# 35 VIEWS OF CYBER RISK

John Donald

john.donald@axiscapital.com
Winter 2019

AXIS

www.axiscapital.com

# Foreword

**Dan Trueman**
Global Head of Cyber and Technology

The need to understand the world we inhabit is not a new issue, but rather the eternal problem. However, what has changed is the pace at which our understanding needs to develop because the pace of change has accelerated and shows no signs of slowing up.

This constant change leaves us in a volatile, uncertain, complex and ambiguous world – often outlined under the acronym VUCA. This is obviously a challenge to those striving for clarity and understanding, a VUCA world cannot be understood with just one simple model, it cannot be evaluated with just one measurement and it cannot be communicated with one simple narrative.

This is the problem we face when trying to evaluate, measure, assess, communicate and respond to the cyber risks we face. The cyber world is constantly changing, seemingly locked in a spiral of increasing reliance on technology that appears ever more vulnerable to outside influences.

We can never claim to fully know the cyber world, but we do believe that it can be better understood. Understanding needs to come from a diverse range of thinking and draw on ideas from wide schools of thought; in a complex world, diversity of thought trumps single models every time. The cyber environment is not a linear one and thus can be much better understood when different parts of the picture come together to form one wider illustration. We need to combine knowledge from several fields and use it to improve our understanding – combinative thinking.

It is such combinative thinking that my inspirational colleague John Donald has brought together here. Taking his own inspiration from the master print maker Katsushika Hokusai, who himself tried to explain the changing nature of his own world by producing 36 different views of the same object, we have tried to bring together the 35 views that we see as we traverse around the cyber risk environment. This journey is never over but I hope it gives food for thought and helps explain away at least some of the volatility, uncertainty, complexity and ambiguity.

# 1. Hokusai’s great wave

If you browse the souvenir shops at Tokyo’s Narita airport you will find yourself surrounded by images of the “Great Wave off Kanagawa” on T-shirts coasters, mugs, fans and jigsaw puzzles. This image is a global icon (you will also find it on tea towels in the British Museum), a universally recognised symbol of Japan. What is less well known is that it is actually an image of Mount Fuji. Look closer and you will see Mount Fuji’s conical tip poking up above the horizon in the background. It forms part of a series of prints by Katsushika Hokusai in 1830 called “36 views of Mount Fuji”. Each image in this series shows a snapshot of daily Japanese life in the foreground with the mountain somewhere in the background.

## Different perspectives, same object

Fuji-San, to give it its proper honorific title, is not just the tallest mountain in Japan but also a sacred site. In the Shinto religion it symbolises Japan’s cultural and spiritual soul. Hokusai, in his series of prints, is saying that to properly comprehend its mystic significance it must be approached obliquely. It must be viewed from all angles. So the 36 different views of daily life are all different perspectives of the same thing. The spiritual soul of Japan as symbolised by the holy mountain shown through 36 different lenses.

* Security Insurance Cyber

(* **S** for security, **I** for cyber insurance and **C** for cyber and information technology)

## Three plates meet

And so to cyber risk. The active volcano of Mount Fuji sits at the intersection point of three tectonic plates. The Eurasian Plate, the North American Plate and the Philippine Sea Plate which form a triple junction at this spot. We can read across from geological risk to the cyber realm. Cyber risk also sits at the nexus of three ‘plates’: the three different disciplines of information technology, security and insurance. Each of these three subjects has its own conceptual framework, its unique acronyms and specialist expertise. To truly understand cyber risk, we must slowly circle through each zone of this triad, gaining a 360-degree view of the half-glimpsed entity at the centre. The three zones are marked by a symbol (*) in the top right-hand corner to help the reader navigate this triune journey; each picture a fragmentary shard in a multifaceted whole.

Hokusai portrayed Japan as a series of 36 prints. We offer our view of cyber risk as a set of 35 views, one less than 36, in due deference to the master.

# 2. The four quadrants of risk

"There are things that we know we know. We also know there are known unknowns; that is to say there are some things we know we do not know. But there are also unknown unknowns, the ones we don't know we don't know"

That was Donald Rumsfeld trying to explain the lack of nuclear weapons after the invasion of Iraq to the assembled press corps in February 2002. The audience was completely bamboozled. Understandably so, because Rumsfeld was using the word "known" to mean two different things; known meaning 'aware of its existence' and also known in the sense of being a quantifiable or predictable entity. Separating out these two senses of the word gives us the four quadrants in the diagram to the right. You will notice that Rumsfeld forgot to mention the fourth quadrant: unknown knowns. These are things your colleagues know, but you don't. Ironically, it was this risk quadrant that cost Rumsfeld his job when he was forced to resign because of the Abu Ghraib prison atrocities.

## Cyber risk moves anticlockwise

Applying these risk quadrants to a cyber context, we can see that cyber risk travels in an anticlockwise fashion. Starting in quadrant three, unknown unknowns – sometimes called 'black swan' events – are events that have never happened before, so we are completely unaware of them. This was generally the case for cyber risk 20 years ago. But once the first attack of a certain type occurred, let's say ransomware for instance, it moved into quadrant two as a known unknown. Awareness dawned that that type of attack vector existed, even if it was hard to quantify when or where such an attack might take place.

Today we stand on the threshold of quadrant one, with cyber-attacks happening so often that there is a reasonable historic data set with which to quantify the risk. So, there is at least some small upside to the increasing frequency of cyber breaches; actuaries have a richer statistical input for their models enabling better pricing of risk.

In summary, each new emergent cyber threat moves from quadrant three to quadrant two and then on to quadrant one. From unthinkable, it passes through uncertainty to end up as commonplace.

Astute readers will notice we have ignored quadrant four, as Rumsfeld did. Fear not, we pick up this thread in view 24.

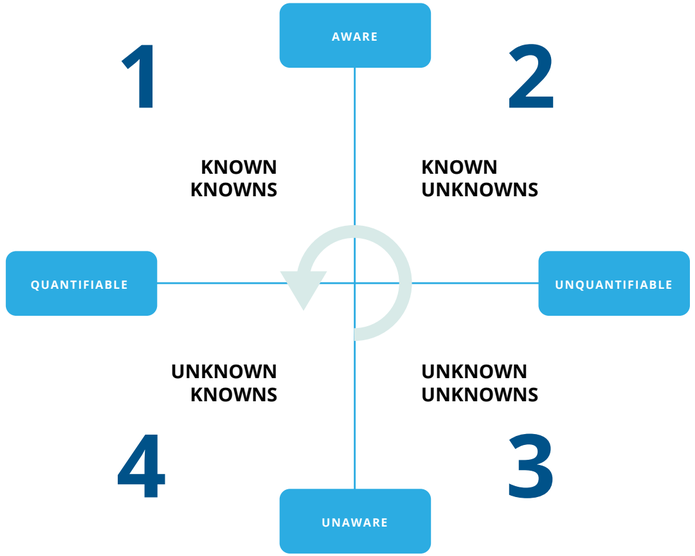

# 3. The four quadrants of security

Security is defined as the degree to which your assets are resistant to threats from adversaries. If we map these elements into our risk quadrants, we end up with the diagram to the right.

**Adversaries:** In the cyber world, we never really know who our adversaries are. They hide behind multiple nodes, proxies and relays so we don't even know where they are. It is unlikely that we will ever know so they belong in quadrant three: the unknown unknowns.

**Threats:** We are aware of the attack vectors and methods that cyber criminals use (see views 17 and 20). That much is known. But it is still hard to quantify when an attack might occur and who will be the target. So threats belong in quadrant two: the known unknowns.

**Assets:** Most companies are aware that they have data that would be valuable to their competitors. These are assets that need to be protected. It is also probably true that many companies may have not done an extensive audit of the data assets that they hold or fully defined their business critical 'crown jewels'. Despite this, assets belong in quadrant one as known knowns, if only in a partial sense at present.

**Impact:** The impact of a cyber incident belongs in quadrant four. Most companies are unaware of what a cyber incident might cost until they actually suffer from one. Much of this information could have been gathered beforehand. It is quantifiable (see view 24 for a suggested model) but often remains unknown.

## Practice (almost) makes perfect

One piece of advice that all security experts agree on is that rehearsing the corporate incident response plan through table top exercises can have a major positive impact. Practice may never quite make perfect, but it will substantially improve cyber resilience. In a table top exercise, crucial information known to one department head becomes known to all. Does the sales director know how long it might take for the IT department to rebuild core systems after an attack? Does the legal department understand the urgent requirements of the corporate communications department in a crisis? What policies can be agreed calmly beforehand rather than hastily thrashed out in an emergency situation?

Much that is unknown can be revealed when practicing a crisis response. Logically speaking, quadrant four is the most productive place to invest time and money. It is far easier to raise awareness than to try to quantify the unquantifiable.

# 4. Cyber risk in context

In February 2019, the great and the good gathered in Davos at the World Economic Forum to discuss the current state of the world. They examined the threats facing the world economy and ranked them according to their likelihood and impact. So, for example, Weapons of Mass Destruction (WMDs), which were a hot topic a decade previously, were not seen as a pressing concern. Though nuclear bombs going off would clearly have a big impact, it was thought unlikely that this would actually occur. Hence their position in the top left corner of in the diagram to the right: high impact. low likelihood. So, what were the issues of greatest concern? Look in the top right-hand corner. You will see the dots up there are either orange or light blue. In other words, the two issues of greatest concern were climate change and cyber risk.

The recognition of these two risks has been fairly recent. In 2015, cyber-attacks did not even make the top 10 list in terms of impact and five years before that the agenda was dominated by collapsing asset prices and failures in global financial governance. Today both of these threats are given their due prominence in the corporate risk registers when companies file their annual reports with stock exchanges.

In insurance terms, this translates into two main product lines – property and cyber. Climate change is driving an increase in natural catastrophes such as wildfires, hurricanes and tornados which in turn is driving up property damage claims. Cyber-attacks have been growing exponentially and show no sign of peaking yet.

## Systemic vs specific risk

We should note, however, that despite having similar prominence in terms of likelihood and impact, property and cyber have an important difference in risk characteristics. This hinges on the difference between specific and systemic risk. Property risk is largely specific and based on geography. The likelihood of three buildings – one in New York, one in London and one in Tokyo – collapsing at the same time is zero. But three computers in each of these locations could easily become corrupted at the same time if they are all connected to the same network. Specific risk can be reduced through diversification, systemic risk cannot. Since the internet connects all computers globally, systemic risk is prevalent in the cyber world.

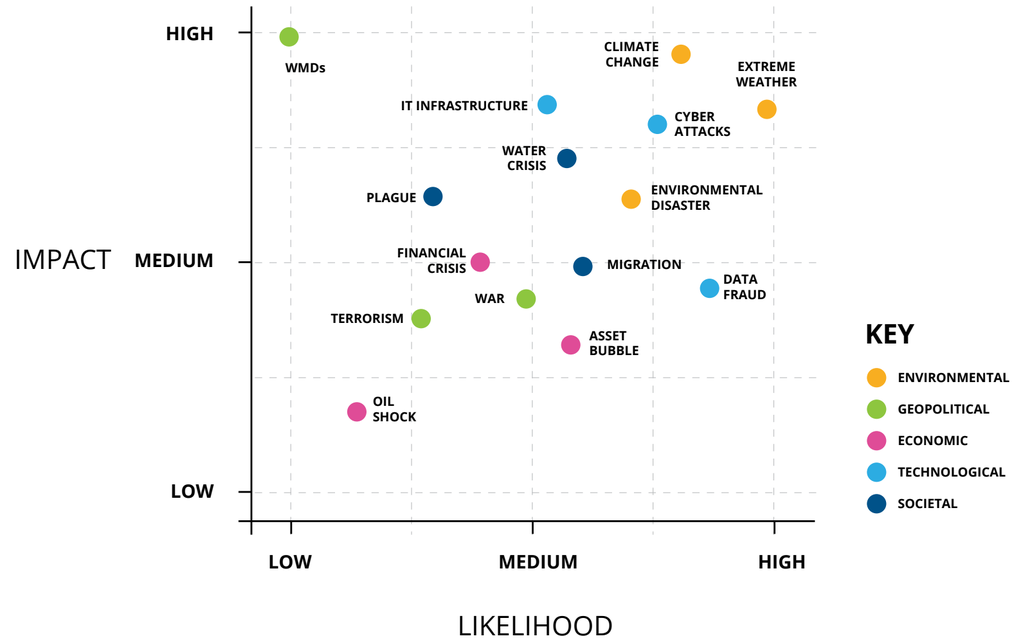


Source: WEF Global Risks 2019

# 5. Systemic risk: The catataxic shift

Earlier, in view four, we explained the difference between specific and systemic risk. The former can be mitigated by diversification. Strictly defined, systemic risk is ‘undiversifiable risk’, in other words, the risk that cannot be mitigated through diversification, though in an insurance context the term is used more loosely. But to divide risk into these two categories is misleadingly simple because it assumes a static environment. As we know, particularly in the cyber world, the environment is always changing (see view 13). So, let’s look at a dynamic model in which systemic risk can emerge over time.

Imagine an experiment in which there are 20 buttons and 20 threads. The ideas behind this approach were first put forward in two papers: On Random Graphs by Erdos & Reyni (1958) and The Origins of Order by Stuart Kauffman (1993) although they were applied in a biological context to explain how life might have originally emerged on earth.

## A sudden change

At the start of the experiment, one thread is randomly attached to two buttons, one at each end. After five threads have been attached randomly in this way, if you were to close your eyes and pick up a thread at random, you would probably pick up two buttons with it. Continuing in this way, after you have added 10 threads picking up a random thread might also lift up three or four buttons. You can see that at some point picking up a thread will pick up all the buttons. The question is when does this point happen?

The answer is when the ratio of buttons to threads is 0.5, in other words very soon after 10 threads have been added to the 20 buttons. Please refer to the original papers for the mathematical proof. You may think “more threads, more buttons, so what?”. The important thing is that it happens very suddenly, in a step function change. At one moment you are picking up a few buttons, the next moment you are picking up the whole thing. This sudden change - the catataxic shift - marks the point that you can no longer view the system as a set of independent elements. You must view all the elements as a single entity. This is analogous to a phase change. In a gas, you can view the particles as independent elements. When the temperature changes and the substance becomes solid you must view all particles as a single object.

Systemic risk arises when the degree of connectivity reaches a certain point. You may have assumed you have a diversified portfolio with all your eggs in different baskets. But a change in the environment, such as temperature or the ‘threads to button ratio’ in these examples, forces you to realise that all your eggs are in fact in the same basket - one big global basket that we call the internet.

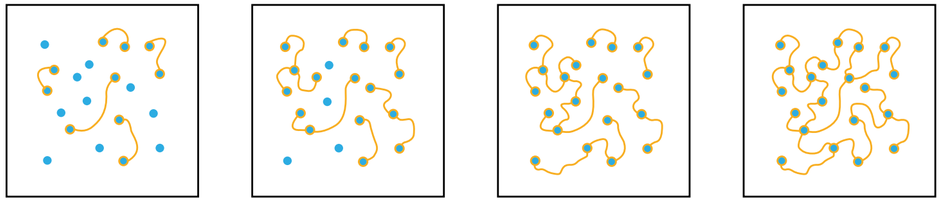

5 THREADS: 20 BUTTONS | 10 THREADS: 20 BUTTONS | 15 THREADS: 20 BUTTONS | 20 THREADS: 20 BUTTONS

See Erdos & Renyi
On Random Graphs (1958) &
Stuart Kauffman
The Origins of Order (1993)

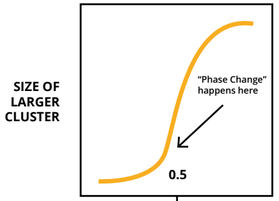


THE **CATATAXIC** SHIFT

A step function change occurs when the ratio of threads to buttons reaches a half. This is the moment at which systemic risk emerges.

# 6. Systemic risk: Industry connectivity

Following on from the 'buttons and threads' model of systemic risk (view five), let's look at connectivity issues and risk aggregation across industries. The diagram shows a simplified layer model of an IT system from hardware at the base to people at the top. At each layer in this pyramid there are connectivity pathways that can bind companies in a particular industry together. This implies that systemic risk is rife throughout different industry sectors at all stages.

**Hardware:** Companies in a particular industry normally use the same hardware. A good example is the point of sale card readers used by retailers. The notorious Target breach that exposed 40m credit card details in 2013 was based on a vulnerability in the RAM memory of the credit card readers. Note also that the machine tool industry is very fragmented. A small German mittelstand company can often have a 100% market share for a particular type of precision milling machine used in an industry vertical. Now that they are being connected to the internet through supervisory control and data acquisition systems and the Internet of Things (IoT), a new vector of systemic industry hardware risk is emerging.

**Networks:** Cloud service providers, internet service providers and the national telecom infrastructure in general all have the obvious risk of being a single point of failure. Amazon Web Services dominates the cloud services market with a 50% market share [1]. Adding the next three biggest providers, Microsoft Azure, Google Cloud and Alibaba Cloud makes this go up to 85% [2]. A major failure at one of these four could cause significant disruption across all industries.

**Application software:** Just as machine tool manufacturers dominate in micro-specific niches for plant and equipment, application software companies tend to dominate in particular small-scale industry verticals. So, for example, a software package specifically designed to help dentists run their practice could have a huge market share amongst dentists but not amongst doctors who have their own preferred software provider. The smaller the size of the specialist niche, the greater the likelihood that a single supplier will dominate it.

**Websites:** Websites are natural industry aggregators. A widely read blog on an Industry Association website is a natural target for a 'waterhole' attack. In the 19th Century, big game hunters would wait at waterholes to pick off the animals that gathered there to drink at dusk. In the 21st Century, hackers do the same thing at popular industry websites.

**People:** Industry conferences are another easy target for waterhole attacks. A list of the emails of all attendees is fairly easy to acquire, providing an excellent starting point for an industry specific phishing campaign.

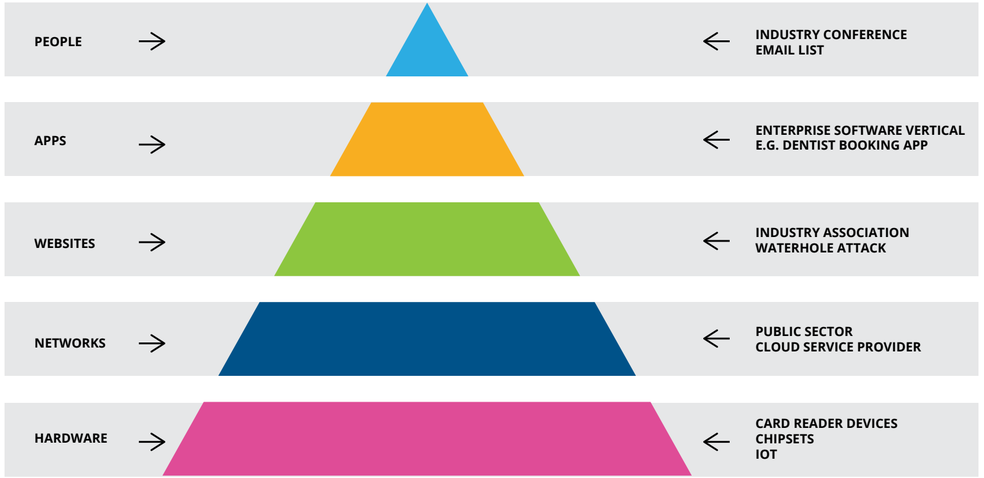

# 7. Systemic risk: Complex adaptive systems

In this year's Davos meeting of the World Economic Forum (view four) the top two threats to the global economy were deemed to be climate change and cyber-attacks. Climate change from a coverage perspective translates into property insurance. So, let's compare and contrast these two types of risk: natural catastrophes and cyber catastrophes. John Foster, in a paper published in the Cambridge Journal of Economics in 2005 (vol 29, 873-892), developed a scheme for describing orders of complexity which has four levels as follows:

**First order complex systems:** These are systems with purely physical components. A tornado, for example, is a complex structure that requires sophisticated non-linear mathematics to model its flow. A supercomputer can crunch the numbers in a natural catastrophe model to estimate the damage caused when a tornado hits a set of buildings in a specific place. But the tornado and the buildings are both physical structures, so this problem is only of first order complexity.

**Second order complex systems:** Second order complexity shifts from the physical to the biological realm. In these systems, there is the extra element of adaption as Darwinian evolution becomes a factor. So, a tornado will not become stronger as a result of buildings becoming stronger, but a beaver's teeth will evolve over time to deal with bigger trees.

**Third order complex systems:** Third order complex systems are one step more adaptive as they involve creativity rather than just natural selection. The adaption is directed by an intelligence that actively tries to change the environment to benefit itself. Humans first using tools marked a transition point from second order to third order complexity. Second order systems progress under the influence of feedback loops. In third order systems, there is a feedforward effect too: an ability to actively create a different future by design.

**Fourth order complex systems:** In fourth order systems, there are two intelligences rather than one. The environment is one of anticipatory change. The protagonist and the opponent are trying to predict the other's moves before they have even happened and retaliating in advance of that. All markets are fourth order complex, where the secret to success is not to figure out what something is worth but how your opponent will value it instead. You are not judging a beauty contest, you are judging the judges of a beauty contest.

The cyber realm is fourth order complex, with hackers anticipating defensive responses and vice versa. So, the key difference between cyber risk and natural catastrophe risk is that the former is fourth order complex while the latter is only first order complex.

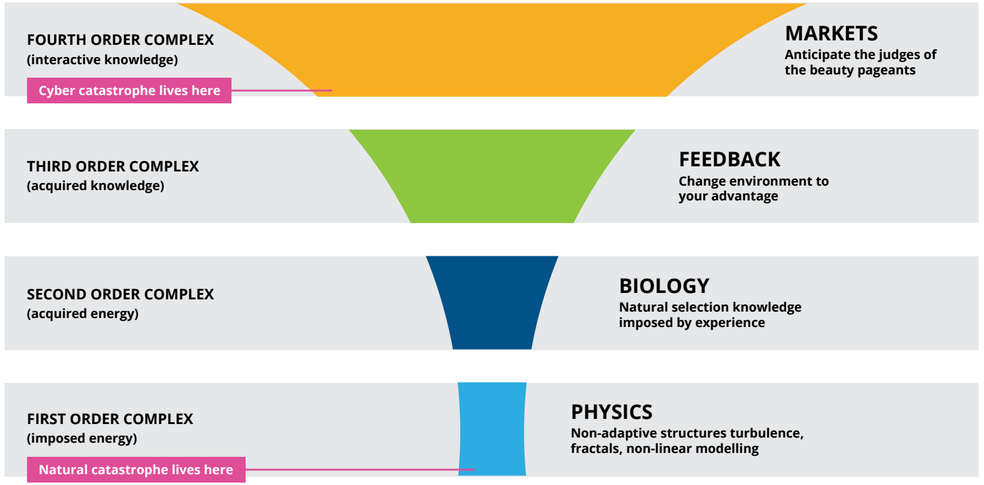


(after John Foster, 2005, 29, 873-892 - Cambridge Journal of Economics

# 8. Medieval castle model

One way to conceive of cyber security is to use the analogy of a medieval castle. Then castle walls and the moat correspond to the firewall with the drawbridge allowing permitted visitors inside. The sentries patrolling the ramparts represent the anti-virus software on the lookout for suspicious events. A Distributed Denial of Service attack would be analogous to a siege engine lobbing boulders.

As in all good Hollywood movies there is a secret way into the castle - maybe a tunnel, a postern gate or a small iron grille over a sewer - which is known only to the people who built it. This is the software 'backdoor' created by the original coders of the system which still exists but has been forgotten about.

## Send out the scouts

Outside the castle walls, an invisible enemy is lurking behind some hills. It's probably worth sending some scouts out to find out what they are up to. In the cyber world this is called cyber threat intelligence. This involves hiring some experts to search the dark web looking for indications that you might be a target of a planned attack. Also bear in mind that the wagons bringing essential goods into your castle are a vulnerable element. In the Hollywood movie, the enemy hijack the wagons and enter the castle in disguise. In a cyber context, this is known as a Trojan attack, named after the Trojan horse that concealed Greek warriors in the Iliad.

## The king in the keep

Inside the castle, there are locked storerooms and maybe some golden treasure in a strong room. The locks on these rooms are called endpoint security software in an IT context, providing an extra level of protection to individual devices. Sitting in the keep is the King who seems impressively well protected by all these defensive layers. Indeed, it would take a powerful 'brute force' attack to break through them all. In the cyber world, a brute force attack is a trial and error method used to crack passwords using automated software to test every possible combination in turn.

Brute force attacks take a lot of time and computing power, just like the siege engines that throw boulders to knock down the castle walls. But there are simpler ways of getting the King to surrender other than destroying the ramparts. In the Hollywood movie, it would be a non-physical attack of some sort - a trusted courtier turning traitor, a letter that causes a change of mind or a psychological trick that completely saps morale. In the cyber world this would be a phishing attack - an email that fooled the CEO into giving up their password.

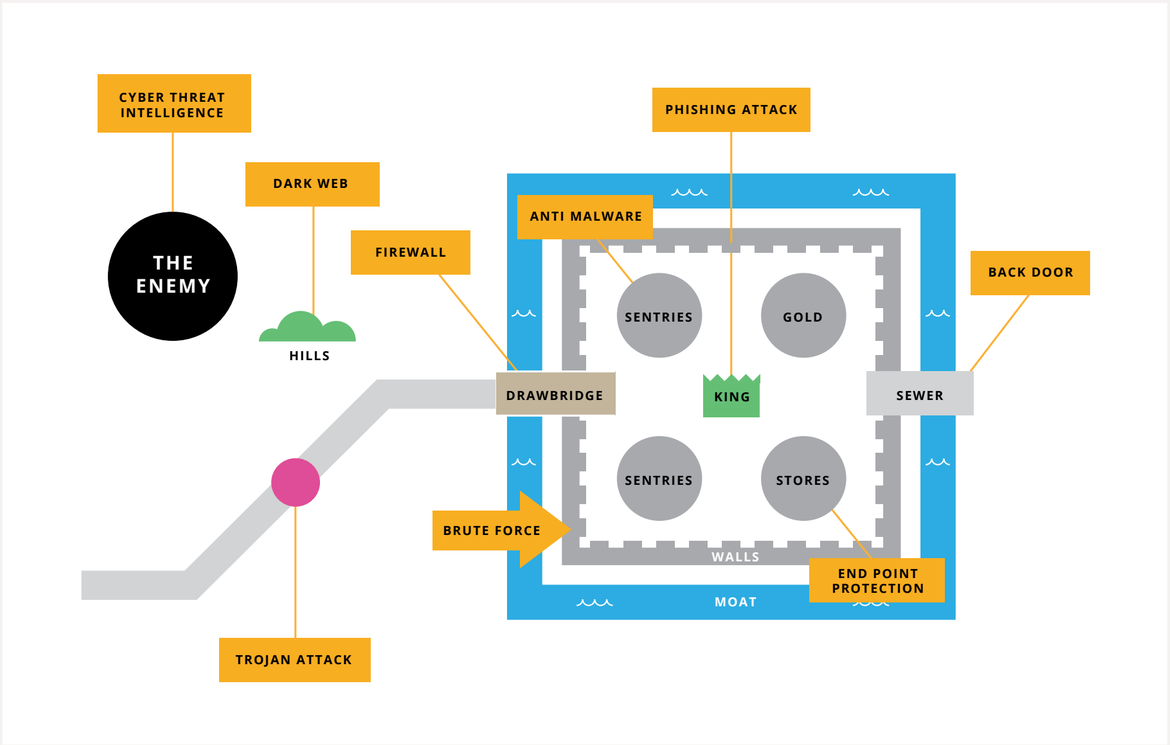

# 9. Where is the wall?

The castle model of cyber security is a useful way of illustrating some simple cyber security concepts, but it has a major flaw which is where do you put the wall? As the world becomes both more mobile and more interconnected, it is increasingly hard to draw the line between inside and outside from a system standpoint. This issue, known rather clumsily as de-perimeterisation, is a big challenge for security professionals.

Most cities in medieval times were surrounded by a wall for their protection. But as global trade flourished, these walls were torn down to improve the flow of goods and services. In London and Paris this happened in the 18th century, in Beijing not until the 1950s. For similar reasons, over reliance on perimeter security and a binary distinction between 'us' and 'them' is becoming an outmoded approach in the cyber realm.

The drawing of this dividing line can be framed as an attempt to find a balance between business drivers and security concerns. In the majority of cases, it is the business drivers that tend to win in the end.

## Inside or outside?

The diagram to the right illustrates some of these issues. It is common for companies to use cloud-based software for client relationship management or accounting. Salesforce and QuickBooks are popular examples of these 'software as a service' (SaaS) packages. However, it is debatable as to whether they should be inside or outside the corporate perimeter. Similarly, most companies use contractors and third parties for software development or for website design; are they insiders or outsiders? Mobile working exacerbates the perimeter problem. Laptops and smart phones used for both office and home purposes blur the dividing line between the internal and the external zones. This issue is described by the acronym BYOD which stands for "Bring Your Own Device" to the workplace.

Lastly, one of the vulnerabilities most often exploited by penetration testers is the physical realm, particularly where an organisation has multiple office locations; it is difficult to maintain the same level of security across all of them. Most offices rely on third party contractors for building management services like cleaning staff and underfloor wiring. Should these individuals be under the same vetting and security regimes as employees?

The conclusion from all these examples is that it is almost impossible to draw a clear line between an organisation's internal and external zones. The concept of a 'castle wall' is a useful but outmoded metaphor. There are other more apt analogies for cyber security (see view 11).

# 10. Why the squid lost its shell

It is not widely known that squid used to have shells. Ancient cephalopods in the Jurassic era, the common ancestor of modern octopus and squid, were creatures like today's nautilus. They relied on a large external shell for defence. However, the seas became more acidic, weakening shells made from calcium carbonate. This meant that the squid gradually evolved three other mechanisms for self-defence. These were:

- **Intelligence** – squid and octopus are amongst the most intelligent creatures in the ocean. If in doubt, look up the video on YouTube of an octopus opening a screw top jar to get at the food inside.
- **Camouflage** – squid have special pigment cells called chromatophores in their skin which enables them to change colour and blend into the background.
- **Agility** – squid are very rapid swimmers using a form of jet propulsion. They fill their internal cavity (unconstrained by a shell) with water and then expel it quickly in a jet enabling them to leap upon their prey.

Looking at the current state of cyber security investment, we are still in the Jurassic era. Most of the spending, some 80%, is on prevention which is a defensive shell strategy (see view eight for the castle wall analogy). In the future, cyber security investment is expected to become more evenly spread with big gains in other areas such as response and monitoring.

This change can be likened to the squid losing its shell. Monitoring is, in effect, a form of **intelligence** gathering. Likewise, investments in improving response time is analogous to the squid's **agility**. The speed with which an organisation responds to a cyber-attack is a critical factor in determining the degree of eventual damage. Investing in table top exercises to rehearse incident response plans is often money well spent.

What about **camouflage?** How does that relate to the cyber realm? Camouflage is the ability to blend into the background; the art of not standing out as an obvious target. Another critical factor in cyber defence is the speed of the patching cadence. Companies who neglect to install patches to upgrade their software to the latest versions are extremely vulnerable to hackers. It's the equivalent of walking down the high street in antique Victorian clothing; you would clearly stand out from the crowd. So patching discipline is the equivalent of a squid's camouflage, an effective way to avoid becoming a target.

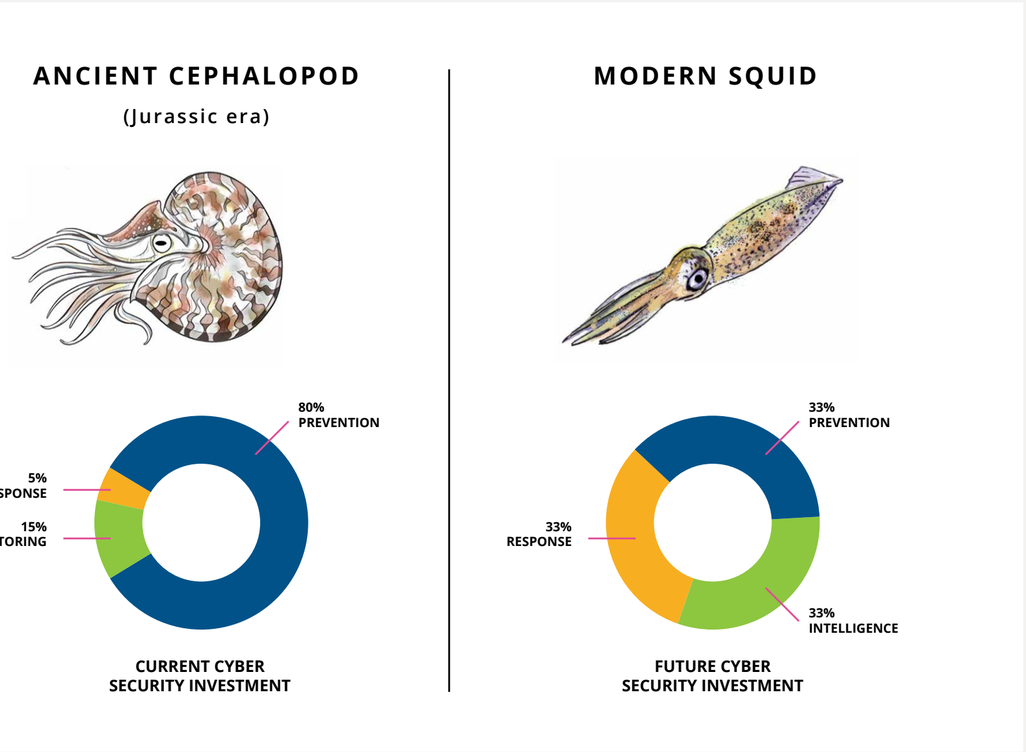

# 11. The immune system model

A more useful model than the static castle one (view eight), is a dynamic model based on the human immune system. The squid analogy (view 10) explained why reliance on an external perimeter for defence is outmoded (see also view nine). The immune system model presupposes that systems are constantly under attack and so the focus is on the speed and effectiveness of the counter attack.

The immune system has two parts – an innate system at the initial stages which is the same for all attacks and an adaptive system that kicks in at a later stage which is a bespoke response to that specific attack. In humans, the innate system consists of barriers to infection such as skin, mucous membranes, saliva (which has antibacterial properties) and the tonsils in the throat. These are designed to deter and delay infection from germs.

## Detect, respond, recover

More interesting is the adaptive immune system. Once a virus enters our bodies it causes local inflammation which is the first warning sign that something is wrong. This in turn acts as a trigger for the production of white blood cells which then go on to produce antibodies that bind with pathogens and killer T cells which destroy the virus. Once these lymphocytes have done their job the body is able to recover.

The five steps in this process are exactly analogous to the five steps required in a cyber incident response plan: deter, delay, detect, respond and recover (see view 23). The innate immune system like skin and tonsils correspond to the cyber security policies and the firewall. The adaptive immune system covers the other three steps which in the cyber world are executed through system monitoring and the security operations centre (SOC). Large organisations may have more than one SOC, smaller companies tend to outsource this function to third parties.

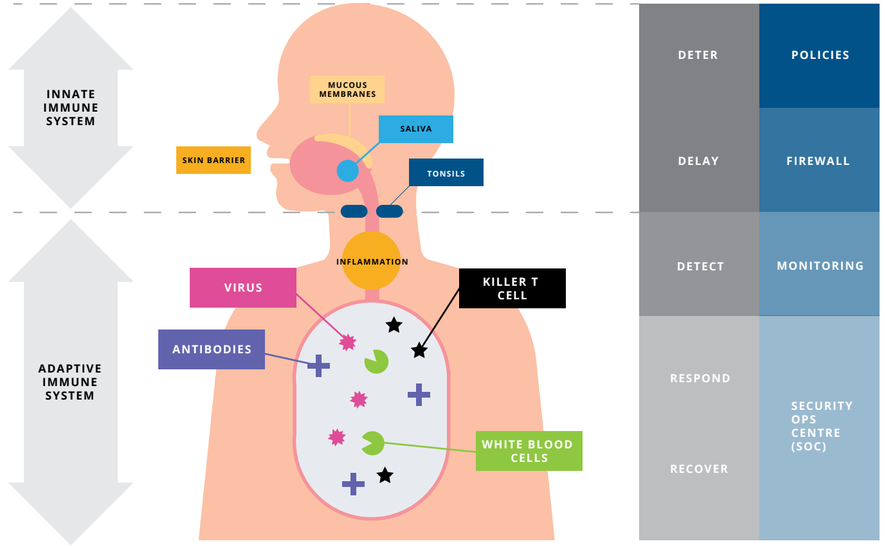

# 12. History of cyber attacks

What is the best way to illustrate the growing threat of cyber attacks? You could count the number of attacks happening worldwide. There are several maps available on the internet showing cyber-attacks happening in real time from companies such as Threatcloud, FireEye and Kaspersky. These make troubling viewing. On a typical day they might log around 16 million attacks [3] happening around the globe. However, these attacks are not all successful, nor is there any sense of the scale or severity of each one.

A different way to monitor the cyber threat is to measure it by counting the number of records breached. This method gives a better sense of scale and only includes successful attacks by definition. The chart to the right shows the cumulative number of records breached in cyber-attacks over the last 15 years based on data from the Identity Theft Resource Center (idtheftcenter.org).

When cyber criminals hacked the WIFI network of a Minnesota store and stole the details of 94 million credit cards from TJ Maxx in 2007, people were shocked by the scale of the breach. But as the chart shows, that breach has paled into insignificance when compared to the scale of the recent events such as the 1.3bn email records [4] breached when the spam operator River City Media was hacked in 2017.

## Not just the USA

Also note that while early breaches tended to be US focussed, they have now spread geographically to become a truly global problem. For example, in 2018:

- Nametests, an online quiz app based in Germany suffered a 120m record breach.
- Some 200m Chinese resumes filled with phone numbers and work place details surfaced on the dark web from an unknown source.
- The personal information of more than a billion Indian citizens were leaked from the government's new national identity database known as Aadhaar. This is the world largest biometric database which can be reportedly bought online on the dark web for as little as $10.

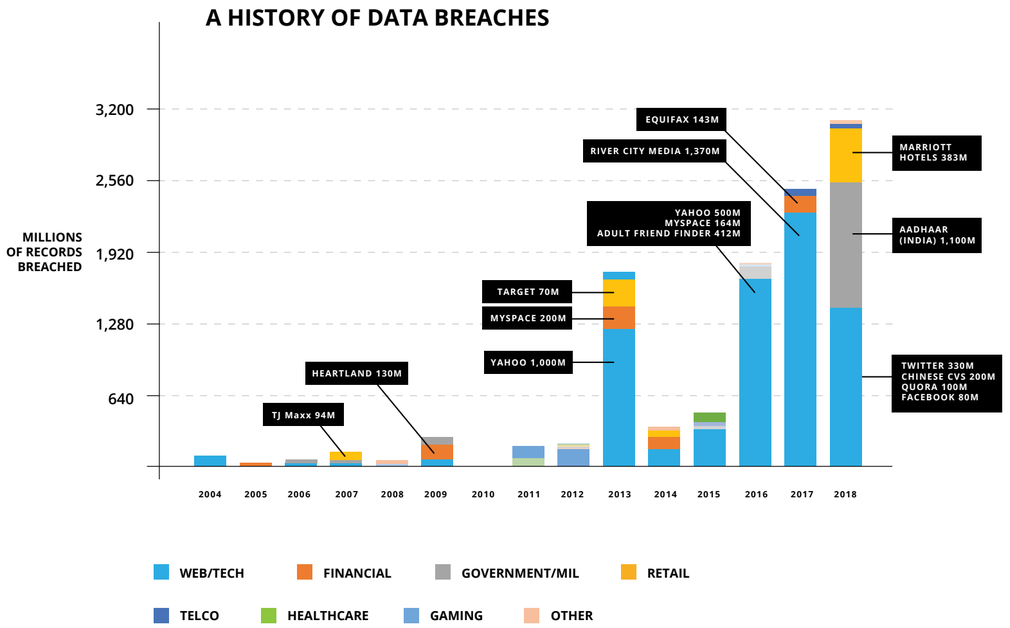


Data source: idtheftcenter.org,datacreaches.net

# 13. The evolutionary arms race

Cyberspace is a Darwinian space. Just as with lions and wildebeests on the Serengeti plains, or taxpayers and taxmen, hackers and technology companies are in an evolutionary arms race. Each side develops in response to changes in the other. The predator develops better attack methods which leads to improved defences from the prey; the two locked in a dynamic equilibrium that may swing to favour one side temporarily before reverting to the mean.

## FAANGs ain't what they used to be

The internet is fundamental to the business model of most tech companies. It represents a commons that must be safeguarded from evil doers. Consider these five companies: Facebook, Amazon, Apple, Netflix and Google. Collectively known as FAANG stocks, they make up 13% of total capitalisation of the US stock market [5]. All rely on the internet being regarded as a trusted and safe domain within which to conduct transactions. So, whilst governments may have concerns over cyber warfare, and regulators drive compliance through the threat of fines, it is actually the private sector - the tech companies like FAANGs and smaller cyber vendors - that are the main driving force behind the evolution of cyber defences.

Technological innovation has resulted in significant improvements in cyber defence mechanisms in recent times. Common attack vectors are being stymied by defensive counter plays. Where consumers once used the same 'easy to guess' password for all logins, password managers are now built into most internet browsers. These automatically suggest strong randomised passwords to users, storing them securely so there is no need to commit them to memory. Likewise, 'two factor authentication' based either on a mobile phone number or on biometrics like fingerprint and face recognition have added an extra layer of security to the process of identity verification.

## Secret SaaS?

Ransomware attacks, which lock up access to data until a ransom is paid, have made the case for cloud computing more compelling. If your accounting or client account management systems are provided by a cloud based third party like Salesforce or Xero, they are no longer sitting on a vulnerable local server in the office. Software as a service (SaaS) moves the data further away from the attacker, reducing the leverage of ransomware.

The most interesting recent development in the evolutionary arms race is the 'disposable' credit card offered by fintech companies such as Final Inc. If you are concerned about making a particular online purchase, why not use a disposable, 'one-off' credit card. A new credit card number can be generated in seconds which is specific to that transaction and will be deleted after use. This is a powerful solution to the problem of credit card theft.

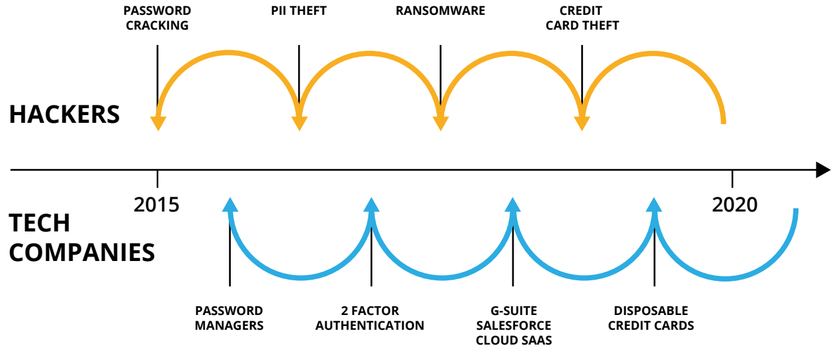

# 14. Cyber vulnerability pyramid

Vulnerability to cyber risk can be thought of as a pyramid with three parts. At the bottom layer is the technology component: the devices, the network, the firewall and all the other parts of the IT infrastructure. The mistake that many companies make is the belief that cyber risk belongs 'down there'; that it is solely an IT issue and something for that department to sort out on its own. But there are two other layers of equal, if not greater, importance.

The next layer up is the process layer. There is always a debate in organisations about finding the best balance between resilience and efficiency. Security measures have an overhead cost, both in time and money, and are often viewed as an obstruction to business rather than an enabler. Process vulnerabilities normally stem from a misalignment between security procedures and business objectives. Processes may have been defined only to be ignored or only partially implemented. This is where executive buy-in is key. Department heads should not only lead by example but also providing constructive feedback to get the security vs efficiency balance right.

## To err is human

At the top of the pyramid is the people layer. This can be seen as the area of greatest vulnerability as 90% of cyber security incidents are reputedly caused by human error [6]. This can be either accidental as in a 'fat finger' error (see view 21) or malicious (see view 19). The vulnerability in the latter case at the people layer is known as social engineering. This is the psychological manipulation of people to get them to divulge confidential information such as log-in credentials. It exploits human curiosity by getting people to click on an interesting looking link which then secretly installs malware in the system. These are known as phishing attacks and come in these varieties:

**Phishing** - an official looking email which looks like it comes from, say, your bank. It's really a scam to get your account passwords or credit card details.

**Spear phishing** - this is a targeted version of phishing directed at a single individual rather than a mass email list. Carefully crafted after research on social media sites, they target high profile people in key roles.

**Vishing** - voice based phishing using the phone rather than email.

**Smishing** - as above but using SMS messaging on a mobile phone rather than email.

The best defence against these attacks is user awareness training; teaching people not to click on dodgy links in emails. We leave the last word on the social engineering threat to the cyber security expert Bruce Schneier and his famous aphorism:

"Only amateurs attack machines, professionals attack people".

*"Only amateurs attack machines, professionals attack people..."*

**SCHNEIER ON SECURITY**

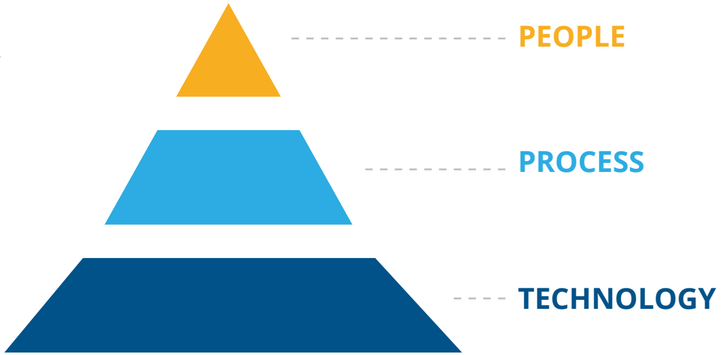

# 15. Hacking return on investment

## Targeted or random?

The dark web has developed so fast that you can now conceive of it as an economy in its own right. If the dark web were a country, then Ireland would be a good proxy with a population of 5m and a GDP of $350bn [7]. The dark web has its own central bank in the form of Bitcoin which has daily transaction volumes of $11bn. Of this, maybe 10% represents illegal activity rather than currency speculation.

As in a developed economy, there is a high degree of specialisation in the dark web with the formation of complex supply chains. So, for example, an expert in crafting phishing emails with a high click rate, sells them on to a ransomware expert who has previously bought a malware kit from another third party. Any data that is successfully exfiltrated can then be sold on again to an outfit that specialises in handling stolen credit card details. This is a mirror of the real economy where the transformation of raw materials into finished goods for consumers is a journey that passes through many intermediaries.

## Honour amongst thieves

Maybe the most surprising thing about the dark web is the discovery that there is in fact honour amongst thieves. A study of the dark web by the London School of Economics in 2017 found that online satisfaction ratings for shops on the dark web were extremely high, with a negativity rate of less than 3%. As with Amazon and eBay, reputation for online shops is everything even if what they are selling is illegal. Even more surprising are the online help desks, if the malware kit you purchased does not work you can call the vendor for support, just as you can with Microsoft.

The sophistication of the dark web economy means that hacking is best comprehended as a return on investment equation. The costs of multipurpose hacking tools - kits that enable you to steal data from a small company with rudimentary cyber defences - have fallen dramatically in recent years. What used to cost thousands of dollars can now be bought on the dark web for little more than $100. Conversely, developments in crypto currencies and an increasingly sophisticated distribution chain have increased demand for stolen data pushing sale prices up. This means the return on investment for a cybercriminal has been improving from both ends.

It is possible we are close to a turning point in this cycle. The huge supply volume of stolen data now offered for sale has meant prices have begun to fall. In addition, the increased capability of cyber defences through biometrics and two factor authentication have required more sophisticated, and therefore more expensive, hacking tool kits.

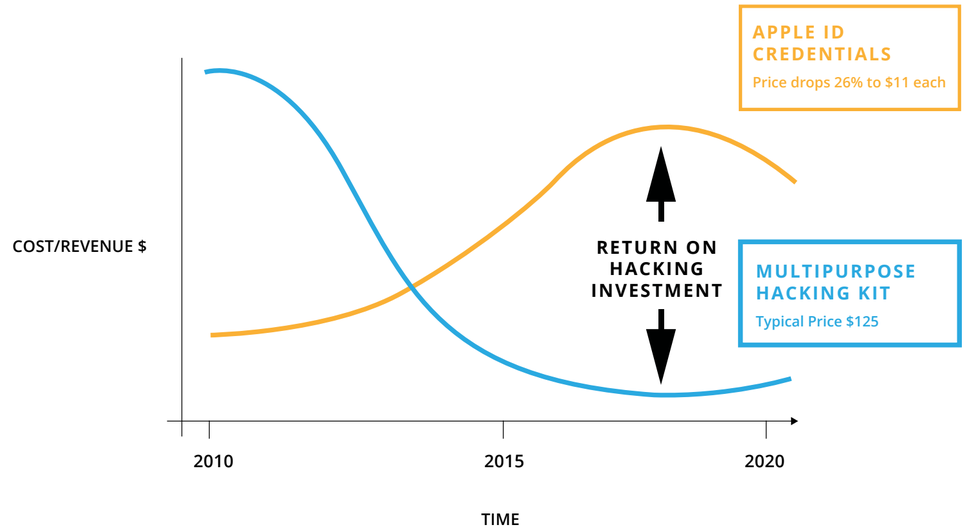


Data source: Top10VPN
Dark Web Market Price Index

# 16. PII prices on the dark web

Let's examine the prices of personally identifiable information (PII) on the dark web. The chart to the right displays data from a Comparitech survey taken in 2018. Look first at the cost of a passport scan if you wanted to buy one from the dark web. They would typically cost $15 each, although as you can see there is a big variation in prices depending on the nationality of the passport involved. Chinese passports are the cheapest. This may reflect supply since China's population is almost five times bigger than the USA which would imply a greater number of passports. It might also reflect demand. The GDP per capita in the USA is six times that of China and wealthier passport holders suggests a greater potential for financial exploitation.

## Identity is cheap

Prices increase fourfold if a proof of address (PoA) is offered in conjunction with a passport scan as shown by the charts to the right. These two pieces of information constitute sufficient proof of identity in most countries to pass identification checks and, say, open a bank account or begin a mobile phone contract. Notice the prices for USA PII of this type seem extremely low. This may be a timing effect as the survey was taken soon after some major cyber breaches dumped 500m PII records into the dark web marketplace. Notice also that French PII prices are comparatively high which could be an indication of the prophylactic properties of a non-English language. Only 3% of Internet use is in French - a smaller and more difficult pool to fish from.

## The real-world premium

Lastly look at the difference between the cost of a passport scan and the cost of an actual passport document. A forged passport costs $1,478 on the dark web. A 'real' passport made from an official blank and acquired through a corrupt bureaucrat costs almost 10 times more. Next time an apparatchik asks you for your real passport document and not a scan, temper your frustration. You can see from the prices shown here that the extra diligence is clearly worthwhile.

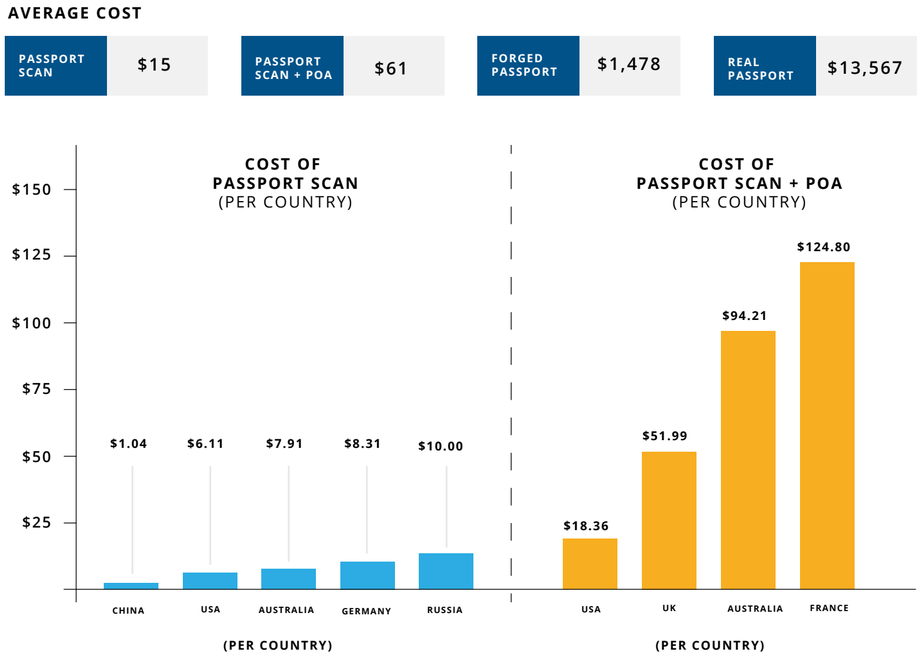


Data source: Comparitech 2018

# 17. Malware infection rates

Most cyber-attacks involve the installation of malware in the victim's system at some stage in the process. Malware is malicious software designed to cause damage through executable code and can take the form of computer viruses, trojans, ransomware, spyware and keyloggers amongst many others. The diagram to the right shows malware infection rates in different countries around the world. The size of the circle is proportional to the number of devices in that country while the colour shows the malware infection rate. For example, countries where less than 8% of devices are infected are shown in green and those over 30% shown in purple. Basically, if it's big and red hued, it's bad news.

## Mobiles outnumber PCs

The diagram has two sections, the top one shows the infection rates for PCs and the bottom one for mobile devices (but not laptops which are counted as PCs). Note that the number of PCs and mobile devices in the USA is broadly similar at around 200m units which is why the circles are the same size in both the top and bottom sections of the diagram. The same is also true for Western Europe and Japan. Now look at China, where the situation is very different. There are eight times more mobile devices in China than PCs and in India and Nigeria this ratio is even greater.

The conclusion is that while the cyber threat is normally seen as a PC based issue for developed economies in the West, in reality it is a much more serious issue for mobile devices in emerging markets. China is the manufacturing centre of the world, as India is for IT and outsourced services. There are few companies in the world that do not use goods or services that originate from one of these two countries. These high malware infection rates should be borne in mind whenever communicating or transacting online with these countries.

## IoT on the Horizon

A further concern is the evolving Internet of Things (IoT) where internet connectivity is being added to physical devices and everyday objects to create smart homes and autonomous vehicles. In a few years' time when this diagram is redrawn, the new IoT section will dwarf the other two in scale. Security is notoriously poor on most IoT devices.

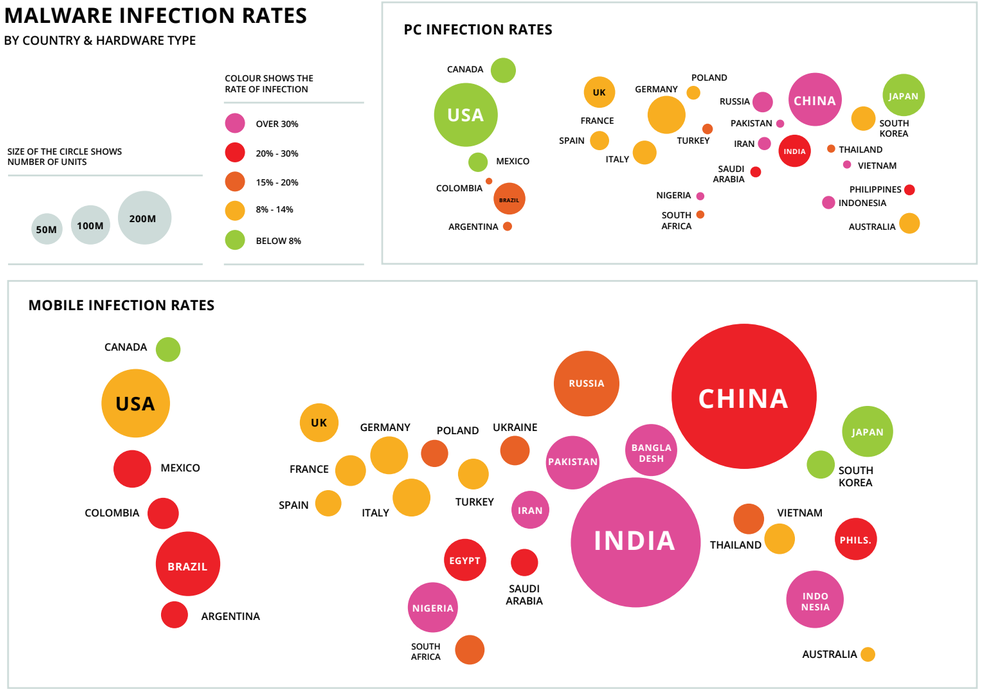


Hardware units data source: Wikipedia. Nationmaster

Infection data source: Kaspersky Labs 2019

# 18. Threat actors: Nation states

As insurers our focus is normally on the commercial sector, but there are some important points to make about cyber conflict between nation states. First, war in the information arena is not new and existed even before nation states did. Stealing information from the enemy with espionage strategies was eloquently outlined in Sun Tzu's Art of War in the 5th Century BC. What is new is how the objectives of cyber conflict between nation states have recently changed. The focus of espionage was simply gathering information. Other objectives are now emerging such as the disruption of infrastructure (e.g. the Stuxnet attack in 2010) or influence over society as a whole (e.g. suspected Russian interference in the US Elections in 2016) which represent an escalation of cyber conflict to new strategic theatres.

## Offence vs defence

A second point to highlight is the distinction between cyber warfare and traditional armed conflict sometimes referred to as 'kinetic' warfare. Cyber conflict is seen as being a constant low-level background occurrence until it reaches such an egregious level that it crosses a tipping point into kinetic warfare. This critical point is known in military circles as LOAC (the level above which armed conflict begins). At this point, the relative advantage between offence vs defence changes dramatically.

A rough rule of thumb in military circles is that in kinetic warfare an attacker needs a three to one advantage in manpower and firepower in order to successfully defeat a defender. Defenders typically have an advantage because it is normally easier to protect and hold than it is to move forwards, to destroy and to take. However, in cyber warfare the opposite is true. Attackers have an enormous advantage, maybe by a factor of as much as 10 to one. Large institutions must defend against many thousands of attacks every day. Only one needs to get through for an attacker to succeed. Generally speaking, offensive cyber-attacks are low cost with a high payoff, where defensive operations are expensive, overstretched and often ineffective.

Cyber warfare is asymmetric. Both the attacker and the defender are in a race to find vulnerabilities. But the number of vulnerabilities grows exponentially with the size and complexity of the system. The defender has little chance of finding every single vulnerability and patching it before the attacker finds one to exploit.

## Response options

How should a nation state respond to a cyber incident? The diagram to the right shows the range of options from a passive defensive strategy which can be expensive and time consuming to a more aggressive retaliatory stance inflicting proportionate damage backed up by evidence. The USA's stated strategy in cyber warfare is one of persistent engagement delivering proportionate counterattacks and 'forward defence' in neutral zones. As in the cold war, the end point will probably be a tacit bargain; an informal understanding that evolves over time as to what is or is not acceptable before LOAC.



### THREAT ACTORS: NATION STATES

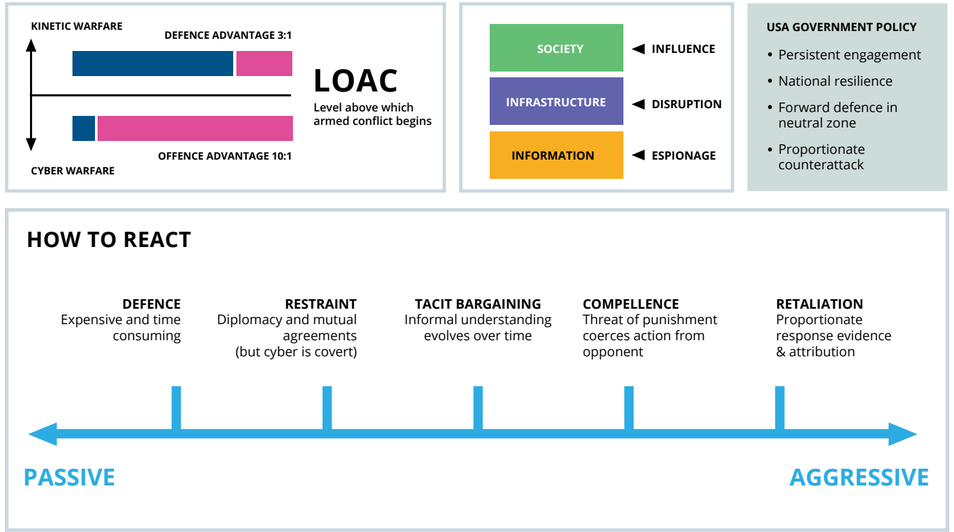

# 19. Threat actors: Rogue employees

A common statistic states that some 90% of all cyber security breaches are caused by human error of some type. At the most basic level, this could just be an employee mistakenly sending confidential information to the wrong recipients - known as a 'fat finger' error from hitting the wrong key on the keyboard. Then there are phishing attacks which exploit human curiosity; an employee clicks on an innocent looking email that piques their interest and so installs malware unintentionally. Drive-by cyber-attacks trick conscientious executives who put in that extra hour of work in the airport lounge over insecure local WIFI connections. (Note to self: keep off the laptop and head for the snacks!)

## Spotting a rogue employee

All the above are examples of unintentional errors, but we should also recognise that there are sometimes rogue employees in organisations who deliberately cause a cyber breach. How would you spot such a rogue employee? Psychologists agree that they share three personality traits. First, they have narcissistic tendencies with an over inflated sense of self and a need to feel superior if their ego is threatened. Second, they have a Machiavellian mind set; a cynical world view where unprincipled behaviour is acceptable because the end justifies the means. Last, they display psychopathic tendencies, impulsively seeking thrills while disregarding other's feelings.

This type of analysis puts all the blame on the individual but note that corporate culture has an important role to play too. There are five factors in the corporate environment that can trigger destructive behaviour in employees as identified by Furnham and Taylor in their book "Bad Apples". The first trigger is an uncaring company atmosphere where bullying is rife and employees feel downtrodden. The second is unmet expectations where promises made during the interview process are not upheld. Third is corporate hypocrisy, a huge rift between the CEO's vision statement on the website and the reality of daily work. In this environment, words clearly don't match deeds. Fourth is a lack of trust, managers are suspicious of workers and vice versa. Last is a high level of inequality where employees receive vastly different treatment; loyalty and diligence is unrewarded while sycophants are promoted.

Any organisation where these five factors are characteristic of the corporate culture is creating a toxic brew that is bound to produce malicious behaviour from rogue employees. Thankfully, all five factors are completely within the compass of corporate control. The best route to reducing rogue employee risk is therefore clear.



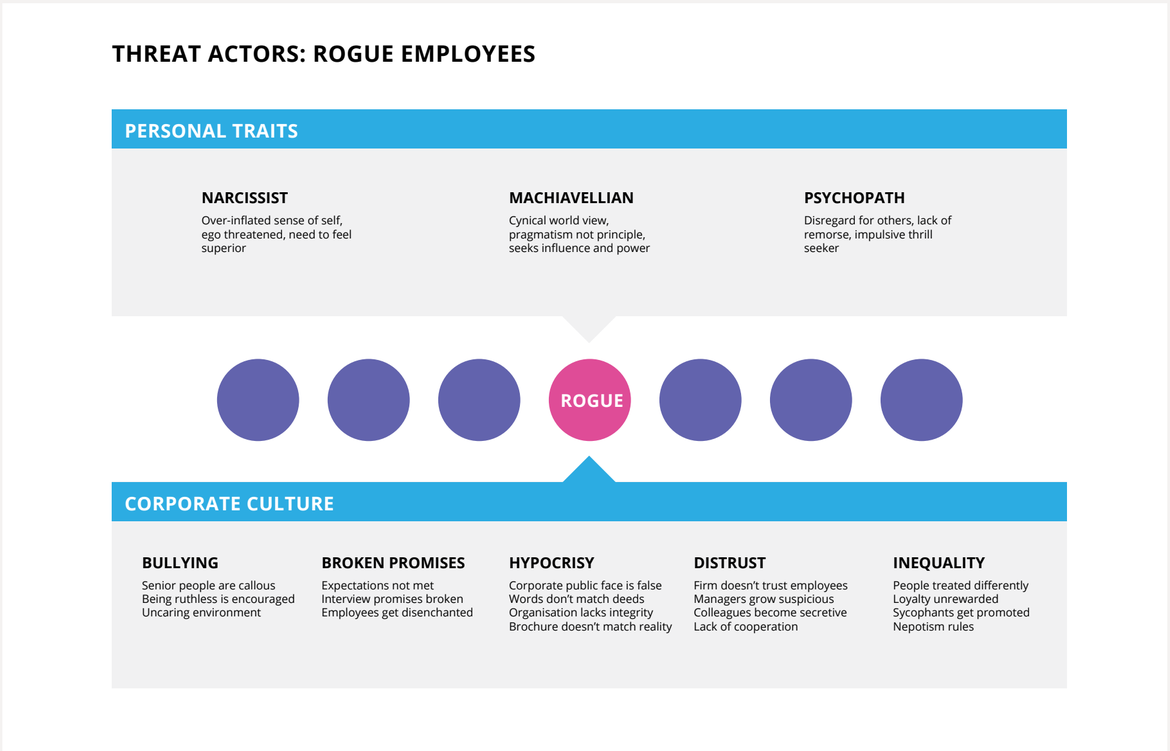

# 20. Threat actors: Botnets

Sometimes you can be involved in a cyber-attack where the intended victim is someone else. Crypto-jacking is a good example of this, where your system has been hijacked and control has been ceded to a third party. A 'bot master' will link together a large number of these hijacked machines and assemble them into a botnet which can then be rented out to other cyber criminals on the dark web.

A botnet can be employed for a variety of purposes. They are often used in Distributed Denial of Service attacks to take down a company's website by overloading it with spurious requests. In the case of crypto-jacking, the processing power of the botnet is used to mine for bitcoin on the internet. Botnets are also commonly employed to spew out phishing emails. A fourth and particularly lucrative exploit is to use a botnet for advertising fraud. The diagram to the right illustrates how this works.

## How ad fraud works

Google pay per click advertising service has four key steps as shown:

1. User clicks on web ad

2. Ad click is registered by Google

3. Google pays money to web owner

4. Advertiser pays Google

However, a cybercriminal can set up a new website and then use a botnet to automatically click on the Google web ads. Since it can be hard to distinguish if a real human being or a botnet is clicking on the ad link, the website owner can fraudulently extract a river of cash from Google and ultimately from advertisers.

A study by Juniper Research in May 2019 estimated that advertisers are losing $42bn every year to this type of ad fraud and that it could grow to $100bn a year by 2023. The same study concluded that as much as 50% of internet advertising could never actually be reaching real humans. Advertisers are understandably concerned about this and suspect that Google is not addressing the problem as vigorously as it should because to do so would dramatically reduce Google's advertising revenues. Google strenuously denies this.

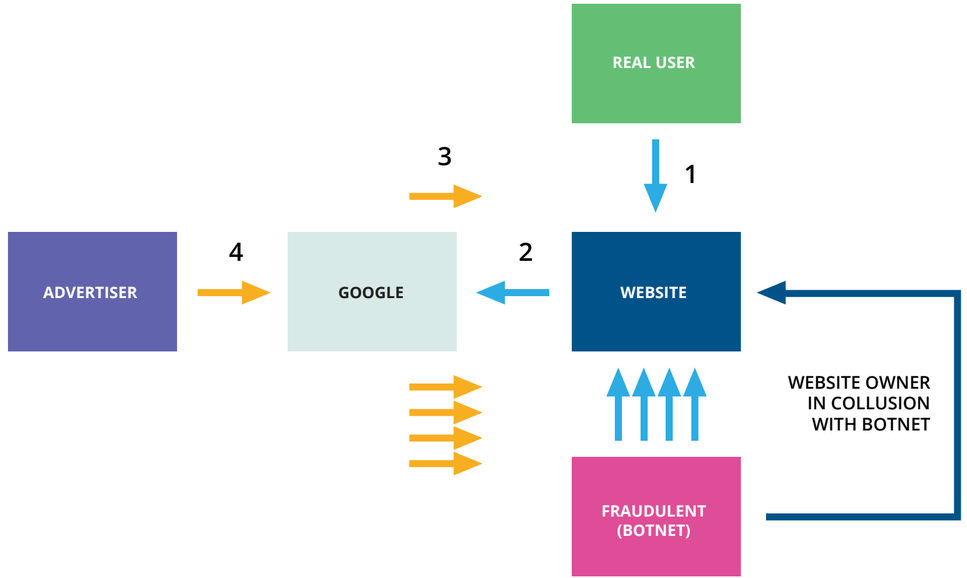

# 21. Why me?

A question asked by many small and medium-sized enterprises (SME) is “Why me?”. What do SMEs have that would be of interest to a cybercriminal and so make them a target? The answer is that you may not have been picked for any particular reason and purely on a random basis. Also, the motivation behind any attack may be moral rather than financial. The two quadrants to the right illustrate this issue.

## Targeted or random?

The left-hand quadrant shows the different types of adversaries. The attack may be targeted or completely random (horizontal axis) and casual or persistent (vertical axis). The most dangerous types of attack are Advanced Persistent Threats (APTs) in quadrant two. If an attacker is determined enough and prepared to invest substantial amounts of time and money, they are likely to succeed in the end. This type of adversary is typically state sponsored, looking to steal sensitive defence secrets or disrupt critical national infrastructure. Countries such as Iran, Russia, China and North Korea are reputedly active in this area. A more typical adversary would be a cyber criminal in quadrant one, randomly selecting targets based on unpatched system vulnerabilities or even an opportunistic amateur in quadrant four casually scanning the internet for victims.

## Financial or moral?

The right-hand quadrant shows the different motivations behind an attack. Cyber criminals are normally out for financial gain and as much money as they can make. But in a crypto-jacking attack, you are not the ultimate target and the harm may be mild. Rather, your spare system capacity has been hijacked to use for another purpose. Also, the motivation for the attack may not be financial at all but driven by moral concerns instead. When Sony Pictures was hacked in 2014, some 170,000 confidential emails were posted on Wikileaks containing many embarrassing disclosures. The aim of the attack was to shame the senior executives in the company and expose their hypocrisy. Ethically motivated hackers - known as hacktivists – are less interested in financial gain and more driven by their political orientation or personal beliefs.

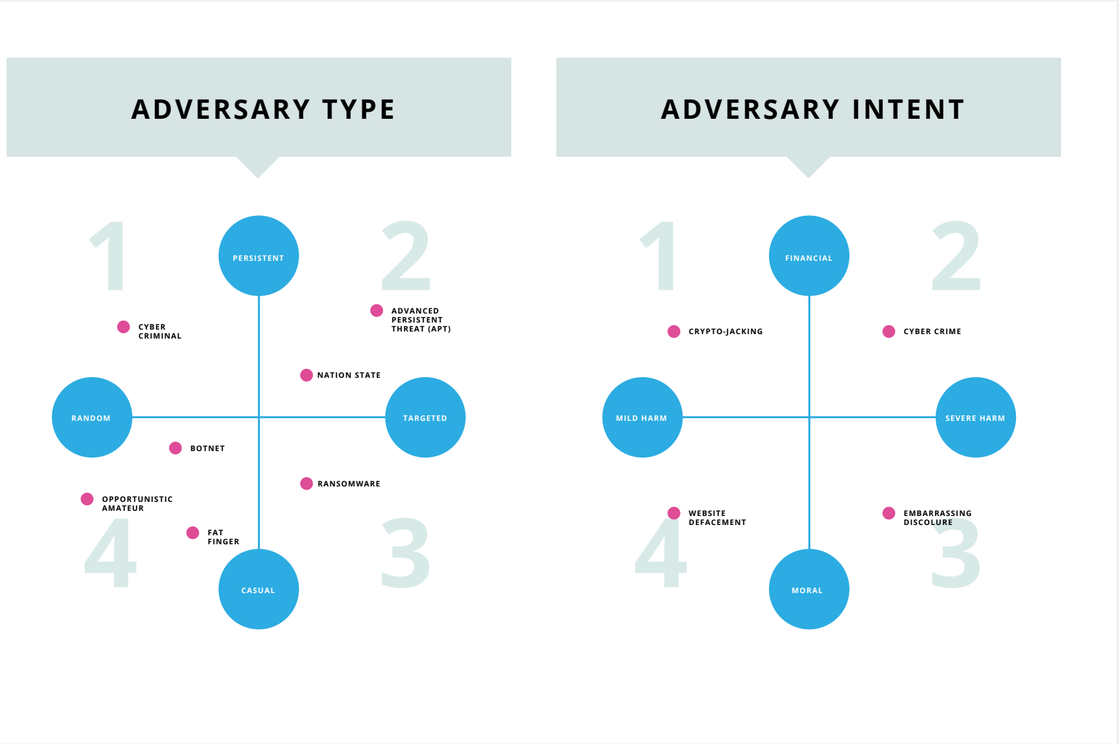

# 22. Breach incident chain

Some companies mistakenly assume that a cyber-attack is just an IT problem. In fact, almost all departments in an organisation need to be involved in the response to such an incident. The diagram to the right gives a very simplified picture of the main stages in how a cyber incident plays out.

First, bear in mind that the response may not be triggered until several months after the systems have become compromised. Times to detection can easily be this long and the notification of suspicious activity will need to be escalated several times up the bureaucratic hierarchy before a formal breach response is triggered.

## Assemble the response team

The first steps are to inform the board and engage an external breach coach. This individual will then coordinate the response with the key stakeholders which typically are the insurer, legal counsel, IT forensics and crisis communications. These key response team members will then in turn engage with other parties to contain the breach and organise the steps towards recovery.

**Insurer:** The insurance provider and the broker need to be notified at a very early stage, not least because many of the costs involved in breach response will be covered as part of the policy. Ideally the range of breach response services will already have been agreed in advance.

**Legal Counsel:** A key conduit for communications with the government regulator, law enforcement and handling potential liabilities with third parties. Sound legal advice is essential given the current complexity of data protection and privacy laws. It's a good idea to use external legal experts for this.

**IT Forensics:** Specialist external IT consultants will be needed to help figure out the source of the breach and the extent of the damage. Once the path to recovery becomes clear, they can then coordinate with in-house IT support to rebuild systems and recover corporate data.

**Corporate communications:** Cyber incidents can cause major reputational damage. Skilful handling of communications with customers, the media and internally with staff will reduce this. Call centres and hotlines will need to be set up to cope with the flood of enquiries from concerned customers.

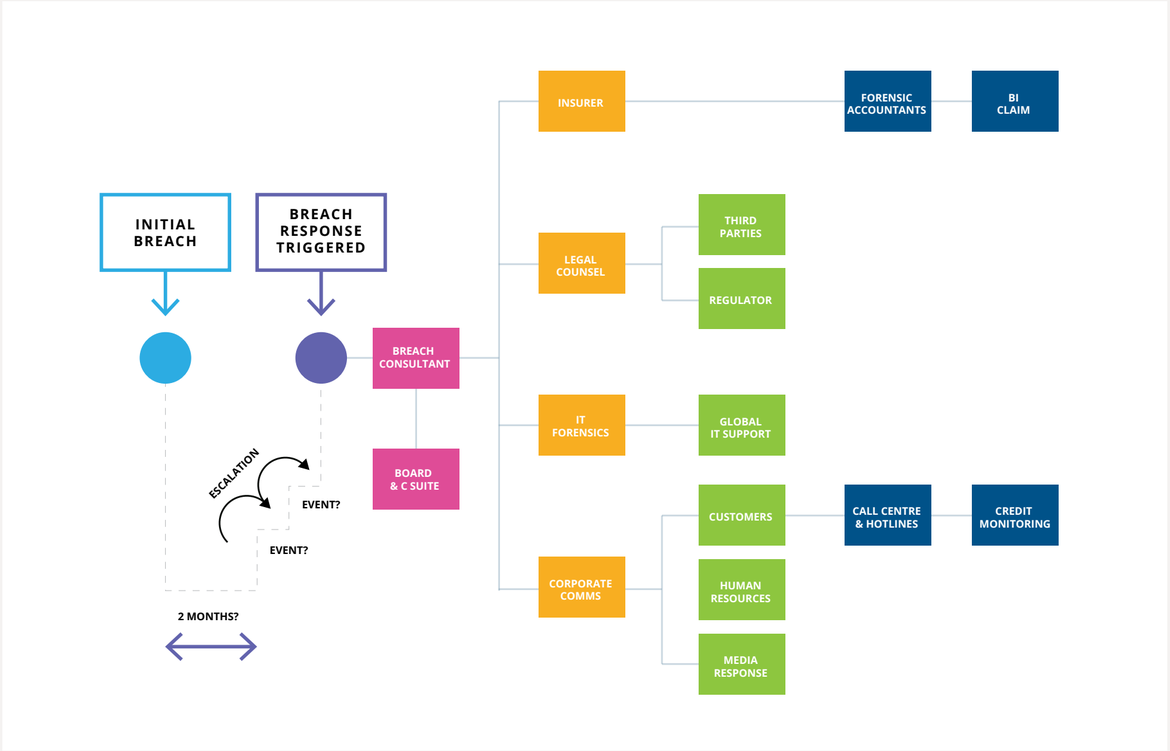

# 23. The incident response plan

Let's look in a little more detail at a typical incident response plan, as illustrated in the diagram to the right. It is divided into several different stages such as detect, assess, isolate, recover and post mortem. At each of these stages the following questions need to be addressed:

- who needs to be involved as part of the decision-making team?
- how will the necessary steps be executed?
- where will these activities take place?
- when and in what sequence will they happen?

## Why do I need one?

An incident response plan brings clarity during times of confusion by providing pre-prepared guidance and instruction. It clearly defines the roles of the people on the response team and spells out in advance a communications plan describing who will inform whom of what, and in what order.

When monitoring systems, a substantial number of false positives should be expected. If there are no false positives, in other words no suspect events that turn out to be harmless, there is something wrong with your level of monitoring. Your employees may not be informing you of suspicious events.

The success of the kaizen system which propelled Japanese car manufacturers to world dominance was based around giving production line workers the authority to hit the stop button if they saw something wrong. Pushing this authority low down the corporate hierarchy was counterintuitive but resulted in significant improvements in quality and efficiency. In a similar vein, in an incident response plan, the decision as to who has the authority to hit the stop button, so triggering the active incident response process, is a key one.

You can see that each stage in the incident response process has some key performance indicators (KPIs) that should be recorded and reviewed on a regular basis. As mentioned previously, false positive are very instructive. Other useful ways of quantifying performance are measures of the mean time to detection (MTTD) and mean time to resolve (MTTR) which we return to later (view 27).

One area that is often overlooked is application of lessons learned. Remedial actions are often listed as desirable in the post mortem phase, but too often these are never applied or followed up on. Time and budget need to be allocated to fix the root causes of the problems.

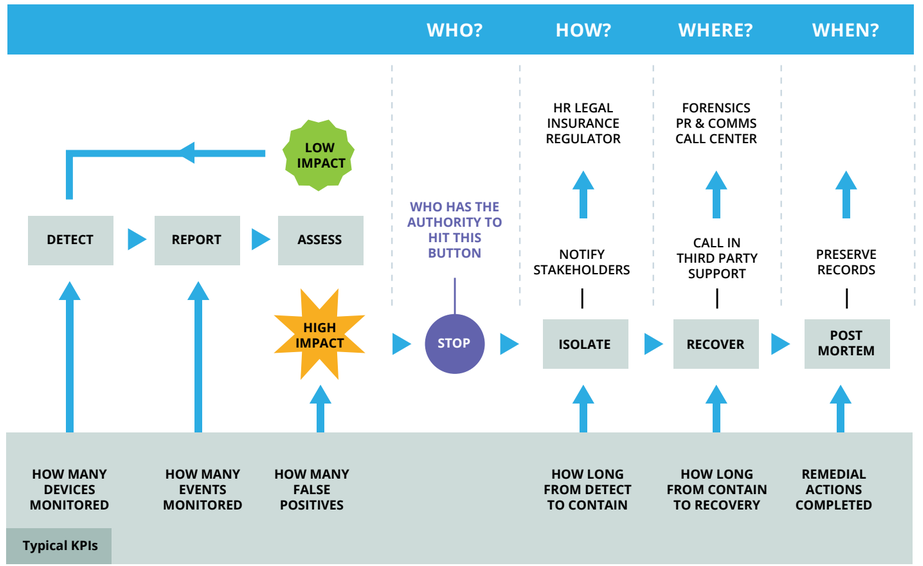

# 24. Costs of a cyber incident

What does a typical cyber breach cost? It depends on a large number of variables and much granular happenstance. Taking an average across 477 cyber incidents, the Ponemon Institute suggested a cost of $4m for a typical breach in 2018. But averages such as this, across many different industry sectors, incident types and corporate sizes can be misleading. A better approach is to build a company specific model from the bottom up.

A crude but workable model can be built using only three company specific inputs: headcount, revenues and customer base. These three factors vary widely for different industries. A subcontractor making clothes for a fashion retailer would have a large manual workforce but only one customer. Conversely, an online retailer would have few employees but a very large customer base.

Some parts of the cost of a cyber breach can be viewed as more or less fixed. That is not to say that a tiny company and a huge multinational will be paying the same bill regardless, but some aspects of cost across, say, the SME segment will be very similar. A forensic specialist hired to investigate a breach will have the same day rate whether the breach is large or small. However, other elements are very volume dependent and directly proportional to three key factors:

**Revenues:** lost sales due to business interruption form a substantial part of a cyber incident's costs. These can be estimated by multiplying average daily revenue by number of days expected outage. Some cyber incidents can take business critical systems out for months.

**Headcount:** restitution costs are proportional on the size of the IT estate. It is not uncommon for a company to replace all its software and PCs post breach to ensure they are restarting with a clean system. The number of PCs in a company is proportional to the number of employees, adjusted for the blue to white collar ratio.

**Customers:** In the USA and Europe, companies that suffer data breaches are likely to suffer fines from the regulator linked to the number of customer records breached. But even excluding these regulatory fines, there are other costs that scale up relative to the number of customers a company has. Customers need to be formally notified that their data has been exposed and the dark web monitored to see where this data is surfacing. Often an external call centre needs to be engaged to handle all the concerned calls from clients. These costs are all proportional to the size of the customer base.

## Policies can cover costs

The last element to factor into the cost model is the insurance policy. Many policies cover some part of the breach costs described. See view 30 for a more detailed discussion of what a typical cyber policy might cover.

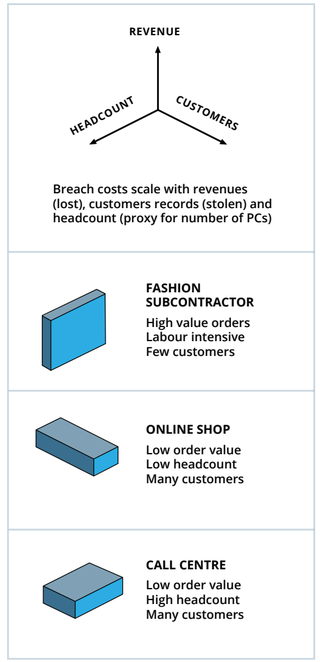


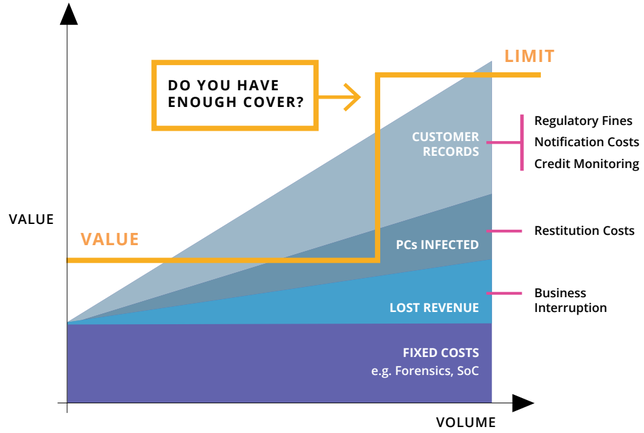

# 25. Cyber insurance: Product or peril

The insurance industry, in the 330 years since its beginnings at Edward Lloyd's coffee shop, has grown organically, evolving from the bottom upwards into the industry we see today. This organic development, in the absence of a top down blueprint, has been very effective in responding to client's demands but has led to a very confusing taxonomy of the classes of cover on offer. Some classes of insurance are defined on industry lines such as marine, energy or aviation. Others are defined on a product basis such as property, kidnap and ransom (K&R), directors and officers (D&O) liability and the like. These two different categorisation methods lead to plenty of definitional overlap. A librarian, believing that classification schemes need overarching coherence, would throw up their hands in horror at this jumble of confusion and retreat, whimpering, to their Dewey Decimal system. Who cares! From an insurance standpoint, it may be ugly but it works...

## Non-affirmative Cyber

Until now. The two defining axes of product and industry are illustrated to the right. The kidnap and ransom of a ship's crew would be at the intersection point of K&R and Marine as shown. But there is a third concept in insurance: the peril. A peril is an event or circumstance that causes a loss such as a fire, floods or tornado. Perils cause damage in both product and industry categories; a tornado can damage buildings (property) and boats (marine). So, perils are represented on a third axis, underlying the horizontal plane that contains the product and industry classes.

That leads us to the key issue: is cyber a product or a peril? Until now, cyber policies have typically been seen as a stand-alone insurance product, represented by the green slice. However, an 'all perils' property policy theoretically covers damage from a cyber-attack even though this is not explicitly spelled out. This is known as non-affirmative or 'silent' cyber and represented by the orange layer in the diagram. A white paper on non-affirmative cyber called "Are we heading for PC&C" published in 2018 by Capsicum Re estimated that non-affirmative cyber was nine times bigger than affirmative cyber. Since this cyber coverage is silent, that's a lot of missing premium for underwriters.

It's a big headache for regulators too because it means that risk may be being mispriced in the marketplace. How will the problem be resolved? The evolution of the terrorism insurance market offers some clues. After 9/11, terrorist perils became explicitly excluded from property policies and silent terrorist cover became explicit. However, as time passed, and data sets improved, terrorist risk was reabsorbed into property policies again. In a similar way, cyber may move from product to peril and then back again in the fullness of time.

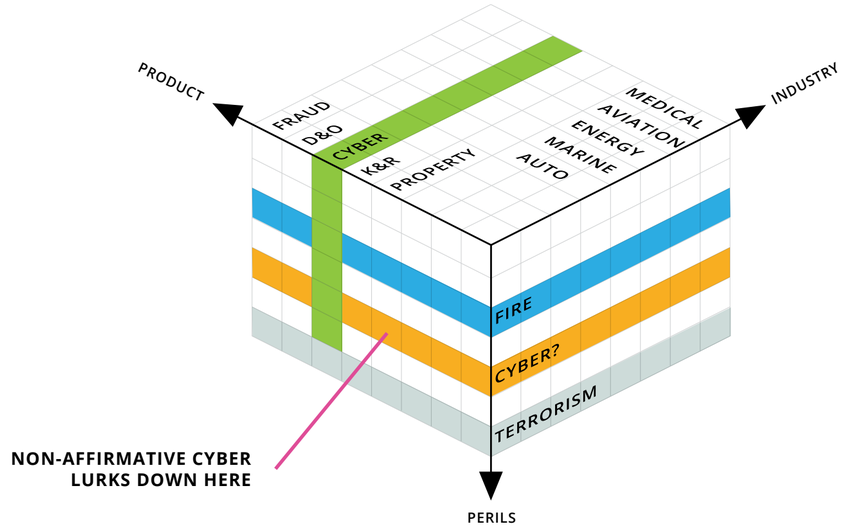

# 26. Where are the peacock's feathers?

Beautiful as they are, a peacock's feathers actually serve a very practical purpose. They have evolved over millennia as a visible sign of mating fitness. A peahen, simply by looking at the splendour of the peacock's tail, can gauge the health and desirability of that individual as a potential mate. A simple visual clue that provides a reliable measure of underlying soundness.

## Smoking or non-smoking?

In some branches of insurance there are similar simple measures that enable the segregation of good risks from bad. Think of health insurance. A question such as "Are you are smoker or a non-smoker?" is an easy place to start. Of course, a more exhaustive questionnaire and some medical tests with blood work will provide a much more detailed assessment. But only a few simple questions will separate the sheep from the goats. This is the insurance equivalent of the peacock's feathers.

So, the question for cyber is "Where are the peacock's feathers?". The sad answer is that they are still evolving. There are some published technical standards such as the NIST Cyber Security Framework, the ISO 27000 series standards and the UK Government's Cyber Essentials certification which set out sensible guidelines. But it is still a fairly open question whether compliance with these guidelines substantially reduces cyber risk. After all, the very best companies still get hacked and even the National Security Agency itself has suffered breaches (e.g. Snowdon).

## The fortuitous loss

The central issue here is the concept of the 'fortuitous' loss; a loss that is beyond the control of the insured. Government regulators in the USA and Europe impose fines for data breaches. But if the insured has followed best industry practice in terms of cyber security, is it really fair to be penalised in this way? What is needed is a widely accepted set of indicators – some peacock's feathers - that can be used as a fortuity test to absolve insured corporates of blame. This is an area where close cooperation between the cyber security and insurance industries could lead to very fruitful results for both parties.

# 27. The units of cyber risk

We know how to measure Mount Fuji. In three-dimensional space, we have the measurements of height, depth and breadth and we can multiply these three together to calculate the mountain's volume. We can take rock samples to calculate density which, when combined with volume, will allow us to estimate mass. We can measure the temperature of lava in degrees centigrade and we can use a seismograph to measure geological tremors on a Richter scale. The units of measurement in the physical world are well defined and understood.

What about the cyber realm? As we discussed before (view 26) there are no peacock's feathers. The industry as a whole is yet to develop standardised, easy, visible markers as to a company's cyber health. But we can at least advance some ideas as to the dimensionality of cyber risk. This indicates a way to describe the issue even if the gradation markings on the ruler have not yet been formalised.

In our proposed scheme, there are three dimensions: hierarchical, lateral and temporal. Just as three axes define physical space, these three dimensions can describe cyber risk in an analogous way. So, the hierarchical dimension has a vertical connotation, the lateral a horizontal emphasis with the last axis being time related. We put these forward as a coherent way to describe cyber risk; the AXIS cyber axes, if you will.

**Hierarchical** - this axis focusses on elements of top down control. It is a measure of cyber governance; the degree to which the rules of cyber hygiene are embedded in the corporate culture. What is the company's patching cadence? Are the password policies implemented well? Is the firewall configuration robust? What level of training and awareness exists at the user level?

**Lateral** - this axis examines connectivity and risk aggregation issues. Computers and mobile phones are communications devices and the internet connects each device to every other into a large lateral landscape of systemic risk (see view six). Network topography issues such as node concentration and network segregation are the key metrics to investigate here. Looking beyond the IT infrastructure to the business as a whole, single points of failure in the supply chain should also be evaluated.

**Temporal** - time is a crucial factor in assessing cyber risk; the faster the response time the better the damage limitation. There are well known metrics to use here such as MTTD (mean time to detection) and MTTR (mean time to resolve). In a wider business sense, the speed of executive decision making or of corporate communications in managing the news flow are other aspects to evaluate.

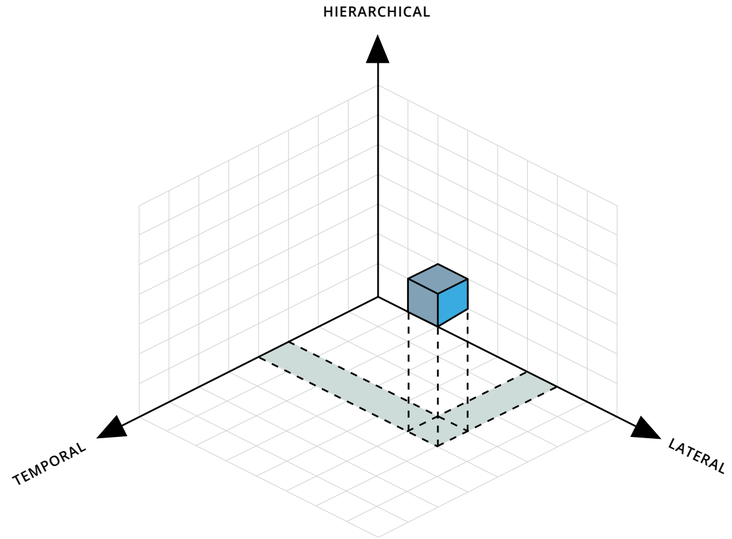

# 28. Assets: The Parkerian Hexad

Security can be defined as the degree to which your **assets** are resistant to **threats** from **adversaries**. We have explored two of these - threats and adversaries - in other diagrams (see views 18, 19 and 20) so it's now time to turn our attention to the third leg of the stool: assets.

In a traditional security model, assets have three attributes that require protection: Confidentiality, Integrity and Availability. This is known as the C-I-A model after the initial letters of this triad. However, in 1988 a cyber security expert called Donn Parker realised that there were other attributes that were important from a cyber perspective that the traditional model overlooked. So he added three more: Possession, Authenticity and Utility. Taken together these six attributes are now known as the Parkerian Hexad.

The diagram to the right shows the six security attributes in the inner blue ring and the method used to protect of each of these in the green outer ring. The green ring defends against attacks symbolised by the pink arrows. So, for example, ransomware is an attack on the availability attribute; access to data is denied until the ransom is paid. One defensive method against such an attack is good data backup discipline, combining both cloud-based storage and air-gapped hard drives.

**Confidentiality** – preventing unauthorised access to sensitive information by using encryption and data classification and clearance schemes (e.g. internal only, restricted and top secret).

**Possession** – retaining control of data and preventing unauthorised copying. Note that encrypted data can be lost without breaching confidentiality.

**Integrity** – ensuring that your data is unadulterated and has not been tampered with in any way. Hashing techniques generate a numerical value (known as a hash) from a string of text to check for data integrity.

**Authenticity** – this is proof of authorship. Did the message really come from that sender? Digital certificates and signatures are the best tools for establishing this.

**Availability** – data is useless if you can't access it when you need it. Establishing good data backup routines and avoiding single points of network failure through multiple firewall clustering are typical strategies here.

**Utility** – data can be theoretically available but still useless if it is in a format that cannot be read. An example is forgetting the password used to protect a spreadsheet, or data stored on an old-style minidisc if you don't have a minidisc player.

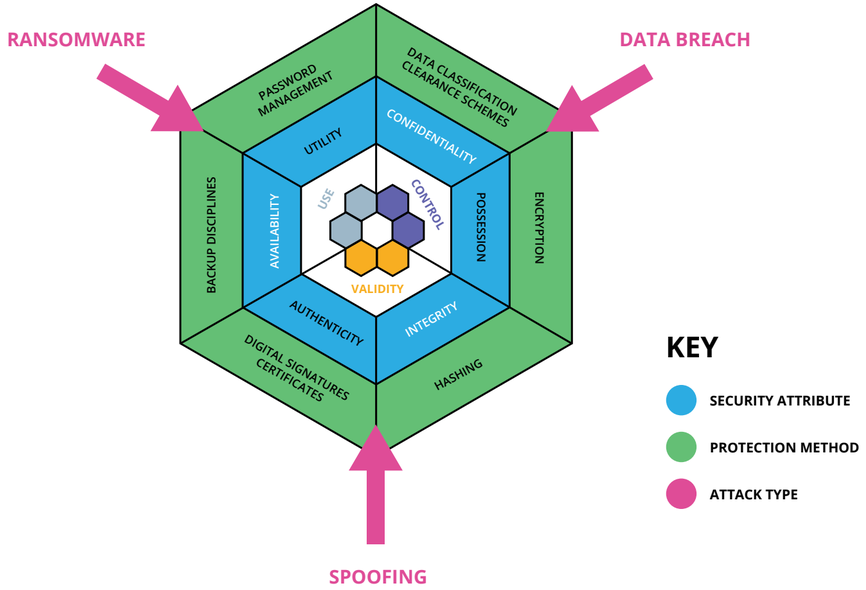

# 29. Assets: Industry variations

Cyber security assets do not only vary with type as in the Parkerian Hexad (view 28) but also across industries. That is to say that the most important assets will be different from one industry to another. The diagram to the right gives a very simplistic illustration of this. The blue circles in the diagram are assets that require protection arranged in the form of a crude corporate model. The vertical axis represents the internal vs external divide. IT systems are focused inwards while reputation is an asset determined externally by the marketplace. Likewise, on the horizontal axis, products sit upstream of customer records which are downstream. In the centre, with a grim inevitability, sits finance.

On the right-hand side of this diagram, there are four contrasting types of business. The most important and least important assets for each are shaded pink and green respectively:

**Commercial law firm:** If IT systems were to go down for a day or two in a law firm, it would be inconvenient but not catastrophic. It would still be possible to conduct business the old-fashioned way with a phone, a pen and legal pad. Law firms, however, are very vulnerable to reputational risk. Mossack Fonseca, a Panamanian law firm, went out of business in 2018 following a data breach that leaked details of widespread tax evasion by its numerous secretive clients. If a law firm loses the trust of is clients, it will cease to be. So reputation is coloured pink and systems coloured green.

**Grommet Manufacturer:** In a mirror image of a law firm, systems are vital but reputation for a B2B company is less of an issue. System failure will halt production, maybe for months; a significant problem. Reputational risk is relatively low. As an industrial parts supplier they will have little brand recognition from consumers and their handful of main industrial customers can be appeased in person.

**Dental practice:** Dentists sit on a large database of sensitive health records which would be extremely damaging if breached. Their core product, essentially an activity requiring precise manual work, is not highly IT dependent. A dentist can still drill and fill your teeth without a computer. So, customer records are pink and product is green.

**Wedding photographer:** If paid in cash on a jobbing basis, a wedding photographer may not need a sophisticated accounting system. The product, however, if the photos are all digital is completely dependent on an IT system. Losing a hard drive and its backups would be catastrophic, destroying the photographer's life's work. So product is pink and finance is green.

This oversimplified model is only intended as a crude illustration of the degree of variation across industry sectors. You will require a proper in-depth analysis of your particular requirements from an appropriately qualified professional.

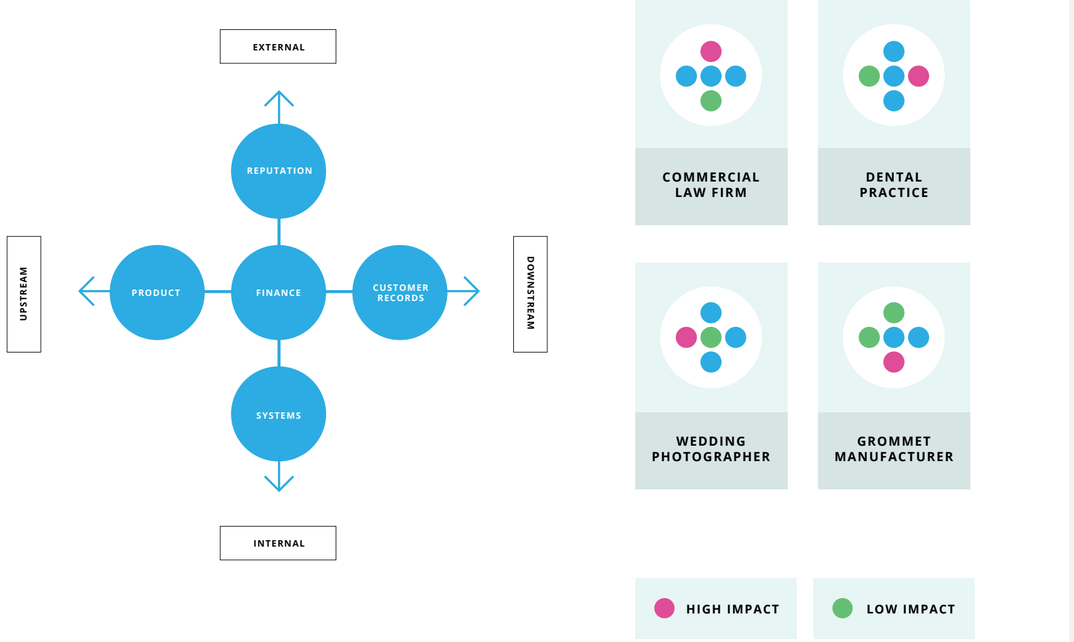

# 30. Parkerian Hexad: Insurance mapping

Insurance policies differ in wording and scope depending on the underwriter. Brokers, having fully comprehended the client's specific needs, provide valuable guidance as to the type of policy and the extent of cover that best fits these requirements. Good communication between the client, the broker and the underwriter at inception is essential to ensure there is no misunderstanding later if and when claims are made.

Below is a list of the types of insurance cover that are generally available from many underwriters in the market. This is not an exhaustive list, nor is it meant as a substitute for full and proper consultation with a qualified expert. In view 28, we explained the Parkerian Hexad model of asset security attributes. In the diagram above, we show how the different types of insurance cover match up with those six attributes.

**Data Restoration** – covers the costs associated with the replacement, repair or restoration of assets damaged in a cyber-attack or an accidental failure. This maps primarily to the **integrity** part of the hexad and to a lesser extent to availability.

**Breach Costs** – these are the costs of employing a forensic team to investigate the breach. They regularly include the costs associated with notifying customers of lost data, credit monitoring, call centre services and PR efforts. It maps to the **possession** attribute of the hexad.

**Privacy Liability** – covers third party liability for settlements arising from the failure to protect confidential information. It maps to the **confidentiality** part of the hexad.

**Cyber Extortion** – coverage for losses incurred from extortion, of which a ransomware attack is a good example. This maps primarily to the **utility** part of the hexad and to a lesser extent to confidentiality.

**Business interruption** – this covers lost income and extra expenses caused by the failure of computer systems and networks. It maps to the **availability** part of the hexad.

**Fraud and e-theft** – this is coverage of costs associated with theft or fraudulent transfer of funds and other property of value from cyber-attacks such as 'man in the middle' spoofing. It maps to the **authenticity** part of the hexad.

There are, of course, many other types of insurance cover available. Examples include cover for failure of a third party's systems like a cloud provider (see view 34) or reputational damage. The intention of the above diagram is simply to demonstrate that the six key security attributes for cyber defence have a direct counterpart in terms of coverage available in the insurance market.

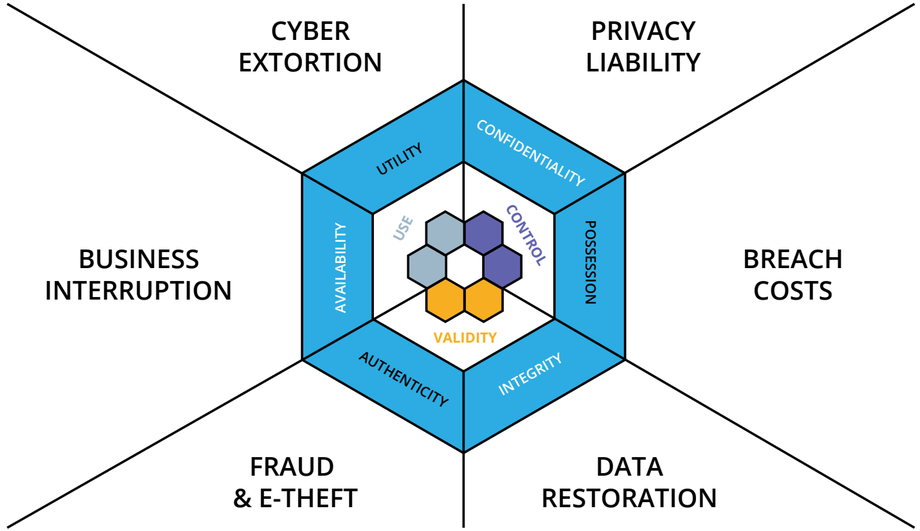

# 31. Capping risk

Companies buy insurance policies to limit their risks. On the other side of the fence, insurance underwriters want to manage the risks that they are accepting. This can sometimes lead to misunderstanding as to what is or is not covered in a policy. Through no fault on either side, the client could believe they have cover for a particular incident when a closer reading of the policy wording would show they do not.

In the interests of reducing this type of confusion, the diagram to the right shows the three main ways that underwriters cap the risks that they take on when writing a policy:

**Value** - a policy will often limit the total amount of money that will be paid out in the event of a claim. This can be is done through aggregate limits which cap the total or sub limits which cap a particular part. Most policies also have an excess or retention clause specifying that initial losses below a certain amount will be borne by the insured.

**Time** - just as with value, there can be time retentions too. This means a claim will only be paid out if the service outage is longer than a certain number of hours. This eliminates minor IT glitches and puts the focus on serious cyber incidents. A second consideration is how the timing of the incident relates to the coverage period. Policies can either be worded to cover losses occurring during that period or claims made. In the former case, the potential liability could extend for years beyond the policy expiry date as was the case with the asbestos settlements. Cyber coverage is typically written on a 'claims made' basis meaning any claim needs to be made before the policy expires. However, a short, extended reporting window post expiry is often included.

**Consequence** - a cyber incident in the upstream part of an industrial supply chain can have knock on effects running all the way downstream. A parts supplier taken down by a cyber-attack might fail to deliver a key component on time, ultimately delaying a major project several steps downstream and causing substantial losses. One way that underwriters seek to mitigate this type of exposure is through the careful wording of the business interruption part of the policy. This will define both the type and extent of the third parties that are covered, both upstream and downstream. See view 34 for a more detailed discussion of this.

Aside from these three axes, the final backstop that underwriters employ to limit their exposure is re-insurance where some part of their risk book is ceded to another insurer. This can either be on a facultative basis with a separate negotiation for each policy or on a treaty basis specifying a particular subset of the whole book.

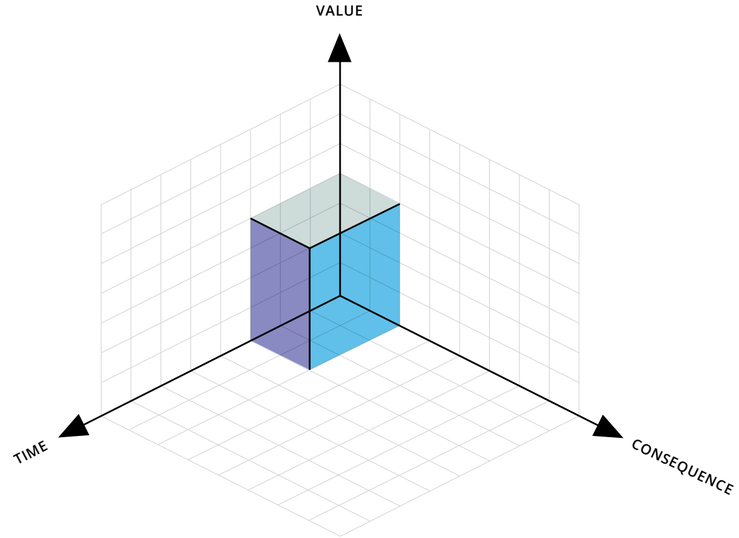

# 32. Cyber incident under-reporting

Many small and medium sized enterprises (SMEs) believe that a cyber incident is unlikely to happen to them because it is only big companies that are targets. Certainly, going by what is reported in the press, this might appear to be the case. However, an interesting white paper published in 2018 by AIR Worldwide, a risk modelling company, refutes this. The paper describes the AIR Probabilistic Cyber Model from which the data in the diagram is derived.

The blue line shows the incidents reported in the press. It is clear that unless the company's revenues are greater than $100m the press will not consider it a newsworthy story. However, just because SME incidents are not hitting the headlines, it does not mean that they are not happening. If you examine the claims data, you can discover the real story which is that plenty of SMEs [8] are getting hit by cyber criminals, even companies with revenues of $1m.

## SMEs do need cyber insurance

Notice that the left-hand scale shows breach likelihood. It seems that the likelihood of suffering a breach is broadly the same whether your revenues are $1m or $100m. Since there are a lot more small companies than big ones, there must therefore be a large number of small companies suffering cyber-attacks. The gap between the orange and blue lines shows the extent of underreporting in the press.

There is one other point to make. For a company to make a claim they must have had a cyber insurance policy in the first place. We also know that SME's are generally underinsured when it comes to cyber. So, the claims data is understating the problem. The message to SMEs is that getting cyber cover is a good idea.

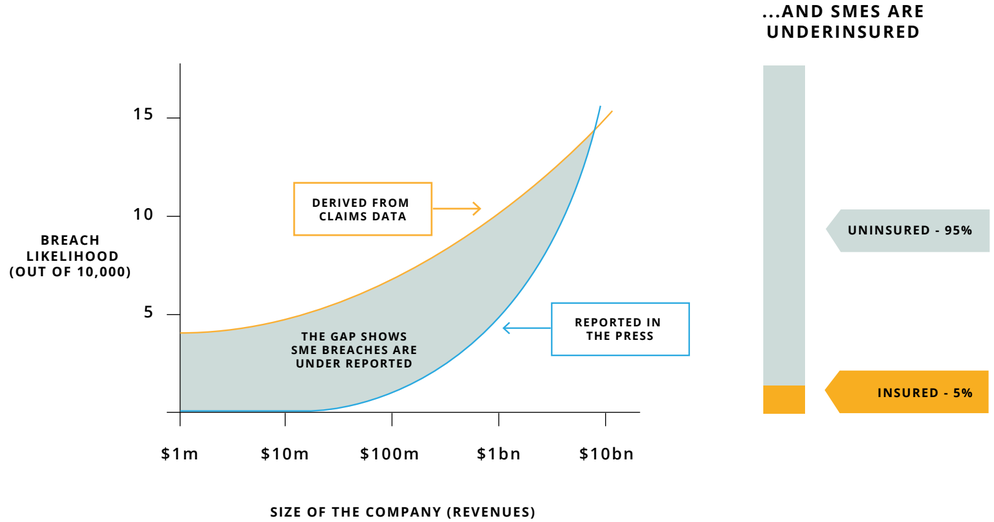


Source: AIR Worldwide

# 33. Cost of capital factors

In view 32 we examined the underreporting of cyber incidents amongst small and medium sized companies (SMEs), which may explain the lack of insurance cover in that segment of the market. Here we explore a different reason why it is a good idea for smaller companies to get coverage. It is an argument based on cost of capital illustrated on the right.

The orange line on the diagram shows, in conceptual terms, how insurance costs change as you move up the tower. An insurance coverage tower is made up of several layers of underwriters. The underwriter at the bottom of the tower (known as the primary) is first in line for any claims that might be made. Underwriters on layers higher up the tower are less likely to have to pay out claims because the other underwriters will be liable before any claim gets to them. Lower claim risk means cheaper pricing. So the orange line is a curve that flattens out; the higher up the tower you go the cheaper the cost to insure.

Large companies with plenty of capital often decide to self-insure, sometimes through a captive in-house entity, because the bottom layers of the tower can be expensive. The decision for large multinationals hinges around how their cost of capital compares with the cost of insurance. At a certain point up the tower, marked by a pink circle in the diagram, it makes sense to start buying insurance because the costs from that point on are relatively cheap.

## Steeper lines mean higher costs of capital

Smaller companies tend to have higher costs of capital than large multinationals. This is represented in the diagram by the fact that the dark blue line is steeper than the light blue one. If an SME has a high cost of capital, the dark blue line might never intersect the orange one meaning that self-insurance would be a poor choice. So, following this line of argument, it makes more sense for a small company to buy insurance that attaches at a lower point than a larger one.

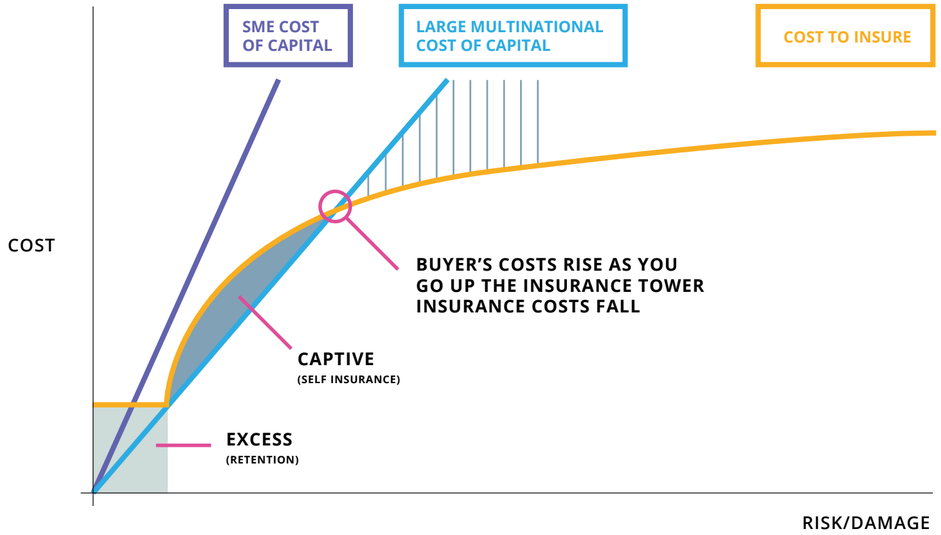

# 34. Cyber supply chain cover

One important way in which the cyber insurance market is maturing can be seen in the way in which coverage is changing. As little as five years ago, business interruption cover was not standard in cyber insurance policies. Back then, the market was mainly focussed on regulatory fines and other costs associated with data breaches in the USA.

Since then coverage has been steadily broadening as shown in the diagram. Business interruption coverage for the insured has become more or less standard in cyber policies now. Following demand from clients, the market is now beginning to offer coverage to mitigate the risks from cyber incidents in different parts of the supply chain both upstream and down. This comes in several different forms:

**Dependent Business Interruption (BI):** Typically, this would be limited to a named supplier. So, for example, a company that depended on Google G-Suite or Amazon Web Services for cloud hosting would want coverage for any system outage from these business-critical suppliers. However, there is a growing demand for full supply chain dependent BI cover where all suppliers to the insured are covered not just a named few.

**Receivers Business Interruption (BI):** This focuses on the downstream part of the supply chain. If a cyber incident means that the insured is unable to fulfil a contractual obligation to a customer and that customer then suffers damages, then receivers BI coverage will compensate.

**Other upstream coverage:** The number of different companies involved in an industrial supply chain from raw materials at one end to finished goods delivered to a consumer at the other can easily exceed 20. Coverage for the full length of this chain would be very unusual, not least because of the problem of tracking the liability across so many entities. Most insurance policies include critical infrastructure exclusions to limit the exposure to an event like the whole power grid going down. Upstream cover is also typically limited to entities with a contractual agreement which limits the coverage chain. Companies separated by a few steps in the chain are unlikely to have contract with each other.

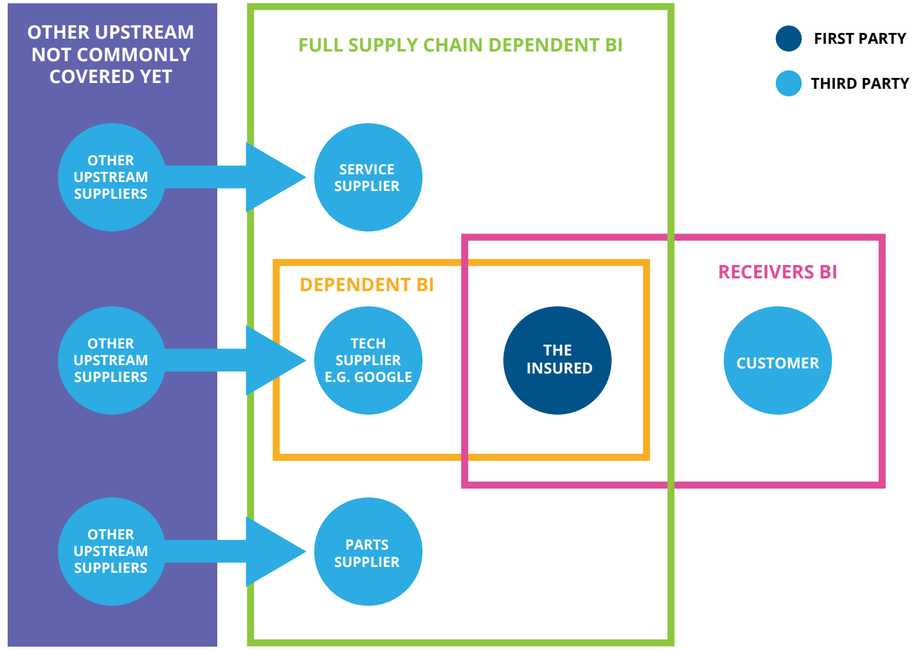

# 35. Return on security investment (RoSI)

One interesting feature of the cyber insurance market, at present, is the unusual ratio between spending on prevention and spending on insurance. In mature markets, such as property, the amount invested in fire prevention in a building and the amount invested in insurance cover are broadly similar. In the cyber arena, the amount spent on cyber security is around 20 times greater that the size of the cyber insurance market. One conclusion could be that the cyber insurance market will grow fast and eventually catch up.

A different conclusion could be that the amount spent on cyber security is disproportionally high and this has prompted some attempts to quantify the return on this investment. For insurance cover, the return on investment is explicit; this much cover for that much premium paid. For cyber security the RoSI is more opaque; how do you quantify the amount you have saved in future costs by spending on a penetration test today?
A white paper by Accenture on "The Costs of Cyber Crime" in 2017 attempted to answer this question and the results are summarised in the diagram to the right. In a survey consisting of 2,182 interviews of 254 companies in seven countries, participants were asked to evaluate the return on investment for nine categories of enabling security technologies. The results of their rankings are shown.

The broad conclusion that can be drawn is that three types of cyber security investments are particularly cost effective: security intelligence systems, machine learning for anomaly detection and user behaviour analytics. These technologies are more to do with monitoring and intelligence than traditional defensive barriers, which chimes in well with the squid model we discuss in view 10.

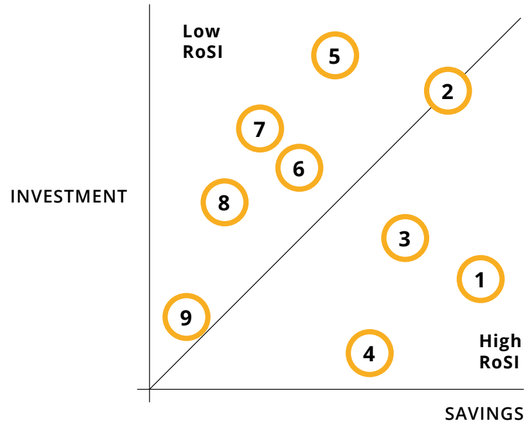


1. Security Intelligence Systems
2. Identity & Access Governance
3. Machine Learning
4. User Behaviour Analytics
5. Advanced Perimeter Controls
6. Encryption Technologies
7. Data Loss Prevention
8. Governance & Compliance
9. Automated Policy Management

Source: Accenture - Costs of Cyber Crime 2017

## About the author

John Donald graduated from Exeter University with an Engineering degree in 1983 and spent the next 20 years in Investment Banking both in the Far East and Europe. During that time he worked for Phillips and Drew, Jardine Fleming and ING Barings. He received the Institutional Investor award for No.1 ranking in Asian Equity Research for five consecutive years.

He left Investment Banking in 2004 and bought a Scotch Whisky company which he successfully sold on to investors in 2008. For the next 10 years, he ran his own risk consultancy advising on geopolitical and cyber risk. He joined Axis Capital as a Cyber Adviser in February 2019.

John is the author of two books. His first book “Catataxis: When more of the same is different” was published in October 2011 by Quartet Books. His second book “Bolt from the Blue: Navigating the new world of corporate crises” was published by Elliott and Thompson in 2013.

## Sources


1. Gartner Group (Public Cloud Services Market: 29 July 2019)
2. Gartner Group (Public Cloud Services Market: 29 July 2019)
3. Axis estimates
4. Idtheftcentre.org
5. Bloomberg News
6. Axis Estimates
7. Axis Estimates
8. AIR Worldwide

www.axiscapital.com

in

**CONTACT**

**AXIS Cyber Centre of Excellence**
**cybercoe@axiscapital.com**
**+44-20-7877-3800**