Tigger Team – A novel methodology to manage business risk

Ian McKinnon

Technical Report
RHUL-MA-2008-13
15 January 2008

Department of Mathematics
Roal Holloway, University of London
Egham, Surrey TW20 0EX, England
http://www.rhul.ac.uk/mathematics/techreports

Tigger team – a novel methodology to manage business risk

Author: Ian D. McKinnon
Logica

Supervisor: Professor Keith Martin
Information Security Group, Royal Holloway,
University of London, Egham, Surrey, U.K

Submitted as part of the requirements for the award of the MSc in Information Security at Royal Holloway, University of London.

I declare that this assignment is all my own work and that I have acknowledged all quotations from the published or unpublished works of other people. I declare that I have also read the statements on plagiarism in Section 1 of the Regulations Governing Examination and Assessment Offences and in accordance with it I submit this project report as my own work.

Signature:

Date: 7 September 2007

Table of contents

1. Introduction 5
2. Fundamental problems with IT security 7
   2.1. IT Security often fails silently 7
   2.2. Suppression of breach information 9
   2.3. Proving a negative 10
   2.4. Low probability – high impact events 11
   2.5. Accurate risk assessment 11
3. Business related problems with IT security 13
   3.1. Business imperative 13
   3.2. Value is created at the edge of what is technically possible 14
   3.3. Foundations necessary to build effective security on 14
   3.4. Demonstrating return on investment 15
   3.5. Building security in from the beginning 15
   3.6. Low marginal cost of service 17
   3.7. Outsourcing 18
4. The evolution of IT security 19
   4.1. Selling security to the business 19
   4.2. Change of focus from confidentiality to CIA 22
   4.3. Migration from risk prevention to risk management 22
5. The evolution of current best practice 24
   5.1. Security as a process 24
   5.2. Problems associated with managing security as a process 26
   5.3. BS7799 / ISO - IEC27001 27
   5.4. The cult of penetration testing 29
   5.5. Problems associated with penetration testing 30
   5.6. Changing nature of attacks 33
6. Tigger team methodology 36
   6.1. Introduction 36
   6.2. Basic methodology 36
   6.3. Problems with the tigger team 38
   6.4. Contractual exceptions for tigger team members 38
   6.5. Safeguards 39
   6.6. Staff risks 39
   6.7. Bomb risks 40
   6.8. Profiling potential attackers 41
   6.9. Motive, opportunity and means 42
      6.9.1. External threats 42
      6.9.2. Internal threats 43
   6.10. Threat agent pyramid 44
   6.11. Capabilities 45
   6.12. Capability does not provide motivation 46
7. Trialling the tigger team 47
8. Conclusion 48
9. Bibliography 50

Executive summary

Security is hard. Security is expensive. Security negatively impacts business function. All of these are bad, but far worse is the difficulty of measuring the effectiveness of security.

IT security over the last decade has become increasingly visible and important to a broad range of businesses. At the beginning of this period the response to IT risk was predominantly focused on technical prevention. Gradually this has evolved into a more business-oriented approach to risk management. This change has come about largely because of the perception that the technical approach to security provided too narrow a view of risk, failed to engage effectively with business and was failing to deliver benefit.

This paper explores a number of the fundamental difficulties that hamper the delivery of effective IT security. It also examines some of the difficulties created because of the conflict between the goals of security and those of business.

This paper describes a methodology that attempts to minimise the impact of a number of these difficulties. The primary goal of this methodology is to provide business with clear justification to support IT security activities and to demonstrate an adequate return on investment.

The methodology proposes the development of offensive and defensive capabilities within an organisation, in order to identify and manage both contextualised business risk and generic technical risk. The defensive capabilities act as both a control and a deterrent, but most importantly they provide concrete evidence of loss, which can be used to justify future activities. The offensive capabilities allow the business to refine an understanding of their specific risk, rather than generic risk. In addition they also allow realistic testing of the defensive capabilities through simulated attacks.

The methodology is cyclic and as it progresses the understanding and management of risks specific to the business should evolve. This will allow security to address increasingly remote and esoteric risks, until it is no longer possible to economically justify deploying mitigation. When this stage is reached the risks will be sufficiently small to fall within the business's risk appetite. The monitoring process should identify exploitation of these risks but no controls would be deployed because they would be uneconomic.

1. Introduction

This dissertation is largely motivated by a deepening dissatisfaction with IT security developed over a period of about 6-7 years working in the industry. I came to security after a period as a consultant specialising in systems integration - a posh way of saying "*plugging it together and making it do what it said it would do in the brochure*". I had been fairly successful at this and had especially enjoyed the sense of satisfaction derived from getting a solution working effectively by the end of a project. This process was clear and simple, even if the solutions themselves were sometimes complex. People wanted to do things with computers and networks to help their business. I helped them design and build systems to achieve their stated requirements – simple and satisfying.

Towards the end of this period I became more involved with security related projects including Firewall and DMZ design and implementation, to allow Internet connectivity, along with various content filtering solutions. In 2001 I moved to a pure IT security role and since that time I have rarely experienced the same feeling of satisfaction that I gained from successfully completing a systems integration project. Worse than that, I have found myself at odds with some of the methods, motivations and tactics that IT security has used to justify and sustain itself over this period.

These feelings have been exacerbated by the fact that IT security is an immature discipline. As a consequence there is precious little that is universally accepted as the definitive way to do things. Because there is so little 'canonical law' in IT security, professionals have to make it up as best they can as they go along. This is especially true in more rigid hierarchical organisations where long service is valued above competence.

There are a number of key questions that I have found very difficult to answer. These include: How do you know when you have the 'right' amount of security? How do you demonstrate value for money for your security expenditure? How can you convince your company, from the board down, that IT security measures are necessary and valuable? How do you know when you have done a good job?

In this paper I will initially explore some of the reasons why I believe it is so difficult to answer these questions and 'right size' IT security within an organisation. I will go on to highlight some of the obvious pitfalls that I have encountered within the security business.

I will then examine some of the existing approaches to security management and how they have evolved. I will highlight how they make it easier to answer the difficult questions presented by IT security. I will also point out where these standards and methodologies fail to address some of the important issues.

Finally I will present a novel methodology for assessing risk and managing security within organisations. This methodology is designed to avoid some of the fundamental problems of applying security outlined earlier in the document. It is based on the existing foundations of risk assessment and security management but re-orders the standard process and re-focuses efforts in order to be able to demonstrate real value to the business.

2. Fundamental problems with IT security

There are a number of fundamental problems associated with IT security. This section explores the most significant ones and examines how they impact on the ability of a business to effectively manage security.

2.1. IT Security often fails silently.

One of the few things that appear to be universally accepted in the field of IT security are the primary goals [1]:

- Confidentiality
- Integrity
- Availability

Of these only availability is by definition obvious when a failure occurs. The fact that a system or service is unavailable is self-evident to the user. If an availability failure occurs silently it is because there are no users attempting to use the service at that point. In the case of infrequently used or emergency functionality the failure will not remain unnoticed indefinitely but it is likely to cause a significant problem once it is discovered.

Breaches of integrity can certainly fail silently. However because the data remains in the hands of the business the failure can be detected using crosschecking, reconciliation or forensic audit. The ability to identify a failure is largely dependent on the nature of the data, its linkage with other known values or logs and the extent to which the system constrains unauthorised changes. For example a balance derived from a number of transactions can be checked by reconciling component transactions. On the other hand a text value for an individual's address provides limited opportunity for checking because it is not highly constrained, or certainly wasn't before the introduction of post codes.

When someone empties your back account of all your money it is easy to identify that something is wrong. In the same way if your wallet mysteriously becomes empty it becomes obvious when you go to pay for something. However a clever thief may only take £20 from a wallet containing £100. In this case it would be perfectly reasonable for the owner to be surprised at the remaining amount but put it down to a forgotten purchase or some other type of mistake on their part. It would be rare for an individual not to have

reasonable doubt that a theft had occurred if only a small proportion of the money was missing.

In Schneier's book "*Secrets & Lies*" [2] he describes the changing nature of attacks possible in an on-line environment and identifies automation as one of the three key differences. Automation allows adversaries to execute attacks that without computerisation would be uneconomic, characterised by *salami attacks that are* described in an article of the same name by M.E. Kabay [3]. A salami attack is where an attacker shaves a fraction from an asset many times. Each individual fraction is too small to arouse suspicion but the aggregate value is worthwhile for the attacker because the attack is automated. For a thief to continue to take low value notes from a wallet involves risk of discovery on every occasion. In a computer based attack the risk is limited to setting up the process and extracting the funds anonymously. Furthermore the fact that the funds are accrued in small amounts enables the attacker to avoid detection.

Computers automate repetitive and standard tasks to allow more customers to be managed by less staff. One consequence is that human intervention and oversight is largely removed from the process. Therefore sanity checking and monitoring for suspicious transactions needs to be programmed into the system in an attempt to provide the level of scrutiny that a human offers. Computers are not good at this type of qualitative assessment and will often simply not identify abnormal behaviour or activity leading to silent failures.

Confidentiality is historically the most important security goal but it is the most difficult failure to identify. This is because a breach in confidentiality is unlikely to result in a situation that provides the data owner with the opportunity to identify the failure. The only time a data owner is likely to be informed of a breach is if the attacker's goal is simply to access rather than exploit the data. This may be the case if the attacker is a journalist or penetration tester.

For example if a company loses control of a mailing list with millions of peoples' personally identifiable information on it they will have failed in their statutory duty to protect the data under the Data Protection Act [4]. However it is neither obvious how this breach would be identified nor whether the loss of confidentiality would actually impact those unfortunate individuals whose privacy had been eroded.

Furthermore if individuals do suffer a loss of privacy as a consequence of the breach it may be very difficult to determine where the data leakage occurred, especially if identical data is stored in numerous locations. If this were the case it would be nearly impossible to identify the guilty data controller as the source of the information.

When the security breach results in data loss how does the data controller know that they have lost control of data? As they do not have access to or visibility of the stolen data set they cannot directly identify the failure. The only way to identify this security failure is to correlate multiple privacy violations experienced by individuals whose data is held by them as evidence of a potential breach. If you were trying to identify the source of a leak you would need to match individuals who have suffered privacy violations against various databases until you find one that contains them all. If all the names were found on more than one database then it would be impossible to determine who lost control of the information. The final possibility to discover the origin of data in this case would be to attempt to identify the origin from minor discrepancies between records in order to link the breach to a single data source.

Seeding data sets is one way that misuse can be detected and is discussed in a white paper by Matthew Eberz [5]. Data controllers need to send subsets of their data to data processors or mailing houses to enable the execution of campaigns. There is a risk that these data processors will retain the data and use it. However if the data is seeded with bogus individuals then the data controller can check if the campaign is executed effectively and also if the data has been used for subsequent unauthorised activity.

The fact that security can fail silently makes it very difficult for security professionals to be confident that they have achieved their goals.

2.2. Suppression of breach information

When security is found to be inadequate and a breach is detected it is often not in a company’s best interest to publicise the fact. The damage to a company’s reputation could easily outweigh the monetary loss, so it is often better to simply write it off. Business would rather quietly learn from the mistake and implement mitigation to reduce the probability of a similar breach in the future. The other potential problem is that news of successful attacks will undoubtedly alert the criminal fraternity to the brittleness of systems, which could lead to copycat attacks.

Dorothy Denning describes an example that explains the reluctance to disclose information about breaches in her book "*Information Warfare and Security*" [6]. In 1994 a Russian hacker successfully stole money from CitiBank, by gaining access to its cash management system. This theft was identified, responded to and much of the funds transferred illegally were recovered. Fearing negative publicity CitiBank suppressed information about the episode to avoid alarm. A year later when CitiBank were presumably fairly confident that they had successfully avoided any negative repercussions from this incident details of the attack were used in open court during the trail of Vladimir Levin on similar but unconnected charges. The furore caused by this disclosure led to a catastrophic loss of confidence and resulted in the withdrawal of client funds, the value of these withdrawals were far in excess of the losses sustained in the attack.

Given this case study it is understandable why companies would wish to carefully manage the disclosure of security breaches. Whilst significant security failures happen infrequently, information about them is invariably suppressed and it is therefore very difficult to develop a clear and accurate understanding of the scale of the problem.

Interestingly California has taken the lead in requiring companies to disclose security breaches where the control of personal information has been lost [7]. This is as a direct consequence of the increase in identity theft crimes. Although it only relates to breaches related to privacy it will be interesting to see whether or not IT security as a whole will benefit from enforced disclosure.

2.3. Proving a negative

Security is basically all about preventing bad things happening. Unlike building systems and demonstrating that they work as specified, the goal of IT security is to ensure that nothing bad can happen. This goal is rather negative. It is also very difficult to achieve as proving a negative is itself fraught with difficulty.

For an IT security manager to stand up in front of their board of directors and say that nothing bad will happen is difficult. It is difficult to prove that something won't happen.

Furthermore if security fails silently, as discussed in a previous section, then it is very difficult for someone to assert with confidence that a breach has not occurred.

Also if information about security breaches is suppressed then the IT security manager does not get a chance to compare what is happening in their organisation with what is happening in peer organisations. This means that it is not possible to assert that the observed level of malicious activity is in fact the expected level because there is no information to base an expectation on.

2.4. Low probability – high impact events

IT security professionals are naturally drawn to focus on low probability – high impact events. This is understandable as they represent the most interesting and significant problems that could occur within an organisation. By focusing on these types of events it is easier to negotiate the case for funding IT security activity to the board.

The problem is the ability to accurately evaluate the risk associated with these types of potential events. Whilst the impact and financial consequences are significant, the probabilities that these events will occur are extremely low. As the probability of an event occurring tends to zero, the human brain has great difficulty thinking about the risk rationally, a problem explored in a paper by Jonathan J. Koehler & Laura Macchi [8]. This is demonstrated by irrational reactions to information about rare adverse health outcomes. This often means that statistics can sound more significant than they really are. For example if a one in a million chance of dying becomes a two in a million chance a headline could conclude that there was double the risk. However the absolute risk remains extremely low. The same argument could be made to support the deployment of a control to manage a very low probability risk.

2.5. Accurate risk assessment

Because companies are reluctant to publicise security breaches there is limited information on which to base judgements about the overall level of risk that business faces. The types of security breaches that make it into the news are the high impact ones that companies find impossible to suppress. This in turn artificially focuses businesses on low probability high impact incidents, making them appear more significant than they really are.

However for every high impact breach that makes it into the news there are likely to be numerous low impact breaches that go unreported, not to mention the presumably countless near misses.

By contrast, in financial risk management the information on which decisions are based is flooding in every second of the trading day. Stock prices are updated in near-real time. The risk manager has a clear idea of what he is protecting and its value because it is a big pile of money. The connection between the asset and profit or loss is direct, immediate and relatively easy to manage. If you would lose large amounts of money if a specific stock or index fell, you are able to mitigate that risk by taking out an option at a relatively small cost. The Black-Scholes [9] model for option pricing, which helped win the 1997 Nobel Prize in economics for its authors, is an excellent example of how risk can be efficiently and effectively quantified in financial markets.

The protection of IT assets is far more difficult. The value of the asset is often difficult to determine. The asset may be a number of interconnected systems and the value may be the correct interoperation of these systems. The protection that can be applied to mitigate risk is often fairly indirect; for example physical security measures to control access to the systems that the asset resides on, or screening of staff.

There is a significant difference between IT risk markets and financial risk markets. In financial risk markets everyone is operating openly and transparently. In the IT security market place risk is frequently associated with incompetent staff, hackers and criminals who, for a number of obvious reasons, have no intention of incriminating themselves. The criminals' activities are covert and it is therefore difficult to understand clearly how they operate. Honey pots and honey nets go some way to providing information in general terms about how these adversaries operate.

In IT risk there is no concept of a *market maker*, which does exist in money markets. In financial terms a market maker is obliged to trade in a given stock no matter what the state of the market is. They can use the strike price to manage their risk as they trade but market makers cannot refuse to trade. Unfortunately IT security does not operate as a market and as a result it is difficult to draw assistance from the similar but more mature discipline of financial risk management.

3. Business related problems with IT security

Some of the problems associated with delivering effective IT security are a result of the way business operates. Resolving these issues is an easier task than addressing the fundamental problems associated with IT security. They can be resolved through careful management of the business, especially by managing the perception of IT security at the board level.

3.1. Business imperative

Getting something working is far easier than getting something working securely. As a consequence wherever time and money are constrained, security frequently ends up de-scoped.

Business people can have a great money making scheme and cobble a solution together without ever stopping to think about the full ramifications of what they have created. It is very difficult to effectively express the risk of a solution to an enthusiastic businessperson. It is their job to see the opportunity, benefit and ultimately profit. In contrast the IT security professional is expected to see the whole picture, warts and all. It is difficult not to get into a situation where colleagues don't ask you to review their proposals because they fear a negative response that they don't comprehend. The IT security professional just becomes the person that says 'No!'

It is important to accept that business people have a responsibility to create new ways to make money. They are optimists and will focus on how their scheme will work. They are very unlikely to spend time thinking about how their new scheme may be defrauded. As a consequence they are pre-programmed to see only profit and never loss. The IT security input to the development of a business opportunity invariably takes a more pessimistic view of how a system will be used, or abused. In almost all cases neither party is able to perform a rigorous cost / benefit, or loss / benefit analysis so it is difficult to accurately assess the accuracy of either position. Suffice it to say that profits rarely meet the optimist's expectations and losses rarely meet the pessimist's fears.

It must be remembered that without a working system there is no value created or revenue stream to protect so it is important to understand the natural order applied to the various elements of the business. Security is always subordinate to revenue generation.

3.2. Value is created at the edge of what is technically possible

In the high technology business sector, where IT security is considered critical, most 'value' is created at the limit of what is technically possible. Whilst technology can be used effectively to boost profits by reducing a company's cost base the biggest profits are derived from technology that enables something very useful which was previously not possible. Mobile phone technology is a good example.

When you are trying to create profit from innovation, getting it working is often very difficult. As a result security is often sidelined in order to reduce complexity and ensure a working product makes it to market in the necessary timeframe.

In these sectors it is also possible that margins are very high once the solution is in place and working. Where there is excess capacity and the marginal cost of service is low, even moderate levels of fraud may not materially impact the business. In this case it may be decided to accept the losses rather than trying to prevent them up front.

3.3. Foundations necessary to build effective security on

IT security can only be built on robust foundations and frequently businesses do not provide adequate support. To be able to manage IT security effectively it is essential to know what you have, know how it works and know how it changes. IT security should not be expected to make up for inadequacies in the quality of IT operations management before they can offer advice on security. It is important to remember that the people in the best position to identify weaknesses should be system owners themselves. They should have an intimate knowledge of what their systems does, how it does it and what the impact would be if the confidentiality, integrity or availability of the system were impaired.

Operational rigor needs to be in place before attempting to apply security. This is especially true in the case of change management to ensure that security is not expected to hit a moving target. One standard that can be used to provide a stable foundation on which to build be secure environment is ITIL, or BS15000, or ISO20000 [10].

Security advice often relies on dubious assumptions i.e. "*Disable any OS / Networking services that are unnecessary*". This instruction requires an intimate and comprehensive knowledge of both the OS and the environment, which is often not available.

3.4. Demonstrating return on investment

As stated above the goal of good security is to ensure bad things don't happen. Unfortunately it is difficult to prove that bad things didn't happen due to chance rather than due to IT security expenditure. The difficulty in demonstrating an adequate return on investment is explored in an Information Security Forum paper by Adrian Davis [11]. Once security is in place the person responsible can simply say we are doing a good job because nothing bad happened. The IT security budget will remain at a specific level or increase to deal with emerging threats. However it is very difficult to determine if the IT security applied is cost effective. That would be easy if you could accurately calculate the losses that would occur if the IT security expenditure were zero. Unfortunately such a calculation is not possible for the same reason it is not possible to determine if it would have been quicker to stay on your original route after deciding to take a detour. In addition such a calculation would rely on the accurate evaluation of low probability - high impact events.

It should be pointed out that the difficulty of performing cost-benefit analyses is not confined to IT security. Business frequently makes significant decisions without clear empirical evidence that the outcome will be profitable. However at least most business initiatives have the advantage of potentially generating income. It is this potential for profit that ultimately drives these projects. Furthermore even if they are not as profitable as initially expected, either due to over optimistic forecasts or failure to accurately assess associated risks, this can always be hidden using creative accounting. IT security has no way to demonstrate a profit. The best that can be achieved is to provide an estimate of the level of reduced losses on the basis of previous figures.

Building security into a project from the beginning can also present a problem. Because the successful outcome for an IT security professional is for nothing bad to happen, it is difficult to demonstrate a causal link between the costs incurred applying security measures and the reduced losses. After all if nothing bad happens is it because of, or in spite of the IT security effort expended?

3.5. Building security in from the beginning

IT security best practice strongly suggests that it is best to build security in from the beginning of a project. Bolting security on after the event is believed to cost twice as

much and is invariably half as effective as if it had been incorporated from the beginning. Clearly this is at odds with the reality of operating at the cutting edge of technology and standard business pressure to get things done quickly and cheaply.

For the reasons stated above it is difficult to develop a clear argument, which clearly justifies the decision to build security into a project from the beginning. This is because it is difficult to project losses against a revenue stream that doesn't exist. It is therefore difficult to formulate a case for a specified level of expenditure on security.

Possibly more important than the difficulty of providing effective cost-benefit analysis for IT security spend, is the lesson of experience. This seems to point to the ability of business to rely on reactive rather than proactive security without incurring significant losses. One example of how business has ignored this 'best practice' advice, without suffering catastrophic failures, is the development of the mobile telephony market. Peter Howard from Vodafone covered this during a guest lecture at RHUL to MSc students [12].

The 1st generation mobile phones were almost totally devoid of security when they were rolled out. This did not stop companies thriving. As the technology was deployed the weaknesses in the system became known and were exploited. It should be noted that initially the exploitation of the weaknesses did not result in losses that were material to the company. As the losses increased, the companies worked on the 2nd generation phone system in an attempt to eliminate these weaknesses. The additional security built into the 2nd generation phones largely succeeded, as demonstrated by the fact that there were limited security enhancements built into the 3rd generation phones. In fact the most significant enhancement is mutual authentication between the handset and the base station to mitigate risks associated with rogue or evil twin base stations. Attacks using rogue base stations have been considered the exclusive domain of security services and law enforcement, although there is some evidence that the attacks are now within the capability of smaller non-governmental organisations and tech-savvy hobbyists [13].

This is a prime example of security best practice being ignored by business and business suffering no ill effects. A consequence of this will be to undermine the security industry and its various pronouncements of doom.

From the mobile phone example, adding security as it impacts the bottom line in order to stem losses was clearly an effective way of doing business. This risk management strategy is founded on a clear and accurate understanding of the various losses that a business incurs in order to take cost effective action to minimise these losses. There is of course a risk that the losses will overwhelm the company before it has a chance to react. There is a possibility that there will not be an effective mitigation for the vulnerability, leaving a company with the option of continuing business whilst accepting the losses or shutting up shop, which is simply a profit / loss calculation.

It is probably worth pointing out that businesses like the mobile phone industry can accept losses because they are so profitable, whereas those operating on smaller margins become unprofitable far more quickly. As a knock on effect, big business has more money to spend on security and smaller businesses have less. However attackers will target big companies in preference to small companies because they also want to maximise their profits.

3.6. Low marginal cost of service

In new technology the costs are primarily for R&D and the roll out and maintenance of infrastructure. The actual cost of providing the service is minimal. When the system capacity far exceeds the normal level of usage, fraudulent use of the infrastructure has an insignificant impact on costs.

In the case of the mobile communications industry, the marginal cost of the product meant that the losses had little to do with call time. In the same way that phone phreaks initially used techniques to make calls to support their hobby rather than to evade charges for calls that they would otherwise have paid for, the losses associated with cloning first generation mobile phones were less to do with the airtime and more to do with the administrative mess of unravelling fraudulent calls from genuine ones. As a consequence the inconvenience and subscribers' loss of confidence in the service became the most significant elements.

It is interesting to note how attack techniques evolved within a small community and then spread to a wider audience. It is probably accurate to say that the vast majority of people would not engage in fraudulent activity simply because they believe it to be wrong. However once a technique becomes so widespread it not only becomes difficult to

control, it may also start to appear legitimate to the law-abiding members of the population. A good example is music downloads. Once public perception is sufficiently altered about the legality and therefore acceptability of an activity, it may be too late to enforce a right that exists.

3.7. Outsourcing

Outsourcing and the use of consultancy services are very common in the IT sector. Whilst there are a number of situations when the use of external consultants is justified, there are some critical business functions that should be retained and managed in-house. IT security is one of these functions.

Consultants can provide very specific technical skills, which may be critical to the success of a project, to quickly fill a gap without the delay of getting staff trained. Consultants can provide short-term resource during busy periods to execute tasks that could be handled in-house if capacity were not an issue.

Longer-term outsourcing can be a cost effective way of releasing a business to focus on its core capabilities. Outsourcing real-time monitoring of significant but sporadic events that could occur 24x7 to an aggregation service is effective. It very difficult to hire and retain appropriate staff to monitor 24x7 when a businesses incident count is low and the skills required to effectively interpret the events is high.

There are some functions within a business that are just so important that you wouldn’t consider passing control of them to an outside organisation. Business strategy, finance and security fall into these categories. These elements are critical to a business and if there is insufficient skill and knowledge in-house to manage them it can only indicate a failure.

Using consultants to fill a skills gap to provide security advice may provide short-term relief but it is likely to result in larger problems in the long term. There is no easy substitute for a permanent and capable in-house team to manage IT security within a business.

4. The evolution of IT security

This section focuses on how various elements of IT security have evolved. The very real problem of how security professionals justify the need for security to the business is covered. This is followed by two sections which describe the changes that have occurred as IT security has moved from a government setting that requires confidentiality to a wider range of commercial organisations that have more diverse security goals.

4.1. Selling security to the business

For the reasons given in the previous sections it is clear that selling security within a business represents a significant challenge. Whilst this challenge has resulted in a variety of tactics being employed over time, none have really acknowledged or addressed the underlying difficulties.

Selling security is like selling insurance. It is difficult to show a good return on the investment in premiums unless there is an accident. Because of this difficulty the initial pitch of pioneering security professionals was to use Fear Uncertainty and Doubt (FUD). Spend this money on security or bad things will happen. By using FUD many security professionals secured funding for the most basic protection – firewalls and content scanning. Unfortunately using FUD has a limited shelf life and like the boy who cried ‘wolf’ business started to ignore the funding requests based on doom scenarios. Worse than this it highlighted to the business the fact that security professionals, like most business activities, were unable to justify the expenditure they were proposing. Unlike most business initiatives however, which have profit as the ultimate goal, security initiatives were unable to demonstrate such a clear-cut benefit.

It was clear that worse than simply no longer working, the use of FUD sowed the seeds of mistrust in security from the top-level management. The feeling that budgets were simply being used to keep techies in a constant supply of rack-mounted appliances with flashing blue lights was possibly not unfounded. The reliance on ‘technical security’ to provide the protection required was coming to an end.

Another nail in the coffin for the ‘technical’ approach to security may have been the anti-climax of the Y2k issue. Along with getting fed up with FUD, Y2k may have helped to develop management resistance to stories of doom and gloom.

Security needed a new way to sell itself and get back into managements good books. Along came the massive financial collapses of Enron and MCI. In the aftermath the Sarbanes-Oxley Act was brought into force in an attempt to reduce the potential for fiduciary malfeasance in US listed companies. A good overview is available in the Wikipedia [14], but most importantly included in this act were segments about the effectiveness and accuracy of financial reporting systems … that is computers. Therefore the integrity and availability of financial data was vital to demonstrate the probity of these companies. IT Security decided that it had a significant role to play in the drive towards good corporate governance. Cleverly IT security recast itself into a corporate governance and compliance role, in an attempt to curry favour with directors who were genuinely scared about their new responsibilities and the personal liability that they brought.

In reality both Enron and MCI were at their heart old-fashioned ‘frauds’ based on extremely complex and creative financial accounting mechanisms that had precious little to do with IT security, despite what the shiny new corporate governance and compliance officers told us. FUD had ended its useful life and had to be replaced with something. As a result SOX has now stepped into the role vacated by FUD as the new pressure point to drive IT security.

This is all faintly ridiculous as SOX only applies to public companies listed on the New York stock exchange (NYSE). Indeed once the implications of SOX became apparent, a number of UK based companies with their primary listing on the London stock exchange (LSE) withdrew from the NYSE, where they had secondary listings, to avoid the cost of compliance. Despite this fact we had a number of guest lecturers on the MSc course at RHUL suggest that simply doing business with a US company required a UK company to achieve SOX compliance. This preposterous position makes me wonder if FUD has really disappeared.

Around the same time that SOX was emerging, a slightly gentler sales pitch was being developed. “*Security is a business enabler*” was a mantra that was glibly trotted out by a number of distinguished speakers who came and dispensed their wisdom to those

studying for their Masters degree. It is a phrase that for me demonstrates a fundamental misunderstanding of the problem by anyone who attempts to use it. Whilst it is a pleasing phrase that does much to try to make security positive, underneath the spin and style there is no substance. As has been stated, security is about stopping bad things happening; it is about defence; it is about mitigating risk to an acceptable level - it needs to be based on a deep understanding of the technology, the business and its associated processes.

At the time the phrase emerged, business in general was dissatisfied with the return from their technical security investment. It was difficult to identify what was being derived from this expenditure. Like a rabbit from a hat the phrase was plucked and then quoted ad nauseam. It is as if the user believes that saying something that is patently false sufficiently often will finally make it true. By hiding the unpalatable truth about security and making business-friendly statements we will retain our budgets for another year. After all bad things don't happen that often so the probability is that we will get away with it without anyone discovering the reality. In the meantime we can continue to order more 'security' appliances with blue flashing lights to keep up the security theatre, rather than needing to interact with the business in a meaningful way.

Ultimately IT security is still in the position where the problems of justifying security spend remain unresolved. We have transferred the reasoning behind security spend, but it is still inadequate. The elephant is still in the room.

What weakens the position further is if management do not see similar companies suffering losses, attracting fines or even going out of business as a result of IT security failures. Unless this happens they are unlikely to sit up and take notice. After all they are in the business of making money and at the heart of the matter, business relies on the simple strategy of charging as much as possible and spending as little as possible.

The reality is all too obvious for top management to see – companies don't appear to be suffering material losses as a result of IT failures. Companies don't appear to be being fined for breaches in data protection. Companies are not going out of business as a result of IT security failures. It is simple economics, in the absence of compelling evidence, to limit spending on IT security.

Even when cases are publicised the financial impact is limited. Shares don't plummet and sales don't either. In fact the only people who pay much attention are the IT security professionals who simply say "*we told you so!*" The example of the TKX privacy breach is a case in point, and is covered extensively in the Register [15].

Whilst there is theoretically a risk of providing "*too much security*" I'm certain that IT security people would not be concerned about this possibility. Unfortunately business people seem much more sensitive to the possibility of "*too much security*". For this reason it is important to determine what is the right amount of security for a given situation to ensure that there is no tension between IT security and the business. This balance is best determined using monetary values of data and the cost of breaches to the business, in order to ensure that the cost of security does not exceed the cost of failure. There is undoubtedly a fundamental tension between the levels of security that IT security people want and the level of security that the business wants.

4.2. Change of focus from confidentiality to CIA

IT Security has evolved out of classically high-security environments – government and military - where confidentiality is the key requirement. With increased computerisation in business the range of organisations that require security and the type of security that they require have both expanded. A bank needs integrity and possibly availability over confidentiality, although confidentiality is still important if only because of obligations under the DPA.

In his paper "*Cryptography and Trust*" [16] Professor Richard Walton discusses this change of focus in relation to trust. Where security has moved away from the controlled government environments that require confidentiality and have the processes, staff and infrastructure to deliver the goal, to far less controlled environments which require a broader range of security services.

4.3. Migration from risk prevention to risk management

In classic secure military and government environments you can be certain that there is an adversary. You can't necessarily see them or know who they are but you know that they are there. They wish to gather sensitive or secret information to give themselves, or the governments that sponsor them, an advantage through intelligence. In business there is no such guarantee. In fact there is no such guarantee that companies' sensitive

information will be of much value whatsoever. Most criminal activity is related to getting money for free. There are of course other attack motivations with varying degrees of risk aversion all the way through to religiously motivated suicide bombers.

As IT security matured as a discipline two distinct groups emerged to undertake the challenges. Those from an IT background moved towards security using their understanding of the technology, how it can be used and how it can be subverted along with various technical mitigations. At the same time auditors and business people moved towards security from a detailed view of the business and a detailed understanding of both profit and loss. As a result these two groups met in the middle.

Initially IT security was considered to be a function of the IT department. As a consequence the technologists took the upper hand in the struggle to control IT security. IT security was firmly anchored in the deployment of technical mitigation to prevent risk.

Over time the importance of IT security to business became apparent at board level. The failures of technical security, the inability to justify expenditure and the inability to take a business-centric approach to risk weakened the technologists' grip over IT security. They have subsequently been gradually replaced by auditors, quality and compliance officers as the custodians of IT security. In addition many IT security groups were migrated from the IT department to Finance departments and re-cast as risk management groups.

It is too early to tell if this migration was primarily the result of dissatisfaction with the apparent failures of the technical approach to security, or a belief that a business-centric approach would be more effective. I believe that it is safe to say that good security will come from a combined approach of managing both technical and business risk.

5. The evolution of current best practice

5.1. Security as a process

As has been alluded to, much of the early security activity was focused on technical risk and technical solutions. These solutions were inevitably point solutions for point risks. The individuals who were developing these solutions were IT staff who were firmly rooted in the technology. They knew how the technology worked and they knew how it could be broken.

This led to the deployment of technologies such as firewalls and anti-virus capability to deal with what were perceived to be the most significant threats in that period.

One of the more interesting descriptions of the realisation of the shortcomings of a techno-centric approach to security is by Bruce Schneier. In his book "*Secrets and lies*" [2] he effectively apologises for his earlier book "*Applied cryptography*" [17]. He apologises because he was deluded into thinking that cryptography was the answer to all security problems. He was seduced by the beauty and perfection of the mathematics. Unfortunately this delusion became apparent as he assessed system after system that contained significant security flaws - despite the beauty and perfections of the maths. It became all too obvious that the mathematics, so solid and so perfect in an abstract setting, became brittle when applied in the real world. Weakness was introduced due to poor implementation, human controlled elements or invalid assumptions such as the availability of effective random number generation.

The final straw for Schneier was when a colleague remarked, "*The world is full of bad security systems designed by people who read 'Applied cryptography'*". This indictment made it clear that an absolutist view of security simply didn't work. An absolutist view says 'deploy this technology and the risk disappears'. This position is adopted by naive technophiles who believe that technology can solve all their problems.

In "*Secrets and Lies*" Schneier firmly asserts that security is a process not a product. It is only by engaging fully in an entire process that security will be addressed end-to-end. This process comprises assessment, protection, detection and response.

The assessment phase forces you to determine the value of the asset or system that you are attempting to protect. It forces systems designers and implementers to focus not only on how their systems will work and be used but possibly more importantly on how their systems could fail or be abused. Once the probable failure modes have been described the impact of each of these, in terms of confidentiality, integrity and availability, is determined. In addition an attempt to estimate the probability of each attack scenario being realised is made.

Using the information from the assessment phase it should be possible to identify appropriate controls to protect each asset. The protection needs to be in proportion to the value of the asset and the probability of an attack.

In addition the principle of defence in depth should be applied. This states that two or more weak but complimentary controls are likely to be more effective than a single strong control. This principle is based on the acceptance that no single control can be 100% effective.

Protection without detection is pointless. If it is accepted that no control can ever be 100% effective then it is an essential part of any defensive posture to identify when breaches have occurred. Where defence in depth is applied this is especially important because a single failure should not result in a breach. In this case action can be taken before the complimentary control is broken and the asset is compromised. In the case of a failure where a single control is deployed, the detection is more likely to be the initial step in a forensic investigation of the breach.

Finally the response stage comprises two elements. The first is the immediate actions necessary to limit the damage, gather more details, assess and manage the impact and recover normal operations. The second element of response occurs after the situation is stabilised. This involves reviewing the details of the attack and passing concrete evidence back into the security process. The assessment phase can be re-entered and reviewed in light of events to enhance quality of the risk assessment in the light of the empirical evidence provided. The existing protection applied can be strengthened or complimented with an additional control in the light of the new understanding of the threat. Finally the detection process can also be modified to better identify similar attacks in the future.

5.2. Problems associated with managing security as a process

This security process is in many ways attractive as it attempts to counter the techno-centric absolutist approach to security.

However there are problems with applying the process, a number of which have been explored earlier in this paper. Accurate assessment of risk is difficult. There are numerous factors why this is so – asset valuation, impact costing (especially impairment of intangibles i.e. damage to reputation), probability calculations due to lack of empirical data, uncertainty surrounding threat agents.

In addition to these difficulties risk assessment can remain skewed to focus on technical risk rather than end-to-end risk. It is important to assess risk within a business context to ensure that a true picture is derived.

There are also a number of fundamental problems associated with detection. First and foremost is the sheer scale of the task of wading through the numerous logs searching for traces of an attack. The task is perceived as tedious and monotonous and as a result is considered as low value / low kudos / low grade work. Finally if attacks occur infrequently then the probability that an operative will spot the signature of the attack in the sea of data is actually quite low. Bruce Schneier discusses this problem in his book "*Beyond Fear*" [18]. For these reasons detection is often done poorly, if at all.

The response step in the process should ideally be invoked when suspicious but ultimately harmless events occur. This way more can be understood about how the systems are being used and provide insight into emerging threats. Clearly response needs to be invoked when an attack occurs and is identified. Because detection is rarely performed in a rigorous way response is often confined to significant events that can't be ignored. In this case response and detection are effectively merged into a forensic investigation to understand what took place rather than to manage events as they happen. One of the additional difficulties that detection faces in this scenario is that those individuals who forensically examine logs for evidence have limited understanding of what they should contain. This makes identifying abnormal entries even more difficult.

The result of these difficulties often leads to the contraction of this cyclic four-stage process. In the best examples assessment is attempted, protections are applied on the basis of these assessments and incident response exists to deal with significant failures. Unfortunately in the worst cases the process contracts to simply applying protection on the basis of the current general consensus on risk.

5.3. BS7799 / ISO - IEC27001

Another example of how security has evolved of from a techno-centric approach can be found in the development of BS7799, an information security management standard.

The development of an information security management standard is one of the most important milestones in the relatively short history of IT security. Of the numerous standards that exist the original and most globally important is currently defined by a combination of ISO/IEC27001 and ISO/IEC17799. These two standards attempt to codify the steps necessary to secure IT within a business. They are generic in that they attempt to cater for businesses of all sizes and in all market sectors.

This standard started life as BS7799. It was developed in the UK under the guidance of the British Standards Institute (BSI) by a host of commercial organisations. The work that underpinned the original standard was the Department for Trade and Industry (DTI) "*Code of Practice for Information Security Management*" published in 1993. Two years later in 1995 the BSI released their "*Part I – Code of Practice for Information Security Management*" as BS7799-1:1995 [19]. The standard was revised in 1999 to remove UK specific references and allow it to be applied globally.

It should come as no great surprise that BS7799-1 contained a list of generic risks grouped into ten categories, a set of controls by which those risks could be managed and advice on the implementation of controls.

Whilst it should be pointed out that these controls were being offered up without any assessment of how appropriate or necessary they were, they did provide a good starting point. They started to be adopted as "*business best practice*", which carried some risk considering they had not been demonstrated to be proportionate for a specific environment. The reality was that organisations simply implemented the controls that they thought were appropriate for them and ignored the ones that they considered irrelevant.

This mix and match approach had a rather unfortunate side effect when the working group tried to develop certification criteria for the standard. It quickly became apparent that due to the nature of the standard it was impossible to certify against it. There was no obligation to implement any of the controls and a business could therefore certify against BS7799-1 by doing nothing and not even have to document their justification for this stance.

To resolve this impasse BSI developed a second security management standard with the primary intention that businesses would be able to be certified as complying with it. BS7799-2:1999 – Specification for the information security management system was derived from the management process Quality Management Standard which ultimately became ISO9001:2000 [20].

This moved the standard from simply saying implement whichever controls you think apply to you from this list, to describing the processes needed to determine risk, select and apply appropriate controls and maintain those controls over time. Security had become standardised as a process.

The process was modelled on the Plan, Do, Check, Act cycle. This is to all intents and purposes identical to the "*assess, protect, detect, respond*" cycle discussed above. A critical part of the Information Security Management System (ISMS) that the part II standard defined was the "*statement of applicability*". This document sets out the rationale for including or excluding specific controls, detailed in BS7799-1, in light of the perceived risks and the risk appetite of the business.

In 2005 the International Standards Organisation (ISO) adopted BS7799 as a global standard in the following standards: ISO/IEC17799:2005 and ISO/IEC27001:2005 to cover parts one and two respectively.

It is easy to see how the development of the information security standard mirrored the move from a techno-centric approach towards a more business-oriented approach to managing security. Unfortunately the standard does not remove some of the fundamental or business problems related to delivering IT security within a business.

5.4. The cult of penetration testing

In contrast to the way that much of the IT security industry has evolved, the penetration testing industry appears to have become more tightly focused on technical minutiae. Despite this penetration testing is considered to be an important part of an organisation's security posture.

Penetration testing is an exercise whereby individuals attempt to break into or subvert electronic systems using 'hacking' techniques. If a penetration tester attempts to compromise a system and fails it is reported as secure. Clearly if exploits are directed at a system are successful, or vulnerabilities are observed during testing, then the organisation has the opportunity to resolve these issues.

Testing can be undertaken at the end of the development cycle, before a system is exposed to the Internet. Most testing is carried out on Internet facing systems because these are most at risk. They are at risk because they can be attacked from anywhere, by anyone with access to the Internet.

Testing can also be carried out on internal systems. This is usually restricted to critical internal servers rather than end user systems. Internal penetration testing is predominantly used to sanity check the patching schedule for vulnerabilities in operating systems and critical services.

The goal of a penetration testing attack is to identify vulnerabilities within the target system. These may have the potential to crash the system leading to an availability issue. It may allow users to manipulate values that should be protected. It may allow users to defraud the system and result in gain for the attacker. The ultimate goal is to find a working remote exploit that results in administrative access to the host operating system. These type of exploits are usually related to buffer overflow vulnerabilities [21].

Whilst identifying buffer overflows exploits has been possible for a considerable period, since 2000 locating buffer overflow conditions within systems has become increasingly automated. The Morris worm is an early example of how poor programming can be used to subvert the operation of a system. The process of finding vulnerabilities and then developing exploits for them has become commoditised.

Fuzzers can be used to automatically discover buffer overflow conditions without the need to review source code. Fuzzers are used to send malformed or illegal inputs to a system in an attempt to identify weaknesses, such as buffer overflow conditions. These tools speed up the discovery of weak points within systems and allow researchers to quickly identify where they should focus their attention. Once a weak point is identified it can be manually investigated to determine whether or not it could be successfully exploited.

This manual investigation often involves debuggers or virtual operating environments that allow researchers to clearly understand how the system fails. This information can then be used to craft malformed packets that exploit the vulnerability in a consistent and useful way.

What we have seen since around 2000 is the systematic and widespread scouring of common operating systems and their services for vulnerabilities. These have led to some fairly dramatic security failures in the shape of worms such as Slammer, Nimda, Blaster. It could however be argued that vulnerabilities in ubiquitous network facing services have been exhausted. As a consequence bug hunters are moving on to more obscure sub-systems and services in an attempt to find more vulnerabilities. The impact of finding a vulnerability in a non-internet facing service, or a service that is only run by a small percentage of users is relatively limited.

5.5. Problems associated with penetration testing

Unlike most other areas in IT security, penetration testing has become more and more focused on arcane technical vulnerabilities over time. This contrasts with the evolution to a more business-centric approach adopted by IT security in general.

Penetration testing has developed into a vulnerability treadmill. In order to continue to demonstrate value the penetration testers need to find new and increasingly obscure technical vulnerabilities with which to scare their clients. This is important for the penetration testers, even if the vulnerabilities identified constitute no real threat to the business. There is a risk that this treadmill leads the business nowhere because the issues identified have limited relevance to the overall security posture of the organisation.

Whilst some large organisations run vulnerability scanning internally, the majority of penetration tests are undertaken by external consultants. By outsourcing these activities a number of issues arise. Firstly individuals outside the organisation know about weakness within the security infrastructure. This knowledge may be of specific vulnerabilities that are present or a more general appreciation of the level of security deployed within the environment. Second the ability of penetration testers to identify vulnerabilities is often not matched by their ability to explain the potential impact of the vulnerabilities discovered. It should also be pointed out that occasionally the client staff members who receive the information are not able to comprehend the explanations. Finally rather than locking knowledge into the internal team and adding value over time, penetration testing provides knowledge to the business in the form of a report. This report, even if it is fully understood, has limited value over time. The only people who really benefit long term from performing the penetration test are the penetration testers who take the knowledge gained away with them.

For many penetration testing activities the target is either Internet facing or will be once the system goes live. In this case the only access requirement that an attacker needs is an Internet connection. However many penetration tests are often carried out on systems that are specifically protected from direct Internet connectivity. Take the case of a classic three-tier architecture comprising web server, application server and database server. Individuals on the Internet have access to the web server sitting in its own screened sub-net, or DMZ, exclusively on ports 80 and 443. The web server communicates through a firewall on a single port to the applications server which also resides in it's own DMZ. The application server in turn communicates through a firewall on a single port to the database server which also resides in it's own DMZ. As a consequence an attacker must compromise two hosts via a maximum of three ports before gaining access to the database server using a fourth port.

Despite all the protection provided by this compartmentalisation it is not uncommon for penetration testers to request a network port on the database DMZ so they can execute their test. Frankly if an attacker gains all ports access to the database server in your environment you have already failed. Penetration testing often ignores defence in depth in order to focus exclusively on the vulnerabilities it is most interested in and able to exploit.

Many testing tools passively identify potential vulnerabilities. This is because to confirm that a vulnerability exists would involve sending exploit code to the system and may result in a crash, or unpredictable system behaviour. On systems that are in development executing active tests may be advantageous because there is limited risk. However if the system under test is live, or contains live data, then it may be decided to restrict activities to those that do not threaten the integrity and availability of the system. The consequence of this cautious approach is that vulnerability reports may contain numerous notifications that are in fact false positives.

In my opinion the most significant failure of the penetration testing paradigm is the inability to determine the relevance of the discovered vulnerabilities to the business. The expense associated with developing a genuine understanding of a system and how vulnerabilities could be exploited for gain makes it unattractive for clients. It is therefore more viable for testing to remain focused on generic technical issues associated with widely adopted operating systems and services rather than business specifics.

Much penetration testing is simply a matter of pointing "*off the shelf*" or open source tools at a target system. The process of discovering and reporting on vulnerabilities is automated and requires limited human intervention or interpretation. Whilst there are exceptions it is rare that a tester gains a genuine understanding of the operation of the system. In many cases the tools may report vulnerabilities that the tester is unfamiliar with and cannot therefore determine the significance to the business.

The probability is that during most penetration tests the most active attack attempted will be value manipulation using an in-line http proxy. This is superficially impressive but of limited significance because this class of attacks is relatively easy to execute when you are controlling both the browser and the proxy but become extremely difficult to execute even slightly remotely when a victim is driving the targeted browser.

Penetration testing is realistic provided real attackers are penetration testers and real attacks look like penetration tests. This is because the capabilities required to exploit the vulnerabilities are finely honed. It is unlikely that an individual without experience of penetration testing would have the capability to execute an attack. Moreover the goals of an attack are different from the goals of a penetration test so it is difficult to assert that a penetration test is representative of a real attack.

This does not stop the vulnerability treadmill alerting. It does mean that business have to be cautious and assess the impact to themselves accurately before making a significant investment in immediate patching. The need to assess the impact of a specific vulnerability within a business context is often where penetration testing falls down.

Penetration testing also reinforces the fallacy that security is about 100% effective mitigations. If a penetration test identifies that no vulnerabilities exist the report indicates that there is nothing to worry about until the next test in six months time. This view of threat is inadequate.

It should also be a concern for the business if the only report that results from a penetration test is a report from the test team. The business should expect a real-time alert from their monitoring team identifying suspicious activity. If this does not happen it means that if someone was maliciously scanning the network for vulnerabilities the business would be blissfully unaware. As vulnerabilities emerge all the time a penetration test is at best a sanity check that patching is up to date. Identification of suspicious activity by contrast is a far more effective way to protect the systems on a day-to-day basis.

5.6. Changing nature of attacks

Although the types of attacks that can be executed remain fundamentally the same, their nature changes as a result of computerisation. Bruce Schneier identifies three key differences in the nature of attacks brought about by the introduction of computers [2]:

1) automation
2) attack at a distance
3) technique propagation

Networking technology provides the opportunity to attack a target from a distance. This could be from inside or outside an organisation. This not only increases the number of potential victims for an attacker, it also provides an additional level of protection for the attacker. This protection is afforded by the relative anonymity of network-based attacks and the difficulty of prosecuting attackers for both practical and jurisdictional reasons.

Exploitation techniques travel very much faster with the introduction of the Internet. One example of the speedy and efficient dissemination of attack techniques is associated with and affiliate marketing scheme. Affiliate marketing schemes allow advertisers and even customers to receive payments from purchases made at online stores. The affiliate network tracks users' activity as they click through an advertisement displayed on a web site to an online store and purchase an item. Once the item is purchased a small percentage of the transaction value is returned to the web site that hosted the successful advertisement. In other schemes the customer making the purchase is entitled to some money back thereby acting as a discount.

The affiliate tracking mechanism is stateless and has limited security incorporated to deter fraud. Within two months of launching a new affiliate marketing scheme fraud started to occur. It would probably be true to say that the fraud perpetrated against this new scheme was likely to have been developed and used on pre-existing schemes and the scams were simply transferred to the new setting. However the use of on-line communities, bulletin boards and web sites made the knowledge of these scams available globally. It would also be true to say that the people who frequented these sites were relatively few in number, because the activities discussed were of a legally dubious nature. As a consequence the overall level of fraudulent activity remained low and did not have a material effect on the ultimate profitability of the service.

Another example of the speed with which attack information can be disseminated is related to a poorly thought out marketing campaign in the same affiliate marketing scheme. An on-line retailer who was engaged in the scheme had a large promotional budget of loyalty points to distribute. This budget was allocated to the organisation when they joined the scheme but had to be used within a set timeframe. As the deadline got closer a promotional campaign was hastily developed. Customers would receive a significant bonus point allocation if they bought inkjet printer cartridges from the retailer. It was noted by some eagle-eyed customer that the cost of the cheapest cartridge was significantly less than the value of the points accrued in the transaction. As a result purchasing cartridges represented a negative cost to the customer.

This is similar to a backwardation in stock quote system. A backwardation occurs when a bid price from one market maker is higher than an offer price from another market maker. The bid price is always lower than the offer price creating a difference called a spread.

This spread is where brokers make money from simply executing trades, regardless of whether a stock increases or decreases in value. When a backwardation occurs it is theoretically possible to buy a large volume of stock at the offer price and immediately sell it at the higher bid price thereby making an instant profit. However the easy way for market makers to avoid a loss due to a backwardation is to not answer the telephone. An online system is unable to make such a choice.

News of the ink cartridge scam was posted to a bulletin board system that specialised in discussions about loyalty card schemes. The thread contained advice about how best to maximise the return from purchases and regular updates from individuals confirming that the promotion was still operational. Whilst multiple purchases of the cheapest cartridges were not in the spirit of the promotion there was nothing in the Terms and Conditions that restricted this activity. The promotion came to a natural conclusion: a fairly small group of individuals had bigger loyalty point balances and hundreds of black ink cartridges appeared for sale on e-bay.

Whilst the old adage which says "*If something looks too good to be true it probably is*" holds, there are always lazy marketing managers who can provide an exception to prove the rule.

6. Tigger team methodology

6.1. Introduction

The paper so far has focused exclusively on the problems that hamper the development of security within an organisation. In this section I will describe a methodology that attempts to overcome some of these problems.

This involves shifting the focus of security activity: setting up a cyclical process that helps develop capabilities within the organisation. The primary goal of this methodology is to identify and address the specific, contextualised, risks that are directly relevant to the business, rather than focus on generic technical risk. Being a cyclical process the capabilities will develop over time, thereby gradually raising the bar against increasingly complex attacks and capable adversaries. One of the critical success factors is to be able to clearly calculate the return on investment for IT security expenditure. As a consequence the methodology should provide empirical evidence that security within the organisation is ‘right sized’.

6.2. Basic methodology

The basic methodology involves two teams operating alternately in offensive / defensive modes. The offensive team adopts the personality of a specified threat agent and tries to develop and execute attacks that are within their grasp. Initially the capabilities of the threat agent will be limited, but as the cycles progress the capabilities of the attack personality will increase. The defensive team is tasked with monitoring the environment to identify abnormal and suspicious activity. The defensive team need to investigate anomalies and respond appropriately. Because the defensive team are operating a live monitoring service within the organisation they could encounter evidence of genuinely malicious activity. One would hope that any suspicious activity encountered would be the result of simulated attacks from the offensive team.

This looks very much like an ongoing penetration test and in many ways it is. However there are very important differences between a tigger team and a standard penetration test. The tigger team is designed to develop both attack and defence capabilities in tandem and therefore incorporates monitoring and response activities. The goal of the

attack team is to develop an understanding of generic technical risk but more importantly specific business risk. This should allow the tigger team to develop a deeper understanding of the business systems, processes and assets and enable them to more accurately assess the genuine risks their company faces by placing them into a business context. By thinking like attackers they can explore a whole range of potential attacks open to the threat agent they are operating as and determine which ones are most likely to result in maximum gain for the attacker, or loss for the company.

The change of focus from protection to monitoring and response in the security process has the potential to bring significant benefits. It would be true to say that monitoring is an extremely difficult part of the security process to manage as explained above. However if the difficulty associated with monitoring is met head on then the rewards can be substantial. It is often said that protection is useless without detection so it is important not to get draw into the trap of abandoning monitoring because it is so difficult.

A large part of the initial tigger team methodology is developing ways in which monitoring can be automated to deal with the volume of events that are generated. This will by necessity result in the development or implementation of an enterprise event management system. It could be developed in-house using existing logging technology such as syslog or be based on a proprietary solution such as Arcsight. Either way getting an organisation-wide logging solution will provide vital information on what is really going on within the organisation.

This process puts the team into an ideal position to develop mitigations and possibly more importantly, develop monitoring and detection techniques that will ensure that even if the attack could be executed it couldn’t be executed without detection and response.

Another advantage with the tigger team methodology is that it incrementally develops crucial capabilities within the company. This is in contrast to paying for outside penetration testing resources where the long-term value is limited. In addition the cyclical nature of the methodology means that capabilities develop over time in a way that is most appropriate for the business context.

Through the monitoring of malicious activity it is possible to extract empirical data on which to base the security response. Although this approach is reactive rather than pro-

active it is possible to determine when the return on investment becomes insufficient to justify further expenditure. As a result it is possible to demonstrate that security is providing an adequate level of protection for the risks identified.

6.3. Problems with the tigger team

The information and capabilities developed through this methodology are potentially very damaging. For this reason it is critical that control is maintained over information, tools and processes. If they fell into the wrong hands they could be used against the company to perpetrate fraud. As a result it should be appropriately marked and handled. It is also worth considering operating the tigger team from a secure area. Because the attack team need to have access to information and software that could be used for offensive activities they should operate from a secure area. Although it does not initially appear as critical, the monitoring capabilities, information and processes of the defensive team are also very useful to an attacker. This is especially true if they wish to try and avoid detection so it is suggested that they also operate from a secure area.

There is a significant risk associated with staff members developing offensive capabilities. It may put them in a position where they have knowledge of vulnerabilities and may be tempted to exploit them for personal gain. To ensure that someone does not get into a position where they can abuse the offensive capabilities they have developed, it is essential to have an absolute minimum of two tigger team members. In an ideal arrangement both the offensive and defensive teams should contain a minimum of two people. This way the execution of frauds against the company would require some level of collusion.

The goal of the tigger team methodology is not to eliminate vulnerabilities but to reduce the number of vulnerabilities in a cost effective way, and to ensure that monitoring will be able to identify fraudulent activity attempting to exploit those vulnerabilities that cannot be eliminated in a cost effective way. The end result would be that tigger team members would be aware of remaining vulnerabilities but also aware that they would not be able to exploit them without being caught because of the monitoring they know is in place.

6.4. Contractual exceptions for tigger team members

It is essential to put in place appropriate contracts for tigger teams members. This is because they are being asked to engage in activities that will put them in breach of

company policy. It may also put them in breach of computer misuse law. There must be exceptions to company policy to allow the offensive teams to do their job. One advantage is that this type of offensive activity can be used to clarify how company policy and / or law is broken and how internal disciplinary procedures would proceed and on what basis. In addition it may give some indication on how to prosecute a fraud through the legal system should that be necessary.

6.5. Safeguards

There is a need to differentiate between offensive activity carried out as part of the remit of a tigger team and genuinely malicious offensive activity carried out by an individual from a tigger team. This distinction could be managed using activity logs. Tigger team members would be required to pre-log all offensive activity that they intend to engage in. The defensive teams monitoring suspicious activity could then cross check the activity they have detected against logged activity. Offensive activity that is not logged is deemed to be malicious and is treated as a disciplinary matter. Any evidence that a tigger team member is engaged in covert activities is a cause for considerable concern and should be dealt with appropriately.

There may come a point in the cycle where an attack develops sufficiently that there may not be effective monitoring in place to identify it. In this case there is a possibility that the attack could be executed without the defensive team being able to respond. However once the teams get together for a debrief session at the end of the cycle they can jointly identify potential mitigation or the monitoring necessary to identify the attack. This also presents an opportunity to launch a forensic investigation, which is what would happen if a genuinely malicious attack were executed successfully. The two teams can get together to examine the digital evidence in an attempt to understand what happened, how and by whom thereby developing their forensic capabilities.

Tigger team logs should be signed and dated regularly. They should be checked and counter-signed by the individual responsible for the management of the tigger team process. Individuals should maintain their own log.

6.6. Staff risks

The knowledge developed by the tigger team members is very valuable. As a consequence the tigger team needs to be remunerated realistically and fairly to ensure

that they do not simply seek better paid work elsewhere or become consultants. This is more likely if they are focused on the type of vulnerabilities favoured by penetration testing. This is because generic technical risk can be applied to any company without having to take the time to understand their processes and procedures. However in the case of the tigger team the knowledge is contextualised for a specific business. This makes it more difficult to simply take those specific skills and apply them elsewhere for more money.

Staff turnover does occur so it is essential to have effective succession planning in place to ensure that the knowledge built up within the team is not dissipated through a few key defections. This is especially difficult because the team has to remain relatively compact to ensure that the knowledge developed is not spread too far within the company. It is also true that the ability to move individuals in and out of the team from and to other parts of the organisation carries risk because of the capabilities developed whilst members. For this reason it is essential to ensure that career development within the tigger team is managed very well.

Attacks don't always occur in the way you expect. For example e-Stores expected a cookie hijack / replacement scam but instead encountered a timing scam. The tigger team may not always attack in the way that an attacker would and may miss attacks as a result.

6.7. Bomb risks

The tigger team methodology is fundamentally based on a reactive process, which responds to losses as they become apparent. In the early stages of development the ability of the business to defend itself against more sophisticated attacks is limited. As a result there remains the potential for bomb risks. Bomb risk is characterised by the complete destruction of a business through the successful exploitation of a vulnerability.

It is important to accurately assess and manage bomb risk. The tigger team is designed to manage risk that leads to gradually increasing losses as they happen. However there is the potential for specific types of risk that could result in a company being put out of business almost immediately. It is these types of risks that must be managed pro-actively.

The tigger team monitors fraudulent activity to manage losses. It also needs to identify and mitigate any 'bomb' risks before they have a chance to do irreparable damage to the

business. There is both a reactive portion and a proactive portion to the tigger teams responsibility. It is quite easy to calculate cost-effectiveness of the reactive part because it is responding to concrete monetary losses. These losses are measurable and the response to them needs to be proportionate to ensure money spent stemming the losses is less than the losses themselves. The pro-active part needs to take the value of the business and top slice a proportion to pay for the identification and mitigation of bomb risk.

Bomb risk should be easy to identify. Whereas it is relatively easy for a business to suppress information about minor frauds, it is very difficult to cover up the fact that you have gone out of business. This means that the number of businesses falling foul of bomb risk should be relatively easy to determine.

There is very little evidence to suggest that even medium sized companies go out of business due to bomb risks. Because of this it becomes quite difficult to determine the value of the top-slice necessary to provide cost-effective bomb risk management.

Clearly the threat from bomb risk is dependant on a number of factors. The smaller the business, the less able it is to withstand a successful attack. The more reliant a company is on a single asset or service the more significant an attack on that asset would be. This is especially true for businesses that exclusively rely on electronic channels to do business. E-business is far more exposed to bomb risk than bricks and mortar companies that shift product.

6.8. Profiling potential attackers

It has been asserted that IT security has migrated from high security environments to business environments. In a military environment it can be assumed with 100% certainty that there will be well-funded and capable forces attempting to gain access to confidential information or disrupt the ability to operate effectively. Because there is no doubt about the existence of threat agents it is reasonable to spend money on protecting the assets and capabilities.

However in business there is no such assurance that threat agents are intent on bringing your company to its knees. If there are no threat agents then it doesn’t really matter how

many vulnerabilities exist in your environment as no one will exploit them and to do anything about them could represent a waste of time and money.

If different classes of attackers are not guaranteed to exist for a specific company then the risk equation for these classes evaluates to zero. Security expenditure defending against these classes of attackers is effectively wasted.

For this reason it is important for business to think carefully about the threat agents their business may face. This involves identifying potential threat agents' goals, capabilities and risk appetite.

6.9. Motive, opportunity and means

For an attack to take place the following elements need to be present - motive, opportunity and means. For any individual there has to be an acceptable balance of these three elements, which is moderated by their risk appetite, before they will undertake an attack. Threat agents can be split into internal and external. IT security initially focused on external threats, but has gradually realised that internal threat agents are responsible for the majority of security incidents.

6.9.1. External threats

The number of individuals that are external to a company runs into the billions. On the plus side the vast majority of these have no knowledge or interest, malicious or otherwise, in the company. However if the company has a connection to the Internet there are millions of individuals who have some type of (hopefully) restricted access to your systems. Of these there will be a small sub-set who may have the skill, motivation and resources to attack. Many will randomly select the company because it appears to provide access to a known vulnerability that the attacker is targeting, rather than any direct interest in that company or its activities. In these instances the motivation for the attack is not related in any way to the victim.

More significant are those attacks that are targeted specifically at an organisation. In this case the motivation is directed and the goal is more than simply finding somewhere that the attacker can break into.

The random attackers once they have gained access to company systems may deface web sites, or use resources to continue their random search for vulnerable systems, or use resources to host illegal activities. They are less likely to spend time and effort attempting to gain a deeper understanding and gather intelligence on the organisation whose perimeter they have pierced.

In contrast a directed attacker will use the penetration to gather information on the additional systems they have access to, set up surveillance on the compromised systems in an attempt to gather sensitive network traffic and passwords and attempt to understand the business processes supported by the systems. This is a laborious process that needs patience and skill. By compromising additional hosts from the initial foothold it is possible to drill deeper into an organisation and gain access to more valuable systems that are not necessarily Internet facing.

6.9.2. Internal threats

Internal staff members have far more opportunity than external staff to cause harm to a company. They have a level of access that allows them to complete their job. In terms of motive, opportunity and means and internal member of staff has opportunity in abundance. Given that greed is a common vice you could argue that there is motive in terms of financial gain. The level of access and understanding that an insider possesses goes a long way towards means. However it is possible that the individual will need to develop some technical attacks skills to execute an attack. In addition unless the attack results in direct access to cash another element of ‘means’ would be how to convert the item of value they have purloined into cash. The skills, tools and contacts necessary to turn a database full of personal details into cash can be sourced from the Internet.

For both Internal and External threats risk aversion is an important factor. Different attackers have different levels of risk that they are willing to accept. At one end of the scale individuals may reject the idea of doing something that would result in disapproval from people around them. Most people would not engage in an activity that could result in them going to jail. At the other end of the scale suicide bombers expect to die as a consequence of executing an attack.

In general terms rational people do bad things in inverse proportion to the probability that they will get caught and be punished. If the probability of getting caught is sufficiently high

then this will deter a rational attacker. It won't deter an attacker who doesn't care if they get caught such as those with mental illness, journalists who can plead a public interest defence or those with no fear of punishment.

As a consequence of this hackers will continue to hack whilst they perceive that there are inadequate resources to bring them to justice. Some may convince themselves that what they are doing is not that harmful, others may believe it is wrong but their chances of being prosecuted are sufficiently low to continue anyhow.

6.10. Threat agent pyramid

It is helpful to view the threat agents arranged in a pyramid.

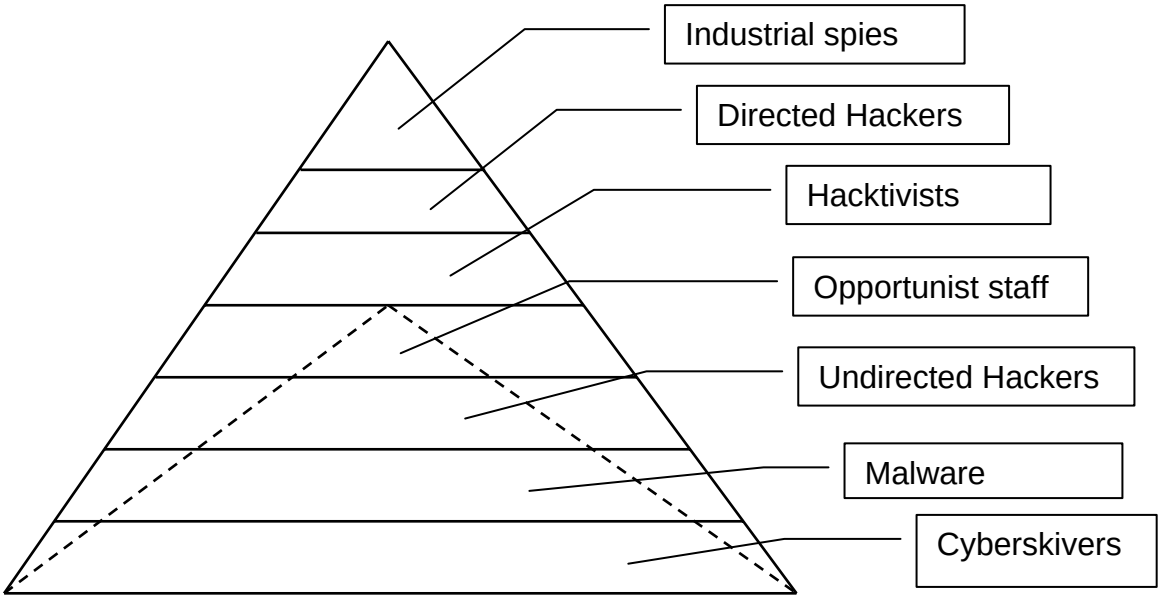

Figure 1 – Threat agent pyramid

At the base of the pyramid are low level risks that are widespread such as *cyberskiving*. Cyberskiving is using the web for non-work activity during the working day thereby reducing productivity and impacting the company. Although each individual engaged in cyberskiving results in a very small loss to the company the number of individuals who are engaged in this activity means that the losses can be significant. What is more, the requirements for motive, opportunity and means are met by all staff.

As you move up the pyramid the probability that adverse events will occur reduces, the probable impact of the event increases, the capabilities required to execute the attacks increase and the risks associated with executing the attack increase.

Industrial spies populate the top of the proposed pyramid. These threat agents are highly motivated, well resourced and funded and have significant capabilities. They are risk averse to a degree in that the value of the information that they seek is often linked to the ability to retrieve it out without being detected. Clearly being fired for breaching company policy on computer misuse would be bad, although not disastrous for an industrial spy as it wasn't their primary job in any case. However being identified as an industrial spy would be very bad indeed.

One of the suggestions of the pyramid is that the overall aggregated cost of individual event classes may be the similar. Worrying about low probability high impact events to the exclusion of high probability low impact events does not make economic sense. Especially when you consider the difficulty an organisation would have managing the high impact events and the comparative ease with which the high probability events could be managed.

The purpose of the tigger team is to try to quantify some of these losses and address them from the bottom of the pyramid up. By doing this it is possible to show a return on investment, develop capabilities (predominantly monitoring and therefore an understanding of what is occurring in the environment) and progress towards tackling lower probability higher impact events. Primarily this progression should be in response to evidence of loss, although addressing bomb risk may be a secondary motivation.

The dotted line indicates the probable extent of the threat pyramid for an ordinary company as opposed to a high security environment. This is flattened out on the assumption that there is insufficient value that could be derived from targeting an organisation with expensive attacks such as industrial espionage.

6.11. Capabilities

This is the level of expertise that is needed to meet the 'means' threshold of the 'motive opportunity and means' requirements. Different attackers will have different capabilities. Some of these capabilities will be a function of their job, others will be personal capabilities. The personal capabilities may be attack capabilities such as hacking techniques and tools. It is important for the tigger team to accurately assess the capabilities of a threat agent and accurately assess the ability of that agent to develop

capabilities if motive and opportunity exist. The availability of attack tools and tutorials on the Internet allows an autodidactic approach to developing capabilities.

When threat agents who are at the top of the threat pyramid operate in organised groups, have their own resources, have the ability to discover and exploit 0-day vulnerabilities then the tigger team is very much on the back foot. The critical difference between a professional and an amateur is the ability to operate covertly. The majority of the skill that a professional exhibits is the ability to operate without detection. As the tigger team understands the lower levels of the pyramid and develops monitoring and mitigation to address risks at this level they migrate capabilities up to more subtle anomalies that a professional attacker expects to go unnoticed.

If a company is targeted and a capable threat agent is installed within that organisation it is unlikely that it would be possible to stop them executing an attack successfully. Clearly this starts off with a pretty high level motivation in the first place. It is probable that the attacking organisation has significant capabilities and resources – means. By inserting an operative into the target organisation they provide a level of opportunity that is unthinkable if they were to attack from outside the organisation.

6.12. Capability does not provide motivation

Whilst we worry about the attacks at the top of the pyramid, capability doesn't always imply motive. For example there are a number of cases where hackers have purloined millions of credit card details and don't know how to convert that into value for themselves. Equally if a member of staff took a copy of the customer database, what is the likelihood that they would be able to convert it into a cash value? Without a benefit for the attacker the motive is weak and the attack becomes improbable.

There is quite a significant threshold before an attacker is in a position to benefit from a security breach. Security people are often concerned about an arch nemesis hiding in a hollowed out volcano waiting to attack them in fiendish and technical ways. The reality is rather different for almost all industries other than government and military. It is important to consider not only who has the resources to execute an attack against your company but also who has the resources to benefit from such an attack?

7. Trialling the tigger team

Clearly these ideas need to be tested to determine whether or not they provide concrete benefits over the current methodologies. I had hoped to run a trial against a test environment using both standard penetration testing and a tigger team approach. However given the limited time available to design, set up and trial a test environment it became obvious that it would be very unlikely to provide meaningful results. The environment would be too limited to be representative of even the smallest company. It would have been difficult to develop the environment in a neutral way so as not to favour the proposed methodology over the status quo. The limited scope of the experiment meant that it would be very difficult not to pre-programme a favourable outcome.

If the tigger team process were to be tested this is how I would expect the test to be structured: Two identical environments complete with systems, physical security, business processes documentation. The two teams would manage the security of these identical systems. One team uses conventional security methodology and the other uses the tigger team approach. Various business events are played out in the environments over time. These events are initially normal transactions and business process. As time progresses a variety of attack events are played out into both environments. With each cycle the capabilities and security posture of the two teams evolve and these differences can be assessed in their ability to resist the emerging attacks. At the end of the cycle the two teams have to present their capabilities, demonstrate to the board why continued support and funding are necessary and assess the risk to the business.

To move the test into a more representative environment it would be necessary for it to be adopted in a company. The capabilities of the IT security team could be base-lined at the start and then tracked as they applied the methodology. The security posture of the company could be base-lined and then tracked as well.

For a controlled experiment you could use two different regions in a large company. This would hopefully present a level playing field in terms of risks. However it is unclear how an accurate evaluation of the losses would be gathered from the control region, as one of the problems associated with existing methodologies is that they are unlikely to provide such data.

8. Conclusion

The purpose of this paper was firstly to identify, express and explore a range of problems that hinder the deployment of effective IT security management and secondly to describe a novel methodology to reduce the impact of these problems. The methodology described is an iterative process that continually develops security capabilities within a company and enables risk to be managed in context. The methodology allows the risk to be controlled in response to empirical data thereby ensuring that the security applied matches the threats faced and demonstrating an acceptable return on investment.

In conclusion I will try to highlight how the proposed methodology mitigates some of the difficulties identified and provides a more sustainable foundation on which to build and maintain effective IT security within an organisation.

One of the key factors in the proposed methodology is migrating the initial focus of the security process to monitoring. Unfortunately simply accepting the primary importance of monitoring does not make doing it effectively any easier. However if enterprise monitoring is put at the heart of IT security activity a number of other difficult issues become more manageable.

Accurate monitoring will provide the empirical data necessary to determine the actual losses incurred by the business. This information can be used to direct the deployment of security controls supported by a meaningful cost - benefit analysis.

By deploying security reactively in response to actual events rather than expected events the business can ensure that benefit is gained from every part of the security spend.

The use of complementary offensive and defensive teams draws on the benefits delivered by penetration testing but provides a number of additional advantages:

- Testing is performed by company staff, thereby developing the knowledge and capabilities of the organisation.
- Tests are performed iteratively, making simulated attacks increasingly complex and realistic.
- Tests are performed with a greater understanding of the business context.
- Tests help develop the monitoring and defensive capabilities of the organisation.
- Tests can employ a wide range of attacks against both technology and business processes from the bottom up.

By cycling the teams between offensive and defensive activities it is possible to more quickly develop skills, understanding and capabilities. Setting up a competitive scenario motivates both teams and improves the probability that simulated attacks will be as realistic as possible.

By adopting the persona of different threat agents the offensive team can more effectively replicate realistic attack scenarios. Developing attacks using the motive, opportunity and means of a threat agent allows the team to start simple, then build in complexity until the risks are deemed acceptable.

The focus on monitoring will reduce the probability that suspicious or malicious activity will remain unnoticed. Over time it will also allow the business to build a clearer picture of events and failures that can be used to help predict future trends.

By placing the focus on monitoring it allows the business to do what it wants, whilst providing evidence when things start to go wrong. Holding the business accountable for the identified losses will provide an incentive to address security issues upfront. This method allows the business to compare the profitability of projects that build security in from the outset and those that can't be bothered.

The methodology provides a more realistic approach to security than the one adopted by high security organisations for whom 'failure is not an option'. It is much more appropriate for commercial organisations where risk is simply another thing to be managed. This change of focus matches perfectly the migration from risk prevention to risk management.

The most important benefit is that it allows security to be effectively sold to the business. Rather than guessing at the threats faced by an organisation this methodology links the losses suffered to protection and provides a clear indication of the return on investment.

I believe that the potential benefits of this novel methodology have been demonstrated in this paper. There is undoubtedly scope to refine the tigger team methodology through further research. However I believe that the best way to test these ideas and their effectiveness is to apply them in a practical setting by creating and operating a tigger team.

9. Bibliography

1 Shon Harris – *CISSP Certification* – McGraw-Hill/ Osborne 2003

2 Bruce Schneier - *Secrets & Lies* - John Wiley & Sons 2000

3 M. E. Kabay – *Salami fraud* - Network World Security Newsletter, July 2002

4 HMG – *Data Protection Act* – HMSO 1998

5 Matthew Eberz - *Protecting Company Data through Data Seeding, an analysis of Personal Data* - Tech-I LLC August 2004

6 Dorothy E. Denning - *Information Warfare and Security* - Addison-Wesley 1999

7 *http://info.sen.ca.gov/pub/01-02/bill/sen/sb_1351-1400/sb_1386_bill_20020926_chaptered.html* - February 2002

8 Jonathan J. Koehler, Laura Macchi - *Thinking About Low-Probability Events. An Exemplar-Cuing Theory* - Psychological Science Vol. 15 Issue 8 August 2004

9 Fischer Black and Myron Scholes - *The Pricing of Options and Corporate Liabilities* - Journal of Political Economy, 81:3, 1973

10 BSI – *ISO/IEC20000-1 IT Service Management* – BSI 2005

11 Adrian Davis - *Return on security investment – proving it's worth it* – ISF December 2005

12 Peter Howard – *Lecture notes: GSM and UTMS security* – Vodafone 2007

13 http://www.binrev.com/forums/lofiversion/index.php?t28559.html - 2007

14 The hive mind - *Sarbanes-Oxley Act* – Wikipedia 2007

15 John Leyden - *http://www.theregister.co.uk/2007/08/16/tjx_charges/* - The Register 2007

16 Richard Walton - *Cryptography and Trust* - Information Security Technical Report II Elsevier 2006

17 Bruce Schneier - *Applied Cryptography* - John Wiley & Sons 1996

18 Bruce Schneier - *Beyond Fear* - Copernicus Books 2003

19 BSI – *BS7799 Information Security Management – part 1& part 2* – BSI 1995

20 BSI – *BS9001 Quality Management System* – BSI 2000

21 Aleph One [pseudonym] – *Smashing the stack for fun and profit* – Phrack 7(49):14 November 1996