Government Chief InformationOffice (GCIO)

A Guide for Government Agencies Calculating Return on Security Investment

Version 2.0

Lockstep Consulting
13 June 2004

Table of Contents

1. Executive summary 3

2. Introduction 4

2.1 Objectives 4
2.2 New tools featured in this Guide 5

3. A survey of security cost-benefit approaches 6

3.1 Qualitative justifications for security investment 6
3.2 Annualised Loss Expectancy 7
3.3 Security Attribute Evaluation Method (SAEM) 8
3.4 Cost-Effectiveness Analysis 9
3.5 Fault Tree Analysis 9
3.6 Limitations of existing approaches 10
3.7 A new ROSI model 10

4. Recap: Information Security Risk Management 11

4.1 Consequences to government of security incidents 11
4.2 Threat & Risk Assessment methods 11
4.3 Types of Countermeasures 12

5. Quantifying likelihood and costs of incidents 13

5.1 Likelihood of security incidents 13
5.2 Direct costs of security incidents 14
5.3 Indirect (opportunity) cost of security incidents 15
5.4 Statistical variability of security incidents 16

6. The proposed ROSI model 18

6.1 The expected annual cost of security incidents 18
6.2 The basic spreadsheet tool 21
6.3 Parameters of the basic tool 21
6.4 The statistical module 22
6.5 Interpreting the statistical results 27

7. Limitations of the tool 29

7.1 Limited iterations in the Monte Carlo freeware 29
7.2 Restricted sources of randomness 29
7.3 Difficult to separate effects of countermeasures 29
7.4 Difficult to predict how countermeasures affect severity 30
7.5 Security incidents are not necessarily independent 31

8. Bibliography 32

8.1 Cost-benefit references 32
8.2 Further reading on cost-benefit 32
8.3 Monte Carlo and other tools 33

1. Executive summary

This Guide and research report is intended to assist government agency IT managers evaluate and quantify the potential Return On Security Investment (ROSI) from implementing perimeter security systems.

The ROSI Guide started as a research project in early 2003, examining available approaches to measuring the cost-benefit of information security. The first version of the Guide proposed a hybrid tool, implemented as an Excel spreadsheet, combining the Annualised Loss Expectancy method with an Australian-standard Threat & Risk Assessment framework.

This latest version of the ROSI Guide describes an extension to the tool introducing "Monte Carlo" statistical analysis of the possible spread in cost-benefit results arising because security incidents vary randomly in their rate of occurrence and their severity. A prototype extended spreadsheet is attached, incorporating freeware Monte Carlo add-ins. Users are able to insert their own values for the expected ranges of incidence and costs for different grades of security incidents, drawing on the actual experience of their respective departments.

This report includes a discussion of how and why statistical variability should be injected into the ROSI model, instructions for running the chosen Monte Carlo tools, example simulations drawn from actual TRA, and an updated reference list to aid with further research into statistical cost-benefit analysis.

2. Introduction

This document was commissioned by the NSW Department of Commerce Office of Information and Communications Technology (OICT), to provide assistance to government agencies seeking to better understand the benefits of investing in information security systems. While most managers appreciate the "strategic" or qualitative returns from good information security, a financial benefits model is often desirable, to round out the business case, or to engage non-technical stakeholders in the investment approval process. As one author puts it, "specialists usually make security decisions, but program managers are left wondering whether their investment in security is well spent" [1].

The tool provided by this Guide may be used to assist in the economic appraisal of proposed information security capabilities in accordance with NSW Treasury Guidelines [2].

This document (with its attachments) provides:

- A pragmatic "how to do it" Guide for appraising proposed expenditure on IT security, with wide applicability.
- A framework for transforming qualitative judgments of likelihood and severity into quantitative estimates of loss expectancy, with and without security.
- A spreadsheet tool that realises the ROSI Model, ready for running "what if" scenarios.
- A survey and critical review of available methods for assessing the cost-benefit of information security systems, together with reference list identifying important contemporary research and reference work, as well as suggested avenues for further research.

2.1 Objectives

The specific objectives of this Guide are as follows:

- Provide a tool and framework for deriving reasonable quantitative security cost-benefit estimates from qualitative judgements of likelihood and severity.
- The tool and framework should be *transparent* (i.e. open to inspection by users) and *parameterised* so that users can input their own empirical data where appropriate.

- The tool should be equally applicable to the acquisition of hardware and software controls, services, new or changed staff positions and activities such as training.

The tool described in this Guide is based on familiar standard ways of characterising security threats and risks, including (a) the expected likelihood and the expected severity without mitigation, and (b) the reduced likelihood and severity to be expected when given security countermeasures are put in place.

Furthermore, the tool can accommodate situations where several countermeasures are selected to address a single threat, contributing in varying degrees to reducing the likelihood and severity of an incident.

2.2 New tools featured in this Guide

The Guide includes two sets of tools, as summarised below, in the form of two Excel spreadsheets. One set of tools produces fixed value forecasts for cost-benefit; the other tools produce statistical forecasts showing the spread of cost-benefit values given the inherent uncertainty in predictions about security incidents.

- **Frequency calculators**: lookup tables for transforming qualitative standardised judgements of the likelihood of incidents into quantified annual rates of occurrence.

- **Cost calculators**: lookup tables for transforming qualitative standardised judgements of the severity of incidents into corresponding dollar costs. The lookup values in the cost tables, as well as the frequency tables above, can be easily customised.

- **Annual Security Cost-Benefit calculators**: tables based on a standard Threat-Risk Assessment methodology, which compute the expected total annual cost of untreated security incidents, and the expected annual cost once selected security countermeasures are implemented. The Cost-Benefit Calculators default to per-incident cost and annual frequency figures from the lookup tables above. But if the user has empirical data for the cost or probability of specific incidents, this data can be easily entered into the cost-benefit calculation instead.

3. A survey of security cost-benefit approaches

In recent times, Return On Investment has been a topic of great management interest for IT infrastructure across the board. ROI for Security – or "ROSI" – has received particular attention, perhaps because security is traditionally such an inscrutable field. ROSI entered the mainstream of IT management thinking in early 2002 with an article in CIO Magazine "Finally, a Real Return on Security Spending" [6]. Several other non-technical articles appeared around the same time providing useful qualitative advice and rough quantitative methods; see [10], [13] and [16].

Computer World magazine has since created a dedicated "knowledge centre" for the topic of ROI [10]. Although little of the content relates specifically to information security, the web site contains a wealth of tutorial material, links to quantitative tools, and useful opinion pieces for and against the state of the art in ROI estimation.

A variety of ways to evaluate cost-benefit of security investment have been used over the past ten years or more. Here we briefly review the main approaches.

3.1 Qualitative justifications for security investment

Traditional justifications for security expenditure are largely qualitative, or "strategic", in that they argue that *without* investing, the organisation will miss out on some broader benefit or other. Points made in favour of information security expenditure include:

- Security represents a cost of doing business.
- Security is akin to insurance costs.
- New e-business revenue streams may depend on proper security.
- Security is one aspect of risk management.
- Legal actions might result from failure to meet a general duty of care manifest as minimum security standards.
- Current resistance to security expenditure will shrink as the information age matures; after all, nobody questions the cost of building security anymore.

All of these are good points. Some may be compelling. Yet a quantitative, financially robust business case for information security investments can still be desirable. Some important stakeholders, such as Finance, may not be swayed by qualitative arguments alone.

3.2 Annualised Loss Expectancy

The ROSI model reported in CIO Magazine in February 2002 [1] drew on work done at the University of Idaho [7] and other institutions on "Annualised Loss Expectancy". The basic approach is to calculate the untreated losses that an organisation would expect to face, and compare those losses to the security investment needed to mitigate them.

The Annualised Loss Expectancy (ALE) from information security breaches is worked out empirically from the organisation's experience and intelligence regarding intrusions, viruses, denial of service attacks and so on. Real dollar losses can be readily ascribed (at least for private sector organisations) to the following outcomes of security incidents:

- revenue lost from e-commerce sites brought down for a time;
- lost custom and good will from clients adversely affected;
- overtime paid to IT staff and additional contractors working to bring systems back on line;
- consulting fees paid to external specialists assisting with data recovery, repairs, forensics, legal work and so on;
- damages to complainants suffering from cyber-crime or privacy infractions;
- repair bills for physical damage that can result from cyber-attacks in certain sectors, such as water and utilities.

To mitigate these expected losses, the organisation should invest some amount in perimeter security, including firewalls to restrict attacker's from gaining access, Intrusion Detection Systems for early warning of the reconnaissance that precedes any concerted attack, and Anti-Virus measures to detect malicious code of various forms. If the organisation decides to establish and operate the security system for itself, then the cost will be made up of:

- Set up costs
  - security product licence fees
  - servers and other hardware
  - (possibly) consulting fees on analysis and configuration

- Recurring costs
  - security product support & maintenance fees
  - IT security staff salary and on-costs
  - hiring costs associated with IT security staff turn-over
  - research costs for ongoing threat assessment and periodic evaluation of new technologies.

It is important to note that no information security system can ever be perfectly effective. To quantify the benefits of a given security installation, we therefore need to calculate the degree to which it mitigates the ALE. Typically, a properly installed and tuned perimeter defence system can achieve around 85% effectiveness in preventing or flagging security breaches (see [7]) so long as it is well maintained and properly patched as necessary.

The net financial benefit of a perimeter security system in the Annualised Loss Expectancy model is given by:

$$\boxed{\text{Annual Savings} = \text{ALE} \times \text{Effectiveness} - \text{Annual Cost}}$$

Note that the ALE model assumes that all security breaches carry the same cost implications, so that if we can stop 85% of breaches, then we assume that we are saving 85% of costs. See further discussion of this point under 3.6 below.

3.3 Security Attribute Evaluation Method (SAEM)

SAEM has been developed at Carnegie Mellon University in an effort to produce robust cost-benefit results for comparing one security architecture against another, in light of the real world difficulty that security specialists rarely have precise data on the benefits of the technologies (see Butler [5]). Not only is security specialised beyond the comprehension of most programme managers, but security specialists themselves often have only their personal experience, intuition and judgement to go on. Intuition of course can be a powerful weapon but it makes it difficult for non-specialist managers to objectively review security recommendations.

The SAEM methodology incorporates peer reviewed likelihood and impact rankings for the particular environment in question, and delivers a multi-variate weighting of the *relative* costs and benefits of alternate security designs.

SAEM does not deliver stand-alone quantitative cost estimates, and so is of limited direct relevance to this Guide. However, we commend Butler for further research and for its useful diagramming of the coverage of security measures against identified threats.

3.4 Cost-Effectiveness Analysis

Dating from defence sector work done in the 1970s, Cost Effectiveness Analysis has more recently been applied in the health sector to study the relative improvement in performance that results from proposed investments in new systems [14]. One strength of Cost Effectiveness Analysis is that it accommodates non-financial performance measures; in clinical practice for example, it can return the percentage improvement in mortality that one type of healthcare intervention might have over another.

The technique looks to be applicable to comparisons of candidate security countermeasures, assessed according to a variety of performance measures, such as system availability. We suggest that Cost Effectiveness Analysis be studied further as an adjunct to the tool presented here.

3.5 Fault Tree Analysis

While not traditionally a cost-benefit tool, Fault Tree Analysis (FTA) and its spin-off methods such as Failure Modes Effects [and Criticality] Analysis (FMEA and FMECA) seems to have promise for studying the root causes of security breaches and the mitigating effects of countermeasures. A Fault Tree is a graphical tool which attempts to trace all failure modes of a complex system back to logical combinations – simply AND and OR relationships – of component failures. If good data is available on the failure rates of all critical components, then FTA can generate the expected failure rate of the overall system.

To apply this technique to IT security, we might produce a tree that portrays the cause-and-effect relationships between attack vectors and system failure. The application of countermeasures would be expected to prune branches of the tree, so that the overall effect with and without treatment could be compared.

Importantly, orthodox FTA is based on the twin assumptions that (1) components fail randomly according to well characterised statistics, and (2) at the lowest level of the tree, component failures are independent of one another. Yet in software and therefore in IT security, failures are not random, but rather are due to systematic design error. Further, it is in the nature of most software that the failure of one line of code can indeed affect other parts of the program. Therefore we believe caution is needed in applying FTA and related reliability engineering techniques to IT security. It was beyond the scope of the present study to explore these issues in more depth.

The fact that IT security incidents are often the result of deliberate actions rather than bottom up component failures is another complication worthy of further study.

3.6 Limitations of existing approaches

The principal weakness of all the aforementioned analyses is that they do not readily uncover the quantified cost/benefit of individual security countermeasures. SAEM and, to a lesser extent, Cost Effectiveness Analysis can help prioritise the *relative* cost/benefit of selected countermeasures, but they do not provide any baseline financial data.

The current best known quantitative approach – the calculation of Annualised Loss Expectancy – rolls up the contributions of all countermeasures into a single "effectiveness" figure. However, no experimental methods for measuring rolled-up effectiveness have been found.

The ALE model is also flawed in that it assumes that all security breaches carry the same cost implications. If the possible annual cost of security failures is for instance $10,000,000 and our security system is thought to be 85% effective, it does not necessarily follow that the security system will save us $8,500,000 p.a. If a particularly expensive type of breach falls into the 15% of incidents against which the system is ineffective, then the ALE result will be overly optimistic.

3.7 A new ROSI model

A hybrid ROSI model, combining Annualised Loss Expectancy and Australian-standard Threat and Risk Assessment (TRA), is proposed, and recommended for four reasons:

1. the proposed model is financially quantitative
2. it separates out the contributions made to the overall cost-benefit by different security countermeasures,
3. it makes use of a widely familiar security tool, making it easy to grasp with minimal new training, and
4. the model is readily extendible to provide statistical modelling of the spread of security costs given the variable nature of likelihood and impact of real life security threats.

The proposed ROSI model augments the TRA table with quantified likelihood and severity estimates, to produce a "bottom up" calculation of expected annual losses with and without treatment by security countermeasures.

4. Recap: Information Security Risk Management

In this section we review some pertinent points in conventional information security risk management.

NOTE: It is assumed that the reader is broadly familiar with information security. If this is not the case, Office of ICT Guidelines should be consulted, such as [8] and others at www.oit.nsw.gov.au/content/2.3.Guidelines.asp.

4.1 Consequences to government of security incidents

In general, security incidents have both direct costs, related to the effort required to restore the system, and indirect costs related to the interim loss of service. In the private sector it is relatively straightforward to count the indirect cost of security incidents. As noted above, a number of financial costs can be readily calculated from lost revenue streams, disenfranchised customers, damaged reputation, and so on. Many of these consequences do not at face value apply to government agencies, either because they represent a monopoly which is invulnerable to customer churn, or because they are delivering services with no direct monetary value. The Office of ICT describes a range of such services, including the community's expectation for improved service (in terms of range, geographic access, tailoring and equity), and enhanced government efficiency [8].

Obviously most government services have some sort of value, even if it is not expressed directly in dollars. An attempt might be made to convert community benefits into dollar figures, but we believe this will always be a controversial exercise, which runs the risk of distracting the ROI exercise. A simpler approach to indirect costs is to look at opportunity cost, as in section 5.3.

4.2 Threat & Risk Assessment methods

In Australia, there are two main standards for Threat & Risk Assessment (TRA). The Defence Signals Directorate's ACSI 33 is commonly used at Commonwealth level, while the Australian risk management standard AS 4360 is perhaps better known at state level and in the private sector. Wide ranging general advice is also available in Standards Australia's Handbook HB231 [4].

For our purposes, the two TRA standards are very similar. They both provide calibrated look-up tables based on "ordinal" (i.e. ordered but non-numeric) values of the likelihood of a given incident occurring and its severity if it does occur. In essence, only the number of grades of likelihood and of severity vary between ACSI 33 and AS4360.

NOTE: The spreadsheet tools in this Guide are based on ACSI 33 but can be readily changed to reflect AS4360.

4.3 Types of Countermeasures

A security countermeasure can have one or two effects on a threat: it can reduce the likelihood of the threat manifesting as an incident, and/or it can reduce the severity of the incident should it actually occur. We can describe countermeasures as ***preventative*** if they reduce likelihood, and as ***curative*** if they reduce severity. These labels are obviously overlapping, as illustrated below.

Figure 1: Schematic representation of countermeasure types

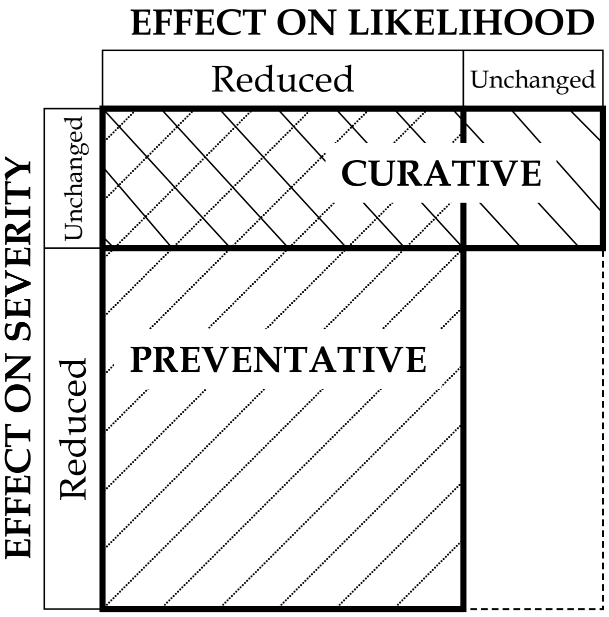

As suggested by the diagram's proportions, countermeasures tend to be more preventative than curative. The following lists provide examples of each type.

Preventative	Curative	Both
− Most Standards, Procedures and Guidelines − Audits, inspections, drills − Firewalls − Intrusion Detection − Virus/Content scanning − Encryption − Data classification	− Redundant systems − Backup regimes	− BCP/DR P − Training

5. Quantifying likelihood and costs of incidents

This section presents frameworks for transforming qualitative descriptions of likelihood and severity into quantitative measures, and for combining these to produce cost-benefit estimates in dollars. The basis for these frameworks is the likelihood and severity definitions from the ACSI 33 Risk Assessment method [1]. The tables in this section are based upon that standard; parts of the tables shaded grey are quoted directly from the standard.

5.1 Likelihood of security incidents

The likelihood of a security incident can be described over seven levels, from "Negligible" through to "Extreme". These levels are quasi-quantitatively defined in ACSI 33 as quoted in the following table. We can convert the quasi-quantitative values to notional annual frequencies as per the right hand column.[1]

Table 1: Likelihood grade transformed to frequency

Likelihood	Description from ACSI 33	p.a.
Negligible	Unlikely to occur	0.05[2]
Very Low	Likely to occur two/three times every five years	0.5
Low	Likely to occur once every year or less	1.0
Medium	Likely to occur once every six months or less	2.0
High	Likely to occur once per month or less	12.0
Very High	Likely to occur multiple times per month or less	50.0
Extreme	Likely to occur multiple times per day	500.0

[1] Please note that some of the notional values in the right hand column are different in this version of the ROSI Guide compared with the first version. We have chosen some slightly different representative frequencies, in order to achieve better consistency within the ranges of frequencies introduced later in the statistical ROSI module.

[2] We assign an arbitrary value for annual frequency corresponding to the "Negligible" likelihood. On advice from OICT, we assigned a rate of once per 20 years. An alternative view is that the frequency could be set to precisely zero. The ACSI 33 TRA method equates Risk (i.e. Likelihood X Severity) to "Negligible" (i.e. do nothing) for all cases where Likelihood is negligible, even if the Severity is Grave, the worst case. We can interpret this to mean that "negligible" likelihood threats are truly not expected ever to occur, and their annual frequency would therefore be zero. In any event, experiment shows that the bottom line result from the ROSI tool is not much affected by setting the frequency for Negligible likelihood to either 0.05 or 0.00.

5.2 Direct costs of security incidents

The severity of a security incident is categorised by ACSI 33 [1] over six levels, ranging from "Insignificant" through to "Grave". These levels are described in the following table (the descriptions are quoted from ACSI 33[1], with emphasis added). Beneath each standard description is a generic interpretation of the impact and the cost to repair. The corresponding order-of-magnitude direct cost is shown in the right hand column.

Indirect costs are handled separately in the model; see section 5.3.

NOTE: that these order-of-magnitude cost reckoners are only indicative, and are based on generic approximations of impact of incidents. The generic cost estimates do not represent NSW Government wide policy. Agencies should assess the estimates provided in this framework, and adjust them where they think it appropriate; see section 6.3 concerning parameters below.

Table 2: Severity grade transformed to direct cost

Severity	**Description /** ***Interpretation***	**Cost**
Insignificant	*Will have almost no impact if threat is realised.*	
	No extra financial impact at all.	$0
Minor	*Will have some minor effect on the asset value.* ***Will not require any extra effort to repair or reconfigure.***	
	No financial impact from incident itself but let us assume that a few hours effort may be required to confirm the nature of what has happened .	$1,000
Significant	*Will result in some tangible harm, albeit only small and perhaps only noted by a few individuals or agencies.* ***Will require some expenditure of resources to repair*** *(e.g. "political embarrassment").*	
	The approximate cost of a press release to redress "political embarrassment".	$10,000
Damaging	*May cause damage to the reputation of management, and/or notable loss of confidence in the system's resources or services.* ***Will require expenditure of significant resources to repair.***	
	Corresponding to the cost of several weeks consultancy to restore resources and services, recover data from backups, conduct investigations and so on.	$100,000
Serious	*May cause extended system outage, and/or loss of connected customers or business confidence.* ***May result in compromise of large amounts of Government information or services.***	
	Corresponding to the cost to totally recover a large set of data, including reverting to original paper files, client interviews and other labor intensive methods.	$1,000,000
Grave	*May cause system to be permanently closed, and/or be subsumed by another (secure) environment.* ***May result in complete compromise of Government agencies.***	
	Corresponding to a typical agency's annual budget.	$10,000,000

5.3 Indirect (opportunity) cost of security incidents

The indirect costs to an agency arising from security breaches – such as the loss of service while an incident is rectified – can be regarded as *opportunity costs*. There are several categories of opportunity costs.

5.3.1 Financial costs

Government agencies tend not to irretrievably lose direct revenue when an electronic service goes down, because payers have no other place to go. For example, if people are using a roads authority website to pay their registration renewals, the monetary effect of an outage is only to delay payments for as long as the site is out of action. In the interests of a simple, conservative and uncontroversial ROSI model, we recommend for the most part ignoring financial opportunity costs.

5.3.2 Interim overhead costs

One type of opportunity cost is probably common to all government agencies, related to the interim overheads of reverting to less efficient manual processing methods while electronic service delivery is unavailable. In most cases, agencies implement ESD at least partly in order to improve efficiency. If a system goes down due to a security breach, then we can expect agency operating costs to temporarily increase, to the same or even higher levels as experienced before the ESD was introduced (higher levels will apply if temporary staff are required to fill in for the absent computer system; additional costs may be expected from the delays experienced while temporary resources are located and put in place).

Consider for example, employee self service (ESD) systems. These are simply cost-justified on the basis of reduced HR overheads, and elaborate spreadsheets have been developed by NSW Government to calculate these savings. Typical net savings over three years are in the order of a million dollars p.a. The *gross* saving is up to $5,000 a day for a medium sized agency. Gross daily dollar saving from ESD is a better indicator of opportunity cost than the net saving, because it reflects the true cost of the sorts of manual methods which need to be employed during the interim.

5.3.3 Other indirect costs

Individual agencies may be able to identify other opportunity costs specific to their own business, for given types of incident. In particular, agencies conducting sensitive transactions carrying high privacy risk – including human services and justice departments – may reckon high opportunity costs arising from incidents which breach client and staff confidentiality. Such costs might factor in legal proceedings, investigations, compensation, intensive public relations campaigns, and the extra implications of perhaps needing to suspend sensitive online services after a breach, while corrective measures are undertaken – and *seen to be* undertaken.

5.4 Statistical variability of security incidents

The simple ROSI model described above requires a significant number of variables to be estimated and given fixed values. The impact and rate of occurrence of security incidents are notoriously difficult to know in advance, and a single point estimate of the cost-benefit outcome will invariably be wrong. So rather than fix the assumed annual frequency and dollar cost of each grade of security incident as per the tables above, a more powerful cost-benefit method will attempt to model the real world variability in these factors.

So-called "Monte Carlo" statistical methods involve introducing variability into one or more parameters of a complex model, re-running the calculations many times and studying the ranges of resulting outputs. Monte Carlo techniques are the accepted method of studying real world randomness in simulations such as ROSI. They are readily implemented using standard spreadsheets.

While there are a number of implementations available, all Monte Carlo spreadsheets in essence involve applying a random number generator to one or more cells, re-running the calculation a number of times, and compiling the spread of results. Good Monte Carlo software will provide user-friendly means for tailoring the statistical characteristics of the randomisation processes, and for presenting the results in meaningful graphical ways. In preparing this version of the Guide, we evaluated several Monte Carlo applications and selected two pieces of freeware for producing a new prototype ROSI tool.

Before implementing a given Monte Carlo application, the key methodological questions to consider are firstly, Which parameters will we choose to vary, and secondly, What statistical characteristics govern the variability?

Reviewing the parameters discussed above in this section – including likelihood, direct costs and indirect costs – we see that potentially all of them may carry statistically uncertainty. Furthermore, when it comes to calculating the overall cost-benefit of security countermeasures, we might also seek to understand the effect of variability in security expenditures.

However, in the interests of a simple, practicable model, it is sufficient to introduce variability into just two parameters:

1. **The frequency of each grade of security incident** can be bounded by the least likely and the most likely frequency per annum. For example, while Table 1 suggests that a Medium likelihood incident (one that is "likely to occur once every six months or less") has a annual frequency of 2.0, it might be more realistic to assume such events occur over a frequency range on either side of the nominal value – say from 1.5 to 2.5 times p.a.. We can assume that likelihood is *uniformly distributed* between the least likely and the most likely frequencies.[3]

2. **The direct cost of each grade of security incident** can be characterised by three values: the most likely cost, the minimum and the maximum. For example, Table 2 suggests that a *Damaging* incident might typically entail several weeks work to repair, with a nominal cost of $100,000. Agencies might have practical experience of this grade of incident with actual costs ranging (for instance) from a minimum of $50,000 to a maximum $200,000. The most likely, minimum and maximum numbers conveniently define a *triangular distribution* for cost.[4]

These idealised probability distributions – uniform for frequency and triangle for direct cost – are simplifications of how real life security incidents occur over time and how they impact the organisation. In trying to translate from the qualitative, ordinal grades of ACSI 33 (or any other TRA standard) to quantitative, continuous measurements of frequency and cost, we see that certain simplifications have been present in the conceptual model from the outset. These will be discussed further in section 6.4.

[3] Individual security incidents probably occur over time according to a Poisson distribution, and their corresponding frequencies would not be uniformly distributed. It is beyond the scope of the current exercise to apply a more sophisticated statistical analysis, and so we will live with the uniform approximation.

[4] Cost in reality is probably normally distributed with a certain mean and standard distribution, but this is much harder for non-statisticians to conceptualise and estimate. Thus the triangular distribution is preferred for its practicality. It will be seen that sophisticated users of the spreadsheets attached to this Guide can readily substitute alternate distribution functions if desired.

6. The proposed ROSI model

In this chapter we explain how a standard TRA table can be modified to include the calibrated likelihood and cost of each security incident, and automatic calculations made of the expected annual cost of all incidents. These calculations, performed using a spreadsheet, form the basis of the proposed ROSI model. Later in this chapter, we extend the basic model to include Monte Carlo simulation with variable incident likelihood and cost.

6.1 The expected annual cost of security incidents

To calculate the annual cost of security incidents with and without countermeasures, we begin with a standard Threat & Risk Assessment (TRA) for the system in question. The major deliverable of a TRA is an systematic tabulation of (1) all threats to the system's information assets, (2) an estimate of the risk level of each threat, being the product of likelihood and severity, (3) one or more countermeasures proposed to mitigate each risk, and (4) the estimate residual risk level after treatment with the countermeasures.

The basic ROSI tool augments the standard TRA table as follows

- For each threat, the tool incorporates the corresponding annual frequency and the per incident cost, as per Table 1 and Table 2 above.
- For each threat, the tool calculates the expected annual *untreated* cost (being simply the product of the annual frequency and the per incident cost).
- For each proposed countermeasure, the user enters the anticipated upfront cost of implementation, the annual cost of maintaining the countermeasure, and the amortisation period (number of years) over which the upfront cost is to be spread when calculating return.
- For each countermeasure – or collected set of countermeasures – the user enters the residual likelihood and severity anticipated after treatment.
- Finally, for each threat, the tool calculates the expected annual *treated* cost.

The following table illustrates the basic tool in action. It is based on an extract from an actual government Threat & Risk Assessment. Parts of the table shaded grey are taken from the original TRA.

Table 3: TRA with calculated annual cost of incidents, untreated

No.	Asset	Potential Incident (Threat to the Asset)	Likelihood	Severity	Estimated Risk	Annual rate of occurrence	Direct Cost per incident	Opportunity Cost per incident	Total UNTREATED Annual Cost
A8	Availability of D-XYZ internet connection	Destruction of key infrastructure e.g. routers, PIX, switches)	Negligible	Serious	Nil	0.05	$ 1,000,000		$ 50,000
A9		Failure of Cooling System	Medium	Significant	Medium	2	$ 10,000		$ 20,000
A10		Mis-configuration of key infrastructure e.g. routers, PIX, switches)	Low	Serious	High	1	$ 1,000,000		$ 1,000,000
A11		Hardware failure of key infrastructure e.g. routers, PIX, switches)	Very Low	Damaging	Low	0.5	$ 100,000		$ 50,000
A12		Incorrect building patching	Low	Significant	Medium	1	$ 10,000		$ 10,000
A13		Denial of service attack on carrier or provider network infrastructure	Very Low	Significant	Low	0.5	$ 10,000		$ 5,000
A14		DNS hardware failure	Negligible	Damaging	Nil	0.05	$ 100,000		$ 5,000
A15	Availability of D-XYZ internet email	Denial of service attack on email system	High	Damaging	High	12	$ 100,000		$ 1,200,000

Table 4: TRA (cont.) with calculated annual cost of incidents, treated

No.	Counter Measures	Upfront Cost per Counter-measure	Recurring Cost per Counter-measure	Residual likelihood	Residual severity	Total TREATED Annual Cost	Saving Per Threat	Notes on mitigations
A8	Business Continuity Plan (1)	$ 50,000	$ 20,000	Negligible	Minor	$ 50	$ 49,950	
	Spare parts (4)	$ 50,000	$ 10,000					
	Service level agreements (5)	$ 0	$ 0					
	Physical security (access control procedures and controls for computer room) (6)	$ 10,000	$ 10,000					
A9	Environmental controls for computer room (2)	$ 30,000	$ 5,000	Very low	Minor	$ 500	$ 19,500	Harm reduced to Minor by BCP; Likelihood to Very Low by Environ controls
	Business Continuity Plan (1)	*Counted*	*Counted*					
	Service level agreements (5)	*Counted*	*Counted*					
A10	Configuration management system (8)	$ 70,000	$ 10,000	Negligible	Serious	$ 50,000	$ 950,000	Likelihood reduced to Negligible by Config Mgt
	Change control procedures (15)	$ 30,000	$ 5,000					
A11	Business Continuity Plan (1)	*Counted*	*Counted*	Very low	Minor	$ 500	$ 49,500	Won't affect the likelihood of an event, but reduces harm by better recovery
	Spare parts (4)	*Counted*	*Counted*					
	Service level agreements (5)	*Counted*	*Counted*					
A12	Standards for cabling including labelling and coding (9)	$ 10,000	$ 0	Very low	Significant	$ 5,000	$ 5,000	
	Physical security (6)	*Counted*	*Counted*					
A13	Large capacity network connection (10)	$ 10,000	$ 10,000	Very low	Minor	$ 500	$ 4,500	Redundancy means minor effect on failover
	Redundant Internet connection (7)	$ 10,000	$ 10,000					
A14	Replication of DNS server (11)	$ 10,000	$ 0	Negligible	Minor	$ 50	$ 4,950	
A15	Network based Intrusion Detection System (NIDS) (12)	$ 70,000	$ 20,000	Low	Significant	$ 10,000	$ 1,190,000	No amelioration of degree of harm
	Use DSD evaluated products (13)	$ 20,000	$ 5,000					
	Deny all unless explicitly allowed firewall rules (14)	$ 0	$ 0					

6.2 The basic spreadsheet tool

Attached to this Guide are two Excel spreadsheets which implement the tables and methods described above. The first spreadsheet – referred to as the "basic tool" – is a simplified version, calibrated for the most likely expected likelihood and cost for each type of security incident. The second spreadsheet includes a statistical module, for calculating the possible spread of cost-benefit given the variability of occurrence and cost of incidents.

The basic spreadsheet has two work sheets as follows:

- **Quantified Risk Grades** contains lookup tables corresponding to Table 1 and Table 2 above. The sheet also contains, for explanatory purposes only, a copy of the ACSI 33 likelihood-times-severity table, and a table of annual incident costs at each point in the risk table.

  Note that we cap the maximum annual cost of a given threat at the cost of a single Grave incident. This is on the basis that if a Grave incident were in fact to strike an agency more than once in a given period, the first occurrence, being catastrophic, renders all subsequent events academic.

- **TRA table with Calculated Costs** contains an example TRA table augmented with the annual cost calculations described above. The table is colour-coded to aid users wishing to modify any of the assumptions or parameters. See the spreadsheet for details.

The basic spreadsheet also calculates the following summary estimates:

- the annual cost of security incidents, untreated
- the residual annual cost of security incidents, after treatment
- the gross annual savings attributable to security treatments
- amortised countermeasure upfront cost
- the annual cost of countermeasures, comprising recurring costs and amortised set-up cost
- the net annual savings.

6.3 Parameters of the basic tool

The spreadsheet allows users to readily modify several parameters of the model (indicated by purple cells) as follows:

- the annual frequency attributed to each likelihood level (sheet one)
- the direct per-incident cost attributed to each severity level (sheet one)

- the opportunity cost of each identified incident should it occur (sheet two)

> *NOTE: Agencies should carefully consider the indirect costs that result from each type of incident should it occur, according to the implications specific to their own business; see section 5.3 above.*

- the upfront cost of each identified countermeasure (sheet two)
- the recurring cost of each identified countermeasure (sheet two)
- the amortisation period for treating up-front countermeasure cost, in years (sheet two).

Users may also wish to over-ride the spreadsheet's calculated annual costs per incident, in sheet two, case-by-case if so desired.

6.4 The statistical module

In order to reveal the expected spread of cost-benefit given the variability of occurrence and cost of incidents, the second spreadsheet attached to this Guide includes a statistical module featuring Monte Carlo methods.

The statistical module is similar to the basic spreadsheet, also comprising two main worksheets: *Quantified Risk Grades* and *TRA table and Calculated Costs*.[5] But in the statistical module, we have modified the lookup tables to feature the statistical spread of likelihood and cost for each grade of incident, rather than specify constant values. For likelihood, the lookup table has the minimum and maximum annual frequencies; for cost, the table has the most likely, minimum and maximum expected costs.

The current version of the tool uses two freeware Excel add-ins: *Quadrant*[6] and the Trial Version of *XLSim*[7]. Quadrant was chosen for its easy-to-use graphing functions, while XLSim offers convenient additional functions for specifying random distributions of various types, over and above Excel's standard features.

> *NOTE: In order for the ROSI statistical module to run properly, it is important that the Quadrant add-in (approx. 128kB) and XLSim executable (approx. 2.5MB) be downloaded and installed from the respective websites, as noted above. Note also that XLSim.exe is to be run, not Excel itself. When XLSim runs, it launches Excel automatically, together with the extended statistical functions, after which the program looks and feels in all respects like Excel.*

[5] Note too that the Monte Carlo add-ins automatically create additional worksheets in addition to these two, when the simulations are run.

[6] See Quadrant home page at http://home.btconnect.com/Farnham/home.htm.

[7] From AnylCorp http://analycorp.com.

6.4.1 The statistical outputs of interest

Once familiar with the Monte Carlo tools implemented in the prototype ROSI statistical module, you will see that a wide variety of "outputs" of the spreadsheet can be studied in terms of they are affected by random variations in security event occurrence and cost. The default spreadsheet attached to this Guide selects the following three particular outputs and graphs their statistical spread:

- **Annual untreated cost of incidents**: the sum of the cost of individual incidents expected over one year without perimeter security countermeasures

- **Annual treated cost of incidents**: the sum of the cost of the same incidents over one year after perimeter security has been deployed as per the TRA (not counting the cost of the security system), and

- **Annual nett savings**: the difference between annual untreated and treated costs, less the cost of investment in security.

These three outputs appear in the summary section of the spreadsheet. To have them included in the Quadrant Monte Carlo simulation, we insert special instruction comments into the respective cells; see section 6.4.3.

6.4.2 Notes on choosing the statistical parameters

As mentioned in section 5.4, we have made certain simplifying assumptions about the probability distribution of frequency and direct cost of security incidents; these are modelled as uniform and triangular distributions respectively. ACSI 33 gives us quasi-quantitative descriptions of the expected annual frequency of each grade of security incident, and in Table 1 we converted those descriptions into notional frequencies. When defining the start and end of the (assumed) uniform distributions of frequency for different grades of incident, the simplest approach is to divide up the range from *Negligible* through to *Extreme* into sub-ranges more or less centred on the notional representative values. See Table 5. These values have been loaded into the attached statistical module spreadsheet (but can be readily changed if you desire).

Table 5: Likelihood grade transformed to *probabilistic* frequency

Likelihood	Min. freq p.a.	Max. freq p.a.	Representative frequency p.a.
Negligible	0.0	0.10	0.05
Very Low	0.1	0.8	0.5
Low	0.8	1.5	1.0
Medium	1.5	2.5	2.0
High	2.5	20.0	12.0
Very High	20.0	120.0	50.0
Extreme	120.0	1000.0	500.0

The triangular distributions were chosen to have their most likely costs aligned with the notional costs previously assigned to each severity grade. The minimum and maximum costs were arbitrarily set to be half as bad and twice as bad as the most likely cost respectively. See Table 6.

Table 6: Severity grade transformed to *probabilistic* direct cost

Severity	Min. Cost	Most Likely Cost	Max. Cost
Insignificant	$0	$0	$0
Minor	$0	$1,000	$2,000
Significant	$5,000	$10,000	$20,000
Damaging	$50,000	$100,000	$200,000
Serious	$500,000	$1,000,000	$2,000,000
Grave	$5,000,000	$10,000,000	$20,000,000

The triangular cost distributions are therefore skewed by default, exhibiting longer tails towards the maximum cost than the minimum cost. We regard this shape as a conservative assumption, and as we shall see, leads to higher cost forecasts from the statistical module compared with the corresponding simple spreadsheet. The default spread from minimum to maximum cost is more arbitrary. While we chose a spread of 50% to 200% of the most likely value, agencies may wish to insert their own statistical parameters, especially if data from real security experiences are available.

6.4.3 Brief instructions for XLSim and Quadrant

To use the ROSI statistical module, first run XLSim (from the *Start→Programs* or *Start→All Programs* menu in Windows 2000 or Windows XP respectively). XLSim will automatically launch Excel, as if you had run Excel from the outset. Then, when running the statistical module for the first time, add in Quadrant from the Excel *Tools→Add-Ins* … menu.[8] Quadrant creates a simple two button toolbar.

The randomly distributed values of likelihood and cost are calculated using the XLSim functions *gen_Uniform()* and *gen_Triang()* respectively. The effect of these functions can be seen directly by opening the *Risk Ratings & Cost Parameters* worksheet and hitting the F9 (Calculate) key a few times. The return values of the functions highlighted in yellow change randomly on each re-calculation, over the ranges specified.

[8] This assumes that the Quadrant add-in file has previously been downloaded from http://home.btconnect.com/Farnham/home.htm and installed into the default Windows folder.

The Monte Carlo simulation is controlled as follows. Quadrant must be instructed as to which spreadsheet cells are to be included in the simulation. In our case, we are interested primarily in the variability of untreated and treated (residual) costs of security incidents. Recall that these outputs are provided in the summary section at the bottom of the *TRA table and Calculated Costs* worksheet. Quadrant makes uses of special key words inserted into Excel comments to tell which cells are to be included; by default, if the first word in a cell's comment is "output" then that cell will be included. Additional lines of comments are picked up by Quadrant and used to form the title of its output graphs.

A dialog box invoked from the Quadrant toolbar allows the number of iterations to be specified (up to a maximum of 250) as well as properties of the output histograms such as the bin size. Full documentation for Quadrant is found at http://home.btconnect.com/Farnham/home.htm.

6.4.4 Example outputs

The diagrams below show the outputs of one Quadrant run over 250 iterations of the prototype statistical module, with settings of 15 histogram bins (the default) and "neat" bin ranges.

NOTE: Each time a simulation is run, the output graphs will change, even if all parameters remain constant, due to the randomisation of incident occurrence rate and cost. This is normal. See also section 6.5 below.

Figure 2: Example outputs from the Monte Carlo prototype

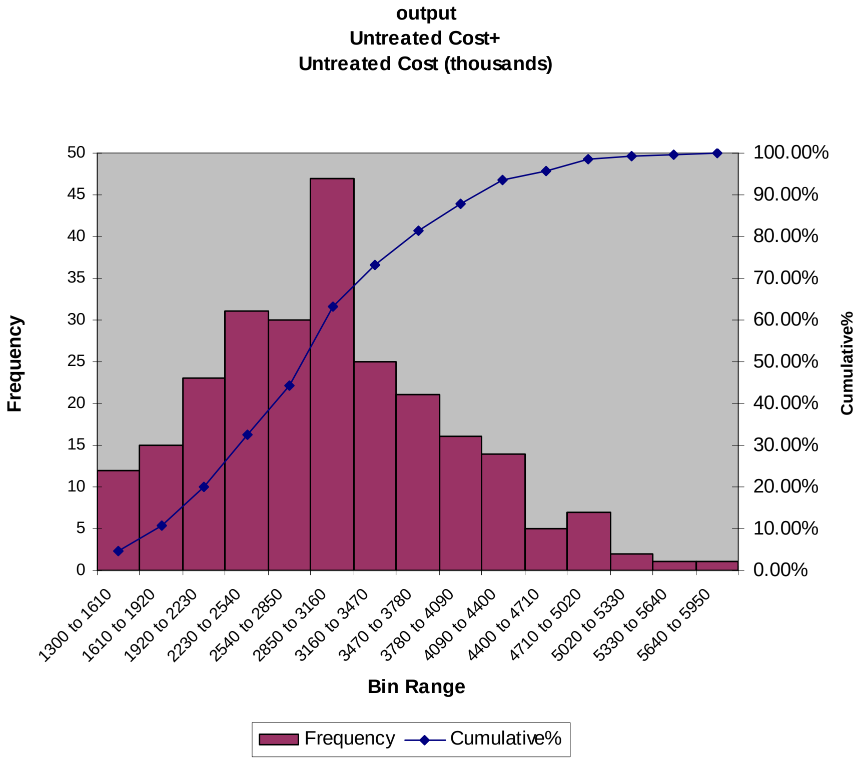

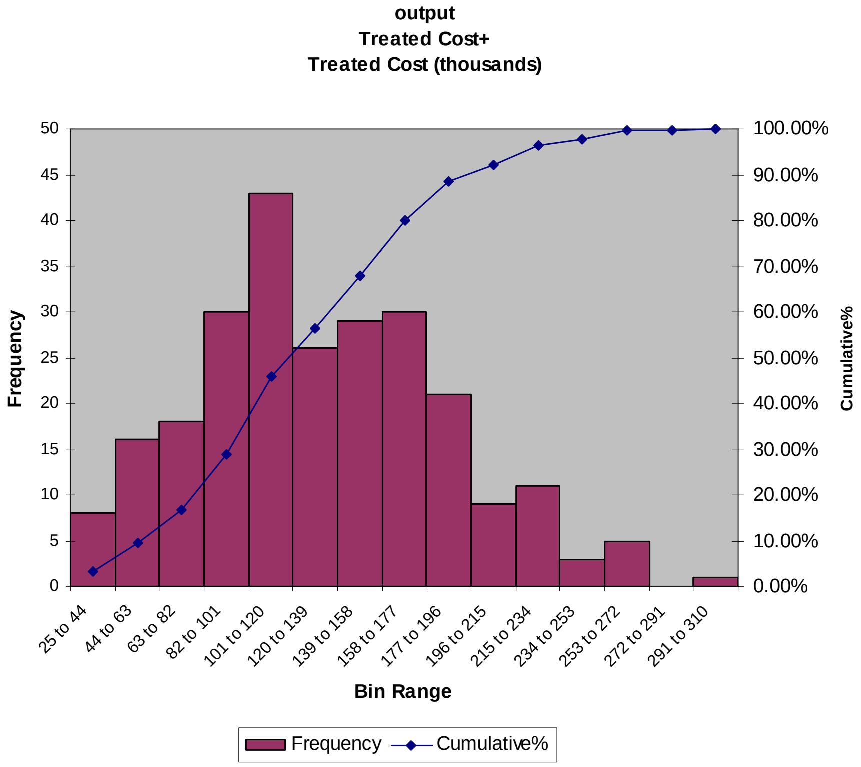
output
Treated Cost+
Treated Cost (thousands)
Frequency
Cumulative%
0
5
10
15
20
25
30
35
40
45
50
0.00%
10.00%
20.00%
30.00%
40.00%
50.00%
60.00%
70.00%
80.00%
90.00%
100.00%
25 to 44
44 to 63
63 to 82
82 to 101
101 to 120
120 to 139
139 to 158
158 to 177
177 to 196
196 to 215
215 to 234
234 to 253
253 to 272
272 to 291
291 to 310
Bin Range
Frequency
Cumulative%

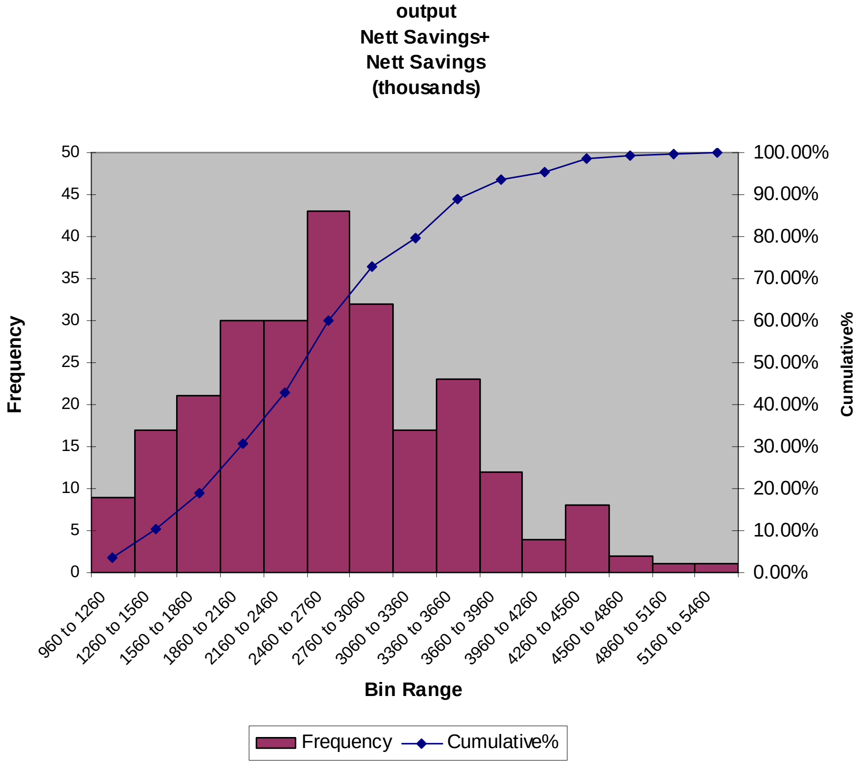
output
Nett Savings+
Nett Savings
(thousands)
Frequency
Cumulative%
0
5
10
15
20
25
30
35
40
45
50
0.00%
10.00%
20.00%
30.00%
40.00%
50.00%
60.00%
70.00%
80.00%
90.00%
100.00%
960 to 1260
1260 to 1560
1560 to 1860
1860 to 2160
2160 to 2460
2460 to 2760
2760 to 3060
3060 to 3360
3360 to 3660
3660 to 3960
3960 to 4260
4260 to 4560
4560 to 4860
4860 to 5160
5160 to 5460
Bin Range
Frequency
Cumulative%

6.5 Interpreting the statistical results

The statistical ROSI module generates two types of graphical output: *histograms* and *cumulative plots*. The graphs can be interpreted in a number of ways, and allow you to extract security cost-benefit forecasts with different degrees of confidence.

A histogram plots the probability that a certain result will fall into a given range of possibilities. Most probability distributions broadly take the shape of a hill, indicating a most likely value at the peak of the hill, and steadily decreasing probabilities on either side. [9] The simulation run in Figure 2 shows that the most likely annual cost of security incidents in our example is somewhere between $2,850,000 and $3,160,000 without mitigation, and between $101,000 and $120,000 with mitigation.[10]

It might be tempting to interpret these figures as indicating that the saving due to security mitigation will be at least $2,730,000 (i.e. $2,850,000 - $120,000). Yet this inference falls into the same trap as the fixed value predictions of the simple ROSI model, because single cost-benefit results are more likely to be wrong than right. Superior ways to look at the range of potential cost-benefit results include – graphically – the forecast distribution of *nett savings*, as per the third plot in Figure 2, and – numerically – the use of *confidence limits*, as will now be discussed.

Firstly, consider cumulative plots. These show what proportion of possible results fall below a given value. For instance, while the most likely untreated cost according to the above histogram is between $2,850,000 and $3,160,000, the cumulative plot shows that 60% of values fall at or below this value range. Therefore a sizable proportion of outcomes – namely 40% – will in fact exceed $3,160,000.

If we wish to make transparently conservative forecasts, we can first pick a confidence limit that suits our purposes, and then read off the cost value corresponding to that confidence limit on the cumulative plot. For instance, let us assume that a 90% confidence limit is sufficiently cautious. From the cumulative plots in Figure 2 we see that it is 90% likely that the untreated cost will be no higher than $4,090,000, while it is 90% likely that the treated cost will be no higher than $196,000.

[9] The histograms plotted by the Quadrant add-in used in the current ROSI prototype are not as smooth as one might expect, due to the limited number of iterations supported in the freeware. A Monte Carlo simulation over thousands of iterations (rather than merely 250) would show smoother probability distributions.

[10] The fixed value prediction from the simple ROSI tool of untreated cost in this example is $2,440,000, which falls outside the statistical model's most likely bin of $2,850,000 to $3,160,000. When the Monte Carlo simulation is repeated many times, we see that the fixed value prediction is frequently well below the most likely bin range. We should not be surprised by this result, for as discussed in section 6.4.2, our default choices for the cost distributions of security incidents is skewed on the high side and will therefore on average lead to higher untreated cost predictions. On the other hand, the simple tool's fixed value prediction for treated cost is $116,600 which does fall inside the most likely bin.

In ROI "debates", critics sometimes challenge a forecast annual cost as being arguably too high, and therefore artificially bolstering a security advocate's case for investment. You can answer such criticism with figures from inverting the cumulative plots. Looking at the 10% levels of the cumulative plots above, the corresponding untreated cost is near $1,610,000 and treated cost is near $63,000. We can therefore state that it is 90% likely (i.e. 100% less 10%) that the untreated cost will be *at least* $1,610,000 and that the treated cost will be *at least* $63,000. Using confidence limits in this manner provides conservative and transparently qualified predictions of the possible return on security investment.

Remember that each time the simulation is run, the outputs will change. This is to be expected, thanks to the randomisation built into the model. The basic shape of the distributions should not change significantly between runs. If a more powerful Monte Carlo package is adopted, such as the production version of XLSim, then with higher numbers of iterations, the variation between runs will be reduced. With greater iterations, the bin sizes can also be reduced to produce smoother curves with fewer artefacts.

7. Limitations of the tool

7.1 Limited iterations in the Monte Carlo freeware

The freeware Monte Carlo add-ins evaluated and selected for the prototype calculator limit the number of iterations performed per run to 200-250. This makes the results rather coarse in resolution. For now, there are sufficient iterations to show the broad shape of the cost characteristics. However, in routine use, software allowing at least 1,000 iterations is recommended, implying either some modest investment in a commercial Monte Carlo application, or else customisation of the Excel add-ins.

7.2 Restricted sources of randomness

As discussed, the statistical module is relatively simple, so as not to overly trade off usability against complexity. The two major simplifications at present are:

1. *All types of incident in each grade of severity are assumed in the lookup tables to have the same statistically variability.* In reality, different threats with the same types of root cause – for example, "user error" – can be expected to exhibit the same variability, even if they result in different levels of severity. The current model applies the same statistical measures to all threats of the same severity, regardless of their root causes. We would suggest that the next update to the ROSI tool investigate whether finer grain randomness can be injected into the model without sacrificing usability.

2. *No randomness is explicitly allowed for in the cost and effectiveness of security countermeasures.* Arguably a more sophisticated ROSI model would allow for variability to be expressed in the countermeasures too. Note however that when we inject statistical variability into the *generic* likelihood and cost of all security incidents, then we automatically capture a degree of uncertainty in the effectiveness of security countermeasures. This is because the TRA table rates the likelihood and cost of each identified security threat after countermeasures have been implemented, and the statistically uncertainty injected into the lookup tables will carry over to the treated risks.

7.3 Difficult to separate effects of countermeasures

It is generally accepted that the calculation of residual risk after a set of countermeasures has been selected, is based on the entire set, and not on individual countermeasures. Separate calculations of reduced likelihood or severity are difficult, and indeed can be meaningless.

To help understand this difficulty, consider the common threat of mis-configuring say a firewall or operating system, such that it is left with security holes. To address such a threat, we would almost always:

- make sure to document the proper configuration
- make sure to implement change control procedures, and
- make sure to protect the system against unauthorised access.
- These countermeasures respectively:
- reduce mistakes being made when the system is built
- reduce mistakes being made over the lifetime of the system (e.g. during maintenance) and
- reduce malicious tampering.

Yet none of these measures on their own will usually reduce the overall risk to satisfactory levels, no matter how good each of them is. Thus, to meaningfully mitigate the risk arising from mis-configuring a computer system, we need to implement the suite of countermeasures. And by the same token, it can me meaningless to try to account for the cost and benefit of each countermeasure separately.

NOTE: Nevertheless, the spreadsheet has been constructed to list all countermeasures implemented against each identified threat. If the user wishes, they can assign residual likelihood and severity against each countermeasure separately, and run a what-if scenario to deduce separate cost-benefits. In this way, the spreadsheet could also be used to produce Cost Effectiveness Ratios as described in [14].

7.4 Difficult to predict how countermeasures affect severity

It is worth noting that in orthodox TRAs it is actually unusual to separately assess the likelihood and severity of each threat after treatment is applied. Typically, a TRA will show the rolled-up residual risk as a single figure, determined largely by the judgement of the security analyst. To calculate residual cost after treatment of course it is necessary for us here to assess likelihood and severity separately.

The reason this separation is not usually attempted has to do with the difficulty of reliably forecasting consequences of a threat once a countermeasure is implemented. Some countermeasures so radically alter the risk profile that it is impossible to predict the type of threat that might still penetrate the defences.

Consider for example firewalls, which clearly reduce the likelihood of penetration. But the consequences of such threats that can in future breach the firewall are very hard to assess. Almost by definition, these events are unanticipated and therefore unknown in nature and severity.

Therefore it is recommended that conservative figures be used for severity estimates after treatment.[11]

7.5 Security incidents are not necessarily independent

A major simplifying assumption featured in all security cost-benefit models researched in preparing this Guide is that security incidents are statistically independent events. Yet clearly this is not always the case. For instance, software bugs – which represent the root cause of so many vulnerabilities – are notoriously correlated due to idiosyncrasies of software design processes. Moreover, the discovery of one bug can lead to a cascade of related vulnerabilities becoming known and exploited in a decidedly non-random way

Yet few if any mature statistical methods appear to be available to accommodate the real world complexity in how security incidents correlate non-randomly with one another. There is little alternative therefore but to live with this acknowledged simplification in the model, and to ensure that enough "fat" is built into the conservative estimation of likelihood and consequence grades that weaknesses in the model are not in fact important.

Another interesting "non-linearity" in the model concerns the cumulative impact of multiple security breaches. We assume in the model that the impact of security incidents simply adds up, so that the annual cost of multiple breaches is the sum of the individual respective costs. In real life however the total impact can exceed the sum of the parts, if for example the user community suffers such a loss of confidence that they choose to abandon the system. To this extent, the spreadsheet could be extended with an exponential factor of some sort to show that cumulative impact worsens faster than the arithmetic sum of the costs. We would caution however that quantitative modelling of these effects remains immature.

[11] In fact, some security analysts take the conservative view that security countermeasures can in principle have *no effect* on severity, but can *only* reduce the likelihood of incidents. Needless to say, we do not subscribe to this view.

8. Bibliography

8.1 Cost-benefit references

[1]. ***Risk Management Handbook 3, Australian Communications – Electronic Security Instruction 33*** V1.0 (ACSI 33) published by Defence Signals Directorate, 2000.

[2]. ***Economic Appraisal Principles and Procedures*** TPP99-1 published by NSW Treasury March 1999, see www.treasury.nsw.gov.au/pubs/tpp99_1/prin_pro.htm.

[3]. ***Risk Management Standards*** AS/NZS 4360:1999 published by Standards Australia, 1999.

[4]. ***Information Security Risk Management Guidelines*** SAA HB 231:2000 published by Standards Australia, 2000.

[5]. ***Security Attribute Evaluation Method: A Cost-Benefit Approach*** Shawn A. Butler, Computer Science Department, Carnegie Mellon University, 2002. See www2.cs.cmu.edu/~Compose/ftp/SAEM-(Butler)-ICSE_2002.pdf

[6]. ***Finally, a Real Return on Security Spending*** CIO Magazine, 15 February 2002; see www.cio.com/archive/021502/security.html.

[7]. ***Cost-Benefit Analysis for Network Intrusion Detection Systems*** Huaqiang Wei, Deborah Frinke et al. Centre for Secure and Dependable Software, University of Idaho. In Proceedings of the 28th Annual Computer Security Conference October 2001. See http://wwwcsif.cs.ucdavis.edu/~balepin/new_pubs/costbenefit.pdf.

[8]. ***Information Security Guideline Part 1 – Risk Management*** NSW Government Office of Information and Communications Technology, June 2003. See www.oit.nsw.gov.au/pdf/4.4.16.IS1.pdf.

[9]. ***Benefits Realisation Register Guideline*** NSW Government Office of Information and Communications Technology, March 2004. See www.oit.nsw.gov.au/pdf/4.4.3.Benefits.pdf.

8.2 Further reading on cost-benefit

[10]. Computer World ROI Knowledge Centre at www.computerworld.com/managementtopics/roi.

[11]. ROI that never arrives: The Devil is in the Assumptions Cost/Benefit Newsletter August 2003 at http://www.solutionmatrix.com/Newsletter35.html.

[12]. ***Calculated Risk*** Scott Berinato in *CSO Magazine*, December 2002. See www.csoonline.com/read/120902/calculate.html

[13]. Secure Business Quarterly, Special Issue on Return on Security Investment, Quarter 4, 2001. See www.sbq.com/sbq/rosi/index.html

[14]. ***Primer on Cost-Effectiveness Analysis*** published by the American College of Physicians' Effective Clinical Practice, September/October 2000. See www.acponline.org/journals/ecp/sepoct00/primer.htm

[15]. ***A Guide to Security Risk Management for Information Technology Systems*** Published by the Government of Canada Communications Security Establishment, 1996. See www.cse.dnd.ca/en/documents/knowledge_centre/publications/manuals/mg2e.pdf

[16]. ***Executives Need to Know: The Arguments to Include in a Benefits Justification for Increased Cyber Security Spending*** Timothy Braithwaite in Information Systems Security, Auerbach Publications, September/October 2001.

[17]. ***Seeking Security Scorecards*** Chris King, Meta Group (File: 9377), Dec 2001.

[18]. ***Security Metrics Guide for Information Technology Systems*** Special Publication 800-55 US National Institute of Standards and Technology Computer Security Research Centre, 2002. See csrc.nist.gov/publications/nistpubs/800-55/sp800-55.pdf.

8.3 Monte Carlo and other tools

[19]. ***Quadrant: The QUick And Dirty Risk ANalysis Tool*** Home page http://home.btconnect.com/Farnham/home.htm.

[20]. ***AnylCorp*** (supplier of XLSim) Home page http://analycorp.com.

[21]. Sam Savage's Monte Carlo page http://www.stanford.edu/~savage/software.htm

[22]. ***SIMTOOLS and FORMLIST*** Home page http://www.kellogg.nwu.edu/faculty/myerson/ftp/addins.htm