SENSE: SDN for End-to-end Networked Science at the Exascale

Lead PI: Inder Monga, Lawrence Berkeley National Lab, imonga@es.net, 510-499-8065
Team: ANL – Linda Winkler, Kate Keahey, Caltech – Harvey Newman, Ramiro Voicu, FNAL – Phil DeMar, LBNL/ESnet – Chin Guok, John MacAuley, LBNL/NERSC – Jason Hick, UMD/MAX – Tom Lehman, Xi Yang, Alberto Jimenez

Abstract: Traditionally, WAN and campus networks and services have evolved independently from each other. For example, MPLS traffic engineered and VPN technologies have been targeted towards the WAN, while the LAN (or last mile) implementations have not incorporated that functionality. These restrictions have resulted in dissonance in services offered in the WAN vs. the LAN. While OSCARS/NSI virtual circuits are widely deployed in the WAN, they typically only run from site boundary to site boundary, and require painful phone calls, manual configuration, and resource allocation decisions for last mile extension. Such inconsistencies in campus infrastructures, all the way from the campus edge to the data-transfer hosts, often lead to unpredictable application performance. New architectures such as the Science DMZ have been successful in simplifying the variance, but the Science DMZ is not designed or able to solve the end-to-end orchestration problem. With the advent of SDN, the R&E community has an opportunity to genuinely orchestrate end-to-end services - and not just from a network perspective, but also from an end-host perspective. In addition, with SDN, the opportunity exists to create a broader set of custom intelligent services that are targeted towards specific science application use-cases. This proposal describes an advanced deployment of SDN equipment and creation of a comprehensive SDN software platform that will help with bring together the missing end-to-end story.

In order to capture and address the variety of end-site network requirements, policy variations, security postures and applications, the proposal will be based on advanced deployment of SDN-capable equipment at two National Labs – Argonne and Fermi, two Universities – Caltech and University of Maryland and two DOE User Facilities – ESnet and NERSC. The proposal recognizes the complex nature of end-to-end problem and proposes to focus its initial deployments on Science DMZ locations at each of the end-sites. In addition, discusses creation of a flexible network operating system platform, based on SDN principles, that can run within the execution environment of a DTN, Science DMZ, campus LAN and/or the WAN. The SENSE Network Operating System (SENOS) provides a base for implementing the workflows and protocols to implement end-to-end orchestration across multiple SDN domains. In addition, our approach takes on the challenge of building the right set of abstractions and feedback mechanisms for the application to interact with the network. We explore the expression of intelligent network services as abstract Intents. Expressing the application requirements as 'Intents' ensures that the applications can specify *what* they need to accomplish, leaving the *how* to the intelligence in the network and availability of resources.

Highly capable networking has become a critical component of the Science Discovery workflow with representative examples that include LHC workflows, Light Source Superfacilities and KBase. Building a feature-rich and programmatic feedback loop between science applications, workflows and network has the potential to dramatically increase the efficiency of distributed science and reduce the time investment scientists spend in order to detect, debug and resolve application, host/disk and network problems.