

# Using the Network to Optimize Data Transfers for Science

Presented by ESnet / Berkeley Lab

## Introduction

Scientists move hundreds of petabytes of data every year to and from collaborators, institutions, and storage and analysis facilities around the world. As an example, the LHC ATLAS experiment transfers data averaging just under 25 TB/day in the United States, as shown in Figure 1. To successfully enable data transfers of this magnitude and scale, computer networks need to be specifically engineered for high-performance, data-intensive science from end to end.

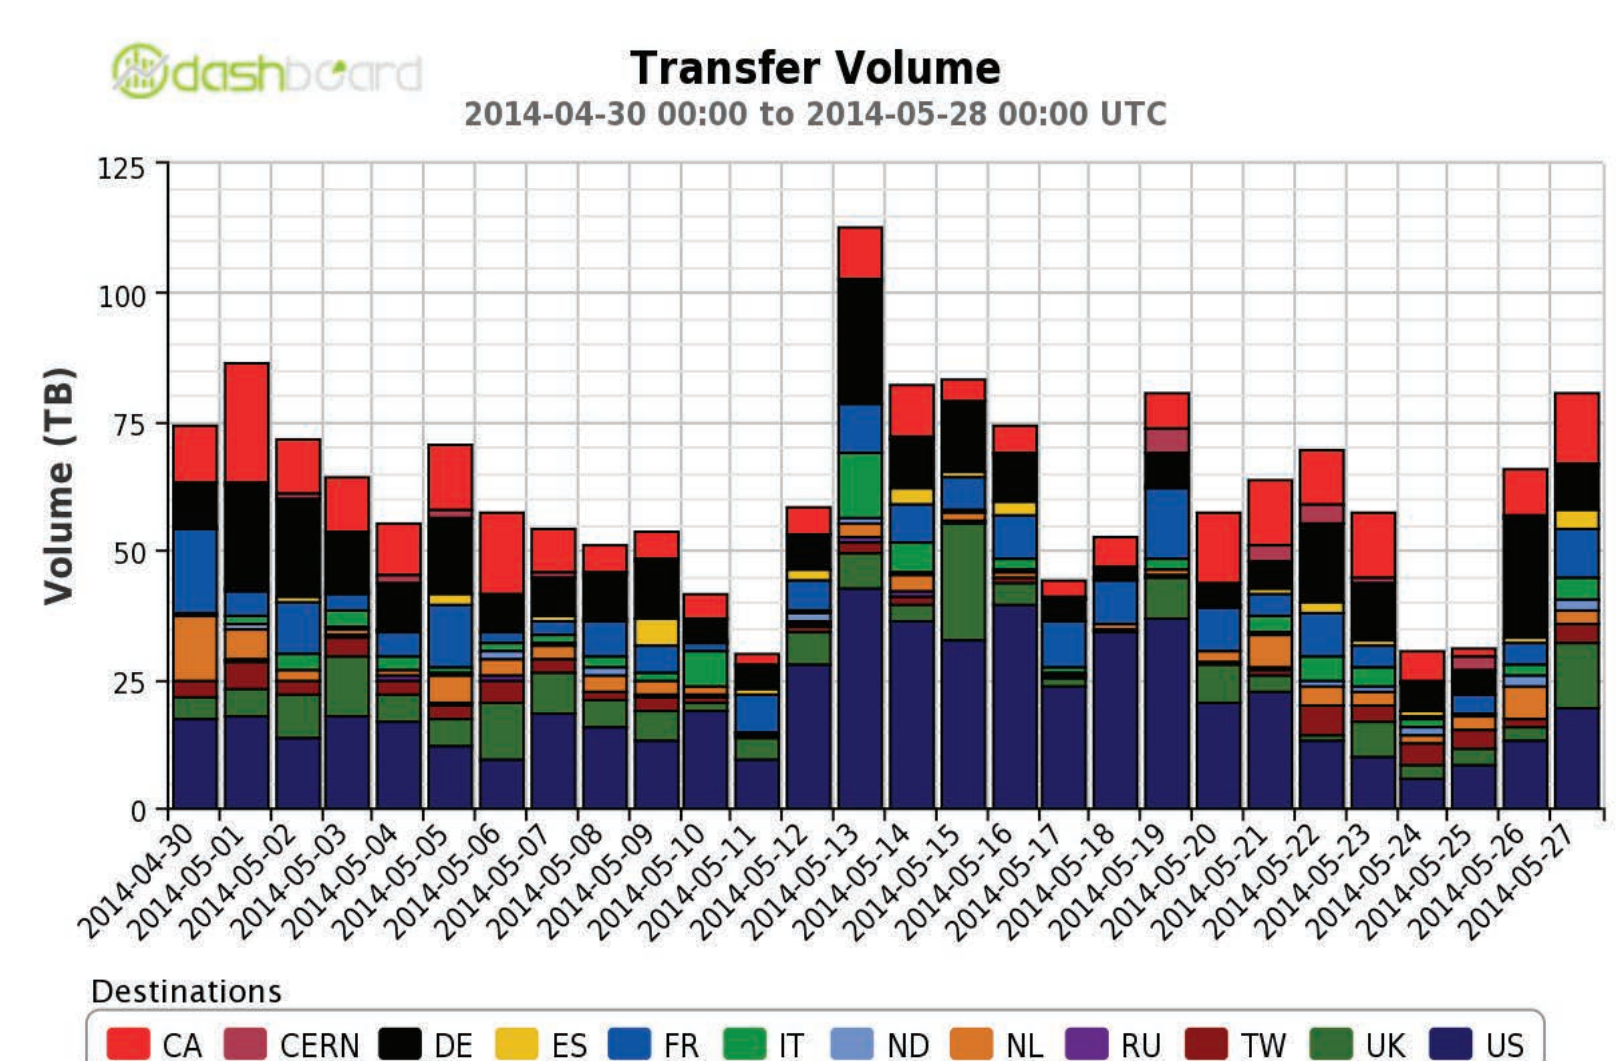


Figure 1: Data transfers from just the ATLAS experiment over a 28-day period. The data transfer volumes are in terabytes vs. days. Note the US data transfers are marked in dark blue and averaging just under 25 TB/day.

## SETTING NETWORK EXPECTATIONS

For benchmarking, the following table shows approximately how long it should take to transfer 1 TB of data across varying network speeds:

10 Mb/sec network	300 hrs (12.5 days)
100 Mb/sec network	30 hrs
1 Gb/sec network	3 hrs
10 Gb/sec network	20 mins

These describe the throughput for data transfers based on end-to-end network bandwidth noted in megabits and gigabits per second (Mb/sec and Gb/sec).

(See <http://fasterdata.es.net/fasterdata-home/requirements-and-expectations/> for more information.)

## USING THE RIGHT TOOLS FOR TRANSFERRING DATASETS

It is important to use the right tools and protocols for data transfers. Different protocols behave differently on the networks and some protocols perform better than others. Some good protocols and tools for science datasets are:

- Globus: <https://www.globus.org>
- GridFTP from ANL has features needed to fill the network pipe
- bbcp: <http://www.slac.stanford.edu/~abh/bbcp/>
- lftp: <http://lftp.yar.ru/>
- axel: <http://axel.alioth.debian.org/>
- FDT (Fast Data Transfer) tool from Caltech: <http://monalisa.cern.ch/FDT/>

## WHEN DATA TRANSFERS DON'T WORK WELL...

There are many causes of poor network performance, but one of the most common causes is packet loss. Packet loss causes the most common network protocol, TCP, to adjust its rate of data transmission creating poor network performance. The graph in Figure 2 shows the large performance impact on data transfers with even a very small amount of packet loss.




Figure 2: Graph that displays the behavior of network throughput when it is dropping 1 in 22,000 packets. With .0046% packet loss, data that traverses network paths more than 10 ms away will see extremely low data transfer performance.

## THE NETWORK AS A CIRCULATORY SYSTEM FOR SCIENCE

Serving as the primary "circulatory system" for the DOE Office of Science, ESnet enables tens of thousands of scientists to transfer data, access remote resources, and collaborate productively around the United States.

In one recent experiment, biosciences researchers at SLAC National Lab's Linac Coherent Light Source (LCLS) ran an experiment that collected 113.6 TB of data over 5 days. As this data was generated at the light source, it traveled over 40 miles of ESnet's 100 Gbps network to the National Energy Research Scientific Computing Center (NERSC) for data processing in near real-time. With ESnet and NERSC expediting the data from the LCLS experiment, scientists increased the quality of the overall experiment and their science results. This experiment's massive data analysis relied on speed and scale provided by ESnet's high-bandwidth networking and NERSC's support of HTC (high throughput computing) batch queuing.




Figure 3: Researchers leveraged LCLS's beamline to investigate the different photoexcited states of photosystem II. By analyzing several tens of thousands of LCLS beamline diffraction images (a), scientists were able to reconstruct the electron density around the atomic centers of the crystallized molecule (b) using the NERSC supercomputing center.




Figure 4: Graphs show the data transferred from the beamline's detector to NERSC's data transfer nodes at over 7.5 Gbps, where the data was then processed on thousands of cores. (Graphs are displayed via MyESnet Portal.)

Kern, Jan, et al. "Taking snapshots of photosynthetic water oxidation using femtosecond X-ray diffraction and spectroscopy." *Nature Communications* 5 (2014).

## About ESnet

ESnet is the Department of Energy's high-performance, high-bandwidth network that provides reliable network connections for scientists at national laboratories, universities and other research institutions, enabling them to collaborate on some of the world's most important scientific challenges including energy, climate science, and the origins of the universe. ESnet is funded by DOE Office of Science and located within the Scientific Networking Division at Lawrence Berkeley National Laboratory.

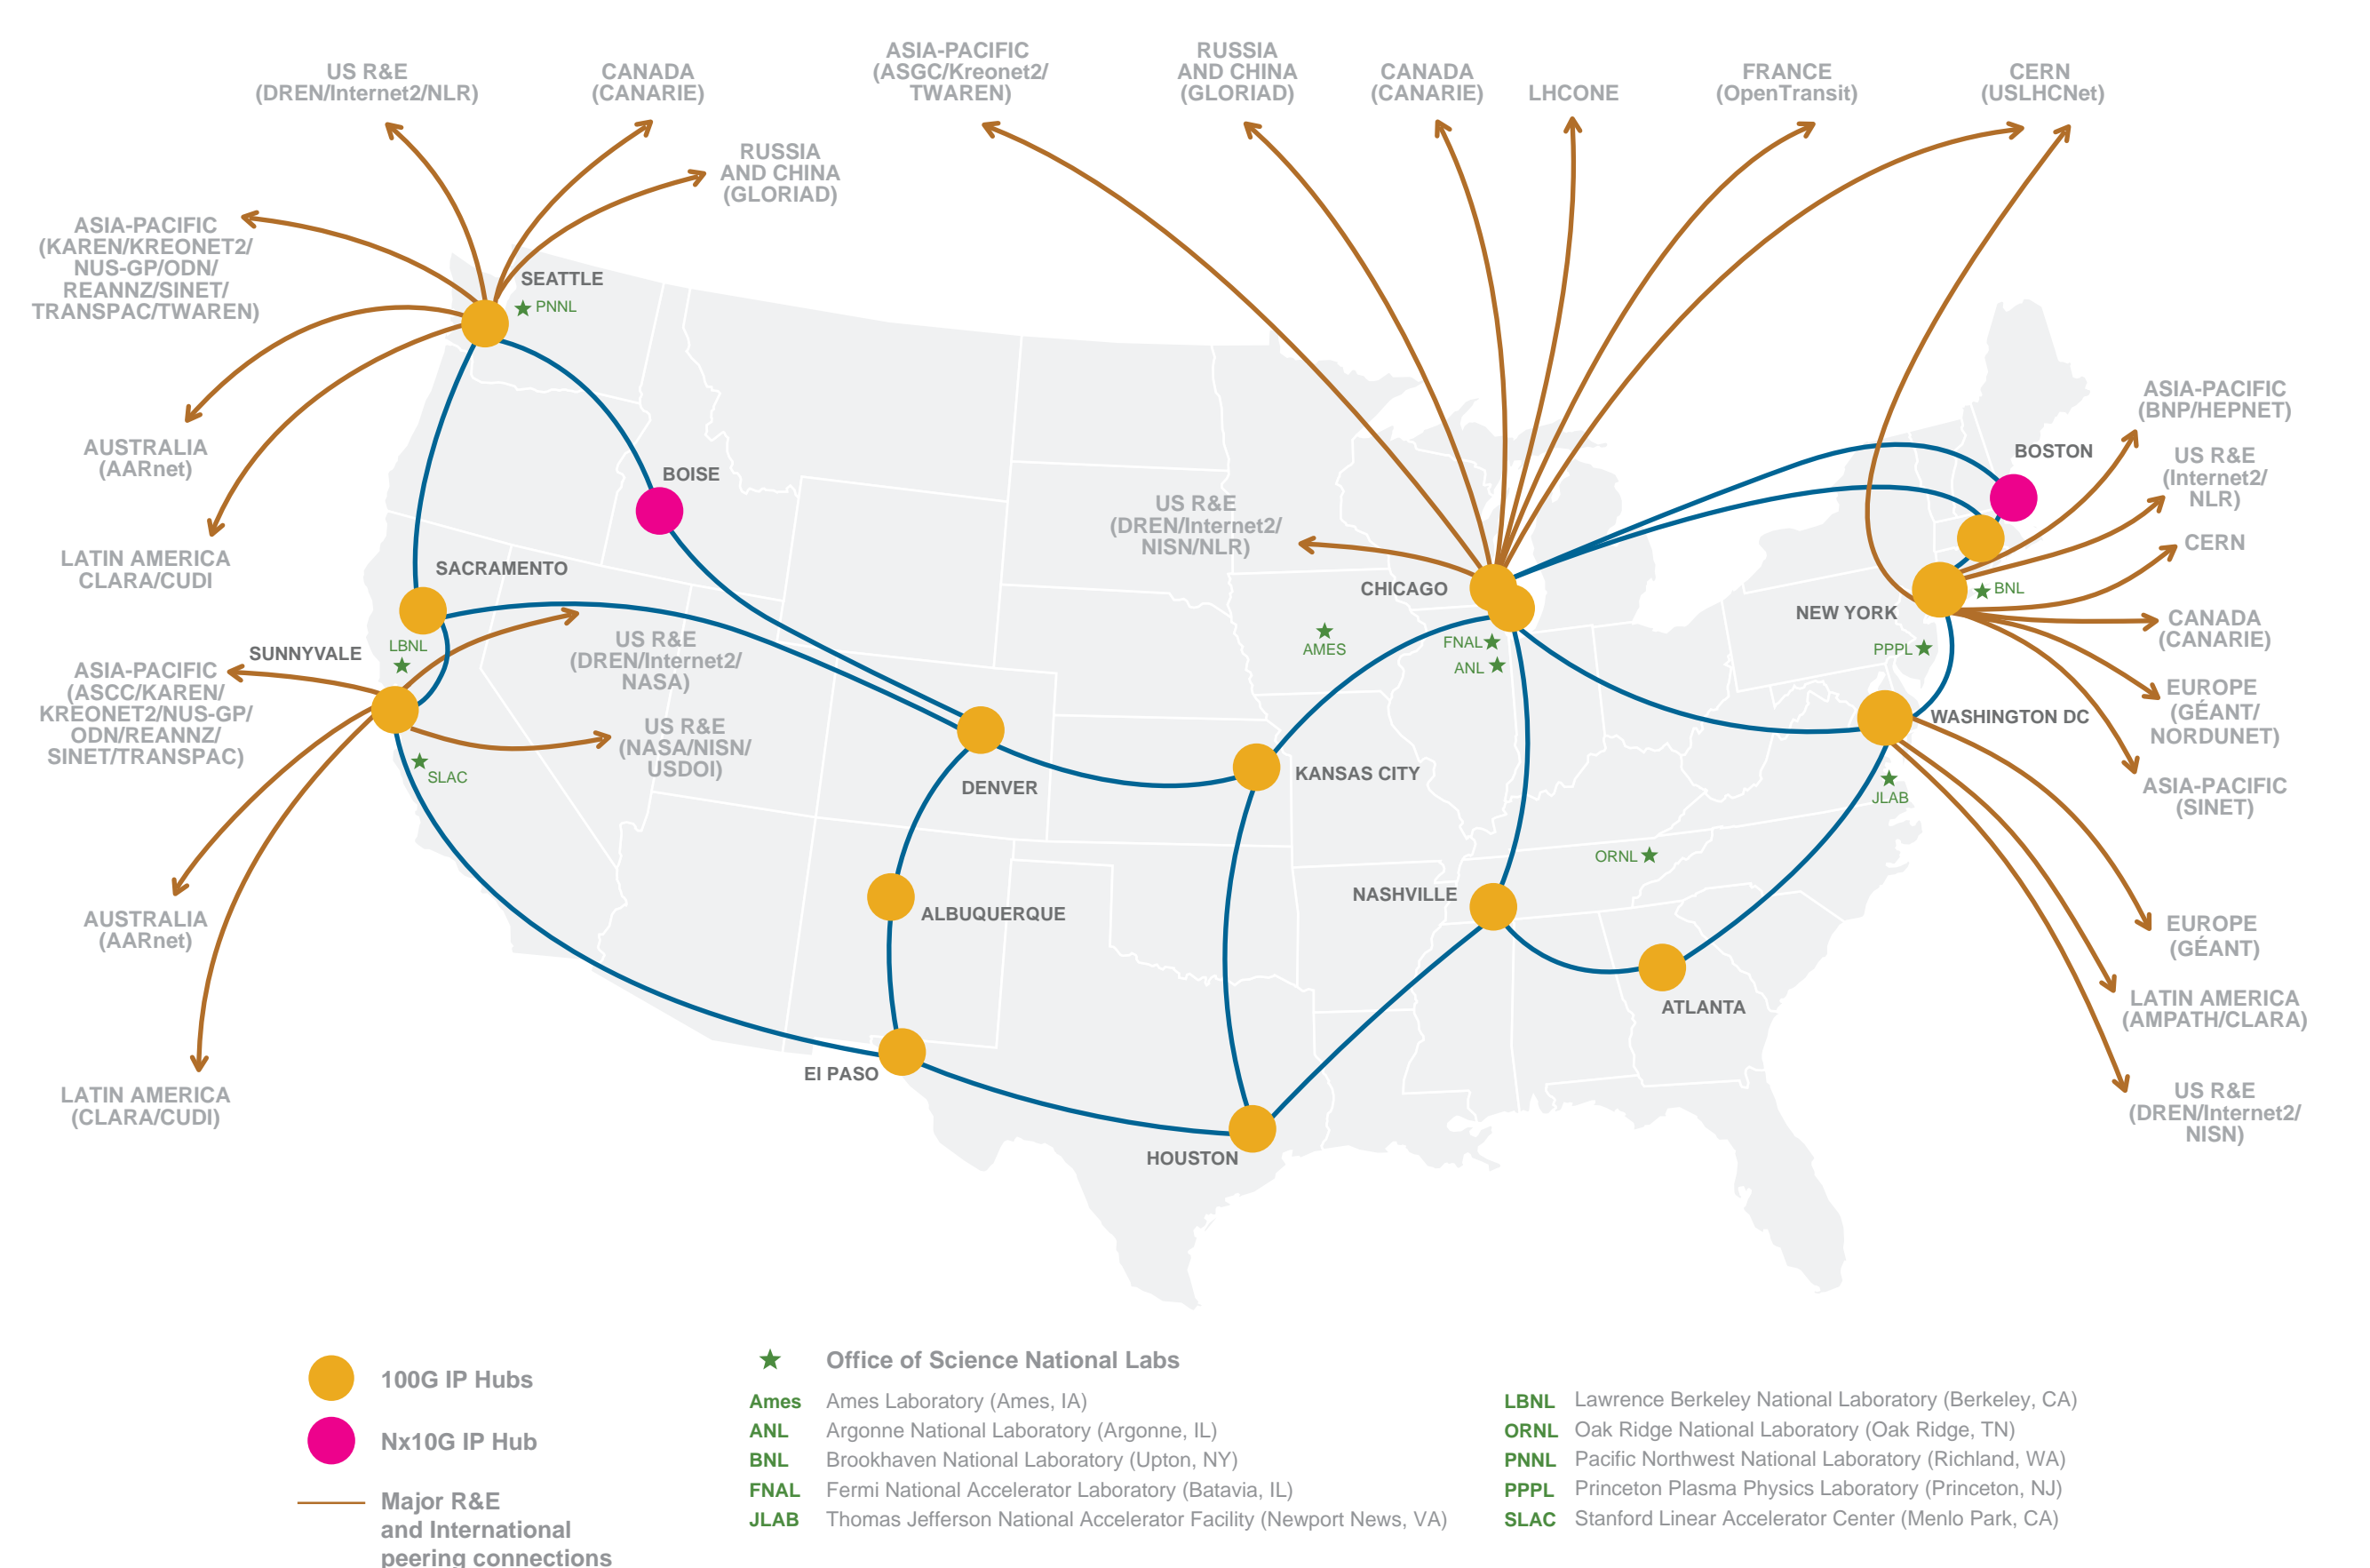


Figure 5: A high-level topology map showing ESnet's backbone network. ESnet serves more than 27,000 postdoctorals, graduate students, undergraduate students, engineers and technicians; 26,000 users of open-access facilities; 300 leading academic institutions; and 17 DOE laboratories.

## Contacts and More Information

Call your local support or IT team if you are seeing any problems!

For other questions or more help, email: [engage@es.net](mailto:engage@es.net)

More information on transfer protocols and best practices visit Fasterdata Knowledge Base: [fasterdata.es.net](http://fasterdata.es.net)

