International Journal of Safety and Security Engineering
Vol. 10, No. 1, February, 2020, pp. 49-57
Journal homepage: http://iieta.org/journals/ijsse

A Risk-Based Framework to Inform Prioritisation of Security Investment for Insider Threats

Daniel Sektas-Bilusich*, Rick A. Nunes-Vaz, Leung Chim, Steven Lord

Defence Science and Technology, Edinburgh, South Australia 5111, Australia

Corresponding Author Email: daniel.bilusich@dst.defence.gov.au

https://doi.org/10.18280/ijsse.100107

Received: 22 January 2019
Accepted: 7 November 2019

Keywords:
insider threat, investment prioritisaition, risk management, risk-based framework

ABSTRACT

Threats to information security from inside an organisation are difficult to manage as insiders, by definition, have legitimate access to the organisation's information, consistent with their roles. Impacts of insider threats range from minor information compromise perhaps through carelessness, to catastrophic financial and reputational damage. Security managers are required to continually upgrade security measures to reduce the risk posed by insider threats, however with so many security controls to choose from, finding optimal security solutions based on benefit-cost is challenging. We have developed a risk-based framework called Security-in-Depth (SiD) where residual risk is the metric that assists the security manager to make informed decisions on which security packages contribute more to the organisation's security objectives. We present a case study to illustrate the way our framework is applied, customised to manage a range of insider threats. Uncertainties about the future threat spectrum and the future effectiveness of controls are included in the framework to inform the decision-making process.

1. INTRODUCTION

Most modern organisations have been conditioned to invest and protect their sensitive information from external physical and cyber-attack, in order to maintain their competitive market edge. However, the threat posed by an organisation's employees or contractors, working within the firewall, can be more problematic to manage as they can perhaps more easily cause financial and reputational damage. Insiders, by definition, are trusted to access at least some, if not all, of an organisation's sensitive information. Barriers, firewalls and other controls intended to stop external adversaries are generally ineffective for insiders, making this a major and complex security challenge [1]. Although attacks from outside the organisation are more frequent, insiders deliberately or inadvertently misusing their knowledge and access to information generate greater overall consequences [2-5]. Estimates of the costs of the average insider incident exceed US$400,000, and many incidents exceed US$1 billion in losses [6]. Major information leakages by Bradley/Chelsea Manning in 2010 and Edward Snowden in 2013 [7, 8] demonstrate the extent of plausible impacts when insiders disclose large amounts, or particularly sensitive information.

Manning and Snowden fall into a class of insiders often referred to as disgruntled insiders. They generally enter the organisation as a loyal employee with good intentions but their inclination changes for one of many reasons and they develop malicious intent. More generally, insider threat activities cover a spectrum from careless, accidental or unintentional compromise, e.g., losing sensitive electronic media, through to someone who was trained and deliberately placed in the organisation (referred to as a mole in this paper) in order to exfiltrate specific information (e.g., espionage) or to cause reputational or financial harm. Although there is no universally accepted taxonomy of insider threat types, there are a number of classifications [8-13] and definitions [1, 8, 9, 14]. Security managers must be aware of the spectrum of possible insider threat types in order to make informed holistic decisions for their security arrangements. This is particularly critical considering that security controls will generally have different levels of effectiveness against the different types of insider threats. Understanding the value of controls in managing an organisation's plausible range of insider activities is therefore a challenging but crucial task.

The insider literature is extensive, although a significant portion [15] focuses on describing specific controls or countermeasures for combating a limited range of insider threat types [9, 16-22]. Controls for insiders range from physical and technical to behavioural and organisational [1, 10]. Recent additions to the toolset of controls available for the security manager to consider include eye-tracking to detect either the motivated insider in the process of planning or exfiltrating electronic data [23], or to detect unintentional insider electronic breaches [24]. Many of the controls discussed in the literature could be described as deterrence, detection or prevention techniques with the majority focusing on addressing malicious insiders only. It is critical for security managers to be aware of these factors and limitations when developing security management solutions, and that they take into account the full spectrum of plausible insider threats.

To assist security managers, identify the opportunities they have to control insiders, insider threat pathways (also referred to as kill-chains or attack vectors) have been developed to reveal the general steps an insider follows to achieve their objectives [6, 9, 12, 25-31]. Examples of threat pathways for both careless and disloyal insiders are shown in Figure 1. The pathways represent scenarios involving a sequence of steps required to achieve the security breach. In the example pathway for the careless insider, benign intent associated with taking files to work on at home may result in the media being

misplaced, later found by someone who profits from selling the media to a third-party publisher. For the disloyal insider, the example pathway may represent an employee who becomes disgruntled and decides to steal information and sell it to a competitor resulting in loss of business advantage. Many pathways (scenarios) are feasible for both unintentional and malicious insiders, but for the purposes of analysis, many can be shown to be equivalent (refer to [32]). The two shown in Fig. 1 are abstracted to a high level of detail to allow the security manager to identify the character of potential interventions to prevent the security breach or reduce its damage. A framework for how to disaggregate pathways into more detailed steps is also available [27] and may be useful to pinpoint where specific controls should be emplaced. Although the majority of literature describes preventative controls to stop the occurrence of a security breach, by assessing the threat pathways (Figure 1) it becomes clear that there are also opportunities to reduce consequences when a security compromise has occurred. Recent literature is starting to articulate this [8].

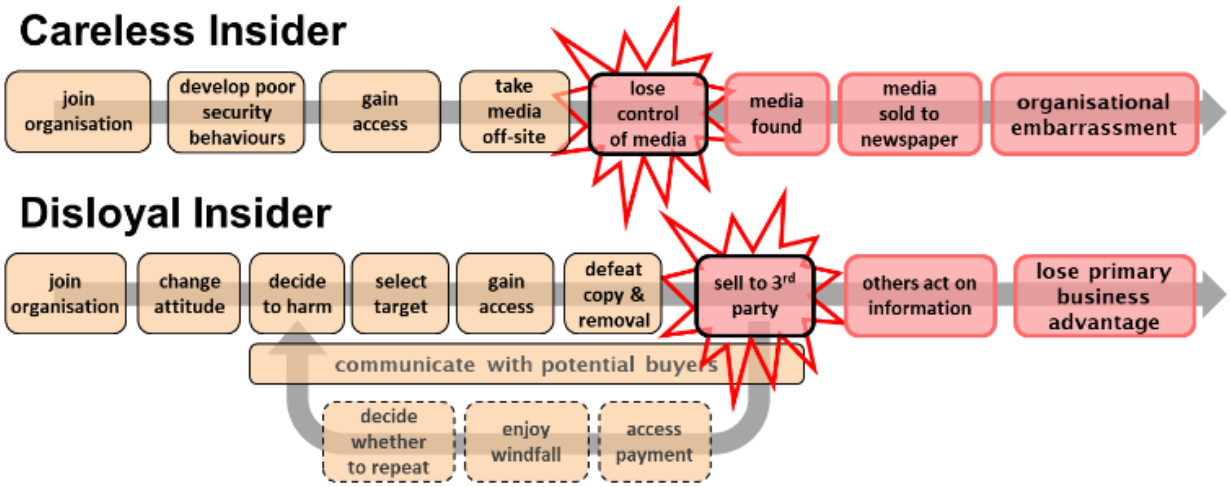

Figure 1. Threat pathways for a carless (unintentional) and disloyal (malicious) insider. The security breach (risk event) is shown by the star in each pathway

The insider threat is complex (with different threat types and various threat pathways to consider). With so many controls available, the task for the security manager to select the best security package to implement is quite difficult. Insider security strategies, or approaches, have been developed to guide the security manager's focus in order to achieve a stronger security system as a whole [6, 9, 33-41]. Some state where along the threat pathway the opportunity to stop the perpetrator is greatest. Others describe what sets of controls working together are effective, or whether the focus should be specific types of controls (such as controls that achieve deterrence) over others (that target prevention as an example). However, upon close examination of these strategies, we identify some limitations that must be taken into account by the security manager when deciding which approach to implement.

One strategy based on experience within the FBI focuses on the disgruntled insider as the main insider threat type [6]. They argue that deterrence is more effective than monitoring behaviours and detecting inappropriate actions. A different strategy, referred to as 'no dark corners', advocates targeting the higher-level activities of moles by empowering workers at the team level to deter, detect and intervene [36]. According to [34], detection of breaches is more crucial than deterrence, as an organisation's efforts to prevent inappropriate insider activity is reasonably likely to stop the unsophisticated or weakly-motivated insider, but will be less effective against a determined or sophisticated attacker, who causes the most damage. For these insiders, they advocate an approach that focuses on detection of attacks to inform damage control, as well as evidence gathering and prosecution to limit damage escalation and deter future potential perpetrators. This approach contrasts with others by being generally reactive, and focuses on security controls that work in the post-breach event timeframe. The final strategy [35] in our summary is based on General Deterrence Theory [42, 43] where the aim is to maximise the effectiveness of deterrence and prevention in order to minimise the need for post-event detection and prosecution. The model strengthens deterrence by making potential offenders keenly aware of the consequences of an inappropriate activity.

In our brief review we note that some authors were quite clear about the range of insider types their approach was targeting, and it was apparent that different security philosophies were needed to address different parts of the insider threat spectrum. Others unfortunately were less clear, and whether they intended their advice to be generally applicable, or applied to more specific threat types, remained ambiguous. We suggest that security managers should be cautious about adopting a particular security strategy without testing its value across the spectrum of insiders their organisation may face. We also note from the cited approaches above that some of the advice is inconsistent. On the assumption that inconsistency means that not all are correct, how does the security manager identify which approach to apply to achieve the greatest benefit-cost for their security investment? For those that advocate a multi-methodology approach [18, 35], how does the security manager determine which strategies deserve investment, and in what relative balance? These two questions are our research questions. To answer these questions, we need a framework that allows the security manager to test the relative effectiveness of alternative security strategies (and combinations) against a variety of insider types. This would enable them to prioritise their investments and customise their security arrangements to meet the organisation's needs. Security solutions should be guided by a conceptual framework that takes into account the various insider types, as well as how security controls interact

to reduce insider risk along the various threat pathways [10, 26, 37, 44, 45].

In the remainder of the paper, we address these issues using a risk-based framework which we have called Security-in-Depth (SiD) [46]. SiD was originally developed to support investment decisions in the physical security domain [47, 48], but was then extended and applied to address all national security threat types [32] and to explore Defence's needs in building cyber security capabilities [49]. We begin by providing a brief introduction to the SiD approach, and then apply it to a case study example where we take the role of the security manager of a hypothetical organisation tasked with improving the security against insider threats with a limited budget. The case study is illustrative only, and guides users through the process of applying SiD and modelling the problem, to make informed investment decisions. We take account of uncertainty about the future and uncertainty in estimating the effectiveness of the security systems against the various insider threat scenarios, and describe how these can be modelled to support investment decisions. The process can be applied by a security manager using data and evaluations specific to their organisation's requirements to inform wise security investment decisions.

2. SID RISK-BASED FRAMEWORK

The SiD framework [32, 46] is a systematic approach to risk analysis, which links risks to organisational objectives, in our case insider threats to information security. The framework is consistent with the ISO 31000:2009 Risk Management Standard [50] and can be used both qualitatively and quantitatively to inform investment decisions by prioritising controls that have the greatest potential for risk reduction. Central to the framework is the concept of a security layer which is defined as an integrated set of controls that can potentially stop a defined event from occurring, or reduce the consequences when an event has occurred [46]. The key to this definition of a security layer is that the layer contains all the controls needed to reduce the likelihood or consequences of a risk event independent of other control measures. Each layer achieves an independent contribution to risk reduction. Security layers are applied to manage specific parts of a threat pathway, and using multiple layers requires the attacker to defeat each and all security layers to be successful. Building multiple layers generates multiple opportunities to stop the perpetrator or limit the harm. Although the concept of multiple security layers is not new, what SiD has been diligent to provide is a clear definition of a security layer in the context of risk management and to support investment prioritisation decisions based on the calculation of residual risk.

Each layer is independent from other layers (although individual controls may contribute to more than one layer) and the effectiveness of each layer can be evaluated. This is a critical aspect of SiD that can be used to support prioritisation decisions and will be demonstrated in the following section. A layer is composed of one or more functions, each of which is performed by integrated sets of security controls. Security controls can range from physical, technical, psychological, to procedural. The hierarchy – layer, function and control – is central to the SiD framework. Controls contribute to the performance of security functions, and several different functions must be integrated coherently to create each security layer.

The term 'security layer' is a commonly used term in literature, as is the need for security systems to employ multiple layers. However, we warn readers that many references to the term 'security layer', are actually referring to individual controls (which on their own may not reduce risk). Others also describe security functions as layers, an example being 'detection'. Detection (which can be composed of several controls such as surveillance cameras, security patrols and motion sensors) must be coupled with a security response function (intervention), to prevent the attack from occurring. Detection without a response does not reduce the risk of the event. When the term layer is used too loosely, it fails to support comparisons of effectiveness, or trade-offs between alternative security arrangements. Mixing the concepts of controls, functions and layers, precludes comparison on the basis of risk reduction effectiveness. Comparison of the risk reduction effectiveness of different layers is the value contribution of the SiD framework to risk-based investment prioritisation in security.

In SiD, the threat attack is modelled as a critical pathway of sequential steps leading to the risk event, followed by fallout steps that generate the undesirable consequences and potential impact (Figure 2). SiD consists of seven types of layers, as shown. The shape, deter and prevent layers affect the likelihood of a risk event occurring while protect, contain and recover/adapt layers affect the consequences and impacts where a risk event has occurred. The investigate layer contributes to prevent future events.

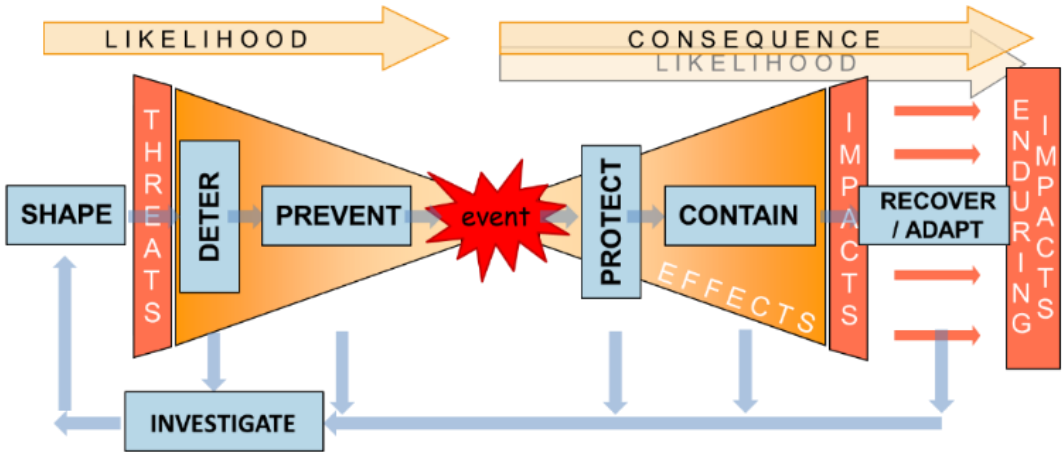

Figure 2. The Security-in-Depth approach of layers along a risk pathway

The shape layer may involve screening potential insiders before they join the organisation, or influencing the mind state of potential perpetrators so they do not develop malicious intent. The deter layer refers to a collection of controls which work to discourage someone with careless or malicious intent from taking actions, which may lead to harm to the organisation. Deterrence is often generated according to the effectiveness, or perceived effectiveness of controls belonging to other layers, i.e., the strength of perimeter barriers, while primarily intended to prevent physical breach, may also contribute to deterrence. If such barriers are perceived as weak, not only do they fail to deny access, but they also fail to deter attempts at access. For those perpetrators not deterred, the prevent layer is a collection of controls and functions, which act to identify and monitor activities to prompt interventions that stop the risk event before the breach occurs. Both the shape and prevent layers are made up of three interlinked generic functions; detect, alert, and respond [31, 46]. As an example, for the prevent layer, the three functions are needed to generate awareness of the perpetrator's actions (detect), raise an alarm and undertake decision making (alert), and execute an intervention to stop the risk event from occurring (response). Each layer alone, if fully effective, could potentially reduce the likelihood of a successful attack to zero.

Where a risk event is not stopped, and information has been compromised, there are still security controls which can reduce some of the consequences and impacts of such occurrence. Post-event layers include protect, contain and recover/adapt. The protect layer consists of passive, impact-specific controls such as encryption to deny exploitation of information that is exfiltrated. The contain layer involves the same active functions as in shape and prevent, that is, detect, alert and response, although the nature of these functions is quite different here, involving detecting the breach and raising actions that manage the harm (such as enforcing laws to deny publication in public domain outlets). The recover/adapt layer stops immediate consequences escalating to greater organisational impacts and is often encapsulated under principles of business continuity or resilience. The investigate layer involves the accumulation of evidence based on previous events (whether successful or otherwise), or from models or experience shared from elsewhere (other similar organisations or security agencies, e.g., Computer Emergency Response Team (CERT)). The investigate layer contributes to the prevention of future events.

3. CASE STUDY APPLICATION OF SID TO THE PRIORITISATION OF SECURITY INVESTMENT TO MANAGE INSIDERS

In this section we illustrate the form of analysis that might be undertaken by an organisation using the SiD framework to assess and invest in insider threat management. We assume that the security manager has a limited investment budget and is tasked with making wise decisions about balancing systems to manage corresponding risks.

As noted earlier, several types of insider belong to accidental or malicious categories. To keep things relatively simple we select two, that is, one from each category, although the approach scales up easily. We describe the extant security arrangement as the 'default', and develop two different enhanced, post-investment security arrangements which we evaluate for their overall effectiveness. Note that in reality the analysis would need to consider all relevant types of insider, and perhaps more than two alternative investment approaches.

In order to apply our approach, the first step is to identify and express our objectives. In this case our objective is to secure sensitive information, ignoring all potential external threats (although, again, the method can be scaled up to include these). The next step involves understanding how the two threat types potentially impact the objective, for which we draw pathways (as shown in Figure 1). There is clearly an infinite number of potential pathways for the exfiltration of data by an insider, however, a small number of pathways at the right level of abstraction can be shown to represent the problem, avoiding duplication and redundant analytical effort.

Pathways allow the security manager to identify potential interventions along them, that is, where controls might be placed and, in terms of their contributions to functions and layers, how effective alternative solutions may be. Working through the layers helps the security manager to identify a comprehensive list of potential control systems. Such qualitative analysis is often superior to other approaches that confuse controls and functions with layers. But, the SiD framework permits quantitative analysis of potential solutions.

In our scenario, the risk event (centre of the bow tie in Figure 2) is the compromise of sensitive information. For careless insiders this may include leaving sensitive information on the printer where someone without access privileges to that information removes it from site. For disloyal insiders it can include direct removal of information from the premises, either physically or electronically. In such cases the left-hand side of the risk event contains security controls to manage the likelihood of information compromise, while the right-hand side involves security controls that reduce the consequences to the organisation where the information is compromised. For this simplified case study, we focus on the risk event being a single compromise of information and exclude the effectiveness of the investigate layer from our analysis.

In Table 1 we show example controls that contribute to the various layers of SiD. This set is for illustration purposes only and is not a complete or recommended set. Some controls are effective against both the careless and disloyal insiders, such as 'good work-life balance' policies where the insider is less likely to become a careless insider by emailing sensitive files to home email to finish on the weekend, and less likely to become a disloyal insider where their good nature is not exploited. Other controls may be more effective for one insider type than the other. We will use this set to represent the current 'default' package of insider security controls in our organisation. The security manager is tasked with adding additional controls to this set to improve security.

We split the analysis into two steps. The first step is to assess the likelihood reduction of each left-hand side layer for each of the insider types. We start by considering the hypothetical default set of controls as described in Table 1 under the shape, deter and prevent layers. Layer effectiveness can be determined through literature reports of the effectiveness of specific security controls, through the personal experience and judgements of the security manager and/or other experts, or some justifiable combination of these.

Table 1. Example security controls against careless and disloyal insiders clustered within security layers

	Shape	Deter	Prevent	Protect	Contain	Recover/Adapt
Careless Insider	• Annual security course • Work-life balance	• Policy of inspections • Prosecutions of minor breaches • Random car searches	• Random car searches • Email keyword monitoring	• Encryption of highly sensitive files	• Org.releases conflicting information	• Org.shifts focus • Org.speeds up investment
Disloyal Insider	• Work-life balance • Good pay	• Random car searches	• Random car searches • Email keyword monitoring	• Encryption of highly sensitive files	• Org.releases conflicting information • Org.pays ransom for return of information	• Org.shifts focus • Org.speeds up investment

Figure 3 shows the hypothetical assessment by the security manager for the effectiveness of the default security controls in the shape, deter and prevent security layers, for the careless insider. The utility of the SiD framework is that the layers are independent which allows the security manager to only consider the interconnection of controls within one layer when assessing its effectiveness, as its effectiveness is independent of controls in other layers. Based on Figure 3, the security manager has assessed that the controls within the shape layer will be 40% effective at stopping employees developing motivation or an attitude, which would lead them to behave in a way that may compromise information security. In 60% of cases these controls will not be effective.

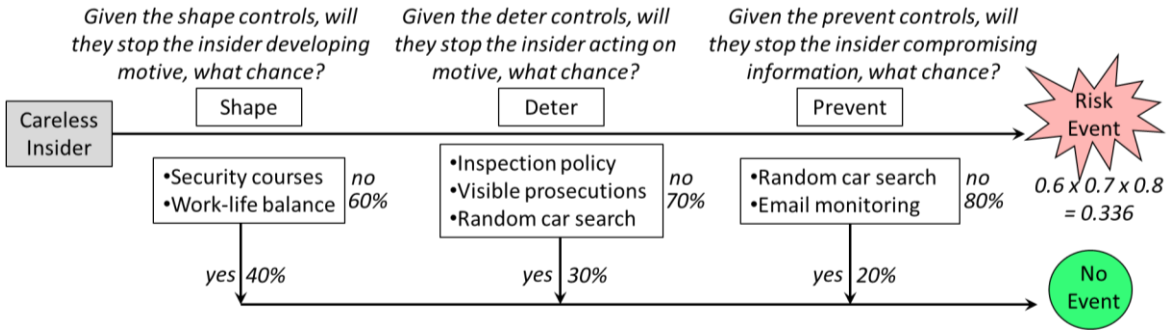

Figure 3. Hypothetical assessment of default security control effectiveness for the shape, deter and prevent security layers for the careless insider

It was also assessed that another combination of controls would deter 30% of potential careless insiders (those whose attitudes to security have not been adequately shaped). Similarly, of those who were not adequately shaped or deterred, extant controls are likely to stop (prevent) 20% of careless actions.

Since security layers are independent, and the effectiveness of one layer is independent of the effectiveness of others, the likelihood that a careless insider will not be stopped from compromising sensitive information is calculated by the product of the ineffectiveness of the layers, in our example 0.6 x 0.7 x 0.8 = 0.336.

However, any particular insider is likely to behave quite differently from any other, and combined with potentially significant epistemic uncertainty associated with control effectiveness (a colleague may be particularly diligent in correcting poor security behaviour one day and quite ambivalent on another), precise assessments of effectiveness are unjustified. Even asking experts to assess against defined likelihood ranges (bins) can be difficult and flawed, such as 'our deterrence is 30-40% effective'. We therefore advocate that the security manager or expert group be given the freedom to express their uncertainties, by distributing their assessments right through the likelihood range (for example, 10% confidence that the controls will be between 0 and 5% effective, 25% confidence that they will be between 5% and 20% effective etc.). Figure 4 illustrates these opinions as distributions of confidence (small graphs displayed under each layer) that the set of controls in question will be effective. Bin sizes for these distributions are not marked, but can be set to achieve discrimination as required. For one security arrangement it might be useful to use equal bins, i.e., 0-20% (very low effectiveness of controls), 20-40% (low), 40-60% (medium), 60-80% (high), 80-100% (very high). In another, it may be prudent to use bins that emphasise extreme values, e.g., 0-2%, 2-10%, 10-50%, 50-90%, 90-98% and 98-100%.

Figure 4 shows hypothetical assessments for both the careless and disloyal insiders with the default security controls. The probability distribution for the risk event is then calculated by applying Monte Carlo simulations. The Monte Carlo simulation approach plays out a single weighted random selection of the effectiveness of each layer (according to the weightings (the distribution in confidence) allocated by the experts) and calculates the combined (three layer) result. It then runs another simulation, randomly selecting effectiveness values again. By running many thousands of such simulations, the statistical effectiveness of stopping breaches, based on expert opinions of the effectiveness of each layer, is determined. Figure 4 shows a hypothetical result for the likelihood of the breach occurring below the risk event star on the right. In this illustration, the compromise of sensitive information will happen more often for careless insiders than it will for disloyal insiders.

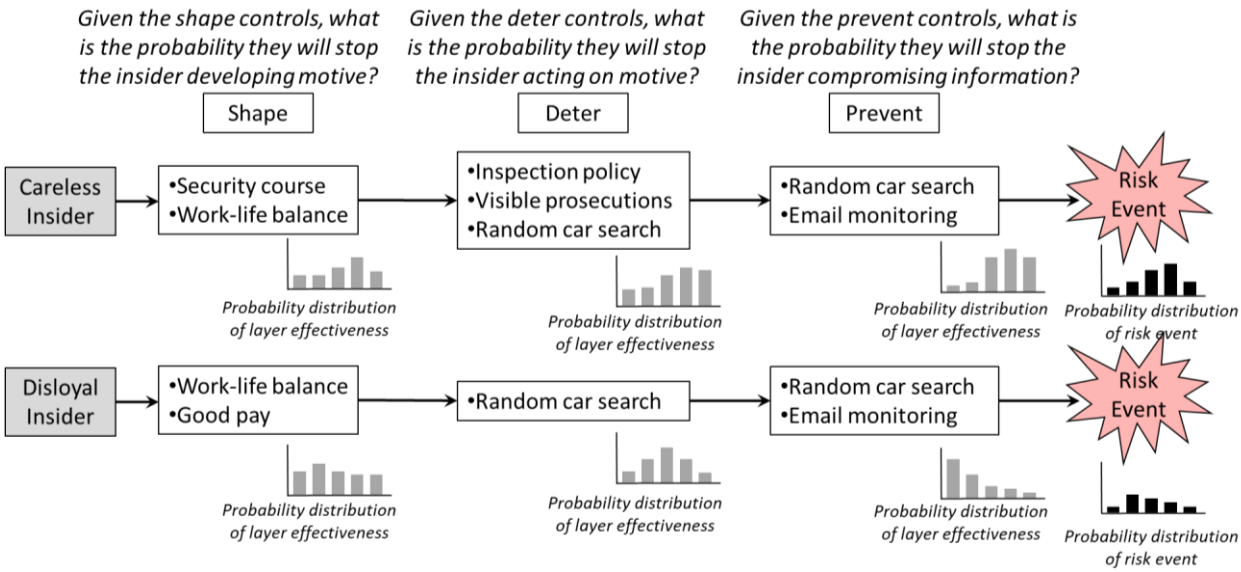

Figure 4. Uncertainty in effectiveness of controls within security layers for the careless and disloyal insiders, to produce the probability distribution of the risk event

Once the analysis has been completed using the default security controls (layers), the security manager should generate enhanced control packages (consistent with the investment budget) and repeat the process to determine comparative outcomes. Alternative enhancement proposals can be tested in this way, following different security design hypotheses (e.g., deterrence is more effective than prevention, or vice versa).

In our hypothetical case study, the security manager wishes to compare the effectiveness of enhancing the default security system by adding package A or package B. Table 2 shows that package A contains three controls that contribute to security layers that reduce the likelihood of information compromise (shape and prevent) while package B consists of four different controls within the same two security layers.

Table 2. Additional security control measures clustered under package A and B that contribute to reducing the likelihood of a risk event

	Shape	Prevent
Package A	• Background checks on employment	• Eye monitoring for fatigue errors • Tattle-taps on sensitive reports
Package B	• Regulat social events • Casual Fridays	• Eye monitoring for fraudulent behaviour • Working in pairs policy

In our hypothetical alternative security enhancements, the outcome of analysis is a total of six probability distributions showing residual distributions of 'likelihood of breach' for two types of insider (careless and disloyal), against each of three different security systems (default, enhanced package A and enhanced package B) as shown in Figure 5. From the figure we can see that package A is more effective in reducing the probability of a breach for the careless insider, while package B is more effective against the disloyal insider.

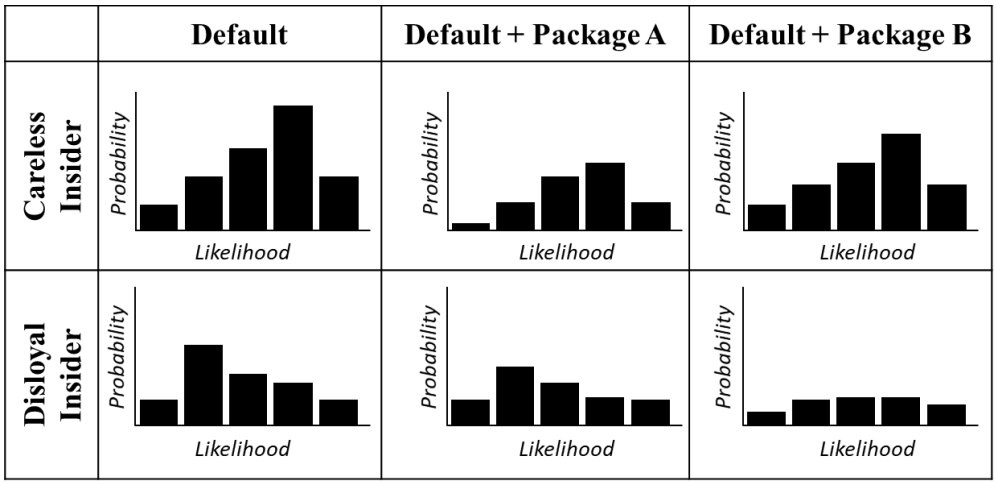

Figure 5. Probability distribution of the likelihood of a risk event for careless and disloyal insiders with default, and additional package A or B security controls

A similar procedure is followed to assess the impacts of security breaches, by seeking effectiveness assessments for each of the consequence management layers (protect, contain and recover/adapt). In this case, experts are asked to assign confidence values to bins that represent the magnitude of impacts (e.g., '0 to $10,000 loss', '$10,000 to $1m loss' and etc., or perhaps 'reputational damage causing 0 to 10% revenue loss', 'reputational damage causing 10 to 25% revenue loss' and etc.). It is at the security manager's discretion to determine how many consequence bins are most relevant to his/her organisation and what kind of impacts are most useful to inform security investment decisions. For our hypothetical case study, we define three consequence or impact bins (notionally 'low', 'medium' and 'high'). Using expert weightings of the effectiveness of protect, contain and recover/adapt layers we again generate six distributions (not shown) representing impacts of breaches when they occur. The process for determining two of these impact distributions (for the careless and disloyal insiders with default controls from Table 1) is shown in Figure 6.

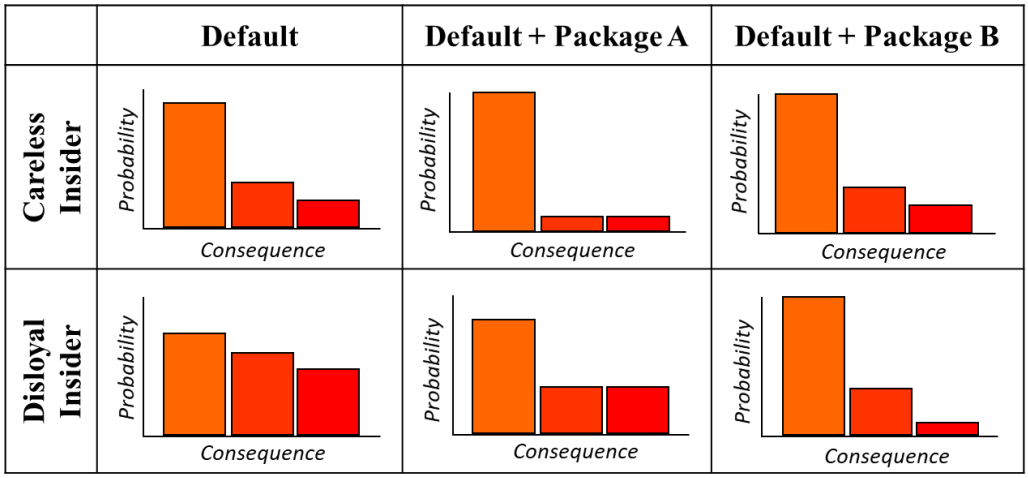

Figure 6. Probability distribution of potential consequences through the protect, contain and recover/adapt security layers for the careless and disloyal insiders with default security controls

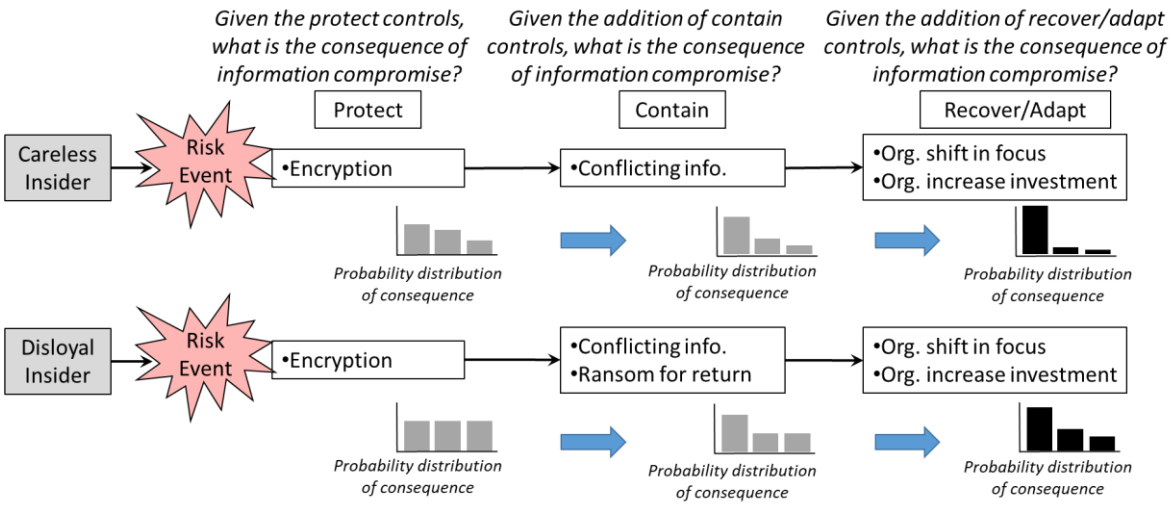

Figure 7. Residual risk for the careless and disloyal insiders with default, and addition of package A and B security controls

The outcome of conducting Monte Carlo simulations when all layers are included, is to determine residual risk ($R = f(P, C)$ where P is the probability and C the consequence of a specific event) associated with each modelled security arrangement. Different security systems will generate different residual risk distributions (see Figure 7). For our illustration, enhancement package A is more effective at limiting the damage from careless insiders, while package B does a better job against disloyal insiders. Note that the relative event frequencies for the careless versus disloyal insiders (say, 50 times more careless than disloyal events, e.g., [24]) has been incorporated into this result. It shows that the nett impact of (relatively infrequent) disloyal insider events is not dissimilar to the nett impact of the (far more frequent) careless insider events (for all package options). A sensitivity analysis considering the ratio of careless to disloyal insider events would then identify when package A is the superior investment choice, and at what ratio that decision switches to package B.

While we have deliberately constrained the scope of analysis, a real-world assessment would necessarily consider all relevant insider types as well as a broader set of upgrade options tailored for the organisation of interest. The SiD framework allows us to integrate the effect of control packages into higher level constructs that are intuitively assessable by security experts. Without this abstraction, experts would need to consider the overall effects of the entire control arrangement to assess security effectiveness: a far more difficult cognitive task. The independence of security layers in the SiD

framework allows the security manager to assess the effectiveness of the smaller, more mentally manageable set of controls within the layer, and combine layer effectiveness to calculate the residual risk, and support prioritisation of investment.

This process has been applied by the authors to inform decisions for Australian Defence and National Security agencies.

4. CONCLUSION

Insiders, by virtue of their legitimate access to their organisation's information, pose a significant risk to integrity of sensitive information both accidentally and deliberately. Security managers are continually required to ensure their organisation has the most effective controls in place to reduce this risk. With limited budgets and an abundance of security controls to choose from, it can be difficult to determine which set of controls is superior to any other choices in the complex security arrangements of moderately sized organisations.

We have developed Security-in-Depth, a risk-based framework, to assist security managers make informed investment decisions using residual risk as the principal metric. By explicitly defining the parameters of a security layer, the effectiveness of the independent security layers can be assessed, and combined to determine the likelihood and consequences of specific threats, and therefore the residual risks. By assessing various combinations of security controls, the security manager is able to make informed decisions on where to invest based on benefit-cost to improve the security system and identify what security strategies are most appropriate for their security requirements. As the future is uncertain, and the exact effectiveness of controls is not always fully understood, we demonstrated how to incorporate this uncertainty into the investment decision process.

REFERENCES

[1] CERT (2016). Common sense guide to mitigating insider threats, fifth edition.

[2] CERT 2011 Cybersecurity watch survey. How bad is the insider threat? (2011).

[3] Hua, J., Bapna, S. (2013) Who can we trust?: the economic impact of insider threats. Journal of Global Information Technology Management, 16(4): 47-67. https://doi.org/10.1080/1097198X.2013.10845648

[4] Upton, D.M., Creese, S. (2014). The danger from within. Harvard Business Review. 95-101, Sep.

[5] Wang, Y.L., Yang, S.C. (2014). A method of evaluation for insider threat. Proceedings of the 2014 International Symposium on Computer, Consumer and Control, pp. 438-441. https://doi.org/10.1109/IS3C.2014.121

[6] Reidy, P., Randal, K. (2013). Combating the insider threat at the FBI: real world lessons learned. Presented at RSA Conference, San Francisco CA.

[7] Greene, R., Kehl, D., Morgus, R., Bankston, K. (2014). Surveillance costs: The NSA's impact on the economy, internet freedom and cybersecurity.

[8] NATO Cooperative Cyber Defence Centre of Excellence. Insider Threat Detection Study. https://ccdcoe.org/insider-threat.html, accessed on 6 Dec. 2018.

[9] Smith, J.A. (2015). Mitigating malicious insider cyber threat. Royal Holloway Univeristy of London, UK.

[10] Hunker, J., Probst, C.W. (2011). 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. https://doi.org/10.22667/JOWUA.2011.03.31.004

[11] Securelist. Recognizing Different Types of Insiders. https://securelist.com/threats/recognizing-different-types-of-insiders/, accessed on 14 Oct. 2011.

[12] Greitzer, F.L., Imran, M., Purl, J., Axelrod, E.T., Leong, Y.M., Becker, D.E., Laskey, K.B., Sticha, P.J. (2016). Developing an ontology for individual and organizational sociotechnical indicators of insider threat risk. Proceedings of the 11th International Conference on Semantic Technology for Intelligence, Defense, and Security, pp. 19-27.

[13] Ismail, I., Hassan, R., Othman, M.R., Ahmad, A.S., Tawfiq, N.E. (2017). Insider risk profile matrix to quantify risk value of insider threat prediction framework. Journal of Theoretical and Applied Information Technology, 95(20): 5595-5608.

[14] Bishop, M., Engle, S., Frincke, D.A., Gates, C., Greitzer, F.L., Peisert, S., Whalen, S. (2010). A risk management approach to the 'insider threat'. In: Probst, C.W., Hunker, J., Bishop, M., Gollmann, D. (eds.) Insider Threats in Cyber Security, Springer, New York NY, 115-137. https://doi.org/10.1007/978-1-4419-7133-3

[15] Ophoff, J., Jensen, A., Sanderson-Smith, J., Porter, M., Johnston, K. (2014). A descriptive literature review and classification of insider threat research. Proceedings of Informing Science & IT Education Conference, pp. 211-223.

[16] Zeadally, S., Yu, B., Jeong, D.H., Liang, L. (2012). Detecting insider threats: solutions and trends. Information Security Journal: A Global Perspective, 21(4): 183-192. https://doi.org/10.1080/19393555.2011.654318

[17] Guido, M.D., Brooks, M.W. (2013). Insider threat program best practices. Proceedings of the 46th Annual Hawaii International Conference on System Sciences, pp. 1831-1839. https://doi.org/10.1109/HICSS.2013.279

[18] Ahmad, A., Maynard, S.B., Park, S. (2014). Information security strategies: towards an organizational multi-strategy perspective. Journal of Intelligent Manufacturing, 25(2): 357-370. https://doi.org/10.1007/s10845-012-0683-0

[19] US Department of Homeland Security, Combating the insider threat. (2014). National Cybersecurity and Communications Integration Center.

[20] Hashem, Y., Takabi, H., GhasemiGol, M., Dantu, R. (2016). Inside the mind of the insider: towards insider threat detection using psychophysiological signals. Journal of Internet Services and Information Security, 6(1): 20-36. https://doi.org/10.22667/JISIS.2016.02.31.020

[21] Almehmadi, A., El-Khatib K. (2017). On the possibility of insider threat prevention using intent-based access control (IBAC). IEEE Systems Journal, 11(2): 373-384. https://doi.org/10.1109/JSYST.2015.2424677

[22] Omar, M., Mohammed, D., Nguyen, V. (2017). Defending against malicious insiders: a conceptual framework for predicting, detecting, and deterring

malicious insiders, International Journal of Business Process Integration and Management, 8(2): 114-119. https://doi.org/10.1504/IJBPIM.2017.083794
[23] Matthews, G., Reinerman-Jones, L., Wohleber, R., Ortiz, E. (2017). Eye tracking metrics for insider threat detection in a simulated work environment. Proceedings of the Human Factors and Ergonomics Society, 61(1): 202-206. https://doi.org/10.1177/1541931213601535
[24] Takabi, H., Hashem, Y., Dantu, R. (2018). Prediction of human error using eye movements patterns for unintentional insider threat detection. Proceedings of the IEEE 4th International Conference on Identity, Security, and Behaviour Analysis. https://doi.org/10.1109/ISBA.2018.8311479
[25] IAEA Preventive and protective measures against insider threats. (20018). International Atomic Energy Agency, Vienna, Austria.
[26] Duran, F.A., Conrad, S.H., Conrad, G.N., Duggan, D.P., Held, E.B. (2009). Building a system for insider security. IEEE Security & Privacy, 7(6): 30-38. https://doi.org/10.1109/MSP.2009.111
[27] Nurse, J.R.C., Buckley, O., Legg, P.A., Goldsmith, M., Creese, S., Wright, G.R.T., Whitty, M. (2014). Understanding insider threat: a framework for characterising attacks. Proceedings of the IEEE Security and Privacy Workshops, 214-228. https://doi.org/10.1109/SPW.2014.38
[28] Kammuller, F., Nurse, J.R.C., Probst, C.W. (2016). Attack tree analysis for insider threats on the IoT using Isabelle. Proceedings of the 4th International Conference on Human Aspects of Security, Privacy and Trust.
[29] Musman, S., Turner, A.J. (2018). A game oriented approach to minimizing cybersecurity risk. International Journal of Safety & Security Engineering, 8(2): 212-222. https://doi.org/10.2495/SAFE-V8-N2-212-222
[30] Chim, L., Bilusich, D., Lord, S., Nunes-Vaz, R. (2017). A risk-based layered defence for managing the trusted insider threat. Journal of Information System Security, 13(3): 151-173.
[31] Bilusich, D., Chim, L., Nunes-Vaz, R.A., Lord, S. (2018). There is no single solution to the 'insider' problem but there is a valuable way forward. WIT Transactions on Engineering Sciences, 121: 135-146. https://doi.org/10.2495/RISK180121
[32] Nunes-Vaz, R., Lord, S., Bilusich, D. (2014). From strategic security risks to national capability priorities. Security Challenges, 10(3): 23-49.
[33] Maybury, M., Chase, P., Cheikes, B., Brackney, D., Matzner, S., Hetherington, T., Wood, B., Sibley, C., Marin, J., Longstaff, T., Spitzner, L., Haile, J., Copeland, J., Lewandowski, S. (2005). Analysis and detection of malicious insiders. Proceedings of the International Conference on Intelligence Analysis.
[34] Cole, E., Ring, S. (2006). Insider threat: protecting the enterprise from sabotage, spying, and theft, Syngress, Rockland MA.
[35] Theoharidou, M., Kokolakis, S., Karyda, M., Kiountouzis, E. (2005). The insider threat to information systems and the effectiveness of ISO17799. Computers & Security, 24(6): 472-484. https://doi.org/10.1016/j.cose.2005.05.002
[36] Catrantzos, N. (2010). Tackling the insider threat, ASIS International Foundation CRISP Report.
[37] Montelibano, J., Moore, A. (2012). Insider threat security reference architecture. Proceedings of the 45th Annual Hawaii International Conference on System Sciences, pp. 2412-2421. https://doi.org/10.1109/HICSS.2012.327
[38] Park, S., Ruighaver, A.B., Maynard, S.B., Ahmad, A. (2012). Towards understanding deterrence: information security manager's perspective. Proceedings of the International Conference on IT Convergence and Security, pp. 21-37. https://doi.org/10.1007/978-94-007-2911-7_3
[39] Australian Government. (2016) Managing the insider threat to your business: A personnel security handbook.
[40] Buckley, O., Nurse, J.R.C. Legg, P.A. Goldsmith, M., Creese, S. (2014). Reflecting on the ability of enterprise security policy to address accidental insider threat. Proceedings of the Workshop on Socio-Technical Aspects in Security and Trust, pp. 8-15. https://doi.org/10.1109/STAST.2014.10
[41] Stavrou, V., Kandias, M., Karoulas, G., Gritzalis, D. (2014). Business process modeling for insider threat monitoring and handling. In: Eckert, C., Katsikas, S.K., Pernul, G. (eds.) Trust, Privacy, and Security in Digital Business, Springer, 119-131. https://doi.org/10.1007/978-3-319-09770-1_11
[42] Forcht, K.A. (1994). Computer security management, Boyd & Fraser, Danvers MA.
[43] Straub, D.W., Welke, R.J. (1998). Coping with systems risk: security planning models for management decision making. Management Information Systems Quarterly, 22(4): 441-469.
[44] Stolfo, S., Bellovin, S.M., Evans, D. (2011). Measuring security. IEEE Security & Privacy, 9(3): 60-65. https://doi.org/10.1109/MSP.2011.56
[45] Pan, L., Tomlinson, A. (2016). A systematic review of information security risk assessment. International Journal of Safety & Security Engineering, 6(2): 270-281. https://doi.org/10.2495/SAFE-V6-N2-270-281
[46] Nunes-Vaz, R., Lord, S., Ciuk, J. (2011). A more rigorous framework for security-in-depth. Journal of Applied Security Research, 6(3): 372-393. https://doi.org/10.1080/19361610.2011.580283
[47] Lord, S., Nunes-Vaz, R. (2013). Designing and evaluating layered security. International Journal of Risk Assessment and Management, 17(1): 19-45. https://doi.org/ 10.1504/IJRAM.2013.054377
[48] Nunes-Vaz, R., Lord, S. (2014). Designing physical security for complex infrastructures. International Journal of Critical Infrastructure Protection, 7(3): 178-192. https://doi.org/10.1016/j.ijcip.2014.06.003
[49] Rowe, C., Seif Zadeh, H., Garanovich, I.L., Jiang, L., Bilusich, D., Nunes-Vaz, R., Ween, A. (2017). Prioritizing investment in military cyber capability using risk analysis. Journal of Defense Modeling and Simulation: Applications, Methodology, Technology, 16(3): 1-13. https://doi.org/10.1177/1548512917707077
[50] ISO/IEC 31000:2009 Risk Management – Principles and Guidelines.

NOMENCLATURE

C	consequence
P	probability
R	risk