# Review of the GENI Research Plan
**January 15-16, 2007**

**Report dated January 28, 2007**

The Computing Research Association GENI Community Advisory Board (precursor to the Computing Community Consortium GENI Science Council, which is in the midst of formation) organized a review of the GENI Research Plan in Seattle on January 15-16 2007. Participants in the review are listed in Attachment A.

Prior to the review, participants read the *GENI Research Plan* (GDD-06-28, Version 4 (4.3) of January 2, 2007), a *Commentary on the GENI Research Plan* written by David Clark, and the GENI *Facility Usage Scenarios* (GDD-06-41). The review itself combined presentations and discussions; the agenda is listed in Attachment B.

This report summarizes our findings and recommendations. At the highest level:

- We became persuaded during the course of the review that there is a research case that justifies GENI.

- However, this case was not laid out sufficiently convincingly (to us) in the GENI Research Plan as written – additional work is needed.

- Also, because of the nature of GENI, it will be easy for the construction effort to go off the rails. Great sensitivity will be required in planning, selecting appropriate partners, and establishing appropriate processes.

**Findings**

The exponential growth of the Internet is a remarkable story, and the Internet has achieved extraordinary impact throughout society. Despite this already great impact, we believe the Internet (and the applications that run on it) will continue to increase in importance as a key component of the technical, commercial, governmental, and societal infrastructure of the United States and the world.

The extent to which the Internet has successfully accommodated this exponential growth – growth that goes far beyond mere numbers of hosts – is a testament to the capabilities of the community that created it and has evolved it. Today, though, the strains are showing. To meet tomorrow's needs (and uses that have ever-greater, society-wide importance), the Internet must be better engineered to provide improvements in areas such as availability, manageability, robustness, etc. The Internet must also support new functions and applications – for example, several orders of magnitude more hosts to accommodate ubiquitous computing, sensor networks, and new communication services. And the Internet must incorporate an increasing range of new technologies.

There are many ideas in the research community on how to address issues such as these. Validating and refining these ideas, however, requires experimentation on a programmable network of realistically large scale and, often, with a large number of real people spread over a wide geography using it. In some cases, the construction of demonstration distributed systems and applications would be required in order to achieve full validation.

Thus, we believe that a facility such as GENI, which supports large-scale experimental research in networking in ways that the production Internet and small-scale testbeds simply cannot, is likely to be a necessary mechanism for furthering essential innovation in networking and the applications that run on top of it.

**The GENI Research Plan**

A great deal of effort by highly talented individuals has gone into the current GENI Research Plan, and it contains a broad collection of very interesting research directions. Nonetheless, we believe that the plan can be and should be made more convincing.

*Depth versus breadth*

The current Plan has been written with the goal of being as inclusive as possible of the enormously broad set of research topics that will benefit from GENI. While it is of course valuable to show many communities of researchers how their work will fit into and be supported by GENI, we believe that the "mile wide, inch deep" approach to making the case for GENI needs to be augmented. We feel that a fuller consideration of a small set of compelling examples is also required, even if this comes at a cost of less comprehensiveness in terms of breadth. Without a "deep dive" into some number of specific research problems, the document was not sufficiently convincing as to whether the GENI facility is truly necessary to any of the included research, versus other approaches to supporting such research.

During the course of our review, we had substantive discussions of several such research topics: for example, rethinking networking from the point of view of extreme manageability, and rethinking how networking can assist in reducing the collateral damage caused by end host security problems. These discussions, which went well beyond the material in the Plan, led us to conclude that GENI construction is likely to be merited (*highly* likely, in the view of most members of the committee), although of course final judgment should be reserved until the case has been made explicit and peer reviewed.

In revising the Plan, perhaps a half dozen examples must be described in detail sufficient to convince an objective expert in the domain, have a potential outcome worthy of large scale funding, and require the facility.

Essentially, what we are proposing is that the scientific rationale in the Research Plan have two components:

1. It should illustrate the breadth of the potential uses of GENI, motivating its wide applicability for a range of scientific/engineering experimentation and helping to clarify the requirements

that it should meet. Due to its breadth, this aspect will be somewhat shallow and won't be convincing to the critical reader.

2. It should go into significant depth on perhaps a half dozen projects that carefully document their research objectives, their potential impact, the reasons that they require experimentation, and an explicit discussion on how and why GENI will provide the appropriate infrastructure for them.

This is the most important recommendation that we have to offer concerning the GENI Research Plan.

*Defining the scope of GENI*

The Research Plan would be stronger if more attention was devoted to the scope of the research to be supported by the GENI facility. Scoping is good because it focuses resources and attention, thereby increasing the chances of succeeding; on the other hand, scoping risks alienating members of the community whose work falls outside the scope.

One possibility would be to define the scope of GENI by defining the Internet (for GENI purposes) as the processors, communication channels, algorithms and protocols that control the processing within the network (between computing endpoints) and the software at the endpoints that connect to it. This definition includes in scope, for example, network security, network management, distributed services, in-network data manipulation, and network node operating systems. It omits from scope, for example, application user interfaces, host security, host operating systems, sensor node design, and purely physical-layer work in wireless and optics. (Notably, research in these latter categories can proceed with experimentation in the small, with individual end nodes.)

*Attracting users*

It is now well recognized that the traffic workloads, network dynamics and social challenges posed by real users are part and parcel of how networked systems succeed or fail. Many of the experiments that researchers aim to conduct with GENI are only valid under *in vivo* scenarios. How can one convincingly talk about scalability of a service without demonstrating use at scale? How can one demonstrate security without exposing oneself to both attackers and naïve users? How can one claim the benefits of increased robustness without validating that users ultimately experience it as such? Experience tells us that things happen in real deployments that the designers failed to anticipate.

While the current plan does touch on the need for attracting real users, we believe that this is a first class issue that is critical to GENI's success and demands additional attention. It is not reasonable to adopt an "if we build it they will come" strategy. Instead, a plan for how to attract and retain a range of real users needs to be addressed. This may require partnership with ISPs, enterprises and/or universities in addition to the development of proxy mechanisms for end nodes.

*Generality: achieving it, or managing the risks of sacrificing it*

GENI is envisioned as a multi-purpose experimental facility – one that can support a wide variety of research, some of it not fully conceived at this time. Indeed, ensuring generality in GENI's components is a significant bulwark against design and execution risk. GENI's programmability at the switch and node levels – its general purpose and FPGA computing platforms – allows a wide range of processing to be performed on data as it passes through the network, and imposes few limitations on experimentation.

By contrast, the portions of GENI that focus on particular physical layer embodiments tend to be, by nature, tied to a particular technology and are therefore riskier. In effect, we are making a "bet" that this particular technology will be the important one, and if that bet proves wrong, we will be unable to adapt to experiment with a different technology without reengineering that new technology into the system.

It is our suggestion that additional effort be devoted to considering if there are not alternatives to increase generality in these technologies. For example, in the wireless aspects of GENI, perhaps software radio technology and software controllable antennas should be reconsidered as a way to maximize research flexibility in an area that is undergoing considerable evolution. In general, in those cases where it is decided that a single technology should be used in GENI, there should be explicit justification about why the risk is warranted.

*Documenting the need for large-scale experimentation*

The review committee readily accepts the need for large-scale experimentation (in addition to, for example, simulation, analysis, and small-scale experimentation). However, we feel that the case should be strengthened for some groups of readers, by including several short examples from previous research. We identify three scenarios:

1. Something unexpected or truly surprising emerges from a deployment. We found the content distribution example highlighted in a "sidebar" in the current Research Plan compelling: a perfectly plausible approach that a research group had in mind was invalidated after running the system atop PlanetLab. Another similar example could strengthen the Plan, particularly if it concretely showed a situation where simulation or emulation was inadequate.

2. Research that cannot be conducted for want of a GENI-like facility. Networked systems often exhibit important behaviors only at large scale, which requires some combination of a sufficient number of real users, a large and diverse mix of traffic, and a large enough number of components. It would help to add an example showing that scale matters and that there are important, innovative, and interesting projects currently being explored in the community that have hit a "brick wall" for want of a large enough infrastructure to test the ideas – "these researchers have tried every option that's available and can't move forward unless we have something as big and diverse as what's envisioned in GENI."

3. Even some relatively small-scale research can be very difficult to conduct because the necessary experimental platforms are painful to construct. A shared hardware or software

infrastructure can speed research advances by fostering synergies where researchers build on one another's systems, by encouraging collaboration, by increasing the community's ability to conduct repeatable experiments, and by lowering the barrier to entry for certain kinds of research. For example, the shared PlanetLab infrastructure has allowed researchers to conduct more ambitious research than they could otherwise accomplish by making use of the tools and services developed by other research projects on PlanetLab. The wireless component of GENI is a prime example of such an opportunity: today it is quite difficult to deploy and operate even a 100- to 1000-node wireless facility, and the number of groups who have access to such networks is tiny (but a large number of research groups write papers using simulation where the results are suspect, and often disappoint when tried out in practice).

Essentially, what we're suggesting is a few more, and more diverse, "sidebars" such as the content distribution example currently on page 11. Perhaps the material currently in Section 4 could be folded in here; it seems disruptive where currently located.

*Complementary approaches*

A major motivation for building the GENI facility is the difficulty of running experiments in today's commercial Internet. For example, a router vendor may not be willing to let a researcher modify the functionality of its routers to evaluate a new idea, and a service provider may be unwilling to enable this feature in its operational network.

That said, there may be creative means by which NSF could engage industry – for example, by supporting cooperative research with the private sector. NSF could facilitate a researcher's working with a service provider to help that researcher evaluate a specific research idea, such as quantifying the performance impact of running routers with much less packet-buffer space on its existing network.

This approach would significantly restrict the range of experiments that could be conducted, and would limit the ability of the researchers to iterate through successive experiments and build on each other's work. Still, it may be a cost-effective way to enable a certain class of experiments.

*Technology transition paths*

We spent considerable time at the review clarifying what GENI *is* and what GENI *isn't*. This needs to be front-and-center in the Research Plan. GENI is not "the next Internet." Rather, GENI is a highly flexible instrument that enables research necessary to design the next Internet – including supporting the systems and applications and users necessary to fully evaluate that research.

There are three possible technology transition paths, and these should be discussed at least briefly:

1. Some discoveries made possible by GENI could possibly be retrofit into the current Internet, yielding significant improvements.

2. GENI could provide some capabilities that are so revolutionary and compelling that GENI grows into a parallel network used alongside the current Internet (which is perhaps improved in various ways along the lines of (1)).

3. Some new application domain adopts some of the ideas demonstrated in GENI, and then becomes wildly popular, essentially becoming an application "pulling" (2) into existence.

**Construction of the GENI facility**

Our panel gathered to review the GENI Research Plan. However, because several members of the panel have deep experience with large commercial and research software efforts, the process for constructing the GENI instrument generated considerable discussion, which we summarize here. These are *not* issues to be addressed in the Research Plan; rather, they are issues that should be front-and-center in the minds of those shepherding the GENI specification and construction effort.

Building GENI will be different from the usual process for building a large scientific instrument, aerospace project, or large software system, for several reasons:

1. GENI will involve an enormous amount of *custom* software development. A great deal of planning and skill will be required.

2. That being said, it will not be possible to lay down a firm set of requirements that have clear acceptance criteria and can't change significantly during the multi-year construction. For instance, it's important to have programmability in various places, but the exact meaning of "programmability" may well change. Fixing it in advance – without understanding the tradeoffs that emerge as the design proceeds, technology evolves, and users' understanding of their needs deepens – will lead to a bad outcome.

3. The technology out of which GENI is built changes rapidly. Processing nodes, FPGAs, optical data transmission, wireless modulation and MAC schemes, and protocols for distributed computing will improve rapidly during the several years it will take to build the instrument. Constructing the instrument and designing experiments to be done with it should proceed hand in hand as much as possible, to ensure both that the most cost-effective components are procured and that the resulting facility accommodates upgrades gracefully.

4. There is more uncertainty about how GENI will be used than is typical for a large scientific instrument. The objects of study are themselves constructed by humans rather than found in nature, and often by the experimenters themselves. It typically takes only a few months or years to build an object of study. Furthermore, what is interesting to build depends on the capabilities of GENI itself as well as on external factors. As a result, there needs to be much more interaction between experimenters and builders throughout the construction than is typical for large projects.

These characteristics do *not* create an impossible situation. Efforts of this sort *can* be successful – cf. PlanetLab, Apache, and many others. But these characteristics *do* create a project that does not lend itself to a development process that has a rigid set of testable requirements. A much better process is the agile one used by successful businesses such as Amazon or FedEx to develop large IT systems: frequent releases of incrementally improved systems, and rapid feedback from users (for GENI, the scientists) to the requirements for the next release. It is essential that:

- GENI be built by an organization that is familiar with this kind of development approach, and can engage effectively and continuously with the user community.

- GENI use a contracting and acceptance process that is friendly to this approach.

Another important consideration is that there are many opportunities to reduce both the cost and the risk of GENI development by partnering with industry. Examples of partners that have extensive experience with various aspects of the GENI instrument are:

- Cisco, Juniper, etc., for networking hardware and in-network software.
- Google, Yahoo, Amazon, Microsoft, etc., for large-scale, easily managed processing nodes.
- Comcast, Verizon, etc., as ISPs.
- IBM, HP, and other companies that operate business telecommunication networks..

The organization building GENI, and the procurement process, should be friendly to this kind of partnering, and the builder should not have major conflicts that would make it in its interest to discourage it, or that would inherently discourage potential partners.

**Summary**

While this document makes many suggestions, even tough ones, we emphasize again that the review committee had the clear sense that GENI would be a necessary facility to many, important kinds of networking and related distributed systems research. We believe the GENI constituency can take our comments into account and produce plans that are convincing and that will yield a clearly valuable, and exciting, project.

**Attachment A: Participants in the review**

**GENI Planning Group and/or GENI Research Coordination Working Group**

Dan Blumenthal
Dean Casey
David Clark
Aaron Falk
Ramesh Govindan
John Heidemann
Nick McKeown
Larry Peterson
Mike Reiter
John Wroclawski

**Reviewers**

Hari Balakrishnan
Gaetano Borriello
Mike Dahlin
Butler Lampson
Ed Lazowska
Peter Lee
Stefan Savage
Jonathan Smith
Alfred Spector
George Varghese

**Dual role**

Tom Anderson
Jennifer Rexford
Scott Shenker
Ellen Zegura

**NSF**

Suzi Iacono
Paul Morton

**Attachment B: Agenda for the review**

**Monday January 15**

**12:00 Setup; box lunches available**

**2:00 Overview**
Ed Lazowska: Welcome / CCC Background (15)
Larry Peterson: GENI Background (15)
David Clark: Research Plan Background (15)
All: Discussion (15)

**3:00 Research Plan: Sections 1 & 2 (The need for transformative research; Is it time to rethink the Internet?)**
David Clark: Presentation (30)
All: Discussion (30)

**4:00 Break**

**4:15 Research Plan: Section 3 (An agenda for research using GENI)**
David Clark: Introduction (10)
John Heidemann: Sensornet research opportunities (20)
Nick McKeown: Security research opportunities (20)
Jennifer Rexford: Network management research opportunities (20)
All: Discussion (20)

**5:45 Research Plan: Section 5 (Requirements for GENI)**
Larry Peterson: Presentation (30)
All: Discussion (30)

**7:00 Dinner / Discussion**

**Tuesday January 16th**

**8:00 Breakfast**

**8:30 Reviewers caucus to set agenda for morning**

**9:30 GENI team available as determined by reviewers**

**12:00 Lunch**

**1:00 Further discussion and report preparation report**
GENI team available as determined by reviewers

**5:00 Deadline for adjournment**