G E N I

Global Environment for Network Innovations

Spiral 2 Overview

Document ID: GENI-INF-PRO-S2-OV-1.1

June 3, 2010

Prepared by:
The GENI Project Office
Raytheon BBN Technologies
10 Moulton Street
Cambridge, MA 02138 USA

Issued under NSF Grant CNS-0741315

TABLE OF CONTENTS

1 Document Scope

This section describes this document's purpose, its context within the overall GENI document tree, the set of related documents, and this document's revision history.

1.1 Purpose of this Document

This document provides an overview of GENI Spiral 2. It describes this spiral's developmental goals and summarizes the active projects. GENI Spiral 2 runs from October 2009 through September 2010. It is the second phase of exploratory rapid-prototyping that will begin to inform technical and operational plans for the envisioned GENI suite of research infrastructure. Spiral 2's primary goal is to begin moving towards continuous experimentation. Key developments for this period include improved integration of GENI prototypes; architecture, tools, and services enabling experiment instrumentation; a movement towards interoperability across GENI prototypes; and improved approaches towards researcher identity management. All GENI projects are expected to orient their efforts to make progress towards these goals during this spiral. Additionally, Spiral 2 features the launch of a "meso-scale" deployment of GENI-enabled commercially available equipment in over a dozen university campuses and within two national research backbone networks.

GENI prototype developers should read this document to assess their place in the overall project and their progress toward furthering the concrete GENI goals; potential GENI contributors (through proposals to GENI or collaboration) should use this to assess what GENI is, what has been done, and where there are gaps waiting to be filled; and potential GENI users should use this document to understand the current state of the infrastructure and the near-term plans for adding additional capabilities.

1.2 Context for this Document

Figure 1-1 below shows the context for this document within GENI's overall document tree.

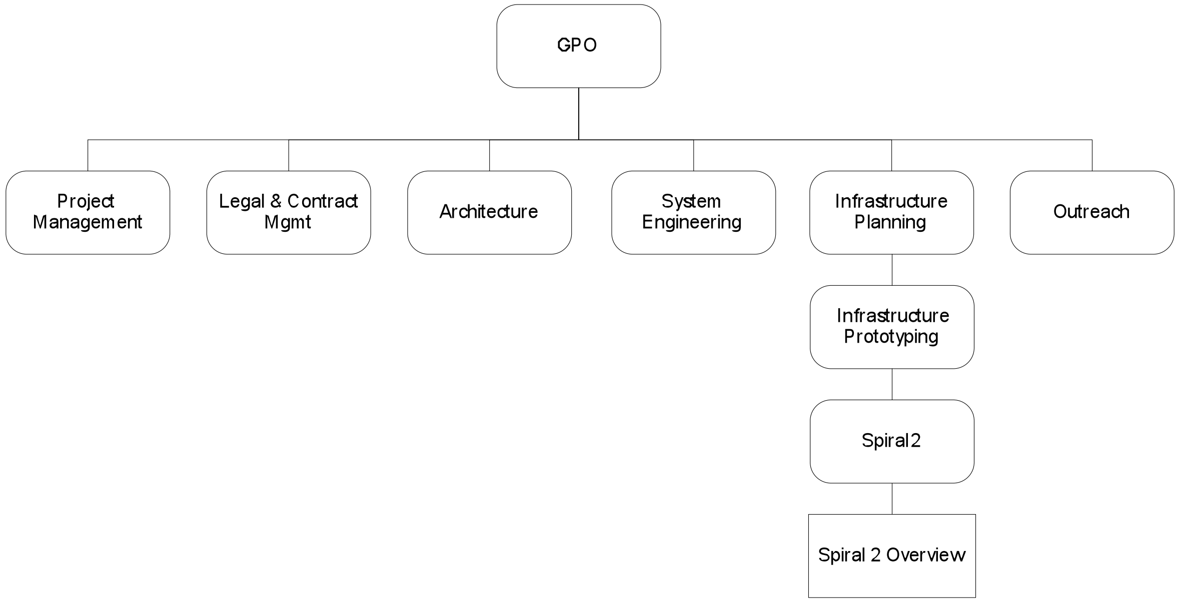

Figure 1-1. This Document within the GENI Document Tree.

1.3 Related Documents

The following documents of the exact date listed are related to this document, and provide background information, requirements, etc., that are important for this document.

1.3.1 National Science Foundation (NSF) Documents

Document ID	Document Title and Issue Date
N / A	

1.3.2 GENI Documents

Document ID	Document Title and Issue Date
GENI-SE-SY-SO-02.0	GENI System Overview (9/29/08)
GENI-INF-PRO-S1-OV-1.12	GENI Spiral 1 Overview (9/29/08)
GENI-SE-SY-RQ-02.0	GENI System Requirements Document (7/7/09)
GENI-SE-SY-SA-TS-S2-00.1	GENI Spiral 2 Security Plan (3/15/10)
GENI-SE-IM-ARCH-0.1	GENI Instrumentation and Measurement Architecture (3/8/10)

1.3.3 Standards Documents

Document ID	Document Title and Issue Date
N / A	

1.3.4 Other Documents

Document ID	Document Title and Issue Date
N / A	

1.4 Document Revision History

The following table provides the revision history for this document, summarizing the date at which it was revised, who revised it, and a brief summary of the changes. This list is maintained in chronological order so the earliest version comes first in the list.

Revision	Date	Revised By	Summary of Changes
1.0	5/27/10	A. Falk	Initial draft.
1.1	6/3/10	A. Falk	Added Spiral 2 map

2 Introduction

GENI Spiral 2 is the second phase of exploratory rapid-prototyping that will inform technical and operational plans for the envisioned GENI suite of network research infrastructure. Each GENI Spiral lasts a year and GENI Spiral 2 runs from October 2009 through September 2010. The primary goal is to help the maturing GENI prototypes migrate towards being capable of supporting continuous experimentation. Key developments for this period include improved integration of GENI prototype aggregates and their associated control frameworks; establishment of the architecture, tools, and services enabling experiment instrumentation; a movement towards interoperability across GENI control framework clusters; and improved approaches towards researcher identity management. All GENI projects are expected to orient their efforts to make progress towards these goals during this spiral.

In the fall of 2009, the GENI Project Office added 33 new design and prototyping projects creating a total of 60 active D&P projects[1] (see Figure 2-1). The new projects were selected to accelerate prototyping of a suite of infrastructure for the GENI project with federation and shakedown experiments that will guide future GENI system design. In addition, the GPO launched a "meso-scale" deployment of several promising GENI technologies into over a dozen locations. GENI "meso-scale" will enable three sets of collaborating academic/industrial research teams to replicate prototype systems that have gained significant traction, based on GENI-enabled commercial hardware, across 13 U.S. campuses and two national research backbones. These prototypes will form a foundation capable of supporting early experiments on an end-to-end suite of GENI infrastructure at a scale significantly larger than has been possible until now.

This document provides an overview of GENI Spiral 2 development. We begin with some background about GENI – its motivation, conceptual design, and a summary of Spiral 1 accomplishments. Then we review the technical goals of Spiral 2 and the resources that we expect to be available to experimenters. The document concludes with a catalog of all active Spiral 2 projects. More information about GENI's development, including Spiral 2 milestones for each project, can be found on the GENI project wiki at http://groups.geni.net.

[1] GENI design and prototyping projects may be funded for multiple years and thus will last for more than one spiral.

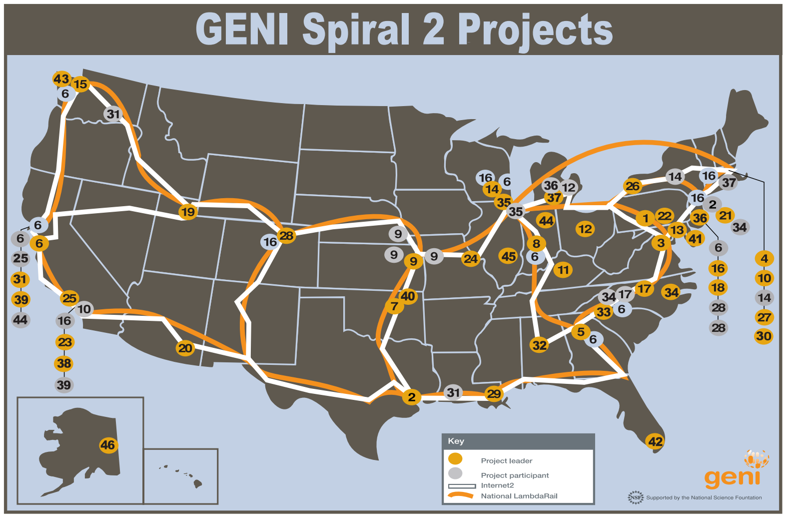

Project Name	Project Lead	Project Participants
1. CMUlab	Carnegie Mellon University	
2. D Meas, LEARN	University of Houston	Columbia University
3. Digital Object Registry	Corporation for National Research Initiatives (CNRI)	
4. CLOUD-CTL, DOME, ViSE	University of Massachusetts Amherst	
5. DTunnels	The Georgia Institute of Technology	
6. EnterpriseGENI, OpenFlow	Stanford University	Princeton University University of California, Berkeley Clemson University Georgia Institute of Technology Indiana University Nicira Networks Princeton University Rutgers University University of Wisconsin University of Washington
7. GENI4YR	Langston University	
8. GMOC, netKarma, K-GENI	Indiana University	
9. GpENI	University of Kansas	Kansas State University, University of Nebraska-Lincoln The University of Missouri-Kansas City
10. GushProto	Williams College	UC San Diego
11. INSTOOLS, ISM Infrastructure	University of Kentucky	
12. KANSEI, OTM	Ohio State University	Wayne State University
13. MAX	University of Maryland	
14. MeasurementSys	University of Wisconsin-Madison	Boston University Colgate University
15. MillionNodeGENI, Security	University of Washington	
16. ORBIT, WiMAX	Rutgers University	UCLA, Los Angeles, CA Columbia University, NY, NY University of Colorado, Boulder, CO Polytechnic University of NYU, Brooklyn, NY University of Massachusetts, Amherst University of Wisconsin, Madison, WI
17. ORCA/BEN	The Renaissance Computing Institute (RENCI)	Duke University
18. PlanetLab, Scaffold, Federation	Princeton University	Universite Pierre et Marie Curie (UPMC)
19. ProtoGENI	University of Utah	
20. PROVSERV	University of Arizona	
21. ERM	Columbia	
22. REGOPT	Pittsburgh Supercomputing Center (PSC)	
23. SECARCH, Distributed Identity	SPARTA, Inc.	
24. SPP	Washington University	
25. TIED	USC Information Sciences Institute	University of California, Berkeley
26. UB_OANets	SUNY Buffalo	
27. UMLPEN	University of Massachusetts Lowell	
28. CR-GENI	University of Colorado Boulder	Radio Technology Systems LLC Rutgers University
29. CRON-T	Louisiana State University	
30. Design of Information Subs	MIT	
31. DSL, HIVE	UC Davis	Batelle CA Labs
32. EXP-SEC	University of Alabama	
33. FPGA-RADIO	Clemson University	
34. GENI IMF	North Carolina State University	The Renaissance Computing Institute (RENCI) Columbia University
35. iGENI	Northwestern University	University of Illinois Chicago
36. LAMP	University of Delaware	Internet2
37. LEFA, Supercharged Planetlab	Internet2	Brown University
38. NLR	Cypress, CA	
39. Open--CIRRUS	HP Labs, Palo Alto	UCSD
40. OKGems	Oklahoma State University	
41. PIGEON-NET	Howard University	
42. PrimoGENI	Florida International University	
43. QUILT	The Quilt	
44. S3-GENI	Purdue University	HP Labs
45. SEC-POL	University of Illinois (NCSA)	
46. VMI	University of Alaska Fairbanks	

Figure 2-1. GENI Spiral 2 Projects

3 GENI Overview

The Global Environment for Network Innovations (GENI) is a novel suite of infrastructure now being designed to support experimental research in network science and engineering.

This new research challenges us to understand networks broadly and at multiple layers of abstraction; from the physical substrates through the architecture and protocols to networks of people, organizations, and societies. The intellectual space surrounding this challenge is highly interdisciplinary, ranging from new research in network and distributed system design to the theoretical underpinnings of network science, network policy and economics, societal values, and the dynamic interactions of the physical and social spheres with communications networks. Such research holds great promise for new knowledge about the structure, behavior, and dynamics of our most complex systems – networks of networks – with potentially huge social and economic impact.

As a concurrent activity, community planning for the suite of infrastructure that will support network science and engineering experiments has been underway for several years. This suite is termed the Global Environment for Network Innovations (GENI). Although its specific requirements will evolve, the initial design is now clear enough to support early planning and prototyping. The core concepts for the suite of GENI infrastructure are as follows:

- **Programmability** – researchers may download software into GENI-compatible nodes to control how those nodes behave;
- **Virtualization and Other Forms of Resource Sharing** – whenever feasible, nodes implement virtual machines, which allow multiple researchers to simultaneously share the infrastructure; and each experiment runs within its own isolated slice, created end-to-end across the experiment's GENI resources;
- **Federation** – different parts of the GENI suite are owned and/or operated by different organizations, and the NSF portion of the GENI suite forms only a part of the overall 'ecosystem'; and
- **Slice-based Experimentation** – GENI experiments will be an interconnected set of reserved resources on platforms in diverse locations. Researchers will remotely discover, reserve, configure, program, debug, operate, manage, and teardown distributed systems established across parts of the GENI suite.

As envisioned in these community plans, the GENI suite will support a wide range of experimental protocols, and data dissemination techniques running over facilities such as fiber optics with next-generation optical switches, novel high-speed routers, city-wide experimental urban radio networks, high-end computational clusters, and sensor grids. The GENI suite is envisioned to be shared among a large number of individual simultaneous experiments, with extensive instrumentation that makes it easy to collect, analyze, and share real measurements.

4 GENI Conceptual Design

To understand how the various projects will fit together, it is useful to briefly review the elements of the GENI conceptual design. This topic is explored in more detail in the *GENI System Overview*[2] and the *GENI System Requirements Document*[3], but we include an abbreviated introduction here for clarity.

It is important to note that this conceptual design is still very high-level and provisional, and may change considerably as we all gain experience. Indeed, the exploratory prototyping being done through the GENI spirals of which this is the second is intended to help drive the GENI design process forward. As the community gains experience with these prototypes, both in building them and in attempting to operate them, this design will become far more concrete. By the *end* of the prototyping stage, several years from now, we should have a very good design; but right now it's still a little too early.

As shown in Figure 4-1, the principal entities in the GENI conceptual design are:

- **Aggregates** are composed of one or more individual **components** under common administrative control. Example aggregates include compute clusters, backbone networks, campus networks, sensor grids, etc. Resources on components are allocated to researchers by way of the GENI **control framework**. Individual components (aggregated or not) may be shared, for example using virtualization, and are expected to be deeply programmable to permit configuration by researchers. The GENI control framework utilizes an **aggregate manager** interface to the control framework to manage authorization and sometimes resource allocation. The use of a common control framework will allow aggregates owned and operated by different organizations to **federate**, allowing access by researchers but retaining local control.
- One or more **clearinghouses** operate **registries** for principals, **slices** and components. GENI researchers obtain credentials from a clearinghouse using an identity management system. Aggregates that have federated with the clearinghouse will recognize these credentials and permit resource allocation. Aggregates also obtain credentials from a clearinghouse (perhaps more than one) when they federate. Clearinghouse registries log resource use so that researchers are accountable for their actions on GENI. Multiple clearinghouse federations may in turn federate (i.e., form a confederation), for example permitting GENI users to access commercial or foreign federations.
- We expect a wide range of GENI-affiliated **experiment support services** to be developed by the GENI community. They may include **storage services** for researchers to archive code, configurations and experiment results; **instrumentation & measurement services**, to make, gather, and archive experiment measurements; and **GENI-Internet gateways**, to permit a controlled exchange of traffic with the Internet..
- **Opt-in End-Users** may join a GENI experiment using a layer 2 VLAN, for example via a campus network, or using IP over the Internet. Additionally, they may contribute resources to GENI such as access to their computers, handsets, or network bandwidth.
- The GENI **Meta-NOC** provides information on the status of GENI aggregates and services such as a **Help Desk** for researchers.

[2] http://groups.geni.net/geni/wiki/GeniSysOvrvw

[3] http://groups.geni.net/geni/wiki/SysReqDoc

Researchers obtain a *sliver* of resources on an aggregate of interest by way of the GENI control framework. The collection of slivers a researcher obtains from one or more aggregates forms a *slice* that is then *stitched* together to obtain end-to-end connectivity. Experiment tools and services are used to remotely discover, reserve, configure, program, debug, operate, manage, and teardown the resources in the slice. Virtualized components can provide resources for multiple experiments at the same time, but keep the experiments isolated from one another.

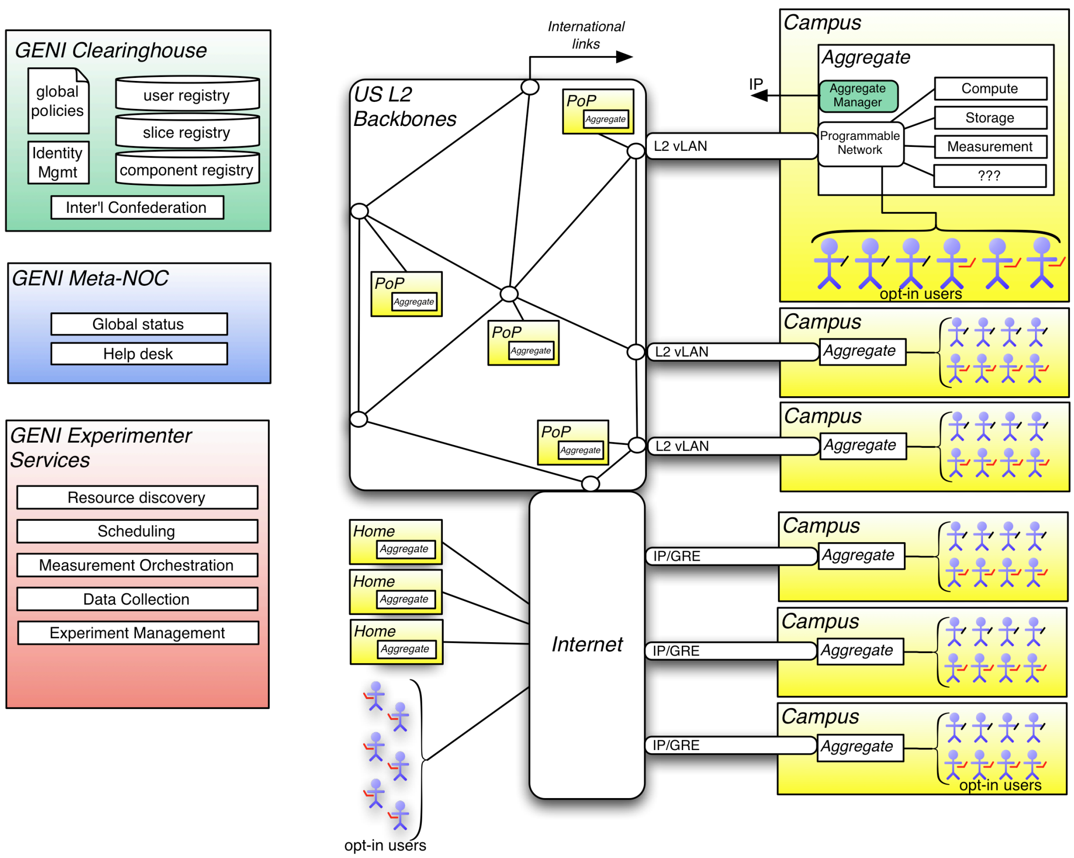

Figure 4-1. GENI System Decomposition Overview

5 GENI Spiral Development

The GPO and community have settled on an engineering approach to building GENI: spiral development. Spiral development is an engineering process that combines “elements of both design and prototyping-in-stages, in an effort to combine advantages of top-down and bottom-up concepts”[4]. In spiral development, system requirements are translated into a system design, and the result is prototyped, tried, and evaluated. Based on experience derived from this prototype, the design is revised (and requirements may be revised as well) and a new prototype is built, tested and evaluated. Each cycle of design, prototype, integrate, test and evaluation, is known as a spiral.

Spiral development is an excellent way to address both technical and community challenges. Early prototyping demonstrates the core GENI concepts, uncovers their strengths and limitations, identifies critical risks, and lets us learn which additional technologies are required to enable programmability, virtualization, and federation. At the same time, working prototypes help focus community discussions of GENI’s scope and construction.

Thus the GENI community’s fundamental strategy (see Figure 5-1) embraces spiral development with development and prototyping (D&P) as the critical means to identify, understand, and reduce project risks. Academic / industrial teams will fill in the critical missing pieces in the GENI efforts to date: security architecture, experiment workflow tools, and instrumentation. These new efforts will integrate into ongoing GENI spiral development in 6-12 months, with broad academic and industrial participation, while encouraging strong competition in all essential aspects of the design.

Each GENI spiral is 12 months in duration. We expect the community’s evolving research agenda, combined with insights and experience from earlier spirals, will set specific goals for future spirals. GENI is being designed and prototyped by the research community, with project management and system engineering provided by the GENI Project Office (GPO). After an open solicitation process and extensive peer reviews, the GPO selects development teams. We expect that a large majority of these teams will continue to participate in subsequent spirals, augmented by additional teams funded by later GPO solicitations. We also anticipate that other projects may join the GENI prototyping effort even if not explicitly funded through GPO solicitations.

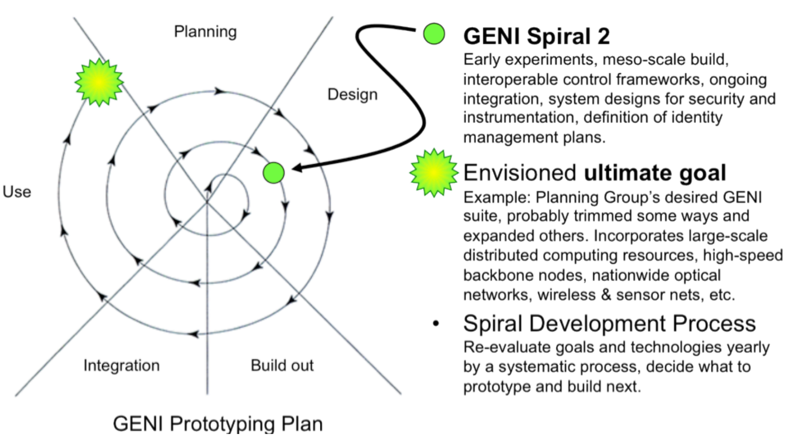

Figure 5-1. GENI Spiral Development Model

[4] See Wikipedia entry for “Spiral model”.

6 Spiral 1 Summary

GENI Spiral 1 ran from October 2008 through September 2009 with multiple academic / industrial teams creating and integrating early prototypes of the GENI suite. Spiral 1's two central goals were to demonstrate: (a) one or more early prototypes of the overarching GENI control framework, and (b) end-to-end slices operating across multiple technologies. It has shown concrete results, with the earliest demonstrations conducted in March 2009 at GENI Engineering Conference 4.

A tremendous amount was accomplished in GENI Spiral 1, and the project moved along far more rapidly than anyone dreamed possible. Energetic collaboration across the GENI community created many accomplishments in Spiral 1 including:

- Multi-project clusters created in Spiral 1 demonstrated excellent results in linking many kinds of GENI prototypes with their control and data interfaces;
- Initial deployment of prototypes was completed with compute resources and programmable network gear now running in campus and backbone networks across the United States.
- A steady march of integrated demos were performed, each showcasing increasing capability and strawman experiments many of which span multiple prototypes.

GENI Spiral 1 launched multiple competing approaches to the GENI control framework. Towards that end, most GENI projects were grouped into one of five control framework clusters. Each cluster implemented key control framework functions. In general, each cluster consisted of one prototype clearinghouse plus some number of prototype aggregates that it will control.

The five clusters organized for Spiral 1 were:

A. **TIED** control framework from USC/ISI, a cluster emphasizing issues around federation, trust, and security.
B. **PlanetLab** control framework, a cluster based on the PlanetLab implementation at Princeton, emphasizing experimentation with distributed virtual machines over the Internet.
C. **ProtoGENI** control framework, a cluster based on the Emulab implementation at the University of Utah, emphasizing network control and measurement.
D. **ORCA** control framework from Duke University and RENCI, a cluster emphasizing development of resource allocation strategies and integration of sensor networks.
E. **ORBIT** control framework from Rutgers University, emphasizing wireless networks.

Spiral 1 control framework clusters had different emphases. Some were more established and concentrated on integrating new approaches, technologies, and projects into their framework. Some investigated characteristics of control frameworks for newer substrates in depth. Some investigated alternative architectures in particular environments.

A single interoperable control framework was not produced in Spiral 1. In fact, we expect that control frameworks will continue to implement the required functions in different ways. This will provide an opportunity to make design decisions on this important topic informed by implementation and operational experience. Over time, member projects may switch clusters; for example, if another cluster will provide easier integration. Clusters may share technology and merge if they desire.

Spiral 1 integrated a wide variety of research testbeds ("aggregates") into these control frameworks. See Table 6-1 for a representative sample. Because we are using multiple competing frameworks, each aggregate has chosen which framework to integrate into. We view this as a valuable way to test the usefulness of the various control frameworks by encouraging aggregate developers to vote with their feet.

With the exception of the regional networks and those nodes going into Internet2 PoPs, Spiral 1 concluded with most aggregates as islands of programmable components residing in campus laboratories. At that time GENI lacked the ability to assemble end-to-end programmable network components between labs, and across campus, regional, and backbone networks. Spiral 2 represents the beginning of such "meso-scale" end-to-end experimentation.

Table 6-1. Representative research testbeds ("aggregates") and tools in GENI Spiral 1.

Compute clusters
- Large clusters at Univ. of Utah and USC/ISI
- Smaller clusters at Georgia Tech, Univ. Wisc, Univ. Kentucky
- Global integration with PlanetLab nodes (1000+)
- Remotely-accessible residential nodes at CMU

Sensors and Actuators
- Steerable radar & pan-and-tilt cameras at UMass-Amherst
- Mote-based sensor testbed (700+ nodes) at OSU

High-speed, optical networks
- Ethernet VLANs (up to 40Gbps) on Internet2 and National LambdaRail (national backbones)
- DWDM optical circuits and Ethernet VLANs on Mid-Atlantic Network, KanREN, BEN (regional optical networks)

Programmable network equipment
- OpenFlow campus deployment at Stanford
- High-speed (10Gbps) programmable routers and NetFPGA systems at Internet2 PoPs

Wireless Testbeds
- FPGA-based link emulator at CMU
- ORBIT 802.11 testbed (400nodes) at Rutgers
- Wireless city bus network in Amherst on 40 buses covering 1- 50 miles at UMass-Amherst
- Programmable WiMax base station at Rutgers

Very early tools and services to support experiment design and management and data collection

BGP route feeds and global IP addresses

7 Spiral 2 Technical Goals

As mentioned earlier, GENI Spiral 2 is the second phase of exploratory rapid-prototyping that will inform technical and operational plans for the envisioned GENI suite of network research infrastructure. Each GENI Spiral lasts a year and GENI Spiral 2 runs from October 2009 through September 2010. The primary goal is to help the maturing GENI prototypes migrate towards being capable of supporting continuous experimentation. Key developments for this period include improved integration of GENI prototype aggregates and their associated control frameworks; architecture, tools, and services enabling experiment instrumentation; a movement towards interoperability across GENI control framework clusters; and improved approaches towards researcher identity management.

Continuous Experimentation: It is critical to the successful design of GENI that researchers use it early on in the spiral development process. The feedback from this early use will be needed to guide design revisions and additional prototyping. Many Spiral 1 prototypes were only sufficiently mature to support demonstrations and thus could only be operated with substantial manual configuration by the developers. An important transition, which we plan to begin in Spiral 2, is to move from demonstrations to limited, continuous operations. Our goal is to find alpha-users -- researchers -- to exercise the early prototypes to perform experiments that are sustained over periods ranging from days to weeks. Besides field-testing the prototypes and the GENI design concept in general, we believe that asking prototypes to support continuous experiments will introduce GENI's development community to the challenges of making their prototypes "operational", e.g., educating new users, diagnosing and repairing failed components and services, and scheduling downtime. Therefore, in Spiral 2 we will solicit GENI users and help them use prototype resources to run early experiments. We will monitor their experiences carefully to assess what additional capabilities are needed for greater use. This topic is discussed further in the next section.

Integration: An important goal in Spiral 1 was to integrate an initial set of programmable, virtualizable components, organized into aggregates, into control frameworks. The GENI control framework will mediate access control and other common functions and, like the components themselves, are based on evolving prototypes. A Spiral 2 goal is to increase this key integration between aggregates and their associated control frameworks. For example, the control framework interface to the aggregate manager is expected to evolve. Improving aggregate-control framework integration will provide a greater, more diverse, and more consistently controlled catalog of resources for GENI researchers to build upon and will make it easier for new aggregates to join GENI.

Instrumentation & Measurement: GENI needs to provide rich instrumentation to experimenters. In Spiral 2, the GPO is leading the establishment of an instrumentation and measurement architecture that will define the major functions and interfaces for collecting, aggregating, and storing measured data. As with other elements of GENI, several excellent systems already exist supporting some of the needed functions and environments. A significant effort in Spiral 2 will be to agree upon strawmen design choices where possible, debate tradeoffs where multiple approaches exist, and identify areas that require further development. This topic is discussed further in Section 7.2.

Interoperability: Spiral 1 started out with five GENI control framework clusters. The GPO has encouraged friendly competition between the control framework projects that anchor these clusters as a way to bring the best ideas forward and maintain an energetic development environment. Nevertheless, it is also important to avoid confusing potential aggregates that might want to join GENI by forcing them to pick between control frameworks and risk being in a cluster perceived as potentially 'losing'.

Further, we want to avoid researchers being constrained to use only the subset of GENI's resources in a particular cluster. Therefore, an objective in Spiral 2 is to establish interoperability between the control framework clusters.

Ideally, a researcher should be able to use any GENI prototype in their experiment, independent of which cluster it is in. Initially we will work to establish a common interface between three Spiral 2 projects: PlanetLab, ProtoGENI and OpenFlow. Work is underway to implement a common API, the *GENI Aggregate API v1.0*[5], that enables credentials and function calls issued by the PlanetLab control framework (and associated client software) to be recognized by ProtoGENI aggregates and vice versa. Additionally, both the PlanetLab and ProtoGENI control frameworks should be able to use the API to establish slices containing OpenFlow switches. The net result is that users and tools developed by the PlanetLab and ProtoGENI clusters will be able to access resources in all three projects. This effort will lay the groundwork for revisions of the API to support needed functions as defined in the GENI Control Framework working group.

Identity Management: The process of identifying researchers who may use GENI prototypes is currently ad-hoc and varies widely between projects. The GPO believes that adoption of common identity management technologies will be very beneficial as GENI grows. In Spiral 2 the goal is to investigate and prototype the use of identity management technologies. These technologies should provide an indication of a researchers' university affiliation and some information about their status, for example, as faculty, staff, or student. A current candidate technology of interest is Shibboleth/InCommon because it is currently used for other NSF cyber-infrastructure and is widely supported across a large number of academic institutions, including many associated with existing GENI projects.

All GENI projects are expected to orient their efforts to make progress towards these goals during this spiral. Where appropriate, project milestones have been established to measure progress towards these goals. Each project's Spiral 2 milestones are listed on the *GENI Wiki* at http://groups.geni.net/geni/wiki/SpiralTwo.

Several new projects are starting in Spiral 2 that fill-in some known "missing pieces": security requirements and architecture, experiment workflow tools and user interfaces, and prototypes for instrumentation and measurement. Additional projects will build upon Spiral 1 achievements to date, including support for international and commercial federations, and will provide several early "shakedown" experiments that will prove critical in guiding system design. Finally, we will launch two related types of campus GENI build-outs – using OpenFlow and WiMax technology – with significant commonality in the technologies employed, and significant overlap in the two sets of campuses. The national backbones will both be built out with OpenFlow. In addition, Internet2 will be built out with GENI-enabled commercial routers. Accelerating the roll-out of this rich, interoperable "meso-scale" GENI infrastructure will:

- Create a compelling infrastructure for entirely new forms of network science and engineering experimentation at a much larger scale than has previously been available.
- Stimulate broad community participation and "opt in" by early users across 13 major campuses, which can then grow to additional campuses as the build-out progresses, with a strong partnership between researchers and campus infrastructure operators.

[5] Available online at http://groups.geni.net/geni/wiki/GAPI_AM_API

- Forge a strong academic / industrial base by GENI-enabling commercial equipment from Arista, Cisco, HP, Juniper, and NEC, with software from AT&T Labs and Nicira.

Internet2 and National LambdaRail are contributing world class engineering and operations expertise, along with handsome donations of bandwidth across their national research networks, to GENI's prototyping, trial integrations, and early operations.

In the following sections we expand on Spiral 2 plans to develop three key aspects of this year's technical work: bringing on users, designing a suitable instrumentation and measurement system, and security infrastructure as the prototypes transition to supporting operations.

7.1 Closing the Loop: Experiments guide GENI development

Frequent and early user feedback is vital to any spiral development process. Because GENI is a virtual laboratory, our user-base is those members of the research community seeking a suitable venue for innovative experiments in network science, security, technologies, services and applications. During the course of Spiral 2, GENI is rapidly increasing in capability and scale, enabling early experimenters to begin using the emerging GENI infrastructure.

The goal for GENI is to 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 also support experiments operating at a wide range of time scales, from individual trials of novel capabilities, perhaps executing over minutes or hours, to long-running, open-ended deployment experiments. In all cases, GENI will provide extensive instrumentation, as well as the requisite data archival and analysis tools. The capabilities needed to support these types of experimentation are still evolving and will be quite rough in Spiral 2. Early use by experimenters will provide the developers with important feedback that will shape their designs.

GENI's current focus is on experiments that highlight and stress three central GENI concepts enabled by Spiral 1 and 2 capabilities.

- **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.

Of particular interest are experiments that emphasize heterogeneous GENI slices, drawing on GENI infrastructure components that cross multiple testbeds, technologies, and control frameworks. The GENI Aggregate Manager Application Programmer Interface (API) is a key Spiral 2 capability, which provides a common point of access to and unified user credentialing for GENI resources across multiple control frameworks. Using this API, GENI experimenters gain access to unique and sophisticated end-to-end slices.

During Spiral 2, GENI is pursuing four approaches to introduce potential experimenters to GENI capabilities. In the first approach, the GPO is providing direct experiment design, development, and integration expertise and/or assistance to four early experiment teams identified at the completion of spiral 1. The GENI Alpha user program, the second approach, is focused on rapid integration and execution of demonstrable experiments, with the goal of highlighting results at the conclusion of Spiral 2. The third approach is an increased emphasis on training opportunities for interested potential experimenters. This process began with Jon Turner's well-received tutorial at GENI Engineering Conference 7 (GEC7) on the Supercharged PlanetLab Platform (SPP) and is expected to continue with expanded tutorial sessions at future GECs. Finally, over twenty experimenter teams have been invited the NSF-sponsored GENI Experimenters Workshop (GEW) at Princeton, June 29-30, 2010. The goal of the workshop was to bring together researchers interested in running experiments on GENI to attend talks and tutorials and to present their experiment plans. The workshop is expected to lead to future projects and requests for NSF funding to conduct these experiments on the GENI infrastructure.

Our expectation is that GENI will support an increasing number of experiments over the course of Spiral 2 and future GENI development spirals. We expect reports of initial results to begin during the summer and fall of 2010, as the first early experimenters complete their initial trials. The experience gained and feedback provided by the user community will help to highlight GENI's unique capabilities, identify design and implementation flaws, and guide GENI's future development path.

7.2 Establishing a GENI Instrumentation & Measurement Architecture

GENI will require instrumentation and measurement (I&M) systems that provide broad data gathering, analysis and archival capability sufficient for GENI's research mission and sufficient for operations. While some aggregate prototypes include tools for instrumentation and measurement, GENI currently lacks a comprehensive instrumentation and measurement architecture that provides common tools and interfaces for measurement acquisition, aggregation, and storage. This will be a key to the success of the GENI infrastructure.

A *GENI Measurement Workshop* was held on June 26, 2009[6] during Spiral 1 that brought together measurement experts to review the following topics: 1) measurement architecture, 2) instrumentation, 3) experiment specification, and 4) data management. The speakers suggested approaches to each topic that would allow GENI to meet its goals, thus beginning the design of an effective GENI measurement architecture. To continue the effort, six GENI I&M prototyping projects were established at the beginning of Spiral 2, complementing three existing I&M prototyping projects.

Additionally, a GENI I&M Working Group (WG) was formed with the goal of establishing an instrumentation and measurement architecture for GENI. In Spiral 2, work has been focused on creating and documenting a GENI I&M architecture and, in Spiral 3, will coordinating the design and deployment of a first GENI I&M system.

---

[6] See http://groups.geni.net/geni/wiki/GENIMeasWS

To start, an early draft of a *GENI I&M Architecture* was completed[7] by the GPO and reviewed at the GEC7 WG meeting. There was general agreement on the draft of the architecture, and it was decided that these architecture topics should be given first priority:

- GENI I&M use cases
- GENI measurement plane
- GENI I&M services
- Interfaces, protocols and schema for measurement data in GENI

In addition, a *GENI I&M Capabilities Catalog*[8] was drafted reviewing each current GENI I&M project, and other selected projects, and lists: architecture components addressed or implemented; implementations in GENI or elsewhere; and uses in GENI or elsewhere.

Based upon this catalog, several key projects were identified that have already implemented comprehensive, end-to-end capabilities:

- *OMF (ORBIT Management Framework),* Max Ott – NICTA, Ivan Seskar – Rutgers
- *Instrumentation Tools,* Jim Griffioen - Univ Kentucky
- *perfSONAR,* Matt Zekauskas - Internet2, Martin Swany - Univ Delaware
- *Scalable Sensing Service (S3),* Sonia Fahmy – Purdue, Puneet Sharma - HP Labs
- *OnTimeMeasure for network measurements,* Prasad Calyam - Ohio Supercomputing Center

In addition, there are projects currently implementing data schema and/or data archives, key parts of I&M systems for GENI:

- GENI Meta-Operations Center, Jon-Paul Herron, Indiana Univ.
- netKarma: GENI Provenance Registry, Jon-Paul Herron, Beth Plale – Indiana Univ.
- DatCat Project (not a GENI project), Brad Huffaker - CAIDA
- Data-Intensive Cloud Control for GENI, Michael Zink
- Experiment Management Service – Digital Object Registry, Larry Lannom - CNRI

Since the GEC7 meeting, the GPO system engineers and the WG Chairs have gathered technical references from and had extended discussions with many of the key projects listed above, and have gained an even better understanding of how they can best contribute to the GENI I&M architecture.

To make progress on the priority architecture components, the GPO has organized a 2nd GENI I&M workshop on June 8-9 2010 in Chicago, IL which will continue work on the I&M architecture document and draft a roadmap for implementations in Spiral 2 and 3.

7.3 Meeting Early Security Needs

GENI faces security challenges that are different from those faced by typical enterprise networks. GENI features that differentiate it from typical enterprise networks include:

- GENI is not owned and operated by one legal entity
- GENI users belong to different organizations
- GENI resources are much more diverse than the typical hosts and network devices on enterprise networks

[7] See http://groups.geni.net/geni/wiki/GeniInstrumentationandMeasurementsArchitecture

[8] See http://groups.geni.net/geni/wiki/GENIIandMCAPCAT

Table 7-1. GENI security must meet the needs of its different stakeholders.

Stakeholder	Security Objectives
NSF/BBN	GENI not used for illegal activity or as launchpad for attacks. GENI availability not compromised by attacks.
Control framework developers	Control framework functions not hijacked to gain unauthorized access to resources.
Aggregate providers	Information needed to enforce resources usage policies is available and trustworthy (who, when, how much). Resources not used for illegal activities or as launchpad for attacks.
GENI Ops	Quickly learn about incidents. Ability to determine scope and severity of threat. Ability to contain/eliminate threat. Procedures for responding to requests from law-enforcement.
Experimenters	Experiment privacy (nature of experiment, data). Resource availability.
Opt-in users	Privacy agreements are honored.
Campus IT	Campus security and privacy policies are not violated.
Federation Partners	Security compromises/attacks don't cross federation boundaries. Resources usage policies are adhered to (who, when, how much). Resources not used for illegal activities or as launchpad for attacks.

- GENI users can program almost any resource on the network, including network devices
- GENI is designed to be much more configurable than any enterprise network, which means it is more prone to errors in configuration.
- GENI connects to the Internet in a large number of places that are controlled by different organizations.

The scale of GENI, in terms of its number of resources and their geographic spread, makes it an attractive launch pad for large-scale attacks. The national and international attention garnered by the GENI project makes it an attractive target for attacks, for bragging rights if nothing else. The use of GENI for virtually any kind of networking experiment, including long-lived experiments, makes it an attractive platform for hiding and distributing illicit content.

GENI has a number of different stakeholders, each with a different objective for GENI security. These stakeholders and their security objectives are summarized in Table 7-1. Adam Slagell of the NCSA has a similar analysis of GENI[9].

It is important, therefore, that GENI have a security plan for confronting the security challenges that will arise as the system becomes available to users. The primary objective of the GENI Spiral 2 security plan is to prepare for growth in GENI resources and scope by the end of the spiral. Growth in Spiral 2 comes from GENI deployments in over a dozen campuses (with about a dozen more planned for Spiral 3) and from over twenty aggregates associating themselves with GENI. The number of

[9] GENI Security Use Cases and Stakeholders , Adam Slagell, January 2010, http://groups.geni.net/geni/attachment/wiki/ComprehensiveSecurityPgm/GENI-Security-Use-Cases-and-Stakeholders.pdf

researchers using GENI is also expected to grow by Spiral 3. Some researchers are using GENI for standing up experimental services for opt-in users; the numbers of these opt-in users too will increase as these services mature and gain in popularity.

As GENI grows in size and popularity, so does the importance of a dedicated operations team tasked with ensuring GENI remains a safe and reliable infrastructure for research and experimentation. The fact that the GENI operations team consists of members from multiple organizations with different priorities and with different levels of experience and expertise with operations presents some unique challenges.

Good policies and procedures with well-defined roles, responsibilities and expectations are essential for a diverse and distributed operations team to be effective. The GENI Spiral 2 security plan therefore focuses on putting together procedures that will allow the GENI operations team to work together to prevent security incidents and to deal with such incidents should they occur. Specifically, the following documents will be developed in Spiral 2:

1. An aggregate provider agreement,
2. Security best practices for aggregate providers,
3. Emergency stop procedures,
4. Procedures for responding to threats of legal action or law enforcement requests.

The scope and objectives of these documents are described further in *GENI Spiral 2 Security Plan*[10].

7.4 Meso-Scale Buildout

GENI will create and operate a “meso-scale” prototype GENI infrastructure built from GENI-enabled commercial hardware across 13 university campuses, linked by compatible build-outs through two US competing national research backbones (Internet2 and NLR). Figure 7-1 provides a high-level view of the GENI-enabled campuses and the two GENI-enabled backbones that will link them across an aggregate national footprint of 40 gigabits/second.

[10] http://groups.geni.net/geni/wiki/SpiralTwoSecurityPlans

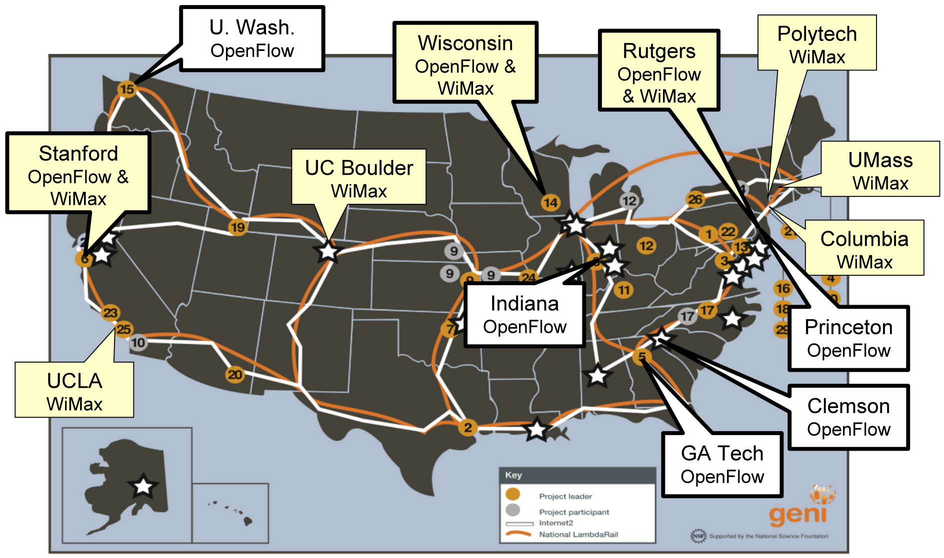

Figure 7-1. Campus buildouts for OpenFlow and WiMax prototype infrastructure.

The Meso-scale deployment includes two related types of campus build-outs – OpenFlow and WiMax – with significant commonality in the technologies employed, and significant overlap between the two sets of campuses. Initially, campus OpenFlow networks will serve as PlanetLab aggregates, while campus WiMax networks will be controlled by ORBIT. It is likely that in coming spirals all major control frameworks will be capable of managing experiments that use both these technologies.

The national backbones will both be built out with OpenFlow as part of the PlanetLab control framework. In addition, Internet2 will be built out with GENI-enabled commercial routers (Juniper M7i) as part of the ProtoGENI control framework.

While starting in Spiral 2, these build-outs will take place over three years. They will GENI-enable more than a dozen campuses of major research universities throughout some or all of their footprint, including both wireline Ethernet and WiFi access, as well as GENI-enabled WiMax for “cellular telephone” style services that are open for researcher experimentation. Each of these campuses has pledged to guide expansion to a further three campuses.

The campuses will be interlinked by GENI-enabled equipment in National LambdaRail and Internet2. These backbones will be thinly built at first, due to funding constraints, but will provide the earliest, native-mode GENI interconnections between GENI-enabled campuses.

Table 7-2 provides a summary list of campuses, Principal Investigators (PIs), and campus infrastructure leadership who will be responsible for these build-outs. As can be seen, the large majority of projects already have campus infrastructure officials committed to the build-out; the GPO will ensure that *all* campuses have such active partnerships. In Spiral 1, 9 of the 13 campuses were already engaged in GENI prototyping activities; all have directly relevant research projects that can benefit from involvement with the GENI project (and vice versa).

Table 7-2. Campus researchers and IT leadership committed to GENI infrastructure build-outs.

Campus	PI & Infrastructure	Build-out	Relevant campus projects (• GENI)
Clemson University	KC Wang & Jim Pepin (CTO)	OpenFlow	○ Wireless and sensor testbeds
Columbia University	Henning Schulzrinne	WiMax	○ PI chairs GENI Opt-in working group
GA Tech	Nick Feamster & Ron Hutchins (CIO)	OpenFlow	• BGPmux
Indiana University	Christopher Small & David Jent (AVP)	OpenFlow	• GMOC network management
Polytechnic Institute of New York University	Thanasis Korakis	WiMax	○ Wireless experimentation
Princeton University	Michael Freedman	OpenFlow	• PlanetLab control framework • SCAFFOLD (Solicitation 2)
Rutgers	Ray Raychaudhuri & Charles Hedrick (CTO)	OpenFlow WiMax	• ORBIT control framework • WiMax
Stanford University	Nick McKeown	OpenFlow WiMax	• Enterprise GENI (OpenFlow)
UCLA	Mario Gerla & James Davis (CIO)	WiMax	○ Wireless & vehicular networks ○ Sensor network testbeds
U. Colorado (Boulder)	Dirk Grunwald	WiMax	• Cognitive radios (Solicitation 2)
U. Mass. Amherst	Mark Corner & Daniel Blanchard (ACIO)	WiMax	• DOME • ViSE
U. Washington	Tom Anderson & Clare Donahue (AVP)	OpenFlow	• Million-Node GENI
U. Wisconsin (Madison)	Suman Banerjee & Perry Brunelli (Dir.)	OpenFlow WiMax	• GENI Measurement ○ Metro wireless testbeds

7.4.1 OpenFlow Buildout

The OpenFlow build-out will install GENI-enabled LAN, WiFi, and backbone equipment through 8 major research university campuses and 10 points of presence in 2 national research backbones in order to create an initial GENI infrastructure across the United States.

What is OpenFlow? OpenFlow is "an open standard that enables researchers to run experimental protocols in the campus networks we use every day. OpenFlow is added as a feature to commercial Ethernet switches, routers and wireless access points - and provides a standardized hook to allow researchers to run experiments, without requiring vendors to expose the internal workings of their network devices."[11]

Within the GENI context, OpenFlow opens up a network so it can be shared between production traffic and research flows, which can then be managed by a GENI control framework to create virtual "slices" for separate research experiments. Such "dual use" approaches are extremely interesting as they

[11] http://www.openflowswitch.org/

enable GENI infrastructure to grow by opening up existing networks for research experimentation. In particular, they GENI-enable existing campus LANs and WiFi access points, greatly accelerating the introduction of students as early opt-in users.

OpenFlow is not the only potential means towards achieving this end. Nevertheless, according to its inventors and advocates, "OpenFlow is a pragmatic compromise that allows researchers to run experiments on heterogeneous switches and routers in a uniform way, without the need for vendors to expose the internal workings of their products or researchers to write vendor-specific control software. If we are successful in deploying OpenFlow networks in our campusses [*sic*], we hope that OpenFlow will gradually catch on in other universities, increasing the number of networks that support experiments. We hope that a new generation of control software emerges, allowing researchers to re-use controllers and experiments, and build on the work of others."[12]

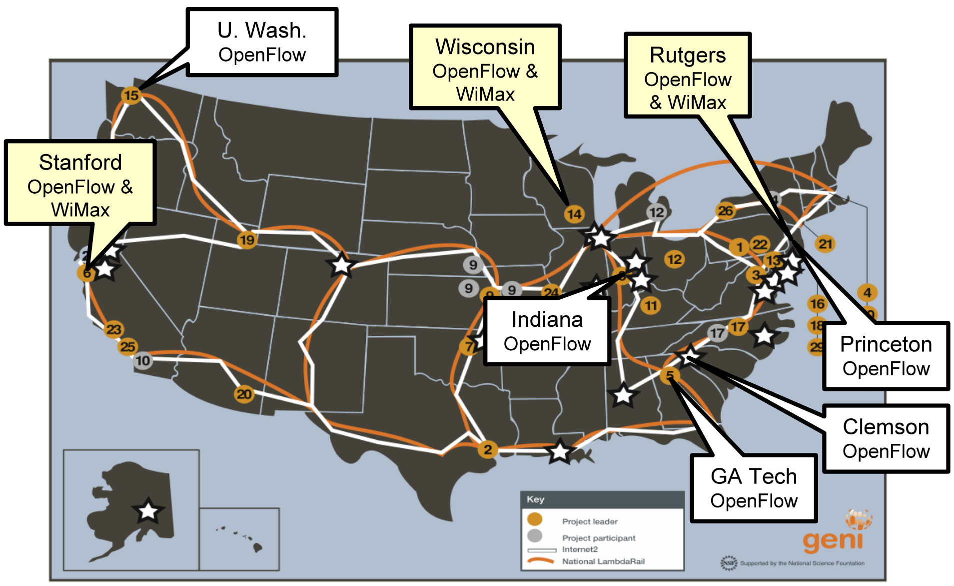

Figure 7-2. OpenFlow campus build-outs.

Figure 7-2 identifies the initial 8 campuses for GENI's OpenFlow build-out. At each campus, researchers are teamed with campus IT staff to plan, deploy, shake down, and operate OpenFlow switches as part of the emerging end-to-end GENI suite of infrastructure.

[12] "OpenFlow: Enabling Innovation in Campus Networks," whitepaper, by Nick McKeown, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Scott Shenker, and Jonathan Turner, March 14, 2008.

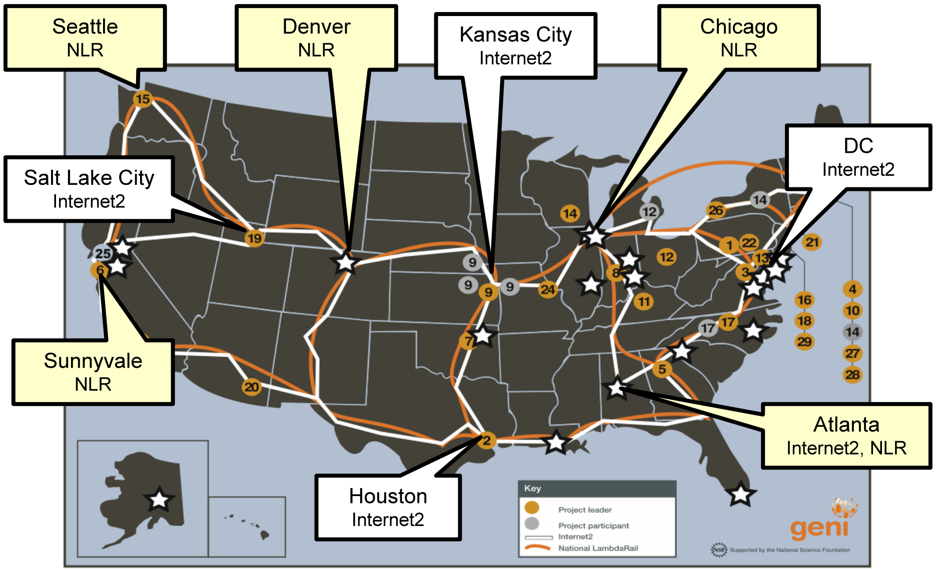

Figure 7-3 OpenFlow backbone build-outs through Internet2 and NLR (notional).

Figure 7-3 shows 10 Points of Presence (PoPs) targeted for installation and operation of OpenFlow switches in the Internet2 and NLR backbones. They will be linked to campus switches to form the emerging end-to-end GENI suite of infrastructure. Note that the NLR PoP locations are currently notional. Internet2 PoPs and prices have already been negotiated.

Table 7-3. Anticipated OpenFlow Equipment with Deployments (eval = in evaluations).

Anticipated equipment	Locations
 HP ProCurve 5400 Switch	• Georgia Tech • Indiana University • Univ. Wisconsin • Univ. Washington • Princeton • Stanford • Internet2 (eval) • National LambdaRail • BBN
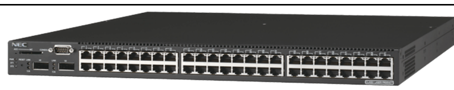 NEC IP8800 Switch	• Georgia Tech • Rutgers • Stanford • BBN
 Toroki Lightswitch 4810	• Clemson • Stanford
 Quanta L4BG Switch	• Stanford • BBN

Who does what for the OpenFlow buildout? The OpenFlow build-out is quite complex and involves a number of campuses, backbone providers, equipment vendors, and software development teams. It is a major undertaking, with many moving parts, and will require intensive project management and engineering integration.

Stanford, BBN I2, NLR, and the campuses are all working together on software development, integration, and testing. Regional networks are also participating, although they don't yet have OpenFlow switches installed. The GPO and Stanford are coordinating the work. More details are available on the *GENI wiki Sprial2 Project OpenFlow pages*[13] and on the *Stanford CleanSlate GENI wiki pages*[14].

7.4.2 WiMax Build-out

The WiMax build-out will create an open, programmable, GENI-enabled “cellular-like” infrastructure through 8 major research university campuses as the first step towards creating an open,

[13] http://groups.geni.net/geni/wiki/SpiralTwo#Meso-ScaleDeploymentandOperations

[14] http://www.openflowswitch.org/foswiki/bin/view/OpenFlow/Deployment/

research-enabled cellular telephony system across the United States. This opens up a path for direct "opt in" of student users in these campuses into GENI research experiments, via WiMax modems and, as they become available, cell phone handsets.

Open WiMax base stations will provide network researchers with wide-area coverage and the ability to support both mobile and fixed end-users. Campus-wide WiMax coverage will make it possible to quickly extend GENI services for "opt-in" by large user populations with a relatively low capital investment.

We note that these WiMax base stations are OpenFlow compatible. Given the ambitious scope of the work already laid out for these campus deployments, we do not propose to undertake this additional level of integration now; however, it does open up many future opportunities.

What is WiMax? WiMax is an inexpensive technology that provides "cellular-like" services, i.e., wireless coverage from a base station to a handset held by a pedestrian or located within a moving vehicle. Sprint and Clearwire are currently using this technology, based on the IEEE 802.16 standard, to build new cellular systems for commercial use. It differs from conventional cellular systems in the exact radio modulations used for transmission, but more importantly, in the RF spectrum in which it runs. A key advantage is that WiMax can be used freely in spectrum already controlled by universities, unlike commercial cellular telephony spectrum which is controlled by companies such as AT&T or Verizon.

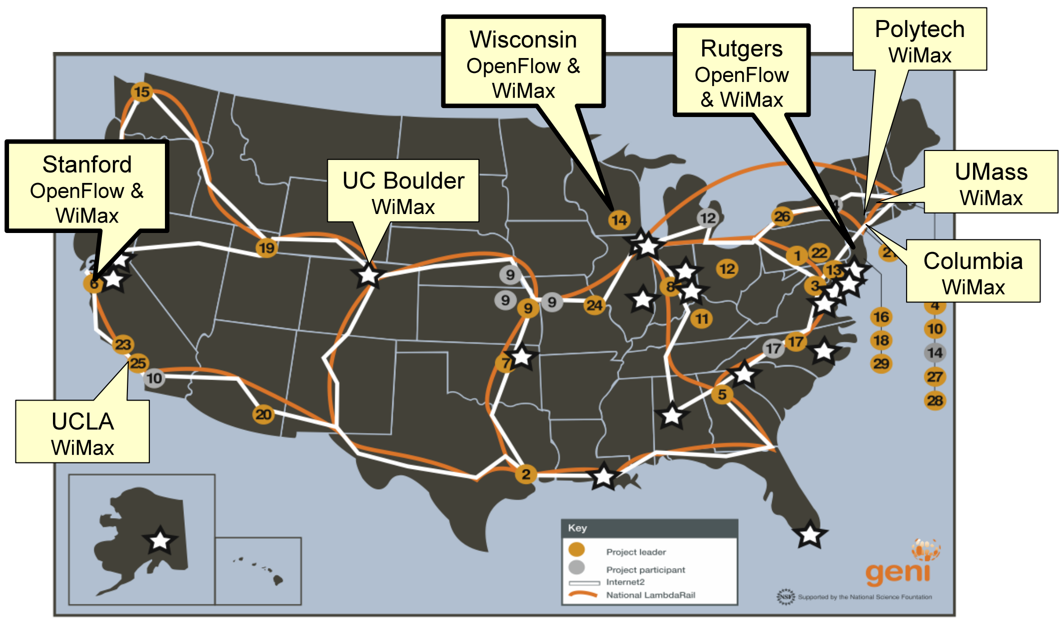

Figure 7-4. WiMax campus buildout

Figure 7-4 depicts the campuses that will participate in the initial WiMax buildout. Each will perform its own local experimentation with the equipment and also make it available as part of the end-to-end GENI infrastructure to researchers across the United States.

Table 7-4. Anticipated WiMax Equipment with Deployments

Anticipated equipment	Locations
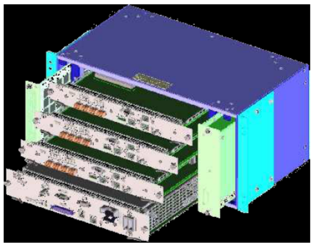 NEC WiMax Base Station (with RF and antenna)	• Columbia • Polytechnic Institute of New York University • Rutgers • Stanford • UC Boulder • UCLA • UMass Amherst • U. Wisconsin (Madison)

NEC's WiMax base-station hardware is a 5U rack-based system which can be populated with up to three channel cards, each supporting one antenna sector for a maximum of three sectors. It operates in the 2.5 Ghz or the 3.5 Ghz bands and can be tuned to use either 5, 7 or 10 Mhz channels; many universities control RF spectrum in these bands, which can thus be used freely for WiMax experimentation. The base station has been tested for radio coverage and performance in realistic urban environments and is being used in early WiMax deployments – typical coverage radius is ~3-5Km, with peak service bit-rates from 15-30 Mbps depending on operating mode and terrain.

The WiMax base station has an external PC controller that runs ASN gateway software plus the ORBIT control framework software to interface the base station to other parts of the GENI suite of infrastructure. This controller provides support for multiple GENI slices. Each slice runs within its own virtual machine and can emulate its own router and perform IP routing, or alternatively implement novel non-Internet protocols.

Who does what for the WiMax build-out? Rutgers will buy additional WiMax equipment from NEC, and will help each of the campuses install the equipment and make it operational. Campuses will buy auxiliary equipment such as WiMax modems. Rutgers will continue to enhance the ORBIT control framework to control these base stations. The GPO will assist as needed with integration and shake down, and will maintain stable software repositories.

7.4.3 ShadowNet Build-out

The ShadowNet build-out will create an open, virtualized backbone research infrastructure through Internet2, using GENI-enabled commercial Juniper routers. It will include extensive measurement services integrated into the ProtoGENI control framework.

What is ShadowNet? ShadowNet is a collaborative academic/industrial partnership between the University of Kentucky, AT&T Research, the University of Utah, the University of Delaware, and Internet2. It will install four GENI-enabled, commercial Juniper M7i routers into the ProtoGENI sites within the Internet2 backbone, with an instrumentation and management system, ShadowBox, created by AT&T, and operate them as part of GENI for researcher experiments.

Figure 7-5 provides a geographic map of the notional ShadowNet build-out. We intend to install ShadowNet nodes into existing and forthcoming ProtoGENI sites. Internet2 sites will be selected so as to incorporate topologically interesting locations, while avoiding sites at which rack space or power may be problematic or too expensive.

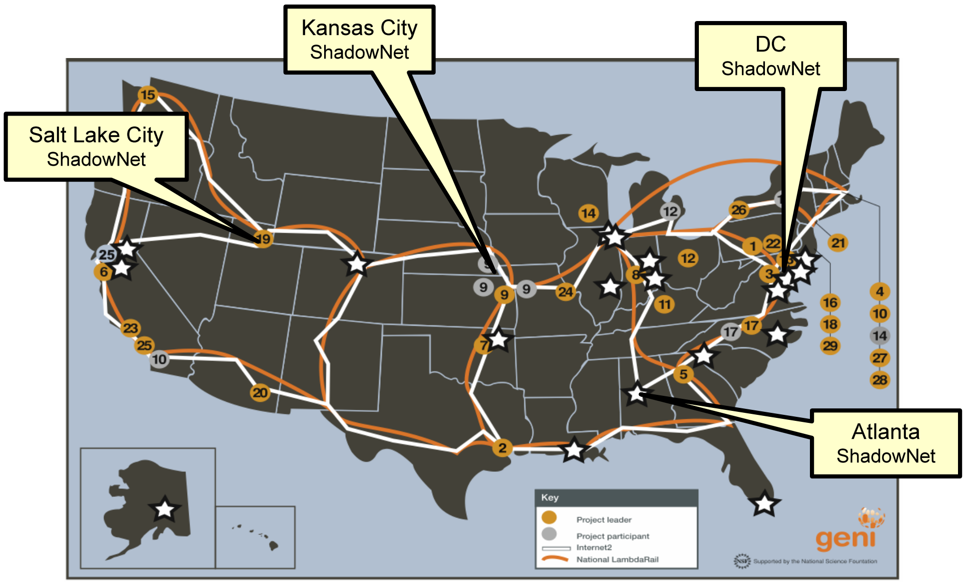

Figure 7-5. ShadowNet build out through Internet2 backbone (ProtoGENI sites)

Table 7-5 depicts basic ShadowNet equipment that will be installed into ProtoGENI backbone sites; measurement and control computers are not depicted. Each physical M7i platform can be sliced into multiple logical routers, which will be separately programmed and separately instrumented as part of a GENI experiment.

Table 7-5. Anticipated ShadowNet Equipment with Deployments

Anticipated equipment	Locations
 Juniper M7i Router	• Internet2 (ProtoGENI)

Figure 7-6, drawn from the ShadowNet proposal, depicts how the Juniper M7i routers will be installed within ProtoGENI sites in the Internet2 backbone.

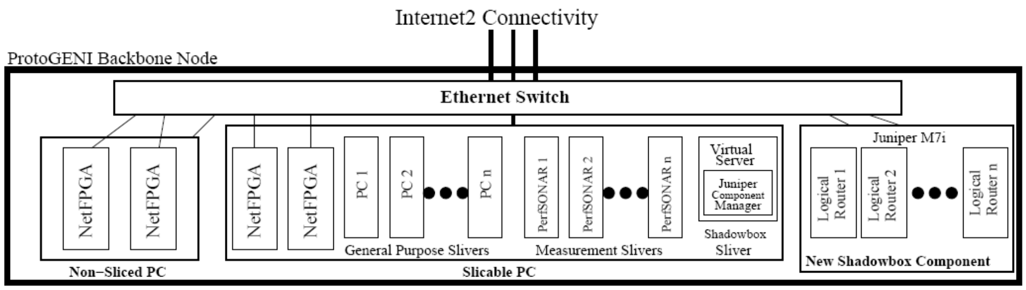

Figure 7-6. ProtoGENI backbone sites augmented with Shadowbox component.

Who does what for the ShadowNet build-out? Although the physical installation of the Juniper routers into ProtoGENI sites in the Internet2 backbone is relatively straightforward, the overall ShadowNet system requires integration of several software systems. Project responsibilities are as follows:

- **U. Kentucky** – Project Lead. Develop tools to extend existing ProtoGENI instrumentation tools and user interfaces to support access to the new measurement data.
- **AT&T** – Integrate Juniper routers into the ProtoGENI framework with support for per-slice monitoring. AT&T will perform their tasks at no cost to GENI.
- **Internet2** – Install and operate Juniper routers. Add Dynamic Circuit Network (DCN) capabilities into the measurement plane to ensure timely access to, collection of, and display of measurement data.
- **U. Delaware** – Deploy measurement point services that incorporate the *perfSONAR* data model to provide a common, extensible format for data storage and exchange.
- **U. Utah** – Integrate new measurement plane into the ProtoGENI framework and APIs.
- **GPO** – Assist with integration and trouble-shooting.

8 Resources Expected to be Available for Experimenter Use in Spiral 2

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 Table 8-1 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 LambdaRail (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.

Table 8-1. Available GENI Resources (expected summer 2010[15])

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			

[15] Academic operations only in Summer 2010 (no central operations). Does not include general resources, such as the commercial Internet.

Resource	Available number	Host Institution	Notes
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	Internet2, 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	
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 and WiMAX deployment	1	Rutgers University	
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			

Resource	Available number	Host Institution	Notes
GUSH Experiment Control and Management Tool	TBD	Williams College	
Raven Provisioning Tool	TBD	University of Arizona	
GENI Measurement Services[16]			
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	

Notes:

[1] Available through GENI Aggregate API v1.0.

[2] May need custom configuration or review / approval by provider. GPO can help.

[16] Does not include measurement tools built into control frameworks like ProtoGENI, ORCA etc.

9 Control Framework Integration Projects

This section provides a brief summary of each of the control framework integration projects active in Spiral 2. GENI projects are, in general, funded for 1 year with up to two additional option years. The projects listed here are in their first or second year. The project scope covers the work planned for the entire period. Information about the specific plans during Spiral 2 can be found on the *Spiral 2 wiki* page at http://groups.geni.net/geni/wiki/SpiralTwo.

Control framework integration projects emphasize development and integration of GENI prototype components and services with one of the candidate control frameworks. Each project has an associated "cluster" which identifies a control framework. Note that many projects choose to integrate with multiple control frameworks. The mapping of projects to clusters is shown in Figure 9-1. More information can be found on the Spiral Two wiki page.

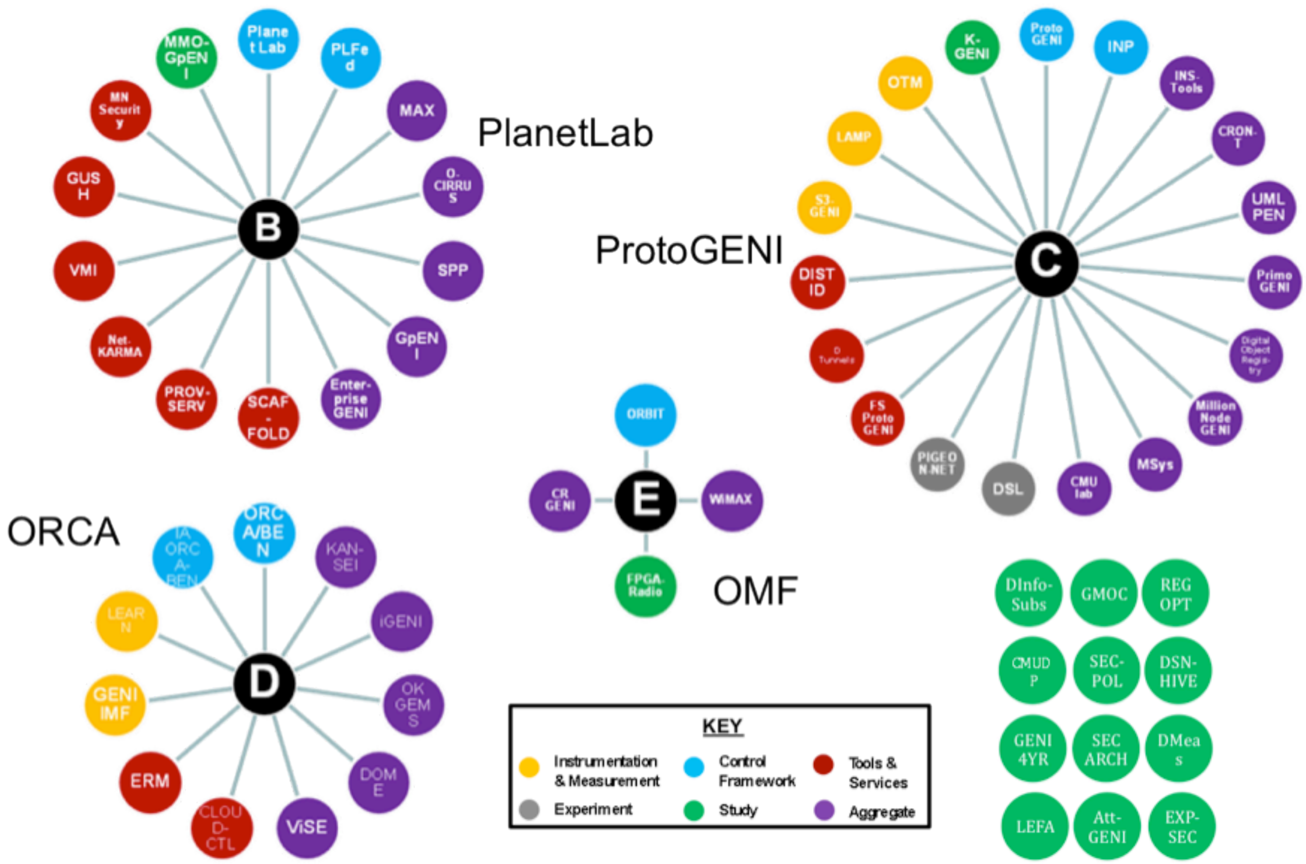

Figure 9-1. Spiral 2 Integration Clusters

9.1 Control Framework Design and Prototyping

Larry Peterson **Prototyping and Validating the GENI Control Framework (PlanetLab)**
Princeton University

The scope of work on this project is to prototype the control framework that logically stitches GENI components and user-level services into a coherent system; integrates a representative set of components and autonomous organizations into a coherent, operational whole; integrates edge clusters, high-performance backbone nodes, enterprise-level nodes, and edge-sitting wireless nodes; and federates across multiple independently-controlled aggregates, including those managed by international and corporate partners. This effort will culminate in a running system that can be accessed by the network research community.

Specific development goals include creating and deploying component manager packages and clearinghouse packages; operating a prototype clearinghouse; making reference clearinghouse and component manager implementations available via download to GENI prototype developers to use in their own aggregate or campus infrastructures; providing limited integration and development support for those users so that reference implementations can be deployed at multiple GENI locations; and providing GENI user access to PlanetLab (1000 nodes, 500 sites).

Rob Ricci **ProtoGENI and PG Augmentation**
University of Utah

This project includes a large-scale integration of existing and under-construction systems that provide key GENI functionality. The integration consists of four key components: a nationwide, high-speed backbone on Internet2's wave infrastructure; a set of sliceable, user programmable components embedded in this backbone, including PCs and "programmable hardware" NetFPGA cards; a set of subnets, including a variety of wireless networks, residential broadband, and programmable edge clusters, that will be connected to this backbone; and software from the University of Utah, based on an enhanced version of the Emulab testbed management software, plus additional software from PlanetLab and VINI, to manage the backbone, many of the subnets, and to provide a common point of integration between them.

The primary contributions of this project to GENI are the development and deployment of the ProtoGENI programmable network on Internet2; deployment and operation of a prototype clearinghouse, control framework, and tools; delivery of reference clearinghouse and component manager implementations available to GENI prototype developers to use in their own aggregate or campus infrastructures; and limited integration and development support for those users so that reference implementations can be deployed at multiple GENI locations.

This augmentation effort will integrate additional projects with the ProtoGENI control framework (Cluster C). This work will support the activities of the Utah ProtoGENI team toward this goal. This proposal will provide new Cluster C projects with the software, standards, and technical assistance required to integrate their projects into the ProtoGENI control framework, and to refine the design of ProtoGENI to make smooth integration possible. These projects span a wide range of the GENI ecosystem, including international collaboration, instrumentation and measurement, hybrid environments involving multiple network technologies, early experiments, and aggregate deployments of sliceable networks, routers, storage, and hosts. The project will also support researchers running experiments on ProtoGENI.

Larry Lannom **GENI Registry Services (DigitalObjectRegistry)**
Corporation for National Research Initiatives

The scope of work on this project is to adapt the Handle System and/or a Digital Object Registry to realize a clearinghouse registry for principals, slices, and/or components in a selected GENI Spiral 1 control framework (possibly using XML RPC), and capable of supporting limited operations in Year 1. Also, to analyze how the Handle System and/or a Digital Object Registry could be used to identify and register GENI software, including experimenter's tools, test images and configurations, and test results. And finally, to define the operational, scaling, security, and management requirements, plus recommended design approaches, for implementing GENI clearinghouse and software registry services.

Ilia Baldine
RENCI

Deploying a Vertically Integrated GENI "Island": A Prototype GENI Control Plane (ORCA) for a Metro-Scale Optical Testbed (BEN)

The scope of work on this project is to extend ORCA (a candidate GENI control framework) to include the optical resources available in [BEN] metro optical testbed. Also, to make a reference implementation of ORCA available to additional GENI prototypes and assist, with limited support, in the integration of ORCA into select wireless/sensor prototypes. The team will operate a clearinghouse for the various testbeds under the ORCA control framework. Also, connect with the NLR backbone network FrameNet service to facilitate end-to-end VLAN's connections into BEN. Finally, make the integrated (VLAN connections and control framework) ORCA/BEN infrastructure available to external researchers by the end of Year-1.

The ORCA Augmentation project augments the ORCA/BEN project which was started at the beginning of Spiral 1, and which provides the code for the ORCA control framework. This project will augment the features of the current ORCA Control Framework prototype and integrate work from other relevant GENI projects to substantially increase the capabilities of the framework in the areas of:

(1) unified measurement and experimenter tools, particularly physical layer measurement
(2) identity and trust management based on Shibboleth and SAML
(3) cloud computing substrates
(4) resource description and allocation mechanisms, policies and algorithms

As the capabilities of ORCA are increased, they will be made available to the associated projects in Cluster D.

Marco Gruteser
WINLAB / Rutgers University

Control, Measurement, and Resource Management Framework for Heterogeneous and Mobile Wireless Testbeds (ORBIT)

The scope of work on this project is to extend the OMF (cOntrol and Management Framework) to design and implement a prototype of a GENI-compliant resource management, experiment control, and measurement framework, so that it can support experiments across heterogeneous testbed resources, with a specific focus on mobile testbeds. GENI control framework functions will be added to the OMF as they become defined. Support for mobile testbeds will include the handling of disconnection, and providing temporal and spatial control. Reference implementations of a GENI-compliant resource management, experiment control, and measurement framework will be provided to another wireless project, and its testbed will be integrated into a OMF environment as a federated testbed. Both testbeds will be connected to a GENI backbone network.

9.2 Network Aggregate Design and Prototyping

Peter O'Neil
Mid-Atlantic Crossroads GigaPOP

Mid-Atlantic Network Facility for Research, Experimentation, and Development (MANFRED) (MAX)

This project will provide access to an operational, regional, multi-wavelength optical network. Key contributions include:

1. Extending DRAGON's open-source GMPLS-based control plane implementation to include edge compute resources and support network virtualization (viewing DRAGON as an aggregate/component manager);
2. Leading the integration of PlanetLab control framework within Cluster B onto the DRAGON test-bed, a physical DWDM and Layer 2 network deployment of open-source GMPLS control plane software development and the deployment of that control plane software over other networks;
3. Enabling backbone connections to resources of substrate components in Cluster B participants across Internet2 in support of end-to-end VLAN connections into the DRAGON test-bed;
4. Making integrated (VLAN connections and control framework) DRAGON infrastructure available to external researchers by the end of Spiral 1;
5. Representing/offering developed DRAGON technology to the various control frameworks selected in Spiral 1 through active participation in GECs and the appropriate working groups.

James P.G. Sterbenz **Great Plains Environment for Network Innovation (GpENI)**

University of Kansas

This scope of this project is to:

1. Lead the integration of the Princeton control framework onto the GpENI testbed.
2. Enable backbone connection (Internet2/NLR) supporting end-to-end VLAN connections into the GpENI testbed.
3. Make GpENI infrastructure running Princeton control framework available to external researchers over VLAN connections by the end of year 1.
4. Present GpENI integration and operational status to the various control frameworks selected in Spiral 1 through active participation in GEC meetings and the appropriate working groups

Nick Feamster **Bringing Experimenters and External Connectivity to GENI (BGPMux)**

Georgia Tech University

This project will add capabilities to the VINI testbed to enable experiments to carry traffic from real users; and increase the experimental use of the VINI testbed by providing familiar experiment management functions. The deliverables for this project all comprise software for supporting external connectivity and flexible, facile experimentation on the GENI testbed. The primary deliverables are a BGP session multiplexer – a software router based on the Quagga software routing suite, software support for virtual tunnel and node creation, and integration of the above functionality with clearinghouse services developed as part of the ProtoGENI project.

In particular, this project contributes to GENI design and prototyping through BGP mux development integration with ISPs; tunnel & topology establishment & management; ProtoGENI clearinghouse integration; and support for isolation & resource swapout.

Seung-Jong Park **Integrating a CRON (Cyberinfrastructure of Reconfigurable Optical Network) Testbed into ProtoGENI (CRON)**

Louisiana State University

This effort provides a reconfigurable optical network emulator aggregate (CRON) connected to the GENI backbone over Louisiana Optical Network Initiative (LONI). The role of optical network emulation in GENI is to provide a predictable environment for repeatable experiments, and to perform early tests of network research experimentation prior to acquiring real network resources. The tools and services developed by this project will integrate with the ProtoGENI suite of tools. The aggregate manager and network connections between LONI and GENI for this project will also allow other LONI sites to participate in GENI.

Jason Liu **PrimoGENI – Developing GENI Aggregates for Real-Time Large-Scale Network Simulation**

Florida International University

This work will integrate a large-scale, real-time network simulator (PRIME) into ProtoGENI, thus enabling slices involving both physical and simulated networked components. The experimenter tools, including experiment design and management tools, developed by this project, will be integrated with the ProtoGENI suite of tools. In addition to deploying and supporting the prototype network simulator service and tools, the project will also facilitate connections between the GENI resources and geographically distributed computational nodes located in Miami at FIU and in Brazil.

Ilia Baldine
RENCI

Deploying a Vertically Integrated GENI "Island": A Prototype GENI Control Plane (ORCA) for a Metro-Scale Optical Testbed (BEN)

The scope of work on this project is to extend ORCA (a candidate GENI control framework) to include the optical resources available in [BEN] metro optical testbed. Also, to make a reference implementation of ORCA available to additional GENI prototypes and assist, with limited support, in the integration of ORCA into select wireless/sensor prototypes. The team will operate a clearinghouse for the various testbeds under the ORCA control framework. Also, connect with the NLR backbone network FrameNet service to facilitate end-to-end VLAN's connections into BEN. Finally, make the integrated (VLAN connections and control framework) ORCA/BEN infrastructure available to external researchers by the end of Year-1.

Joe Mambretti
Northwestern University

iGENI: A Distributed Network Research Infrastructure for the Global Environment for Network Innovation

This project will define, design and implement iGENI, a distributed network research infrastructure; integrate it with current and planned GENI resources; and operate it for the use of GENI researchers conducting experiments that involve multiple aggregates (at multiple sites). The iGENI infrastructure will be defined in collaboration with the GPO and other GENI projects to expand the controllable transport services available to GENI researchers, and make GENI available to more research communities. The iGENI infrastructure will connect existing resources with iCAIR involvement (e.g. StarLight?) with current GENI backbone transport resources (e.g., Internet2 and NLR layer 2/Ethernet VLANs); current and planned GENI regional transport resources (e.g., BEN, CENIC, and others); and additional available connections to other research networks (e.g. FIRE).

Deniz Gurkan
University of Houston

Programmable Measurements over Texas-based Research Network: LEARN

This project will integrate the Lone-star Education And Research Network (LEARN) regional optical network, located in Texas, into the ORCA control framework (GENI Cluster D). This will enable GENI researchers to utilize the LEARN network for L2 (VLAN) transport. This project will also work with the Embedded Real-time Measurements (ERM) project to develop methods to include the use of existing commercial transport and measurement equipment for real-time measurements. Specifically, it will work with the ERM project to define the API and implement Measurement Handler software to utilize an existing interface (TL1 over SSH) into the Infinera Digital Transport Node (DTN) to make optical measurements.

9.3 Programmable Network Node Design and Prototyping

Nick McKeown
Stanford University

Enterprise GENI: Prototype, Demonstrate and Create Kits for Replication (E-GENI)

The scope of work on this project is to illustrate how GENI can be deployed on local networks, such as campus and enterprise networks, and to develop a kit that allows the work demonstrated in this project to be easily replicated elsewhere. Priority activities that contribute to this scope include the following items:

1. Deploy Enterprise GENI technology in a campus (wireline and wireless).
2. Integrate an OpenFlow network with a GENI control framework in an Aggregate Component Manager.
3. Provide access to Enterprise GENI testbeds for GENI users.
4. Define an Enterprise GENI deployment kit for research and potential commercial transition.

Jon Turner
Washington University, St. Louis

Internet Scale Overlay Hosting (SPP)

The objective of the project is to acquire, assemble, deploy and operate five high performance overlay hosting platforms, and make them available for use by the research community, as part of the emerging GENI infrastructure. These systems will be hosted in the Internet 2 backbone at locations to be determined. We will provide a control interface compatible with the emerging GENI control framework that will allow the network-scale control software provided by Princeton to configure the systems in response to requests from research users. The project will leverage and extend our Supercharged PlanetLab Platform (SPP) to provide an environment in which researchers can experiment with the kind of capabilities that will ultimately be integrated into GENI. We also plan to incorporate the netFPGA to enable experimentation with hardware packet processing, in the overlay context.

David Andersen
Carnegie Mellon University

Prototype Support for Heterogeneous Testbed Resources: Integrating cluster, broadband, and wireless emulation nodes into the "Proto- GENI" Framework (CMULAB)

This project will build upon CMU's existing cluster, neighborhood wireless/broad-band, and wireless emulation testbeds to concretely identify—and build prototypes of—the authentication, resource arbitration, and node management primitives needed to coherently deal with this very diverse set of resources. The project will integrate these testbeds with the ProtoGENI effort from the University of Utah, which is itself based upon that group's Emulab software.

Yan Luo
University of Massachusetts, Lowell

Programmable Edge Node (UMLPEN)

scope of work on this project is to work on the design and performance evaluation of a Programmable Edge Node with x86 Multi-Core Processors and Network Processors. Specific tasks include:

1. Developing a PEN virtual router templates, set-up scripts and traffic shaper as elements of the component manager or user help tools within a specified GENI control framework.
2. Developing a measurement and diagnostic tools as elements of GIMS and user help tools within a specified GENI control framework .
3. Deploying an operational PEN in spiral-1 infrastructure under a specified control framework.
4. Representing position taken on PEN architectures and design by active participation in GEC's and appropriate working groups.

9.4 Compute Aggregate Design and Prototyping

Rick McGeer **GENI/Ecalyptus Federated Resource Allocation (GENICloud)**

HP Labs

The GENI facility envisages researchers constructing, deploying, and executing experiments on a variety of test infrastructures. Various large-scale tests, including wide-area storage, content distribution, and some forms of sensing with analysis, require the simultaneous use of both widely distributed sensing and actuating nodes and analysis nodes in large-scale cloud clusters. Eucalyptus is popular, open-source software for building a cloud computing infrastructure. The GENICloud project will build a GENI federation interface for compute clusters running Eucalyptus. This interface will allow Eucalyptus clusters to federate via the Slice-based Federation Architecture (SFA), and will enable experimenters to seamlessly deploy and configure slices spanning PlanetLab and multiple Eucalyptus installations.

The project anticipates devoting approximately 32 nodes at the HP Labs OpenCirrus facility long-term to hosting a Eucalyptus cluster, and opening this to approved GENI researchers. This cluster will be tightly instrumented by HP Labs in order to gather data on cluster usage and resource usage

---

Thomas Anderson **A Prototype of a Million Node GENI (MNG)**

University of Washington

The scope of work on this project is to specify, build, and demonstrate a prototype end host deployment platform consisting of three components: a) a general-purpose yet safe execution environment for experimenter code running on end hosts, b) a light-weight, programmable and customizable proxy for end hosts for redirecting end-user traffic into GENI, and c) a clearinghouse for resources contributed by end users who have opted into the system. Specific development goals include:

1. Development of a viable end-host virtual machine
2. Control-plane integration, including resource discovery and reservation; experiment debug and control
3. End-user configuration and management tools
4. Deployment strategies such as inclusion in Bit Tyrant or web plug-in
5. End-user policy & awareness issues and mechanisms

Michael Zink **Data-Intensive Cloud Control for GENI (DICLOUD)**

UMass, Amherst

This project will develop a complete environment for researchers to conduct data-intensive experiments in GENI from start (the data collection point) to finish (processing and archiving). To do so, this project will extend the GENI/ViSE sensor network (sensornet) testbed at UMass-Amherst and augment GENI Cluster D's Orca control framework with capabilities for researchers to

(i) obtain data-centric slices that span core sensornet nodes, data center nodes, and, importantly, storage volumes "in the cloud,"

(ii) deploy popular cloud computing programming paradigms to enable simple, but powerful, distributed data processing, and

(iii) execute experiment workflows to explicitly control experiment data flow and resource allocation across a network of components/aggregates.

The project will build on existing software artifacts in the GENI "ecosystem" and tailor them to the distinct requirements of data-intensive experiments. While the enhanced software artifacts will generalize to any high-bandwidth data-intensive experiments, the GENI/ViSE sensornet testbed, which collects high-bandwidth data from multiple high-power (virtualized) sensor/actuators, will be the initial data source. Our goal by year one is to incorporate commercial cloud computing services, including storage services, as GENI substrates available for researchers.

9.5 Wireless Aggregate Design and Prototyping

David Andersen
Carnegie Mellon University

Prototype Support for Heterogeneous Testbed Resources: Integrating cluster, broadband, and wireless emulation nodes into the "Proto- GENI" Framework (CMULAB)

This project will build upon CMU's existing cluster, neighborhood wireless/broad-band, and wireless emulation testbeds to concretely identify—and build prototypes of—the authentication, resource arbitration, and node management primitives needed to coherently deal with this very diverse set of resources. The project will integrate these testbeds with the ProtoGENI effort from the University of Utah, which is itself based upon that group's Emulab software.

Brian Levine
UMass Amherst

Slivers and Slices in a Diverse, Outdoor, Mobile Network Testbed (DOME)

The scope of work on this project is to extend the DOME outdoor, mobile network testbed to support slivering and utilize a GENI candidate control framework (ORCA). This work includes the following items. 1) A virtualized operating system on the vehicular computers using Xen, which will allow for the isolation of experiments and customization of the operating environment, as well as permitting each experiment to begin from a known configuration. 2) Virtualized access to attached devices, including a 802.11 radio, GPS, and 900 MHz radio on each bus. 3) Standardized, general GENI mechanisms to authenticate users, install customized software and experiments, schedule DOME resources, and to log results. 4) Integration, deployment and access to GENI users.

Prashant Shenoy
UMass Amherst

Sensor Virtualization and Slivering in an Outdoor Wide-Area Wireless GENI Sensor/Actuator Network Testbed (ViSE)

The scope of work on this project is to extend an outdoor, wide-area sensor/actuator network testbed to support slivering and utilize a GENI candidate control framework (ORCA/Shirako), and then bring it into an environment of GENI federated testbeds. This includes: 1) Virtualization of a sensor/actuator system. 2) Integration with GENI-compliant Software Artifacts, including the use of Shirako software (part of the ORCA project) as the base for the control framework. 3) Making the testbed publicly available to GENI users, starting in year 1, and integrate it into an environment of GENI federated testbeds by the end of year 2. 4) Providing documentation for testbed users, administrators, and developers.

Anish Arora
Ohio State University

GENI-fying and Federating Autonomous Kansei Wireless Sensor Networks (KanseiGenie)

The scope of work on this project is to adapt the large-scale prototype wireless sensor networks of Kansei Consortium to provide GENI-compliant interfaces and principles, to provide publically available support for programmability, virtualization, and slice-based experimentation, and to deploy three geographically dispersed sites in a federated arrangement. Also, provide periodic demonstrations, code releases and installation packages. Furthermore, provide an experiment interaction dashboard as well as a scripting environment for composing long-running, complex/phased experiments. And finally, provide the Kansei Doctor, whose multi-level invariant-based detector-corrector modules will help make Kansei more self-repairing and autonomic.

Xiaolin (Andy) Li
Oklahoma State University

OKGems: A GENI-Federated Cyber-Physical System with Multi-Modalities

This project will integrate sensor nodes, robotic platforms and PCs into the ORCA control framework (Cluster D), and make them available for use by GENI researchers. The sensor networks will be deployed in a laboratory setting and in a farmland setting. The sensor nodes will include various kinds of motes, such as Tmote, MicaZ, Telosb, Jennic nodes, and Micadot, and high-end nodes, such as Stargates and Imote2. The robotic platforms include Pioneer 3 DX/AT all-terrain autonomous robots, iRobot Create autonomous robots, and Surveyor robots with cameras and laser range finders.

Dipankar Raychaudhuri
Rutgers University

Open Virtualized WiMAX Base Station Node for GENI Wide-Area Wireless Deployments (WiMAX)

The scope of work on this project is to leverage a commercial IEEE 802.16e WiMAX base station product to prototype an open, programmable and virtualizable base station node that could work over a metropolitan area and connect with off-the shelf WiMAX handsets and data cards. This includes the following steps:

1. Develop an open/virtualizable WiMAX base station with external control and data API's.
2. Import a control framework from the Orbit project, and implement GENI-specific software on an external Linux-based PC controller for the WiMAX base station, demonstrating basic GENI compliance: virtualization, slice setup & resource management, L2/L3 programmability, and opt-in for off-the shelf terminals.
3. Integrate the GBSN with backhaul options to validate and demonstrate the total access networking solution, including VLAN connections to a GENI backbone network.
4. Demonstrate the GBSN capabilities in a phased manner, starting with controlled lab tests of each GENI feature and then migrating to a small-scale outdoor trial with off-the-shelf WiMAX terminal equipment.
5. Develop a deployment plan for multi-cell WiMAX services in GENI, including guidelines for site selection, frequency allocation, network backhaul and equipment/operating costs.

Dirk Grunwald
Univ. of Colorado, Boulder

CR-GENI - Cognitive Radios for GENI Spiral 2 (COGRADIO)

The scope of this project is to develop a set of ruggedized, expandable and configurable multiradio cognitive radio systems that will facilitate experimentation by GENI researchers who have only limited experience with hardware development tools. Two platform designs will span the range of requirements for cognitive radio systems:

- stand-alone cognitive radio system, and
- infrastructure-class cognitive radio system.

Year 1 of the project will emphasize hardware development of a single FPGA system. The software interface provided during year 1 will be a simple "remote radio" system compatible with the Gnu Software Radio framework. This will enable wide-band sampling but may not be suitable for tightly synchronized real-time projects due to the limited interface.

9.6 Instrumentation & Measurement Design and Prototyping

Kara Nance
University of Alaska Fairbanks

Virtual Machine Introspection and Development of a Model Federation Framework for GENI (VMI-FED)

This work develops secure monitoring mechanisms (Virtual Machine Introspection) for virtual machines in GENI slices, and also brings University of Alaska resources into GENI to help develop working implementations of federation and provide a diverse environment for evaluating VMI and other GENI software or experiments. Deployment begins with the Advanced System Security Education Research and Training (ASSERT) Lab, with software based on their previous work with Xen virtual monitoring. The UA CIO is directly involved in the project, and will provide realistic requirements and support for GENI deployments in environments that range from supercomputer centers to geographically isolated communities operating at the extremes of the network's thin edge.

James Griffioen **Instrumentation Tools for a GENI Prototype (INSTOOLS)**

University of Kentucky

The scope of work on this project is to create a GENI-enabled testbed based on the existing University of Kentucky Edulab, and to implement and to deploy instrumentation capabilities that will enable GENI users to better understand the runtime behavior of their experiments.

Priority activities that contribute to this scope include the following items: 1) Provide an early operational federated testbed with administrative support and active student use integrated with a GENI control framework, allowing time for studying and improving the testbed over the period of the contract. 2) Provide an implementation of instrumentation capabilities for researcher and student helper tools that give GENI users the ability to better understand the runtime behavior of their experiments.

Paul Barford **Instrumentation and Measurement for GENI (MeasurementSystem)**

University of Wisconsin, Madison

The scope of work on this project is to develop, test and deploy a prototype implementation of network instrumentation and measurement systems for the GENI infrastructure. Priority activities that contribute to this scope include the following items:

1. Develop an instrumentation and measurement implementation for GENI that provides an interface for researchers, as well an interface to a GENI control framework.
2. Develop an implementation that gives GENI users access to a shared measurement and instrumentation infrastructure that is extensible to several different kinds of measurement devices.
3. Implement systems consistent with the GENI Instrumentation and Measurement Systems (GIMS) Specification, to help reduce risks in this as-yet-unimplemented area of GENI.

Prasad Calyam **OnTimeMeasure: Centralized and Distributed Measurement Orchestration Software**

Ohio Supercomputer Center/OARnet

This effort provides GENI with an on-demand measurement service used in forecasting, anomaly detection, and fault-location diagnosis in GENI experiments and GENI operations. The on-demand measurement service will be integrated into the ProtoGENI control framework and GENI operations (GMOC). The project will deploy a prototype measurement service to support operations and early experimenters use in the first year, and will revise the service in each development spiral to improve services and integration, based on GENI community feedback. The project will also analyze GENI privacy and security requirements for measured data, and prototype service to address appropriate requirements in each development spiral. The project will collaborate with related GENI measurement and security projects (e.g. University of Wisconsin's Instrumentation and Measurement for GENI) on a common GENI instrumentation and measurement architecture. The option years are for possible refinement and integration of OnTimeMeasure into the ProtoGENI control framework and tools.

Martin Swany **Leveraging and Abstracting Measurements with perfSONAR (LAMP)**

University of Delaware

The project will create an instrumentation and measurement system, based on perfSONAR, for use by experimenters on ProtoGENI. The project will collaborate on a plan to develop a common GENI instrumentation and measurement architecture with an emphasis on providing a common but extensible format for data storage and exchange. The project will work to develop a representation of GENI topology to be used to describe measurements and experiment configuration, also based on perfSONAR. The project will collaborate with related GENI measurement and security projects (e.g. University of Wisconsin's Instrumentation and Measurement for GENI) on a common GENI instrumentation and measurement architecture.

Sonia Fahmy **Scalable, Extensible, and Safe Monitoring of GENI (S3MONITOR)**

Purdue University

This effort will develop a prototype shared measurement service based on the existing S3 service, integrate it with ProtoGENI and deploy it for GENI experimenter's use. This shared measurement service will emphasize scalability and safety to best utilize network resources associated with measurements. The project will also analyze GENI privacy and security requirements for measured data, and prototype the service to address appropriate requirements in each development spiral. The project will collaborate with related GENI measurement and security projects (e.g. University of Wisconsin's Instrumentation and Measurement for GENI) on a common GENI instrumentation and measurement architecture.

Keren Bergman **Embedding real-time substrate measurements for cross-layer communications (ERM)**

Columbia University

Ensure GENI includes the technology to support cross-layer communications, specifically, the ability to incorporate a diverse set of real-time measurements in networking protocols. The project addresses the GENI challenge of architectural experimentations across diverse heterogeneous technologies by supporting real-time cross-layer communications and measurements. Our objective is to develop networking capabilities within the GENI infrastructure that enable deeper exposure of cross-layer information and user access to real-time measurements.

Deniz Gurkan **Programmable Measurements over Texas-based Research Network: LEARN**

University of Houston

This project will integrate the Lone-star Education And Research Network (LEARN) regional optical network, located in Texas, into the ORCA control framework (GENI Cluster D). This will enable GENI researchers to utilize the LEARN network for L2 (VLAN) transport. This project will also work with the Embedded Real-time Measurements (ERM) project to develop methods to include the use of existing commercial transport and measurement equipment for real-time measurements. Specifically, it will work with the ERM project to define the API and implement Measurement Handler software to utilize an existing interface (TL1 over SSH) into the Infinera Digital Transport Node (DTN) to make optical measurements.

Rudra Dutta **GENI IMF: Integrated Measurement Framework and Tools for Cross Layer Experimentation**

North Carolina State University

This project will develop and integrate the GENI Integrated Measurement Framework (IMF) for optical communication substrates into the ORCA control framework prototype, and integrate the FIND SILO framework into the ORCA control framework prototype, and IMF and SILO with each other, to enable cross-layer experimentation involving the physical layer of an optical network.

9.7 Experiment Workflow Tools Design and Prototyping

Jeannie Albrech **GENI Experiment Control Using Gush**

Williams College

The scope of work on this project is to design and implement a prototype to build Gush (the GENI User Shell), a robust experiment control and management framework for GENI. The final product of this work is expected to be a fully functional framework that supports experiment control through three user interfaces, including a graphical user interface (GUI), command line interface, and a programmatic interface.

This work will include design of an API in Gush for interacting with GENI Clearinghouses. Gush will integrate with GENI Clearinghouse prototypes provided by the control framework developers to test the Gush API and the use of slice interaction functions. Initial work will be with the PlanetLab control framework using XML-RPC communication with the PlanetLab Central Database (PLC). Development will focus on getting the Gush command line and programmatic interfaces working before moving on to the graphical interface, with an emphasis on detailed error reporting that will simplify debugging.

John Hartman **A Provisioning Service for Long-Term GENI Experiments (Raven)**

Univ. Arizona

The scope of work on this project is to prototype a provisioning service that provides the infrastructure required to develop, deploy, monitor, and maintain long-term and short-term experiments on GENI.

Priority activities that contribute to this scope include the following items:

1. Develop a GENI-specific provisioning service that manages software deployment for slices and also interfaces with GENI clearinghouses. Make the service available to projects beginning in Spiral 1 for at least one clearinghouse, and expanding availability in subsequent spirals.
2. Provide configuration management and resource management for longer-term experiments where software and components can change over the lifetime of the experiment.
3. Provide other researcher helper tools such as monitoring that make it easier to manage experiments.

Beth Plale **netKarma: GENI Provenance Registry**

Indiana University

The project will collect provenance of the data generated by GENI. The provenance of an experiment is its lineage or historical trace that can capture experiment conditions, time ordering, and relationships within the experiment and across the experiment and infrastructure layer. The project will develop a GENI Provenance Registry (netKarma) that will capture the workflow of GENI slice creation, topology of the slice, operational status and other measurement statistics and correlate it with the experimental data. NetKarma will allow researchers to see the exact state of the network and store configuration of the experiment and its slice.

NetKarma is based on the Karma provenance architecture that has been used to collect scientific workflows in diverse domains, such as meteorology and life science.

Michael Freedman **A SCAFFOLD for GENI-based Distributed Services**

Princeton University

The project will build and deploy SCAFFOLD (Service-Centric Architecture For Flexible Object Localization and Distribution), a platform for designing and operating wide-area distributed services on GENI. SCAFFOLD will be a distributed service on GENI that serves as a platform for deploying user-facing services, thereby lowering the barrier for others creating new distributed services. SCAFFOLD will leverage GENI-related prototyping efforts such as OpenFlow switches, the NOX controller, the VINI backbone, the BGP multiplexer and the PlanetLab control framework. The project will use SCAFFOLD to build three services for GENI: A reliable DNS, a port of CoralCDN and a Flash-based video-on-demand service.

Rob Ricci — **Experiment Workflow Tools and Services for GENI (PGTools)**

Univ. Utah

This work will extend the University of Utah Emulab experiment lifecycle management software to become a “full service” front end for experiments on the ProtoGENI control framework and potentially for other frameworks based on the GENI control APIs.

9.8 Security Design and Prototyping

Justin Cappos — **Exploiting Insecurity to Secure Software Update Systems (SECUREUPDATES)**

Univ. Washington

This proposed effort will create a framework that secures the software update systems that operate on GENI. The work will define and implement a security layer that can operate over many different application-specific installation environments, thus providing secure update functions for diverse GENI nodes and clients. The proposal plans to leverage the VM and the redirection proxy from the Million Node GENI project to support multiple platforms. The effort provides secure key management support for software update system developers, allowing software updates to be signed, validated, and distributed securely.

Xiaoyan Hong — **GENI Experiments for Traffic Capture Capabilities and Security Requirement Analysis (ExptsSecurity)**

Univ. Alabama

This effort will help define GENI security requirements based on investigations through ProtoGENI experiments using at least two aggregates, one of which will be a wireless aggregate. The investigations will use network traffic capture and analysis under both normal and misbehaving situations. The well-articulated experiments will analyze the security and privacy characteristics, test GENI traffic recording and capture tools (following the results from Spiral 1 and 2), and make suggestions for improvements to the ProtoGENI control framework with special emphasis on security improvements.

Stephen Schwab — **Distributed Identity and Authorization Mechanisms (ABAC)**

Cobham

This effort will develop and prototype Attributed-Based Access Control (ABAC) extensions that allow the distinct security mechanisms of the various control frameworks to share security information within a single control framework, as well as with each other, starting with ProtoGENI and proceeding to ORBIT and ORCA according to their integration readiness. (Support for DETER’s use of ABAC is already well-established.) The work will support trust management functions, including identity definitions and authentication mechanisms, and distributed authorization and access control mechanisms. Existing ABAC prototype software from SPARTA and other available open-source software will be leveraged to provide critical GENI functions. SPARTA will continue to collaborate with other GENI projects on analyzing and documenting security requirements for each spiral as part of this effort.

Sean Peisert — **The Hive Mind: Applying a Distributed Security Sensor Network to GENI**

UC Davis

This work will develop and prototype a security layer underlying GENI that will allow providers of the system to collaboratively identify and defend against attacks and misuse of GENI resources. Specifically, the effort will develop prototypes for security monitoring, evaluating, and reporting software that could be useful to both GENI experimenters and GENI operations. The effort will also perform experiments using decentralized security algorithms that communicate between sensors in a federated system and evaluate and improve the security layer’s usefulness to potential security experimenters. The project team will collaborate with other teams to develop a security architecture for GENI.

9.9 Early Experiments

Jiang Li **Experimenting on Opportunistic Mobile Wireless Networks in GENI (OppWireless)**

Howard Univ.

This effort will perform early experiments on ProtoGENI with the goal of identifying issues with setting up, running and instrumenting experiments on ProtoGENI/CMU wireless emulator. The project team will document its experiences and provide program examples and libraries that will be useful to other researchers wishing to conduct experiments using ProtoGENI.

S. Felix Wu **DSL (Davis Social Links) on GENI**

UC Davis

This work will port and deploy the Davis Social Links (DSL) system prototype as an early GENI experiment of ProtoGENI. DSL will support a number of advanced communication features such as social keyword based routing, reputation and trust control along the route paths, community-oriented networking, anonymity, and application dualism (a separation between the end2end semantics and social-network based access control). The project will address privacy and other issues related to making traffic from DSL available to other experiments on GENI.

10 Meso-Scale Deployment and Operations

This section provides a brief summary of each of the meso-scale deployment and operations projects active in Spiral 2. Meso-scale deployment and operations projects are focused on establishing operational GENI infrastructure on university campuses and GENI's national backbones.

GENI projects are, in general, funded for 1 year with up to two additional option years. The projects listed here are in their first or second year. The project scope covers the work planned for the entire period. Information about the specific plans during Spiral 2 can be found on the Spiral 2 wiki page at http://groups.geni.net/geni/wiki/SpiralTwo.

10.1 OpenFlow Meso-Scale Deployment

Nick McKeown — **Enterprise GENI: Prototype, Demonstrate and Create Kits for Replication (E-GENI)**
Stanford University

The scope of work on this project is to illustrate how GENI can be deployed on local networks, such as campus and enterprise networks, and to develop a kit that allows the work demonstrated in this project to be easily replicated elsewhere. Priority activities that contribute to this scope include the following items:

1. Deploy Enterprise GENI technology in a campus (wireline and wireless).
2. Integrate an OpenFlow network with a GENI control framework in an Aggregate Component Manager.
3. Provide access to Enterprise GENI testbeds for GENI users.
4. Define an Enterprise GENI deployment kit for research and potential commercial transition.

Kuang-Ching Wang — **OpenFlow Campus Trials at Clemson University (OFCLEM)**
Clemson University

This project will deploy an OpenFlow testbed on the Clemson University campus and connect with wireless mesh access points and mobile terminals. This trial will conduct OpenFlow experimentation focused on OpenFlow enabled network operation solutions as a precursor to deployment into Clemson research and production networks.

Nick Feamster — **OpenFlow Campus Trials at Georgia Tech (OFGT)**
Georgia Tech University

This project will deploy an OpenFlow testbed in the CS building on the Georgia Tech Campus. This project will conduct OpenFlow experimentation focused on OpenFlow enabled next generation network admission control system as a precursor to deployment into campus residential network. The testbed will replicate a campus deployment required to handle thousands of end hosts.

Christopher Small — **OpenFlow Campus Trials at Indiana University (OFIU)**
Indiana University

This project will deploy an OpenFlow testbed between two campuses. This trial will conduct OpenFlow experimentation focused on OpenFlow enabled dynamic provisioning and distributed monitoring as a precursor to deployment into the Indiana University campus production network. It will also examine and propose ways for network operators to manage and support their networks using Openflow enabled devices.

Ivan Seskar **OpenFlow Campus Trials at Rutgers University (OFRG)**

Rutgers University

This project will deploy OpenFlow switches between wireless research networks and Internet2 to enable researcher programmability and virtualization in the wired network connecting the Rutgers wireless assets to the GENI backbone. OpenFlow switches in the ORBIT indoor mesh testbed will also enhance performance and flexibility and provide support for grid virtualization.

Aditya Akella **OpenFlow Campus Trials at University of Wisconsin (OFUWI)**

University of Wisconsin, Madison

This project will deploy an OpenFlow testbed throughout the University of Wisconsin campus. This project will develop OpenFlow extensions focused on the creation of flexible, heterogeneous slices, the development of a generic load balancer and a captive portal for authentication. This work is a precursor to the deployment of OpenFlow into campus production networks.

Arvind Krishnamurthy **OpenFlow Campus Trials at University of Washington (OFUWA)**

University of Washington, Seattle

This project will deploy a hybrid OpenFlow and RouteBrick testbed within the computer science and engineering department. This project will develop building blocks allowing researchers to investigate the placement of middlebox functionality in enterprise networks. This work is a precursor to the deployment of OpenFlow and RouteBricks into campus production networks.

Michael Freedman **OpenFlow Campus Trials at Princeton University (OFPR)**

Princeton University

This project will deploy an OpenFlow testbed within the CS self-operated IT department. This trial will conduct OpenFlow experimentation focused on OpenFlow enabled management of experimental and production services as a precursor to deployment into their department's production network.

Nick McKeown **OpenFlow Campus Trials at Stanford University (OFSTAN)**

Stanford University

This project will deploy an OpenFlow, testbed throughout the Stanford University campus. This project will develop OpenFlow extensions focused on the creation of flexible, heterogeneous slices, the development of a generic load balancer and a captive portal for authentication. This work is a precursor to the deployment of OpenFlow into campus production networks.

Heidi Dempsey **OpenFlow Campus Trials at BBN (OFBBN)**

Raytheon BBN/GPO

This project will deploy an OpenFlow testbed for development, experiment, and integration testing at BBN. The testbed will include multiple OpenFlow switches and wireless networks, and will support both experimental and production uses.

Martin Casado **Nicira Networks Development and Support for OpenFlow Campus Trials (OFNOX)**

Nicira Networks

Nicira Networks will support enhancements to the NOX software base used in OpenFlow, and provide trouble-shooting support for OpenFlow shakedown in campus and backbone environments in GENI. Campuses and backbones are expected to use NOX for at least three years, with Nicira's help (BBN and Stanford will also help support campuses). The primary enhancements to NOX are 1) Upgrades for NOX compatibility with OpenFlow 1.0 in the GENI environment, 2) Design improvements that make it easier to install and run NOX in GENI and 3) a new console interface (SNAC) intended primarily for Campus IT staff to use to manage OpenFlow on their campuses.

Eric Boyd **GENI OpenFlow Backbone Deployment at Internet2 (OFI2)**

Internet2

Internet2 will implement Stage 0 of the Internet2 OpenFlow test network by creating a five node OpenFlow test network using small enterprise switches (HP ProCurve 5406 switches proposed, pending evaluation). The initial design has 10GE WAN interfaces, but that is subject to negotiation of the MOU Internet2 has with GENI for wide-area connectivity. Stage 0 comprises the appropriate OpenFlow controller(s), interfaces to connect several enterprise OpenFlow campuses, and interfaces to connect other projects such as ProtoGENI and SPP. The Stage 0 network will allow members and non-members of Internet2 to connect to GENI OpenFlow services.

Internet2 will create OpenFlow nodes in five cities: Atlanta, Chicago, Houston, New York, and Salt Lake City. Internet2 will also explore integrating the OpenFlow nodes with OpenFlow deployments on ProtoGENI nodes. In this second, complementary-to-ProtoGENI scenario, Internet2 will create OpenFlow nodes in five cities: Atlanta, Chicago, Los Angeles, New York, and Seattle.

Glenn Ricart **GENI OpenFlow Backbone Deployment at National LambdaRail (OFNLR)**

National LambdaRail

Deploy and operate OpenFlow-enabled HP Procurve 6600 switches at 5 NLR POPs and interconnect NLR's FrameNet network to the GENI OpenFlow-enabled backbone, permitting members and non-members of NLR to connect to GENI OpenFlow services.

10.2 WiMAX Meso-Scale Deployments

Dipankar Raychaudhuri **Open GENI WiMAX Base Station Kit for Campus Deployments (WIMXRG)**

WINLAB, Rutgers University

This project is aimed at the development, procurement and software support of a GENI-enabled WiMAX base station equipment kit for deployment on seven campuses. This is follow-on to the current Spiral 1 project on the prototyping and trial deployment of an open GENI WiMAX 802.16e base station. The project leverages a commercial 802.16e base station from NEC, replacing the standard WiMAX controller with an open GENI software implementation that supports virtualization and layer 2/3 programmability.

Henning Schulzrinne **GENI WiMax Base Station Kit - Columbia University Deployment (WIMXCOLUM)**

Columbia University

This project will install a GENI WiMax Base Station Kit provided by Rutgers on the Columbia University campus, and operate it with the ORBIT Management Framework (OMF). It will enhance our WORKIT nodes with WiMax? clients and GPS capabilities, add remote control and virtualization capabilities, and operate them with OMF. It will also build teaching modules to explore wide area wireless network issues and applications.

Thanasis Korakis
Polytechnic Institute of New York University

Campus deployment of an open GENI WiMAX 802.16e Testbed in Polytechnic Institute of NYU (WIMXNYUPT)

This project builds an open/programmable WiMax testbed on the campus of Polytechnic University of NYU. It utilizes the "Open GENI WiMax Base Station Kit for Campus Deployments" provided by an associated project (Rutgers and NEC). It provides for the installation, maintenance and integration of the WiMax base station kit with existing experimentation facilities at the university.

Mario Gerla
UCLA

C-VeT --UCLA Campus Vehicular Testbed: An Open Platform for Vehicular Networking (WIMXUCLA)

This project will add WiMax capability to a Campus Vehicular Testbed (C-VeT) to support the deployment and testing of vehicular services and applications over both WiFi and WiMax. It will feature resource virtualization in order to enable shared access by multiple experimenters using the ORBIT Management Framework (OMF). To establish the WiMax capability, we will procure a "GENI WiMax base station kit".

Dirk Grunwald
University of Colorado at Boulder

Deploying GENI WiMax in Colorado (WIMXUCO)

This proposal will install and maintain a programmable WiMax base station at one of the existing University of Colorado Wide-Area Radio Testbed (WART) wireless network sites. It will use a "GENI WiMax Base Station Kit" provided by Rutgers.

Mark Corner
UMass Amherst

WiMax Deployment in the DOME and ViSE Testbeds and Cluster-D Framework (WIMXUMASS)

This project will deploy an NEC WiMax base station (provided by Rutgers through an associated WiMax base station proposal) on a building located at UMass Amherst, and deploy WiMax client stations on mobile vehicles to enhance our existing GENI DOME testbed.

Suman Banerjee
University of Wisconsin, Madison

A Programmable Facility for Experimentation with Wireless Heterogeneity and Wide-area Mobility (WIMXUWI)

This project will integrate the currently operational, heterogeneous WiNGS-outdoor testbed spanning 60 square miles in and around Madison, WI, called MadBed, into GENI, and provide systematic low-level access to PHY/MAC features using the ORBIT Management Framework (OMF). It will add a "GENI WiMax Base Station Kit", provided by an associated proposal, to the testbed to complement an existing Cisco/Navini WiMax base station. It will add a mobility services engine to the testbed, and enhance the user nodes, to allow the testbed to support a wide range of wireless experiments. It will provide a unique trial mechanism for opt-in users in the broader Madison area.

Tony Michel
Raytheon BBN Technologies

GENI WiMax Base Station Kit installed for GPO at BBN Technologies (WIMXBBN)

This project will install a GENI WiMax Base Station Kit provided by Rutgers at BBN Technologies in the GENI Testbed for use by the GENI Project Office (GPO), and operate it with the ORBIT Management Framework (OMF). It will enable the GPO to better understand WiMAX technology, and assist others in planning and conducting WiMAX experiments.

10.3 Regional Networking

Jen Leasure **Realizing GENI in the Three-Tier R&E Infrastructure (QUILT)**

The Quilt

This project will convene a workshop to explore issues and ideas for using GENI in regional and state networks that can interoperate with early spiral GENI campus and backbone network projects. The Quilt will solicit a wide range of participants from the three network tiers, as well as others in the R&E community, and in GENI to get a wide range of inputs and ideas for current and future GENI development.

10.4 Meso-scale Measurement

James Griffioen **A ShadowBox-based ProtoGENI Instrumentation and Measurement Infrastructure (ShadowNet)**

University of Kentucky

This proposal will deploy GENI-enabled commercial routers into the ProtoGENI backbone, develop S/W to enable per- slice monitoring and measurement, and develop tools and interfaces to control and access the measurement infrastructure. This proposal leverages measurement software developed at AT&T Research and combines it with measurement software developed at the University of Kentucky

10.5 Operations

Jon-Paul Herron **GENI Meta Operations (GMOC)**

Indiana University

The scope of work on this project is to facilitate the sharing of operational and experimental information among GENI experimental components. This effort has both technical development and operational requirements. Technically, the GENI Meta Operations Center (GMOC) would require a well-defined protocol to help generalize the operational details of GENI prototypes and for the providers of prototypes to send those details to an operational data repository. These requirements suggest a modular approach, with a generalized protocol rather than a restricted set of hardware and software that GENI prototype participants would be required to run. In other words, it would be largely up to the GENI Spiral project investigators to decide what data to share and how to collect this data from their prototype infrastructure. The GMOC would provide the standardized format for this data and the systems required to share, monitor, display, and act on this data. In addition, the GMOC could be used to help provide a repository for data collections passing into and out of GENI prototypes for the purpose of discovering and isolating prototypes that have caused problems. This might require additional instrumentation at the connection points and substrate elements between prototypes. This would be accomplished with the help of the other prototypes that are part of GENI Spirals. The GMOC will work with these other projects to develop the operational data formats, processes, and systems needed for a consistent and useful suite of GENI infrastructures. During the project, participants will investigate how a Meta Operations Center might interact with various prototype participants to accomplish operations functions.

10.6 Security

Stephen Schwab **GENI Security Architecture Toolkit (SECARCH)**
Cobham

This effort will define a GENI Security Architecture, in support of the broad goals for GENI Spirals 1-3, including (a) working with teams prototyping multiple control frameworks and (b) demonstrating end-to-end slicing across a range of technologies, including Ethernet VLANs as an initial universal service offered by GENI. By interacting continuously with testbed prototyping efforts, we will jointly refine the security requirements, reflect those requirements within the security architecture, and validate through feedback gleaned from our collaborator's rapid deployment cycles that our security architecture concepts are indeed aligned and addressing the needs of the GENI testbed community.

Adam Slagell **Developing a Comprehensive GENI Cyber Security Program (CompSec)**
NCSA

This effort will develop an evolving GENI Cyber Security Program through detailed analysis of GENI architecture documents, specific implementations, and project work, using an open process the encourages community feedback. The project team will leverage experience with security planning for the National Center for Supercomputing Applications, the NSF TeraGrid? and Open Science Grid projects, the Blue Waters Project and the National Center for Digital Intrusion and Response, broadening the experimental community's interactions with GENI. The project will produce interim and longer-term GENI security plans, guidelines and procedures, and will collaborate closely with other GENI security projects, particularly SPARTA's.

Kenneth J. Klingenstein **Leveraging Emergent Federated Activities (LEFA)**
Internet2

The work will focus on enabling GENI for federated identity, developing options for supporting virtual organizations within the GENI community, brokering GENI's international requirements at the middleware layer, providing white papers on critical GENI topics in which the Internet2 middleware community has established expertise, requirements gathering for federated authentication/authorization and attribute aggregation across the GENI clusters, and on working closely with other peer proposals investigating related themes.

10.7 Federation

James Williams **K-GENI: Establishment of Operational Linkage between GENI and ETRI/KISTI-Korea for International Federation**
Indiana University

This effort will support early trial linkage of research network operators as part of understanding management of a global-scale, federated infrastructure suite that can support GENI. GMOC engineers and ETRI/KISTI engineers will use a donated 1 Gbps international connection between Korea and Indiana University to test methods of interoperability between GMOC and the dvNOC system. The dvNOC is being developed as a part of the GLORIAD-KR project funded by Korean government and is in use by ETRI/KISTI. Examples would include testing a data stream of "externally significant" operational data between GMOC and dvNOC, including things like outage information and maintenance notification.

Kara Nance
University of Alaska Fairbanks

Virtual Machine Introspection and Development of a Model Federation Framework for GENI (VMI-FED)

This work develops secure monitoring mechanisms (Virtual Machine Introspection) for virtual machines in GENI slices, and also brings University of Alaska resources into GENI to help develop working implementations of federation and provide a diverse environment for evaluating VMI and other GENI software or experiments. Deployment begins with the Advanced System Security Education Research and Training (ASSERT) Lab, with software based on their previous work with Xen virtual monitoring. The UA CIO is directly involved in the project, and will provide realistic requirements and support for GENI deployments in environments that range from supercomputer centers to geographically isolated communities operating at the extremes of the network's thin edge.

Larry Peterson
Princeton University

Understanding Federation (PLFED)

This effort will integrate PlanetLab (PLC), PlanetLab Europe (PLE), PlanetLab Japan (PLJ)-along with other testbeds in Korea, Brazil, Europe, Japan, and the US-into an international federated research infrastructure. We will deploy and use the federation mechanisms developed as part of the PlanetLab cluster currently funded by the GPO, and focus our attention on the policy issues that arise when autonomous organizations federate their networks together. Funding is requested for Princeton (the prime) along with Université Pierre et Marie Curie (UPMC) and the University of Tokyo (subcontractors), which will drive the effort, but multiple testbeds will be federated with the support of this proposal. The effort will focus on resolving the policy issues related to federation-Princeton, UPMC, and the University of Tokyo already have funding to perform the necessary engineering work to build the mechanisms required by federation.

John Wroclawski

TIED: Trial Integration Environment in DETER

The scope of work on this project is to develop and evangelize a control framework that particularly emphasizes usability across different communities, through federation, rich trust/security models, and similar enabling mechanisms.

11 Studies

This section provides a brief summary of each of the GENI studies active in Spiral 2. Studies are small projects that tend to produce analyses that are of general use to the GENI development effort.

GENI projects are, in general, funded for 1 year with up to two additional option years. The projects listed here are in their first or second year. The project scope covers the work planned for the entire period. Information about the specific plans during Spiral 2 can be found on the Spiral 2 wiki page at http://groups.geni.net/geni/wiki/SpiralTwo.

Pierre F.Tiako
Langston University

Use Case Scenarios for promoting GENI at four-year Colleges and Minority Institutions (GeniFourYearColleges)

Working with the GPO during the first year, Langston University (LU) team involved in this project will assess the substrate resources offered in spiral-1 and help to define potential experiment scenarios which could serve as topics of interest to four year colleges and minority institutions. We will also use the capabilities offered in GENI spiral-1 prototypes to shape and propose use-case scenarios. Within the second year, LU team involving undergraduate research assistant students will apply one of the use-case scenarios and actually conduct a (limited scope) experiment on GENI prototypes using one of the prototype GENI control frameworks.

Matt Bishop
University of California at Davis

Attribution for GENI

The project will produce a report describing the requirements of an attribution infrastructure that will support experiments to test the effectiveness of attribution features, with special focus on requirements for GENI and for providing a platform for attribution systems to be incorporated into ongoing GENI project work.

James Sterbenz
University of Kansas

Measurement, Monitoring, and Outreach in the Great Plains Environment for Network Innovation (MMOGpENI)

This project will analyze development and prototyping requirements in a federated international network. Specifically, the team will work with collaborators in Europe to determine how they can use existing international research and education infrastructure (e.g. DANTE, NORDUnet) to affiliate with GENI clusters, contribute resources and software, and run early GENI experiments. The team will coordinate with other funded GENI federation projects to keep the requirements results consistent with those projects.

Kuang-Ching (KC) Wang
Clemson University

A Mobile Programmable Radio Substrate for Any-layer Measurement and Experimentation (MPRADIO)

This project will study and report on the capabilities recommended for a programmable radio substrate in GENI to best support wireless networking innovations. It is expected that a key capability should be to provide programmability and measurement at all layers. This project will also recommend the capabilities that should be included in the cognitive radio systems that are being developed in the “Cognitive Radios for GENI Spiral 2” project.

Optical Access Networks
SUNY at Buffalo

Optical Access Networks

This project will study how to make an optical edge network, such as a campus-scale optical access network based on passive optical network (PON) technologies, a part of the GENI infrastructure. The facility to be studied shall initially be a campus/city wide Passive Optical Network (PON) based on hybrid TDM/WDM technologies. The study shall involve issues related to both hardware and software, with a focus on its programmable and virtualizable aspects per need of GENI. It shall also address issues related to the applicability of one or more GENI control frameworks. The usage scenarios shall also be investigated.

Karen Sollins
Massachusetts Institute of Technology

Design of an information substrate for measurement, monitoring, and inference (InfoSubstrate)

This effort will produce two study papers. The first will describe the increasingly complex set of system design issues involved in building an information plane for both experimenters and the various players providing network management. The second will be a systems design paper. This problem is challenging given that information may be collected or inferred in one domain in a federated environment, such as GENI, where there may be conflicting objectives across either the boundaries of federation of the infrastructure or between the GENI control and management plane and the experimental environments.

12 Acronyms

12.1 Acronyms

AVP	Assistant Vice President
CIO	Chief Information Officer
CTO	Chief Technology Officer
GEC	GENI Engineering Conference
GENI	Global Environment for Network Innovations
GPO	GENI Project Office
I&M	Instrumentation & Measurement
NOC	Network Operations Center
O&M	Operations and Management
SE	System Engineer
TBD	To be determined
TBR	To be reviewed
TBS	To be specified
WiMax	Worldwide Interoperability for Microwave Access, Inc
WG	Working Group

12.2 Key GENI Terms

Term	Explanation
Aggregate	An *aggregate* is an object representing a group of components, where a given component can belong to zero, one, or more aggregates. Aggregates can be hierarchical, meaning that an aggregate can contain either components or other aggregates. Aggregates provide a way for users, developers, or administrators to view a collection of GENI nodes together with some software-defined behavior as a single identifiable unit. Generally aggregates export at least a component interface, i.e., they can be addressed as a component, although aggregates may export other interfaces, as well. Aggregates also may include (controllable) instrumentation and make measurements available. This document makes broad use of aggregates for operations and management. Internally, these aggregates may use any O&M systems they find useful.

Term	Explanation
Clearinghouse	A *clearinghouse* is a, mostly operational, grouping of a) architectural elements including trust anchors for Management Authorities and Slice Authorities and b) services including user, slice and component registries, a portal for resource discovery, a portal for managing GENI-wide policies, and services needed for operations and management. They are grouped together because it is expected that the GENI project will need to provide this set of capabilities to bootstrap the facility and, in general, are not exclusive of other instances of similar functions. There will be multiple clearinghouses that will federate. The GENI project will operate the NSF-sponsored clearinghouse. One application of 'federation' is as the interface between clearinghouses.
Components	A *component* might correspond to an edge computer, a customizable router, or a programmable access point. A component encapsulates a collection of resources, including physical resources (e.g., CPU, memory, disk, bandwidth) logical resources (e.g., file descriptors, port numbers), and synthetic resources (e.g., packet forwarding fast paths). These resources can be contained in a single physical device or distributed across a set of devices, depending on the nature of the component.
End-user Opt-In	An important feature of GENI is to permit experiments to have access to end-user traffic and behaviors. For examples, end-users may access an experimental service, use experimental access technologies, or allow experimental code to run on their computer or handset. GENI will provide tools to allow users to learn about an experiment's risks and to make an explicit choice ("*opt-in*") to participate.
Experiment	An *experiment* is a researcher-defined use of a slice; we say an experiment runs in a slice. Experiments are not slices. Many different experiments can run in a particular slice concurrently or over time.
Federation	Resource *federation* permits the interconnection of independently owned and autonomously administered facilities in a way that permits owners to declare resource allocation and usage policies for substrate facilities under their control, operators to manage the network substrate, and researchers to create and populate slices, allocate resources to them, and run experiment-specific software in them.
Resource	*Resources* are abstractions of the sharable features of a component that are allocated by a component manager and described by an RSpec. Resources are divided into computation, communication, measurement, and storage. Resources can be contained in a single physical device or distributed across a set of devices, depending on the nature of the component.
Slices	From a researcher's perspective, a *slice* is a substrate-wide network of computing and communication resources capable of running an experiment or a wide-area network service. From an operator's perspective, slices are the primary abstraction for accounting and accountability—resources are acquired and consumed by slices, and external program behavior is traceable to a slice, respectively. A slice is the basis for resource revocation (i.e., shutdown). A slice is defined by a set of slivers spanning a set of network components, plus an associated set of users that are allowed to access those slivers for the purpose of running an experiment on the substrate. That is, a slice has a name, which is bound to a set of users associated with the slice and a (possibly empty) set of slivers.

Term	Explanation
Slivers	It must be possible to share component resources among multiple users. This can be done by a combination of virtualizing the component (where each user acquires a virtual copy of the component's resources), or by partitioning the component into distinct resource sets (where each user acquires a distinct partition of the component's resources). In both cases, we say the user is granted a *sliver* of the component. Each component must include hardware or software mechanisms that isolate slivers from each other, making it appropriate to view a sliver as a "resource container."