

Review of Terabit/sec SDN demonstrations at Supercomputing 2015 and plans for SC16

Tuesday 11 October 2016 11:15 (15 minutes)

The HEP prototypical systems at the Supercomputing conferences each year have served to illustrate the ongoing state of the art developments in high throughput, software-defined networked systems important for future data operations at the LHC and for other data intensive programs. The Supercomputing 2015 SDN demonstration revolved around an OpenFlow ring connecting 7 different booths and the WAN connections. Some of the WAN connections were built using the Open Grid Forum's Network Service Interface (NSI) and then stitched together using a custom SDN application developed at Caltech. This helped create an intelligent network design, where large scientific data flows traverse various paths provisioned dynamically with guaranteed bandwidth, with the path selection based on either the shortest or fastest routes available, or through other conditions. An interesting aspect of the demonstrations at SC15 is that all the local and remote network switches were controlled by a single SDN controller in the Caltech booth on the show floor. The SDN controller used at SC 15 was built on top of the OpenDaylight (Lithium) software framework. The software library was written in Python and has been made publicly available at pypi and github: pypi.python.org/pypi/python-odl/.

At SC 16 we plan to further improve and extend the SDN network design, we plan to enhance the SDN controller by introducing a number of higher level services, including the Application-Layer Traffic Optimization (ALTO) software and its path computation engine (PCE) in the OpenDaylight controller framework. In addition, we will use OpenvSwitch at the network edges and incorporate its rate-limiting features in the SDN data transfer plane. The CMS data transfer applications PhEDEx and ASO will be used as high level services to oversee large data transactions. The scale of storage to storage operations will be scaled up further relative to past demonstrations, working at the leading edge of NVMe storage and switching fabric technologies.

Secondary Keyword (Optional)

Network systems and solutions

Primary Keyword (Mandatory)

High performance computing

Tertiary Keyword (Optional)

Primary author: MUGHAL, Azher (California Institute of Technology (US))

Co-author: KCIRA, Dorian (California Institute of Technology (US))

Presenter: MUGHAL, Azher (California Institute of Technology (US))

Session Classification: Track 6: Infrastructures

Track Classification: Track 6: Infrastructures