

11 March, 2002

C. Edward Oliver
Associate Director of Science
for the Office of Advanced Scientific Computing Research
U.S. Dept. of Energy, SC-30
19901 Germantown Road
Germantown, MD 20874

Dear Dr. Oliver,

In April, 2001, James Decker (then Acting Director, United States Department of Energy (DOE) Office of Science), presented this charge to the Advanced Scientific Computing Advisory Committee (ASCAC), advisory panel to the Office of Advanced Scientific Computing Research (ASCR):

For ASCR facilities such as NERSC, ESnet, Chiba City at ANL, and the CCS at ORNL

- (a) What is the overall quality of these facilities relative to the best-in-class in the US and internationally?
- (b) How do these facilities relate and contribute to Departmental mission needs?
- (c) How might the roles of these facilities evolve to serve the missions of the Office of Science over the next three to five years?

To address this charge, the ASCAC formed a Subcommittee on Facilities, with members: John Connolly (U.KY); James Coronos (Krell Institute); Helene Kulsrud (IDA); Greg McRae (MIT), Paul Messina (Caltech); Warren Washington (NCAR); Stephen Wolff (Cisco); and myself (Jill Dahlburg (GA), Subcommittee Chair).

The essential finding of the Subcommittee is that each of the four diverse and complementary ASCR facilities is among the best worldwide in its respective category. It is the opinion of the Subcommittee that these ASCR facilities and the related spin-off research efforts contribute outstandingly to the mission needs of the DOE, and profoundly and positively impact high performance computing activities worldwide.

Looking ahead, the Subcommittee offers five recommendations for the ASCR future:

1. ASCR should retain focused commitment to high end computing in the service of DOE Office of Science missions.
2. ASCR should build on its present plan to develop a strategic plan for the next generation high end 21st Century multi-user mission-driven computing environment.
3. ASCR should develop an integrated allocation strategy for its computational resources. This allocations process should seek to ensure that each machine is filled to the greatest extent practicable with high priority DOE Office of Science jobs that are not feasibly run on smaller machines.
4. DOE Office of Science researchers should be encouraged by ASCR to procure both mid-range and lower end machines/clusters with individual program funds.
5. ASCR should embrace a cohesive networking and resource allocation approach for computing infrastructure as a way to provide the most uniform interface to all the types

of computing facilities that are encompassed by ASCR. To this end, ASCR should continue to incorporate advances in networking and resources integration as they develop, and should, further, encourage enhanced research efforts in new architectures and networking capabilities.

I. Background

During the period April through December 2001 the Subcommittee gathered information for the above-noted findings and recommendations, making primary use of two sources: presentations to the ASCAC Subcommittee on Facilities and associated institutional information, and query response.

The first, facilities presentations to the ASCAC and/ or to the Subcommittee, included briefings: (a) on May 2-3, 2001, by: Walter Polansky, *DOE Office of Science* [overview of the MICS Facilities]; Horst Simon, *National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory (LBNL)* [overview of NERSC]; Richard Stevens, *Argonne National Laboratory (ANL)* [overview of Chiba City]; Thomas Zacharia, *Oak Ridge National Laboratory (ORNL)* [overview of ORNL CCS]; and, James Leighton, LBNL [overview of the ESnet]; (b) on August 16, 2001, by: David Schissel, *General Atomics (GA)* [overview of the National Fusion Collaboratory]; and, Paul Messina, *Caltech*, [overview of the Grid, internationally]; and, (c) on October 25-26, 2001, by: Richard Stevens, *ANL* [on high-performance computing (HPC) facilities: Grids, petaflops, and usage]; Robert Ryne, *NERSC* [on requirements of a HPC system user]; Dalton Schnack (talk given by Jill Dahlburg, *GA*), *Science Applications International Corporation (SAIC)*, [on a mission-driven perspective on the status and needs for DOE computing]; and, Stephen Wolff, *Cisco* [on networks, with particular emphasis on ESnet]. The second source of information was obtained as ASCR facilities' answers to posed questions.

With regards the first aspect of the Decker charge, facilities evaluation, the Subcommittee considered the facilities in isolation and comparatively. Findings from this study are summarized in Sec. II.

The second aspect of the charge, relation to and contribution of the facilities to DOE mission needs, is addressed in Sec. III. Several of the speakers to the Subcommittee described the fundamental role of advanced computing in solving DOE mission-relevant problems (*e.g.*, Schnack), the importance to their research of state-of-the-art ASCR HPC resources, and the excellent responsiveness of the ASCR facilities. The more subtle subtext of mission driven research as a 'requirements pull' for enabling fundamental advances was also indicated directly or indirectly by several of the speakers (*e.g.*, Stevens noting that much grid software has been developed for the solution of specific DOE problems). The essential observation is that mission driven research, which inspires some of the best basic research extant, is exemplified superbly by the DOE ASCR.

Twenty years ago, high end scientific computing was performed on manufactured-on-request vector mainframes for which 30 Mflops sustained was considered to be good

performance. Vector FORTRAN was the standard programming language for which few debugging and optimizing tools were available, and most jobs were submitted by remote batch processing using dumb terminals and 9600 baud telephone connections. In contrast, high performance computing of 2001 was typified by commodity massively parallel platforms on which 30 Gflops of sustained performance was easily possible using group-developed object orientated FORTRAN/C software that was coded with the assistance of automated parallel debugger and development tools on versatile desktop workstations. The higher speed connectivities such as ATM OC-12 enabled mainframe-driven visualization systems for tasks ranging from debugging to large database results processing. With this advent of the ability to routinely perform highly resolved, multi-dimensional, engineering-class computations, advanced scientific computing has become an enabling tool for first principles exploration. Computational science now is considered by most researchers within the DOE Office of Science, and in the physics and engineering communities at large, to be a third arm of research, ranking as equal with theory and experiment as a tool for discovery. Continuity of these advances will require a careful planning of resources: of facilities capabilities, of allocations, and of connectivity. Strategic directions for the next three to five years, part (c) of the charge, are topics of Sec. IV.

II. Facilities: *What is the overall quality of these facilities relative to the best-in-class in the US and internationally?*

The charge to this Subcommittee includes as first element the evaluation of four facilities funded by the OASCR. These are NERSC, ESnet and the Chiba City and CCS facilities at ANL and ORNL respectively. Each of these four facilities has a very different vision, and all serve complementary purposes within the DOE Office of Science.

NERSC is a computational production center providing state of the art access to computational resources to applications researchers. These researchers are interested in a stable productive environment where software can be optimized without the added complication of a constantly changing environment. The role of NERSC is to concentrate on the current generation of supercomputer, and provide service to a wide variety of computational scientists. Whether or not the facility should venture beyond the current generation, and overlap its role with the other facilities is a matter of ASCR policy. This Center, situated at LBNL, is an excellent, classic production center, with a long history of providing first class service to large scale applications users by means of state of the art commercial machines. With the shift to massively parallel platforms (MPP), NERSC consulting requirements became more involved and demanding. NERSC responded successfully, by using the personnel resources to provide in-depth collaboration, algorithm and visualization support. The current, well-used, high end machine at NERSC is ranked Number 3 in the world, in the November 2001 version of the Top 500 list, with the DOE classified machine ASCI White ranked top and the Pittsburgh Supercomputer Center Compaq machine ranked second. NERSC has made a proposal to continue to move up the terascale ladder, which has drawn good reviews from the user base. It proposes to increase the facility's capacity by an order of magnitude over the next

seven years. The NERSC allocation process gives priority to Office of Science “grand challenge” problems, but is also available to a large number of DOE grantees from all branches of the Office of Science. In FY2000 NERSC delivered 7,846,244 machine hours to production projects.

ESnet is a facility that provides networking and no cycles. It serves a similar purpose within the Office of Science as does NERSC. At first blush, it is a small, reliable ISP. It is governed by and is extraordinarily responsive to its users. It has relieved the user sites of the need to provide networking expertise and user services. In terms of performance, ESnet’s connectivity is at present adequate for most current users. It is a lean, cost-effective operation, with good management tools and user services, but no central capability for networking research. In February 2001, ESnet carried a tremendous 45 terabytes of data, and continues to experience a 100% growth each year. The danger is that the net is required to meet the increasing demand even as the governance structure, which primarily supports current connectivity stability, is not well constituted to deal with the explosive growth. In order to meet the growing demand, ESnet needs to take on research tasks in strategic directions. New applications could add 1 Gbytes/sec or 300 Tbytes/month.

The ANL Chiba City Project, is a testbed for generating computer science tools. The goal of the project is to provide a series of parallel hardware and software testbeds for the computer science and applications community aimed at supporting research in software scalability. The work supported by the initial testbed has had significant effect on the high-performance computing community, and as the project evolves and expands, researchers at ANL believe that it will have increasingly broad impact by enabling more rapid progress in realizing the dream of scaleable systems. This impact will only be achieved if the community can exploit the capabilities of the testbed on a routine basis, an objective which is embraced by the project personnel.

The Oak Ridge Center for Computational Sciences (CCS) is a high-end capability computing center focusing on a few key timely topics in support of the Department’s science mission. In support of this mission, the CCS is currently focused on providing specialized services to the biology, climate and materials sciences communities. In addition, the CCS is the principal resource for SciDAC projects. Another important role is to evaluate new architectures through specific applications benchmarks. This Oak Ridge system was set up on the model of a few groups of highly sophisticated users running very large and/or long-running jobs on large computer systems to “push the envelope” in computational science by performing tuned calculations that could not otherwise be easily carried out. This model has allowed the CCS to often take delivery of emerging, and unproven architectures, such as the Intel Paragon (or more recently the IBM Power4) to drive computational sciences at the leading edge. Today, CCS continues to cater primarily to a small number of applications groups that use the systems. These groups have both very sophisticated users and technician level users. The sophisticated users develop the codes to a state where they can be turned over to the technicians to make many runs of the code to study a parameter space. The computers and the problems are so complex that it is nearly impossible for one person to understand both

the science and the computers at a “world class” level; consequently, teams of users are required to effectively utilize the systems. It is anticipated that this trend will increase in the next five years. With movement to clusters of larger and larger SMP nodes, and, in particular to next generation "cellular" petascale machines, the applications will be forced to adopt more and more levels of parallelism to take advantage of the resources. The CCS is encouraging teams that, as they grow, include more specialists.

One way to characterize (and perhaps oversimplify) the three computational facilities is that NERSC is a computational production center providing access to multi-teraflop state-of-the-art computational resources to a broad set of applications researchers; the Center for Computational Sciences at ORNL is a topical computational sciences facility providing multi-teraflop, focused resources on advanced architectures for a few key science topics; and Argonne is a teraflop-range computer sciences facility that is focused on enabling rapid progress in realizing the dream of scaleable systems. Each is a first class facility of its type and purpose when considering both US facilities such as those within the National Partnership for Advanced Computational Infrastructure and systems in the DOE Advanced Scientific Computing Initiative complex, and also international capabilities.

III. Mission Relevance: *How do these facilities relate and contribute to Departmental mission needs?*

In this section is examined the relationship between mission needs and facilities of the Department of Energy’s Office of Science. Large-scale computational modeling and simulation have become central to most scientific and engineering research and the missions of the DOE Office of Science are no exception. A number of mission statements to this effect have been issued by the Secretary, by the Office of Science, and Advanced Scientific Computing Research Program. A brief review of some of the most relevant mission statements is given below:

1) In the FY02 request to Congress the ASCR program articulated the following mission statement: “The research and facilities supported by ASCR are critical to the success of all the missions of the Office of Science because computational modeling and simulation have become an important contributor to progress in all SC scientific research programs. Modeling and simulation is particularly important for the solution of research problems that are insoluble by traditional theoretical and experimental approaches, hazardous to study in the laboratory, or time-consuming or expensive to solve by traditional means. All of the research programs in the U. S. Department of Energy’s Office of Science—in Basic Energy Sciences, Biological and Environmental Sciences, High Energy and Nuclear Physics—have identified major scientific challenges that can best be addressed through advances in scientific computing.”

2) At the joint DOE and National Science Foundation (NSF) National Workshop on Advanced Scientific Computation of July 30-31, 1998, there were a set of recommendations developed a long-term strategy for the two agencies. The report that

contains the recommendations is often referred to as the Langer report, which is named after the chair, James S. Langer. “ The impact of Advanced Scientific Computing on industry, government, and national labs has been growing for several decades. However, in the future, we believe that will be a very rapid expansion of such techniques across a far broader segments as Advanced Scientific Computing will become an indispensable tool in understanding and managing our ever more complex and inter-related world. In industry, it will move beyond crash simulations, airplane design, and drug design to a whole new world of data intensive computing such as financial risk management, fraud detection, and supply chain optimization. In government, computing in the service of national defense will be extended to decision support for such societal issues as disaster planning and management, infrastructure investments in protection, and environmental and energy security. As one of the largest producers of data and reports, the digital age will employ scalable computers to help organize and deliver more cogent information products to our citizens. National labs will extend their missions from use of high-end computing in the service of national defense to national decision support for policy issues involving the environment on the energy economy.”

The above broad statements allude to the future for scientific computing with the Office of Science. Through the materials that have been presented to this Subcommittee and the testimony of many experts, we see a future of increasing reliance of scientific computation to accomplish the missions. The Associate Directors of each of the major components of the Office of Science strongly concurred with the need for a range of computing resources for their respective science missions.

As important as providing access to large-scale computing environments (powerful computers, large data archives, advanced visualization technologies) are efforts aimed at advancing the state of the art of computing technologies and making it easier to use them effectively and efficiently. The ASCR program and its predecessors in the Office of Science have a long and distinguished history of developing mathematical models, algorithms, software libraries, and software tools for high-end computing on advanced architectures.

The current manifestation of such efforts is found in the recently developed Scientific Discovery Through Advanced Computing (SciDAC) program, which has as its goal to produce the scientific computing, networking, and collaboration tools for DOE science. The program goals of SciDAC are aimed at addressing the computation needs more effectively across the major Office of Science programs and in collaboration with other government programs. DOE Office of Science has also contributed to the creation of new computer architectures that have led to major advances. These activities are also highly relevant to the Office of Science mission, because without them the available computer systems would be less capable and more difficult to use, and the Subcommittee urges that they be continued.

More recently, Secretary Abraham said in a speech on homeland defense, "Our world class scientific and engineering facilities and creative researchers helped make our nation more secure for over 50 years. These same resources have been trained on the threats

posed by terrorism for some time and because of this foresight, technologies such as those are in deployment today." Also, the Secretary has stated that "program like the Human Genome Project, or the President's National Climate Change Technology Initiative support our mission." DOE senior management and OMB have made performance, planning, and accountability high priority. In addition, the House Committee on Science has held hearings December 5th on how the nation's research federal establishment can contribute to the war on terrorism. The Office of Science and the Office of Budget and Management is presently conducting an inventory assessment of federal research related to terrorism. The urgent current activities are likely to lead to more mission responsibilities for the Office of Science, some of which will involve the need for advanced computing capability. As examples of how advanced computing can help, there is a new need for more rapid DNA sequencing of microbial pathogens used in biothreat agents and for faster and more detailed holographic imaging devices. Over the next few months it is expected that this new mission responsibility will be better articulated.

Many of these mission statements and requirements will involve increased computing capability. The Subcommittee sees the role of advanced computing in the Office of Science's research program growing significantly in the future and becoming more integrated into the Office's program and projects.

IV. Strategic Directions: *How might the roles of these facilities evolve to serve the missions of the Office of Science over the next three to five years?*

In addressing possible future roles of ASCR facilities, there are two strategic points to consider. First, mission directed research, from basic to highly applied, is the orientation of ASCR. Second, high end computing is the unique charge of ASCR within the DOE Office of Science.

Testimony from researchers within the Office of Science computing constituency has established that DOE Office of Science missions require computing resources that range from a small number of processors (local clusters) to the highest end (petaflops and beyond). These resources are to a large part already available within the Office, and are successful at satisfying mission needs.

However, in the absence of appropriately evolving usage and procurement strategies for existing resources, HPC hardware and funding can easily be misdirected for computing which either neglects the ASCR mission of high end computing in the exclusive service of mission needs, or neglects the missions in the interests of a few very large computational science projects. The balance of using high end resources effectively both from the perspective of flop rate and also from the perspective of Office of Science missions requires significant attention to allocation procedures and strategies. This balance, filling each ASCR machine to the greatest extent practicable with mission-relevant jobs that are not feasibly run on smaller machines, implies that (1) a range of machines are required within the ASCR portfolio, and (2) allocations will need to be

considered globally across that portfolio. The evolution of the internet with grid technologies will further enable such a strategy, allowing ready integration of technologies from the highest end HPC resources to few-processor clusters across geographically diverse communities of interest.

In summary, the Subcommittee applauds ASCR for its current success, and urges the Office to mandate the future. With mission-directed research from the basic to the highly applied as the orientation and driving excitement of the research fostered by ASCR, and high end computing the unique charge of that Office within the DOE Office of Science, the Subcommittee believes that ASCR provides an ideal environment in which to develop a plan to harvest a future that spans the spectrum from research and development of new architectures to the deployment of the multi-user computing facilities of the 21st Century. The Subcommittee strongly encourages the Office to develop such a blueprint that will provide the framework for the synthesis of research from applications people, computer scientists, and computational scientists in a way which, when achieved, will produce the paradigm shift to the next generation of High Performance Computing.

Please do not hesitate to contact me for further information regarding this report.

Yours truly,

Jill Dahlburg, Chair, ASCAC Subcommittee on Facilities
General Atomics
P.O.Box 85608
San Diego, CA 92186-5608