OMIS WG

System Engineering Report

Michael Patton
MAP@GENI.net
Working Group System Engineer

groups.geni.net
GENI working group wiki

Outline

- Introduction to the talk
- Addresses and routing in general
- GENI context
- Sample projects
  - Intro
  - 1599: DOME
  - 1578: Internet Scale Overlay Hosting
  - 1613: Enterprise GENI
- Conclusion

Motivation

Integration: the I in OMIS

- How projects connect to one another
  - Control Plane
  - Data Plane (Layer 2, sometimes Layer 3)
- How the CF clusters talk
  - to each other
  - to the outside
- Concentration on addressing and routing
  - Focus in IPv4 (IPv6 similar, but not scarce)

Milestones for integration ***this summer***

Connectivity Classes

What are we talking about

- Three classes of communication
  - Project Internal
    - Little impact outside project
    - L2 and/or L3
  - GENI Control plane
    - External to project, Internal to GENI
  - Experiment data plane
    - Layer 2 (goal)
    - Layer 3

Project analysis will use this same organization

Terminology

- "Clearinghouse" may include services, too
- Careful to distinguish
  - Project
    - Project internal control

      Used only to refer to a GENI funded D&P Project
  - GENI control plane

    Not the only control plane, but the main one in this talk
  - Experiment
    - Setup, loading, monitoring (GENI Control Plane)
    - Experiment Internal
      - May or not be IP
    - Experiment External
      - Traffic between an experiment **in** GENI and the global Internet

“I need a GENI address”

In the GENI context addresses are:

- Communication with Clearinghouse
- Control plane access to CM and/or AM
  - Component controlled internally by AM would not need one (for GENI, maybe for AM to Component)
- Tunnel end points over legacy Internet
  - For components without direct L2 connection
  - Opt-In users
- Experiments running legacy IP
  - L2 experiments do not need L3 addresses
- Operations and Management

What you need it for determines where you get it.

Where do you get an address?

- Primarily depends on routing
  - An address that's not advertised in routing isn't useful
- Today's Internet routes on address prefix
  - For two nodes to communicate:
    - Addresses at both ends in valid prefix
    - Those prefixes must both be advertised from origin
    - Routing info must propagate both ways to other end

Where do you get an address?

That means:

- Equipment connected to a campus IP network needs an address from the campus prefix
- Equipment connected directly to an ISP needs an address from the ISP prefix
- Equipment connected directly to a regional network needs an address from the regional network's prefix

In general:

- Equipment connected via *whatever* needs an address from *whatever*'s prefix

GENI addresses?

Addresses in a GENI prefix are useful to:

- Equipment getting IP connectivity from “GENI backbone”
  - Doesn't exist, yet, I2 and NLR in interim
  - Care needed to avoid circular dependency
- Experiments inside GENI that need external connection (e.g. for Opt-In users)

or

- Projects that are working on providing connectivity to legacy Internet
  - DTunnels, Regional Opt-In

Example projects

- Chosen to show various “interesting” concerns
- Emphasis for pedagogical purposes
  - Low probability; contingency plans
- These are not complete project descriptions
  - Things irrelevant to external connections left out
  - Sometimes only one of several options described
- I don't speak for the projects
  - I may have some things wrong, let me know
- Any criticisms not project specific
  - Often these projects are the least problematical
  - May be just to make a point

The projects

- 1599: DOME
- 1578: Internet Scale Overlay Hosting
- 1613: Enterprise GENI

I picked these as described.

Anyone want to be another example?

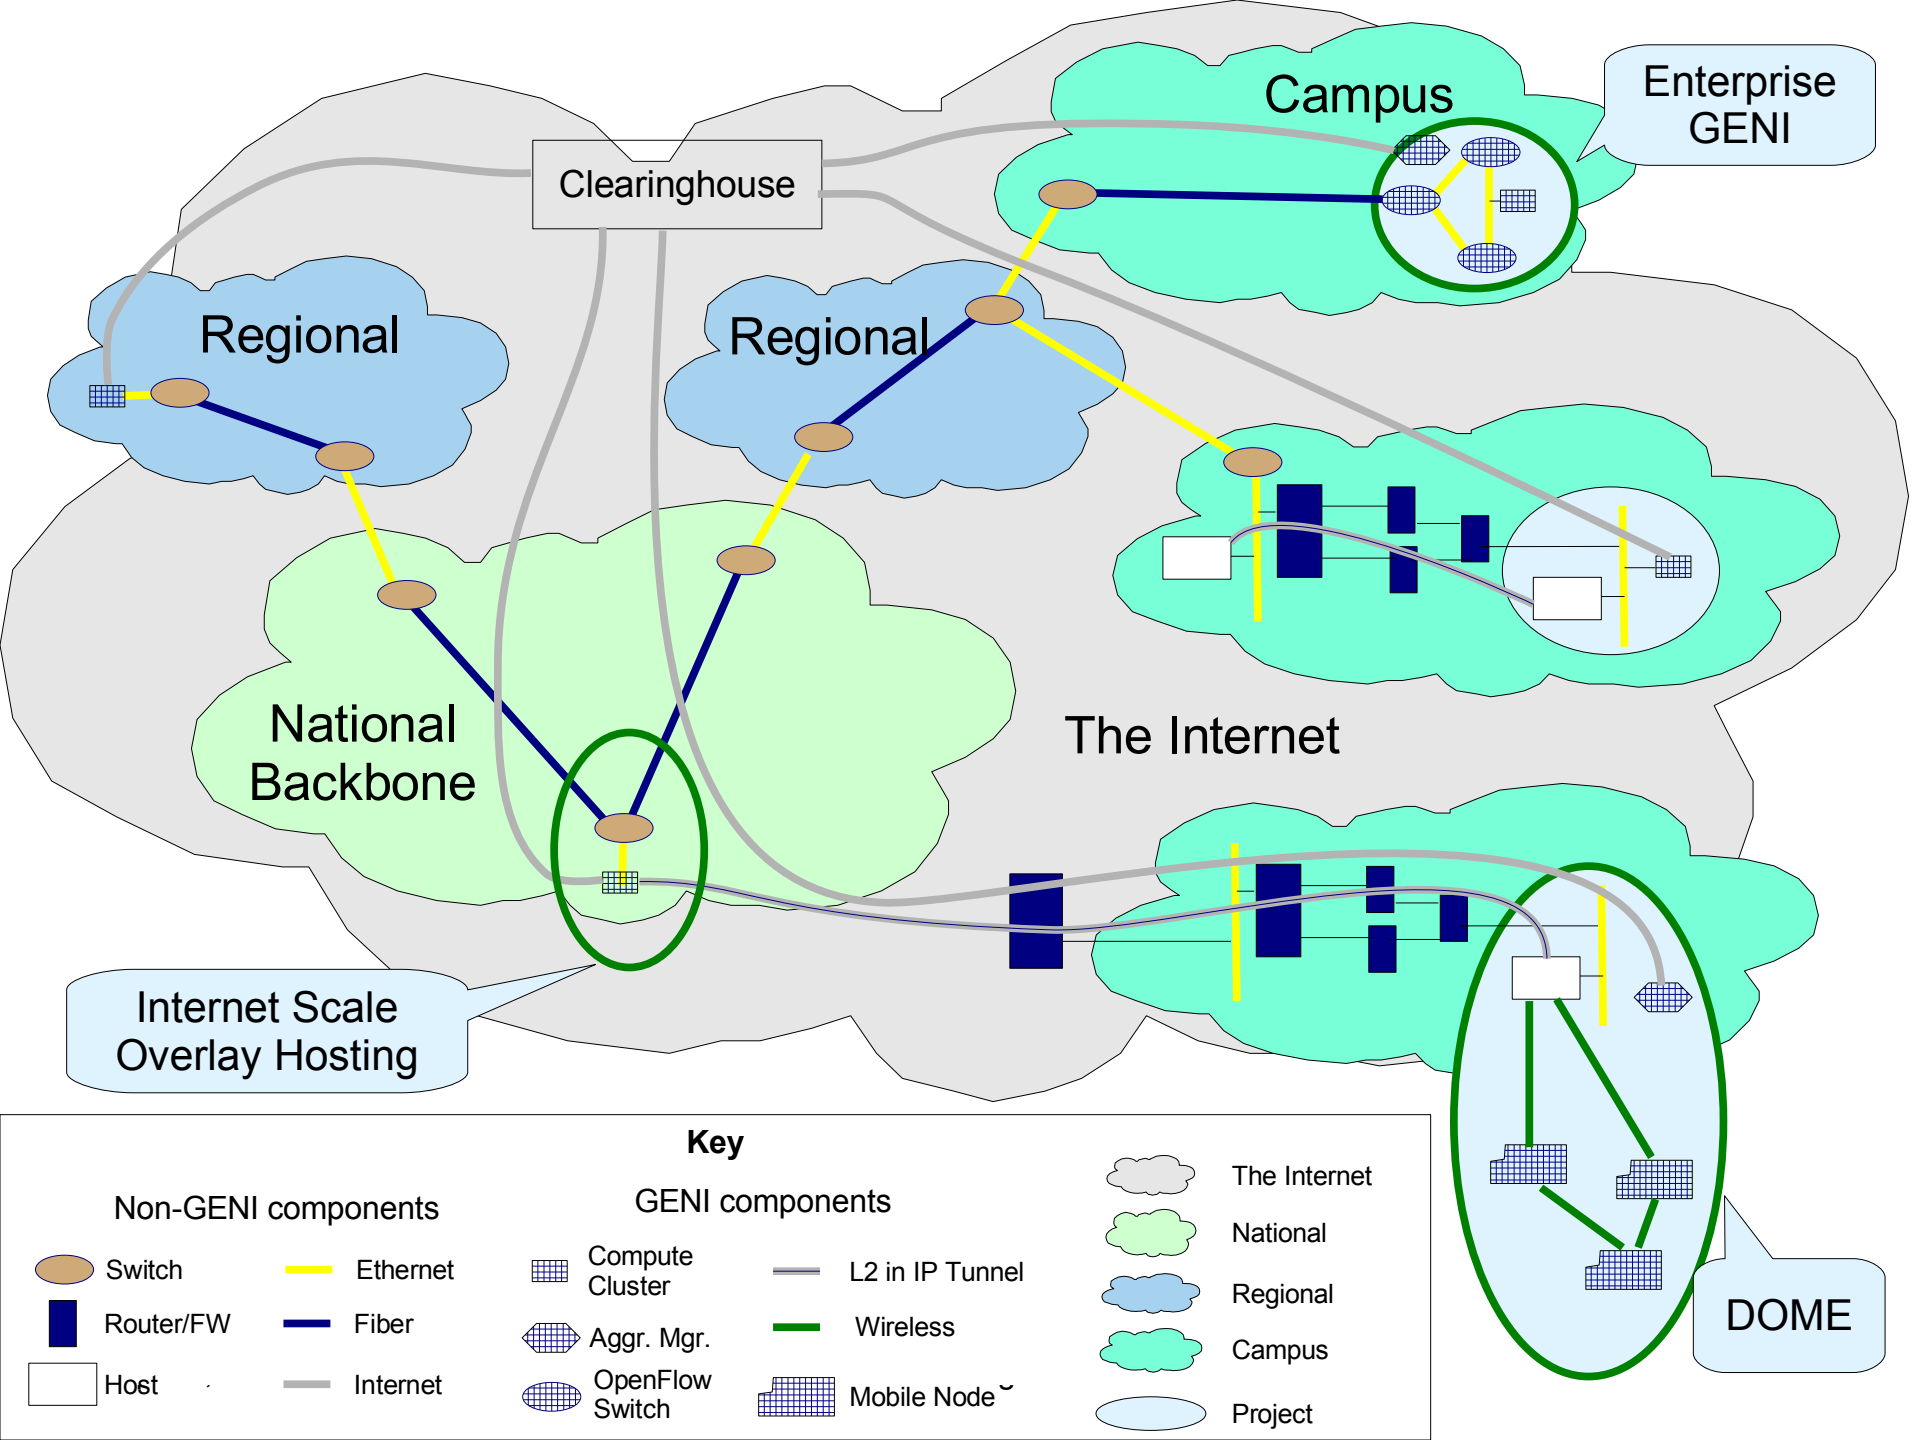
Campus
Enterprise GENI
Clearinghouse
Regional
Regional
National Backbone
The Internet
Internet Scale Overlay Hosting
DOME
Key
Non-GENI components
Switch
Router/FW
Host
Ethernet
Fiber
Internet
GENI components
Compute Cluster
Aggr. Mgr.
OpenFlow Switch
L2 in IP Tunnel
Wireless
Mobile Node
The Internet
National
Regional
Campus
Project

1599: DOME

- Intermittently connected mobile nodes
  - Layer 2 available in a few locations
    - Bus to bus only in first phases
  - Only commercial IP access in most locations
- Layer 3 connections
  - Campus Network for management and control
  - Wireless (Wi-Fi, 3G) for mobile nodes
- Layer 2 connections
  - Expecting L2 connectivity to rest of GENI
  - Most wireless is IP only, L2 via tunnels
  - Some wireless nodes might do L2, eventually

Original diagram

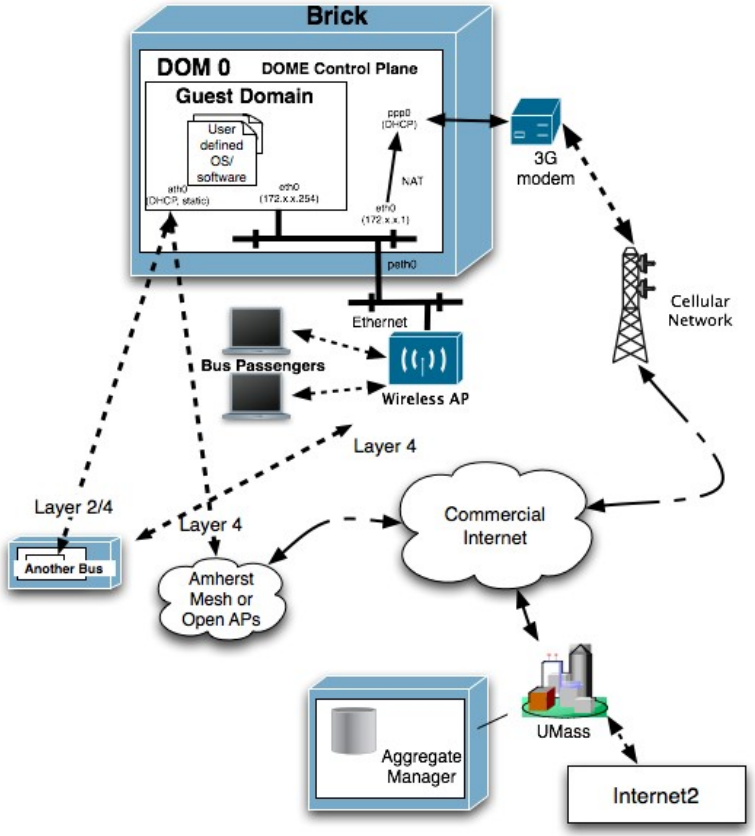

Updated diagram

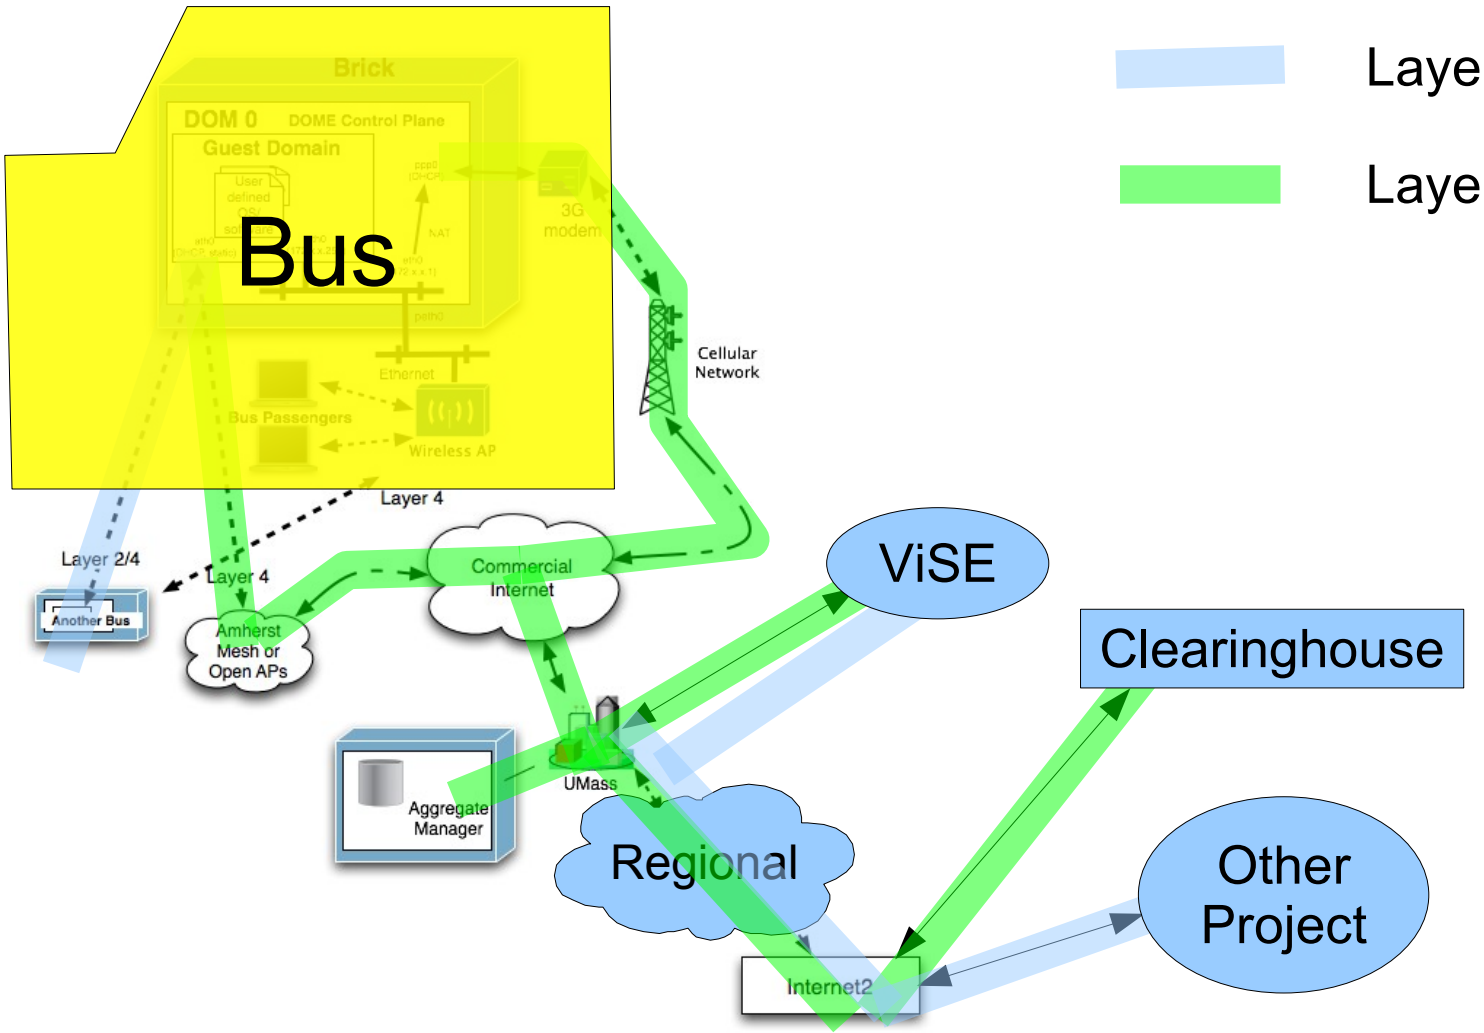

1599: DOME

- Project internal
  - 3G wireless IP for internal control
  - IP over open AP and/or commercial Internet
- GENI control plane
  - AM to control plane elements via campus network
- Experiment data plane
  - L3 native
  - L2 via regional to I2 or via tunnels (internally and externally)

1578: Internet Scale Overlay Hosting

- SPP nodes at multiple Internet2 locations
- Connects to Internet2 IP switches
  - Control Plane
  - IP-only Experiment Plane
  - Measurement and Ops, too
- Connects to GENI wave on Internet2
  - IP only at present
  - Layer 2 Ethernet, soon

Project Diagram

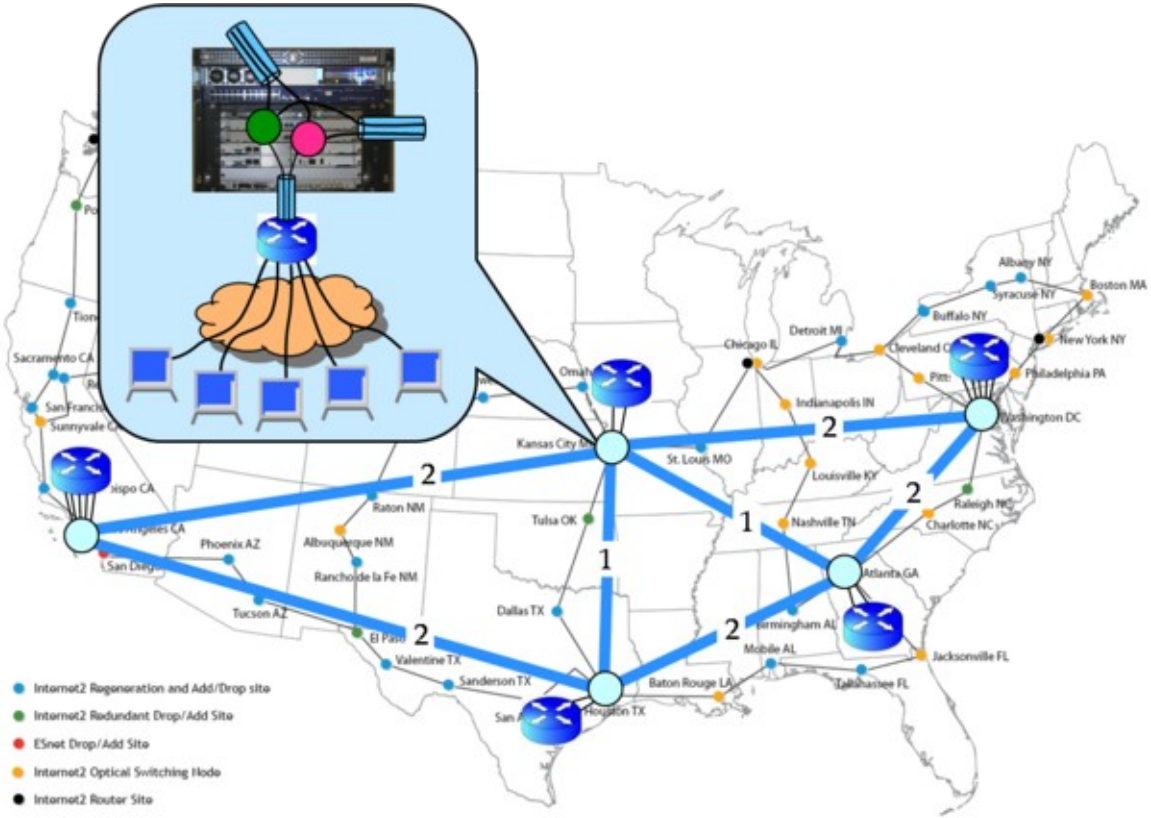

Diagram

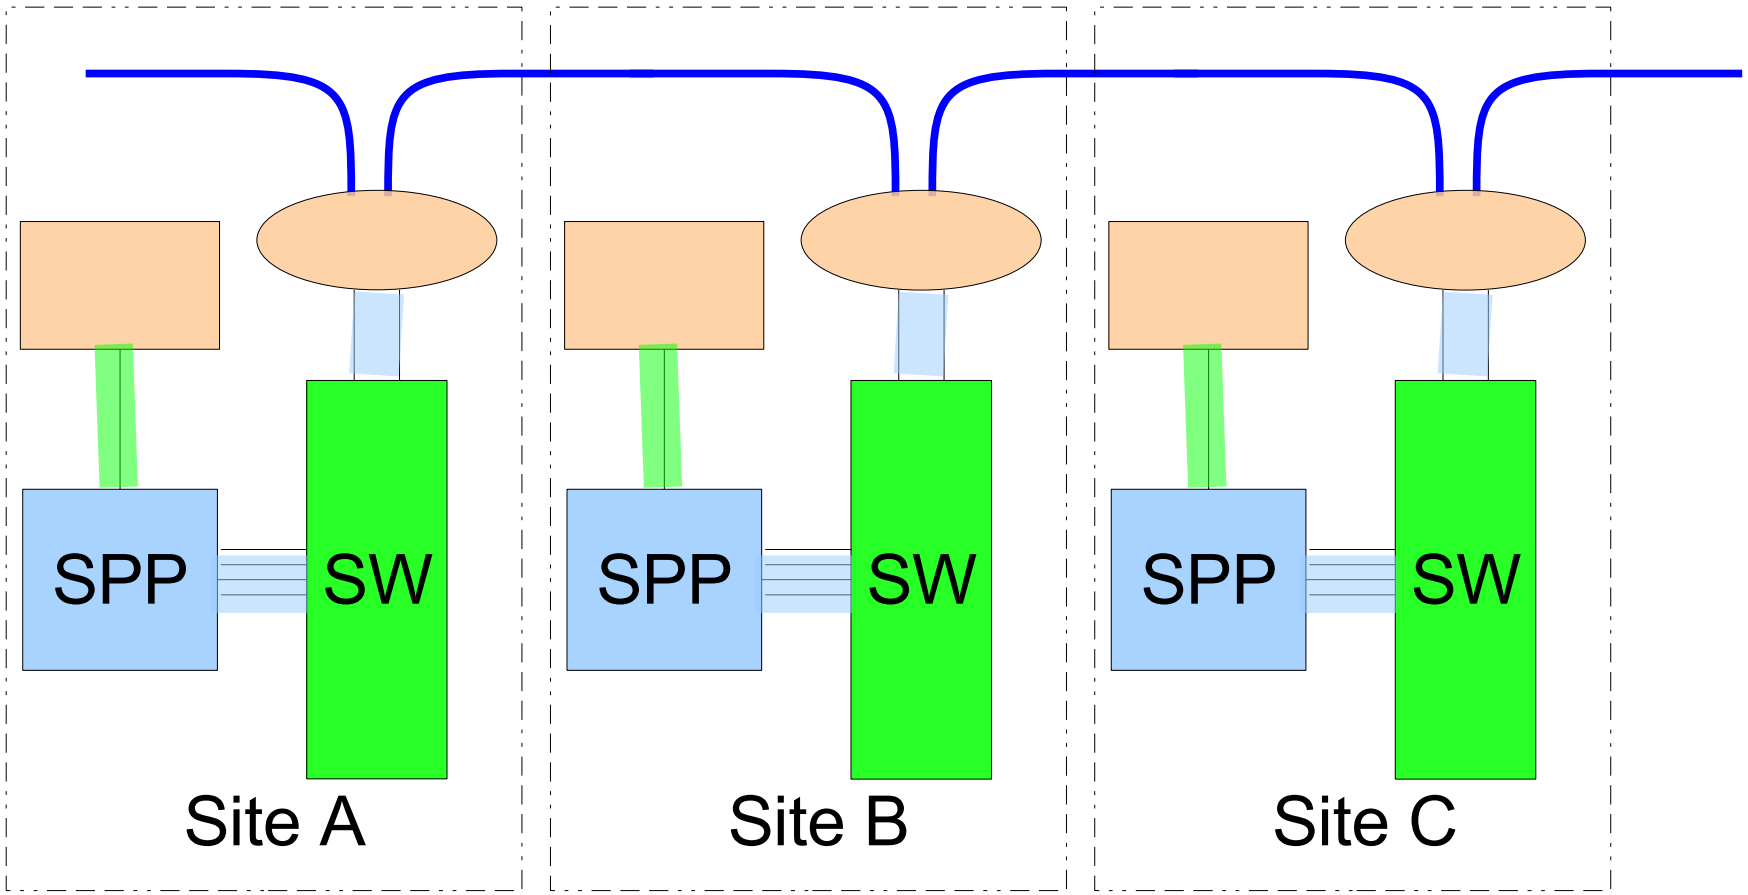

Addressing and Routing

- Project internal
  - Internal to SPP
- GENI control plane
  - One address for CM/AM control connections
  - Routed through I2 router, so I2 routed address
- Experiment data plane
  - IP initially, Layer2 soon
  - Engineered I2 connections for SPP to SPP
  - IP routing for outside (tunnels possible for Layer2)
  - Layer2 connections outside I2 envisioned

1613: Enterprise GENI

- OpenFlow Ethernet switches on campus
- OpenFlow Controller (AM for switches)

Diagram

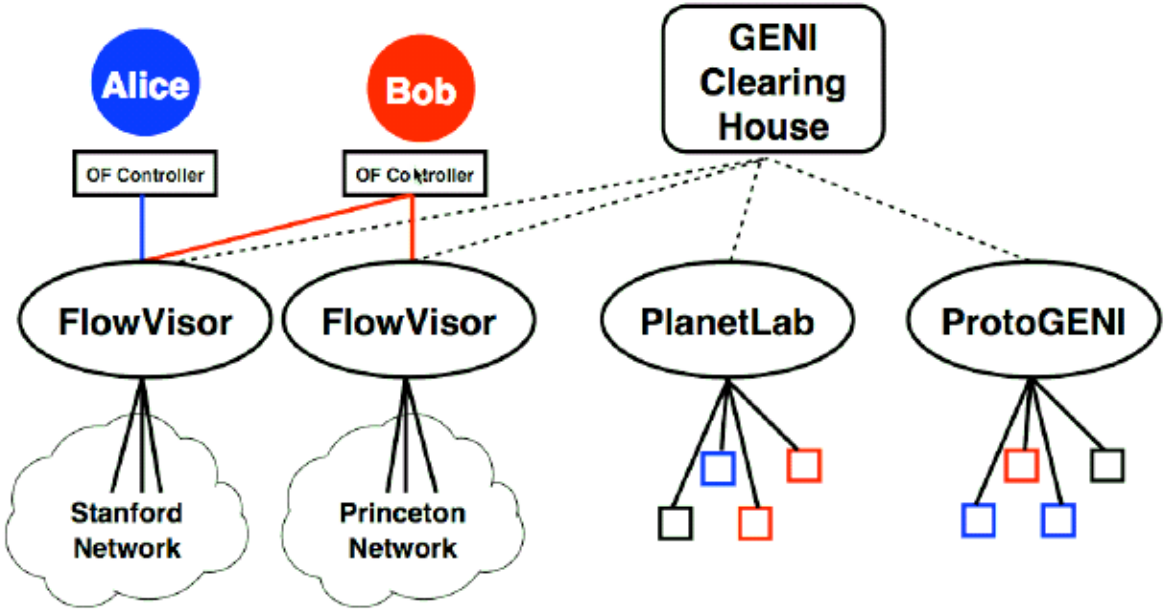
Alice
Bob
GENI Clearing House
OF Controller
OF Controller
FlowVisor
FlowVisor
PlanetLab
ProtoGENI
Stanford Network
Princeton Network

Addressing and Routing

- Project internal
  - Connections between controller and switches uses IP and requires addresses
  - Could be “private use”, but probably from campus
- GENI control plane
  - OpenFlow Controller (AM) needs routed address, just another connection on the campus network
- Experiment data plane
  - Layer 2 connection through regional to backbone

Wrap Up

- Most projects have a milestone for interoperation this summer
- Need to start working on connectivity (especially routing) early
- Integration can often hit snags
  - Start exploring early
  - Try linkage to all other projects in CF(s)
- Connectivity diagrams on wiki
- *Read about other projects: as I was looking at just these three, I noticed some useful cross fertilization between apparently unrelated projects. Use what others develop.*

Conclusion

- IP addresses come from connectivity
  - Use addresses in block routed to campus, ISP, regional, *whatever* where connected
  - Until GENI is its own ISP, GENI addresses won't work
    - Meanwhile use campus, I2, NLR, Regional, *whatever*
- Need addresses with suitable routing
  - Two projects working on external connectivity for GENI will start to provide this
- Sites without L2 connection can tunnel

Follow-up

- Notes, slides, actions, etc will be sent to the working group mail list and posted on the wiki page: http://groups.geni.net/geni/wiki/GeniOmis
- Any project that wants addressing and/or routing analysis can contact me directly or through your SE
- We are working on direct GENI addresses
  - Interim from I2, ISP, Regional, Campus, etc.
  - Long term block for GENI

    If you have good contacts in ARIN, we'd like to hear...