CSD : GENI Security Use Cases & Stakeholders

This page last changed on Jan 15, 2010 by slagell.

This document has two main objectives. First, it identifies key stakeholders in GENI's success and how they are affected by the security or insecurity of GENI. The second objective is to generate uses cases for discussion.

These use cases are meant to generate discussion about the types of security incidents GENI will have to prepare for when it goes into the operational phase of Spiral 3. In some sense it may be better to call these incident response scenarios. The discussion here will allow us to dig deeper into a more technical discussion about specific threats against GENI, a part of our next milestone. Some of these threats fall out rather directly from these scenarios, but there is not a strict one-to-one correspondence between threats and the scenarios that follow.

We have tried to break these use cases into a few categories. First, we discuss generic, GENI agnostic incidents that are not specific to GENI's implementation in any significant way. Second, we discuss those incidents targeted at GENI specifically by external parties. Third, we consider incidents that are completely internal to GENI, between internal actors. Last, we consider GENI itself being used to launch attacks, where it is not just the target. These are not completely mutually exclusive categories, and hence some very similar scenarios can appear in multiple categories.

Related Documents

These related documents all contributed to the development of this document, some of them providing existing use cases and others an understanding of GENI's organization and structure.

Document ID	Document Title and Issue Date
GENISESYSO02.0	"GENI System Overview", September 29, 2008. http://www.geni.net/docs/GENISysOvrw092908.pdf
GDD 06-23	"GENI Facility Security," by Thomas Anderson and Michael Reiter, GENI Design Document 0623, Distributed Services Working Group, September 2006. http://groups.geni.net/geni/attachment/wiki/GENISecurity/GDD-06-23.pdf
GDD 0610	"Towards Operational Security for GENI," by Jim Basney, Roy Campbell, Himanshu Khurana, Von Welch, GENI Design Document 0610, July 2006. http://groups.geni.net/geni/attachment/wiki/OldGPGDesignDocuments/GDD-06-10.pdf
GENI-FAC-PRO-S1-OV-1.12	"GENI Spiral 1 Overview", September 29, 2008 http://groups.geni.net/geni/attachment/wiki/SpiralOne/GENIS1Ovrvw092908.pdf
GENI-SE-SY-TS-UC-LC-01.2	"Lifecycle of a GENI Experiment", April 30, 2009 http://groups.geni.net/geni/attachment/wiki/ExperimentLifecycleDocument/ExperimentLifeCycle-v01.2.pdf
GENISECARCH0.55	"GENI Security Architecture", July 31, 2009 http://groups.geni.net/geni/attachment/wiki/GENISecurity/GENI-SEC-ARCH-0.55.pdf

Stakeholders in GENI's Security

This is a list of stakeholders who have an interest in GENI's security. We have tried to discuss the particular interest or "stake" that each of these parties have. There is no significance to the ordering.

National Science Foundation

The NSF has a large stake in GENI's success being the primary means by which GENI is funded. They not only want to provide the best and most stable environment for researchers whom they will later fund to utilize this infrastructure, but they are also concerned about image and reputation. The NSF has to report to the tax payers most directly through Congress, and an embarrassing security incident (e.g., if GENI was used to attack federal resources and this made the news cycle) could harm their reputation and future finances.

The Network Science and Engineering (NetSE) Council

The NetSE council has a research agenda for the nation, and GENI is a big part of achieving their vision. The availability of GENI is very important to them as well as the ability of researchers to trust GENI. Trust in GENI is especially important if GENI is to be federated and NSF GENI connected with similar testbeds worldwide---a part of their vision.

Aggregate / Component Owners

The researchers and students running specific components that serve as GENI resources will have varying degrees of interest in the security of these components. They may be very interested if they will depend upon some unique resource they are sharing and also need themselves, but they may be more interested in GENI as a whole from a researcher perspective.

The incident response teams and security admins at the various institutions (e.g., universities) hosting these resources will be very interested in the security of GENI, but not so much from an availability perspective that the researchers would care about. These security teams would be more interested that GENI resources do not introduce new vulnerabilities into their infrastructure and that they would not have to spend significant time dealing with false positive alerts created by the GENI devices. These security teams would also like to know that the GENI resources have people maintaining them who could be quickly contacted if there is a security incident.

National Lambda Rail & Internet2

Both of these institutions are dedicating significant network resources to GENI, namely, large data pipes. They are not contributing "components" to be broken into slivers for GENI. They are likely concerned about GENI traffic affecting other traffic, though they can probably provide physical isolation that prevents GENI from over-consuming resources. Still, if they need to install special GENI software or firmware for measurement or other purposes, vulnerabilities in such software could affect their other operations.

Federation Partners

GENI is envisioned with at least two layers of federation. At the first level there will be an NSF-funded clearing house(s) connecting resources for NSF funded researchers. At a layer higher, it is expected that this "NSF GENI" will connect to other GENI-like testbeds from other nations. It is this latter type of federation partner we are discussing in this section.

Such partners are definitely interested in GENI security because experiments will cross such boundaries if the federation is to be useful. It is almost certain they will need to trust GENI's authentication mechanisms as it would be untenable to collaborate if they did not delegate the authentication of NSF GENI researchers. They will also be making explicit connections between resources and need to trust that GENI does not open up new vulnerabilities that would exploit special trust relationships over these connections. Lastly, they must trust NSF funded researchers to run experiments, and hence code, on their systems. Therefore, they are likely to be concerned that GENI has appropriate procedures and mechanisms to handle users that behave inappropriately.

Researchers

These are the PI's, staff and students running experiments on GENI. They may or may not also be aggregate or component owners. As researchers, they are most concerned with security in how it affects

availability of resources. Systems that are constantly taken offline are of little use to them. This includes a concern of runaway experiments that over-consume resources. Second, they may be very concerned about privacy of their experiments. This is especially likely for those working with the private sector on development of new telecom or networking products. Lastly, it is very important to researchers that other experiments are isolated such that they do not affect the validity of their own experiments.

Opt-in Users

Opt-in users from the Internet, whose traffic could be part of an experiment, probably have little interest in GENI. If GENI increased their exposure to some sort of threat, they would likely no longer opt in. They have little to lose by not using GENI. One of the more likely concerns opt-in users would have is about privacy. For example, many would not use an experimental GENI ISP service if they thought it would track all their behaviors online. Furthermore, if GENI was insecure and it allowed people other than the known experimenters to watch their behaviors, they would likely be concerned.

GENI Management & Operations Center (GMOC)

Much of the role of incident response, especially coordination between sites, will depend on the GMOC as it is an organizational bridge between the aggregate owners. Security is thus a very high concern for them. Even if a separate group or just a subgroup of the GMOC is involved with incident response, security is important from the aspect of providing availability to services. Therefore, security impacts the ability of the GMOC to do their job in a fundamental way.

BBN Technologies

BBN's concern for the security of GENI will be similar to that of NSF's. They are a partner with NSF and have a vested interest in the success of GENI. While they do not have to report to Congress, they do report to NSF who reports to Congress. So embarrassing failures affect them, their image and their funding as well.

Tax Payers

Most tax payers will not be directly aware of GENI, and their interests are supposed to be represented by Congress which funds the NSF. Most directly then it is the NSF that represents their interests. Their interest is that GENI is not a waste of money, which requires its success; and security is necessary, but not sufficient, to the success of GENI.

Operational Security Use Cases

Non-specific Use Cases

The following use cases describe issues incident response teams will have to deal with when involved in any federated project such as GENI with shared computational resources.

Malware, Worms & Botnets

Even with no opt-in experiments running, GENI will have several resources exposed to the Internet. This includes the GMOC systems, aggregate managers, clearing houses and hosts running the virtual machines and virtual routers. All of these will be exposed to worms and other forms of malware common on the Internet. They will also be targeted with the same types of brute force attacks against common services (e.g., SSHD brute force password guessing) that all machines on the Internet face. Furthermore, some of these attacks will almost certainly be successful, and GENI resources will become parts of botnets.

There is nothing about the GENI architecture that makes success of these attacks more likely than for many other machines. However, the organizational and management structure of GENI means that there will be a long tail of resources that are more loosely managed. It is expected that many universities and institutions will contribute a small number of components likely managed by graduate students. Experience with hosting Planet Lab machines and similar resources has shown us that these machines are usually not hardened, kept up to date or closely monitored. Often these machines get infected, and the person who was running the host is not even associated with the university or research group anymore.

GENI's operational security team will need to be responsive and proactive to minimize this problem. Otherwise GENI's reputation will be at risk, and it will become more difficult to convince universities to host these resources. Coordinating the cleanup and recovery from these incidents could easily consume most of the GENI incident response team's time. Therefore, it is to their benefit to be proactive and detect these sorts of problems before the security teams at the hosting institutions discover them. Even more beneficial would be education and standardization in the management of these resources to reduce incidents.

It is worth noting that if an experiment has opt-in users, this greatly increases the chance of virtual resources within a slice being compromised in these ways. In that case, the GENI incident response team will likely have to disable the experiment. They will certainly need to be in communication with the experiment's PI and possibly the aggregate and component owners. It would be useful if there was a way that the incident response team could freeze the state of an experiment before turning it off. However, they may need to use the emergency stop mechanism depending upon the seriousness of the incident.

Credential Harvesting

One of the most common problems incident response teams deal with in federated systems (e.g., the grid community) is credential compromise and re-use. In the case of grid resources, the attackers are usually not even using grid-aware techniques (e.g., they don't know about Globus or how to submit jobs). Similarly, we expect attacks by people who don't know anything about running GENI experiments. They will just exploit exisiting trust relationships and the fact that the same credential (often a password or private key) is used at multiple sites and enter hosts through commodity technologies they do understand (e.g., SSHD on Linux hosts).

Unless One Time Passwords (OTP) are used everywhere in GENI, this tactic of harvesting credentials in an ever expanding network is likely to affect GENI. When such incidents occur, the GENI security team will likely be involved if for no other reason than to provide coordination between sites. If GENI specific credentials are harvested, the security team will need to revoke them.

Requests for Private Data

Another issue that we frequently deal with at the NCSA are requests from PI's or professors to access data from experiments run by a student who has left the University. Straightforward privacy policies should be in place so that these kinds of requests can be handled efficiently. A clear process for handling these requests will reduce the load on GENI operations.

This may or may not be a problem depending on how accounts are allocated in GENI. If a PI must be the owner of a slice and students can only be delegated rights, then we can largely eliminate this issue.

Illicit Materials Hosted on a GENI Resource

Illegally copied data/software, child pornography or other illicit materials could be stored and served from GENI resources. This is less likely to be from a GENI user or even a GENI aware attacker. The most likely scenario is that a resource is compromised as discussed in the section above on "Malware, Worms and Botnets", and it is discovered that there is this extra dimension of illegal materials being hosted. In this case GENI operations, the security teams at the institutions with compromised hosts, and law enforcement may all be involved in the incident. GENI operations will again be providing a coordinating role.

Researcher Laptop or Workstation Compromised

Without being specifically targeted towards GENI it is highly unlikely that an infection of malware on a laptop connecting to a GENI portal or some other resource could infect GENI. However, a key-logger on the researcher's system could collect passwords, keys or other GENI credentials. In this case, it relates to the discussion of credential harvesting above.

In any case, there little involvement GENI's incident response team would have except possibly revoking a user's credentials and disabling their account until their machine is cleaned up. The researcher's affiliated institution will likely need to be the one to verify that the system as been restored to a safe state, and depending upon the authentication system they may be the one GENI operations has to work with to disable the account.

Aggregate Owner Missing or Unknown

This sub-scenario can occur along with most any of the above scenarios. Just as it can be difficult to identify the actually machine admins for tier 3 resource providers in Open Science Grid, maintainers of small aggregates or components may be difficult to identify in an emergency. While the majority of the physical resources may be provided by small community, the majority of the resource providers may be these small sites. Managing these connections and relationships will be important so that when an incident does come to the attention of GENI operations, they can quickly identify all the components involved and the true owners of these components.

Access Control bypassed by Social Engineering

Researchers, aggregate owners, and GENI operators could all fall prey to social engineering attacks of various kinds. This is hardest to prevent for users or researchers, but they are also the least worrisome case for the integrity of GENI as the effect would be most contained. There is little other than typical user education that can be done to remedy this.

The most serious case would be someone at GENI operations or someone with control over a clearing house or CA being socially engineered in a way to compromise a significant piece of GENI infrastructure. In this case, policies and procedures can protect against social engineering. For example, one would always call back GENI operations at a specific number and not trust a call from someone claiming to be from the operations center. After the discovery of such a failure, the security operations team would have to evaluate two things: whether procedures were being followed or not and whether revised or additional procedures are needed to protect against this threat in the future. If the problem was a failure to follow procedures, audit mechanisms may need to be evaluated to determine how well the procedures are being followed more generally.

GENI Specific Incidents

In this section we discuss incidents involving third parties that are specific to GENI and not just the result of shared, federated computational resources.

Vulnerability in GENI Software

GENI will have a software stack that is used by aggregate owners, at the clearing houses and by GENI operations. If a security vulnerability is discovered in this software stack (e.g., the control framework) or a dependency (e.g., in OpenSSL), then patches must be created and distributed. The GENI incident response team will likely need to evaluate the criticality of the vulnerability, determine how it affects GENI, create a short-term work-around if needed, and ensure that all sites update the software in a timely manner. Pushing out patches will also have to be done with consideration for long-term experiments already running, and only GENI operations is likely to have a broad enough view to understand these issues.

In a federated environment, it is likely to be a challenge to get the different aggregate owners to patch software, especially if there is not an active exploit and the patch is cumbersome to apply. While the security team will likely be checking for compliance, it is not clear who would have the authority to demand that the aggregate owners patch a particular vulnerability or to set the schedule. This must be defined early in the life of GENI and addressed in the agreements signed by aggregate owners and other owners of GENI infrastructure. Furthermore, the consequence for non-compliance must be clear and enforceable (e.g., a component may be disconnected from GENI or accounts may be disabled).

Vulnerability in Experiment Software

A vulnerability may also be found not in the GENI software stack, but software specific to an experiment. This could be in the underlying OS of the virtual machines or virtual routers, in code written by the experimenter, or in utilities or services used by the experimenter. Common to all of these is that the software is in a virtual environment. Therefore, a single exploit is likely bound to a particular slice, but that does not mean there cannot be multiple exploits of the same vulnerability if the software is common to many experiments. The use of opt-in users also makes such an exploit much more likely as otherwise the system is closed to the outside world.

GENI's incident response team will have several roles in such a scenario. First, they will need to determine how wide-spread the vulnerability is and where it is most exposed. They will also have to determine where it has actually been exploited. In the cases where it is exploited, they may need to shut down experiments. Last, they will need to facilitate a recovery process and help push out updates to VM images that are commonly reused, which feature the vulnerability.

False Positive Alerts

Being a networking test bed, GENI will be producing peculiar traffic that may set off false positive alerts on anomaly-based intrusion detection systems. Because GENI will likely use VPN tunnels between nodes, this may not be as much of an issue as it has been with previous test beds. Encryption would keep the IDS from being able to inspect the details of the traffic. However, we must again consider the case when opt-in users are a part of an experiment. Encryption and tunneling will have some end point where it meets the Internet traffic, and here IDSs could still be triggered.

This use case again brings to bear the importance of establishing a good relationship between the aggregate owners, the security teams at the institutions hosting GENI resources, and the GENI incident response team. The security teams at these various institutions should be aware of GENI hosts and who they can contact at GENI operations to determine whether or not an experiment is valid and behaving properly. If an experiment is misbehaving or causing problems on a hosting organization's network, the GENI incident response team may have to initiate an emergency stop of an experiment, once they identify the offending experiment.

Certificate Authority Compromise

Certificate Authorities (CAs) are a significant point of trust in all the currently discussed GENI authentication and authorization systems proposed. Compromise of a CA could compromise the integrity of all or much of GENI. In addition to strict guidelines for hardening and auditing CAs, the GENI security team should have a recovery processed developed in advance. Creating them on-the-fly or after the fact will significantly affect recovery, complicating the process and keeping GENI offline for much longer than necessary.

GENI is DoS'ed

GENI does have centralized points that could suffer a Denial of Service (DoS) attack. While this seems an unlikely threat, the GENI incident response team would be primarily responsible for recovery from such an attack. A DoS against GENI would most likely attack the clearing house(s), but it could also hit portals, the control plane or certificate authorities. Replication or real-time backups of these services would be prudent.

More likely, but much less disruptive, a DoS could take out an aggregate manager because of an attack against the hosting institution. GENI security would be much less likely to be involved in this case, and it could be handled by the same mechanisms that handle a failure of any component or aggregate.

Lastly, such an attack could come from another testbed that GENI is federated with abroad. For this reason and others, there should be a process for deciding when to disconnect from a partner (likely the responsibility of GENI operations), and the GENI incident response team should have a process to quickly make such a disconnection.

GENI Worm

While a more generic worm spreading through the Internet is more likely, a worm could be specifically created to spread through a slice. A more difficult but more useful worm would exploit vulnerabilities in virtualization software to spread vertically through different experiments as well as laterally through slices. Such an attack is unlikely to be an end to itself but would be trying to either DoS GENI, gain control of GENI to DoS another entity or gain control of GENI to build something (e.g., an anonymizer network).

With all the instrumentation and researchers paying attention, it would be difficult for all but the slowest and stealthiest worms such as this to go unnoticed. Most likely, if GENI operations did not notice such an incident, they would get strange reports piling up from researchers quickly. GENI's incident response team would then work with researchers, aggregate owners and others to determine how GENI was exploited. They may also have to shut down many experiments depending on the nature of the incident.

Recovery would be challenging and depend upon specifics of the incident, such as, whether it exploited a vulnerability in the GENI software stack, GENI authentication and authorization, or vulnerabilities in clearing houses. There would not be one single recovery process for a GENI worm. Also, if NSF GENI was federated with other partners, this would be additionally complicated from a political perspective. Whether NSF GENI was a source of a worm spreading to other testbeds or got infected from other testbeds, international coordination would complicate incident response. It may in that case be necessary to temporarily disconnect from these partners.

Other GENI Malware

As discussed above, generic malware on a researcher's work station is unlikely to directly harm GENI in any way expect by harvesting credentials. However, it is conceivable that someone with enough knowledge of GENI could infect a PC with malware specifically intended to infect an experiment and spread from there. This is highly improbable in our estimation. Furthermore, the most likely case would still be that the malware is jailed to a specific experiment.

If such an incident did occur, the GENI security team would want to evaluate the malware and verify that it was contained to a particular experiment. They would also seek a signature so they could examine other experiments for signs of similar malware. They would have to work closely with the researcher and the security team at their institution to investigate such a case. They would also need to work with the developers of the GENI software suite to close any holes that may be exploited by such malware.

Intra-GENI Incidents

These are incidents that do not involve any third party attackers and only actors within GENI.

Resource Contention

Either maliciously or accidentally an experiment could take more than its share of resources and negatively affect other experiments. This should be a rare case since the use of virtualization in GENI's architecture should mitigate many of these problems.

GENI operations will be involved in either case as the offending experiment will need to be investigated and possibly halted via the emergency stop mechanism. The incident response team will need to determine if this is a case of an accidental or runaway experiment or if something more malicious has happened. For example, credential theft could be used to gain access to more resources, or at least to reduce usage by competing experiments.

Data Theft or Leakage

It is possible that companies with competing proprietary technologies will be sharing resources on common GENI infrastructure. In this case, there will be strong interest in preserving confidentiality in regards to an experiment and its results. We neither want researchers stealing data nor inferring it from shared resources with another experiment.

It is extremely difficult to detect when such a breach has occurred, especially if it is an inference rather than someone exceeding their privileges and viewing a data set that they should not. Much of this has to be addressed in the architecture. The strong use of virtualization in hosts and the network do provide some strong isolation properties that makes this kind of incident much less likely. In fact, the most likely way this would happen is not by inference but by credential theft allowing one to monitor another experiment. While it is theoretically possible to break out of the virtualization into the hypervisor, that is a much more sophisticated attack.

In any case, accusations of data theft will need to be examined by the GENI security team if they are made. They will need cooperation from aggregate and component owners and from the researchers involved.

GENI as an Attack Platform

This section describes GENI being used for nefarious purposes to attack others or promote illegal activities. In general, we see these as being the least likely security use cases.

GENI Used as an Anonymizer

It is conceivable, though unlikely, that GENI could be of interest to attackers as a way to create a mixing network like TOR. They could then use GENI to hide the origin of attacks. While this is possible given enough understanding of GENI, it is likely the more difficult path. TOR already exists, and attackers often already use their botnets for this purpose. Creating specialized GENI software for the task would be the hard way to accomplish the goal.

If this did happen, it would likely be contained to certain slices. The incident response team would need to stop those slices. Also, they would want to investigate the intrusion and make sure that nothing outside the virtual resources was infected. It is conceivable that the infection could be in the hypervisor layer and more persistent, or it could be utilizing an unknown exploit in the control framework. In either case, cleanup and investigation would need coordination between several aggregate owners and the GENI incident response team.

GENI used to DoS

A large GENI slice could be illegitimately created or hijacked to perform a DoS against some other site. This would not be much different from a botnet, except that it would require the attacker to have specialized knowledge of GENI and how to create GENI experiments. Furthermore, it does not seem like they could even create something as large as current botnets. Therefore, we see this threat as unlikely. It is possible for a valid GENI user to also perpetrate such an attack, but that is also unlikely due to the fact that it would be quickly detected and most likely easily traced to that user.

However, if GENI is used in such a way, the GENI incident response team will need to use the emergency stop mechanism to halt the offending experiment. It will also need to coordinate with aggregate owners to perform the proper forensic investigation to make sure the incident was contained to a particular slice. Compromised credentials, some sort of exploit of the GENI control framework, or a compromise of the clearing house could all be involved in such an incident.

GENI used as a Cracker

While GENI is designed for its networking research capabilities, it is still going to be a significant computational resource. In theory, it could be used to do anything current botnets do. This includes cracking keys and passwords. Further, it is not unprecedented to see such utilities running on research systems, especially in the HPC community.

The investigative responsibilities and procedures for the GENI security team would be similar to the use case already discussed for GENI being used as an anonymizer. As with that use case, we would most likely expect the activity to be contained to a single slice or experiment. However, a full investigation would be necessary to make sure it did not spread to other slices than the one where the detection was first made.

Lastly, it is not clear that an emergency stop would be necessary if no party is being actively harmed and the experiment was still running. In fact, it may be desirable to turn this slice into a honeypot to discover how initial compromise was made. Such decisions would have to be made on an individual basis and with the consent and knowledge of all parties involved: experimenters, GENI operations and aggregate owners.

Incident Response Foundation

Several policies, processes and technologies are needed as a prerequisite for a successful incident response program for GENI. We can see many of these things discussed repeatedly in the scenarios above.

One of the most important things to establish in a federated environment such as GENI is a list of points of contact at the different organizations. The GENI incident response team will need points of contact for each experiment, for each aggregate/component owner and the security teams at the institutions hosting GENI components. GENI's incident response team will need to be able to quickly identify these persons and coordinate communication between these different parties for many types of incidents.

Policies must also set forth the proper expectations of responsibility between all the parties listed above, especially in regards to incident response. Questions, such as, "What maintenance is expected of aggregate owners?", "Who is allowed to talk with the media in the event of a security incident?", "What information is private and what will hosting institutions share with GENI operations?", will all have to be answered.

Secure communication channels between the various security and incident response teams will have to be established, whether that is simple PGP email keys or a centralized wiki and file storage space. Furthermore, it will be the responsibility of the GENI incident response team to establish the communication between different institutions involved in the same incident.

Lastly, processes and procedures to recover from the different types of incidents must also be established and agreed upon by the GENI community. This is necessary to minimize interruption to the services provided by GENI and is a benefit to all parties.