An Information Security Ontology Incorporating Human-Behavioural Implications

Simon E. Parkin*
School of Computing Science
Newcastle University
NE1 7RU
Newcastle-upon-Tyne, UK
s.e.parkin@newcastle.ac.uk

Aad van Moorsel
School of Computing Science
Newcastle University
NE1 7RU
Newcastle-upon-Tyne, UK
aad.vanmoorsel@newcastle.ac.uk

Robert Coles
Merrill Lynch Europe Limited
1 Swan Lane
EC4R 3TN
London, UK
robert_coles@ml.com

ABSTRACT

Security managers often regard human behaviour as a security liability, but they should accommodate it within their organisation's information security management procedures. To further the comprehension of human-behavioural factors we develop an information security ontology. This ontology is intended for organisations that aim to maintain compliance with external standards (in this case ISO27002) while considering the security behaviours of individuals within the organisation.

We demonstrate use of our ontology with an applied example concerning management of an organisation's password policy, and how it may be perceived by individuals in the organisation. We formally represent information security controls and findings regarding human behaviour, and relate these to each other and the accomplishment of standards compliance. In doing so we provide a model that information security managers can use to consider the impact of their security management decisions.

Categories and Subject Descriptors

H.1.2 [**Models and Principles**]: User/Machine Systems—*Human Factors*

General Terms

Security, Human Factors

Keywords

information security ontology, human behavioural implications, password policy

*Corresponding Author.

SIN'09, October 6–10, 2009, North Cyprus, Turkey.

1. INTRODUCTION

Increasingly organisations are looking to external, industry-recognised best-practice standards for guidance on how to manage their information security infrastructures (such as the ISO27K series e.g., [1, 2]). By seeking compliance with standards, an organisation can demonstrate that their information is adequately secured, and illustrate to customers and business partners alike that they can be trusted to protect important information.

One shortcoming of applying information security standards in a "one-size-fits-all" manner is that there is no outward consideration of the security priorities and working culture of individual organisations [25]. Organisations may for instance differ in their willingness to exchange data to leverage business opportunities, and the behaviour they wish to encourage within employees towards the data that they have access to [16].

Information security managers need to understand the usability requirements of employees [3], as well as the potential risks posed by employee behaviour. Efforts have been made to categorise how employees perceive information security within the workplace [7], but not necessarily how they behave in reaction to security [9] or how security managers can learn to respond appropriately.

Previous work examining the use of removable USB data-storage devices within organisations [5, 9, 13] has shown a need to consider human-behavioural factors when managing information security policies and security mechanisms. Here we seek to provide a standardised information model for representing behavioural factors and how they relate to an organisation's security needs.

We achieve this goal by augmenting the use of best practice information security standards with a structured definition of the associated *human-behavioural implications*, encapsulated within an ontology (i.e. an information model). This informs the decision-making process, allowing security managers to account for the identifiable effects (be they direct or indirect) that information security mechanisms have upon individuals within an organisation.

As an example, a password authentication policy may mandate that employees use complex passwords to reduce the chance of passwords being guessed or cracked, and thereby provide security. However, the mandated password complexity may push those employees struggling to remember their passwords to write them down in an unsecured location. This could conceivably result in a less secure environment than would have been experienced had simpler (and more

easily memorised) passwords been employed. In this case an information security manager would have benefited from being able to see the consequences of their security decisions. This and other examples are covered in the paper.

We frame our work within the context of information security standards compliance, referencing the ISO27002 standard [2]. Security ontologies based upon the ISO27K standards already exist but do not incorporate human-behavioural factors. We have selected specific ISO27002 guidelines and identified potential policy decisions that may be made during their enactment. We focus on decisions that may impact upon both the security of an organisation's information assets and the behaviour of employees as they use those assets. This information is then structured within the ontology.

Throughout this work we have consulted an industrial partner representative (a senior information security manager within a large, international financial organisation), who is responsible for the computer-system accounts of 50000 staff and 20000 contractors. The organisation has successfully applied ISO27K standards, and so our consultations have provided an insight into how our work might be used in practice.

Section 2 discusses the background to our work, including appreciation of human behaviour in information security, how an ontology can help, and related work. Section 3 introduces our ontology, followed in Section 4 by an example of its use in password policy management. Section 5 provides an evaluation of the suitability and prospective use of the ontology, followed by concluding remarks in Section 6.

2. BACKGROUND

2.1 Information Security & Human Behaviour

Information security managers have traditionally relied on technical controls (e.g. firewalls, e-mail filtering) to secure the information that an organisation values. However they can ill afford to ignore the human element within the organisation [4, 16].

Organisations often aim to employ individuals who have a willingness to take risks and exploit opportunities to benefit the organisation itself. At the same time this behaviour can create *human vulnerabilities*, as security incidents may occur as a result. Ill-informed or inappropriate attitudes to information security can cause a great amount of damage, be it through careless talk, excessive or accidental distribution of documents, or simply failing to adhere to security procedures [20]. Employee behaviour can at times be knowingly insecure, intentionally malicious towards the organisation, or a source of accidental security breaches [15].

When managing the human element in information security it is necessary to consider both the impact that security mechanisms will have upon the workforce and how they will choose to react to those mechanisms. An individual will be less willing to comply with security policies if they perceive them as having a detrimental effect upon their primary work tasks [10].

The burden placed upon employees can be measured in different ways, such as a restriction of work capabilities, delayed tasks, or additional processes and information to remember and recognise. This can in turn have differing effects upon an individual's attitude towards security, for instance instilling a sense of frustration or futility [19]. Inadvertently influencing a person's attitude in such a negative way only fuels the conscious choice to circumvent security controls, or otherwise hold them (or the organisation) in contempt [16].

2.2 Taking Control

Organisations must cultivate an awareness of the human-behavioural implications of their internal information security decisions. The person best positioned to do so would be the internal information security manager (CISO, CIO, etc.) and, if applicable, any members of their team that are normally included in the decision-making process. We refer simply to the 'CISO' as a collective term for such an individual.

Ideally those involved in the decision-making process would have an awareness of the business goals of the organisation (and with this the associated legal and financial concerns [24]). The propensity for risk, as communicated by senior management, must be adequately represented in the information security policies of the organisation. Individual employees can (and arguably should) be capable of behaving in a 'risky' manner to further the goals of the organisation. CISOs are in a position to ensure that the right risks are free to be taken at the right time, and that risks that senior management do not want being taken are prevented all of the time. This provides a challenge to clearly represent the directives of senior management.

Predicting the usability needs of employees and tailoring the information security infrastructure around them should be a priority [25]. Understanding information security from a human-behavioural perspective could also in turn be used to promote desirable security behaviours, and pacify negative perceptions of information security within the organisation [4]. Such an understanding could help in identifying, managing, and potentially stemming the causes of persistent problems (e.g. staff forgetting passwords), as opposed to perpetually reacting to the symptoms [16].

2.3 How an Ontology Can Help

We now discuss a potential solution to the problem identified in the previous sections.

Organisations are driven to follow security standards by a need to structure their security management processes and furthermore demonstrate these processes to auditors and customers. It has been observed that external standards must be tailored to respect the organisational culture and business priorities, and the usability requirements of employees [25]. Doing so serves to make standards compliance a more viable endeavour [30]. We assume here that a CISO is free to balance standards compliance and employee usability requirements.

A fully-informed organisational view of information security must to some degree include internal controls, external standards, and usability concerns. Provision of a structured information model would go some way towards achieving a holistic view of information security management in terms of both its intent and its impact upon members of the organisation. For example if a standard recommends deployment of a password authentication system to protect access to valuable data, a model can be used to relate the CISO's password policy decisions to the projected end-user experience.

Representing such a view in an ontology would be appropriate for a number of reasons:

- By providing a taxonomy of information security terminology security engineers can broaden their knowl-

edge of related concepts [34], in this case the human-behavioural implications of security decisions.

- By encapsulating a standardised taxonomy an ontology can provide a common language [35], which may facilitate improved communication of information security needs and outcomes [34]. This may serve to bridge differing "auras of understanding", as may be seen between senior management and CISOs.

- Ontology content can be reappropriated for other uses, and developed over time [35].

- An ontology provides opportunities for inter-operability e.g., between different assessment methodologies or software tools [36]. This has the potential to generate new knowledge.

- Creating ontology content requires that a diverse array of terms, concepts and relations be reduced into a more refined, structured information model. This serves to organise and make precise any contained knowledge.

An ontology can thus serve as a framework within which to associate human-behavioural research findings with the requirements of a CISO. This would require an appropriate structure to represent the complex relationships and dependencies between concepts within both human behaviour research and information security management.

2.4 Related Work

A number of ontologies, taxonomies and models have been developed for purposes relating to information security. It is useful to review notable works to better understand what is required of our ontology.

Work by Magklaras & Furnell [26] provides a tool for estimating the level of threat originating from an organisation insider. The argument here is that "all actions that constitute IT misuse lead back to human factors", and that individuals within an organisation have greater access capabilities than those outside. A taxonomy was developed to represent properties of users, including behavioural motivations e.g., intentional and unintentional behaviour. There are also basic representations of the system-level consequences of insider threats. It is proposed that the taxonomy be used to profile individual users, and that these profiles be correlated with related system activity to determine the threat posed by each system user. This work supports the point that a particular pattern of security behaviour can have a number of causes and consequences, and that a CISO needs to consider these.

Another investigation into the misuse of information is described in the work of Braz et al [32]. This work discusses using high-level policies to mitigate procedural threats. For instance, "verify source of information" would counter "customer provides false info". The model also describes the combination of policies and their implementation as technical security controls. This work demonstrates the concept of composing controls to influence potential user behaviours. Furthermore it provides examples of security controls that can be used to align user behaviour with organisational policy.

A further study of the threat posed by malicious insiders is found in [41]. The authors provide a model representing the behavioural and technical indicators of malicious insider activity (specifically sabotage and espionage), and the causal relationships between these events (e.g., stressful events as precursors to malicious activity). Experts used the model to encapsulate knowledge of insider threats and convey advice on limiting malicious insider activity. Behavioural and technical indicators of malicious activity are formalised in the model such that organisations can apply it to better manage their own workforce. With our ontology we also seek to encapsulate and communicate expert knowledge of human factors in security, and provide a model that is of potential use to security managers when making policy decisions.

The ROPE methodology [39] and related security ontology [14] provide organisation-wide evaluation of IT security management, with a focus on business processes and risk-management. The ontology encapsulates well-known information security concepts such as assets, vulnerabilities, threats and controls. These inter-related concepts form a framework for arranging organisation-specific knowledge, used for both high-level decision-making and as input to a risk assessment process. The work in [39] and [14] illustrates use of organised infrastructure knowledge as a tool in holistic security management decision-making. Our work aims to provide a similarly 'global' view of organisational priorities. The security ontology in [14] stresses the need of IT managers to represent and communicate qualities of the IT infrastructure to provide justification for security decisions and reduce the reliance on intuition. We aim to utilise an ontology for similar reasons, however we introduce qualities that represent the behaviour of the organisation's employees.

A security ontology incorporating external standards is described in the work of Fenz et al [33]. Here individual guidelines from the ISO27001/2 standards are related to tangible security control implementations within an organisation, providing a means of evaluating internal security policies against the ISO standards. Our work also informs per-organisation security management with a consideration for external standards, but also considers human factors.

Work by Lee et al [38] describes the derivation of security requirements from external standards (e.g. US Department of Defense guidelines). A process is developed for determining interdependencies across content from different standards, and the derivation of criteria for use in adapting standards to internal security configurations. This work demonstrates adaptation of natural-language security standards to internal security infrastructures by way of an information model. This approach is used to help predict and understand how standards will function when they are applied to an organisation's technical infrastructure. In our work we use an ontology to relate a standard to the security compliance of employees, through the effects of security infrastructure upon working practices.

We take the approach of formalising knowledge of human factors in information security, as also demonstrated in [26], [32] and [41]. We represent these factors in an ontology alongside recommendations from security standards and policies (building on evidence from [33] and [38]). Human behaviour is arguably too rich and varied to be reduced to an ontology. However by adapting familiar concepts within information security (e.g., asset, vulnerability, etc., as also seen in [14]) we provide a framework that allows us to relate aspects of both human behaviour and information security.

2.5 Requirements

The previous discussions regarding information security ontologies and human behaviour have highlighted a number of requirements that we must consider when developing an ontology:

- Information security mechanisms are guided by policies. These policies are increasingly informed by external standards. Standards should be recognised and suitably accommodated in the ontology.
- The usability and security behaviours of staff must be considered as part of an organisation's information security policies. This includes identification and representation of:
  - the information assets that IT users have access to;
  - the vulnerabilities that IT users create;
  - the intentional or unintentional threats user actions pose to organisational assets and the information security infrastructure, and;
  - the potential process controls that may be employed to manage user behaviour, and their identifiable effects upon that behaviour.
- CISOs must be able to relate ontology content to the security infrastructure they manage. Ontology elements representing both human factors and external standards should be:
  - presented clearly and unambiguously, and;
  - related logically to one another so as to illustrate interdependencies.

3. DESIGN

Here we introduce an ontology that represents attributes of external information security standards, and aligns these with the potential *human-behavioural implications* of their implementation. This requires separating the content of each standard guideline to identify individual security processes and controls, which can then be associated with the usability concerns that they may raise in practice.

We use the international ISO27002 standard [2] as a framework. ISO27002 establishes guidelines and guiding principles for information security management [2], built upon a prescribed security management process detailed in the ISO27001 standard [1]. The ISO27002 guidelines were determined by a joint technical committee, based on accepted information security management requirements. ISO27002 is intended as a framework for the development of internal security policies and practices.

We assume that CISOs should have the capacity to intentionally or otherwise stray beyond the guidance of the standard if they believe it to be beneficial to the organisation. The goal is then a balance of security, user policy compliance and business process constraints. The ontology in its present form does not model the business processes within which a particular user behaviour may manifest, nor the probability of a behaviour manifesting. It can however record the owner of an information asset and the perceived cost to the organisation of a potential threat or control, and the causal relationships between different behaviours.

We do not consider technical- or configuration-level concerns in our ontology (e.g. the mandated length of passwords or the required character sets). New or alternative technologies emerge with greater regularity than security standards, and so we focus on high-level, procedural controls. Other works have made a similar distinction (e.g., [22, 32, 35]).

During development of our ontology we followed the advice for creation of ontologies as outlined in [40]. By adhering to these recommendations we ensured that the structure of our ontology would be consistent, robust and ultimately usable. We also built upon ontology concepts previously developed by Fenz et al. [33] to align ISO27K standards with components of an organisation's security infrastructure.

Our ontology has been implemented in the Ontology Web Language (OWL) [17]. We use OWL as it is a well-supported ontology language, and there are code libraries available to facilitate building software applications on top of the ontology in future. We used the Protégé Ontology Editor application [18] to construct the ontology and enter data for our applied example (described in Section 4).

3.1 Overview

The concepts represented in our ontology are shown in Figure 1. Each individual concept has a relationship with one or more other concepts. The objects `Chapter`, `Section`, `Guideline` and `Guideline Step` represent content from the ISO27002 standard. An individual `Guideline` can be associated with a particular information `Asset` by way of the '`hasSubject`' relationship. Otherwise if a `Guideline` has been broken down into more refined `Guideline Step`s it will be these that are linked to an `Asset`. We represent those information `Asset`s identified in a `Guideline` or `Guideline Step` that either must be secured or which are crucial to an information security management process. In our ontology an `Asset` can be '`ownedBy`' someone that has an identified `Role`, who is then responsible for its maintenance.

The ontology also represents the security and usability weaknesses of an `Asset` that may promote or inhibit certain employee behaviours. It is with the `Vulnerability` concept that we introduce these human-behavioural factors into the ontology.

A `Vulnerability` may be '`exploitedBy`' a `Threat` (e.g., if "memorisation of password is difficult" it may follow that "password is forgotten"), which renders the `Asset` unusable or insecure. Note that when a `Vulnerability` is 'exploited', this may be intentional or accidental.

A `Threat` may be either an `Infrastructure Threat` or a `Procedural Threat`. The former represent activities that directly affect security mechanisms, whereas the latter represent security events that impact upon an individual and their behaviour. A `Threat` may also affect productivity, for instance if an employee forgets a password and is unable to access a system until it is changed. For each `Procedural Threat` we record the `Behavioural Foundation`, as a means to classify behaviours and indicate the concerns that they raise within an organisation (e.g., a person's memory capabilities or attitude towards security).

A `Vulnerability` may be '`mitigatedBy`' a `Behaviour Control`. A `Behaviour Control` represents a procedural activity that can be enacted by a CISO to manage the interactions between humans and organisational security controls.

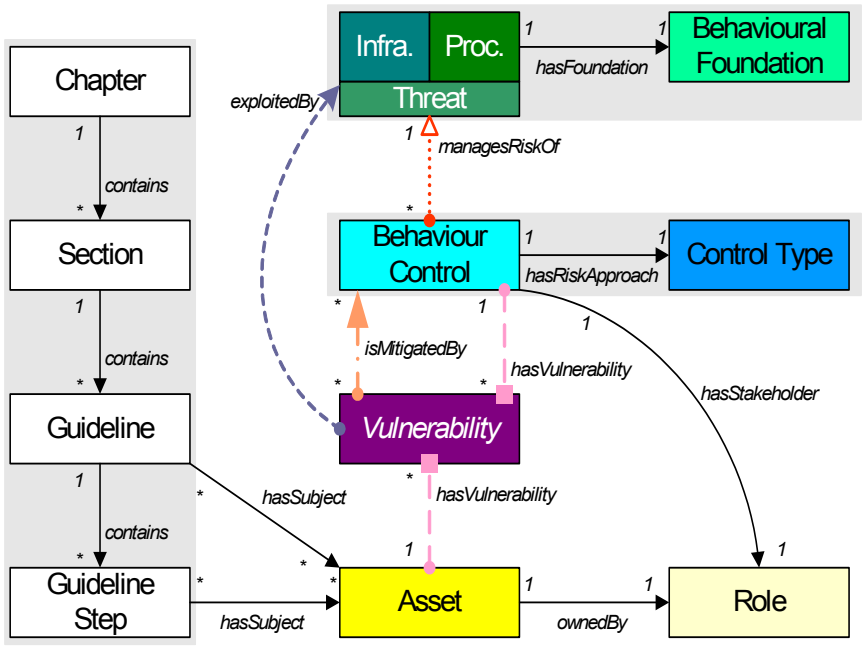

Figure 1: Overview of the concepts and relationships in the information security ontology.

Each `Behaviour Control` has a `Control Type` which indicates the associated risk management approach, such that a `Behaviour Control` '`managesRiskOf`' a specific `Threat`.

Further details regarding each of these concepts and their relationships are described in the following sections.

3.2 Chapter, Section, Guideline & Guideline Step

In the ISO27002 standard content is arranged into `Chapters`. Each `Chapter` refers to a general area of information security management (e.g. "Access Control", "User Training & Education"), and has a number of `Sections` which each address a specific area of the `Chapter`'s subject matter. Each `Section` contains a number of `Guidelines` detailing specific procedural concerns (e.g. "User Password Configuration", "Notifying New Employees of Terms of Use"). Each `Guidelines` contains 'Implementation Guidance', which in some cases is broken down into discrete steps (referred to here as `Guideline Steps`).

We resolve the discourse between external standards and infrastructure concepts (such as information assets) by associating infrastructure components with specific `Guidelines` or `Guideline Steps`. We accept that the standards structure we have chosen is tightly-bound to the structure of the ISO27002 document. However, it is not inconceivable for other external standards to be modelled also (owing to the natural extensibility provided with use of an ontology).

3.3 Asset

Within the ontology an `Asset` is an identifiable information artifact which is of value to the organisation (be it monetary or as a means to further the organisation's business goals). By identifying `Assets` and the controls that can be employed to secure them (as described in the ISO27002 standard), it is possible to begin informing the development of internal policies for their protection.

3.4 Vulnerability

An `Asset` may exhibit some weakness that makes it susceptible to exploitation. Such a weakness is referred to as a `Vulnerability`. We choose to concentrate our efforts on identifying those vulnerabilities exposed in security processes that may be influenced directly or indirectly by human behaviour. That behaviour may be intentionally malicious or an unintentional mistake or oversight. An example would be asking an employee to remember a password when they already have any number of work-related passwords to remember, which may cause that individual to forget or confuse one or more old and new passwords.

At this stage we rely on the judgement of the CISO to determine whether a given `Vulnerability` is exhibited by their information security infrastructure. Since each `Vulnerability` we identify is based in a potential pattern of behaviour, and less in a technical configuration, this essentially requires a judgement on the capabilities and working culture of the organisation's employees.

We consider both *behaviour that affects security* and *security that affects behaviour*, as one can influence the other. A `Threat` (Section 3.5) to a `Vulnerability` constitutes the former, and a `Behaviour Control` (Section 3.6) the latter. For example if an organisation has encouraged a culture of trust among its employees, deploying a password-authentication system might make individuals question who and what they trust.

3.5 Threat

An `Asset` may be perceived as vulnerable to some form of exploitation, but that is not to say that it will necessarily be exploited. A `Vulnerability` becomes a problem when there is some means of exploiting it (a `Threat`) and a perceived probability of that `Threat` manifesting.

In our ontology we make a distinction between those threats that affect the infrastructure of an organisation (e.g. "IT

Table 1: Types of Behavioural Foundation

Cultural	Different cultural practices may exist across geographic (or perhaps even social) boundaries
Ethical	Basic ethical considerations should be noted, e.g. personal privacy
Temporal	Conditional changes may exist based upon the time of day or the duration of an event (e.g. employees may lose focus on their work or have diminished patience at the end of the working day)
Mindset	Someone's disposition could indicate that they may behave maliciously or opportunistically with respect to the organisation's assets
Capability	There may be individuals within the workforce who have some form of physical impairment. This may affect their ability to interact with security mechanisms

help desk too busy to answer password-reset requests"), and threats that affect the human-oriented procedures and usability requirements inherent in using a particular `Asset` (e.g. "user has forgotten system log-on password"). We refer to the former as `Infrastructure Threat`s, and the latter as `Procedural Threat`s.

A `Threat` may hinder user productivity and adversely affect an individual's attitude to security. If security measures are not attuned to the usability requirements of the organisation's employees, they may feel inclined to sidestep security measures that they regard as cumbersome in order to "get the job done". It may then be said that the ontology identifies potential `Threat`s to the *desirable security behaviour* of an organisation.

We record the potential consequences of each `Threat` should it occur (an approach also seen in e.g., [36, 37]). The ontology permits any text-based record of the potential impact of a `Threat` upon individuals and infrastructure within the organisation.

Regarding the probability of a Threat manifesting, we assume that a CISO would base this on either their own experience or on the output of some methodologies external to the ontology. The ontology also as yet has no representation of the passage of time, and as such no measure of how compounded security behaviours may affect an individual (i.e. how the human-behavioural implications of co-dependent `Threat`s may influence each other as they are aggregated over time).

3.5.1 Behavioural Foundation

`Procedural Threat`s are essentially events that constitute a conflict of usability and security instigated by an individual within an organisation. It is beneficial to have a basic comprehension of the individual or societal behaviour that drives this conflict. We refer to this as the `Behavioural Foundation`. The `Behavioural Foundation` allows us to classify a `Procedural Threat`, and assess it in terms of the organisation's desired security behaviour. Table 1 shows the basic types of `Behavioural Foundation` that we have considered in our work so far.

Knowing the `Behavioural Foundation` also informs the level of sensitivity required when managing a `Procedural Threat`. An individual who tells a 'trusted' colleague their password in confidence fearing they might forget it themselves should be approached differently to an individual who, without prompting, tells their password to a colleague simply because they do not acknowledge the security value of that password. The solution changes based upon the identified behaviour. Forgetfulness may (or may not) be accommodated by security training workshops. Ignorance might be stemmed through careful use of sanctions (thereby forcing an individual to address the source of their own ignorance). The focus is not to address the root cause of behaviours that cause incidents (e.g. why some people forget passwords, or why they act maliciously), but to provide guidance to CISOs so they may reason about security decisions to manage complex or conflicting behaviours.

The scope of a `Behavioural Foundation` can be refined. One potential application of this would be in establishing which aspects of the information security policy to concentrate on during staff education programmes. In the case of a person's security 'Mindset' (see Table 1) we have defined different actor behaviours (e.g. "Traitor", "Foe", "Opportunist" etc.), based on those in [9]. These relate motives to behaviour, allowing CISOs to approach altering or countering a person's perceived intentions. Refined scopes for other kinds of `Behavioural Foundation` can also be defined.

3.6 Behaviour Control

A CISO may have methods available to manage a human-behavioural `Vulnerability`. We therefore approach the mitigation of each `Vulnerability` in terms of the potential system-wide policy or high-level infrastructure changes that can be applied. We refer to these changes as `Behaviour Control`s.

Each `Behaviour Control` describes a procedural solution, and leave it to the CISO to decide how to enact the control in practice. We would for example consider a mechanism of policy change at a level of "make passwords more complex" over "include one or more punctuation characters not in succession within the mandated 8-12 character password".

Work by Neubauer et al [6] discusses representing "countermeasure side-effects". Our ontology represents the usability-oriented side-effects of deploying a `Behaviour Control` by way of its human-behavioural implications. In our ontology it is possible for a `Behaviour Control` to introduce further `Vulnerability` concerns. For instance, an unwieldy password authentication system could be mitigated by relying on token-based security - it is then conceivable that employees could lose their tokens or pass them to other people. In this sense our ontology can inform CISOs of the human-behavioural implications of the security mechanisms or policies they intend to implement.

A stakeholder role can also be associated with each `Behaviour Control`, to clarify who should be consulted should it be enacted.

3.6.1 Control Types

Decisions relating to information security management must be explicitly agreed by senior management if they are to have the resources they need. This is not simply a case of mitigating every `Vulnerability` or neutralising every `Threat`, as senior management may regard the associated costs or resultant restrictions on user behaviour as too great to be justified. The security management choices that

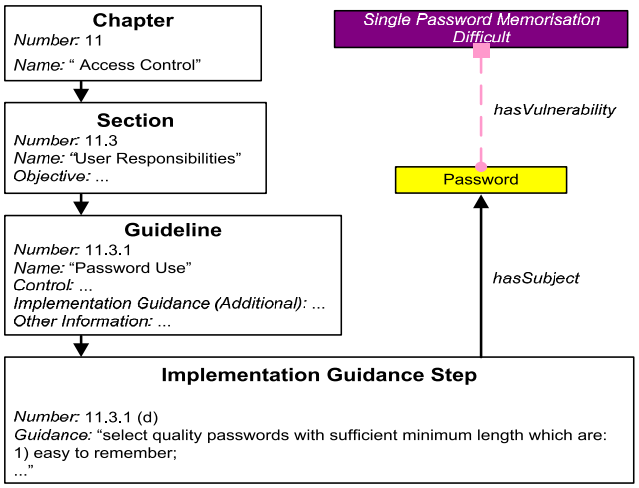

Figure 2: Sample of ontology content relating to password policy.

are made must reflect both the organisation's risk propensity and its working culture.

To this end we associate a `Control Type` with each `Behaviour Control`, as a record of the risk management approach that the control offers toward a particular `Threat`. We use the category of risk approaches described in Section 4.2 of [2] i.e., *risk retention*, *risk reduction*, *risk transfer* and *risk avoidance*.

An example of the use of `Control Type`s would be of a password policy that mandates complex passwords, which individuals may find difficult to remember. There is a `Threat` that staff may forget their passwords, which could be targeted by making passwords simpler (but less secure as a result). Another `Behaviour Control` might be to maintain the same policy but employ IT help-desk staff to manage password resets. This does not make passwords any easier to remember (thereby side-stepping the `Threat`), but manages the implications of the `Vulnerability`. The choice of control should be aligned with the risk approach of the organisation. In this sense a `Behaviour Control` can be used to promote those working practices that fit the organisation's idealised security behaviour.

4. EXAMPLE - PASSWORD POLICY

To assess the efficacy of our ontology we created content based on specific ISO27002 guidelines. Here we examine an example based on one of those guidelines.

Researchers have already investigated the usability issues associated with passwords (e.g., [11]). Here we provide a means for a CISO to consider some of the concepts of this research alongside compliance with the ISO27002 standard.

We consulted experts at the Human Centred Systems Group at University College London (UCL) [42] and examined existing research to identify issues that users typically raise when using passwords (e.g., "forgetting password").

We selected an ISO27002 guideline relating to password policy, specifically `Guideline Step` (d) of `Guideline` 11.3.1, "Password Use". This is shown, in abbreviated form, in Figure 2 (note that Figures 2-5 follow the style conventions as in Figure 1). We identified a user's Password as the `Asset` requiring protection (since a Password provides access to business data). The example ontology content then represents possible behaviours resulting from use of passwords, from the perspectives of both an employee and a CISO.

`Guidance Step` 11.3.1(d) stipulates a number of qualities that a password should have, but here we focus on one quality (that passwords should be "easy to remember") and one `Vulnerability` that directly relates to it. The other qualities that passwords should exhibit and further `Vulnerability` types are found in the extended `Guideline Step` example.

4.1 Applying the Ontology

To derive ontology content, we chose an individual ISO27002 guideline that related to information security (as opposed to physical security). We then identified the information `Asset`s to which the guideline refers. Through consultation of 'human factors' experts and relevant research material we derived the `Vulnerability` types that could potentially manifest when the `Guideline` was enacted. This involved the identification of those research findings stating that a particular security quality of the `Asset` would be undermined by a human behaviour that may be instigated by the security process. Each event where a security quality was placed at risk would form the basis of a `Vulnerability`.

With a set of distinct `Vulnerability` types defined, the `Threat`s that might apply to each `Vulnerability` (and the `Behaviour Control` instances that then applied to each `Threat`) were identified, through a similar process of consultation.

It is worth noting that the consultation process also served to define knowledge precisely. As an example, the `Threat` "Password Stored Externally to Avoid Recall" was initially referred to as "Recall Aids Used". This was perceived by contributors not only as a reference to external storage of a password, but also to the use of memory training techniques to aid recall of a password, and so the definition was refined.

The process of creating ontology content for a particular guideline may be regarded as complete once `Threat` and `Control` instances have been derived upon which the CISO cannot exert influence. For instance, the CISO cannot affect the likelihood of an individual successfully completing a user-training programme.

Within the password policy example we identify a number of concerns linked to password policy, as described in the following discussions. These examples illustrate that a CISO must understand the impact that their information security decisions have upon the behaviour of employees.

4.2 Example Password Policy Decisions

Password Memorisation: For the `Asset` Password identified in Figure 2 there is a `Vulnerability` "Single Password Memorisation Difficult". This represents the possibility that an individual could have difficulty remembering a password with the complexity required by the organisation's password policy.

There are a number of behavioural `Threat`s which may exploit this weakness. Figure 3 illustrates that a Password may be forgotten, perhaps due to an individual's Capability (or more precisely their memory). The consequence of this would be that the individual would be without system access until they could remember their password or have it changed.

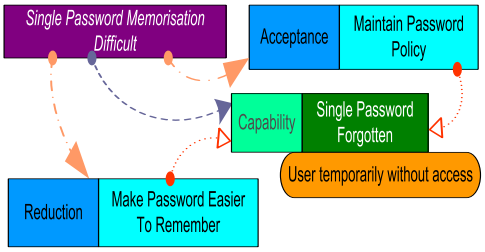

Figure 3: Ontology content for password memorisation.

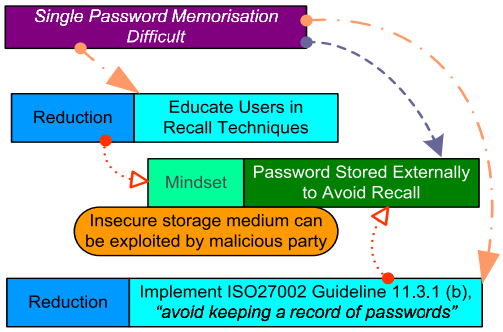

Figure 4: Ontology content for management of password recall methods.

The organisation may accept that this is a possibility but choose to "Maintain Password Policy" as a `Behaviour Control` (valuing perceived security over usability). Otherwise a choice may be made to "Make Password Easier To Remember", thereby reducing the chance of someone forgetting their Password (though potentially at the cost of perceived security).

Managing Password Recall Methods: If an individual is having difficulty remembering a Password it is possible that they will for instance write it down, or use some other record of the Password to avoid the need to recall it themselves. An insecure record of the Password (e.g., on a piece of note paper next to the user's workstation) could be exploited by a malicious party.

One method of reducing the need to record a Password would be to "Educate Users in Recall Techniques" (as in Figure 4), thereby teaching users how to remember complex Passwords. A different approach to the problem of insecure Password records would be to enforce a 'clear desk' policy, whereby users are not permitted to leave any artifacts at their desks which may contain or allude to secured information. This `Behaviour Control` does not solve the problem of users forgetting passwords, but would strictly limit the range of insecure behaviours that may impact upon protection of passwords.

Password Reset Function: To balance the security and usability of passwords organisations can utilise an internal "Helpdesk Password Reset Management" `Control` function [8], as shown in Figure 5. If a member of staff forgets a Password, they can then make a request to the help desk to have their Password changed.

If an IT helpdesk team must carry the cost of users forgetting their `Password`s, this cost may be measured by the number of reset requests that have to be resolved. If many system users are having difficulty with their passwords the helpdesk may be too busy to answer requests promptly. This would prolong the time for which callers are unable to access the system. One solution is then to employ "Additional Helpdesk Staff" in an effort to expedite call resolution.

Helpdesk staff must typically be provided with details verifying the identity of a caller, to ensure the communication of passwords to the correct recipients. A malicious party within the helpdesk team may exploit these details to gain access to an individual's system account. A CISO must then consider that employing a helpdesk function to increase usability introduces security concerns of its own.

An individual may find the "Password Reset Process Laborious", growing impatient with the time it requires. A consequence of this may be a reduced willingness to comply with the process, especially if it detracts from the caller's primary work tasks. One answer to this would be to introduce an "Automated Password Reset System" (assuming here that such a system would answer requests quicker than a helpdesk operative).

5. EVALUATION

5.1 Applicability

During development of our ontology and example content we consulted a large IT consultancy and human factors researchers. This provided insight into the requirements and usability needs of prospective users of the ontology.

With reference to the ontology evaluation criteria outlined in [27], consultation with potential users of the ontology has provided us with assessment of its syntax, usability and content requirements. In consulting human factors researchers regarding our password policy example, we made effective use of expert knowledge to guarantee the completeness, correctness, and consistency of the ontology content that we have developed.

By following guidance on ontology development [40] and encoding our ontology in OWL, we have served to demonstrate that the concepts in the ontology are well-defined and that the relationships between those concepts are precise. The use of OWL also provides us with application-based evaluation of the ontology content [29].

5.2 Envisaged Use of the Ontology

Ideally organisations would have both an IT security analyst and a human-behavioural expert in their employ. With this we envisage that case studies would be formulated to concentrate on specific information security mechanisms and associated human-behavioural concerns. We have demonstrated in our example how research relating to human factors in password authentication procedures can be appropriately incorporated into our ontology, by associating expert knowledge with relevant ISO27002 guidelines.

External standards can act as a foundation for an organisation's internal security policies [8], to be translated by the security team into workable security mechanisms [24]. Here we envisage that our ontology would provide a perspective on how policy directives can affect employee behaviour before they are realised, and allow the CISO to consider the behaviours that they need to promote or prohibit.

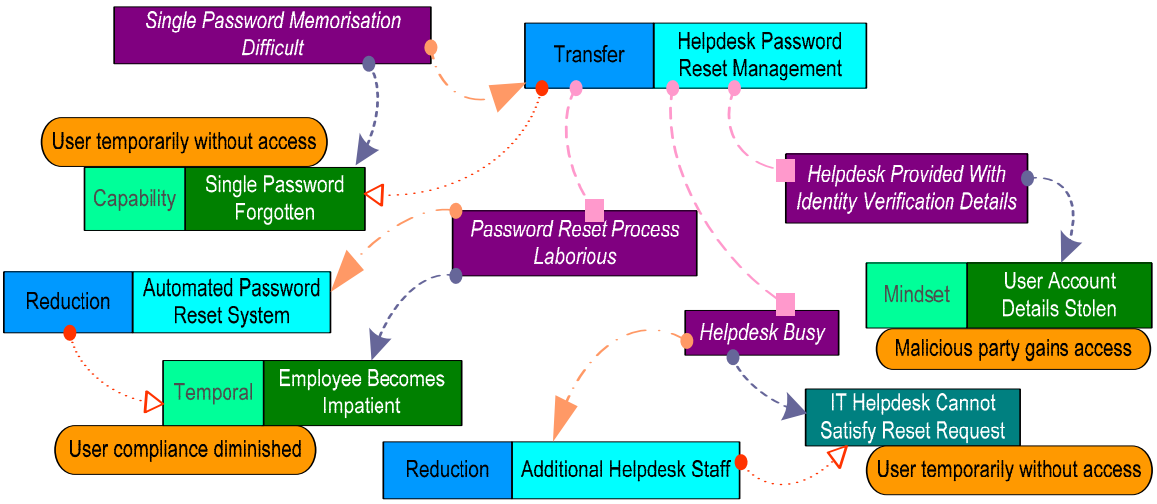

Figure 5: Ontology content for a password reset function.

It is logical to include representatives from all areas of the organisation affected by information security in the compliance process [23]. Our ontology is well-placed to provide natural language recommendations in a structured manner, promoting an inclusive approach to information security management.

An approach of consultation similar to that described in [31] would facilitate effective use of our ontology, wherein the business processes and security properties of an organisation are identified (as found in e.g., resource manifests). Use of an ontology is recommended during such a process.

A CISO may augment use of the ontology with other forms of risk/benefit calculations. We assume that the ontology would be integrated into a broader framework of standards and assessment tools.

6. CONCLUSION

We have investigated the need to understand the usability requirements and security behaviours of employees within an organisation to ensure the effectiveness of information security measures. We have provided an ontology that combines content from information security standards with representations of the potential human-oriented security concerns that their implementation may create. The ontology provides a framework for investigating the causal relationships of human-behavioural implications resulting from information security management decisions, before security controls are deployed.

We conclude that it is possible for organisations to consolidate standards compliance with a consideration for human factors, and have shown that our method of associating security infrastructure properties with their human-behavioural implications can identify potential user behaviours or effects upon behaviour for the benefit of information security managers. We demonstrated use of the ontology to encode expert knowledge of human factors in password policy management, and the ontology can be applied to other aspects of information security management.

Alongside this work we have been actively investigating the usability issues of the ontology from the perspective of an information security manager. For this purpose we have been progressively integrating the ontology into a prototype 'Knowledge Base' application [28].

7. ACKNOWLEDGEMENTS

The authors are supported in part by EPSRC grant EP/F066937/1 ("Economics-inspired Instant Trust Mechanisms for the Service Industry") and UK Technology Strategy Board (TSB), grant nr. P0007E ("Trust Economics").

We are grateful to Adam Beautement and M. Angela Sasse (University College London) for their contributions.

We would like also to thank the industrial and academic partners on the Trust Economics project [12] for their continued feedback and support.

We are also grateful for the insights offered by Howard Smith (Sunderland City Council, UK), Maciej Machulak and Daria Stepanova.

8. REFERENCES

[1] British Standards Institution, "BS ISO/IEC 27001:2005 - Information Technology – Security Techniques – Information Security Management Systems – Requirements", 2005

[2] British Standards Institution, "BS ISO/IEC 27002:2005 - Information Technology - Security Techniques - Code of Practice for Information Security Management", 2005

[3] R. Briggs & C. Edwards, "Skills for Corporate Security", The Business of Resilience, Demos, 2006

[4] KTN Human Factors Working Group, "Human Vulnerabilities in Security Systems: White Paper", Cyber Security Knowledge Transfer Network (KTN), 2007

[5] S.E. Parkin, R. Yassin Kassab, A. van Moorsel, "The Impact of Unavailability on the Effectiveness of Enterprise Information Security Technologies", In Service Availability. 5th International Service Availability Symposium (ISAS 2008), Springer, pp 43-58, 2008

[6] T. Neubauer, A. Ekelhart, S. Fenz, "Interactive Selection of ISO 27001 Controls under Multiple Objectives", Proceedings of the 23rd International Security Conference (SEC 2008), Springer-Verlag GmbH, p. 477-492, 2008

[7] R. Coles, G. P. Hodgkinson, "A Psychometric Study of Information Technology Risks in the Workplace", Risk Analysis, 28(1), pp 81-93, Society for Risk Analysis, 2008

[8] F. N. do Amaral, C. Bazilio, G. M. Hamazaki da Silva, A. Rademaker, E. H. Haeusler, "An Ontology-Based Approach

to the Formalization of Information Security Policies", Proceedings of the 10th IEEE on International Enterprise Distributed Object Computing Conference Workshops (EDOCW), pp 1, 2006

[9] A. Beautement, R. Coles, J. Griffin, B. Monahan, D. Pym, M.A. Sasse, M. Wonham, "Modelling the Human and Technological Costs and Benefits of USB Memory Stick Security", Workshop on Economics in Information Security (WEIS), 2008

[10] A. Beautement, M. A. Sasse, and M. Wonham. "The Compliance Budget: Managing Security Behaviour in Organisations", In Proc. 2008 Workshop on New Security Paradigms, 2008

[11] A. Adams, M. A. Sasse, P. Lunt, "Making Passwords Secure and Usable", Proceedings of HCI on People and Computers XII, pp 1-19, 1997

[12] Newcastle University UK, "Trust Economics Website", `http://www.trust-economics.org/`, last viewed 24/02/09

[13] R. Coles, J. Griffin, H. Johnson, B. Monahan, S.E. Parkin, D. Pym, M.A. Sasse, A. van Moorsel, "Trust Economics Feasibility Study", In 38th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2008), IEEE Computer Society, pp A45-A50, 2008

[14] A. Ekelhart, S. Fenz, M. Klemen, E. Weippl, "Security Ontologies: Improving Quantitative Risk Analysis", pp.156a, 40th Annual Hawaii International Conference on System Sciences (HICSS'07), 2007

[15] Cyber Security Knowledge Transfer Network (KTN) & Economic & Social Research Council (ESRC), "The Economics of Information Security", ESRC Seminar Series, 2008

[16] ISACA, "An Introduction to the Business Model for Information Security", ISACA, 2009

[17] W3C, "OWL Web Ontology Language Overview", `http://www.w3.org/TR/owl-features/`, 2004, last viewed 24/02/09

[18] Stanford Center for Biomedical Informatics Research, "The Protégé Ontology Editor and Knowledge Acquisition System", `http://protege.stanford.edu/`, last viewed 24/02/09

[19] P. Dourish, R. Grinter, J. Delgado de la Flor, and M. Joseph, "Security in the Wild: User Strategies for Managing Security as an Everyday, Practical Problem". Personal and Ubiquitous Computing, 8(6), pp 391-401, 2004

[20] Information Security Awareness Forum (ISAF) & Information Assurance Advisory Council (IAAC), "Creating a Strong Information Handling Culture", `http://www.iaac.org.uk/Portals/0/23176\_DIAN\_A5\_PEOPLE\_15\_4.pdf`, last viewed 24/02/09

[21] A.S. Patrick, "Human Factors of Security Systems". Invited presentation to the HTCIA Atlantic IT Security Professional Development Day, Sept. 30, 2008

[22] H. Mouratidis, P. Giorgini, G. A. Manson, "An Ontology for Modelling Security: The Tropos Approach", 7th International Conference on Knowledge-Based Intelligent Information and Engineering Systems (KES), pp 1387-1394, 2003

[23] B. Karabacak & I. Sogukpinar, "A Quantitative Method for ISO 17799 Gap Analysis", Computers & Security, 25(6), pp 413-419, 2006

[24] N. Nagaratnam, A. J. Nadalin, M. Hondo, M. McIntosh, P. Austel, "Business-Driven Application Security: From Modeling to Managing Secure Applications", IBM Systems Journal, 44(4), pp 847-868, 2005

[25] P. Skidmore, "Beyond Measure", Demos, 2003

[26] G. Magklaras & S. Furnell, "Insider Threat Prediction Tool: Evaluating the probability of IT misuse", Computers & Security, vol. 21, no. 1, pp 62-73, 2002

[27] A. Gómez-Pérez, "Towards a Framework to Verify Knowledge Sharing Technology", Expert Systems with Applications, Vol. 11, No. 4, pp 519-529, 1996

[28] D. Stepanova, S. Parkin, A. van Moorsel, "A Knowledge Base for Justified Information Security Decision-Making", In Proceedings of the 4th International Conference on Software and Data Technologies (ICSOFT), 2009

[29] J. Brank, M. Grobelnik, D. Mladenic, "A Survey of Ontology Evaluation Techniques", In Proceedings of the Conference on Data Mining and Data Warehouses (SiKDD 2005), 2005

[30] R. Gururajan & V. Gururajan, "An Examination into the Role of Knowledge Management and Computer Security in Organizations", The 7th International Research Conference on Quality, Innovation & Knowledge Management, 2005

[31] T. Neubauer & J. Heurix, "Objective Types for the Valuation of Secure Business Processes", Proceedings of the Seventh IEEE/ACIS International Conference on Computer and Information Science (ICIS 2008), pp 231-236, 2008

[32] F. A. Braz, E. B. Fernandez, M. VanHilst, "Eliciting Security Requirements through Misuse Activities", Proceedings of the 2008 19th International Conference on Database and Expert Systems Application (DEXA), Pages 328-333, 2008

[33] S. Fenz, G. Goluch, A. Ekelhart, B. Riedl, and E. Weippl, "Information Security Fortification by Ontological Mapping of the ISO/IEC 27001 Standard", Proceedings of the 13th Pacific Rim International Symposium on Dependable Computing (PRDC2007), IEEE Computer Society, pp 381-388, 2007

[34] D. G. Firesmith, "A Taxonomy of Security-Related Requirements", International Workshop on High Assurance Systems (RHAS'05), 2005

[35] A. Vorobiev & N. Bekmamedova, "An Ontological Approach Applied to Information Security and Trust", ACIS 2007 Proceedings, 2007

[36] L. A. F. Martimiano & E. S. Moreira, "The Evaluation Process of a Computer Security Incident Ontology", 2nd Workshop on Ontologies and their Applications (WONTO'2006), 2006

[37] J. Undercoffer, J. Pinkston, A. Joshi, T. Finin, "A Target-Centric Ontology for Intrusion Detection", In Proceedings of the IJCAI-03 Workshop on Ontologies and Distributed Systems, 2004

[38] S. Lee, R. Gandhi, D. Muthurajan, D. Yavagal, G. Ahn, "Building Problem Domain Ontology from Security Requirements in Regulatory Documents", Proceedings of the 2006 international workshop on Software engineering for secure systems, pp 43-50, 2006

[39] G. Goluch, A. Ekelhart, S. Fenz, S. Jakoubi, S. Tjoa and T. Mueck, "Integration of an Ontological Information Security Concept in Risk Aware Business Process Management", Proceedings of the 41st Hawaii International Conference on System Sciences (HICSS 2008), IEEE Computer Society, pp 377-385, 2008

[40] N. F. Noy & D. L. McGuinness, "Ontology Development 101: A Guide to Creating Your First Ontology", Stanford Knowledge Systems Laboratory Technical Report KSL-01-05 and Stanford Medical Informatics Technical Report SMI-2001-0880, 2001

[41] S. R. Band, D. M. Cappelli, L. F. Fischer, A. P. Moore, E. D. Shaw, R. F. Trzeciak, "Comparing Insider IT Sabotage and Espionage: A Model-Based Analysis", CMU/SEI-2006-TR-026, Carnegie Mellon University, 2006

[42] University College London, "Human Centred Systems Group", `http://hornbeam.cs.ucl.ac.uk/hcs/index.html`, last viewed 19/05/09