Contribution ID: 18

Orchestrating Dynamic High-Speed Network Flows for Secure Research Data Transfers Using an SDN-enabled Infrastructure

Thursday, 4 April 2019 17:00 (30 minutes)

Big data is now key to nearly all research disciplines. Even research areas historically unrepresented in high performance computing must now cope with vast data sets that need to be analyzed, processed, and transferred over the network. Network choke points can create significant delay during transmission of these large data sets to and from the cloud, where they often reside. Campus and enterprise networks depend on middleboxes (e.g., firewalls, NAT, load balancers, IDS/IDP) to provide essential services or enforce network policies, yet these middleboxes often contribute to the negative network performance affecting big data transfers. The transmission delays can cause troubling workflow design issues for researchers and raise a dire need for high-throughput networks as part of the campus cyberinfrastructure.

The traditional fix to the above problem is to create a science DMZ, a strictly-administered segment of the campus network that is positioned outside the firewalls and other policy-enforcing middleboxes. Privileged research systems needing high-speed data transfer capabilities are moved from the general campus network to the science DMZ. This exposes the entire machine to attack, and necessitates careful bastion-host configurations. Researchers are confronted with the dilemma of choosing between high-speed networking with the risks of the DMZ, or the negative performance impact that comes with the services and security provided by the general purpose campus network.

To meet this challenge, we proposed a new approach to the design of campus networks based on software defined networking (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. Using OpenFlow, flows can be redirected to a new SDN core. The SDN core then forwards packets directly to our campus edge router, bypassing all middleboxes in the campus infrastructure. A benefit of this design is that individual flows from a given 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. 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.

Here we will present a system we call "VIP Lanes" that leverages the SDN-enabled network to provide the authentication, delegation, authorization, and orchestration to dynamically create and teardown trusted research flows either transparently or as requested by researchers. Unlike a science DMZ where privilege is granted to individual machines, VIP Lanes authorization is given out on a per-flow basis. VIP Lanes provides the ability for pre-authorized, trusted users or applications to create flows that bypass the normal campus route, thereby enabling those flows to achieve substantially better performance while maintaining security and policy compliance for other network traffic. The system dynamically calculates routes that bypass bottlenecks using a graph database approach that computes custom paths and inserts OpenFlow rules that modify packet forwarding hop-by-hop. Additional functionality is added to the path when needed using network function virtualization (NFV) techniques (e.g., NAT). We present the VIP Lanes abstraction and services. We describe our current production implementation that not only shows the viability of the VIP Lanes approach, but also demonstrates the types of performance improvements achieved - in some cases approaching a two order of magnitude reduction in transmission times.

Summary

We present the VIP Lanes system that leverages SDN-enabled networks to provide the authentication, delegation, authorization, and orchestration to dynamically create and teardown trusted research flows either transparently or as requested by researchers. Unlike a science DMZ where privilege is granted to individual machines, VIP Lanes authorization is given out on a per-flow basis. VIP Lanes provides the ability for pre-authorized, trusted users or applications to create flows that bypass the normal campus routing, thereby achieving substantial performance improvements while maintaining security and policy compliance for other network traffic. **Primary authors:** Mr PIKE, Charles (University of Kentucky); Dr GRIFFIOEN, James (University of Kentucky); Ms HAYASHIDA, Mami (University of Kentucky); Mr RIVERA, Sergio (University of Kentucky); Dr FEI, Zongming (University of Kentucky)

Presenter: Mr PIKE, Charles (University of Kentucky)

Session Classification: Networking, Security, Infrastructure & Operations

Track Classification: Network, Security, Infrastructure & Operations