

A Cloud System for Teaching and Research on P4 Programmable Data Plane

Elie Kfoury, Jose Gomez, Ali AlSabeih, Jorge Crichigno
College of Engineering and Computing
University of South Carolina
<http://ce.sc.edu/cyberinfra>

KNIT 5: A FABRIC Community Workshop
Northwestern University, Wieboldt Hall
September 20 – 22, 2022
Chicago, IL



Supported by NSF 2118311
“CyberTraining: Implementation: Small: Cybertraining on P4 Programmable Devices
using an Online Scalable Platform”



Agenda

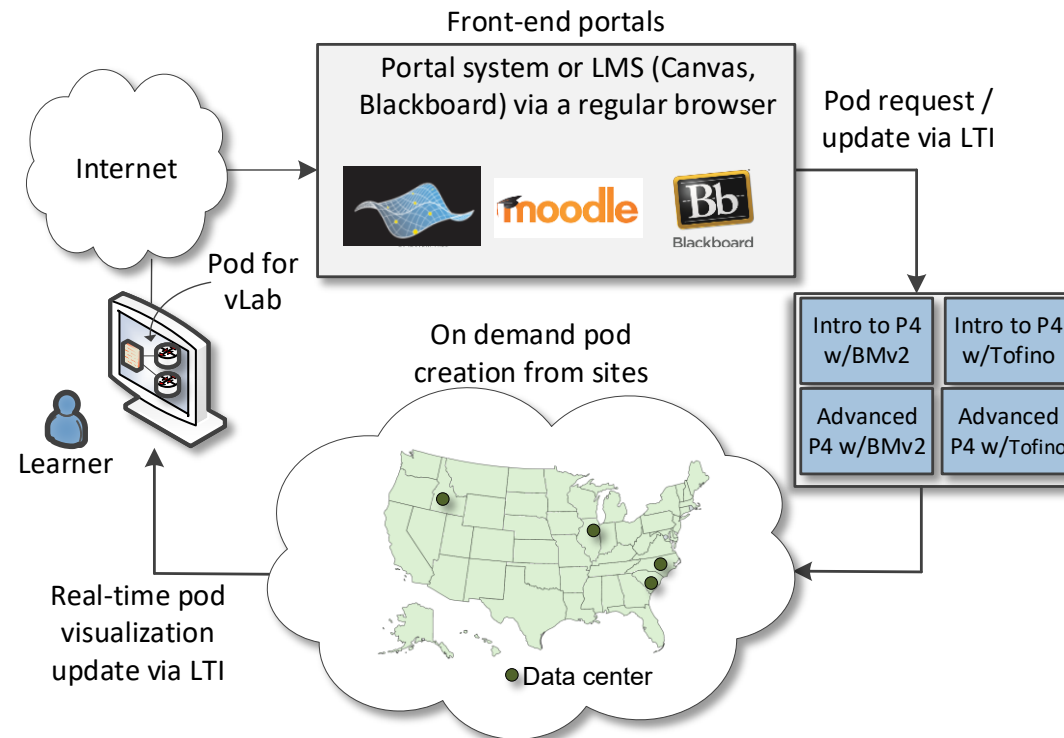
- Motivation for virtual labs
- Platform
- Libraries
- P4 Tofino Library
- Topologies used by the research community
- Conclusion

Motivation for Virtual Labs

- IT curriculum should emphasize “learning IT core concepts with authentic practice¹”
 - “It is not enough to simply attend courses and read books”
- Disadvantages of physical labs
 - Difficult to scale
 - Expensive (space, maintenance, staff)
 - Since COVID-19 emerged, the capacity of labs has been further reduced

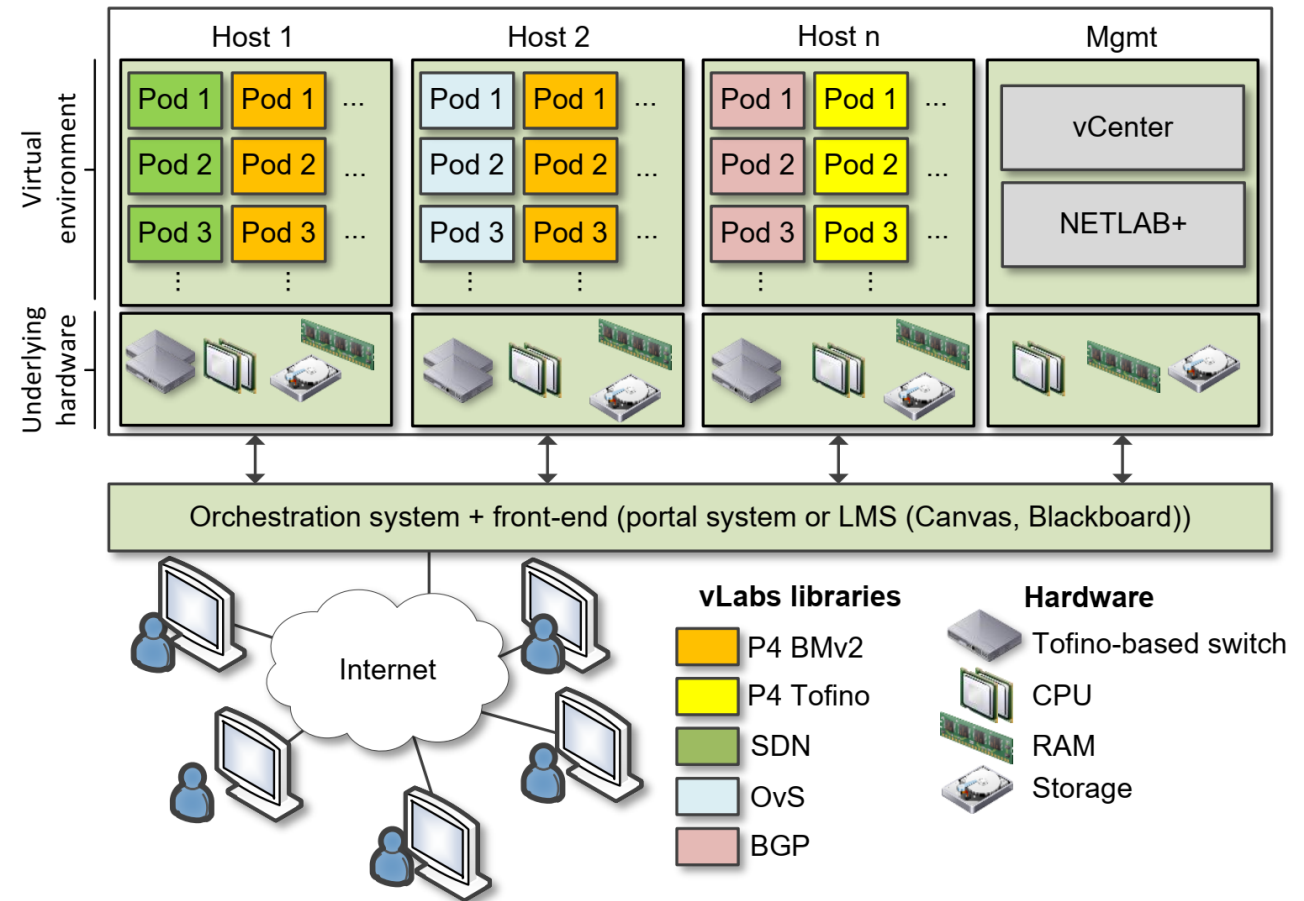
Motivation for Virtual Labs

- The University of South Carolina (USC) (SC), the Network Development Group (NDG) (NC), and Stanly Community College (SCC) (NC) have deployed an Academic Cloud
 - Virtual labs on P4, routing, high-speed networks (USC)
 - Remote-access capability to lab equipment via Internet
 - Shared resources (CPU, memory, storage) from four data centers



USC Data Center

- Hosts 1-n store virtual machines (VMs) for virtual labs
- Management server runs vCenter, Management Software (NETLAB+)
- Partnership with Network Development Group (NDG)¹



Libraries

- A library consists of between 10-20 lab experiments
- Each lab experiment includes a detailed, step by step manual
- Once a learner completes all experiments, the learner acquires significant knowledge and hands-on expertise, and may earn an academic credential or certificate
- Information about libraries are available at <http://ce.sc.edu/cyberinfra/cybertraining.html>

Library on Introduction to P4 with BMv2

Experiments

- Lab 1: Introduction to Mininet
- Lab 2: Introduction to P4 and BMv2
- Lab 3: P4 Program Building Blocks
- Lab 4: Parser Implementation
- Lab 5: Introduction to Match-action Tables (Part 1)
- Lab 6: Introduction to Match-action Tables (Part 2)
- Lab 7: Populating / Managing Match-action Tables
- Lab 8: Checksum Recalculation and Deparsing

Exercises

- Exercise 1: Building a Basic Topology
- Exercise 2: Compiling and Testing a P4 Program
- Exercise 3: Parsing UDP and RTP
- Exercise 4: Building a Simplified NAT
- Exercise 5: Configuring Tables at Runtime
- Exercise 6: Building a Packet Reflector

Library on P4 Applications, Stateful Elements, and Custom Packet Processing


Experiments

- Lab 1: Introduction to Mininet
- Lab 2: Introduction to P4 and BMv2
- Lab 3: P4 Program Building Blocks
- Lab 4: Defining and processing custom headers
- Lab 5: Monitoring the Switch's Queue using Standard Metadata
- Lab 6: Collecting Queueing Statistics using a Header Stack
- Lab 7: Measuring Flow Statistics using Direct and Indirect Counters
- Lab 8: Rerouting Traffic using Meters
- Lab 9: Storing Arbitrary Data using Registers
- Lab 10: Calculating Packets Interarrival Time w/ Hashes and Registers
- Lab 11: Generating Notification Messages from the Data Plane

Library on P4 Programmable Data Plane with Tofino

Experiments

- Lab 1: Introduction to P4 and Tofino
- Lab 2: Introduction to P4 Tofino Software Development Environment
- Lab 3: Parser Implementation
- Lab 4: Introduction to Match-Action Tables
- Lab 5: Populating and Managing Match-Action Tables at Runtime
- Lab 6: Checksum Recalculation and Packet Deparsing



Library on P4 Programmable Data Plane with Tofino

- <https://netlab2.cec.sc.edu/>

The screenshot shows a web browser window with the address bar displaying "Not secure | https://10.173.85.50/my-netlab-i.cgi". The page header includes the University of South Carolina logo and navigation links for Home, Schedule, Manage, Help, and a user profile for jcrichigno. The main content area is titled "Scheduled Lab Reservations" and contains a message: "You have no scheduled lab reservations. Select from the Schedule menu above to add reservations." A blue button labeled "+ New Lab Reservation" is visible at the bottom. A dropdown menu is open over the "Schedule" link, listing options: "View or Cancel Lab Reservations", "Schedule Instructor-Led Training", "Schedule Lab for Myself", "Schedule Lab for a Student", and "Schedule Lab for a Team". An orange arrow points from the message area to the "Schedule" dropdown menu.

Library on P4 Programmable Data Plane with Tofino

Not secure | https://10.173.85.50/makeres-i.cgi?res_type=l

UNIVERSITY OF
SOUTH CAROLINA

Home jcrichigno

Select Content

Multiple course topics are available in this class. Please select one.



Intro to P4 Programmable Data Planes
This pod is an introduction to P4 programmable data planes

Introduction to P4 with Tofino
This pod provides P4 training using the Intel Tofino switch

P4 Applications and Custom Processing
This lab series presents P4 applications, stateful elements, and custom packet processing


Previous Cancel






Library on P4 Programmable Data Plane with Tofino

Not secure | <https://10.173.85.50/makeres-i.cgi>  

UNIVERSITY OF SOUTH CAROLINA [Home](#) [jcrichigno](#)

Select Content (Introduction to P4 with Tofino) > Select Lab

 This pod provides P4 training using the Intel Tofino switch

Lab Name	Action
Lab 1: Introduction to P4 and Tofino	
Lab 2: Introduction to P4 Tofino Software Development Environment (SDE)	
Lab 3: Parser Implementation	
Lab 4: Introduction to Match-action Tables	
Lab 5: Populating and Managing Match-action Tables at Runtime	



Library on P4 Programmable Data Plane with Tofino

Not secure | <https://10.173.85.50/resplan.cgi>

UNIVERSITY OF SOUTH CAROLINA

Home jcrichigno

Select Content (Introduction to P4 with Tofino) > Select Lab (Lab 2: Introduction to P4 Tofino Software Development Environment (SDE)) >

Reserve Pod

Pod Scheduler

September 2022

Sun	Mon	Tue	Wed	Thu	Fri	Sat
28	29	30	31	1	2	3
4	5	6	7	8	9	10
11	12	13	14	15	16	17
18	19	20	21	22	23	24
25	26	27	28	29	30	1

Selected Day: September 17, 2022

Current Time: 16:55 Eastern Time (US & Canada)

	Tofino_H1_pod1	Tofino_H1_pod3	Tofino_H2_pod4	Tofino_H2_pod5
16:00				
17:00				

Library on P4 Programmable Data Plane with Tofino

Not secure | <https://10.173.85.50/resbook.cgi>

UNIVERSITY OF SOUTH CAROLINA

Home jrichigno

Select Content (Introduction to P4 with Tofino) > Select Lab (Lab 2: Introduction to P4 Tofino Software Development Environment (SDE)) > Reserve Pod (Tofino_H2_pod4) > Settings

Add Reservation

Pod Tofino_H2_pod4

Reservation Type Instructor Private Reservation


Reserve For Jorge Crichigno

Lab Exercise Lab 2: Introduction to P4 Tofino Software Development Environment (SDE)

Lab Device Configurations load preset configurations for exercise [Settings](#)

Time Zone Eastern Time (US & Canada)

Start Time 2022-09-17 16:56

End Time 2022-09-17 20:00 






Length of Reservation 2 hrs., 53 mins.

Library on P4 Programmable Data Plane with Tofino

Not secure | <https://10.173.85.50/my-netlab-i.cgi> 🔍 📄 ☆

UNIVERSITY OF SOUTH CAROLINA 🏠 Home 📅 Schedule ▾ 🔧 Manage ▾ 🔗 Help 👤 jcrichigno ▾

📅 Lab Reservations Search

ID	Date/Time	Description	Pod
298	 2022-09-17 16:57  2022-09-17 20:00  2 hrs., 52 mins. 	Class: P4 Course Lab: Lab 2: Introduction to P4 Tofino Software Development Environment (SDE) Type: Instructor User: Jorge Crichigno	Tofino_H2_pod4 

Showing 1 to 1 of 1 items

[+ New Lab Reservation ▾](#)


Library on P4 Programmable Data Plane with Tofino

Not secure | <https://10.173.85.50/lab.cgi> 🔍 🔗 ☆

UNIVERSITY OF CAROLINA Home Pod Reser

Tofino_H2_pod4 > Reservation 298 > Lab 2: Introduction to P4 Tofino Software Development Environment (SDE)

Content Status Tofino Switch PC 1 PC 2 Tofino Model



Library on P4 Programmable Data Plane with Tofino

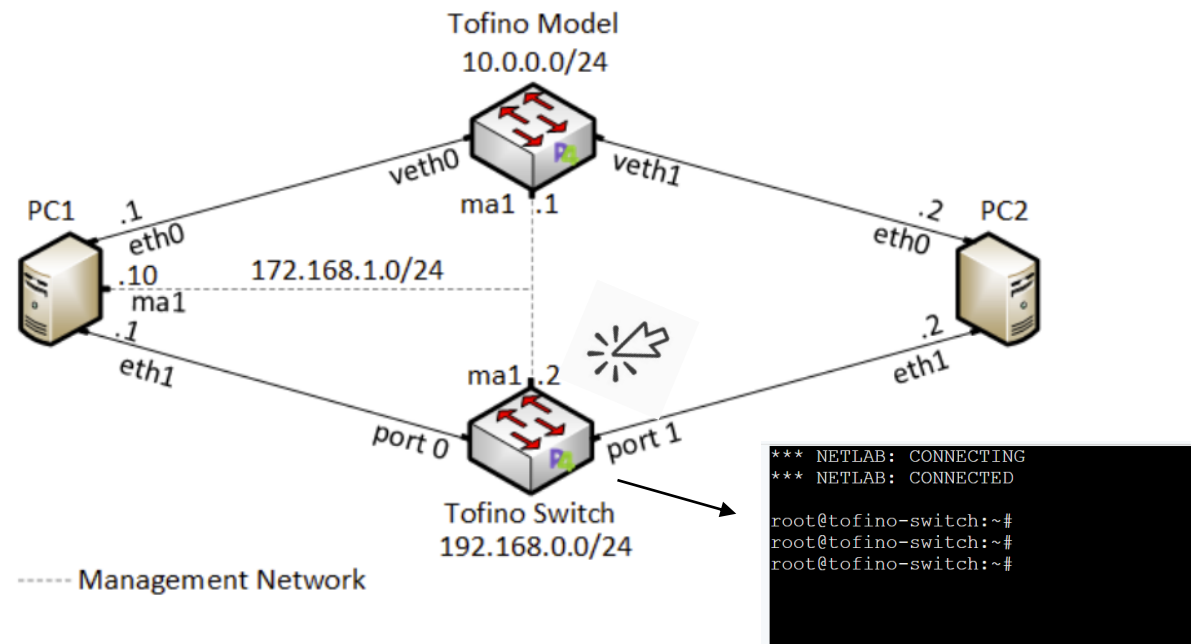
Not secure | <https://10.173.85.50/lab.cgi>

UNIVERSITY OF
CAROLINA

Home Pod Reser

Tofino_H2_pod4 > Reservation 298 > Lab 2: Introduction to P4 Tofino Software Development Environment (SDE)

Content Status Tofino Switch PC 1 PC 2 Tofino Model



Library on P4 Programmable Data Plane with Tofino

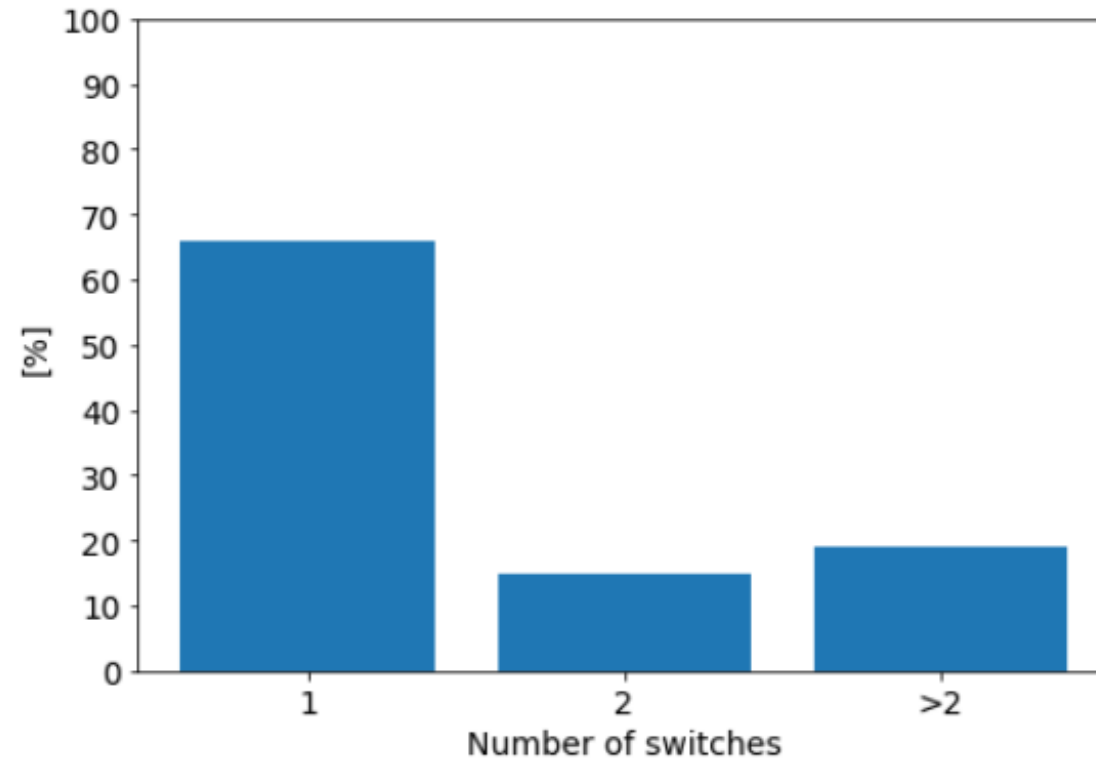
- Topology complexity
 - 6.4 Tbps Tofino programmable switch
 - Tofino model for debugging (trace execution in the data plane)
 - Servers to send/receive data to/from the switch/other servers
 - Multi-mode fiber
 - QSFP28+ transceivers
 - Open Network Linux (ONL) (control plane)
 - Software Development Environment (SDE) from Intel (control plane)
 - Sample P4 codes for each lab (data plane)
 - Laboratory experiments with step-by-step directions (thousands of development hours)
- Logistic challenges
 - NDA with Intel, lawyers' agreement
 - Procurement process
 - Physical hardware, rack space, data center, etc.
 - Software tools, SDE, operating system, etc.

Literature Survey: Topologies used for Research

- 293 papers from 2014 to 2022
- Determine the topologies that are commonly used by experimenters
- Determine the devices that are used

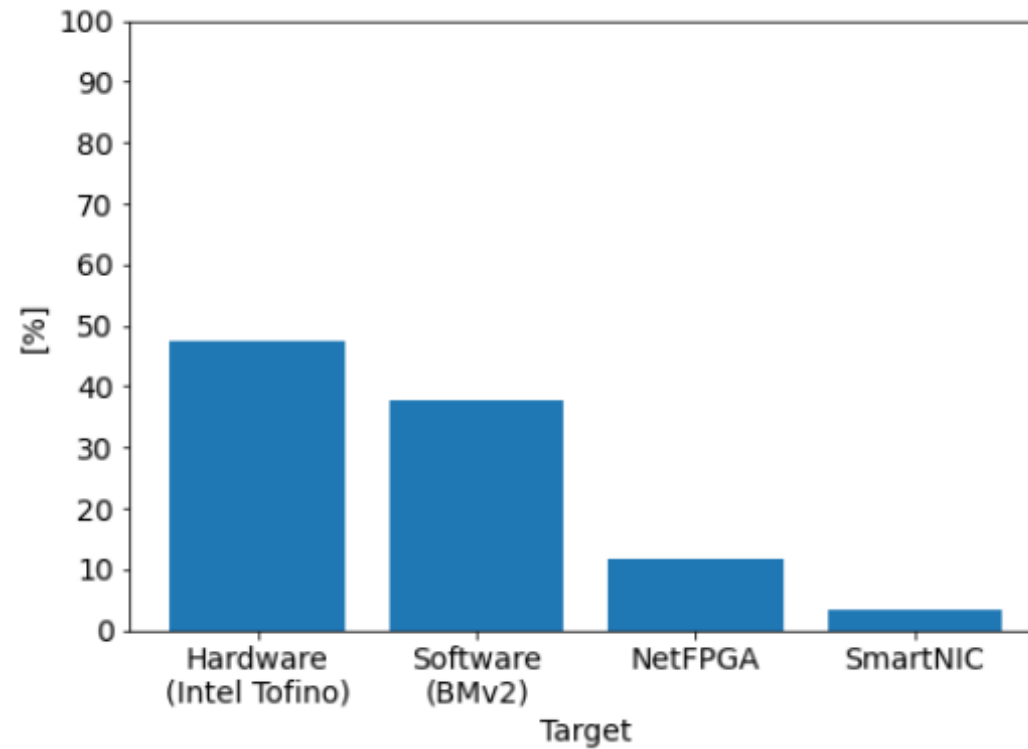
Literature Survey: Topologies used for Research

- Number of switches



Literature Survey: Topologies used for Research

- Targets



Literature Survey: Topologies used for Research


- Number of switches and hosts

		Number of hosts				
			0	1	2	3
Number of switches	1	8	2	52	7	31
	2	1	2	11	1	8
	3	0	1	0	1	3
	≥ 4	0	0	4	0	19

Literature Survey: Topologies used for Research

- Papers that use a 2-host, 1-switch topology

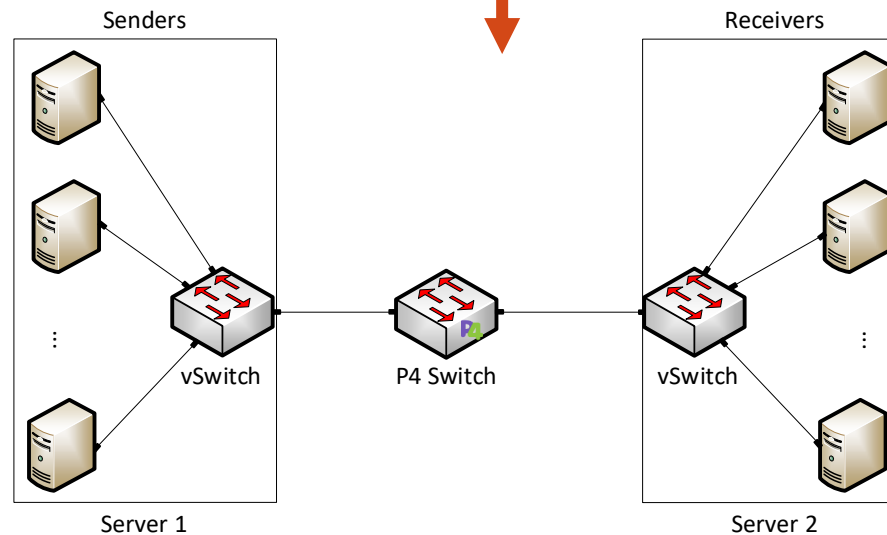
		Number of hosts					
		0	1	2	3	≥ 4	
Number of switches	1	8	2	52	7	31	
	2	1	2	11	1	8	
	3	0	1	0	1	3	
	≥ 4	0	0	4	0	19	



Literature Survey: Topologies used for Research

- 25% of the papers rely on virtualization to add end devices

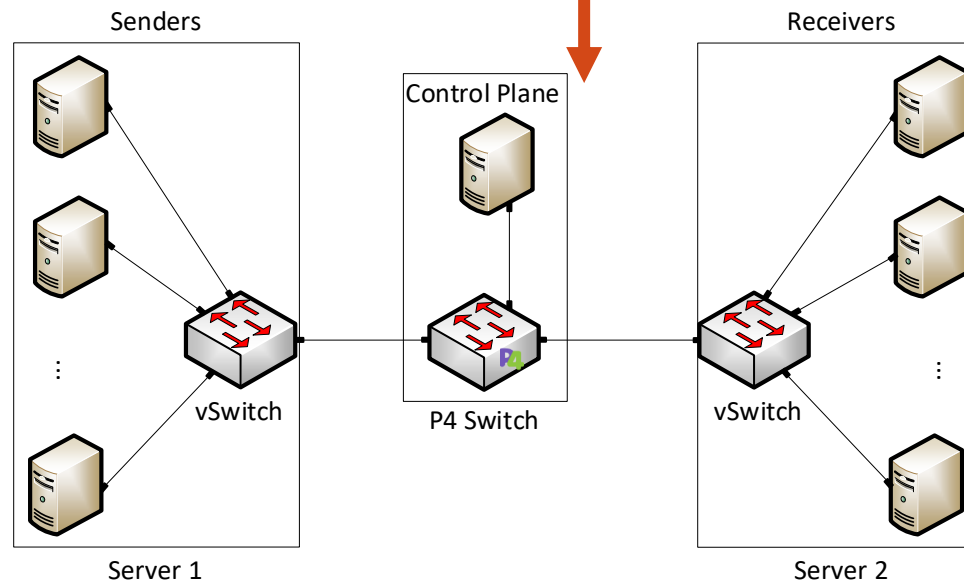
		Number of hosts					
Number of switches	switches	0	1	2	3	≥4	
	1	8	2	52	7	31	
	2	1	2	11	1	8	
	3	0	1	0	1	3	
	≥4	0	0	4	0	19	



Literature Survey: Topologies used for Research

- Topology supported by the current Tofino pod

		Number of hosts				
		0	1	2	3	≥ 4
Number of switches	1	8	2	52	7	31
	2	1	2	11	1	8
	3	0	1	0	1	3
	≥ 4	0	0	4	0	19



USC, UTSA, USF Research Work

- “INC: In-Network Classification of Botnet Propagation at Line Rate”
- “P4DDPI: Securing P4-Programmable Data Plane Networks via DNS Deep Packet Inspection”
- “Dynamic Router's Buffer Sizing using Passive Measurements and P4 Programmable Switches”
- “On Offloading Network Forensic Analytics to Programmable Data Plane Switches”
- “Coarse Estimation of Bottleneck Router's Buffer Size for Heterogeneous TCP Sources”
- “Offloading Media Traffic to Programmable Data Plane Switches”
- “Towards a Unified In-Network DDoS Detection and Mitigation Strategy”
- “Enabling TCP Pacing using Programmable Data Plane Switches”
- “An Exhaustive Survey on P4 Programmable Data Plane Switches: Taxonomy, Applications, Challenges, and Future Trends”
- “A Survey on TCP Enhancements using P4-programmable Devices”
- “A Survey on Security Applications of P4 Programmable Switches and a STRIDE-based Vulnerability Assessment”

Conclusion

- The primary use of the platform is for teaching
- The platform is used at USC, UTSA, and USF for research
 - Time to research is shortened
 - Scalable, cost efficient
 - Resources are shared

UNIVERSITY OF SOUTH CAROLINA [Home](#) [administrator](#)

[Admin](#) > [Usage](#) > [Community Usage](#) > List

Community Usage

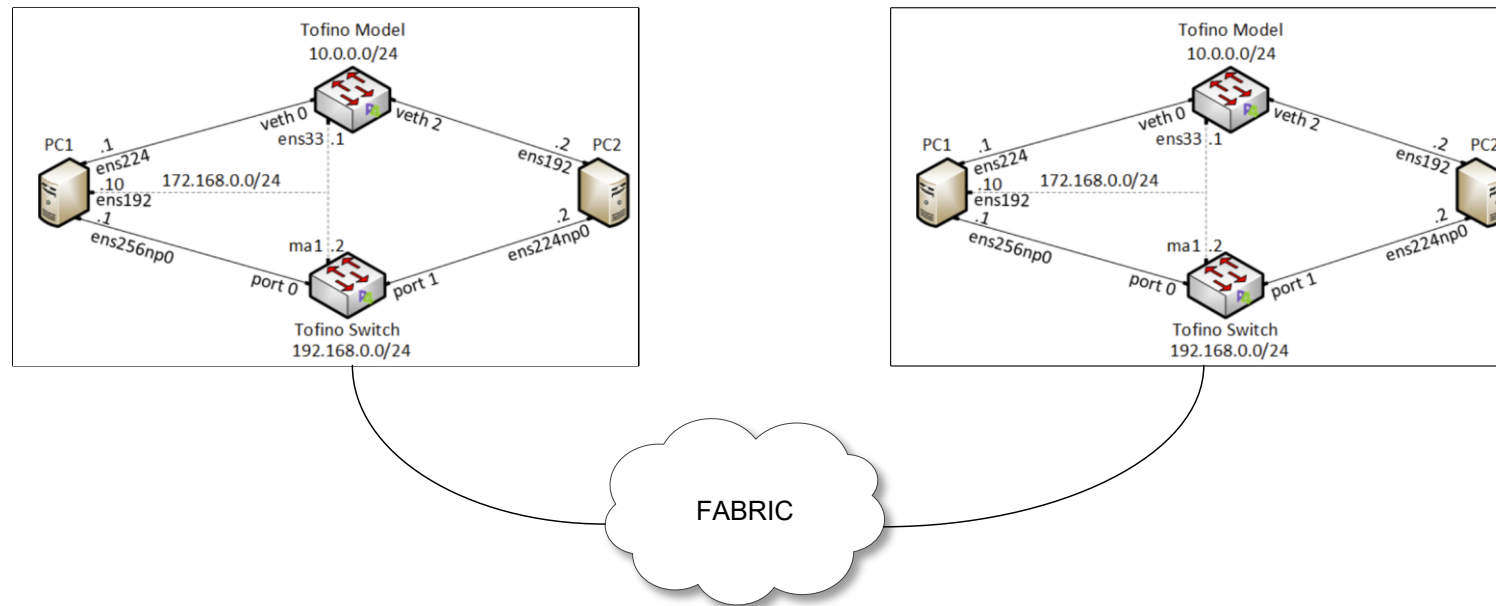
ID	Name	Reservations Made	Labs Attended	Hours Reserved	Hours Attended
1	default	21578	20577	233195.61	105614.35
Page Total:		21578	20577	233195.61	105614.35
Table Total:		21578	20577	233195.61	105614.35

Showing 1 to 1 of 1 items

[Dismiss](#) [Export](#)

Conclusion

- Future work includes exploring potential use cases when pods are connected via a wide-area network
 - Visibility
 - Accurate real-time measurements
 - Data plane processing speed



Conclusion

- Demo, as the time permits
- Application examples in the Poster Presentation session

Thank you



UNIVERSITY OF
SOUTH CAROLINA

Platform Features

Feature	Comments
Allocation of resources	Pod granularity
Custom pods	Easy to create custom pods
Cost	Cost-effective when used extensively
Presentation layer for pedagogy	Topology is graphically presented to the learner using a regular browser
Time sharing	Easy to implement time-sharing policies
IP addresses	Pods have the same topology / IP addresses (overlapping addresses w/o conflict)
Functional realism	Virtual labs have the same functionality as real IT hardware in a real deployment
Traffic realism	Devices generate/receive real, interactive network traffic to/from the Internet

Motivation for Virtual Labs

- “The Missing Millions”(NSF report – Oct. 2021. <https://tinyurl.com/5awhdazy>)
- What can be done to reach out those who are yet to be engaged in STEM workforce?
 - “The present research computing and data ecosystems look impenetrable to many...”
 - “Lower barriers to entry...”
 - “Invest in CI and community laboratories at the edge, enabling broader and more diverse participation”
 - “Explore investments in research computing ... that are easily accessible (such as GUIs, ...)