

UC Davis

UC Davis Previously Published Works

Title

Case Studies of an Insider Framework

Permalink

<https://escholarship.org/uc/item/7dw4m2w6>

ISBN

9780769534503

Authors

Bishop, Matt
Engle, Sophie
Peisert, Sean
et al.

Publication Date

2009

DOI

10.1109/hicss.2009.104

Peer reviewed

Case Studies of an Insider Framework

Matt Bishop, Sophie Engle,
Sean Peisert, Sean Whalen
University of California, Davis
Davis, CA

{bishop,engle,peisert,whalen}@cs.ucdavis.edu

Carrie Gates
CA Labs
Islandia, NY
carrie.gates@ca.com

Abstract

Many groups are interested in the insider threat problem, but the model generally used by all of these groups is implicitly binary—one is either within a perimeter or not. There is another model, however, that employs a graduated approach to defining insiders. This approach gives greater flexibility for considering many threats that are not traditionally captured by a model, such as the impact of social engineering attacks. This new definition enables more accurate and useful security policies to be implemented so that well-defined insiders can be deterred, detected, and analyzed. We examine the flexibility of this model in this paper through case studies, showing how the model captures both traditional insiders and social engineering attacks.

1 Introduction

Studies of the insider problem predate computers by hundreds of years. Most of those studies have considered insiders in a binary way—someone is either inside some defined perimeter, or they are not. However, the concept of a well-defined perimeter is disappearing—particularly with regard to computer systems [13]—as organizations hire contractors, outsource to other organizations, partner and merge, and use software-as-a-service (SaaS) offerings such as Salesforce.com as part of their core business. As a result, a strict definition of insiders based on a defined perimeter is becoming less meaningful, in the

same way that traditional computer security defenses such as firewalls and network intrusion detection systems no longer protect against attacks (including mistakes) occurring on the inside, including Trojan horses, phishing, and client-side, cross-site scripting attacks.

The use of a binary definition of an insider does not allow for the prioritization or categorization of insiders; all insiders are considered equal. Such an approach does not allow for an efficient division of resources, and the focus is diverted away from protecting an organization’s most valuable assets or intellectual property. As a result, traditional models of insiders are often not useful components of a meaningful risk analysis that allows management and security administrators to prioritize resources and effort.

We have previously presented a definition of an insider, using a concept that we call Attribute-Based Group Access Control [4], and have extended that model into a usable framework for categorizing insiders, using a hierarchy of policies [3]. In this paper, we demonstrate the utility of that approach through two case studies. The first case study illustrates how the framework applies to the traditional notion of an insider. To demonstrate that this framework also applies to other circumstances not commonly considered during the creation of insider threat definitions and models, we present a case study where we apply our model to an embezzlement scandal and a “spear-phishing” attack.

The rest of this paper is organized as follows. In Section 2 we describe related work. We then describe

the initial definitions of insiders in Section 3, followed by a framework based on this approach in Section 4. We provide a case study analysis of how this model performs given traditional insiders in Section 5, while in Section 6 we describe how the model captures social engineering attacks. We then provide some concluding remarks in Section 7.

2 Related Work

Many researchers have investigated the problem of insider threat. However, most publications do not precisely define an insider, instead proceeding to talk about defenses while assuming that the user inherently understands the term. But without a consistent definition of an insider, researchers have developed their own definitions particular to their own data sets, situations, biases and assumptions.

As a result, research into the detection or mitigation of insider threats can not necessarily be applied from one domain to another as the underlying model does not necessarily translate between domains. The situation is further complicated by definitions which are not only ambiguous but contradictory. For example, a RAND report defines an insider as “an already trusted person with access to sensitive information and information systems” ([5], p. xi). Elsewhere it defines an insider as “someone with access, privilege, or knowledge of information systems and services” ([5], p. 10), omitting the need for that person to be trusted. A different report implicitly defines the insider as anyone operating inside the security perimeter ([14], p. 3), again ignoring trust and also knowledge of the systems.

The problem of defining an insider complicated by the assumption that a perimeter can even be defined to begin with, such that someone inside the perimeter is therefore an insider. However, the concept of distinct borders around an organization are blurring with the increased usage of mobile computing, outsourcing and contracting. Even in those cases where a distinct border can be defined, many definitions focus on technology borders and fail to consider physical borders and the ability to circumvent borders (for example, by social engineering).

Consider the definition given in the 2005 U.S. Secret Service and CERT report which defines an insider as “individuals who were, or previously had been, authorized to use the information systems they eventually employed to perpetrate harm” ([11], p. 3). This presents a narrow view of an insider and implicitly hints at the existence of a security perimeter separating those who are authorized (insiders) from those who are unauthorized.

Using the USSS/CERT definition, we are unable to capture the insider threat posed by a disgruntled janitor with physical access to information systems but without the authorization to use them. We are also unable to capture individuals with no access defined in the organization at all. For example, consider a discontent teen with access to a parent’s laptop. Social engineering further complicates the picture. What if an attacker is able to trick an insider into revealing sensitive information? These examples blur the security perimeter separating an insider from the outside.

These issues surface in other definitions as well. In their 2005 model of insider threat assessment, Chinchani et al. [7] define insiders as “legitimate users who abuse their privileges.” However, despite the narrow definition of an insider, they acknowledge that the perimeter between insider and outsider is “fuzzy” and that insider threat involves both “computational elements and human factors.”

The notion that defining a perimeter separating insider from outsider is difficult at best is not new. For example, Schultz in 2002 touches on this in his framework for understanding insider attacks [18]. Despite this knowledge, the community still implicitly relies largely on perimeter-based definitions of who is an insider.

When attempting to model the insider threat, we must examine both who is the “insider” and what is the “threat” presented by the insider. Before we present our definition of an insider, we explore why insiders pose a threat to institutions and organizations. This allows us to explore different categories of insider threat, and how this affects the notion of an “insider” in different scenarios.

Understanding the inherent limitations of a security policy is critical to understanding the extent of

the threats posed by insiders. The first limitation is that some policies cannot be implemented on a computer system at all, leaving the organization vulnerable due to lapses in procedures or worse, due to misunderstandings about what the policy actually says. A second limitation is with policies: while following the principle of least privilege is the ideal situation [17], individuals typically have more privileges than necessary to complete an authorized task. For example, as we indicated earlier, employees need to be able to do their jobs, and it is rare that their privileges are exactly equal to their needs. Thus privileges are nearly impossible to get “just right.” Privileges are almost always either too restrictive (which frustrates the employee and impedes work), or not restrictive enough (which introduces unnecessary privileges that can be abused). We now explore the limitations of policy and its relation to insider threat.

3 Definition

Bishop [2] proposed a definition of an insider as: “a trusted entity that is given the power to violate one or more rules in a given security policy ... the insider threat occurs when a trusted entity abuses that power.” This definition hints at the need to recognize that an insider must be determined with reference to some set of rules that is part of a security policy. We expanded that definition [4, 3] by arguing that a security policy is inherently represented by the access control rules employed by an organization. So, an insider is defined with regard to two primitive actions:

1. violation of a security policy using legitimate access, and
2. violation of an access control policy by obtaining unauthorized access.

In the first case, the insider uses their legitimate access to perform some action that is contrary to the security policy, such as might be observed when sensitive data is leaked to some third party or when access to a resource is given or blocked. Here the insider has legitimate access to the data or resources, but uses that access to provide the information to someone who does not themselves have access (or to deny

access to someone who does have access). In the second case, the insider uses their access to extend their privileges in a manner that breaks both the access control and security policies. An example of such a breach occurs when a user might have a legitimate capability to log into a particular system, but then abuses that privilege to gain illegitimate superuser-level access to the system (e.g., by exploiting some system vulnerability such as a buffer overflow or race condition).

As we discussed earlier, previous definitions gave rules or descriptions intended to allow the reader to determine who is an insider, resulting in a binary distinction: an entity is either *an insider* or *not an insider*. We argue that a non-binary approach is required, to indicate degrees of “insiderness,” and that the access control rules for an organization can be used to develop these degrees. We define someone as an insider *with respect to access to some data or resource X*.

4 Framework

In order to capture this notion of insider as a function of access to data or resources, and develop it into a usable framework. To do this, we initially developed a model called *Attribute-Based Group Access Control (ABGAC)* [4], which, when merged with Carlson’s Unifying Policy Hierarchy [6], becomes a usable and implementable framework.

ABGAC is a generalization of *role-based access control (RBAC)*. Whereas RBAC largely assigns rights based on specific job functions that a person has within an organization, ABGAC assign rights based on general attributes, which might include elements of person’s job function, but may also include a variety of other sources regarding the person or their environment.

We define attributes as descriptions of the protection domain of entities. The protection domain can include access rights to objects and resources, including systems, printers, documents, buildings, and generally any other object to which a user can have access. The protection domain can also include procedural access rights such as physical presence, or the

ability to block access.

Once defined, the protection domains need to be partially ordered by value. The organization must do a cost/benefit analysis to assign a value to the protection domains. For example, an organization might specify that access to financial documents, the email of senior level executives, and source code for specific products represents the information potentially of greatest value and, therefore, represents the greatest damage if leaked or compromised. The value of a protection domain of a user should not be defined solely by a systems administrator, but rather as a joint effort between the senior executives and the security administrators. Once ordered, the protection domains can be combined into groups, where the group indicates the threat level a particular set of attributes represents. These are called *pd-groups* to distinguish them from groups of users.

Paired with each protection domain is the group composed of the users to which that protection domain applies. In other words, groups are created based on the protection domains of the associated users, rather than on the job functions of the associated users (as in a role-based system). The users with access to the *pd-groups* with the highest value then represent those users who pose the greatest risk for insider threat. Given this pairing, we can create a lattice based on the ordering of protection domains and the ordering of groups. Given two pairs, we can determine which indicates the greatest risk by their ordering.

The creation of such a lattice requires a two-stage approach: determining the important components of the protection domain relevant to some privilege and identifying all users. It is not necessary to provide all components and privileges, but rather only those that are relevant to the well-being of the organization and therefore at risk due to insider threat. Initial users include not only direct employees, but also all contractors and out-sources (e.g., technical, clerical, janitorial), and any “special case” accesses (such as facility visitors or guest logins). Once the protection domains and users have been identified, the two are mapped together based on the access the users have. This creates an ordered group of users who represent insiders, where the ordering is based on the value

of the resources to which they have access. Thus a security administrator can focus their attention on those insiders who pose the greatest threat, and an insider is thus defined with respect to the resources to which he has access.

Degree of insider abuse and job function may be very different, but we assert that the actual level of threat is much more important than the level of threat implied by a job function (which may or may not be true). So, by employing attributes and protection domains for the lattice rather than roles we are able to specify and group disparate users who might have equal access in terms of insider abuse rather than simply by job function. For example, assume that the CEO of a company has identified the customer contact and purchase information as high-priority information that requires protection. Users who might have access to this information, and hence be placed in a group together, would include not only the sales representatives for the company, but also potentially external entities, such as the system administrators for Salesforce.com (assuming that the organization uses Salesforce.com). This is an example of the disappearance of a well-defined perimeter for an organization, and how ABGAC is able to still capture potential insider threats.

We have since extended the ABGAC model to include the notion of using policy discrepancies to identify initial locations for insiders [3]. The extended model builds on Carlson’s Unifying Policy Hierarchy [6], which defines four levels of policy (from highest to lowest): oracle policies, Feasible Policies, Configured Policies, and Real-Time Policies, as shown in Fig. 1. The Oracle Policy is a policy that assumes perfect knowledge (including, for example, intent of a transaction). The Feasible Policy implements the Oracle Policy as best it can given real-world constraints (for example, a Feasible Policy will only be able to allow or disallow a transaction, but cannot determine the actual intent behind a transaction). The Configured Policy is the actual policy that has been implemented, since some Feasible Policies might be configured in multiple ways; the Configured Policy represents the choices made in implementing the Feasible Policy. Finally, the Real-Time Policy represents the policy decisions being made as they are

UNIFYING POLICY HIERARCHY		
Level	Domain	Description
ORACLE POLICY	all possible (s, o, a, e) tuples	Captures notion of an “ideal policy” even if such a policy isn’t explicitly defined.
FEASIBLE POLICY	system-definable (s, o, a) tuples	Represents what can in practice can be captured on an actual system.
CONFIGURED POLICY	system-defined (s, o, a) tuples	Represents the policy as configured on an actual system.
REAL-TIME POLICY	system-defined (s, o, a) tuples	Represents what is possible on an actual system.

s : subject o : object a : action e : environment/intent

Figure 1: Four levels of Carlson’s Unifying Policy Hierarchy [6].

made, and take into consideration other issues such as system constraints.

We argue that the ability to perform an insider attack comes from a discrepancy in the expressiveness (and therefore enforceability) between two policy layers. As given in the example above, an Oracle Policy might specify the intent behind a transaction as determining if a transaction should be allowed, however the Feasible Policy has no ability to determine intent. Thus an insider might have the capability to perform a particular transaction, and might decide to abuse that capability as part of some insider attack. The discrepancy between the Feasible Policy (where the transaction is allowed) and the Oracle Policy (where the transaction is allowed only for specific reasons) allows this attack to occur. These two aspects of the model combine to provide a framework for discussing the insider threat problem, and for defining who is an insider.

From a practical, implementation perspective, we can merge the ABGAC model into a model of attacks for improved forensic analysis, simply by monitoring the use of credentials on sensitive documents. While security policies must identify the sensitive documents and high levels of access to begin with, the forensic model helps to determine what is needed to understand the path to the objects, and the actions taken on them. The two models can also jointly iden-

tify where we cannot easily enforce policies by logging information, and thus provides a measure of how well an attempted threat to security can or cannot be determined in a post mortem analysis.

5 Case Study 1: Union Dime Embezzlement

In the years from 1970 to 1973, the Union Dime Savings Bank lost U.S. \$1.5 million to embezzlement at the hands of Jérôme Kerviel, their chief teller. The scheme would likely have lasted longer had there not been an unrelated arrest of Kerviel’s gambling bookie, whose records resulted in his investigation and eventual conviction.

As chief teller, Kerviel was able to issue an “error correction” to accounts that reduced the digitally recorded account balance. He then pocketed the remainder in hard cash. When the time came for interest calculations, he would move money from other accounts into the account he ‘corrected’ so the balance would appear as expected and interest was properly calculated.

The embezzlement was enabled by several of the bank’s practices. There were two types of accounts whose interest was calculated on different days, allowing money to be shifted from one account type to

the other on the day interest was calculated. This allowed account records to appear balanced despite the teller's pocketing of money after issuing corrections. In addition, customers received no monthly statements. An account's balance was recorded on a customer's booklet stamped at the time of deposit. Any adjustments to the bank's records would not be reflected until the customer's next withdraw, making low activity accounts an attractive target.

First, we fit this into our policy hierarchy. We emphasize that the following is one reasonable interpretation of the policy hierarchy. Others give similar results. We assume that the Oracle Policy states that "the chief teller can issue error corrections to accounts to correct errors in data entry." The feasible Oracle Policy cannot distinguish between an "error in data entry" and "an error arising from illicit withdrawal." Thus, the feasible Oracle Policy eliminates the motivation behind the error correction, and simply says "the chief teller can issue error corrections to accounts to correct errors." Here, the chief teller is complying with the feasible Oracle Policy (because he is authorized to issue error corrections to accounts to correct errors) but not with the Oracle Policy (because the error being corrected is not related to an error in data entry; it arises from an illicit withdrawal).

Now consider a variant of the feasible Oracle Policy that says "the chief teller can issue error corrections to accounts to correct errors, and shall record the reason for each error correction in a log." Now, when an auditor checks the accounts, the auditor can determine whether the chief teller issued the correction to fix a data entry error. But consider the next layer of policy, where the system is configured to record the log. If the log can only be made writable and not append-only, the chief teller can erase entries to hide that a change was made (and thus suppress the need to enter a reason). So, if the configuration policy says "the chief teller can write (edit) the log associated with error corrections to accounts," then there is a discrepancy between the Configured Policy (which says that the chief teller can change anything in the log) and the feasible Oracle Policy (which says the chief teller's reason for changing the account must be recorded in the log).

Note also the discrepancy between the Configured

Policy, the feasible Oracle Policy, and the Oracle Policy. The Oracle Policy asserts that the chief teller's reason for changing the account is known, at least to the oracle, which can then pass on whether the reason and the change comply with its policy. But the feasible Oracle Policy says nothing about motive, merely that the teller record the reason for change. Similarly, the Configured Policy simply says the chief teller can write to the log, and nothing about *what* he must write. Hence there is a discrepancy on multiple levels: the chief teller can lie. Underlying this assertion is the Oracle Policy's ability to discern the actual reason for an act, and the inability of policies at other layers in the policy hierarchy to know the actual reason.

Given all this, we can integrate the ABGAC model to determine where insiders might arise. Let us assume the oracle, feasible oracle, and Configured Policy as above. The resources involved in this episode of the Union Dime Bank are the cash in the bank, and the ability to take it physically from the bank; and the error correcting function and the ability to execute it. The tellers have access to the cash in the bank, as do those with access to the bank vault. Assuming the tellers are not searched when they leave, they also have the ability to take the cash physically from the bank. The question of who can execute the error correcting function limits the set of tellers to the chief teller, assuming correct implementation (a point we shall touch on in a moment). Thus, the set of people who can perform the above insider attack, namely embezzle funds in the manner described, is one: the chief teller.

The above analysis makes two assumptions. The first is that only one person is involved. The execution of the error correcting function requires the chief teller to act, so he must be involved in this compromise. But he need not be the one who takes the cash out of the bank. He could be in cahoots with one or more other tellers, who will remove the money for him. Such a compromise is feasible, but less likely to succeed due to Benjamin Franklin's claim, "three may keep a secret, if two of them are dead" [9].

The second assumption is that the implementation of the Configured Policy is correct. For example, suppose there is a bug in the software managing the er-

ror correction routine. Then the Real-Time Policy is that anyone with access to the system on which that routine resides can change the amounts in accounts, thereby performing the same function as the chief teller. Thus, the ABGAC analysis captures those attackers who exploit implementation bugs (or, more properly, discrepancies between the Real-Time Policy and the configuration policy) in the same way it captures those who can exploit discrepancies between the higher layers of policy abstraction.

The application of ABGAC to this scenario provides a basis for identifying the threat. In doing so, it also provides a basis for mitigating the threat ahead of time, as well as instrumenting a system in a way that [15, 16] enables targeted logging of potential violations of the security policy, and therefore, more efficient analysis the whether a violation was attempted and successful. Specifically, the ABGAC model gives a benchmark of where logging is (a) feasible, and (b) useful. Where logging is not feasible, we can place bounds on the possible gaps in our levels of knowledge and attempt other forms of monitoring (e.g., physical security). Where logging is not useful, we can avoid taxing computer and network resources collecting useless data. In the place of the bank teller example, we can certainly isolate the bank teller and the systems and accounts that the bank teller has access to, and by generating attack graphs, starting with the teller’s likely, ultimate goals (and/or the largest threats)—embezzlement—we can develop metrics that might help the other possible paths to accomplish those goals (e.g., collaborating with other bank employees), and monitor and protect accordingly.

Embezzlement is a particularly good demonstration of the insider model, and is broadly applicable in other such situations. For example, one might imagine that the French bank Société Générale wished they had been able to perform a better risk analysis by classifying insiders, threats, and targets using such a system before losing U.S. \$7.1 billion [8].

6 Case Study 2: Social Engineering

Phishing as a security issue has traditionally been viewed as a social engineering attack and identity theft threat. However, it can also be viewed as a special case of the insider problem. In a phishing attack, the adversary sends an email to a target group soliciting them to perform some action that will reveal the target’s credentials (to some target location) or sensitive information [1]. For example, the adversary might have set up a fake web site emulating a popular bank. He then sends email to some large number of email addresses, where the email appears to be an official communication from the bank. The email might encourage the user to “follow the link below” to log in and perform some action. As the link is actually to the fake web site, the adversary is then able to capture the credentials of the target user. Note that in this case the attack is indiscriminant and succeeds due to the large number of users targeted.

Related to phishing attacks are spear phishing and whaling [12]. Spear phishing refers to attacks that are targeted at particular individuals or companies, rather than indiscriminant as in generic phishing attacks. Whaling is a special case of spear phishing that is aimed at company executives.

In this section we examine two examples of phishing. First we consider the more generic (and prevalent) forms of phishing, and then we provide an example specific to spear phishing.

Given these descriptions, the insider attack occurs with respect to the target *organization* and not the individual. In the case of phishing, for example, the adversary might be trying to gain an individual’s credentials in order to log into that individual’s bank account and transfer funds to the adversary’s PayPal account or make online purchases. Here the bank would be the target organization. The Oracle Policy in this case might be “Only the owner of an account can access and perform transactions using that account.” In contrast, the Feasible Policy would state, “Only someone presenting the credentials of the owner of an account can access and perform trans-

actions using that account.” Note that the Feasible Policy can not distinguish the owner of an account except through the use of his credentials. Thus any person providing those credentials is assumed to be the owner of the account. In this case the Configured Policy is the same as the Feasible Policy.

Translating this into the ABGAC model we have as resources the account at the bank and the money in that account, while the users are the person owning that account, the employees of the bank, and the adversary. In this case the account owner has access to the account and the money in that account, as do the bank employees, while the adversary does not. The insiders for this account are therefore the bank employees and the account owner. When the adversary obtains the account owner’s credentials, he is the account owner from the perspective of the bank. Thus the adversary is also now an insider.

The examples for spear phishing (or whaling) are slightly different from the more general forms of phishing. In these cases often the email exchange aims to gather the trust of the target, enticing him to install some piece of software [10]. Often it achieves this trust by providing a sufficient amount of identifying information that the target believes the adversary is who he claims to be. The software installed is generally some form of malware, such as a keylogger, that then sends information back to the adversary. This is interesting as an insider threat problem because, technically, the adversary *never* operates inside the perimeter of the target organization, however he does receive that organization’s information (e.g., logged keystrokes, particular files). However, using the policy discrepancy and ABGAC models, this case can be represented as an insider problem.

Assuming the case of a keylogger that sends information back to the adversary, the Oracle Policy might state “this computer can send information to other machines on the network upon explicit approval of the user.” However, the Feasible Policy might be less restrictive: “This computer can send information to other machines on the network if the user is logged in.” The assumption here might be that if the user is logged in (which is easy to determine—at least that *some* user is logged in, using this user’s credentials) then he approves of the communication being

sent since he is (presumably) the one sending it. In contrast, requiring explicit approval for all communication requests would likely be onerous for the user (e.g., needing to approve all web browsing activity, every email sent, etc.). Thus the computer program, which logs the keystrokes of the user and then sends the information back to the adversary, acts as an insider activity. The adversary, although always external to the host machine and network, is the insider as he gains insider knowledge based on the keystroke logging he receives. This is possible due to the gap due to assuming that a user who is logged in is therefore explicitly approving all communications between his computer and any outside systems.

This can be translated into the ABGAC model by stating the resources as being access to the particular system, the information on that system, and the new information entered and activities performed on that system. The users for this system are the user with legitimate access to this system (e.g., the employee, CEO, etc.) and the adversary (along with others, such as system administrators). In this case, due to the presence of the keylogger, while the two users are separate with regards to access to the system and the information on that system, they can be grouped together regarding the new information gathered and activities performed on that system. Thus the adversary in this case is an insider.

7 Conclusion

In our previous work, we proposed an approach to defining insiders that takes the the current state of organizations, computer networks, and mobile computing into account by re-evaluating perimeters and binary distinctions. The approach provides a basis for gaining traction toward measuring, monitoring, and mitigating the insider threat. While the majority of research implicitly defines an insider as a binary condition (one is either an insider or not), we take the approach of defining an insider based on their access attributes. More specifically, we have defined a lattice consisting of protection domains on one axis and users (not roles) on the other axis. By ordering protection domains based on their value, we can then

group them by their value. By then grouping users according to the value of their protection domains, we can provide a continuum of insiders. This allows researchers and security personnel to focus on the insiders who can cause the greatest amount of damage to an organization, and to develop policies and solutions for reducing the threat of those insiders.

In this paper we summarize our previous work and apply the approach to two case studies that demonstrate the utility of our approach. While our previous work has been largely theoretical, we focus here on the utility of our approach and its capability to represent multiple forms of insiders. We demonstrate how the model can be applied in both traditional insider cases (e.g., embezzlement) as well as in social engineering threats (e.g., phishing). Old threats are increasing in severity due to the speed at which they can occur, and the ability of computer networks to provide rapid, anonymous communication, whereas traditional threats required in-person, human contact, which was slower, more complicated, and more dangerous. In both of these case studies, we have shown the model allows for finer-grained and more appropriate classification of the threat scenarios. We also briefly demonstrated how the classification and identification of the threats can be merged with a model of attacks to guide a post mortem analysis.

Acknowledgements

Matt Bishop was supported by grant CND-0716827 from the National Science Foundation to the University of California at Davis. Sophie Engle was supported by grant H98230-07-1-0234 from the Department of Defense to the University of California at Davis. The views and conclusions expressed in this paper are those of the author, and not necessarily those of any funding agency.

Sean Peisert was supported by grant 2006-CS-001-000001 from the U. S. Department of Homeland Security, under the auspices of the Institute for Information Infrastructure Protection & I3P research program. The I3P is managed by Dartmouth College. The views and conclusions contained in this document are those of the authors and should not be inter-

preted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security, the I3P, or Dartmouth College.

References

- [1] Binational Working Group on Cross-Border Mass Marketing Fraud. Report on phishing. Technical report, US Justice Department and the Ministry on Public Safety and Emergency Preparedness Canada, October 2006.
- [2] M. Bishop. Position: “Insider” is Relative. In *Proceedings of the 2005 New Security Paradigms Workshop (NSPW)*, pages 77–78, Lake Arrowhead, CA, October 20–23, 2005.
- [3] M. Bishop, S. Engle, S. Peisert, S. Whalen, and C. Gates. We Have Met the Enemy and He Is Us. In *Proceedings of the 2008 New Security Paradigms Workshop (NSPW)*, Lake Tahoe, CA, September 22–25, 2008.
- [4] M. Bishop and C. Gates. Defining the Insider Threat. In *Proceedings of the 2008 Cyber Security and Information Infrastructure Research Workshop*, Oak Ridge, TN, 2008.
- [5] R. Brackney and R. Anderson. Understanding the Insider Threat: Proceedings of a March 2004 Workshop. Technical report, RAND Corporation, Santa Monica, CA, March 2004.
- [6] A. Carlson. The Unifying Policy Hierarchy Model. Master’s thesis, University of California at Davis, Dept. of Computer Science, 1 Shields Ave., Davis, CA, June 2006.
- [7] R. Chinchani, A. Iyer, H. Ngo, and S. Upadhyaya. Towards a Theory of Insider Threat assessment. In *Proceedings of the International Conference on Dependable Systems and Networks (DSN)*, pages 108–117, June/July 2005.
- [8] N. Clark and D. Jolly. Fraud Costs Bank \$7.1 Billion. *New York Times*, January 25, 2008.
- [9] B. Franklin. *Poor Richard’s Almanack*. 1735.

- [10] iDefence. Spear Phishing and Whaling Attacks Reach Record Levels. *iDefense Press Release*, June 2008.
- [11] M. Keeney, E. Kowalski, D. Cappelli, A. Moore, T. Shimeall, and S. Rogers. Insider Threat Study: Computer System Sabotage in Critical Infrastructure Sectors. Technical report, U.S. Secret Service and CERT, May 2005.
- [12] J. Markoff. Larger Prey are Targets of Phishing. *New York Times*, April 16, 2008.
- [13] A. P. Moore, D. M. Cappelli, and R. F. Trzeciak. The “Big Picture” of Insider IT Sabotage Across U.S. Critical Infrastructures. Technical Report CMU/SEI-2008-TR-009, CERT, May 2008.
- [14] J. Patzakis. New Incident Response Best Practices: Patch and Proceed is No Longer Acceptable Incident Response. Technical report, Guidance Software, Pasadena, CA, September 2003.
- [15] S. Peisert, M. Bishop, and K. Marzullo. Toward Models for Forensic Analysis. In *Proceedings of the 2nd International Workshop on Systematic Approaches to Digital Forensic Engineering (SADFE)*, pages 3–15, Seattle, WA, April 2007.
- [16] S. P. Peisert. *A Model of Forensic Analysis Using Goal-Oriented Logging*. PhD thesis, Department of Computer Science and Engineering, University of California, San Diego, March 2007.
- [17] J. H. Saltzer and M. D. Schroeder. The Protection of Information in Computer Systems. *Proceedings of the IEEE*, 63(9):1278 – 1308, September 1975.
- [18] E. E. Schultz. A Framework for Understanding and Predicting Insider Attacks. *Computers and Security*, 21(6):526–531, 2002.