



Live Help
Chat Now

Preview of Award 1341039 - Annual Project Report

[Cover](#) |
[Accomplishments](#) |
[Products](#) |
[Participants/Organizations](#) |
[Impacts](#) |
[Changes/Problems](#)

Cover

Federal Agency and Organization Element to Which 4900
Report is Submitted:

Federal Grant or Other Identifying Number Assigned by 1341039
Agency:

Project Title: CC-NIE Networking Infrastructure: 100 Gb/s
Science DMZ

PD/PI Name: Bradley R Smith, Principal Investigator
Mary Doyle, Co-Principal Investigator

Recipient Organization: University of California-Santa Cruz

Project/Grant Period: 01/01/2014 - 12/31/2015

Reporting Period: 01/01/2014 - 12/31/2014

Submitting Official (if other than PD\PI): N/A

Submission Date: N/A

Signature of Submitting Official (signature shall be
submitted in accordance with agency specific
instructions) N/A

Accomplishments

* What are the major goals of the project?

- 100 Gbps WAN connection to the Corporation for Educational Network Initiatives in California's (CENIC's) High Performance Research (HPR) Layer 3 (IP) network.
- Build a Science DMZ, located in the campus's data center, connected to this new 100 Gbps WAN connection and the existing L2 connection to CENIC's HPR L2 network.
- Installation of perfSONAR and a data transfer node (DTN) on the Science DMZ for researcher use.
- Extend the Science DMZ to buildings housing network intensive research activities to allow direct connection of research systems to the Science DMZ.
- Work with researchers to adapt their activities to make use of the Science DMZ and new network connections.
- Integrate SDN/OpenFlow technology into the Science DMZ to improve the capabilities available to researchers.

* What was accomplished under these goals (you must provide information for at least one of the 4 categories below)?

Major Activities: We convened a Science DMZ project team composed of the following:

- Brad Smith; Principal Investigator and Director Research & Faculty Partnerships (RFP) in the Information Technology Services Division (ITS).

- Jim Warner; Senior Network Engineer, RFP.
- John Haskins; Manager Network Operations Center (NOC), ITS.
- Mark Boolootian; Senior Network Engineer, NOC.
- George Peek; Network Engineer, NOC.
- Josh Sonstroem; Systems Administrator, Data Center Operations, ITS.
- Shawfeng Dong; Project Scientist for the Hyades supercomputer cluster (<http://hipacc.ucsc.edu/Hyades/>).
- Alan Lin; student (originally undergraduate, now Computer Engineering MS).

This team continued dialogues with the participating projects from the proposal on the integration of the Science DMZ into their research activities, and on-going collaboration in meeting the needs of these important research project. Participating projects include: Astrophysics and Earth & Planetary Sciences communities using the Hyades cluster, Center for Biomolecular Science and Engineering (now the UC Santa Cruz Genomics Institute), and the Santa Cruz Institute for Particle Physics (SCIPP).

CENIC has convened a Technical Advisory sub-Committee of CENIC research institutions customers who have a CC-NIE award. A subset of the team participates on the monthly CENIC CC-NIE TAC calls.

Alan Lin (the student) is working on a project to analyze the high-bandwidth data transfer tools using innovative network testbed where two Linux machines are configured as data transfer client and server, and a third machine is configured to simulate network conditions (e.g. bandwidth, delay, loss rate, etc.).

Specific Objectives:

To date, we have completed the following specific tasks:

- Acquired, installed and configured the Science DMZ router (a Brocade MLXe).
- Working with CENIC, we have installed the 100 Gbps link to CENIC's HPR L3 network, and configured routing on the Science DMZ router to use this new link for all destinations advertised from the HPR L3 networks.
- Installed a perfSONAR node on the new DMZ, connected at 10 Gbps.
- Installed a DTN on the new DMZ, conected at 10 Gbps, with 10TB of disk capacity configured to handle disk-to-disk high-bandwidth transfers.
- Integrated UCSC Cruz ID authentication (based on InCommon) in all Science DMZ systems.
- Configured UCSC authentication systems to work with Globus Online, allowing UCSC users to easily register with Globus Online using their UCSC Cruz ID.

Significant Results:

- Installed and configured 100 Gbps link to CENIC HPR L3.

Key outcomes or Other achievements:

- Established 100 Gbps-connected Science DMZ.
- Launched proof-of-concept DTN service for broader campus use.

* What opportunities for training and professional development has the project provided?

The HPC Network Engineering staff, project student, and the NOC and RFP network engineering staff all attended the February 27-28, 2014 OIN workshop presented by ESnet on the Science DMZ architecture, perfSONAR, DTNs, and Software Defined Networking.

The student participating on the project is using his work analyzing high-bandwidth data transfer tools (analyzing the performance of high-bandwidth data transfer nodes) as the topic of his Master's thesis.

* How have the results been disseminated to communities of interest?

We have presented or are scheduled to present an overview of the Science DMZ and its services to a number of campus

groups: the Academic Senate Committee on Computing and Telecommunications (requested by the committee; scheduled for 1/21/2015), the Advisory Committee for Information Technology (an advisory committee to the CIO composed of Assistant Deans, the campus Librarian, and faculty representatives; presented 12/17/2014), the IT Senior Management Group (scheduled for 1/14/2015), and the technical support staff for the academic Divisions (presented 12/1/2014). The initial response has been very positive; we have begun receiving requests for accounts on the systems.

The PI was an invited participant on NSF CC-NIE panel at the Internet2 Technology Exchange in Indianapolis in October, 2014.

The CIO (and Co-PI) wrote an article in the campus newsletter describing the project and expected benefits (<http://news.ucsc.edu/2014/11/doyle-keynote.html>)

One of the team member's (Josh Sonstroem) proposal for a presentation titled "Of Mice and Elephants (Or the Science DMZ and You)" has been accepted for the the CENIC Annual Conference (March 9-11, 2015; <http://cenic2015.cenic.org>).

* What do you plan to do during the next reporting period to accomplish the goals?

We have the following goals for the next year of the project:

- Select a product for use as the "Science DMZ switch," and acquire and deploy on at the first site (in School of Engineering for the UC Santa Cruz Genomics Institute).
- Begin dialog with researchers and sites to extend the Science DMZ to use the new Science DMZ switches.
- Continue spreading the word of our proof-of-concept DTN and perfSONAR services among campus researchers.
- Work with CENIC and the UC Santa Cruz Genomics Institute researchers to experiment with L2 circuits, as well as other means, for improving the ease and rate of data transfer between UCSC and the Cancer Genomics Hub (cghub.ucsc.edu) located at the San Diego Supercomputer Center.
- Begin developing a more robust architecture for L2 services, perhaps employing SDN capabilities included in the Science DMZ.
- Explore other possible uses of the SDN functionality available from the DMZ.
- We are participating in the "Cyberinfrastructure for Big Data Research Applications" project that is being coordinated by the Center for IT Research in the Interest of Society (CITRIS; www.citris-uc.org) and the California Institute for Telecommunications and Information Technology (CalIT2; <http://www.calit2.net>), and involves CENIC and a number of UC campuses. The goal of this project is to create a UC-wide working group on cyberinfrastructure for big data research charged with developing an inventory of data-intensive research programs at UC campuses and national labs that require moving Big Data between campuses, completing a comprehensive "gap analysis" to determine where further investments are needed to create reliable high-speed connections throughout the UC system, and draft a planning document for integrating all 10 campuses and affiliated entities for a statewide Science DMZ.

Products

Books

Book Chapters

Conference Papers and Presentations

Inventions

Journals

Licenses

Other Products

Other Publications

Mary Doyle, Vice Chancellor, Information Technology Services (2014). "Researchers will benefit from faster, more robust network connectivity". UCSC's CIO, and this project's CoPI, wrote an article to the campus community describing the Science DMZ project and the benefits it offers to researchers.. Status = PUBLISHED; Acknowledgement of Federal Support = No

Patents

Technologies or Techniques

Thesis/Dissertations

Websites

Participants/Organizations

What individuals have worked on the project?

Name	Most Senior Project Role	Nearest Person Month Worked
Smith, Bradley	PD/PI	1
Doyle, Mary	Co PD/PI	1
Boolootian, Mark	Other Professional	1
Haskins, John	Other Professional	1
Peek, George	Other Professional	2
Sonstroem, Josh	Other Professional	3
Warner, Jim	Other Professional	3
Dong, Shawfeng	Staff Scientist (doctoral level)	3
Lin, Alan	Graduate Student (research assistant)	5

Full details of individuals who have worked on the project:

Bradley R Smith

Email: brad@soe.ucsc.edu

Most Senior Project Role: PD/PI

Nearest Person Month Worked: 1

Contribution to the Project: Served as Principal Investigator. Invited participant on NSF CC-NIE panel at the Internet2 Technology Exchange in Indianapolis.

Funding Support: No funding other than the award.

International Collaboration: No

International Travel: No

What other organizations have been involved as partners?

Name	Type of Partner Organization	Location
CENIC	Other Nonprofits	16700 Valley View Ave. #400 La Mirada, CA 90638

Full details of organizations that have been involved as partners:

CENIC

Organization Type: Other Nonprofits

Organization Location: 16700 Valley View Ave. #400 La Mirada, CA 90638

Partner's Contribution to the Project:

Facilities

Collaborative Research

Personnel Exchanges

More Detail on Partner and Contribution: CENIC collaborates with its member organizations in the development and implementation, and coordinates member organization forums to support collaboration among members.

Have other collaborators or contacts been involved? Yes

Impacts

What is the impact on the development of the principal discipline(s) of the project?

The focus of this project is to enhance the IT infrastructure for network-intensive research at UCSC. The impact we have had so far is primarily in adapting the Science DMZ architecture to the specifics of UCSC's environment, and implementing the resulting architecture. Through our participation in the CENIC HPR Technical Advisory Committee (TAC), we have shared the results of our work with the broader CENIC research community as well as learned from the lessons other CC-NIE projects in the HPR-TAC community have learned.

Examples of architecture adaptations we've addressed include 1) focus on minimizing costs for future upgrades to the infrastructure, and 2) support scalability to allow easy expansion of the DMZ as new disciplines and building evolve to make use of it.

UCSC (and many of the institutions in the region) employ a campus architecture based on "distribution-layer" firewalls where the issue of avoiding firewall restrictions for the Science DMZ doesn't exist (or is a minimal impact). Instead, the focus of our architecture has been to implement the Science DMZ as close to the off-campus link to minimize the costs involved in future connectivity upgrades (e.g. to 400 Gbps and then 1 Tbps).

UCSC is a large campus, with Departments distributed among many geographically dispersed buildings. The Science DMZ architecture we have developed allows for the easy expansion of the Science DMZ to new buildings as new disciplines evolve to take advantage of Science DMZ services. In addition, referring to the previous bullet, this architecture allows for a cost effective evolution of the link technology from 10 Gbps to 40 Gbps, 100 Gbps and beyond.

What is the impact on other disciplines?

Nothing to report.

What is the impact on the development of human resources?

This project has provided a topic for Alan Lin to work on and use in completing his BS degree in Computer Engineering, and

to now use it as the topic for his MS CE thesis.

What is the impact on physical resources that form infrastructure?

To date, this project has resulted in two significant enhancements to the campus's network infrastructure. First, it has resulted in the establishment of a site (rack space and power) in the campus communications center for the Science DMZ (including fiber paths to the campus border routers, to existing research equipment in the data center, and to the facilities housing the 100 Gbps demarc), and the acquisition, configuration, and installation of the Science DMZ router to feed all of these connection. And second, it has resulted in the provisioning of a DTN and perfSONAR system in the Science DMZ

What is the impact on institutional resources that form infrastructure?

Nothing to report.

What is the impact on information resources that form infrastructure?

The availability of the Science DMZ-connected DTN provides a significant enhancement of infrastructure for accessing information resources. This capability is currently being "socialized" across the campus.

What is the impact on technology transfer?

Nothing to report.

What is the impact on society beyond science and technology?

While the potential for such impact is clear, with the Science DMZ and related services (e.g. the DTN) having just been established, to date there has been no such impact.

Changes/Problems

Changes in approach and reason for change

We engaged two staff at 25% time each for the 50% HPC Network Engineer position. The reason for this is we had two strong staff available to help with the project who offered different, and complementary skillsets to the project.

Josh Sonstroem is a member of the campus's core systems administration group. He has in depth understanding of the campus's systems environment, of UCSC's InCommon/Shibboleth implementation (UCSC Cruz ID), and a good background in networking. We felt Josh would be a big help in implementing and managing the systems side of the Science DMZ, including integrating it with the existing campus operation environment.

Dr. Shawfeng Dong is a PhD Astrophysicist, the architect and administrator of the two generations of supercomputing clusters (Pleiades and Hyades) used extensively by researchers in the Astrophysics, Earth & Planetary Sciences, and Applied Math and Statistics Departments. He has an in-depth understanding of the challenges researchers face in moving large research data sets, as well as an in-depth understanding of the research world, and the issues and priorities researchers have in dealing with their computing infrastructure. We felt Shaw would be invaluable in keeping the Science DMZ services relevant to researchers needs in general, and specifically in designing the data transfer systems and working with researchers to adapt their research flows to integrate the Science DMZ tools.

So far we believe this decision has been a big success.

Actual or Anticipated problems or delays and actions or plans to resolve them

The campus is undergoing the first major communications infrastructure upgrade since its founding. As a part of this project, the border routers are being upgraded. The Science DMZ architecture we developed includes redundant connections between the Science DMZ router and the border switches (to provide robust, production grade paths to the Internet). The old border routers are full, and do not have any spare 10 Gbps ports to connect the Science DMZ router to. This will be addressed by the new border routers. The one timing mis-match that we've encountered is the Science DMZ router was installed and connected to the new 100 Gbps circuit before the border routers could be upgraded (currently estimated for Q1 2015). To work around this we are using VLANs from the Science DMZ router, through an existing router, to the borders for

this connectivity. We will upgrade these links to fiber when the border router upgrade is complete.

Changes that have a significant impact on expenditures

The evolution of SDN-capable switches, and “whitebox” switches in specific, will likely result in a significantly reduced cost for the Science DMZ switches we will be acquiring. This will free up funding to either acquire more Science DMZ switches and extend the Science DMZ to more research buildings, or to increase funding in other parts of the project.

Significant changes in use or care of human subjects

Nothing to report.

Significant changes in use or care of vertebrate animals

Nothing to report.

Significant changes in use or care of biohazards

Nothing to report.