

**Keynote Address to
EuroView 2007
Würzburg, Germany
July 23, 2007**

“GENI at Age Three: Origins, Objectives, Outlook”

**Peter A. Freeman
Emeritus Dean and Professor
Georgia Institute of Technology**

Note: The views expressed herein are the personal views of the author and are not necessarily those of any organization with which he is or has been associated.

Good morning.

Thank you, Mr. Chairman.

It is a pleasure to be here this morning and to talk with a group of people who are as concerned about the future of networking, as am I.

SLIDE 1

My congratulations to the organizers of this workshop for focusing on future networks, and for inviting a number of us from outside Europe. Truly international interaction of this type is absolutely essential to move our collective scientific understanding forward; I expect to learn a lot from being here and trust that you will learn something useful from me as well.

Specifically, I want to share with you my perspective on the GENI Project being developed under the sponsorship of the U.S. National Science Foundation. My role in GENI is that I was the executive at NSF that started the project and guided it through its first two and one half years, and now as a private citizen, am trying to assist its development in several ways.

SLIDE 2

My brief remarks this morning will cover the origins, objectives, and outlook of the Project, and are intended to form a basis for your penetrating questions. I would note that beyond these short verbal remarks and visuals, I am preparing a more thorough history of the Project to date for later publication. The ever-expanding body of information about the plans for the GENI Facility can be found in several places that I have referenced. The GENI Research program is just getting underway, but I am confident that coherent compendiums of research results will be written in due time.

Let me also give credit to a number of people for many of the visuals. GENI is a community project and that is reflected in the origin of the visuals.

Let me begin by reminding you of two points I believe to be a common understanding by all who have looked deeply at the subject of networking:

- The Internet has changed everything (in a manner of speaking) in a remarkably short period of time;
- When we envision the future, it is increasingly clear that we may well not be able to get *there* from *here* by continuing to innovate new networking technology in the same manner as the past.

SLIDES 3-6 (illustrating the two points above)

Similarly, there is a set of technical/operational characteristics of a future network that most of us can agree on, perhaps with different weightings of importance depending on your viewpoint.

SLIDE 7

Namely, a future global network, at a minimum, should:

- Be worthy of society's trust
- Provide a bridge between physical and virtual worlds
- Support pervasive computing
- Enable further innovations in research and commerce
- Create a world in which we would want to live.

Yet, many of the challenges that concern most people in our societies are beyond purely technical issues. If these broader challenges are not dealt with soon, "remedies" may be imposed by society that will limit the rich and expanding digital future that many of us can envision. While these broader challenges cannot be completely dealt with by technical means alone, it is our responsibility as scientists and engineers to make sure that the best technical basis exists for meeting these challenges without unduly limiting future possibilities. These include issues that are:

- Social (e.g. children's use)
- Political (e.g. posting false information)
- Policy (e.g. access)
- Legal (e.g. copyright)

These fundamental issues are closely intertwined and must be addressed if we are to realize the opportunities before us.

SLIDE 8

Understanding these points, some of us have come to the conclusion that entirely new networking mechanisms may well be needed and that those mechanisms can not be

realized only by continued incremental changes to the existing fundamental networking structures on which the Internet is built. As with any significant judgment of the future, of course, this conclusion is open to question.

SLIDE 9

As you may know, NSF does not directly perform research but has the responsibility to provide financial support to a large fraction of the U.S. fundamental computing research community. To discharge that responsibility, NSF routinely and continuously engages the broad research community, as well as practitioners where appropriate, to ascertain where to invest its research funding. Although by law NSF support must be limited to those working through U.S.-based organizations, NSF tries to engage the international scientific community as much as possible; indeed, I'm sure several of you have participated in reviews or workshops for NSF or talked with various representatives of NSF informally.

Through these processes over the past seven or eight years, several of us at NSF came to the conclusion around 2002 that careful, empirically based research is needed to inform the design of future networks and set out to start to do something about it. As we looked more deeply and talked with more networking experts of all sorts, an interesting set of ideas emerged from several well-established network researchers in early 2004 (Professors Larry Peterson, Scott Shenker, Jon Turner, and Tom Anderson), pulled together by one of the NSF program directors at that time, Dr. Guru Parulkar. When presented to my deputy, Dr. Deborah Crawford, and me we immediately understood the potential importance of the ideas in the broadest possible context. Bearing in mind that NSF does not perform research, make policy, or build production networks, we set about to put into motion a program of research to *enable* the design of future networks.

Part of the broad responsibility of NSF is to help provide the infrastructure needed for research. Indeed, we had already instituted in early 2003 the Experimental Infrastructure Networks (EIN) funding program which supported several smaller scale networking testbeds including PlanetLab, ORBIT, Emulab, DETER, and so on (some of which pre-existed the EIN program and were supported in other ways). These testbeds played a very important role in shaping the concept of the GENI Facility and continue to do so as the GENI design matures. The original ideas that led to GENI developed partly in that context; the idea of a dedicated, large experimental testbed on which researchers could prototype and demonstrate their ideas was part of the thinking brought forward to us.

Our response was to establish the GENI Program.

SLIDE 10

The GENI Project, then, is about research and the infrastructure needed to support that research:

SLIDE 11

NSF has supported networking and systems research for many years prior to the advent of GENI and that support will undoubtedly continue indefinitely. That provides the general context for GENI. The next visual shows the general timeline of GENI activities so far – a formal planning activity that started in mid-2004 that resulted in a preliminary design for a Facility by the end of 2005. Since that time a variety of activities to further detail and improve the design have been underway; it must be frozen before final funding is approved. Current estimates are that construction will not begin for several years. On the diagram I have shown this indeterminacy with a large question mark.

SLIDE 12

This brief description of the origins of the GENI Project has implicitly articulated the objectives, but let me make them explicit.

SLIDE 13

First and foremost, the objective is to encourage and support research that will produce and explore, in a systematic and scientific manner, innovative ideas for future networks that will meet the criteria for future networks that I mentioned earlier. Those are, of course, very high-level criteria so that there are a plethora of detailed questions that need to be answered – many more than have been thought through at this time. While there are efforts underway to articulate some of the major research questions to be answered, for our purposes here I want to focus on the *type* of research that is being encouraged in the GENI Program.

A key aspect of the research that needs to be conducted is that it will permit exploration of new constructs under realistic loads – something that is for the most part impossible on the current Internet. Given that the Internet has rapidly become a critical infrastructure for so much activity throughout the world, it can never be experimented with in any way that would possibly compromise it. Thus, from the start it was understood by all that a separate, large-scale test-bed would be needed.

A second aspect of the needed research is that a good bit of it should be quantitatively based. That is, it needs to permit detailed measurement that will permit deep evaluations and characterizations of new networking mechanisms; those measurements must be reproducible in good scientific fashion. Again, the commercial Internet will not permit sufficient instrumentation to support such investigations in general.

A third aspect is the strong need for better theoretical constructs that can be used to explain observed network behavior and/or make predictions about future operations and performance. Developing and validating such constructs requires a close collaboration between designers (of new constructs), experimenters, and theoreticians.

Beyond research topics that are directly contributory to future network design there are a number of related topics that can be encouraged and that can make use of the GENI Facility, at least implicitly. The theoretical constructs just mentioned may be developed to explain digital network behavior, but they may be generalizable to apply to other types of networks. Conceivably, a specialized test network – for example, a

network representing connections between a diverse group of people - could be stood up and used to validate a more generalized model's application to this situation.

The third objective listed – provide an instrument for at-scale experimentation – follows from the research characteristics listed above. While this objective will be the most expensive – at least in terms of a single, large expenditure of funds – and broadly compelling since it will necessitate a major construction project involving not only large sums but also many people, it must be remembered that the primary objective is the research.

SLIDE 14

Before commenting on the outlook for the GENI Project, let me illustrate some of the research questions that need exploring.

The next visual lists four challenges, among many, for future networking:

- Security and robustness
- Pervasive computing with mobility
- Bridging physical and cyber space
- Realizing the potential of opto-electronics

along with a capability that will be needed to meet the associated challenge:

- Information access with high availability and trust
- Seamless information access anywhere and anytime
- Access to information about the physical world in real-time
- Access to bandwidth-on-demand with low latency and QOS guarantees.

The GENI facility will be able to support experimentation on new mechanisms that will provide these and other capabilities in a realistic environment in which performance can be monitored and measured in detail.

SLIDE 15

The initial support for research focused on Internet design is called FIND (Future Internet Design – see reference at end) and is being provided by NSF as a part of its on-going support for network research - the NeTS funding program that has annual calls for proposals (called 'solicitations' by NSF). The FIND program is in the midst of its second full cycle and is the current place to look for research activity specifically aimed at GENI-type research.

SLIDE 16 & 17 (provide information taken directly from the NSF solicitation for NeTS)

An example of the theoretical research that is envisioned can be found in the SING part of the Theoretical Foundations solicitation of NSF (see reference at end).

SLIDE 18 & 19 (provide information taken directly from the NSF solicitation that includes SING)

The FIND and SING programs were developed by astute, innovative NSF program officers (many of whom are “rotators” on loan from their university to NSF for a few years). They were encouraged by NSF/CISE leadership in order to “prime the pump” and encourage a new level of thinking. At the same time, we realized that there were many questions that need to be articulated before large sums are spent on building an experimental facility on which they can be explored (compare to the decades-long process of developing fundamental questions in particle physics and then specifying and building large accelerators and detectors such as the LHC at CERN). Consequently, parallel efforts were started to develop a more comprehensive “GENI Science Plan.”

The first public version (there were several earlier iterations with limited circulation) was posted this past April on www.geni.net. As part of the responsibilities of the GENI Science Council, this plan will be further developed over the coming months and years.

SLIDE 20

Finally, let me comment on the outlook for the GENI Program. In a word, it is *indeterminate* and cannot be accurately forecast. The reasons for this are many – it is about research, technology, and the future, after all! – but, the following visual may help you understand the situation.

SLIDE 21

This diagram is an expanded version of an earlier one this morning, but now includes two additional timelines, purposely shown as above the basic GENI Facility timeline because they will have impact and in some cases control over the GENI Program through funding and partnering decisions.

One of the first things to note is that from very early on, a variety of discussions have taken place – and continue – regarding partnering opportunities with the GENI Program. These include international discussions, conversations with interested industry groups, potential for collaboration with other U.S. Government agencies, formation of a GENI Project Office (external to NSF and just announced this Spring) to manage completion of the design of the GENI Facility and eventually its realization, and formation of the Computing Community Consortium (for wider leadership objectives) that includes the GENI Science Council that will provide community leadership and oversight for the GENI Program. Obviously, any of these potential partnering arrangements can and probably will impact the course that this Project takes.

The second thing to note is that a number of entities will eventually be involved in funding decisions that will have very direct impact on the outlook for the GENI Program, most especially the GENI Facility. Without going into detail here on the scope and interaction between these entities, let me simply note that the larger the amount of money under consideration, the more scrutiny there will be and the higher the level of the deciding authority will be. So, for example, decisions on small, individual research awards are made by a small number of scientific officers at NSF

who organizationally are at a fairly low level of authority. At the other end of the spectrum, appropriation of funds in the range of several hundred million dollars for the complete construction of the GENI Facility requires the approval of the U.S. Congress.

Thus, predicting timelines and even final results of a project with so many and such high-level authorities eventually involved is a hopeless task! All that can be said with certainty at this point is that there is strong support for the Project at the highest levels within NSF and growing and widespread support not only in the networking research community but the broader computing research community.

Please note, however, that all of these “outlook” comments refer to the GENI Project at NSF. Given various articles that have appeared in the popular press (NY Times, The Economist, Congressional Quarterly) and the interest that the Project is generating among researchers and others not only within the U.S. but internationally, I believe it is safe to say that:

- the “Geni is out of the bottle,” and
- the subject of *how* to develop future networking will differ significantly from the current, incremental approach.

Perhaps with less certainty of what will happen, but with even more conviction of the need for it, I believe that in the future more truly scientific approaches to network design involving serious, at-scale experimentation will be used.

Rather than try to summarize what I’ve said this morning, I think our time will best be spent in answering a few questions.

Let me again congratulate all of you on organizing and participating in this workshop. I am confident that your research and development, coupled with that of others, will help advance the frontiers of networking for everyone, to benefit the peace and advancement of all mankind. Thank you!

For More Information (starter list):

Technical

- www.geni.net (primary source for updated information on GENI)
- www.nsf.gov/pubs/2007/nsf07507/nsf07507.htm (FIND)
- www.nsf.gov/pubs/2007/nsf07525/nsf07525.htm (SING)

Policy

- www.oii.ox.ac.uk/research/
- www.issues.org/22.3/p_nelson.html
- www.oecd.org/topic/0,2686,en_2649_37441_1_1_1_1_37441,00.html
- www.caida.org/home/ (also technical)

Legal

- www.cyber.law.harvard.edu/home/
- www.cyberlaw.stanford.edu/