CSD : Asset Valuation and Risk Assessment Report

This page last changed on Apr 12, 2010 by slagell.

Document Name: Asset Valuation Report ~~and Risk Assessment Report~~
Version: 0.2
Date: Apr. 12, 2010

This document focuses on asset identification and assigning qualitative values to those assets. This is a prerequisite to determining the impact of any threat. Version 0.2 incorporates feedback on this initial version from BBN. Version 0.3 will incorporate even more feedback from a broader audience that will have commented on version 0.2. Notice that the full title is somewhat deprecated with our new milestones from the April 2010 SOW revision, hence the crossed-out text.

The first version of this asset report (v0.1) included soft assets like reputation and relationships with stakeholders. However, it has been noticed that these relationships will only be affected by loss of some more tangible asset or capability within GENI (If an asset is of no concern to any stakeholder, then it hardly seems worth considering in any future risk analysis). So these soft assets are indirect and can be removed from this report, as long as this list of assets specifies the *most* relevant stakeholders. Of course, different stakeholders are more or less important to the project, and thus the associated stakeholders impact the value of these assets. Please refer to the *Catalog of Relevant Use Cases* (Milestone 1) for a complete list of stakeholders.

The structure of this document is as follows. Section 1 lists the GENI documents that have had the most impact on this document, though dozens have probably influenced this document in minor ways. Section 2 lists, categorizes and describes the different assets, briefly discussing how security incidents may affect those assets. Section 3 describes the methodology for assigning values to assets and maps assets to these values in a large table. Section 4 describes the changes between versions.

1. 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://groups.geni.net/geni/attachment/wiki/GeniSysOvrvw/GENISysOvrvw092908.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
GENI_Concept_of_Operations-final	"GMOC: GENI Concept of Operations", Oct. 1, 2009 http://gmoc.grnoc.iu.edu/uploads/8i/Gu/8iGu80-LqQB37VU4ZE1i5g/GENI_Concept_of_Operations-final.pdf

2. Asset Descriptions

Data & Communications

Experimental Data & Templates

The results of experiments are going to be very valuable to the individual researchers, but templates of experiments are likely to also be very useful to future experimenters as well. It is expected that many experiments will not be created completely from scratch, but build off of similar past experiments to make development more rapid. At the very least, there will be system images and VMs used as a starting point by many experimenters.

These assets can be affected by denial of service, deletion or corruption. The stakeholders most directly affected are the researchers.

Communication Channels

This includes email lists and their archives right now. In the future it could include other things like RSS or Twitter feeds. However, right now important information is mostly sent via email lists.

These can be affected by denial of service attacks as well as incidents that destroy email archives. Currently the stakeholders who depend most on these assets are the aggregate owners and the GMOC.

Software

There are two main categories of software that fall under this label. First, there is the software that is developed as a part of GENI infrastructure (e.g., control framework and clearing house software), but there is also software developed by the researchers for their experiments.

DoS attacks that affect code repositories, bug trackers or software update systems can affect these assets. Malware can also threaten this software, as well as newly discovered vulnerabilities in libraries or other dependencies. The stakeholders most directly affected by attacks on these assets are the researchers, aggregate owners and GMOC.

Documentation

There will be copious documentation created as GENI goes operational for both its use and for development of its software. Much of this is on the wiki right now, or at least indexed there.

Documentation is threatened mostly by incidents that cause its deletion or DoS its distribution channels (e.g., the wikis). The stakeholders most dependent upon this asset are clearly the researchers and experimenters.

Backups

This includes backups of experimental data, email list archives, software; basically anything mentioned above that GENI takes responsibility for backing up.

Incidents that damage backups are the primary threats. We are assuming that there is no highly confidential or proprietary data on GENI whose simple exposure would be detrimental. Researchers are most directly affected by loss of these backups, though aggregate owners could be as well.

Credentials / tickets / tokens

These are all assets that need to be protected. Some of them have a very short life-time (e.g., tickets and tokens), but other authorization and authentication credentials will be longer term.

The primary threat against these assets are that they are captured and reused in inappropriate ways by those other than their owners. Credential theft is a security problem in many environments, but it is too early to say how much of a threat it will be to GENI as the authentication and authorization systems have not converged to a single solution yet.

Researchers are clearly affected by loss of credentials, but it also exposes aggregate owners to risk as well. Certainly the security teams at the organizations hosting such aggregates care about this even if aggregate owners do not themselves.

Equipment & Services

Much of the infrastructure of GENI is distributed across several organizations, and this will be even more true in the future. We have thus broken down these assets into three main categories: centralized assets owned by NSF or federated partners of GENI, common aggregates or components whose loss individually is minor, and special aggregates or components whose loss individually is major. For example a cluster of generic hosts may itself not be very important, but in aggregate ALL of the hosts are critically important. GENI cannot operate without ANY hosts. Other special aggregates, like those providing the backbone network, are very important and their loss individually can affect all of GENI. Also, some of the following assets are proposed in various GENI documents but may never come to fruition. So there are many "potential" assets described here as well.

It should also be noted that the owners of these assets are only the best guess at this time. The assignment to owners could change in the future for some of these. Furthermore, some of these assets may have centralized components and decentralized components at the various aggregates. For example, any sort of logging infrastructure will probably have a centralized component as well as components at each aggregate.

GENI owned assets

Presumably not everything will be federated and there will be some GENI/NSF owned resources, which will be administered through organizations with NSF contracts. The GMOC is a prime example of such an entity now, and their infrastructure is part of what will be considered in this section.

The main threat to most of these services are incidents that cause a denial of service, though they could be affected by other kinds of incidents as well. Threats to these assets can be very serious because they can affect all experimenters, though threats whixh could affect ALL aggregates at once are even more serious.

Web Portals

While not well-defined at this phase, there will be several tools presented to researchers to control and run their experiments. Discussions to date have hinted that these will be provided through one or more web portals. These "portals" will support several functions, tools for **provisioning, resource discovery, experiment control, instrumentation control, virtual network configuration,** and **archiving/ sunsetting**.

The most relevant stakeholder here is the researcher.

Clearing Houses

The clearing houses are critical infrastructure for resource discovery in GENI. At a minimum, there will be an NSF sponsored clearing house in the beginning, though there may be more in the future as GENI federates with other GENI-like testbeds. This is a critical piece of infrastructure that can be a central point of failure.

The most direct stakeholders are the researchers and federated partners.

Slice Authorities may be a sub-component of these clearing houses that deserve special attention in future plans.

Certificate Authorities

All of the proposed mechanisms of authentication and authorization for GENI involve some sort of certificates. This means there will be a need for certificate authorities and OCSP or CRL servers. At a minimum, there will be some CAs associated with the NSF clearing house and GENI affiliated research organizations. It is unclear how distributed this PKI infrastructure would be at this stage.

In addition to DoS risks, incidents that compromise the integrity of a CA can have great impact. Researchers, aggregate owners and their host organizations (typically universities) are most concerned with threats to these assets.

Authentication Servers

Authentication will likely be federated, but this does not mean there will be no commonly held databases, directories or centrally owned front-ends to get credentials. Some of this will be at the GENI affiliated research organizations, but not necessarily all. It is too early now to say much about these assets.

However, they would be threatened by credential theft as well as the common threat of denial of service. The researchers and GENI affiliated research organizations are the biggest stakeholders here.

Collaboration Tools

GENI users/researchers may desire collaboration tools in the future. Some have envisioned GENI supported **chat** servers, **whiteboards**, and **visualization** tools. These are non-critical assets. The stakeholders would be the researchers.

Log Hosts & Accounting Services

The GMOC will likely need to keep some statistics particularly in regards to accounting for NSF. This will require some reporting to a centralized log server at the GMOC.

Threats to the integrity of these assets must also be carefully considered. NSF and the GMOC are the most direct stakeholders, though the researchers and aggregate owners may depend on some of this infrastructure to debug problems.

GENI web presence

GENI will have **web servers** and **list servers** as a part of its public face. NSF and the NetSE Council are the primary stakeholders.

Update Servers

GENI software running on different aggregate managers and other GENI control framework software will likely need to be updated from time to time. Update servers for this software will may need to be maintained by a central organization like the GMOC.

Aggregate owners are probably the biggest stakeholder here, the researchers will be concerned if there becomes an inconsistency in software versions and features across GENI because of failings here.

Backup Services

GENI will have the ability to archive experiments. This information and all the GENI developed software (e.g., for control frameworks) will need to be backed up by some system. Integrity and availability of these assets are critical, especially integrity.

The main stakeholder here is the experimenter. As it currently does not look like GENI operations will maintain source code repositories for control framework software, the aggregate owners are less of a stakeholder.

Internet Gateways

GENI will need some servers to provide connectivity to opt-in users. It is unclear if these will be run by aggregate owners or some central GENI resource. Possibly both will be done.

The main stakeholders are the researchers and the opt-in users.

Intrusion Detection Systems

GENI may have some sort of network-based IDS to detect runaway experiments. Almost certainly some of this infrastructure would have to be located at partner sites providing the network backbone or at the major aggregate providers. However, we would also need some centralized place to analyze the data, perhaps at the GMOC.

There may also be a host-based IDS distributed among the aggregates. Stakeholders are most directly aggregate owners, GENI ops, NLR & Internet2, and the operations teams at the GENI-affiliated research organizations hosting aggregates.

Measurement Equipment

Measurement equipment will need to be distributed, but also some of the infrastructure for presentation will have to be centralized. For example, GENI may provide **visualization** services for the measurement data.

The most direct stakeholders are the researchers.

Aggregate Owned Assets

As GENI will be federated, many of the assets will be owned by researchers at GENI affiliated research institutions. This includes at a minimum all the aggregates/components.

It must also be pointed out that while none of these resources are critical in isolation, in aggregate they are GENI. So when we discuss the "value" of these assets later, we must be clear whether we are talking about a single such example of an asset, a set of them or all of them.

Aggregates / components

There are many types of resources that aggregate owners will make available for slices to run over. Slivers could exist on **clusters**, **virtual hosts**, **routers, wireless testbeds, sensor networks, regional networks, storage systems** and systems and networks not yet envisioned.

The main stakeholders are the particular aggregate owners and the experimenters interested in the resources.

Aggregate/component Managers

Each aggregate or component needs an aggregate manager to communicate with the clearing house about the availability of its resources. This is critical infrastructure at the level of the aggregate, but it is distributed so that one down aggregate manager doesn't affect all of GENI.

The main stakeholders are the researchers and the particular aggregate owners.

Operations portal

There has been discussion of some sort of operations portal, perhaps on the aggregate manager that could be used for various tasks (e.g., emergency stop mechanism and exporting logs).

The most direct stakeholder is GENI operations and the NSF.

Measurement tools

There is likely to be instrumentation and measurement infrastructure at least at some of the aggregates. In particular, this will be important for wireless test beds.

The experimenters are the most direct stakeholders.

Host-based IDS

There may be host-based IDSs used at the aggregates and correlated with data at some centralized GENI server. The aggregate owners, GENI ops and the GENI-affiliated research organizations are the most direct stakeholders.

VPNs and Encapsulating Tunnels

Experimenters may wish to tunnel traffic between slivers in an experiment for privacy and security. This may be done in many different ways which may require special hardware assets.

The most direct stakeholder in such resources are the experimenters, but GENI ops and the affiliated research organizations have an interest as well because such experiments will be better isolated and less likely to be hijacked by attackers.

Special / Critical Aggregates

There are other partners providing assets beyond the average aggregate owner. For the typical aggregate, it's loss is not critical. But for these aggregates, loss of service can cause problems for all or most experimenters.

Network backbone

NLR and Internet2 provide some critical GENI infrastructure currently in the form of network back bones. Threats to their availability as well as inappropriate use of these resources must be considered.

The stakeholders include NLR, Internet2, the experimenters, the GMOC and probably NSF and BBN to some degree.

Opt-in user hosts

Projects like Million Node GENI (MNG) will be using vast quantities of opt-in user hosts, actually installing software on these hosts. They are supposed to be protected by sand boxes, so the code running on those hosts could not do anything malicious. However, there are no 100% guarantees, and thus there is a risk of attack against or with these opt-in user hosts.

The main stakeholders here are the opt-in users, and NSF whose reputation would be on the line if millions of opt-in user machines were misused.

3. Asset Values

We have *four* qualitative classifications of value based on whether or not compromise or unavailability of the asset causes:

- **Critical:** *immediate inability to run experiments for 25%+ of users, or complete system outage for 24+ hours;*
- **Important:** *immediate inability to run experiments for fewer than 25% of users, or complete system outage for < 24 hours;*
- **Normal:** *reduced service from long queues or fewer types of experiments can be run because unique equipment is out, or loss / compromise of the asset leads to a critical failure if unaddressed for more than a week; and*
- **Non-essential:** *user frustration, loss of efficiency or other problems that do not prevent experimentation and do not affect the majority of users.*

Below is our initial attempt to classify the above assets into these categories. This will be an on-going process that will need feedback from the various GENI stakeholders. Over time, assets may be combined, removed or reclassified.

Asset	Critical	Important	Normal	Non-essential
Data & Communications				
Experimental Data & Templates			X	

Communication Channels				X
GENI Software	X			
Documentation				X
Backups			X	
Credentials / Tickets / Tokens			X	
GENI Equipment & Services				
GENI Owned				
Web Portals	X			
Clearing Houses	X			
Certificate Authorities	X			
Authentication Servers	X			
Development Servers			X	
Collaboration Tools			X	
Log Hosts & Accounting Services			X	
GENI Web Presence				X
Update Servers			X	
Backup Services			X	
Internet Gateways		X		
Intrusion Detection Systems			X	
Measurement Equipment	X			
Aggregate Owned*				
Aggregates/ Components		X		
Aggregate Managers		X		

Operations Portal		X		
Measurement Tools		X		
Host-based IDS				X
VPNs & Tunnels		X		
Special / Critical Aggregates				
Backbone network	X			
Opt-in user hosts			X	

* In this evaluation we are assuming not all aggregates or even most of them are unavailable. In particular, we assume that this affects less than 25% of the experimenters.

Revision History

Version 0.2

- Updated introduction to reflect this document's place in the revised SOW.
- Removed proprietary commercial intellectual property asset.
- Removed generic resource availability assets.
- Tied assets to associated stakeholders.
- Treated NLR & Internet2 a special aggregate owners rather than "partners".
- Removed development servers: cvs, svn, wikis, etc. It does not see like this will ever be a GENI owned asset, but probably just maintained at the major aggregate providers like Emulab and PlanetLab. It is unclear that there is much we can do to protect these resources. This does not mean that there is 0 risk that important code to GENI could be lost, though.
- Added special aggregate assets by opt-in users