



# GENI Project Update

IETF-78 MOBOPTS Research Group  
Maastricht, Netherlands

Aaron Falk  
GENI Engineering Architect  
July 29, 2010  
[www.geni.net](http://www.geni.net)



- GENI – Exploring future internets at scale
- GENI's status and plans
  - GENI Spiral 2
  - Meso-scale buildout
  - Starting experimentation
- Looking ahead

Global networks are creating  
extremely important new challenges

## Science Issues

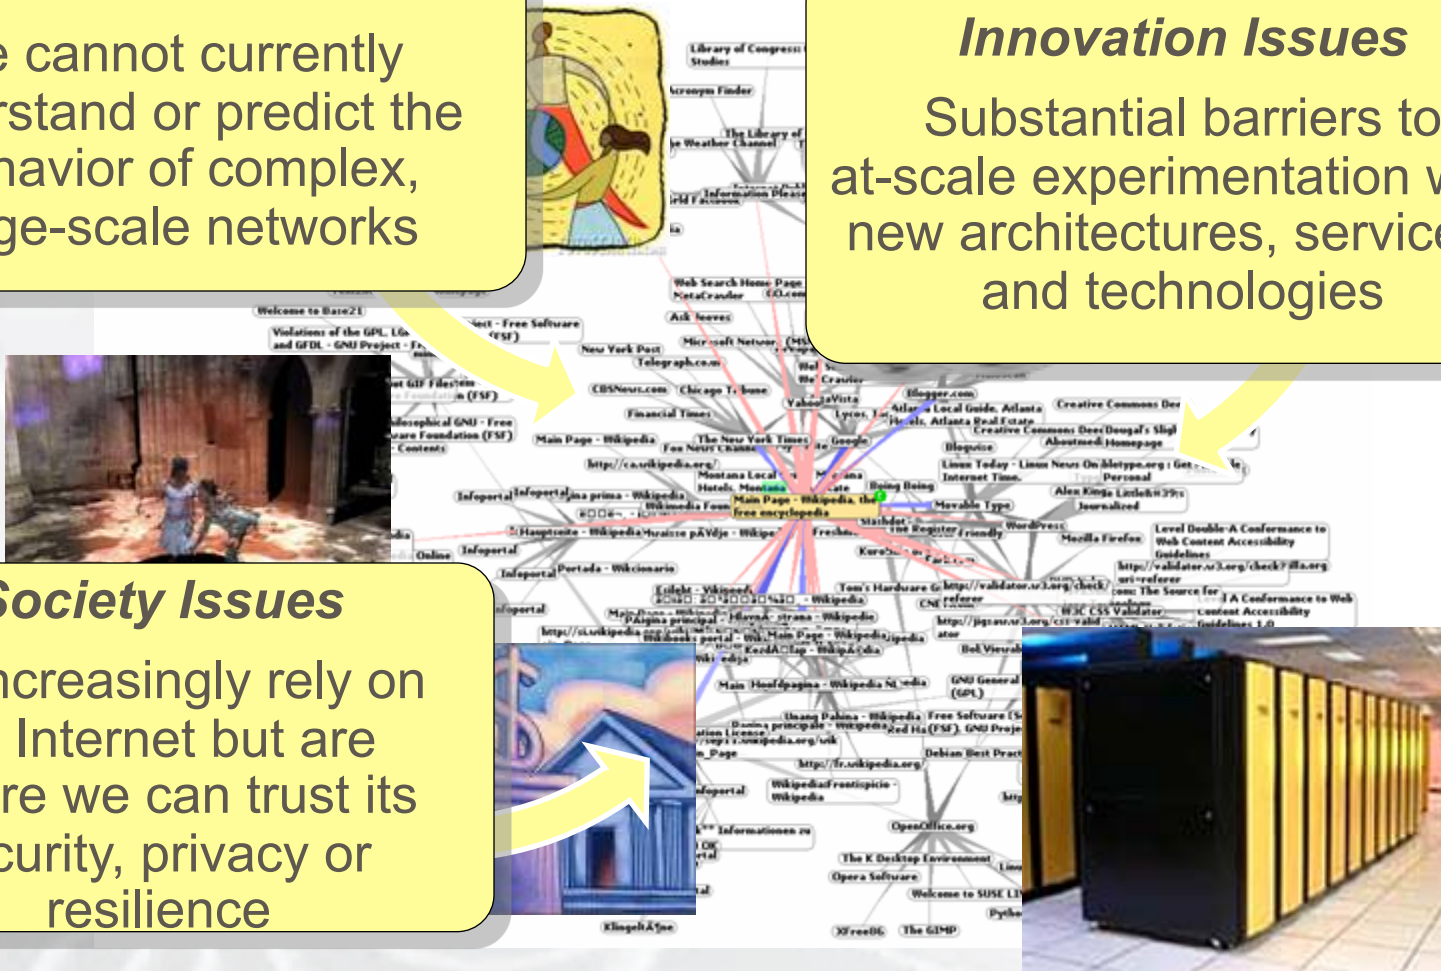
We cannot currently understand or predict the behavior of complex, large-scale networks

## ***Innovation Issues***

Substantial barriers to  
at-scale experimentation with  
new architectures, services,  
and technologies


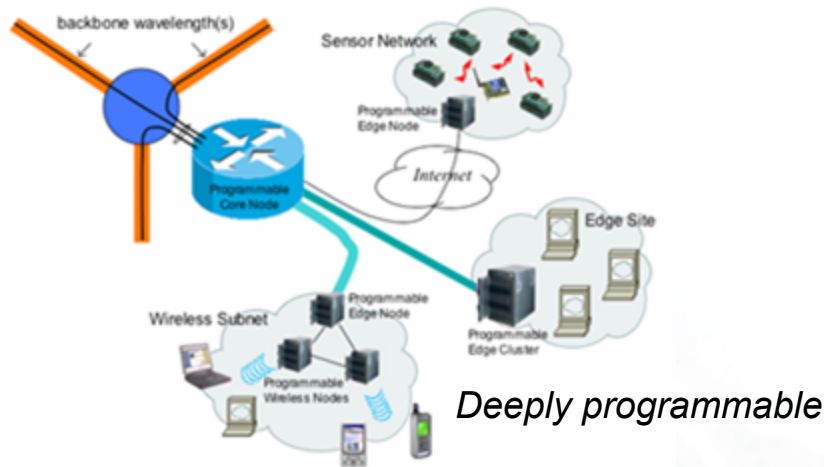
## Society Issues

We increasingly rely on the Internet but are unsure we can trust its security, privacy or resilience




# GENI Conceptual Design

## Infrastructure to support at-scale experimentation



***Programmable & federated, with end-to-end virtualized “slices”***





# How We'll Use GENI

Note that this is the “classics illustrated” version – a comic book!

Please read the Network Science and Engineering Research Agenda to learn all about the community's vision for the research it will enable.

Your suggestions are very much appreciated!

# A bright idea



I have a great idea! The original Internet architecture was designed to connect one computer to another – but a better architecture would be fundamentally based on PEOPLE and CONTENT!

*That will never work! It won't scale!  
What about security? It's impossible  
to implement or operate! Show me!*




# Trying it out



My new architecture worked great in the lab, so now I'm going to try a larger experiment for a few months.

And so he poured his experimental software into clusters of CPUs and disks, bulk data transfer devices ('routers'), and wireless access devices throughout the GENI suite, and started taking measurements . . .



He uses a modest slice of GENI, sharing its infrastructure with many other concurrent experiments.

# It turns into a really good idea


Boy did I learn a lot! I've published papers, the architecture has evolved in major ways, and I'm even attracting real users!



*Location-based social networks are really cool!*



His experiment grew larger and continued to evolve as more and more real users opted in . . .



His slice of GENI keeps growing, but GENI is still running many other concurrent experiments.

# Experiment turns into reality



My experiment was a real success, and my architecture turned out to be mostly compatible with today's Internet after all – so I'm taking it off GENI and spinning it out as a real company.

*I always said it was a good idea, but way too conservative.*





## Meanwhile . . .



I have a great idea! If the Internet were augmented with a scalable control plane and realtime measurement tools, it could be 100x as reliable as it is today . . . !

And I have a great concept for incorporating live sensor feeds into our daily lives !




If **you** have a great idea, check out the  
**NSF CISE Network Science and Engineering** program.

- GENI – Exploring future internets at scale
- GENI's status and plans
  - GENI Spiral 2
  - Meso-scale buildout
  - Starting experimentation
- Looking ahead

# Spiral Development

GENI grows through a well-structured, adaptive process



GENI Prototyping Plan



## GENI Spiral 2

Early experiments, meso-scale build, interoperable control frameworks, ongoing integration, system designs for security and instrumentation, definition of identity management plans.



## Envisioned **ultimate goal**

Example: Planning Group's desired GENI suite, probably trimmed some ways and expanded others. Incorporates large-scale distributed computing resources, high-speed backbone nodes, nationwide optical networks, wireless & sensor nets, etc.

- **Spiral Development Process**  
Re-evaluate goals and technologies yearly by a systematic process, decide what to prototype and build next.

# Current GENI Status

GENI-enabling testbeds, campuses, and backbones







PLANETLAB

An open platform for developing, deploying, and accessing planetary-scale services

ProtoGENI



UIC



ARISTA

Battelle



ciena



FUJITSU



infinera

JEFFREY HUNKER ASSOCIATES LLC  
Technology • Government • Global business • Insights with impact

JUNIPER NETWORKS



NETRONOME

Microsoft

NEC



Radio Technology Systems





# Infrastructure examples



**DRAGON core nodes**  
**Mid-Atlantic Crossroads**



**WAIL, U. Wisconsin-Madison**



**DieselNet, U. Mass Amherst**



**ViSE,**  
**U. Mass Amherst**



**SPPs, Wash U.**



**ORBIT, Rutgers WINLAB**




## National LambdaRail

Up to 30 Gbps nondedicated bandwidth

## Internet2

10 Gbps dedicated bandwidth

ProtoGENI  
& SPP



**40 Gbps capacity for GENI prototyping on two national footprints  
to provide Layer 2 Ethernet VLANs as slices (IP or non-IP)**



# Nationwide Meso-scale Prototype

Current plans for locations & equipment

## OpenFlow

Stanford  
U Washington  
Wisconsin U  
Indiana U  
Rutgers  
Princeton  
Clemson  
Georgia Tech

## ShadowNet

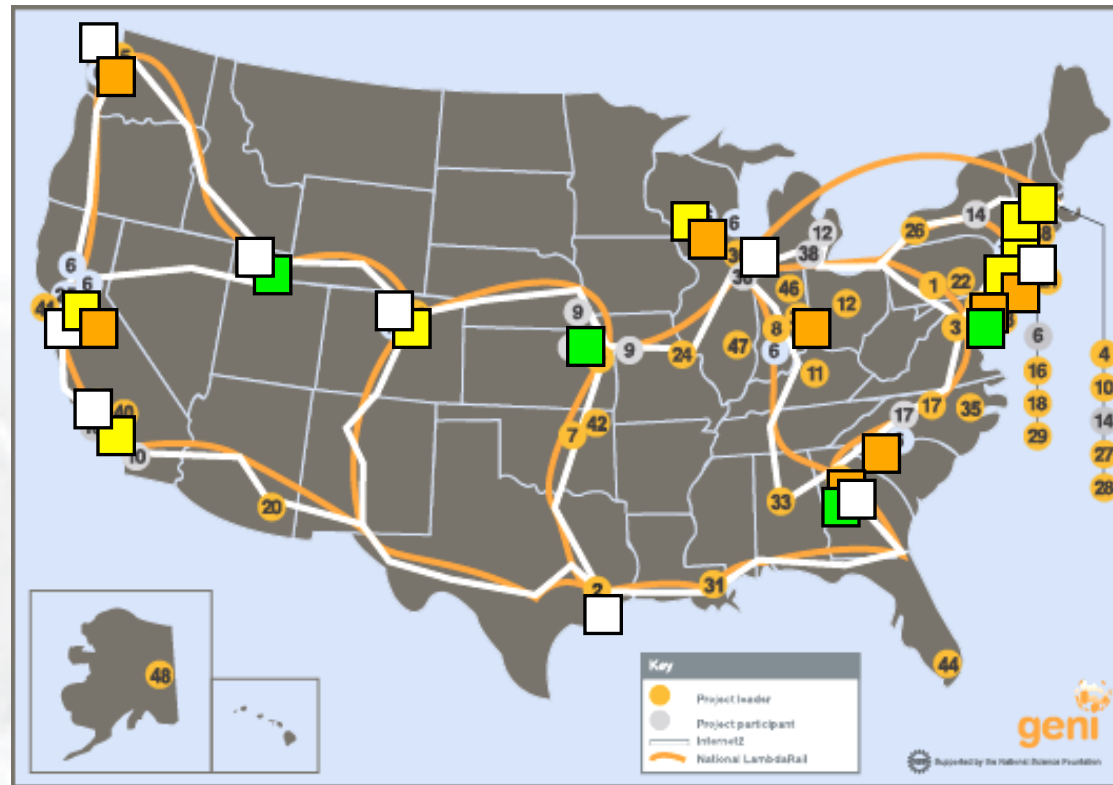
Salt Lake City  
Kansas City  
Washington, DC  
Atlanta

## WiMAX

Stanford  
UCLA  
UC Boulder  
Wisconsin  
Rutgers  
NYU Polytech  
UMass  
Columbia

## OpenFlow Backbones


Seattle  
Salt Lake City  
Sunnyvale  
Denver  
New York City  
Houston  
Chicago  
Los Angeles  
Atlanta



HP ProCurve 5400 Switch




Juniper MX240 Ethernet  
Services Router



NEC WiMAX Base Station



Toroki LightSwitch 4810

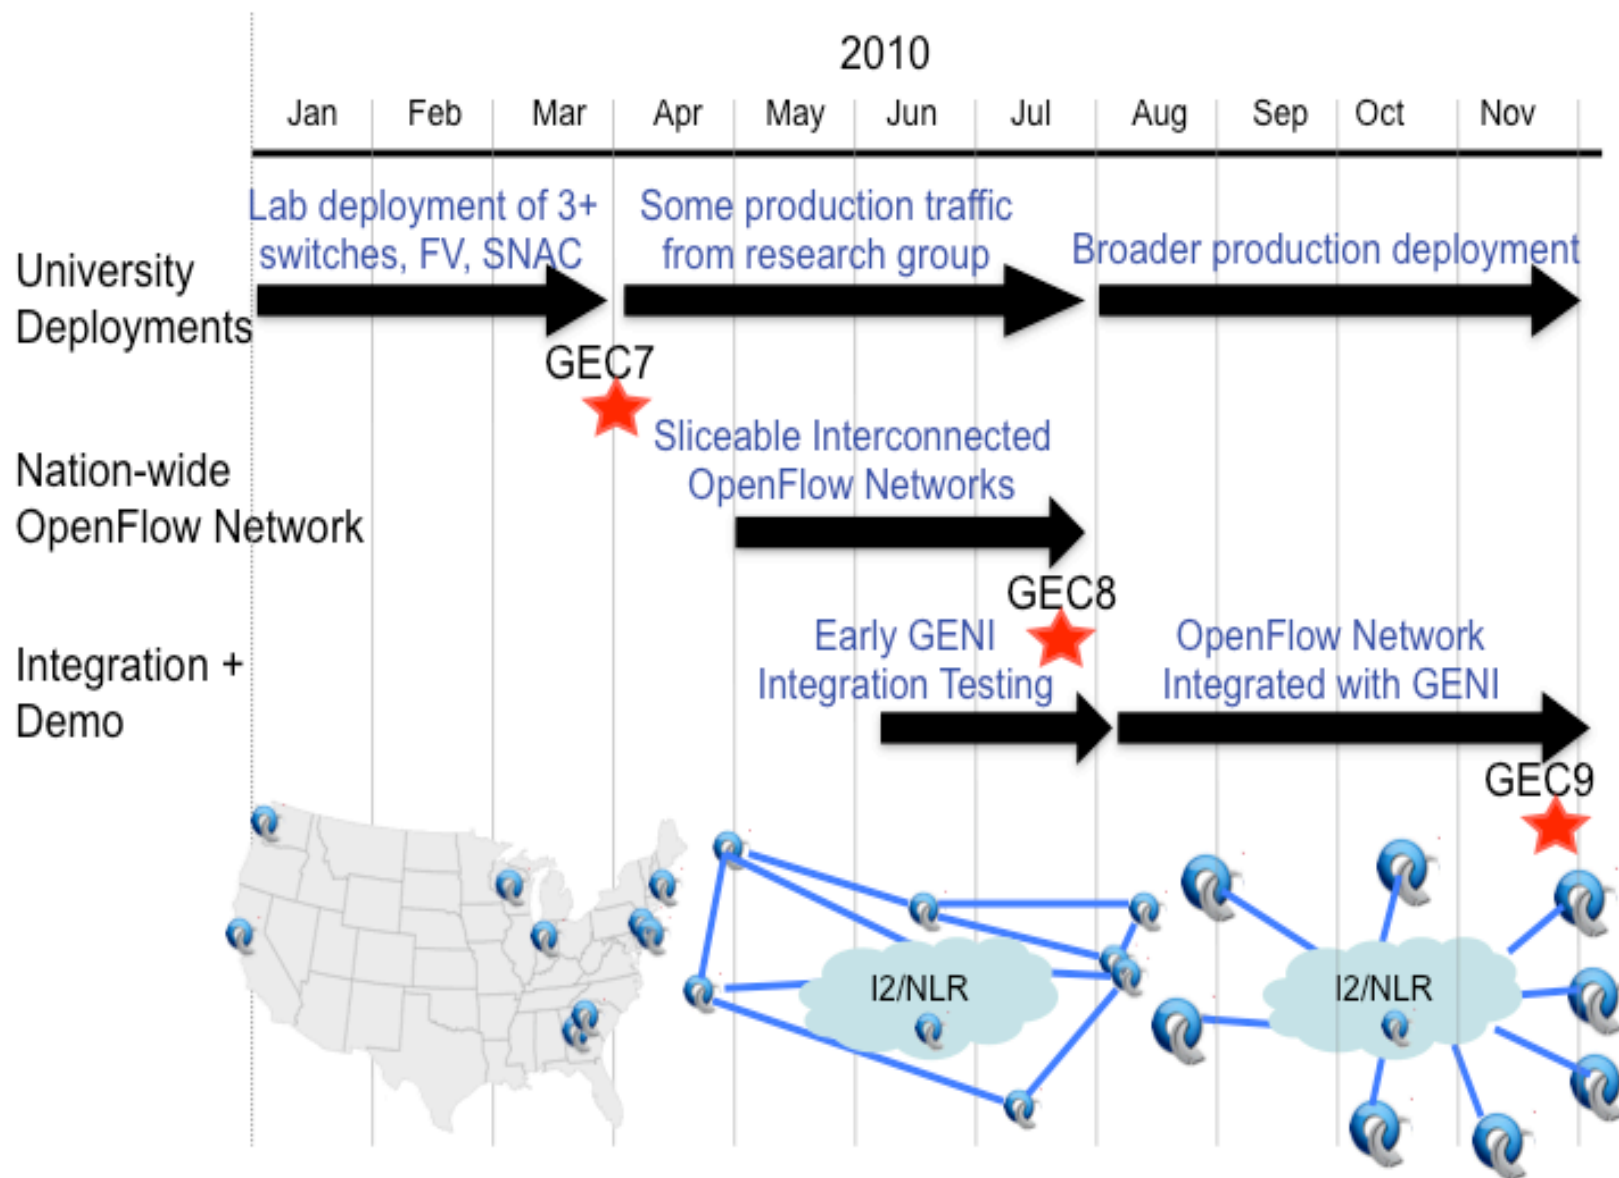


Arista 7124S Switch

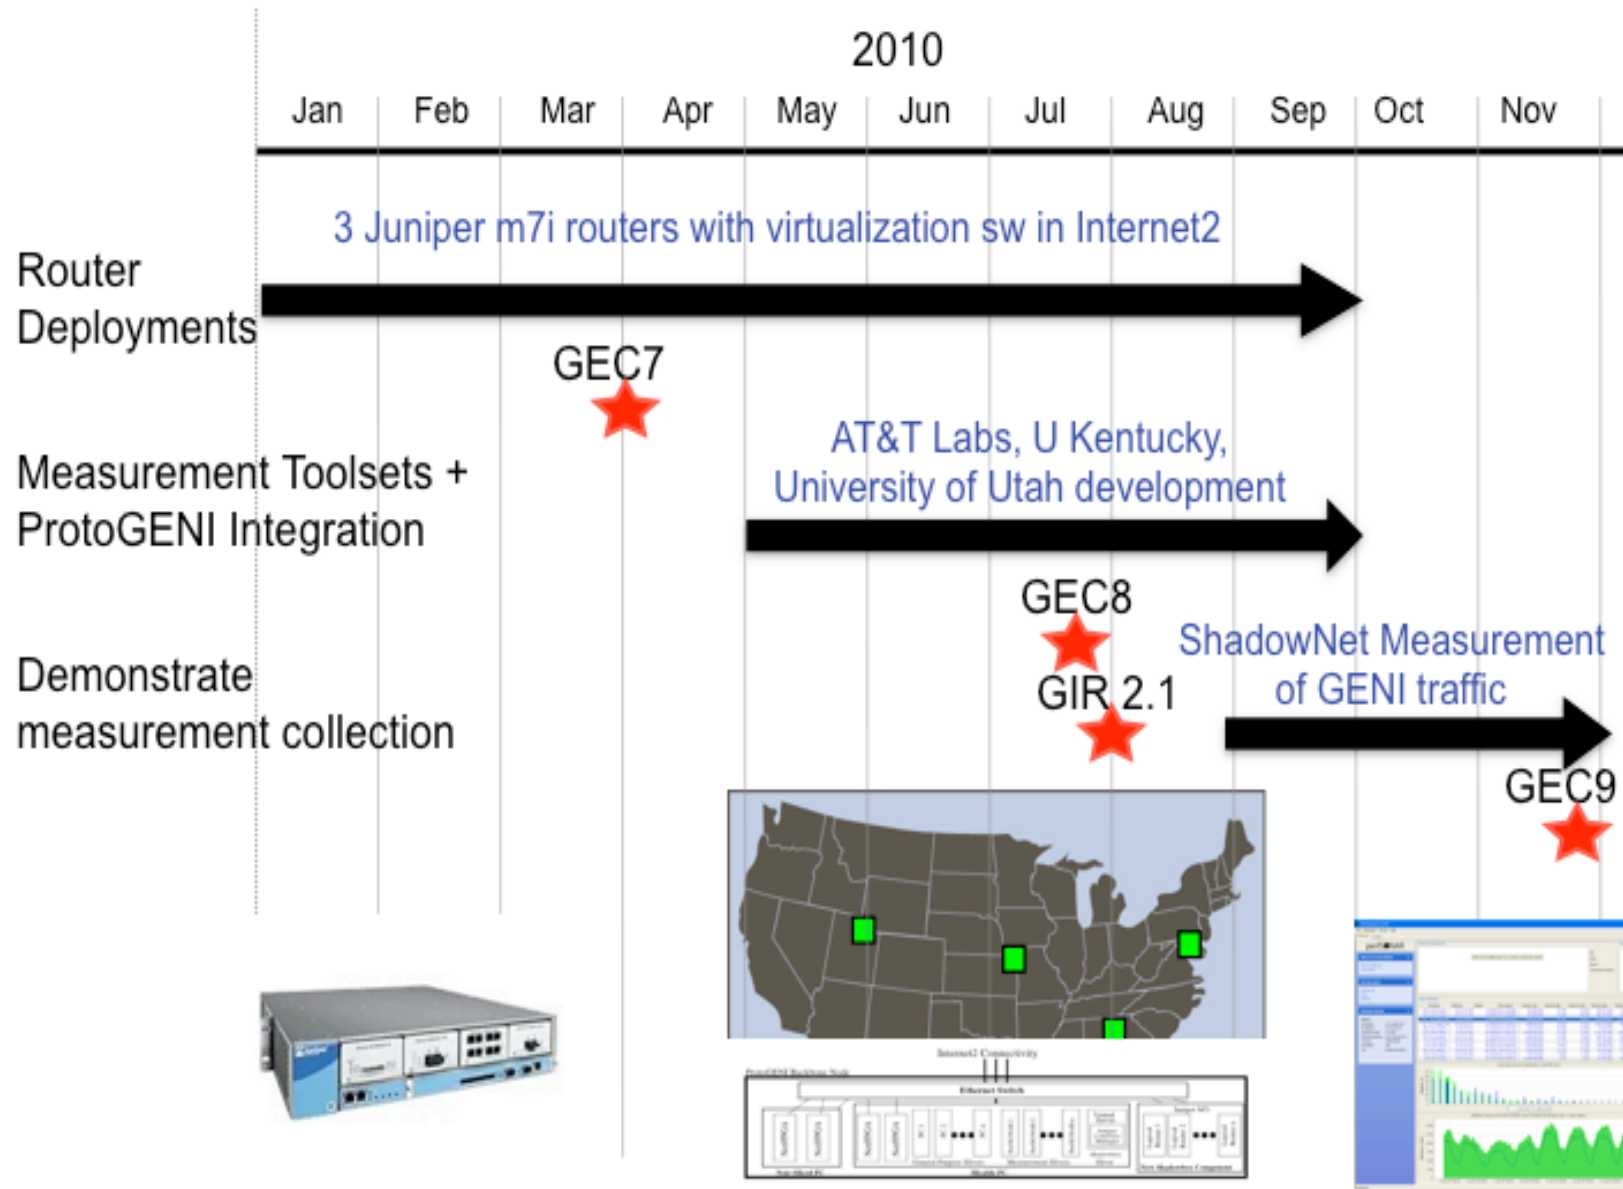


NEC IP8800 Ethernet Switch

# OpenFlow Deployment Roadmap

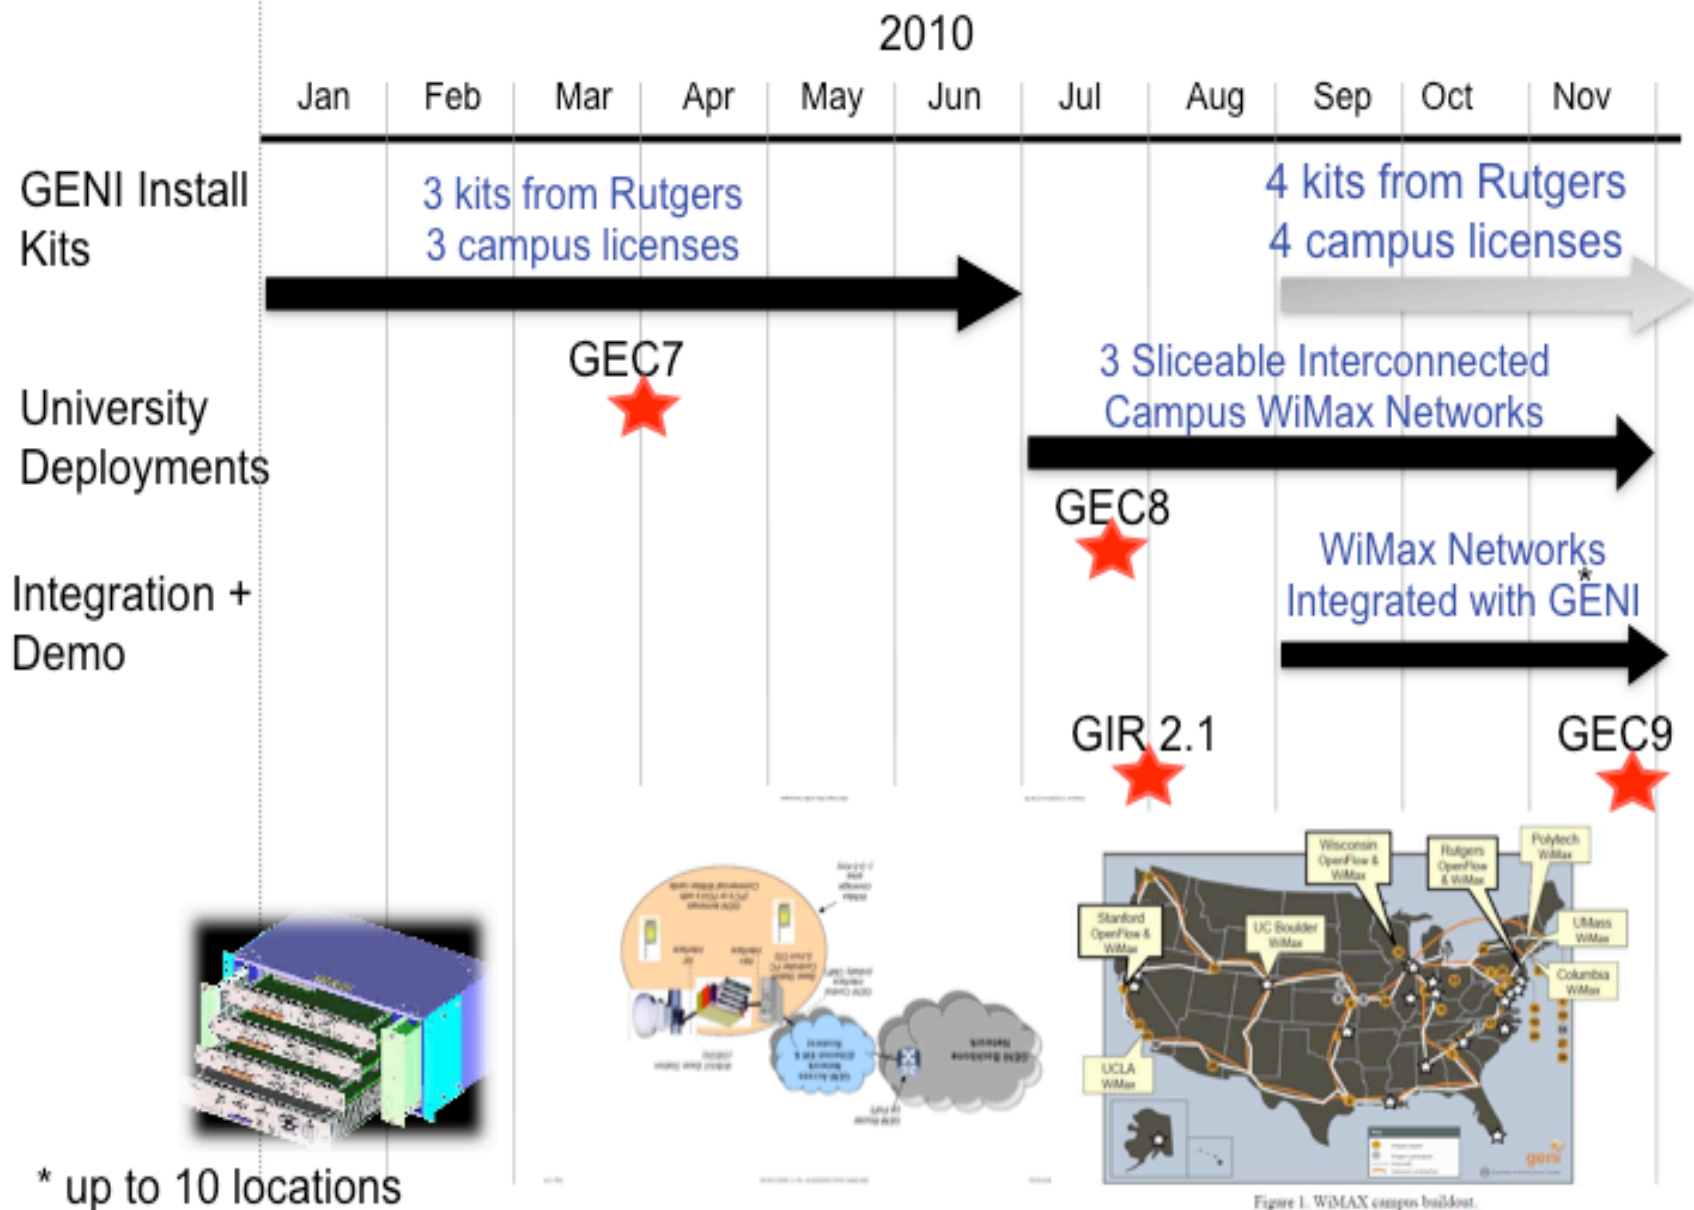


# ShadowNet Deployment Roadmap





# WiMAX Deployment Roadmap



- To succeed as a virtual laboratory, GENI must support a wide variety of experiments.
- Early GENI goals include support for
  - Repeatable and/or “in the wild” behavior
  - Large-scale infrastructure
  - Novel network architecture
  - Deep programmability
  - Programmable switches and routers
  - Opt-in users
- These capabilities are rapidly taking shape
  - GENI will continue to increase in capability, scale, and interoperability

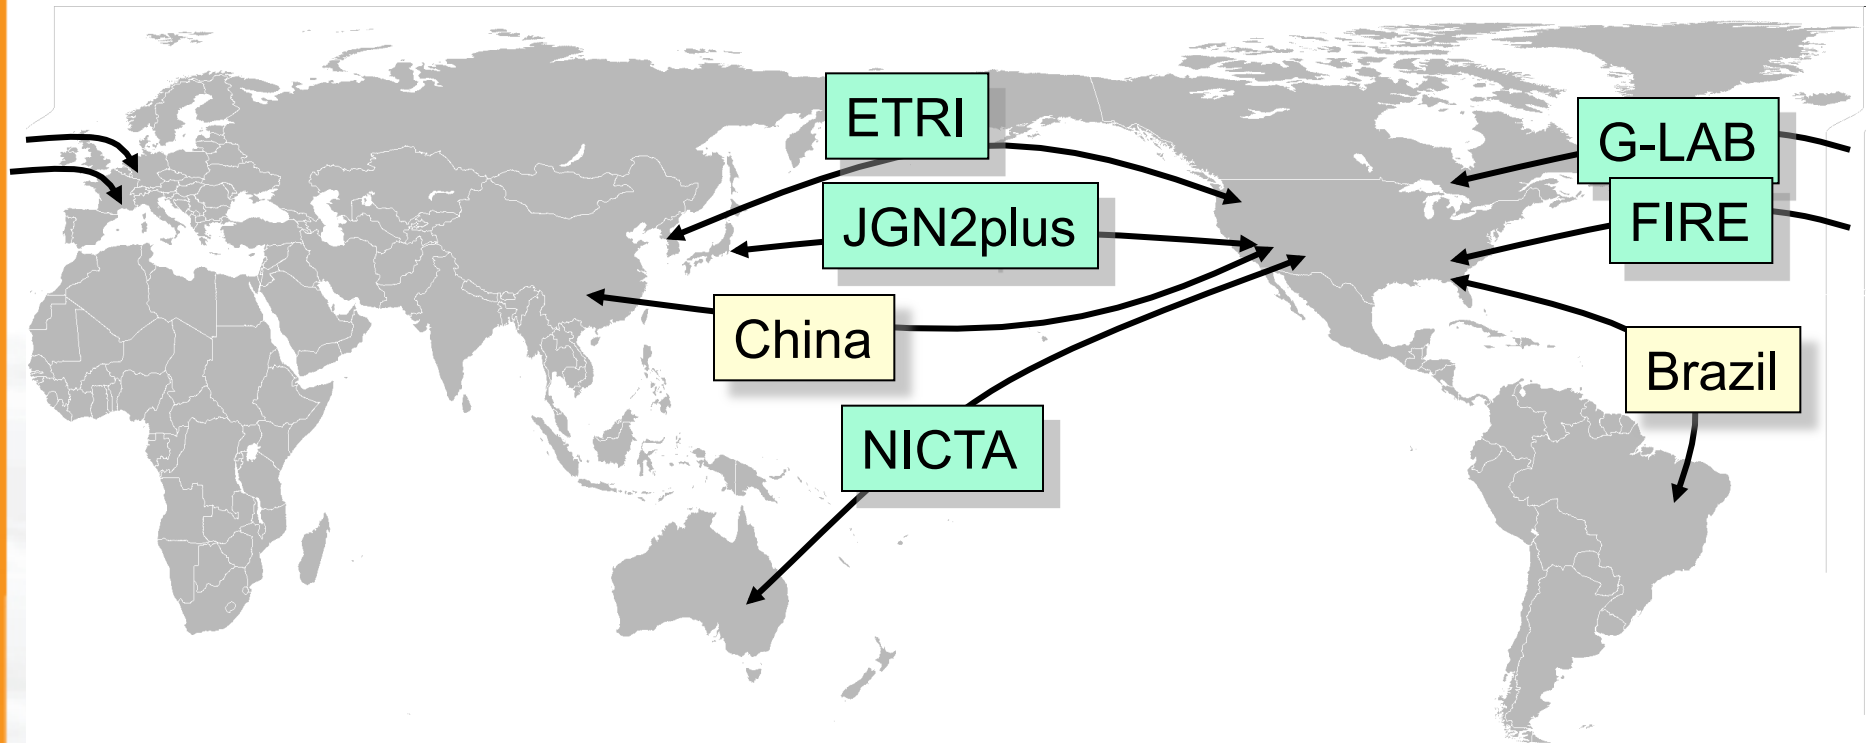
- Today:
  - GENI backbones connect ProtoGENI, SPP, BEN
  - Other resources connect via IP using tunnels as needed
  - Four control frameworks
  - GPO-assisted, manual stitching of VLANs
  - Limited tools for discovery, management, measurement
- By Fall 2010
  - OpenFlow & WiMax campuses
  - Interoperability between PlanetLab and ProtoGENI control frameworks
  - Improved tools
- In 2011
  - Prototype I&M system
  - Common control framework API
  - End-users

- GENI – Exploring future internets at scale
- GENI's OpenFlow status and plans
  - GENI Spiral 2
  - Meso-scale buildout
  - Starting experimentation
- Looking ahead

- Solicitation areas
  - 1. Aggressively grow meso-scale build** (next slide)
    1. Enhanced regional & backbone buildouts
    2. More WiMAX sites
    3. New “GENI Racks” (eg rack of PCs with OpenFlow switch)
  - 2. GENI Instrumentation system** (build & deploy)
  - 3. Experiment support / training / education & curriculum development**
- Solicitation document: see [www.geni.net](http://www.geni.net)
- Proposal deadline: August 20, 2010



These are exciting times  
all around the world!



The GENI project is actively collaborating with peer efforts outside the US, based on equality and arising from direct, “researcher to researcher” collaborations.

# GENI Engineering Conferences

Meet every 4 months to review progress together

- **9th meeting, open to all:  
November 2-4, 2010, Washington, DC**
  - Team meetings, integrated demos, Working Group meetings
  - Also discuss GPO solicitation, how to submit a proposal, evaluation process & criteria, how much money, etc.
  - **Travel grants** to US academics for participant diversity
- **Subsequent Meetings, open to all who fit in the room**
  - Held at regular 4-month periods
  - Held on / near university campuses (volunteers?)
  - All GPO-funded teams required to participate
  - Systematic, open review of each Working Group status (all documents and prototypes / trials / etc.)
  - Also time for Working Groups to meet face-to-face
  - Discussion will provide input to subsequent spiral goals

# Thanks!

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

# Backup Slides

# Useful GENI Links & Documents

- GENI wiki: <http://groups.geni.net>
- GENI System Overview (describes conceptual design):  
<http://groups.geni.net/geni/attachment/wiki/GeniSysOvrwv/GENISysOvrwv092908.pdf>
- GENI Spiral Two Overview (summarizes current goals and projects):  
<http://groups.geni.net/geni/attachment/wiki/SpiralTwo/GENIS2Ovrwv060310.pdf>
- Spiral Two Project Pages: <http://groups.geni.net/geni/wiki/SpiralTwo>
- Experimenter Page (summarizes available resources):  
<http://groups.geni.net/geni/wiki/GeniExperiments>



# GENI Concepts & Terminology

- **Researcher**
  - someone who wishes to run an experiment or service on GENI.
- **Clearinghouse**
  - A collection of trust anchors, identifying researchers and resources
  - A collection of operational services that facilitate the GENI control framework
    - Researcher account and resource utilization recordkeeping
    - Resource discovery services
    - Federation-wide policy implementation
    - Operations and management services
  - GENI currently includes multiple clearinghouses which are beginning to federate with each other.
- **Aggregate**
  - a collection of resources available for GENI researchers under common ownership and administration
- **Aggregate Manager**
  - The entity responsible for resource discovery, experimenter authorization, resource allocation, and coarse control at an aggregate
  - Exports a standard interface, *the GENI Aggregate API*
- **End-User**
  - A principal participating in GENI who is not a GENI researcher
  - End-users may generate traffic that passes through GENI resources or be measured by GENI experiments
  - End-users may also contribute computational or networking resources for GENI researchers to use, e.g., Million-node GENI

## GENI Concepts & Terminology (2)

- **Sliver**
  - The resources in an aggregate allocated to an experiment
  - May be allocated virtually or physically
- **Slice**
  - A collection of slivers
  - The primary abstraction for accounting and accountability
  - The basis for resource revocation (i.e., shutdown).
  - Slice = slivers + authorized researchers
- **RSpec**
  - Resource specification
  - Represents all GENI resources that can be bound to a sliver within an aggregate.
  - Describes both the resources available, advertised or allocated at a component and the relationships between those resources, and perhaps other resources.
- **Credentials**
  - Authenticated documents which describe privileges held by a principal and are cryptographically signed
  - Currently, the format is an XML structure containing X.509 certificates issued by a Clearinghouse
- **Clusters**
  - An organizational construct used for rapid integration of GENI resources with a control framework
  - GENI currently has 4 clusters around the PlanetLab, ProtoGENI, ORCA, and ORBIT control frameworks
  - The importance of clusters for interoperability will decline as common APIs and tools are sActiveported

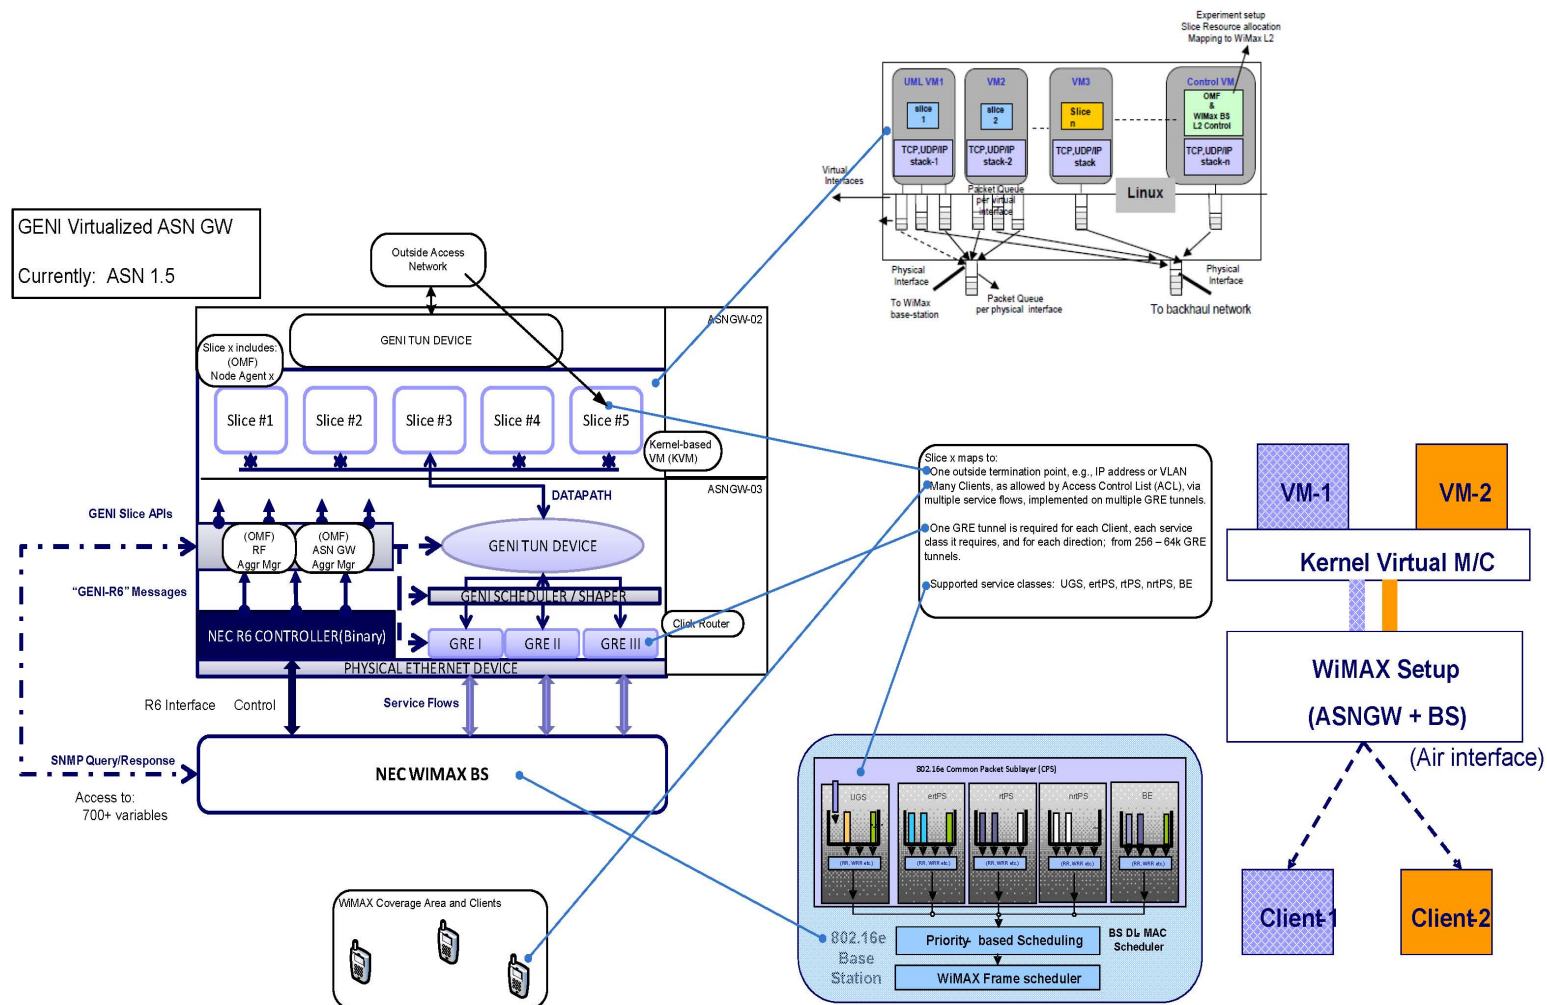
# Programmable WiMax Base Stations

# WiMAX kit (software)

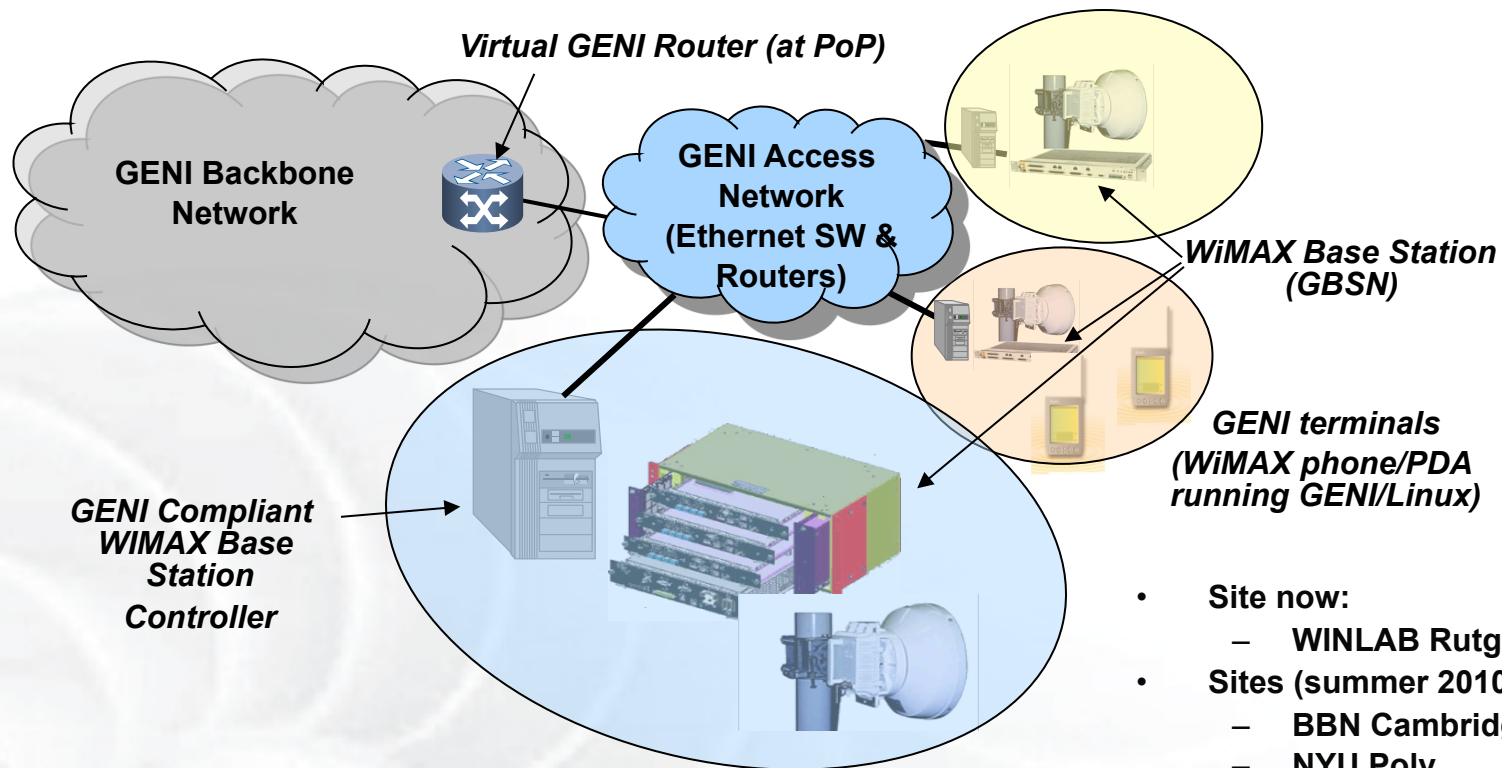
111309b WiMAXSystemOverview

GENI WiMAX Campus Deployment Kit: Fig 1-4) Virtualization

Page 6 of 9



# Programmable WiMax Base Stations



- **Site now:**
  - WINLAB Rutgers
- **Sites (summer 2010):**
  - BBN Cambridge
  - NYU Poly
- **Sites (late 2010):**
  - Columbia
  - UMass Amherst
  - Univ Wisconsin
  - Univ Colorado Boulder
  - UCLA




# WiMax Experiment Setup Steps

- Use OMF instance (GENI AM) at a site:
  - (1) Get account from site Admin
  - (2) Login to site
  - (3) Access other sites, as desired (later)
- Do basic OMF admin functions:
  - (1) Initialize grid services
  - (2) List all running slices
  - (3) Create your slice
- Use your slice:
  - (1) Configure and program slice
    - Add data path to GENI backbone network
  - (2) Start/Stop Slice
  - (3) Add Client
    - Registers a client with the slice
    - Currently adds default service flow settings for the client
    - Adds mapping to the datapath controller on ASN-GW
  - (4) Configure measurements with OML
  - (5) Conduct experiment

# Mock WiMax Experiment Sequence

- Mobile associates, gets added to default slice, starts UL traffic
- Slice user starts a new slice, adds the mobile to its slice
- Datapath switch from (Mobile – VM0) → (Mobile – VM1)

Default Slice (VM-0)



User Slice (VM-1)