Incorporating a Knowledge Perspective into Security Risk Assessments

Piya Shedden
Department of Information Systems
University of Melbourne
Victoria 3010
p.shedden@pgrad.unimelb.edu.au

Rens Scheepers
Department of Information Systems
University of Melbourne
Victoria 3010
r.scheepers@unimelb.edu.au

Wally Smith
Department of Information Systems
University of Melbourne
Victoria 3010
wsmith@unimelb.edu.au

Atif Ahmad
Department of Information Systems
University of Melbourne
Victoria 3010
atif@unimelb.edu.au

Abstract

Purpose – Many methodologies exist to assess the security risks associated with unauthorized leakage, modification and interruption of information used by organisations. This paper argues that these methodologies have a traditional orientation towards the identification and assessment of technical information assets. This obscures key risks associated with the cultivation and deployment of organisational knowledge. The paper aims to explore how security risk assessment methods can more effectively identify and treat the knowledge associated with business processes.

Design/ methodology/ approach – The argument was developed through an illustrative case study in which a well-documented traditional methodology is applied to a complex data backup process. Follow-up interviews were conducted with the organisation's security managers to explore the results of the assessment and the nature of knowledge 'assets' within a business process.

Findings – It was discovered that the backup process depended, in subtle and often informal ways, on tacit knowledge to sustain operational complexity, handle exceptions and make frequent interventions. Although typical information security methodologies identify people as critical assets, this study suggests a new approach might draw on more detailed accounts of individual knowledge, collective knowledge and their relationship to organisational processes.

Originality/ value – Drawing on the knowledge management literature, we suggest mechanisms to incorporate these knowledge-based considerations into the scope of information security risk methodologies. A series of propositions and a knowledge protection model are presented. This model outlines ways in which organisations can effectively identify and treat risks around process knowledge critical to the business.

Keywords Information security, security risk management, asset identification, knowledge protection.

Paper Type Case study.

Introduction

Information security is of paramount importance to organisations. Information security risk assessments (ISRAs) enable organisations to identify their key information assets and risks in order to develop effective and economically-viable control strategies (Baskerville, 1991; Roper, 1999; Peltier, 2001; den Braber et al, 2007). Various popular ISRA methodologies are used in industry, including OCTAVE, CRAMM, NIST SP800-30 and

the AS/NZS 4360 standard (Stoneburner et al, 2002; Yazar, 2002; Alberts & Dorofee, 2004; AS/NZS, 2004). These ISRA methodologies ensure that critical assets are identified through a rigorous process involving various stakeholders, including senior management, operational managers and technical staff (Alberts & Dorofee, 2004). This facilitates protective action for valuable or critical assets whilst also ensuring that resources are not wasted on protecting lesser risks or unimportant assets (Visintine, 2003). Significantly, these risk assessments tend to be *asset-focussed* in that they are based on identifying information as objects of value that are threatened, have vulnerabilities and require protection (Shedden et al, 2006).

Distinct from the discourse on information security, the literature on knowledge management has emphasised the strategic value of tacit and explicit organisational knowledge. Organisational knowledge is seen as having an inherent competitive and commercial value; any loss or degradation of it will adversely affect a company's ability to operate in a normal manner (Alavi & Leidner, 2001; Davenport & Prusak, 1998; Hansen et al, 1999). Further, such losses can harm employee morale, customer confidence or have a direct impact on competitive performance. Some commentators have suggested that the protection of knowledge has not received adequate attention in the organisational security arena, despite its criticality (Ruighaver & Maynard, 2003; Bloodgood & Salisbury, 2001; Gold et al, 2005; Holsapple & Jones, 2005).

The next section of this paper reviews current ISRAs and explains their focus on information assets. Following this, there is an exploration of the crucial but overlooked role of knowledge which, as noted by Davenport (1998), is 'baked into' organisational processes. In we explore and unpack The connection between risk and knowledge through an illustrative case study is established, whereby we apply a typical information-security risk assessment methodology – OCTAVE-S – to a core organisational process. Finally, we discusses the deficiencies of security assessment in OCTAVE-S in terms of process knowledge and considers how a knowledge perspective could be incorporated in security-risk assessments in general.

Information Security Risk Assessment: The Asset Focus

An information-security risk assessment (ISRA) is a systematic method by which organisations can identify and protect information assets to achieve a desired level of security (Lichtenstein, 1996; Blakely et al, 2002). Risks to assets are identified in terms of confidentiality, integrity and availability. The criticality of each risk is rated according to potential impact and likelihood of occurrence. Risks to organisational assets are collated and then prioritised on the basis of criticality for further action (Roper, 1999; Alberts & Dorofee, 2004). An inadequate risk assessment process, or the absence of one, can lead to severely adverse consequences for organisations, such as a loss in reputation, legal issues or direct financial impact.

There are a number of popular information-security risk assessment methodologies in wide use in industry in Europe, the US and Australasia, at least, including FRAP, CRAMM, COBRA, OCTAVE, OCTAVE-S and CORAS (Peltier, 2001; Yazar, 2002; Alberts & Dorofee, 2004; den Braber, 2007; Dhillon, 2007). Although they differ in their make-up, order and depth of activities, they generally follow a three-stage pattern: context establishment, risk identification, and risk analysis (Whitman & Mattord, 2005; Shedden et al, 2006; Dhillon, 2007).

Context establishment considers the organisation's industry, structure, current security status, overall goals and long-term strategy. This stage allows for the scoping and focus of the rest of the risk assessment process for maximum effectiveness and to ensure that any risks inherent in the organisation's industry or line of business are identified. Activities that define the evaluation criteria for the assessment (eg., what constitutes a 'high' impact to 'reputation'?) are also performed during this stage (Alberts & Dorofee, 2004).

Risk identification concerns the systematic discovery and selection of the organisation's most critical information assets and then the identification of the threats and vulnerabilities of each of these assets. *Information assets* are typically defined as the IT resources, of infrastructure or information, that are of value to the organisation, including: data, hardware, systems and applications (Visintine, 2003; Alberts et al, 2003; AS/NZS, 2004). A *threat* is defined as a category of events which may present some form of danger to the critical information asset (Whitman & Mattord, 2005). *Vulnerabilities* are those security 'holes' or weaknesses inherent in the system that may present an avenue of attack against the information asset (Otwell & Aldridge, 1988).

Risk analysis concerns the determination of probability (chance of the threat event occurring) and impact (the cost of compromising the asset). The integration of the probability and impact will present the level of risk. Risks are then prioritised to allow assessors to determine whether each risk should be avoided, mitigated, transferred or accepted (Whitman & Mattord, 2005). Analysis and assessment of impact and probability can be performed through either quantitative or qualitative means, offering a range of metrics derived from mathematical equations and statistical modelling or qualitative indicators such as interviews and documentation (Alberts & Dorofee, 2004; Roper, 1999).

At the heart of all ISRA methods is the risk identification stage with its focus on the world in terms of assets that might experience threats or vulnerabilities and that might ultimately suffer loss or damage. This is the nature of

the risk identification that is conducted and it underpins the preceding context establishment, and severely constrains the following analysis of risk. It must also be stressed that there is a focus in current ISRA methods on technical assets, eg. hardware and software (Salmela, 2008) rather than the richer organisational elements of information systems that include people, knowledge and practice (Dhillon & Backhouse, 2001).

Organisational Knowledge and Knowledge Protection

Organisational knowledge has long been recognised as a resource of strategic significance (eg. Davenport & Prusak, 1999) and the importance of knowledge management is now well established (Hansen et al., 1999; Zack, 1999). Here we will examine some of its aspects that are pertinent to conceptualising knowledge protection. First, we draw on Davenport & Prusak's widely cited definition of knowledge as 'a fluid mix of framed experience' that 'often becomes embedded not only in documents or repositories, but also in organisational routines, practices and norms' (Davenport & Prusak, 1999).

Behind this view is Polanyi's (1962) distinction between explicit knowledge (that which can be articulated and encoded) and tacit knowledge (that which remains in the 'minds of knowers'). Building on this distinction in his well-known SECI model, Nonaka (1994) describes four processes through which organisational knowledge is created via conversions between tacit and explicit knowledge: socialisation, externalisation, combination and internalisation. Socialisation is the process of sharing tacit knowledge (eg. via means apprenticeships). Externalisation involves the articulation of tacit knowledge into explicit knowledge (this corresponds to what some sources term codification). Combination is the process of converting explicit knowledge into sets of more complex explicit knowledge. Lastly, internalisation involves the process of converting explicit knowledge back into tacit knowledge.

While some have argued that knowledge is inherently personal and contextually-bound, others have explored the idea of knowledge as distributed phenomenon beyond the single individual human mind. In this regard, knowledge and knowing are seen as attributes of groups of individuals. The notion of communities of practice (Brown & Duguid, 1991) builds on the idea of knowledge as an attribute of a collective. Indeed, the theory of distributed cognition proposes that knowledge extends beyond the mind of an individual and includes interactions between people, material and environmental resources (Hutchins, 1991; Hollan et al, 2000). In this regard, individuals transfer some of the cognitive load onto the resources and environment, embedding information and knowledge in artefacts (eg. displays, notice boards, documents).

Despite the wide recognition of the value of knowledge to organisations, research into knowledge protection has been described as 'thin' (Gold et al, 2001). A possible reason for this is an inherent conflict between the premises of knowledge sharing (typically assumed in the knowledge management literature) and those of knowledge protection. On the one hand, knowledge sharing is viewed as a valuable activity that gives individuals access to knowledge that will assists in completing tasks (Alavi & Leidner, 1999; Fischer & Ostwald, 2001). On the other hand, from an information security standpoint, such sharing activities bring new risks of knowledge falling into the 'wrong minds' and providing a means to inflict harm on the organisation or on its customers and partners (Whitman & Mattord, 2005).

Some commentators have explored the possibility that current knowledge management philosophies are inherently insecure. The need for a supportive 'knowledge control' process has been outlined – ensuring that required knowledge processors and resources are available in sufficient quality and quantity, subject to security requirements (eg., by Holsapple & Jones, 2005). However, there are few suggestions in the literature as to how this could be accomplished. One of the few proposals, for example, is 'data sanitisation' of codified knowledge (Oliveira & Zaiane, 2003). Holsapple & Jones (2005) also argue that using security risk management standards may be an option and that access controls and monitoring facilities should also be provided. This would aim to limit availability of the knowledge and increase confidentiality from a security perspective. Bloodgood & Salisbury (2001) further discuss the merits of limiting access to knowledge and ensuring that there is no single point-of-failure (eg., no single employee knows all).

Gold et al (2001) have broadly suggested several means to protect knowledge, including asset-based ISRAs. However, asset-based risk assessment methods typically deal only with explicit knowledge (i.e., that which can be codified or articulated in documents, databases, etc). Yet, as argued above, not all knowledge can be rendered explicit, nor can explicit records of knowledge often be understood and applied without related tacit knowledge. Indeed, it has been argued that the two forms of knowledge are 'mutually constituted' (Tsoukas, 2004) and explicit encodings must be actively re-contextualised on application.

In summary, it appears that is not sufficient to augment current ISRA methodologies merely by including the identification of 'knowledge assets' in the form of databases, or even key people. Indeed, a complex organisational process tends to rely on both explicit and tacit knowledge of various individuals and networks of experts. Therefore, understanding the full spectrum of risks associated with a particular process extends considerably beyond individuals and information assets alone. This line of thought suggests that if we wish to consider knowledge as a possible source of risk, the asset-based risk identification approach is likely to be insufficient.

Issues in Asset Identification: An Illustrative Case Study

In this section, we present a brief case study that illustrates the limitations of current information-security risk assessment methodologies that result from neglecting from the knowledge aspects of processes and their relationship with possible risks.

Research Methodology

Our approach was first to apply a traditional information-security risk assessment methodology to identify critical information assets in a knowledge-intensive work process involving IT. Importantly, we then went beyond the traditional method - through the use of narrative and a qualitative analysis based on holistic interviews – to probe the role of critical process knowledge that was not explicitly addressed by the methodology. The use of an exploratory case study is appropriate in this under-researched area (Eisenhardt, 1989; Benbasat et al, 1987; Yin, 2003) and its small size is sufficient to illustrate the knowledge gap in current ISRA methodologies without attempting to fully chart that gap.

A company called SoftCo (a pseudonym) was selected for study on the basis that it (a) had a need for ISRA because security was of significant concern to its managers, and (b) presented a sufficiently complex yet understandable process for analysis.

OCTAVE-S was selected as an ISRA methodology for study. Developed by Carnegie Mellon University and applied throughout industry, OCTAVE-S is a variant of the OCTAVE (Operationally Critical Threat and Vulnerability Evaluation) method, geared specifically for small-medium enterprises. Consistent with our review of such methodologies, the OCTAVE-S risk assessment model flows through the three phases of context establishment, risk identification and risk analysis of the desired risk treatment plans (Alberts et al, 2003). OCTAVE, applied notably in the healthcare industry and the US military (West et al, 2002), is representative of ISRAs used by industry more generally.

Data was collected from the security managers of the organisation using a workshop approach consistent with the OCTAVE-S methodology. The participants for these workshops were the security managers for the firm, described here as the 'Security Head' and the 'Co-Security Manager'. The Security Head is a team leader of the Managed Services division which maintains the infrastructure and environment. He holds in-depth knowledge of SoftCo's systems, security awareness and culture of the organisation. The Co-Security Manager is a systems administrator and supports the Security Head. These participants were selected due to their involvement in the area of study and because they provided different levels of domain knowledge and input based on the requirements of the OCTAVE-S methodology. For the purposes of this study, we focused upon a core business process for the assessment. The backup process was selected as it is common in organisations and is an extremely important security process designed to ensure that losses of data and hardware can be recovered (Stair & Reynolds, 1999). The workshops consisted of the participants answering structured questions and filling forms as per the OCTAVE-S method. Interviews were recorded and transcribed. A follow-up interview with the participants was conducted after the OCTAVE-S results were analysed in order to further explore points of interest and participant knowledge. This interview involved specific questioning and probing around the information assets discovered through OCTAVE-S and the knowledge requirements of the process.

The SoftCo Case Study: the Backup Process

SoftCo is an Australian software house and service provider employing 60 people. It provides business-to-business 'middleware' services for large manufacturing and retail firms. These firms integrate their back-end systems with SoftCo's so that customer orders and invoices to suppliers are routed through SoftCo's 'production environment' where they are translated into formats readable by the relevant other party's systems. Of particular importance are SoftCo's large set of 'translation maps' that perform the conversions of various document formats between customer and supplier organisations. These 'irreplaceable' translation maps are the core of SoftCo's business.

SoftCo's vital backup processes store copies of both customer data and the translation maps. Every document that has ever passed through SoftCo is preserved. The backup is largely automated through scripts which push data from the live production environment to an off-site backup server. The Security Head creates workable versions of these scripts initially; but, significantly, they must be continually 'tweaked' thereafter, meaning polished and refined for efficiency and effectiveness. The backed-up data (including historical backups going back to the founding of SoftCo) is of high ongoing value given that problems occur that require data restoration weekly.

Asset Identification Using OCTAVE-S

The OCTAVE-S assessment identified five critical information infrastructure assets for SoftCo's backup process, including three data assets, a personnel asset and an application asset. The results of the asset identification process are summarised in Table 1.

Asset Category	Asset Identified	Description
Data	Production data	The data that passes through the live Production Environment, including scripts, spreadsheets and directory structures. This is the data to be backed up.
	Backup files – Translation maps	The backed-up translation maps (text-based scripts that 'translate' the contents of invoices and order forms into a format readable by the customer organisation's partner company). Stored off-site at a remotely-managed data centre.
	Backup files – Live data	The backup files of the data that flows through the Production Environment. Includes order forms, invoices, SQL database back ups, scripts and e-mails. Stored off-site at a remotely-managed data centre.
People	Security Head and Co-Security Manager	The information security managers of the organisation. If they are not on-hand, the back up process will fail.
Applications	Backup scripts	The backup scripts are proprietary software that points the data from the live Production Environment to the backup storage locations.

Table 1. Critical Asset Identification Results

These assets satisfy the OCTAVE-S requirements for asset identification and critical asset selection, and are a common series of results from typical asset identification processes.

Case Study Results: Critical Process Knowledge Identification

The OCTAVE-S asset identification process suffered from the problem examined previously in that it did not identify a crucial area of vulnerability, which is knowledge about the backup process. Though questions were asked by the method on what knowledge a person might have ('what special skills or knowledge are provided by this person?'), it serves as a justification for the person to be considered a critical asset. The knowledge itself is not the target of an asset identification process nor does it surface any information about what those people know. For SoftCo, the security managers were identified as being information assets whose availability needed to be preserved. However, after analysis of the workshop transcriptions and follow-up interviews, a deeper understanding of the backup process and its key assets emerged related specifically to the key knowledge required in order to keep the process operational. It became apparent that while OCTAVE-S did identify the higher-level information assets critical to the ongoing operations of the backup process, it did not identify process knowledge. Key individuals were identified and described by the methodology, but not the richness and variety of their knowledge, including that which is tacit and explicit, and also the security managers' network of knowledge. If this knowledge was lost or not on hand for application, the process would fail with the potential for a 'catastrophic' (as reported by the participants) failure of the backup process.

Key Individuals

OCTAVE-S identified key personnel as part of its information asset identification process. As has been described, the security managers (who we refer to as their job titles Security Head and Co-Security Manager) are responsible for the monitoring and maintenance of the backup process to ensure its continued operation. Despite extensive automation of the process through the backup scripts, the security managers' interventions were still required. The Security Head outlined that their maintenance of the process was important.

Specifically, the application of OCTAVE-S enabled the risk-evaluation team to identify that the security managers were essential for the process. Otherwise the backup process itself would fail. Again reported by the Security Head:

> *If we weren't here, the backups would continue to function and everything would continue to work – if nothing stopped – it would run for about three weeks, and then stop because the server we back up to fails. It will get full. So there's some things where we actually manually go in and delete because you're not prepared to trust another process to automatically delete them.*

OCTAVE-S identified that in order for the backup process to remain in operation, the availability of the security managers would need to be preserved. The security managers would need to be in contact with and remain in

control of the process through monitoring its functions, which they currently do through the use of e-mail and mobile phone alerts. Additionally, they have enacted an informal mechanism of staggering leave and absences.

Distributed Knowledge (Held Collectively)

However, OCTAVE-S did not consider that knowledge is not only held independently but shared between a work team following the theory of distributed cognition. The security managers do not work as individuals with this process: rather, they exist in a partnership, sharing knowledge and expertise. For example, when we asked about their documentation practices, the Security Head responded:

> *Yeah, very little documentation – I set up a lot of it, [the Co-Security Manager] maintains it and sets up stuff now. I trained [the Co-Security Manager] into it and he's worked it out and so it's pretty much the result of both of our thinking.*

This actually states that they have applied mutual thought, knowledge and expertise to the process. While either manager independently is important for the backup process, the current methods and operations are a product of both of their minds. Therefore, their distributed knowledge is just as important as the individually-held knowledge of each manager.

Individually-Held Tacit Knowledge

As individuals, the security managers have knowledge that is very difficult to articulate. They have built the process and its components through incremental customisation of the backup scripts over the course of many years. When discussing the nature of the backup process itself, the Security Head outlined that while it would be possible to document some elements of the process, it would not be possible to capture everything due to its 'messy' complexity:

> *Either one of us is critical. But if both of us weren't here, it would be difficult for someone else to come in and work out what we're doing, because it's so scripted and script-based and this process interacts with that process, so it's almost like spaghetti, to a degree. I mean it's very robust spaghetti, but if you had to flow-chart it all out, it would be a very messy flowchart.*

This is further reinforced by the Co-Security Manager's comments on their attempts at documentation and capturing their knowledge in codified form. When discussing how much could actually be clearly articulated into procedures, the Co-Security Manager explained that documentation could cover the general perspective, as he explained:

> *You still need to apply a level of understanding of the environment to understand how it works before; documentation will only tell you, like… There's only so much you can do, maybe 70% of it but then you need the experience outside of it to understand how it works.*

The additional 30% comes from their own experience and would be difficult to articulate. Though this is a routine process, the participants surprisingly reported that the adhoc nature of this process' development and the organic organisational environment it resided in left a number of tacit subtleties that would be difficult to articulate clearly or describe.

Therefore, there exists a great deal of tacit knowledge within the backup process that relates to an understanding of its complexities and how it actually operates. OCTAVE-S does not capture this tacit component. Consequently, methods to increase availability of this knowledge (to prevent process failure) are not suggested by the methodology either.

Individually-Held Explicit Knowledge

Despite the tacit nature of much of the critical process knowledge, the security managers also expressed their desire to attempt codification of explicit knowledge and the importance of doing so. For example, the Security Head discussed the explicit knowledge of the backup process held by members of the organisation that have left.

> *… so very little of it was ever written down, [because] there was no need to formalise it because you knew what everybody was doing anyway. Then someone would leave and they'd take all that, what should have been written down and documented away with them. Less of a problem now because documentation actually takes place because the lines of communications are poorer, so it has to be documented for people to know what's going on, if you get what I mean…*

As the Security Head explained, individuals retain key explicit knowledge. This explicit knowledge can and should be codified to ensure the ongoing availability of critical process knowledge and subsequently the operation of the backup process.

As has been demonstrated SoftCo's backup system is complex, built upon backup scripts that have been extensively altered and customised over a long period. However, while OCTAVE-S identified the security

managers as critical information assets, it is actually their knowledge and their network of knowledge that should be considered key. While the managers hold individual knowledge, they actually operate within a network, where various aspects of the backup process are, as stated by Security Head '*a product of both our thinking*'. However, their individual knowledge is also of critical importance, including that knowledge which is tacit and difficult to qualify and explain ('*it would be a very messy flowchart*') and that knowledge which is explicit and can be codified ('*it has to be documented for people to know what's going on*').

Discussion

Using the SoftCo case as illustration, we now return to the central issue of the limits of traditional ISRAs like OCTAVE-S, and consider the possible contribution of a knowledge management perspective. According to the analysis of the backup process produced by OCTAVE-S (see Table 1), the risks facing SoftCo can be neatly partitioned into the threats and vulnerabilities associated with five separable assets: data (production, translation maps, and live); people (Security Head and Co-Security Manager); applications (backup scripts). If any one of these assets is lost or damaged the organisation will suffer, and the total risk becomes the summation of the risks across the five categories. In this view, the role of knowledge is recognised only indirectly through the identification of 'key individuals' as assets that might suffer loss. In contrast, a knowledge management perspective offers a very different view. First, a risk analysis might identify the various types of knowledge on which the process depends. Second, and following on, it is important then to distinguish between explicit and tacit knowledge because they present different kinds of risk, one being documentable while the other is not. Third, it is also productive to recognise the way that these tacit and explicit forms of knowledge are 'mutually constituted', one depending on the other for its meaning and value. Fourth, attention must be given to the way knowledge is embedded in practices; so it might not be observable or transferable away from the relevant task. And fifth, attention must also be paid to the distributed nature of knowledge, or the way it is typically co-produced and co-utilised by a collective. Following these insights, a knowledge management perspective exposes the way that the different asset categories produced by OCTAVE-S are not separable but instead are deeply intertwined in the way they contribute to ongoing work and present risks.

An overall observation in our case study, and consistent with the knowledge management perspective, was that knowledge only emerged as a key area of vulnerability when focusing on an organisational process. Without a process viewpoint, asset-focused ISRA naturally down-play the subtlety and value of each individual's knowledge. Neglect of the role of knowledge seems to be right at the heart of the way ISRAs construct the world and then observe it. Although Australian security risk management guidelines and handbooks do suggest that knowledge should be identified as an extension of a 'staff' or 'people' asset category, this view is not uniform across ISRA methods or organisational implementations of these methods (AS/NZS, 2004; Shedden, 2005; DSD, 2007). In particular, tacit knowledge, which is impossible to articulate or codify, but which inherently forms part of successful execution are only likely to emerge from a process perspective in the analysis.

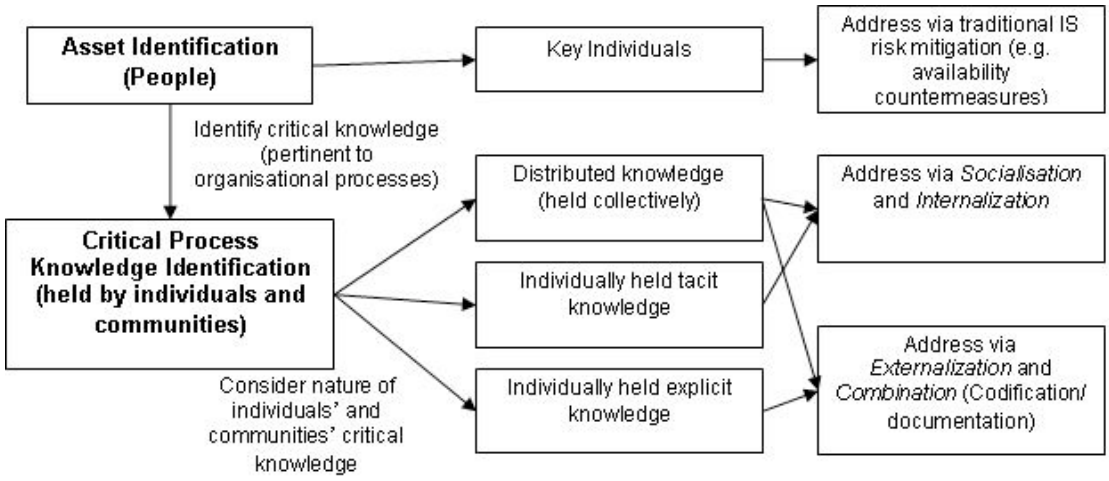

Figure 1: Incorporating a knowledge perspective in information-security risk assessment

Moving beyond the identification of key people and traditional information security risk mitigation (the asset based approach), an augmented methodology should identify critical knowledge that is pertinent to typically core or at least critical organisational processes. In this regard, tacit and explicit individual knowledge and distributed knowledge that is pertinent to the process should be identified. In the case study, the two managers could be considered as a small community of practice. As with asset-based approaches, risk mitigation strategies for critical knowledge should form part of a knowledge-sensitive risk assessment. This would include strategies

such as socialisation and internalisation (cf. Nonaka, 1994), which is the basis by which a degree of tacit knowledge could be shared. For example, having an apprentice 'shadowing' the key managers could be one strategy to spread the intimate knowledge they possess of the backup process to more people. Explicit knowledge relating to the process should also be considered. In this regard, strategies such as documentation and systemisation could be appropriate mitigation mechanisms to codify and combine explicit knowledge about the process. This would make it easier for a newcomer to gain insight about the backup process. Indeed, documentation of the backup process, at least those aspects which could be articulated explicitly, should be available from a quality assurance and general risk management perspective.

It should, however, be recognised that tacit and explicit knowledge are not easily separated (Polanyi, 1967; Tsoukas 1996). To codify the knowledge of the backup process into procedures or documents would lead to a '*very messy flowchart*'. Given the managers tacit knowledge, an outsider to this process would still find it difficult to know when and why to follow the said procedures and when interventions in the automated process would be necessary. This underpins the need for a holistic knowledge perspective as part of an overall ISRA.

A combined information security and knowledge perspective would identify that knowledge of SoftCo's backup process (and the organisation's systems and security) is concentrated into only two minds (ie. too few people know too much). Few asset-based ISRA methodologies would identify this as a key vulnerability. If either of these employees grew disgruntled with their organisation, retired, or resigned, this would amount to a major threat for SoftCo's systems and processes generally, including the backup process. A combined perspective on this problem would facilitate action to control this risk. This may include tradeoffs between increasing the availability of knowledge (through knowledge sharing and externalisation, on a 'need-to-know' basis) cognisant of the risk of confidentiality breaches.

A practical example to establish the need for this model can be seen in the shift in the nature of the Australian electrical industry. Electrical power networks underpin modern society's integrated and interdependent infrastructure and have therefore been identified as critical infrastructure by most western nations. Much has been said regarding the vulnerability of power networks to cyber terrorism as they are typically controlled by Supervisory Control and Data Acquisition (SCADA) computing systems. Although it is quite possible that a computer virus may attack SCADA networks and subsequently cause a power outage, manual 'backup' control of power machines can always restore services. In the past, the SEC employed a small team of electrical inspectors across Australia that maintained power networks. These inspectors, by virtue of their intimate and long-standing association with the power systems and their knowledge of the vulnerabilities in the design of the system and its various safeguards, acted as a countermeasure to security threats (accidental and malicious) and the consistent operation of the system. Their small number and the tacit nature of the knowledge also made the spread of this knowledge less likely. Although the electrical inspectors were unlikely to have military style 'clearances', their small number and mutual associations kept critical knowledge in-house. However upon the privatisation of the electrical industry in Australia, the power infrastructure was divided and sold. Electrical inspectors came under the employ of different companies. This prevented electrical inspectors from sharing knowledge of potential security threats and countermeasures of the system as a whole. Electrical inspectors with holistic knowledge of the old system are being replaced by newcomers with little or none of the invaluable 'bigger picture'.

As a practical example, the adoption of a knowledge-based approach as seen in Figure 1 would have beneficial properties when running security risk assessments of critical infrastructure. Such assessments are run frequently given the highly important nature and national security impacts of utilities and services bound by the 'critical infrastructure' term. However, given the intense level of in-house knowledge and difficulties in current methods providing information and knowledge sharing infrastructures (NARUC, 2007), an ISRA operated under Figure 1 would highlight and preserve the intricate tacit knowledge to assist in the education of newcomers under the new system.

In summary, we suggest that a knowledge perspective could, and should, be incorporated into ISRAs. We believe the identification of core knowledge can occur through a business process-based focus. As illustrated by our second analysis of the case study, this can occur through conducting qualitative interviews with relevant staff members in the context of key business processes. Tools such as business processes mapping and rich process descriptions could further help to facilitate the identification of core knowledge and key knowledge workers. Such a more inclusive approach to security risk assessment would thus help to identify what knowledge is drawn upon in the process and what knowledge needs to be protected to control operations.

Conclusion

Information-security risk assessments (ISRAs) are important for organisations as they are the means by which critical information assets are identified, their threats and vulnerabilities assessed and a level of risk assigned and prioritised for future action. ISRA methodologies, however, mostly adopt an asset-based focus. By means of an illustrative case study, we have demonstrated the inherent limitations of a typical risk assessment methodology by exploring what it misses: critical knowledge of staff in the context of a complex but important organisational process.

To address such shortcomings in ISRAs, we argue for a focus on organisational processes and the inclusion of a knowledge perspective as a more encompassing approach for organisations to assess their overall security risk. Existing methods identify key individuals as 'critical information assets'. However, we argue that the individual's knowledge is of critical importance, as well as that knowledge which is distributed among a team or held collectively. We propose that protective techniques, for example the codification of explicit knowledge, or purposeful knowledge-redundancy and succession planning for tacit knowledge should be considered as part of the remedial action to protect critical organisational knowledge in a systematic and transparent manner.

We have illustrated that current ISRA methods do not identify security risks associated with knowledge in organisations. This implies that the unauthorised disclosure, modification and interruption or destruction of critical knowledge does not figure in the practice of security risk assessments. Although this paper has focused on availability of knowledge assets, further research must be conducted on the confidentiality and integrity of knowledge in security risk assessment methodologies.

References

Alberts, C. and Dorofee, A. (2004). *Managing Information Security Risks*. Pittsburgh, Mellon Software Engineering Institute.

Alberts, C., Dorofee, A., Stevens, J. and Woody, C. (2003). *Introduction to the OCTAVE Approach*. Pittsburgh, Carnegie Mellon Software Engineering Institute.

Alavi, M. and Leidner, D. E. (1999). "Knowledge Management Systems: Issues, Challenges and Benefits." *Communications of the Association for Information Systems*, vol.1, no.7.

Alavi, M. and Leidner, D. E. (2001). "Knowledge management and knowledge management systems: Conceptual foundations and research issues." *MIS Quarterly*, vol.25, no.1, pp.107-136.

AS/NZS (2004). *Information Security Risk Management Guidelines*. Sydney, Australia/ Wellington, New Zealand, Standards Australia/ Standards New Zealand.

Baskerville, R. L. (1991a). "Risk analysis: an interpretive feasibility tool in justifying information systems security." *European Journal of Information Systems*, vol.1, no.2, pp.121-130.

Benbasat, I., Goldstein, D. K. and Mead, M. (1987). "The case research strategy in studies of information systems." *MIS Quarterly*, vol.11, no.3, pp.369-386.

Blakely, B., McDermott, E. and Greer, D. (2002). "Information Security is Information Risk Management." *NSFW '01*, Clourcroft, New Mexico, USA.

Bloodgood, J. M. and Salisbury, W. D. (2001). "Understanding the influence of organizational change management strategies on information technology and knowledge management strategies." *Decision Support Systems*, vol.31, no.1, pp.55-69.

Brown, J. S. and Duguid, P. (1991). "Organizational learning and communities of practice: toward a unified view of working, learning and innovation." *Organization Science*, vol.2, no.1, pp.40-57.

Davenport, T. H. and L. Prusak (1998). *Working knowledge: how organizations manage what they know*. Boston, Massachusetts, Harvard Business School Press.

den Braber, F., Hogganvik, I., Lund, S., Stolen, K. and Vrallsen, F. (2007) "Model-based security analysis in seven steps – a guided tour to the CORAS method." *BT Technology Journal*, vol.25, no.1, pp.101-117.

Dhillon, G. and J. Backhouse (2001). "Current directions in IS security research: towards soci-organizational perspectives." *Information Systems Journal*, vol.11, no.2, pp.127-153.

Dhillon, G. (2007). *Principles of Information Systems Security: Text and Cases*. Hoboken, NJ, John Wiley & Sons, Inc.

Dubin, R. (1969). *Theory building*. New York, Free Press.

DSD. (2007). *Australian Communications-Electronic Security Instruction 33 (ACSI 33)* Handbook 3, RISK MANAGEMENT.

Eisenhardt, K. M. (1989). "Building Theories from Case Study Research." *The Academy of Management Review*, vol.14, no.4, pp.532-550.

Fischer, G. and Ostwald J. (2001). "Knowledge Management: Problems, Promises, Realities, and Changes." *IEEE Intelligent Systems*, January/ February 2001, pp.60-72.

Gold, A. H., Malhotra, A. and Segars, A.H. (2001). "Knowledge Management: An Organizational Capabilities Perspective." *Journal of Management Information Systems*, vol.18, no.1, pp.185-214.

Grover, V. and Davenport, T.H. (2001). "General perspectives on knowledge management: Fostering a research agenda." *Journal of Management of Information Systems*, vol.18, no.1, pp.5-21.

Halliday, S., Badenhorst, K. and von Solms, R. (1996). "A business approach to effective information technology risk analysis and management." *Information Management & Computer Security*, vol.4, no.1, pp.19-31.

Hansen, M. T., Nohria, N. and Tierney, T. (1999). "What's your strategy for managing knowledge?" *Harvard Business Review*, March-April, pp.106-116.

Hollan, J., Hutchins, E. and Kirsh, D. (2000). "Distributed Cognition: Toward a New Foundation for Human-Computer Interaction Research." *ACM Transactions on Computer-Human Interaction*, vol.7, no.2, pp.174-196.

Holsapple, C. and Jones, K. (2005). "Exploring Secondary Activities of the Knowledge Chain." *Knowledge and Process Management*, vol.12, no.1, pp.3-31.

Hutchins, E. (1991). Chapter 13: The Social Organization of Distributed Cognition. *Perspectives on Socially Shared Cognition*. L. B. Resnick, Levine, John M. and Teasley, Stephanie D. Washington, DC, American Psychological Association, pp.283 - 307.

Lichtenstein, S. (1996). "Factors in the selection of a risk assessment method." *Information Management & Computer Security*, vol.4, no.4, pp.20-25.

Maynard, S. and Ruighaver, A.B. (2003). Development and Evaluation of Information System Security Policies" in M.G. Hunter and K.K. Dhanda (Eds), *Information Systems: The Challenges of Theory and Practice*. Las Vegas, The Information Institute.

NARUC (2007). *Information Sharing Practices in Regulated Critical Infrastructure States; Analysis and Recommendations*, US Department of Homeland Security.

Nonaka, I. (1994). "A dynamic theory of organizational knowledge creation." *Organization Science*, vol.5, no.1, pp.14-37.

Oliveira, S.R.M. and Zaiane, O.R. (2003). "Protecting Sensitive Knowledge by Data Sanitization". *Third IEEE Conference on Data Mining*.

Otwell, K. and Aldridge, B. (1988). "The Role of Vulnerability in Risk Management." *1988 Computer Security Risk Management Model Builders Workshop*.

Parker, M., R. Benson, et al. (1988). *Information economics: Linking business performance to Information Technology.* New Jersey, Prentice-Hall Inc.

Peltier, T.R. (2001). *Information Security Risk Analysis*. Boca Raton, Auerbach.

Polanyi, M. (1962) *Personal Knowledge in M.Polanyi and H.Prosch*, Meaning. Chicago, University of Chicago Press.

Polanyi, M. (1967). *The Tacit Dimension*. London, Routledge & Kegan Paul Ltd.

Roper, C.A. (1999). *Risk management for security professionals*. Butterworth-Heinemann.

Salmela, H. (2008). "Analysing business process losses caused by information systems risk: a business process analysis approach." *Journal of Information Technology*, vol.23, no.3, pp.185-202.

Shedden, P. (2005). *Security Risk Management in Organisations*. Department of Information Systems. Melbourne, University of Melbourne.

Shedden, P., T. Ruighaver, A.B and Ahmad, A. (2006). "Risk Management Standards - the Perception of Ease of Use". *The 5th Security Conference*, Las Vegas, Nevada, USA.

Spears, J. (2006). "A Holistic Risk Analysis Method for Identifying Information Security Risks", in *Security Management, Integrity, and Internal Control in Information Systems*. Boston, Springer Boston, pp.185-202.

Siponen, M. T. (2005). "An analysis of the traditional IS security approaches: implications for research and practice." *European Journal of Information Systems*, vol.14, pp.303-315.

Stair, R. M. and Reynolds, G. W. (1999). *Principles of Information Systems*. Cambridge, MA, Course Technology.

Stoneburner, G., Goguen, A. and Feringa, A. (2002). *Risk Management Guide for Information Technology Systems*, National Institute of Standards and Technology.

Thompson, M.P.A. and Walsham, G. (2004). "Placing Knowledge Management in Context." *Journal of Management Studies*, vol.41, no.5, pp. 725 - 747.

Tsoukas, H. (1996). The Firm as a Distributed Knowledge System: A Constructionist Approach. *Strategic Management* Journal, vol.17, pp.11-25.

Tsoukas, H. (2004). *Complex Knowledge: Studies in Organizational Epistemology*. Oxford, Oxford University Press.

Visintine, V. (2003). *An Introduction to Information Risk Assessment*, SANS Institute.

West, S., Crane. L.S. and Andres, A.D. (2002). *OCTAVE-DITSCAP Comparative Analysis*. Fort Detrick, Fredrick, U.S. Army Medical Research and Material Command.

Whetten, D.A. (1989). "What Constitutes a Theoretical Contribution?", *The Academy of Management Review*, vol. 14, no.4, pp.490-495.

Whitman, M. E. and Mattord, H. J. (2005). *Principles of Information Security*, Thomson Course Technology.

Yazar, Z. (2002). *A qualitative risk analysis and management tool - CRAMM*, SANS Institute.

Yin, R. (2003). *Case Study Research, 3rd Edition*. Thousand Oaks, Sage Publications.

Zack, M. (1999) "Developing a knowledge strategy", *California Management Review*, vol. 41, no. 3, pp.108-145.