10341 Abstracts Collection
Insider Threats: Strategies for Prevention, Mitigation, and Response
— Dagstuhl Seminar —

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. From August 22 to 26, 2010, the Dagstuhl Seminar 10341 "Insider Threats: Strategies for Prevention, Mitigation, and Response" was held in Schloss Dagstuhl – Leibniz Center for Informatics. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Insider Threat, Security Policies, Threat Modelling

10341 Report – Insider Threats: Strategies for Prevention, Mitigation, and Response

The Dagstuhl seminar "Insider Threats: Strategies for Prevention, Mitigation and Response" was held on August 22 – 26, 2010 (Seminar #10341) to advance our understanding of ways of reducing insider threats.

The 2010 seminar built on the results of its predecessor from 2008 (Countering Insider Threats, #08302, Seminar Homepage, Seminar Report). In this seminar we developed a shared, inter-disciplinary definition of the insider 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.

Dagstuhl Seminar Proceedings 10341
Insider Threats: Strategies for Prevention, Mitigation, and Response
http://drops.dagstuhl.de/opus/volltexte/2010/2904

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, 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.

Joint work of: Bishop, Matt; Coles-Kemp, Lizzie; Gollmann, Dieter; Hunker, Jeffrey; Probst, Christian W.

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2010/2903

The Policy Gap and the Influence of Organisational Learning

Lizzie Coles-Kemp (RHUL - London, GB)

This presentation considers the types of gap that occur between security policy and the practice of information handling, management, and processing in the workflow. It explores how the lens of systemic organisational learning can help to explain the dimensions of the gap. In doing so, both the process of organisational learning and the nature of faulty learning are considered. A case study is used to evaluate how far organisational learning might influence the dimensions of policy gap and possible responses to reduce the gap.

Joint work of: Lizzie Coles-Kemp, Jose-Rodrigo Cordoba-Pachon

Analyzing multi-domain insider threats

Trajce Dimkov (University of Twente, NL)

The security goals of an organization are realized through security policies, which concern physical security, digital security and security awareness. An insider is aware of these security policies, and might be able to thwart the security goals by combining physical, digital and social means. A systematic analysis of such attacks requires the whole environment where the insider operates to be formally represented. This paper presents Portunes, a framework which integrates all three security domains in a single environment. Portunes consists of a high-level abstraction model focusing on the relations between the three security domains and a lower abstraction level language able to represent the model and describe attacks which span the three security domains.

Using the Portunes framework, we are able to represent a whole new family of attacks where the insider is not assumed to use purely digital actions to achieve a malicious goal.

Keywords: Insider threat, physical security, digital security, social engineering

Testing Insider Threat Detection Systems

Carry Gates (CA Labs -New York, US)

Algorithms and systems for detecting insider threats—particularly those performed by traitors as opposed to masqueraders—have special requirements when it comes to proper experimental design for testing their accuracy. This is particularly true given the lack of data from actual attacks. Due to the practical nature of detection algorithms, we define insiders with respect to the resources or information they can access. We further scope the problem by defining a perimeter as being those systems that you can monitor. Given this, experiments using real data can be performed in one of three ways: (1) through a red team exercise, where specific insiders are given example scenarios but not detailed instruction, (2) through performing malicious activities in a sandbox and inserting this data into the real data stream, and (3) through installing the detection system and waiting for real attacks.

Keywords: Insider threat detection

Conceptualising social engineering attacks through system archetypes

Jose J. Gonzalez (University of Agder - Grimstad, NO)

At the highest abstraction level, an attempt by a social engineer to exploit a victim organisation either attempts to achieve some specific target (denial of service, steal an asset, tap some particular information) or it wishes to maximise an outcome, such as to disable the organisation by a terrorist attack or establish a permanent parasitic relationship (long-term espionage). Seen as dynamic processes, the first kind of exploit is a controlling (“balancing”) feedback loop, while the second kind is a reinforcing feedback loop. Each type of exploit meets a first line of defence in control processes or in escalating (“reinforcing”) processes of resistance.

The possible combinations of the two modes of attack and the two modes of defence yield four archetypes of exploit and natural defence.

Predictably, the social engineer would seek to outsmart the first line of defence; it is shown that each archetype implies a particular strategy to do so.

Anticipation of these modes of attack must be the starting point for an effective multilayered defence against social engineering attacks.

Keywords: Social engineering; Insider threat; System archetypes

Joint work of: Jose M. Sarriegi, Jose J. Gonzalez

Full Paper:
http://www.inderscience.com/storage/f109451131286712.pdf

Research on Employee Behavioral Monitoring for Insider Threat: Tyrannical Oppression or Responsible Science?

Frank L. Greitzer (Pacific Northwest National Lab., US)

There is ample evidence that for many types of insider threat, observable "concerning" behaviors are indicated and recognized well before the insider crime. But because of lack of processes or knowledge or mechanisms to deal with such indicators, organizations are ill-equipped or unable to take proactive steps to alleviate problems or prevent abuses. Thus, there is a need for research on identification of and responsible mitigation strategies based on observable behavioral precursors of insider crimes. In our research, we have encountered obstacles to publishing scientific works on research with this focus in part because of two apparently prevailing views: (1) management and human resources personnel are "clueless" about relevant behavioral "precursors" or problems in the workplace; and (2) the research should not be performed because of the potential misuse and ethical pitfalls in application of the technology.

The presentation concerns the proposition that developing a model based on psychosocial "markers" or "precursors" (indicators) is a legitimate research endeavor, particularly given careful attention to and accommodation for privacy and ethical constraints. Questions for discussion concern whether or not the workshop participants consider this research thrust to be appropriate and viable, how to address privacy and ethical concerns, and what additional considerations might be needed to help advance the development, testing and refinement of predictive models that integrate behavioral and IT/cyber data.

Keywords: Employee monitoring, insider threat, privacy

About the role of the Insider in Crime Science

Pieter H. Hartel (University of Twente, NL)

Crime Science prevents crime by changing the opportunity structure of crime. The conceptual framework of Crime Science includes the Routine Activity Approach, Crime Pattern Theory (CPT) and the Rational Choice Perspective. Of these three, only the second theory provides a handle for reasoning about insiders through the notion of a local offender, who commits crimes in his local neighbourhood. The notion of an insider is probably too general to be of much use in the crime specific approach of Crime Science.

Keywords: Criminology

Joint work of: Pieter Hartel, Marianne Junger, Roel Wieringa

Creeping Failure – the overall failing of G-8 cyber security policy

Jeffrey Hunker (Jeffrey Hunker Associates LLC, Pittsburgh, US)

The Internet is really a city, a vast tuple of streets, highways, and addresses. It is a city like London in the early 1800's—rapidly growing, based on new technologies—but lacking the technical, social, and political infrastructure to deal with the new problems that have emerged, problems of our own creation.

These problems are foundational to our inability to deal with insider threats—namely, the lack of incentives and infrastructure to address security on networks. Addressing these problems requires not just technological but economic and political action:

- reforms in regulations (command requirements),
- reforms in incentives for software developers (where should liability be put),
- a new social contract among users (we all wash our hands, why not the same social contract for Internet use),
- new frameworks (like for dealing with fire risks, where regulation, insurance, response, and research work together), and finally
- we should migrate to a more secure "new Internet".

Keywords: Creeping failure, policy failure, insider threats

Full Paper:
http://www.mcclelland.com/catalog/display.pperl?isbn=9780771040245

Striking back at the Privileged: Application-based Monitoring for Detection, Policy, Specification, Forensics, and Automated Response

Karl Levitt (University of California, Davis, US), Lenore Zuck (NSF - Arlington, US), Sean Peisert (University of California, Davis, US)

"Insiders" often have privileges that enable them to successfully launch their attacks; prevention is likely impossible. In principle, a system should have a security policy in place to preclude such activity, but for complex systems, it is difficult to formulate such a policy to render it enforceable or to envision all possible undesirable scenarios. Since perfect attack prevention is impossible, it is at least necessary to detect insecure activity after the fact, and possibly to enable actions to mitigate damage and make the attacker accountable.

We strongly believe that monitoring application-level data—rather than host and network data alone—can lead to successful detection. Further, we argue that machine learning can aid in identifying "good behaviour". We suggest the use of anomaly detection as the engine for detection, relying in part on relational database systems and their structure as partial descriptors for feature selection. Combined with specifications, monitors can provide a powerful tool to detect attacks by privileged users.

Joint work of: Karl Levitt, Lenore Zuck, Sean Peisert

LUARM and ITPSL - Sensing and specifying insider threats

George Magklaras (University of Plymouth, GB)

Specifying insider threats is a non trivial task. There is absence of realistic data sets about how insider threat signs are manifested at system level and lack of general semantics to describe user actions in IT systems. LUARM and ITPSL are two tools/methodologies that target these issues. LUARM is a bespoke logging engine that logs user actions at network, file and process execution level. ITPSL is a Domain Specific Language XML based construct to make sense of the logged data. Together, these two proposed tools should enable researchers to enrich the systemic study of insider misuse and replay attacks at system-level to derive conclusions. The intersection of insider misuse specification and detection with data forensics is also examined.

Keywords: Insider threat specification, insider threat detection, Insider threat prediction, computer forensics, data logging, action logging, luarm, itpsl

Fraud prevention as a preventive strategy

Jörg Meyer (KPMG - Köln, DE)

Analyzing data in search for anomalies and misbehaviour can be a part of an organization's control system to detect non-compliant activities. With a proper communication strategy this serves as a preventive control as well and is usually more efficient then fine-graining process controls. Additionally it is well known that detective controls allow for more flexibility in the business process then preventive controls. However, data analysis needs to be performed very carefully and in full compliance with any applicable data protection or privacy regulations.

The presentation shows different approaches including K-Trace, Continuous Monitoring, individual analytics and Digital Evidence Recovery.

Keywords: KPMG Forensic Technology Data Analytics

Summary of workshop on accelerated learning to mitigate insider threat

Andrew P. Moore (Carnegie Mellon University - Pittsburgh, US)

I will present the goals, structure, and preliminary findings from a workshop on accelerated learning to mitigate insider threat held in June of 2010. I will present some of the knowledge and research gaps identified during the workshop in both the instructional design domain and the insider threat domain.

Keywords: Accelerated learning, insider threat, instructional design, cyber security, enterprise security

Bringing the Customer into Audit

Steven Murdoch (University of Cambridge, GB)

In this talk, I present three examples of disputed card transactions (both ATM and point-of-sale). In each of them, the customer was motivated to discover how these transactions had occurred; in particular, whether the correct card was used and whether the correct PIN was entered. However, they had not been given sufficient information to establish what had happened and whether it was fair for them to be held liable. In cases like this, customers could be empowered to act as auditors, mitigating the known weaknesses of existing audit processes. I propose three different ways in which the customer-auditor could be facilitated. First, standard procedure should be for customers to retain cards for which there is a dispute (not to destroy them). Second, the log of events which occur on an account should be processed to form a hash-chain, and this should be provided on customer statements. Finally, receipts for transactions should include enough information to allow the corresponding transaction MAC to be verified.

Full Paper:
http://prezi.com/1t-vhqdeeonk/bringing-the-customer-into-audit/

Insider Threats Considered Holistically

Peter G. Neumann (SRI - Menlo Park, US)

In this talk I take a total system view of how to cope with insider threats. They summarize my 2008 position paper, recent developments relating to insider misuse in elections (including some convictions for election fraud perpetrated by insiders), and some potentially useful directions for the future:

- revisit layered and distributed trustworthy system architectures,
- reconsider multi-level security and multi-level integrity,

- exploit novel uses of cryptography, such as homomorphic encryption and functional encryption, and
- consider risks of insiders in the clouds.

Keywords: Insider threats holistic trustworthiness encryption

See also: 1. P.G. Neumann, Combatting Insider Misuse, with Relevance to Integrity and Accountability in Elections and Other Applications, Dagstuhl Workshop on Insider Threats, Schloss Dagstuhl, Germany, July 2008. 2. P.G. Neumann, Combatting Insider Threats, Chapter 2 in Insider Threats in Cybersecurity – and Beyond, C.W. Probst and J. Hunker and D. Gollman and M. Bishop (editors), Springer Verlag, 2010. 3. P.G. Neumann, Computer-Related Risks, Addison Wesley, 1995. 4. Craig Gentry, A Fully Homomorphic Encryption, PhD Thesis, Stanford University, 2009, and Fully Homomorphic Encryption Using Ideal Lattices, STOC, 2009. http://crypto.stanford.edu/craig/ 5. B. Waters, S. Hohenberger, A. Lewko, A. Sahai, et al., Functional Encryption.

Can you attack a system if it is not real?

Richard Overill (King's College - London, GB)

This talk extends a positioning article from 2008 and analyses the nature of the insider threat to an ISMS, evaluates how technologies typically used to model adversarial relationships might be used to assess the robustness of an ISMS to withstand an insider attack and concludes with a discussion of the practical implications of these insights for the management of insider risks.

Joint work of: Richard Overill, Lizzie Coles-Kemp

The insider threat – a practical view

Sachar Paulus (FH Brandenburg an der Havel, DE)

Based on the experience of running a corporate security department at a large software manufacturer, we summarize insights on dealing with the insider threat.

In a business context, insiders are clearly defined by "people having access to information that may have (share) market impact." Consequently, managing this group of people is a major task for a corporate security, compliance, and risk management function.

There are a number of problems with increasing cost of implementation:

- identifying insider,
- assign the appropriate authorizations,
- detecting anomalies, and

- consequence management.

Security policies are used for governing the behaviour of insiders, working with critical assets of the organization. But they have similar implementation issues, among others:

- there are too many policies,
- there is an important trust component that needs to be taken into account when enforcing policies, and
- minimizing risks is not in the focus of a business organization, control measures need to pay off.

Key success factors were finally presented that led to a successful, world-wide rollout and implementation of a security management framework.

Keywords: Experience report, insider threat

Catching spies by math

Dusko Pavlovic (University of Oxford, GB)

Spies have been a problem for kings and governments for a very long time. With the spread of networks, spies are becoming a technical problem for computer scientists. Networks are normally protected from malicious users through the processes of trust building and authentication. But some spies manage to build up trust, get authenticated, and penetrate the system. What defenses remain when the normal trust and authentication processes fail? I consider this question within the formalism of network theory, which lies at the intersection of computer science and social sciences. Spies can be viewed as a form of "dark matter" in a network. I shall describe some mathematical methods for mining dark matter, and discuss their applicability to the task of recognizing malicious insiders.

Keywords: Trust, treason, adverse selection, SVD, privacy

Optimistic Access Control and Anticipatory Forensic Analysis of Insider Threats

Sean Peisert (University of California, Davis, US)

Insiders are inside our network and must be allowed to do their jobs, by definition. Also by definitions, those jobs may entail taking potentially harmful actions, which must be allowed. Our best defense is usually analyzing what they have been doing in the past. We propose an approach for making analysis more efficient through behavioral analysis of users and "optimistically" allowing (but constraining) a user's future actions when we see that the most likely paths based on execution history are not being followed.

Keywords: Insider threat, forensic analysis, forensics, optimism, access control, intrusion detection, epsilon machines, machine learning

Joint work of: Sean Peisert, Matt Bishop, Christian W. Probst

Actor-Network Security: Why Humans Don't Count in Threat Analysis

Wolter Pieters (University of Twente, NL)

In determining vulnerabilities of organisations against insider threats, a model for threat analysis is needed that

- avoids relying on containment of assets within perimeters, and
- includes humans as entities mediating access.

I propose to use actor-network theory as a framework for such analysis. This approach focuses on the behaviour of networks of actors from the perspective of access, and treats humans and non-humans symmetrically. In this sense, a human opening a door by means of a key is isomorphic to a dongle using a human to get into a computer. The model is formalised in terms of hypergraphs. Policies express which entities allow which other entities to access what. Based on the policies, attack can be generated that represent the possible insider attacks. Future work includes tool support, case studies, dynamic policies, and probabilistic policies. Also, I hope to understand which human properties do need to be included explicitly to make the model sufficiently expressive.

Keywords: Actor-network theory, containment, hypergraphs, security modelling, threat analysis

Full Paper:
http://eprints.eemcs.utwente.nl/17809/

How many Insiders does it take to change an Election?

Peter Y. A. Ryan (University of Luxembourg, LU)

In this talk I use the example of secure voting systems to illustrate a number of issues and challenges in designing such systems. The challenges are highly socio-technical in nature, as illustrated by various known attacks that subtly exploit human factors, insider threats, etc., including chain voting, "Italian" attacks, and social engineering (e.g. turning cut and choose into choose and cut).

I argue that the insider/outsider distinction is far to coarse, and its utility is unclear.

I discuss the approach of "end-to-end" verifiability as a way to address the challenge of dealing with such socio-technical vulnerabilities. I also speculate as to possible approaches to analysing systems in a way that takes account of human factors as well as the purely technical aspects.

Keywords: Secure voting systems, socio-technical threats

When you should trust your employees - and when you shouldn't: A systematic approach to applying knowledge of risk and incentive structures to designing security

M. Angela Sasse (University College London, GB)

Social science researchers agree that the question of trust arises in situations of risk and uncertainty. It has been shown to have economic benefits, arising from two sources

- trusting rather assuring (which is what security mechanisms do) saves resources required for monitoring and enforcement, and
- trust between people builds social capital in organisations and societies.

Organisations thus face a dilemma—not trusting its employees will be expensive, and reduce goodwill and creativity; trusting too much will make the organisation vulnerable to insider attacks. To design effective safeguards against insider attacks, an organisation has to identify its specific risks and vulnerabilities – which can be part of the IT infrastructure, or part of the business workflow – and consider the motivations and incentive structures of different types of attackers, namely professional attackers, opportunistic attackers, and emotional attackers.

This provides the basis for a taxonomy of vulnerabilities, attackers and incentive structures. Different technical and non-technical security mechanisms are appropriate for different risk and incentive structures, and can be mapped accordingly.

Keywords: Trust, insider

Insider Threats to Voting Systems

Alec Yasinsac (University of South Alabama, US)

Insider attacks are particularly insidious threats to electoral integrity. Traitors that misuse the trust that is placed in them often have system access that facilitates malicious acts themselves and their subsequent cover-up efforts.

In this paper, we define what it means to be an insider and we identify several classes of elections insiders. We also categorize the threats that each insider class has relative to the electoral functions.

Beyond specifying well-known elections insiders such as poll workers and local elections officials, we address several insider categories that are rarely, or never, mentioned in considering election insider threats. For example, we have not previously seen members of the judiciary identified as prospective elections insiders and we give a concrete example of how judges can accomplish insider

attacks on elections. Similarly, we identify the impact that policy makers can have on the electoral process and show how malicious legislators may be able to influence a broad spectrum of elections through the laws that they propose and promote.

Insider attacks are real and imminent threats to electoral integrity. By identifying insiders and categorizing the threats that they pose allows us to create policies and procedures that better ensure sound elections and to ensure the integrity of our way of government at local, state, and federal levels.

Keywords: Election Threats, Voting System Security, Risk Assessment, Secure Software