Understanding cyber risk management vs uncertainty with confidence in 2017

"When I use a word,' Humpty Dumpty said in rather a scornful tone, 'it means just what I choose it to mean — neither more nor less." [1]

Prepared by:

Stephen Bailey, Associate Director

Jeff Bennison, Senior Consultant

Shanne Edwards, Executive Principal Consultant

Matt Field, Managing Consultant

Lee Hazell, Managing Consultant

Chris Hilder, Associate Director

Ted Ipsen, Regional Vice President

Patrick McCloskey, Managing Consultant

Tim Rawlins, Director

John Rostern, Vice President

Reuben Sinclair, Associate Director

Ollie Whitehouse, Chief Technical Officer

Table of contents

1. Introduction……..3

2. Defining cyber risk……4

3. Understanding the threat.....4

4. Understanding the vulnerability……5

5. Understanding the consequence……6

6. Understanding the impact……6

7. Understanding the likelihood……6

8. Understanding the risk rating……7

9. Determining inherent risk rating (RI)……8

10. Risk registry……8

11. Risk treatment……8

12. Risk appetite & risk tolerance……8

13. Determining residual risk (RR)……9

14. What direction should the industry take?.....10

15. References……11

1. Introduction

Every organisation faces uncertainty and this is often a key challenge in achieving its objectives. Naturally, much of this uncertainty comes from the lack of ability to accurately predict every future event. In attempting to plan for the future, organisations try to identify these potential future events, gather data to analyse them, make and communicate decisions based on this analysis and monitor them. Generally, we can define a potential future event that could affect an organisation's objectives as a 'risk' and the process of forecasting and responding to these potential future events as 'risk management'.

Many existing risk management methodologies attempt to improve the process of understanding and responding to potential future events. None are perfect (in fact multiple flaws can be found in all of them) but many provide a framework for organisations to more effectively and efficiently deploy resources in the pursuit of their objectives by considering risk. The first decision for an organisation will be how much risk management to do to strike the right balance between the amounts of effort spent on risk management activity versus the benefits brought about by the insight it provides. For example, endless analysis into understanding the interconnected likelihood of every risk might be beneficial but it will be costly. Equally, risk management activity for the sole purpose of demonstrating said activity to auditors might leave an organisation making ill-informed resourcing decisions.

Despite various industry and international standards, there is no universally accepted risk management method or universal acceptance of risk nomenclature. Some standards have clearly defined taxonomies and frameworks, while others are intentionally very loose. As such, a key part of effectively discussing risk management lies in first stipulating these definitions, ensuring effective communication and consistency. In our experience, organisations can do much more in this respect to bring greater structure and clarity to internal risk-management discussion.

For the sake of this discussion, throughout the remainder of this paper we shall use a set of definitions that have been used in many organisations to effectively manage risk. These may be different to your organisation but the concepts and insights will still apply.

There is no universally accepted risk management method or universal acceptance of risk nomenclature.

2. Defining cyber risk

Enterprise Risk Management (ERM) functions typically define risk ‘types’ within an organisation in accordance with the areas of the organisation that are best equipped to understand and manage them. They might define numerous risk types such as ‘financial’, ‘strategic’, ‘legal and regulatory’ and ‘people’. One such type that is becoming increasingly important is ‘cyber risk’ and the complexity of this often warrants its own risk management approach. This short whitepaper explores the concepts around the topic and suggests how organisations can evolve their thinking about cyber risk while also outlining some challenges.

When defining a cyber risk it’s critical to include the full scenario of the potential event in question, otherwise the risk will be ill-defined and cause communication, analysis, aggregation and comparison issues. As such, all risk descriptions should include at least a ‘threat’ exploiting a ‘vulnerability’ to cause a ‘consequence’. Each risk will then have an associated ‘likelihood rating’ and ‘impact rating’.

3. Understanding the threat

The threat is the cause of the event and in this sense it is often best to consider the threat actor as the ultimate cause of any event. They may be internal (dishonest/disgruntled/unintentional) or external (ex-employee/supplier/hacker/nature itself).

In discussing the threat we are primarily considering our threat actors’ capabilities and motivations. Are they a ‘lone wolf’ operator looking for glory? Is the intention to cause maximum commercial damage? Are the actors capable of attempting political change? Their capability may be both technical and operational and, therefore, considering their motivation will provide valuable insight into both the scenario in question and relationship with other risks.

An interesting factor to consider is that an external threat actor or adversary’s capability typically develops over time, either organically or through theft or acquisition. This was seen in the development of the former NSA exploit Eternal Blue which was weaponised and released by other threat actors, as the WannaCry ransomware and NotPetya data destroyer in 2017. The four main threat actor types are generally seen as cyber-criminal activity, lone or group ‘hactivism’, state sponsored actors and, finally, insider threats. However, a threat actor’s motivation can change in an instant due to a variety of factors, many of which are outside of an organisations control.

The whole concept of threat assessment is, unfortunately, an art not a science. That is you can give two people the same information and they will make different assessments, based on a number of factors including their experience and biases. The world of cyber risk and the threat assessments it relies on are no different. In the end the assessment of the threat against an organisation relies on knowledge, experience and a good guess.

4. Understanding the vulnerability

The vulnerability is a weakness that can be exploited by the threat. Looking at the vulnerability one generally considers an organisation's technical vulnerability. However, this view is both incomplete (given the unknown latent issues) and too shallow (given complex supply chains). The true cyber vulnerability of an organisation needs to consider people, processes and technology while also being able to factor in a percentage of 'unknown unknowns' until there is a more complete understanding of the environment.

To add to this is the density of vulnerabilities. If few are known in the organisation, this might make managing them viable, or if there a very large number then vulnerability management may become ineffective. This can be mitigated, however, by understanding how the primary threats against the known vulnerabilities can be controlled, in particular by grouping threats and vulnerabilities into manageable units.

Not all vulnerabilities will be introduced into an organisation by direct actions of that organisation itself, for example due to the adoption of a new technology stack. Through mergers and acquisitions, or as an indirect consequence of a new technology stack, new and unseen vulnerabilities can enter an organisation that may have a profound impact on its risk profile. Understanding the impact that the new vulnerabilities may then have, in particular the density of these new vulnerabilities, should be seen as a key element of the business decision process.

The vulnerability is a weakness that can be exploited by the threat.

5. Understanding the consequence

The consequence is a description of the result that would face the organisation should a risk be realised, i.e. the vulnerability exploited by threat. Often this result is a combination of costs (and benefits) broken down into categories, such as reputational damage, service downtime, information compromise, financial loss, etc. In the case of simple scenarios this is generally trivial to understand, however, with complex scenarios the impact can become less clear due to unknown dependencies.

6. Understanding the impact

The impact is a rating of the consequence. Organisations choose to apply qualitative rating categories or quantitative rating ranges when assigning impact ratings to consequences.

In our experience, organisations often attempt to quantify risks, to facilitate understanding, to permit prioritisation to be performed by non-subject matter experts and to allow a financial value to be applied to them. However, the world of risk in cyber security is evolving and at times very challenging in this respect.

The reality is that organisations quantifying risk using various subjective and objective inputs; we would posit that the methods for doing so need defining and continual refinement.

7. Understanding the likelihood

The likelihood is a probability rating of the threat exploiting the vulnerability in order to cause the consequence. Again this can be done using qualitative categories or quantitative ranges.

Organisations often look at the probability of a single cyber event happening in isolation to the rest of the information technology or operational technology systems. No other industry does this. They understand the connected nature of events that may flow from one to another or cascade to cause a multiplier effect so that the outcome is far more significant than the initial incident. If you take most current cyber risk assessments to an insurance actuary they may

challenge many of the assumptions it makes. Yet this approach has become the norm with regards to corporate governance and risk management of cyber.

For example, one can calculate what the risks of an office fire on the 3rd floor of an eight storey office block in London are. This is because data has been recorded and can be analysed by the actuaries. This well-developed data set enables the insurance premium to be calculated. While moving towards a more cyber centric set of risks to include in the risk management exercise requires new input vectors to be analysed and the greater speed and breadth of the attack vector to be taken into account, the basic premise of cause and effect remain the same.

While organisations like NCC Group have access to a wealth of experience and data, which helps us to make valid assessments of cyber risk, many organisations in the cyber world do not. At present most organisations simply cannot answer the question "what is the actual risk, one that can be quantified?" with anything other than a very loose estimation.

8. Understanding the risk rating

To allow for the prioritisation of risks, organisations will look to assign a single rating for a risk that encompasses the impact and likelihood. Again, this can be quantitative or qualitative but typically involves assigning a risk rating based on a matrix specifically designed around the organisations approach to risk. It is nearly always an oversimplification of the reality of a risk, as any single risk could have varying impacts, each with an associated likelihood, but very rarely is any complex statistical analysis performed. The picture becomes even more complex when the interaction between different risks are considered when determining the risk rating.

Despite its drawbacks, adopting a qualitative approach to rating risks has its benefits. Critically, it enables organisations to aggregate risk ratings and, therefore, provide much greater insight into threats, vulnerabilities and consequences, along with any other data collected about risks. For example, an estimate into the risk exposure reduction (of documented risk) that implementing a new process or technology will achieve in comparison to another control can be powerful insight to inform decision making. This does come with a warning though, as with any information, it is critical to understand the assumptions that have been made to come up with those estimates. They will be numerous.

The impact of a cyber event is manifested in the effect of that event on the business processes and functions of the organisation. Having a cyber event is not a risk in and of itself, instead rather the cyber event has a deleterious effect on the ability of the organisation to conduct business.

9. Determining inherent risk rating (R^1)

The inherent risk rating is the rating of risk in the absence of controls. This can be particularly useful in helping an organisation to understand its key controls and therefore ensuring that adequate testing and monitoring for key controls is in place. By looking at the difference between the inherent risk rating and the current risk rating attributed to each controls across all risks, an organisation can get an estimate into the total risk exposure reduction that is estimated for each control.

10. Risk registry

The risk registry is often an output of the risk assessment process and is a critical artefact in the process of risk management. The registry provides a consistent reference point as a repository for detailed descriptions of the risks specific to the organisation. As such, the registry must be kept up to date but also must be kept updated in conjunction with changes in the business model, along with the physical or technological environment.

11. Risk treatment

Treatment of risk is a governance function for the organisation. Organisations typically categorise the actions taken towards risk into four classifications:

1. **Accept**: Accept the risk as part of the business model of the organisation.
2. **Avoid**: Eliminate the risk through changing business processes, standards, practices, etc.
3. **Mitigate**: Take steps to mitigate the risk to the desired level that the business is comfortable with. Typically through the implementation of controls specifically designed to address the specific risk.
4. **Transfer**: Transfer typically relates to the assumption of the risk by a third party such as an insurer.

Treatment of risk is a governance function for the organisation.

12. Risk appetite & risk tolerance

The net effect of all treatments applied should equal the difference between the inherent and residual risk. There is often much debate over the definitions for risk appetite and risk tolerance. Typically, risk appetite is defined as how much risk the organisation is prepared to take in order to achieve its objectives. Risk tolerance is the acceptable deviation from the risk appetite. The amount of risk designated for mitigation becomes the basis for the design and implementation of controls.

In our experience, having a risk appetite definition can really help an organisation to work towards managing long term organisational risk exposure. Unfortunately, these are often statements that have weak practical application. However, NCC Group believe some of the best definitions will outline risk exposure targets over periodic intervals and take account of changes in business profitability. In some other, non-cyber risk management areas, an organisation should try to get as close to its appetite as possible, otherwise it may be able to allocate these resources in other areas more effectively. Cyber risk management is much more of an art, with less clarity and more estimation. Therefore, while deploying all resources towards a small proportion of risk is inadvisable, finding that a risk has been mitigated below the organisation's risk appetite is usually not an indication of inefficient use of resources.

13. Determining residual risk (R^R)

The residual risk (R^R) is that which remains after all treatments other than acceptance have been applied. The risk acceptance threshold of the organisation should ideally be >= to R^R.

For risks mitigated through the implementation of controls, the value of risk reduction is the product of the design (E^D) and operating (E^O) effectiveness of the controls. As noted earlier the risk registry provides the reference point for the design of controls to mitigate a specific risk. The audit of design effectiveness validates that the controls are, in fact, designed to mitigate the risks described in the registry. However, this is a 'test of 1' or 'control walk through' that does not provide assurance beyond a point in time.

The operating effectiveness of controls is measured through an audit/assessment of evidence produced by the controls over a period of time (reporting period). This process provides assurance that the controls operate effectively to mitigate the associated risks during the specified time period.

Beyond the difficulties of defining it there is another challenge with cyber risk, in that at times it is like water. That is to say it will find its way in wherever there is a crack in the defences, be that across technology, people and process. The likelihood of a very specific risk scenario materialising is likely infinitesimal but, when combined, the likelihood of one of them leading to a specific consequence can be considerably higher. As such, in order to minimise inherent risk within an organisation (even without any external threat) there are some quite simple base principles that should be adopted around people, process and technology. These will be discussed in a future NCC Group whitepaper.

14. What direction should the industry take?

Collectively industry and government should be doing more to gather data and statistics to form the basis of a heuristic approach to analysis, while aligning with wider business thinking on probability as a key factor in determining risk.

As an industry we need to consistently inform those we serve on the basic tenet of risk assessment as the driver to ensure the correct to controls are implemented. Risk assessments provide the basis for the controls assessment process by establishing the risk frame for the organisation and identifying those risks specific to it. This then leads to the design and implementation of controls, which is then the subject of the controls assessment. Failing to provide this understanding across the industry severely impacts our ability as a profession to positively affect outcomes with regards to cyber security. With one of the primary roles of the information security profession being to support the business, providing the relevant frameworks to allow best practice decisions to be made is vital if we are to protect our clients, both internal and external, public and corporate.

There is a dearth of real actuarial data on which to base these types of empirical probability ratings. However, even in organisations that try to be quantitative in their cyber risk analyses, if you look deeply enough, at some point they fall back on a quantitative element that involves a person making a determination on a scale of 1 to 1+N.

NCC Group expects that as more clients turn to cyber insurance as part of their overall risk mitigation strategy, the cost of underwriting the exposed risks will focus businesses on developing the correct mitigation plans. These will be based on the risk and control methodologies developed in conjunction with increased knowledge of the threat actors and vulnerabilities. In much the same way as advances in technology within vehicles (that have both direct and indirect impacts on safety) enable a reduction in vehicle owner/user insurance costs. Conducting well managed risk assessments that drive well defined and implemented controls will lead to reductions in operating costs as defined by that insurance industry [2]. And, therefore demonstrating that security is there to support the business.

15. References

[1] 'Through the Looking Glass and what Alice found there' - Lewis Carrol 1871

[2] "Cost of car insurance to plunge with rise of driverless vehicles" - https://www.ft.com/content/8718f37a-21d9-11e6-9d4d-c11776a5124d