AT HOME WITH ENGINEERING EDUCATION

JUNE 22 - 26, 2020

Asee's Virtual Conference

#ASEEVC

#ASEEVC

ASEE'S VIRTUAL CONFERENCE

JUNE 22 - 26, 2020

At Home with Engineering Education

Training and Teaching Students and IT Professionals on High-throughput Networking and Cybersecurity using a Private Cloud

- Jorge Crichigno[1]
- Elias Bou-Harb[2]
- Elie Kfoury[1]
- Jose Gomez[1]
- Antonio Mangino[2]

[1]University of South Carolina
[2]University of Texas at San Antonio

Training and Teaching Students and IT Professionals on High-throughput Networking and Cybersecurity using a Private Cloud

Jorge Crichigno[1], Elias Bou-Harb[2], Elie Kfoury[1], Jose Gomez[1], Antonio Mangino[2]

[1]University of South Carolina, [2]University of Texas at San Antonio

- Building a private cloud
- Virtual labs on high-speed networks and cybersecurity
- Real protocol stacks and live traffic experimentation
- Rates of 50 Gbps, professional tools
- Scalable platform, hundreds of users simultaneously

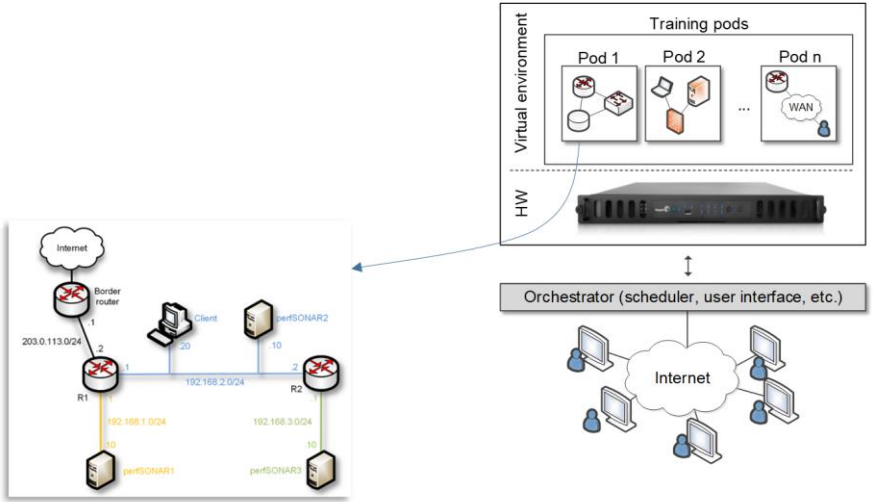

Agenda

Introduction

Motivation

Private cloud

Virtual labs

Comparison private vs public clouds

Conclusion

Introduction

General-purpose enterprise networks can transport basic data, e.g., emails, multimedia, and web content

However, these networks face many challenges when moving terabytes (TB) of scientific data, e.g., genomic, climate, imaging, and high-energy physics

As the popularity of high-speed networks (e.g., Science DMZs) moving data at tens / hundreds of terabits per seconds surges, there is a need for teaching material

The need for trained engineers with the skills to condition these high-performance cyberinfrastructures (CIs) has increased tremendously

Introduction

Concerns raised during the 2017's NSF Campus Cyberinfrastructure PIs meeting

#	Concerns by PIs, Co-PIs, and attendees of 2017 NSF CC* meeting
1	"Very difficult to find, or nonexistent - difficult to retain (CI engineers)"
2	"Largest challenge was in the area of time to hire... ended up taking 10 months… (difficult to find CI engineers with the right skills)"
3	"Candidates should have hands-on knowledge of networking, at least bachelor degree, and certifications in networking and security"
4	"Combination of education and experience"
5	"At least one tour of duty as an intern or apprentice"
6	"System & network engineering, user support experience, good communication (written and presentation)…"
7	"Training in routing and switching (e.g., Juniper, Cisco), a minimal knowledge and/or training in security (e.g., Palo Alto or similar), cabling"
8	"Working knowledge of theory and practice underlying VLAN/LAN/WAN network operations"
9	"Working with researchers to identify areas where their research can benefit from high-end technologies such as HPC, Science DMZ, Data Transfer Node (DTN), Big Data platforms"
10	Difficult to find, preferred qualification: combination of "Bachelor degree" and "certifications in networking and security"

http://www.thequilt.net/wp-content/uploads/NSF-2017-PI-Workshop-CI-Engineer-Survey_v4.pdf

Motivation

Performance of enterprise devices

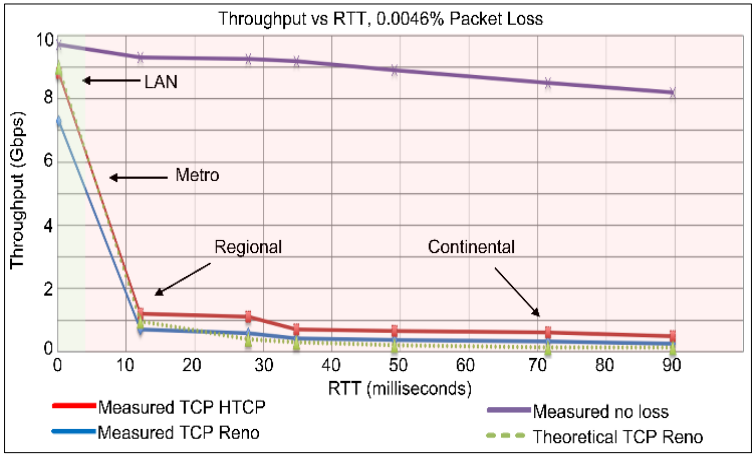

Throughput vs RTT, two devices connected via a 10 Gbps path

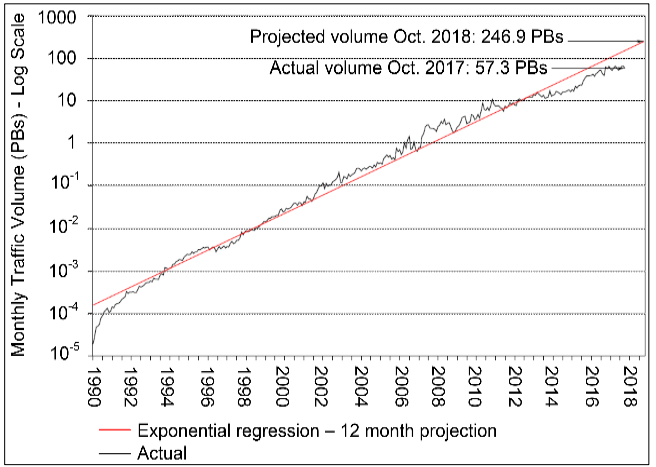

Monthly average traffic volume, ESnet

Motivation

Elements of a high-speed network (science DMZs) to transfer data across a wide are network (WAN)

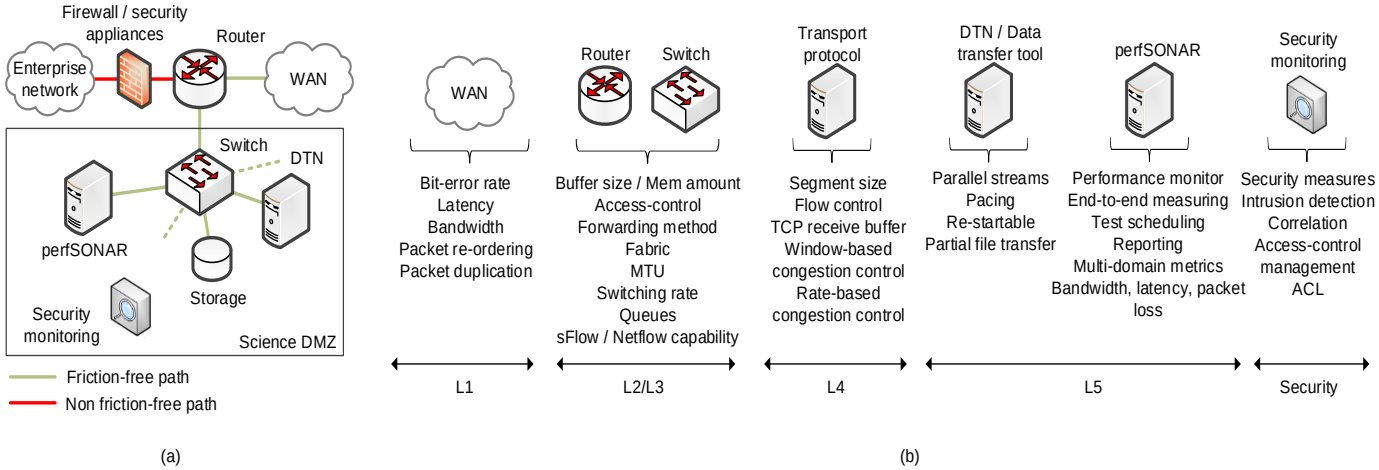

Private Cloud

The project relies on a private cloud

The framework used to develop and run virtual labs is NETLAB, from the Network Development Group (NDG)[1]

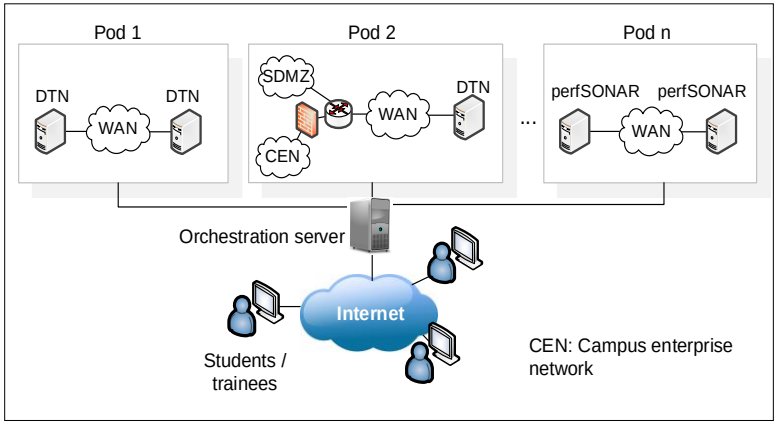

[1] www.netdevgroup.com

Private Cloud

The private cloud used relies on physical servers

The physical resources can be classified in compute capability (CPU cores), storage (non-volatile memory), and RAM memory

Device	Cores	Storage (TBs)	RAM (GB)	Notes
Server 1 (management server)	20	4.8	128	Hosts orchestration server
Server 2 (hosting vLabs pods)	32	4.8	512	Hosts pods' VMs
Server 3 (hosting vLabs pods)	32	1.92	768	Hosts pods' VMs
Server 4 (hosting vLabs pods)	32	1.92	768	Hosts pods' VMs
Total	116	8.08	2,176	

Virtual Labs

Multiple pod types

A pod is a set of equipment need to complete a lab experiment: (1) Network Tools and Protocols (NTP), (2) perfSONAR, (3) Zeek

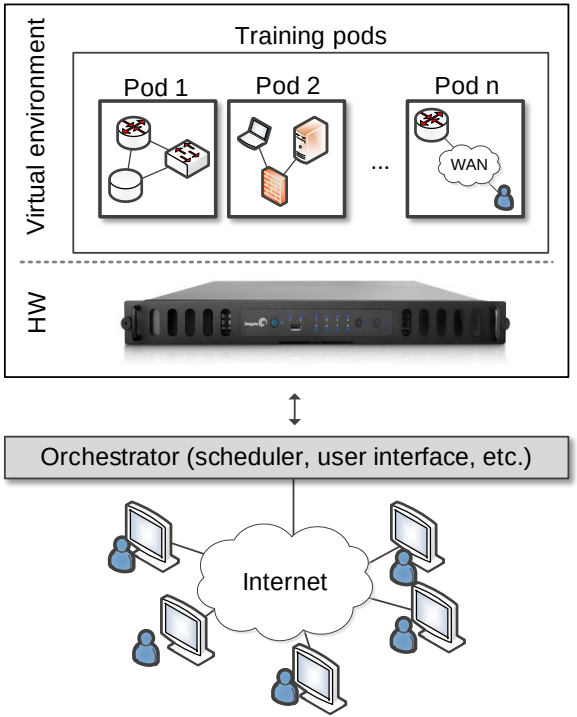

Virtual Labs

Multiple pod types

A pod is a set of equipment need to complete a lab experiment: (1) Network Tools and Protocols (NTP), (2) perfSONAR, (3) Zeek

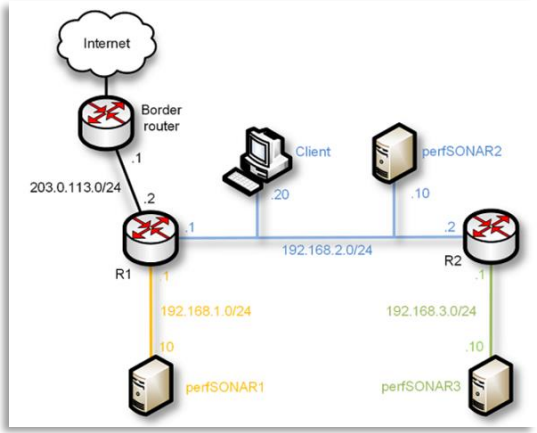

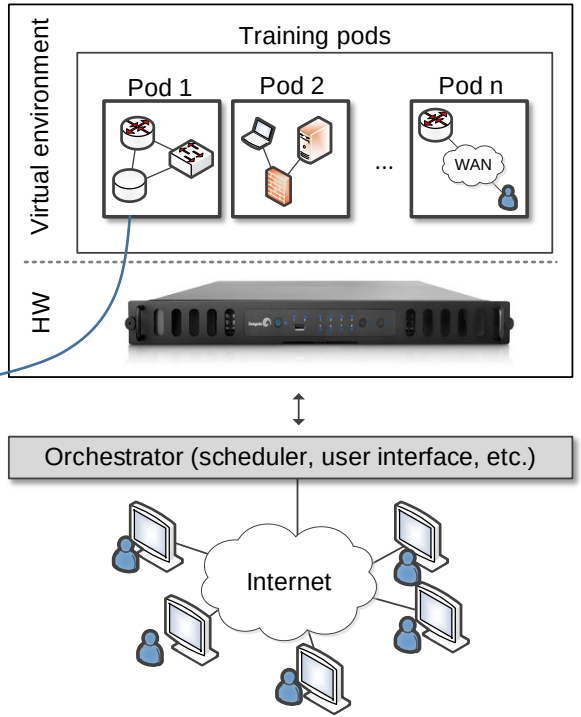

Virtual Labs: NTP

The lab series provides learners an emulated WAN infrastructure operating at high speeds, up to 50 Gbps, and devices running real protocol

Lab 1: Introduction to Mininet
Lab 2: Introduction to iPerf
Lab 3: WANs with latency, Jitter
Lab 4: WANs with Packet Loss, Duplication, Reordering, and Corruption
Lab 5: Setting WAN Bandwidth with Token Bucket Filter (TBF)
Lab 6: Traditional TCP Congestion Control (HTCP, Cubic, Reno)
Lab 7: Rate-based TCP Congestion Control (BBR)
Lab 8: Bandwidth-delay Product and TCP Buffer Size
Lab 9: Enhancing TCP Throughput with Parallel Streams
Lab 10: Measuring TCP Fairness
Lab 11: Router's Buffer Size

Lab 12: TCP Rate Control with Pacing
Lab 13: Impact of Maximum Segment Size on Throughput
Lab 14: Router's Bufferbloat
Lab 15: Analyzing the Impact of Hardware Offloading on TCP Performance
Lab 16: Random Early Detection
Lab 17: Stochastic Fair Queueing
Lab 18: Controlled Delay (CoDel) Active Queue Management
Lab 19: Proportional Integral Controller-Enhanced (PIE)
Lab 20: Classifying TCP traffic using Hierarchical Token Bucket (HTB)

Virtual Labs: NTP

For the Network Tools and Protocols lab series, pods are embedded into Mininet

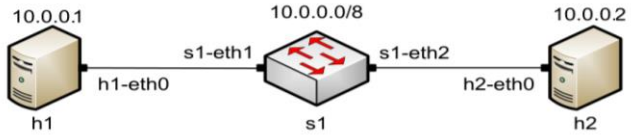

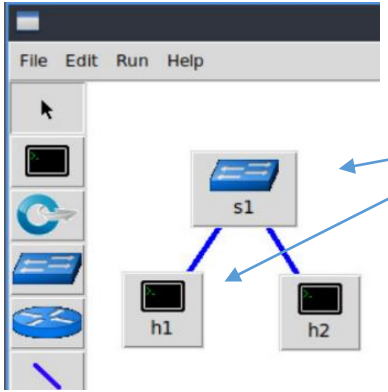

Intel Xeon Gold 6130
CPU, 2.1 GHz

Virtual Labs: perfSONAR

perfSONAR is a network measurement app. designed to monitor end-to-end paths

The lab series enables users to learn perfSONAR on a multi-domain internetwork

Lab 1: Configuring Admin. Information Using perfSONAR Toolkit GUI

Lab 2: PerfSONAR Metrics and Tools

Lab 3: Configuring Regular Tests Using perfSONAR GUI

Lab 4: Configuring Regular Tests Using pScheduler CLI Part I

Lab 5: Configuring Regular Tests Using pScheduler CLI Part II

Lab 6: Bandwidth-delay Product and TCP Buffer Size

Lab 7: Configuring Regular Tests Using a pSConfig Template

Lab 8: perfSONAR Monitoring and Debugging Dashboard

Lab 9: pSConfig Web Administrator

Lab 10: Configuring pScheduler Limits

Virtual Labs: perfSONAR

POD with multiple networks and live traffic to the Internet

Devices running real protocol stack

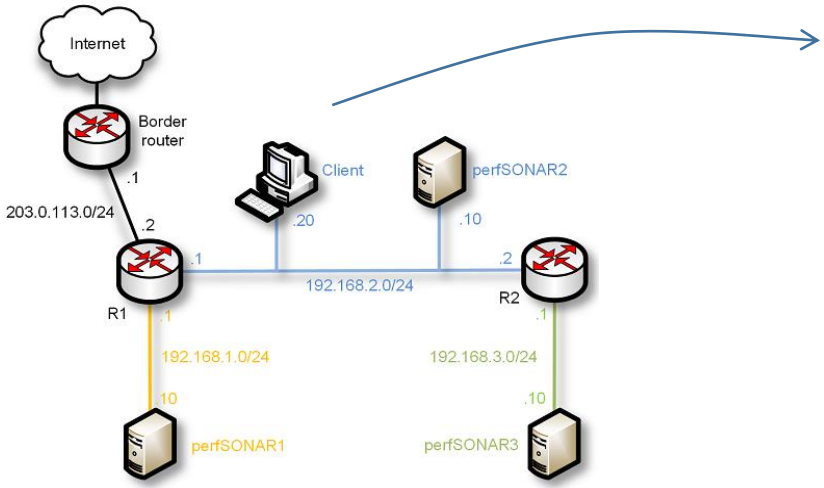

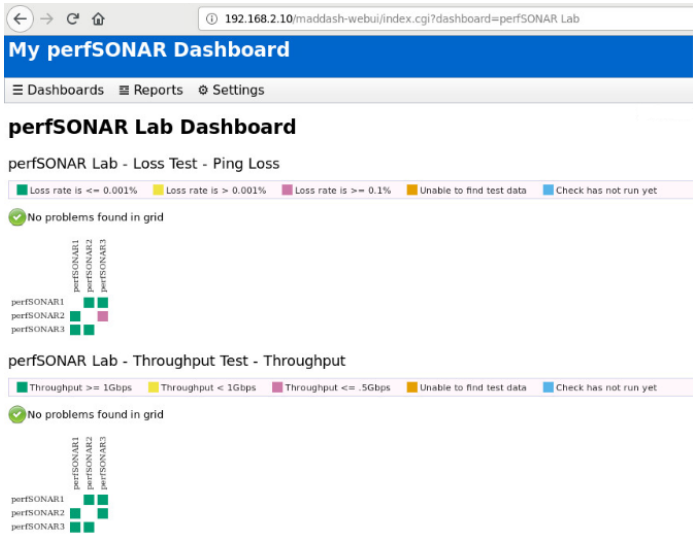

Virtual Labs: Zeek

The Zeek lab series enables users to conduct cyberattacks in a controlled environment, and secure networks using techniques suitable for high speeds

Lab 1: Introduction to the Capabilities of Zeek
Lab 2: An Overview of Zeek Logs
Lab 3: Parsing, Reading and Organizing Zeek
Lab 4: Generating, Capturing and Analyzing Network Scanner Traffic
Lab 5: Generating, Capturing and Analyzing DoS and DDoS-centric Network Traffic
Lab 6: Introduction to Zeek Scripting
Lab 7: Introduction to Zeek Signatures
Lab 8: Advanced Zeek Scripting for Anomaly and Malicious Event Detection
Lab 9: Profiling and Performance Metrics of Zeek
Lab 10: Application of the Zeek IDS for Real-Time Advanced IDS Protection
Lab 11: Preprocessing of Zeek Output Logs for Machine Learning
Lab 12: Developing Machine Learning Classifiers for Anomaly Inference and Classification

Virtual Labs: Zeek

Labs use routers, end devices, and Zeek nodes used as intrusion detection system and intrusion prevention system

Experiments include capturing and analyzing network attacks such as scanner traffic, DoS and DDoS, and more

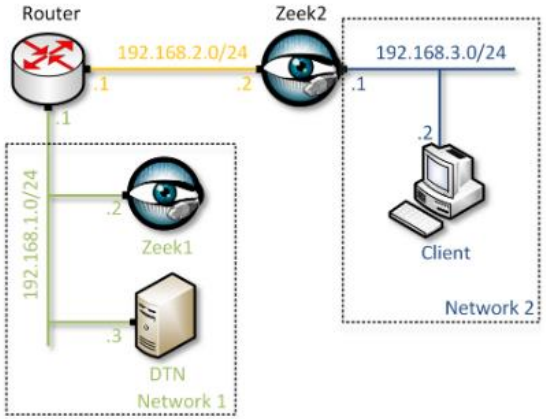

Virtual Labs: Platform Use

Use of the platform – January 1 2019 – December 31 2019

Admin > Usage > Community Usage > List

Community Usage

Search

ID	Name	Reservations Made	Labs Attended	Hours Reserved	Hours Attended
1	default	4690	4483	73739.17	23863.83
	Page Total:	4690	4483	73739.17	23863.83
	Table Total:	4690	4483	73739.17	23863.83

Showing **1** to **1** of **1** items

Comparison Private and Public Cloud

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 and 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)

Conclusion

This presentation describes a project that implemented a private cloud

Virtual labs running on the private cloud are developed for teaching, training, and research on high-speed networks and cybersecurity

The platform supports customized pod and lab designs that emulate complex internetworks operating at up to 50 Gbps

The scalability of the platform permits the simultaneous on-demand deployment of hundreds of emulated WANs / LANs, serving hundreds of users at the same time

The material and platform have been used to support academic courses, self-pace training of professional IT staff, and workshops across the country (30+ states)