Enterprise GENI

Status and Outlook - June 2009

Guido Appenzeller
Stanford University

AGENDA

- Status Spiral 1
  - Progress Review
- Outlook Spiral 2
  - Objectives
  - Observations

AGENDA

- Status Spiral 1
  - Progress Review
- Outlook Spiral 2
  - Objectives
  - Observations

MILESTONES

Milestone	Status
• Implement “Verison 0” Aggregate Manager	**Done**
• Initial small deployment in CS Building	**Done**
• Define Protocol integration with Clearing House	**Done**
• Start integration with Clearing House	**Done**
• Integrate with Switched VLAN infrastructure in I2	**Target: GEC6**
• Hold OpenFlow Workshop for interested Universities	**Scheduled 8/27/09**
• Demonstrate use of OpenFlow Switches with Clearing House	**Target: GEC6**

AGGREGATE MANAGER

Basic functionality implemented, work on opt-in and policy ongoing

Functionality Today:

- Speaks lightweight protocol as defined in Denver meeting
- Automatically discovers and reports switches and network topology and reports to Clearing House via RSpec
- Can virtualize Stanford OpenFlow network based on reservation RSpec received from Clearing House

Work In Progress

- Opt-In Mechanism for local users
- Connection to other networks
- Policy management and user feedback

STANFORD OPENFLOW DEPLOYMENT

Phase 1 on schedule, small delay expected for Phase 2

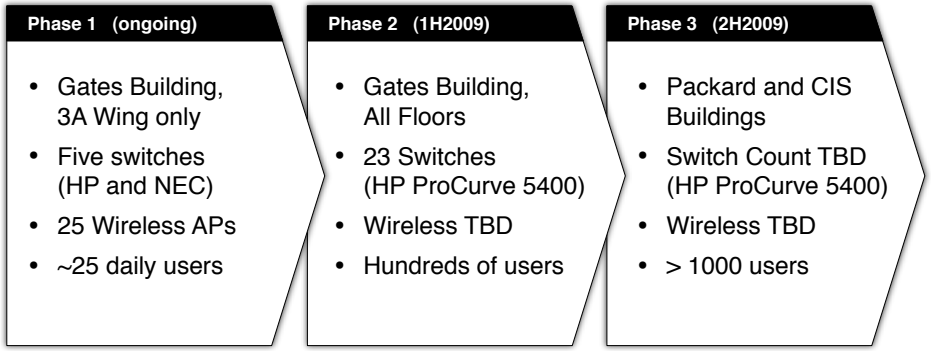

Stanford Enterprise GENI

- Phase 1 production deployment can be virtualized today
- Phase 2 expected to be ready for GEC6 demo

STANFORD OPENFLOW DEPLOYMENT TODAY

5 Switches, 25 users

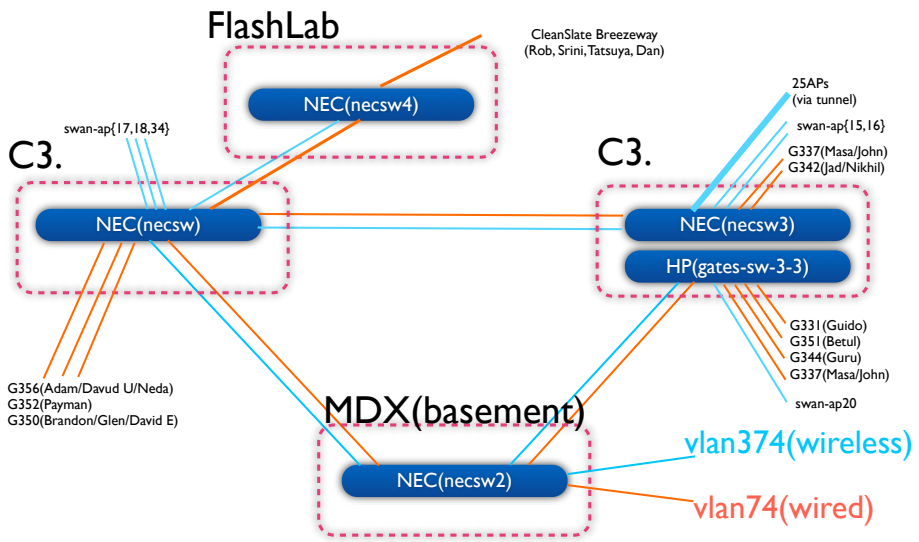

INTEGRATION WITH CLEARING HOUSE

"geniLight" protocol as defined in Denver

Simplified GENI Protocol

- All communication via CH
- Simplies Security Model
- No more PKI required (other than SSL certificates)
- Implemented in Enterprise GENI Aggregate Manager
- Used by "Toy Clearinghouse" to set up experiments on aggregate

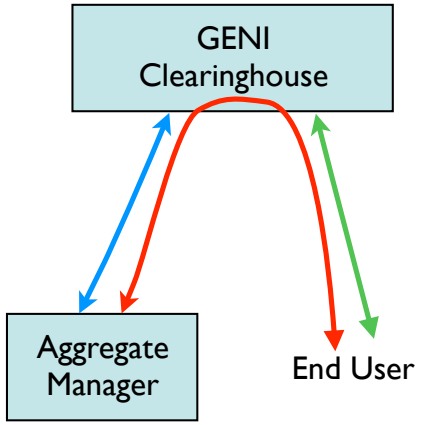

INTEGRATION AND DEMO

Planed until end of August, Target is Demo at GEC6

Integration

- Integrate with VLAN infrastructure at Internet 2 POP
- If not available, use tunnels over IP
- Starting this process now, completion by 8/31 seems low risk

Demo at GEC6

- Using at least two types of aggregates
- Possibly using OpenFlow aggregates at multiple universities
- Likely will require use of more than one control framework
- Exact Scope TBD

AGENDA

- Status Spiral 1
  - Progress Review
- Outlook Spiral 2
  - Objectives
  - Observations

MAIN FOCUS: OPENFLOW CAMPUS TRIALS

Expand OpenFlow Substrate to 7-8 other Campuses

- Multiple Vendors have comitted to participate including Cisco, Juniper, HP, NEC, Arista, Toroki
- Goal is to virtualize production infrastructure and allow experiements on this infrastructure via GENI
- Wired and Wireless
- Goal is a total of > 100 switches, 5000 ports

Connect via OpenFlow Internet2 Backbone

- Currently NetFPGA and Juniper, final platforms are TBD

GOAL FOR OPENFLOW SUBSTRATE IN SPIRAL 2

7-8 Deployments plus Backbone

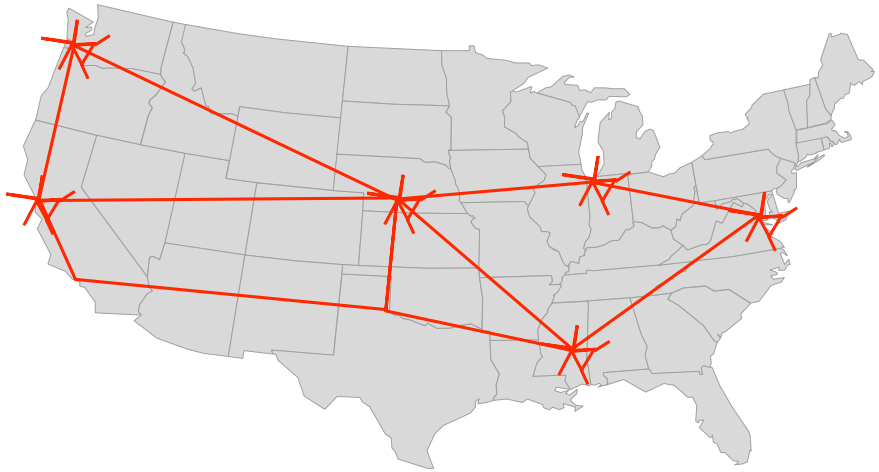

AGENDA

- Status Spiral 1
  - Progress Review
- Outlook Spiral 2
  - Objectives
  - Observations

CLEARING HOUSE INTEGRATION

Substrates drive Clearing House Functionality

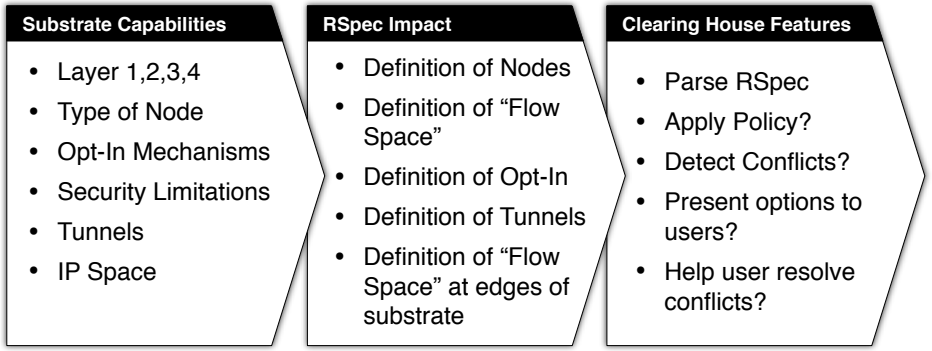

- How can Clearing Houses manage this complexity?
- How can Clearing Houses keep up to date with the rapid evolution of substrates?

OBSERVATIONS

Substrates drive Clearing House Requirements

- Substrates are still evolving (at least our is)
- At this point, we couldn't define a stable RSpec, as we don't know what we need yet
- At this point, it would be impractical to maintain a clearing house that can leverage our substrate

Possible Solutions

- Short Term: Use individual clearinghouses
- Long Term: How do we manage heterogeneity in substrates?
  - Modular Clearinghouse with plug-ins for substrates?
  - Distributed Clearinghouse?

OBSERVATIONS

What is the scope of the clearing house?

- Collects information from Aggregates
- Help user with reserving slices
- Manage accounts?
- Enforce Policy?
- Audit Trail?
- Command Line?
- HTML GUI?
- Visualize Topology?