Reflecting on the Ability of Enterprise Security Policy to Address Accidental Insider Threat

Oliver Buckley, Jason R. C. Nurse, Philip A. Legg, Michael Goldsmith, Sadie Creese
Cyber Security Centre, Department of Computer Science, University of Oxford, UK
{*firstname.lastname*}*@cs.ox.ac.uk*

***Abstract*—An enterprise's information security policy is an exceptionally important control as it provides the employees of an organisation with details of what is expected of them, and what they can expect from the organisation's security teams, as well as informing the culture within that organisation. The threat from accidental insiders is a reality across all enterprises and can be extremely damaging to the systems, data and reputation of an organisation. Recent industry reports and academic literature underline the fact that the risk of accidental insider compromise is potentially more pressing than that posed by a malicious insider. In this paper we focus on the ability of enterprises' information security policies to mitigate the accidental insider threat. Specifically we perform an analysis of real-world cases of accidental insider threat to define the key reasons, actions and impacts of these events – captured as a grounded insider threat classification scheme. This scheme is then used to perform a review of a set of organisational security policies to highlight their strengths and weaknesses when considering the prevention of incidents of accidental insider compromise. We present a set of questions that can be used to analyse an existing security policy to help control the risk of the accidental insider threat.**

***Index Terms*—Security policy, risk, insider threat, accidental, unintentional, benign, case studies**

I. Introduction

The security of an organisation has always been considered to be a constantly evolving challenge. Traditionally the emphasis has been placed on ensuring that an enterprise was guarded against external threats, using perimeter defences such as firewalls and layered security mechanisms. However, the current reality is that there is a far greater threat posed from those within the organisation. This is especially true from unwitting employees that through human error, oversight, or the poor design of the security controls, they are expected to use, pose a threat to their enterprise (commonly known as the accidental or unintentional insider threat). This is an issue that has long been recognised as needing to be addressed by the research community, with recent evidence being reported by PricewaterhouseCoopers and the Department for Business Innovation & Skills [1].

A security policy provides an exceptionally important statement of the expectations of the staff within an organisation, and will have direct consequences for the budgets that are available to control risk, including both organisational culture and awareness campaigns. Many researchers consider the need and effectiveness of security awareness campaigns in addressing this risk ([2], [3]). The work we present in this paper is complimentary to those contributions as we examine the effectiveness of enterprise security policies against the risk of accidental insider compromise. We have developed a classification scheme for accidental insider threats grounded in an assessment of 60 documented cases, using our threat characterisation framework [4] to identify the central components. We then use the resulting classification scheme as a basis for assessing the policies. We have used this method to consider 10 publicly available enterprise information security policies, from a cross-section of industries, and 5 information security policy templates.

The remainder of this paper is structured as follows. In Section II we provide a review of the related work covering both accidental insider threat and also the design, analysis and effectiveness of enterprise security policies. Section III details the methodology used to conduct our analysis. Section IV reflects on the effectiveness of organisational security policies when compared to the results of our analysis. Finally in, Section V, we conclude by providing a reflection on our analysis and also a discussion of the further work in this area.

II. Related Work

In this section we present a review of the related literature; we begin with work relating to enterprise security policies, following this we provide a discussion of security policies in relation to human error and non-compliance. Finally, we provide a review of material that focuses on the notion of the accidental (or unintentional) insider threat.

The information security policy is one of the most important controls within an organisation. It provides an enterprise with the basis of a strategy that will define the working culture and the behaviour that is expected of the organisation's employees [5]. Doherty, Anastasakis and Fulford [6] present a study that analyses the structure and content of university security policies. Their paper provides a comparative study of academic security policies, focusing on four key areas: the university details (e.g. name, country, ranking); the policy structure (e.g. policies freely available); the policy administration details (e.g. the date the policy was created); and the policy coverage (e.g. physical security or internet access). The work is focused on security policies from a single sector, in contrast to our own work where real-world cases are used to highlight the relative strengths and weaknesses of security policies across a range of different industries.

Höne and Eloff [7] provide an analysis of the common issues with organisational security policies, as well as suggesting

the key factors that will contribute to the effectiveness of a security policy. The paper suggests that the biggest barrier to the success of an information security policy is that it often fails to adequately communicate with employees for a variety of reasons, including: users are not aware of the policy, it is too long or technical and there is often a disconnect between the policy and day-to-day tasks. Höne and Eloff conclude that the most effective security policies will be focused on the users, where the content is secondary to the way in which it is communicated to the users.

A large number of incidents that could be attributable to an accidental insider are often the result of policies within a company that are either poorly defined or are not properly disseminated and regularly reinforced among the employees. Boss *et al.* [8] looks at the degree in which security policies and procedures are followed by individuals, and introduces the idea of 'mandatoriness'. Mandatoriness is used as a metric for the degree to which individuals believe that compliance with security policies is compulsory. The findings, by Boss *et al.*, show that if an employee believes that they are being observed by their managers, then they are more likely to follow security policies and procedures.

The notion of policy non-compliance is particularly relevant when considering incidents of accidental insider compromise. Pahnila, Siponen and Mahmood [9] provide an empirical study of employees' behaviour in relation to security policy compliance. Their work proposes a model to provide insight into why an employee may or may not comply with a security policy. Herath and Rao [10] present an Integrated Protection Motivation and Deterrence model of security policy compliance, based on literature, protection-motivation theory, deterrence theory and organisational behaviour. Whilst it is recognised that user behaviour plays a part in a number of security incidents, it is too easy to place the blame entirely on the employee. Sasse, Brostoff and Weirich [11] argue that it is counter-productive to place the blame on the users and that security policy needs to be designed with human behaviour in mind to maximise user buy in and to minimise non-compliance.

Another area that is strongly related to our own research is the idea of human error, in particular, human error in the workplace. As such there are a number of relevant papers that cover workplace accidents and human error. Ganguly [12] provides an overview of human failure, including a brief categorisation, along with a description of important factors that influence human behaviour in relation to workplace safety. Human failure is broadly broken down into two categories, intentional and unintentional errors, where unintentional errors are described as 'actions that were not as planned'. Liginlal, Sim and Khansa [13] provide a discussion on privacy breaches that were a direct result of human error. This analysis of privacy breaches and their causes is then used to develop a 'defense-in-depth' solution designed to avoid, intercept and correct errors. Their work looks at developing a taxonomy of incidents that were a result of human error and malicious acts.

In cases of malicious insider threat the motive is usually for reasons of personal gain or revenge [14], whereas in the case of an accidental insider incident there is often no motivation to attack [15]. It could be argued that the motivation of an accidental insider is, in most cases, to carry out their role and as such their motivations are for the most part positive and well intentioned. For example, a reported case in Salt Lake City concerns an employee of Good Data Systems was fired after losing a USB stick containing 6000 medical records, despite being regarded as a 'terrific employee' [16]. In this example the employee was very well thought of, but had made a copy of the medical records to a USB stick, in the course of her job, without realising that this violated company policy.

Work presented by Wall [17] focuses on the issue of the accidental insider and further divides the category into 'negligent' and 'well-meaning' insiders. The negligent insider refers to an insider whose 'eyes are not always on the ball', that is to say, an individual who is perhaps less risk averse and is willing to bend the rules to ensure that things get done. A negligent insider will generally accept the broad organisational goals, but will adhere to policies to achieve the organisational goals only as long as they do not create additional work. The well-meaning insider is typically a valued employee who is dedicated to pursuing the primary goals of the organisation. For a well-meaning insider the active pursuit of the organisation's goals will often mean that some security policies will be a secondary concern. The work in Wall's paper is focused on the human aspects of accidental insider threat, without providing details or classifications on the causes or outcomes of these incidents.

The information security policy is a critically important document for an organisation, which must be well defined and well communicated to give the best chance for it to be properly understood and adhered to by employees within the organisation. However, our reflection on the literature has highlighted that even with a well written and communicated security policy, accidental insider compromise is still a pressing issue. It is this area that we aim to address with the work presented in this paper.

III. Method

In this section we describe the method that we have used to collect and evaluate the effectiveness of enterprise security policies in dealing with the risk of accidental insider compromises. We begin by discussing our method for collecting cases of accidental insider compromise, then detail the security policies that were collected, and finally, the assessment of the cases (drawing on our attack characterisation framework).

A. Accidental Insider Cases

The definition of an insider is both well defined and well understood [14]. This description is refined by Grietzer *et al.* [18] to provide definition of what is meant by an accidental (or unintentional) insider threat. It is this definition that we have used to support the collection of our accidental insider cases.

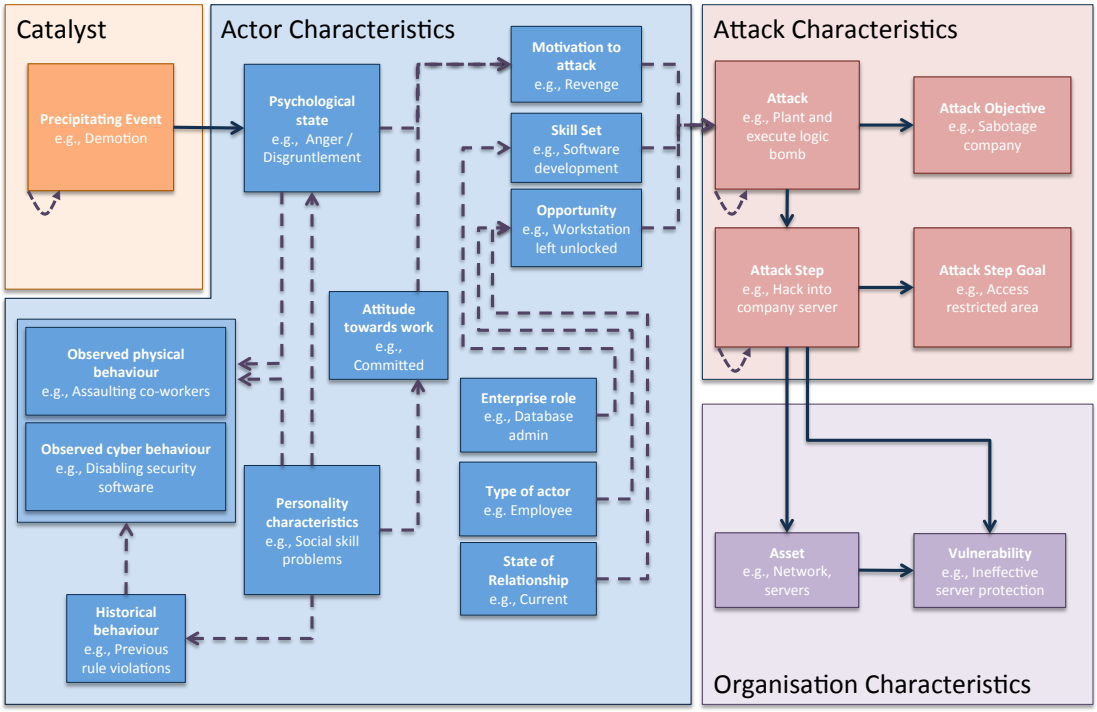

Fig. 1. A framework for capturing insider threats containing several key elements: the 'Catalyst' refers to the overarching reason for the incident, 'Actor Characteristics' capture the state of the actor, 'Attack Characteristics' detail the elements relating to the attack and finally, 'Organisation Characteristics' that includes organisational assets and the vulnerability [4]

We have collected 60 cases of incidents that meet the characterisation of an accidental insider threat from a number of different sources, including: news articles, official reports and other relevant articles. We have not placed a restriction on the scale of the cases that have been collected, that is to say we have identified and collected cases where the size and scope of the incident was not one of our selection criteria. We have collected cases with the aim of providing a cross-section of the types of incident and also across a range the industries that are affected by accidental threat. The cases that we have collected all reflect incidents that have been discovered and reported within the last 10 years.

We have used our previous insider threat characterisation framework [4] to provide a basis for the analysis of the cases collected. The framework has been created to capture all of the key data-points that are associated with cases of insider threat (both malicious and non-malicious). This framework has facilitated the capture of the key fields within an incident, and has helped to highlight the most pertinent factors involved in the cases that we have collected. This has provided us with a consistent dataset across all of our collected cases. An example of the framework can be seen in Figure 1. The characterisation framework was used to code our set of collected cases, using a single coder.

B. Security Policy Collection

We have collected a sample of 10 information security policies that were publicly available on the Internet, across a range of different sectors: academia (3 policies), local government (3 policies), health (1 policy), finance (1 policy), science and technology (1 policy) and law enforcement (1 policy). The security policies have been selected to provide a reflection of the distribution industries in the collected cases of accidental insider threat.

In addition to the real-world policies that have been collected, we have also used template policies, which were publicly available online. The reason for the inclusion of template policies in our analysis was that it is reasonable to expect that these templates would be representative of the policies that would be used by several organisations in reality. Due to a lack of availability, the policies selected are not associated with the organisations represented by our accidental insider cases. All of the information security policies that have been used in our study have been anonymised.

C. Classification of Cases

The framework in Figure 1 is designed to capture both malicious and non-malicious insider incidents. Mapping the collected case, and through information in the wider literature, revealed the need to expand the scope of some of the components within the framework. This was to ensure that all of the relevant information was captured for cases of accidental insider threat. Below we briefly discuss the modifications that were required.

1) Motivation: One of the most obvious changes that was identified was the lack of malicious intent for the incident across all of the cases that were analysed. This is further emphasised by Jones and Ashenden [15] who state that the

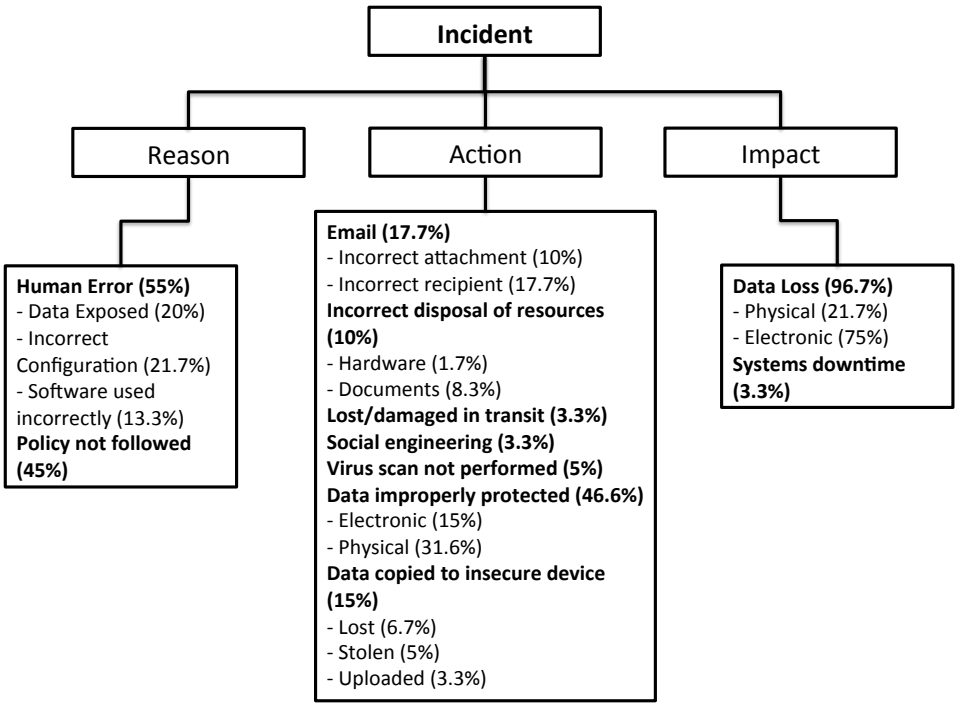

Fig. 2. Break down of the reasons, actions and impacts of accidental insider threat

absence of malice or motivation is one of the key factors that separates accidental insider incidents from intentional acts of insider threat.

2) Precipitating Event: Our previous work [4] describes the precipitating event as the key event that has the potential to cause an insider to become a threat to their employer. This notion of a 'tipping point' does not directly apply when we are considering an accidental insider compromise. The cases studied revealed that there was very rarely a single event that caused an individual to negatively impact their organisation. Magklaras and Furnell [19] suggest reasons for the accidental misuse of computer systems: inadequate system knowledge, factors that can affect work related performance (e.g. excessive workload), and finally a lack of awareness of security training. When considering accidental insider compromise we are more interested in why the employee made a mistake in the first instance and so we extend the precipitating event to consider the overarching reason for the incident, for example leaking sensitive data via email might be due to inadequate training or policy awareness [20].

3) Attack Objective: When considering a malicious insider threat, the attack objective is very clearly defined. For instance, an employee copying sensitive organisational data to pass on to a competitor has the defined objective of leaking the information to a rival organisation. In contrast to the precipitating event, when considering the attack objective for the incident we are interested in what the insider was attempting to achieve when the incident occurred (e.g. a user attempting to upload information to a secure file server [21]). This classification of collected cases provides the basis for a comparison of our findings to a number of widely available information security policies. During the classification process it became apparent that there were a number of reoccurring themes, across all of the cases of accidental insider threat. Figure 2 illustrates the distillation of the collected cases, that will be used to analyse the coverage and utility existing security policies, broken down into three major categories:

- Reason – What was the cause of the incident?
- Action – What was the thing that was done incorrectly?
- Impact – What was the impact of the incident?

The categorisation, listed above, provides a focus of the information collected that directly relates to the kinds of clauses that are seen in information security policies. For example, we consider the impact of physical and electronic data loss as an organisation will often have separate sections of the policy to control the management of data held physically or electronically.

IV. Analysis of Security Policies

The results in Table I provide an overview of the coverage of each of the security policies when compared against the reasons and actions of accidental insider threat, as seen in our case analysis. Here, the Reason is the cause of incident (e.g. incorrect use of software) and the Action is the thing that was done incorrectly (e.g. an email was sent to the wrong recipient). In Table I the policy numbers correspond to the following industries: 1–3 Academia, 4–6 Local government, 7 Medical, 8 Finance, 9 Science and technology, 10 Law enforcement and 11–15 Template policies. A green tick indicates

	Policy														
Reason	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**
Human Error (55%)	✗	✗	✓	✓	✓	✓	✗	✓	✓	✓	✗	✓	✓	✓	✓
- Data exposed (20%)	✗	✗	✗	✗	✓	✓	✗	✓	✓	✓	✗	✓	✓	✓	✓
- Incorrect Configuration (21.7%)	✗	✗	✓	✗	✗	✗	✗	✓	✓	✗	✗	✓	✓	✗	✓
- Software used incorrectly (13.3%)	✗	✗	✓	✗	✓	✗	✗	✗	✗	✗	✗	✓	✓	✗	✗
Policy not followed (45%)	✗	✓	✓	✓	✓	✓	✓	✗	✓	✗	✗	✗	✓	✓	✗
Action															
Email (17.7%)	✓	✓	✓	✗	✓	✓	✗	✓	✗	✓	✓	✗	✓	✗	✗
- Incorrect Attachment (10%)	✓	✗	✗	✗	✗	✗	✗	✓	✗	✗	✓	✗	✗	✗	✗
- Incorrect Recipient (17.7%)	✗	✗	✓	✗	✓	✗	✗	✓	✗	✗	✓	✗	✗	✗	✗
Incorrect disposal of resources (10%)	✗	✓	✓	✓	✓	✓	✗	✗	✓	✓	✓	✓	✓	✓	✓
- Hardware (1.7%)	✗	✗	✓	✓	✓	✓	✗	✗	✓	✓	✓	✓	✓	✓	✓
- Documents (8.3%)	✗	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✓	✓	✗	✗
Lost/damaged in transit (3.3%)	✓	✗	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
Social Engineering (3.3%)	✓	✗	✓	✓	✗	✗	✗	✗	✗	✗	✓	✗	✓	✗	✗
Virus/malware scanning (6%)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Data improperly protected (46.6%)	✗	✓	✓	✗	✓	✗	✗	✓	✓	✓	✓	✓	✓	✗	✓
- Electronic (15%)	✗	✓	✓	✗	✓	✗	✗	✓	✓	✓	✓	✓	✗	✗	✗
- Physical (31.6%)	✗	✓	✓	✗	✓	✗	✗	✗	✓	✓	✓	✓	✓	✗	✓
Data copied to insecure device (15%)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✓	✓	✓	✗
- Lost (6.7%)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✓	✓	✓	✗
- Stolen (5%)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✓	✓	✓	✗
- Uploaded (3.3%)	✗	✓	✗	✗	✗	✓	✗	✗	✓	✓	✗	✓	✓	✗	✗

TABLE I
THE COVERAGE OF SECURITY POLICIES WHEN COMPARED AGAINST THE REASONS AND ACTIONS OF ACCIDENTAL INSIDER THREAT

that the policy contained clauses to mitigate against the reason or action listed, and a red cross denotes the lack of coverage within a policy.

The remainder of this section will provide a detailed look at our policy analysis, including discussion of relevant clauses from the security policies collected.

A. Human Error

80% of the security policies analysed contained clauses that considered the risks of incidents where human error was the main contributing factor, and provided some form of controls in place to protect against it. For example, one of the academic policies surveyed contains a clause that states controls should be applied to protect against "the vulnerability to error, e.g. incorrect addressing or misdirection, and the general reliability and availability of the service". This is in contrast to controls on the prevention of the incorrect use of software within an organisation, where 80% of the policies reviewed did not mandate that any kind of software training be undertaken. However, this did appear more frequently in the template security policies than in the enterprise policies. For instance, one of the policies taken from local government requires that, in order to ensure the integrity of organisational data staff should have "received training on any application that they would be required to access and any software package they will be required to use". However, none of the policies document how a judgement of suitability or effectiveness of such training would be made.

B. Policy not followed

When considering 'Policy not followed' we are focusing on whether there is any stipulation that employees must receive information security training. Our case analysis highlighted two main reasons for policy non-compliance: the policy was incomplete or poorly defined, or the employee was not aware of the security policy. The wider literature suggests other potential reasons why this may occur; for example, the policy was not well communicated or that the policy made it more difficult for the employee to carry out their role [22].

C. Email

While the majority of the policies surveyed (60%) contained some consideration for the correct use of email, this often focused on how email should be used (e.g. "personal use of e-mail is permitted provided such use is only during free time and is not of significant volume"). Very little consideration is given to protecting against the incorrect use of email either by attaching a file erroneously (20%) or by sending to the incorrect recipient (27%).

D. Incorrect disposal of resources

The majority of the policies reviewed (73%) contained guidance about the safe disposal of resources, however, there were marginally fewer policies that contained information about the safe disposal of hardware (67%) compared to the safe disposal of paper-based resources (74%). Typically the safe disposal of paper-based resources was as simple as placing the documents into the correct place, according to their classification (e.g. "restricted and confidential material can be placed in confidential waste bins, whereas secret material must be shredded"). Conversely, when considering the safe disposal of hardware, the policies were often less prescriptive. For example, one of the academic security policies collected requires that employees must ensure that data "is thoroughly and securely cleansed from that equipment when they leave". This stipulation is placing the emphasis on the employee to

ensure that sensitive data is safely removed from any hardware, which can be viewed as an inappropriate apportioning of responsibility. It is unlikely that a general employee will be well informed as to what thorough and secure cleansing is.

E. Transporting data

The term 'transporting data' is used to describe the physical process of transporting data (including tapes, discs, USB devices and paper-based data) to different areas of an organisation (e.g. to another site) or to a partner organisation. The safe transport of data was something that was, for the most part, not well covered within the security policies reviewed. Only 13% of the policies contained clauses to control this action. For example, an academic policy provided protection against this by stipulating that "reliable transport or couriers should be used". There were a number of incidents recorded where the items being transported were insufficiently packaged and as such were damaged in transit [23]. Only one of the policies that we reviewed considered this aspect of transporting data with the requirement that "packaging should be sufficient to protect the contents against any physical damage likely to arise during transit".

F. Social engineering

When comparing this type of incident against the collected set of information security policies we were analysing the policies for clauses that provided guidance on the trustworthiness of information. For the most part this is something that was not well covered in the policies surveyed, with only 33% containing guidance on preventing social engineering attacks, however, two of the three policies from academia contained information on this issue. The most common information provided on social engineering was focused on the trustworthiness of an information received over email (e.g. 'you should not necessarily trust what you receive in an email – in particular, you must never respond to an email request to give a username or password').

G. Virus/malware protection

The protection against malware was one of the areas that was covered by all of the policies in our sample. The controls for protecting against malware were well prescribed, with all of the policies requiring "appropriate anti-virus software is installed and maintained". In some instances this was something that would be handled by the organisation's IT department, in others there was a list of acceptable providers of anti-virus software, and finally some organisational policies only required the software be installed and maintained with no mandate on what should be installed. In these cases, the emphasis was placed onto the employee to maintain an integral security control, which can be considered an unreasonable expectation to place on an employee, as they are not likely to be familiar with the best approach for preventing against this kind of issue. Indeed it is entirely possible that an employee could download malware posing as anti-virus software.

H. Data improperly protected

When considering whether data has been improperly protected, we are referring to whether data has been erroneously exposed. Many of the security policies (73%) analysed contained clauses that considered the protection of sensitive data. The most simplistic of these clauses typically referred to ensuring that the organisation, and its employees, adhere to the Data Protection Act [24], for example, when referring to the confidentiality of sensitive data employees are tasked with "ensuring compliance with the Data Protection Act".

Other policies were more detailed about the controls and measures that should be used to ensure that data (both physical and electronic) was adequately protected. The policies provided guidance for protecting data in a number of ways. The accidental insider cases that were analysed highlighted that the incorrect configuration of a web server was often to blame for data that was wrongfully displayed online. A little over half (53%) of the policies analysed mitigated against this kind of loss, for example, an academic security policy required that the issues relating to "security and data protection implications of publishing directories entries" be considered.

I. Data copied to insecure device

The protection against data loss, by copying data to an insecure device, was one of the issues that was covered by all of the security policies that were reviewed. There were three approaches that the policies took when considering the use of removable devices. The first was the most straightforward, and prohibitive in that it was forbidden for data to be copied to a device and removed from the organisation. The second approach was to require that all removable devices, which contained sensitive data, to be encrypted before they were removed from the organisation. Finally, the third approach was to mandate that any removable devices, laptops or other hardware be secured or hidden from view when removed from the organisation and not in use.

J. Summary

The analysis presented above highlights the fact that there was a great deal of variability across all of the policies in our sample. However, there was a degree of commonality for certain areas within the policies analysed. For example, all of the policies surveyed contained guidance on the prevention of malware infection and also on safe practices for copying data to a removable device. Conversely, whilst there were no areas that were entirely ignored by all of the policies there were certainly some areas that were very sparsely represented. For instance, very little consideration was given to mitigating against data being lost or damaged in transit. This is also true when considering the correct use of email, for instance; when considering the issue of sending an email to the wrong recipient or erroneously attaching a file to an email.

The results in Table I illustrate that, for the most part, the policies included in our analysis provide better coverage of the 'Actions' than the 'Reasons' involved in an accidental insider incident. Two of the security policies surveyed failed

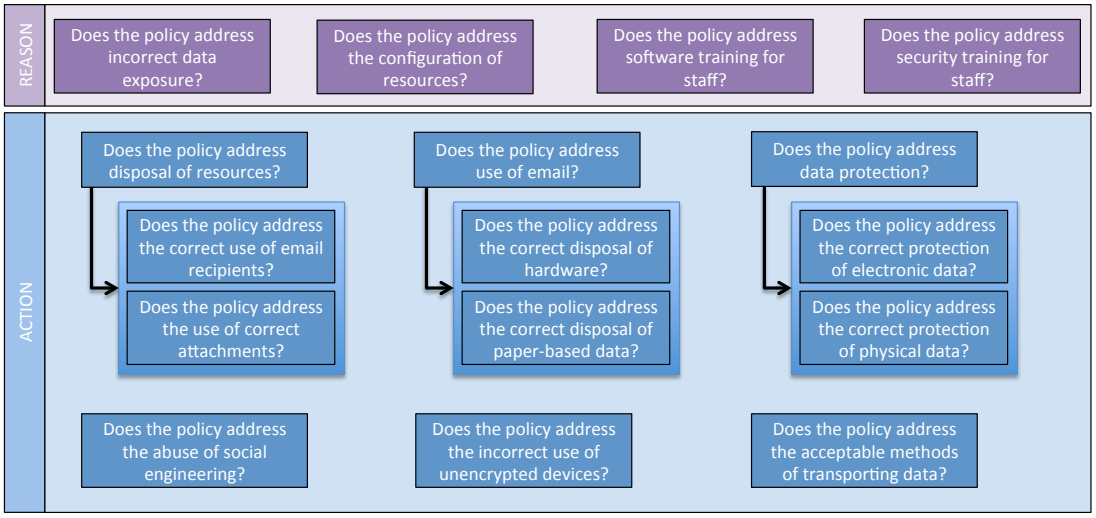

Fig. 3. Questions that can be used to assess the coverage of a security policy against an accidental insider threat

to provide any clauses that would prevent against the reasons for cases of accidental insider compromise, which we have identified. The reasons that we discovered are largely related to employee education, for example, providing training that covers the organisation's information security policy or how to correctly use various software packages. The actions were, generally, associated with the control of technology or physical processes and as such are arguably easier to mandate within a security policy.

It is perhaps worth noting that in our sample set of security policies, there was not a policy that provided complete coverage, against all of the potential accidental insider threats that our case analysis revealed. Even the policy with the best coverage (policy 3 in Table I), only provided guidelines that would help to prevent against 80% of the reasons and 86.7% of the actions associated with accidental insider threat.

Below we provide an analysis of the coverage of the policies when compared to the accidental insider cases collected, based on the policy's industry sector. This analysis assumes that policies were perfectly implemented, which is of course not the case in reality. Policies 1–3 are all from the academic sector and even when combined there is still a gap in the policies with respect to cases of human error, where data was exposed, which amounts to 20% of the cases studied. The policies collected from local government (policies 4–6) contained gaps in the reasons for incidents in the form of incorrect configuration as a result of human error, which amounts to 21.7% of cases. The actions do not provide any mitigation against a file erroneously attached to an email (10%) or data that is lost or damaged in transit (3.3%). The policy collected from the medical sector only provides guidance on cases where the reason was lack of policy awareness, which meant a vulnerability to 55% of the cases surveyed. There was also only consideration given to a limited number of actions (malware protection and copying to data to an insecure device), which meant the policy was vulnerable to 83.3% of cases. Policy 8 (financial sector) lacked information to protect against 58.8% of the reasons for accidental insider incidents, as well as being open to 51.5% of the cases surveyed. The policy from the science and technology sector lacked clauses that accounted for 65% of the reasons for the accidental insider incidents, but provided better coverage against the actions involved in cases with only 32.6% not covered. Finally, policy 10 (law enforcement) provided information that could potentially have protected against 20% of the reasons for the cases studied, but was potentially vulnerable to only 24.3% of the cases surveyed.

V. Conclusions and Future Work

An organisation's information security policy is a very important statement of the expectations of its employees, and the acceptable behaviour and culture of that enterprise. The threat of accidental insider compromise are a real concern to all organisations and carry a significant threat to the systems, data and reputation of an organisation, perhaps more so that those carried out by a malicious insider.

This paper has highlighted the common areas, within organisational information security policies, that were lacking when considering a policy's ability to facilitate the prevention of accidental insider incidents. We surveyed 15 instances of information security policy to determine the key areas of coverage with respect to accidental insider threat. It was apparent that some incident types were more commonly handled than others by our sample set of security policies. For example, clauses were seen in all of the policies surveyed to protect against malware infection and also to prevent the copying of data to an insecure device. It could be argued that these are two controls that are technically straightforward to implement and monitor, whereas a large number of the other potential threats are more complex to manage with technological controls alone. For example, it would be a non-trivial task to ensure that emails only contained the correct attachments or recipients.

The work presented in this paper is used to establish the risk posed by the accidental insider within organisation, and the degree to which enterprise security policies help to mitigate this risk. We propose a specialisation of an existing framework to capture pertinent data points which can then be reduced down to a set of causes, attack vectors, and impacts to an organisation. Currently we have used this model and information to provide a review of publicly accessible security policies, to highlight their strengths and weaknesses when it comes to handling this category of threat. In our future work we will consider a larger sample of cases of accidental insider threat and a wider sample of information security policies to determine whether our model can be used to identify policy refinements and risk control options that need to be addressed.

Research has shown that accidental insider issues are a significant problem, and our research emphasises the need for a strong policy when considering the mitigation of these incidents. In the future we would look at the ways in which a security policy can be designed to directly address the accidental insider threat.

As a result of our policy analysis we were able to create a set of questions that can be asked of a security policy to help perform a self-assessment of the policy coverage, with respect to accidental insider incidents. The set of questions can be seen in Figure 3. In future work we will look to further develop this set of questions to provide guidelines with a distinct flow that could be utilised to create or analyse an existing information security policy. Of course a good security policy is only one piece of the puzzle and there are many other factors that contribute to a policy's impact and effectiveness.

We recognise the limitations of our approach in using only publicly available security policies. In future work we would look to expand on our initial study to provide a more detailed analysis of corporate security policies that are not public, and which do relate to accidental insider cases. In addition to this we would look to expand our work to include details of the impact and severity of the cases of accidental insider threat.

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] PricewaterhouseCoopers LLP and Department for Business, Innovation and Skills. (2013) Information security breaches survey. [Online]. Available: https://www.gov.uk/government/publications/information-security-breaches-survey-2013-technical-report

[2] M. E. Thomson and R. von Solms, "Information security awareness: educating your users effectively," *Information management & computer security*, vol. 6, no. 4, pp. 167–173, 1998.

[3] J. D'Arcy, A. Hovav, and D. Galletta, "User awareness of security countermeasures and its impact on information systems misuse: a deterrence approach," *Information Systems Research*, vol. 20, no. 1, pp. 79–98, 2009.

[4] J. R. C. Nurse, O. Buckley, P. A. Legg, M. Goldsmith, S. Creese, G. R. T. Wright, and M. Whitty, "Understanding insider threat: A framework for characterising attacks," in *Security and Privacy Workshops (SPW), 2014 IEEE*. IEEE, 2014.

[5] M. E. Palmer, C. Robinson, J. C. Patilla, and E. P. Moser, "Information security policy framework: best practices for security policy in the e-commerce age," *Information Systems Security*, vol. 10, no. 2, pp. 1–15, 2001.

[6] N. F. Doherty, L. Anastasakis, and H. Fulford, "The information security policy unpacked: A critical study of the content of university policies," *International Journal of Information Management*, vol. 29, no. 6, pp. 449–457, 2009.

[7] K. Höne and J. Eloff, "What makes an effective information security policy?" *Network Security*, vol. 2002, no. 6, pp. 14–16, 2002.

[8] S. R. Boss, L. J. Kirsch, I. Angermeier, R. A. Shingler, and R. W. Boss, "If someone is watching, I'll do what I'm asked: mandatoriness, control, and information security," *European Journal of Information Systems*, vol. 18, no. 2, pp. 151–164, 2009.

[9] S. Pahnila, M. Siponen, and A. Mahmood, "Employees' behavior towards is security policy compliance," in *System Sciences, 2007. HICSS 2007. 40th Annual Hawaii International Conference on*. IEEE, 2007, pp. 156b–156b.

[10] T. Herath and H. R. Rao, "Protection motivation and deterrence: a framework for security policy compliance in organisations," *European Journal of Information Systems*, vol. 18, no. 2, pp. 106–125, 2009.

[11] M. A. Sasse, S. Brostoff, and D. Weirich, "Transforming the weakest linka human/computer interaction approach to usable and effective security," *BT technology journal*, vol. 19, no. 3, pp. 122–131, 2001.

[12] S. Ganguly, "Human error vs. work place management in modern organizations," *International Journal of Research in Management and Technoloy*, vol. 1, no. 1, pp. 13–17, 2011.

[13] D. Liginlal, I. Sim, and L. Khansa, "How significant is human error as a cause of privacy breaches? an empirical study and a framework for error management," *Computers & Security*, vol. 28, no. 3, pp. 215–228, 2009.

[14] D. M. Cappelli, A. P. Moore, and R. F. Trzeciak, *The CERT guide to insider threats: how to prevent, detect, and respond to information technology crimes*, 1st ed. Addison-Wesley Professional, 2012.

[15] A. Jones and D. Ashenden, *Risk management for computer security: Protecting your network & information assets*. Butterworth-Heinemann, 2005.

[16] Naked Security. (2014) 1 "terrific employee" + 1 thumb drive + 6,000 lost medical records = fired! [Online]. Available: http://nakedsecurity.sophos.com/2013/01/21/1-terrific-employee-1-thumb-drive-6000-lost-medical-records-fired/

[17] D. S. Wall. (2011) Organizational security and the insider threat: Malicious, negligent and well-meaning insiders (white paper). [Online]. Available: https://www4.symantec.com/Vrt/offer?a_id=108920

[18] F. L. Greitzer, J. Strozer, S. Cohen, J. Bergey, J. Cowley, A. Moore, and D. Mundie, "Unintentional insider threat: Contributing factors, observables, and mitigation strategies," in *System Sciences (HICSS), 2014 47th Hawaii International Conference on*. IEEE, 2014, pp. 2025–2034.

[19] G. Magklaras and S. Furnell, "Insider threat prediction tool: Evaluating the probability of IT misuse," *Computers & Security*, vol. 21, no. 1, pp. 62–73, 2001. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0167404802001098

[20] The Register. (2013) MoJ fined 140k for EMAILING privates of 1,000 inmates. [Online]. Available: http://www.theregister.co.uk/2013/10/22/inmate_detail_mailout_data_breach

[21] Forensicon. (2012) Chicago board of election catastrophic security breach. [Online]. Available: http://www.forensicon.com/forensicon-news/chicago-board-of-election-website-catastrophic-security-breach

[22] A. Beautement, M. A. Sasse, and M. Wonham, "The compliance budget: managing security behaviour in organisations," in *Proceedings of the 2008 workshop on New security paradigms*. ACM, 2009, pp. 47–58.

[23] LA Times. (2012) Data for 700,000-plus home-care workers, recipients goes missing in mail. [Online]. Available: http://latimesblogs.latimes.com/california-politics/2012/05/california-data-breach.html

[24] Data Protection Act. (1998). [Online]. Available: http://www.legislation.gov.uk/ukpga/1998/29/data.pdf