Final Report of the Workshop

2005 Cybersecurity Summit Tysons Corner, VA December 12-13, 2005

2005 Cybersecurity Summit Program Committee

James J. Barlow	National Center for Supercomputing Applications
Tom Bettge	National Center for Atmospheric Research
Mark S. Bruhn	Indiana University
Craig Foltz	National Science Foundation
Victor Hazlewood, Chair	San Diego Supercomputer Center
Jody Malik	Lawrence Livermore National Laboratory
Jeffrey C. McCabe	Texas A&M University
Corbin Miller	NASA Jet Propulsion Laboratory
Andrea Norris	National Science Foundation
Rodney J. Petersen	EDUCAUSE
Peter M. Siegel	University of Illinois of Urbana-Champaign
Howard Walter	Lawrence Berkeley National Laboratory

Table of Contents……2

Executive Summary

The 2005 Cybersecurity Summit workshop brought together a broad range of people from the research and academic community, who are mainly sponsored by NSF, to focus on the security issues that are shared by this community and is a continuation of the work from the 2004 Cybersecurity Summit workshop brought about by the security incident events of 2003 and 2004. The workshop was sponsored by the National Science Foundation (NSF) and over 110 professionals representing systems, network and security administrators, center management and staff, and others participated in an invitation-only workshop held in Tyson's Corner, VA in December 2005. The main goals of the workshop included

- Sharing of Information and Ideas
- Understanding our Communities Diverse Perspectives
- Discussion of Our Communities Strengths and Weaknesses
- Identify Our Community Security Needs

The loosely held theme of the workshop, *Data Security*, was represented in the keynote presentation and one of the breakout sessions. Five breakout sessions were selected for specific and focused discussion in breakout groups among the participants. The breakouts were

- Policy, Standards, Guidelines and Security Plans
- Awareness and Training
- Intrusion Detection and Incident Response
- Data Integrity Assurance in the Research Environment
- Authentication

As a result of the discussions, each breakout developed a set of observations, conclusions and recommendations which are detailed in this report. The highlights of these from the breakout sessions include:

Breakout Session 1: Policy, Standards, Guidelines and Security Plans

- recommendation for an additional paragraph that could be added to the guidelines for large sites and FFRDCs that addresses security for NSF awardees.
- A recommendation for clarification of several issues regarding the upcoming NSF security guidelines/requirements.
- Concerns of NSF awardees potential requirement for reporting of security incidents and the associated concerns regarding Freedom of Information Act inquiries.

Breakout Session 2: Awareness and Training

- Promote and leverage existing educational materials and opportunities
- Encourage NSF to be more aggressive in providing security awareness assistance (e.g., Guidelines for IT Security of NSF's Large Facilities).

- Encourage institutions to include technology support (IT Security) in grant proposals

Breakout Session 3: Intrusion Detection and Incident Response

- There is no silver bullet for security of NSF sponsored projects.
- Encourage and support a Defense-In-Depth security strategy for NSF projects.

Breakout Session 4: Data Integrity Assurance in the Research Environment

- User awareness and education needs to be improved regarding data integrity assurance issues in the research environment
- Organize a workshop(s) specific to data security and integrity, perhaps with NSF sponsorship.
- The evaluation and specification of a data security service.

Breakout Session 5: Authentication

- OTP deployment has largely stalled due to concerns regarding conflicts with HSPD-12 and usability issues if large numbers of sites rolled it out in an uncoordinated manner.
- It was observed and stated that authentication is an arms race.
- Several specific authentication issues were discussed and recommended for further research

Additionally, Educause was employed this year to coordinate and support the workshop. This contributed to the smooth operation and success of the workshop by employing a professional organization to handle logistics, provide support to the program committee and providing web-based tools for program committee support and the web hosting of the program, registration and survey.

1. Overview

1.1 Motivation

Following on the work done in the 2004 Cybersecurity Summit workshop, the NSF sponsored the 2005 Cybersecurity Summit workshop held in Tyson's Corner, VA on December 12-13, 2005. This workshop was intended to go beyond the focus of any particular incident. The workshop was by invitation only and was intended to bring together stakeholders from the research centers and academic community sponsored by or related to the NSF supported community to discuss security related topics. The goals of this workshop were to

- **Share Information and Ideas:** By sharing information and ideas the community can understand the common issues and problems that affect security in the research community and 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, 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. Discuss and analyze the strengths and weaknesses related to security of these environments.

- **Identify Our Community Security Needs:** 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. Strive for 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

The NSF enlisted the help of Victor Hazlewood of the San Diego Supercomputer Center (SDSC) to form and chair a program committee that represented the nation's diverse research and academic community, NSF's interests, and some representation from other federal agency centers. The program committee met regularly on Thursday from September 15, 2005 to December 8, 2005 by teleconference. With the support of EDUCAUSE and the dedicated work by the program committee, this represented the right amount of time spent to adequately prepare for the conference and not miss or leave any details for the last minute. This led to a very successfully organized and planned workshop. The EDUCAUSE staff made a program committee website available whereby

documents could be shared, minutes posted and tasks tracked. The program committee members were:

James J. Barlow	National Center for Supercomputing Applications
Tom Bettge	National Center for Atmospheric Research
Mark S. Bruhn	Indiana University
Craig Foltz	National Science Foundation
Victor Hazlewood, Chair	San Diego Supercomputer Center
Jody Malik	Lawrence Livermore National Laboratory
Jeffrey C. McCabe	Texas A&M University
Corbin Miller	NASA Jet Propulsion Laboratory
Andrea Norris	National Science Foundation
Rodney J. Petersen	EDUCAUSE
Peter M. Siegel	University of Illinois of Urbana-Champaign
Howard Walter	Lawrence Berkeley National Laboratory

This committee developed a program to cover many aspects of security for the research and academic community and with a loose theme of *Data Security*. This was different from the incident focused 2004 Cybersecurity Summit. The program developed by the committee is included in the Appendix.

1.3 Participation

Participants were invited from the nation's research centers, universities and other locations that the program committee and NSF program managers felt would significantly add to or benefit from the discussions of the workshop. Invitations were issued to a broad collection of people mainly, or in some form, associated with NSF sponsored research. A few representatives from other federal agencies including law enforcement were pursued. As with the composition of the program committee, the invitations were issued to a broad audience in order to have participation from system, network and security administrators, project managers, and center management.

There were 110 people who participated in the workshop. Out of the total number of attendees, 108 were from the United States, while two were from other countries: one from Chile and one from Japan. There were 35 attendees (32.4%) from the Virginia/Maryland/DC area. California had the second-largest number of attendees (18; 16.67%), followed by Colorado (10, 9.26%).

Additional demographic information was requested as part of an online survey. There were 95 individual responses to the survey, compared to the 110 actual attendees. It should be noted that respondents were allowed to select more than one option for the first 2 questions (see Figures 1 and 2).

Figure 1 shows the area of science most closely related to the participants' job or interest. Close to 40% of the participants have a job related to or an interest in the Office of Cyberinfrastructure (OD/OCI). Nearly 15% of the participants have a job related to or an

interest in Geosciences – Atmospheric Sciences (GEO/ATM), while over 10% have a job related to or an interest in Math and Physical Sciences – Astronomical Sciences (MPS/AST). Of the participants surveyed, 12% selected "No direct science area".

Figure 1. Area of Science Most Closely Related to Participants' Job/Interest

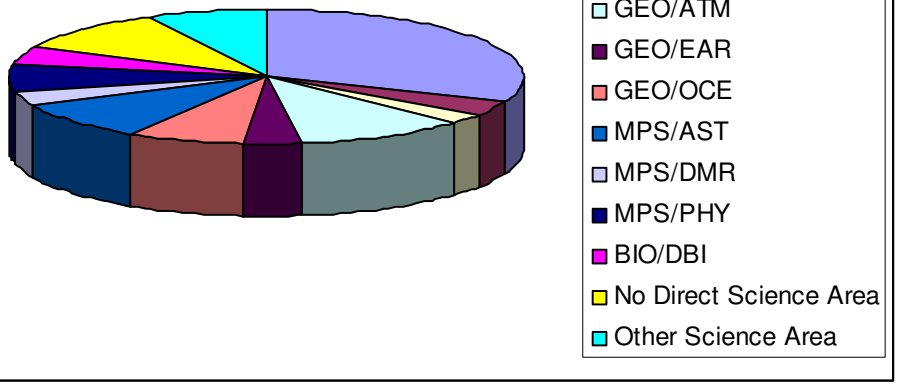

Figure 2 shows the function most closely related to the participants' job or position. Almost 50% selected Network or Computer Security Engineering, while 40% selected IT Security Management and nearly 35% chose IT Security Policy. Close to 30% of the participants selected Facilities Operation and Management as most closely related to their job, and about 20% selected Facility User.

Figure 2. Function Most Closely Related to Participants' Job/Position

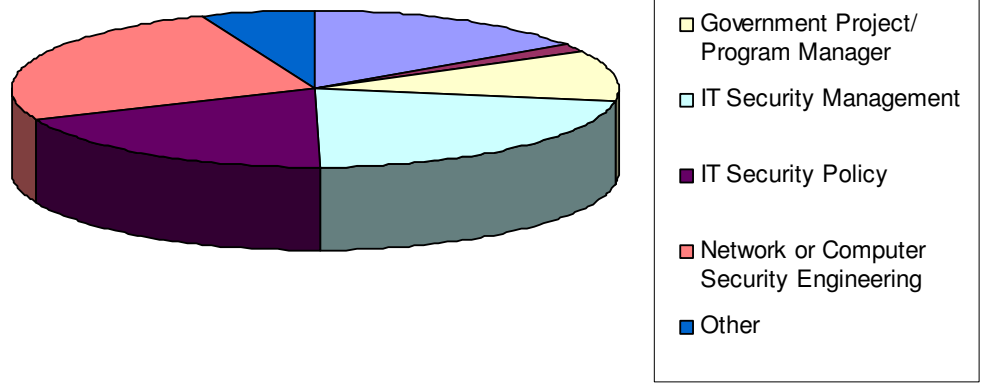

Figure 3 shows the distribution of attendees by agency. Almost 40% of the participants represented an academic institution. Nearly 20% represented an NSF large facility; 11% represented NSF. The Department of Defense, the Department of Energy, and a few additional government agencies were also represented.

Figure 3. Participant Distribution by Agency

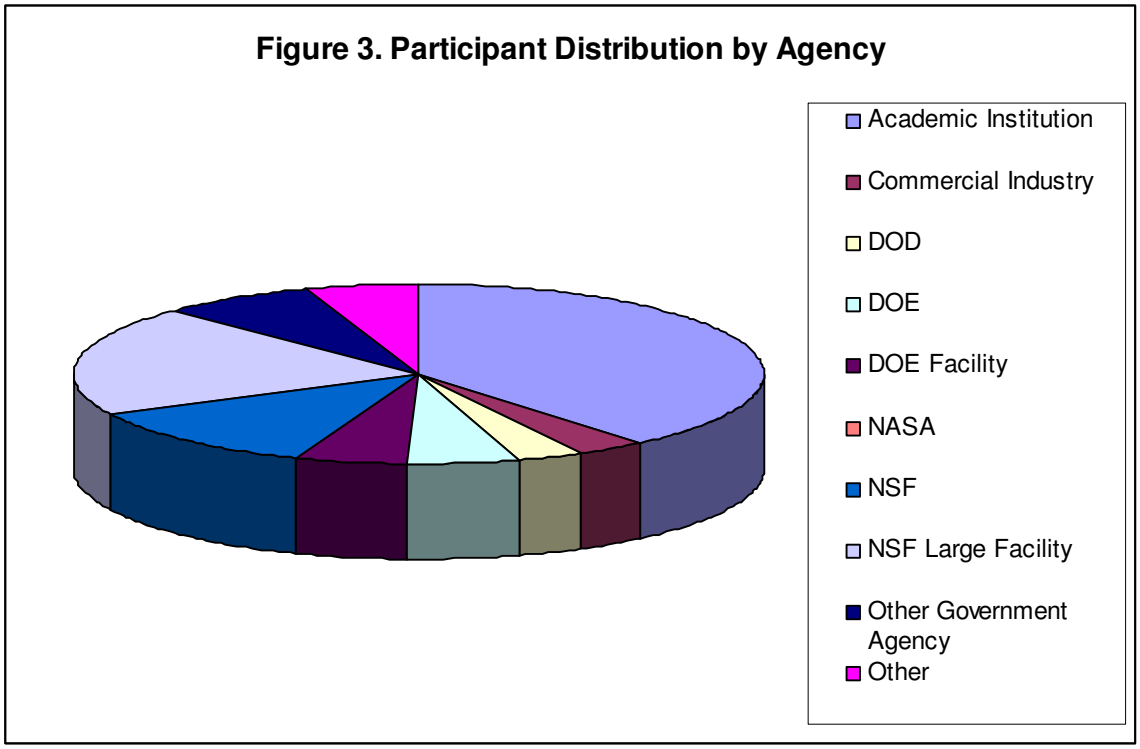

2. Plenary Sessions

The entire 2005 Cybersecurity Summit program is included in the Appendix and the plenary session presentation materials are available on the summit website at http://www.educause.edu/cyb05. The following provides a brief overview of the plenary presentations.

Victor Hazlewood, the program chair, gave an introductory welcome message to the workshop and introduced Margaret Leinin, Assistant Director of Geosciences, who gave an NSF welcome message and framed NSF's perspective of the workshop. She emphasized that security is taken seriously by NSF – through direct funding of cybersecurity research, through collaboration and cooperation between the NSF grantees (as demonstrated by the CSS workshops), through following government guidelines within the NSF itself, and in the strategic planning process. She invited the 2005 Cybersecurity Summit attendees to provide comments to the NSF Strategic Plan for Cyberinfrastructure, which is currently under development and will include security as a theme embedded within the plan.

Victor Hazlewood and Jim Barlow gave a review and current status of the Case 216 security incident that brought the 2004 Cybersecurity Summit workshop together. In

addition, they gave a message to put the 2005 Cybersecurity Summit in a transitional context. The context was for this workshop to go beyond the specific security incident from 2004 and to broaden the scope of talks and discussions to include all aspects of security in the research laboratory and academic environment.

The workshop had a loose theme of *Data Security* and Reagan Moore gave an interesting presentation on security requirements for shared data collections from the perspective of the research community. As data management requirements and complexity in collaborative science explodes, Dr. Moore pointed out the security implications in collaborative data management in illustrative Storage Resource Broker examples. Security requirements and issues discussed included authentication, trust virtualization, access control, data integrity of terascale and petascale collections, replication, and multi-site interoperability.

Peter Siegel, moderated a panel discussion on architecting security into research projects. The panel was made up of George Strawn, Von Welch and Scott McCauley who represented the different perspectives of security for research projects of program management/agency administration, security professionals, project staff/principal investigator, respectively. The panel gave their introductory comments in regard to the topic and, through discussion and question and answer, explored the challenges of including security requirements into the mix [domain] of preparing, winning and executing a large research project. Key issues included: the growing importance of addressing security in the proposal submission, program execution and program review stages; the need for a vetted security plan within research programs; the importance of effective lab/campus-level training on security issues, appropriate to PIs, researchers, and technical staff; and appropriate partnerships between research labs and campus security groups.

Ron Ross of the National Institute of Standards and Technology gave a very well received presentation on how the research and academic community can take advantage of the standards, guidelines and controls spelled out by the NIST publications that are being developed as a result of FISMA requirements for the federal government. Though FISMA does not apply to the systems operated by awardees of NSF sponsorship, the documents cover many of the security issues of interest to the research community. The NIST project includes the development of security categorization standards; minimum security requirements standards; and guidelines for the selection of minimum or baseline security controls for information systems, assessing the effectiveness of security control, and the security certification and accreditation of information systems. The talk covered comprehensively how the research and academic community could take advantage of these documents and use them in a comprehensive, risk-based enterprise information security program.

On the second day, Jim Marsteller gave the day's only plenary talk which covered an overview of the people and processes involved in Teragrid incident response. In his presentation he gave the background for the forming of the Teragrid Security Working Group and the people, policies, procedures and documents put together by this group. He

detailed the incident response related contact list, secure communications capabilities, processes, documents, and teleconference calls created and maintained by the Teragrid Security Working group. Some of these capabilities are documented on the Teragrid Security Working group website, http://security.teragrid.org/ .

3. Breakout Session Summaries

Five breakout sessions were selected by the program committee for specific and focused discussion in breakout groups among the participants. The breakouts were

- Policy, Standards, Guidelines and Security Plans
- Awareness and Training
- Intrusion Detection and Incident Response
- Data Integrity Assurance in the Research Environment
- Authentication

Description of the breakout sessions are detailed in the following sections.

3.1 Breakout Session 1: Policy, Standards, Guidelines and Security Plans

This breakout was lead by Abe Singer of SDSC and David Seidel of Purdue University. The breakout group consisted of about 50 people who met over two days of the workshop.

This breakout group's initial goal was to discuss general issues of policies, standards, and guidelines in a research environment. However, as the group started discussion it found a common interest in the proposed NSF requirements for security programs for large facilities and FFRDCs. Approximately 50 people were in attendance, including several NSF program officers, which made for a very good interactive discussion between all participants. There was much concern over the proposed language changes. Many participants believed that the NSF would be mandating particular security requirements on them. The result of the breakout session was suggested changes to the proposed language, identifying information and issues that need to be communicated to the NSF and/or awardees.

The initial language proposed by the NSF is in the FATC supplement, and is as follows:

> The awardee is responsible for all information technology (IT) systems security and associated equipment and information, funded directly or indirectly by this award. The awardee shall present to the cognizant NSF Program Officer and Grants and Agreements Officer a written plan addressing policies and procedures for review and approval within 60 days of award.

The group proposed adding the following text to the paragraph above:

The plan shall describe the information security program appropriate for the project, including but not limited to roles and responsibilities, risk assessment, technical safeguards, administrative safeguards, physical safeguards, policies and procedures, and awareness and training. The plan should include evaluation criteria that will measure the successful implementation and deployment of the plans, policies and procedures.

3.1.1 Issues and recommendations

In addition to the proposed language change, the group noted several issues about the requirement that needed to be clarified:

- The requirements are simply an extension of the business plan that is currently required of awardees. The initial reaction by many was that the requirements were an entirely new set of hurdles to overcome, but NSF program managers pointed out that a business plan is already a requirement, and these requirements would be treated in the same manner.
- The requirements are for Large Facilities and FFRDCs only, not for all NSF awardees
- The goal of the requirements are to goal is to engage the program manager in a dialogue with the PI, awardee, research office, and local organization security in order to put together a security program that is satisfactory.
- There was wide concern that they are not guidelines, but mandated set of requirements or checklist for compliance. The proposed language changes are intended to reflect this concern.

Some issues that still need to be addressed:

- How is the NSF to communicate back to awardee on guidelines and program requirements? Possible suggestions were outreach to PIs and contract offices.
- How should sites manage notification of security incidents back to NSF? There was consensus that not all incidents should require notification, in which case, some guidelines need to be provided which give thresholds for notification. Suggested thresholds suggested were:
    - When business continuity is affected or potentially affected
    - Potential impact on community as a whole
    - Likelihood of bad PR and political/reputation ramifications
- There was also some concern about the confidentiality of information that is reported to the NSF. Many sites have legal confidentiality requirements, in addition to institutional privacy policies. It is not clear whether information provided to NSF becomes public information, or whether it is subject to FOIA requests. This issue should be clarified by NSF so that sites can properly comply with all requirements.

3.2 Breakout Session 2: Awareness and Training

The breakout was lead by Mark Bruhn of Indiana University and Corbin Miller of the NASA Jet Propulsion Laboratory. The breakout group consisted of about 8 people who met over two days of the workshop.

3.2.1 Methodology

The breakout group oriented the discussion toward these outcomes:

Identification of Constituencies
Identification of Challenges
Inventory of Existing Programs and Products
Identification of Gaps
Identification of Gap-Fillers
Recommendations

3.2.2 Constituencies

We enumerated constituencies that are involved in some form with research and science. The list is long (of course) and not necessarily exhaustive (of course):

- Researchers
    - Scientists
    - Research Faculty
    - Research Assistants
    - Graduate students
    - Undergraduates
- Funding agencies
- Institutional Review Boards/Human Subjects Committees
- Visitors / affiliates
- Faculty
- Librarians
- Students (resident versus non-resident)
    - Undergraduate
    - Graduate
    - Teaching Assistants
- Administrators
    - Senior executives, CIO -- decision makers
    - Policy/compliance officers
    - Staff, employees, email users, basic users
    - Power users (tinkers, meddlers)
    - Data custodians
    - Auditors
    - Archivists

  - Human resources
  - Student affairs
- Technicians
  - Security Professionals
  - System administrators
  - Database administrators
  - Network administrators
  - Web administrators
  - Helpdesk/support staff
  - Programmers (Coders)
- Guests/Visitors/Transients
- Collaborators
  - Onsite
  - Visiting
  - Members of existing community
- Remote push/pull
  - Local
  - Regional
  - National
  - International
- Private service partners
  - Contractors
  - Vendors
  - Consultants
- Law enforcement
  - Internal
  - External
- University services
  - Outreach
  - Alumni

Existing Offerings and Opportunities

The group spent considerable time thinking about and identifying existing education and awareness, and collaborative, offerings, along with their general orientation/audience:

EDUCAUSE/Internet2 TF Security Education/Awareness Working Group
    Chief Information Officers and IT Professionals
National Cyber Security Alliance
    General Student Body
Centers for Excellence in Information Assurance Education (CEIAE)
    Variety (~62) of recognized security programs
    Curriculum development
    Self-paced training for IT Professionals
    Self-paced training for Researchers?
Colloquium for Information's Systems Security Education (CISSE)

- Faculty Boot Camp
- SANS Institute
  - Technicians
  - Certifications
- USENIX
  - Graduate Students
  - Computer Science Faculty
  - Opportunities for Training (cont)
- Institute of Electrical and Electronic Engineers (IEEE)
  - Graduate Students
  - Engineering Faculty
- Association of Computing Machinery (ACM)
  - Special Interest Group on Security, Audit, and Control (SIGSAC– o)
  - Online digital references and journal
  - Computer Science Faculty
  - Students
- Vendors
  - Certifications for IT staff
  - Free training for faculty
- Open Courseware Initiative (give and take)
  - Source for Curriculum
- Government online training (NIH, NSF, NOAA, etc.)
  - NSF Annual Security Awareness Training
  - Administrative staff
- National Security Telecommunications and Information Systems Security Committee (NSTISSC)
  - Curriculum Standards
- Information Systems Audit, and Control Association (ISACA)
  - Audit and security professionals
  - Certifications
- Information Systems Security Association (ISSA)
  - Information security professionals

3.2.2 Challenges

Inasmuch as the list of participants is long and the time was short, the group concentrated on discussing challenges related to the single audience of *research scientists and research faculty.*

We agreed that this group above others is hard to reach. They are typically very focused on their particular science and individual research projects. They generally fear barriers to their research activities and goals. This is a community in which openness is part of the culture, and security has not traditionally been considered a necessary part of the activity -- though it is interesting that those researchers are very (and rightly) concerned about the integrity of their data and intellectual property value of their research product. They do not traditionally think broadly about the impact that lax security of their systems

has on unrelated local or other systems, nor do they think about the potential impact of lax security on systems outside of their organizations on systems supporting their projects. Researchers, especially those outside of engineering or other technical sciences, have limited access to formal and trained technicians. Most often, that support is provided by graduate or even undergraduate students, or by the researcher him/herself

Researchers are subject to sometimes conflicting requirements or specifications affecting their activities, from external funding bodies, local campus or organizations, classified research sponsors, and others.

3.2.3 Fundamental Recommendations

There are many approaches to increasing (or attempting to increase) the awareness about security issues. They are being discussed and addressed in a variety of other venues and activities. However, for the particular base community being addressed – scientists and researchers – the discussion seemed to come back to a set of approaches that might be the most effective;

- Ensure that applicable aspects of security are considered at the institutional level
    - In the deliberations and considerations of Institutional Review Boards and Human Subjects Committees and Research Compliance Committees
    - In employee job descriptions
    - In faculty, staff, and student orientation sessions
    - Along with other already-required compliance training
- Find and engage external organizations that have influence over and/or the respect of the research community:
    - Higher education Presidential associations
    - Professional organizations
    - Academies
    - Accreditation boards
    - NSF and other funding agencies
- Promote and leverage existing educational opportunities (listed above).
- Encourage NSF to be more aggressive in providing security awareness assistance (e.g., Guidelines for IT Security of NSF's Large Facilities).
- Encourage institutions to include technology support (IT Security) in grant proposals, especially graduate students (future researchers).

3.3 Breakout Session 3: Intrusion Detection and Incident Response

The breakout session was led by Jim Barlow of NCSA and Stephen Lau of LBNL. The breakout group consisted of approximately 25 people.

The Intrusion Detection and Incident Response breakout session covered several topics, including host based and network based intrusion detection systems (IDS), intrusion prevention systems (IPS), network flow tools, incident response and research topics.

There were over 25 people who participated in this breakout session, representing over 20 different organizations. When surveying participants regarding the tools in use at their sites, it became apparent that many sites utilized common tools and techniques for IDS/IPS. More information about these tools, additional tools discussed, and also information resources can be found in the Appendix.

Most sites had some sort of IDS/IPS/HIDS in place. A key point was made that an IDS should not be considered a complete solution, but is one tool in a suite of tools that needs to be deployed to protect a site. IPS's were a new area covered this year. The state of IPS's last year was either not mature enough, or too cost prohibitive for sites to look at deploying. Sites this year are still reluctant to deploy an IPS because of cost or because of the potential impact on performance and reliability.

Several open areas were brought up in regards to IDS/IPS. These included IPv6 support, the fact that many attackers are encrypting their traffic, the ability for systems to handle high rates of network traffic, proper handling of dropped traffic, and whether or not an IDS/IPS fails closed or open.

Data correlation was an area that was explored in the session. This was something new from last year when sites were not discussing correlating their log data for incident analysis. Many sites had multiple sensors deployed, such as a centralized syslog host, network flow monitors, and multiple IDS/IPS. Some sites were just beginning to explore correlating this data to handle incidents and detect intrusions. There was also talk about correlating data from multiple sites.

3.3.1 Incident Response

Incident response varied from site to site. About half of our participants had formalized Incident Response guidelines and procedures in effect with some of these guidelines developed in response to Case 216. Sites reported that incident response guidelines and procedures assisted incident handling. Documentation of incidents was also reported as useful for handling subsequent incidents and also for historical analysis. Documentation methods, however, varied from site to site. Some sites utilized formal trouble ticket systems, such as RT while others used more freeform tracking such as a Twiki. Encryption of incident data was also considered to be essential.

Most sites ignore the general everyday scanning against their site, and did not consider those as "incidents". When responding to a crosssite incident, participants relied upon trust relationships previously developed, and tended to work with people they were familiar with, such as those within the same project or virtual organization.

Many sites had the ability to securely communicate incident information amongst their security staff. Encrypted email was used as well as instant messaging services, such as encrypted Jabber servers (which several sites use). Some virtual and real organizations had dedicated 24x7 contact methods, such as a telephone number or an email address that is monitored.

Another group of issues brought up during the session revolved around communication and response. Many sites reported difficulties with sites and countries that were essentially "black holes". Language barriers also proved problematic in some instances. Cross site/agency incident response coordination was still somewhat ad hoc in that incident responders relied upon previously developed trust relationships at other sites.

Who to report to was also brought up as an issue facing incident responders. In many cases it was unclear if a site should contact a user directly, or a system administrator at the site the user was coming from, or both. In cases where a user used a compromised host to access another site, it is unclear whose responsibility it is to inform the remote site of their compromised system. Should the user or security staff?

The impact of regulations and requirements on incident response was also discussed in our session. Some sites had formal chain of custody procedures for evidence preservation. It was brought up by FBI attendees that sites were not necessarily held up to the same evidence collection procedures as law enforcement. The FBI recommended that sites should consider creating MD5 hashes or checksums of data collected for evidence needs. Sites should also consider their own site requirements for data preservation outside of law enforcement needs.

Some participants fell under HIPPA, DMCA, and Freedom of Information Act which impacted their incident response procedures. Some participants were concerned that although their responsibilities did not include handling information impacted by regulations such as HIPPA they were unsure if some of the projects utilizing their resources were handling this type of data. One of the suggestions included requiring reporting of data that may be impacted by regulations such as HIPPA.

3.3.2 Open Research Areas

The group also discussed open research areas for IDS/IPS/HIDS and incident response. These ideas were mainly brought up since an indepth discussion of them was beyond the scope of the breakout session. The main ideas are listed as follows:

- Distributed intrusion detection systems
- Virtual IDS for overlay networks, such as a virtual organization.
- Correlation of multiple sensors
- IPv6 and NIDS
- Proper failure of IDS/IPS.
- High speed intrusion detection/prevention systems.
- Placement of IDS and security in depth.

3.3.3. Key issues and recommendations

Realizing that *no one recommendation is going to be the silver bullet for site security*, a defense in depth, or layered security approach needs to be taken. Sites have to balance the

limited monetary resources they have, along with their personnel resources, to choose what security measures can/should be applied. Because of the diversity of sites, even within the 20+ organizations represented, it was recognized that each individual site needs to evaluate what security measures are best for them. Sites with a strong control of host end systems may be able to focus on HIDS, whereas other sites may have to take a more network centric approach with NIDS or flow analysis. For each site it is very important to "Know thy network".

Regardless of the approach, the following items can be investigated. Keep in mind when reading these recommendations that sites are different with different needs. Most sites are not going to be able to do every one of these, so will have to choose which ones they are able to deploy.

- Intrusion detection systems were determined to be very useful, and sites not deploying them should evaluate utilizing them.
- Flow tools are useful as a complement to an intrusion detection system and sites should start collecting and analyzing flows.
- Syslog data is useful as a host intrusion detection system and sites should set up a centralized syslog server.
- Data correlation from different sensors is proving to be a valuable tool (such as syslog and IDS data). There are a couple tools mentioned in the appendix that are worth looking into.
- Sites need to have an Incident Response plan and procedure in place. This includes knowing ahead of time who is involved on the incident response team.
- Sites should have an out of band communication method (i.e. encrypted email, Jabber servers) which was found very useful in incident response.
- Sites need to know requirements and laws they fall under.
  Reporting needs, data collection procedures, legal aspects
- Sites need to know their researchers and resources.
  Assists in incident response and prevents incidents.

3.4 Breakout Session 4: Data Integrity Assurance

The breakout session was led by Tom Bettge of NCAR and Victor Hazlewood of SDSC. The breakout group consisted of approximately 10 people.

3.4.1 Overview

The research cyberinfrastructure contains a wide variety of data stored on computing systems, SANs, archival systems, and data collection services. Access to this stored data requires authentication and access control. Data integrity assurance involves the mitigation of the risk of loss of data from a variety of sources including disasters, media corruption, software and hardware failures, operational errors and unintentional and/or unauthorized activity. These data integrity and security assurance issues were examined in detail, and the discussion resulted in three broad recommendations.

3.4.2 Research Data Defined

Generally speaking, malicious activity during 2005 has centered upon theft of personal data – social security numbers, identification materials, financial records, etc. The awareness of data security has accordingly risen, and a number of research information technology professionals have justifiably questioned the state of the security and integrity of data within the research community. While research data is not as high profile as personal data (which, when stolen, can be used for criminal activities), it is nevertheless an important and vital component of the nation's scientific research agenda. For example, the recent report from the National Science Board on *Long-Lived Digital Data Collections: Enabling Research and Education in the 21st Century*, September 2005, assigns responsibility to data managers to "provide for the security" and to "provide for the integrity, reliability, and preservation" of the data "by developing and implementing plans for backup, migration, maintenance, and all aspects of change control."

Examples of research data can be characterized by type, such as:

- Observational
  - direct from sensors
  - derivative from direct
- Experimental - laboratory data
- Computational - simulation data
- Visual/Audio
- Source Repositories
- Textual
- Metadata

Attributes of research data need to be determined and itemized in order to form a basis for formulating security requirements. These attributes include, but are not limited to:

- Value
  - long-term storage needed?
  - duplicate, offsite copy needed?
- Reproducibility
  - is the data rare (e.g., historical in some aspect)?
  - is it cheaper to reproduce the data?
- Integrity Assurance
  - is it sensitive or dependent upon unintentional activity, such as transfer, media type, archival technology upgrades, etc.?
- Accessibility
- Confidentiality
  - categorized as private (with release date) or public data
  - need for encryption
- Precision
  - acceptable error rate (stated rate: 100 bit errors per 1 PB of stored data)

  - checksum needed for verification (accuracy of checksums, security of checksums themselves)
- Special Owner Requirements (self-imposed, or externally imposed)
  - dictated by funding agency
  - conform to government regulations (federal or state)
  - needs or goals of the PI

3.4.3 Data Security and Integrity Goals

While recognizing that it is important to maintain secure data, and to provide a mechanism to detect degradation of its integrity either through capricious fate or willful misconduct, one needs to understand that varying degrees of security can be applied to widely diverse types of data, perhaps via the a categorization of the above data types and/or attributes. One example of a categorization of data which could be used as a model for security controls can be found at the University of Texas: http://wwwtest.utexas.edu/its/policies/opsmanual/dataclassification.html

A reasonable goal is to balance security with usability……the key here is to determine where the balance bar is located. There appear to be very few mechanisms to assist the owners or the data mangers to identify the balance point. Data managers need to make security easier than the consequences.

What is the threat to data security or integrity?

It is also helpful for determining requirements by knowing the specific threat to the data security and/or integrity. Recognized threats, with no attempt to assign probability, are:

- Self-inflicted wounds
  - owners of data accidentally destroy or modify data files
  - owners of data fail to adhere to existing, basic protection best practices (file permissions, password protection, etc.)
- Lack of geographical separation of valuable duplicate data copy
- Background treatment of data by compute, data, web, or storage server
  - data is transferred under the covers more often than owners are aware
- Destruction
  - random (hacker)
  - targeted (disgruntled employee)
- Alteration
  - accidental modification (user, system administrator)
  - clever modification (deception, political action)
  - bit rot loss (inherent error characteristics)
- Theft
  - inappropriate copying of the data

3.4.4 Recommendations

1. **User Awareness and Education**. Consistent with most areas of cybersecurity, the effort to keep owners and users of data aware of the threats, and methods to mitigate the threats, must be maintained and increased. Communication between major data centers to discuss, develop and update best practices should be encouraged. **Recommendation:** Centers should publish minimum user best practices for data security and integrity, and encourage adherence to best practices.

2. **Collaboration.** Opportunities for detailed discussions of how to determine user requirements (user survey?), how the find the balance between security and usability, and how to implement measures to assure data security and integrity, should be sought. **Recommendation:** Organize a workshop(s) specific to data security and integrity, perhaps with NSF sponsorship.

3. **Data Security Service.** One of the primary unanswered questions within this breakout dealt with determining and alerting users, by monitoring in an automated fashion (enterprise level), if data has been destroyed or modified. Services could be offered which enhance the level of security required by the user, at the discretion of the user – such as duplication, replication, access control, change detection, encryption, metadata security tracking, etc. Because of varying user needs, a method to categorize the level of security should be developed. **Recommendation:** Explore the demand/difficulty in creation/enhancement of enterprise data security service.

3.5 Breakout session 5: Authentication

The breakout was lead by Von Welch of the National Center for Supercomputing Applications and Ken Klingenstein of the University of Colorado at Bolder. The breakout group consisted of 12-16 people who met over two days of the workshop. There was good carry over in the group from the first day to the next.

3.5.1 Challenges surrounding Authentication

The group identified a number of issues which make authentication a challenging area:

- Authentication is broadly deployed in a variety of forms and has as large of an install base as any other aspect of security. This means authentication has a fairly high resistance to change.
- Authentication is the part of security that non-IT professionals interact with most often and hence has a strong usability component.
- Authentication, Attribute establishment, Authorization and Identity vetting are strongly entwined are hard to full separate.
    - Attribute establishment: establish what I am, e.g. U.S. citizen

  - Identity vetting: vet my email address or postal address, mapping from identifier in cyberspace to real world.
  - Authorization – what you can do may depend on how strongly I believe you are who you say you are (e.g. OTP required for root access)
- Authentication is used for both authorization and auditing/incident response.
  - Requirements of each are different. Authorization is regarding mapping me to a set of permissions in a database. Auditing and incident response usually means mapping me to a physical presence.

3.5.2 Main Discussion Topics

Status of One Time Passwords rollout

At last year's summit, there was a lot of discussion about moving to one time passwords (OTP) to replace our current static passwords. This was largely motivated by the Trojan-based password sniffing methods used in Case 216.

For the most part we didn't. Why not? The consensus of the group was that some form of hardware token is needed to prevent authentication spoofing attacks such as used in Case 216. The reasons why we have not moved to OTP en mass were discussed and are presented here in no particular order.

- OTP deployment brings high costs in terms of dollars both for initial roll-out and ongoing support.
- OTP also has high costs in terms of usability, particularly if all sites rolled it out without coordination and each user had to have a token per site.
- There was fear that attackers were already using methods that would not be effected by OTP deployment, namely session hijacking. Sites are concerned that they would roll out hundreds of thousands of dollars of infrastructure only to have hackers change their tactics without reducing the number of compromises.
- OTP cannot be easily deployed ubiquitously across all applications (login, email, web, etc.). So users would still need to have a normal password in addition the OTP token, increasing usability issues.
- HSPD-12 emerged, causing uncertainty at DOE sites about their future authentication requirements.
- There was a feeling that some form of federated authentication is a solution for some of the usability concerns since it would allow users to have a single token for multiple sites. Sites didn't want to deploy a new authentication solution (OTP) only to have to replace it with a federation mechanism shortly there after, so sites wanted to wait and see if a federated OTP solution emerged.

The general thinking of sites seems to be one of selective OTP deployment where it makes sense from a risk management point of view. Typically this meant deploying OTP for system administrators and for use in critical systems. Sites were also more focused on incident response, namely having plans in place for responding to incidents to mitigate their effectiveness.

It was also noted that the FDIC will require banks to have two-factor authentication by end of 2006. We will be able to see how they handle this and if there is anything from their solution(s) we can leverage.
Following the workshop, OTP deployment experiences were solicited from a number of sites. The following non-exhaustive list of experiences were gathered (listed in alphabetical order):

- ESNet: Deployed a prototype cross-site RADIUS experiment between ESNet, NERSC and ORNL (http://www.es.net/raf/). This effort has stalled due to HSPD-12.NCSA: Deployed OTP for privileged access (e.g. root accounts and non-public machines). Further deployment deferred due to concerns about cost, but a plan to rapidly deploy OTP tokens to key users is in place in the event of a repeat of Case 216.
- NCAR: Fully deployed OTP to all their users are a cost of approximately $350k. Estimates of ongoing user support costs are approximately .5 FTE/year.
- ORNL: Deployed OTP for shell-based logins. Further deployment deferred due to HSPD-12.
- SLAC: Deferred OTP deployment due to HSPD-12.
- SDSC: 80% deployment of RSA OTP capability for SDSC critical infrastructure systems, including supercomputer management nodes, data servers, dns servers, console concentrators. SDSC has spent approximately $20K for systems, software and tokens. Current user base is 82 staff members.

Authentication is an arms race

The participants agreed that improving authentication technologies is basically an arms race with intruders. The comparison was made to frauds and banks – as long as the cost of fraud is below some monetary level, banks are happy and will not improve technology. Could we have similar goals and metrics? One suggestion was that our costs as computing sites involved having an incident response team that managed compromises; as long as that team handled incidents in such a manner that they did not come to management's or NSF's attention, we are content.
The question was raised when we decide to upgrade our authentication technology. It was pointed out that upgrading technology (or practices) inevitably leads attackers to improve their techniques and tools, so in a sense we are making the attackers stronger.
The example of SSH was brought up as the last time the community upgraded their authentication technology en mass. The upgrade to SSH from telnet (and rsh, rlogin, etc.) was precipitated by a clear threat, network-based password sniffing, and a solution that was easily deployed at the system level and did not incur much of a monetary or usability cost. In our current scenario, we also have a clear threat, but OTP incurs much higher costs.

Shell access vs provisioned services

One difficulty faced by our community was that fact that our standard paradigm is to grant our users access to shells on our computing systems. In contrast, campuses often grant access via provisioned services (often web-based). This allows the campuses to tailor the authentication of each service to the level of that service's sensitivity. For shell

access it's impossible to tell when authentication is done what the user's intention for that session will be.

Desire for a Federated Authentication System

The participants agreed that having a federated authentication scheme seems to be the way to solve the usability issues with strong authentication alternatives. It was discussed that we would need to roll out new infrastructure to deploy stronger authentication, so it would be best to combine a federated system at the same time. Several possible mechanisms for federated authentication were discussed: RADIUS, Shibboleth and Online CAs (e.g. the Fermilab KCA deployment). (Editor's note: Kerberos was also mentioned to the facilitators after the session and seems reasonable.)

There seems to be a chicken and egg problem in this area, as until there is a critical mass of sites running a compatible federated authentication mechanism the motivation to do so is low.

Privacy was also mentioned as a possible concern in such an environment. Shibboleth has focused a lot of energy on this issue. For any sort of internationally scoped solution, there are laws for other countries to consider. Are users concerned about privacy? The participants thought they probably weren't, but it wasn't clear if they should be.

Authentication vs Attributes

The question arose of authentication versus the providing of attributes. Authentication would provide an identifier about a user, but a site might need a number of other attributes about a user (e.g. citizenship).
A subtle point was raised that authentication may not mean the same thing to all sites. For example, some sites expect an identity to persist over time while other sites recycle an identity once a user has left. The EduPerson scheme defined a persistent identifier as an attribute just to solve this issue.

3.5.3 Areas for further work

The breakout participants came up with a number of areas that could use further work. They are listed here in no particular order.

- Hardware token usability: Current tokens are not very usable, they could support multiple keys to support multiple sites and could also have a connection to the user's workstation (while still providing a trusted display and keypad) to make them more usable. This will probably come up with the FDIC mandate for banks.
- What can we do to be less concerned about hackers? Several options were mentioned.
    - Having systems based on virtual machines to create "disposable" servers that could be easily discard to be compromised.
    - Using provisioned services, or chroot() environments to allow for the matching of service capabilities to the strength of authentication.

- The use of "creative second factors" in place of hardware tokens was mentioned. For example banks are using cookies in web browsers or the IP address of their customers as the second factor. These could reduce deployment costs.
- Preventing session hijacking on compromised end station to address concerns that hackers will trivially move to hijacking if we strengthen authentication.
- How can we strengthen our validation of remote systems to provide greater assurance of authentication done from that system. For example, how can one tell if the client system is up-to-date on patches and anti-virus signatures? Or determine if the client system is already compromised?
- How we determine the level of sensitivity of a particular request so that we can determine the strength of authentication that should accompany it. Depending on users to consistently and accurately provide this information seems fragile. Can this be automated somehow?
- The issue of authenticating network traffic was raised as being another aspect of authentication – what packets belong to what users? This was mentioned in the context of CALEA.
- The higher level issue was raised about what level of coordination is required for our upgrading of authentication? SSH allowed for very independent authentication, however OTP seems to require more coordination. Each site acting independently would seem to create a large usability issue for users. The possibility of having some coordinated diversity for competition and attack resistance was also mentioned.
- Namespace management, name federation, and identity management were discussed as areas needing work to manage the binding of different user names together in a federated environment.

4. Conclusions and Final Recommendations

The NSF sponsored 2005 Cybersecurity Summit was a valuable follow on to the 2004 Cybersecurity Summit. It continued the important work of raising security awareness for participating NSF sponsored sites and provided a valuable resource for security discussion and recommendations to the NSF and summit attendees. Results of the 2004 Cybersecurity Summit can already be seen in the attention to security issues in NSF program reviews that have occurred in January 2006. The committee believes the Cybersecurity Summit has and can continue to be a valuable resource to the NSF to provide a forum for security discussions and recommendations of particular interest to the NSF and the NSF community related to security issues surrounding NSF proposals, programs and projects.

In addition to the recommendations spelled out in the breakout session sections above, the 2005 Cybersecurity Summit program committee recommends that NSF consider sponsoring a 2006 Cybersecurity Summit. A 2006 Cybersecurity Summit could be used

to continue to raise security awareness for NSF sponsored program staff, point out existing work to the community that has solved a previously identified security problem set and/or NSF could identify a set of security issues of interest to be discussed and addressed at the next summit.

5. Participant Evaluation Summary

Educause will be providing the evaluation summary. Contact Rodney Peterson of Educause for the summary.

Appendix A
Program

Appendix B
Security Tools

During the site survey, a number of tools and procedures were mentioned as being used at the various sites represented. Here is a compilation of those mentioned.

Network IDS/IPS
· Bro, Snort, Tipping Point
Flow based tools
· Netflow, Argus, inMon, RNA, nfdump/nfsen
· Sampling issues with flow analysis.
· Lots of custom flow analysis tools.
· Flow data handling tools (data management).
Host Based Systems
· Tripwire, Osiris, tcpwrappers, prelude, BlackICE, centralized log servers, custom tools for checking permissions and configurations.
Other Techniques
· Border firewalls, ACLs, Darknet
Scanning
· Nessus, Nmap, one site had a self scan portal.

Some of the breakout sessions produced discussions related to various security information tools and resources. These sites are listed here for potential interest to conference attendees.

Bro Intrusion Detection System
· http://broids.org/

Darknet/Internet Motion Sensor webpages
· http://www.cymmu.com/Darknet/index.html
· http://ims.eecs.umich.edu

(FIRST) Forum for Incident Response and Security Teams
· http://www.first.org

Nessus Vulnerability Scanner
· http://www.nessus.org

Nmap Network vulnerability scanner
· http://www.insecure.org/nmap/

NSPSEC Email List
· http://puck.nether.net/mailman/listinfo/nspsecurity

Rancid Infrastructure configuration management and monitor
· http://www.shrubbery.net/rancid/

SGUIL GUI for incident handling and response
· http://sguil.sourceforge.net/

Security Conferences and Additional Resources
· Defcon http://www.defcon.org/
· EDUCAUSE Security Task Force http://

www.educause.edu/security
· Joint Techs Workshops http://jointtechs.es.net/
· SANS http://www.sans.org/
· Shmoo Group http://www.shmoo.com/
· USENIX Security Conference http://www.usenix.org/