NERSC NEWS

News from the National Energy Research Scientific Computing Center August 2005

NERSC Is First Production Site on ESnet's Bay Area MAN

On August 23, the NERSC Center became the first of six DOE research sites to go into full production on ESnet's new San Francisco Bay Area Metropolitan Area Network (MAN). Once completed, the new MAN will provide dual connectivity at 20 to 30 gigabits per seconds (10 to 50 times the current site bandwidths, depending on the site using the ring) while significantly reducing the overall cost.

The connection to NERSC consists of two 10-gigabit Ethernet links. One link will be used for production scientific computing traffic, while the second is dedicated to special networking needs, such as moving terabyte-scale datasets between research sites or transferring large datasets which are not TCP-friendly.

"What this means is that NERSC is now connected to ESnet at the same speed as ESnet's backbone network," said ESnet engineer Eli Dart.

(continued on page 3)

DOE Office of Science Director Raymond Orbach (left) and LBNL Director Steven Chu made the ceremonial connection between NERSC and ESnet in June. After testing, the full production connection was launched this month.

NERSC Launches Linux Networx Supercomputer into Production

On Monday, Aug. 1, NERSC announced that the new Linux Networx cluster system had entered full production mode for NERSC users. The new system has a total of 702 processors, including 640 processors for computing and 40 for storage.

NERSC's newest cluster, Jacquard, was opened to users on Aug. 1.

Named "Jacquard," the Linux Networx system will provide computational resources to scientists who run jobs on up to 124 processors, freeing up more resources on Seaborg for jobs which scale to 512 or more processors.

The Jacquard system is one of the largest production InfiniBand-based Linux cluster systems and has met rigorous acceptance criteria for performance, reliability and functionality that are unprecedented for an InfiniBand-based cluster. Jacquard is the first system to deploy Mellanox 12X InfiniBand uplinks in its fat-tree interconnect, reducing network hot spots and improving reliability by dramatically reducing the number of cables required.

The system has 640 AMD 2.2 GHz Dual Opteron™ processors devoted to computation, with the rest used for I/O, interac-

(continued on page 3)

Climate Model Developed on NERSC System Shows Faster CO_2 Emissions Will Overwhelm Capacity of Land and Ocean to Absorb Carbon

One in a new generation of computer climate models that include the effects of Earth's carbon cycle indicates there are limits to the planet's ability to absorb increased emissions of carbon dioxide.

If current production of carbon from fossil fuels continues unabated, by the end of the century the land and oceans will be less able to take up carbon than they are today, the model indicates, according to Inez Y. Fung, a professor at UC Berkeley, director of the Berkeley Atmospheric Sciences Center and NERSC user. She is also the lead author of a paper describing the climate model results that appeared in the Aug. 9, 2005 edition of the Proceedings of the National Academy of Sciences (PNAS). Fung developed the model on a NERSC supercomputer.

Fung was a member of the National Academy of Sciences panel on global climate change that issued a major report for President Bush in 2001 claiming, for the first time, that global warming exists and that humans are contributing to it.

(continued on page 2)

Climate Models

(continued from page 1)

"If we maintain our current course of fossil fuel emissions or accelerate our emissions, the land and oceans will not be able to slow the rise of carbon dioxide in the atmosphere the way they're doing now," said Fung. "It's all about rates. If the rate of fossil fuel emissions is too high, the carbon storage capacity of the land and oceans decreases and climate warming accelerates."

Currently, the land and oceans absorb about half of the carbon dioxide produced by human activity, most of it resulting from the burning of fossil fuels, Fung said. Some scientists have suggested that the land and oceans will continue to absorb more and more CO_2 as fossil fuel emissions increase, making plants flourish and the oceans bloom.

Fung's computer model, however, indicates that the "breathing biosphere" can absorb carbon only so fast. Beyond a certain point, the planet will not be able to keep up with carbon dioxide emissions.

"The reason is very simple," Fung said. "Plants are happy growing at a certain rate, and though they can accelerate to a certain extent with more CO_2, the rate is limited by metabolic reactions in the plant, by water and nutrient availability, et cetera."

While the modeling runs cited in the paper were run at the National Center for Atmospheric Research in Colorado, Fung notes that the code was developed on NERSC's Cray T3E supercomputer and cites NERSC in the paper.

A description of her research using NERSC can be found at <http://www.nersc.gov/news/annual_reports/annrep99/56sh_fung.html>.

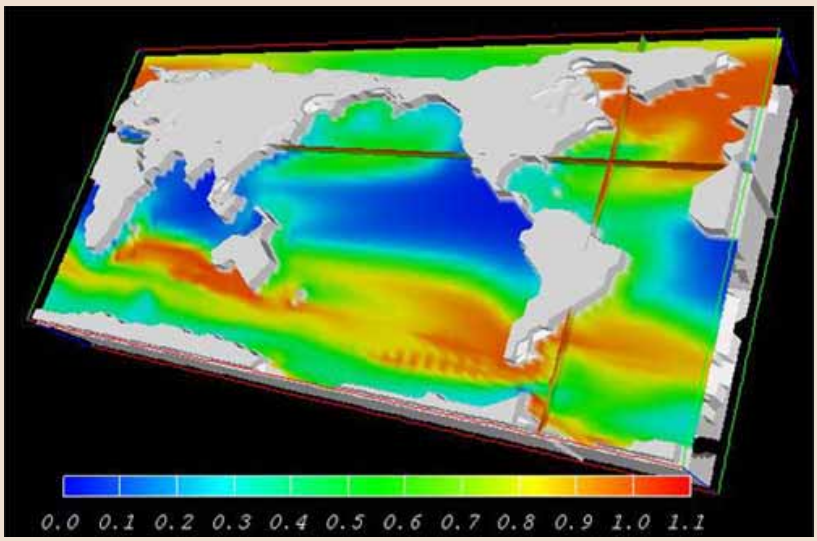

The three-dimensional distribution of a hypothetical inert surface tracer through time demonstrates the pathways this tracer uses to go from the surface into the deeper ocean. The model was run for 50 simulated years; this frame shows the values of the passive tracer at the 11th vertical level.

New Data Analysis, Mathematics and Visual Analytics Server Enters Production

In mid-August, NERSC put into production a new server specifically tailored to interactive visualization and data analysis work. The 32-processor SGI Altix, called DaVinci, offers interactive access to large amounts of large memory and high performance I/O capabilities typically required to analyze the large datasets produced by the NERSC high performance computing systems (Jacquard and Seaborg).

"With its 192 gigabytes of RAM and 25 terabytes of disk, DaVinci's balance is biased toward memory and I/O, which is different from the other systems at NERSC," said John Shalf of LBNL's Visualization Group. "This design gives us expanded capabilities for data analysis and analytics, along with interactive visualization."

DaVinci has 6 gigabytes of memory per processor, compared to 1 gigbyte per processor on Seaborg and 4 gigabytes on the new Linux Networx cluster, Jacquard.

Users can get interactive access to all 192 gigabytes of memory from a single application, whereas the interactive limits on production NERSC supercomputing systems restrict interactive tasks to a far smaller amount of memory.

With DaVinci now in production, NERSC will retire the previous visualization server, Escher, on Sept. 15. The new server will also replace Newton, the math server.

The new server will run a number of visualization, statistics and mathematics applications including IDL, Mathematica, Star-P (a parallel implementation of MatLab), AVS/Express, LLNL VisIT (a parallel visualization application), and CEI Ensight. Many users depend on IDL and MatLab to process or reorganize data in preparation for visualization. The large memory will particularly benefit these types of jobs, Shalf said.

While DaVinci will be available for interactive use by day, by night the system will be set up to run batch jobs, especially those jobs that are data intensive.

DaVinci is already connected to the HPSS and ESnet networks at NERSC by two independent 10 gigabit Ethernet connections and is expected to be integrated in the Facility Wide File Sharing System later this year.

NERSC News

NERSC News, highlighting achievements by staff and users of DOE's National Energy Research Scientific Computing Center, is published every other month via email and may be freely distributed. NERSC News is edited by Jon Bashor, JBashor@lbl.gov or 510-486-5849.

ESnet's Bay Area MAN

(continued from page 1)

The new architecture is designed to meet the increasing demand for network bandwidth and advanced network services as next-generation scientific instruments and supercomputers come on line. Through a contract with Qwest Communications, the San Francisco Bay Area MAN will provide dual connectivity to six DOE sites—the Stanford Linear Accelerator Center (SLAC), Lawrence Berkeley National Laboratory, the Joint Genome Institute (JGI), NERSC, Lawrence Livermore National Laboratory and Sandia National Laboratories/California. The MAN will also allow high-speed access to California's higher education network (CENIC), NASA's Ames Research Center and DOE's R&D network, Ultra Science Net.

All sites are scheduled to be connected by September. The Bay Area MAN will connect to both the existing ESnet production backbone and the first segments of the new Science Data Network backbone.

New Jacquard Cluster (continued from page 1)

tive work, testing and interconnect management. Jacquard has a peak performance of 3.1 trillion floating point operations per second (teraflop/s). Storage from DataDirect Networks provides 30 terabytes of globally available formatted storage.

Workers install the Jacquard cluster cabinets in NERSC's machine room. The system provides 640 processors for scientific computing.

The acceptance test included a 14-day availability test during which a select group of NERSC users were given full access to the Jacquard cluster to thoroughly test the entire system in production operation. Jacquard had a 99 percent availability uptime during the testing while users and scientists ran a variety of codes and jobs on the system. The thorough acceptance testing by NERSC ensures Jacquard is ready for a production environment for thousands of scientists and researchers across the nation.

The installation also includes a smaller development cluster called Jacdev. This system consists of 20 processors.

Following the tradition at NERSC, the system was named for someone who has had an impact on science and/or computing. In 1801, Joseph-Marie Jacquard invented the Jacquard loom, which was the first programmable machine. The Jacquard loom used punched cards and a control unit that allowed a skilled user to program detailed patterns on the loom.