Tofino Pods for Teaching and Research

Jorge Crichigno, Elie Kfoury, Jose Gomez
University of South Carolina
http://ce.sc.edu/cyberinfra
jcrichigno@cec.sc.edu, ekfoury@email.sc.edu, gomezgaj@email.sc.edu

February 16th, 23rd, 2022
Online

Virtual Platform

- The Academic Cloud provides remote-access capability to laboratory equipment via the Internet
- It pools and shares resources (CPU, memory, storage, switch) needed to run virtual labs
- USC works closely with the Network Development Group (NDG)[1]
- Libraries incorporate performance and realism along with NETLAB's features

1. Network Development Group (NDG). https://netdevgroup.com

Virtual Platform

- Features

Feature	Description
Performance	High-performance (e.g., 40/100 Gbps)
Functional realism	Same functionality as IT hardware in a real deployment; execution of real code
Traffic realism	Devices within the virtual lab generate and receive real, interactive network traffic (including traffic to/from the Internet)
Presentation layer	Navigating through an experiment is easy for an inexperience learner
Topology flexibility	Easy to create topologies for more complex experiments, including inter-connecting heterogeneous VMs
Physical devices	Physical hardware can be integrated into pods
Timesharing / calendar	Reservations are made through NETLAB calendar interface

Inside the Datacenter

- Hosts 1-n store virtual machines (VMs) for virtual labs
- Management server runs vCenter, NETLAB+
- Partnership with NDG (NETLAB+)[1] and VMware[2] (ESXi, vCenter)

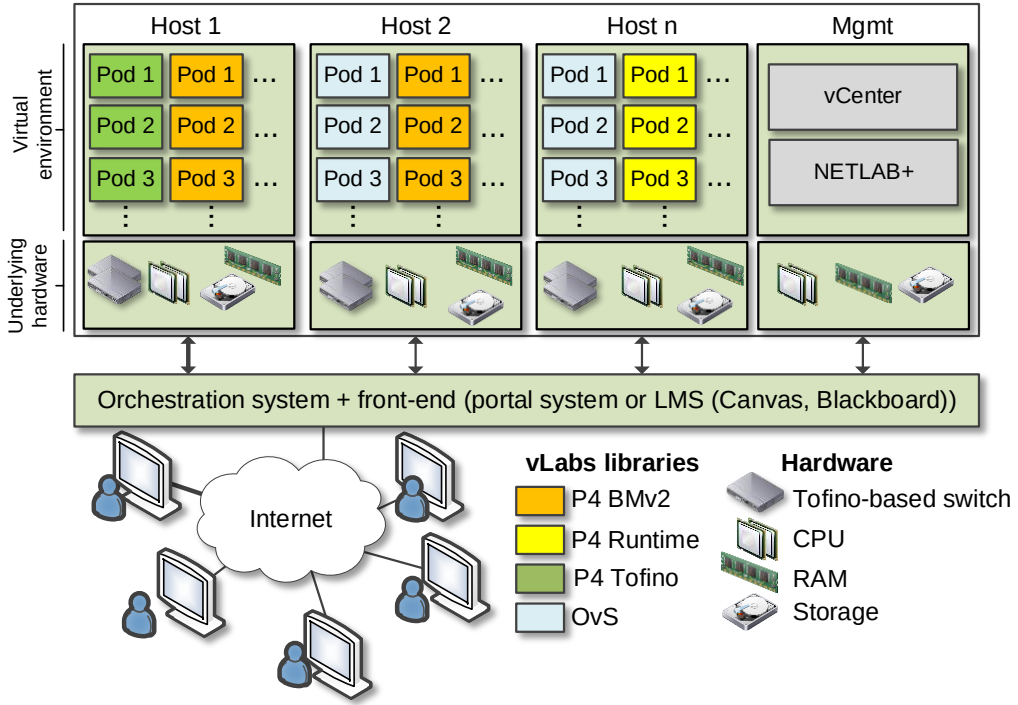

1. Network Development Group (NDG). https://netdevgroup.com
2. VMware. https://www.vmware.com

Virtual Platform

- The environment is fully integrated (topology, lab experiments, etc.)
- Tofino switches come “empty”
- Considerable effort is required to
  - Acquire, Install, and learn the operating system (e.g., Open Network Linux)
  - Acquire, install, and learn Capilano compiler / Software Development Environment
  - Deploy and test fibers
  - Configure topologies
  - Agreement with Intel
- Tofino pods help USC to onboard new students and researchers
  - They can immediately focus on learning how to program the switch
  - Share environment

Tofino Programmable ASIC

- Tofino uses the Tofino Native Architecture (TNA)
- P4 programs are written in $P4_{16}$
- The switch model is Wedge 100BF-32X from Edgecore
- This switch has 32 x 100G QSFP28 switch ports

Tofino Model

- Tofino Model is a software switch used for testing and troubleshooting P4 programs
- The same code that runs on a Tofino model can be ported to a physical switch
- The model allows tracking the lifecycle of a packet traversing the pipeline
- The model has the same purpose as the BMv2 switch

Development Environment

- The user reserves a pod through the web calendar interface
- The pod consists of a physical switch, Tofino Model, and two virtual machines

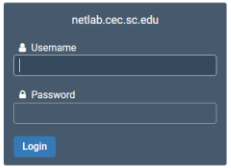

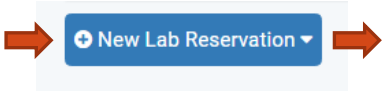

Cyberinfrastructure Lab @ UofSC

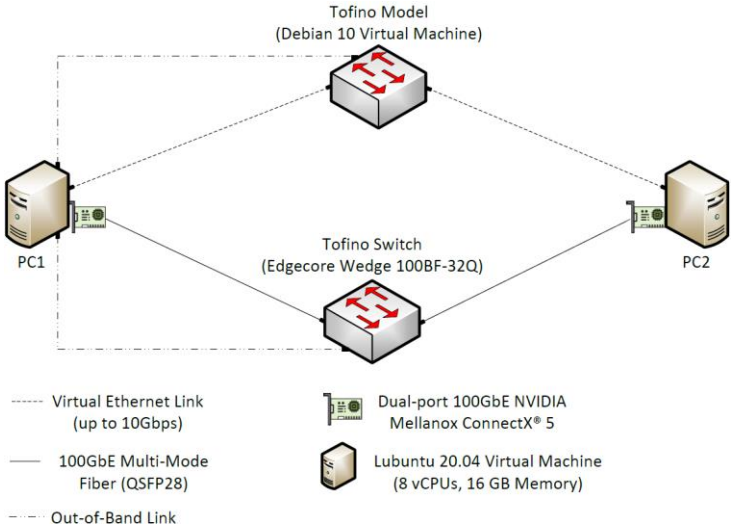

[1]www.netdevgroup.com

Introduction to P4 on Tofino

Lab experiments

Lab 1: Introduction to P4 and BMv2
Lab 2: P4 Program Building Blocks
Lab 3: Parser Implementation
Lab 4: Introduction to Match-action Tables (Part 1)
Lab 5: Introduction to Match-action Tables (Part 2)
Lab 6: Populating and Managing Match-action Tables
Lab 7: Checksum Recalculation and Packet Deparsing

Exercises

Exercise 1: Compiling and Testing a P4 Program
Exercise 2: Parsing UDP and RTP
Exercise 3: Building a Simplified NAT
Exercise 4: Configuring Tables at Runtime
Exercise 5: Building a Packet Reflector

Demo

Demo

- Tofino Model and ASIC running on NetLab
- Programmer can select the target (Tofino model for debugging; physical switch for performance)