

Report of the GENI User Opt-In Workshop

Cambridge, MA

22 July 2008

GENI prototyping is being sponsored by the National Science Foundation to support experimental research in network science and engineering. A key goal of the GENI effort is to enable, indeed, to encourage users to participate in experiments running across the evolving suite of GENI infrastructure. Users could choose to opt into experiments of interest to them. Enabling this “user opt-in” has been recognized as a challenge, as described later in this report.

This workshop was organized by the GENI Project Office, with the generous support of the National Science Foundation, to solicit expert advice and ideas on the topic of user opt-in for GENI experiments. The hope and expectation was that by reaching outside the networking and distributed systems research communities, we could gain valuable insights from additional perspectives.

Brief Meeting Overview

The workshop was structured to solicit ideas from experts interested in empirical understanding of the behavior and dynamics of the Internet and its users, including Internet traffic, media economics, and incentives to participate in online communities and Internet retailing. Many are interested to better understand how to design Internet applications and architecture for increased privacy, security, innovation and fair auctions, for example. To encourage discussion about what GENI should *do* (vs. how it could be prototyped), participation from the GENI prototyping community was limited.

Invited participants submitted short white papers in advance of the workshop. The ideas in the white papers were divided into very broad sessions, where participants were asked to give short presentations to stimulate discussion. The workshop concluded with a session focused on trying to extract the key ideas and challenges for further study and potentially for future workshop agendas.

Participants were: Alessandro Acquisti (Carnegie Mellon University), Yan Chen (University of Michigan), Aaron Falk (BBN & GENI Project Office), Donna Hoffman (University of California – Riverside), Suzi Iacono (National Science Foundation), Joe Konstan (University of Minnesota), Bill Lehr (MIT), Jeff MacKie-Mason (University of Michigan), Harry Mussman (BBN & GENI Project Office), Thomas Novak (University of California – Riverside), Andrew Odlyzko (University of Minnesota), David Parkes (Harvard University), Craig Partridge (BBN & GENI Project Office, chair), Henning Schulzrinne (Columbia University), Michael Schwartz (Yahoo), Barbara van Schewick (Stanford University), Rahul Telang (Carnegie Mellon University), David Waterman (Indiana University) and Christopher Yoo (University of Pennsylvania).

What We Learned

The day was filled with ideas, several of which were refined through group discussion. Overall, most of the ideas fell into two broad thrusts, which were just inconsistent enough to cause creative tensions and yet consistent enough that both approaches to users could co-exist in the GENI infrastructure suite.

We also observed that users would not understand that GENI is a federated suite of infrastructure that supports research. They may not even know it exists at all (they'll simply opt into some interesting experiment and won't even worry about how it runs). At best, our expectation is that users will know GENI as something akin to a brand name that is somehow related to better networking. (That realization carries through to this report, which treats GENI as the abstract concept the users will see rather than worrying about how GENI will actually be composed).

Thrust 1: Think Applications First

A surprising realization in the meeting was that we needed to focus less on users and more vigorously on applications. Users are going to opt-in to a GENI experiment because there is an application that they want to try. Indeed, some applications (such as gaming), if occurring as experiments within GENI slices, would likely bring along their existing user base.

If we focus on serving applications, several issues immediately come to the fore, of which the most important seem to be:

- *How do we make GENI valuable for application developers to try out new or improved applications?* We discussed the possibility of using GENI's measurement infrastructure to give application developers a better understanding of how their application was behaving and how it was being used (with due considerations for privacy, both for users and to allow application authors, many of which may be corporations, to protect proprietary data).
- Another vital point is that application developers may try creating new GENI experiments, or leveraging existing experiments, to use new communications technologies that improve application performance. Different applications may desire different experimental communications technologies. This observation suggests it may be useful to create a place (marketplace, dating forum) where a communications researcher running experimental technology in a GENI slice can be matched with application authors looking for specific networking capabilities (and vice-versa). It was observed that such matchmaking has virtues for both sides, as researchers get to show the benefits of their research to a nifty application and applications get improved performance.
- *What support will application authors need for running in a GENI slice?* If hosting an experimental application within a GENI slice is different from

hosting an application on the Internet (and it presumably will be very different), how do we make that transition easy for networked application developers to understand? We may have to build application toolkits and API libraries. We may need a group specifically charged to provide training and support for application authors.

- *Commercialization.* There were two issues. One (as hinted in the first bullet), we need to figure out how to support commercial application developers who may wish to offer novel applications running as experiments within GENI slices. Two, we need a plan for successful experiments to migrate from GENI into the larger, commercial world, painlessly. If GENI becomes an incubator where great ideas get stuck due to commercialization issues, we've missed our goal.

If we can solve these problems, there's a sense that GENI may become a place where application developers and users who feel limited by the Internet can come and incubate their ideas and create "markets of opportunity" – visible communities that are economically desirable to serve and which, by virtue of their existence, create pressure for technical evolution.

Thrust 2: Certain Users Matter A Lot

The other central observation from the meeting was that, in some situations, users matter a lot. The most visible example was Hoffman & Novak's identification of *emergent users*, defined as users with "the ability to process information in a synergistically experiential and rational thinking style, and exhibit a unique set of personality traits such that the users are able to recognize which product concepts mainstream consumers will find more appealing and useful, compared to mainstream, lead user, or even innovative consumers."

But we also found ourselves talking about lead users and about close-knit communities that viewed themselves as cutting edge and might seek to use experimental applications within GENI slices simply because the GENI "image" was seen as new and novel.

We talked a bit about how we make the GENI experience attractive to these different classes of desirable users and identified a few key ideas:

- *Keep the cost of joining low.* We may have to subsidize access to some portions of the GENI infrastructure (which suggests we need to understand which users are important to the GENI effort, so that subsidies go to the right parties).
- *Appeal to collective/public service aspects.* One of the ways to nurture close-knit communities is to make clear that we value their contributions and their service. (Indeed, as was noted, there are studies showing that, in some cases, paying them is demotivational!).

- *Weed out bad experiments/applications.* Successful experiments and applications will presumably grow up and migrate out of GENI (see the discussion in applications above). Less successful applications may linger and give GENI a reputation for mediocrity. We need to work to ensure that the GENI experience is more often positive and that weak/bad applications are not allowed to persist or accumulate. Note that if we weed out too vigorously we put tremendous stress on experimenters and may stifle their ability to “iterate to a solution that works” but if we fail to weed out weak stuff, we lose vitality.
- *Calibrate and communicate privacy rules and expectations.* Researchers would like to collect a considerable amount of measurement information via GENI. It seems likely that smart researchers could find unique correlations to reveal private user information (which is doubly damaging, as users lose privacy and also, often, behave differently because they know they are being observed). How do we assure appropriate privacy and collect the information needed by legitimate experiments? How do we properly inform users about what levels of privacy they can expect? Will the levels of privacy that GENI offers be sufficient that users will be willing to opt into GENI experiments?

Other Ideas

In addition to the two thrusts discussed above, there were a number of freestanding ideas, including:

- *Employees as users.* Many companies use their own employees as beta-testers of new products. Might we facilitate corporate beta testing with corporate users in the GENI infrastructure suite?
- *Institutional Review Boards (IRBs).* Workshop participants noted that emerging research experiments involving opt-in users would likely open up unfamiliar territory for IRBs, and consequently the group had a number of interesting discussions and opinions on the potential relationships between IRBs and such experiments. It was noted that IRBs vary greatly from university to university, with a wide range of policies, procedure, and core expertise, and there was much discussion of how various potential approaches to experiments with opt-in users might interact with IRBs. These issues are important, and should be explored further through another workshop concentrating on the relationship between IRBs and network science and engineering experiments involving opt-in users.
- *Externalities and tipping points.* Innovation in the marketplace comes when enough users/customers are motivated to use an alternative service, such that providing the service becomes viable. How do we exploit network externalities to reach this tipping point for GENI-based experiments?
- *Leading edge applications.* It was observed that leading edge application development is often driven by products that many people disapprove of

(e.g. porn played a role both in VHS adoption and adding video to the Internet). Is there a less objectionable leading edge that can help drive GENI? (The quick guesses were the distributed gaming industry or broadband entertainment).

- *Need to clearly articulate what GENI does that current networks do not.* Application developers will need a crisp statement of what the GENI suite of infrastructure enables for them.
- *Once someone has joined one GENI experiment, can we make it easy for him or her to join another?* Once someone has shown a willingness to try experimental applications, how do we encourage him or her to find and join other experiments?

Thank You

The GENI Project Office would like to thank the participants for their insights and the NSF for funding the workshop. We would also like to thank Kathleen Rich for handling the local arrangements.