High-speed Networks, Cybersecurity, and Software-defined Networking Workshop

Jorge Crichigno
University of South Carolina

2020 Western Academy Support and Training Conference
Summer Conference
June 15 – June 19

Motivation for High-speed Networks

Motivation

- Science, engineering, mobile applications are generating data at an unprecedented rate
- From large facilities to portable devices, instruments can produce hundreds of terabytes in short periods of time
- Data must be typically transferred across high-throughput high-latency Wide Area Networks (WANs)
- The Energy Science Network (ESnet) is the backbone connecting U.S. national laboratories and research centers

Applications

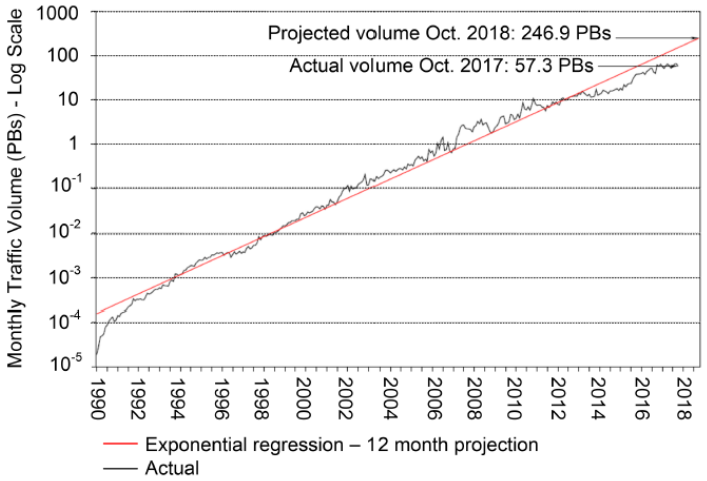

ESnet traffic

Motivation

- A biology experiment using the U.S. National Energy Research Scientific Computing Center (NERSC) resources

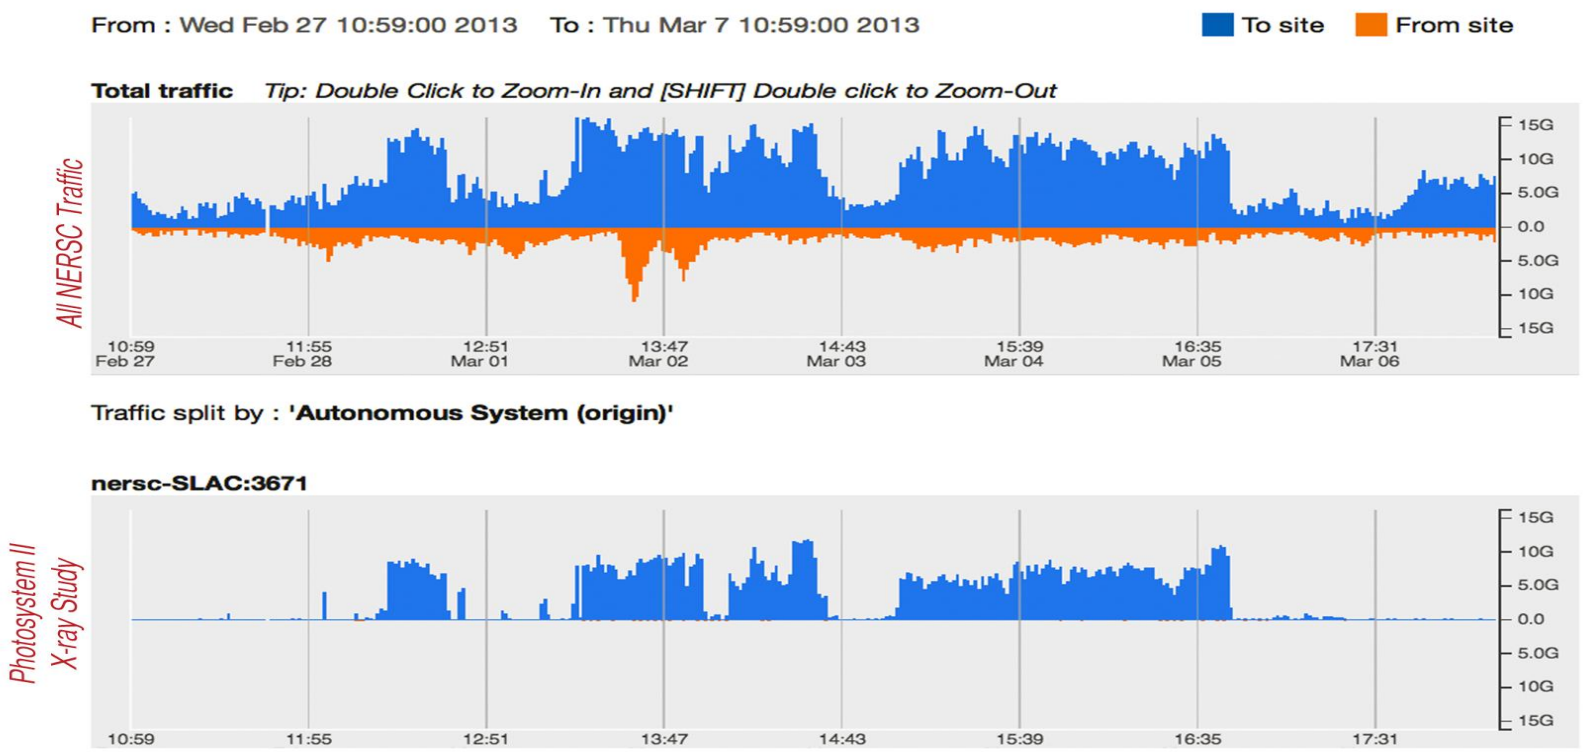

Motivation

- A biology experiment using the U.S. National Energy Research Scientific Computing Center (NERSC) resources

SnapChat Data produced per day worldwide by millions of people

= 38 TB

One Biology experiment by a team of nine scientists:

= 114 TB

(Photosystem II X-Ray Study)

http://www.nature.com/articles/ncomms5371

Motivation

Enterprise network limitations:

- Security appliances (IPS, firewalls, etc.) are CPU-intensive
- Inability of small-buffer routers/switches to absorb traffic bursts
- At best, transfers of big data may last days or even weeks

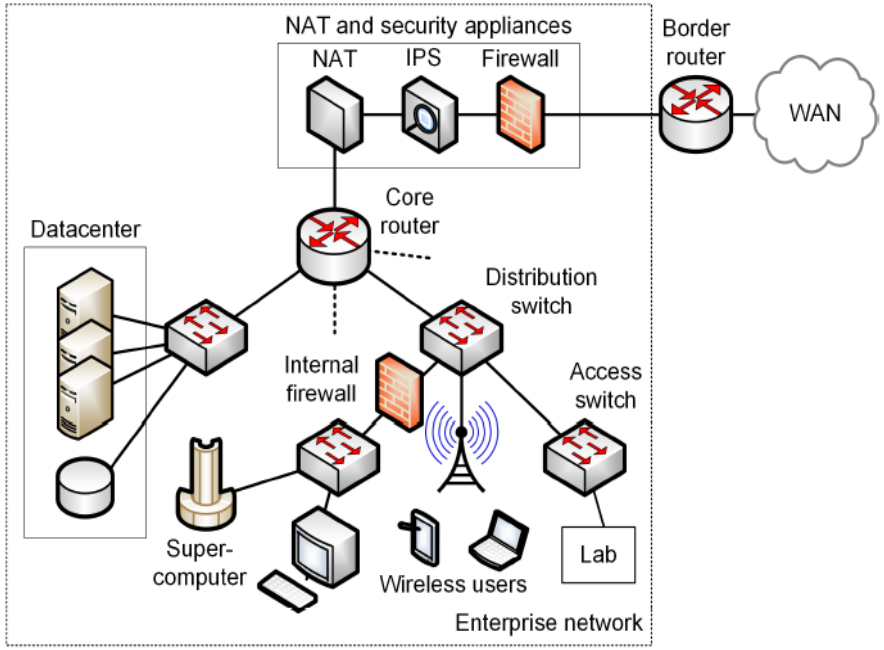

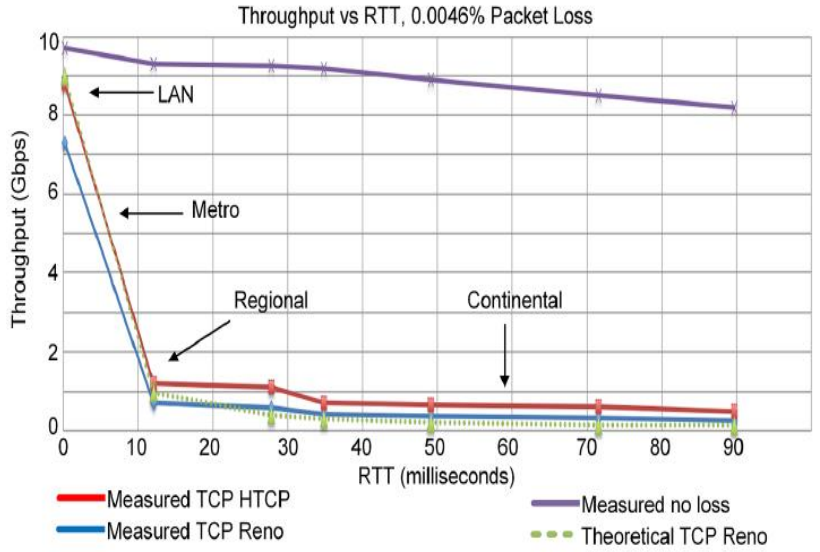

Two devices exchanging data on a 10 Gbps network
Packet loss rate is 1/22,000, or 0.0046%

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

Motivation

There are features in network devices that are important for high performance

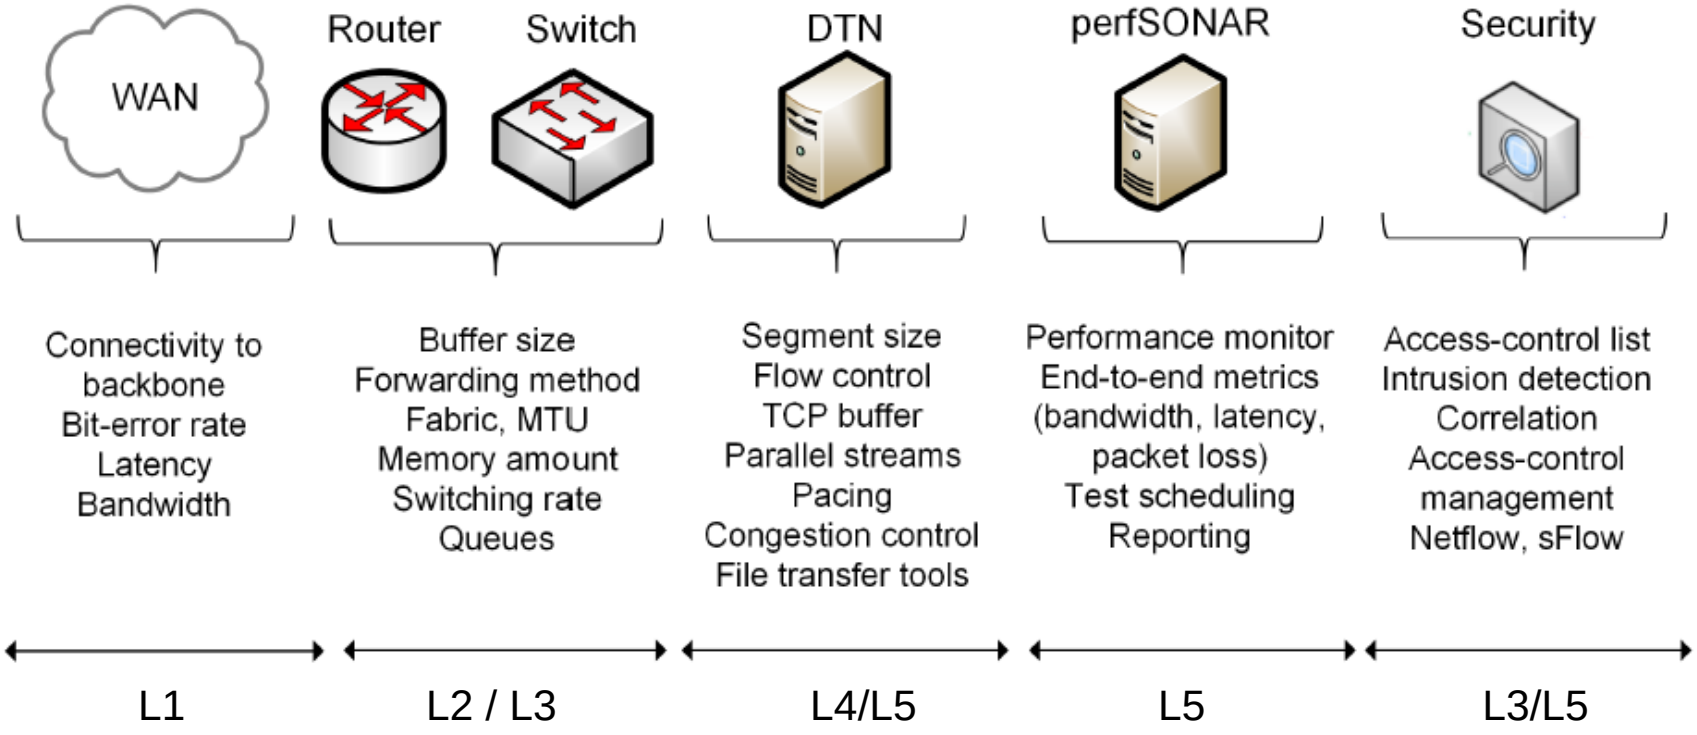

Motivation

How can we teach these topics using a scalable environment?

Requirements

- High-speed networks; speeds of at least 10 Gbps
- Scalability; platform must be capable of cloning pods and expand capacity easily
- Real protocol stack, no simulation
- Free tools

Motivation

Partnering with the Network Development Group (NDG)

Feature	Private Cloud	Public Cloud
Granularity to allocate physical resources	Very granular	Not granular (access to the physical resources requires additional fees)
Easy to create custom pods	Easy	More difficult; hard to design complex topologies
Cost	Cost effective when used extensively	Cost effective for individual / small virtual machines; costly for large virtual machines over time
IT Staff	Higher cost	Lower cost
Application layer for pedagogy, presentation of virtual scenarios	Very flexible	Not flexible; limited to providers' interface, e.g., command-line interface
Time-sharing compute resources	The owner controls who can access resources. Easy to implement time-sharing policies	Cloud provider controls who can access resources (typically, a fee is required per user accessing resources)