

Defining the Insider Threat

Matt Bishop
UC Davis
Davis, CA
bishop@cs.ucdavis.edu

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

1. INTRODUCTION

Many diverse groups have studied the insider threat problem, including government organizations such as the Secret Service, federally-funded research organizations such as RAND and CERT, and university researchers. In addition, many industry participants are interested in the problem, such as those in the financial sector. However, despite this interest, no consistent definition of an insider has emerged.

We argue that the lack of a consistent definition of an insider hinders research in the detection of threats from insiders. In particular, a definition of an insider is first required in order to ensure that the research is, in fact, detecting threats of the desired type. Further, through the development and use of a consistent definition of insiders, it is possible to then compare different detection approaches to determine the best approach for detecting particular types of insiders.

In this paper we propose a definition of an insider that can be extended across various domains and that takes into consideration both cyber and physical security issues. While the majority of papers consider a binary approach to an insider — an attacker is an insider if he is inside some definable perimeter — we adopt a lattice approach that combines the access that is required to be considered an insider within a particular domain. As access control domains can be mapped to the degree of potential damage and therefore the level of threat, the result is a more nuanced definition of an insider that indicates both where detection should focus and the degree of insider threat any one person presents.

2. JUSTIFICATION

Many researchers have investigated the problem of insider threat, however most papers have not precisely defined an insider, instead assuming that the user inherently understands the definition. Without a consistent definition of an insider, each researcher develops their own definition that is particular to their own data set, situation, biases and assumptions.

As a result, research into the detection of insider threats can not necessarily be applied from one domain to another as the underlying model does not necessarily translate between the domains. This situation is further complicated by the existence of definitions that can even be contradictory. For example, a RAND report defines an insider as “an already trusted person with access to sensitive information and information systems” ([2], p. xi). Elsewhere it defines an insider as “someone with access, privilege, or knowledge of information systems and services” ([2], 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 ([3], p. 3), again ignoring trust and also knowledge of the systems.

The problem of defining an insider is further complicated by the assumption of a perimeter that can be defined, 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.

3. OUR APPROACH

We extend the definition of an insider presented by Bishop [1]: “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 argue that a security policy is represented by the access control rules employed by an organization. An insider can thus be 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 root-level access to the system (e.g., by exploiting some system vulnerability such as a buffer overflow or race condition).

In previous definitions, a rule-based system has been used to determine who is an insider. This results 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. For example, Alice and Bob might both have privileges within an organization, and thus both represent “insiders” in the binary sense. However, for some resources Bob might have more privileges than Alice, and so therefore is *more* of an insider than Alice with respect to those resources. At the same time, Alice may be more of an insider than Bob with respect to other resources. This definition extends to physical as well as cyber security. For example, if there is a concern that printed documents might leave a building, then the rules used to define an insider would include access to paper printouts. A janitor is therefore an insider, while Alice and Bob, both of whom work remotely, are not. Thus we call someone an insider *with respect to access to some data or resource X*. By using such a definition, both researchers and security personnel can focus their efforts on detecting those insiders that are likely to cause the most damage to an organization by focusing on those resources of greatest value.

In order to capture this notion of insider as a function of access to data or resources, we propose a model that we call *group-based access control (GBAC)*. This model is a generalization of role-based access control (RBAC). GBAC assigns rights based on general attributes that may or may not be included in a person’s job function, rather than on the specific job functions a person has within an organization. For example, one “group” might be the set of people who come to work after 5:00PM. A second “group” might be the set of all system administrators (in which case this group is also a role). An insider attack may arise from attributes other than job function (such as being in the building after 5:00PM). GBAC can capture entities with the same attributes. RBAC would require the attributes to be job functions to do so.

The attributes that concern us are descriptions of the protection domain of entities. Here, we mean “protection domain” in its broadest sense, not simply a technological listing of rights from a C-List or ACLs. So, the protection domain can include access rights to resources (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. The organization must do a cost/benefit analysis to assign a value to the protection

domains. This might be a single number (producing a linear ordering) or a vector (producing a partial order). 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. Note that, once ordered, the protection domains can be combined into groups (containing a *contiguous* set of access control settings so that the order is maintained), where the group indicates the threat level a particular set of attributes represents. Call these *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. There is a natural ordering of groups based on set containment.

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 if they are ordered; if not, we can determine how they compare by establishing the distance of each from the least upper bound using some suitable metric.

Note that the groups of users differs from those users with the same job function, as use of RBAC would imply. Perhaps such an aggregation can be performed once the lattice has been created. However this is not done *a priori* as it is often the case that some users may have exceptions to their role (e.g., they either have access to additional resources, or they lack access to some resources). Additionally, users may be found to be even more diverse than their roles would imply as attributes beyond RBAC resources are added to the lattice.

The creation of such a lattice requires a two-stage approach. The first is to determine what are the important components of the protection domain relevant to some privilege (including physical access, or lack thereof). 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. For example, access to a particular printer or computer system might not be important, however the ability to print a particular document on that system might have value. A quick initial approach to determining the relevant parts of the protection domain for a system are to combine determining what a person can do (e.g., using standard access control systems) with determining the accesses that are relevant to compromising some resource.

Concurrent with determining the protection domains for the lattice is the requirement to identify all users. Initial users include not only direct employees, but also all contractors and out-sourcers (technical, clerical, janitorial, etc.), as well

as any “special case” access (such as facility visitors or guest logins).

Once the protection domains and users have been identified, the second stage is to map the two together based on the access the users have. This can take either (or both) of two approaches:

1. Determine what a person can do. This thought process is similar in nature to that used when creating capabilities.
2. Determine who has access to a resource. This thought process is equivalent to that used in creating an access control list.

Given these two approaches to determining access, an initial lattice can be created using any existing access control or capability systems. However, it needs to be recognized that this alone is insufficient to define any given threat.

This approach has advantages when applied to computer forensics, in addition to demonstrating where insider threat detection should focus. In the first case — determining what a person can do — can be used for those investigations where the person who accessed a resource needs to be determined. In the second case — determining who has access to a resource — allows you to determine what an attacker did.

4. CONCLUSION

In this extended abstract we presented an initial approach to defining insiders, and hence the insider threat problem. While the majority of research implicitly defines an insider as a binary condition (one is either an insider or not), this paper takes 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 those insiders who can cause the greatest amount of damage to an organization.

5. REFERENCES

- [1] Matt Bishop. Position: Insider is relative. In *Proceedings of the New Security Paradigms Workshop*, 2005.
- [2] R. Brackney and R. Anderson. Understanding the insider threat: Proceedings of a march 2004 workshop. Technical report, RAND Corporation, Santa Monica, CA, March 2004.
- [3] J. Patzakis. New incident response best practices: Patch and proceed is no longer acceptable incident response. Technical report, Guidance Software, Pasadena, CA, September 2003.