it 4/2011

Identifying and Mitigating Insider Threats

Insiderbedrohungen erkennen und kontrollieren

Christian W. Probst, DTU Informatics, Kongens Lyngby, Denmark

Summary Organisations face many threats that coarsely can be separated in inside threats and outside threats. Threats from insiders are especially hard to counter since insiders have special knowledge and privileges. Therefore, malicious insider actions are hard to distinguish from benign actions. After discussing new definitions of insiders and insider threats, this article gives an overview of how to mitigate insider threats and discusses conflicting goals when dealing with insider threats. ►►► **Zusammenfassung** Organisationen sind mit vielfältigen Bedrohungen konfrontiert, die grob in Innenbedrohungen und Außenbedrohungen eingeteilt werden können. Bedrohungen durch Innentäter sind besonders schwer zu verhindern, da Innentäter besondere Kenntnis und Rechte haben. Daher sind bösartige Handlungen von Innentätern schwer von gutartigen Handlungen zu unterscheiden. Im Anschluss an die Definition der Begriffe Innentäter und Innenbedrohung gibt dieser Artikel einen Überblick über Maßnahmen gegen Innenbedrohungen und diskutiert widersprüchliche Ziele bei der Bekämpfung von Innenbedrohungen.

Keywords K.5 [Computing Milieux: Legal Aspects of Computing]; insider threats, organisational security, sustainable process security ►►► **Schlagwörter** Innentäterbedrohung, Organisationssicherheit, nachhaltige Prozesssicherheit

1 Introduction

When considering security of organisations we must distinguish between *threats* and *attacks*. A threat is a menace against the organisation, that can be caused by, for example, insufficient security precautions, incorrect configuration of security devices, or human error. An attack, on the other hand, represents a realisation of a threat; an attacker has identified a threat and uses it to his advantage. Clearly organisations can face many threats and attacks at the same time. While only an attack can potentially cause "real" damage, threats can for example cause loss of reputation. To minimize the likelihood of a certain attack to happen, organisations must therefore minimize the number of threats they are exposed to, or at least they must ensure that they can identify and are aware of threats.

Threats against an organisation can come from a multitude of sources, which can be classified based on different metrics. If we consider the origin of threats, they can coarsely be divided into threats from the *outside* and threats from the *inside* of the organisation's perimeter.

The usual goal of IT security precautions is to contain threats from the outside of an organisation; protecting the organisation's perimeter and assets against outside threats is fairly well understood, and a multitude of techniques exist. These approaches usually assume that the inside of the organisation is well-behaved or trusted.

Threats from the inside, on the other hand, pose a completely different set of problems. Actors inside of an organisation are legitimated to perform certain actions as part of their work. It is therefore difficult to distinguish benign and malicious activities by insiders. Insiders are trusted to adhere to the policies governing their work, and only to break them in certain situations [3].

In 2008 and 2010, Christian W. Probst has co-organized two Dagstuhl Seminars related to insider threats, the first on how to counter insider threats, and the second on strategies for prevention, mitigation, and response. The aim of both seminars was to bring together the different, quite heterogeneous communities involved in defining insider threats, and initiate the process of developing a common understanding of concepts. In the first seminar a shared, inter-disciplinary definition of the insider and a good formulation for a taxonomy that characterizes insider threats was developed, and this process was continued in the second seminar. To do so it is important to consider many non-IT factors such as the economics of insider threats or the role of law for prevention and penalisation. The goal of this work is to create and test alternative integrated frameworks, so that practitioners and researchers can make informed choices as to combinations of actions targeted at insider threats, and to evaluate the effectiveness of these actions.

In principle, insider threats are easy to contain. One needs some components, such as a concise model of human behaviour and its dependencies on the surroundings, a sufficiently precise surveillance system, and an evaluation system, that can draw the necessary conclusions from its input [9]. None of these components are "easy" to realise, or in any form desirable. We lack techniques to model human behaviour and how it relates to its surroundings; surveillance systems depend on legal boundaries defining what may be monitored; and the amount of data collected even in law-abiding systems is by far too voluminous to allow meaningful analysis.

One can argue that some insider threats are facilitated by a loss of identification with the organisation. This development certainly has its roots in the development of huge companies with several thousands employees, and in lockstep has its roots in the creation of the position of a manager, who increasingly does not have roots in the organisation but is hired because of qualifications in managing other, possibly completely diametral businesses. This leads to a loss of visibility of the individual employee, and a loss of the personal identification of employees with the organisation. Not surprisingly, spies can be seen as the very first "inside attackers". Clearly they were deliberately placed in organisations to obtain specific assets, or damage the organisation's operations.

The question of how to identify insider threats, and how to contain them, is an area of ongoing research [1; 2; 5; 7–10]. Answering it requires insights from many different communities, including but not limited to sociology, law, public policy, criminology, economy, and IT security. All these communities contribute by defining factors that influence an insider's actions, the penalties to be expected, the policies governing what is admissible and what is not, the policies and workflows used in the organisation, detection mechanisms in the real domain and the cyber domain. This list can be continued arbitrarily, but the important point is not only is the problem of insider threats complex, but also the solution.

The problem is further engraved by the need to ensure that the different communities involved have a common understanding of the issues in question. Different approaches to insider threats often use different definitions of what an insider is, and what constitutes an insider threat, and so do different communities. Not only is the definition often implicit or vague, it also is to a certain degree adapted to the concrete case considered.

To conclude, insider threats pose a complex problem that not only has many different facets, but also requires many different approaches and knowledge from many different communities to be solved. Before discussing some of these issues, we will first try to highlight factors that identify an insider threat.

2 Defining Insiders and Insider Threats

Before we are able to target insider threats, we must first understand how to define them. As discussed above, there exist many different aspects of threats that may be relevant, depending on the focus of the approach chosen.

2.1 Insiders

Definitions of insiders and insider threats often concentrate on IT systems. One common definition of an insider is that "[an insider] is defined as an individual with privileged access to an IT system"[1]. This focus on IT systems is not surprising, since they often are used to realise insider attacks, and often are the source of insider threats. On the other hand organisations are so much more than their IT systems, and if already the fundamental definition is putting focus on a specific technique, it is dubious how well it will be useable to address the whole spectrum of issues.

In recent years two developments have had a huge impact on factors defining insiders versus outsiders. On the one hand network access is now ubiquitously available, making all kinds of resources available from almost anywhere. On the other hand, the boundary between inside and outside of an organisation's perimeter seems like an ill-defined concept to distinguish between different actors. In times of joint ventures, out sourcing, cloud computing, and mergers, it is becoming increasingly unclear, what is inside an organisation, and what is outside.

What however is shared by various definitions is their emphasis on the following attributes: *access to the system*, *ability to represent*, *knowledge*, and *trust by the organisation*. Some of the attributes of course are harder to measure and concretize than others, but in general they can serve as a solid basis for defining interesting aspects of insiders.

In 2008, a cross-disciplinary workshop on "Countering Insider Threats" [7] concluded that "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 organisation's structure."

While this definition in principle is helpful by moving the focus from an IT bias to organisational assets, it still is people-centric, which leaves out a large set of problems actually created by IT systems, where processes may execute with rights of an insider, or where an outsider may actually obtain rights that enable executing such processes.

To take account of this problem, we therefore suggest a people-neutral definition of insiders based on the above:

> "An insider is an entity having the ability to access, represent, or decide about one or more assets of the organisation's structure."

Some of the aspects defined clearly require a human actor, but in principle it can be any element in an organisation,

[1] This is the definition used by the Department of Homeland Security (US) research project "Human Factors, Awareness, and Insider Threats", 2007–2009.

from programs and processes, over employees, to visitors, third-party workers, and so on.

2.2 Insider Threats

Once we have a definition of what constitutes an insider, we can address the question of how to define insider threats. The obvious definition is that "an insider threat is the threat that an insider can perform a malicious act." Being parametric in the definition of an insider, this can be used together with any of the definitions above. If we use the last definition based on entities, it therefore also covers programs or outsider with the ability to perform malicious acts.

Of course this definition immediately leads to the question how to define malicious acts. Organisations usually strive to define what is admissible and what not by means of policies, which can be both explicit and implicit. Policies, however, are not necessarily what governs actors' actions, especially not if the actor is malicious. It therefore seems a much better approach to define insider threats with respect to policies or, even better, with respect to the organisation's expectations and goals, which can be quite different from what is expressed in the policies. As a consequence we suggest to define insider threats similar to insiders:

> "An insider threat is the threat that an insider can perform an action that violates the organisation's goals or expectations."

Just like the insider definition above, this definition separates the defined concept from technical aspects. Goals and expectations can be realised as policies, but they can just as well be implicitly expected behaviour.

It is worthwhile to investigate how this definition deals with, e. g., the case of an outside attacker hacking a system and then using it to obtain data. Both these actions constitute an attack, with the former enabling the latter. The threat of the first attack clearly constitutes an insider threat; the outsider is able to access the resource. The threat of obtaining the data also constitutes an insider threat, if the data is not public.

3 Mitigating Insider Threats

As the previous section illustrates, defining the terms *insider* and *insider threat* is not a simple matter, since they depend on many factors related to, e. g., the situation in question or the expected attacks. Many organisations also choose to deliberately ignore insider threats, as long as they can guarantee that enough information is collected to enable and guarantee later retaliation.[2]

Containing insider threats requires three major components: *identification* of potential inside attackers and threats, *monitoring* of operations, and *training* of employees.

[2] This is especially the case when insider attacks would requite shutting down operations, and where this might effectuate unacceptable interruption of service for the organisation.

All of these techniques pose interesting research questions in the light of insider threats, mainly because of the fact that inside attacks are executed with the rights and privileges of legitimate actors. It is therefore so difficult to identify an attack once it started. It is therefore so important to develop techniques to identify insider threats *before* an attack.

3.1 Identifying

Identification provides the signals that other components need to check when monitoring for insider attacks and for the level of insider threats. When dealing with insider threats we need identifying techniques in a number of areas, including, e. g., legal frameworks, policies, and human behaviour. The main goal with these techniques is to provide classifications of events and observations to decide whether or not an insider attack is in progress or to be expected. Taxonomies such as [4; 6] are important components in identifying insider threats, since they establish guidelines for whether a combination of factors should be considered a threat or not.

The analysis of legal frameworks for the public level and policies for the private level allows to determine short-comings, contradictions, inconsistencies, and loopholes in policies and regulations, and their implementation. These distortions often are exploited to realise insider attacks.

Analysis of human behaviour can also help in identifying potential insider threats, or an increased risk for insider threats, by identifying relevant factors in previous cases.

3.2 Monitoring

Monitoring analyses the events in an organisation for signals identifying inside attacks or insider threats. It is important to adapt the level of monitoring to the severeness of threats being faced, and to the value of assets for the organisation. Technically, the main problem with monitoring is to ensure that relevant parameters, as for example determined from identification taxonomies, actually are monitored, and that the amount of data can be handled promptly.

The challenge is to separate legitimate actions by legitimate users from illegitimate actions by illegitimate users. Approaches from intrusion detection have been used for some time, but are limited in applicability due to the wide area of actions that are in principle admissible.

3.3 Training

Finally, training is an important component in containing insider threats. One approach is for organisations to conduct specialized internal exercises [7; 8]. These exercises may be conducted with different goals, all contributing to rising awareness for insider threats. Typical goals include to streamline policies and detect distortions, or to sharpen employees' alertness to insider threats. Thus, training can be used to extract employees' knowledge to

improve policies and workflows, and to identify insider threats or mitigation techniques.

We come back to the interaction with employees below, and we have before discussed employees' identification with a company; clearly the suggested internal exercises, which take up insider threat-related issues from different angles, can also help in identifying potentially weakening identification.

4 Conflicting Goals

In the area of insider threats we have to deal with a number of confliction goals in different areas, where realising either extreme results in severe consequences; these might, e. g., be violation of laws or serious damage of employees' compliance with policies. Here we consider the two most serious ones, namely dealing with surveillance and privacy, and with regulation and autonomy.

4.1 Surveillance vs. Privacy

The agonism between surveillance and privacy can be seen as the most important conflict between goals. When monitoring as much as needed, this will with big probability violate privacy rights of employees and visitors, to name a few. When monitoring as little as possible to protect privacy, the monitored data will often not be meaningful to determine threats and attacks.

Finding the sweet spot between surveillance and privacy requires fine-tuning policies, and means to assure that this level is neither exceeded nor undercut. The problem is that the acceptable level can very well oscillate over time, depending on legal requirements, policies, and so on. Over-monitoring can result in an increased feeling of surveillance and distrust, resulting in a weakening identification of employees with an organisation.

4.2 Regulation vs. Autonomy

As stated above, policies and regulations are important since they determine which actions are admissible for employees. Organisations can therefore be tempted to regulate as many aspects of workflows as possible, leaving little leeway for employees. Actually, many certification processes *require* organisations to demonstrate that they have policies in place to regulate large parts of the organisation's workflows according to certain rules. As research shows, the number of regulations, and the gravity or ease of obeying regulations, has a direct influence on employees' compliance with policies.

This means that also in the areas of regulation and autonomy, there exists the need to identify the sweet spot between the number of policies and the level of compliance. Of course, also this conflict is hard to resolve, especially as compliance is highly depending on external factors, and therefore close to unpredictable.

5 Beyond Insider Threats

In the networked world we live in, insider threats will become less and less important. This is mostly because there no longer exists a real difference between a malicious inside attacker and an outside attacker. To the contrary, we probably need to assume that outside attackers know much more about an organisation's infrastructures and its vulnerabilities, and they know about the latter usually much earlier than the organisation. The problem is that the exploration of an organisation's infrastructure can happen in a very long timeframe, and will therefore be hard to identify by monitoring techniques. Since usually infrastructure does not change quickly, or if it changes does so following certain, often predictable upgrade paths, the knowledge obtained over time is easily available when an attack should be launched.

Because of this vanishing distinction between insiders and outsiders, it seems necessary to also consider the sustainability of business processes. Realising that technical protections alone often are insufficient, or unable to deal with insider threats in a promising way, focus should shift from *mitigation after the fact* to *sustainability of business processes*. Sustainable business processes 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 organisational processes.

6 Conclusion

The term "insider threat" describes a complex problem. Organisations face a multitude of threats, and how to protect against threats from outsiders is fairly well understood. Threats from the inside, however, are much harder to deal with. This is because insiders have special knowledge about the organisation they work in, and they have special privileges because of their status within the organisation. The same holds for IT infrastructure, which executes with certain privileges that an outside attacker does not have immediately.

The definitions for "insider" and "insider threat" given above seem easily applicable. However, there are many issues to be considered, including how to identify insiders who are on the verge of turning malicious, and how to monitor and evaluate actions. To support this, we need techniques for combining analysis of policies and legal restrictions, we need monitoring techniques to collect the necessary data, we need an understanding of the compliance to policies.

To understand and mitigate insider threats one needs to develop a common understanding of the involved concepts between many different communities. This article describes some steps in this direction, and the findings from two recent seminars on the topic.

References

[1] Robert H. Anderson. Research and development initiatives focused on preventing, detecting, and responding to insider misuse of critical defense information systems. Technical Report RAND CF-151-OSD, RAND Corporation, 1999.

[2] Robert H. Anderson, Thomas Bozek, Tom Longstaff, Wayne Meitzler, Michael Skroch, and Ken Van Wyk. Research on mitigating the insider threat to information systems #2. Technical Report RAND CF-163-DARPA, RAND Corporation, 2000.

[3] Matt Bishop. The insider problem revisited. In: *Proc. of the 2005 New Security Paradigms Workshop*, pages 75–76, New York, NY, USA, 2005. ACM.

[4] Matt Bishop, Sophie Engle, Sean Peisert, Sean Whalen, and Carrie Gates. Case studies of an insider framework. In: *Proc. of the 42nd Hawaii Int'l Conf. on System Sciences*, pages 1–10, Washington, DC, USA, 2009. IEEE Computer Society. ISBN 978-0-7695-3450-3. doi: 10.1109/HICSS.2009.104.

[5] Michelle Keeney, Eileen Kowalski, Dawn Cappelli, Andrew Moore, Timothy Shimeall, and Stephanie Rogers. Insider threat study: Computer system sabotage in critical infrastructure sectors. Technical report, U.S. Secret Service and CERT Coordination Center, 2005. Available from http://www.cert.org/insider_threat/insidercross.html (last viewed March 2011).

[6] Joel Predd, Shari Lawrence Pfleeger, Jeffrey Hunker, and Carla Bulford. Insiders behaving badly. In: *IEEE Security and Privacy*, 6:66–70, July 2008. ISSN 1540-7993. doi: 10.1109/MSP.2008.87.

[7] Christian W. Probst, Jeffrey Hunker, Matt Bishop, and Dieter Gollmann. Countering insider threats. *Dagstuhl Seminar Proceedings*, 2008. Available from http://drops.dagstuhl.de/opus/volltexte/2008/1793.

[8] Christian W. Probst, Jeffrey Hunker, Matt Bishop, Lizzie Coles-Kemp, and Dieter Gollmann. Insider threats: Strategies for prevention, mitigation, and response. *Dagstuhl Seminar Proceedings*, 2010a. Available from http://drops.dagstuhl.de/opus/volltexte/2010/2903.

[9] Christian W. Probst, Jeffrey Hunker, Dieter Gollmann, and Matt Bishop, editors. *Insider Threats in Cybersecurity*. Springer, 2010b.

[10] Salvatore Stolfo, Steven Bellovin, Schlomo Hershkop, Angelos Keromytis, Sara Sinclair, and Sean W. Smith, editors. *Insider Attack and Cyber Security: Beyond the Hacker*. Springer, 2008.

Received: May 12, 2011

Prof. Dr. Christian W. Probst is an Associate Professor in the Department for Informatics and Mathematical Modelling at the Technical University of Denmark, where he works in the section for Language-Based Technologies. He holds a Diplom (M. Sc.) and a Dr.Ing. (Ph. D.) in Computer Science from the Universität des Saarlandes, Saarbrücken. The motivation behind his research is to realize systems with guaranteed properties. An important aspect of his work are questions related to safety and security properties, most notably insider threats. He is the creator of ExASyM, the extendable, analysable system model, which supports the identification of insider threats in organisations.

Address: DTU Informatics, Richard Pedersens Plads B. 322, R. 117, 2800 Kongens Lyngby, Denmark, Tel.: +45-45 25 75 12,
e-mail: probst@imm.dtu.dk