Advanced Scientific Computing Advisory Committee

Committee of Visitors

INCITE Report

COV date: April 23-24, 2008

COV membership:

Marsha Berger (chair), New York University
Roscoe Giles, Boston University
Cray Henry, HPCMP, Department of Defense
James Kinter, COLA, George Mason University
Dimitri Mavriplis, University of Wyoming
Gopal Shenoy, Argonne National Laboratory

Executive Summary

The COV commends the DOE Office of Science for creating the INCITE program, which has transformed the nation's ability to conduct massive computational science. The INCITE program puts unparalleled computing power into the hands of scientists working on the cutting edge. It has created some of the most advanced computing facilities in the world. It also makes available support structures and staff to help scientists make the most of these massive resources. INCITE has already transformed the computational science landscape.

INCITE is a relatively new program with a correspondingly new and evolving implementation. It is off to a good start but, as with all such path-finding efforts, adjustments can be made to make it more effective.

The COV's most important finding is a tension between supporting DOE mission science and the objective of supporting national leadership class computing. The INCITE mission statement and the HEC-RTF report envision a leadership class supercomputer dedicated to supporting a few high-risk and high-payoff projects. The authors of the HEC-RTF report expected mission agency capability class work to be accommodated on agency-specific capability class supercomputers. Instead, the DOE mission researchers also turn to INCITE for high-end capability class computing. The COV supports this and recommends that it be institutionalized by allocating a significant portion, but less than half, of INCITE resources to DOE mission needs, while reserving the majority for leadership class science. This approach is cost effective, is consistent with DOE's leadership role in driving high-end computing and will continue the beneficial interactions between the two groups (which are not always distinguishable).

The COV also recommends clarifying the criteria used in INCITE resource allocation. Today, because of the mixed objectives, the allocation process juggles conflicting requirements in a way that lacks transparency. Scientific merit and computational readiness are most important, but the allocation decisions also rely on unevenly collected project priority information, along with requests for specific supercomputer resources. Scientific reviews are accomplished differently for different science disciplines. This makes the comparison of scores from the different groups problematic, and creates opportunities to question the outcome. The COV found no evidence that the outcome was unfair, but because of the lack of transparency applicants question the decisions. Greater transparency and more feedback would improve user confidence in the process.

The readiness evaluation process also needs to be re-examined. A more flexible approach that provides more in-depth analysis than ready/not-ready would be beneficial. The COV also recommends instituting an appeals process, and staggering the timing of the call for proposals to better fit with other relevant calls. Project data should be collected so that in approximately five years a panel can be convened to assess the progress made on those projects to validate (or invalidate) the effectiveness of INCITE allocations on scientific progress.

Overall, the COV commends DOE Office of Science for creating the INCITE program and for its initial implementation, while recommending some changes to the process going forward.

1. Introduction

1.1 COV Charge and Process

The DOE Office of Science Advanced Scientific Computing Advisory Committee (ASCAC) was charged by Dr. Ray Orbach, Director of the Office of Science, in a letter dated Oct. 17, 2007 as follows:

"To help the research communities tap into the capabilities of current and future supercomputers, the Office of Science launched the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program in 2003. I am very pleased with the response from the community and the numbers and the quality of the proposals we have been receiving. Given the high visibility of the INCITE program and the importance of this program to national competitiveness, I now ask the ASCAC to form a COV to review the process used to manage the INCITE program. A report from ASCAC is expected by the Fall 2008 ASCAC meeting."

The letter is in Appendix 1. A Committee of Visitors (COV) was formed. It consisted of a diverse group representing a variety of fields and backgrounds: public and private universities, national laboratory, government agency, supercomputer users and directors. A list of the six COV members and their affiliations is in Appendix 2.

The group met on April 23-24 at the Germantown Headquarters of DOE. Before the meeting, the COV had on-line access to the 2008 INCITE dockets (and 2007 for the renewals), as well as summary spreadsheets prepared by INCITE program manager Barb Helland and staff. User feedback was solicited from a representative subset of users, including researchers from universities, laboratories, and industry. Some were new users and some were users whose projects had been renewed.

During the visit, the Committee heard presentations from the ASCR office, the Directors of the Oak Ridge, Argonne and NERSC (National Energy Research Scientific Computing Center) Computing Facilities, and had conversations with three DOE program officers with interest in INCITE. Each committee member read a selection of proposals in an area of his or her competence, including a few that received allocations and at least one that did not. The agenda is in Appendix 3.

1.2 INCITE Program History

The INCITE program was formally announced by the Director of the Office of Science in 2003. Its mission was to provide computing resources to a small number of computationally intensive large-scale research projects needing Leadership Class computing. The earlier HEC-RTF report had defined capability and capacity computing, and recommended the establishment of Leadership Class computing systems.

> *"The INCITE program was conceived specifically to seek out computationally intensive, large-scale research projects with the potential to significantly advance key areas in science and engineering."* (http://www.er.doe.gov/ascr/incite/index.html)

INCITE was originally based at NERSC, with 10% of the cycles given to three projects. In 2004, the DOE Office of Science received funding to create the Leadership Class Facilities (LCF) at Oak Ridge (OLCF) and Argonne (ALCF). In 2006, 10% of the Oak Ridge and Argonne facilities were added to the INCITE allocation pool. In response to OMB's request for a greater share of the new resources, 80% of the OLCF and ALCF cycles were allocated in 2007 and 2008, in addition to the 10% still at NERSC. (NERSC allocates 70% of its cycles to non-INCITE DOE production computing through a completely different process called ERCAP). Finally, 5% of the cycles at Pacific Northwest's Computing Facility were allocated by INCITE starting in 2006.

Table 1 contains a brief summary of the growth of the program. It is open to national and international researchers, including industry, with no requirement for DOE funding. Multi-year requests are now allowed. At each facility, 10% of the total cycles are reserved for Dr. Orbach's discretionary use and 10% are at the Center Director's discretion.

Year	Proposals Submitted	Proposals Approved	Processor-Hours Requested	Processor-Hours Allocated
2004	52	3	130M	5M
2005	23	3	28M	6M
2006	43	15	95M	18M
2007	107 (incl. 20 renewals)*	45	250M	95M
2008	112 (incl. 24 renewals)	55	600M	265M

* 9 were INCITE renewals, 11 were previously allocated DOE renewals.

The next sections discuss the different aspects of the INCITE process in detail, along with the committee's findings in each category. Since the INCITE program as well as the review process has changed substantially over the years, the COV focused on the most recent INCITE process and statistics.

2. The INCITE Program

2.1 Overall Picture

It is clear from the growth that INCITE has been a success; it provides the main domestic opportunity to apply a large fraction of a national class supercomputing resource to a specific scientific challenge. DOE is to be commended for taking the lead and accepting the risk inherent in pursuing big computational science projects across a variety of disciplines.

Finding 1: The INCITE program provides the largest national leadership class computational capability for the national science community, and is the primary high-end capability class support for the DOE community. The user perspective is largely positive. There are multiple examples of science successes made possible by the INCITE program. We commend the DOE Office of Science for fostering these successes.

The COV noted the energy and enthusiasm that the INCITE group demonstrates in the accomplishment of their work. The program managers have made a heroic effort in assembling diverse inputs to achieve a balanced program. INCITE is a new concept with a correspondingly new and evolving implementation. It is off to a good start but, as with all such path-finding efforts, adjustments can be made to make it more effective.

Finding 2: **INCITE has become the umbrella program that encompasses both national leadership class capability computing and high-end capability computing that directly supports DOE science. There is tension between meeting the DOE science mission objectives for high-end computing and the national objective of providing leadership class computing for the broader scientific community.**

The initial HEC-RTF report envisioned a leadership-class supercomputer dedicated to supporting a few high-risk and high-payoff projects. The authors of the initial recommendation expected mission agency capability-class work to be accommodated on agency-specific capability-class supercomputers. The leadership-class computing resources would be reserved for a few select projects judged to be of strategic national importance.

The COV found that although the program is still relatively young, the mission has evolved over time and is now considerably larger than originally envisioned; success naturally engenders some growth. However, this presents a tension with the program's stated goal of focusing on only a small number of leadership-class computing projects. In particular, current INCITE projects include both leadership class applications ready to run with large computational requirements as well as projects that reflect key capability computing needs of DOE science program areas.

Thus, the DOE mission applications and the wider OMB directive that INCITE serve as a national resource are now in conflict, with each group having widely different expectations. Historically the DOE has been in the forefront in supporting high-end computational science. It is inevitable that any merit and impact based allocation on the part of INCITE would include projects coming from within DOE. A natural expectation of DOE funded researchers is that they have access to the cycles they need on DOE machines. Ideally the scientific merit of these proposals would not need to be reviewed a second time, and only computational readiness would need to be assured. DOE program officers are in an uncomfortable position if they have funded projects that rely on LCF cycles, which are then not available to produce the work that has been funded. On the other hand, members of the national user community, should they get an INCITE award, expect to have exceptional access to generate a scientific breakthrough. The impact of this conflict is ameliorated in part by the increase in capacity of DOE computing. However the allocation processes, as well as the DOE user community, are showing signs of stress.

Both the DOE science mission objectives and the INCITE stated mission would be better served if this tension were reduced. One way might be to revise the distribution of cycles, with both DOE and INCITE users getting 30-50% of the cycles. While the same selection process might serve both groups, separate and unique selection processes tailored to each community should be considered. DOE might also consider other possibilities, such as connecting a DOE portion of the allocation to SciDAC and other agency programs, as well as the possibility that such allocations could be done at a different time of year. The COV recommends that this

restructuring be a high priority, but leaves it to the program offices to devise a suitable mechanism and implementation.

Finding 3. The INCITE portfolio has evolved so that there are a large number of projects taking up a smaller fraction of the facilities.

The INCITE program has grown over the years to the point where there are 55 projects being supported. While there is no obvious calculus to derive the optimal number of projects, the COV felt that 55 is too many. The highly-skilled support staff at the LCFs is struggling to cover all these projects. One of the strengths of INCITE is that each INCITE allocation includes the support of the application specialists assigned to work closely with the project (typically two projects per staff member). All three Centers indicated that they were overly stretched due to the growth in the number of awards over the last few years.

Through an examination of a subset of the proposals (both successful and declined), the COV felt that there were a number of submissions that either did not qualify or were too computationally immature to justify dedicated use of LCF. Several projects received relatively small allocations that put them outside the leadership class category. Data from NERSC indicates that their INCITE projects are not near the top in terms of size of jobs executing on the system, as they have been in previous years. Only 11 of 55 projects receiving 2008 allocations requested more than ten million processor hours.

This is not to suggest that an arbitrary minimum cut-off should be imposed. But projects that don't require a significant fraction of an LCF should be supported through Agency allocations and not INCITE allocations. The DOE may want to set community expectations on the amount of computing resources a typical INCITE project would require.

The COV also felt that the mix of proposals was slightly tilted towards DOE program priorities. This could be a result of the INCITE host being DOE, it could be due to inadequate awareness of the program outside DOE and DOE-sponsored activities, or perhaps that DOE scientists were in the best positions to apply.

As described in many contexts, a balanced portfolio of computing resources would include a few LCFs which operate in dedicated mode on a handful of large scientifically significant problems, complemented by a significant array of capacity facilities dedicated to support of the broader computing program. Since INCITE allocations are designed to speed up existing research projects by providing large amounts of computing resources over a short period of time we expect few projects to continue through multiple selections, thus over time INCITE will be able to promote more breakthroughs across more high impact activities. 'Continuations' of existing projects should be less common in the INCITE program than in other programs supporting DOE mission computing.

Finding 4: The INCITE program nicely balances the computing and support resources needed to effectively make use of complex computing environments inherent in leadership class supercomputers. The emphasis is appropriately split between operating the very

large supercomputers and in providing direct project level support to the individual science teams to improve their ability to make the best use of the supercomputer resources.

2.2 INCITE Process

2.2.1 Description of Review and Allocation Processes

INCITE proposals are either new or renewal proposals. Submissions received for new INCITE projects undergo a two-phase review, with the first phase focusing on scientific and technical merit, and the second phase focusing on computational readiness. Renewals (those projects in the second or third year of a multi-year proposal) are asked to submit a progress report to obtain the next year's allocation. They undergo only a readiness review.

The science review process for new proposals conforms to 10 CFR – Energy Part 605.10 (entitled *Application Evaluation and Selection*), and the procedure is coordinated through the Oak Ridge Institute for Science and Education (ORISE). The review process addresses four criteria for grading: *a*. Scientific and /or technical merit of the project, *b*. Appropriateness of the proposed method or approach, *c.* Competency of the personnel, and *d.* Reasonableness and appropriateness of the proposed request for computational resources. Science reviews are performed either by individual mail-in reviewers from the community at large (minimum three reviewers per proposal) selected according to their expertise in the specific science domain areas, or by convened review panels in specific disciplinary areas. In 2007, review panels were held in the fields of biology, climate, and combustion, while mail-in reviews were used for all other scientific disciplines. By all accounts, adequate precautions are taken to prevent any conflicts of interest. The science reviews provide overall ratings on a scale of 0 to 10 for four areas: scientific and technical merit, appropriateness of the approach, competency of the personnel, and justification of the request for computational resources. For proposals with overall ratings of 7 or higher, a justification of the importance of the research to the scientific field must also be supplied by the evaluator, and a second review for computational readiness is performed.

The computational readiness reviews for new as well as renewal projects are performed at the centers (ANL, ORNL, LBNL). All proposals are reviewed by one or more computational experts at each site, who provide input on scalability, reasonableness and appropriateness of the request for computational resources, processor hours required, appropriateness of computational approach, and overall technical readiness of project. A new process is in place for requesting additional information from PIs if needed for clarification of the review. The computational readiness review results in a rating of 1 (ready) or 0 (not ready), which must be justified by the reviewers. If a project is scientifically worthy but not computationally ready, it can be referred to a SciDAC Center, or to the Director's reserve, to help it get ready.

New project proposals are also made available to DOE science (SC) program offices to identify projects that are relevant to their program, and to rank these according to program priorities. SC program managers are given access to the INCITE proposals, but not the reviews.

An initial recommendation for project selection and computational resource allocation is then made by the INCITE program manager, based on the science and readiness reviews, the SC

program office prioritizations, and the availability of processor hours at the requested sites. Awards are made as much as possible on requested systems, but adjustments are made in consultation with the principal investigator to use alternative facilities if insufficient resources are available at the primary site. Renewal requests are awarded the full amount requested, presumably because their estimates of resource requirements are based on prior INCITE experience and hence more robust. These recommendations are submitted to the Director of the Office of Science for final decision on all allocations.

2.2.2 Process Issues

Several concerns with the review process were presented to the COV by DOE program personnel. The selection of qualified reviewers and the timeliness of their responses for mail-in reviews were described as problematic. The highly varied and multidisciplinary nature of the proposal submissions also makes it difficult to find knowledgeable program managers to help set priorities.

Members of the COV reviewed a subset of the 2008 proposals including both accepted and declined proposals for both new projects and renewals. The COV found some of the reviews to be of high quality and well balanced. However, many reviews were deemed short on scientific detail, and the readiness reviews were generally rather perfunctory.

The committee found that the review process was not entirely uniform or transparent, and that program guidelines had been stretched. As said above, the current INCITE portfolio, particularly at NERSC, contains projects that are smaller than leadership class. Some projects that did not appear to be computationally ready, principally from industrial applicants, were nonetheless given allocations. It was explained that the potential high impact of these projects was considered in making the allocation decisions. Additionally, different science disciplines were reviewed in different ways. Some used panel reviews, while others relied exclusively on mail-in reviews. This made it hard to compare scores from different groups. Similarly, projects of relevance to SC programs were prioritized by SC program offices, with little programmatic prioritization available for proposals in areas outside of DOE program interests. This was ascribed to the difficulty of obtaining outside input.

For renewal proposals, the review process almost always has resulted in continuing project allocations. On the one hand, this has been justified by arguing that these projects are more mature and their resource requests are more predictable and reliable, given their experience from past years. On the other hand, some projects were renewed in spite of minimal previous year accomplishments.

The computational readiness reviews provide critical information on the potential for projects to use the leadership class facilities effectively. However, they are not completely effective in identifying those computing projects ready to compute on day one. In addition, the binary *not-ready – ready* outcome of this review process was found to be overly restrictive. This is particularly the case when multiple reviews yielded different outcomes and the final recommendation necessarily considered an average of these recommendations. The COV cautions that simple speedup studies place certain applications at a disadvantage, and overall

time to completion among other factors should also be considered in the review. The readiness reviews also are non-optimal in that the individual LCFs conduct the review of their own Center's proposals, which creates the appearance of a conflict of interest, although the COV found no evidence that any such conflict had occurred. Additionally, dissimilar readiness review procedures are used at the various centers.

The actual allocation mechanism has many steps, some of which are less than transparent and many of which are not documented. The COV was impressed with the dedication, professionalism and energy level of the INCITE staff. In particular the INCITE program manager has the difficult task of organizing and collecting the numerous and diverse inputs from the review process, formulating the recommendations in a timely manner, and forwarding these initial recommendations to Dr. Orbach. However, the considerable latitude and level of discretion exercised by the program manager may not be a sustainable policy and could lead to difficulties in the future. In view of the unevenness of the review process across disciplines, and the diverse set of factors used in making project decisions, a clear articulation of the importance of these various factors and a more transparent and formal decision process is warranted.

The review does not currently have an appeal process. The COV felt that the institution of an appeals process would benefit both the user community and the INCITE program office. On the one hand, an appeals process would provide a mechanism for declined project PIs to obtain clarification of the review process, while on the other hand providing the INCITE program office with useful feedback on the quality of the review process.

The DOE runs several scientific research facilities other than INCITE that allocate access through competitive processes. The INCITE program staff, as they revise the INCITE allocation procedures, should investigate the procedures at these other facilities to investigate alternative models of operation.

Finding 5: Because of the mixed objectives, the INCITE selection process attempts to balance conflicting requirements and lacks a transparent set of selection criteria. While scientific merit and computational readiness are the primary criteria in the selection, the decision also considers unevenly collected project information and specific supercomputer system requests. The scientific reviews are accomplished differently for different science disciplines making the comparison of scores from the different groups problematic and creating the opportunity for proposers to question the outcome. We found no evidence that the outcome was unfair or biased but user comments did highlight a concern.

2.4 Proposal Timing

INCITE proposals are solicited annually from the science, engineering, and industrial community at large, both nationally, and internationally. The request for proposals (RFP) is posted on the DOE web site at http://hpc.science.doe.gov/. The RFP outlines the important characteristics of the INCITE program, such as the emphasis on Leadership Class computing, high impact science, and unique computational opportunities offered by the program. Additionally, the various hardware platforms are described in detail, with links to the appropriate lab web pages.

The frequency of the call for proposals was discussed during the COV interviews with DOE personnel. On the one hand, ORNL, ANL, and NERSC all felt that multiple (perhaps two) calls per year would be better at generating a more even flow of new projects. This concern was also echoed in some of the user feedback. One of the issues with a single yearly RFP is that declined INCITE proposals must wait a full year to reapply for computational resources, with potentially adverse impact on their scientific progress. On the other hand, the program office felt an increased frequency of calls would result in additional bureaucratic burden, complicate the review process, and dilute the impact of INCITE program announcements. Currently, the INCITE program is considering staggering the timing of the RFP in order to permit declined projects to re-apply for computational resources through other programs, such as the NERSC allocation process, *within the same year.*

The COV finds merit in the notion of more frequent allocations, despite the increased administrative load. As the allocation process is streamlined, documented and partially automated, this concern would diminish. Furthermore, it is conceivable that more frequent allocations would reduce the proposal pressure on the individual calls, thereby reducing the administrative load. The COV also strongly endorses the idea of staggering RFP schedules, and makes the general recommendation that INCITE RFPs should be coordinated as much as possible with other allocation RFPs such as ERCAP, as well as SciDAC.

2.5 Metrics of Success

As with many areas of the scientific enterprise, measuring the success of the INCITE program is difficult. There are several metrics that could help define success for INCITE, including for example the number, citation index, and impact factor of resulting publications; the degree to which the results could not be obtained with other than INCITE support, as well as a variety of performance statistics and computational methodologies developed on the LCFs. The individual LCFs collect anecdotal information about projects – the so-called "nuggets" – and these can be very valuable to help define the success of INCITE. However, these are insufficient to fully describe and monitor progress supported by INCITE. Reliance on journal publications is also problematic, since these often appear long after the work is complete, and the full impact of a given publication can sometimes become apparent only after years have passed. Most of the data that could quantify INCITE progress is collected by the individual Facilities, making it difficult to compare and compile these data across all of INCITE. A mechanism to obtain more continuous feedback from users is not yet in place. The ASCR office is putting in place a database to gather some of this information. This should be done on a regular basis, for example as part of the annual renewal process. INCITE is large enough and expensive enough that more thought should be given to its assessment.

3. Recommendations

Based on the discussion of the findings presented above, the COV's recommendations are as follows:

Recommendation 1: The selection processes for leadership class and DOE capability class computing should be separated. A significant portion, but less than half, of INCITE computational resources should be allocated to high-end DOE capability-class computing using a similar INCITE-type process.

Recommendation 2: INCITE awards should be fewer in number and larger in size with the expectation of demonstrated concurrency across a very large number of cores. To allow for projects deemed important but not ready, some resources could be reserved for development. Renewal should meet an achievement threshold below which projects are rejected, or referred to additional technical support on smaller platforms to make way for more promising new projects.

Recommendation 3: INCITE should continue to provide robust expert assistance to the science teams performing leadership class computing.

Recommendation 4: Review Process:

a. The selection process should be made as transparent and as uniform across disciplines as is practical. Selection criteria should be formulated and published. These should include scientific promise and importance, appropriateness of the computational technique, and potential impact on overall technological capability.
b. When this has been accomplished, consideration should be given to increasing the frequency of INCITE calls for proposals or at least staggering the annual call with other relevant calls such as ERCAP.
c. The computational readiness review process should adopt a more descriptive outcome, for example an overall grade (0-5) could be used for computational readiness. The readiness review would also benefit from a more systematic process such as a panel review performed by a group of computational experts from all the leadership class facilities, as well as from high-end computing experts outside DOE.
d. An appeals process for allocation decisions should be implemented.

Recommendation 5: The COV recommends that in approximately five years a formal review panel be convened to assess the impact of the INCITE program. This would be done through examination of project final reports, publication records, and the assessment of other types of impact, and by collecting feedback on what worked and what didn't from past users. Such information needs to be collected more systematically to be able to measure more precisely the scientific impact of INCITE.

Appendix 1: Charge to ASCAC from Ray Orbach

Under Secretary for Science
Washington, DC 20585
October 17, 2007

Dr. Jill P. Dahlburg, Chair
Naval Research Laboratory, Code 1001
4555 Overlook Avenue
Washington, DC 20375

Dear Dr. Dahlburg:

Thank you for the excellent Committee of Visitors review of the Scientific Discovery through Advanced Computing program. The Office of Advanced Scientific Computing Research (ASCR) has already undertaken changes to respond to the recommendations of the COV and improve the management of this important program. The full program response and action plan will soon be posted on the ASCR website.

To help the research communities tap into the capabilities of current and future supercomputers, the Office of Science launched the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program in 2003. I am very pleased with the response from the community and the numbers and the quality of the proposals we have been receiving. Given the high visibility of the INCITE program and the importance of this program to national competitiveness, I now ask the ASCAC to form a COV to review the process used to manage the INCITE program. A report from ASCAC is expected by the Fall 2008 ASCAC meeting.

I appreciate ASCAC's willingness to undertake this important activity.

Sincerely,

Raymond L. Orbach

Appendix 2: INCITE COV Members

Marsha Berger (chair)
Courant Institute, NYU
251 Mercer St.
New York, NY 10012
212-998-3305
berger@cims.nyu.edu

Roscoe Giles
Boston University
Dept. of Electrical and Computer Engg.
8 St. Mary's St.
Boston, MA 02215
617-353-9590
roscoe@bu.edu

Cray Henry
High Performance Computing Modernization Program
10501 Furnace Rd., Suite 101
Lorton, VA 22079
703-812-8205
cray@hpcmo.hpc.mil

James Kinter
Center for Ocean-Land-Atmosphere Studies
4041 Powder Mill Road, Suite 302
Calverton, MD 20705
301- 595-7000
kinter@cola.iges.org

Dimitri Mavriplis
University of Wyoming
Department of Mechanical Engg.
1000 E. University Ave.
Laramie, WY 82071
(307) 766-2868
mavripl@uwyo.edu

Gopal Shenoy
Argonne National Laboratory
Bldg. 401
9700 S. Cass Ave.
Argonne, IL 60439
630-252-5537
gks@aps.anl.gov

Appendix 3: INCITE COV Agenda

Meeting at DOE Headquarters
19901 Germantown Rd., North Entrance
2nd floor, E wing, Room E243

Wednesday April 23

8:45-9:00 Welcome to Committee, review charge — Marsha Berger

9:00-11:00 Description of INCITE program and process — Barb Helland, INCITE Prg. Dir.
Michael Strayer, ASCR Dir.

10:15-10:30 Break

11:00-12:30 INCITE Centers via access grid
- 11:00-11:30 Oak Ridge
- 11:30-12:00 Argonne
- 12:00-12:30 NERSC

12:30-1:30 (Working) Lunch

1:30-3 Committee reads proposals, internal discussions

3-4:30 Input from other DOE science offices
- 3:00-3:20 John Mandrekas (FES)
- 3:20-3:40 Susan Gregurick (BER-Bio)
- 3:40-4:10 Anjuli Bamzai (BER-Climate)
- 4:10-4:30 Open

4:30-6 Committee continues to read proposals and discuss. Make adjustments to Thursday schedule.

Dinner

Thursday April 24

8:30-2 Committee convenes, discusses, reads proposals, makes findings and recommendations.

10-11 Follow up questions to Barb and DOE office as needed, now or during the day.

Kathy Yelick may only be available Thursday if she wants to add anything via the grid.

12:30-1:30 Lunch

1:30-3:00 Initial Drafting of Report.

3:00-3:30 Closeout meeting with ASCR.