GEC16 – Dell GENI Rack Update

An Open Innovation Platform using the GENI-Rack

Rajesh Narayanan
Dell Networking

2013

Agenda

- GENI at Dell
- GENI Rack status and 'Roadmap'
- Research – Deep Programmable Switches
- SDN Innovation Framework
- Collaboration

Dell @ GENI

- GENI Rack
  - On Existing Dell SDN Enabled infrastructure
    - Plugfest participation at Indiana University - InCentre
  - Rack Collaboration with Clemson University
  - Integrated OpenStack and BBN-GRAM

- GEC Participation
  - Attending GENI since Nov 2010 (GEC9)
  - Demonstrated SDP Enabled Dynamic Flow Encryption with University of Houston (UoH)
  - Demo at GEC16 – Application Innovation Framework

Dell GENI Rack Architecture

OpenFlow Switch
Force10
S4810

Five R620 Servers

Management Switch
PowerConnect
7048

TBD

Power

RPS-720

5600W UPS

KMM

Rack frontal view at Clemson

Software and Network Configurations

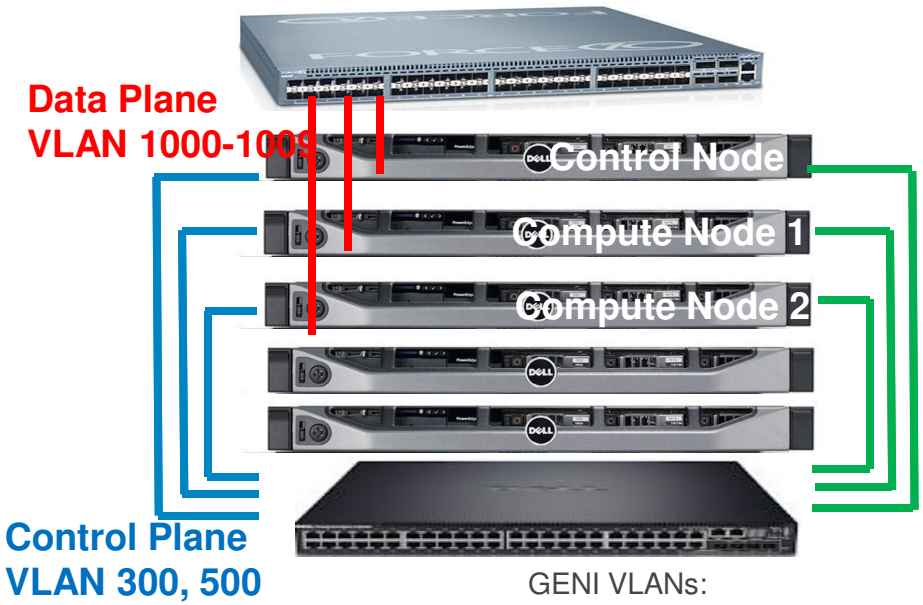

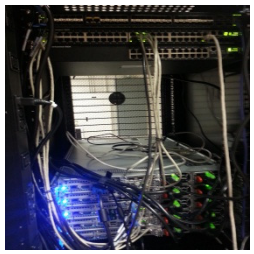

Rack back view at Clemson

GENI VLANs:

- 917 (NetServ)
- 916 (Wisc)
- 3705 (core)
- 3707 (Stanford)
- 3711 (BBN)
- 3715, 3716 (core)
- 973 (GT Regional)
- 1751, 1755-1759 (GENI rack)

Clemson local mapping: http://groups.geni.net/geni/wiki/SiteClemson

- Software
  - OpenStack Folsom (With Quantum)
  - Aggregate Manager
    - GRAM, Foam
  - Server provisioning
    - For Virtual (/Physical) Hosts:
      - GRAM
    - For VMs: OpenStack
      - Nova
  - Federation
    - GPO: Standard procedure
  - Network Connectivity
    - Start with local testing, then GENI rack vlan
  - Monitoring
    - GMOC

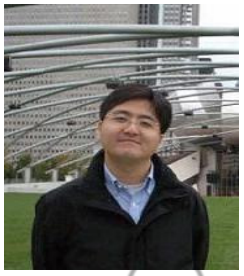

S4810 – OpenFlow Enabled Hybrid Switch

purpose-built for high-performance data center and computing environments

S4810 Advantages

- **Low latency (~700ns)**
- **Low power consumption (280W)**
- **Scalable using front port stacking up to 3 units**
- **Feature rich Layer 2 and Layer 3 protocols**
- **Low cost**

- 1RU high-density 10/40GbE top-of-rack (ToR) switch with 48 dual-speed 1/10GbE (SFP+) ports
- Four 40GbE (QSFP+) uplinks (totaling 64 10GbE ports with breakout cables)
- Multiple OF Instances with separate DPIDs

OpenFlow 1.0 protocol support

The Dell GENI Rack – Vision*

Version 1.0

First Version: S4810, 5Servers Mgmt. Switch

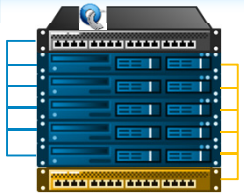

SDN Enabled Storage + Enhanced L2/L3 Access

S4810, 5-10 Servers Mgmt. Switch+ SDP Storage

** This is a vision for GENI-Rack. Not a Dell Product Roadmap.*

S4810, 5-10 Servers Mgmt. Switch + SDP

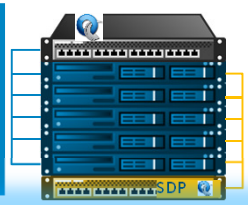

SDP + Innovation Framework

Every network component is SDN configurable

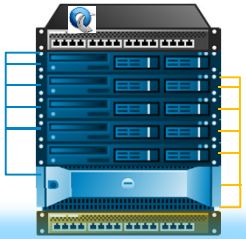

Fully Integrated

SDP Split Data Plane Module


S4810: OpenFlow Enabled Switch


PowerConnect 7024 Switch


R620 Servers. Maximize Memory and Storage

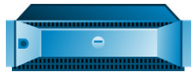
EqualLogic Storage 1-3TB (?)

Research

All Flows are Not Created Equal

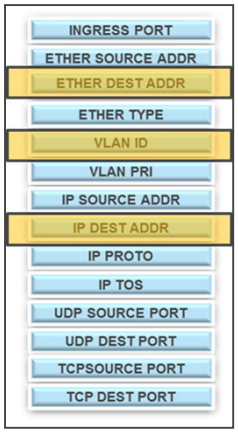

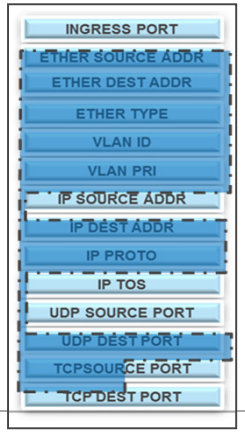

❑ Coarse Flows

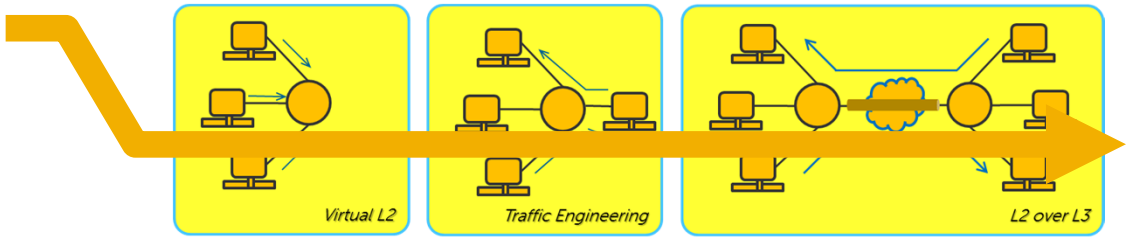

Allow, Drop, Reroute, Multicast

(Essentially L2/L3/PBR)

❑ Granular Flows

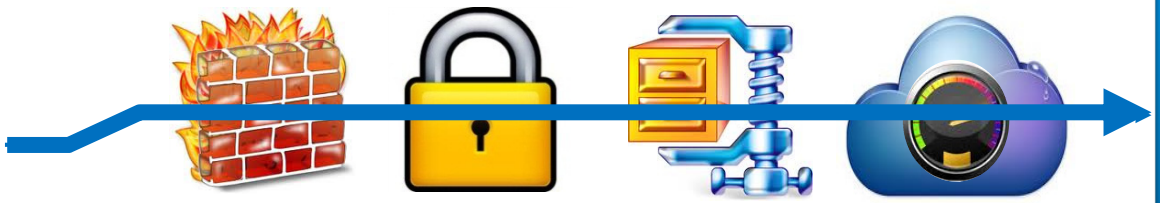

Encrypt, Compression, Acceleration, Buffering, Watermarking, Custom

(Rich Action-set Missing)

****All Flows are not Created Equal***

Split SDN-Data-Plane (Divide and Rule)

- Coarse Flows → Merchant Silicon Switch
- Granular-Flows → Programmable multicore NPU
- Controller manages both OF-Switches
- Coarse-Flows identified for granular processing are redirected to module
- FlowBursting → Apps Insert many 10K flows

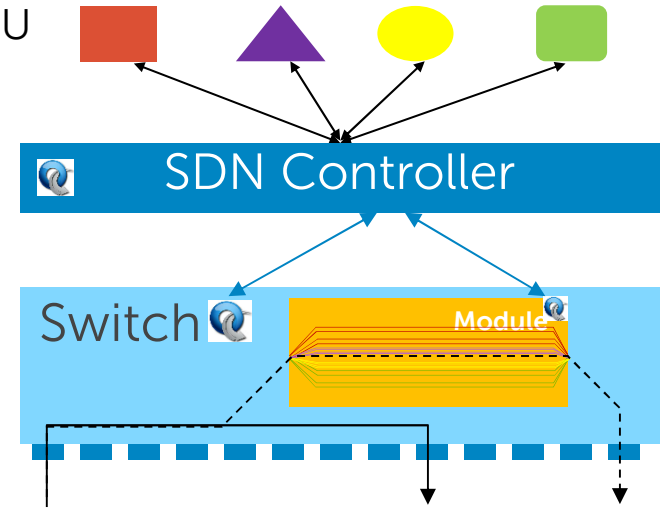

- ***http://www.ewsdn.eu/presentations/EWSDN_Rajesh.pdf***
- ***The SDP modules are still experimental and not generally available***

SDN Application Innovation Framework

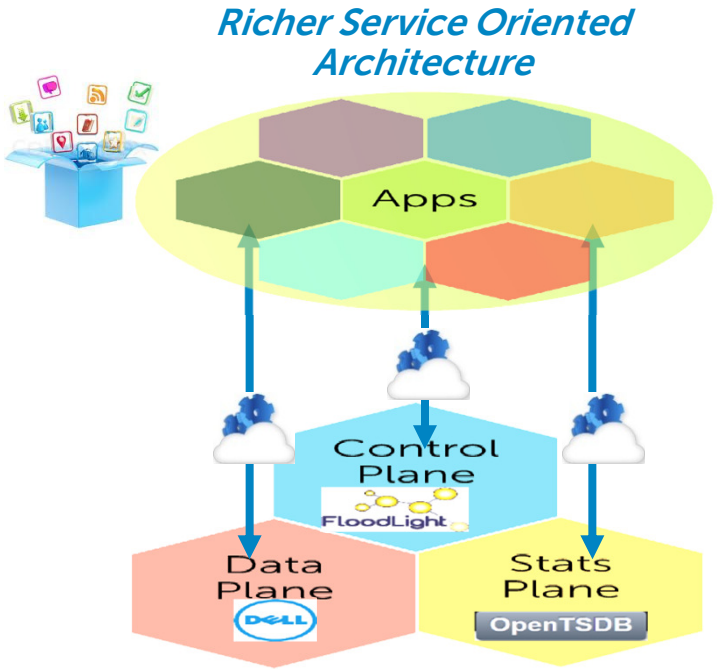

- Open Control Plane
  - New APIs Leverage Extensible data-plane
  - Beyond Policy Based Routing
- Extensible Data Plane
  - Open Data Plane architecture (Split Data Plane)
  - New packet processing pipelines
  - Dynamically insert data-plane apps w/o Extensions
- Statistics Plane
  - Asynchronous Statistics
  - Unlock Volumes of Statistics
  - Big Data, Correlate, Visualize

Time Series Database

Application Framework

- Extensibility Without Extensions
  - No controller-extensions needed
  - Network Functions Appear as virtual interfaces
    - E.g. ICMP proxy → **`'icmp0'`**
  - Extensions become path-property
- Statistics Plane
  - Asynchronously Externalize Network stats
  - Configurable Time-Slice to Capture
    - Assumes – 90% of Northbound Apps should be fine with 1 second granular
  - REST-Interfaces to Northbound Apps

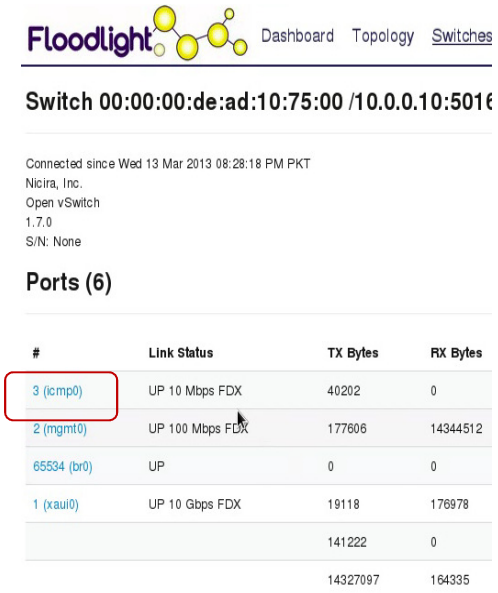

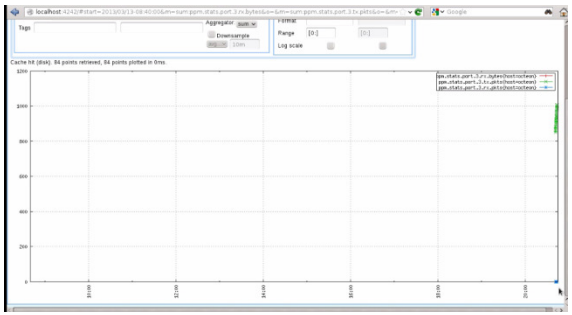

UNIVERSITY of HOUSTON
YOU ARE THE PRIDE
COLLEGE of TECHNOLOGY

SDN Innovation Platform

Rajesh Narayanan (Dell), Fahd Gilani (XFlow Research), Deniz G and Leven : (UoH)

Project sponsored by DELL
XFLOW RESEARCH
Special thanks to Michael Weaver from CAVIUM

What is SDN Innovation Platform?

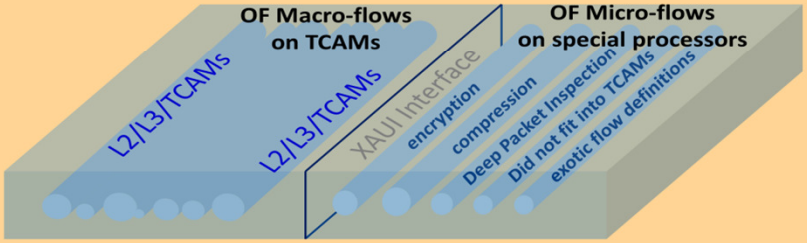

Developing Applications

- As easy as socket programming – *not at end points but on network nodes*
- Popular programming languages
- Hands-on experimentation
- *Sky is the limit*

Deploying on GENI

- Booting capabilities are similar to GENI
- Better way to generate applications and experiments on the physical network nodes
- Network visualization
- Opening up new research fields

ICMP-Proxy

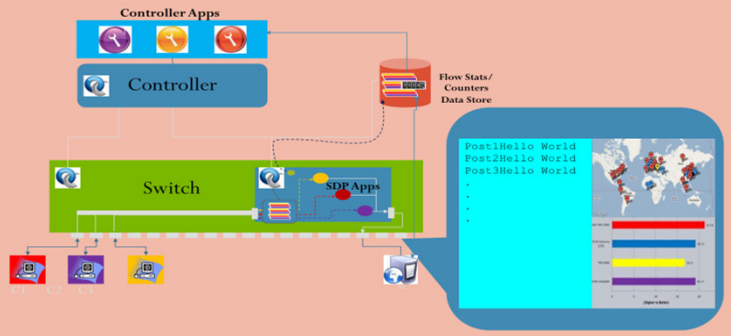

Collaborating with Universities

- GENI – Rack
  - Extending Rack capabilities
  - Opportunities for Storage
  - L2/L3 Extensions
- Application Innovation Framework
  - Build new packet processing functions
  - Contribute to an SDN App Store
    - Community Apps Development
    - OpenSource
  - Use-Cases to Enhance GENI