Information Security is Information Risk Management

Bob Blakley
Tivoli Systems, Inc.
blakley@us.ibm.com

Ellen McDermott
J.P. MorganChase

Dan Geer
@Stake

ABSTRACT

Information security is important in proportion to an organization's dependence on information technology. When an organization's information is exposed to risk, the use of information security technology is obviously appropriate. Current information security technology, however, deals with only a small fraction of the problem of information risk. In fact, the evidence increasingly suggests that information security technology does not reduce information risk very effectively.This paper argues that we must reconsider our approach to information security from the ground up if we are to deal effectively with the problem of information risk, and proposes a new model inspired by the history of medicine.

1. INFORMATION RISK

Information security is required because the technology applied to information creates risks. Broadly, information might be improperly disclosed (that is, its confidentiality could be compromised), modified in an inappropriate way (that is, its integrity could be compromised), or destroyed or lost (that is, its availability could be compromised).

Compromise of a valuable information asset will cause dollar losses to the information's owner whether acknowledged or not; the loss could be either direct (through reduction in the value of the information asset itself) or indirect (through service interruption, damage to the reputation of the information's owner, loss of competitive advantage, legal liability, or other mechanisms).

1.1 What is Risk?

In business terms, a risk is the possibility of an event which would reduce the value of the business were it to occur. Such an event is called an "adverse event."

Every risk has a cost, and that cost can be (more or less precisely) quantified. The cost of a particular risk during a particular period of time is the probability of an adverse event occurring during the time period multiplied by the downside consequence of the adverse event. The probability of an event occurring is a number between zero and one, with zero representing an event which will definitely not occur and one representing an event which definitely will occur. The consequence of an event is the dollar amount of the reduction in business value which the event will cause if it occurs [Har]

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
NSPW'01, September 10-13th, 2002, Cloudcroft, New Mexico, USA.
Copyright 2002 ACM 1-58113-457-6/01/0009...$5.00.

1.2 Measuring Risk

A common measure of the cost of risk is "Annualized Loss Expectation," or ALE. ALE is the expected cumulative cost of risk over a period of one year as estimated in advance. For example, a chemical company estimates the probability of an explosion at one of its plants during the year 2001 as one in a million. If an explosion occurs, it will cost the company 150 million dollars in direct and indirect expenses, (for example, repair costs, legal costs, or lost business).

The ALE created by the risk of a plant explosion for the year 2001 is simply:

$$ALE = \$15{,}000{,}000 \times (1/1{,}000{,}000) = \$150$$

It's important to understand that the actual cost of this risk will never be that of the ALE, i.e., it will never be $150 during a particular year – it will be either $0 or $150 million. In less certain situations, the probability or the cost may be ranges rather than point estimates. If the probability of the explosion is between one in five hundred thousand and one in a million while the cost varies between 100 million and 200 million, the ALE would be:

$$ALE = (\$100M..\$200M) \times (1/500{,}000..1/1{,}000{,}000 = \$100..\$400$$

It may be possible to estimate the probability distribution of expected loss within the range (so for example, the ALE for the example above might be uniformly distributed between $100 and $400). ALEs can also be figured based on inequalities, as is doubtless obvious.

2. MANAGING RISK

Businesses routinely manage risk as part of their day-to-day operations. Risks can be managed using a variety of mechanisms, including liability transfer, indemnification, mitigation, and retention.

2.1 Liability Transfer

A business can transfer liability for an adverse event to another party. This takes the risk off the business's books. Liability can be transferred in two ways: by disclaimer and by agreement.

- A business disclaims liability when it undertakes an activity with the explicit understanding that it will not be held responsible for the consequences of certain adverse events, but without specifying who will be responsible for those consequences.
- A business transfers liability by entering into an agreement; to do this the business engages in an activity with counter-party after they both agree that the counter-party will be responsible for the consequences of certain adverse events.

2.2 Indemnification

A business can indemnify itself against the consequences of an adverse event. There are two major types of indemnification: pooling and hedging.

- In pooling schemes, several businesses share the cost of certain risks. If adverse events are unlikely to happen simultaneously to a meaningful fraction of the businesses in the pool, pooling will decrease the cost of risk to each organization in the pool while increasing the predictability of the cost of risk for each business in the pool. Insurance policies are the most common type of risk-pooling scheme.
- In hedging schemes, a single business essentially places a bet that an adverse event will happen to it. If the event is improbable, other organizations or individuals are likely to take the bet, because the probability is high that they will win the bet. If the adverse event does not happen, the business will pay off the bet. If the adverse event does happen, the bettors will have to pay the business. In this case, the business uses the money it collects from winning the bet to defray the costs of the adverse event. The key to a successful hedging scheme is getting the odds right on the bet. Being better than others at estimating the true odds of an adverse event can enable a business or an individual to make money on hedging schemes in the same way as casinos make money on card games. Options are the best-known example of risk-hedging scheme.

2.3 Mitigation

A business can try to reduce the expected cost of a risk, either by reducing the probability of the adverse event occurring, or by reducing the consequences if it does occur.

- The probability of an adverse event can be reduced by redesigning systems or processes to eliminate the event's known or suspected causes. In the extreme case, the probability of an event can be reduced to zero by entirely avoiding the activity which creates the risk. In business terms this might mean foregoing an opportunity which has potential rewards but also carries substantial risk.
- The consequences of an adverse event can be reduced by taking steps to limit the damage the event causes. These steps either prevent the damage caused by the adverse event from spreading, or they shorten the time during which the event causes damage by accelerating detection and recovery. Building codes that anticipate earthquakes do nothing to prevent earthquakes but they do lessen the damage that would otherwise be inevitable and uncontrolled.

2.4 Retention

If an adverse event is not very costly or not very likely to occur, or if the benefits to be realized from taking a risk are great, a business may choose to retain the risk which the adverse event creates.

- If the business chooses to set aside funds to offset the cost of retained risks, it is said to self-insure against these risks. Cyclical industries often approach inherent sector risk in this way, storing up funds in fat years against the lean.
- A business which retains risks without setting aside funds to offset their costs is said to accept retained risks. Many large companies do this with respect to the travel risks their employees incur, for example when they rent automobiles.

3. INFORMATION SECURITY

Up to this point we have used examples unlrelated to information risk to illustrate risk management. Failures of information security are clearly adverse events which cause losses to business; therefore, information security is a risk management discipline, whose job is to manage the cost of information risk to the business.

3.1 What is Information Security?

Where information risk is well enough understood and at least in broad terms stable, information security starts with policies. These policies describe "who should be allowed to do what" to sensitive information.

Once an information security policy has been defined, the next task is to enforce the policy. To do this, the business deploys a mix of processes and technical mechanisms. These processes and mechanisms fall into four categories:

- Protection measures (both processes and technical mechanisms) aim to prevent adverse events from occurring.
- Detection measures alert the business when adverse events occur.
- Response measures deal with the consequences of adverse events and return the business to a safe condition after an event has been dealt with.
- Assurance measures validate the effectiveness and proper operation of protection, detection, and response measures.

The final information security task is an audit to determine the effectiveness of the measures taken to protect information against risk. We say "final" but, obviously, the job of information risk management is never done. The policy definition, protection, and audit tasks are performed over and over again, and the lessons learned each time through the cycle are applied during the next cycle.

3.2 What's wrong with information security?

It's increasingly evident that information security as defined above simply isn't doing the job. Every day, newspapers and trade journals carry stories of the latest virus, denial-of-service attack, website defacement, or bug in an important security product. The public is getting the message even if the only sensible reaction is dread.

Why is information security failing? We posit two reasons: information security focuses on only a small part of the problem of information risk, and it doesn't do a very good job of protecting businesses against even that small part.

3.2.1 Focus

Information security technology focuses primarily on risk mitigation. Information security risk analysis processes are geared toward imagining and then confirming technical vulnerabilities in information systems, so that steps can be taken to mitigate the risks those vulnerabilities create. In some cases management will be asked to sign a risk acceptance

(that is, to retain a risk) after a risk analysis. A risk acceptance will typically include either a plan for future mitigation or a justification of the economic rationale for choosing not to mitigate.

Information security as a discipline is often biased

- toward technological mechanisms rather than process mechanisms,
- in favor of logical (that is, computer hardware and software) mechanisms, and
- against physical mechanisms (such as locks, walls, cameras, etc...)

Even within the category of risk minimization activities, information security focuses more on reducing probability of an adverse event than on reducing its consequences. And where consequence reduction is implemented, it tends to focus much more strongly on quick recovery (for example, by using aggressive auditing to identify the last known good state of the system) than on minimizing the magnitude of a loss through measures to prevent damage from spreading.

Information security activities rarely include any discussion of indemnity or liability transfer, although some organizations do address these issues in an "operational risk" organization separate from the information security organization.

The following chart organizes information security products and processes according to the risk management activities they implement. The chart clearly illustrates the problem.

Table 1.

		risk management activity							
		transfer liability		indemnify		mitigate		retain	
		shift	share	pool	hedge	reduce prob.	reduce cons.	self-insure	accept
security method	process	notice	Empl. Agree.			ACL, firewall, PKI, crypto, pre-hire screen, filter web	anti-virus,		risk accept.
	data					audit, snoop email,	intrus. detect,		
	required						disast. recov., back-up, key recov.,		

3.2.2 Effectiveness

The annual FBI/CSI computer crime surveys and the CERT coordination center annual summaries [CERT] have shown substantial increases in the number of security incidents and in the dollar losses resulting from incidents in each of the past five years.

The year 2000 FBI/CSI survey [CSI] nevertheless reports that use of information security technologies is very widespread – close to 100% of companies responding to the FBI/CSI survey use antivirus, firewall, and access control technologies.

The combination of nearly universal deployment of security technology with rapidly and steadily rising losses strongly suggests that security technologies (and processes, although these are not covered in the FBI/CSI survey) do not prevent losses – in other words, they don't work!

Further, as Arbaugh, Fithen, and McHugh have shown [AFM], identification of a vulnerability and its exploitation are both separated in time. Furthermore, risks arising from a vulnerability are often multiplied both by scripting of the attack and by the haphazard deployment of patches even when they are easily available.

4. QUANTIFICATION OF INFORMATION SECURITY RISK

Risk analysis has been recognized as an important information security discipline for a long time. Information security risk analysis methodologies were developed long ago, and some of these methodologies have been included in formal information security standards. The large majority of these standards have been qualitative – that is, their assessment of probability and consequence of risks is based on a "low/medium/high" characterization rather than on a specific probability and a specific dollar amount of loss. Qualitative information security risk management standards include the US Federal standards [FIPS31] and [FIPS191]. Recent guidelines which recommend qualitative risk analysis techniques include [GAO] and the newly issued draft [NISTRMG].

Quantitative information security risk management standards have been developed, including the now withdrawn [FIPS65]. The authors are not aware of any current information security standard which mandates the use of a quantitative risk analysis method, though the Australian national standard for risk analysis [AS] permits the use of either qualitative or quantitative analysis. Methodologies for quantitative and mixed quantitative/qualitative information security risk analysis have been published; see for example [Pelt]. Quantitative risk analysis is used extensively in disciplines other than information security, including finance, healthcare, and safety (see [KBPS] for a number of examples). There is a large body of literature on methods for quantitative risk analysis in these fields; sources include [Koll] and [Vose].

Good data is a prerequisite to qualitative risk analysis, and the lack of good data may be the main reason qualitative analysis of information security risk is not usually performed. [GAO] explicitly acknowledges this: "Reliably assessing information security risks can be more difficult than assessing other types of risks, because the data on the likelihood and costs associated with information security risk factors are often more limited and because risk factors are constantly changing."

Insurers seem to agree that data is lacking. The National Underwriter Company's guide to risk in the wired world [ERisk] warns: "The lack of historical data presents one of the most difficult challenges when trying to analyze online exposures... the insurance industry typically depends on large bodies of actuarial data collected over long periods of time to develop pricing models for insurable exposures. But in the Wired World exposures are so new and are growing so rapidly in terms of frequency and severity that this is not an easy task."

Despite the lack of actuarial data, many insurers (including AIG, Lloyd's, Chubb, Zurich, and others; a partial list can be found in [ECov]) are offering policies which cover losses due to failures of information security. But the actuarial basis for these policies is unclear, as the National Underwriter Company [ECov] explains: "The insurance industry has worked closely

with actuaries and financial analysts to map out the calculations for the probabilities of loss, the probable costs for various scenarios of loss, and for rates and rating structures for acceptable risks. After years, decades, and even centuries of study, calculations for various probabilities have been developed. These tables and charts typically deal in a world where the events that tend to cause the damages have been identified previously and provide a basis for which the future can be predicted. The new economy has disrupted this equilibrium. New risks are emerging, and the insurance industry has had only a brief period of time to scratch the surface for potential liabilities. So far there have been relatively few claims that have materially affected the technology industry. It is too early to establish actuarial tables to quantify technology risks... Because the actuaries don't have the data needed to predict losses, the financial analysts are hampered in predicting the financial viability of insuring technology risks."

The finance industry certainly sees a lack of information security risk data. The revised Basel accord [Basel], which governs the amount of capital that banks must set aside as a hedge against risk, requires for the first time that banks set aside capital to offset operational risk (which includes information security risks). Banks which can demonstrate, starting in January 2005 and based on 3 years of auditable data, that their risk exposure is lower than the Basel accord's estimate, can reduce their capital setaside from the very substantial amount required by the accord as a baseline. A large number of financial institutions have commented [BCom] on the revised accord. The American Banker's Association wrote "...only a few institutions appear to be actively modeling operational risk, and modeling is very much in development... whether attempting to model operational risks or not, most banks have not captured the data necessary to evaluate operational risk, even at a theoretical level". Bank of America wrote "We do not believe that operational risks are measurable using methods and data that are available at this time... Only a handful of banks have implemented quantitative approaches for measuring operational risk and the models are largely untested." The Richmond Federal Reserve wrote: "We are concerned about the lack of data on operational risk, and acknowledge that banks have been very reluctant to publicize details of losses from such problems as deficiencies in internal controls, human error, or system failure."

In order to quantify information security risk, and the effectiveness of information security risk control measures, the following information needs to be collected. Some is already in good supply, some is not. There will be temptations to extrapolate from available data to less-available data, and to apply risk-measurement methods which are already understood outside of their appropriate domains of use; the authors caution that these temptations should be avoided.

4.1 Vulnerabilities

A comprehensive list of information security vulnerabilities needs to be developed. For each vulnerability, information needs to be gathered and regularly updated about the ease and frequency of exploitation, and ease and speed of recovery from exploitation. This information must be collected and made available in a way that demonstrably minimizes the probability of exploitation in an economically harmful way

4.2 Incidents

Information needs to be gathered about security incidents experienced by businesses worldwide. This information must include what vulnerabilities were exploited and how response and recovery were handled. Incidents that are traceable to vulnerabilities already known are one thing and will be a matter of discussion between insurers and victims if in no other situation. Incidents that highlight previously unknown vulnerabilities must be fed back to that catalog. This information needs to be collected and made available in a way which does not create additional liabilities for the reporting organizations (and hence incentives to avoid reporting).

4.3 Losses

For each incident identified, information needs to be collected about direct monetary losses caused by the incident and about indirect losses (for example, reputation damage or lost business) with an estimate of the monetary losses resulting from these indirect losses. The calculation of losses needs to be done using a uniform methodology, and the information needs to be collected and made available in a way which does not create additional liabilities for the reporting organizations.

The National Underwriter Company [ECov], recognizing the lack of this kind of actuarial information about information-security-related losses, has solicited the aid of the technology staff of the insurance industry itself in fixing the problem: "Even though insurance IT staffers can revert to the same techie talk that technology clients use, they are often required to explain technological advancements and enhancements to upper management of the insurance company, especially when discussing IT expenditures. If they can do that, why can't they be used to help underwriters develop assessment and underwriting tools and train claims professionals in the intricacies of IT losses."

We ask a similar question: if the IT security industry can design countermeasures and counsel clients on how to defend their systems, why can't we help underwriters develop assessment and underwriting tools and train claims professionals in the intricacies of IT losses? Do we have something more important to do?

4.4 Countermeasure Effectiveness

A comprehensive list of available security measures needs to be developed, together with information about about the cost of acquiring, managing, and maintaining each security measure. For each incident identified, information needs to be collected about which security measures were in use at the time of the incident, which security measures were bypassed, which security measures were defeated, and how much time and effort were required to circumvent or defeat the security measures in place. Some mechanism must be put in place to combat the obvious temptations to distort pre- and post-event readiness and protection postures and event details in order to obscure or conceal the occurrence of events, to embellish war stories, or to avoid personal or corporate accountability.

5. WHAT DOES THE CURRENT SITUATION LOOK LIKE?

We have described a world in which we have very little information about frequency of occurrence of adverse events and about the seriousness of their consequences. We also

know very little about the effectiveness of the measures we take to prevent adverse events or alleviate their consequences. The people to whom these events happen have few incentives to report them; conversely, they have many incentives to suppress information about them. Finally, the system we are attempting to protect (roughly composed of the global Internet and everything attached to it) is far too complex to be understood in detail.

This situation looks to the authors very much like the state of medical practice in the 19th century (for a good general treatment of the development of scientific medicine, see [Por], which includes an extensive bibliography). Medical practitioners had a poor understanding of the prevalence, , and likely outcomes of illness causes (the 1899 first edition of the Merck Manual [Mer1] contains no information about causes, symptoms, or mortality rates of the conditions it describes; it consists entirely of lists of preparations which could be administered for each condition, with no advice on how to choose among the many options), and the safety and effectiveness of treatments (The 1900 edition of the Old Farmer's Almanac includes an advertisement for Wistar's Balsam of Wild Cherry, which claims that "It is the most reliable preparation in the world for the cure of Coughs, Influenza, Bronchitis, Whooping Cough, and all Throat and Lung Troubles, and in many well attested cases, Consumption [i.e. Tuberculosis] has yielded to its wonderful influence" [OFA]). The public feared medical treatment (for good reasons, despite frequent outbreaks of serious diseases), and widely considered medicine to be ineffective. And of course, the human organism was too complex to really understand.

The world of medicine today is very different – even though the human organism is still too complex to understand. Today, drug advertising is heavily regulated, and advertisements are required to provide extensive information on side effects, effectiveness as measured in clinical studies, contraindications, interactions with other medications, considerations for use in children and pregnant women, and so on.

The 2000 Centennial Edition of the Merck Manual [Mer17] lists, for each condition it describes, the cause or causes, etiology and pathology information, related or similar conditions together with methods for distinguishing between them, syptoms, signs, and methods of diagnosis, laboratory tests and findings, and prognosis and treatment regimens. Much of this information is based on quantitative studies of outcomes.

The 2002 edition of the Prentice-Hall Health Professional's Drug Guide [HPDG] includes, for each listed medication, information on action and pharmacodynamics, uses (including unlabelled uses), pregnancy risk category, routes of administration and dosages, pharmacokinetics, contraindications and precautions, adverse reactions and side effects, interactions with drugs and medicinal herbs, assessment of patients during the course of therapy, and patient and family education. Again, this information is based on strict quantitative studies of use of the medications included.

What has made all this possible is the increased professionalism of medical practice, based in large part on the collection and study of quantitative data about prevalence and outcomes of illnesses and treatments. Three critical developments helped modernize western medicine:

- Mandatory professional education and licensure of practitioners
- Systematic collection and study of public health data
- Systematic observational studies of safety and effectiveness of treatments

We propose that these same developments would put information risk management on a sound footing. In the next three sections, we make specific proposals which could drive these developments into the practice of information risk management.

6. HOW SHOULD INFORMATION RISK BE MANAGED?

Today, information risk management professionals have training but often no formal information risk management education. They don't hold revocable licenses (or any licenses at all). They have no formally recognized ethical obligation to use only safe, effective risk management treatments for the problems they encounter. No professional body exists which could discipline ethical lapses if they occurred. There is no ethical obligation imposed on information risk management professionals to avoid the use of ineffective or even harmful treatments. There is no obligation of confidentiality to the organizations they treat – other than those negotiated on a case-by-case basis in employment agreements or consulting contracts. There is no obligation whatsoever to report information which might have "public health" or "public safety" implications to an established authority (and in fact sometimes the aforementioned employment agreements and consulting contracts explicitly forbid such disclosures).

The authors posit that in the future, information risk should be treated by professionals with the characteristics of a physician. A physician has:

- A specialized professional education
- A revocable license to practice
- An ethical obligation to treat patients appropriately and keep their private information in confidence
- A professional obligation to control (through the power of prescription) the use of potentially harmful treatments
- A professional obligation to report important public health information to the proper authorities.

Information risk professionals should have all these things too. Particularly important in our view are the ethical obligation to apply only appropriate treatments and protect confidentiality of those treated, and the professional obligation to report information to "public health" authorities.

The information risk management professional's obligation to treat appropriately, and to control the use of potentially harmful treatments, will require assessing the costs and benefits of all risk treatment options – liability transfer, indemnification, and retention as well as mitigation, detection and response as well as prevention, and procedural as well as technical treatments. Choice of treatment options should be based on the welfare of the "patient" – which will be maximized by optimizing cost of risk to the business rather than on minimizing probability of occurrence of adverse events. Needless to say, the information risk professional will

be obligated to avoid the use of risk treatments whose effectiveness is demonstrably low.

Professional training in management of information security risk should present a broad and integrated view treatments (including, for example, risk transfer and indemnification), rather than the one-dimensional, vulnerability-mitigation focus common today. At the simplest level, this means that information security risk education should include financial and legal disciplines in addition to the technical disciplines taught today. Some risk-management experts have begun to describe how risk management activities can be integrated across the entire spectrum of business risks [Shim]; information security education should be built on this kind of comprehensive framework.

6.1 Reporting

Today, almost all information security risk assessments use qualitative rather than quantitative methods. Some risk analysis methodologies and standards already incorporate rudimentary loss-expectation estimation methods, but these are usually limited to a "low/medium/high" categorization with arbitrary dollar ranges assigned to the categories. Some industries already quantify intellectual property risk in financial terms and take steps to manage risk using financial instruments.

Risk assessment findings are essentially never shared with anyone except the business being assessed, and possibly its external auditors.

In the future, the authors believe that information security risk assessments should focus not just on identifying risks, but also on quantifying them. Specifically, information security risks should be characterized in financial terms, as annualized loss expectations

Once risks are identified and quantified, the resulting data should be reported (by the information risk management professionals, in a way that respects their ethical obligation to protect the privacy of those they treat) to the information risk equivalent of a public health service. The next section discusses this service at more length.

7. HOW SHOULD INFORMATION RISK BE STUDIED?

Today, some data on risk prevalence and severity is collected by the US FBI, CERT, and other organizations. However, reporting to these organizations is voluntary, and only a small sample of businesses even receive the questionnaires which these bodies use to collect their summary information. Furthermore, no standard taxonomies of vulnerabilities, incidents, losses, or countermeasures are used in the collection or reporting of this information.

In the future, collection of data on information risk needs to be much more regular, formal, and comprehensive. Information risk should be studied by an independent body with the characteristics of a public health service. This "Public Security Service" should collect from information risk management professionals, in a way which protects the privacy of the organizations those professionals treat, data on the prevalence of losses, the causes of losses, the effects of losses, and the effectiveness of information risk treatments. The Public Security Service should analyze this data and publish the results of its analyses as a way to improve the state of information risk management practice, and to inform public policy decisions about information risk management.

Obviously, the advanced research which drives the development of new treatments and deeper understanding of the causes of risks will continue to be carried out in the academic and business communities, just as advanced medical research into new drugs and the causes of disease is carried out by academic medical schools and pharmaceutical research labs today.

8. HOW SHOULD INFORMATION SECURITY TECHNOLOGY BE EVALUATED?

Today, information security technologies are subjected to design and implementation analyses defined by a number of assurance regimes (most notably the Common Criteria [CC]). Businesses can also submit voluntarily to "seal" programs, whose certifications are based on deployment of popular technologies, and on contract, process and system configuration audits. Businesses can contract for penetration testing, but the authors are not aware of any certification regime which requires penetration testing, or any other explicit measure of the effectiveness of security protection measures, as a condition of granting certification.

No systematic effectiveness testing of information security measures is done by any independent body, and the results of effectiveness testing done by vendors and their contractors are almost never published. Information risk management professionals have no training in the design of experiments to test effectiveness of the measures they design, and no training in publishing or reviewing the results of such experiments.

A workshop participant pointed out that the information security industry has no equivalent of the white laboratory mouse which can be used to test the effectiveness of security mechanisms without having to subject business' production systems to unethical levels of risk. This is an important, and true, observation.

The authors observe also, however, that medicine has not always had white laboratory mice as models either, and we urge research into the development of an appropriate "security mouse analog" for use as an effectiveness testbed for security measures.

In the future, the authors believe that the effectiveness of information security technology would be most effectively evaluated by an impartial body following a process similar to the one used by the US Food and Drug Administration (FDA) to approve medical treatments for use. The FDA's process is based on systematic, quantitative observational studies of actual outcomes, and includes an ongoing monitoring phase which updates safety and effectiveness information after treatments have been approved and are in use by the medical community.

Security technology development and selection should be based on quantitative observational studies of effectiveness, not on synthetic a priori assurance of vulnerability avoidance. Probabilities of explotation must be balanced with consequences. ALEs (that is, observed outcomes) must rule, not the emotion of a good story and the fear, uncertainty and doubt that continues to be the selling proposition for most security technology.

While assessment of technical vulnerabilities and the likelihood of their exploitation should and will remain a part of information technology risk management, assessment must include the overall risk control process, including personnel, physical, and technical measures. It must be sensitive to the rate of change in each of these parameters.

A determined effort should be made to evaluate all kinds of protection, detection, and response measures (both technical and non-technical) to quantify how each measure the affects annualized loss expectation arising from many specific kinds of risks.

The impartial body which carries out evaluations could be a government agency (such as the US NCSC) or government-sponsored security laboratory (such as the CERT Coordination Center), a commercial organization through a seal program, an industry consortium such as IT-ISAC, an insurers' consortium similar to Underwriters' Laboratories, a consumer organization similar to Consumers' Union, or a combination of some or all of the above.

Information risk management professionals should, as stated in the previous section, be professionally obligated to avoid the use of demonstrably ineffective treatments.

8.1 Tracking and Reporting

Today, no equivalent of The Lancet or Journal of the American Medical Association exists to enable publication and review of information about the effectiveness of information risk treatments, and information risk management professionals do not have training in technical writing or review of other practitioners' results. We note in passing that journals of this sort are useful to, and used by working practitioners (not just academics) in some disciplines; for example, police laboratory personnel regularly publish in and read the Journal of Forensic Science.

The effectiveness of information risk treatments will change over time as the technical environment and the risk environment "in the wild" evolve. Information risk management professionals should be required to report regularly to the evaluation body on the effectiveness of the treatments they "prescribe" to their "patients". The evaluation body should continually update its assessments of treatment effectiveness based on the information it receives, and should distribute these updates to the community of information risk management professionals.

9. A WORD ABOUT THE ETHICS OF RISK QUANTIFICATION

A review of an earlier draft of this paper questioned whether quantification of certain types of risks (particularly risks to human life and safety) in financial terms is ethically acceptable.

The first point to be made in this context is that systems which pose known or suspected risks to human life or safety should be treated using techniques for managing risk in safety-critical systems, even if they also require information security risk treatment (see for example [Leve] or [Stor] for full treatments of risk in safety-critical systems). The authors do not claim that information security risk management techniques do, or should, protect against safety risks.

The second point to be made is that society must take risks it considers unacceptable out of the realm of economic justification by imposing mandatory control regimes. Serious safety risks should be controlled using a regime which is not voluntary and is not based on a cost/benefit analysis. If a society concludes that a certain safety risk is sufficiently serious that controlling it is mandatory, that society should use legal and regulatory mechanisms to mandate control of that risk.

At least in capitalist societies, any risk for which there is no legally required control regime will be controlled only to the extent that the cost of control can be economically justified. The economics of controlling risks can be distorted by competition. Risk-tolerant firms may gain temporary competitive advantage against risk-averse firms by spending less on control (especially for risks with low probability of occurrence) as long as they are lucky and the risks do not cause them losses. The authors maintain that cost-justifying risk controls can only be effective if the risks can be quantified.

The third point which needs to be made is that accurate quantification of the costs of risks to human life and safety might in fact provide powerful incentives for control. Putting a price tag on a human life is certainly fraught with ethical dangers. On the other hand, if NASA had had a realistic estimate of the probability that the Space Shuttle Challenger would be destroyed, and had also had an accurate estimate of the financial and reputation costs of this event, there seems little doubt that the Challenger launch would have been delayed and the ship saved.

One argument against this point of view might be that the real cost of the loss of a life to the organization which causes the loss is not very great in some cases. Estimates of the total cost of the Union Carbide Bhopal plant to the Union Carbide corporation vary, but the direct cost of the legal settlement ($US 470 million) represents only about $US 12,400 for each of the roughly 3800 people killed by the accident, and this does not include consideration of the more than 2700 people permanently disabled. Twelve thousand dollars for a human life is an uncomfortably low figure. Does this mean that quantifying this risk is ethically irresponsible? The authors think not – the fact that a life costs a major corporation only $12,000 looks to us like a call for reform of the liability system.

In summary, while the authors do not believe that every risk should be controlled using a monetary cost/benefit framework, we do believe that all risks should be quantified to the greatest extent possible, regardless of the anticipated control regime. We also believe that information security risks will be poorly understood until we do a much better job of quantification of economic losses. Finally, we believe that information security countermeasures will continue to be difficult to justify in voluntary control regimes until their effectiveness can be expressed as a quantifiable reduction of economic losses.

10. ACKNOWLEDGEMENTS

The authors are grateful to the anonymous reviewers, and specifically to the reviewer who called the issue of the ethics of quantification of losses to our attention.

The first author would like to acknowledge the support of the Open Group in providing an ongoing forum for his investigations into the topic of security and Risk Management.

11. REFERENCES

[AS] Standards Australia, "AS/NZS 4360:1999 Risk Management", 1999.

[CSI] Computer Security Institute and US FBI, "Computer Security Issues & Trends", CSI, 2000.

[AFM] Arbaugh, W., Fithen, W., and McHugh, J., "Windows of Vulnerability, a Case Study Analysis", IEEE Computer, IEEE, December, 2000.

[Basel] Bank for International Settlements, "The New Basel Capital Accord", Basel: Bank for International Settlements, 2001.

[Bcom] Comments on New Basel Capital Accord, http://www.bis.org/bcbs/cacomments.htm

[CERT] CERT, CERT Annual Reports, http://www.cert.org/annual_rpts/index.html

[ECov] TechRisk.Law, "e-Coverage", Cincinnati, OH: National Underwriter Company, 2000.

[ERisk] Lang, S., Davis, J., Jaye, D., Erwin, D., Mullarney, J., Clarke, L., and Loesch, M., "e-risk: Liabilities in a Wired World", Cincinnati, OH: National Underwriter Company, 2000.

[FIPS31] US Department of Commerce/National Bureau of Standards, "Guidelines For Automatic Data Processing Physical Security and Risk Management", 1974.

[FIPS191] US Department of Commerce/National Institute of Standards and Technology, "Guideline for the Analysis of Local Area Network Security", 1994.

[GAO] US General Accounting Office, "Information Security Risk Assessment: Practices of Leading Organizations", 1999.

[Har] Harrington, S., and Niehaus, G., "Risk Management and Insurance", Boston, Irwin/McGraw Hill, 1999.

[HPDG] Shannon, M., Wilson, B., and Stang, C. (eds.), "Health Professional's Drug Guide", Upper Saddle River, NJ, Prentice Hall, 2002.

[Koll] Koller, G., "Risk Assessment and Decision Making in Business and Industry", Boca Raton, Fla.: CRC Press, 1999.

[KBPS] Kolluru, R., Bartell, S., Pitblado, R., and Stricoff, S., "Risk Assessment and Management Handbook for Environmental, Health, and Safety Professionals", Boston: McGraw-Hill, 1996.

[Leve] Leveson, N., "Safeware: System Safety and Computers", Reading, Mass.: Addison-Wesley, 1995.

[Merl] Merck & Co., "Merck's 1899 Manual", New York, Merck & Co., 1899.

[Mer17] Beers, M., and Berkow, R. (eds.), "The Merck Manual of Diagnosis and Therapy", 17th ed., Whitehouse Station, NJ, Merck Research Laboratories, 1999.

[NISTRMG] US National Institute of Standards and Technology, "Special Publication 800-30: Risk Management Guide" (Draft), 2001.

[OFA] Thomas, R. (ed.), "Old Farmer's Almanac", William Ware & Co., Boston, 1900.

[Pelt] Peltier, T., "Information Security Risk Analysis", Boca Raton, Fla: Auerbach Publications, 2001.

[Por] Porter, R., "The Greatest Benefit to Mankind", New York, W.W. Norton & Company, 1997.

[Shim] Shimpi, P., "Integrating Corporate Risk Management, New York, Texere, 1999.

[Stor] Storey, N., "Safety-Critical Computer Systems", Reading, Mass.: Addison-Wesley, 1996.