SC Project Management Plan, schedule and proposed budget appear adequate to ensure successful completion of the project scope that was presented to the committee.

US CMS – Findings and Evaluation

<u>Project Scope</u>

The SCP comprises two WBS elements: user facilities (UF -- WBS 1.X.X.X.X), and core application software (CAS -- WBS 2.X.X.X.X). The scope of the UF component covers the U.S. regional Tier 1 center at FNAL, 2 prototype Tier 2 centers, and ultimately 5 Tier 2 centers. The scope of the CAS component covers two main software deliverables. These are (i) WBS 2.1, the IGUANA visualization environment, and (ii) WBS 2.3, the Distributed Data Management and Processing (DDMP) subsystem. The DDMP comprises of one set of software modules to be built and delivered “on-project” using US CMS SCP resources. A second set of software modules is *expected* to be provided as deliverables from the several U.S.- funded Grid computing programs: PPDG supported by DOE; GriPhyN and iVDGL supported by NSF.

The scope of the CAS tasks is defined relative to the overall international CMS CAS effort: the US CMS team is contributing 25% of the FTEs to the effort. Although the project has developed a well-defined WBS for the software, the approach to software development is level-of-effort build to cost. In addition, the US CMS effort will rely on delivery of robust (“industrialized”) grid tools produced by the U.S. Grid projects.

The project has a budgetary profile imposed on it by the funding agencies. Further, the time scale for software development is paced by the US CMS detector hardware development and its deployment schedule at CERN. Therefore, scope contingency is the one management tool available to the project to manage risk.
In light of the constrained fiscal and schedule realities, the committee believes that US CMS SCP has made good progress towards finalizing the project scope definition. The development of a CMS Collaboration Note defining the requirements of Grid-project delivered software for CMS applications is a significant step towards providing concrete design inputs to the U.S. Grid projects. If the Grid projects are able to deliver the specified applications kits on time, US CMS SCP will be able to take advantage of these resources without having added their development to its own project scope.

The committee was provided a summary of what CAS components US CMS expects to be provided by the various Grid projects. This is shown Table 1.

Table 1: Matrix of CMS deliverables expected from various Grid projects

Grid Project	Architecture	Job Schedulers	Data Management	Testbeds	System Services
PPDG	Common Architecture Development Release 2002	MOP 2001 Prototype	GDMP V2.0 Product	-	Authentication Development 2002
GriPhyN		DagMan 2001 Prototype	VDT yearly release; first in 2001	-	Monitoring Tools
IVDGL	-	-	-	Early 2002 starting	
EU DataGrid	-	WP-1 Grid Scheduler 2002	WP-2 GDMP V2.0 Product	??	WP-4 Monitoring

At present the definition at both the CERN/LHC level and at the international CMS levels of common project software and grid environment specifications remains fluid. The newly constituted LHC Program Office Software and Computing Committee (SC2) and Program Execution Board (PEB) will be in a position to provide standards and to ensure their adoption by all four LHC experiments. Until these are formulated however, there is a concern on the part of this committee that the specifications imposed on LHC experiments by these newly formed CERN/LHC management structures may affect the final CAS products. If standards imposed by the LCGP differ significantly from the current US CMS CAS design, then this outcome will affect the scope of CAS development. The SCP will be forced to adapt to and to adopt new requirements and specifications after the project has made substantial progress toward its own design.

The committee notes that the lack of well defined scope in the CAS tasks will make it difficult for the US CMS SCP management to assess progress and to take appropriate management-level actions in a timely fashion. This is especially so for the Grid software packages for which the US CMS SCP must rely on other projects and resources that not under its immediate control.

Cost

The committee at first had a difficult time of constructing the overall budgetary scenario for the US CMS SCP. This was due to the complexity of the funding provisions for this project. The budget for the computing and software component of the US CMS effort was not included in the cost to fabricate the US contribution to the CMS detector. As a second step, DOE and NSF agreed to jointly support these activities along a budgetary partition that was apportioned according to the level of participation in US CMS SCP of national laboratories and university research programs. Budgetary guidance has been imposed as a constraint by the funding agencies and the US CMS SCP is required to be executed within these constraints. While DOE budgetary guidance was provided via a definitive communication to US CMS, the corresponding NSF portion of the budget cannot defined prior to approval by NSF of a peer reviewed proposal submitted by US CMS. Such a grant request covering the five year period FY2002 - FY2006 is now pending at NSF and it is being reviewed as part of this annual review. If the proposal is funded at the requested level, then the total SCP budget funded by DOE and NSF for the period FY2000 – FY2006 will amount to $71.491M. The breakdown of this figure is shown in Table 2. An additional $5.966M or 8.3% of the SCP budget has been made available by NSF for grid project tasks that will provide deliverables to US CMS.

The budgetary guidance has been provided by NSF and DOE and SCP management has budgeted the scope to meet the it. There is no management reserve (contingency) in this budget. As it was discussed earlier, the programmatic elasticity is provided by an adjustable scope of the CAS and UF deliverables. Under the assumption that de-scoping in case unexpected work is incurred at some future date is acceptable to US CMS, the budget is consistent with current scope.

Table 2: Breakdown of US CMS SCP costs by year and by soruce of funds

Source	FY2000 Actual $M	FY2001 Actual $M	*FY2002 Proposed $M*	*FY2003 Proposed $M*	*FY2004 Proposed $M*	*FY2005 Proposed $M*	*FY2006 Proposed $M*	*FY2007 Proposed $M*	*Cumulative Actual + Proposed $M*
DOE	1.1647	1.785	*2.53*	*4.53*	*4.91*	*10.02*	*11.35*	*11.67*	*48.9597*
NSF*	0.31	0.48	*1.202*	*1.455*	*2.943*	*4.691*	*5.22*	*5.73*	*22.031*
US CMS Detector (DOE Loan)		0.500							
US CMS SCP Totals	**1.4747**	**2.765**	***3.732***	***5.985***	***7.853***	***14.711***	***16.57***	***18.39***	***71.491***
Off-projects funds to be expended in support of US CMS SCP									
PPDG (DOE)		0.399	*0.399*	*0.399*					*1.197*
GriPhyN (NSF)		0.582	*0.582*	*0.582*	*0.582*				*2.328*
IVDGL (NSF)			*0.466*	*0.528*	*0.545*	*0.447*	*0.455*		*2.441*
"Off-project" Totals		0.981	*1.447*	*1.509*	*1.127*	*0.447*	*0.455*		*5.966*

* **NSF totals do not include O&M budget. They are escalted 3% per year after first year.**

Schedule

Project management has revised its spending profile to match reasonably well the agency guidance funding profiles. The project schedule is fixed by the expected LHC turn-on date. Project milestones and data challenges follow from this LHC milestone. The project needs to be ready with a commissioned data analysis system sufficiently before the LHC milestone in order to identify performance issues and to provide scientists with sufficient "hands-on" experience prior to the science run. US CMS SCP has been able to take advantage of the currently projected LHC schedule delay in order to better match the funding profile constraint. The committee believes that the schedule presented at the review is credible. There is, however, little, if any, elasticity in the schedule. Project management will need to keep a careful eye on software development and respond accordingly whenever delays threaten.

Project Management

The committee was pleased to see that the project management that was put in place over the past year has demonstrated the leadership and organization that is required to successfully complete this project. Project management has developed a credible budgetary profile and the schedule to deliver the US CMS SCP deliverables in a timely manner consistent with current projections of the LHC turn-on milestone.

The project manager has implemented a reporting system that provides adequate visibility to progress on a quarterly basis. The committee reminds the project management and the funding agencies that the quarterly status reports can also provide an important vehicle for communicating U.S. programmatic concerns to the CERN management and to the international CMS project management. In addition, the reports should be sent to management of the various grid projects.

Project coordination takes place at multiple levels. The ASCB meets approximately every two months. It convenes upon the request of project management. Examples of issues that would appear on its agenda include the recent NSF proposal, definition of its scope, seeking of advice on management issues. The PMG also meets on a similar schedule, namely every two months. An internal Project Technical Board convenes every two weeks to consider progress and issues. In addition, there is a biweekly US CMS Collaboration meeting at FNAL. There are also frequent informal meetings between the project manager and the WBS Level 2 task leaders. The project manager provides biweekly reports to FNAL Computing Division and monthly reports to FNAL management. Lastly, the project manager participates in the weekly GriPhyN meetings.

The Project Management Plan (PMP) provides a detailed description of how the project management is organized and how it functions. The project appears to be following the plan. However, the committee noted that the project has not yet instituted the change control process in order to track significant changes that could affect budget or schedule.

The PMP describes how the project negotiates and controls work to be provided by collaboration institutions by means of detailed Memoranda of Understanding (MOUs). A draft MOU that was provided to the committee looks reasonable. Placement of institutional MOUs and accompanying Statements of Work (SOWs) is in progress and appears to be functioning well. The MOUs are expected to have permanence throughout the course of the US CMS SC Project, while the attached SOWs are updated on a yearly basis.

The committee appreciates the constraint imposed by the funding agencies on the budgetary profile to which the project must adhere. This budgetary constraint has led the project to adopt and enforce a build-to-cost approach to executing the project. Without budgetary freedom and with no real schedule elasticity, the only management tool the project manager has at his disposal is scope contingency. Moreover, the committee also understands that the details of the annual budgetary expenditures may, at times, come into conflict with the source of funds that are available even though the overall budget profile may be consistent with the guidance. The funding agencies should make every effort to coordinate their funding resources in order to ensure that US CMS SCP can proceed with continuity at all times.

The committee appreciates how budgetary constraints have required closer coordination between both US LHC projects and the US GRID computing R&D and implementation projects (GriPhyN, iVDGL, PPDG). This coordination is very evident in the progress reported over the past year.

The committee endorses the approach which US CMS SCP has taken in this regard. However, the committee is concerned because reliance of SCP on GRID software that is being developed as an R&D program outside its immediate oversight and programmatic control presents a new challenge for project management. The committee feels that SCP has taken an excellent first step in meeting this challenge by defining which services the project *requires* from GRID-developed software. SCP should be commended for having taken the initiative to develop a set of CMS-endorsed requirements.

The multiplicity of independent GRID projects, both in the U.S. and internationally, poses concern that as these projects evolve, there is likely to take place differentiation among product behavior and redundancy of functionality. Eventually, it is expected that LHC experiment management will impose a level of uniformity by identifying and adopting a set of specifications and standards for grid software. Depending on the outcome of these decisions, there is risk that some part of the US CMS SCP GRID-supported effort will have to be replaced.

CMS has described how they can insulate themselves against either failure of one or more of these GRID software efforts, or the need to accommodate CERN-imposed standards. In such a case, SCP would reallocate project manpower from lower priority tasks to develop the undelivered products -- perhaps with less generality or reduced functionality --themselves.

Last, the committee notes its concern that these mission-critical GRID software products are being delivered by primarily R&D programs. These programs have neither the mandate nor resources to ensure enduring support and maintenance of these software products throughout the life of the CMS experiment.

US CMS -- Recommendations

1. The committee finds that the US CMS SC Project Management Plan, schedule and proposed budget appear adequate to ensure successful completion of the project scope that was presented. **The committee recommends that the project has reached a level of maturity that warrants making the transition to a project-imposed configuration, schedule, and change control procedures.** These procedures are needed in order to ensure that SCP will be able to keep on track over the course of the next five years.

2. SC project management expressed concern that, while its budgetary spending profile is consistent with the overall budgetary envelope, there may be periods of time when SCP may lack flexibility in which funds it can spend. **The committee urges NSF and DOE to coordinate their resources at agency levels in order to minimize impact on project progress.**

3. **The committee recommends that US CMS SCP should formally integrate the task of tracking the required grid software products into its WBS.** It should include in the tracking activities a set of milestones that are mutually agreed upon between US CMS SCP and the grid software developers. SCP should also develop as milestones GO/NOGO decision points for each product that it is expecting to receive from grid projects.

4. **The committee recommends that CMS estimate its exposure to reliance on GRID projects for required software modules.** It should produce a list of all the projects and their promised deliverables. It should estimate the risk of non-delivery, and estimate the SCP manpower needed to complete the task. This should be done on a product by product basis because some tasks are better defined than others. The levels of risk for GRID deliverables vary widely since some tools are very mature and have well established performance records for delivery while others are being developed by new teams that, in some cases, are just being formed.

5. **The committee urges NSF, DOE, CERN to consider how to ensure continuing support and maintenance for GRID software throughout the life of the LHC experiments.** Failure to provide for this aspect of the grid tools development program will lead to chronic problems these tools become integrated into the experiment software systems.

6. The US collaborations could be negatively impacted by any subsequent decisions made by CERN LCDG on the standards and specifications of GRID applications. Used by the various LHC experiments. The implications of an

unfavorable choice of standard by CERN could lead to large parts of the U.S. – provided software having to be replaced or modified in order to bring them into compliance. **For this reason, it is very important that US CMS SCP establish and maintain close interactions and interrelationships with other projects who are participating in GRID technology development. In addition, it is critical that US CMS SCP be adequately represented on the various technical and oversight boards being constituted by CERN to address these issues.**

6 Appendices

Appendix A – Charge to Committee

To: John Huth, US ATLAS Associate Project Manager for Software and Computing
Lothar Bauerdick, US CMS Project Manager for Software and Computing
Date: 10/31/01
Re: Review of the U.S. LHC Software and Computing Projects

An independent peer review of the U.S.LHC Software and Computing (S&C)Projects will be conducted at Fermi National Accelerator Laboratory on November 27-30,2001. This review will continue systematic oversight of the U.S. LHC Research Program.

The scope of this review is to include both the individual U.S.ATLAS and U.S.CMS S&C Projects and the common projects, which provide software resources to both efforts. The goal of this review is to assess the scope, cost and schedule baselines for the U.S.LHC Software and Computing Projects, and their proposed management structures. Both US ATLAS and US CMS should present self-consistent baseline plans targeted to the funding guidance received from DOE and NSF, and separately address how they would use incremental funds. Due to the dynamic nature of the software and computing fields, we do not expect that complete long-term (5-year)baselines could or should be set at this time. Thus, we are requesting a detailed technical, cost, schedule and management review of only the near-term project efforts, up through Fiscal Year 2003. However,the review committee will make its best effort to gauge whether these near-term efforts can reasonably be extrapolated to the long-term requirements of the Research Program.

The charge for this review will be to assess:

• The overall scope of the U.S.LHC S&C efforts, and their connections to both the international LHC S&C efforts and the CERN LHC Computing Project ;
• The risk to U.S.LHC S&C schedule or scope given current funding profiles and overall LHC project schedules;
• The contributions of each of the U.S. collaborations in providing and supporting "core " and detector--specific software deliverables to the international ATLAS and CMS computing efforts;
• The function, scope and structure of the national ("Tier 1 ")U.S.LHC computing facilities, and their relationship to any smaller regional and university facilities;
• The integration of computing infrastructure efforts (such as networking and Grid computing)into the planning and execution of the US LHC S&C projects ;
• The plans of the U.S. collaborations to provide computing resources to users and their success in integrating them into the software development process;
• Existing and possible common computing projects which could benefit both ATLAS and CMS; and
• The Project Management Plans, organizational structures, and adequacy of personnel for each of the U.S.LHC S&C Projects.

We appreciate your assistance in this matter. These reviews are an important element of the Department of Energy/National Science Foundation joint oversight of the U.S.LHC Project and help ensure that the U.S.meets our commitments on cost and schedule.

Appendix B – Members of Review Committee

Glen Crawford, Chair	**DOE**
Joe Boudreau	**U. Pittsburgh**
Joel Butler	**Fermilab**
Aesook Byon-Wagner	**Fermilab**
Richard Dubois	**SLAC**
Lawrence Gibbons	**Cornell**
Albert Lazzarini	**Caltech**
Matthias Messer	**BNL**
Chip Watson	**Jefferson Lab**

Appendix C – Review Agenda

U.S. LHC Software and Computing Review
Nov. 27-30, 2001
Fermilab

Tuesday November 27

Executive Session		Comitium
8:30am Executive Session		**Full Committee**

International and Common Projects Overview		1 WEST
9:30	**Welcome**	**M. Witherell**
9:35	**CERN's Plans for LHC Computing**	**L. Robertson**
10:15	**LHC S&C Steering Committee**	**M. Kasemann**
10:45	**Coffee**	
11:00	**GriPhyN and LHC Computing**	**P. Avery**
11:30	**SciDAC and Networking Initiatives**	**V. White**
12:00	**Lunch**	

US CMS Overview Session		1 WEST
1:00 pm	**CMS Overview**	**M. della Negra**
1:20	**U.S. Project Overview**	**LAT Bauerdick**
1:50	**CMS and U.S. CMS Physics Plans**	**J. Branson**
2:10	**U.S. CMS Detector Project**	**D. Green**
2:30	**CMS Computing and Core Software**	**D. Stickland**
3:00	**Coffee**	
3:15	**U.S. CMS Core Application Software**	**I. Fisk**
3:45	**U.S. CMS User Facilities**	**V. O'Dell**
4:15	**U.S. CMS Data Grid**	**H. Newman**
4:35	**CMS Software and Computing Demo.**	**I. Fisk, J. Amundson, G. Graham**

Executive Session		Comitium
5:00	**Formulation of Questions for US CMS**	**Full Committee**
6:00	**Adjourn**	

Dinner		User Center
6:30	**Drinks**	
7:00	**Dinner at Chez Leon**	

Wednesday November 28

Executive Session		Comitium
8:30am	**US CMS Response to Questions**	**Full Committee, CMS Management**
Breakout Session CMS Project Management		Beauty Parlor, WH12NW and Neptune VR
9:30	**Project Organization Planning and Budget**	**LAT Bauerdick**
Breakout Session CMS Core Software		West Wing , WH10NW and Sun VR
9:30	**Core Application Software**	**I. Fisk**
Breakout Session CMS User Facilities		Racetrack, WH7X and Liberty VR
9:30	**Project Organization Planning and Budget**	**LAT Bauerdick**
Executive Sessions		Breakout Rooms, Comitium
11:00	**Subcommittee Executive Session**	**Subcommittees**
12:00 pm	**Working Lunch at Comitium**	**Full Committee**
1:30	**Executive Session with U.S.CMS Management**	**Full Committee**
2:00	**Subcommittee Sessions/Writing**	**Subcommittee**
3:30	**Coffee Break**	
4:00	**Closeout Dry Run**	**Full committee**
Closeout CMS		1 West
5:30	**Closout with U.S. CMS**	
	General Conclusions	**G. Crawford**
	Core Application Software	**R. Dubois**
	User Facilities	**C. Watson**
	Project Management	**A. Lazzarini**
6:30	**Adjourn**	

Thursday November 29

Executive Session		Comitium
8:30am Executive Session		**Full Committee**
US CMS Overview Session		1 WEST
9:00 am	**Project status, budget**	**J. Huth**
9:40	**U.S. ATLAS Core Application Software**	**T. Wenaus**
10:25	**Coffee**	
10:45	**U.S. ATLAS User Facilities**	**B. Gibbard, R. Baker**
11:45	**Networking**	**S. McKee**
12:05	**GRID planning**	**T. Wenaus, J. Schopf**
12:35	**Lunch**	
Breakout Session ATLAS Project Management		Beauty Parlor, WH12NW and Neptune VR
1:30 pm	**ATLAS overview**	**N. McCubbin**
	Project Management	**B. Ernst/H. Gordon**
	US computing management details	**J. Huth**
Breakout Session ATLAS Core Software		West Wing , WH10NW and Sun VR
1:30 pm	**Architecture/Framework**	**D. Quarrie**
	Data Management	**D. Malon**
	Sub-system software	**J. Shank**
Breakout Session ATLAS User Facilities		Racetrack, WH7X and Liberty VR
1:30 pm	**Tier 1 details**	**B. Gibbard/R. Baker**
	Tier 2	**R. Gardner**
	More on grids	
Executive Session		Comitium
3:00	**Subcommittee Executive Sessions**	**Breakout Rooms**
4:00	**Break**	
4:30	**Formulation of Questions for US ATLAS**	**Full Committee**
6:00	**Executive Session with US ATLAS Mgmt**	
6:30	**Adjourn**	

Friday November 30

Executive Session		Comitium
8:30am	**US ATLAS response to questions**	**Full committee, ATLAS management**
9:00	**Subcommitee Sessions/writing**	**Subcommittees**
10:30	**Executive Session**	**Full committee**
11:00	**Closeout Dry Run**	**Full committee**
12:30 pm	**Working Lunch**	**Full committee**

Closeout CMS		1 West
1:30pm	**Closout with U.S. ATLAS**	
	General Conclusions	**G. Crawford**
	Core Application Software	**J. Boudreau**
	User Facilities	**C. Watson**
	Project Management	**J. Butler**
	Concluding Remarks	**G. Crawford**
2:30	**Adjourn**	

Appendix D – Cost Tables

Agency Funding Guidance for each experiment

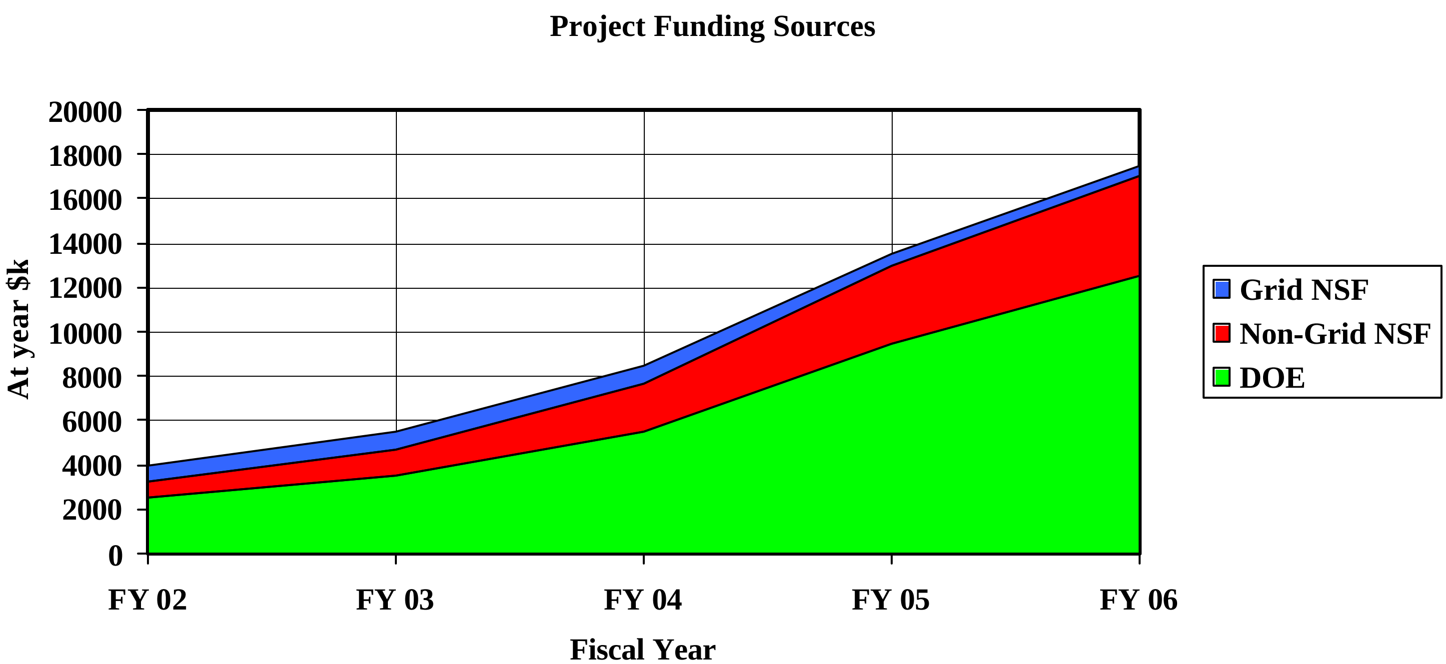

US CMS

The following table shows the profile for US CMS contribution (in FTEs) to the Core Software effort, compared to the total CMS Core software effort. The actual current contribution is 6 FTEs.

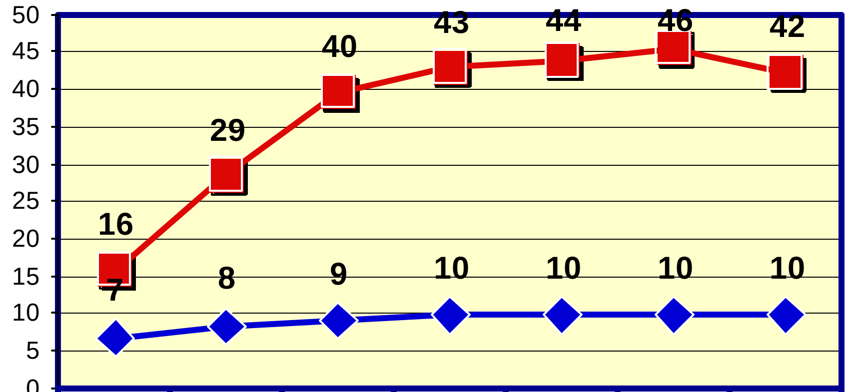

	2000	2001	2002	2003	2004	2005	2006
Total CMS Offline + Online S/W FTE	16	29	40	43	44	46	42
U.S. CMS Core-SW contribution	7	8	9	10	10	10	10

The funding for US CMS for FY2000 and FY2001 is shown in the table below.

All funds in AY$ x 1000	**FY2000**	**FY2001** Requested	**FY2001** Received	**Total** Received
DOE	1164.7	2000.0	1785.0	2949.7
NSF	310.0	1500.0	480.0	790.0
Loan from U.S. CMS Detetector Project			500.0	500.0
Total	**1474.7**	**3500.0**	**2765.0**	**4239.7**

The NSF funding for US CMS Software and Computing, through Grid Projects, is shown in the table below.

NSF Grid R&D Funding for CMS	2002	2003	2004	2005	2006
GriPhyN					
Total, including CS and all Experiments	2543	2543	2241		
CMS Staff	582	582	582		
iVDGL					
Total, including CS and all Experiments	2650	2750	2750	2750	2750
CMS Equipment	232	192	187	57	65
CMS Staff	234	336	358	390	390

Included in the table below is the non-GRID funding for Software and Computing for US CMS, requested from the NSF.

Budget Profile NSF

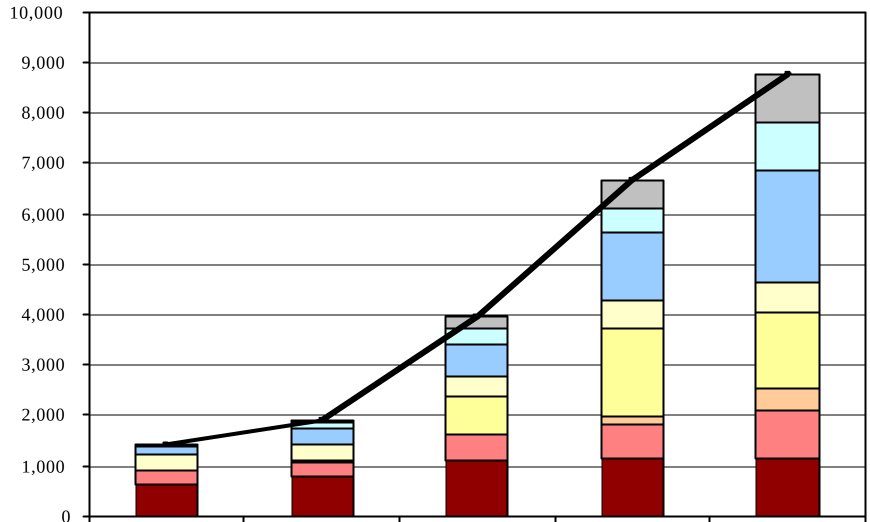

	2002	2003	2004	2005	2006
Escalation (3%/year)		56	227	567	980
Upgrade R&D	60	120	324	498	948
Detector M&O	162	327	630	1319	2219
PO, E/O, Reserve	274	293	417	555	587
Tier-2 Equipment	0	48	750	1750	1500
Central Staff				155	465
Tier-2 Staff	308	284	494	697	928
Software Staff	620	787	1113	1135	1158
Grand Total, escalated	1424	1916	3955	6676	8784

The following table shows the budget profile for the US CMS User Faciltiies and how that translates into computing capability is shown in the next table.

Fiscal Year	2002	2003	2004	2005	2006	2007	Total	2008 (Ops)
1.1 T1 Regional Center	0	0	0	2,866	2,984	2,938	8,788	2,647
1.2 System Support	29	23	0	35	0	0	87	15
1.3 O&M	0	0	0	0	0	0	0	0
1.4 T2 Regional Centers	232	240	870	1,870	1,500	1,750	6,462	1,250
1.5 T1 Networking	61	54	42	512	462	528	1,658	485
1.6 Computing R&D	511	472	492	0	0	0	1,476	0
1.7 Det. Con. Support	84	53	52	0	0	0	189	0
1.8 Local Comp. Supp.	12	95	128	23	52	23	333	48
Total	**929**	**938**	**1,584**	**5,306**	**4,998**	**5,239**	**18,992**	**4,446**
Total T1 only	697	698	714	3,436	3,498	3,489	12,530	3,196

Installed Capacity Tier-1 Facility

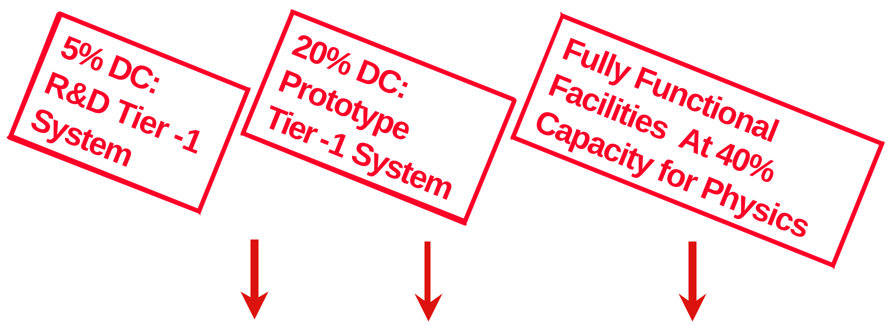

Fiscal Year	2002	2003	2004	2005	2006	2007
Simulation CPU (Si95)	2,000	3,000	4,000	7,200	28,800	72,000
Analysis CPU (Si95)	750	2,100	4,000	8,000	32,000	80,000
Disk (TB)	16	31	46	65	260	650
Server CPU (Si95)	50	140	270	1,500	6,000	15,000

The requested funding profile for the entire US CMS Software and Computing project is shown below, together with current agency guidance.

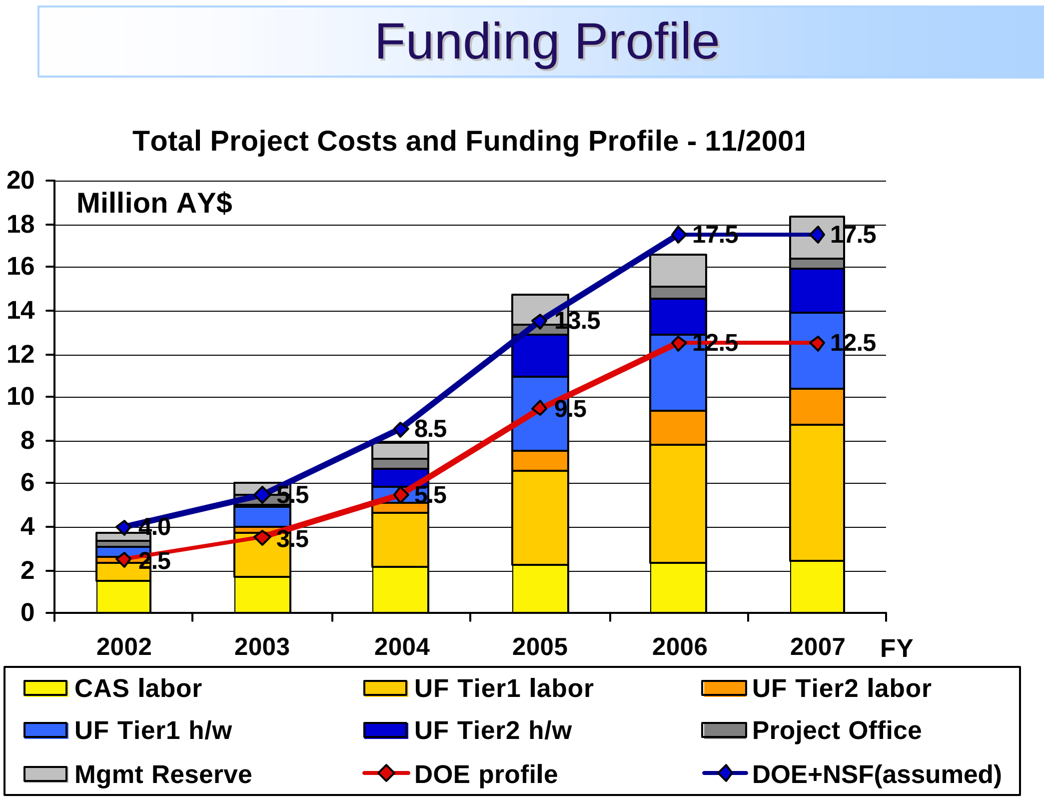

Total Project Costs

In AY$M

Fiscal Year	2002	2003	2004	2005	2006	2007
Project Office	0.32	0.48	0.49	0.51	0.52	0.54
DOE	*0.15*	*0.31*	*0.32*	*0.33*	*0.34*	*0.35*
NSF	*0.17*	*0.17*	*0.18*	*0.18*	*0.19*	*0.19*
Software Personnel	1.49	1.72	2.14	2.25	2.36	2.48
DOE	*0.87*	*0.91*	*0.96*	*1.01*	*1.06*	*1.11*
NSF	*0.62*	*0.81*	*1.18*	*1.24*	*1.30*	*1.37*
UF Personnel	1.14	2.26	3.00	5.26	6.99	7.89
for Tier-1 DOE	*0.83*	*1.97*	*2.48*	*4.33*	*5.42*	*6.28*
for Tier-2 NSF	*0.31*	*0.29*	*0.52*	*0.93*	*1.57*	*1.62*
UF Equipment	0.45	0.75	1.51	5.35	5.19	5.52
for Tier-1 DOE	*0.45*	*0.70*	*0.71*	*3.44*	*3.50*	*3.49*
for Tier-2 NSF	*0.00*	*0.05*	*0.80*	*1.91*	*1.69*	*2.03*
Management Reserve	0.34	0.77	0.71	1.34	1.51	1.96
DOE	*0.23*	*0.64*	*0.45*	*0.91*	*1.03*	*1.44*
NSF	*0.11*	*0.13*	*0.27*	*0.43*	*0.47*	*0.52*
Total Costs	**3.73**	**5.98**	**7.86**	**14.71**	**16.57**	**18.39**
Total DOE	***2.53***	***4.53***	***4.91***	***10.02***	***11.35***	***12.67***
Total NSF	***1.20***	***1.45***	***2.94***	***4.69***	***5.22***	***5.73***

Budget Issue in FY2003

Shortfall of $1M in FY2003

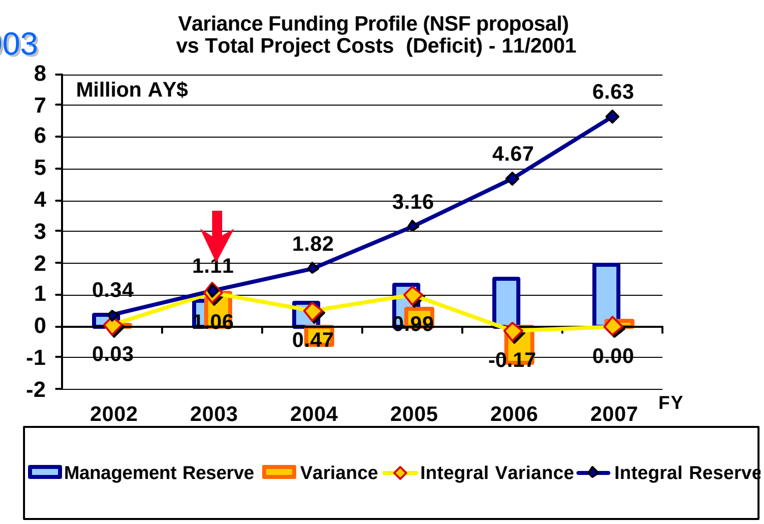

- Not fully covered by the reserve, as we have to pay back $500k loan
- Request to get this from DOE to help with the R&D phase!
- Or "advance" from NSF (original guidance was $2M!)
- Otherwise need to delay some of the hiring for UF some more

US ATLAS

The chart below shows the profile for building capacity at the Tier 1 US ATLAS computing center and how it has been changed to delay purchases until FY07.

CPU Capacity (kSi95)

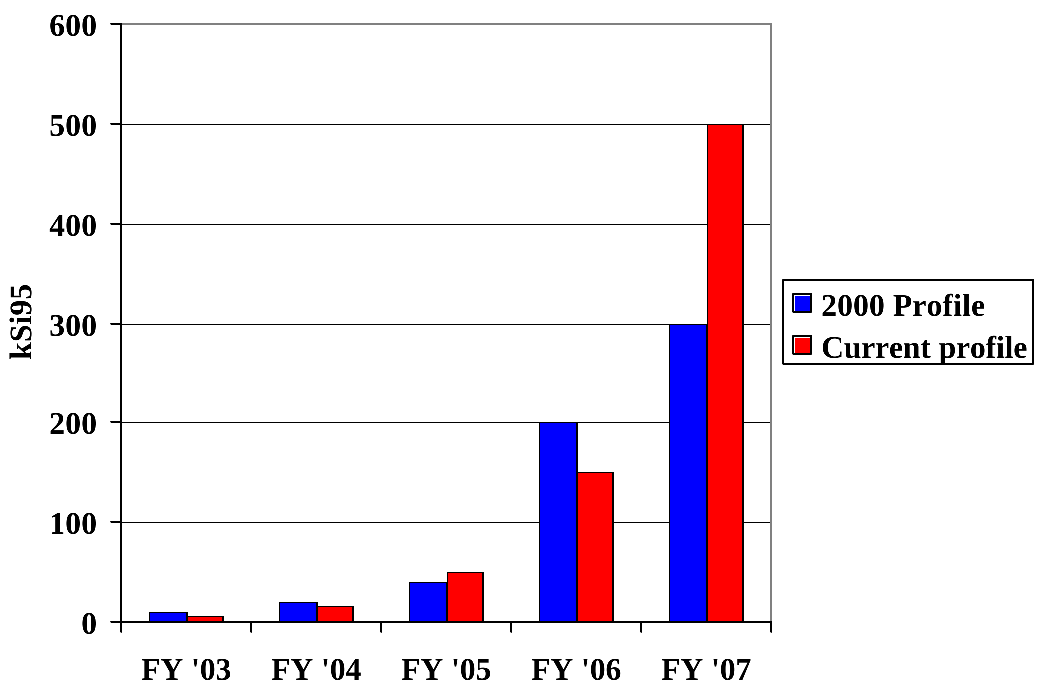

"2007" Capacities for U.S. Tier 1 Options

	Tape Based Model	3 Tier 1 Disk Model	Standalone Disk Model
CPU (SPECint95)	209	329	500
Disk (TBytes)	365	483	1000
Tape (PBytes)	1.85	1.85	1.85
Disk (GBytes/sec)	18.3	18.3	18.3
Tape (MBytes/sec)	802	185	185
WAN (Mbit/sec)	4610	9864	9864
		1/3+1/6 of ESD on disk	Add other 2/3 of ESD
	ESD pass each month	ESD pass per group each day	

- "3 Tier 1" Model (Complete ESD found on disk of U.S. plus 2 other Tier 1's)
  - Highly dependent on the performance of other Tier 1's and the Grid middleware and network (transatlantic) used to connect them
- "Standalone" Model (Complete ESD on disk of US Tier 1)
  - While avoiding above dependencies, is more expensive

Associated Labor Profile

	FY '01	FY '02	FY '03	FY '04	FY '05	FY '06	FY '07	FY '08
11/00 Projection (FTE's)	5	7	10	15	25	25	25	25
11/01 Projection* (FTE's)	2.7	4.2	6.5	11	16	22	25	25
Labor Cost (@Yr $K)	419	677	1090	1918	2901	4149	4903	5099
Support Costs (@Yr $K)	50	66	91	141	199	271	313	322
Total Cost (@Yr $K)	469	743	1181	2058	3100	4420	5216	5421

* Not including .5 FTE of PPDG in FY '02-'04

Summary Tier 1 Cost Profile
(At Year $K)

	2001	2002	2003	2004	2005	2006	2007	TOTAL	2008
CPU	$ 30	$ -	$ 59	$ 117	$ 305	$ 565	$ 1,316	**$ 2,392**	
Disk	$ 100	$ -	$ 118	$ 263	$ 564	$ 1,058	$ 2,446	**$ 4,549**	
Tertiary Storage	$ 55	$ 6	$ 45	$ 140	$ 120	$ 225	$ 305	**$ 896**	
LAN	$ 79	$ -	$ 20	$ 20	$ 90	$ 100	$ 250	**$ 559**	
Other Infrastructure	$ 40	$ -	$ 11	$ 26	$ 53	$ 90	$ 207	**$ 427**	
Sftwr, Lic. & Maint.	$ 50	$ 89	$ 128	$ 165	$ 215	$ 307	$ 443	**$ 1,398**	
Overhead	$ 35	$ 19	$ 47	$ 80	$ 136	$ 228	$ 455	**$ 999**	
Hardware	$ 389	$ 114	$ 428	$ 811	$ 1,484	$ 2,573	$ 5,422	**$ 11,220**	$ 2,572
Labor	$ 469	$ 743	$ 1,181	$ 2,058	$ 3,100	$ 4,420	$ 5,216	**$ 17,187**	$ 5,421
Total	**$ 857**	**$ 857**	**$ 1,609**	**$ 2,869**	**$ 4,584**	**$ 6,992**	**$ 10,638**	**$ 28,407**	**$ 7,993**
Guidance	$ 855	$ 839	$ 1,600	$ 2,500	$ 4,600	$ 7,000	$ 10,700	**$ 28,094**	$ 8,000

- Current plan violated guidance by $370k in FY ’04, but this is a year of some flexibility in guidance
- Strict adherence to FY ’04 guidance would …
  - reduce facility capacity from 3% to 1.5% or staff by 2 FTE’s

Tier 1 Capacity Profile

	2001	2002	2003	2004	2005	2006	2007
CPU (SPECint95)	3	3	6	15	50	150	500
Disk (TBytes)	2	2	8	30	100	300	1,000
Disk (MBytes/sec)	40	40	200	600	2,000	6,000	20,000
Tape (PBytes)	0.01	0.02	0.05	0.09	0.15	0.65	1.85
Tape (MBytes/sec)	10	10	20	20	48	106	212
WAN (Mbits/sec)	155	155	622	622	2488	9952	9952

Tier 1 Capacity Profile

The table below shows the breakout of funds by activity

Budget Profile by Item

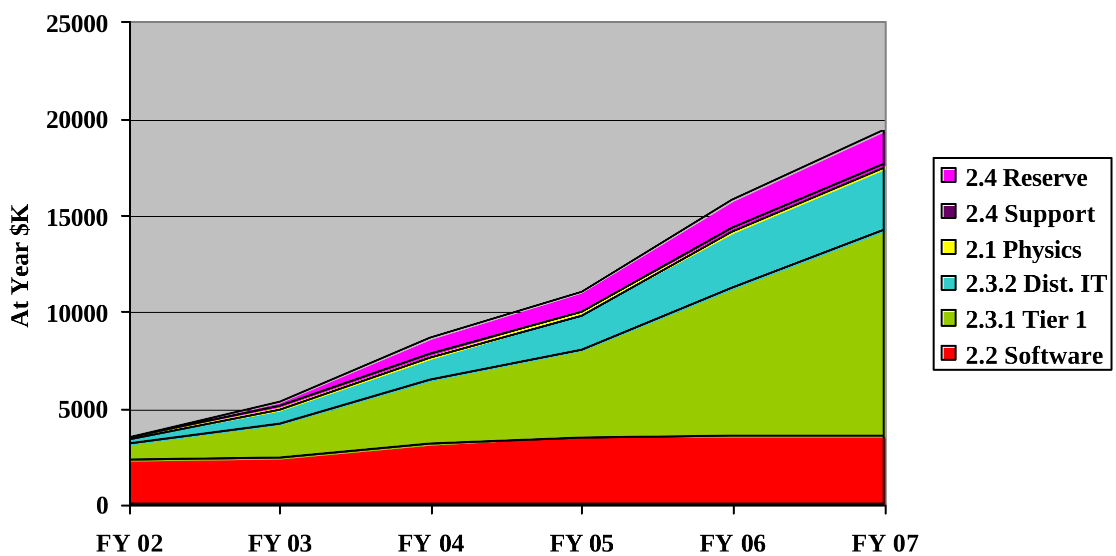

The following table shows the same information in tabular form.

WBS Number	Description	FY 02	FY 03	FY 04	FY05	FY06	Total	FY07
		Fiscal Years AY k$'s						
2	US Atlas Computing	3,581	5,328	8,201	10,123	14,457	41,690	17,755
2.1	Physics	100	147	196	210	215	868	215
2.2	Software Projects	2252	2400	3043	3446	3547	14688	3500
2.3	Computing Facilities							
2.3.1	Tier 1	839	1701	3392	4467	7575	17974	10615
2.3.2	Distributed IT	290	780	1120	1850	2970	7010	3265
2.9	Project Support	100	300	450	150	150	1150	160
	Management Reserve	0	250	820	1,012	1,446	3528	1,776
	US ATLAS Computing w/reserve	3,581	5,578	9,021	11,135	15,903	45,218	19,531

The following table shows the funding request to the NSF.

	FY 02	FY 03	FY 04	FY 05	FY 06
Physics	124	150	200	200	200
Software	820	844	870	896	925
Tier 2					
Local staff	542	620	653	697	930
Central staff	0	0	0	300	465
Hardware	240	240	870	1870	2000
Reserve	259	278	389	594	678
Sum	1985	2132	2982	4257	4733
Related Projects					
iVDGL	403	532	550	449	457
GriPhyN	139	139	139	0	0
Grid Telemetry	167	167	167		
Sum Related	709	838	856	449	457
Total	2694	2970	3838	4706	5190

US CMS

CPT Level 1 Milestones

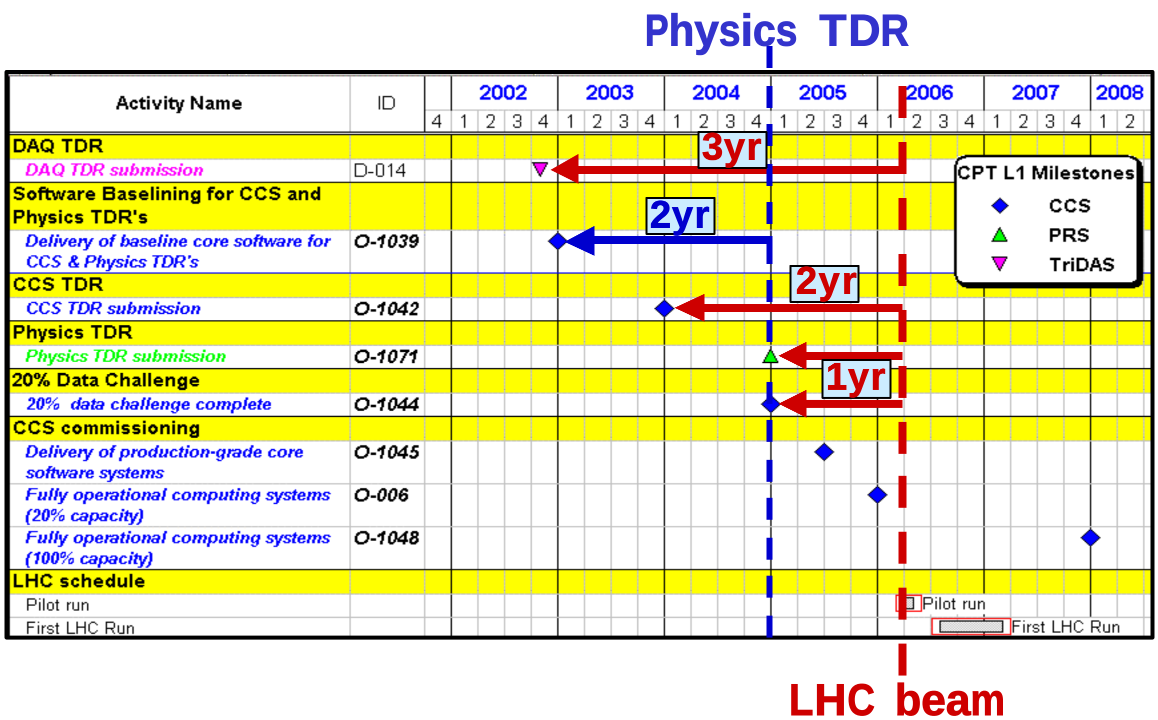

Major Milestones

1 Tbyte database prototype	1/1/99	(Done)
Release of Athena pre-alpha version	5/9/00	(Done)
Tier 1 Processor Farm Prototype	9/29/00	(Done)
Athena alpha release	9/29/00	(Done)
Geant3 digi data available	10/30/00	(Done)
Athena Beta Release	12/29/00	(Done)
First definition of regional centers	1/1/01	(Done)
Decide on database product	6/29/01	Delay
Tier 1 Storage Prototype	10/1/01	Part
MDC0 Completed	12/12/01	
Full validation of G4 physics	12/31/01	Delay
MDC 1 Completed	7/30/02	
Computing TDR Finished	11/29/02	
Tier 1 Upgrade (for MDC2)	12/31/02	
Tier 1 Large Scale Test (MDC2)	9/30/03	
Physics readiness report completed	6/30/04	
Full software chain in real envioron.	7/30/04	
Full DB infrastructure available	12/31/04	
20% Processing Farm Prototype	9/30/05	Delay
Tier 1 Full scale	10/2/06	Delay

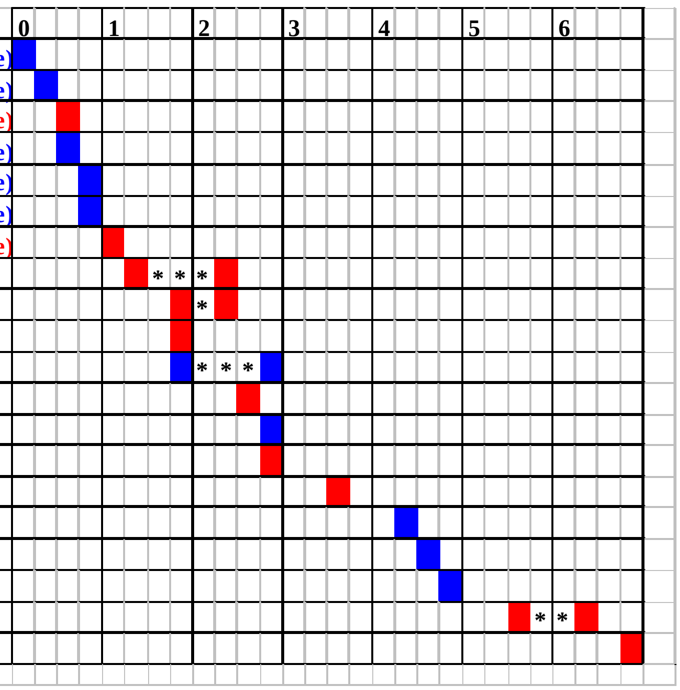

Appendix F – Organization Charts

U.S. CMS

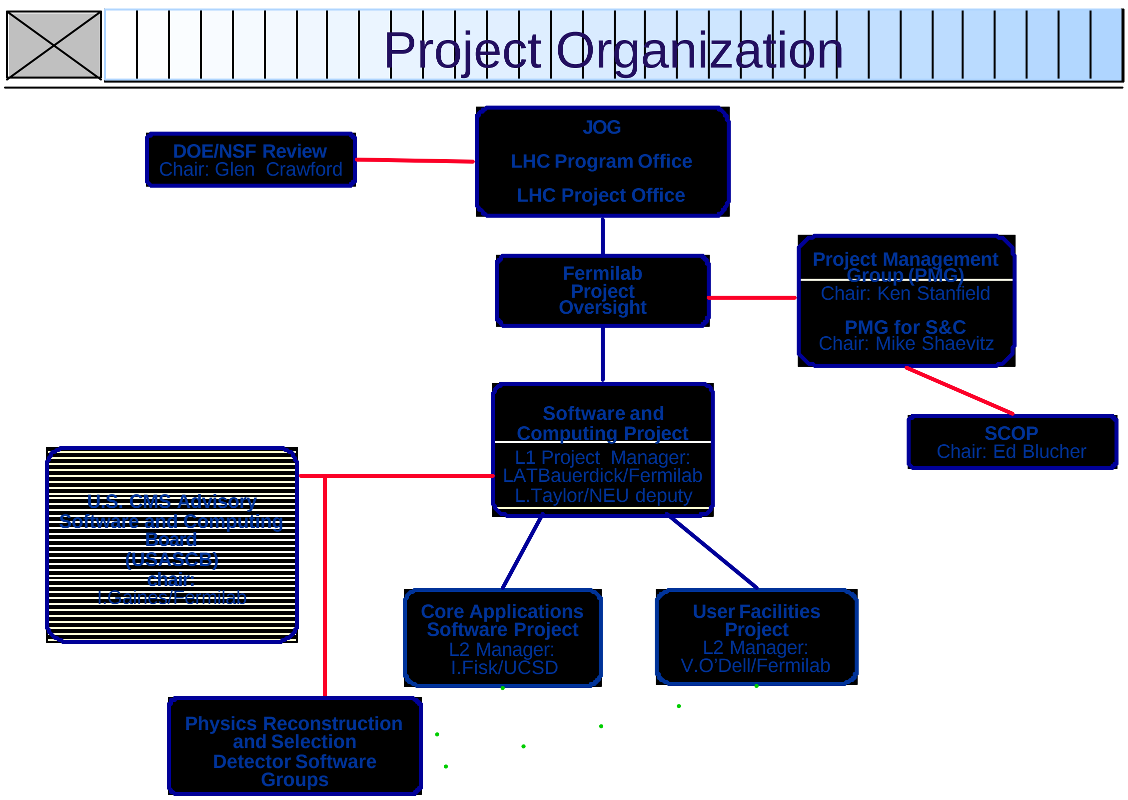

U.S. ATLAS

U.S. ATLAS Project