U.S. ATLAS Western Tier 2 Center

May 15, 2006

Executive Summary

The Stanford Linear Accelerator Center (SLAC) proposes a Western Tier 2 Center for ATLAS. This proposal has strong support from university-based ATLAS collaborators in the Western region, and from the Lawrence Berkeley National Lab (LBNL). SLAC brings long experience in managing large, data-intensive experiments and international collaborations. Substantially more computing power is available on a shared basis at SLAC than can normally be provided by a Tier 2 center, along with robust network connections and a large and experienced staff. We feel that this will make a valuable contribution to the ATLAS experiment.

1 The Western Tier 2 and the Western Tier-2 Community

1.1 Introduction

The exciting challenges of the ATLAS physics program demand an approach to computing and physics analysis that is coherent on the international, national and regional scales. The Western Tier 2 Center must be an integral part of the overall ATLAS computing system. We intend that it should also provide a regional facility encouraging the growth of a strong community of physicists and computing experts able to brainstorm on the most fruitful approaches to ATLAS physics analysis. The Western Tier 2 Center will have a strongly synergistic relationship with the LBNL Analysis Center, providing responsive data analysis resources to the Analysis Center's community.

LBNL, University of Arizona, UC Santa Cruz, UC Irvine, University of Oregon, University of Wisconsin Madison, and University of Washington have been actively involved in shaping this proposal, and are keen to participating in the scientific oversight of the Western Tier 2. Participation in this oversight will be extended to additional university groups in U.S. ATLAS that have a strong interest.

The Western Tier 2 Center will be located within SLAC Scientific Computing and Computing Services (SCCS) at the Stanford Linear Accelerator Center and will leverage existing and planned investments in the programs at SLAC such as *BABAR*. This leverage will provide computing that is good value for money, but much more importantly, it will build on a successful history of international collaboration in computing, and will bring together physics and computer science intellect at SLAC, LBNL and the Western universities to facilitate outstanding physics analysis.

This proposal builds on a history of close collaboration with other institutions in the San Francisco Bay Area and more broadly, with the universities principally, but not exclusively, in the Western United States. SLAC's history of partnership with LBNL significantly strengthens the link to the ATLAS Western Regional Analysis Support

Center. The physics groups at the Universities participating in the oversight of the Tier 2 provide strong coupling to the requirements for university groups and vital feedback on operations and performance issues to ensure that the center provides outstanding service to the U.S. ATLAS community.

The long-standing partnership between LBNL and SLAC is the foundation for a close collaboration between the Analysis Support Center at LBNL and the Tier 2 Center at SLAC. The high-speed network infrastructure between SLAC and LBNL will support the analysis needs of physics users at the Analysis Support Center. We note also that the DOE's leading computer-science funded data management group is at LBNL and that this group has an excellent track record of collaboration with the high-energy physics community in PPDG and other projects.

Tier 2 centers have a well-defined role in complementing the ATLAS Tier 1 center at Brookhaven National Lab (BNL). The aggregate storage and computing power at the Tier 2 centers will be required to meet the needs of ATLAS analysis. However, a Tier 2 can have a much more important role than just filling gaps in aggregate capacity. The Western Tier 2 Center aims to be a responsive focal point for its community, and to leverage significant local computing and intellectual resources.

The Tier 2 Center will be designed in accordance with specifications of the ATLAS Computing TDR[1] and the U.S. ATLAS Distributed Computing Facility call for proposals. The ATLAS computing model calls for Tier 2 centers to provide facilities for:

- Fast and detailed Monte-Carlo event generation, simulation and reconstruction;
- Data processing for physics analysis, generally this will be done in a chaotic way;
- Data processing for calibration and alignment tasks, code development and detector studies.

The Western Tier 2 will provide these facilities for U.S. ATLAS, to be used under the overall direction of U.S. ATLAS management. SLAC believes that by promoting community involvement in the operation of the Tier 2, its effectiveness in potentially chaotic physics analysis and in calibration, alignment and detector studies, will be maximized.

1.2 Viewpoint of a typical university-based group

Computing support is essential for university groups to make a significant contribution to the physics analysis in ATLAS. In particular the following tasks must be supported:

- Extraction/selection of events of interest from a pool of simulated and experimental Event Summary Data (ESD, likely standard ATLAS data);
- Production of specific Analysis Object Data (AOD) from these events;

[1] ATLAS Computing TDR: http://doc.cern.ch//archive/electronic/cern/preprints/lhcc/public/lhcc-2005-022.pdf

- Reasonably quick simulation response to solve/understand very specific analysis problems not addressed in the (possibly already existing) data in the standard ATLAS simulation pool;
- Capabilities to (at least temporarily) store larger amounts of events outside of the official ATLAS data streams at ESD and AOD level for analysis (both formats) and reconstruction software development support (mainly ESD);

Clearly, the computing support must include access to a batch processing farm for running very high statistics analysis jobs, in addition to significant storage capabilities and high speed networking access around the clock.

Close proximity in distance and time to a Tier 2 center is important for most university groups. Some of the obvious reasons are:

- Quick response to technical problems due to largest overlap of typical working hours (maximum difference in time zones two hours ahead, which produces more overlap compared to up to three hours behind the East Coast, and nine hours behind Central Europe);
- Contributions to a User Advisory Group or any managing, advising, or guiding body of the center are important for smaller groups as well. They are possible within a relatively small travel effort;
- Expected better recognition of specific physics interests and the corresponding computing needs within a smaller and more regional user community, including faster and generally more positive response to occasional short term requests for more-than-usual resources for specific simulation production, for example, as those could be presented and discussed efficiently in (regular) personal resource planning meetings;
- Close collaboration with specialists resident at the Tier 2 center in software development issues helps university groups to contribute to general physics analysis support software and even important framework software development.

Some of these use cases and usage scenarios are based on long standing practical experiences in all aspects of software development. In particular, intense collaborations between university and laboratory physicists and software engineers have led to efficient development cycles for important improvements of the present day ATLAS framework and physics analysis software. For example, a long-standing collaboration between University of Arizona and LBNL has developed ATLAS framework software, mainly in the areas of data model, object relations and navigation, and job configurations. The same intense and efficient working relations are not easily possible when time zone differences and travel effort are significant.

A Tier 2 center situated at SLAC will also benefit from the synergistic interactions between SCCS and university personnel in the areas of core software design and implementation. Responding to *BABAR*'s needs to foster better communications among its collaborators worldwide, handle large volumes of data and reduce systematic uncertainties in its detector Monte-Carlo simulation, several areas of collaboration between SCCS and university-based members of *BABAR* were formed resulting in:

- The installation, customization and ongoing support for *BABAR*'s web-based communication system, Hypernews. SCCS now also supports ATLAS, CMS, the Linear Collider Detector community and the Geant4 User Forum with the same technology.
- The development and implementation of xrootd to manage *BABAR*'s Root-Based data files
- Major contributions to the development, implementation and validation of the Geant4 detector simulation toolkit upon which the *BABAR* detector Monte-Carlo simulation program is based
- The ongoing development of data distribution software based on Grid technology.

Recognizing the need to reduce the systematic uncertainties in its detector Monte-Carlo simulation, *BABAR* became the first experiment to make use of the Geant4 toolkit as part of its production simulation in 2001. The cooperative development of Geant4 began in 1996 and is ongoing between *BABAR,* SCCS and the Geant4 collaboration. This collaboration will continue in the future and the developed expertise will benefit the ATLAS experiment.

2 Proposal Details

In this section we specifically call out our responses to the criteria set out in the solicitation. Points 1 though 6 in the solicitation are answered by 2.1 to 2.6 below.

2.1 Principal Investigator

Richard P. Mount is the Principal Investigator. He is SLAC's Director of Scientific Computing and Computing Services (SCCS) and has long experience of responsibility for developing, providing and managing computing solutions for high-energy physics experiments. Ten percent of his time will be spent explicitly on the Tier 2 management; 50% of his time overall is devoted to managing similar scientific computing enterprises. In addition Chuck Boeheim, SCCS Assistant Director, and Randy Melen, head of the High Performance Computing and Storage team, will devote 15% and 20% of their time respectively to managing the Tier 2 center. These efforts will be provided by SLAC at no cost to the Tier 2. This team approach provides a depth of experience in managing scientific computing and continuity of coverage at all times. Together this team has a century of experience with high performance computing, and has worked on numerous experiments, including *BABAR*, SLD and L3.

Wei Yang will serve as the principal ATLAS liaison. Wei has a PhD in experimental high energy physics, has worked extensively on computing for *BABAR*, and has been implementing the infrastructure for OSG at SLAC. He will have as his highest priority the smooth running of ATLAS computing at SLAC, and will have access to all necessary resources to carry this out. Wei is listed officially as 30% time on Tier 2 support in section 3.2, however he will be able to devote 100% effort when needed to ensure timely response to ATLAS issues.

If, at any time, there would be a benefit from assigning a SLAC ATLAS physicist as "ATLAS Coordinator for the SLAC Tier 2", SLAC will make such an assignment.

Figure 1 is a functional organization chart of SCCS showing the depth of expertise available in the various support areas. Only functional areas relevant to the Tier 2 are shown.

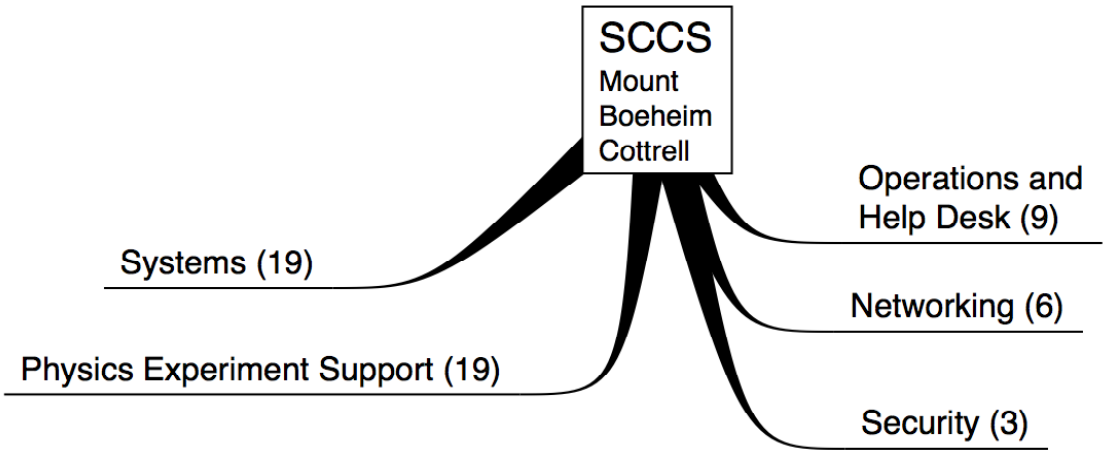

Figure 1: SCCS Simplified Organization Chart. The numbers in parentheses are the FTEs relevant to the ATLAS Tier 2.

2.2 Leveraging Existing Infrastructure

The Western Tier 2 Center will be located within SLAC Scientific Computing and Computing Services (SCCS) at the Stanford Linear Accelerator Center. With one of the most data-intensive experiments to operate to date, SLAC has developed expertise in efficiently storing, moving, and processing enormous amounts of data. The computing center has long emphasized low-overhead lights-out operations to enable it to operate very large numbers of processors and disk servers with minimum overhead and maximum reliability. The *BABAR* experiment was a pioneer of wide-area distributed production and analysis of data even before the advent of grid tools. The infrastructure established by these previous efforts provides access to all needed facilities at minimal incremental cost to U.S. ATLAS.

In terms of computing systems and their direct support, the principal leverage assumed for the period through 2010 is the existence of *BABAR* computing at SLAC. *BABAR* will acquire data until September 2008 and *BABAR* analysis will continue at SLAC for several further years. We assume that equipment bought for the Western Tier 2 will be identical to that acquired each year to provide computing for *BABAR* and therefore calculate hardware costs based on extrapolations of current costs for the equipment acquired for *BABAR*. Given the firm commitment to *BABAR* through 2010, we propose to charge to Tier 2 funding only the additional technical and operations effort needed to extend *BABAR* support to also cover the ATLAS systems – this amounts to one FTE of effort spread over the SLAC Systems Group and Operations Group. As described above, SLAC's established mode of "lights out" operation and support for high uptime computing will allow this minimal effort to provide strongly supported and highly available services to ATLAS at the hardware and systems level.

For most of the period through 2010, *BABAR* will have many times more resources at SLAC than the Western Tier 2. When even a small fraction of *BABAR*'s computing resources are temporarily unused, they will be available to the Tier 2 and are likely to add significantly to its aggregate capacity. SLAC's batch system (Load Sharing Facility (LSF) from Platform Computing) is fully capable of assuring communities guaranteed access to the hardware they have contributed, while allowing other applications access to unused resources. In addition to *BABAR*, the SLAC science program is expected to require a major expansion of scientific computing facilities in the next five years. Many of the new facilities are expected to be usable by ATLAS on an opportunistic basis.

BABAR and the ongoing program of particle astrophysics and photon science will require that SLAC maintain and develop its tape storage systems and mass storage software. The marginal cost of making the mass storage system available to the Tier 2 will be extremely small, but we expect that its value, for example to archive old analysis datasets just in case they need to be reexamined, will be considerable. The ease of archiving and recovery in SLAC's robotic systems will make more ATLAS disk space available for active analysis.

2.2.1 Batch Processing Capacity

SLAC operates a growing batch farm currently consisting of 1500 Compute Nodes, with a total of over 3700 processor cores and an aggregate of approximately 3800 kSpecInt2000 capacity. This is scheduled via LSF as a single cluster serving all SLAC science, with allocation targets for each major activity. This allows each activity to be guaranteed a target capacity, and also the opportunity to use resources beyond that if other groups have not fully used their allocations.

2.2.2 Mass Storage

SLAC has six STK Powderhorn silos with 40 STK 9940B (200GB/cartridge) drives and a number of older drives. This system holds about 2.5 petabytes of data and has a capacity of 6 petabytes with the current technology. *BABAR* requirements are likely to result in upgrades to the mass storage system in 2006/2007 that will increase its maximum capacity to about 15 petabytes.

SLAC's mass storage system uses HPSS to manage the tape store and perform data movement to and from the disk system. Locally-written systems provide means for a number of disk access methods to perform dynamic staging of data on demand, and policy-based migration of data into the tape system. Approximately ¾ of a petabyte of disk storage is currently managed by this mass storage system.

2.2.3 Networking

SLAC is connected to the ESnet Bay Area MAN at 10 gigabits/s and has access to an additional 10 gigabits/s wavelength targeted at data-intensive DOE science. SLAC is also connected at 1 gigabit/s to Internet2 via Stanford University and CalREN2. This connection will be upgraded to 10 gigabits/s in the near future. SLAC owns a 96 fiber pair cable between SLAC and Stanford, making it possible to take advantage of any network that reaches Stanford.

The Internal SLAC network will have a central core linked by 10-gigabit trunks by the end of 2006.

2.2.4 Grid Infrastructure

SLAC has already installed and has operating the basic OSG Grid software, and has interfaced it to local batch and storage systems. Members of different organizations, including ATLAS and CMS, have used the SLAC Grid facility for production scale applications. This prior investment in Grid infrastructure will allow SLAC to begin processing ATLAS workload almost immediately.

2.3 Leveraging Existing Expertise

The technical and operational expertise of SCCS ensures the highest quality service for the U.S. ATLAS physics program. Thirty staff members are devoted to the operation, system administration, networking, user support, and security of the scientific computing complex. Highly automated tools ensure the efficient installation and management of hardware and software resources for client groups. One FTE is included in the proposal as the incremental cost of the support of the servers and storage to be dedicated to ATLAS.

SLAC has built a strong team of computer scientists who have focused on the data-intensive computing needed by the SLAC program. In collaboration with Western U.S. ATLAS partners, SLAC will tune its computing systems to optimize ATLAS data-access performance and will also explore new ways to use computing for ATLAS analysis. SLAC will seek non-ATLAS funds to make general advances in data intensive computing that will be highly relevant to ATLAS. The PetaCache project, an example of SLAC's attraction of non-HEP funding, is itself likely to be of direct relevance to LHC physics analysis. SLAC also intends to exploit the synergies between the many data-intensive parts of its science program such that, for example, work done to develop novel data analysis approaches for astronomy can also benefit the experimental HEP program and ATLAS.

Both ATLAS as a whole and the Western community will also require strong support for the relevant Grid applications and for more ATLAS-specific applications at the Tier 2. SLAC (and LBNL) are founder members of the Particle Physics Data Grid collaboration and are deeply involved in ensuring the success of the Open Science Grid as the prime resource for U.S. high energy physics in the future. In addition to being a functional OSG site, SLAC has led the OSG initiative to develop the accounting software that is required for any Grid to which funding agencies will commit mission-directed resources. SLAC is also a national leader in issues related to computer security.

SLAC has a very strong team of physicists and computer scientists who have worked together to ensure the success of *BABAR* in the face of unprecedented volumes of data. The SCCS Physics Experiment Support Group comprises 19 people with expertise in physics and computer science. As the requirements of *BABAR* for such support decrease, it is SLAC's intention that most of this team turn its attention to the challenges of ATLAS and to the specific needs of the analysis and detector studies being performed by the Western ATLAS community. The ATLAS physicist team includes two former *BABAR* Computing Coordinators and a former Deputy Computing Coordinator. These individuals bring additional experience in coordinating the computing efforts of a geographically distributed collaboration.

SCCS is expected to provide the following support for SLAC science including ATLAS:

2.3.1 OSG development and operations

SLAC currently has a team of three full-time people who support the OSG services at SLAC and contribute to the development of OSG. In addition, the SLAC computer security team has been, and will continue to be, a long-term contributor to the design of Grid security architecture and policies.

2.3.2 Simulation development and user support

SLAC is a major center for development and support of the Geant4 toolkit used for ATLAS simulation. The SLAC contribution to Geant4 amounts to over 4 FTEs, and includes the chief architect/deputy spokesperson, hadronic working group coordinator and several other working group coordinators. SLAC hosts the Geant4 User Forum (Hypernews) service and already provides significant support to the LHC community. SLAC expects that members of its Geant4 team will contribute to the ATLAS simulation code and to developments, such as parameterized showers, that will benefit all LHC experiments.

2.3.3 Data access and data analysis

SLAC is a contributor to the ROOT software. In particular SLAC's xrootd data access software is now part of the ROOT distribution and is a candidate for use by LHC experiments. SLAC's physics experiment support group is exploring data-access technologies based on databases such as MySQL, and on non-database approaches such as xrootd, to meet the needs of the several data-intensive studies that are already part of SLAC's future.

2.3.4 PetaCache

The PetaCache[2] project aims at freeing data-analysis from the huge latency and limited access rate of disks by caching data in low-cost memory. The project has built a 1-terabyte prototype system with very promising performance, and will seek 2006 funding to move to a "development system" of up to 30 terabytes that will be able to cache the most intensely accessed *BABAR* data. In the longer term, we believe that even larger systems, taking advantage of future low-cost memory technologies, will be highly effective for many aspects of ATLAS data analysis. However, as *BABAR* is already demonstrating, taking full advantage of low latency data access requires significant software development and will impact both production and analysis activities.

2.3.5 Network Monitoring

For more than a decade, the Internet End-to-end Performance Monitoring (IEPM) Group[3] at SLAC has been monitoring network connectivity and End-to-end performance for sites involved in High Energy Nuclear and Particle Physics. The members of the group work closely with the Energy Sciences Network (ESnet), Internet2, the U.S. DOE funded

---

[2] http://www-user.slac.stanford.edu/rmount/public/PetaCache-Report-to-june2005-v14.pdf
[3] http://www-iepm.slac.stanford.edu/

laboratories, laboratories throughout the world, and Institutes and Universities throughout the world involved in data intensive science.

2.3.6 Long-Term Goals

The SLAC programs of high-energy physics, particle astrophysics, and in the future "ultrafast" photon science, will all require or benefit from excellence in the analysis of daunting volumes of complex data. Pursuing this long-term goal requires exploiting the synergies between the programs, and seeking funding for new developments that will benefit all programs. As just one example, research into petabyte-scale relational databases for potential use by Large Synoptic Survey Telescope (LSST) is being followed by a number of other research projects for their own use.

SLAC's PetaCache project, funded partly by DOE MICS is one example of how such funding can benefit the SLAC program and advance the science of scientific computing.

Although the scale of the Western Tier 2 is relatively modest, there is nothing modest about the potential for ATLAS analysis driving the development of innovative approaches to the use of computing in this analysis. SLAC views computing for ATLAS analysis as a "Grand Challenge" and will seek opportunities to allow SLAC and the Western ATLAS community to make revolutionary contributions to the success of the experiment.

SLAC intends to expand its scientific computing activities in accelerator simulation, particle astrophysics and cosmology, photon science, and the computer science of data-intensive computing. It is expected that by 2010, the aggregate activity, for example as measured in hardware dollars per year, will be several times greater than the present *BABAR*-dominated activity.

2.4 Physical Infrastructure

2.4.1 SLAC Infrastructure Relevant to the Tier 2

For the period covered by this proposal, SLAC will maintain its current approach of providing substantial components of the power, cooling and space infrastructure at no cost to ATLAS. The required water-cooled racks and the power distribution within the racks will be charged to the Tier 2. All other components of the power, cooling and space infrastructure, including the utility bills for power, will be carried by SLAC. We estimate that the total power load from the Tier 2 will approach 200KW asymptotically, to be compared with SLAC SCCS current total power availability of over 3MW.

The SLAC computer center (building 50) has about 16,000 square feet of raised floor area. The space is about 75% used, and this is expected to remain approximately constant as old equipment is replaced.

Building 50 draws power from two substations providing up to 1MW and 4MW respectively at 480V. Only just over 1MW is currently distributed to the raised floor. Upgrades are planned to use the substations to support 1MW of high availability equipment (dual redundant power sources, UPS backup on one source) and up to 2MW of unprotected equipment (primarily farm CPUs). Currently 0.45MW is available as UPS protected power.

Building 50 is served by an 8-inch chilled water feed from the main SLAC chiller plant. This feed is expected to be adequate to remove about 1.5MW of heat. New installations on the raised floor will be cooled either by installing nearby water-cooled air handlers or, for the densest farms, by using water-cooled rack enclosures. Upgrades are planned to install a dedicated Building 50 chiller system to handle at least 2MW of heat, using the SLAC main chillers as backup.

The Building 50 computer rooms have electronic access control and video surveillance commensurate with modern standards for computer rooms. SLAC's physical and cyber security plans and procedures have passed all relevant DOE reviews, and have been cited by DOE reviewers to other sites as examples of effective management.

In order to provide a cost-effective infrastructure for the large expansion of computing activities that is expected late in this decade, SLAC is actively engaged in planning for a major new facility that may benefit from similar emerging needs of Stanford University.

2.5 Operations Model

Round the Clock Operation. SLAC aims that its computers should operate 24 hours a day, 365 days a year. The track record of achieving this goal is excellent: service availability averaged over ten years has been 99.4%. The dominant cause of outages is power issues, divided equally between planned and unplanned outages (e.g. the 60-hour outage in 2005 due to a tree falling in the Santa Cruz mountains and cutting SLAC's main power feed at the same time as the backup feed was disconnected due to construction external to SLAC). Figure 2 shows the impact on services, including recovery time after the primary outage.

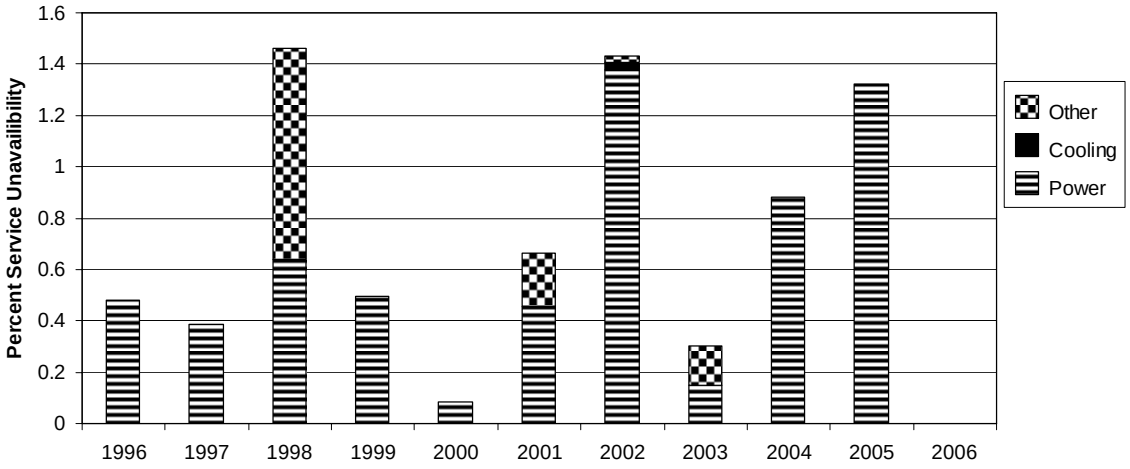

Figure 2: SLAC's ten-year record of outages impacting user services.

SCCS provides minimal guaranteed out-of-hours support (only one person officially on call outside normal working hours). However, this is complemented by explicitly engineering services for appropriate redundancy and remote operability (including remote power cycling) and by a very high level of commitment from staff to respond to out-of-

hours problems. Loss of service availability due to staff response times has been almost negligible in the last ten years.

Support is via an online ticketing system available via email and the web. Response commitments are one business day for routine requests, and best possible efforts for emergency situations. Provisions are made for escalation of priority of requests, up to an immediate paging of an on-duty coordinator at any time.

2.6 Resource Projections for U.S. ATLAS

Table 1 shows the projected resources that would be dedicated to ATLAS and those that would be available for opportunistic use. The assumptions used are

- In response to the call for proposals, we list resource projections through 2010. It is SLAC's intention to continue involvement with and development of the Tier 2 effort for the life of the ATLAS experiment.
- Only one incremental FTE of systems/operations support will be required, such that 2/3 of the funding can be spent on hardware each year.
- For the purposes of this projection, we assume that 2/3 of the hardware funding will be spent on disk and 1/3 on CPUs. This yields approximately the ratio of SpectInt2000s of CPU to TB of disk that is in the TDR. The actual purchases in each year can be adjusted at the direction of U.S. ATLAS to meet actual needs.
- The costs of disk and CPU are extrapolations from the prices paid in 2005 for equipment at SLAC. The baseline CPU-cost assumption is that a 4-core 2.0 GHz Opteron system purchase price was $3000 in 2005 and delivers 4*1274 SpecInt2000. The baseline disk cost assumption is a purchase price of $3800 per useable terabyte in 2005 including the required fileservers. Both assumptions are in line with current experience delivering high-reliability computing.
- Costs for CPU and Disk are increased above the purchase price to include all the hardware needed to be able to offer the disk or CPU services, and SLAC's G&A overhead rate of 6.8% on hardware purchases. Experience with *BABAR* computing costs over more than six years is the basis for estimating the ancillary costs detailed in the following table.

Ancillary Costs	CPU Server	Disk Server
Racks	16.7%	1.6%
Power Distribution	1.7%	0.7%
Networking	16.4%	5.0%
Console Server	1.4%	0.1%
G&A	6.8%	6.8%
Total	34.0%	14.3%

- Unit costs are assumed to decrease by a factor 2 every 24 months as directed in the clarifications to the call for proposals.

- We assume a 4-year lifecycle for hardware. Hardware acquired in year one of the table will begin retiring in year five. A similar retirement cycle has been in use for *BABAR*, and is reflected in the available capacity in line c.

Table 1: Projected Resources at the Western Tier 2

		2006	2007	2008	2009	2010
a. Dedicated U.S. ATLAS Capacity using Tier 2 Funds	CPU	111	427	872	1503	2505
	Disk	64	244	498	858	1430
b. Dedicated U.S. ATLAS Capacity using Local Funds	CPU	0	0	0	0	0
	Disk	0	0	0	0	0
c. Useful Non-dedicated Capacity	CPU	880	1456	2443	3877	6155
	Disk	0	0	0	0	0
d. Dedicated U.S. ATLAS FTEs using Tier 2 Funds	FTE	1	1	1	1	1
e. Dedicated U.S. ATLAS FTEs using Local Funds	FTE	0	0	0	0	0
f. Useful Non-dedicated FTEs	FTE	5	5	5	5	5
g. WAN Connectivity	Gb/sec	10+10+1	10+10+10	10+10+10	10+40+10	10+40+10

a. CPU is in units of kSpecInt2000, and disk is in units of usable terabytes after subtracting RAID and hot spare overheads. CPU and Disk acquisitions follow a constant-dollar investment each year, split so as to purchase 2*n* kSpecInt2000 for each *n* terabytes, resulting in an approximately twice as much being spent on storage as on CPU. Such a split is consistent with our understanding of Tier2 requirements and with *BABAR* experience.

b. Additional capacity will be available, but not dedicated. See c.

c. **CPU:** SLAC follows a model of shared computing capacity. When other groups are not using their full allocations, the capacity is available to all other groups. On the average, approximately 20% of the installed capacity has been available for opportunistic use. This figure shows the 15% of the expected installed capacity for SLAC science, exclusive of U.S. ATLAS acquisitions, that is likely to be available for opportunistic use by ATLAS. While continued investment for *BABAR* is not planned for FY2009 and 2010, it is expected that other SLAC science programs will drive a major expansion of capacity, and that at least the capacity shown will be opportunistically available to and suitable for ATLAS computing.

Disk: SLAC's disk allocation scheme does not provide for the same degree of capacity sharing, but use of the HPSS mass storage system and demand staging can provide a similar degree of expansion when required. Our experience has been that a mass storage system can give the appearance of up to twice the amount of actually installed disk, not counting use of the mass store for archival purposes.
Access to data: SLAC's computer center network allows unimpeded access from any batch machine to any storage server, thus opportunistically available CPU resources will have full access to ATLAS data at the Tier 2.

d. It is expected that the incremental cost of installing and maintaining the ATLAS dedicated capacity will require one FTE of ongoing effort, leveraging the high degree of automation already established.

e. Additional support will be available, but not dedicated. See f.

f. In as much as the expected scale of U.S. ATLAS Tier 2 computing will be 10–20% of the overall SLAC scientific computing load, this figure represents a similar fraction of the overall staff dedicated to scientific computing facilities. U.S. ATLAS computing, as an integral part of SLAC's scientific mission, will be treated as a priority activity by the staff. This model provides much greater continuity of effort and depth of expertise than dedicating individuals solely to the support of a single effort.

g. The bandwidth numbers show connectivity on the three principal network paths into SLAC: ESnet, ESnet Scientific Data Network, and via the Stanford campus into CENIC/CalREN2 and Internet2.

3 Governance

The Western Tier 2 must engage with U.S. ATLAS so that it plays a productive role in ATLAS simulation, detector studies and calibrations and physics analysis. The Tier 2 should also engage with its constituency of Western physicists to offer a flexibility to pursue their new ideas or to perform in-depth studies of the detector components for which they are responsible.

At a formal level, the SLAC Director will have responsibility to ensure that the resources provided by the funding agency are used correctly and effectively. The Director will ensure that the commitments to provide resources to International ATLAS and to U.S. ATLAS are met in full. It is anticipated that the SLAC Director and SLAC line management, with the agreement of U.S. ATLAS management, will seek the advice of a Western Tier 2 Advisory Board on which U.S. ATLAS management and the university-based and lab-based Western physics groups will be represented. The Advisory Board will participate setting technical directions for the Tier 2, and in deciding how best to use these resources to meet the challenges and opportunities of physics analysis and detector studies.

It is also anticipated that more informal interactions with the user community will occur in regular open meetings and phone calls.

3.1 Proposed user-community membership of the Advisory Board

University of Arizona	Peter Loch
University of California, Irvine	Andrew Lankford
University of California, Santa Cruz	William Lockman
University of Oregon	David Strom
University of Washington	Gordon Watts
University of Wisconsin, Madison	Sau Lan Wu
LBNL	Stewart Loken
SLAC	Charles Young

3.2 Tier -2 Personnel

Name	Appointment and Functions	Time on Tier 2 (2007%)	Time funded by U.S. ATLAS
Richard Mount	Principal Investigator Director, SLAC Scientific Computing and Computing Services	10	
Chuck Boeheim	Assistant Director, SLAC SCCS	15	
Randy Melen	High Performance Team Leader	20	
Wei Yang	ATLAS Tier-2 Contact; Systems Software Developer (Compute Element)	30	30
Booker Bense	Information Systems Specialist (OSG support)	20	20
Lance Nakata	Information Systems Specialist (Storage Element)	20	20
Systems Support Staff		20	20
Operations Support Staff		10	10

4 Proposal Summary

We propose a U.S. ATLAS Western Tier 2 center to be located within SLAC Scientific Computing and Computing Services at the Stanford Linear Accelerator Center. The infrastructure provides access to all needed facilities at minimal incremental cost to U.S. ATLAS. In addition, the technical and operational expertise of SCCS ensures the highest quality service for the U.S. ATLAS physics program. The Western Center is complemented by expertise at the Lawrence Berkeley Laboratory. An ESnet Metropolitan Area Network (MAN) operating at 10 Gbps ensures transparent and timely exchange of data.

This proposal is strongly supported by U.S. ATLAS groups, principally in the Western United States and will be closely coupled to the Western Regional Physics Analysis Center at LBNL. The governance of the Tier 2 Center includes university users to ensure that the center meets ATLAS needs. The operational experience within SLAC in working with grid-based collaborators in the U.S. and Europe and the history of bringing *BABAR* data to publication guarantee high value to the U.S. ATLAS program.

5 Budget Request

Table 2

		2006	2007	2008	2009	2010
Requested Funding $k	Equipment	200	400	400	400	400
	Operating		200	200	200	200