Contribution ID: **4** Type: **Oral Presentation**

# A Simplified SDN-Driven All-Campus Science DMZ

*Friday, 23 March 2018 10:50 (30 minutes)*

Data-intensive computational techniques such as machine learning, data analytics, and visualization, increasingly require data sets at unprecedented scale - massive sizes that are orders of magnitude larger than previous work. This presents challenges for computer networks. This problem is particularly acute for universities where researchers are increasingly using big data in their research, but the campus network infrastructure is not designed for high-throughput communication. Moreover, north/south traffic to cloud storage providers such as Amazon and Google is growing at an explosive rate, and it is now common for researchers to move terabytes of data to/from the cloud. The result is nothing short of a dire need for high-throughput campus network infrastructure.

To complicate matters, campus networks, which are designed to support common institutional business functions (web browsing, email) are also intended to support high-end research endeavors. Is it possible for the network to do both? As designed, these networks are almost always littered with so-called middle-boxes that perform services such as intrusion detection, rate limiting, firewalling, and other forms of deep packet inspection. These middleboxes play an important role in ensuring security and stability of the campus network, but sacrifice network performance. Though the resulting decline in network performance may be acceptable for common low-bandwidth applications, large data transfers that rely on higher bandwidth suffer greatly.

To meet this challenge, we proposed a new approach to the design of campus networks based on software defined networks (SDN) — specifically OpenFlow. We began by replacing certain building distribution routers with OpenFlow-enabled switches that operate in a hybrid mode, providing normal routing and switching to our standard campus core by default. However, using OpenFlow, we are able to redirect high-throughput flows from approved researchers to an all-new SDN core. The SDN core then forwards packets directly to our campus edge router, bypassing all middleboxes north of the standard campus core infrastructure. A major benefit of our approach is that individual flows from a machine can receive high-speed (middlebox free) paths while all other flows from the same machine travel the standard campus path through policy-enforcing middleboxes. Consequently, transparent to the end user, a host can perform a high-speed file transfer to the cloud while at the same time streaming rate-limited video content. This effectively creates a virtual all-campus DMZ, granular to protocol port level, that can be turned on or off programmatically as needed by researchers.

## Summary

University researchers increasingly must choose between connectivity to science DMZs or to the traditional campus network. The campus network provides convenient access to campus services and security policy enforcement whereas the science DMZ allows unfettered network throughput needed for large data sets. We present an SDN-driven approach that permits individual flows from a machine over either path, effectively creating a granular all-campus DMZ that can be enabled or disabled programmatically as needed by researchers.

**Primary author:** Mr CHAPPELL, Jacob (University of Kentucky)

**Co-authors:** Dr SEALES, Brent (University of Kentucky); Mr PIKE, Charles (University of Kentucky); Dr BUMGARDNER, Cody (University of Kentucky); Dr GRIFFIOEN, James (University of Kentucky)

**Presenter:** Mr PIKE, Charles (University of Kentucky)

**Session Classification:** Networking, Security, Infrastructure & Operation Session

**Track Classification:** Networking, Security, Infrastructure & Operations