

Perspectives on SDN

Roadmap to SDN Workshop, LBL

Philip Papadopoulos

San Diego Supercomputer Center

California Institute for Telecommunications and Information Technology

University of California, San Diego


14 July 2015

What motivates SDN activities @ SDSC/Calit2?

- Can SDN/Openflow simplify and/or improve network management tasks within/to the data center?
 - External Layer 2 connectivity depends on which research network is being utilized
 - Dynes (ESNet) and Internet2 AL2S (+ Coordination with CENIC, California's education network)
 - Must already deal with three different mechanisms
 - Internal networking to the center is multi-vendor (Juniper, Arista, Cisco, Mellanox)
 - Production, high-performance, NSF-funded systems will be hosting high-performance virtual clusters. Isolation via VLANs and Single Root I/O virtualization (SRIOV)
- Being utilized for experiments by particular research activities
- Overall: evaluation of technology, use for research. Not yet pervasive to be considered a solution


A subset of Activities Related to
SDN

A Specific Research Activity: Community of Practice Scientific Expeditions and Infrastructure Experiments for Pacific Rim Institutions and Researchers



PRAGMA Experimental Network Testbed (ENT) Activities


- ENT Backbone
 - TWAREN connected through Pacific Wave
- OpenFlow Network Slicing
 - AutoVFlow deployed and tested (demo)



ENT Activities – International Openflow

Testing: Control Plane

- Distributed implementation via AutoVFlow
- Support for multiple experiments




See Kohei Ichikawa at this Workshop for more details


Prism@UCSD: Campus-scale DMZ. Openflow on some edge devices, but not enabled at core switch



- Climate Modeling
- Ocean Observing
- 3D Electron Microscopy
- Chemistry
- LHC Physics
- Proteomics
- Computer Science
- High Performance Computing
- High Performance data




Prism in Relation to Other Campus Networks



See Tom Hutton at this Workshop

Prism in Relation to Other Campus Networks



“wiring” to LHCONe crosses multiple administrative domains.

Doable, but not scalable

This is not the only usage:
Connections to StarLight,
New activities with SDX
(Joe Mambretti, PI), Pacific
Research platform (PRP), ...


See Tom Hutton at this Workshop



HPC for the 99%


R. L. Moore, C. Baru, D. Baxter, G. Fox (Indiana U), A Majumdar, P Papadopoulos,
W Pfeiffer, R. S. Sinkovits, S. Strande (NCAR), M. Tatineni, R. P. Wagner, N.
Wilkins-Diehr, M. L. Norman
UCSD/SDSC (except as noted)

Comet Architecture (1944 Compute Node)



Notable: Bridging IB Fabric to Ethernet. All nodes are potential 40GbE Sources

Virtual Clusters – Overlay Physical Cluster with User-Owned High Performance Clusters



Isolation of Virtual Clusters

- SRIOV (Single Root I/O Virtualization) for Independent Infiniband Subnets
- Standard VLANs to isolate private Ethernet broadcast domain
 - Virtual cluster nodes can PXE-boot from their virtual frontends

Ethernet switches are openflow enabled:
reticent about using this for first revision of production.

Takeaways

- Activities range from computer science experimentation (actively using openflow) to isolation of virtual clusters (using VLANs, evaluating Openflow)
- A variety of activities utilize Layer2 connectivity options provided by ESNET/Internet2
 - Is there any unification of mechanism on the horizon?
- New Activities (not detailed in this talk) push on high-performance data at distance
- Openflow is really still in the “experimental” networking bin
 - Not in existing network equipment, new purchases are often capable
 - Core equipment lasts 5+ years, it will take a while to be pervasive.
- Personal view: terrific potential, need to be careful to not reinvent capabilities already existing (Routing is a good thing)