



Deploying OpenFlow

Srini Seetharaman
Deutsche Telekom
(with help from Paul and Masa)

Open Networking Summit
April 2012

Goals


- Understand the challenges
- HOWTO for deployment
- Using examples, get tips on deployment issues

Outline

- Deployment challenges
- LAN deployment
 - Howto
 - Example: Stanford
 - Intro
 - Lessons learned
- WAN deployment
 - Howto
 - Example: GENI
 - Intro
 - Lessons learned
- Summary

Illustration of topology discovery

- OpenFlow controller view is not always complete. For instance, what does the controller see here?



Deployment challenges

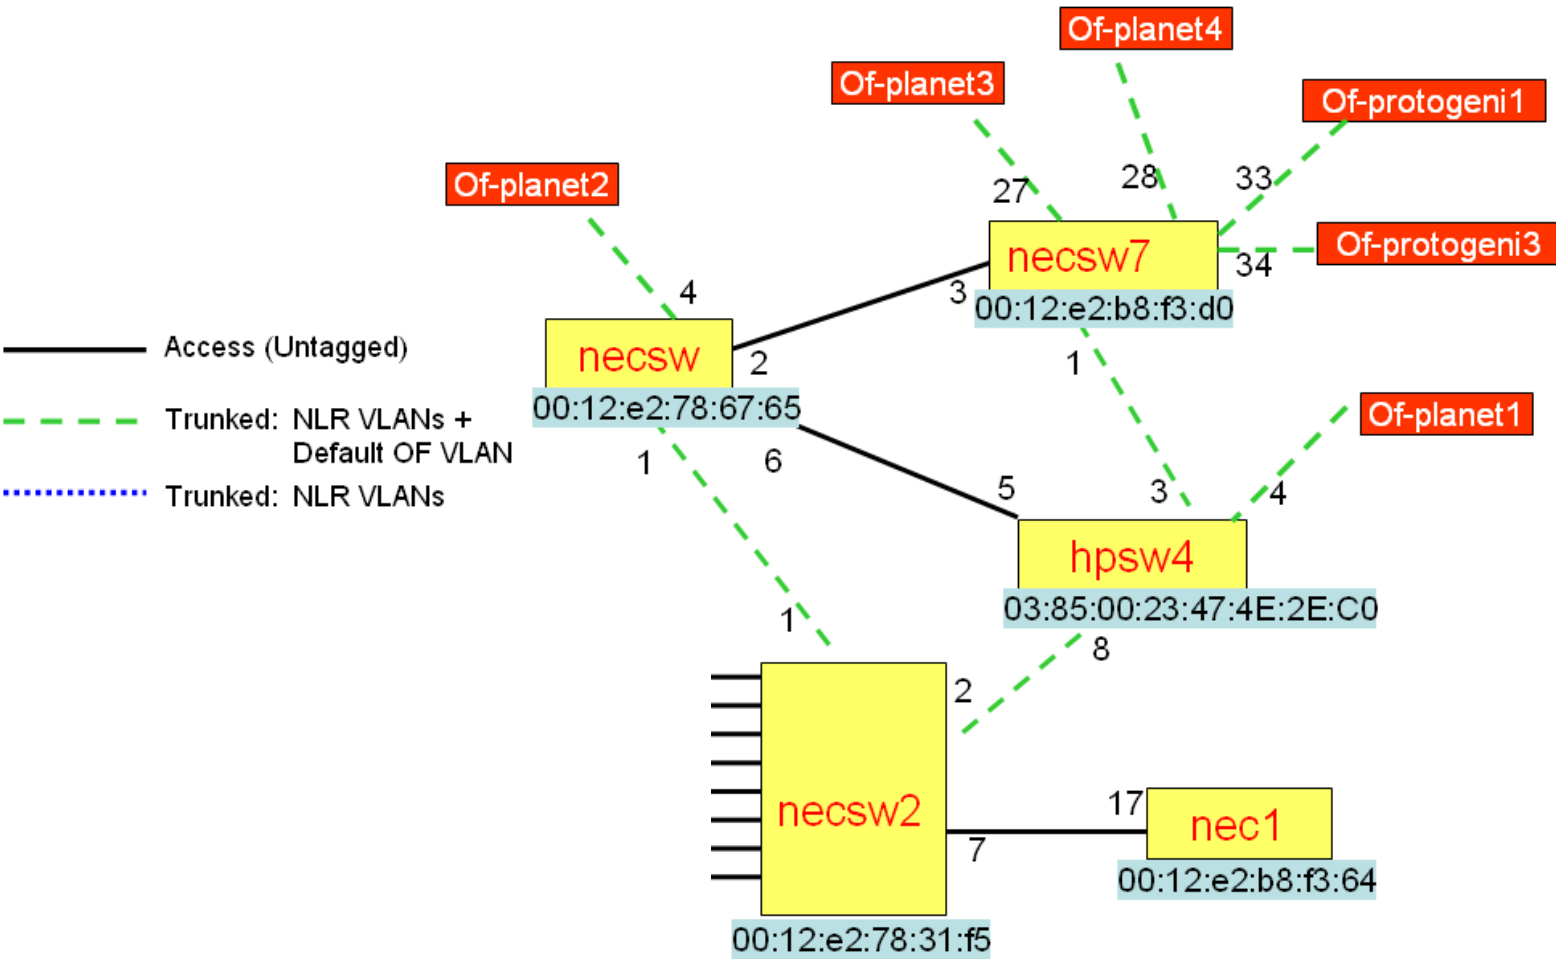
- Production, experimentation, or both?
- LAN vs WAN?
- Network design will need to accordingly handle the control + dataplane virtualization, topology discovery and other functions based on the purpose of the deployment.
 - How to structure the data plane and the control plane?
 - How to perform topology discovery over OpenFlow-enabled switches?
 - What happens when you have a non-OpenFlow switch inbetween?
 - What if there are two islands connected to same controller?
 - What legacy functions need to be used (e.g., STP)?

LAN deployment – Howto

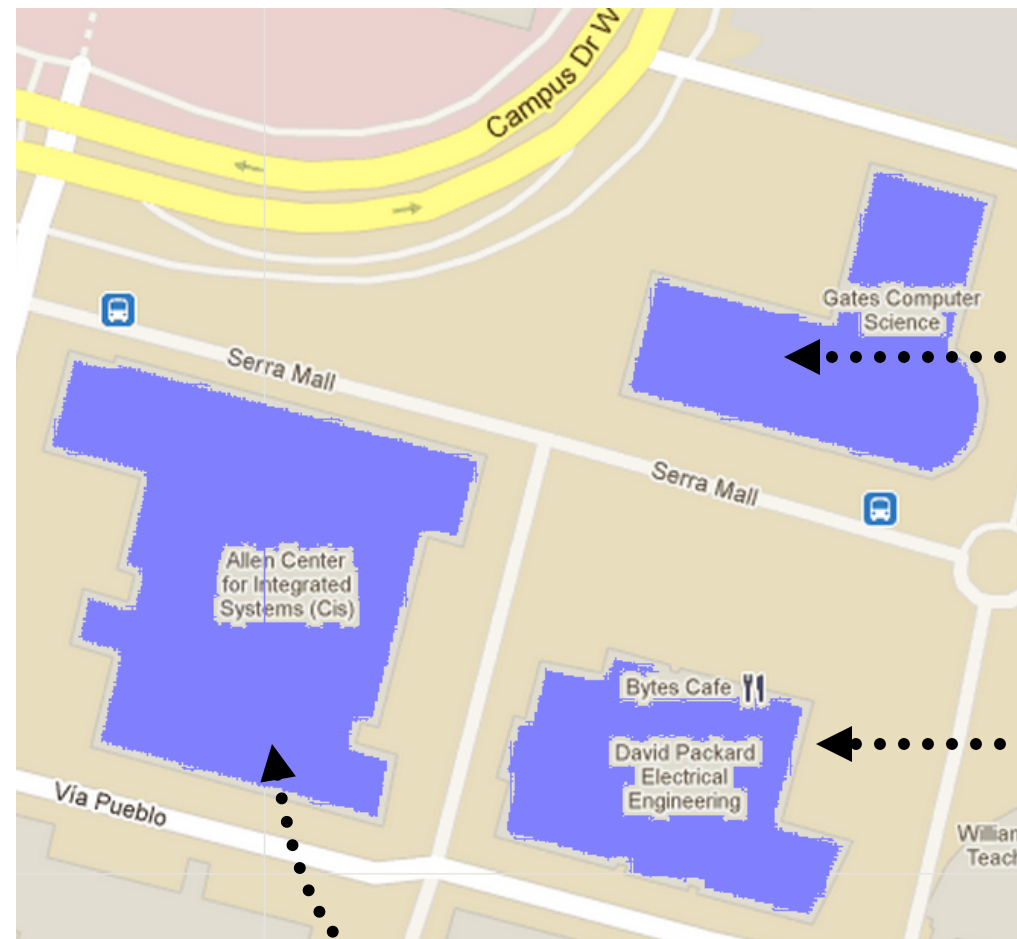
Staged Deployment of OpenFlow

1. Add experimental VLAN
2. Enable OpenFlow for Exptl VLAN
3. Configure Controller for new network
Verify correctness and performance
4. Add new Production subnetwork
5. Gradually add/move users to new subnet
Verify reachability
6. Enable OpenFlow for new subnet
7. Slice the network

Example #1: Stanford demo network



Example #2: Stanford production network



- Wired network in Gates 3A for 18 users over 7 rooms
- Open wireless network with 35 APs and max of 20 users in Gates bldg
- Demo network (and GENI substrate) with 7 switches and few NetFPGA nodes
- VICCI cluster with 70 nodes

- Upcoming wifi deployment in Packard building

- Production deployment in CIS bldg with 7 switches + 17 APs, serving 50 PCs + max of 15 wifi users

Lessons learned at Stanford

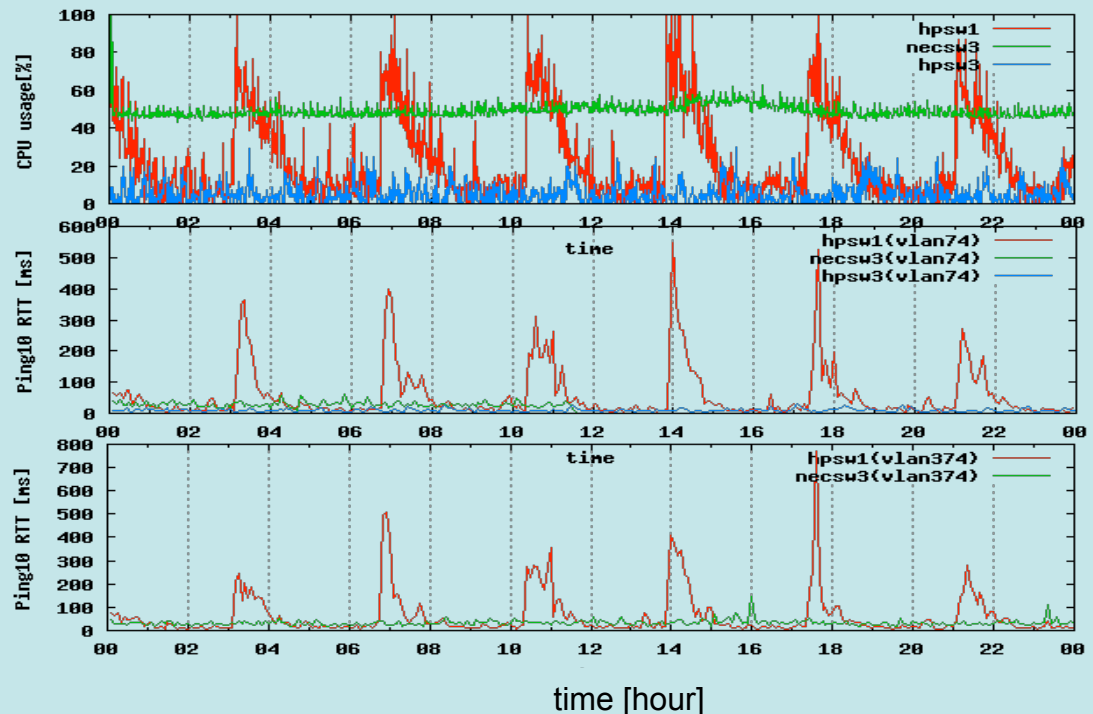
- Legacy configs and protocols:
 - STP is used in Enterprise networks
 - OF vendor implementations interact badly with STP
 - Other protocols like CDP, LLDP confuse debugging

- Debugging is a dark art because of less understood components and correlated effects

CPU Usage [%]

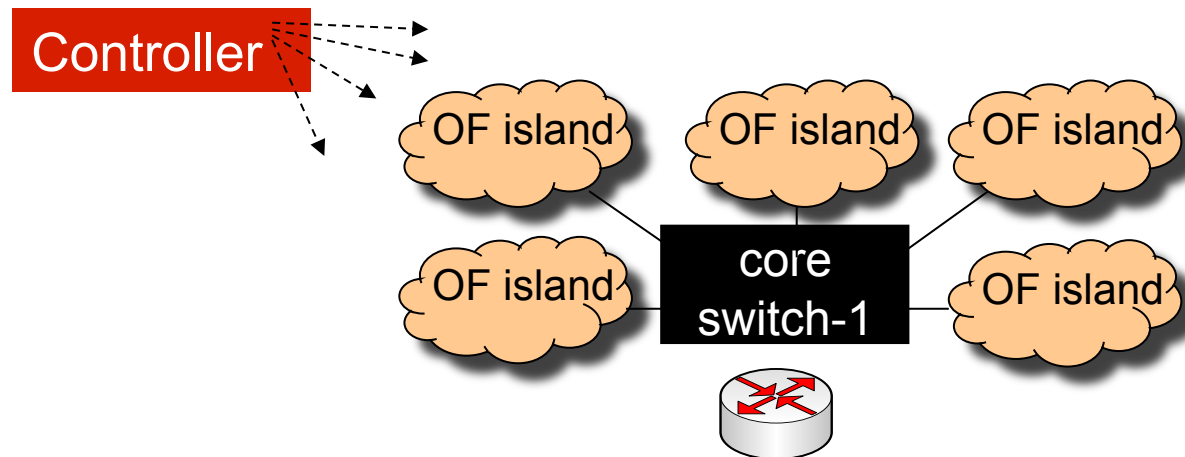
5min Avg Flow Setup Time [ms] (VLAN74)

5min Avg Flow Setup Time [ms] (VLAN374)



Lessons learned at Stanford (contd.)

- Controller issues:
 - No performance issues. Just beta quality software that crashes or generates too much log
- Switch issue:
 - CPU problems on switches (weak CPUs)
 - Flow rate on core cannot be managed by today's OF switches:
 - Workaround: No OpenFlow on core switch!



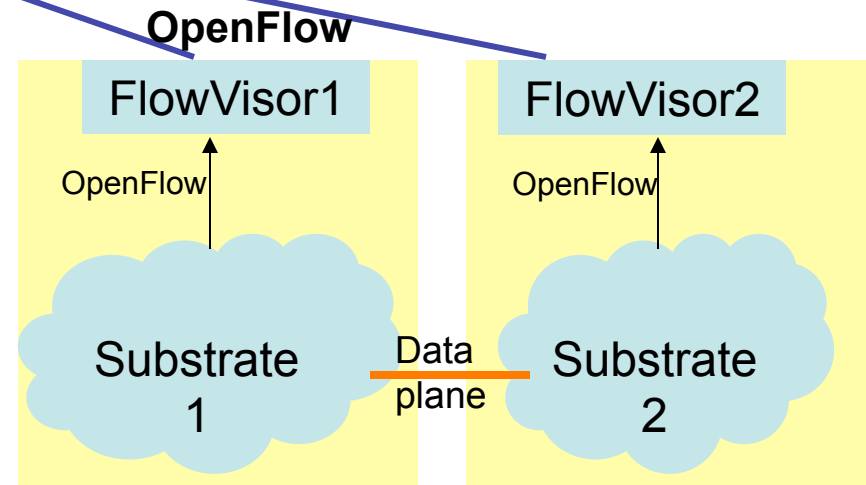
WAN deployment – Howto

- Diverse needs, e.g.,
 - Flat L2 network for experimentation
 - Inter-DC connectivity over L3VPN



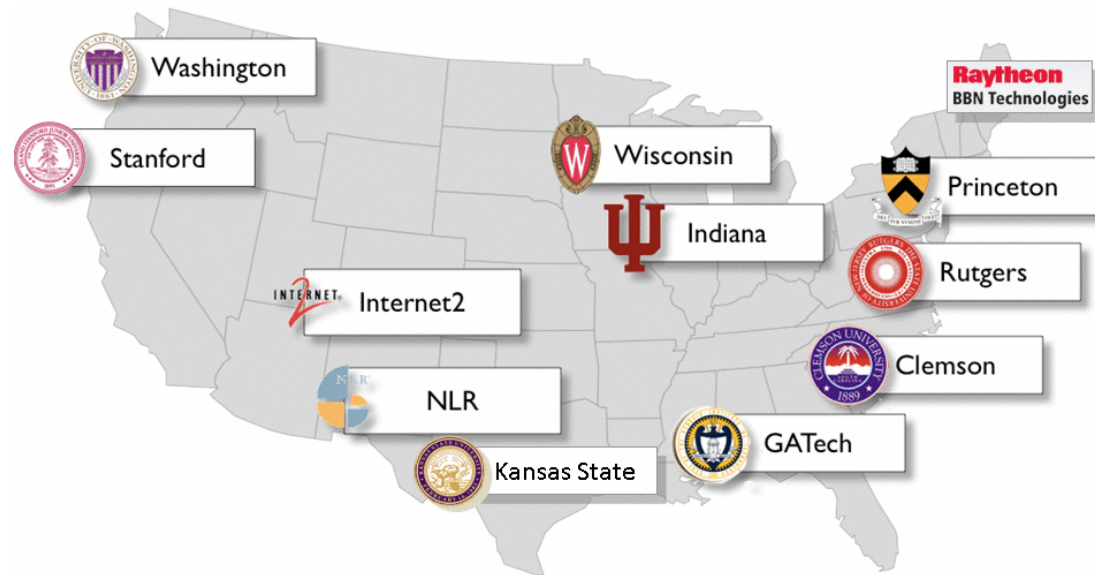
Controller

- Mixed vendors, esp with different admin entities
- Segmented control plane;
Stitched data plane

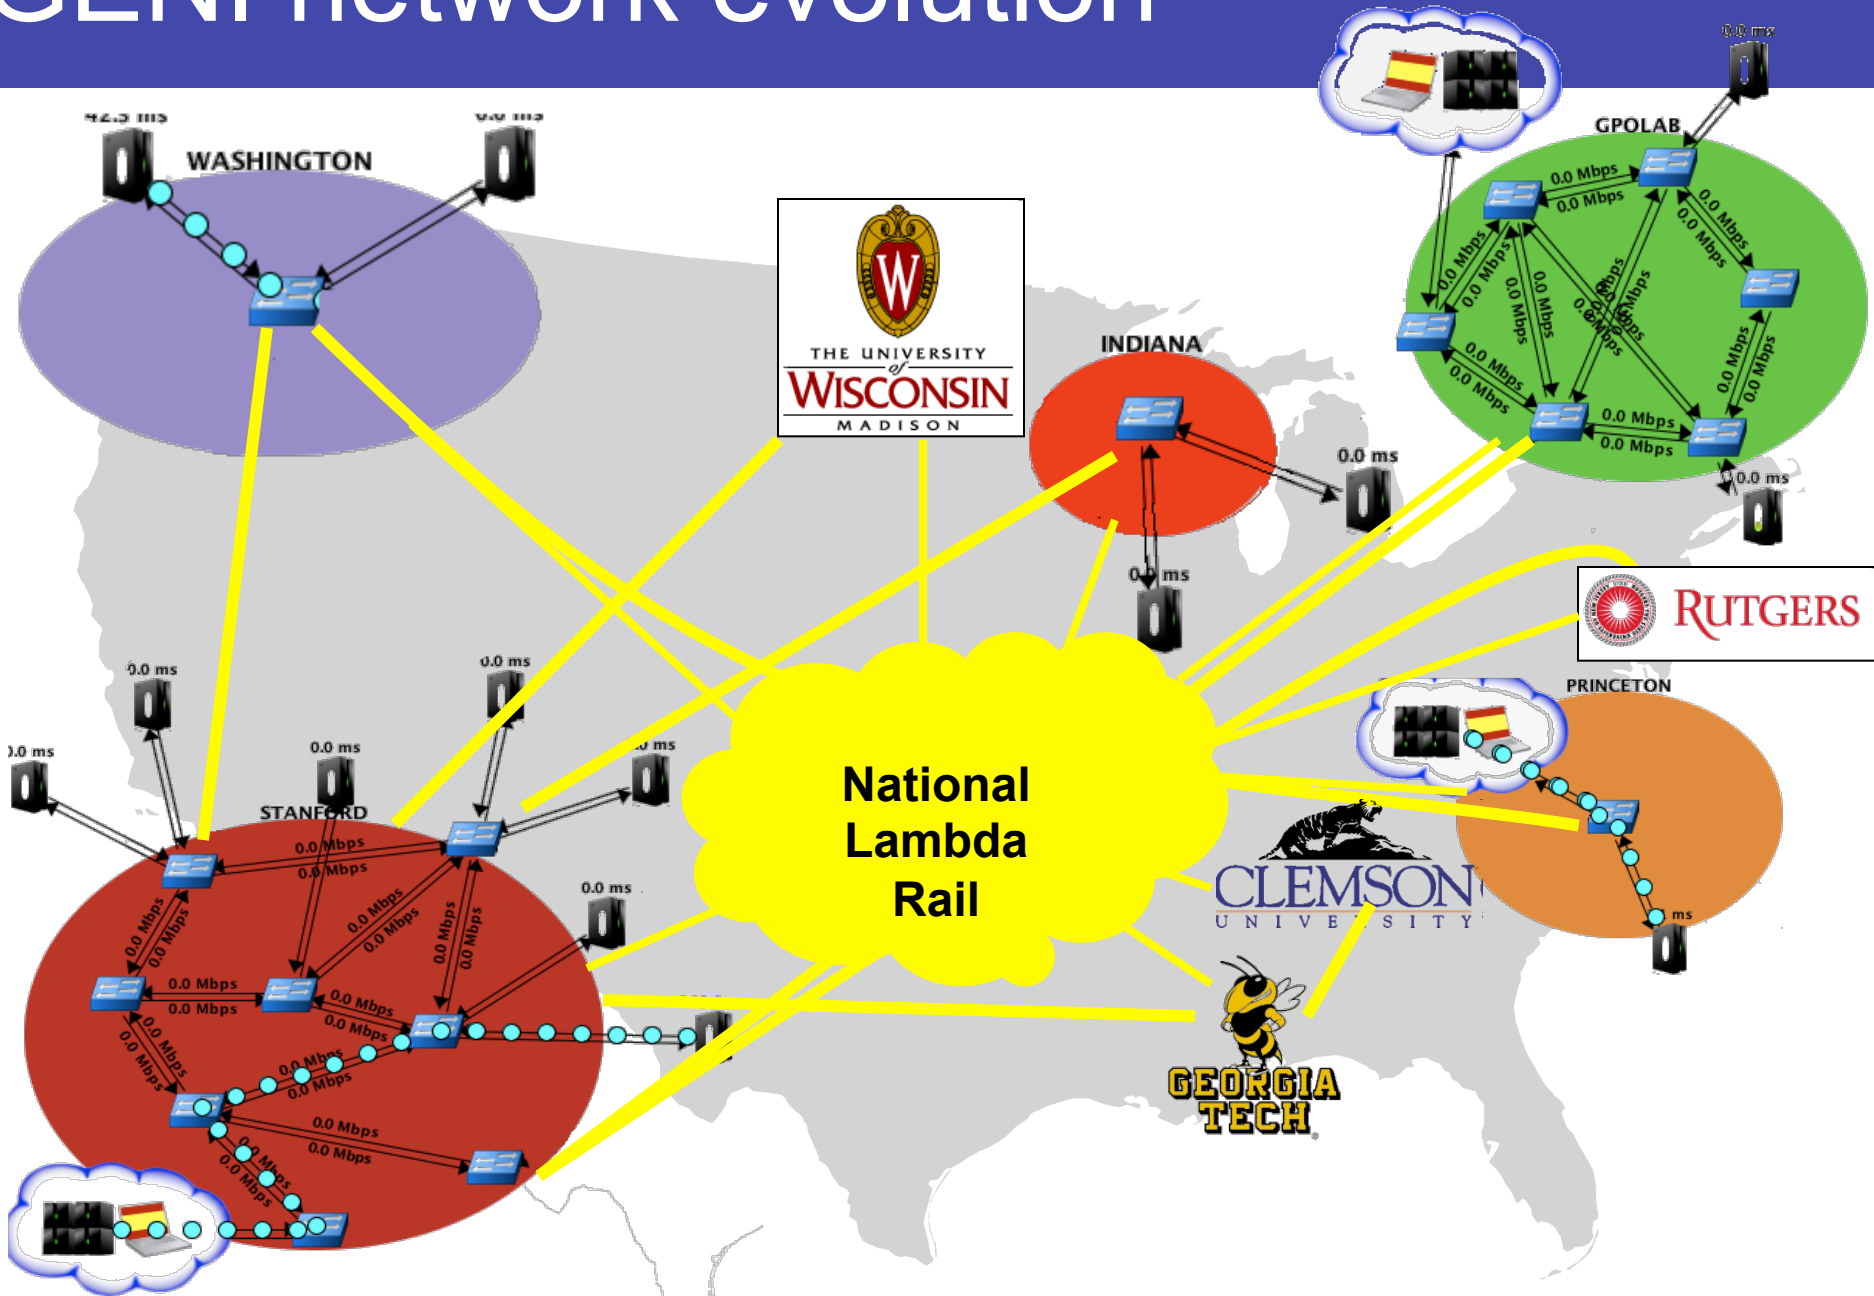


Example #3: GENI network

- Global Environment for Networking Innovations
- Since 2008
 - 8 campuses
 - US-wide network
 - Internet2/NLR
- Several demonstrations on distributed substrate
 - Lots of issues (OF and non-OF)



GENI network evolution



Lessons learned in GENI network

- OpenFlow over Q-in-Q
 - OpenFlow routing is unaware and sends traffic with same MAC address in both direction, causing perpetual learning and CPU inflation
- Moving uplinks for 1 switch, while being pointed to the same controller (causing two islands)
 - Causes controller learning to oscillate between the 2 uplinks
- Bad interaction with legacy protocols
 - LLDP and STP are treated differently with different switches
- Loop in OpenFlow network being exposed to non-OF side
 - Miscommunication between the aggregate operator and the experimenter during testing phase
- Loop across backbones
 - Same campus connected over NLR and Internet2

Summary

- OpenFlow moved to commercial world 1 year ago
 - Better and more support forthcoming
 - More deployments coming up!
- Announced proof-of-concept deployments:
 - Production networks at Stanford
 - Distributed, US-wide GENI substrate
- Lessons:
 - New network management/control paradigm
 - Many issues not foreseen

