# A critical reflection on the threat from human insiders – its nature, industry perceptions, and detection approaches

Jason R.C. Nurse[1], Philip A. Legg[1], Oliver Buckley[1], Ioannis Agrafiotis[1], Gordon Wright[2], Monica Whitty[2], David Upton[3], Michael Goldsmith[1], and Sadie Creese[1]

[1] Cyber Security Centre, Department of Computer Science, University of Oxford, UK
{firstname.lastname}@cs.ox.ac.uk
[2] Department of Media and Communications, University of Leicester, UK
grw9@leicester.ac.uk, mw229@leicester.ac.uk
[3] Saïd Business School, University of Oxford, UK
david.upton@sbs.ox.ac.uk

**Abstract.** Organisations today operate in a world fraught with threats, including "script kiddies", hackers, hacktivists and advanced persistent threats. Although these threats can be harmful to an enterprise, a potentially more devastating and anecdotally more likely threat is that of the malicious insider. These trusted individuals have access to valuable company systems and data, and are well placed to undermine security measures and to attack their employers. In this paper, we engage in a critical reflection on the insider threat in order to better understand the nature of attacks, associated human factors, perceptions of threats, and detection approaches. We differentiate our work from other contributions by moving away from a purely academic perspective, and instead focus on distilling industrial reports (i.e., those that capture practitioners' experiences and feedback) and case studies in order to truly appreciate how insider attacks occur in practice and how viable preventative solutions may be developed.

**Key words:** insider threats, human factors, technical and psychological indicators, detection approaches, survey reports

## 1 Introduction

Corporations today face an increasingly difficult task when it comes to their computer security. On the one hand, there are a plethora of threats (e.g., criminals, hackers, hacktivists) keen to penetrate defences and compromise systems and data. On the other hand, internal (or insider) threats appear to be on the increase and can be particularly debilitating given their privileged access to the enterprise. The insider-threat problem is especially concerning because corporations' defences are arguably still focused on external threats, resulting in inadequate consideration of attacks originating from those with inside knowledge of and access to systems, security processes, and precious company secrets.

To explore this problem further, and to better understand the various elements involved, this paper engages in a critical reflection upon the threat posed by insiders. We adopt a novel perspective that moves away from a purely theoretical discussion and instead concentrates on distilling the range of industrial reports, which capture broad experiences and feedback from practitioners [1, 2, 3, 4]. We also look at case studies of insider-threat (our own [5] and those from CMU-CERT [6]), in order to further understand how and why insider attacks occur, and how effective detection tools can be developed and deployed.

Our reflection on the insider-threat problem is split into three broad sections. Firstly, we consider the nature of human insider-threats. This includes an investigation into the types of attacks actually being launched against enterprises, an analysis of the motives and psychological aspects surrounding these attacks, and the impact that new technologies may have on the future of insider attacks. We move on to study many of the industry reports that have been published (e.g., [2, 7, 8]), in order to assess how corporations perceive and are responding to this type of risk. Our findings suggest that there is an underestimation of the risks associated with these threats, particularly evidenced by the minimal investment being made. Finally, we describe techniques that are currently used for detecting insider threats, and explore the state-of-the-art research that is currently being conducted in this area, discussing the effectiveness of techniques and what limitations may exist. To conclude, we discuss own research within the Corporate Insider Threat Detection project (CITD), which aims to address the interdisciplinary nature of insider threat, to provide an enhanced detection tool that addresses both technical and human dimensions of insider threat.

## 2 The nature of insider threat

In order to understand the nature of the insider-threat problem, there are several fundamental questions of interest. For instance, what exactly is the threat, and what are the most prevalent types? What motivates insiders to attack? Are some insiders more susceptible to becoming a threat? What behaviours may be indicative of an (impending) attack? What is the effect, if any, of new technologies on the problem? These are the questions which we seek to discuss in this section, with a special focus on real-world cases, feedback and reports.

### 2.1 Types of insider threat

There have been many definitions of insider threat throughout the years [9]. Some of these definitions emphasise the active misuse of insider privileges, while others broaden the scope and consider the negative impact of such misuse on the confidentiality, integrity and availability of the organisation's systems and data [6]. The essence of most definitions, however, is that an insider threat is a member of trusted personnel (e.g., employee, contractors, business partners) that used their privileged access for some unauthorised purpose such as revenge

or financial gain, and to the detriment of their enterprise. CMU-CERT [6] identifies three types of threat based on observation of typical patterns and on the attacker's purpose and motivation – namely, fraud, theft of Intellectual Property (IP), and sabotage of infrastructure.

Insider fraud is regarded as one the most frequent kinds of attack [2]. Incidents of fraud can range from direct theft of company funds, to complex cases where company services or data is illegitimately traded for personal financial gain. Kroll Advisory's recent fraud report emphasises the strong link between fraud and insiders, in that, of the companies hit by fraud in the last year, more than 67% identified an insider as a leading perpetrator, signalling yet another increase from previous years' studies [10]. While this is concerning, an even more disturbing aspect looking forward is that according to the *Risk of Insider Fraud* report [2], practitioners continue to believe that their enterprises are at a high risk of insider fraud. This is clearly a serious and prevalent problem in companies today and, as hinted above, financial gain is one of the most common motives.

Another threat that causes great concern is IP theft. In this attack, insiders use their access to steal valuable company data, including trade secrets, business information, source code and customer information [11]. There are several key features of this type of attack. First, the target tends to be product information, proprietary software and source code (these are clear targets in CMU-CERT studies [12]). Also, attacks appear more likely to be conducted by technical personnel (e.g., scientists and engineers) [6] and using technical means (54% of insiders used either email, remote access channel or network file transfer [11]) rather than physical theft of prototypes, for example. Finally, a majority of these thefts are committed by employees with legitimate access to the stolen IP; almost 75% stole material they had authorized access to [12]. Although 75% is a strong statistic and it is therefore very tempting to monitor only these individuals for this attack, yet as other articles have highlighted (e.g., the case of the foreign national who stole Ford secrets worth in excess of $50 million [13]), insiders with no legitimate access are also causing a great deal of harm.

Incidents involving IT sabotage, as one might imagine, tend to be more technically sophisticated. These attacks often require privileged access to systems and networks, or particular knowledge of how they are configured. Examples of specific insider attacks range from insertion of malware (most commonly, logic bombs) to tampering and disrupting system hardware components. Moore *et al.* [14] provide one of the more comprehensive points of reference for data on these types of attack. Amongst their findings, some of the most significant include the high proportion of attackers who had system-administrator privileges (90%) and the crucial role of unmet expectations, disgruntlement and stress in the pathways to an attack (for instance, 92% of all the insiders in their sample attacked enterprises following a negative work-related situation or event). In terms of real-world cases, the attempted attack on Fannie Mae [15] is a perfect example of the sabotage threat. Presumably aggrieved after being dismissed, the insider in this case used the last hours of his legitimate access to upload

malicious code set to auto-execute 7 days later and designed to erase essential company data on finances, securities and mortgages.

In addition to the focus on malicious insiders (covered above), emphasis on benign or accidental insiders has also grown [16]. These individuals have legitimate access to systems, but through carelessness, neglect or accident introduce a form of insider attack. These accidental attacks have become more important to organisations and researchers because, as studies such as the Credant [17] and Clearswift [18] surveys point out, they occur significantly more often than their malicious counterparts. Unwise email activities and loss of storage devices or laptops are some of the most common sources of these breaches. Further analysis on the different types of benign insiders can be found in several reports, particularly the Symantec's *Data Loss Prevention* white paper [19] where the author distinguishes a number of categories of negligent insiders.

### 2.2 The psychology of the insider

Researchers have argued that insiders have specific psychological traits and characteristics. Turner and Gelles [20], for instance, believe the following types of behavioural indicators need to be considered when examining insider risk: self-centredness, arrogance, risk-taking, manipulativeness, coldness, self-deception and defensiveness. Others have suggested that insider threats score high on the personality traits that make up the 'Dark triad': narcissism, Machiavellianism and psychopathy [11, 12, 14, 20]. The UK's Centre for the Protection of National Infrastructure (CPNI) have identified a number of other personality characteristics they believe are typical of an insider, including: immaturity, low self-esteem, amoral and unethical perspective, superficiality, proneness to fantasy, restlessness and impulsivity, and lack of conscientiousness [21].

If it is indeed the case that insider threats possess specific psychological traits and characteristics, then it might aid detection if employers were able to be privy to their employees' psychological make-ups. However, there is also the possibility that specific personality characteristics are linked to specific attacks rather than all attacks. For example, an insider who scores high on narcissism and Machiavellianism and is a risk taker might be more likely to commit IP theft but less likely to deface Web sites. Moreover, psychological characteristics on their own are clearly not enough to predict that someone is likely to become a malicious insider, and also that there are other personal attributes that should also be considered.

It has been argued that shorter-term psychological or emotional states can also help identify the type of individual who is more likely to attack their organisation. Such psychological states might include stress, depression or anxiety, for instance. It has been theorised, for example, that those under extreme stress are more likely to become threats [11, 20]. It might be that the insider instigates the attack to help alleviate the stress that they are encountering. It is argued here, however, that consideration of psychological states in isolation is not sufficient. As is often the case, an external event can trigger a psychological state. Take the case of a person who has experienced financial hardship – such an event

may well cause extreme stress; however, in addition, the individual might see an opportunity at work to conduct fraudulent activities which will help them out of their problems. In contrast, someone who is under extreme stress because of marital problems (exhibiting the same behaviours as in the previous case) might be far less likely to conduct fraudulent activities. These examples illustrate the importance of developing a more holistic model on insider-threat psychology.

In addition to external events, psychological disorders have been reported to make some employees more of a risk to an organisation. CPNI have found that those with a gambling or drug addiction are more likely to attack an organisation than those without such addictions [21]. Of course, if an individual is identified as having such a problem, then an organisation might find ways to provide support for that individual, which in turn might reduce the risk they pose.

In considering the psychology of the insider we might want also to consider their attitude towards the workplace. For example, a person who scores high on the dark triad traits and is highly stressed might be less likely to attack an organisation if they have a strong affinity to their workplace. CPNI have found that those who do not follow established procedures, or read or follow announcements and instructions issued by their organisation, are more likely to attack an organisation [21]. Others have identified the 'disgruntled employee' as a real potential risk [22]; that is, someone who believes they have not been fairly treated by their organisation (e.g., missing out on a promotion). Our belief is that those who have a strong identification with their workplace, and then experience an event which leads them to disgruntlement, pose a greater risk. Whilst our preliminary findings have identified important psychological factors in the context of insider-threat, it becomes quite apparent that there is much more work to be done in this space, by considering a more complete view of the attributes that are associated with identifying potential insider-threats.

### 2.3 The impact of new technologies

As new technologies evolve within organisations, so does the potential insider-attack surface [3, 18]. Bring Your Own Device (BYOD) is becoming increasingly popular within many organisations, and yet in the survey by Ponemon [2], almost half of the 700 participants state that BYOD has resulted in a significant increase in fraud risk. The same study also reports significant challenges in securing corporate data and networks that are now being accessed through this growing gamut of personal devices. There is a definite trade-off being experienced between the convenience and cost-savings of BYOD, as against the security implications and attack vectors that this also introduces, which organisations will need to consider carefully in the future. Cloud services also introduce difficulties regarding security of information. Credant expands on the risks associated with the cloud, and highlight that although this distributed approach has benefits, it translates into a direct loss of control for the business [17]. This introduces yet another possible attack vector, and could also be exploited as part of an attack by existing employees or by the third parties involved. Again, this raises

the trade-off of convenience and cost-savings against maintaining and managing both data and security from within the walls of the organisation.

Social-media use is also generating complex new challenges for enterprises [8, 23]. Through sites such as Facebook, Twitter, blogs and forums, sensitive information (e.g., trade secrets, organisation plans and IP) can be leaked much more easily than before and publicised to anyone, anywhere in the world. The literature is full of cases of this happening, and its affect on both private and governmental organisations [24, 25]. Malicious or careless insiders are not the only concern either. As a result of the amount of information freely shared on these sites, external entities can now exploit social media to identify, target or recruit prospective insider threats [8]. As social media continue to expand in popularity, organisations appear to underestimate the power and reach that they can have. However, the ethical and legal concerns about monitoring personal communications, and whether this is a breach of privacy, remain to be resolved.

## 3 Insider threat from the organisational perspective

From the previous section, it is clear that the threat from insiders is real and significant. Despite this fact, however, reports suggest that corporations continue to underestimate the associated risks, as especially evidenced by minimal investment. For example, the findings in the *State of Security* report [7] show that many companies allocate between 11-14% of their annual revenue to their total IT budget, and of this, they spend 10-14% on security-related issues in general. Investment in detecting and preventing insider threats is therefore likely to be much lower. Of course, the appropriate amount to invest must be determined contingently, by individual companies, depending on their circumstances. But there is evidence of general underinvestment in mitigating this risk at the board level. Another article [8] reports that 25% of respondents stated that there was no regular formal review of cybercrime threats by the Chief Executive Officer and the Board. This suggests that security in some corporations still has not reached the level of importance that it warrants, and again, this obviously has knock-on effects for any hope of adequately managing the risk of insider threat.

More specifically, Ponemon's survey concludes that a large number of companies are not attributing the appropriate priority to the risk of insider fraud, while also noting that it is becoming more of a challenge [2]. One of their main observations as it pertains to organisations' views on risk is that, although 61% of respondents rated the threat of insider fraud within their enterprise as very high or high, only 44% believed that their company viewed the prevention of insider threats as a top priority in security. This highlights that even though organisations view themselves as somewhat unprepared, there does not appear to be an overwhelming impetus to address the risks. These findings mirror those in earlier studies such as McAfee's report [7], where 68% of companies recognise insider threat in their security plans but only 48% have actually addressed it.

Another indication that companies may be underestimating insider threat is the lack of awareness demonstrated by employees and the dearth of training

programmes offered. In one report [23], it was found that 42% of large companies surveyed do not conduct on-going security awareness training sessions with staff and, worse yet, 10% fail to brief staff on induction. This trend of poor awareness in organisations can also be seen more globally, as highlighted in the *Global State of Information Security* survey [3]. The issue here is that due to a lack of training, personnel may be unaware of new risks that insider crimes may present to the company or, indeed, may have forgotten about the risks they used to be aware of. Due diligence is also a particularly salient point, as we continue to see evidence (e.g., [1]) of a considerable number of companies not conducting personnel background checks on their employees.

Companies' views on insider risk can also be understood from how they treat them once detected. The first aspect to note is that they are typically under-reported [8, 26]. In Kaspersky's article [26], for instance, respondents reported that in 59% of the cases nobody outside the company was notified. PwC's survey [8] supports this point, but also found that for very serious fraud offences, some only issued a warning (18% of respondents) and, in a few incidents, organisations did nothing at all (4% of the cases). While we might assume that failure to report incidents is linked to the fear of negative publicity, it is unclear why, even in the case of serious insider incidents, stricter measures are not undertaken. This might further emphasise an underestimation of the problem within corporate culture, but could equally be due to a dearth of solid evidence.

## 4 Detecting insider threats

As the problem of insider threat continues to escalate, there is a growing focus on how to detect such attacks. Here, we explore the current techniques for detection, and where state-of-the-art research is moving towards in the future.

### 4.1 Techniques in use

A variety of approaches have been proposed to mitigate the risk of insider attacks, focusing on prevention, detection and response. Best practices from CMU-CERT include: considering threats from insiders and business partners in enterprise-wide risk assessment; logging, monitoring, and auditing employee's online actions; anticipating and managing negative workplace issues; and developing insider incident-response plans [6]. While a number of these are in common use, the Malicious Insider Threats report notes that many more could be adopted [1]. As discussed in Section 3, what is required is improved education and awareness within enterprise, to encourage active use of such practices.

A key point that arises from published sources (e.g., [12]) is that many attacks are detected by non-technical means (e.g., co-workers noticing suspicious behaviour). Kaspersky's survey article on insiders also identifies reporting by co-workers as the main detection resource as well (indicated in 47% of cases), but also notes the contribution of IT staff in discovering irregularities in system activity logs (41% of cases) [26]. PwC's cybercrime survey identifies three approaches

that organisations use to detect threats: corporate controls (e.g., suspicious-transaction monitoring), corporate culture (e.g., whistle-blowing systems), and those beyond the influence of management (e.g., discovering by accident or a third-party) [8]. They found that the effectiveness of corporate-culture methods has declined compared to previous years. From the detection methods reported, the only noteworthy increase in effectiveness compared with previous years was in automated suspicious-transaction monitoring (up from 0% in 2005 to 18% in 2011). It was observed, however, that whistle-blowing and tip-offs are still an important part of detection, contributing to suspicious behaviour being reported rather than overlooked. This does not stop at employees alone, since reports of suspicious behaviour may come from law enforcement, business partners, and even from customers [12, 26].

Activity logs are becoming more widely used for detecting suspicious activity conducted on organisations' systems [26]. These can provide detail on a range of activities that employees conduct, from entering buildings and logging-on to systems, through to the e-mail communications that they make and the files that they access on a data server. This mass of data provides a wealth of information on employee usage patterns, including any potentially malicious activity that they may choose to carry out. However, due to the large amount of data that can potentially be logged, actually analysing this can quickly become a laborious and error-prone task. There is growing interest around the notion of automated detection of insider threat, and more recently there have been commercial software tools such as SpectorSoft's Spector360, SureView by Raytheon, and DarkTrace. The *Risk of Insider Fraud* report emphasises this desire for automated tools for detecting and analysing insider risk [2].

Many anomaly-based approaches [27, 28] aim to establish what an employee's normal activity may look like, and then analyse how their current behaviour differs from this normal. This opens up a number of challenges, such as how to establish what is actually normal behaviour within an organisation, particularly given that there may already be malicious activity present, and how much of a deviation causes an employee to be classified as a potential insider threat. All organisations will operate differently, as do all humans, and so there will exist many forms of what is deemed to be normal. Likewise, the routine that employees will perform activities on a daily basis will often vary based on their current workload, their personal life, and their mindset, as well as demands made of them by supervisors and co-workers. An employee may well be asked, or need, to perform activities that are outside of their expected normal in order to fulfil their job, and yet this would be flagged as anomalous behaviour. For a system to automatically determine whether an employee is posing a threat or not requires very careful management by the system analyst. An excess of false-positives results in a burden of cases that require investigation, and could result in high resentment by employees. On the other hand, a false-negative would render such a system a failure and could allow the organisation to be severely damaged. It is clear then, that there are many challenges still left to overcome in terms of both detecting, and also analysing, the threat posed by an employee's actions.

### 4.2 State of the art in research

Given the severity of insider threat within many organisations and the strong desire to detect and prevent future attacks, there has naturally been a wealth of research around the problem. Here, we shall examine some of the most notable contributions in the literature and address issues that are currently present.

Brdiczka *et al.* [29] present an approach for proactive detection of insider threats. Their method incorporates structural anomaly-detection, which consists of four stages: graph-structure analysis, graph embedding, dynamic tracking, and anomaly-detection. As they address, this identifies anomalies within the data, not necessarily threats. In order to assess the potential of a threat, they conduct psychological profiling using the Big-5 model, with behavioural, text analysis, and social-networking information as the data used for their profiling. For experimentation, they detect malicious insiders in World of Warcraft data as a proof-of-concept. As acknowledged by the authors, however, in-game malicious behaviour is much more obvious than that of an insider threat in the workplace, who aims to be discrete in their malicious intent. Therefore it would be of great interest to know how the approach copes with more realistic data.

Greitzer *et al.* [30, 31] discuss the use of psychological factors for identifying potential insider threats. They propose a Bayesian Network model that consists of a variety of binary observable behaviours (e.g., engagement, accepting criticism, confrontation, performance, stress, absenteeism). Each behaviour has a prior probability that estimates how frequently it occurs, and a weighting term that specifies how significant the behaviour is with regard to monitoring threats. They derive conditional probabilities through a training process, using expert judgement to assess the threat that an employee exhibits based on particular parameters being set to true. Due to the qualitative nature of the behaviours that are modelled, there remains a need for a human observer to assess whether the employee in question is exhibiting such characteristics. The authors note that future work is necessary to develop methods for automatically extracting and inferring psychological factors from employee-data analysis, rather than using subjective behavioural assessment, which is clearly a non-trivial task to achieve.

Kandias *et al.* [32] also present a prediction model that consists of psychological profiling and real-time usage profiling. These two aspects serve as input to a decision manager that determines whether the user is a potential threat, based on scoring their motive, opportunity and capability. Each user is categorized by their system role, their capability, their predisposition and their stress level. The psychological profiling is conducted by questionnaires that cover user sophistication, predisposition and stress level, whilst the usage profiling consists of monitoring system calls, intrusion-detection systems, and honeypots. The authors state that their future work will focus on the implementation of the model, and so there is currently no indication of how well this performs. The use of questionnaires for psychological assessment raises issues such as the accuracy of the answers provided by participants. In addition, a sophisticated insider may well be capable of circumventing traditional monitoring tools as part of their attack.

As we have seen, there are many proposals for managing insider threat. These approaches draw on a wide range of tasks, such as monitoring, detection, prevention, and prediction. Yet still the insider-threat problem persists. One reason for this is the difficulty of implementing such approaches in real-world environments. Proposals that rely on psychological profiling, for instance, may require compliance from the insider at some stage (e.g., accurate completion of questionnaires). Similarly, gathering data on psychological and behavioural factors within a workplace is a challenging task, as it also requires the attention and compliance of other employees (e.g., reporting suspicious behaviour), while also appreciating the related legal and ethical considerations with such monitoring.

Regarding the development of prototype detection systems, the lack of realistic testing data representing the activities monitored still remains a difficult hurdle to overcome. There has been work on the development of synthetic-data generation, such as that by CMU-CERT [33], where malicious-insider threat data is inserted within normal employee-monitoring data. However, they acknowledge that even these datasets lack the noise and variation that would be present in any real-world data. Undoubtedly, however, and as stressed in [1], there is certainly more that could be done by organisations in order to help support and develop the research surrounding insider threats. Previously, we have proposed a conceptual model for insider-threat detection [34]. As part of our on-going research, we have developed an initial system that is capable of reasoning about the threat posed by an individual, based on their observed activities in the technical domain, whilst also incorporating behavioural analysis and psychological assessment. Whilst the system performs well in preliminary experimentation, we are currently at the stage of requiring more complete data, either synthetic or real-world, in order to truly evaluate its effectiveness.

## 5 Conclusions

Our research in the CITD project recognises the multi-disciplinary nature of insider threat, covering research into the psychological and behavioural aspects that motivate an individual, development of detection systems and analysis tools, and education and awareness-raising within organisations. As a means to detect, prevent, and deter insider threat, the collaboration between these developments is fundamental for addressing the problem effectively. What is clearly apparent, though, is that the insider-threat problem is evident in all types of organisations, can originate in a variety of individuals, ranging from low-level employees through to high-ranking business partners, and can escalate into an attack in many different ways. In this paper, we provide a study on the problem, with the intention of allowing for a better understanding of the nature of insider threats, industry views on the risks faced, and prevention and detection techniques in practice and research. With this critical reflection on current findings and developments, we believe that this serves as an important stage in understanding the ever-persistent and ever-evolving threats that are increasingly occurring within organisations of today.

## Acknowledgements

This research was conducted in the context of a collaborative project on Corporate Insider Threat Detection, sponsored by the UK National Cyber Security Programme in conjunction with the Centre for the Protection of National Infrastructure, whose support is gratefully acknowledged. The project brings together three departments of the University of Oxford, the University of Leicester and Cardiff University.

## References

1. Computer Economics: Malicious insider threats. http://www.computereconomics.com/page.cfm?name=Insider_Threats (2010)
2. Ponemon Institute and Attachmate Corporation: The risk of insider fraud second annual study: Executive summary. http://www.attachmate.com/resources/analyst-papers/bridge-ponemon-insider-fraud-survey.htm (2013)
3. PricewaterhouseCoopers: The global state of information security® 2014. http://www.pwc.com/gx/en/consulting-services/information-security-survey/index.jhtml (2013)
4. PricewaterhouseCoopers: US state of cybercrime survey. http://www.pwc.com/us/en/increasing-it-effectiveness/publications/us-state-of-cybercrime.jhtml (2013)
5. Whitty, M., Wright, G.: Deliverable 3.1 Short report of findings from Case Studies (Corporate Insider Threat Detection project), Leicester University Report (2013)
6. Cappelli, D.M., Moore, A.P., Trzeciak, R.F.: The CERT Guide to Insider Threats. Addison-Wesley (2012)
7. McAfee and Evalueserve: State of security. http://www.mcafee.com/us/resources/white-papers/wp-state-of-security.pdf (2011)
8. PricewaterhouseCoopers: Cybercrime: Protecting against the growing threat. http://www.pwc.tw/en/publications/events-and-trends/e256.jhtml (2012)
9. 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) (2011) 4–27
10. Kroll Advisory Solutions and Economist Intelligence Unit: The global fraud report 2012/13. http://www.kroll.com/library/KRL_FraudReport2012-13.pdf (2012)
11. Shaw, E.D., Stock, H.V.: Behavioral risk indicators of malicious insider theft of intellectual property: Misreading the writing on the wall (Symantec Report) (2011)
12. Moore, A.P., Cappelli, D.M., Caron, T.C., Shaw, E., Spooner, D., Trzeciak, R.F.: A preliminary model of insider theft of intellectual property. Technical report, CMU-CERT (2011)
13. Kaspersky: Threatpost series: Insider threats. http://usa.kaspersky.com/resources/knowledge-center/threatpost (2011)
14. Moore, A.P., Cappelli, D.M., Trzeciak, R.F.: The "big picture" of insider IT sabotage across U.S. critical infrastructures. Technical report, CMU-CERT (2008)
15. FBI: Fannie Mae corporate intruder sentenced to over three years in prison for attempting to wipe out fannie mae financial data. http://www.fbi.gov/baltimore/press-releases/2010/ba121710.htm (2010)

16. Booz Allen: The accidental insider threat: Is your organization ready? (expert voices panel). http://www.boozallen.com/media/file/Accidental-Insider-Threat-Panel-Discussion-Transcript.pdf (2012)
17. Credant: Insider threat. http://go.credant.com/campaigns-insider (2011)
18. Clearswift: The enemy within: an emerging threat... http://www.clearswift.com/blog/2013/05/02/enemy-within-emerging-threat (2013)
19. Wall, D.S.: Organizational security and the insider threat: Malicious, negligent and well-meaning insiders. Technical report, Symantec (2011)
20. Turner, J.T., Gelles, M.: Threat assessment: A risk management approach. Routledge (2003)
21. CPNI: CPNI insider data collection study – report of main findings. http://www.cpni.gov.uk/Documents/Publications/2013/2013003-insider_data_collection_study.pdf (2013)
22. Holton, C.: Identifying disgruntled employee systems fraud risk through text mining: A simple solution for a multi-billion dollar problem. Decision Support Systems **46**(4) (2009) 853–864
23. The Department for Business, Innovation and Skills (BIS) & PricewaterhouseCoopers: 2013 Information security breaches survey (2013)
24. Sky News: MoD secrets leaked onto the Internet. http://news.sky.com/story/753966/mod-secrets-leaked-onto-the-internet (2010)
25. Harrysson, M., Metayer, E., Sarrazin, H.: How not to unwittingly reveal company secrets (Harvard Business Review blog network). http://blogs.hbr.org/2012/12/how-not-to-unwittingly-reveal/ (2012)
26. Kaspersky: Threatpost's insider threats survey. http://usa.kaspersky.com/resources/knowledge-center/threatpost (2011)
27. Patcha, A., Park, J.M.: An overview of anomaly detection techniques: Existing solutions and latest technological trends. Computer Networks **51**(12) (2007) 3448–3470
28. Salem, M., Hershkop, S., Stolfo, S.: A survey of insider attack detection research. In Stolfo, S., Bellovin, S., Keromytis, A., Hershkop, S., Smith, S., Sinclair, S., eds.: Insider Attack and Cyber Security. Volume 39 of Advances in Information Security. Springer US (2008) 69–90
29. Brdiczka, O., Liu, J., Price, B., Shen, J., Patil, A., Chow, R., Bart, E., Ducheneaut, N.: Proactive insider threat detection through graph learning and psychological context. In: IEEE Symposium on Security and Privacy Workshops. (2012)
30. Greitzer, F.L., Hohimer, R.E.: Modeling human behavior to anticipate insider attacks. Journal of Strategic Security **4**(2) (2011) 25–48
31. Greitzer, F.L., Kangas, L.J., Noonan, C.F., Dalton, A.C., Hohimer, R.E.: Identifying at-risk employees: Modeling psychosocial precursors of potential insider threats. In: 45th Hawaii International Conference on System Science, IEEE (2012)
32. Kandias, M., Mylonas, A., Virvilis, N., Theoharidou, M., Gritzalis, D.: An insider threat prediction model. In: Trust, Privacy and Security in Digital Business. Springer (2010) 26–37
33. Glasser, J., Lindauer, B.: Bridging the gap: A pragmatic approach to generating insider threat data. IEEE Symposium on Security and Privacy Workshops (2013)
34. Legg, P., Moffat, N., Nurse, J.R.C., Happa, J., Agrafiotis, I., Goldsmith, M., Creese, S.: Towards a conceptual model and reasoning structure for insider threat detection. Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications **4**(4) (2013) 20–37