

**ADVANCED SCIENTIFIC COMPUTING ADVISORY COMMITTEE
COMMITTEE OF VISITORS REPORT**

Advanced Scientific Computing Research

Date: September 18, 2007

Date of COV: July 17-18, 2007
Program: Scientific Discovery through Advanced Computing – 2
(SciDAC-2)
Office: Advanced Scientific Computing Research (ASCR)
Agency: United States Department of Energy

Committee of Visitors Membership:

Robert Voigt, SAIC (Chair)
James Hack, National Center for Atmospheric Research
Merrell Patrick, Consultant
Ralph Roskies, Pittsburgh Supercomputing Center
Warner Miller, Florida Atlantic University

Executive Summary

The Committee of Visitors (COV) for the Department of Energy (DOE) Office of Science (SC) Office of Advanced Scientific Computing Research (ASCR) for the Scientific Discovery through Advanced Computing – 2 (SciDAC-2) program met at the DOE Germantown location on Tuesday, July 18 and Wednesday, July 19, 2007.

The COV is extremely grateful to the program officers and other ASCR staff who gave graciously of their time and knowledge to help the committee in its deliberations enabling the review process to proceed smoothly and effectively. We also want to extend our appreciation for the many program managers from the other offices in SC who are participants in SciDAC-2 and who provided extremely useful and candid briefings on their role in the process.

Findings:

The SciDAC-2 program is unique given the computational science goals of integrating science and simulation at the petascale level. But it is equally unique because of its broad intellectual scope and a broad administrative scope that cuts across multiple offices within the SC, and includes financial and intellectual participation by the National Nuclear Security Administration (NNSA) and the National Science Foundation (NSF). For these reasons the management of such a program is an extremely complex undertaking, particularly the competition phase, which had an extremely aggressive timeline.

Nevertheless, it is the opinion of the COV that the process was very successful despite time pressures and the ambitious and complex nature of the solicitation. One of the clear messages delivered to the COV was that there was a remarkable level of coordination amongst the various programs and offices within the Office of Science and partnering agencies. It was clear that the existence of regular communication across the Office of Science that was enabled by a Coordinating Committee, which met regularly throughout the process, was critical to the successful management of this very complex solicitation. As a result, SC has in place a strong scientific portfolio that is well positioned to address the goals of the SciDAC-2 program.

This report contains many recommendations which are offered in the spirit of improving the program and any future similar solicitations. Perhaps of most immediate relevance is the following recommendation:

ASCR should institute an annual review of SciDAC-2 awards.

It is the view of the COV that such a review would greatly strengthen the program and would facilitate defending it in the future. Given the size of the program, it also constitutes good stewardship of the taxpayers' money.

TABLE OF CONTENTS

1. Introduction	Page 4
2. Call for Proposals	Page 5
3. Review Process	Page 8
4. Selection Process	Page 12
5. Award and Declination Documentation	Page 14
6. Management of awards	Page 15
7. Recommendations	Page 17
Appendix 1: Charge	Page 19
Appendix 2: SciDAC Coordinating Committee	Page 20
Appendix 3: Ten Important Things to Remember	Page 21

1. Introduction

The Advanced Scientific Computing Advisory Committee (ASCAC) for the Office of Science (SC), United States Department of Energy (DOE), was charged by Dr. Raymond Orbach, Director of the Office of Science, with assembling a Committee of Visitors (COV) to review the SciDAC efforts within the Office of Advanced Scientific Computing Research (ASCR). The ASCR 2007 COV met at the DOE Germantown location on Tuesday, July 18 and Wednesday, July 19, 2007, where the focus of the activity was on the process that led to the recent SciDAC-2 awards.

The COV meeting opened on July 18th with a series of presentations by DOE program officers across the Office of Science. Walt Polansky reviewed the overall SciDAC program, starting with the SciDAC management principles and goals for the SciDAC-2 research program. As was noted in the solicitation, the program sought proposals that contributed to the creation of a comprehensive, scientific computing software infrastructure, integrating applied mathematics, computer science, and computational science in the physical, biological, and environmental sciences for scientific discovery at the petascale level, *and* a new generation of data management and knowledge discovery tools for the large data sets obtained from large experimental facilities and from high end simulations. A broad range of science application domains were called out in the SciDAC-2 solicitation, including Accelerator Science and Simulation, Astrophysics, Climate Modeling and Simulation, Computational Biology, Fusion Science, Groundwater Reactive Transport Modeling and Simulation, High Energy Physics, Nuclear Physics, Data-intensive High Energy and Nuclear Physics, Materials Science and Chemistry, QCD, Radiation Transport, and Turbulence.

Clearly, the SciDAC-2 program is unique given the computational science goals of integrating science and simulation at the petascale level. But the program is equally unique because of its scientific scope, engaging all programs within the Office of Science, and partnering with the NNSA and NSF. For these reasons the management of such a program is an extremely complex undertaking, particularly the competition phase which had an exceptionally aggressive timeline. The call for proposals was released at the end of December 2005. Letters of Intent (LOI) were due one month later with full proposals due in early March 2006. Panel reviews were conducted in April, and the development and justification for the SciDAC-2 portfolio was completed during the period between May and August. More than 350 Letters of Intent were received, and 268 of these were encouraged to submit a full proposal. Pre-proposals that were encouraged included 148 in the area of Scientific Applications (SAs), 44 in Scientific Application Partnerships (SAPs), 35 in Institutes, and 42 in Centers for Enabling Technologies (CETs). A total of 240 proposals were ultimately received with 109 of the submissions led by a Laboratory and 131 led by Universities. Of the 223 proposals received 114 were in the area of SAs, 30 were in the area of SAPs, 44 in the area of CETs and 33 were Institute proposals. After a very thorough review and vetting process 30 awards were announced in early September 2006, only eight months following the call for proposals. The awards included 17 Science Application and Scientific Application Partnerships, 9 Centers for Enabling Technologies, and 4 Institutes.

After reviewing the overall SciDAC-2 award process from the ASCR perspective, the COV was given a series of presentations from across SC (BER, HEP, NP, FES) on the way in which these programs participated in the SciDAC-2 competition. Following the formal presentations, the COV proceeded to review proposal materials, and internal materials related to the awards and declinations under the SciDAC-2 program. The remainder of the time was spent in discussions with staff about specific aspects of the award process and outlining and discussing the findings presented in this report.

One of the clear messages delivered to the COV was that there was a remarkable level of coordination amongst the various programs and offices within the Office of Science and partnering agencies, with the establishment of a Coordinating Committee (See **Appendix 2**) which met regularly throughout the award process. It was clear that the existence of regular communication across the Office of Science was critical to the successful management of this very complex solicitation. Indeed, it is the opinion of the COV, that the process was very successful despite the ambitious and complex nature of the solicitation. The process has led to a very strong scientific portfolio that is well positioned to address the goals of the SciDAC-2 program.

Overall, the COV's interactions with the Office of Science were excellent. The presentations were excellent and the interactions included direct and candid responses to any questions that arose. The committee was provided complete access to all materials related to the SciDAC-2 process, which was an overwhelming amount of information to digest. The committee received excellent support in response to questions, and received the full and enthusiastic support of the Office of Science personnel both prior to the committee's arrival and during the two day visit. The committee strongly commends the staff for their outstanding efforts in supporting the COV.

2. Call for Proposals

Because SciDAC is a project with both a broad intellectual scope and a broad administrative scope, cutting across multiple offices within the Office of Science, and including also funding from NNSA and NSF, crafting the RFP must have been a daunting task. Doing it within the imposed time pressures made it even more daunting. What emerged required compromise from all sides, resulting in a document clearly produced by committee, and therefore with some inconsistencies. The text for the RFP can be found at www.science.doe.gov/grants/FAPN06-04.html. Nevertheless the overall process has led to an excellent set of activities.

Reflecting ASCR priorities, the overall goals of the RFP were to attract applications that contribute to¹:

¹ Material in italics is taken from the RFP.

- *The creation of a comprehensive, scientific computing software infrastructure that fully integrates applied mathematics, computer science, and computational science in the physical, biological, and environmental sciences for scientific discovery at the petascale level, and*
- *A new generation of data management and knowledge discovery tools for the large data sets obtained from large experimental facilities and from high end simulations.*

But probably reflecting the priorities of other offices, the solicitation goes on to explain that it seeks applications that

- *Address obtaining significant insight into, or actually solve, a challenging problem of National scientific or engineering significance clearly related to DOE missions through computational science,*
- *Create scientific simulation codes that: achieve high single node performance; scale to thousands of nodes and tens-of-thousands of processors; and can be readily ported to other computer architectures,*
- *Develop applied mathematics and computer science methodology focused on computational science at the petascale and work with application teams to apply innovations,*
- *Integrate computational science with discipline-driven applications through teaming and partnerships with computer scientists and applied mathematicians,*
- *Engage experimental and observational data-intensive science, and/or*
- *Empower new scientific communities to achieve scientific discovery through computational science.*

Note that the emphasis on the petascale in applications is now downplayed or even lost. This is further reinforced where the general review criteria

- 1) *Scientific and/or Technical Merit of the Project,*
- 2) *Appropriateness of the Proposed Method or Approach,*
- 3) *Competency of Applicant's Personnel and Adequacy of Proposed Resources,*
- 4) *Reasonableness and Appropriateness of the Proposed Budget.*

are refined as follows (note the absence of emphasis on petascale, or even computational science):

In considering item 1 particular attention will be paid to:

- a) The potential of the proposed project to make a major scientific advance in a specific domain or to have a significant impact in the effectiveness of SciDAC applications researchers;*
- b) The demonstrated capabilities of the applicants to perform basic research and transform these research results into software that can be widely deployed;*
- c) Knowledge of and coupling to previous efforts in scientific simulation;*
- d) For enabling technology applications, the likelihood that the algorithms, methods, mathematical libraries, and software components that result from this effort will have*

impact on or is extensible to science disciplines outside of the SciDAC applications projects;

e) Identification and approach to software integration and long term support issues, including component technology, documentation, test cases, tutorials, end user training, and quality maintenance and evolution; and

f) Extent to which the application incorporates broad community (industry/academia/other federal programs) interaction;

In considering item 2, particular attention will be paid to

a) Quality of the plan for effective coupling to emerging advances in enabling technology or to applications researchers;

b) Quality and clarity of the proposed work schedule and deliverables;

c) Quality of the proposed approach to intellectual property management and open source licensing;

d) Quality of the plan for effective collaboration among participants; and

e) Quality of the plan for ensuring communication with other advanced computation and simulation efforts or enabling technology efforts.

Thus, a significant opportunity to prepare for and advance the use of petascale computing was lost. Nevertheless, very valuably what emerged was a set of applications that each integrated significant computer science and/or mathematics, and a set of enabling technologies that were tied to very specific applications.

The proposal called for 4 categories of proposals:

- Science Applications (SA),
- Centers for Enabling Technologies (CET),
- Institutes, and
- Science Application Partnerships (SAP).

The descriptions of the subject matter for science applications were clear and well done. However, it was not clear whether accompanying computer science and mathematics should be embedded in the science application, and/or identified as a separate partnership. For those which were separated, it was not clear to the reviewers what would happen if only the application but not the partnership were funded, or vice-versa. As it turned out, every funded science application had a corresponding funded partnership, and every partnership served at least one science application. Likewise, the distinction between Centers for Enabling Technology and Institutes was unclear, (although it was clear that Institutes have to be university led). At the COV meeting, it was stated that Institutes had to have a significant education and training aspect, but the solicitation merely suggests that they may have such a component. It appears that the program offices sometimes reclassified proposals into a different category to fit their crisper understanding of what was intended than was evident in the RFP.

In spite of the ambiguities of the RFP, the result is an impressive set of proposals that integrate science with computer science and mathematics, and the various offices involved are to be commended for their achievement.

With the above points in mind, the COV makes the following recommendations:

Any future SciDAC RFP should assure that:

- **Key goals such as petascale computing are included in the review criteria to help focus proposers and reviewers on the important issues;**
- **Partnerships are an integral part of applications so that it is clear how they are to be presented and judged;**
- **The distinctions between CETs and Institutes are clear;**
- **Training of graduate students is a criterion for evaluating Institutes**

In a future SciDAC RFP, consideration should be given to delaying the Centers for Emerging Technologies (CET) and Institute competitions until after the Science Applications (SA) have been selected.

The second recommendation is included for two reasons. First, delaying the competitions would make it possible to focus an RFP on those items in the SAs that require additional, focused research, which might also reduce the number of proposals requiring review. Second, such a delay would simplify both RFPs, and spread out the burden of review on the offices and the community.

3. Review Process

The office was faced with the task of creating a review process that could deal with both the complexity of the proposals and the fact that many different offices had a stake in the program, both financial and intellectual. This began with consideration of the Letters of Intent, followed by review of the proposals, both via submitted written reviews and panels, and culminating in a crosscut panel of experts with broad expertise in simulation based science.

The RFP required a two page Letter of Intent (LOI) including information on the PI and all senior personnel, projected funding information and a summary of the research. ASCR received 350 LOI. Of these, 268 were encouraged to submit full proposals while 82 were discouraged. In a few cases the LOI helped locate the activity in a more appropriate program. They also were used to identify potential reviewers. Requesting more technical information in the LOI might have made it possible to discourage more activities whose proposals would not fit the program, thus saving the offices and the community unnecessary work. With this in mind the COV makes the following recommendation fully understanding that a “discouraged” designation is not binding and the PI may in fact submit a full proposal:

Consider requesting a more detailed letter of intent that could be used to discourage proposals that do not address important aspects of the RFP and thus will not be competitive.

In the reduction from 350 LOI to 242 proposals, it is clear that discouragement had the desired effect on the community, and anything that could be done to sharpen the use of discouragement would help both the offices and the community.

ASCR received 242 full proposals in response to the RFP. Both the number and diversity of the proposals created a challenge for the review process. Simply finding sufficient numbers of reviewers who were not conflicted was a difficult task. The reviewers were selected by program directors from all programs involved. All reviewers provided written reviews via PeerNet; in addition, most were organized into 15 panels. These panels focused on a particular sub-area of the RFP such as Centers for Emerging Technologies or a specific application such as ground water. In the case of the application panels, there was usually at least one panel member with expertise in computational science. All panels were given essentially the same charge (see **Appendix 3**) so there was reasonable consistency of focus across the panels, although based on the reviews, the consideration of petascale computing and computational science, prominent factors in the RFP, seems to have received inconsistent treatment, and confusion over the categories, particularly SAPs, CETs and Institutes, was evident.

In order to provide focus on computational science and petascale computing, it may be necessary to convene additional panels that could focus on those topics. Given the future importance of petascale computing to the Office of Science, it is crucial that ASCR identify and pursue research activities that effectively exploit emerging computing capabilities. Thus the COV recommends:

Following review of the applications in similar technical areas, consider convening “computing” panel(s) to address the need and preparedness for high performance computing in the proposals considered appropriate for funding by the application panels.

The COV appreciates the fact that such panels have the disadvantage of introducing another level of review, but they have the advantage of providing uniform focus on topics of importance to ASCR. Furthermore, they could be used to eliminate from further consideration proposals that are deemed deficient in these areas, thus making the overall review process more efficient.

Given the complexity of the proposed activities, panels were an effective means for evaluation. Using panels allowed discussions across disciplines and led to many panelists’ changing their opinion on a proposal and in turn changing their review in PeerNet.

Program managers then built a list of approximately 130 fundable proposals using the following criteria²:

² Material in italics is quoted from the presentation by Dr. Walter Polansky to the COV on July 17,2007.

- *Clearly contributes to the SciDAC mission (e.g. advancing science through modeling and simulation at tera- to peta-scale; data management tools for extremely large data bases);*
- *Peer and/or panel review findings support proposal score;*
- *Panel score 7, or above;*
- *Other factors (affordability, program office mission, etc.)*

Ultimately the list included 133 proposals as a few meritorious proposals with scores slightly below 7 were added. This was further refined with a “yes”/”no” vote by program managers resulting in 71 of the 133 considered to be of more interest.

The review process now turned to a cross-cut panel made up of 17 individuals with broad backgrounds in and experience with simulation based science and engineering. The goals of the panel are summarized in the following charge³:

- *Identify and comment on proposals that could be candidates for funding, based on an analysis of reviewer comments, on full consideration of the review criteria, within the context of the SciDAC mission and other factors, including program manager ratings, when provided.*
- *Offer a perspective on the quality, the breadth and the balance represented by the top rated proposals within each major SciDAC category (Science Applications & Partnerships, Centers for Enabling Technologies, and SciDAC Institutes)*
- *Comment on the overall quality and synergism of a SciDAC portfolio comprised of top rated proposals.*

A cross-cut panel to look at the overall collection of proposals for potential awards is an excellent concept; however, the execution was flawed and thus the value of the panel, particularly with regard to building a final SciDAC-2 portfolio, is questionable.

The panel met for over three days to assess the 133 proposals. There was no opportunity to look at material beforehand and the deliberations were supposed to be based on the proposal abstracts and on the reviews, although full proposals were available. The 133 proposals were considered in a complex process involving subpanels that did a down select and swapped the stronger proposals with other subpanels for further consideration. The full panel then met for an extended discussion of the strongest proposals.

Based on the reports submitted by the panel members, there were several items that led to confusion and reduced the potential of the panel:

1. There were too many proposals to deal with;
2. There was no advanced material made available;
3. The review criteria were not clear and lacked specificity;
4. Proposals with a similar focus were not reviewed as a group which complicated comparisons;

³ Material in italics is quoted from the Charge to the Cross-Cut Panel that was given on April 25, 2007.

5. The “Yes/No” designation given by program managers prior to the panel meeting caused confusion;
6. There was a lack of understanding of the distinction between Institutes, CETs, and SAPs;
7. Since first round reviews were weak on computational needs/plans, without reading the proposal it was difficult to assess that aspect of a proposed effort.
8. Specific information on the performance of SciDAC-1 activities that were re-competing was not provided

In our discussions with program managers it became clear that they did not understand the purpose of the cross-cut panel and were concerned that it would negatively impact decisions that they had already made. In the end it appears that the panel’s deliberations had little impact on the makeup of the final SciDAC-2 portfolio, and program managers seemed universally pleased with the final outcome. In the future the purpose of a cross-cut panel should have clearly defined goals that are discussed with all program managers.

Conceptually a crosscut panel for a program of this size and complexity is an excellent idea; in this case the problems were in the implementation and would be relatively easy to correct in a future competition.

Items 1. and 7. could be addressed by introducing a “petascale” panel prior to the crosscut panel as recommended above. After the proposals are screened for application quality and relevance, a panel could examine them for their need and preparedness for high performance computing. This should be factored into a down select process that would result in a manageable number of proposals passed to the cross-cut panel.

Item 2. is easy to correct by making full information available to the panel prior to the meeting. This does require sufficient time for execution which, as stated elsewhere, was an issue with this review.

The difficulty raised in item 3. is traceable back to the RFP, which lacked specificity in the review criteria and was not clear on the importance of some aspects of the program such as petascale computing.

For item 4. many reviewers felt that it would have been easier and fairer to be able to discuss proposals with a similar thrust in a given sub-panel. This might be viewed as contradictory to the goals of a cross-cut panel that is charged with looking at a balanced portfolio, and the issue should be given some serious thought for any future review.

Perhaps item 5. caused the most confusion at the crosscut panel as it was not clear to panel members whether the “yes/ no” votes were in effect binding. These were provided by program managers prior to the panel review, but were apparently not used to down select proposals. It was never clear what useful purpose the vote served. In the end, some proposals with a “no” vote were funded. It should either be eliminated or used to reduce the number of proposals under consideration.

Item 6. again reflects an issue with the RFP. Proposers and reviewers alike did not understand the distinctions among CETs, SAPs and Institutes. SAPs were a particular problem as it was unclear whether they should be reviewed alone or in concert with the application to which they were connected. This was further complicated by the fact that some were submitted as stand-alone proposals and some were part of an application proposal. It was never clear what would happen to an application proposal that was dependent on a SAP that was not funded. These ambiguities should be corrected in any future RFP that offers funding for multiple, integrated activities.

Item 8. reflects a serious lack of important information. SciDAC-1 awardees had a clear advantage in this competition simply because of their experience, yet there was no information on to what degree they had fulfilled the promises in their SciDAC-1 proposals. This should be corrected in any future competition as discussed in **6. Management of Awards.**

With these points in mind, the COV makes the following recommendation:

Any future SciDAC effort should include a crosscut panel to assess the overall breadth and effectiveness of the portfolio, but it must be organized so as to address Items 1. through 8. above.

4. Selection Process

The breadth and complexity of SciDAC-2 and the requirement that projects receiving awards had to integrate activities across science applications, computer science and mathematics complicated the process of getting consensus on awards. This integration could be accomplished by SA projects including computer science and/or applied mathematics efforts with their proposed activities; SA projects partnering with SAPs or vice versa; and CETs targeting SA or SAPs projects for specific interactions. The selection was further complicated by the fact that a significant number of fundable proposals could not be selected because of lack of funds within both ASCR and some science programs. Finally, a compressed timeline further stressed the selection process. (See section **7. Recommendations**, for a recommendation on this issue.)

The Review Process is described in the section above, and activities conducted during this process were an integral part of the selection process. Some of these are obvious, for example, panel and cross-cut reviews and the choice of an average panel rating of 7 as a breakpoint led to a down selection of proposals for final consideration. Less obvious are steps taken before the review by panels began, for example, panel selection and assignment of proposals to panels. After the review process some proposals were moved to different categories as the selection portfolio was being finalized. A consistent theme during presentations to the COV was that the compressed timeline impacted panel selections in terms of the number of reviewers for individual proposals. Furthermore some programs had less money than others to fund the proposals deemed fundable in

their area so assignment of proposals to categories and panelists was important. The COV has no way of assessing the fairness of these actions, but points them out only to emphasize the importance of an adequate timeline and the care that needs be taken in such actions in order to build an excellent program of the size, breadth, and complexity of SciDAC-2.

As to the final phase of the selection process, as noted earlier, a SciDAC Coordinating Committee with members representing each of the programs was formed to help execute the SciDAC-2 program. The willingness of the members of this committee to work together and the leadership they provided within their constituents programs was key to relaxing the tensions created by the above mentioned stresses. Impressively, many of the tensions were overcome by extended 1-1 meetings between ASCR and SC program directors and in only one case did final resolution require higher management involvement. The COV was very impressed with this, the extent of negotiations and the give and take that took place in arriving at the selection portfolio. It is important to note that these intense negotiations allowed ASCR to contribute funds to every award in the selected portfolio except for one.

It is noted that only ASCR members serving on the Coordinating Committee selected CET and Institute awards since they were funded only by ASCR funds. Thus, the selection process in this case was less difficult. Since a major goal of CETs and Institutes is to produce technologies and education and training activities that meet the needs of SAs and SAPs, involving program managers representing those areas in the selection process could have led to stronger ties across the offices and might have resulted in stronger integration of mathematics and computer science. Delaying the CET and Institute competition as recommended earlier, could make it easier to involve other offices in their selection.

The COV had concerns about the final selection of Institutes as it was one of the areas in which proposals were moved after the review process, that is, some were reviewed in one category and moved to another for awards. The COV was specifically concerned about the lack of emphasis within the proposals on the training of graduate students. During the discussion of the selection process, the COV learned that the decision to fund one of the Institutes was made primarily by a single program director. This is well within the prerogative of program directors, but when such decisions are made, extra care must be taken to document the decision. This is discussed more thoroughly in the next section.

In summary, the selected portfolio consisted of 30 awards, 17 in SAs and SAPs, 9 in CETs and 4 in Institutes. Eighteen (18) were new awards meaning that they were not funded by SciDAC-1. Fourteen awards (14) were lead by PIs from Universities, 14 by PIs from the DOE laboratories, 1 by NASA Ames and 1 from industry. The DOE Laboratory breakdown of awards was LBNL (3), ORNL (3), LLNL (1), LANL (1), PNNL (1), and NREL (1). The COV commends ASCR and the other offices for achieving this excellent balance. Two application areas did not receive awards. Given the goals of the program to encourage the use of advances in computing technologies across all application areas, SciDAC-2 and application area program managers decided to release a call for proposals

in 2007 that will address climate change and accelerator physics. The COV commends the program managers on these extra efforts in achieving an even more balanced portfolio.

Regarding funding levels, SAs and SAPs received essentially the same level as CETs with Institutes receiving roughly one-third of the level of the other two. The level of funding for Institutes more than met pre-competition goals. Achieving this balance in funding across the three areas is commendable. It is noted that the level of funding across application areas within SAs was dependent on the funds available within individual applications programs so funding level balance there is not considered.

The COV notes that many fundable proposals in all categories could not be supported due to insufficient funds. This is true across many government programs, but we want to emphasize the need for increased level of funding for excellent programs like SciDAC that are attempting to take advantage of emerging technologies that address applications of national and international importance.

A few application proposals were not funded because the enabling technologies proposed as apart of the proposal were not considered to be fundable. Some of these were funded from offices' base programs, which is an indication of the quality of proposals received. Given the focus of SciDAC-2, such action was the appropriate way to support this work.

Finally the COV was very impressed by the role ASCR played in the selection and funding process. Their determination and willingness to negotiate one-on-one and to help fund all awards except one was a major factor in the ultimate breadth and balance of the portfolio.

5. Award and Declination Documentation

The COV conducted an essentially random survey of the jackets for both awards and declinations covering all components of the program; this resulted in the review of approximately 25 jackets. Throughout this process the staff of ASCR and the other Offices were extremely helpful. In some cases jackets had to be retrieved from program managers in other offices and in some cases program managers met with a COV member to explain further the decision process. The COV is very appreciative of this support and the openness of these discussions.

The information in the jackets was not consistent across offices and even within ASCR. For awards, discussion of the decision process was spotty. There was no analysis of awards in which the PI was from a DOE laboratory, but even for university awards the information ranged from essentially a restatement of the reviews to a careful analysis of the strengths and weaknesses. For declinations, there was no analysis.

For a complex, multidisciplinary, integrated program like SciDAC-2 that cuts across a number of offices, there should be consistency in the presentation of information in the

jackets, both awards and declinations. In many cases the funding decision involved input from several organizations including the Advanced Simulation and Computing (ASC) campaign of NNSA. This process must be documented and included in the jacket. It is not sufficient to track down a program manager and have him/her explain the process using a collection of saved e-mails. What happens to that information trail when the individual leaves the office? With this in mind we make the following recommendation:

Every jacket, both awards and declinations, with both lab and non-lab PIs, must have an analysis of the reviews that justifies the decision. This is particularly true for an award chosen from equally fundable proposals by a single program director.

There does not appear to be consistent treatment of communication with the PIs of declined proposals. Some offices return the reviews, some do not. In all cases the PI does not see an analysis of the reviews since none is done. Consequently, we recommend:

Reviews should be sent to all PIs whether the proposal is funded or not.

Since SciDAC-2 is an ASCR program, it is surprising that ASCR does not maintain a copy of every award jacket. Ideally, the program should use electronic jackets so that there would be a consistent copy of the jacket for all of the offices involved to access. This would assure that all offices are looking at an up-to-date copy of the jacket. Consequently, we recommend:

ASCR should maintain a copy of the jacket for every award regardless of what office has the lead role. Preferably this would be an electronic jacket available to all participating offices.

6. Management of Awards

As part of the COV's charter to evaluate the quality and relevance of the SciDAC-2 portfolio we also explored the proposed management approach to maximize the productivity from each of the thirty awards. This portfolio integrates and involves almost every facet within the Office Science. We can very well believe that the Program Managers will require some novel management strategies. We see a need for a tighter management of this complex and diverse portfolio. In particular, annual reviews would be useful with special attention given to each of the four Institutes.

The SciDAC-2 program represents an extraordinary opportunity for the Office of Science, and hand-in-hand with this prospect come unique challenges for effective review and management of the program. We provide some broad-brushed suggestions for consideration on managing the portfolio, taking into account the severe staffing and marginal travel budget the office faces.

We support the decision of the Office of Science to make Cooperative Agreement Awards to the vast majority of the PI's as this allows one to steer the effort, through

periodic changes in the statement of work, to better integrate the application areas to the HPC needs of the ASCR Program Office. We also understand that the standard procedure for many of the program managers is to perform a review midway through the award.

We see a distinct need for tighter office management of the awards, and stress the need for annual review of the SAs; these reviews should involve relevant SAP, Institute and CET personnel and include Program Manager and Laboratory representation. In addition, the longer-term education-based Institutes should be closely monitored so that their workshops and foci are tightly correlated with the SA areas. The COV has a deep appreciation of the severe staffing shortages as well as the tight restrictions on the travel budget; however, such reviews are, in our opinion, mission critical. One suggestion we offer in order to alleviate some of the staffing restrictions is to have an existing Office of Science contractor (e.g. Krell Institute) charged with implementing the panel reviews. We also could imagine great utility in the hire of an IPA versed in HPC from one of the DOE labs who would be dedicated to the close tracking and management of the Cooperative Agreements and who would coordinate with the contractor charged with implementing the reviews.

Past experience with a similar program (the NNSA ASC Alliance Program) has demonstrated the importance and utility of annual evaluations and involvement of both lab personnel and non-lab peers. The review panel consisting of Program Managers, Laboratory personnel, and peers identified early on a problem with the cross disciplinary/cross departmental integration within each of the five universities. In a few cases incentives were offered by the ASC Program Director to facilitate changes. This decision is widely viewed as a critical component of the success of this program. In addition, annual reviews help the PI's better manage the project. The benefits of annual reviews should be especially important with the loosely coupled and independently funded SAPs, CETs and Institutes.

The most challenging problem the COV identified in this area is the effective utilization of the four Institutes. Each of these Institutes has a more broadly-defined and longer-term mission, and has an important education requirement as well. Special attention, and annual reviews would be especially useful to help their PI's better focus on the longer-term needs of the seventeen SAs, and to support the ASCR Mission in petascale HPC.

With these points in mind, we offer the following recommendations:

ASCR should institute an annual review of the SAs, CETs and Institutes. Reviews of the SAs should include relevant SAPs; reviews of the CETs and Institutes should include relevant SAs. Given the severe staffing issues in ASCR, this could be organized and conducted by an independent contractor.

Reviewers for future SciDAC competitions should be given access to reviews of existing efforts that are participating in the new competition.

The second recommendation addresses a major concern raised by many reviewers, both in focused panels and the crosscut panel, and discussed in section **3. Review Process**.

7. Recommendations

In this section we include first a general recommendation and then for convenience collect all of the recommendations made in the preceding text.

For future programs of the complexity and breadth of SciDAC-2, careful attention must be given to develop a timeline that addresses the following points while providing for awards in a timely fashion:

- **The preparation of an RFP that is a clear, concise statement of the goals and objectives, the review criteria, the selection process, and the competition requirements;**
- **The selection of highly qualified panels of sufficient number that cover all aspects of the program;**
- **Changes that may need to be made after proposals are received; for example, the selection of additional reviewers to cover a topic more effectively;**
- **A selection period that allows for the negotiations among program offices that must take place in order to select the most appropriate portfolio;**
- **Preparation of consistent, high quality selection and declination documentation;**
- **Negotiation of awards.**

Any future SciDAC RFP should assure that:

- **Key goals such as petascale computing are included in the review criteria to help focus proposers and reviewers on the important issues;**
- **Partnerships are an integral part of applications so that it is clear how they are to be presented and judged;**
- **The distinctions between CETs and Institutes are clear;**
- **Training of graduate students should be a criterion for evaluating Institutes**

In a future SciDAC RFP, consideration should be given to delaying the Centers for Emerging Technologies (CET) and Institute competitions until after the Science Applications (SA) have been selected.

Consider requesting a more detailed letter of intent that could be used to discourage proposals that do not address important aspects of the RFP and thus will not be competitive.

Following review of the applications in similar technical areas, consider convening “computing” panel(s) to address the need and preparedness for high performance

computing in the proposals considered appropriate for funding by the application panels.

Any future SciDAC effort should include a crosscut panel to assess the overall breadth and effectiveness of the portfolio, but it must be organized so as to address the items outlined in section 3. Review Process.

Every jacket, both awards and declinations, with both lab and non-lab PIs, must have an analysis of the reviews that justifies the decision. This is particularly true for an award chosen from equally fundable proposals by a single program director.

Reviews should be sent to all PIs whether the proposal is funded or not.

ASCR should maintain a copy of the jacket for every award regardless of what office has the lead role. Preferably this would be an electronic jacket available to all participating offices.

ASCR should institute an annual review of the SAs, CETs and Institutes. Reviews of the SAs should include relevant SAPs; reviews of the CETs and Institutes should include relevant SAs. Given the severe staffing issues in ASCR, this could be organized and conducted by an independent contractor.

Reviewers for future SciDAC competitions should be given access to reviews of existing efforts that are participating in the new competition.

Appendix 1: Charge

The following charge was provided by Dr. Raymond Orbach to the ASCAC on February 22, 2007:

The August 15 2003 charge to the ASCAC instituted a Committee of Visitors (COV) to assess the program management of major elements of ASCR program every two or three years. The first two COV reviews - of the research program and the facilities efforts - resulted in a number of improvements of the processes. Following on these reviews I now ask ASCAC to conduct a COV review of the SciDAC efforts within ASCR. A report to ASCAC should be planned for the Fall 2007 ASCAC meeting.

Appendix 2: SciDAC Coordinating Committee

The following individuals represented their offices on the SciDAC-2 Coordinating Committee. The committee met regularly throughout the process from preparation of the RFP to decisions on awards.

Christine Chalk, ASCR
Anil Deane, ASCR
Fred Johnson, ASCR
Thomas Ndousse-Fetter, ASCR
May Anne Scott, ASCR
Mark Sears, ASCR
Yukiko Sekine, ASCR
Walt Polansky, ASCR
Craig Anderson, BER
Anjuli Bamzai, BER
John Houghton, BER
David Thomassen, BER
Richard Hilderbrandt, BES
Roger Klaffky, BES
Dale Koelling, BES
John Mandrekas, FES
Rostom Dagazian, FES
Steve Eckstrand, FES
Glen Crawford, HEP
Craig Tull, HEP
Sid Coon, NP
Njema Frazier, NNSA

Appendix 3: Ten Important Things to Remember

High Level Issues to Consider when Discussing Proposals

1. Do the proposals for research in the scientific application domains have a plan for partnerships that integrates advanced applied mathematics and computer science technologies with the proposed domain-specific efforts?
 - If not, is there a recommended approach to provide this linkage?
2. Do all proposals for Centers of Enabling Technologies include a plan for developing partnerships with science applications?
 - If not, is there a recommended approach to provide this linkage?
3. Do the proposals for research in the scientific application domains anticipate a path to the petascale regime, in either data or computational power?
4. Do all proposals, particularly for Centers of Enabling Technologies and SciDAC Institutes, address a plan for supporting their software over the long-term?

(The first two questions must be answered in the affirmative to be considered “SciDAC-like”)

Conflict of Interest Protocol

5. If you are conflicted on a particular proposal, please leave the room and take a well-deserved break while that discussion is in progress. The facilitating DOE representative will let you know when to return.

Guidance for Standards on Scoring (1-10) and Review Comments

6. During the day as you listen to and participate in discussions, you can alter or add to the scoring in PeerNet. You can also add comments for any proposal discussed in the panel, even though you may not have been assigned to review that proposal.
7. Please provide clear, but comprehensive, review comments aligned with scoring. The comments need to be concise with strengths and weaknesses identified and discussed in the context of the evaluation criteria.
8. Guidelines on what the scores mean are
 - Only proposals judged to be “SciDAC-like” will be considered for funding.
 - Score of 9 should be considered “cream of the cream.”
 - Score of 8 should be considered “top echelon.”
 - Score of 7 should be considered “fundable.”
 - Score of 6 should be considered “on the margin.”
 - Score of 4 or below should be considered “don’t consider.”

...and lastly

9. The SciDAC Director, the DOE Program Managers, and ORISE staff are here to help you in any way needed to complete this important task.
10. Thank you for your considered, valuable opinions.

