Requirements for Wholesale Opt-In

Matthew Mathis
Pittsburgh Supercomputing Center
mathis@psc.edu
Version 0.3

Friday 20th February, 2009

Abstract

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. It does not require "selling" GENI to application developers or anybody other than network researchers.

Since the ISPs bear a disproportionate share of the risk with this technique, the ISPs are granted additional controls, implemented in an "ISP Daemon" that serve to isolate the ISPs' critical core routers from GENI.

When this approach is fully mature we expect it to be sufficiently robust where NSF might include GENI opt-in language in research solicitations across the entire foundation. This could bring as much traffic to GENI as is currently carried in aggregate by all of today's production research and education backbones.

1 Introduction

This document outlines requirements of a strong wholesale opt-in mechanism for GENI. By "wholesale", we mean opting-in large numbers of ordinary users who did nothing on their own to participate in GENI. It is "strong" because we assume full engagement of the Institutional Review Board (IRB) process for review of human subject experiments, and have structured the mechanism such that it supports all stakeholders in meeting their goals and responsibilities. The mechanism described here has the property that nearly all participants are motivated by their own self interests to do the right thing to bring the masses to GENI in a very tightly controlled way.

The exceptions are the people who run the network, who have little to gain by participating in GENI and often suffer the majority of the blame when things go wrong. Without loss of generality, we will refer to the people who run the network as the "Internet Server Provider" (ISP), independent of how they are organized or whether they are responsible for one university building, a national backbone network or some intermediate scale. We believe that the issues described in this document apply at all ISP types and scales, although weaker solutions might be be acceptable to certain small scale ISPs, such as departmental network support staff. The success of wholesale opt-in is critically dependent on supporting the ISP and its primary mission, supporting its users.

The IRB has a special role: it is mandated by law[1] to supervise all U.S. federally funded research involv-

ing human subjects. Its fundamental objective is to balance the gains from doing the research against the risks born by the experimental subjects.

Computer Science and Network researchers have sometimes had very bad experiences because their institution's IRB is focused on Biomedical and Social Sciences research and is not prepared to assess the risk and merits of Networking and Computer Science research. By their very nature, IRBs are cautions about evaluating risks that they do not understand, and thus have sometimes denied permission to conduct important research in these areas.

We believe that the optimal strategy is to strengthen the IRB's ability to evaluate GENI research by recruiting additional IRB members with backgrounds in Networking and Computer Science Research. At institution that have different IRB review tracks for different research fields, this might include the creation of an NCS review track within the IRB. At institutions that rely on combined IRB's, this might entail the creation of an NCS advisory panel, responsible for evaluating some of the more subtle risks and gains that might result from GENI experiments. In either case we refer to these as an "NCS panel" of the IRB.

We want to engage the IRB because it can be quite powerful. For example under some conditions it can supervise the collection of data sets that are protected from subpoenas, typically to study such things as drug addiction and crime rates. Although we are not lawyers, we suspect that a carefully planed IRB NCS track might be able to trump telecommunication law under some conditions.

Our small proposal to investigate the creation of an NCS panel was not funded. For this reason, the material in this document on strengthening the IRB process is rather limited, and only mentions some of the key points.

The requirements presented in Section 3 are derived from goals listed in the rest of this section and a survey of the stake-holder's agendas, responsibilities and constraints, described in Section 2. Appendix A describes some experimental scenarios that are useful for framing the problem.

1.1 Goals

We see the following goals for large scale GENI opt-in.

1.1.1 Interpose GENI Experiments

We want to create an opt-in mechanism that can interpose GENI infrastructure and experiments in the path of all traffic passing through a major campus interconnection point (e.g. a campus backbone, DMZ, GigaPoP, Regional Optical Network or national backbone). This will enable experiments requiring completely authentic user traffic. Note that although this traffic is by definition authentic, it is also intrinsically based on legacy protocols.

1.1.2 Strongest Possible Position

The Opt-in policies, procedures and techniques have to be sufficiently robust where we can imagine the NSF ultimately including GENI opt-in language in research solicitations across the entire foundation. If NSF is going to spend many millions of dollars on GENI infrastructure, it is not unreasonable for NSF to encourage potential long-term benefactors of GENI to also participate in the experiments as innocent users, even if they are not themselves a party to the experimental services or results.

We anticipate this to be an evolutionary transition. When GENI is in it's infancy, the only GENI users might be people who have individually explicitly opted-in. As GENI matures, it would make sense for other NSF Computing and Network Systems (CNS) solicitations to include language about GENI participation. Of course NSF cannot mandate that their awardees opt-in to GENI experiments, but they can mandate that awardees participate in the opt-in conversation and justify any decisions not to opt-in. After further maturation, this might be extended to all NSF CISE solicitations and perhaps eventually to all institutions accepting funding from any NSF division. Evolving opt-in to this scale would require that all risks be well managed.

1.1.3 Fully engage the IRB process

We want to strengthen the IRB to better support Networking and Computer Science experimentation and to encourage researchers to engage and rely on the IRB as an ally for addressing subtle policy issues. The IRB must be better equipped for evaluating GENI opt-in experiments. In particular it has to be able to accurately gage the benefits and risks of unintended consequences of GENI experiments.

1.1.4 ISP coverage

The ISP is ultimately responsible to it's users and the quality of service that they receive. For a variety of reasons nearly all ISP's are hypersensitive to complaints from their users. They do not want to be responsible for things that they can not control and they do not want to be blamed for events that are not their responsibility. As a consequence the lines of responsibility and control have to be crystal clear such that the ISP can explain them to any users who are complaining to the wrong people.

ISPs must have a strong and accurate way to transfer the blame for any unfortunate events (such as a GENI experiment crashing) to appreciate parties, otherwise the ISPs will elect not to participate by not providing the access needed to implement opt-in.

1.1.5 Fine Grained Control

We want the opt-in mechanism to have sufficiently fine grained control to precisely balance the needs of the researchers and the users, such that GENI can support full scale experiments while controlling the extent to which services become ossified by addicted users. Some relevant scenarios are described in Appendix A. The opt-in mechanism must permit individual GENI researchers to independently manage the balance between supporting a large user community and agile experimentation.

1.1.6 A Foundation for Experimentation

We want to lay the foundations for formal policies and procedures to experiment with the Internet.

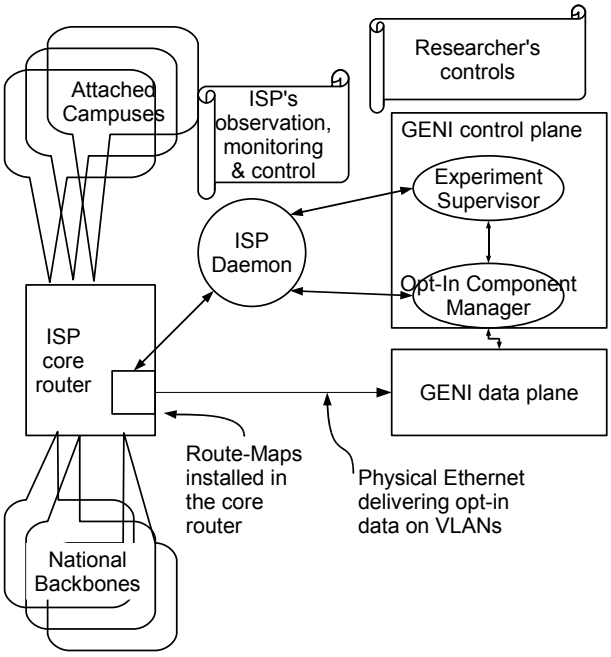

Figure 1: Opt-In Diagram

When the Internet was young, it was itself experimental, and everyone who used it had, by that act, opted-in to an experiment. Today, the Internet is no longer an experiment, but it has yet to develop a healthy culture of supervised experimentation, as it already exists in the biomedical and social science fields.

1.2 Justification and Approach

We envision an ISP intercepting and rerouting traffic through a core router as shown in Figure 1. Without loss of generality, the diagram assumes a typical "Regional interconnect" setting, exemplified by Three Rivers Optical Exchange.

We claim that the discussion and requirements presented here could be applied to any key router at any level of the Internet (e.g. Internal departmental or campus core, within a backbone, or national interconnect). Some locations, such as a router within a single department, may be able to streamline these

procedures and techniques, but only because the distinction between the ISP's role and other roles can be blurred in small local settings.

The actual traffic interception is done by matching all transit traffic against a pattern and invoking some action to reroute the traffic through the GENI infrastructure. The mechanism might be Open-Flow [2], the the "route-map" command on Cisco routers[3] or the "firewall filter" command on Juniper routers. At high speeds, all of these mechanisms normally rely on Ternary Content Addressable Memory (TCAM) to match packet headers at line rate. In many routers, the TCAM is a heavily constrained critical resource that must be managed carefully.

For most mechanisms, the pattern match is quite general can match many IP header fields, including IP addresses, DSCP byte, IP options, etc.

The intercept occurs in one or more of the ISP's core routers. These devices are expensive and mission critical to the ISP and all of their users. They are the "Jewels of the Kingdom" as far as the ISP is concerned.

Intercepted traffic has to be eventually be delivered to the remote destination, otherwise it is guaranteed to disrupt user applications. The easiest way to do this is to return the traffic to the same core router, but via a path that is exempt from a potential second intercept. Additionally, outbound traffic (from campuses to the wider Internet) can be delivered via experimental backbones to other locations, presumably near to the destination network.

In the GENI control plane, the opt-in control is subdivided into two separate elements that function at two different levels. The "Opt-In Component Manager" resembles other component managers for link like devices, except unlike other links, when properly invoked and enabled, the opt-in link provides traffic from the real world into GENI. The Opt-In Component Manager would be responsible for the managing the "plumbing" to connect the opt-in intercept to the other components of the GENI slice.

The "Experiment Supervisor" portion of the opt-in control is responsible for assuring the integrity of the entire slice. It determines if the entire slice is healthy and ready for the researchers experiment. The Experiment Supervisor would also be responsible for most opt-in shutdowns.[1]

The GENI data plane might provide an additional layer of filtering, beyond that provided by the interception point. This additional filtering might for example include deep packet inspection or an unconstrained instance of OpenFlow[2]. If there is a second opt-in filtering step within the GENI data plane, the interception point would probably have a fairly general match (e.g. all hosts in a given address block) and the traffic rejected by the second match would be returned to the core router to be delivered as normal.

The "ISP daemon" isolates GENI from the ISP's core router. There are a number of reasons why this is needed, described in Section 2.3.1. Suffice it to say that in general ISPs are likely to refuse to allow anybody or any thing not under the ISP's direct control manipulate the configuration in it's core routers. We envision the ISP daemon being mostly GENI provided code, with perhaps plug-ins to do ISP specific policy and access to the core router. It must enforce the ISP's policies, which should encompass enforcing the IRB authorizations. It also must expose all currently invoked and pending Op-In intercepts to the ISP and present the ISP with manual overrides to individually or collectively disable them.

The ISP daemon should run on ISP manged hardware in the ISP's premises. The ISP must have permission to inspect (and potentially alter) the code as they see fit. Note that the ISP daemon may expose GENI key material to the ISP, but it must not expose ISP key material to other GENI components.

1.3 3ROX Implementation

To get a better sense of the scale of Figure 1, we describe the existing components at Three Rivers Optical Exchange (3ROX).

We have two core routers running in tandem in geographically distributed locations. The dual routers provide redundant backup services for each other. They are based on carrier grade Cisco 6500 chassis,

[1]The ISP Daemon would also be able to preemptively shut down experiments, for emergencies or other reasons, but under normal operation it would be preferable for the Experiment Supervisor to shutdown experiments itself.

with multiple 1 Gb/s and 10 Gb/s Ethernet interfaces. Fully populated, they cost about a quarter of a million dollars each.

Our user base consists of the Pittsburgh Supercomputing Center itself, four large universities (CMU, PITT, PSU, and WVU), several "Intermediate Unit" networks connecting most K-20 schools in western Pennsylvania, and a number of smaller institutions and businesses such as libraries, museums and the CERT. All told we estimate that we have roughly two hundred thousand users. The vast majority of our traffic is for research and educational, but we do carry some commercial traffic.

We connect our users to two different commercial backbones (Sprint and Global Crossing) and to both major research backbones (Internet2's Abilene and National Lambda Rail's Transit Rail). We also have connection to various NLR experimental services and the potential to connect to the similar Internet 2 services.

Downtime is very expensive: 3ROX customers pay in aggregate more than $1 per second for our services, averaged over the long term. The opportunity cost of disruption to the users during to prime time down time is probably a couple of orders of magnitude higher. How would it color your perception of risk if any simple error that caused a router reboot during prime time potentially cost your customers $60,000?

2 Stakeholders

Each group of stakeholders has their own interests and constraints. The trick to making opt-in work well is to make sure that all of the stakeholders are self motivated to do the right thing. If GENI requires cajoling people to do something that is really not in their self interest, it can not succeed at really large scales.

2.1 GENI and Individual Researchers

The researchers are the driving force behind the entire process. They design and implement experiments, involving the IRB and opt-in mechanisms as needed to recruit users. We distinguish between two classes of users: "active users", who actively choose to participate in an experiment, for example by using an experimental service or installing experimental software. In contrast "innocent users", are ones who did nothing to participate in the experiment. An innocent user might be somebody who's normal Internet service is somehow redirected through GENI infrastructure. The point of wholesale opt-in is to have a mechanism and proper policy controls and procedures such that very large pools of innocent users can be enlisted to participate in GENI experiments.

Without loss of generality we assume that most experiments will be hybrids, involving both types of users. For example a typical experiment might include a small pool of active users using an advanced service and a large pool of innocent users to provide authentic background traffic and a contrasting service class.

By law, federally funded research involving people must be reviewed by an IRB. This clearly applies to opting-in a pool of innocent users. Although the process differs by intuition, this generally involves the researcher submitting a formal application with a description of the experiment to the IRB. The application is issued a unique identifier, and sometimes the researcher and the IRB iterate on the the description. Ultimately the IRB process results in a final description of the experiment, bearing an "approved" designation. The likely "credential" is a letter from the IRB to the researcher and to the funding agency, bearing the application's unique identifier with the negotiated final experimental protocol (procedure) as an attachment.

2.1.1 Modeling Experiments

The experiments themselves are likely to utilize an number of different opt-in scenarios as described in Appendix A. These scenarios can be implemented by a applying sequences of traffic matching rules and actions at the intercept point.

To facilitate implementation and control, the traffic matching rules and actions can be modeled as a set of named template rules, which the researcher applies in the proper order to implement the experi-

ment. They are templates in the sense that they may not be fully qualified, for example they may not include the individual IP addresses for the developers (See Section A.1.3) or the destination slice for the traffic.

Different portions of the scenarios might require different authorizations: for example graduate students should probably be able to opt-in their own workstations, but only senior personnel might be authorized to invoke large scale wholesale opt-in.

The approved IRB application (the human readable document) must be translated into a set of template rules for matching traffic. This translation is likely to involve engineering insight in to both the experiment and the network, and is likely to require some negotiation between the ISP engineering staff and the researcher. It is unreasonable to expect that the template rule mechanism be able to fully represent all of the complexity that researchers might request and successfully negotiate with the IRB. When there are selection criteria that can not be represented by matching rules, we distinguish between "ISP enforceable" and "GENI implemented" portions of the IRB approved experimental procedure. It is assumed that the ISP's opt-in pattern match will select a minimal super-set of the desired traffic, and further selection will be done down stream, within the GENI infrastructure.

The representation of these template rules is beyond the scope of this document. However the point of the template rules is to permit the ISP daemon to automatically match requests from GENI against the ISP enforceable template rules which were derived from the experiment described in the human readable IRB application.

We envision a couple of ways in which this might be done:

- Under "explicit construction", requests from the GENI control plane come in the form "I am userQ, please invoke Rule2 of IRB Plan9 with arguments x, y, and z". The ISP daemon is responsible for the final authorization check, constructing the actual pattern match and installing it at the interception point.

  In the short term we will be exploring this approach.

- Under "implicit validation", GENI sends fully formed standard requests to the ISP daemon, which matches them against the rule set and forwards them to the core router as appropriate. This approach is a better fit with standardized protocols, such as openflow[4], but has some implementation difficulties, particularly from ambiguities associated with matching IRB authorizations against requests.

  We do not expect to further investigate of this approach.

- Under "explicit validation", standard requests are tagged with authorization information of the form "I am userQ please invoke the attached OpenFlow message under Rule2 of IRB Plan9". The ISP daemon would then perform the final authorization check, verify that the OpenFlow message is consistent with the named template rule, and then forwards it to the core router as appropriate.

  This is probably the strongest approach of the three, but it may require protocol adaptation to associate the authorization information with already standard OpenFlow requests. This might be done, for example, by wrapping OpenFlow messages in opt-in request messages.

2.2 Institutional Review Board

This section is a vague since it summarizes the anticipated results of an unfunded GENI small proposal.

The laws defining Institutional Review Board and related process are specified by CFR Title 45, Part 46[1]. This is one portion of the legislation defining the role of the United States Department of Health and Human Services. It defines HHS as the lead agency for the protection of human subjects, but also specifically encourages other government agencies to adopt their own versions.

The NSF version of the code[5], is nearly identical in content to subpart A of the HSS code[2], although

[2]NSF reverts to the HSS version of the code for subparts B, C, and D which cover pregnant women, prisoners and children

due to formatting differences the HSS version is easier to read. The NSF maintains some excellent web pages[6] that summarize the IRB rules as they apply to NSF research.

The ultimate goal of the IRB can be paraphrased as balancing the risk to the subjects against the gains from the research. It is explicitly responsible for considering the rights and interests of the experimental subjects, and acts as a proxy for them.

The IRB process expressly requires protecting the confidentiality of Personally Identifiable Information (PII). A substantial portion of the required training for conducting research on human subject covers PII and the precautions that must be taken prevent unintended information leaks. We believe that if a competent panel of Network and Computer Scientists certify that the data collected by a specific experimental procedure does not expose PII, then this data is unlikely to be of interest to anybody who is concerned with enforcing telecommunication law, even though it is not expressly exempt. (Note that we are not lawyers).

Put another way, rather than trying to construct a "one-size-fits-all" anonymization standard for network packet traces, as a community we are better off developing an IRB based process such that experimenters can design their own anonymization (or choose from a menu of established techniques) for a specific experimental process, and have the data collection and anonymization process certified as properly protecting PII.

This approach might permit a researcher who wants detailed data about one end of each conversation to use a fixed anonymization of one endpoint IP address in exchange for much less information about the other endpoint[3]. It would be up to the IRB to determine if a particular data collection design sufficiently protected PII.

"Informed Consent" [5, §.116] is problematic for network opt-in as envisioned in this document. Fortunately, there are provisions for waiving this requirement, and a clear analogy from another field: Consider an experiment to evaluate a new design for highway detour signs. Obtaining consent, for example by photographing license plates and mailing consent forms to driver's home addresses, incurs far greater risk than the experiment itself, since doing so would leak PII. People may not want others at their home address to know that they traveled a particular road at some unexpected time. For most GENI experiments, waiving informed consent would probably entail locally publishing a description of the experiment and providing network users a means to Opt-Out.

IRB rules require that the reviewers have relevant qualifications for the research at hand. In particular CISE proposals must be evaluated by IRB's that include Computer Science and Networking people. In our opinion, rejected reviews by IRBs with insufficient expertise should be challenged. [5, §.107].

We need to consider what needs to be done to permit an IRB at one institution to supervise a GENI wide experiment involving users from multiple institutions.

An IRB can also take the initial steps to obtain a "Certificate of Confidentiality", which is used to protect research data from subpoenas[7]. Although the law was intended for use by HSS agencies for conducting research into drug abuse and other crime, the wording of the law is completely general and might be applied in other areas, for example, to study Internet phishing.

2.3 Internet Service Provider

All good ISP's are driven by their responsibility to provide quality services to their users, often on a very limited budget. However, users frequently blame "the network" for all sorts of problems associated with geographically distributed applications and services, even if when they are not really network problems. As a consequence ISP's are generally hypersensitive to risk — especially from things that might go wrong in ways that they can not not anticipate, control or diagnose. A popular personal motto at many ISP's is "Paranoia Pays".

Managing the ISP's perceived risk is absolutely

[3]Fixed anonymization, such as a fixed cryptographic hash of IP addresses, isn't normally considered strong enough because an attacker can use brute force searching over 32 bit IP addresses to invert the hash. However, if only one endpoint is present, this is not sufficient to identify conversations and thus much less interesting from a legal standpoint.

critical to the success of large scale opt-in. One way to understand the ISP's perception of risk it to recognize that they act as a proxies for all of their users, especially including the ones who are not selected to participate in a GENI experiment. Another model of would be to consider the ISP staff to be an unwilling participant in any experiment which has the potential to cause a user to complain to the ISP.

To get some insight into how ISP's normally address risk, consider procedures most ISPs use to schedule configuration changes. The first step is always to estimate of the risk associated with this sort of change, based on past experience. The level of risk is then used to select one of half a dozen or so different procedures for implementing the changes. A small number of (well tested) types of changes are permitted during normal operational hours. At the other extreme, significant service disruptions have to be scheduled and announced a well in advance. Most changes are only permitted during certain hours (e.g. early morning[4]) and have varying requirements for advance announcements.

To some extent judging risk is subjective. One area that does not have a good reputation are manual changes that affect the TCAM (e.g. all TCAM changes except normal routing updates). This is because in many routers the TCAM is an optional hardware acceleration for a function that is normally performed in software. In these devices overflowing the TCAM results in functionally correct operation, but possibly at an order of magnitude reduced performance. Such an event might be complete disaster during peak load but undetectable out of peak hours.

GENI opt-in does does not naturally fit well with typical ISP's conventions for managing risk. In the interim we expect all ISPs to want the ability to constrain the timing on invoking GENI opt-in. For example it is likely the ISP will want to limit first invocation of a new experiment to their normal testing window. Since ISPs will also want the ability to withdraw an experiments under any conditions at any time, disabling an experiment must be low risk in all operational states.

[4]This is an east coast convention. On the west coast, maintenance time is likely to be late in the evening since morning outages might affect east coast users.

It is likely that manipulating opt-in template rules can be decomposed into finer steps: for example installing a traffic matching pattern with a null action and then later amending the action to redirect traffic into GENI. This is useful if the different steps have different risks associated with them, and can be subjected to different restrictions. It could be used to mitigate most of the risk associated with manually manipulating the TCAM, as described above.

The ISP is also likely to want to disable all opt-in in the presence of serious failures, even if they are completely unrelated to GENI. For example if there is an outage of one of the major national backbones, there is likely to be a significant number of people trying to do their own debugging. It would be a very unpolitical for traceroutes to show that the traffic was going via GENI, even if GENI had nothing to do with the failure.

2.3.1 Isolated Management

The ISP's normal tools for managing routers are also subject to two very strong constraints. First, since there is a possibility that any significant ISP failure will cause all upper layer services to fail, including DNS, kerberos and other multi-party authentication services, there must be reliable mechanism to manage routers under crisis conditions. This typically means not relying on services beyond those existing entirely on the router itself. Thus at the lowest level, routers are usually managed using relatively crude but very robust techniques, such as reconfigured ssh keys. It would be very risky for opt-in to be managed with any protocols that might possibly be disrupted by the GENI experiments themselves. For this reason many ISPs will be very reluctant to use any multi-party authentication protocol to manage opt-in.

Since there has been a rash of attacks against routers in the public portion of the Internet, it is becoming more common for ISPs to use a private network to manage their routers. Routers can be configured with internal firewalls such that they will only respond to routing and signaling protocols on their primary interfaces, and only to their chosen management protocols on their management interface. Furthermore the management network is often

"net 10" or other RFC 1918 address space to provide an additional level of isolation and security auditing. Although these management isolation techniques are only common in large ISPs they are under consideration or in planing stages at nearly all medium ISPs such as 3 ROX, due to the recent attacks on routers.

These two techniques for securely and robustly managing routers under crisis conditions work against the GENI ever having direct access to core routers. GENI would have to have access to router cryptographic keys (which do not typically provide provide fine grained controls) and the private management networks, which would enable a GENI insider to disturb the routers in ways that the ISP could not observe and have great difficulty diagnosing.

One solution to this security and robustness problem is the "ISP daemon" as described in Section 1.2. It would have one interface on the ISP's private management network and have the necessary cryptographic credentials to manage the ISP's router configurations. To protect the ISP, it has to be under the ISP's control. It would instrument and log all actions it permitted on the core router and enforce all of the ISP's policies and the ISP enforceable portion of the IRB application. It would also provide the ISP with manual and/or automatic overrides to disable experiments.

2.3.2 ISP Procedural Requirements

We can summarize the ISP operational requirements as follows:

- Opt-in must not significantly impact the services that the ISP provides to other users.
- The opt-in mechanism must not leak privileges to potential rogue actors even if they have authorized access to GENI resources. The ISP must retain the ability to monitor, log and control every privileged access to it's core infrastructure, under all conditions and events.
- The ISP must be able to clearly identify and document the association between IRB certifications and potentially affected innocent users and track them as they change through time.
- When experiments go wrong, the ISP must be able to publicly document the connection between the failure, the experiment and the IRB authorization, such that it can fully defend it's decision to re-route user traffic. To the extent that the failure was disruptive, it needs to be documented to improve the process in the future. All parties involved in the failure need to work together to develop a post mortem report, which needs to be forwarded to the NCS panel of the IRB. This report should specifically address the question of how to better mitigate risk from this sort of experiments in the future.
- The ISP must have the ability to preempt opt-in experiments that are causing or merely appear to be causing service problems for users.

3 Technical Requirements

At this stage we are agnostic about some of the details about how the functions are partitioning between the ISP Daemon, Opt-In Component Manager and the Experiment Supervisor. We are assuming that many functions will be partially duplicated between the ISP Daemon and GENI control plane, because they need be implemented in slightly different ways to reflect the different perspectives of their constituents. For example, while it is important that the ISP know what experiments (and IRB authorizations) are active, the ISP is less interested in knowing the details of the roles within the experimental team, or which experimental states are active (See Appendix A). On the other hand the research team wants to precisely control the transitions between experimental states, possibly including fine grained control over who is authorized to invoke particularly risky state transitions.

Both the ISP Daemon and Experiment Supervisor need to have some sort of "operators console", such that humans have ultimate control over experiments. Again, these are likely to be quite different for the ISP and GENI researchers,

We are also not completely clear on how functions should be divided between the Component Manager and the Experiment Supervisor, since this is likely

to depend how similar functions are divided in other parts of the control plane.

3.1 Intercept

These are requirements for the opt-in intercept, a pattern match and action, installed into a core router or switch.

A.1 The Opt-In intercept must not compromise the ISP's reliability in uncontrolled ways.

A.2 The Opt-In intercept must be able to reliably remove or disable opt-in interceptions under heavy (much more than normal) load. Ideally this means no less reliable than normal routing table updates by routing protocols.

A.3 Assuming opt-in pattern insertion can be decomposed into different steps (e.g. separately install a pattern match, and then enable it), the last step to enable a particular opt-in rule should be reliable enough to be permitted during prime time under normal ISP operational rules. Ideally this means no less reliable than normal routing table updates by routing protocols. Steps other than the last step should be as reliable as possible, but not so unreliable that they have to be announced to the public, or planed more than 24 hours in advance (The worst case that we would expect would be restricted to a routine morning "minor update" window).

A.4 The Interception point must be able to minimally match on the fields supported by "Type 0" openflow [2, Table 1]. Ideally it would support matching on all link layer and IP header fields.

A.5 The available actions must include "No-op" (no opt-in, do normal forwarding) and "forward to some (logical) interface". Ideally it should be safe to change the actions while under full load.

A.6 All pattern matches should have counters which are independent of other actions. These are critical for debugging and validating configurations. These counts should be exported via the ISP daemon to both the ISP's console and to the GENI control plane for the experimenter's console.

3.2 ISP Daemon

The ISP daemon is responsible for providing a level of isolation between the ISP's core assets and GENI, as described in Section 2.3.1.

B.1 It must be feasible for the ISP to inspect the code, or rely on an independent certification by somebody else.

B.2 The ISP Daemon is responsible for implementing the IRB enforceable portion of the IRB approved experimental procedure.

B.3 The ISP Daemon must provide complete logging of all interceptions, and the associated IRB authorizations.

B.4 The ISP Daemon must provide both summary and detailed status information about all experiments in progress, showing traffic volumes for active experiments and time since last observed counter increment for inactive experiments.

B.5 Must provide the ISP a console interface to disable individual experiments or to refuse new experiments.

B.6 It is unclear at this point where to implement GENI authentication and authorization. It may be feasible for the ISP daemon to participate in the same authentication and authorization protocols as the rest of the GENI control plane, however we suspect that the ISP is not interested in the this level of control and would be satisfied relying on the authentication and authorization present in the Experiment Supervisor as long it is connected to the ISP daemon via a secure private channel.

B.7 The ISP daemon should accept input from multiple types of experiment liveness tests. Although we anticipate that the researcher would be primarily responsible for monitoring the

health of the experiment (e.g. monitoring connected to the Experiment Supervisor), the ISP may want to run backup tests, for example by sending traffic from inside of one of the campuses through the experiment. The ISP daemon should also monitor the health of the Experiment Supervisor. In particular if the GENI control plane becomes unresponsive, all associated experiments should be disabled.

B.8 The ISP daemon itself needs to be extremely robust and needs to have a failsafe mechanism that would disable all opt-in experiments if it ever fails.

3.3 GENI Opt-In Control

The GENI opt-in control conceptually decomposes into a low level Opt-In Component Manager, and a higher level Experiment Supervisor, as described in Section 1.2. Since this separation is still fluid, we list the requirements together.

C.1 The low level Opt-In Component Manager is responsible for connecting the VPNs (or other interfaces) carrying opt-in traffic to other GENI components. It is anticipated that the Opt-In Component Manager would be quite similar to other link like devices in the GENI tool kit, except it bring in traffic from the production Internet.

C.2 The opt-in Experiment Supervisor is responsible for managing the integrity of the experiment as a whole, including verifying that all other GENI components are proper allocated an configured before enabling opt-in.

C.3 The opt-in Experiment Supervisor must provide some sort of experimenter's console, such that the person running the experiment can monitor it's operation and has direct control over it.

C.4 The opt-in Experiment Supervisor must support multiple type of experiment specific liveness tests. e.g. researcher provided plug-ins to verify that the service is functioning properly.

C.5 The opt-in Experiment Supervisor must have some interlocks with the GENI slice and sliver allocation system, such that the experimenters GENI resources can not be deallocated or shutdown without first disabling opt-in. If it is necessary to do an emergency shutdown, opt-in must be the first component to be shutdown.

C.6 If the ISP Daemon is using "Explicit Construction" to build intercept pattern match rules (see Section 2.1), the GENI Opt-In Component Manager must provide all of the arguments to fill out the template rule. Alternatively, if the ISP Daemon is using "Explicit Validation", the GENI Opt-In Component Manager must provide the full pattern match and action.

4 Conclusion

This document takes a rather extreme position on opt-in. Although the policies and techniques described may seem to be excessive, they are appropriate for implementing truly large scale opt-in. For example our existing equipment could redirect all CMU, PITT and PSU traffic traversing 3ROX through GENI. The hard part would be doing so reliably, with the proper controls and permissions. The approach outlined here has the property that all nearly all participants are motivated by their own self interests to do the right thing. The one participant not so motivated is the ISP, 3ROX. This can be offset by giving the ISP full visibility and control over the opt-in, and not exposing any of the ISP's assets to tampering by rogue actors, even if they have GENI authorization.

In GENI's infancy, it is appropriate to streamline the policies and techniques described here. However, it will be easier to reach really large scales if we start with a full scale vision, and then simplify it for the early implementations.

As GENI evolves and grows, wholesale opt-in can evolve and grow to provide as much authentic traffic as GENI can possibly carry. To some extent the fate of GENI and large scale wholesale opt-in are linked: large scale opt-in can only be justified when it is used to feed shared research infrastructure. Experiments

that require large scale opt-in will be infeasible unless GENI succeeds.

5 Glossary

Opt-in: Selecting people to use GENI infrastructure, to provide researchers with users other than the researchers themselves.

Active users: Users who took some action to participate in a GENI experiment. The actions might be as simple as visiting a website, or as complicated as installing an entire operating system. If properly informed, active users have implicitly given consent to participate in an experiment.

Innocent users: Users who did nothing to participate in an experiment, for example by having their ISP redirect their traffic through GENI.

Wholesale opt-in: Opting-in a large number of presumably innocent users, for example by redirecting all traffic to or from a subnet or IP address block though GENI.

ISP: Internet Service Provider is used to refer to the people responsible for running the network, independent of their size, organization or business model.

Interception: Rerouting an ISP's regular transit traffic through GENI.

Interception point: The element of the production infrastructure where the traffic is intercepted, typically a packet matching filter running within the ISP's core switch or router.

hybrid experiments: Experiments that include both active and innocent users.

IRB: Institutional Review Board. Required by law to evaluate all research on human subjects[1]

NCS panel: A proposed Networking and Computer Science panel that might constitute one IRB approval track or advise a combined IRB.

ISP enforceable: Those portions of the IRB approved experimental procedure that can be implemented or enforced by pattern matching at the interception point. Portions of the experimental procedure that can not be implemented at the interception point, would have to be implemented elsewhere in GENI.

A Opt-in Scenarios

In this appendix we describe some example experiment scenarios. The opt-in mechanism, policies and procedures should support these and presumably additional scenarios. This material is not really part of the Opt-in requirements, but it is necessary to set the context and provide concrete examples for the discussion.

A.1 Simple Opt-in

The simplest op-in experiments involve rerouting traffic through GENI on a temporary basis. These experiments involve no long term commitments to either the users or the researchers. These are most suitable for one time or short run experiments or for debugging and prototyping larger longer running experiments that will ultimately fall into one of the other scenarios.

The simple scenarios also model the primitives needed to build up the more complicated scenarios described in the following sections. They are listed in order of increasing complexity.

A.1.1 Individual IP address

In the simplest case, the experiment applies to only a one specific IP address, probably in conjunction with other conditionals such as port numbers or service types. This is most likely to be used for testing experimental code, for example by opting-in an individual researchers own workstation. This case should not normally need IRB process or protection, however since there is not a general mechanism to distinguish between single user and shared systems, at the very least there has to be a way to white list systems that are authorized to opt-in themselves.

An individual opt-in might also be used to assure that an in-band slice liveness test is always routed via the GENI, even if no other users are opted-in. This might prevent false positive test results, caused by the liveness test being inadvertently routed around the experiment.

A.1.2 IP subnets

The experiment applies to all of the hosts on a single subnet, probably in conjunction with other conditionals. This could be a research project team performing a shared debugging session, in which case it would be considered self opt-in. At the other extreme, the subnet might be extended to include the entire network.

This might also be a common case for some types of experiments. This use clearly requires explicit IRB supervision.

A.1.3 Listed IP addresses

Same as Section A.1.2, except using an explicit list of address. This requires additional mechanisms to manage the list of addresses.

A.1.4 IP subnet with static exclusions

Combining Sections A.1.2 and A.1.2 except the static list is of excluded IP addresses. This might be used to exclude specific systems that have known operational or technical incompatibilities with an experiment.

A.1.5 IP subnet with dynamic exclusions

Same as Section A.1.4 except with the addition of a web page or other user tool to opt-out, so that anybody can opt out in real time.

This is beyond the scope of our current work, but since it potentially introduces some important constraints on the opt-in control, and has to be at least minimally included in the design.

A.2 Implementing version agility

One of the common recurrent problems with supporting real users on experimental services is that an installed user base makes it hard to make changes to an experiment. If the experimental service is useful, active users can become addicted to it, and any changes are likely to be disruptive and create strong disincentives for the active users.

In the section we describe a technique that might be used to provide strong version agility while avoiding creating strong disincentives for adoption by active users.

To illustrate the technique we assume a hybrid experiments, where active users invoking advanced network features have to load a custom application on their workstation.

The opt-in technique is as follows: Allocate two complete slices to the project. For sake of discussion name them “blue” and “green”. Assume the loyal user base is wholesale opted in on the blue slice in a production like configuration and the developers are alpha testing the new version on the green slice (presumably using listed IP addresses). When the new networking code is ready, notify the users that at a specific time they will need to stop, reload and restart their applications. At the appointed time toggle the mapping between the opt-in filters and the slices, such that the wholesale opt-ed in users are on the green slice, and the blue slice it available to the developers for the next development cycle.

This can be done with even less disruption if you add some additional constraints. For example if there is an explicit version number in the protocol that can be used to key the intercept, then you can do something like route all odd versions to the green slice and even versions to the blue slice. The user can then be migrated to alternate slices by upgrading their software at natural stopping points in their own work cycle.

There are many possible variants on this theme. As long as the researchers pay attention to version coexistence within their own research code, some variant of this technique can be used to maintain a large pool of real users while having the capability to make fairly frequent change to the experimental service.

A.3 Weaning Users

Once in a while experiments have suffered from success disasters: if the service is really useful, it can

spread by word of mouth to the point where too many people opt-in.

What typically happens is there are some high profile user who becomes addicted to the experimental service, and they influence the research project to extend the experiment beyond the point where it is useful to the researchers. Meanwhile, the good word continues to spread, and even more users discover the service, further increasing the user base making it harder to decommission. This self marketing is especially problematic if there is not a good way to identify who is using the services.

If the experimental service includes an opt-in network component (presumably in addition to a downloaded component), the opt-in mechanism can be used to manage the user experience. This addiction cycle can be broken by toggling the network from wholesale opt-in to some form of listed IP addresses to opt-in, such that the users have to have a individually request to continue to use the service. The researchers can exclude all new users while progressively tightening the criteria for ongoing use by legacy users. In this way fairly gracefully wean everybody off of the service, without causing undue disruption to any one user.

References

[1] Code of federal regulations title 45 - public welfare department of health and human services, part 46 protection of human subjects. 45 CFR 46, Obtain via: http://www.hhs.gov/ohrp/documents/OHRPRegulations.pdf.

[2] Nick McKeown, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Jonathan Turner, and Scott Shenker. Openflow: Enabling innovation in campus networks, March 2008. Obtain via http://www.openflowswitch.org/documents/openflow-wp-latest.pdf.

[3] Inc. Cisco Systems. Catalyst 6500 series switch command reference, 2003.

[4] Brandon Heller. Openflow switch specification, version 0.8.9, December 2008.

[5] National Science Foundation. 45 CFR part 690: Federal policy for the protection of human subjects. Obtain via http://www.nsf.gov/bfa/dias/policy/docs/45cfr690.pdf.

[6] National Science Foundation. Human subjects. Obtain via http://www.nsf.gov/bfa/dias/policy/human.jsp.

[7] The public health service act 301(d), 42 u.s.c. section 241(d).