


# Department of Energy's ESnet and NERSC blaze 400G production network path

November 10 2015



This map shows the 400G production link between Wang Hall at the main Berkeley Lab site and the Oakland Scientific Facility. Credit: Berkeley Lab This map shows the 400G production link between Wang Hall at the main LBNL site and the Oakland Scientific Facility.

The Department of Energy's Energy Sciences Network (ESnet) and the National Energy Research Scientific Computing Center (NERSC) have built a 400 gigabit-per-second (Gbps) super-channel, the first-ever 400G production link to be deployed by a national research and education network.

The connection, nicknamed the BayExpress, will provide critical support for NERSC's 6,000 users as the facility moves from its current location in Oakland, Calif. to the main campus of Lawrence Berkeley National Laboratory in Berkeley over the next year. Both ESnet and NERSC are DOE Office of Science User Facilities and have a track record of deploying leading-edge technologies to advance scientific research.

The project is important to the mission of the Office of Science, the single largest supporter of basic research in the physical sciences in the United States. That research is increasingly data-driven as scientists generate massive datasets at experimental facilities and through modeling and simulation at supercomputing centers like NERSC. ESnet provides the high-speed network connectivity for moving that data, enabling the scientists to move, analyze and share the data effectively. ESnet currently operates a 100G backbone network connecting DOE sites and universities, but the growth in data is threatening to oversaturate that capacity, especially as supercomputers approach exascale capabilities.

To develop the 400 Gbps production connection as a step toward the next level of networking, ESnet and NERSC joined with Ciena, a global supplier of telecommunications networking equipment, software and services; and Level 3 Communications, a telecommunications and Internet service provider. The production link is deployed as a pair of 200 Gbps per wavelength and allows for a simple-to-deploy dual subcarrier 400 Gbps solution.

As part of the project, the team also set up a 400 Gbps research testbed for assessing new tools and technologies without interfering with production data traffic and allowing staff to gain experience operating a 400 Gbps optical infrastructure. For the testbed, the optical frequencies are more narrowly tuned to maximize use of the fiber spectrum, known as "gridless superchannels," said Chris Tracy, co-leader of the project

and a member of ESnet's Network Engineering team. The testbed uses a dark fiber link provided by Level 3 to ESnet and NERSC for six months. The team will begin conducting field trials on the testbed within the next two months.

"We had two goals for this project," said Tracy. "We wanted to demonstrate a real-world production deployment of 400 Gbps optical connection linking two distinct, large-scale data centers and to integrate scientific workflows, which will make efficient use of this bandwidth. We also wanted to conduct research in the areas of fiber-optic signal transport, data movement and management techniques to support the execution of those workflows, which are important for DOE's research mission."

The team did advance work in the lab with spools of fiber and became familiar with the technology before deploying it in the field. The first goal was met by deploying a production-quality 400 Gbps connection over the 11 km stretch between the Oakland Scientific Facility where NERSC is located and the Berkeley Lab campus in the hills overlooking UC Berkeley. This meant adding more capabilities, then more testing to get the connection up to speed and reliable enough for production use. Then the team moved to begin testing and make the system work over longer distances.

As part of the longer distance testing, the team experimented on a 450 km long link that is a reverse loop connecting Berkeley Lab to Oakland via Sacramento and Sunnyvale. The team gained significant experience as they strove to achieve 400G performance over the longer segment. An additional testbed research goal is to use spooled fiber to extend the 11 km link to 91 km to conduct further optical networking research and test next-generation network architectures.

According to Jason Lee, the networking deputy for NERSC, the project

was the first of its kind, mainly because the necessary equipment is in short supply. The key is having the 400 Gbps network interface cards, which Ciena first demonstrated at the SC13 conference with ESnet and SCinet support two years ago. At the time, only a limited number of the cards had been produced.

The 400 Gbps link will allow NERSC to continue providing resources to users as systems are temporarily taken out of service so they can be moved into the computing center in recently completed Wang Hall at Berkeley Lab. During two previous moves in 1996 and 2000, the installation of new system by NERSCs and connectivity provided by ESnet ensured that users had access to both computing and data storage systems even while equipment was temporarily taken out of service for relocation.

In the current situation, the NERSC facility in Wang Hall ingests data from the computing and [data storage systems](#) in Oakland over four 100 Gigabit Ethernet router links, which are then combined to act as a single 400 Gbps optical transport connection between the two sites.

## **A History of Strong Connections**

The user facilities known today as NERSC and ESnet both date back to the 1970s, when a direct link was set up to allow fusion researchers at the Princeton Plasma Physics Lab in New Jersey to access a high-performance computer at Lawrence Livermore National Laboratory in California. As more powerful computers were installed at the center and the number of users and institutions accessing the systems grew, a full-fledged network was created. Since 1990, the amount of traffic carried by ESnet has doubled about every 18 months and networks staff have regularly rolled out new technologies and greater bandwidth to stay ahead of the growth curve.

During the same time, NERSC has become a "net importer" of data as the center has evolved into a central data repository for projects ranging from cosmology to genomics to materials, in addition to generating massive datasets from scientific modeling and simulation. And with the nation's largest user community using its computing and data resources 24/7, NERSC relies on ESnet for uninterrupted connectivity.

"We expect the 400G prototype network to provide us with many valuable insights with regard to future network architectures and be useful to others in the research community," said ESnet Chief Technologist Inder Monga. "Other universities, national laboratories and regional and commercial networks will be planning their next-generation networks on this time scale and our goal is to publicize not only our experiences from an optical networking perspective, but also from a scientific workflows perspective in which large-scale datasets are moved between facilities coupled by 400 Gbps of dedicated network bandwidth."

Provided by Lawrence Berkeley National Laboratory

Citation: Department of Energy's ESnet and NERSC blaze 400G production network path (2015, November 10) retrieved 18 April 2024 from <https://phys.org/news/2015-11-department-energy-esnet-nersc-blaze.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.