

NSF CC* Networking Infrastructure: Building a Science DMZ for Data-intensive Research and Computation at the University of South Carolina

Jorge Crichigno
College of Engineering and Computing
University of South Carolina
<http://ce.sc.edu/cyberinfra>

2022 NSF CC* PI Workshop
September 19 – 22, 2022
Minneapolis, MN



UNIVERSITY OF
SOUTH CAROLINA



NSF CAMPUS CYBERINFRASTRUCTURE PI WORKSHOP

SEPTEMBER 19 – 22, 2022

Minneapolis, MN



Quad Chart for:

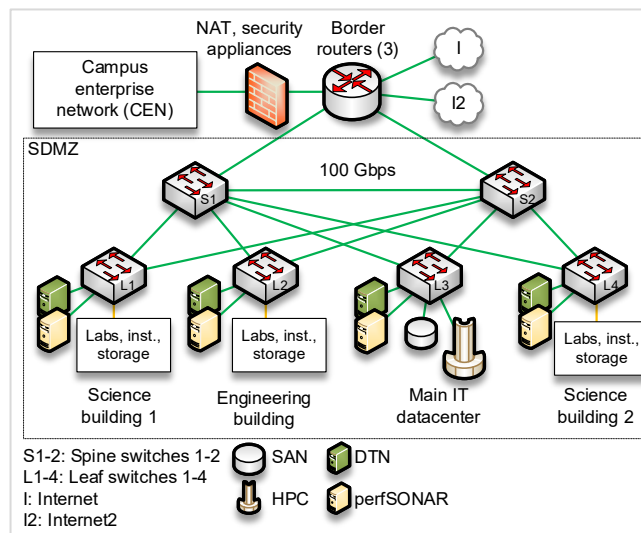
Building a Science DMZ for Data-intensive Research and Computation at the University of South Carolina

Challenge Project Seeks to Address:

- 35,000 students on main campus and 50,000 statewide
- Very high research activity: external and internal science data transfers
- Limited capacity to Internet and Internet2, approximately 10 Gbps
- Research and commodity traffic on the same network
- Need for an efficient approach to move big science data

Solutions and Deliverable:

- A 100 Gbps Science DMZ, co-located to campus network
- Increased bandwidth to Internet2 from ~10 Gbps to 100 Gbps
- Increased bandwidth from key research laboratories on campus to Science DMZ, 1 Gbps to 40-100 Gbps
- Faculty members and IT's Research Computing (RC) group now participate in the CC* community

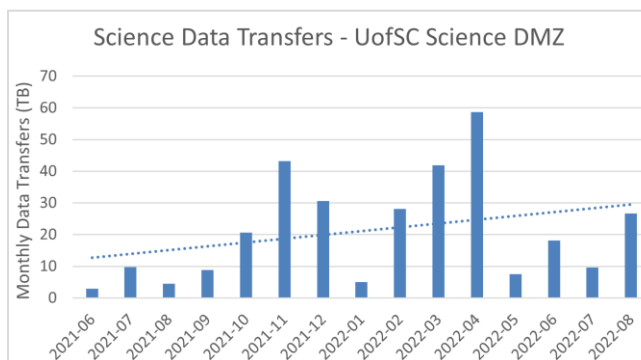


Scientific Impact:

- Enhanced research capability
- Increased number of science data transfers
- Research activities on areas including cognitive processes, WWII film digitization, chemical engineering, nuclear physics, environmental nanoscience, genomics data mining, digital image processing, and others
- Four UofSC XSEDE campus champions
- Findings related to the project published in journals, conferences, and a book

Workforce Development:

- Research testbed deployed for proof-of-concept ideas, prior to production deployment
- 43 undergraduate students completed one or more semesters working on CI topics
- Dissemination of best practices via workshops conducted with ESnet, EPOC, RENs
- Agreements with industry, e.g., Barefoot Networks / Intel, VMware, Palo Alto Networks, Cisco Systems






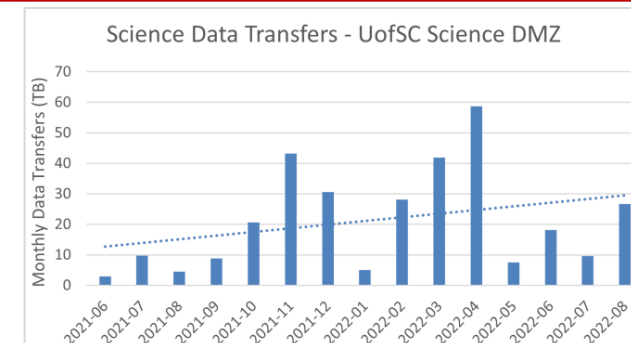
NSF CAMPUS CYBERINFRASTRUCTURE PI WORKSHOP

SEPTEMBER 19 – 22, 2022
Minneapolis, MN

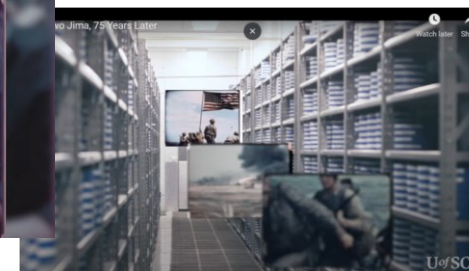


Outcomes arising from CC* award

- Increased number of science data transfers
- Research activities on areas including cognitive processes, WWII film digitization, chemical engineering, nuclear physics, and others
- Campus infrastructure suitable for an R1 institution
 - 100 Gbps Science DMZ
 - 100 Gbps connection to Internet2
 - Direct connection to providers (AWS)
- Strengthened collaboration between IT and faculty
- Partnership with agencies
 - National Laboratories
 - SC State Administration



WWII film digitization
14,000+ cans of film



<https://tinyurl.com/59cwtbs3>



CIO SRNL, interns, PI





NSF CAMPUS CYBERINFRASTRUCTURE PI WORKSHOP



SEPTEMBER 19 – 22, 2022

Minneapolis, MN



Outcomes arising from CC* Award – workforce development

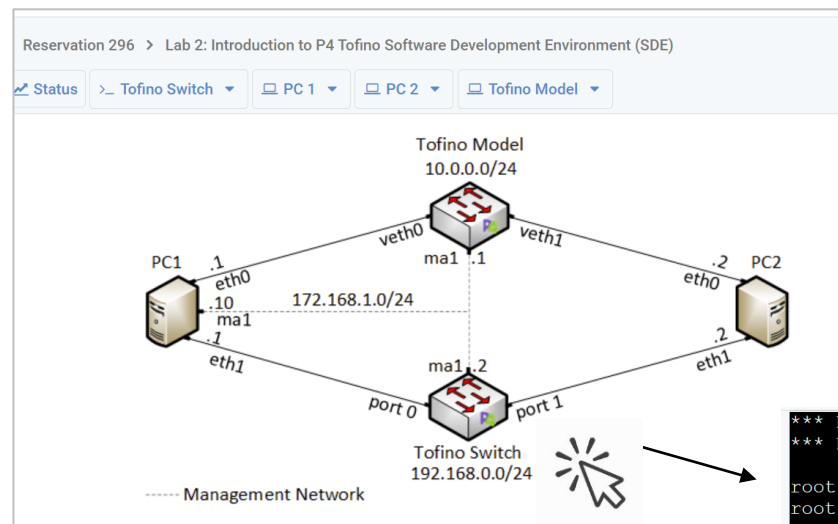
- Research testbed was deployed for proof-of-concept ideas
- Undergraduate research on CI: 43 undergraduate students completed one or more semesters working on CI topics



Undergraduate researchers - Summer Research Symposium



Testbed – 100 Gbps programmable network accessible via regular browser



```

*** NETLAB: CONNECTING
*** NETLAB: CONNECTED

root@tofino-switch:~#
root@tofino-switch:~#
root@tofino-switch:~#

```



NSF CAMPUS CYBERINFRASTRUCTURE PI WORKSHOP


SEPTEMBER 19 – 22, 2022

Minneapolis, MN




Outcomes arising from CC* Award – workforce development

- Graduate research on CI: 3 conference papers, 3 journal articles, and 1 book were published during the 3-year project
- New relations were established
 - Amazon – AWS “Direct Connect” to resources (via I2)
 - Barefoot Networks / Intel
 - Juniper Networks
 - VMware



Computer Communications
Volume 161, 1 September 2020, Pages 212-224



An emulation-based evaluation of TCP BBRv2 Alpha for wired broadband

Elie F. Kfoury ^a, Jose Gomez ^a, Jorge Crichigno ^a, Elias Bou-Harb ^b

Show more

Journals & Magazines > IEEE Communications Surveys & Tutorials > Volume: 21 Issue: 2


A Comprehensive Tutorial on Science DMZ

Publisher: IEEE Cite This PDF

Jorge Crichigno ; Elias Bou-Harb ; Nasir Ghani All Authors

14 Paper Citations 1200 Full Text Views

Abstract: Science and engineering applications are now generating data at an unprecedented rate. From large facilities such as the Large Hadron Collider to portable DNA sequencing devices, these instruments can produce hundreds of terabytes in short



Conferences > 2022 18th International Conference on Cloud Computing

Enabling P4 Hands-on Training in an Academic Cloud

Publisher: IEEE Cite This PDF

Jose Gomez ; Elie F. Kfoury ; Jorge Crichigno All Authors

Abstract: This paper describes a cloud infrastructure and virtual laboratories on P4 programmable data plane switches. P4 programmable data planes emerged as a technology that enables innovation in networking. P4 is a programming language used to describe how network packets are processed. This paper explains an entry-

Journals & Magazines > IEEE Access > Volume: 9

An Exhaustive Survey on P4 Programmable Data Plane Switches: Taxonomy, Applications, Challenges, and Future Trends

Publisher: IEEE Cite This PDF

Elie F. Kfoury ; Jorge Crichigno ; Elias Bou-Harb All Authors

10 Paper Citations 3876 Full Text Views

Open Access Comment(s)

Under a Creative Commons License



NSF CAMPUS CYBERINFRASTRUCTURE PI WORKSHOP

SEPTEMBER 19 – 22, 2022

Minneapolis, MN



Sustainability

- Strengthen relation among faculty, Research Computing (RC), and IT leadership
- RC, IT: communicate continuously with faculty, researchers: webinars, trainings
- Faculty, researchers: provide input to RC/IT on research needs, opportunities
- RC, IT, faculty: implement work study and undergraduate research positions at RC/IT
- Faculty: promote positions and recruit students early –freshman, sophomore
- Share ideas with the community, receive feedback, learn best practices. E.g., CI Lunch and Learn Series; Intel's Connectivity Research Program, FABRIC
- Proactively seek funding opportunities, donations, common projects with industry, agencies

CUSTOMER STORY

vmware IT ACADEMY



Uof SC South Carolina

INDUSTRY
UNIVERSITY OF SOUTH CAROLINA
COLLEGE OF ENGINEERING
AND COMPUTING

LOCATION
COLUMBIA, SOUTH CAROLINA

KEY CHALLENGES

- Needed to educate students who were located in multiple academic and military institutions for high-demand

The University of South Carolina partners with VMware IT Academy to help students learn digital technology skills to fill high-demand jobs