OpenFlow in production networks.

Paul Weissmann, Srini Seetharaman
30.08.2011 CHANGE/OFELIA Summer school

Erleben, was verbindet.

OpenFlow in production networks. Agenda.

1. Introduction
   - Current status
   - Ecosystem
2. Stanford
   - Infrastructure
   - Issues
3. GENI
   - Infrastructure
   - Problems & experiences
4. Discussion

Introduction.
Current OpenFlow status.

- OpenFlow is deployed in hundreds of networks/campuses

- OpenFlow specification:
  - Current: OpenFlow 1.1
  - Implemented: OpenFlow 1.0
  - Forthcoming: OpenFlow 1.2 (just being standardized)
  - Standardization moved to ONF

Introduction.
Ecosystem.

- Ecosystem is growing
  - Moving to commercial
- Vendors
  - Better support
  - More vendors
  - Commercial controllers
- Standardization
  - Stanford → ONF

Introduction.
Controller vendors.

- Controllers/software
  - Enterprise
    - BigSwitch Pico
  - Data centers
    - Nicira NVP/ONIX
    - NEC PFC
  - Many experimental ones
    - Beacon, NOX, Maestro, Trema

Introduction.
Switch vendors

Juniper MX-series

NEC IP8800

WiMax (NEC)

HP Procurve

Netgear 7324

PC Engines

Pronto 3290/3780

Ciena Coredirector

More coming soon...

Introduction.
Standardization.

- Standardization: ONF
  - Founded 2011
  - 6 founding members
  - 43 members, including vendors

As a nonprofit, mutual benefit corporation, the purpose of the Corporation is to promote the development and use of Software-Defined Networking

- Kernel WG:
  - First WG to start
  - Take-over of OpenFlow specification from Stanford
  - Focus on extensibility & modularity

Introduction.
Deployment.

Network design will need to accordingly handle the virtualization, topology discovery and other legacy functions.

- 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?

Stanford. Overview.

- Production network since 2008 - CS network as live testbed
- OpenFlow early versions, many vendors and systems
- Since we left (2011): growing and expanding

Stanford.
Infrastructure.

- Computer science network
  - HP ProCurves (+ NEC, Quanta)
  - 1,000 users, 3,000 ports
  - OpenFlow into production
- Wireless
  - Parallel WiFi deployment
  - Open network, 35 APs
- Electrical Engineering network
  - Extension

Stanford.
Topology (CS).

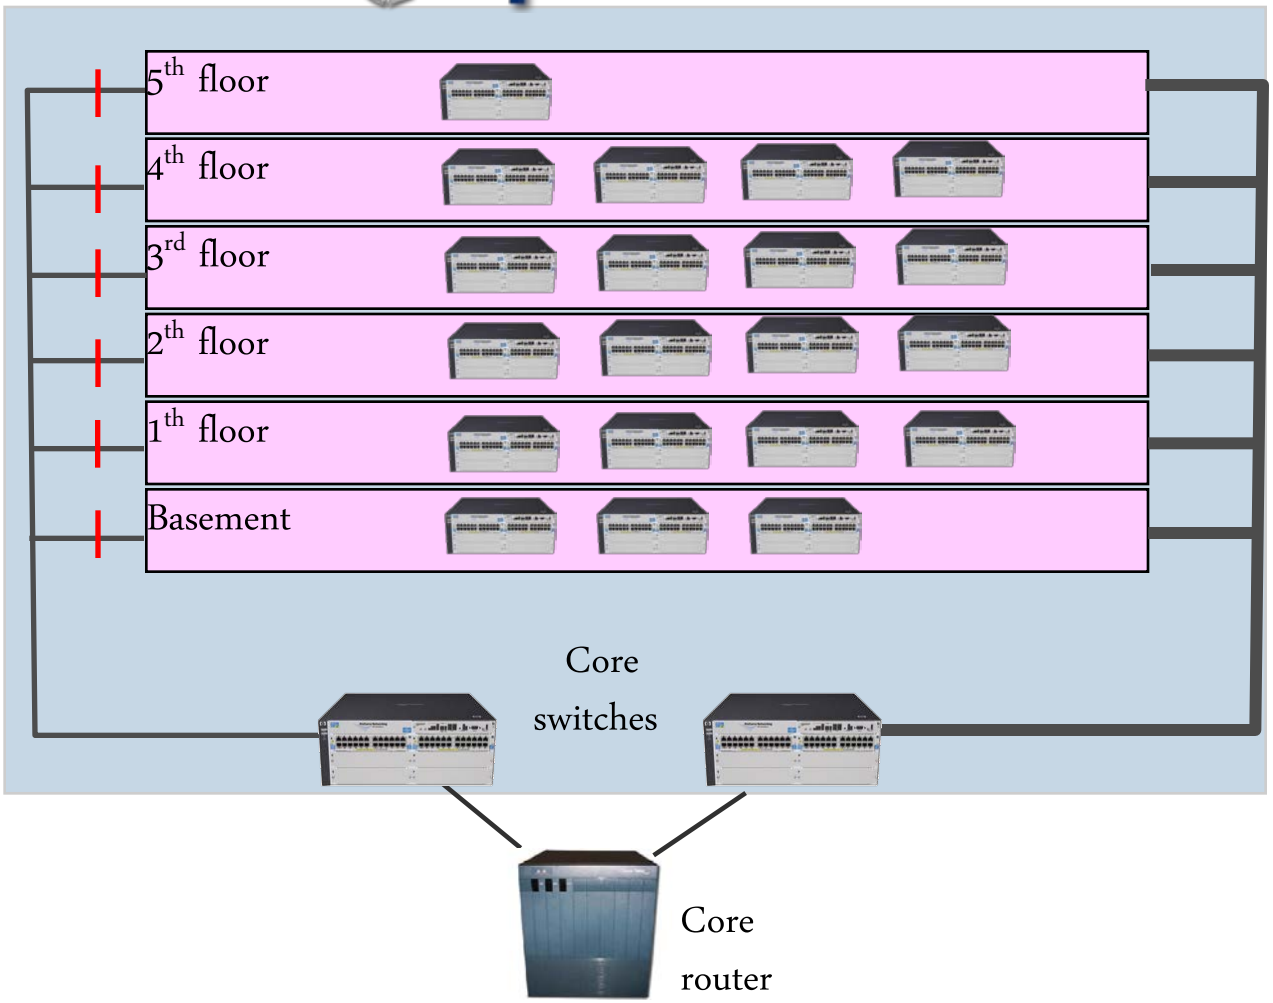

Each switch has 1+3 uplinks

1G backup

LACP 3 x 1G

Stanford.
Plan.

- Phased deployment plan
- Four different networks
  1. Test 2. Demos 3. Group 4. CS
- Close collaboration with IT
- Phases
  - Test/demo networks
  - Group network (30 nodes)
  - More VLANs/switches
  - Core switches
  - Whole building

Stanford.
Network engineering.

- Running system had to be upgraded
- Firmware updates (~20 rounds)
- Network not in best shape
- 1,000 users affected
- Demos/group affected

- Took almost a year to get right

Stanford.
Issues.

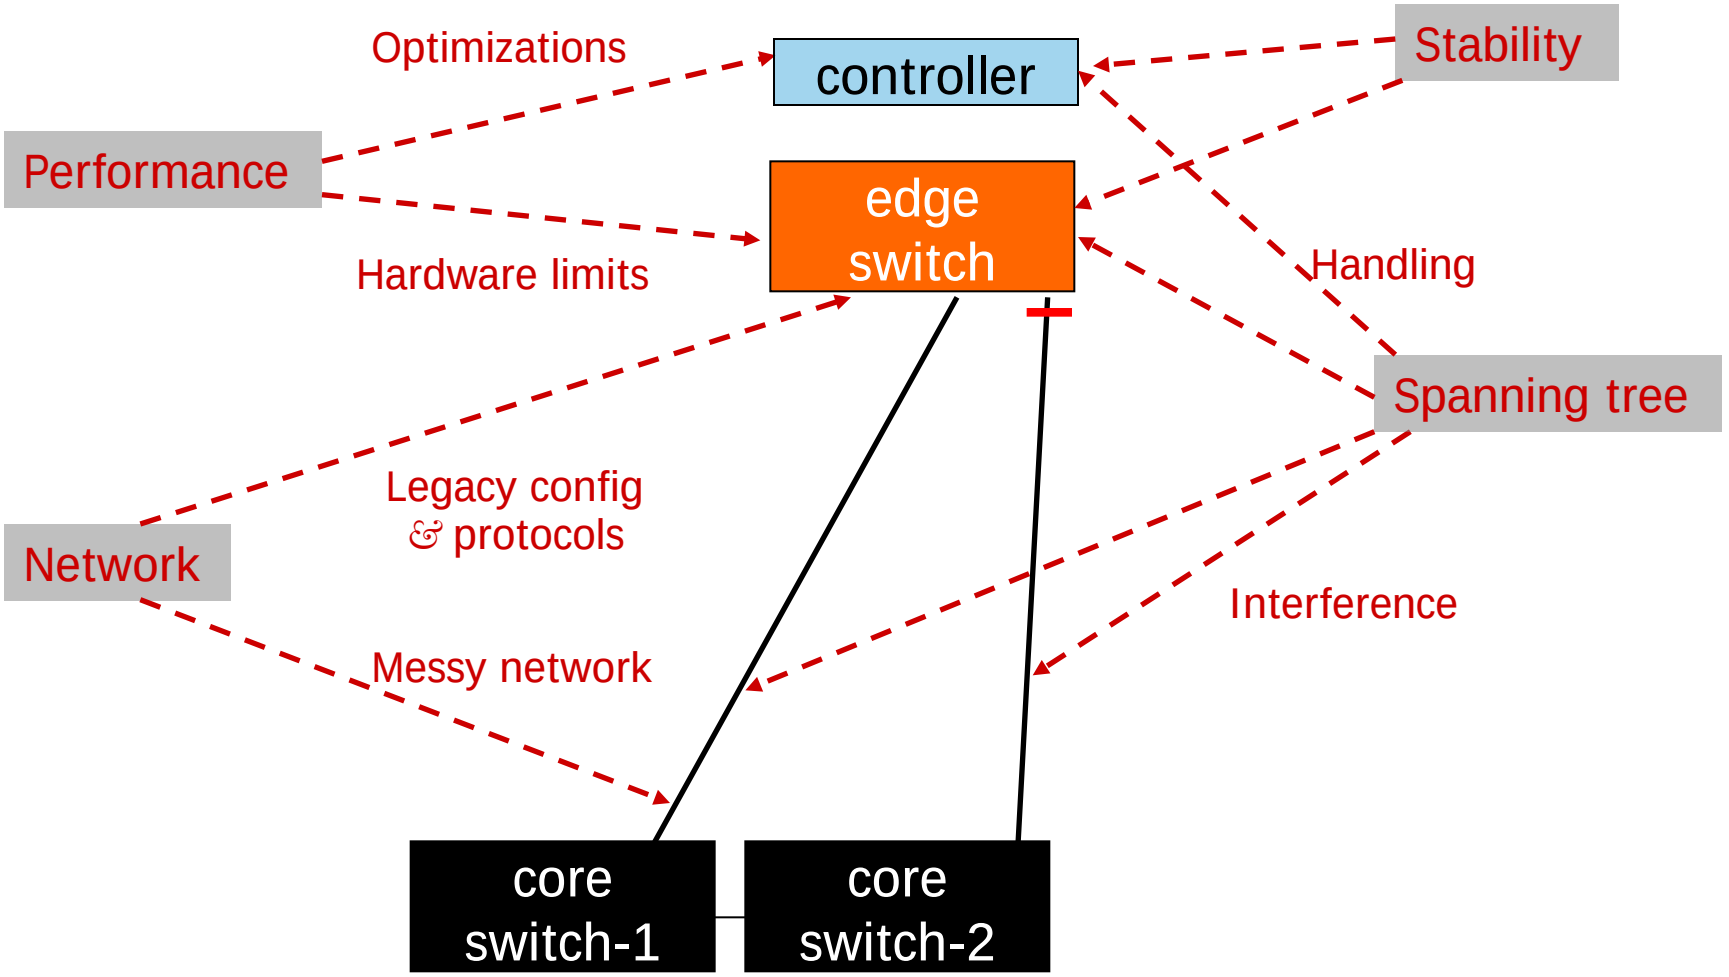

Stanford.
Performance.

- Flow setup rate
  - Switch can handle 150 flows/sec
  - Enough for edge
  - Insufficient for core
  - Optimizations!

    (pro/reactive)

- No performance problems on controller (just beta software quality)
- CPU problems on switches (weak CPUs)

Stanford.
Performance (contd).

- Workaround for flow rate limit on core:

  No OpenFlow on core!
- Controller can handle it, but:

Plan

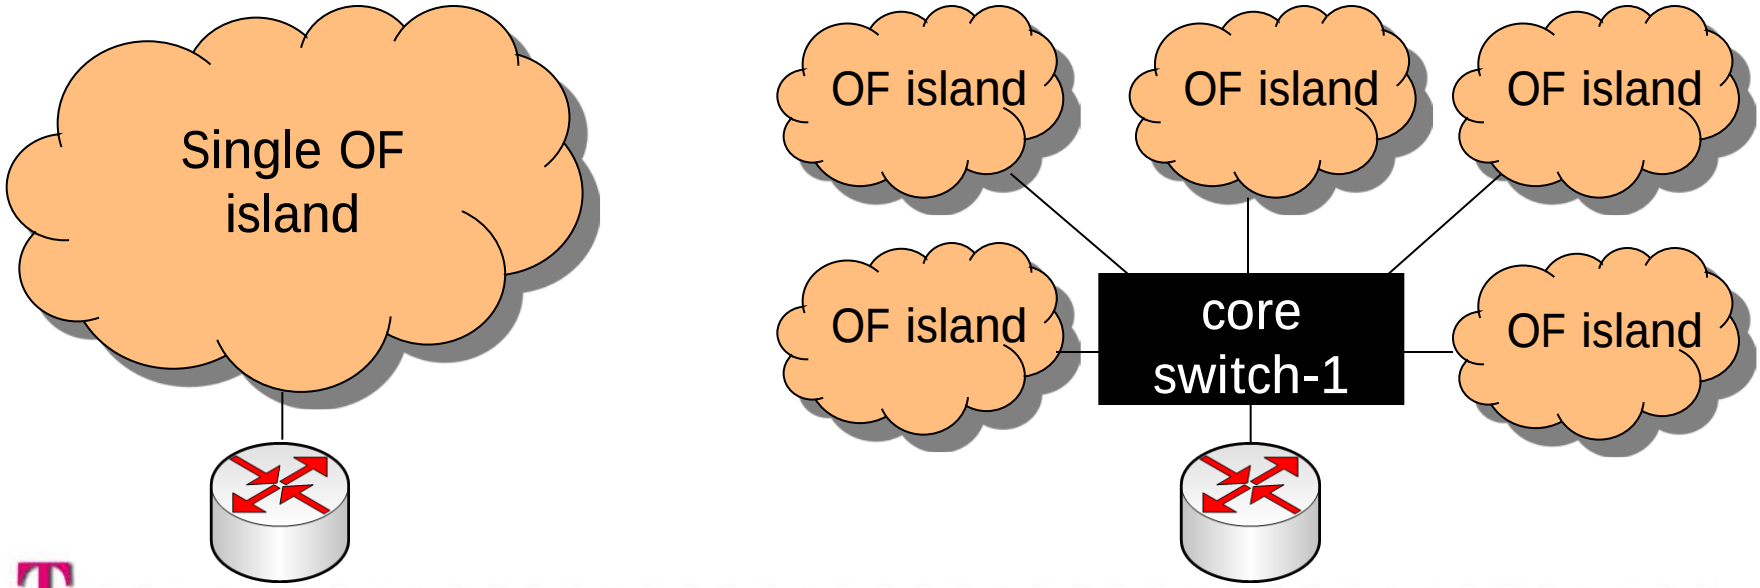

Debugging is still a dark art

CPU Usage [%]

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

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

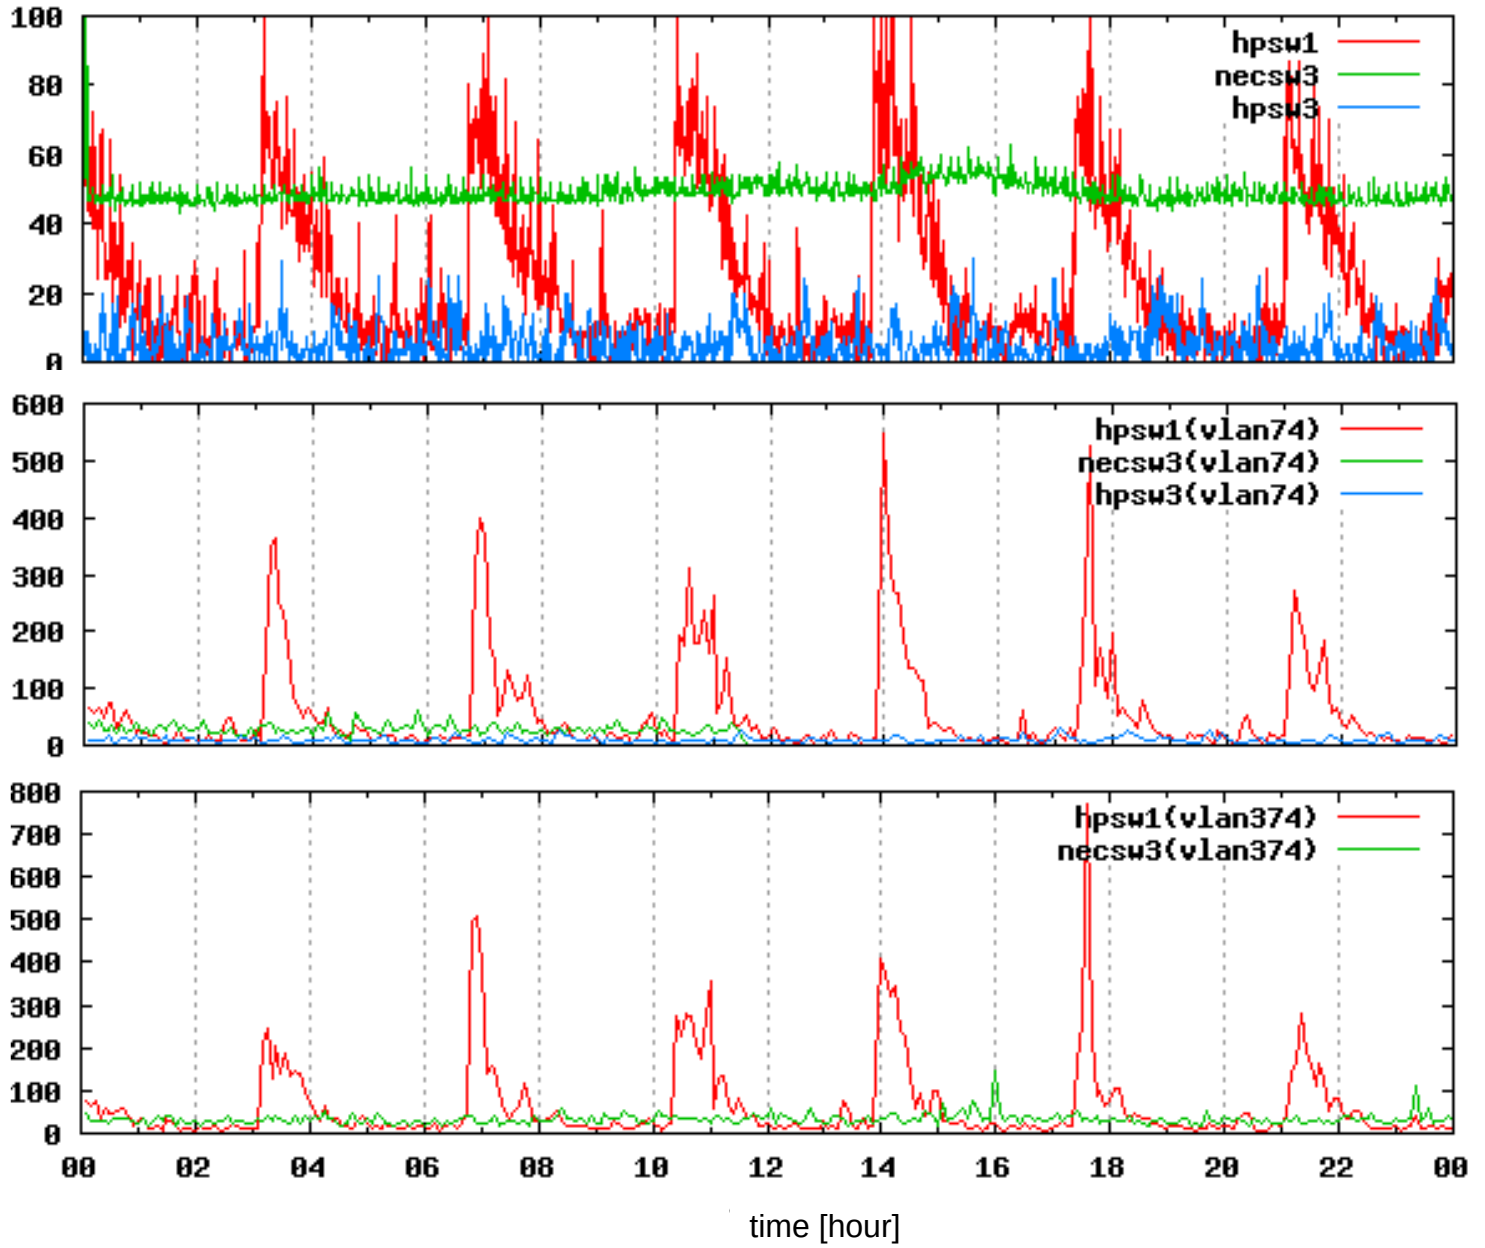

Stanford.
Spanning Tree.

- Legacy networks use STP
- Everyone hates (touching) STP

- HP ProCurve
  - OpenFlow gets packet *after* STP module
  - OpenFlow uses that topology
- “What’s that?” BPDU, LLDP, CDP ...

- Still no convergence on right behavior

Discussion.
Security.

- *Controller is the big red button*
  - Information on all traffic

- Control channel not encrypted
  - Controller↔switch
- Controllers/switches do not authenticate

- MitM, DoS, eavesdropping

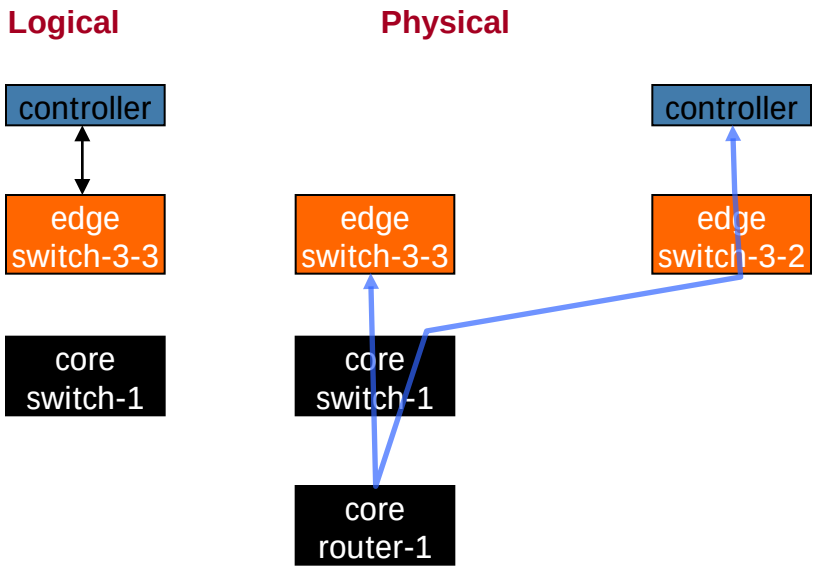

Stanford.
Wrap up.

- Fixed numerous bugs
  - Firmware, controllers, software, OF/logic

- 30 new firmware releases
- Five controllers, dozens releases
- Network engineering, complicated network

- Next steps:
  - IT department takes charge (now)
  - Extension to core and DCs, WiFi, Residences

GENI.
Overview.

- Global Environment for Networking Innovations
- Since 2008
  - 8 campuses
  - US-wide network
  - Internet2/NLR
- 5-6 vendors

- Demonstrations on distributed substrate
  - Lots of issues (OF and non-OF)

GENI. Network.

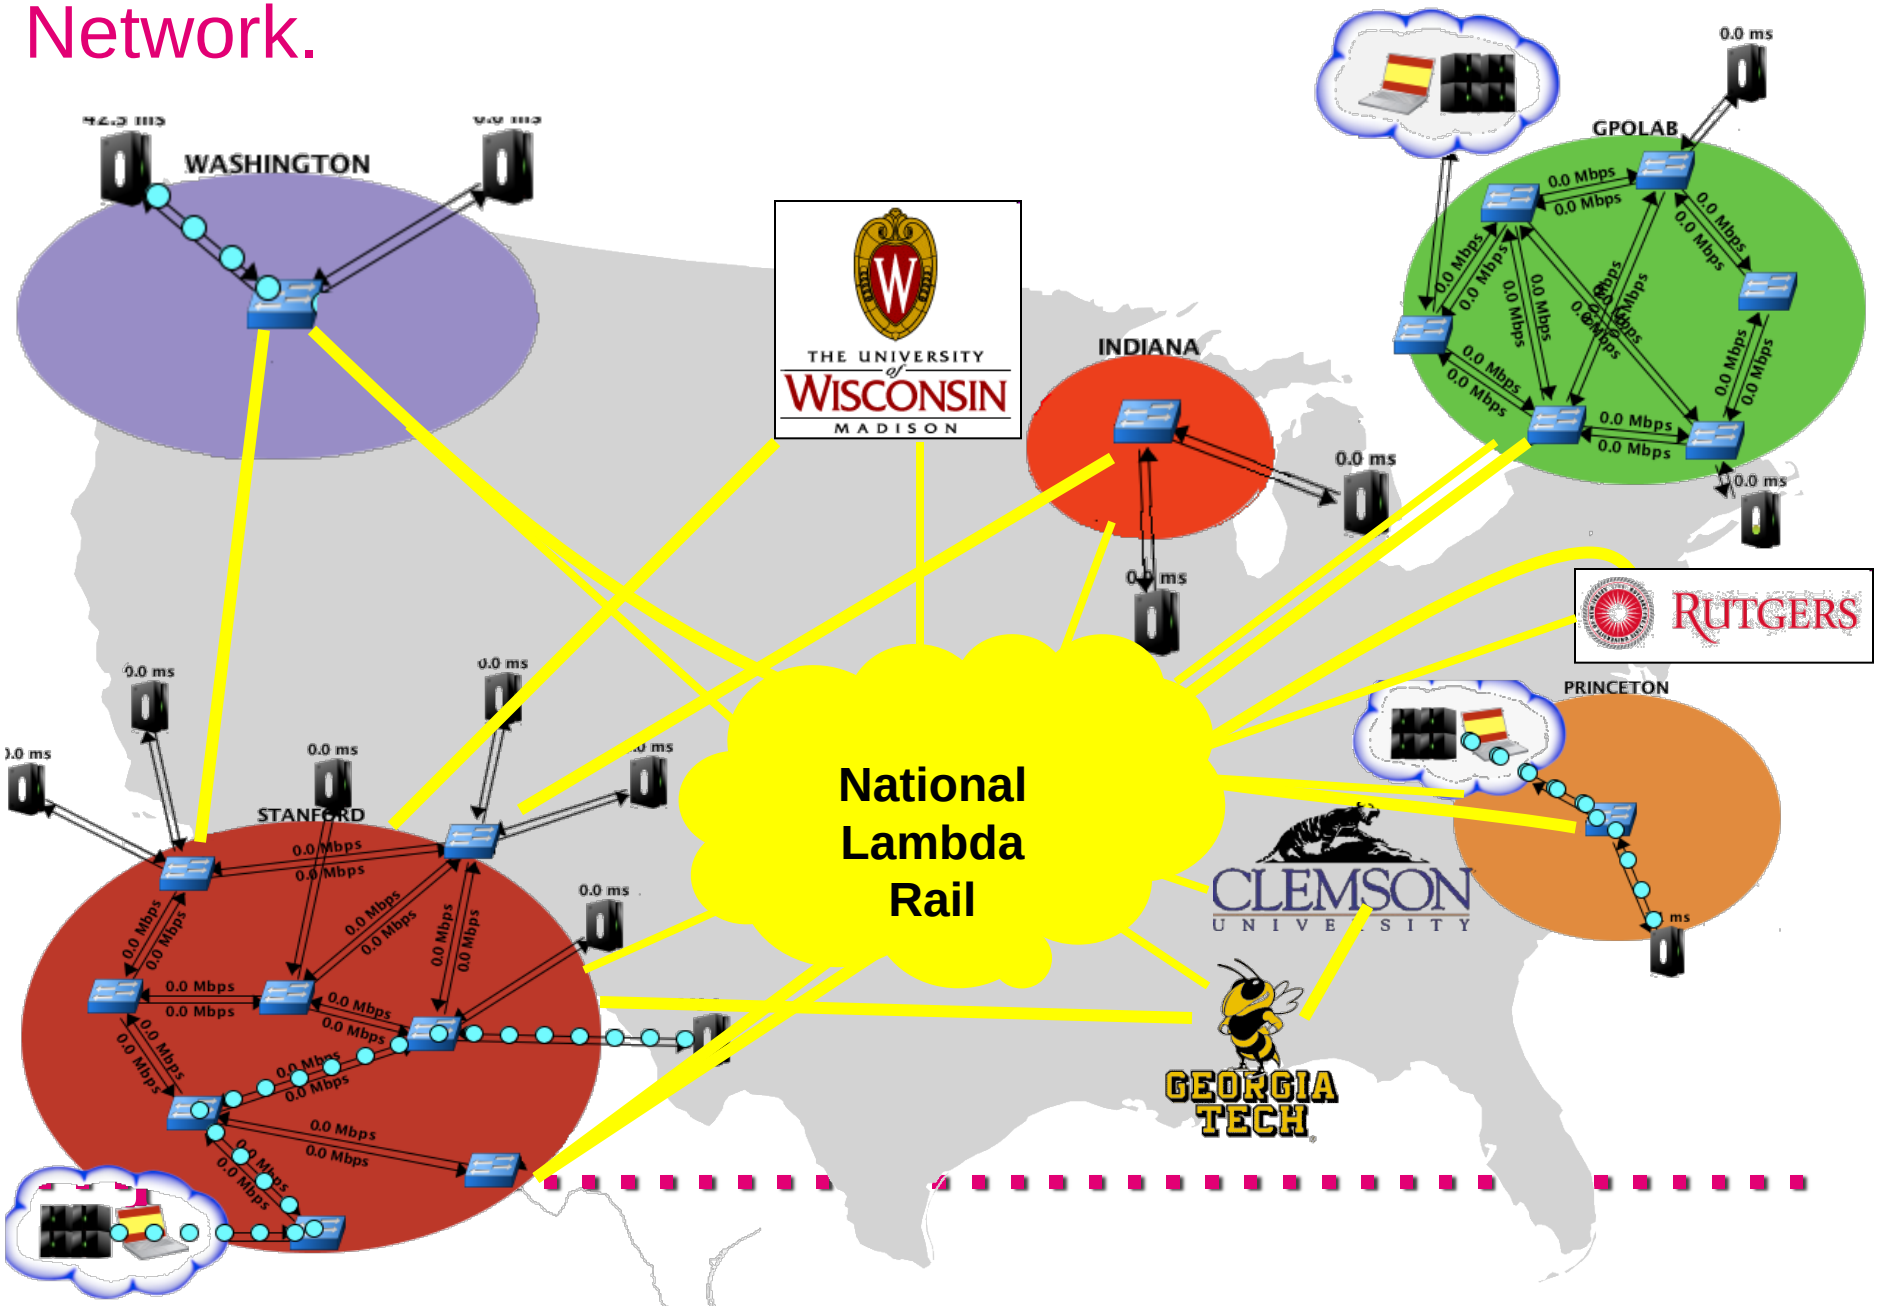

GENI. Integration.

- FlowVisor
  - Slicing control
- Expedient
  - Experimenter's portal for slice management
- Opt-in Manager
  - Network admins' portal to approve/ deny expt requests for traffic

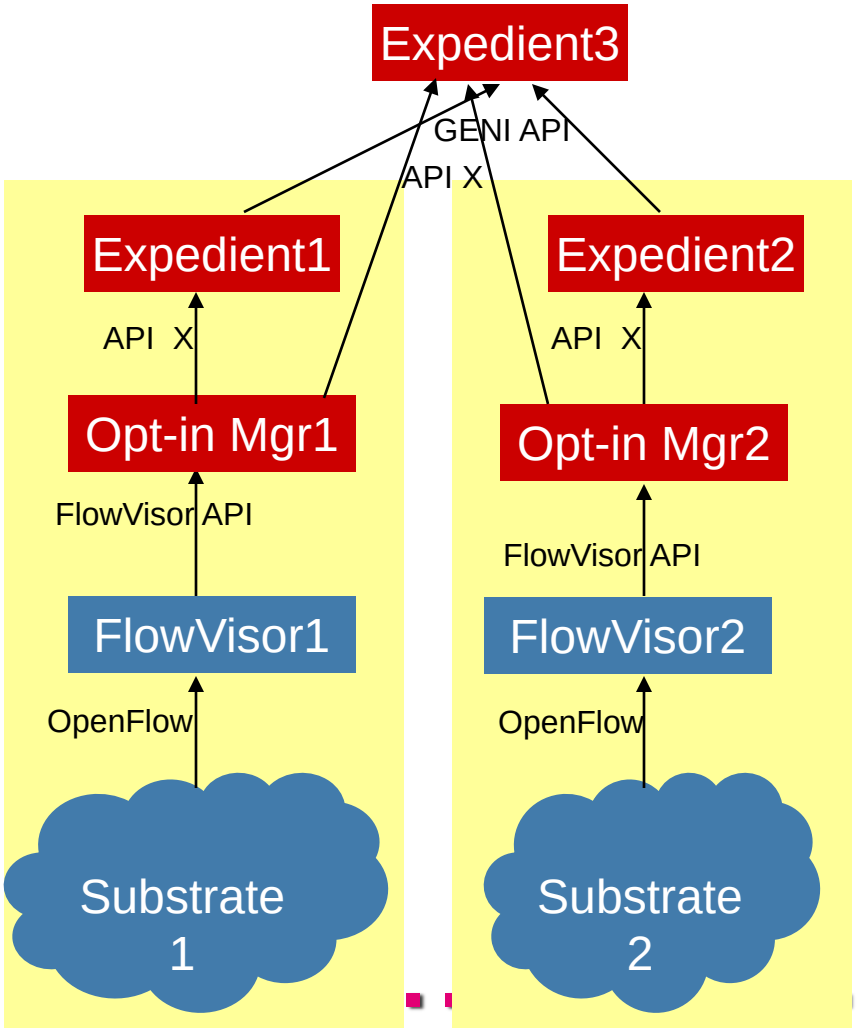

GENI.
Problems and experiences.

- 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

GENI.
Internet2: 35+ 100G POPs, nationwide.

GENI.
Next steps

- GENI
  - Building monitoring infrastructure for computing and networking resources
  - Expanding to regional networks
  - GENI racks: Scaling up computing

Discussion.
Done (almost).

- OpenFlow moving to commercial world 2011
- Better and more support forthcoming

- Proof of concept:
  - Distributed, US-wide substrate
  - Production networks at Stanford

- Lessons:
  - New network management/control paradigm
  - Tons of issues not foreseen

Discussion.

Questions?