Position Paper for the DOE Network 2025 Challenges Workshop

M. Veeraraghavan, University of Virginia, mvee@virginia.edu

This paper identifies several network-layer and transport-layer challenges for the DOE Scientific Computing community in the upcoming decade. Anticipated services include: (i) Terabit/sec (Tb/s) transport-layer throughput for large data movement, (ii) low-latency, predictable-service paths for workflows involving instruments such as photon/neutron sources, and (iii) isolated virtual network slices in a shared-resource environment. These services should be supported across multiple domains with high levels of reliability, performance, automation, and security. Finally, challenges will be posed by the significant changes occurring in the networking community at large, namely Software Defined Networking (SDN) and Network Function Virtualization (NFV).

First, new data-plane protocols are needed to support Tb/s transport-layer throughput. As scientists' laboratories are geographically separated from super-computing facilities, data-movement will inevitably be required. The main challenge is to design network-layer protocols that can operate over parallel physical-layer communications channels, and transport-layer protocols that use the services of these new network-layer protocols. As data rates per channel are approaching theoretical limits, advances in optical communications are relying on parallel solutions such as multi-core fibers, large bundles of fibers, and flexible switching solutions such as colorless, directionless, contentionless (CDC) reconfigurable optical add/drop multiplexers (ROADMs) and FlexiGrid solutions. These technologies can help address the energy consumption challenge as well.

Second, scientists will require the network to offer a choice of paths with different characteristics, e.g., shortest-propagation-delay (shortest-distance) path for low-latency applications. Remote instrument control requires users' applications running on HPCs to process data collected from the instruments in real-time, and allow for configuration changes for the next set of experiments. Workflows involving expensive instruments on a which any single user can obtain only limited time requires predictable performance from the network, and low-latency paths.

Third, the need for isolated virtual network slices will increase. The use of Cloud Computing, in which resources are shared widely among large numbers of users through virtualization, will likely grow in the next decade. Solutions for isolating network bandwidth among multiple users of virtual machines or bare-metal servers will be required. Next-generation solutions for making advance reservations of bandwidth will require advances to support more users and more service options. Further, the multi-domain dimension adds complexity to these control-plane protocols.

Fourth, scientists will require an increasing level of reliability, performance, automation, and security. If the network is not reliable, or if the network performance is unpredictable, scientific discovery will be hindered. As scientists share common resources, authentication methods such as InCommon could be useful. On the other hand, some super-computing facilities prefer more fine-grained authentication.

Finally, the introduction of SDN and NFV create new opportunities as well as challenges for high-performance scientific-computing networks. Should the DOE/SC community create new protocols that are tailored for its needs? If so, how would a network with its own protocols inter-operate with the rest of the Internet? Will other new architectures such as Named Data Networking (NDN) offer new opportunities as well as challenges for scientific applications? Can new transport-layer protocols that leverage WAN accelerators and NDN nodes improve throughput when compared to today's HTCP?

In summary, top-down needs of scientific applications offer new challenges, while the bottom-up opportunities provided by new architectures and technologies will help us develop Tb/s, low-latency, predictable services for the next decade.