

08302 Abstracts Collection
Countering Insider Threats
— Dagstuhl Seminar —

Matt Bishop¹, Dieter Gollmann², Jeffrey Hunker³ and Christian W. Probst⁴

¹ Univ. of California - Davis, USA
bishop@cs.ucdavis.edu

² TU Hamburg-Harburg, D
diego@tuhh.de

³ CMU - Pittsburgh, USA
jhunker@andrew.cmu.edu

⁴ Technical University of Denmark, DK
probst@imm.dtu.dk

Abstract. From July 20 to July 25, 2008, the Dagstuhl Seminar 08302 “Countering Insider Threats ” 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

08302 Summary – Countering Insider Threats

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

This article summarizes the objectives and structure of a seminar with the same title, held from July 20th to July 25th, 2008, at Schloss Dagstuhl, Germany. The seminar brought together researchers and policy-makers from all involved communities, to clarify what it is that identifies an insider threat, and to develop a common vision of how an insider can be categorized as well as an integrated approach that allows a qualitative reasoning about the threat and the possibilities of attacks. This report gives an overview of the discussions and presentations during the week, as well as the outcome of these discussions.

Keywords: Insider threat, executive summary

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2008/1793>

Combatting Insider Misuse, with Relevance to Integrity and Accountability in Elections and Other Applications

Peter G. Neumann (SRI International, US)

Various risks of insider misuse arise at different layers of abstraction. This observation leads to a perspective on insiders that is both hierarchical and context-dependent. This position paper examines systemic approaches that might be most useful in overcoming the risks. It applies these approaches to the problems of developing and operating computer-related systems that would be suitable for use in applications requiring trustworthy systems and networking, such as critical infrastructures, privacy-preserving database systems, voting systems, and so on. It also examines the relevance of the Saltzer-Schroeder security principles to elections.

Ultimately, insider misuse cannot be sensibly addressed unless significant improvements are made in system and networking trustworthiness, architecturally, developmentally, and operationally.

Some of the distinctions presented here are intentionally not all clear-cut. There are nuances that must be considered, because blurrings exist among what some people might superficially think are dichotomies. These subtleties can be quite significant in assessing how we should approach insider misuse within the more general context of system and network trustworthiness.

This position paper draws on and extends an earlier one, *The Challenges of Insider Misuse*, that the author wrote for the 1999 RAND Workshop on Preventing, Detecting, and Responding to Malicious Insider Misuse, 16-18 August 1999, of which he was a co-organizer. Sadly, many of the conclusions of that earlier analysis are still relevant today.

Keywords: Insider threat, security policy, access control

What is the Insider Problem?

Matt Bishop (Univ. of California - Davis, US)

The insider problem has long been considered one of the most serious threats to computer security, and one of the most difficult to combat. But the problem has never been defined precisely. Our work presents a definition of the insider problem and shows how the definition enables an analysis of the set of problems traditionally lumped into the “insider problem”. It introduces a hierarchy of policy abstractions, and argues that the “insider” is defined in discrepancies between the different layers of abstraction. It also presents a methodology for analyzing the threat based upon our definitions.

Joint work of: Matt Bishop; Carrie Gates; Sophie Engle; Sean Peisert

The Insider Threat and Security Management

Lizzie Coles-Kemp (RHUL - London, GB)

An organisation often regards an insider as a member of the organisation who is, in some way, empowered to represent it. Traditionally information security management frameworks and systems have a specific formal process for accepting an individual into an organisation and for removing an individual from an organisation but rely on policy compliance, incident management and audit processes for security management of the insider.

A study of 36 organisations and their security management structures concluded that an organisation has multiple views and beliefs about information security and how it should be managed. Compliance is only one of those views. Moreover the study concluded that there are aspects of organisational culture and structure which affect both how those views are expressed and the security management forms they give rise to.

Organisational sociology and theories of organisational management and structure help us to better understand how these different views of information security are formed. They also provide different perspectives through which to evaluate traditional security management methods for responding to the insider threat. This talk presents a number of those theories and summarises the conclusions from case studies in which those theories were applied.

Modeling and Classifying Security Vulnerabilities

Sophie Engle (Univ. of California - Davis, US)

This work began by observing that most classification schemes today are ad-hoc and classify bugs which tend to lead to vulnerabilities, but not classify vulnerabilities themselves. We ask two fundamental questions: (1) what is the difference between a bug and vulnerability, and (2) is it possible to build a non ad-hoc vulnerability classification framework based on characteristics? This has led to the creation of a formal model of security, policy, and vulnerabilities based on the Turing machine. By applying this model to the problem of insider threat, we are able to offer different categories of insider threat, and provide insight to where, why, and how insiders may pose a threat to an organization.

Keywords: Security, policy, vulnerabilities, insider threat, insider problem, formal model, classification

The Insider Threat from the Business Perspective

Ulrich Flegel (SAP Research Center Karlsruhe, DE)

We present our view on the insider threat from the business perspective of our customers. This encompasses two parts.

In the first part we summarize the state of the art of fraud detection in practice and show the relations between the technology for fraud detection and intrusion detection. We identify prospective directions for further investigation and imminent challenges.

In the second part we investigate the privacy dimension of collaborative fraud detection envisioned for outsourcing scenarios. Firstly, we investigate the privacy requirements derived from privacy law and present the resulting judicial argument for pseudonymizing audit data generated for the purpose of fraud detection. Second, we summarize the requirements for such pseudonymization derived from the requirements of the misuse detection approach for fraud detection. Third, we describe our approach for pseudonymization of audit data and two approaches for hiding timestamps in audit data.

Keywords: Insider threat, occupational fraud, privacy law, PET, logical clocks, pseudonyms

Joint work of: Flegel, Ulrich; Kerschbaum, Florian; Wacker, Richard; Vayssière, Julien; Bitz, Gunter

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2008/1794>

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2008/1795>

Detecting Potential Insider Threats in Policies for E-Systems

Michael Huth (Imperial College London, GB)

Increasingly, access to electronic systems or system resources is regulated through the enforcement of policies. These policies are often composed of policies that reside at different locations and are written with local concerns in mind. Such complex policies, as abstract and potentially remote access interfaces, can thus blur the boundaries between insiders and outsiders, and may enable outsiders to become insiders without entering inside a physical domain.

The semantics of real policies is typically given by “the compiler” or “the standard”, creating an obstacle to formal and automated analysis. The academic literature, as another extreme, features toy policy languages with formal semantics such that writing policies, let alone analyzing them, often requires the expertise of a PhD.

In this talk we ponder the feasibility of narrowing this gap between practical but meaningless policy-driven systems and theoretical but inadequate policy languages such that better static or dynamic detection techniques for potential insider threats can be developed.

Keywords: Policies, access control, policy implementation, formal methods

Software Assurance for Countering Insider Threats

Jan Jürjens (The Open University - Milton Keynes, GB)

Compliance frameworks, laws and regulations such as Sarbanes Oxley, Basel II, Solvency II, HIPAA etc. demand from companies in a more and more rigorous way to demonstrate that their organisation, processes and supporting IT landscape implement and follow a set of guidelines at differing levels of abstraction, which have the goal to limit the risks arising from insider threats. This work aims to contribute to a software engineering process which is driven by security, risk and compliance management considerations.

We concentrate on a part of this approach that focusses on the question how one can use software engineering methods and tools to enforce that the configuration of a system enforces the security policies that arise from business compliance regulations, with an emphasis on security policies that aim to counter insider threats. We present tool support for Model-based Compliance Engineering, i.e. for the model-based development and analysis of software configurations that ensures compliance with security policies. It allows one to check UML models of business applications and their configuration data for adherence to security policies and compliance requirements. The tool is based on standardized data formats, such as UML and XML, which makes its integration into existing business architectures as efficient as possible.

More information at: <http://mcs.open.ac.uk/jj2924/research>

Full Paper:

<http://mcs.open.ac.uk/jj2924/publications/papers/fase08.pdf>

See also: J. Jürjens, J. Schreck, and Yijun Yu: Automated Analysis of Permission-Based Security using UMLsec. FASE 2008 (ETAPS)

Countering Insider Threats - a governmental view

Volker Kozok (Bundesministerium der Verteidigung - Bonn, DE)

“Governmental” Definition: Insiders are individuals who are or previously have been authorized to use the information systems they eventually employed to penetrate harm.

Insiders in the governmental area are soldiers, officials, employees, are laborers. The important aspect of governmental insiders is that they have the duty of allegiance, are mostly security checked and trained, working with information that is classified “official secret”.

Cases include unauthorized or illegal attempts to view, disclose, retrieve, delete, change, or add information; exceed or misuse an authorized level of network, system or data access in a manner that affected the security of data, systems or daily business operations of the Federal Armed Forces. Most cases are the use of IT-systems for private purposes.

The reactions consist of disciplinary or labor law sanctions from the internal organization, or in the criminal cases from law enforcement.

Keywords: Critical infrastructure, legal aspects, incident management, motivation

Systematizing Insider Threat Mitigation

George Magklaras (University of Oslo, NO)

The talk outlines a range of techniques that facilitate predictive insider threat modeling. The results of a specific Insider Misuse survey are presented followed by a system-oriented taxonomy that classifies misusers according to system-level detectable consequences. It is important to relate a threat to an insider. Various Insider Threat predictive models are discussed and an Insider Threat Prediction Specification Language expressing system-level specific threat metrics is introduced.

The talk includes extensive references to relevant research areas and should be used mainly as reference material. The research is conducted at the Information Security and Network Research Group, University of Plymouth, UK.

Keywords: Insider Threat, predictive modeling, Domain Specific Languages (DSL), Insider Threat Taxonomy

Beyond Dolev-Yao

Jan Meier (TU Hamburg-Harburg, DE)

This talk proposes to analyze security protocols in the face of insider attacks. In contrast to Dolev-Yao attackers, insiders have initial knowledge. By allowing for insiders, the focus of the protocol analysis shifts. It is no longer the prime goal to verify if an attacker is able to obtain certain information. What becomes more important is what insiders can do with such information. In order to limit the impact of insiders, they should not be able to use information from one protocol run in another protocol run. To verify this additional protocol goal, this talk proposes to model actions insiders can perform. Having a way to measure the impact of insiders enables protocol designers to build more insider-resistant protocols.

Some examples from real life

Vebjorn Moen (GE Money Bank - Stavanger, NO)

A short overview of which kind of insider problems we experience in GE Money Bank, and how we work to handle the risk.

Making Attack Analysis as Simple as Possible, and no Simpler

Sean Peisert (Univ. of California - Davis, US)

Anomaly detection and misuse detection can't stop the insider problem before it happens. The techniques can reduce the insider problem in some ways, but they can't come close to eliminating it. For example, misuse detection doesn't work, by definition. Insiders are doing their jobs, and so how could we recognize them as intruders? In many cases, the best that we can do is analyze events after the fact. However, traditionally audit trails have been some combination of overwhelming, misleading, imprecise, or lacking in detail or scope. We assert that there are methods of addressing this issue, starting with the idea that there are some elements of intrusion analysis which are an art and some which are a science, and that although the proportion will undoubtedly become more scientific over time, there is a portion of the discipline which science can't and shouldn't try to address.

Thus we have started to work on the problem of intrusion analysis, and thus insider analysis, in two ways. First, we have developed a model of forensic logging, which helps to direct and understand the data necessary to log, and also helps to correlate that data once logged. But the amount of data can still be overwhelming to a human analyst who will perform their "analysis art," and so we are also developing methods to help significantly prune away unrelated data that an analyst does not need to look at, and thus, the potential sources of attack or failure are significantly reduced.

We hypothesize that forensics has focused nearly exclusively on analysts rather than presenting meaningful information to the people that actually care and can understand the significance and value of objects and events, and thus is best able to prioritize how to handle them. This makes the best use of human time by spreading the analysis effort, and it allows the lay-analysts to determine the appropriate level of necessary logging and then analyze that data in the context of their area of expertise.

However, to return to our original notion, if all one is doing is generating data, then we've already failed. We simply can't look through that much data on a regular basis. Thus our combined approach is to formalize a process for understanding how to prune data that we know is irrelevant (science), and then present it to the people who understand it best in a way that they can understand it and proceed accordingly (art).

Keywords: Insider, forensics, analysis, attack graphs, optimistic access control, transparent society

Towards a taxonomy of insider action

Joel Predd (RAND - Pittsburgh, US)

A framework is developed for organizing the four key dimensions of insider threats: the organization, the individual, the system, and the environment. The frame work is used as a basis for a taxanomy of insider actions that can be used to characterize different insider threats, and to identify distinguishing features of different insider actions. Implications of the taxonomy for response are discussed.

Joint work of: Joel Predd; Jeffrey Hunker; Shari Lawrence Pfleeger

Modelling the Real World

Christian W. Probst (Technical University of Denmark, DK)

Analysing real-world systems for vulnerabilities with respect to security and safety threats is a difficult undertaking, not least due to a lack of availability of formalizations for those systems. Many approaches to assurance of (critical) infrastructure security are based on (quite successful) ad-hoc techniques. We believe that they can be significantly improved beyond the state-of-the-art by pairing them with static analysis techniques.

We present an approach that formalises real-world systems as directed graphs, which in turn are mapped on a process calculus with support for access control. The process calculus provides the underlying semantics, which allows for easy development of analyses for the abstracted system. The system model is extensible, and allows modular composition of larger scenarios.

Keywords: Insider threat, formal model, access control

Joint work of: Christian W. Probst; René Rydhof Hansen

Overview of the Human Factors Division, Department of Homeland Security, Science and Technology Directorate

Colby Raley (Strategic Analysis, Inc. - Washington, US), Patricia Wolfhope (Department of Homeland Defence, US)

The Human Factors Division of DHS S/T is moving far beyond classical human factors research and development to explore the “human factor” in all DHS-relevant problem spaces.

Human-machine optimization, biometrics, radicalization analysis and observational methods for the detection of suspicious behavior are being explored.

Insider Threat—What can we learn from Human Factors

M. Angela Sasse (University College London, GB)

- systems that are difficult to use lead to mistakes and a negative perception of security
- in an organisation where many users make mistakes, the behaviour of an inside attacker is hard to detect
- need to get honest users “on-side”—human intelligence is only way of detecting novel attacks
- divide and conquer—control in areas associated with high risk, trust elsewhere
- any detected breach of trust must be dealt with, and seen to be dealt with

Insider DoS on Access Control Systems

Ludwig Seitz (Axiomatics AB - Kista, SE)

Many modern access control systems support delegation. Naive implementations of the delegation chain verification algorithm allow for an insider denial of service attack against the access control system by introducing policies that trigger the worst-case complexity of the verification algorithm. This kind of attack is possible even if the insider only has been delegated very limited rights.

Keywords: Access Control, Delegation, Insider Attack, Denial of Service

Insider Threat Detection: Host and Network Monitoring Techniques

Salvatore Stolfo (Columbia University, US)

We presented a definition of insider attack differentiating between masqueraders (attackers who impersonate another inside user) and traitors (an inside attacker using their own, legitimate credentials).

Our research at Columbia is sponsored by the Dartmouth College-hosted I3P organization with funding from the US Department of Homeland Security.

We presented background work on modeling user actions in order to detect abnormal user behavior indicative of masquerade attacks. Intent is modeled by categorizing commands. In particular search actions are especially important to identify likely masquerade attacks. The reasoning is that a masquerader lacks knowledge and would first search to learn what is available on the target system.

Our other line of work is of an offensive nature to confuse and deceive a traitor by leveraging uncertainty, to reduce the knowledge they ordinarily have.

We presented trap-based technology, “decoys” that look like real objects but are purposely planted to detect nefarious acts. An operational example was presented by way of decoy documents with embedded “beacons”. When the decoy is opened, an alert is generated and emailed.

Keywords: Insider threat, decoy, deception, host monitoring, masquerade detection

Criminology theories and the insider treat—A social/psychological perspective

Marianthi Theoharidou (Athens University of Economics and Business, GR)

Modern crime prevention theories are presented and compared with the predominant theory in the Information Systems field—General Deterrence Theory. We examined the incentive mechanisms of these theories and also publications that attempt to describe psychological profiles of insiders. Finally, some implications for practice are discussed.

Insider-free Computing

Alec Yasinsac (University of South Alabama, US)

For centuries, the Maginot Line model was the standard security representation of choice. In this model, defenses focused on protecting the security perimeter from malicious intruders. While modern day armed forces long ago abandoned the Maginot Line model as obsolete in the rapidly changing horizontal and vertical battlefields of today, information security researchers have not been able to replace our perimeter defense posture. Thus, insider attacks continue to garner research attention and to threaten our networks and applications.

In this talk, we consider an emerging model that allows only two types of relationships between subjects (OWNER and NONE), thus eliminating insiders from the environment, and two types of objects (OWNED and UNOWNED). We discuss the security properties of this model, its practical limitations, and the prospective impact of its implementation. The main takeaway from the talk is to encourage minimizing and isolating transferred trust.

Keywords: Isolating trust, Maginot Line