2023 Internet2 Technology Exchange

Science DMZs and Networking for All

Science DMZ concepts, campus network deployments, perfSONAR

Jorge Crichigno
University of South Carolina
https://research.cec.sc.edu/cyberinfra/

University of South Carolina (USC)
Energy Sciences Network (ESnet)

September 18, 2023

Workshop Website

- All material is posted on the website of the tutorial

  https://research.cec.sc.edu/cyberinfra/workshop-techex2

Motivation for a High-Speed Science Architecture

- Science and engineering applications are generating data at an unprecedented rate
- Instruments produce hundreds of terabytes in short time periods (“big science data”)
- Data must be typically transferred across high-bandwidth high-latency Wide Area Networks (WANs)

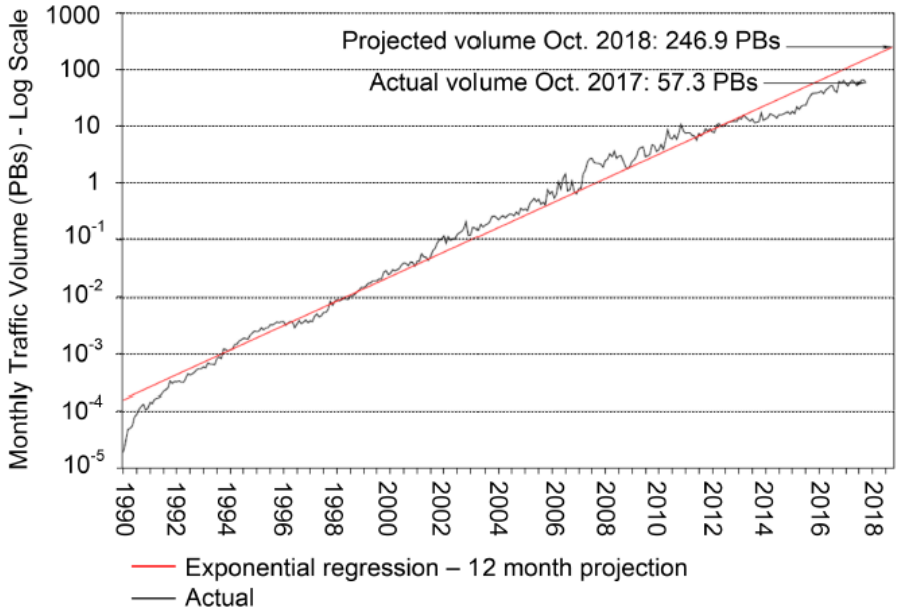

The Energy Science Network (ESnet) is the backbone connecting U.S. national laboratories and research centers

Enterprise Network Limitations

- Security appliances (IPS, firewalls, etc.) are CPU-intensive
- Inability of small-buffer routers/switches to absorb traffic bursts
- End devices incapable of sending/receiving data at high rates
- Lack of data transfer applications to exploit available bandwidth
- Many of the issues above relate to TCP

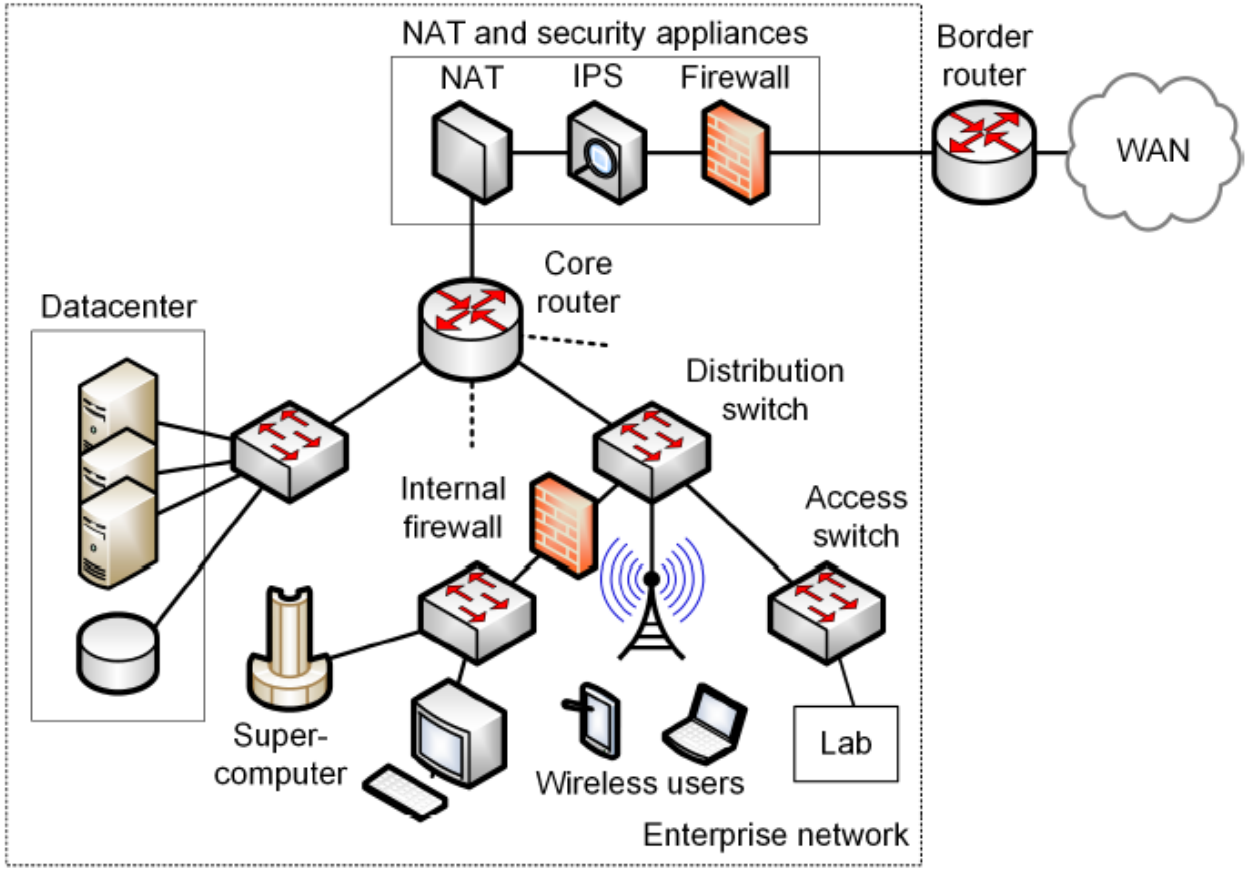

Enterprise Network Limitations

- Effect of packet loss and latency on TCP throughput

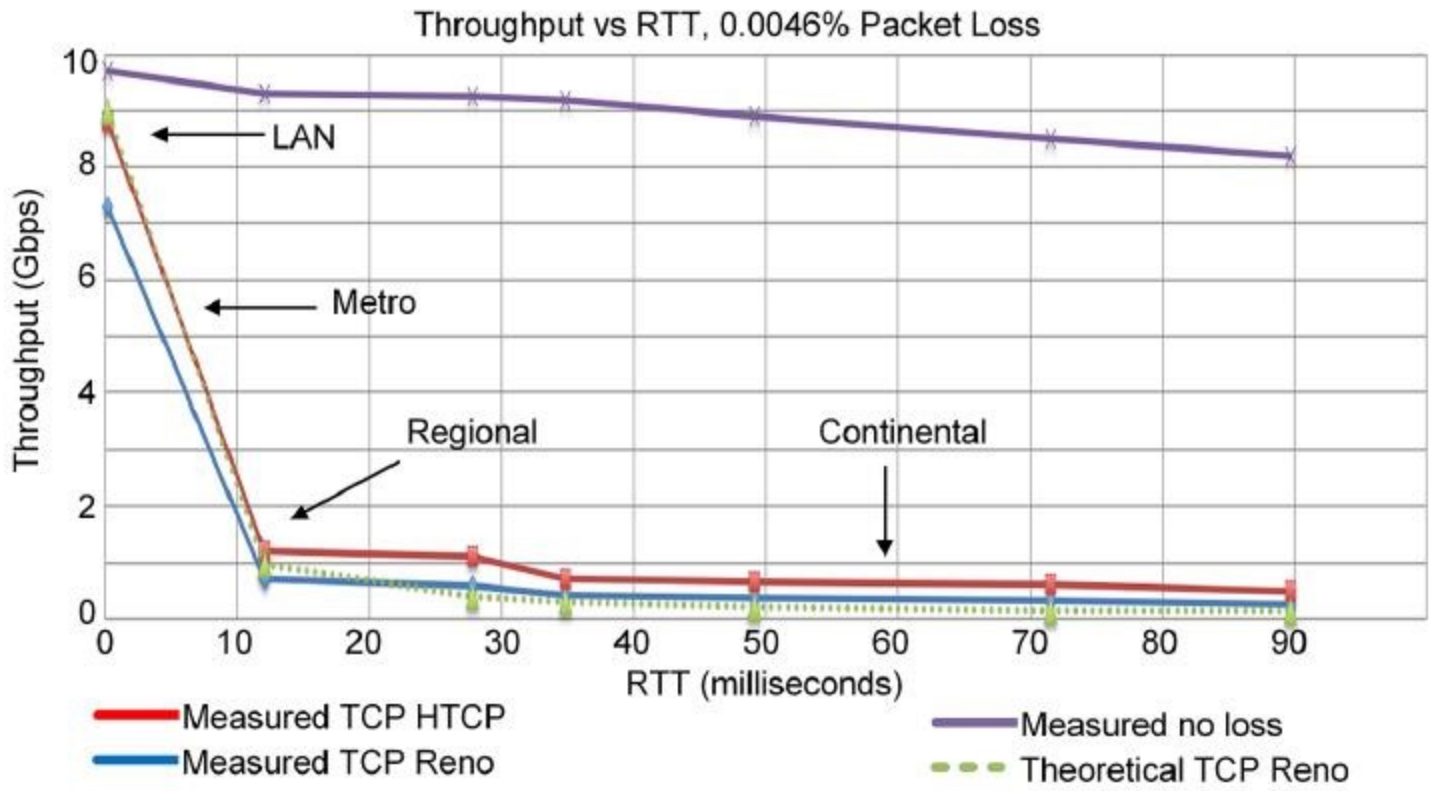

E. Dart, L. Rotman, B. Tierney, M. Hester, J. Zurawski, "The science dmz: a network design pattern for data-intensive science," *International Conference on High Performance Computing, Networking, Storage and Analysis*, Nov. 2013.

Science DMZ

- The Science DMZ is a network designed for big science data
- Main elements
  - High throughput, friction free WAN paths
  - Data Transfer Nodes (DTNs)
  - End-to-end monitoring = perfSONAR
  - Security tailored for high speeds

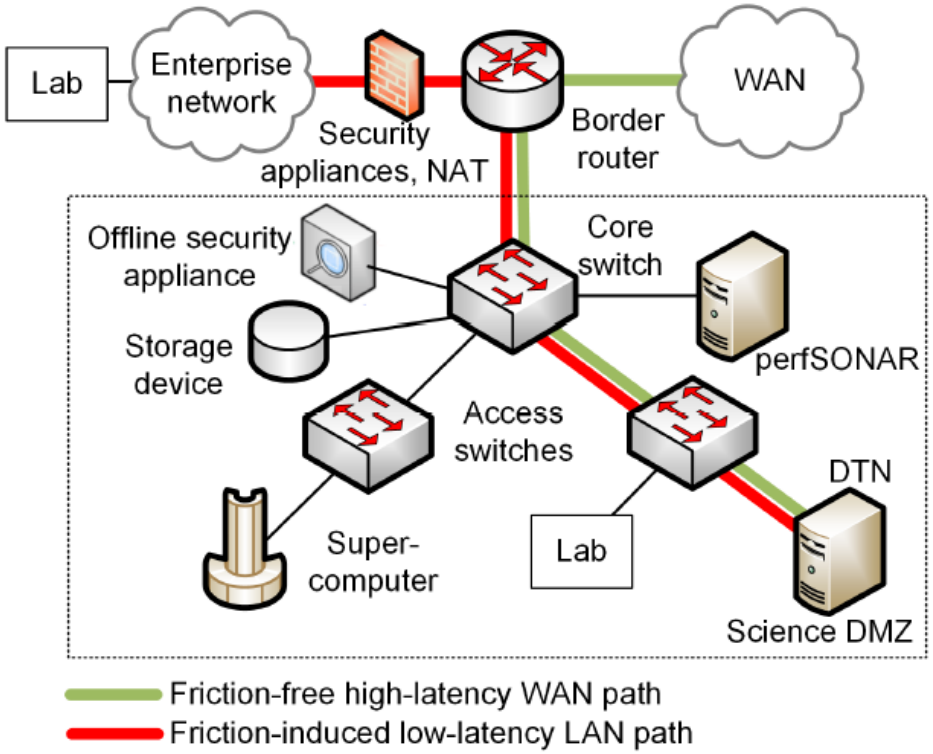

Science DMZ

- The Science DMZ is a network designed for big science data
- Main elements
  - High throughput, friction free WAN paths
  - Data Transfer Nodes (DTNs)
  - End-to-end monitoring = perfSONAR
  - Security tailored for high speeds

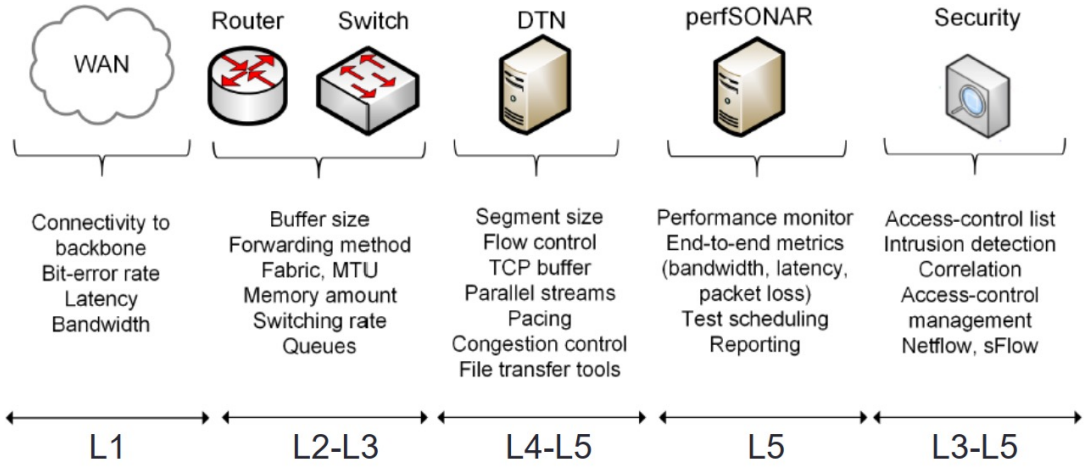

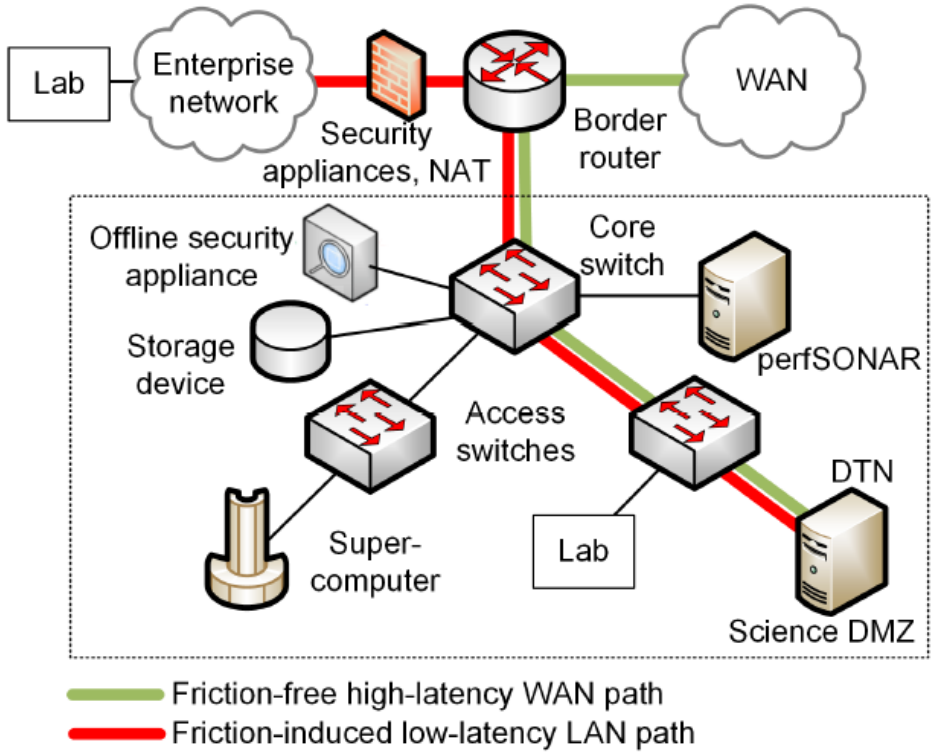

Science DMZ

- Science DMZ deployments, U.S.

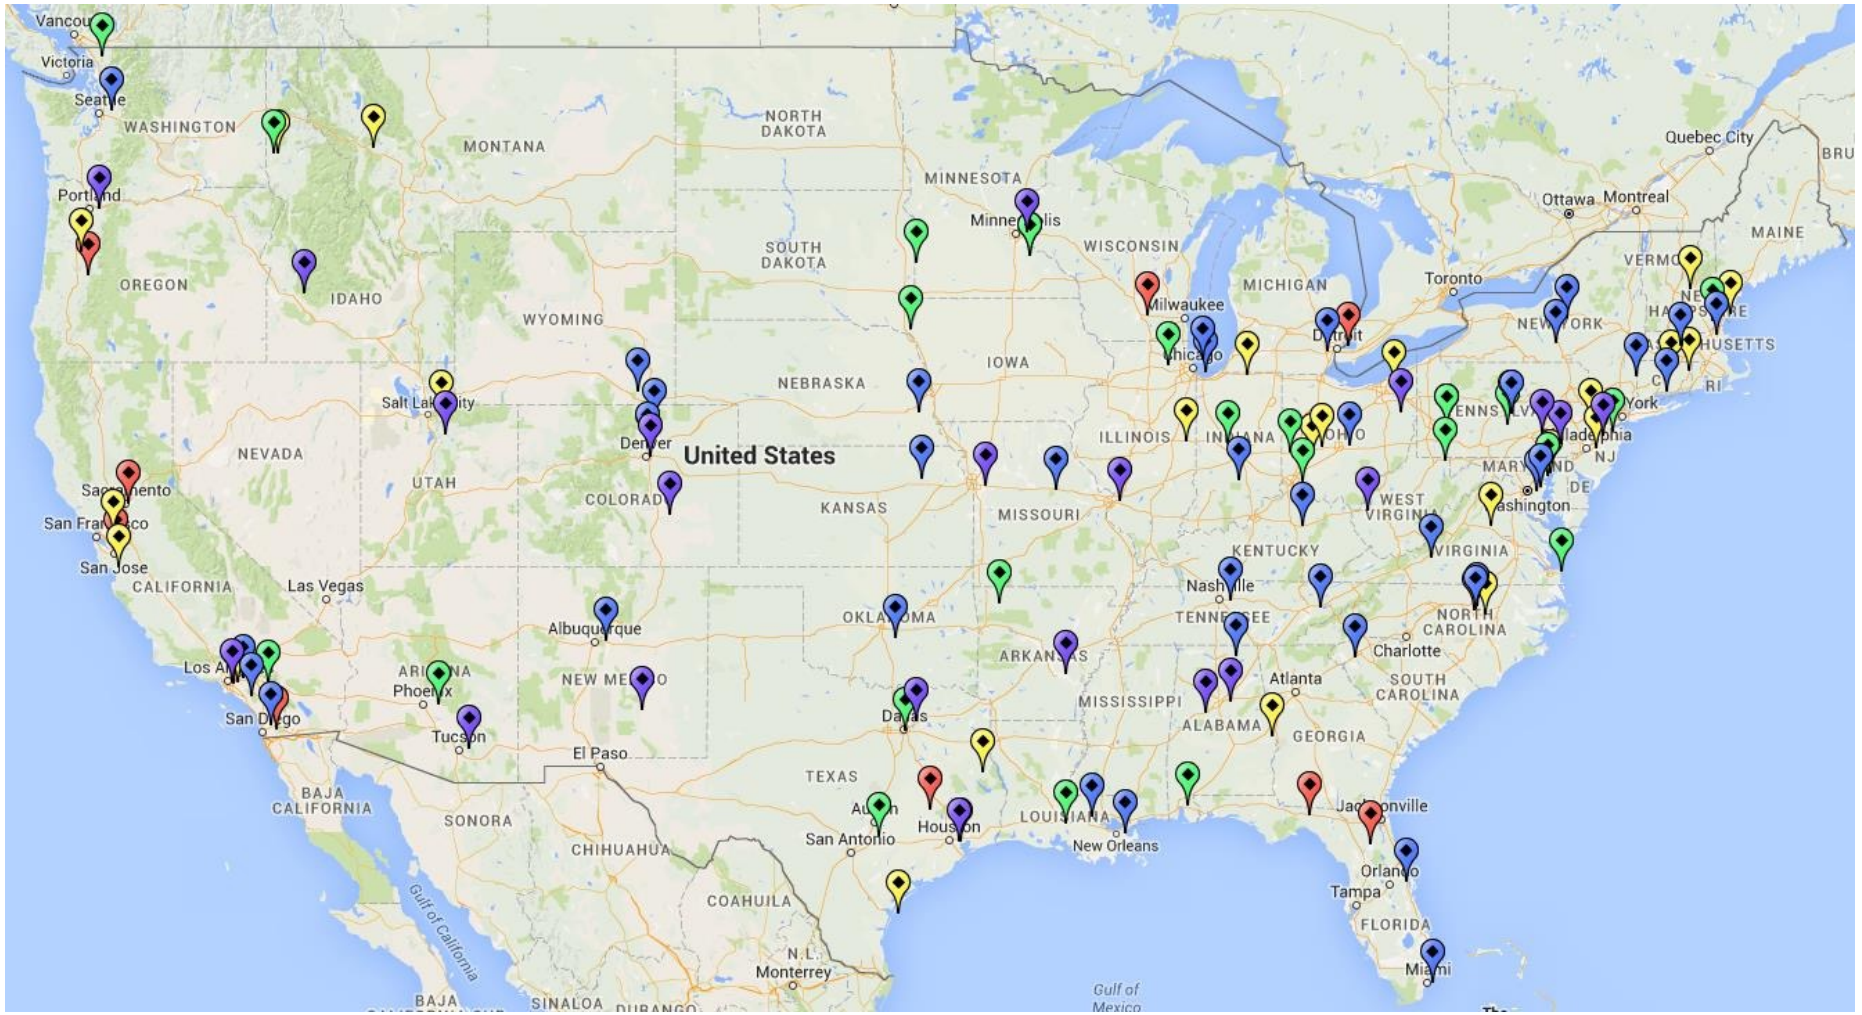

perfSONAR

- The components of the Science DMZ can eventually fail
- Networks and systems comprise many components that can experience bugs, misconfigurations, and other problems
- Such problems result in total failure or the underutilization of resources

perfSONAR

- These network issues are classified as:
- Hard network failures: when a component stops working (e.g., fiber cut, power outage, hardware malfunction)
- Soft network failures: performance issues that cause performance degradation (e.g., long delays and packet losses)
- Hard network failures are easy to detect, whereas soft network failures can remain undetected for a long time

perfSONAR

- perfSONAR is a network measurement tool designed to provide federated coverage of paths and help achieve end-to-end usage expectations
- It facilitates diagnosing, visualizing, and troubleshooting network performance issues such as soft failures and hard failures
- perfSONAR is a key component of the Science DMZ

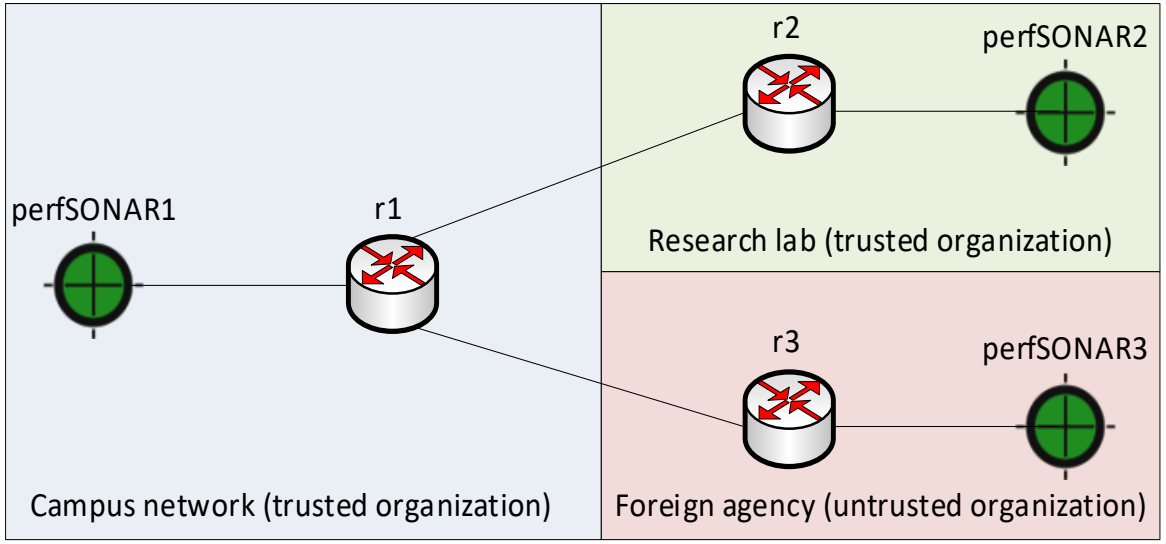

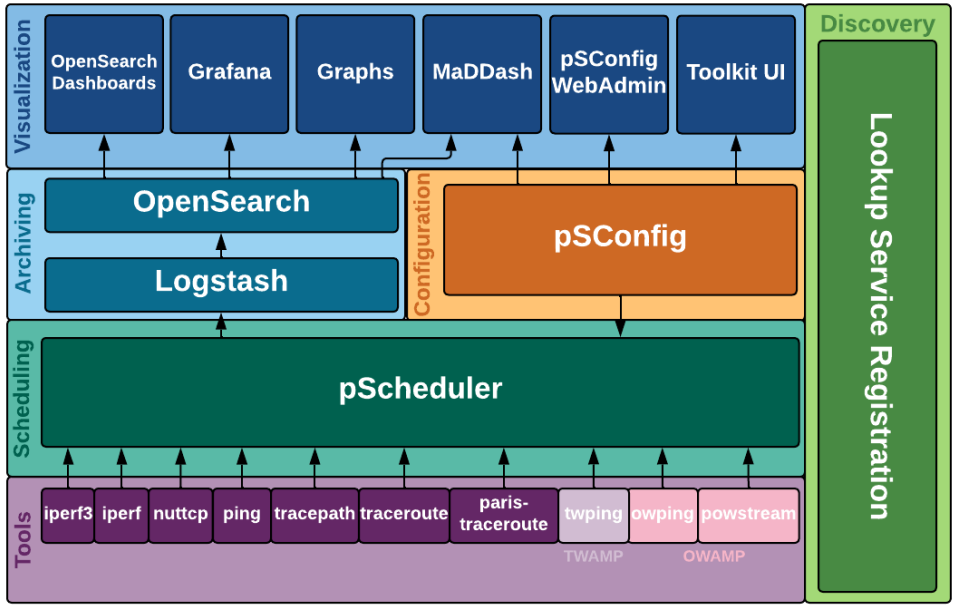

perfSONAR

- perfSONAR consists of a suite of tools that facilitate running, storing, and visualizing network measurements
- Network administrators can schedule regular tests and display them on a dashboard
- perfSONAR provides information about network failures (where, when, what)

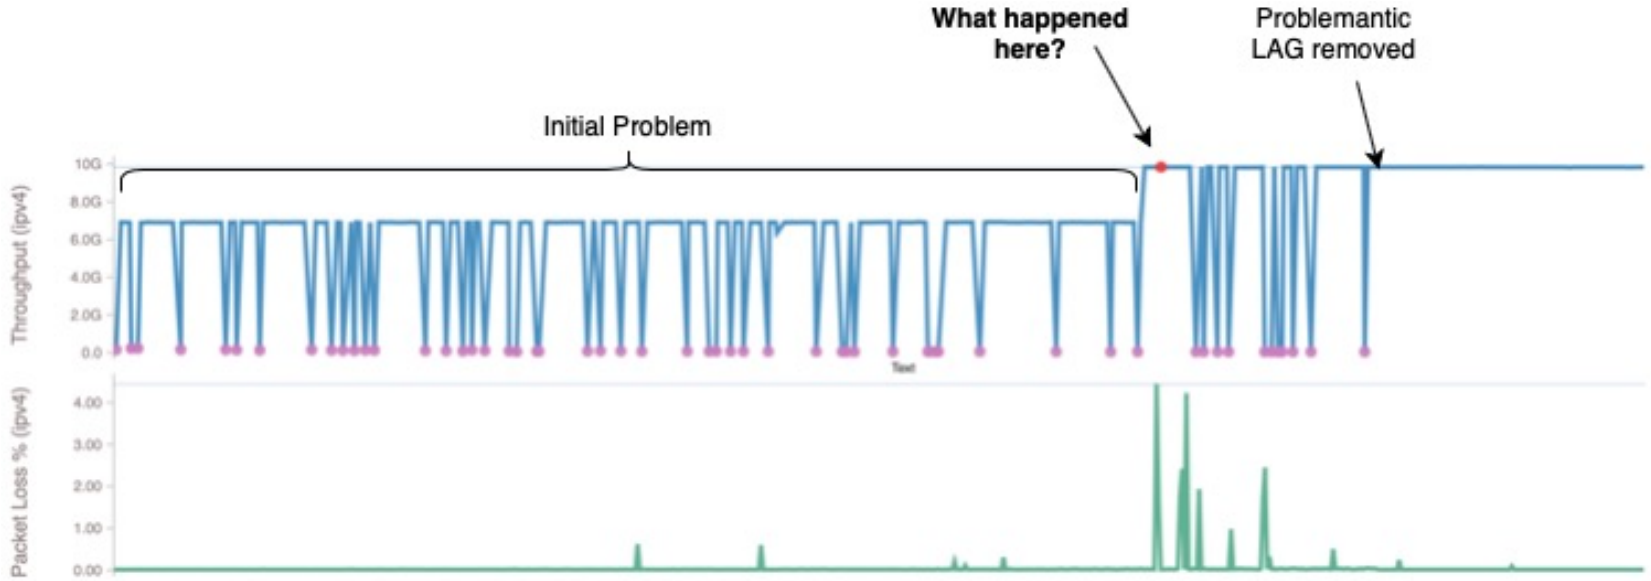

perfSONAR Use Cases. “Pan-STARRS Hawaii to Queen’s University Belfast.” [Online]. Available: https://tinyurl.com/3jjtypfr

Summary

- The Science DMZ is an efficient approach to enable large data transfers
- Such data transfers experience limitations over regular enterprise networks
- Key elements of a Science DMZ: friction-free WAN path, DTNs, perfSONAR, security tailored for the Science DMZ
- Problems may occur within a Science DMZ
- perfSONAR offers the resources to visualize and troubleshoot network problems

Additional Slides

- BBR performance on FABRIC
- Performance measurements for a single flow, 0.0046% packet loss rate

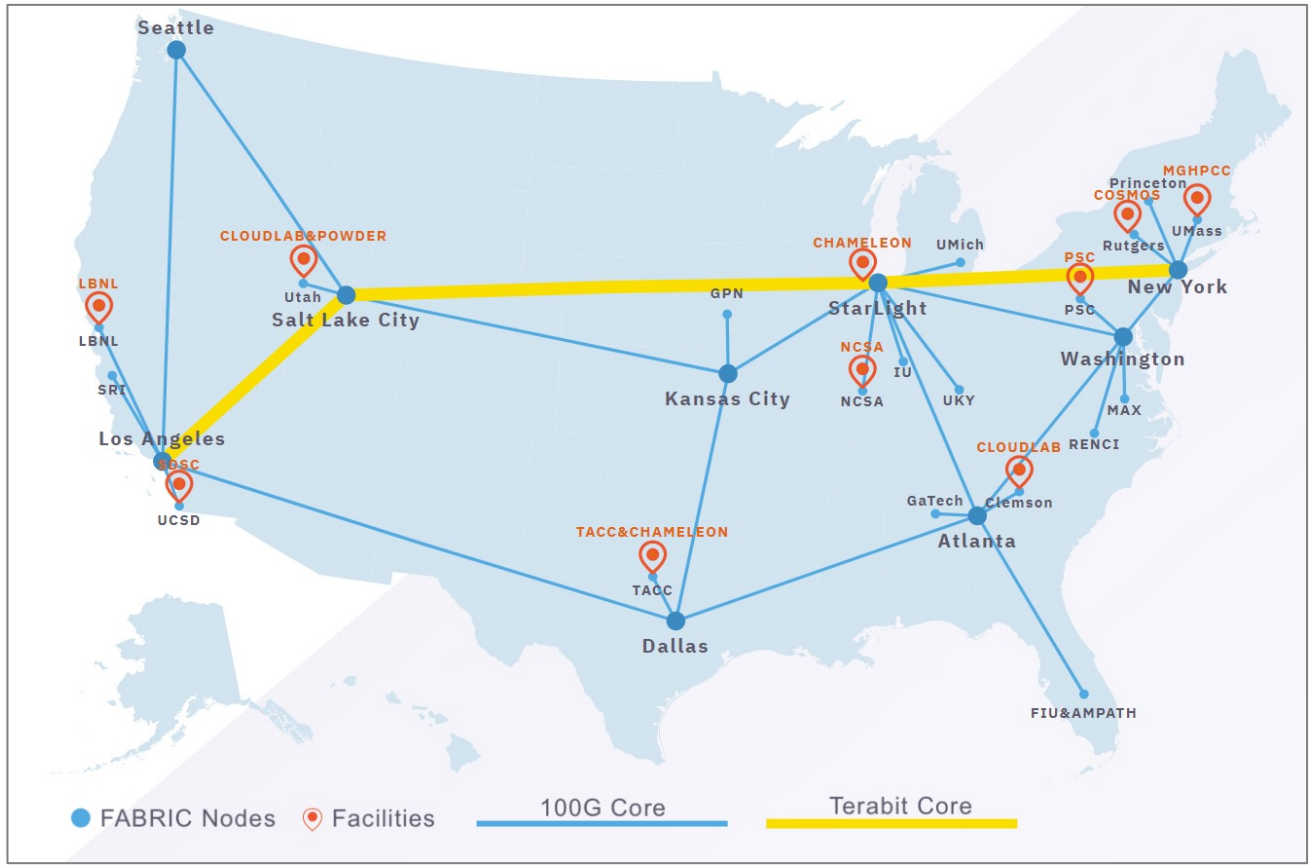

Site 1	Site 2	RTT
TACC (TX)	TACC (TX)	0.3ms
DALL (TX)	TACC (TX)	6ms
DALL (TX)	WASH (DC)	27ms
SALT (UT)	FIU (FL)	44ms
GPN (MO)	DALL (TX)	61ms
UTAH (UT)	WASH (DC)	72ms
GPN (MO)	FIU (FL)	113ms

Additional Slides

- BBR performance on FABRIC
- Performance measurements for a single flow, 0.0046% packet loss rate

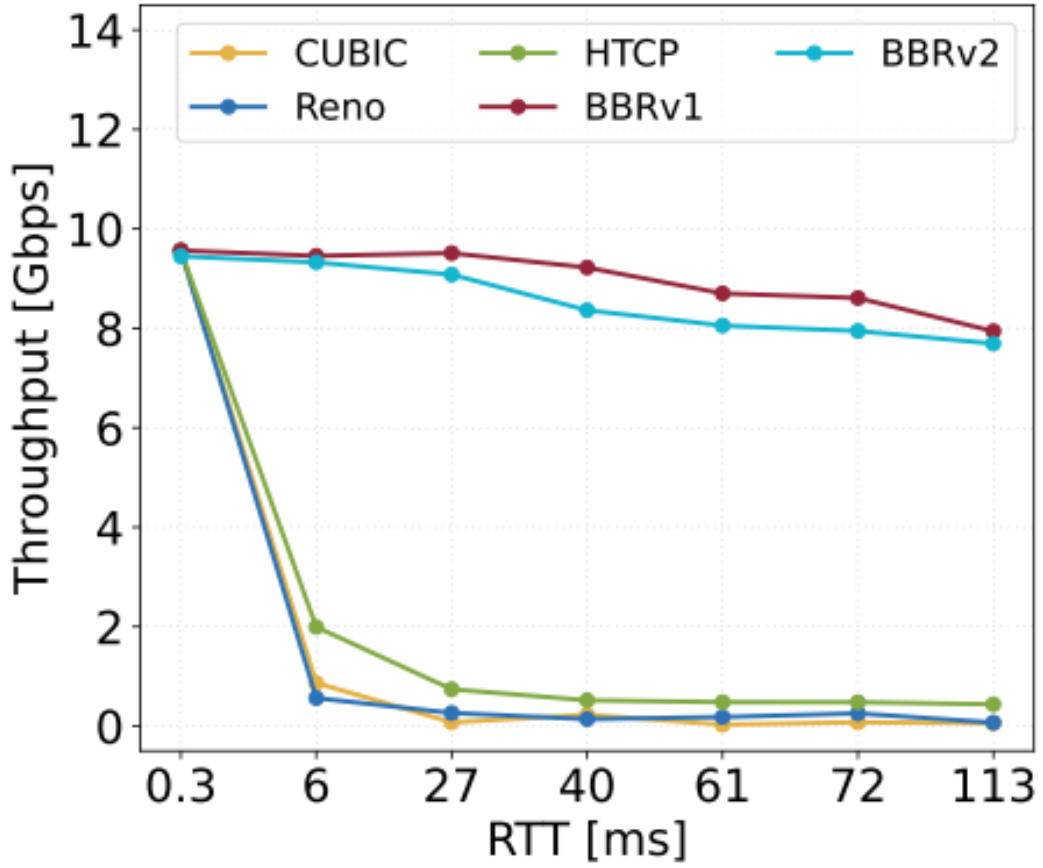

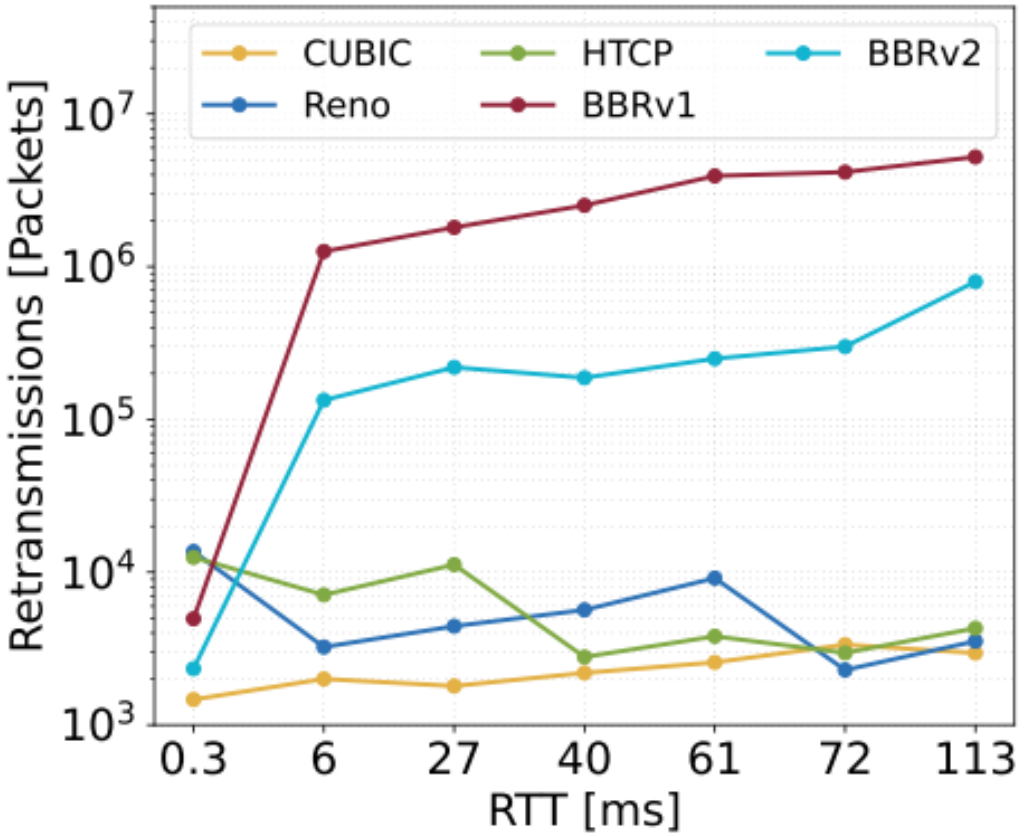

BDP

- Bandwidth = 1Gbps
- RTT = 30ms
- BDP (bytes) = 3,750,000 bytes
- BDP (MB) = 3.57MB