Insider Threats: Strategies for Prevention, Mitigation, and Response

Dagstuhl Seminar 10341
August 22nd to 26th 2010

Matt Bishop[1], Lizzie Coles-Kemp[2], Dieter Gollmann[3], Jeffrey Hunker[4], Christian W. Probst[5]

[1] University of California, Davis
bishop@cs.ucdavis.edu
[2] Royal Holloway
Lizzie.Coles-Kemp@rhul.ac.uk
[3] Technical University Hamburg-Harburg
diego@tu-harburg.de
[4] Jeffrey Hunker Associates LLC
hunker@jeffreyhunker.com
[5] Technical University of Denmark
probst@imm.dtu.dk

Abstract. This article summarizes the objectives and structure of a seminar with the same title, held from August 22nd to 26th, 2010, at Schloss Dagstuhl, Germany. The seminar brought together researchers and policy-makers from quite diverse communities, to make progress towards an integrated framework for understanding insider threats and their interaction with organizations and policies. During the seminar, social and organizational factors relevant to insider threats, were discussed, as well as urgent questions in four areas: synthesizing social science and technical research, metrics and assurance, language formulations and ontology, and the threats facing intangible systems. This report gives an overview of the discussions and presentations during the week, as well as the outcome of these discussions.

1 Overview

The Dagstuhl seminar "Insider Threats: Strategies for Prevention, Mitigation and Response" was held on August 22 – 26, 2010 (Seminar #10341, [2]) to advance our understanding of ways of reducing insider threats. The insider threat is cited in many studies as the most serious security problem facing organizations. Insider threats are particularly difficult to deal with because insiders have legitimately empowered knowledge of the organization and its systems, and therefore malicious and benign actions by insiders are hard to distinguish.

The 2010 seminar built on the results of its predecessor from 2008 (Countering Insider Threats, #08302, [1,3]). In this seminar we developed a shared,

Dagstuhl Seminar Proceedings 10341
Insider Threats: Strategies for Prevention, Mitigation, and Response
http://drops.dagstuhl.de/opus/volltexte/2010/2903

inter-disciplinary definition of the insider [1] and a good formulation for a taxonomy or framework that characterizes insider threats. The seminar also began to explore how organizational considerations might better be incorporated into addressing insider threats.

The purpose of the 2010 seminar was to make progress towards an integrated framework for selecting among and evaluating the impact of alternative security policies against insider threats. An integrated framework, we recognized, needs to include issues not considered in insider work before, such as the economics of insider threats [4], and the role of law as both a preventative and punitive instrument. We saw the need for creating and testing alternative integrated frameworks so that practitioners and researchers could make informed choices as to combinations of actions targeted at insider threats, and also the need for methods to evaluate the effectiveness of these actions.

In our proposal for the 2010 seminar we expected the seminar to develop:

- A taxonomy for identifying insider threats;
- An integrated approach that allows a qualitative reasoning about the threat and the possibilities of attacks;
- A deeper understanding of security policies and how to evaluate them.

In fact, as a result of the 2010 seminar we made significant progress in each of these topics, and perhaps even more. We developed a better appreciation of social and organizational factors relevant to insider threats, and addressed urgent questions in four areas: synthesizing social science and technical research, metrics and assurance, language formulations and ontology, and the threats facing intangible systems.

At the end of the seminar we concluded that the management of insider threat problems fundamentally is about dealing with the impact of changes in organizational processes on the effectiveness of internal controls. The nature of 'trust' and the changeability of human behavior make prediction difficult, and create unique security problems:

- In *prevention*: Does the organization trust the insider, or not? There is a duality in the meaning of trust. Trust can strengthen bonds between an insider and an organization; conversely that which is trusted can hurt you.
- In *mitigation*: There is evidence that the majority of people are prepared to cheat if they feel unobserved [5]. Also, over the Internet people are prepared to do things that they would never do face to face. Does this mean that all insiders have to be monitored all of the time, and if so, what is monitored, and does monitoring itself increase the insider threat? Can we really detect signs of changes in the relationship of an insider to an organization?
- In *response*: Flawed reaction can make the problem worse, so how do we ensure to react in an appropriate way?

[1] From the 2008 seminar: An insider is a person that has been legitimately empowered with the right to access, represent or decide about one or more aspects of the organization's structure [3].

Prediction is further complicated by the lack of reliable data. As insider attacks are a sensitive topic, potentially revealing an organization's secrets or damaging its reputation, we have little information about insider attacks, limiting the use of statistical analysis for predicting and recognizing them. Another problem is that even with security mechanisms in place, they will fail every now and then, but it is poorly understood how to detect when they do, and it is not clear either how to best react when they do. In addition, human engagement with technology is emergent and the nature of the emergence is often unpredictable.

Thus, there appear to be limits to our ability to predict and prevent insider attacks. We are resigned to conclude that insider attacks will occur despite our best efforts. The challenge is then to determine what we are able to predict, define the extent of the exposure, and limit it.

We finally concluded that insider threats should be treated as a subset of the larger class of what we call *informed threats*. Informed threats are malicious individuals or groups with special knowledge of the organization and its systems, however acquired, and this larger class increasingly poses the threats of concern. This is because the boundaries of organizations and the dividing line between those with and without system knowledge or trust are ever more dynamic and unclear. However, it is very complex to adapt formal procedures to dynamic, unclear scenarios, resulting in de-facto procedures, knowledge, and processes that are not covered by formal policies and procedures. This is hard to capture in policy development and implementation, as it often is based on internal organizational culture.

The same can be said of how security policies are forced upon an organization. Top down promulgation of security policies often does not match with reality. Effective policy development needs to focus on the relationship between the insider and the organization. How policies are learnt (and unlearnt) by organizations and individuals is a key emerging issue, and also poorly understood.

It is interesting to note that some organizational processes are more 'sustainable' with respect to security than others—that is, they are more resilient with respect to insider threats and more capable of limiting the damage from insider attacks. Resiliency appears to stem from usable, effective, and efficient security having been built into the organizational processes, and having been accepted by the organization's staff.

2 Seminar Structure

In addition to the presentation of papers and subsequent discussion, about half of the seminar's schedule was centered around highly interactive small group discussions, either in working groups organized around specific topics, or in a structured exercise. The plenary presentations and talks thematically supported the working groups.

Working Groups

Four working groups were established on the first day of the seminar. These groups met throughout the week and reported to the entire seminar on at least two occasions. Each group focused on one of the following issues:

- How can we synthesize social science and technical research output to respond to insider threat problems?
- How can we advance language formulations and ontology for specifying insider threat policy?
- What are meaningful and useful metrics? How do we provide assurance that our actions are working? How do we account for and reduce unintended impacts?
- Can intangible systems be subject to insider attacks? What are the problems to modeling these systems, providing detection approaches and solutions?

The working group results are presented in the next section.

Structured Exercise

A key part of the seminar was a structured exercise in which seminar participants were divided into small groups and asked to respond to several different insider threat case studies while alternatively playing different roles in the organization—the security policy maker, the ordinary staff members that have to work with these policies, and the malicious insider.

During the exercises, three main points became obvious—standardized forms for each group to record their recommendations would assist the process, there should be a clearly defined and observed time schedule, and finally we came to appreciate that each case study should either be based on the organization of the participants (if played at a single organization) or, as was the case at Dagstuhl, the case studies should be real life. The latter enables participants to obtain further information through on-line searches. Made-up case studies often lack the richness of detail and context needed to make the exercise effective.

The structured exercise served several purposes. We believe that in an organizational context exercises like this can be an effective means of extracting 'tacit knowledge' about threats and procedures from an organization's staff. Participants agreed that the exercise also encouraged everyone to think about the insider threat from different perspectives. The hands-on nature of the exercise enabled many of the participants to reflect on aspects of the insider threat problem glossed over by more abstract discussions.

3 Key Themes

The discussions at the seminar can be grouped into a few key themes:

- Redefining the problem and approach (working group 1),

- Organizations, staff, and organizational processes (working groups 1 and 3),
- Policies and language specification and ontologies (working groups 2 and 3), and
- Intangible systems (working group 4).

A key theme throughout the seminar was the need to step back and fundamentally redefine both the problem and our approaches to it. While the term 'insider' has value in defining who has legitimate rights within an organization, the once bright line and relatively stable framework for dealing with insider threats does not capture the reality that individuals without legitimate rights nonetheless have system or organizational knowledge sufficient to be security threats, and that organizational boundaries become increasingly blurred and developments in socio-technical engagement enable organizational processes to become more dynamic. This more nuanced characterization of the threat also focused the discussion on the limitations of monitoring and detection, and consideration of other frameworks for thinking about role of security policy.

Closely related to the first theme, discussions about organizations, staff, and organizational processes focused on the fact that insiders are people too, and the organizations affected by insider threats are complex human constructions. As a consequence, social science perspectives ranging from psychology and criminology to economics need to be brought into the analytical framework.

Expressing and enforcing policy remains an art rather than a science—how can we systematize our expressions of security intent? As a foundation for expressing policy, the structure and ontology of languages for policy expression and enforcement remains an underdeveloped field. The measurement and assessment of performance through useful metrics remains a perennial challenge, as does modeling of complex or modeling of non-physical, human systems.

3.1 Redefining the Problem and Approach

Responses to the insider threat have historically been built on a notion of organizations with fixed perimeters. This seminar explored the different ways in which organizational perimeters are now subject to change and the ways in which organizational processes have to adapt to respond to these changes. Two related questions were a common thread throughout many discussions and presentations. First, what threat are we really concerned about? Second, are monitoring and detection the best approaches for addressing these threats, or should we rethink in some fundamental way our approaches?

In the 2008 seminar we concluded that an insider is defined as a person that has been legitimately empowered with the right to access, represent or decide about one or more aspects of the organization's structure [3]. This definition is useful: as one presenter noted, it captures the fact that an organization does have a boundary defined by who does and does not have legitimate rights. An insider can be monitored by the organization; the organizational perimeter defines what can and cannot be monitored. The term has both operational and legal meaning.

Nonetheless, there was a strong sense of many that the security threat of interest was something more. Many successful 'outsider' attacks use insider connections. Because of data leakage or porous organizational boundaries, in some instances outsiders have knowledge equivalent to insiders. Some examples were cited (e.g., the Christstollen stolen and replaced with credit card data from the Landesbank Berlin in 2008), illustrating that it is only after a successful attack that it becomes clear that seeming outsiders have quite considerable insider privileges. Several key concepts emerged during this discussion:

- Organisational perimeters are fluid and can be defined in many ways. The physical perimeter is no longer static with technology developments increasing the likelihood of the construction of ad hoc work contexts (such as mobile working), temporary partnerships between organizations and the dispersal of data across many geographical locations. The hierarchical perimeters within an organisation are also more fluid than previously. The fluidity of physical perimeters contributes to the often ephemeral nature of organizational hierarchies. Technology can also enable greater diversity of organizational purpose and the increased likelihood of product diversity. As a result hierarchical structures are often less static and more likely to undergo changes as the organizational purpose and outputs re-form.
- It was noted that 'the [inter-personal] relationship was everything' and that information flows and data sharing, including that which would officially be considered part of policy-granted legitimate access, should be considered in light of how relationships within and across organizational boundaries worked.
- Trust has different meanings in security and social sciences. In security, trust means dependability and assurance. In the social sciences, trust connotes the 'willingness to be vulnerable based on positive expectations or the actions of others.'
- Relationships change, and thus threats are 'emergent.'

To summarize: individuals gain knowledge about an organization and its systems through relationships with staff and others. These relationships do not map neatly into formal structures and policies, and thus knowledge about the organization that can be used to launch attacks equivalent to insider attacks is not limited to those who are insiders. Furthermore, relationships are dynamic and threats emerge (are emergent) as relationships change.

We subsequently have concluded that a special and important class of security threats is from 'informed actors' (our term) operating within or at the perimeter of dynamic organizations. The class of informed actors includes insiders but also captures those whose special knowledge of the organization and its systems derives from changing relationships, aided by the increasing porosity and dynamism of the organization's boundaries. A software engineer at a commercial vendor who worked on the systems subsequently incorporated into the organization's systems would be an informed actor; he would not be an insider.

A second intertwined discussion concerned the goal of effective policy. If interpersonal relationships drive information flows in ways not captured by formal

organizational policies and boundaries what should policy focus on? One key conclusion is that there are real limitations on the effectiveness of monitoring and prediction. Most people appear willing to cheat if they believe that they can do so without consequence; many apparently are willing to act on-line in ways less ethical than they would in person. 'Traitors' are an important class of insider threats. For all of these types of behaviors, monitoring to detect violation of privileges may not be effective, because motivation rather than violation of privileges may be the discriminate between threat or benign. If changes in relationships are an important basis for security threats of interest, detecting these changes through monitoring remains imperfect at best. In fact, we concluded that there are some fundamental limits on how effective any security policies can be. Security mechanisms fail. As one participant noted how do we even detect when a security mechanism has failed? Predicting threats is also highly imperfect. Insider threats appear to be rare events, and statistical methodologies do not work well with rare events. In any event, the general lack of good data also confounds our ability to create solid and operationally useful predictive models.

The challenge therefore is to find policies that more effectively help organizations recognize when interpersonal relationships change in ways that may pose security threats. As discussed in Section 3.3, there is evidence that some proactive approaches may have benefits over detect and respond controls. The sustainability or long-term viability of an organization (particularly in the business world) is linked to its security. This is a very important conclusion. Some organization processes are more 'sustainable' than others. That is, they are more resilient with respect to informed threats and more capable of limiting the damage from informed threat attacks. Resiliency appears to stem from usable, effective, and efficient security having been built into the organization processes. In order to make organizational processes sustainable in this organizational scenario where perimeters are ad hoc and volatile, their design must be grounded in an understanding of the organizational culture. In order to continue to be of benefit, organizational processes must be capable of detecting and adapting to organizational change.

3.2 Organizations, Staff, and Organizational Processes

A second main discussion thread in this area was how to make sense of the available data, and how much data actually is and may be available. Challenges beyond making sense of data include how to deal with it effectively, how to fit it into a model, and how to understand and validate that model. What is normal for one individual, based on history, maybe helpful in detecting deviation, but might also just cause false positives.

The fundamental challenge involves integrating behavioral (non-cyber) data and cyber data into a reasoning system capturing knowledge representation and synthesis. It must be noted that these systems have to comply with legal restrictions, for big organizations even many, probably contradicting ones. While the system's knowledge (both cyber and psychosocial) can be modeled using a formal ontology language that captures expert/domain knowledge in a computational

accessible form, representation of social science knowledge is more difficult; such representations are informed by inferences about WHY the exploit occurs—i.e., the motivation. By integrating both into the same system, we may have a better chance to logically synthesize them; for example, if we sense highly suspicious psychosocial indicators, we may better know where to look for cyber indicators.

Behavioral (psychosocial) data tend to be sparse and asynchronous compared to cyber data, which is characterized as massive, streaming, and real-time, particularly when used to predict events. A motivating example here might be the Fort Hood shooter in 2009. If we are examining his perusal of terrorist web sites, we might initially dismiss this behavior as reasonable research activities related to his caseload. If we subsequently find that he had more serious terrorist ties, then this new information should trigger a re-evaluation of the data (in light of new information). The architecture of the reasoning system must support the re-evaluation of data-streams given triggering events.

Another important discussion centered on the integration of human behavior into security analysis. Traditionally, threat identification in information security focuses on computers and networks. Two main approaches can be distinguished. First, formal methods can be used to analyze the exchange of messages in a security protocol, thereby determining whether an attacker can manipulate the protocol in such a way that the desired outcomes do not hold. This can be called the interaction-oriented approach. Second, model checking and static analysis can be used to analyze which sequences of actions in a network an attacker (insider) can exploit to achieve a certain goal. Physical infrastructure can be included in such models as well. This can be called the access-oriented approach.

In both approaches, humans have not been included explicitly as actors in the protocols or in the system models. There are several reasons to investigate options in this direction. As explained by Peter Ryan a malicious device could fool humans into executing a slightly different protocol than prescribed, thereby invalidating its security properties. Also besides exploiting vulnerabilities in a network, an attacker can use social engineering to get closer to her goal (borrowing a credential, asking for a password, asking someone to open a door).

In order to include humans in these models we need to decide which human properties are to be included. These properties will normally be probabilistic (most of the people most of the time). Social science results can be used to determine these probabilities. Social science can also be used to develop postulations as to why some patterns of behaviour are more likely than others. This is particularly important when developing a sophisticated and adaptable regulatory and cultural response to the insider threat. It then becomes possible to analyse possible attacks according to the probability of every necessary action being successful. Heuristics are needed to keep the analysis scalable for larger systems / networks.

For enhancing our understanding of insider attacks, it may also be worthwhile to consider results from criminology. Criminology and crime science focus on the offender and his motive and on the environment, to reduce opportunity. They may thus be instrumental in developing tools that can select effective counter

measures and aim at reducing opportunities in business processes. It remains unclear though what the environment in an insider scenario is. The practical question is how to make the five principles of Situational Crime Prevention (Increase effort, Increase risk, Reduce reward, Reduce provocation, Remove excuses) work to prevent insider related offences. For example to make the principle "increase the risk" work through deterrence, the problem is that offenders generally do not know what the punishment for crimes is going to be—and they do not think they will get caught either [5]. So neither severity of sanctions nor certainty of sanctions is guaranteed to have a large deterrent effect. It could be argued that this is different for insiders, since they should know about sanctions, but this must be investigated.

We need to consider usability and acceptability of any tools and procedures we may devise to counter with insider threats . There are several ways in which bad usability and low acceptability can reduce effectiveness of tools and procedures. Regarding the adoption of tools and procedures we note that tools and procedures that are difficult or cumbersome to use will not be widely adopted (previous examples from security such as PKI provide lots of examples), and even if they are adopted (in the sense of being bought and deployed by some organizations), they may not be effective because they are found to difficult to use, or are not effective since users make mistakes. People can, of course, be forced to work with unusable security tools/procedures, through monitoring/compliance functions. But forcing people in this way reduces personal and organisational productivity—meaning it consumes resources beyond the cost of monitoring/compliance—and can lead to disgruntlement with the organisation—given that disgruntlement is a starting point for insider attacks [6], this can be extremely counterproductive. Effective security tools and procedures must therefore be usable and make it easy for users to perform well, and "do the right thing" by security.

The problem is that epistemological differences between social and computing disciplines hinder collaboration by both parties. As a generalization computer scientists aim on solutions, produce new things (processes, methods, algorithms, products), and like to generalise and experiment. Social scientists, on the other hand, like to explore things (theories, concepts, techniques) and describe problems on a long-term scale., and are likely to produce social and economic theories that might be used to develop interventions. Each field often finds it difficult to appreciate the value of the research approaches and contributions of the other fields. Computer science researchers are able to design better technologies if they are given clearer social and economic theories about insider behaviour. There is a tendency to focus strictly on the technical aspects of defending against insider attacks. While this is an important aspect, understanding the nature of the problem requires social science approaches of insider threat case study investigations. We need to understand differences between insider and outsider behaviors (as well as different classes of insider attacks, such as insider theft, insider fraud, and insider IT sabotage). Evidence indicates that the patterns of these types of crimes vary substantially [7]. Understanding their differences before understand-

ing their commonality is critical to developing effective countermeasures to the full range of potential compromise.

3.3 Policies and language specification and ontologies

Evidence indicates that certain proactive approaches to security management (prevention controls, screening, establishing a trust environment) may have benefits over reactive approaches (detect and respond controls, punitive approaches such as sanctions).

In general proactive approaches are needed to complement reactive approaches (generally help in short term, but lower performance in long term) for effective security in both short and long terms. Theory suggests that attempts to improve the proactivity of an organization fails not because of any inherent deficiency in the techniques themselves, but because of how the introduction of the more proactive techniques interacts with the physical, economic, social and psychological structures in which implementation takes place.

What was identified as a need in the area of policy languages are methods for automatically generating policy rules from a data set. This would also enable systems to trace contradictory policies back to the underlying specification. Especially in merging or extending organizations, where policies from different sources are merged and need to be unified, a possibility of disambiguating or at least flagging contradicting policies is in high demand. The same holds for sanity checks of policies such as consistency or completeness.

There was an ongoing debate about properties of an "ideal" policy language, for example whether it should include first or second order logic, or whether it should be enforceable.

Ideally, one also would like to express the likelihood of a certain threat to occur and the potential impact of a certain threat, possibly with input of socio-technical tools as described above. The problem in designing this kind of language clearly is finding the sweet spot between usefulness and expressiveness.

However, often the best help against an attack is information, not a policy language or a policy. Even worse, it is unclear how policies beyond a social level can prohibit insider attacks. The working group concluded that a policy cannot prevent attacks of insiders by definition, but it can be layered over post mortem analysis data. Policies should be platform-independent—ontologies should make them platform-specific.

From an organisation's point of view, detection and deterrence must be auditable and traceable to evaluate policies effectiveness. Especially in case of emergency, the only possible action is to audit. Critical infrastructure is an example where we only can audit and try to enforce physical countermeasures, like two keys in nuclear submarines. In other cases we may be able to specify the threat, but may not be able to prevent it, e.g., in the power grid.

Simply measuring the number of incidents and cost of losses is not sufficient to capture the costs and benefits of a particular program of controls and procedures. The concepts of least privilege and separation of duties are excellent guides in discussing insider protection properties that can be measured.

Categories of measures include threat level, vulnerability (perhaps excess privilege for given user roles/functions), cost of expected consequence of incident type or recovery to specified incident type, cost of attribution (speed of attribution, size of set of potentially responsible individuals), error-proneness (likelihood of security-critical errors), costs of the program, including impact on function, internal compliance: policy, ratio of constrained data items (CDI) to trusted parties, employee job satisfaction and attitudes toward the company, number of individuals required to collude in order to achieve a specified security breach (detectably / indetectably), or number of insider incidents reported. It was noted that in other areas, an increase in the number of reported (say safety) incidents resulted in a reduction in the overall impact of loss due to (safety) incidents.

Measures of awareness should be included—although they are indirect measures, they should be considered as sine qua non for a sound insider strategy. Shaw *et al.* [8] describe three levels of user awareness: 1) perception (the user is able to detect threats in the environment), (2) understanding (the user is able to combine information from different sensors, interpret them and use the resulting knowledge to reduce risk in the environment), and (3) prediction (the user is able to predict future attacks and proactively change own behaviour to reduce or remove risk).

3.4 Intangible Systems

Most of the systems considered in insider threat research are "real" systems. One thread of the discussion of the seminar tried to identify threats against "unreal" or intangible systems. Two groups of systems were identified, non-developed, not yet tangible systems (*i.e.* from concept, through design, to the point of implementation), and functionally non-tangible systems that extend outside the boundary of an organization. During discussions the focus was further centered on the first category, represented by systems such as supply chain attacks, specifically looked at ATMs, voting machines.

Insider attacks on intangible systems can be characterized by the fact that vulnerabilities are put into the system for exploits at a later time. This makes attribution even harder than in "real" systems, because it may be impossible to trace the source of an attack back to its origin, or to detect before realizing the system. If applicable, one therefore should have verifiable system specifications, and throughout the development and deployment process apply testing, validation, and auditing. The development process can be supported by application business policies such as separation of duties; if someone is designing a new ATM machine, the components should be designed by more than one individual, or possibly by more than one company. Of course, introducing 3rd party designers necessitates a need for more rigor in evaluations. At the same time, a mitigation of these threats seems almost impossible, so focus should be on disaster recovery.

Future development of such a topic could include aspects of the system that are socially constructed as a result of organizational culture, the belief systems of the people engaged with the system, the experiences of the system users

and system designers etc. This line of discussion would accept the existence of multiple realities present within a tangible system.

4 Conclusions

The Dagstuhl seminar on strategies for prevention, mitigation, and response with respect to insider threats explored all these areas through discussions and presentations based on input from different and divert communities.

The purpose of the 2010 seminar was to make progress towards an integrated framework for selecting among and evaluating the impact of alternative security policies against insider threats. An integrated framework, we recognized, needs to include issues not considered in insider work before, such as the economics of insider threats [4], and the role of law as both a preventative and punitive instrument. We saw the need for creating and testing alternative integrated frameworks so that practitioners and researchers could make informed choices as to combinations of actions targeted at insider threats, and also the need for methods to evaluate the effectiveness of these actions.

The goal of the seminar was to develop a taxonomy for identifying insider threats and an integrated approach that allows a qualitative reasoning about the threat and the possibilities of attacks. We expected this to result allow us to develop a deeper understanding of security policies and how to evaluate them.

During the seminar, all these issues were inspected and scrutinized, resulting in a better appreciation of social and organizational factors relevant to insider threats, and addressing important questions in related areas.

The main conclusions of the seminar are

1. that we need to further investigate the impact of changes in organizational processes on the effectiveness of internal controls. The nature of 'trust' and the changeability of human behavior make prediction difficult, and create unique security problems in preventing, mitigating, and responding to insider threats.
2. There are real limitations on the effectiveness of monitoring and prediction. Monitoring to detect violation of policies and privileges may not be effective, because motivation rather than violation of privileges may be the discriminate between threat or benign. If changes in relationships are an important basis for security threats of interest, detecting these changes through monitoring remains imperfect at best. In fact, we concluded that there are some fundamental limits on how effective any security policies can be.
3. that a special and important class of security threats is from 'informed actors' operating within or at the perimeter of dynamic organizations. The class of informed actors includes insiders but also captures those whose special knowledge of the organization and its systems derives from changing relationships, aided by the increasing porosity and dynamism of the organization's boundaries.
4. The sustainability or long-term viability of an organization is linked to its security, and some organization processes are more 'sustainable' than others.

That is, they are more resilient with respect to informed threats and more capable of limiting the damage from informed threat attacks. Resiliency appears to stem from usable, effective, and efficient security having been built into the organization processes. In order to make organizational processes sustainable in this organizational scenario where perimeters are ad hoc and volatile, their design must be grounded in an understanding of the organizational culture. In order to continue to be of benefit, organizational processes must be capable of detecting and adapting to organizational change.

To summarize we believe that a general research area in the direction of how to develop sustainable, resilient processes in organizations. Based on the seminar's results we expect these processes to be vital in mitigating informed threats, which we see as a superset of insider threats, covering both attacks by "real" insiders and informed outsiders with some access to an organization's assets.

We would like to thank all participants of the seminar for making it a fruitful and inspiring event—and especially Dagstuhl's wonderful staff, for their endless efforts, both before and during the seminar, to make the stay in Dagstuhl as successful as it has been.

References

1. Homepage of Dagstuhl Seminar 08302: "Countering Insider Threats". Available from http://www.dagstuhl.de/08302, last visited December 13, 2010 (2008)
2. Homepage of Dagstuhl Seminar 10341: "Insider Threats: Strategies for Prevention, Mitigation, and Response". Available from http://www.dagstuhl.de/10341, last visited December 13, 2010 (2010)
3. Probst, C.W., Hunker, J., Bishop, M., Gollmann, D.: Countering insider threats. Dagstuhl Seminar Proceedings (2008)
4. Probst, C.W., Hunker, J.: The risk of risk analysis and its relation to the economics of insider threats. In Moore, T., Pym, D., Ioannidis, C., eds.: Economics of Information Security and Privacy. Springer (2010) 279–299
5. Karstedt, S., Farrall, S.: Law-abiding majority? the everyday crimes of the middle classes. Technical report, Centre for Crime and Justice Studies, Kings College London, UK (2007)
6. Cappelli, D.M., Desai, A.G., Moore, A.P., Shimeall, T.J., Weaver, E.A., Willke, B.J.: Management and education of the risk of insider threat (merit): System dynamics modeling of computer system sabotage. In: Proceedings of the 24th International Conference of the System Dynamics Society. (2006)
7. Cappelli, D.M., Moore, A.P., Trzeciak, R.F., Shimeall, T.J.: Common sense guide to prevention and detection of insider threat (3rd ed.). Technical report, CERT Program, Software Engineering Institute, and CyLab of Carnegie Mellon (2008)
8. Shaw, R.S., Chen, C.C., Harris, A.L., Huang, H.J.: The impact of information richness on information security awareness training effectiveness. Computers & Education **52** (2009) 92–100