

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

REPORT

Advanced Scientific Computing Research

April 27, 2005

Date of COV: April 4-5, 2005
Programs: ASCR Facilities Program (including the National Energy Research Scientific Computing Center, the Advanced Computing Research Testbeds, Leadership Class Computing, and the Energy Sciences network), and, network testbeds and research
Office: Advanced Scientific Computing Research (ASCR)
Agency: United States Department of Energy

Committee of Visitors Membership:

Jill Dahlburg, Naval Research Laboratory (Chair)
James Hack, National Center for Atmospheric Research
Kevin Mills, National Institute of Science and Technology
Philip Snyder, General Atomics
James Tomkins, Sandia National Laboratory

Executive Summary

The Committee of Visitors (COV) for the Department of Energy (DOE) Office of Science (SC) Office of Advanced Scientific Computing Research (ASCR) programs in facilities and network testbeds and research met April 4-5, 2005 at the DOE facility in Germantown, MD.

The COV is extremely grateful to the program officers and other ASCR staff who gave unstintingly of their time and knowledge to help the committee in its deliberations, and who thus enabled a review process that was both streamlined and effective.

Findings:

Based on presentations by, and interviews with program officers and management, and on examination of project information of the facilities and research programs under consideration, **the COV deems the reviewed ASCR facilities and networking research programs to be highly effective and well managed.**

ASCR is one of the principal players in U.S. high performance computing (HPC); many world-wide view the Office as the home of high performance computational science. The ASCR 2005 COV observed that the Public Law 108-423 (Nov. 30, 2004): ‘.. to require the Secretary of Energy to carry out a program of research and development to advance high end computing’ (where ‘the term Secretary means the Secretary of Energy acting through the Director of the Office of Science of the DOE’) is well-enacted per the facilities and network research programs of ASCR.

ASCR provides substantial value to the computational science community, with the cohesive portfolio of facilities and research activities under direction of that Office. Further, ASCR contributes significantly to areas of broad U.S. applications science research, for example, with 10% of National Energy Research Scientific Computing Center (NERSC) computational resources provided for the DOE SC Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program, and with a similar amount proposed for Leadership Class Computing (LCC) block time. The COV found that ASCR management applies vision and judgment throughout the areas under review, even with respect to the proactive handling of recent budget cuts.

The COV noted in particular that services to users provided by the ASCR facilities – e.g., the Energy Sciences network (ESNet) and NERSC – are outstanding.

Comment:

While funding for the overall ASCR program may be shrinking, the needs for HPC are growing, and HPC is becoming an ever more important part of the DOE SC portfolio. Thus, it is critical for ASCR to pay due regard to continuing a vibrant and synergistic program portfolio in the face of potential budget realities.

TABLE OF CONTENTS

	<i>page</i>
Executive Summary	2
1 Introduction	4
2 Network Testbeds and Research	5
3 Energy Sciences network (ESnet)	7
4 National Energy Research Scientific Computing Center (NERSC)	9
5 Advanced Computing Research Testbeds (ACRTs)	11
6 Leadership Class Computing (LCC)	11
7 Need for Capacity, Capability Planning	12
A1 List of Acronyms	13
A2 Committee of Visitors (COV) Resumes	14
A3 Committee of Visitors (COV) Instructions Letter	20

1. Introduction

The Advanced Scientific Computing Advisory Committee (ASCAC) for the Office of Science (SC), United States Department of Energy (DOE), was charged by Dr. Raymond Orbach, Director of the Office of Science, with assembling a Committee of Visitors (COV) to review the facilities program of the Office of Advanced Scientific Computing Research (ASCR) that includes the National Energy Research Scientific Computing Center (NERSC), the Advanced Computing Research Testbeds (ACRTs), the Leadership Class Computing (LCC) initiative, the Energy Sciences network (ESnet), and ASCR network testbeds and research. The ASCR 2005 COV met at the DOE Germantown location on Monday, April 4 and Tuesday, April 5, 2005.

The COV meeting opened on April 4th with a series of presentations by the DOE program officers. Walter Polansky provided a summary of the charge to the ASCR 2005 COV, and discussed the DOE conflict of interest policy. Michael Strayer presented an overview of the ASCR Facilities Program. Each program manager then presented hour-long program overviews. In each of the program briefings, emphasis was placed on describing the processes used to shape and manage the portfolios, and how these processes affect portfolio elements and impact portfolio standing in both national and international communities. The program managers who provided briefings to the COV were [with topic]: Gary Johnson [Leadership Class Computing (LCC) and ACRTs]; David Goodwin [NERSC]; Mary-Ann Scott [ESnet]; and, Thomas Ndousse-Fetter [Network Testbeds and Research]. Following these overviews, the COV deliberated upon the materials presented in closed session, and then met with the program managers to discuss preliminary findings and to convey a series of questions to be addressed on the morning of April 5.

On the second day of the review, the requested additional information was presented by the program managers and considered in depth. The additional material included: discussion of the resources leveraged to fund the Ultra-Science Network Testbed; the range of projects and categories in the network research portfolio; indicators of network research program success and achievements; overview of the approach considered by ASCR staff to provide the 100 Gbps backbone and 20-40 Gbps redundant site connectivity required for future DOE science; examples of capabilities provided by ESnet that are not available commercially; a discussion of the processes per allocating facilities resources among capacity and capability computing at NERSC; clarification of ACRT and LCC budgetary issues; overview of the relationship between the ACRTs and the LCC; discussion of the allocation mechanisms and guidelines for the ACRTs and for the LCC; and, ASCR plans for standardizing allocation policies across all of the DOE SC ASCR facilities.

The review findings and comments of the ASCR 2005 COV are presented in the body of this report. A list of acronyms is appended as Attachment 1. Resumes of COV members are provided in Attachment 2. The instruction letter for the ASCR 2005 COV is reproduced in Attachment 3.

2. Network Testbeds and Research

ASCR conducts a small (approximately \$5M/ year) research program focused on meeting future requirements of DOE science programs for transferring high volumes of experiment and computation data over long distances. This research program is organized into four main elements: (1) research, development, and engineering (about 35% of the funding); (2) evaluation of emerging technologies (about 20% of the funding); (3) experimental test beds for future high-capacity networks (about 25% of the funding); and, (4) experimental application prototypes within network test beds (about 20% of the funding). Of the approximately 50 projects identified by the program manager, Thomas Ndousse-Fetter, about 1/3 extend over a 3-6 year period, while 2/3 extend over a shorter, 2-3 year, period. The longer-term projects tend to focus on either fundamental research or on experimentation within high-capacity network test beds. The shorter-term projects tend to focus on either evaluation of emerging technologies or on experimentation with application prototypes and pilots.

Over the period of FY01 to FY04, 57% of funds for networking research were distributed to DOE laboratories, while 43% went to universities. During the period of FY04 to FY06, distribution of funds shifted somewhat, so that 70% were distributed to DOE laboratories, while only 30% went to universities. This shift in funding distribution appears to reflect a planned progression in the network-research program. During FY01 to FY04, the program emphasized fundamental research to develop technologies for high-speed transport protocols, network monitoring, and cyber security, which could improve the data throughput and network security available to DOE science programs. In the period FY04 to FY06, the program is shifting toward an emphasis on deploying, prototyping, applying, and evaluating the research results from the earlier period. Still, some funds are being devoted to investigate mechanisms for on-demand network connections with guaranteed quality of service. After FY06, one could expect funding distribution to shift again to address additional needs for fundamental research that are identified during the period of trial deployments and application pilots. This cycling of funds back-and-forth, between fundamental research and testing and evaluation in trial deployments, appears to be a sensible strategy for a network-research program with such a small funding level.

Overall, the program is well managed, yielding the DOE an exceptional return on its small investment. The exceptional return on investment (ROI) derives from several factors. First, the program focuses on the future requirements of DOE science programs, which concentrates funding on areas of most need. Second, the program has strong ties to other, related, research programs within ASCR. For example, the program jointly funds some projects with the ESnet program. Since ESnet is the main vehicle for providing networking services to DOE scientists, such joint research allows the networking research program to gain visibility into current operations of DOE scientists and to informally elicit unmet or looming requirements. In addition, DOE scientists and ESnet managers gain visibility into developments emerging from the networking-research program, and thus become potential volunteers to deploy and test new networking technologies. The networking-research program also jointly funds and otherwise interacts (e.g., through joint meetings of principal investigators) with ASCR research programs in collaboratories

and grid computing. Such interaction increases the visibility of network researchers into the needs of future applications, and encourages science programs to deploy application pilots over prototype networks developed in the networking-research program. These strong ties increase the synergy among research programs within ASCR.

In addition to strong ties with other ASCR programs, the networking research program exhibits significant interactions with other networking research programs in the U.S. Government, and most notably the NSF. The ASCR program manager serves as a co-chair of an interagency network research team composed of program managers from NSF, NSA, NASA, NOAA, NIH, and NIST, *etc.* These connections enable ASCR to identify research funded elsewhere that might be adapted for use within DOE networks. In fact, on more than one occasion, research initially funded by NSF has been extended by ASCR for deployment, testing, and further development within DOE network test beds. Engaging with other agencies also allows ASCR to relay future networking requirements of DOE scientists, which sometimes encourages NSF independently to fund fundamental research to meet such needs.

The ASCR networking-research program uses two strategies to increase ROI. First, the program manager leverages a significant amount of other funding. For example, ASCR is investing \$3.2M to fund an Ultra-Science Network test bed. The program manager identified additional funding, totaling \$14.6M, contributed by laboratories, universities, and corporations participating in the test bed. Second, the program manager focuses on transferring technology from the networking research program into other DOE activities. Approaches to technology transfer cover a wide spectrum: (1) distribution of open-source software; (2) creation of standards and specifications; (3) direct implementation in DOE science projects; (4) publication of results; (5) researcher exchanges and sabbaticals; and, (6) adding network requirements to procurement solicitations.

Beyond specific program-management strategies and interactions with other government research agencies, the program manager is proactive in developing a wide community of interest related to high-speed networking. For example, upon noting diminished interest and contributions to IEEE INFOCOM related to high-speed networking, the program manager volunteered to organize a high-speed networking research symposium in the upcoming INFOCOM conference. Creating and maintaining a wide community of interest should increase the research base upon which ASCR can leverage its own investment in networking research.

During the COV review, the program manager indicated that ASCR is planning by mid-year to set up an advanced-technology committee within ESnet. The committee will provide consulting to DOE science programs that use the network. The program manager suggested that a committee should be formed at the level of the DOE Office of Science in order to establish formal communications between ASCR networking research and DOE science programs. Such a committee could help ASCR identify future requirements for networking and would also serve as an enabler for the DOE SC community to find out about emerging networking technologies. The COV recommends that ASCR place a high priority on establishing a committee within the Office of Science, to promote formal

interactions between ASCR networking research and DOE applications science programs.

A headquarters level committee is being formed to give all SC program offices direct involvement in the management and policies governing a networking infrastructure that supports their programmatic needs at present, and into the future. The initial charter for this committee, called the Network Advisory Committee (NAC), has been developed. The target date for establishing the NAC is October 30, 2005. A key responsibility of the NAC will be to address the coordination of networking research activities with projected needs of SC's research programs.

Upon questioning by the COV, the program manager identified five programs of high-speed networking research being conducted in other countries (four in Europe and one in Canada). Most of these research programs appeared derivative of research underway in the U.S. ASCR provided little information or insight about the substantial high-speed networking developments underway in Asian countries, most particularly Korea, Japan, and China. The COV recommends that ASCR take steps to gain a comprehensive understanding of the posture of high-speed networking research within Asia.

ASCR agrees that such steps need to be taken. An early activity for the NAC mentioned previously will be to serve as a catalyst to further understand high-speed networking research activities throughout the world, with an emphasis on Asia, and their possible relevance to SC's mission.

From ASCR's perspective, our current understanding of research and education networks in Asian countries is that their main goal is to provide connectivity to Starlight in Chicago, which is the rendezvous point for international research networks.

3. Energy Sciences network (ESnet)

The Energy Sciences network provides reliable communications infrastructure and leading-edge network services to DOE scientists and their collaborators around the world. ESnet provides direct connections to all major DOE Office of Science sites, enabling effective use of DOE research facilities and computing resources independent of geographic location.

ESnet and its predecessors at DOE have provided advanced networking services to the DOE science community since the early days of wide-area networking. ESnet was created as a result of a survey of networking needs initiated in FY85 by Dr. Alvin Trivelpiece, then Director of the Office of Energy Research. The survey results demonstrated a need for enhanced networking capabilities, and as a consequence, existing networks including MFEnet and HEPnet were combined and extended into what would become ESnet, with the ESnet Steering Committee (ESSC) formed to provide

oversight and guidance. ESnet began providing networking services across the Energy Research community in January 1988, starting with X.25 backbone lines at 56 and 256kbps, and growing to T1 lines (1.5 Mbps) in 1989, and T3 lines (45 Mbps) in 1993. More recent deployments have upgraded the ESnet backbone to OC3, OC12, OC48, and OC192 (10 Gbps).

ESnet planning and operation is carried out with extensive input from the DOE laboratories and the broader DOE science community. Funding is provided by ASCR, with underpinning telecommunications services competitively procured from Sprint (1994-2001) and Qwest (1999-present). Guidance from SC programs and other stakeholders is provided through the ESSC, and network operation is a shared activity within the community facilitated by the ESnet Site Coordinators Committee (ESCC). A retrospective peer review is conducted every 3-4 years, and operational reviews take place annually in August.

The ESnet community has presented a detailed vision of future plans and needs in its Roadmap to 2008 and other planning documents. This vision focuses on developing a seamless, high-performance network infrastructure in which science applications and advanced facilities are n-way connected to terascale computing, petascale storage and advanced visualization capabilities. Future networking requirements have been quantified for a number of science areas, indicating 1000+ Gbps estimated throughput on a 5-10 year time scale. The needs both for high bulk throughput, and for high bandwidth time-critical throughput are expected to increase dramatically.

The existing ESnet architecture is not expected to be able to address increasing reliability requirements, in particular needs for remote operation of experiments, and long term bandwidth needs of the DOE science community. In response, a new ESnet architecture has been proposed with a goal of high speed (20+ Gbps), fully redundant connectivity for every site. The new architecture would consist of a set of Metropolitan Area Networks (MANs), a Science Data Network core, and a high-reliability IP core, together providing redundancy and high-speed site and core connectivity.

The ESnet is a world class facility that has earned a very strong reputation for supplying reliable, high bandwidth networking to the DOE science community. ESnet has a long history of efficiently and cost-effectively meeting or exceeding the needs of a very demanding user community. Its record of rapidly implementing state-of-the-art networking technology to stay ahead of user demand is extremely impressive, as are its reliability and security record. We note in particular the effective response to security threats such as the Sapphire/Slammer worm, and the early implementation of technologies such as ATM and IPv6. The recently conducted study comparing 4 year life-cycle costs of the current centralized ESnet with centralized or decentralized procurement of commercial ISP services, strongly reinforces the cost-effectiveness of the current ESnet organization, and underscores its unique ability to efficiently meet the specific needs of the DOE's large scale science community. The comprehensive suite of services provided by ESnet, including a full suite of network services, comprehensive user support, overlay networks, videoconferencing, grid middleware, and full internet

service, are highly valued by the user community. The close relationships and effective collaboration with national (particularly Abilene) and international peers is commendable.

The COV finds the management of ESnet to be efficient and highly effective. The ESSC and ESCC have successfully enabled strong community involvement in management and operations and strong responsiveness to the needs of DOE scientists. The COV notes that some restructuring of these committees is under consideration. The need for modification was not completely clear to the COV, and we encourage ESnet management to maintain the key elements of the streamlined structure and clear lines of responsibility that have been so successful in the past, when making any needed organizational changes. We note that the division of ESnet and High-Performance Network Research appears to have successfully led to focused management in both areas, and we encourage continued close collaboration between network research (discussed in Sec. 2 of this report) and ESnet to ensure coordination of goals, and efficient implementation of new technologies in ESnet.

In anticipation of the formation of the Network Advisory Committee (NAC) mentioned above, the annual ESnet Operations Review conducted on August 1, 2005, included participation by HEP, BES, and FES. Once formed, the NAC will help the ESnet program manager optimize the governance of ESnet by effectively addressing R&D coordination, technology transfer, operations and planning for networking throughout the Office of Science research enterprise.

The COV would like to highlight and encourage the continued organization of workshops which have successfully brought members of the networking and scientific communities together, and have broadly outlined future bandwidth requirements on 5 and 10 year time scales. A detailed and impressive roadmap has been assembled, which lays out expected future needs, and a comprehensive plan to meet those needs. The plan involves significant leveraging of outside resources (particularly National Lambda Rail), and appears to be cost effective given the substantial challenges posed by the expected very rapid growth in bandwidth needs, along with technological challenges associated with advancing single channel bandwidth beyond 10 Gbps.

The substantial upgrade of ESnet to accommodate the expected increase in the rate of traffic growth from 100%/yr to 300%/yr over the next 5-10 years has been rated 'priority number 7' of 20 in DOE's "Facilities for the Future of Science – A Twenty Year Outlook." The committee is in strong agreement with the high priority placed upon this upgrade, and we note that the success of many of the large experimental projects on the 'top 20' list will depend substantially on the availability of the high bandwidth networking services provided by ESnet. For example, the top priority project, the ITER fusion experiment, will be located internationally, and will require both very high bandwidth communication to multiple SC Office of Fusion Energy Sciences sites for data analysis, and also highly reliable, secure, dedicated networking for real-time analysis and remote operation of experimental systems.

The COV commends ESnet management on its successful identification of key future needs, and development of a clear vision for meeting those needs. We reiterate our strong support for the plans for the ESnet upgrade, and emphasize its importance to the overall mission of the Office of Science. We recognize that in the expected budget environment, it may be difficult to obtain funding necessary for the planned upgrades, and we encourage ESnet and ASCR management to work hard to obtain these funds. However, it is also important to plan for the possibility that these funds may not be forthcoming. We encourage management to consider backups plans, and to prepare for a situation that ESnet has not recently faced, in which user demand strongly and consistently exceeds available bandwidth. We note that such a situation will lead to fragmentation, as users look for other solutions to their networking needs, and lead to inefficient use of resources as individual facilities work to separately procure bandwidth. We encourage management to develop a backup plan that will allow for continued high quality service for high priority research needs while minimizing fragmentation, if this unfortunate situation arises, and to promulgate/ publicize the plan with sufficient lead time for the ESnet user community to absorb the implications of this situation.

ASCR is encouraging a higher level of integration of ESnet with the Internet2/NLR activities as those infrastructures move to a hybrid model for providing production IP services as well as circuit on demand services. This offers the possibility for cost efficiencies as well as assuring good performance between labs and the university community. This is timely since the ESnet backbone contract will be competed during the coming year. Contingency plans for that activity will be finalized by December 30, 2005.

In an environment of constrained resources, the Network Advisory Committee will be a valuable resource for prioritizing activities and for communicating plans.

4. National Energy Research Scientific Computing Center (NERSC)

The ASCR 2005 COV commends the DOE Office of Science's ASCR program for its sustained support of computational science. A major part of this support comes by means of its NERSC capability. NERSC has a long history of providing exemplary service to the computational science community. We are particularly pleased by the broad spectrum of computational science research supported and by the wide range of research organizations supported. We believe that NERSC is truly unique in its level and quality of service provided to the computational science community.

The COV believes that processes used for oversight of NERSC by ASCR are reasonable and appropriate.

The availability of sufficient supercomputer resource allocations is very important to the success of many research projects. Therefore, it is critical that the process used to allocate

resources be fair and also consistent with programmatic needs. We believe that the process used to allocate the NERSC resources is fair and consistent and commendable for its support of projects not funded directly by DOE. We also want to commend DOE for the INCITE program.

A 3 TF system (NCS) was procured for the NERSC program in FY04, for production in FY05, and there are plans to purchase a 6 TF system (NCS-b) in FY05 for production in FY06. We believe that these systems will significantly add to the ability of NERSC to serve the ASCR HPC user community and that they are consistent with the FY04 job mix in which over 50% of the processor hours went to jobs that used less than 512 processors. However, if the trend toward larger jobs (< 26% of processor hours for jobs under 512 processors) reported for the first quarter of FY05 continues or increases, the mix between capacity and capability resources may need to be re-evaluated.

Plans for future NERSC supercomputer acquisitions include the 30+ TF (NERSC-5) procurement in FY06, for a first year of production in FY08. Plans for the NERSC-5 machine are not complete at this time; however, they will need to account for the balance between capacity and capability computing that will be required in FY08. Plans also include a NERSC-6 acquisition for FY09. While the NERSC-6 procurement is not fully defined, it will need to address the balance between capacity and capability computing.

ASCR agrees that there needs to be continued diligence to maintain the appropriate balance between capability and capacity computing at NERSC.

As such, ASCR now emphasizes that NERSC is a center for high-performance capacity and capability computing. Furthermore, the scheduling scheme at NERSC that partitions the machine so that several jobs can execute simultaneously, thereby maximizing the utilization of the processors, allows operators to dynamically balance capability and capacity computing, as requested by its 2,000 users. Plans for future acquisitions will emphasize this feature.

In summary, NERSC is a well-managed program that is providing excellent service to its users. We believe that it is very important for the future of computational science in the U.S. that the DOE Office of Science continue its emphasis of the mission of NERSC to “accelerate the pace of scientific discovery (as) the principal (premier) provider of high-performance computing to SC programs.”

5. Advanced Computing and Research Testbeds (ACRTs)

The Office of Advanced Scientific Computing Research has historically supported an Advanced Computing Research Testbed activity to provide early evaluations of new computing technologies of interest to the DOE SC. Historically, this has been a well-supported and successful activity. It was, for example, central to the selection of the CRAY-X1e in the recent Leadership Class Computing competition. The evaluation of

the CRAY-X1 also provided valuable feedback to the manufacturer with regard to architecture, machine organization, operating system, compilers, and libraries.

As valuable as these activities have been, experience with the CRAY-X1 illustrates the growing complexity of this type of evaluation. It has moved far beyond the original architectural evaluations that took place in universities (e.g.: Caltech hypercube; University of Illinois/ Urbana Cedar Project). For example, it is most likely that future evaluations will need comprehensively to explore changes to current programming paradigms that may be required to exploit new architectures. In a climate of flat or declining budgets, the level and nature of future support for such programs should be carefully weighed against other programmatic priorities. One option would be to explore mechanisms for conducting these evaluations using more of a distributed community approach. The COV encourage exploration of thoughtful strategies for providing insight into evolving high-performance architectures with the goal of heavily leveraging the direct investments made by the Office of Science.

ASCR accepts this advice. Such a strategy is now in development for implementation in FY 2006.

One additional observation is related to the life cycle of high-performance computing systems that are acquired for evaluation purposes. The IBM Power-4 system acquired by the ACRT program is an immediate example. This system remains a significant computational resource, but the operation of such a system can place a non-trivial demand on the ASCR budget. We encourage ASCR to explore and adopt formal mechanisms to transition experimental hardware to production status, which may require the identification of alternate funding sources for the operation of the computational resource.

ASCR accepts this advice. While there appear to be substantial bureaucratic obstacles to overcome, mechanisms are being explored with the hope that a transition strategy can be developed and implemented soon.

6. Leadership Class Computing (LCC)

The establishment of the Leadership Class Computing facility, based on a merit review competition, is a very positive accomplishment as well as a very appropriate activity for the DOE Office of Science. Many of the SC computational resources over the last decade have been deployed to satisfy growing demand for high performance computing cycles, whereas the LCC provides the opportunity for capability-limited computational science to move forward. The strategy of providing a home for the most demanding of scientific activities is applauded by the COV. We also commend the partnerships with industry, which will enhance the likelihood that the LCC will be highly successful. We expect that the existence of the LCC will accelerate the advancement of scientific

understanding in many mission-critical areas of science, and will inspire a broadened scientific scope at the high end of computational science within the Office of Science. This will leverage many long-term SC investments, and thereby will open the door to wide areas of scientific progress. Along similar lines, we are also pleased to hear that ASCR has plans to allocate a percentage of LCC machine resources to research scientists pursuing world-class capability-limited science.

It is very important that the Office of Science does everything appropriately possible to help the LCC live up to expectations as driven by the science application needs. We recognize that the funding for this capability is only a fraction of what was first anticipated, but encourage ASCR to continue to be nimble in responding to future funding opportunities that will allow further augmentations of the LCC initiative.

ASCR is committed to establishing the LCC as a capability high end computational resource driven by the needs of the scientific applications. Further, ASCR recognizes that the current funding level for the LCF is not optimal. ASCR will continue efforts to identify the computational requirements and technical barriers of DOE's mission-related scientific applications so that the office has the information needed to successfully respond to future funding opportunities,

Deleted: current

Deleted: for mission-critical

7. Need for Capacity, Capability Planning

The establishment of a Leadership Class Computing capability brings new challenges for ASCR. There are two worth noting, which are tightly coupled to each other. The first is the need to establish the metrics by which ASCR's portfolio will be judged, both internal to the Office of Science and by the external computational science community. The second challenge is to allocate these resources optimally, in order to realize the best return on the large NERSC and LCC investments. This may have implications for the allocation mechanisms for each of these programs.

ASCR agrees that it is necessary to review the mechanisms for allocating our high end computing resources and develop a set of allocation guidelines. Beginning with computer allocations for FY2006, a portal will be available for use by the other SC program offices to manage their resource allocations on the Office of Sciences computing resources at NERSC, LCC, ANL and PNNL. The functionality of the portal will be increased to accept computer allocation requests at NERSC, LCC, and EMSL's MSCF.

Formatted: Font: Bold, Font color: Blue

Deleted: accounts

Formatted: Font: Bold, Font color: Blue

Deleted: C

Formatted: Font: Bold, Font color: Blue

Formatted: Font: Bold

ASCR needs to be proactive with regard to the development of relevant metrics that emphasize the delivery of new scientific capabilities and understanding, as compared with the more common (and irrelevant) peak computational performance measures. This should be a high priority for the Office, so that it is not left to others in the high performance computing community to quantify the value of ASCR's unique computational investments. It is important that these metrics differentiate between

capability-limited science and capacity-limited science so that the spectrum of ASCR NERSC and LCC computational resources are appropriately allocated to the user community.

ASCR agrees that it important to establish relevant metrics to quantify the value of ASCR's unique computational investments. **In the first quarter of FY06, ASCR will solicit information from other agencies with high end computational resources such as NSF, DOD, to determine the applicability of their metrics to the ASCR resources. In addition, ASCR will conduct a series of performance evaluations workshops to identify relevant metrics for both capability and capacity science. Once ASCR has established a set of metrics that will provide performance measurements along the office's computational resource continuum, we solicit feedback on these metrics from the scientific community prior to their implementation.**

Deleted: al

The COV notes that this discussion comes full circle to the central comment offered in the Executive Summary at the beginning of this report: well-thought ASCR allocation metrics that span the continuum from capability through capacity simulation will provide a framework for portfolio planning and prioritization, for network requirements and developments, for new computational platform procurements, and, overall, for the future of high performance scientific computing in the United States.

A1. List of Acronyms

ACRT	Advanced Computing Research Testbed
ASCAC	Advanced Scientific Computing Advisory Committee
ASCR	Office of Advanced Scientific Computing Research
ATM	Asynchronous Transfer Mode
COV	Committee of Visitors
DOE	Department of Energy
ESSC	ESnet Site Coordinators Committee: coordinates among the individual institutions that benefit from the use of ESnet's resources. The ESSC includes representatives from each of the major ESnet 'backbone' sites and other DOE-funded activities. ESSC is appointed by ASCR from within the ESSC; the ESSC charters the ESSC.
ESnet	Energy Sciences network
ESSC	ESnet Steering Committee: was originally formed in 1986 with representation from the Energy Research science program offices. ESnet services now extend to the Offices of Energy Efficiency, Nuclear Energy, The Energy Information Agency, and Human Resources, and the ESSC membership has expanded to ensure adequate breadth of program representation.
FY	Fiscal Year
Gbps	Gigabits per second
HEPnet	High Energy Physics network
HPC	High Performance Computing
IEEE	Institute of Electrical and Electronics Engineers, Inc
INCITE	Innovative and Novel Computational Impact on Theory and Experiment program
IP, IPv6	Internet Protocol, Internet Protocol version 6
ISP	Internet Service Provider
LCC	Leadership Class Computing
kbps	kilobits per second
MAN	Metropolitan Area Network
MFEnet	Magnetic Fusion Energy network
NASA	National Aeronautics and Space Administration
NERSC	National Energy Research Scientific Computing Center
NIH	National Institutes of Health
NIST	National Institute of Science and Technology
NOAA	National Oceanic and Atmospheric Administration
NSA	National Security Agency
NSF	National Science Foundation
OC [1, 12, 48, 192]	Optical Carrier levels; OC-48 is 2.488 Gbps
ROI	Return on Investment
SC	DOE Office of Science
TF	teraflop, or trillion floating point operations per second
T [1,3]	dedicated telephone transmission path connection

A2. Committee of Visitors (COV) Resumes

Jill Dahlburg, Naval Research Laboratory (chair)
James Hack, National Center for Atmospheric Research
Kevin Mills, National Institute of Science and Technology
Philip Snyder, General Atomics
James Tomkins, Sandia National Laboratory

Jill P. Dahlburg, PhD

Naval Research Laboratory, CODE 1001.2, Washington, DC 20375
(202) 767-1312 [office]; Jill.Dahlburg@NRL.NAVY.MIL

Jill Potkalitsky Dahlburg is the Senior Scientist for Science Applications at the Naval Research Laboratory (NRL) in Washington, DC, reporting to the NRL Director of Research, Dr. John Montgomery. Prior to her appointment to this position, she was director of the Division of Inertial Fusion Technology and co-director of the Theory and Computing Center in the Energy Group at General Atomics (GA) in San Diego, from February 2001 to June 2003.

Dr. Dahlburg, who received her PhD in theoretical plasma physics from the College of William and Mary in May, 1985, is recognized as a technical authority on: high performance computing algorithms and techniques; fluid and gas dynamics; and, experimental data analysis, interpretation, and integration. Dr. Dahlburg joined the civil service research staff at NRL in June, 1985, as a computational physicist. As a member of the NRL Nike KrF Laser Program from its inception, and Head of the Laser Plasma Hydrodynamics Section for that Program, she contributed to laser matter interaction research, with emphasis on the understanding of the Raleigh-Taylor instability, implosion and coronal hydrodynamics, and laser beam imprinting. In particular, she spearheaded the development of RAD3D, the first three-dimensional multi-group radiation transport hydro-code appropriate for laser-plasma modeling. RAD3D has remained a premier simulation code in that field for more than a decade.

In 2000, Dr. Dahlburg joined the Tactical Electronic Warfare Division (TEWD) of NRL as Head of the Distributed Sensor Technology Office. From TEWD she co-proposed and then served as Co-Principal Investigator on the Office of Naval Research / Marine Corps Warfighting Lab - sponsored SECNAV Small UAV initiative, Dragon Eye, for its first year of development. Dragon Eye numbered among the Popular Science Top 100 Best of What's New 2001, and is currently seeing active duty in Iraq.

During her 2001-2003 tenure at General Atomics, Dr. Dahlburg's IFT Division grew to employ more than seventy full time technical and support personnel and ten students, including approximately forty GA permanent employees. IFT Division activities included: target and component fabrication for inertial confinement fusion research; mass production techniques development for inertial fusion energy; cryogenics and target injection systems development and deployment; high intensity laser matter interaction research; efficient production of hydrogen; and, sources of iodine for high power gas lasers.

Dr Dahlburg's technical collaborations have included scientists in both the national and international physics and engineering communities. She has served on numerous review and other committees for the Department of Defense (DoD), the Department of Energy (DOE), the National Research Council, the National Science Foundation, the American Physical Society (APS), and a number of DOE National Laboratories, with membership in the DOE Office of Science FESAC and ASCAC federal advisory committees, the Lawrence Livermore National Laboratory Defense and Nuclear Technologies Director's Review Committee, and as federal advisor to the 2004 Defense Science Board Task Force on the National Ignition Facility. Jill is 2005 Chair of the APS Division of Plasma Physics (DPP), and is a member of the APS Nominating Committee. She served as Divisional Associate Editor (Plasma Physics) for The Physical Review Letters (1996-2000), and is presently an editor of Fusion Engineering & Design. Her professional honors include six NRL Alan Berman Research Publication Awards for scientific publication excellence, and being named APS Centennial Speaker (1998-99) and APS/DPP Distinguished Lecturer (1999-00). She is a Fellow of the American Physical Society.

James J. Hack, Ph.D.
National Center for Atmospheric Research
1850 Table Mesa Drive
Boulder, Colorado 80305
303-497-1387; jhack@ncar.ucar.edu

James J. Hack is a Senior Scientist at the National Center for Atmospheric Research (NCAR) in Boulder, Colorado, Deputy Director of the Climate and Global Dynamics Division, and Head of the NCAR Climate Modeling Section. He also holds an Adjoint Professor position at the University of Colorado, Boulder. He received his Ph.D. in Atmospheric Dynamics from Colorado State University in 1980, after which he became a Research Staff Member at the IBM Thomas J. Watson Research Center, in Yorktown Heights, NY. While at IBM Research his work included detailed performance analysis and modeling of computationally-intensive scientific applications, emphasizing vector multiprocessor architectures. He was also engaged in research on the design and simulation of memory sub-systems for advanced, high-performance, scalar machine organizations targeted for grand challenge scientific applications.

Hack joined NCAR in 1984, where he has been a lead developer of the NCAR global atmospheric model, currently called the Community Atmosphere Model (CAM). His primary scientific activities have been in the areas of tropical dynamics, the parameterization of moist convection, cloud processes and their modulation of radiative transfer, and on diagnostic methodologies for evaluating simulation quality. He has worked on all aspects of large-scale global modeling, including the development and evaluation of numerical methods, the development of analysis frameworks, and the implementation of global models on high-performance computer systems.

Hack is a member of the U. S. Water Cycle Scientific Steering Committee, a member of the U.S. GEWEX Cloud Systems Study Scientific Steering Group, and Co-Chair of the NCAR Clouds and Climate Program. He has also served as a member of numerous NSF and Department of Energy national panels on high-performance scientific computing, as well as served in the capacity of Editor for the Journal of Climate. He is currently a member of the DOE Office of Science Advanced Scientific Computing Advisory Committee, the Oak Ridge National Laboratory Computer Science and Mathematics Division Advisory Committee, the DOE Computational Science Graduate Fellowship Steering Committee, and Chairs the LLNL PCMDI Advisory Committee.

Kevin Mills, PhD
Senior Research Scientist
National Institute of Standards and Technology

Dr. Mills holds the position of senior research scientist within the networking-research program of the National Institute of Standards and Technology (NIST). Currently, Dr. Mills conducts research regarding methods to identify and elicit emergent behaviors in large-scale distributed systems and global networks. From 2001 to 2005, Dr. Mills managed a research program focusing on networking for pervasive computing. From 1999 to 2001, Dr. Mills conducted research into a means to represent processing-time requirements among heterogeneous computer platforms distributed throughout a network. From 1996 to 1999, Dr. Mills served as program manager for the intelligent collaboration and visualization program conducted within the Defense Advanced Research Projects Agency (DARPA). Prior to joining DARPA, Dr. Mills conducted PhD research in software engineering at the George Mason University (GMU). Before matriculating at GMU, Dr. Mills managed the NIST program in open-systems interconnection (OSI) networking, where he focused on techniques to measure protocol efficiency and to improve performance of transport-layer protocols. Prior to joining NIST in 1982, Dr. Mills worked in commercial companies to develop performance-measurement products (at Tesdata Systems Corporation) and to apply formal methods to design network protocols (at the System Development Corporation). From 1973 to 1977, Dr. Mills was an automated data systems officer with the United States Marine Corps, where he helped to develop automated systems for air defense and air-traffic control. Dr. Mills has authored approximately 50 publications in the fields of networking, software engineering, distributed systems, and collaboration technology.

Concurrently with his government service, since 1996, Dr. Mills has served on the adjunct faculty within the GMU School for Information Technology and Engineering, where he teaches courses in computer networks and software engineering. In addition, Dr. Mills advises research programs within several other government agencies, including the National Science Foundation, DARPA, the Department of Homeland Security, and the Department of Energy.

Contact information and additional details regarding the qualifications of Dr. Mills may be obtained from the following URL: <http://www.antd.nist.gov/~mills/index.html>

Dr. Philip B. Snyder

Principal Scientist (Scientist VI), General Atomics, Theory and Computational Science

B.S., Computational Physics, Yale University, 1993

Ph.D., Plasma Physics, Princeton University, 1999

Dr. Snyder is a recognized expert on computational physics and the plasma physics involved in core turbulence, edge localized modes (ELMs) and the H-mode pedestal. He specializes in employing both large scale nonlinear simulation and analytic theory to gain physical insight into complex plasma behavior. He conducted his doctoral research at Princeton University, working with G.W. Hammett on electromagnetic turbulence and transport in tokamaks. He developed a new theoretical model to efficiently describe the evolution of electromagnetic drift waves and Alfvén waves, and implemented the model in realistic nonlinear simulations using a massively parallel gyrofluid code. These simulations predicted fluctuation spectra and turbulent transport which match experimental trends, and revealed significant finite beta effects on turbulent transport. Dr. Snyder joined General Atomics in September 1999, and is currently a Principal Scientist in the Theory and Computational Science Division of the Fusion Group. At present he is engaged in studies of the stability and nonlinear dynamics of the edge plasma, aimed at developing an understanding of ELMs and constraints on the H-mode pedestal. In collaboration with H.R. Wilson, he has made pioneering contributions to the peeling-ballooning model of ELMs, and developed the ELITE code, which has facilitated the quantification and experimental confirmation of the model. He is also conducting (with X.Xu) studies of the nonlinear dynamics of edge instabilities using the BOUT code, which have identified poloidal narrowing, explosive growth, and fast radial propagation of filaments during the early stages of the ELM crash.

Dr. Snyder has presented invited talks at several conferences including APS/DPP 2004, Sherwood 2004, IAEA Fusion Energy Conference 2002, APS/DPP 2001, Festival de Theorie Workshop 2003, Plasma Edge Theory Workshop 2001, and Sherwood 1999. In addition Dr. Snyder has presented seminars on ELM and pedestal physics at the DOE Office Seminar (1/04, 2/02), U. of Wisconsin (1/05, 10/03), Culham Science Centre (7/03), UCSD (1/03), UCLA (2/02), MIT (10/01), and LLNL (12/04, 1/01). In 2004, Dr. Snyder received the Rosenbluth Award for Fusion Theory in recognition of his contributions to ELM and pedestal physics.

Honors and Awards

Rosenbluth Award for Fusion Theory (2004)

DOE Fusion Energy Postdoctoral Fellowship (declined)

National Science Foundation Graduate Fellowship

Princeton University Merit Prize

Henry Prentiss Becton Prize for Exceptional Achievement, Yale University

Phi Beta Kappa

Tau Beta Pi

B.S. degree earned *summa cum laude*, with distinction in the major

Professional Service

Edge Coordinating Committee (7/04-present)

Sherwood Executive Committee Chair (11/04-present), Vice-Chair (11/03-11/04)

Sherwood Executive Committee (11/02-present)

DIII-D Research Council (8/01-present)

DIII-D Pedestal and ELM Research Thrust, Deputy Leader (11/02-present)

ITPA Pedestal Working Group

PI of NERSC project on Simulation of Magnetically Confined Fusion Plasmas (mp94)

James L. Tomkins

Distinguished Member of Technical Staff, Org 9220, MS 1109
Sandia National Laboratories

Professional Experience: I have over thirty years of experience in high performance computing for science and engineering. I spent the first 15 years of my professional career doing application code and model development and analysis for nuclear reactor safety. During this period I was the code architect and project leader for a major reactor safety code development project at Sandia (MELPROG). In 1989 I transferred from reactor safety work to advanced computing with the goal of demonstrating the applicability of massively parallel computing to a broad set of science and engineering applications. In my first assignment I lead a successful project to demonstrate the feasibility of using massively parallel computing to track thousands of objects in space in real-time for the Strategic Defense Initiative (SDI). In the mid '90s I became the project leader for Sandia's ASCI Red computer system project. I wrote the specifications for ASCI Red. ASCI Red has been a very successful computer system and was the world's first TeraFlop supercomputer. Currently I am the chief architect and project leader for the Red Storm supercomputer project. I have a unique combination of experience in application code development and supercomputer architecture design.

Sandia National Laboratories, Aug 1987 to Present:

- Red Storm supercomputer system chief architect and project leader. Sandia ASCI platforms element lead.
- Project leader for the ASCI Red supercomputer system, the world's first general purpose TeraFlops computer system.
- Leader of a Strategic Defense Initiative project that demonstrated the feasibility of tracking thousands of objects in space in real-time using massively parallel computing.
- Code architect and leader of MELPROG Light Water Reactor simulation code development project. MELPROG was designed to simulate the complex physics involved in severe accidents. This included meltdown of the reactor core such as happened in the Three-Mile Island accident.

Formatted: Bullets and Numbering

Science Applications International Corporation (SAIC), Aug 1982 to Aug 1987. (Full time contract at Sandia):

MELPROG code development and physics model development. I designed the overall code architecture and developed adaptive mesh fuel rod models for heat transfer, melting, oxidation, and relocation. I also developed a radiation heat transfer model that adapted to changes in geometry caused by material relocation during the simulation.

Los Alamos National Laboratory, September 1976 to August 1982:

I worked on the fast reactor safety analysis code SIMMER doing experiment analysis and model development. This work included putting a radiation heat transfer model in SIMMER. I developed a fuel rod failure model for light water reactor fuel rods based on the LAFM fast reactor fuel rod model.

General Atomic Company, June 1974 to September 1976:

At General Atomic I worked on both fast and thermal reactor safety analysis and on development of models for fast reactor safety for General Atomic gas cooled reactor designs.

Education: I received an M.S. degree in Nuclear Engineering from the University of Illinois in Urbana, Illinois in May of 1974, and a B.A. degree in Physics from the University of Northern Iowa in Cedar Falls, Iowa in January of 1972.

A3. Committee of Visitors (COV) Instructions Letter

Date: Fri, 18 Mar 2005 13:41:48 -0500
From: "Polansky, Walt" <Walt.Polansky@science.DOE.gov>
Subject: COV Review of ASCR Facilities & Network research- April 4-5:
Administrative and Logistical Information
To: Jill Dahlburg <jill.dahlburg@nrl.navy.mil>,
Phil Snyder <snyder@fusion.gat.com>, Kevin Mills <kmills@nist.gov>,
Jim Tomkins <jltomki@sandia.gov>, Jim Hack <jhack@ncar.ucar.edu>
Cc: "Strayer, Michael" <Michael.Strayer@science.DOE.gov>,
"Baker, Melea" <Melea.Baker@science.DOE.gov>,
"Hiegel, Jane" <Jane.Hiegel@science.DOE.gov>,
"Hitchcock, Daniel" <Daniel.Hitchcock@science.DOE.gov>,
"Kreisman, Norman" <Norman.Kreisman@science.DOE.gov>,
"Oliver, Ed" <Ed.Oliver@science.DOE.gov>,
"Polansky, Walt" <walt.polansky@science.DOE.gov>

Members of the Committee of Visitors -

Thank you for making the commitment to serve as a member of the Committee of Reviewers (COV roster contact information-final-031105.doc) for the April, 2005 review of the facilities and network research activities supported by the Advanced Scientific Computing Research program.

Your chair, Jill Dahlburg, has asked me to communicate with you on administrative and logistical matters associated with this COV review, which consists of NERSC, Leadership Class Computing, the Advanced Computing Research Testbeds, ESnet, network testbeds and network research. The COV review will be held at the DOE offices in Germantown, MD on April 4-5, 2005. A draft agenda is attached (COV-agenda-031705.doc).

In the context of its overall charge (cov-charge.pdf), and with an emphasis on the FY2002-FY2004 time-frame, this COV is asked to:

(1) assess the efficacy and quality of the processes used by ASCR for its activities in high-performance computing facilities, high-performance networking facilities, network testbeds and network research:

- (a) to identify promising research directions,
- (b) to solicit proposals (from laboratories and from universities),
- (c) to review, recommend, and document actions on proposals and
- (d) to monitor active projects and the vitality of the overall portfolio.

(2) comment, in recognition of the DOE mission and available funding, on how current ASCR processes have affected:

- (a) the breadth and depth of the elements within a portfolio, and
- (b) the national and international standing of the portfolio.

(3) comment on future directions proposed by ASCR management and on opportunities that might not have been presented,

(4) comment on how the process for these reviews might be improved.

Attached are conflict of interest statements. Please review the appropriate statement, provide the requested information and fax the completed form to Melea Baker at 301-903-4846, prior to the review. If you have any questions about these statements, please contact me at your earliest convenience.

Your travel and lodging arrangements will be handled by personnel affiliated with the Oak Ridge Institute for Science and Engineering (ORISE). ORISE staff will be working directly with you to arrange flights, rental cars, lodging and reimbursements for local travel.

Lunch will be provided on both days. On April 4th, a COV-only dinner is being planned at a local restaurant.

To help ensure a timely, and successful, transition through our security procedures, please do not bring either your personal, or an officially-assigned, laptop computer into the building. A laptop will be provided to you for your use during the review. Network access will not be available to you within the building. If you desire, we can provide storage media for you to download your notes, analyses, etc., should you decide to work further that evening at your hotel on your laptop.

Cell phones and Blackberry devices are permitted within the building, however.

We are looking forward to your visit on April 4-5, 2005 and to helping you successfully complete your review.

Best regards,
Walt Polansky
(301-903-5935)
COV Liaison
Office of Advanced Scientific
Computing Research