



Institute
and Faculty
of Actuaries

Cyber operational risk scenarios for insurance companies

Research project

By the Institute and Faculty of Actuaries' Cyber Risk
Investigation Working Party

Disclaimer: The views expressed in this publication are those of invited contributors and not necessarily those of the Institute and Faculty of Actuaries. The Institute and Faculty of Actuaries do not endorse any of the views stated, nor any claims or representations made in this publication and accept no responsibility or liability to any person for loss or damage suffered as a consequence of their placing reliance upon any view, claim or representation made in this publication. The information and expressions of opinion contained in this publication are not intended to be a comprehensive study, nor to provide actuarial advice or advice of any nature and should not be treated as a substitute for specific advice concerning individual situations. On no account may any part of this publication be reproduced without the written permission of the Institute and Faculty of Actuaries.

Title

Cyber operational risk scenarios for insurance companies

Authors

This paper was written by the Institute and Faculty of Actuaries' Cyber Risk Investigation Working Party. Membership of the working party and contributing authors are set out below.

R. Egan*	V-J. Jaeger
S. Cartagena	D. Katz
R. Mohamed	P. Meghen
V. Gosrani	M. Silley
J. Grewal	S. Nasser-Probert
M. Acharyya	J. Pikinska
A. Dee	R. Rubin
R. Bajaj	K. Ang

Abstract

Cyber Operational Risk

Cyber risk is routinely cited as one of the most important sources of operational risks facing organisations today, in various publications and surveys (Hubmann 2018) (Osborn 2018). Further, in recent years, cyber risk has entered the public conscience through highly publicised events involving affected UK organisations such as TalkTalk, Morrisons and the NHS. Regulators and legislators are increasing their focus on this topic, with General Data Protection Regulation (“GDPR”) a notable example of this.

Risk actuaries and other risk management professionals at insurance companies therefore need to have a robust assessment of the potential losses stemming from cyber risk that their organisations may face. They should be able to do this as part of an overall risk management framework and be able to demonstrate this to stakeholders such as regulators and shareholders.

Given that cyber risks are still very much new territory for insurers and there is no commonly accepted practice, this paper describes a proposed framework in which to perform such an assessment. As part of this, we leverage two existing frameworks – the Chief Risk Officer (“CRO”) Forum cyber incident taxonomy, and the National Institute of Standards and Technology (“NIST”) framework – to describe the taxonomy of a cyber incident, and the relevant cyber security and risk mitigation items for the incident in question, respectively.

Summary of Results

A table summarising the findings on each of the three scenarios investigated is below:

Table 1. Summary of scenario results

Scenario	Threat vectors	Most relevant security/risk control categories	Main cost components	1 in 200 Loss (£m, % of annual revenue)
Employee leaks data at a general (non-life) insurer	Insider attack, social engineering	Protect & respond	Compensation, regulatory fines	£210.5m (2%)
Cyber extortion at a life insurer	External attack, social engineering	Detect, respond & recover	Business interruption, reputational damage	£179.5m* (6%)
Motor insurer telematics device hack	External attack, software vulnerabilities	Identify, protect & detect	Remediation (device replacement)	£70.0m (18%)

*Note that further costs for this scenario have been explored in Section 3.5 although these do not form part of the Solvency Capital Requirement.

Limitations

The following sets out key limitations of the work set out in this paper:

- Whilst the presented scenarios are deemed material at this point in time, the threat landscape moves fast and could render specific narratives and calibrations obsolete within a short time frame.

- There is a lack of historical data to base certain scenarios on and therefore a high level of subjectivity is used to calibrate them.
- No attempt has been made to make an allowance for seasonality of renewals (a cyber event coinciding with peak renewal season could exacerbate cost impacts).
- No consideration has been given to the impact of the event on the share price of the company.
- Correlation with other risk types has not been explicitly considered.

Conclusions

Cyber risk is a very real threat and should not be ignored or treated lightly in operational risk frameworks, as it has the potential to threaten the ongoing viability of an organisation. Risk managers and capital actuaries should be aware of the various sources of cyber risk and the potential impacts to ensure that the business is sufficiently prepared for such an event.

When it comes to quantifying the impact of cyber risk on the operations of an insurer there are significant challenges. Not least that the threat landscape is ever changing and there is a lack of historical experience to base assumptions off.

Given this uncertainty, this paper sets out a framework upon which readers can bring consistency to the way scenarios are developed over time. It provides a common taxonomy to ensure that key aspects of cyber risk are considered and sets out examples of how to implement the framework.

It is critical that insurers endeavour to understand cyber risk better and look to refine assumptions over time as new information is received. In addition to ensuring that sufficient capital is being held for key operational risks, the investment in understanding cyber risk now will help to educate senior management and could have benefits through influencing internal cyber security capabilities.

Keywords

Cyber risk; Operational risk; Costs; NIST; Scenario

Correspondence details

*Correspondence to: Rory Egan, Chair of IFoA Cyber risk working party, c/o Institute and Faculty of Actuaries, 7th Floor · Holborn Gate · 326-330 High Holborn · London · WC1V 7PP, UK.

E-mail: rory.j.egan@me.com

1. Introduction

1.1 Aims and Terms of Reference

The Cyber Risk Investigation Working Party is a subgroup under the Institute's ERM committee. The group was established as a forum for actuaries to share insight and research, and to respond to cyber risk developments within the industry.

The group aims to support actuaries working on realistic capital calculations and/or within enterprise risk management for life and general insurers. In particular, the purpose of the research is to provide insight into setting out potential impacts of cyber events and the measures available to mitigate such risks.

The initial research conducted by the group focussed around deriving specific cyber risk scenarios that can be referred to when determining operational risk capital requirements for insurance companies. This was deemed to be a significant emerging issue given the ever-increasing dependency on data and information technology to support the business operations of insurers. Given the multitude of possible permutations for insurer type vs scenario narrative, the group quickly began to focus more generally on developing a proposal for a framework within which to build appropriate scenarios.

This paper aims to drive greater awareness of cyber as an operational risk for insurers through a proposed framework for scenario development and three worked examples. The three worked scenarios modelled within this paper are as follows:

- employee leaks data at a general (non-life) insurer (set out in Section 3.4);
- targeted ransomware attack on a life insurer (set out in Section 3.5); and
- motor insurer telematics device hack (set out in Section 3.6).

1.2 Definition of Cyber Risk

Cyber Risk is the risk of any financial loss, disruption or negative reputational impact because of a failure in information technology systems; whether through people, process or technology. According to the CRO Forum (CRO Forum 2016) cyber risk covers:

- any risks emanating from the use of electronic data and its transmission, including technology tools such as the internet and telecommunications networks;
- physical damage that can be caused by cyber-attacks;
- fraud committed by misuse of data;
- any liability arising from data use, storage and transfer; and
- availability, integrity and confidentiality of electronic information – be it related to individuals, companies or governments.

The risk is dependent upon the malicious (or non-malicious) threats the organisation faces and how organisations mitigate the risks through business and strategic decisions.

This paper does not consider cyber underwriting risk but rather the cyber risks that an insurance organisation is exposed to (i.e. operational risk).

2. Methodology

To drive greater awareness of cyber as part of an operational risk for insurers it is important to define and introduce a framework of analysis within which scenarios can be developed in a consistent manner. This section of the report proposes such a framework.

Each scenario set out in Section 3 has been designed and assessed in a consistent manner within this framework.

2.1 Defining a Common Taxonomy

A common taxonomy is of critical importance in ensuring consistency in the design and parameterisation of scenarios relating to cyber risk. There is a range of publicly available material aiming to bring consistency to this discussion. This paper highlights two specific sources of material:

- CRO Forum Concept Paper on a proposed categorisation methodology for cyber risk (CRO Forum 2016).
- NIST framework (National Institute of Standards and Technology 2018).

The taxonomy used within this framework has been created by leveraging information from these two sources.

2.1.1 Cybersecurity assessment taxonomy



The National Institute of Standards and Technology Cybersecurity Framework (“NIST framework”) has been developed to provide standards, guidelines and best practices to manage cybersecurity-related risk. It provides a guide for US private sector organisations to assess and improve their ability to identify, prevent, detect, respond and recover from cyber-attacks. A Gartner report cited that 30% of US companies have adopted the NIST framework with 50% expected by 2020 (National Institute of Standards and Technology 2016).

Given the NIST framework is focussed on providing guidance for ensuring cybersecurity resilience this research group has leveraged this work to define the cyber security vulnerabilities taxonomy.

The Securities and Exchange Commission “SEC” has stated its preference that NIST should be used as the standard for Cyber Security assurance for organisations which contribute to critical national infrastructure (Clayton 2017). It has expectations that companies meet the basics of this framework for regulatory purposes.

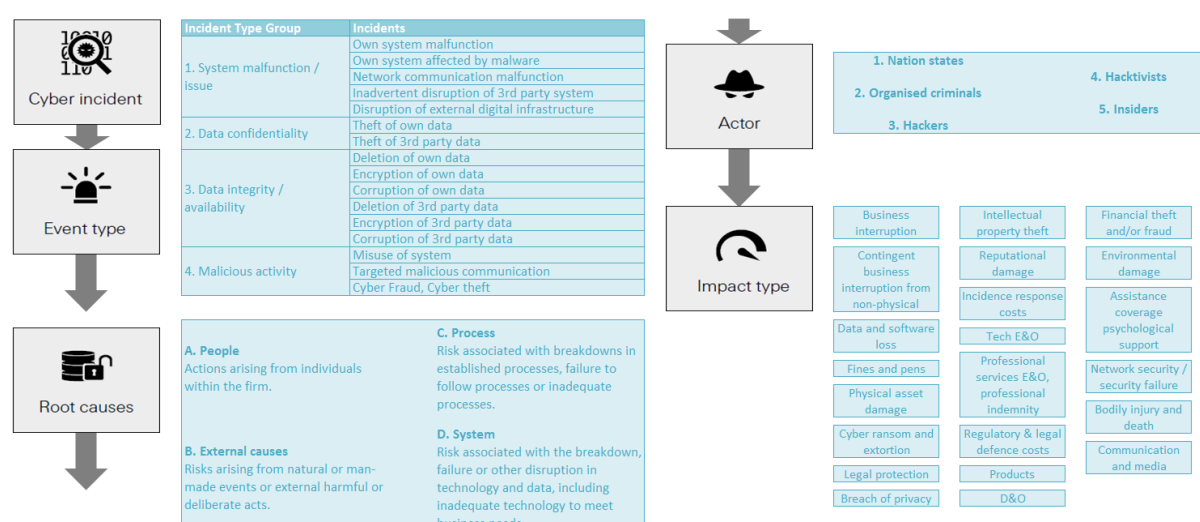
Within this framework of analysis, we have relied upon v1.0 of the NIST framework released in February 2014. It is worth noting that v1.1 was released in April 2018. The working party has reviewed the ‘Notes to Readers on the Update’ section of the accompanying report and determined that the updates do not have a material impact on this paper.

2.1.2 Cyber incident taxonomy

The CRO Forum concept paper proposes a methodology for a categorisation of cyber risk. The aim of the paper is to assist with data capture for cyber incidents. In particular the concept paper proposes categorisations for:

- cyber incident;
- event type;
- root causes;
- threat actors; and
- impact type.

Figure 2. CRO Forum concept paper; a proposal for cyber categorisation



Although the original aim of the concept paper was to support claims data capture, the categorisations have been useful when considering the design and corresponding economic impact of the operational scenarios presented in this paper. The CRO Forum categorisations have therefore been leveraged as the basis for the cost / impact taxonomy used within this research group's work.

2.2 Designing a Scenario

Given agreement of a common taxonomy (as set out in Section 2.1), operational risk scenarios can be developed consistently within a simple framework. A proposal for such a framework is set out in the remainder of this section. It is worth noting that this framework is independent of any individual scenario; examples of how to implement this framework are detailed in Section 3.

When defining a scenario, the organisation should first define their view of cyber risk (see Section 1.2 for the working party definition) and consider how any tangible or intangible losses could arise from failures in their cyber related processes. A key part of this assessment for an insurance organisation is to consider high value assets and/or or key weakness/dependencies that could lead to a significant business impact if a cyber risk were to materialise. A precursor to defining a cyber-operational risk scenario is having an accurate understanding of organisational maturity across all the fields in the NIST framework. Once the key tangible and intangible assets of the organisation are defined, relevant scenarios can be developed to understand the impacts of the key threats to the company. Some of these key considerations are discussed in the following sub-section.

2.2.1 Scenario selection

When designing an operational risk scenario, it is important to think through a range of factors relevant to the scenario including, but not limited to:

- structure and size of the company e.g. national/global;
- types of insurance products written;
- IT systems used within the business including dependencies/contingencies in place and third-party dependencies;
- volume and use of data stored within the company including internal data warehousing process and maintenance (e.g. are old records deleted/duplication of records, etc);
- type of data records stored (e.g. Personally Identifiable Information or 'PII', Payment Card Information or 'PCI', Protected Health Information or 'PHI');
- assessment of the company's current cyber resilience (useful to reference the NIST framework);
- current global cyber threat landscape e.g. active threat actors and prevalent threat vectors if applicable. Consider who and why different threat actors may want to attack you directly or whether you may be indirectly exposed to collateral damage from attacks on others e.g. NotPetya;
- company specific cyber threat landscape i.e. existence of factors which increase the motivation for a cyber-attack; and
- legal and regulatory framework the company is governed by.

Given the uncertainty, changing landscape and complexity of cyber risk it is recommended that key stakeholders from around the business should be consulted when considering the design and materiality of scenarios. This might take the form of workshops. The following is a non-exhaustive list of stakeholders who might be included:

- | | |
|----------------------|-------------------------|
| - ERM; | - HR; |
| - head of IT; | - board members; |
| - CISO; | - internal audit; |
| - procurement; | - supplier manager; |
| - cyber underwriter; | - COO; and |
| - legal; | - business dept. heads. |

The scenarios selected for quantification within this paper are detailed in Section 3 of this report. A useful position to start is to consider near missed events such as NotPetya/insider data leaks and consider how these could have caused a significant impact on the organisation.

2.2.2 Assessment against the NIST framework

Each scenario is assessed against an aggregated NIST framework which includes a total of 22 control categories across the 5 core functions; Identify, Protect, Detect, Respond and Recover (details of the control categories used are set out in Appendix 2). For a given scenario, the following steps are then taken for each control category:

1. Consideration is given to whether or not a control category is relevant to the scenario.
2. Assessment of cost types which could be impacted by failure of the given cost category.
3. Qualitative assessment of potential impact of the event of failure of a control; consideration is given to both frequency of event and severity of event.

This exercise is uncertain by nature given the subjectivity involved. The purpose of this assessment is to help focus the outcome of the scenario; in particular the potential for scalable costs and areas of mitigation.

2.2.3 Costs estimation approach

Once the cost types impacted by the scenario have been identified the next step is to quantify an estimate of each loss amount. Estimation is completed through group discussion with reliance placed on members' own experience and understanding of losses. For each identified cost type the following sources of information have been used to inform the calculations:

- database of prior events (e.g. NetDiligence, Ponemon, Verizon);
- publicly available reports; and
- expert judgement.

2.2.4 Mitigation assessment approach

There is no quantitative assessment of the impact of potential risk mitigation mechanisms due to the uncertainty associated to the cost estimates and likelihood of breaches. However, a qualitative assessment is performed against the NIST framework to identify which areas and controls would be most relevant to focus mitigation efforts to ensure reduction of the potential risk of the event.

The approach taken within this exercise is to identify the high-risk control areas and summarise what reasonable mitigation attempts would look like. A more detailed assessment would include quantification of the impact each mitigation mechanism would have on each cost estimate. An assessment would also need to be completed to understand the cost benefit analysis of these techniques against alternative risk transfer mechanisms such as insurance policies.

3. Scenario analysis

Section 2 of this paper set out the working party's proposal for the framework within which cyber operational risk scenarios can be developed. Section 3 provides working examples of implementing this framework; detailing 3 scenarios including narrative of the event and estimated costs.

It is worth highlighting that there is a vast range of potential cyber operational incidents and some resulting costs are largely untested and therefore uncertain (e.g. GDPR fines). The example scenarios set out in Sections 3.4 - 3.6 should be seen as illustrative examples rather than a robust model for readers to use blindly.

Each scenario team worked independently during the parameterisation process which highlighted differences in views around impacted cost types and quantification. Whilst an exercise has been conducted to ensure reasonable consistency between scenarios, any apparent differences represent the underlying uncertainty inherent in this risk and the fact there is currently no clear single industry-wide consensus on how the risk should be approached.

3.1 Scenarios selected

The following scenarios were selected for the purpose of this paper:

- employee leaks data at a general (non-life) insurer (see Section 3.4);
- targeted ransom attack on a life insurer (see Section 3.5); and
- motor insurer telematics device hack (see Section 3.6).

The three selected scenarios were selected from an original set of seven through group discussion. They were selected as being the most relevant scenarios to the insurance industry from an operational risk perspective given the current risk climate. All scenarios considered are detailed in Appendix 1.

3.2 Return Period

For the scenarios in this paper we have chosen to target a 1 in 200 year event measure for each hypothetical company given that operational risk scenarios for capital purposes would generally aim for an event at this return period (in line with Solvency II). Given the significant uncertainty in estimation (lack of historical/public data) we consider the events discussed to be extreme but plausible and that the range around the estimate would be significant depending on the company, jurisdiction and market conditions.

The working party would encourage the reader not to place sole focus on the specific numbers reported in the following sections. The key takeaway is intended to be the framework and methodology for constructing such scenarios with the intent of equipping the reader to produce scenarios relevant to their own business.

The working party also recognises the difficulty in rationalising a 1 in 200 year scenario and thus readers should also consider creating and analysing scenarios at more frequent return periods, and then extrapolating.

3.3 Expected cost calculations

Estimated costs have been derived using a combination of research of current consultant rates, historical events and expert judgement. References have been provided where appropriate and it can be assumed that expert judgement was applied where no reference is provided.

It is worth noting that some costs are likely to be variable by the size of the company (e.g. compensation depends on customers exposed) while some other costs may be considered more fixed (e.g. some regulatory fines or consultancy costs dealing with the incidence & response). Readers of this document should assess the appropriateness of each cost estimate given the characteristics unique to their business.

There is significant potential for economic impacts on insurers beyond those which would form part of the operational risk capital charge (e.g. loss of future sales). While this report focusses on those costs forming the capital charge, Scenario 3.5 looks in more detail at some of these other costs due to the potential materiality to the insurer in that given scenario.

3.4 Employee leaks data at a general (non-life) insurer

3.4.1 Scenario

A general (non-life) insurer writing a diverse business including a large motor portfolio is hacked by an internal staff member. Details of all motor insurance policyholders are leaked onto an internet website and are widely available.

3.4.2 Description of the insurer

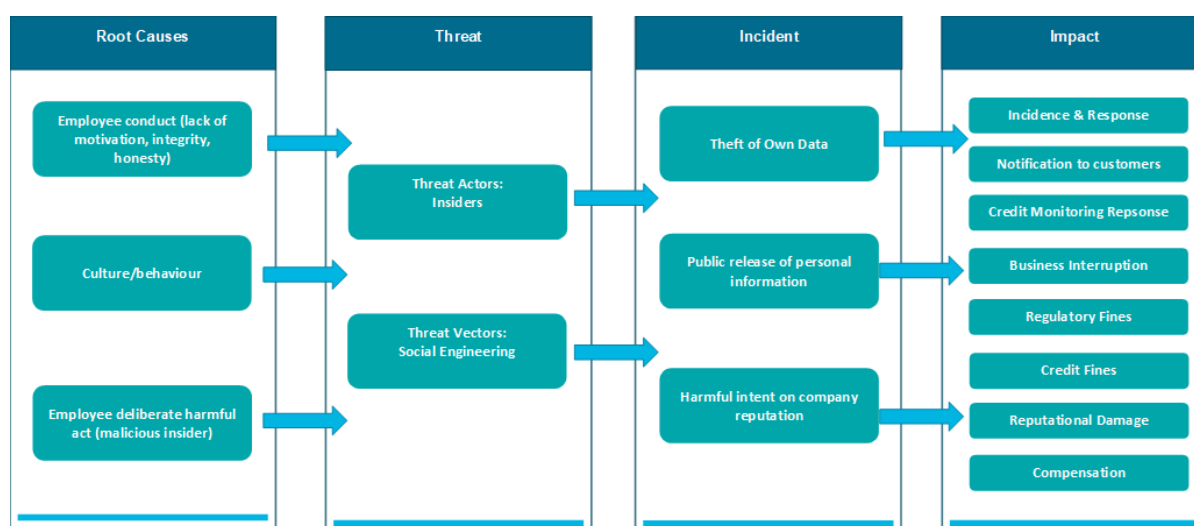
The insurer has a global presence, with over £10bn in revenue. The UK motor insurance book is a major unit of the insurer, with £1bn annual premium. The UK motor insurance portfolio contains 4m data records, with 3m policyholders on risk and 1m legacy records.

3.4.3 Event narrative

An employee had a poor working relationship with their manager. Low morale led to resentment and the employee decided to take harmful action. The employee published all motor insurance policyholder data online, both financial and non-financial. They accessed financial data including credit card information by persuading other employees to give access a few weeks' earlier using social engineering techniques. The data leak was noticed by a policyholder who called the emergency claims team. This did not get escalated appropriately and it took another day before key staff members were aware of the data breach.

Slow response and poor communication with the public led to a backlash from policyholders who took to social media to vent their anger. Employees also shared their opinion on social media around poor working practices. Investors, concerned at the poor controls in place and potential reputational damage to the remainder of the business, sold shares resulting in a 5% drop in share price overnight.

Figure 3. Incident summary for employee leak scenario

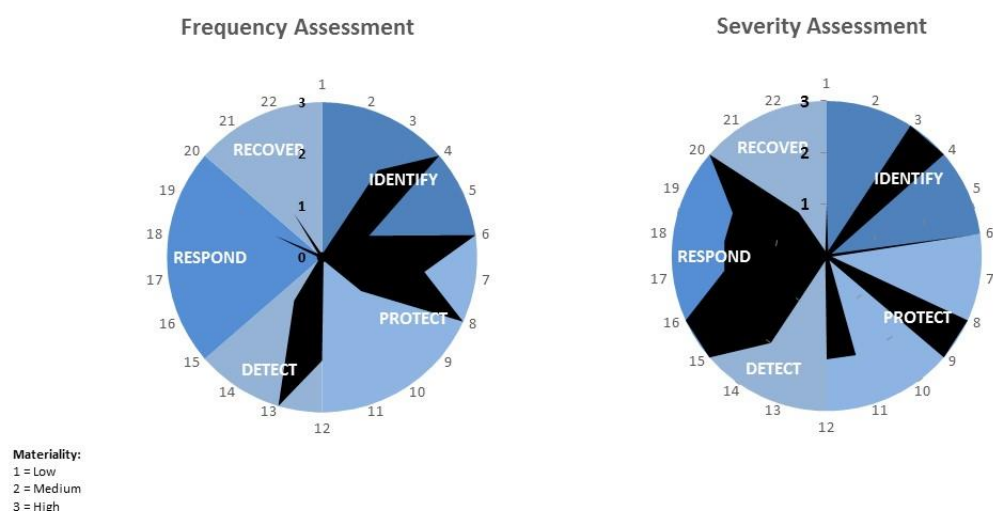


3.4.4 Security assessment and mitigation

The following charts display the assessment of this scenario's vulnerability across the NIST framework for the impact on frequency and severity of the event and indicate that the following control areas are expected to be the key vulnerabilities for this scenario:

- protection e.g. access controls, data security and information protection processes; and
- respond e.g. response planning, communication and improvements.

Figure 4. NIST framework assessment for employee leak scenario



3.4.5 Expected Costs

The table below summarises the expected costs considered relevant to this scenario. The costs presented are only estimations of the potential magnitude given the specified parameters of this scenario.

Table 5. Expected cost summary for employee leak scenario

	Cost type	Scenario cost	Approximate cost (gross)	Rationale
1	Incident response costs	External consultants used to investigate data breach.	£1.0m	1-month consultancy fee for detection/escalation, forensic costs of 2 months for tracking activity of user(s) and understanding extent of access / breach. Assume approx. £5,000 per day for consultancy fees and load for charged expenses. PR response (possibly performed in house for large companies), assumes 3 months of PR help on an assumed hourly rate of £220 (Gould + Partners 2014).
2	Incident response costs	Notification costs - people resource cost to notify parties affected by incident.	£5.5m	Number of customers affected combined with assumed average notification cost per customer (£1.40 per policy, based on Net Diligence findings (eRiskHub n.d.)) Includes - emails / letters, call centre & response team.
3	Incident response costs	Credit monitoring services offered to all customers for one year.	£6.5m	Credit monitoring costs associated to the PCI/PII data lost. Anthem (Wikipedia 2015) agreed cost is used as a benchmark but we have assumed each affected customer in this scenario would be an approximate cost of \$2 per person based on expert insight. No allowance is made for economies of scale.

	Cost type	Scenario cost	Approximate cost (gross)	Rationale
4	Business interruption	Business interruption – systems taken offline for maximum two days.	£0.5m	Two days of profit impacted assumed with a 95% combined ratio on 1bn annual revenue. There is uncertainty as unknown seasonality impact i.e. timing of the BI could have very different impact throughout the year based on when policies are renewed, assuming minimal impact for motor business. Assumed no contingent business interruption impact but applied an increased cost of working load of 50%.
5	Regulatory Fines	Fine for loss of customer exposure data – assumed failure to comply with GDPR rules.	£40.0m	£10bn revenue x 0.4%. Largest fine in UK to date is Facebook at £500,000 = the maximum possible, pre-GDPR. Assuming 80 times fine level under GDPR, then the max would be 80 * 500k = £40m. Under GDPR, can fine up to 4% of revenue; however this may seem too extreme a step change, especially as there has only recently been the first instance of a maximum fine under pre-GDPR data protection laws.
6	Fines	PCI breach fine and non-compliance fine - all fines incurred through non-compliance with PCI data security standards requirements.	£1.0m	Assumed a fixed cost of £100k each for PCI Forensic Investigator (“PFI”) investigation and Qualified Security Assessor (“QSA”) assessment (IT Security Expert 2017). Average PCI fine per lost record * number of customers affected but capped at £1m.
7	Regulatory Fines	Financial ombudsman fine.	£25.0m	Assume 1% of policyholders complain to Ombudsman with average cost of £600 to company.
8	Compensation	Liability compensation to policyholders and claimants - loss of claims data and with it health information.	£130.0m	1% of customers suffer financial loss of £1,000, plus £30 voucher given as compensation to all customers. Assumed 75% usage of vouchers.
9	Regulatory costs	S166 into how breach occurred and validity of actions taken to remediate weaknesses and avoid future occurrences.	£1.0m	The costs of S166’s have ranged from £30k to £1.3m in 2017. Given the nature of the event we assume this would be at the higher end.
Total			£210.5m	

This scenario represents a cost of approximately 2% of the company’s total revenue. Following an employee data leak we would expect there to be a reputational impact to the company that would impact future business and potentially the share capital. For motor insurance we consider it unlikely that there would be significant lapses for in-force policies following the event, however there may lower renewal and new business rates at renewal period. Hence reputational damage may occur and will depend on the PR handling by the company and/or remediation efforts following the event but, for this scenario, we have not quantified any short term reputational damage to premium volumes.

The key drivers of expected loss within this scenario are regulatory fines and compensation. It is worth highlighting the heightened uncertainty around the GDPR fines given that the legal and regulatory environment is currently untested. For the purposes of this scenario a worst-case outcome was assumed and hence the mitigation actions proposed would help to manage the risk.

3.4.6 Mitigation

The impact and ability to mitigate the risk is dependent on the following key areas (as labelled in the NIST framework):

- protect; and
- respond.

The table below summarises some of the possible mitigating actions that could be taken to limit the potential risk associated with this type of scenario. For this scenario, the protection controls are those likely to have the greatest mitigating impact (in terms of both the likelihood and the severity) on the potential losses facing the company.

Table 6. Proposed mitigation approach for employee leak scenario

NIST function	Mitigation type	Examples	Mitigating benefit
Protect	Control access	Password controls for all databases (policy, claims).	Each employee only given access to data that they need. For example actuarial staff do not need access to personal details. This makes such a widespread data breach more difficult.
		Limit access to all (physical and digital) assets.	
		Access (within the office or remotely) is managed, monitored and audited.	All access is monitored and managed - this makes breaches more "trackable" providing disincentives for employees to directly or indirectly be involved with potential data misuse.
	Staff training	Training relating to data protection laws and corresponding penalties for breaches.	Establishing a culture where each employee understands that they have a role to play in reducing cyber risk can also mitigate the risks associated with this type of scenario.
		Cyber security training for those who monitor network usage.	
		Incentives for reporting problems, concerns and whistleblowing.	Educating employees so that they are able to spot potential "warning signs" (e.g. line managers/other team members/IT staff) as well as the importance of accountability (e.g. the importance of following correct procedures
		Personnel screening during recruitment processes for "cultural fit". Breaches to confidentiality agreements included in staff contracts.	

NIST function	Mitigation type	Examples	Mitigating benefit
	Secure networks	Incident response plan preparation and training (including Board level).	especially when relating to data access and system permissions).
		Adequate information protection processes and procedures in place.	Securing networks sufficiently to make mass data access, downloading and transferring difficult; thereby reducing the frequency of potential data breaches.
		Removable media is protected (e.g. no USB ports available for use). Access to personal emails/websites restricted.	
		Logged use of company networks.	Introduce tighter email restrictions to include filters that block the sending of non-encrypted data e.g. national insurance ("NI") numbers.
		Networks monitored with automatic notifications in events of potential misuse taking place.	
	Data security	Regular reviews of the controls around systems and access.	Ensuring that all data regulations are being adhered to (and any changes to regulations are monitored on a regular basis) to avoid amplifying the potential costs of such a scenario by the exposure of non-compliance.
		PCI standard must be complied with, including anonymising and tokenising data. All data should be encrypted on transit and at rest.	
Respond	Effective response plans	Effective incident response plans with employees knowing their roles and the order of operations.	To reduce the risks relating to business disruption as well as regulatory action (e.g. of not informing within 72 hours of breach).
		Incidents are reported (to all relevant stakeholders) in a timely manner in line with response plans and regulations.	
		Purchase of cyber insurance.	
	Containment of event	Business continuity plans in place.	Work has already been done prior to the event (as part of business continuity planning) to understand which systems are required for the business to continue operating and which can go down (i.e. to limit the risk of further breaches whilst
		Consultants have already identified "choke points" in the organisation to understand how quickly systems can be back up and running.	

NIST function	Mitigation type	Examples	Mitigating benefit
			investigations are ongoing) with no significant revenue impact.
	Analysis and improvement	Automatic notifications from detection systems set-up. For example: - monitoring of data access with detection systems in place to notify when large amounts of data has been downloaded/uploaded; and - monitoring of employees' login and logout times especially during out of hours.	This makes detection of potential breaches easier thus allowing for appropriate response plans to be triggered. Time spent after the incident regarding "lessons learned" and potential improvements that can be made to processes to minimise costs in the event of a similar scenario occurring in the future.
		The impact of the incident is understood as well as lessons learned. Response strategies are reviewed in response.	

It is worth commenting that data breaches could occur in several different ways, such as an external hack. It is likely that these scenarios would produce different loss estimates, and different recommendations on how to mitigate the risk (such as the need for penetration testing and security around third party vendors). Although less likely, internal threats may have a greater financial and reputational impact to a company, as evidenced by the Morrison's case (Paatz 2018). At a 1 in 200 return period, we would want to consider more extreme events and hence have focused on internal threats.

3.5 Cyber extortion at a life insurer

3.5.1 Scenario

A life insurer is subject to a ransomware attack following a successful targeted spear-phishing campaign by hackers.

3.5.2 Description of the insurer

The insurer is a subsidiary of a FTSE100 listed financial services group. It has gross written premiums of £3bn, and an annual profit of £300m.

The company has historically relied on legacy IT systems to manage its customer portfolio data, but has recently begun an IT transformation programme to modernise its systems. It has agreed an outsourcing arrangement with a data services company to develop, test, maintain and support new technology applications, both during and after the transformation phase. Back-up systems are linked to the core systems to allow for continuous back-ups.

3.5.3 Event narrative

A group of hackers carry out a co-ordinated series of attacks against the insurance companies via a sophisticated and tailored spear-phishing campaign. This allows them to obtain employee logins and passwords for corporate systems. The insurer in question is one of the targets. For this company, the hackers go undetected for several months, during which they use these credentials to move laterally throughout the corporate network and are able to identify the new back-up procedures and stored backup files.

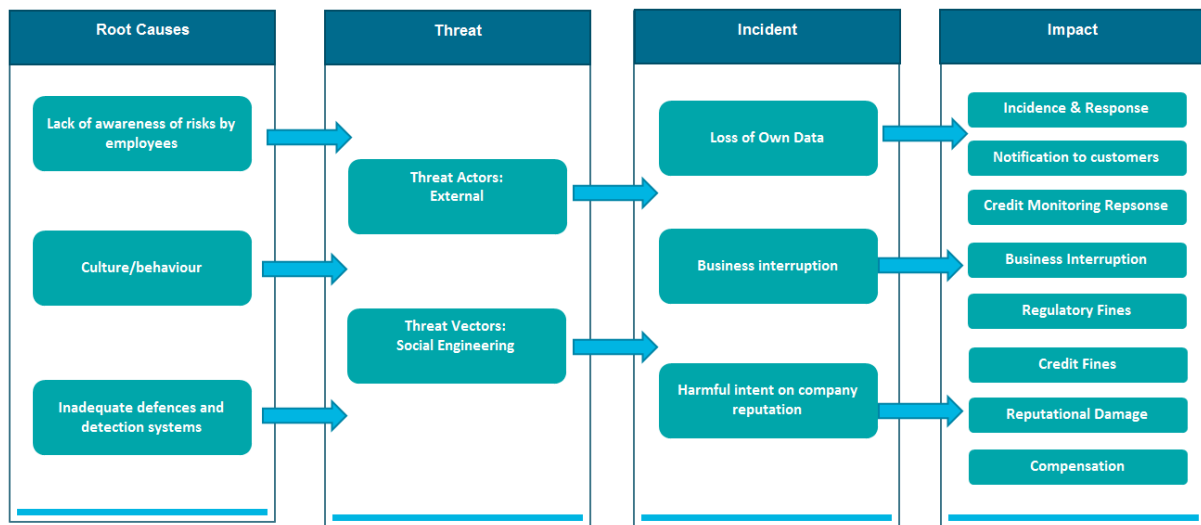
The ransomware worm is then delivered covertly and infects almost all of the insurance company's systems including both production and backup environments.

Upon launching the attack, operating systems become unavailable; critical systems and services are inaccessible and data is encrypted. In effect all operations grind to a halt. A request for a ransom payment of £15m is received to unlock all systems.

The firm calls an emergency management meeting and decides that given the dire situation of all systems and data including backups, being subject to the ransom the best course of action is to pay the ransom. Following investigations, the company identifies the critical systems held to ransom and a revised ransom figure of £7.5m is paid to the hackers. However, unexpectedly; the payment of the ransom does not result in the decryption of data. It is not known whether that was the intention of the hackers or not, but the resulting impact is that a huge data recreation, malware decontamination and IT systems restoration effort is needed. As the insurer is still in the middle of the IT transformation project, the restoration work is far more complex.

The incident has a huge impact on the firm's business through interruption and increased cost of working as many employees cannot do their jobs and are sent home. The media focuses on the poor internal controls of the firm, in particular that the lack of network segregation led to the ransomware worm spreading quickly across the network. The reputational fallout is catastrophic as many customers are not able to check their balances, let alone conduct any transactions, and the firm suffers a significant drop in sales as well as regulator scrutiny.

Figure 7. Incident summary for cyber extortion scenario

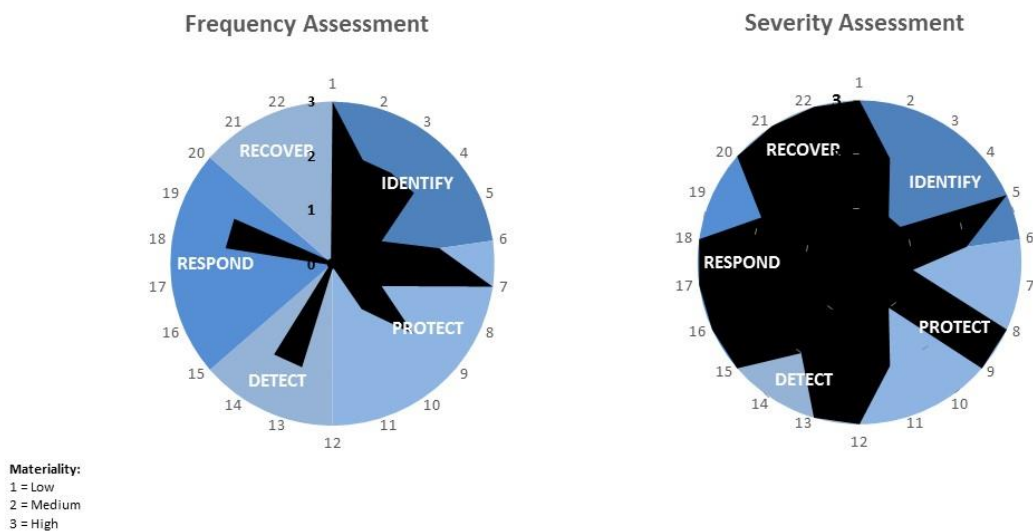


3.5.4 Security assessment and mitigation

The following charts display the assessment of this scenarios vulnerability across the NIST framework for the impact on frequency and severity of the event and indicate that the following control areas are expected to be the key vulnerabilities for this scenario:

- detect e.g. security continuous monitoring and detection processes;
- respond e.g. analysis, mitigation and improvements; and
- recover e.g. recoverability and communications strategy.

Figure 8. NIST framework assessment for cyber extortion scenario



3.5.5 Expected Costs

The table below summarises some of the expected costs for this type of scenario. These costs are only indications of the potential magnitude of each cost area for the specified parameters of this scenario.

Table 9. Expected cost summary for cyber extortion scenario

	Cost type	Scenario cost	Approximate cost (gross)	Rationale
1	Ransom Costs	Payment of ransom.	£7.5m	Recent demand on HBO was \$6m. Uplift for 1 in 200 scenario.
2	Incident response costs	IT forensics, crisis management, communications.	£1.5m	Based on UK consulting fees for IT and PR experts.
3	Data restoration	Restoration project (malware decontamination, data restoration / recreation, system rebuild).	£10.0m	Influencing factors include number of employees (number of workstations to fix) and complexity of IT (more servers, more complex networks, more outsourcers etc. to a bigger clean up job).
4	Business interruption	Expense risk, including productivity loss due to data centre outage, transaction delays, which require rectification, unbudgeted overtime and temporary staff costs.	£33.0m	We have assumed 2 weeks of full outage, and further 2 weeks at 50% outage before systems are fully restored in this severe event, with reference to the NotPetya attack which crippled companies' operations for several weeks (Novet 2017). Ponemon 2016 Cost of Data Center Outages report (Ponemon Institute 2016)suggests an average cost of \$9000 per minute during an unplanned outage. We have used this but removed the component relating to incident response and data restoration costs to avoid double counting.
5	Business interruption	Increased liability due to delays with processing.	£1.5m	2 weeks delay for processing of claims over period, with a small minority seeking substantial compensation.
6	Regulatory fines	PRA and FCA regulatory fines for operational resilience failures.	£5.0m	RBS fines in 2012 were £56m (BBC 2014) for a significant system outage. For a large life insurer, there would be a lower impact on the daily lives of customers, so a smaller but still significant fine could be expected, due to recent increased focus on cyber security.
7	Regulatory costs	S166 into how breach occurred and validity of actions taken to remediate weaknesses and avoid future occurrences.	£1.0m	The costs of S166's have ranged from £30k to £1.3m in 2017. Given the nature of the event we assume this would be at the higher end.
8	Business interruption	Lapses on in-force policies, reducing own funds through loss of net present value of future profits.	£120.0m	40% lapses per 'mass lapse event' approach in Solvency II lapse risk calculation (Boros 2014). 40% of revenue x 10% assumed profit margin foregone.
		Total	£179.5m	

This scenario represents a risk capital charge of approximately 6% of the company's total revenue. However, it is important to note that this excludes any impact from a data breach scenario, which is dealt with in Section 3.4, though hackers could steal as well as corrupt data.

The key driver of expected loss within this scenario is business interruption combined with regulatory fines and compensation costs, this scenario could give rise to severe losses. For this scenario, significant improvements in the ability to segment critical systems, improve defences and promptly detect unauthorised behaviour are critical to the outcome. The mitigation actions proposed would help to manage the risk.

As well as the losses above, the reputational damage resulting would give rise to loss of future sales in addition to those losses that typically make up the Solvency Capital Requirement. Nonetheless these result in significant additional economic impacts on the insurer which have been explored in the table below.

Table 10. Additional expected cost summary for cyber extortion scenario

	Cost type	Scenario cost	Approximate cost (gross)	Rationale
1	Reputational damage	Loss of future sales and goodwill	£150m	Assuming loss of 50% profit due to length of time incident was undetected

3.5.6 Mitigation

The impact and ability to mitigate the risk is dependent on the following key areas (as labelled in the NIST framework):

- detect;
- respond; and
- recover.

Key mitigation actions include network segmentation, patch controls, vulnerability scans, i.e. having appropriate detection processes and testing in place to help to identify the leak early on, ensuring the situation can be tackled as it arises and therefore reducing the impact of any attack. In addition, it is important to have an incident response plan in place, covering areas such as a decision tree for payment of ransom, a communications strategy and consideration for any external support which could be required to assist with the resolution of any incident.

Circuit breaker back-ups could help to mitigate impacts. This works through one of a pair of back-up systems being connected to main systems, with the other not being connected at all; then switching over. This stops the back-up system becoming infected.

Staff should receive training to make them aware of phishing attacks and assist them in identifying and flagging potential attacks. I.T. systems should scan incoming communications to try to eliminate or quarantine potential attacks.

The table below summarises some of the possible mitigating actions that could be taken to limit the potential risk associated with this type of scenario

Table 11. Proposed mitigation approach for cyber extortion scenario

NIST function	Mitigation type	Examples	Mitigating benefit
Detect	Anomalies and events	Model trends in standard behaviour, to incorporate detection processes which identify anomalies from this trend, which could identify unauthorised activities.	If it is detected in a timely manner it is highly likely that the company can take appropriate actions to stop it from spreading to wider networks/backups.
		Monitor unusual access requests.	
	Security continuous monitoring	Ensure logs are reviewed in real time (i.e. 24/7 monitoring).	Ongoing near real time analysis helps with early detection of security threats and enables companies to respond to security attacks quicker thereby reducing their impact on the business.
Respond	Detection processes	Run penetration testing at least annually to identify any vulnerabilities in security.	Having an appropriate detection processes and testing in place can help to identify the leak early on, ensuring the situation can be tackled as it arises and therefore reducing the impact of any attack.
		Carry out frequent vulnerability scanning activities to detect emerging security weaknesses.	Carrying out regular testing ensures that new vulnerabilities are detected and managed throughout the year.
	Response planning	Establish a decision tree for settlement of ransomware should an event occur.	Spread of the ransomware throughout the network could be limited by quickly executing the response plan.
		At a minimum agree T&C's for a cyber expert on retainer to be available immediately should an incident occur.	Leveraging external resources when required provides a balance between having experts onsite to support complex incidents without retaining them within the organisation full time.
	Communications	Identify who will handle media/PR and member communications.	Effective communication is vital during the response to ensure the plan is coordinated effectively to limit the damage.
		Establish alternative means for communication, for example if email is compromised.	
	Analysis	Identify and quantify key risks for the business, using the expected costs analysis.	Early analysis of the issue will help reduce the cost of the response to the incident.
	Mitigation	Consider cyber and crime insurance.	Risk mitigation ensures further aggravating occurrences of the incident are avoided.
	Improvements	Regularly test and improve incident response plan.	Lessons learned may be key for limiting the damage caused by future incidents.
		Include all senior managers (IT/Risk/Finance etc.) in tabletop exercises to run through a simulated incident.	Cyber risk spans across the organisation and therefore requires buy-in and response holistically rather than relying on one department to manage cyber risks.

NIST function	Mitigation type	Examples	Mitigating benefit
Recover	Recovery planning	Embed recovery protocols in the organisation and regularly test these (including third party validation).	Faster recovery reduces the impact of the incident.
	Improvements	Establish feedback protocols for review by management and improving processes for future incidents.	Lesson learned may be key for increasing recovery time for future incidents.
	Communications	Establish a clear communications plan, covering PR, internal and external messages.	The speed of recovery will be dependent on the public's perception of the way the business has handled the incident.

3.6 Telematics device hack at a motor insurer

3.6.1 Scenario

A motor insurer deploys telemetry in customer vehicles for measuring driver patterns using a specific telemetry device. A security researcher publicises a hack on this device that allows anyone with internet access to remotely access images from the camera of the telemetry device as well as the location and PII data on them. The insurer needs to recall / replace / replenish the device with each of its clients.

During the course of the recall, a number of hostile hackers break into the devices and publish data including locations, pictures and journeys of high profile policyholders who have installed the devices in their vehicles.

3.6.2 Description of the insurer

For this scenario, we have assumed it will affect a medium sized UK only motor insurer with many motor insurance policies issued using telematics devices.

The insurer has premiums of £400 million p.a. with a fleet of 500,000 cars using its telematics device. There is an average premium of £500 per annum per client for the telematics product, resulting in c£250m premium p.a. for the telematics product.

3.6.3 Event narrative

All 500k telematics devices get hacked, rendering the devices (costing c£50 each) unusable or untrustworthy. Every device needs to be recalled and replaced.

Sensitive data from the devices is compromised and published online; including places visited, camera images and policyholder names. The data held by the devices is deleted or inaccessible and ongoing driver usage is not captured, resulting in 3 to 6 months' driving data being unavailable. This data would normally be used by the insurer to determine the risk charges / premiums for the insurance product. (Note that an alternative adverse scenario could have involved the manipulation of data to make it unreliable on a policy by policy basis. This type of exercise could have continued for many months or years before detection.)

Compromised devices are used as part of a Botnet to launch a distributed denial of service attack. Such an attack would result in the attackers having control of the devices and being able to hire out the devices for others to perform attacks or doing them themselves. No costs are assumed, since at present litigation has not been directed towards those whose networks have been taken over by attackers. However, this is still mentioned as part of this in the scenario, as it is plausible that litigation to recover costs for the cybersecurity negligence of organisations whose networks are used for distributed denial of service ("DDoS") attacks could result in additional costs in the future.

The attack published by the researcher highlights the fact that a web service is enabled by default on the telemetry device. The administrative interface to this web service is accessible using a default username and password combination (Admin/Admin). When logged into the web service with administrative credentials, the user can visit a page on the web site which provides the location of the device, a recent history of previous locations, the home address of the driver, driver's license and a live feed of images coming from the camera. The web server also allows the administrative user to remotely wipe the device and upload new device management software on it for upgrade/support purposes. In addition, the device has an old version of Apache web server software which is susceptible to a buffer overflow attack leading to unauthorised remote access to the device.

A few weeks after the researcher's results were published, a malicious botnet was created that automatically exploited the vulnerability and replaced the software on the devices with an image that ran DDoS attacks as part of a DDoS botnet.

Timelines

Week 0: Hack occurs

Week 3: A problem is detected in the devices. Investigation of the cause of the issue is identified; no information is coming out of the devices due to the hack. To rectify, the insurer needs to replace the product or fix it "over the air".

Week 5: After investigation, the insurer finally realises that the problem is caused by a hack on the devices which need to be replaced. (Fixing over the air would typically reduce the costs of the scenario, and thus for the sake of a remote scenario this is not considered possible.) At the same time, data from the devices is being published online.

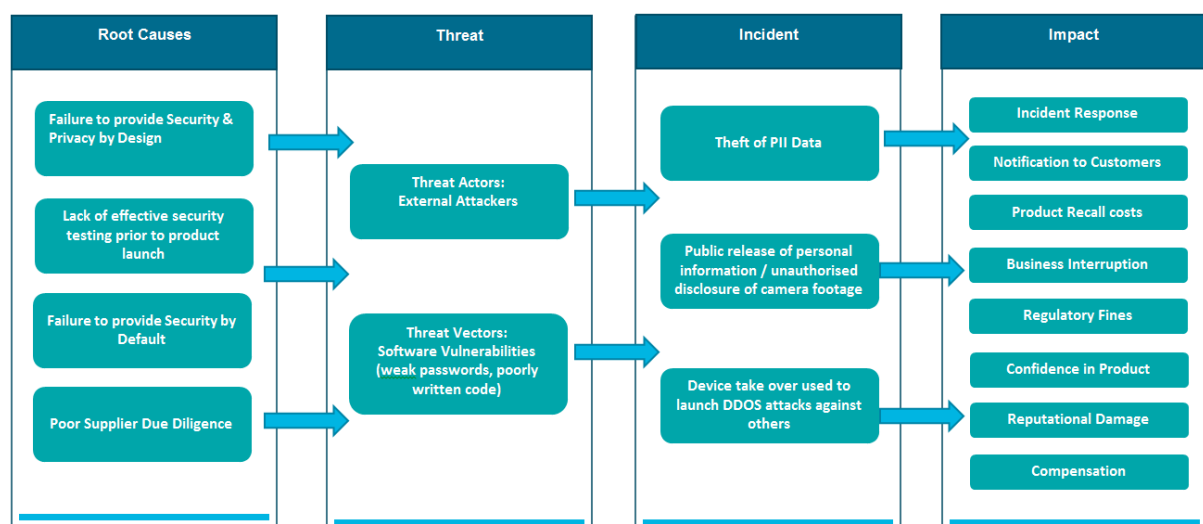
Week 10 - 20: To replace devices, the insurer needs to produce new devices and ship them to UK.

End of year 1: The Information Commissioner's Office ("ICO") applies a fine due to loss of customer data resulting from device security weaknesses.

Years 3 – 5: Damages incurred from complaints cases, reputational damage remains (uptake in new insurance products integrated with telemetry devices is slower compared with competition) and sales are reduced.

Year 5: Incident now in past and reputation restored

Figure 12. Incident summary for telematics device hack scenario



Examples of Internet of Things (“IoT”) devices used by insurers

IoT products measuring behaviours and driving down premiums are exposed to this type of hack. There are a growing number of IoT devices being used by insurers for the insurance products. Some examples are shown below for different insurer types.

Healthcare:

- fitness measurement devices; and
- monitoring devices such as heart monitors.

Home insurance:

- gas meters / electric meters to insurer to reduce premium;
- smart smoke/ heat alarm; and
- smart water detection.

Ship / cargo insurance:

- telematics / GPS keeping track of ships / cargo / shipments .

Car insurance:

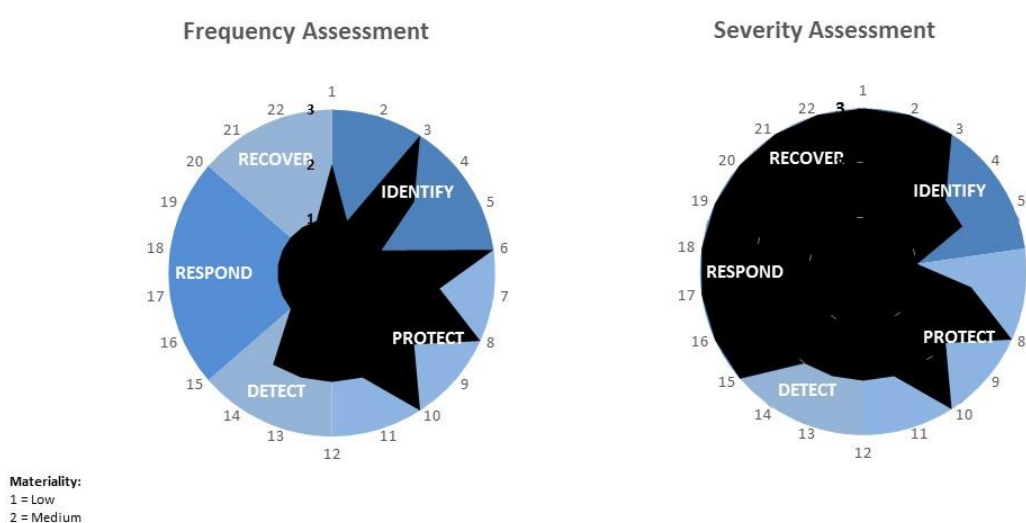
- devices used in cars to measure driving habits/behaviours and encouraging good behaviour premium.

3.6.4 Security assessment and mitigation

The following charts display the assessment of this scenarios vulnerability across the NIST framework for the impact on frequency and severity of the event and indicate that the following control areas are expected to be the key vulnerabilities for this scenario:

- identify e.g. asset management and inventory;
- protect e.g. access controls, data security, remote management and information protection processes; and
- detect e.g. anomalies and events.

Figure 13. NIST framework assessment for telematics device hack scenario



3.6.5 Expected Costs

The table below summarises some of the expected costs for this type of scenario. These costs are only indications of the potential magnitude of each cost area for the specified parameters of this scenario.

Table 14. Expected cost summary for telematics device hack scenario

	Cost Type	Scenario cost	Approximate Cost (gross)	Rationale
1	Incident response costs	External consultants used to investigate data breach.	£0.5m	This is expert judgement given the uncertainty of the scenario. This is expected to be a concentrated effort for 2 weeks - at £20k a day (Big 4 consultancy team of 5 people with senior support being significant) for 12 days, this is £240k. This is then followed by further support averaging £50k per week in weeks 5 to 10 to attempt to obtain the data and also to ensure that the new devices have independent eyes on their security.
2	Physical damage	Physical Device - product replacement, labour costs to install new devices and customer outreach programme costs.	£42.5m	(£50 device cost + £25 installation cost + £10 customer outreach cost) x 500k devices. Above is expert judgement based on scenario as there are no direct precedents. The outreach cost is greater than the costs in other scenarios to coordinate customers to having their devices placed in centralised centres eg supermarket. It would include an incentive eg a £5 gift voucher to spend whilst having the device replaced.
3	Business interruption	Premium income – loss of future premium income.	£14.0m	Give 25% credit to historical data to all customers for lost data (i.e. assume all had metrics resulting in 25% lower metrics for 3 to 6 months, resulting in 15% lower premium) - 15% premium credit * Ave(3,6) months / 12 months x £250m annual premium Note it may not seem intuitive as to why a 25% credit to driving history does not result in a 25% reduction in insurance costs. Telematics insurance is based on car usage, driving habits and other

	Cost Type	Scenario cost	Approximate Cost (gross)	Rationale
				policy details but there are a number of fixed expenses and even a car that is not driven is exposed to an insurable loss.
4	Regulatory Fines	Fine for loss of customer exposure data; assumed failure to comply with GDPR rules.	£2.0m	<p>If results published on-line including personal details (e.g. home address) and driving habits. Data privacy fine from FCA / ICO £400m revenue x 4% x 10%.</p> <p>Note the 10% could be as high as 100%. This is not higher because the exploit was only exposed weeks before the attack. However, it is not nil because tighter controls could have been in place. The newness of the GDPR regime makes this figure very uncertain.</p>
5	Regulatory costs	S166 into how breach occurred and validity of actions taken to remediate weaknesses and avoid future occurrences.	£1.0m	The costs of S166's have ranged from £30k to £1.3m in 2017. Given the nature we presume this would be at the higher end.
6	Compensation	Ex-gratia Payments: complaints due to sensitive data disclosure (home address, trips, etc.) leading to customer losses and ex-gratia offers to compensate customers.	£10.0m	Assuming 1% complain with an average award of £2k each to 5,000 customers (all 1% that complain) skewed to lower end with a few high value offers.
		Total	£70.0m	

The majority of the costs estimated for this scenario are caused by the product replacement cost for all the cars. The scenario overall results in a cost of c18% of annual premium. It is possible that some portion of the scenario costs could be recovered e.g. from the manufacturer of the devices or a separate insurance policy, however this has not been assumed for this scenario.

Business interruption costs and reputational damage have not been considered relevant for this scenario. There may need to be some system downtime for investigative work but it is not considered that it would be significant and thus normal operations would not be greatly impacted. Also, the type of consumer buying these policies is likely to be saving money by using such a device. This will require consumers to either switch away from such a device or switch provider; it is not clear whether consumers would believe that switching away would solve the issue.

3.6.6 Mitigation

The impact and ability to mitigate the risk is dependent on the following key areas (as labelled in the NIST framework):

- identify;
- protect;
- detect; and
- respond.

The devices need to have better security and may require some security upgrades (software and hardware) to reduce their vulnerability to a hack. In addition, the devices should be monitored for unauthorised access, and regular security testing put in place to ensure they are safe.

Table 15. Proposed mitigation approach for

NIST function	Mitigation type	Examples	Mitigating benefit
Identify	Asset management	Maintaining an asset inventory of devices that have been deployed to customers.	Keeping track of assets in the field and having the ability to control / remotely manage these devices if required.
	Risk assessment	Carrying out a risk assessment prior to the acquisition and deployment of the devices to identify potential risks and exposures and put in place mitigating actions to reduce risks of device deployment. In addition, the threat environment should be considered on an ongoing basis in order to put relevant procedures in place to protect against it.	Security risks could have been anticipated ahead of the incident that occurred and additional controls may have been considered including better passwords, encryption of PII data and firmware device integrity checking. This will ensure that ongoing procedures are in place to avoid threats that need ongoing attention. Some are more routine such as patching software vulnerabilities, others may develop over time, examples being the assessment of new types of cyber threats.
	Risk management strategy	Assessing project risks such as the IoT deployment project and also risks of third party suppliers such as the ones who provided the devices to the insurer.	Early identification of security risks can help companies implement controls on new projects (security by design) and also identify red flags with suppliers providing software / hardware to the client which may have security holes within them.
Protect	Access control	User and administrative accounts are well managed from creation through use and deletion.	Strict control over user access accounts to devices can significantly reduce risk of unauthorised access to devices including password policies, removal of default accounts and passwords.
	Data security	Data at rest adequately protected. Integrity checking in place on firmware.	Use of encryption and access control over sensitive data stored on devices could have reduced the risk of this incident escalating the way it did. An ability to check the integrity of firmware running on a device would make it harder for hackers to install new versions of software that enabled them to launch DDOS attacks.

NIST function	Mitigation type	Examples	Mitigating benefit
	Information protection processes and procedures	<p>Security baseline configuration created and maintained.</p> <p>A systems development life cycle ("SDLC") is implemented and managed which includes security design within it.</p> <p>A vulnerability management programme for security testing and remediation is in place to detect and mitigate vulnerabilities identified.</p>	<p>Establishing a strong security baseline including changes to default passwords and stronger enforcement of access controls to PII Data would have assisted.</p> <p>Ensuring that security has been embedded in the full SDLC of the device software would have identified security risks and vulnerabilities prior to the devices being sold and deployed to the insurer.</p> <p>A vulnerability management programme throughout the SDLC and in production environments would assist in catching security vulnerabilities before external attackers do.</p>
	Maintenance	<p>Regular maintenance in place on device health and management.</p> <p>Remote maintenance is performed in a manner that minimises unauthorised access.</p>	<p>Providing a facility for remote device management and health checks ensures that the integrity of devices remains intact.</p> <p>Ensuring remote management is securely implemented helps achieve the first goal above without compromising the security of the device being managed. Without effective implementation of remote management, it becomes another attack vector to target.</p>
	Protective technology	<p>Penetration testing to understand how devices can be exploited and what can be achieved with exploits. Bug bounty programmes achieve a similar goal.</p>	<p>Better understanding of the potential for damage resulting from vulnerabilities in devices.</p>
Detect	Anomalies and events	<p>Baseline of events established to analyse events and detect unauthorised access.</p> <p>Malicious code detected.</p>	<p>Monitoring can help detect unusual activities on devices and identify anomalies quicker to reduce the impact of attacks should they occur.</p> <p>Monitoring device behaviour can be used to detect malicious code and activities should the device integrity be compromised.</p>
Respond	Security intelligence gathering	<p>Identifying security threats through open source information and responding to them before they escalate.</p>	<p>Proactive identification of security research activity may have helped to detect at an early stage that a threat (vulnerability within the device) was moving from theory to practical as the security researcher published their results and vulnerabilities which others exploited.</p>

Acknowledgements

The authors would like to thank Jonathan Evans, Patrick Kelliher, and Edward Pocock for their invaluable feedback as well as the IFoA staff for their continued support and assistance.



References

- BBC. "RBS fined £56m over 'unacceptable' computer failure." *BBC News*. 2014.
<https://www.bbc.com/news/business-30125728> (accessed September 15, 2018).
- Boros, D. "On Lapse risk factors in Solvency II." *KTH*. 2014.
<https://www.math.kth.se/matstat/seminarier/reports/M-exjobb14/140611a.pdf> (accessed September 15, 2018).
- Clayton, J. *Statement on Cybersecurity*. Statement, U.S. Securities and Exchange Commission, Washington, DC: U.S. Securities and Exchange Commission, 2017.
- CRO Forum. *CRO Forum Concept Paper on a proposed categorisation methodology for cyber risk*. Concept Paper, Amsterdam: CRO Forum, 2016, 28.
- eRiskHub. *NetDiligence Mini Data Breach Cost Calculator*. n.d. <https://eriskhub.com/mini-dbcc> (accessed August 2018).
- Gould + Partners. *PR Agency Industry 2014 Billing Rates & Utilization Report*. New York: Gould + Partners, 2014.
- Hubmann, C., Polke-Markmann, H., and Vanheyden, P. *Allianz Risk Barometer - Top Business Risks for 2018*. Allianz, Munich: Allianz Global Corporate & Specialty (AGCS), 2018, 20.
- IT Security Expert. *PCI DSS Penalties & Fines? Cyber Insurance? How to Estimate the Cost of a Payment Card Breach*. March 9, 2017. <https://blog.itsecurityexpert.co.uk/2017/03/pci-dss-fines-cyber-insurance-how-to.html> (accessed September 2018).
- National Institute of Standards and Technology. *Cybersecurity "Rosetta Stone" Celebrates Two Years of Success*. News, Information Technology Laboratory, Gaithersburg, Md: National Institute of Standards and Technology, 2016.
- National Institute of Standards and Technology. *Framework for Improving Critical Infrastructure Cybersecurity*. U.S. Department of Commerce, Gaithersburg, Md. and Boulder, Colo.: National Institute of Standards and Technology, 2018, 55.
- National Institute of Standards and Technology. *Framework for Improving Critical Infrastructure Cybersecurity*. Gaithersburg, Md. and Boulder, Colo.: National Institute of Standards and Technology, 2014.

Novet, J. "Shipping company Maersk says June cyberattack could cost it up to \$300 million." *CNBC*. 2017. <https://www.cnn.com/2017/08/16/maersk-says-notpetya-cyberattack-could-cost-300-million.html> (accessed September 15, 2018).

Osborn, T., Campbell, A., Marlin, S., Isa, A., and Woodall, L. *Top 10 operational risks for 2018*. Report, London: Risk.net, 2018.

Paatz, M. "Morrisons data breach sounds warning on vicarious liability." *Personnel Today*. 2018. <https://www.personneltoday.com/hr/morrisons-data-breach-sounds-warning-on-vicarious-liability/> (accessed September 15, 2018).

Ponemon Institute. *Cost of Data Center Outages*. North Traverse City: Ponemon Institute, 2016.

Wikipedia. *Anthem medical data breach*. Talk, Wikipedia Foundation, Inc, Wikipedia: The Free Encyclopedia, 2015.

Appendix 1 – Scenario Selection

The following 7 scenarios were discussed by the working party. These scenarios were originally conceived through brainstorming based on known events and events considered to be plausible given the knowledge of the cyber threat environment at the time. Care was taken to consider scenarios relevant to insurance organisations and across the whole industry regardless of area of business focus. The final selection of scenarios to focus on was based on a group vote to determine the scenarios which the group considered the most relevant and interesting to explore in greater detail.

Scenario 1: A general insurance business with a diverse business including a large motor portfolio is hacked by an internal staff member. Details of all motor insurance policyholders are leaked onto an internet website and are widely available.

Scenario 2: A large life insurance business is targeted by a spear phishing email to their CFO, apparently from their CEO. This results in a large transfer of funds intended for an investment portfolio, into a rogue bank account.

Scenario 3: A Lloyd's syndicate has a large portfolio of risks in the USA. The internet in the East Coast of the United States is attacked by cyber anarchists, resulting in no internet connectivity for 2 weeks.

Scenario 4: A large insurer is in the process of migrating its data centre operations to the cloud. A member of their IT team extracts a large volume of data containing Personally Identifiable Information client data onto a high capacity disc to transfer to the new data centre. During the physical transfer of this disc, the disc gets stolen.

Scenario 5: A broker for a general insurer gets infected with ransomware on their computer. The ransomware spreads within the company and encrypts a major file share containing client records. The company is unable to access these records as they are encrypted by the malware. The online backup of the file share is also affected by the malware as it automatically backed up encrypted files. The insurer experiences an inability to process client requests due to lack of availability of important client information.

Scenario 6: An insurer employs a third party to print and send invoices and statements to all their customers. Large volumes of client data are shared monthly with the service provider to carry out necessary print and invoice operations. The insurer gets notified by the third party that they have experienced a data breach and customer records have been stolen.

Scenario 7: A motor insurer deploys telemetry in customer vehicles for measuring driver patterns using a specific telemetry device. A security researcher publicises a hack on this device that allows any internet user to access the camera of the telemetry device as well as the location and PII data on it. The insurer needs to recall / replace / replenish the device with each of its clients.

Scenarios 1, 5 and 7 were selected as being the most relevant to the insurance industry from an operational risk perspective and the following amendments were suggested.

Scenario 1: Ensure that the data breach focus is retained but expand the narrative of the scenario to include both personal lines (volume focus) and commercial lines/London market (sensitivity focus e.g. high net worth, K&R, M&A).

Scenario 5: The focus of the scenario should be on business interruption e.g. ransomware/cloud downtime.

Scenario 7: In researching the scenario consider IoT and the potential impact of this area of technology more broadly.

Appendix 2 – Detailed NIST Framework

The following table sets out the 5 core functions proposed within the NIST framework to ensure a company responds to cyber risk. We have assessed each scenario against the 22 control categories within each of these core functions as set out in v1.0 of the NIST framework paper (National Institute of Standards and Technology 2014).

Function	ID	Control category
IDENTIFY (ID)	1	Asset Management (ID.AM): The data, personnel, devices, systems, and facilities that enable the organization to achieve business purposes are identified and managed consistent with their relative importance to business objectives and the organization's risk strategy.
	2	Business Environment (ID.BE): The organization's mission, objectives, stakeholders, and activities are understood and prioritized; this information is used to inform cybersecurity roles, responsibilities, and risk management decisions.
	3	Governance (ID.GV): The policies, procedures, and processes to manage and monitor the organization's regulatory, legal, risk, environmental, and operational requirements are understood and inform the management of cybersecurity risk.
	4	Risk Assessment (ID.RA): The organization understands the cybersecurity risk to organizational operations (including mission, functions, image, or reputation), organizational assets, and individuals.
	5	Risk Management Strategy (ID.RM): The organization's priorities, constraints, risk tolerances, and assumptions are established and used to support operational risk decisions.
PROTECT (PR)	6	Access Control (PR.AC): Access to assets and associated facilities is limited to authorized users, processes, or devices, and to authorized activities and transactions.
	7	Awareness and Training (PR.AT): The organization's personnel and partners are provided cybersecurity awareness education and are trained to perform their cyber security related duties and responsibilities consistent with related policies, procedures, and agreements.
	8	Data Security (PR.DS): Information and records (data) are managed consistent with the organization's risk strategy to protect the confidentiality, integrity, and availability of information.
	9	Information Protection Processes and Procedures (PR.IP): Security policies (that address purpose, scope, roles, responsibilities, management commitment, and coordination among organizational entities), processes, and procedures are maintained and used to manage protection of information systems and assets.
	10	Maintenance (PR.MA): Maintenance and repairs of industrial control and information system components is performed consistent with policies and procedures.
	11	Protective Technology (PR.PT): Technical security solutions are managed to ensure the security and resilience of systems and assets, consistent with related policies, procedures, and agreements.
DETECT (DE)	12	Anomalies and Events (DE.AE): Anomalous activity is detected in a timely manner and the potential impact of events is understood.
	13	Security Continuous Monitoring (DE.CM): The information system and assets are monitored at discrete intervals to identify cybersecurity events and verify the effectiveness of protective measures.
	14	Detection Processes (DE.DP): Detection processes and procedures are maintained and tested to ensure timely and adequate awareness of anomalous events.
RESPOND (RS)	15	Response Planning (RS.RP): Response processes and procedures are executed and maintained, to ensure timely response to detected cybersecurity events.

	16	Communications (RS.CO): Response activities are coordinated with internal and external stakeholders, as appropriate, to include external support from law enforcement agencies.
	17	Analysis (RS.AN): Analysis is conducted to ensure adequate response and support recovery activities.
	18	Mitigation (RS.MI): Analysis is conducted to ensure adequate response and support recovery activities.
	19	Improvements (RS.IM): Organizational response activities are improved by incorporating lessons learned from current and previous detection/response activities.
RECOVER (RC)	20	Recovery Planning (RC.RP): Recovery processes and procedures are executed and maintained to ensure timely restoration of systems or assets affected by cybersecurity events.
	21	Improvements (RC.IM): Recovery planning and processes are improved by incorporating lessons learned into future activities.
	22	Communications (RC.CO): Restoration activities are coordinated with internal and external parties, such as coordinating centres, Internet Service Providers, owners of attacking systems, victims, other CSIRTs, and vendors.

It is worth noting that an additional control category was added to the 'Identify' function in v1.1 of the NIST framework (National Institute of Standards and Technology 2018). As mentioned in section 2.1 of this paper this is not deemed to have a material impact on the conclusions of the paper. For completeness, the additional control category has been included below:

Function	ID	Control category
IDENTIFY (ID)	-	Supply Chain Risk Management (ID.SC): The organisation's priorities, constraints, risk tolerances, and assumptions are established and used to support risk decisions associated with managing supply chain risk. The organization has established and implemented the processes to identify, assess and manage supply chain risks.

Appendix 3 – Glossary of terms

Attacker: Malicious actor who seeks to exploit computer systems with the intent to change, destroy, steal or disable their information, and then exploit the outcome.

Botnet: A botnet is a collection of internet-connected devices, which may include PCs, servers, mobile devices and internet of things devices that are infected and controlled by a common type of malware. Users are often unaware of a botnet infecting their system.

Breach: An incident in which data, computer systems or networks are accessed or affected in a non-authorized way.

Brute force attack: Using computational power to automatically enter myriad value combinations, usually in order to discover passwords and gain access.

Bug bounty programmes: A bug bounty program is a deal offered by many websites and software developers by which individuals can receive recognition and compensation for reporting bugs, especially those pertaining to exploits and vulnerabilities.

CISO: A chief information security officer (CISO) is the senior-level executive within an organisation responsible for establishing and maintaining the enterprise vision, strategy, and program to ensure information assets and technologies are adequately protected.

CRO Forum: The CRO Forum is a group of professional risk managers from the insurance industry that focuses on developing and promoting industry best practices in risk management. The Forum consists of Chief Risk Officers from large multi-national insurance companies. It aims to represent the members' views on key risk management topics, including emerging risks.

Cyber resilience: Cyber resilience refers to an entity's ability to continuously deliver the intended outcome despite adverse cyber events.

Cyber underwriting risk: Cyber underwriting risk is defined as the set of risks emanating from underwriting insurance contracts that are exposed to losses resulting from a cyber-attack.

Data at rest: Describes data in persistent storage such as hard disks, removable media or backups.

Data warehousing: Data warehousing is a technology that aggregates structured data from one or more sources so that it can be compared and analysed for greater business intelligence.

DDoS: A distributed denial-of-service (DDoS) attack is an attack in which multiple compromised computer systems attack a target, such as a server, website or other network resource, and cause a denial of service for users of the targeted resource. The flood of incoming messages, connection requests or malformed packets to the target system forces it to slow down or even crash and shut down, thereby denying service to legitimate users or systems.

Device hack: Embedded device hacking is the exploiting of vulnerabilities in embedded software to gain control of the device. Attackers have hacked embedded systems to spy on the devices, to take control of them or simply to disable them. Embedded systems exist in a wide variety of devices including Internet and wireless access points, IP cameras, security systems, pace makers, drones and industrial control systems.

ERM: Enterprise risk management (ERM) is the process of planning, organizing, leading, and controlling the activities of an organization in order to minimize the effects of risk on an organization's capital and earnings.

Firmware: In electronic systems and computing, firmware is a specific class of computer software that provides the low-level control for the device's specific hardware. Firmware can either provide a standardized operating environment for the device's more complex software(allowing more hardware-independence), or, for less complex devices, act as the device's complete operating system, performing all control, monitoring and data manipulation functions.

GDPR: The General Data Protection Regulation (GDPR) is a legal framework that sets guidelines for the collection and processing of personal information of individuals within the European Union. The GDPR sets out the principles for data management and the rights of the individual, while also imposing fines that can be revenue-based. The General Data Protection Regulation covers all companies that deal with data of EU citizens, so it is a critical regulation for corporate compliance officers at banks, insurers, and other financial companies. GDPR came into effect across the EU on May 25, 2018.

IoT: Internet of Things (IoT) is the network of physical devices, vehicles, home appliances, and other items embedded with electronics, software, sensors, actuators, and connectivity which enables these things to connect and exchange data, creating opportunities for more direct integration of the physical world into computer-based systems, resulting in efficiency improvements, economic benefits, and reduced human exertions.

Malware: Malware, is defined as the malicious software file or program harmful to a computer user which can execute different malicious functions like encrypting, stealing or deleting sensitive data, hijacking or altering core computing functions and monitoring computer activities of users without their permission.

Network segmentation: Network segmentation in computer networking is the act or practice of splitting a computer network into subnetworks, each being a network segment. Advantages of such splitting are primarily for boosting performance and improving security.

NIST Framework: The NIST Cybersecurity Framework provides a policy framework of computer security guidance for how private sector organizations in the United States can assess and improve their ability to prevent, detect, and respond to cyber-attacks.

Operational Risk: Operational Risk is the risk of loss resulting from inadequate or failed internal processes, people and systems, or from external events. Operational Risk is the residual risk not covered by other categories of risk, including insurance, financial, credit and liquidity risk.

Patch controls: Patch management is an area of systems management that involves acquiring, testing, and installing multiple patches (code changes) to an administered computer system. Patch management tasks include: maintaining current knowledge of available patches, deciding what patches are appropriate for particular systems, ensuring that patches are installed properly, testing systems after installation, and documenting all associated procedures, such as specific configurations required.

Petya / Notpetya: Petya is a family of encrypting ransomware that was first discovered in 2016. The malware targets Microsoft Windows-based systems, infecting the master boot record to execute a payload that encrypts a hard drive's file system table and prevents Windows from booting. It subsequently demands that the user make a payment in Bitcoin in order to regain access to the system. Variants of Petya were first seen in March 2016, which propagated via infected e-mail attachments. In June 2017, a new variant of Petya was used for a global cyberattack, primarily targeting Ukraine. The new variant propagates via the EternalBlue exploit, which is generally believed to have been developed by the U.S. National Security Agency (NSA), and was used earlier in the year by the WannaCry ransomware. Kaspersky Lab referred to this new version as NotPetya to distinguish it from the 2016 variants, due to these differences in operation. In addition, although it purports to be ransomware, this variant was modified so that it is unable to actually revert its own changes.

Penetration test / Pentest: An authorised test of a computer network or system designed to look for security weaknesses so that they can be fixed.

PFI: PCI Forensic Investigators (PFIs) help determine the occurrence of a cardholder data compromise and when and how it may have occurred. These PCI Forensic Investigators are qualified by the Council's program and must work for a Qualified Security Assessor company that provides a dedicated forensic investigation practice. They perform investigations within the financial industry using proven investigative methodologies and tools. They also provide relationships with law enforcement to support stakeholders with any resulting criminal investigations.

PII: Personally identifiable information (PII) is any data that could potentially identify a specific individual. Any information that can be used to distinguish one person from another and can be used for de-anonymizing anonymous data can be considered PII.

QSA: Qualified Security Assessor is a designation conferred by the PCI Security Standards Council to those individuals that meet specific information security education requirements, have taken the appropriate training from the PCI Security Standards Council, are employees of a Qualified Security Assessor (QSA) company approved PCI security and auditing firm, and will be performing PCI compliance assessments as they relate to the protection of credit card data.

Ransomware attack: Ransomware is a type of malicious software from cryptovirology that threatens to publish the victim's data or perpetually block access to it unless a ransom is paid. While some simple ransomware may lock the system in a way which is not difficult for a knowledgeable person to reverse, more advanced malware uses a technique called cryptoviral extortion, in which it encrypts the victim's files, making them inaccessible, and demands a ransom payment to decrypt them.

S166: A s166 notice is a notice issued by the Financial Conduct Authority (FCA) under s166 of the Financial Services and Markets Act 2000 requiring a firm to carry out a "skilled person review". The FCA serves around 50 a year.

SDLC: Software Development Life Cycle (SDLC) is a process used by the software industry to design, develop and test high quality softwares. It is also called a Software Development Process. SDLC is a framework defining tasks performed at each step in the software development process.

Social engineering: Social engineering, in the context of information security, refers to psychological manipulation of people into performing actions or divulging confidential information. A type of confidence trick for the purpose of information gathering, fraud, or system access, it differs from a traditional "con" in that it is often one of many steps in a more complex fraud scheme.

Software vulnerabilities: In computer security, a vulnerability is a weakness which can be exploited by a Threat Actor, such as an attacker, to perform unauthorized actions within a computer system. To exploit a vulnerability, an attacker must have at least one applicable tool or technique that can connect to a system weakness. In this frame, vulnerability is also known as the attack surface.

Spear-phishing: Spear phishing is an email-spoofing attack that targets a specific organization or individual, seeking unauthorized access to sensitive information. Spear-phishing attempts are not typically initiated by random hackers, but are more likely to be conducted by perpetrators out for financial gain, trade secrets or military information.

Telemetry: Telemetry is an automated communications process by which measurements and other data are collected at remote or inaccessible points and transmitted to receiving equipment for monitoring.

Vulnerability scans: Vulnerability scanning is an inspection of the potential points of exploit on a computer or network to identify security holes. A vulnerability scan detects and classifies system weaknesses in computers, networks and communications equipment and predicts the effectiveness of

countermeasures. A scan may be performed by an organization's IT department or a security service provide, possibly as a condition imposed by some authority.

Worm: A worm is a standalone malware computer program that replicates itself in order to spread to other computers. Often, it uses a computer network to spread itself, relying on security failures on the target computer to access it.

London

7th Floor · Holborn Gate · 326-330 High Holborn · London · WC1V 7PP
Tel: +44 (0) 20 7632 2100 · Fax: +44 (0) 20 7632 2111

Edinburgh

Level 2 · Exchange Crescent · 7 Conference Square · Edinburgh · EH3 8RA
Tel: +44 (0) 131 240 1300 · Fax +44 (0) 131 240 1311

Oxford

1st Floor · Park Central · 40/41 Park End Street · Oxford · OX1 1JD
Tel: +44 (0) 1865 268 200 · Fax: +44 (0) 1865 268 211

Beijing

6/F · Tower 2 · Prosper Centre · 5 Guanghai Road · Chaoyang District · Beijing · China 1000020
Tel: +86 (10) 8573 1000

Hong Kong

2202 Tower Two · Lippo Centre · 89 Queensway · Hong Kong
Tel: +11 (0) 852 2147 9418 · Fax: +11 (0) 852 2147 2497

Singapore

163 Tras Street · #07-05 Lian Huat Building · Singapore · 079024
Tel: +65 6717 2955

www.actuaries.org.uk

© 2016 Institute and Faculty of Actuaries