The University of Kansas

Technical Report

GpENI: Great Plains Environment for Network Innovation (Proposal)

James PG Sterbenz, Deep Medhi, Greg Monaco,
Byrav Ramamurthy, Caterina Scoglio,
Baek-Young Choi, Joseph B Evans,
Don Gruenbacher, Ronqing Hui, Wesley Kaplow,
Gary Minden, and Jeff Verrant

ITTC-FY2009-TR-0061349-01

October 2008

Project Sponsor:
National Science Foundation

GpENI: Great Plains Environment for Network Innovation

James P.G. Sterbenz*,
Deep Medhi†, Greg Monaco‡, Byrav Ramamurthy¶, Caterina Scoglio§,
Baek-Young Choi†, Joseph B Evans*, Don Gruenbacher§, Ronqing Hui*,
Wesley Kaplow‖, Gary Minden*, and Jeff Verrant°

ABSTRACT

This technical report is an edited version of the funded proposal to the National Science Foundation: *GpENI: Great Plains Environment for Network Innovation*, which is a collaboration among four universities (The University of Kansas, Kansas State University, University of Missouri – Kansas City, University of Nebraska – Lincoln), three research networks (Great Plains Network, Kansas Research and Education Network, Missouri Research and Education Network), and industry (Ciena, Qwest). The core proposed activities are included as written and numbered in the original proposal (with financial and schedule information redacted) as a service to the GENI community and to provide the GpENI project a vehicle for very early dissemination of information. ***It is important to note that the scope of the project has and will continue to change.*** Current project information and references to publication are available on the GpENI wiki: `www.gpeni.net`.

Proposed Activities

1. Vision and Scope of Work

The vision of the Great Plains Environment for Network Innovation (GpENI) is to provide a flexible, scalable, programmable, and sliceable network infrastructure at all layers from the physical through the network to the application, in both the data and control planes, and to support its management toward the greater goal of converging to GENI. We propose to leverage optical network research infrastructure that has recently been deployed between Kansas, Missouri, and Nebraska and made available to GpENI, thus distributing GpENI across four universities in a regional multi-wavelength deployment. Each node will be connected by a Ciena optical switch, and consist of a cluster of computing nodes running open-source software, including PlanetLab, XORP (eXtensible Open Router Platform), Click, and IDC (inter-Domain Controller). This permits rapid deployment of GpENI nodes with maximum flexibility at all layers without significant software development. In the first six months of the project, we will open GpENI to the research community to allow experimental research and further development of GENI. Over the rest of the project we will develop and deploy automated management and control interfaces and software to better integrate the programmable components, and to

*The University of Kansas, †University of Missouri – Kansas City,
¶University of Nebraska – Lincoln, ‡Great Plains Network, §Kansas State University, ‖Qwest, Ciena°

increase their virtualization, sliceability, and isolation of resources to permit sharing in time and space. We will make all of our results and software publicly available, and provide real-time visualization to the operation of the infrastructure and experiments. We will engage in active outreach to increase the scale of GpENI in number of nodes and reach, and to engage and interoperate with other funded GENI initiatives. We will be available as a testbed for other GENI projects that do not have their own distributed infrastructure. The primary goal of GpENI to support the goals of the GENI spiral development process and convince the nation that GENI is a viable strategy that should be funded by the US Congress.

1.1 Project Goals, Milestones, and Deliverables: The project goal is to rapidly deploy a regional network infrastructure based on a wavelength-switched optical backbone, that provides programmability and sliceability at all layers in the protocol stack, and is made openly available as a resource to the GENI community. The main milestones are:

1. Deploy GpENI nodes at four Midwestern universities connected by GPN optical network
2. Ability for research community to run simple GENI experiments based on PlanetLab
3. Integrate / interface GpENI programmable components (PlanetLab, XORP, Click, IDC)
4. Automate experimental access and implement virtualization and sliceability
5. Support the GENI spiral development process by incorporating ideas from the community and interconnecting with other GENI funded prototypes and trial integrations

GpENI will produce five deliverables: physical infrastructure, software infrastructure, GENI development platform, federation with other GENIs, wiki and forum, and project reports.

1.2 Prototyping and Development Activities: The primary goal of this project is to deploy infrastructure quickly by leveraging as much existing open software as possible. Thus much of the development activity will consist of gluing interfaces together, for example using the PlanetLab management and slicing software to interface to the programmable router and optical switch parts of GpENI. There will be some development of management software to manage and automate access to GpENI, and we will also develop Web-based tools to log, document, and visualize the operations of GpENI; this will be publicly linked from the GpENI wiki for the GENI community.

1.3 Breakdown of activities and tasks by institution: While GpENI is a highly collaborative cross-sector (academia/vendor/operator) project, each institution will take the lead in a particular part of the deployment, development, and operation of GpENI:

- The University of Kansas will manage the overall project, and take the lead on the GpENI management and control infrastructure.
- Kansas State University will take the lead on using PlanetLab, will administer MyPLC on their node and will explore how to extend MyPLC management and control to other layers in GpENI – particularly to the network topology and programmable router (layer 3) and the lightpath topology and Ciena optical switches (layers 2 and 1).
- The University of Missouri – Kansas City will take the lead on integrating the flexible router platforms, and on how to virtualize the platform and provide sliced access to XORP and Click configurations.
- The University of Nebraska-Lincoln will take the lead on the control and management of the Ciena optical switches, including integration of IDC (Inter-Domain Controller).
- Great Plains Network will coordinate the operation of the various network entities, including MOREnet (Missouri Research and Education Network), KanREN (The Kansas Research and Education Network), Internet2, and the service providers including Qwest, and coordinate outreach and interoperation with other GENI projects including the logistics of interconnecting over Internet2.
- Ciena will provide technical support and insight on the deployment and continued operation of the optical switches, and access to their control and management.
- Qwest will provide support for the measurement and analysis of their part of the fiber-optic network infrastructure, and continuing access to additional fiber resources.

2. GENI Relevance

2.1 Evolvability and Scalability: The primary thrust of the proposed GpENI infrastructure is evolvability and scalability. Evolvability is accomplished by constructing a cluster of programmable computing platforms that can easily adapt to the needs of the GENI community in both the short and long terms. For example, while we will initially deploy XORP over Click on the programmable router platforms, other flexible router software and programmable network interfaces can be loaded on one or more of these Linux computers. We are even agnostic to operating system; if it makes sense to insert, for example, a NetBSD box, this may easily be done. In the longer term, the Native GENI Spiral Target platforms are intended to run various (perhaps competing) stages in the evolution to a GENI platform, at least from the software and control perspective. GpENI is scalable by design, both in number of nodes (initially four) and in the configuration of each node (initially five computing platforms in addition to the optical switch). Additional platforms can and will be added to the base configuration, as shown by the dotted boxes in Figure 2. These can either be more platforms of a given type (e.g. PlanetLab, programmable router, or native GENI nodes), or specialized nodes, such as to experiment with wireless and optical capabilities (as is already planned with the KU Agile Radios example shown).

2.2 Outreach: While an important goal of this proposal is to provide flexible infrastructure to support the GENI spiral development process, we recognize that a critical aspect of this that the GENI community be driving this development. We actively outreach to the GENI and FIND communities at three levels.

2.2.1 Interconnection with Other Funded Prototypes and Testbeds: We will seek to interconnect and engage with other funded GENI prototypes and testbeds whenever possible. Between direct access via Internet2 as well as a more general external Internet interface, we will seek cross-pollinate ideas, and run joint experiments within the GENI community. We will provide accounts to those in the GENI community who wish to use GpENI as an experimental and development platform.

2.2.2 Growth of GpENI within Midwest: The initial configuration of GpENI is limited to four nodes due to the realities of available funding and physical topology of existing fiber connectivity within Kansas, Missouri, and Nebraska. We will seek to expand GpENI in the Midwest by pursuing additional funding and encouraging others to do so to increase the reach of the physical infrastructure (the primary expense) and adding additional nodes. In particular, we will pursue EPSCoR funding in Kansas and Nebraska for this purpose, and work with universities in other states to increase the reach of GpENI. We expect that the existing GPN (Great Plains Network) consortium and our relationship with Internet2 will help facilitate this goal, and the high visibility of GENI will make other institutions eager to join.

2.2.3 Open Access to Research Community: We intend to provide open access to the NSF (and FIND) research community in particular, and the broader networking research community in general by providing accounts and slices to run experiments on GpENI. We hope that this will fuel the development of GENI, and get the broader research community exited and ultimately help provide the NSF with ammunition to support the funding of the longer term GENI program. On the local level we will leverage the two NSF-funded FIND projects at the University of Kansas and run PoMo (PostModern Internet Architecture in conjunction with the Universities of Kentucky and Maryland) and CogNet (Cognitive Radio Protocols in conjunction with Rutgers WINLAB and CMU) on GpENI. Furthermore, we will bring in European partners that collaborate with the University of Kansas, in particular Lancaster University in the UK and ETH Züruch in Switzerland.

2.3 Risk Reduction Rationale: GpENI reduces risk to the GENI program by the rapid deployment and demonstration of a regional "mini-GENI" that will be able to demonstrate the GENI ideal. The underpinning of this strategy is a flexible node architecture based initially on existing open-source infrastructure and software, including PlanetLab, XORP, Click, and IDC. We anticipate demonstrating the first GENI experiments based on PlanetLab within the first six months of funding. While the inital management and control will be rather manual (except for existing PlanetLab MyPLC management); this is calculated to prevent dependencies on software development. As the project progresses, the PIs and graduate students will develop management and control interfaces and software to automate GpENI access and operation, borrowing these capabilities from PlanetLab when practical.

Risk is further reduced by a GpENI node architecture that consists of a scalable architecture of commodity PCs running Linux (by default) interconnected by a Gigabit Ethernet switch. This allows the arbitrary configuration and expansion of GpENI components, and allows each GpENI site to attach unique resources for experimentation (such as wireless nodes and radios, optical equipment). The architecture is designed to be fundamentally scalable and flexible.

4. Technical Approach

4.1 Physical Infrastructure and Topology: The physical topology and fiber interconnection is shown in Figure 1. New proposed GpENI infrastructure is depicted by grey blocks; infrastructure that is currently deployed or scheduled for imminent deployment and will be made available for use by GpENI is shown by white blocks. Each of the four university nodes interfaces into the GpENI backbone via a Ciena CN4200 or CoreDirector switch. The rest of the GpENI node infrastructure for each site is labeled "GpENI node cluster" (this is expanded in the Figure 2).

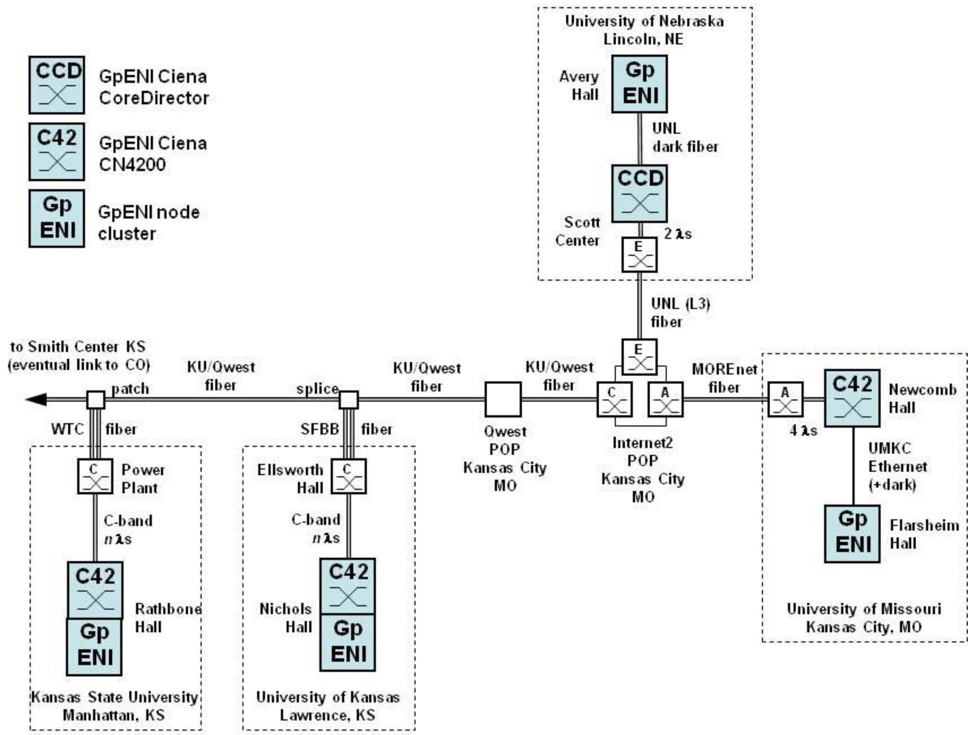

Figure 1. GpENI Physical Topology

The main fiber run between KSU, KU, and Kansas City is Qwest Fiber IRUed (leased) to KU, proceeding through the Qwest POP, and continuing to the Internet2 (Level3) POP, which will provide access to GpENI from Internet2. A chunk of C-band spectrum is available providing multiple wavelengths at Ku and KSU. UMKC is connected over MOREnet (Missouri Research and Education Network) fiber to the Internet2 POP, with four wavelengths anticipated. UNL is also connected to the Intenet2 POP over fiber IRUed from Level3 with two wavelengths committed. Local fiber in Manhattan and Lawrence is leased from Wamego Telephone (WTC) and Sunflower Broadband (SFBB), respectively. There is abundant dark fiber already in place on the KU, KSU, UMKC, and UNL campuses to connect the GpENI nodes to the switches (existing or under deployment) on the GPN fiber backbone. For reference, the UNL link is terminated by Ekinops, the UMKC link will be terminated by ADVA, and the KU/KSU link will be terminated by Ciena switches.

4.2 Node Architecture: Each GpENI node consists of several components, as shown in Figure 2. The components are physically interconnected by a Gigabit Ethernet switch to allow arbitrary and flexible experiments. GpENI will use a private IP address space; access to the facility will be via dual-homing of the Node Management and Experiment Control Processor. The node is designed to be as flexible as possible at every layer of the protocol stack. The arrow overlaid on the figure shows a conceptual flow of an experiment in which the GENI experiment controls the configuration of the PlanetLab, which in turn configures a custom routing protocol, which in turn configures the optical switch configuration.

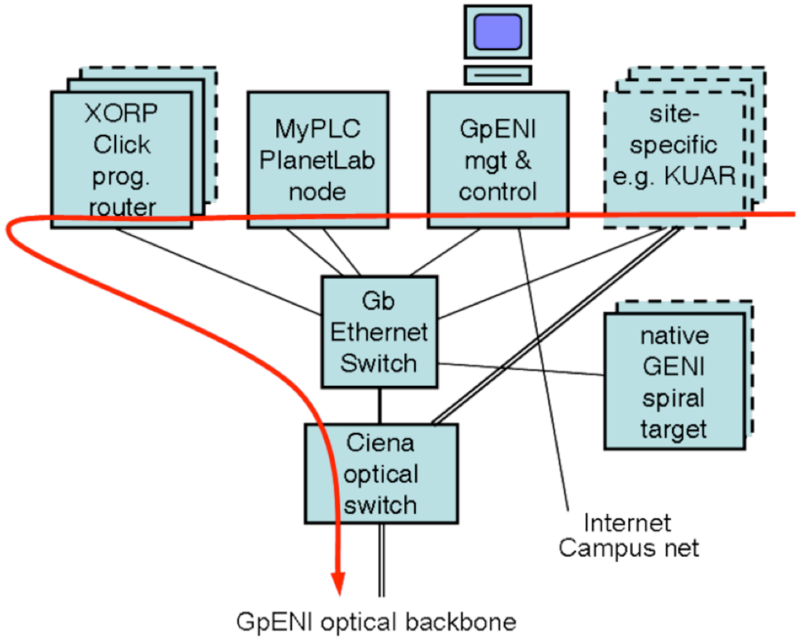

Figure 2. GpENI Node Cluster

4.2.1 Ciena Switch: The proposed GpENI network testbed will facilitate experimentation at all protocol layers over a programmable network architecture. In particular, the network will deliver any mix of connectivity to include fiber, spectrum, wavelengths, GbE or SONET circuits across the four university campuses interconnected through the Kansas City Internet2 POP. This will provide unprecedented capability to create and embed multiple virtual/logical topologies (ring, mesh, linear configurations), switch or add/drop circuits at any of the granularities mentioned above and assign bandwidth for specific functions (such as control channel on a particular circuit). All of these functionalities are enabled by the bleeding-edge carrier-grade network platforms provided by Ciena Corporation. In particular, critical components of GpENI are Ciena's CN4200 advanced services platform (multi-service optical and data network switching, aggregation, and transport system) and CoreDirector® CI Multi-Service Switch (an O/E/O switch that provides 160 Gbps of switching capacity through its OC-3, OC-48, OC-192 optical interfaces, STM-1e electrical interfaces and Gigabit & 10Gb Ethernet interfaces). To provision connections across this regional network testbed, we plan to use the Inter-Domain Controller (IDC) API interface being developed by the GLIF control plane working group, which is already being used in managing the DRAGON, Internet2 and ESNet SDN network infrastructures.

4.2.2 Programmable Router: The programmable router platform will enable GENI experimentation with new network-layer mechanism, including control and routing algorithms. By default we will provide a XORP over Click environment, but permit the loading of arbitrary programmable router software. Each GpENI node will have at least two Linux-based computing platforms, to allow immediate slicing (albeit primitive at first).

XORP is a open-source extensible, modular software that can be loaded to a PC to turn into a functional router. XORP also has a command line interface to configure routing protocols and interfaces. It supports different routing protocols to run in different sandboxes, making it ideal to serve as an experimental research platform. Currently, XORP does not implement its own forwarding engine, depending instead on the underlying operating system. However, it can use Click modular router that has sixteen elements on its forwarding path. We plan to integrate XORP-over-Click to provide a programmable router platform for experimental work for the community. Our goal is to explore how we can make the platform sliceable, especially in regard to integrating it for controlling experiments at all levels. In this regard, our development effort will proceed in close interaction with MyPLC, and control and management at the optical level.

4.2.3 PlanetLab Node: PlanetLab is an overlay testbed designed to allow researchers to experiment with network applications and services. The PlanetLab open-source software (MyPLC) consists of the core system that runs at the PlanetLab Central manager, and a client system running on each of the individual nodes. Each GpENI node will include one PlanetLab sub-node running the client software of MyPLC, which will allow flexible resource allocation at both the application layer and the transport layer. The PlanetLab manager of MyPLC will will reside at KSU and perform resource allocations in each GpENI node involved in a given experiment. The MyPLC manager will be connected to all the GpENI PlanetLab nodes taking part to a given experiment through the dedicated fiber network or through the Internet. The initial configuration will consist of one PlanetLab box per GpENI node, but this is easily expandable if needed. The local GpENI management and control node will request resources at the layers involved in a given experiment by contacting the respective central managers and, upon acknowledgement, will implement the information flow in the respective sub-node set, by specification of proper socket sequences.

4.2.4 Node Management and Experiment Control Processor: The Node Management and Experiment Control Processor is dual homed. The interface to the private GpENI address space is used to load, monitor, and control GENI experiments. The other interface connects to the public Internet (through routers and firewalls based on the policy of each university) and allows access by the research community to load and run experiments. Local campus access for the GpENI PIs, students, and staff will be through this interface, via the public Internet or Internet2. These processors will be used to control and manage the other components in each cluster. The processor at KU will serve as the overall management and control interface to GpENI. After initial deployment, monitoring, control, management, and visualization software will be developed for these nodes and made available to the research community.

4.2 5 GENI Spiral Development Platform: These Linux-based processors allow the development and test of GENI platforms. As design and development of the GENI platform proceeds in the community, these processors are available for the testing and analysis of new GENI software. As the GENI platform begins to gel, successive releases will be loaded on this machine so that the GpENI nodes evolve into the target GENI platform. It will be possible to do this in parallel with experiments on the other components in the GpENI cluster. Initially one processor will be provided at each site, but this can be expanded as necessary.

4.2.6 Local Facilities: Each GpENI node will interconnect local facilities to the node to conduct site-specific experiments requiring additional hardware. For example, KU will connect to its optical lab as well as to KUARs (KU agile radios) to bring in wireless networking experiments. At UNL, projects such as the US Compact Muon Solenoid (CMS – for which UNL is one of seven Tier 2 sites in the U.S.) will benefit from GpENI. At UMKC, UNL, and KSU, wireless sensor networks will be interconnected.