

Aspects of Insider Threats

Christian W. Probst, Jeffrey Hunker, Dieter Gollmann, and Matt Bishop

Abstract The insider threat has received considerable attention, and is often cited as the most serious security problem. It is also considered the most difficult problem to deal with, because an “insider” has information and capabilities not known to external attackers. The difficulty in handling the insider threat is reasonable under those circumstances; if one cannot define a problem precisely, how can one approach a solution, let alone know when the problem is solved? This chapter presents some aspects of insider threats, collected at an inter-disciplinary workshop in 2008.

1 Introduction

The “insider threat” or “insider problem” has received considerable attention [2, 13], and is cited as the most serious security problem in many studies. It is also considered the most difficult problem to deal with, because an “insider” has information and capabilities not known to other, external attackers. However, the term “insider threat” is usually either not defined at all, or defined nebulously.

The difficulty in handling the insider threat is reasonable under those circumstances; if one cannot define a problem precisely, how can one approach a solution, let alone know when the problem is solved? It is noteworthy that, despite this imponderability, definitions of the insider threat still have some common elements. For

Christian W. Probst
Technical University of Denmark, e-mail: probst@imm.dtu.dk

Jeffrey Hunker
Jeffrey Hunker Associates, e-mail: hunker@jeffreyhunker.com

Dieter Gollmann
Hamburg University of Technology, e-mail: diego@tu-harburg.de

Matt Bishop
University of California, Davis, e-mail: bishop@cs.ucdavis.edu

example, a workshop report [4] defined the problem as malevolent (or possibly inadvertent) actions by an already trusted person with access to sensitive information and information systems. Elsewhere, that same report defined an insider as someone with access, privilege, or knowledge of information systems and services. Another report [12] implicitly defined an insider as anyone operating inside the security perimeter—while already the assumption of only having a single security perimeter may be optimistic.

In 2008, a Dagstuhl seminar on insider threats brought together researchers and practitioners from different communities to discuss in a multi-national setting what the problems are we care about, what our response is, which factors influence the cost of dealing with insider threats and attacks, and so on. In a time where we barely understand which factors cause insider threats, and our solutions are scattered all over communities, areas, and instruments, this coordinated action between the involved communities seems to be needed more than ever.

This chapter presents some of the results of that workshop, where also the idea for this book was born. Many of the aspects identified in this introductory chapter are touched upon throughout the book. An earlier version of this chapter, as well as more information on that seminar, is available from [9].

2 Insiders and Insider Threats

One of the most urgent quests for communities dealing with insider threats is identifying the characteristic features of an insider. One approach for doing so is to look at recent insider threat cases, and try to find individual or common properties. This is an important step, since insider threat cases can be rather diverging—take for example cases such as Binney vs. Banner [1], a message flood created as consequence of a security bulletin [11], spies that stole secrets for the Chinese Army [7], or a tax authority employee who used her influence to embed backdoors into taxation software [10] (see boxes below for short summaries). While these cases could not differ more, they serve the purpose of illustrating the widely differing characteristics of insider threats.

The wide range of properties that can characterize insider threats recently has led to the development of taxonomies, which are discussed in Section 2.2. Especially Case 2 (message flood) is interesting, as it seems unclear whether this really is an insider case, and if yes, whether it was the deed of a single insider, or a confluence of several actions by insiders. Case 4 (taxation software), on the other hand, seems typical for an employee who is an insider, but needs to “break into” the system to reach certain goals.

To be able to deal with cases so divergent, one clearly needs 1) a common vision of how insiders can be categorized; and 2) security policies for countering insider threats, and ways to evaluate the impact of alternative security policies.

From analyzing cases such as the above, several approaches to identifying an insider can be developed:

Example 1: The Hard Disk Example: Naive user and absent policy

On April 5, 2003, Banner Therapy (a small privately owned company in North Carolina, USA) employee Christina Binney was discharged from her position for “misconduct”, and instructed not to return to the office.

Christina Binney was also a co-founder of Banner Therapy. According to Banner, there were two reasons for Binney’s dismissal. First, the company disputed her assertion of copyright interest in the company catalogue and website. Second, the company claimed she impermissibly removed from her work computer a hard drive that she took home over the weekend to prepare for a client meeting. The company claimed that the disk drive removal crippled Banner’s operations and placed vital company data at risk. Binney explained that a Banner customer requested a meeting on a Friday for the following Monday morning. To prepare for the Monday meeting, Binney chose to physically remove the entire hard drive from her work computer to use with her compatible home computer, rather than take time to transfer the files to a disk.

At the time, Banner Therapy had neither company policy about taking work equipment home nor established computing protocols. When Binney attempted to return to work on Monday, she was denied access; this inability to enter the workplace prevented her from returning the hard drive as she claimed she intended to do.

Example 2: The Email Example: Ordinary user generates an extraordinary amount of email

In early October 2007, Alex Greene was changing jobs. In preparation for the switch, he wanted to update his subscription to a Department of Homeland Security intelligence bulletin by changing his designated email address. In doing so, he mistakenly hit “reply all”, and touched off a listserv free-for-all when his request arrived in the electronic mailboxes of several thousand government and private sector security specialists. The result was what commentators described as a mini-distributed denial of service attack. There were more than 2.2 million emails pinging among approximately 7,500 recipients before the email server was forced to shut down.

The information contained in the bulletin is unclassified, but nevertheless, the decision to respond inadvertently compromised classified contact and departmental information. Individual subscribers with security classifications remained anonymous until they also hit reply, responding from work accounts that included automatically generated signatures. Indeed, one poster pointed out that, armed with the information contained in auto-signatures, he was one fake letterhead away from impersonating a Department of Defense employee.

Example 3: The Trade Secret Example: Malicious user steals trade secrets

On June 16, 2007, FBI agents, using a sealed grand jury indictment, entered two luxury homes in Silicon Valley and arrested a pair of engineers. Both Lan Lee (an American citizen) and Yuefei Ge (a Chinese national) had worked for NetLogic Microsystems (NLM) until July 2003. The two men used money from mainland China to create and incorporate a company for the sole purpose of exploiting the secrets they stole.

Lee and Ge downloaded sensitive NLM documents onto their home computers. NLM data sheets are “top-level confidential technical descriptions of their products”, including information described in enough specificity to enable someone to produce the technology. Together, the men accumulated the information needed to design and produce their own lines of microprocessors and microchips. To finance the business they were creating, the men contacted Beijing FBNI Electronic Technology Development Company Ltd, and entered into an agreement to develop and sell micro-processor chips. Both men were able to access proprietary information without exceeding their individual authorizations.

By late September investigators had uncovered evidence that the venture capitalist had ties to the Chinese government and military.

Example 4: The Tax Fraud Example: Perimeter definition and system design

The District of Columbia (as of summer 2008) is pursuing a case against Harriette Walters and her co-conspirators, for perpetrating the biggest fraud in the city's history. Until her arrest, "Walters was a 26-year tax employee known among her colleagues as a problem solver with a knack for finding solutions by using the department's antiquated and balky computers or finding a way around them." She allegedly used her position to produce fake checks for bogus refunds with fictitious names; the total is said to exceed (USD) \$50 million.

The scheme involved Washington's new Integrated Tax System. During design phase, Walters "contributed to the decision that her unit, which handled real estate tax refunds, be left out of it." At the time, the decision seemed to make sense. D.C. had spent \$100 million to implement the business and income parts of the system, and it had only \$5 million remaining for implementing the real estate tax portion. So the system's perimeter was defined to omit real estate tax processing. That design decision allowed Walters and her co-conspirators to create bogus tax refunds with fictitious names that were not checked against actual real estate records. Some refunds were issued multiple times; the recipient (often someone's boyfriend) would claim that the check was never received, and a new one was issued—with interest to compensate for the long delay! The schemes exploited several loopholes: each check was under the \$40,000 threshold for requiring a supervisor's approval, and no action was taken to cancel the first check or confirm that it had not already been cashed.

- An insider is defined with respect to a resource, leading to "degrees of insider-ness";
- An insider is somebody with legitimate access to resources;
- An insider is a wholly or partially trusted subject;
- An insider is an individual who has or had access to resources;
- An insider is a system user who can misuse privileges;
- An insider is an individual with authorized access who might attempt unauthorized removal or sabotage of critical assets or who could aid outsiders in doing so; and
- An insider is a person or company whom we trust.

These definitions immediately lead to a series of discussions on what is meant by "access" (code, credentials, timing of access rights), whether an insider is sufficiently defined based on resources or whether a definition should take the system into account, and how the definition relates to a masquerader, namely an outsider being able to trick a system into believing he is an insider.

Exploring these aspects enables us to reason about what makes a good insider:

- Knowledge, intent, motivation;
- Possesses power to act as agent of the business;
- Knowledge of underlying business IT platforms;
- Knowledge/control over IT security controls; and
- Ability to incur liability in pecuniary terms or in brand damage or other intangible terms.

The skill of insiders is also an important a factor defining the threat posed by malicious insiders, or non-malicious insiders just trying to get their job done. "Motivation" in general is an important question when dealing with insider threats and their consequences. This can cover the whole range from "innocent action", "fun",

“technical challenge”, “criminal intentions”, to “espionage”, or a combination of each of these factors. Surprisingly, even though one would expect the contrary, the effect of actions can be equally devastating for each of these motivations. This, of course, makes detecting a threat even more important—but also more complicated. A key observation is that the definition of an insider for threat purposes is different than the definition for business purposes.

Based on the aspects defined above, one can in turn decide how to define an insider, namely in terms of someone with:

- Knowledge: Implies an open system, one that remains secure (if at all) even with full knowledge of the system operation; alternatively, security through obscurity; or
- Trust: An individual is empowered by the organization to be an insider; or
- Access: An insider is in possession of a credential giving access to the system — an IT centric perspective, since the system in general does not know who possesses the credential.

At the end of the Dagstuhl seminar [9], a trust-based definition of an insider was proposed:

“An insider is a person that has been legitimately empowered with the right to access, represent, or decide about one or more assets of the organization’s structure.”

The rationale behind this definition is that it removes any specific IT bias from the definition, it focuses on organizational assets rather than a narrow approach based on system credentials, and while people that constitute threats may not be entrusted access to credentials, they might still have the ability to decide (based on policies) and to represent the organization.

The ability to represent is rather important, as the policies imposed on an actor inside an organization in general are not known to the outside, where the same actor can pretend to be subject to completely different policies—a factor rarely ever being checked.

“Knowledge” by an individual (*e.g.*, the knowledge of the person who originally designed the system, but is not part of the organization or in any way associated with the organization anymore) is not a good way of capturing what is an insider.

2.1 Insider Threats

A natural consequence of having defined the term “insider” is to consider the term “insider threat”. As discussed in the previous section, again many different aspects can be important:

- Risk to organization or organizational resources posed by actions (the entity as agent of the business);
- Access type or system role;
- Aim or intentionality or reason for misuse;

- Level of technical expertise;
- Type of malicious insider behaviour or system consequences; and
- Threats from masquerader, traitors, and naïve insiders.

The last point of masqueraders (individuals pretending to be legitimate insiders but without valid access), traitors (legitimate insiders acting in malicious ways), and naïve insiders (who cause damage without malicious intent) is very closely related to the question of motivation discussed above. The problem of interest is dealing with the “real real insider”—an individual deeply embedded in an organization (*e.g.*, a high level executive, or a systems administrator). Detection techniques for each of these types of insider threats will vary.

The impact of insider threats can occur in multiple dimensions: financial loss, disruption to the organization, loss of reputation, and long term impacts on organizational culture. These impacts can be highly nuanced, and are neither well or easily measured or accounted for. An important aspect here is that “bonus round” insider threats (actions taken in anger, revenge, spite regarding bonuses, compensation) can have severe consequences on all levels of an organization. Thus a rather “small” or meaningless motivation for the individual can have a rather huge impact for the organization. Equally, the impact may not depend on motivation—an innocent act can have as devastating an effect as a maliciously motivated attack. The goal of detecting insider threats may therefore be to avoid catastrophic consequences regardless of the motivation. These aspects as well as other risk accelerants should be represented in threat models, to acknowledge their importance.

2.2 Taxonomies

In order to allow identification of insiders and insider threats across organizational boundaries, an effective taxonomy of insider threats is a necessary foundation for further work by both researchers and practitioners. Taxonomies provide a means to order problems; in the problem considered here such ordering is necessary both to differentiate types of insiders and types of insider threats, and to make explicit the key dimensions which serve as the basis of the differentiation. By identifying the key dimensions, we can then begin to systematically build prevention and response strategies.

At least some common acceptance of the dimensions of such a taxonomy is needed; experts can disagree about how the dimensions are applied (as in defining who is an insider), but by forcing an explicit discussion of different interpretations, taxonomies can serve a vital role.

Examples for taxonomies specifically aimed at insiders can be found in [14] and [3]. Pred *et al.* [14] observe that insider threats can be defined in terms of four reference perspectives, namely the organization, the individual, the information system, and the environment. This results in a “holistic” or top-down taxonomy. In contrast, Bishop *et al.* [3] define insiders with respect to a resource. Resources are defined as pairs of the resource itself and access privileges, and insiders are defined

with respect to these pairs. With this structuring, it is possible to define degrees of insiderness.

An important question when considering taxonomies is whether a single taxonomy should be adopted by the community. In some cases, like the taxonomy for speciation (kingdom, phylum, ...) a more or less universal adoption has provided great value, but it can also limit a frameworks expressiveness.

Key factors important as determinants of the insider threat may be difficult to categorize a priori. As noted elsewhere, knowledge of the insider's intent is desirable, but requires all-embracing knowledge. Each determining factor for an insider can be used for defining a taxonomy, for example based on:

- The distinction between malicious and accidental threats;
- The distinction between doing something intentionally (for malice, or good reasons which nonetheless may result in damage) versus events that occur accidentally.
- The distinction between obvious and stealthy acts.
- Acts by masqueraders, traitors and naïve or accidental use that results in harm.
- A combination of factors such as access type; aim or intentionality or reason for misuse; level of technical expertise; and the system consequences of insider threats.

3 Detection and Mitigation

Forensics appears to be highly undeveloped when addressing insider threats. Insider behaviour may be close to the expected behaviour, and the still often used audit trail is generally inadequate (redundant, misleading, missing data), and often lacks time correlation. The number of appropriate characteristics to observe may be large, resulting in overwhelming amounts of data. While we have decent tools as the result of a large body of work on intrusion detection, it is unclear how these tools help with insider threats. Current forensics tools often require assumptions such as “only one person had access” or “the owner of the machine is in complete control”. Therefore, forensics remains an art, and as an art questions such as what to log or determining the relevance of log data elude clear answers.

Detection, forensics, and response must also wrestle with how to distinguish motive and intent. Malicious acts may be equivalent to acts due to accidents or naïveté. Insiders may legitimately use domains in unexpected ways that might trigger false alarms. Outsiders, insiders acting with malicious intent, insiders acting without malicious intent, and accidental behaviour may all result in similar effects on the organization. Hence there are always going to be gray areas in how security policies define both insider misuse and proper behaviour. Furthermore, actions are context bound, but most security policies only inadequately capture the nuances of context. **Monitoring.** While monitoring can help with technical aspects, it does potentially worsen behavioural aspects. The deciding factor is how much monitoring is acceptable (both ethically and legally), and whether it is at all beneficial. The noteworthy

point here is that this question arises at all levels in an organization, from individual actors, to groups, to companies, to the society as a whole. The problem is that not only may the same actor have different opinions depending on at which level he is asked, but also that different answers for different individuals may exist at the same level.

An interesting observation is that in certain settings with significantly enhanced monitoring, the number of identified incidents has stayed almost constant. At the same time, and even more worrying, cases such as Kerviel and the Liechtenstein case [8, 15] had in common that the attacker intimately knew the monitoring system and knew how to play it. It is often hypothesized that malicious insiders seek to avoid setting off monitoring alarms by slowly adjusting their profiles, but it seems unclear how easy current behavioural systems can be tricked.

In summary, trust in insiders is a behavioural expectation that still needs to be controlled. While the easy solution to reducing the number of insider cases would be to remove all restrictions (making the illegal actions legal by changing the semantics of the term “legal”), we aim for making the monitoring as efficient as possible, where in different situations the term “efficient” may have different interpretations. An important aspect that can not be underestimated are legal restrictions and privacy aspects of data collection, which may be even harder to follow in multi-national settings.

The goal of monitoring (or observing in general) should be to only monitor what is needed to identify the threat in question. Since currently trust can often be transferred, for example by handing over a code card, it is important to isolate transferred trust as much as possible, not least to allow the result of monitoring to be used to bind actions to actors.

In large complex systems risk analysis and focused detection require a significant effort. Not only may monitoring affect trust within an organization, it also is highly nuanced what to look for [16]. For example, an inordinate amount of searching may indicate a masquerade attack (the masquerader is less familiar than the legitimate insider about data structures)—or a forgetful mind. Observations on “higher levels of behaviour” not observed by systems monitoring may be useful, for example, by using human intelligence to pick up novel attacks and signals that are out of the system. CERT data suggests [5] that in most cases someone else knew about the insider threat actions going on, so it is important to find ways to encourage reporting of “suspicious activity” by others. Looking for suspicious (different) behavioural patterns by insiders is appealing, but difficult to systematically apply; behavioural patterns include cyber activity, physical movements, physiological signals, and many more. Employment screening data and self/organizational reported data might be useful here, but any screening for behavioural changes is bound to produce false positives from otherwise innocent factors like individual predispositions or lifestyle changes. Fundamentally, the attributes key to insider threat identification will be largely context bound.

From a company point-of-view it turns out to be often preferable to not mitigate ongoing insider attacks for numerous reasons. Here optimistic access control is seen as a viable option, *i.e.*, allowing insider actions to happen until there is no way back,

or even letting them happen unhindered, at the same time ensuring that enough evidence is collected through monitoring. It seems often more ruinous to take systems down because of an ongoing attack than to accept the losses and prosecute after the fact.

Outsourcing. Even more difficult is the handling of outsourced parts of a company, both for technical and legal reasons. On the one hand data may be much harder to obtain (or be less trustworthy), at the other hand the data protection laws regulating data collection in other countries may be vastly different from the laws in the home country.

Another problematic area is outsourcing the auditing itself; while in the “regular” outsourcing scenario it may be difficult to obtain the data in the first place, it now is paramount to protect the already collected data. This means the data should be anonymized as much as possible, revealing only as much data as necessary to allow external auditors to produce meaningful results, but at the same time hindering them from drawing unwanted inferences from the data. One example is to anonymize timestamps to preserve relative ordering between events, but blurring them such that the exact order and timing is lost. It was noted that formal methods can and should be applied in these settings, and some were presented, but at the same time often require significant resources, such that for example before and after the fact application is often feasible, but not online detection.

4 Policies

Policies obviously play an important role with respect to insider threats, as they define the boundaries between permissible and not permissible behaviour, both on a technical and non-technical level, and tie together insiders, insider threats, and detection and mitigation.

Policies not only define proper behaviour, but implicitly also define the notion of insider. It is problematic that policies often are only specified implicitly, possibly leading to large differences between *de facto* policies and “real”, intended policies. To support the externalization of these intended policies, policy languages have been developed, which are usually quite well suited to support technical issues, and at the same time try to add support for non-technical aspects of policies.

When considering policies, one needs to pay special attention to the notions of “context” and “dynamicity”. For example, a given actor might be an insider in one situation, but would be considered an outsider in another. Similarly, for some policies violations in special, emergency cases might be acceptable, but in the general case they should be observed. In this case it would be the insider’s margin of discretion to decide for or against breaking a policy rule. This ties policies as well as their specification and enforcement tightly to human factors, which are discussed in the next section.

Policy Hierarchy. Policies themselves are developed based on three sources: 1) legal and regulatory (so-called best practices); 2) business requirements; 3) security

requirements. All of these sources can result in implicit or explicit policies, establishing a grey zone where behaviour is neither good nor bad. This potential gap is extended and formalized in the Unifying Policy Hierarchy [6], which established four different levels:

- Oracle Policy. Given perfect knowledge, what would policy be? Deals with inherent vulnerabilities.
- Feasible Policy. With imperfect knowledge, implement oracle as good as possible. Deals with configuration vulnerabilities.
- Configured Policy. What the system implements via configuration. Deals with real time vulnerabilities.
- Real Time Policy. Add in security vulnerabilities.

Gaps and conflicts in policies can be considered as a principle factor in allowing insider threats to occur; in some cases because gaps/conflicts create confusion among insiders in terms of “what is right” or “how do I get my job done”; in other instances because gaps/conflicts create opportunities that malicious insiders can exploit.

To acknowledge the risk of gaps between policies, we need an analysis of specifications for gaps and conflicts. Reasoning about insider threats it becomes apparent that policies normally do not make explicit who an insider is—an obvious requirement if we want to be able to analyse their hierarchies and fine-tune their impact. If we have policy-language support for specifying the roles of actors, then one may classify certain requests as coming from insiders, or in general build in context-dependent handling of insiders versus outsiders. For example one might want to be able to express that certain requests may only come from insiders, or on an even more context-dependent level, what degree of insiderness is required for a certain behaviour to be permissive.

Policy languages. A gap of a different kind exists for policy languages. Here the gap exists between the existing capabilities to specify system (and more broadly, organizational) policies, and the needed qualities of policy to adequately prevent insider threats. One of the most urgent needs, already mentioned above, is for policies to be aware of behavioural aspects and context, thereby being able to handle and regulate abstract events. While a “zoo” of policy languages exists, with a vast overlap in terms of what they can achieve, the user often may not be able to write policies, let alone read, understand, and follow them. This is expected to gain growing importance in future interactions between previously independent parties.

As mentioned above many systems define the notion of insider relatively to system boundaries. In the long run we may therefore need domain-specific policy languages, in which for example actions would be allowed only if discretionary circumstances justify their execution.

From a research perspective we often seem to be unaware of the different levels that the same policy may exist in, but instead take for granted that an oracle security policy is provided; given “the” security policy, it is often assumed that this resolves all tensions between organizational culture, work flow, and compliance by (implicitly) enforcing for compliance with security practices for the sake of security.

However, having some policy is not enough, that is deploying security technology and policies does not automatically help in achieving security. This is especially true due to the above-mentioned context-dependency of policy rules.

5 Human Factors and Compliance

When considering human factors and compliance, it seems that most insider threat policies are based on a set of incorrect assumptions, namely, that 1) once someone is vetted, audit and incident management processes will pick up policy violations; 2) that risk assessment processes will pick up changes to individual and group values; and 3) that training and awareness-building education programs will instill desired security culture. However these assumptions never fit to how people actually want to pick up information, or act in the course of doing their jobs within the organization. The inadequacy of many existing security solutions to address real life human behaviour presents us with a set of challenges on how to better incorporate human factors into solutions.

Second, an important problem is to align security policies with organization workflow, or, stated simply, security should support people doing their jobs. Often technological security approaches are not accepted and in fact actively subverted, because they interfere with work flow (*e.g.*, an iris reader with an “unacceptable” delay before allowing access resulted in staff finding other ways of gaining access). Compliance with security policies is hard; to make compliance easy for insiders is absolutely necessary for any successful effort to constrain insider threats. Compliance (defined as efforts users will make for purposes they do not understand or agree with) is limited; getting compliance gets more expensive the closer you get to the limit of peoples’ tolerance for disruptions to their work flow and social interactions. Successful security policies need to demonstrate to insiders the value of security, not just the requirement for security.

Another important observation is that motive and intent matter a great deal, but multiple motivations may map into a single intent. As a simple example consider the act (intent) to prop a door open. The motive for this action might be benign (being lazy, or carrying large packages into the room) or malicious (propping the door open to allow unauthorized persons to enter). Observables may be able to capture the intent but not the motivation. This has important implications on the limitations of monitoring, and highlights again the need to establish context for specific actions.

Understanding and Integrating Human Factors. Criminology can inform insider threat understanding, and within criminology are several theories relevant to insider threats. Earlier theories of deterrence, social bonds, and social learning have been integrated into a theory of planned behaviour: for a crime to be committed a person must have both motive and opportunity. As just noted, motive matters; for the “cold intellectual attacker” when the possibility of punishment is high and the sanction severe potential criminals will be deterred from committing illegal acts, especially when their motives are weak. More generally, the goal of “situational” crime pre-

vention is to 1) make the criminal act appear more difficult; 2) make the criminal act more dangerous; 3) reduce the benefit a potential criminal is expecting to recover; and 4) remove the excuses available to the potential malefactor.

Organizational purpose and management structures affect both security structure and policy. In discussing organizational factors relevant to the insider threat a number of questions must be considered:

- How does trust grow in organizations? In some organizations for example there is lots of trust at the base of the organization but it does not necessarily rise up.
- How can organizations adjust management processes to engender a more positive environment for security? Specifically, how can organizations develop a “reflexive view” that looks at the whole person rather than just as a work resource?
- Whistleblowing: When are organization members comfortable with whistle blowing? Is there a role for technology in extending the whistle blowing capabilities of staff?
- Policy conflict within organizations: It seems reasonable to assume that all organizations have implicit tradeoffs about what is more and less important in their expressions of policy. How can these be made more explicit so that policy and security architectures can more effectively capture these values? Doing so might require a hierarchy of organizational needs like the Maslow hierarchy of individual needs.
- Organizational clustering: how much do organizational units cluster in their values? Are there psychological contracts by group clusters within organizations that can be mapped by looking at risk behaviours?
- How can we build robust policy so that when conflicts do arise they can be resolved efficiently in the best interests of the organization?

Insiders (or people in general) will act unexpectedly. Thus, flagging potential insider threats based on departures from “normal” patterns may lack reliability; monitoring for “out of normal” actions may generate too many false positives. There will also always be “gray areas” in drawing the line between insider misuse and proper behaviour.

Hence, context of an activity matters a great deal in accurately characterizing abusive insider actions. Context is defined in terms of physical, social, cultural, and temporal dimensions. In adding context into the shaping of security policies, the implication is that there are no standard solutions for workable security — what is usable is what fits. We need security policies appropriate for a specific domain (context) but what happens when insiders use domains in unexpected ways? Security controls must be characterized in the context of the activity, and because there will always be gray areas, those defining security controls must resist the temptation to believe that controls can eradicate all risk. Those defining controls must do this with full participation of management. Those enforcing controls must be willing to accommodate managerial discretion in certain settings.

The unpredictability of human behaviour has its implications for the role of trust in an organization. As stated earlier, trust is a behavioural expectation, and trust is only necessary in an organization when behaviours cannot be controlled in all

dimensions. Trust is also transitive, so one could argue that reducing insider threats would require environments where no one is trusted, or at worst only a few people are trusted; in any event transferred trust relationships should be eliminated.

Policies and Human Factors. Policies need to be shaped and evaluated in terms of their human impact. How specific should policies be? There is a perception that there are too many policies. The psychological contract with employees generally means that 1) policies need to be made more manageable, and 2) that there is a need to find a way of testing policies to remove redundant policies. The ideal would be a small set of consistent security policies related to behaviours, and fit with business processes, and organizational values and norms.

A common theme is the need to link the user community (the insiders in the organization) with the policies being developed and enforced. Failing to engage staff in security may be the norm, but this lack of engagement weakens security. Security will only work in organizations where people feel that they are part of a larger community. Organizations could conduct specialized internal exercises with most or all the insiders to identify both the set of useful and acceptable policies, and unique contexts which may result in generalized policies in conflict with organizational needs. Equally it will be key to monitor the implementation of policies “on the ground” by engaging staff and managers on whether policies are appropriate, or interfere with their workflow or culture. Sustained discourse with insiders can help highlight positive examples (of senior executives, for example) and in myth busting; an important goal here is to remove frequently made excuses.

Issues of what we can measure, what is relevant to measure, and how and when we intervene when suspecting threatening insider actions need to take human factors into account. Consider the impact of false accusations of insider threats on both the individual and the organization. Many suspicious activities which can be observed are correlated with insider threat behaviour, but not causally linked. False accusations have multiple deleterious effects: investigative resources are spent, the individuals so accused may quit, seek legal or other recourse (including becoming a real insider threat!), or be affected psychologically, the organization’s culture may be affected, possibly for extended periods. There is, therefore, a need for decision processes to decide when to intervene and how.

6 Conclusion

The Dagstuhl seminar on Countering Insider Threats improved our understanding of what different communities mean by “insider”. As stated above, this knowledge has already during the seminar been used to develop integrated approaches towards qualitative reasoning about threats and possible attacks. Beyond this shared definition of what constitutes an insider, the most prominent outcome of the seminar is the beginning of a taxonomy or framework for categorising different threats. The seminar identified the need for:

- A framework or taxonomy for distinguishing among different types of insider threats;
- Methodologies for assessing the risk and impact of insider threat incidents;
- Incorporating human factors into the development of solutions;
- Better formulations for specifying useful policy at both systems and organizational levels—policy that would be meaningful and applicable to the insider threats deemed most important.

There were some cross-cutting conclusions that emerged from the seminar. The role of trust was discussed in a number of different contexts. In one sense, the ideal security framework for addressing insider threats would eliminate the need for trust—all behaviours would either be defined permissible, or else made impossible to execute. But this model ignores two realities. In any but the simplest settings, context of actions is highly determinative in shaping what is appropriate or needed behaviour. Further, many (most?) organizations would not accept a working environment so rigidly defined as to eliminate the need for trust. Hence, we emerge with the conclusion that trust relationships will be present in most organizations; how to best factor trust into security policies and frameworks remains, however, unclear.

Security, moreover, is context dependent. Security is not achieved by deploying generic (context free) controls. However, the importance of context in addressing insider threats poses a number of challenges. Capturing qualitatively the various situations that might arise in an organization is itself probably impossible, though effective dialogue between those defining security controls and those working as insiders in the organization will certainly help. Hence, insider threat prevention and response has to deal with the reality that controls will not adequately capture all of the behaviours that might be appropriate in a given context. Even if all contexts could be qualitatively described, policy languages and controls are inadequate at the current time to fully capture the range of contexts identified.

Motivation and intent clearly are important in defining insider threats and defining appropriate detection, forensics, and mitigation strategies. While intent (the purpose of actions) is at least partially observable, motivation (the incitement to action) is not. The intent to, for instance, obtain certain data may reflect malicious motives, or may reflect positive motives (as in a hospital emergency where certain information is desperately needed regardless of legitimate access). Devining motivation highlights the need for context-aware policies, but even with context motivations may be difficult to determine. We conclude that approaches for understanding motivation a priori are still highly immature.

Each of these observations emphasizes the conclusion that security will not be achieved solely by deploying security technology. Most people are not entirely logical or consistent in their behaviour, and this confounds our ability to formulate measures to reliably prevent or detect malicious insider behaviour.

References

1. Binney v. Banner Therapy Products, 631 S.E. 2d 848, 850. North Carolina Court of Appeals (2006)
2. Bishop, M.: The Insider Problem Revisited. In: Proceedings of the New Security Paradigms Workshop 2005. ACM Press, Lake Arrowhead, CA, USA (2005)
3. Bishop, M., Engle, S., Peisert, S., Whalen, T., Gates, C.: Case studies of an insider framework. In: Proceedings of the 42nd Hawaii International Conference on System Sciences (HICSS) (2009)
4. Brackney, R.C., Anderson, R.H.: Understanding the insider threat : proceedings of a March 2004 workshop. RAND, Santa Monica, CA : (2004)
5. Cappelli, D.M., Moore, A.P., Shaw, E.D.: A Risk Mitigation Model: Lessons Learned From Actual Insider Sabotage. In: Computer Security Institute, 33rd Annual Computer Security Conference and Exhibition (2006)
6. Carlson, A.: The unifying policy hierarchy model. Master's thesis, Department of Computer Science, University of California, Davis (2006)
7. Cha, A.E.: Even spies embrace china's free market. Washington Post, February 15, 2008. Available from <http://www.washingtonpost.com/wp-dyn/content/article/2008/02/14/AR2008021403550.html>, last visited March 2010.
8. Hawley, C.: The liechtenstein connection—massive tax evasion scandal in germany. Spiegel Online International, 18 February 2008. Available from <http://www.spiegel.de/international/business/0,1518,535768,00.html>, last visited March 13, 2009.
9. Homepage of Dagstuhl Seminar 08302: “Countering Insider Threats”. Available from <http://www.dagstuhl.de/08302>, last visited December 4, 2008 (2008)
10. Keating, D.: Tax suspects guidance on software left d.c. at risk. Washington Post (2008)
11. Kirk, J.: Homeland security e-mail server turns into spam cannon. InfoWorld.com, October 4, 2007. Available from <http://www.infoworld.com/d/security-central/homeland-security-e-mail-server-turns-spam-cannon-924>, last visited March 2010.
12. Patzakis, J.: New incident response best practices: Patch and proceed is no longer acceptable incident response procedure. White Paper, Guidance Software, Pasadena, CA (2003)
13. Pfleeger, S.L., Stolfo, S.J.: Addressing the insider threat. IEEE Security and Privacy 7, 10–13 (2009). DOI <http://doi.ieeecomputersociety.org/10.1109/MSP.2009.146>
14. Predd, J., Pfleeger, S.L., Hunker, J., Bulford, C.: Insiders behaving badly. IEEE Security and Privacy 6, 66–70 (2008). DOI <http://doi.ieeecomputersociety.org/10.1109/MSP.2008.87>
15. Schwartz, N.D., Bennhold, K.: A trader's secrets, a bank's missteps. New York Times, 5 February 2009, New York, USA.
16. Probst, C.W., Hunker, J.: *The Risk of Risk Analysis-And its relation to the Economics of Insider Threats*, Proc. of the Eighth Workshop on the Economics of Information Security (WEIS 2009), June 2009.