Operational Security Requirements for Large Collaborative Compute Infrastructures

Himanshu Khurana[1,2], Jim Basney[1], Von Welch[1,2], Roy Campbell[2]
[1]National Center for Supercomputing Applications (NCSA)
[2] Information Trust Institute (ITI)
University of Illinois, Urbana-Champaign
{hkhurana, jbasney, vwelch, rhc}@uiuc.edu

Abstract

Large collaborative infrastructures that span multiple organizations such as those that enable grid computing and scientific experimentation are being deployed and used today. In order to secure these infrastructures a comprehensive requirements study is needed that takes into account the novel risks, threats, and operational issues brought on by the large-scale, distributed nature of these systems. In this paper we argue that gaps in security policies and procedures combined with organizational autonomy are the primary drivers motivating a set of requirements that go beyond those observed today. With three example infrastructures in mind, namely, Teragrid, LHC Grid, and GENI, we explore the novel risks, threats, and operational issues to compose a set of operational security requirements; the satisfaction of which will be essential for securing such large collaborative infrastructures.

1. Introduction

Collaborative compute applications combined with unique organizational skills and resources are motivating the deployment of Large Collaborative Compute Infrastructures (LCCIs). These LCCIs have several interesting characteristics: a common goal, multiple organizations, autonomous administration, and an open, evolutionary nature. A common goal motivates their need; establishment over multiple organizations' (or sites') physical infrastructures enables resource sharing to achieve the common goal with each organization still administering its resources autonomously; an open nature allows dynamic addition of organizations and users who desire to utilize these infrastructures at a global scale; and, an evolutionary nature allows for development of new services over the life of the infrastructure. The TeraGrid[1] is an example of a deployed LCCI infrastructure while the LHC (Large Hadron Collider) Computing Grid[2] and GENI[3] (Global Environment for Networking Innovations) are examples of planned LCCIs. While all three LCCIs have different objectives, they share the characteristics identified above

Just like any other major compute infrastructure, security of LCCIs is crucial for their success. The enormous effort in terms of both manpower and funding result in LCCIs being viewed as major, representative technological capabilities of the associated scientific communities, funding agencies, and nations. Consequently, they become attack targets for determined adversaries and script kiddies alike. The risks of insecure LCCIs that can lead to successful intrusions and attacks must therefore be weighed carefully. Such attacks can bring disrepute to the associated entities and question their credibility. Furthermore, they can even cause delays in the deployment and use of the LCCIs (an analogous example is the recent delays in the NASA Space Shuttle program due to safety issues) or, at the very least, have a negative impact on future collaborative endeavors involving these entities.

In this paper we argue that LCCIs face new and magnified threats, which when combined with the identified risks, lead to a new set of operational security requirements. While each organization has its own set of security policies and procedures for preventing, detecting, and responding to attacks, conflicts and gaps between policies and procedures of different organizations in a LCCI create new threats. For example, a LCCI organization may have minimal authentication procedures for registering users allowing potential adversaries easy access to other organization's systems (via the connected LCCI) that otherwise would not allow users access to their system without extensive authentication. In addition, the presence of multiple network paths in the connected LCCI and a large number of users with potentially untrusted end systems and desktops seriously magnifies the threats to LCCIs and their host organizations. These magnified threats further strain the limited manpower for security

[1] http://teragrid.org

[2] http://lcg.web.cern.ch/LCG

[3] http://www.geni.net

administration leading to missed attacks and coarse-grained response such as disconnecting entire organizations from their networks.

In the face of these new and magnified threats we explore the kinds of security policies and procedures that would be needed to secure LCCIs. Based on this exploration we propose the following requirements: (1) specification of an *operational security architecture* that defines organization boundaries and major resources, (2) an *agreement* on the *baseline security policies and procedures* that each organization must implement, (3) an *agreement* on *incident handling and response* that guides the organizations in case of an incident, (4) specification of an *implementation integration plan* to implement the architecture and the agreements, and (5) establishment of a *security management authority* that oversees the development and modification of these policies, procedures, agreements, and plans. The security management authority is needed to address the evolutionary nature of LCCIs and must be compromised of individuals who represent the diverse, collaborative nature of LCCIs. While some of these requirements are already being addressed in practice (e.g., TeraGrid and LHC have incident handling and response guidelines in place) we believe that addressing all of the identified requirements is essential for securing LCCIs.

The rest of this paper is organized as follows. In Section 2 we give examples of LCCIs. In Section 3 we abstract the architecture elements of LCCIs. In Section 4 we discuss risks and threats. In Section 5 we explore security policies and procedures. In Section 6 we propose requirements for operational security of LCCIs and conclude in Section 7.

2. Example LCCIs

In this section we review the characteristics of three example LCCIs, namely, TeraGrid, LHC Grid, and GENI. Other infrastructures and test-beds that share characteristics with these ones but perhaps not at the same scale include PlanetLab [1], [8], Deter [2], X-Bone [11], and Optiputer [10].

TeraGrid is a LCCI for science experimentation enabled by resources provided by nine sites/ institutions. Currently, the integrated resources include more than 100 teraflops of computing capability and more than 15 petabytes of storage with all of the resources being connected via high-performance connections. These resources are coordinated by the Grid Infrastructure Group (GIG) led by the University of Chicago and are being used by scientists worldwide in over a 1000 projects. The TeraGrid infrastructure comprises a high-speed National Lambda Rail backbone, compute clusters at sites connected to the backbone and to the Internet, and software systems at sites to manage the resources and to enable scientists to conduct experiments. Typical experiments are long running compute jobs that use resources of multiple sites orchestrated by advanced scheduling systems.

The LHC Grid is currently being designed and deployed to support scientific experiments that utilize data generated by the LHC. By 2008 integrated resources are expected to provide 140 million SPECint2000[4] compute capabilities, 60 petabytes of disk storage, and 50 petabytes of mass storage. To enable worldwide access to the data, a four-tier architecture with over 100 sites is being planned with Tier-0 being located at CERN, Tier-1 and Tier-2 being large compute facilities, and Tier-3 being university departments. The infrastructure will comprise high-performance network connections (between Tier-0 and Tier-1 sites), large compute clusters, and massive disk storage systems (at least at the higher tiers). In addition, there will be extensive software systems managing access to data and enabling experimentation. A typical experiment may include a complex simulation running on a high-performance compute cluster and accessing massive amounts of data from multiple sites.

The GENI LCCI 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. The current PEP (Project Execution Plan)[5] identifies GENI facility infrastructure components that are likely to span two dozen sites and will 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 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 networks, tail circuits for Internet inter-connections and tunneling, 802.11-based mesh networks, 3G/WiMax radio networks, cognitive radio networks, and sensor networks. A management framework embeds and manages *slices* in the GENI substrate, where each slice comprises a set of GENI resources and is assigned to a research experiment. This

[4] In current terms this would be equivalent to about 80,000 powerful Pentium processors.

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

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. Typical experiments will be virtual networks deployed and managed by scientists that run distributed systems and applications developed for a new and more robust Internet.

3. Architecture Elements

LCCI architectures are very different from each other as well as evolving in nature. However, as indicated by the three examples above, they have common architectural elements. To both simplify the discussion and focus on security, we identify architectural elements in four different categories, namely, organizations/domains, processing elements, networking elements, and other 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 hosts and devices that provide processing capabilities for LCCI experiments. Networking elements are the fibers, cables, wireless/radio/sensor subnets, routers, switches, Internet exchanges, and gateways that connect the processing elements. Other elements include storage nodes and scientific instruments.

Organizational autonomy will play an important role in LCCI 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 LCCIs where LCCI resources are hosted within an organization's IP space this leads to issues with autonomy because while both the organization and the LCCI will be responsible for incidents originating in those resources, the organization will be legally responsible. In this regard, the nature of a federated system in LCCIs 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 any LCCI.

- **Organizations**
  - *Edge Sites*. These organizations host LCCI processing and networking elements and have researchers and users that set up, maintain, and participate in the LCCI experiments. These organizations may be connected to the LCCI via physical links or virtual ones over the Internet.
  - *Participating Sites*. These organizations have researchers and users that set up, maintain, and participate in LCCI experiments. These organizations may be connected to the LCCI via physical links but are most likely to be connected with virtual links over the Internet.
  - *LCCI Core Organization*. This logical organization hosts processing and networking elements that comprise the LCCI core facility; e.g., the backbone and infrastructure services. In practice, the Core Organization will likely be distributed over multiple edge sites; however, since it needs to secure its resources autonomously we consider it to be a single (logical) organization.
  - *Other Organizations*. In the overall LCCI system there will be other organizations that provide necessary services. For example, those that provide leased bandwidth to the LCCI in the backbone or network connections at Points-of-Presence (PoPs). However, it is unlikely that the LCCI 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 LCCI experiments and belong to a particular organization. For example, edge site commodity clusters that run distributed applications, and Core Organization servers that run management functions.
  - *Devices*. When LCCIs include wireless, radio, and sensor subnets, then these are the processing elements of wireless, radio, and sensor subnets that belong to a particular organization. For example, these devices include GENI's edge site radio nodes in the WiMax subnet.
- **Networking Elements**
  - *LCCI Core Elements*. These are the fibers as well as routers and switches that comprise the LCCI backbone managed by the Core Organization. These elements also include Internet Exchanges that are either part of the backbone or provided at PoPs.
  - *Edge Site Wired Elements*. These are the fibers/cables as well as routers and switches that connect hosts, the LCCI backbone (via PoP) and the commodity Internet (via site Internet connections).
  - *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.

- **Other Elements**
    - *Storage Nodes.* These are the disks and tapes as well as the backup and replication systems used to manage access to experimental data. These storage elements are rarely the source of security failures but they can be the targets of an attack.
    - *Instruments.* These are scientific instruments (e.g., the LHC collider) that are used in LCCIs for generation of experimental data.

4. Risks and Threat Analysis

LCCI facility security entails unique risks. As discussed in the Introduction, successful attacks against LCCIs can bring disrepute to national efforts as well as to the associated scientific communities, the funding agencies, and industrial collaborators. These risks will only get compounded when the LCCI considers integration with other large test-beds and distributed systems as part of its evolutionary process; e.g., those part of the military networks and international projects. The compounded risks include the potential of LCCI resources being used to attack other networks and the challenges of responding to attacks crossing organizational and national boundaries (including regulatory and law enforcement obstacles). All of these risks make it imperative that the technological threats to LCCIs be well understood and necessary policies and procedures be established to minimize the possibility of a catastrophic attack.

We now look at the some of the major security threats and concerns that will be faced by LCCIs. We map these threats and concerns on to the architectural elements and identify specific points of vulnerability. Note that this is not meant to an exhaustive threat analysis but only a representative one.

At a high-level a sample potential attack on a LCCI (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 LCCI host or device. This exploit would grant him certain privileges that provide access to services running on the host as well as to networking elements that the host is connected to. The adversary can also attempt privilege escalation attacks at this compromised host. Using these privileges the adversary would attempt to compromise other hosts on the LCCI 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 have greater network connectivity. At these hosts the adversary may again attempt privilege escalation. The attack may continue to spread. At some point in time, which could be considerably later than the intrusion making detection challenging, the adversary can launch an attack with a significant impact that concerns the LCCI community at large; e.g., compromise of data on storage nodes, denial-of-service on LCCI 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.

While this general attack scenario is applicable to any LCCI, we note that specific attacks may also be guided by the adversary's desire to utilize specific LCCI resources. For example, LHC Grid provides access to very high capacity compute power and data resources leading to potential misuse for cryptoanalysis while GENI offers access to router algorithms, virtual networks and slices, and future network security provisions

The above scenario outlines the steps an attacker may take against a LCCI or against the Internet via the LCCI. 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 LCCI applications 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. Misconfigurations can also mask actual attacks by classifying events incorrectly or flooding monitoring services with too many flagged events. Errant experiments can result in processes and services using their privileges to direct networking traffic and requests towards the LCCI, 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 LCCI community.

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

- **Processing Elements**. The various LCCI hosts and devices are the primary points of vulnerability because they might be infected with exploitable software vulnerabilities. Different kinds of LCCI 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 LCCI network for setting up, maintaining, and participating in LCCI experiments at both edge sites and participating sites. Compromise of these machines may give the adversary access to credentials (e.g., username/ password) for LCCI accounts; i.e., lead to account compromise. These machines are often user administered and connected to open networks (e.g., at universities) making it difficult to protect them from occasional compromise (e.g., with no mechanism for ensuring up-to-date patching).
  - *Experiment Hosts*. These hosts will run LCCI experiments and applications. Compromise of such a host can lead to compromise of the experiment and may provide opportunities to attack any other resources (including storage elements) accessible via connected networking elements. 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 LCCI network; e.g., provisioning and account management services. Compromise of these hosts can lead to the widespread compromise of LCCI services. In practice, these hosts should only provide 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 LCCI 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 through wired connections to the larger LCCI network (if present).
- **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 LCCI network as well as enable connection to site networks and the Internet. Compromise of these machines can lead to the adversary having direct control over the networking paths of which the machines are a part. In practice, these machines should be well-configured services 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 LCCI network, the site network, or the Internet. Controlling and monitoring connections on these network paths can protect LCCI from significant attacks but faces challenges of system and personnel costs.

5. Security Policies and Procedures

In this section we explore security policies and procedures for preventing, detecting and responding to LCCI incidents. This is not a comprehensive list but instead a representative one geared towards the eventual establishment of acceptable LCCI security policies and procedures. Policies are documented guidelines that need to be specified by organizations to define the overall approach for securing the LCCI. 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 a LCCI is relatively unique and challenging primarily because any LCCI will be a large federated system. In such a multi-site system resources are federated between sites (that have administrative and legal responsibilities for all LCCI resources that lie in their IP space) and LCCI Core (that connects LCCI site resources with the larger LCCI networks and has the responsibility to ensure its availability). Even the LCCI Core is likely to comprise several edge sites and it is essential that the LCCI Core have the administrative and legal support of these sites to ensure secure operation of LCCI 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 LCCI 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 LCCI 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. For example, TeraGrid and LHC Grid have already taken initial steps in defining such security policies[6].

The federated nature of LCCIs makes it important for organizations to not only carefully craft the policies and procedures but to also carefully decide who should craft these policies. In other words, there needs to be an authority (e.g., the TeraGrid Security Working Group), ideally with representative membership, that ratifies security policies and procedures, negotiates with participating sites, enforces the policies, and oversees change to these policies.

5.1 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 [6], software updates, patch management, configuration, and assignment and separation of privileges to accounts and processes with least privilege in mind.
- *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.

5.2 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). 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 LCCI 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 experiment hosts, subnet devices, and Core Organization hosts that run management services.

[6] http://www.teragrid.org/basics/security.html; http://lcg.web.cern.ch/LCG/activities/security/security.html

- *NIDS on Networking Elements*. A few NIDS that provide rule-based signature and anomaly detection and are placed at strategic points on the LCCI facility can be successful in detecting attacks against the facility. Examples of strategic points include those (1) between LCCI resources and non-LCCI resources at edge sites and (2) at inter-connections within the LCCI backbone. The first will detect attacks from the LCCI subnet to the site or to the Internet from the site network and vice versa. The second will detect attacks against the LCCI facility from a compromised edge site LCCI subnet. Since at least a subset of LCCI 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.

5.3 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, in GENI if the backbone NIDS detects a denial-of-service 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 LCCI 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. For example, the TeraGrid logs tens of gigabytes of network traffic and host events daily, and the GENI management framework is already envisioned as providing some of these logging capabilities that would be very useful in intrusion detection and response.
- *Collaboration*. In the example above, a potentially 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 outside normal business 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 LCCI sites is essential in enabling collaboration. For example, such agreements include the memorandum of understanding between the TeraGrid sites[7] and incident handling policies of the LHC Grid [7]. Additional incident response policy issues that require collaboration include funding agency notifications, media handling (e.g., should the cites contact the media or should the LCCI) 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. These capabilities must require strong authentication and authorization to ensure that they are not misused.

6. Requirements

Based on our threat analysis and exploration of security policies and procedures for a LCCI facility, we identify the following requirements to ensure a comprehensive approach for securing LCCIs.

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

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

7. Conclusions and Economics

In this work we argue that LCCIs face new and magnified threats. The new threats emerge from the federated nature of LCCIs while magnified threats emerge from the large scale of LCCIs. To deal with these threats we identify requirements for operational security of LCCIs. We argue that these requirements must be met in order to minimize risks of successful intrusions. We give examples of security policies and procedures that would satisfy the identified requirements.

An important aspect of LCCI operational security is financial costs, both for preventive measures and for responding to incidents. These are the costs for proactive and reactive security measures, respectively [9]. We argue that by addressing the requirements identified in this work, both costs will go down significantly[8]. For example, (1) if a site cannot afford high-performance IDS systems it can leverage those of other sites in the LCCI with appropriate security agreements, and (2) by instating an agreed upon incident response plan the time and cost for reacting to incident will go down. A quantitative analysis of such costs can provide further insights on how to formulate cost-effective security policies and procedures.

Acknowledgements

We would like to thank Jim Barlow, Tim Brooks, Aashish Sharma and Tony Rimovsky for insights into day-to-day security operations at computational facilities. We would also like to thank numerous members of the TeraGrid Security Working Group and the Open Science Grid Security Technical Group for insights into operational security issues for LCCIs. This material is based upon work supported by the Office of Naval Research (ONR) under Grant No. N0001404-1-0562 and the National

[8] Other cost reduction measures include establishment of Computer Security Incident Response Teams - *http://www.cert.org/csirts/*

Science Foundation under Grant Nos. SCI-0438712 and CNS-0524695. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the ONR or NSF.

References

[1] M. Bavier, B. Bowman, D. Chun, S. Culler, S. Karlin, L. Muir, T. 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.

[2] T. Benzel, A. Joseph, D. Kim, C. Neuman, R. Ostrenga, S. Schwab, and K. Sklower. Experience with DETER: A Testbed for Security Research. *Tridentcom*, Barcelona, Spain, 2006.

[3] P. Brett, M. Bowman, J. Sedayao, R. Adams, R. Knauerhase, and A. 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, 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] M. Gaff and K. van Wyk. *Secure Coding: Principles and Practices*. O'Reilly, 2003

[7] 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

[8] 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)*, 2002.

[9] B. Rowe and M. P. Gallaher. Private Sector Cyber-Security Investment: An Empirical Analysis. *The Fifth Workshop on the Economics of Information Security (WEIS).* June 2006.

[10] L. Smarr, A A. Chien, T. DeFanti, J. Leigh, and P. 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.

[11] J. Touch, Y. Wang, V. Pingali, L. Eggert, R. Zhou, and G. Finn.. A Global X-Bone for Network Experiments. *IEEE Tridentcom* 2005, Trento Italy, Mar. 2005.