



Catching Inside Attackers: Balancing Forensic Detectability and Privacy of Employees

Ephraim Zimmer, Jens Lindemann, Dominik Herrmann, Hannes Federrath

► To cite this version:

Ephraim Zimmer, Jens Lindemann, Dominik Herrmann, Hannes Federrath. Catching Inside Attackers: Balancing Forensic Detectability and Privacy of Employees. International Workshop on Open Problems in Network Security (iNetSec), Oct 2015, Zurich, Switzerland. pp.43-55, 10.1007/978-3-319-39028-4_4 . hal-01445792

HAL Id: hal-01445792

<https://inria.hal.science/hal-01445792>

Submitted on 25 Jan 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Catching Inside Attackers: Balancing Forensic Detectability and Privacy of Employees

Ephraim Zimmer, Jens Lindemann, Dominik Herrmann, and Hannes Federrath

University of Hamburg, Computer Science Department, Germany
{firstname.lastname}@informatik.uni-hamburg.de

Abstract. IT departments of organisations go to great lengths to protect their IT infrastructure from external attackers. However, internal attacks also pose a large threat to organisations. Despite detection and prevention of insider attacks being an active field of research, so far such techniques are rarely being deployed in practice. This paper outlines the state of the art in the field and identifies open research problems in the area. The lack of unified definitions and publicly available datasets for evaluation is detrimental to the comparability of published results in the field and hinders the continual improvement of technology. Another important problem is that of data protection: On the one hand, the data captured for insider attack detection could also be used for surveillance of employees, so it should be anonymised. On the other hand, anonymisation may make some attacks undetectable, leading to a trade-off between detectability of attacks and privacy.

1 Introduction

Nowadays, an organisation’s IT infrastructure will typically be connected to the Internet. Internal and external communications rely on digital technology and employees often require the Internet to research information necessary for their day-to-day work. To protect against attacks from the Internet, which have increased dramatically in recent years [3], IT security departments rely on established security mechanisms, such as firewalls, Intrusion Detection Systems (IDS), honeypots and so-called De-Militarised Zones (DMZ). These measures are intended to stop external attackers who are trying to interfere with the execution of business processes or obtain internal assets, such as trade secrets or confidential customer data. However, the danger arising from attacks originating from within the organisation itself is often overlooked.

A workshop on *Countering Insider Threats* in 2008 defined an *inside attacker* as “a person that has been legitimately empowered with the right to access, represent, or decide about one or more assets of the organization’s structure” [26]. However, the term is not defined consistently throughout the scientific literature. According to Pfleeger, an inside attacker can also be “anyone properly identified and authenticated to the system including, perhaps, someone masquerading as a legitimate insider, or someone to whom an insider has given access (for example by sharing a password)” [25]. The two examples include regular employees as well

as sophisticated system administrators, but give different context variabilities of an inside attacker. The former implicitly exclude outside attackers masquerading as insiders, whereas the latter includes those scenarios into its definition.

Insider attacks are a major threat for organisations. As insiders typically have extensive access rights (especially if they are system administrators) and possess detailed knowledge about the IT infrastructure of their organisation, they know where to strike for maximum impact and are capable of hiding their activities. In a survey of the CERT Insider Threat Center among US companies 47 percent of companies acknowledged that they were knowingly affected by insider attacks throughout the years 2004 to 2013 [5]. The dark figure might very well exceed this number significantly, as insider attacks often are either not detected by organisations or withheld from the public due to the high risk of reputation loss. Less than half of those surveyed companies have deployed defined mechanisms and procedures to deal with insider attacks [23]. Experts estimate the impact of insider attacks on the German economy at about 50 billion Euros per year [32].

The security mechanisms that are currently in practical use cannot adequately detect attacks by insiders. Due to their access rights and inside knowledge, insiders can hide malicious activity significantly easier than external attackers, e. g. by deactivating security systems or manipulating log files. Moreover, insider attacks are hard to detect at the network perimeter, where traditional security mechanisms are typically located.

Even though several technical detection and prevention mechanisms have been proposed by various researchers, those mechanisms have not reached widespread practical implementation and deployment yet. Currently, insider threats are mainly being countered by organisational measures, such as by imposing a two-man rule for actions having a high impact on security (e. g. disabling a firewall or modifying log files) [4,30].

In discussions about insider attack detection, data protection is often overlooked. However, comprehensive logs for detecting and attributing insider attacks can reveal a lot of information about the behaviour of employees, thereby invading their privacy. On the other hand, if too much anonymisation is performed before passing the data to a detection algorithm, some attacks may not be detectable anymore.

The remainder of the paper is structured as follows. In Section 2, we describe the state of the art in detection and prevention mechanisms for insider attacks. Section 3 shows some avenues for future research. In Section 4, challenges for the development of the field are presented, before the paper is concluded in Section 5.

2 State of the art in detection and prevention mechanisms

Researchers have proposed to counter insider threats by means of technical and non-technical mechanisms [9,27]. With the help of technical solutions, attack detection and monitoring data can be collected, correlated and analysed for insider activities during or after an insider attack. Non-technical solutions and best practices serve the goal of insider detection and prevention by providing

strict policies and evaluating information about social behaviour, an employee's productivity or insights from Human Resources (HR), like imminent employee terminations. (Semi-)automatic collection and assessment of this kind of data as well as the establishment of corporation-wide guidelines, employee training and formal policies (focussing on insider threats) are supposed to effectively unveil and eliminate insider activities and attacks. In the following, we review recent insider threat research and related work.

2.1 Non-technical means of protection

In terms of organisational or structural protection mechanisms against the insider threat, the literature focuses mainly on motivation- and opportunity-based countermeasures. Examples are the destruction of incentives for insider attacks [24], training to change employees' mindsets as well as a close cooperation with HR to obtain deep insight into employees' projects and groups to identify employees who need to have access to sensitive data [14].

Silowash et. al [30] took a more structured approach and developed a common sense guide to mitigating insider threats including various practices to prepare an organisation for correctly dealing with insider threats. Among those practices, the guide considers the following non-technical countermeasures:

- Consider threats from insiders and business partners in enterprise-wide risk assessments.
- Clearly document and consistently enforce policies and controls.
- Incorporate insider threat awareness into periodic security training for all employees.
- Beginning with the hiring process, monitor and respond to suspicious or disruptive behaviour.
- Anticipate and manage negative issues in the work environment.
- Know your assets.
- Enforce separation of duties and least privilege.
- Develop a comprehensive employee termination procedure.
- Develop a formalised insider threat program.

This guide as well as other proposed non-technical countermeasures are mainly derived from control domains specified in Annex A of ISO 27001, as Coles-Kemp and Theoharidou showed [6], and deal with the insider threat on a very high and abstract level. It takes a lot of effort and the involvement of a whole corporation to realise and run them in practice. Furthermore, due to the corporation-specific execution of practices it is difficult to transfer them to another corporation or environment.

2.2 Technical means of protection

At first sight, most, if not all, traditional technical countermeasures that are used to protect organisations against cyber attacks (like Intrusion Detection Systems (IDSs) and log data analysis) may also be employed to detect and prevent insider activities. However, the domain and circumstances of an insider attack are

fundamentally different: On the one hand, this leads to significantly more alerts and increased false positive rates. On the other hand, those mechanisms might possibly be tricked or circumvented by insiders with the help of their specific and internal knowledge.

Behavioural profiling of users

Research on technical protection mechanisms against insider attacks has devoted a lot of attention to profiling employee behaviour. Here, the objective is to learn the legitimate characteristics of users in order to perform (semi-)automatic detection of potentially anomalous insider activities. These approaches also strike the threat of masqueraders, who are outside attackers possessing stolen credentials of employees and therefore have access to inside resources and systems.

Schonlau et. al [28] studied possibilities of detecting insider attacks by profiling Unix shell commands. Over several months, they collected shell commands of 50 different users and additionally simulated insider activities by injecting commands of users who played the role of masqueraders. Based on this dataset, they tried to evaluate different methods of anomaly detection. The results showed a rather high rate of false alarms as well as false negatives. Later, other researchers re-used the Schonlau dataset, applying improved detection methods [21]. Although promising, the Schonlau dataset does not provide a very good base for evaluating insider attack detection mechanisms as the masquerader simulation is rather artificial and Unix is only a small part of employees' production environments.

Other approaches considered user profiling in the context of the graphical user interface of Microsoft Windows. Goldring [10] evaluated user profiling by periodically collecting data from the Windows process table in short intervals. This data shows the lifecycle and additional information (such as owner and CPU usage) of all programs that have been or are running on a system. To filter out operating system noise, Goldring exploited the fact that each user interaction with the system takes place in a window. Therefore, he additionally took window titles into account. The resulting concept looks promising, but associated evaluation results have not been published.

Li and Manicopoulos [17] also studied profiling of Windows users. They created a dataset with simulated insider attacks (similar to the Schonlau dataset) and applied a one-class Support Vector Machine (SVM) to build models of legitimate user behaviours. With this model, a binary classifier could be used to test new models for compliance or deviation. However, their technique achieves only moderate accuracy in terms of detection and false alarms rates and their dataset entails the same deficiencies as the Schonlau dataset.

Network-based approaches

Besides host-based user behaviour profiling, corporate network traffic comprises a great source of information about employees' IT activities and thus valuable

data for insider attack detection and prevention mechanisms. Spitzner [31] applied the now widespread knowledge and application of honeypots and honeytokens from the domain of outsider attack countermeasures and perimeter threats to the insider threat. The idea is to stimulate the interest of inside attackers, who are looking for some kind of valuable information, in specialised honeytokens. Whenever attackers access a honeytokens, they are automatically redirected to a honeypot, where the interaction can be monitored and analysed in a secure environment. Although interesting as a means to decrease false positive rates in insider detection mechanisms, the concept has not been evaluated, which by design is very hard to conduct. Only empirical evaluation of practically deployed systems could provide reliable results, as simulated insider attacks are not suitable. Further, the effectiveness of honeytokens and honeypots in insider attack detection and prevention is highly dependent on several attributes of an inside attacker, like knowledge of countermeasures, technical skill level and level of suspicion.

Maloof and Stephens [20] also concentrated their work on network traffic collection and analysis. They created a system called ELICIT, which aims at the detection of inside attackers, who try to access information they do not need to know according to their job description and similar additionally acquired information. The ELICIT system consists of four parts: First, network traffic is collected and prepared in the form of events. Secondly, events are enriched with additional contextual information about employees and alerts are issued. Thirdly, a threat score is calculated based on a Bayesian network, which takes these alerts as input. Finally, the scores are presented for further examination by security personnel. The authors evaluated their system by collecting internal data of an organisation over several months, replayed activities from publicly known past insider attack cases and applied the dataset offline to ELICIT afterwards. The evaluation showed very good results in terms of detection rates and remarkably low false positive rates. However, the system strongly relies on the presence of machine-interpretable contextual information about employees, job descriptions, need-to-know domains and such. For this information to be present, a corporation needs to have strict HR policies and comprehensive procedures in place, which may hinder adoption in practice. Additionally, the simulated execution of past insider attacks occurred over few days in contrast to real world insider attacks, which more likely occur over several weeks, months or years [27].

2.3 Integrated approaches

Recent proposals go one step further and try to integrate non-technical approaches with technical countermeasures to combine their advantages and create comprehensive insider attack detection and especially prevention systems. Gritzer and Frincke [12] took psychological data in addition to classical security audit data into account. The objective was to create possibilities in predicting insider activities of employees by means of a set of predictive indicators and an integration and analysis framework for organisational, social and cyber security

data. Costa et al. [7] created an ontology-based approach by studying 800 real-world insider attack cases, allowing them to identify entities involved, insider actions conducted, assets targeted and events triggered. This information was translated into an insider threat indicator ontology and combined and enriched with (semi-)automatically processable operational context data from HR. With the help of a semantic reasoner, which monitors current activities and responsibilities of a corporation and evaluates this information against the ontology, potential insider activities could be identified and alerted.

Major challenges for all integrated approaches were found to be the lack of reliable testing and evaluation datasets, no operational evaluation, privacy and ethical issues, and the need for extensive training and awareness of employees.

3 Avenues for future research

The measures outlined in the previous Section aim to effectively detect or prevent insider attacks. However, these measures ignore some fundamental problems, which will be outlined in the following.

3.1 Unified definition of terms, motives and tools

Research on insider attacks is only meaningful in the context of a concrete adversary model that describes capabilities and motives. For instance, system administrators are more powerful than regular employees due to their extensive access to all IT infrastructure and monitoring devices. Moreover, strategically-acting intentional attackers have to be treated differently than users who bring their own devices to work and infect the corporate network with malware inadvertently. With modern forms of e-commerce and outsourcing of IT services, even third parties may act like insiders [13]. Customers who rent infrastructure or software as a service can either interfere with the underlying infrastructure or exploit it to launch attacks against others [18].

Despite experts from science, industry, the financial sector, and the US government concluding that there is a lack of standardised definitions for insiders and insider attacks during a workshop on *Insider Attack and Cyber Security* in 2007 [25], there are still no such definitions. This leads to scientists using different definitions for their research, typically choosing a definition that is beneficial to their research project and expected outcome. As already stated with the two examples of a definition in Section 1, different foci on the context variabilities of an inside attacker for example provide different, sometimes even competing results. This leads to a delusive comparison of countermeasures, which seem to provide solutions for the same insider problem, when in fact the problem domain is significantly different. Even four years later, scientists aiming to establish unified definitions came to the conclusion that additional, more detailed definitions would have to be established [13]. The authors state, that current definitions lack the reflection of two recent developments. First, the new capabilities and applications of networked environments. And second, the increasingly

indistinct separation of corporation boundaries. Further, the characterisation of inside attackers becomes progressively multidimensional, emphasising different capabilities or circumstances of an inside attacker or an insider attack.

As a conclusion, this leads to diverging results being published by different research groups on the one hand. On the other hand, the comparability of published results is impeded by the lack of unified definitions, which leads to continuous improvement of scientific results not taking place.

3.2 Generation of datasets

Technical approaches for detection and prevention of insider attacks published by scientists – such as [7] – have so far failed to find widespread use in practice [27]. This can partly be attributed to the unsatisfactory effectiveness of these approaches, as is evident by their high false positive and false negative rates. Another problem is that it is difficult for researchers to evaluate their proposed solutions in a way that approximates their behaviour in a real production environment and allows a comparison between results published by different research groups. This is due to the lack of comprehensive datasets of insider attacks captured in a real production environment, which could be used for realistically evaluating effectiveness and efficiency [13,27]. Existing datasets have either been taken from a different context (e.g. the Schonlau dataset [28]) or have been derived from simulations that made special assumptions about the attacking insider and the organisation (such as time restrictions or existing formal regulations [20]). As these methods for generating datasets cannot provide a realistic approximation of insider attacks [27], they are not useful for evaluating detection mechanisms.

3.3 Software implementations and their evaluation

Existing work	False positive rate	Detection rate
Honeypots [31]	-	-
ELICIT [20]	1.5	84.0
Unix commands [28]	6.7	69.3
Unix commands [21]	1.3	61.5
MS Windows [10]	-	-
MS Windows [17]	22.0	67.7
Psychology [12]	-	-
Ontology [7]	-	-

Table 1. Percentage of false positive and detection rates of a selection of insider attack detection mechanisms

Software implementations of scientifically proposed solutions for the detection and prevention of insider attacks as well as the (semi-)automatic collection

and evaluation of additional information sources – as proposed by Costa et al. [7] and Maybury et al. [22] – are not publicly available and can thus not be evaluated or verified by others. This is detrimental to the use of these solutions in practice. Even the evaluation of proposed techniques by their authors is missing in many publications, as is illustrated by a selection of techniques shown in Table 1. Even if an evaluation exists, the numbers for false positive and detection rates are not directly comparable, since most authors used different, often newly constructed, simulated data on insider activities. Furthermore, in addition to not having been evaluated using realistic datasets, there is also a lack of evaluation of their use in production environment, e.g. through field tests, which in some cases, like the system of honeypots and honeytokens by Spitzner [31], is the only way to properly evaluate a proposed countermeasure.

3.4 Post-mortem detection

The long-term objective consists in designing *preventive security mechanisms* against insider attacks. Unfortunately, it is questionable whether effective protection is achievable at all. However, *detecting attacks* (post-mortem) and identifying the culprit (attribution) might be sufficient to deter insiders in practice.

Thus, a potential avenue of research is to focus on detecting insider attacks. Existing technical approaches, as described in Section 2.2, focus on data collection, user profiling and anomaly detection. However, inside attackers typically have extensive access rights (especially if they are system administrators) and possess detailed knowledge about the IT infrastructure of their organisation, which endows them with the capability of changing data, manipulating user profiles and hiding their activities. A (debatable) approach is to deploy a comprehensive logging infrastructure that monitors the behaviour of all users and systems from a number of vantage points. This allows the examination of specific events from many different point of views of a system and makes it harder for attackers to cover their tracks, as they would have to manipulate logs in many different places in a consistent manner. Manipulating logs of one vantage point would in turn generate traces in other logs, leading to detectable inconsistencies that make it possible to verify the veracity [11] of the information presented in logs.

4 Challenges

In the following, we will show some challenges which will have to be solved in order to create an effective detection and prevention system for insider attacks.

4.1 Techniques for post-mortem detection of insider attacks

One of these challenges is the *automatic detection of insider attacks* based on log data and contextual information collected by the system. In addition to data collected from IT systems, data available from physical security systems

could also be considered, such as biometrical access control systems or motion detectors, which could provide information about suspicious “offline” activities (e. g. access to server rooms at unusual times) and help with attributing activity to specific users. Algorithms from the field of anomaly detection will have to be adapted to this scenario. Special attention needs to be paid to the number of false positives generated by the anomaly detection algorithms, as even seemingly low false positive rates can lead to the number of false positives vastly exceeding the number of true positives [2], which may cause true positives to be shrugged off as “yet another false alarm”.

4.2 Data protection

The most important challenge for the implementation of a comprehensive logging system (as described in Section 3.4) is *sufficient data protection*. As the system would continuously monitor and log activities of all employees, the data produced by it would have to be protected in order to comply with data protection laws concerning employment. Insufficient protection may lead to the data being misused for surveillance of employees. Additionally, the system might also log sensitive data related to customers. Therefore, any information that could be used to identify persons (customers or employees) should be obfuscated (e. g. by pseudonyms as proposed and argued in the context of internal fraud screening by Flegel [8]) or removed altogether. It should only be possible to reverse this in case of a suspected security incident and it should not be possible for a single person to link log entries to persons.

A possible way of ensuring the consent of multiple parties before information is de-anonymised would be the use of a threshold decryption scheme. These schemes require a minimum number of private keys – but not necessarily all that are part of the scheme – to be present to decrypt data previously encrypted using a public key [29]. A similar approach has been developed by Armknecht and Dewald [1] in the context of digital forensics on sensitive e-mail data.

In some cases, partly de-anonymising data before applying anomaly detection algorithms to it may be necessary to be able to detect attacks at all. An example of this are login attempts on a server, where the IP address or at least information about the geographical location of the computer trying to connect would be relevant for detecting anomalies. On the other hand, even incomplete de-anonymisation may lead to linkability of certain types of behaviour to individual employees, invading their privacy. This shows that there is a trade-off between improved detectability of some attacks and user privacy related to how thoroughly the data used for anomaly detection is anonymised. Experiments will have to be performed to evaluate the impact of different forms and extents of anonymisation on the detectability of insider attacks. A similar trade-off exists in intrusion detection on network traces and has been discussed by Lakkaruja and Slagell [16] as well as Lundin and Jonsson [19]. Compared to detection of incoming attacks by intrusion detection systems, privacy is significantly more important in insider attack detection, as the data used for it will focus on the contextual information and activities of an organisation’s employees, which are

using the system over a long period of time, thus making it relatively easy to build profiles of them for illegitimate purposes.

Before systems for the detection and prevention of insider attacks can be used in practice, a sweet spot on this trade-off will have to be found. If no anonymisation was performed, these systems could not be used in practice at all in many legislations, as the privacy of employees and/or customers would be invaded. On the other hand, total anonymisation will likely remove too much information for the system to be of any help in tackling insider attacks. One way to solve this dilemma may be to adapt solutions that allow anonymity to be revoked under certain conditions [15].

4.3 Realistic datasets for evaluation of detection techniques

Another challenge is to create publicly available datasets which can be used to evaluate and improve detection and prevention techniques for insider attacks. As outlined in Section 3.2, there is currently a lack of such datasets. One possible way of obtaining a dataset would be to capture it in a real production environment that may be affected by insider attacks, i.e. a corporate or government network. However, it is unlikely that such organisations would be willing to let researchers capture a comprehensive dataset, as this would reveal all activity conducted within the organisation’s IT infrastructure. This would make it necessary to remove or replace all confidential information from the dataset, which in itself poses a challenging research problem. Even if one was to find a way to achieve this, it is likely that the anonymisation would impact detection and prevention techniques, leading to the dataset potentially not resembling the real environment closely enough, rendering it useless for evaluation purposes.

With capturing data within an organisation’s network being unrealistic, simulation remains as a way of generating a publicly available dataset. The challenge here is to make sure that a synthetically generated dataset resembles real world environments closely enough to allow evaluation of detection and prevention techniques as well as to provide means of robust comparability between different countermeasures.

5 Conclusion

In this paper, we analyse the current state of the art in insider attack detection and prevention and show some potential avenues of future research as well as challenges in the field. Our analysis has shown that existing security mechanisms cannot prevent insider attacks reliably. Detection and attribution is complicated by the ability of insiders to cover their tracks and fabricate evidence. Therefore, designing effective preventive, reactive and forensic techniques seems to be a fruitful area of future research.

Advances towards more effective techniques are hindered by a lack of unified definitions in the field and no datasets being publicly available that resemble real production environments closely enough to allow a comparative evaluation

of techniques. Furthermore, previous research often ignores data protection and does not take the trade-off between detectability of insider attacks and protection of employee data into account.

References

1. Armknecht, F., Dewald, A.: Privacy-preserving email forensics. *Digital Investigation* 14, Supplement 1, 127 – 136 (2015), <http://www.sciencedirect.com/science/article/pii/S1742287615000481>, the Proceedings of the Fifteenth Annual DFRWS Conference
2. Axelsson, S.: The base-rate fallacy and the difficulty of intrusion detection. *ACM Trans. Inf. Syst. Secur.* 3(3), 186–205 (2000)
3. Bundeskriminalamt: Polizeilich erfasste fälle von cyberkriminalität im engeren sinne in deutschland von 2000 bis 2014. Statista — Das Statistik-Portal (2014), <http://de.statista.com/statistik/daten/studie/295265/umfrage/polizeilich-erfasste-faelle-von-cyberkriminalitaet-im-engeren-sinne-in-deutschland/>
4. Centre for the Protection of National Infrastructure: Ongoing personnel security: A good practise guide (2014), <http://www.cpni.gov.uk/documents/publications/2014/2014006-ongoing-personal-security.pdf?epslanguage=en-gb>
5. CERT Insider Threat Center: 2014 U.S. State of Cybercrime Survey (2014), <http://resources.sei.cmu.edu/library/asset-view.cfm?assetID=298318>
6. Coles-Kemp, L., Theoharidou, M.: Insider threat and information security management. In: *Insider Threats in Cyber Security*, pp. 45–71. Springer (2010)
7. Costa, D.L., Collins, M.L., Perl, S.J., Albrethsen, M.J., Silowash, G.J., Spooner, D.L.: An ontology for insider threat indicators. 10th International Conference on Semantic Technology for Intelligence, Defense, and Security (STIDS) (2015), <http://resources.sei.cmu.edu/library/asset-view.cfm?assetid=426803>
8. Flegel, U.: Privacy compliant internal fraud screening. In: Pohlmann, N., Reimer, H., Schneider, W. (eds.) *ISSE 2010 Securing Electronic Business Processes*, pp. 191–199. Vieweg+Teubner (2011)
9. Flynn, L., Huth, C., Trzeciak, R., Buttles-Valdez, P.: Best practices against insider threats in all nations. Tech. Rep. CMU/SEI-2013-TN-023, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA (2013), <http://resources.sei.cmu.edu/library/asset-view.cfm?AssetID=59082>
10. Goldring, T.: User profiling for intrusion detection in windows nt. In: *Proceedings of the 35th Symposium on the Interface* (2003)
11. Gollmann, D.: Veracity, plausibility, and reputation. In: Askoxylakis, I.G., Pöhls, H.C., Posegga, J. (eds.) *Information Security Theory and Practice. Security, Privacy and Trust in Computing Systems and Ambient Intelligent Ecosystems - 6th IFIP WG 11.2 International Workshop, WISTP 2012, Egham, UK, June 20-22, 2012. Proceedings. Lecture Notes in Computer Science*, vol. 7322, pp. 20–28. Springer (2012)
12. Greitzer, F.L., Frincke, D.A.: Combining traditional cyber security audit data with psychosocial data: towards predictive modeling for insider threat mitigation. In: *Insider Threats in Cyber Security*, pp. 85–113. Springer (2010)
13. Hunker, J., Probst, C.W.: Insiders and insider threats – an overview of definitions and mitigation techniques. *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications* 2(1), 4–27 (2011)

14. Kaplan, J.M., Bailey, T., O'Halloran, D., Marcus, A., Chris, R.: Beyond Cybersecurity: Protecting Your Digital Business. John Wiley & Sons (2015)
15. Köpsell, S., Wendolsky, R., Federrath, H.: Revocable anonymity. In: International Conference on Emerging Trends in Information and Communication Security (ET-RICS'06). pp. 206–220. Springer (2006)
16. Lakkaraju, K., Slagell, A.J.: Evaluating the utility of anonymized network traces for intrusion detection. In: Levi, A., Liu, P., Molva, R. (eds.) 4th International ICST Conference on Security and Privacy in Communication Networks, SECURECOMM 2008, Istanbul, Turkey, September 22-25, 2008. p. 17. ACM (2008)
17. Li, L., Manikopoulos, C.N.: Windows nt one-class masquerade detection. In: Proceedings from the Fifth Annual IEEE SMC Information Assurance Workshop. pp. 82–87. IEEE Computer Society (June 2004)
18. Lindemann, J.: Towards Abuse Detection and Prevention in IaaS Cloud Computing. In: Processings of the 10th International Conference on Availability, Reliability and Security (ARES 2015). IEEE Computer Society (2015)
19. Lundin, E., Jonsson, E.: Anomaly-based intrusion detection: privacy concerns and other problems. *Computer Networks* 34(4), 623–640 (2000), [http://dx.doi.org/10.1016/S1389-1286\(00\)00134-1](http://dx.doi.org/10.1016/S1389-1286(00)00134-1)
20. Maloof, M.A., Stephens, G.D.: Elicit: A system for detecting insiders who violate need-to-know. In: Kruegel, C., Lippmann, R., Clark, A. (eds.) Recent Advances in Intrusion Detection, Lecture Notes in Computer Science, vol. 4637, pp. 146–166. Springer Berlin Heidelberg (2007), http://dx.doi.org/10.1007/978-3-540-74320-0_8
21. Maxion, R., Townsend, T.: Masquerade detection augmented with error analysis. *Transactions on Reliability* 53(1), 124–147 (March 2004)
22. Maybury, M., Chase, P., Cheikes, B., Brackney, D., Matzner, S., Hetherington, T., Wood, B., Sibley, C., Marin, J., Longstaff, T.: Analysis and detection of malicious insiders. Tech. rep., DTIC Document (2005), <http://oai.dtic.mil/oai/oai?verb=getRecord&metadataPrefix=html&identifier=ADA456356>
23. Michelberg, K., Schive, L., Pollard, N.: Us cybercrime: Rising risks, reduced readiness — key findings from the 2014 us state of cybercrime survey (2014), <https://www.pwc.com/us/en/increasing-it-effectiveness/publications/2014-us-state-of-cybercrime.html>
24. Neumann, P.G.: Combatting insider threats. In: Insider Threats in Cyber Security, pp. 17–44. Springer (2010)
25. Pfleeger, C.P.: Reflections on the insider threat. In: Stolfo, S.J., Bellovin, S.M., Keromytis, A.D., Hershkop, S., Smith, S.W., Sinclair, S. (eds.) Insider Attack and Cyber Security, Advances in Information Security, vol. 39, pp. 5–16. Springer US (2008), http://dx.doi.org/10.1007/978-0-387-77322-3_5
26. Probst, C.W., Hunker, J., Gollmann, D., Bishop, M.: Aspects of insider threats. In: Probst, C.W., Hunker, J., Gollmann, D., Bishop, M. (eds.) Insider Threats in Cyber Security, Advances in Information Security, vol. 49, pp. 1–15. Springer US (2010), http://dx.doi.org/10.1007/978-1-4419-7133-3_1
27. Salem, M.B., Hershkop, S., Stolfo, S.J.: A survey of insider attack detection research. In: Stolfo, S.J., Bellovin, S.M., Keromytis, A.D., Hershkop, S., Smith, S.W., Sinclair, S. (eds.) Insider Attack and Cyber Security, Advances in Information Security, vol. 39, pp. 69–90. Springer US (2008), http://dx.doi.org/10.1007/978-0-387-77322-3_5
28. Schonlau, M., DuMouchel, W., Ju, W.H., Karr, A.F., Theus, M., Vardi, Y.: Computer intrusion: Detecting masquerades. *Statistical Science* 16(1), 58–74 (2001), <http://www.jstor.org/stable/2676780>

29. Shoup, V.: Practical threshold signatures. In: Preneel, B. (ed.) *Advances in Cryptology - EUROCRYPT 2000*, International Conference on the Theory and Application of Cryptographic Techniques, Bruges, Belgium, May 14-18, 2000, *Proceeding. Lecture Notes in Computer Science*, vol. 1807, pp. 207–220. Springer (2000), http://dx.doi.org/10.1007/3-540-45539-6_15
30. Silowash, G., Cappelli, D., Moore, A., Trzeciak, R., Shimeall, T., Flynn, L.: Common sense guide to mitigating insider threats. Tech. Rep. CMU/SEI-2012-TR-012, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA (2012), <http://resources.sei.cmu.edu/library/asset-view.cfm?AssetID=34017>
31. Spitzner, L.: Honeypots: catching the insider threat. In: *Computer Security Applications Conference, 2003. Proceedings. 19th Annual.* pp. 170–179. IEEE Computer Society (December 2003)
32. Zimmermann, S.: *Wirtschaftsspionage – Gefahr im eigenen Haus* (2015), <http://dw.de/p/1FPAo>