The Great Plains Environment for Network Innovation (GpENI): A Programmable Testbed for Future Internet Architecture Research

James P.G. Sterbenz[1,5], Deep Medhi[2], Byrav Ramamurthy[3], Caterina Scoglio[4], David Hutchison[5], Bernhard Plattner[6], Tricha Anjali[7], Andrew Scott[5], Cort Buffington[8], Gregory E. Monaco[9,4], Don Gruenbacher[4], Rick McMullen[1], Justin P. Rohrer[1], John Sherrell[4], Pragatheeswaran Angu[3], Ramkumar Cherukuri[2], Haiyang Qian[2], and Nidhi Tare[4]

[1] The University of Kansas, Lawrence KS, USA
{jpgs,mcmullend,rohrej}@ittc.ku.edu
http://www.gpeni.net
[2] University of Missouri, Kansas City MO, USA
{dmedhi,rc5gd,hqian}@umkc.edu
[3] University of Nebraska, Lincoln NE, USA
{byrav,pangu}@cse.unl.edu
[4] Kansas State Univ., Manhattan KS, USA
{caterina,grue,ntare}@k-state.edu
[5] Lancaster University, UK
{dh,acs,jpgs}@comp.lancs.ac.uk
[6] ETH Zürich, Switzerland
{plattner,jpgs}@tik.ee.ethz.ch
[7] Illinois Institute of Technology, Chicago IL, USA
tricha@ece.iit.edu
[8] KanREN, Lawrence KS, USA
cort@kanren.net
[9] Great Plains Networks, Kansas City MO, USA
greg@greatplains.net

Abstract. The Great Plains Environment for Network Innovation – GpENI is an international programmable network testbed centered on a regional optical network in the Midwest US, providing flexible infrastructure across the entire protocol stack. The goal of GpENI is to build a collaborative research infrastructure enabling the community to conduct experiments in future Internet architecture. GpENI is funded in part by the US National Science Foundation GENI (Global Environments for Network Innovation) program and by the EU FIRE (Future Internet Research and Experimentation) Programme, and is affiliated with a project funded by the NSF FIND (Future Internet Design) Program.

Keywords: programmable Future Internet testbed, GENI, FIND, FIRE, PlanetLab, VINI, DCN, ResiliNets, PoMo, ResumeNet.

1   Introduction and Motivation

The Great Plains Environment for Network Innovation – GpENI (pronounced with accent on the middle syllable and rhyming with GENI) is an international programmable network testbed centered on a regional optical network between The University of Kansas (KU) in Lawrence, Kansas State University (KSU) in Manhattan, University of Nebraska – Lincoln (UNL), and University of Missouri – Kansas City (UMKC) within the Great Plains Network, supported with optical switches from Ciena interconnected by Qwest fiber infrastructure, in collaboration with the Kansas Research and Education Network (KanREN) and Missouri Research and Education Network (MOREnet). GpENI is undergoing significant expansion to Europe and Asia. The goals of GpENI are to:

- Build a collaborative research infrastructure in the Great Plains region among GPN and other institutions
- Construct an international programmable network infrastructure enabling GpENI member institutions to conduct experiments in Future Internet architecture, supporting projects such as PoMo: PostModern Internetwork Architecture [1] and ResumeNet [2]
- Provide programmable optical infrastructure to the GpENI Midwest optical backbone, and expand optical connectivity to selected international sites
- Provide flexible infrastructure to support the GENI program as part of the PlanetLab control framework cluster B
- Deploy tools developed by GpENI and the GENI community such as Gush for experiment control and Raven for code deployment
- Provide an open environment for networking research community experiments

1.1   Previous Testbeds

It is important to remember that the idea of large-scale testbeds on which to conduct networking research is not new with GENI and FIRE. This section summarises a few of the most relevant previous efforts.

Gigabit testbeds. A set of testbeds was constructed in the early 1990s to further the state of high-speed networking research, funded by the US NSF and DARPA (Defense Advanced Research Projects Agency), managed by CNRI (Corporation for National Research Initiatives). Five separate testbeds were constructed: Aurora, Blanca, Casa, Nectar, and Vistanet [3], later supplemented by MAGIC [4].[1] The Gigabit Testbeds were a platform for research in high-speed networking, new bandwidth-enabled applications [5], and networked supercomputing (before the term *grid* was coined).

Active network testbeds. In the late 1990s, testbeds were constructed in the US and Europe to support active networks research. *Active networks* are programmable networks in which one of the programming modalities includes

---

[1] GpENI is currently using the research fiber originally installed for MAGIC at the University of Kansas.

capsules of mobile code that can dynamically program network nodes. In the US, the ABone [6] was constructed as part of the DARPA-funded Active Networks Program[2][7], to permit experimentation on programmable network languages, management and control [8], node operating systems [9], and security mechanisms [10]. The ABone had the goal of open access to the research community, with a separation into core nodes provided by program fundees, and edge nodes attached by the larger community. In Europe, the EU FP5 FAIN (Future Active IP Networks) [11] and related projects (e.g. LARA++ [12]) also investigated active and programmable networks, with testbeds constructed for experimentation. These active network architectures and testbeds permitted sharing of infrastructure by the simultaneous execution of active applications (AAs) in execution environments (EEs) on a network node operating system (NodeOS). While not generally recognised in this manner, the active network testbeds had many goals similar to that of GENI, and should be considered a conceptual precursor.[3]

Modern large-scale testbeds. More recently, two large scale testbed infrastructures have been constructed with the explicit goal of permitting open access for networking research. PlanetLab [13] is a large-scale worldwide infrastructure that permits users to run networked experiments at large scale. The infrastructure is shared using the *slice* paradigm. It is important to note that while PlanetLab permits experimentation in networked applications and end-to-end protocols, the network itself is not programmable, and experiments in lower-layer protocols can only be performed on overlays. VINI [14] is a version of PlanetLab that allows control over the network topology.

Emulab [15] is a network testbed consisting of a cluster of computing nodes interconnected by flexible network infrastructure, which permits researchers to experiment with network protocols and applications with complete root access to the systems. A number of Emulab facilities are located throughout the world, some of which provide access to external researchers in addition to the main facility at the University of Utah. Both PlanetLab and Emulab are the basis for GENI control frameworks; GpENI uses the PlanetLab GENI control framework.

1.2   Current Future Internet Initiatives

While a number of researchers were proposing alternatives to the Internet architecture as early as the 1980s (including research programs such as DARPA Next Generation Internet – NGI), there is now a general consensus in the research community that the current architecture is limiting in scale and support for emerging application paradigms such as mobile and nomadic computing and communications. Recent research initiatives include NSF FIND (Future Internet Design) [16] in the US, EU FP6 SAC (Situated and Autonomic Communications) [17], and the research component of FP7 FIRE (Future Internet Research

---

[2] Program managers were Gary Minden, Hilarie Orman, and Doug Maughan.
[3] Three of the GpENI PIs (Sterbenz, Plattner, Hutchison) performed research in these programs.

GpENI Layer		Programmability
	experiment	Gush, Raven
7	application	PlanetLab
4	end-to-end	
3	router	Quagga, XORP, Click
	topology	VINI
2	VLAN	DCN
	lightpath	
1	photonics	site-specific

Fig. 1. GpENI Programmability Layers

and Experimentation) [18]. These research initiatives aim to investigate clean slate (greenfield) as well as incremental (brownfield) architectures to evolve the Future Global Internet architecture.

A key problem remains how to experiment with Future Internet architectures at reasonable scale. For this reason, the NSF GENI (Global Environments for Network Innovation) program [20], experimental component of the EU FP7 FIRE programme [18] , and Japanese JGN2plus [19] testbed plan to deploy large-scale programmable testbeds for experimentation of the Future Internet research. GpENI is a large-scale international programmable testbed currently under construction as part of the GENI, FIRE, and FIND programs.

1.3   GpENI Programmability and Flexibility

The defining characteristic of GpENI is programmability of *all* layers, as shown in Figure 1, implemented on a *node cluster* of general- and special-purpose processors, described in detail in Section 3. At the top layer Gush provides experiment control and Raven distributes code; both are software developed as part of the GENI program. Layer 7 and 4 programmability are provided by the GENIwrapper version of PlanetLab. At layer 3, programmable routers are implemented in Quagga, XORP, and Click, supplemented by any other technology for which GpENI institutions should choose to deploy.[4] Flexible network-layer topologies are provided by VINI. At layer 2, dynamic VLAN configurations are provided by DCN-enabled managed Gigabit-Ethernet switches at the center of each GpENI node cluster. GpENI institutions directly connected to the optical backbone use DCN-enabled Ciena switches to provide dynamic lightpath and wavelength configuration. At layer 1, the architecture even permits programmability at the photonic layer for switches that provide such support. Furthermore, each GpENI institution can connect site specific networking testbeds; plans include wireless, sensor, and cognitive radio testbeds (e.g. KUAR [22]). External users in the broader research community may request GpENI accounts with which to run network experiments.

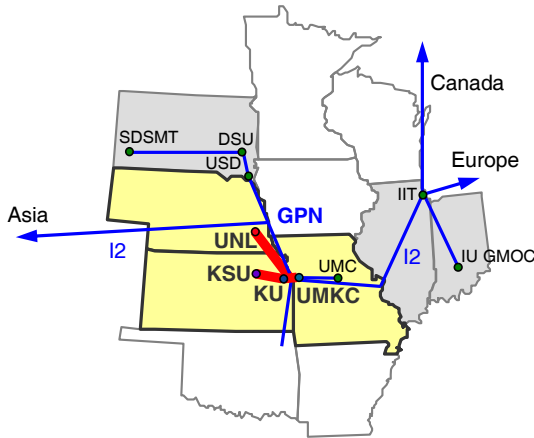[4] We are investigating including OpenFlow [21] as a core GpENI technology.

Fig. 2. GpENI Midwest Topology

2   Topology and Network Infrastructure

The core of GpENI is the regional optical backbone centered around Kansas City. This is extended by KanREN (Kansas Research and Education Network) to various GPN (Great Plains Network) institutions in the Midwest US. Connectivity in Kansas City to Internet2 provides tunneling access to the Canadian, European, and Asian GpENI infrastructure. Optical connectivity is currently in place in the UK between Lancaster and Cambridge, and will replace other GpENI L2TPv3 and IP tunnels as available. GpENI is growing, currently with about 200 nodes at 40 institutions in 20 nations. Institutions may connect to GpENI if they are interested in becoming part of the GpENI community, and have the resources to install, connect, and manage a node cluster.

2.1   Midwest US GpENI Core and Optical Backbone

GpENI is built around the core GpENI optical backbone centered in the Midwest, shown in Figure 2, among the principal institutions of KU, KSU, UMKC, and UNL, currently under extension to additional institutions, including the GMOC (GENI Meta-Operations Center). The optical backbone consist of a fiber optic run from KSU to KU to the Internet2 PoP in Kansas City, interconnected with dedicated wavelengths to UMKC and UNL, as shown in Figure 5.

Each of the four core institutions will have a node cluster that includes optical switching capabilities provided by a Ciena CoreDirector or CN4200, permitting flexible spectrum, wavelength, and lightpath configurations.[5]

2.2   GpENI European Topology

GpENI is extended to Europe across Internet2 to GÉANT2 and NORDUnet and then to regional or national networks, as shown in Figure 3. Currently,

---

[5] GpENI optical infrastructure is currently undergoing phased deployment; the UNL switch is installed and the KU switch is under procurement.
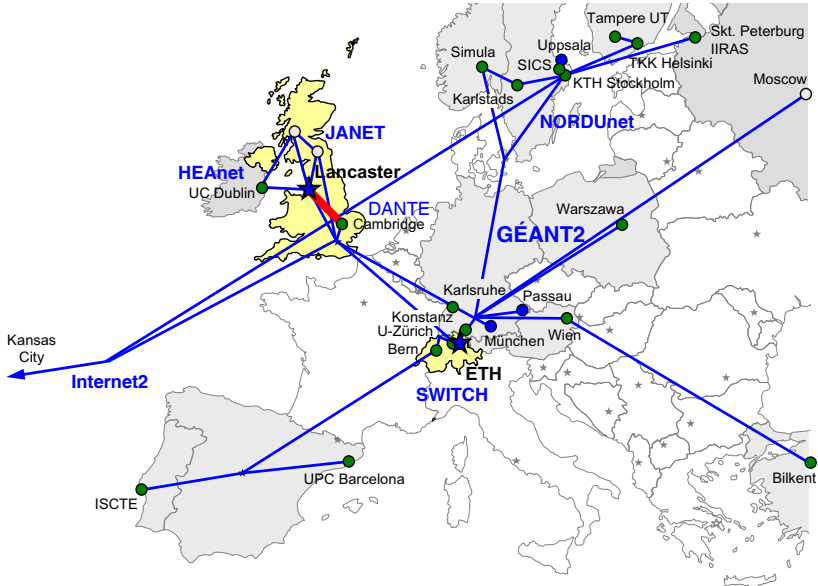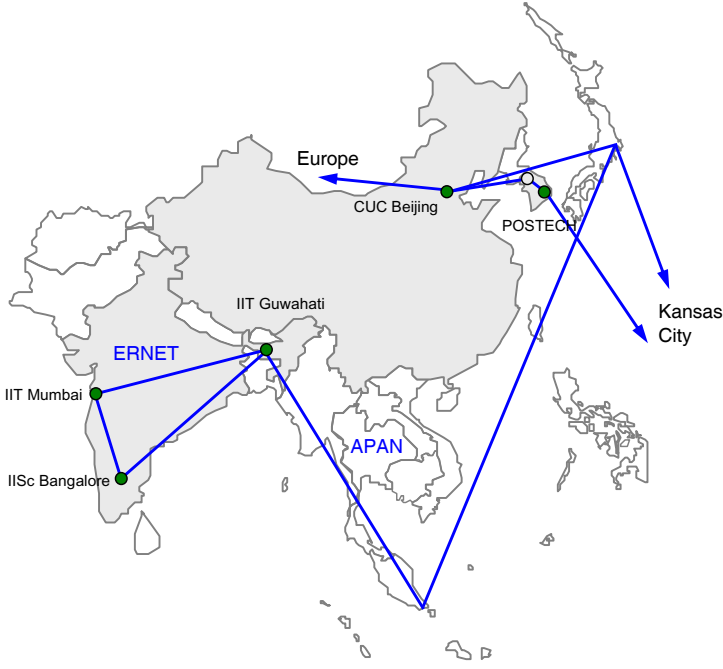
Fig. 3. GpENI European Topology

Fig. 4. GpENI Asian Topology

connectivity is achieved using L2TPv3 and IP tunnels. A direct fiber link over JANET is deployed between Lancaster and Cambridge Universities.[6] The principal European GpENI institutions are Lancaster University in the UK and ETH Zürich in Switzerland.

2.3   GpENI Asian Topology

Similarly, GpENI is extended to Asia across Internet2 to APAN, then to national research network infrastructure including ERNET. as shown in Figure 4.
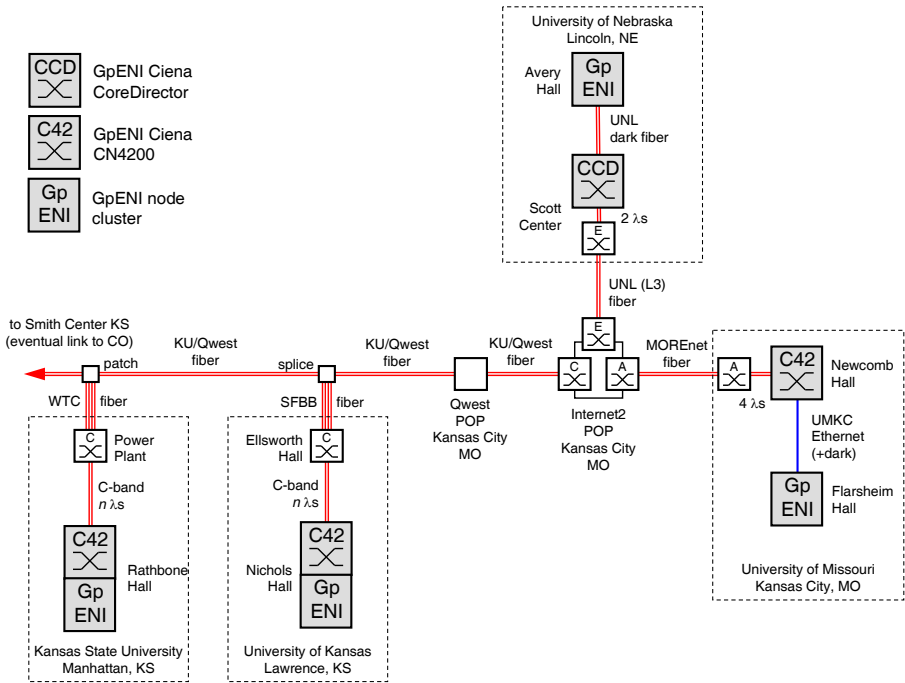

Fig. 5. GpENI Midwest Optical Backbone

3   Node Cluster Architecture

Each GpENI node cluster consists of several components, physically interconnected by a managed Netgear Gigabit-Ethernet switch to allow arbitrary and flexible experiments. GpENI uses a KanREN 198.248.240.0/21 IP address block within the gpeni.net domain; management access to the facility is via dual-homing of the Node Management and Experiment Control Processor. The node

---

[6] GpENI is currently working with research network providers to increase direct optical connectivity.

cluster is designed to be as flexible as possible at every layer of the protocol stack, and consists of the following components, as shown in Figure 6:

- GpENI management and control processor: general-purpose Linux machine
- PlanetLab control framework consisting of aggregate managers: MyPLC with GENIwrapper SFA (at KSU), myVINI (at UMKC), and DCN (at UNL)
- PlanetLab programmable nodes (enabling layer 4 and 7 experimentation)
- VINI-based programmable routers (providing flexible network topologies), with Quagga and other extensions such as XORP and Click (enabling layer 3 experimentation), as well as the ability for GpENI partners to install their own programmable routers
- Site-specific experimental nodes and testbeds, including software defined radios (e.g. KUAR), optical communication laboratories, and sensor testbeds
- Managed Gigabit Ethernet switch, providing L2 VLAN programmability and connectivity to the rest of GENI
- Ciena optical switch running DCN providing L1 interconnection among GpENI node clusters on the Midwest US optical backbone
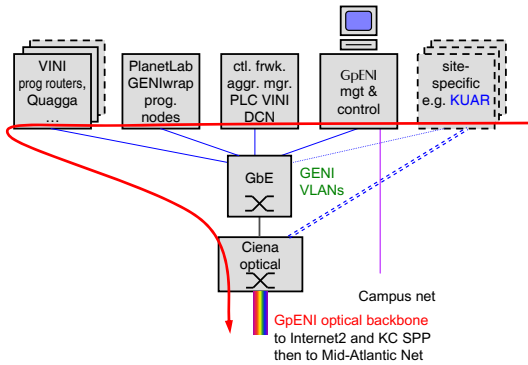
Fig. 6. GpENI Node Cluster

3.1   GpENI Management and Control

The GpENI management and control services are distributed across the Linux machines dedicated for the purpose at each of the node clusters. Open-source tools are used wherever possible to minimise the amount of GpENI-specific software development and maintenance required. Some of these services are installed at every node cluster, for example the Cacti monitoring tool [23] is used to monitor the per-port network usage on each of the Netgear Gigabit-Ethernet switches. Nagios [24] is used to monitor the status of individual nodes and services across all the clusters. Zenoss Core [25] is also being evaluated as an alternative to Nagios. On the other hand, some functions are specific to one site, such as the GpENI-Planetlab demo [27], running on Apache and hosted only on the management node at the KSU node cluster.

The control node for each cluster also provides firewall and NAT services using Firestarter [28] for that cluster's private subnet, thereby protecting insecure devices, such as the Netgear switch telnet and SNMP management interfaces, from direct exposure to the Internet.

3.2   Experiment Control

To ease experimenters role in resource discovery and preparation of experiments, the GENI User Shell (Gush) [29], provides a robust experiment control and management framework. Gush extends the PlanetLab control framework to implement an API that interacts with the GENI Clearinghouse. With Gush deployment on the GpENI aggregate, a user can reserve both PlanetLab and GpENI resources by downloading agents on the selected nodes and deploy their experiments subsequently with the help of the controller that communicates with the agent about node status. The Raven [30] provisioning service provides services such as the proper execution environment, software packages, configuration information and computational resources.

3.3   PlanetLab Control Framework Sub-aggregate

The PlanetLab Control Framework provides the control software to implement control plane, data plane, management plane, and operations plane functionalities in GENI Cluster B.[7] Cluster B contains a number of distinct *aggregates*: PlanetLab nodes, VINI nodes, Supercharged PlanetLab Platform (SPP) backbone nodes [31], OpenFlow switches [21], and GpENI node clusters. The GpENI aggregate consists of three sub-aggregates: the PlanetLab control framework (described in this section), the routing and topology sub-aggregate, and the DCN sub-aggregate (discussed in the following sections). A *slice* is a fraction of resources allocated to a particular experiment throughout an aggregate. GpENI has deployed the PlanetLab GENIwrapper slice facility architecture (SFA) [32] to export three interfaces: registry, slice manager (SM), and aggregate manager (AM). These GENI interfaces are accessible via the slice facility interface (SFI) implementing functions to get slice details, node details, and user accounts. Using the SFA, GpENI has *federated* with the public PlanetLab, which means that GpENI resources are available for authorised PlanetLab researchers using the public PlanetLab interface. In the long term, there are plans for federations among GENI, as well as to other Future Internet testbeds, such as OneLab [33].

3.4   Routing and Topology Sub-aggregate

For layer 3 programmability, GpENI provides programmable topologies using VINI and an arbitrary number of programmable routers in each node cluster.

Virtual Topology Control using VINI. VINI enables arbitrary virtual topology creation on top of physical networking infrastructure. VINI is

---

[7] GENI is divided into several clusters, each of which uses a distinct control framework.

essentially a flavor of PlanetLab with a set of enhancements to the PlanetLab kernel and tools called *Trellis*. Trellis allows users to create their own topology, either automatically using the IIAS (Internet in a slice) toolkit or manually designating links between *slivers*. MyVINI is the private VINI implementation within the GpENI nodes, permitting full control of virtual topology.

Programmable Routers. Trellis can host existing routing software such as Quagga [34], XORP [35], and Click [36]. Therefore, MyVINI with programmable routing software installed enables the implementations of layer 3 programmable routers for GpENI nodes. The routing softwares support a wide range of existing routing protocols. For example, Quagga supports RIPv1, RIPv2, OSPFv2, BGP-4, RIPng, OSPFv3. However, these programmable routers have very limited processing power and can only handle moderate size forwarding tables compared with realistic routers in backbone networks since they are running in commodity PCs. GpENI is initially running Quagga, with planned modifications to run XORP and Click.

The ultimate goal is to design a toolkit with a GUI permitting users to plug-in arbitrary research protocols. This will permit new routing protocols to be tested on the arbitrary layer-3 topologies provided by VINI, and on arbitrary layer-2 topologies provided by the DCN sub-aggregate described in the next section. Two areas of emphasis will be to study network resilience (in the context of the EU-FIRE ResumeNet project) and routing among heterogeneous realms (in the context of the NSF-FIND Postmodern Internet Architecture project).

3.5   DCN Sub-aggregate

GpENI uses DCN for control of VLAN interconnection among L2TPv3 tunneled node clusters as well as optical switches connected directly to the core backbone.

Internet2 ION Service and DCN Optical Control Plane Framework. In recent years, the Internet2 network has evolved from a pure IP-based packet-switching network into an advanced hybrid optical and packet network. Currently, the Internet2 network consists of three different, but related physical networks: Advanced IP network (provided by Juniper routers), Virtual Circuit network (provided by the multiservice switching capabilities of the Ciena CoreDirectors), and the core optical network (provided by the Infinera platform). Apart from the traditional IP service, the new Internet2 network offers a virtual circuit service to provision dedicated bandwidth across the network, called the Internet2 Interoperable On-demand Network (ION) [37]. The ION service can be used to set up on-demand, dedicated optical paths between endpoints in standard SONET bandwidth increments up to 10 Gbps. The ION service is dynamic and can thus be used to set up short term connections by a requestor or an application through a Web interface. The control plane software that automates the set up and tear down of the circuits was developed for the Internet2 Dynamic Circuit Network (DCN) research prototype and leverages technology developed by DRAGON (USC/ISI East, MAX, and George Mason University), GÉANT2,

and the DOE ESnet (OSCARS project). The DCN software suite is used to control the Internet2 ION service and is also available as open source to other institutions to use in creating their own DCN.

DCN in GpENI. GpENI has modified the DCN software suite to support the Netgear GSM7224 Ethernet Switch [38] used at the center of each GpENI node cluster. This will enable the creation of on-demand circuits at the required bandwidth for specified durations using the DCN software suite. There are two options to deploy DCN across the GpENI testbed [39]. The first option involves setting up static VLANs and Q-in-Q at the GPN switch in the Kansas City Internet2 PoP. The second option involves setting up a dedicated switch for DCN in the GpENI testbed. Deploying DCN across GpENI will also facilitate setting up VLAN circuits across the Ciena CoreDirectors located at various locations in Internet2. The CoreDirector Component Manager Interface [40] describes the use of the CoreDirector in the GpENI testbed. As additional GpENI optical switches are deployed, a common GpENI-wide DCN testbed will emerge over a multidomain network with CoreDirectors forming the optical domain and Netgear switches forming the Ethernet VLAN domain at each GpENI institution.

4   Early Demonstrations and Experiments

GpENI is currently operational at the core Midwest US nodes, principal European nodes, and several other institutions.[8] We have just begun experiments with GpENI, and are permitting limited accounts to researchers in the GENI, FIRE, and larger community.

The first public demonstration[9] creates a GpENI PlanetLab slice that generates an overlay loop among GpENI nodes; a screenshot example is shown in Figure 7. The next public demo will use the routing sub-aggregate to create a VINI topology slice giving control to the underlay in addition to the PlanetLab overlay.

Very preliminary performance experiments have begun using the current overlay capabilities, measuring throughput and delay for various topologies among the KU, KSU, UMKC, UNL, Bern, and Cambridge node clusters.

Table 1 shows the average delay between a set of 5 GpENI node clusters, measured on 23 Nov. 2009, beginning at 09:00. For each cluster pair, a `ping` command was issued 6 times at 2 hour intervals. Each command used 64KB packets, with a repeat count of 1000 and an interval of 1 sec. Thus the averages for each pair are computed for 6000 values over an 11 hour period. This table clearly indicates that the PlanetLab overlay topology is not representative of the layer-2 underlay, since traffic between Cambridge and Bern is backhauled to Kansas in the current tunnel confuguration. With the further deployment of L2TPv3 tunnels within Europe and DCN control of the L2 topology, we will be able to control both logical and physical topologies.

---

[8] As of this writing, KU, KSU, UMKC, UNL, Lancaster, Cambridge, ETH Zürich, and Bern nodes are operational, with many others under installation.

[9] Available at `http://control-1.ksu.gpeni.net/demo/#sites`.
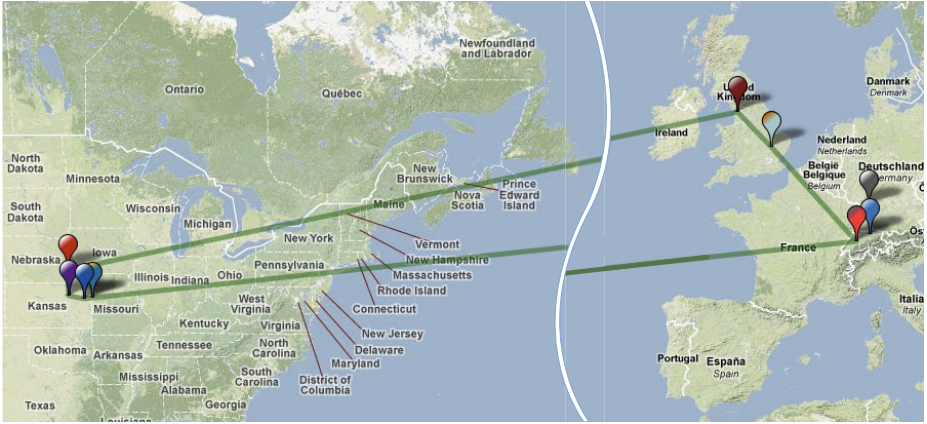
Fig. 7. GpENI-Planetlab Overlay Path Demonstration

Table 1. Average ping times in msec. between GpENI sites

Location	KSU	KU	UNL	Bern	Camb.
KSU	–	2	9	145	155
KU	2	–	6	143	153
UNL	9	6	–	152	158
Bern	145	143	152	–	295
Camb.	155	153	158	295	–

5   Future Outlook

GpENI is a testbed programable at every layer, intended to enable research in Future Internet architectures. The first GpENI nodes are operational, and we believe that it is the first such international testbed spanning the US GENI and EU FIRE programs, with plans to connect to JGN2plus in Japan, as well as Korean and Chinese testbeds. GpENI is currently under expansion to over 200 nodes at 40 institutions in 20 nations. We will use GpENI for experiments in the NSF-FIND Postmodern Internet Project at the University of Kansas, and the EU-FIRE ResumeNet project among KU, Lancaster University, ETH Zürich, TU München, Universität Passau, and Uppsala Universitet. GpENI has direct interconnection to ProtoGENI [41] linking GENI clusters B and C. Furthermore, efforts are underway to connect the GpENI to the MAX (Mid-Atlantic Crossroads) regional optical network at Layer 2 through the ProtoGENI network. The MANFRED project is in GENI cluster B with GpENI. In the longer term, we are under discussion with regional and national research networks to extend the optical substrate domestically and internationally.

Acknowledgements

References

1. Bhattacharjee, B., Calvert, K., Griffioen, J., Spring, N., Sterbenz, J.P.G.: Postmodern Internetwork Architecture. ITTC-FY2006-TR-45030-01 (February 2006), `http://wiki.ittc.ku.edu/pomo`
2. ResumeNet, `http://www.resumenet.eu/project/index`
3. CNRI. The Gigabit Testbed Initiative (1996), `http://www.cnri.reston.va.us/gigafr/Gigabit_Final_Rpt.pdf`
4. MAGIC Gigabit Testbed, `http://www.magic.net/`
5. Partridge, C., Davie, B., Campbell, R., Catlett, C., Clark, D., Feldmeier, D., McFarland, R., Messina, P., Richer, I., Smith, J., Sterbenz, J.P.G., Turner, J., Tennenhouse, D., Touch, J.: Report of the ARPA/NSF Workshop on Research in Gigabit Networking. National Science Foundation, Washington (July 1994)
6. Braden, B., Ricciulli, L.: A plan for a Scalable ABone – A Modest Proposal. TR, USC/ISI (January 1999)
7. Active Networks Program, `http://www.darpa.mil/sto/strategic/an.html`
8. Jackson, A.W., Sterbenz, J.P.G., Condell, M.N., Hain, R.R.: Active Network Monitoring and Control: The SENCOMM Architecture and Implementation. In: Proc. DARPA DANCE, pp. 379–393. IEEE, Los Alamitos (2002)
9. Shalaby, N., Peterson, L., Bavier, A., Gottlieb, Y., Karlin, S., Nakao, A., Qie, X., Spalink, T., Wawrzoniak, M.: Extensible Routers for Active Networks. In: Proc. DARPA DANCE, pp. 92–116. IEEE, Los Alamitos (2002)
10. Murphy, S., et al.: Security Architecture for Active Nets. In: DARPA Active Networks Security Working Group (July 1998)
11. Galis, A., Plattner, B., Smith, J.M., Denazis, S.G., Moeller, E., Guo, H., Klein, C., Serrat, J., Laarhuis, J., Karetsos, G.T., Todd, C.: A flexible IP active networks architecture. In: Yasuda, H. (ed.) IWAN 2000. LNCS, vol. 1942, pp. 1–15. Springer, Heidelberg (2000)

12. Chart, T., Schmid, S., Sifalakis, M., Scott, A.: Active Routers in Action: Evaluation of the LARA++, Active Router Architecture in a Real-Life Network. In: Wakamiya, N., Solarski, M., Sterbenz, J.P.G. (eds.) IWAN 2003. LNCS, vol. 2982, pp. 215–227. Springer, Heidelberg (2004)
13. PlanetLab, http://www.planet-lab.org/
14. VINI: A Virtual Network Infrastructure, http://www.vini-veritas.net/
15. Emulab: Network Emulation Testbed, http://www.emulab.net/
16. NSF NeTS FIND Initiative, http://www.nets-find.net
17. FP6 Situated Autonomic Communications,
    http://cordis.europa.eu/fp7/ict/fire/future-internet-projects_en.html
18. FIRE: Future Internet Research & Experimentation,
    http://cordis.europa.eu/fp7/ict/fire/
19. JGN2plus Testbed, http://www.jgn.nict.go.jp
20. GENI: Global Environments for Network Innovations, http://www.geni.net/
21. The OpenFlow Switch Consortium, http://www.openflowswitch.org/
22. Minden, G.J., et al.: KUAR: A flexible software-defined radio development platform. In: Proc. DySPAN, pp. 428–439 (April 2007)
23. Cacti, http://www.cacti.net/
24. Nagios, http://www.nagios.org/
25. Zenoss, http://www.zenoss.com/
26. GpENI wiki, http://wwww.gpeni.net
27. GpENI Demonstration, http://control-1.ksu.gpeni.net/demo/
28. Firestarter, http://www.fs-security.com/
29. Gush: Geni User Shell, http://gush.cs.williams.edu/trac/gush
30. Raven Provisioning Service, http://raven.cs.arizona.edu/
31. Supercharged PlanetLab Platform (SPP) Hardware Components,
    http://wiki.arl.wustl.edu/index.php/SPP_Hardware_Components
32. Peterson, L., Sevinc, S., Lepreau, J., Ricci, R., Wroclawski, J., Faber, T., Schwab, S.: Slice-based facility architecture. Technical report, Princeton (November 2007)
33. OneLab: Future internet test beds, http://www.onelab.eu/
34. Quagga routing suite, http://www.quagga.net/
35. XORP: Extensible Open-Source Routing Platform, http://www.xorp.org/
36. The Click Modular Router Project, http://read.cs.ucla.edu/click/
37. Internet2 DCN/ION Software Suite,
    https://wiki.internet2.edu/confluence/display/DCNSS/
    DRAGON+Supported+Switches
38. Lake, A., Tracy, C.: Dragon Supported Switches,
    https://wiki.internet2.edu/confluence/display/DCNSS/
    DRAGON+Supported+Switches
39. Angu, P., Ramamurthy, B.: DCN/ION in GpENI Investigation Report (October 2009), http://www.gpeni.net
40. Angu, P., Ramamurthy, B.: Ciena CoreDirector Component Manager Interface Document (October 2009), http://www.gpeni.net
41. ProtoGENI, http://www.protogeni.net
42. Mid-Atlantic Crossroads, http://wiki.maxgigapop.net