A Distributed Academic Cloud and Virtual Laboratories for Information Technology Education and Research

Jorge Crichigno[1], Elie Kfoury[1], Kelly Caudle[2], Paul Crump[2]
[1] University of South Carolina
[2] Stanly Community College

2021 IEEE Telecommunications and Signal Processing Conference (TSP'2021)
Submission ID 52

July 26-28, 2021

Agenda

- Motivation for the Academic Cloud and virtual labs
- Academic Cloud
- Virtual lab libraries
- Impact
- Conclusion

Motivation for Virtual Labs and Academic Cloud

- Over the last decade, the information technology (IT) sector has experienced a significative increase of job needs in areas including cybersecurity, systems administration, cloud computing, and others
- The IEEE and ACM are the main societies that provide guidance for IT education. According to them[1], the IT curriculum should emphasize “learning IT core concepts combined with authentic practice” and “use of professional tools and platforms”
  - “It is not enough to simply attend courses and read books. Hands-on learning is essential for information technology”
- Using physical laboratories has been challenging
  - Difficult to scale
  - Expensive (space, maintenance, staff)
  - Since COVID-19 emerged, their capacity have been further reduced (distance requirements)

1. Information Technology Curricula 2017, ACM/IEEE Joint Committee. Online: https://tinyurl.com/4nqqwa5m.

Motivation for Virtual Labs and Academic Cloud

- The University of South Carolina (USC) in South Carolina (SC), Stanly Community College (SCC) in North Carolina (NC), and the Network Development Group (NDG) in NC have deployed the Academic Cloud
- The Academic Cloud provides remote-access capability to laboratory equipment via the Internet
- It seamlessly pools and shares resources (CPU, memory, storage) from four data centers, needed to run virtual laboratories

1. Information Technology Curricula 2017, ACM/IEEE Joint Committee. Online: https://tinyurl.com/4nqqwa5m.

Academic Cloud

- Data center locations: USC (South Carolina), SCC (North Carolina), Illinois data center, and Idaho National Laboratory (Idaho)

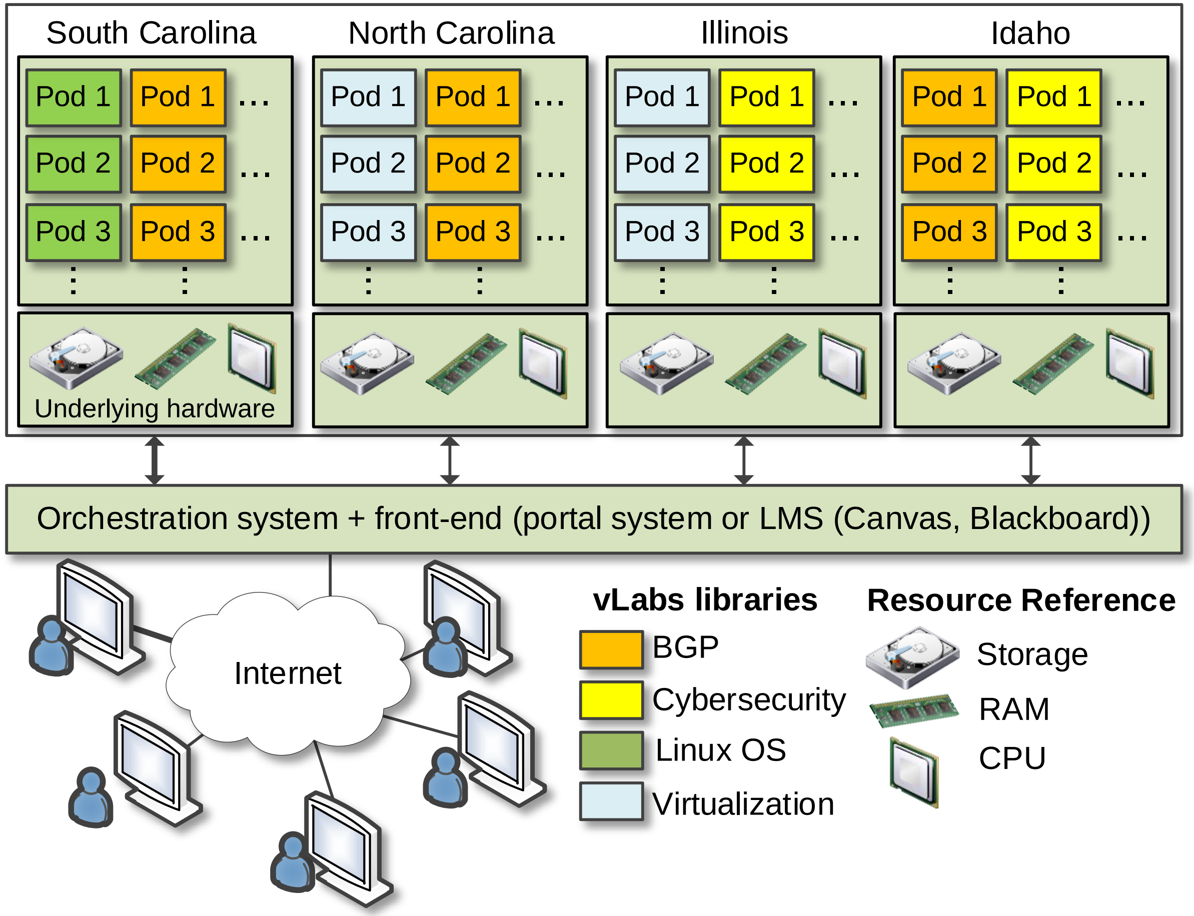

Inside a Data Center

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

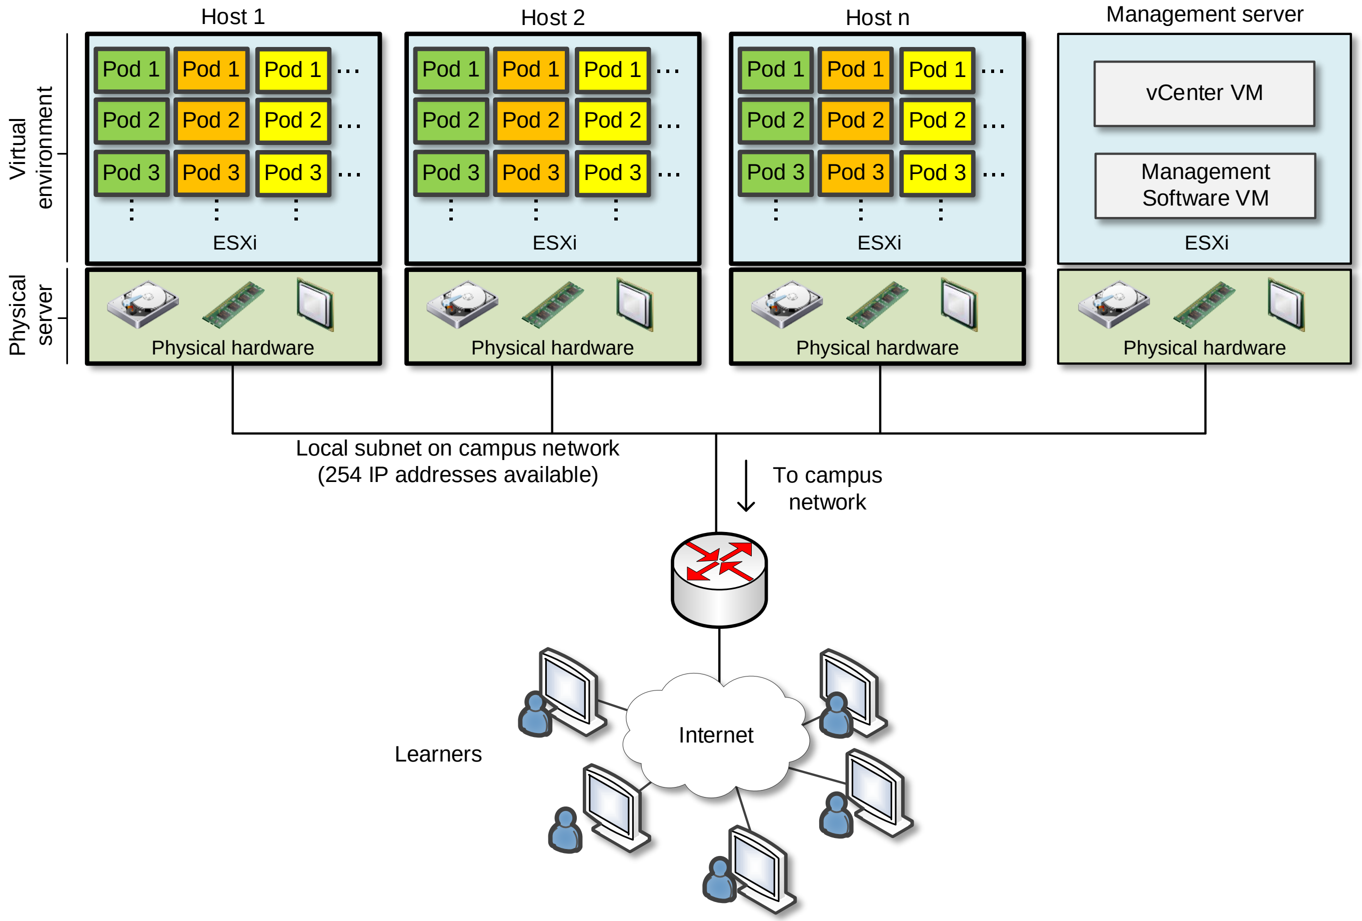

1. Network Development Group (NDG). Online: https://netdevgroup.com

Inside a Data Center

- Example: Stanly Community College

Device	Cores	Storage (TBs)	RAM Memory (GB)
Server 1 (management SCC)	20	12	264
Server 2 (hosting vLabs pods)	32	4	768
Server 3 (hosting vLabs pods)	32	4	768
Server 4 (hosting vLabs pods)	32	4	768
Server 5 (hosting vLabs pods)	32	4	768
Server 6 (hosting vLabs pods)	32	4	768
Server 7 (hosting vLabs pods)	48	1.92	768
Server 8 (hosting vLabs pods)	48	1.92	768
Server 9 (hosting vLabs pods)	48	1.92	768
TOTAL	**324**	**37.76**	**6408**

POD Design

- A virtual laboratory experiment requires a **pod** of devices, or simply pod
- Example: perfSONAR library

POD for perfSONAR labs

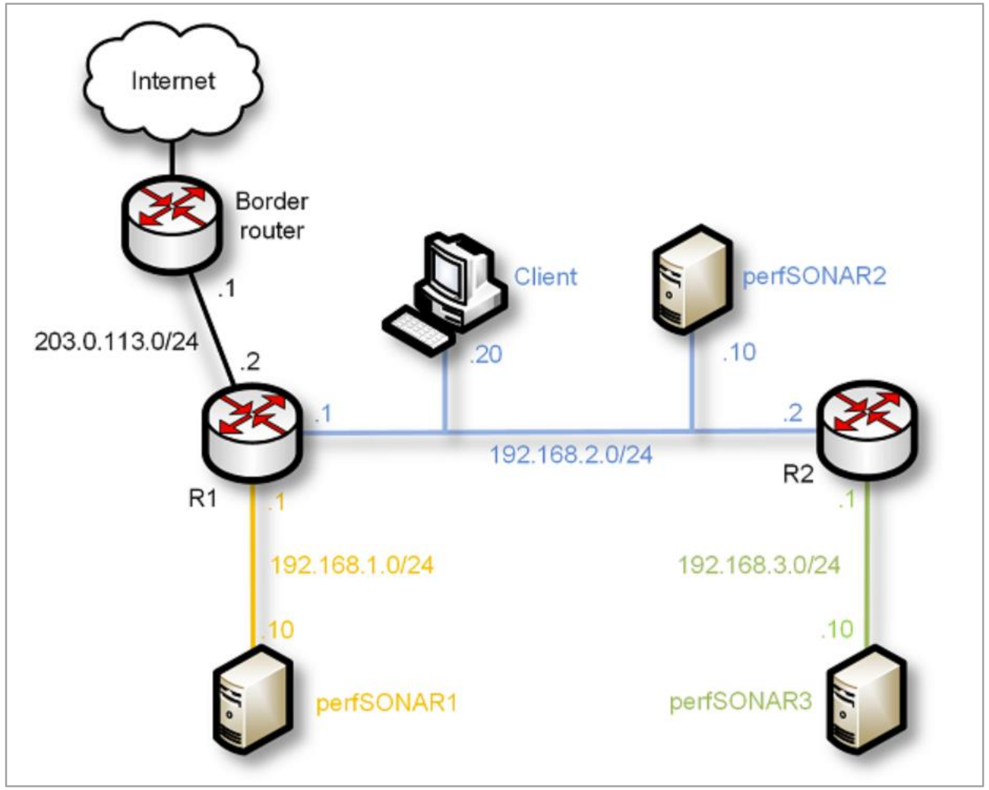

POD Design

- A virtual laboratory experiment requires a **pod** of devices, or simply pod
- Example: perfSONAR library

POD for perfSONAR labs

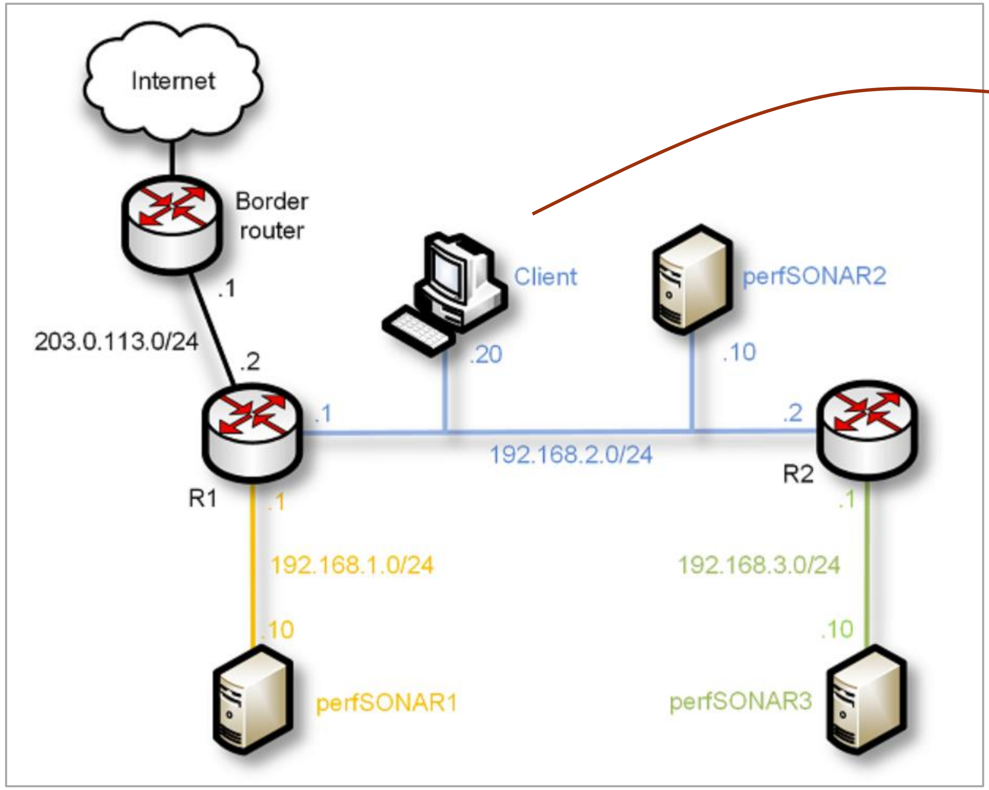

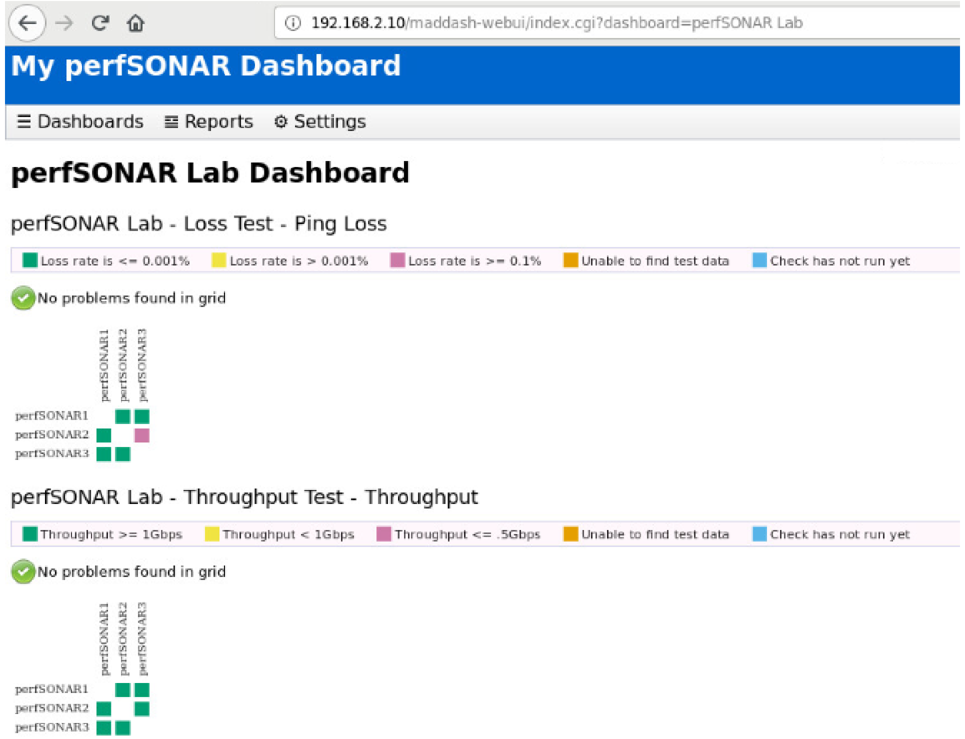

POD Design

- A virtual laboratory experiment requires a **pod** of devices, or simply pod
- Example: perfSONAR library

POD for perfSONAR labs

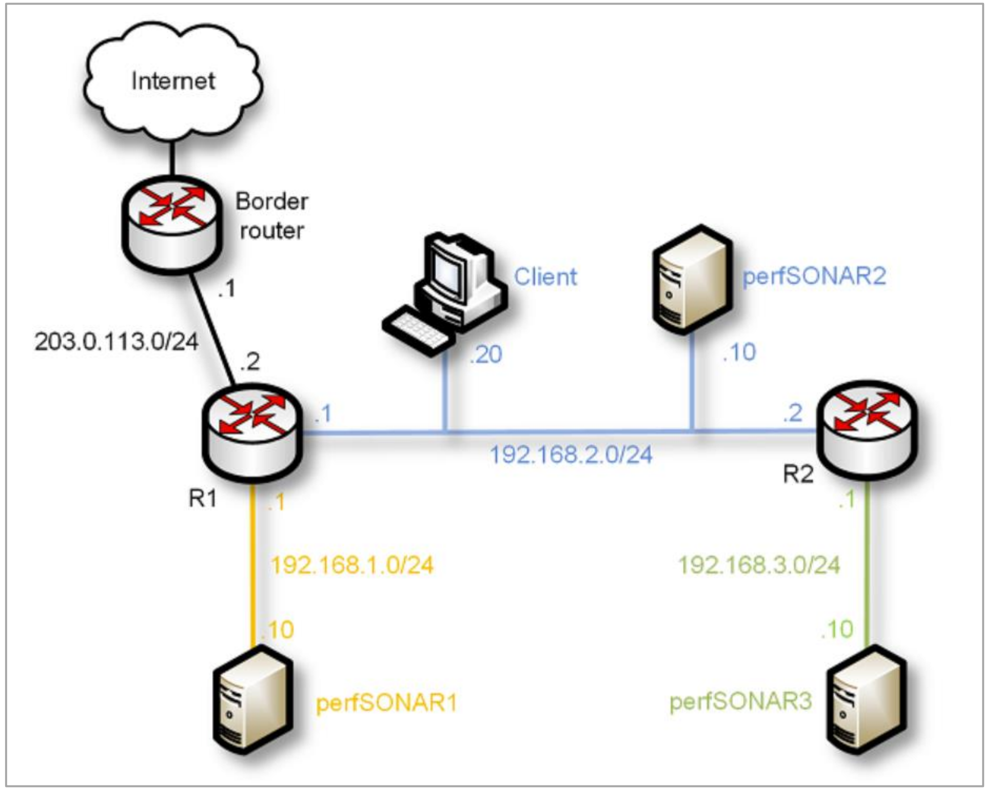

perfSONAR labs

Lab 1	Configuring Administrative 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

POD Design

- Details of perfSONAR pod
  - Four networks
  - Three servers
  - One client
  - Three routers
  - Connectivity to the Internet
  - Total of 7 heterogeneous VMs

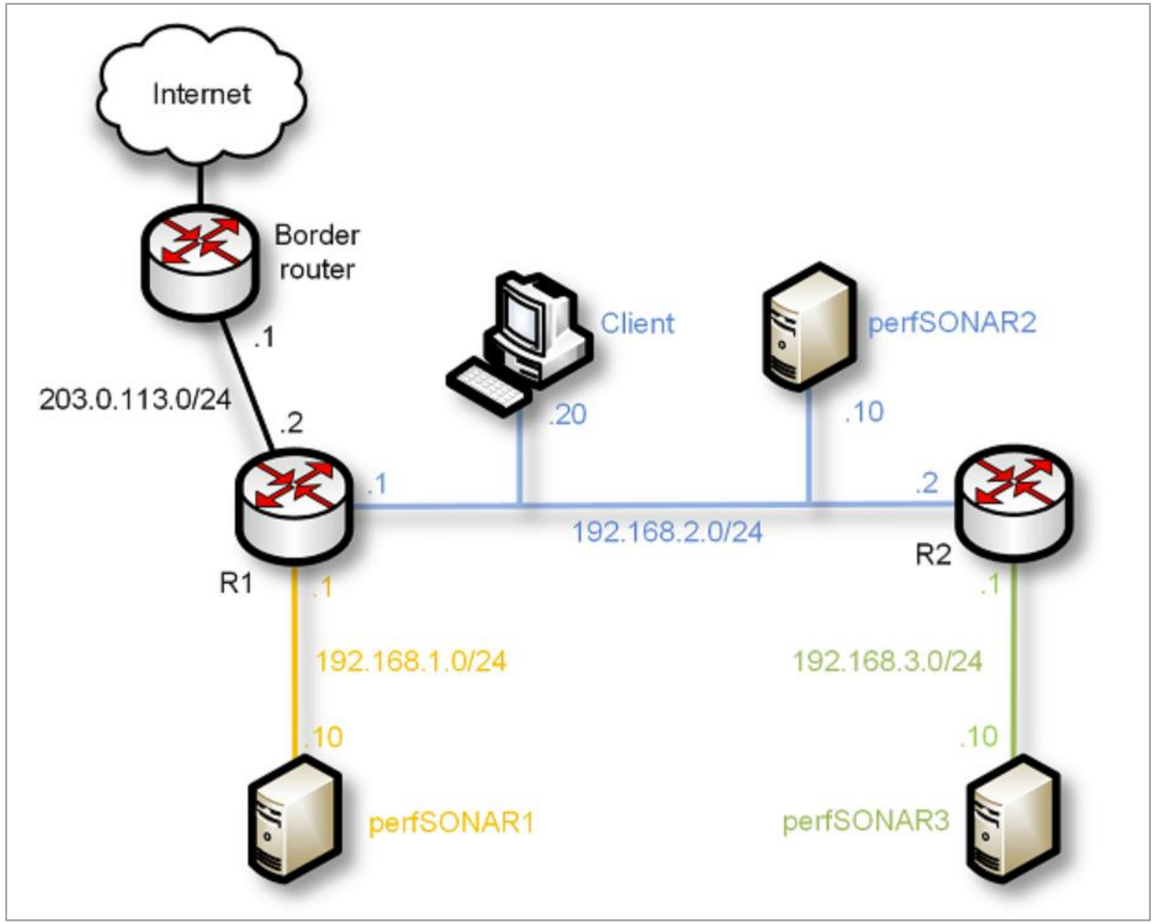

POD Design

- Details of perfSONAR pod
  - PODs running simultaneously use the same block of IP addresses
  - Lab manuals are uniform
  - “Local NAT” is performed by the device connected to the campus network
  - There is a master pod in the system
  - Linked clone VMs are created from the master pod VMs

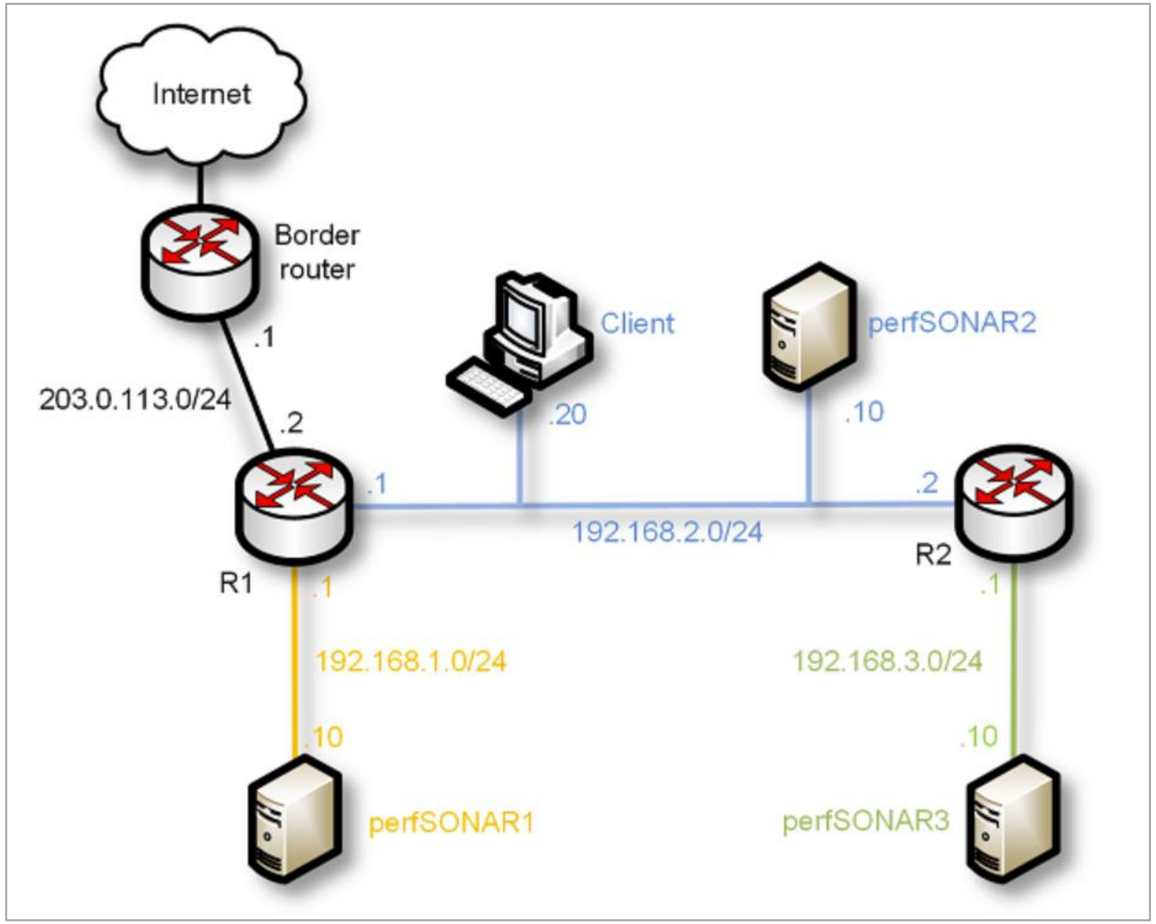

Lab Libraries

- Network Tools and Protocols

Lab 1	Introduction to Mininet
Lab 2	Introduction to Iperf3
Lab 3	Emulating WAN with NETEM I: Latency, Jitter
Lab 4	Emulating WAN with NETEM II: Packet Loss, Duplication, Reordering, and Corruption
Lab 5	Setting WAN Bandwidth with Token Bucket Filter (TBF)
Lab 6	Understanding Traditional TCP Congestion Control (HTCP, Cubic, Reno)
Lab 7	Understanding 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 MSS 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)

Lab Libraries

- Network Tools and Protocols, Lab 14: “Router’s Bufferbloat”

Description of the problem

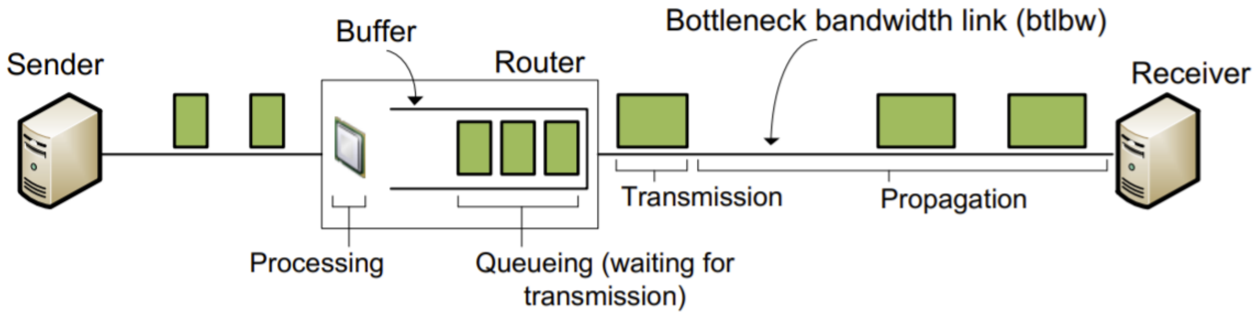

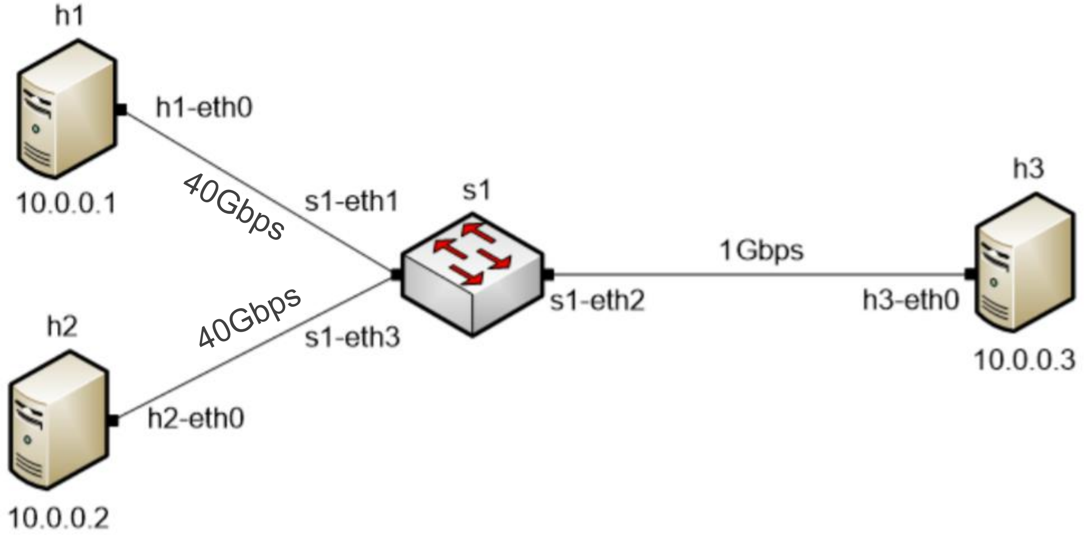

Topology

Lab Libraries

- Introduction to SDN

Lab 1	Introduction to Mininet
Lab 2	Legacy Networks: BGP Example as a Distributed System and Autonomous Forwarding Decisions
Lab 3	Early efforts of SDN: MPLS Example of a Control Plane that Establishes Semi-static Forwarding Paths
Lab 4	Introduction to SDN
Lab 5	Configuring VXLAN to Provide Network Traffic Isolation
Lab 6	Introduction to OpenFlow
Lab 7	Routing within an SDN network
Lab 8	Interconnection between Legacy Networks and SDN Networks
Lab 9	Configuring Virtual Private LAN Service (VPLS)
Lab 10	Applying Equal-cost Multi-path Protocol (ECMP) within SDN networks

Lab Libraries

- Introduction to SDN, Lab 8: “Interconnection between Legacy Networks and SDN Networks”

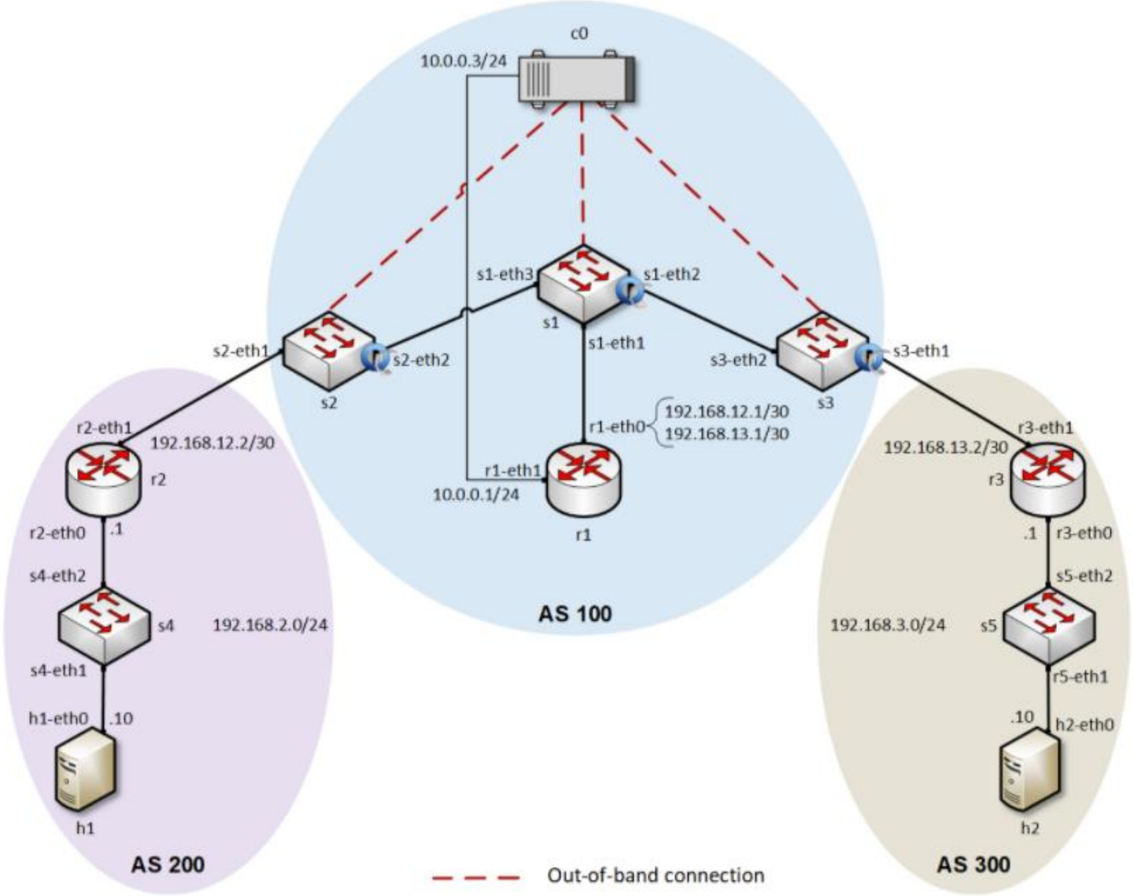

Lab Libraries

- Introduction to BGP

Lab 1	Introduction to Mininet
Lab 2	Introduction to Free Range Routing (FRR)
Lab 3	Introduction to BGP
Lab 4	Configure and Verify EBGP
Lab 5	BGP Authentication
Lab 6	Configure BGP with Default Route
Lab 7	Using AS_PATH BGP Attribute
Lab 8	Configuring IBGP and EBGP Sessions, Local Preference, and MED
Lab 9	IBGP, Next Hop and Full Mesh Topology
Lab 10	BGP Route Reflection

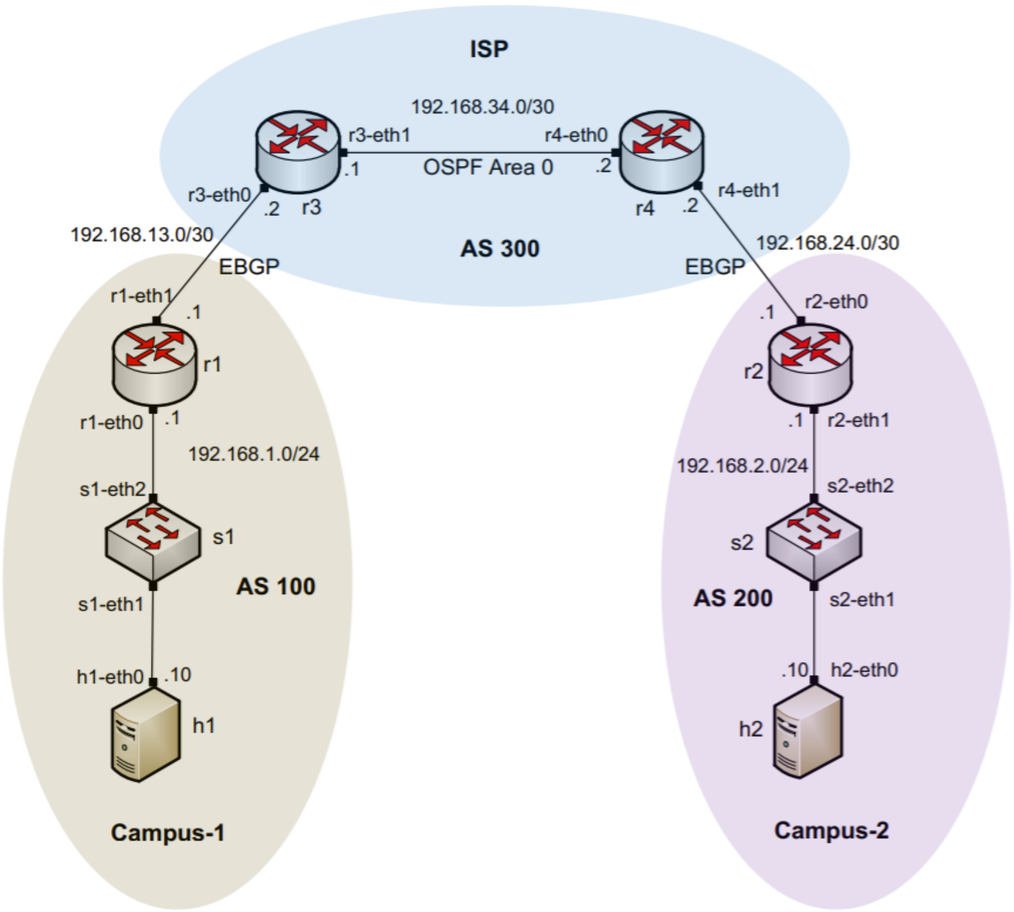

Topology for Lab 4

Cloud Features

Feature	Comments
Allocation of resources	Granular allocation of physical resources
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	The owner controls who can access resources; easy to implement time-sharing policies
IP addresses	Pods (and learners) can have the same topology and IP addresses (overlapping addresses w/o conflict)
Functional realism	Virtual labs have the same functionality as real IT hardware in a real deployment, and execute the same code
Traffic realism	Devices generate/receive real, interactive network traffic to/from the Internet, or to/from other devices within the lab environment

Partnership with Industry

- The IEEE and ACM are the main societies which guide IT education
  - IT curriculum should emphasize “learning IT core concepts combined with authentic practice” and “use of professional tools and platforms”[1]
- USC works with the Network Development Group (NDG)[2], VMware, Palo Alto Cybersecurity Academy, Cisco, Juniper, and others to virtualize labs

1. “Information Technology Curricula Guideline 2017 (IT2017),” report by the ACM / IEEE Task Force on Information Technology Curricula, Dec. 2017. Online: https://tinyurl.com/yxauot8w
2. Network Development Group (NDG). Online: https://netdevgroup.com

Partnership with Industry

- These labs enhance the student's understanding of how modern firewalls work, referred to as Next-generation Firewalls (NGFWs)

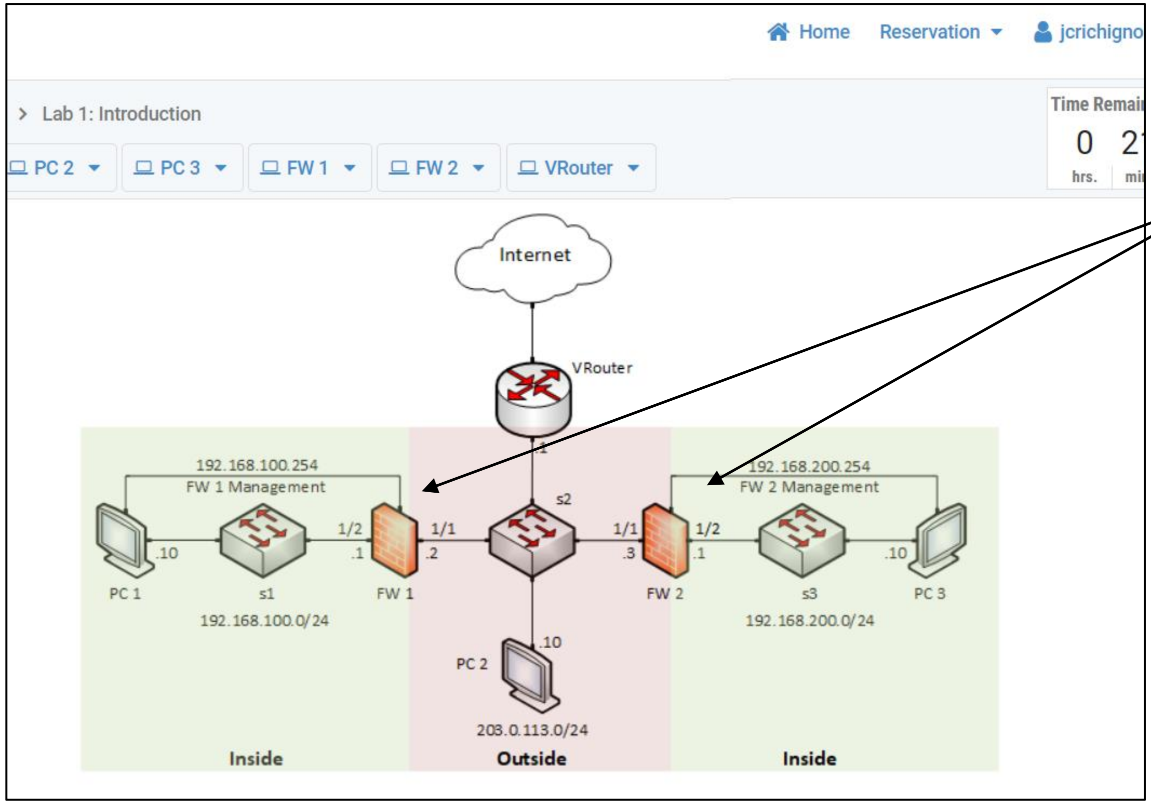

Pod deployed in private cloud

Next-generation Firewall Virtual Machine + licenses

Impact

- The Academic Cloud has served over 100,000 learners
- Academic institutions (colleges, universities, high-schools) and training centers
- Self-pace learners
- Usage example from one institution supporting one academic program (~300 students, January 1, 2020 – December 30, 2020)

Community Usage

Reservations Made	Labs Attended	Hours Reserved	Hours Attended
6909	6534	42446.59	25158.03
6909 6909	6534 6534	42446.59 42446.59	25158.03 25158.03

Conclusion

- USC, SCC, and NDG implemented the Academic Cloud, a platform to run virtual labs
- The cloud is managed by a customized software that connects four data centers distributed across the U.S.
  - It controls the resources in those data centers (servers, CPUs, RAM, VMs)
- The system has shown to be scalable
  - It has served over 100,000 learners in one year
- Due to the positive feedback, the system is expanding with more resources and virtual labs

UNIVERSITY OF
SOUTH CAROLINA