



Figure 1. Franklin, named in honor of America's first scientist Benjamin Franklin, is a 19,496-processor Cray XT4.

Science-Driven SUPERCOMPUTING

As DOE's flagship center for unclassified supercomputing, the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory is a world leader in accelerating scientific discovery through computation.

Established in 1974, NERSC provides computational resources to hundreds of research projects, including those funded under SciDAC. In 2006, SciDAC project teams were allocated 17,366,000 processor-hours on NERSC systems, accounting for 28% of the total production allocation by DOE. Among the achievements from these SciDAC allocations are:

- The SciDAC Advanced Computing for 21st Century Accelerator Science and Technology project used NERSC systems to model electron cloud instabilities that can disrupt the main proton accelerator beam, with the goal of improving the performance of existing and planned accelerators.
- Another accelerator modeling team ran simulations of plasma wakefield accelerator runs, which helped laboratory scientists better understand their experimental results, bringing them closer to the ultimate goal of building powerful

accelerators which are just a few meters—rather than kilometers—in length.

- Using algorithms developed under SciDAC and running simulations at NERSC and other centers, members of the Applied Differential Equations Center were able to create the first-ever 3D simulation of a laboratory-scale turbulent flame from first principles.
- Seeking insight into the best methods to inject frozen hydrogen fuel pellets into the 100 million-degree Celsius plasma of fusion reactors, members of the Center for Extended Magneto-hydrodynamic Modeling discovered that injecting pellets from the inside of the donut-shaped fusion chamber was much more effective than adding fuel from the outside of the ring.

Because of the high demand for computing resources from SciDAC investigators and other researchers, NERSC has developed a three-part strategy to ensure that its systems and services ensure the highest level of scientific productivity. Science-Driven Systems represent a balanced introduction of the best new technologies for complete computational systems—computing, storage, networking,

“While the theoretical peak speed of supercomputers may be good for bragging rights, the real measure of a machine's performance is the degree to which it advances scientific discovery. This is the primary driver behind NERSC's efforts to best meet the needs of our 2,500 users as they carry out research central to DOE's missions.”

DR. HORST SIMON
Director, NERSC Division,
LBNL




Figure 2. These images are isosurface contours of the accelerating electric field from a 3D particle-in-cell (PIC) simulation of a plasma wakefield accelerator. The simulation was done at NERSC and it used 14 million grid cells and 56 million particles.



Figure 3. Bassi, named for 18th century Italian professor Laura Bassi, is an IBM POWER 5 system with 888 processors. The system has a theoretical peak performance of 7.4 teraflop/s and 100 terabytes of disk space.



Figure 4. Seaborg, named in honor of LBNL Nobel Laureate Glenn Seaborg, is a 6,756-processor IBM supercomputer. Scheduled to be retired in 2008, Seaborg has a theoretical peak speed of 10 teraflop/s.

visualization, and analysis—coupled with the activities necessary to engage vendors in addressing the DOE computational science requirements in their future roadmaps. Science-Driven Services are 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 concentrates on resources needed to realize the

promise of the new highly scalable architectures for scientific discovery in multidisciplinary computational science projects. Science-Driven Analytics include 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 simulations and experiments.



Figure 5. The High Performance Storage System (HPSS) at NERSC.

This strategy has guided NERSC's procurement of large computing systems, in particular the newest machine, Franklin, a 19,496-processor Cray XT4 supercomputer (figure 1, p36). The system will deliver sustained performance of at least 16 trillion calculations per second—with a theoretical peak speed of more than 100 teraflop/s—when running a suite of diverse scientific applications at scale. The system will have over 400 terabytes of high-performance, parallel disk space. Initial installation began in 2006 and the full system is scheduled to go into production service in late summer 2007.

Other computing systems at NERSC are shown in figures 3–7 (p37, 38). Additionally, visualization and data analysis capability is provided by DaVinci, a 32-processor SGI Altix 350 server (figure 7).

To provide archival data storage for its 2,600 users, NERSC operates a High Performance Storage System (HPSS; figure 5) with a current capacity of 22 petabytes, or three times the volume of information in the Library of Congress. NERSC is one of five DOE research institutions which joined with IBM in 1992 to develop HPSS.

While theoretical peak speeds provide one indicator of a system's performance, a more realistic measure is how well the system performs when running scientific production codes. As part of its competitive procurement process, NERSC evaluates systems from a number of vendors using the Sustained System Performance (SSP) metric. The SSP metric, developed by NERSC, measures sustained performance on a set of codes designed to accurately represent the challenging computing environment at the center.

To further allow users to spend more time on their science and less time managing and accessing data, in 2005 NERSC deployed the NERSC Global Filesystem (NGF) into production, providing seamless data access from all of the center's computational and analysis resources. NGF facilitates users sharing data between machines and/or other users without doing explicit copies or making extra versions. For example, if a project has multiple users who must all access a common set of data files, NGF provides a common area for

Because of the high demand for computing resources from SciDAC investigators and other researchers, NERSC has developed a three-part strategy to ensure that its systems and services ensure the highest level of scientific productivity.



Figure 6. Jacquard is a 740-CPU Linux Network cluster running a Linux operating system. Theoretical peak speed for the system is 3.1 teraflop/s. The machine is named in honor of inventor Joseph Marie Jacquard, whose loom was the first programmable machine, using punch cards to control a sequence of operations.



Figure 7. DaVinci is named after the famous Italian artist and scientist Leonardo da Vinci, due to its intended combination of visual imagery of technical information.

those files. Alternatively, when sharing data between machines, NGF eliminates the need to copy large datasets from one machine to another. Because NGF has a single unified namespace for all systems, a user can run a highly parallel simulation on the 6,000-processor IBM SP (Seaborg), followed by a serial or modestly parallel post-processing step on the Linux Network cluster (Jacquard), and then perform a data analysis or visualization step on the SGI visualization cluster (DaVinci)—all without having to explicitly move a single data file.

To ensure that the systems and services meet users' needs, NERSC conducts an annual survey of its users, who often make comments and suggestions which are used to further improve services. Many also express their satisfaction, such as the user who recently noted, "Powerful and well maintained machines, great mass storage facility, and helpful and responsive staff. What more could you want?" ●

Contributor: Written by Jon Bashor, LBNL