

# UC BERKELEY SCIENCE DMZ & 100 Gb/s BORDER

Enabling extreme data for research and education.

## What Is a Science DMZ?

Consists of **three key components**, all required:

- “Friction Free” network path
  - Highly capable network devices (wire speed, deep queues)
  - Virtual circuit connectivity option
  - Security policy and enforcement specific to science workflows
  - Located at or near site perimeter if possible
- Dedicated, high-performance Data Transfer Nodes (DTNs)
  - Hardware, operating system, libraries all optimized for transfer
  - Includes optimized data transfer tools such as Globus Online and GridFTP
- Performance measurement/test node
  - perfSONAR

Further details are available at <http://fasterdata.es.net/science-dmz>

## Campus Network Improvements

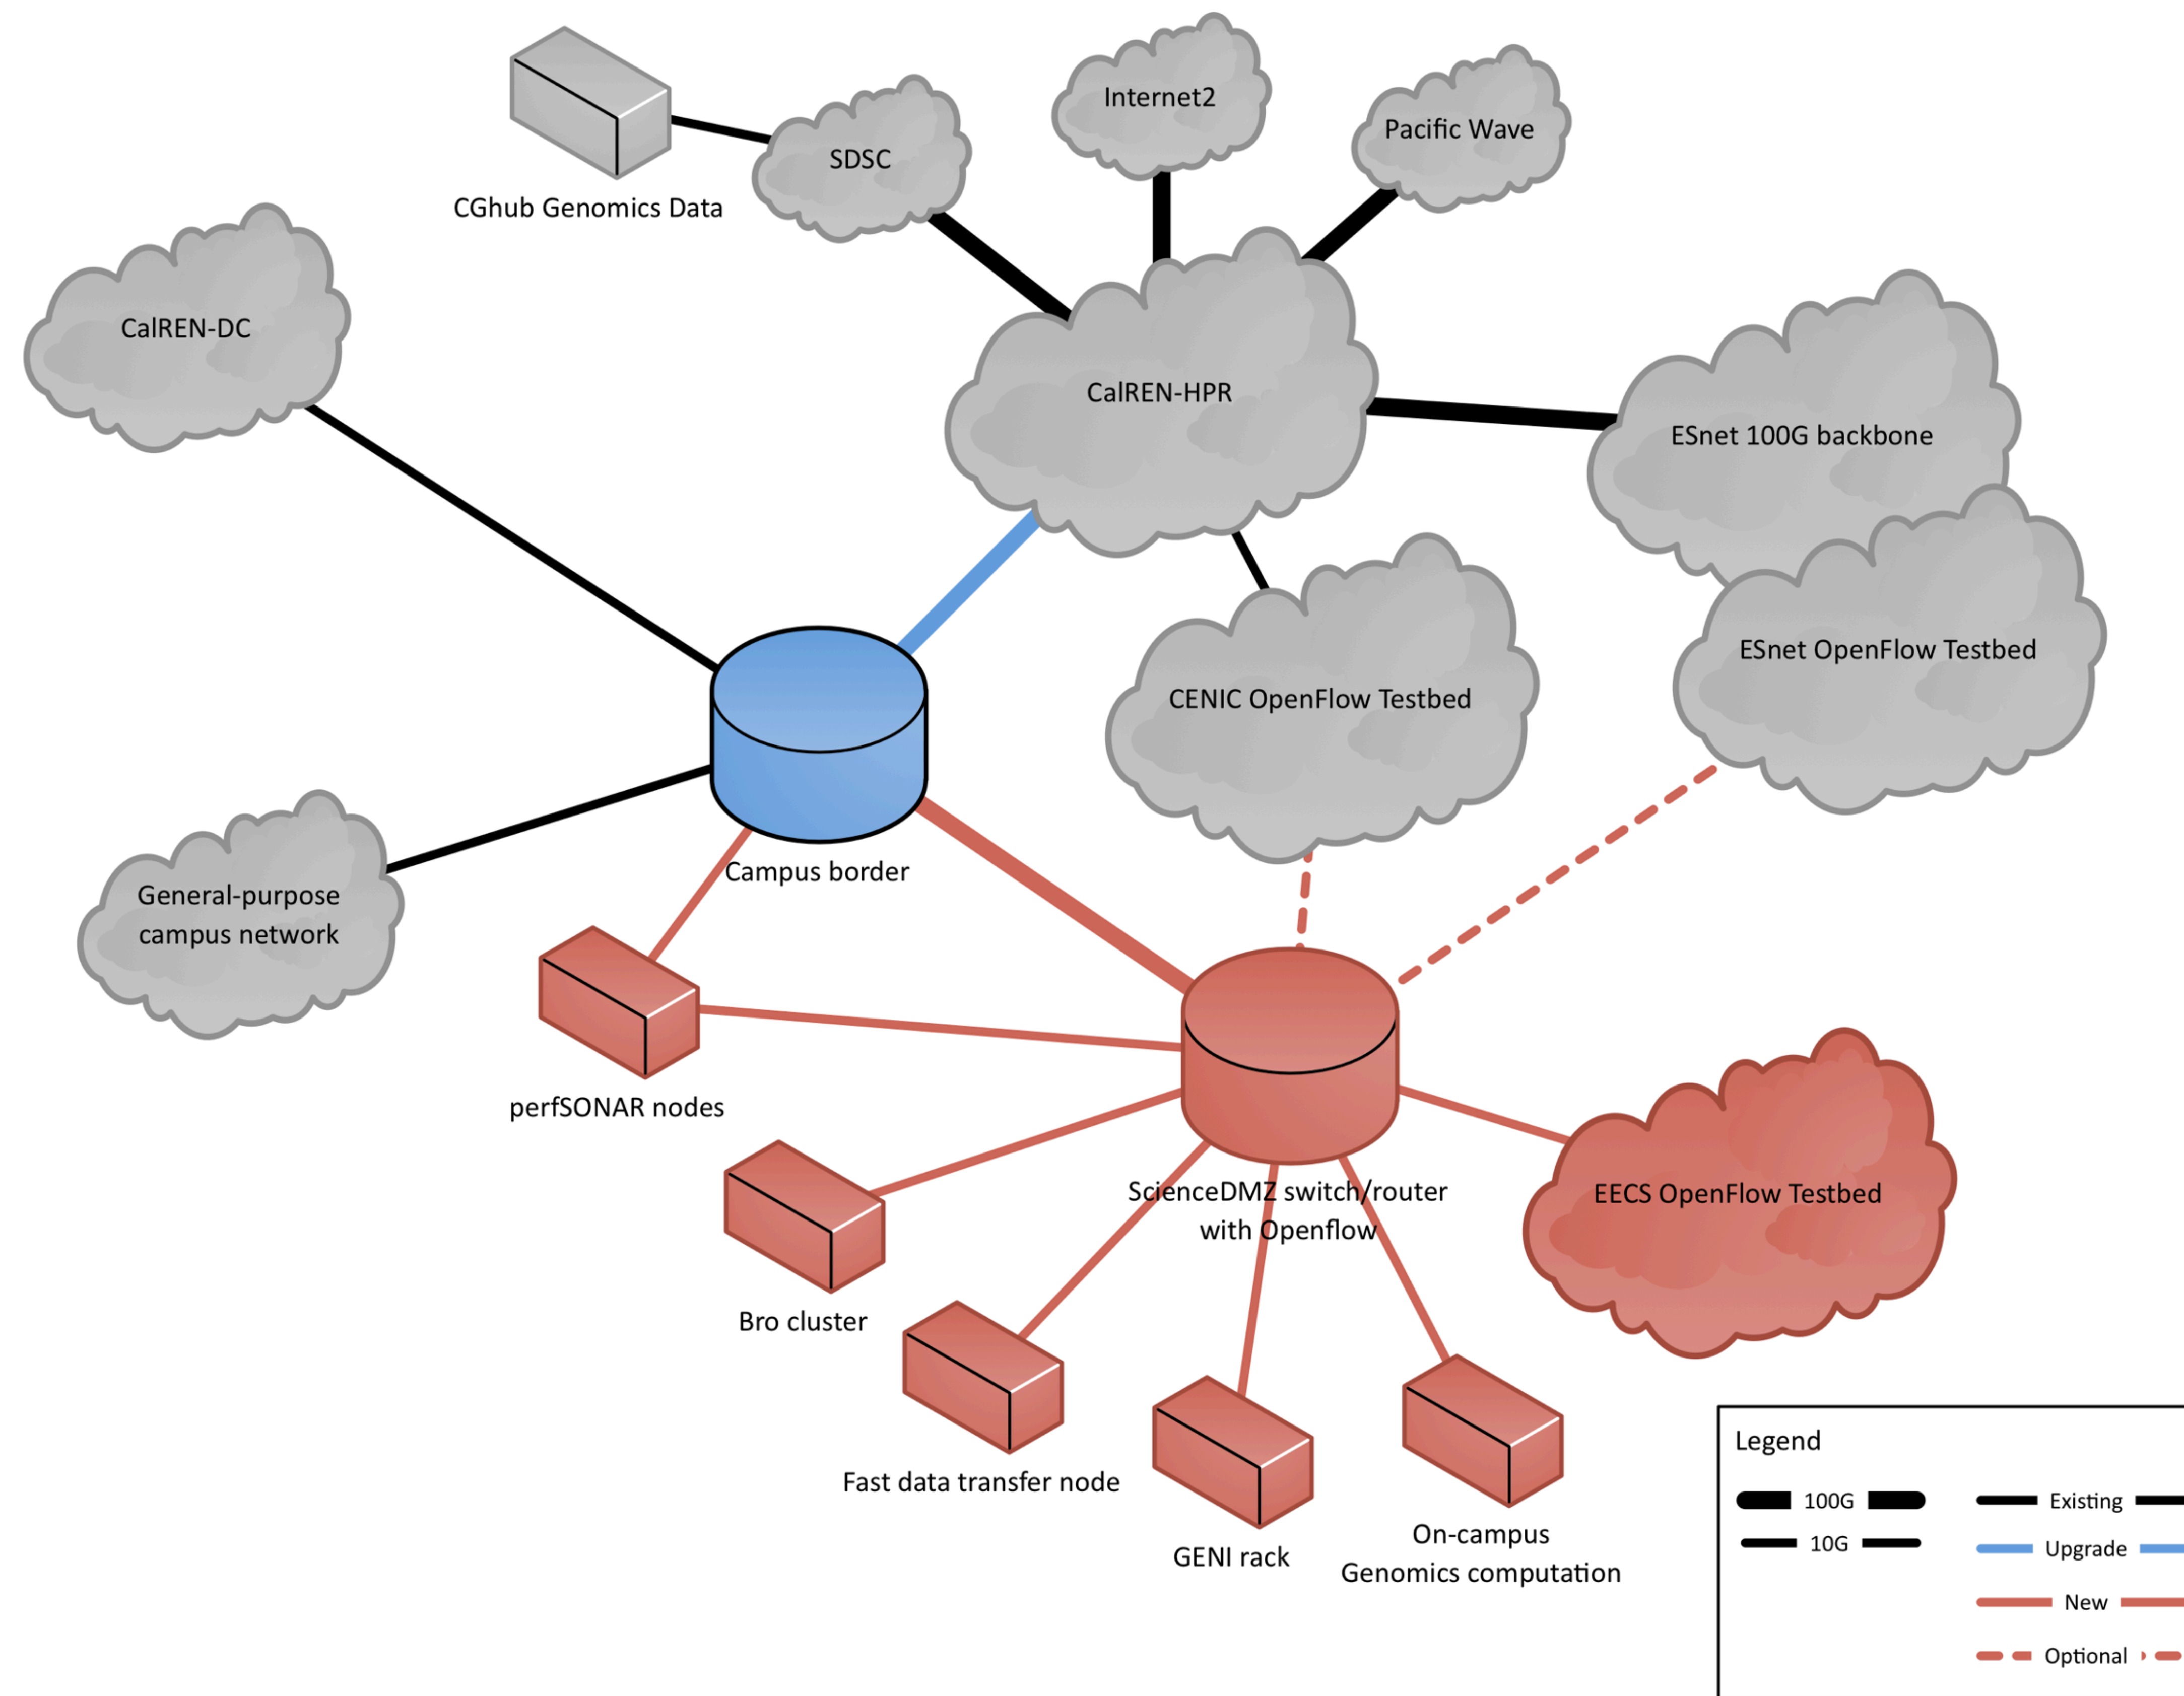
The implementation of the “friction free” networking component of UC Berkeley’s Science DMZ will include the following:

- Upgrade of a campus border router to provide 100Gb/s networking capabilities.
- Implementing a new 100Gb/s connection between UC Berkeley and CENIC’s CalREN-HPR network, which provides high speed connectivity to other institutions in California, nationally and globally.
- Implement a high speed switching infrastructure to carry campus Science DMZ traffic. Data Transfer Nodes, performance monitoring systems and other users of the Science DMZ will connect directly to this switching infrastructure to take advantage of its performance.

## Performance Monitoring

The Science DMZ architecture includes a test and measurement hosts based on perfSONAR. This host helps with fault diagnosis on the Science DMZ, and with end-to-end testing with collaborating sites if they have perfSONAR installed. The perfSONAR host can run continuous checks for latency changes and packet loss using OWAMP, as well as periodic throughput tests to remote locations using BWCTL. If a problem arises that requires a network engineer to troubleshoot the routing and switching infrastructure, the tools necessary to work the problem are already deployed - they need not be installed before troubleshooting can begin.

## UC Berkeley Science DMZ Implementation



## EXCEEDS – Extensible Cyberinfrastructure for Enhancing Extreme-Data Science

UC Berkeley’s Science DMZ is being developed as part of the EXCEEDS project. EXCEEDS was proposed by Sylvia Ratnasamy, an Assistant Professor in Computer Science at UC Berkeley and Scott Shenker, a Professor in the Electrical Engineering and Computer Science Department at UC Berkeley. The project has received funding through a grant from the National Science Foundation. The aim of the EXCEEDS project is to provide UC Berkeley with an advanced infrastructure for extreme-data science. Planned uses of the infrastructure are the testing of innovative software defined networking systems and supporting access to genomics data for medical research. An example data source is USCS’s Cancer Genomics Hub (CGHub), currently located at SDSC on the UC San Diego campus, which currently contains 450TB of data and is expected to expand to 5PB of data. Planned projects could include up to 100PB of data. The infrastructure developed as a result of EXCEEDS will enable researchers at the UC Berkeley campus to access this significant volume of data.

## Data Transfer Nodes

The computer systems used for wide area data transfers perform far better if they are purpose-built and dedicated to the function of wide area data transfer. These systems, which we call Data Transfer Nodes (DTNs), are typically PC-based Linux servers built with high-quality components and configured specifically for wide area data transfer. The DTN also has access to local storage, whether it is a local high-speed disk subsystem, a connection to a local storage infrastructure such as a SAN, or the direct mount of a high-speed parallel filesystem such as Lustre or GPFS, or a combination of these. The DTN runs the software tools designed for high-speed data transfer to remote systems – typical software packages include GridFTP and its service-oriented descendent Globus Online, discipline-specific tools such as XRootd, and versions of default toolsets such as SSH/SCP with high-performance patches applied.

DTNs typically have high-speed network interfaces (10Gbps currently, though experiments with 40Gbps DTNs are already underway), but the key is to match the DTN to the capabilities of the site and wide area network infrastructure. So, for example, if the network connection from the site to the WAN is one gigabit Ethernet, a 10 gigabit Ethernet interface on the DTN may be counterproductive.

## Colaboration

The following groups are involved in the EXCEEDS project and the development of a Science DMZ at UC Berkeley:

### UC Berkeley EECS:

The PIs who proposed EXCEEDS are faculty of EECS. EECS will be developing DTNs and an OpenFlow (software defined networking) testbed to utilize the Science DMZ infrastructure.

### UC Berkeley IST:

The IST Telecommunications department is responsible for developing the campus network infrastructure for the Science DMZ.

Multiple IST groups contributed to the development of a Campus Cyberinfrastructure plan which provided the context for the Science DMZ within the campus’s IT infrastructure.

### Energy Sciences Network (ESnet):

ESnet provides the high-bandwidth, reliable connections that link 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. Funded by the DOE Office of Science, and managed and operated by the ESnet team at Lawrence Berkeley National Laboratory, ESnet provides scientists with access to unique DOE research facilities and computing resources.

ESnet originated the concept of a Science DMZ and assisted UC Berkeley in developing the EXCEEDS project proposal.