
3rd GENI Engineering Conference
Opt-In Working Group Meeting Minutes

Prepared by Opt-In Systems Engineer:
Harry Mussman at GENI Project Office hussman@bbn.com,
November 4, 2008

Content:

The agenda for the conference can be found at: <http://www.geni.net/GEC3/GEC3-Agenda.pdf>

All slides from the conference can be found at:
<http://groups.geni.net/geni/wiki/presentations>

On the first day of the conference, there were two talks in plenary session that summarized NSF-sponsored workshops, including one on the “User Opt-In Workshop”. Notes on this workshop are presented first in this report.

Then, on the second day of the conference, the Opt-In WG met. It heard one short “lightning” talk, three related Spiral 1 project talks, the system engineering report, and had a period of open discussion. Notes from this meeting conclude this report.

One action item was identified:

Agree: WG should consider producing: “Best practices for involving an Institutional Review Board when starting a GENI networking experiment”. How can we get this started?

Plenary Session: Workshop Reports
Tuesday, October 28, 4:15pm – 5pm.
Building 20 Auditorium, Hewlett Packard, Palo Alto, CA
For an audio recording of this session, go to: TBD

“Report from the GENI User Opt-IN Workshop”

Speaker and workshop chair: GENI Outreach Director Craig Partridge at GENI Project Office craig@bbn.com

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

Related workshop report at: TBD

This talk reviewed a one-day workshop (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 users is (mostly) through getting great applications into GENI slices. Thus, we need to think about how to ensure GENI is application-friendly.

Opt-In Working Group Breakout Session

Wednesday, October 29, 2pm – 5pm.

Bldg 5M Spyglass Auditorium, Hewlett Packard, Palo Alto, CA

For an audio recording of this session, go to: TBD

1) WG Co-Chair: Henning Schulzrinne at Columbia hgs@cs.columbia.edu

Review of agenda

Introductions

(Note: WG Co-Chair Helen Nissenbaum was unable to attend.)

2) “Lightning talk” relevant to WG. (Invited by WG Chairs)

a) “PlanetLab Policies: Learning on the Job”

Speaker: Larry Peterson at Princeton llp@cs.princeton.edu

Slides:

http://groups.geni.net/geni/attachment/wiki/presentations/OptIn%20WG%20%202a%20%20llp_policy.ppt

This talk summarized PlanetLab experience, starting with the actors and how experiments lead to complaints. It ended with a review of suggested policies and advice.

Here is a summary of the slides:

- PlanetLab Policies:
 - Learning on the Job
- Acronym Soup
- CALEA
 - CFAA
 - CMA
 - DMCA
 - ECPA
 - SCA
 - Acronym Soup
- CALEA
 - Communication Assistance for Law Enforcement Act
- CFAA
 - Computer Fraud & Abuse Act
- CMA
 - Computer Misuse Act (UK)
- DMCA
 - Digital Millennium Copyright Act
- ECPA
 - Electronic Communications Protection Act
- SCA
 - Stored Communications Act

- Actors
- PlanetLab Central
 - testbed operator / trusted intermediary
- Hosting Sites
- Researchers
- Third Parties
 - end-users
 - content providers
 - attackers
 - ISPs
 - Experiments Complaints
- Measurement Studies
 - network probes
 - application probes
 - edge probes
- Deployment Studies
 - opt-in strategy
 - privacy
 - abusing services
 - bandwidth shifting
 - financial gain
- Policies/Advice
- Do not police content
 - DMCA is relevant
- Prefer opt-out to opt-in
 - IRB oversight (of human studies) not required
- Protect privacy of log files
 - ECPA & SCA are relevant
- Research is distinctive
 - CALEA is not relevant
- Adhere to best practices
 - avoid random probes (CMA - UK)
 - prevent service abuse (CFAA)

Questions from the audience:

Q: What should you do if asked to capture more per CALEA?

A: Argue that it is not appropriate, or shut down the experiment.

Q: (Craig Partridge) we need to protect privacy of log files; is there a way to make them publically accessible?

A: Keep them private, or modify data to be anonymous.

Q: What does opt-out mean?

A: Make sure that you are not included or touched by the experiment.

Comment: Matt at Internet2: We are looking at a process for releasing data.

Q: (Heidi Dempsey) what are EU requirements? The need to be able to delete their own records? Do you do it?

A: Not yet.

3) Talks by related Spiral 1 projects:

a) “Seattle: Building a Million-Node Testbed”.

PIs: Thomas Anderson and Arvind Krishnamurthy at the University of Washington – Seattle

Speaker: Thomas Anderson at the University of Washington – Seattle

Slides:

<http://groups.geni.net/geni/attachment/wiki/presentations/OptIn%20WG%20%203a%20%20seattle.pptx>

This talk presented the case that existing testbeds don't model the Internet. Instead, we should use resource donation and end user applications to build a million node testbed to realize: platform diversity; network diversity; and scale. See

<https://seattle.cs.washington.edu>

Questions from the audience:

Q: (Henning Schulzrinne) what do you suggest for incentives?

A: Diversity of opinion: 1) Give away app, but take 10% of resources. 2) Make it easy to write an app.

Q: (Larry Peterson) How does GENI workflow fit this model?

A: Not well. User has to do it; it is not driven from the network.

Q: (Rob Ricci) this is good if research app provides a benefit to end users. But, does this exclude research apps that do not benefit end users, e.g., a measurement app?

A: Measurement works only when it is integrated into another application.

Q: (Henning Schulzrinne) But, DIMES is a client downloaded to your PC to make measurements, and it is used.

A: Yes, but there are only a few 1000s of users, via volunteerism.

Comment by John Wroclawski: Consider not just volunteers, but also viral marketing.

Q: (Henning Schulzrinne) Can user benefit include access to I/O devices?

A: Yes, but not important.

Q: (Max Ott) how do you know if users are representative?

A: We don't know.

b) “Bringing External Connectivity and Experimenters to GENI”.

PI and speaker: Nick Feamster at Georgia Tech.

Slides:

<http://groups.geni.net/geni/attachment/wiki/presentations/OptIn%20WG%20%203b%20b%20gp-mux-gec3.ppt>

This project will provide for connectivity to and from the Internet, via a BGP mux platform that provides BGP interconnects and tunnels. It is based on the Quagga routing suite.

Questions from the audience:

Q: With a big AS, you can have multiple BGP sessions, at multiple locations. How will mux do this?

A: Like any AS does it.

Comment by Heidi Dempsey: We will not have any upstream connectivity until Internet2 agrees to provide it.

Q: Will mux be used only for BGP routing, or in a more general case?

A: Whenever you want upstream connectivity, reachable from this network, you need upstream via this mux.

c) "Regional Opt-In".

PIs: Matt Mathis and Ken Goodwin at Pittsburg Supercomputing Center and Three rivers Optical Exchange

PI and speaker: Matt Mathis at Pittsburg Supercomputing Center and Three rivers Optical Exchange

Slides:

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

This talk presented a technique to allow regional Internet traffic flows to be directed into a GENI experiment, and reviewed how such a process could be managed on a day-to-day basis, including failure or success scenarios.

Questions from the audience:

Q: (Heidi Dempsey) Open flow switching is another way to do large-scale opt-in. Can you compare?

A: Yes, but opt-in via open flow switch is local; is L2; and is specifically selected. This is L3, and can be used at anywhere from campus exit router (relatively easy) to backbone router (very difficult). They are complementary in functions

4) "Opt-In System Engineering Report"

Speaker: Harry Mussman at GENI Project Office hmussman@bbn.com

Slides:

http://groups.geni.net/geni/attachment/wiki/presentations/Opt_in%20WG%20%204%20%2020102908%20%20SE_Report_OptInWG_GEC3.ppt

This talk first provided an introduction to the role of the Opt-In WG system engineer:

- Harry Mussman
 - Current: Senior Systems Engineer in the GPO at BBN
 - Last: Voice-over-IP architect at BridgePort Networks (a startup) and GTE Internetworking/Genuity
 - BSEE Univ Michigan, MSEE Northwestern Univ, PhD Stanford Univ
 - hmussman@bbn.com
- Role of Opt-in WG SE
- Frame technical issues from top-down
 - Collect issues from WG, organize and revise
 - Use to identify and structure WG documents
- Synthesize input from bottom-up
 - Collect input from WG, compile and distribute
 - Look for and summarize consensus (or lack of it)
- Draft WG documents...
 - Manage process to completion
- Assist WG communications
 - Take and distribute notes
 - Maintain wiki

Then, a possible way to summarize opt-in scenarios and capabilities was presented:

- Overview of GENI
 - Opt-in Scenarios and Capabilities
- “GENI end-user opt-in is defined by scenarios where end users (not researchers) become involved with GENI experiments, and by the capabilities necessary to support them.”
- Each opt-in scenario involves:
 - One or more use cases.
 - A set of players, with differing motivations.
 - Unique issues.
- Each opt-in scenario requires:
 - A set of policies and best practices.
 - One or more underlying GENI capabilities to make it work.
- Opt-in Scenarios
- Have currently identified four distinct scenarios (which can overlap).
- Scenario 1: An experiment on GENI that provides a service to users and/or others
- Scenario 2: An experiment on GENI that can affect and disrupt users and/or others
- Scenario 3: When users (or others) contribute (or associate) resources to (with) experiments on GENI
- Scenario 4: When experiment data involving users is gathered on GENI, and made available to researchers, and possibly others
- Other scenarios?
- Opt-in Capabilities
- Have currently identified three basic capabilities that will be required.
- 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, and managing their distribution
- Other capabilities?
- Current Spiral 1 projects are focused primarily on providing a capability.
- Capability 1: Gateway from GENI to another network, e.g., the Internet
 - 1601 Virtual Tunnels
 - 1650 Regional Opt-In
- Capability 2: Contribution (or association) of a user's node to an experiment on GENI
 - 1645 Million Node GENI

Next, it presented an overview of related Spiral 1 projects.
(See slides)

Then, current opt-in issues were reviewed:

- Current Opt-in Issues and Tasks
- Scenario list:
 - Review, and affirm or change.
- For each scenario:
 - Formulate in more detail, and sub-divide as needed.
 - Understand the players, their relationships and their motivations.
 - Formulate proposed policies and best practices.
 - Clearly identify and define key issues.
 - Determine required capabilities.
 - Reference current implementations, research and projects.
- Who in WG wants to contribute?
- continued (2)
- Capability list:
 - Review, and affirm or change.
- For each capability:
 - Formulate in more detail, and sub-divide as needed.
 - Formulate requirements.
 - Suggest design approaches.
 - Reference current implementations, research and projects.
- Who in WG wants to contribute?
- Work towards defining and understanding scenarios is underway.
- Scenario 1: An experiment on GENI that provides a service to users
 - See Opt-in Workshop report by Craig Partridge.
- Scenario 2: An experiment on GENI that can affect and disrupt users and/or others
 - Consider experience from PlanetLab, by Larry Peterson.
- Scenario 3: When users contribute (or associate) resources to (with) experiments on GENI
 - Studied as part of Spiral 1 project, 1645 Million Node GENI.

Finally, the documents planned for the next year were reviewed.

- Planned Opt-in Documents
- Architecture:
 - Opt-in Architecture, v1 DRAFT due 3/1/09
- Scenarios?
- Capabilities?
- References?
- Who in WG wants to be an author? Editor?
- Who in WG wants to be an editor?
- Who in WG wants to be a reviewer?

- How WG Creates Document
- SE drafts document, with input from WG
- GPO does internal review
- SE posts first draft
 - On wiki (to start); repository up RSN
- WG discusses document on WG list
 - Possible one-on-one follow-ups
- SE assembles changes, and revises
- SE posts revision
- (Repeat, until document completed)

- Next...
- Notes, slides, action items, etc will be sent to the working group mail list and posted on the wiki page:

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

Question:

Q: (Harry Mussman) we heard about the PlanetLab issues; what about Emulab and DETER?

A: (Ted Faber) Emulab has many fewer issues, because more self contained. DETER has no issues, since completely isolated.

5) Discussion, including comments and questions from the audience:

Q: (Henning Schulzrinne) who is audience for documents?

A: Future builders and experimenters. Those who provide required software and hardware.

Comment from Josh Bers at BBN: As a member of the operations community, would like you to consider use of GENI to try out novel techniques to monitor network.

Topic for discussion: The need to involve an Institutional Review Board when starting a GENI networking experiment.

Q: (Matt Mathis) How much of the personality of an IRB is institution specific?

A: Quite a bit.

Comment by?: Had to go through an IRB to survey students; went very badly, but eventually it was judged to not be needed.

Comment by Matt Mathis: Had submitted a short proposal that has not been funded: Produce a white paper titled “Networking for IRB Members”, Work has not yet started.

Comment by Heidi Dempsey: IRB interface should be responsibility of experimenters.

Agree: WG should consider producing: “Best practices for involving an Institutional Review Board when starting a GENI networking experiment”. How can we get this started?

Q: (Henning Schulzrinne) How much disclosure is due users in the million-node scenario?

A: Based on Sony activity that left active software on platform without user’s consent, you must at least inform them.

Comment by Heidi Dempsey: OMIS WG is working of security policies needed for Spiral 1 operation; this overlaps with User Opt-In.

Q: (Henning Schulzrinne) Do we need architecture? Data format work? Standard disclosure?

A: Yes. Yes. Yes.

Q: (Harry Mussman) How many are on Opt-In WG mailing list?

A: Small minority of audience.
