


# NERSC Programmatic Review

May 17-19, 2005



BERKELEY LAB



# National Energy Research Scientific Computing Center

## SUMMARY REPORT

### NERSC Programmatic Review

May 17-19, 2005

#### Introduction

On May 17-19, 2005, a program review of the National Energy Research Scientific Computing (NERSC) Center took place at the Lawrence Berkeley National Laboratory (LBNL).

The charge to the program review committee was the following:

1. Planning for the NERSC-5 upgrade and beyond, assuming a dollar flat, \$38 million per year budget from FY05 thru FY10.
2. The Request for Proposals (RFP), the proposal review, and source selection process prior to the NERSC-5 RFP.
3. The operating budget (which includes the upgrades), especially the staffing levels and the skills mix for NERSC-5 and beyond.
4. The allocation process; and
5. The Department of Energy (DOE) management.

In compliance with the Federal Advisory Committee Act, DOE required the peer reviewers to provide individual reports and DOE prepared this summary report.

The individual reports from the following peer reviewers are in the appendices:

1. Professor Frank Williams, program review chair, Director of the Department of Defense (DoD) Arctic Region Supercomputing Center at the University of Alaska Fairbanks.
2. Dr. Cray Henry, Director of the DoD High Performance Computing Modernization Program.
3. Dr. José L. Muñoz, Deputy Director of the Division of Shared CyberInfrastructure at the National Science Foundation (NSF)
4. Dr. Walter Brooks, Division Chief of the National Aeronautics and Space Administration (NASA) Advanced Supercomputing Division, NASA Ames
5. Dr. Lawrence Buja, NERSC user from the National Oceanographic and Atmospheric Administration (NOAA) National Center for Atmospheric Research (NCAR).
6. Mr. Robert Meisner, Deputy Director, National Nuclear Security Administration (NNSA) Advanced Simulation and Computing.
7. Dr. Tomasz Plewa, NERSC user and faculty member of the University of Chicago.

## **Executive Summary**

This summary report represents the Chairman's distillation of comments provided by a panel of seven external reviewers invited by Department of Energy Program Manager, David Goodwin, to consider five major aspects of the National Energy Research Scientific Computing Center (NERSC) during May 17 – 19, 2005, at the Lawrence Berkeley National Laboratory, University of California (LBNL, UC). The panel considered following areas: Planning; NERSC-5 Procurement and Process; Operations, Budget and Staffing; Allocations; and DOE Management. Panel members were Walt Brooks, NASA; Lawrence Buja, NCAR; Robert Meisner, NNSA; Jose Munoz, CISE NSF; Tomasz Plewa, University of Chicago; and Frank Williams, ARSC UAF.

Review panel members consistently expressed two common themes:

- 1) NERSC is a strong, productive, and responsive science-driven center that possesses the potential to significantly and positively impact scientific progress by providing users with access to high performance computing systems, services, and analytics beneficial to the support and advancement of their science.
- 2) The impact of DOE direction (for example, in system over-allocation and prescribing the number of jobs above an arbitrary large size), as well as duplication of compliance obligations and oversight by DOE and LBNL, UC, may decrease effectiveness of NERSC's scientific mission and impact NERSC's ability to provide the highest level of responsiveness and efficiency to users.

Observations and comments in the five major review areas included:

### **Planning**

- NERSC's strategic plan for 2006-2010 is impressive in its comprehensive approach to provide systems, services, and analytics capable of supporting research of interest to DOE's Office of Science.
- The center's facility-wide file system contributes to functionality of the high-performance storage system and to its excellent record of cyber security, capabilities that were highly praised by users.
- Investment of resources in analytics is prudent in anticipation of the needs of users and in the interests of DOE.

### **NERSC-5 procurement and process**

- NERSC-5 is a much-needed and justified acquisition that will support community-wide, high-end computing needs. It was apparent that NERSC performed a thorough appraisal and evaluation of center needs, that the procurement process is thorough and appropriate, and the procurement is endorsed and can be supported by LBNL staff.
- The panel urges acceleration of the NERSC-5 procurement process to make the new system available as soon as possible.

- The process may be improved through consideration of time value of cycles and dollars.
- If DOE intercession to constrain selection of HPC systems will occur, it should be made at the initiation of the procurement process.
- There may be value gained by competitive negotiations regardless of DOE constraints.

### **Operations, Budget, and Staffing**

- Flat funding will have negative impacts in face of increasing requirements for high-end computing.
- NERSC appeared understaffed compared to other sites, and will become more so particularly considering anticipated increased diversity of architectures.
- Demands on NERSC staff require flexibility to accommodate dynamic shifts in effort.

### **Allocations**

- Activities such as Innovative and Novel Computational Impact on Theory and Experiment (INCITE) are exceptionally valuable, and continuation of NERSC's participation is encouraged.
- Over-allocation is detrimental to users and the center's long-term reputation, effectiveness, and utility.
- Allocation recommendations from NERSC are based on scientific merit, and computational readiness and efficiency, however actual allocations appear disconnected in the allocation processes, with undesirable consequences.
- Allowing NERSC a computational allocation allowance to use at its discretion to meet contingencies, and to take advantage of fortuitous opportunities, is in the best interest of the center and Office of Science.

### **DOE Management**

- Excessive and duplicated (DOE and LBNL UC) oversight and guidance could contribute to a decreased efficiency in meeting NERSC's scientific mission.
- A holistic approach is needed across Office of Science computational centers to meet computational requirements of the DOE user community.
- DOE's architecturally-diverse set of computer resources should continue with both ORNL and NERSC supporting high-end computing capacity and capability.

Members of the review panel each report that NERSC is extremely well run with a lean and knowledgeable staff. The panel members saw evidence of strong and committed leadership, and staff who are capable and responsive to users' needs and requirements. Widespread, high regard for the center's performance, reflected in such metrics as the high number of publications supported by NERSC, and its potential to positively impact future advancement of computational science, warrants continued support. Clarification and streamlining of oversight and compliance (especially in areas of duplication) will increase the center's efficiency and utility, and further its long-term reputation.

## **Background**

A panel of seven external reviewers was asked to individually comment on five major aspects of NERSC. The panel was also given freedom to look into any aspect of NERSC the members wished to examine, and comment accordingly. The external reviewers were asked to independently form their opinions and report their comments to the NERSC Program Manager (PM), David Goodwin.

NERSC and David Goodwin provided advance support for the review. The information included annual reports; user surveys and responses; copies of materials to be presented during the site visit; links to the web page set up for the review to provide electronic access to information; notes from another, recent review-like activity; and especially:

NERSC's strategic plan for 2006-2010,

A draft RFP, Detailed Acquisition Plan, Benchmark Instructions and Acceptance Plan for the procurement of NERSC-5, and

An agenda for the two-and-a-half day review.

During the site visit the reviewers engaged NERSC management and staff in an ongoing dialog to gain clarity and express opinions on matters as they came up in presentations and in response to questions. Therefore, executive sessions were not taken as scheduled or for the purposes as expressed in the agenda. The panel privately discussed reporting methodology, their charge as expressed by the NERSC PM on behalf of DOE's Office of Science, their constraints and prohibition from purposefully coming to consensus among themselves or with NERSC staff, and coordinating further engagement with NERSC staff and LBNL leaders during any remainder of the site visit. Accordingly, each member of the review panel has submitted his own observations, comments and recommendations to the NERSC PM. Those reports are appended to this summary.

## **General Comments of Panel Members**

Throughout the on site review, the excellent quality of leadership and staff at NERSC was abundantly evident and further solidified the conclusion drawn from examination of NERSC's record of accomplishment. The mutual appreciation for each other's contributions at NERSC provides a solid fabric supporting the prodigious accomplishments that have brought widespread, high regard for the center. The self-initiated reviews (technology check) are purposeful and appropriate use of time spent to ensure quality. This feature is integral to successful advancement of opportunities throughout the computing science community.

Increased NERSC User Group (NUG) activities strengthen the viability of NERSC. This reflects NUG's commitment to serving users as they carry out their computational research. This group provides guidance to both NERSC and the DOE Offices of Science regarding the current services offered by NERSC and the direction of future

development. The emphasis on user surveys and appropriate response to them are exemplary. Frankly, the User Services aspects of NERSC are enviable.

NERSC should be commended for employing self-initiated reviews as checks on technology approaches being considered for deployment, and for periodic development of the “Green Book” to define and help quantify user requirements.

Location specific leveraging is apparent in that the intellectual community of LBNL and the academic community of the University of California Berkeley provide a world class, science driven environment for a science driven center. And, there is the highly qualified pool of talented people in the San Francisco bay area that can be drawn upon for staffing the center.

### **Planning**

NERSC’s strategic plan for 2006-2010 is impressive in its comprehensive approach to providing systems, services, and analytics for a science driven center supporting research of interest to DOE’s Office of Science, whose goal is to advance the fundamental science knowledge base and train future scientists. The message is compelling and the plan viable within the constraints specified by DOE. Every presentation and interaction on site furthered the panel’s confidence in the ability for NERSC to accomplish the plan. First, the plan calls for implementation of a facility wide file system – in clear response to NERSC’s evolutionary role as a computation and data center and promises to provide efficiencies to users and relieve data congestion across the center. Second, recognition of the importance of analytics to the understanding of data and computational results has led NERSC to plan for investment through redirection of staff work emphases and incremental funding to meet the challenge. The strategic plan is filled with evidence of detailed, deep thinking and careful choices.

The strategic plan was well organized and appeared accurate in its assessments of NERSC’s mission, projected technology trends that will impact their ability to succeed, and organizational structure needed to accomplish their mission in a resource-constrained environment.

NERSC’s move to analytics is a good one however, caution in limiting focus on visualization is recommended. Visualization gains much attention because it is easier to demonstrate and it has a certain inherent appeal. Rather than focus on feature extraction, data programming has the potential of benefiting a greater number of users and is complementary to any visualization activities (provide hints to users as to where one might look for interesting data). Staff should be acquired and trained to address this area.

NERSC’s recognition of the growing need for increasing staffing for software support is applauded, as this is an area that is typically under-staffed.

Creation and maintenance of a unified center-wide file system allows for increased usability and functionality of the high-performance storage system, and maintains the excellent record of cyber security at NERSC; remarkable achievements and elements

highly praised by users. Providing a sustained level of funding and development in these areas are essential for the success of NERSC operations.

Directing extensive testing at vendor sites or purchasing small test bed systems (if not available at other supercomputing centers) prior to acquiring future main computing systems may mitigate possible risks. Such a policy would help avoid problems, such as those currently experienced with the NCS machines (Jacquard).

LBL Director Steve Chu informally asked the reviewers about possible metrics for NERSC, those that may supplement traditional measures applied to supercomputing centers. Two possibilities could be: number of publications acknowledging NERSC resources (systems, services, analytics) and breakthroughs enabled by utilization of NERSC resources. A metric currently used by DOE (level of utilization and amount of time used by large jobs), does not necessarily reflect actual user/science needs. Smaller jobs and projects requiring specific resources (e.g. long runs) might be delegated to other systems leaving the Center's main computing system available to the most demanding applications. NERSC leadership and staff should give careful consideration of any proposed metric before implementation to avoid unintended consequences.

#### **NERSC 5 (year plan) Procurement and Process**

Based upon the draft, the continually updated Green Book, a user requirement document, which generally acknowledges understanding of the community wide requirement for more high-end computing resources, NERSC 5 is a much-needed procurement. The panel urges accelerating the NERSC 5 procurement process to make the system available much sooner. Delay of this process would be detrimental to NERSC. NERSC 5 instituted on May 29, 2003, is drawing near a close while the best value procurement process portrayed to the review committee seems to be exactly the right approach, acceleration would be advisable. It is doable in the LBNL context because of the valuable support and experience provided by the LBNL procurement officer.

The prolonged length of the process may be moot if NERSC/DOE cash flow cannot sustain the NERSC 5 procurement sooner than the projected acceptance, but ignoring that issue, the process seems very protracted. If there is need to delay the process in order to align the procurement in time with anticipated vendor technology releases, then the start of a faster process could be set accordingly. If there is concern over meeting milestones for performance of the acquisition, and if the process is driven by those considerations, then I think the process is being driven by the wrong considerations. Thus, the case for completely delegating the procurement authority to LBNL or UC would be made.

The process is extremely thorough, perhaps overly so, in that it may not be necessary to test many more technical performance aspects than are necessary to differentiate between technical solutions proposed. Having a comprehensive suite of tests for most desired performance features may be helpful in extremely close decisions, but what is realistically in the realm of possibilities for a very large, general purpose supercomputing system? The staff should be commended for putting together an understandable, rational, technical approach to the technical evaluation.

The time value of cycles provided to users might be looked at once again. Concerns raised after quick review by the team may obviate the need to include the costs of porting and optimizing codes, among other issues.

The endgame of negotiating is helped when there is clear competition. Without full appreciation for the constraints in the NERSC environment, the panel members suggest that NERSC think of ways to force competitive offers in a realistically limited field of potentially viable vendors. Certainly, NERSC's experience in the California and DOE context has evolved into a process that seems to bring good value in procurements.

Direction from DOE's Office of Science should be expressed early in the process allowing NERSC to include them in the best value evaluation and save unnecessary evaluation. For example, consider efficiency of center operation with only one main architecture to support at each facility, Office of Science-wide balance of architectures required, or national interest based limitation to domestic procurements.

The NERSC Team has created a benchmark suite that is representative of 85% workload at NERSC. The Team may want to use that collection of benchmarks to measure performance of several more systems currently installed or to be installed in the near future to gather performance and efficiency information and collect first-hand experience (portability).

NERSC may consider switching their hardware profile from a single large machine to two high-end systems installed at the same time. In such a scenario, hardware updates would take place more often offering users access to the most recent computing technology. The two systems could be consigned into a single (perhaps heterogeneous) installation through a dedicated internal fast network, given the common file system and archive storage already available at NERSC.

Aspects of the NERSC acquisition planning process were laudatory. NERSC pays close attention to systems balance, not only in terms of processors, memory and interconnect attributes, but also for systems storage and networks. Further, NERSC capitalizes on the benchmarking work of other centers from across the high-end computing community to increase the likelihood that they will receive best value in their acquisitions for the applications they support. The NERSC Procurement Team may want to seek additional information about vendors and their products from sister governmental agencies (DOE, DoD, NASA, NSF). Such information appears to be of great value in negotiations with vendors.

### **Operations, Budget and Staffing**

Evidence of potential negative impacts of flat funding in the face of increasing requirements for High End Computing (HEC) were clear. This will not allow the advancement necessary for success within this venue, even though they were not given emphasis in the documents provided or presentations on site.



Panel members repeatedly pointed to the need for more, not less, at NERSC:

NERSC makes a robust contribution to DOE's computationally based scientists and thereby the HEC community with as few staff as possible. The recent reductions in staffing seem to leave NERSC understaffed. Increasing the diversity of architectures may further stretch the staff. Given the constraints, one needs to appreciate the dynamic shifts in work and flexibility required of staff to meet the mission of NERSC.

NERSC management presented a thorough overview. The overview showed its current operation and its evolution over the next five years. They are running a lean, service-oriented, organization. In fact, when compared to other sites, current support staffing appears too lean. The current lean staffing will be stretched further due to the higher levels of support that will be required by the predicted increase in the architectural diversity of future NERSC computing platforms. This will stress the consulting group particularly hard right at the time when much greater proactive involvement is needed to ensure that the complex community simulation codes are able to run efficiently on the new systems.

Current NERSC services are highly rated by user surveys. NERSC has been proactively responding to the recent budget contractions by centralizing web services and automating its project administration to require less staff.

Reductions in NERSC's benchmarking and new technology evaluation staff are a regrettable loss and may have contributed to the rocky Jacquard acceptance. It was noted that cutting future architecture research/analysis investment is basically a "going out of business" posture. If NERSC does leave this area unstaffed, they should create strategic partnerships with other federal computing centers that are maintaining these types of groups to retain access to expertise in this area.

One minor caution was raised in reviewing the cost scenarios shown in the strategic plan. Escalating personnel costs due to even a low 3% inflation factor will cut into available services more significantly than shown in the table.

One-on-one discussions with NERSC employees showed a high level of staff moral, confidence and goodwill towards management.

### **Allocations**

Over-allocation is detrimental to the users and to the center's long-term reputation and usefulness. Cutting the allocations back is a hard decision that the Office of Science needs to make. The DOE Allocations Managers hold responsibility for managing the DOE Base Program and SciDAC allocations and should be providing input to the centers allocation management to ease the flow of process.

The allocation process itself appears disconnected from actual allocations given to users. The Office of Science needs to find ways to more obviously incorporate assessment of

computational readiness of applicants in the overall determination of allocations, or reduce the requirement for NERSC reviews of the proposals for HEC use. Scores given to the proposal by reviewers seems completely uncorrelated with amount of computing resources eventually granted. This trend seems persistent over at least the last few years, indicating not only that the review process does not work, but also the issue is not being addressed. One would expect more weight given to scientific merits of proposals.

The long-term health of the NERSC, and with appreciation for Office of Science nimbleness, could be dramatically assisted by giving the center a percentage of the computational allocation to use at its discretion to meet contingencies and fortuitous opportunities. Most center directors have this discretion, allowing them to take advantage of situations that most often bring high returns to the center. For example, trading runs of each other's benchmark codes would double the knowledge of system performance for participating centers and that would eventually benefit all users at the centers.

Setting metrics for fractions of jobs run above a specific job size can be efficacious if the fraction and size are carefully set, and consequences are predictable and/or acceptable. The recent, dramatic increases in time waiting for jobs to run, as reported by the NERSC user group, seems to call for re-evaluation of the apparently arbitrarily set requirement for the majority of jobs being larger than 512 processors. Having a goal of running applications that use greater than 512 processors more than 40% of the time is an appropriate goal for a center with a computational mission of enabling new science. But, in the long term this might not be sufficiently demanding. As computational models and methods progress, one would expect them to demand increasing numbers of processors and/or memory; thereby, putting increasing demand on capability resources. There will always be a tension between available resources (constrained funding) and application demands. The program needs to address this balance in a constrained resource environment, and possibly adjust the run time goal. This should be done as part of a more comprehensive approach to building a strategic plan for complex-wide use of systems for scientific discovery.

Activities such as Innovative and Novel Computational Impact on Theory and Experiment (INCITE) are valuable, exciting and holds promise for big returns in terms of advancement of the science it enables. Continuation of NERSC's participation is strongly encouraged because NERSC can support that kind of activity effectively, and it is an element of the portfolio of science support that sustains a center.

The high level of utilization (one of criteria used by the Office of Management and Budget in evaluating NERSC) seems to be a mere by-product of a long waiting queue of small jobs. Although such a long waiting queue allows to patch holes between large jobs and therefore achieve exceptional utilization level, waiting queue times reaching several (five) days do not really allow for advancing science at a reasonable pace. This is especially true in case of projects involving relatively small jobs where quick turnaround time is essential (e.g. steering required in parameter studies).

The unmet requirements for HEC availability could be improved by installing additional computing resources at NERSC in the near future.

### **DOE Management**

The burden of compliance, oversight, and guidance could be decreasing efficiency towards meeting NERSC's scientific mission. Though if well managed and thought out this could be an opportune time to finesse the procurement process within the DOE structure in its dealings with NERSC. This is a multi-faceted issue, perhaps doubled because of DOE and UC oversight, and, some may be driven as part of widespread increases in Federal and State oversight/compliance roll downs. However, cost effective responses could include taking care of reporting at DOE headquarters or trusting LBNL/UC oversight and eliminating DOE oversight/guidance.

A strategic approach, across the Office of Science community, to meet computational requirements appears to be called for to best meet the unmet demands of computational scientists. There seems to be a lack of computational resource provisioning for the middle and large classes of computing, those just shy of Leadership Class.

There appears to be considerable angst over the issue of who is the "capacity" and who is the "capability" resource provider at DOE. This is counter-productive. Clearly, both ORNL and NERSC are providing capability-computing cycles. Capability computing is multi-dimensional and DOE's currently architecturally diverse set of computing resources should continue.

Two concerns became apparent during the review that make it apparent that DOE has crossed the line from telling the Lab *what* it wants it to do to *how* it should do it's job:

Direction to stop work on emerging cluster technology represented by systems such as Alvarez resulted in a lack of experience with vendor products and the challenges of cluster systems that are a major contributing factor to the slow deployment of Jacquard system.

It appears that DOE strongly steered the NERSC team towards significant cuts in personnel and there appears to be a drive and suggestion that even deeper cuts are possible. As a result, during a period in which the complexity of architectures and systems is increasing, the panel was briefed on cuts to users services by a factor of 2X. Based on my knowledge of HEC center operations, NERSC appears to be extremely well run and has a lean/knowledgeable staff. Budget pressures may require a reduction in both hardware and personnel, but NERSC is NOT an inefficient group with a bloated staff. A cut from 66 to 60 FTE will seriously impact user services and further cuts will jeopardize the viability of NERSC as a full service simulation center.

DOE should work directly with NERSC to streamline review and reporting requirements to help mitigate the effects of the 10% staff reduction. There seems to be a number of reviews (DOE, LBNL, Univ. California) taking place at NERSC. These activities consume a substantial amount of time (25% according to rough estimates) and energy of

NERSC's management and staff. The burden associated with reviews can be lowered by either decreasing the frequency of reviews (one in two years, rather than once per year), limiting them and/or by scheduling joint reviews. The latter combination seems most attractive.

The requirements and effect of broad cyber security policies directed primarily at desktop computer systems and management and administrative personnel can have dire consequences if applied in a broad brush to research facilities and research teams. DOE top management may help by creating a board that balances at the CIO level the extreme requirements in networks, supercomputers and data storage that are pushing the state of the art.

Security was recognized as an increasingly important topic and planned increased staffing reflects this. Given the natural tension between productivity and security in a center supporting so many remote users, NERSC has done an exemplary job to date maintaining an open, yet secure site. NERSC's expertise in this area will be tested by the substantial security challenges that are still to come. On the horizon are far-reaching regulations on foreign national access to Federal computers that may have significant impact on the NERSC community. A potential risk noted by the panel is NERSC's exposure to becoming involved in security incidences from having users unwittingly violating ITAR statutes on open NERSC machines or from program PI's failing to sufficiently verify their researcher's citizenship or visa status. NERSC should elevate current user security awareness by implementing security education programs for users that are in place at other DOE centers such as LANL. At a higher level, to reduce the danger of overly restrictive security directives coming down from the DOE CIO with little or no input from NERSC, NERSC should take a leadership position in the development of DOE-wide computer strategies.

### **General Comments**

The Panel Members provided a wide range of general comments:

Quantify "savings" through code speed-ups – most centers have tried to do this, the NERSC staff may be up to the challenge, and it certainly would be a measure of value added by NERSC.

Consider sharing the security model and practice developed at NERSC with the HEC and IT community.

Seek computational system cost information from other Federal Agencies (cost benchmarks) to add confidence to your procurement process.

Ensure staff and NERSC participate in advancing HEC technology through ongoing activity in development, test and evaluation. This is a necessary condition for sustaining a center and its users.

Strategic collaborations and HEC community activities are valuable in providing enhanced systems and services to users, (i.e. SP-SSL, SC0X, among many others you engage in).

In this time when the centers are being asked to do more with less, efforts to increase the efficiency of resource utilization can have large, measurable benefits. NERSC should automatically profile running codes to identify inefficiencies and provide incentives, such as software engineering support, to help those codes make better use of the resource. Incorporate an automated code analysis mechanism, perhaps using IPM, which would look at the jobs executing and identify jobs that are in need of improvement with respect to efficiency.

Incorporate check point/ restart in order to provided greater flexibility in job preemption. This is apparently already being looked at and should continue.

Newer users should automatically be identified for assistance to get them started along the right direction. This should be started ruing the STARTUP allocation process where users are given 20K hours.

NERSC has a large base of experienced users conducting research in several diverse fields of scientific computing. It is important to define and use in the Center's evaluation metric that goes beyond already stated qualitative "excellent grade". A common practice adopted by several countries to evaluate performance of their scientific institutions is the number and impact of publications that acknowledge support (funding, computing time) obtained from such organizations. NERSC tracks and collects information of such publications, the number of which exceeded 2,000 in 2004. The sheer number of publications supported by NERSC is impressive. However, this metric can be improved and made quantitative by providing a cumulative number that weights each publication by its corresponding impact factor (as defined by the Science Citation Index). Other more involved, metrics can be constructed, but including the ISI impact factors would offer significant improvement over currently used total number of publications.

In some cases, however, no single metric can be applied. In particular, one would hope that among such a large volume of published work a few truly exceptional breakthrough discoveries are present. One such example is the Center's support for Tokamak research that has resulted in construction of high-quality numerical algorithms and application codes allowing for engineering-type design studies of these extremely complex and expensive devices. Identifying three new highly stable beam configurations resulted in funding and building three new experimental facilities that will study actual stability of these theoretically predicted configurations.

The INCITE program is one more example of an aggressive approach to computations. This high-risk high-return large-scale computational program allows a narrow group of carefully selected researchers to break barriers preventing them from exploring regimes in which existence of new phenomena is anticipated or significantly extending our current knowledge. One additional measure of success of the INCITE program might be

awards given to foreign (non-US based) institutions. It is conceivable that a number of such applications will be growing with time. Presence of such awards would indicate that the idea of INCITE has a truly global dimension.

The whole panel was most impressed with the esprit-de-corps observed throughout the visit. Without exception, everyone that I had the pleasure of talking to had positive attitudes and genuinely liked their work. They were proud of NERSC's contributions to enabling new science and excited about being part of the team. Maintaining that vigor in an era of constrained resources will be a demanding leadership challenge, but they have the talent to succeed.

Special thanks for excellent support in preparation for the review (i.e., the read aheads), logistics, and individual accommodations for panel members.

### **Follow-up Actions**

On May 25, 2005, the DOE SC staff members of the Office of Advanced Scientific Computing and Research (ASCR) were briefed on the program review. As stated in the attached thank-you letters to the peer reviewers: "Out intent is to use your findings and recommendations to improve the operations of NERSC".

On May 27, 2005, the Office of Science (SC) Supercomputer Allocations Committee (SAC) for NERSC, with members from all of the SC research Program Offices, were briefed on the program review. The SAC members were requested to review a draft of this report and the NERSC response to the NERSC Users Group (NUG) Greenbook (user's needs). The SAC members were advised that future computational reviews of the proposals to use NERSC would no longer include peer reviewers external to NERSC. Comments on this report were also requested from the NUG chair and the NUG Greenbook Editor.

For the June 6 weekly report to DOE Secretary Bodman, a summary of the program review was prepared.

On June 16, the ASCR management decided to prevent future over-allocations by not allocating a computer until it has passed acceptance testing.

For his June 24 visit to NERSC, a draft of this report was provided to the Director of the Office of Science (Dr. Ray Orbach).

To try to reduce the NERSC-5 procurement time, the DOE Office of Science (SC) has prepared project management guidance for Information Technology (IT) projects.

For the Fiscal Year 2006 allocations and beyond, SC decided to resolve the "capability computing versus capacity computing" issue by allocating most of the 40 teraflops at the Oak Ridge National Laboratory Leadership Class Facility to only about 13 Principal Investigators (PIs). The NERSC "capability computing" capacity will continue to be

allocated to over 240 PIs (NERSC has 10 teraflops as the time of the review, was upgraded to 13 teraflops on August 1, and is being upgraded to 20 teraflops in 2006).

It is anticipated that a briefing on the program review will be provided at the next (to be scheduled) meeting of the ASCR Federal Advisory Committee Act (FACA) Advanced Scientific Computing Advisory Committee (ASCAC).

Appendix 8 provides the NERSC response to the peer reviewer reports. This response was approved by the LBNL Director and will be used to implement most of the recommendations.

### **Appendices**

**Appendix 1:** Report from Professor Frank Williams, program review chair, Director of the Department of Defense (DoD) Arctic Region Supercomputing Center at the University of Alaska Fairbanks

**Appendix 2:** Report from Dr. Cray Henry, Director of the DoD High Performance Computing Modernization Program

**Appendix 3:** Report from Dr. José L. Muñoz, Deputy Director of the Division of Shared CyberInfrastructure at the National Science Foundation (NSF)

**Appendix 4:** Report from Dr. Walter Brooks, Division Chief of the National Aeronautics and Space Administration (NASA) Advanced Supercomputing Division, NASA Ames

**Appendix 5:** Report from Dr. Lawrence Buja, NERSC user from the National Oceanographic and Atmospheric Administration (NOAA) National Center for Atmospheric Research (NCAR)

**Appendix 6:** Report from Mr. Robert Meisner, Deputy Director, National Nuclear Security Administration (NNSA) Advanced Simulation and Computing

**Appendix 7:** Report from Dr. Tomasz Plewa, NERSC user and faculty member of the University of Chicago.

**Appendix 8:** NERSC response to the peer reviewer reports.

**Appendix 9:** June 16 thank-you letters to the peer reviewers.

# **Appendix 1**

## **National Energy Research Scientific Computing Center**

### **NERSC Programmatic Review**

**May 17 – 19, 2005**

#### **Report**

#### **By**

**Frank Williams, Director  
Arctic Region Supercomputing Center**

A panel of seven external reviewers was asked to individually comment on five major aspects of NERSC. The panel was also given freedom to look into any aspect of NERSC the members wished to examine, and comment accordingly. The external reviewers were asked to independently form their opinions and report their comments to the NERSC Program Manager, David Goodwin.

NERSC and David Goodwin provided advance support for the review. The information included annual reports; user surveys and responses; copies of materials to be presented during the site visit; links to the web page set up for the review to provide electronic access to information; notes from another, recent review-like activity; and especially:

NERSC's strategic plan for 2005-2010,

A draft RFP, Detailed Acquisition Plan, Benchmark Instructions and Acceptance Plan for the procurement of NERSC-5, and

An agenda for the two-and-a-half day review.

During the site visit, the reviewers engaged NERSC management and staff in an ongoing dialog to gain clarity and express opinions on matters as they came up in presentations and in response to questions. Therefore, executive sessions were not taken as scheduled or for the purposes as expressed in the agenda. We did privately discuss our reporting methodology, our charge as expressed by the NERSC PM on behalf of DOE's Office of Science, our constraints and prohibition from purposefully coming to consensus among ourselves or with NERSC staff, and coordinating our further engagement with NERSC staff and LBL leaders during any remainder of the site visit. Accordingly, each member of the review panel will submit his own observations, comments and recommendations to the NERSC PM.

Copies of all the documents including the agenda were given to the PM, and therefore they are assumed to be appropriately incorporated in the summary report. The review team members were Walt Brooks, Division Chief, NASA Advanced Supercomputing Division; Lawrence Buja, National Center for Atmospheric Research; Cray Henry,



Director, DoD High Performance Computing Modernization Program; Robert Meisner, Office of Integrated Computing Systems, National Nuclear Security Administration; Jose Munoz, Deputy Division Director, Shared Cyber Infrastructure, CISE, NSF; Tomasz Plewa, The ACS FLASH Center, U of Chicago; and Frank Williams, Director, Arctic Region Supercomputing Center.

### **General Comments**

Throughout the on site review, the excellent quality of the leadership and staff at NERSC was abundantly evident and further solidified the conclusion drawn from examination of NERSC's record of accomplishment. The mutual appreciation for each other's contributions at NERSC provides a solid fabric supporting the prodigious accomplishments that have brought widespread, high regard for the center.

User Group activities strengthen NERSC and reflect commitment to serving users as the users carry out their computational research. The emphasis on User Surveys and appropriately reacting to them are exemplary. Frankly, the User Services aspects of NERSC are enviable.

NERSC should be commended for employing self-initiated reviews as checks on technology approaches being considered for deployment, and for periodic development of what is known as a Green book to define and help quantify user requirements.

Location specific leveraging is apparent in that the intellectual community of LBL and the academic community of the University of California Berkeley provide a world class, science driven environment for a science driven center. And, there is the wonderful pool of talented people in the San Francisco bay area that can be drawn upon for staffing the center.

### **Planning**

NERSC's strategic plan for 2006 -2010 is impressive in its comprehensive approach to providing Systems, Services, and Analytics for a Science Driven center supporting research of interest to DOE's Office of Science. I found the message compelling and the plan viable within the constraints given by DOE. Every presentation and interaction on site furthered my confidence in the ability for NERSC to accomplish the plan. Rather than repeat comments that others are sure to make, or that may fit better elsewhere in this report, I'd point to two pieces of the plan that support my position. First, the plan calls for implementation of a facility wide file system – in clear response to NERSC's evolutionary role as a computation and data center and promising to provide efficiencies to users and data congestion across the center. Second, recognition of the importance of analytics to the understanding of data and computational results has led NERSC to plan for investment through redirection of staff work emphasis and incremental funding to meet the challenge. The strategic plan is filled with evidence of detailed, deep thinking and careful choices.

LBL Director Steve Chu informally asked the reviewers about possible metrics for NERSC, those that may supplement traditional measures applied to supercomputing

centers. Two possibilities could be: number of publications acknowledging NERSC resources (systems, services, analytics) and breakthroughs enabled by utilization of NERSC resources. NERSC leadership and staff should give careful consideration of any proposed metric before implementation to avoid unintended consequences.

### **NERSC 5 Procurement and Process**

Based upon the draft, updated green book and generally acknowledged understanding of the community wide requirement for more high-end computing resources, NERSC 5 is a much-needed procurement. My comments should be taken as supportive of the procurement and the process proposed in general, but colored by my anxiety for meeting the demand as soon as possible.

The process is extremely thorough, perhaps to the point of being too thorough in that it may not be necessary to test many more technical performance aspects than are necessary to differentiate between technical solutions proposed. Having a comprehensive suite of tests for most desired performance features may be helpful in extremely close decisions, but what's realistically in the realm of possibilities for a very large, general purpose supercomputing system? The staff should be commended for putting together an understandable, rational, technical approach to the technical evaluation.

The best value procurement process portrayed to the review committee seems to be exactly the right approach. It is doable in the LBL context because of the valuable support and experience provided by the LBL procurement officer.

Concerns for parts of the process did emerge. A) The length of the process may be moot if NERSC/DOE cash flow cannot sustain the NERSC 5 procurement sooner than the projected acceptance, but ignoring that issue, the process seems very protracted. If there is need to delay the process in order to align the procurement in time with anticipated vendor technology releases, then the start of a faster process could be set accordingly. If there is concern over meeting milestones for performance of the acquisition, and the process is driven by those considerations, then I'd think the process is being driven by the wrong considerations. Thus the case for completely delegating the procurement authority to LBL or UC would be made. B) The time value of cycles provided to users might be looked at once again. Concerns raised after quick review by the team may be obviated with more reflection including the costs of porting and optimizing codes among other things. C) In my experience, the end game of negotiating is helped when there is clear competition. NERSC experience in the California and DOE context has evolved into a process that seems to bring good value in procurements. Without full appreciation for the constraints in the NERSC environment, I'd only suggest that NERSC think of ways to force competitive offers in a realistically limited field of potentially viable offerors. D) If there is direction from DOE Office of Science based upon considerations of efficiency of center operation with only one main architecture to support at each facility, or Office wide balance of architectures required, or national interest based limitation to domestic procurements, then expressing those interests now would allow NERSC to include them in the best value evaluation and save unnecessary evaluation.

### **Operations/Budget/Staffing**

Potential negative impacts of flat funding in the face of increasing requirements for HEC were clear even though they weren't given emphasis in the documents provided or presentations on site.

I think NERSC makes a robust contribution to DOE computationally based scientists and thereby to the HEC community with as few as staff as possible. The recent reductions in staffing seem to leave NERSC understaffed. Increasing the diversity of architectures may further stretch the staff. Given the constraints, one needs to appreciate the dynamic shifts in work and flexibility of staff to meet the mission of NERSC.

### **Allocations**

Over-allocation is detrimental to the users and to the center's long-term reputation and usefulness. Cutting the allocations back is a hard decision that the Office of Science needs to make.

The allocation process itself appears disconnected from actual allocations given to users. I'd suggest the Office of Science find ways to more obviously incorporate assessment of computational readiness of applicants in the overall determination of allocations, or reduce the requirement for NERSC reviews of the proposals.

Long term health of the NERSC, and appreciation for Office of Science nimbleness, could be dramatically assisted by giving the center a fraction of the computational allocation to use at its discretion to meet contingencies and fortuitous opportunities. Most center directors have this discretion, allowing them to take advantage of situations that most often bring high returns to the center. For example, trading runs of each other's benchmark codes would double the knowledge of system performance for participating centers and that would eventually benefit all users at the centers.

Setting metrics for fraction of jobs run above a specific job size can be efficacious if the fraction and size are carefully set, and consequences predictable/acceptable. The recent, dramatic increases in time waiting for jobs to run, as reported by the NERSC user group, seems to call for re-evaluation of the apparently arbitrarily set requirement for the majority of jobs being larger than 512 processors jobs. There may be other factors, but something should be done.

"INCITE like" activity is valuable, exciting and holds promise for big returns in terms of advancement of the science it enables. Continuation of NERSC's participation is strongly encouraged because NERSC can support that kind of activity effectively, and it is an element of the portfolio of science support that sustains a center.

### **DOE Management**

The burden of compliance, oversight and guidance could be decreasing efficiency towards meeting NERSC's mission. This is a multi-faceted issue, perhaps doubled because of DOE and UC oversight. And, some may be driven as part of widespread increases in Federal and State oversight/compliance roll downs. However, cost effective

responses could include taking care of reporting at DOE headquarters or trusting LBL/UC oversight and eliminating DOE oversight/guidance.

A strategic approach, across the Office of Science community, to meet computational requirements appears to be called for to best meet the unmet demands. There seems to be a lack of computational resource provisioning for the middle and large classes of computing, those just shy of Leadership Class.

**Comments**

Quantify “savings” through code speed-ups – most centers have tried to do this, the NERSC staff may be up to the challenge, and it certainly would be a measure of the value added by NERSC.

Consider sharing the security model and practice developed at NERSC with the HEC and IT community.

Seek computational system cost information from other Federal Agencies (cost benchmarks) to add confidence to your procurement process.

Ensure staff and NERSC participate in advancing HEC technology through ongoing activity in development, test and evaluation. This is a necessary condition for sustaining a center and its users.

Strategic collaborations and HEC community activities are valuable in providing enhanced systems and services to users. (i.e. SP-XXL, SCOX, among many others you engage in)

A special thanks for excellent support in preparation for the review (i.e., the read aheads), logistics and individual accommodation for panel members.

## **Appendix 2**

### **National Energy Research Scientific Computing Center**

#### **NERSC Programmatic Review**

**May 17 – 19, 2005**

#### **Report**

**By**

**Frank Williams, Director**

**Arctic Region Supercomputing Center**

A panel of seven external reviewers was asked to individually comment on five major aspects of NERSC. The panel was also given freedom to look into any aspect of NERSC the members wished to examine, and comment accordingly. The external reviewers were asked to independently form their opinions and report their comments to the NERSC Program Manager, David Goodwin.

NERSC and David Goodwin provided advance support for the review. The information included annual reports; user surveys and responses; copies of materials to be presented during the site visit; links to the web page set up for the review to provide electronic access to information; notes from another, recent review-like activity; and especially:

NERSC's strategic plan for 2005-2010,

A draft RFP, Detailed Acquisition Plan, Benchmark Instructions and Acceptance Plan for the procurement of NERSC-5, and

An agenda for the two-and-a-half day review.

During the site visit the reviewers engaged NERSC management and staff in an ongoing dialog to gain clarity and express opinions on matters as they came up in presentations and in response to questions. Therefore, executive sessions were not taken as scheduled or for the purposes as expressed in the agenda. We did privately discuss our reporting methodology, our charge as expressed by the NERSC PM on behalf of DOE's Office of Science, our constraints and prohibition from purposefully coming to consensus among ourselves or with NERSC staff, and coordinating our further engagement with NERSC staff and LBL leaders during any remainder of the site visit. Accordingly, each member of the review panel will submit his own observations, comments and recommendations to the NERSC PM.

Copies of all the documents including the agenda were given to the PM, and therefore they are assumed to be appropriately incorporated in the summary report. The review team members were Walt Brooks, Division Chief, NASA Advanced Supercomputing

Division; Lawrence Buja, National Center for Atmospheric Research; Cray Henry, Director, DoD High Performance Computing Modernization Program; Robert Meisner, Office of Integrated Computing Systems, National Nuclear Security Administration; Jose Munoz, Deputy Division Director, Shared Cyber Infrastructure, CISE, NSF; Tomasz Plewa, The ACS FLASH Center, U of Chicago; and Frank Williams, Director, Arctic Region Supercomputing Center.

### **General Comments**

Throughout the on site review, the excellent quality of the leadership and staff at NERSC was abundantly evident and further solidified the conclusion drawn from examination of NERSC's record of accomplishment. The mutual appreciation for each other's contributions at NERSC provides a solid fabric supporting the prodigious accomplishments that have brought widespread, high regard for the center.

User Group activities strengthen NERSC and reflect commitment to serving users as the users carry out their computational research. The emphasis on User Surveys and appropriately reacting to them are exemplary. Frankly, the User Services aspects of NERSC are enviable.

NERSC should be commended for employing self-initiated reviews as checks on technology approaches being considered for deployment, and for periodic development of what is known as a Green book to define and help quantify user requirements.

Location specific leveraging is apparent in that the intellectual community of LBL and the academic community of the University of California Berkeley provide a world class, science driven environment for a science driven center. And, there is the wonderful pool of talented people in the San Francisco bay area that can be drawn upon for staffing the center.

### **Planning**

NERSC's strategic plan for 2006 -2010 is impressive in its comprehensive approach to providing Systems, Services and Analytics for a Science Driven center supporting research of interest to DOE's Office of Science. I found the message compelling and the plan viable within the constraints given by DOE. Every presentation and interaction on site furthered my confidence in the ability for NERSC to accomplish the plan. Rather than repeat comments that others are sure to make, or that may fit better elsewhere in this report, I'd point to two pieces of the plan that support my position. First, the plan calls for implementation of a facility wide file system – in clear response to NERSC's evolutionary role as a computation and data center and promising to provide efficiencies to users and data congestion across the center. Second, recognition of the importance of analytics to the understanding of data and computational results has led NERSC to plan for investment through redirection of staff work emphasis and incremental funding to meet the challenge. The strategic plan is filled with evidence of detailed, deep thinking and careful choices.

LBL Director Steve Chu informally asked the reviewers about possible metrics for NERSC, those that may supplement traditional measures applied to supercomputing centers. Two possibilities could be: number of publications acknowledging NERSC resources (systems, services, analytics) and breakthroughs enabled by utilization of NERSC resources. NERSC leadership and staff should give careful consideration of any proposed metric before implementation to avoid unintended consequences.

### **NERSC 5 Procurement and Process**

Based upon the draft, updated green book and generally acknowledged understanding of the community wide requirement for more high-end computing resources, NERSC 5 is a much-needed procurement. My comments should be taken as supportive of the procurement and the process proposed in general, but colored by my anxiety for meeting the demand as soon as possible.

The process is extremely thorough, perhaps to the point of being too thorough in that it may not be necessary to test many more technical performance aspects than are necessary to differentiate between technical solutions proposed. Having a comprehensive suite of tests for most desired performance features may be helpful in extremely close decisions, but what's realistically in the realm of possibilities for a very large, general purpose supercomputing system? The staff should be commended for putting together an understandable, rational, technical approach to the technical evaluation.

The best value procurement process portrayed to the review committee seems to be exactly the right approach. It is doable in the LBL context because of the valuable support and experience provided by the LBL procurement officer.

Concerns for parts of the process did emerge. A) The length of the process may be moot if NERSC/DOE cash flow cannot sustain the NERSC 5 procurement sooner than the projected acceptance, but ignoring that issue, the process seems very protracted. If there is need to delay the process in order to align the procurement in time with anticipated vendor technology releases, then the start of a faster process could be set accordingly. If there is concern over meeting milestones for performance of the acquisition, and the process is driven by those considerations, then I'd think the process is being driven by the wrong considerations. Thus the case for completely delegating the procurement authority to LBL or UC would be made. B) The time value of cycles provided to users might be looked at once again. Concerns raised after quick review by the team may be obviated with more reflection including the costs of porting and optimizing codes among other things. C) In my experience, the end game of negotiating is helped when there is clear competition. NERSC experience in the California and DOE context has evolved into a process that seems to bring good value in procurements. Without full appreciation for the constraints in the NERSC environment, I'd only suggest that NERSC think of ways to force competitive offers in a realistically limited field of potentially viable offerors. D) If there is direction from DOE Office of Science based upon considerations of efficiency of center operation with only one main architecture to support at each facility, or Office wide balance of architectures required, or national interest based limitation to

domestic procurements, then expressing those interests now would allow NERSC to include them in the best value evaluation and save unnecessary evaluation.

### **Operations/Budget/Staffing**

Potential negative impacts of flat funding in the face of increasing requirements for HEC were clear even though they weren't given emphasis in the documents provided or presentations on site.

I think NERSC makes a robust contribution to DOE computationally based scientists and thereby to the HEC community with as few as staff as possible. The recent reductions in staffing seem to leave NERSC understaffed. Increasing the diversity of architectures may further stretch the staff. Given the constraints, one needs to appreciate the dynamic shifts in work and flexibility of staff to meet the mission of NERSC.

### **Allocations**

Over-allocation is detrimental to the users and to the center's long-term reputation and usefulness. Cutting the allocations back is a hard decision that the Office of Science needs to make.

The allocation process itself appears disconnected from actual allocations given to users. I'd suggest the Office of Science find ways to more obviously incorporate assessment of computational readiness of applicants in the overall determination of allocations, or reduce the requirement for NERSC reviews of the proposals.

Long term health of the NERSC, and appreciation for Office of Science nimbleness, could be dramatically assisted by giving the center a fraction of the computational allocation to use at its discretion to meet contingencies and fortuitous opportunities. Most center directors have this discretion, allowing them to take advantage of situations that most often bring high returns to the center. For example, trading runs of each other's benchmark codes would double the knowledge of system performance for participating centers and that would eventually benefit all users at the centers.

Setting metrics for fraction of jobs run above a specific job size can be efficacious if the fraction and size are carefully set, and consequences predictable/acceptable. The recent, dramatic increases in time waiting for jobs to run, as reported by the NERSC user group, seems to call for re-evaluation of the apparently arbitrarily set requirement for the majority of jobs being larger than 512 processors jobs. There may be other factors, but something should be done.

"INCITE like" activity is valuable, exciting and holds promise for big returns in terms of advancement of the science it enables. Continuation of NERSC's participation is strongly encouraged because NERSC can support that kind of activity effectively, and it is an element of the portfolio of science support that sustains a center.



## **DOE Management**

The burden of compliance, oversight, and guidance could be decreasing efficiency towards meeting NERSC's mission. This is a multi-faceted issue, perhaps doubled because of DOE and UC oversight. And, some may be driven as part of widespread increases in Federal and State oversight/compliance roll downs. However cost effective responses could include taking care of reporting at DOE headquarters or trusting LBL/UC oversight and eliminating DOE oversight/guidance.

A strategic approach, across the Office of Science community, to meet computational requirements appears to be called for to best meet the unmet demands. There seems to be a lack of computational resource provisioning for the middle and large classes of computing, those just shy of Leadership Class.

## **Comments**

Quantify "savings" through code speed-ups – most centers have tried to do this, the NERSC staff may be up to the challenge, and it certainly would be a measure of the value added by NERSC.

Consider sharing the security model and practice developed at NERSC with the HEC and IT community.

Seek computational system cost information from other Federal Agencies (cost benchmarks) to add confidence to your procurement process.

Ensure staff and NERSC participate in advancing HEC technology through ongoing activity in development, test and evaluation. This is a necessary condition for sustaining a center and its users.

Strategic collaborations and HEC community activities are valuable in providing enhanced systems and services to users. (i.e. SP-XXL, SCOX, among many others you engage in)

A special thanks for excellent support in preparation for the review (i.e., the read aheads), logistics, and individual accommodation for panel members.

## **Appendix 3**

### **National Energy Research Scientific Computing Center**

#### **NERSC Programmatic Review**

**May 17 – 19, 2005**

#### **Report**

#### **By**

**José L. Muñoz, Deputy Director  
Division of Shared CyberInfrastructure  
National Science Foundation**

#### **Overview**

NERSC is a world-class computational resource provider facility. NERSC provides a rich collection of computing and storage resources; more importantly they provide the necessary services to enable the scientists and engineers to fully execute their research agenda. This is all accomplished with a very lean staffing level in an atmosphere of reduced budgets and the need to respond to external pressures making demands that are at times burdensome and provide no real value.

NERSC has managed to attract and retain a world-class staff that provides to the user expert insight into how their code could be altered to better take advantage of NERSC's computer assets.

#### **Planning**

NERSC does an excellent job of planning and equally important in keeping its staff and users aware of what potential plans are. Presented were NERSC's strategic plan for the next five-year (2005-2010). These plans focused on addressing DOE's request to reduce their staffing level of the current 66 FTEs.

NERSC has an active user group activity and annually request user inputs, via surveys, which is then used by the staff to plan for changes in their current modes of operation in order to respond to user issues.

#### **Findings/Recommendations**

1. In light of reduced staffing and flat budgets it is imperative that NERSC increase its collaborative activities. There is the perception that NERSC is isolated from the community outside its very loyal, and well served, user community and this must be corrected.
2. Form strong (stronger?) alliances with the other DOE Office of Science laboratories, the laboratories from the DOE NNSA organization, with other large supercomputing centers such as those found at NSF and HPCMO and with international supercomputing facilities. This is especially necessary in light of reduced staffing. Strategic collaborations should be formulated.

### **Acquisition planning**

NERSC presented their very detailed plan for acquisition of their next large-scale system, NERSC 5. This author was very impressed with their approach to system acquisition and how that system is to be targeted to meet the needs of their science and engineering community. I was especially intrigued with the use of quantifiable metrics required by potential bidders in order to make for a fair selection process. However, this reviewer has some concerns with the process as presented (this does not detract from this activity, indeed I'd like to see the NSF centers incorporate many of the NERSC ideas).

#### Findings/Recommendations

1. At 2+ years, the timeline is too long. NERSC should look at what can be done to significantly reduce the time required to acquire a system to less than one year. This should be doable if the system to be acquired exists in one form or another. In its current formulation there will be significant loss opportunities as a result of the long time lag.
2. While the acquisition process is very detailed (perhaps overly so), it fails to capture a key consideration: the cost of converting from their current programming paradigm to perhaps a completely different one. Effort should be made to “somehow” explicitly factor this into the evaluation. All things being equal this could be the determining factor in selecting one system over another.
3. Acquisitions should look at DOE architectural diversity and acquire a system that increases that diversity so that the DOE science and engineering community is better served. For example, don't buy another X1 as one already exists at ORNL.
4. There should be increased system prototyping/testbed activities and not reduction in those efforts as currently planned. In addition there should be increased use of simulation and modeling to explore the architecture and system configuration space.
5. Running both capability and capacity jobs on the same machine is causing undo tensions, typically at the expense of the capacity computing community. There needs to be acquisition planning for capacity systems in order to handle “smaller” users. Capacity user needs might best be met by clusters and do not require the tailored systems required to address capability users.
6. Their hybrid approach to system design is to be commended. However, it could run the risk of acquiring a system that might be adequate to a large number of users but not excellent for a specific community. This should be monitored.
7. NERSC should consider negotiating with more than one vendor at the same time in order to create an atmosphere of competition and gain greater value for the money.

### **Staffing/operations budget**

At DOE's request, NERSC is, unfortunately, being required to reduce staff. As it currently stands the NERSC operation is extremely lean (perhaps too lean), as a result any staff reductions must result in a change services that will eventually impact the quality of services the user sees.

#### Findings/Recommendations

1. Staff reductions must be made strategically, and this appears to be the case. However, I don't feel the staffing should be reduced below the current 66 FTEs. Doing so will have a negative impact on provided services.
2. NERSC is increasing the number of systems while at the same time facing a reduction in staffing. This is a situation that cannot be sustained.
3. The user community should be informed of any staffing refocusing prior to implementing these changes so that they can prepare accordingly.
4. NERSC's strategy of reducing the number of domain specific experts is not the right way to go. Those experts provide keen insight into specific codes and are able to have fruitful dialogs with their user community because they all speak the same language. If anything the number of domain experts should be increased.
5. The move to analytics is a good one, however I feel that focusing on visualization is not the best move. Visualization gains much of the attention because it is easier to demonstrate and it has a certain inherent appeal. However, I feel that rather than focus on the visualization aspects of analytics NERSC should focus on feature extraction and data mining. This has the potential of benefiting a greater number of users and is complementary to any visualization activities (provide hints to the users as to where he might look for interesting data). Staffing should be acquired/trained to address this area.
6. I applaud NERSC's recognition of increasing the staffing for software support as this is an area that typically gets under-staffed. At NERSC about 50% of the support time is currently in the area of software support so increased staffing there is correct.

### **DOE Management**

DOE headquarters has to not only provide NERSC oversight, but they must also respond to federal government guidelines (e.g. OMB/OSTP) and must do so in an era of reduced funding an external pressures.

### **Findings/Recommendations**

1. There appears to be considerable angst over the issue of who is the "capacity" and who is the "capability" resource provider at DOE. This is counter-productive. Clearly both ORNL and NERSC are providing capability-computing cycles. Capability computing is multi-dimensional and DOE's currently architecturally diverse set of computing resources should continue.
2. Something must be done about the current set of OMB metrics and the artificial need to increase the number of processors used by jobs. This is counter-productive in that it has resulted in increased job queues and as a result jobs are waiting days for access to their cycles. A metric that more accurately reflects the NERSC mission of providing the necessary resources to enable the execution of science should be used. NERSC and DOE should look at metrics used by other organizations such as the HPCMO and use or modify those metrics.
3. Additional metrics that might be considered are major publications and citations.
4. DOE needs to address the issue of providing capacity computing. While capacity computing may not always grab the headlines, it is a necessary element of the entire computing resource fabric. Using so-called capability machines as capacity providing engines is not cost effective.

5. DOE reviews should be minimal and focused in order to minimize disruptions at NERSC.

### Allocation process

Considerable time was spent discussing the DOE/NERSC allocation process. The area of allocation is very complex and covers a very broad spectrum of issues ranging from how many cycles particular users might be allocated to how long jobs might way in a queue to security issues.

### Findings/Recommendations

1. The allocation process is currently opaque and this situation must be remedied. It needs to be documented and publicized so that the user community is fully aware of what are the steps that are being followed to address their requests.
2. DOE and/or NERSC must do a better job of making sure that users pass the necessary security test for access to the system. Randomly selecting PI's and making sure that the necessary credentials are in place could address this.
3. How does DOE and/or NERSC make sure that ITAR guidelines are being followed? Mechanisms should be developed to provide for such a check.
4. At the NSF supercomputing centers we allow the center director to set aside 10% of the resources to be used at his discretion. This could be used to target specific domain areas, provide resources to industry, meet unexpected needs, etc.
5. As described, the CORP step of the process is not effective and should be eliminated. There was no correlation between the award score and provision of cycles. The NERSC recommendation of having NERSC staff review the request **and** to use that recommendation should be seriously considered.
6. Look at other scheduling algorithms in order to reduce the long wait times on Seaborg. HPCMO apparently has mechanisms they use and should be considered.
7. Incorporate checkpoint/restart in order to provide greater flexibility in job preemption. This is apparently already being looked at and should continue.
8. Incorporate an automated code analysis mechanism, perhaps using IPM, that would look at the jobs executing and identify jobs that are in need of improvements with respect to efficiency.
9. New users should automatically be identified for assistance to get them started along the right direction. This should be started during the STARTUP allocation process where users are given 20K hours.

## Appendix 4

### National Energy Research Scientific Computing Center

NERSC Programmatic Review  
May 17 – 19, 2005

Report  
by  
Walter Brooks  
National Aeronautics and Space Administration NASA  
Advanced Supercomputing Division  
NASA Ames

#### I Overall Planning Strategy

1. User Interaction: Approach of steady contact with users, surveys and the Greenbook is to be commended
2. Vision: as stated in the strategic plan is very good science driven services, architecture- View that the complexity is increasing and as a result need tighter integration of the computer scientist and the “physical scientist” – large programming is becoming more and more heroic
3. Facility Evolution: Overall view of the challenges that will be faced by scientist and engineers in DOE who rely on their systems is very good. This is complimented by a deep understanding of the SOA in HEC
4. Acquisition: strategy is too conservative see detailed comments below
5. User support: NERSC recognizes that the increase in complexity of the codes and concurrent increase in the size of the teams and the computing challenges will require that the computational scientist become directly involved in the science teams.
6. Collaboration: Initial presentation did not have a strong emphasis on collaboration. Later a document was produced that listed 40+ areas of collaboration but many of these seemed like information exchange. Believe that NERSC should consider strategic alliances with critical mass teams to try to “import “ technology from other centers. This will be especially critical with the significant reduction in staff in many areas.
7. Testbeds: Advanced testbeds whether at NERSC, other centers or other agencies are critical to the rapid and effective acquisitions
8. Facilities: New building appears to be an excellent long-term plan that is well integrated into the overall strategy for the lab but the flat budget may remove one of the drivers for the building. / The plan appears to be missing budgets for reallocation of the offsite assets, and building unique costs associated with large supercomputers
9. Security: Need to pay attention to emerging trends in sensitivity of the allowed users on the system. Right now the PI has responsibility for certifying CO-I’s nationality and other key data. Recommend spot check on users. Education of

user base on applications and data sensitivity (ITAR data). NERSC should try to get in front of this problem

10. Metrics-publications are one measure of output; you have to distinguish between supporting and enabling. There were at least a dozen examples given of truly enabling science calculations that could not have been done without a system like NERSC need to find a way to highlight this to stakeholders as well as the general public.
11. Mission Statement: After some dialogue we were given a statement of DOE vision for NERSC as its capability-computing center. If this is true, the work load being delivered by the programs to the facilities is overly driven by a capacity work load that might more appropriately be run on a loosely coupled cluster but the current DOE ecosystem of computing does not appear to include enough of these “in expensive” assets

## **II Please Comment on the acquisition planning- philosophy and approach**

Acquisition process briefed has been used effectively on the two previous NERSC procurements and at LLNL (original model?) and will result in an effective acquisition. However there were a number of areas in which NERSC may want to take a look at streamlining or modifying the best value process

1. Current acquisition team has no Viz/analytics team members. Emerging trends in modeling and simulation and vendor architectural roadmaps indicate that the next acquisition could consider the inclusion of embedded data analysis and visualization systems in the NERSC5 systems. This is not necessarily required by the possibility will not be considered unless the team has the proper composition.
2. The current process appears to be very protracted. At its most extreme starting from the formation of the original team it can be as long as 3.5-4 years until the proposed system in place in it's “highest impact” configuration. This means that the current Seabourg system if maintained until then will represent an extremely high \$/Flop system.
3. Consider developing /or using a “BAA” like vehicle that allows rolling yearly buys of smaller systems that are clearly available and vetted. I believe an analysis of the use of this approach (similar to the DoD mod office approach) may prove that this actually produces more integrated TFlop-years over a given 5 year period than the lumpy buys of large systems with 5 year lifetimes which often have extremely poor price performance in their last 2 years
4. As presented to our team, the current targets that were labeled as “aggressive “ (requirement of looking for 7TF sustained on the SSP benchmark in the FY07 time frame) seem to overly conservative. Based on 10% efficiency the current NASA Columbia system would deliver a sustained performance of 5-6TFlops in mid FY05. Consider a goal of 15 to 30 flop sustained.
5. Reexamine how long assets are maintained. –Not sure that using a sunk costs rationale for continuing to support 5 year old machine is valid – two years of maintenance might replace current assets and significantly reduce power and floor space requirement

### **III Please Comment on the operating budget and staffing and skill mix**

There are some disturbing trends in budget and staffing

1. Current flat line budget result in an imbalance in personnel and H/W costs that make it impossible to implement NERSC-6
2. The Current system is being run extremely cost effectively the staffing costs are approximately 40% of the overall budget which compares well with the top US and international computer centers and is significantly better than many.
3. I believe that the cuts in users services are too deep considering the challenges that lie ahead with NERSC 5 as well as the deployment of the NSC systems
4. The investment in Visual Analytics is visionary and will be required to meet the challenge of providing true integrated simulation environment over the next 5 years with the dramatic increase in code and data complexity
5. Operations: The most troubling trend is the growth in the average wait time in the NERSC queues. The “phase change” is driven by multiple factors, over allocation, late acceptance of the cluster, pressure to run big jobs at high priority to hit artificial processor metrics and too large a user base for the system to name a few. If this result is not corrected a large percentage of the user base will find that they can deploy a relatively small cluster in their lab and get a faster turnaround than the supercomputer, although this is the right local optimization the sum of all the users doing this results in a proliferation of low utilization- high maintenance (full cost) mini clusters.

### **IV Please comment on the Allocation Process**

1. Current allocation process which includes 30% over allocation, highest priority for large jobs and creation of a large number of relatively small users and accounts has led to excessively long “rot times” in the system queues. 4-7 day wait time in relatively for relatively small jobs will drive the customer base to build their own local clusters since a 50-100p lab system dedicated to these small jobs could outperform the current Seaboard by factors 2 or more.
2. DOE appears to be ignoring or deeply discounting the technical proposal evaluation process that NERSC provides. There appears to be no correlation between the technical score given and the selection of the final proposals that will be allocated time on the NERSC systems.
3. The allocation process appears to be following a methodology that gives full allocation to modest and small jobs and on average decreases the allocation of the largest jobs by factor of 3-4X. There is a basic conflict in the instruction to increase the number of jobs that run over 512p and than to significantly decrease request allocation for these users. Users who have scaled their science from 64-640p processors require a 10 X increase in allocation in order to allow and equivalent level of exploration of the problems at this scale

### **Please Comment on DOE /NERSC interaction DOE management:**

Two example became apparent during the review that make it apparent that DOE has crossed the line from telling the Lab *what* it wants it to do to *how* it should do it’s job.

1. Direction to stop work on emerging cluster technology represented by systems such as Alvarez resulted in a lack of experience with vendor products and the challenges of cluster systems that is a major contributing factor to the slow deployment of Jacquard system.



2. It appears that DOE strongly steered the NERSC team towards significant cuts in personnel and there appears to be a drive and suggestion that even deeper cuts are possible. As a result during a period in which complexity of architectures and systems is increasing we were briefed on cuts to users services by a factor of 2X. Based on my knowledge of HEC center operations NERSC appears to be extremely well run and has a lean / knowledgeable staff. Budget pressures may require a reduction in both H/W and personnel but NERSC is NOT an inefficient group with a bloated staff. My personal feeling are that the cut from 66 to 60FTE will seriously impact user services but any further cuts will jeopardize the viability of NERSC as a full service simulation center.

**-Mission statement:** Conflicting statements as to whether the NERSC mission is capability or capacity computing made it difficult for the review team to get a clear understanding of NERSC role in servicing the DOE science community. It might be better to couch goals in terms of mission critical science and engineering computing  
-Clear uncertainty as to whether any aspect of advanced technology development, prototyping or research is allowable.

**-Allocation:** DOE appears to be ignoring or deeply discounting the technical proposal evaluation process that NERSC provides. There is no correlation between the technical score given and the selection of the final proposals that will be allocated time on the NERSC systems.

**-Allocation:** The allocation process appears to be following a methodology that gives full allocation to modest and small jobs and on average decreases the allocation of the largest jobs by factor of 3-4X. There is a basic conflict in the instruction to increase the number of jobs that run over 512p and than to significantly decrease request allocation for these users. Users who have scale their science from 64-640p processors require a 10 X increase in allocation in order to allow and equivalent level of exploration of the problems at this scale. It appears that there is no documented approved allocation process. This makes it hard to determine if the queuing approach and job size and priority are being developed in a way that supports the underlying objectives of the DOE allocation approach.

**Budget-**DOE/OMB flat budget to 2010 puts the program out of balance and will insure that the NERSC -6 procurement does not provide the necessary increase in computing hours and scale to keep pace with the projected growth in user requirements.

**Reviews** – estimate that they spend 1—15 % of their budget preparing just for reviews – OMB EVM compliance as well as other security mandates indicate that it is plausible to estimate that a majority of the staff spends over 25% of their time on these review and reporting activities.

Recommend that DOE work directly with NERSC to streamline review and reporting requirement to help mitigate the effects of the 10% staff reductions

**Cyber security:** The requirements and effects of broad cyber security policies directed primarily at desktops and management and administrative personnel can have dire consequences if applied in a broad brush to research facilities and research teams.

Suggest that DOE top management attempt to create a board (at the CIO level) that

attempts to create appropriate security controls the NERSC environment, which includes advanced networks, supercomputers and data storage that are pushing the state of the art.

## Appendix 5

### National Energy Research Scientific Computing Center

#### NERSC Programmatic Review

May 17 – 19, 2005

#### Report

by

Lawrence Buja

National Oceanographic and Atmospheric Administration

National Center for Atmospheric Research

Conclusions from May 2005 NERSC Review  
Lawrence Buja (southern@ucar.edu)

#### *Q: How well is NERSC performing its mission*

NERSC is doing an excellent job in fulfilling its mission as the capacity computing center for the Department of Energy's scientific research efforts. NERSC management was found to be highly effective and should be recognized for its fine leadership through which it has developed a world-class computing center supported by a reliable infrastructure and an enviable array of user services.

NERSC management and operations has shown themselves to be very responsive to their users. NERSC has developed a very strong user community that is actively involved with center planning via the NERSC Users Group and annual user satisfaction surveys. The broad DOE research community has come to regard NERSC as a high productivity center that provides its users with extensive support and high reliability.

High-end, INCITE-like, activities that stretch NERSC's capabilities were viewed very positively. The breakthrough science resulting from the INCITE program is speeding up scientific discovery and enabling science domains to begin working on the advanced research topics that will be their future.

Looking ahead, their new "Science Driven" theme resonates well with their mission. NERSC management is clearly aware of the significant challenges ahead and presented a well thought out roadmap for addressing them in a resource-limited environment.

#### *Q: 1. Planning for the NERSC-5 upgrade and beyond, assuming a dollar-flat, \$38 million per year budget from FY05 thru FY10*

NERSC managers have developed a detailed plan for maintaining their exceptional level of service in a resource-constrained environment. The NERSC operation is clearly well reviewed, both externally and by internal teams. For example, NERSC substantially

altered its design for the Facility Wide File System in response to findings from NERSC commissioned review.

NERSC management has built a strong organization that is capable of confronting all of the NERSC's future challenges. However, achieving the goal of petaflop computing capability in a flat funding scenario will pose a significant challenge. If current budget trends continue, it is not clear that there will be sufficient resources for a significant NERSC-6 upgrade in 2010. .

Security was recognized an increasingly important topic and this is reflected by increased staffing planned for this area. Given the natural tension between productivity and security in a center supporting so many remote users, NERSC has done an exemplary job to date maintaining an open, yet secure site. NERSC's expertise in this area will be tested by the substantial security challenges that are still to come. On the horizon are far-reaching regulations on foreign national access to Federal computers that may have significant impact on NERSC community. A potential risk noted by the panel is NERSC's exposure to becoming involved in security incidences from having users unwittingly violating ITAR statutes on open NERSC machines or from program PI's failing to sufficiently verify their researcher's visa status. NERSC should elevate current user security awareness by implementing security education programs for users that are in place at other DOE centers such as LANL. At a higher level, to reduce the danger of overly restrictive security directives coming down from the DOE CIO with little or no input from NERSC, NERSC should take a leadership position in the development of DOE-wide computer strategies

***Q. The Request For Proposals (RFP), the proposals review, and source selection process prior to the NERSC-5 RFP.***

The NERSC team is well poised to carry out the NERSC-5 Procurement. The procurement process is extremely rigorous and is being carried out by technical and contract specialists with a great deal of experience from previous acquisitions. There is sound emphasis on acquiring a "balanced" system and the choice of a very wide range of user applications as the benchmark suite. The procurement team appears to have realistic expectations from vendors in terms of their ability to benchmark the application codes.

While the Jacquard acceptance has been rocky, it shows that NERSC is pushing the technology envelope with its second tier systems, without exposing its primary mission-critical systems to excessive risk. NERSC is using the lessons learned in the Jacquard acquisition for the NERSC-5 RFP. Items noted by the panel include:

While the N5 acquisition needs to be carried out with great care, the length of the procurement seems much too long, introduces extra variables/risk. Both NERSC and DOE management should seek ways to drastically reduce this time.

NERSC would benefit by researching negotiating strategies with similar large computer acquisitions at NSF, DOD and NASA supercomputing centers. Strategies that were mentioned include soliciting proposals based on maximum performance for a fixed price rather than minimum price for a fixed performance level. It was noted that the current Flops target should be reevaluated in light of the price/performance delivered with the NASA Ames Columbia acquisition. The timed delivery schedule SSP measure should take into account the concept of present value and the time value of cycles, making cycles delivered (or not delivered) early in the project more valuable than the same cycles later in the project. In the end-game, consider changing negotiation to play vendors against each other. Finally, it was suggest that NERSC seek final costing information for any similar system from other Federal Agencies.

It is important ensure that the NERSC-5 acquisition is (and is perceived by the community as) a very open and non-predetermined process. While it is imperative that DOE headquarters is actively involved throughout the selection process, there was concern about the possibility of last-minute intervention circumventing NERSC's careful selection process.

Finally, while both of the NERSC-5 contract and technical evaluation teams appear strong, NERSC management must ensure that there is a very close working relationship and open communications between the two groups

***Q. The operating budget (which includes the upgrades), especially the staffing levels and skills mix for NERSC-5 and beyond***

NERSC management presented a thorough overview of their current operation and how this will evolve over the next five years. They are running a lean, service-orientated, organization. In fact, when compared to other sites, current support staffing appears too lean. Current NERSC services are highly rated by user surveys. NERSC has been proactively responding to the recent budget contractions by centralizing web services and automating its project administration to require less staff.

NERSC has done an impressive job balancing future needs with the available resources and the current skills mix appears reasonably well balanced. One-on-one discussions with NERSC employees showed a high level of staff moral, confidence and goodwill towards management.

However, the current lean staffing will be further stretched due to the higher levels of support that will be required by the predicted increase in the architectural diversity of future NERSC computing platforms. This will stress the consulting group particularly hard right at the time when much greater proactive involvement is needed to ensure that the complex community simulation codes are able to run efficiently on the new systems.

Reductions in NERSC's benchmarking and new technology evaluation staff are a regrettable loss and may have contributed to the rocky Jacquard acceptance. It was noted that cutting future architecture research/analysis investment is basically a "going out of

business” posture. If NERSC does leave this areas unstaffed, they should create strategic partnerships with other federal computing centers that are maintaining these types of groups to retain access to expertise in this area.

NERSC’s new Analytics services holds great promise for assisting DOE researchers needing to condense, interpret and present results from their data intensive simulations. However, with the current budget realities, the return on investment of this group should be measured and evaluated carefully in relation to other programmatic priorities.

***Q: The allocation process.***

The technical execution of NERSC’s allocation review process is very thorough. Proposals are vetted through several layers of review and NERSC goes to great lengths to ensure that the projects running on their machine are both appropriate to NERSC’s mission and using codes that run efficiently on NERSC’s computing platforms. The allocation process was recognized to be evolving. It was suggested that the NERSC director consider adopting a practice common at other centers that gives the director oversight over a small reserve allocation to be used at his discretion. Actual publication of the draft NERSC allocation document, with special attention called to new or changed policies, will complete a fine effort by the services staff.

However, on a higher level, the allocation award process that occurs subsequent to the proposal review process appears to be flawed and disconnected from the allocation review results. The results from careful technical scrutiny of a proposals merit seem to have little correlation with the distribution of final awards. The allocation award process should be altered to correct this.

This year’s over-allocation of the NERSC computing resource is detrimental to NERSC. The problem is being compounded by the delays in the Jacquard acceptance. Policies linking following year allocation awards to previous year full allocation use will have to be modified if the current high level of computing resource contention continues. This is likely to have long term negative impact on the NERSC community as researchers become frustrated with long queues and move to alternative sites or build their own clusters that give them better time to completion.

***Q: 5. The DOE management.***

It is clear that the ASCR strategic plan needs to be revisited to define a unified set of goals, metrics and allocation processes that makes the best use of the leadership and capability center resources for meeting the DOE science communities leadership, capability and capacity computing needs. A coordinating operational plan that addresses the complete range of DOE high-performance computing needs will relieve the leadership and capability centers from having to struggle to define convoluted individual

allocation procedures in an attempt to locally optimize what is really a global resource management problem.

Between reviews by the University of California and the Department of Energy the burden of compliance oversight being placed on NERSC appears to be large, perhaps too large. There was also the sense that DOE Management is being too prescriptive with NERSC operations, making tactical level decisions about the specific NERSC staffing mix, hardware choices and production machine queue priorities in addition to providing strategic guidance.

### *Other Items*

NERSC managers indicated that more accurate metrics are needed to define mission success. This topic is a difficult one that all centers struggle with and it is well beyond scope of this group. That said, our discussions concluded that the general philosophy must center around some measurement of science outcome such as publication or citation volume or numbers of breakthroughs enabled. Unfortunately, the time-scale of such metrics can be very long. Another measure of the return on investment of taxpayer dollars could be a weighted sum of: stated scientific goals achieved, % resource utilized, job turnaround goals met and user satisfaction.

While NERSC was able to demonstrate extensive external collaborations, there is a lingering impression and management sensitivity that NERSC is still viewed as being isolated. NERSC should make efforts to develop significant strategic level collaborations and build strong ties with computing, networking and security groups in other DOE programs and in other federal agencies. NERSC should endeavor to adopt whole products suites from these other groups whenever possible. This will be of particular importance as new hybrid systems arrive and need to be integrated in a tight budget climate.

In this time when the centers are being asked to do more with less, efforts to increase the efficiency of resource utilization can have large, measurable benefits. NERSC should automatically profile running codes to identify inefficiencies and provide incentives, such as software engineering support, to help those codes make better use of the resource.

## Appendix 6

### National Energy Research Scientific Computing Center

NERSC Programmatic Review

May 17 – 19, 2005

Report

by

Bob Meisner

#### Observations Resulting from a Visit and Discussions with NERSC Facility Managers

Allow me to preface my observations with a comment that I was most impressed with the esprit-de-corps I observed throughout my visit. Without exception, everyone that I had the pleasure of talking to had positive attitudes and genuinely liked their work. They were proud of NERSC's contributions to enabling new science and excited about being part of the team. Maintaining that vigor in an era of constrained resources will be a demanding leadership challenge, but they have the talent to succeed.

#### Planning and Strategic Plan

The Strategic Plan was well organized and appeared accurate in its assessments of their: mission space, projected technologies and trends that will impact their ability to succeed, and organizational structure needed to accomplish their mission in a resource-constrained environment.

Under Paragraph 1.3, A Key Resource for the DOE Office of Science, I don't believe the statement that "NERSC provides focused support for these teams, with the goal of bridging the software gap between currently achievable and peak performance on terascale platforms, as was explicitly stated in the SciDAC plan" is possible. I would suggest changing the statement to something they plan on attaining, such as, "...the goal of vastly narrowing the software gap between currently achievable and peak performance on terascale platforms by doubling sustained performance every three years on applications requiring greater than 2048 processors." I don't profess to know that these are the right numbers to have as a goal and suspect that even this more modest objective statement might not be achievable with the resources provided. Regardless, I offer them as a more realistic illustration of what NERSC might obligate itself to do.

#### Acquisition Planning

If I understand the purpose of NCS-class systems, they are acquired to focus on a subset of the center's workload, which can be off-loaded from the NERSC-x systems, and still achieve best value for sustained system performance. This approach is reasonable and good. I would suggest that the program should track some measure that would indicate whether the systems acquired for these purposes meet their stated purposes. My intent here is not so much to prove that NERSC made the right decision on a single particular



procurement, but to document lessons learned that could be used to improve price-performance when matching system architectural attributes and applications over time.

The NERSC-5 strategy of focusing on hybrid supercomputers is a reasonable approach. The acquisition approach is a trusted and proven process that will result in good value in support of the strategy. However, the length of the process until final delivery introduces significant risks in terms of on-time delivery and projected costs. A judgment as to the acceptability of these risks is the purview of the program. I don't believe they will cause the program to fail, but may lead to technical and political challenges. Regardless, it is not clear to me that this process must be changed for the NERSC-5 procurement; but, I would recommend that the program evaluate alternative strategies and approaches for NERSC-6.

Two aspects of the NERSC acquisition planning processes were laudatory. They pay close attention to systems balance, not only in terms of processor, memory and interconnect attributes, but also for systems storage and networks. Further, they capitalize on the benchmarking work of others across the high-end computing community to increase the likelihood that they will receive best value in their acquisitions for the applications that they support.

#### **Operations Budget, Skills Mix and Staffing**

Over the recent past, NERSC center has operated with a lean staff, and has proudly and effectively served their constituency. Recent planning assumptions for a flat budget through 2010 have caused the leadership to reevaluate and adjust their operating skills mix. Their plan going forward into 2010 appears to be a reasonable one for focusing on their mission.

I would add one minor caution in reviewing the cost scenarios shown in the strategic plan. Escalating personnel costs due to even a low 3% inflation factor will cut into available services more significantly than shown in the table. Aggregating staff salaries and overhead for FYs 2000-2005 and 2006-2010 makes it appear that personnel costs are holding steady as a percentage of the total NERSC budget, yielding 40.4% and 40.8%, respectively. Such reporting masks an annual trend from FY2005 thru FY 2010 of 37.2, 38.4, 39.9, 40.7, 41.9, 43.2.

#### **Allocation Process**

The review process that precedes allocations needs reengineering. Analysis shown during presentations indicated that evaluations do not impact allocation decisions. Consequently, the considerable amount of time spent conducting evaluations does not appear to add value to the allocation process.

Taking into consideration the above observation, it suggests that processes outside the scope of this review might also benefit from reengineering. Not knowing the details of the process that determines what research projects get funded, but taking into account the above observation, one wonders if the process for allocating computational resources is

essentially disconnected from process for approving research proposals. If it is, should it be?

Presentations showed over-allocation of available cycles during the past year. This had a negative impact on the center's reputation for superior service and an adverse effect on center moral, but will probably not be devastating if it is corrected in the coming year.

Having a goal of running applications that use greater than 512 processors more than 40% of the time is an appropriate goal for a center with a computational mission of enabling new science. But, in the long term this might not be sufficiently demanding. As computational models and methods progress, one would expect them to demand increasing numbers of processors and/or memory; thereby, putting increasing demand on capability resources. There will always be a tension between available resources (read funding constrained) and applications demands. The program needs to address this balance in a constrained resource environment, and possibly adjust the run time goal. This should be done as part of a more comprehensive approach to building a strategic plan for complex-wide use of systems for scientific discovery (see second paragraph under DOE Management)

### **DOE Management**

DOE management of NERSC appeared very effective. In an era of constrained resources, the program office and NERSC leadership worked synergistically to redirect limited resources to accomplish the mission of enabling new science. While reductions and redistribution of personnel, and constrained spending for capability computing platforms could potentially reduce services available to the science community, the plans presented to the COV were intelligent and will likely result in excellent mission support given the working constraints.

Over the past 3-5 years the Office of Science computational complex has made great progress in providing quality cycles for scientific computing. Examples presented included results from the INCITE program and large calculations accomplished on NERSC machines. However, programs that produced these results are apparently maturing and changing, including the addition of Leadership Computing. As the complex enters a new era of leadership/capability/capacity computing it is not clear how the elements of the complex will interact to provide an integrated environment for enabling new science. This suggests the need for a strategic plan for the computing complex.

## **Appendix 7**

### **National Energy Research Scientific Computing Center**

**NERSC Programmatic Review  
May 17 – 19 , 2005**

**Report  
by  
Tomasz Plewa  
Computational Physics and Validation Group  
ASC Flash Center  
University of Chicago  
Chicago, IL 60637  
tomek@flash.uchicago.edu**

Chicago, May 23, 2005  
Reviewer Report Letter

#### **Overview**

The Review has taken place May 17-19, 2005. The first two days of the Review included a number of very detailed, high quality presentations. Presentations have focused on several aspects of the NERSC operations: mission, science, planning, budget, acquisition of new computing systems, and management. I would like to express my appreciation for all the effort the NERSC management and staff has put in presenting the amount of accomplishments, problems and issues that they deal with in their everyday work and thinking about the Center's future.

The NERSC mission is to help advancing science by offering access to integrated supercomputer infrastructure and high-end high-quality computational services. Scientific projects supported by NERSC are numerous (more than 200 each year) and originate within DOE Office of Science, other governmental agencies, as well as from across scientific community including U.S. and foreign academic institutions, groups, and individual researchers. The latter characteristic asks for the Center's resources being accessible to participants located at remote sites, providing excellent network services and appropriate security protocols, and in some cases extends to making software components (licenses) available to the NERSC users.

The NERSC user community anticipates and expects the highest quality of services (computing resources, consulting, training) from the Center. My overall assessment is that NERSC has achieved exceptional level of user services and serves as a prime example of science- and user-oriented supercomputing facility. NERSC has created a unique, in my opinion, group of users (organized in the NERSC User Group) that offers NERSC invaluable feedback regarding its daily and long-term operations. User surveys

are extensive and the number of responses clearly indicates the strength and activity of this unique group of computational scientists.

### **Science**

NERSC has a large base of experienced users conducting research in several diverse fields of scientific computing. It is important to define and use in the Center's evaluation metric that goes beyond already stated qualitative "excellent" grade. A common practice adopted by several countries to evaluate performance of their scientific institutions is the number and impact of publications that acknowledge support (funding, computing time) obtained from such organizations. NERSC tracks and collects information of such publications, the number of which exceeded 2,000 in 2004. I am very impressed by the sheer number of publications supported by NERSC. However, this metric can be improved and made quantitative by providing a cumulative number that weights each publication by its corresponding impact factor (as defined by the Science Citation Index). Other, more involved, metrics can be constructed, but including the ISI impact factors would offer significant improvement over currently used total number of publications.

In some cases, however, no single metric can be applied. In particular, one would hope that among such a large volume of published work a few truly exceptional breakthrough discoveries are present. One such example is the Center's support for tokamak research that has resulted in construction of high-quality numerical algorithms and application codes allowing for engineering-type design studies of these extremely complex and expensive devices. Identifying three new highly stable beam configurations resulted in funding and building three new experimental facilities that will study actual stability of these theoretically predicted configurations.

INCITE program is one more example of aggressive approach to computations. This high-risk high-return large-scale computational program allows a narrow group of carefully selected researchers break barriers preventing them from exploring regimes in which existence of new phenomena is anticipated or significantly extending our current knowledge. One additional measure of success of the INCITE program might be awards given to foreign (non-US based) institutions. It is conceivable that a number of such applications will be growing with time. Presence of such awards would indicate that idea of INCITE has a truly global dimension.

Overall, NERSC has achieved excellent level of support for scientific computing that can be given as an example to other supercomputing centers.

### **Planning**

Creation and maintenance of unified center-wide filesystem, high usability and performance of the high-performance storage system, and excellent record of cybersecurity at NERSC are remarkable achievements and elements highly praised by users. Providing a sustained level of funding and development in these areas are essential for success of NERSC operations.

Regarding acquisition of future main computing systems, direct extensive testing at vendors site or purchase of small testbed systems (if not available at other supercomputing centers) seems desirable and will mitigate possible risks. Such policy would help avoiding problems currently experienced with the NCS machines (jacquard).

Evaluation of the Center's performance and planning of its future operations should reflect its strong emphasis on science. Currently used metric (level of utilization and amount of time used by large jobs) does not necessarily reflect actual user/science needs. Smaller jobs and projects required specific resources (e.g. long runs) may be delegated to other systems leaving the Center's main computing system available to the most demanding applications. A limited queue backfill procedure, already in place on seaborg, may be used to keep utilization level at acceptable high value without offering users a 300-long list of waiting jobs.

### **Acquisition of NERSC-5**

To meet its mission goals, NERSC has to periodically update their computer systems. That includes not only their major computing system but also the support structure (software, data storage, network) that has to match the new resource and allow for efficient support and management of user applications.

The existing main computational resource, seaborg IBM SP2 system, is currently number 21 on the TP500 list and is quickly moving down the list. It has to be noted that some desktop machines offer computational power equivalent to that of a single node of Seaborg making NERSC flagship system much less attractive for computations. I anticipate that despite excellent record of NERSC, users will quickly start seeking access and eventually migrate to other supercomputer centers. Such erosion of the user community appears as a serious problem. Commitment to funding, acquisition, and availability of NERSC-5 combined with continued high-quality of user services will prevent such negative trends from developing. Also, motivation and performance of the NERSC staff is a certain function of the Center's appreciation among community of computational scientists. Current high level of recognition cannot be sustained without retaining NERSC's status of a leading supercomputer facility.

As the user community matures with time, demand for computing cycles constantly grows. This is not only due to a constant increase in the number of users, but also due to constant growth of complexity of scientific applications and ability of the user codes to use computer resources more aggressively (i.e. parallelization, data growth). This trend clearly emerged from the user jobs statistics presented during the Review. Overall, there is a strong need for immediate and substantial increase in computer resources available at NERSC.

The acquisition procedure of NERSC-5 is one of major current operations under way at NERSC. The operations is led by the Procurement Team. The Procurement Team is following a set of procedures aimed at guaranteeing that an optimal decision will be made and the selected machine will offer best value. The following elements may need to be considered to improve effectiveness of the current procurement process.

The Procurement Team may want to seek additional information about vendors and their products from sister governmental agencies (DOE, DoD, NASA, NSF). Such information appears to be of great value in negotiations with offerors.

The Team has created a benchmark suite that is representative of 85% workload at NERSC. The Team may want to use that collection of benchmarks to measure performance of several more systems currently installed or to be installed in the near future to gather performance and efficiency information and collect first-hand experience (portability).

NERSC may consider conducting negotiations with several offerors at the same time rather than with one vendor at a time. If such alternative procedure is adopted, potential increase in workload seems only temporary and will be offset by benefits of having up-to-date information that can be instantaneously exploited in negotiations.

The procurement process seems very long (several years). Extending the process so much does not appear necessary or justified, makes the process less focused with precious information about vendors and available computer installations quickly becoming outdated. Moreover, some elements of the current RFP document are based on extrapolated data or are too precise. The requirement that the system has to deliver a certain number of Tflops (sustained) seems unnecessary and may actually decrease the level of competition between offerors. Providing only minimal and essential information to the offerors appears as a strategy offering NERSC desired advantage in negotiations.

Finally, NERSC may consider switching their hardware profile from a single large machine to two high-end systems installed at the same time. In such scenario, hardware updates are taking place more often offering users access to the most recent computing technology. The two systems could be combined into a single (perhaps heterogeneous) installation through a dedicated internal fast network given the common filesystem and archive storage are already available at NERSC.

### **Budget**

Related to the above mentioned NERSC-5 acquisition issue is the proposal of a constant level of funding for several next years. Adoption of such policy is likely to affect overall perception of NERSC. One possible conclusion following from adopting a flat-budget policy seems to be that excellent service to wide community of computational scientists may not be a future priority of NERSC. In addition, lack of funding is likely to have a negative effect on morale of the Center's staff that naturally expects funding level being tied to their performance.

Budget constraints, and in particular decrease in level of funding, forced NERSC to reorganize and modify profile of its user support operations. This process appears to be conducted very aggressively with certain services (i.e., support to select science applications) being replaced by more widely accessible services. This seems to be the right strategy.

Although NERSC seems to be successful in preparing for future budgetary challenges, it has to be noticed that further reduction of user consulting services (from 4.5 FTE in the past to 4.0 FTE currently and 3.0 FTE in the near future) will have a widespread negative impact on regular users. This conclusion is supported by opinion expressed by NERSC consultants who feel they has already reached their maximum capacity. One may expect a growing dissatisfaction of users with consulting services, one of the most visible and appreciate activities of NERSC.

### **Resources allocation process**

Several observations followed from the NERSC Users Group report and discussion with NERSC management.

NERSC resources are currently overallocated. The current prime computing facility, seaborg, is not offering enough cycles to address the needs of all major group users.

The high level of utilization (one of criteria used by the Office of Management and Budget in evaluating NERSC) seems to be a mere by-product of a long waiting queue of small jobs. Although such a long waiting queue allows to patch holes between large jobs and therefore achieve exceptional utilization level, waiting queue times reaching several (5) days do not really allow for advancing science at reasonable pace. This is especially true in case of projects involving relatively small jobs where quick turnaround time is essential (e.g. steering required in parameter studies).

This situation could be improved by installing additional smaller computing systems. Delays in deployment of New Computing Systems (i.e. jacquard) is truly unfortunate. Every measure should be taken to prevent such delays from occurring in the future. Given high and still growing demand for computing cycles, investment in proven technology seems to be optimal at this point in time.

There seems to be a disconnect between computing proposal review process and final allocation decision. Scores given to the proposal by reviewers seems completely uncorrelated with amount of computing resources eventually granted. This trend seems persisting over at least last few years indicating that not only the review process does not work but the issue is not being addressed. One would expect more weight being given to scientific merits of proposals.

It is recognized that NERSC tries to meet certain externally defined performance criteria such as aforementioned high level of machine utilization level and amount of computing time consumed by large jobs. The latter criterion in particular offers a potential to skew statistics, obscure naturally developing trends, and lead to pathological situations (users running inefficient jobs on large partitions). Deployment of automatic runtime system of parallel and floating point performance monitoring system seems highly desirable. Such system may help users and consultants in porting and optimizing applications for NERSC computing systems. Startup allocations of 5,000 SUs may not be sufficient for

development of large applications and some flexibility in regard of granting additional time is advisable. Mid-year progress review of startup allocations seems desirable.

The data from the aforementioned proposed runtime performance monitoring tool may help reviewers to eliminate or otherwise promote certain groups of users. Such information may also help identifying projects requiring different type of hardware than available at NERSC (e.g. parallel vector processor systems at ORNL).

### **DOE management**

There seems to be a number of reviews (DOE, LBNL, Univ. California) taking place at NERSC. These activities consume substantial amount of time (25% according to rough estimates) and energy of the NERSC's management and staff. The burden associated with reviews can be lowered by either decreasing frequency of reviews (one in two years rather than once per year), limiting their and/or by scheduling joint reviews. The latter combination seems the most attractive.

Tomasz Plewa, Group Lead



## Appendix 8

August 20, 2005

Dr. David Goodwin  
Mathematical, Information, and Computational Sciences Division  
SC-21  
U.S. Department of Energy  
Germantown Building  
1000 Independence Ave., SW  
Washington, DC 20585-1290

Dear Dave,

First, thank you for arranging the programmatic review of NERSC in May. We found it extremely useful from several points of view. First, it caused us to do serious self-evaluation of NERSC's contribution to the DOE computational community and also how we currently deliver our services. Second, it caused us to look at the challenges, requirements and roles for NERSC over the next five-year period. Rather than just responding to the reviewers' specific questions, we decided it was best to present a more comprehensive view of NERSC's performance and plans. We greatly appreciate the reviewer's time, insight and comments. We found their advice to be invaluable.

In preparation for the review, NERSC produced a five-year plan that covers our efforts from 2006 to 2010. This plan extends our previous plan (2001-2006) and includes new activities. It assumes a stable budget at the FY 05 level, and provides deliverables to DOE in terms of systems and services.

This memo addresses the comments and advice of the review team. Although the comments were all individual and not a consensus view, many of the comments overlapped or were similar. Thus, in Appendix A, we summarized and condensed the comments in the individual reviewers' letters, indicating how many reviewers made comments in similar veins. Indeed, as you see, some of the areas had conflicting advice from the reviewers.

Many of the individual comments are specific and can be handled at a low level. So, rather than address every comment, this response will focus on what we believe are the strategic and core issues raised by the review as well as the issues that we believe need high priority and visibility.

- 1) NERSC has achieved an exceptional level of user services and serves as a prime example of a science and user oriented supercomputing facility.**  
The reviewers were unanimous in the view that NERSC has done an outstanding job addressing the DOE/SC's computational requirements and helping the DOE/SC computational community use large-scale resources to address very challenging problems. NERSC appreciates this recognition and will continue to strive to provide outstanding systems and services.

**2) NERSC's cyber security record is exemplary and NERSC should teach others how to do it.**

NERSC agrees. NERSC will continue to strive to provide excellent Cybersecurity will providing a flexible and productive range of services for the science community. NERSC has shared our experiences, practices and technology and will continue to do it within the limits of resources. NERSC staff presented security tutorials at the last three SC conferences, at the Global Grid Forum and other venues. NERSC also co-hosted, along with NCAR, a workshop on recent cyber security incidents. The LBNL tools NERSC uses have been made available to the entire community, and NERSC continues to help improve these tools.

**3) Advanced testbeds, whether at NERSC, other centers or other agencies are critical to rapid and effective acquisitions.**

NERSC agrees. Hands-on experience is essential to making good long-term decisions. There are two areas of experience that are key to being able to rapidly acquire and field new technology. One is application conversion and performance and the other is system management. Both need test beds of moderate scale but not full scale.

NERSC has used its own testbed systems and information from testbeds at other sites for application evaluation. Application evaluation can be done effectively with remote systems. On the other hand, system management evaluation requires hands-on experience and attempts to implement on the testbed the production functionality NERSC systems need. Because of the nature of this work, and the fact there has been little system management evaluation information coming from other testbed sites, this requires testbeds at NERSC. Over the past nine years, DOE/SC has not supported any testbed activity at NERSC. Indeed, in 2000 LBNL invested nearly \$1M in non-DOE funds to implement a Linux cluster testbed. The experience gained in this testbed turned out to be critical in NERSC's ability to select NERSC-4, NCSb and NCSa, and was invaluable to the successful implementation of the recent NCS system. Unfortunately, LBNL's ability to totally underwrite other testbeds is essentially gone, particularly given the Lab's focus on funding a new computer facility for NERSC.

Hence, it is critical that DOE provide testbed resources at LBNL/NERSC that are tightly integrated with the NERSC production systems and staff.

**4) NERSC's 2006–2010 Plan is appropriate, comprehensive, compelling and "impressive."**

NERSC appreciates the reviewers' endorsement of our five-year 2006–2010 plan.

**5) The overall view of the challenges that will be faced by scientists and engineers in DOE who rely on their systems is very good.**

NERSC agrees there are serious technical and non-technical challenges to face in the future. We believe our plan will help us successful meet those challenges.

**6) NERSC's 2006–2010 Plan represents a deep understanding of the state of the art in high-end computing.**

NERSC appreciates the reviewers seeing the connections of our plan to the overall HPC landscape.

- 7) Analytics is an important new service — NERSC is “visionary” in identifying this area as being key.**

NERSC agrees.

- 8) The focus on hybrid supercomputers is correct .**

NERSC agrees but is not precluding a completely commodity or custom solution if those architectures demonstrate better value for the NERSC workload in the future.

- 9) The Facility Wide File System is important.**

NERSC agrees.

- 10) A flat budget with rising support costs is not viable in the long term (NERSC-6).**

This level budget will keep NERSC as a world-class facility for NERSC-5, but it is clear the level of funding will be insufficient to support a world-class system for NERSC-6, in large part due to growing infrastructure costs (power, cooling, etc.).

- 11) Staffing is too low.**

NERSC staff are highly effective, but it is clear that NERSC has operated for nine years on the very lowest edge of staff exhaustion. Even small reductions can have unexpectedly large impacts and limit the flexibility NERSC and DOE have to address future computational science needs.

- 12) Shifting staff from science support endangers the quality of service.**

NERSC agrees. NERSC has always valued a balance between science support and raw flop/s-bytes.

- 13) Shifting staff from domain experts is not the best solution to staffing caps.**

NERSC both agrees and disagrees. While loathe to eliminate any service, NERSC cannot take on new efforts without reducing staffing levels in others. Other programs (e.g., SciDAC) have made significant strides in this area, which is why NERSC choose it for reduction. On the other hand, NERSC would be happy to continue to resource these areas if increased funding was allowed.

- 14) NERSC is minimally staffed to understaffed now.**

NERSC agrees. See comments above. NERSC notes that several of the reviewers indicated their sites, with similar missions and scope, had about 85 to 90 FTEs to carry out the responsibilities NERSC does.

- 15) The cut from 66 to 60 FTEs will seriously impact user services, but any further cuts will jeopardize the viability of NERSC as a full-service simulation center.**

NERSC agrees there will be an impact on some services at 60 FTEs, but believes most services will remain intact. The main area of impact is how many projects can be done in depth.

- 16) NERSC cannot sustain increasing the numbers of systems and keeping staff so lean.**

NERSC is very concerned that the direction from DOE to increase the number of systems will severely stress the staff, particularly at the reduced levels.

- 17) Aggregating staff salaries and overhead for FYs 2000–2005 and 2006–2010 makes it appear that personnel costs are holding steady as a percentage of the total NERSC budget, yielding 40.4% and 40.8%, respectively. Such reporting masks an annual trend from FY2005 thru FY 2010 of 37.2%, 38.4%, 39.9%, 40.7%, 41.9%, and 43.2%.**

Aggregating staff costs by 2000–2005 and then 2006–2010 shows the portion of the NERSC budget for staff remains about the same, but the NERSC budget did increase between those periods, so the absolute costs have increased somewhat each year. NERSC uses standard DOE escalation factors for salary and other costs, and indeed the trend pointed out is driven only by those planning factors. Despite the comments and the knowledge that more staff would provide valuable services, NERSC is not proposing more staff.

It is also important to note that the staff cost increases are not the major contributor to the increases in overall support. Rather, electrical costs, consumables, and other costs play a more significant role.

- 18) The approach of steady contact with users, surveys and the Greenbook is to be commended.**

NERSC will continue in these efforts and finds great value in them. We are also thankful to be blessed with an active and helpful user community.

- 19) There is no relationship between review of proposals and allocation decisions.**

This is really a DOE issue that must be resolved within the SAC and/or Office of Science. NERSC proposes eliminating the CORP review because it is time consuming for the science community and it does not appear to have any impact in the allocation decisions. A simple staff review, along with the required IPM performance data, should provide information as to the readiness and appropriateness of the code to program managers as they decide allocations.

- 20) Allocations are too small compared to the metrics of system.**

NERSC agrees and would like to work with DOE/SC to make allocation amounts for projects align with the ability of the project to run at scale. One way to do this is to allocate enough time to projects that have demonstrated the ability to scale to the required metric levels. NERSC believes that if the metric is “N% of time is used by applications that use  $\geq K$  processors,” then the DOE must allocate at least 1.5N% of time to applications that have demonstrated the ability to scale to K processors. For AY 06, this means that 60% of the time should be allocated to projects that have demonstrated the ability to run at  $\geq 768$  CPUs on Seaborg.

The other alternative, one that NERSC has successfully employed to reach the required metrics, is to discount and give preferred processing. This technique, although successful, is detrimental to some projects that do not scale and results in some level of overallocation.

- 21) Overallocation is harmful.**

NERSC agrees. In order to prevent the situation we have this year, NERSC will only commit to deliver time for systems already in production and will wait until a system is accepted and in production before augmenting the committed level to DOE. This approach worked very well in the past and prevented severe overallocation.

**22) The OMB metric is not sufficient to assure effective use of system for science.**

NERSC agrees. The single metric is not effective in guiding such a complex facility. NERSC will investigate metrics used at other sites and recommend one or more replacement metrics by AY 07.

**23) The NERSC-5 process is effective and will produce a good result.**

NERSC agrees.

**24) The NERSC-5 team does not have a visualization/analytics staff member on it.**

NERSC agrees, but notes that analytics is new service thrust and was not explicitly part of NERSC's plan at the time the NERSC-5 team was formed. The NCSa procurement team did include a staff member with visualization experience because that system was conceived to handle an analytics like workload. Since less than 2.5 FTEs are funded for visualization, it is not practical to have visualization staff on all procurements. As more staff focus on analytics, it will be easier to have a staff member from that area involved.

**25) Two aspects of the NERSC acquisition planning processes were "laudatory." They pay close attention to systems balance, not only in terms of processor, memory and interconnect attributes, but also for systems storage and networks. Further, they capitalize on the benchmarking work of others across the high-end computing community to increase the likelihood that they will receive best value in their acquisitions for the applications that they support.**

NERSC agrees and will continue along this path for NERSC-5. NERSC (NERSC-5) and HPCMP (TI-06) are the first HPC organizations trying to coordinate procurement and benchmarking activities. Vendors are extremely pleased to see this effort. NERSC believes this once again demonstrates HPC leadership and that the experience gained will help other agencies as well. We expect to expand on this coordination in the future.

**26) NERSC requiring performance for NCS is the right thing.**

NERSC agrees and continued to work with LNXI to assure the system was capable of reliable production. The system was partially accepted after completing all the acceptance tests, including an available test with users in late July. The system went into production service August 1.

**27) Continue to do factory testing.**

It is NERSC's intention to do factory tests for all major computational systems, but it is not always possible. A factory test is required for NERSC-5, and was for NCSb. However, in the NCSb, due to a backlog and limited infrastructure in IBM's assembly floor, requiring a factory test would have delayed delivery almost six months. In this case, the system will be built at NERSC and then heavily tested. NERSC intends to do a factory for NERSC-5.

**28) Align UC procurement with the NERSC-5 technical team.**

NERSC disagrees that procurement and the NERSC-5 technical team are not aligned. The procurement officer is an integral part of the procurement team and works on a daily basis with all technical staff and aspects of the procurement. There are weekly and sometimes semi-weekly meetings with the integrated team.

NERSC believes the comment was more due to the organization of the presentation than any real misalignment.

That being said, LBNL just hired a new chief procurement officer who is very familiar with best-value and HPC acquisitions. Indeed, he was involved with the NERSC-1 and NERSC-2 systems. We expect his appointment will create even better integration between NERSC staff and procurement staff.

**29) Consider being more aggressive on measured performance.**

NERSC data indicates the target of 7.5 TF SSP is aggressive, but NERSC will work with DOE to gather costing data from other agencies to verify this. An “apples to apples” comparison is difficult to make without actually reviewing the contracts. For example, systems have different balances and may then be upgraded, maintenance may or may not be included, and if it is included, it may have different terms and durations. Other mitigating conditions may apply, such as less stringent testing or providing a share of the system for vendor testing and benchmarking.

NERSC data represents the full system cost with three years of total support and aggressive production performance and reliability metrics. Even so, NERSC believes the data shows we get as good a deal as other major sites and we are aggressive but not unreasonable in performance goals.

Finally, regardless of the goal in the RFP, NERSC’s BVSS process allows us to achieve higher performance if there is an underestimate, and the measurements indicate higher performance is possible within the budget.

However, this is an important issue and NERSC will attempt to accumulate more data, particularly from the review sites, to align NERSC-5 goals. NERSC is also attempting to run the NERSC-5 benchmarks on NASA Ames’ Columbia system to assess the SSP value. This is one of the more aggressive systems in terms of price per peak TF, so we will have what might be considered a lower bound.

**30) Negotiate in parallel with multiple vendors.**

Nothing in NERSC’s process prevents this possibility. NERSC is not sure this is effective, and there are several significant procurement policy issues (e.g., auctioning, leveling) that have to be carefully reviewed if this were to happen. Nonetheless, NERSC will investigate this concept in more detail to understand the processes at other sites and determine if this is effective. We will do the review once the new Chief Procurement Officer starts in the fall.

**31) The NERSC-5 procurement process is too long.**

NERSC presented the entire acquisition project timeline that conforms to DOE and OMB guidance — from the first initial RFP planning meeting to final system acceptance.

NERSC starts early and the procurement team works part time while handling their usual responsibilities. It is feasible to condense the schedule since we work part time for the RFP creation as NERSC does not have dedicated performance or procurement staff.

While the schedule presented is long, it does compare favorably to other organizations, particularly when viewed in the context of NERSC using stringent performance metrics and assuring vendors deliver systems that operate to expectations. Further, it is important to point out that there is only a short time (mostly vendor build time) from the time of actual system specification (SOW) to the delivery, so the system being specified and its cost are extremely current.

For example, Cray Henry (HPCMP) estimated that the DOD Mod TI-06 performance team costs about \$1M to craft the benchmark and performance part of the RFP, the usability team consumes about 4 FTEs (~\$1M), and we believe there are some other costs. This gets DOD Mod to a contract, but does not include installation, testing acceptance, facility preparation, or final production readiness, all of which are in the NERSC-5 DME project.

In contrast, for the same order of magnitude systems, NERSC spends less than \$1M for RFP/benchmark creation, vendor marketing information and proposal evaluation, with another \$0.75M estimated for installation<sup>1</sup>, testing, acceptance and productization. NERSC is known to be very thorough in its testing to assure problems are detected early rather than by users.

Other sites may do it faster or may compare the time to acquire evaluation or limited use systems, but to have a comprehensive procurement for a production system, NERSC is extremely efficient.

Finally, it must be recognized the NERSC-5 procurement has a great deal of oversight and review. Following normal business practice for a system of this size, DOE contracting and DOE program review the RFP and, unless waived, the contract. NERSC-5 is a DME/EVM project that reports regularly to DOE and OMB. The NERSC-5 Project Execution Plan is also reviewed and approved by DOE. It is the only LBNL project selected for DOE's internal business audit, which requires an additional review of the acquisition plan which is not normally done. Finally, each review for NERSC-5 is done by the Berkeley Site Office, the Chicago Office and DOE Headquarters. Each review takes multiple months. All these reviews are included in the project flow time presented, extending the overall flow time by six or more months.

**32) Acquisitions should take into account architectural diversity so if another site has a system, NERSC should not get a similar system.**

NERSC disagrees. NERSC's job is to get the best possible system for sustained performance across the spectrum of the workload. During the era when Cray Research dominated the field, almost every site had the same architecture (Cray-1s, XMPs, YMPs, C-90s) with no major loss of effectiveness.

**33) Consider awarding two smaller systems at the same time.**

---

<sup>1</sup> Note: Site preparation is very system dependent, and major facility work, if needed, is not included in the \$.75M.

NERSC disagrees. When NERSC-5 is awarded, NERSC will have NERSC-3, NCSa, NCSb, the PDSF and visualization servers. Adding another system increases operational costs and does not significantly improve the ability to run the NERSC workload.

**34) It is important for NERSC-5 to be truly open and competitive.**

NERSC agrees and points out every major computational system at NERSC has been procured with wide open competition.

**35) The best value procurement process portrayed to the review committee seems to be exactly the right approach.**

NERSC agrees.

**36) Wait times are too long on NERSC systems and this risks pushing many users (those not “preferred”) to other solutions.**

NERSC agrees, but long wait times are the result of the overallocation discussed above. Until the allocation of time is aligned with the needed turnaround/productivity goals, it is not possible to have good throughput to all types of jobs. Despite the heavy load, NERSC is providing very good throughput and turnaround for the jobs DOE identifies as priority jobs.

**37) NERSC needs to review foreign national access and not rely on PIs.**

NERSC realizes this is important and will investigate ITAR rules, current and potential. After analysis and in consultation with DOE/SC, NERSC will adjust its policies and business practices accordingly.

**38) NERSC should automatically profile all codes and provide incentives for code improvements.**

NERSC is considering this, but it has implications for performance and user access to performance information for their codes. We will likely try system-wide profiling on NCSa to determine the appropriate information and tools to use.

**39) Using NCSa systems to offload the major NERSC-x systems is good.**

NERSC agrees, but is concerned with the number of systems the fixed staff can support to the quality of service NERSC users expect.

Again, we appreciate the reviewers’ suggestions and remain dedicated to making NERSC the best resource for large-scale computation.

Sincerely,

Horst Simon  
Associate Laboratory Director for Computing  
Sciences, LBNL

William Kramer  
NERSC General Manager  
LBNL



## Informal Aggregation of Reviewers' Comments

### General Comments — No need to address

- NERSC is highly effective. — Henry, Baju, Munoz, Plewa, Williams
  - NERSC has achieved exceptional level of user services and serves as a prime example of a science- and user-oriented supercomputing facility. — Plewa
  - Frankly, the User Services aspects of NERSC are enviable. — Williams
- NERSC is well managed. — Henry, Baju, Munoz
- NERSC has excellent quality of leadership and staff. — Brooks, Buja, Williams
- NERSC has an outstanding record of accomplishment. — Williams
- Good “esprit-de-corps.” — Meisner
- NERSC provides excellent support for science. — Plewa
- NERSC has the right balance of systems and services. — Henry
- INCITE support is very effective and important for science. — Buja, Plewa, Williams
- The emphasis on user surveys and appropriately reacting to them are exemplary. — Plewa, Williams
- NERSC should be commended for employing self-initiated reviews. — Williams
- The mutual appreciation for each other's (staff and management) contributions at NERSC provides a solid fabric supporting the prodigious accomplishments that have brought widespread, high regard for the center. — Williams
- NERSC's cyber security record is exemplary and NERSC should teach others how to do it. — Henry, Buja, Williams
- Advanced testbeds, whether at NERSC, other centers or other agencies, are critical to rapid and effective acquisitions. — Brooks, Plewa, Williams
- New building appears to be an excellent long-term plan that is well integrated into the overall strategy for the lab, but the flat budget may remove one of the drivers for the building. — Henry, Brooks
  - The plan appears to be missing budgets for reallocation of the offsite assets and unique building costs associated with a large supercomputer. — Brooks

### Allocations

- Allocation process is “broken” and needs “re-engineering.” — Henry, Buja, Meisner
  - NERSC proposal management and review is thorough. — Buja
  - No relationship between review and allocation. — Henry, Buja, Brooks, Munoz, Plewa, Williams, Meisner
    - Eliminate the CORP and use only staff reviews. — Munoz
  - No relationship between science funding and allocation. — Henry
  - Allocations too small vs. metrics of system. — Henry, Brooks, Plewa, Williams
  - Director should have ability to allocate some time. — Buja, Munoz, Williams
  - Current process is opaque — should be transparency. — Henry, Munoz
  - Automatically provide all startup users with extra help. — Munoz
- Processes outside of the review need re-engineering. — Meisner
- Overallocation is harmful. — Buja, Brooks, Plewa, Williams, Meisner
  - Adds to user frustration. — Buja

### DOE Management

- DOE management appears effective. — Meisner
- DOE needs a clear, transparent strategic plan for computing. — Henry, Meisner

- NERSC does not get clear guidance as to role and mission. — Henry
- Needs to address goals, metrics and allocation process for leadership and high-end capability. — Buja
- DOE needs to provide inexpensive capacity computing as part of the plan. — Brooks, Munoz, Williams
- Need to monitor amount of capacity resource so there are enough funds for true capability resources. — Munoz
- Conflicting statements as to whether the NERSC mission is capability or capacity computing — may be better to couch goals in terms of mission-critical science and engineering computing. — Brooks
- Counterproductive to worry about capability and capacity descriptions. — Munoz
- Clear up uncertainty as to whether any aspect of advanced technology development, prototyping or research is allowable. — Brooks
- DOE management too “prescriptive.” — Buja
- DOE “crosses the line” in telling the lab how to do things. — Brooks
  - Directed staffing cuts in advanced technologies. — Brooks

### **Metrics**

- OMB metric is not sufficient to assure effective use of system for science. — Henry, Munoz, Brooks, Plewa, Williams
  - NERSC should develop new metrics. — Henry, Munoz
- OMB metric of 40% at 512 way is appropriate — and may not be sufficiently demanding in the long run. — Meisner
- Very impressed with the number of publications supported by NERSC. — Plewa, Williams,
  - Improve by weighing numbers with ISI impact factor. — Plewa
- DOE reviews: look under “Resourcing — non-staff”
- NERSC should develop a throughput metric rather than utilization metric. — Henry, Munoz
  - Relevant to science outcome. — Buja
  - Papers and citations. — Munoz, Williams
- Numerous enabling science calculations that could not have been done without a system like NERSC need to find a way to highlight this to stakeholders as well as the general public. — Brooks, Plewa, Williams

### **NERSC-5 and General Procurement**

- Very impressed with approach to system acquisition and how that system is to be targeted to meet the needs of the science and engineering community. — Munoz
- Intrigued with the use of quantifiable metrics required by potential bidders in order to make for a fair selection process. — Munoz
- Like to see the NSF centers incorporate many of the NERSC ideas (for procurements). — Munoz
- NERSC-5 process is effective and will produce a good result. — Henry, Buja, Brooks, Munoz, Meisner
- Team is well qualified. — Buja
- Two aspects of the NERSC acquisition planning processes were “laudatory.” They pay close attention to systems balance, not only in terms of processor, memory and interconnect attributes, but also for systems storage and networks. Further, they capitalize on the benchmarking work of others across the high-end computing

community to increase the likelihood that they will receive best value in their acquisitions for the applications that they support. — Meisner

- Team does not have a visualization/analytics staff member on it. — Brooks
- System balance the right goal. — Buja, Plewa
- NERSC requiring performance for NCS is right thing. — Buja
- Learning from NCS is good. — Buja
  - Continue to do factory testing. — Plewa
- Align UC procurement with technical team. — Henry
  - Contract and technical team appear to be working very closely. — Buja
- Use SSP with net present value. — Henry, Buja
- Consider being more aggressive on measured performance. — Buja
  - 15–30 TF/s for SSP is a reasonable goal. — Brooks
- Run benchmarks on more systems before release. — Plewa
- Negotiate in parallel with multiple vendors. — Buja, Munoz, Plewa, Williams
- Process too long. — Henry, Buja, Brooks, Munoz, Plewa, Williams, Meisner
  - Maybe due to OMB measures. — Henry
  - Will not cause a failure, but may lead to some more challenges. — Meisner
    - Do not need to adjust for NERSC-5, but consider for NERSC-6 — Meisner
  - Adjust start time based on when the system should arrive (funding, technology, etc.) and do it in a more concentrated and faster manner. — Williams
- Have code conversion cost as an explicit review criteria. — Munoz, Williams
- Acquisitions should take into account architectural diversity so if another site has a system, NERSC should not get a similar system. — Munoz, Williams
  - DOE has to express their plan clearly to be included in the procurements — Williams
- DOE too involved — DOE should either buy the systems or provide only general guidelines. — Henry
- Consider awarding two smaller systems at the same time. — Plewa
- Important for NERSC-5 to be truly open and competitive. — Buja
- Consider having a yearly purchase as in DOD Mod. — Brooks
- The staff should be commended for putting together an understandable, rational, technical approach to the technical evaluation. — Williams
- Not clear all the tests are real discriminators of the best systems. — Williams
- The best-value procurement process portrayed to the review committee seems to be exactly the right approach. — Williams
- Valuable support and experience provided by the LBNL procurement officer is key to success. — Williams.
- Consider getting cost information from other agencies. — Williams, Plewa

### **Operations**

- NERSC has an excellent relationship with its users. — Henry, Buja, Plewa
- Wait times are too long on NERSC systems and this risks pushing many users (those not “preferred”) to other solutions. — Brooks
  - Continue to emphasize checkpoint/restart as part of the solution. — Munoz
  - Look at other scheduling schemes — e.g., DOD Mod. — Munoz
  - Add more smaller systems for smaller jobs. — Plewa
- Need to review foreign national access, not rely on PIs. — Henry, Munoz

- Risk is from uses violating ITAR regulations. — Buja, Brooks, Munoz
- NERSC should automatically profile all codes and provide incentives for code improvements. — Buja, Munoz, Plewa
  - NERSC may be able to do what others have not been able to — quantify savings/impact of code speeds. — Williams
- Using NCS systems to offload NERSC-x systems is good. — Meisner
  - Should have a way to track whether having NCS systems is moving smaller work off the large system. — Meisner
- Despite many collaborations, NERSC is perceived as isolated and should seek out collaborations and make them more visible. — Buja, Brooks, Munoz, Williams
  - Collaborations are more important as staff decreases. — Williams, Brooks, Munoz

### **Planning/Strategic Plan**

- NERSC does an excellent job planning and keeping DOE and users informed. — Munoz
- NERSC Strategic Plan is appropriate, comprehensive, compelling and “impressive.”— Henry, Buja, Brooks, Plewa, Munoz, Williams, Meisner
  - Overall view of the challenges that will be faced by scientists and engineers in DOE who rely on their systems is very good. — Brooks, Meisner
    - Deep understanding of the SOA in HEC. —Brooks
  - Analytics is an important new service — “visionary.”— Henry, Buja, Brooks, Munoz, Williams
- Focusing on visualization is not the right thing for analytics — rather focus on data mining and feature extraction. — Munoz
- Should change the wording for supporting the highly parallel performance. — Meisner
  - Need careful consideration if this is more important than evaluation and future technology. — Buja
  - Focus on hybrid supercomputers is correct. — Meisner
  - Facility Wide File System important. — Plewa, Williams
  - NERSC staff effort in supporting software is appropriate. — Munoz
- NERSC recognizes the increase in complexity of the codes and concurrent increase in the size of the teams and the computing challenges. — Brooks

### **Resourcing — Nonstaff**

- Reconsider how long older assets are maintained. — Brooks
- Current facility is being run extremely cost effectively. — Brooks
- Flat budget with rising support costs is not viable in long term. — Henry, Buja
  - It means no NERSC-6 unless corrected. — Brooks
- NERSC spends a significant staff effort in reviews and audits. — Henry, Buja, Brooks, Munoz, Plewa, Williams
  - Too much oversight. — Buja
  - Some reviewers estimated NERSC spend 25% of resources in audits, compliance and reviews. — Brooks, Plewa, Buja
  - No more than 10% of time should be spent in compliance, reviews and audits. — Brooks

### **Resourcing — Staff**

- NERSC attracts and retains world-class talent. — Munoz

- Location-specific leveraging is apparent in that the intellectual community of LBNL and the academic community of the University of California, Berkeley provide a world-class, science-driven environment for a science-driven center. — Williams
- There is a wonderful pool of talented people in the San Francisco Bay area that can be drawn upon for staffing the center. — Williams
- Cost of FTEs lower than in many other places. — Henry
- NERSC staff and management have a healthy and mutually appreciative relationship. — Buja
- In light of budget limitations, the choice to focus on “widely accessible services” and sacrifice “support for select science areas” is the right strategy. — Plewa
- Staffing is too low. — Henry, Buja, Brooks, Munoz
  - DOD Mod MRSCs are about 85 FTEs. — Henry
  - Long-term damage. — Henry
  - Staffing costs are very efficient compared to other national and international centers. — Brooks
- Shifting staff from science support endangers the quality of service. — Henry, Baju, Brooks, Munoz, Plewa
  - Shifting staff from domain experts is not the best solution to staffing caps. — Munoz
- NERSC is minimally staffed to understaffed (now). — Henry, Buja, Brooks, Munoz, Williams
  - NERSC is NOT an inefficient group with a bloated staff. — Brooks
  - The cut from 66 to 60 FTEs will seriously impact user services, but any further cuts will jeopardize the viability of NERSC as a full-service simulation center. — Brooks, Williams
  - Cannot sustain increasing the numbers of systems and keeping staff so lean. — Munoz, Williams
  - Lack of funding will have a negative morale impact on NERSC staff. — Plewa
- Aggregating staff salaries and overhead for FYs 2000–2005 and 2006–2010 makes it appear that personnel costs are holding steady as a percentage of the total NERSC budget, yielding 40.4% and 40.8%, respectively. Such reporting masks an annual trend from FY2005 thru FY 2010 of 37.2%, 38.4%, 39.9%, 40.7%, 41.9%, and 43.2%. — Meisner

### **Users**

- NERSC has created a unique group of users (organized in the NERSC User Group) that offers NERSC invaluable feedback regarding its daily and long-term operations. — Plewa
- Approach of steady contact with users, surveys and the Greenbook is to be commended. — Henry, Buja, Brooks, Plewa
- User Group activities strengthen NERSC and reflect commitment to serving users as the users carry out their computational research. — Williams

## Appendix 9

### June 16 Thank-You Letters for the Peer Reviewers

Dr. Frank Williams  
Arctic Region Supercomputing Center  
University of Alaska Fairbanks  
West Ridge Research Building  
909 Koyukuk Dr., Suite 105  
Fairbanks, AK 99775-6020

Dear Dr. Williams:

Thank you for your participation in the May 17-19, 2005, program review of the National Energy Research Computing (NERSC) Center at the Lawrence Berkeley National Laboratory.

Your contributions to the program review are sincerely appreciated.

Our intent is to use your findings and recommendations to improve the operations of NERSC.

Sincerely,

C. Edward Oliver  
Associate Director of Science  
for the Office of Advanced Scientific Computing Research  
Office of Science

cc:

R. Orbach, SC-1

S. Chu, LBNL

M. Strayer, SC-21.1

D. Goodwin, SC-21.1

Dr. Tomasz Plewa  
Computational Physics &  
Validation Group  
The ASC Flash Center  
The University of Chicago  
5640 South Ellis Avenue, RI 475  
Chicago, IL 60637

Dear Dr. Plewa:

Thank you for your participation in the May 17-19, 2005, program review of the National Energy Research Computing (NERSC) Center at the Lawrence Berkeley National Laboratory.

Your contributions to the program review are sincerely appreciated.

Our intent is to use your findings and recommendations to improve the operations of NERSC.

Sincerely,

C. Edward Oliver  
Associate Director of Science  
for the Office of Advanced Scientific Computing Research  
Office of Science

cc:

R. Orbach, SC-1

S. Chu, LBNL

M. Strayer, SC-21.1

D. Goodwin, SC-21.1

Dr. José L. Muñoz  
Deputy Director  
Division of Shared Cyberinfrastructure  
National Science Foundation  
CISE/SCI Suite 1145  
4201 Wilson Blvd.  
Arlington, VA 22230

Dear Dr. Muñoz:

Thank you for your participation in the May 17-19, 2005, program review of the National Energy Research Computing (NERSC) Center at the Lawrence Berkeley National Laboratory.

Your contributions to the program review are sincerely appreciated.

Our intent is to use your findings and recommendations to improve the operations of NERSC.

Sincerely,

C. Edward Oliver  
Associate Director of Science  
for the Office of Advanced Scientific Computing Research  
Office of Science

cc:

R. Orbach, SC-1

S. Chu, LBNL

M. Strayer, SC-21.1

D. Goodwin, SC-21.1



Mr. Robert Meisner  
Deputy Director  
Advanced Simulation and Computation  
National Nuclear Security Administration  
NA-114.2  
1000 Independence Avenue, SW  
Washington, DC 20585-0104

Dear Mr. Meisner:

Thank you for your participation in the May 17-19, 2005, program review of the National Energy Research Computing (NERSC) Center at the Lawrence Berkeley National Laboratory.

Your contributions to the program review are sincerely appreciated.

Our intent is to use your findings and recommendations to improve the operations of NERSC.

Sincerely,

C. Edward Oliver  
Associate Director of Science  
for the Office of Advanced Scientific Computing Research  
Office of Science

cc:

R. Orbach, SC-1

S. Chu, LBNL

M. Strayer, SC-21.1

D. Goodwin, SC-21.1

Dr. Cray Henry  
Director  
High Performance Computing  
Modernization Program  
Office of the Director  
Defense Research and Engineering  
1010 North Glebe Road, Suite 510  
Arlington, VA 22201

Dear Dr. Henry:

Thank you for your participation in the May 17-19, 2005, program review of the National Energy Research Computing (NERSC) Center at the Lawrence Berkeley National Laboratory.

Your contributions to the program review are sincerely appreciated.

Our intent is to use your findings and recommendations to improve the operations of NERSC.

Sincerely,

C. Edward Oliver  
Associate Director of Science  
for the Office of Advanced Scientific Computing Research  
Office of Science

cc:

R. Orbach, SC-1

S. Chu, LBNL

M. Strayer, SC-21.1

D. Goodwin, SC-21.1

Dr. Lawrence Buja  
National Center for Atmospheric Research  
1850 Table Mesa Drive  
Boulder, CO 80305

Dear Dr. Buja:

Thank you for your participation in the May 17-19, 2005, program review of the National Energy Research Computing (NERSC) Center at the Lawrence Berkeley National Accelerator Laboratory.

Your contributions to the program review are sincerely appreciated.

Our intent is to use your findings and recommendations to improve the operations of NERSC.

Sincerely,

C. Edward Oliver  
Associate Director of Science  
for the Office of Advanced Scientific Computing Research  
Office of Science

cc:

R. Orbach, SC-1  
S. Chu, LBNL  
M. Strayer, SC-21.1  
D. Goodwin, SC-21.1

Dr. Walter F. Brooks  
Division Chief  
NASA Advanced Supercomputing Division  
M.S. 258-5  
NASA Ames Research Center  
Moffett Field, CA 94035

Dear Dr. Brooks:

Thank you for your participation in the May 17-19, 2005, program review of the National Energy Research Computing (NERSC) Center at the Lawrence Berkeley National Laboratory.

Your contributions to the program review are sincerely appreciated.

Our intent is to use your findings and recommendations to improve the operations of NERSC.

Sincerely,

C. Edward Oliver  
Associate Director of Science  
for the Office of Advanced Scientific Computing Research  
Office of Science

cc:

R. Orbach, SC-1

S. Chu, LBNL

M. Strayer, SC-21.1

D. Goodwin, SC-21.1