GENI: Opening up new classes of experiments in global networking

What is GENI?

The Global Environment for Network Innovations – GENI – is a suite of research infrastructure rapidly taking shape in prototype form across the United States. It is sponsored by the National Science Foundation, with the goal of becoming the world's first laboratory environment for exploring future internets at scale, thereby promoting innovations in network science, security, technologies, services and applications.

GENI will allow academic and industrial researchers to perform a new class of experiments that tackle critically important issues in global communications networks:

- Science Issues - We cannot currently understand or predict the behavior of complex, large-scale networks.
- Innovation Issues – We currently face substantial barriers to innovation with novel architectures, services, and technologies.
- Society Issues – We increasingly rely on the Internet but are unsure that we can trust its security, privacy, or resilience.

GENI will support two major types of experiments: (a) controlled and repeatable experiments, which will greatly help improve our scientific understanding of complex, large-scale networks; and (b) "in the wild" trials of experimental services that ride atop or connect to today's Internet, engaging large numbers of human participants. GENI will provide extensive instrumentation for both forms of experiments, as well as the requisite data archival and analysis tools.

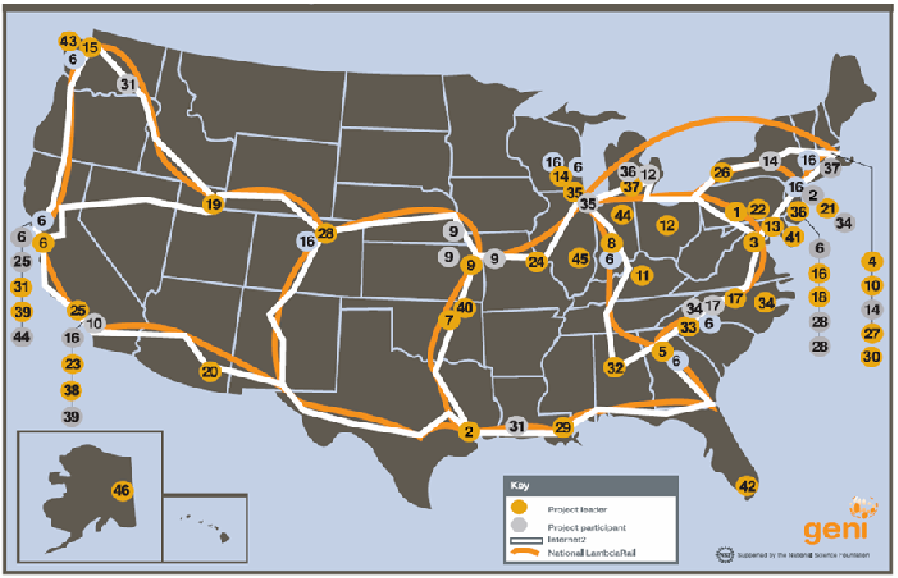

Figure 1: GENI Spiral 2 Sites

Building GENI via rapid prototyping

GENI is being created as a series of rapid prototypes via "spiral development" so that hands-on experience with early experimentation and trials can drive its evolution. Planning and prototyping of GENI brings together capabilities of PlanetLab, Emulab, OpenFlow, ORBIT, ORCA, and DETER, plus a variety of other, innovative research and measurement tools. Industrial research teams including AT&T, CA Labs, HP Labs, IBM, NEC, and

SPARTA are also engaged with academic researchers, as are Internet2 and NLR, the two US national research backbones, and several regional optical networks.

Two basic GENI concepts are "GENI-enabling" and "federation." The first involves enhancing experimental infrastructure, including existing testbeds, campuses networks, regional and backbone networks, cloud computation services, experimental switches and commercial equipment, so that they can be identified and incorporated into experiments via a common access mechanism ("control framework"). The second refers to the process by which heterogeneous types of infrastructure are incorporated into GENI. Federation is key to the ability to grow GENI by incorporating existing GENI-enabled infrastructure.

It is envisioned that "at-scale" experimentation of new technologies, protocols and services will ultimately grow to involve large numbers of human participants and computers which, if successful, may be smoothly migrated out of the GENI infrastructure and into production use. In October 2009, the GENI project started the process of paving the way for such experiments by initiating a "meso-scale" ("intermediate-scale") build-out involving more than a dozen US campuses, two national backbones, and several regional networks. It is expected that this step will provide a path towards a more substantial build-out.

Key GENI Concepts

- **Deep programmability** is the capability to understand and program behavior deep inside the network, not just at the network edge. GENI capabilities include a number of instrumented and programmable devices physically co-located and interconnected at key points in major network backbones; known as Points of Presence (PoPs). These devices create rare opportunities for novel experiments.

- **Slicing** is the assembly of a diverse collection of resources for use in an experiment. GENI uses a variety of techniques, including end-to-end network and computer virtualization, along with combinations of dedicated and shared hardware resources to create slices suitable for the needs of a particular experiment. For example, a slice created for an experiment with strong interest in reproducibility may rely on dedicated resources and VLANs for increased isolation. By contrast, another experiment's slice may seek higher exposure to "real world" conditions through direct connection to the commercial Internet.

- **Control frameworks** are interfaces and tools available to the experimenter to create and manage GENI slices. By accessing GENI through a control framework, the experimenter has a unified API and/or GUI that facilitates identification, reservation, configuration, and interconnection of diverse GENI resources, significantly reducing the burden of experiment setup.

What GENI resources are available for your experiment?

Because of the rapid pace of GENI spiral development, the suite of GENI resources is constantly growing. GENI resources that are currently available or expected to be soon available for use by experiments are geographically distributed throughout the nation and include the following (see Appendix for additional details):

- Tools for end-to-end experiment development
    - Control frameworks that facilitate identification, reservation, configuration, and interconnection of diverse GENI resources
    - Graphical and API-based experimenter tools supporting experiment design, configuration, and deployment
- High-speed connectivity, at rates up to 10 Gbps, and programmable devices at multiple points within the Internet2 and National Lambda Rail (NLR) national backbones and at many endpoint site networks. Because these resources are located within the network, they support experiments requiring deep programmability. Typical resources may include

    - Supercharged PlanetLab Platform (SPP): a high performance programmable router, developed at Jon Turner's lab (Washington University, St. Louis), compatible with PlanetLab, and providing sliceable fast and slow paths co-located at PoPs.
    - NetFPGA: an open hardware platform for prototyping high-speed, hardware-accelerated networking systems.
    - OpenFlow networks to enable innovative network architecture capabilities in control, management, and data planes.
- Computing and network resources from ongoing GENI spiral 2 deployments at over a dozen leading research universities nationwide. Resources may include:
    - OpenFlow networks to enable innovative network architecture capabilities in control, management, and data planes.
    - A programmable, sliced WiMax base station, enabling long-range broadband connectivity to mobile devices.
    - General-purpose compute resources (computers, cluster, and/or cloud), suitable for application hosting, data collection, etc.
- Testbeds
    - PlanetLab distributed testbed,
    - ProtoGENI: Emulab-based network and distributed computing testbeds
    - OpenFlow networks at 8 campuses interconnected by OpenFlow networks in Internet2 and NLR
    - CMULab wireless networking testbed
    - ORBIT wireless networking testbed
    - Diverse Outdoor Mobile Environment (DOME): virtualized mobile networking environment deployed on a regional bus system
    - Breakable Experimental Network (BEN): programmable optical network experiment environment
    - KanseiSensorNet: Extreme Scale Motes (XSM) based sensor network testbed
    - ViSE: outdoor wide-area sensor/actuator network testbed
- Instrumentation
    - Initial instrumentation system based on passive packet capture at selected points in the GENI network.
    - Access to instrumentation capabilities provided by backbone network, including NLRView (NLR) and perfSONAR (Internet2).
    - Access to regularly collected network statistics such as traffic and routing log data.
    - Specialized data collection configuration and/or software provided by the experimenter.

Although an individual experiment is likely to use only a modest subset of the diverse GENI resources, GENI's slices and control frameworks promise significant benefits, particularly for experiments that will profit from executing under a number of different resource configurations or scales.

Potential research areas

GENI seeks to open up new areas of experimental research at the frontiers of network science and engineering – fields with significant potential for socio-economic impact. These experiments may be fully compatible with today's Internet, variations or improvements on today's Internet protocols, or indeed radically novel "clean slate" designs.

Research interests, and thus experiments, will evolve significantly over the coming decade. Early types of GENI-based experiments might involve research in the following areas. This list is not all-inclusive; decisions on experiments using GENI are up to you, the researcher.

- Content distribution services – As the Internet is increasingly used to distribute high-bandwidth content (e.g. video and virtual worlds), many researchers have focused on new, more scalable architectures for such

services. GENI is well-suited to such experiments with its emphasis on deep programmability, clouds, and GENI-enabled campuses.

- Disruption Tolerant Networks (DTNs) – GENI is specifically aimed to enable large-scale, well-instrumented, repeatable experiments on novel protocols and architectures. DTNs are a perfect case in point since many DTN architectures are independent of today's TCP/IP architecture. We expect several DTN experiments to begin on GENI within the coming months.
- Novel mobility architectures – Many networking researchers have proposed novel protocols to improve support for mobile devices in the Internet architecture. GENI's near-term emphasis on wireless support throughout campuses allows real-world experimentation with these new protocols.
- Novel routing architectures – As concerns have grown over the scalability of the global Internet routing architecture, particularly with the rise of multi-homing, a number of research teams have proposed alternative global routing architectures. Although GENI will not be "as big as the Internet," it may offer sufficient scalability so that such approaches can be tried out in a realistic, well-instrumented suite of infrastructure.
- Reliable global networks – All of us increasingly rely on the Internet, but are increasingly uncertain that we can trust its security, privacy or resilience. There is now growing interest in experimental efforts that will help ensure an Internet that is solid and reliable in the critically important role it now plays for society.
- Experimental methodologies – Because of the difficulties in conducting at-scale networking research, the community lacks a rich set of common methodologies, standards, and benchmarks for experiment design. GENI is well positioned to support the development and validation of such methodologies and their promulgation through research and educational channels.
- Experimental exploration of theoretical models - As the network science and engineering discipline works to develop theoretical models of networks, challenges arise in determining the predictive power of these models and understanding their large-scale implications. GENI is a powerful venue for the study of these models, particularly their ability to support dynamic, non-engineered networks, such as those arising from societal and economic pressures in P2P and social networks.
- Network science and design - Gaining confidence in a network design is a challenging problem, either at the level of architecture and protocol design or at the level of validating the design of a specific network instance. GENI provides an opportunity for theoretical / experimental collaborations where network design principles may be amenable to validation through a combination of mathematical and experimental analysis.
- Social and behavior aspects of global networks - Exploration of the social, economic, and legal aspects of future global communications systems.
- Virtualization architectures – GENI's own architecture is based on end-to-end virtualization, which is now becoming an area of keen interest and study to networking researchers. Indeed, GENI prototyping teams are actively experimenting with new network architectures based on virtualization; in addition, we expect future virtualization experiments to run within the GENI infrastructure as it comes online.

Appendix: Available GENI Resources (expected summer 2010*)

GENI is a suite of research infrastructure rapidly taking shape in prototype form across the United States, with additional capabilities becoming available on a frequent basis. Diverse GENI capabilities are composed into interoperable slices accessed through the GENI Aggregate API. Version 1.0 of the API, available in summer/fall 2010, supports interoperation of PlanetLab, ProtoGENI, and OpenFlow, with interoperation of additional capabilities expected in future releases.

Resource	Available number	Host Institution	Notes
GENI-enabled Compute Nodes			
PlanetLab nodes	200 +	Planet Lab Consortium and participants (100+ US locations)	[1]
ProtoGENI backbone nodes	5	University of Utah, Internet2	3 nodes now, 2 more in August.
ProtoGENI host nodes	500+	University of Utah, University of Kentucky, plus several additional sites	[1]
Home/office computers (P2P hosting platform)	TBD	University of Washington and volunteer participants	[2]
SPP nodes	3 now +2 planned	Washington University, St. Louis and Internet2	[2], training required.
Programmable Edge Node (virtual routers)	1	University of Massachusetts, Lowell	
Eucalyptus cluster nodes (cloud computing)	32	HP Labs Palo Alto	[2]. Available in September.
GENI-enabled Networks			
Up to 1 GbE dedicated links between ProtoGENI nodes in Internet2 core network	TBD	University of Utah	
OpenFlow networks in Internet2 and NLR backbones	10+ nodes	I2, NLR	[2]
Access to resources including OpenFlow networks in eight campuses	TBD	8 campuses (Stanford, Clemson, Georgia Tech, Indiana University, Rutgers, University of Wisconsin Madison, University of Washington, Princeton)	[2]
NLR shared links (up to 10Gbps)	TBD	NLR	[2]
Internet2 shared VLANs (up to 1 Gbps)	TBD	Internet2	[2]
Regional VLANS	TBD	various throughout the US	[2]
Breakable Experimental Network testbed (optical)	1	RENCI, Duke	Available to GENI users in April
DRAGON Testbed (GMPLS)	1	University of Maryland	Mid-Atlantic region testbed as well as end-to-end connections with GENI participants
Great Plains Environment for Network Innovation (GpENI Testbed)	1	University of Kansas, University of Missouri, University of Nebraska, Kansas State, KanREN, Lancaster University, ETJ Zurich	
ORBIT wireless Testbed	1	Rutgers University	

Resource	Available number	Host Institution	Notes
and WiMAX deployment			
TIED testbed (DETER)	1	ISI	
BGP Multiplexer	1 or more	Georgia Tech	Available in May to GENI participants. Some integration with VINI.
MetaVPN (OpenVPN network manager)	TBD	Carnegie Mellon University	
GENI-enabled International Connections			
South Korea (1 Gbps)	1	ETRI/KISTI Indiana University	Integration trials for network management and operations only
Various European locations, up to 10GBE	TBD	Great Plains Network	[2]
Australia/US VLANS	TBD	Rutgers, NICTA, Internet2	[2]
GENI Experimental Tools			
GUSH Experiment Control and Management Tool	TBD	Williams College	
Raven Provisioning Tool	TBD	University of Arizona	
GENI Measurement Services**			
Instrumentation and Measurement System	1	University of Wisconsin Madison	
LAMP (perfSONAR)	1	University of Delaware, Internet2	Available to GENI users through ProtoGENI in July
OnTime Measure (on-demand measurement system)	1	Ohio Supercomputer Center	Available to GENI users through ProtoGENI in June
S3 Measurement service	1	Purdue University, HP Labs	Available to GENI users through ProtoGENI in July
GMOC Operations data collection	1	Indiana University	
Other			
Digital Object Registry Service	1	CNRI	
PRIME real-time network simulator	1	Florida International University	Available in July
CMULab Emulator	1	Carnegie Mellon University	

* Academic operations only in Summer 2010 (no central operations)
Does not include general resources, such as commercial Internet

** Does not include measurement tools built into control frameworks like ProtoGENI, ORCA etc.

Notes:
[1] Available through GENI Aggregate API v1.0.
[2] May need custom configuration or review / approval by provider. GPO can help.