

GENI Operational Security Plan

Document Name: GENI Operational Security Plan

Document Type: MOU

Version: 0.5.1 (Draft)

Date: Jul. 21, 2011

Goals & Scope

GENI, being the large cooperative effort that it is, relies upon the collaboration of several stakeholders (including, but not limited to the GENI Project Office, the National Science Foundation (NSF), the aggregate providers at many GENI-affiliated research organizations and campuses, and the engineers developing GENI's software and infrastructure) to provide a secure and stable research environment for scientists and engineers running experiments. The goal of this document is to (1) recommend the structure of an operational security team, (2) set forth the basic processes for incident response, and (3) recommend actions to mitigate known risks. This is necessarily a living document as GENI is rapidly evolving, with the potential for priorities and risks to quickly change.

Roles & Responsibilities

In order to even discuss the processes to follow in the event of a security incident, it is necessary to reference roles and entities that have yet to be established. Therefore, the discussion of responsibilities below references some new roles within GENI. If these roles are not created, named differently or somehow combined, this document will of course need to be revised. We do not address the issue of who has authority to create these roles since no long term governance structure for the GENI federation has been established to date.


GENI Computer Security & Incident Response Team (GENI-CSIRT)

To effectively address any inter-site security issues for GENI, there must be a team with primary responsibility for investigating and responding to security incidents for the GENI community, which we will abbreviate as the *GENI-CSIRT*. **If there is no such team, then tickets created around security issues will likely be passed around like a hot potato.** Hence, there needs to be some party responsible for security who will take ownership of these issues and provide the very necessary role of coordination between aggregate operators, experimenters and campuses within the GENI community.

It is vital, for redundancy, that there be more than one person on this team. GENI will not take a vacation, and so we must have backups. It is also important for at least one person on this team to be funded 50% or more on GENI if they are to take ownership of security and give GENI priority. This team should be created with members from multiple stakeholders in GENI. For projects the size such as GENI is evolving into, a security team (in aggregate hours) should consist of at least 3 full-time employees (FTEs). However, GENI is still ramping up operationally and may not need so many resources dedicated to security operations in Spiral 4.

The leader of this team, which we will call the *GENI Security Officer (GSO)*, needs to be at least 50% dedicated to GENI, and they need to have prior incident response experience. The GSO will be responsible for assigning investigation leads for security incidents and also reporting back to the GENI governance body. While the incident response team will collaborate closely with the GENI Meta-Operations Center (GMOC) at times, it is not necessary that members of the GMOC be on this team. However, communication with the GMOC and a clear point of contact with the GENI governance body will be necessary. The security team will likely also rely on the GMOC for services like an incident investigation wiki, a network filesystem (for incident response data), email lists for users to report abuse and incidents, and some sort of ticketing system, (which the GENI-CSIRT will probably share with a GENI help desk).

In addition to the GSO role, there should be a *Deputy GSO (DGSO)* funded at least 25%. The DGSO will fulfill all the duties of the GSO when the GSO is unavailable, but otherwise act as regular member of the GENI-CSIRT. During Spiral 4, GENI may be able to get along with just a GSO and DGSO on the GENI-CSIRT, depending upon the percentages funded.

 *Currently, there is no established governance body for GENI, and in the interim the role is most closely fulfilled by the GPO, who has significant input on project direction and funding allocations. Without knowing the name of the future organization that will deal with federation level issues, we will refer to it as the GENI Oversight Group (GOP) throughout the rest of this document. Keep in mind though, that the GOP does not yet exist and its duties as discussed here may be fulfilled by the GPO in the interim.*

GENI-CSIRT Responsibilities

The GENI-CSIRT will be responsible for resolving all help tickets related to security. The details of the incident response process for investigating such events will be described later in more detail. Because of the decentralized nature of GENI, much of this effort will be developing security guidelines, coordinating between organizations, verifying security problems have been resolved before bringing equipment back online and making sure issues are followed-up upon and resolved. It will usually be the individual sites hosting aggregates who will physically investigate

systems and restore them to a clean and safe state, though the GENI-CSIRT will have a role in verification. For centrally owned GENI resources, such as those at the GMOC, the GENI-CSIRT may be involved more hands-on, (e.g. hardening hosts, providing monitoring and developing security architectures).

The GENI-CSIRT will also have longer term responsibilities of planning and auditing. They will be responsible for maintaining and updating security agreements, auditing compliance with agreements and providing appropriate security training and education to the GENI community. Lastly, if GENI is to run a Certificate Authority (CA), the GENI-CSIRT will at a minimum draft recommended policies regarding its administration and configuration and need the ability to revoke compromised certificates. Though, actual day-to-day system administration of the CA would not likely be their responsibility.

Lastly, the GSO may also be the GENI Legal, Law Enforcement and Regulatory point of contact, as described in the *Responding to Inter-organizational Legal, Law Enforcement & Regulatory Requests in the GENI Federation MOU*.

Aggregate Authorities and Operators

Those providing resources to GENI, the Aggregate Authorities (AAs), have responsibilities to **prevent**, **report** and **respond**. These 3 categories of responsibility are spelled out more specifically in the *GENI Aggregate Provider Agreement* to be signed by AAs hosting "GENI-approved" aggregates.

Prevention

For major resources, especially those common across many campuses (e.g. GENI Racks, WiMAX base stations and OpenFlow deployments), the GENI-CSIRT may provide specific hardening guidelines. However, it is always the responsibility of the hosting organization to protect their own aggregates and components. Therefore, they should keep systems up-to-date and patched against security threats, run the minimal number of services, keep good logs with accurate timestamps, give administrative access to only those who need it, and use strong authentication for administrative accounts. This is important not only as their components have trusted connections to others, but credentials and certificates from GENI users will likely exist, if only ephemeral, on their systems. Aggregate operators have a responsibility of due diligence to protect those credentials and experimental data.

Reporting

It is expected for those providing GENI resources to report when there are incidents on their equipment, if that equipment is providing a GENI resource. This follows as the problems may affect other parts of the GENI infrastructure, but also it will help the GENI-CSIRT see the bigger picture and trends. For example, an incident at one site may just be the first of a new exploit being tried out on many other pieces of GENI infrastructure or a piece of a larger attack.

Reporting a problem to the GENI-CSIRT (through the help desk at the GMOC most likely) does not preclude a site from reporting it to their own security team or CERTs. In fact, it is expected that security teams at their own institutions will lead the investigation on their side. What AAs are precluded from doing is sharing information they may learn about a similar incident at another AA (usually in the course of an active investigation) without permission from the other AA. The privacy of federation members must be respected in the course of investigating incidents, and both the GENI-CSIRT and the AAs have a responsibility to protect that privacy.

Response

If one is maintaining resources in the GENI ecosystem, they need to be responsive to inquiries from the GENI-CSIRT or GMOC, as noted in the Aggregate Provider Agreement. Response here may just mean acknowledging an issue, or it could mean shutting down slices, slivers, components or whole aggregates temporarily. It could also mean demonstrating that a host has been cleaned-up and properly recovered.

Part of being responsive is providing up-to-date contact information for someone who is responsive, knowledgeable of the aggregate's operation and able to shutdown that aggregate or block an offending component. This, too, is noted in the Aggregate Provider Agreement.

GENI Oversight Group (GOP)

The GOP, through agreements with federation members, gives the GENI-CSIRT the necessary authority to do their job. They provide oversight and clarify the requirements for security and operations. The GOP is also the connection back to the NSF who will expect to hear about major incidents from the GOP before any side-channels. To achieve all of this, there needs to be a clear escalation path from the GSO to a specific person at the GOP up to the relevant NSF program officers.

The GOP representative is also the only one who is to talk to the media about a GENI incident, and they will do so protecting the privacy of individual federation members. Of course, individual organizations can talk about their own site in isolation, but if they are talking about an incident involving a GENI resource they must not make mention of GENI or anything they learn about the incident at other sites.

Other Members

As mentioned, there are numerous stakeholders and most any of them could be needed to help with an investigation. Security is a collaborative effort, especially in GENI. For example, an incident involving an exploit of control framework code would require the help of one of the developers of the control framework to be properly addressed. Therefore, it is possible for any large incident that a special team may have to be formed from a diverse group of GENI participants.

Incident Classification

There are three classifications that will be used for security incidents in GENI: *high*, *medium*, and *low*.

High

The incident could lead to exploitation of the trust fabric, i.e. user and host identities; or the incident could lead to instability or misuse of the overall GENI ecosystem. This type of incident will require the GSO to assign an investigation lead and likely require members external to the GENI-CSIRT to be pulled into the investigation. This type of incident must also be reported to the GOP immediately and to the NSF from there.

Medium

The incident affects an instance of a GENI service, but GENI stability is not at risk; or a denial-of-service affects one replica of a given GENI service; or a local attack compromised a privileged user account. This type of incident may require an investigation lead if wide-spread enough. It must also be reported to the GOP immediately.

Low

A local attack compromised an individual user, non-privileged credentials exposed; or a denial-of-service attack or compromise affects only local GENI resources. This type of incident does not require an investigation lead. It also need not be reported to the GOP out of the normal reporting schedule.

Information Classes

Different types of information that the GENI-CSIRT may come across during an incident investigation have different levels of sensitivity. Accordingly, we define 3 information classifications below.

Sensitive Personal Information

This type of information has challenging security requirements and is incidental to planning, provisioning or using GENI. Examples are user names, passwords, social security numbers (SSNs), and credit card numbers.

Some of these, like SSNs and credit card numbers, are legally protected. While users/experimenters should never intentionally store legally protected personal identifying information on GENI systems, it is always possible that such information could end up on an aggregate inadvertently due to the activities of a cyber criminal or by an experiment accidentally harvesting such data from places like torrent sites.

GENI Restricted Data

This is any non-public GENI information that is not personally identifying. Most information generated as part of an incident investigation will be of this type until the final, public report is made.

GENI Public Data

This is data that has no privacy requirements. Data that is not explicitly classified is presumed to be public.

Data Classification Process

The GSO will determine the classification of any information from a security incident based upon the definitions above. These same definitions could (with some modification) be used by other groups in the GENI community, but it is up to the GOP to determine whether these information classifications should be standardized across domains.

It is also expected that most everything done in GENI will be public, and that should be the default level for information not otherwise classified. The exception is that the default for an *active* GENI security incident is *GENI Restricted*. This is important to encourage cooperation and information sharing that is critical to prompt incident response.

Incident Response Process

The incident response process used to investigate all GENI security incidents has 7 steps: *Discovery and Reporting*, *Initial Analysis &*

Discovery & Reporting

Incidents may be discovered through a variety of means including reports from users, system administrators, engineers, and peers; operations center monitoring of infrastructure, services, and resources; and through monitoring of intelligence channels.

When an incident is discovered that relates to GENI resources, services or identities, it should be reported to the local institution's incident handling process, and the discovering/reporting party needs to report the incident to the GMOC or GENI-CSIRT directly. This can be done through special *security* or *abuse* email lists, which must be advertised to all GENI federation members and partners. The help desk, upon receiving such a report, must then create a ticket which will alert the GENI-CSIRT to the new incident. Secure internal communication can then be made through the ticketing system.

Initial Analysis & Classification

It is critical to understand the scope and severity of an incident. Therefore, as a new ticket is received, the GSO will evaluate the incident and classify its severity by using the 3 definitions discussed above for high, medium and low severity incidents.

Classification will likely require some discussion between the reporter, the GENI-CSIRT and the sites involved. Also, the classification may change as more is learned about an incident. Still, it is important to try to classify the incident early to know the amount of attention to give it and whether or not it must be reported to higher levels.

The second goal of this step is to create the right team if necessary (it may only be necessary to decide which GENI-CSIRT member to assign a *low* severity incident to). For example, if it is clear that a GENI software or firmware vulnerability is being exploited in an attack, then team being setup for this investigation needs to include GENI developers as well as some sites who will beta test the patch or remediation solution.

Containment

As a general matter, the level of response must take into account a number of factors, including, but not limited to:

- whether the resource/service has been compromised or is it just under attack?
- the kind of attack — DOS or user or privileged user compromise?
- the importance of the resource/service locally?
- the importance of the resource/service to the operation of the whole GENI test bed?
- the importance of the resource/service to various PIs and projects?

As an operational principle for an aggregate *under* attack from another GENI resource, the normal response should be to block access from GENI during the initial stages of dealing with an intrusion – only opening access as is prudent and justified, without extraordinary risk. Of course, AAs need to inform the GMOC of actions they take affecting availability GENI resources/services, as they have agreed to advertise their services as accurately as possible in the Aggregate Provider Agreement.

If it is determined that the attack comes from a specific user account or slice, the more reasonable response would be to temporarily deny access to the resource or service through the appropriate local control for authorization. Again, such action should be reported back to the GENI-CSIRT who will follow up with the appropriate project leader corresponding to the offending slice or user.

We must also consider what to do about an aggregate from which an attack is *launched*. A problematic aggregate or service might be reported by another GENI AA, an aggregate operator or ISP on another test bed federated with GENI, or it might be discovered by generic infrastructure monitoring capabilities.

While one might expect that the GMOC or GENI-CSIRT would have the ability to block a site or service that was misbehaving, which might be true in cases for specific centrally controlled services, it will not be true for the vast majority of services on GENI nor will it be true for federated GENI-like testbeds that have their own operations centers.

Depending on the severity of the attack and based on the other aggregates and slices potentially affected, the GENI-CSIRT will attempt to notify the affected aggregate operators so they can take appropriate action to protect themselves. The GENI-CSIRT may ask the offending site to disable specific slivers, whole components or in extreme cases (when an emergency stop is required) disconnect aggregates from the GENI backbone.

The second phase of containment is the process of narrowing down the things that are blocked to the specific components, resources, services and users which were compromised. Incident response teams or aggregate operators for the AAs, in communication with their peers, are expected to restore normal operations as quickly as the problem areas can be identified and isolated.

Notification

Any incident requiring a team leader **MUST** be reported to the GOP. Furthermore, a report on the incident will be sent to the GOP upon resolution of the incident.

Any incident rated as *high*, must be reported back to the appropriate NSF program officers through the GOP. It is important that the NSF first hear about an incident though the GOP rather than the media.

Analysis & Response

This is the step where an in-depth analysis is done and a path to resolution is established.

Resource Tracking

Since the total cost of the incident is often important for legal action, the GENI-CSIRT and any aggregate owners involved should track:

- the number of man-hours spent on response;
- the extent of the damage;
- what resources were made unavailable; and
- how long said resources were offline.

This is also a good time to assign attributes to incidents for categorization in annual reports. In addition to just the severity of an incident, it is useful to track the organizations involved, the type of exploits used, and the types of vulnerabilities used for entry for future planning.

Evidence Collection

All supporting data for an incident should be treated consistent with the AA's institutional rules for maintaining and storing such materials in non-GENI incidents. Collection and coordination of evidence between aggregate operators is expected to be handled by the appropriate law enforcement agencies, and not the GENI-CSIRT.

Removal & Recovery

Determine the extent of known and potential compromise of user and host credentials and passwords. Did the initial containment step treat the entire scope of the compromise? Work with contacts to revoke/suspend credentials, keys and passwords.

Point of entry should also be established, and if a vulnerability was exploited to gain access, that vulnerability must be patched. If this vulnerability is likely to affect other aggregates, they must be notified by the GENI-CSIRT. If applicable, any IDS or log monitoring system should have rules updated to detect this threat automatically, if possible, in the future.

Any malware installed must be removed, and components must be shown to be clean before the GENI-CSIRT verifies the reconnection of "GENI-approved" aggregates.

Post-Incident Analysis

At the end of an incident, the investigation lead (if applicable) schedules a conference call to review the lessons learned and formulate feedback to appropriate groups (e.g. GENI participants, management, developers). A close-out report must be completed within 1 month following the incident. This report will be given to the GOP but also made public in a sanitized form to protect the identity of individual persons and organizations. This will help prevent the same mistakes from being repeated.

Auditing

Any security plans developed and monitoring infrastructure rolled out by the GENI-CSIRT needs to be audited and evaluated regularly. Reviewing security plans is especially important as the GENI ecosystem is very dynamic. Therefore, security plans must be reviewed and updated annually.

The audit schedule for the security controls (whether preventative, detective or reactive) must be developed as such infrastructure is rolled out. In addition to services managed directly by the GENI-CSIRT, they will also want to regularly audit the GENI-specific authentication and authorization systems.

GENI will be relying upon services provided by many parties, and it is prudent to test the security of the major aggregates. The GENI-CSIRT should, with approval from appropriate AAs, randomly test the security of major aggregates, developing reports and recommendations from their findings. Furthermore, the GENI-CSIRT should perform exercises to test incident response procedures ahead of time and before a real major incident occurs. Exercises involving several different parties will help to discover cracks in the system and potential problems or bottlenecks.

Recommended Threat Countermeasures

This final section of the security plan recommends several countermeasures intended to control risks that we have identified. Many of the threats being addressed are from an initial threat and risk assessment performed during Spiral 2 and is hence heavily weighted towards the OpenFlow and WiMAX build-outs. As this document evolves, and GENI evolves as well, it will need to be updated and will likely become more balanced.

Several recommendations are made. Some require funding of new development efforts, and some require new operational services to be implemented. However, many of these can be addressed without additional funding by the groups developing software and systems for GENI, if they take these security implications into consideration during the early architecting and prototyping phases.

High Priority Threat Mitigations

Here we discuss policies, procedures and technologies that must be implemented to protect GENI's assets in the near term.

Official Project Leader and Experimenter Agreements

While there is a GENI MOU that discusses some of the expectations of the users, in particular for the PIs or slice owners, there is no formal agreement. This opens up GENI to liabilities that it most likely does not want. In developing other agreements, it has become clear that user agreements are needed which assert that at a minimum slice owners and project leaders:

- take responsibility for the behavior of their slices/projects;
- will not knowingly perform any illegal activities with GENI resources;
- will not knowingly interfere with the experiments of others;
- will respond to inquiries about problems with the behavior of their slice/project within a reasonable timeframe;
- make opt-in users aware that GENI takes no responsibility for the software they install on the opt-in user's devices;
- provide contact info to opt-in users that are installing software on their hosts;
- registers any software installed on opt-in user devices with the GENI LLR representative; and
- are aware that GENI makes no guarantees about privacy and is not responsible for the loss of any of the experimenter's data.

Communication between aggregate owners and GMOC

As the GMOC developed the Emergency Stop Procedure, they correctly realized the importance of having up-to-date contact information for each AA. They have been collecting this information, and the need for it has been called out in the Aggregate Provider Agreement. In addition to having contact info for the person in charge of an aggregate, it would also be useful for AAs or aggregate operators to provide a security contact at their institution. This would be the party to whom they report a local security incident. This information is useful to have as the GENI-CSIRT may want to make the local institution's security team aware of activities when coordinating an investigation. Furthermore, the aggregate operator will probably not be a 24/7 on-call person, but someone from the security team at the aggregate provider's home institution is probably on call 24/7 for emergencies.

In addition to having contact info so that the GMOC can reach an aggregate, the aggregate needs contact info for the GMOC, both via phone and email. It is important for the GMOC to have a callback number known to all aggregates so that they can verify the GMOC's call, should someone try to impersonate the GMOC with a fake request. **The GMOC should provide such a number to aggregates as they collect contact information from them.** Furthermore, the GMOC and GENI-CSIRT should publish their GPG public keys on their web pages and use them to sign email messages.

Law Enforcement Requests

The GENI federation needs to have a plan and process in place for handling Law Enforcement, Legal and Regulatory (LLR) requests. At a minimum, GENI will likely have to deal with DMCA threats, much like PlanetLab and other networking testbeds have. Therefore, an LLR plan spelling out the responsibilities of different parties and the process for forwarding requests from LLR entities to the appropriate places was drafted during Spiral 3.

A key component of this LLR plan is the establishment of a GENI LLR facilitator to help parties reroute requests. This person could be on the GOP, or they could be delegated by the GOP to either a GMOC or GENI-CSIRT representative. Regardless, this position should be established by Spiral 4, and all federation members must be aware of this person and how to contact them.

Log Retention

If there is a security incident on GENI infrastructure, it will be important to have good logs at the aggregates for both their internal investigation and any incident analysis done with the assistance of the GENI-CSIRT. As stated in the Aggregate Provider Agreement, AAs or aggregate operators should follow their own institutional policies for log retention, retaining security relevant logs for GENI components as long as they can. It is just as important that logs have accurate timestamps and that they are archived with the relevant time zone information to make cross-correlation possible.

Of course, aggregate logs are not the only ones to consider. It is just as important that the clearinghouse and the GMOC maintain detailed logs as long as possible. Since clearinghouse will provide a slice registry service, it will have logs that are very important for mapping slices to slice owners. Lastly, if the GENI-CSIRT ever deploys a security monitoring infrastructure, they will need to take care to store accurate logs.

Hardening Guidelines for GENI aggregates

While there is a long tail of small, heterogeneous aggregates, a large proportion of shared resources are provided by major build-outs, such as, OpenFlow, WiMAX, ShadowNet, and GENI Racks. These core pieces of infrastructure can and should be configured and hardened in a standard way. While not feasible to develop hardening guidelines for all resources, recommendations could probably be developed for these.

Since any hardening guidelines require intimate knowledge of these resources, each development team should have an involved security consultant who is familiar with the equipment and software in question. For example, the team that develops the WiMAX kits that other groups deploy needs a security expert who will advise on the architecture, development and deployment phases. That way they can make these aggregates more secure out of the gate as opposed to retrofitting security considerations with just hardening guidelines after the fact. While this may not be possible for more mature projects, calls for proposals for new infrastructure should require such a security component to the project which must also be adequately funded.

Because we do not have access to or intimate knowledge of these resources beyond what can be browsed on the GENI wiki, we can only speak in generalities about hardening these hosts. Furthermore, these devices are under active development and any low-level guidelines would quickly become out of date. Therefore, specific configuration guidelines should come from the teams developing these kits as they are being deployed, and they should be regularly updated as software and configurations change.

With this caveat, we recommend the following general considerations for developing hardening guidelines:

- A separate interface should be used for administration of these hosts as opposed to the network interface used for components to reach others within a slice.
- Administrative accounts should use two-factor authentication.
- Any local accounts on these hosts should use unique passwords, not shared across other unrelated devices or hosts.
- Unnecessary services and accounts should be disabled.
- Security patches need to be kept up-to-date. The developer of the aggregates are in the best position to monitor for updates to the software and should utilize email lists to notify the appropriate aggregate operators when they need to update software.
- The number of setuid programs should be minimized as they can be used for privilege escalation along with new zero-day exploits.
- File integrity checkers should be utilized to detect new software installed or changes to configuration files. This would be especially important on infrastructure serving both GENI experiments and production services at the host institution, such as OpenFlow. For example, one would want to know immediately if any files that control the separation of research from production traffic were altered in FlowVisor.
- Most of these aggregates utilize virtualization to separate one physical unit into multiple logical components. To maintain the security, stability privacy and integrity of individual experiments, this virtualization should strongly compartmentalize different slices. Therefore, consideration to virtualization technologies used is very important (e.g., UML vs. Xen), and tests should be run to determine the degree to which a malicious slice could affect others.
- A secure configuration checklist should be utilized in situations where configurations are likely homogenous between deployments (e.g., WiMAX kits).

Hardening Critical Infrastructure

No less important than the health of the aggregates, we must consider critical infrastructure without which experiments are threatened. Therefore, we recommend that several services (listed below) be hardened and monitored by the GENI-CSIRT. Additionally, these should be made redundant to make them resistant to denial of service or simple failure. It is expected that these services will be run by the GMOC, but even if they are not, it is not unreasonable for GENI to expect their security team to monitor these services. In either case, it will be important to have the GENI-CSIRT work early on with the providers of these services to decide how best to secure them.

- Clearing House(s) and Certificate Authorities; (There should be agreements between trusted CAs about their operation such as the Grid community uses)
- Update Services for common aggregate software and GENI framework software;
- Logging hosts that aggregates may submit to for centralized monitoring of availability/status, performance or security purposes; and
- Backup servers used for experiment templates and sunseting experiments.

Education and Policy Adoption

It is vital in a federated community like GENI to get buy-in from all the major stakeholders for any security policies or plans. This can be achieved by allowing all stakeholders to provide input into the process. The regular GECs are one of the only opportunities to bring people together to discuss these issues, and we should use the GECs to gain consensus on all security plans and policies, especially when users, AAs or developers are being asked to take action.

Education of security policies and procedures are a major component of any good security plan. Simply posting documents on the web and having AAs, users and others acknowledge that they have read them is not enough. The GECs should also be used as a venue to disseminate this information, as broadly as possible to all relevant parties. A security track at these meetings could be used to achieve this, but that would give too many the opportunity to miss the presentations. Rather, it is probably better to integrate security into discussions of existing working groups, like we have done with COMIS.

The GOP must make clear the importance of security to the project. Education is part of that, but it is also important to provide incentives for buy-in to security policies as well as to integrate security across all projects. The latter can be done through CFPs if security is a funded, non-optional part of most large projects.

Medium Priority Threat Mitigations

These policies, procedures, and technologies should be implemented within the Spiral 4 to protect GENI's assets.

Automated Emergency Stop

An emergency stop procedure has been created by the GMOC, which will be used to stop a badly misbehaving or out of control slice. For example, an experiment may be co-opted for a DoS attack or some other illegal activity. Alternatively, a slice may not be attacking or actively performing any illegal activity, but due to misconfiguration it could threaten the stability of GENI and other experiments.

The problem of course with the current mechanisms is that they depend upon the rapid collaboration of several people being coordinated over phone and email. There are no guarantees about the effectiveness or timeliness of an emergency shut down. If functionality were built into the control framework, then that could be used shutdown/pause a slice rapidly or at least isolate offending components.

Aggregate Authorities may object to an interface that would allow slivers (though not whole components) to be shut-off. While such concerns

about someone exploiting the GENI control framework to DoS a specific slice are probably unwarranted, a compromise could be made by only providing the interfaces needed to surgically disconnect slivers from other parts of GENI. Otherwise, in an emergency, it is more likely that whole aggregates would be blocked.

Secure Software/Firmware Update System

As security vulnerabilities are found on common GENI infrastructure (e.g. OpenFlow firmware, control framework software, etc) and patches become available, Aggregate operators for the AAs must be notified and have a simple way to update their systems. The GMOC could provide an email list service to which developers would publish these announcements and aggregate operators would subscribe.

However, notification is half the issue. Updates should be simple and automated when possible, and there should be a way to get them in a secure manner where their integrity can be verified. One GENI project is working on development of such a system. However, even if they are successful, someone must fund its wide-scale implementation, maintenance and operation. If shared infrastructure operations is to be maintained through the GMOC, it makes sense for them to run this service.

Secure communication for the GENI-CSIRT

The GENI-CSIRT will need a way for its team members to communicate securely. This could be through a help desk ticket system, though they will likely need more than that. They will need a wiki to collaborate on documents/reports as well as some sort of content repository to store things like exploit code and malware discovered. If this content repository could utilize WebDAV in a way that they could link from the wiki to specific files, that would be ideal, but other secure network file systems like AFS could be used.

If internal email lists are used, they must be private and unlisted. Furthermore, encryption through PGP or some other means should be used. Forcing communication through a secure ticketing system would likely be simpler than trying to secure group email, though.

Relatedly, the GENI-CSIRT must have a way to communicate securely with the aggregates, PIs and other users. Therefore, they must have both a publicly available callback number and a GPG key for official emails.

Log Correlation

For commonly deployed and large pieces of infrastructure, it would be useful for the GENI-CSIRT to have forwarded logs to correlate and analyze for problems. While it is unlikely that forwarding could be mandated, there should be a central logging infrastructure maintained by the GMOC or GENI-CSIRT for those AAs that can and would like to forward their system and security relevant logs.

Such an effort should start with a single build-out, like OpenFlow, and begin by determining the useful types of logs that could be gathered and horizontally correlated. In the case of OpenFlow, it would be very useful to map new OpenFlow rules to slices or users. Then the infrastructure could be put in place to consume and analyze this type of data. If successful, which will depend on both the type of data and adoption rates, then they should move onto provide security monitoring for other aggregate types.

Hardening and Scanning for Secondary & Tertiary Services

There are several services, probably to be centrally provided, that are not critical to the continuation of experiments. While their interruption would be annoying, they are probably not critical enough to require redundancy. Still, the providers of these services (presumed to be the GMOC at this time) should follow standard best practices and be regularly scanned/audited by the GENI-CSIRT for vulnerabilities. These services include, but are not limited to:

- GENI web site;
- Documentation resources if separate from above;
- Experiment control portals;
- Clearinghouse; and
- collaboration tools (e.g., email list servers, wikis, and code repositories)

Private Key Management

While SSH public keys and certificates have overcome many issues with passwords in the grid community, they still suffer from the problems associated with users managing these credentials insecurely. In particular, users tend to store the associated private keys in the clear on their workstations and many remote hosts. This makes them very susceptible to harvesting attacks.

Such issues can be mitigated in GENI by removing control of managing these credentials from the user. However, this must usually be architected into the authentication system early on with proxies being setup to manage credentials. It is hard to implement after the fact without hacking protocols or dramatically altering user behavior.

In particular, this is handled poorly in ProtoGENI. They utilize python scripts to store unencrypted private keys in the filesystem for ease of use, and users must remember to run a clean-up script after they are done or this key remains. First, it should never be written to disk. Second, clean-up scripts should be automatic. The more elegant solution would take something like ssh-agent or gpg-agent and modify it to store private keys for certificates in memory as well.

It is unclear how well ABAC implementations would mitigate these issues. If ABAC avoids them, it should be mandatory. If one can always fall back to the insecure and easier way to do something that they have always used, they will. Therefore, we suggest that as GENI's authentication and authorization system is being developed that consideration of this type of attack be given along each step of the way.

Low Priority Threat Mitigations

These policies, procedures, and technologies should be considered and debated within the community as to their utility. We are not recommending their implementation at this point. The costs of many of these are unclear, and those costs would need to be determined to decide if these countermeasures are worth implementing, as well as funding sources determined. Others may cost little but still need buy-in from many stakeholders to implement.

Code Audits

Important and trusted software/firmware needs a high level of assurance. Unfortunately code prototyped and developed by academic research groups often have a poor track record of security, and without an external code review for security vulnerabilities, we cannot justifiably put much trust in our critical software components. Therefore, we recommend a professional code review of the of any widely deployed software platforms developed in-house (e.g., OpenFlow firmware, FlowVisor, control & measurement frameworks, authN/Z software, etc). This could be done as a new project activity in Spiral 4 if called out in the new CFP. Alternatively, centers like MIST ([Middleware Security & Testing](#)) at the University of Wisconsin-Madison could be funded to perform such activities.

WiMAX encryption for experimenters

Just as WiFi has 802.11i, WiMAX has PKM (Private Key Management). Using PKMv2 to manage keys along AES encryption provides robust protection against eavesdropping and integrity checking that will detect any packet manipulation. While an experimenter should in theory be able to implement PKM themselves, it is wasted effort for each experimenter to re-invent the wheel and figure out a way to do this.

It would be more ideal to provide tools that allow an experimenter using the WiMAX base station to simply and transparently turn on encryption in a single config file. This would encourage more experimenters to use encryption for purposes such as protecting the privacy of opt-in users. Ideally, this would be implemented by those developing the WiMAX kits.

GENI Intrusion Detection System

As GENI grows, in the future it would be helpful to utilize a distributed IDS to detect threats that could not be noticed easily with the perspective of an individual aggregate. Such an IDS would likely be administered by the GENI-CSIRT. The appropriate configuration and architecture is not obvious at this early point in GENI's evolution, though the afore mentioned log correlation engine could be a starting point. The obvious next step would be to add distributed network traffic sensors that would feed into a centralized event engine.

Development of such an IDS would likely require an entire GENI project devoted to its creation. Further, the GENI-CSIRT would probably need some operational experience with GENI to develop the requirements for such an IDS.

Glossary

- **Aggregate Authority:** is responsible for some subset of substrate components: providing operational stability for those components, ensuring the components behave according to acceptable use policies, and executing the resource allocation wishes of the component owner.
- **Component:** encapsulates a collection of resources, including physical resources (e.g., CPU, memory, disk, bandwidth) logical resources (e.g., file descriptors, port numbers), and synthetic resources (e.g., packet forwarding fast paths).
- **Aggregate:** is a collection of components under the authority of a single management authority
- **Aggregate Manager:** a service that exports a well-defined remotely accessible control framework interface to an aggregate.
- **Aggregate Operator:** is appointed by the Aggregate Authority to operate the Aggregate Manager and any components of the aggregate. This may be a the MA itself, or someone from another organization operating on its behalf.