Executive Summary

ATLAS and CMS will be large, general-purpose detectors for observation of very high energy proton-proton collisions at the Large Hadron Collider (LHC). This facility is currently 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 US investment in constructing the LHC and its detectors, the US LHC software and computing projects must be successful in empowering the U.S. community of scientists and students to perform frontline physics analysis. The US LHC Computing Projects must contribute an appropriate share of computing resources and software effort and must ensure that US physicists can fully and immediately perform analysis of the data and contribute meaningfully to the research work of the physics groups. A key goal is to enable US participants to analyze data while continuing to reside in the US.

A peer review of the US LHC Software and Computing Project was held on 14-17 January 2003 at the Lawrence Berkeley National Laboratory in Berkeley, California to assess the current understanding of the scope, cost and schedule for the US LHC Software and Computing projects and the operation of their management structures. Both US ATLAS and US CMS were asked to provide self-consistent project plans targeted to the funding guidance provided from the U.S. Department of Energy (DOE) and National Science Foundation ~SF), to address separately how incremental funds could be used, and to present contingency plans that could deal with funding levels as much as *25%* below the guidance. Because of the dynamic nature of software and computing, the review emphasized plans for the next two years, and concentrated on issues arising from major changes during the past year that have caused changes to the initial baselines adopted at the November 2001 review. These include changes to the LHC schedule, the expanding worldwide effort to facilitate Grid-computing, including the advent of the LHC Computing Grid (LCG) project, funding guidelines from the agencies, and the development of the US LHC Research Program. In addition, the review addressed how external efforts, particularly from Grid R&D and deployment projects, are integrated into the LHC Software and Computing projects. Lastly, since the projects have been operating for roughly one year since the initial baselining, an assessment of technical progress was made. The complete charge to the review panel is shown in Appendix A.

The review consisted of plenary talks and breakout sessions in four topical areas:
Facilities, Grids and Common Projects, Software, and Project Management. The expert reviewers provided comments throughout the review directly to the U.S. LHC collaboration members as well as to the DOE and NSF representatives present. They asked questions of the collaborations both orally and in writing, and deliberated issues and impressions in executive session with only agency representatives present. The comments, observations and assessments of the reviewers form the basis of this report. Because of the large overlap in subject, the committee combined the Grids and Common Projects breakout sessions with those of Facilities. In this Executive Summary these two sections are kept separate for US ATLAS, but for US CMS they have been combined.

Comments and Recommendations applicable to both experiments

The committee commends both experiments for their detailed presentations and clear answers to the many questions that were posed. Both groups are adjusting well to tight budgets by stretching out some activities (taking advantage of the delay in LHC startup), leveraging software development efforts, particularly those associated with grid developments, adopting common tools, from the LCG and other external sources, and making opportunistic use of external hardware resources.

In the area of Facilities, both experiments achieved impressive success in the completion of recent Data Challenges, despite funding shortfalls at the Tier 1 centers. This was done by making opportunistic use of off-project resources. Funding must now stay on or above the present guidance if the US groups are to succeed in future, much more ambitious data challenges, since it will difficult to make up larger shortfalls with outside resources.

The committee congratulates both ATLAS and CMS for integrating grid developments into their work in a serious way. Both have shown good acceptance of the LCG- selected software as their ultimate platform for grid computing. The committee applauds their efforts to influence and contribute to LCG decisions and encourages vigorous input to the prioritization of LCG products to ensure completion of those critical to US CMS and ATLAS. The committee encourages efforts to push for the ability to use heterogeneous configurations within their computing facilities and to avoid insistence on a homogeneous, overly rigid, configuration. The current arrangement of a single US member on the Grid Deployment Board at CERN gives US CMS and ATLAS a coherent approach to grid issues and is viewed by the committee to be advantageous to the US.

The US funding agencies need to make commitments to provide on-going, long-term maintenance and development support of the grid middleware that is developed in their R&D projects if they expect them to be adopted by LCG. Given the uncertain state of some of the grid components and their support, the committee encourages ATLAS and CMS to develop a matrix of grid software that articulates the interdependencies and places them within the overall software framework. We further encourage them to develop contingency plans for the various components.

The committee encourages both experiments to continue efforts to get the prototype Tier 2 centers into operation and to establish MOU's with the iVDGL Tier 2 centers concerning various "production" deliverables. However, the prototype Tier 2 centers are also excellent sites for testing prototype software developments and deployments. They should be allowed to operate for part of the time in research mode", which is consistent with their charter. This research will facilitate opportunistic use of non-ATLAS/CMS owned resources in the future.

There are good efforts underway to pursue grid monitoring. The committee encourages continued efforts to develop integrated monitoring capabilities that provide end-to-end iriformation and cooperation with the LCG to develop common monitoring capabilities.

Comments and Recommendations concerning US ATLAS

Facilities:
ATLAS presented the ramp-up profile for development of facilities, aiming to meet the goals of future data challenges within the current agency funding guidance. The ATLAS plan includes reduced staffing of the Tier 1 center to 20 FTE[1]s from last year's plan of 25, based on experience obtained with DC1 and the continued synergy with the RHIC Computing Facility. Since personnel costs dominate within the Facilities project, this savings allows achievement of the remaining goals. With the current agency guidance, the minimal ATLAS requirements will be met, but with almost no headroom.

The fully disk resident computing model was discussed. Implementation would improve access to the data for analysis. However, this scenario does not fit into the current agency guidance and would require modest additional fluids in 2007 and 2008.

The committee has the following recommendations:
- In view of the lean staffing of the Tier 1 center, plans should be developed on how the required grid and networking support and maintenance requirements can be covered.
- Continued attention must be paid to providing adequate network bandwidth and support.
- Every effort should be made to identify the additional fluids necessary in order to realize the fully disk-resident ESD model.

Grids:
The committee encourages US ATLAS to develop a coherent plan and dedicated hardware platform for testing of prototype configurations and new software. The committee is concerned by the loss of base funding for the Grid/Athena integration. This should likely be a priority if additional funding becomes available.

Software:
US Atlas reported on their contributions to core software, database management, software support for development and analysis, and subsystem software related to hardware responsibilities. It is encouraging that US ATLAS has consolidated leading roles in ATLAS core software areas

ATLAS is well-coupled to LCG activities. ATLAS has established a substantial architectural design presence in the LHC-based grid projects, while its Project DB activities have been adopted by/transferred to LCG. Architecture/framework effort is being sustained.

ATLAS goals and definition of success in the DC 1 effort have been scaled back substantially, to the point where some of the accomplisliments do not reflect scalable accomplishments for the LHC science; success is also heavily dependent of the state of the POOL project. ATLAS base program support from the national

laboratories has diminished at an alarming rate, and continuing lack of US ATLAS and ATLAS presence at CERN is a serious concern to the committee.

The committee has the following recommendations:

- US ATLAS must take care to properly plan across its operation the impact on US-specific efforts of taking on new central tasks.
- We endorse the US ATLAS request for three software engineers to be resident at CERN

Project Mana~ement:

The committee has the following recommendations:

- The scope should be more formally defined with software agreements. These can be modified as things evolve, perhaps by using the change control procedure.
- US ATLAS should continue to be wary of taking on additional projects from International ATLAS (same as recommendation #2 from last year). Scope creep should not be allowed to jeopardize the US deliverables.
- The project managers could benefit from increased use of traditional project management tools to aid in decision making and presentations. US ATLAS should consider assisting this effort with additional personnel to aid the computing project managers.
- It is important to have some personnel at CERN to assist the US ATLAS members who are moving into larger roles in ATLAS computing.

Comments and Recommendations directed to US CMS

Grids and Facilities

CMS has dealt with its current funding shortfall by ramping up more slowly than planned so that the target of 13 FTEs planned for the beginning of FY03 will not be achieved until much later in the year. To offset shortfalls in hardware at the Tier 1 center, US CMS has made opportunistic use of available facilities from US base programs and the US grid program. This has allowed them to meet their obligations to the recent data challenge. Provided the funding based on current guidance is achieved, Tier 1 and 2 procurements will be able to deliver the needed capacity for the next data challenge, DCO4.

US CMS has a strategy of supporting separate production, integration and development grid platforms, which is very good. US CMS has a strategy for selecting and deploying "final" production Tier 2 centers, but funding will come from the as yet not completely funded US LHC Research Program. US CMS is contributing to a number of R&D topics. There is some concern that these projects could drain US CMS resources if not kept under control.

Fermilab uses Kerberos encryption for security while many other LHC sites use a different scheme (PKI). Dealing with Kerberos interoperability is underway, including

evaluations of existing tools. The committee is concerned about scalability of possible solutions to all of International CMS.

Networking to Fermilab is lagging that of US universities and may not provide adequate support for a CMS Tier 1 center. It must provide bandwidth commensurate with other Tier 1 centers. Either ESNET must be more aggressive at deploying bandwidth, or FNAL should consider an alternative supplier.

The committee has the following recommendations:

- US CMS should ensure that FNAL has networking plans which meet CMS[9]s needs as a function of time from DCO4 to turn-on (in the presence of other FNAL demands).
- The issue of authentication in heterogeneous environments (including Kerberos credentials) for all of International CMS should be studied further.

Software

The SW team successfully supported the September 2002 production runs for the Trigger and Data Acquisition Technical Design Report. They have a good reconstruction program on which to build. They have been able to engage 200 people worldwide, 50 in the US. The committee was told that the US CMS baseline program has fully adopted a reliance on emerging grid technology for its data distribution and analysis model. To date, US CMS has built a subscale production grid for the collaboration and has used it for both demonstration, in the Supercomputing 2002 conference, and production. US CMS is now contributing 25% (7.5 FTE's) to the International CMS computing effort.

CMS has made a formal decision to replace Objectivity to provide data persistency. They are committed to replace it with POOL deliverables from LCG. POOL progress is partially on project.

In general, the software project is making extensive use of the Work Breakdown Structures (WB S) and other formal tools to manage the program. They are now tracking required external product developments as well, e.g., POOL and grid products. NonCMS POOL leaders regularly participate in CMS internal status reviews. CMS will rely on the results of external grid R&D, e.g., NSF Information Technology Research projects in the US, to adopt interactive distributed analysis capability for the experiment. They are committed to a non-local focus for most of the DCO4 analysis, e.g., alignment or calibration. ~Project management does not want a local DC analysis on local workstations. They Want to push the distributed analysis model.

The committee has the following recommendations:

- CMS should not allow current manpower allocation to erode and should push for the needed ramp up.
- CMS should promote tightest possible coupling between key CMS software efforts with their specific needs to off-project (quasi external) efforts that will benefit CMS.

Pro lect Mana~ement

The committee finds that: management has good control of the project, aided by close collaboration between US CMS SC management and that of the CMS Core Software Project(CCS). Close interaction with US CMS management is also helpful. Con servative projection of resources available and level-of-effort commitments help in managing the project. Schedule and scope control mechanisms are in place and appear to be working adequately.~Changing totals, profiles and arrival of flinds at the end of the fiscal year greatly complicate the management task. Flexibility from the agencies in shifting support for tasks in different parts of the project is a helpful tool for addressing this.

Grid and application software from LCG are now in the WBS to aid in project planning. CMS expects LCG to deliver, but this implies some risk However, there is concern whether LCG will deliver all the functionality needed for DCO4 in time.

US CMS devotes *25%* of their computing FTE's to US user support. The committee commends this step and encourages continued efforts to maximally support analysis work
in the US.

Development of a remote control room and a US-based physics analysis center at Fermilab for CMS should provide useful ways to support collaboration and maximize US physics impact.

The committee has the following recommendations:

- The project management plan for the current era of research project be written and approved.
- The committee considers it important to continue to allow flexibility in moving fluids among projects in the coming era of funding as a research project.
  CMS should closely monitor delivery, schedule and performance of externally provided software.

Table of Contents

1 Introduction

This report is the product of the DOE/NSF review of US LHC Computing Projects held at Lawrence Berkeley National Laboratory on 14-17 January 2003. 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 (next two years) plans of both collaborations where plans are more concrete in developing software and user facilities for the LHC experiments, ATLAS and CMS, that are scheduled to begin taking physics data in 2007.

Twelve outside experts were selected jointly by officers of the Division of High Energy Physics of the US Department of Energy and officers of the EPP Program in the Division of Physics of the National Science Foundation. The outside experts reviewed presentations made by both collaborations on their individual projects. There were also presentations by representatives of International ATLAS, International CMS, the CERN Lab6ratory where the experiments are located, and the European Data Grid Project. The evaluations of the experts are contained in this report. At the review, they asked many questions of the presenters and provided many recommendations to the collaborations and to the agencies.

In addition to plenary sessions, four breakout sessions were planned in which more detailed presentations of a highly technical nature were made to four sub-committees, each consisting of various members of the review committee. The sub-committees were concerned with: (a) Grids and Common Projects, (b) Facilities, (c) Software, and (d) Project Management. However, after noting the large overlap in subject matter, the review committee decided to combine the breakout sessions on Grids and Common Projects with those on Facilities. This report also reflects that combination.

Observers from the funding agencies were present at all meetings and participated freely both in the open sessions and in the executive sessions. The review was chaired by Alexar~der Firestone of the NSF with much assistance froin Irwin Gaines, Moishe Pripstein, and Michael Procario of the DOE and James Whitmore of the NSF. Anna Smith, Cathy Thompson and Stu Loken of LBNL 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 shown in Appendix C. The first day was devoted to presentations on Grids and Common Projects. Reports were heard from the CERN Research Director, several representatives of the CERN LHC Computing Grid Project, and the European Data Grid Project. Separate presentations were made for the US ATLAS and US CMS computing efforts totaling one-and-a-half days for each experiment. Additional material prepared by US ATLAS and US CMS, including cost and schedule tables, status of milestones, and slides of many

of the talks, may be found at the experiment web sites for software and computing. These are:

http://www.usatlas.bnl.gov/computing/meet/030 1 AgencyReview/
http://www.uscms.org/sandc/reviews/doe-nsf/2003-0 1/

This report, including its recommendations, represents the views of the committee members on issues raised during the review, but it does not attempt to portray the personal opinions of every reviewer nor to 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 US 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 start of operation in 2007. It will provide two colliding 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 most interesting proton-proton collisions. They will be among the largest and most complex devices for experimental research ever built, and the events that they record 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 types of sub-atomic particles.

The US scientific community strongly and repeatedly endorsed US involvement in the LHC program. Numerous groups of US 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 (JOG) to coordinate and manage these efforts and to negotiate an appropriate US role in the LHC program.

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

US physicists are participating in many aspects of the detectors, including important management roles. With more than 300 physicists from over 30 US universities and four national laboratories working on each of the two large detectors, the US groups comprise over 20% of the internati~onal collaborations, and the US 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 US LHC research program must be successful if the

US HEP community is to reap the scientific benefits of the US investment in the LHC. In addition, the international scientific community is depending on the US to contribute 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 US LHC Research Program, including software and computing projects that will be central to generating the physics results over the lifetime of the experiments.

A second comprehensive review of the US LHC software and computing projects was undertaken in November 2001 that concentrated on technical, cost, schedule, and management plans for the near4erm efforts up through fiscal year 2003 of both collaborations in developing software and user facilities for ATLAS and CMS.

The US LHC Research Program is a joint effort of DOE and NSF, utilizing the oversight structures established for the US LHC Construction Project, as detailed in the DOE/NSF Memorandum of Understanding concerning US participation in the LHC program, as modified to include the US LHC Research Program. In particular, this report is the result of the third formal "baseline" review of the Software and Computing Projects of both US ATLAS and US CMS. This review was conducted in a manner analogous to the DOE/NSF reviews of the US ATLAS and US CMS Detector Construction Projects. However, due to the dynamic nature of the software and computing fields, and the rapidly evolving technology, 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 (2-year) project efforts, and to evaluate the technical accomplishments in the first full year of baselined operations.

3 US ATLAS

Grids

The US ATLAS grid projects are fully engaged in the process of developing a stable, reliable, and manageable grid environment. They are working with the LHC Computing Grid Project (LCG) to coordinate development of a common grid software suite that can be used by the LHC experiments. The first release of LCG is not yet available, and its components are not yet fully specified, but it already includes key packages from the existing US Grid implementations (VDT for instance). Since LCG is placing a priority on stability rather than functionality, the first release is likely to be conservative. In the meanwhile, the US ATLAS developers have been pursuing grid development and depl6yment with other US grids research projects. They are planning to create a level-2 mana~ger position for distributed computing applications.

The US ATLAS team has developed many tools over the last two years including GridView, Gripe, Magda, pacman, Grappa, GRAT, Gridsearcher, GridExpert, and a VO toolkit. These tools are aimed at monitoring and managing the Grid environment and managing the data and metadata. In this work they are coordinating with a large number of groups including GriPhyN, iVDGL, and PPDG. The Virtual Data Toolkit (\[DT) is the standard installation of Grid tools and software US ATLAS is using and inclusion of this toolkit is currently planned for LCG. US ATLAS is also beginning to integrate the Chimera Virtual Data System.

A US ATLAS team member, T. Wenaus, holds a key position as LCG Applications Area Manager. US ATLAS also has representation on the Project Overview Board (POB), Grid Deployment Board (GDB), and the Software and Computing Committee (5C2). They are not completely happy with their lack of individual representation on the GDB. US ATLAS personnel expressed concerned that there has been a divergence between the US and EU grid services and that there may be some upcoming issues with regard to LCG. In particular there is a concern that the LCG might define a single software suite and installation. If this leads to a requirement that all sites participating in a data challenge must run that set of software, then the flexibility to use available off-project resources might be diminished. There is also significant uncertainty in the requirements LCG- 1 will place on Tier 1'5, and there has been a divergence in the site requirements.

US ATLAS participated in data challenge I (DC 1), phase 1 in August/September 2002 and generated 200,000 events using 30,000 CPU hours and 100 GB of storage. They also generated 75,000 SUSY and Higgs events for DC 1. They are currently working on phase 2 of the DC 1. They are moving away from the current low-level Globus calls they originally used and are migrating to the higher-level calls available from DAGman, Condor-G and Chimera. During DC1 US ATLAS was able to make opportunistic use of

several university and lab resources outside of the Tier 1 center. They have also completed some interoperability tests with CMS and DO grid sites and EDG sites.

Findin~s and Evaluation

The project is aggressively pursuing evaluation, testing, and deployment of grid middleware components. US ATLAS is leveraging existing products within the grid effectively and has shown good acceptance of the LCG selected software as the ultimate platform for the US ATLAS grid software. We applaud their efforts to influence and contribute to LCG selected software in terms of developments and decisions, and note in particular the inclusion of pacman into LCG- 1. We congratulate the experiments for integrating the grid developments into their work in a serious way. The current direction of the LCG selected software and configuration is good but its motion toward definition of a homogeneous, potentially rigid, environment has many possibly negative repercussions in terms of site configuration.

The integration of grid middleware and grid products into the ATLAS software suite introduces a dependence on that software. The committee noted that the funding agencies need to be encouraged to make monetary commitments to provide on-going maintenance and development support of the grid middleware.

It was noted during the presentation that there has been a loss of base funding at LBNL for the Grid/Athena integration efforts. This integration effort is a critical component of the software development effort. The committee expressed concern regarding the loss of this effort.

There were several talks and discussions mentioning that US ATLAS feels that they would prefer to see two US representatives on the Grid Deployment Board, one from ATLAS and one from CMS, instead of the current one. It was the feeling of the committee that although there are advantages to having a voice in the room advocating specifically for the US ATLAS prograrn, there is also an advantage to having a single voice advocating for both US ATLAS and US CMS. The current coherent US CMS and US ATLAS approach to the GDB was viewed by the committee to be advantageous to the US.

Distributed computing applications are notoriously hard to tune and debug since there is not one place to monitor and track progress but instead there are many. Efforts are currently underway to develop grid monitoring capabilities to allow performance tuning and debugging. The committee applauds these efforts to develop monitoring capabilities at both the application and resource levels.

US ATLAS is making good progress toward establishment of the prototype Tier 2 centers. Efforts to establish MOU's with the iVDGL Tier 2 centers are underway currently. The prototype Tier 2 centers have shown a strong commitment to supporting the needs of the ATLAS project. We encourage continued efforts to get the prototype

Tier 2 centers into operation and are very happy to see and encourage the efforts to establish MOU's with the iVDGL Tier 2 centers. We were also impressed by the commitment to interoperability between the Tier 2 and the Tier 1 site and we encourage continued efforts in this direction. Interoperability between the ATLAS Tier 1 center and RHIC Computing Facility should be retained and has the potential to help US ATLAS contribute to data challenges. Interoperability between Tier 1 and Tier 2 centers and smaller university sites and interoperability between US ATLAS and US CMS Tier 1 centers also offers the opportunity to leverage scarce resources during data challenges, if the data challenges for ATLAS and CMS can be scheduled so that they don't overlap in time.

US ATLAS would like to have a few people based at CERN to be directly involved in the software development work there. We feel that this should be encouraged, as it will make the US involvement more efficient and would allow a faster and more accurate communication between US-ATLAS and Europe.

Recommendations

- US ATLAS should continue efforts to push for the ability to use heterogeneous software configurations ~articularly with respect to grid middleware) within the ATLAS computing facilities. We recommend continued efforts to push the goal of interoperability rather than homogeneous Tier 0 / Tier 1/ Tier 2 sites.
- US ATLAS needs to form a coherent strategy toward development and deployment of grid software taking into account the need to eventually run the LCG selected software and configuration. We encourage the development of a matrix of grid software identifying the interdependencies and each component's place within the overall software framework. A coherent contingency plan for each component should be developed.
- The Athena software is a critical component of the ATLAS software and to have it not integrated with the Grid environment will soon become a major problem. Funding for the Grid/Athena integration should likely be a priority if additional funding becomes available.
- The prototype Tier 2 centers provide an excellent location for some of the prototype software developments and deployments. Consideration should be given to allowing these centers to operate in a less than production mode consistent with their charter.
- We encourage development of an explicit policy and a testbed consisting of dedicated hardware for testing of prototype configurations and new software before it is used in production.
- We recommend continued efforts to develop integrated monitoring capabilities that provide end-to-end information. We also would like to encourage ATLAS to advocate the development of common monitoring capabilities with the LCG.

Facilities

Summary

The US ATLAS user facilities have to provide two forms of deliverables. In keeping with the overall LHC computing plan for three tiers of computing facilities they provide a cache for the event summary data (ESD) and compute cycles for individual users and production tasks for International ATLAS. In addition they support the physics analysis of US ATLAS physicists. Currently, the facilities comprise the prototype Tier 1 computing center at Brookhaven National Laboratory and two prototype Tier 2 centers at Indiana University and Boston University, which are funded through the iVDGL project. A third prototype Tier 2 center at the University of Chicago will be added in FY03. The selection of five final permanent Tier 2 centers will begin in 2004.

The group presented the ramp-up profile for the development of the facilities, aiming to meet the goals of future data chaHenges (DC), within the current agency guidance for funding. The ATLAS plan incorporates a reduced staffing of the Tier 1 center of 20 FTE's, relative to last year's plan of 25. This reduction is based on experience gained during DC 1 and depends on continued synergy with the RHIC computing facility. This significant reduction of personnel cost within the facilities project allows achievement of the goals for the upcoming DC's. However, within this plan the minimal ATLAS requirements will be met with almost no headroom. No grid or networking support can be provided at these staffing levels.

In the past year, the US ATLAS group achieved an impressive success in the completion of Phase 1 of data challenge 1 (DC 1). US ATLAS facilities processed 14 % of the total number of events. This was achieved by opportunistically taking advantage of the RHIC computing facilities and by making use of the prototype Tier 2 centers. The group appears to be on track preparing for the next major goal, DC2, now scheduled for late 2003/early 2004 with a scope of 10 M events, but they expect added CPU and storage requirements due to the first use of GEANT 4. Due to budgetary constraints the ramp up of both equipment and staffing level are behind the original November 2000 plan by 1.5 - 2 years and 1 year respectively. Nevertheless, there is every reason to believe that the data requirements of DC2 will be met by making use of resources in the grid test bed. However, the goal to test the Tier 1 center at levels of increased complexity will not be met.

The fully disk resident computing model that was first presented during last year's review was discussed again. This model promises a large increase in physics productivity but does not fit into the current agency guidance. It requires additional funds in 2007 and 2008. In contrast to the large increase in physics productivity, the expected increase in overall project cost is very modest.

Findin~s and Evaluation

Status of Tier 1 Facility at BNL

The Tier 1 facility is being hosted at the Brookhaven National Laboratory adjacent to the RHIC computing facility (RC F). While expected to be of comparable size to the RCF in 2008, it is currently operational at 1% of the level required in 2008. The major hardware resources available as of November 2002 are 11.4 TBytes of FibreChannel Raid disks and 68 Intel Dual Pentium systems running Linux and one HPSS server with one 9940B tape drive connected to the RCF tape silo.

Currently, the facility provides a near complete Tier I functionality (at much reduced scale and complexity) with 4.5 FTE~s on project. This is possible by taking advantage of the synergistic relationship with the RHIC computing facility. In practice this happens through a great deal of shared expertise and shared operational activities in the areas of operating systems, mass storage, and networking and cyber security. There are also shared infrastructure components such as the tape robots, the firewall, and the WAN connection.

The facility is most notably used for simulation of neutron backgrounds and for detector studies, in particular for simulation of the EM calorimeter response. It participated successfully in data challenge 1 (DC 1) phase 1. DC 1 phase 2 is currently underway. After completion of DC lit will serve as one of a handful of primary data repositories for the produced data and consequently it will also serve as a primary analysis site. The resulting utilization may strain the budgetary constraints but is clearly an excellent exercise for the facility.

Based on the experience from two cycles of production operations for RCF and on two years of operating the combined RCF/US ATLAS Tier 1 facility and relying on continued synergy between the two facilities the staffing level estimate for the full facility has been significantly reduced from 25 FTE's down to 20 FTE's. The ramp up to this level is expected to be slower than linear

The committee commends the group on the reworked final staffing plans (20 down from 25). However, the committee believes that this plan is very lean. The committee is concerned that there will be minimal ability to do network and grid troubleshooting.

Status of Tier 2 Facilities

In addition to the Tier 1 facility at BNL the US ATLAS computing grid will comprise five permanent Tier 2 facilities, which are scheduled for selection at the beginning of 2004. Currently there are two prototype Tier 2 sites at Indiana University and Boston University, which are funded through the iVDGL project. These centers concentrated on prototyping and on participation in the grid production demos at the 5C2002 cbnference. This demonstration used the standard ATLAS DC production environment but was not

part of the official DC production in order to avoid schedule conflicts. Major hardware purchases are taking place now in order to be ready for full participation in DC2.

At Boston University ten nodes are currently available for prototyping. Thirty-two nodes are being purchased now and another 32 nodes will be purchased at the end of 2003 just in time for DC2. One hundred nodes can be borrowed for periods of 3 - *5* weeks. One TB of raid disks and 150 TB of tape (shared) are available. The WAN network connection is
0C12.

At Indiana University 16 nodes are currently used for prototyping and tests. Another 64 nodes will be purchased in 2003 for DC2. Another 100 nodes can be shared. 2 TB of raid disks and HPSS mass storage (shared) are available. The WAN network connection is
0C12.

Data Challen~es

ATLAS is currently conducting data challenge I (DC 1), the second in a series of three originally planned data challenges narned DCO, DCl and DC2. The first phase of DC1 has recently been completed and consisted of the generation of 10 million simulated events for high level trigger (HLT) studies and for physics analyses. Of these 10 million events US ATLAS with 7% of the available CPU capacity generated 14% of the events. Phase 2 of DCl is currently underway, during which the US ATLAS Tier 1 center will serve as one of a handful of data repositories. These data are expected to be accessed regularly by approximately 100 physicists working on the HLT TDR and various physics topics.

The next data challenge, DC2, originally scheduled to start early in 2003 will start at the end of 2003 or early 2004. It has been descoped from 100 M events to 10 M events, the same number of events as in DC1. These adjustments reflect the overall slippage of the LHC schedule. Following the recommendation of last year[1]s review, two new data challenges have recently been defined, DC3 (500 M events) starting late 2004/early 2005 and DC4 (1000 M events) starting late 2005/early 2006.

For the first time GEANT4 will be used in DC2 possibly leading to an increase in CPU and storage requirements. The US ATLAS group very conservatively estimates these increases to be a factor of two. Since the budgetary constraints will not allow further scaling up of Tier 1 facilities in 2003, the required additional resources will have to come from other participants in the US ATLAS grid test bed. While the group sees the use of grid connected resources as an important additional goal of DC2, the fact that the Tier 1 center cannot be tested at levels of increased complexity is seen as a problem.

The committee applauds the group on the impressive execution of DC1. There is every reason to believe that the group will be able to execute subsequent data challenges within the current funding guidance. In particular, DC2 is on track since its scope has been reduced from 100 M to 10 M events.

Analysis Model/Tier 1 Architecture

The default facilities model for International ATLAS assumes that each Tier 1 site will store one third of the current year's and one sixth of the previous year's ESD on disk. This allows access to the complete data set of the current year by accessing three Tier l's and of the previous year by accessing six Tier l's. Already during the previous review in 2002 the US ATLAS group proposed to move from this model to a model where 100% of the current year's ESD is available on disks.

This transition requires additional disk and CPU capacities but a reduced I/O bandwidth for tapes. The US ATLAS group estimates a 33% increase in overall disk capacity, a 50% increase in CPU capacity and an 80% decrease in tape bandwidth. A factor 2-3 increase in the network bandwidth requirement is expected from the exploitation of the model. The latter may be covered by a faster development in networking technology than previously estimated (0C768 in 2008). Costing this in the US Tier 1 context produces the following cost balance: $576k + $53 8k - *$650k* = $364k. It should be kept in mind, however, that these numbers strongly depend on current assumptions for trigger rates and luminosity. Since there is no headroom in these estimates, e.g., a 20% increase in trigger rate 6r luminosity would mean that ATLAS would have to fall back on the 1/3 data on disk model.

The committee strongly believes that the analysis model with 100% disk resident ESD would greatly enhance the physics capabilities of the US ATLAS group. It would be achievable with a very moderate funding increase of roughly 10% in integral Tier 1 cost.

Plans for FYO3/04

The primary focus for FYO3/04 is the completion of DC1 and the preparations for DC2. At the current budget guidance the staffing level will remain at 4.5 FTE's. Any Tier 1 equipment growth depends on supplemental end-of-year funding. Staffing growth is currently 1.5 - 2 years and equipment growth is 1 year delayed relative to the November 2000 plans. This delay will flirther increase in 2003.

Recommendations

- In view of the reduced staffing levels that are now planned for the Tier 1 center, plans should be developed for how the required grid and networking support and maintenance requirements can be covered.
- Continued attention must be paid to providing adequate network bandwidth and support.
- Every effort should be made to identify the additional funds necessary in order to realize the fully disk-resident ESD model.

Software

Observations

US ATLAS contributes to core software, database management, software support for development and analysis, and subsystem software related to hardware responsibilities.

ATLAS is well-coupled to LCG activities. ATLAS has established a substantial architectural design presence in the LHC-based grid projects, while its Project DB activities have been adopted by/transferred to LCG. Architecture/framework effort is being sustained. US ATLAS is delivering baseline core software to ATLAS.

Torre Wenaus, as manager of the LCG Applications Area, will have considerable influence over these common developments. David Quarrie is about to become Software Project Leader, further cementing the group's high profile in core software. Similarly David Malon is the Database Leader for ATLAS. Much of ATLAS's expertise on the hybrid ROOT/relational database persistency solution resides in US ATLAS. The nightly build system, pioneered by US ATLAS, remains a valuable day to day tool for the developers and has been adopted by ATLAS.

Torre Wenaus is the Planning Officer for ATLAS, resident at CERN. In addition, he is the software manager for US ATLAS. It appears that Torre will hand off the Planning Officer position this year.Some 17.9 FTE's are involved in the software effort, with 11.50 coming from Project funds.

We have seen good evidence of improved decision making in ATLAS, an example being the termination of the FADS/Goofy project - which was a significant diversion at the last review.

Since the last review, the LHC schedule has been delayed by two years. This has eased some of the pressure on US ATLAS software related tasks, allowing some stretch of their own schedule.

The R&D program has been initiated and the M&O and Computing elements have been consolidated into a single "research program", sharing the overall budget. The division of funding is determined by the US Project Manager.

The leadership role in analysis is thought to be contingent on strong US database management skills. Leadership in core software leads to the good of ATLAS as a whole. US ATLAS continues to be a source of firefighters. This continues to result in a broadening of scope for the US ATLAS team.

The US ATLAS base program support from the national laboratories has diminished:

- ANL:-0.5FTE BNL:-1.OFTE

LBNL: complete loss of base support threatened

US ATLAS and ATLAS are lacking CERN presence of software engineers. The problem could be relieved by adequate funding for CERN-based US ATLAS personnel. The LCG SEAL project (Core Libraries and Services) is supposed to grow from 3 to 8 FTE's, of which 0.5-3 FTE's would be from US ATLAS.

Findin~s

The core software activities will be noticeably impacted by LCG work. The SEAL project is likely to be an ATHENA-like (component) framework replacement, while POOL (hybrid ROOT/relational database) will be the common persistency mechanism. POOL is in fact an outgrowth of the ATLAS persistency solution. These are of course opportunities and problems - opportunities in leveraging LHC manpower to ATLAS problems, but the startup pain of migrating to them.

It is encouraging that US ATLAS has consolidated leading roles in the ATLAS core software areas. This will hopefully ensure a closer coupling and alignment between international project activities and US interests.

The adoption of the ATLAS framework for DB (POOL) and the adoption of pacman as the grid software deployment and configuration tool helps to guarantee that US ATLAS needs can be partially met by LHC-wide cooperative efforts.

US ATLAS appears to be well-coupled to LCG activities. This is critical because the US ATLAS data analysis effort is particularly sensitive to the quality of data management by ATLAS. The US needs extensive ESD access on this side of the ocean in order to ensure that US physicists will be able to participate and share in the discoveries promised by LHC. Yet, there is concern that the ATLAS overall effort is not well represented within the decision-making organs of LHC and a number of critical ATLAS milestones and tasks have been transferred to LCG's scope (e.g., the adoption of the DB framework by LCG and the replacement for ADL).

The consolidation of both M&O and Software & Computing elements into a single research program represents a potential source of tension in the manner in which any budgetary shortfalls will be apportioned between these two critical activities.

ATLAS's aggressive pursuit of alternative/additional sources of funding options is to be commended. Hopefully, successful proposals will lead to support that can supplement agency funding levels. The concern here is that constraints and limitations on how funding from different sources may be used will unduly inhibit solving US ATLAS problems. The exigencies of using grid funding have tended to distort US ATLAS priorities for the delivery of project-specific software.

ATLAS goals and definition of success in DC 1 effort has been scaled back substantially, to the point where some of the accomplishments do not reflect scalable accomplishments

for the LHC science run - e.g., DC 1 still used the Zebra persistence framework for data (instead of the Athena model).

ATLAS base program support from the national laboratories has diminished at an alarming rate. The core software activities are not attractive to academic institutions and require the support of professionals and dedicated software physicists that are typically only available from national laboratories.

Continuing lack of US ATLAS and ATLAS presence at CERN is a serious concern to the committee. The core software effort led by US ATLAS needs to provide EU time-zone support to EU users. Lack of such support hinders progress and could threaten unity of approach as frustrated users start to invent their own solutions.

Scope has continued to increase for US ATLAS contributions to ATLAS. This threatens the ability of US ATLAS, and more generally ATLAS, to accomplish the required SW milestones in order to support a successful 2007 start up of LHC ATLAS data-taking. Lack of sufficient data management software will affect US ATLAS to a greater extent than EU-based collaborators in the broader International ATLAS collaboration. If things continue to deteriorate, the US ATLAS failback position may need to be to situate US scientists at CERN in order to analyze the data.

Recommendations

- US ATLAS acts as a fire-fighting organization within ATLAS. This is viewed as good and necessary, but US ATLAS must take care to properly plan across its operation the impact on US-specific efforts of taking on new central tasks.
- ATLAS in general is under-represented at CERN, and US ATLAS in particular. We endorse the US ATLAS request for three software engineers to be resident at CERN. These people would act both as support for David Quarrie and as ambassadors for the US-created products (e.g., ATHENA) to the ATLAS collaborators in Europe. This would enhance the likelihood of these products receiving general acceptance in the collaboration.
- While the committee concurred that the emergence of a common projects initiative within CERN to support all LHC experiments was a good development, nonetheless, it is concerned about the necessary reliance of ATLAS on the newly formed LCG. This risk is ackhowledged by all parties, and there seems to be the will to make it work. There must be continued detailed tracking of the LCG progress and contingency plans in place to meet major milestones if the LCG effort flags.
- Continued search for NSF funds through the ITR program appears necessary in order for ATLAS to achieve its programmatic goals and address the erosion of effort from the base program.
- Last, the committee is still not clear on the confusing relationships that exist between projects within ATLAS and external grid. We have some concern that ATLAS needs may not be realized by this fluid relationship. How is this effort "guided" (not just tracked) to maintain alignment with US ATLAS needs? US

ATLAS need to prioritize grid projects and divert people from those less important in order prevent schedule slip. There should not be any diversion of individuals assigned to high priority projects to grid projects.

Project Management

Findin~ and Evaluation

Or~anization

The US ATLAS Physics and Computing Project is part of the US ATLAS project. The Project Manager for US ATLAS Physics and Computing is an Associate Project Manager in the US ATLAS organization. This tight and explicit connection is even more important now that the Computing project has been combined with experiment Maintenance and Operations into a single Research Program. Coordination of the US computing effort within the full ATLAS collaboration is managed through a small number of individuals from US institutions who have key roles in the International ATLAS computing organization (Chief Architect D. Quarrie, LBNL and Planning Officer T. Wenaus, BNL/CERN). US ATLAS Computing is also represented in the Computing Steering Group (J. Shank, Boston University and D. Malon, ANL). In addition, the Project Manager for US ATLAS Physics and Computing (J. Huth, Harvard) is a member of the ATLAS National Computing Board.

The US ATLAS Computing Project also interacts with several groups working to develop tools for Grid-based computing. Liaisons for several of these groups are named in the US ATLAS Computing Project Management Plan. This plan, dated 21 November 2001, does not mention interaction with the recently-formed LHC Computing Grid (LCG) project. Coordination with LCG is through T. Wenaus (BNL/CERN) who is LCG Applications Area manager and a member of the LCG Project Execution Board. Also, J. Huth is a member of the LCG Project Overview Board.

The ATLAS Collaboration has recently appointed a new Computing Coordinator, D. Barberis. He presented a proposal for a new organization for ATLAS computing. The proposal includes a new Software Project Leader position. D. Quarrie is proposed to fill this role. The committee would like to emphasize that it is a great benefit to the US ATLAS Computing Project to have personnel from US institutions in key positions in the International ATLAS computing organization.

Managing the US ATLAS Computing project, connected as it is to all these other entities, is challenging. US ATLAS is fortunate to have a group of talented and dedicated people who are working hard to make the effort succeed. The committee strongly supports US collaborators taking on visible leadership roles in the international computing effort, despite the increased risk of strain on US resources. These US ATLAS members will help coordination with the international computing effort, and ensure that US efforts are directed toward the ultimate goals of analysis of first LHC data and support of US

pbysicists in ATLAS. The US ATLAS Computing management sets the assignment of additional personnel to CERN as a high priority, and the Committee agrees with this.

Overall, the management appears to be functioning well. In particular, Committee members who have served on previous reviews note that there has been noticeable improvement in this area. The Committee, however, felt that there was room for improvement in the presentations prepared for this review. In many cases, the presented material was too long (leaving insufficient time for questions) and was not focused on the critical issues. This made it difficult for the Committee to evaluate progress.

Scone

The US ATLAS Physics and Computing project has a work breakdown structure with three Level 2 subprojects: Physics, Software and Facilities.

Physics subproject

The Physics subproject supports event generators, simulations and algorithms for physics objects. This task also contributes effort to the ATLAS Data Challenges. The committee did not hear a presentation on these activities due to lack of time, but slides from the Physics Manager were posted to the review web site. A support person for Physics activities was hired in May 2002. Work on support of generators appears to be going well. Contributions of US ATLAS to the Data Challenge effort are mentioned in the slides as an issue.

Software subnrolect

The software task consists of providing deliverable software packages to US ATLAS, International ATLAS and CERN. The scope is specified in software agreements that will eventually be combined into a Memorandum of Understanding. There is currently one software agreement in place for the control framework. This project uses about 30% of the software FTE effort. US ATLAS software personnel are also working in the areas of data management, grid software, event model and simulation, and reconstruction infrastructure. The scope of these efforts is in some flux because of dependence on LCG deliverables. US ATLAS is well connected to the LCG process, efforts, and leadership, but does not have complete control over LCG software development.

Allocation of US ATLAS software personnel is decided by US ATLAS managers with roles in the international ATLAS computing organization (Wenaus, Malon and Quarrie).

Facilities subi,roject

The Facilities subproject supports the US ATLAS data analysis effort and also provides computing power for the ATLAS Distributed Virtual Off-line Computing Center. Current plans call for a Tier 1 center at Brookhaven National Laboratory and five Tier 2

centers throughout the US. There are two prototype Tier 2 centers active now as well as seven Tier 3 centers in the testbed. The analysis model currently has all ESD on disk.

Cost and Fundin~

The project managers presented a profile indicating a level of funding desired for the project. The sum is $66.5M (not counting money in iVDGL) spread over six years (FY 03 - FY 08). The project managers have also constructed a profile based on September, 2002 agency guidance. The sum is *$52.6M* (also without iVDGL) for FY 03 - FY 081. The difference between the two profiles has been divided roughly equally between Facilities and Software. Confronted with a potential funding shortfall, the project managers plan to keep existing personnel supported and to support as much growth as possible in facilities in order to participate fully in the data challenges. The impact of the difference between the larger profile and the guidance would be:

- Software - loss of some university-based FTE's. This would result in reduced subsystem-specific support.
- Facilities - slower ramp and reduced level in FY 07 and FY 08.

The project managers also indicated that cuts in infrastructure at CERN (support help desk, librarian, etc.) would be required with the guidance profile.

We also note that the US ATLAS Software and Computing project has some software personnel funded from the "base program" at national labs. There have been some problems sustaining this effort due to shortfalls in the "base program" funding.

The committee did not examine the US ATLAS Computing Project cost estimate in any detail.

Schedule

The project has extensive detailed schedule information including 600 milestones, but it is not clear how the information is used to monitor progress. Overall, the schedule is constructed to be consistent with the start-up of the LHC. There are trade-offs between the scope and the schedule. In particular, the schedule is built around a series of Data Challenges. If problems are encountered, the managers may choose to delay the Data Challenge, or to proceed with it with reduced functionality. It was difficult to see how these trade-offs would impact the overall schedule. The US ATLAS Computing web pages contain lists of milestones with completed milestones marked, but there is no baseline information on these pages 50 one can't tell whether or not the milestones were met as scheduled

This was an area where the material presented was difficult for the Committee to evaluate. Gantt charts may not be essential for planning for US ATLAS Computing, but

We note that there was some confusion at the time of this review about the magnitude of the September 2002 agency guidance.