A Longitudinal Study of Information System Threat Categories: The Enduring Problem of Human Error

Ghi Paul Im
Georgia State University

Richard L. Baskerville
Georgia State University

Acknowledgement

The authors gratefully acknowledge the assistance of Edward Chang in data collection and coding of the original 1993 study. His excellent records, coding numbers and notes greatly facilitated the further analysis of human error incidents of this older study.

Abstract

Taxonomies of information security threats usually distinguish between accidental and intentional sources of system risk. Security reports have paid a great deal of attention in recent years to the growing problem of hacking and intentional abuse. The prevalence of these reports suggests that hacking has become a more severe problem in relation to other security threats, such as human error. In this paper, we report on research that addresses this question: "How have changes over time in the frequency of hacking and other intentional forms of security threats affected the validity of information systems risk management taxonomies?" We replicate a simple study of the proportions of categories of security threats that was originally completed in 1993. Comparing the results from the replicated study with the results from the original study, we find that the proportions of threat categories have, in contradiction with the popular perception, remained relatively stable over the past decade. These results indicate that human error remains a significant and poorly recognized issue for information systems security. We propose and validate an elaborated taxonomy of information security threats that provides additional insight into human error as a significant source of security risk.

ACM Categories: C.4, H.1.2, H.2.0, K.4.1, K.6.5

Keywords: Computer Security, Information Infrastructure Protection, Information Security, Information System Threat, Human Error, Information System Threat Taxonomy, Software Defects, Software Quality and Reliability

Introduction

As businesses, governments, and educational organizations rely more on computers and the internet, our valuable computing resources are exposed to greater security threats and thus protecting our properties is becoming more of a prime concern to individuals and businesses. According to the Internet domain survey (Internet Systems Consortium, 2004) (www.isc.org/ds/), the Internet grew from 171 million computers in January 2003 to more than 233 million in January 2004. This trend demonstrates that computers and the Internet are becoming part and parcel of our day-to-day operations and business activities. Although this trend promises many benefits, it also poses many security risks.

Protection against threats is often described as computer security. In defining computer security, some authorities (Amoroso, 1994; Howard, 1997) narrowly focus on the activities that lead directly to system vulnerabilities. Howard's (1997) definition includes the *means* used to gain unauthorized access (such as a virus) and the *ends* of the attacks (such as corruption and fraud), but excludes indirect incidents such as computer equipment theft, environmental threats and unintentional events. However, others (Loch et al., 1992; Neumann, 1995) include in their accounting of security incidents, those that can directly or indirectly lead to system vulnerabilities.

In this research, we consider security threats broadly to be security incidents that can directly or indirectly lead to system vulnerabilities. Thus information security threats include accidental (indirect) ones such as natural disasters and human errors (Baskerville, 1996; Cohen, 1997; Loch et al., 1992; Neumann, 1995). Natural disasters and human errors create vulnerabilities that can be exploited, leading to security problems. For example, errors in system designs that provoke erroneous entries by users will also provide vulnerabilities that can be intentionally exploited by attackers (Norman, 1983). Power failure due to natural disaster disrupts normal operations of the system and requires emergency response to prevent security threats. The breadth of this definition also encompasses the notion that the security safeguards that address accidental threats will often lend help in the prevention of intentional attacks.

The profiles of current security threats are reported in several industry surveys. For example, of 530 organizations surveyed for the 2003 CSI/FBI survey (Richardson, 2003), 56% reported unauthorized use, compared to 60 percent for the 2002 survey. Of those, 38% reported from 1 to 5 incidents, and 16% reported 11 or more incidents. The top three types of attack or misuse detected in the last year include virus (82% of all the respondents were affected), insider abuse of net access (80%), and laptop theft (59%). The 2003 Australian Computer Crime and Security Survey (AusCERT, 2003) similarly reports that virus/worm/Trojan (80% of respondents were affected), insider abuse of Internet access (62%), and laptop theft (53%) are the major three security incidents. In addition to these well analyzed incidents, the emergence of cyber-terrorists poses new security risks. Terrorist groups are now developing Internet sites and using Internet technologies (Warren & Hutchinson, 2001) for propaganda, publicity, fundraising, and information dissemination, as well as attacking government and civilian computers (Berkowitz, 2003).

These reports suggest that intentional security threats such as hacking, computer viruses, and computer theft are becoming a more severe problem in relation to other security vulnerabilities. Accidental and indirect incidents, such as erroneous entries by users or power failures, are neglected vulnerabilities in these surveys. Such neglect leads to a gap in our understanding of security incidents and their corresponding threats because only some types of threats are covered. The use of these surveys as foundations for scholarly works creates a corresponding gap in the research literature, a gap driven by the inferential limits to analyses based on partial and unrepresentative data (Shimeall & Williams, 2002). Our work seeks to address this gap by providing a more complete taxonomy for analysis of security incidents.

A complete taxonomy is fundamental to the completeness in the threat inventory process. The taxonomy keeps the threat inventory complete and representative. These threat inventories are also an important component in the risk analysis stage of most security design methodologies, such as CRAMM (CCTA Risk Analysis and Management Methodology) or BDSS (Bayesian Decision Support System).

Taxonomy refers to the theory and practice of classification, arising in a branch of science known as systematics (Mckelvey, 1982). Unlike nomological science with its focus on uniformity, the taxonomies of systematics focus on diversity. Classification separates phenomena into different kinds. In this research, we focus on the durability and validity of threat taxonomies such as Loch et al. (1992), Neumann (1995) or Baskerville (1996) in light of the shifting proportions of intentional threats to accidental threats. We ask, "How have changes over time in the frequency of hacking and other intentional forms of security threats affected the validity of information systems risk management taxonomies?"

By adopting a broad view of security threats in our risk management taxonomies, human error becomes a consideration that is often overlooked by popular crime surveys. Human errors are an important issue in information systems. David Parnas recognized the severity of human error in information systems in general (and security threat in particular) as early as the 1980s. He argued that building reliable software systems is problematic because of the massive number of different states in software systems. Despite his recognition that the next 20 years of research would not change that fact (Parnas, 1985), human error factors have received sporadic attention from IS (security) researchers. Prior literature consistently reports human errors as top-ranked security threats (Loch et al., 1992; Whitman, 2004). While researchers recognize the severity and importance of human errors, the body of work addressing this problem understanding remains limited, and the issues have not been thoroughly investigated. This is especially true in the area of threat taxonomies, where threats arising from human error have been largely neglected.

Experimental evidence shows that human errors are inevitable (Brown & Patterson, 2001). Supporting this evidence, two surveys have studied the causes of failures due to operators, hardware failures, software failures, and overload (Patterson et al., 2002). One collected failure data on the U.S. Public Switched Telephone Network (PSTN) and one from three Internet sites. Surprisingly, the surveys are consistent in their suggestion that operators are the leading cause of failure: 59 % for the PSTN and 51% for the Internet sites. The present research seeks to analyze such human errors in greater detail.

This research reviews examples of published threat taxonomies, focusing on one well-documented taxonomy from 1993. In order to determine if this taxonomy retains its validity (its durability), we replicate the study. We will examine the changing threat constellation and how this has affected the validity of this older taxonomy. We then provide a simple elaboration of the taxonomy and demonstrate the effectiveness of this elaborated taxonomy.

Threat Taxonomy

An information systems threat is the danger that an information system vulnerability can actually lead to undesirable consequences (Neumann, 1995). A threat requires a vulnerability or the undesirable consequences cannot arise. Security problems can take place at any stage of the information system life cycle from systems conceptualization and requirements definition to operation, evolution and decommission. For example, one must be careful to incorporate security concepts, even in the systems conceptualization stage, in cases where there is a strongly likelihood of dangerous risks (like an Internet-based system). In the analysis and design stage, one also can complete the analysis based on false or inappropriate assumptions for the computing environment and human behavior. Problems also can occur during systems operation and use. The major causes of security threats include natural environments, infrastructure (like electrical power), hardware malfunction, software misbehavior, communication media failures, and human errors (Neumann, 1995).

Given these potential problems and consequences, information systems security management is founded on security risk assessment, which in turn depends on threat identification. Comprehensive approaches to threat identification are typically supported with some form of threat classification system. There are different kinds of taxonomies that have been used to classify threats. These include: asset groupings that use characteristics of IS assets as the primary criterion for dividing the spectrum of threats into categories (Parker, 1981), impact groupings that use characteristics of losses when threats occur as the primary criterion for dividing the spectrum of threats into categories (Courtney, 1977), convenience groupings that do not use analytic criteria (Forcht, 1994), and multi-dimensional groupings that add dimensions and complexity to the models (Neumann, 1995). These different kinds of taxonomies are not equally effective with regard to the well-known taxonometric criteria of parallelism, mutual exclusivity, and completeness (Baskerville, 1996).

Baskerville (1996) developed the multi-dimensional threat taxonomy based on a comparative study of previous taxonomies. This taxonomy is shown in **Figure 1**. He classifies threats into two major classes: accidental and deliberate. Since accidental and deliberate threats are fundamentally different in their characteristics, they are further expanded using different classification schemes. Accidental threats are those not intentionally posed by humans. They are further classified as a simple two-class, one-dimensional taxonomy, either *catastrophes* or *errors*. Deliberate threats are caused by the people who interact with information systems and IS development. They are analyzed with a multiple category, two-dimensional taxonomy: *mode* (the person's basic approach to creating the threat) and *motive*. The modes of deliberate threats include *physical assault*, *falsification*, *malicious code*, and *cracking*. The motives of deliberate threats include *fraud*, *espionage*, and *vandalism*. Although other, deeper motives can be imagined, such as information warfare and cyber terrorism, we subsume the effects of these underlying motives into the more direct classes of threats such as vandalism (seeking to destroy systems, systems integrity or to deny systems services) and espionage (seeking to compromise the confidentiality of systems). For example, cyber terrorists may seek to vandalize systems and information warriors may engage in computer espionage, etc.

Baskerville sought to demonstrate the validity and usefulness of the taxonomy by using the model to analyze the vignettes published in the "Risks to the Public" column edited by Peter Neumann in *Software Engineering Notes* for a two-year period in 1992 and 1993 (January, 1992, 17:1 through October, 1993, 18:4) (Neuman, 1992 ~ 1993). This column lists computer-related risks reported through the related Internet newsgroup in that period. There were a total of 147 distinct incidents of threats that affected people and organizations and appeared to arise from the use of computer-based systems.

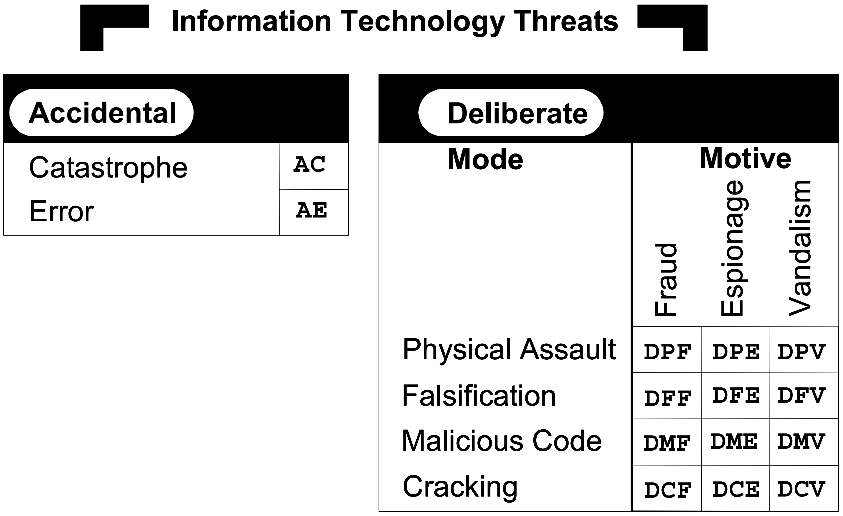

Information Technology Threats

Accidental	
Catastrophe	AC
Error	AE

Deliberate			
Mode	Motive		
	Fraud	Espionage	Vandalism
Physical Assault	DPF	DPE	DPV
Falsification	DFF	DFE	DFV
Malicious Code	DMF	DME	DMV
Cracking	DCF	DCE	DCV

Figure 1. The 1992-1993 Information System Threats Taxonomy (Baskerville, 1996)

There are limitations in the development and validation of this taxonomy. The crude action research validation process is not well justified or described. The published vignettes are clearly a convenience sample affected by current events and uncontrolled editorial influence. The coding process is not described, and appears to be rather subjective. However, the claims are equally modest. The application demonstrated that the model satisfied the taxonometric criteria of parallelism, mutual exclusivity, and completeness. The analysis of the vignettes also indicated four high priority threats for risk management: errors, fraud, cracking, and vandalism. The outcome of the original 1993 analysis is show in **Table 1**.

In this analysis of 1993 data, it would appear that the most important threat category was human error. Intentional attacks through malicious code (viruses and worms) have since risen into the millions annually. Systems cracking for the purpose of web site defacement has also become commonplace. Such vandalic events now occur every few minutes along every major firewall on the Internet (Pethia, 2003). Perhaps fraud has increased, such as identity theft or spoofing. If this study were repeated today, would the taxonomy still remain valid, and would the recommended proportional attention of managers shift dramatically?

Accidental	Catastrophe	Error	Total
Total (%)	17 (19%)	72 (81%)	**89**

Deliberate	Fraud	Espionage	Vandal	Total (%)
Physical	3	0	0	3 (5%)
False	14	5	0	19 (33%)
Malicious	3	1	8	12 (21%)
Cracking	9	2	13	24 (41%)
Total (%)	29 (50%)	8 (14%)	21 (36%)	**58**

Table 1. **1993 Study** (Total incidents = 147)

Accidental	Catastrophe	Error	Total
Total (%)	13 (17%)	65 (83%)	**78**

Deliberate	Fraud	Espionage	Vandal	Total (%)
Physical	5	2	1	8 (20%)
False	9	1	2	12 (29%)
Malicious	0	0	3	3 (7%)
Cracking	6	5	7	18 (44%)
Total (%)	20 (48%)	8 (20%)	13 (32%)	**41**

Table 2. Replicated Study (Total incidents = 119)

Revisiting Threat Taxonomy

The present study considers whether the proportion of hacking and other intentional forms of security threats increased relative to the accidental forms emphasized in 1996. Fortunately, Peter Neumann has continued to publish the vignettes in the "Risks to the Public" column in *Software Engineering Notes* (Neuman, 2001~2003). We can repeat Baskerville's original two-year analysis, for the more recent two-year (March, 2001, 26:2 through March, 2003, 28:2).[1]

We replicated the original study method in which the taxonomy was applied to classify the vignettes. Vignettes that were not directly related to computer-based information systems were excluded. We also excluded stories about unrealized threats, that is, vignettes that did not actually result in damage or harm, because the field of such potential threats is unbounded. One researcher coded each of the vignettes using a simple database tool. A second researcher then separately re-coded the same data. There was an inter-coder reliability rate of 0.95. The researchers conferred on the disputed items and resolved each disagreement. This process resulted in a classification of the vignettes that was at least as reliable as the original 1993 study. Simple descriptive statistics of the results are shown in **Table 2**.

Comparing **Tables 1** and **2** does show a smaller population of vignettes despite the slightly longer sample period. This shift may be because the vignettes in the replicated study contain many unrealized incidents from electronic voting systems, which may have eclipsed reports of realized incidents. However, the proportional relationships between the various classifications are remarkable for a number of reasons. Most importantly, the proportion of events (65 of 119) represented as human error is 55% of the population. The proportion of error-driven events has risen from 49% to 55%. Most surprising and unexpected are the proportions of events classified as malicious code (about 3% of the population) and cracking (about 15% of the population). Although firewall logs and Internet incident monitors number these incidents today in the millions annually, cracking is proportionally similar in this model and malicious code has fallen from 8% to 3% of the population.

One explanation for this surprising result arises from the nature of the underlying data. The risks column is not a simple log of security incidents. These are vignettes, stories often taken from the press, about computer threats; threats that have resulted in significant harm to an organization or a group of persons. As mentioned earlier, reports among the incidents that were about unrealized threats (imagined concerns) were excluded from our data. In order to validate the framework empirically (in reality) our data only included the realized threats.

Thus this population of incident reports, albeit framed in convenience, represents reports of *unmanaged* risks that actually harmed stakeholders. Unmanaged risks are those incidents for which the systems had been left vulnerable and unprepared. These risks arise because there are no commonplace safeguards that easily protect against the risks, or else the available safeguards were not in place. These unmanaged risks are the "interesting" incidents, events of public interest; the interest arises because these threats actually resulted in some form of unusual havoc. It is these interesting incidents that were selected from the published stories for use as vignettes in our study.

Managed risks are incidents for which there are safeguards to easily make systems protected and

[1] We extended the coverage by one issue in order to achieve a larger vignette population. *Software Engineering Notes* published several issues during this more recent period entirely dedicated to conference proceedings.

prepared, and these safeguards were in place. The millions of incidents of malicious code and cracking have become *managed* risks because of their sheer prevalence through the Internet and the widespread placement of safeguards (such as firewalls, automatic software patch distribution and virus software) to protect systems against such code. Managed risks are not interesting because these do not result in unusual havoc. For example, malicious code is mostly a managed risk these days. Mostly, but not entirely, as the data indicate, some 3% of unmanaged risks related to malicious code. These are the incidents of malicious code for which commonplace safeguards didn't work, or else the safeguards were not in place, and some interesting havoc resulted.

An Elaborated Model

The simple replication of the original 1993 study suggests that the model presented in that study is inadequate in at least two ways. First, the model purports to be useful and valid in modeling threats to information systems security for the purpose of risk management. This original claim is probably too simple. The model's usefulness and validity is narrower: it is useful and valid in modeling "unmanaged" threats, such as those for which there are no easily available tool-based safeguards.

Second, the proportion of unmanaged threat arising from human error appears to be increasing. While the "old" model does identify this class of problem as important (in terms of frequency), it does not provide any indications of what kinds of errors make up this class, and therefore little useful practical value as a management aid for preventing or compensating for the associated risks. We therefore seek to elaborate the 1993 model to explore sub-classification of human error and to similarly validate this elaborated model in the set of recent vignettes.

Human Error

There are different views of human error. On the one hand, human error can be viewed as a complex, socially constructed behavior. Argyris and his colleague (1986; 1978) point out that individuals create error and misunderstandings by unconsciously, but faithfully, following their theories-in-use, a form of skilled incompetence. Errors are among the human factors that socially influence information technology implementation, and require more sophisticated investigations into barriers to information technology implementation (Levine & Rossmoore, 1993; Schenk et al., 1998).

On the other hand, human error can be viewed in a rational, instrumental way: As one of many kinds of error. Levine and Rossmoore (1993) argue that scholars traditionally follow publicly espoused values for the conduct of science, and consequently model individual managers as independent, rational decision makers. Neumann (1995) defines errors simply as deviation from expected behavior (p. 12). However, this definition extends to computer errors, network errors, and other machine and electrical failures. The 1993 taxonomy regards errors as human errors, and adopts this rational view. Accordingly, our research will continue to adhere to this rational view of error.

The 1993 study failed in an attempt to elaborate errors because of ambiguity in the reports. It attempted to distinguish human errors, design errors and machine error, but failed because most errors described as computer error can indeed be traced to an operator or programmer error. Many operator and programmer errors can be further traced to problems with errors in the software designs. Consequently the 1993 study lumped errors into one category as a whole as distinct from catastrophes or acts-of-nature (that are clearly not human artifacts).

The failure of the 1993 study to achieve elaboration of the error model is probably due to a classification scheme that lacked the vaunted mutual exclusivity claimed for the rest of the model. For example, human error and design error are not mutually exclusive. We believe a better sub-classification could be developed to elaborate this critical element of the overall taxonomy.

Reason (1990) defines human error as a failure of some planned sequence of mental or physical activities to achieve its intended outcome (p. 9). He goes on to distinguish slips from mistakes. Slips or lapses are unintended actions, i.e. the actions do not go as planned. Slips are execution failures. Mistakes arise when the intended action proceeds as planned but fail to achieve their intended outcome. Mistakes are planning failures.

Reason builds an error taxonomy on Rasmussen's (1986) three levels of performance (skill-rule-knowledge) framework well known in systems reliability community. At the skill-based level (e.g. data input errors), human performance is governed by stored patterns of preprogrammed instructions represented as analogue structures in a time-space domain. At the rule-based level (e.g., design flaws), human performance relates to tackling familiar problems in which solutions are governed by stored rules of the type *if then*. At the knowledge-based level (e.g., integration problems), actions must be planned on-line for novel situations using conscious analytical processes and stored knowledge.

Performance Level	Error Type	Examples
Skill-based level	Skill-based slips	• Data input errors • Clerical errors • Pressing the wrong button
Rule-based level	Rule-based mistakes	• Stupid defaults • Incorrect or unreliable data • Simple design flaws • Misapplication of valid rules (e.g., day light saving time) • Application of invalid rules (e.g., truncation or rounding errors, distinguishing identical names, leap year errors)
Knowledge-based level	Knowledge-based mistakes	• System crashes • Software upgrades and crashes • Severe software malfunctions • Integration problems • Procedural flaws with the system

Table 3. Human Error Taxonomy

From this framework, Reason builds a generic human error modeling system that integrates the error mechanisms operating at each of the three levels of performance. Thus there are three basic error types: skill-based slips (and lapses), rule-based mistakes, and knowledge-based mistakes. Skill-based slips are attributable mainly to monitoring failures. These involve *inattention*, not monitoring critical nodes, and *overattention*, monitoring at an inappropriate moment during a routine action sequence. Rule-based mistakes arise in the misapplication of good rules (i.e., rules of proven worth) and application of bad rules. Knowledge-based mistakes are caused by bounded rationality and the fact that knowledge relevant to the problem space is nearly always incomplete and often inaccurate.

From this analysis, we propose to elaborate the 1993 security risk model to include the three types of errors discussed above. These are summarized in **Table 3.**

Applying the Elaborated Taxonomy

Figure 2 illustrates the elaborated taxonomy of information systems threats. Human error threats are detailed into skill-based, rule-based, and knowledge-based error.

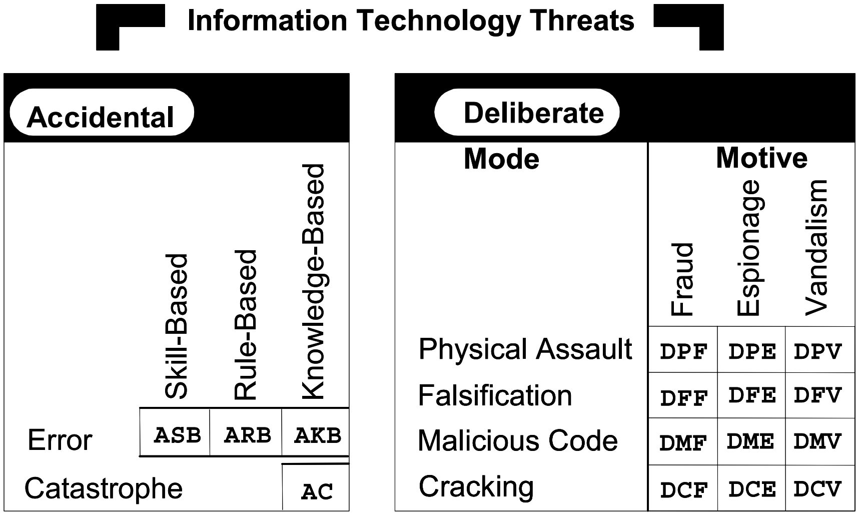

Figure 2. Elaborated Information Systems Threats Taxonomy

	Error	Skill-based	Rule-based	Knowledge-based	Total
1993 Study	Total (%)	4 (6%)	26 (36%)	42 (58%)	**72**
Present Study	Total (%)	10 (15%)	14 (22%)	41 (63%)	**65**

Table 4. Elaborated Study

In a manner similar to the original and the replicated studies, we sought to demonstrate the validity and usefulness of the elaborated taxonomy by using the human error model to further analyze the vignettes from the replicated study. One researcher coded each of the vignettes, another reviewed the coding results, and the researchers conferred on disputed items and resolved each disagreement. This process resulted in an elaborated classification of the human error in the vignettes. Simple descriptive statistics of the results are shown in **Table 4**.

Table 4 presents the number and percentage of incidents for each error types. The results indicate that of the 65 human error incidents, 63% involved knowledge-based error, 22% involved rule-based error, and 15% were skill-based error. The original 1993 study indicated that of the 72 human error incidents, 58% involved knowledge-based error, 36% involved rule-based error, and 6% were skill-based error. There is an interesting fall in proportion of rule-based errors and rise in the proportion of skill-based errors. No definitive rationale for this trend stands out in the data. Speculation is possible, for example, the trend might reflect the entry of a large number of inexperienced, but well-supervised new professionals in the field during the early part of the 2000's. While such a shift in the profession's demographics would explain higher proportions of naïve skill-based mistakes among newcomers while more experienced members of the profession were freed to reduce the number of rule-based errors, such conclusions are highly speculative and require further research.

Discussion

The analysis above suggests that the major source of unmanaged risks to information systems continues to be accidental in nature. Most of these accidents result from human errors. Most of these errors arise at the knowledge base error level. These various errors are not readily recognized as security lapses, yet these are the major wellspring of security failures.

This logic brings us to the boundary of systems and software engineering, and the disciplines concerned with software reliability, bug fixing and avoidance. Clearly, addressing these well-studied issues is beyond the boundary of a single paper. However, this issue also regards how information systems security management ought to regard the management of human error as a security risk.

The problem is not new. Parnas raised these concerns in information (security) systems in 1985, recognizing human error as one of the salient factors of security threats from the IS and the computer science community. After conducting a cross-sectional survey of 131 organizations, Loch et al. (1992) reported that "accidental entry of bad data by employees" and "accidental destruction of data by employees" were among top 3 out of a list of 12 security threats (p. 180). These proportions do not seem to have changed in a decade. Whitman (2004) describes that the top 3 security threats out of 12 lists include the items related to human error: "act of human error or failure" and "technical software failures or errors" that are attributable to human (p. 51).

Solving the problem is not impossible. There are two basic approaches for handling human errors as a security threat. The first one is **avoidance**, for example, by improving the system interface, providing better security policies and procedures, or providing better training (Brown & Patterson, 2001). As early as the 1970s, Salter and Schroeder (1975) suggested that the security system needs to be psychologically acceptable to the user by having *ease of use,* easy to use interfaces, making the user's mental image of her protection goals match the mechanisms available. They argued that this psychological acceptability by the user would minimize mistakes and errors. Despite this early recognition, this design principle has not received attention, whereas other, more functional principles (such as least privilege and fail-safe defaults) have become guiding dogma in the field. Traditionally, secure systems have been indifferent to the user's needs. For example, the typical authentication mechanism, the password, must not be too easy-to-use, or it can be guessed. Most secure systems concepts have strong roots in the military. Military users are often trained to follow the policies and rules no matter how arduous. This early tradition of secure systems and training processes in the military decreased the necessity of user-friendly secure systems (Salter & Schroeder, 1975).

Another avoidance mechanism that involves ease of use is *user-centered security* (Zurko & Simon, 1996). This mechanism synthesizes usability and security with reference to security models, mechanisms, systems, and software that have usability as a primary motivation or goal. They then suggested several potential approaches to achieving user-centered security: (1) applying usability testing (Nielsen, 1994) and usability techniques to the development of existing or new secure systems; (2) integrating security services that have software with a strong usability component, such as groupware; and (3) considering user needs as a primary design goal at the start of secure system development. These concepts are quite applicable for avoiding human error in using secure systems. When security applications are difficult or confusing to use, these will not be effectively used or may even be abandoned by the users. When the systems are user friendly, it is expected that the users are less likely to make skill-based slips or rule-based mistakes. User-centered security is a legitimate goal for secure systems.

Another way of avoiding human error is to providing better *security policies and procedures* (Dhillon, 2001a). A security policy is defined as the set of rules and practices that regulate how an organization manages, protects, and distributes its key asset, information (Walker, 1985; Woodward, 2000). Such security policies and procedures may help streamline complicated organizational functions in order to resolve ambiguities and misunderstandings within organizations (Dhillon, 2001a). Thus they would prevent the misinterpretation of data and misapplication of rules that otherwise would lead to security problems (Dhillon, 2001b). Pethia (2003) recommends that the threat of worms and viruses is best met when organizations avoid reactive solutions, and instead adopt proactive security policies and practices. The users following these proactive policies should be quite clear in their procedures when conducting tasks using a secured system. These clarifications and guidelines are likely to reduce skill-based slips and rule-based mistakes.

Finally, a mechanism for avoidance of human error is embodied in *training* programs. These programs most often focus on security awareness. Such training can lead to less security errors because individuals are equipped with better knowledge and are expected to be more attentive in solving problems. To tackle the threat of worms and viruses, system operators need to keep their skills and knowledge current, help educate the users of their systems. Also technology vendors need to dramatically reduce systems implementation errors (Pethia, 2003). Reducing implementation errors is crucial because most vulnerabilities in products come from these errors. Such vulnerabilities are latent in products and the errors are fixed only after they are discovered while in use. Identical flaws are often adopted into new versions of products without being fixed. Most of these vulnerabilities are caused by low-level design or implementation errors due to incomplete knowledge of software developers. Training facilitates the transition of systems analysts from novice to expert. Novices experience more difficulties in accurate problem definition and problem analysis than do experts, and proper training expands the limited knowledge base and enhances the problem-solving skills of such novices (Schenk et al., 1998). Empirical research also shows how managers trained in security planning techniques tend to exploit these in their planning processes (Straub & Welke, 1998).

Reason (1990) models how errors arise in settings of incomplete knowledge. There is interplay between two important aspects. The first aspect is the degree of expertise humanly available. That is to say, how much knowledge does the person bring to the problem setting? The second aspect is the degree of specificity in the cues presented to the person from the problem setting. That is to say, how much information is available to the person about the immediate problem setting? When the person brings insufficient knowledge into the problem setting, or cannot discover exactly how their knowledge applies to the problem, they engage in frequency-gambling solutions. Simply put, they gamble that the problem setting is like one they know to be frequently encountered, or that the knowledge they have frequently used will apply also to this setting (p. 147). Put differently, incomplete knowledge can lead to the following kinds of failure: (1) to recognize the existence of a problem, (2) to define the correct problem, (3) to use available information, (4) to recognize or question assumptions, (5) to consider a wide range of alternatives, and (6) to address implementation issues (Couger, 1995).

Security awareness training acts on both aspects of this error-making scenario. First, it does increase the stock of knowledge that trainees will take into security problem settings in the future. Second, it raises sensitivity (awareness) of the cues found in security settings that will help the trainee recognize that the security solutions from their knowledge base apply to the setting. From the perspective of knowledge-based human errors, security awareness training is indeed a very appropriate approach to managing these problems, both in theoretical and practical terms.

With regard to this approach, Norman’s (1988) four design principles for reducing error opportunities are notable: (1) promote good conceptual modeling, (2) make actions and their effects visible, (3) exploit natural mappings between intentions, actions, and effects, and (4) use feedback providing users with information about effects. He also suggests designing for error reduction should allow actions to be reversible and that systems should be consistent in structure to avoid memory problems (Norman, 1983). In addition to awareness training, such design principles should be promoted as important elements of organizational security programs. This activity represents a considerably expanded scope for most security programs.

While improving the system interface, providing better security policies and procedures, or providing better training may address the security risks to some extent, there is a second approach for reducing error: **tolerance** (Brown & Patterson, 2001). We can design systems so that they are fault-tolerant with respect to human errors, and thereby minimize the effects of human errors as security threats. Gray (1999) calls for Trouble-Free Systems, which can serve millions of people each day but require only minimal oversight by management. Hennessy (1999) mentioned that performance needs to share the spotlight with availability, maintainability, and scalability. Even Bill Gates has set "trustworthy computing" as the highest priority for his company, which means the computing must be available, reliable and secure as electricity, water services and telephony (Gates, 2002).

Although the idea of building tolerance systems to cope with human errors is new in information systems security community, we can borrow some ideas from recent recovery oriented computing (ROC) movement in computer science for possible application in security. ROC takes the viewpoint that "hardware faults, software bugs, and operator errors are facts to be coped with, not problems to be solved" (Patterson et al., 2002). ROC techniques include: Recovery experiments to test corrections, diagnostic aids, partitioning to contain incidents and enable rapid recovery, reversibility (undo) and safety margins, and redundancy to survive incidents and provide multiple lines-of-defense. ROC focuses on failure recovery time (MTTR). Applying this concept to security management would suggest a focus on minimal recovery time following a security breach. In cases where a large portion of system administration would be engaged in security incident management, an ROC orientation could reduce total cost of ownership.

Dhillon and Backhouse (2000) recognizes that computer security in itself is not a technical problem, but a social and organizational problem that involves people that operate the technical systems. He further suggests that the traditional information security principles such as confidentiality, integrity and availability need to be expanded to incorporate some additional ones such as responsibility, integrity of people, trustworthiness and ethicality. Security should not simply be viewed as means of protecting something concrete, but need to broaden its horizon by taking into account individuals and their social relationships (Dhillon & Backhouse, 2001). Any complete security solution ought to engage a security culture within the organization. A product of such a culture would be normative controls (Dhillon, 1999), organizational norms that promote integrity throughout an organization. Such controls include open communications and informal monitoring for behavioral changes and group conflicts.

Conclusion

This paper reviewed examples of published threat taxonomies, focusing on one well-documented taxonomy from 1993. We described a replication of this study intended to determine if the changing threat constellation has affected the validity of this older taxonomy. We discovered a need for a simple elaboration of the threat taxonomy to include a new focus on human error as a class of information systems threat. Our initial application demonstrated the potential effectiveness of this elaborated taxonomy.

Our study in threat taxonomy has been narrowly limited to the demonstration of the validity of one such elaborated threat taxonomy. We chose a well-documented taxonomy and repeated its application in order to reveal its explicatory value and its shortcomings in the present time. This shortcoming centered mainly on a lack of treatment for threats arising from human error. We succeeded in elaborating the taxonomy in this fashion, and demonstrating that it could be usefully applied in analyzing a real population of known security risks.

Another limitation of our research regards the convenience sample employed in applying the model. Our goal was to demonstrate the taxonometric validity by applying the taxonomies to actual threat populations to verify the functionality of the taxonomy in terms of completeness, mutual exclusivity and parallelism. The statistical value of the results in terms of threat proportions is modest at best. While the available data itself may lack statistical validity, it is certainly sufficient to satisfy our purposes in validating of the elaborated framework as an analytic tool.

Further research is needed to validate the statistical indications from the convenience sample, viz., that the proportions of threats in the selected population indicate human error is the most prevalent unmanaged threat. Further research is also needed to pragmatically validate the elaborated taxonomy in risk management methodologies and security design practice. Following this line, more research could investigate security safeguards that improve organizational performance in avoiding losses caused by human error. Also, rigorous statistical research is needed into more representative sampling frames of security threats.

Among the prevalent implications of this elaborated taxonomy is the serious need for research and practical knowledge about the management of human error in secure information systems. Security management needs to focus attention on new safeguards that protect systems from human error. Such safeguards could include specialized training, work aids, economic frameworks, and other motivational schemes.

References

Amoroso, E. G. (1994). *Fundamentals of Computer Security Technology*, Upper Saddle River, NJ: Prentice-Hall PTR.

Argyris, C. (1986). "Skilled Incompetence," *Harvard Business Review*, Vol.64, No.5, pp. 74-79.

Argyris, C. and Schon, D. A. (1978). *Organizational Learning: A Theory of Action Perspective*, Reading, MA: Addison-Wesley.

AusCERT (2003). Australian Computer Crime and Security Survey: Australian Federal Police.

Baskerville, R. (1996). "A Taxonomy for Analyzing Hazards to Information Systems," in Katsikas, S. and Gritzalis, D. (Eds.), *Information Systems Security: Facing the Information Society*, London: Chapman & Hall, pp. 167-176.

Berkowitz, B. (2003). *The New Face of War: How War Will Be Fought in the 21st Century*, New York: The Free Press.

Brown, A. B. and Patterson, D. A. (2001). "To Err is Human," First Workshop on Evaluating and Architecting System Dependability (EASY '01), Goteborg, Sweden.

Cohen, F. (1997). "Information System Attacks: A Preliminary Classification Scheme," *Computers & Security*, Vol.16, No.1, pp. 29-46.

Couger, J. D. (1995). *Creative Problem Solving and Opportunity Finding*, Danvers, MA: Boyd and Fraser.

Courtney, R. (1977). *Security risk assessment in electronic data processing*. AFIPS Conference NCC, Arlington, VA, pp. 97-104.

Dhillon, G. (1999). "Managing and controlling computer misuse," *Information Management & Computer Security*, Vol.7, No.4, pp. 171-175.

Dhillon, G. (2001a). "Challenges in Managing Information Security in the New Millennium," in Dhillon, G. (Ed.), *Information Security Management: Global Challenges in the New Millennium*, Hershey, PA: Idea Group Publishing, pp. 1-8.

Dhillon, G. (2001b). "Principles for Managing Information Security in the New Millennium," in Dhillon, G. (Ed.), *Information Security Management: Global Challenges in the New Millennium*, Hershey, PA: Idea Group Publishing, pp. 173-177.

Dhillon, G. and Backhouse, J. (2000). "Information system security management in the new millennium," *Communications of the ACM*, Vol.43, No.7, pp. 125-128.

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

Forcht, K. A. (1994). *Computer Security Management*, Danvers, MA: Boyd & Fraser.

Gates, W. (2002). Microsoft email on January 15, www.wired.com (accessed in May 2003)

Gray, J. (1999). "What Next? A dozen remaining IT problems," Turing Award Lecture.

Hennessy, J. (1999). "The Future of Systems Research," *Computer*, Vol.32, No.8, pp. 27-33.

Howard, J. D. (1997). *An Analysis of Security Incidents on The Internet 1989 - 1995* unpublished doctoral dissertation, Carnegie Mellon University, Pittsburgh, PA.

Internet Systems Consortium, Inc. (2004). ISC Internet Domain Survey, http://www.isc.org/ds/ (accessed August 2004)

Levine, H. G. and Rossmoore, D. (1993). "Diagnosing the human threats to information technology implementation: A missing factor in systems analysis illustrated in a case study," *Journal of Management Information Systems*, Vol.10, No.2, p. 55.

Loch, K. D., Carr, H. H., and Warkentin, M. E. (1992). "Threats to Information Systems: Today's Reality, Yesterday's Understanding," *MIS Quarterly*, Vol.16, No.2, p. 173.

Mckelvey, B. (1982). *Organizational Systematics: Taxonomy, Evolution, Classification*, Berkeley, CA.: University of California Press.

Neumann, P. (1992~1993). "Risks to the Public," *Software Engineering Notes*, Vol.17, No.1 - Vol.18, No.4.

Neumann, P. (2001~2003). "Risks to the Public," *Software Engineering Notes*, Vol.26, No.2 - Vol.28, No.2.

Neumann, P. G. (1995). *Computer-related Risks*, New York: ACM Press.

Nielsen, J. (1994). *Usability Engineering*, San Diego, CA: Academic Press.

Norman, D. (1983). "Design rules based on analysis of human error," *Communications of The ACM*, Vol.26, No.4, pp. 254-258.

Norman, D. (1988). *The Psychology of Everyday Things*, New York: Basic Books.

Parker, D. (1981). *Computer Security Management*, Reston, VA: Reston Publishing.

Parnas, D. L. (1985). “Software Aspects of Strategic Defense Systems,” *Communications of the ACM*, Vol.28, No.12, pp. 1326-1335.

Patterson, D. A., Brown, A. B., Broadwell, P., Candea, G., Chen, M., Cutler, J., Enriquez, P., Fox, A., Kiciman, E., Merzbacher, M., Oppenheimer, D., Sastry, N., Tetzlaff, W., Traupman, J., and Treuhaft, N. (2002). Recovery-Oriented Computing (ROC): Motivation, Definition, Techniques, and Case Studies: UC Berkeley Technical Report UCB/CSD-02-1175.

Pethia, R. D. (2003). “Viruses and Worms: What Can We Do About Them?”, *Congressional testimony Before the House Committee on Government Reform*: CERT Coordination Center.

Rasmussen, J. (1986). *Information Processing and Human-Machine Interaction*, Amsterdam: North-Holland.

Reason, J. (1990). *Human Error*, Cambridge: Cambridge University Press.

Richardson, R. (2003). CSI/FBI Computer crime and security survey: Computer Security Institution, http://www.gocsi.com (accessed August 2004)

Salter, J. H. and Schroeder, M. D. (1975). “The protection of information in computer systems,” *Proceedings of the IEEE*, Vol.63, No.9, pp. 1278-1308.

Schenk, K. D., Vitalari, N. P., and Davis, K. S. (1998). “Differences between novice and expert systems analysts: What do we know and what do we do?,” *Journal of Management Information Systems*, Vol.15, No.1, pp. 9-50.

Shimeall, T. and Williams, P. (2002). “Models of Information Security Trend Analysis,” SPIE Aerosense Conference, Orlando, FL.

Straub, D. W. and Welke, R. J. (1998). “Coping with systems risk: Security planning models for management decision making,” *Mis Quarterly*, Vol.22, No.4, pp. 441-469.

Walker, S. T. (1985). “Network Security Overview,” *IEEE Symposium on Security and Privacy*, Oakland, CA.

Warren, M. and Hutchinson, W. (2001). “Cyber Terrorism and the Contemporary Corporation,” in Dhillon, G. (Ed.), *Information Security Management: Global Challenges in the New Millennium*, Hershey, PA: Idea Group Publishing, pp. 53-64.

Whitman, M. E. (2004). “In defense of the realm: understanding the threats to information security,” *International Journal of Information Management*, Vol.24, No.1, pp. 43-57.

Woodward, D. (2000). “Smart Security,” *The British Journal of Administrative Management*, Vol.18, pp. 22-23.

Zurko, M. E. and Simon, R. T. (1996). “User-centered security,” *ACM New Security Paradigms Workshop*, Lake Arrowhead, CA.

About the Author

Richard L. Baskerville’s research and authored works regard security of information systems, methods of information systems design and development, and the interaction of information systems and organizations. He is a Chartered Engineer, holds a B.S. *summa cum laude*, from The University of Maryland, and the M.Sc. and Ph.D. degrees from The London School of Economics, University of London.

Ghi Paul Im is a doctoral student in the department of Computer Information Systems at Georgia State University. His research interests center on computer security, organizational learning, knowledge management, and innovation enabled by information technology. He holds an M.S. in Information Systems from NYU Stern School of Business.