G E N I
Global Environment for Network Innovations

Towards Operational Security for GENI

Jim Basney
National Center for Supercomputing Applications (NCSA)
University of Illinois, Urbana-Champaign

Roy Campbell
Information Trust Institute (ITI),
University of Illinois, Urbana-Champaign

Himanshu Khurana
(Document Coordinator)
University of Illinois, Urbana-Champaign

Von Welch
University of Illinois, Urbana-Champaign

GDD-06-10

July 2006

Status: Draft

Towards Operational Security for GENI

Jim Basney[1], Roy Campbell[2], Himanshu Khurana[1,2] (Document Coordinator), Von Welch[1,2]

[1]National Center for Supercomputing Applications (NCSA)
[2]Information Trust Institute (ITI)
University of Illinois, Urbana-Champaign

{jbasney@ncsa.uiuc.edu, rhc@uiuc.edu, hkhurana@ncsa.uiuc.edu, vwelch@ncsa.uiuc.edu}

Draft dated June 30, 2006

NOTE: This document is a draft that has been posted online for comments from the GENI community at large. We seek inputs on major errors and misrepresentation, if any, as well as additional topics that should be included in the final version of the document or in future editions. Comments can be sent to us or to the main GENI discussion mailing list. We will upload the final version soon after the San Francisco Town Hall Meeting (July 13, 14).

Abstract

Ensuring operational security of the GENI facility is critical to its success and the current GENI Design Document identifies some important security threats and concerns. Addressing these threats and concerns requires the implementation of a comprehensive set of operational security policies and procedures. This working document identifies necessary steps for ensuring operational security of GENI by detailing major threats and exploring kinds and examples of necessary policies and procedures. We identify three next steps, namely, development of (1) a security architecture, (2) a set of agreements that ensure conformance to security requirements by participating sites and organizations, and (3) an implementation integration plan for implementing and enforcing the agreements.

1. Introduction

The GENI facility is envisioned to provide a test-bed for experimentation that researchers can use to evaluate new network technologies with a goal towards deployment and transition to industrial development. Currently, researchers study their technologies with limited simulations and experiments on small-scale prototypes, which are insufficient to demonstrate effectiveness for deployment purposes. GENI will fill this gap and allow execution of the complete research cycle. The end goal of the individual technologies studied with GENI is the deployment of a new Internet that is secure and robust, supports new network and computing technologies, and supports new distributed applications and systems.

An extensive GENI facility that can support the entire research cycle needs design, implementation, and maintenance. The current PEP (Project Execution Plan)[1] identifies facility components that fall into two broad categories, namely, the physical network substrate and the global management framework. The physical substrate will be built on a nation-wide high-speed backbone (e.g., using one or more NLR – National Lambda Rail – lambdas), which connects *edge sites* (e.g., Universities) that host computational nodes. The building blocks of the physical substrate include clusters of commodity PCs that are capable of hosting virtual machines, customizable high-speed routers, optical fiber lambdas and switches that enable high-speed

[1] http://www.geni.net/GDD/GDD-06-07.pdf

networks, tail circuits for Internet inter-connections and tunneling, 802.11-based mesh networks, 3G/WiMax radio networks, cognitive radio networks, and sensor networks. The management framework will embed and manage *slices* in the GENI substrate, where each slice comprises a set of GENI resources and is assigned to a research experiment. This framework comprises component managers for allocating and controlling embedded slices, a GENI management core for instantiating and remotely managing slices across building blocks, a set of infrastructure services that allow researchers to manipulate and interact with GENI, and a set of underlay services that allow management and control of experiments. Furthermore, the PEP outlines specific deployment tasks for setting up the various GENI components over a period of time taking into consideration availability of technologies, costs, and per-task effort.

Ensuring security of the GENI facility is critical to its usefulness and success. Because of its visibility and resources, GENI is a high risk target for attacks. Thus, various threats and security concerns have already been identified. Identified threats include experimental services that may be vulnerable to attacks or may be the source of attacks, and the resource sharing nature of GENI making resource depletion attacks possible. Identified security concerns include the security of the GENI building blocks, the security of the experimenter environment, the security from errant experiments, and the security from malicious GENI nodes.

Addressing these security threats and concerns, as well as those that arise with time, requires the establishment of operational security policies and procedures for all organizations that participate in GENI. Similar to the plan of deploying GENI components over time, the security policies and procedures will evolve over time as well. This working document attempts to explore the kinds and examples of necessary policies and procedures. The responsibility of securing GENI lies with the core GENI administration team, the edge sites that host GENI resources, the experimenters that use GENI resources, and users that play a part in the experiments. This document attempts to answer the following question: how should organizations (e.g., Universities), researchers, and users that wish to participate in GENI prepare for securing the GENI network? Though GENI is clearly unique in its envisioned collective capabilities, several individual capabilities have commonalities with existing test-beds and distributed computing systems; e.g., PlanetLab [1, 7], Deter [2], X-Bone [10], TeraGrid[2], and Optiputer [9]. We look at several of these systems and incorporate the lessons learned from securing them into the policies and procedures.

Based on our analysis we recommend efforts in the following three directions. First, there is a need to develop a *security architecture* for GENI, which will include security perimeters, requirements, and technological solutions for satisfying the emerging set of requirements. Second, there is a need for developing *agreements* that will be signed by GENI participants to enable operational security. We argue that because GENI is envisioned to be a large federated system with GENI resources being hosted by edge sites (e.g., Universities), agreements for securing these resources are needed. This is analogous in some ways to the networking agreements between these sites for connecting GENI resources to the GENI backbone, edge site networks, and the Internet. Two primary agreements are those that specify policies and procedures for baseline security and for incident handling and response. Third, there is a need to develop *an implementation integration plan* that integrates the implementation and enforcement of security policies and procedures with the overall GENI implementation plan. In order to specify, obtain sign-offs, and control changes to these documents and plans, there is a need to establish a *security management authority*. The authority should ideally comprise a representative set of individuals that can ensure a community-wide consensus-based process.

[2] http://www.teragrid.org/

The intent of this document is to identify the major and unique threats to GENI, explore the kinds and examples of operational security policies and procedures that would be needed to address these threats, discuss the challenges in implementing the policies and procedures, and recommend a path forward to define, agree, and implement the necessary policies and procedures. The focus of this document is restricted to GENI's physical substrate and management framework components. It does not include a discussion of additional security services that may be deployed as part of experiments or perhaps even integrated as facility components to support a number of experiments. Such services will be part of the evolving security architecture much like the facility architecture itself. However, previously identified underlay services that provide security are included in the discussions.

The rest of this document is organized as follows. Section 2 identifies elements of the GENI architecture. Section 3 identifies major threats and points of vulnerability in these architectural elements. Section 4 explores security policies and procedures needed to prevent, detect, and respond to attacks at these points. Section 5 makes recommendations for next steps in ensuring operational security for GENI.

2. GENI Architecture Elements

Though the GENI architecture is currently in its design phase and, furthermore, will be an evolving one even after initial deployment, the primary elements can already be identified. To both simplify and focus on security, we identify architectural elements in three different categories, namely, organizations/domains, processing elements, and networking elements. Organizations are entities of autonomous security administration that are responsible for securing the processing and networking elements that lie within their boundaries. Processing elements are GENI hosts and devices that provide processing capabilities for GENI virtual networks and experiments. Networking elements are the fibers, cables, wireless/radio/sensor subnets, routers, switches, Internet exchanges, and gateways that connect the processing elements.

Organizational autonomy will play an important role in GENI security in general, and managing incidents in particular. The cardinal rule of operational security is that if an incident originates from an organization's IP space then it's the organization's problem. In case of GENI where GENI resources are hosted within an organization's IP space this leads to issues with autonomy because while both the organization and GENI will be morally responsible for incidents originating in those resources, the organization will be legally responsible. In this regard GENI's nature of a federated system will likely require the support of every site's administrative and perhaps even legal teams.

The following is a technology-independent list of architectural elements that we expect to see in GENI.

- **Organizations**
    - *Edge Sites*. These organizations host GENI processing and networking elements and have researchers and users that set up, maintain, and participate in GENI virtual networks/experiments. These organizations may be connected to the GENI backbone via physical links or virtual ones over the Internet.
    - *Participating Sites*. These organizations have researchers and users that set up, maintain, and participate in GENI virtual networks/experiments. These organizations may be connected to the GENI backbone via physical links but are most likely to be connected with virtual links over the Internet.

- *GENI Core Organization*. This logical organization hosts processing and networking elements that comprise the GENI core facility; e.g., the backbone and infrastructure services. In practice, the Core Organization will likely be distributed over multiple edge nodes; however, since it needs to secure its resource autonomously we consider it to be a single (logical) organization.
- *Other Organizations*. In the overall GENI system there will be other organizations that provide necessary services. For example, those that provide leased bandwidth to GENI in the backbone or network connections at Points-of-Presence (PoP). However, it is unlikely that GENI can rely on these organizations for security; therefore, we do not consider them in this discussion.

- **Processing Elements**
  - *Hosts*. These are physical machines that run GENI virtual networks and management functions and belong to a particular organization. For example, edge site commodity clusters than run virtual machines and component managers, and Core Organization servers that run the GENI Management Core (GMC) on behalf of a set of GENI edge sites and gateways for enabling virtual links between participating nodes and the GENI backbone.
  - *Devices*. These are the processing elements of wireless, radio, and sensor subnets that belong to a particular organization. For example, edge site radio nodes in the WiMax subnet.
- **Networking Elements**
  - *GENI Core Elements*. These are the fibers as well as routers and switches that comprise the GENI backbone managed by the Core Organization. These elements also include Internet Exchanges that are either part of the backbone or provided at PoP.
  - *Edge Site Wired Elements*. These are the fibers/cables as well as routers and switches that connect hosts, the GENI backbone (via PoP) and the commodity Internet (via site Internet connections). In particular, the following distinct network connections will be needed:
    - Between GENI hosts to form a host subnet.
    - Between GENI host subnets and PoP. This physical connection to GENI may be limited to a subset of edge sites.
    - Between GENI host subnets and local site networks. This connection is needed for several reasons; e.g., to enable (1) site administrators to remotely administer GENI host subnets, (2) virtual links to GENI backbone gateways via the Internet (typically commodity Internet access is available at the periphery of site networks), and (3) site researchers and users access to host subnets via high-speed links (if needed).
  - *Edge Site Wireless Elements*. These are the wireless, radio, and sensor networks and the corresponding routers and gateways. In addition, these are also the elements that connect the wireless/radio/sensor networks with the host subnets. Note that some radio nodes will connect to both wired and wireless routers so there is an overlap between these and the wired elements.

3. Risks and Threat Analysis

An insecure GENI facility is at risk. In order to understand these risks we must first look at what GENI will represent to the nation and to the world: GENI will be seen as a significant scientific research infrastructure that is funded by the US Government for developing the next generation of

advanced information technologies. Though the infrastructure is envisioned to be an International one, the core infrastructure including the backbone will be housed in the US. Therefore, successful attacks against this infrastructure can bring disrepute to this national effort as well as for the scientific community, the funding agencies, and industrial collaborators. Furthermore, successful attacks can cause a shift in funding priorities leading to long pauses in GENI's progress and delaying the deployment of researched technologies. A recent analogous example is the delays in the NASA Space Shuttle program because of safety issues. These risks will only get compounded when GENI considers integration with other large test-beds and distributed systems; e.g., those part of the US Military networks and International projects. The compounded risks include the potential of GENI resources being used to attack Military networks and lack of effective regulations and law enforcement in foreign countries leading to bolder attacks. All of these risks make it imperative that we understand the technological threats to GENI and enforce necessary policies and procedures to minimize the possibility of a successful attack.

We now look at the some of the major security threats and concerns identified in the PEP as well as others that have been faced by large multi-site test-beds and distributed computing systems in the recent past. We map these threats and concerns on to the GENI architectural elements and identify specific points of vulnerability. Note that this is not meant as an exhaustive threat analysis but only as a representative one and, furthermore, one that will evolve with time.

At a high-level, a sample potential attack on GENI (or on any other large distributed system) can be characterized as follows. The adversary would begin an attack by exploiting software vulnerabilities at a particular GENI host or device. This exploit would grant him certain privileges that provide access to services running on the host as well as to the networking elements to which the host is connected (e.g., ports). The adversary can also attempt privilege escalation attacks at this compromised host in order to gain more privileges. Using these privileges the adversary would attempt to compromise other hosts on the GENI network that can be contacted via the networking elements connected to the originally compromised host. Some of the targeted hosts in this case would be those that run more critical services or that are connected to bigger networking elements. At these hosts the adversary may again attempt privilege escalation. This process may continue for a while depending on the adversary's success. At some point in time the adversary can launch an attack with a significant impact that concerns the GENI community at large; e.g., denial-of-service on GENI resources as well as on the Internet. In this significant attack the adversary would use all the services, processes, and accessible networking elements available to him at the compromised nodes (with associated privileges). Clearly, the greatest threat comes from distributed attacks where an adversary compromises a large number of hosts before launching the "significant" attacks. In rare cases, the adversary may also succeed in compromising networking elements such as routers to cause even bigger problems.

The above scenario outlines the steps an attacker may take against GENI or against the Internet via GENI. Another source of attacks that remote adversaries may attempt are via the use of viruses and worms. In these attacks, worms can be programmed to quickly corrupt systems and propagate themselves throughout the network by exploiting software vulnerabilities and using available networking elements for the propagation.

Misconfigurations and errant GENI experiments are clearly not malicious but they can potentially lead to attacks in the following way. Misconfigurations can grant processes and services additional privileges or access to networking elements that they don't need. Errant experiments can result in processes and services using their privileges to direct networking traffic and requests towards GENI, organizational or Internet resources via accessible networking elements that would not be sent under correct operating conditions. Individually or combined together,

misconfigurations and errant experiments can lead to significant attacks that also concern the GENI community.

Based on this high-level description of attacks we now identify points of vulnerability in the GENI architectural elements that can lead to significant attacks.

- **Processing Elements**. One could say that the various GENI hosts and devices are the primary points of vulnerability because they might be infected with exploitable software vulnerabilities. Different kinds of GENI hosts, if compromised, can lead to different kinds of attacks.
    - *User Desktop Machines*. These machines will be used by researchers and users to connect to the GENI network for setting up, maintaining, and participating in GENI experiments at both edge sites and participating sites. Compromise of these machines may give the adversary access to credentials (e.g., username/password) for GENI accounts; i.e., lead to account compromise. These machines are often user administered and connected to open networks (e.g., at most Universities) making it difficult to protect them from occasional compromise (e.g., ensuring up-to-date patching).
    - *Virtual Machine Hosts*. These hosts will run virtual machines at edge sites that form the virtual networks for GENI experiments. Compromise of a virtual machine on such a physical machine can lead to compromise of the virtual network/experiment and may provide opportunities to attack other virtual machines running on the same physical machine, the component managers running on the physical machine, the physical machine itself, and any other resources accessible via connected networking elements. Further compromise of the host, for example, the component manager, can enable control over slice resource allocations. Challenges in protecting these machines include experimental services that may have software vulnerabilities.
    - *Core Organization Hosts*. These hosts run important management and infrastructure services for the entire GENI network; e.g., GMC and provisioning and account management services. Compromise of these hosts can lead to the compromise of crucial GENI services. In practice, these hosts should only offer software services that are well defined and trustworthy (e.g., approved by a vetting process) and should be closely monitored.
    - *Devices*. These devices are part of the wireless, radio, and sensor subnets and run GENI experiments on those subnets. They may be compromised by an adversary that is either within the wireless range of the subnet or can access the subnet remotely through wired connections to the larger GENI network (if present). Compromise of these devices can lead to effects similar to the compromise of virtual machine hosts.
- **Networking Elements**. These architectural elements have vulnerabilities of two types, namely, machines (e.g., routers) and network paths.
    - *Machines*. Routers, switches and gateways form the core of the GENI network as well as enable connection to site networks and the Internet. Compromise of these machines can lead to the adversary having direct control of the networking paths of which the machines are a part. In practice, these machines should be well-configured and should be closely monitored.
    - *Network Paths*. A major source of vulnerability here are the network paths available to an adversary (or errant experiment and worms) that connect him to other hosts and systems for further compromise or attack. These network paths can lead to the GENI network, the site network, or the Internet. Controlling and monitoring connections on these network paths can protect GENI from

significant attacks but this activity faces challenges of system and personnel costs.

4. Security Policies and Procedures

In this section we explore security policies and procedures for preventing, detecting and responding to GENI incidents. This is not a comprehensive list but instead a representative one geared towards the eventual establishment of acceptable GENI security policies and procedures. Policies are documented guidelines that need to be specified by organizations to define the overall approach for securing GENI. Procedures are steps for implementing the policies that involve human administrators and instrumented tools, technologies, and mechanisms.

Developing and enforcing a comprehensive set of operational security policies and procedures for GENI is relatively unique and challenging primarily because GENI is envisioned to be a large federated system. In such a multi-site system resources are federated between sites (that have administrative and legal responsibilities for all GENI resources that lie in their IP space) and GENI Core (that connects GENI site resources with the larger GENI networks and has the responsibility to ensure its availability). Even the GENI Core is likely to comprise several edge sites and it is essential that GENI have the administrative and legal support of these sites to ensure secure operation of GENI Core resources. The contention that arises as a consequence is deciding who is responsible for preventing, detecting, and responding to security incidents; e.g., who should manage patch updates, who should be contacted when attacks are detected, and should an incident be considered a site incident or a GENI incident or both. Note that at each edge site resources such as clusters and server farms will need and require site administrative support for installation and maintenance. The need arises because of the expertise involved and the requirement arises from legal responsibilities of the site. Therefore, resolving this contention is not simply about some researcher "owning" these resources and granting GENI Core complete control over them. These issues must be resolved with site administrative and perhaps even legal teams, which stresses the need for a collaborative, community-wide effort in establishing the necessary policies and procedures. These challenges are faced today by other multi-site distributed systems[3] and while the GENI environment may be different from these systems, efforts in securing them can provide useful guidelines and lessons learned.

4.1 Prevention

The aim of prevention policies and procedures is to minimize the (1) presence of vulnerabilities (e.g., via patching), (2) ability of an adversary to exploit the vulnerabilities (e.g., via firewalls/filters and appropriate authentication and authorization measures), and (3) limit the scope of attacks if vulnerabilities do get exploited (e.g., via rate limitation). Often these are part of the site security policies [5]. In general, there is an array of best practices that need to be followed for preventative policies and procedures with the following being some of the primary elements of such policies and procedures.

- *Host Protection* including secure software assurance practices [8], software updates, patch management, configuration, and assignment and separation of privileges to accounts and processes.

[3] Examples of multi-site distributed system security policies are those defined for TeraGrid: http://www.teragrid.org/basics/security.html and for the LHC Grid in Europe: http://lcg.web.cern.ch/LCG/activities/security/security.html

- *Network protection* including configuration, ingress and egress filtering, routing protocols, service/port blocking and restrictions, and rate limitations.
- *Authentication and authorization* including mechanisms that provide security in accordance with the privileges associated with an account (e.g., username/password or Public Key Infrastructure (PKI) certificates for user accounts while hardware token based One-Time-Passwords (OTP) for administrator accounts), associated trust mechanisms (e.g., Certificate Authority (CA) policies and root certificate distribution), and protection of credentials over the network (e.g., prohibiting cleartext passwords).
- *Security audits and drills* including periodic internal and external reviews and exercises that document weaknesses and suggest improvements. Such reviews go a long way towards ensuring security. Several test-beds and systems have undertaken such internal and external reviews voluntarily with documented results that serve as important lessons learned [3, 4].

Several policies and procedures for prevention may need to be agreed upon between the sites to ensure adequate levels of protection. For example, one site may desire to apply patches as soon as they are available while another site may want to make sure that the patches break no applications before applying them; how should GENI deal with this?

4.2 Detection

IDS policies and procedures are geared towards *signature detection* (where the "unusual" is well-defined and anything belonging to this category is considered an event), *anomaly detection* (where the "usual" is defined and anything out of the ordinary is considered an event), or a combination of these two. Signature and anomaly detection can be done at both processing elements (i.e., via host based IDSs or HIDS) and networking elements (i.e., via network based IDSs or NIDS) though current IDSs face challenges in dealing with false positives and false negatives. As a result, effective monitoring of large systems that successfully detects intrusions requires a combination of instrumented IDS hardware and software systems and system administrators. The IDSs usually provide a large number of alerts and human administrators use intuition and experience to follow up on "meaningful" alerts. Administrators use several tools to aid in the follow up efforts including, for example, visualization tools and logging techniques. The following are some important points on the GENI facility that need HIDS and NIDS.

- *HIDS on Processing Elements*. HIDS that provide, for example, integrity checking, process monitoring, and virus/worm detection can be successful in detecting intrusions at virtual machine hosts, subnet devices, and Core Organization hosts that run management services.
- *NIDS on Networking Elements*. A few NIDS that provide rule-based signature and anomaly detection and are placed at strategic points on the GENI facility can be successful in detecting attacks against the facility. Examples of strategic points include those (1) between edge site resources and edge site GENI resources and (2) at inter-connections within the GENI backbone. The first will detect attacks from the GENI subnet to the site or to the Internet from the site network and vice versa. The second will detect attacks against the GENI facility from a compromised edge site GENI subnet. Since at least a subset of GENI networks will be high-speed in nature, novel hardware-based NIDS may need to be deployed that are costly and may need additional staff training.

4.3 Forensics, Collaboration, and Response

The detection of a successful intrusion triggers an iterative process of forensics and response where, depending upon the understanding of the attack, appropriate response mechanisms are used in each iteration with the final iteration being the complete restoration of services. For example, if the GENI backbone NIDS detects a DoS attack from a particular virtual network then the first response might be to take the entire virtual network offline. As the forensic investigation begins it might discover that only a few virtual machine hosts at a couple of edge sites are responsible for the traffic. In that case the response is modified to bring the virtual network back online but keep the compromised virtual machine hosts offline. Further investigation might reveal that a particular software vulnerability exists at the hosts and was used by the adversary to compromise the machines via the network. Now the response will be to patch the machines, remove any malware that may have been installed on those machines and completely restore the virtual network.

There are three crucial components of this forensic discovery and response process, namely, logging, collaboration between GENI sites, and remote command and control capabilities.

- *Logging*. In order investigate an attack administrators need data that can help analyze the path taken by the adversary. This data is provided by logs; e.g., those generated by NIDS and HIDS as well those generated by networking and processing elements including router logs and syslogs. The GENI management framework is already envisioned as providing some of these logging capabilities that would be very useful in intrusion detection and response; e.g., at GMC and in infrastructure services.
- *Collaboration*. In the example above, a crucial component of the forensic discovery and response process is the collaboration between the edge sites where the virtual machine hosts were compromised. Without effective collaboration the detection will either take longer or not succeed at all. Effective collaboration requires sites to (1) know who to contact in case of an incident (especially after hours), (2) securely communicate with responders (if vulnerable to eavesdropping or impersonation these channels may allow the adversary access to sensitive information and may make the system vulnerable to social engineering attacks), (3) share incident data and logs, and (4) work together to eliminate the adversary's advantage and restore services (e.g., to clean up and patch user desktop machines). Often organizations are reluctant to share incident data because of reasons of privacy and negative publicity. Therefore, an agreement between the GENI sites is essential in enabling collaboration. Examples of such agreements include the memorandum of understanding between the TeraGrid sites[4] and incident handling policies of the LCG/EGEE Grids [6]. Additional incident response policy issues that require collaboration include funding agency notifications, media handling (e.g., should the cites contact the media or should GENI) and dealing with law enforcement (e.g., ensuring evidence gathering and sharing).
- *Remote command and control*. Effective response to an incident requires remote command and control capabilities at multiple levels of granularity; e.g., shut down a virtual network or virtual machine, take a physical machine off the network and modify NIDS/HIDS policies or firewall/filter rules. Many such capabilities are already envisioned as part of the GENI management framework; e.g., via component managers and GMC. These capabilities must require strong authentication and authorization to ensure that they are not misused.

[4] http://security.teragrid.org/docs/Security-MOU.txt

5. Recommendations

Based on our threat analysis and exploration of security policies and procedures for the GENI facility, we make the following recommendations to ensure a comprehensive approach for securing the facility before it is made available to researchers and users. As the GENI community progresses in ensuring operational security, additional steps may need to be defined.

- Develop the *Operational Security Architecture* for GENI (or, identify security components in the GENI Architecture), which will define at least the following:
    - Organizational boundaries and security perimeters
    - Requirements for securing each class of GENI resource
    - Tools, technologies, and mechanisms for satisfying requirements; e.g., network and host-based IDSs, authentication mechanisms, logging and remote command and control mechanisms
    - An analysis of risks in the architecture; e.g., threats that cannot be addressed due to cost limitations
- Develop *Agreements* that will be signed by all sites to enable operational security in the multi-site GENI system. Two primary agreements are:
    - A Baseline Operational Security Document that will define at least the following:
        - Minimum acceptable level of preventive policies and procedures to ensure overall GENI operational security
        - Additional security requirements for sites contributing critical resources; e.g., a site that provides account management
    - An Incident Handling and Response Procedures Document that will define at least the following:
        - Information on contact personnel (especially after hours)
        - Secure communication requirements and solutions between responders
        - Steps for incident detection, collaborative forensics, containment and response, and service restoration
        - Policies for communicating with media and funding agencies as well as working with law enforcement
- Develop an *Implementation Integration Plan* to implement and enforce operational security policies and procedures that will define at least the following:
    - Estimates of staff and training needs
    - A budget of costs for staff as well as for necessary tools and mechanisms
    - Timelines and support for implementing policies and procedures at sites
    - Periodic audits and drills to ensure conformance
    - An operational maintenance plan
- Establish a *Security Management Authority* that comprises a representative group of individuals that will
    - Specify the above documents
    - Obtain agreements on them from GENI participants
    - Guide and control changes to the documents including, for example, deployment of additional security services into the GENI facility to supporting experiments
    - Specify additional vetting procedures; e.g., certifying trustworthiness of hardware and software for core/critical services

Acknowledgements

This document has benefited greatly from discussions with Jim Barlow, Nikita Borisov, Carl Gunter, Ravi Iyer, David Nicols, Bill Sanders, and Tony Rimovsky.

References

1. Bavier, M. Bowman, B. Chun, D. Culler, S. Karlin, S. Muir, L. Peterson, T. Roscoe, T. Spalink, and M. Wawrzoniak. Operating System Support for Planetary-Scale Network Services. *First Symposium on Networked Systems Design and Implementation (NSDI)*, March 2004, 253-266.

2. Benzel, Braden, Joseph, Kim, Neuman, Ostrenga, Schwab, and Sklower. Experience with DETER: A Testbed for Security Research. *Tridentcom* 2006, Barcelona, Spain, March 2006.

3. Paul Brett, Mic Bowman, Jeff Sedayao, Robert Adams, Rob Knauerhase, and Aaron Klingaman. Securing the PlanetLab Distributed Testbed: How to manage security in an environment with no firewalls, with all users having root, and no direct physical control of any system. *Proceedings of LISA '04: Eighteenth Systems Administration Conference*, Atlanta, GA: USENIX Association, November, 2004.

4. J. Clem, B. Badgett, T. MacAlpine. X-Bone: Automated System for Deployment and Managament of Network Overlays. Security Assessment Report. Information Design Assurance Red Team, Sandia National Laboratories. April 21, 2003.

5. B. Fraser (Editor). Site Security Handbook. IETF Network Working Group. RFC 2196. September 1997.

6. Joint Security Policy Group. LCG/EGEE Incident Handling and Response Guide. Technical Report. LHC Computing Grid. June 2005. Available at: https://edms.cern.ch/document/428035

7. L. Peterson, T. Anderson, D. Culler, and T. Roscoe. A Blueprint for Introducing Disruptive Technology into the Internet. *First Workshop on Hot Topics in Networking (HotNets-I)*, October 2002.

8. Samuel T. Redwine, Jr. (Editor). *Secure Software Assurance: A Guide to the Common Body of Knowledge to Produce, Acquire, and Sustain Secure Software* version 0.9 (Draft). US Department of Homeland Security, January 9, 2006

9. Larry L. Smarr, Andrew A. Chien, Tom DeFanti, Jason Leigh, and Philip M. Papadopoulos. The OptIPuter. Special issue on *Blueprint for the future of high-performance networking, Communications of the ACM*, Volume 46, Issue 11, November 2003, pp. 58-67.

10. J. Touch, Y. Wang, V. Pingali, L. Eggert, R. Zhou, G. Finn.Proc. A Global X-Bone for Network Experiments. *IEEE Tridentcom* 2005, Trento Italy, Mar. 2005, pp. 194-203.