GENI

Exploring Networks of the Future

Status update for the CCC
March 22, 2010

Chip Elliott
GENI Project Office
www.geni.net

Sponsored by the National Science Foundation

Outline

- GENI – Exploring future internets at scale
- Current status and plans: GENI Spiral 2
  - System integration and refinement
  - Meso-scale buildout
  - Starting experimentation
- GPO program activities
  - Retasking & reorganization
  - GENI Solicitation 3
- Wrap-up

What is GENI?

- GENI is a virtual laboratory for exploring future internets at scale.
- GENI creates major opportunities to *understand, innovate, and transform* global networks and their interactions with society.
- GENI opens up new areas of research at the frontiers of network science and engineering, and increases the opportunity for significant socio-economic impact.

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

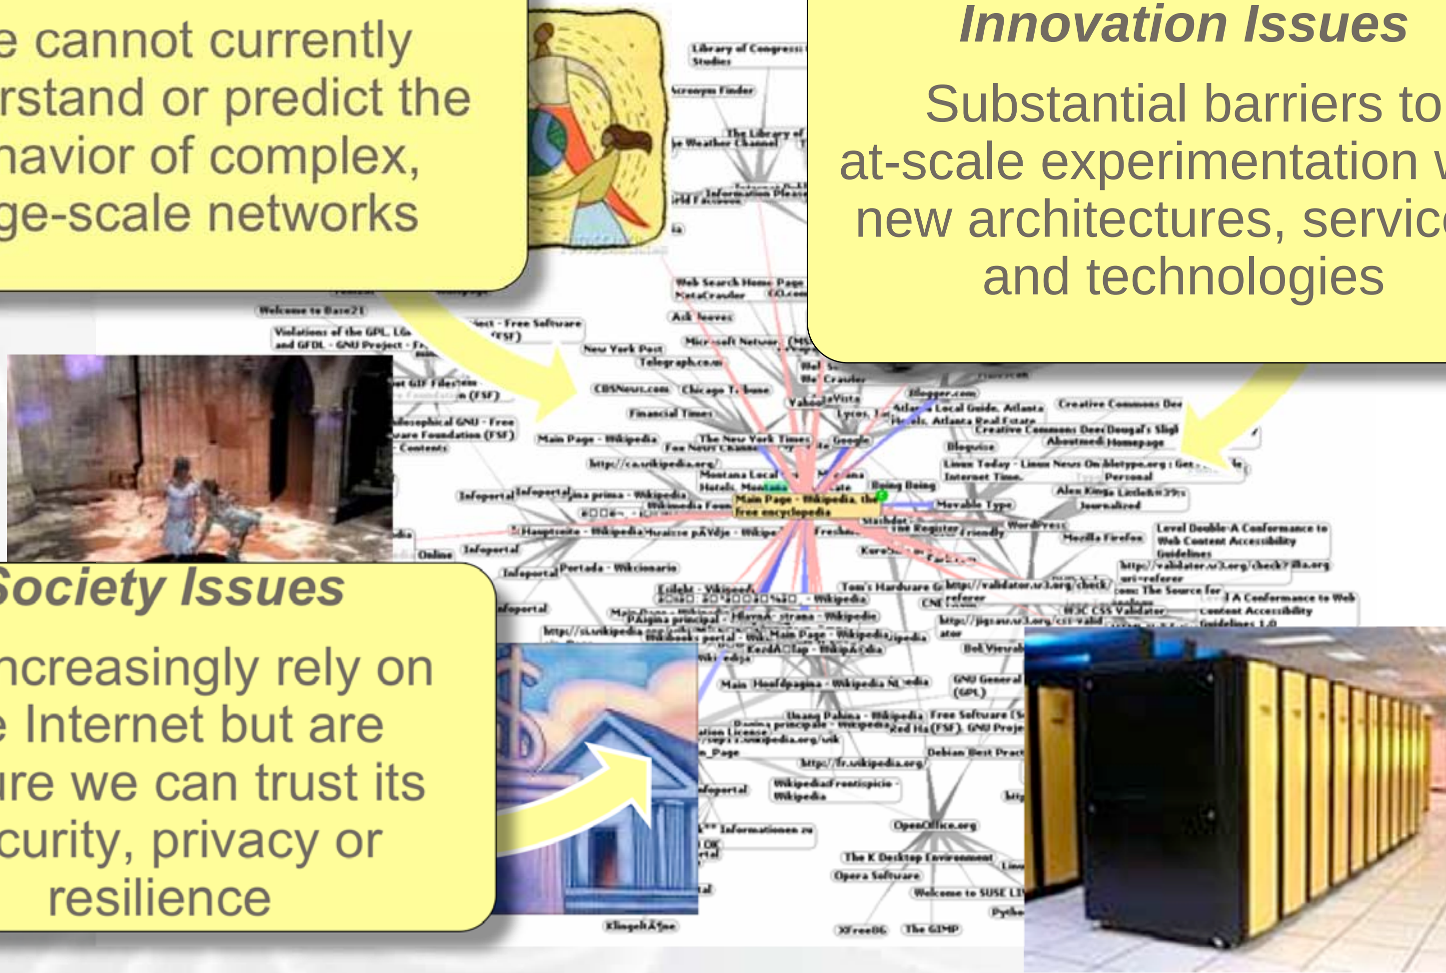

Society Issues

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

National Science Foundation
Network Science & Engineering (NetSE)

Science → Understand the complexity of large-scale networks → **Network science and engineering researchers**

- Understand emergent behaviors, local–global interactions, system failures and/or degradations
- Develop models that accurately predict and control network behaviors

Technology → Develop new architectures, exploiting new substrates → **Distributed systems and substrate researchers**

- Develop architectures for self-evolving, robust, manageable future networks
- Develop design principles for seamless mobility support
- Leverage optical and wireless substrates for reliability and performance
- Understand the fundamental potential and limitations of technology

Society → Enable new applications and new economies, while ensuring security and privacy → **Security, privacy, economics, AI, social science researchers**

- Design secure, survivable, persistent systems, especially when under attack
- Understand technical, economic and legal design trade-offs, enable privacy protection
- Explore AI-inspired and game-theoretic paradigms for resource and performance optimization

GENI Conceptual Design
Infrastructure to support at-scale experimentation

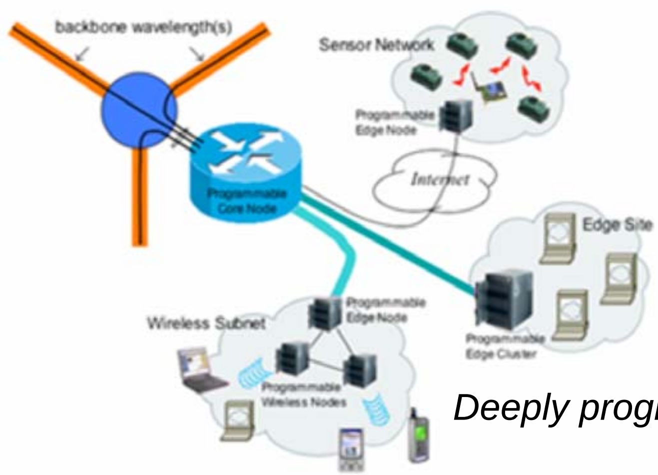

Deeply programmable

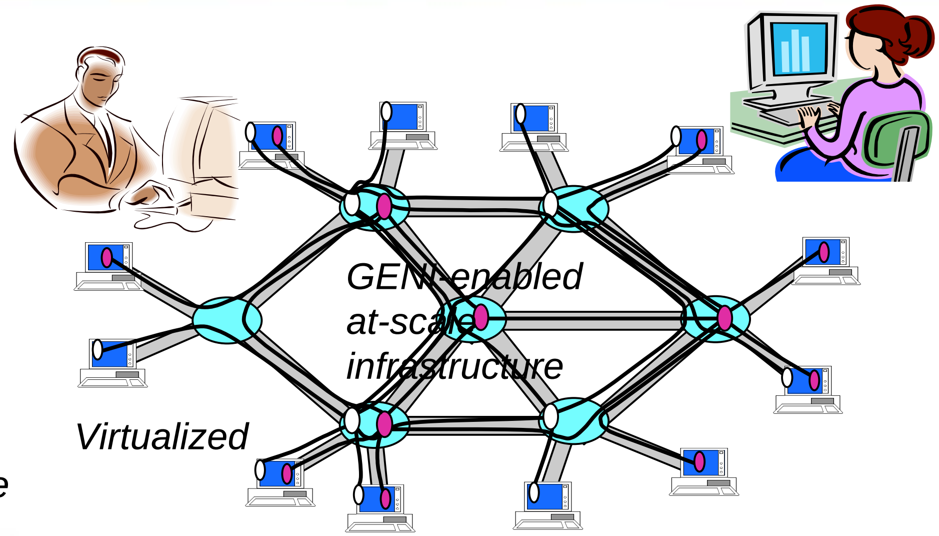

Programmable & federated, with end-to-end virtualized "slices"

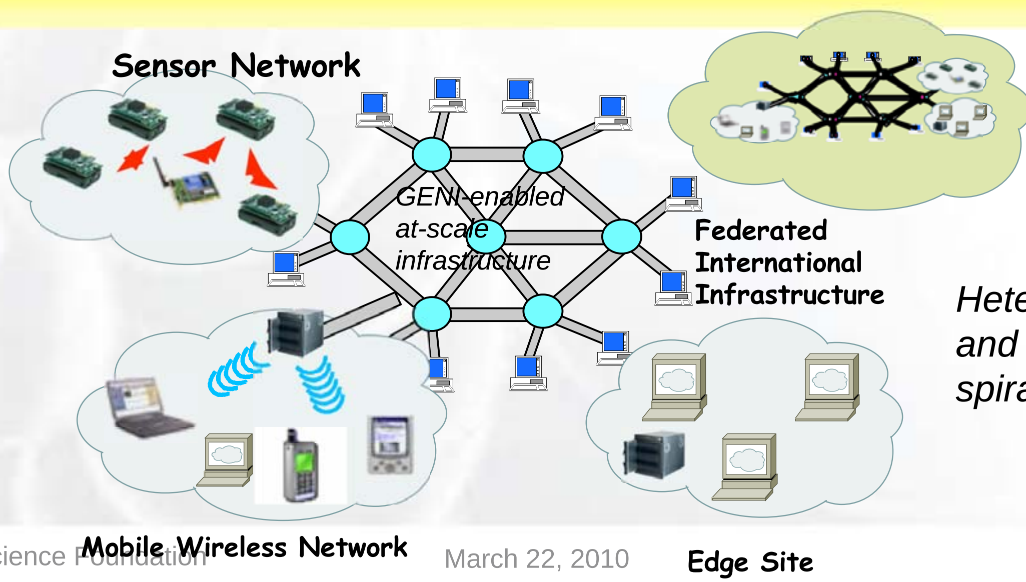

Heterogeneous,
and evolving over time via
spiral development

Enabling "at scale" experiments

- **GENI is enabling two classes of "at scale" experiments:**
  - **Controlled and repeatable** experiments, to help improve scientific understanding of complex, large-scale networks; and
  - **"In the wild" trials** of services that piggyback or connect to today's Internet and engage large numbers of participants.
  - With instrumentation and data archival / analysis tools for both
- **How can we afford / build GENI at sufficient scale?**
  - Clearly infeasible to build research testbed "as big as the Internet"
  - Therefore we are "GENI-enabling" testbeds, commercial equipment, campuses, regional and backbone networks
  - Key strategy for building an at-scale suite of infrastructure

Outline

- GENI – Exploring future internets at scale
- Current status and plans
  - GENI Spiral 2
  - Meso-scale buildout
  - Starting experimentation
- GPO program activities
  - Retasking & reorganization
  - GENI Solicitation 3
- Wrap-up

Current status & plans

- "GENI is a virtual laboratory for exploring future internets at scale"
  - How are we doing?

- Rapid progress to date
  - GENI community appears highly energized and surprisingly happy
  - System architecture is taking shape via spiral development
  - Meso-scale build has considerable buy-in from PIs, campus CIOs, national backbones, regionals
  - We are executing plans for getting a number of research experiments started on the GENI suite
- What are the next steps (Spiral 3) ?
  - Converge upon interoperable control frameworks & tools
  - Aggressively grow the meso-scale build, adding "GENI racks"
  - Transition to "operations" to support large-scale, continuous experiments

Spiral Development

GENI grows through a well-structured, adaptive process

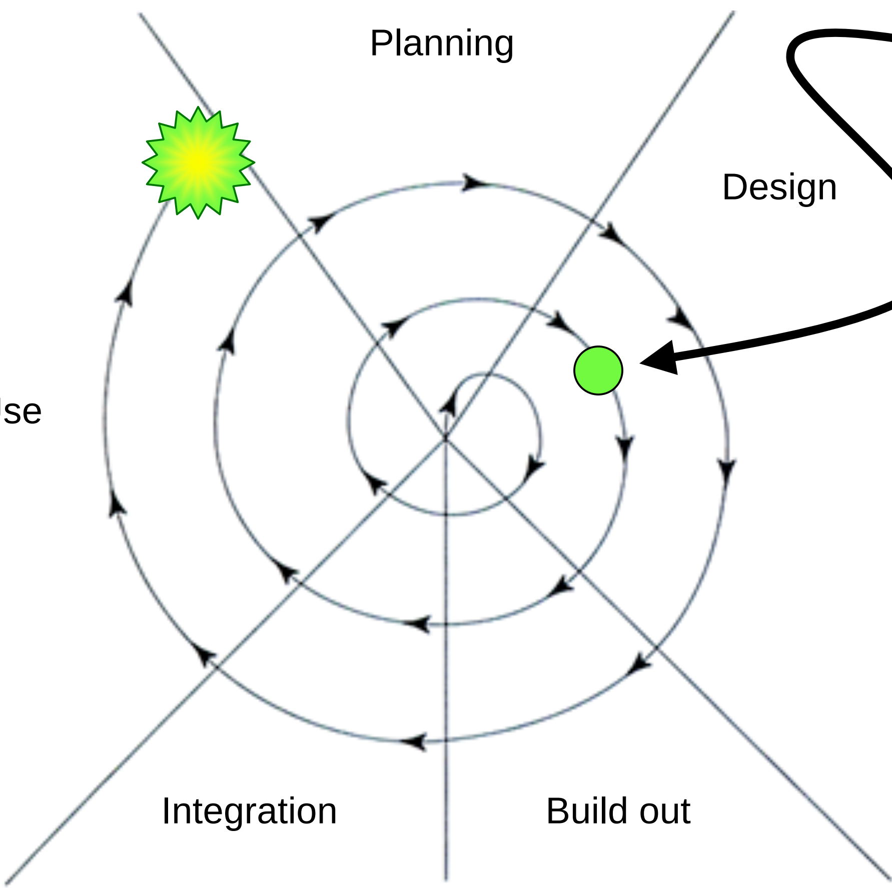

GENI Prototyping Plan

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.

Key goals for Spiral 2

- Overarching goal
  - Get real experiments up and running
- Technical emphases
  - Integration, particularly of the meso-scale prototype
  - Interoperability
  - Instrumentation
  - Identity management

Federation

GENI grows by “GENI-enabling” heterogeneous infrastructure

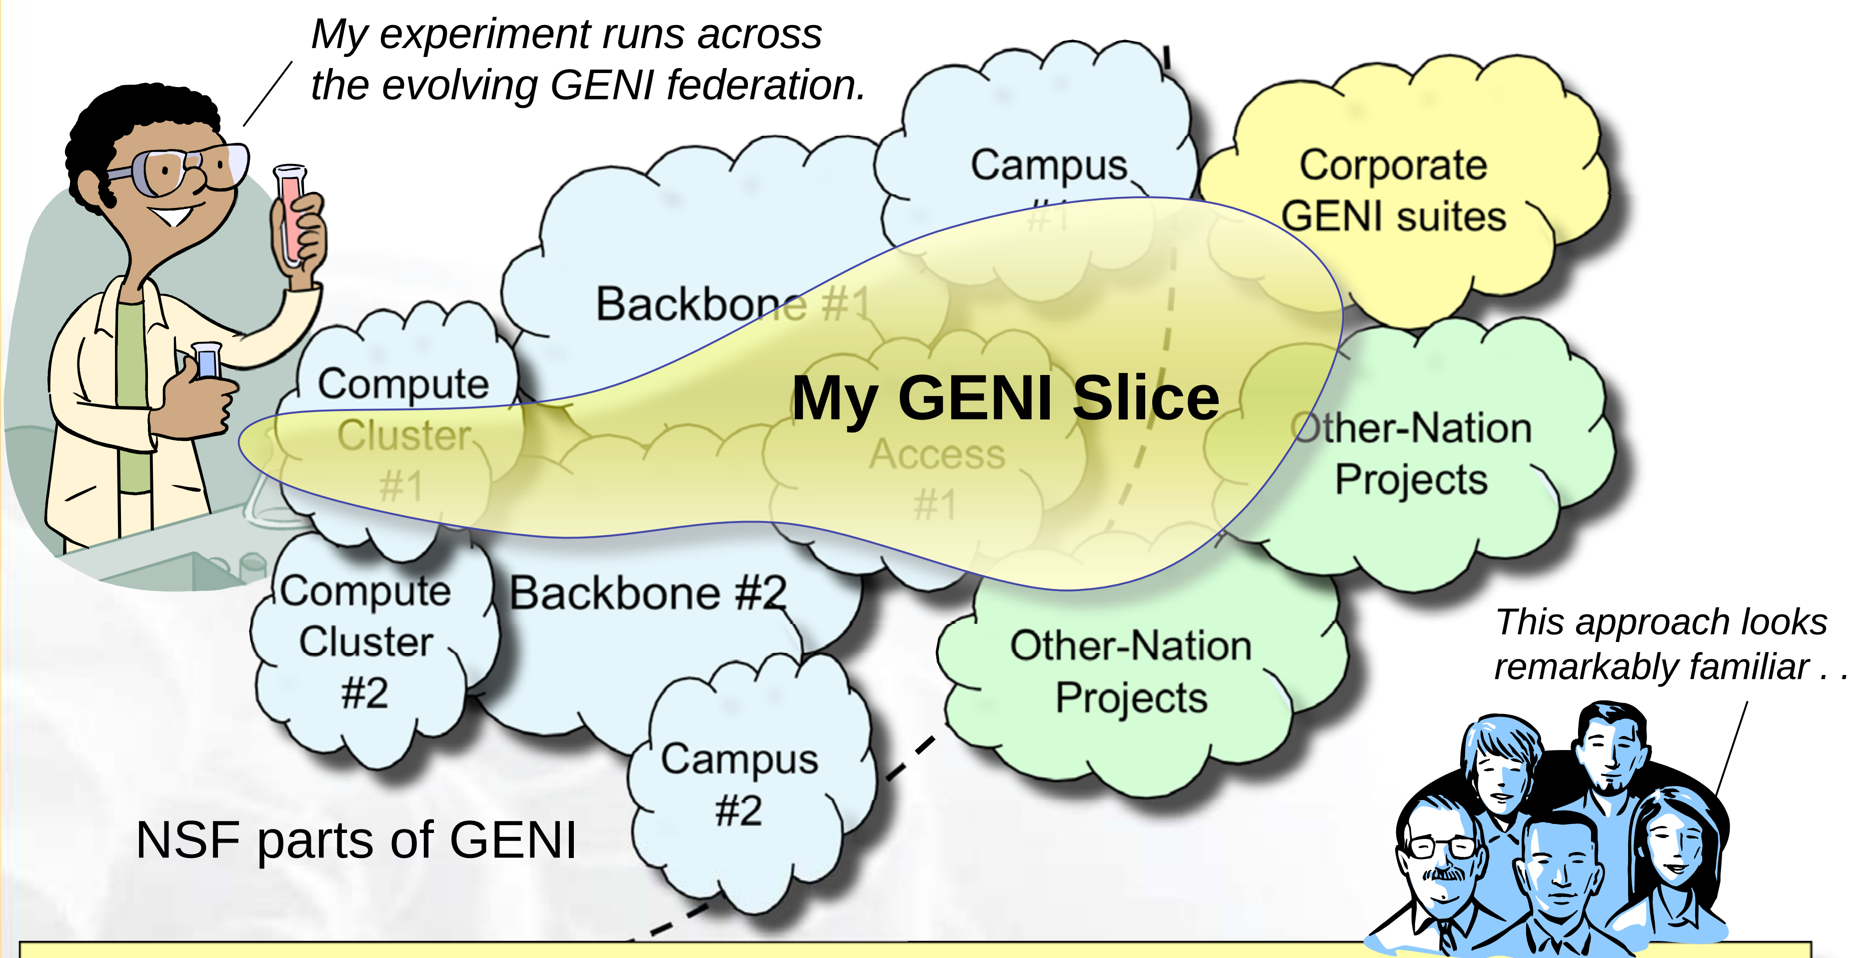

Goals: avoid technology “lock in,” add new technologies as they mature, and potentially grow quickly by incorporating existing infrastructure into the overall “GENI ecosystem”

Current GENI Status

GENI-enabling testbeds, campuses, and backbones

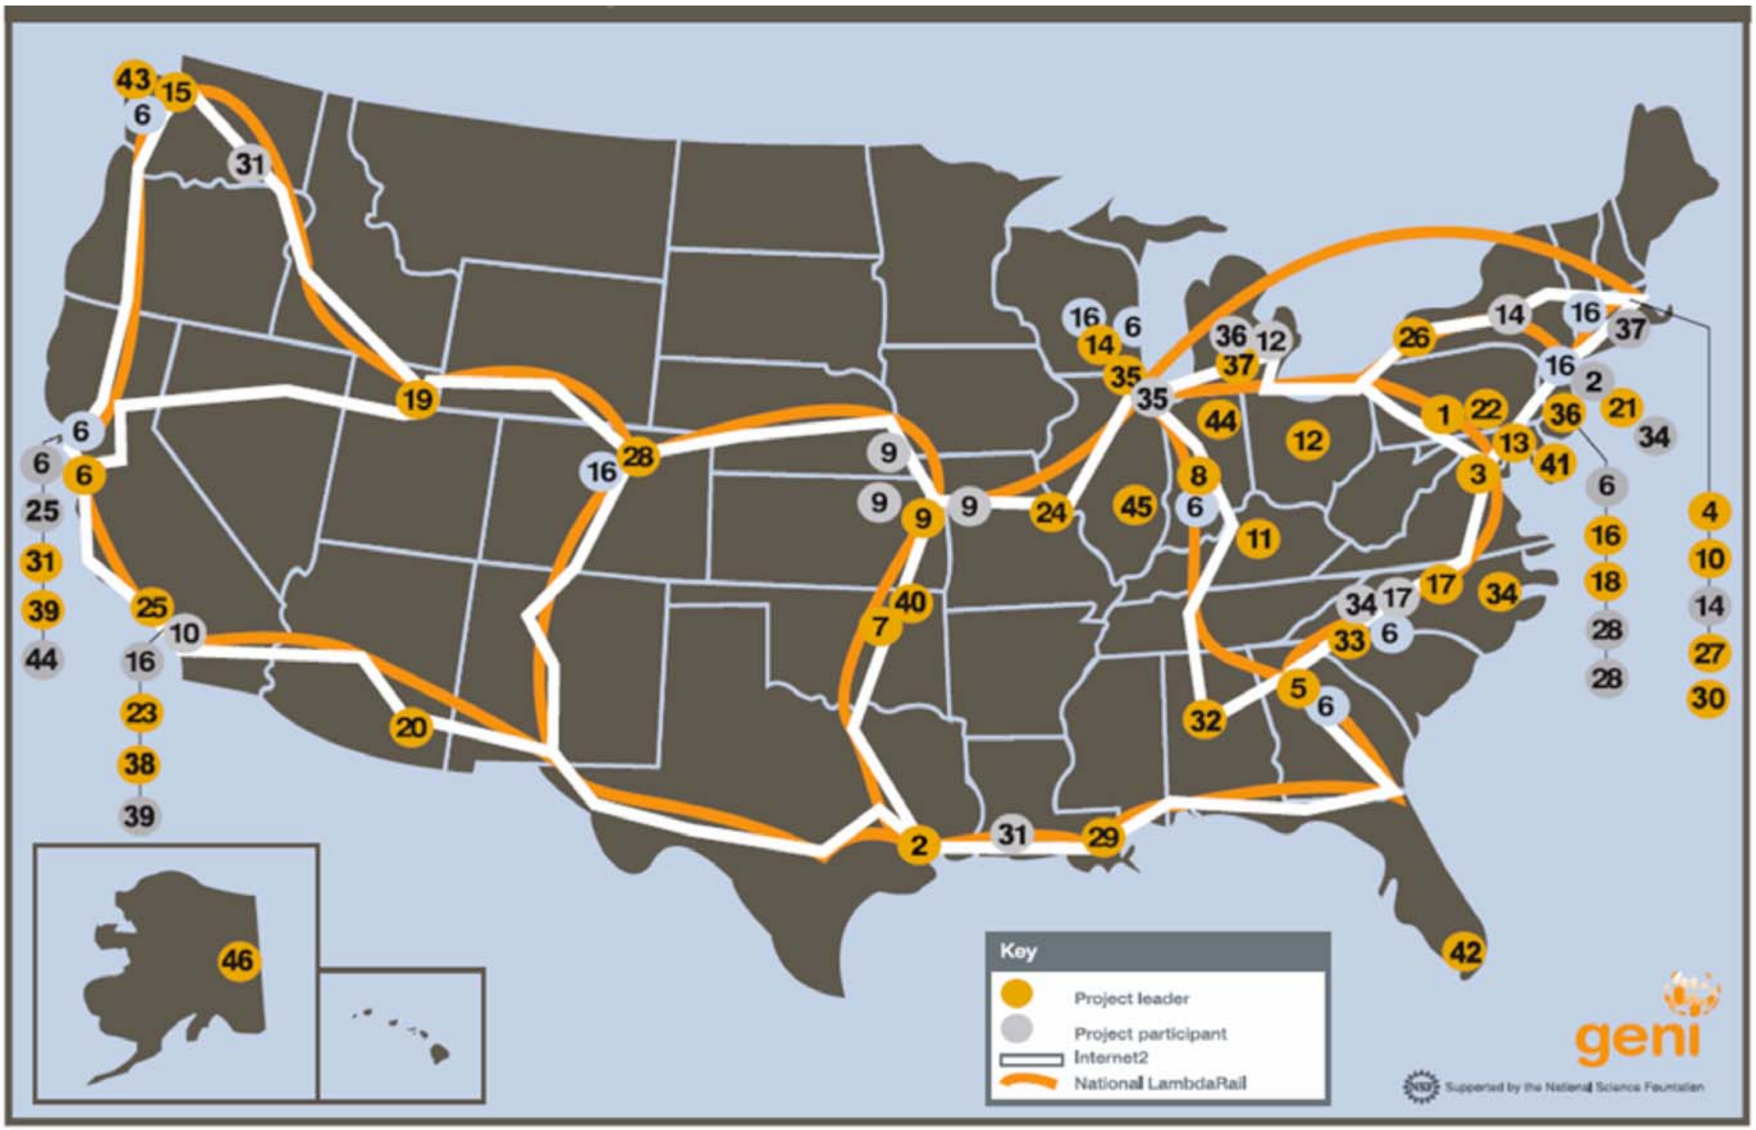

BEN
Breakable Experimental Network

ORBIT

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

ProtoGENI

UIC

UMASS LOWELL

OKLAHOMA STATE UNIVERSITY

THE QUILT

UPMC
PARISUNIVERSITAS

---

at&t

ARISTA

Battelle

ca

ciena

cisco

CNRI

FUJITSU

hp
invent

IBM

infinera

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

JUNIPER NETWORKS

nicira

NETRONOME

Microsoft

NEC

Qwest

Radio Technology Systems

SPARTA

Spiral 2 Control Framework Teams

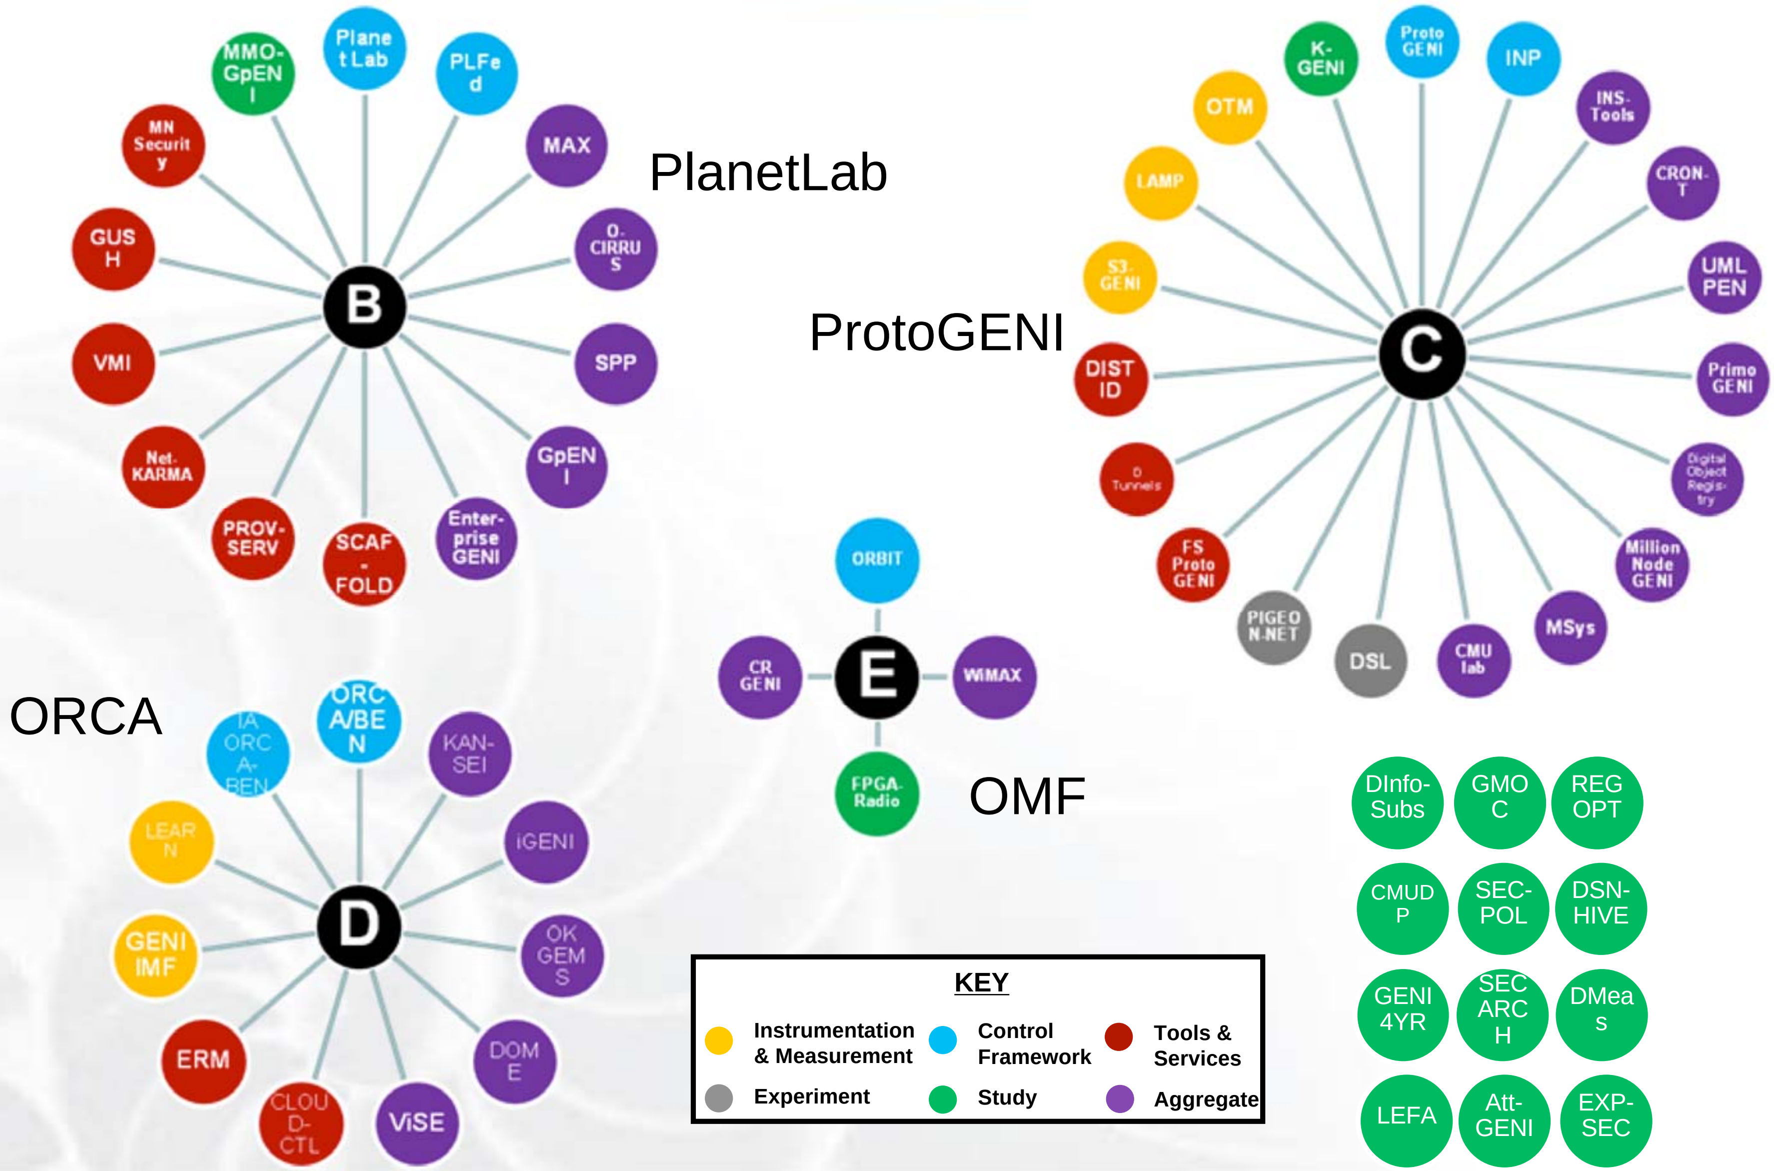

Infrastructure examples

DRAGON core nodes
Mid-Atlantic Crossroads

WAIL, U. Wisconsin-Madison

DieselNet, U. Mass Amherst

ViSE,
U. Mass Amherst

SPPs, Wash U.

ORBIT, Rutgers WINLAB

World-class expertise in GENI Partners

Internet2 and **National LambdaRail**

National LambdaRail

Up to 30 Gbps nondedicated bandwidth

Internet2

10 Gbps dedicated bandwidth

ProtoGENI & SPP

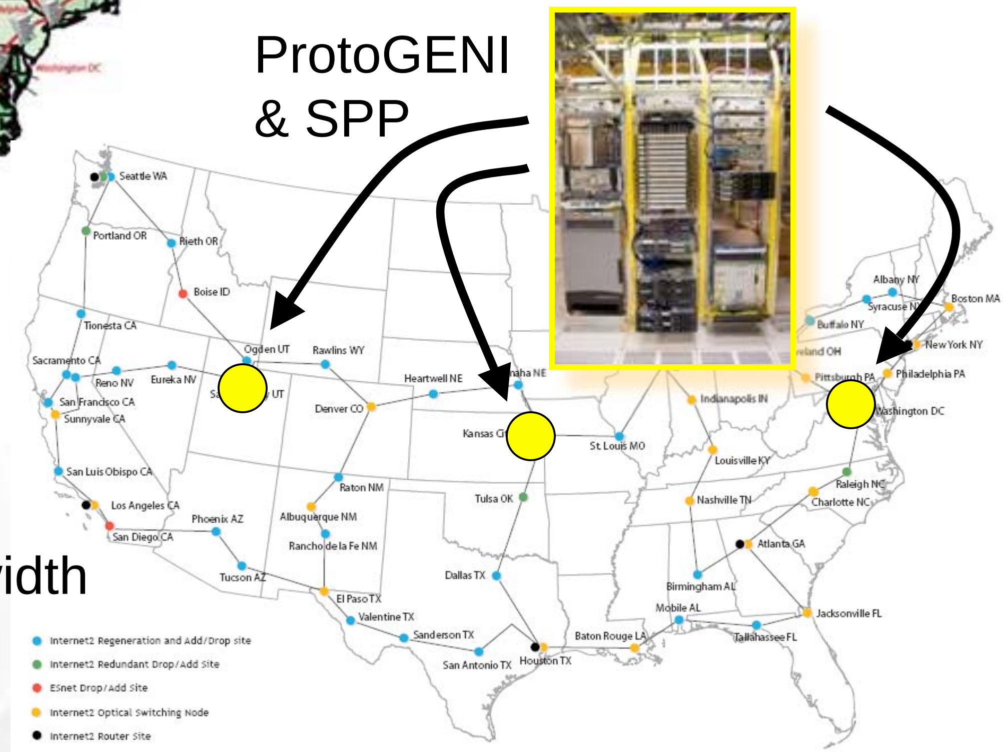

40 Gbps capacity for GENI prototyping on two national footprints to provide Layer 2 Ethernet VLANs as slices (IP or non-IP)

Photo courtesy of Chris Tracy

Outline

- GENI – Exploring future internets at scale
- Current status and plans
  - GENI Spiral 2
  - Meso-scale buildout
  - Starting experimentation
- GPO program activities
  - Retasking & reorganization
  - GENI Solicitation 3
- Wrap-up

Building the GENI Meso-scale Prototype

Current plans for locations & equipment

OpenFlow (orange)

Stanford
U Washington
Wisconsin
Indiana
Rutgers
Princeton
Clemson
Georgia Tech

ShadowNet (green)

Salt Lake City
Kansas City
DC
Atlanta

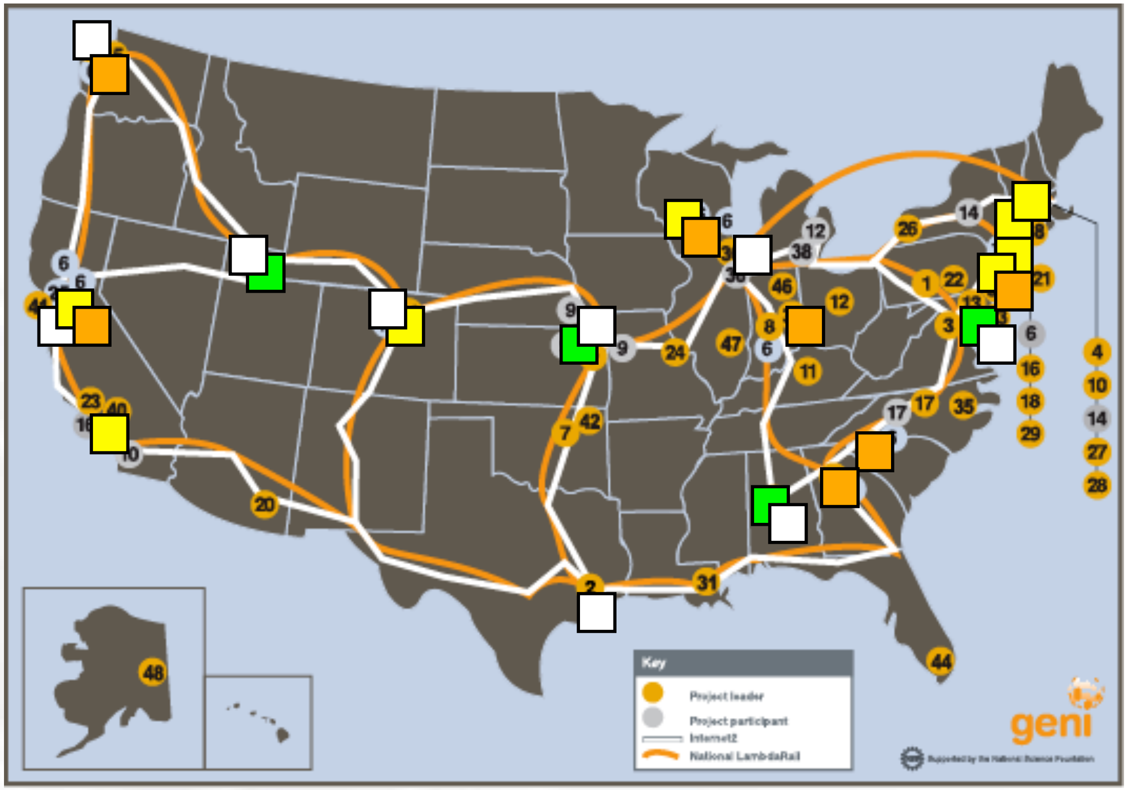

WiMAX (yellow)

Stanford
UCLA
UC Boulder
Wisconsin
Rutgers
Polytech
UMass
Columbia

OpenFlow Backbones (white)

Seattle
Salt Lake City
Sunnyvale
Denver
Kansas City
Houston
Chicago
DC
Atlanta

HP ProCurve 5400 Switch

Juniper MX240 Ethernet Services Router

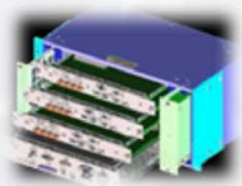

NEC WiMAX Base Station

Cisco 6509 Switch

Arista 7124S Switch

NEC IP8800 Ethernet Switch

OpenFlow Deployment Roadmap

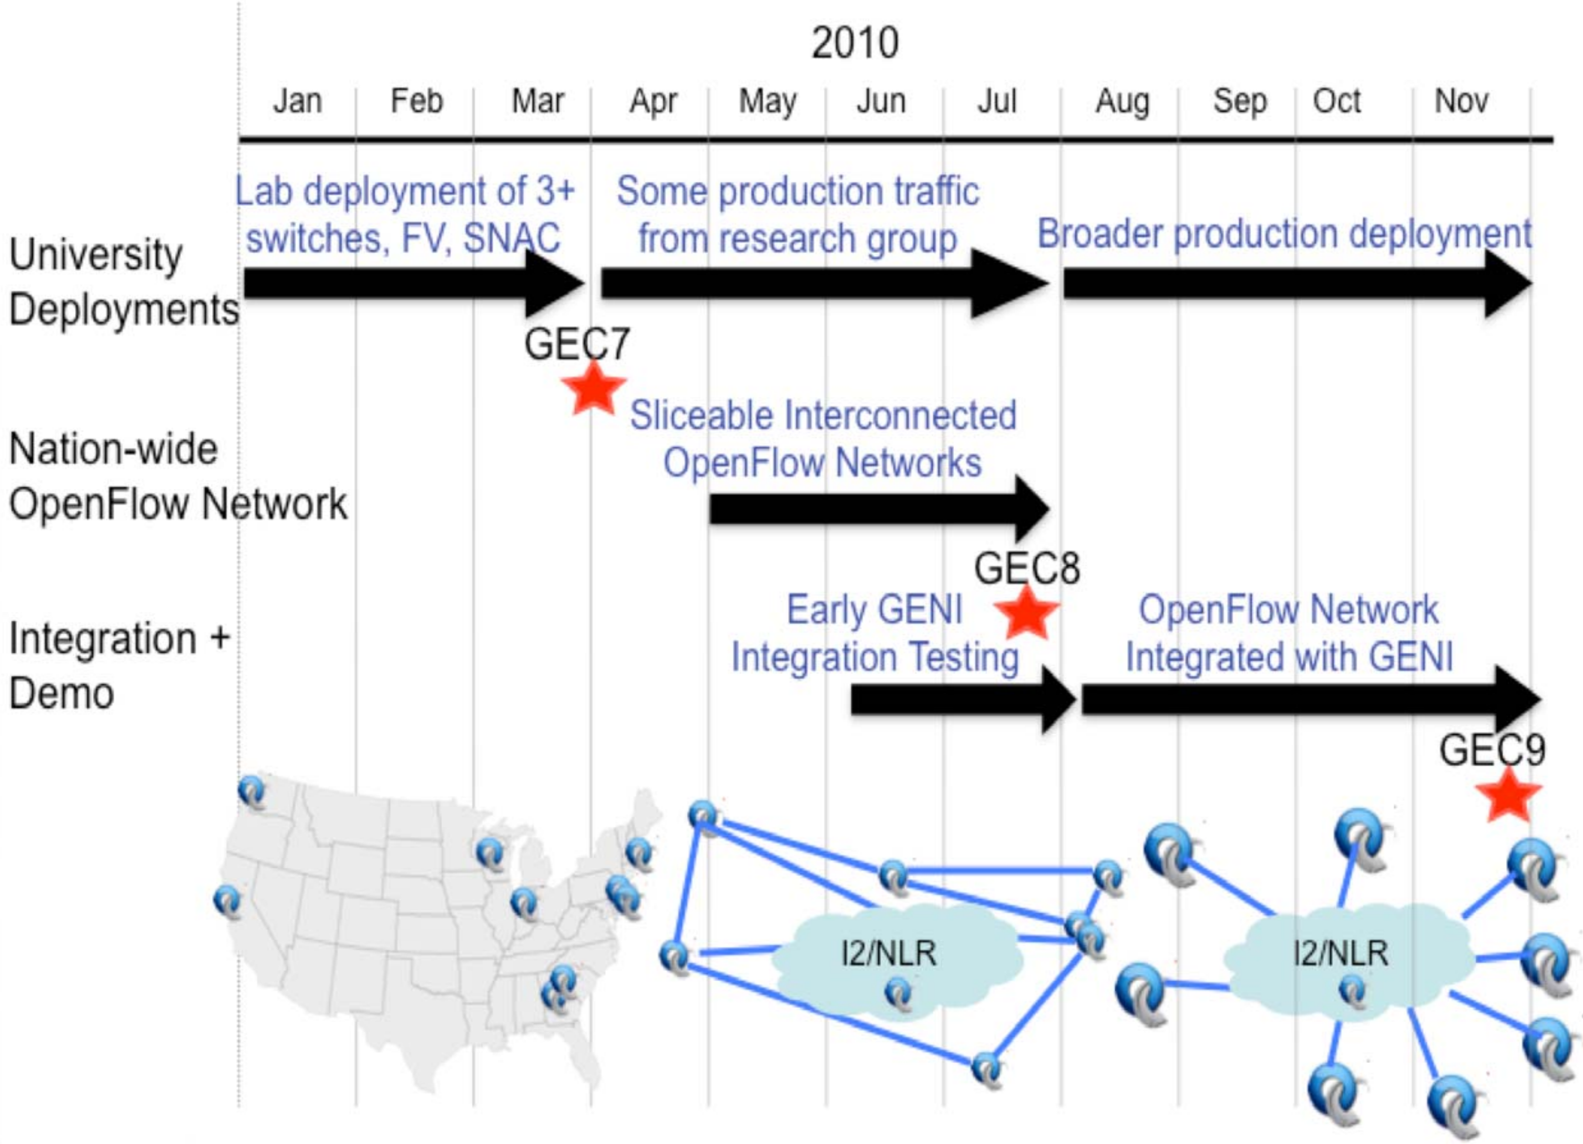

WiMAX Deployment Roadmap

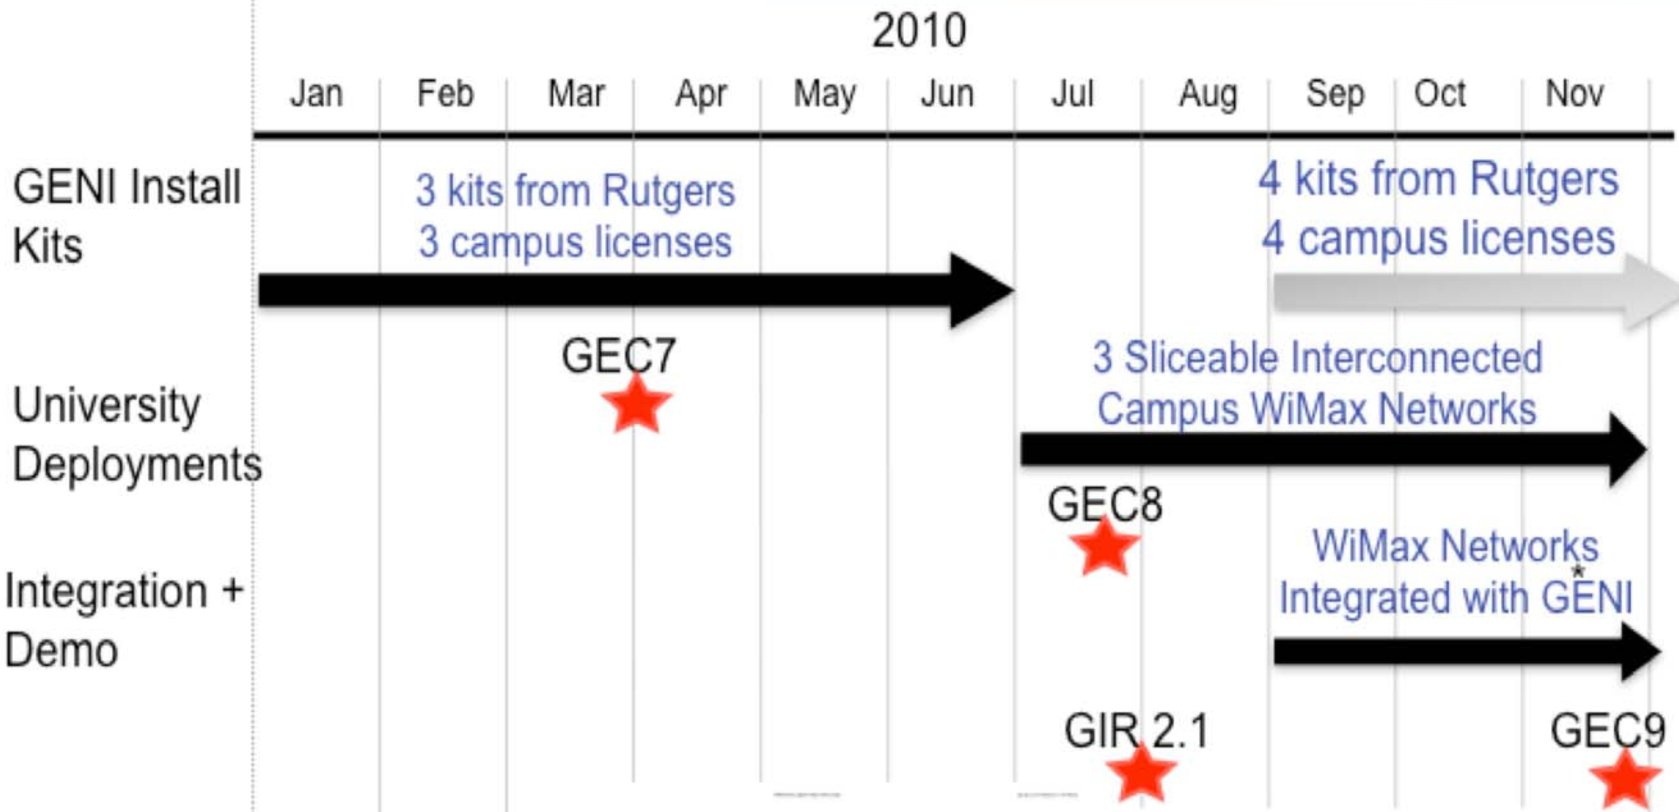

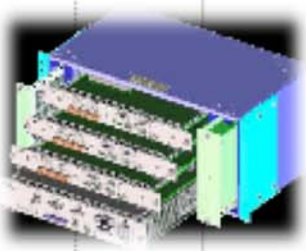

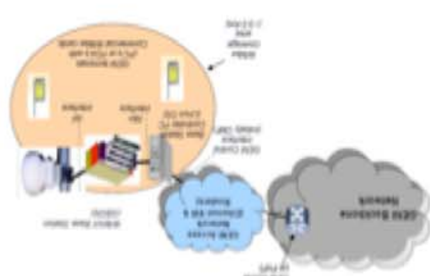

Figure 1. WiMAX campus buildout.

* up to 10 locations

ShadowNet Deployment Roadmap

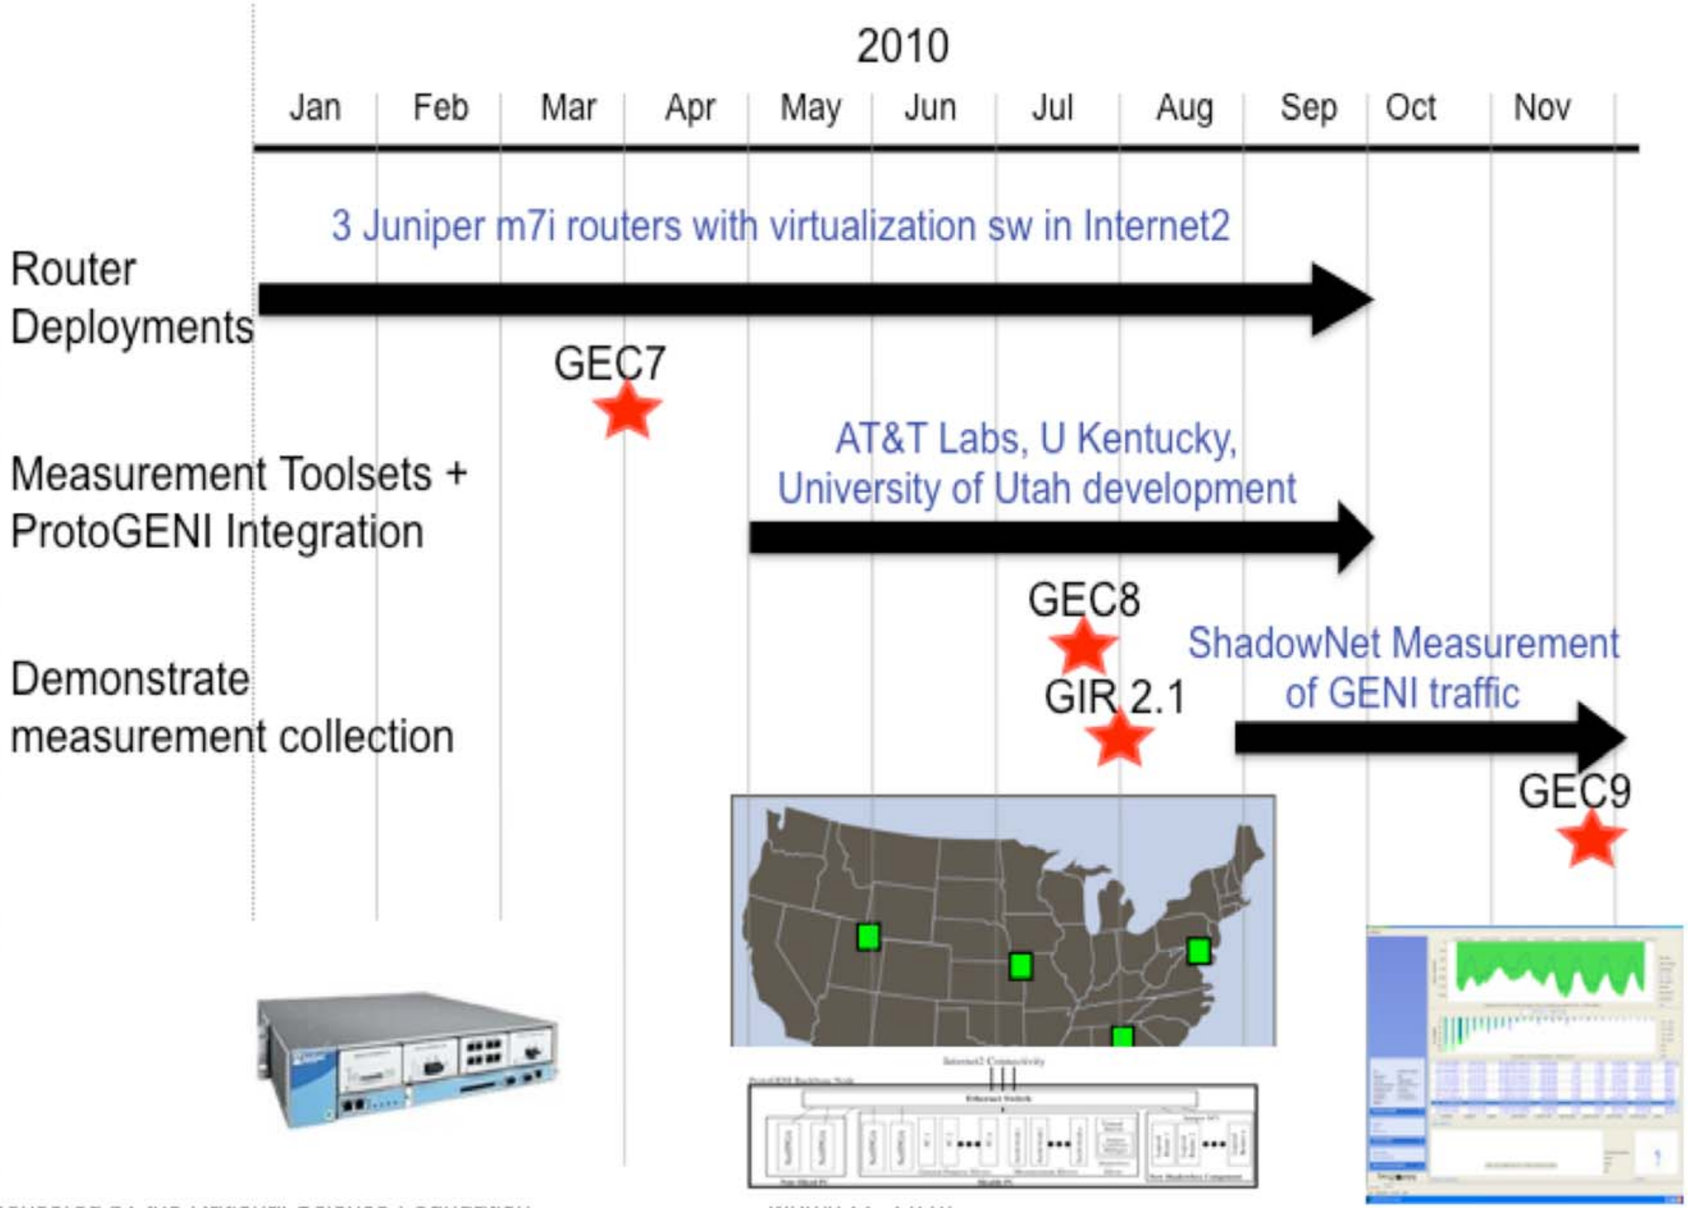

From Roadmap to Infrastructure

- Coordinate deployment plans (wiki pages)
- Set up GPO integration labs for joint project use, integration (e.g. OpenFlow)
- Track individual projects wikis, repositories, tags, releases, kits (external and internal), hardware, documentation, software, campus security policies, support groups, escalation procedures
- Trial GENI Integration Release (GIR) process on Spiral 1 software and resources (example VISE, Enterprise GENI doc, sw and configurations).
- Plan Spiral 2 GIR in February 2010 for software and systems successfully integrated between 2 GENI projects (e.g. ProtoGENI and Instrumentation Tools)
- Test layer2 VLAN data connections with GENI sites (ION, FrameNet, QinQ, EGRE tunnels, OpenVPN) and provide reference configurations for switches, routers
- Prepare net maps, tickets, GENI RUP, procedures (e.g. emergency shutdown), logs, mailing lists for Meso-scale GENI Response Teams
- Coordinate demos at GEC8 and GEC9

Stanford GENI Network

Nick McKeown, PI

Guido Appenzellar Guru Parulkar

- OpenFlow production traffic ***now***
- OpenFlow 1.0 ref implementation ***now***
- Early integration with campus trials HP, NEC, Toroki, Quanta, and OpenWRT switches
- OF sw devel/support
- WiMAX deployment

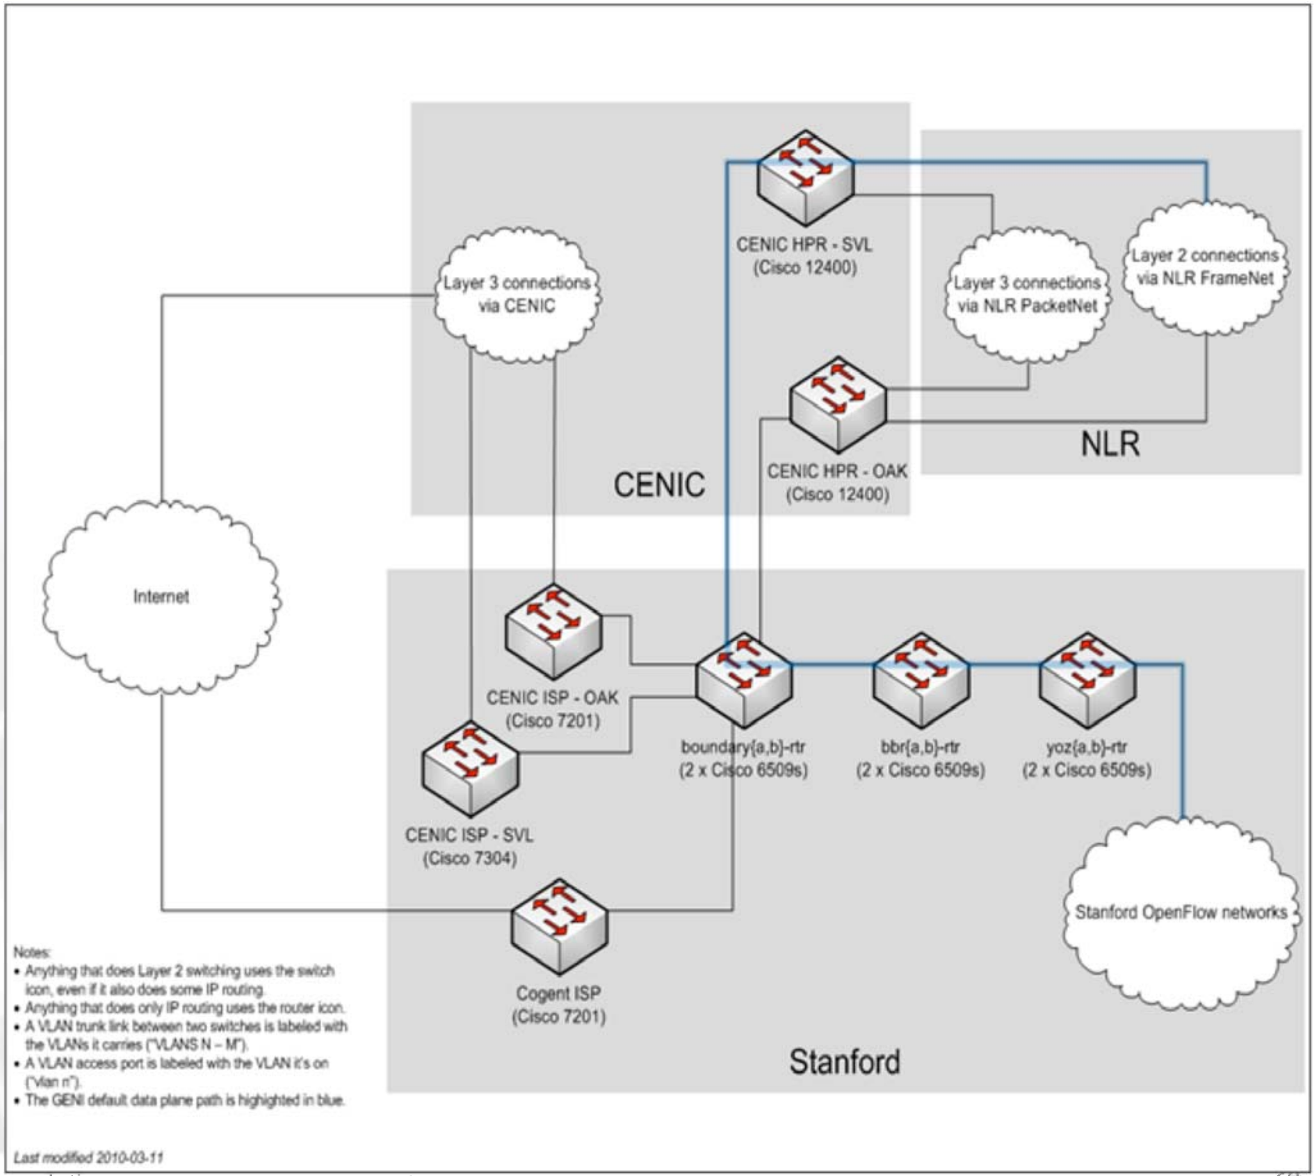

Clemson GENI Network

Kuang-Ching Wang, PI

Jim Pepin, CTO

Dan Schmiedt, Chief Network Engineer

- OpenFlow in ECE lab (wireless mesh and ethernet) ***now***
- Expanding to more campus buildings
- Early integration with campus network operation center

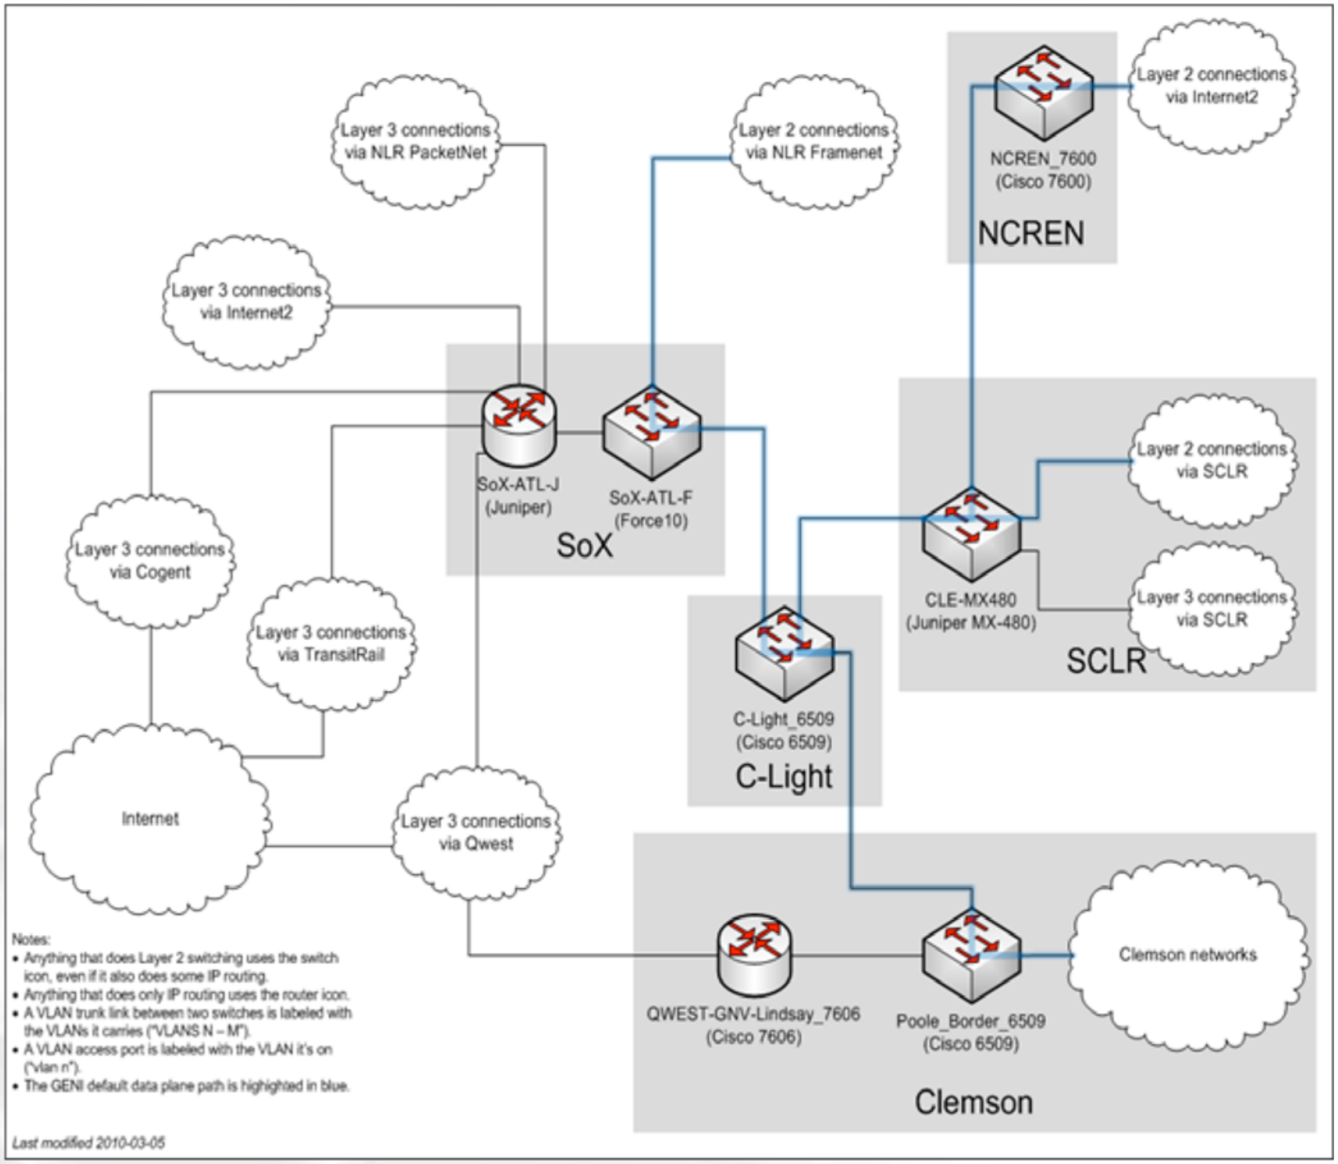

Georgia Tech GENI Network

Nick Feamster, PI

Ellen Zegura

Russ Clark, GT-RNOC

Ron Hutchins, OIT

- OpenFlow in 2 GT-RNOC lab bldgs ***now***
- OpenFlow/BGPMux coursework ***now***
- Dormitory trial
- Access control, authentication focus

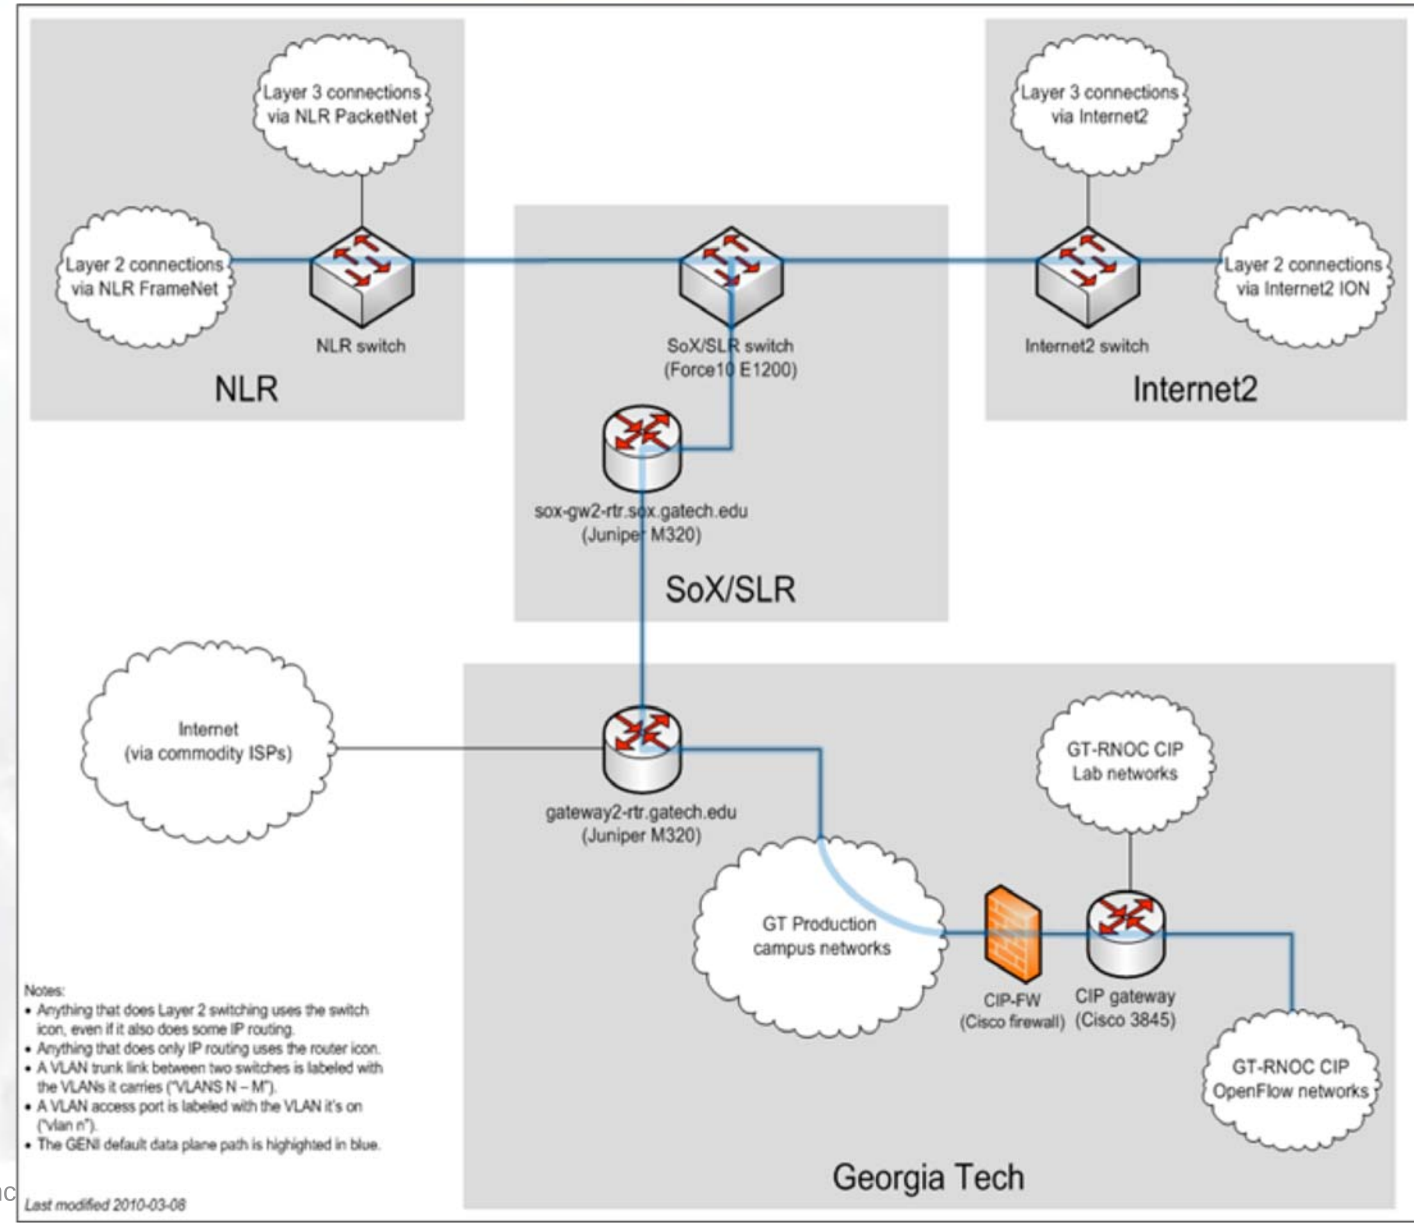

Indiana University GENI Network

Not shown: Christopher Small, PI

Matthew Davy,
Chief Network
Architect

Dave Jent,
IU IT

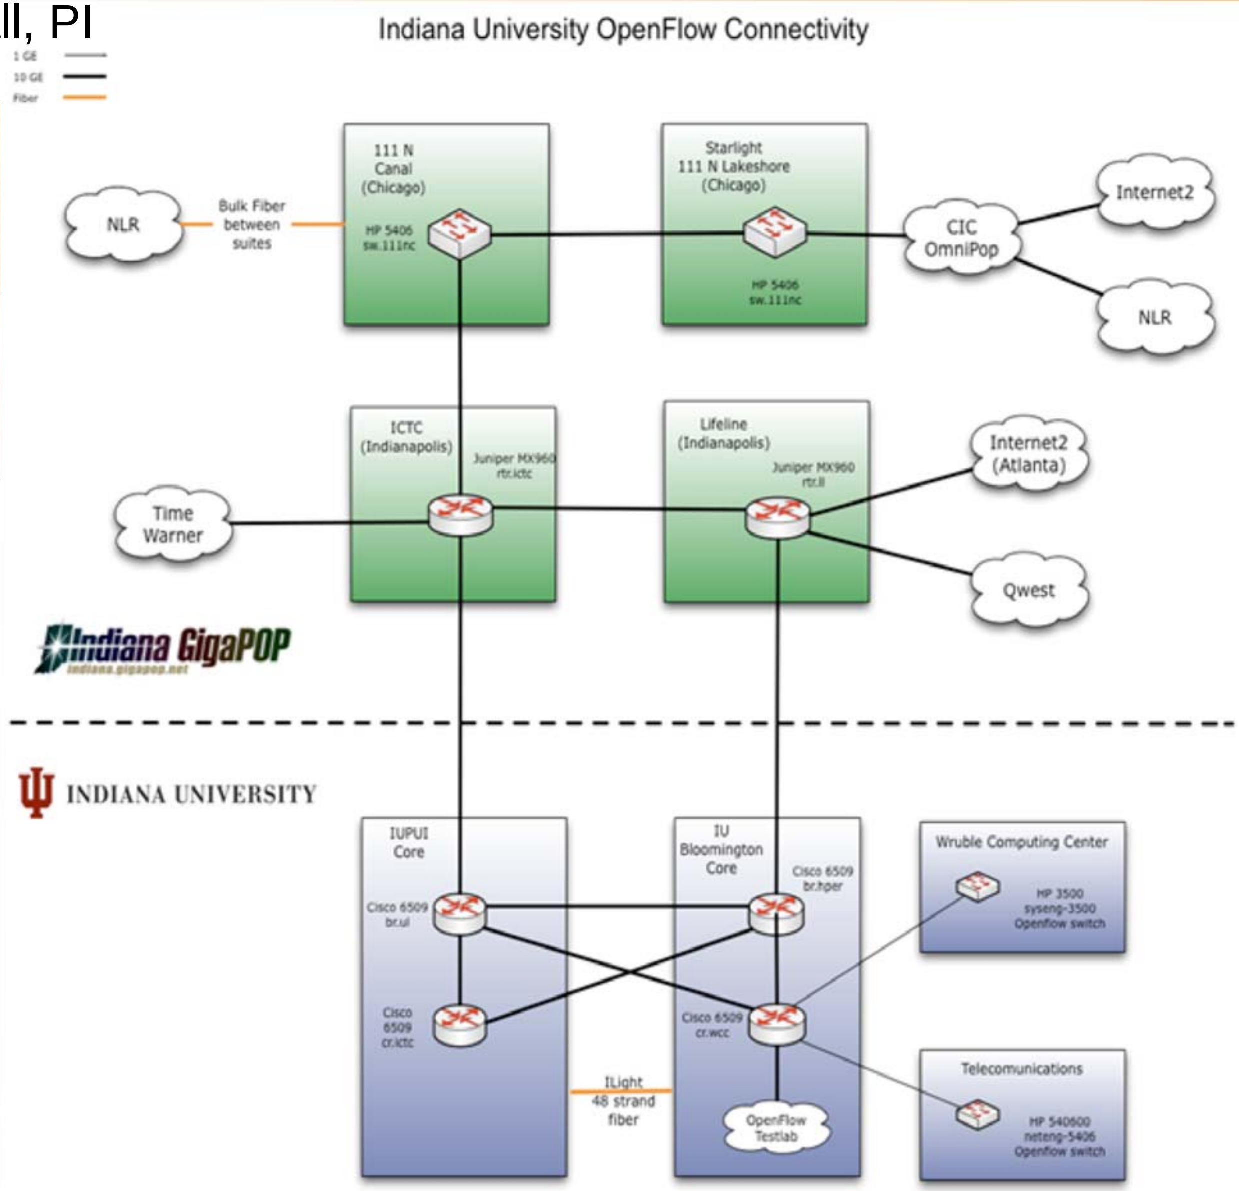

- OpenFlow in IU Testnet Bloomington Data Center ***now***
- Integration with IU GpENI cluster ***now***
- Expanding to other campus locations
- Focus on operations: campus and GMOC

Princeton GENI Network


Michael Freedman, PI


Larry Peterson


Jennifer Rexford


Scott Karlin, Computing Facilities


Chris Tengi, System Admin

- Switch evaluation in progress
- Trial deployment to CS Dept
- Fine-grain opt-in required
- Tools for infrastructure management

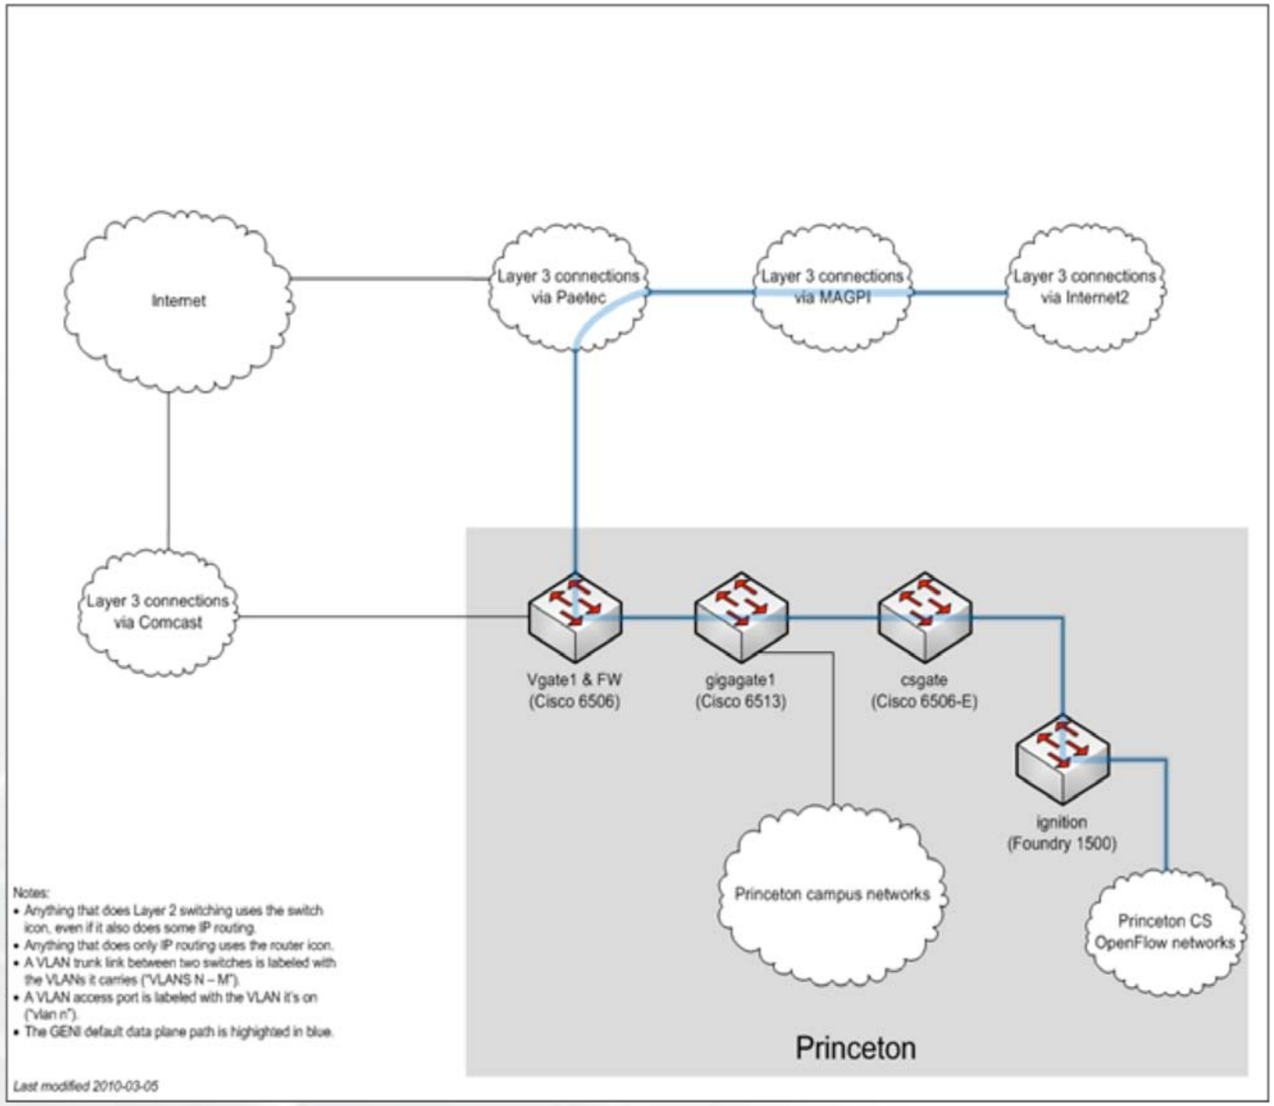

University of Washington GENI Network

Arvind Krishnamurthy, PI

Tom Anderson

Not shown: Clare Donahue, OUWT

- Switch evaluation in progress
- Trial deployment to CS Dept
- Hybrid OF/RouteBrick testbed
- Tools for network function partitioning

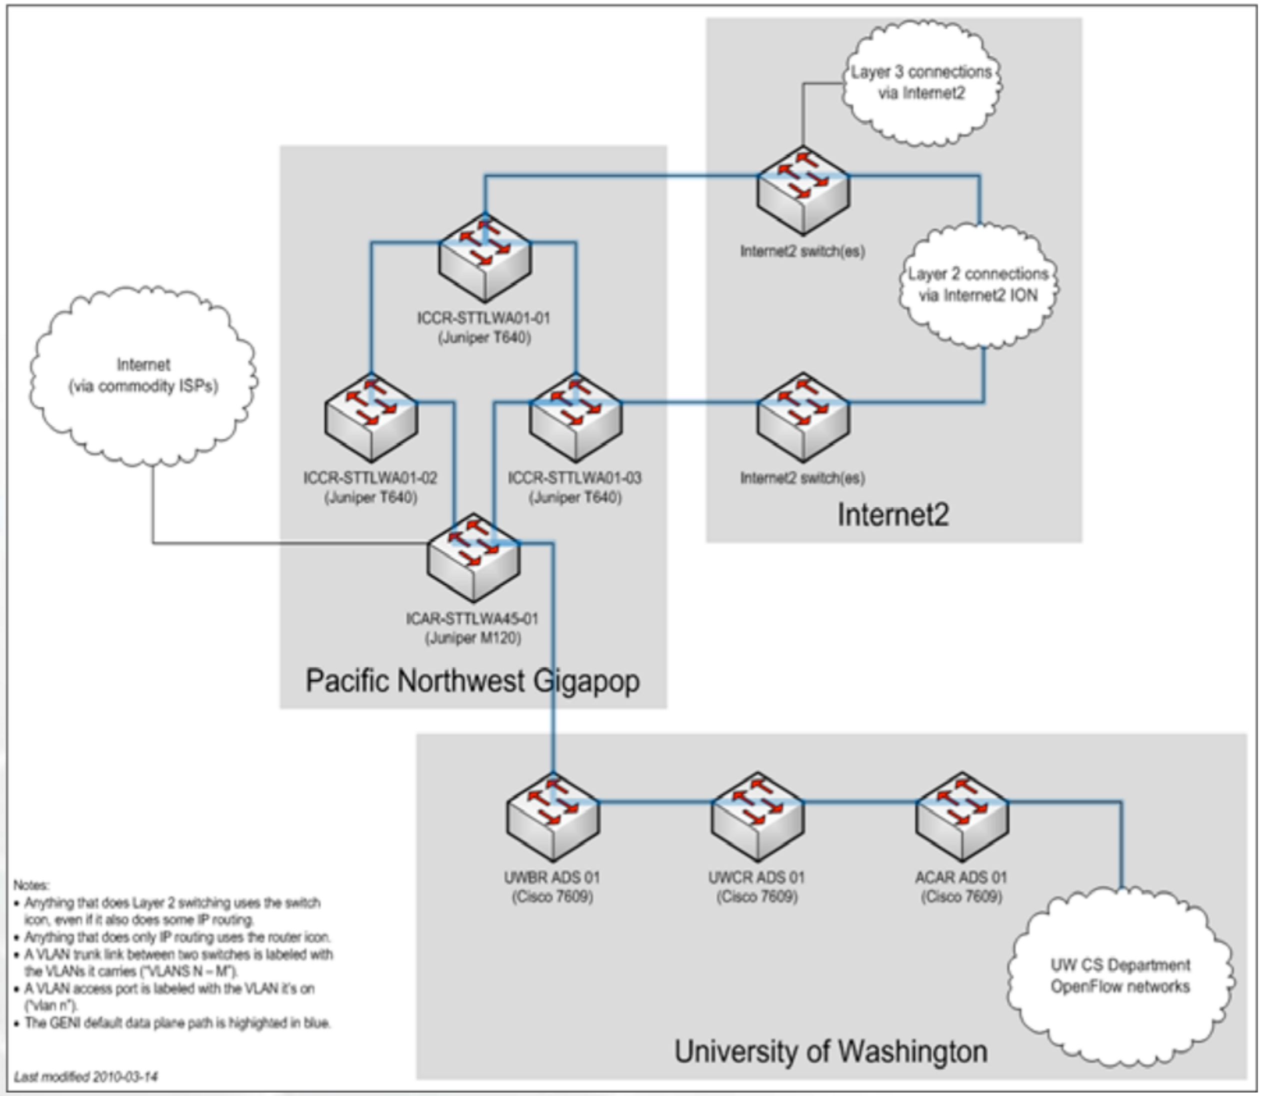

University of Wisconsin, Madison GENI Network

Aditya Akella, PI

Not shown: Perry Brunelli, IT
Hideko Mills, IT

- Switch evaluation in progress
- OpenFlow + Emulab integration ***now***
- Trial deployment to CS Dept
- OpenFlow + trusted computing development
- WiMAX deployment (hybrid WiFi/WiMAX)

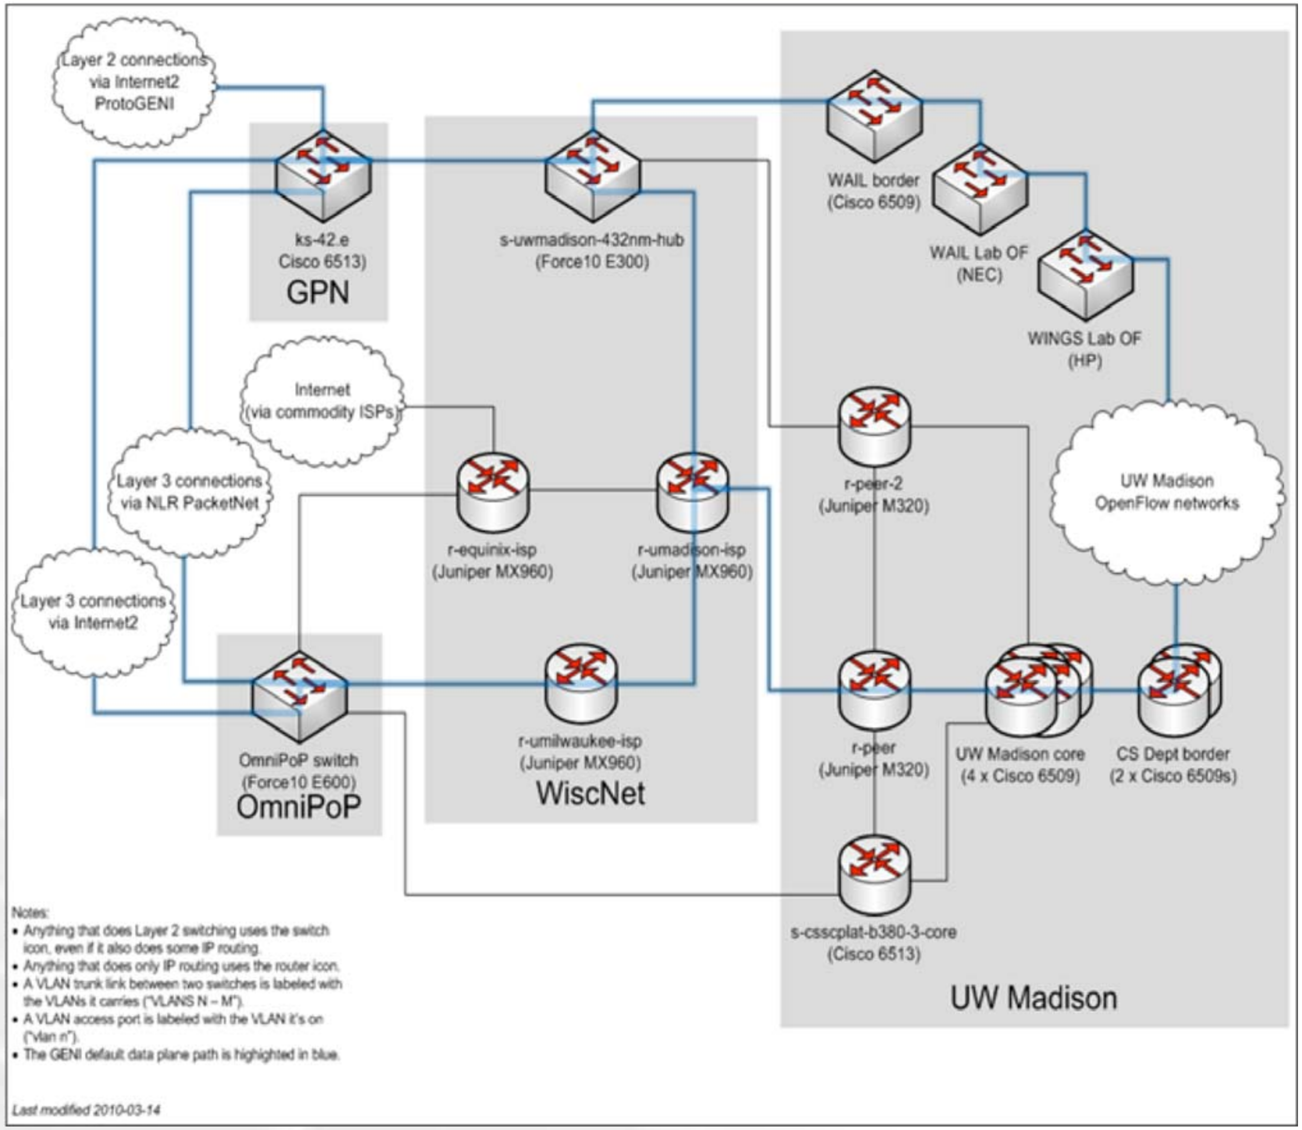

BBN GENI Network

Manu Gosain, John Williams, Josh Smift, Chaos Golubitsky, Nidhi Tare, Heidi Dempsey

Mike Gribaudo, IDSG

- Integration testbeds
- OpenFlow/Campus VLAN integration
- WiMAX integration
- GENI API Agg Mgr
- Support for early use and experiments

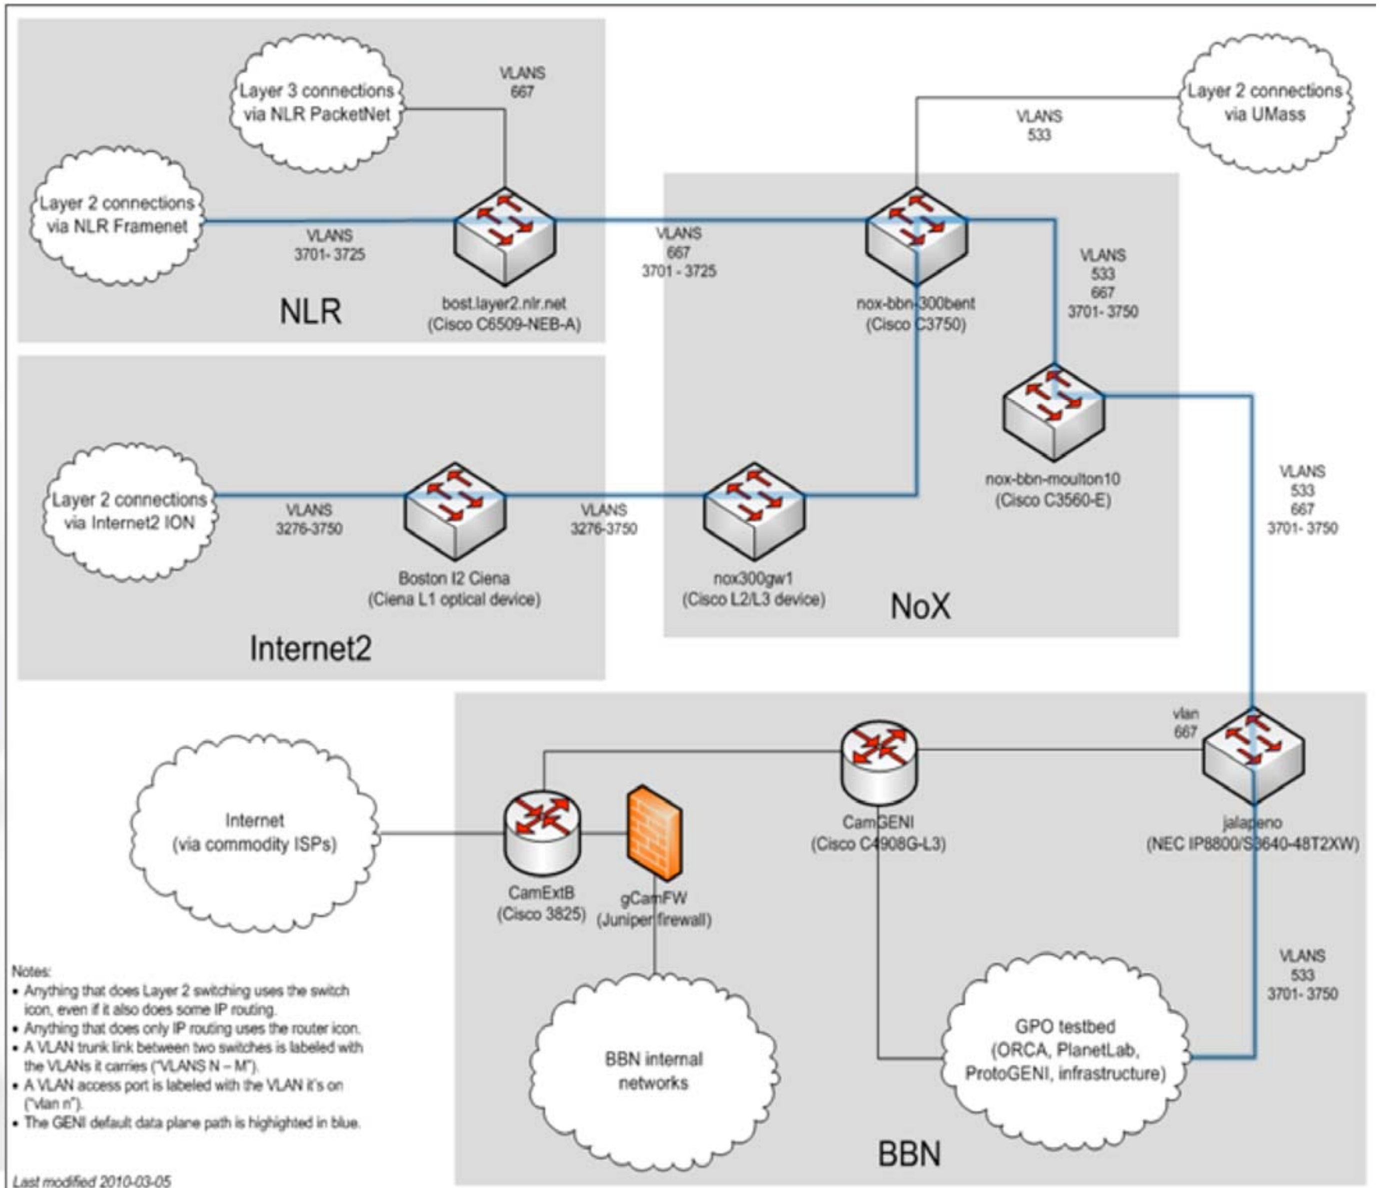

WiMAX deployments

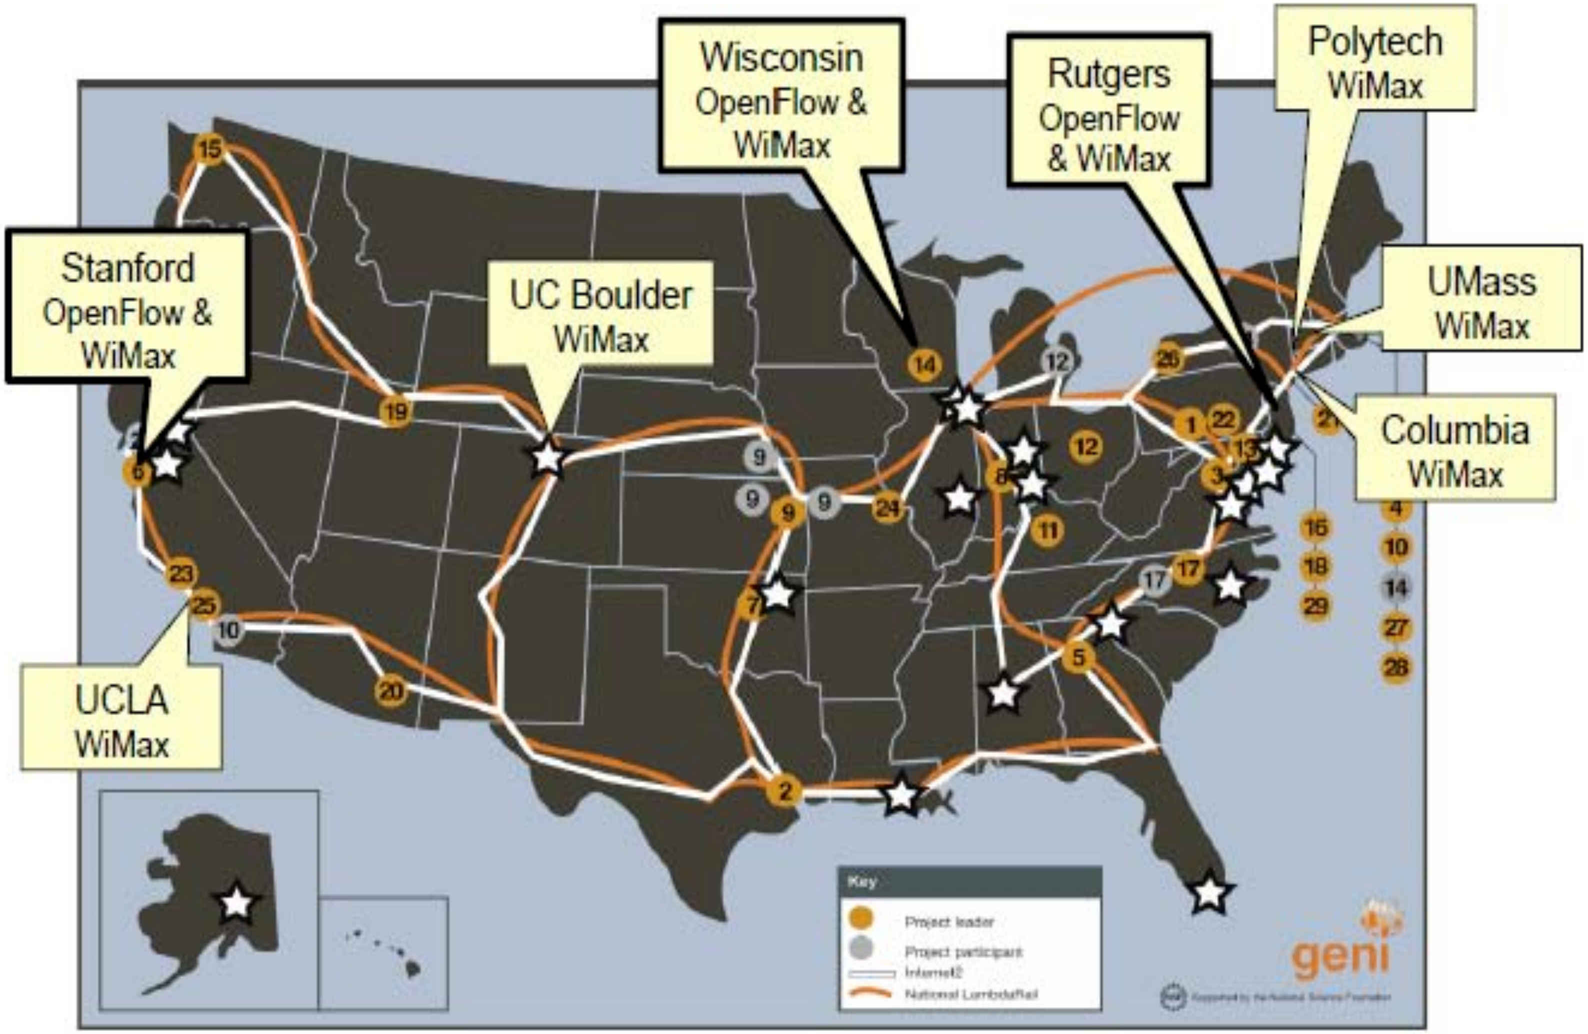

Figure 1. WiMAX campus buildout.

Rutgers / NEC WiMAX kits

NEC Base Station (BTS)
Equipped with:
(blank)
(blank)
radio module #1
control module

NEC Outdoor Unit (ODU)
(shown indoors, but intended fro mounting with antenna)

NEC Base Station (BTS)

Profile A

Supported service classes:

rtPS	real- time polling service
ertPS	enhanced real- time polling service
nrtPS	non real- time polling service
UGS	unsolicited grant service
BE	best effort

PHY	Access mode	SOFDMA/TDD
	Frequency	2535 ~ 2605 MHz
	DL:UL ratio	35:12, 26:21, 29:18
	Channel BW	10 MHz , 8.75 MHz
	FFT size	1024, 512
	Frame duration	5ms
	TX output Power	35dBm (max)
	# of sectors	3
MAC	Head compression	PHS
	ARQ	HARQ/CC, ARQ
	MBS support	Single BS, multiple BMBS
	Resource management	Power control, mode control (idle, sleep etc.)
Networking	IP protocols	IPv4, IPv6
	Bridging/Routing	Transparent L2 switch, Bridging
	Packet handling	802.1Q VLAN, PHS**)

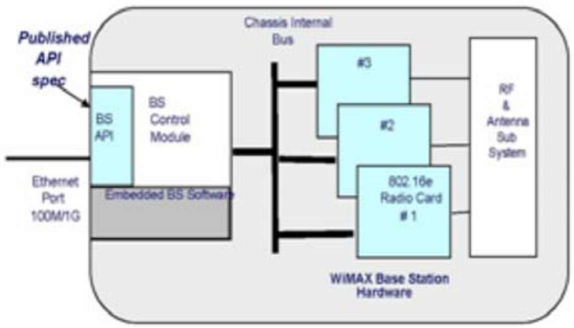

Omni-directional antenna
(elev. < 6ft above roof)

WiMAX coverage area:
3-5km

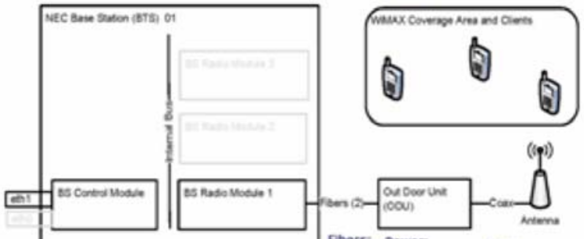

WiMAX kit (hardware)

Item	Supplier	Part No	Description	Qty
1.1	NEC	NWA-025035-001	2.5GHz SECTOR ANT: Antenna, dual-polarization, omnidirectional antenna, for operation in the 2.5GHz range.	1
2.1	NEC	NWA-027932-001	NEC 2.5GHz ODU: Base Station Outdoor Unit, for connection to one antenna.	1
3.1	NEC	NWA-024297	IDU: Base Station Indoor Unit, equipped equipped for 1-Sector Configuration, with one Network Interface (NW INTFC) card and one Channel Card (CHC).	1
3.2	NEC	-	Network Interface (NW INTFC) card	1
3.3	NEC	-	Channel Card (CHC)	1
4.1	NEC	-	100m (approx 300ft) of dual fiber cable, rated for outdoor mounting, yellow SM, with connectors	1
5.1	-	-	1m (approx 3ft) (or as needed) antenna cable, LMR300 (or LMR400) coax, with N connectors	
5.2	-	-	Lightning arrestor, for use with ODU antenna cable connection	
6.1	-	-	10m (approx 30ft) outdoor power cable, two conductors, 12AWG, rated for outdoor use, with ODU Circular Connector, soldered onto cable	
6.2	Mouser: Hirose	JR25WP-4S71	ODU Circular Connector	
6.3	-	-	Additional outdoor power cable, as needed, two conductors, 12AWG, rated for outdoor use, to extend from roof down to indoor testbed equipment room	
7.1	-	-	10m (approx 30ft) indoor power cable, two conductors, 10AWG, with one IDU Rectangular Power Connector and four IDU Rectangular Power Connector Contacts , two for each of the positive and negative DC power rails, crimped onto cable with a special tool, or soldered.	
7.2	Mouser: Tyco/Amp	1-917807-2	IDU Rectangular Power Connector	
7.3	Mouser: Tyco/Amp	316041-2	IDU Rectangular Power Connector Contacts	
8.1	-	-	Server, loaded with ASN-GW and ORBIT Management Framework software, for use during installation and initial checkout, and perhaps for final operating configuration	

9.1	-	-	Outdoor site, with pole for mounting of antenna and Base Station Outdoor Unit, grounded to building ground for lightening protection	1	-	-	IMPORTANT: In some locales, the installation and grounding will need to be certified by a professional engineer
9.2	-	-	Pole for mounting of antenna and Base Station Outdoor Unit	1	-	-	Pole diameter between 48mm and 120mm
9.3	-	-	Ground cable, as needed.	1	-	-	-
10.1	Mouser: Mean Well	SP-240-48	Power supply for ODU, Mean Well SP-200-48, 110VAC input, -48VDC output, rated 4.2A or 200W, mounted in testbed equipment room, or outdoors on roof, in weathertight enclosure (see UMass Amherst)	1	-	-	-
10.2	-	-	When available, one weatherproof container on roof, to mount ODU power supply.	0 - 1	-	-	See UMass Amherst for typical design.
11.1	-	-	Indoor site, with racks, for mounting of Base Station Indoor Unit and Linux Servers	1	-	-	-
11.2	Mouser: Mean Well	SP-320-48	Power supply for IDU, Mean Well SP-320-48, 110VAC input, -48VDC output, rated 6.7A or 320W, mounted in testbed equipment room	1	-	-	-
11.3	-	-	Servers, for loading with ASN-GW and ORBIT Mangement Framework software, for final operating configuration	0 - 3	1,600	0 - 4,800	Depends on final configuration
11.4	-	-	Ethernet switch ports, for multiple VLANs	-	-	0 - 2,500	OpenFlow? compatible switch desired. PC-based design underway that utilizes FPGA card, total cost: $1300 + $400 = $1700

WiMAX kit (software)

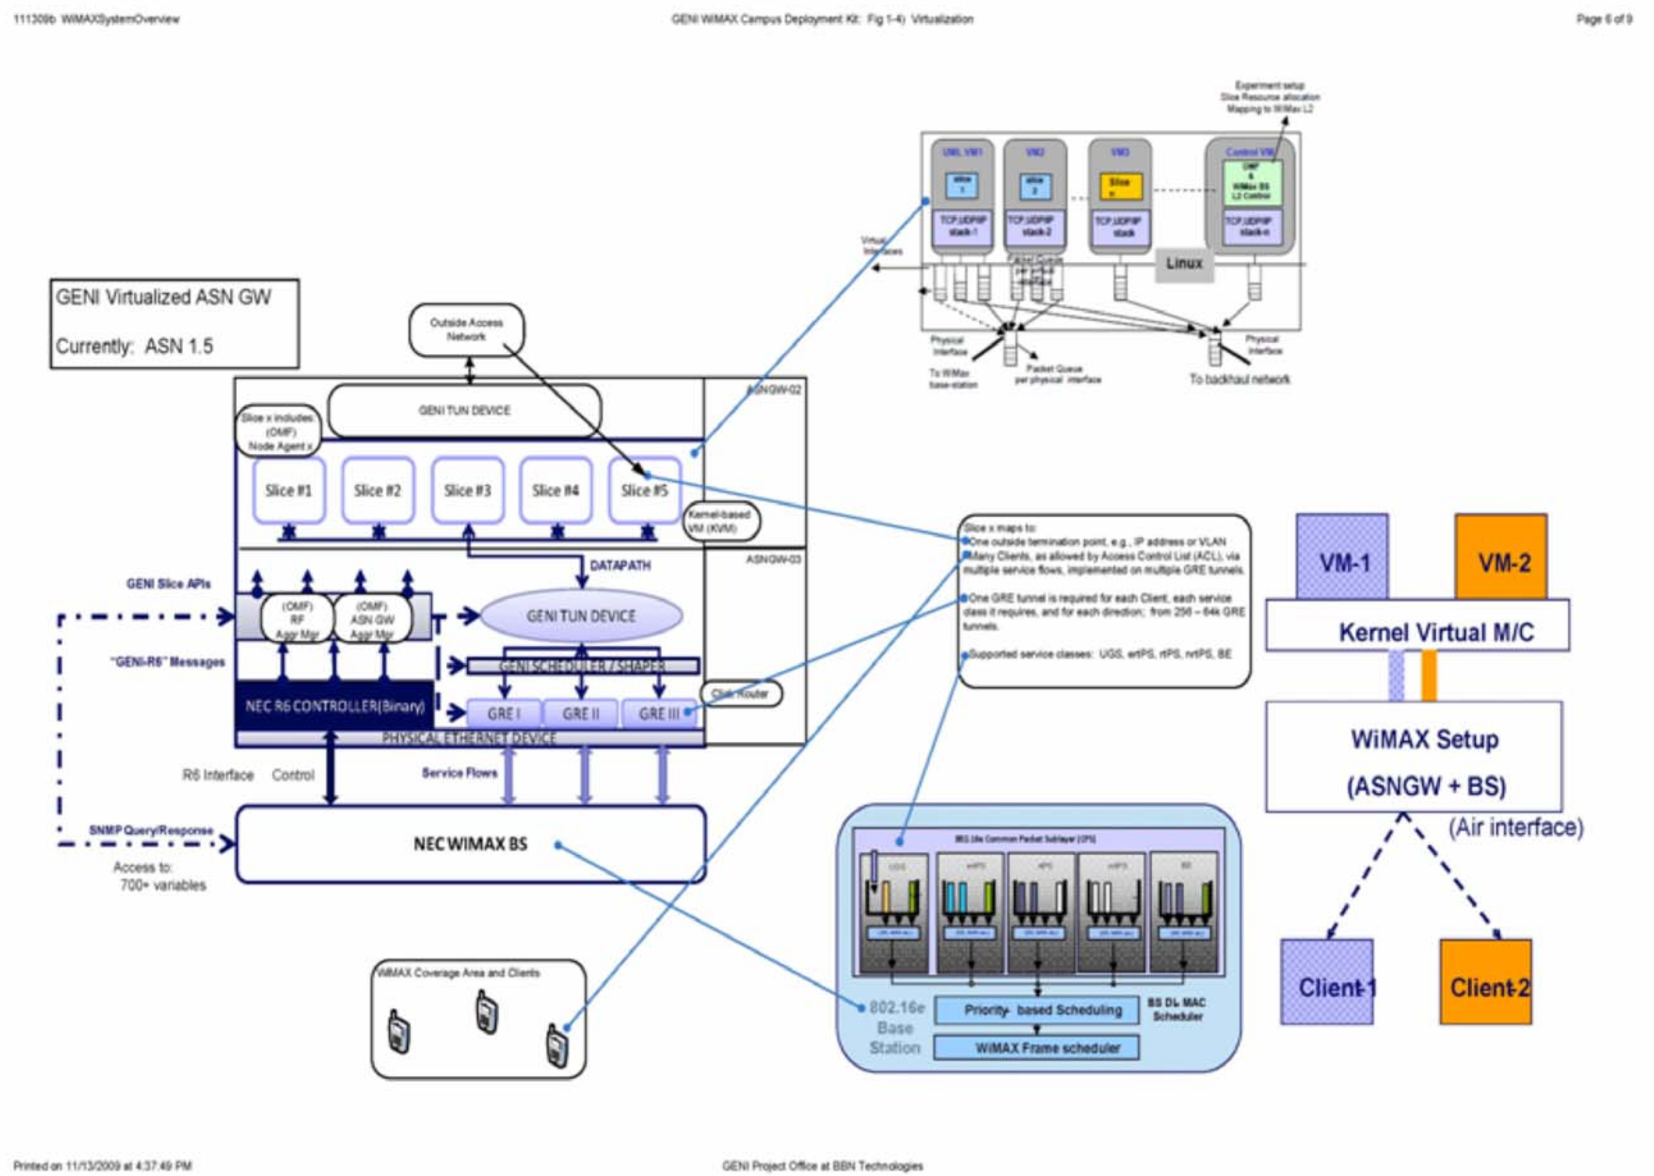

Kits will be delivered to Rutgers this May, then rolled out to campuses

Step	Description	Stanford	Rutgers WINLAB	Rutgers Bush	NEC	Columbia	Poly NYU	UCLA	Colorado	UMass	Wisc	BBN
1	Identify Outdoor Site on Campus	-	Completed	Completed	-	Completed	Completed	Started	Started	Completed	Completed	Completed 1/28/10
2	Obtain FCC License	-	Completed	Completed	-	Applied 2/25	Applied 2/22	Started	Started	Completed 3/2010	Completed 1st Started 2nd	Completed 1st 2/24/10 Applied 2nd 1/14/10
3	Identify Indoor Equipment Room on Campus	-	Completed	-	-	-	Completed	-	-	Completed	-	Completed 7/09
4	Install Wiring at Outdoor Site and to Indoor Equipment Room	-	Completed	-	-	-	-	-	-	-	-	-
5	Install Pole at Outdoor Site and Ground	-	Completed	-	-	-	-	-	-	-	-	
6	Install Antenna at Outdoor Site	-	Completed	-	-	-	-	-	-	-	-	-
7	Install Base Station Outdoor Unit	-	Completed	-	-	-	-	-	-	-	-	-
8	Complete Wiring at Outdoor Site and to Indoor Equipment Room	-	Completed	-	-	-	-	-	-	-	-	-
9	Install Indoor Equipment Racks and Ground	-	Completed	-	-	-	-	-	-	-	-	-
10	Install 120VAC and - 48VDC Power Feeds and/or Supplies	-	Completed	-	-	-	-	-	-	-	-	-
11	Install Base Station Indoor Unit	-	Completed	-	-	-	-	-	-	-	-	-
12	Complete Wiring of Base Station Indoor Unit	-	Completed	-	-	-	-	-	-	-	-	-
13	Initial Check -Out of Antenna and Base Station Outdoor and Indoor Units	-	Completed	-	-	-	-	-	-	-	-	-
14	Install Ethernet (VLAN) Switch, or Identify Existing Capacity	-	Completed	-	-	-	-	-	-	-	-	-
15	Install Linux Server(s) and Software	-	Completed	-	-	-	-	-	-	-	-	-
16	Connect Backbone Network	-	Started	-	-	-	-	-	-	-	-	-
17	Provide Reference Client Platform	-	Completed	-	-	-	-	-	-	-	-	-
18	Initial Check -Out of Campus WiMAX Testbed	-	Completed	-	-	-	-	-	-	-	-	-

MesoScale in the backbones

- GENI build-out spreading through the US research networks (backbones now, regionals in Spiral 3)
- Internet2, NLR backbones are installing 5 HP ProCurve OpenFlow switches in each backbone
- Internet2, NLR will interconnect GENI Layer 2 data planes
- ShadowNet: installs 3 Juniper M7i routers for measurements in I2 PoPs this year
- Preliminary investigations into ProtoGENI/OpenFlow
- Regional networks will have workshop at GEC 8, leading to GENI build-out through regionals in Spiral 3

Outline

- GENI – Exploring future internets at scale
- Current status and plans
  - GENI Spiral 2
  - Meso-scale buildout
  - Starting experimentation
- GPO program activities
  - Retasking & reorganization
  - GENI Solicitation 3
- Wrap-up

Building the experiment pipeline

- Three pipeline efforts to encourage experiments
  - Getting earliest experiments running (now)
  - Organizing / running training sessions (startup)
  - NSF-sponsored experimentation workshop (June)
- NSF Future Internet Architectures program

Experiments now in progress

- Four experiments now in progress
  - Davis Social Links (Felix Wu, UC Davis)
  - Floating Cloud-Tiered Internet (N. Shenoy, RIT)
  - DTN for Space Networks (Ed Birrane, JHU APL)
  - Pigeon Net (Jiang Li, Howard)
- General process
  - Very early in GENI – PIs need active help
  - GPO engineers help PI get experiment ported to GENI cluster (ProtoGENI currently the favorite)
  - Now shaking them down in GPO lab
  - Will then put it out into the national GENI infrastructure
  - Hope to present research experiment results at GEC 9

Training sessions

- Goals
  - Train students / young PIs on how to use GENI
  - Help establish & grow active student cohorts
- Training sessions
  - Currently oriented by cluster / toolset
  - Next year should be more “generic GENI”
  - Jon Turner gave SPP training last week at GEC 7, about 20 participants
  - GEC 8 (July) will host a number of different training sessions, each ½ day long

NSF Experimentation Workshop

- Current plans
  - Chairs: Jex Rexford, Guru Parulkar
  - To be held at Princeton in late June
  - About 40 participants (20 profs, each with student)
  - Active GPO participation to say “what’s possible”
- Goals
  - A number of solid research experiment proposals . . .
  - . . . that can be run as GENI experiments starting in the Fall 2010 timeframe . . .
  - . . . and which can be promptly funded by supplements or EAGER grants

NSF Future Internet Architectures

- Loosely coupled to GENI
  - Experimentation / trials required
  - Can use GENI, National Cyber Range, or purpose-built infrastructure
- Programmatics
  - Proposals due late April 2010
  - Teams appear to have strong overlap with GENI prototyping teams
  - GPO will be fair & even-handed, will not participate in any FIA proposals
  - (Note that FIA budget appears substantially bigger than the total GENI budget)

Outline

- GENI – Exploring future internets at scale
- Current status and plans
  - GENI Spiral 2
  - Meso-scale buildout
  - Starting experimentation
- GPO program activities
  - Retasking & reorganization
  - GENI Solicitation 3
- Wrap-up

GPO "hands on" focus for making Spiral 2 a success

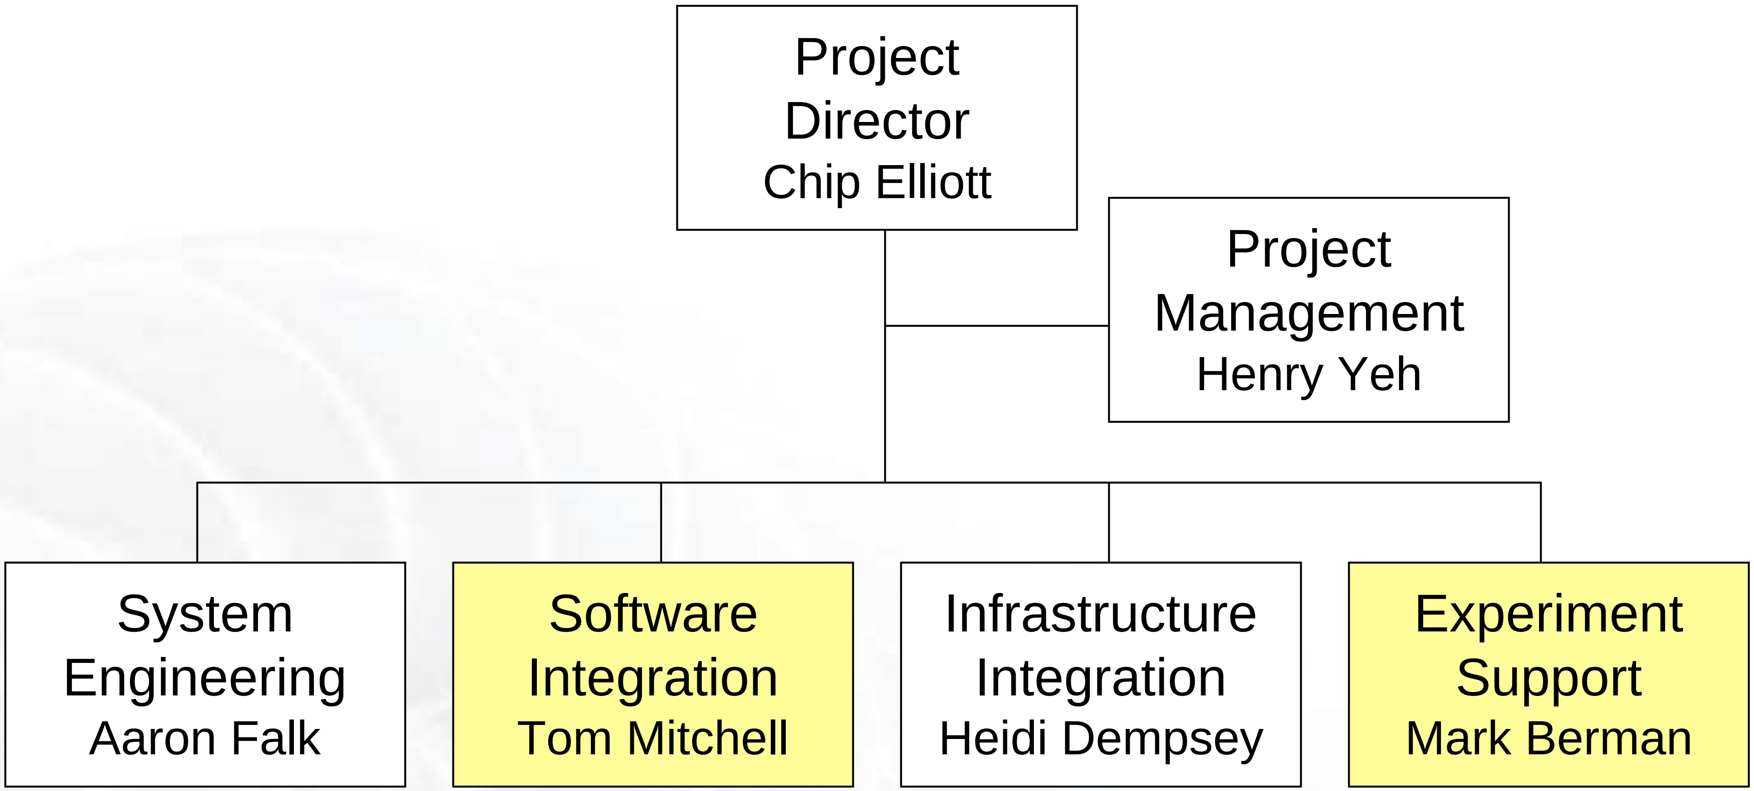

Emphasizes integration & experimentation.

Retasking / reorg started in November 2009; now almost complete.

Solicitation 3 ideas

We seek your suggestions and feedback

- Actively pondering GPO Solicitation 3
  - Notional schedule: Issue solicitation in late spring, with proposals due in mid-late summer
  - Notional funding level: similar to Solicitation 1
  - Talk to us now about your ideas
- Solicitation areas as currently envisioned
  1. **Aggressively grow meso-scale build** (next slide)
     1. More campus, regional, & backbone sites
     2. New "GENI Racks" (eg rack of PCs with OpenFlow switch)
  2. **GENI Instrumentation system** (build & deploy)
  3. **Operations / experiment support / training / education & curriculum development**
  4. **Interesting new ideas**

Aggressively grow meso-scale build

- This is just a concept – we seek your input
- **Accelerate and expand meso-scale build** started in Spiral 2 (add more campus, regional, backbone sites)
- Inject "**GENI Racks**" throughout to beef up computation / storage
  - 1 high end Rack = basic unit of computation / storage
  - Notionally a rack of 1U computers with OpenFlow switch (eg)
  - Highly sliceable, programmable, virtualized, & significant storage
  - Deploy into network's topologically significant points (eg, backbones, regionals, campuses, near WiMAX)
  - Wide range of possible research uses
  - Eg, programmable routers, content distribution, …

Outline

- GENI – Exploring future internets at scale
- Current status and plans
  - GENI Spiral 2
  - Meso-scale buildout
  - Starting experimentation
- GPO program activities
  - Retasking & reorganization
  - GENI Solicitation 3
- Wrap-up

Summary

- "GENI is a virtual laboratory for exploring future internets at scale"
  - How are we doing?
- Rapid progress to date
  - GENI community appears highly energized and surprisingly happy
  - System architecture is taking shape via spiral development
  - Meso-scale build has considerable buy-in from PIs, campus CIOs, national backbones, regionals
  - We are executing plans for getting a number of research experiments started on the GENI suite
- What are the next steps (Spiral 3) ?
  - Converge upon interoperable control frameworks & tools
  - Aggressively grow the meso-scale build, adding "GENI racks"
  - Transition to "operations" to support large-scale, continuous experiments