

GENI End-User Opt-In Working Meeting

1. Announcement

A working meeting of the [GENI End-User Opt-In WG](#) will be held:

Tuesday, April 14, 2009
10am - 4pm

Columbia University
Computer Science Building, Room 453
500 West 120th St
New York City
Dial-in number: 212-939-7127

Local contact: Henning Schulzrinne

Conference bridge: 866-453-5550 651-3886#

This meeting is by invitation only from the Opt-In WG Co-Chairs: Henning Schulzrinne at hgs@cs.columbia.edu and Helen Nissenbaum at hfn1@nyu.edu. If you have any question, please contact either of the co-chairs, or the WG System Engineer: Harry Mussman at hmussman@bbn.com

Or send mail to the WG mailing list at opt-in-wg@geni.net

Directions are at <http://www.cs.columbia.edu/resources/directions>, and here is a summary:

The department office is at the entrance to the Computer Science Building, itself located in the lobby of the 4th floor of the Mudd Engineering Building. There are doors to Mudd on the 1st floor at 500 West 120th Street (corner with Amsterdam Avenue). The doors on the inside of the campus are on the 4th floor. Enter through the main entrance to the campus at Broadway and 116th Street, walk along College Walk and turn left up the steps. At the top of the steps is a large building with a dome; Low Library. Take the walkway to the right of Low Library, and the Mudd building is at the end of this walk. Upon entering Mudd, turn right and then right again. This is the entrance to Computer Science.

2. Attendees

Attending: (9)

Henning Schulzrinne – Columbia University and WG Co-Chair hgs@cs.columbia.edu
Salman Baset – Columbia University (sa2086@columbia.edu)

Helen Nissenbaum – NYU and WG Co-Chair hfn1@nyu.edu

Solon Barocas – NYU solon@nyu.edu

Joe Lorenzo Hall – NSF ACCURATE Postdoctoral Research Fellow, UC Berkeley School of Information / Princeton University Center for Information Technology Policy joehall@gmail.com

Aaron Burstein – University of California, Berkeley (School of Law)? (Washington, DC) aaron.burstein@gmail.com

Larry Peterson – Princeton University, PlanetLab PI, llp@cs.princeton.edu

Matt Mathis – Pittsburg Supercomputing Center mathis@psc.edu

Harry Mussman – GPO, End-User Opt-In WG SE hmussman@bbn.com

Providing contributions and/or interviews: (2)

Justin Cappos – University of Washington justinc@cs.washington.edu (Interview scheduled 4/15/09)

David Kotz – Dartmouth College, CRAWDAD project kotz@cs.dartmouth.edu (Email received 3/30/09)

Contact later: (2-4)

Mark Allman and/or Vern Paxson – ICSI <http://www.icir.org/mallman/> and <http://www.icir.org/vern/>

Wendy Kellogg and/or Martin Wattenberg – IBM, Many Eyes project NSF (?)

3. Objective

The objective is to expand the use-cases in the DRAFT document: “GENI End-User Opt-In Overview” (working title), and finish with a document that is ready for a first review by the full WG.

To date, five end-user opt-in use cases have been identified:

- Use Case 1: User opt-in to GENI experiment for service
- Use Case 2: Wholesale opt-in to GENI of traffic
- Use Case 3: A disruptive GENI experiment
- Use Case 4: Opt-in of user resources to a GENI experiment
- Use Case 5: Gathering data involving opt-in users in a GENI experiment

These use cases then drive requirements on capabilities that are being included into GENI, including:

- Capability 1: Gateway from GENI to another network, e.g., the Internet
- Capability 2: Contribution (or association) of a user’s node to an experiment on GENI

- Capability 3: Gathering logs and experiment data on GENI (some possibly user-identifiable) and managing their distribution

Finally, we expect various common policy themes to arise from an examination of the policies and best practices required for each use case, such as:

- Privacy policies for end-users
- Disclosure policies

For each use case, we want to better understand and write down:
the players, their motivation and relationships;
major issues (incentives, risks to consider, etc.)
and a first cut at policies and best practices when operating in a GENI environment.

For example, here is a possible way for an experimenter and an end user to work together:

An experimenter wants to have opt-in participation on an individual basis. They publish a document that answers a set of standard questions so that participants (end users) can judge whether they want to play along and provide “informed consent”. Questions might include:

- What kind of experiment is this?
- Who is responsible for the experiment?
- What kind of data is being gathered by the experiment?
- How long will the data be retained?
- Will participation in this experiment affect other applications I'm running?
- Will the experiment prevent my PC from going to sleep?
- How much extra traffic will the experiment generate on my access link?
- Should I run this application on a cell phone or other low-bandwidth device?

In particular, we want to finish this meeting with a better understanding, and a few pages of text for each use case. Then, we hope to have a baseline view of GENI opt-in that can be expanded in the coming months.

4. Agenda

Before the meeting:

- a) Update current DRAFT document "GENI End-User Opt-In Overview" (working title) with revised use cases, and distribute to the attendees. (Harry)
- b) Take GENI opt-in references and submitted contributions, and distribute to the attendees. (Harry)
- c) Provide text for DRAFT document. (All attendees)
- d) Update DRAFT document "GENI End-User Opt-In Overview" (working title), with text provided by attendees, and distribute to the attendees. (Harry)

The current agenda for the meeting is:

- e) Summarize use cases. (Harry)
- f) For each use case, better understand and write down:
 - the players, their motivation and relationships;
 - major issues (incentives, risks to consider, etc.)
 - and a first cut at policies and best practices when operating in a GENI environment. (All attendees)
- g) Identify "cross-cutting policy themes", and map each to the appropriate use case. (Helen to lead)

5. References

- 1.1) [End-User Opt-In WG Meeting at the GEC-1: Meeting minutes \(October 2007\)](http://groups.geni.net/geni/wiki/GeniOptInMinneapolis) at <http://groups.geni.net/geni/wiki/GeniOptInMinneapolis>
- 1.2) End-User Opt-In presentation at the GEC-2:
<http://groups.geni.net/geni/attachment/wiki/GeniOptIn/geni-opt-in-0803.ppt>
- 1.3) [End-User Opt-In WG Meeting at the GEC-3: Meeting minutes, including all presentations \(October 2008\)](http://groups.geni.net/geni/wiki/OptInWGGEC3) at <http://groups.geni.net/geni/wiki/OptInWGGEC3>
- 1.4) [End-User Opt-In WG Meeting at the GEC-4: Meeting minutes, including all presentations \(April 2, 2009\)](http://groups.geni.net/geni/wiki/GEC4OptInWGAgenda) at <http://groups.geni.net/geni/wiki/GEC4OptInWGAgenda>
- 1.5) GENI End-User Opt-In Scenarios and Capabilities Summary, by Harry Mussman - GPO at [http://groups.geni.net/geni/attachment/wiki/GEC4OptInWGAgenda/102208b%20OptInWG Scenarios Capabilities Issues Documents.pdf](http://groups.geni.net/geni/attachment/wiki/GEC4OptInWGAgenda/102208b%20OptInWG_Scenarios_Capabilities_Issues_Documents.pdf)

A one-day workshop was held on July 22, 2008 in Cambridge, MA, that was structured to solicit ideas from experts interested in technology, and its impact on society. Conclusions include: the way to attract end-users is (mostly) through getting great applications into GENI slices. Thus, we need to think about how to ensure GENI is application-friendly.

2.1) White papers were submitted by participants in the workshop :

<http://groups.geni.net/geni/attachment/wiki/GeniOptIn/White%20paper%20from%20cycle%201%20participants.pdf>

2.2) Summary report from the GENI User Opt-IN Workshop (July 22, 2008) at

<http://groups.geni.net/geni/attachment/wiki/presentations/Partridge-Opt-In-Workshop-v1.0.ppt>

2.3) Final report at <http://groups.geni.net/geni/attachment/wiki/GeniOptIn/Opt-In%20Workshop%20Report-v21.5.pdf>

2.4) Contributions on Many Eyes visualization service from Wendy Kellogg and/or Martin Wattenberg – IBM.

2.5) Many Eyes is a long-running visualization service

<http://services.alphaworks.ibm.com/ManyEyes/> that is hosted on the IBM alpha Works platform.

2.6) The IBM alpha Works platform <http://www.alphaworks.ibm.com/> hosts Many Eyes and many other IBM services.

3.1) GENI Project on Regional Opt-In at

<http://groups.geni.net/geni/wiki/RegionalOptIn>

3.2) Regional Opt-In presentation at GEC3 at

<http://groups.geni.net/geni/attachment/wiki/presentations/OptIn%20WG%20%203c%20%20Mathis-regOPT.pdf>

3.3) Regional Opt-In presentation at GEC4 at

<http://groups.geni.net/geni/attachment/wiki/GEC4OptInWGAgenda/OptInGEC4.pdf>

3.4) Opt-In Requirements doc <http://groups.geni.net/geni/wiki/OptInReqs> and

<http://groups.geni.net/geni/attachment/wiki/OptInReqs/OptInReqs.pdf>

This document outlines requirements for a strong wholesale opt-in mechanism for GENI. When fully deployed it would permit GENI experimenters to request that ISPs redirect traffic from a huge population of innocent users through GENI

infrastructure. These users are innocent in the sense that they do not have to do anything at all to participate, and might not even be aware that they are doing so. Key to wholesale opt-in is that it fully engages the Institutional Review Board (IRB) process and that all participants are motivated by their own self interests to do the right thing.

4.1) GENI Project on Million Node GENI at
<http://groups.geni.net/geni/wiki/MillionNodeGENI>

4.2) SIGCSE 2009 paper on Seattle: A Platform for Educational Cloud Computing at
<http://groups.geni.net/geni/attachment/wiki/GEC4OptInWGAgenda/Cappos-SIGCSE2009.pdf>

4.3) Slides on Seattle from GEC3 at
<http://groups.geni.net/geni/attachment/wiki/presentations/OptIn%20WG%20%203a%20%20seattle.pptx>

4.4) Justin reviewed million-node GENI project, also known as Seattle, and presented a live demo of a user contributing a resource to Seattle at GEC4. See
<http://groups.geni.net/geni/attachment/wiki/GEC4OptInWGAgenda/Seattle%20-%20GEC%204.pdf>

4.5) From Justin, here is some of the information we have that explains how Seattle works. There are a few holes in this (notably how the SeattleGENI website provides resources to users), but I think this might be a good starting point. This wiki page has a very pure view of Seattle (no Seattle GENI).
<https://seattle.cs.washington.edu/wiki/SeattleComponents>

4.6) The node manager API also has a pure view (no Seattle GENI).
<https://seattle.cs.washington.edu/wiki/NodeManagerDesign>

4.7) GENI Project by David Anderson, Patrick Gunn, on “Prototype Support for Heterogenous Testbed Resources: Integrating cluster, broadband, and wireless emulation nodes into the “Proto- GENI” Framework”, at
<http://groups.geni.net/geni/wiki/CmuLab>

5.1) Contributions on Dartmouth Internet Security Testbed
<http://www.cs.dartmouth.edu/~dist/> and on CRAWDAD project
<http://crawdad.cs.dartmouth.edu/>, provided by David Kotz – Dartmouth College, and summarized at http://groups.geni.net/geni/wiki/CRAWDAD_Experience

5.2) Paper on “Legal Issues surrounding Monitoring During Network Research” by Paul Ohm, Douglas Sicker and Dirk Grunwald at
<http://portal.acm.org/citation.cfm?id=1298307>

5.3) References on “Sharing Internet Measurements” by Mark Allman and/or Vern Paxson at ICSI, summarized at <http://www.icir.org/mallman/research/proj-simr.html>

5.4) Particularly, a talk at ACM SIGCOMM/USENIX Internet Measurement Conference, October 2007, on Issues and Etiquette Concerning Use of Shared Measurement Data. Copies of the paper are at <http://www.icir.org/mallman/papers/etiquette-imec07.pdf>.

5.5) See also the slides at <http://www.icir.org/mallman/papers/etiquette-imec07-talk.pdf>

The slides are especially instructive to GENI opt-in policy.

5.6) Also, a related paper on The Devil and Packet Trace Anonymization at <http://www.icir.org/mallman/papers/devil-ccr-jan06.pdf>

5.7) Also, an earlier (2002) paper by Allman (when he was at BBN) that describes “A Scalable System for Sharing Internet Measurements”, which can be found at <http://www.icir.org/mallman/papers/>. It seems like a straightforward approach to solving the general problem.

5.8) Aaron Burstein, “Issues Relating to Data Acquisition, Retention, Use and Disclosure” at <http://groups.geni.net/geni/attachment/wiki/041409NYCOptInWGAgenda/051509%20%20burstein%20%20dataretention.pdf>

6.1) Understanding and Resolving Conflicts in PlanetLab, by Larry Peterson - Princeton Univ, at <http://www.cs.princeton.edu/~llp/policy.pdf>

6.2) Slides from GEC3 at http://groups.geni.net/geni/attachment/wiki/presentations/OptIn%20WG%20%202a%20%20llp_policy.ppt

6.3) A description of the PlanetLab Consortium at <http://www.planet-lab.org/consortium>

7.1) GENI Project on Bringing Experimenters and External Connectivity to GENI at <http://groups.geni.net/geni/wiki/DTunnels>

7.2) Presentation at GEC3 by Nick Feamstra, on “Bringing External Connectivity and Experiments to GENI”, at <http://groups.geni.net/geni/attachment/wiki/presentations/OptIn%20WG%20%2003b%20bgp-mux-gec3.ppt>

8.1) GENI Project by Paul Barford, on “Instrumentation and Measurement for GENI”, at

<http://groups.geni.net/geni/wiki/MeasurementSystem>

9.1) GENI Workshop on security, with slides from GEC4 at

<http://groups.geni.net/geni/attachment/wiki/Gec4presentations/GEC%204%20-%20202.%20Matt%20Bishop%20-%20security%20wkshp.ppt>

10.1) Aaron Burstein, "Toward a Culture of Cybersecurity Research" at

http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1113014

10.2) Aaron Burstein. (2008). Conducting Cybersecurity Research Legally and Ethically. USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET '08). Retrieved July 24, 2008, from

http://www.usenix.org/event/leet08/tech/full_papers/burstein/burstein.pdf

10.3) Aaron Burstein. (2008). Amending the ECPA to Enable a Culture of Cybersecurity Research. Harvard Journal of Law & Technology, 22(1), 167--222. Retrieved October 14, 2008, from

<http://jolt.law.harvard.edu/articles/pdf/v22/22HarvJLTech167.pdf>

10.4) Creative Commons provides free tools that let authors, scientists, artists, and educators easily mark their creative work with the freedoms they want it to carry:

<http://creativecommons.org/>