

## GENI: An Introduction

Harry Mussman<sup>1</sup>

GENI Project Office at Raytheon BBN Technologies  
Cambridge, MA 02138, USA  
hmussman@bbn.com  
<http://www.geni.net>

**Abstract.** The Global Environment for Network Innovations (GENI) is a suite of infrastructure being designed to support “at scale” research in networking and innovation in services. Efforts are underway to extend the GENI infrastructure within the next 3 years by: a) adding approximately 54 GENI racks of compute and network resources; adding L2 connectivity with OpenFlow features to regional and backbone networks; c) introducing tools and improving integration so that large-scale experiments can be setup and run with a reasonable effort; and d) adding WiMAX base stations to extend the reach of GENI. The goals are to support large-scale GENI-wide research and service experiments within 3 years, and to permit campus, city and WiMAX-connected users from many locations to opt-in to the same service experiment.

### 1 Introduction

The Global Environment for Network Innovations (GENI) [1] is a suite of infrastructure now being designed to support “at scale” research in networking and related disciplines, and rapid innovation in novel, large-scale services.

Key GENI concepts include: a) the deep programmability of all resources, including compute and network resources; b) a virtualized (shared) infrastructure, where each GENI experimenter (or other GENI user) gathers resources into their own isolated “slice” of resources, and then configures them to support their experiment; c) a federated heterogeneous infrastructure, that is evolving over time.

GENI resources are gathered into aggregates which include: a) compute aggregates, that offer servers and Virtual Machines (VMs); b) network aggregates, that offer Layer 2 (L2) connections with optional OpenFlow (OF) [2] features; c) wireless aggregates, that establish on-the-air WiMAX sites [3], with programmable Base Stations [4] and programmable Mobile Stations (MSs); and d) sensor aggregates, that provide access to programmable sensors, such as the ViSE system. [5]

GENI entered its prototyping phase in 2007, when NSF awarded the GENI Project Office (GPO) role to Raytheon BBN Technologies. The GPO subsequently initiated a community-based design and planning process, and issued a series of academic and industrial subcontracts to build, integrate, and operate a prototype GENI suite of infrastructure utilizing a “spiral” or iterative process. The current GENI suite includes the efforts of over 92 projects, spanning three years of effort in Spirals 1, 2 and 3.

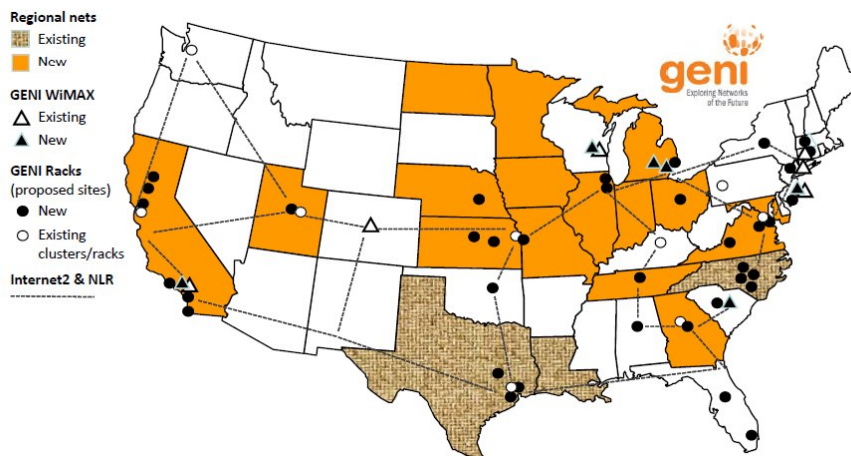
---

<sup>1</sup> This introduction is based on a research paper entitled “Extending GENI to Support Large-Scale Research and Service Experiments,” by the same author.

The building of “meso-scale” prototype GENI infrastructure was funded by the NSF [6] and started in October 2009 at the beginning of Spiral 2. At the end of Spiral 3 GENI-enabled commercial hardware was deployed across 13 university campuses [7], including 10 clusters of servers and VMs, plus 8 WiMAX sites, linked by build-outs through two US national research backbones (Internet2 and NLR).

In October 2011 at the beginning of Spiral 4, new “extended meso-scale” projects were approved by the NSF [8] to greatly extend the prototype GENI infrastructure within the next 3 years. This paper will summarize the planned extensions, and identify the tasks that must be completed, so that the GENI infrastructure can support large-scale research and service experiments by the end of Spiral 6.

The extended meso-scale projects will add approximately 54 pre-packaged GENI racks to campuses at sites shown in Fig. 1; deployment starts in Spiral 4, and will be completed in Spiral 6. Each rack includes a compute aggregate and a network aggregate. The compute aggregate includes servers and VMs that provide programmable, virtualized compute resources to GENI experimenters. The network aggregate includes a programmable L2 switch with OpenFlow features to provide connections to the GENI backbones via regional access networks.



**Fig. 1.** GENI rack and WiMAX site build-outs by extended meso-scale projects

In addition, this effort will grow the OpenFlow footprint in 6 major regional research networks across 16 states and in the Internet2 and NLR national backbones.

Section 2 of this paper explains the configuration of the racks and their connectivity, and shows how they will be used for large-scale research and service experiments.

The extended meso-scale projects will also grow the GENI WiMAX footprint by extending 3 existing WiMAX sites to include multiple base stations, and by adding 3 new WiMAX sites, as shown in Fig. 1. Deployments are planned to be completed in Spiral 4, and fully integrated with GENI compute and network aggregates by the end of Spiral 6.

The use of the latest WiMAX base-station technology is expected to extend coverage at all of these sites, both on campus and in nearby city areas. Furthermore, those sites with multiple base stations are being designed to support vehicular mobile experiments.

Section 3 of this paper explains the configuration of these new and extended WiMAX wireless aggregates, and shows how they can be used for experiments by local and/or remote GENI experimenters.

As these WiMAX sites become fully integrated with GENI compute and network resources, resources from multiple WiMAX sites can be combined to support GENI-wide research and service experiments.

## **2 Adding GENI Racks to GENI-enable a Campus**

Currently, 10 clusters of compute (servers and VMs) and network aggregates are deployed on “GENI-enabled” campuses, and connected via the regional access and national backbone networks as shown Fig. 1. These deployments include 2 PlanetLab clusters with Internet Layer 3 (L3) connectivity and 8 ProtoGENI clusters with L2 and L3 connectivity.

The extended meso-scale projects will add approximately 54 pre-packaged GENI racks of compute and network resources to campuses at sites shown in Fig. 1.

Pre-packaging the racks greatly simplifies their installation and operation in a campus environment, and makes it much easier to “GENI-enable” a campus.

One reason for installing the racks on campuses is to allow campus “opt-in users” to be directly connected to service experiments running in GENI slices. For example, the traffic from users on campus could be routed by the campus IT network to the network aggregate in the GENI rack, where it could be routed into a GENI slice running a service experiment. Because a GENI slice can include resources from multiple campus racks, this allows “opt-in users” from several campuses to participate in a single GENI service experiment at the same time.

In addition, several US cities have expressed an interest in deploying GENI racks so that their citizens can participate in service experiments running in GENI slices. This would allow thousands of “opt-in users” to be included in a GENI service experiment. Furthermore, it could allow a service experiment to be tailored to the needs of an individual city. Thus, it is expected (but not yet confirmed) that some racks will be deployed in US cities.

### **Configuration of a GENI Rack**

Each rack includes a compute aggregate and a network aggregate, accessed via one Aggregate Manager (AM) service. A typical rack configuration from the ExoGENI project [9] at RENCi is shown in Fig. 2. The compute aggregate includes VMs on worker nodes that provide programmable, virtualized compute resources to GENI experimenters. In addition, there are storage resources that can be assigned to a slice.

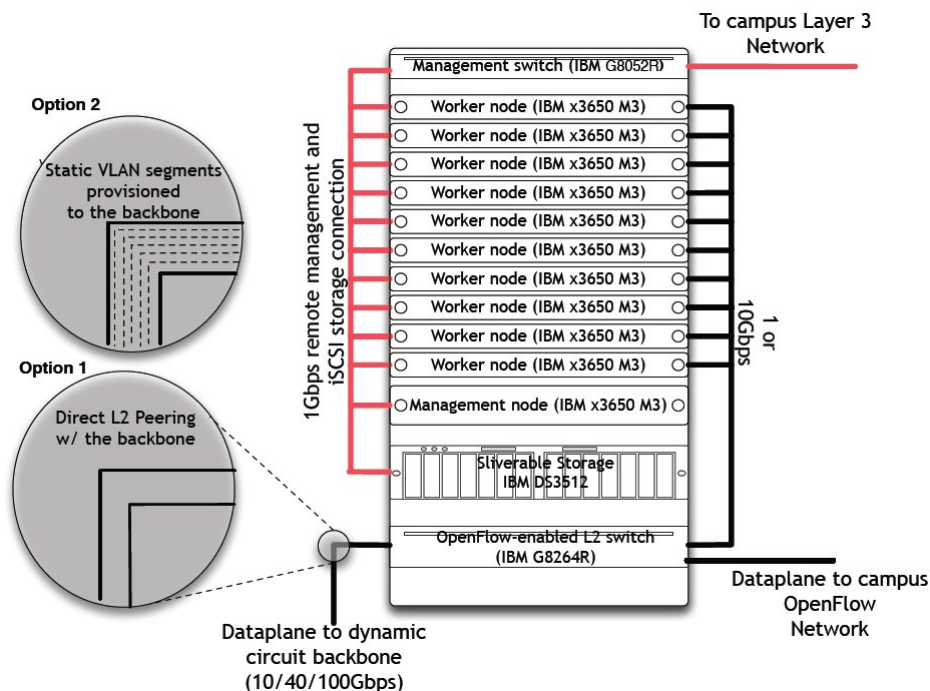


Fig. 2. GENI rack configuration from ExoGENI project [9] at RENCi

This rack configuration includes a management node that provides resource assignment plus rack management and monitoring functions. Resource assignment is done by an AM service running on the management node, exporting the GENI AM API [10].

The management node and the worker nodes are connected via a management switch to the Internet for GENI Control Plane functions.

The network aggregate includes a programmable L2 switch with OpenFlow features to provide connections to the GENI backbones via regional access networks for GENI Data Plane functions. In addition, the L2 switch will also accept L2 traffic from a campus (or city) network serving “opt-in users”.

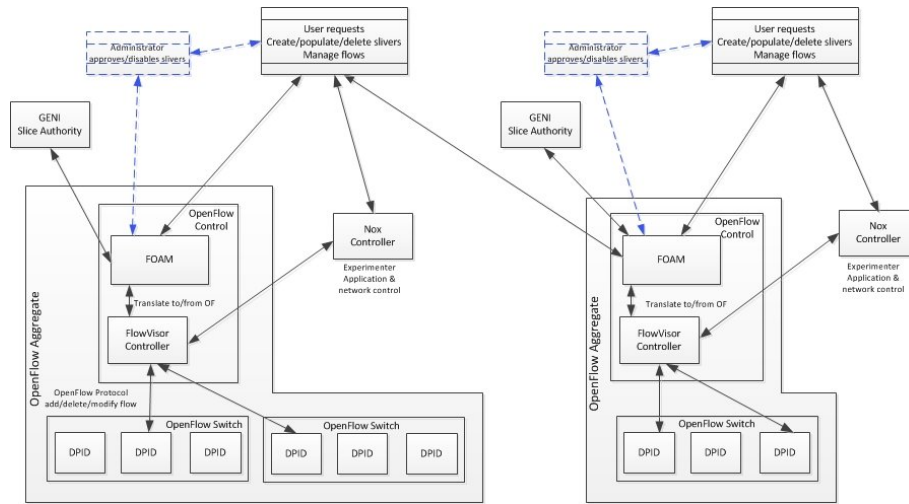
The L2 switch connects to the OpenFlow footprint that is being established in 6 major regional research networks across 16 states, and in the Internet2 and NLR national backbone networks, as identified in Fig. 1.

Network connections are setup by the AM service. It can directly establish L2 connections through the L2 switch as part of a long L2 path, e.g., a L2 VLAN that spans multiple racks.

When a path with OpenFlow features is desired, the OpenFlow-enabled switches are configured to check with the FlowVisor and Nox controllers before a flow is established, as shown in Fig. 3. These controllers are programmed and configured through the AM to realize the desired flows.

Three types of GENI racks are planned to be deployed, utilizing the cluster software developed by three different GENI projects: PlanetLab [11], ProtoGENI [12] and ORCA [13]. All now support the common GENI AM API for assigning resources, in addition to their own AM APIs. This will enable a GENI experimenter to assemble a slice of resources from multiple compute and network aggregates of different types using only one API.

## GENI: An Introduction



**Fig. 3.** Controlling an OpenFlow enabled L2 switch

### Assembling a GENI Slice for a Large-Scale Experiment

In GENI, an experimenter gathers resources from GENI aggregates into their own dedicated “slice” of resources using experiment control tools that utilize the GENI AM APIs exported from the AM on each aggregate. The experimenter then programs and/or configures the resources in their slice to conduct their experiment.


For a small-scale experiment, this may involve a limited number of VMs and L2 paths, perhaps from only one rack. Then, an experimental tool that uses a graphical interface to describe the desired topology and observe the experiment can be used; one example is the FLACK tool [14].

For a large-scale research experiment, an experimenter may gather into a large-scale slice: a) multiple VMs from multiple racks; b) connected using multiple L2 paths, without or with OpenFlow features; c) using network resources on multiple racks, in multiple regional network and in a backbone network; d) with paths “stitched” together to provide the desired topology. A representative large-scale slice is shown in Fig. 4.

As GENI grows its meso-scale resources from the current 10 to 50 or more sites, there are many challenges to setting up such a slice and running a large-scale experiment.

First, script-driven experiment control tools are required so that the experimenter can efficiently and repeatably specify a complex topology, acquire all desired resources, program and configure them, and then orchestrate the resultant experiment. Graphical displays of topology and status are required to allow the experimenter to best observe and run their experiment. Continuous tool development is expected through Spiral 6 to accomplish these tasks.

Second, stitching the L2 paths together to form the desired topology is currently a manual process, and is particularly difficult when multiple types of cluster/tracks must be stitched together. Much effort is underway to introduce automated stitching tools [15] by Spiral 5, and then these will have to be extended through Spiral 6 to instantiate larger and more complex topologies.



**Fig. 4.** Assembling a large slice from resources in multiple racks and in multiple networks

Third, the initial quality of a large-scale slice must be high, and it must remain stable over a long period of time. Benchmarking experiments of slice quality and stability [16] have been started by GENI operations staff, who gather resources into a large slice, run a test application, and keep the slice for an extended period of time. Based upon these benchmarking experiments, continuous improvements must be made to raise the quality and stability of slices through Spiral 6, even as their size and complexity is increased.

Finally, for a large-scale service experiment, the experimenter also includes in their slice: e) traffic from “opt-in users”, also shown in Fig. 4. This may include users from multiple campuses and/or cities, to provide the experiment with diverse traffic. If the service experiment continues for a long period of time, and the users begin to depend on the service, then the GENI “service experiment” becomes a real “service trial”. A task is underway to introduce traffic from opt-in users by Spiral 5.

### 3 Adding and Extending WiMAX Wireless Sites

The GENI WiMAX deployments allow researchers to experiment with open and programmable “4G” wireless systems. Many researchers believe that user interaction with “future internets” will take place primarily through wireless devices, and this is a practical approach to making such experimentation widely available to the research community.

The GENI WiMAX sites are setup so that they can be used by local and/or remote GENI experimenters. Furthermore, they will be fully integrated with GENI compute and

network resources by the end of Spiral 6, so that resources from multiple WiMAX sites can be combined to support GENI-wide research and service experiments.

The initial meso-scale projects installed sliceable and programmable WiMAX base station kits at 8 sites, including 7 campuses and at Raytheon BBN Technologies in Cambridge, MA. The extended meso-scale projects will grow the GENI WiMAX footprint using next-generation WiMAX radio hardware, by extending 3 existing WiMAX aggregates to include multiple base station kits, and by adding 3 new WiMAX aggregates. All existing and new sites are shown in Fig.1. New base stations are planned to be installed by the end of Spiral 4.

All sites operate in the 2.5GHz Educational Broadband Services band using an experimental license obtained from the FCC or on spectrum “on loan” from Clearwire. WiMAX technology was chosen for GENI wireless deployments because of the availability to campuses of this spectrum, plus: a) commercial WiMAX base station radio hardware operating in the 2.5GHz range and b) commercial WiMAX modems operating in the 2.5GHz range, for use in assembling programmable mobile stations.

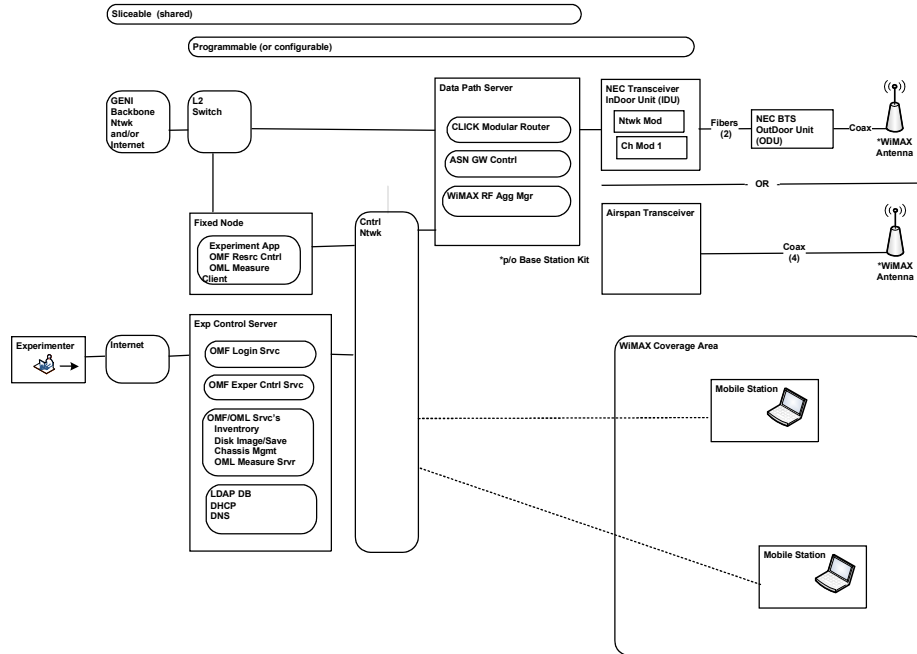
### **WiMAX Base Station Kit**

Each WiMAX base station kit includes commercial WiMAX radio hardware and interface software, plus servers for data-path and control software developed at Rutgers WINLAB [17].

See Fig. 5 for the configuration of the base station kits [18] used in the initial WiMAX deployments. They utilize NEC WiMAX radio hardware, including a rack-based InDoor Unit (IDU) connected with fiber to an OutDoor Unit (ODU), which is connected with a coaxial cable to a 120-degree sector antenna. They are tuned to operate at a frequency in the 2.5 GHz Educational Broadband Services band with a 10 MHz channel width. Urban testing has given a typical range of 1Km with peak downlink bit-rates up to 15 Mbps.

The base station kits in the extended WiMAX deployments will utilize Airspan WiMAX radio “Profile C” hardware, where all hardware is contained in one outdoor unit, also shown in Fig. 6. This radio hardware supports a 2 (transmit) x 4 (receive) MIMO arrangement, and the all-in-one outdoor unit is connected with 4 coaxial cables to a 120-degree sector antenna. Because of the MIMO arrangement, a longer range is expected than with the initial kits, both on campus and extending to nearby city areas.

## GENI: An Introduction



**Fig. 5.** Configuration of the WiMAX base station kit

At each site, the WiMAX data path server bridges the WiMAX radio hardware to a local Layer-2 switch, which is then connected to a GENI backbone network, e.g., Internet 2. The WiMAX data path server performs frame switching using only Layer 2 information and can provide isolation between simultaneous experiments [4].

At each site, one (or more) fixed nodes are also connected to the local Layer-2 switch, so that they can be used for making data path measurements, or for locally sourcing content. For example, a fixed node can run an iperf server, so that network performance measurements can be done with a mobile station running an iperf client.

Mobile stations within the WiMAX coverage area can connect to the WiMAX base station if they are configured into its access list.

By configuring the base station, a group of mobile stations can be designated using the control path so that their traffic is mapped by the data path software into designated VLAN tag, and connected by the local switch to a designated VLAN in the backbone network.

Multiple groups can be defined and mapped into multiple VLAN tags, thus virtualizing the base station [4]. Thus, each base station can simultaneously support multiple, diverse experiments, e.g., a mobility research experiment and a social networking experiment.

The default is for one group of mobile stations to be mapped into one VLAN tag, which is then mapped to the Internet; in this case, a DHCP service provides each mobile station with an IP address.

The WiMAX control server includes ORBIT Management Framework (OMF) software [19] to configure and monitor the base station for each experiment, and ORBIT Measurement Library (OML) software [20] to gather measurements from the base station and the mobile stations.

An experimenter gains access via the Internet to the WiMAX control server using an OMF Login service. Thus, remote access is provided to the base station as a WiMAX



aggregate so that it can simultaneously support experiments by local researchers and by GENI researchers from across the United States. A task is underway to provide WiMAX resource assignment using the GENI AM API starting in Spiral 5.

The extended meso-scale projects will grow the GENI WiMAX footprint, by extending 3 existing WiMAX sites to include multiple base stations, and by adding 3 new WiMAX sites, as shown in Fig. 1; new base station installations are planned to be completed by the end of Spiral 4.

A total of 5 sites will be equipped with multiple base stations, and 4 sites are being designed to support vehicular mobile experiments: Wisconsin, UCLA, Wayne State and Clemson. One site (Wisconsin) is fielding a multi-vendor WiMAX network, with mobile nodes carried on metro transit busses. Both Wayne State and Clemson plan to provide connectivity to a vehicle as it is moving along a highway; they will be working with their state Departments of Transportation on innovative WiMAX applications.

In the WiMAX sites where there are multiple base stations providing a “multi-cell” topology covering a geographic area, there will be a need to automatically “handover” WiMAX connectivity as a mobile station moves from one cell to the next. Thus, it is planned to include a mobility management or handover server at such sites as shown in Fig. 6. The mobility management server may follow commercial WiMAX practice, or research may drive servers that include new functions. A task is underway to define this server and deploy it starting in Spiral 5 or Spiral 6.

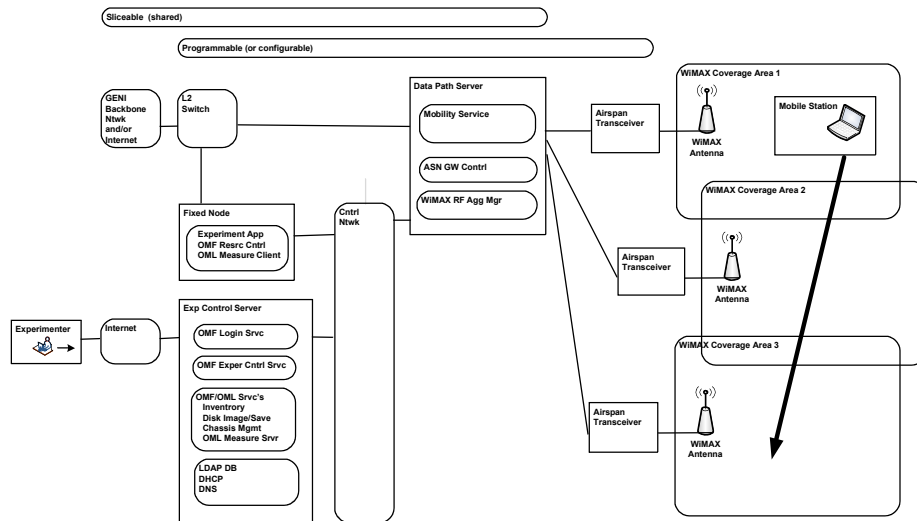


Fig. 6. A WiMAX site with multiple base stations and mobility management

### WiMAX Mobile Stations

Any WiMAX Mobile Station (MS) that can broadcast the GENI identifier can connect to a GENI WiMAX base station, and be used in a research or service experiment. In the simplest case, the base station provides only basic Internet connectivity, and the MS runs a client that utilizes the Internet connectivity.

A variety of MSs have been connected to a GENI WiMAX site, including: a) PCs with USB-connected WiMAX modem “dongles”; b) PCs with built-in WiMAX modem cards; and c) WiMAX-equipped handsets.

USB-connected WiMAX modem “dongles” [21] operate with both Windows and Linux PCs, but do not have open-source drivers. They are most useful for PCs running applications that utilize basic Internet connectivity.

A more capable reference MS has been assembled using a Linux netbook that is equipped with a built-in Intel 6250 WiMAX modem card [22]. Such a MS has been used for a variety of experiments. The Linux environment supports a wide range of experiments, including those utilizing OMF for control and OML for instrumentation and measurements.

An open-source driver is available for the Intel 6250 WiMAX modem card [23], which opens an avenue for research into the lower levels of the protocol.

As an option, the Intel 6250 WiMAX modem card can be mounted in a small box, connected to the Linux PC via a USB cable, and connected to an external antenna of any type. This provides a longer range than the internal antenna, and is particularly useful in vehicular experiments.

WiMAX-equipped handsets can be connected to a GENI WiMAX site if they are “unlocked”, but most available handsets are “locked” to Sprint/Clearwire and cannot be used. An effort is underway to identify an unlocked handset for use in the GENI environment.

## 4 Summary

The Global Environment for Network Innovations (GENI) is a suite of infrastructure being designed to support “at scale” research in networking and innovation in services. Efforts are underway to extend the GENI infrastructure within the next 3 years by: a) adding approximately 54 GENI racks of compute and network resources; adding L2 connectivity with OpenFlow features to regional and backbone networks; c) introducing tools and improving integration so that large-scale experiments can be setup and run with a reasonable effort; and d) adding WiMAX base stations to extend the reach of GENI.

The goals are to support large-scale GENI-wide research and service experiments and to permit campus, city and WiMAX-connected users from many locations to opt-in to the same service experiment.

**Acknowledgments.** This paper summarizes the efforts of the many talented and dedicated people who have contributed to the projects that have built the current GENI prototype over the past 3 years; many are identified in the References below. Within the GPO, Chip Elliott has provided the overall vision for GENI that is reflected here. A special thanks to Ray Raychaudhuri and Ivan Seskar at Rutgers WINLAB who have contributed so much to the WiMAX effort. Also thanks to those who have worked on the WiMAX experiments, including: Gautam Bhanage (Rutgers WINLAB); Manu Gosain (GPO); Thanasis Korakis (NYU Poly); Hamed Soroush (UMass Amherst); and Kiran Nagaraja (Rutgers WINLAB)

## References

1. Global Environment for Network Innovations (GENI), [www.geni.net](http://www.geni.net)
2. OpenFlow at Open Networking Foundation, <https://www.opennetworking.org/>
3. H. Mussman. A GENI WiMAX Infrastructure for Research Experiments, WiNTECH 2010.
4. G. Bhanage, I. Seskar and D. Raychaudhuri. A Service Oriented Experimentation Framework for Virtualized WiMAX Systems. TridentCom 2011.
5. D. Irwin, N. Sharma, P. Shenoy and M. Zink. Towards a Virtualized Sensing Environment. TridentCom 2011.
6. NSF award #0944089, for “GENI D&P Infrastructure”.
7. C. Elliot, GENI, Exploring Networks of the Future. TridentCom 2010.
8. NSF award #1128122, for “GENI Solicitation 3 D&P Proposal”
9. I. Baldine, J. Chase, C. Heermann and B. Viviano. ExoGENI Rack Architecture. GENI Engineering Conference 12.
10. GENI Aggregate manager API, <http://groups.geni.net/geni/wiki/GeniApi>
11. GENI PlanetLab Project, <http://groups.geni.net/geni/wiki/PlanetLab>
12. GENI ProtoGENI Project, <http://groups.geni.net/geni/wiki/ProtoGENI>
13. GENI ORCA Project, <http://groups.geni.net/geni/wiki/ORCABEN>
14. Introduction to GENI using Flack and the Instrumentation Portal, <http://groups.geni.net/geni/wiki/GEC12GENIusingFlack>
15. GENI Network Stitching, <http://groups.geni.net/geni/wiki/GeniNetworkStitching>
16. Tango GENI Project, <http://groups.geni.net/geni/wiki/TangoGENI>
17. GENI WiMAX Node Development Project, <http://groups.geni.net/geni/wiki/WiMAX>
18. GENI WiMAX Deployments, <http://groups.geni.net/geni/wiki/WiMAX>
19. T. Rakotoarivelo, M. Ott, G. Jourjon, and I. Seskar. OMF : A Control and Management Framework for Networking Testbeds. ACM SIGOPS Operating Systems Review, Vol. 43, pp. 54–59, 2010.
20. J. White, G. Jourjon, T. Rakotoarivelo, and M. Ott. Measurement architectures for network experiments with disconnected mobile nodes. In Proc. of TridentCom 2010, Berlin, May 2010.
21. WiMAX USB Clients, <http://groups.geni.net/geni/wiki/WiMAXInteg/ClientPlatforms#a7.2.1USBPlatforms>
22. Linux PC with Intel 6250 WiMAX Modem Card, <http://wimax.orbit-lab.org/wiki/WiMAX/30#a12.1LinuxPCwithIntelWiMAXModemCard>
23. Intel driver, <http://wimax.orbit-lab.org/wiki/WiMAX/30#a12.1.4DownloadingCompleteImage>
24. J. Huang, Q. Xu, B. Tiwana, Z. M. Mao, M. Zhang, and P. Bahl, “Anatomizing Application Performance Differences on Smartphones,” in Proceedings of Mobisys, 2010.
25. S. Teerapittayanon. Replacing WiMAX’s HARQ/ARQ with Network Coding: an Implementation, [http://groups.geni.net/geni/attachment/wiki/GEC12WiMaxDeploymentAndExperimentation/GEC12Slides\\_plusreadme.ppt](http://groups.geni.net/geni/attachment/wiki/GEC12WiMaxDeploymentAndExperimentation/GEC12Slides_plusreadme.ppt)
26. T. Korakis. ParkNet Demo <http://groups.geni.net/geni/attachment/wiki/WiMAXExper/110311%20%20ParkNet%20%20WednesdayMorning.pptx>
27. H. Soroush. WiMAX Throughput Experiment Using OMF/OML, <http://groups.geni.net/geni/wiki/OMFWiMAXExperiments#GENIWiMAXExperimentsUsingOMFOML>
28. K. Nagaraja. MobilityFirst Architecture and Protocol Evaluation on GENI, <http://groups.geni.net/geni/attachment/wiki/GEC12WiMaxDeploymentAndExperimentation/MF-GEC12-PlenaryDemo-10-3-2011.pdf>