

Will GENI Be a Magic Lamp or a Dim Bulb?



Greg Goth

Without much fanfare, a sweeping new networking research initiative was announced at the ACM Special Interest Group on Data Communications (SIGCOMM) conference at the end of August.

Presented as part of the works-in-progress session, the US National Science Foundation's (NSF) Global Environment for Networking Investigations (GENI; www.nsf.gov/cise/geni/) was so unheralded beforehand that even conference coordinator Joe Touch, director of the Postel Center at the University of Southern California's Information Sciences Institute and inventor of the X-Bone overlay technology, didn't know about the initiative.

Most presentations at the work-in-progress sessions fall under the "My grad student and I have this idea; what do you think?" category, Touch says. So he was surprised to see the potentially groundbreaking presentation under the NSF's "dealing with the Internet's core architecture" banner for that session.

"That's not where you expect to hear this kind of stuff come out," Touch says.

How revolutionary might GENI be? Its stated goals are to research the development of nothing less than a "clean slate" Internet. According to the NSF, the GENI initiative will support the research, design, and development of new networking and distributed systems capabilities by

- creating new core functionality –

going beyond the existing paradigms of datagram, packet, and circuit switching and designing new network-management paradigms and naming, addressing, and identity architectures;

- developing enhanced capabilities – building security into the architecture; designing for high availability; balancing privacy and accountability; and designing for regional differences and local values;
- deploying and validating new architectures – incorporating emerging technologies (wireless and optical, for example) and new computing paradigms enabled by pervasive devices;
- building higher-level service abstractions – using, for example, information objects, location-based services, and identity frameworks;
- building new services and applications – developing principles and patterns for distributed applications as well as making large-scale distributed applications secure, robust, and manageable; and
- developing new network architecture theories – investigating network complexity, scalability, and economic incentives.

To broaden participation in the initiative, the NSF says its Computer and Information Science and Engineering (CISE) board will work with industry, other US agencies, and international groups.

New Era for Funding?

Over the long term, GENI might succeed in some or all of its goals. But in the near term, the initiative signals a potential tectonic shift in the way the US research community proposes funding projects. Computer science researchers and US Congressional science experts have been concerned about several trends under Darpa, which has funded and overseen the bulk of large-scale Internet architecture research since the network's incubation days. For example, the proposed fiscal year 2006 budget for the US Networking and Information Technology Research and Development (NITRD) program is US\$2.15 billion, down slightly from last year's \$2.2 billion (www.aas.org/spp/rd/06pch23.htm). In addition, some are critical of what they view as Darpa's new emphasis on short-term research, which they believe will denude the pipeline of "blue sky" ideas, such as the packet-switching technology of the original Arpanet, that will reach fruition years down the road. Darpa Director Anthony Tether disputes these claims. In testimony before the Science Committee in the US House of Representatives (www.house.gov/science/hearings/full05/may12/), Tether said the agency is emphasizing a new multidisciplinary approach, and computer scientists must learn to live with the new era.

"The major complaint about reduced university funding seems to be coming from one discipline – computer science," Tether testified. "But if Darpa's overall funding of universities is more

News in Brief

The **Internet Society** has launched the *IETF Journal*, a new online publication that offers readers an overview of Internet standards activities, focusing on **IETF working groups**. The journal, published in cooperation with the IETF, will also cover important issues from IETF meetings and mailing lists. The first issue includes an IETF overview, general news, and reports from the recent IETF meeting in Paris, including articles on meeting activities related to IPv6, routing, and wireless technologies.

The journal is online at <http://ietfjournal.isoc.org>.

More than 100 regulators from **15 West African nations** have hammered out an agreement on a common **regulatory framework for information and communication technology (ICT)**. The market-based framework for ICT service provision is aimed at spurring ICT investment and development in West Africa, which is seeking to create a single EU-style market.

A full text of the regulatory guidelines is available at www.itu.int/ITU-D/treg/Events/Seminars/ITU-EC-Project/Ghana/Ghana.html.

The **W3C** has launched a **Mobile Web Initiative** to address usability and interoperability problems related to mobile Web access. The effort's initial focus is on four key areas: best practices, a Web site trustmark, the device information required for content adaptation, and marketing and outreach efforts. The initiative involves content providers, authoring tool vendors, handset manufacturers, browser vendors, and mobile operators. The **Mobile Best Practices working group's** first draft of a **Scope of Mobile Web Best Practices** is available at www.w3.org/TR/mobile-bp-scope.

More information on the initiative is at www.w3.org/mobile/.

In advance of November's **World**

continued on p. 9

or less constant, then other disciplines must be the recipients of Darpa's research funding. The key question becomes, 'What other discipline has grown significantly over the past five years at the expense of computer science?' As part of my investigation, I reviewed several dozen university Web sites to see if I could determine the new discipline that was on the rise.

"The answer is surprising on the one hand and obvious on the other: no single discipline has been taking over. Every university Web site I visited advertised that they had created centers for multidisciplinary research and professed that these centers were the harbinger of the future. What must be happening is that, while computer science is always part of the multidisciplinary efforts, its past dominance and, hence, respective share [of funding] must be decreasing in relationship to the other disciplines involved in the effort."

Some computer scientists say the result of the new order at Darpa is a dog-eat-dog atmosphere in which the struggle for continued backing trumps collaboration as funds diminish. If the NSF can fund GENI (GENI backers' preliminary figures totaled \$300 million over five years), Touch says he believes the zeitgeist among researchers will be one of cooperation.

"It may be very competitive in the proposal phase, but generally, once you get support from NSF there's just no need to compete with each other," he says. "They are much more collaborative and cooperative, encouraging and assuming that of their researchers."

Nonetheless, a shift in network research funding from Darpa to NSF would be epochal.

"I would be shocked," Touch says. "A typical program out of NSF is a couple million dollars. A typical program out of Darpa is tens of millions. If they were to get \$300 million to do this, that would be encouraging — it would be astounding, because it would be more than Darpa put toward a single program. I recall their programs being clos-

er to \$50 to \$60 million, never a couple hundred million. They're talking about three to four times what Darpa spends, and the NSF typically spends a tenth of what Darpa spends."

Touch also says computer science hasn't been a darling of NSF administrators.

"A lot of times in the hard sciences, they've been given that money and [said], 'Please do biology, please do chemistry.' In computer science, when it came to those big projects, they basically said, 'Please do enough networking so the biologists can get their work done.' Basically, we are the plumbers to them. They're not funding big projects to do big infrastructure revolutions. The GENI initiative talks about doing that, but the NSF has not historically been able to get the critical mass or funding together in a way that is decoupled from the support for the hard sciences."

However, Jennifer Rexford, professor of computer science at Princeton University and one of the movers behind GENI, says the universal dependence on the Internet might be a factor in bringing more people to believe major improvements — up to and including a "clean slate" initiative — need to happen, no matter who funds such a program.

"In the last couple of years, the security and network management problems have become so serious there's an erosion of trust in the Internet itself," Rexford says. "There's also a growing sense we're at risk of losing our edge in IT competitiveness globally. In talking to deans and administrators, there's a lot of talk that other countries are making a bigger push and our own funding situation is grimmer. There's a feeling that it's ripe for some big agenda to help drive us to do something more efficient."

Dave Andersen, an assistant professor of computer science at Carnegie Mellon University (CMU), is optimistic about GENI's potential and says he perceives some heightened interest among other researchers.

"It generates buzz and a little bit of

pushback skepticism, but I think a lot of people think this may have ... more impact than 'business as usual,' Andersen says. "I think [in] the networking community, in particular, we've had some amazing successes. But there's been some frustration in recent years that certain kinds of research, on the core architecture, for example, — 'Knowing what we know now, how would you build the Internet right?' — it's hard for that to have an impact because the Internet is a success, its out there, and you'd have to convince 300 million people to change. So I think part of the feeling behind GENI is, to overcome inertia, you're going to take a hit in the short term to get something better. But in order to do that, you have to have something big."

Where Will New and Old Pieces Fit?

If GENI's backers can't point to specific ways in which the initiative can leverage the millions of dollars already spent trying to improve the "old" Internet, getting R&D money for a "clean state" initiative in a time of funding cutbacks is doubtful.

Although the NSF's GENI announcement stated that, "the GENI Facility will leverage the best ideas and capabilities from existing network testbeds, such as PlanetLab, ORBIT, WHYNET, Emulab, X-Bone, DETER and others," Touch says nobody, including the NSF, has contacted him about how the X-Bone testbed (www.isi.edu/xbone/) might dovetail with the other technologies in the larger GENI picture. In fact, he says his efforts at integrating X-Bone and the PlanetLab technologies have been unsuccessful thus far, due more to matters of policy than technological incompatibility.

"PlanetLab is intended to be a testbed for people to do network experiments, and has perceived, I think, our stuff to be one of the experiments that could use their testbed as opposed to being one of the pieces that might be used to build or augment their test-

bed," Touch says. "So we have literally not been able to put our stuff into PlanetLab."

However, CMU's Andersen says even concerted discussion about GENI might help the research community rethink how the pieces fit together, including some orphan technologies such as TRIAD (Translating Relaying Internet Architecture integrating Active Directories; www-dsg.stanford.edu/triad/), the networking concept developed by Stanford University professor David Cheriton, which has gotten no traction in the community at large.

"TRIAD has very deep changes it wants to make, and so getting it deployed on the basis of one SIGCOMM paper is far more difficult than if you could push it out there in the field and really play with it," Andersen says. "TRIAD is really a classic example of the kind of thing I think GENI is aiming at."

Touch says he thinks there is a place in network research for serious appraisal of technologies like TRIAD, but he also wants to encourage efforts to ensure that existing investments and infrastructure don't get lost in a wave of change for its own sake.

"People have used the phrase *ossification* for rigidity," he says. "One man's ossification is another man's stability. The Internet is what it is precisely because it is the lingua franca of so many different things. It never was intended to do anything perfectly, but it was intended to do everything 'well enough.'"

"In the process of looking for these new architectures and testbeds, I want to build something that's more than what we have. I don't want to start with a blank sheet of paper; it'll be too hard to get back to what we have. Your clean slate here is the stable of researchers you can get. They're there. They're hungry. They want proposals. They want to do research. They don't like doing the 'development of the week.'" □

Greg Goth is a freelance technology writer based in Connecticut.

News in Brief

continued from p. 8

Summit on the Information Society in Tunis, the **US** has rejected a proposal — backed by the **European Union** — to transfer **ICANN oversight** from the US to a more "representative body" within the **United Nations**. The rejection came at **PrepCom-3**, a pre-summit meeting in Geneva in early October. In addition to the EU-backed proposal, at least three others are under consideration: strengthen ICANN's Governmental Advisory Committee, but otherwise keep ICANN as it is; establish an International Internet Council outside the UN and limit ICANN to a narrow, technical role; or create three new bodies in ICANN's place: a debating forum, a group to coordinate work on public policy issues, and a group to take over the Internet addressing system oversight and other ICANN duties.

More information on PrepCom-3 is available at www.itu.int/wsis/preparatory2/pc3/index.html; Internet Society delegates' commentary on the meeting is at <http://geneva.isoc.org/blogs/wsis/>.

Information on the upcoming WGIS, to be held 16–18 November, is at www.itu.int/wsis.

The number of **US Internet users** has increased 5 percent to a total of **68 percent** over the past year, according to the **Pew Internet and American Life Project's** latest survey. Among those adults least likely to go online are those over the age of 65 and those who didn't graduate from high school. The survey found that 53 percent of Internet users now have **broadband access**, up from 21 percent in 2002. Although 10 percent of those who don't use the Internet have been online at least once, 22 percent, or 1 in 5 American adults, have never used the Internet — a figure down only a single percentage point since 2002.

The full report is available at www.pewinternet.org/PPF/r/165/report_display.asp.