

Preface

On behalf of the Organizing Committee, I am pleased to present to you the proceedings of the first Workshop on Insider Attack and Cyber Security held in Washington DC in June 2007. This book serves to educate all interested parties in academia, government and industry and that helps set an agenda for an ongoing research initiative to solve one of the most vexing problems encountered in securing our critical IT infrastructure, the insider threat. In some sense, the insider problem is the ultimate security problem. Insider threats, awareness and dealing with nefarious human activities in a manner that respects individual liberties, and privacy policies of organizations, while providing the best protection of critical resources and services that may be subjected to insider attack, is a very hard problem requiring a substantial effort by a large research community. We hope this book helps establish a community of researchers focused on addressing the insider problem.

The book contains a number of invited papers authored by attendees of the workshop. We believe the material that has been selected is of wide interest to the security research community. Papers have been invited that help define the nature and scope of the insider attack problem. Several papers provide an overview of technical solutions that have been proposed and discuss how they fail to solve the problem in its entirety. An essential theme of the workshop was to educate researchers as to the true nature of the problem in real-world settings. Papers are provided that describe the nature and scope of the insider problem as viewed by the financial industry. The book concludes with technical and legal challenges facing researchers who study and propose solutions to mitigate insider attacks.

We wish to thank Cliff Wang of the Army Research Office, Daniel Schutzer of the Financial Services Technology Consortium and Eric Goetz of the Institute for Information Infrastructure Protection for supporting our effort and sponsoring the Workshop, and Shari Pfleeger of Rand Corporation for providing the venue for our meeting and assistance in organizing the Workshop. We also thank the reviewers who served anonymously to vet the technical papers included here. Finally, we are especially grateful to Shlomo Hershkop and Sara Sinclair for their remarkable effort to organize and format the individual papers to produce a final cohesive manuscript.

January 2008

Salvatore J. Stolfo

Reflections on the Insider Threat

Charles P. Pfleeger

Pfleeger Consulting Group

Abstract This paper reports on a workshop in June 2007 on the topic of the insider threat. Attendees represented academia and research institutions, consulting firms, industry—especially the financial services sector, and government. Most participants were from the United States. Conventional wisdom asserts that insiders account for roughly a third of the computer security loss. Unfortunately, there is currently no way to validate or refute that assertion, because data on the insider threat problem is meager at best. Part of the reason so little data exists on the insider threat problem is that the concepts of insider and insider threat are not consistently defined. Consequently, it is hard to compare even the few pieces of insider threat data that do exist. Monitoring is a means of addressing the insider threat, although it is more successful to verify a case of suspected insider attack than it is to identify insider attacks. Monitoring has (negative) implications for personal privacy. However, companies generally have wide leeway to monitor the activity of their employees. Psychological profiling of potential insider attackers is appealing but may be hard to accomplish. More productive may be using psychological tools to promote positive behavior on the part of employees.

1 Introduction

In June 2007 the U.S. Army Research Office, the Financial Services Technology Consortium (FSTC) and the Institute for Information Infrastructure Protection (I3P) sponsored a workshop on insider attack and cyber security. The two-day event featured participants from academia, research institutions, consulting firms, industry, and the government. The security researchers, practitioners and vendors in who attended shared insights and frustrations.

Reflecting on the presentations, discussions and comments, I am documenting in this paper some high level observations that came as a result of that meeting.

2 Who Is an Insider?

Who is an insider?

This question seems straightforward and easy to answer. But as with other fundamental terms in computer security (such as integrity, availability, or even security) the definition of insider is not well established. There are several possible uses of the term.

An insider can be:

- an employee, student, or other “member” of a host institution that operates a computer system to which the insider has legitimate access
- an associate, contractor, business partner, supplier, computer maintenance technician, guest, or someone else who has a formal or informal business relationship with the institution
- anyone authorized to perform certain activities, for example a bank’s customer who uses the bank’s system to access his or her account
- anyone properly identified and authenticated to the system including, perhaps, someone masquerading as a legitimate insider, or someone to whom an insider has given access (for example by sharing a password)
- someone duped or coerced by an outsider to perform actions on the outsider’s behalf
- a former insider, now using previously conferred access credentials not revoked when the insider status ended or using access credentials secretly created while an insider to give access later

This rather broad range of interpretations of the term insider is by no means exhaustive. But it does point out the potential for confusion both inside and outside the computer security profession.

2.1 *Motive*

The motives for an insider attack are similarly diverse. In fact, the term “attack” may be overly harsh for certain types of insider actions:

- making an unintentional mistake
- trying to accomplish needed tasks—for example, in a case in which the system does not support a particular action or the insider is blocked from accessing certain data, the insider may try workarounds to accomplish the same thing
- trying to make the system do something for which it was not designed, as a form of innovation to make the system more useful or usable

- trying innocently to do something beyond the authorized limit, without knowing the action is unauthorized
- checking the system for weaknesses, vulnerabilities or errors, with the intention of reporting problems
- testing the limits of authorization; checking the system for weaknesses, vulnerabilities or errors, without the intention of reporting problems
- browsing, killing time by viewing data
- expressing boredom, revenge or disgruntlement
- perceiving a challenge: treating the system as a game to outwit
- acting with the intention of causing harm, for reasons such as fame, greed, capability, divided loyalty or delusion

We obviously react differently to these different motivations, sympathizing with the employee who wants to get work done in spite of the system, but deploring agents with malicious intent. Unintentional errors are usually seen as unfortunate but inevitable, and malicious behavior is usually seen as something heinous that should be prevented. But the area between these two ends is grey.

Unfortunately for research purposes different people include different ones of these cases in the definition of insider behavior. A given action may be classified as an insider attack in one study but not in another, which complicates assessing the severity and frequency of insider “attacks.” Because different projects use different definitions, comparing results or statistics between projects can be difficult for analysts and misleading to the public.

As one participant pointed out during the workshop, two interesting cases arise: when bad things happen even though system privileges are not exceeded, and when good things happen even though system privileges are exceeded. We might initially say we want to prevent the former, but blocking acceptable behavior risks limiting a system’s usability. We also tend to excuse the latter if the good dominates. These two cases show how difficult it is to separate acceptable insider behavior from unacceptable. With a murky policy definition, enforcement becomes problematic.

2.2 *Effect*

Another way to analyze the insider threat is to look at the effect insiders’ actions have had. Here are some impacts of insider attacks:

- making available data or computer services to people who would otherwise not have had them—either because the people were not authorized or because the system failed to deliver as expected or intended
- receiving data for which the user was not authorized because such data fell outside the user’s job requirements

- obtaining data or services for fraudulent purposes

The first impact here would sometimes be considered positive, and the last is usually negative. The middle impact can be mixed, depending on what use the user made of the data. The impact of a insider can thus range from positive to negative.

2.3 Defining the Insider Threat

Two major points stand out: First, we need standard definitions of insiders and insider behavior so studies and discussions can compare like entities. These definitions need to be used not just in the computer security research community but also by commercial security professionals (such as chief security officers and other management) and the press. (Convincing the press to use these terms precisely may be challenging.)

Second, we need to recognize that, unlike the “outsider” threat, insider behavior with the potential to do harm ranges from human nature (unintentional errors) through positive intentions (getting the job done in spite of an uncooperative system) and finally to all kinds of malice. Threat is the common term in computer security for an action with the potential to cause harm. But because the word “threat” has a negative connotation, some people would understandably not ordinarily use it to describe unintentional or non-malicious behavior. We must be especially careful when using the term “insider threat” to be sure our meaning is not misconstrued and insiders are not offended.

2.4 Context

Distinguishing acceptable from unacceptable insider behavior is difficult in part because of context. A disclosure may be acceptable only to certain people, at a certain time, in a certain location, in the presence (or absence) of certain other people, if certain other conditions obtain, for one time, and only if the recipient has not already obtained certain other data. Although such complex access control rules can be modeled and implemented, these rules go well beyond the subject–object–mode paradigm traditionally used for access control. These complex rules reflect the factors employed daily in personal data sharing decisions (between people, not involving computers), computer scientists do not even know the full set of parameters on which access control decisions are based outside of computers; thus it is premature to expect their implementation in most computing systems.

In fact, physical security recognizes a need for two kinds of systems: automated, mechanical systems that are unforgiving (such as gates and badge readers),

and human overrides that can exercise judgment (such as dealing with the lost or forgotten badge or allowing emergency access by medical personnel). Acceptable behavior can be similarly rigidly determined by a system. But the working of some organizations is highly nuanced and sensitive data are communicated under subjective access control regimes.

These rich, context-based human access control decisions pose a problem for insiders: To share computerized data in those ways may require going outside or around the system, that is, violating system access controls. This is one of many examples in which insiders need to go outside or around the system's controls in order to accomplish needed goals.

3 Insider Threat Issues

Research on insider threats has several limitations. First, there is only meager data on inappropriate insider activity. Second, it would be very useful to probe the minds of insiders to determine what makes an insider good or bad. In part because of limited data, and in part because of limitations of current psychology, success in this avenue may be narrow. Third, the way to determine what insiders are doing is to monitor them, but monitoring of users or employees has privacy implications. Finally, technology is important in many areas of computer security, but the insider threat may be one for which the uses of current technology are somewhat incomplete.

3.1 Data

Research on the insider threat is hampered by the definitional problems described above. An even more serious limitation is the scarcity of data.

Scarcity of data seems puzzling in light of various comments at the workshop. One participant said the majority of insider attacks are undetected. That statement seems self-contradictory: If the majority of insider attacks are undetected, how can we know those attacks constitute a majority? Another researcher reported on a study to try to analyze behavioral intent using host-based sensors. The researcher acknowledged that the work had both false positives and false negatives. But here again, knowing or asserting that a system produces false positives and false negatives almost implies that we know the true positives and true negatives in order to be able to classify other events as false. Another participant noted that people use USB devices to transport data avoiding access controls. When another participant asked if there were studies to back up that assertion, the first replied that the report was merely anecdotal.

As a community we assert certain points, but in the realm of insider threat and insider behavior some of our assertions are hunches. Repeated enough times, hunches become accepted as fact.

Obtaining accurate data on the insider threat is difficult for several reasons, including

- imprecise definitions (as previously discussed)
- unclear policy of what constitutes welcomed or allowable insider behavior versus what constitutes behavior to be discouraged or prohibited
- massive amounts of data: assuming that the number of acceptable insider actions is far larger than the number of potentially negative insider threat actions, large amounts of uninteresting data will have to be filtered out
- reticence to share: because of laws, image, morale, and other factors, some organizations who collect data on insider activity are unwilling to share data with other organizations
- privacy concerns that limit data collection and sharing

The absence of good data limits researchers' ability to analyze, hypothesize and validate. One researcher went so far as to say that researchers need data to address problems; if organizations are not serious enough to supply researchers data, they (the organizations) aren't treating their problem as serious.

One source of data are police and court records. Cases are usually in the public record and details of the crime are reasonably complete. However, these records present a biased view of insider threat. First, police are involved only in crimes. As described earlier, insider behavior can sometimes be positive, so the police will not be involved. Even when the behavior is negative in some cases the companies will let the insider off with a warning or at most a dismissal. And some kinds of insider malicious activity are not criminal offenses. Second, some companies choose not to prosecute even in the case of a crime, fearing the negative publicity. Furthermore, faced with many crimes, district attorneys sometimes put computer crime cases low on their priority list, especially cases in which the loss is intangible and not huge, because of the complexity of prosecuting the case and the consequently low probability of winning. Finally, crime statistics typically cover only a single country (or other jurisdiction, such as a city or district), meaning that insider attacks against multinational companies may be hard to track. For all these reasons, criminal data must be viewed in context.

3.2 *Psychology*

In the workshop several speakers cited a need for a psychological component to insider threat study. There were basically two directions to the work involving psychology: profiling and motivating.

Some participants wanted a psychological profile of an insider who was likely to offend (and preferably before the offense). More than one person wanted to know how to identify potential insider attackers (before they attack, and ideally before they are hired).

For years the criminal justice system has unsuccessfully sought the profile of the criminal. Criminologists are not even close to identifying criminals reliably in advance. It seems as if criminals are varied in their motivation and psychological makeup. We may be able to identify some very antisocial personalities, but other criminals elude advance detection. The possibility of false positives hampers these efforts. If we have been unable to identify serious criminal intent or behavior, why should we expect to be able to identify insider threats?

Complicating psychological identification is that we send mixed signals to insiders. We praise creative individuals, ones who are adept at making a recalcitrant system work. Initiative, industriousness, and problem solving are positive traits on employee reviews. So we should not be surprised when an insider uses these traits to be more productive.

We do not know if insiders expand their threat activity, first for nonmalicious purposes and then gradually to more malicious ends. Consequently we do not know if our rewarding unorthodox system use actually starts insiders on a path to malicious behavior. The situation is probably far more nuanced than this description.

Psychological screening would be ideal before an employee is hired. The typical job interview lasts no more than one day, and it involves both trying to get a sense of whether to hire the potential employee and at the same time convincing the employee to accept a job if offered. An intense psychological evaluation rigorous enough to identify potential inside attackers might be off-putting to non-attackers who should be hired. And time spent evaluating the candidate psychologically reduces the time to assess whether the person would be an asset to the organization. So, even if a psychological exam were available, its use might be counterproductive.

Prospects do not look good for developing psychological profiles. We have too little data (too few cases) with which to work, we do not have a good understanding of the norms of acceptable behavior, we are not sure where is the boundary between acceptable and unacceptable behavior, and we must be able to address many different motivations for unacceptable behavior. Perhaps when we understand general human behavior better we will be able to develop useful profiles.

The other major use for psychology is positive: developing ways of reinforcing good behavior. Some participants wanted to understand how to use psychology to keep insiders acting in positive ways. The prospects seem more promising for this use of psychology than for profiling.

The difference between profiling and motivating is that we want profiling to be precise, generating few false positives and false negatives (because the risk of a false positive is not hiring a potential good employee or holding back or dismissing someone who has not yet—and might never—exhibit harmful behavior, and the risk of a false negative is failing to prevent or detect an attack). If a motivating

technique is largely effective, meaning that it serves its desired purpose on a significant enough proportion of people, it is deemed successful. We can afford to use several motivational techniques that work for different people.

3.3 Monitoring and Privacy

Privacy concerns significantly limit data collection and psychological modeling. Again, the definition of insider becomes important.

When the insider is an employee, privacy rights are subordinated to business rights. The courts have consistently upheld the right of a company to monitor employees' behavior, as long as there is a reasonable business purpose for the monitoring and the monitoring does not violate basic human and civil rights. Thus, companies can generally capture and analyze an employee's email and other communications that use company equipment, log all files and other resources an employee accesses, and retain copies of programs and data an employee creates under the company's auspices. A company is far more free in tracking its employees' system activities than would be law enforcement, for whom probable cause and a search warrant are needed.

But not all insiders are employees. Some definitions of insider include people such as account holders who access their bank accounts, patients who use an electronic system to communicate with medical professionals or view or manage their medical records, students at universities, customers who use online shopping systems, and similar users. Each of these users has certain authorized access to a system. Privacy for some classes of users is covered by laws, such as HIPAA for patients in the United States or the European Privacy Directive for many accesses by Europeans. In other cases, the privacy regulations allow monitoring, data collection and retention, and even data sharing if it is documented in a privacy policy (and sometimes even if not). In these cases, then, privacy rights vary.

Regardless of whether the company has the right to monitor its users' actions, some companies choose not to monitor because of possible negative public opinion.

Another type of insider is the business partner, consortium member, subcontractor, or the like. In these cases, privacy rights are even weaker than for the category of users. The contract between partners may spell out rights to track behavior, although not all such relationships are covered by a contract.

So, is monitoring of insiders' activity permissible? Perhaps and sometimes. Is it desirable for the organization to perform? Perhaps and sometimes. The other important question is whether the monitoring is effective.

3.4 Detecting Insider Attacks

Insider attacks are difficult to detect, either by human or technical means. One workshop participant observed that most insider attacks are detected only because of some reason to suspect: the insider may have talked (bragged) about the act, for example. In other kinds of crime police investigators sometimes profit from a perpetrator who does something to draw suspicion.

An insider attack recognition tool would be useful to flag attacks or suspicious behavior in time to limit severity. Clearly most insider activity is not malicious; otherwise organizations' computer systems would be constantly broken. Thus, the volume of nonmalicious insider activity far outweighs that of malicious activity. Such volume of data is hard to analyze.

A similar example is an intrusion detection system protecting a system from malicious network access: Most network traffic is benign. Intrusion detection technology is improving all the time. However, intrusion detection systems are best at finding specific examples of inappropriate access, either because the access fits a pattern of known malicious activity or because the access touches specific sensitive resources in unusual ways. The hardest attack for an intrusion detection system to recognize is one composed of pieces spread across a long period of time. For those attacks the intrusion detection system has to collect and correlate pieces of data over time, which implies a long window of data comparison.

Inside attackers presumably will perform both normal and malicious acts, which complicates the search for anomalous activity beyond that performed by an intrusion detection system.

One important question raised, then, about monitoring to identify inappropriate behavior is whether the monitoring is effective. Intrusion detection techniques may be of some value. But because there is so little published research on insider attacks, it is impossible to tell whether monitoring helps. Monitoring is useful to confirm a suspected case of insider attack. There is controversy as to whether monitoring serves as a deterrent; that is, if insiders know their activity is being monitored are they less likely to engage in inappropriate activity? The answer to that is unknown, although one workshop participant noted that monitoring is not effective to deter retail theft by employees. Another participant said that detection of a data leak is unlikely unless there is some trigger that makes the leak prominent.

3.5 Technology

What technology is available to detect, deter, or prevent insider attacks?

Most existing computer security technology is based on the concept of a perimeter defense. The attackers are outside the line, the defense blocks the attackers, and the sensitive resources inside are safe. Firewalls are the classic perimeter

defense. With insider attacks, drawing the protection line is more difficult because the attacker and the sensitive resources are on the same side of the line.

Intrusion detection systems may be of some value in detecting insider attacks. As previously discussed, these systems need to analyze a large amount of data, correlate pieces of data potentially spread over a long period of time, and distinguish malicious from nonmalicious intent. These three factors exceed the current demands on intrusion detection systems.

Operating systems, access controls and audit logs offer little support for controlling insider threats, because the insider is acting within limits of authorized behavior, just doing inappropriate things with that allowed access.

But more to the point, technological approaches may be wrong for dealing with the insider threat. The basic element of the insider threat is human: a perpetrator has abused a position of trust. The insider is part of the organization and has some loyalty to it. Capitalizing on the human aspect includes determining what kinds of people are likely to abuse trust or creating an environment in which people would not want to abuse.

4 Conclusions

Little is known about the insider threat. Even computer security professionals use different definitions for insider, and so it is not surprising that the general computing field, as well as the general public, can offer little insight into the problem of insider attacks. Add to this organizational reticence to be embarrassed in public and it is not surprising that there is little valid measurement and reporting on insider attacks or the insider threat.

Several significant points evolved at the insider threat workshop:

- The term insider must be clearly defined. That definition must be communicated to security professionals, computer professionals outside of security, management, and the general public.
- The term insider attack must be clearly defined. That definition must be communicated to security professionals, computer professionals outside of security, management, and the general public.
- Data from reliable measurement of insider activity—malicious and not—must be gathered and shared within the security research community.
- Cooperation with industry is necessary: Industry has the insider attack cases, but security researchers have the tools and inclination to analyze the data. Each side needs the other.
- Developing a psychological profile of a likely attacker is an attractive goal. Because of variation among human motivations, and limitations in the knowledge of psychology, such a profile may prove elusive.

- Psychology may be more effective at finding positive controls: conditions that make it less likely that an insider will want to harm the organization.
- Technical controls to prevent, detect, or deter malicious insider behavior will be difficult to develop. The insider exploits legitimate access. Limiting such access may have a negative effect on nonmalicious employees' productivity.

Acknowledgments

I would like to thank the conference organizers for inviting me to attend this workshop and to write this overview paper for the proceedings.

Many of the basic ideas presented here come from the workshop participants, although not quoting from them directly, I have not identified them by name, because they have not had the opportunity to view or rebut anything I write about their ideas here. Conclusions and elaborations are my own.