GENI

Exploring Networks of the Future

Aaron Falk & Mark Berman

GENI Project Office

June 2010

www.geni.net
groups.geni.net

Sponsored by the National Science Foundation

Outline

- Introduction
- Some Simple Examples
- Resources
- Being a Pioneer
- Appendix: Concepts & Terminology

GENI's Unique Advance

- Today
  - Lots of specific testbeds
  - Mostly homogeneous
  - Require separate accounts, tools
  - Interconnected via Internet
- GENI
  - End-to-end, controlled interconnection
  - Shared toolset
  - Common authentication, access control
  - Direct L2 access to end-users
  - Lots of stuff (quantity and diversity)

Resource discovery

Aggregates publish resources, schedules, etc., via clearinghouses

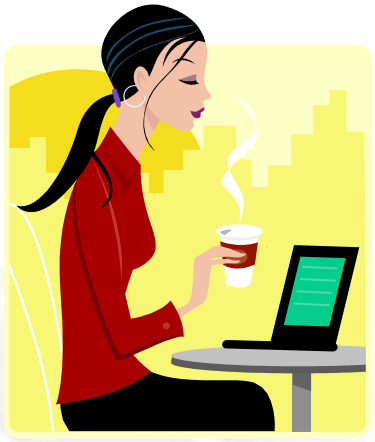

Researcher

What resources can I use?

GENI Clearinghouse

These

Components

Components

Components

Aggregate A
Computer Cluster

Aggregate B
Backbone Net

Aggregate C
Metro Wireless

Slice creation

Clearinghouse checks credentials & enforces policy
Aggregates allocate resources & create topologies

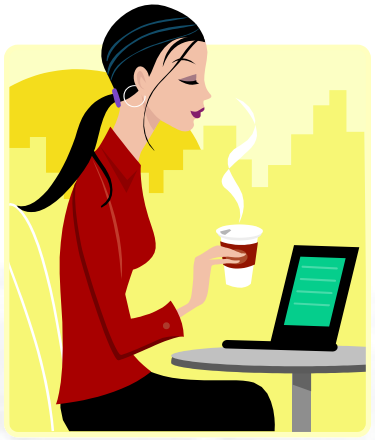

Experimentation

Researcher loads software, debugs, collects measurements

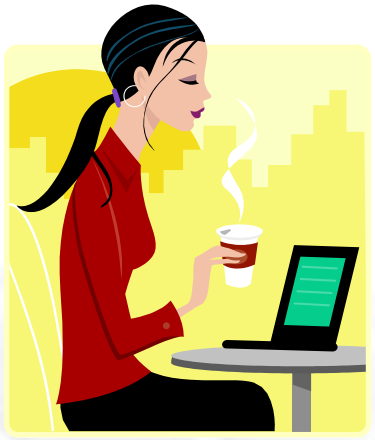

Experiment – Install my software, debug, collect data, retry, etc.

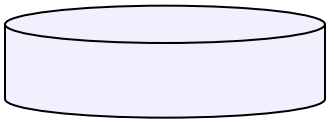

GENI Clearinghouse

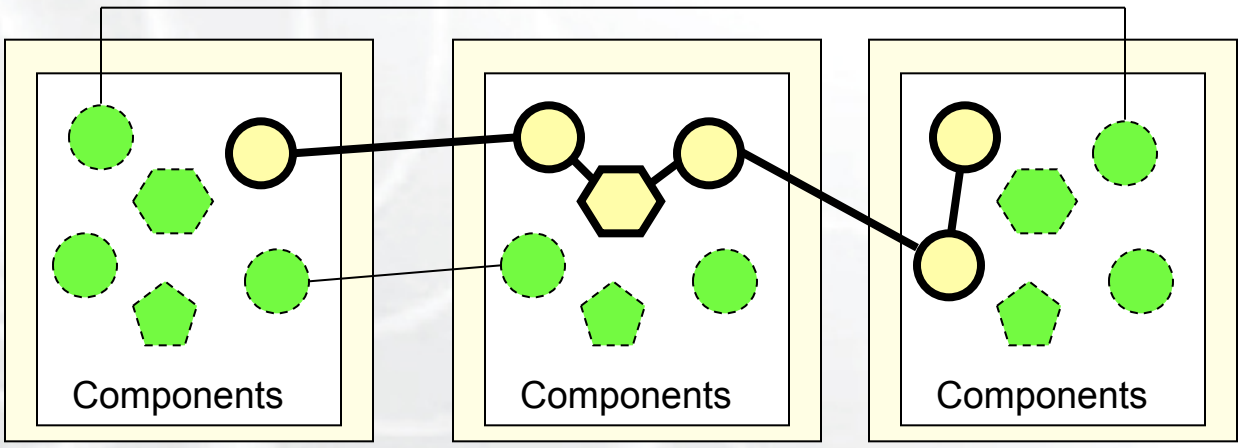

Slice growth & revision

Allows successful, long-running experiments to grow larger

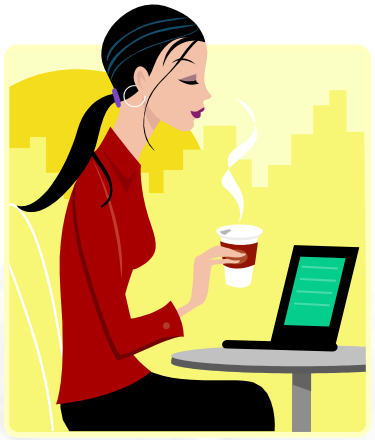

Federation of Clearinghouses

Growth path to international, semi-private, and commercial GENIs

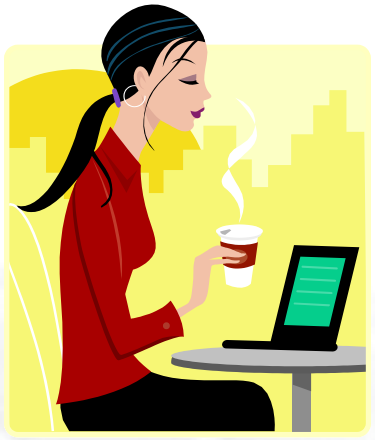

Make my slice even bigger !

GENI Clearinghouse

Federated Clearinghouse

Components

Components

Components

Components

Aggregate A
Computer Cluster

Aggregate B
Backbone Net

Aggregate C
Metro Wireless

Aggregate D
Non-NSF Resources

Operations & Management

Always present in background for usual reasons
Will need an 'emergency shutdown' mechanism

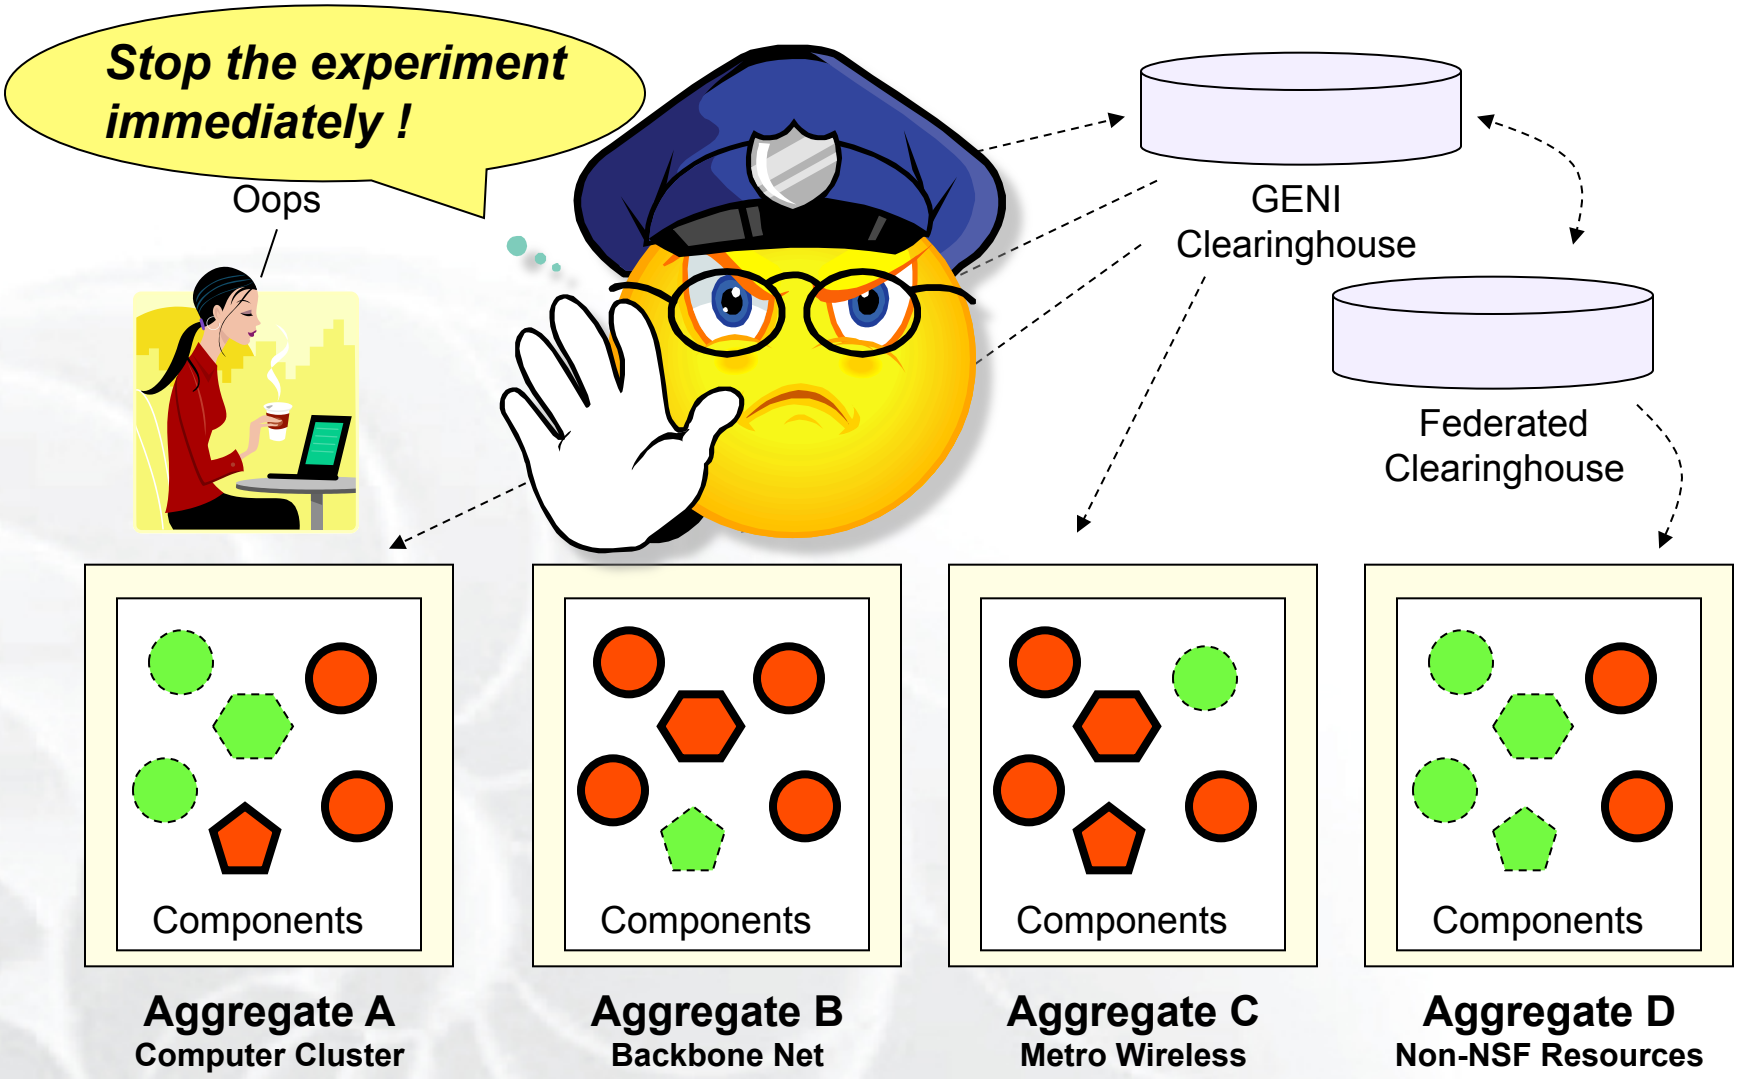

Currently, Clusters are Tightly Linked to Resources

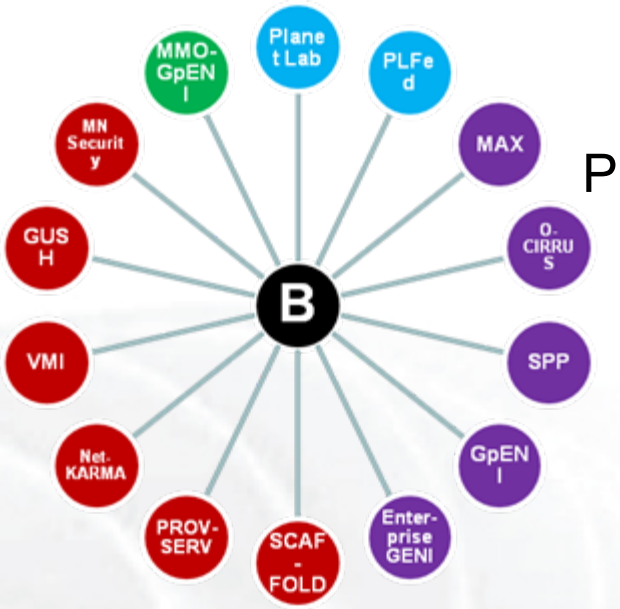

PlanetLab

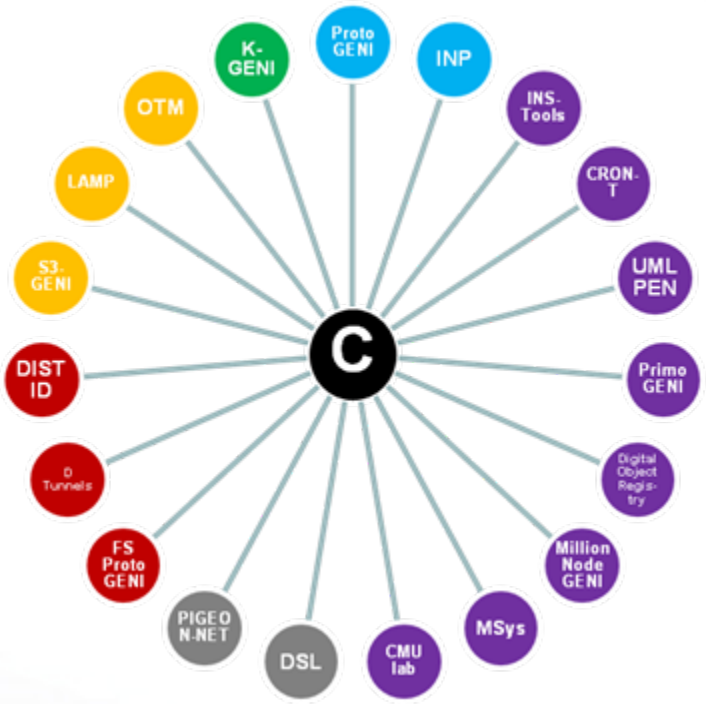

ProtoGENI

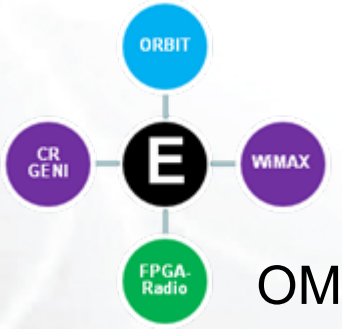

OMF

ORCA

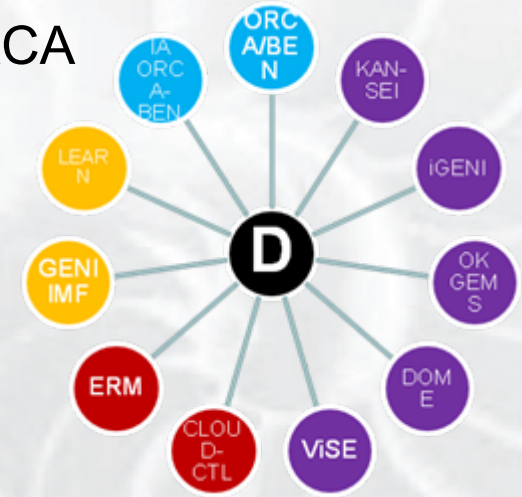

…but this will start to decouple by Fall 2010

DInfo-Subs | GMOC | REG OPT
CMUDP | SEC-POL | DSN-HIVE
GENI 4YR | SEC ARCH | DMeas
LEFA | Att-GENI | EXP-SEC

Outline

- Introduction
- Some Simple Examples
- Resources
- Being a Pioneer
- Appendix: Concepts & Terminology

ProtoGENI Resources

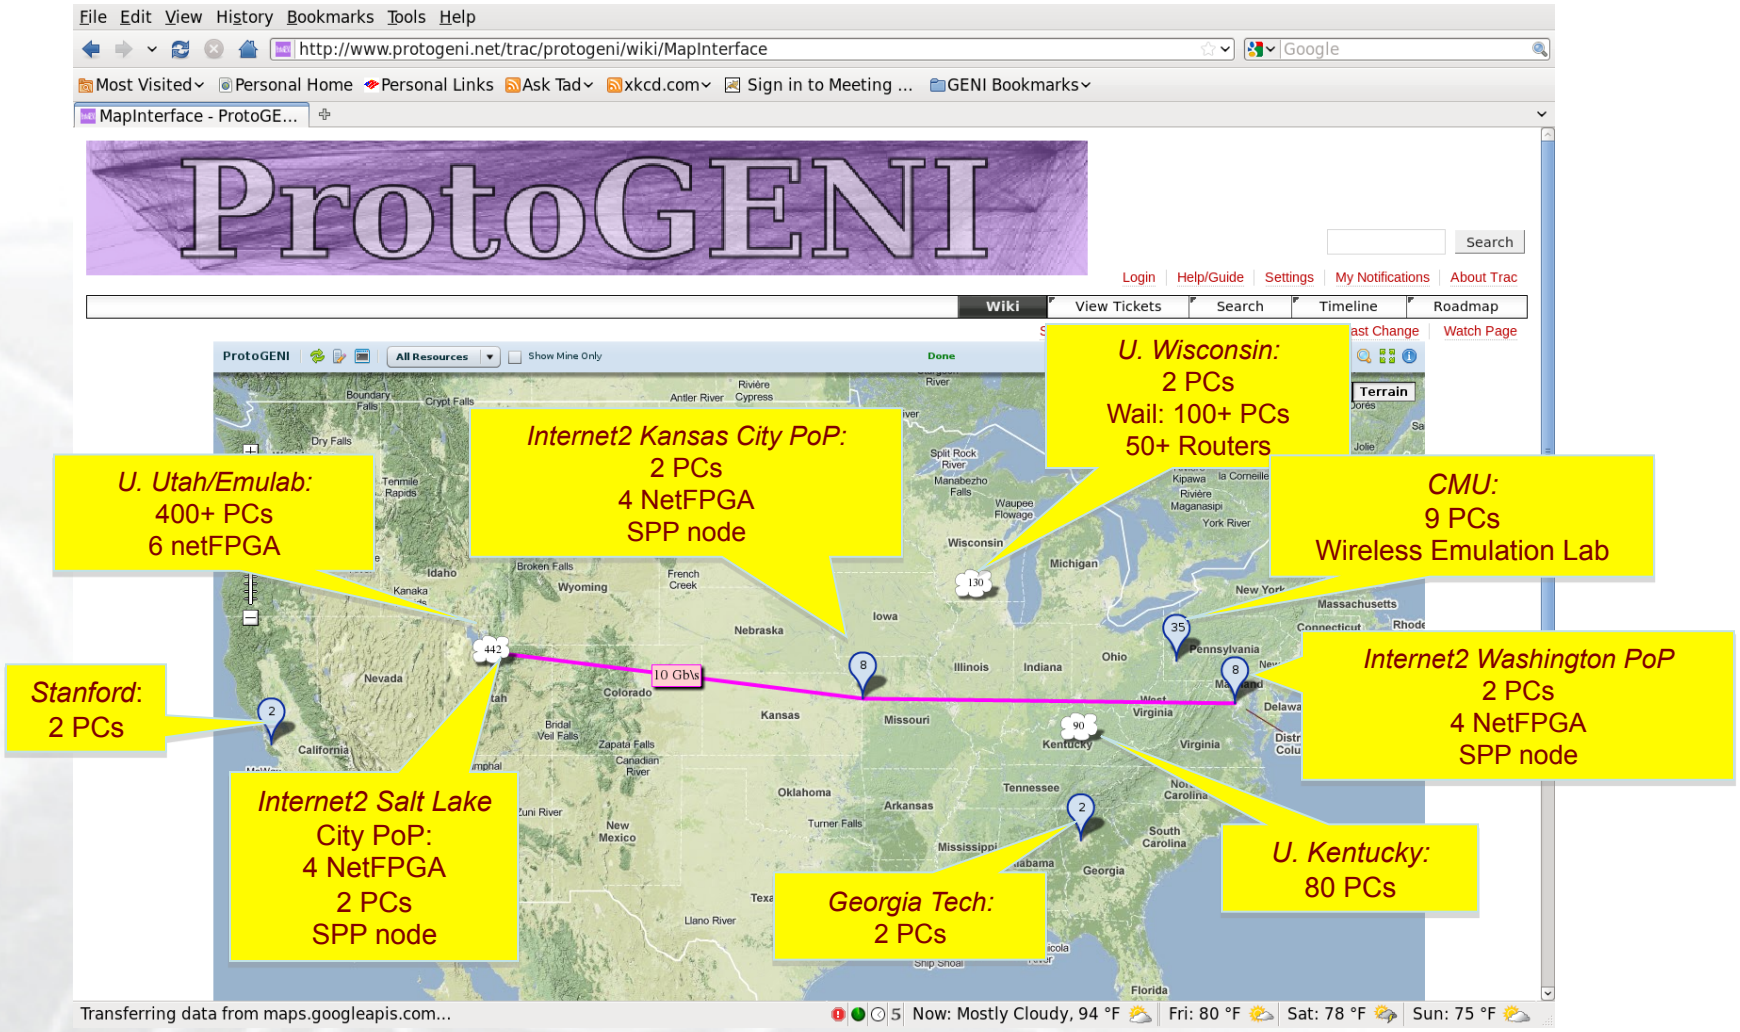

Creating an Experiment via ProtoGENI

- A top level reference about ProtoGENI offered by Utah
  - http://www.protogeni.net/trac/protogeni/wiki/Tutorial

1. Register as a ProtoGENI user and get ProtoGENI credentials
2. **Create a slice** for your experiment
3. **Discover resources** via the web GUI, select (virtual) hosts and topology
4. **Create slivers** reserves resources, installs OS & your keys & configures the network
5. **Boot slivers** (re)starts your machines
6. Log in and start programming!

Programmable WiMax Base Stations

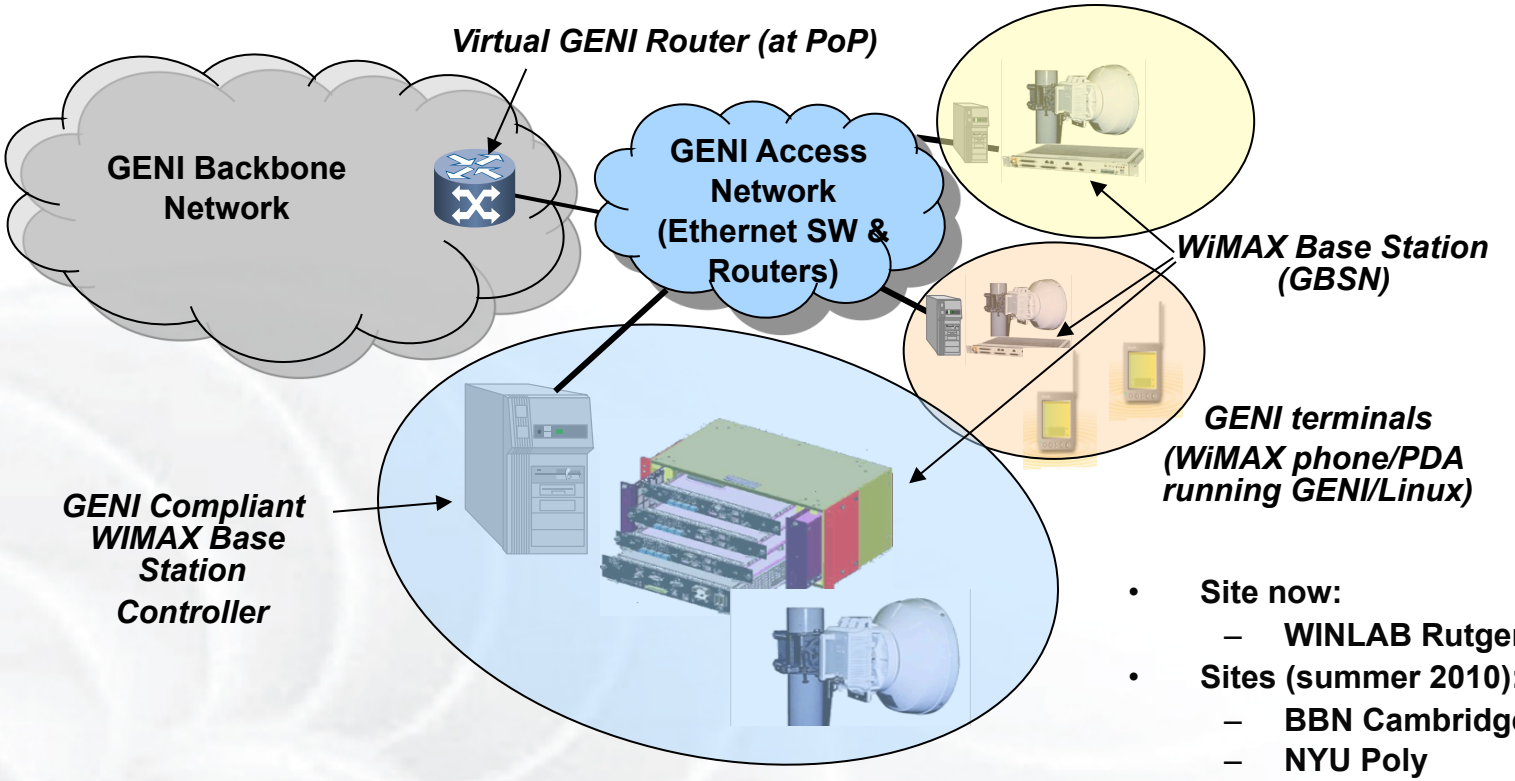

- Site now:
  - WINLAB Rutgers
- Sites (summer 2010):
  - BBN Cambridge
  - NYU Poly
- Sites (late 2010):
  - Columbia
  - UMass Amherst
  - Univ Wisconsin
  - Univ Colorado Boulder
  - UCLA

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 uplink traffic
- Slice user starts a new slice, adds the mobile to its slice
- Datapath switch from (Mobile – VM0) ➔ (Mobile – VM1)

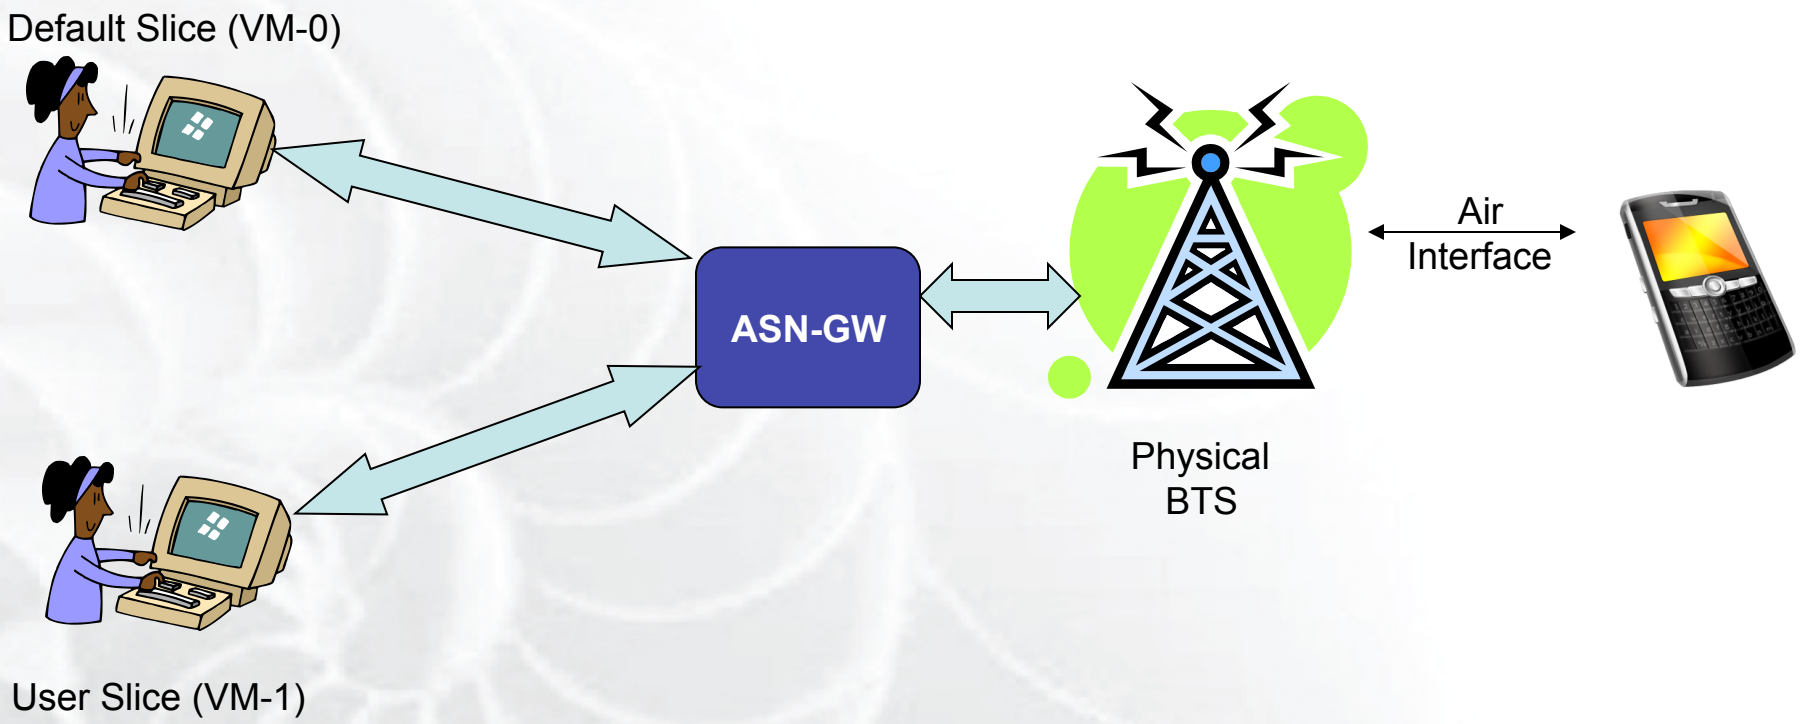

PlanetLab & SPP

- PlanetLab Central: 1000+ nodes worldwide
- SPP: programmable router in 5 Internet2 PoPs
- Other sites running local versions of PlanetLab:
  - GpENI high-speed network in Kansas
  - BBN

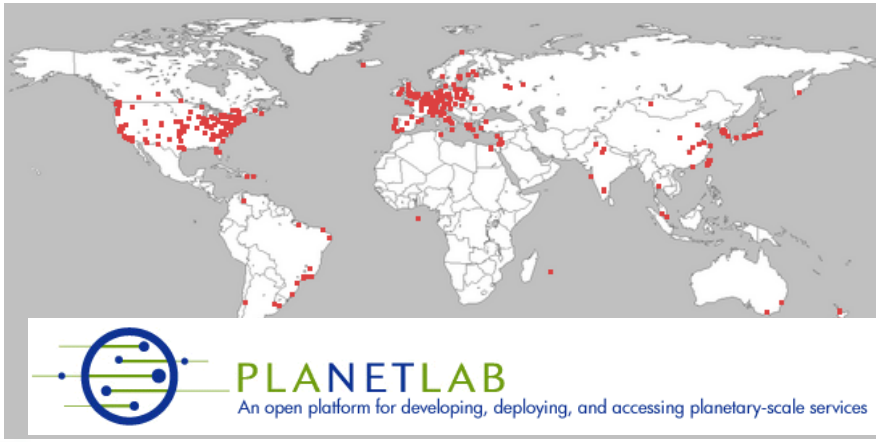

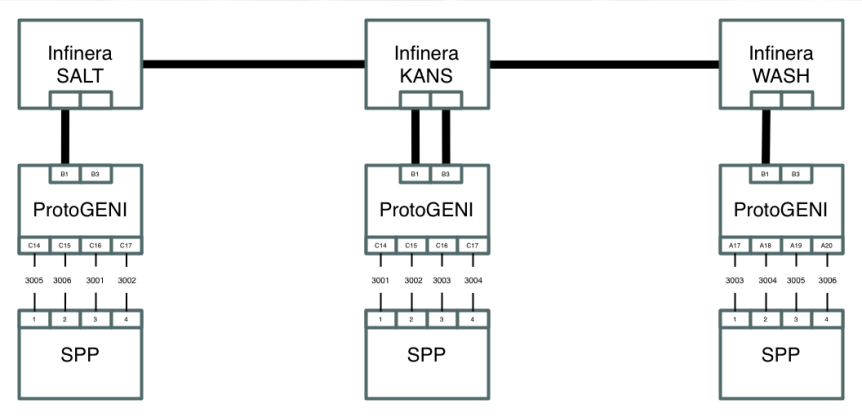

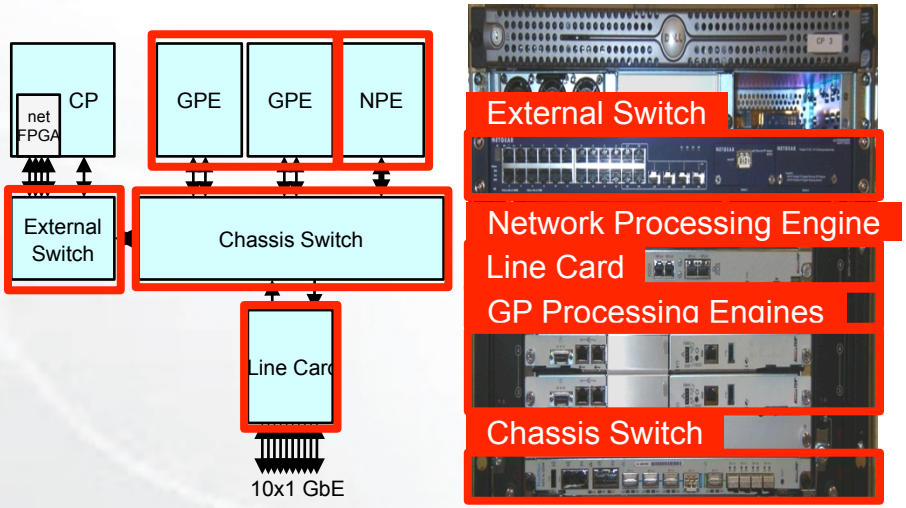

OpenFlow Campus: Stanford GENI Network

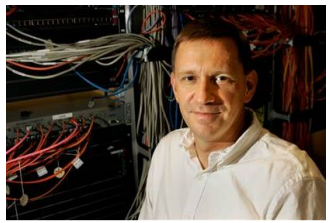

Nick McKeown, PI

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/ sActiveport
- WiMAX deployment

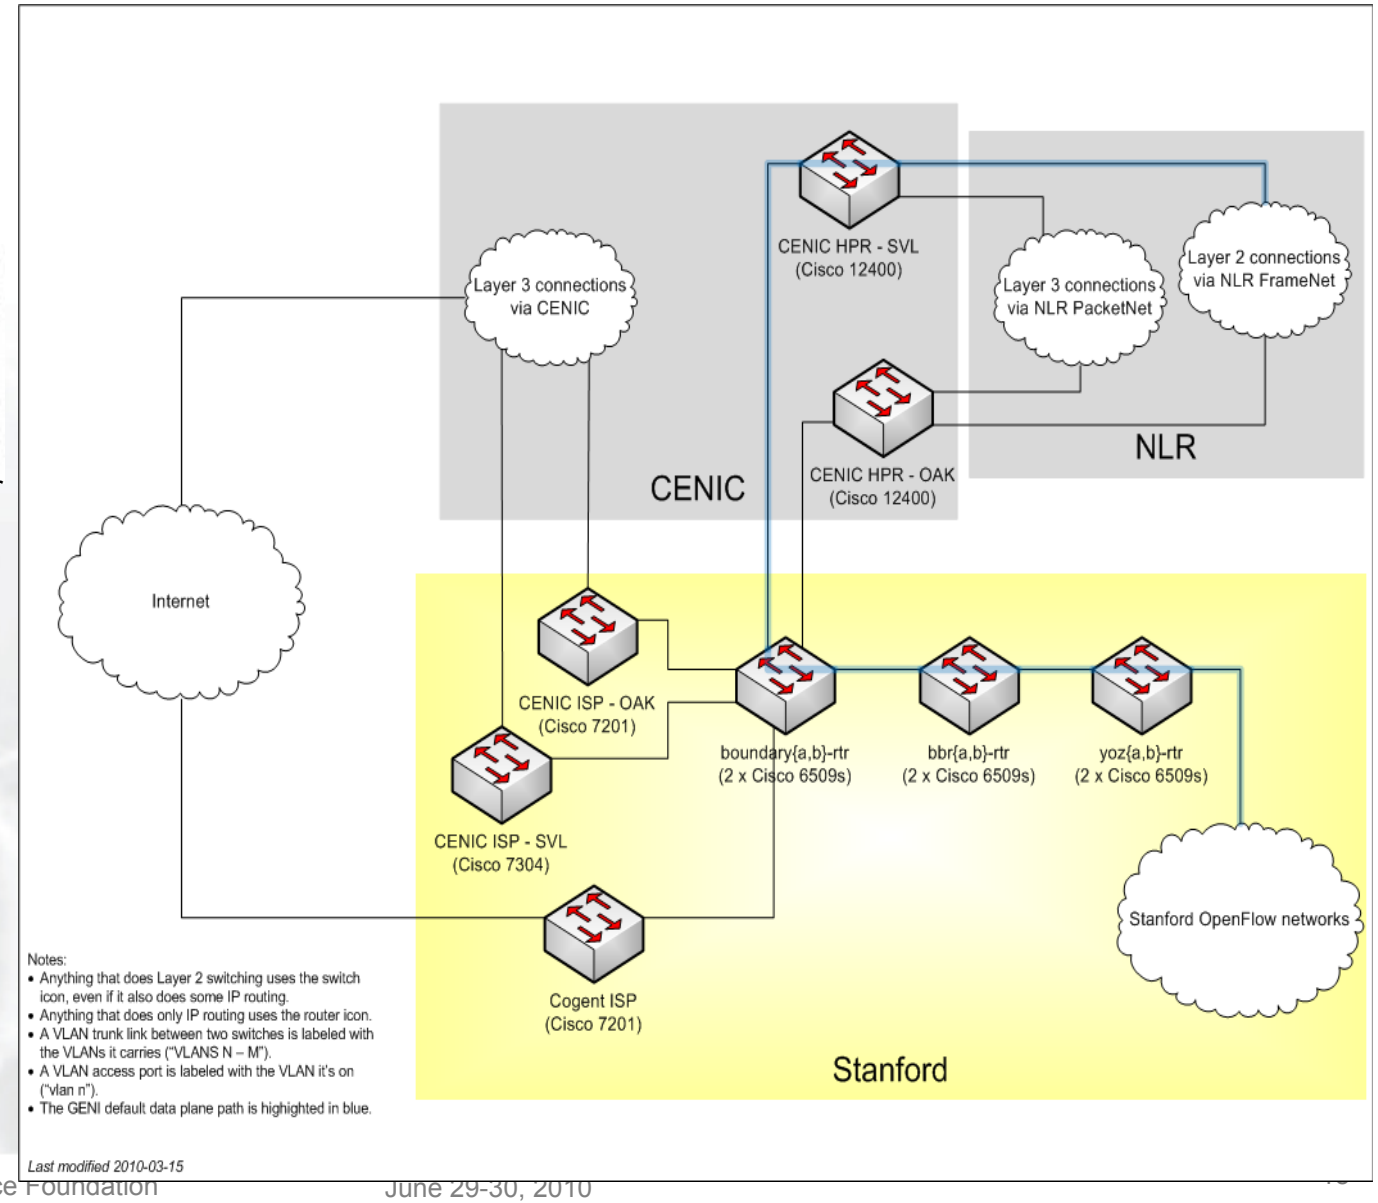

OpenFlow Campus: Rutgers GENI Network

Ivan Seskar, PI

Dipankar Raychauduri

Richard Martin

Charles Hedric, OIT

- OpenFlow and WiMax in ORBIT wireless testbed now
- Campus OF and WiMAX deployments
- NetFPGA OF links
- Connections to Australia (NICTA)

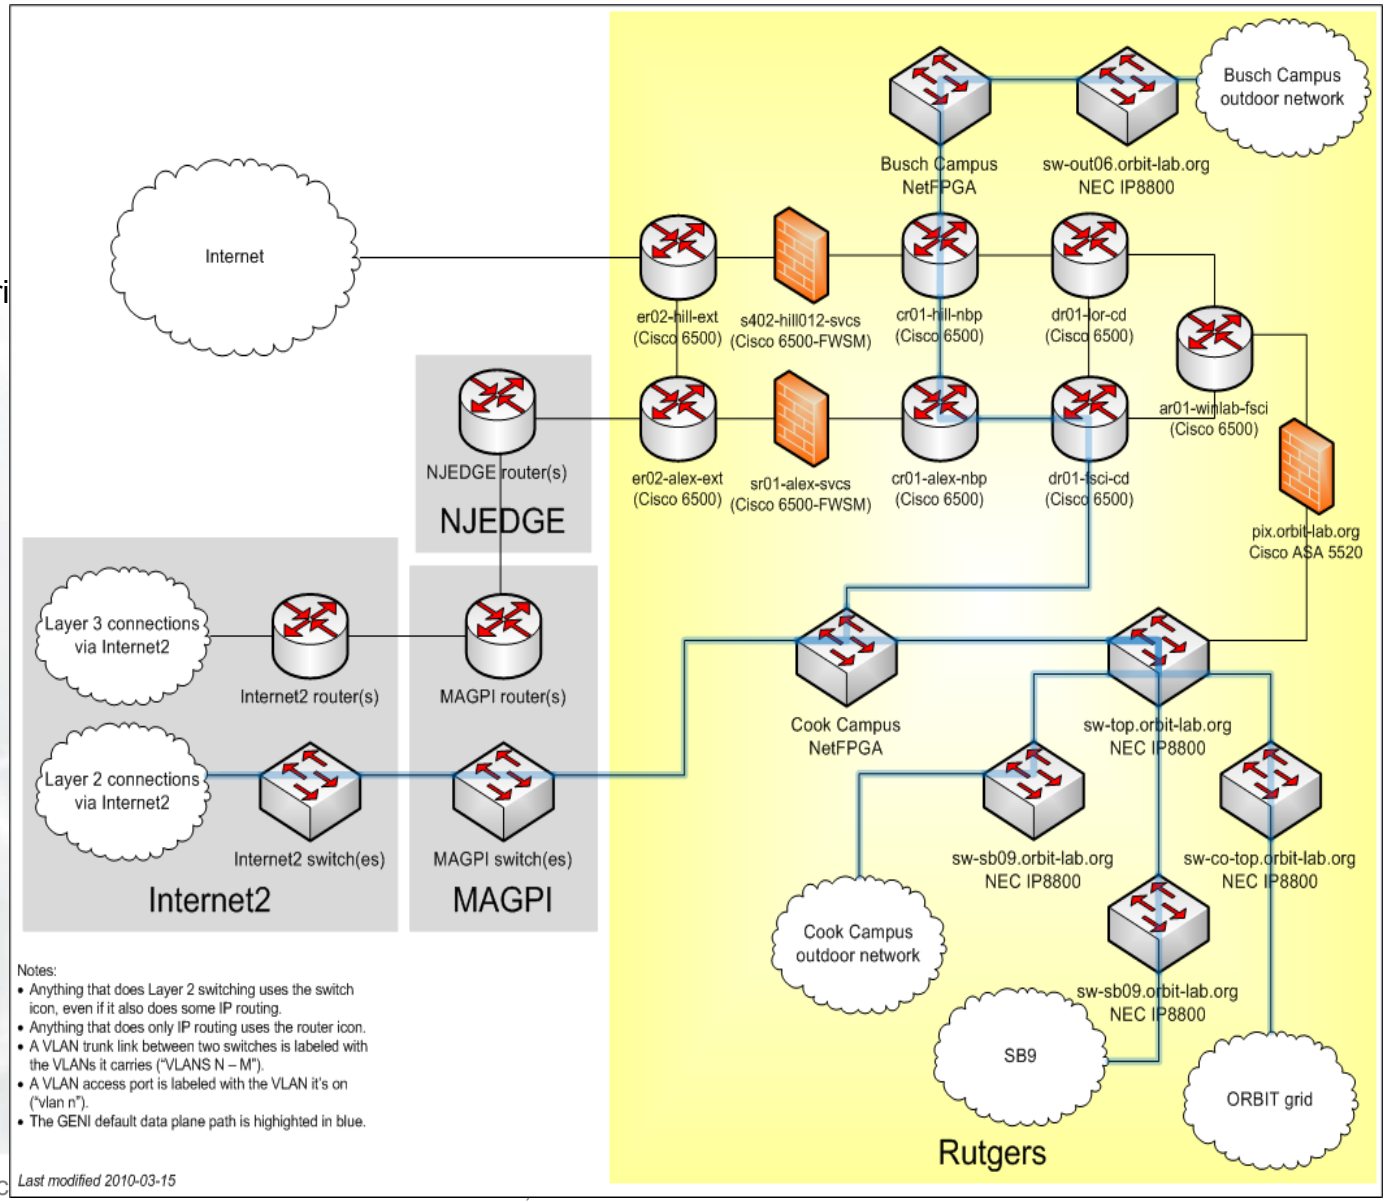

Sponsored by the National Scienc

OpenFlow Campus: BBN GENI Network

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

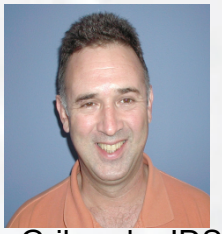

Mike Gribaudo, IDSG

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

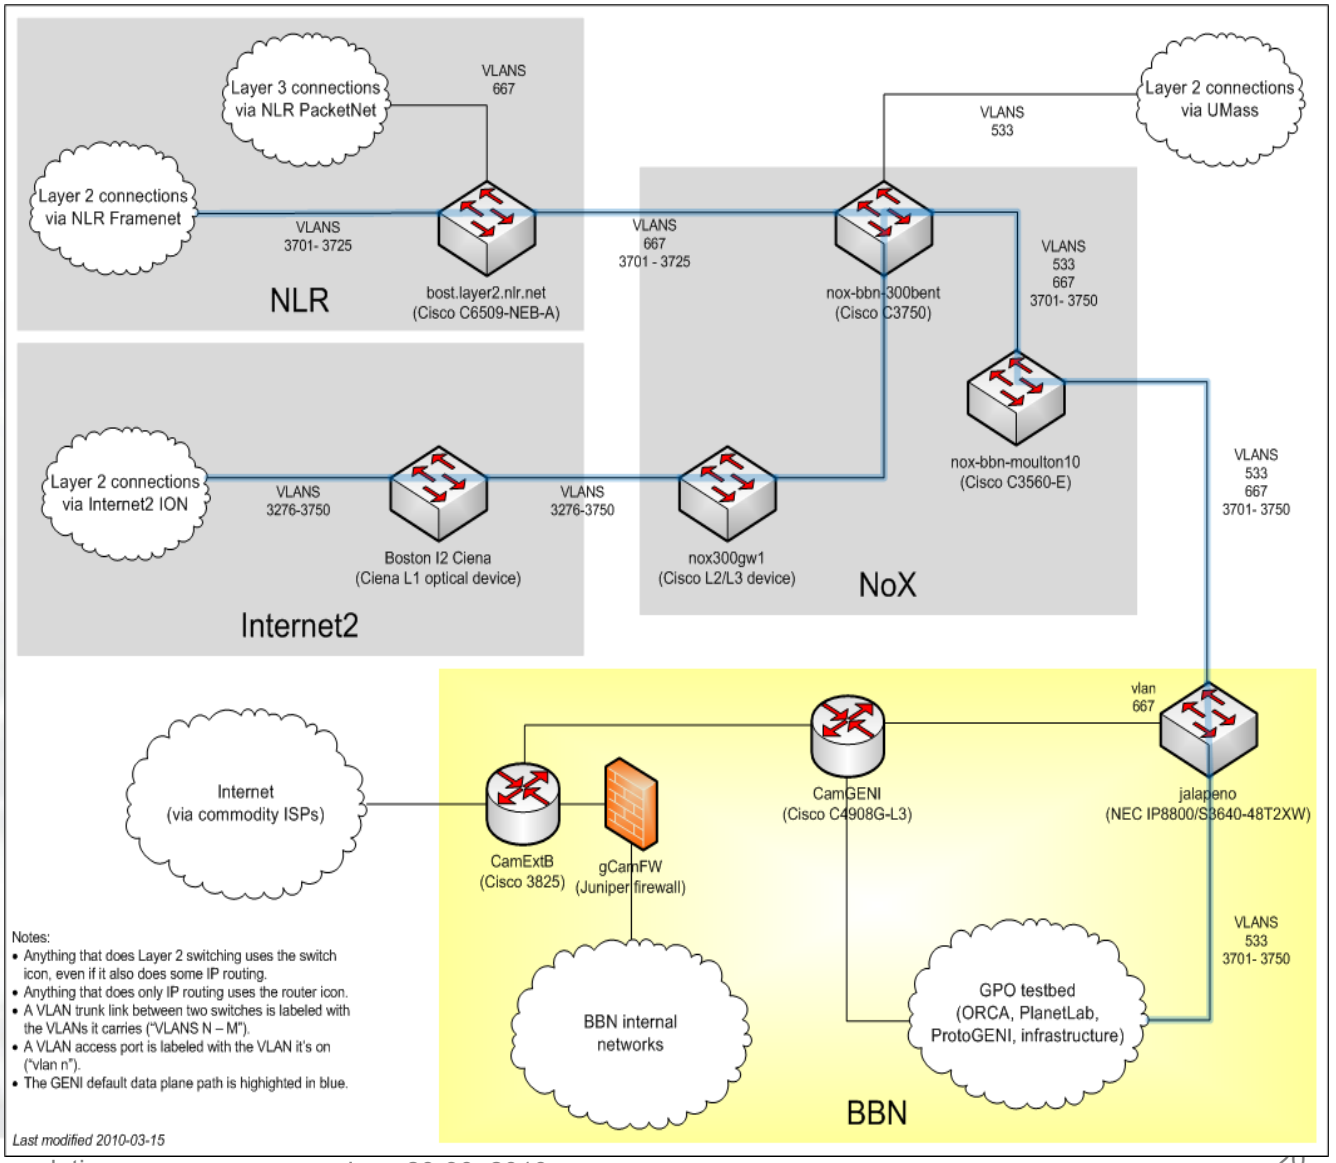

ORCA/BEN

- ORCA: a broker-capable control framework
- BEN: highly-configurable optical network
- Integrated with NLR backbone
- Cluster also has integrated
  - ViSE steerable weather radars
  - DOME VMs on city buses
  - Kansei mote network

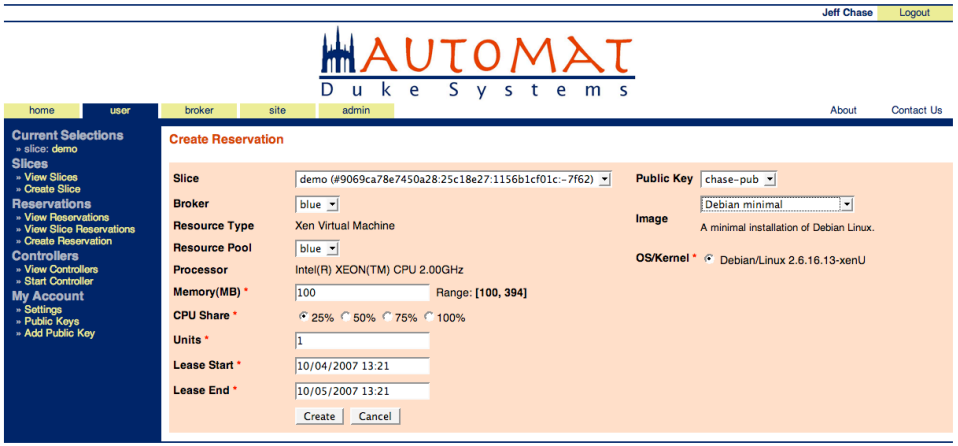

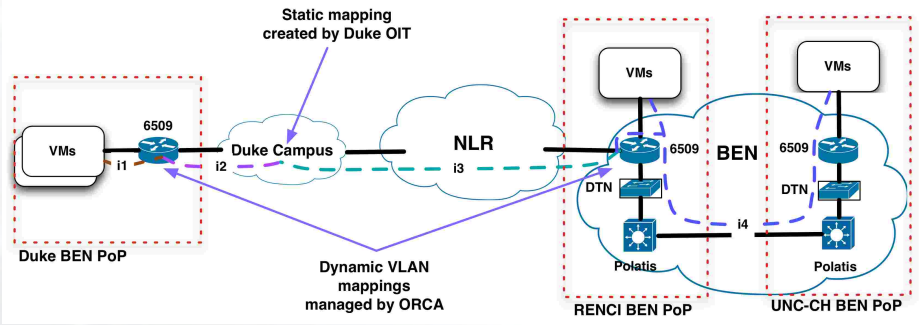

- Our advice on how you should proceed:

- Pick one control framework
  - Take advantage of the common tools, experience

- GPO can help
  - Advice on best match to your goals
  - Establishment of end-to-end VLANs
  - Some software support

Outline

- Introduction
- Some Simple Examples
- Resources
- Being a Pioneer
- Appendix: Concepts & Terminology

Evolution

- Today:
  - GENI backbones connect ProtoGENI, SPP, BEN
  - Other resources connect via IP using tunnels as needed
  - Four control frameworks
  - Manual stitching of end-to-end VLANs; GPO-assist needed for backbones
  - Limited tools for discovery, management, measurement
- By Fall 2010
  - Small number OpenFlow networks on multiple campuses connected to Internet2 & NLR backbones
  - WiMax at more 3+ locations
  - PlanetLab, ProtoGENI, and OpenFlow control framework interoperability
  - Improved tools
- In 2011
  - Prototype I&M system
  - Broader control framework interoperability
  - End-users

Early Integration Focus: Spiral 2 Control Framework Clusters

	PlanetLab Cluster (B)	ProtoGENI Cluster (C)	ORCA Cluster (D)	ORBIT Cluster (E)
Control Framework Design and Prototyping	PlanetLab	ProtoGENI DigitalObjectRegistry PGAugmentation	ORCA/BEN ORCA Augmentation	ORBIT
Network Aggregate Design and Prototyping	Mid-Atlantic Crossroads GpENI	BGPMux CRON PrimoGENI	ORCA/BEN iGENI LEARN	
Programmable Network Node Design and Prototyping	EnterpriseGeni Internet Scale Overlay Hosting	CMULab ProgrammableEdgeNode		
Compute Aggregate Design and Prototyping	GENICloud	MillionNodeGENI	Data Intensive Cloud Control	
Wireless Aggregate Design and Prototyping		CMULab	DOME ViSE KanseiSensorNet OKGems	ORBIT WiMAX D&P COGRADIO
Instrumentation & Measurement Design and Prototyping	VMI-FED	InstrumentationTools MeasurementSystem OnTimeMeasure LAMP ScalableMonitoring	ERM LEARN IMF	
Experiment Workflow Tools Design and Prototyping	GushProto ProvisioningService (Raven) netKarma SCAFFOLD	PGTools		

From http://groups.geni.net/SpiralTwo

GENI-enabled Compute Nodes

Resource	Available number	Host Institution	Notes
PlanetLab nodes	200 +	Planet Lab Consortium and participants (100+ US locations)	[1]
ProtoGENI backbone nodes	5	U of Utah, Internet2	3 PoPs now, 2 more in August.
ProtoGENI host nodes	500+	U of Utah, U of Kentucky, plus several additional sites	[1]
Home/office computers (P2P hosting platform)	TBD	U of Washington and volunteer participants	[2]
SPP nodes	3 now +2 planned	Washington U, St. Louis and Internet2	[2], training required.
Programmable Edge Node (virtual routers)	1	U of Massachusetts, Lowell	
Eucalyptus cluster nodes (cloud computing)	32	HP Labs Palo Alto	[2]. Available in September.

Resources available now except where noted.

[1] Common programmatic interface to be available through GENI Aggregate API v1.0.
[2] May need custom configuration or review / approval by provider. GPO can help.

GENI-enabled Networks

Resource	Available number	Host Institution	Notes
OpenFlow switches in Internet2 and NLR backbones	10+ nodes	I2, NLR	[2] – NLR online in July; I2 online in early fall
Access to resources including OpenFlow networks in eight campuses	TBD	8 campuses (Stanford, Clemson, Georgia Tech, Indiana U, Rutgers, U of Wisconsin, Madison, U of Washington, Princeton) + BBN	[2]
VLANs on NLR shared links (Active to 10Gbps)	varies	NLR	[2]
Internet2 shared VLANs (Active to 1 Gbps)	varies	Internet2	[2]
Regional VLANS	varies	various throughout the US	[2]
Breakable Experimental Network testbed (optical)	1	RENCI, Duke	
DRAGON Testbed (GMPLS)	1	U of Maryland	Mid-Atlantic region testbed
Great Plains Environment for Network Innovation (GpENI Testbed)	1	U of Kansas, U of Missouri, U of Nebraska, Kansas State, KanREN, Lancaster U, ETJ Zurich	
ORBIT wireless Testbed and WiMAX deployment	1	Rutgers U	
TIED testbed (DETER)	1	ISI	
BGP Multiplexer	5	Georgia Tech	Some integration with VINI.
CMULab Wireless Link Emulator	1	Carnegie Mellon U	

GENI Backbone Connectivity Status

Campus	Location	Regional	I2 GENI Wave	I2 ION	NLR FrameNet	Other	Public Internet
BBN	Cambridge, MA, USA	NoX	planned	UP	UP	planned	planned
UMass Amherst	Amherst, MA, USA	NoX	planned	--	UP (via BBN)	--	--
Univ. of Utah	Salt Lake City, UT, USA	UEN	planned	--	--	--	--
RENCI	Chapel Hill, NC, USA	BEN	--	--	UP	--	--
I2 SALT Node	Salt Lake City, UT, USA	--	UP	--	--	--	--
I2 KANS Node	Kansas City, MO, USA	--	UP	--	--	--	--
I2 WASH Node	McLean, VA, USA	--	UP	--	--	--	--
Stanford	Stanford, CA, USA	CENIC	--	--	UP	--	Planned (GRE)
Rutgers	North Brunswick, NJ	MAGPI	--	UP	--	--	--
Indiana	Bloomington, IN, USA	GigaPop	--	planned	planned	--	planned
Clemson	Clemson, SC, USA	SoX, SCLR	--	planned	--	--	planned
Georgia Tech	Atlanta, GA, USA	SoX,SLR	--	planned	--	--	planned
U. Washington	Seattle, WA, USA	P.N. gigapop	--	planned	--	--	--
U. Wisconsin	Madison, WI, USA	WscNet, GPN, OmniPop, SCLR	--	planned	--	--	--
Princeton	Princeton, NJ, USA	Paetec, MAGPI	--	--	--	--	planned

GENI-enabled International Connections

Resource	Available number	Host Institution	Notes
South Korea (1 Gbps)	1	ETRI/KISTI Indiana U	Integration trials for network management and operations only
Various European locations, Active to 10GBE	TBD	Great Plains Network, StarLight	[2]
Australia/US VLANS	TBD	Rutgers, NICTA, Internet2	[2]

GENI Experimental Tools

Resource	Available number	Host Institution	Notes
GUSH Experiment Control and Management Tool	TBD	Williams College	
Raven Provisioning Tool	TBD	U of Arizona	
MetaVPN (OpenVPN dynamic tunnel mgr)	TBD	Carnegie Mellon U	
Digital Object Registry Service	1	CNRI	
PRIME real-time network simulator	1	Florida International U	Available in July

GENI Measurement Services

Resource	Available number	Host Institution	Notes
Instrumentation and Measurement System	1	U of Wisconsin Madison	
LAMP (perfSONAR)	1	U of Delaware, Internet2	Available to GENI users through ProtoGENI in July
OnTime Measure (on-demand measurement system)	1	Ohio SActiveercomputer Center	Available to GENI users through ProtoGENI in June
S3 Measurement service	1	Purdue U, HP Labs	Available to GENI users through ProtoGENI in July
GMOC Operations data collection	1	Indiana U	

Outline

- Introduction
- Some Simple Examples
- Resources
- Being a Pioneer
- Appendix: Concepts & Terminology

GENI Is a Virtual Laboratory

- 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

Your Ideas for Experiments

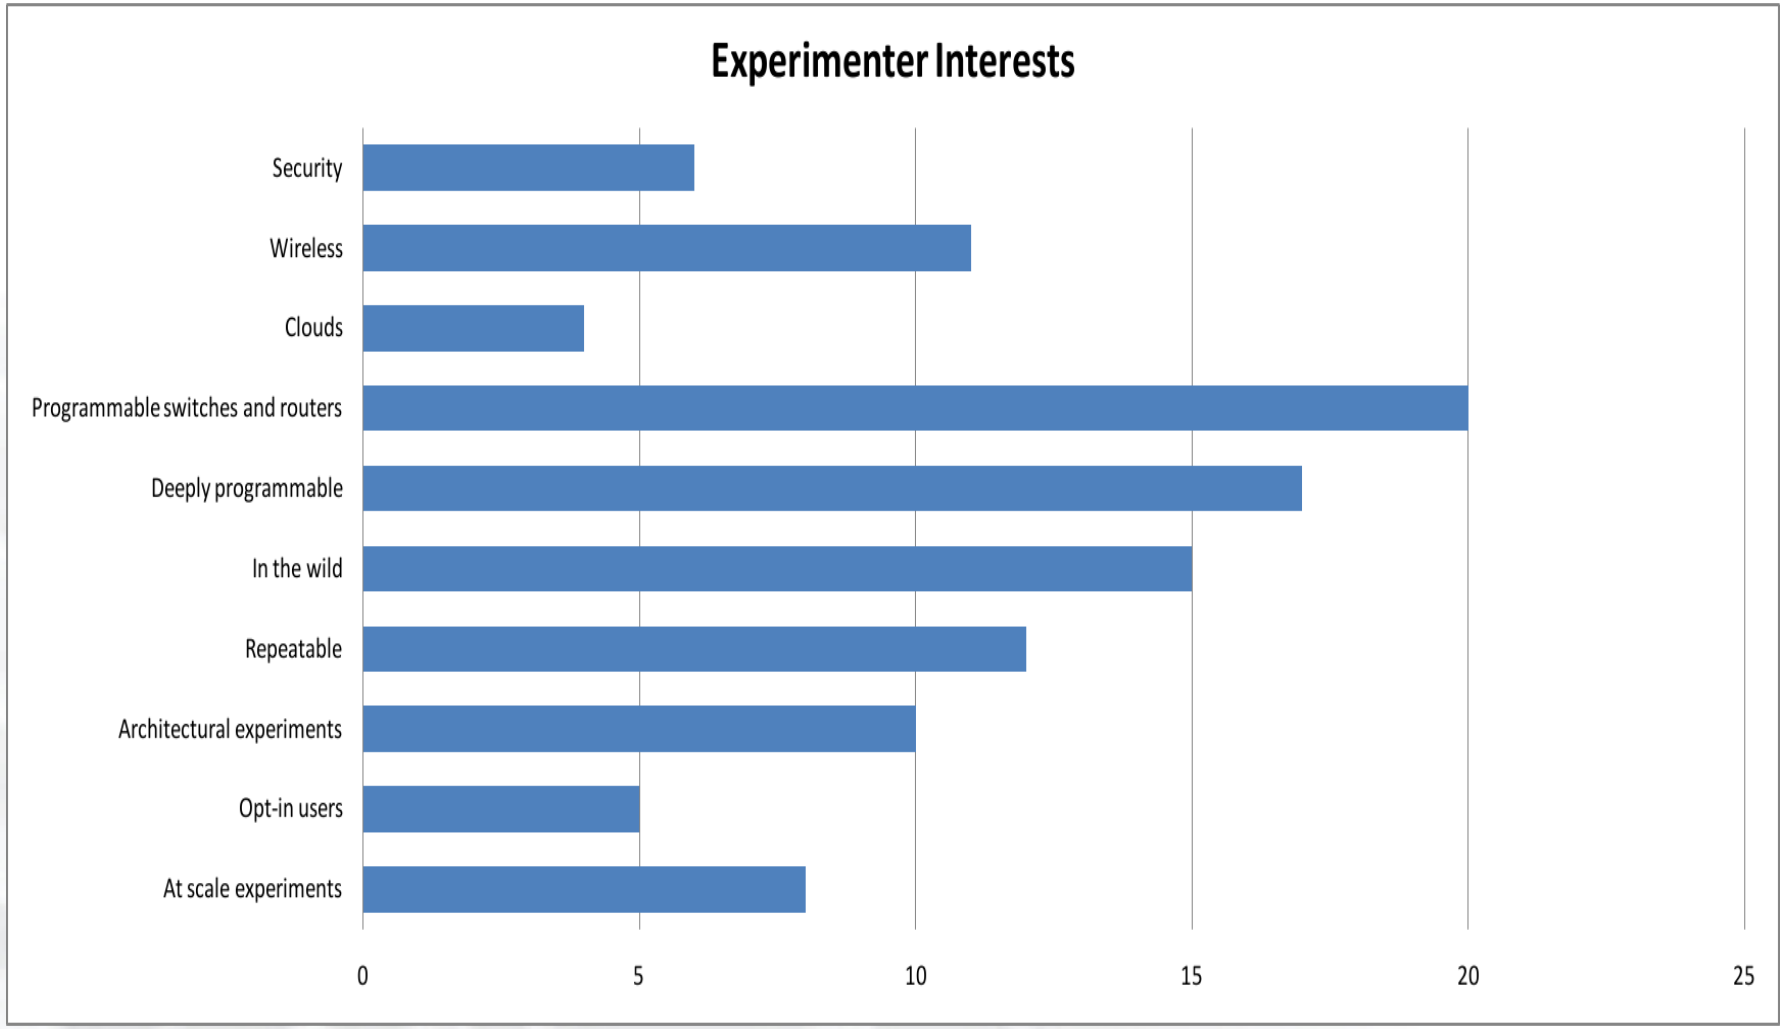

Experiments Guide GENI Development

- GENI needs your feedback
  - As experimenters, you are the GENI user community
  - What works? Doesn't work? Hasn't been built yet?
- GENI Solicitation 3 addresses some key needs
  - Place more GENI-enabled switches in backbone and regional networks
  - Additional WiMax deployments
  - "GENI racks" for increased in-network storage and computation
  - Instrumentation
  - Experiment Support

Conclusion

- GENI is entering an exciting phase!
- Nobody's done this before
- The GPO is here to help

Outline

- Introduction
- Some Simple Examples
- Resources
- Being a Pioneer
- Appendix: Concepts & Terminology

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 Manger
  - 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