Report: Security Experimentation on GENI

Sean Peisert
University of California, Davis
and Lawrence Berkeley National Laboratory
http://www.cs.ucdavis.edu/~peisert/

September 4, 2012

1 Introduction

What kind of security experiments can be run on GENI? Are any of these new classes of experiments that are particular to GENI? What modifications might be needed to GENI policies or infrastructure to enable such experiments to be run? If or how should GENI Operations be involved with security experiments? What is considered a "security"-relevant experiment? These questions form the theme of this report, as we examine what the concerns of GENI are, what the concerns of other, similar facilities are; and what makes GENI different.

The Global Environment for Networking Innovation (GENI)[1] is a distributed testbed in which users can request resources for networking experiments. It is intended to be "*open and broadly inclusive*," providing collaborative and exploratory environments for academia, industry and the public to catalyze groundbreaking discoveries and innovation in these emerging global networks." (Emphasis added.) As such, unlike Emulab, in which testbed resources are located within a single location, GENI users can request resources called "slices" that cut across several testbeds, simulating "real" network variability in terms of latency, bandwidth, and reliability. As such, GENI is subject to most of the same restrictions as Emulab, with the added restriction of the fact that GENI can cross a variety of interconnections between the testbed resources, including the ordinary Internet. As such, many restrictions which would also apply in law or policy to the Internet would also apply to GENI.

At the "Security Experimentation on GENI" session at GEC13, Stephen Schwab gave a presentation[2] on the ethics of performing security experiments in a live, shared network infrastructure, such as GENI. This presentation both introduced the draft code of ethics and helped to prompt discussion on allowable and unallowable classes of security experiments. Indeed, it may well serve to guide what is acceptable on GENI. The draft version 0.9 code of ethics currently reads as follows:

1. Avoid conducting experiments that are harmful to other GENI participants, the GENI infrastructure, or the Internet.

2. Coordinate security experiments in advance with GENI operators.

3. Respect the privacy and confidentiality of other GENI participants and users.

4. Access GENI resources and services only when authorized.

[1] GENI: http://www.geni.net/

[2] Stephen Schwab's Draft 0.9 GENI Experimenter's Code of Ethics: http://groups.geni.net/geni/attachment/wiki/GENISecurity/experimenters-code-of-ethics-draft-0.9.pdf

The discussion about this code of ethics indicated that the code was relatively uncontroversial. However, there was considerable discussion in the vein of "Well, what about this?" The discussion highlighted the community's discomfort with "rules" that are not both precise and exhaustive. That said, a GENI AUP that is both precise and exhaustive may not be possible. The specific problem is that the resources are not under the GPO's control, and differing institutions may have very different requirements. The GPO may co-ordinate this through a clearinghouse, but if the policies of the constituent resource owners are not precise or exhaustive, neither will those of GENI be. Further, such policies are also constantly changing, so keeping them up to date would be impractical. Thus, contacting operators and other experimenters (as indicated in the code of ethics) may be the most feasible guideline.

2 Comparisons

Other, existing testbeds have explored both minimal and more extensive acceptable use policies. Given that GENI is not alone in its place as a network testbed, we now explore the policies imposed by other systems to gain insight into what has worked or not for others, and therefore what may work for GENI.

Emulab Emulab[3] provides a set of policies that are quite general:[4]

> "Abuse" of the facility or its other users, in any form, will of course result in termination of access. Abuse includes using the facility for other than a project's stated purpose.

These constraints clearly prohibit disruption of other experiments even though the experiments do not cross the Internet. Therefore, one might imagine that policies for GENI would be even more strict. Note that ProtoGENI[5] currently runs on top of the Emulab testbed codebase.

PlanetLab Like GENI, PlanetLab[6] is a network that spans many domains across the world as well. PlanetLab's acceptable use policy indicates:[7]

> A good litmus test when considering whether an experiment is appropriate for PlanetLab is to ask what the network administrator at your organization would say about the experiment running on your local site. If the experiment disrupts local activity (e.g., uses more than its share of your site's Internet bandwidth) or triggers complaints from remote network administrators (e.g., performs systematic port scans), then it is not appropriate for PlanetLab. It is your responsibility to ensure that your use of PlanetLab falls within these constraints. This means you should debug your code in a controlled environment so you have confidence that you understand its behavior.
>
> PlanetLab is also designed to allow experimental services to run continuously, thereby supporting an end-user community. As a consequence, PlanetLab could indirectly support users that have not officially registered with PlanetLab, and may even be unknown to you (the service provider). It is your responsibility to ensure that your users do not cause your service to violate the terms of this AUP. In particular, service providers

[3]Emulab: `http://www.emulab.net/`
[4]Emulab Administrative Policies: `https://users.emulab.net/trac/emulab/wiki/AdminPolicies`
[5]ProtoGENI: `http://www.protogeni.net/`
[6]PlanetLab: `https://www.planet-lab.org/`
[7]PlanetLab Acceptable Use Policy (AUP): `https://www.planet-lab.org/aup`

should ensure that their users are not able to hijack the service and use it to attack or spam other nodes or network users.

PlanetLab is designed to support network measurement experiments that purposely probe the Internet. However, we expect all users to adhere to widely-accepted standards of network etiquette in an effort to minimize complaints from network administrators. Activities that have been interpreted as worm and denial-of-service attacks in the past (and should be avoided) include sending SYN packets to port 80 on random machines, probing random IP addresses, repeatedly pinging routers, overloading bottleneck links with measurement traffic, and probing a single target machine from many PlanetLab nodes.

...

PlanetLab provides absolultely no privacy guarantees with regard to packets sent to/from slices. In fact, users should assume packets will be monitored and logged, for example, to allow other users to investigate abuse (see previous paragraph).

While "abuse" is more specifically defined by PlanetLab, some "disruption" is in fact allowed (e.g., "probing"). Interestingly, some complaints from external network administrators are expected and possibly even tolerated in PlanetLab experiments.

Internet2 Internet2 is a national network in the United States that connects universities and research laboratories.[8] As with Emulab, its policy largely prohibits disruption, although it does not specifically define disruption. Its AUP says:[9]

> The Internet2 Network can be used for any legal purpose, so long as it does not interfere with or adversely affect the operation of the Internet2 Network or any network user, as may be determined by Internet2.

ESnet ESnet is a U.S. Department of Energy network managed and operated by Lawrence Berkeley National Laboratory that connects national laboratories, universities, and other research institutions.[10] Its Acceptable User Policy (AUP) outlines general guideoines, acceptable use, unacceptable use, and also enforcement with regard to violations.[11] The ESnet AUP addresses a number of points relevant to security experimentation. Among other things, items that are unacceptable include:

- Use of ESnet to gain unlawful access to computational, information, or communications devices or resources.
- Intentional introduction of malicious code, such as computer worms or viruses. Transmission of material, in violation of applicable copyright laws or patents.
- Attempting to intercept, redirect, or otherwise interfere with communications intended for others

What is notable about ESnet's policy in contrast to Internet2's policy is that "legal purpose" is more specifically defined (e.g., copyrighted material). As with many such lists, this has the benefit of making such activities more explicit for potential users, but the danger of seeming to exclude things if they are not listed.

[8] Internet2: `http://www.internet2.edu/`

[9] Internet2 Network Acceptable Use Policy (AUP): `http://www.internet2.edu/network/aup.html`

[10] ESnet: `http://www.es.net/`

[11] ESnet Acceptable User Policy: `http://www.es.net/about/governance/ESnet-Acceptable-Use-Policy/`

NERSC The National Energy Research Scientific Computing Center (NERSC) is a U.S. Department of Energy supercomputing center managed and operated by Lawrence Berkeley National Laboratory.[12] In contrast to some other government facilities, NERSC is an entirely open and unclassified facility. Its policy says:[13]

> Users must not intentionally introduce or use malicious software such as computer viruses, Trojan horses, or worms.
>
> ...
>
> Users may not deliberately interfere with other users accessing NERSC or other system resources.

In this case, "disruption" is more specifically defined (and "prohibited") than that for PlanetLab and Internet2.

DETERlab DETERlab also runs on top of the Emulab software suite, but with additional requirements that seek to constrain malware for the purposes of experimentation. When requesting an account on DETERlab,[14] the questions are specifically asked:

> "Privacy and Threat Level" "Do you plan to use any malicious code?" "Do you need external connectivity?" "Can you ensure that your code does not generate traffic on our control net? (192.168.0.0/16)" "Please list any concerns about eavesdropping from other experiments"

Though DETERlab asks these questions, the implication is that answering "yes" to most of these questions simply results in additional scrutiny, rather than a rejection of an experiment, because DETERlab has the ability to monitor, manage, and contain such experiments. This is particularly notable given that DETERlab has recently been connected to cyber-physical devices [**?**]. However, although the addition of such devices expands the potential consequence of attack against DETERlab resources, the devices are currently simple and downside of attack is relatively benign.

3 Summary

At GEC13, we held a session on security experimentation on GENI[15] in which we attempted to come up with answers for the following five questions posed in the beginning of this document. The recommendations that came out of it were roughly as follows:

1. The types of security experiments that can be run on GENI are still open to debate, but largely, anything is appropriate as long as it is ethical, does not conflict with other experimenters, and is acceptable and legal with regard to interactions with external users. We note that interactions with external or *opt-in users*[16] are a separate ethical and legal subject that have been covered at GEC14 in a discussion session[17] on "GENI Opt-In Users," and will undoubtedly continue to be explored further in the future.

[12] NERSC: `http://www.nersc.gov/`
[13] NERSC Computer User Policies: `http://www.nersc.gov/users/accounts/user-accounts/nersc-computer-use-policies-form/`
[14] DETERlab: `http://www.isi.deterlab.net/`
[15] GEC13 Session on Security Experimentation: `http://groups.geni.net/geni/wiki/GEC13Agenda/SecurityExperimentation`
[16] About Opt-In Users - GENI: http://groups.geni.net/geni/wiki/AboutOptInUsers
[17] GENI Opt-In Users: `http://groups.geni.net/geni/wiki/GEC14Agenda/OptIn`

2. There are indeed new classes of experiments that are particular to GENI. Specifically, the diverse set of aggregators means that security policies must be accounted for, and experiments considering diverse security policies may have a home on GENI. Also, experiments that can take advantage of heterogeneous systems and networks may find particular value in running on GENI.

3. GENI policies or infrastructure would probably not need to be modified much other than policies relating to external users. A code of ethics may address the rest. An acceptable use policy may still be necessary if a code of ethics is determine to be insufficient.

4. GENI operations should be informed of security experiments, per a code of ethics.

5. A "security relevant experiment" is in the eye of the beholder: if the experimenter's ethical hackles are raised, then it is probably a security experiment. More specifically, a security relevant experiment is probably one that explicitly impact others (experimenters, aggregators, control framework, ordinary Internet, users). A pure networking experiment uses bandwidth, but the goal is clearly not explicitly to impact others. Tests of DOS attacks clearly are different. An exhaustive list of what is a security experiment will never be made. An exhaustive list of what is not a security *threat* may be easier. Instead of focusing specifically on security experiments, it is probably most effective to develop AUPs and codes of ethics for *all* experiments. Then, security experiments do not need to be treated as "special," and issues such as using too much bandwidth are seen as asocial ot antisocial, regardless the experiment is related to "security" or not.

These recommendations still seem to hold up. GENI is not designed first and foremost with security experiments in mind. Clearly some testbeds are designed in such a way, such as DETERlab and the National Cyber Range (NCR).[18] Such systems have the ability to contain out-of-control malware and do not even rely on external resources that would be harmed by malware-generated traffic transiting their systems and networks, for example, in a distributed denial of service attack experiment, or in an experiment about Internet "kill switches" such as those believed to have been used in Egypt in early 2011.

As such, for now (and the forseeable future) "open science" security experiments that are primarily malware-based or are otherwise heavily potentially disruptive to external environments will be limited to the DETERlab testbed that is federated with GENI.[19] What this means is that one can use a GENI account to stitch together an experiment that involves DETERlab resources and other GENI resources. Malware will however have to be contained within DETERlab.

But there are classes of security experiments that do not involve malware. For example, the Hive Mind[20] project is using GENI to experiment with a distributed monitoring architecture that could be used for detecting security events. The Attribute-Based Access Control (ABAC) project[21] could perform human factors experiments to examine the usability of the ABAC mechanisms. Davis Social Links (DSL)[22] can conduct human factors experiments regarding reputation-based, trust-based and relationship-based email exchange. Some have proposed doing wireless security experiments using the wireless resources in GENI. Other possibilities include using GENI to get

[18] National Cyber Range (NCR): `http://www.darpa.mil/Our_Work/STO/Programs/National_Cyber_Range_(NCR).aspx`

[19] Trial Integration Environment in DETER project: `http://groups.geni.net/geni/wiki/TIED`

[20] Hive Mind project: `http://groups.geni.net/geni/wiki/HiveMind`

[21] Attribute-Based Access Control (ABAC) project: `http://abac.deterlab.net/`

[22] Davis Social Links (DSL): `http://groups.geni.net/geni/wiki/DavisSocialLinks(DSL)`

access to real background traffic and using this to train anomaly detectors, for example. As pointed out in the AUP for PlanetLab, "minimally" disruptive experiments may be allowable, as long as the GENI Project Office can monitor those experiments and cope with the administrative and political liabilities of allowing them, particularly given that such policies are not precisely defined.

Clearly this list is not and cannot be exhaustive, but the common thread among these experiments is that they are both controllable and minimally disruptive to external systems and networks. When GENI projects are opened to external users, an additional level of importance will be placed on the issue of security experiments and GENI that also involves privacy. Setting hard resource limits on network usage, processor load, disk space, memory space, and other system and network metrics may go a long way to deal practically and automatically with disruptive experiments, albeit at the cost of flexibility of experiments, particularly over a period of many years and meaningful processor speed, disk space, and network bandwidth statistics evolve. Other experiments, such as those involving privacy, seem more likely to need to rely somewhat on a combination of a code of ethics and/or an acceptable use policy and/or some level of human security monitoring that cannot be automated. That is, to paraphrase the late U.S. Supreme Court Justice Potter Stewart, the GENI Project Office would have to put some energy into "knowing it when they see it."

Acknowledgements

Many thanks to Stephen Schwab (USC Information Sciences Institute) and Gale Pomper (U.S. Department of Defense) for their presentations as well as all of the participants at the GEC 13 GENI Security Experimentation Session. Sincere appreciation to Matt Bishop and Stephen Schwab for their comments on this draft report as well.

This work was supported in part by the National Science Foundation and the GENI Project Office under Grant Number CNS-0940805. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author and do not necessarily reflect those of any of the sponsors of this work.