Cooperative Information Security Knowledge:
Content Validation and incentives to contribute

Florian Stahl, Simon E. Parkin, Aad van Moorsel
School of Computing Science
Newcastle University
Claremont Tower
NE1 7RU
Newcastle-upon-Tyne, UK
florian.stahl@alumni2010.ncl.ac.uk, {s.e.parkin, aad.vanmoorsel}@ncl.ac.uk

Abstract

The aim of this work was to examine the attitude of Chief Information Security Officers (CISOs) towards sharing knowledge with other CISOs in general and in particular the potential characteristics of a collaborative security knowledge sharing tool, which would simplify and encourage knowledge sharing. Building on this, this study intended to establish which key features such a knowledge sharing tool should provide in order to be accepted and actually used by security managers and potentially improve business performance. In this regard, content validation and incentives to contribute to such a system have been identified as key issues. They were investigated by interviewing three CISOs regarding the current state of knowledge organisation and sharing. The interviews were then transcribed and analysed using an explorative method. The analysis identified learning from each other as most important incentive for knowledge sharing and authorship as the main factor regarding credibility of contribution. From this it followed that such a sharing tool should demand users to register with the system in order to give credibility to their knowledge. However, since potentially sensitive business data would be shared, users should be given a choice of whom they would like to be able to access their contribution. Furthermore, it became clear that different levels of information detail should be provided for managers and technicians. Finally, the whole system needs to be managed in order to administrate users and maintain security and data integrity.

Suggested keywords

INFORMATION SECURITY ONTOLOGY,
CONTENT VALIDATION,
KNOWLEDGE SHARING INCENTIVES

1 Introduction

Over the last decade much research has been conducted into the need for and the characteristics of Information Security Knowledge Ontologies (ISKOs or SKOs). For SKOs have only relatively recently been recognised as research subject, it is of no surprise that many SKOs and SKO Editors are still in early stages and mostly domain specific (Blanco et al. 2008). Generally, these prototypes require the user to know how the underlying ontology is organised in order to be able to use it correctly. To overcome this shortcoming Mace, Parkin and van Moorsel (2010) developed an ontology editor which can be used more intuitively and without the need to understand the underlying concept of ontologies. Their aim of building what is referred to here as a "Collaborative Information Security Knowledge Ontology Editor" (CISKOE), which could potentially be used by security personnel across organisational boundaries, inspired this study to further examine how CISKOEs could provide more value to businesses by offering functionality to work jointly over the internet.

Comprehensive research on tools enabling collaborative work over the internet, namely Collaborative Information Tools (CITs), has been conducted. One of the most common forms of CITs are Wikis, and particularly Wikipedia (the largest Wiki on the Internet) has been extensively examined. However, it has not been examined, whether findings on Wikis, such as motivations of contributors and content validation, can be transferred to a highly domain-specific CIT such as one for security knowledge. Particularly interesting is the fact that CISKOEs are intended to be used across organisational boundaries, as this might imply biased usage.

In light of this, the present study intended to examine the attitude of CISOs towards sharing knowledge with other CISOs in general and the potential characteristics of a CISKOE in particular. To do this, interviews were conducted which targeted the current working practices of information security managers and their opinions about a potential knowledge sharing tool. Gaining insights into the working life of CISOs revealed what they need from or expect of a knowledge organising and sharing tool. If such a CISKOE could be developed not only can a CISO's day-to-day work be advanced, but also the IT security management decision-making can be enhanced owing to a more comprehensive knowledge base.

The rest of this work has the following structure: Section 2 discusses related works and adds background information. Thereafter, Section 3 will describe how the research was designed. Section 4 then presents and discusses the findings discovered. Lastly, the work will be concluded in Section 5 by giving recommendations for the development of a CISKOE and outlining possible future research into this subject area.

2 Related work

This section expounds the current literature relevant to this research. It is structured in three parts. The first part covers Security Knowledge Ontologies (SKOs), which are models that depict the knowledge of a certain area of expertise in a structured way. In the case of information security knowledge, this is achieved by modelling security relevant objects, their attributes and relationships between each other into an ontology (Blanco et al. 2008; Raskin et al. 2001). Different objects or assets have been determined as relevant by different authors. For example, Parkin, van Moorsel and Coles (2009, p. 5) define an asset as: "identifiable information artefact which is of value to the organisation". As a practical example passwords can be identified.

Some ontologies have been developed with the aim of being usable collaboratively by many users. For this reason part two discusses Collaborative Information Tools (CITs), which enable multiple users to work jointly on documents over the Internet. At first, this will be generally and later, in more detail, covering wikis (Wikipedia and wikis in a corporate context) as specific CITs. Lastly, the research questions for this work are derived from the discussions in a concluding section.

2.1 Security Knowledge Ontologies

2.1.1 Review of Ontologies

In 2008, Blanco et al. (Blanco et al (2008)) reviewed 28 works concerned with IT security ontologies, identifying 17 as IT security ontologies. They concluded that there was a need for a holistic SKO in the IT-security community that had not been fulfilled, and that the situation could only change if the community combined its efforts and ontologies. In this work the characteristics of the CISO community are examined to support these aims.

There are numerous works that have sought to produce a general SKO (e.g., Tsoumas and Gritzalis (2006), Fenz and Ekelhart (2009), Parkin, van Moorsel and Coles (2009)). Tsoumas and Gritzalis (2006) were motivated to gather security knowledge from various sources in one place. They considered "direct" sources (e.g. standards, risk analysis reports) and "indirect" sources (e.g. best practices). The work further characterised IT-security knowledge into security requirements (the "*what"*), technical implementation (the *"how"*), and required actions (the *"do"*).

Another SKO has been developed by Fenz and Ekelhart (2009). Their ontology takes company specific information into account and is organised in such a way that assets (tangible or intangible) get assigned security relevant attributes and have certain vulnerabilities. The German "IT Grundschutz Manual" has been used as knowledge base for this ontology. The authors suggest utilising their ontology in the area of risk management.

An ontology which apart from organisations also takes human behavioural factors into account has been developed by Parkin, van Moorsel and Coles (2009). Potentially, this extends the realm of knowledge to human factors expertise as well as IT security. Their ontology aligns human behaviour and technical controls with the ISO IT security standard. To investigate the deployment of this ontology, Mace, Parkin and van Moorsel (2010) conducted semi-structured interviews with CISOs, finding justification for a collaborative ontology editor tool, since knowledge is currently gathered from various sources at various organisations to solve similar IT-security problems.

2.2 Collaborative Information Tools

There has been minimal examination of whether ontologies as described previously would be accepted and used by IT security managers. CITs aim to enable users to work jointly on documents utilising web technologies. In that way a CISKOE can be seen as a CIT, so much can be learnt from revisiting studies of existing CIT user communities.

2.2.1 Tools for online collaboration

In a comprehensive report for the German Federal Ministry of Education and Research, Schulzki-Haddouti and Lorenz-Meyer (2008) described the current developments in online collaboration. They state that the power of collaboration derives from the fact that a problem can be solved more effectively collectively by a group than by an individual (p 82). Surowiecki (2005) describes this phenomenon as the wisdom of the crowds. In this regard the openness of collaborative technologies appears to be particularly important. Here openness is referred to in two ways. Firstly technical openness, i.e. an open interface which can be used intuitively, secondly an openness regarding the users' possibilities to communicate, coordinate and manage, as well as share content (Schulzki-Haddouti & Lorenz-Meyer (2008), p. 20).

Schulzki-Haddouti and Lorenz-Meyer (2008) identified four families of CITs: *"Questions and Answers tools"*, *"Blogs"* (Web logs, specifically those used in a cooperative environment); *"Wikis"*; and *"Shared Workspaces"* (shared software that enables users to work collaboratively online). Blogs offer only supportive functions, whereas "question and answer" tools exemplify the use of collective intelligence rather than shared manipulation of a knowledge base. Shared workspaces, though likely to grow in use, offer far more functionality than is actually needed for an ontology editor. Wikis, websites for collaborative editing of articles, however, are comparable to CISKOEs, as both systems are similar in user perception, despite having different underlying technological concepts. For this reason, the next section will discuss research conducted around Wikipedia (as the largest Wiki in use).

2.2.2 Wikipedia

The online encyclopaedia Wikipedia is the largest Wiki and also the largest CIT on the internet. All of its content is generated by the voluntary work of contributors referred to as Wikipedians (Johnson (2008); Kittur at el (2007); Wagner & Prasarnphanich (2007); Kuznetsov (2006); Forte & Bruckman (2005)). Owing to its immense size, much of the research on CITs has been conducted by examining Wikipedia.

Conducting a statistical analyses of different user groups' contribution on Wikipedia, Javanmardi et al. (2009) found that 80% of revisions are made by as little as 7% of the users, the majority of which are registered. Furthermore, they found a positive correlation between the quality of an article and user registration. This supports the hypothesis that Wikipedia seemingly relies on a great many users with good intentions to neutralise the negative effects of vandals. These results certainly show that it is fruitful and worthwhile to develop an understanding of user communities, just as in this work.

2.2.2.1 Incentives to contribute to Wikipedia

There have been many examinations into the relevance of incentives as relate to Wikipedia. Forte and Bruckman (2005) found that the incentives in Wikipedia resemble those of the scientific community as contributors want to *"collaboratively identify and publish true facts about the world"*. Kuznetsov (2006) found that over 50% of research survey respondents would add new content if something were missing (either create an article or add to an existing one) and more than 81% would correct errors if they found any. The willingness to contribute was also found to correlate with the frequency of a participant's use of Wikipedia. Many reasons for contributing to Wikipedia were also uncovered: "educate humanity", "make a difference", and "give back to the community". Less than 5% of the Wikipedians studied named selfish reasons such as bragging and building a reputation.

Reasons for not contributing discovered were: “lack of time”, “missing qualifications”, “unfamiliarity with the editing system”, and “laziness”. Building on her results, Kuznetsov extracted underlying principles such as altruism, reciprocity, community, reputation and autonomy (where Wiki-technology supports the last three points particularly well). Wagner and Prasarnphanich (2007) found that within in the Wikipedian community, altruistic motivations outweigh selfish motivation. These works of research build a picture of how the incentives for contributing to a shared knowledge repository can be varied and detailed, and in some cases personal. As such, the motivations of community members must be considered if any CIT is to be deployed with an eye to its sustainability.

Wagner and Prasarnphanich (2007) argue that the main reason for Wikipedia’s growth is its technology, which enables easy collaboration and makes even minor contribution worthwhile. The fact that Wikipedia offers several communication channels is, according to Forte and Bruckman (2005), another factor for the success of Wikipedia, where these channels enable Wikipedians to interact and establish credibility in various ways with greater efficiency.

Wagner and Prasarnphanich (2007) go so far as to introduce the term “wiki-magic” in reference to the relatively quick positive feedback on a user’s contributions. This “wiki-magic” was identified as one reason for contributing. This sense of “wiki-magic” implies that a CIT needs to offer users the right rewards if it is to be successful. Forte and Bruckman (2005) conclude that in order to sustain involvement, people should be given more than simple rewards. They believe that deeper commitment to the community can be achieved by offering the possibility for higher levels of efficiency, responsibility and influence within it.

2.2.2.2 Conflict and content validation in Wikipedia

Different communication channels are also important when resolving conflicts pertaining to the correctness of information and whether it should be published. In 2010, Kittur and Kraut (2010) analysed the effects of user communication (article talk and inter-user-communication), common mind sets (possibly achieved through policies) and workgroup structure on conflicts. To examine these correlations, reverts (the action of undoing changes to an article) were statistically analysed. The study examined over 6800 publicly available wikis, finding that Wikipedia was comparable to other Wikis, differing only in terms of policies and conflict correlation. This suggests that the results could be transferable to other CITs such as CISKOEs.

With regard to the conflict mediating mechanisms, Kittur and Kraut (2010) found that a group structure with a core set of leaders is most likely to produce quality articles but also more likely to revert previous work, i.e. cause conflict. However, they also found that group structure and communication reduce conflicts, chiefly because these mechanisms support the building of common mind sets.

2.2.3 Wikis in a corporate environment

Wikis have been successfully implemented by many companies (Pfaff & Hasan 2006), and studying Wikis in corporate environments offers insights into how relatively closed, specialised and competitive business environments can influence knowledge-sharing activities.

Incentives to contribute are critical to the viability of corporate Wikis. White and Lutters (2007) found that wikis would not be used by employees until their managers refused to answer questions and referred their staff to the

Wiki. When forced to, employees started to use the Wiki and eventually became editors themselves by correcting errors and starting to contribute. Similarly, Swee-Lin Tan and Zou (2009) found that two key factors increase the usage and contribution of corporate Wikis: users must be aware of their possibility to contribute, and minimal effort should be required to search and match knowledge content with their own. Given that familiarity with the CIT was also identified as a critical factor, it was concluded that staff training can also influence the acceptance of CITs in a corporate environment. In 2006, Pfaff and Hasan (2006) examined the reasons why one particular company did not use a Wiki as a knowledge management tool, identifying reasons such as fear of vandalism (i.e. corruption of information), lack of rewards for work, and legal concerns regarding intellectual property. These works demonstrate that a corporate Wiki needs to be considered not only in terms of the knowledge it contains, but also the business processes it is seen to support or hinder.

By statistically evaluating questionnaires sent to corporate Wiki users, Majchrzak, Wagner and Yates (2006) categorised the contributors of corporate Wikis into three groups: synthesisers (add content to existing articles), adders (create new articles) and commentators (comment to improve content). Furthermore, they named different incentives for the different groups. According to this, synthesisers are more affected by the impact that they can have, whereas the other two groups want to help the organisation and make their work easier. These findings illustrate that the benefits need not be personal, but that the benefits to the organisation are equally relevant, as may also be the case with a CISKOE.

When studying collaborative tools used between organisations, Schulzki-Haddouti and Lorenz-Meyer (2008) drew parallels between technologies which enable cooperation and public goods. In this sense, people who do not contribute might nonetheless benefit from the work of others or people might harm public welfare by acting selfishly. An issue closely related to this is the fear of stronger competition between cooperation partners and the question of who owns intellectual property.

2.3 Research Questions

Despite SKOs and CITs having been examined comprehensively in previous research, there has been no investigation of any tool that constitutes a combination of both. By considering the work of Mace, Parkin and van Moorsel (2010) and how it may be deployed, a set of research questions were identified:

RQ1: *What are the incentives for Chief Security Information Officers (CISOs) to share their knowledge and contribute to a CISKOE?*

RQ2: *What constitutes a credible piece of knowledge for a Chief Security Information Officer (CISO)?*

Answering these research questions informs our view of the knowledge-sharing activities of CISOs, potentially helping to produce recommendations to consider when developing a CISKOE. Establishing such an editor in a corporate environment can potentially improve business performance owing to optimised processes and more informed decisions.

3 Methodology

To answer the research questions, which concluded the last section, it was chosen to conduct interviews with CISOs as it was of the highest importance to get first hand information. For this explorative study individual semi-structured interviews were chosen for a number of reasons. Firstly, Interviews generally allow insights into an interviewee's opinion regarding a certain subject. Individual interviews were chosen because potentially highly confidential information was being discussed. Additionally in this way biased answers because of the presence of other CISOs could be excluded. Secondly the loose structure allowed for easy comparison between the interviewees' answers. It has to be mentioned though that the interviewees were given freedom in what to say when and they have only been interrupted when they went far off topic. Thirdly the required user base was a reason to choose individual interviews over group interviews. Senior management staff were required in order to provide the needed information. These managers, however, are generally highly engaged with their own business. For this reason it was chosen to work more intensively on fewer sources to allow an in-depth analysis of a few interviews rather than a shallow analysis of many. However, since generalisations can only be drawn from quantitative research, this study has to be seen as preliminary work for an even deeper analysis of this subject area in future work.

3.1 Sample

All CISOs interviewed were familiar with the research topic of Security Knowledge Ontologies (SKOs), which was advantageous when talking about possible features for a Collaborative Information Security Knowledge Ontology Editor (CISKOE). Nonetheless, when asked about their recent practices, interviewees might have tended to answer with the SKO in mind. During the interview and analysis process, no bias could, however, be identified.

The CISOs have a combined work experience of nearly 80 years. All of them have worked in IT security roles in large organisations in various different sectors, and some sectors were covered by more than one interviewee. The sectors included local government, higher education, military and defence, and the financial sector. In addition to this, some of them have worked as consultants and thereby gained insights into different companies in a way that would otherwise not have been possible.

It has to be confessed that three interviewees per se do not seem to be broad data. However, this was deemed a sufficiently large knowledge base for the purpose of this explorative study, given the breadth of experience amongst the participants.

3.2 Process Overview

Adopting the sociological approach the first thing to do was to review analytic and cultural categories to design the interview. To do this the current literature was reviewed, results of this review can be found in Section 2 related work. The review had to be done critically in order to not acquire preconceptions, i.e. the literature read had to be scrutinised to prevent assumptions about possible interview results, which then could potentially influence the analytic process. In the present study this was achieved by deriving the research questions from the literature. This is valid because deriving the research questions, in itself, is questioning the literature and finding questions not answered by it. Analysing cultural categories describes the action of the researchers, to think of

their own expectations and experiences. Therefore, in this case it will be referred to as expectations review. The expectations have to be analysed to realise preoccupations and prevent these from influencing the later analysis. At the same time, the expectations review prepares for the data analysis as it clarifies what is actually being looked for. In this case, potential recommendations were conceptualised, which helped to develop the interview questions as much as they helped to realise expectations. Knowing these expectations helped to actively analyse the content as objectively as possible.

The next step was the creation of the interview questions from the literature and expectations review. There are different opinions regarding the structure of semi-structured interviews. While some (King (1997)) recommend a loose structure to allow for a maximum of flexibility, others (McCracken (1988, p. 24)) recommend a structured questionnaire to organise the interview. The slightly more structured interview was chosen for a better interview organisation, but interviewees were neither interrupted nor redirected to other questions to maintain their flow of speech.

During the interviews it was ensured that all questions were understood correctly by visualising them on PowerPoint slides. The problem of status differences between interviewee and interviewer (King (1997); McCracken (1988, p. 25) was addressed by never questioning the status or the knowledge of the participants. To prevent the problem of a possibly low-status interviewer, two interviewers were present.

All interviewees allowed anonymous recording of their interviews. Recording is generally considered crucial to be able to gather all relevant information (King (1997), Morton-Williams (1993, p. 70); McCracken (1988, p. 41)). Nonetheless, the need for privacy was important as potentially sensitive matters were discussed. This is why the authors committed themselves to treat all information given with great confidentiality.

3.3 Analysing Content

Having a relatively small sample size allowed an in depth analysis. To achieve this, an approach suggested by McCracken (1988 p. 41 et seq.) was utilised to examine the data. Whilst it formalises the process it allows the researcher more freedom to work in a truly explorative manner, compared to more stringent methods of content analysis (e.g. grounded theory).

The actual process of analysing content according to McCracken (1988) consists of five steps. Firstly the interview transcripts are read carefully and relevant pieces of information are identified by approaching the transcript with a *"disingenuous wonder"*. This means to read as objectively as possible: even as if one had no assumptions about what has been said and its importance. Pieces of information relevant to the research questions then become *observations.* The information given is examined without relating it to other information in the same transcript, let alone other transcripts. The second step itself is three fold, analysing each observation on its own, in relation to the transcript and in relation to the literature reviewed. This step was slightly amended in that way that the relationship to the literature was omitted, for most of the observations were completely explorative. However, links to the literature are pointed out in Section 2. Thirdly, the observations are re-examined with particular focus on their interrelation. The fourth step is to find themes that cluster the observations of individual interviews. Lastly the themes generated in the previous step are clustered into interview theses. In this study the last two steps have been combined, as the groupings became evident in the

process of analysing the interrelations of all observations. Figure 1 depicts the analytical process adapted from McCracken (1988 p. 43)

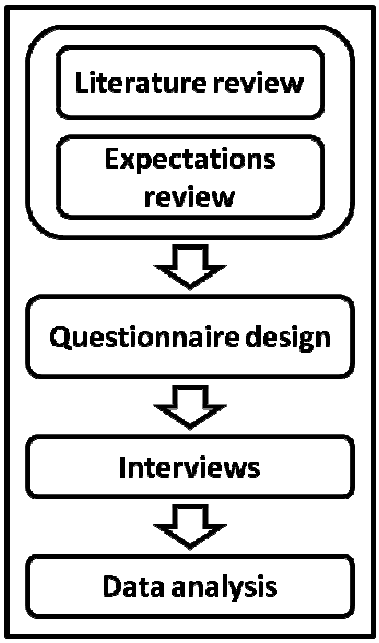

Figure 1: Research process adapted from McCracken (1988, p. 30)

4 Findings

During the analysis process 72 unique observations were discovered. These were grouped into 11 groups, which are listed below:

1) Knowledge organisation has to follow the organisation
2) Current state of knowledge organisation
3) Sources of knowledge
4) Credibility (of source / knowledge)
5) Anonymity
6) Factors for knowledge sharing
7) Factors against knowledge sharing
8) Role of senior management
9) Need for management of a collaborative tool
10) Relevant information about a contributor's organisation
11) Miscellanea

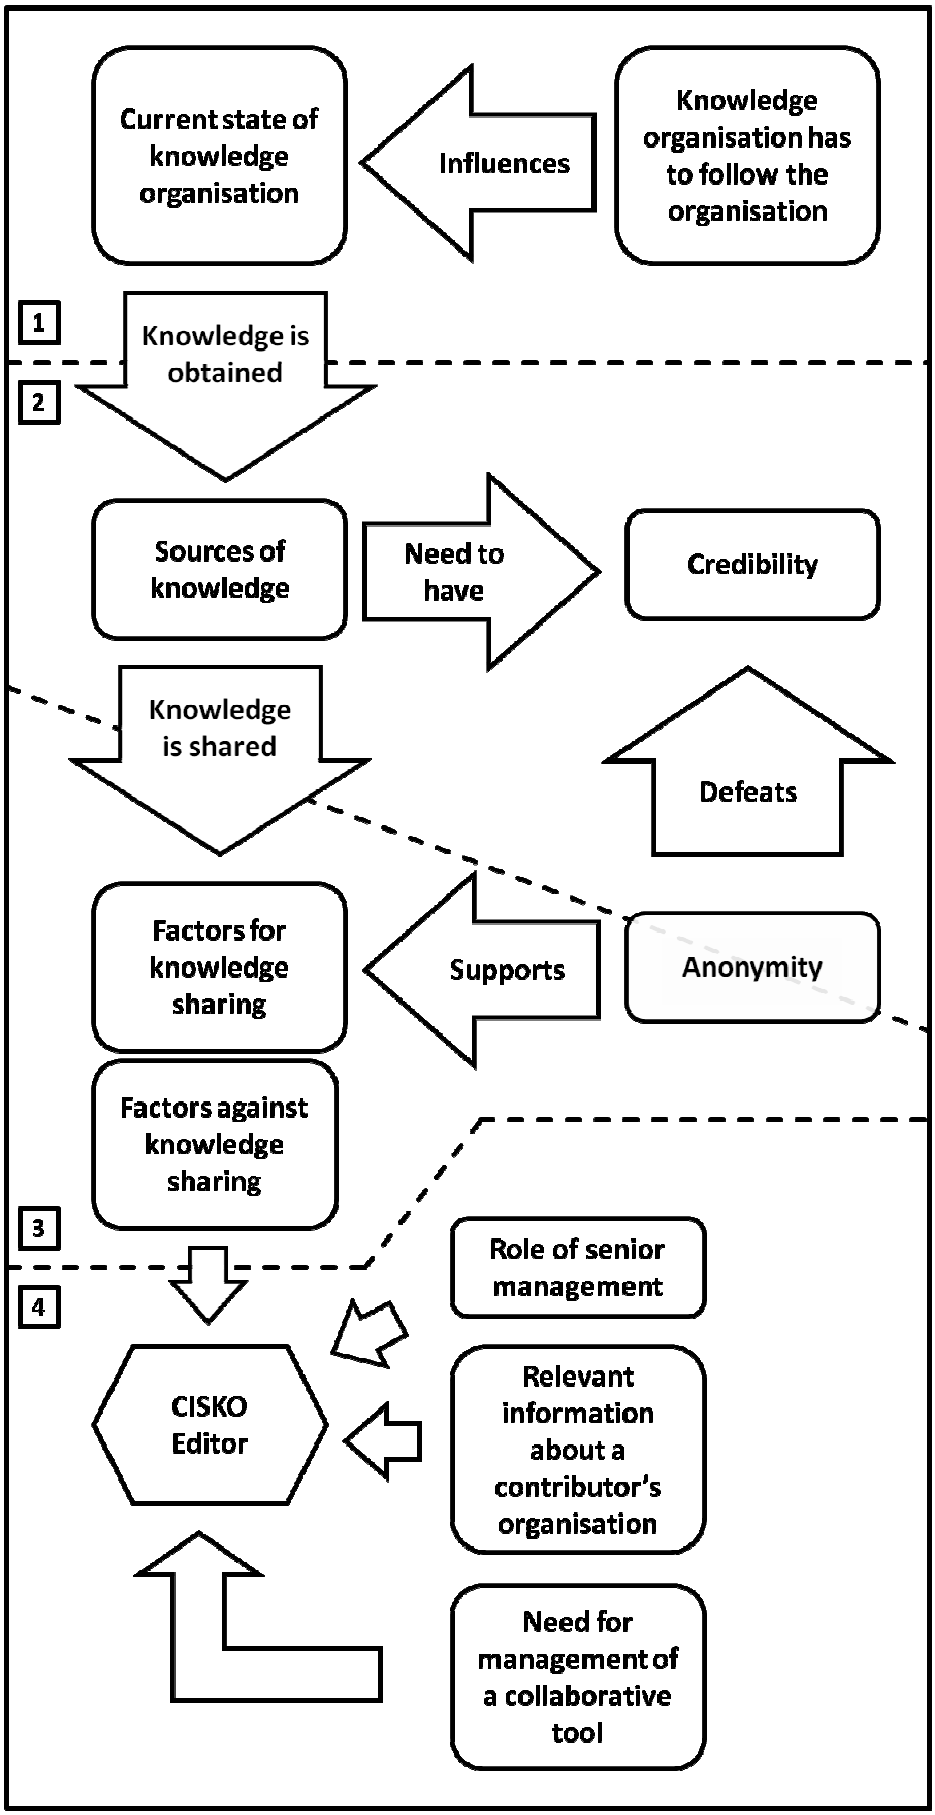

Figure 2: Interrelation of interview theses

The rest of this section is structured according to the group interrelations, which are depicted in Figure 2. Above all it was discovered that the organisation of knowledge has to follow the organisation. By building on this the current state of knowledge organisation could be examined. Further important findings include identification of the sources from which knowledge is obtained, and factors that constitute credibility of knowledge. As a particularly critical aspect, the issue of anonymity will be further examined. This is because the requirement for anonymity seems to defeat the need to assess the credibility of knowledge, while at the same time it facilitates knowledge-sharing. Building on this, factors which encourage or discourage knowledge-sharing will be discussed. Lastly, this section discusses the importance of senior management, relevant information about a contributor's organisation, and the need for management of collaborative tools with regard to Collaborative Information Security Knowledge Ontology Editors (CISKOEs). The "Miscellanea" group contains 4 observations that could not otherwise be categorised.

Note that within the following sections, transcribed quotes from participants appear in italics.

4.1 Current state of knowledge organisation

This section refers to Part 1 of Figure 2 and describes knowledge organisation and the influence of the type of business on it. Regarding knowledge organisation, three distinctly different approaches could be identified. They can be classified as "hardly-structured", "semi-structured" and "structured". From interviewee statements, it can be concluded that these are only three distinct forms on the scale from "not structured" to "structured".

The least structured organisation of knowledge can be described as "people knowledge". This way of organising knowledge is characterised by the fact that there are "*specialists within their own area"* who share their knowledge within the organisation when needed. Furthermore, in this case there are hardly any documents describing formalised standard procedures.

The semi-structured way of organising security knowledge is characterised by file-based organisation. These files contain standards (e.g. ISO 27001), best practices and catalogues of potential risks to IT Security. To maintain integrity to some degree, files are cross-referenced between each other. This has the disadvantage that one particular piece of information is potentially hard to find.

A completely structured approach could be found in the military and the financial sector where every process and procedure has to be documented in great detail. Generally, it seems that larger companies have tighter procedures and more formalised knowledge. One CISO stated that the advantage of structure is *"a) better security and b) more system up time",* while disadvantages could be identified in terms of flexibility.

Tools that support structured security knowledge organisation are integrated databases, which centralise all security relevant information of a company. This includes best practices, standards and evidence of how standards are enacted. It is in the nature of database knowledge organisation that all the information is cross-referenced and easily searchable. This integrated solution, however, has the disadvantage that the security on it has to be very high since all relevant information can be stolen with one successful attack. Potentially, this can be one of the biggest issues with a Collaborative Information Security Knowledge Ontology Editor (CISKOE) as well, for such a system centralises security knowledge of not only one but of various enterprises, making such a collaborative knowledge repository an attractive target for attackers.

During the research process it also became evident that knowledge organisation is *"not a case of one size fits all."* This includes the need for different layers of knowledge organisation for different roles or different types of organisations. Furthermore, it shows that a possible CISKOE has to address the needs of various kinds of businesses. The biggest differences are probably between the public and the private sector. Ostensibly the public sector (excluding the military) tends to lean more towards the less organised end of the spectrum whereas the private sector tends to lean more towards the more structured end.

4.2 Sources and credibility of knowledge

As is evident form Part 2 of Figure 2 knowledge has to be obtained from sources and these sources need to be credible in order to be accepted.

4.2.1 Sources of knowledge

Uniformity could be found in the types of sources used by IT Security mangers. The following not comprehensive set seems to include the most relevant sources:

- Standards such as ISO 27001
- Best practices repositories such as COBIT
- Offline and online articles
- Suppliers of infrastructure and software
- Conferences, informal meetings and peers

The last source is of particular relevance as it is used not only to gain but also to share information, which will be discussed more detailed in section 4.3. Suppliers are consulted with care (as they are likely to be biased), but it seems that especially peers and participants of informal meetings are highly regarded. For online sources it can be said that they are gaining reputation. The way of using these sources is relatively similar amongst the interviewees. This means all CISOs use their sources continuously rather than case base. Nonetheless, time is seen as critical issue. Utilising the internet, a web based Collaborative Information Security Knowledge Ontology Editors (CISKOEs) can provide functionality to share content which so far mainly has been done offline. In this way, a CKSOE can be beneficial as time can be saved.

4.2.2 Credibility of knowledge

One of the basic factors in regard to credibility is personal judgement, which can be achieved by simply looking at the information and seeing how relevant it is. Relevance itself refers to what actually is needed. Moreover, the quality of information is often judged according to personal feeling. Credibility can also be achieved if information can be verified from another independent source. Following this idea, it would be an option for an implementation of a CISKOE to ask users to provide test procedures alongside their knowledge. Similar to studies on Wikipedia, it could be found that open discussions of the content increases its credibility, as the impacts can be judged by the reader in more detail. Additionally, the participating CISOs considered a content quality ranking system, suggested by the interviewer, to be useful. However, one interviewee expressed a wish to see who rated the content. With regards to referencing external sources, this study found that it would add credibility. Nonetheless, doubt was raised concerning the capability of proper referencing within the IT Security community. For a CISKOE this means that common internet practices, such as comments and ratings, as well as

the academic habit of properly citing sources would add credibility. A rather surprising discovery was that too much contribution can limit credibility as one would presume that IT personnel, who have a high level of knowledge, would be highly engaged with their work. It would then be implied that highly engaged IT personnel should not have too much (if any) available time to commit to sharing their knowledge with the rest of the community.

Trust has been identified as the most important characteristic of credibility, as all interviewed CISOs mentioned it in connection to credibility. Trust can be achieved by the following factors:

- knowing a source in person
- knowing the organisation that a source stems from or that a person works for
- contractual assurance
- positive previous experiences

Contractual assurance explicitly includes non-disclosure agreements (NDAs) Knowing the author and their reputation, as well as that of the organisation they work for can build trust. To take advantage of this trust-building factor, a knowledge-sharing system would have to provide functionality, where the author is known to the reader. Potentially, this causes problems as two CISOs said they wanted some content to be anonymised or to only allow chosen users to see what they have contributed.

This would mean that one important way of generating trust is not initially available and also positive experiences cannot be made, since one does not know who the author is.

As authorship has been identified as important to credibility-building, the opposite, i.e. anonymity, causes doubts as it lacks the *"courage of conviction"*. Nonetheless anonymous content could be *"interesting"* but would be treated with scepticism. For a CISKOE this means that even content whose author is not known provides value to some degree.

4.3 Factors for and against knowledge sharing

Whilst anonymity defeats credibility it also is a factor encouraging the sharing of information that would not otherwise be shared. This is evident in Figure 2 (Part 3). Nonetheless, one interviewee said they would not share any sensitive information (e.g. risk registers) that require anonymity, while the other two would be willing to if protective measures were in place. These mechanisms (NDAs for example) are required to protect the reputation of an author or organisation. Furthermore, they are needed as the impact on a business could be severe because the knowledge shared can potentially be used to compromise an organisation's IT security defences.

Apart from anonymity, protection can be gained through the aforementioned NDAs. This was particularly emphasised by two interviewees who were willing to share sensitive information only on condition of provision of an NDA or *"Chatham House Rules"* (an informal agreement that serves a similar purpose) respectively. For the possible realisation of a CISKOE this has a huge impact, as it essentially would require an organisational procedure to register users under an NDA. Were this done, the system could rely on authorship as a mechanism to build credibility and also enable users to share information which is dear to the organisation. Additionally, it is worth mentioning that even the CISO not willing to share sensitive data, did so in the past with a *"trusted*

party". Again this shows that a CISKOE would have to be developed in such a way that CISOs can put trust in it.

As stated above, information that would not be shared freely is data which is crucial for a company's security. This also includes best and in particular worst practices. From this it follows that if protective measures are in place, there generally is a willingness to share knowledge in order to learn from each other and to optimise work in such a way that the same work is not repeated in multiple organisations. Knowledge that would be shared freely, would concern standards and how to enact them. Furthermore, ostensibly CISOs share knowledge for the good of the community of security managers, working collectively *"against the dark forces"*. Seemingly this outweighs the unwillingness to give information to competitors.

Therefore, the sharing of knowledge is in place in various forms. The most important form is knowledge sharing in small groups (10 to 50 members), by gathering for informal meetings. Informality is seen to be the key issue in simplifying the exchange of knowledge and benefiting from the knowledge contributed by others. The development of a group thinking mentality could, however, be identified as a danger of sharing knowledge. Nonetheless, this is an issue that applies to all areas of knowledge-sharing and is not specific to security knowledge.

Similar to consulting sources, conferences and informal meetings are highly time consuming and potentially cost-intensive. Most notably, time was identified as being a key restrictive factor of knowledge sharing because *"more important"* work could be done in the same time. Conversely, the public sector has a duty to share knowledge with other public sector organisations and even with the private sector as long as it is beneficial for the country or region. Regarding budgets, the matter seems to be reversed, i.e. the public sector appears less willing to spend money on knowledge sharing than the private sector. These shortcomings in knowledge-sharing practices could be overcome with relative ease through use of a web-based collaborative tool such as a CISKOE.

Trust, derived from personal relationships, seems to be of similar importance with regards to knowledge sharing as it is for credibility of knowledge. This is evident by one interviewee saying, when speaking of individuals they would be willing to share knowledge with, *"I am not talking about somebody I know. I mean a proper friend."* Similar to what has been found in research on Wikis it could be confirmed that reasons for contribution are mainly non-selfish, i.e. not monetary or reputational. However, there is a reciprocal aspect to it: *"I'd expect to be able to make a phone call to them if I needed help from them as well."* In building a knowledge sharing platform, the case could be made that it would be a good idea to build a system which can be used by these established socio-economic relationships and which encourages contribution as much as possible, as is examined by this work.

All of the participating CISOs generally share their knowledge within limited groups, and expressed a willingness to consider contributing to a knowledge sharing platform. However, some further issues were identified. One CISO stated that the security of any centralised knowledge repository would need to be verified by an independent source before they would consider using it. This is achievable, but the technical implementation of the security on such a system, as well as its costs are yet to be examined. One potential solution to explore is to distribute components of the knowledge repository: one part could be on-site in the organisation to store highly sensitive and restricted data, and another on the Internet to host freely shared data.

4.4 Further issues

This section of findings correlates to Part 4 of Figure 2, and presents the last three identified areas listed in the introduction of Section 4 and their relevance to the development of a CISKOE. Controversial opinions could be identified with regard to senior management. One CISO stated that senior management buy-in is critical as ultimately decisions are approved by senior management, rather than by the CISO. Conversely, another CISO stated that decisions ultimately lie with him. These statements indicate that different levels of autonomy can be seen across the CISO community. This again shows the need for a layered system that can represent data in such a way that it can be understood by non-technical senior management if the need arises.

There was also a further examination of the attributes that would be helpful to CISOs when evaluating the relevance of a knowledge contribution, the focus being on the properties of a contributor's organisation that inform the applicability of their contributions to the working environments of peers. While one interviewee said this information would not be relevant, the other two had clear ideas of what information would be of interest. It could be deduced that generally organisation sizes (share capital and/or employees), the location and organisation sector are potentially relevant. On a technical level the number of (web-)servers and the type of network protection (i.e. firewall) have been identified as relevant. For CISKOEs this means that users should be asked to provide these details as soon as they register.

Uniformly, all CISOs interviewed stated that it is important to manage a collaborative knowledge management system in order to assure the quality of the content and limit redundancies, One CISO actually signalled willingness to pay for such management services. These payments could then be used to cover the costs of maintaining, administrating and securing the system.

5 Conclusion

This Section reviews the research questions outlined in Section 2.3, reproduced here for convenience:

RQ1: What are the incentives for Chief Security Information Officers (CISOs) to share their knowledge and contribute to a CISKOE?

RQ2: What constitutes a credible piece of knowledge for a Chief Security Information Officer (CISO)?

Answering the first research question (RQ1), learning from peers in the community has been recognised as the driving factor for knowledge-sharing. For the second research question (RQ2), trust in the sources of knowledge could be identified as key to building credibility.

The need for a layered system has been pointed out on several occasions in the findings section. This layered structure applies to various areas of a Collaborative Information Security Knowledge Ontology Editor (CISKOE). From what has been discussed with the participants of this study, it is recommended to have three distinct groups of CISKOE users (Unregistered users; Registered users; Registered users with NDA).

Unregistered users would get access to the system, but would only see content that authors share with the general public. To be able to trace the originator of a contribution, it makes sense that only registered users should be able to contribute. Registered users and registered users with NDAs would only be differentiated in that information published under an NDA would not be viewable to those users who have not arranged an NDA. This implies that authors have to state whom they want to share a contribution with. To extend the control over who will be able to read their content it is recommendable to contributors to restrict the access to users or user categories of their choosing. Additionally, contributors should have the option to publish content anonymously if the recipient is unknown, or if the author wants to stay anonymous. Similarly users should be given the option to define what information about themselves (company details, relevant technical information, job title, contact details), can be accessed by other users.

For reasons of simplification, it is worth considering whether to only have registered users with NDA. However, NDAs cause a significant administrative effort, since offline administration would be needed. Additionally, NDAs could potentially discourage some users, who would not want to expend the perceived effort associated with using the tool. Nevertheless, by utilising NDAs the requisite trust could be established. A further way of building credibility or validating content is to use the established methods of content rating, content discussion and referencing. These features should only be useable by registered users to avoid misuse. Again, commenting users should be given the option to choose who will be able to see their comments in case a comment includes sensitive data.

Another matter regarding layers of access is the issue of detail of information. It became clear during the interview process that senior management and technicians need a different view of the same data. Having only interviewed CISOs a first suggestion is to enable authors to flag their content as either technology-orientated or related to administrative processes. However, this would even increase the identified need for content curation. Therefore, a CISKOE should be managed by someone or some organisation to maintain the data integrity.

If a sufficiently large proportion of the CISO community can be convinced of the benefits of the system, it is more likely that users will be willing to contribute and that both the user base and body of knowledge will eventually grow. This then would also increase the value of the system. However, the identified discouraging factor of security on the system has to be addressed by future work.

All ideas mentioned in the previous section are potentially costly. Therefore it is recommended that the costs and benefits of a CISKOE should be examined in greater depth. Furthermore, the acceptability and utilisation of fees for registering with such a system have to be investigated.

The major limitation of this research project was the relatively small sample size which did not allow for generalisation. To further investigate this matter on a larger scale, a next step could be quantitative research utilising questionnaires aimed at a wider audience. Such a questionnaire would be informed by the findings of this work.

Another limitation of this study is that only CISOs have been interviewed. To further explore the applicability of a CISKOE all potential users should be questioned about their opinions. Further groups of interest, as identified within this work, include non-technical senior management, technicians and domain experts such as specialists in human behaviour.

Acknowledgements

The authors would like to thank the interview participants, who wish to remain anonymous. The authors are supported in part by UK Technology Strategy Board (TSB), grant nr. P0007E ("Trust Economics") and HP Labs Innovation Research Program, award ID 2009-1052-1-A ("Prediction and Provenance for Multi-Objective Information Security Management").

References

Blanco, C., et al. (2008), "A systematic review and comparison of security ontologies", **ARES 2008 - 3rd International Conference on Availability, Security, and Reliability, Proceedings,** pp. 813–820.

Fenz, S. & Ekelhart, A. (2009), "Formalizing information security knowledge", **Proceedings of the 4th International Symposium on ACM Symposium on Information, Computer and Communications Security, ASIACCS'09,** pp. 183–194.

Forte, A. & Bruckman, A. (2005), "Why do people write for Wikipedia? Incentives to contribute to open-content publishing", **GROUP 05 workshop: Sustaining community: The role and design of incentive mechanisms in online systems. Sanibel Island, FL.**

Javanmardi S., et al. , "User contribution and trust in Wikipedia", **Collaborative Computing: Networking, Applications and Worksharing, 2009. CollaborateCom 2009. 5th International Conference**, Vol. 2009.

Johnson, B. (2008), "Incentives to contribute in online collaboration: Wikipedia as collective action", **Paper presented at the International Communication Association 58th Annual Conference, Montreal, Quebec, 2008**.

King, N. (1997), "The qualitative research interview" In C. Cassell & G. Symon (Eds.), **Qualitative methods in organizational research. A practical guide,** London: Sage, pp. 14–36.

Kittur, A. & Kraut, R. (2010), "Beyond Wikipedia: Coordination and conflict in online production groups", **Proceedings of the ACM Conference on Computer Supported Cooperative Work, CSCW,** pp. 215–224.

Kittur, A., et al. (2007), "He says, she says: Conflict and coordination in Wikipedia", **Conference on Human Factors in Computing Systems - Proceedings,** pp. 453–462.

Kuznetsov, S. (2006), "Motivations of Contributors to Wikipedia", **ACM SIGCAS Computers and Society**, Vol. 36, (2).

Mace, J., Parkin, S. & van Moorsel, A. (2010), "A Collaborative Ontology Development Tool for Information Security Managers", **Newcastle University Computing Science Technical Report Series**, No. CS-TR-1212.

Majchrzak, A., Wagner, C. & Yates, D. (2006), "Corporate wiki users: Results of a survey", **Proceedings of WikiSym'06 - 2006 International Symposium on Wikis**, Vol. 2006, pp. 99–104.

McCracken, G. (1988), **The long interview,** London: Sage.

Morton-Williams, J. (1993), **Interviewer approaches,** Aldershot: Dartmouth Publishing Company.

Parkin, S., van Moorsel, A. & Coles, R. (2009), "An information security ontology incorporating human-behavioural implications", **SIN'09 - Proceedings of the 2nd International Conference on Security of Information and Networks,** pp. 46–55.

Pfaff, C. & Hasan H. (2006), "Overcoming organisational resistance to using wiki technology for knowledge management", **PACIS 2006 Proceedings**,
paper 110.

Raskin, V., et al. (2001), "Ontology in information security: A useful theoretical foundation and methodological tool", **Proceedings New Security Paradigms Workshop,** pp. 53–59.

Schulzki-Haddouti C. & Lorenz-Meyer L. (2008), **Kooperative Technologien in Arbeit, Ausbildung und Zivilgesellschaft**. Available from: http://www.cio.bund.de/SharedDocs/Publikationen/DE/E-Government/plugin_studie_kooperative_technologien_download.pdf?__blob=publicationFile (accessed 16/02/2011).

Surowiecki, J. (2005), **The wisdom of crowds**. **Why the many are smarter than the few,** London: Abacus.

Swee-Lin Tan, S. & Zou, X. (2009), "An empirical study of the mediating mechanisms of knowledge contribution", **Special Interest Group on Human-Computer Interaction SIGHCI 2009 Proceedings**.

Tsoumas, B. & Gritzalis, D. (2006), "Towards an ontology-based security management", **Proceedings - International Conference on Advanced Information Networking and Applications, AINA**, Vol. 1, pp. 985–990.

Wagner, C. & Prasarnphanich, P. (2007), "Innovating collaborative content creation: The role of altruism and wiki technology", **Proceedings of the Annual Hawaii International Conference on System Sciences**.

White, K. & Lutters, W. (2007), "Midweight collaborative remembering: Wikis in the workplace", **Proceedings of the 2007 Symposium on Computer Human Interaction for the Management of Information Technology, CHIMIT '07**.