Authors: Inder Monga (imonga@es.net), Chin Guok (chin@es.net), Eric Pouyoul (lomax@es.net)

End-to-end, multi-domain SDN with SDI bookends

Background: 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 integrating the compute/storage in terms of SDI. In addition, with SDN, the opportunity exists to create a broader set of custom intelligent services, with SDI/SDX that are targeted towards specific science application use-cases.

Research approaches: The approach is to design network abstractions and an operating framework to allow host, LAN, and WAN auto-configurations based on infrastructure policy constraints designed to meet end-to-end service requirements. We propose a comprehensive approach with a focus on usability, performance and resilience through: a) policy-guided end-to-end orchestration of resources; b) auto-provisioning of network devices and Data Transfer Nodes (DTNs); c) intent-based application interfaces providing intuitive access to intelligent SDN services; and d) network measurement, analytics and feedback to build resilience. We argue that this research approach will significantly improve the end-to-end performance of science application workflows, enable creation of end-site driven intelligent services through intelligent use of SDI, and enable new infrastructure solutions for science such as creation of dynamic, distributed 'Superfacilities'.

Benefits and Science Applicability: 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. New approaches to programmability and auto-configuration of resources will get around the current limitations of manual provisioning for the last mile (end-site networks, Science DMZ resources). Additional benefits include a framework for development of custom end-to-end network services.

About the authors: Inder Monga is the CTO is the of Energy Sciences Network (ESnet) that provides innovative network services to National Labs and Universities with the express goal of enabling distributed science research. He has been leading the research and development effort around SDN in ESnet, has strong collaborations with industry to showcase multiple impactful demonstrations like that of Transport SDN[1] and multi-layer SDN[2], is appointed Research Associate by ONF and is on the technical leadership council there, as well has organized multiple workshops on Operational SDN for DOE and NSF.
Chin Guok is Senior Network Engineer in ESnet and is the technical lead of the ESnet On-demand Secure Circuits and Advanced Reservation System (OSCARS) project which enables end-users to provision guaranteed bandwidth virtual circuits within ESnet. He also serves as a co-chair of the Open Grid Forum On-Demand Infrastructure Service Provisioning Working Group.
Eric Pouyoul is Senior System Engineer and his interests include all aspects of high performance big data movement, networking, hardware, software and distributed systems. He has been ESnet lead for designing Data Transfer Nodes (DTN) as defined in the Science DMZ architecture as well as ESnet's work in Software Defined Networking.

---

[1] Abhinava Sadasivarao, Sharfuddin Syed, Chris Liou, Ping Pan, Andrew Lake, Chin Guok, Inder Monga, "Open Transport Switch - A Software Defined Networking Architecture for Transport Networks", HotSDN, ACM SIGCOMM, August 17, 2013

[2] Henrique Rodriguez, Inder Monga, Abhinava Sadasivarao , Sharfuddin Sayed, Chin Guok, Eric Pouyoul, Chris Liou,Tajana Rosing "Traffic Optimization in Multi-Layered WANs using SDN", 22nd Annual Symposium on High-Performance Interconnects, **Best Student Paper Award**, August 27, 2014