2023 INTERNET2 TECHNOLOGY exchange

Hybrid SDN Campus Architectures for Agility and Enhanced Cybersecurity

Presented By:

Charles Kneifel, PhD, Chief Technology Officer, Duke University
William Brockelsby, PhD, Chief Network Architect, Duke University and Health System

Agenda

- Motivation
- Approach
- Architecture
- Use Cases

Motivation

Campus Networks:

- Explosion of network connected devices
- Increasing awareness of the need for enhanced cybersecurity
- Delivery of enhanced cybersecurity at network speed
- The need to transfer increasingly large research data sets
- Compliance requirements
- Cost reduction

- **We must improve the performance and security of computer networks to make it easier and less costly for scientists to move, store, and analyze their data**

Motivation – SDN Paradox

The origins of the SDN movement were in campus networks over a decade ago and yet SDN is typically non-existent or is only deployed for specialty use cases on campus...

Potential Reasons:

- Many optional features in OpenFlow specification
- OpenFlow bolted onto legacy platforms
- Early SDN “capable” devices with poor support, scale and/or performance
- Low-level abstraction introduces operational risk
- Concerns over a format war
- Concerns over wholesale replacement of existing infrastructure

Approach => Archipelago

- Leverage software defined networking
- Use best-of-breed equipment
- Deploy a hybrid architecture
- Augment existing infrastructure instead of replacing it
- Leverage an existing high-performance substrate
- Reduce complexity
- Minimize risk

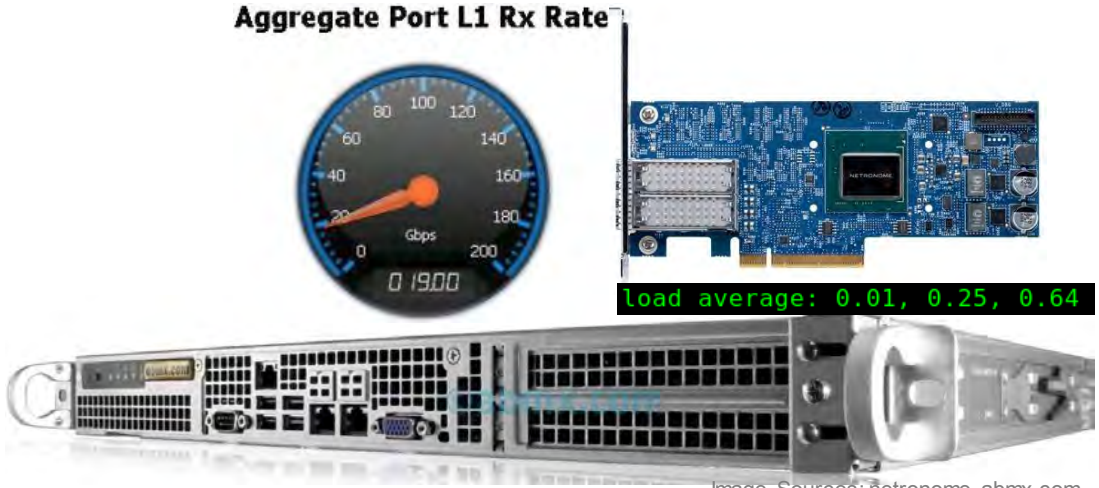

Image Sources: netronome, abmx.com

Archipelago: Equipment Evaluation

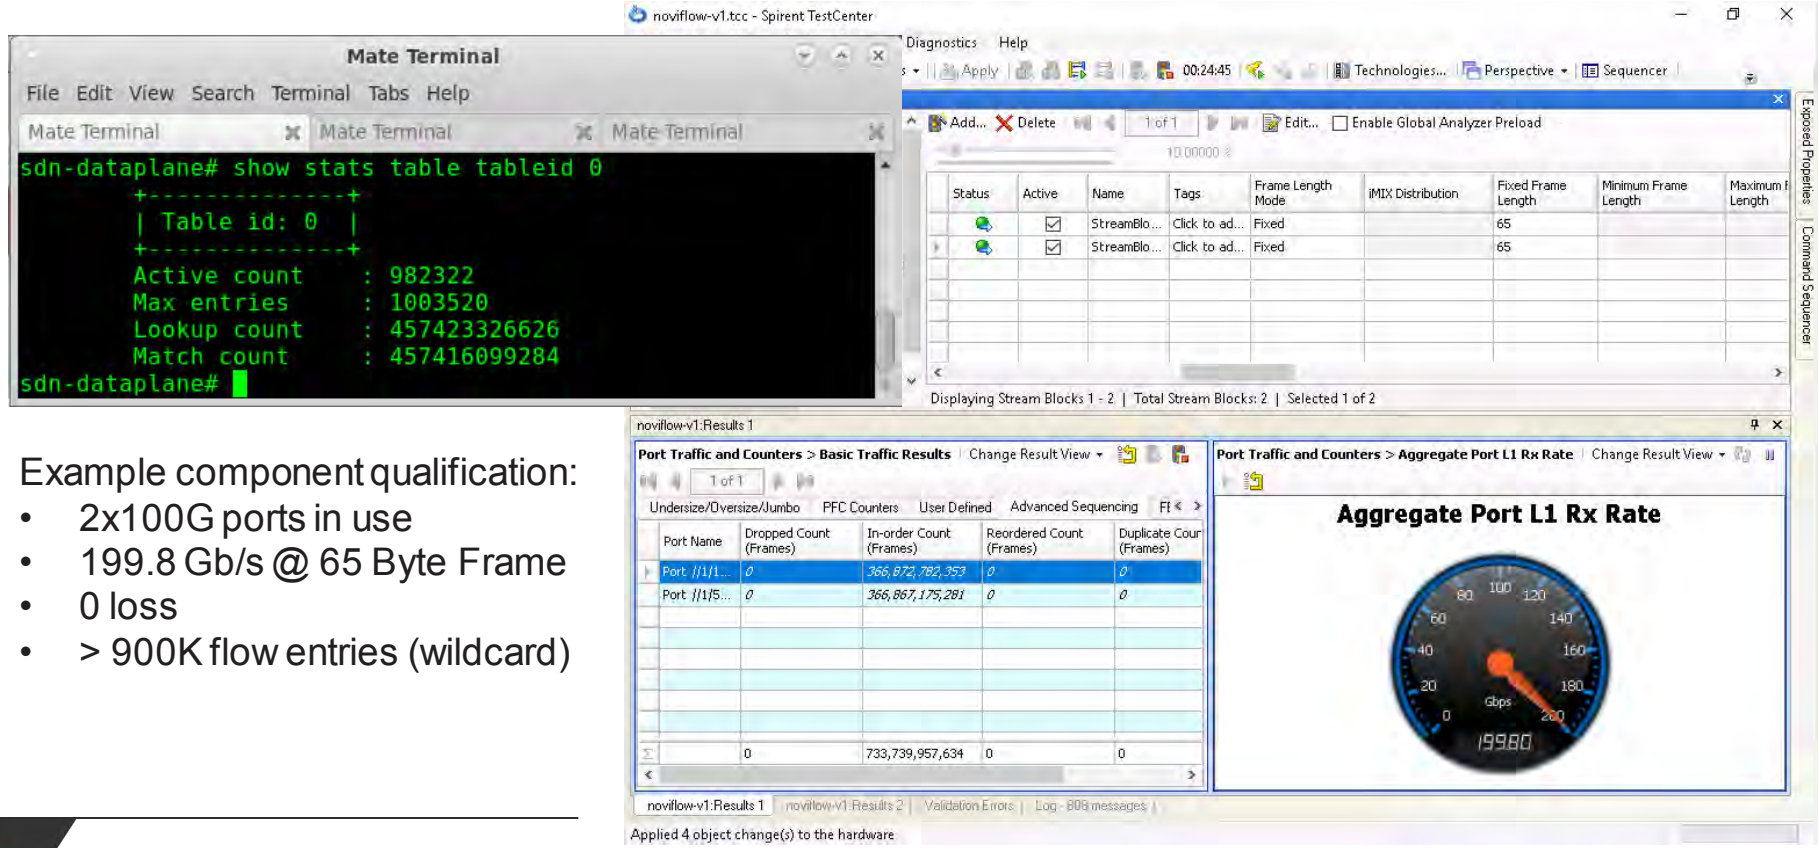

Example component qualification:

- 2x100G ports in use
- 199.8 Gb/s @ 65 Byte Frame
- 0 loss
- > 900K flow entries (wildcard)

Archipelago: Equipment Evaluation

Intel x710 (control)

Netronome CX 2x10G

Netronome CX 2x25G

Netronome CX 2x40G

NoviFlow NS2122 2x100G 20x 10G
NVIDIA/Mellanox NP-5 NPU

Edge-Core WEDGE 100BF-32X
IntelTofino32x 100G

Brockelsby, William, and Rudra Dutta. "Performance implications of problem decomposition approaches for SDN pipelines." GLOBECOM 2020-2020 IEEE Global Communications Conference. IEEE, 2020.

Image Sources: noviflow, edge-core

Archipelago: Architecture

Building Architecture:

- Traffic Study: Across multiple facilities at Duke University and North Carolina State University
- Move SDN insertion point to the distribution layer with enhanced SDN data planes
- Reduces cost, risk and is horizontally scalable
- Insertion at layer-2/3 boundary permits rich policies with access to all supported headers

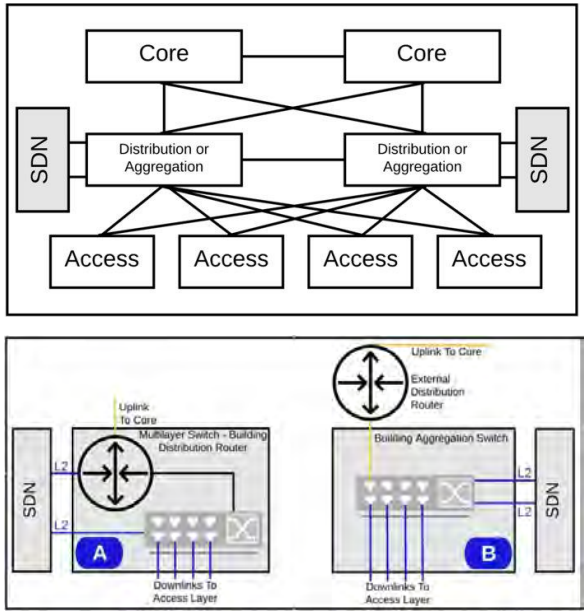

Brockelsby, William, and Rudra Dutta. "Traffic analysis in support of hybrid SDN campus architectures for enhanced cybersecurity." 2021 24th Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN). IEEE, 2021.

Archipelago: Architecture

Hybrid Campus/Community Architecture:

- Native augmentation with SDN forwarding elements yields islands
- Overlays can (but are not required) to connect islands of SDN yielding an Archipelago
- Hybrid use of approaches allows gradual insertion on campuses, between campuses and within the cloud

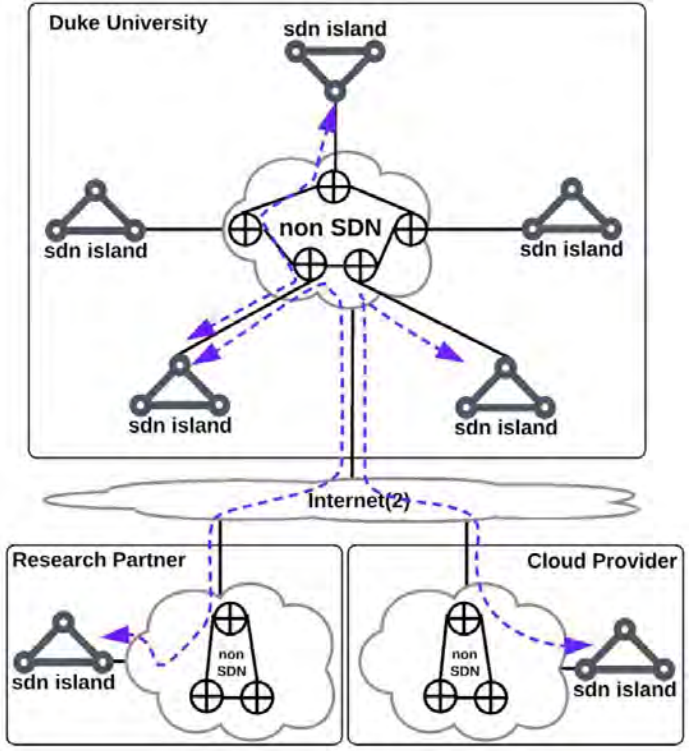

Brockelsby, William, and Rudra Dutta. "Archipelago: A Hybrid Multi-Node Campus SDN Architecture." 2023 26th Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN). IEEE, 2023.

Archipelago: Architecture

Control Plane:

- Building upon our hybrid theme, we again take a hybrid approach with regard to controller placement:
- Couple a low-cost appliance with purpose-built SDN data planes for use as a local controller
- SmartNIC based data planes are already housed in a server appliance - use that as a local controller

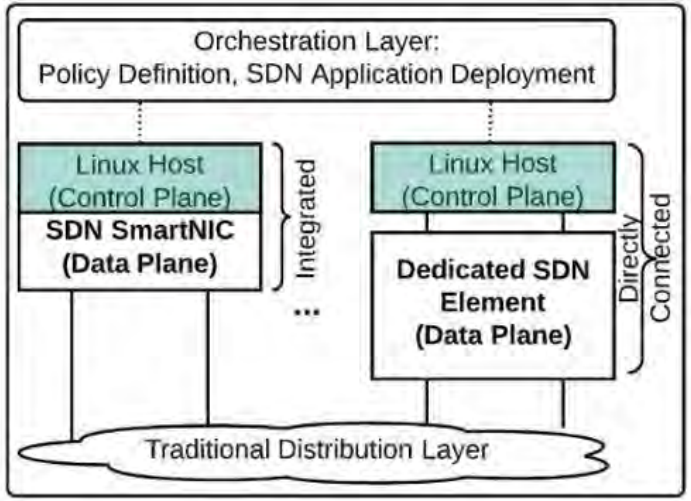

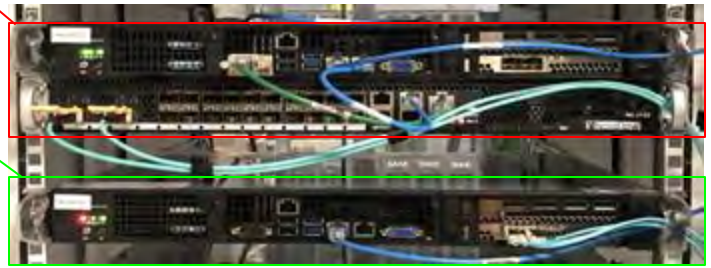

https://sites.duke.edu/archipelago/
https://sites.duke.edu/archipelago/files/2023/01/archipelago-final-appliance-evolution-v1.pdf

Archipelago: Context and Goals at Duke

- Friction free policy enforcement WITHIN Virtual Routing and Forwarding (VRF) instances
- Allow authorized flows, such large research data transfers, to bypass existing sources of friction
- Expand beyond a single point of insertion: Islands of SDN nodes form an Archipelago

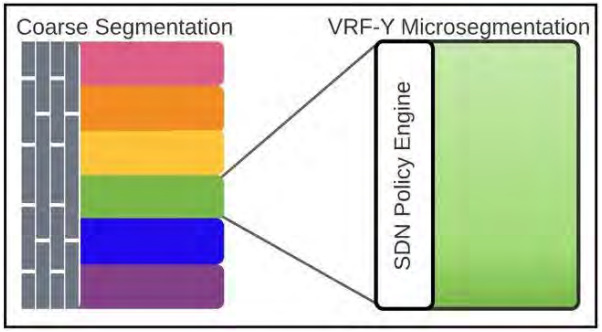

Archipelago: Example – Friction Bypass

- By combining two types of SDN: the table/pipeline abstraction for policy enforcement with overlay networks, it is possible to bypass points of friction within a campus network as illustrated
- The green hosts bypass the inter-VRF firewall (snail) with a 7Gb/s limit by intercepting green traffic and placing it into an overlay network between SDN nodes within the friction-free underlay network
- We plan on leveraging Archipelago to facilitate friction-free inter-organization connectivity for research cyberinfrastruture separated by multiple firewalls and other sources of friction

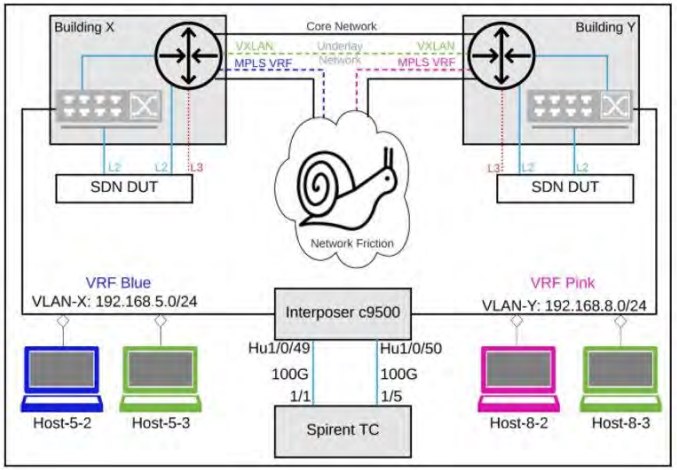

TABLE VII: Multi-Node - Friction Bypass via SDN Overlay

Host Flow	TX Gb/s	RX Gb/s	Latency (µs)	Loss %
5-2 =>8-2	19.5	7	*1125.7*	*64.2*
8-2 =>5-2	19.5	7	*1126.03*	*64.2*
5-3 =>8-3	19.5	19.5	57.31	0
8-3 =>5-3	19.5	19.5	58.38	0

Brockelsby, William, and Rudra Dutta. "Archipelago: A Hybrid Multi-Node Campus SDN Architecture." 2023 26th Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN). IEEE, 2023.

Archipelago: Duke Deployment at French Science

French Family Science Center (FFSC) is a 280,000 square foot building housing wet and dry labs for research and teaching, classrooms and faculty offices for chemistry, genomics, materials science and nanoscience.

The facility leverages a toxic gas monitoring system that was regularly generating false alarms due to undesirable network traffic.

Archipelago isolated this and other sensitive infrastructure within the facility. The insertion was transparent to IT, facilities and lab management personnel.

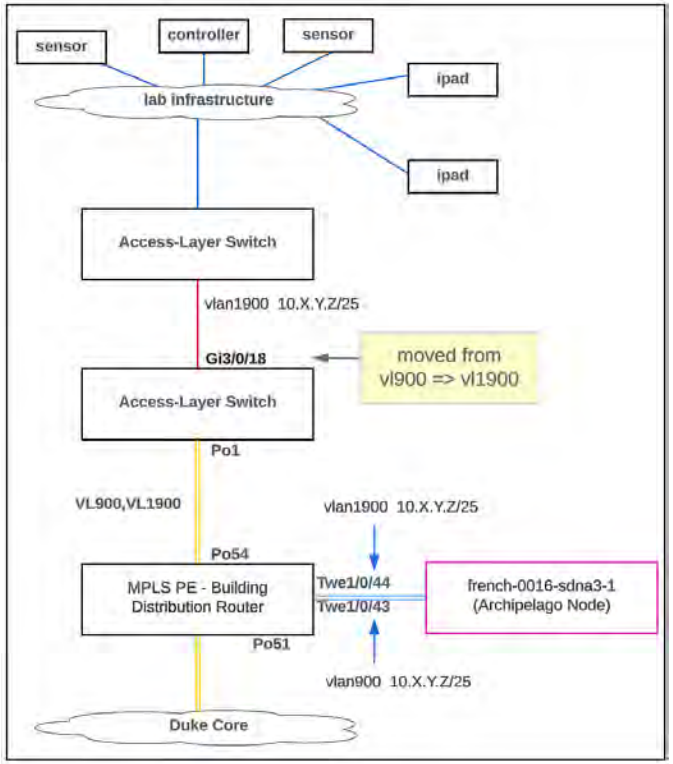

Image Sources: Duke University

Archipelago: Duke Deployment at CIEMAS

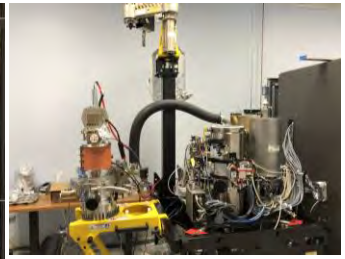

The Fitzpatrick Center for Interdisciplinary Engineering, Medicine and Applied Sciences (FCIEMAS) was designed to support major advancements in the fields of: bioengineering, photonics, communications, materials science and materials engineering.

We leveraged Archipelago to provide enhanced friction-free cybersecurity for instruments including a Thermo Fisher Cryo-Transmission Electron Microscope (Cryo-TEM) outfitted with multiple specialized sensors.

Archipelago isolated this and other sensitive infrastructure within the facility. The insertion was transparent to IT, facilities and lab management personnel and introduced NO degradation in performance when transferring data to the Duke Compute Cluster (DCC).

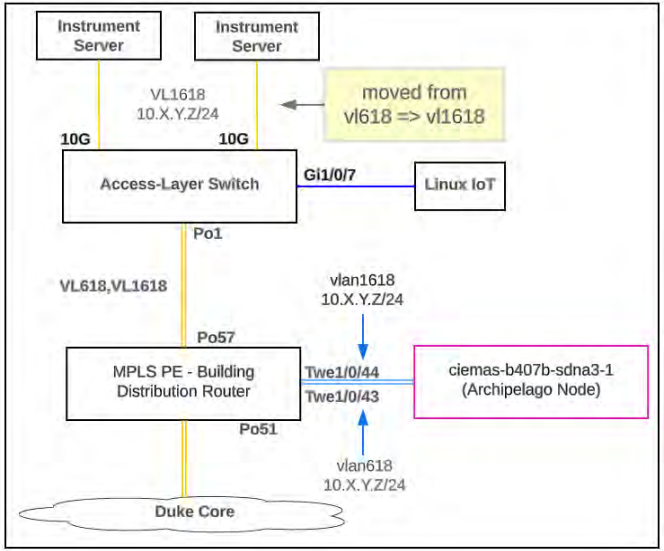

Image Sources: Duke University

Archipelago: Duke Deployment at Chesterfield

The Chesterfield building is a recently renovated 286,000-square foot multi-use facility on Main Street originally built for manufacturing with around 100,000 square feet leased to Duke for a variety of research labs.

We are currently planning an Archipelago deployment within Chesterfield for the Pratt School of Engineering. The deployment will allow legacy operating systems required to control and operate scientific instruments to remain on the network in an isolated manner with full visibility to our security team thereby preserving the utility and life of these instruments.

Archipelago will transparently regulate traffic in/out of isolated VLANs supporting this control infrastructure with minimal changes to the instruments themselves (no IP addresses or other settings within the instruments will need to be changed as some of these changes require vendor support). While the deployment is currently underway, a representative example is shown in the photo to the right.

Representative Deployment Example

Image Sources: Duke University, Triggs Photography

Archipelago: Duke Horizontally Scalable Firewall

- Increasing traffic loads often force organizations to purchase large vertically integrated monolithic network appliances, such as firewalls, to handle the load.
- This approach introduces fiscal and operational concerns. As an alternative, we proposed deploying a horizontally scalable cluster of medium-sized appliances to allow us to rightsize the deployment as our needs change.
- Purpose-built SDN switches, evaluated as a part of our Archipelago research, were selected to steer traffic from traditional switches to the appropriate appliance cluster by policy.
- The SDN switches themselves were deployed in a stateless mode and yet continue to provide flow affinity to the appropriate firewall cluster. This works by making the firewall selection decision by utilizing aggregate blocks of IP space and/or VLAN tags vs individual flows.
- Interestingly, the topology permits traffic to bypass the firewalls altogether if needed by directing specific flows through the SDN switches and completely away from the firewalls. We are currently extending this architecture to our perimeter security layer for deployment this fall.

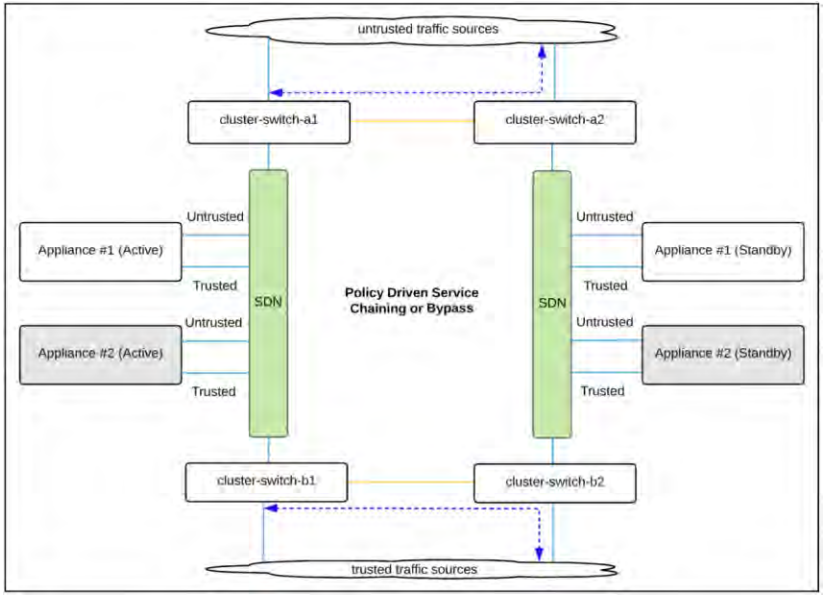

Brockelsby, William. "Hybrid SDN Campus Architectures for Agility and Enhanced Cybersecurity." Dissertation: North Carolina State University, 2022.

Image Sources: Duke University

Archipelago: Duke SDN-Driven Service Chaining

- At Duke, students have the opportunity to have full root access to virtual machines with public IP addresses as a part of the Virtual Computing Manager (VCM) environment
- Although traffic associated with this environment originates WITHIN the Duke network, we want to treat it as if it were external/untrusted and subject all traffic to/from this environment to multiple forms of inspection
- We leverage service chaining with SDN to transparently insert additional forms of inspection to the architecture described previously (only a portion of which are shown in the diagram to the right)

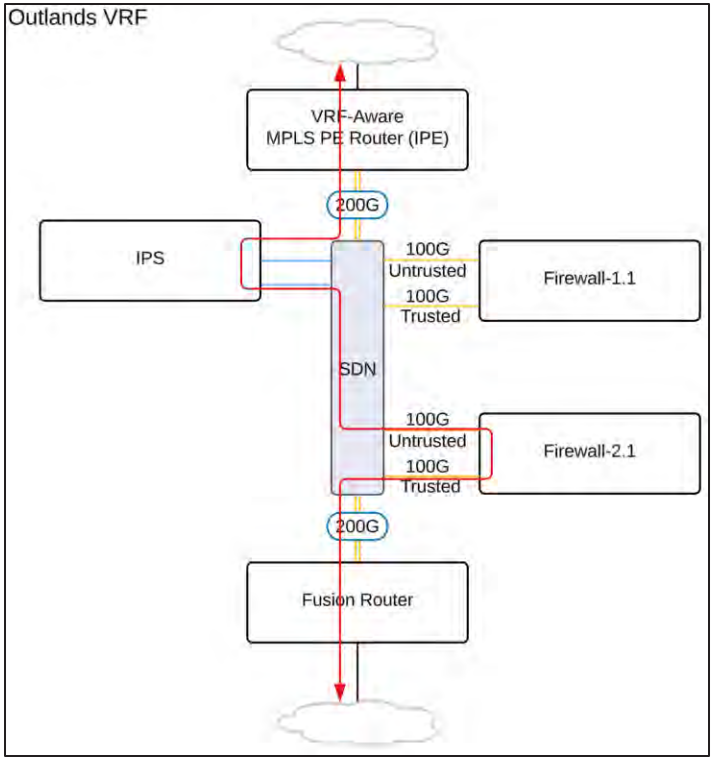

Image Sources: Duke University

Acknowledgement

A portion of this research is based upon work supported by the National Science Foundation under Grant No. 1925550 at Duke University.

Thank you!

Relevant Publications

Brockelsby, William, and Rudra Dutta. "A graded approach to network forensics with privacy concerns." 2019 International Conference on Computing, Networking and Communications (ICNC). IEEE, 2019.

Brockelsby, William, and Sean Dilda. "Tactical network automation with NetZTP and one shot." 2019 IEEE 40th Sarnoff Symposium. IEEE, 2019.

Brockelsby, William, and Rudra Dutta. "Performance Implications of Problem Decomposition Approaches for SDN Pipelines." GLOBECOM - 2020 IEEE Global Communications Conference. IEEE, 2020.

Brockelsby, William, and Rudra Dutta. "Traffic Analysis in Support of Hybrid SDN Campus Architectures for Enhanced Cybersecurity." 2021 24th Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN). IEEE, 2021.

Brockelsby, William. "Hybrid SDN Campus Architectures for Agility and Enhanced Cybersecurity." Dissertation: North Carolina State University, 2022.

Brockelsby, William, and Rudra Dutta. "Augmenting Campus Wireless Architectures with SDN" (In Press) International Conference on Computing, Networking and Communications (ICNC). IEEE, 2023.

Brockelsby, William, and Rudra Dutta. "Archipelago: A Hybrid Multi-Node Campus SDN Architecture" (In Press) 26th Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN). IEEE, 2023.