NERSC NEWS

NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Berkeley Lab

Office of Science
U.S. DEPARTMENT OF ENERGY

SEPTEMBER 2008

Global Reach: NERSC Helps Manage and Analyze LHC Data

Over 15 million gigabytes of data per year will need to be stored, processed, backed up, and distributed to researchers across the world, when the Large Hadron Collider (LHC) begins smashing together beams of protons to search for new particles and forces, and beams of heavy nuclei to study new states of matter. Managing this mountain of data requires an international effort, with portions of the results sent to a variety of supercomputing centers across the globe.

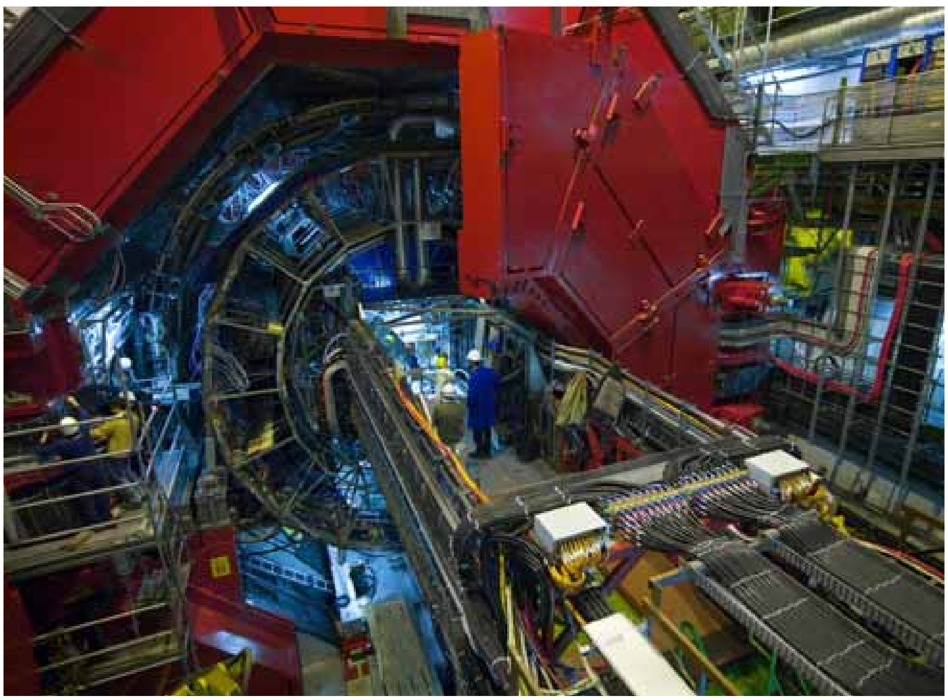

When One Door Closes, Another Opens: A magnet door on the LHC's ALICE Experiment closes, and an unprecedented opportunity to explore the realm of terascale physicis opens.

Large amounts of data from two of the LHC's detectors, ATLAS and ALICE, will be sent to the Department of Energy's National Energy Research Scientific Computing Center (NERSC), where two systems will be utilized to manage it. The High Performance Storage System (HPSS) will archive the raw and processed data. Meanwhile, the Parallel Distributed Systems Facility (PDSF) will process and distribute results to thousands of scientists across the globe.

"You can build a world class particle accelerator like the LHC, but if you don't have a way to store and analyze the data, it's all for nothing," says Peter Jacobs, of the Lawrence Berkeley National Laboratory's Nuclear Science Division, who contributed to the construction of the ALICE detector, one of the four large experiments at the LHC.

The LHC particle accelerator is located in Geneva, Switzerland, and is managed by the European Center for Nuclear Research (CERN).

PDSF Comes Full-Circle

When the LHC begins full operation next year, it will be the most powerful

continued on page 5

New Line of Security for NERSC Computers

The supercomputers at NERSC have a new line of defense against hackers with the installation of a specially instrumented version of SSH. This version of SSH allows NERSC's intrusion detection systems to analyze user activity while maintaining the security and privacy advantages of using SSH.

"In the past few years SSH has become a standard and is required for all users logging into NERSC systems," says Craig Lant, NERSC Security Analyst.

SSH is a security program that provides secure communication over insecure channels by encrypting keystrokes and system responses, including passwords. This means that users can remotely log onto computers over a network, execute commands, and transfer files from one system to another, with the utmost privacy.

"Because SSH encrypts commands that are typed into NERSC computers by remote users, it makes traditional intrusion detection less effective," says Lant.

As a precaution against hacking, scientists at NERSC began work modifying SSH to allow NERSC's adaptive intrusion detection system, called Bro, to monitor user activity and respond to unusual activity in real time. Bro will then alert security personnel if any suspected hacker activity is detected.

The new version of SSH was developed by NERSC scientists over the last year and will be released publicly toward the end of 2008.

NERSC Releases Software Test for Its Next Supercomputer

The Department of Energy's National Energy Research Scientific Computing Center (NERSC) is looking for a new supercomputer, but is not willing to spend millions of dollars on just any machine. The computer scientists and engineers want to know that their new supercomputer can reliably handle a diverse scientific workload, so they've developed the Sustained System Performance (SSP) Benchmarks, a comprehensive test for any system they consider.

The benchmarks were released in

continued on page 4

NERSC NEWSBYTES

Workshop on Risk Management Techniques and Practice for HPC Centers

On September 17 and 18, managers and key staff from high performance computing centers (HPCC) across the globe arrived in San Francisco, Calif. for a DOE workshop on risk management techniques and practice. Participants discussed current and emerging techniques, practices, and lessons learned for effectively identifying, understanding, managing, and mitigating risks associated with acquiring leading-edge computing systems at HPCCs.

"High-performance computing, by its very nature, is an exercise in multi-level risk management," says William Kramer, General Manager at NERSC and a member of the workshop organizing committee. "Every aspect of stewarding HPCCs into the petascale era, from identification of the program drivers to the details of procurement actions and simulation environment component deployments, represents unprecedented challenges and requires effective risk management."

Kramer notes that the fundamental purpose of this workshop was to go beyond risk management processes and learn how to weave effective risk management practices, techniques, and methods into all aspects of migrating HPCCs into the next generation of leadership computing systems and to develop risk portfolios and approaches that apply specifically to Petascale computing. This workshop was a follow-up of the Petascale System Integration workshop hosted by NERSC last year.

Bill Kramer and NERSC are also organizing a Birds of a Feather meeting on the topic of Petascale System deployment at the SC08 conference in November.

Ninth Annual ACTS Workshop Tackles "Building Robust, Scalable, and Portable Software"

Forty applications and tool developers from around the globe met at the Berkeley Lab's Oakland Scientific Facility, home of NERSC, on August 19–22, for the 9th annual Advanced CompuTational Software (ACTS) Collection workshop on "Building Robust, Scalable and Portable Software."

The four-day workshop introduced the ACTS Collection to scientists whose research demands include large amounts of computation, complex software integration, distributed computing, the use of robust numerical algorithms, or combinations of these. Participants were given 13 tutorials on a range of ACTS tools and other software developed by the Department of Energy's (DOE) SciDAC Program. Using NERSC's Jacquard computer, each student also received hands-on practice using ACTS tools to solve simple problems.

"This year, we had a terrific group. The workshop participants were bright, showed great interest in the tutorials and hands-on sessions, they worked diligently through hands-on sessions, despite the long days," says Tony Drummond of the Computational Research Division (CRD), who organized the workshop.

Visualization Team Develops Benchmark for Scientific Graphics Software

The NERSC Analytics Team has developed a benchmark for the high performance graphics industry standard OpenGL, called svPerfGL. This benchmark focuses on measuring OpenGL rendering performance in the presence of extremely heavy graphics payload with relatively few OpenGL state changes, which is a typical of a workload incurred by scientific visualization applications.

In contrast, industry standard benchmarks like SPECviewperf generate workloads typical of CAD and gaming applications, which do not always apply to science visualization. svPerfGL is written in a combination of C and C++ and uses calls consistent with the OpenGL 1.2 specification. Therefore, it is very conservative in its use of the OpenGL API and should run on any modern platform that supports OpenGL. It has undergone formal LBNL technology review and release and is currently available under an open source license. For more information and a link to download, please http://www-vis.lbl.gov/Research/svPerfGL/

Delegation from France's Strategic Council for HPC Visits LBNL

On Wednesday, August 27, a delegation from the French Strategic Council for High Performance Computing visited with Berkeley Lab Computing Sciences managers and staff to discuss research in scientific data management, algorithm development, computer architectures, collaborations with the University of California Berkeley, and other issues ranging from staffing to budgeting.

The council was established by the French Ministry of Research to advise the government on investments and research programs in supercomputing. Members of the delegation included Strategic Council Chair Olivier Pironneau, a professor at the University of Paris, and council members Michel Kern and Laura Grigori (formerly a researcher in LBNL's Scientific Computing Group).

This summer, the delegates also visited several American universities to learn about cutting-edge projects in software development for high performance computing.

NEW @ NERSC PROFILES

Shane Canon, Group Leader for Data Systems Group

Shane Canon "camping" at the Oak Ridge National Laboratory.

Shane Canon credits the cool, breezy climate and cosmopolitan cultural diversity of California's Bay Area for bringing him back to Berkeley and NERSC, after a three-year stint at the Oak Ridge National Laboratory in Tennessee. And, the new job was nice too. He will be leading the new Data Systems Group, which will be responsible for the NERSC Global Filesystem (NGF) and associated activities within the center.

"I like to tell people at NERSC that going to Oak Ridge was like going away to a three-year 'summer camp,' and now I'm coming home," says Canon, whose interest in supercomputers emerged when he managed computer clusters for the nuclear sciences laboratory at Duke University, in Durham, N.C., as a graduate student.

From 2000–2005, Canon served as Team Lead for the Parallel Distributed Systems Facility (PDSF) at NERSC. He then left for Oak Ridge to be closer to family in Spartanburg, South Carolina. However, his wife and kids missed the Bay Area, and the job opportunity at NERSC convinced him to move back to California. Canon is kept busy outside of working hours keeping up with soccer games for his daughters Rachel and Halle and keeping his two-year-old son Casey out of trouble.

Jen Jasper, Webmaster and Software Engineer

Jen Jasper strikes a pose near a giant termite mound in Australia. A wombat is nearby, but out of sight.

Jen Jasper's journey to NERSC began on a 60-acre Pennsylvania farm near the Susquehanna River, where she harbored a deep love for both art and nature. In college at Pennsylvania State University, she combined these interests by studying art, biology and teaching. She continued her education at Columbia University in New York, pursing a Master's in International Education, then spent the next few years at the City University of New York working as an international student advisor. When state budget cuts forced her to manage her own computer systems, Jasper discovered a new talent and hobby. Eventually, further funding cuts forced her westward to California.

During the dot-com era, Jasper found herself in the Silicon Valley developing and designing websites for WombatNet (former Internet service provider for Mac-users), Hotmail, Equinix, and Visto. As the dot-com boom turned to bust, she moved on to NASA Ames, where she did shell scripting and image processing to display remote field images from the Arctic on Future Flight Central's simulator screen. She also worked on sensor network and fault detection management for NASA's Mars Underground Mole (MUM), evaluated a web interface for the space shuttle's Integrated Vehicle Health Management system (IVHM), and designed a testbed and user interfaces for collecting and displaying remote sensing instrument data for a Mars analog site in Andalusia, Spain (MARTE). Just before arriving at NERSC, Jasper completed an Image Reduction and Analysis Facility (IRAF) data reduction and image processing pipeline for the astronomy department at City College of San Francisco.

David Stewart, High Performance Computer Network Engineer

David Steward cruises the Santa Fe Channel in Port Richmond, California.

David Stewart's fascination with computer networks and the Internet can be traced back to his freshman year at the University of California, Davis, when he was introduced to UNIX. "That was the hook for me," says Stewart.

He decided to go into computer networking, hoping to one day work with scientists and supercomputers. His goal was realized on April 15, 2008, when he started at NERSC. Before joining the Berkeley Lab, Stewart spend 15 years in industry and served as a senior network engineer for a variety of organizations, including BitTorrent Inc., Get Active Software, Robert Louis Stevenson School, Excite@Home Inc., and Navisite Inc.

A Bay Area native, Stewart spent most of his life in Oakland and Monterey, and has now settled in Emeryville. On his free time, Stewart enjoys eating sandwiches, taking black-and-white photos, bicycle riding, sailing, playing music, and general fiddling with electromechanical thingies.

Trever Nightingale, NERSC Systems Administrator

Trever Nightingale hikes though Pinnacles National Monument, near Salinas, California — home of the endangered California condor.

For Trever Nightingale, computers began as a hobby. In graduate school he focused on tackling the mind-brain problem — can science explain how a brain creates conscious experience? He eventually concluded that this was a philosophy problem that couldn't be solved by science.

"While that stuff was most interesting to me, and still is, it wouldn't pay the bills; there were no philosophy jobs in academia," says Nightingale. "Computers were a hobby but became a living

continued on page 3

NERSC Profiles: Trever Nightingale *continued from page 3*

because it made sense to do so."

For more than 10 years, Nightingale has worked as a full-time UNIX administrator. His career began at Minnesota Public Radio, where he contributed to the webcast for Garrison Keillor's Prairie Home Companion. Prior to starting at NERSC in February, he built many workgroup-sized Beowulf clusters at the Naval Research Laboratory in Monterey, Calif. Here, he also worked with high-speed interconnects such as Myrinet and Infiniband and gained some exposure to problems and solutions in parallel computation.

A Minnesota native, Nightingale now lives in Berkeley, Calif. On his free time, he enjoys playing with computers and hiking.

NERSC Releases Software Test for Its Next Supercomputer *continued from page 1*

conjunction with the NERSC-6 request for proposals on September 4. This version of SSP marks the first time that both vendors and the performance research community can easily access all applications and test cases. The NERSC-6 system will be the next major system acquisition to support the DOE Office of Science computational challenges.

"SSP is a key part of NERSC's comprehensive evaluation for large-scale systems," says William Kramer, General Manager at NERSC. "We look for balanced systems that expand our computational and analytics capability by assessing the systems' abilities to provide sustained performance, effective work dispatching, reliability, consistency and usability, for the entire range of the Office of Science computational challenges."

Instead of peak performance estimates, that is, the number of teraflop/s that could potentially be performed, NERSC scientists and engineers are concerned with the actual number of teraflop/s that the system will achieve in tackling a scientific problem. NERSC staff refer to this concept as sustained performance, and measure it using the SSP.

The SSP suite consists of seven applications and associated inputs, which span a wide range of science disciplines, algorithms, concurrencies and scaling methods. Kramer notes that this benchmark provides a fair way to compare systems that are introduced with different time frames and technologies. The test also provides a powerful method to assess sustained price/performance for the systems under consideration.

"NERSC-6 will be the system that provides the best value overall for supporting the DOE computational workload, taking into account Performance, Effectiveness, Reliability, Consistency and Usability, summed up in the acronym PERCU," says Kramer.

An updated version of the Effective System Performance (ESP) test, developed to encourage and assess improved job launching and resource management, was also released with the request for proposals, as are the other tests and benchmarks NERSC uses to assess large scale systems.

The new SSP suite can be downloaded from http://www.nersc.gov/projects/ssp.php

Specific information about the NERSC-6 request for proposals is located at http://www.nersc.gov/projects/procurements/NERSC6

Information about ESP can be found at http://www.nersc.gov/projects/esp.php

ASCAC Visits The Oakland Scientific Facility

NERSC General Manager Bill Kramer (right) describes NERSC's 100 teraflop/s Cray XT4 to Gordon Bell and Roscoe Giles (center) of the Advanced Scientific Computing Advisory Committee (ASCAC) during a tour held in conjunction with the August ASCAC meeting in Berkeley.

NERSC Helps Manage and Analyze LHC Data

continued from page 1

particle accelerator in the world. Capable of smashing together protons at an unprecedented 14 tera-electron volts of energy, it opens up a vast new landscape called the "terascale" for exploration. One exciting possibility is the discovery of the Higgs boson, a fundamental particle predicted to give mass, or weight, to all matter in the Universe; another is the discovery of a new family of "supersymmetric" particles, which are predicted to exist by theories that unify all the forces of nature. The LHC will also be the world's highest energy collider of heavy nuclei, generating matter under the extreme conditions that existed a few microseconds after the Big Bang.

Terrascale physics may be the new frontier, but the LHC is not the first collider designed to explore it. The U.S. had ambitions to investigate this realm of science in the 1980s, when its scientists began constructing the Superconducting Super Collider (SSC) in Waxahachie, Texas. In fact, the original PDSF was built to analyze SSC data. However, the project was cancelled in the mid-90s, and PDSF was transferred to NERSC, where it underwent multiple upgrades and expansions. Now it is one of the most flexible computing facilitates in the U.S., and has supported the majority of large nuclear and high energy physics projects undertaken by the country's leading scientists.

"PDSF is designed from the beginning to be able to support a wide range of nuclear science and high energy physics, from terascale phsics accelerators to experiments in the wastes of Antarctic to space experiments, in different ways. With the LHC, it comes full-circle to support ultra-large collider experiments again," says Jay Srinivasan, PDSF Systems Lead.

According to Jacobs, the PDSF architecture is ideal for processing high energy physics data because the different nodes in the cluster do not really need to communicate with each other. Each particle collision is taken as a single event, and only one node is required to process that event. "All we really need is a large set of processors that can access tremendous amounts of data," he says.

PDSF will also provide hundreds of terabytes of disk storage for the LHC experiments, allowing users to leverage NERSC's expertise in deploying high-performance parallel filesystems.

"At NERSC we will provide scientists the processing capabilities of PDSF, along with user support from consultants who are experts in computational science and performance tuning, visualization assistance, training, customized support, and other services," says Srinivasan. "And, because PDSF uses commodity technology, or hardware that is in the marketplace, we offer a very cost efficient service."

In addition to PDSF, data from the LHC will also be archived by HPSS, a system that has been used for archival storage since 1998. HPSS has over 6 petabytes of data stored in more than 70 million files.

"Both PDSF and HPSS have processed and stored data from experiments similar to the LHC, so we know what to expect," says Srinivasan. "However, we are all really excited to be a part of a major project that will advance our knowledge of the world around us."

PDSF is funded by the Computing Sciences, Physics, and Nuclear Sciences divisions at the Lawrence Berkeley National Laboratory.

ALICE is one of the four large LHC detectors, and is the only one optimized to study the very complex collisions of heavy nuclei. These collisions will occur inside the LHC for approximately four weeks every year, and will generate a fireball about 100,000 times hotter than the core of our Sun. Such temperatures have not existed in nature since a few millionths of a second after the Big Bang, the event that created our Universe. At these temperatures, scientists expect to see an extended fluid of quarks and gluons, which are typically confined inside of subatomic protons and neutrons. The fluid state is called the quark-gluon plasma.

A group of US institutions has joined ALICE to carry out this research. As part of this collaboration, the institutions are obligated to provide computing resources to the project in proportion to the fraction of authors on ALICE papers that they represent. PDSF is the main site for these US computing resources, and scientists suspect that by 2010 there will be about 400 processors at PDSF devoted to this experiment. ALICE is exploiting GRID capabilities extensively for its large-scale computing tasks. PDSF and LBNL are active participants in this development.

ATLAS is one of two general-purpose detectors that will investigate a wide range of physics, when particles collide inside the LHC, ATLAS will record the birth of new particles formed by the smashup — their paths, energies, and identities. ATLAS will play an important role in the searches for the Higgs boson, a fundamental particle predicted to give mass, or weight, to all matter in the Universe; and a new family of "supersymmetric" particles, which are predicted to exist by theories that unify all the forces of nature.

In order to efficiently plan and execute experiments, scientists must fully understand ATLAS detector capabilities before-hand. They will achieve this by running the Monte Carlo simulations of detectors on PDSF. LBNL has been designated as an Analysis Support Center for U.S. ATLAS and will host a series of workshops to help physicists in the U.S. prepare for early physics.