

Lawrence Berkeley National Laboratory

Lawrence Berkeley National Laboratory

Title

Supporting National User Communities at NERSC and NCAR

Permalink

<https://escholarship.org/uc/item/2f8300b9>

Authors

Killeen, Timothy L.
Simon, Horst D.

Publication Date

2006-05-16



LBLN-60166

Supporting National User Communities at NERSC and NCAR

Timothy L. Killeen
National Center for Atmospheric Research, NCAR
Boulder, Colorado 80301

and

Horst D. Simon
NERSC Center Division
Ernest Orlando Lawrence Berkeley National Laboratory
University of California
Berkeley, California 94720

May 16, 2006

This work was supported by the Director, Office of Science, Office of Advanced Scientific Computing Research of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

NCAR is operated by the University Corporation for Atmospheric Research under sponsorship of the National Science Foundation.

1. Introduction

The National Energy Research Scientific Computing Center (NERSC) and the National Center for Atmospheric Research (NCAR) are two computing centers that have traditionally supported large national user communities. Both centers have developed responsive approaches to support these user communities and their changing needs, providing end-to-end computing solutions. In this report we provide a short overview of the strategies used at our centers in supporting our scientific users, with an emphasis on some examples of effective programs and future needs.

2. Science-Driven Computing at NERSC

2.1 NERSC's MISSION

The mission of the National Energy Research Scientific Computing Center (NERSC) is to accelerate the pace of scientific discovery by providing high performance computing, information, data, and communications services for research sponsored by the U.S. Department of Energy (DOE) Office of Science (SC). NERSC is the principal provider of high performance computing services for the capability needs of Office of Science programs — Fusion Energy Sciences, High Energy Physics, Nuclear Physics, Basic Energy Sciences, Biological and Environmental Research, and Advanced Scientific Computing Research.

Computing is a tool as vital as experimentation and theory in solving the scientific challenges of the twenty-first century. Fundamental to the mission of NERSC is enabling computational science of scale, in which large, interdisciplinary teams of scientists attack fundamental problems in science and engineering that require massive calculations and have broad scientific and economic impacts. Examples of these problems include global climate modeling, combustion modeling, magnetic fusion, astrophysics, computational biology, and many more. NERSC uses the Greenbook process (see ^{1,2} for more details) to collect user requirements and drive its future development.

Lawrence Berkeley National Laboratory (Berkeley Lab) operates and has stewardship responsibility for NERSC, which, as a national resource, serves about 2,400 scientists annually throughout the United States. These researchers work at DOE laboratories, other Federal agencies, and universities (over 50% of the users are from universities). Computational science conducted at NERSC covers the entire range of scientific disciplines, but is focused on research that supports DOE's missions and scientific goals.

2.2 A SCIENCE-DRIVEN STRATEGY TO INCREASE SCIENTIFIC PRODUCTIVITY

Since its founding in 1974, NERSC has provided systems and services that maximize the scientific productivity of its user community. NERSC takes pride in its reputation for the expertise of its employees and the high quality of services delivered to its users. To maintain its effectiveness, NERSC proactively addresses new challenges. We observe three trends that NERSC needs to address over the next several years:

- the widening gap between application performance and peak performance of high-end computing systems
- the recent emergence of large, multidisciplinary computational science teams in the DOE research community

¹ *Science Driven Computing: NERSC's Plan for 2006–2010* (Horst D. Simon et al., Berkeley Lab report LBNL-57582, May 2005), <http://www.nersc.gov/news/reports/LBNL-57582.pdf>.

² *DOE Greenbook: Needs and Directions in High Performance Computing for the Office of Science* (S. C. Jardin, editor, Princeton Plasma Physics Laboratory report PPPL-4090, June 2005), <http://www.nersc.gov/news/greenbook/2005greenbook.pdf>.

- the flood of scientific data from both simulations and experiments, and the convergence of computational simulation with experimental data collection and analysis in complex workflows.

NERSC's responses to these trends are the three components of the science-driven strategy that NERSC will implement and realize in the next five years: *science-driven systems*, *science-driven services*, and *science-driven analytics* (Figure 1). This balanced set of objectives will be critical for the future of the enterprise and its ability to serve the DOE scientific community.

- **Science-Driven Systems:** Balanced introduction of the best new technology for complete computational systems — computing, storage, networking, visualization and analysis.
- **Science-Driven Services:** The entire range of support activities, from high-quality operations and user services to direct scientific support, that enable a broad range of scientists to effectively use NERSC systems in their research. NERSC will concentrate on resources needed to realize the promise of the new highly scalable architectures for scientific discovery in multidisciplinary computational science projects.
- **Science-Driven Analytics:** The architectural and systems enhancements and services required to integrate NERSC's powerful computational and storage resources to provide scientists with new tools to effectively manipulate, visualize, and analyze the huge data sets derived from both simulation and experiment.

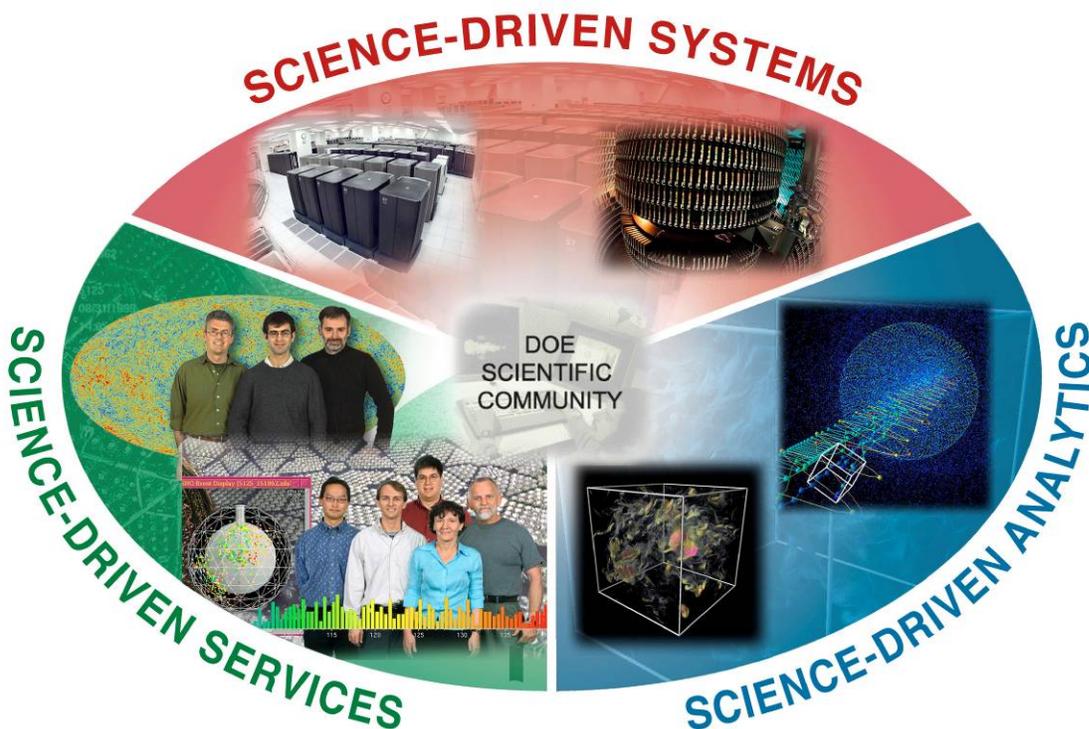


Figure 1. Conceptual diagram of NERSC's plan for 2006–2010.

Science-Driven Systems

Applications scientists have been frustrated by a trend of stagnating application performance relative to dramatic increases in claimed peak performance of high performance computing systems. This trend has been widely attributed to the use of commodity components whose architectural designs are unbalanced and inefficient for large-scale scientific computations. It was assumed that the ever-increasing gap between theoretical peak and sustained performance was unavoidable. However, results from the Earth Simulator (ES) in Japan clearly demonstrate that a close collaboration with a vendor to develop a science-driven solution can produce a system that achieves a significant fraction of peak performance for critical scientific applications.

Realizing that effective large-scale system performance cannot be achieved without a sustained focus on application-specific systems development, NERSC has begun a science-driven systems strategy. The goal of this effort is to influence the vendors' product roadmaps to improve system balance and to add key features that address the requirements of demanding capability applications at NERSC — ultimately leading to a sustained petaflop/s system for scientific discovery. This strategy involves extensive interactions between domain scientists, mathematicians, computer experts, as well as leading members of the vendors' research and product development teams.

NERSC must be prepared for disruptive changes in processor, interconnect, and software technologies. Obtaining high application performance will require the active involvement of NERSC in understanding, driving, and adopting these technologies. The move towards open-source software will require additional efforts in software integration at NERSC.

The goal of the science-driven systems strategy is to enable new scientific discoveries, and that requires a high level of sustained system performance on scientific applications. The NERSC approach takes into account both credibility and risk in evaluating systems and will strike a balance between innovation and performance on the one hand and reliability on the other. While the discussion often focuses on the high-end platforms, NERSC will continue to emphasize maintaining Center balance, that is, improving all the systems at NERSC — storage, networking, visualization and analysis — commensurately with improvements in the high-performance computing platforms.

Science-Driven Services

The DOE computational science community, in all its disciplines, has been organizing itself into large multidisciplinary teams. This trend was driven by the DOE Scientific Discovery through Advanced Computing (SciDAC) initiative, but has reached beyond the SciDAC teams. This trend has been driven by necessity as well as opportunity. The transformation became most apparent after massively parallel computers came to dominate the high end of available computing resources.

Technology trends indicate that the gap between the peak performance of next-generation systems and performance that is easily attainable could increase even more. NERSC has been focused on working with computational scientists to close this gap and help them scale their applications efficiently to current platforms. NERSC has formulated a science-driven services strategy that will address the requirements of these large computational science teams even more

so than in the past, while at the same time maintaining the high level of support for all of its users.

Science-Driven Analytics

A major trend occurring in computational science is the flood of scientific data from both simulations and experiments, and the convergence of experimental data collection, computational simulation, visualization, and analysis in complex workflows. Deriving scientific understanding from massive datasets from major experimental facilities is a growing challenge.

In recent years NERSC has seen a dramatic increase in the data arriving from DOE-funded research projects. This data is stored at NERSC because NERSC provides a reliable long-term storage environment that assures the availability and accessibility of data for the community. NERSC has helped accelerate this development by deploying Grid technology on all of its systems and by enabling and tuning high performance wide-area network connections to major facilities, for example the Relativistic Heavy Ion Collider at Brookhaven National Laboratory.

Now NERSC must invest resources to complete an environment that allows easier analysis and visualization of large datasets derived from both simulation and experiment. Our third new thrust in science-driven analytics will enable scientists to combine experiment, simulation, and analysis in a coordinated workflow. This thrust will include activities enhancing NERSC's data management infrastructure, expanding NERSC's visualization and analysis capabilities, enhancing NERSC's distributed computing infrastructure, and understanding the analytical needs of the user community.

2.3 A KEY RESOURCE FOR THE DOE OFFICE OF SCIENCE

In "Facilities for the Future of Science: A Twenty Year Outlook," the Office of Science has identified the need for creating new and/or improving on the current computational capability as a critical aspect of realizing its advanced scientific computing research vision.³ It identified the NERSC upgrade as a near-term priority to ensure that NERSC, DOE's premier scientific computing facility for unclassified mission-critical research, continues to provide high-performance computing resources to support the requirements of scientific discovery.

As a high-end facility that serves all the DOE SC programs with capability and high-end capacity resources, NERSC is a key resource in SC's portfolio of computing facilities. NERSC has established a reputation for providing reliable and robust services along with unmatched support to its users. Because of investments such as SciDAC, and the important role that computation will play in Genomics:GTL (formerly Genomes to Life) and the Nanoscale Science Research Centers, demands for computational resources in SC will continue to grow at a rapid rate, and NERSC's growth must keep pace. NERSC supports a large number (200–300) of projects of medium to large scale, occasionally needing a very high capability resource, that fall within the mission of the Office of Science. The scientific productivity enabled by NERSC is

³ *Facilities for the Future of Science: A Twenty-Year Outlook* (Washington, DC: U.S. Department of Energy, Office of Science, November 2003).

demonstrated by the 3,654 papers in refereed publications from 2003 to 2005 that were based at least in part on work done at NERSC.⁴

In NERSC's experience there is a continuum of scientific computing systems and facilities, as represented in Figures 2 and 3. There are a few research groups with experienced users and very high computational requirements who are in a good competitive position to use a leadership class facility. There are a much larger number of PIs and projects with high-end requirements who are best served by NERSC's high-end systems and comprehensive services, both of which distinguish NERSC from leadership computing and midrange computing centers, such as institutional or departmental clusters. Capability users include both single-principal-investigator (PI) teams and community science teams. NERSC's science-driven services are important for both types of high-end users.

NERSC supports large-scale teams working on advanced modeling and simulation "community codes" whose development is shared by entire scientific research communities. These codes employ new mathematical models and computational methods designed to better represent the complexity of physical process and to take full advantage of current computational systems. NERSC provides focused support for these teams.

NERSC also supports single-PI teams consisting of a lead researcher and his or her group of collaborators, postdocs, and students, usually concentrated at a single location. For this class of users, NERSC's science-driven service is important because they usually are less knowledgeable about computational technologies and they lack the resources to establish in-depth collaborations with computer science or mathematics experts. Computing at NERSC not only produces important scientific insights but also gives these users and teams the opportunity to advance to the leadership computing level for their most challenging computations.

NERSC, as a centralized facility, properly staffed and managed, provides the best possible mechanism for technology transfer between the computational efforts of different research programs. Moreover, a concentration of computing resources provides a more flexible mechanism to address changing priorities. SC's priorities for its programs sometimes change quickly because it is a mission agency. A general-purpose facility like NERSC, with a staff prepared to support the broadest possible array of scientific disciplines, allows DOE to switch priorities and quickly apply its most powerful computing resources to new challenges.

NERSC's role as a general scientific computing facility requires it to provide resources that are of common utility to the programs of the Office of Science. However, NERSC must be responsive to the specific needs of each program. Specific support for different programs, tailored to their varying needs, has been a key to the success of the Center. Examples range from the collaborative effort of NERSC staff in scaling INCITE applications to 2,048 and 4,096 processors, to the operation of the PDSF cluster for the high energy and nuclear physics communities. The breadth of NERSC's support is best expressed by Figures 4 and 5, which summarize NERSC usage by discipline and institution.

⁴ <http://www.nersc.gov/news/reports/>

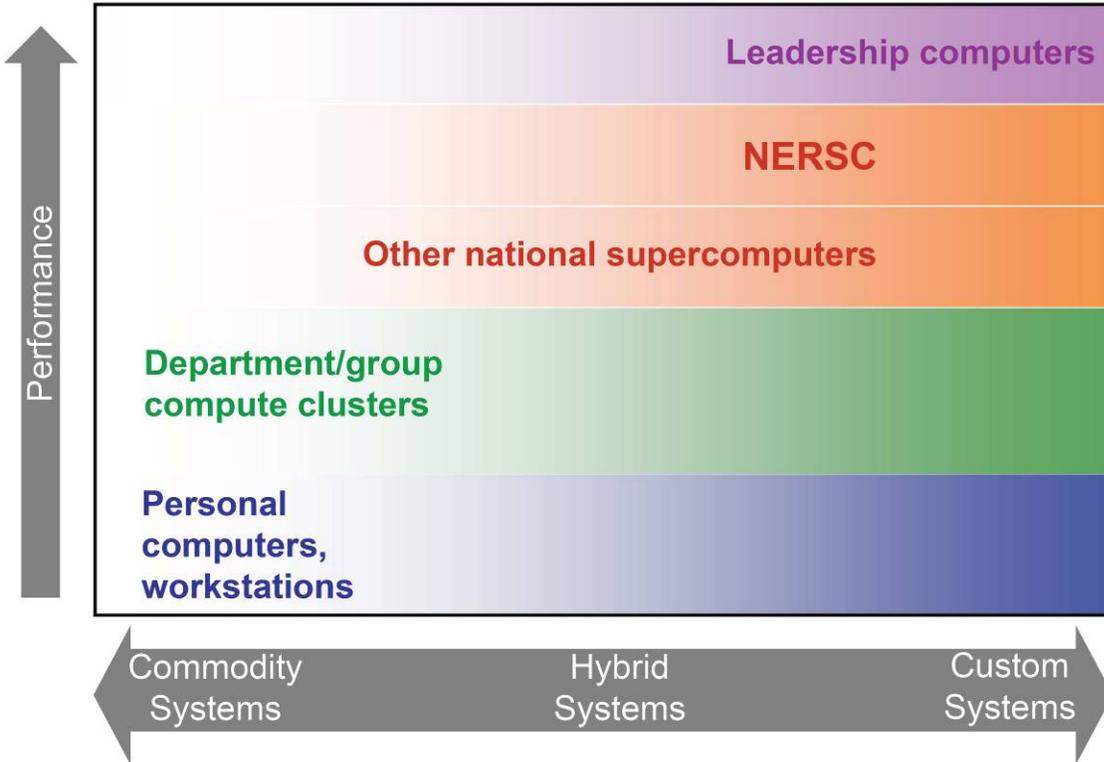


Figure 2. The continuum of scientific computing systems.

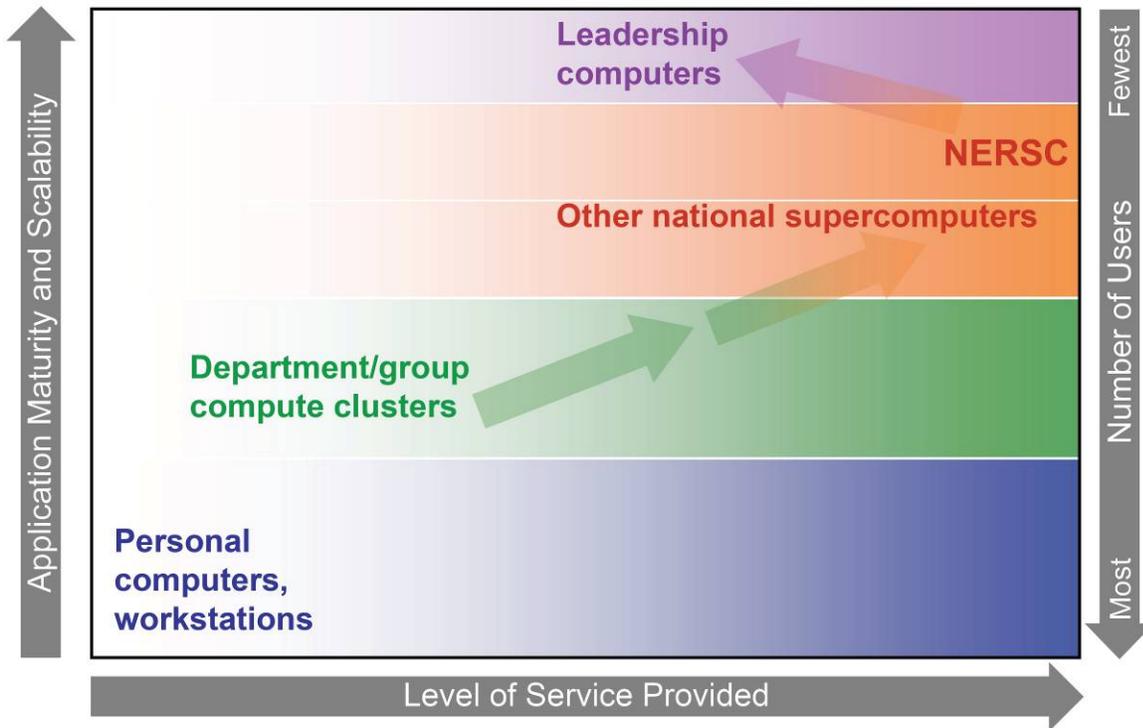


Figure 3. The continuum of scientific computing facilities.

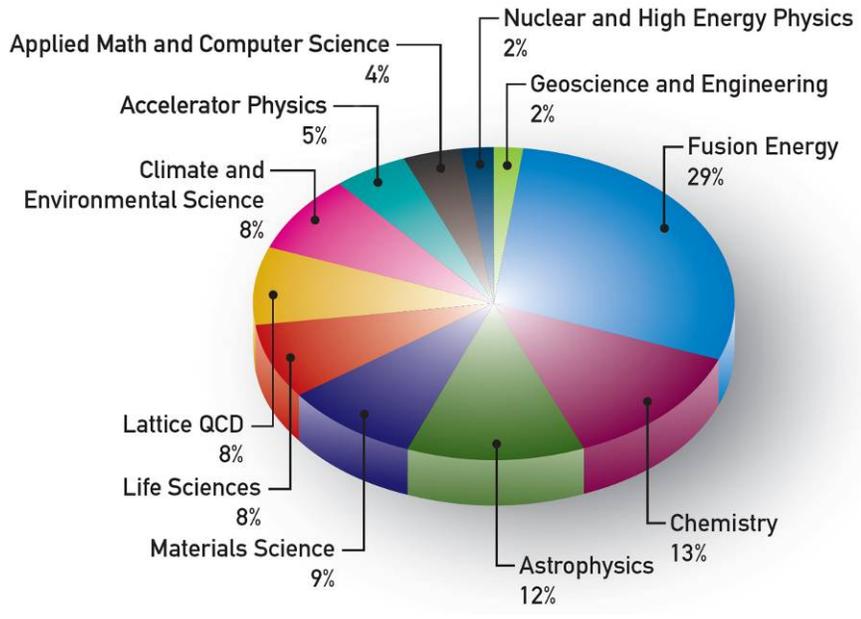


Figure 4. NERSC usage by scientific discipline for FY2005.

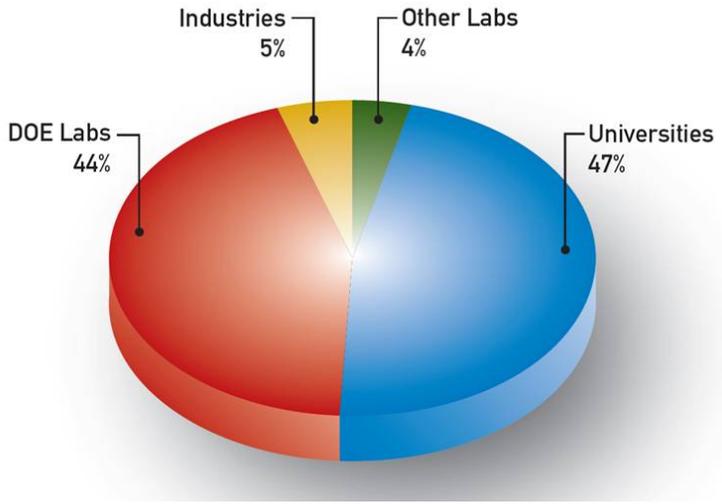


Figure 5. NERSC users by institution type for FY2005.

3. Science-Driven Computing at NCAR

3.1 NCAR's MISSION

The mission of NCAR is to support, enhance, and extend the capabilities of the university community nationally and internationally to understand the behavior of the atmosphere and the global environment and to foster the transfer of knowledge and technology for the betterment of life on Earth.⁵ NCAR is a principal provider of high performance computing services for the academic geosciences community in the United States and has a 48-year record of providing community supercomputing services. Over the years, NCAR and the community it serves have contributed centrally to the scientific underpinnings of numerical weather forecasting and climate modeling as well as the understanding of:

- the coupled ocean/atmosphere climate system
- the detailed chemistry of the stratosphere and troposphere
- solar magnetism, helioseismicity, and solar coronal mass ejections
- the dynamics and chemistry of the upper atmospheres of earth and other planets
- the microphysics of clouds and convective processes
- the socioeconomic impacts of climate change and severe weather
- the role of human activities in causing and responding to large-scale Earth system change.

Computing continues to be an essential part of NCAR's work, and the center has a commitment to end-to-end services, spanning high performance computing, application development and user support services, data management and data curation, visualization, networking, middleware, and all the components of what is commonly referred to as *cyberinfrastructure*. The emphasis at NCAR is on solving computing problems related to the geosciences, and the NCAR computational architecture acquisition and system support decisions are centered on the needs of this large but finite scientific domain. Human capital development is an essential part of this commitment.

In a similar fashion to NERSC, NCAR favors a balanced approach to high performance computing, stressing robust operational performance of diverse computing platforms with regular upgrades paths (see Figure 6), sophisticated application development, attention to software reuse and application portability with careful verification pathways, computational efficiency, redundant mass storage, and secure data management systems. NCAR has experienced many of the same trends and challenges reported by NERSC above, including the move to larger and more interdisciplinary teams of investigators, the need to close the gap between "sustained" and "peak" performance, and the requirement for matching the data system performance with application needs.

⁵ *NCAR as Integrator, Innovator, and Community Builder: A Strategy-Implementation Plan for the National Center for Atmospheric Research* (January 2006), <http://www.ncar.ucar.edu/stratplan/2006/stratplan06.pdf>.

Estimated Sustained GFLOPs at NCAR (with ICESSES)

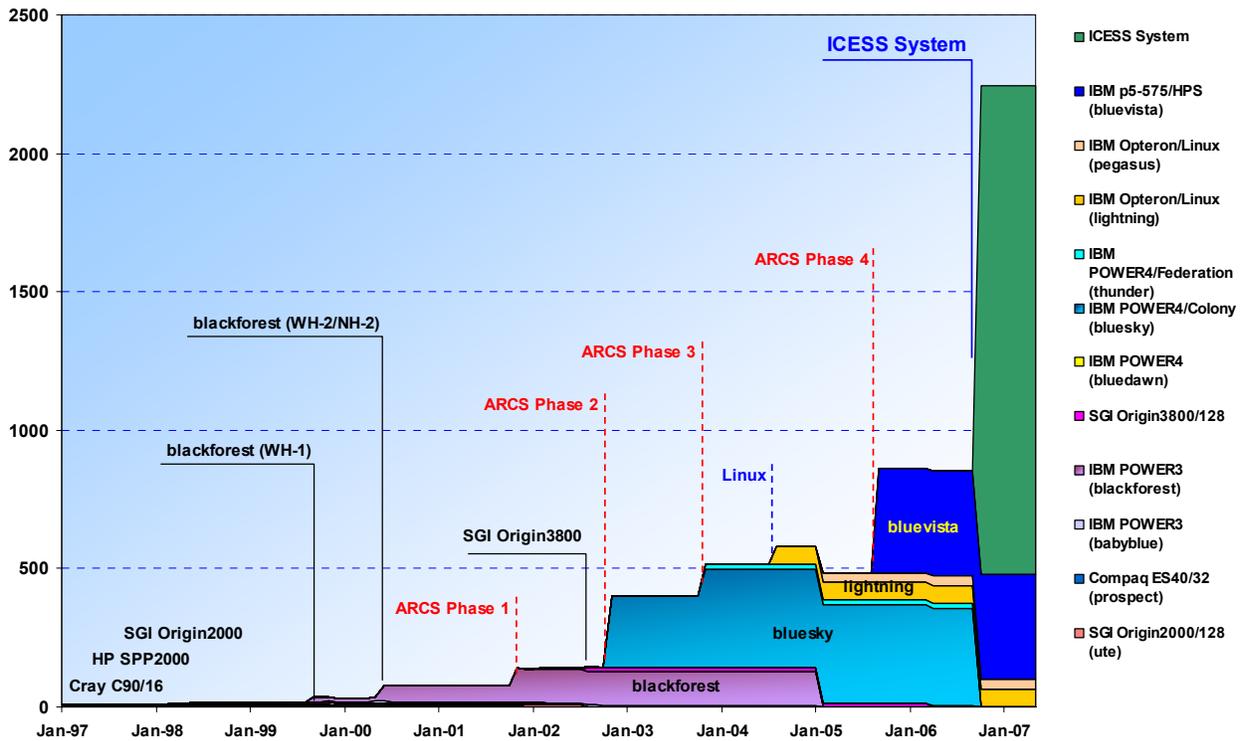


Figure 6. Sustained performance of applications running on NCAR computing platforms over the past 9 years. ICESSES stands for the NCAR Integrated Computing Environment for Scientific Simulation, an ongoing procurement effort.

NCAR supports a *community model* approach that is perhaps unique among the large computational centers in the United States. This approach involves the development of well supported, open-source, large-scope codes that have lifetimes of years to decades, are regularly enhanced and updated to reflect emerging scientific needs, and are managed and driven by the broad academic community, with NCAR playing the key coordinating role. NCAR's community models are freely available to all and are supported with help desks, version control systems, extensive documentation, regular user tutorials and workshops, and a significant body of peer-reviewed publications describing both computational and scientific aspects. Important examples of NCAR-managed community models include the Community Climate System Model (CCSM) and the Weather Research and Forecasting Model (WRF). Another important community product managed by NCAR is the Earth System Modeling Framework (ESMF). Brief descriptions of these three community science activities at NCAR are provided below to illustrate how NCAR supports national user communities.

3.2 THE NCAR COMMUNITY CLIMATE SYSTEM MODEL PROGRAM

The Community Climate System Model (CCSM)⁶ is a comprehensive system for studying the past, present, and future of the Earth. In contrast to traditional weather-forecast models that focus only on the atmosphere, the CCSM includes components that simulate the evolution and interactions among the atmosphere, ocean, land surface, and sea ice. The principal objectives of the CCSM program are to develop a comprehensive numerical model with which to study the Earth's present climate, to investigate seasonal and interannual variability in the climate, to explore the history of the Earth's climate, and to simulate the future of the environment for policy formation.

CCSM has been designed with input from a broad community of climate scientists, computer scientists, and software engineers. This community also shares the scientific code and results produced by the model. In fact, CCSM is the only climate model that is developed as open-source code and is distributed via the Web to the worldwide climate community. CCSM is funded with support from the National Science Foundation (NSF), the Department of Energy (DOE), the National Aeronautics and Space Administration (NASA), and the National Oceanographic and Atmospheric Administration (NOAA). The CCSM community includes some 900 members, located at universities and laboratories throughout the world.

In order to support a broad community, CCSM must operate both as a research and an operational climate model, and therefore must be easily portable to a wide range of computational platforms. CCSM or its components can be run "out of the box" on a variety of Linux clusters, Apple servers, SGI Origin and Altix systems, and IBM and Intel clusters. It has also been enabled on NEC and Cray vector supercomputers, IBM Power-series clusters, and Cray clusters of scalar processors. The developers are now exploring modifications to CCSM to ensure efficient execution on other massively parallel architectures. The CCSM team has developed a comprehensive suite of tests to insure that the model algorithms work reliably and transparently across such a heterogeneous computing environment.

CCSM is designed to be flexible and extensible, an important characteristic since it will serve as a basis for the development of a more complete Earth System Model over the next several years. This Earth System Model will simulate the chemical, biogeochemical, and physical state of the climate system. The CCSM development effort is managed by a Scientific Steering Committee (SSC) with membership from the broad academic research community, as well as from NCAR.

The CCSM results for the Intergovernmental Panel on Climate Change (IPCC) provide a sobering look into the future of the planet and are being documented in more than 200 peer-reviewed scientific publications. Figure 7 shows projections of the time evolution of summer Arctic ice area for several IPCC greenhouse gas forcing scenarios. Note that summer ice is projected to disappear from the Arctic toward the latter part of this century under the IPCC "A2" scenario for socioeconomic development.

⁶ <http://www.cesm.ucar.edu>

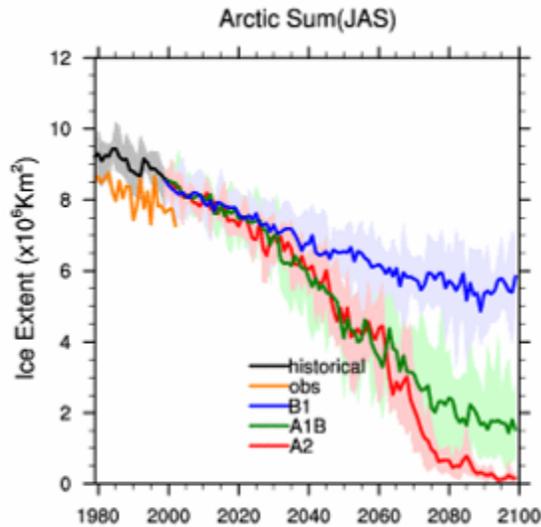


Figure 7. CCSM IPCC ensemble simulations of Arctic ice extent for the next century. The individual curves represent IPCC scenarios, and the shaded regions provide the uncertainty bounds from the multiple realizations.

3.3 THE WEATHER RESEARCH AND FORECASTING MODEL

A longtime focus in numerical modeling of the atmosphere has been the development and improvement of capabilities that can simulate the conditions that dictate the weather. Such systems are typically called *mesoscale* atmospheric models, where *mesoscale* refers to the spatial dimension over which most of the weather that daily influences human activity occurs. NCAR has been developing a new numerical weather prediction (NWP) model which is now coming into its own: the Weather Research and Forecasting Model (WRF).⁷ WRF is employed worldwide, with the largest number of registered users (over 3,700) for any such model today.

The WRF model is different from existing NWP technologies in a number of ways. Rather than being created by a single researcher, institution, or agency, WRF was developed in the U.S. through a partnership of both research and operational (i.e., official weather forecasting) groups. The initial development began in 1997, and the partners have been NCAR, the U.S. National Centers for Environmental Prediction (NCEP), the U.S. Air Force Weather Agency, the U.S. Navy's Naval Research Laboratory, the National Oceanic and Atmospheric Administration's Earth System Research Laboratory, the Federal Aviation Administration (FAA), and Oklahoma University. The goal was to create an NWP tool for use by both the operational and research meteorological communities. A key motivation was having a vehicle that, with relative ease and rapidity, could make the latest in research advances available to public forecasting.

The WRF modeling system features a software framework that is modular, plug-compatible, and allows portability to a wide range of computer architectures. It runs on hardware from laptops to desktop workstations to PC Linux clusters to high performance supercomputers. WRF is parallelized and is efficient in massively parallel distributed-memory environments. The

⁷ <http://www.wrf-model.org>

software framework permits ease of coupling with other earth system numerical models (e.g., ocean circulation codes or air chemistry modules). WRF also provides sophisticated *data assimilation* — the incorporation of observed meteorological information from satellites and other observing systems.

WRF is currently being used for official forecasting in the U.S. by NCEP, which provides NWP model guidance for the forecasters of the National Weather Service. On the research side, WRF's applications range from study of atmospheric processes and weather from the tropics to the poles. Targets of special interest for WRF so far have been severe thunderstorms and powerfully damaging hurricanes, given their enormous societal impacts in the U.S. For the past three hurricane seasons, for example, WRF has been run at NCAR in real time to offer high-resolution (i.e., detailed) forecasts of storms that have threatened landfall. Figure 8 offers an example of how well WRF can depict one of these monsters. Successes such as this are demonstrating that WRF is fulfilling its promise as the preeminent next-generation numerical weather prediction model.

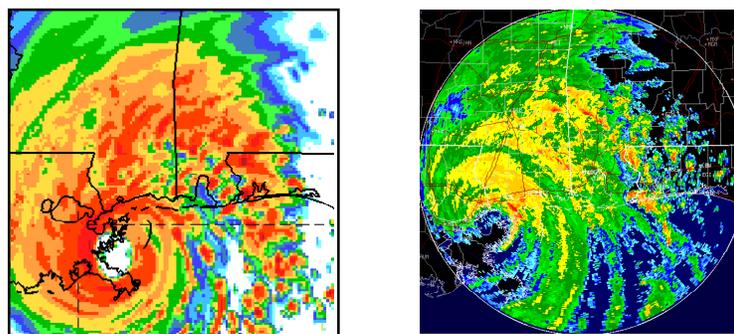


Figure. 8: WRF simulation of Hurricane Katrina computed three days before landfall (left), compared with later radar observations of the actual landfall (right).

3.4 THE EARTH SYSTEM MODELING FRAMEWORK

In another example of NCAR-supported community systems, the Earth System Modeling Framework (ESMF)⁸ provides high performance common modeling infrastructure for climate and weather models and is widely available as a community-owned and -managed product. It is in active use by groups working with hydrology, air quality, and space weather models. ESMF is the technical foundation for the NASA Modeling, Analysis, and Prediction (MAP) Climate Variability and Change program and the Department of Defense Battlespace Environments Institute (BEI). It has been incorporated into the CCSM, the WRF model, and many other applications.

The key concept that underlies ESMF is *software components*. Components are software units that are “composable,” meaning they can be combined to form coupled applications. These components may be representations of physical domains, such as atmospheres or oceans; processes within particular domains such as atmospheric radiation or chemistry; or

⁸ <http://www.esmf.ucar.edu>

computational functions, such as data assimilation or I/O. ESMF provides interfaces, an architecture, and tools for structuring components hierarchically to form complex, coupled modeling applications. ESMF components may be run sequentially, concurrently, or in a mixed mode on computers ranging from laptops to the world's largest supercomputers. The ESMF project encourages a new paradigm for geosciences modeling, one in which the community can draw from a federation of many interoperable components in order to create and deploy modeling applications. The goal is to enable a rich network of collaborations and a new generation of models that can simulate the Earth's environment and predict its behavior better than ever before.

ESMF is an open-source project that is actively reaching out to universities, national laboratories, industry, and the international community. ESMF is funded by a collection of agencies, and its development priorities and direction are set by multi-agency management bodies. Although the core development team is located at NCAR, the ESMF code has a growing number of contributors from collaborating sites. The project has been remarkably successful in its ability to bring disparate groups together, from the developer level all the way up to the agency level, and to get them working towards the common goal of better models.

Because of the success of CCSM, WRF, ESMF, and other similar community projects, NCAR is considering an overarching effort to develop an "Earth System Knowledge Environment." This environment would combine the key functions of all these programs and would lead to a fully supported and integrated workspace for modeling, computation, analysis, data management, data assimilation, and end-user diagnostics for the international community of geoscientists and societal decision makers charged with understanding the Earth system and its variability.

4. Summary

A strong emphasis on community involvement and governance has been critical to the success of NERSC and NCAR and is also central to plans for the future for both centers. NERSC and NCAR both support broad communities that are poised to make major breakthroughs in knowledge and understanding in very important scientific fields. Careful optimization of resources and capabilities will undoubtedly require continued attention and creativity as new computational systems develop and propagate. Both centers are ready to meet the challenge.

Acknowledgements

One of the authors (TLK) acknowledges important assistance from Al Kellie, Jordan Powers, Cecelia DeLuca, Marika Holland, Bill Collins, Jim Hack, and Veda Emmett in the development of this report.

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.