HOME

Organization Chart
Jobs and Fellowships
Berkeley CSE
Safety and Health
CS Newsletters

NERSC

News & Publications
Image Gallery
Annual Reports

CRD computational research division

News & Publications
Visualization Gallery

ESnet

ESnet News
Publications

CS Staff Only

# Bandwidth Challenge Highlights DOE Science Services

**November 6, 2009**



**Contacts: Linda Vu, lvu@lbl.gov, 510-495-2402**

It typically takes about two days to move 10 terabytes of climate data between U.S. Department of Energy (DOE) computing facilities. But during SC09, a collaboration of researchers and engineers from Argonne, Lawrence Berkeley and Lawrence Livermore National Laboratories will attempt to transfer more data than this in approximately two hours. The occasion is the annual SC09 Bandwidth Challenge.

"As the climate community prepares for the Intergovernmental Panel on Climate Change Fifth Assessment Report (AR5), which will be published in 2013, the ability to quickly and reliably move terabytes of data across the world is crucial," said Dean Williams, a computer scientist at the Lawrence Livermore National Laboratory's (LLNL) Program for Climate Model Diagnosis and Intercomparison (PCMDI). "We hope that the lessons learned in this experiment will help us to do this better."

The data will be streaming into the SC09 showroom in Portland, Ore., from three DOE sites — the Argonne Leadership Computing Facility (ALCF), the National Energy Research Scientific Computing Center (NERSC) at the Lawrence Berkeley National Laboratory (Berkeley Lab) and LLNL — carried by DOE's Energy Sciences Network (ESnet).

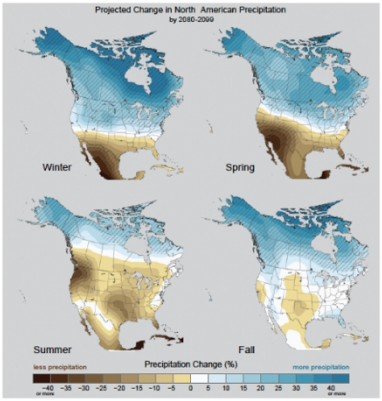


**Precipitation map developed by Berkeley Lab's Michael Wehner shows, among other things, a substantial reduction in springtime rains in California, and summertime rains in the Pacific Northwest.**

Once the data arrives at its destination in the University of Utah's SC09 booth, it will be stored on disks provided by DataDirect Networks Inc., processed in real time with Climate Data Analysis Tools and the Visualization Streams for Ultimate Scalability (ViSUS) visualization tool, and then publicly displayed along with graphs depicting the demo's transfer rates. Researchers from DOE's Visualization and Analytics Center for Enabling Technologies (VACET), who contributed to the development of ViSUS, will be on hand to help with visualizing the climate data.

Information about the transfer throughput will be captured regularly with monitoring software called NetLogger, which will be running at the three DOE sites and in the University of Utah's booth. NetLogger was developed and is maintained by collaborators in DOE's Center for Enabling Distributed Petascale Science (CEDPS), who are providing the key middleware for this bandwidth competition.

## Why Climate?

The idea to transfer climate data as part of this challenge came from researchers at the Argonne National Laboratory (ANL), who are leading this multi-lab effort. The data comes from the World Climate Research Program Coupled Model Intercomparison Project, Phase 3 which currently stores more than 35 terabytes of data at LLNL. The data is accessed by more than 3,000 users worldwide via the Earth System Grid (ESG), the leading infrastructure for accessing and distributing climate model data. ESG is funded through DOE's Scientific Discovery through Advanced Computing (SciDAC) program, with support from the offices of Advanced Scientific Computing Research and Biological and Environmental Research.

"Climate is a discipline that is highly collaborative, and its datasets are distributed across the globe. Moreover, the climate community practices bulk transfers, so it benefits from the tools that we wanted to highlight in this challenge," said Rajkumar Kettimuthu, a staff scientist at ANL and technical lead of this bandwidth challenge entry.

"An interesting feature of climate data is that the actual file size is not very large compared to that of other sciences. Climate researchers, however, need to move hundreds or thousands of files in a single transfer, so the volume of data that is moving across the network is massive. This is what we mean by bulk transfers," explained Alex Sim, a member of the Berkeley Lab's Scientific Data Management Group. Sim leads a project called Climate 100, which seeks to help the research community effectively use emerging 100-gigabit-per-second networks.

## The Right Tools and Network

According to Kettimuthu, the transfers typically initiated by the climate community can be between a client and a server or between two remote servers initiated by the user from a third machine. He explained that GridFTP is an ideal tool for these types of transfers because it offers high bandwidth, built-in security and a host of other functionalities. Unfortunately, most of the existing GridFTP client tools require the users to closely oversee the transfer or are difficult to operate. These difficulties inspired Argonne researchers to create an online hosted service, called Globus.org, for reliable, high-performance, fire-and-forget data transfers.

"With this hosted data-movement client, users can easily initiate a third-party data movement without the need to install, run, and monitor anything on their own systems. They get an e-mail when their transfer is completed," said Kettimuthu. In fact, software from Globus.org was used to replicate a portion of ESG data from NERSC to ALCF.

For the challenge, the team will be using GridFTP with all the new features developed by the CEDPS and Bulk Data Mover software developed by the DOE Earth System Grid-Center for Enabling Technology (ESG-CET) as the GridFTP client to initiate data transfers from the three DOE sites. The data will then flow across ESnet's circuit-oriented Science Data Network (SDN), which is optimized to handle massive scientific datasets, to the convention center. To ensure that the data arrives within the two-hour challenge timeframe, the team reserved bandwidth on the SDN with ESnet's On-Demand Secure Circuit and Advance Reservation System (OSCARS).

"This challenge is not just about having the right tools and applications. We needed a good network that would reliably move these datasets at high speeds to the convention center. ESnet was the perfect fit," said Kettimuthu. "The combination of SDN and OSCARS guarantees that we will have a dedicated circuit on the network for the duration of the challenge and don't have to compete with anyone else for bandwidth."

The wide-area network infrastructure for this challenge will be based on a system of dedicated data transfer nodes deployed at NERSC and the ALCF. The nodes, which are optimized for such transfers, have been moving large datasets between the DOE computing facilities at a rate of 200 megabytes per second. The infrastructure was developed and is being maintained by a collaboration of engineers from ALCF, ESnet, NERSC and the Oak Ridge National Laboratory Leadership Computing Facility. CEDPS, ESG-CET and VACET are all supported by DOE's SciDAC Program.

## About NERSC, ESnet and Berkeley Lab

The **National Energy Research Scientific Computing Center (NERSC)** is the primary high-performance computing facility for scientific research sponsored by the U.S. Department of Energy's Office of Science. The NERSC Center currently serves thousands of scientists at national laboratories and universities across the country, researching problems in combustion, climate modeling, fusion energy, materials science, physics, chemistry, computational biology, and other disciplines.

**ESnet,** formally known as the Energy Sciences Network, is already one of the world's most advanced research networks. Primarily supported by DOE's Office of Science, ESnet provides high-bandwidth network connections to more than 40 sites conducting DOE-funded research, including some 20 large-scale experimental facilities and large supercomputing centers used by thousands of DOE scientists generating massive amounts of data.

Both NERSC and ESnet are managed by **Berkeley Lab**, which is a U.S. Department of Energy national laboratory located in Berkeley, California. It conducts unclassified scientific research and is managed by the University of California for the DOE Office of Science. For more information about computing sciences at Berkeley Lab, please visit: **www.lbl.gov/cs**

## About Argonne and the ALCF

Argonne National Laboratory seeks solutions to pressing national problems in science and technology. The nation's first national laboratory, Argonne conducts leading-edge basic and applied scientific research in virtually every scientific discipline. Argonne operates the ALCF for the DOE Office of Science as part of the larger DOE Leadership Computing Facility strategy. DOE leads the world in providing the most capable civilian supercomputers for science. Argonne researchers work closely with researchers from hundreds of companies, universities, and federal, state and municipal agencies to help them solve their specific problems, advance America's scientific leadership and prepare the nation for a better future. With employees from more than 60 nations, Argonne is managed by UChicago

Argonne, LLC for the U.S. Department of Energy's Office of Science.

## About the Lawrence Livermore National Laboratory

Founded in 1952, Lawrence Livermore National Laboratory is a national security laboratory, with a mission to ensure national security and apply science and technology to the important issues of our time. Lawrence Livermore National Laboratory is managed by Lawrence Livermore National Security, LLC for the U.S. Department of Energy's National Nuclear Security Administration.

## About SC09

SC09, sponsored by ACM (Association for Computing Machinery) and the IEEE Computer Society, offers a complete technical education program and exhibition to showcase the many ways high performance computing, networking, storage and analysis lead to advances in scientific discovery, research, education and commerce. This premier international conference includes a globally attended technical program, workshops, tutorials, an exhibit area, demonstrations and hands-on learning. The SC conference series is among Tradeshow Week magazine's Top 200 events. For more information on SC09, please visit **http://sc09.supercomputing.org/**.

---