A utilitarian re-examination of enterprise-scale information security management

Andrew Stewart

Technical Report

RHUL–MA–2014– 8
01 September 2014

Information Security Group
Royal Holloway, University of London
Egham, Surrey TW20 0EX,
United Kingdom

www.ma.rhul.ac.uk/tech

Andrew Stewart; student number 090390922

A utilitarian re-examination of enterprise-scale information security management

Andrew Stewart

March, 2013

Supervisor: John Austen

Submitted as partial requirement for the degree of M.Sc in Information Security from Royal Holloway, University of London

I declare that this assignment is entirely my own work and that I have acknowledged all quotations from published or unpublished work created by other parties. I also declare that I have read the statements regarding plagiarism in section one of the Regulations Governing Examination and Assessment O↵enses. In accordance with these regulations I submit this project report as my own work.

Signature:

Date:

Table of contents

A utilitarian re-examination of enterprise-scale information security management[⇤]

Andrew Stewart[†]

March, 2013

Abstract. An action is utilitarian when it is both useful and practical. In this paper we examine a number of traditional information security management practices in order to ascertain their utility. That analysis is performed according to the particular set of challenges and requirements experienced by very large organizations. Examples of such organizations include multinational corporations, the governments of large nations, and global investment banks. We identify a number of information security management practices that are considered to be "best practice" in the general case but that are suboptimal at the margin represented by very large organizations. A number of alternative management practices are proposed that compensate for the identified weaknesses.

Keywords and phrases: Information security; management; best practices; budgeting; spending; security management frameworks; BS 7799, ISO/IEC 27001, NIST special publication 800-100, end-user security awareness training; security policies; organizational structures; configuration management; change management; entitlement management.

1 Introduction

There is a danger when activities that are considered to be "best practice" become tradition – when they are applied in all cases regardless of the specific circumstances. Human beings are prone to taking shortcuts and it can be easier to say "we do this because it's a best practice" than to assess the pros and cons. A "best practice" that is applied in an unthinking way could conceivably do more harm than good.

It is therefore a useful exercise to periodically assess where "best practices" are best and where they are not. Our goal in this paper is not to critique current information security management practices necessarily but rather to examine them through the eyes of a particular set of constituents, namely very large organizations. A best

[⇤] Submitted as partial requirement for the degree of M.Sc in Information Security from Royal Holloway, University of London.

[†]

practice might be useful in the general case but perhaps not useful at the margin represented by such organizations.

This is not a survey paper nor one with a mathematical proof. This is a discussion paper based on observation, review of existing literature, and analysis. Our goal is to identify weaknesses in current approaches and thereby identify new, potentially more beneficial ways of working.

The analysis we perform in this paper is *utilitarian* in the sense that our focus is on assessing management practices to determine their utility, i.e. the degree to which they are they are both useful and practical rather than simply being attractive. The paper is a *re*-examination because there is an existing doctrine of thought regarding the management of information security.

We believe that it is always a worthwhile exercise to examine first principles and revisit conventional wisdom with a skeptical eye. From that analysis, new insights and approaches can flow.

The paper is organized as follows. We begin in §2 by defining terms. §3 contains a description of some of the key challenges faced by very large organizations along with some common approaches to those challenges. A review of the existing literature is performed in §4. §5 provides an analysis of a select number of topic areas within the field of information security management. §6 proposes a number of new approaches to information security management within very large organizations. Opportunities for future work are identified in §7. Acknowledgments are listed in §8. The conclusion is presented in §9 and references are provided in §10.

2 Definition of terms

In this paper we are primarily concerned with the requirements of very large organizations. It is from the point of view of such organizations that we will evaluate a number of traditional information security management practices. We must begin therefore by describing how we define a 'very large organization'. We will also provide the definitions of some additional terms that we will employ within the paper.

Various institutions employ terminology to describe and categorize organizations of different sizes. For example, the European Union, the World Bank, and the World Trade Organization all use the term 'small and medium enterprise' (or SME). They collectively define an SME as an organization with less than 250 employees and either a turnover of less than 50 million euro or a balance sheet total of less than 43 million euro [1]. The U.S. Small Business Administration defines a 'small business' as an independent business having fewer than 500 employees [2]. Also in the U.S., the Inland Revenue Service defines a 'large corporation' as one that had taxable income of at least one million dollars for any of the three tax years immediately preceding the current tax year [3].

We can see from the above that an organization with operations in both the U.S. and Europe can be classified as an SME under the European Union scheme and simultaneously classified as a large corporation under the U.S. Inland Revenue scheme. This situation occurs because the two schemes are independent; they are separated by both geography and legal jurisdiction. Perhaps recognizing the challenges associated with the act of classification, the approach employed by the U.S. Census Bureau is not to define terms, but simply to provide data that enables the categorization of businesses in several ways. (In other words, the onus is on the user.) That data includes the number of employees that work in the organization, the industry type, receipt size, and many other items of interest [4].

The field of finance might appear to provide the tools that could be used to create a categorization scheme that will function globally. Financial markets are certainly global and foreign-exchange mechanisms allow the value of one country's currency to be described in terms of the currencies used in other countries. Financial markets can therefore be used to compare companies in terms of their market capitalization with those values normalized into one currency. A cut-o↵ can then be employed to identify companies with a certain minimum value. This is essentially the method employed for defining 'large-cap' stocks (meaning large capitalization). In the U.S. a large-cap stock is typically defined as one where the company has a market capitalization of more than five billion dollars [5].

It is not necessarily the case however that a certain *economic* value means that a company is large. A company with a high economic valuation could conceivably have a relatively small number of employes and be housed in a single building in a single city. A high-technology start-up that has had a recent public o↵ering of shares is one example of how such a scenario can occur.

A second issue with attempting to use a categorization scheme that employs a method based on an economic valuation is that numerous large organizations are not present in the market for stocks. Examples of such organizations are the governments of large countries, private companies, and non-profits. None of these organizations are public companies and therefore none are represented within equity markets. (Government bonds represent a proxy for the inherent value of the economy of a country and not for the government itself.) We want to include organizations of these various types in our analysis and therefore a finance-based scheme is not sufficiently inclusive for our purposes.

There is no widely agreed terminology for defining very large organizations, nor is there an existing scheme that would enable the sorts of organizations in which we are interested to be both categorized and identified. We prefer not to narrow the scope of our analysis and therefore the approach we will employ is to define our own terms for the purposes of this paper.

We will define a 'very large organization' as one that satisfies the following two

conditions. First, that the organization has workers numbering in the tens of thousands, with those workers being present on a number of continents. Second, that the organization is dependent on the capabilities provided by its computing infrastructure as an operational necessity. This second condition is not to say that the organization could *never* recover in some fashion from losing its information technology base, but that the impact would be both severe and long-lasting.

Information security is sometimes said to require a consideration of three fundamental aspects: people, process, and technology [6]. In our definition above we have addressed the people and technology aspects directly. We will assume that any organization that meets the two criteria we have described will necessarily employ policies and procedures in some manner and not note that aspect in our definition explicitly. Using our definition, most multinational corporations, global investment banks, governments of large nations, and large non-profit organizations are all considered to be very large organizations. The definition is therefore sufficiently broad for it to encompass a useful variety of organizations upon which to perform our analysis.

We will borrow two other terms that are in common use. The first is 'enterprise-scale'. This term is used colloquially within the information technology field and refers to the considerations that must be made when developing software that will operate within very large organizations. Enterprise-scale software typically has requirements that include the software being highly available, having the ability to handle a large number of simultaneous users, having a user population that is geographically diverse, and so on [7] [8]. We will use this term so that we can more easily refer to 'enterprise-scale' rather than 'the requirements imposed by operating within a very large organization'. Following from our definition of enterprise-scale we will also employ the term 'enterprise environment' as a synonym for 'the underlying computing environment that supports a very large organization'.

Having now defined terms we proceed in *§*2 with a description of some fundamental challenges faced by very large organizations.

3 The implications of operating at enterprise-scale

The management of an enterprise environment involves the consideration of a particular set of requirements and constraints. These primarily emerge from the matter of *scale*. At enterprise-scale a number of challenges are present that typically do not exist in smaller organizations. (If any are present in a smaller organization they will be at a much lower degree of difficulty.)

In this section we explore these challenges and the broader implications. We do so in order to provide a context for later sections that discuss the management of information security within very large organizations. We will first describe some of the key challenges then describe some common responses. Note that these topics are

not normally a direct security issue. Rather, they are technological or managerial challenges that have a security *implication*.

3.1 Key challenges

Within an enterprise environment there are simply a *lot* of 'moving parts'. There are a large number of data centers, servers, workstations, applications, network-aware devices, networking equipment, and so on. There are a large number of gateways where the internal network connects to external networks such as the Internet. A large amount of private network connectivity with partners, customers, and subsidiaries is also to be expected. Configuring the rule set of just a single firewall to avoid errors is known to be a challenging task [9] [10]. As a firewall rule set grows, the number of configuration errors increases sharply and the performance of the firewall degrades [11]. One of the challenges in an enterprise environment is to configure *hundreds* of firewalls in a manner that ensures their policies are logically consistent and that they do not contain errors that could compromise security. In addition to the challenge of managing the network perimeter there are tens of thousands of computers within the environment that must receive regular software updates such as security patches. These large numbers of computers must also be monitored to ensure they have not been compromised by hostile code or by other attacks. These are just a few examples of areas where security e↵orts are made more challenging by the complicating e↵ects of scale.

There are various factors that can drive the enterprise environment towards a high degree of heterogeneity in the computer hardware and software that is in use. It is difficult to impose purchasing requirements globally, especially when suppliers do not sell into every country in which the organization operates. So-called 'legacy' systems (meaning old) can become entrenched because of the high cost to replace them with more modern systems. The organization might also acquire companies whose technology environments then have to be integrated, and those companies may have employed a very di↵erent strategy in their past technology purchasing decisions.

From a security point of view, a heterogeneous environment might be perceived to have some advantages. A virus or worm that infects a particular flavor of operating system would not spread as rapidly or a↵ect as many computers as in a homogeneous environment [12]. The challenge with a heterogeneous environment is that it imposes a logistical burden in tracking newly announced vulnerabilities and understanding their potential impact. The task of testing and deploying security patches is also made more difficult because of the large amount of variety in the types of software being used.

Within the enterprise environment there is a very large amount of data. We can distinguish between "data at rest" (meaning data that is stored), "data in use" (meaning data that is actively being used by a computer program), and "data in

motion" (where the data is in transit across the network). All three types of data exist in large quantities within enterprise environments. A further distinction is sometimes made between "structured data" that is organized and stored according to a schema (such as in a relational database or a directory such as LDAP) and "unstructured data" that is found in presentations, spreadsheets, text files, and a large number of other formats. Unstructured data within an enterprise environment is usually stored in shared file directories, on wiki sites, and in document management systems such as Microsoft SharePoint.

Data has a tendency to proliferate and this has the effect of making the implementation of certain security controls and security processes more difficult. For example, a Chinese wall that is intended to restrict two classes of worker from interacting might be subverted by file systems or other data stores being used as a 'drop box'. Another example is where regulators require evidence that access to certain data is restricted. This can be a challenging task if that data has 'escaped' into unstructured formats that can then be freely and easily copied by workers.

In addition to there being a large number of technological components and a large quantity of data within enterprise environments, there are a large amount of human beings interacting with those various components and ultimately creating and altering the underlying data. Workers in very large organizations come in a variety of flavors: full-time employees who work on-site, employees who are full-time but work remotely (telecommuters), independent contractors, contractors provided through agencies, workers who are employed by another organization but are collaborating on a joint venture and are present on-site, and many other instances and permutations of working arrangement. Each worker and each category of worker has a set of resources within the organization to which they require access and that access must typically be limited to only those specific resources.

Workers may have remote access to the organization's internal network via the Internet or via dial-up connectivity. The ability of a worker to interact with the computer systems within the environment is therefore not restricted by the need to be physically present. Crucially, the set of workers who are permitted access to internal systems changes on a daily basis as workers join or leave the organization or when contracts begin or end. For these reasons it is very valuable to have the ability to quickly suspend all of the access associated with a particular individual. This is useful in the situation where a worker is fired in acrimonious circumstances or where the credentials used by a worker to access the environment remotely are believed to have become compromised in some way.

Per our earlier definition, a very large organization has operations that span a number of continents. Depending on the nature of the organization, in some countries in which the organization operates there may be a large number of small offices in order to provide a distributed local presence. An example is retail banking with

those offices being local branches of the bank. In 'hub' cities such as London, Tokyo, and New York, operations are probably restricted to a relatively small number of buildings, albeit with large numbers of workers inside those buildings. The result is a high degree of variation in the size of each computing facility, ranging from major data centers down to offices with just a handful of desktop machines. Of note here is that the available bandwidth for communicating with satellite offices might be very small. Connectivity might also only be available to those offices on an unpredictable and intermittent basis. This poses challenges for monitoring the security of those distant computers and in the remote installation of security patches and security software.

A large geographic span can also introduce challenges of internationalization and localization. Text within software such as instructions for the user must be converted into the native language in each country where English cannot be used. This could be because English is not spoken as a primary or secondary language in those countries or because country-specific laws require the native language to be used (this is the case with the French language in Quebec for example). The need for internationalization and localization can a↵ect timelines for the development and use of security software and the development and delivery of security training.

When an organization spans multiple countries it is a↵ected by the laws in each of those nations. There are particular implications for information security here, such as in the permissible use and potential 'export' of encryption technologies. Other laws that are known to vary between legal jurisdictions and that can a↵ect security controls include those relating to the monitoring of employees, retention requirements for digital records such as log files, and the laws surrounding the handling of personally identifiable information (also known as PII). For examples of the latter see the Swiss Federal Data Protection Act [13] and the UK Data Protection Act [14].

Two consistent themes in all of the above are the complicating e↵ect of scale and the e↵ect of a high rate of change over time. In summary, the enterprise environment is characterized by a torrent of new workers, new computer hardware, new software, and new policies and procedures.

3.2 Common responses

Having now described some of the key challenges that exist within very large organizations we will note some common responses to those challenges. These approaches are not widely documented in the literature, nor in the popular press. In most cases this is either because the approach was developed in-house, or because the approach is not novel per se, but rather a collection of standard practices that have been combined to produce a novel result. Not all enterprise-scale organizations have implemented these techniques and this could be for various reasons: they might be unaware of them, the techniques might not be suitable for use in their specific environment, or

they might not have completed the work or are struggling with the implementation. Additionally, commercial products that provide these capabilities might simply not yet exist or the current products might not be able to function at the scale required.

An important technique for attempting to address challenges of scale, of heterogeneity, and of a high rate of change, is to place an emphasis on collecting, maintaining, and using metadata. Metadata is "data about data" (or see the more formal definition from ISO/IEC [15]). The responsibility for maintaining metadata may rest with a single dedicated group within the organization or the responsibility for maintaining metadata of particular types might be held by the relevant internal teams.

The type of data that is of particular interest from a security perspective includes inventories of computer devices such as servers and workstations, human resources data such as join-dates and leave-dates for workers, and data concerning application software within the environment – who the 'owner' of the application is, what type of data is processed by the application, and so on.

By combining these various types of data, a type of triangulation can be performed. This allows control logic to be defined such as "ensure that applications that are identified as being in production only run on registered servers" or "ensure that only full-time employees are allowed access to confidential data" [16]. Enforcing that control logic at a higher layer within the environment is in some regards a more tractable problem than trying to enforce the same control at lower layers.

As the scale of the computing environment increases, so does the degree of difficulty in tackling even basic engineering and operational tasks. To administer tens of thousands of computer servers by hand is simply not practical, just as one example. In a similar manner to the way in which mechanization was essential in the transition from craft production to mass production, so has automation become central to enterprise-scale computing. We will note three different types of automation that are used to help tackle the previously described challenges and discuss each in turn. These are: configuration management, change management, and entitlements management.

> *Configuration management* is focused on ensuring that computing devices such as servers, workstations, and network devices such as switches and routers are all configured as desired. "As desired" is normally defined by a policy statement that describes specific goals. Configuration management is not just concerned with the *initial* configuration of a computing device but also with the maintenance of that desired configuration over time. Computer systems that are used day-to-day tend to experience a type of decay where their configuration drifts from its initial state and that drift can introduce security vulnerabilities [17]. The goal of configuration management is to ensure that the initial state is correct and then to help stop that configuration from slipping over time.
>
> *Change management* is concerned with testing changes before they are imple-

mented within the environment. The goal is to avoid making changes that would cause problems such as slowing a system down or causing a crash. Most enterprise-scale organizations create multiple computing environments labeled as either a 'production' or 'non-production' environment. Production environments contain live systems such as those used by customers. Non-production environments contain systems used for development and testing. To 'promote' code from non-production to production requires a formal process to be followed where proof of adequate testing must be provided. These safeguards also help to reduce security risk by limiting the set of workers that can make changes to production systems.

Over time, security vulnerabilities are identified in software and the software vendors who created that software release security patches for those vulnerabilities. Change management is therefore an important process when testing and deploying those security patches. When considering the installation of security patches there must be a consideration of the need to balance 'patch risk' (the risk that a patch will break a functioning system) against 'vulnerability risk' (the risk that an attacker will exploit the vulnerability before the patch is applied). For an in-depth exploration of this topic see the paper "Timing the Application of Security Patches for Optimal Uptime" [18].

Entitlement management is the set of technologies and processes that attempt to ensure that only those workers who should have access to computer systems are able to access those systems. There are two fundamental types of entitlement: business entitlements and technology entitlements. Business entitlements enable workers to access business applications. An example is an equities trader within an investment bank who is permitted to access a system that provides pricing information regarding equities. Technology entitlements enable technologists to access the underlying computing infrastructure that supports business applications, such as a Unix machine or the administrative interface to an application.

As workers move around the organization they tend to accumulate entitlements – their old access is not taken away when they transition into their new role. As a result they gain more and more access to resources over time. This situation creates risk and can violate segregation of duties policies. When workers leave the organization their access can also persist even though that access is no longer required. Enterprise-scale organizations face a particularly bad calculus with regards to this problem. A large number of systems and applications combined with a large number of users (workers) can create a large amount of redundant entitlements over time. To tackle this problem some enterprise-scale organizations have built systems and put processes in place to remove access that is no longer required. This is the key role played by entitlement

management. (The task of *granting* access to workers is typically the domain of identity management and provisioning systems.) On a periodic basis such as every three or six months, an 'entitlement review' is carried out that requires system owners and application owners to review the access to their system or application. These entitlement reviews are typically automated using a web-based workflow. Participants have a certain window of time to perform their reviews and entitlements that are identified as no longer being required are purged. If an entitlement is not flagged as being required then it is removed – a type of default-deny. Auditors and other interested parties can review the results of the review to verify compliance.

In the area of entitlement management some commercial products have begun to enter the marketplace that provide similar capabilities to systems that enterprise-scale organizations have built for themselves [19].

We have now identified some of the key challenges faced by very large organizations plus some common responses to those challenges. In later sections that examine security management practices we will consider the degree to which those practices help to mitigate these problems or perhaps exacerbate them.

4 Literature review

There are a relatively small number of academic journals that focus directly on information security. Those journals also tend to focus on the technical aspects of information security or on the mathematical aspects (such as with cryptography) rather than on the management of information security.

We performed a search using the Microsoft Academic Search service to identify journals relating to information security. Microsoft Academic Search is an online search engine service provided by Microsoft that covers more than 48 million academic papers [20]. This search identified 19 journals that were then reviewed to identify those with a history of published articles relating to security management. Most journals contained no articles relating to security management, such as the *Journal of Cryptology* and the *Journal of Computer Security*. Some journals contained only a small number of articles relating to the management of information security, such as *Information Management & Computer Security*.

However, one of the most significant journals in the field does have an ongoing history of publishing articles relating to various aspects of security management. *Computers & Security* has been continuously published since 1982 and has a 5-year impact factor of 1.075 [21].

We then performed a search for articles published in Computers & Security relating to the management of information security. Articles that included the word 'management' in their title but where the topic of the article was not some aspect of

security management were not included, such as with articles relating to certificate management or database management. There were 59 articles meeting these criteria published between 1982 and 2012. The eight year period from 1982 to 1990 contained 18 articles relating to the management of information security. From 1990 to 2000 there were 26 articles. From 2000 to 2012 there were 15 articles. The number of articles can be seen to peak in the 1990s, but this might simply reflect the desires of the editor to publish papers of a particular type. We can also examine the subject matter of each article and in doing so we can identify what appears to be three periods in which articles about security management share similar themes.

Articles published in the initial ten year period from 1984 to 1994 stress the need for information security to be considered a key concern of senior management. The articles also propose strategies for requesting funding for security projects. See for example: “Data security is a management responsibility” [22], “How security fits in – a management view” [23], “Discussing security with top management” [24], “Security systems: getting management to shell out” [25], and “Getting management buy-in to IT Security” [26].

The concerns regarding management support and funding for security projects appear to be reduced by the beginning of the second period. Beginning in 1997 and running for six years until 2003, the majority of articles in this period focus on the *frameworks* employed to deliver information security goals. See for example: “New organizational forms for information security management” [27], “A conceptual framework for information security management” [28], “Information security management: a hierarchical framework for various approaches” [29], and “An integral framework for information systems security management” [30].

Articles post-2003 tend to focus on particular aspects of security management and this trend continues to the current day. Examples include “Improving the ROI of the security management process” [31] and “Holistic security management framework applied in electronic commerce” [32].

5 Gap analysis

In this section we perform an analysis of a number of security management practices. The various security management practices that we will consider are commonplace. Our examination is focused on the question of whether these practices are both useful and practical when employed within very large organizations.

5.1 Security management frameworks

There are a number of well-established frameworks that are often used to structure and guide the implementation of a security management program within organizations. COBIT and BS 7799 / ISO/IEC 27001 are two examples [33] [34]. Over

time, new frameworks emerge and become popular for various reasons. For example, frameworks are sometimes used to satisfy specific legislative requirements that are placed on organizations. The development of COBIT was largely the result of the U.S. Sarbanes-Oxley act of 2002 and in particular section 404 of the act that requires top-down risk assessments and the use of internal controls around IT systems.

The academic literature contains a number of papers that attempt to determine the strengths and weaknesses of certain security management frameworks, whether frameworks are compatible or can be integrated, and various other types of analysis [35] [36]. Some studies have directly investigated the efficacy of security management frameworks, i.e. whether the use of a particular framework can deliver measurable improvements within an organization. Gene Kim studied 25 performance measures across a number of different organizations [37]. Organizations were then examined that followed some or all of the 63 practices recommended in the version of COBIT that was available at the time of the study. The finding was that 21 of the 63 COBIT practices had a major impact on performance with 4 being key. If this finding is correct then two thirds of COBIT recommended practices have no demonstrable value and only 4 are critical.

It may be that the benefit received from using a security management framework such as COBIT is as much in the signal that it sends to other parties that interact with the organization, versus any direct benefits received through the use of the framework itself. An organization might find the use of a well-known framework to be a convenient way to demonstrate to auditors and customers that it is taking a structured approach to the management of information security. The use of a common framework can also create a shared view and a shared vocabulary. There may be a network effect in the use of particular frameworks: as a framework becomes more popular it becomes more widely recognized, and as it becomes more widely recognized it becomes more popular, with this feedback loop feeding in on itself. It is conceivable that a security management framework could become widely used because of network effects such as this, even if it is not effective at improving security![1]

As noted above, COBIT was designed primarily to address specific legislative requirements. Other frameworks are concerned with the management of information security more generally. BS 7799 is considered to be the first significant attempt to formalize security management guidance. BS 7799 was well-received and gathered a great deal of momentum within the commercial world through its certification ecosystem. BS 7799 was later used as the basis for the creation of ISO/IEC standard 27001 which is also in widespread use. It is therefore a useful exercise within the context of this paper to consider the extent to which BS 7799 / ISO/IEC 27001

[1] This thought brings to mind undergraduate education. Other than in some specialized fields of study, the value in an undergraduate degree is perhaps not in what is learnt in pursuing the syllabus but rather the signal to prospective employers that the person is able to complete a multi-year period of study at the undergraduate level; i.e. as a proxy for the ability to learn.

can add value to the task of implementing security management within a very large organization.

BS 7799 is primarily focused on the process of constructing what it describes as an ISMS or *Information Security Management System*. The purpose of the ISMS within an organization is to define the various strategic and operational processes that will deliver the security capabilities that are required by the organization. In addition to describing the core ISMS mechanism, BS 7799 lists a number of potential controls. These are described as an "*information security starting point*" where the listed controls "*are either based on essential legislative requirements or considered to be best practice for information security.*"

The authors of the first version of BS 7799 were a combination of UK government departments such as the Department of Trade and Industry and commercial companies such as Midland Bank plc and Marks & Spencer plc. There was no academic representation within the authors at the level of a university or an academic working group. Perhaps as a result of this fact, the BS 7799 document does not cite any references and there is no mapping between the controls recommended within BS 7799 and the academic literature. The guidance within BS 7799 is essentially the best guess of the authors as to what is desirable and e↵ective. This is a "brain trust" approach rather than one that is based on findings from controlled studies. Such an approach might be a reasonable one in order to overcome the bootstrap problem. It would be an interesting exercise however, to determine what percentage of organizations that follow the BS 7799 or ISO/IEC 27001 controls guidance are aware that those recommendations were created almost 20 years ago. The list of controls within BS 7799 must necessarily have reflected the pressures and circumstances of the time at which the initial version of the document was created – in 1995. The technology landscape and the threat environment have both evolved substantially within the last 20 years and the list of recommended controls within BS 7799 has only received relatively minor updates during that 20 year period.

Similar observations can be made when examining other security management frameworks. A more recent security management framework than BS 7799 / ISO/IEC 27001 is NIST Special Publication 800-100: Information Security Handbook: A Guide for Managers [38]. Intended for use by U.S. federal agencies, the document describes a number of 'program elements' that agencies should focus their security management e↵orts upon. Those program elements are selected "*based on the laws and regulations relevant to information security, including the Clinger-Cohen Act of 1996, the Federal Information Security Management Act (FISMA) of 2002, and Office of Management and Budget (OMB) Circular A-130*". It is not clear how the specific guidance in the document was derived from those laws. For example, NIST Special Publication 800-100 recommends that an organization implement a security awareness and training program (section 4; page 26 of the document). The various laws noted above make no

such recommendation. The authors of NIST Special Publication 800-100 appear to have attempted to translate the spirit of those various pieces of legislation into lower level activities that are actionable. The process by which this was accomplished is not specified. There is no description of the discussions that must have taken place around prioritization and the tradeo↵s that led to the specific guidance being selected. Why for example was incident response included as a top-level category but not security monitoring? The act of security monitoring would seem to be a prerequisite for detecting security incidents that could then be responded to.

The approach taken in the creation of NIST Special Publication 800-100 appears to be the same as with BS 7799 – the specific guidance represents the best guess by the authors of the document. The guidance is not based on findings from within the literature, or at a minimum the sources used to construct the guidance are not quoted as references within the document.

A second key observation that can be made regarding security management framework documents such as COBIT, BS 7799 / ISO/IEC 27001, and NIST Special Publication 800-100 is that they are designed to satisfy the general case, i.e. to benefit the average organization. This is deliberate on the part of the authors so that their work can have the largest possible audience and – in theory – provide the most widespread value. It is unlikely that such generic guidance can apply equally to all types and sizes of organization. It is more likely that the usefulness of generic guidance degrades as the target organization moves away from the platonic form that the framework authors had in mind. At some point the target organization is sufficiently di↵erent from the average organization that the guidance stops adding value and has a negative e↵ect. A simple example would be a control that is more expensive to implement in the target organization than the maximum loss that could occur if no control existed.

There are numerous axis upon which an organization might be di↵erent from the type of organization that the authors of management framework documents envisioned. The organization might be very small or very large, it might be highly centralized or highly distributed, it might operate in a country with very particular regulatory compliance requirements or have no pressing requirements at all in those areas.

In this paper we are specifically concerned with the requirements of very large organizations. It can be seen that much of the fundamental guidance provided in security management framework documents is very difficult to accomplish within such large organizations. One example is section 4.2.1 of ISO/IEC 27001 which describes the risk assessment process that will dictate what controls should be implemented: "*Identify the assets within the scope of the ISMS and the owners of those assets the threats to those assets the vulnerabilities that might be exploited by the threats [and] the impacts that losses of confidentiality, integrity, and availability may have on the assets.*" §3.1 has described the very significant challenges within an enterprise

environment that make accomplishing this 'basic' risk assessment process extremely difficult. This is true even if the scope of the ISMS is relatively small. A number of other examples are provided below in §5.2.

There are several problem areas that have been identified through experience as common reasons why implementations of BS 7799 / ISO/IEC 27001 fail. These include the scope of the implementation being too big, an inability to identify key assets, a failure to identify shared assets such as shared infrastructure, and not identifying dependencies such as third parties and outsourced services [39]. For the reasons described in section §3.1, these specific tasks are some of the *most* difficult to accomplish in an enterprise-scale environment. The authors of ISO/IEC 27002 which is a supplementary document to ISO/EIC 27001 list a 'critical success factor' for implementation of ISO/IEC 27001 within an organization as "*visible support and commitment from all levels of management*". It is difficult to imagine what management project would *not* be successful if that guarantee was met. As such, it is a rather obvious 'get out of jail free' clause on the part of the authors.

Lastly, if a manager within a very large organization were to rely primarily on the guidance within security management frameworks such as BS 7799 / ISO/IEC 27001 and NIST Special Publication 800-100, they would miss the critical importance of activities such as configuration management, change management, and entitlement management, as described in §3.2. The challenges described in §3.1 are some of the most pressing and most difficult faced by enterprise-scale organizations and framework documents do not address them directly.

We continue our analysis below by reviewing guidance that is common across various security management frameworks.

5.2 End-user security awareness training

The first specific security management practice that we will evaluate is security awareness training. The goal of security awareness training is to provide guidance to workers so that they do not perform actions that would expose the organization to information security risks.

Techniques to defend against 'social engineering' are an example of a topic that is traditionally taught in a security awareness training class. Social engineering in the context of computer security refers to an attacker manipulating a worker into performing an action that would assist the attacker in carrying out an attack. An example is where an attacker persuades a worker to reset the password of a legitimate user and then uses that credential to access the system [40] [41].

In the popular press it is a well-worn trope that human beings are the "weakest link" when considering security measures at the unit of an organization. See for example the articles 'Humans: The Weakest Link In Information Security' [42] and 'How to toughen the weakest link in the security chain' [43]. The idea that people are the

weakest link is attractive, perhaps because the sophistication of security technologies is ever-increasing but human beings are viewed as being endlessly fallible.

The logic of end-user security awareness training is as follows: workers may unwittingly perform actions that weaken security because they do not understand the security *implications* of those actions; therefore workers should be educated so that they have the knowledge to not perform those actions. If human actions allow technical security controls to be circumvented or weakened then it follows that there should be an element within a security program to attempt to influence or control those human actions. The canonical example here is with user passwords. A password cannot be considered secure if the user writes their password on a post-it note and attaches it to the underside of their keyboard. Security awareness training would seek to educate the user about the security implications of writing down their password so that they do not introduce that risk into the environment.

Security awareness training is a staple of framework documents that provide guidance regarding the development, structure, and operation of a security program. In performing the research for this paper we were unable to identify *any* general frameworks for managing information security that do not include the recommendation to employ security awareness training.

All of the following notable frameworks for managing information security contain the recommendation to develop and employ a security awareness and training program: NIST Special Publication 800-100, COBIT, and BS 7799 / ISO/IEC 27001. The authors of NIST special publication 800-100 are particularly exclusive in describing the level of importance they place on security awareness training. Quoting from the document: "*The security awareness and training program is a critical component of the information security program. ... In terms of the total security solution, the importance of the workforce in achieving information security goals and the importance of training as a countermeasure cannot be overstated.*" We can consider such arguments in support of security awareness training as having one of two forms: a strong-form and a weaker-form. The strong-form argument of which the above text from NIST special publication 800-100 is one example, is that security awareness training is both essential and non-fungible. In other words, that security awareness training should always take place and that it has no substitute. Policies of this type are in widespread use. An example is the U.S. Office of Management and Budget who have published a legal requirement that all U.S. government agencies ensure that security training is delivered to all workers before they are granted access to computer systems [44].

The weaker-form of the argument is that security awareness training continues to be important but the most value that an organization can receive is when the training is focused on specific populations of workers within the organization such as software developers [45]. Typically, the weaker-form of the argument also intro-

duces the concession that security awareness training will reduce but not eliminate instances where workers make decisions that impact security. The downgrade from the strong-form to the weaker-form is perhaps a response to a number of studies that have shown that security awareness training is not fully effective at removing user error. This is generally understood to be because workers sometimes simply make mistakes. The finding also exists that workers are not particularly motivated to learn about security (viewing it as a secondary task that is separate to their main job responsibilities) and that it is difficult to teach workers how to identify security threats without also increasing their tendency to misjudge non-threats as threats [46]. There are many reasons why workers might make mistakes even after they have received training. People can experience stress and tiredness, they can forget to eat or become distracted, and for lots of other reasons. As a result, the amount of protection that security awareness training can deliver is asymptotic – it never reaches 100%.

Consider this asymptotic property in conjunction with the asymmetric nature of attacking and defending. An organization must defend against all attacks but an attacker requires only one attack to be successful. It follows that over an adequate time period an attacker will always be successful in subverting any particular control that is not 100% effective. Certainly, it is a rare security control that is actually expected to be 100% effective. The use of a defense in depth strategy, the use of compensating controls, and the approach of balancing security efforts across the categories of protection, detection, and response all exist because it is unlikely that any single control will be sufficient to provide adequate security. This approach aligns with the modern view of information security as a risk management function [47] [48]. The desire is to understand the costs and benefits of individual controls. The objective is to maintain an appropriate level of security rather than overspending to create an unnecessarily high level of security. (An excessive number of security controls might also inhibit the ability of the organization to function efficiently.) We must therefore consider the costs and benefits of security awareness training for our particular scenario of interest, namely very large organizations.

Given the nature of the types of attack that end-user security awareness training is designed to protect against, and given the challenges just described regarding the fallibility of human actors, it is certain that some attacks will slip through. In a very large organization the pure numeric quantity of attacks that are successful could be substantial. Studies have shown that attacks that target workers such as phishing attacks have a success rate of around 5% [49] [50]. (Phishing attacks are where an attacker attempts to acquire information by pretending to be a trustworthy entity over email or some other form of electronic communication.) This means that a very large organization with 20,000 workers can expect a failure rate of perhaps 1,000 cases. 50,000 workers increases the failure rate to 2,500 cases. These numbers are a simplification because they assume that every worker will receive the phishing email,

which is unlikely. An attacker might perform multiple rounds of attack however, where the same user has the opportunity to make the incorrect decision on multiple occasions.

Security failures at this scale – where there are thousands of instances where the security measure failed – place a very large burden on the security operations and incident response capabilities within the organization to detect, investigate, and clean up after successful attacks. If a security technology such as a firewall or a piece of remote access software failed with a similar frequency requiring a similar amount of clean up, the organization would likely look to replace it with an alternative technology or approach. Worse still, when a worker provides their password over the phone to an attacker there is no record of this event.

In a small organization a small per-worker failure rate is itself a small number and might be considered manageable. In a very large organization a small per-worker failure rate translates into thousands of successful attacks which then sap resources.

We have described in §3.1 the costs associated with internationalizing content. These costs are also applicable to the development of security awareness training – to support required languages, ensure correctness with regards to local laws, and to accommodate cultural norms in the various countries in which the enterprise has a physical presence. Security awareness can therefore also be expensive to develop and deploy within an enterprise-scale organization. Given these observations, it is unclear if the cost/benefit for end-user security awareness training must always be worthwhile.

5.3 Security policies

As with security awareness training, the doctrinal view is that the development and use of security policies is a key component in the management of information security within an organization. Michael Whitman and Herbert Mattord make the claim that: "*Quality information security begins and ends with quality corporate policies*" [51].

Fundamentally, policies attempt to dictate behavior. Kerry-Lynn Thomson and Rossouw von describe the purpose of security policies as follows: "*The main aim of any policy, whether for information security or not, is to influence and determine decisions, actions, and other issues, by specifying what behavior is acceptable and what behavior is unacceptable*" [52].

At enterprise-scale the same types of problem that we have just described with regard to security awareness training also exist with regard to security policies. A small percentage of workers within a large organization that are unaware of the security policy and therefore do not act in accordance with the policy translates into a large number of potential policy violations (on the order of several thousand workers given our earlier definition of the size of a very large organization).

There are a number of other challenges that concern the use of security policies

within an enterprise environment. A worker cannot comply with a security policy unless they have read the policy document or are somehow made aware of the content of the policy. The logistics of putting a policy in front of tens of thousands of workers – with all of the associated localization challenges and costs – is non-trivial. The opportunity cost of having every worker spend several minutes reading a policy and perhaps taking an online quiz to demonstrate their understanding of the policy must also be considered significant. An organization with 50,000 employees that requires all employees to spend 15 minutes reviewing a policy consumes more than 12,000 hours of worker time.

A peculiar aspect of a number of papers written about security policies within the academic literature is that they make the statement that *every* worker within the organization is responsible for information security. For example, Basie von Solms states: "*Information Security involves everyone in a company - from the Chairman of the Board right through to the data entry clerk on the shop floor and the driver of the vehicle delivering products to the customers*" [53].

It is surprising that such a sweeping claim survived the peer review process. Realistically it cannot be true that the level of information security within an organization depends in any meaningful way on the junior cook who works in the company cafeteria, or on many other types of worker. Perhaps on balance it is a more efficient strategy to make a sweeping statement, a generalization, rather to have to specify which workers do or do not have responsibility for information security. The side effect of that approach is probably to make workers view policy statements less seriously, and this might undermine other, more important security guidance. (A type of 'crying wolf' perhaps.)

Rather than create security policy documents from scratch it is common for organizations to use widely available policy templates. See for example the content made available by the SANS organization [54] and CSO magazine [55]. The "best practices" as represented in these templates is often not applicable to very large organizations. For example, the advice to "confront people in the office you do not recognize" might be useful in a small or medium-size office but it is not practical in a downtown Manhattan office with a high amount of foot traffic and thousands of employees on various floors. This is another example of the observation made in the introduction to §5 that "best practice" guidance that might be reasonable in the general case might simply not be practical at enterprise-scale.

5.4 The security organization

In §4 we identified a ten year period spanning from 1984 to 1994 where published articles on the topic of the management of information security focused largely on the need for information security to be considered a key concern of senior management. A number of articles within that period also proposed strategies for requesting funding

for information security projects.

These various articles likely performed a useful service. It was valuable at that time to highlight the fact that organizations were introducing themselves to various emerging security risks if they did not recognize the need for information security measures and take a structured approach to managing information security. We can perhaps consider this period as the time when information security as a discipline was recognized and brought up to the level of other important topics that an organization should consider. The emphasis was still largely on the technical however, with information security usually being a responsibility of the IT group within the organization.

Today it can reasonably be expected that most enterprise-scale organizations have a senior manager with responsibility for information security. As noted in §3.1, large organizations have been assailed in recent years by a raft of legislation such as Sarbanes Oxley [56], HIPAA (the Health Insurance Portability and Accountability Act) [57], PCI (Payment Card Industry Standards) [58], and many others. Many of the efforts to meet these various standards within the enterprise involve information security measures, and so it is to be expected that there will be a large number of employees working on compliance, on IT Security, and on classical information security projects. These employees might not roll-up under the same top-level manager however.

Much has been written about which organizational structures are "best". A perennial criticism of upper and middle management within large organizations is that they are said to tinker with the organizational chart too frequently, trying to tweak it to improve performance but perhaps leaving employees constantly unsettled by those frequent changes. Such is life within a large organization. In the commercial press there continues to be articles published that discuss various organization structures for security. See for example the article "Who should the CISO report to?" in CSO Magazine [59].

In addition to assigning managerial responsibility for information security, it is safe to say that the majority of enterprise-scale organizations have implemented the standard set of IT Security technologies, that they have responded to legislation such as Sarbanes-Oxley by implementing compliance programs, and that they continue to invest in maintaining those efforts. It is unlikely that there are many enterprise-scale organizations that have not invested in the various practices that are now considered standard. Battles within organizations might be fought over the *amount* of spending on information security, or on the specific technologies and practices upon which money will or will not be spent, but it is unlikely that there are internal battles over whether some funds *should* be spent on information security. Some amount of spending on information security within very large organizations is considered a cost of doing business.

For the above reasons we can say that considerable progress has been made since the early period. This observation is reflected in the change in the types of journal articles published after 2003. As noted in §4 there was a progressive shift in focus from articles focused on funding and establishing the security organization to articles that proposed how to optimize particular aspects of security management.

We will now evaluate three key structural aspects of the management of information security. The first is who should have management responsibility for information security within a very large organization – for convenience we will refer to that person as the 'security leader'. The second is the question of the optimal location for the security leader within the organizational chart. The third is the makeup and structure of the security organization that reports into the security leader.

As noted above, within a very large organization there are likely to be hundreds of employees working on information security, IT Security, and compliance-related initiatives. This suggests that the security leader should have a strong general management skill set. In some organizations the security leader achieved that position simply because they were the first person to champion information security within the organization. Such a person might have strong security knowledge and skills but perhaps might not have strong general management skills. Some people are able to develop management skills over time as the result of being in a leadership role, but some people are unable to do so.

A person with strong security knowledge and skills can identify the security weaknesses in a technology or system, either intuitively or through the application of skill. This is a useful ability for a security expert to have. Perversely however, in a security leader those same skills could potentially be detrimental. The so-called 'security mindset' enables a person to perceive weaknesses, threats, and risks. When not tempered by pragmatism, the security mindset can obscure other important factors that must be considered such as the cost of defensive measures and the opportunity costs of implementing them. The security leader within a very large organization must operate on a plane above the detail of individual vulnerabilities and the threat de jour. Their focus must necessarily be on the strategic rather than on the operational and tactical.

The need to manage hundreds of members of sta↵ and the need to function at a high-level within the organization suggests that a security leader with a strong general management skill set is preferable over a security specialist.

The second key topic is the location of the security leader within the organizational hierarchy. Conventional wisdom holds that the security leader should not report into the CIO (the Chief Information Officer and leader of the IT organization). The justification for this guidance is that if the security leader were to report into the CIO then the security leader might have a disincentive to report weaknesses in the technology that is in use within the organization. This is probably correct. The

IT audit function is typically held distinct from the IT organization for the same reasons. There is no apparent reason why the security leader and the CIO cannot report into the same manager however. The question of precisely where the security leader should report into within the organizational hierarchy is probably best dictated by the requirements of each individual organization.

Related to the question of where the security leader should report is the question of their title. A person with management responsibility for a specific and important area within an enterprise-scale organization is typically given a corporate or 'C-level' title. Examples include CIO (Chief Information Officer), COO (Chief Operating Officer), and CFO (Chief Financial Officer). There has been a proliferation of titular C-level titles in recent years with the existence of somewhat dubious titles such as CVO (Chief Visionary Officer) and CWO (Chief Web Officer) [60].

Within a very large organization there can realistically only be a limited number of individuals who sit on the management committee and report directly to the CEO. In a very large organization the security leader is very unlikely to be a member of that group, even with the CSO (Chief Security Officer) title. The responsibilities of managing information security cannot reasonably be considered to be at the same level of importance as managing the entire IT organization (CIO) or managing all of the finances of the organization (CFO).

The third key question concerns the structure of the security organization. The security organization within most enterprise-scale organizations is discipline-based. Workers in this model are aligned according to their core skill set or role, such as security engineering, security operations, security compliance, security consulting, security investigations, and so on. The granularity at which these various teams are defined is determined by the needs of the particular organization – the security investigations team could simply be part of the security operations team for example.

This approach is fundamentally oriented to deliver services to the IT organization – the security team engineers security solutions, there is a security operations group to operate them, and so on. It perhaps reflects the fact that in many organizations the security team grew out of the IT department. The concern here is that most organizations are not pure technology companies such as software vendors. Typically the role of the technology side of the organization is to support the 'business' side of the organization, which is the side that actually performs the specific work that allows the organization to act as a going concern. It is considered a cardinal sin if the technology organization does not listen-to and support the needs of the business. (The ability to deliver what the business wants is a different question that requires the consideration of what is feasible. The business might also not know what it wants, or what can reasonably be accomplished within a certain time and budget.)

The discipline-based security organization is not best structured to identify and understand the needs of the business. This is because the discipline-based security

organization is technology-oriented. As an example: it is improbable that a member of a security engineering or security operations team is meeting regularly with an investment banker in the Equities trading division within an investment bank. In order to close this gap some enterprise-scale organizations have established their own network of contacts inside the business organization. A BUISO or Business Unit Information Security Officer is a worker who has been given certain responsibilities relating to information security in addition to their day-to-day role. The job of being a BUSIO in most cases is not full-time and might constitute only a fraction of the worker's overall responsibilities. A BUSIO provides information to the security organization on a regular basis so that the security organization can understand the requirements of the business. A BUISO also disseminates security information into the business unit and acts as a focal point for queries relating to security. The BUSIO model can be compared to the practice of embedding a press reporter within a military unit. Without such a model the press is reduced to filing reports from their local bureau, potentially thousands of miles from the front line.

In the BUISO model we see in one sense a decomposition of the security organization. The responsibility for information security is partly taken on by the various business units themselves, i.e. by the BUISOs within each business unit. It is useful to consider extending this idea and the approach more broadly.

It can be seen that the skill set required to operate and manage many technologies shifts over time from initially being highly specialized to eventually becoming commoditized. This progression usually goes hand in hand with the technology becoming more widely understood and easier to manage. No longer the sole purview of highly trained specialists, the technology can be understood and operated by workers with a much lower level of expertise in that particular skill. One example is operating system administration. To administer a mainframe computer in the 1980s required a specialist who was dedicated to that topic whereas the skills required today to administer a Linux server are relatively common today even amongst general programmers.

The same progression can be seen in security technologies. It is unlikely within a very large organization that firewall management today is performed by the IT Security team. It is much more likely that a firewall is viewed as simply a piece of networking equipment and is managed by the network operations team. Network security monitoring is no different from a structural perspective to general network monitoring. Indeed, security monitoring is considered by some organizations to be such a well-understood and commoditized activity that they outsource it to an external vendor. The industry category of 'managed security services' is well-defined and includes activities such as network perimeter vulnerability scanning, monitoring of intrusion detection systems, and such like [61]. A number of large technology vendors provide managed security services such as IBM, Dell, and Symantec [62] [63] [64].

Division of labor is a fundamental principle that underpins all management ef-

forts. The computer security field is now too broad for any one individual to retain meaningful expertise in all of its di↵erent facets. Where a skill can be commoditized it makes most sense for the worker with those responsibilities to roll-up into the appropriate functional team. The most e↵ective location within the organization for a Unix security engineer is inside the Unix engineering team; the most e↵ective location within the organization for a policy author is inside the team with responsibility for policy and compliance e↵orts.

These observations suggest that the dissemination of some of the traditional responsibilities of a security team could be a legitimate management strategy. An organization is by definition an integrated system. As such, security responsibilities should be integrated throughout the organization and not retained behind the castle walls of a centralized security team. If we look to history we can see processes such as 'Romanization' where independent states that had been conquered by the Roman Empire were slowly assimilated and ultimately adopted Roman cultural norms. This process of acculturation where one culture takes on the practices of a second culture is the fundamental mission of a security team: to inculcate security-conscious ways of working into the organization at-large. It may be that the most e↵ective way to accomplish this is not for these ideas to be imposed as a doctrine, but rather for them to be grown from within.

The IT industry has seen a number of oscillations over the last forty years between centralized and decentralized computing. From mainframe computers came the PC revolution, the client-server model, and now cloud computing and the proliferation of tablet devices. It is remarkable given the degree of change within the technology industry that the standard approach for managing information – building and operating a central security team – has remained so static.

5.5 Budgeting

A central and important managerial responsibility is budgeting. There are typically a number of projects that could potentially be carried out at any one time. With only a limited amount of budget funds available, a decision must be made about which projects to pursue. Budgeting is the process of evaluating the various projects that could be performed, weighing their pros and cons, and ultimately identifying the projects that will receive funds. In enterprise-scale organizations budgeting is typically a complex and lengthy exercise because it involves the whole organization. As a result the budgeting process tends to be carried out on an annual basis.

Budgeting decisions have essentially three possible outcomes: the decision to *spend* on a particular item, the decision to *delay* spending on the item until some future date or when certain future conditions arise, or the decision to *forgo* spending on the item altogether [65]. In addition to evaluating potential new projects there are most likely existing projects that also require consideration as to whether they will

continue to receive budget funds. The decision to fund a particular budget item can result in *capital spending* such as on IT hardware, on *operating spending* such as on software licenses and personnel, or more likely a mix of both capital and operating spending.

Budgeting decisions are by nature multivariate decisions. They require close consideration of the operating environment of the organization including matters such as the availability of funds, significant recent or upcoming business events, the macroeconomic climate, the current risk tolerance of the organization, and many other factors.

Budgeting decisions should be *rational* decisions and yet it is not unusual to see attempts to play on emotions in the commercial world during the sales process. Some information security professionals use the term 'FUD' (meaning *fear, uncertainty and doubt*) to refer to efforts by commercial vendors to try and sell security products and services by exploiting the fear of attacks [66] [67]. The seemingly never-ending procession of new vulnerabilities and the availability of exploits that take advantage of those vulnerabilities provides the ammunition for these sales tactics.

FUD works because human beings are not purely rational actors. It is understandable that the security leader within an organization that has just suffered a major security breach would feel the need to act quickly in response. The security leader might also feel the need to be *seen* to be acting decisively. The risk here is that spending decisions might not be made with the necessary consideration of other factors and that over-spending could result. Spending in an attempt to prevent the recurrence of a particularly painful event such as a security breach might also lead to "fighting the last war " – a military adage that refers to investing in solutions that would have been useful in the past but might not necessarily be useful in the future. There are a number of ways in which fear and other emotions can negatively impact decision making [68] [69].

Within the academic literature a number of papers have provided an economic perspective on the management of information security and on related matters that fall under the banner of economics [70] [71] [72]. Examples of this type of work includes studies that investigate the incentives of the actors involved in the security spending decision and the externalities created by spending decisions. The primary venue for published papers on these topics to-date has been the *Workshop on the Economics of Information Security* [73].

Out of that body of work various models and strategies have been proposed for assisting in budgeting decisions that concern information security. Lawrence Gordon and Martin Loeb present an approach for identifying the optimal amount to invest to protect a single asset [74]. Their finding is memorable in that their recommendation is that an organization should not spend more to protect an asset than 37% of the potential costs that would be generated through the loss of that asset. This result is due to the observation that any spending to protect an asset has diminishing marginal

returns. Subsequent work has suggested enhancements to the Gordon & Loeb model and has extended the model [75] [76]. W. Sonnenreich et al. introduce the term ROSI or Return on Security Investment [77]. ROSI is similar to the traditional accounting technique of ROI or Return on Investment but is focused on security spending. K. Hoo provides an analysis of decision making within information security including the topic of security spending decisions [78]. Hoo analyzes decision making through the lens of risk management. Adrian Mizzi describes an approach to security spending using accounting techniques that attempts to identify an upper bound on expenditures [79]. M. Al-Humaigani and D. Dunn propose a model that focuses on costs in which the maximum return that can be obtained from spending on information security is achieved when the total operational cost of security measures is minimized [80]. A number of other models exist. Rachel Rue et al. provide a useful review and assessment of various published models [81].

Security management within a very large organization is necessarily a practical exercise: budget decisions *must* be made and usually within a specified timeframe (the period of the year sometimes colloquially referred to as 'budget season'). Although the increasing number of academic papers being published on the topic of security spending is welcome, there are a number of hurdles that have not yet been crossed that would allow these various models to be used in a real-world setting. We will describe these challenges below.

The first issue is that the published work tends to focus on models that implement spending strategies but not on how to *select* the most appropriate spending model. A manager must select one model from amongst the many models that have been published. (The manager could potentially use a number of different models and compare their output, but even using that approach they would likely still have to select a subset of all available spending models.) This has not been a focus of published work to-date, with the exception of one paper by Lukas Demetz and Daniel Bachlechner [82].

The second issue concerns the fact that many of the models for implementing spending strategies assume the availability of data to populate the model. Examples include data describing the probability of an attack, the probability of an attack being successful, the cost of a successful attack, and so on. Actuarial data that describes these various parameters is not yet widely available [83]. Other parameters that are required for various models are perhaps less straightforward to obtain than it might initially appear. An example is in attempting to determine the cost of security measures. Within an enterprise-scale organization the initial purchase price represents only one part of the total cost of ownership, as noted in §3.1.

The third and perhaps most significant issue is that the various published models tend to use relatively complex mathematics. Unless the model is incorporated into some usable abstraction such as a spreadsheet it is unlikely that the average general

manager will be able to use it in a practical setting. Ph.D-level mathematics is not a typical skill within the ranks of general management. It might be possible to locate those skills within the total population of workers within a very large organization, but this poses its own budgeting challenges such as with cross-department billing.

Given the presence of these issues it appears that it will be some time before spending models from the academic literature will be widely employed in the professional world. There are however a number of approaches to the challenge of security spending that are in active use. One approach is to try and match the spending level within the organization to other organizations that operate in the same vertical market [84]. The rationale for this approach is the belief that if the organization matches common levels of spending then the organization will not be at a disadvantage in comparison to competitors. This approach does not work if the various organizations are all spending their funds on different things, but the question of on *what* to spend can potentially be resolved through further benchmarking or perhaps by consulting with analyst firms. This strategy of matching spending levels to your peers is perhaps only useful if the number of 'free riders' is low. In the situation where every organization attempts to free ride there is the possibility of mean reversion. The effect would be that every organization settles on the same spending level, but that spending level is inefficient for every organization! [85].

There are a number of other structural problems that affect security budgeting decisions and that we will note.

A key structural problem with spending decisions is in proving a negative – how can it be shown that a security breach did *not* occur because of security spending? When security measures fail this might indicate that there is insufficient spending on security within the organization. On the other hand, a breach might simply indicate that it has become easier to detect security failures. Perversely, spending on certain security technologies such as security monitoring makes it easier to detect security problems – which then increases the known number of failures! The effect of this uncertainty as to whether spending on security is creating value could conceivably drive a cycle of over-investing then under-investing in security. For example, a security breach creates the desire to spend; time passes after the breach and the memory of the event begins to recede; the amount of spending on security decreases as the organization becomes complacent, which results in weakened controls; another breach occurs and the cycle repeats.[2]

A second and perhaps more subtle set of structural problems emerge from the incentives that surround the security leader within the organization. As noted in §5.4, the security leader might employ the 'security mind set' in a manner that does not take into consideration important factors such as the cost of security measures. The extreme example here is spending a large amount of budget to eliminate a small

[2]Based on personal experience, the period with which this oscillation tends to occur is approximately three to five years if the budget process occurs annually.

risk. Some people are diverted by the niggling fact that a risk exists, no matter how small that risk might be [68].

It would be an atypical security leader that makes the claim that their organization is *overspending* on information security. We intuitively expect that the security leader will request an ever-increasing amount of budget funds; that they will hold the view that the amount spent on security should increase. We should confront this assumption. The key security challenges that are faced by enterprise-scale organizations and that are described in §3.1 are largely not addressed through pure security spending. Rather, they are challenges of managing people and technology at scale. As a result, it would be a wise security leader that makes the statement that tackling those fundamental problems is a better use of budget funds than purchasing yet another security product or service.

There are two apparent reasons why in the general case this does not seem to happen. First, it is probably more difficult for the security leader to 'sell' projects such as configuration management, entitlement management, and change management than to sell the need for a new security product. New security products and services have marketing to promote them. They might also claim to protect against a vulnerability or a threat that is currently receiving attention in the popular press. Large projects that attack fundamental problems are more likely to be more expensive, cut across functional silos, and be multi-year endeavors. Purchasing and deploying a new security product or service can provide the security leader with quick, visible achievements and therefore might be inherently more attractive. Along similar lines, the act of doing something can sometimes be perceived as having intrinsically more value than the decision to *stop* doing something. Consider the parallel example of the healthcare industry. The most profitable portion of the healthcare industry is the production of pharmaceutical products (drugs) for conditions such as diabetes and heart disease. There is considerably less money to be made by telling people to stop eating unhealthy food, even though that measure would have more overall benefit for substantially less cost. There is less glory to be had in trying to stop bad habits than in producing something that allows people to continue their bad habits.

There are an absence of models and techniques that would assist in enabling good security budgeting decisions. There are perverse incentives that push on the security leader within an organization. These factors would seem to have the effect of guiding security budgeting decisions in directions that might not be the most beneficial. Such problems are compounded in very large organizations because it is much easier to successfully deliver projects that are contained within one functional silo compared to projects that cut across multiple silos or across the entire organization. Implementing a pure IT security project such as deploying a new security monitoring product is much less challenging than changing a core business process that involves disparate groups within the organization such as human resources, software development teams,

the legal department, and so on. Because cross-silo projects are more difficult they are less likely to succeed. The institutional memory of those failures then makes the organization less likely to pursue those types of project in the future – even if they represent the best choice.

The changes that an organization would need to make to enable better security budgeting decisions are non-trivial. Those changes would require the consideration of questions such as the short-term versus long-term incentives and compensation model for managers, the balance between quantitative versus qualitative techniques for making budget decisions, matters of cross-silo collaboration, and questions of culture: how does the organization treat managers that attempt high-risk but high-reward projects and fail in the attempt?

We began §3 with an evaluation of a number of specific topics within the management of information security: the use of security management frameworks, the use of security policies, and the use of end-user security awareness training. In §5.4 we moved up one conceptual level and considered a number of questions at the level of the security organization. The budgeting discussion above has now moved us up to the final level which is the unit of the entire organization. As such, we have seen that there are challenges and opportunities within security management at the level of individual policy decisions all the way up to the macro level of the organization. The study of security management would seem to provide the opportunity to substantially improve the level of information security within organizations and therefore within the world at large. It is also an area that appears quite immature and yet paradoxically also strangely neglected. This is certainly true in comparison to the hubbub that surrounds the discovery of a new vulnerability or the release of the latest security technology. As a field we must carefully consider why this is the case and try to rationally apportion our e↵orts where they can provide the most value. In the next section we provide a number of short guidelines that summarize the analysis above and that can be used by enterprise-scale organizations in their security management e↵orts.

6 New directions

In this section we propose a number of aphorisms regarding the management of information security within very large organizations. These observations and recommendation are derived from the analysis provided in the previous sections.

i. As described in §3.1, an enterprise-scale organization faces a number of structural challenges. Many of these challenges have their root in the complexity that exists within enterprise environments and in the high rate of change within those environments. §3.2 described a number responses to those challenges, specifically *configuration management*, *change management*, and *entitlement management*.

These various activities have technical components but also managerial aspects and therefore should be considered as key activities within the portfolio of security management activities.

ii. In §5.4 it was observed that it is challenging for a strongly centralized security team to understand the needs of the business. This is the reason why some security teams within large organizations employ a coverage model in which Business Unit Information Security Officers (BUISOs) are embedded in business units to act as the eyes and ears of the security team. There are potential benefits to extending this approach to other functional areas. For example, a firewall engineer can probably be most effective when working within the part of the organization that has responsibility for the network infrastructure. A Unix security engineer is likely to be most effective working directly within the Unix engineering team within the organization. Skills such as firewall and Unix engineering that perhaps begin by being owned by the security team can be transitioned into the relevant functional area as they become more widely understood and commoditized. In this regard, one goal of security management is to *shrink* the security team!

iii. §5.1 discussed the use of security management frameworks such as BS 7799 / ISO/IEC 27001 and NIST Special Publication 800-100. It was identified that general frameworks such as these attempt to please all of the people all of the time. This is by-design on the part of the authors of such documents in order to broaden the base of potential users to the greatest possible extent. An organization that believes itself to be substantially different from the norm – such as with a very large organization – should be wary of the guidance provided within such security management framework documents. The further an organization is from the general case that is assumed by the authors of the framework document, the higher the probability that the value received from following the guidance in the document shrinks to zero or even passes through zero and becomes negative. An example of the latter case is where the expense of implementing a generic security control within the specialized environment is greater than the possible losses that would occur if no control existed. The focus should ultimately be on how to achieve the desired *outcome* and not in the first instance on the *mechanism* employed.

iv. We have seen in §5.2 and §5.3 that common security management practices that might be useful in the general case can encounter various problems when attempts are made to implement those practices within very large organizations. At the root of these problems is the fact that a small fraction of a very large number is itself a large number. A small failure rate at enterprise-scale can result in a large number of failures that must then be handled by compensating controls or supporting processes. Those secondary measures might be both expensive

and time-consuming to carry out, and therefore the value of the original practice becomes diminished in comparison to alternative approaches.

v. As identified in §5.4, within a large organization the security leader who has management responsibility for information security will require a strong general management skill set. This is necessary to effectively manage the hundreds of workers within the security organization, optimize the structure of that organization, manage the finances, and so on. Within a very large organization it is unlikely that the security leader will be part of the management committee and report to the CEO, even with the CSO title.

vi. §5.5 described an absence of models that could assist with making budgeting decisions and that are practical for use in a real-world setting. Alternative approaches such as attempting to set the level of spending to the level that exists within peer organizations are clumsy and risk inefficient spending. There are also a number of factors that can make it hard for the security leader to think rationally about security spending decisions. An example is the emotional effects of reacting to recent events such as a security breach. As noted in §3.1, the complicating factor of *scale* makes tasks such as identifying projects upon which to spend funds more difficult within enterprise-scale organizations. For example, it is much more challenging to perform projects that cut across multiple silos within a very large organization. Such projects might provide the opportunity to deliver the most value, but they can be forsaken for projects that are considered more tractable. Inside the information security field there is the well-known idea of 'defense-in-depth' where one security control buttresses another. Organizations should consider applying the idea of defense-in-depth to the process by which security spending decisions are made, with layers of decision-making to flush out the effects of perverse incentives and to enable better decisions at the unit of the organization as a whole.

7 Opportunities for future work

There are a number of opportunities for further work to be carried out.

It would be useful to further investigate the process and timeline by which security technologies transition from being operated by specialists to becoming more widely understood and able to be used in a production setting by generalists, even when those generalists do not have security expertise. There are a number of implications here for the management of information security technologies within organizations of all sizes.

The decision to use a security management framework is perhaps not a binary one, i.e. to either use a framework or not. An organization might determine that there are

certain aspects of a security management framework that provide value for the specific circumstances of their environment. Most security management frameworks such as BS 7799 / ISO 27001 and COBIT are flexible in terms of the specific controls that can be implemented but are not flexible about the core process within the framework (the ISMS process and the 'Plan-Do-Check-Act' cycle in the case of BS 7799 / ISO 27001, for example). It would be interesting to research the flexibility (or rigidity) of individual frameworks and determine to what extent they can bend to function in atypical environments before they lose their 'conceptual integrity' [86].

If security management practices that are common today such as the use of end-user security awareness training and security policies have problems at enterprise-scale because of their error rate, then this suggests the need to develop alternative approaches. Identifying techniques that have a smaller error rate is an obvious research direction to pursue. It would also be useful to investigate other types of solution such as systems that fail closed or that activate a fail-safe mode when an error occurs. Think of permissive action link systems such as those used in nuclear missile silos, for example. These systems require two separate individuals to perform certain actions simultaneously – such as each opening a separate safe to which only they have the combination. Such a system reduces the possibility that any single person will perform a damaging action.

An even more fundamental way to solve the problems that can occur when end-users make decisions that negatively affect security is to not put end-users in situations where they can make those decisions! To use an intentionally dramatic example: when we enter an elevator we do not worry that pushing the wrong button will send us plummeting to our deaths. Technology should ideally only expose those aspects that allow the user to operate it and not enable the user to compromise the security of the entire organization by making a bad decision.

We have noted the need for models that can help managers make better budgeting decisions and that are practical. This latter aspect is as important as the former. The 37% rule provided by Gordon & Loeb is a good example of a heuristic that could be easily applied by a manager. In addition to a lack of usable models there is also an absence of guidance concerning how to differentiate between various models and therefore select a spending model that is most appropriate for the circumstances of a particular organization. The opportunity exists for work to be carried out on this topic.

A very large organization is one particular context in which the management of information security takes place. It would be valuable to identify and categorize the various other major use-cases of security management techniques and the extent to which they are served by published research. Small businesses seem to receive relatively good amounts of research, such as the papers by Debasis Bhattacharya [87] and Ladislav Beranek [88]. Small businesses are probably attractive to researchers in

comparison to larger organizations because it is presumably easier to obtain access to managers and other decision makers within a small business. The task of obtaining relevant data from within a smaller organization must also be easier because the types of challenges described in §3.1 are less present in smaller organizations.

When considering various use-cases for security management techniques there is the interesting question as to what extent the techniques that are shown to have value in one use-case are applicable to other use-cases. The analysis in this paper would seem to suggest that techniques that represent "best practice" in the general case become problematic when an attempt is made to use them within very large organizations. In other words, security management techniques that are effective within small businesses are probably not portable to very large organizations and vice-versa. Our focus in this paper has been on the dimension of size but there are multiple other dimensions in which organizations can differ and this would benefit from further study.

The most significant threat to validity for the findings presented in this paper is that the specific security management practices that were discussed were self-selected. There was *a priori* knowledge based on professional experience that these practices break down when an attempt is made to implement them at the scale of a very large organization. Subsequent research by separate authors should attempt to duplicate these findings and therefore validate them (or not). There is therefore an opportunity for other researchers to carry out that work.

We will note that in addition to opportunities for further academic work there are also potential commercial opportunities. Given the results of the gap analysis performed in this paper, organizations might conceivably pay a consulting firm for guidance regarding these various aspects of their security management activities. Reviewing the description of consulting services provided by major security companies, major technology consulting firms, and major management consulting firms, there appears to be minimal provision of consulting services that would assist an organization with these types of challenges [89] [90] [91] [92] [93].

8 Conclusion

In this paper we have examined the topic of the management of information security from a new perspective, namely the point of view of very large organizations.

The fiduciary duty of management is to maximize shareholder value. (In the case of non-profit organizations and governments the goal of management is to be the best possible steward of the resources of the organization. We can consider this to be essentially the same goal for the purposes of our argument.)

To that end, each professional field has "best practices" that represent the conventional wisdom as to how to achieve this goal. A best practice is intended to be an

activity that adds value in the general case. We have seen in our analysis within this paper that some best practices can experience what economists refer to as diminishing marginal utility. As the target organization drifts from the typical use-case the amount of value-added declines and can potentially enter negative territory.

We have also examined the degree of innovation in the practice of security management and the extent to which the literature can support practical, real-world activities. In both areas we have identified a number of opportunities to perform further work.

As a field we should be wary of the false sense of comfort that comes from doing things the way they have always been done. Organizations are dynamic, living entities and like humans they can become entrenched in particular ways of working. A preference for the status quo can lead to stagnation and worse. A living thing can only survive if it evolves, if only because it is surrounded by an environment that is itself evolving. If there are inefficiencies they should be identified and rooted out.

It is in this spirit of self-improvement that this paper is offered. Our hope is that this contribution can help to make the management of information security more successful.

9 Acknowledgements

I would like to thank my project advisor John Austen. I would also like to thank the module leader, tutors, and students from the winter 2010 IC01 Security Management module from the M.Sc in Information Security course at Royal Holloway, University of London.

10 References

[1] European Union, “Small and medium-sized enterprises (SMEs); What is an SME?”, available: *http:// ec.europa.eu/ enterprise/ policies/ sme/ facts-figures-analysis/ sme-definition/ index en.htm*, 2003.

[2] U.S. Small Business Administration, “Statistics of U.S. Businesses, Business Dynamics Statistics, Business Employment Dynamics, and Nonemployer Statistics”, available: *http:// archive.sba.gov/ advo/ research/ data.html*, 2012.

[3] Inland Revenue Service, “Instructions for Form 1120-W”, available: *http://www.irs.gov/pub/irs-pdf/i1120w.pdf*, 2012.

[4] U.S. Census Bureau, “Statistics about Business Size (including Small Business)”, available: *http:// www.census.gov/ econ/ smallbus.html*, 2012.

[5] Christopher Avery, Judith Chevalier, and Richard Zeckhauser, “The CAPS

Prediction System and Stock Market Returns", John F. Kennedy School of Government, Harvard University, Faculty Research Working Papers Series (RWP09-011), available: *http://www.hks.harvard.edu/fs/rzeckhau/CAPS.pdf*, April 2009.

[6] Laree Kiely and Terry V. Benzel, "Systemic security management", *IEEE Security & Privacy*, pp. 74-77, Nov-Dec, 2006.

[7] Microsoft, "MSDN Library, Considerations for Enterprise-Scale Applications", available: *http://msdn.microsoft.com/en-us/library/↵647734.aspx*, 2012.

[8] IBM, "Sametime 7.5.1 – Best Practices for Enterprise Scale Deployment", available: *http://www.redbooks.ibm.com/abstracts/sg247410.html*, 2012.

[9] Avishai Wool, "A quantitative study of firewall configuration errors", *IEEE Computer*, Volume 37, Issue 6, pp. 62-67, available: *http://waterfallsecurity.com/wp-content/uploads/2009/11/computer2004_Firewall MisconfigurationJune04.pdf*, June 2004.

[10] MyungKeun Yoon, Shigang Chen, and Zhan Zhang, "Reducing the Size of Rule Set in a Firewall", *IEEE International Conference on Communications 2007* (ICC), pp. 1274-1279, 2007.

[11] MyungKeun Yoon, Shigang Chen, and Zhan Zhang, "Minimizing the Maximum Firewall Rule Set in a Network with Multiple Firewalls", *IEEE Transactions on Computers*, Volume 59, Issue 2, pp. 218-230, February 2010.

[12] Daniel Geer, Rebecca Bace, Peter Gutmann, Perry Metzger, Charles P. Pfleeger, John S. Quarterman, and Bruce Schneier, "CyberInsecurity: The Cost of Monopoly; How the Dominance of Microsoft's Products Poses a Risk to Security", available: *http://cryptome.org/cyberinsecurity.htm*, 2003.

[13] Swiss Federal Data Protection and Information Commissioner (FDPIC), "Swiss Federal Data Protection Act (DPA)", available: *http://www.edoeb.admin.ch/org/00828/index.html?lang=en*, 1993.

[14] United Kingdom, "Data Protection Act", available: *http://www.legislation.gov.uk/ukpga/1998/29*, 1998.

[15] International Organization for Standardization / International Electrotechnical Commission, "ISO/IEC 11179 Metadata Registry (MDR) standard", available: *http://metadata-standards.org/11179/*, 2004.

[16] Phil Venables, "Directories, inventories and the power of triangulation", *Information Security Now*, a publication of The British Computer Society,

Autumn 2006.

[17] Alec Mu↵et, “WAN-hacking with AutoHack – Auditing security behind the firewall”, *5th USENIX Unix Security Symposium*, available: *http://131.106.3.253/publications/library/proceedings/security95/fullpapers/mu↵ett.pdf* , 6th June, 1995.

[18] Steve Beattie, Seth Arnold, Crispin Cowan, Perry Wagle, Chris Wright, and Adam Shostack, “Timing the Application of Security Patches for Optimal Uptime”, *USENIX 16th Systems Administration Conference* (LISA 2002), Philadelphia, PA, available: *http://www.homeport.org/~adam/time-to-patch-usenix-lisa02.pdf*, December 2002.

[19] Gery Menegaz, “Enterprise Entitlements Management: Moving beyond authentication”, *ZDNet*, available: *http://www.zdnet.com/blog/btl/enterprise-entitlements-management-moving-beyond-authentication/79516* , June 7, 2012.

[20] Microsoft Academic Search, available: *http://academic.research.microsoft.com/*, 2013.

[21] Elsevier, available: *http://www.journals.elsevier.com/computers-and-security/*, 2012.

[22] Rolf Moulton, “Data security is a management responsibility”, *Computers & Security*, Volume 3, Issue 1, pp. 3-7, February 1984.

[23] James A. Schweitzer, “How security fits in – a management view”, *Computers & Security*, Volume 6, Issue 2, pp. 129-132, April 1987.

[24] William A. J. Bound, “Discussing security with top management” , *Computers & Security*, Volume 7, Issue 2, pp. 129-130, April 1988.

[25] Larry Stevens, “Security systems: getting management to shell out”, *Computers & Security*, Volume 7, Issue 2, pp. 220-221, April 1988.

[26] Martin Plant, “Getting management buy-in to IT Security”, *Computers & Security*, Volume 12, Issue 7, pp. 623-626, November 1993.

[27] Richard Baskerville, “New organizational forms for information security management”, *Computers & Security*, Volume 16, Issue 3, 1997.

[28] Thomas Finne, “A conceptual framework for information security management”, *Computers & Security*, Volume 17, Issue 4, pp. 303-307, 1998.

[29] M.M. Elo↵a and S.H. von Solms, “Information security management: a hierarchical framework for various approaches”, *Computers & Security*, Volume 19, Issue 3, pp. 243-256, March 2000.

[30] Denis Trèek, “An integral framework for information systems security management”, *Computers & Security*, Volume 22, Issue 4, pp. 337-360, May 2003

[31] Steve A. Purser, “Improving the ROI of the security management process”, *Computers & Security*, Volume 23, Issue 7, pp. 542-546, October 2004.

[32] Albin Zuccato, “Holistic security management framework applied in electronic commerce”, *Computers & Security*, Volume 26, Issue 3, pp. 256-265, May 2007.

[33] Information Systems Audit and Control Association, “COBIT (Control Objectives for Information and Related Technologies)”, available: *http:// www.isaca.org/* , 2012.

[34] British Standards Institution, “Information Security Management Systems – Specification with guidance for use”, (BS 7799-2:2002), 2002.

[35] Basie Von Solms, “Information Security governance: COBIT or ISO 17799 or both?”, *Computers & Security*, Volume 24, Issue 2, pp. 99-104, March 2005.

[36] Boehmer, W., “Appraisal of the E↵ectiveness and Efficiency of an Information Security Management System Based on ISO 27001”, *Second International Conference on Emerging Security Information, Systems and Technologies* (SECURWARE), Cap Esterel, France, August 25-31, 2008.

[37] Gene Kim, “Prioritizing Processes and Controls for E↵ective and Measurable Security”, CERIAS Security Seminar, available: *http://www.youtube.com/watch?v=38eQMQLJwn8*, September 20, 2006.

[38] National Institute for Science and Technology, “NIST Special Publication 800-100 (Information Security Handbook: A Guide for Managers)”, available: *http: //csrc.nist.gov/publications/nistpubs/800-100/SP800-100-Mar07-2007.pdf* , 2007.

[39] Mike Usher, “Complying with BS7799” presentation from M.Sc course in Information Security, Royal Holloway, University of London, 2011.

[40] Mikko T. Siponen, “A conceptual foundation for organizational information security awareness”, *Information Management & Computer Security*, Volume

8, Issue 1, pp. 31-41, 2000.

[41] H.A. Kruger and W.D. Kearney, “A prototype for assessing information security awareness”, *Computers & Security*, Volume 25, Issue 4, pp. 289-296, June 2006.

[42] Eric Savitz, “Humans: The Weakest Link In Information Security”, *Forbes*, available: *http:// www.forbes.com/ sites/ ciocentral/ 2011/ 11/ 03/ humans-the-weakest-link-in-information-security/* , 2011.

[43] Douglas Schweitzer, “How to toughen the weakest link in the security chain”, *Computerworld*, available: *http://www.computerworld.com/s/article/77360/*, January 8, 2003.

[44] U.S. Office of Management and Budget, “Appendix III to OMB Circular No. A-130 Security of Federal Automated Information Resources”, available: *http://www.whitehouse.gov/omb/circulars a130 a130appendix_iii*, 1996.

[45] Gary McGraw, “Data supports need for security awareness training despite naysayers”, available: *http:// searchsecurity.techtarget.com/ news/ 2240162630/ Data-supports-need-for-awareness-training-despite-naysayers*, 2012.

[46] Ponnurangam Kumaraguru, Steve Sheng, Alessandro Acquisti, Lorrie Faith Cranor, and Jason Hong, “Teaching Johnny Not to Fall for Phish”, *ACM Transactions on Internet Technology* (TOIT), Volume 10, Issue 2, available: *http:// www.heinz.cmu.edu/ ~acquisti/ papers/ johnny paper.pdf* , May 2010.

[47] Bob Blakley, Ellen McDermott, and Dan Geer, “Information security is information risk management”, Proceedings of the *2001 Workshop on New security Paradigms* (NSPW), pp. 97-104, Cloudcroft, New Mexico, USA, September 2001.

[48] Lawrence D. Bodin, Lawrence A. Gordon, and Martin P. Loeb, “Information security and risk management”, *Communications of the ACM*, Volume 51, Issue 4, pp. 64-68, April 2008.

[49] Rachna Dhamija, J.D. Tygar, and Marti Hearst, “Why Phishing Works”, proceedings of the *2006 Conference on Human Factors in Computing Systems*, Montreal, Quebec, Canada, available: *http: //www.cs.berkeley.edu/ ~tygar/ papers/ Phishing/ why phishing works.pdf*, April 22-27, 2006.

[50] Cisco, “Email Attacks: This Time Its Personal”, Cisco Whitepaper, available: *http:// www.cisco.com/ en/ US/ prod/ collateral/ vpndevc/ ps10128/ ps10339/*

ps10354/targeted_attacks.pdf, 2011.

[51] Michael Whitman and Herbert Mattord, *Principles of Information Security*, Course Technology, 2011.

[52] Kerry-Lynn Thomson and Rossouw von Solms, "Information security obedience: a definition", *Computers & Security*, Volume 24, Issue 1, pp. 69-75, 2005.

[53] Basie von Solms, "Information Security – The Fourth Wave", *Computers & Security*, Volume 25, Issue 3, pp. 165-168, May 2006.

[54], SANS Institute, "Information Security Policy Templates", available: *http://www.sans.org/security-resources/policies/*, 2012.

[55] Joan Goodchild, "Security Tools, Templates, Policies", *CSO Magazine*, available: *http://www.csoonline.com/article/486324/security-tools-templates-policies*, 2013.

[56] U.S. Securities & Exchange Commission, "Sarbanes Oxley Act of 2002", available: *http://www.sec.gov/about/laws/soa2002.pdf*, 2002.

[57] GPO, "Health Insurance Portability and Accountability Act of 1996", available: *http://www.gpo.gov/fdsys/search/pagedetails.action?granuleId= CRPT-104hrpt736&packageId=CRPT-104hrpt736*, 1996.

[58] Payment Card Industry Security Standards Council, *https://www.pcisecuritystandards.org/*, 2012.

[59] John Kirkwood, "Who should the CISO report to?", *CSO Magazine*, available: *http://www.csoonline.com/article/702330/who-should-the-ciso-report-to-?*, March 16, 2012.

[60] Ruth Walker, "Life at C-level: too many chiefs?", *Christian Science Monitor*, available: *http:// www.csmonitor.com/ The-Culture/ Verbal-Energy/ 2012/ 1109/ Life-at-C-level-too-many-chiefs*, November 9, 2012.

[61] Edward Ferrara, "Nine Managed Security Service Providers (MSSPS) Compete in the North American Market", *Forrester Research*, available: *http://blogs.forrester.com/edward_ferrara/12-04-04-nine_managed_security_services_providers_mssps_compete_in_the_north_american_market*, April 4, 2012.

[62] Online listing of IBM managed security services, available: *http://www-935.ibm.com/services/us/en/it-services/managed-security-services.html*, 2013.

[63] Online listing of Dell managed security services, available: *http://www.secureworks.com/enterprise/managed_security/*, 2013.

[64] Online listing of Symantec managed security services, available: *http:// www.symantec.com/ managed-security-services*, 2013.

[65] Christos Ioannidis, David Pym, and Julian Williams, "Fixed Costs, Investment Rigidities, and Risk Aversion in Information Security: A Utility-theoretic Approach", *2011 Workshop on the Economics of Information Security* (WEIS), available: *http: // homepages.abdn.ac.uk/ d.j.pym/ pages/ IoannidisPymWilliams-Utility.pdf*, 2011.

[66] Daintry Du↵y, "The FUD Factor: Fear, uncertainty and doubt (FUD) may help scare your company into short-term compliance, but CSOs say that's a shortsighted strategy", *CSO Magazine*, available: *http:// www.csoonline.com/ article/ 217983/ the-fud-factor*, April 2003.

[67] Bruce Schneier, *Schneier on Security*, Wiley, September 2008.

[68] Barry Glassner, "The Culture of Fear: Why Americans Are Afraid of the Wrong Things", Basic Books, March 2000.

[69] Andrew Stewart, "On risk: perception and direction", *Computers & Security*, Volume 23, Number 5, pp. 362-370, available: *http://www.andrewinfosec.com/AndrewStewart-02.pdf*, July 2004.

[70] Ross Anderson, "Why information security is hard – an economic perspective", proceedings of the *17th IEEE Annual Computer Security Applications Conference*, pp. 358-265, available: *http://www.cl.cam.ac.uk/⇠rja14/Papers/econ.pdf*, December 10-14, 2001.

[71] Ross Anderson, Rainer Boehme, Richard Clayton, and Tyler Moore, "Security Economics and the Internal Market", European Network and Information Security Agency, available: *http://www.enisa.europa.eu/publications/archive/economics-sec*, January 31, 2008.

[72] Ross Anderson and Tyler Moore, "The Economics of Information Security: A Survey and Open Questions", available: *http://www.cl.cam.ac.uk/⇠rja14/Papers/toulouse-summary.pdf*, 2006.

[73] *Workshop on the Economics of Information Security*, available: *http://econinfosec.org*, 2013.

[74] Lawrence A. Gordon and Martin P. Loeb, “The economics of information security investment”. *ACM Transactions on Information and System Security* (TISSEC), Volume 5, Issue 4, pp. 438-457, available: *http:// ns1.geoip.clamav.net/ ~mfelegyhazi/ courses/ BMEVIHIAV15/ readings/ 04 _GordonL02economics _security investment.pdf* , November 2002.

[75] Kanta Matsuura, “Productivity Space of Information Security in an Extension of the Gordon-Loeb’s Investment Model”, *2008 Workshop on the Economics of Information Security* (WEIS), available: *http://weis2008.econinfosec.org/papers/Matsuura.pdf*, 2008.

[76] Jan Willemson, “On the Gordon & Loeb Model for Information Security Investment”, *2006 Workshop on the Economics of Information Security* (WEIS), available: *http:// weis2006.econinfosec.org/ docs/ 12.pdf* , 2006.

[77] W. Sonnenreich, J. Albanese, and B. Stout, “Return On Security Investment (ROSI) – A Practical Quantitative Model”, *Journal of Research and Practice in Information Technology*, volume 31, issue 1, pp. 239-252, available: *http:// www.ra.cs.uni-tuebingen.de/ lehre/ uebungen/ ss09/ introsec/ ROSI-Practical Model.pdf*, 2005

[78] K. J. S. Hoo, “How Much Is Enough? A Risk Management Approach to Computer Security”, PhD thesis, Stanford University, 2000.

[79] Adrian Mizzi, “Return on information security investment – the viability of an anti-spam solution in a wireless environment”, *International Journal of Network Security*, Volume 10, Issue 1, pp. 18-24, *http:// ijns.femto.com.tw/ download _paper.jsp? PaperID= IJNS-2006-07-16-1&PaperName=ijns-v10-n1/ ijns-2010-v10-n1-p18-24.pdf* , January 2010.

[80] Al-Humaigani M. and Dunn D. B. “A model of return on investment for information systems security”, proceedings of the *46th IEEE International Midwest Symposium on Circuits & Systems*, Volume 1, pp. 483-485, 2003.

[81] Rachel Rue, Shari Lawrence Pfleeger and David Ortiz, “A Framework for Classifying and Comparing Models of Cyber Security Investment to Support Policy and Decision-Making”, *2007 Workshop on the Economics of Information Security* (WEIS), available: *http://weis2007.econinfosec.org/papers/76.pdf*, 2007.

[82] Lukas Demetz and Daniel Bachlechner, “To invest or not to invest? Assessing the economic viability of a policy and security configuration management tool”, *2012 Workshop on the Economics of Information Security* (WEIS),

available: *http:// weis2012.econinfosec.org/ papers/ Demetz _WEIS2012.pdf* , 2012.

[83] Adam Shostack and Andrew Stewart, *The New School of Information Security*, Addison-Wesley Professional, 238 pp., March 2008.

[84] Jeremy Kirk, "How much should you spend on IT security?", *InfoWorld*, available: *http://www.infoworld.com/d/security-central/ how-much-should-you-spend-it-security-306* , September 2010.

[85], Andrew Stewart, "Can spending on information security be justified? Evaluating the security spending decision from the perspective of a rational actor", *Information Management & Computer Security*, Volume 20, Issue 4, pp. 312-326, available: *http://www.andrewinfosec.com/AndrewStewart-06.pdf* , July 2012.

[86] Fred Brooks, "The Mythical Man-Month: Essays on Software Engineering", Addison-Wesley, 1975.

[87] Debasis Bhattacharya, "Leadership styles and information security in small businesses", *Information Management & Computer Security*, Volume 19, Issue 5, pp. 300-312, 2011.

[88] Ladislav Beranek, "Risk analysis methodology used by several small and medium enterprises in the Czech Republic", *Information Management & Computer Security*, Volume 19, Issue 1, pp. 42-52, 2011.

[89] IBM web site, available: *http://www-935.ibm.com/services/us/en/ it-services/ consulting-security-services.html* , 2013.

[90] Symantec web site, available: *http:// www.symantec.com/ it-consulting-services* , 2013.

[91] McAfee web site, available: *http://www.mcafee.com/us/services.aspx*, 2013.

[92] Accenture web site, available: *http:// www.accenture.com/ us-en/ Pages/ service-consulting-risk-management-overview-summary.aspx* , 2013.

[93] Booz Allen Hamilton home page, available: *http://www.boozallen.com/consulting*, 2013.