

STRUCTURED RISK ANALYSIS

Neil McEvoy and Andrew Whitcombe

Consult Hyperion, 8 Frederick Sanger Road, Surrey Research Park, Guildford, Surrey,
GU2 7EB, United Kingdom
neil@chyp.com, andrew@chyp.com

Abstract. Decisions concerning the security of information are frequently taken on a different basis to other business decisions. Necessary investments may not take place on the grounds that the returns are not easily quantifiable, when compared with other investments, or in the blithe hope that “it will never happen to us”. Conversely, “fear, uncertainty and doubt” can lead organisations to divert funds towards information security that might be better spent elsewhere.

This paper presents Structured Risk Analysis (SRA), a method to help organisations take rational steps to improve their information security. Using SRA, an organisation can place information security in the context of the business as a whole, to determine the appropriate level of resources to be directed toward improvements. Given this budget, SRA enables the organisation to identify how it should be spent to generate the best possible security ‘return’.

Whilst other risk analysis methods exist, we believe SRA to be more closely integrated with other business processes, systems development and operations, giving a much firmer context for the analysis and more confidence in the results. It is non-proprietary, does not require specialist software tools and is inherently ‘tuneable’ in the level of detail that is applied: all of which means that SRA is particularly cost-effective.

We believe SRA gives business owners the information they need to make truly informed decisions about the security, viability and future direction of their Information Systems.

1 Background

Information security is progressively becoming headline news, particularly as ordinary people start to rely on electronic means to carry out routine transactions such as making payments from a bank account or filing a tax return. Barely a week goes by without scares about the latest virus, or a report of an embarrassing lapse by a large organisation, hitting the front page and radio and television new bulletins. In the week this paper was drafted, the biggest noise was about the UK Inland Revenue suspending their online self-assessment service following reports that users had seen snippets of information relating to other taxpayers [1].

All of this understandably has the tendency to leave the lay user confused or, worse, frightened. What is more surprising is that so many large organisations leave themselves wide open to appear in tomorrow’s headlines. That is not to say that money and effort is not being spent. For example, in the wake of September 11th,

many airports and airlines are deploying biometric technology to screen travellers—despite informed academic comment that many systems are incapable of delivering the claimed benefits.

A number of organisational cultures can lead to irrational stances towards information security. The most negligent would be the assumption that “it could never happen to us”. A more prevalent response to the issue is to appoint a person or team to be responsible for information security, but without the breadth of experience, or the ‘clout’ to identify and solve the most serious security issues and to successfully deploy and promote the solutions. For example, a team with a strong software background may understand the security flaws with the server operating systems and database management systems, but may under- or over-estimate the impact such flaws might have on the business. Conversely, a business person charged with the same responsibility may have an acute awareness of potential damage to the organisation, without understanding how a hacker might be able to penetrate the systems to realise his worst fears.

The successful management of information security risks has many components. Policies must be developed and applied that define who can have what kind of access to which information and infrastructure components. Procedures must define the controls around such access. Technical means must be deployed to enforce policies and procedures. Active monitoring must detect serious or systematic attempts to circumvent the policies and procedures.

None of the above should be defined in a vacuum. Otherwise, (for example) overly bureaucratic procedures might be defined to protect access to equipment which is not critical to maintaining information security; or sensitive traffic on a wireless network may be available, un-encrypted, on the street outside the building. The key to taking optimal decisions in all these areas is to understand the risks to which the organisation is exposed by doing nothing, and the degree to which these could be mitigated by the deployment of particular countermeasures.

Of course, recognition of the value of risk analysis in this area is not new, and methods to conduct such analyses have been proposed. However, uptake of these methods has been slow, leading to *ad hoc* analyses, where they take place at all. We find that most methods for risk analyses have one or more of the following deficiencies:

- They are not rooted in the wider disciplines of information analysis and system development, which leads to doubt about the completeness of any analysis.
- They do not provide a means whereby information security risks can be compared with other business risks (or indeed business opportunities), which would allow business people to take rational decisions about the application of funds.
- They fail to distinguish clearly between *threats* to the business and the *vulnerabilities* of systems, leading to a muddled analysis
- They are either too superficial (nothing which is not obvious is revealed) or too detailed (can’t see the wood for the trees)
- They are ‘static’ in terms of the foreseen vulnerabilities or countermeasures, whereas in reality information security is a fast-moving discipline
- They are inflexible or prescriptive
- They are proprietary, or rely on specific software packages.

2 Requirements for SRA

Recognising these deficiencies, we decided to define our own method, Structured Risk Analysis (SRA) that would address these issues.

What then were the requirements for SRA?

- *Business context.* The method must enable business people to take business decisions in the real world, where there are other calls on resources, and not the 'ideal' world of limitless resources to improve information security.
- *Technical grounding.* The method must be rooted in 'best practice' for systems analysis and design, with a defined place in the systems development lifecycle, whilst not dependent on any particular techniques, methods or tools in this area.
- *Separation of concerns.* The method must clearly distinguish between business and technical concerns, and bring both together leading to a robust analysis.
- *Support for quantitative analysis.* While not appropriate in all circumstances, the method must be capable of quantifying exposure to risk, in order to determine an overall information security budget and the optimal allocation of that budget.
- *'Tuneable' analysis.* It must be possible to control the level of detail of the analysis, appropriate to particular corporate circumstances.
- *Evolution.* It must be possible for the method to evolve, for example to recognise changes in methods of attack and countermeasures
- *Maintainability.* It must be possible to generate a 'living' risk analysis model, which will allow changing circumstances to be tracked without undue expenditure.
- *Openness.* It must be possible to use the method without payment of licence fees or purchase of specific software.

3 How Does it Work?

3.1 Overview

Figure 1 below outlines the basic steps undertaken in order to carry out a successful Structured Risk Analysis.




Fig. 1. The basic steps undertaken during a Structured Risk Analysis

SRA builds on existing Systems Analysis methods and best practice to ensure that the service and systems under consideration are represented clearly, in a way which is understandable to stakeholders and hopefully in a manner that is consistent with the rest of the organisation. This enables the results of the SRA to be viewed in an overall business context, allowing security decisions to be taken as part of a complete picture of organisational structure, spend and development. These representations, or service and system ‘models’ are then used to identify the threats and vulnerabilities that exist within the system.

All the threats and vulnerabilities identified are then cross-referenced to ascertain whether it is possible that a particular threat might be realised by a particular vulnerability. If this is possible, then we have discovered a real risk.

A number of factors are considered in order to help identify the extent to which such risks expose the organisation to damage, leading to the selection of appropriate countermeasures.

3.2 Understanding the Environment

In order to provide a complete picture it is important to understand both the service being offered and the system via which this service is provided. The representation of this understanding, through standard systems analysis modelling techniques, provide the model service and model system that form the foundation of the SRA process.

The ‘Model Service’ process in Figure 1 identifies the information assets that the service must manage. This is achieved by drawing up a Logical Data Structure (LDS), to identify data entities, the keys that uniquely identify each instance of the entity, the entities attributes, and the relationships between the entities. Ideally, an LDS will have been produced during the system development process in the context of a struc-

tured method such as SSADM [2] or one of the flavours of the ‘Yourdon’ method [3]. We find that since an LDS can be presented as a set of tables, it is particularly suited to review and discussion by a wide variety of stakeholders. Alternatively, an object model can be derived if Object-Orientated Analysis (OOA) [4] is the preferred method of analysis within an organisation. Whatever method is used, it is crucial that the outcome represents ‘real world’ entities about which the system holds information, and is uncluttered by any detail concerning the implementation. Therefore, even for a large information system, the number of such entities is usually limited to 20 or so.

The ‘Model System’ process is used to understand the physical architecture of the system within which this information is stored and manipulated. It must identify processing elements, networking elements and storage elements. Normally, such models will be generated during the ‘system architecture’ phase of system development process. The depth, or level of detail, of such a model is of course variable. In our experience, it is best to work (at least initially) at a high level, as long as the model is complete at that level. In this way, the later stages of analysis will remain tractable and the ‘big picture’ will emerge. If necessary, particular features of the big picture can be examined in greater detail later.

Sometimes obtaining the service and system models is straightforward, as the logical and physical architectures have been laid down as part of the overall design of the system, and have been well adhered to during implementation. However this is not always the case, and the key challenge at the start of any SRA is often producing models that bridge the gap between the system an organisation thinks it has, and what is actually in place.

As a final step before undertaking specific security analysis it is useful to construct a cross-reference between the two models. This will identify for each data entity (or, at a more detailed level, for each attribute), which physical system elements are responsible for its capture, processing, transmission and storage.

3.3 Assess the Threats

Threats are things that might damage the business or organisation providing the service. This damage might be indirect, in the sense that most immediate loss might be inflicted on (say) a customer or employee. Threats are, in themselves, nothing to do with the technology which provides the system. A threat might be ‘the reputation of the business may be damaged by revealing sensitive customer information’. On the other hand, ‘an attacker outside the building might eavesdrop on 802.11b wireless traffic’ is definitely not a threat in SRA’s taxonomy. (As we shall see, it is a good example of a ‘vulnerability’.)

It follows that threats are derived from the Service Model. This is done by examining every entity in the model from the familiar ‘CIA’ angles: confidentiality, integrity and availability. For each entity, it is necessary to determine the worst consequences of the loss of Confidentiality, Integrity and Availability. To do this it is necessary to examine all entities for an attribute. For example, imagine a medical database system. For the ‘patient’ entity, the attribute ‘HIV status’ might be most sensitive from the confidentiality viewpoint, whilst ‘blood group’ might be most important from the in-

tegrity viewpoint. However, rather than carry forward the security characteristics of every attribute for every entity, it is generally more convenient to carry forward the 'worst cases' for each of C, I and A, at the entity level. Thus if there are IE information entities, the number of potential threats to carry forward in the analysis is $(3 \times IE)$.

Of course to perform this assessment, it is necessary to place the threat on some kind of scale. The key factor here is the Damage, D , which would be inflicted on the organisation by a single incident. Of course, this could be quantified fully and measured in pounds or some other currency. This is usually necessary in a commercial organisation where it is required to derive a security budget which is proportionate in terms of other potential use of funds. The direct cost is usually straightforward to estimate. However, indirect costs could predominate, for example the loss of custom which might arise from unfavourable publicity, and the marketing effort required to re-establish customer trust. The indirect costs are much harder to estimate. In some cases, for example where human safety is concerned, it may be considered inappropriate to quantify D in hard cash (for example by putting a price on a human life). For these reasons, in the interest of quickly identifying the big picture, we are often asked to perform semi-quantitative analyses: for example D might be assessed on a three-point scale: high, medium or low.

The second factor to assess for each threat is the Gain, G , which an attacker could derive from making a threat a reality. Of course the first step to doing this is to identify what types of attacker might have something to gain. Once again, this assessment could be made in hard cash (say if the most likely attackers are credit card fraudsters or an unscrupulous business competitor) or in a less quantitative way (say if the attackers are teenage hackers, motivated only by perverse self-satisfaction or kudos amongst their 'peers').

3.4 Assess the Vulnerabilities

A vulnerability is a property of a system component, which might be exploited to mount an attack. It has nothing to do with what that component might be used for in any particular system, at any particular time. As with threats, the 'CIA' classification is used, in this case to characterise the security properties of each component.

For each vulnerability, there are two properties to be considered. Firstly, the Cost, C , to an attacker of exploiting the vulnerability must be assessed. In considering this, it is important not only to include the marginal cost to a particular person of mounting an attack, but also to gauge the total amount of resources which need to be deployed. For example, an attack against a smart card chip might require a scanning electron microscope such as might be found in a university electronics department. Although a student might be able to gain access to such a device for minimal personal outlay, the fact that specialist (and expensive) equipment is required serves to limit the opportunity for an attack.

Secondly, the Likelihood, L , that an attacker will be caught during or after an attack should be assessed. For example, an attack that can be mounted using web-mail from a cyber-café, is much less risky for the attacker than breaking into an office (or bribing an employee) to plant a line monitor.

3.5 Identifying and Assessing the Risks

At this point in any SRA we now have depressingly long lists of threats and vulnerabilities to consider, but how many of them actually result in real risk to the organisation?

A threat presents no risk if it can never be realised, and a vulnerability only presents a risk if it can be exploited in a way that exposes an asset to a threat. Restated, a real risk exists if an identified threat can be realised by a vulnerability.

For a system managing IE information entities utilising *SC* system components, we have potentially $3 \times IE \times SC$ risks. For most systems, not all components handle all entities, so the cross-reference described above can be used to reduce the number of combinations to be considered.

For each risk, we wish to identify the degree to which the organisation is exposed, known as Exposure (*E*). This we take to be given by the Damage (*D*) done by any single incident, multiplied by the Probability (*P*) of an incident taking place in a given period:

$$E = D \otimes P \quad (1)$$

We use the symbol \otimes to denote a function that behaves like multiplication, to allow semi-quantitative approaches to be formulated. For example, if *D* or *P* is zero, then so will *E*. We will see this in the worked example presented below. Clearly, straightforward multiplication can be used when a monetary calculation is being performed.

We will already have assigned a value for *D*; as we have seen it is a property of the threat associated with the risk. How do we assess *P*? In SRA, we make the assumption that the probability of an attack being carried out is proportional to the attacker's expectation of gain. This we define to be the Profit (*Pr*) he might derive from a particular attack multiplied by the probability that he is not caught (*PNC*).

$$P = Pr \otimes PNC \quad (2)$$

Pr is calculated from the attacker's gain (*G*) and his costs (*C*). Using the same symbolic convention as before:

$$Pr = G \Xi C \quad (3)$$

where the symbol Ξ is used to denote an operator with similar properties to 'minus' (and which is 'minus' in the fully quantitative case). Notice that *G* and *C* are already known: they are properties of the threat and vulnerability respectively.

Of course *PNC* is just the inverse of the likelihood that the attacker will be caught:

$$PNC = 1 \Xi L \quad (4)$$

Notice that *L* is defined as a property of the vulnerability. We have now defined the exposure to each risk in terms of basic properties of the threat (damage to organisation and gain to the attacker) and vulnerability (cost to the attacker and his probability of being caught). When applied to every risk, we can generate an ordered list of information security risks. When the analysis is fully quantitative, the exposure to each risk can be compared to other types of risks run by the organisation.

3.6 Countermeasures

For the greatest risks, the organisation will wish to select countermeasures to reduce its exposure. Countermeasures can be technical (for example, the use of encryption to protect a communications link) or procedural (for example, the implementation of dual controls) or physical (for example better locks on doors). A countermeasure usually acts to reduce a vulnerability: either by increasing the attacker's cost or increasing the chance that he is caught. Given new estimates for these characteristics, the new (reduced) exposures can be calculated. In the fully quantitative case, it can be checked that the selected countermeasures fulfil the organisation's standard 'rate of return' criteria. Where no suitable countermeasure exists, insurance can be considered. Insurance is generally specific to a threat rather than a vulnerability; but the underwriters will of course be interested to know how each threat might be realised. In any case, we believe that presentation of a thorough SRA will help underwriters to provide competitive premiums.

4 Worked Example

The following is a worked example of the application of SRA and is provided for illustrative purposes only. In this example, a hypothetical service whereby an organisation offers its customers up-to-the-minute company share price and business news via satellite broadcast.

No information service serves any other purpose than to make information which is both accurate and timely, available to legitimate users only. If the information is not accurate or timely, or is made available to people with no right to access the information, then the system has been in some sense subverted.

Therefore, to understand the threats which apply to the service, it is essential to have a model of exactly what information the service is designed to maintain. A Logical Data Model contains the following elements:

- Entities, usually a 'real-world' item, which can be distinguished from other entities of the same type, for example 'person'.
- Attributes, which are a property of an entity, for example 'eye colour' or 'height'
- Relationships, which are links between entities.

In practice, it is the one-to-many relationships that are most important to data modelling. A one-to-one relationship is trivial, and indicates that the two entities can really be considered as one. Many-to-many relationships are complex and should be analysed further. They can always be resolved into two one-to-many relationships via a 'link entity'. Figure 2 and Table 1 identify the Information Entities used in our example, the relationships between them and their key attributes.

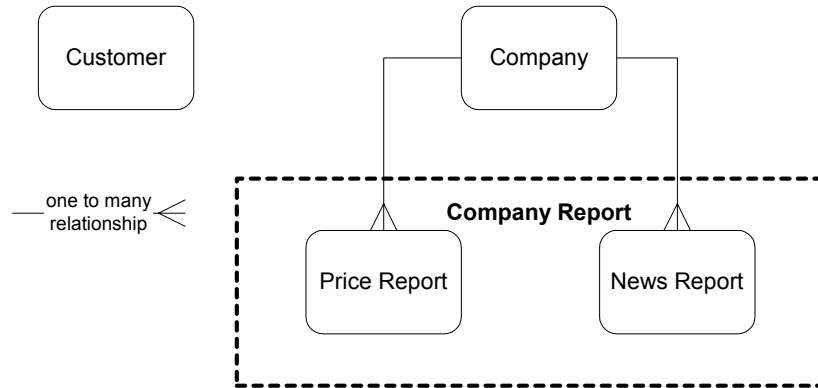


Fig. 2. Logical Data Model of Service

Table 1. Information Entities

Entity	Major Attributes
Customer	Customer ID Customer Name
Company	Company ID Company Name
Company Report	Company ID Date Time Price/News Item

For the purposes of this paper an extremely simple model is used, one that assumes that all ‘Customers’ receive the same generic service and get data relating to all listed ‘Companies’. Also it has been assumed that the ‘Price Report’ and ‘News Report’ entities are so similar in terms of the types of threat they may be exposed to and how they are physically processed, that they are considered as one data entity, ‘Company Report’.

The Service model is used to produce a Threat Catalogue identifying each threat, the damage its realisation would cause the organisation, and the maximum gain that any attacker could achieve. In some cases it is important to fully quantify the damage and gain as part of a budget allocation process. In this worked example a semi-quantitative assessment of the impact of a realised threat is used, based on a three-point scale of ‘High’, ‘Medium’ and ‘Low’.

This worked example is merely an illustration of how SRA is performed, and not a definitive statement of risk in such a system. Indeed it may be that the reader disagrees with the values allocated in this example. If this is the case, the value of such a structured approach becomes all the more clear, as it can be seen how simple assessments can be altered, and the effect rippled through the whole process to show the ultimate impact a small change has on the exposure of the organisation to risk.

Table 2. Threat Catalogue

Information Entity	Threat Type	Damage (<i>D</i>)	Gain (<i>G</i>)	Likely Attacker
Customer	Confidentiality	Medium	Medium	Business Rival
Customer	Integrity	Medium	Low	Hacker
Customer	Availability	Medium	Low	Business Rival
Company	Confidentiality	Low	Low	Hacker
Company	Integrity	High	Low	Business Rival
Company	Availability	Medium	Low	Hacker
Company Report	Confidentiality	Medium	Medium	Business Rival
Company Report	Integrity	High	High	Market Manipulator
Company Report	Availability	High	Low	Hacker

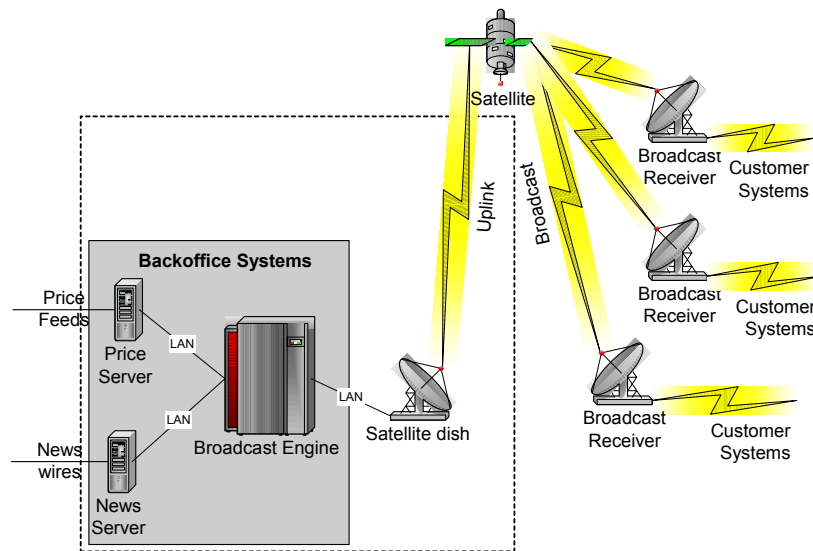


Fig. 3. Physical Model

Figure 3 shows the processing nodes, network links and data storage devices used within our example system. Each of these physical entities is considered in turn to produce a vulnerabilities catalogue such as that shown in Table 3. For the purposes of this simple example, the News Wire/Price Feeds are considered to have the same vulnerabilities, as are the Price Server, News Server and Broadcast Engine. Also the Satellite itself and broadcast receivers are not considered within this example, however, depending upon the implementation (e.g. is the broadcast receiver provided to the customer as part of the service?) these may require further attention if this were a real SRA.

Vulnerabilities are associated with the data links and processing nodes defined in the physical model. In both cases, they are categorised to correspond with the three different classes of generic threat—to confidentiality, integrity or availability.

For each vulnerability, the cost to the attacker of exploiting the vulnerability, and his probability of being caught are estimated. As with the damage and gain parameters associated with threats, these are given on a three point scale—high, medium and low.

For cost, these have the following definitions:

- HIGH - requires the resources of medium or large corporation, or organised crime syndicate
- MEDIUM - requires the resources of a small corporation or criminal gang
- LOW - within the resources of an individual hacker.

For the probability of the attacker being caught, the following definitions apply:

- HIGH - 0.5 to 1.0
- MEDIUM - 0.2 to 0.5
- LOW 0 - to 0.2.

In this example, it is assumed that no special security precautions are taken. Appropriate precautions will be derived following the risk analysis, when considering economic counter measures.

Table 3. Vulnerabilities catalogue

Physical Entity	Vulnerability	Cost to attacker	Likelihood of capture
Price Feed/News Wire	Confidentiality	Medium	Low
Price Feed/News Wire	Integrity	Medium	Low
Price Feed/News Wire	Availability	Low	Low
Back Office Systems	Confidentiality	Medium	Medium
Back Office Systems	Integrity	Medium	Medium
Back Office Systems	Availability	Medium	Low
LAN	Confidentiality	Medium	Low
LAN	Integrity	Medium	Low
LAN	Availability	Low	Low
Satellite Dish	Confidentiality	Medium	Low
Satellite Dish	Integrity	Medium	Low
Satellite Dish	Availability	Low	Medium
Uplink/Broadcast Channel	Confidentiality	Medium	Low
Uplink/Broadcast Channel	Integrity	Medium	Low
Uplink/Broadcast Channel	Availability	Medium	Low

Once Threats and Vulnerabilities have been catalogued, the risk analysis finds areas where vulnerabilities of the system can be exploited to realise a business threat, and ranks each such risk in terms of *exposure*.

As a precursor to this analysis, a cross-reference (e.g. Table 4) needs to be built between the logical and physical models defined earlier. This cross-reference shows which information entities are handled (i.e. stored, processed or transmitted) by which system components.

Table 4. Service/System model cross reference

CROSS REFERENCE		INFORMATION ENTITIES		
		Customer	Company	Company Report
SYSTEM COMPONENTS	Price Feed/ News Wire	✗	✓	✓
	Back Office Systems	✓	✓	✓
	LAN	✓	✓	✓
	Satellite Dish	✓	✓	✓
	Uplink/ Broadcast Channel	✓	✓	✓

For each valid combination of entity and component, risks are identified (classified, as usual, by confidentiality, integrity and availability) from the corresponding threats and vulnerabilities. The severity of each risk is calculated as follows:

Firstly calculate the ‘profit’ to an attacker (Pr) as defined in Equation (3) previously. In a semi-quantitative analysis, such as this example, Pr can be calculated using a lookup table:

Table 5. Attacker Profit lookup table

ATTACKER PROFIT		ATTACKER GAIN		
		HIGH	MEDIUM	LOW
ATTACKER COST	HIGH	Low	Negligible	Negligible
	MEDIUM	Medium	Low	Negligible
	LOW	High	Medium	Low

This amounts to subtraction of the cost from the gain to calculate the profit.

Next Assess the *probability that an attack will take place (P)*, as defined in Equations 2 and 4.

Table 6. Attack Probability lookup table

ATTACK PROBABILITY		ATTACKER PROFIT		
		HIGH	MEDIUM	LOW
DETECTION PROBABILITY	HIGH	Low	Negligible	Negligible
	MEDIUM	Medium	Low	Negligible
	LOW	High	Medium	Low

This amounts to multiplying the attacker profit by the likelihood that he does *not* get caught in committing the attack.

This then enables us to construct a Risk Catalogue, where each valid combination of entities and components on the model cross-reference (Table 5) is assessed once again under the three categories of Confidentiality, Integrity and Availability.

Table 7. Risk Catalogue

Entity	Component	Attacker Profit (<i>Pr</i>)	Attack Probability (<i>P</i>)	Damage (<i>D</i>)
CONFIDENTIALITY				
Customer	Back Office Systems	Low	Negligible	Medium
Customer	LAN	Low	Low	Medium
Customer	Satellite Dish	Low	Low	Medium
Customer	Uplink/ Broadcast Channel	Low	Low	Medium
Company	Price Feed/News Wire	Negligible	Negligible	Low
Company	Back Office Systems	Negligible	Negligible	Low
Company	LAN	Negligible	Negligible	Low
Company	Satellite Dish	Negligible	Negligible	Low
Company	Uplink/Broadcast Channel	Negligible	Negligible	Low
Company Report	Price Feed/News Wire	Low	Low	Medium
Company Report	Back Office Systems	Low	Negligible	Medium
Company Report	LAN	Low	Low	Medium
Company Report	Satellite Dish	Low	Low	Medium

Entity	Component	Attacker Profit (Pr)	Attack Probability (P)	Damage (D)
Company Report	Uplink/Broadcast Channel	Low	Low	Medium
INTEGRITY				
Customer	Back Office Systems	Negligible	Negligible	Medium
Customer	LAN	Negligible	Negligible	Medium
Customer	Satellite Dish	Negligible	Negligible	Medium
Customer	Uplink/Broadcast Channel	Negligible	Negligible	Medium
Company	Price Feed/News Wire	Negligible	Negligible	High
Company	Back Office Systems	Negligible	Negligible	High
Company	LAN	Negligible	Negligible	High
Company	Satellite Dish	Negligible	Negligible	High
Company	Uplink/Broadcast Channel	Negligible	Negligible	High
Company Report	Price Feed/News Wire	Medium	Medium	High
Company Report	Back Office Systems	Medium	Low	High
Company Report	LAN	Medium	Medium	High
Company Report	Satellite Dish	Medium	Medium	High
Company Report	Uplink/Broadcast Channel	Medium	Medium	High
AVAILABILITY				
Customer	Back Office Systems	Negligible	Negligible	Medium
Customer	LAN	Low	Low	Medium
Customer	Satellite Dish	Low	Negligible	Medium
Customer	Uplink/Broadcast Channel	Negligible	Negligible	Medium
Company	Price Feed/News Wire	Low	Low	Medium
Company	Back Office Systems	Negligible	Negligible	Medium
Company	LAN	Low	Low	Medium
Company	Satellite Dish	Low	Negligible	Medium
Company	Uplink/Broadcast Channel	Negligible	Negligible	Medium
Company Report	Price Feed/News Wire	Low	Low	High
Company Report	Back Office Systems	Negligible	Negligible	High
Company Report	LAN	Low	Low	High
Company Report	Satellite Dish	Low	Negligible	High
Company Report	Uplink/Broadcast Channel	Negligible	Negligible	High

Finally exposure to the risks is calculated by combining the Damage (*D*) (a function of the threat) with the attack probability:

Table 8. Exposure lookup table

EXPOSURE		DAMAGE (<i>D</i>)		
		HIGH	MEDIUM	LOW
PROBABILITY OF ATTACK	HIGH	Extreme	High	Medium
	MEDIUM	High	Medium	Low
	LOW	Medium	Low	Slight

In essence, the Damage is multiplied by the Probability of attack to generate the Exposure. It is exposure that is the key to aiding Organisations in making business decisions as to how they can make their systems economically secure. Table 6 indicates the Risks that open up our example service to medium or high exposure. The interesting point of note from this process is that the confidentiality of the data is relatively unimportant, as the kind of customer that requires this service is unlikely to go to the trouble of tapping it for free, as they can well afford the fees charged considering the value of the information to them. The integrity of the information provided about companies to customers is shown to be vitally important, as any actions taken on incorrect information from this trusted source is bound to have serious recriminations, both in monetary terms and in terms of reputation. In addition, false information inserted by an attacker could be used to create a false market, which is likely to lead to recriminations from financial regulatory bodies.

Table 9. Greatest Exposure

Entity	Component	Risk Type	Exposure
Company Report	Price Feed/News Wire	Integrity	High
Company Report	Back Office Systems	Integrity	Medium
Company Report	LAN	Integrity	High
Company Report	Satellite Dish	Integrity	High
Company Report	Uplink/Broadcast channel	Integrity	High
Company Report	Price Feed/News Wire	Availability	Medium
Company Report	LAN	Availability	Medium

5 Conclusions

Before describing the SRA method we listed the requirements that such a methodology should meet to be truly valuable as a risk assessment tool. Now let us consider how SRA meets these requirements.

- *Business context.* By providing the means for quantifying the exposure a particular risk opens an organisation up to, SRA enables business owners to make informed decisions as to whether the benefits of security measures justify their expense.
- *Technical grounding.* SRA uses existing systems analysis techniques to ensure the information service, and the physical system via which it is delivered are completely and openly defined.
- *Separation of concerns.* By separately considering the logical service to identify threats, and the physical architecture to identify vulnerabilities, SRA enables technical and business issues to be considered clearly and separately before the results are cross-referenced to identify where they combine to present a risk.
- *Support for quantitative analysis.* Depending on the information available, the time for analysis and the complexity of the system, SRA enables empirical estimations of cost in order to allow determination of an overall information security budget and the optimal allocation of that budget.
- *'Tuneable' analysis.* Implementing SRA need not be a massive undertaking. Depending on corporate size, circumstances and budget, performing such an analysis, even on simple high-level models of system and service, will provide some shape and direction to security expenditure and focus.
- *Evolution.* As the foundation of the method is based upon best practice in the field of systems analysis it is possible for SRA use the most appropriate techniques to model the logical service and the physical system, as such practice evolves and improves. Also because the method is not based on static logic (e.g. in a proprietary software package) but relies on the real knowledge and experience of real people, this enables dynamic response to advances in IT in general and IT security in particular.
- *Maintainability.* One of the benefits of performing a Structured Risk Analysis in this way is that when services and systems change, it is possible to tweak the models and see whether simple changes to information models, or physical architecture, effect the exposure to risk faced by the organisation.
- *Openness.* SRA is based on a unique combination of open standards in the area of systems analysis and security, requiring no payment of licence fees or purchase of specific software.

Security controls, and the actual dangers presented by the risks these controls are introduced to protect against, are fundamental to the economic success of an IT system. It is common sense, therefore, to give the deployment of such security controls the same sort of rigorous analysis and consideration as other areas of systems analysis and design. Such a statement is hardly controversial, however, the fact remains that quantifying risk is not easy, and is, therefore, often put aside in favour of more easily defined and controlled economic factors, such as development cost and revenue return.

SRA provides an approach that allow the same disciplines of structure and objective observation that are well established in other areas of systems analysis and design to be applied to security, thus placing security costs and savings on a level playing field with other factors that can effect the ongoing success of an information service.

As indicated in the introduction of this paper, we believe the key to taking optimal decisions in all these areas is to understand the risks to which the organisation is exposed by doing nothing, and the degree to which these could be mitigated by the deployment of particular countermeasures.

We believe SRA gives business owners the information they need to make truly informed decisions about the security, viability and future direction of their Information Systems.

References

1. Rogers, James, Simons, Mike.: Revenue's security glitch a warning for e-government, Computer Weekly Magazine, Thursday 6th June 2002.
2. Weaver, Philip L, Lambrou, Nicholas, Walkley, Matthew.: Practical SSADM Version 4+, Financial Times Pitman Publishing, ISBN 0-273-62675-2, 2nd Edition, 1998.
3. Yourdon, Inc.: Yourdon Systems Method: Model-Driven Systems Development, Prentice Hall, ASIN 013045162,: 1st edition (March 4, 1993)
4. Yourdon, Edward, Coad, Peter.: Object-Orientated Analysis, Yourdon Press, ISBN 0-13-629981-4, 2nd Edition, 1991