7 Elements
Independent information security consultancy

Information Security Assurance from a Resilience Perspective

Table of Contents

Introduction

Today the global business environment is more complex and interconnected than ever before. Organisations rely on electronic data as their lifeblood, and the systems that enable the storage, transport, access and manipulation of this data have become critical. Even simple spreadsheets can become mission critical systems in their own right and this has resulted in an era where networks and the applications sitting within them have become the very backbone of every organisation regardless of their size and market sector. As a result, networks and applications are a primary channel for businesses and one that they must protect if they are to meet their businesses objectives and in the end, to survive.

For many organisations, their approach to information security results in a fortress mentality that focuses on the implementation of defences and preventing an attack. It is increasingly acknowledged however, that we cannot build sufficient defences to be 100% secure while allowing our organisations to effectively carry out their business, and as such, for many this siege based approach is no longer acceptable. A more resilient approach to the management of information security is therefore needed. This approach should not only take into account the mentality that organisations cannot be 100% secure but also acknowledge that the cost of securing our organisations can be large. A risk based approach should therefore be adopted which takes a more holistic approach to managing information security that accepts that the risks cannot be fully mitigated and adopts a resilient approach. Doing so will therefore place greater emphasis on the importance of gaining an appropriate level of assurance.

Risk Management in Information Security

Over the last decade the structures within which we approach information security have changed little. Information security has developed centred around the principles and constructs of risk management and continues to be so. Risk management has introduced a structured approach to managing to information security that has resulted in the consistent and effective comparison of risks across an organisation. This ultimately allows for the prioritisation of risks and the commensurate allocation of budgets and resources. It also importantly allows for information security to be embedded into the wider business structures and not treated in isolation. This permits an organisation to have a more holistic picture of the risks they face.

The core objective behind risk management is to identify those things that may harm you and put in place preventative measures. For information security this has resulted in a heavy focus on preventing attacks by implementing security controls. The standards adopted by the industry to manage information security reflect this focus on defence. Both COBIT and ISO27001, widely adopted standards, focus on the implementation of preventative controls. As a result, the focus of organisations' information security programmes has gone into building defences and implementing controls in a belief that this will deliver the required level of protection.

Under the current structures, other aspects of information security have not been entirely neglected. The requirement to prepare a response in the event of an attack is catered for through Disaster Recovery and Business Continuity measures, which have been widely adopted. Organisations also acknowledge the need to detect an attack and have adopted monitoring measures such as log analysis, security alerting and intrusion detection (IDS) systems. However, the focus and bulk of activity of information security programmes remains on preventing attacks by building defences.

The Need for a Resilient Approach

The world of information security is fast changing. New technologies emerge, threats evolve and different vulnerabilities materialise. The ability of organised criminal gangs and motivated attackers to target organisations via the Internet has increased to a level where they are capable of executing attacks that result in huge financial gain for them, while causing both financial loss and reputational damage for the targeted organisation. Even non-targeted attacks can have catastrophic consequences and result in down time and financial loss. The continual battle to maintain a fortress of defences is no longer sufficient to meet the current challenges faced by most organisations and it is becoming increasingly accepted that to do business in today’s world you cannot be 100% secure. In response, our approach to information security also needs to evolve.

If we accept that we can no longer be 100% secure, then we must develop our approach to include an acceptance that organisations will suffer compromises. Our approach should therefore be developed to enable organisations to withstand a compromise and therefore minimise the impact of a breach of defences. Our approach to information security should enable organisations to be resilient to successful attacks.

To combat the evolution of information security threats organisations need to build on their existing risk management structures to include resilience in their approach. Many organisations have already widened their approach to tackling information security threats to encompass not only defence, but also an ability to detect an attack, and react to it before finally recovering. This should be more widely adopted and consideration should be given to the aspects of information security that enable an organisation to cope when they are attacked.

A Resilient Model

A resilient model should encompass four core aspects that together form a more holistic approach to information security and afford organisations a more resilient approach to tackling information security threats. Obviously prevention is better than cure and putting in place defences against attacks should always be a priority. However, to build a resilient approach to information security the three additional aspects of Detect, React and Recover should also form part of an organisation's information security posture.

At a high level the model covers the following four core aspects; a defensive posture, an ability to detect, ability to react and finally an ability to recover.

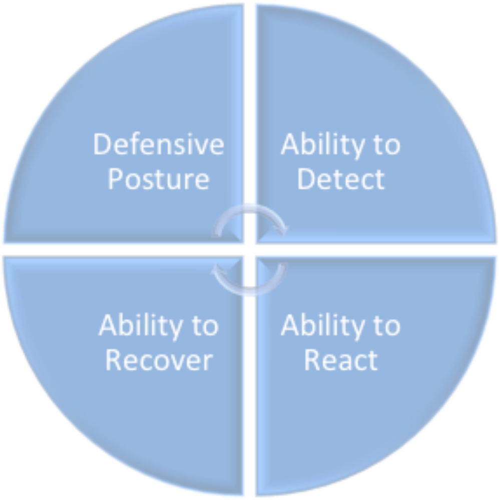

Defensive Posture

The main purpose of a positive defensive posture is to build and deploy security into systems and applications that make them robust to attack. Ideally security will be a core aspect of the initial design and will be built in to products and servcies from the outset. Changes within the threat landscape will need to be addressed through regular reviews and updates, so that controls can be maintained in line with the organisation's risk appetite.

Ability to Detect

The focus of this activity will be on the detection of malicious attacks. In order to effectively respond to an attack an organisation must first know that it is happening. Media stories continue to highlight that these measures are frequently lacking in organisations, as compromises are discovered months after the initial attack and often by external entities. The closer to real time an organisation can detect a compromise the more likely that its response to an attack will be effective. This approach is often referred to as protective monitoring. To be effective, protective monitoring controls must be used as part of a whole system, not in isolation of each other and must be actively managed.

Ability to React

This phase considers the organisation's ability to react to a compromise, often laid out in the organisation's incident management or response measures. Many organisations have these measures in place but they are rarely effectively tested and when they are it is often assumed that the attack has been detected effectively. In reality it is often not clear at the outset of an incident that the detection of a compromise has been truly effective and the real extent of a breach may take days if not weeks to determine. The ability to respond must therefore effectively engage with the ability to detect phase.

Ability to Recover

The recovery phase looks at identifying the required resource and effort needed and the potential implications on recovery ability and time. To manage this, many organisations currently undertake business continuity and disaster recovery planning activity. However, it is vital that this activity takes into account events that are likely to exceed a stated recovery time objective (RTO). The RTO is often seen as the maximum time that an adverse event can occur before the impact to the organisation becomes intolerable. While one off incidents and failures within a system can often be rectified within an established RTO, complex and sustained attacks against an organisation are likely to exceed this and as such an organisation must look beyond isolated metrics when planning for recovery. In the event of a significant compromise of an organisation's infrastructure, recovery time may be measured in months and will likely exceed an organisation's internal resource pool. Therefore, it is key that organisations recovery plans are based upon the current and anticipated threat landscape.

The key elements of this model are not new and many organisations undertake many if not all of these activities already. However, each aspect tends to be implemented in isolation of the others, rather than considering them as individual parts of a single system. As is demonstrated the latter three elements are dependent on the preceding function to work effectively. The system must therefore act as whole and organisations must take steps to gain assurance that it does so.

Requirement to Take into Account the Threat Landscape

Ensuring that information security measures are proportional and in line with an organisation's risk appetite is a key challenge. Risk management assists an organisation in determining a proportional response. By assessing the impact of an information security risk on the business, combined with the likelihood of attack an organisation is able to determine the amount of budget and resource it is willing to commit to tackle an information security risk. However, this approach frequently only focusses on elements of the risk equation that are known within the organisation, such as the impact on the business and the extent of the vulnerability. As a result, a key component of any risk equation is frequently missed out, the threat.

There are many definitions of a threat, but within the context of security risks we will use the following:

> *An actor that has the intent and capability to carry out an act that could cause harm to an organisation.*

A threat must possess both the intent and capability to carry out the act and these two elements can be used to assess the size of a threat to an organisation. In this context, the threat is a wilful actor that chooses to undertake the threat.

Threat forms the key part of any risk equation, as a threat actor must exploit a vulnerability in order to cause an event that has an impact on the business. Threat is a frequently ignored component though as the knowledge about the threat sits outside of the organisation it may impact[1]. However, information security threat actors can be grouped into broad categories that can be used in determining an approach to information security.

[1] More detail on threat in risk management can be found here: http://www.7elements.co.uk/resources/blog/threat-the-missing-component/

The Current Threat Landscape

The current threat landscape is reasonably varied for many organisations and the capability of potential attackers differs greatly. The following diagram demonstrates the differing threat capability against the volume of attackers.

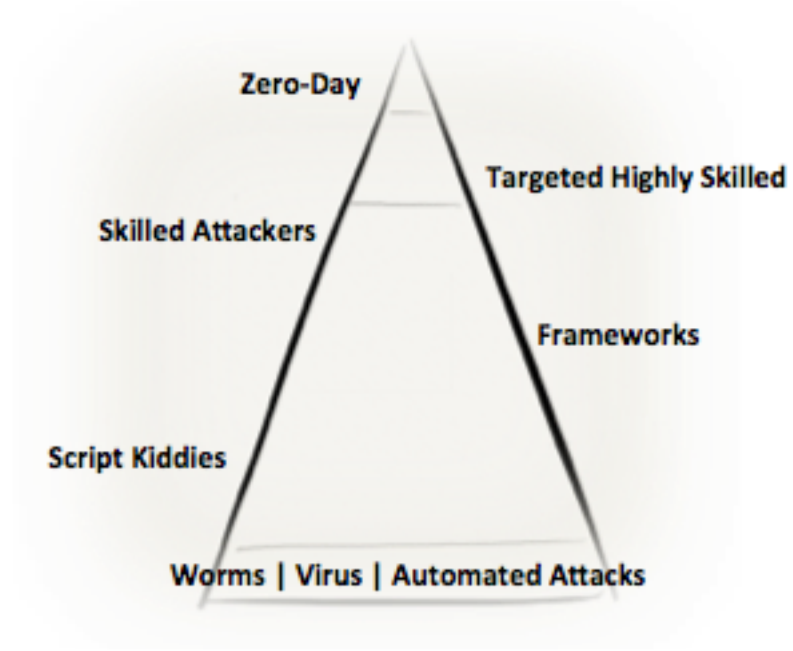

The pyramid represents the volume of capability, with those at the more technical end being of a smaller population when compared to the widespread availability of automated attacks (worms and viruses).

Fully automated attacks such as worms, viruses and automated scripts continually exist on the Internet and form part of the background noise that makes up Internet facing network traffic. These attacks target large IP address ranges looking for specific conditions or vulnerabilities to exploit.

Script Kiddies are low skilled attackers using pre-written simple attack tools. They lack the skill to make changes to code or conduct complex multi-stage attacks. This population of attacker can either be opportunistic in the same way as the automated attacks are, or more targeted, such as those conducting social or politically motivated attacks, often referred to as hacktivism.

Frameworks provide mid-level skilled attackers with access to hacking platforms with greater capability and functionality than those found in individual exploit code. Using frameworks requires an understanding of code, the ability to make changes and select the most appropriate tool for the objective at hand. Exploitation frameworks exist for both professional security testers as well as those designed and traded by malicious groups. Frameworks often have the capability to provide the attacker with post exploitation tools. These tools enable an attacker to move from one compromised system to other internal systems of the compromised environment, therefore extending the potential depth of the breach.

Skilled Attackers are capable of writing attack code and tailoring existing code to meet their requirements. Attackers within this grouping are often opportunistic in nature and will search for instances of a particular vulnerability to exploit. This could be part of hacktivism or just for the status that this brings.

Targeted Highly Skilled attackers are capable of writing attack code and tailoring existing code to meet their requirements. The key differentiator for this group is that they are motivated to target a specific organisation or group of organisations (such as the financial services) and are often seen as part of serious organised crime or more recently as part of economic or state sponsored espionage, sometimes referred to as 'Advanced Persistent Threats' (APT). Highly capable threat actors within both areas are highly organised, well-motivated and funded. This makes both of these actors a real threat. The key difference between an attack being classified as APT or cybercrime is the intention or driver behind the attack. At a high level, cybercrime has a focus on making money by stealing data to commit fraud. APT is looking to gain a commercial advantage through the information that they gain access to and are motivated to maintain access to compromised environments for long periods of time. More detail around this area can be found here:
http://www.7elements.co.uk/resources/blog/advanced-persistent-threat/

Zero-Day attacks relate to a new vulnerability that was previously unknown and has yet to be addressed via patching. These are often found through security research, both by malicious parties and security professionals. There is also a market for the purchase of zero-day attacks, therefore creating a financial barrier to access. In terms of use, highly skilled attackers capable of writing exploit code that make use of new zero-day vulnerabilities are able to deploy attacks at this level. This therefore limits the use to a specific population due to the technical knowledge required. However, it should be noted that over time a new zero-day may become part of a framework and then move down through simple scripts and become an automated attack.

It is also vital to understand that threat actors at a higher capability level can and will utilise lower skilled techniques to compromise an organisation.

Using Threat to Determine a Resilient Approach

Implementing a resilient approach to information security could be deployed across the whole organisation against all threat actors. However, this is not practical or realistic for organisations with limited funds and resources. A proportional response to the implementation of the resilience model is therefore required. Understanding the threat environment can assist an organisation in determining a proportional resilient approach to information security. An organisation must set and agree a risk appetite (determined through their risk management system) that will be used to determine their overall approach to information security and therefore the implementation of the model. This risk position, combined with an understanding of the threat environment will enable an organisation to make informed decisions on how to set the correct level of organisational response required for each of the core components of the model.

To implement a resilient approach an organisation must start with the threat actor groups discussed above and understand the overall risk appetite for the organisation. From the threat landscape described above, the main threat actors can be summarised as:

- **TA1** – Fully automated attacks (worms, viruses and automated scripts).
- **TA2** – Low skilled attackers using pre-written simple attack tools.
- **TA3** – Mid level skilled attackers with access to hacking frameworks and with greater capability and functionality.
- **TA4** – Opportunistic skilled attackers capable of writing attack code.
- **TA5** – Skilled attackers capable of writing attack code and motivated to target the organisation.

The threat actors and risk appetite will be used to determine at what level to set each element of the model. Organisations will determine which threats they wish to combat for each layer of the model. This will initially begin with the defensive posture. An organisation will decide at what level to set the defensive posture required by the organisation. This will then drive the level of robustness required for controls that sit within this area to meet that level of defensive posture. This is illustrated in the diagram below, where a defensive position has been adopted to address fully automated attacks and low skilled attackers, with the understanding that attackers with a mid level ability could potentially be able to breach the systems.

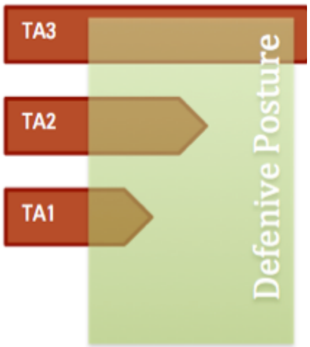

As such, this will then dictate the level of activity required within the detect and response control areas. Finally an organisation would then need to agree how robust a recovery position is required. The organisation's risk appetite will be used to determine how much resource the organisation is willing to commit to combatting the various threat actors against the risk the organisation is willing to hold.

Example of the Model in Action

This is demonstrated across the whole model in the example below.

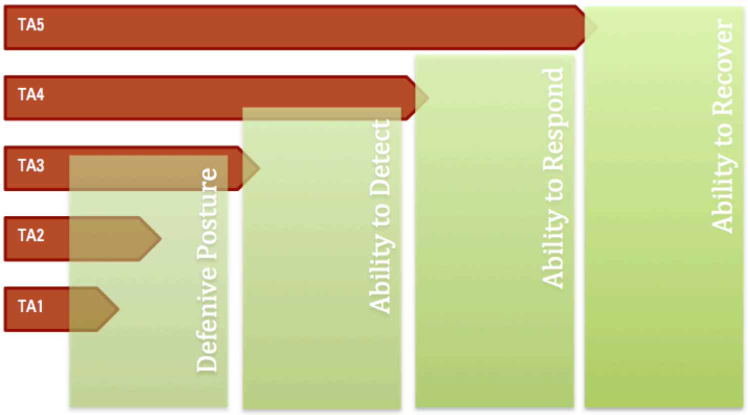

The defensive posture has been set in the example given above to defend an organisation against attack by the majority of mid level skilled attackers (TA3) down to fully automated attacks (TA1). This would have been a risk informed decision by the organisation and the required security controls and the rigour with which they are implemented are proportionate to the level of robustness required. Technical controls such as Web Application Firewalls and Intrusion Prevention Systems (IPS) configured to actively block malicious traffic sit within the defensive arena and should be used to aid in the overall defensive posture.

Controls such as Intrusion Detection Systems (IDS) and log analysis can then be used to enable the identification of mid level skilled (TA3) to opportunistic skilled attacks (TA4). It is accepted that highly skilled attackers (TA5) are likely to be able to evade these measures. As both controls do not prevent a breach, they sit within the ability to detect phase. Viewing controls within the context of the model clearly enables an organisation to understand their relative merits and overall support to a defence in depth approach.

By articulating this clearly, the organisation is then able to identify the need to provide a robust response to deal with any security breach (ability to respond). Finally given this model, the organisation has identified that targeted and highly skilled attacks may happen and that the ability to recover the situation will be the only approach able to deal with this situation. As such a cost and risk balance will have been articulated and the appropriate security spend made within each area.

With any new approach come challenges to change. Accepting that the defences that have been the focus of the information security programme will not always work can be a challenging position given the heavy financial investment and organisational buy-in required to deploy defences. Nonetheless, many high profile attacks continue to show that defences are not impenetrable. As a result, organisations need to consider having an ability to detect an attack. Once an attack is detected they will then need to be able to react to it and recover from it.

Assurance for a Resilient Approach

The core means by which organisations gain assurance of their information security posture has changed little over the last decade. The spectrum of security testing, from vulnerability scanning through to penetration testing, continues to be used as the principal method for assessing the effectiveness of our organisations' information security measures. However, this assurance focusses on assessing organisations' defences alone.

As the additional areas of detect, react and recover have evolved, some organisations have sought assurance that they are effective. Organisations conduct IDS reviews, or carry out business continuity exercises. Up until now though, this activity is limited and each area is tested in isolation. Yet, we expect the system to work as a whole when it is needed for real. The isolated assurance that is gained over these additional elements does not effectively assess how the system will act in its entirety and therefore whether an organisation will be able to effectively respond to an attack.

Security testing alone will not tell an organisation how well it will respond to an attack. The system as a whole needs to be tested. Security testing will of course form a key component to this but the additional measures also need to be encompassed in information security assurance. If we were to start to approach information security from the perspective of resilience, the system as a whole will be considered and the assurance measures undertaken will reflect the more mature information security posture adopted by most organisations. As such, it is now time to start testing for resilience and not just the 'security' of defences.

Resilient Security Testing

Resilient Security Testing is about providing assurance to an organisation from the point of view of business resiliency, which assesses an organisation's position against the four core areas as set out earlier. Using a tailored assurance approach that makes use of technical security testing, self-assessment, and audit style questions, each of the four core areas are assessed individually and a score derived. These discrete scores are then aggregated to create an overall organisational Resiliency Indicator (RI) score. The following high level model outlines such an approach.

Defensive Posture

An organisation should still strive to deploy and maintain a defensive posture. The review will assess how robust an organisation is to attack. This is captured through an assessment against the following two areas:

- **Technical Ability** – Using a manual testing approach complimented by automated and custom tools, the level of technical ability required to compromise an organisation is undertaken. The level of ability is based upon the threat actor groups TA1-4.

- **Depth of Compromise** – An important component of any compromise is the depth of access that is gained. During stage one and in defining an organisation's defensive posture, the depth of compromise is also assessed. Testing involves the exploitation of found vulnerabilities to gain further access to internal systems.

Ability to Detect

An organisation should have the ability to detect malicious activity. The review will assess the capability of the organisation to proactively monitor the environment and identify malicious activity. This is captured through an assessment against the following three areas:

- **Log Availability** – An assessment of the extent to which logging is in place, from having no logging in place through to logs being stored in a centralised area and away from the compromised environment.

- **Log Coverage** – An assessment of the extent to which logs provide coverage of the system, from limited log coverage that is restricted to minimal devices or services through to comprehensive logging in place on all services and devices.

- **Log Monitoring** – An assessment of the extent to which logs are monitored, from having no process in place for repeatable log monitoring or analysis through to a proactive (real time) review process for all log monitoring and analysis.

Ability to React

An organisation should have the ability to react effectively to a compromise event. The review will assess the capability of the organisation to react to a compromise event. This is captured through an assessment against the following three areas:

- **Incident Invoked** – An assessment of the extent to which an organisation is made aware of an incident, from being made aware of the incident after the event by an external agent, through to being aware of an incident in real time through internal controls.

- **Response Options** – An assessment of the extent to which an organisation responds to an incident, from the only course of action being to turn off all access to the network or systems, through to having the ability to trap specific network activity in real time, for example with the use of Intrusion Protection Systems (IPS).

- **Business Continuation** – An assessment of the extent to which an organisation can continue normal business activities, from a total cease of business activity until resolved, through to the negligible disruption to services or the business as a whole.

Ability to Recover

If an organisation is unable to recover from a compromise event, then the organisation will cease as a viable entity, therefore the review will assess the capability of the organisation to recover from a compromise event. This is captured through an assessment against the following three areas:

- **Recovery Time** – An assessment of realistic RTO based upon compromise scenarios, from an organisation not being able to establish an RTO, through to the time to recover measured in months, weeks, days or hours.

- **Resource Required** – An assessment of the extent to which an organisation would have access to the resource required to manage the recovery from an incident, from the organisation requiring specialised external resource (with no contract in place), through to being capable of sourcing resource as required.

- **Recovery Cost** – An assessment of the extent to which an organisation understands and has planned for associated financial costs based upon compromise scenarios, from a position of unknown cost exposure, through to having an estimated cost exposure within available budgets.

Conclusion

Little has changed in the way in which we approach information security as an industry over the last decade. Information security has developed centred around the principles and constructs of risk management and continues to be so, resulting in a focus on defensive controls. However, the threat landscape faced by organisations and the advancement of information security threats has seen the effectiveness of our defensive centred approach dwindle in the face of these challenges. As an industry our approach to information security should also evolve to meet these new challenges.

Building on our existing risk management structures, organisations should consider information security from the perspective of resilience. A resilient approach enables organisations to prepare for the inevitable breach of defences and respond effectively. In addition, by bringing a realistic assessment of the threat into a resilient approach organisations can ensure their approach is proportional and within risk appetite. By utilising this model and implementing a resilient security testing model, organisations will be able to ensure that the system responds effectively as a whole.

Consequently it is now time to start approaching information security from a resilience perspective and not just implementing defences.