

2007 NSF Cybersecurity Summit Final Report

Crystal City, VA
February 22-23, 2007

2007 Cybersecurity Summit Program Committee

James J. Barlow, Chair	National Center for Supercomputing Applications
Tom Bettge	National Center for Atmospheric Research
Paul Dokas	University of Minnesota
Christopher L. Greer	National Science Foundation
Victor Hazlewood	Oak Ridge National Laboratory
Steve Lau, Jr.	University of California San Francisco
Corbin Miller	NASA Jet Propulsion Laboratory
Rodney J. Petersen	EDUCAUSE
Don Petravick	Fermi National Accelerator Laboratory
Nigel Sharp	National Science Foundation
Dane D. Skow	Argonne National Laboratory

Table of Contents.....2

Executive Summary.....3

Overview.....4

Plenary Sessions Summary.....5

Breakout Session Summary.....8

 Breakout Session 1: Policies, Standards, and Guidelines8

 Breakout Session 2: Identification and Classification.....10

 Breakout Session 3: Authentication, Authorization, and
 Accounting.....11

 Breakout Session 4: Incident Handling/Response.....16

 Breakout Session 5: System Administration Practices/Policies/
 Education.....19

Conclusions and Final Recommendations.....21

Participant Evaluation Summary.....22

Appendix A: Program.....23

Executive Summary

The 2007 Cybersecurity Summit workshop built upon the previous two workshops by bringing together security professionals from many of the NSF-funded large research facilities. This was an invitation-only event whose theme was “Providing security throughout the research data lifecycle.” The National Science Foundation sponsored the workshop, which was attended by around 100 participants. There was a broad range of security professionals represented, from program officers and CIOs down to network and system administrators. This brought a good, diverse perspective from the number of different organizations who sent attendees to the event.

As in previous years, the main goals of this workshop included:

- Sharing of information and ideas
- Understanding our communities’ diverse perspectives
- Discussion of our communities’ strengths and weaknesses
- Identifying our Communities’ security needs

The plenary sessions and breakouts were loosely tied to the “Providing security throughout the research data lifecycle” theme. Five breakout sessions were chosen for participants to have detailed discussions:

- Policies, Standards, and Guidelines
- Identification and Classification
- Authentication, Authorization, and Accounting (AAA)
- Incident Handling/Response
- System Administration Practices/Policies/Education

The goal of the breakout discussions was to develop recommendations for research and education sites, as well as recommendations the NSF can use when determining proposals for funding. These recommendations are detailed in the breakout session summaries.

A workshop of this size and scope would most likely not be possible without the help of EDUCAUSE. The EDUCAUSE staff helped coordinate many of the logistics, from location specifics (hotel, conference rooms, etc.) to hosting the workshop Web pages, registration, on-site conference coordination, and after-conference participant surveys. They did another excellent job this year in coordinating logistics for this workshop.

1. Overview

1.1 Motivation

The third NSF-sponsored Cybersecurity Summit workshop was held in February 2007. This workshop continued the success of the previous years' workshops. The first workshop, held in November 2004, was initiated because of a security incident that affected numerous NSF-funded organizations, as well as research and educational facilities around the globe. The second workshop, held in December 2005, was recommended in the 2004 final report because of the high interest and value that participants perceived from the first year's workshop. As with the second workshop, this one did not focus on a particular incident, but brought people together to focus on data security and related topics.

The following goals have held for this and the previous workshops:

- **Share information and ideas:** By sharing information and ideas, participants can understand the common issues and problems that affect security in the research and education communities. They can learn how others have solved these problems and/or identify problems that need further discussion and attention in securing the research cyberinfrastructure.
- **Develop understanding of our communities' diverse perspectives:** While balancing security and usability in the research environment, workshop attendees discuss and analyze the similarities and differences between small to large computing/research facilities.
- **Discuss our communities' strengths and weaknesses:** The research and education environment has specific and somewhat unique requirements for providing open, collaborative environments. Participants discuss and analyze the strengths and weaknesses related to security of these environments.
- **Identifying our communities' security needs:** Attendees explore the competing needs of an open, collaborative research environment and protecting the security and integrity of the nation's research computing and data assets. They strive to describe a secure computing environment that minimizes any negative impact on (a) researchers and their productivity, and (b) computer and network performance.

1.2 Program Committee and Program

James Barlow of the National Center for Supercomputing Applications was asked by the NSF to chair this year's program. A program committee was formed from many different research and educational institutions, as well as from other federal agencies. The committee met bi-weekly for 6 months leading up to the conference, which was an adequate amount of time to prepare for a conference of this size and scope. The conference was initially intended to be held three months earlier, but it was quickly realized that more time was needed for coordination. EDUCAUSE helped out immensely in many of the tasks that led to a successfully coordinated workshop. The EDUCAUSE

2007 NSF Cybersecurity Summit – Final Report

workshop Web pages helped keep all members of the program committee up-to-date on the bi-weekly calls and other planning items.

The program committee members were:

James J. Barlow, Chair	National Center for Supercomputing Applications
Tom Bettge	National Center for Atmospheric Research
Paul Dokas	University of Minnesota
Christopher L. Greer	National Science Foundation
Victor Hazlewood	Oak Ridge National Laboratory
Steve Lau, Jr.	University of California San Francisco
Corbin Miller	NASA Jet Propulsion Laboratory
Rodney J. Petersen	EDUCAUSE
Don Petravick	Fermi National Accelerator Laboratory
Nigel Sharp	National Science Foundation
Dane D. Skow	Argonne National Laboratory

The program is included in the Appendix.

1.3 Participation

This was an invitation-only event of NSF program officers, program committee members, some previous years' attendees, and other recommendations made to the program committee. A diverse group of participants were sought to contribute to this year's workshop from research, educational, and other federal agencies including law enforcement.

About 100 people participated in this year's workshop. Two of the attendees were from outside the United States, one from New Zealand and the other from Chile. The other U.S. attendees were from 18 different states, Washington, D.C., and Puerto Rico.

2. Plenary Sessions Summary

The 2007 Cybersecurity Summit program is included in the Appendix; all of the plenary session presentations are available on the summit Web site at <http://www.educause.edu/cyb07/>. Following is an overview of each of the plenary sessions.

Kathie Olsen, Deputy Director of the NSF, opened the workshop with a welcome from the National Science Foundation. Kathie covered some of the NSF's progress with respect to the last two Cybersecurity Summits and encouraged participation from this year's attendees. Jack Suess, CIO of the University of Maryland, Baltimore County, then gave an overview of the EDUCAUSE/Internet2 Computer and Network Security Task Force. This covered the organization of the higher education sector, security discussion

2007 NSF Cybersecurity Summit – Final Report

groups, the Research and Education Network Information Sharing and Analysis Center (REN-ISAC), and other working groups within EDUCAUSE/Internet2.

Jim Barlow of the National Center for Supercomputing Applications, this year's program chair, gave a welcome and explained how this year's workshop came to be. He gave a quick summary of the first two workshops, then said how this one was organized to not just focus on a particular incident but to consider how all institutions can work together to provide security throughout the research data lifecycle. Jim then introduced Bart Bridwell of the NSF, who gave a presentation on the NSF Cooperative Agreement Security Language.

The cooperative agreement language was added to the NSF's Cooperative Agreement Supplemental Terms and Conditions in September 2006 after recommendations from previous Cybersecurity Summits. The cooperative agreement language can be found at <http://www.educause.edu/ir/library/pdf/CYB07001d.pdf>.

The cooperative agreement sets forth obligations for the awardees to provide a secure information technology environment; a summary of the site's security program; evaluation criteria to measure the success of the security program; and periodic self assessments.

The keynote speaker for this year's event was Peter Gutmann from The University of Auckland. Peter's presentation, The Commercial Malware Industry, was the most well-received talk of the workshop. In his talk Peter went into various details of how the malware industry has changed since the time where many miscreants and script-kiddies were downloading and using much of the malware distributed. Nowadays we are seeing very well developed malware created by professional programmers and managed by international criminal organizations. This malware is used in various ways, such as for spam distribution, carding, phishing, and other targeted methods. Organizations are hiring various professionals, such as psychologists and linguists, to make sure that the malware being developed is targeting the right people and will produce the desired results. It was quite an eye-opening talk for many of the attendees, and the presentation generated lots of good discussion.

The next speaker was William Cook from Wildman Harrold Attorneys and Counselors. This presentation focused on some of the legal aspects of data security. Mr. William Cook covered some of the areas that researchers need to be concerned with, such as policies and practices, doing offsite work, and competitor practices. Some case studies were then presented that covered these issues.

The next set of speakers presented real-world scenarios of the different approaches to data management and data security. The three speakers were Don Petravick from Fermi National Laboratory representing the Open Science Grid, Frank Siebenlist of Argonne National Laboratory representing the Earth Science Grid, and Dane Skow of Argonne National Laboratory representing the TeraGrid.

2007 NSF Cybersecurity Summit – Final Report

The Open Science Grid is a distributed computing infrastructure for large-scale scientific research. It was built by a consortium of universities, national laboratories, scientific collaborators, and software developers. The core security within the OSG is not the security of the sites or the virtual organizations (VOs), but the security of the OSG organization, including data flows like accounting information, its VDT-based software stack, and its configuration management methods. The OSG's role is to not bear the security responsibility of the sites or VOs, but to facilitate it by trying to standardize the discussion. Examples of this were given during discussion of the PANDA architecture and various acceptable-use policies.

The Earth Science Grid is making climate simulation data available globally. Currently around 160TB of data is available from 876 different data sets and 840,000 files. Frank went over some of the current security architecture process, including . an example of delivering data through a data portal and the security processes behind it for access controls and access policies. ESG's goal is to enable and not limit access to its data sets , which involves a number of complex challenges and interoperability requirements to overcome.

Dane started with the TeraGrid mission, which is to provide integrated, persistent, and pioneering computational resources that will significantly improve our nation's ability and capacity to gain new insights into our most challenging research questions and societal problems. Dane covered some of the growth within the TeraGrid as well as current and new initiatives such as the science gateway initiative. He then covered data storage resources more thoroughly, such as disk, tape, GPFS-WAN, and SRB. There were also lessons learned over the current life of the TeraGrid project , one arising from the incident that instigated the first NSF Cybersecurity Summit: one of the most valuable results of that incident was the coordination team building developed. Other lessons learned are that ease of use and ubiquity are essential to adoption of any technology. Work is still needed on distributed group authorization management tooling and the security triad: who you are, where you can go, and what you can do.

Dr. Ron Ross from the Computer Security Division of the National Institute of Standards and Technology (NIST) gave the presentation Information Systems Under Attack – Enterprise Risk in Today's World of Sophisticated Threats and Adversaries. Because of the current state of affairs in computer security, especially in regard to our critical infrastructure, protection of these resources is immensely important. Through some of the legislative policy drivers, such as the Federal Information Security Management Act (FISMA), we can build a solid foundation of information security by establishing a fundamental level of "security due diligence." FISMA characteristics were covered in more detail through a risk management framework and information security program. These standards should not drive the mission of an organization, but rather support the mission. The policies and procedures developed from these are a corporate commitment for protecting the critical enterprise.

2007 NSF Cybersecurity Summit – Final Report

Dr. Ross had a few quick tips for combating particularly nasty adversaries:

- Reexamine FIPS 199 security categorizations.
- Remove critical information systems and applications from the network whenever possible.
- Change the information system architecture; obfuscate network entry paths, and employ additional subnets.
- Use two-factor authentication, especially at key network locations.
- Employ secondary storage disk encryption.

On the second day of the summit the opening speaker was FBI Special Agent Mike Butler. Mike spoke on the FBI Counterintelligence Domain Program, which included the Academic Alliance Program. His talk started with the changing counterintelligence paradigm. Many of these changes are because of technology advances and its uses in organized crime. Because of this the FBI needs partnerships in corporate and educational institutions. So within the educational area the FBI formed a National Security Higher Education Advisory Board (NSHEAB), whose membership includes 16 different university presidents and chancellors. The goal of this partnership is to promote understanding and advice on the culture of higher education and to foster the discussion of matters pertaining to national security. Another initiative that the FBI has started is the Research and Technology Protection Special Interest Group (RTP-SIG). The RTP-SIG's mission is to provide actionable and relevant information to contractors, private industry, and academia to better enable them to protect their research and technology.

3. Breakout Session Summaries

The program committee selected five breakout sessions for detailed discussion in smaller groups among the participants. The breakouts were

- Policies, Standards, and Guidelines
- Identification and Classification
- Authentication, Authorization, and Accounting
- Incident Handling/Response
- System Administration Practices/Policies/Education

Descriptions of the breakout sessions are detailed in the following subsections.

3.1 Policies, Standards, and Guidelines

This breakout group included about 25 participants and was led by Victor Hazlewood of Oak Ridge National Laboratory and Kim Milford of the University of Rochester.

Background: Earlier in the program, Bart Bridwell reviewed the new language regarding IT security, which was developed by the NSF and incorporated into supplemental terms and conditions for Cooperative Agreements (CAs) used to fund large facilities and

2007 NSF Cybersecurity Summit – Final Report

Federally Funded Research and Development Centers (FFRDCs). The new language was adopted by the NSF with input from awardees, program officers, and NSF staff and included recommendations from last year's Cybersecurity Summit.

The language is intentionally broad to allow significant flexibility, reflecting the spectrum of different needs in awardee organizations. It is intended to ensure that information security be considered as a key element to ensure continuity of the funded research for an awardee organization in the face of increasing cybersecurity threats.

Observations:

- Including information security language in the CA is a good first step in protecting research data and systems.
- Policy is not a one-size-fits-all solution. Each organization may have different policy needs and different rules that govern how policy is promulgated. This can be even more complex in interinstitutional and interdisciplinary environments.
- Policy can take years to develop and approve. The development of policies, standards, and guidelines is an iterative process, requiring periodic review and updates.
- Development of information security controls requires expenditures. The adoption of the new clause may require that information security measures be funded through the award.
- Preserving the integrity of research and staying out of the headlines are two motivating factors for including security plans in research projects. If there is a high-profile, high-cost incident, awardees could potentially be placed in the position of meeting very restrictive controls.
- Awardees need additional guidance in developing an information security plan (e.g., templates); program officers need guidance in understanding key elements of the plan (e.g., checklists); and having some consistency across organizations would be helpful in assessing plans.

Recommendations:

- NSF should provide additional guidance as to what they want in the plan. It may make sense to include different recommendations and checklists for different types of organizations (small sites, large sites, interagency sites, etc).
- NSF should engage cybersecurity experts (and cybersecurity experts should make themselves available) to help develop models and templates. A review of available security frameworks and best practices should be undertaken. Seeing what others are doing and sharing with others can provide benefit and help them stop reinventing the wheel. Examples: EDUCAUSE, NIST, ISC2.
- Awardee organizations should engage in risk-based prioritization for information security planning.
- A central site should host models, examples, resources and training events on a single site. EDUCAUSE has a significant collection of relevant information and may be the optimal place for hosting this information.
- NSF should develop a list of cybersecurity experts to provide assistance in assessing plans and during program reviews.

2007 NSF Cybersecurity Summit – Final Report

- NSF should continue to encourage dialogue between the program officers and awardees on developing and refining security plans.
- The community (with NSF sponsorship) should hold workshops and forums for sharing technical cybersecurity tools and techniques.
- Language in CA provides a suggestive list of priorities. Awardee organization should initially focus on the two foundational areas of security planning and risk assessment, using the CA language to guide their planning.
- The breakout group found these additional areas to be high priorities. The NSF may want to begin by developing these frameworks and best practices.
 1. Acceptable use policy
 2. Incident response planning guide
 - NSF CA language calls for notification procedures regarding how the awardee notifies the NSF of incidents. The NSF should develop a consistent high-level protocol about what they need to know and when.
 - Institutions and organizations should develop incident handling and management guidelines specific to their own institutions and consistent with NSF's notification protocols.
 - Examples of incident response flowcharts: Teragrid Incident response flowchart; Yale flowchart; EDUCAUSE blueprint.

3.2 Data Identification and Classification

This breakout session was led by James Marsteller of the Pittsburgh Supercomputing Center and Andrea Nixon of Carleton College and the Internet2/EDUCAUSE Security Task Force. The breakout group consisted of four participants.

The Identification and Classification group was established in keeping with the recommendations from last year's Data Security and Integrity Goals. The Identification and Classification group began with a discussion of current practices at the participants' home institutions. The group also worked from sample data identification practices from Carnegie Mellon University as well as a sample categorization policy from the University of Texas at Austin. While there was recognition of the challenges associated with security of all data stored at our institutions, the group paid particular attention to challenges associated with protecting research data. The group identified a series of issues relevant to securing data as well as distilling a series of recommendations.

Recommendations:

- Sites looking for data identification assistance may want to review the Carnegie Mellon University Ownership of Administrative Data section of the "Data and Computer Security (Confidentiality of Administrative Data)" policy <http://www.cmu.edu/policies/documents/DataSecurity.html>. It can be used as a starting place for identifying data types as well as the data owners that should be involved in the data identification process.
- Since each institution maintains different data, they need to define data classifications that best address the needs of their organization. For guidance, sites may find

2007 NSF Cybersecurity Summit – Final Report

established classifications such as the “Data Classification Guidelines” developed by the University of Austin at Texas a useful resource to begin their own classification process. This resource was also recommended in the 2006 NSF Cybersecurity Summit report for its usefulness in protecting data integrity.

- There are some useful tools that sites can deploy to assist in identifying specific types of data (such as Social Security and credit card numbers). Cornell University’s IT Security Office offers open-source forensics tools <http://www.cit.cornell.edu/computer/security/tools/> that can aid the identification process. While these tools can produce results quickly, there is a nontrivial effort in reviewing the results to eliminate false positives and the hands-on effort needed to remove or properly protect the data. Continued use of these tools may serve ongoing efforts to identify protected data that occurs in predictable formats.
- The group recommends that funding agencies, primary investigators, and resource providers establish an understanding of how research data are classified. Once this classification is established, proper controls can be implemented for appropriate access, storage, and transmission.

3.3 Authentication, Authorization, and Accounting

The Authentication, Authorization, and Accounting breakout session was led by Von Welch from the National Center for Supercomputing Applications and involved approximately a dozen participants. There was significant carryover from the first to the second day of the summit.

Von Welch started out with a summary of the 2005 Cybersecurity Summit Authentication Breakout Group report (<http://www.educause.edu/LibraryDetailPage/666?ID=CYB0525>). Conversations then began with a discussion of changes in the past year. A broad range of topics was covered. The group also discussed an additional high-level topic, Auditing (aka the 4th “A”), which was seen as entangled with the other three topics.

Throughout the discussions, the group found it useful to refer back to the list of threats which AAA systems protect against. The threats the group came up with are:

- Vandalism/petty - e.g. Web sites
- Vandalism/serious - e.g. data
- Stealing information
- Stealing computing cycles
- Stealing storage/distribution - e.g. warez
- Launching attacks on other sites
 - “Enclaves” - e.g. TeraGrid
- Embarrassment

Changes since Previous Cybersecurity Summit

The group decided there hadn’t been any major changes in authentication since the prior cybersecurity summit, in particular:

2007 NSF Cybersecurity Summit – Final Report

- User end systems are still untrustworthy and likely to be compromised by attackers who could deploy Trojans or steal secrets (e.g. passwords, private keys) stored on local disk.
- There had been no major migration either to or from one-time passwords (OTP) by organizations. Cost of OTP, both up-front and for ongoing maintenance, along with concerns about session hijacking, were given as two reasons why sites had not moved to OTP-authentication. Leveraging OTP deployments by banks (given the FDIC mandate for two-factor authentication a little over a year ago) did not seem to be a viable alternative, since many banks were coming up with creative alternatives to OTP as a second factor and those who were deploying OTP seemed unwilling to allow its use by third parties.
- Federated identity was continuing to progress in various forms.
- Authentication of nonhuman entities, e.g. computational jobs, services, sensors, instruments, is an issue and may be growing with large sensor deployments. This discussion was captured subsequently as its own topic area.
- The federal government seemed to be backing off on HSPD-12 requirements (Policy for a Common Identification Standard for Federal Employees and Contractors), so a smaller and smaller number of High-Performance Computing (HPC) users apparently will be affected by Homeland Security Presidential Directive-12.

Operational Issues with AAA

A discussion of operational issues occurred, focusing on problems in the current revocation infrastructure for public key infrastructures (PKIs). These revocation systems (CRLs and OCSP) are too labor intensive and not tied into registration authority or human resource databases, making them unreliable. Additionally, local sites need to have their own, local deauthorization mechanisms to block users they don't like.

Provisioning of trust roots as a growing issue was also discussed. There are a growing number of identity providers (mainly certification authorities) and attribute authorities. Management of these trust roots is becoming increasingly difficult.

Recommendation: A system for revocation of users' rights needs to be developed that takes into account the different parties (identity provider, resource provider, and virtual organization) requirements and allows for time-sensitive operation given the communication overhead between these parties.

User Education

The group noted there are now a number of tools available to help users manage multiple passwords (e.g. Password Safe, Apple Keychain, various Firefox plug-ins). These tools, if used widely, could cut down on weak and/or replicated passwords across sites. User education on the availability and use of these tools could do much to strengthen existing security. It was also noted that existing password policies would need updating to allow for the use of these tools – e.g. many policies tell users not to write passwords down, but that's what these tools allow users to do in a secure manner.

Recommendation: Develop training materials for NSF users on these tools.

2007 NSF Cybersecurity Summit – Final Report

Session Hijacking

The HPC community has generally relied on sessions that authenticate the user initially and then give them a connection (which may or may not time-out) in which they can operate without further authentications. This approach is vulnerable to “session hijacking” attacks in which attackers take over an existing session after authentication has transpired, hence bypassing any strong authentication. Session hijacking is typically done using a modified kernel on the client’s system.

The group discussed two defenses against session hijacking attacks. The first was to get rid of sessions altogether, moving towards a transactional model used by banks and by middleware systems such as Unicore and Globus. This would allow each user’s action to be authenticated, with the potential for stronger authentication for more sensitive actions. The second approach was to focus effort on intrusion detection based on the user’s actions inside of a session in order to detect possible subversion of the session by a third party.

Recommendation: Research into methodologies for preventing session hijacking is needed, particularly when the client system is compromised.

Issues with One-Time Passwords

One-time passwords (OTP) was discussed as a potential solution to user identity spoofing by replacing static passwords that are currently the norm for the majority of the NSF HPC community. A number of issues were discussed:

- OTP deployment and support costs are still an issue.
- Would a system where users receive one-time passwords via an SMS message to their cell phones be a possible alternative to a dedicated hardware device? RSA, SecureComputing, and a number of European banks are using this approach. Cell phones seem to be nearly ubiquitous. An issue raised here was the cost to the end users of receiving SMS messages.
- Some banks, notably Paypal, are moving to OTP tokens, but there is concern that these banks don’t seem interested in allowing use of those tokens by third parties. It was also noted that a number of banks are moving to two-factor authentication based on other means than OTP.
- The issue of usability was raised in that currently users are required to have an OTP token for each site requiring OTP. This led to the discussion of alternatives that don’t require dedicated hardware tokens. The issue also serves as a motivation for federated identity approaches.

Recommendation: The community needs to continue to explore OTP deployment and its alternatives.

Federated Identity

The group discussed federated identity and reached general agreement that this area is progressing, although there were a number of issues that need to be addressed:

2007 NSF Cybersecurity Summit – Final Report

- With respect to level of assurance (LoA), what degree of assurance does a relying party have that the source of a person's identity information has vetted the person sufficiently and is correctly maintaining that information? There does not seem to be consistent handling of these issues among U.S. colleges and universities, though eAuthentication is starting to drive this.
- What policies are required for issues such as incident response, liability, and privacy?
- What are the requirements for citizenship in terms of operating HPC resources? While this is more of a requirement for DOE sites, there are some issues for NSF sites. When does a site have to be cognizant of the citizenship of a user they are serving? There seem to be a number of opinions about this, but none of the participants could identify an authoritative source of policy outside of their own institution.
- In P2P versus federated identity, OpenId (<http://openid.net/>) is starting to emerge as a more user-centric form of identity management (the analogy was made between OpenId and SAML as compared to PGP and PKI).

Recommendation: The community needs to continue to be encouraged to push on these policy issues and do real-world multi-institutional deployments to work through these issues.

Authentication of Nonhuman Devices

Authentication of nonhuman devices was raised as an issue. Examples of this discussed were dynamic services running on computational systems as well as small computational devices such as sensors. Services running on computational devices have the issue that any secret they store typically has to exist unencrypted, particularly if automated start-up of the device is a requirement.

The SRP and TLS-PSK protocols were discussed as a good method for a client of a service to authenticate that service. Patent issues were mentioned as an issue for some implementations.

Authorization

The group discussed four issues related to authorization:

- Debugging authorization failures, particularly in distributed systems, is often very difficult. Often it is the case that it is unknown why a particular event should have succeeded or failed, so it is hard display relevant debugging information.
- There is a desire to be able to match a level of authorization to a level of authentication. The point was made that we, as a security community, lose creditability if we require difficult authentication for pedestrian tasks.
- Getting attribute information about the user in a federation identity setting and determining standards for citizenship (as discussed in the section on federated identity) can be difficult, as is setting levels of certainty regarding attributes.
- A common authorization factor is whether a user has signed an acceptable user policy (AUP) of some sort. The possibility of a standard AUP across grids and sites was discussed as beneficial.

2007 NSF Cybersecurity Summit – Final Report

Auditing and Accounting

The following topics related to auditing and accounting were discussed:

- A growing driver of requirements for auditing is the need to share information across sites for incident response forensics. In addition to the policy issues involved in this, there are a number of technical issues including the standardization of log formats and the semantics of the information.
- The question was raised of how one instruments a virtual organization (VO) that spans multiple sites. Since the VO comprises subsets (typically) of each site, the aggregation of information related to just the VO is nontrivial.
- Given that sites typically have different names/user IDs for the same user, matching users from logs from different sites is a difficult problem.
- It was pointed out that auditing and accounting could serve as the basis for the intrusion detection based on anomalous behavior discussed in the section on session hijacking.
- Debugging of distributed applications was also discussed as a customer for audit information.

Browsers as the User Interface for HPC

The growing use of Web browsers as a important user interface for HPC was discussed, with TeraGrid science gateways given as an example. The group suggested that educational and policies activities regarding security need to begin to take this into account, as these activities have been traditionally directed more towards command-line interfaces and do not take Web browsers into account.

The issue of lack of compatibility between Web single sign-on systems was also discussed as an issue, as well as best practices for Web applications developers in how to write their applications to be compatible with such systems.

Systems for Virtual Organizations

The difficulty in applying AAAA to virtual organizations (VOs) was discussed. In a standard organization, there typically is a match between the organization's network topology and the boundaries of the organization that is lacking a VO. This match is often taken advantage of by AAAA services, leading to difficulties when they are applied to a VO. A prime example of this is firewalls, which work well to match policies to network topology but lack support for enforcing VO policies.

A key question is how responsibility boundaries are defined in VOs – i.e., who is responsible for what and when? The question was raised whether NIST, which has developed a number of useful guidelines for computer system security, could produce similar guidelines for distributed systems. This was argued against on the grounds that we are still learning in this space today, citing the different VO models adopted by TeraGrid, OSG, and ORION, and that it is undesirable to nail down policies too soon. Instead, some path to reaching community consensus is needed.

3.4 Incident Handling/Response

2007 NSF Cybersecurity Summit – Final Report

This breakout section was lead by Abe Singer from the San Diego Supercomputing Center and Scott Campbell from the National Energy Research Scientific Computing Center. The breakout section consisted of approximately 17 people.

Overview

After an initial discussion about a number of shared technical issues surrounding general intrusion detection and incident response issues similar to last year, the group began focusing on issues involving communication, assistance and education.

The Incident Response breakout section covered several topics, focusing specifically on topics involving inter-site communication, standardizing incident response procedures, providing an organized assistance to smaller institutions, and developing a framework for cooperating on international science collaborations.

Issues and Recommendations

Extending from the survey of intrusion detection and incident response techniques developed last year, a number of issues were explored and recommendations made for developing solutions.

The general thrust of the recommendations rests with the idea of increasing inter-site communications while providing better support and education for the security community at large. This is particularly important for sites that currently lack resources and the experience to cope with the current level of compromised hosts (and security issues).

The individual recommendations follow.

(1) NSF should fund a formal inter-site notification organization.

In the current environment there is no official mechanism to share time-sensitive security-related information that crosses inter-site boundaries. We recommend that the NSF fund a currently recognized operational organization to fulfill this need rather than create an entirely new system. Specific details include:

- Use the OSG Security Response group, REN-ISAC, or the Security Incident working group at I2 as functional models.
- The current standard schema for transferring computer security information (based on IODEF/RFC 3016) is too complex and difficult to understand and implement. A subset of this should be left to describe a minimum data set.
- Define policy mandating the exchange of security data with the inter-site mechanism/NSF. As inferred above, this data exchange should be made simple to use and understand, as the skill set of potential users will vary greatly.

In addition, the following services should be provided:

- Archived mailing list(s) - both regular and PGP-encrypted for sensitive information.
- Web page for general security alerts and specific case content.
- Encrypted and authenticated chat services for real-time communication between the notification organization and individual users.

2007 NSF Cybersecurity Summit – Final Report

(2) Common incident response procedures should be created.

As described in the 2005 NSF Cybersecurity Summit report, the quality of incident response varies considerably among the representative group. We propose that NSF fund the development of standardized incident response procedures, both in terms of detailed site reports as well as a simple procedure designed to be used in the case of system compromise.

To maximize the ease of implementation and to minimize problems with cross-departmental conflicts, this development should be based on currently accepted and implemented standards throughout the large (and small) system community. As some sites may not have adequate resources to digest and implement as thorough a document as (for example) the complete NIST 800-61, we suggest the following resources:

- Base this on a *simplified* version of (for example) NIST 800-61.
- DOE has IPWAR (DOE M 205.1-C, Incident Prevention, Warning, and Response)

In addition to a general incident response guide, an extremely concise playbook should be created, which would be used by local system administration and security staff immediately during and after a compromise. The playbook would provide a list of immediate short-term actions that would prevent potential problems later.

(3) Training should be provided.

In addition to standardizing incident response policy and procedures, a series of training courses should be offered to interested system administration and security staff that focuses on incident response and forensic analysis. This recommendation is closely tied to the next recommendation, which addresses the more general problem of dealing with smaller institutions that lack sufficient training and resources to deal with ongoing security issues. While some concern was expressed regarding cost and effectiveness, the general consensus was that SANS-based classes could provide the basis for a significant portion of the training.

(4) Fund a workshop designed to solve the “small facility” problem.

While discussing current issues facing the large-facility community, it became apparent that a significant proportion of overall security issues are located in the small-facility arena. What we mean by small facility would be universities and projects that lack a computer security presence and resources to help protect their own resources. Given that these organizations still take part in collaborative research with larger sites, the problems and issues encountered by the small-facility community is a problem shared by all of us.

The issue of how to get institutions that might lack staff and resources for dealing with what would otherwise be rather mundane issues is complex and well beyond the scope of this group. This problem is something that we strongly recommend be addressed in a forum similar to this one. In particular, we would like to note that at many smaller institutions and sites there is no significant staff or resources being applied to the overall

2007 NSF Cybersecurity Summit – Final Report

security environment. As long as this is the case, such institutions represent a significant opportunistic threat not only to other NSF facilities but also the Internet at large.

(5) Develop an agenda for increasing international security cooperation to support international science.

International science collaborations such as ITER and the HDC project require a tremendous degree of cooperation between the involved organizations. While organizational cooperation between the different research branches has advanced steadily, there has been little or no improvement in communications between the security groups of these different organizations.

An example of this is that security data from events at CERN tend to propagate to U.S. computational facilities via high-energy physics mailing lists. While this is useful, it would be much more efficient to have tools and agreements in place so that the computer security groups could exchange information without worrying about (for example) local privacy rule violations.

We propose that a workshop be funded that addresses the impact of security issues on global science. It would not only be designed to answer basic questions regarding how to respond to international security issues but also would be a platform to develop better relations and communications between the different organizations involved in these collaborative efforts.

Given the cross-departmental nature of this research, we suggest involvement by Internet2, ESnet, FIRST and their EU counterparts, and the OSG/HEP community, which seems to be grappling with many of these issues already.

(6) Focus security efforts on high-risk/high-impact threats.

In both the Incident Response discussions, as well as the presentations, a number of topics arose that are more disciplinary in nature, but describe feedback for the general NSF program.

The following observations were made regarding the changing nature of incidents:

- Attackers seem to be better organized and have better resources. Their motivation is now financial gain more than notoriety. This is embodied by the movement of organized crime into large-scale computer crime.
- Credential theft and insider attacks have become serious problems.
- Counterintelligence is increasingly used in determining how our security systems operate. This includes watching DNS activity and seeing what behavior causes a scanning host to lose access.
- Distributed denial of service does not represent a significant threat to the research community at large. This is a generalization, but there was no indication that the same threat shared by many commerce sites was even seen by the representative audience.

3.5 System Administration Practices/Policies/Education

The session on System Administration Practices, Policies, and Education was well attended. We had a cross-section of systems administrators and cybersecurity specialists from larger sites as well as those with a more dedicated focus on a specific project. The discussions included sharing of practical solutions for day-to-day problems, user and data management, user agreements, configuration management, system monitoring, and larger policy issues.

During the discussions it became clear that even though the specifics of our sites and facilities are different, we have many similarities as well. It would be very useful if we had a place to share solutions that we have found and be able to ask questions in hope that others have found solutions used elsewhere. This breakout, and the hallway discussions that continued long after the actual sessions ended, were found to be very useful.

One of the biggest discussion areas surrounded users, accounts, and data management. A key observation from the system administrator's point of view is to keep the user community happy and productive. This can raise a number of issues with maintaining the user facility. The variations in providing services to primary investigators and their core team of collaborators to those of the important visiting scientist all need to be considered. Each site has its own set of problems and solutions. It is clear that having a documented policy known to all users and followed by the site administration and operation staff goes a long way to solving these issues. The area we see of growing concern is that of access by non-U.S. citizens. The list of things on the export control list is growing, and in some cases may be related to some of what is happening at our sites. It would be useful if NSF could provide guidance on where to find the compliance information related to the open science we are working on.

There is concern from the systems administrators groups on how Web-based applications may be used. We recognize that many of the grid-based applications provide a way to map usage down to individual remote users. There are some projects that need to make improvements in these areas. It's also the case that many of these middleware tools are too complex for both the users and administrators to easily understand and use.

Configuration management issues were discussed. It was clear that most sites had something that worked for them, but would like to have something better. Most sites have clear build and install procedures. Patch management also appears to be something most sites have under control. There was a discussion on how sites monitor for changes in the configuration. Tools like ganglia, tripwire, aide, bcfg2, and redhat enterprise server were discussed. It is clear that more work can be done in this area.

Data management is a big concern at many sites. As computing capabilities continue to grow, the amount of data needing to be maintained is growing as well. We recognize this

2007 NSF Cybersecurity Summit – Final Report

as a problem and wish to convey to NSF that data as well as computing capacity needs to be recognized. A number of topics were discussed on data integrity, archiving, backups, and confidentiality. As with user management, it was clear that well-publicized site policies on these issues are important.

We recognized that several of the recommendations below were also discussed in part at other workshop breakout sessions. We add our voice to the collected set to emphasize the need for action in these areas.

We commend NSF for recognizing the many differences in the projects to which they provide funding and in putting the responsibility for the security of these projects with the site. We would like to see more specific guidance from NSF on how to produce the security plans. Since NIST already provides standards used by other agencies, a simple recommendation that sites should consider following these standards would be a start. This does not require sites to change their existing documentation, but does provide a template to work from. We also recognize that as sites collaborate and become more interconnected, working from a common base to provide security documentation will make it easier to validate that all sites in the collaboration are working from a similar security posture. A long-term goal would be to have a standard that could be used at all sites.

It was recognized that Secure Sockets Layer (SSL) certificates are becoming more important to the general operation of a site. Many sites currently use self-signed certificates because of the costs involved. It would be beneficial if NSF could provide a CA service or negotiate a discount that NSF-sponsored projects could use to obtain certificates from a well-known provider.

It would be useful if NSF, like NIST, were to provide a location for best-practices documents. Many sites are running into the same problems and reinventing solutions. Having a place to find how others have solved these problems would benefit all. Options discussed included mailing lists with searchable archives or a wiki for site comments or discussions.

We were also glad to hear during the opening day's presentation that the NSF has a Cyber Training program for all employees. A number of sites indicated their training may be lacking in some areas or falling behind and in need of updating. It was agreed that it might be useful if we could obtain copies of the training used at NSF for their people. Even if this did not apply directly to our individual sites, it should provide a template on how to construct or update what we have. We discussed having this on the proposed wiki as well.

All of the attendees were concerned about personal identification information (PII). There is concern on exactly what constitutes PII and how it needs to be protected. We deferred most of the discussion to the breakout section on data identification and classification. We would like to see guidance from NSF on these topics.

2007 NSF Cybersecurity Summit – Final Report

We found this meeting very useful on multiple fronts. We found ways that we could make recommendations to NSF on how they can provide useful guidance back to our sites. We also found that we could share solutions for problems we have in common. We would like to see this meeting held in future years.

4. Conclusions and Final Recommendations

In conclusion, this year's summit seemed to be an overall success. There was excellent participation from a range of different NSF-funded large-scale facilities and some good recommendations from the breakout sessions. There is an overwhelming consensus to have the NSF sponsor another cybersecurity summit next year. However, next year's chair and program committee will need to determine what the focus should be so that it does not just duplicate the last three cybersecurity summits. As the breakouts for this year were being developed, the program committee and the session leaders wanted to make sure the breakouts were not just rehashing the previous two summits' breakout sessions. One of the breakouts mentioned that not much has changed in the last couple of years, so they were hard pressed to come up with new actionable items in their session. Something to consider is who the next cybersecurity summit should target. One of the recommendations from a breakout is that the problem of security in small facilities is an area that needs addressing, as well as the move toward international science projects.

One of the other requests expressed during this year's summit was the desire to see the NSF start addressing some of the previous year's recommendations. In the introduction and the closing this year, the NSF representatives also expressed their desire to take action on those recommendations. They also said that even though it may not look like anything is being done, there is progress at the NSF on those items. As an organization, the NSF appears to move slowly. These changes are making their way through the proper channels so that appropriate action can be taken in the correct programs.

From the recommendations above come a number of actionable items for the NSF to consider, as well as sites to implement at their own facilities. The NSF and remote sites can take this year's report, as well as the previous years' reports, and gauge how well they are doing to meet these recommendations as they work to improve the overall security of their respective organizations in the coming years.

5. Participant Evaluation Summary

EDUCAUSE will be providing the evaluation summary. Contact Rodney Petersen of EDUCAUSE for the summary.

Appendix A Program

Thursday, February 22, 2007

8:30 a.m. - 9:00 a.m.

General Session: Welcome and Introduction

9:00 a.m. - 9:50 a.m.

General Session: The Commercial Malware Industry

- Peter Gutmann, Security Researcher and Professional Paranoid, The University of Auckland

9:50 a.m. - 10:30 a.m.

General Session: Legal Perspectives on Data Security

- William Cook, Partner, Wildman Harrold

10:30 a.m. - 10:50 a.m.

Refreshment Break

10:50 a.m. - 12:15 p.m.

General Session: Real-World Data Applications from Different Perspectives

- Don Petravick, Head CCF Dept., Fermi National Accelerator Laboratory

- Frank Siebenlist, Security Architect, Argonne National Laboratory

- Dane D. Skow, Deputy Director, NSF TeraGrid Project, Argonne National Laboratory

12:15 p.m. - 1:15 p.m.

Lunch

1:20 p.m. - 2:00 p.m.

General Session: Information Systems Under Attack: Managing Enterprise Risk

- Ronald Ross, Project Leader, FISMA Implementation Project, National Institute of Standards and Technology

2:00 p.m. - 2:05 p.m.

General Session: Breakout Session Overview and Guidelines

2:10 p.m. - 5:00 p.m.

Breakout 1: Policies, Standards, and Guidelines

Breakout 2: Identification and Classification

Breakout 3: Authentication, Authorization, and Accounting (AAA)

Breakout 4: Incident Handling/Response

Breakout 5: System Administration Practices/Policies/Education

2007 NSF Cybersecurity Summit – Final Report

Friday, February 23, 2007

8:30 a.m. - 9:10 a.m.

General Session: The Law Enforcement Academic Alliance Program
- Mike Butler, Assistant Section Chief, CI Strategy/Domain Section,
Counterintelligence Division, Federal Bureau of Investigation

9:10 a.m. - 10:15 a.m.

Breakout 1: Policies, Standards, and Guidelines

Breakout 2: Identification and Classification

Breakout 3: Authentication, Authorization, and Accounting (AAA)

Breakout 4: Incident Handling/Response

Breakout 5: System Administration Practices/Policies/Education

10:15 a.m. - 10:30 a.m.

Refreshment Break

10:30 a.m. - 12:00 p.m.

General Session: Breakout Session Reports

12:00 p.m. - 12:15 p.m.

General Session: Concluding Remarks