

DOE/NSF U.S. LHC Program Office

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

Brookhaven National Lab
November 14-17, 2000

Executive Summary

CMS and ATLAS will be large, general-purpose detectors used to observe 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.

A peer review of the U.S. LHC software and computing efforts was held on November 14-17, 2000, at the Brookhaven National Lab in Upton, NY. The primary purposes of this review were to evaluate the technical plans of the collaborations for the near-term (defined as FY2001-2002), to assess their estimates of cost and schedule, to review their management structures, and to help set priorities for near-term funding allocations. The expert reviewers provided comments during the review, both to the U.S. LHC collaborations and to the DOE and the National Science Foundation (NSF). These comments, and those provided in writing from the reviewers, form the basis of this report.

The two U.S. collaborations have each proposed funding levels totaling \$3 million in FY2001 and \$4 million in FY2002 for their software and computing projects, based on funding guidance received from the agencies. They both noted that this level of effort is below what they feel is adequate in order to meet the current LHC schedule. Support at this level has a non-negligible impact on the rest of U.S. high-energy physics (HEP) program, which is already strapped for resources. Still, this funding level represents only ~1% of the total current U.S. HEP program and will be a crucial component in supporting the research for roughly one third of the U.S. HEP community in the future. Members of the review committee expressed general support for this level-of-effort, and urged the funding agencies to do their best to meet these requests.

The committee found that considerable progress has been made on software and computing by the U.S. LHC collaborations since the previous review of this activity in January 2000. The overall strategies of both the U.S. ATLAS and U.S. CMS software and computing projects were found to be sound. The U.S. groups have a strong competence in the areas of software and computing and they are making significant contributions, leading the international effort in several key areas.

The committee applauded both the U.S. ATLAS and U.S. CMS groups for taking leading roles in important areas of software that will be used by the full collaborations. Both groups now have a core of experienced software professionals working in teams to develop vital infrastructure. The need for maintaining (and in some areas,

expanding) this core effort to deliver the planned US software contributions was identified as the highest priority. At the same time, the reviewers cautioned the US project managers to carefully delineate the scope of their contributions through MOUs or other tools to avoid undue expectations about the level of US effort from their international partners.

An important feature of the LHC Computing model is its reliance on software that can run in a distributed environment, including plans for a eventual “Data Grid” that would enable data and resource discovery, allocation and brokering on an unprecedented scale. If this vision is achieved, it will have a significant impact not only on HEP computing but also on the IT field in general. The promise of “grid computing” is sufficiently high that the proposal from the GriPhyN collaboration (a consortium of computer scientists and physicists, including CMS and ATLAS representatives) won a recent major NSF IT competition. At the review, we heard of the initial efforts to organize and align the deliverables of GriPhyN with those of the US LHC computing projects. The committee recognized both the potential benefits GriPhyN could bring to these projects as well as the significant management challenges involved in doing so, and encouraged the managers to structure their project plans accordingly.

Present plans call for the U.S. CMS and U.S. ATLAS national computing facilities to be located at the Fermi and Brookhaven National Laboratories, respectively, as part of the distributed LHC Computing model. In this model, major countries participating in LHC physics would have a national (“Tier 1”) center that is the primary repository for data received from CERN. In the U.S., both collaborations envision a second tier of regional computing centers which provide both local high-speed access to a reduced set of the data and significant computing resources. In this report we refer generically to both kinds of computing centers as “user facilities.”

Both Tier 1 plans were found to be detailed and well conceived. Both experiments were encouraged to accelerate staffing for these facilities and procurement of equipment to develop these sites for use by the U.S. LHC community. However, policies and planning for the “Tier 2” regional centers are still in a preliminary phase; early prototypes and testbed facilities are being deployed or proposed. Continuing uncertainties include the detailed technical requirements for these centers, the site selection process, and the sources of funding. The committee recommended that the projects refrain from large investments in this area until these questions are resolved.

The two U.S. efforts employ somewhat different management structures to carry out their software and computing efforts, but both appear to be working effectively. Both collaborations are beginning to coordinate their work with their respective Project Offices that were created for the detector construction projects, and plan to produce regular progress reports for the agencies. Detailed resource-loaded schedules and cost books were presented. Full integration of project milestones into the work schedules and coordination with international project planning still need to be done. Drafts of the U.S. CMS and U.S. ATLAS Project Management Plans (PMPs) are well along, but will need some further refinement over the coming months.

Given the current uncertainties in funding levels, the committee was concerned about impact on the overall project scope and schedule. Both project managers stated that they would adhere to a “build-to-cost” model where the scope of the software and hardware deliverables becomes the contingency in a project with fixed cost and modest schedule float. There was general agreement that, while not ideal, this model was appropriate and adequate for these large software and computing projects.

The near-term activities proposed by the two U.S. collaborations seem reasonable, and the U.S. LHC software and computing efforts will need a robust ramp-up in the near term in order to be ready for physics data taking in 2006. There appears to be a serious need in FY2001 and FY2002 for more resources than the Agencies can currently provide for the U.S. LHC software and computing projects. The collaborations and their project managers should carefully consider how best to use their flexibility in the near-term to optimize the overall U.S. LHC efforts.

Table of Contents

1 INTRODUCTION	1
2 PROGRAM OVERVIEW	2
3 CORE SOFTWARE.....	4
US ATLAS – FINDINGS AND EVALUATION.....	4
US ATLAS -- RECOMMENDATIONS	5
US CMS – FINDINGS AND EVALUATION	6
US CMS -- RECOMMENDATIONS	7
4 USER FACILITIES.....	10
US ATLAS -- SUMMARY.....	10
US ATLAS -- FINDINGS AND EVALUATION	11
US ATLAS -- RECOMMENDATIONS	13
US CMS – SUMMARY.....	14
US CMS – FINDINGS AND EVALUATION	14
US CMS -- RECOMMENDATIONS	17
5 PROJECT MANAGEMENT	18
US ATLAS – SUMMARY	18
US ATLAS – FINDINGS AND EVALUATION.....	19
US ATLAS – RECOMMENDATIONS.....	24
US CMS – SUMMARY.....	25
US CMS – FINDINGS AND EVALUATION	26
US CMS -- RECOMMENDATIONS	30
6 ACTION ITEMS	31
7 APPENDICES.....	32
APPENDIX A – CHARGE TO COMMITTEE	32
APPENDIX B – MEMBERS OF REVIEW COMMITTEE	34
APPENDIX C – REVIEW AGENDA	35
APPENDIX D – COST TABLES.....	39
APPENDIX E – SCHEDULES.....	42
APPENDIX F – ORGANIZATION CHARTS	45

1 Introduction

This report is the product of the joint DOE/NSF review of U.S. LHC Computing Projects held at Brookhaven National Lab on November 14-17, 2000. This review was charged with examining the technical scope, cost, schedule, and project management of these efforts, focussing on the near-term (through FY2002) plans of both collaborations in developing software and user facilities for the LHC experiments ATLAS and CMS that are scheduled to begin taking physics data in 2006.

A team of nine outside experts reviewed detailed presentations made by both collaborations on their individual projects as well as common projects that promise to deliver software products useable by both experiments. Their evaluations are contained in this report. At the review, they provided many recommendations to both collaborations and the agencies at the closeout.

Many 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 considerable assistance from: Pepin Carolan, Dan Lehman, Mike Procaro, Tim Toohig, Kathy Turner, Vicky White, Jim Yeck (DOE); Alex Firestone and Marv Goldberg (NSF); Andreene Witt (ORISE), Jackie Mooney, Linda Feierabend, and Vanessa Langhorn (BNL) provided invaluable local support.

The charge given 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 on different days were made for the U.S. ATLAS and U.S. CMS computing efforts, and a half-day was devoted to common projects. Cost tables for both efforts for FY2001-2002 are given in Appendix D, and milestones and work schedules in Appendix E. Organization charts are collected in Appendix F.

This report and its recommendations represent 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 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 used to observe 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 2005. 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 decided that it is now time for an assessment and 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 will be 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 first 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.



3 Core Software

US ATLAS – Findings and Evaluation

There has been an impressive amount of progress within ATLAS since the last DOE/NSF review in January. This was very evident from the US-ATLAS WBS for software development, which now has a web accessible interface. This web interface has made it much easier for the management to track the software projects. The US-ATLAS WBS has a substantial amount of detail through 2001 with milestones (31 at a high level) and deliverables. We were told that the US-ATLAS planning efforts are the most advanced and sophisticated within the experiment.

US-ATLAS has two core software responsibilities which build on their expertise. These are the control framework and overall offline software architecture along with the database and data management. In addition, there is subsystem reconstruction and simulation code which will be done by the physicists and which will require some core support.

During 2000 the framework group has delivered the ATHENA package, and has met almost all of their milestones for the year. They seem well poised to meet those for 2001 as long as they preserve their critical mass of developers in this area. ATHENA has had good acceptance within ATLAS and has already been adopted by many ongoing efforts. The group has given several tutorials on this framework and has an active developer community. ATHENA is based on the LHCb GAUDI package and a good collaboration has been established with LHCb.

The second core software responsibility is the data management and database effort. The US has expertise in this area that ATLAS wants to use but this effort is substantially underfunded resulting in a shortage of positions. The effort has been further complicated by the request from ATLAS to also support test beam data and to continue the investigation of alternate persistency packages. This critical core project clearly needs more people to succeed and needs to concentrate on a single database target at least in the near term. They also need to give the test beam support load to other ATLAS workers. At this point, we should point out that US-ATLAS is being asked to produce greater than their 20% share of the overall core software. The US team has the expertise and wants to do the work and it would seem that the project management should be allowed the flexibility to make this happen.

In addition to these core efforts, there is a real need to support the US reconstruction and physics analysis efforts with two professionals to help with the transition to OO and C++ and the interfacing to the new framework. One of these would work with the physics generator group and the other as a general consultant with the subsystem developers.

The committee would like to reiterate that the ATLAS briefings were very effective in presenting the status of current software efforts, enabling the review committee to rapidly address critical issues. Furthermore, the committee obtained a strong sense of coherence between participants in the ATLAS core software effort. There was an overall sense that the work is truly being driven from an integrated development plan.

US ATLAS – Recommendations

- Although it is laudable that the U.S. Atlas software effort is moving to an object-oriented paradigm, it will be a large challenge to shift to object-oriented technology while concurrently delivering on the specified milestones. The training in C++, OO, and packages is good investment. **However, the overall cost of developers and users shifting to an object-oriented paradigm needs to be tracked and rolled into cost estimates.**
- There was very little presentation of software quality assurance techniques and tools. Although some techniques and tools are utilized in parts of the project, there was no presentation of an overall coherent software quality assurance plan. **A Software Quality Assurance plan needs to be written and a set of key SQA processes needs to be implemented immediately. This must include metrics to track progress against Software Quality targets.**
- Although there were many discussions of plans, roadmaps, and milestones, there were few examples of working software. Given the scope and urgency of upcoming milestones and the time necessary for adoption of tools in the community, **it is advised that the ATLAS software effort adopt a software development process which allows for intermediate rapid prototyping so some capabilities are made evident to the community and feedback may be obtained.**
- We heard a need from physics for a generators expert and a C++/OO subsystems guru to help with architecture and design (located at CERN). Since this is an investment that will have a large impact on all future development, the committee recommends that the hiring of this expert be made a priority.

- The committee would like the ATLAS core software program management to execute a continuing assessment of the correct balance between frameworks and database efforts – especially in the light of possible shortfalls.
- The committee urges the project management to evaluate the overall impact of grid-enabling the ATLAS software environment and assess the appropriate ramp of grid-deployment activities relative to other priorities in the program.

US CMS – Findings and Evaluation

We heard presentations on the CMS core software projects identified in the WBS as 2.1 through 2.4. These were software architecture (2.1), Interactive graphical user analysis (2.2), distributed data management and processing (2.3) and support (2.4). US resources have gone into the functional prototypes which are now done and which are now being evaluated for future work. The US responsibility to the overall software effort is defined by a specific level of effort expressed in the project overview document and which is “sized” to be “constrained to a level-of-effort contribution of 25%”. There is considerable US leadership within the software effort. There is also a software support requirement specifically for US-CMS physicists.

The US-CMS software effort is covered by a WBS that is distinct from the CMS WBS. This effort currently involves 8 software engineers which is the appropriate fraction (25%) of the total projected CMS need even though there is a 10 FTE shortfall for the experiment overall. The US has recognized the up front need for these engineers and is ahead of the rest of the experiment in this regard.

The software architecture effort (WBS 2.1) has documented the current software architecture and a set of tools to manage documentation within CMS. The Interactive graphical user analysis project (WBS 2.2) provides a general set of visualization tools with emphasis on modularity and standards. The IGUANA functional prototype has been successfully completed and is now being documented. It is currently capable of displaying the full GEANT3 CMS detector and many of the reconstructed objects from ORCA 4.3.0. Plotting and statistical packages are also accessible within the package. There is a fully versioned set of all the IGUANA code and the related documentation that includes examples and tutorials. The documentation system is based on the public domain DOXYGEN and perl scripts.

The CMS distributed computing model is indeed complex and many tools will be needed to make it work. Much prototype work has already been done in this area (WBS 2.3). A prototype now exists for the distributed process management as well as the distributed database management. There are production support tools which take advantage of existing production facilities and tools to archive and transfer data between centers world wide including a functional database replication prototype.

These tools are being used for the current ORCA production. Based on this database replication prototype tools have been designed to automatically record and archive results of production at remote sites and transferred to the CERN mass storage system. There already exist prototype tools to monitor these production facilities.

The committee commends the CMS core software effort for its early adoption of object-oriented technologies. With many deliverables now in flight, the arduous task of learning a new paradigm has been accomplished so researchers can focus on development. Actual prototype and code results were also demonstrated - in particular for IGUANA. This gives the committee confidence that good progress is being made against project milestones.

US CMS – Recommendations

Project Planning

Although the overall US CMS software project efforts appear to be in overall good shape, there is concern as to how the project planning effort is being utilized to coordinate and direct activity. Although the committee is very impressed with the project progress that has been accomplished to date, to the committee, it appears that the project planning efforts are being used to communicate efforts rather than coordinate and direct efforts.

- For this effort to succeed more leadership must be exerted by project management and efforts need to be accountable to the priorities and needs specified by an integrated project plan.
- The project needs to tighten the WBS milestones and deliverables in terms of what can be controlled by the US project management.
- The prioritization of tasks does not appear correlated with the proposed budget reductions. Although a very detailed budget breakdown could be articulated for a proposed reduction, the prioritization needs to reflect similar detail and have better correlation.
- The usage of software engineers is ahead of the overall CMS curve – “mission creep” must be contained.

Distributed Computing

Although a great deal of work has been accomplished, the distributed computing components of the WBS are amorphous in scope and the actual deliverables are not clear. The collaboration relies on many external network grid-oriented, distributed computing initiatives. This technology might greatly ease the complexity of managing the distributed computing challenges articulated by the CMS researchers. However, most of the grid computing efforts appear to be moving forward in directions that are not directly accountable to the CMS project management. This is not necessarily a problem, but could become so if one or more of these efforts falls short of delivering useful results in some critical area.

A prototype for distributed processing and database management has been delivered along with monitoring tools – this is a significant delivery. The project is beginning to assess load balancing and automation strategies. Utilization of modeling tools (MONARC) will play an important role in balancing the overall distributed system architecture.

- The committee encourages the continued interaction and collaboration with the Grid Community and looks forward to a successful integration of Grid technology into the CMS distributed computing architecture.
- It is important for CMS to understand the “minimal” functionality in each area of distributed computing to have specific milestones or trigger points at which the situation in external programs will be evaluated, and to have a reserve contingency to explore or in worse case to develop options where the minimal functionality will not be realized. Of course, if the collaboration has the capability to pursue backup strategies for essential functions, this will reduce risk even further.
- It was clarified that the 50% of overall Distributed Computing responsibility staffed with 3 FTEs is consistent with the difficulty of this task since most of the effort is being accomplished with off-project resources. This raises a serious concern as to deliverable dependencies which are not accountable to CMS Core Software program management

Software Quality Assurance

The Software Quality Assurance efforts are making good progress with many important processes and tools in place.

- The committee would encourage the further development of processes (e.g. daily regression tests) and the integration of tools into an overall capability (e.g., regression testing integrated with bug tracking and version control).

- In March, it was recommended that a general QA/QC plan for software with well-defined metrics be formulated. Although good progress has been made in SQA processes, this overall plan has yet to materialize.

IGUANA

The development of IGUANA appears solid in scope and in its relation to other software components (e.g., from CERN, HEP, and commercial). In order to gain wide acceptance, software products such as Iguana have to be easily accessible at each collaborating institution. These products involve several commercial programs, which may have complicated licensing issues, along with internally developed and freeware components. It is critical to its acceptance that this product is very easy to deploy throughout the collaboration. Licensing arrangements should be straightforward so individual physicists can deal with them with a minimum of effort. Other code should be packaged efficiently and be provided in a way that minimizes each user's effort. This concern was highlighted by some US problems associated with the distributed of LHC++, which continues to be a part of IGUANA's infrastructure.

- For the commercial products, it is crucial that the appropriate licenses at the correct (i.e. supported) version of supported operating systems be available and affordable by all collaborators with each new release.
- It has also been demonstrated that users will reject an analysis product that does not have assured support throughout the project's lifetime. It is essential that the CMS collaboration make a commitment to support this product throughout its useful life.
- To justify the IGUANA effort, it needs to be shown how the rest of the experiment is on track to use the technology and the strategy for deployability and long-term maintainability (productization) of IGUANA made clear.

4 User Facilities

US ATLAS – Summary

The User Facilities review subcommittee met with representatives of US-ATLAS Software and Computing on the afternoon of November 16 at Brookhaven National Laboratory. The discussion of the US-ATLAS facilities components of the project proceeded in a cordial and informal manner. The reviewers appreciated the openness and straightforwardness of the presenters in response to the committee's questions throughout the discussions.

The committee noted language from the draft US-ATLAS Computing Project Management Plan stating:

“The goals of the facilities subproject [are] to provide the basis for the support of U.S. ATLAS physicists in the analysis of data from the ATLAS experiment, and to carry out specific computing tasks for the International ATLAS experiment as per agreement between the two.”

This language served as a guidepost to evaluate whether or not the overall progress of the user facilities subproject was moving in a direction with sufficient funding and on a time-scale to meet these goals.

Although the subcommittee reviewed the overall assessment of the US-ATLAS User Facilities, the group focused on a 12 to 24 month time period from the present when examining proposed detailed plans, costs estimates, and schedules. The report is divided into sections on technical scope, cost, schedule, and management. Within each section, there are findings and evaluation. Recommendations are collected at the end.

The overall summary is that at this time, the committee believes this subproject is on track to meet the stated user facility goals.

US ATLAS -- Findings and Evaluation

Technical Scope

The beginnings of a tier 1 facility are currently operating at BNL, serving a portion of the computational requirements of US-ATLAS. The facility is deployed adjacent to the RHIC Computing Facility (RCF), and is sharing / leveraging a certain amount of RCF hardware, including robotic storage, RAID controller, and networking. The staffing level is currently two FTEs, paid by US-ATLAS. Plans call for staff levels to increase from two to five FTEs in FY01, and then to seven FTEs in FY02.

The project would like to procure sufficient hardware in FY01 to increase computing by roughly a factor of two and de-couple the tier 1 hardware from the RCF. It is also planned to deploy a prototype tier 2 center in FY01, and a second prototype tier 2 center in FY02. These two tier 2 prototype centers would take part in the 5% mock data challenge in 2003.

Several GRID R&D projects aimed at producing the middleware for the multi tier network were part of the presentations in the context of the user facilities. While the significance of related products was not explicitly stated, such projects should be followed with moderate and justifiable effort. Resources for grid and prototype tier 2 development must be balanced with other US-ATLAS computing needs. Care should be taken to not compromise the plan strictly related to the US part of the ATLAS experiment.

Long Range Plans

U.S. Atlas plans to double tier 1 computing and storage by a factor of two in each of '02, '03, and '04, at which point it will be at 20% of requirements. The remaining 80% will be purchased in '05, in time for '06 running.

Staffing is planned to ramp up smoothly from five FTE in '01 to 25 FTE in '05, so that all staff are on board well ahead of running in '06.

In response to committee questions about priorities in view of budget shortfalls, plans were presented to slip many tier 1 procurements in '01 and '02 by approximately 1 year, putting pressure on the '03 budget (just in time for the mock data challenge).

The short and long range plans for hardware and manpower for tier 1 are all reasonable. Dedicated acquisition for US-ATLAS computing at BNL include 3 kSI95 in 2001 to be ramped up to 4 kSI95 in 2002, 4 TB disk storage in 2001 to be ramped up to 7 TB in 2002 and 20 TB tape storage in 2001 to be doubled in 2002 representing about 1.5% of the total capacity to be installed in 2006. Staffing plans

(absent triage due to budget constraints) are viewed by the committee as reasonable.

The plan for late procurement of the bulk of tier 1 hardware in '05 is to be commended as it achieves the best price/performance, and maps well to the funding profile. In addition, because technological advancement in the next few years has large uncertainties, we endorse this approach, which is based on a scalable design. In case of a budget shortfall, the effort can be scaled back to "build to cost".

We strongly agree to US-ATLAS planning to build a prototype Regional Center in FY01 and FY02 as a proof of concept and for testing. The projected size of ~1.5% of the full tier 1 size in FY 01 should be sufficient to achieve relevant results. Tier 1 plans should be re-visited prior to the FY 04 time frame when the tier 1 facility is planned to be at ~10% of full scale and support of the growing tier 2 network demands will have to be addressed. This checkpoint will be an excellent opportunity for verification of assuming

- ◆ commodity computing will satisfy the demands when running in full production
- ◆ costs for commodity computing will fall according to the assumed projections

According to US-ATLAS computing facilities planning, the majority of simulated events will be produced at the tier 2 sites. Only a small fraction of the foreseen CPU resources at tier1 (9 kSI95 out of 209 kSI95) will be devoted to MC production.

Given the fact that the tier 2 architecture is subject to the results of a yet to be completed R&D project leads the committee to recommend devoting a larger fraction of the installed CPU capacity – especially in the ramp-up phase before 2006- to the simulation effort.

The deferrals of tier 1 hardware (HPSS and some disk) and staffing (HPSS) in `01 is viewed as acceptable and should not result in tier 1 becoming a critical path item. Nevertheless, staffing (two FTE) is currently too tight to achieve tier 1 missions.

Tier 2 hardware plans for `01 and `02 are larger than needed for the research and development of distributed computing for US-ATLAS. To the extent that this hardware does not come from DOE and NSF guidance funds for these years, this is not a problem.

Cost

Hardware:

RHIC Computing Facility experience has been used to develop cost estimates for all hardware.

Personnel:

RCF experience has also been used to develop cost estimates. Staffing is expected to reach 25 FTEs for full operation by the end of FY05.

All cost estimates are well done. Cost estimates for computing components according to Moore's law is the best one can do today. The collaboration, however, has to make sure, all required components, e.g. infrastructure, are included. Manpower estimates of 25 FTE are roughly 20% lower than similar estimates for CMS, and should therefore not be considered over-provisioned.

Schedule

We strongly endorse US-ATLAS' approach to push back significant production tier 1 hardware procurement to 2005, and believe doing so will have no serious negative impact upon the project.

Management

As an action item from the January 2000 review, a US-ATLAS User Facility Director has been hired.

The selection process for the two initial tier 2 R&D sites, and the final production tier 2 sites is well stated.

Bruce Gibbard brings valuable expertise to this project related to his management of the RHIC Computing Facility. The committee also welcomes Rich Baker joining the project as the US-ATLAS User Facility Director. The committee is pleased to see that Rich has rapidly familiarized himself with the project and is fully engaged in the US-ATLAS user facilities at this time.

US ATLAS – Recommendations

- Build a prototype Regional Center in FY01 and FY02 as a proof of concept and for testing. The projected size of ~1.5% of the full tier 1 size in FY 01 should be sufficient to achieve relevant results.
- In the case of budget shortfall, defer HPSS hardware and staffing from `01 to `02, and to defer additional CPU, disk, and other hardware from `02 to early `03 (and similarly, corresponding staffing).
- Increase the staff at tier 1 from 2 to 4 FTE as soon as possible this fiscal year to better support US-ATLAS physicists and to take part in distributed computing R&D.

- Rich Baker must have additional planning support in FY01, as part of the increase to a minimum of four FTE at tier 1.
- A description of the requirements (including infrastructure) for the tier 2 sites should be developed and published as soon as possible. The project management plan should be expanded to include a plan for tier 2 selection and operation.

US CMS – Summary

The User Facilities review sub-committee met with representatives of US-CMS on the afternoon of November 14 at the Brookhaven National Laboratory. The discussion of the US-CMS facilities components of the project proceeded in a cordial and informal manner

The mission of the user facility sub-project is to provide the enabling infrastructure to permit the CMS collaborators to fully participate in the physics program of CMS from their home institutions. The user facility is necessary because CERN has stated as policy that a large fraction of LHC computing must be done by the participating institutes and laboratories.

Although the subcommittee reviewed the overall assessment of the US-CMS User Facilities, the group focused on a 12 to 24 month time period from the present when examining proposed detailed plans, costs estimates, and schedules. The report is divided into sections on technical scope, cost, schedule, and management. Within each section, there are findings and evaluation. Recommendations are collected at the end.

At this time, the committee believes that the overall user facilities subproject is on track to meet the stated user facility goals for US-CMS.

US CMS – Findings and Evaluation

Technical Scope

A modest amount of hardware is currently deployed at FNAL as a tier 1 prototype. The facility is currently supporting US-CMS groups performing higher-level trigger studies and test beam studies. Software developers are also using the facility for testing distributed systems.

Hardware:

For this stage of the project, US-CMS has a good handle on the parameters driving the hardware of the multi-tier system such as:

- raw data volume
- processing power per event

- storage requirements
- networking requirements from tier 0 to tier 1 centers (not quite as well defined)
- cost and size extrapolations (like Moore's Law) of processing power, storage, etc. to the final system

Personnel:

The top-down and bottom-up evaluations of personnel requirements have been completed and compared to each other.

The review committee feels that the US-CMS collaboration has made good progress since the January 2000 review on the design requirements for a multi-tier architecture for data analysis and simulation.

The top-down and bottom-up evaluations of personnel requirements are consistent and credible.

Certain items significant to the multi-tier architecture are not yet as well defined or understood:

- task distribution between tier 1 and tier 2
- optimal size and number of tier 2 centers
- remote object de-referencing patterns, and hence networking requirements between tier 1 and tier 2

Some progress has been made in addressing the recommendation of the January 2000 review to "*Continue development of detailed plans for the Tier 2 Centers with a complete model of how these centers will actually operate and service the user community, including details of sizing, cost and functionality*". However, additional work needs to be done to justify the chosen parameters for the tier 2 centers (number, size, etc.).

US-CMS is in the process of deploying a distributed prototype tier 2 center in California, centrally managed and distributed between Caltech and UCSD, with cost sharing from UC Davis. While the reviewers agree to building a prototype in general, it is felt that funding levels and priorities and requirement for other components of the US-CMS activities do not justify the resulting complexity of the Tier 2s.

In light of the current budget constraints, the reviewers feel that smaller centers would have been sufficient for studying several critical issues such as object access patterns and questions of batch or interactive/real-time remote object access. While US-CMS is correct in stating that a large number of systems will be needed to test scalability, the current emphasis should be upon software (manpower) more than hardware.

In contrast, the tier 1 center appears to be too small for FY01 given the anticipated production goals, and needs additional funding to support US physics activities.

Cost

Hardware:

Historical information has been used to develop cost estimates for hardware. These costs are well understood, with ample historical data for backup justification to project future year hardware costs

Personnel:

For FY01 US-CMS is requesting 7.5 FTE in total for tier 1 related effort. The personnel requirements can be categorized as follows:

- 4.0 FTE are currently working for tier 1 and are supported by US-CMS
- 2.0 FTE are foreseen in the project scope for various WBS items, e.g. System and User Support, Operations and Infrastructure, Networking etc, which are not covered by the above mentioned 4 FTE. These two FTE are currently contributed by Fermilab CD out of the base program.
- 1.5 FTE are missing in the area of technology investigation and deployment (1 FTE) and for data import/export (0.5 FTE). Those 1.5 FTE are not covered, by project funding or by FNAL contributions.

The review committee agrees with the hardware cost estimates presented, with only a minor qualification that the SMP costs may be overestimated in view of anticipated industry developments (this is a small effect compared to extrapolation uncertainties).

As is already planned, production level hardware should be procured at the latest possible time to achieve best price/performance, consistent with insuring that the facilities will be operational on schedule.

Assuming FNAL's current contributions of manpower, the reviewers feel that there exists a minimum level of personnel for user facilities to function at this time. As noted in the findings, there are personnel shortfalls as listed. If the current staffing plans remain fixed or if there are no other alternatives for funding personnel, then the continuation of these FNAL contributions are essential for the project to move forward over the next 12 months.

Schedule

US-CMS currently plans to expand the Tier 1 center at FNAL in FY01 and FY02, with particular emphasis upon supporting computing & software R&D (WBS 1.6) and construction activities (WBS 1.7).

Two significantly sized tier 2 prototype centers (called a single prototype, but existing at two locations) are currently being deployed to support studies of tier 2 activities. An additional tier 2 center is proposed for FY01.

The build-up of the tier 1 center is necessary to support the expanding US-CMS activities. The hardware profile presented for the Tier 1 as compared to the hardware for the multiple tier 2 sites seems somewhat unbalanced.

It is unclear from the material presented why there is a need for the deployment in FY01 of a 6-terabyte disk array, and an additional 12 terabytes in FY02. If these procurements could be scaled back somewhat (deferred until the 5% data challenge in FY03), some savings could be realized. Similarly, the need for a third, significantly sized tier 2 prototype site in FY01 does not seem to be compelling. Deferring these hardware expenses into FY03 could yield additional savings.

Management

US-CMS has moved to appoint Lothar Bauerdick as the US-CMS L1 project manager.

The User Facility sub-project is well managed. Under the leadership of Matthias Kasemann as the Acting L1 project manager, the US-CMS project has made substantial progress. The committee welcomes the appointment of Lothar Bauerdick as the permanent L1 project manager for US-CMS.

US CMS – Recommendations

- Additional effort is needed to understand the tier 2 centers. In particular, US-CMS should develop the design and specification documents for the functionality of a tier 2 center, motivated by specific requirements and optimizations of physics data analysis. Critical tier 2 related software components should be prototyped to validate and better understand the essential design parameters of these centers.
- Stable funding sources for required US-CMS personnel levels for the next 12 to 24 months need to be identified to assure adequate staffing for current and anticipated/upcoming projects.
- Re-evaluate computing hardware procurements for tier 1 disks and additional tier 2 prototyping, with a goal of realizing savings through delayed acquisitions.

5 Project Management

US ATLAS – Summary

The committee thanks the US ATLAS computing project managers for their clear and candid presentations of the US ATLAS Physics and Computing Project (PCP). The goals of the US ATLAS Physics and Computing Project are to provide the software, computing and support resources needed to enable collaborating US physicists to fully participate in the ATLAS physics program and to contribute to the International ATLAS Computing project.

The organization of the International ATLAS Computing Project is still in the process of being developed but the project covers the areas of framework, architecture, tools and computing facilities for the ATLAS experiment as well as the simulation and reconstruction needs. At present, the International ATLAS Computing Project is in the process of defining the scope, cost and schedule of the project. The US ATLAS PCP is closely coordinated with the International ATLAS Computing Project and is responsible for the software for control & framework, data management and part of the detector reconstruction. The US ATLAS PCP is divided into three principal areas: the support of physics analysis (simulation, event generators etc); software; and facilities in support for US collaborators. The project organization is set up to carry out or oversee the tasks associated with these areas and to coordinate the activities.

The committee was impressed with a substantial progress made on establishing the cost and schedule of the project. The draft Project Management Plan (PMP) has been refined but a section on the overall project funding profiles needs to be included. US ATLAS PCP is structurally sound and will be capable of providing the leadership and technical resources needed to complete the project.

The physics subproject is well defined but “projectized” organizational structure seems to be lacking.

The area related to the user facility subproject is clearly defined and the associated cost estimate and schedule appear to be well understood. The committee endorses the emphasis on the Tier 1 facility over Tier 2 activities during FY01-FY03.

The scope of the software subproject is well defined. A very detailed task list and associated resource list for next 3 years has been compiled. However, a clearly defined set of deliverables for the control/framework and data management software should be agreed to with the International ATLAS Computing Project as soon as possible.

US ATLAS – Findings and Evaluation

Project Scope

The mission is clearly defined in the WBS:

1. To provide software, computing, and support resources to enable US physicists to fully participate in and make significant contributions to the physics program of ATLAS;
2. To contribute to the overall ATLAS computing effort to a degree that is both commensurate with the proportionate scale of the US contributions to the detector construction and is well matched to the expertise of the US physicists specializing in computing.

The scope of the US ATLAS PCP is captured in the Project WBS and includes:

- Physics (WBS 2.1)
 - Support of Monte Carlo event generators, physics simulation, specification of physics aspects of facilities support
 - Coordination of mock data challenges
- Software Projects (WBS 2.2)
 - Core software: control and framework, event model, data management
 - Detector specific subsystems: pixel, TRT, LAR, Tilecal, Muons, Trigger/DAQ
 - Collaborative tools
 - Software support
 - Training
- Facilities (WBS 2.3) - Hardware , networking and software support of U.S. collaborators in data analysis and in computing contributions to the Atlas collaboration
 - Regional center (Tier 1)
 - Distributed IT infrastructure (Tier 2)
- Common projects: ATLAS is involved in several common projects, usually with entities other than CMS. They are attempting to borrow as much code from others as possible as and then customize it to ATLAS requirements as a

method of reducing the need for software development resources. They do seem to be “collaborating” on development rather than just taking software

- Participation in the construction of grid software

The mission is clearly defined.

The technical scope of the project is well defined insofar as the software and physics projects are concerned. The regional center is also well defined. The scale of the Tier2 centers seems to be established at a management level but there are many issues related to its definition and to the understanding and acceptance of that definition throughout US ATLAS. The participation in development of grid software is a relatively new goal and the scope of the effort here and its relation to the International plan was unclear to us.

In the various software areas, we note the following:

- Physics Software: There is need for an additional software professional in a support role starting next FY.
- Control Framework: The US is playing a leading role here with BNL and LBNL committing significant resources and forming a critical mass for the project. The project currently has 5.5 software engineers.
- Subdetector software activities: The goal was to feed support for the core software into the detector subgroups so that they could be used as testbeds for the development.
- Databases and Data management: This is a critical area. ATLAS has not yet made a choice of database management system. The US wishes to take a large role in this and has a strong voice in management and design issues but does not have enough developers to produce the deliverables it should be capable of bringing to the project. It was stated that there was a need for a threshold of three people in databases and that they now had only 1.5.
- Facilities: There is only a small amount of hardware at BNL.
- Software support: There is a full-time software librarian at BNL. This is a very important function and should be maintained.

While US ATLAS seems to be ahead of International ATLAS in planning its software, there does seem to be general agreement between them on the scope of these projects. ATLAS told us that they thought that it would be highly desirable to have grid-like transparency between Tier 1 and Tier 2 centers (and throughout ATLAS Distributed resources world-wide) but that they thought they could live without this for a while (i.e. it is not mandatory). This probably explains why their R&D effort is just beginning, pushed forward by Rob Gardner at Indiana, and the scope of the U.S. commitment is not very well defined.

The participation in the grid software development seems to be a less well-developed and well-defined project for both US ATLAS and International ATLAS.

International Issues

We got a very detailed summary of the relationship between International ATLAS and US ATLAS in matters concerning computing. The Computing Steering Group (CSG) functions as a software policy board, deals with technical issues, and will be interacting with the associate project manager (APM) for computing. The National Computing Board (NCB) deals with resource issues at a high level – at the national level.

Agreements between International ATLAS and US ATLAS computing will begin with Software Agreements (SA) which are akin to detector construction agreements.

There are three types

- 1) Deliverable – e.g. core software
- 2) Level of effort – in fixed FTE's
- 3) Detector or analysis specific – things with no ownership but have “coordinators”

The software Memoranda of Understanding (MOU's) will be signed after the SA's.

There is a schedule for getting these software agreements done which is:

- i. Core software end of 2000;
- ii. Database, March 2001; and
- iii. Hardware about one year from now.

Cost

US ATLAS presented a cost estimate for the project from 2001 to 2006. The costs for 2007 and beyond are considered operating costs. The costs were broken down into four main cost categories (in Then Yr\$): project office (\$0.5M); core software (\$17M); user facilities (\$22M hardware for Tier 1 and Tier 2, \$23M manpower); management reserve (\$6M). The total project cost (2001-2006) is \$66M. They noted the uncertainties grow in the outyears.

The manpower costs for the user facilities were estimated using detailed costs for the various types of manpower adjusted for varying institutional rates. The manpower costs for the core software consist of two types, as defined by two types of software agreements: one with deliverables using an indicative level of effort, and a second with a level of effort defined in terms of a fixed quantity of FTE's. The hardware costs were estimated using cost projections based on observed cost trends. They note that there can be significant uncertainties in the hardware costs. The management contingency (10%) was not based on a risk-based analysis, but is used to provide flexibility to react to problems. The collaboration noted that the funding guidance they received from the DOE and the NSF falls significantly short of the cost estimate, both in the near term and in

the integral. The shortfall for FY01 is \$1.4M, and \$2.3M in FY02. The total shortfall through 2006 is about \$15M. If this shortfall persists in FY01, the US ATLAS APM has prioritized the request to nearly match the profile. The collaboration states that they will, however, build to cost by reducing scope as required.

The methodology used to estimate the costs seems appropriate and is thoroughly documented in their cost book. There is some concern that the 10% management reserve is not sufficient to cover the uncertainties. There is, however, the mitigating fact that there is a commitment to build to cost, and there is some flexibility in scope.

There are two major concerns: first, the near term funding profile falls significantly short of the cost estimate (\$1.4M in FY01, and \$2.3M in FY02), and second, the integral falls far short of the total through 2006. If this is the funding they receive in FY01, they have identified the items that would not be funded in FY01, some of which will lead to a delay in the core software milestones. Since the project has not been fully resource loaded, it is difficult to be quantitative about the impact of budget shortfalls on the overall software effort. The funding agencies need to give further guidance to the collaboration in order for them to properly prioritize the tasks, scope, and work with them to expand and leverage resources.

Schedule

The collaboration is starting to use the project management tools developed for the construction project. They have identified about 20 major milestones starting in 1999 and extending to 2006 which ends with a full scale Tier 1 facility. The milestones are relatively uniformly distributed over that period. The first five milestones have been met.

The collaboration is in the process of fully implementing the resource loading into their project management tools, so we believe that the links and dependencies are not yet implemented in MS Project, so it is difficult at the moment to understand how slippages propagate through the project. In the list of major milestones, there do not appear to be any Tier 2 milestones. These should be included in the list of major milestones tracked by the project manager.

Project Management

US ATLAS PCP contains three subproject areas: physics, software and facilities. The project is leading the "projectisation" of, and closely coordinated with, the International ATLAS Physics and Computing Project. US ATLAS PCP and International ATLAS PCP are developing a common project plan by using a comprehensive project tool provided by US ATLAS. These two organizations have implemented a common WBS structure and working to same milestones.

The draft of the PMP has been refined. It includes the description of project organization, management structure, and a detailed description of relations between US ATLAS PCP and the various entities (ATLAS, ATLAS Physics and Computing Organization, CSG, NCB, COB, US ATLAS, US funding agencies, BNL and Columbia University, BNL PAP, and PCAP).

Since the US ATLAS PCP has been brought into the existing management structure for the US ATLAS detector construction project, same project reporting line and change control procedure will be used. Using already existing and mature project tools from the US ATLAS detector construction project, the US ATLAS PCP has produced a very detailed and resource loaded WBS, schedule and milestones. Various review and reporting plans will be initiated.

Agreements with ATLAS on the scope of US ATLAS PCP have been under discussion. However, since it will take some time to complete the project scope and produce MOU's, it was agreed that the 'software agreements (SA)' would be generated in order to fill some urgent and obvious vacancies.

Although details are in the process of being compiled, the annual obligation profile for the project was generated. The overall project has 10% contingency assignment with an assumption of "build to cost".

The committee commends US ATLAS PCP for leading the projectisation of the ATLAS Physics and Computing Project. Developing a common WBS structure and working to same milestones with the International ATLAS PCP will ensure a coherent program in the whole ATLAS Computing Project.

The draft Project Management Plan has been refined. The project is well organized and will be capable of providing the leadership and technical resources to complete the project. Although relations between US ATLAS PCP and the various entities are well described, the lines of communication were not very clearly defined (e.g. how does the APM interact with the CSG?).

The project management has made impressive progress on the detailed analysis and planning of the project cost and schedule for FY01-FY03. Using existing and mature project tools and procedures has many advantages. However, it might need some tuning to get useful monitoring and reporting, since reports that are useful in the construction project may be less useful in the software project. The committee is concerned that the large number of interfaces is burdensome and perhaps inconsistent. It is important to have a single mechanism for disseminating progress reports to all interested parties.

The plan for the interim software agreements with the ATLAS collaboration while the MOU is being developed is a good step. The US ATLAS PCP management should work expeditiously on producing the software agreements as well as moving toward the final MOU with the ATLAS collaboration. The US ATLAS

PCP APM should also work the on prompt implementation of the various US institutional MOU's and SOW's.

The presented obligation profile appears to be missing the cost for the Tier 2 facilities. The total project profile should include all funds needed in current plan. The scope definition and funding profile for R&D phase of the project (FY01-FY03) still need to be finalized.

US ATLAS – Recommendations

- We consider it crucial for the final definition of the project scope that these Software Agreements with International ATLAS be completed as expeditiously as possible.
- The role of US ATLAS in grid R&D needs to be defined. The committee feels that some additional effort is warranted but it should take a lower priority than urgent needs in core software and databases.
- While purchasing the main production hardware as late as possible makes sense, the hardware purchases and user support at the Regional (Tier 1) Center should be adequate enough so that BNL is “on the map” as a provider of computing resources during the design, development and startup phase of ATLAS. We recommend that the agencies exercise flexibility to achieve this objective.
- ATLAS should continue to identify areas where they can leverage the efforts of others and should perhaps look towards developing common projects with CMS.
- Over the next year, US ATLAS should define the scope (role and number) of the Tier2 centers in the US and should begin to define the requirements and the selection process (including the number of centers) and schedule for selection. This will help structure the discussion, control expectations, and help the funding agencies in obtaining support for these centers.
- In considering build to cost and level of effort constraints, areas of scope contingency should be defined and prioritized. Possible areas could be supporting multiple database products and the degree of grid-enablement required at startup due to uncertainty in the results of R&D.
- US ATLAS Physics & Computing Project management should work expeditiously on producing software agreements document with the ATLAS collaboration.

- US ATLAS P&C Project manager should work on prompt placements of US institutional MOU's.
- Work with the funding agencies and host laboratory on the obligation profile, so that it matches the proposed funding profile.
- Upon finalizing the project scope definition and funding profile for FY01-03, promptly complete the PMP and update the cost and schedule file.
- Include the cost of the Tier 2 facilities in the project cost and schedule file.
- Establish a procedure to distribute a periodic written progress report of the project to all interested parties.
- Develop priorities based on the present funding guidelines for FY01 and 02. Include priorities for those items that have been omitted, in case additional funding can be found.
- Include Tier 2 milestones in the list major milestones
- Implement the tasks and milestone dependencies by fully resource loading MS Project

US CMS – Summary

The goal of the US CMS Software and Computing Project is to provide the software and computing resources needed to enable US physicists to fully participate in the physics program of CMS. The other goal is to enable the US CMS collaboration to contribute to the international CMS Software and Computing project as a participating country.

The international CMS Software and Computing project, established and coordinated by the CMS collaboration, is responsible for delivering framework, architecture, tools and facilities needed for the computing and software of CMS experiment. The US CMS S&C Project is closely coordinated with the international CMS S&C project and involved in the CMS core framework and infrastructure software (referred as the core application software subproject), and hardware to support the reconstruction, simulation, and physics analysis (referred as the user facility subproject). The project organization is set up to carry out or oversee the tasks associated with these areas and to coordinate the activities and resources to meet the goals mentioned above.

The committee was impressed with a considerable progress made by the US CMS on the definitions of the US CMS S&C project scope and its organizational

structure. The draft Project Management Plan has been substantially completed and the permanent Level 1 Project Manager has been appointed. Structurally, USCMS S&C Project is well organized and will be capable of providing the leadership and technical resources to complete the project.

The area related to the user facility subproject is clearly defined and the associated cost estimate and schedule appear to be well understood. However, the necessity for the accelerated prototype program at Tier 2 centers is in question, given the limited funding availability in FY01-FY03. The scope of the core application software subproject is not as well defined and therefore the evaluation of its cost estimate and schedule was not possible.

US CMS – Findings and Evaluation

Project Scope

The project consists of two elements: user facilities and core application software. The method used to do resource accounting (FTEs) for these elements are quite different. For user facilities, a well-defined set of deliverables are defined in the WBS and manpower is assigned to meet these deliverables. The deliverables fall into two categories: development, implementation and operation of a Tier I facility at FNAL and development, implementation and operation of five Tier II facilities throughout the U.S. For the core applications, an overall level-of-effort is specified relative to the anticipated experiment-wide total need for CMS core software applications (the U.S. anticipates providing 25% of the software professionals). Although the core application software project has identified areas of activity (Software Architecture, Interactive Graphics, and Distributed Data Management: WBS items 2.1, 2.2 and 2.3), the management plans to use a “rolling” approach to defining contributions to the software.

The project management has indicated that the anticipated FY01 funds are insufficient to meet their milestones. They have indicated that in the absence of additional funding, the highest priorities in the CAS area are Detector and Event Visualization and Distributed Database Management. Consequences of the budget inadequacy would be: inability to hire 2 CAS engineers, one for Iguana and one for Oscar and G4 simulation; inability to hire an additional 3.5 user facility engineers; inability to buy \$140k of additional Tier I hardware; and loss of the requested 10% management reserve, exposing the project to additional risk.

The committee had significant concerns with the current model of resource accounting in the core software project. These concerns are centered in the following areas:

- ◆ The number of FTEs that US CMS is expected to provide is coupled to the overall CMS need. If CMS has underestimated this need, the US is subject to additional manpower requests, independent of whether US activities are in areas of the overrun.
- ◆ The lack of clearly defined deliverables makes it difficult for the project to avoid “mission creep.”
- ◆ The lack of clearly defined deliverables makes it difficult for the L1 Project Manager to appropriately track progress on the core software. In addition, it places significant oversight load on the Level 2 Core Software Project Manager.

Cost

US CMS presented a cost estimate that totals a cost estimate for the project from 2001 to 2006. The costs for 2007 and beyond are considered operating costs. The costs were broken down into four main cost categories (in then-Yr\$): project office (\$2M); core software (\$11M); user facilities (\$20M hardware, \$26M manpower); management reserve (\$5M). The total project cost (2001-2006) is \$64M. Given the uncertainties in making estimates for the outyears, particularly software manpower, the group adopted an approach in which the CAS costs for this FY were based the detailed WBS, and at less detail in the ensuing years.

The manpower costs for the user facilities were estimated using bottoms up approach and checked against Fermilab experience with the CDF and D0. Both methods yielded comparable totals (although differing significantly in the details). The manpower costs for the core software are limited to a 25% level of effort (based on the scale of US participation in CMS). The hardware costs were estimated using cost projections based on observed cost trends. They note that there can be significant uncertainties in the hardware costs. The management contingency (10%) was not based on a risk-based analysis, but is used to provide flexibility to react to problems. The collaboration noted that the funding guidance they received from the DOE and the NSF falls significantly short of the cost estimate, both in the near term and in the integral. The shortfall for FY01 is \$1.57M, and ~\$2.3M in FY02. The total shortfall through 2006 is about \$12M. If this shortfall persists in FY01, they will not fund 2FTE core software engineers and 3.5FTE user facility software engineers and some of the Tier 1 hardware, as well, of course, as the management contingency. The collaboration states that they will, however, build to cost.

The methodology used to estimate the costs seems appropriate. The user facility cost estimate is well developed. There is some concern that the 10% management reserve is not sufficient to cover the uncertainties. There are, however, two points that mitigate that concern: the commitment to build to cost; and that the host laboratory is Fermilab, which has some ability to backstop potential manpower

shortages. There are two major concerns. The first is that the near term funding profile falls significantly short of the cost estimate (\$1.6M in FY01, and \$2.3M in FY02), and that the integral falls far short of the total through 2006. If this is the funding they receive in FY01, they have identified the items that would not be funded, some of which will lead to a delay in the core software milestones. The funding agencies need to give further guidance to the collaboration in order for them to properly prioritize the tasks and scope.

The second major concern is that since the core software milestones are not in place, then the 25% level of effort for the core software for international CMS is not well defined. Hence, there is considerable uncertainty in the cost for this task. The committee was somewhat confused about how to reconcile a fixed level of effort when the scope of the US contribution is not formally defined. It is not yet clear how this will be resolved, and there is a significant potential cost exposure.

Schedule

The collaboration presented some high level milestones - 1000 CPU/20 TB in 2001, 5% test in 2002, and 20% in 2004. They also showed a project schedule. The committee did not get see a detailed list of milestones. The schedule tools are not yet fully implemented. The GANNT plots are now simply used to identify the tasks and milestones. The dependencies for these tasks are not included, so it is difficult at the moment to understand how slippages propagate through the project.

Project Management

USCMS S&C Project has made a significant progress on identifying and appointing a various key management roles, completed the final draft of Project Management Plan (PMP), and defined the role and composition of the US CMS Advisory S&C Board (ASCB). The permanent Level 1 Project Manager for the S&C project is now on board and ASCB is in place. Appointment of permanent Level 2 and Level 3 Project Managers will follow this review. Relations to various entities (CMS, USCMS, US CMS Construction Project, US funding agencies, Fermilab Computing Division, and Fermilab directorate) and associated lines of communication are defined in the draft of PMP along with the description of project organization and management structure. The draft of the PMP was endorsed by the US CMS Collaboration Board.

Upon the request from DOE and NSF as project sponsors, Fermilab Directorate is providing the management oversight. Fermilab Directorate established a Project Management Group and a standing external oversight panel to provide monitoring of both management and technical progress of the USCMS S&C Project.

USCMS S&C Project has defined and documented the Project Plan, which includes a detailed and resource loaded WBS, schedule and milestones. Various review and reporting plans are now in place and Project Office at Fermilab is being established. Agreements with CMS on the scope of US CMS S&C Project are being formalized and CMS is working towards producing CMS Computing MOU's in the collaboration by year 2003.

Annual budget requests were submitted to funding agencies. Based on the availability of funds, Level 1 Project Manager will review, negotiate and approve SOW's from the participating US institutions. More general (long-term) responsibilities such as deployment of T2 centers will require approved MOU's.

Structurally, USCMS S&C Project is well organized and will be capable of providing the leadership and technical resources to complete the project. The responsibilities and organizational structures of the project as well as various other entities are well defined in the draft PMP and found to be reasonably sound. The draft of the PMP is well along and should soon be approved by the proper authorities (which need to be defined – no signature page). Because the more precise scope of the project still needs to be defined, some details of PMP, such as Change Control procedure, will have to be worked out later time and the revision of the associated planning documents will be necessary.

Various oversight and review processes by both internal and external organizations, which are now in place, seem adequate. Because of the necessity of extended lines of communication (CMS, USCMS, CERN, Fermilab, DOE, NSF, JOG, ASCB, PMG, Fermilab oversight panel etc), submission of periodic written reports to all parties involved will be very valuable tool.

USCMS S&C Project management should work expeditiously on producing overall project MOU with the CMS collaboration. The proposed date of year 2003 seems to be quite late. The Level 1 Project Manager should also work on prompt placement of various US institutional MOU's and SOW's.

The scope definition and funding profile for R&D phase of the project (FY01-FY03) still need to be finalized. When this has been accomplished, the PMP and its associated planning documents should be iterated to reflect the plan. The PMP should then be submitted to the funding agencies and host laboratory management for their approval.

US CMS – Recommendations

- USCMS S&C Project management should work expeditiously on producing MOU with the CMS collaboration.
- Level 1 Project Manager should work on prompt placements of US institutional MOU's and SOW's.
- USCMS S&C Project management should work with the funding agencies and host laboratory on an obligation profile, which matches with the proposed funding profile.
- Upon finalizing the project scope definition and funding profile for FY01-03, promptly complete the PMP and update other planning documents.
- Establish a procedure to distribute a periodic written progress report of the project (quarterly report) to all organizations, which are on the extended lines of communication.
- Evaluate the overall schedule alignment of the project with the overall CMS S&C Project schedule and progress.
- Developed funding priorities based on the present funding guidelines for FY01 and 02. Include priorities for those items that have been omitted, in case additional funding can be found.
- Provide a clear list of milestones.
- Implement the tasks and milestone dependencies into the project management tools.

6 Action Items

- The agencies should conduct a mini-review of both US ATLAS and US CMS software and computing projects in about 6 months, a full review of both projects in about a year.

7 Appendices

Appendix A – Charge to Committee

To: DOE/NSF LHC Program Office

Re: Review of the U.S. LHC Software and Computing Projects

The Joint Oversight Group for the U.S. LHC Program requests that an independent peer review of the U.S. LHC Software and Computing Projects be conducted at Brookhaven National Laboratory on November 14-17, 2000. This review will initiate 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 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. 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 2002. However, the review committee should 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 should be to assess:

- The overall strategy and scope of the U.S. LHC software and computing efforts;
- 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;
- Other aspects of the U.S. LHC computing structures, such as networking, “grid” computing, and the relationship between U.S. Tier 1 facilities and any smaller regional and university facilities;
- 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 Software and Computing Projects.

Please provide a report on the review to this office by January 15, 2001.

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
Joel Butler	Fermilab
Aesook Byon-Wagner	Fermilab
Pat Dreher	MIT/LNS
Michael Ernst	DESY
John Reynders	Sun Microsystems
Terry Schalk	U.C. Santa Cruz
Marjorie Shapiro	U.C. Berkeley/LBNL
Michael Tuts	Columbia Univ.
Chip Watson	Jefferson Lab

Appendix C – Review Agenda

U.S. LHC Software and Computing Review
Nov. 14-17, 2000
Brookhaven National Lab

Tuesday Nov. 14

8:00am Executive Session (Full Committee) Rm 2-160 Physics

US CMS Overview Session- Large Seminar Room

9:00	Welcome	T. Kirk
9:05	Introduction	G. Crawford
9:20	International CMS Computing	M. Pimia
9:50	CMS Software: Status, Plans, and Milestones	D. Stickland
10:10	Break	
10:30	CMS and US CMS Physics Plans	J. Branson
10:50	US CMS Software and Computing	M. Kasemann
11:20	US CMS Core Applications Software	L. Taylor
11:50	User Facilities	V. O'Dell
12:20pm	Lunch (Cafeteria)	
1:30	<u>Breakout Sessions:</u> -Project Management	Rm 2-160
	-Core Software	Rm 2-187
	-User Facilities	Rm 3-192
3:00	Subcommittee Executive Sessions/Writing	

4:00	Break	
4:30	Executive Session (Full Committee)	Rm 2-160
6:00	Executive Session w/CMS Mgmt	
6:30	Adjourn	

Wednesday Nov. 15

8:00am	CMS Response to Committee Questions	Rm 2-160
8:30	Subcommittee Sessions/Writing	various
10:00	Executive Session (Full Committee)	Rm 2-160
10:30	Closeout Dry-Run	
11:30	Working Lunch (Full Committee)	
12:30pm	Closeout with US CMS	Rm 2-160
1:30	Break	

Common Projects Overview – Large Seminar Room

1:50	Common Projects Intro	G. Crawford
2:00	CERN Response to LHC Computing	H. Hoffmann
2:30	Major Grid Computing Initiatives	I. Foster
2:50	Particle Physics Data Grid	R. Mount
3:05	MONARC	H. Newman
3:20	Networking Needs for LHC	L. Price
3:40	Break	

4:00	Tier 2 Centers and Configuration	P. Avery
4:15	Software and Middleware for Tier 2 Centers	R. Gardner
4:30	US ATLAS Tier 2 Plans	R. Gardner
4:45	US CMS Tier 2 Plans	M. Kasemann
5:00	Executive Session (Full Committee)	Rm 2-160
6:30	Adjourn	

Thursday Nov. 16

8:30am	Executive Session (Full Committee)	Rm 2-160
---------------	---	-----------------

US ATLAS Overview Session – Large Seminar Room

9:00	International ATLAS	N. McCubbin H. Meinhard
9:30	US ATLAS Software and Computing	J. Huth
10:00	Break	
10:20	Software Subproject	T. Wenaus
11:00	User Facilities	B. Gibbard R. Baker
11:40	Physics Support	I. Hinchliffe
12:00	Budget and Schedule Review	J. Huth
12:30pm	Lunch (Cafeteria)	

1:30	<u>Breakout Sessions:</u> -Project Management -Core Software -User Facilities	Rm 2-160 Rm 1-189 Rm 3-192
3:00	Subcommittee Executive Sessions	
4:00	Break	
4:30	Executive Session (Full Committee)	Rm 2-160
6:00	Executive Session w/US ATLAS Mgmt	
6:30	Adjourn	

Friday Nov. 16

8:30am	ATLAS Response to Committee Questions	Rm 2-160
9:00	Subcommittee Sessions/Writing	various
10:30	Executive Session (Full Committee)	Rm 2-160
11:00	Closeout Dry-Run	
12:30pm	Working Lunch (Full Committee)	
1:30	Closeout with US ATLAS (incl. Common Projects)	
3:00	Adjourn	

Appendix D – Cost Tables

US CMS

The following tables summarize the project manpower and cost.

WBS item	FY	2001	2002	2003	2004	2005	2006	2007
1.1	Tier 1 Regional Center	0	0	0	13	17	20	18
1.2	System and User Support	1	2.5	3	3	5	4.5	4.5
1.3	Operations and Infrastructure	1	1	1.5	2	2.5	5	5
1.4	Tier 2 Regional Centers	3	3	4.5	4.5	6	7.5	7.5
1.5	Networking at Tier 1	0.5	1	2	2.5	3	3	3
1.6	Computing and Software R&D	2.5	4	4.5	1	1	1	1
1.7	Detector Construction Phase Computing	2	2	2	0.5	0	0	0
1.8	Support for FNAL based computing	0.5	1	1	3	2	2	2
1	User Facilities (total FTE)	10.5	14.5	18.5	29.5	36.5	43	41
2	Core Application Software (total FTE)	10	11	12	13	13	13	13
3	Project Office	1	2	2	2	2	2	2
	Total	21.5	27.5	32.5	44.5	51.5	58	56

Personnel requirements for all WBS items User Facilities, CAS and the Project Office. The Table shows the full number of FTE's on the project, excluding physicists, but including engineers, technicians, and support staff.

WBS item	FY	2001	2002	2003	2004	2005	2006	Total	2007
1	User Facility Project								
1.1	Tier 1 Regional Center	0.0	0.0	0.0	2,648.9	2,838.7	2,595.4	8,083.0	2,553.3
1.2	System and User Support	0.0	29.0	23.2	0.0	34.8	0.0	87.0	15.4
1.3	Operations and Infrastructure	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
1.4	Tier 2 Regional Centers	282.0	470.0	590.0	1,120.0	1,750.0	1,750.0	5,962.0	1,250.0
1.5	Networking at Tier 1	24.7	26.1	65.3	543.9	449.2	458.8	1,567.9	482.0
1.6	Computing and Software R&D	348.4	622.7	479.4	0.0	0.0	0.0	1,450.5	0.0
1.7	Detector Construction Phase Computing	157.2	40.0	56.3	7.0	0.0	0.0	260.4	0.0
1.8	Support for FNAL based computing	11.6	71.9	139.2	23.2	52.2	23.2	321.3	48.4
	Total hardware cost in FY01\$ [\$k]	823.8	1,259.8	1,353.3	4,343.0	5,124.8	4,827.4	17,732.2	4,349.2

The Table shows the hardware cost for the User Facility including Tier2 centers without escalation (in FY01 k\$).

FY	2001	2002	2003	2004	2005	2006	Total	2007
Project Office (WBS item 3)	0.16	0.32	0.32	0.32	0.32	0.32	1.74	0.32
CAS Total (WBS item 2)	1.58	1.74	1.90	2.05	2.05	2.05	11.38	2.05
UF Total (WBS item 1)	2.41	3.48	4.17	8.89	10.74	11.43	41.13	10.64
UF personnel	1.59	2.22	2.81	4.55	5.62	6.61	23.39	6.29
<i>for Tier 1</i>	<i>1.19</i>	<i>1.82</i>	<i>2.21</i>	<i>3.95</i>	<i>4.82</i>	<i>5.61</i>	19.59	<i>5.29</i>
<i>for Tier 2</i>	<i>0.40</i>	<i>0.40</i>	<i>0.60</i>	<i>0.60</i>	<i>0.80</i>	<i>1.00</i>	3.80	<i>1.00</i>
UF hardware	0.83	1.26	1.36	4.34	5.13	4.82	17.74	4.35
<i>for Tier 1</i>	<i>0.55</i>	<i>0.79</i>	<i>0.77</i>	<i>3.22</i>	<i>3.38</i>	<i>3.07</i>	11.77	<i>3.10</i>
<i>for Tier 2</i>	<i>0.28</i>	<i>0.47</i>	<i>0.59</i>	<i>1.12</i>	<i>1.75</i>	<i>1.75</i>	5.96	<i>1.25</i>
Total costs in Million FY01\$	4.15	5.53	6.38	11.26	13.11	13.80	54.24	13.01
Total costs escalated, including managment res. [M\$]	4.57	6.27	7.44	13.54	16.24	17.60	65.65	17.09

Project Overview Table 1: Budget Summary of US CMS Software and Computing Project. The costs shown for Tier2 centers in this table are for staff and hardware located at the Tier 2 centers. Amounts are given in units of Million \$ FY01, the last line shows totals costs escalated including management reserve.

Summary Information FY 01 & FY 02



Funding Categories	FY01	FY02
Estimate of DOE Funding Requirements	3,493	4,323
Estimate of NSF Funding Requirements	460	617
USAILAS Core Computing Project Subtotal	3,953	4,940
Distributed IT Infrastructure (Tier II)	412	1,255
Summary Total	4,365	6,195

	FIE	
Proposed Project Funded FIEs	FY01	FY02
Physics	1	1
Software	129	15.5
Facilities		
Tier I	44	5.8
Tier II	1.1	1.7
Proposed Infrastructure Funded FIEs		
Distributed IT Infrastructure (Tier II)	1	5
Summary FIE Total	204	29

	AY1s\$	
	FY01	FY02
Core Funded Facilities Purchases	639	579
Distributed IT Facilities Purchases	324	565

Appendix E – Schedules

US CMS

The following tables summarize the high level milestones for the User Facility and the CAS subproject.

Milestones UF	Date
	16-Nov-00
1 M Events Data Challenge (~0.1%)	Jun-00
1st Prototype Tier 2 Center	Sep-00
5 M Events Data Challenge (~0.5%)	Dec-00
Trigger TDR	Dec-00
Create AOD for Physics Group	Jan-01
Choice of 2nd Prototype Tier 2 site	Apr-01
Start of 2nd Prototype Tier 2 implementation	Jun-01
10 M Events Data Challenge (~1%)	Jun-01
DAQ TDR	Dec-01
1-2% Data Challenge with full hierarchy	Dec-01
2% Data Challenge with full hierarchy	Jun-02
Choice of 3rd Prototype Tier 2 site	Oct-02
Start of 3rd Prototype Tier 2 implementation	Dec-02
CMS Software and Computing TDR	Dec-02
Tier 1 5% Prototype	Mar-03
5% Data Challenge	Jun-03
Begin Tier 1 Center Implementation	Oct-03
Decision Plan for Tier 2 Center locations	Oct-03
Start 1st full Tier 2 Center Implementation	Dec-03
Start 2nd full Tier 2 Center Implementation	Dec-03
Physics TDR	Dec-03
10 % Tier 1 installed	Apr-04
Start 3rd full Tier 2 Center implementation	Dec-04
Start 4th full Tier 2 Center implementation	Dec-04
20% Data Challenge	Dec-04
40 % Tier 1 Installed	Apr-05
Begin Pilot LHC Run	Oct-05
Start of 5th full Tier 2 Center implementation	Dec-05
Begin LHC running	Apr-06
100% Tier 1 installed	Apr-06

US CMS Core Application Software Milestones For 2001-2002

2.1 Architecture

2.1.2.1.2	Detector Description Conceptual Design	(Mar 2001)
2.1.3.1.5	First Version of CAFÉ Tools Complete	(Mar 2001)
2.1.3.2.1.4	Top-level Architectural Design Document Complete	(Mar 2001)
2.1.2.4.3	Analysis Architecture / Kernel Defined for FFP [i.e. IGUANA Fully Functional Prototype]	(Oct 2001)

2.2 Interactive Graphical User Analysis

2.2.2.2.4.3	First Version of 2D Browser	(Oct 2001)
2.2.2.3.3.	GEANT4 Visualization Program	(Oct 2001)
2.2.3.4.8.5	Data Browsers for CARF/ORCA/OSCAR	(Oct 2001)
2.2.1.4.3	Reconsideration of Baseline GUI Tools	(Oct 2001)
2.2.2.4.3	Reconsideration of Baseline 2D/3D Graphics Choices	(Oct 2001)

2.3 Distributed Data Management and Processing

2.3.1.1.4	Basic Functionality of Distributed Task Scheduler at Tier2 Prototype Center	(Feb 2001)
2.3.1.1.6	Distributed Task Scheduling First Prototype Complete	(Jun 2001)
2.3.1.2.2	Distributed Task Scheduling Available for ORCA Production at a Single Site	(Sep 2001)
2.3.1.2.3	Distributed Task Scheduling Between Tier1 and Tier2 Prototype Centers	(Dec 2001)
2.3.1.2.6	Distributed Task Scheduling between CERN, Tier1, & Tier2 Prototype Centers	(Jun 2002)
2.3.2.2.4	File Format Independent Data Replication in GDMP Tools	(Feb 2001)
2.3.2.2.6	GDMP Globus Based Data Replicator First Prototype Complete	(Jul 2001)
2.3.2.4.4	Implementation of AMS Security Protocol Plug-in	(Feb 2001)
2.3.5.1.9	Update of Estimated CMS Computing Needs	(Jan 2001)

Top-level closely-related International CMS milestones

"Fully Functional Prototype" GEANT4 Simulation of CMS	(Jun 2001)
ORCA Production for HLT studies, 10M events, 20TB output	(Jun 2001)
"Fully Functional Prototype" reconstruction/analysis framework	(Dec 2001)
Choice of ODBMS supplier	(Dec 2001)
Data Acquisition TDR	(Dec 2001)
"Fully Functional Prototype" detector reconstruction	(Jun 2002)
"Fully Functional Prototype" physics object reconstruction	(Dec 2002)
"Fully Functional Prototype" user analysis environment	(Dec 2002)
Data Challenge (5%)	(Dec 2002)
Software and Computing TDR	(Dec 2002)



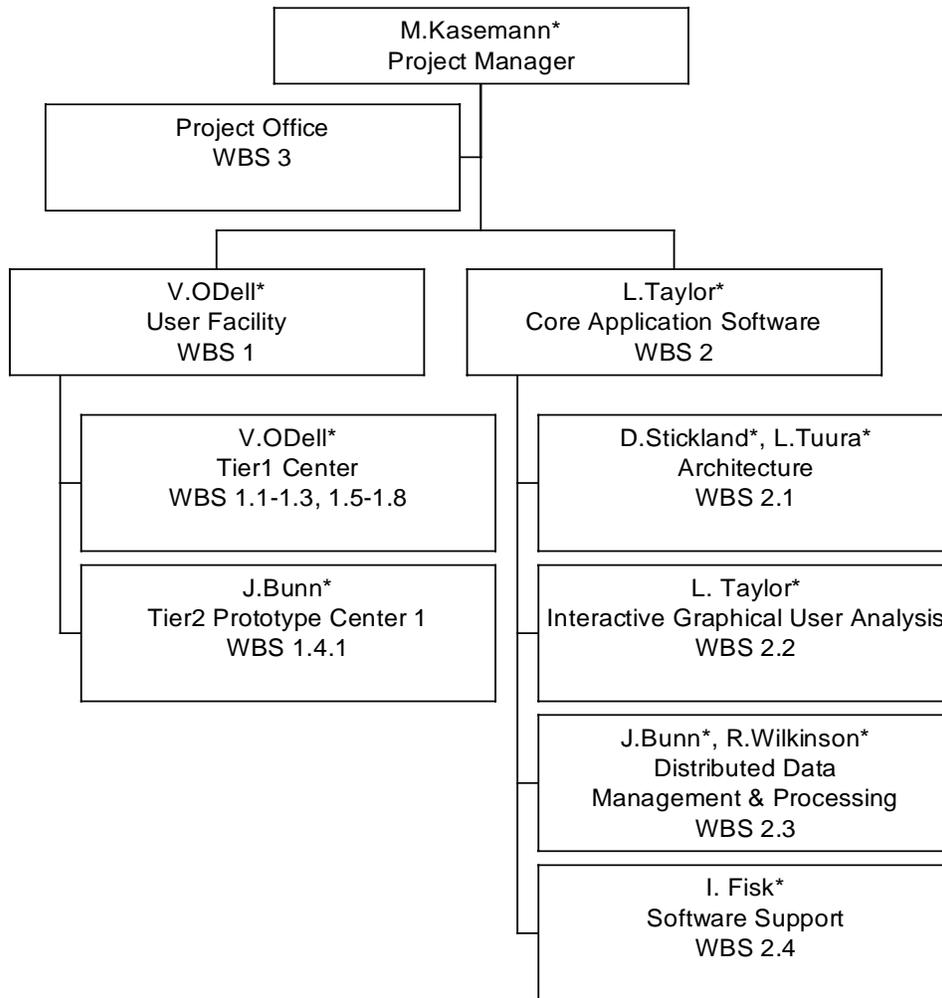
Schedule Milestones

WBS	Description	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008
2.2.1.3	1 Tbyte database prototype	Jan 1									
2.2.1.2	Release of Athena pre-alpha version		May 9								
2.3.1.1.2.1	Tier 1 Processor Farm Prototype			Sep 29							
2.2.1.2	Athena alpha release		Sep 29								
2.2.1.4	Geant3 digi data available		Oct 30								
2.2.1.2	Athena Beta Release			Dec 29							
2.3.2.1	First definition of regional centers			Jan 1							
2.3.1.1.4.1	Tier 1 Storage Prototype				Oct 1						
2.2.1.3	Decide on database product				Jun 29						
2.4.5	MDC 0 Completed					Dec 12					
2.2.2.1.2	Full validation of G4 physics					Dec 31					
2.4.5	MDC 1 Completed					Jul 30					
2.4.2	Computing TDR Finished						Nov 29				
2.3.1.1.2.1	Tier 1 Upgrade (for MDC2)						Dec 31				
2.3.1.1.2.2	Tier 1 Large Scale Test (MDC2)							Sep 30			
2.4.3	Physics readiness report completed							Jun 30			
2.2.2	Test full software chain in real environment							Jul 30			
2.2.1.3	Full DB infrastructure available								Dec 31		
2.3.2.2	20% Processing Farm Prototype									Jun 30	
2.3.1.1.2.3	Tier 1 Full scale										Oct 2

Appendix F – Organization Charts

U.S. CMS

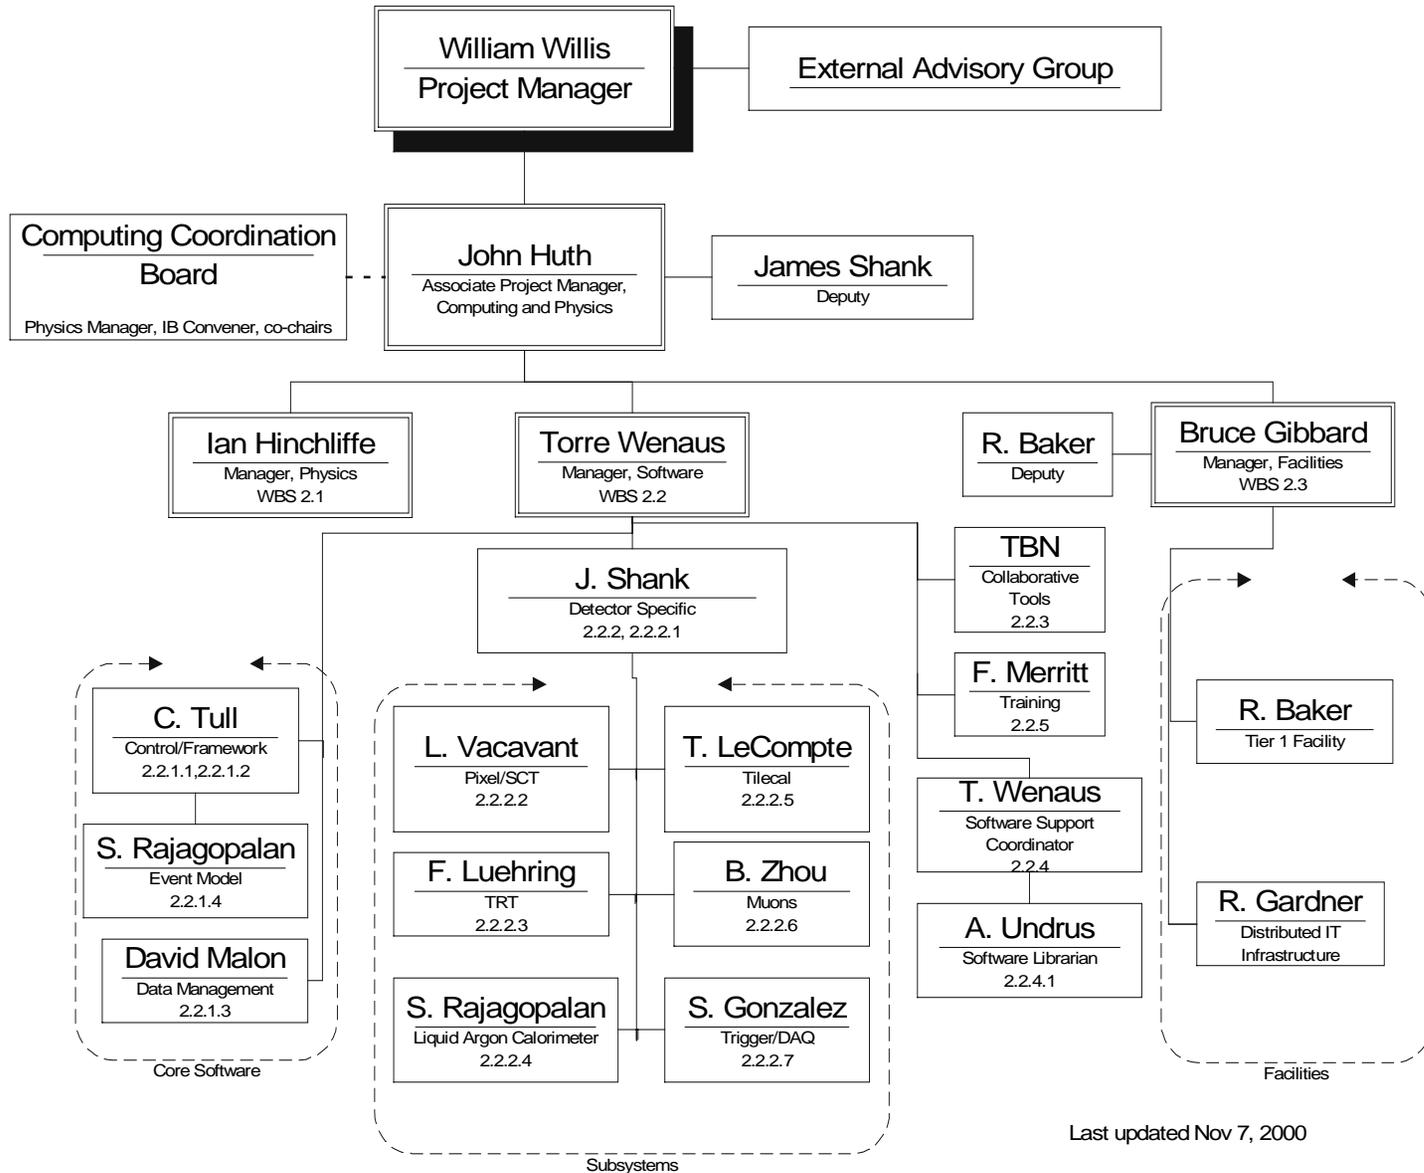
US CMS Software and Computing
Organizational Chart
Version 2.13
11/15/2000



* acting positions

The final appointments will be done by L.A.T. Bauerdick and final project managers.

U.S. ATLAS PCP Organization



46

Last updated Nov 7, 2000