



LA-UR-24-27897

Author(s): Robin Simpson, Anvitha Ramachandran, Dohyun Lee

Mentors: Doug Egan, Alex Lovell-Troy, David Rich

Title: Cluster Management with Containerization on Switches

Abstract:

Network switches, such as those from Arista and Mellanox, often have underutilized computational resources in the form of built-in processors and memory. By leveraging these untapped resources, we can optimize functionality and efficiency of computational cluster networks. Our research focuses on deploying containers directly onto these switches to execute various auxiliary tasks ranging from metric logging to system-wide management via post-boot configuration. By doing so, we can significantly enhance the capabilities of the cluster without the need for additional dedicated hardware. Our research involved five distinct scenarios where switch utilization could have a profound impact on HPC Clusters: run cloud-init services via link-local connection; configuring a Telegraf container to export metrics; deploying a caching proxy; creating a reconfigurable IPv6 DHCP/DNS provider for VLAN; and implementing a client detection with Magellan discovery. These scenarios were containerized with podman and docker, and tested both physically on the switch virtually on a QEMU VM both running SONiC OS. Testing and findings indicate that network switches can indeed be used for these scenarios. They offer a wide range of possibilities beyond these applications. They run as expected as containers on the switches, and although there were some minor issues, work-arounds were implemented. Overall, this is a positive result that can be further explored with more scenarios.