



## Future directions in scientific supercomputing for computational physics

**C. William McCurdy, Horst D. Simon<sup>\*</sup>, William T.C. Kramer,  
Robert F. Lucas, William E. Johnston, and David H. Bailey**

*Lawrence Berkeley National Laboratory, Mail Stop 50B-4230, Berkeley, CA 94720, U.S.A.*

**Elsevier use only:** Received date here; revised date here; accepted date here

---

### Abstract

NERSC, the National Energy Research Scientific Computing Center, is a leading scientific computing facility for unclassified research, and has had a significant impact on computational physics in the U.S. Here we will summarize the recent experience at NERSC, and present the four key elements of our strategic plan for the next five years. Significant changes are expected to happen in computational science during this period. Supercomputer centers worldwide must continue to enhance their successful role as centers that bridge the gap between advanced development in computer science and mathematics on one hand, and scientific research in the physical, chemical, biological, and earth sciences on the other. Implementing such a strategy will position NERSC and other centers in the U.S. to continue to enhance the scientific productivity of the computational physics community, and to be an indispensable tool for scientific discovery.

© 2001 Elsevier Science. All rights reserved

Keywords: high performance computing; supercomputer center; scientific challenge teams; computational and data grids;

---

<sup>\*</sup> Corresponding author. Tel.: +01-510-486-7377; fax: +01-510-486-4300; e-mail: hdsimon@lbl.gov.

## 1. Introduction

NERSC, the National Energy Research Scientific Computing Center, is one of the premier scientific computing facilities for unclassified research in the physical sciences in the U.S. It delivers high-end capability computing services and support to a large segment of the computational physics community. This paper presents NERSC's vision for its activities and new directions over the next five years. Significant changes are expected to happen in computational science during this period. Traditional supercomputer centers will have to anticipate these changes, and be ready yet again to change their model of service.

### 1.1. Emerging Changes in Computational Science

In the 1980s physicist Ken Wilson and others described computing as the new "third component" of the scientific method, achieving parity with experiment and theory. There is abundant evidence that computational physicists have already realized this vision of the critical role of computing. At the present time, two additional trends are apparent: (1) the emergence of large, multidisciplinary teams for scientific research, and (2) the convergence of computation, experiment, and theory on an interactive, real-time basis. These changes will form the basis of new directions and priorities for NERSC.

The first of these two trends, typified by what we call *Scientific Challenge Teams*, represents a break from the single-principal-investigator (PI) model for computational science that has dominated computing in many areas of the natural sciences for most of the last 50 years. The emergence of these teams is primarily motivated by the increased complexity of the science being investigated, but also by the need to close the growing gap between the sustained performance obtained by straightforward porting of single-processor codes and the peak performance capability of these systems. On vector supercomputers of the 1980s and early 1990s, many scientific codes realized 30% to 50% of peak performance. By contrast, scientific codes typically realize only 5% to 15% of peak performance on

massively parallel supercomputers today. It is now clear that exploiting the full potential of massively parallel computers as scientific tools will require much greater effort and a wider range of skills than can be brought to bear by a single PI working with a few collaborators, trained in a single scientific discipline.

The second change that is occurring is the simultaneous convergence of computing, experiment, and theory in scientific research. One example, familiar at NERSC because the Center has been supporting it since 1996, is the Supernova Cosmology Project. To find Type 1A supernovae, telescope images are analyzed using supercomputers, and the results are used to quickly direct telescopes around the world to observe and measure spectra of candidate objects. Simulations and predictions of the spectra based on the abundance of each elemental candidate are performed simultaneously to refine the identification. The entire effort requires the availability and scheduling of resources throughout the world.

It was for such applications that the concept of Computational and Data Grids (discussed in detail in Section 5) was invented. In order to exploit this technology, NERSC must become an active node on the Grid. Such a transformation is not accomplished by merely connecting to the Grid: it must be accomplished by the staged development of a new architecture for high-end computing that incorporates, schedules, and manages the computing and data resources at NERSC and elsewhere into a new and evolving national and international infrastructure.

### 1.2. Principal Components of the Strategy

The goal for the next five years is to address these changes in the context and requirements of the traditional user spectrum, while continuing to strengthen the established user-oriented approach. To accomplish this goal, we have defined four components of a strategy, two ongoing and two new ones (see Figure 1). The two ongoing components are:

*High-End Systems* — NERSC will continue to focus on balanced introduction of the best new technologies for complete computational and storage

systems, coupled with the advanced development activities necessary to wisely incorporate these new technologies.

*Comprehensive Scientific Support* — NERSC will continue to provide the entire range of support activities, from high-quality operations and client services to direct collaborative scientific support, to enable a broad range of scientists to effectively use the NERSC systems in their research.




Fig. 1. The four principal components of the next-generation NERSC.

The new components are:

*Support for Scientific Challenge Teams* — NERSC will concentrate its resources on supporting these teams, with the goal of bridging the software gap between currently achievable and peak performance on the new terascale platforms.

*Unified Science Environment (USE)* — NERSC will enhance its architecture and systems as required to make NERSC the most powerful computational resource on the Department of Energy's (DOE's) Science Grid. Over the next five years, NERSC will use Grid technology to deploy a capability designed to meet the needs of an integrated science environment, combining experiment, simulation, and theory by facilitating access to computing and data resources, as well as to large DOE experimental instruments.

## 2. High-End Systems

*Providing the most effective and most powerful high-end systems possible.* This is the foundation upon which NERSC builds all other services in order to enable computational science for the DOE/SC community. High-end systems at NERSC mean more than highly parallel computing platforms — they also include a very large-scale archival storage system, auxiliary and developmental platforms, networking and infrastructure technology, system software, productivity tools for clients, and applications software. Our high-end system strategy includes advanced development work, evaluating new technologies, developing methodologies for benchmarking and performance evaluation, and acquisition of new systems.

NERSC plans to introduce the NERSC-4 system in 2003 and the NERSC-5 system in 2006, each with a factor-of-4 increase in capability over the previous-generation system, bringing NERSC-5 to approximately 45 teraflop/s peak performance. Figure 2 shows the proposed peak computing power of NERSC. But computing power is only one measure of capability. NERSC will continue to increase the capacity of its storage system, reaching at least 25 petabytes in 2006 (Figure 3).

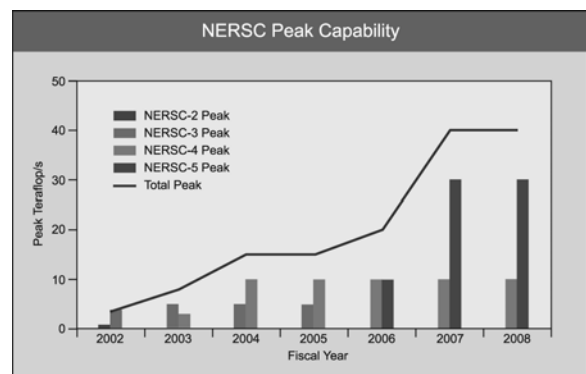


Fig. 2. Computational capability growth of the NERSC Facility.

A major effort will be made in developing and deploying a Global Unified Parallel File System (GUPFS). The development of GUPFS will take existing work as its point of departure, including the Global File System work initially developed at the

University of Minnesota and being developed further at Sстина Corporation, developments within ASCI Pathforward in high-performance computer file system technology, and the IBM Global Parallel File System. This capability will not only increase scientific productivity by simplifying file access, but will also be one of the foundations for USE, which is described in Section 5.

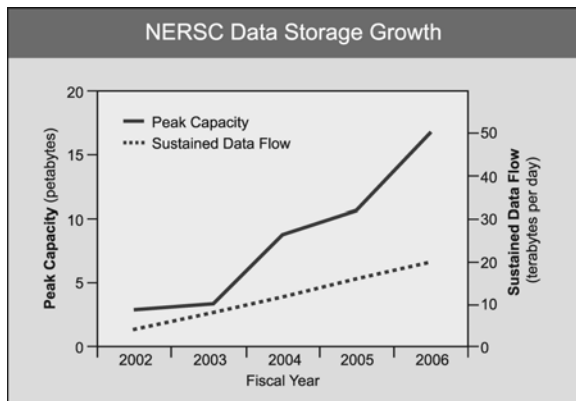


Fig. 3. Predicted NERSC data storage growth, 2002–2006.

### 3. Comprehensive Scientific Support

The goal of NERSC’s Comprehensive Scientific Support function is to make it easy and practical for DOE computational scientists to use the NERSC high-end systems by:

- Providing consistent high quality service to the entire NERSC client community through the support of the early, production quality, large-scale capability systems.
- Aggressively incorporating new technology into the production NERSC facility by working with other organizations, vendors, and contractors to develop, test, install, document, and support new hardware and software.
- Ensuring that the production systems and services are the highest quality, stable, secure, and replaceable within the constraints of budget and technology.

- Participating in other work to understand and address the unique issues of using large-scale systems.

### 4. Support for Scientific Challenge Teams

The arrival of large, highly parallel supercomputers in the early 1990s fundamentally changed the mode of operation for successful computational scientists. In order to take full advantage of the new capabilities of these parallel platforms, scientists organized themselves into national teams. Called “Grand Challenge Teams,” they were a precursor to the “Scientific Challenge Teams” that NERSC anticipates as its leading clients in the next decade. These multidisciplinary and multi-institutional teams engage in research, development, and deployment of scientific codes, mathematical models, and computational methods to maximize the capabilities of terascale computers.

In March 2000 DOE launched a new initiative called “Scientific Discovery through Advanced Computing” (SciDAC). SciDAC defines and explicitly calls for the establishment of Scientific Challenge Teams. These Teams are characterized by large collaborations, the development of community codes, and the involvement of computer scientists and applied mathematicians. In addition to high-end computing, teams will also have to deal increasingly with issues in data management, data analysis, and data visualization. The expected close coupling to scientific experiments supported by the USE environment (described in the Section 5) will be an essential requirement for success for some teams. Scientific Challenge Teams represent the only approach that will succeed in solving many of the critical scientific problems in the DOE Office of Science’s research programs. These teams are the culmination of the process of users moving to ever-higher computing capability, and NERSC’s new structure enables that entire process (see Figure 4).

One primary motivation for the establishment of the SciDAC Scientific Challenge Teams is to bridge the software gap between currently achievable and peak performance on terascale platforms. With the previous generation of vector supercomputers, many scientific codes realized 30% to 50% of the peak

performance of the supercomputer. By contrast, with the current generation of microprocessor-based parallel supercomputers, scientific computing codes often realize only 5% to 15% of the potential peak performance of the computer. This gap will continue to increase with increasing use of parallelism. Closing this potential usage gap represents a major challenge to the scientific computing community.

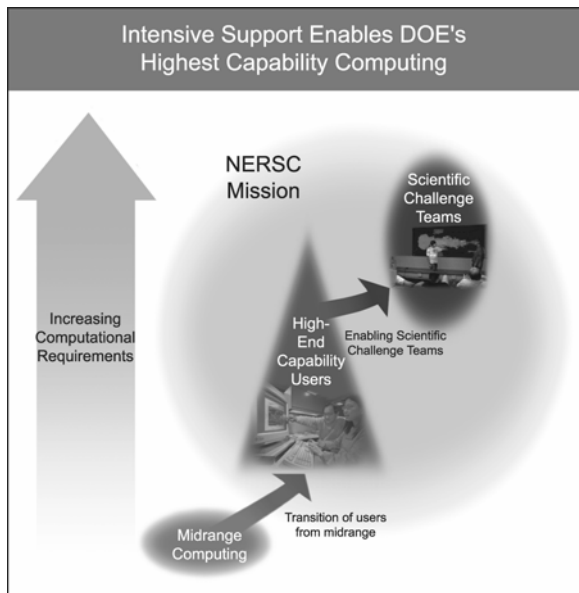


Fig. 4. NERSC facilitates the transition to high-end capability computing, and enables Scientific Challenge Teams through intensive support.

NERSC's proposed focus on Scientific Challenge Teams anticipates these new teams and their special requirements. In addition, we expect that many other scientists will organize themselves into similar collaborations in the coming years. SciDAC will be the catalyst for a fundamental shift in computational science from the principal-investigator model to the collaborative-team model. SciDAC also proposes to establish a set of new centers in computer science and applied mathematics, called Integrated Software Infrastructure Centers (ISICs), with the goal of supporting research, development, and deployment of software in order to

- accelerate the development of and protect long-term investments in scientific codes

- achieve maximum efficiency on terascale computers
- enable a broad range of scientists to use simulation in their research.

NERSC's strategy for the next five years is to build a focused-support infrastructure for the Scientific Challenge Teams consisting of four components:

- integrated support and collaboration from the NERSC staff
- deployment of tools developed by ISICs
- deployment of grid and collaboration technologies (USE)
- building the software engineering infrastructure.

Of these four components, only the software engineering support function is new; all the others are based on either leveraged or redirected existing resources. NERSC's main role here will be to advocate for software engineering, and to provide the tools and training. It is important for NERSC to evaluate the needs of the Challenge Teams first and not push "solutions." NERSC will engage the scientists and motivate them to consider software engineering practices.

## 5. Unified Science Environment

A new paradigm is emerging in which multiple national assets, both computational and experimental, are simultaneously employed in the process of scientific discovery. These systems are geographically distributed and managed by different organizations. Projects such as the DOE Science Grid are constructing an infrastructure of software and services that will automate the task of managing these resources. NERSC calls this environment, integrated with NERSC's high-end computational resources, the "Unified Science Environment" (USE).

Inclusion of NERSC in the DOE Science Grid will make high-end services available to NERSC computational scientists through the uniform Grid environment (see Figure 5). The resulting combination of Grid access to desktop, midrange, and high-end services creates the USE.

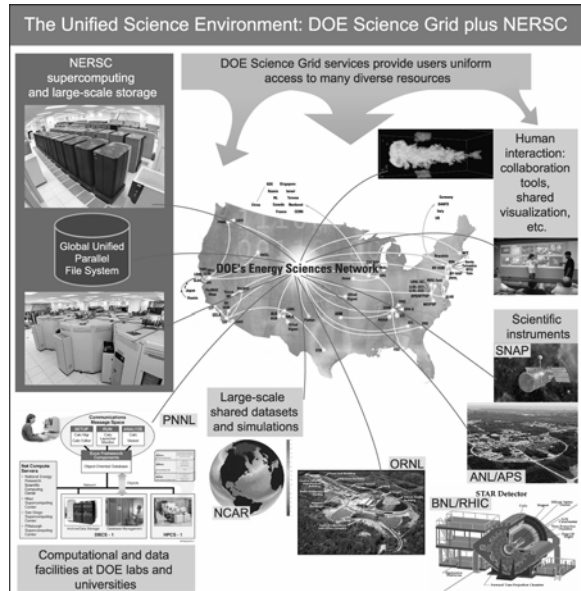


Fig. 5. The role of NERSC as the largest computational resource in the DOE Science Grid.

The model for USE derives from changes that are, of necessity, occurring in how science is done: computing, collaboration, and data must be tightly coupled with theory and experiment across many organizations in order to enable the next generation of science. Examples of this unified approach to computing and science may be found in many of DOE's large-scale science projects, such as accelerator-based science, climate analysis, collaboration on very large simulation problems, and observational cosmology. These activities occur in widely distributed environments and under circumstances that are constrained by the timing of the experiments or collaborations.

This time-constrained use of supercomputing in the heart of the observational cycle represents the new aspect of computing integrated with science. Implicit in these computing models is the assumption that coordination and operation of both these large central facilities and extended environments, within which instruments and researchers are distributed across the globe, are straightforward and easily managed. They are not easily managed now, and improving their ease of management is the sort of

capability that integrating NERSC and Grids will address.

Diverse development activities throughout the global Grids community, coordinated through work in the Global Grid Forum, will provide a steady flow of technology that will form the basis of USE at NERSC. NERSC, in turn, must adapt these development efforts, and forge them into a robust supercomputing environment.

## 6. Conclusions

In 2001 two new trends are apparent: (1) the emergence of large, multidisciplinary teams for computational physics simulations, and (2) the convergence of computation, experiment, and theory on an iterative, real-time basis. We have presented a vision of how a supercomputer center will change its activities and set new directions over the next five years to address these changes.

Supercomputer centers worldwide must continue to enhance their successful role as centers that bridge the gap between advanced development in computer science and mathematics on one hand, and scientific research in the physical, chemical, biological, and earth sciences on the other. Implementing the above strategy will position NERSC and other centers in the U.S. to continue to enhance the scientific productivity of the computational physics community, and to be an indispensable tool for scientific discovery.

## Acknowledgments

We would like to thank John Hules for his expert help in editing this paper, and Flavio Robles for the creation of the artwork. This work was supported by the Director, Office of Advanced Scientific Computing Research, Division of Mathematical, Information, and Computational Sciences of the U.S. Department of Energy under contract number DE-AC03-76SF00098.