




GENI in the Classroom



# Introduction to **GENI** Network Architecture



[www.geni.net](http://www.geni.net)



Sponsored by the National Science Foundation

# GENI: Infrastructure for Experimentation



## Regional nets

-  Existing
-  New

## GENI WiMAX

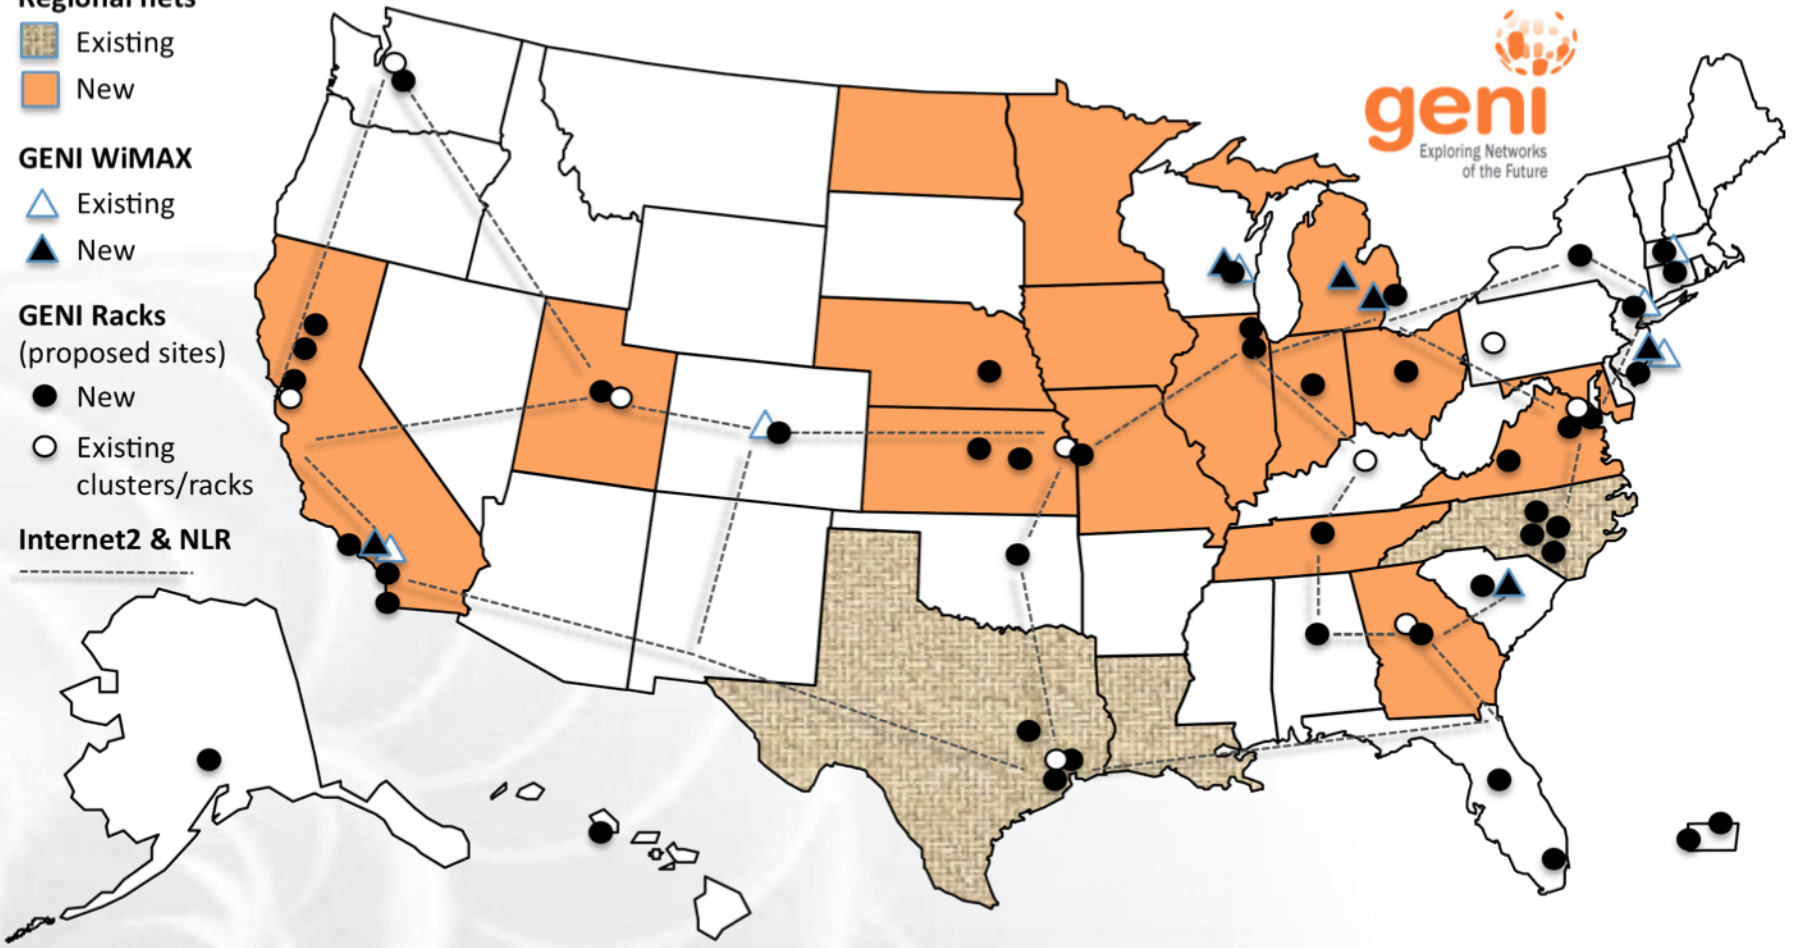
-  Existing
-  New

## GENI Racks (proposed sites)

-  New
-  Existing clusters/racks

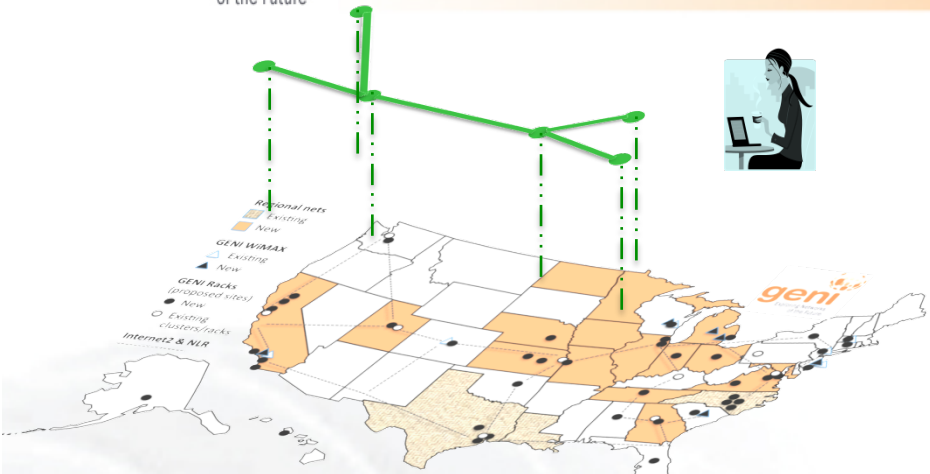
## Internet2 & NLR

-----

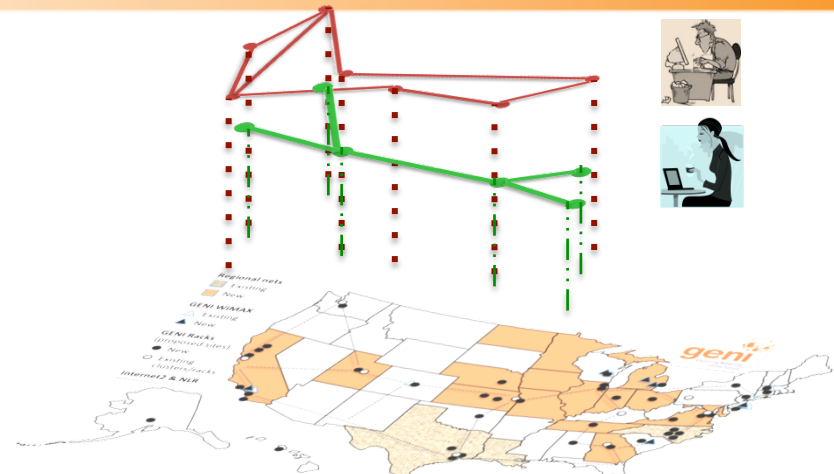


GENI provides **geographically distributed** compute resources that can be connected in **experimenter specified Layer 2 topologies**.


# GENI Key Concepts



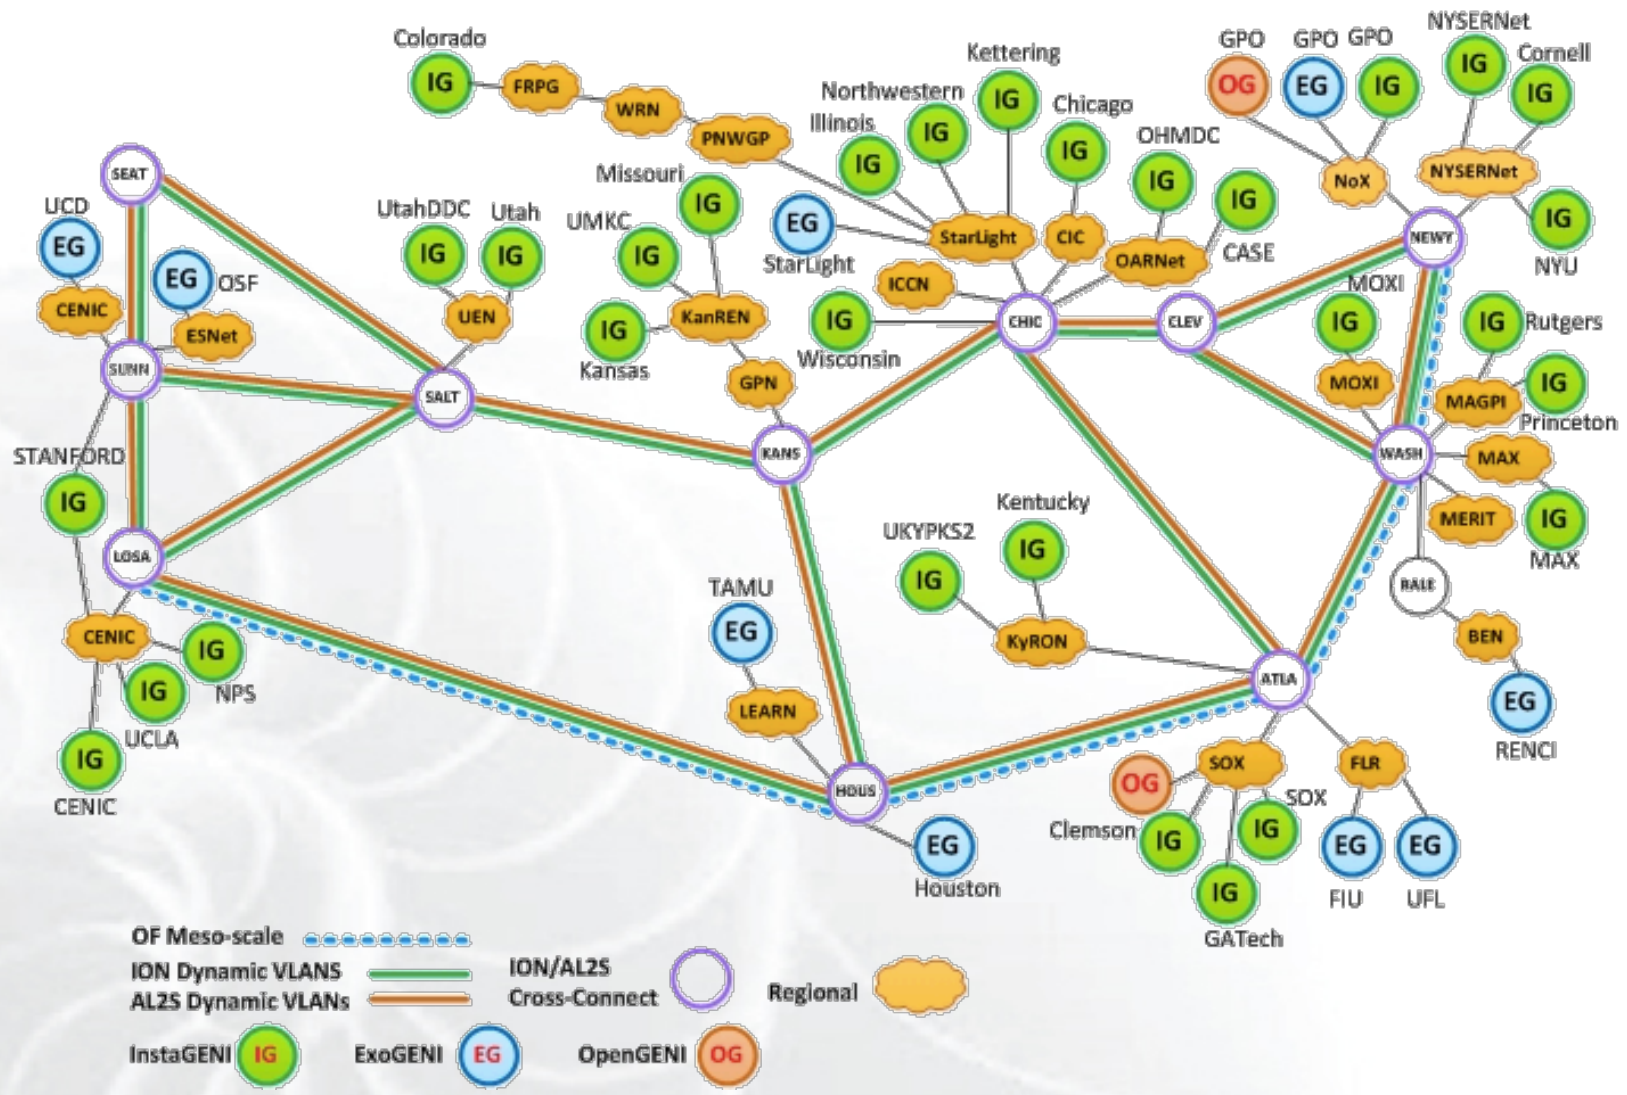
**GENI is deeply programmable** – experiments control forwarding



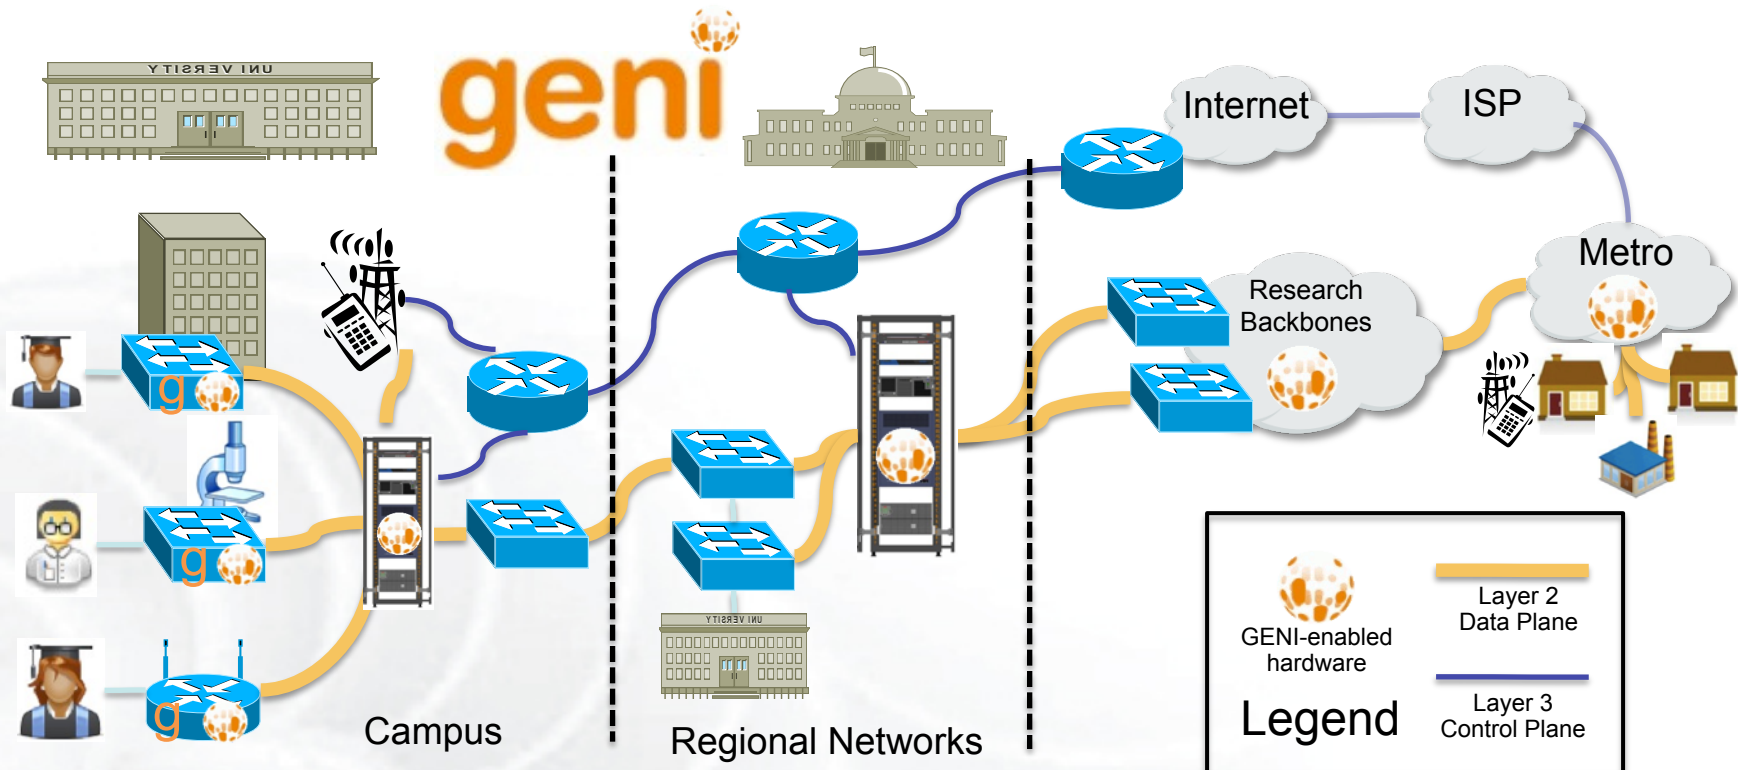
**GENI is sliceable** – supports concurrent experiments



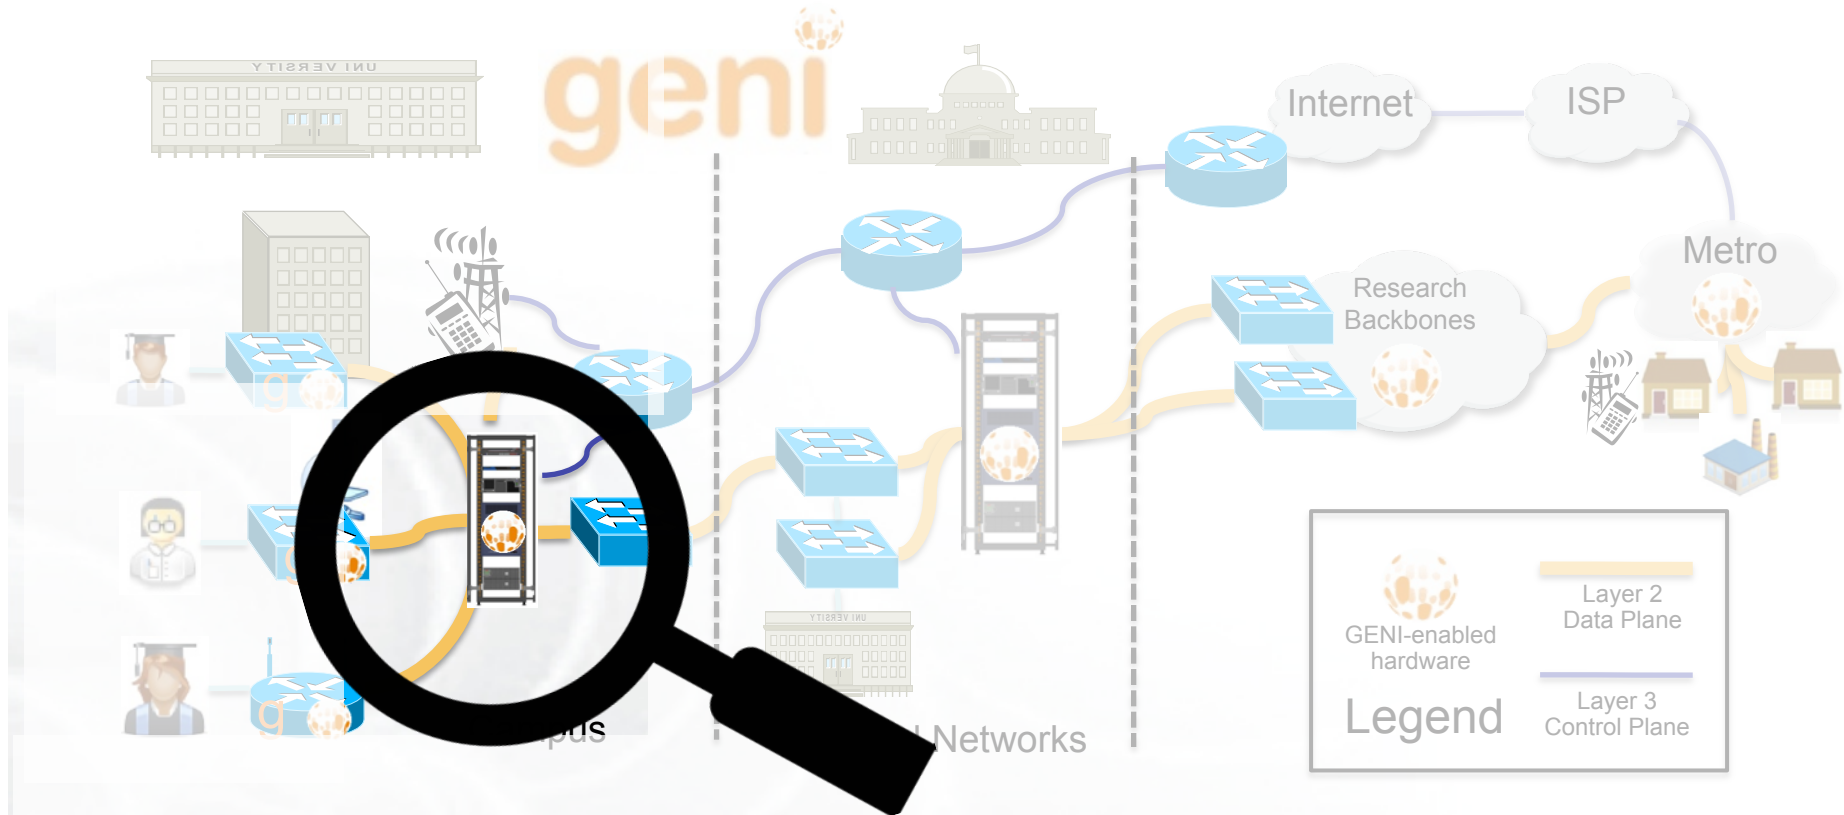
**GENI is a Federation** – to the user appears as one testbed



# GENI Network Architecture Overview

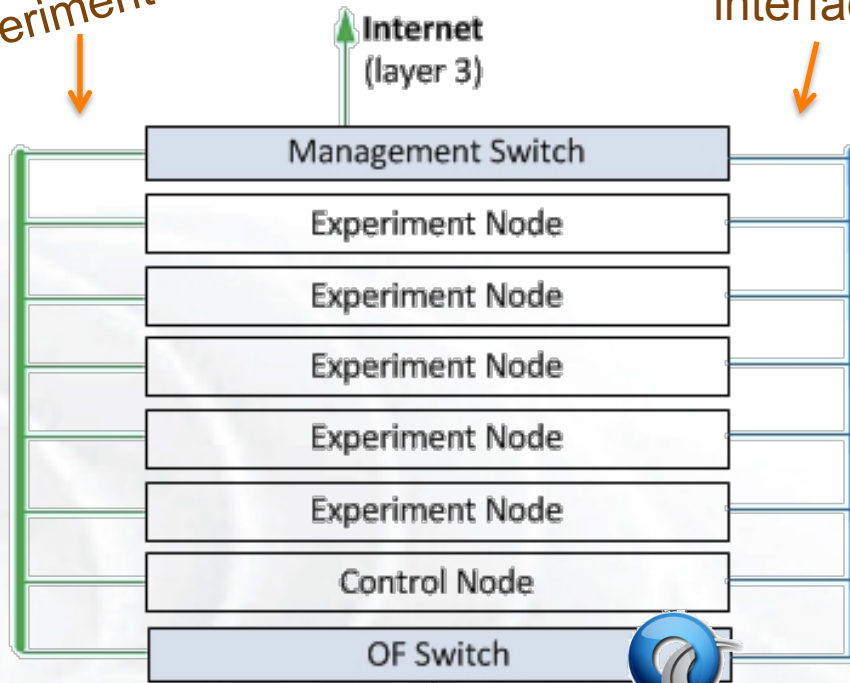


- Spans campus/metro, regional, and nationwide networks
- GENI relies on participants to contribute compute and network resources
- Includes VMs, bare metal nodes, SDN switches (OpenFlow 1.0), WiMAX/LTE base stations and clients, L2 network access




Access to experiment nodes

Multiple data interfaces



*sliceable SDN data switch can be controlled by the experimenter*



*GENI network is sliced by VLAN*

***Embed a variety of topologies within one rack sliced by VLAN***

# GENI Rack OpenFlow switch

## Different OpenFlow modes

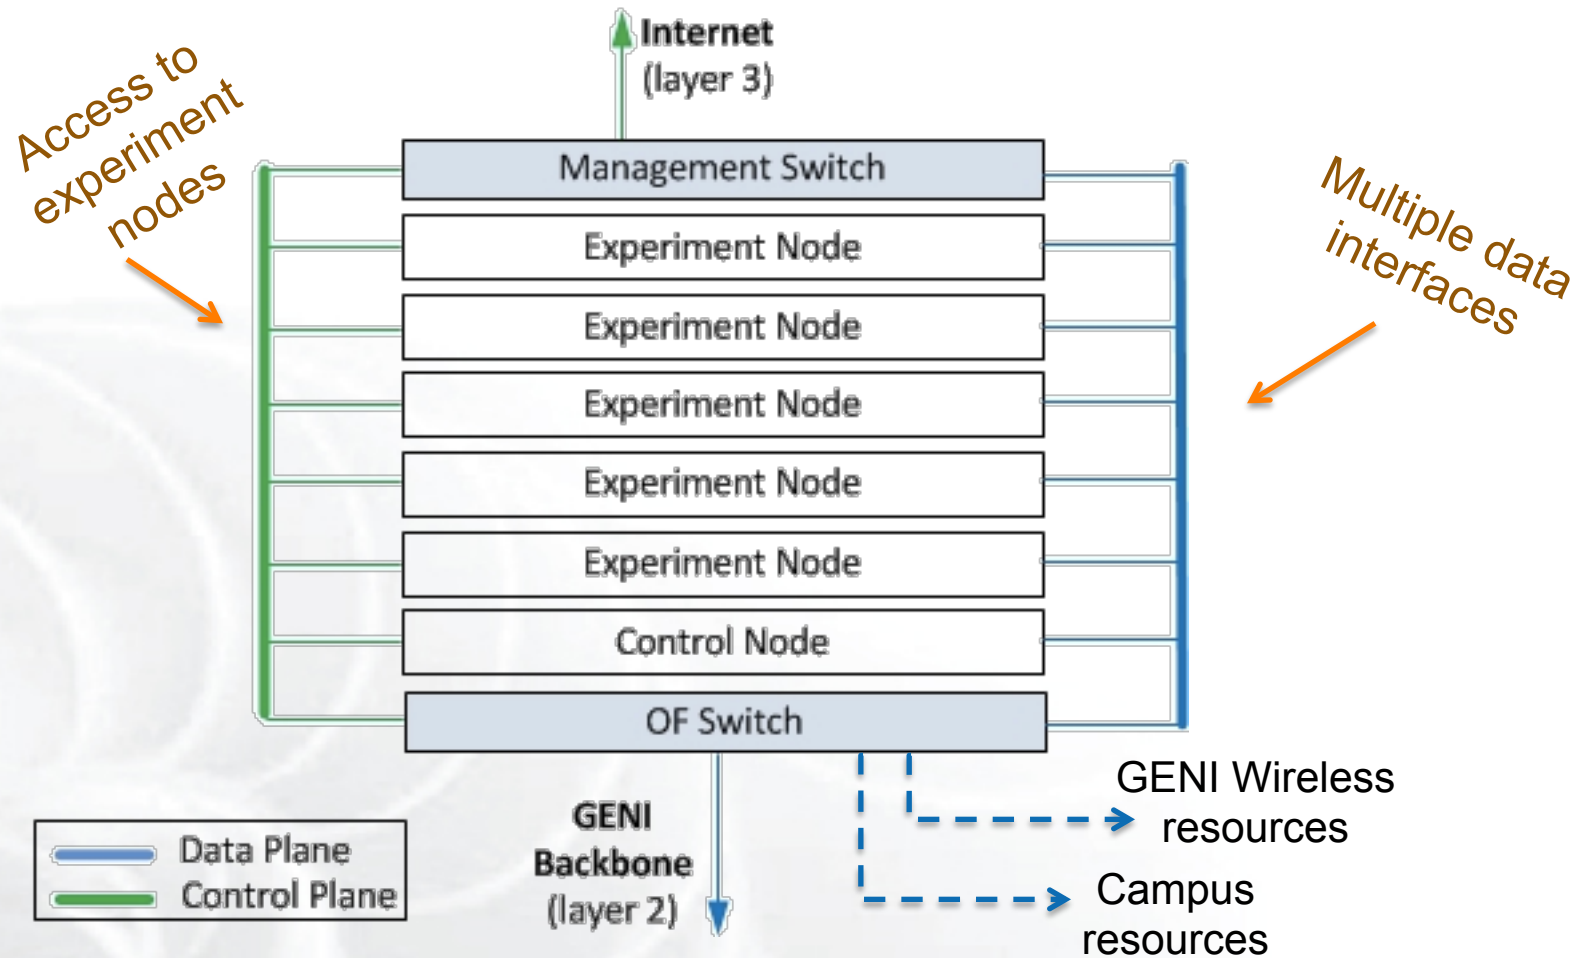
- switches in **pure OF** mode are acting as one datapath
- **Hybrid VLAN switches** are one datapath per VLAN (*sliced at the switch*)
- **Hybrid port switches** are two datapaths (one OF and one non-OF)
- Need slicer for **pure OF** and **port hybrid**
  - FlowVisor, FlowSpace Firewall



OF switch

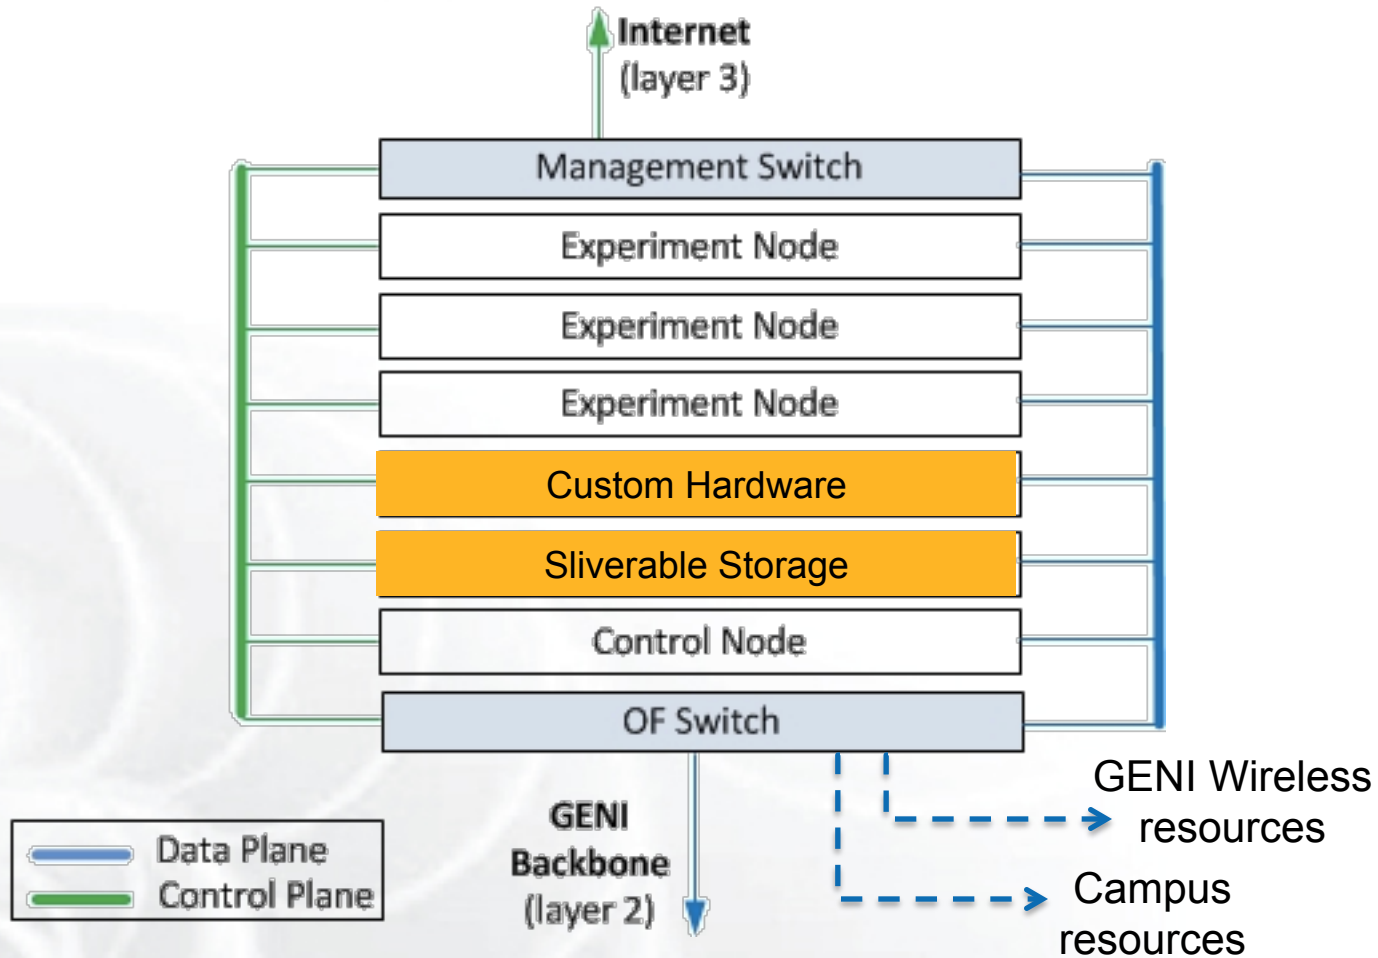
*Embed a variety of topologies within one rack **sliced by VLAN***



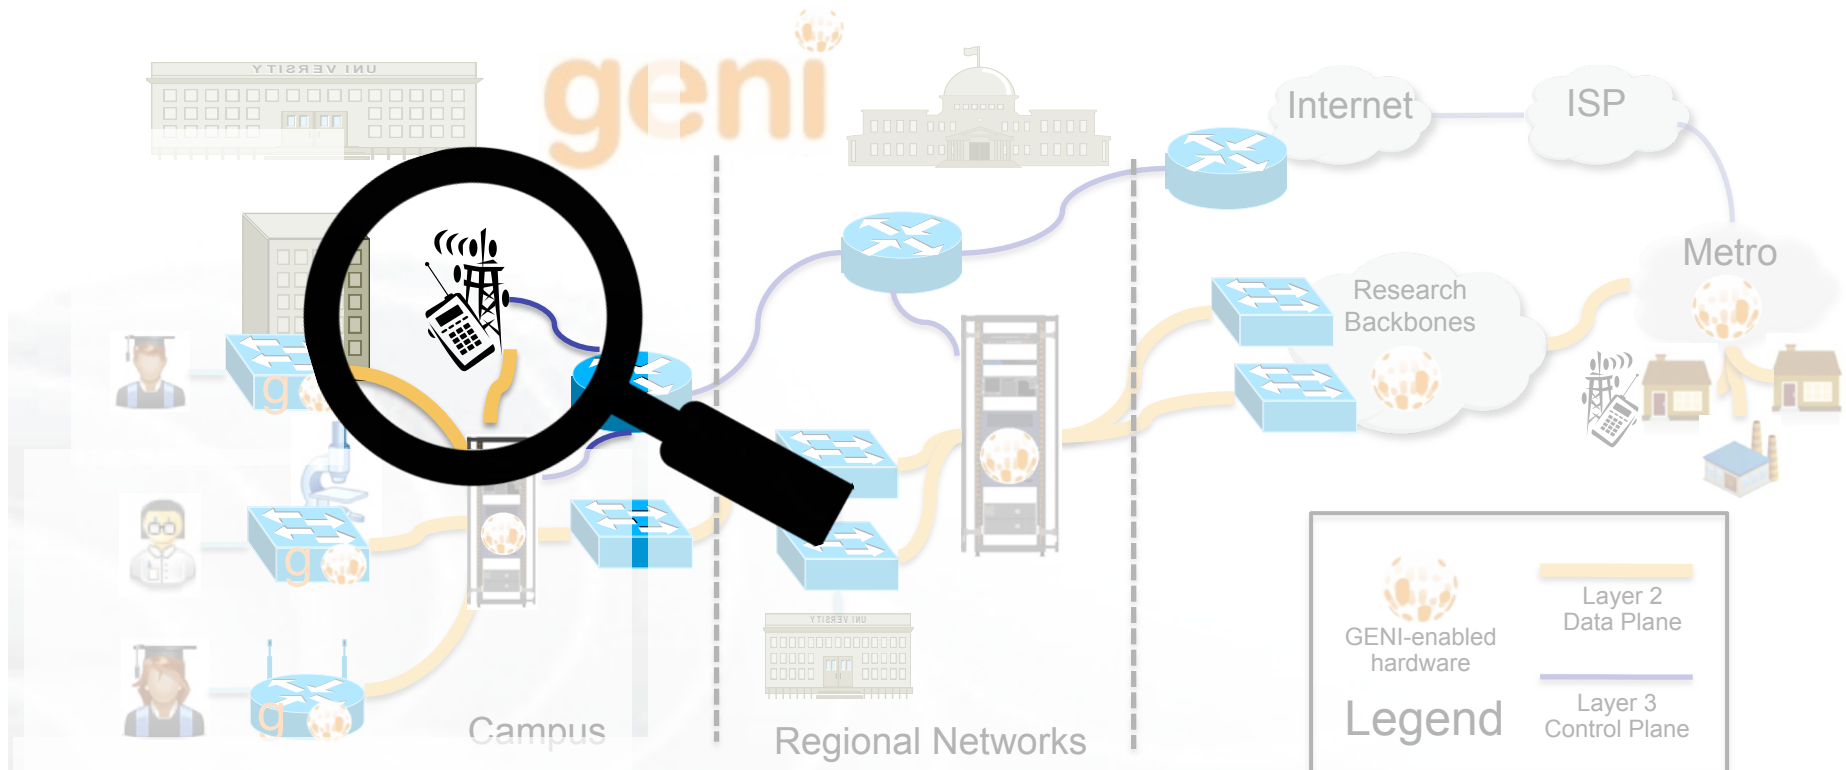


Each Rack is a separate AM that ensures that experiment nodes are connected to the appropriate external resources

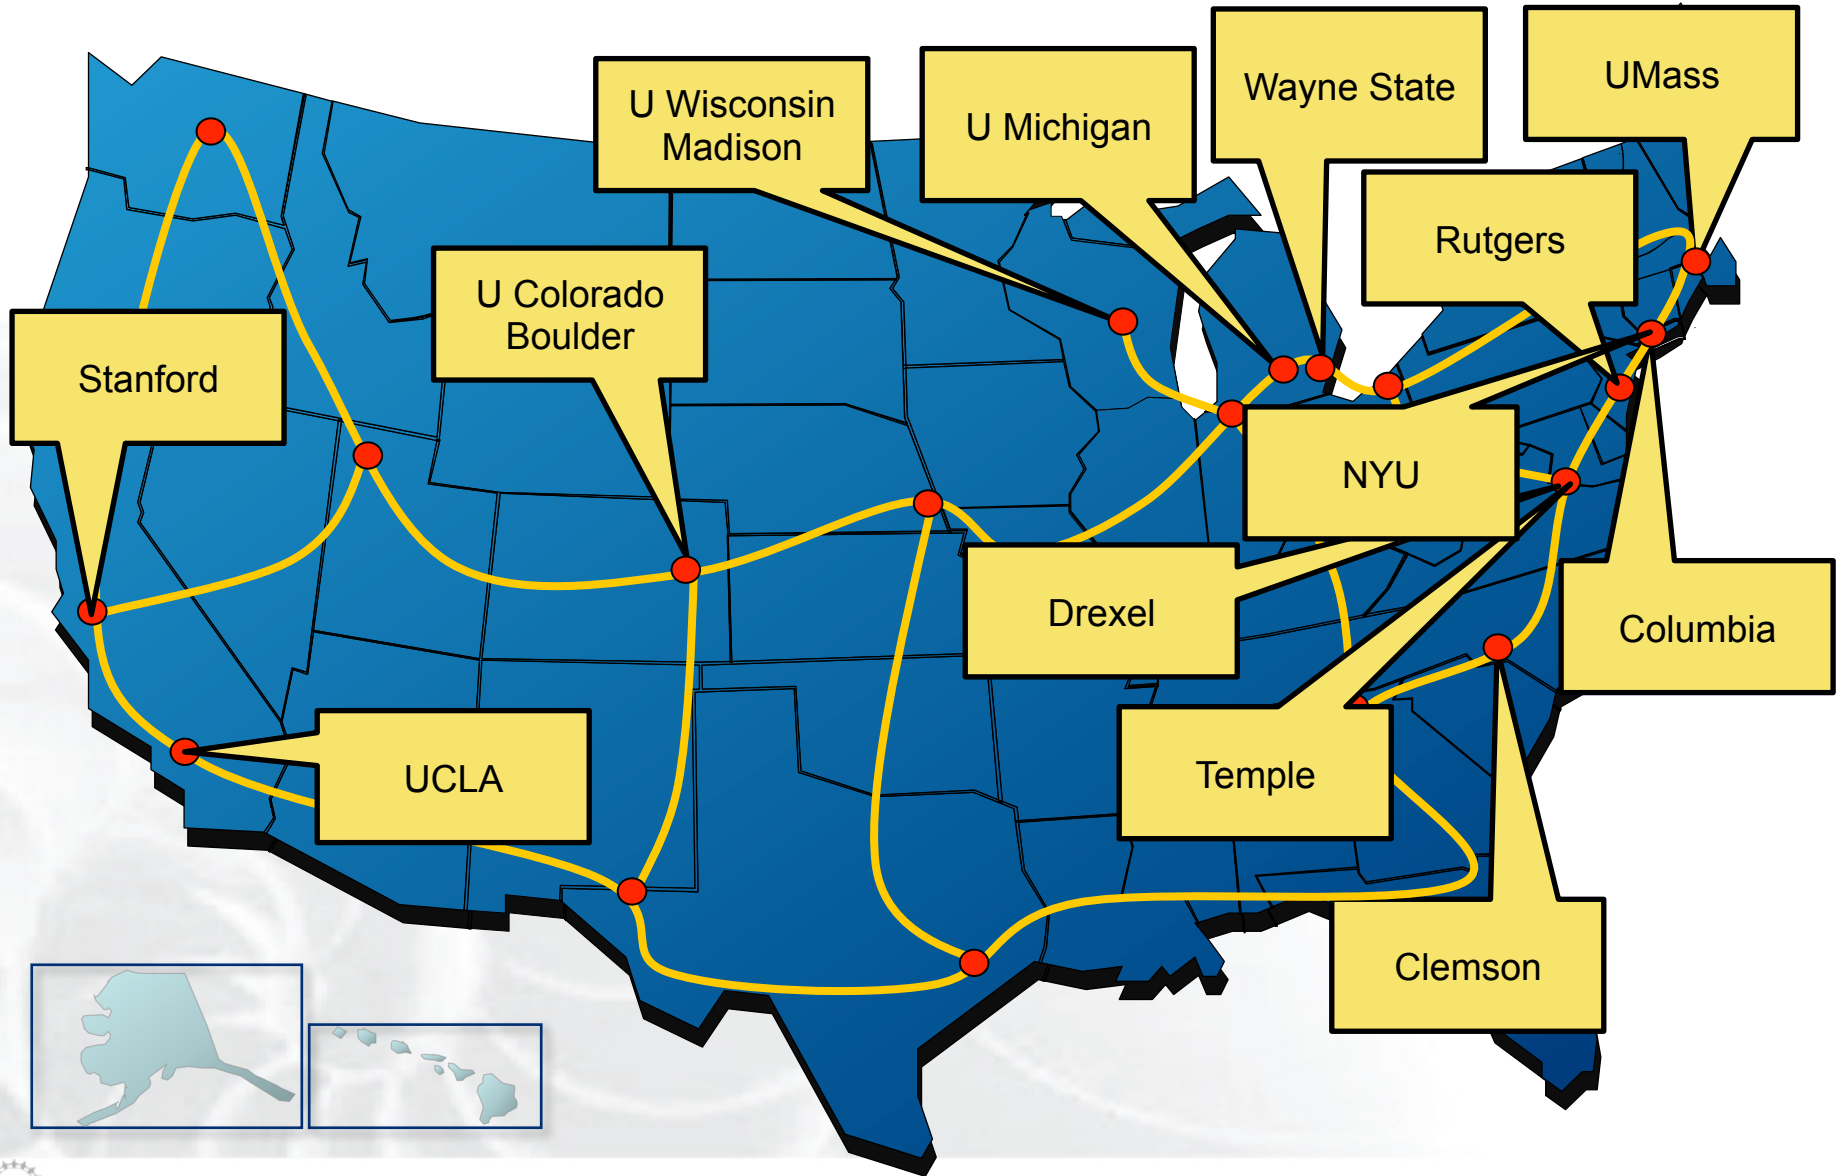
# Extra Devices in a Rack



AM may control extra devices or just provide connectivity



# GENI WiMAX Deployment 2014



- **Agreement with Sprint**


- Sprint and Rutgers University have signed a **master spectrum agreement**
- encompassing all WiMAX sites, to ensure **operation in the EBS Band**.
- An **emergency stop procedure**, in case of interference with Sprint service, has been agreed upon.

- **SciWinet GENI Mobile Virtual Network Operator (MVNO)**

- Partner with Sprint and Arterra (a Sprint partner) to create and operate an (**MVNO**) that serves the academic research community
- The effort is led by Jim Martin and Ivan Seskar, to learn more:  
<http://sciwinet.org>

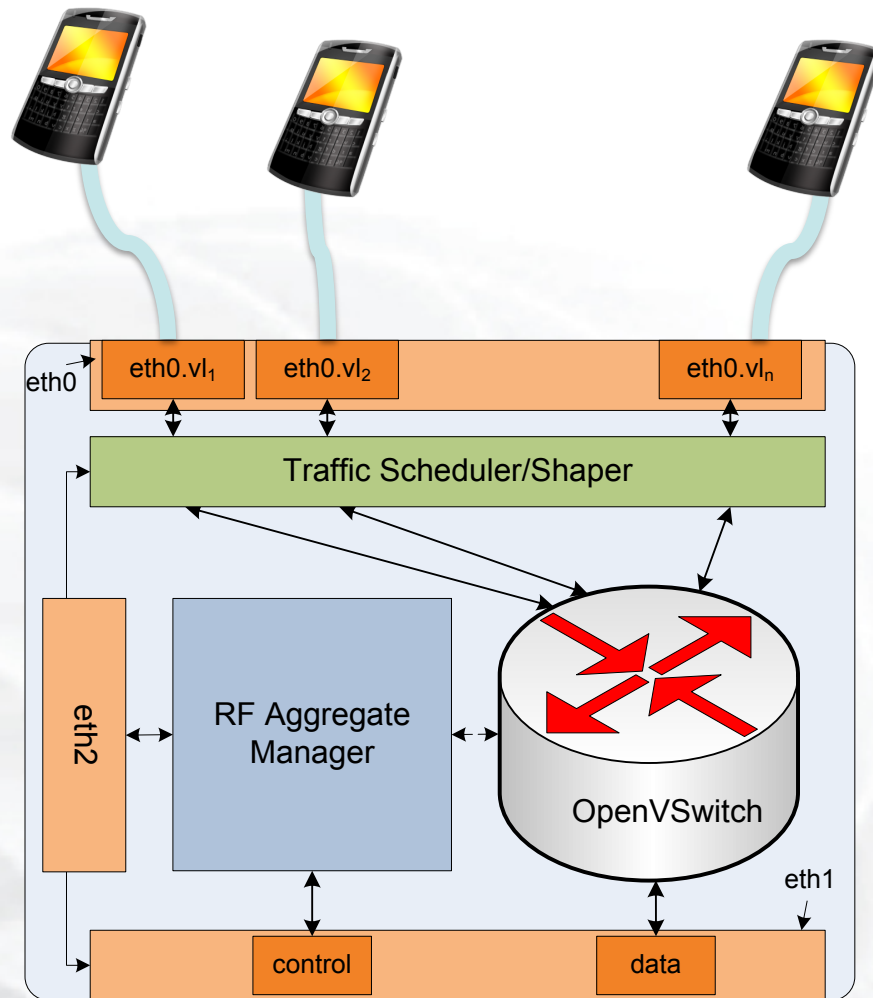
WiMAX Developers session  
*Wed: 11:30am – 12:30pm*

# GENI WiMAX Site Network Architecture



- WiMAX and Wifi edge networks.
- Layer 2 dataplane connectivity to GENI racks.
- Multi-point VLAN interconnecting all WiMAX sites via racks.

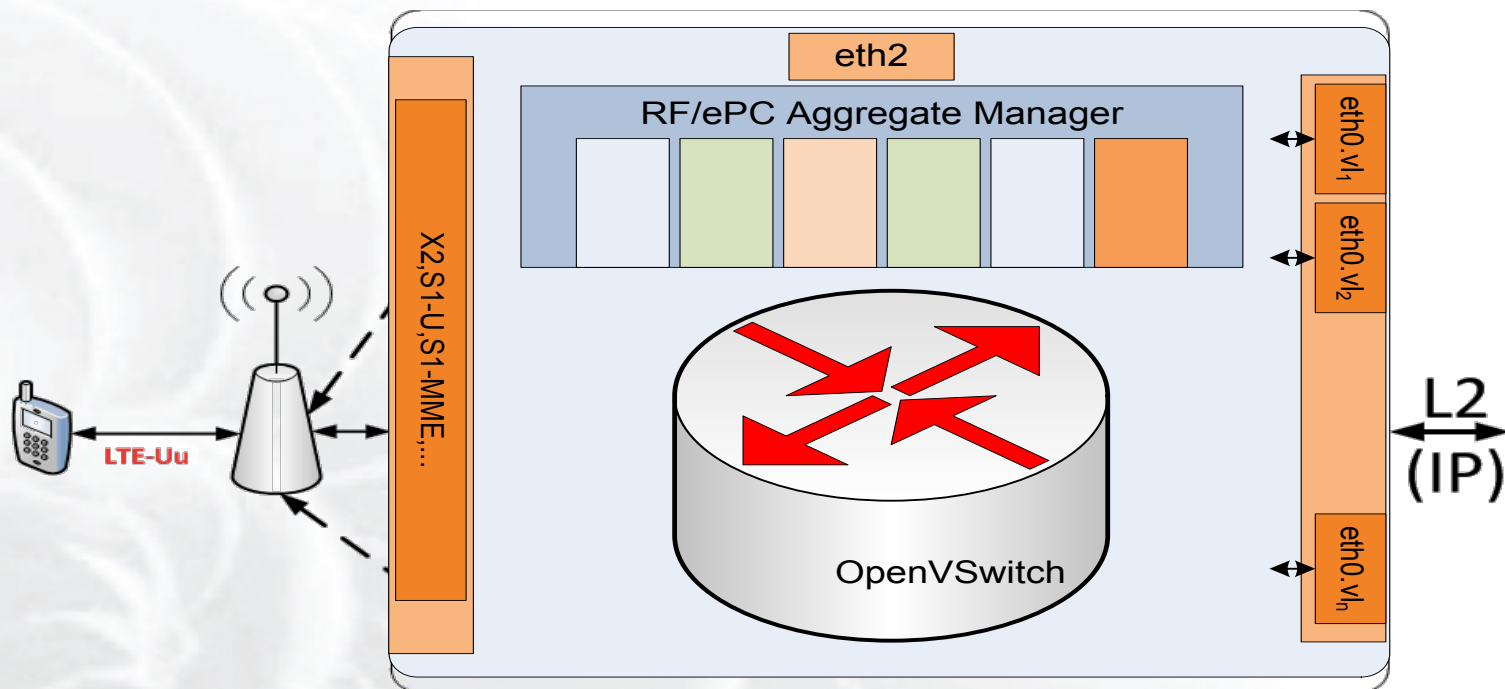
# OPEN BTS Software: WiMAX



- OMF AM REST Interface
- Each slice contains a sw router (click or OVS)
- Slicing of OFDMA resource blocks by mapping client MAC to VLANs
- Clients are mapped to VLANs


# OPEN BTS Software: LTE

New LTE Base station has it own AM to manage the slicing





# The GENI-enabled Campus



- Layer 2 transport through the campus
  - Usually statically provisioned
  - Might provide access to programmable switches
  
- Layer 2 connectivity to:
  - to GENI wireless resources
  - local resources (university labs, specialized equipment)
  - to local community tesbeds (e.g. ORBIT)
  - to metro networks

# Wisconsin: a great example



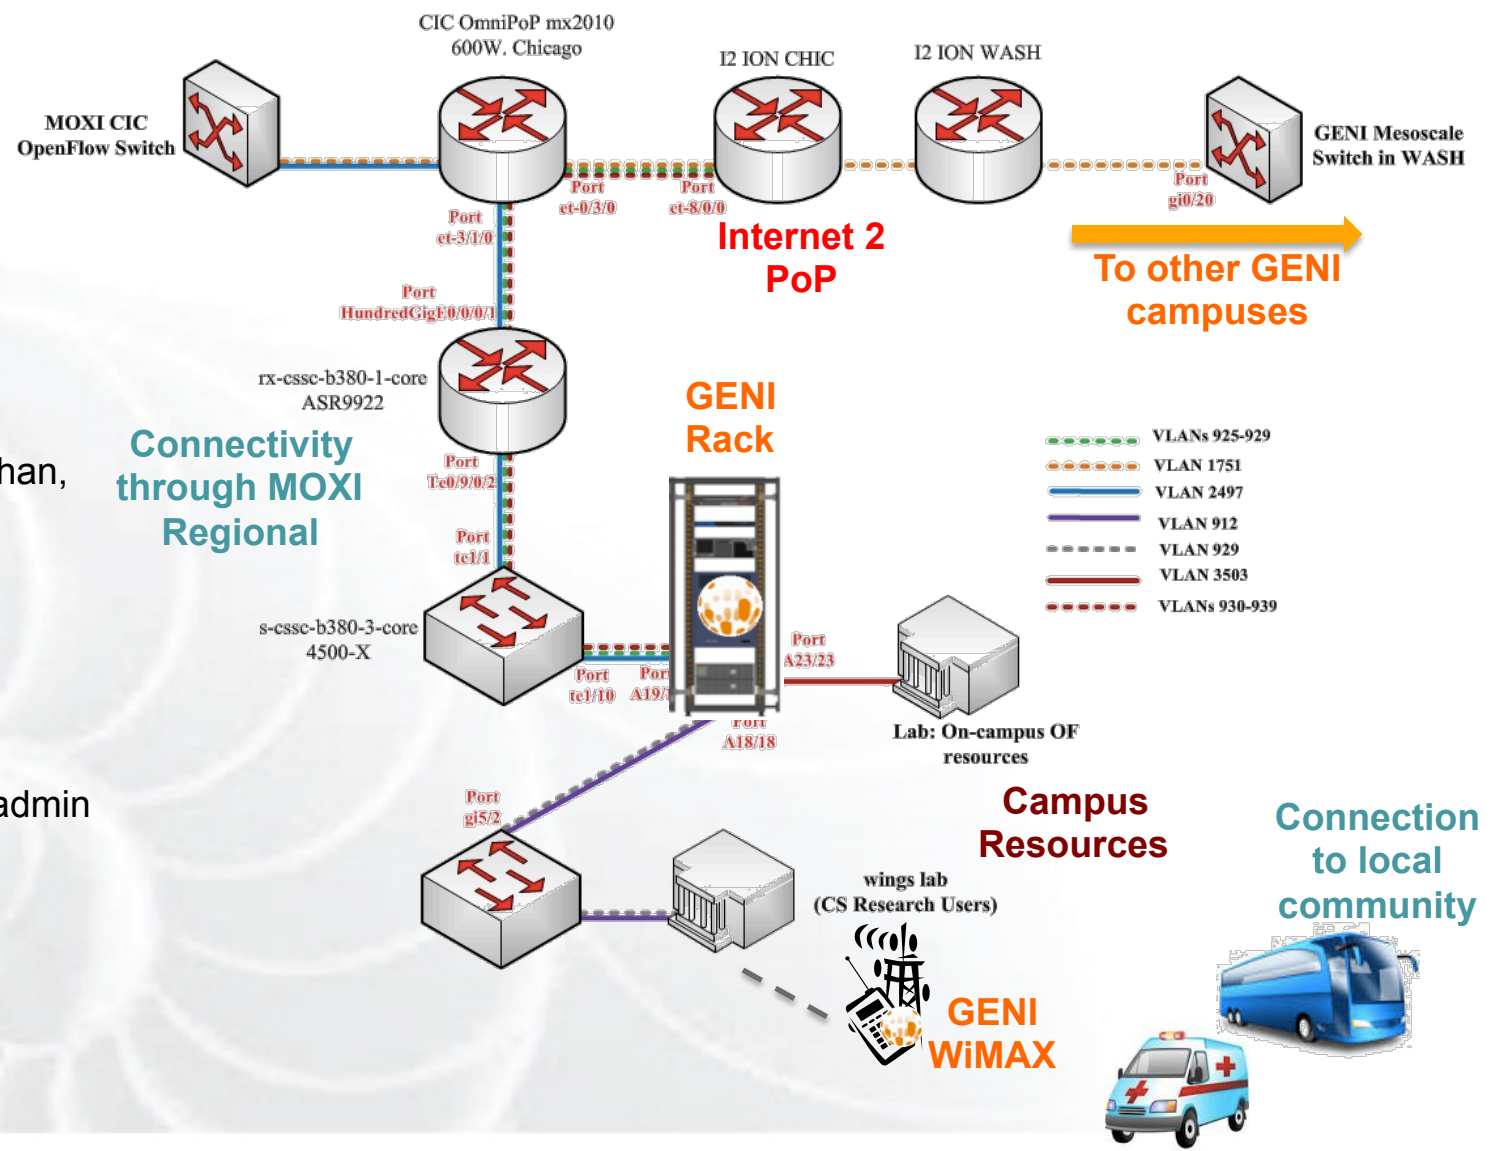
Suman Banerjee, PI



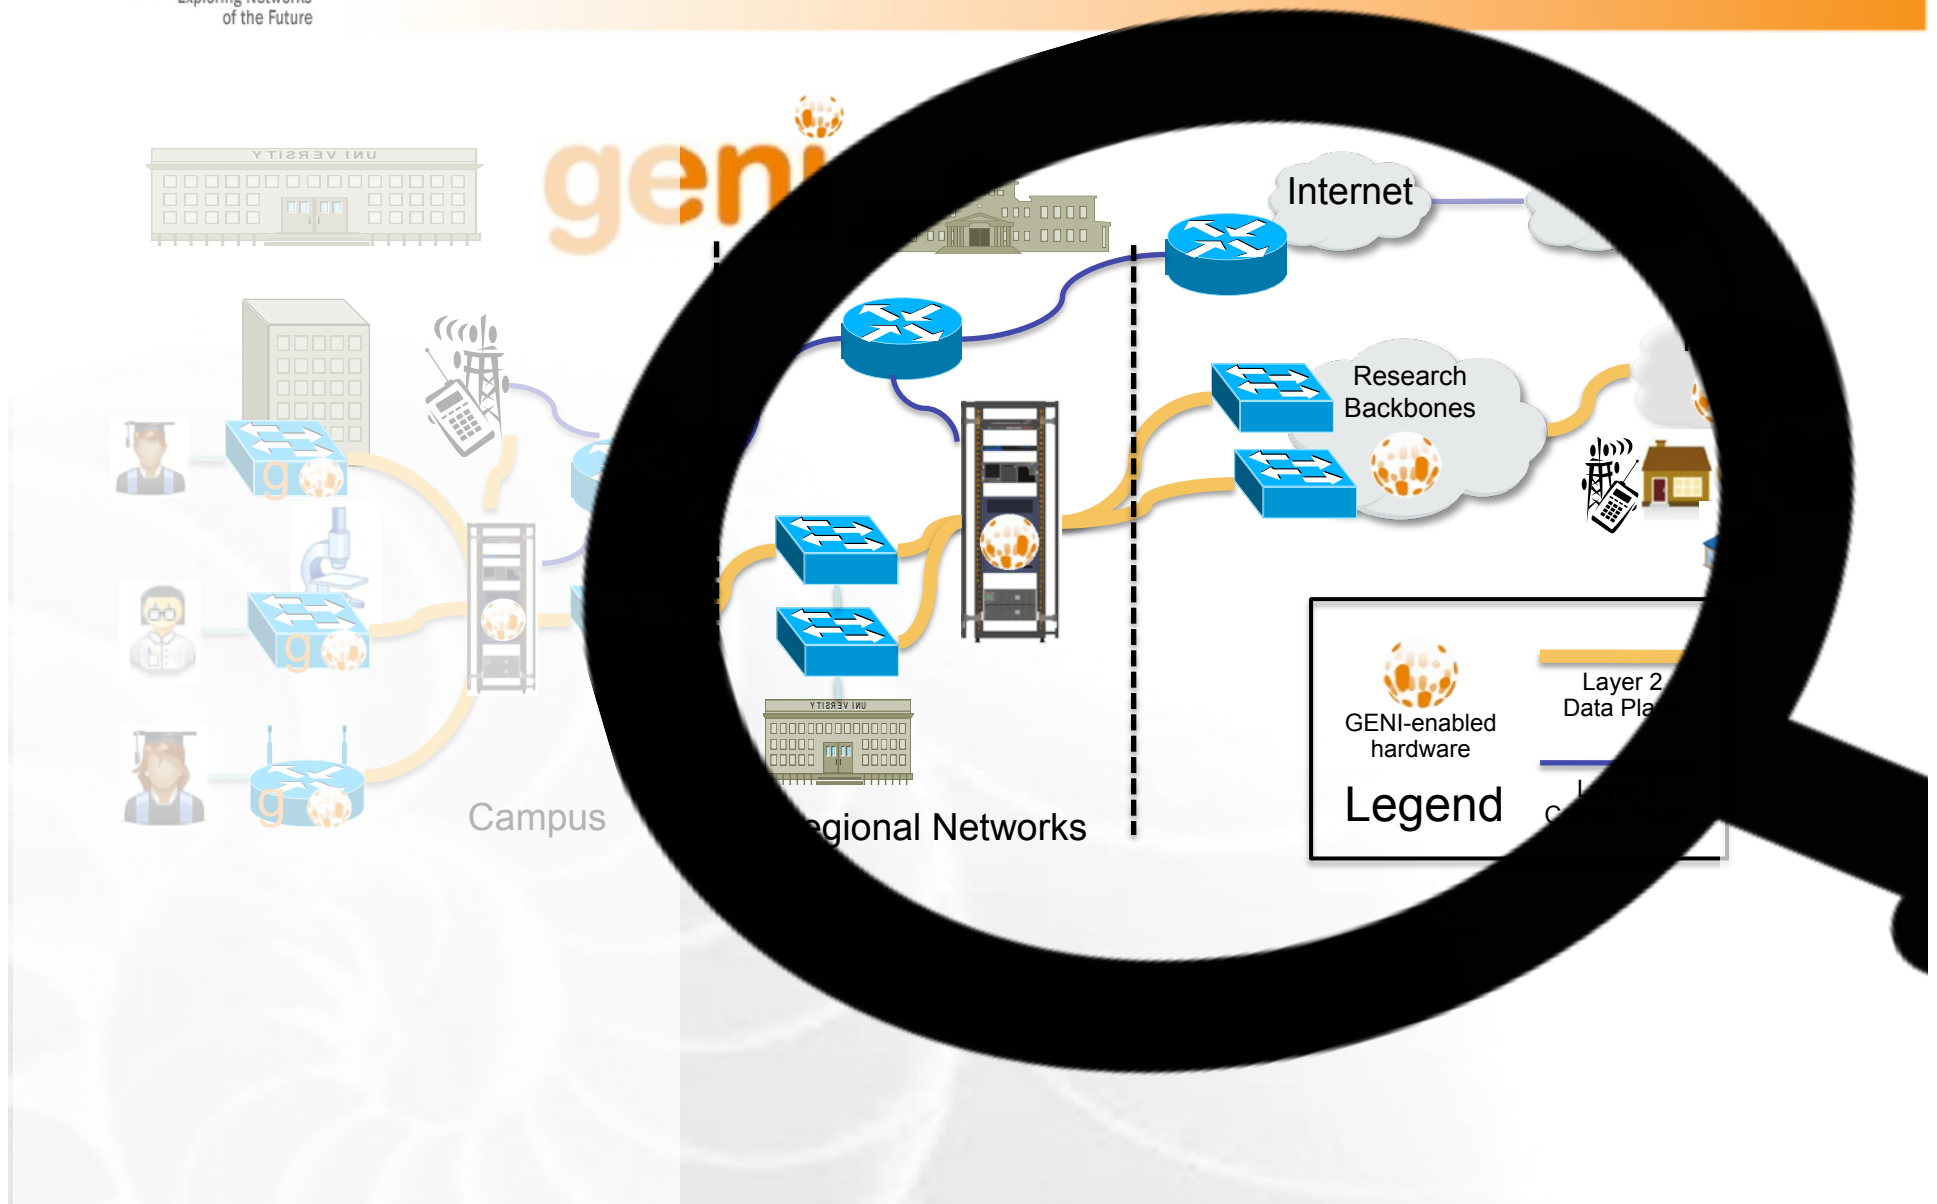
Parmesh Ramanathan, PI



Dale W. Carter, Campus admin

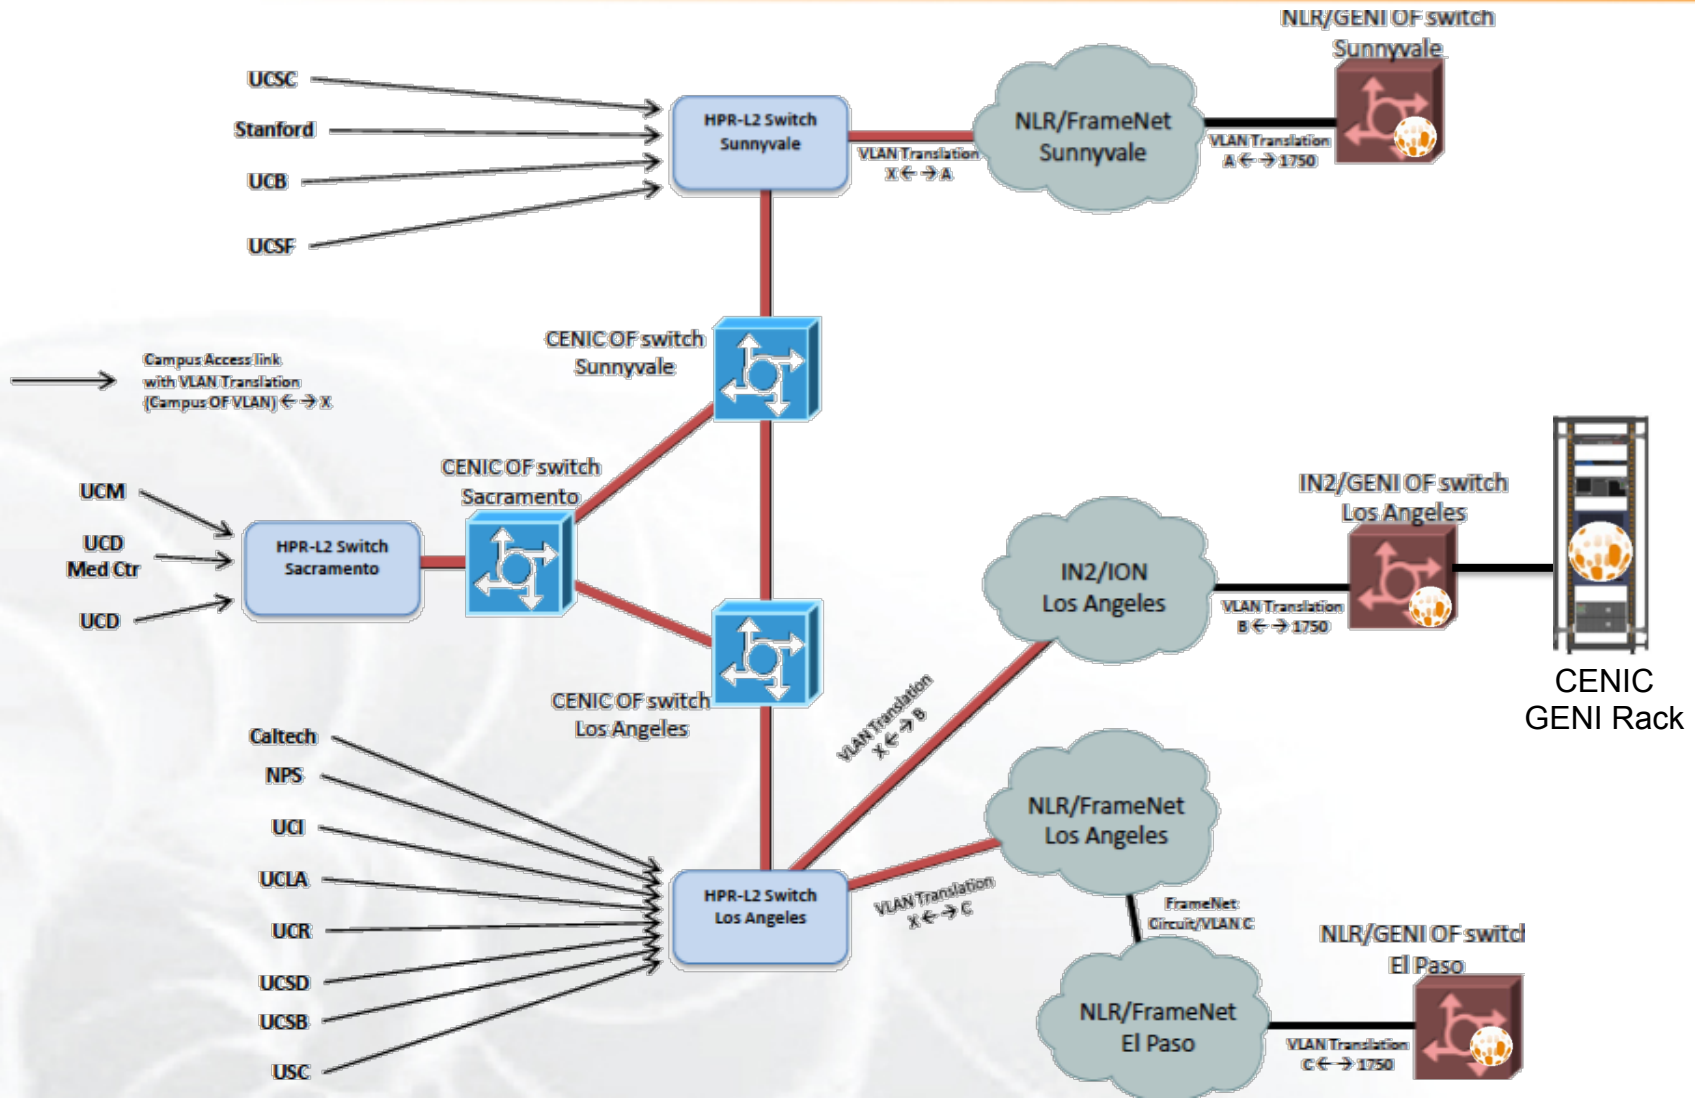


# The GENI Core Network




- Provide simple transport
  - static network configuration of multiple connections managed by edge Aggregate Managers (e.g. NOX)
  - May provide dynamic circuit provisioning (e.g. MAX)
- Provide access to routing to experimenters through SDN capable devices (e.g. StarLight)
- Provide access to compute and storage through a co-located GENI Rack (e.g. Cenic)

# Example regional network CENIC OpenFlow buildout

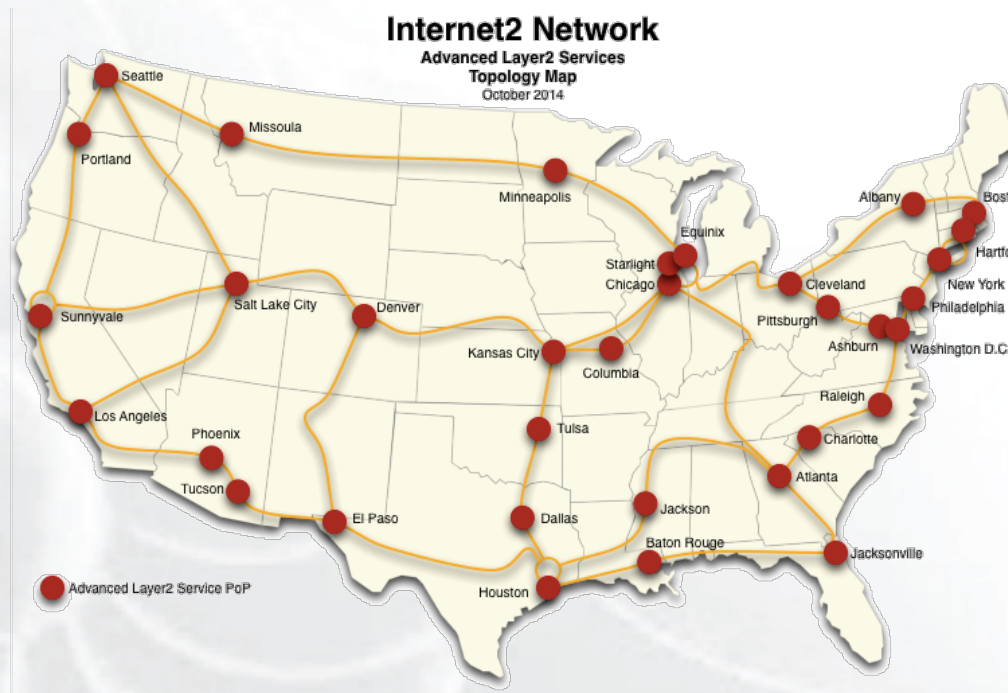


- Non-programmable dynamic Layer 2 transport through ION
- Prototype static programmable backbone with SDN capable devices (decommissioning)



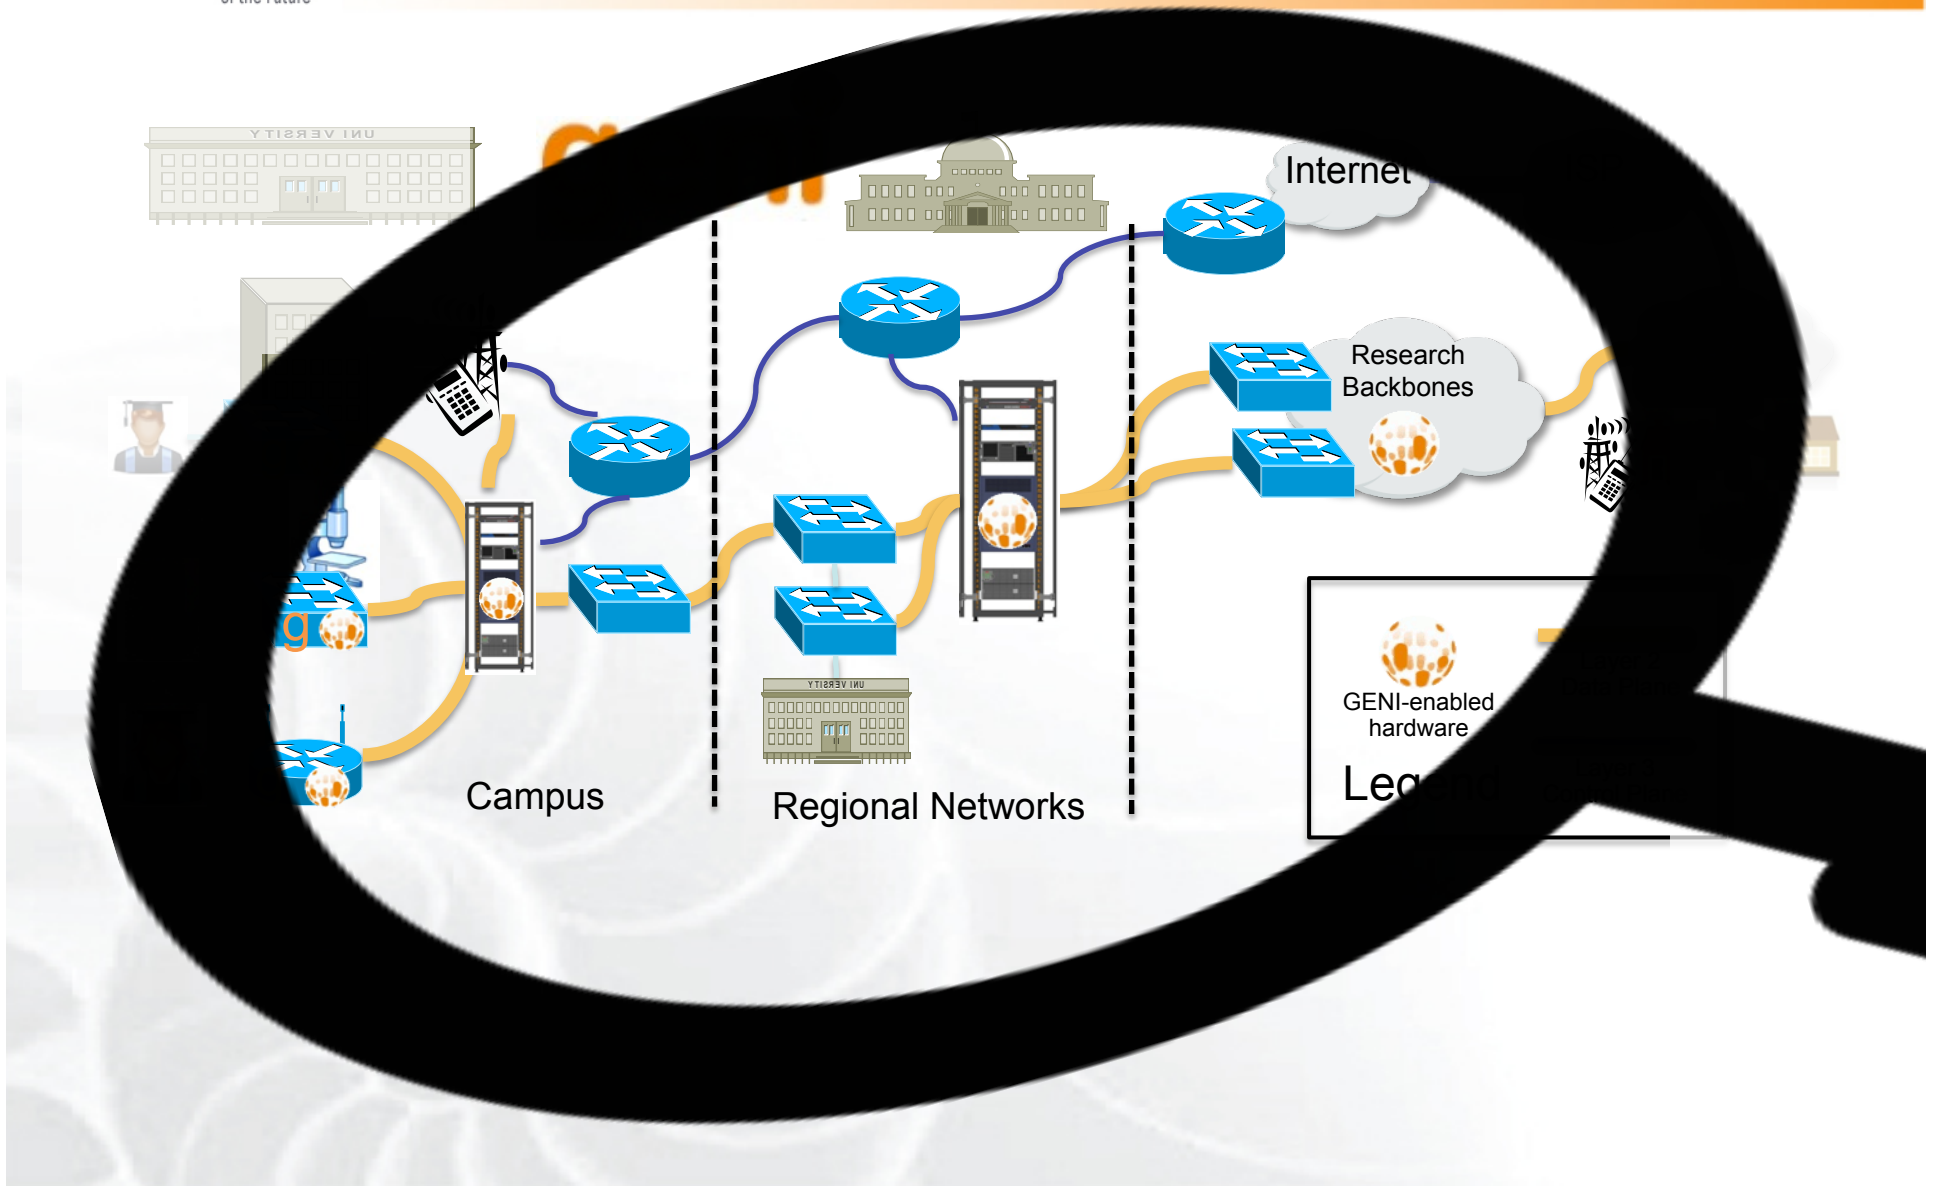
*Prototype GENI Backbone over I2 and NLR with 10 OpenFlow Switches*

- In-progress migration from “prototype GENI” to AL2S production system
  - Testing GENI dynamic provision on AL2S AM
  - Testing Experimenter OpenFlow controllers on AL2S






# Creating Inter-Aggregate Experiments



- Setup point-to-point VLANs
  - Between hosts on different Ams
    - One host/per AM/per stitch
  - Not a broadcast domain
- Dynamic, real-time setup
  - Need to coordinate multiple AMs
  - Takes time
  - Can fail
- Provides traffic isolation and bandwidth constraints



*A common concept used in other networks, applied to GENI, e.g. OSCARS, GLIF*

## How does GENI Stitching Work?

### 1. Rack Configuration (network admins)

- Long process (~weeks, months)
- Done once in advance
- Manual




### 2. Inter-aggregate link reservations (experimenters)

- Automated (tools can make them)
- Quickish (usually a few minutes)
- Live, Easy
- Repeatable





# Example “Pre-work” for End-to-End Data Connections

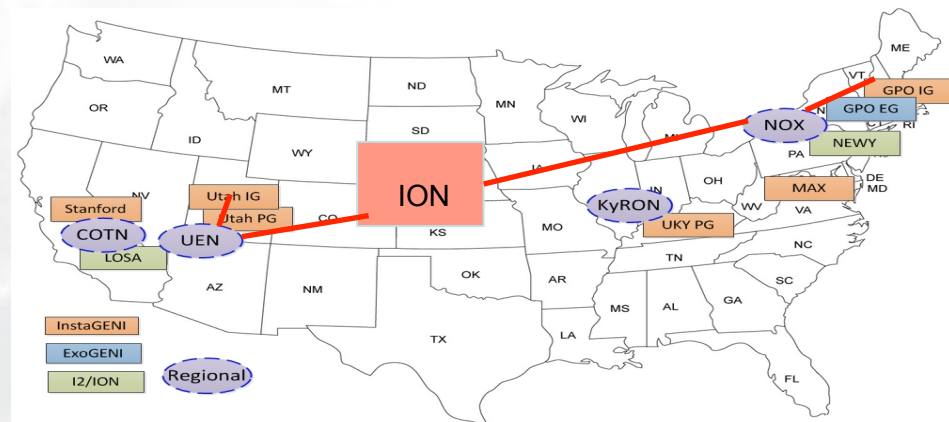


- Identify paths from a rack to GENI core
  - Identify the network providers
    - Typically a campus, a regional, a nationwide provider(GATech, Missouri, Stanford, MAX, SOX, ION, and AL2S)
  - Identify **endpoints** and allocate **VLAN to GENI** that can be used to connect to the rack
- Configure racks with VLANs for GENI stitching
- Test for connectivity

# Stitching Computation Service

Finding a workable path, and the right reservation order can be hard.

- **Stitching Computation Service (SCS) for path and workflow computation**
  - Tom Lehman and Xi Yang wrote this optional service
  - Includes many heuristics to optimize path, chance of success
  - Allows excluding particular connection points, VLANs
- **Other tools may use different heuristics**
  - Stitcher uses the SCS



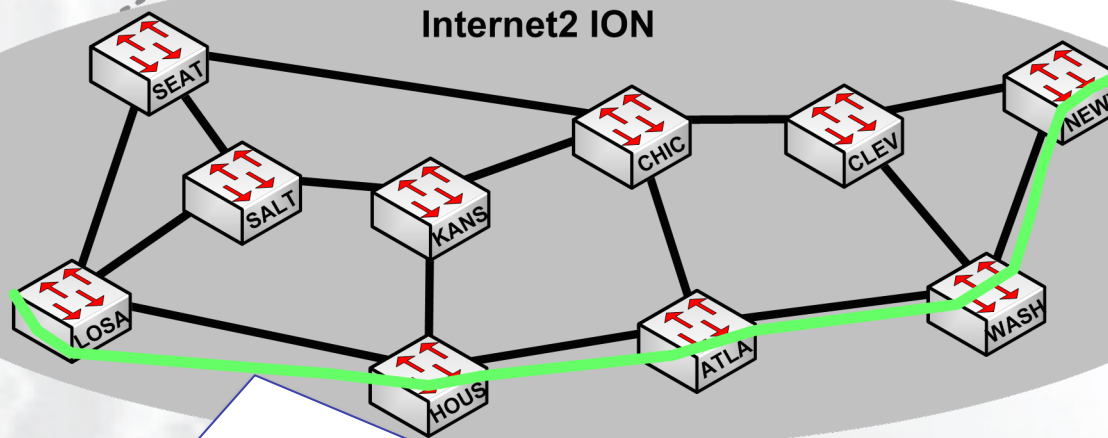
<https://wiki.maxgigapop.net/twiki/bin/view/GENI/NetworkStitchingAPI>

## OSCARS (On-Demand Secure Circuits and Advance Reservation)

- Dynamically provisions circuits on the Internet2 production Juniper routers
- Uses JunOS scripts


OSCARS

Internet2 ION



- **L2VPN**, i.e. point-to-point Ethernet circuit, i.e. Ethernet tunneled **via MPLS**
- Bandwidth guaranteed circuits


# Internet 2 ION Aggregate Manager



Java (called by SFA with MAX patch)

JunOS Script to provision Juniper Routers  
 -L2VPN  
 -Ethernet tunneled via MPLS

- ION AM maps GENI calls to OSCARS calls
- Runs the code from MAX



# Coming Soon: Internet2 OESS

- Upcoming replacement for ION
  - Different dataplane technology: OpenFlow VLANs vs MPLS L2VPN
- Supports multipoint circuits
  - We hope to support that in GENI at some point
- AL2S Aggregate Manager
  - Will control OESS as the ION AM controls ION
  - Developed by Internet2 based on FOAM and OESS
- Can use GENI AM API to stitch between ION and AL2S






# Experimenter: Creating a Circuit

**Automated  
by the tool**

1. Simple Request



7. Manifest Back

2. Send Path Request to Stitching Computation Service (SCS)



3. Get Expanded Request

4. Send Request to Aggregate 1



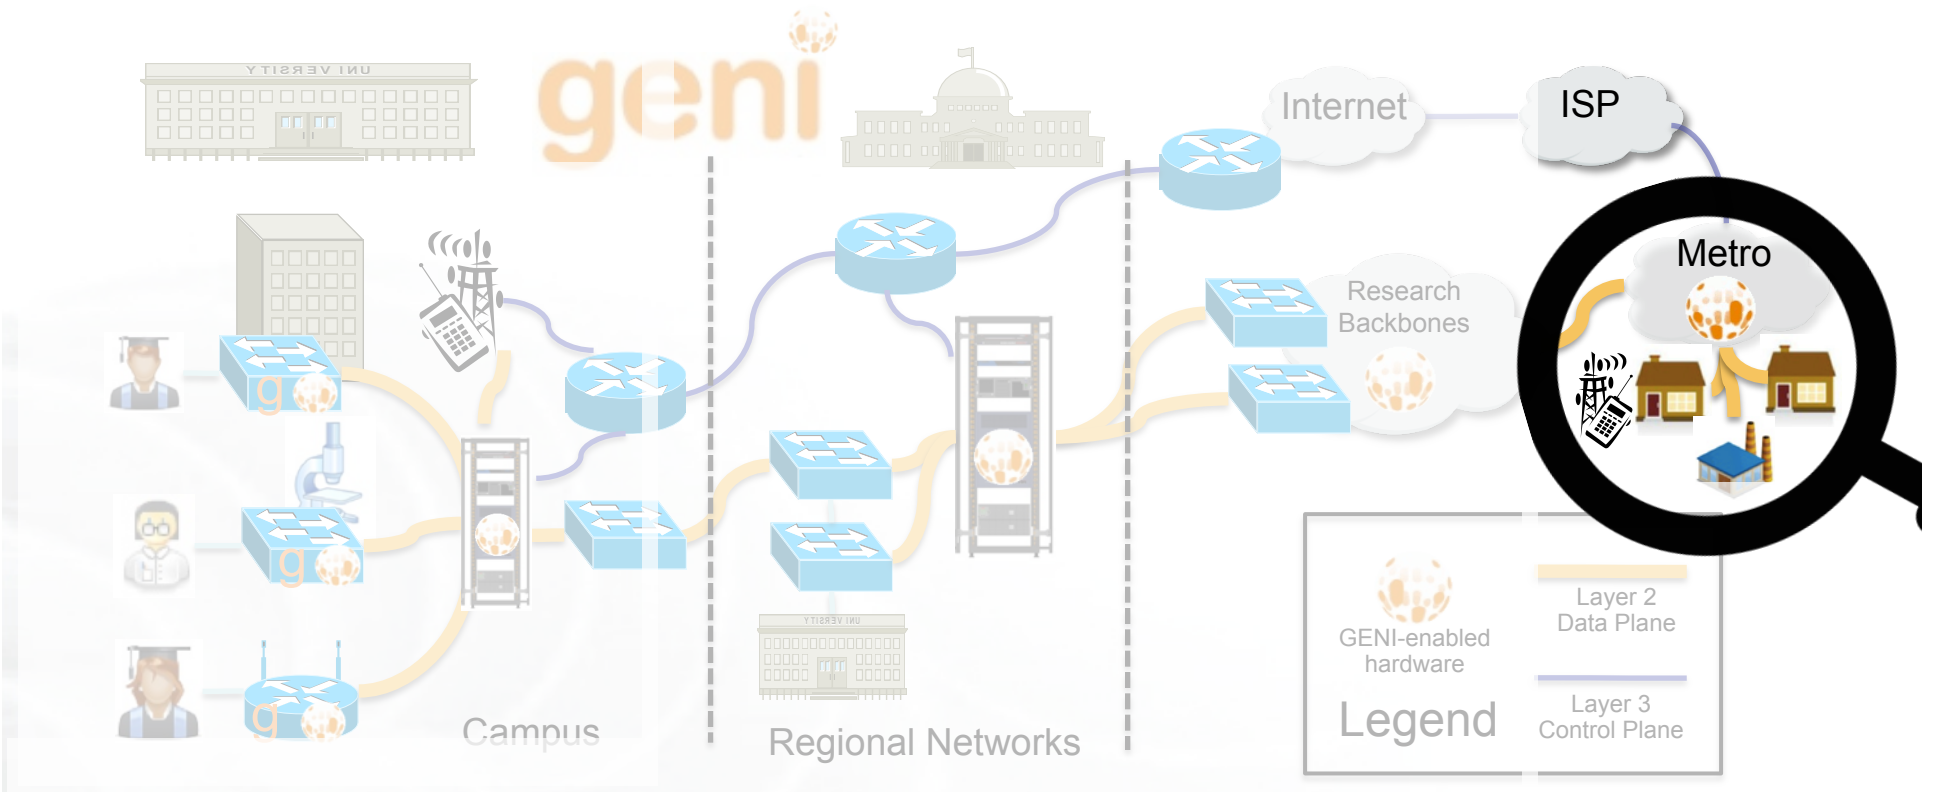
5. Get Manifest



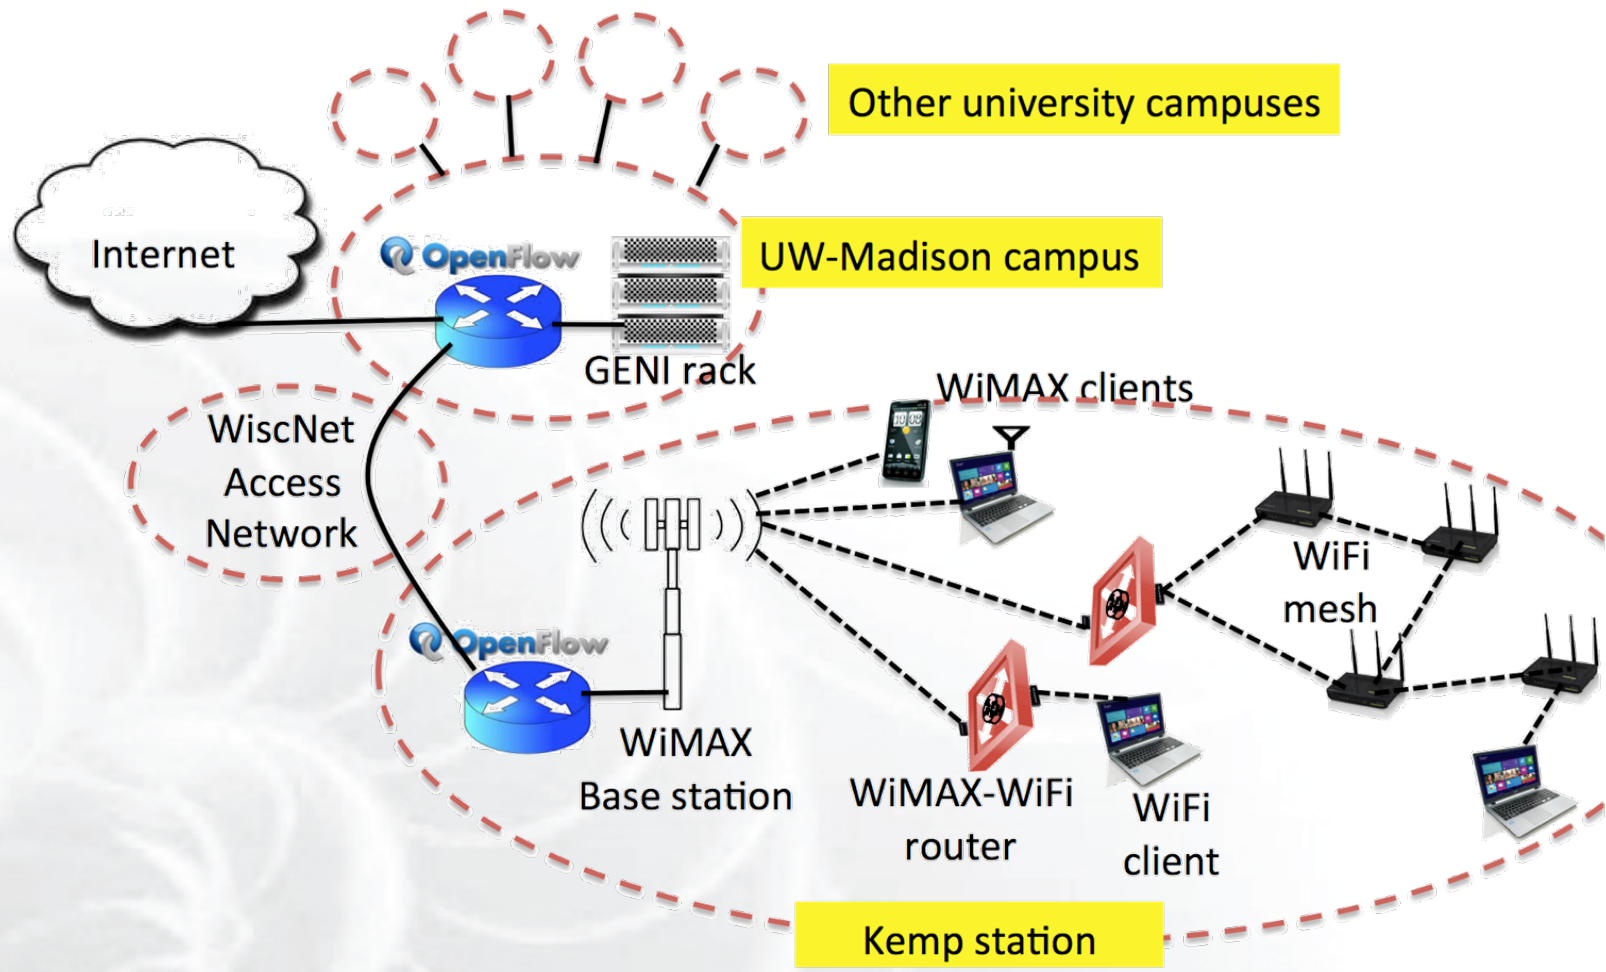
6. Repeat for Other Aggregates

*Over 16,000 stitched slivers since November, 2013*

# Running Services for non-GENI users



# Access to the local community through GENI resources



Parmesh Ramanathan  
U. Wisconsin

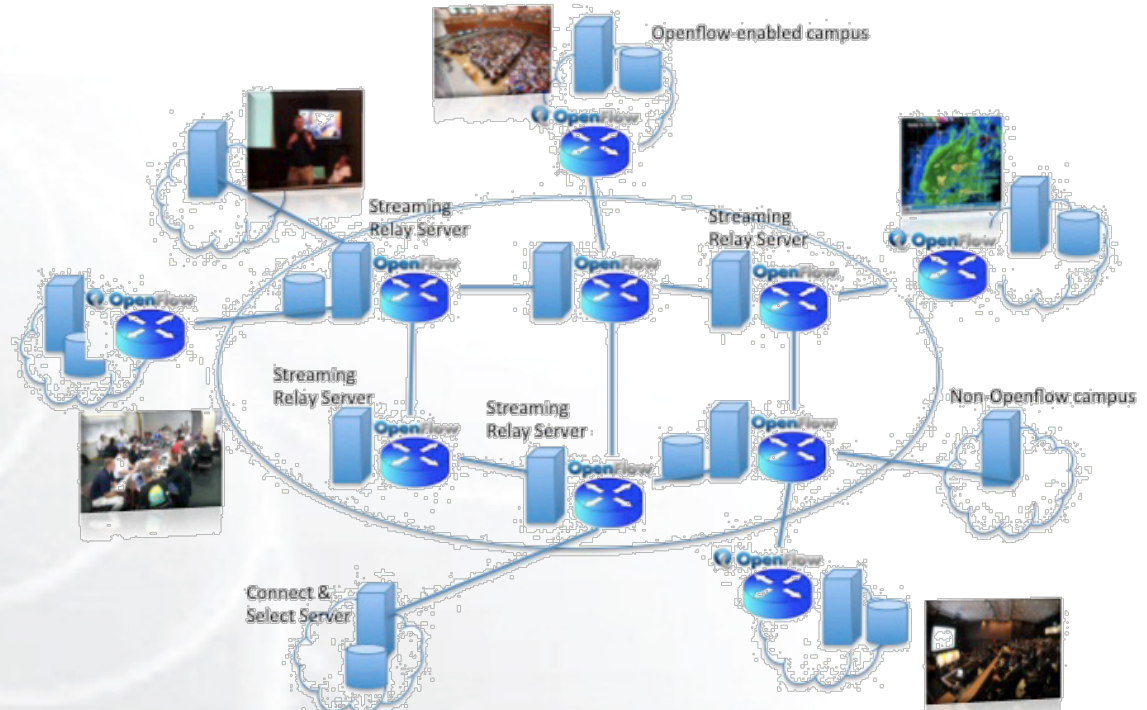


KC Wang  
U. Clemson

Use the commodity Internet

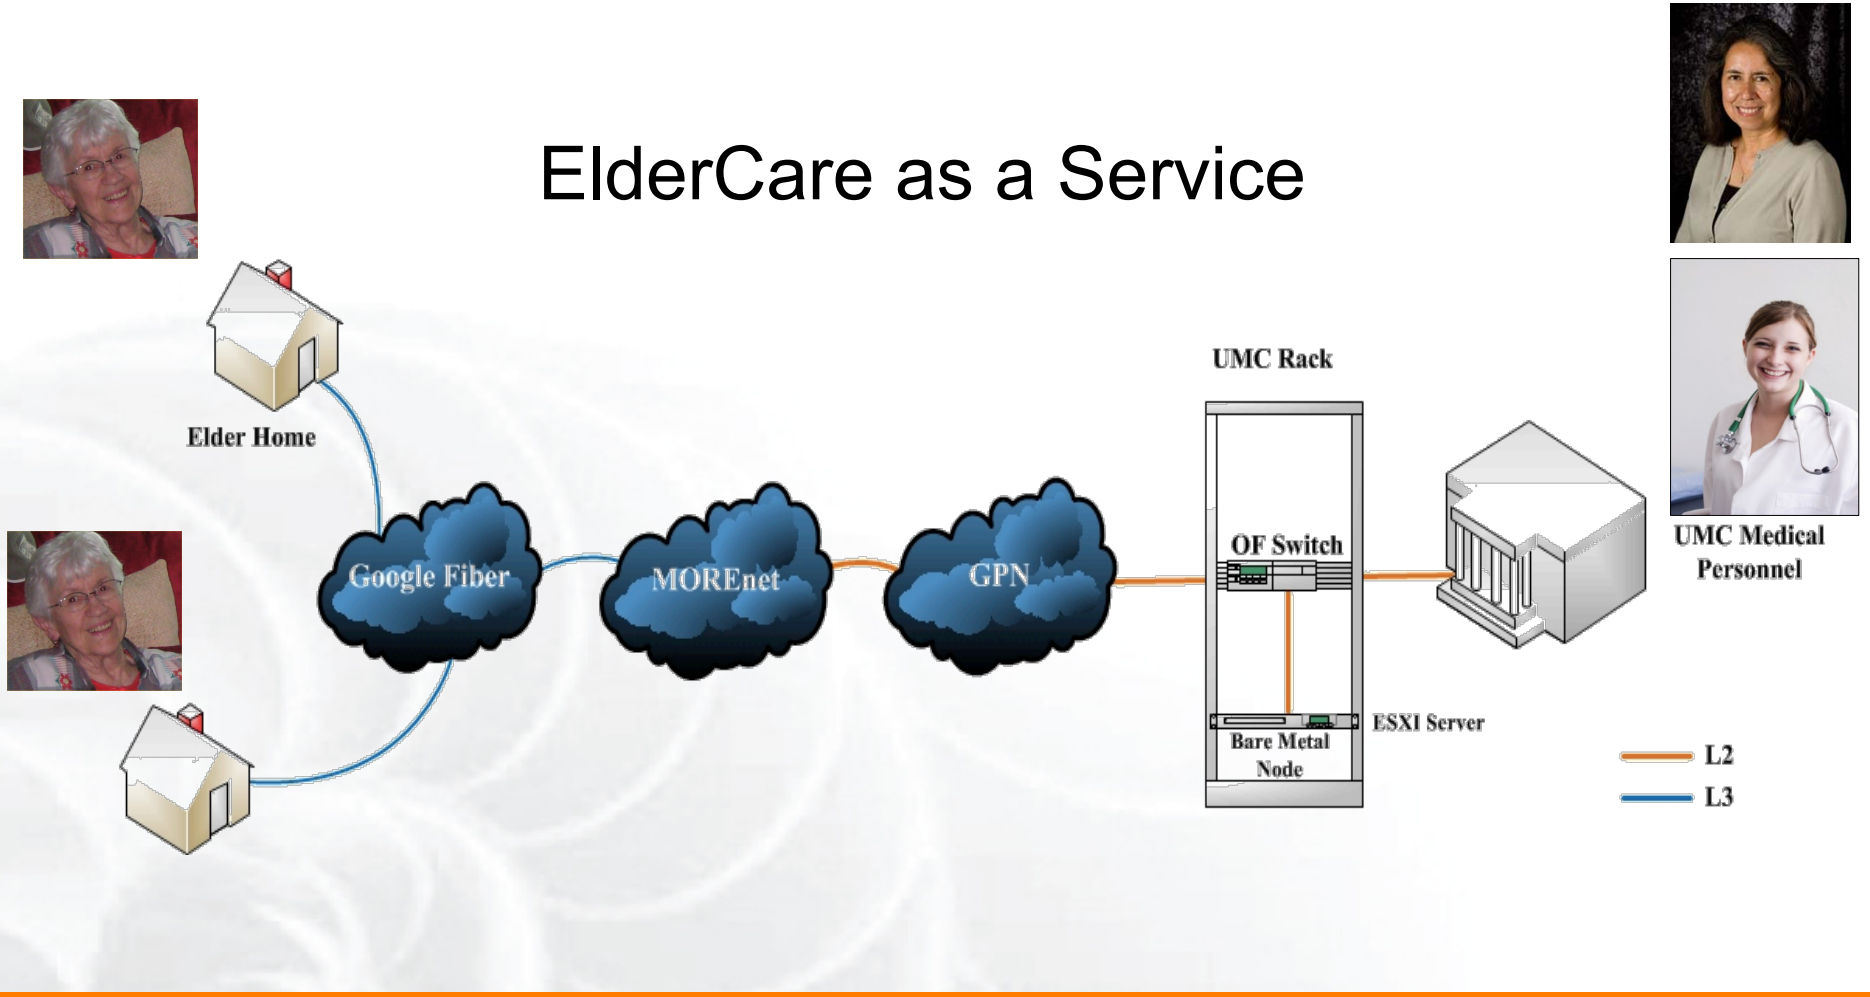
# GENI Cinema

## Persistent live video streaming service over GENI



Opt-in users can view and source live streams

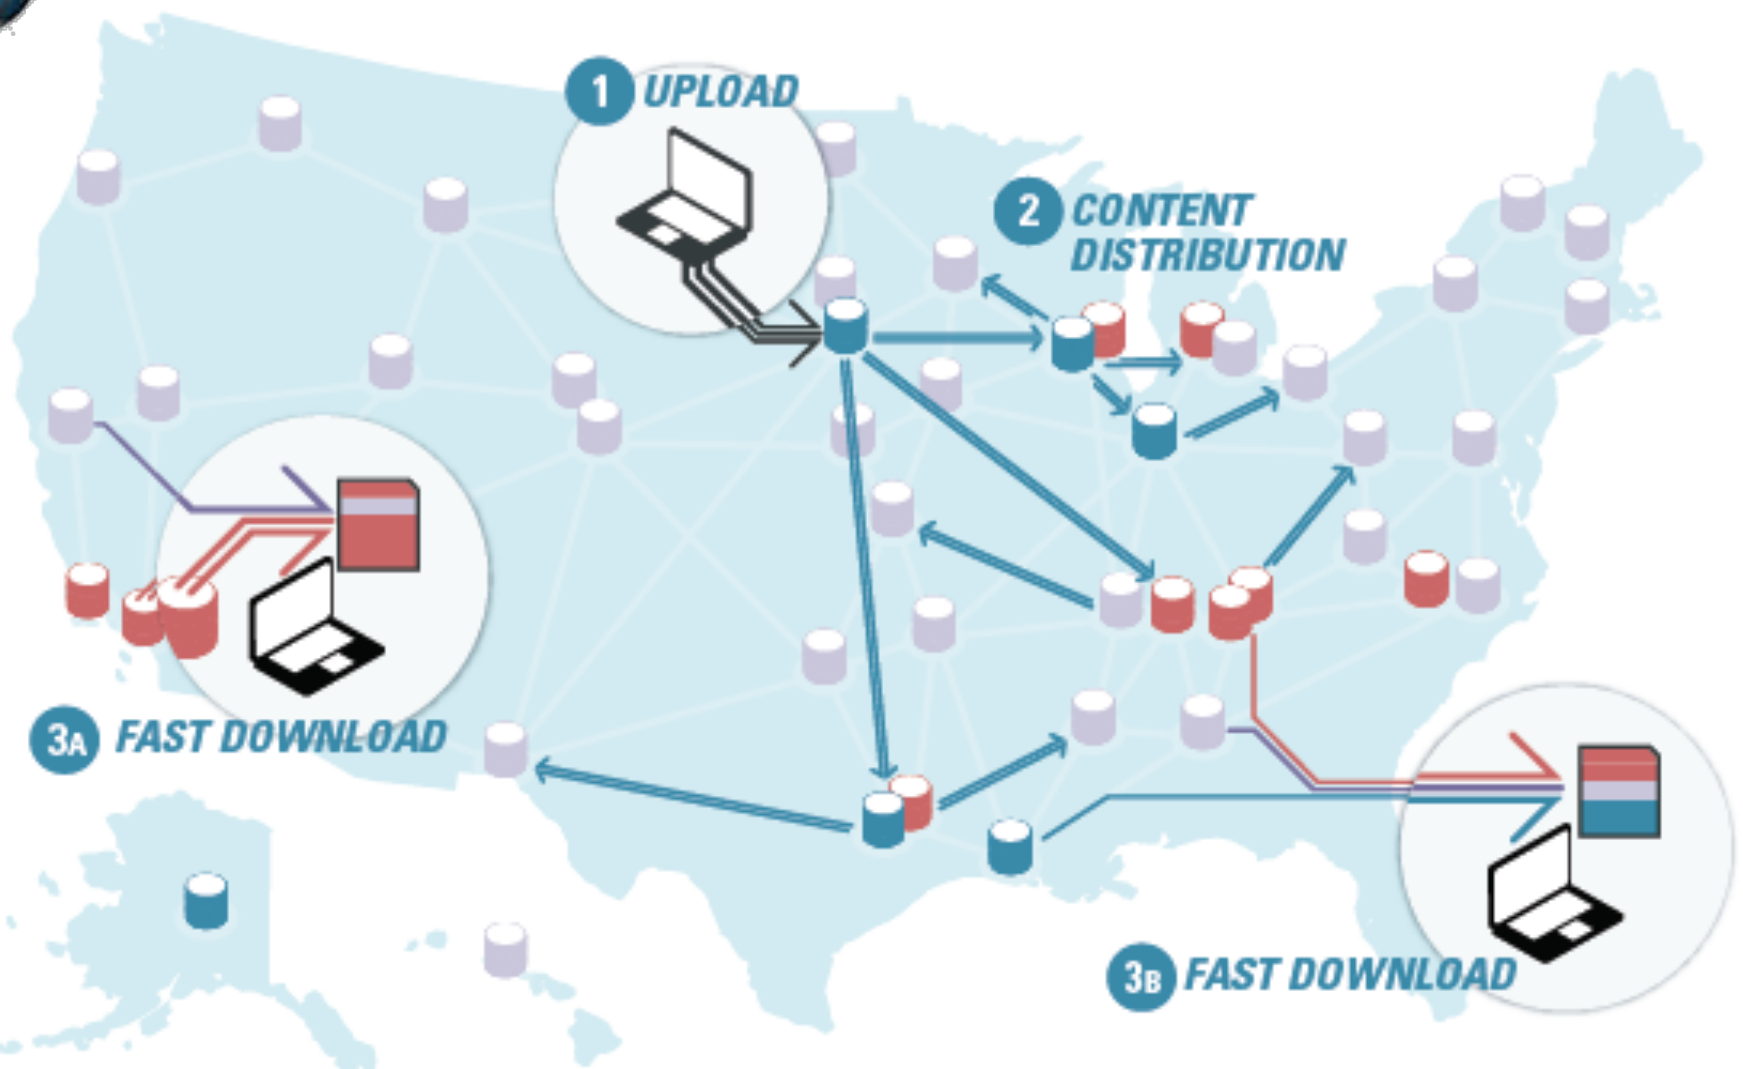
## ElderCare as a Service


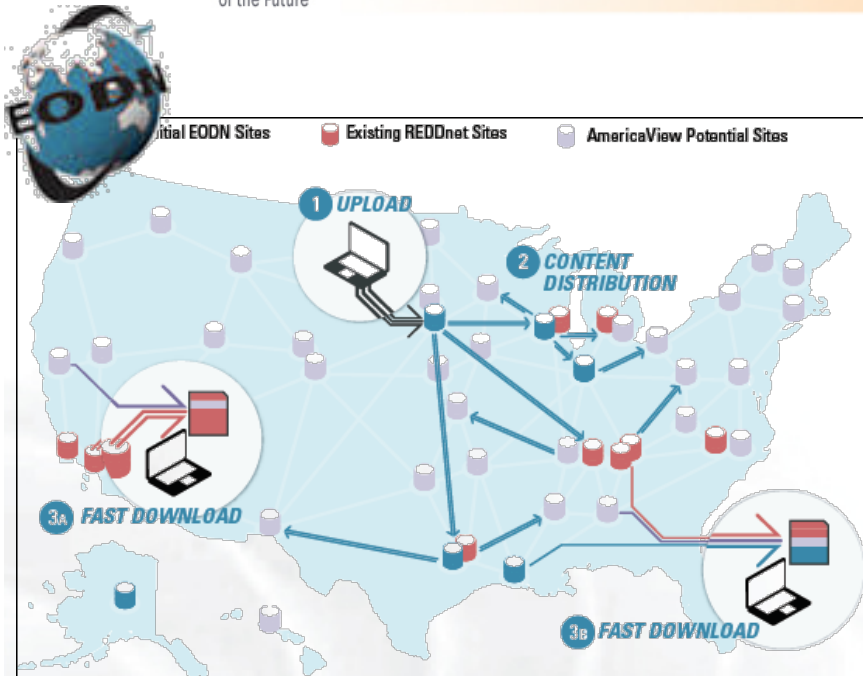


Providing service to the community.




■ Initial EODN Sites    
 ■ Existing REDDnet Sites    
 ■ AmericaView Potential Sites





# GENI's International Collaborations



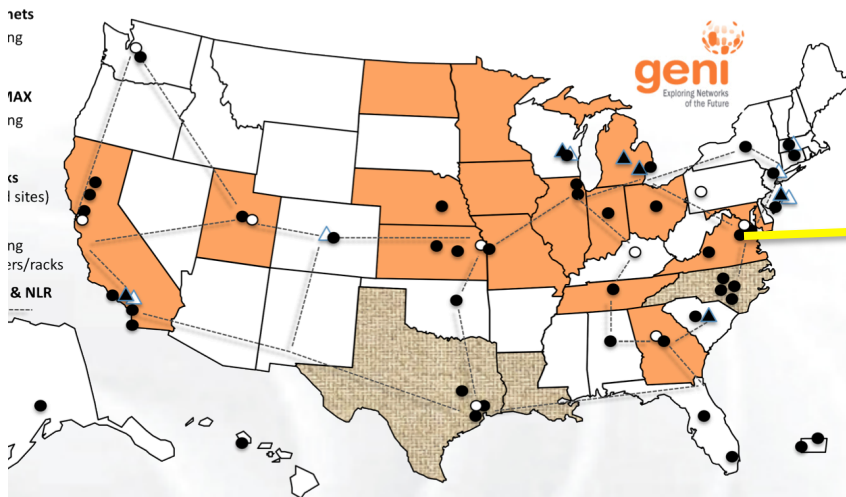


# GENI's International Collaborations



GENI is working actively with peer efforts on five continents to define and adopt common concepts and APIs.

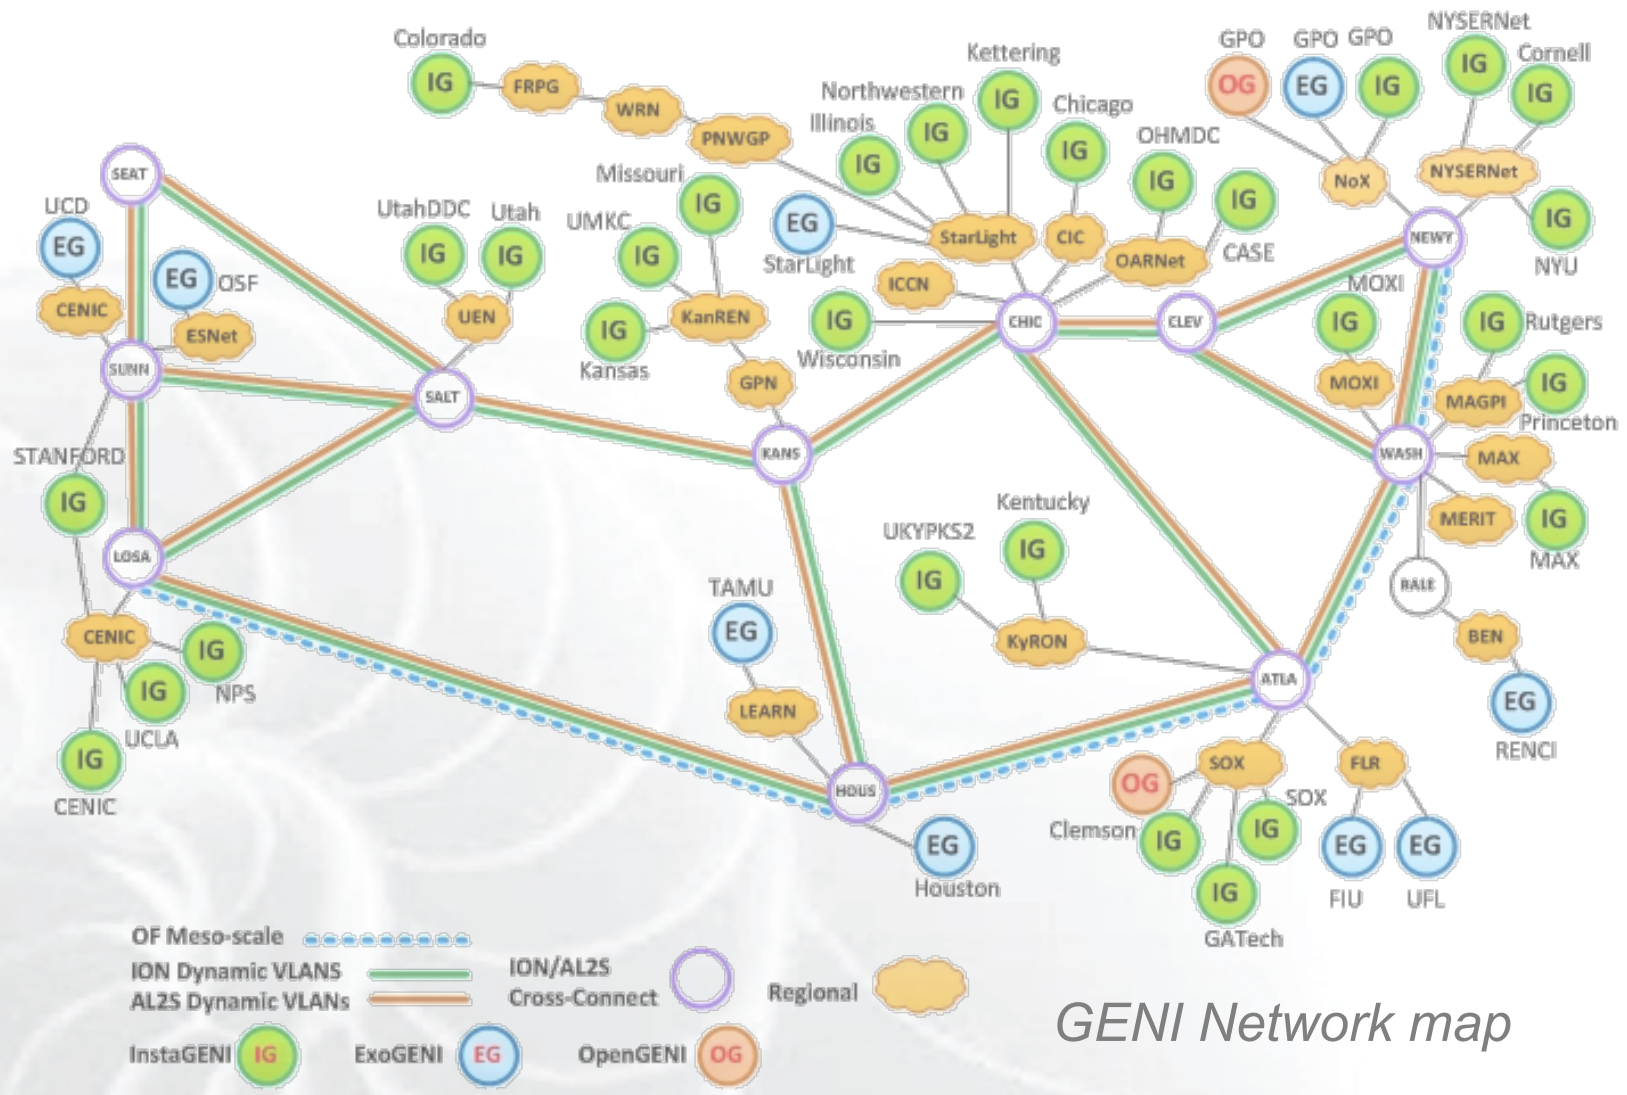




50 VLANS



- 50 Vlans between I2 PoP@ NY and iMinds
- use stitching with FIRE or GENI account
  - look for the demo on Tuesday evening



*GENI Network map*