Counting the cost
Cyber exposure decoded

CYENCE

Lloyd's of London Disclaimer

This report has been co-produced by Lloyd's and Cyence for general information purposes only. While care has been taken in gathering the data and preparing the report, Lloyd's does not make any representations or warranties as to its accuracy or completeness and expressly excludes to the maximum extent permitted by law all those that might otherwise be implied.

Lloyd's accepts no responsibility or liability for any loss or damage of any nature occasioned to any person as a result of acting or refraining from acting as a result of, or in reliance on, any statement, fact, figure or expression of opinion or belief contained in this report. This report does not constitute advice of any kind.

About Lloyd's

Lloyd's is the world's specialist insurance and reinsurance market. Under our globally trusted name, we act as the market's custodian. Backed by diverse global capital and excellent financial ratings, Lloyd's works with a global network to grow the insured world – building the resilience of local communities and strengthening global economic growth.

With expertise earned over centuries, Lloyd's is the foundation of the insurance industry and the future of it. Led by expert underwriters and brokers who cover more than 200 territories, the Lloyd's market develops the essential, complex and critical insurance needed to underwrite human progress.

About Cyence

Cyence empowers the insurance industry to understand the impact of cyber risk in the context of dollars and probabilities. It's unique approach combines economic/risk modeling, cybersecurity and big data analytics to create an economic cyber risk modeling platform. Cyence Platform and analytics are leveraged by leaders across the insurance industry to help understand and manage cyber risk as well as to roll out new transformative insurance products.

Key Contacts

Trevor Maynard
Head of Innovation
trevor.maynard@lloyds.com

For general enquiries about this report and Lloyd's work on innovation, please contact
innovation@lloyds.com

About the authors

Trevor Maynard PhD, MSc, FIA has degrees in pure maths and statistics and is a Fellow of the Institute of Actuaries. He is Head of Innovation at Lloyd's including responsibility for horizon scanning and emerging risks. Subjects covered in recent years include: the economic and social implications of a food system shock; the effects of cyber-attacks on the US energy grid and an exploration of aggregation modelling methods for liability risks.

He is co-chairman of OASIS, an open modelling platform for catastrophe models and sits on the Board of the Lighthill Risk Network.

George Ng, a founder and Chief Technology Officer, leads major research projects and initiatives at Cyence. Previously, he was the Chief Data Scientist at YarcData. George has also worked as a Research Scientist at DARPA and US-CERT and as faculty at American University. He received his PhD from UC Irvine and B.A. from UC Berkeley, both in Economics.

Acknowledgements

Contents

Executive summary

The aim of this report is to provide insurers who write cyber coverage with realistic and plausible scenarios to help quantify cyber-risk aggregation. The understanding of cyber liability and risk exposures is relatively underdeveloped compared with other insurance classes.

By understanding cyber risk exposure, insurers can improve their portfolio exposure management, set appropriate limits and gain the confidence to expand into this fast-growing insurance class.

The report is designed for risk managers whose businesses are exposed to the types of cyber-attacks described in the report's two scenarios: a hack that takes down their cloud-service provider or an attack that causes the failure of a particular operating system across their own company, customers, suppliers and/or business partners.

Each of these scenarios encompasses a range of variables including possible risk mitigation and cyber-attack response. This means organisations can consider the impact on their own operations.

Methodology

This report was developed collaboratively by Lloyd's and Cyence, who brought together a multidisciplinary team of experts in cyber security, economic risk modelling and cyber insurance.

Cyence undertook a structured, seven-stage research process to generate the scenarios and produce the loss estimates in this report. The seven stages were:

1. Review of widely adopted technologies used across industries
2. Review of other non-technical factors
3. Data collection and processing for the exposures
4. Analysis of the exposure accumulation paths
5. Selection of scenarios, frequency and severity models
6. Discussion and review with insurance and cyber security experts
7. Loss calculations and final review

Lloyd's worked with the Lloyd's Market Association on a series of collaborative workshops involving cyber underwriters from the Lloyd's market to discuss and include feedback in the report, and identify the implications and considerations for the insurance industry.

Cyber-attack – an increasing threat

Cyber risk is a growing global threat. While digitisation is revolutionising business models and transforming daily lives, it is also making the global economy more vulnerable to cyber-attacks.

As a result, the economic and insurance consequences of cyber-crime are increasing. In 2016, cyber-attacks were estimated to cost businesses as much as $450 billion a year globally *(Graham, 2017)*. Increasingly, insurers are helping policyholders manage these events; everything from individual breaches caused by malicious insiders and hackers, to wider losses such as breaches of retail point-of-sale devices, ransomware attacks such as BitLocker, WannaCry and distributed denial-of-service attacks such as Mirai.

The cyber threat is increasing and is expected to continue to do so as the world economy continues to digitise operations, supply chains and businesses transactions, as well as employee and customer services.

Challenges for insurers

As the cyber threat grows so the demand for cyber insurance increases. Today, Lloyd's Class of Business team estimates that global cyber market is worth between $3 billion and $3.5 billion *(Stanley, 2017)*; by 2020, some analysts estimate it could be worth $7.5 billion *(PwC, 2015)*. Property/casualty insurers wrote $1.35 billion in direct written premium for cyber insurance in 2016, a 35% jump from 2015, according to reports by Fitch Ratings and A.M. Best *(A.M Best, 2016)*.

Despite this growth, insurers' understanding of cyber liability and risk aggregation is an evolving process as

experience and knowledge of cyber-attacks grows. Insureds' use of the internet is also changing, causing cyber-risk accumulation to change rapidly over time in a way that other perils do not.

Traditional insurance risk modelling relies on authoritative information sources such as national or industry data, but there are no equivalent sources for cyber-risk and the data for modelling accumulations must be collected at scale from the internet. This makes data collection, and the regular update of it, key components of building a better understanding of the evolving risk.

How the report can deepen understanding of cyber-risk aggregation

This report is designed to increase insurers' and risk managers' understanding of cyber-risk liability and aggregation. It analyses aggregation through the prism of six trends that contribute to digital vulnerability. Understanding these trends is crucial to understanding cyber aggregation.

These trends are:

1. Volume of contributors: The number of people developing software has grown significantly over the past three decades; each contributor could potentially add vulnerability to the system unintentionally through human error.
2. Volume of software: In addition to the growing number of people amending code, the amount of it in existence is increasing. More code means the potential for more errors and therefore greater vulnerability.
3. Open source software: The open-source movement has led to many innovative initiatives. However, many open-source libraries are uploaded online and while it is often assumed they have been reviewed in terms of their functionality and security, this is not always the case. Any errors in the primary code could then be copied unwittingly into subsequent iterations.
4. Old software: The longer software is out in the market, the more time malicious actors have to find and exploit vulnerabilities. Many individuals and companies run obsolete software that has more secure alternatives.
5. Multi-layered software: New software is typically built on top of prior software code. This makes software testing and correction very difficult and resource intensive.
6. "Generated" software: Code can be produced through automated processes that can be modified for malicious intent.

The report also uses scenarios to quantify the wide variety of damages that can occur as a result of two different cyber events.

Scenario 1: Cloud service provider hack

A sophisticated group of "hacktivists" sets out to disrupt cloud-service providers and their customers to draw attention to the environmental impacts of business and the modern economy. The group makes a malicious modification to a "hypervisor" that controls the cloud infrastructure. This causes many cloud-based customer servers to fail, leading to widespread service and business interruption.

Scenario 2: Mass vulnerability attack

A cyber analyst accidentally leaves his bag on a train that contains a hard copy of a report on a vulnerability that affects all versions of an operating system run by 45% of the global market. This report is traded on the dark web and is purchased by an undetermined number of unidentified criminal parties who develop system exploits[a] and begin attacking vulnerable businesses for financial gain.

[a] An exploit is the use of software, data or commands to "exploit" a weakness in a computer system or program in order to carry out some form of malicious intent.

Key findings

The report makes five important key findings:

- The direct economic impacts of cyber events lead to a wide range of potential economic losses. For the cloud service disruption scenario in the report, these losses range from US$4.6 billion for a large event to US$53.1 billion for an extreme event; in the mass software vulnerability scenario, the losses range from US$9.7 billion for a large event to US$28.7 billion for an extreme event[b].

- Economic losses could be much lower or higher than the average in the scenarios because of the uncertainty around cyber aggregation. For example, while average losses in the cloud service disruption scenario are US$53 billion for an extreme event, they could be as high as US$121.4 billion or as low as US$15.6 billion[c], depending on factors such as the different organisations involved and how long the cloud-service disruption lasts for.

- Cyber-attacks have the potential to trigger billions of dollars of insured losses. For example, in the cloud-services scenario insured losses range from US$620 million for a large loss to US$8.1 billion for an extreme loss. For the mass software vulnerability scenario, the insured losses range from US$762 million (large loss) to US$2.1 billion (extreme loss).

- The scenarios show there is an insurance gap of between US$4 billion (large loss) and $45 billion (extreme loss) in terms of the cloud services scenario – meaning that between 13% and 17% of the losses are covered, respectively. The underinsurance gap is between US$8.9 billion (large loss) and $26.6 billion (extreme loss) for the mass vulnerability scenario – meaning that just 7% of economic losses are covered.

- When assessing current estimated market premiums against the forecasted cyber scenario insurance loss estimates set out in the report, it is apparent that a single cyber event has the potential to increase industry loss ratios by 19% and 250% for large and extreme loss events, respectively. This illustrates the catastrophe potential of the cyber-risk class.

[b] These figures represent the mean values of simulated loss year severities for large and extreme loss events, and take into account all expected direct expenses related to the events. Impacts such as property damage, bodily injury, as well as indirect losses such as the loss of customers and reputational damage are not taken in to account.

[c] These are illustrated as 95% confidence ranges – the range of values that act as good estimates to cover known and unknown parameters.

Conclusion

As the cyber threat increases so too does the demand for cyber insurance.

Despite this growth, insurers' understanding of cyber liability and risk aggregation is an evolving process as their experience of cyber-attacks increases. It is therefore important that risk understanding, including technical premium calculations and capital models, keeps pace with the changing cyber risk knowledge base.

In some other insurance classes insurers' understanding of liability and risk aggregation is more developed. It is widely accepted, for example, that natural catastrophes can trigger multiple claims from multiple policyholders, dramatically increasing insurers' claims costs. Natural catastrophe insurance policies usually take this into account and reinsurance is commonly used to reduce the impact of risk aggregation.

This report's findings suggest economic losses from cyber events have the potential to be as large as those caused by major hurricanes. Insurers could benefit from thinking about cyber cover in these terms and make explicit allowance for aggregating cyber-related catastrophes. To achieve this, data collection and quality is important, especially as cyber risks are constantly changing.

For the insurance industry to capitalise on the growing cyber market insurers would benefit from a deeper understanding of the potential tail risk implicit in cyber coverage.

Risk managers could use the cyber-attack scenarios to see what impacts cyber-attacks might have on their core business processes, and plan what actions they could take to mitigate these risks.

1. Introduction

The ability to write code and turn functions into software to complete complex tasks has brought efficiency in business administration, enabled advanced manufacturing, and is transforming industry and our day-to-day lives.

These coding outputs – along with the mass availability of programmable consumer devices to run this software – have given rise to new business models. Traditional players in capital-intense industries such as taxi livery and hotels are being challenged by peer-to-peer economy businesses such as ride-hailing and home-sharing companies.

This disruption is occurring as a result of applications of code and customers who are willing to engage with the new services and distribution models they create

However, the pervasiveness of digitisation means the global economy is now heavily reliant on a technology that can be vulnerable, and resulting software failures can have business, economic and insurance consequences. The complexity of the technology is increasing and with it the potential for vulnerabilities.

The cyber insurance marketplace

The economy is becoming increasingly digital, and this is seen in the estimation that 95% of companies in the Organisation for Economic Co-Operation and Development (OECD) have an online presence *(OECD, 2012)*. Organisations are increasingly aware of the reality of cyber risk in the 21st century and the World Economic Forum put cyber-attacks as the 12th largest risk to doing business in 2017 - ahead of natural catastrophes, which ranked 20th (World Economic Forum, 2017).

Organisations are responding to this greater risk awareness through their purchase of cyber liability insurance protection. In turn, the insurance industry is looking to develop solutions to protect those insurance risks at a time when there is limited publically available information on the potential range and scale of cyber events.

Today, Lloyd's Class of Business team estimate that global cyber market is worth between $3bn and $3.5bn *(Stanley, 2017)*; by 2020, some estimate it could be worth $7.5bn *(PwC, 2015)*.

Property/casualty insurers wrote $1.35 billion in direct written premium for cyber insurance in 2016, a 35% jump from 2015, according to reports by Fitch Ratings and A.M. Best. These figures represent a fraction of the US $528.2bn net written premiums for the whole insurance market that domestic carriers wrote in 2016 *(Weisbart, 2017)*.

Aon Benfield Analytics' report Cyber update: 2016 Cyber Insurance Profits and Performance, estimates roughly 85% of premium today is for US risk *(Aon Benfield, 2017)*, but this risk is shared with insurers in Bermuda and London insurers as well as the US *(Laux and Kerman, 2017)*.

Demand for cyber insurance is also anticipated to increase penetration in Europe as a result of the General Data Protection Regulation coming into force next year, with the threat of penalties for breaches driving coverage *(Ralph, 2017)*.

In some ways, the cyber insurance market can be considered in the same light as underinsurance in the natural catastrophe space – risks are growing and insurance penetration figures are low.

Sources of vulnerabilities

It is commonly accepted within the software development community that code is never released error free *(Chelf, 2009)* and industry average number of bugs for every 1,000 lines of code range from 15 to 50 bugs, *(McConnell, 2004)*

These errors, or bugs, generally occur due to trade-offs in areas such as time, features and cost *(Atwood, 2007)*. These bugs are frequently the mechanisms leading to vulnerabilities through which malicious actors can obtain the ability to bypass safeguards or misuse systems outside the intended purpose.

There are a number of common trends influencing the emergence of security issues and vulnerabilities in software. Understanding these trends to identify sources of cyber risk has never been more important, and these have been outlined in the 'house of cards' model of vulnerability (*see Figure 1, below*).

The house of cards model of vulnerability

Figure 1: Six software challenges

Source: Cyence

1. Volume of contributors

In 2011, Linux Kernel – an open source software project – had an estimated 1,400 separate contributors working on curating and developing a project with 15 million lines of code *(Corbet, Kroah-Hartman and McPherson, 2012)*. Proprietary software systems developed by commercial entities also use teams and outsourced contractors who are spread across the world, developing the code between them. The number of contributors collectively developing software has grown significantly over the past three decades.

2. Volume of software

In addition to the number of contributors amending code, the amount of code around continues to grow in volume and complexity.

The original Apollo 11 mission – a mission that landed people on the moon, utilised an estimated 145,000 lines of code *(Johnson, 2012)*. In comparison today's cars run more than 100 million lines of code *(AGC, 2017; Levine, 2012; Gelles, Tabuchi and Dolan, 2015)*.

> *"An increase of software volume (i.e. the number of lines of code) implies that more components are executed by different computers and connected through networks using specific protocols. This increasing use of software also increases its complexity; interconnected components perform various functions, potentially at different criticality levels."*
> *- Delange et al., 2015*

3. Open source software

Open source software has come a long way over the years and active coding communities such as GitHub are one of the primary reasons for its development and uptake. Open source software and other collaborative projects benefit through development or advocacy or because of the community *(Open Source Initiative, 2017)*.

As a result of these features, popular open source software projects are often described by industry experts as being "on the cutting-edge of technology" *(Noyes, 2010; Zivtech, 2015)*. More than half of the software acquired over the next several years is predicted to be open source *(Rowley, 2017)*.

New code may also be a potential point of vulnerability, and open source software is impacted by both the benefits and risks of using a collaborative development approach. Most third-party and open source components do not undergo the same level of security scrutiny as custom-developed software. Many of the open source software projects uploaded to Github and elsewhere are presumed to be reviewed for functionality and security, but in fact no standards definitions exist for this purpose.

4. Old software

The aging of software over time is also a concern. Running older operating systems has been proven to increase risk not only for the organisations using them, but even those outside of the network *(Mutton, 2015)*. One reason cited when individuals and companies are asked why they are reluctant to upgrade their software is that they are comfortable with the existing version's features and how the existing integration interacts with other legacy systems *(Tufekci, 2017)*.

As has been seen in the recent Wannacry attack, the longer software is out in the market, the more time malicious actors have to find and exploit vulnerabilities *(Ralph, 2017)*. It is true that new software can have bugs or weaknesses, so to counteract this, the majority of software vendors release new versions of software to provide greater functionality and security.

Despite the latter improvement, many individuals and companies run obsolete software that has more secure alternatives. For example, as of August 2015, Netcraft reported that: "More than 600,000 web-facing computers — which host millions of websites — are still running Windows Server 2003, despite it no longer being supported" *(Mutton, 2015)*.

5. Multi-layered software

Software is multi-layered and new software is typically built on top of prior software code, which creates many upstream/downstream inputs and dependencies that could track back to a defective line of code. This makes software testing and correction very difficult; especially considering that each line of code is likely to be part of an overall software system, as well as an individual entity itself. Any testing must ensure that all parts function correctly across the many layers.

The result of this is that there are more programmers today dedicated to the maintenance of existing code than there are programmers working on new implementations *(Jones, 2006)*.

6. "Generated" software

"Generative programming" is a style of computer programming that uses automated software creation – rather than being written by programmers – through generic frames, classes, prototypes, templates, aspects and code generators to improve programmer productivity.

Examples of this style of code generation include: plug-ins, device drivers, software extensions, program launchers and automatic software updates that can use a common legitimate framework. These elements can be modified for malicious intent, and purposely obscuring sources and code to make the program difficult to understand can further complicate detection *(Keizer, 2015)*.

Research approach
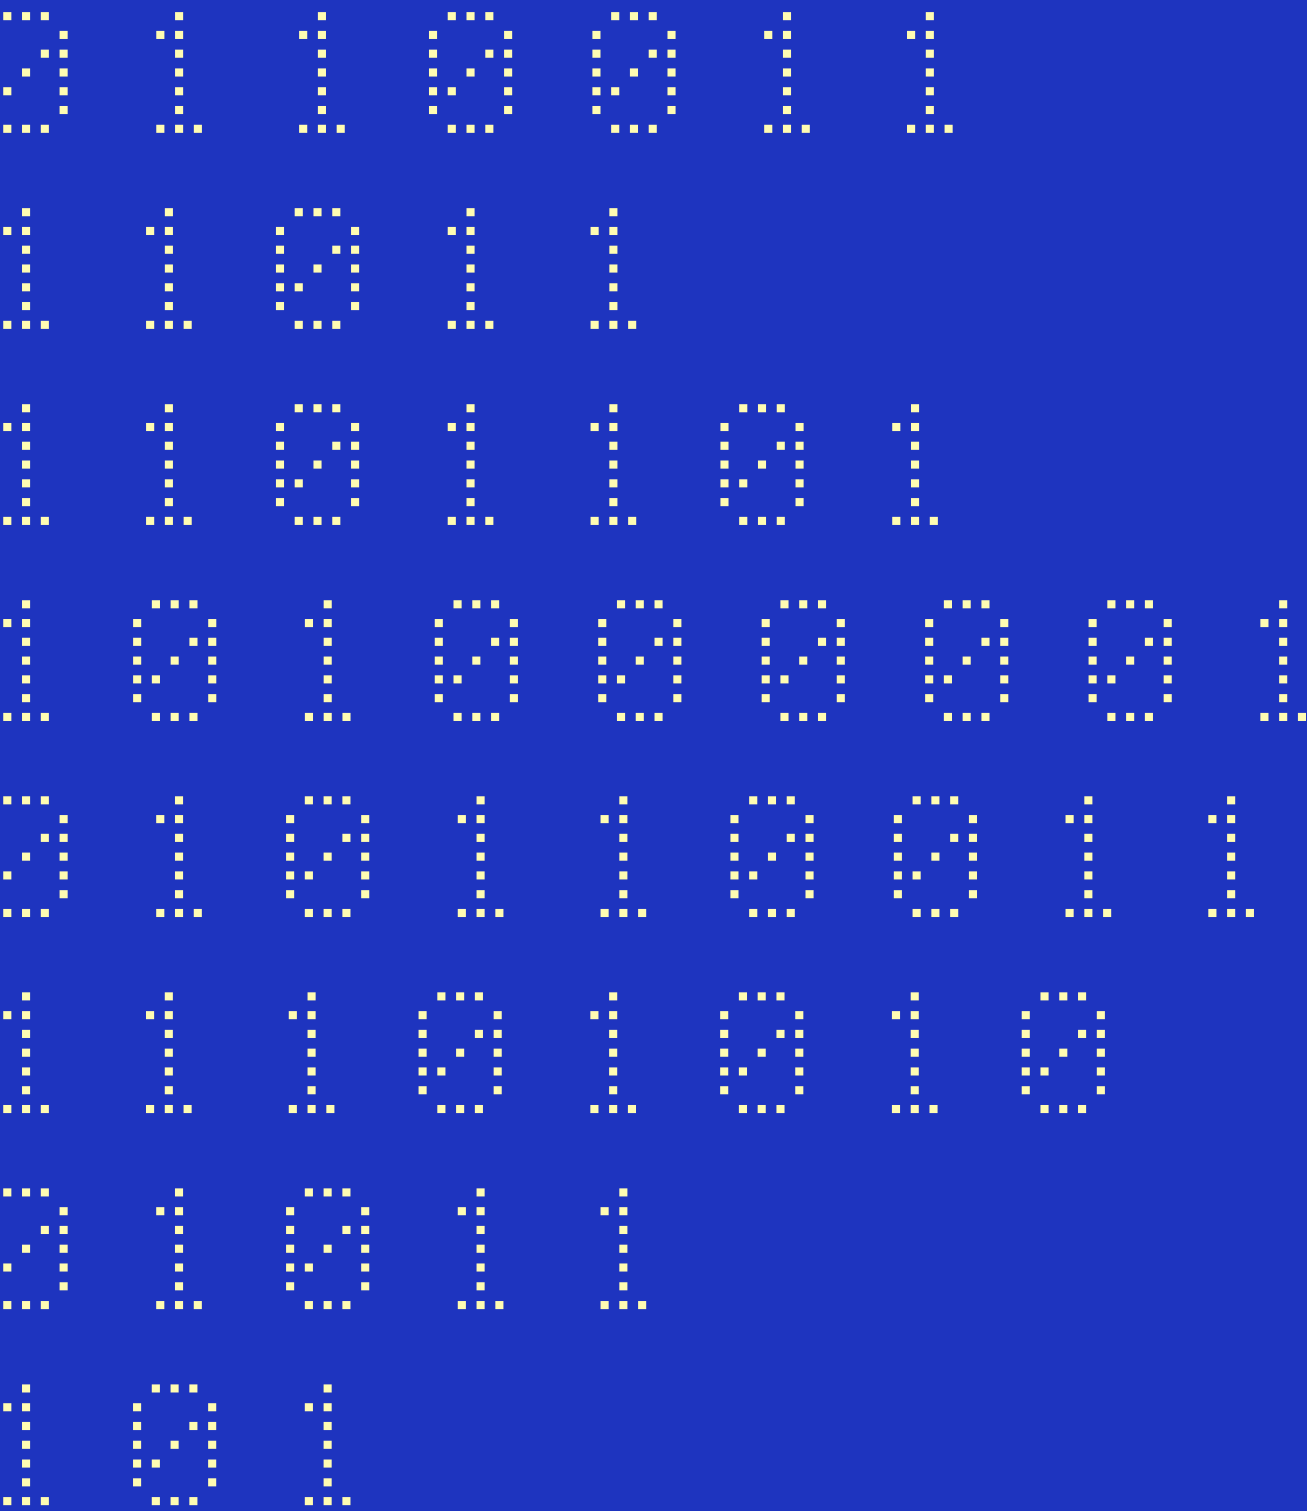
2. Research approach

This report was developed through a structured research process, across seven key stages:

1. Review of widely adopted technologies used across industries

A comprehensive review was undertaken to identify widely adopted technologies that could create exposure to aggregated cyber-loss events. Many technologies were identified which were grouped into two major categories:

– Service providers such as internet service providers, cloud-service providers, domain-name services, content delivery networks and payment processors

– Software – operating systems, web servers, database software, web applications, remote access, etc.

2. Review of other non-technical factors

Cyber risk factors are often thought of as encompassing the security measures a company has implemented. People also play a role in creating vulnerabilities to make attacks possible. This is why companies employing state of the art technology and risk management plans may find themselves subject to attacks.

When modeling cyber risk, it is important to consider both aspects. A review of these factors was undertaken as the probability of event occurrence, processes around mitigation and incident response can often differentiate between a low severity event and a high severity event.

Examples of human factors may include unintended acts such as lost or unattended devices, weak passwords or clicking on malicious links. Intentional acts may include insider attacks from past or present employees or external contractors who may have access to privileged information regarding the company or its security protocols. Training, establishing protocols and the adoption of industry standards across a company and any third parties can also act as factors affecting how an event may unfold.

3. Data collection and processing for the exposures

Traditional insurance risk modelling relies on authoritative information sources such as national or industry data, but there are no equivalent sources for cyber-risk and the data for modelling accumulations must be collected at scale from the internet. Insureds' use of the internet is also changing, causing cyber-risk accumulation to change rapidly over time in a way that other perils do not. This makes data collection, and the regular update of it, key components of building a better understanding of the evolving risk.

Cyence has conducted a comprehensive assessment of a range of companies to determine potential cyber security risk factors. This data set allowed for an assessment of the most common technologies used across industries. Working with industry experts to determine the usage of specific service providers and software systems allowed for specific modelling of two scenario's loss potential to organisations that have exposure to the affected system. The two scenarios are a cloud service provider hack and a mass vulnerability event.

4. Analysis of the exposure accumulation paths

Each identified aggregation path was assessed by Cyence risk modelers, insurance experts and the Cyence cyber team for its concentration and potential to cause economic loss.

5. Selection of scenarios, frequency and severity models

In collaboration with Lloyd's, initial economic losses were assessed against current market coverage offerings to select the two scenarios. For example, outages of key internet service providers would have the potential to cause significant economic losses; however, internet service provider outages are consistently excluded under standard public infrastructure exclusions and thus less relevant to the report. The loss model used was a stochastic probability maximum loss model that

addresses the spectrum of possible probabilities of an insurable direct loss, as well as a range of potential impacts.

Confidence intervals have been included for the "All Industry" loss category to provide a view of the variability of the projected losses based on the level of data available. These values are based on the range of values derived from stochastic simulations of the scenario parameters and include the relevant severity factors for each scenario. For the cloud service provider scenario this covers parameters such as system interruption duration, business dependencies, and the effectiveness of business continuity contingency planning, and for the mass vulnerability scenario this incorporates percentages of affected organisations that experience a breach and the size of the breach.

6. Discussion and review with insurance and cyber security experts

The scenarios were reviewed by a host of experts within the Lloyd's market, as well as variety of cyber experts with cyber-security backgrounds to ensure the plausibility of the scenarios and applicability to common insurance coverage offerings, given the variability in applied coverage terms and conditions.

7. Loss calculations and final review

Once the scenarios were finalised Cyence set out to model the direct economic impacts according to the methodology outlined in this study, leveraging real-world organisation exposure data in the process. Using this data allowed the exploration and understanding of the subtle variations that can influence which organisations may see economic losses.

Cyence then used probabilistic scenario modelling to predict the chances of exposed organisations seeing a breach, as well as the potential severities of such incidents based on historical loss costs. These loss estimates were calculated across a global spread of organisations, and are listed in the report by industry class and revenue size.

Experts from cybersecurity, loss modelling, and cyber and commercial insurance sectors provided a final review of the modelled scenario loss figures as a final review step.

Research approach outcomes

Lloyd's and Cyence hope this study deepens insurers' and risk managers' understanding of cyber insurance risk, stimulates new ideas and raises new research questions. The two scenarios should help guide all interested stakeholders with an interest in ensuring that online processes continue to function as intended to fuel the modern economy. Continued innovation, analysis and collaboration across sectors and industries are critical to address current and future vulnerabilities and build more resilient, flexible digital networks.

The current state
of cyber coverage

3. The current state of cyber coverage

Despite being a consistent top-three risk in many prominent risk surveys for potential buyers, cyber insurance has relatively low penetration rates, especially among SME and middle-market customers, as well as in several industry verticals. According to a report by Deloitte *(Friedman and Thomas, 2017)*, buyers lack understanding of cyber risk about what is and isn't covered under existing insurance policies. There is also a general lack of standardisation around cyber insurance offerings in the marketplace, which makes it hard for risk managers to choose which product to buy. Brokers and insurance companies must do more to address these educational gaps to drive further growth of this important business line.

For the purposes of this report, Cyence provides two scenarios that would trigger most cyber insurance policies. However, the relative newness of the coverage and pace of innovation in the space has resulted in a lack of standardisation in terms and conditions.

For example, the definition of a computer system can vary among policy forms, with some policy wordings including systems owned by outsourced cloud providers who may hold data, while others may cover systems that can be considered strictly under the insured's care, custody and control.

Variation also occurs across the market in the policy limits coverage to notifications that are required by law or regulation, while others may offer voluntary notification cost coverage in the event of a breach.

While there are nuances to policy language and service offerings, there are emergent structures that most cyber policies adhere to. The following section illustrates a number of aspects that are important to consider when exploring the potential impacts and severity of future events – whether through the two scenarios in this report, or through alternative cyber scenarios developed subsequently as approaches evolve.

Common coverage types

Security and privacy liability

In the policies assessed as part of the review stage, nearly all monoline cyber policies included a third-party liability coverage section (Advisen, 2016). These would be triggered by data security and privacy-related litigation, and include defence costs, settlements and judgements.

Data breach costs

Data breaches involving sensitive information of individuals and companies are increasingly driving the introduction of legal obligations to notify the affected individuals. In the US, 47 states have data breach notification laws *(NCSL, 2017)* and other countries including EU nations and Australia have passed similar regulations.

While there is significant variation in the specific contractual coverage provided under this policy section, examples include:

- Hiring of IT forensic consultants to identify the scope of the breach
- Legal representation to advise on obligations based on the geographic and regulatory specifics of the situation
- Letters to notify potentially affected individuals (some policies may require a regulatory requirement to notify, with others covering voluntary notification costs)
- Funds to retain a public relations consultant to help manage publicity around the event
- Credit monitoring, identity theft monitoring, or identity theft insurance for potentially affected individuals (the length of monitoring covered can vary from policy to policy)
- Call centres to respond to queries from potentially affected individuals

Network business interruption

Generally cyber policies provide business interruption coverage that responds when there are outages and disruptions to a company's digital environments. Policies typically cover necessary ongoing operational expenses, any additional expenses and business income lost as a result of network interruption.

In addition to the set self-insured retention value, this coverage aspect is seen to involve a time-based waiting period ranging from eight to 12 hours that must be satisfied for coverage to be triggered.

Most businesses are reliant on their information technology and cyber footprints in some form or another to conduct business and interact with their supply chain. Policies often restrict coverage for contingent business interruption events caused by service providers.

There is a wide variety of coverage offerings for contingent business interruption – some policies exclude these events completely, others sublimit the coverage, others provide an hourly sublimit and some carriers offer options for full limit coverage. A network outage at a cloud service provider is likely to translate into a business interruption for the companies using that service provider.

There are also important coverage differences between policies as to what constitutes an interruption and how that loss is calculated. Some markets offer expanded coverage triggers called "system failure" or "administrative errors" coverage that aims to cover network outages caused by internal errors and omissions. This is a significant expansion of coverage over the ordinary trigger, which is limited to malicious acts from outsiders like a Distributed Denial of Service (DDoS) attack.

Regulatory action costs

Governmental regulatory bodies across many jurisdictions can bring actions against organisations for failure to comply with laws and regulations regarding information security and privacy. This coverage will continue to be increasingly important as more countries adopt data-breach regulations.

Examples include:

–   US healthcare organisations are held to standards of care for safeguarding sensitive patient information by the Health Insurance Portability and Accountability Act, which is enforced by the Office of Civil Rights.

–   Organisations in the US with smartphone applications may be required to disclose details of the types of information collected, how that data is protected and whether it is shared with third parties by the various consumer protection agencies, including the Federal Trade Commission.

–   Most recently, the European Union passed the General Data Protection Regulation, (which requires organisations operating in EU jurisdictions to notify the authorities in the event of a data breach. Additionally, this regulation includes potential fines up to 4% of annual worldwide turnover for non-compliance *(General Data Protection Regulation).*

Amongst stand-alone cyber policies, this coverage is often sublimited in the amount of coverage provided and restricted to the type of losses covered. Some policies are seen to cover the cost of fines and penalties to the extent that they are insurable by law, while others will only cover the cost to defend and comply with the regulatory investigation. This is particularly relevant to large global organisations that are subject to many regulations across various jurisdictional borders.

Additionally, some fines like those levied by the Payment Card Industry (PCI) can sometimes be considered a contractual obligation and not a "fine" since the PCI isn't considered a governmental body. These details can affect coverage, depending on the language in the various policy forms.

Extortion

Cyber extortion has risen extensively overt the past few years. One common extortion method is implemented using ransomware – a malicious software that when installed disrupts the computer, sometimes encrypting or corrupting files and demanding payment to remedy the infection and decrypt the files.

Stand-alone cyber policies will typically offer coverage for such demands, subject to the insurer's prior approval. One consideration for coverage includes whether the affected system could be remediated for a lesser amount than the demand, a value that may be derived from the type of records involved.

Digital asset replacement

This coverage is intended to assist in the replacement of digital assets that are damaged or destroyed as the result of a malicious act. If adequate backups are not available and system data has been destroyed, this coverage will cover the cost to restore, recreate or re-collect the digital assets specified where possible.

Cyber endorsements:

Some insurers identified during the project offer cyber coverage via endorsement onto standard general liability insurance policies. Cyber coverage offered by endorsement will have low sublimits of coverage and may only cover liability and breach response costs.

In late 2016, the International Organisation for Standardisation released a standard affirmative coverage endorsement, The Information Security Protection Endorsement, BP 15 07 03 15, which looks to create standardisation of cyber coverage across the industry *(Insurance Services Office, 2016)*. There are three tiers of coverage specified:

– Tier 1 - Breach response only

– Tier 2 - Breach response and liability

– Tier 3 - Breach response, liability, business interruption and extortion

Attacker types and motivations:

Cyber is an insurance business line that aims to cover a loss caused by an adversarial opponent. One of the challenges is that there are many potential factors that can shape a cyber event. When modelling risk, it is important to consider who the attacker is and their motivation, as well as the resilience of the affected organisation.

Attackers can range from white hat hackers[d], to cyber criminals, hacktivists, nation states and malicious insiders. White hat hackers may be motivated by the challenge of overcoming security systems. Criminal organisations will typically be motivated by financial gain, hacktivists by notoriety, and nation states by espionage and political ends.

Attacks from rogue employees who are likely to have access to insider information surrounding an organisations' IT system architecture and security protocols, as well as some form of privileged access by virtue of their employment, can be particularly problematic. Gartner's Understanding Insider Threats report *(Chuvakin, 2016)*, sets out the top three reasons for an insider attack:

– To generate a second stream of money from selling the stolen assets

– To take advantage of knowledge before changing employers, or

– To sabotage their employer

---

[d] Computer security specialist who break into protected systems to test and assess security

Cyber incidents may occur without any malicious intent. As noted in the introduction, software and IT systems are increasingly complex and mistakes can occur. This was highlighted in February 2017 when Amazon Web Services had a four-hour outage as a result of a typing error by one of their employees *(Amazon Web Services, 2017)*.

TRIA, war exclusions and complications with attribution

In order to provide background context in both scenarios, we have specified particular attackers – a hacktivist for the cloud service provider hack scenario, and criminal organisations for the mass vulnerability scenario. Specifying attackers that are not engaging in terrorism or warfare assists our analysis by making clear that the Terrorism Risk Insurance Act (TRIA) would not be triggered and also avoids potential insurance coverage issues as described below.

It is difficult to attribute a cyber-attack to a particular group or actor. This is especially true for sophisticated actors who are concealing their identities using virtual private network providers established for malicious reasons, Tor browsers[e] and other complex techniques to operate in anonymity.

The latest set of WikiLeaks documents included the CIA's previously secret anti-forensic "Marble Framework". Marble plants "false flags"[f] by obfuscating the code of cyber weapons to appear as though they were created by other nations. Test examples of code in the leaked documents included samples in Chinese, Russian, Korean, Arabic and Farsi *(WikiLeaks, 2017)*. Anonymity and misdirection are possible for an attacker(s) if they have the appropriate skill sets and resources.

Due to these factors, confirmed attribution can be difficult to achieve and is often based on an estimation of capabilities rather than evidenced fact. This has important implications for the cyber insurance marketplace, as most policies exclude acts of war, and the issues with attribution described above may seriously hamper an insurer's ability to assess whether claims are payable.

---

[e] Tor is short for The Onion Router and was initially established as a worldwide network of servers developed with the U.S. Navy that enabled people to browse the internet anonymously. Today, the initiative is a non-profit organisation engaged in research and development of online privacy tools *(Klosowski, 2014)*.

[f] The term has naval origins and is used to describe a situation when in times of war, ships would sometimes change the national flag they flew in order to fool other ships. The term can be applied to cyber attacks with a nation state using technology to hide their identity behind an identity that would point to another nation.

Courts typically adjudicate ambiguity in insurance contracts in favour of the insured. There are cases going back more than 100 years on standard insurance coverages that debate where lines in coverage are drawn *(Massman, 2001)*. The majority of cases have sided with the insured, toward a strict interpretation of what constitutes "war"; however policies have become more specific recently, especially after the attacks of 11 September 2001.

After this event the US government enacted the TRIA, which provides a financial backstop for insurance claims arising out of declared acts of terrorism. There has been political uncertainty on the continuation of TRIA and it lapsed for the first week of 2015. TRIA has been reauthorised through to 2020 through the Terrorism Risk Insurance Program Reauthorization Act and the US Department of Treasury recently deemed cyber liability as contemplated under TRIA *(Department of the Treasury, 2016)*.

To further complicate matters, cyber coverage has been written on blended forms that include professional liability coverage, which is specifically excluded from TRIA *(Kalinich, 2017)*. Developing policy language and challenges with attribution like false flag operations should be considered by insurers going forwards in relation to their exclusions or the TRIA backstop. The fast-evolving world of cyber attacks, the available evidence and the certainty of perceived facts will influence how these points are decided in law courts going forwards.

The scenarios described in this report are not considered to trigger any war or terrorism exclusions.

The scenarios

Technology
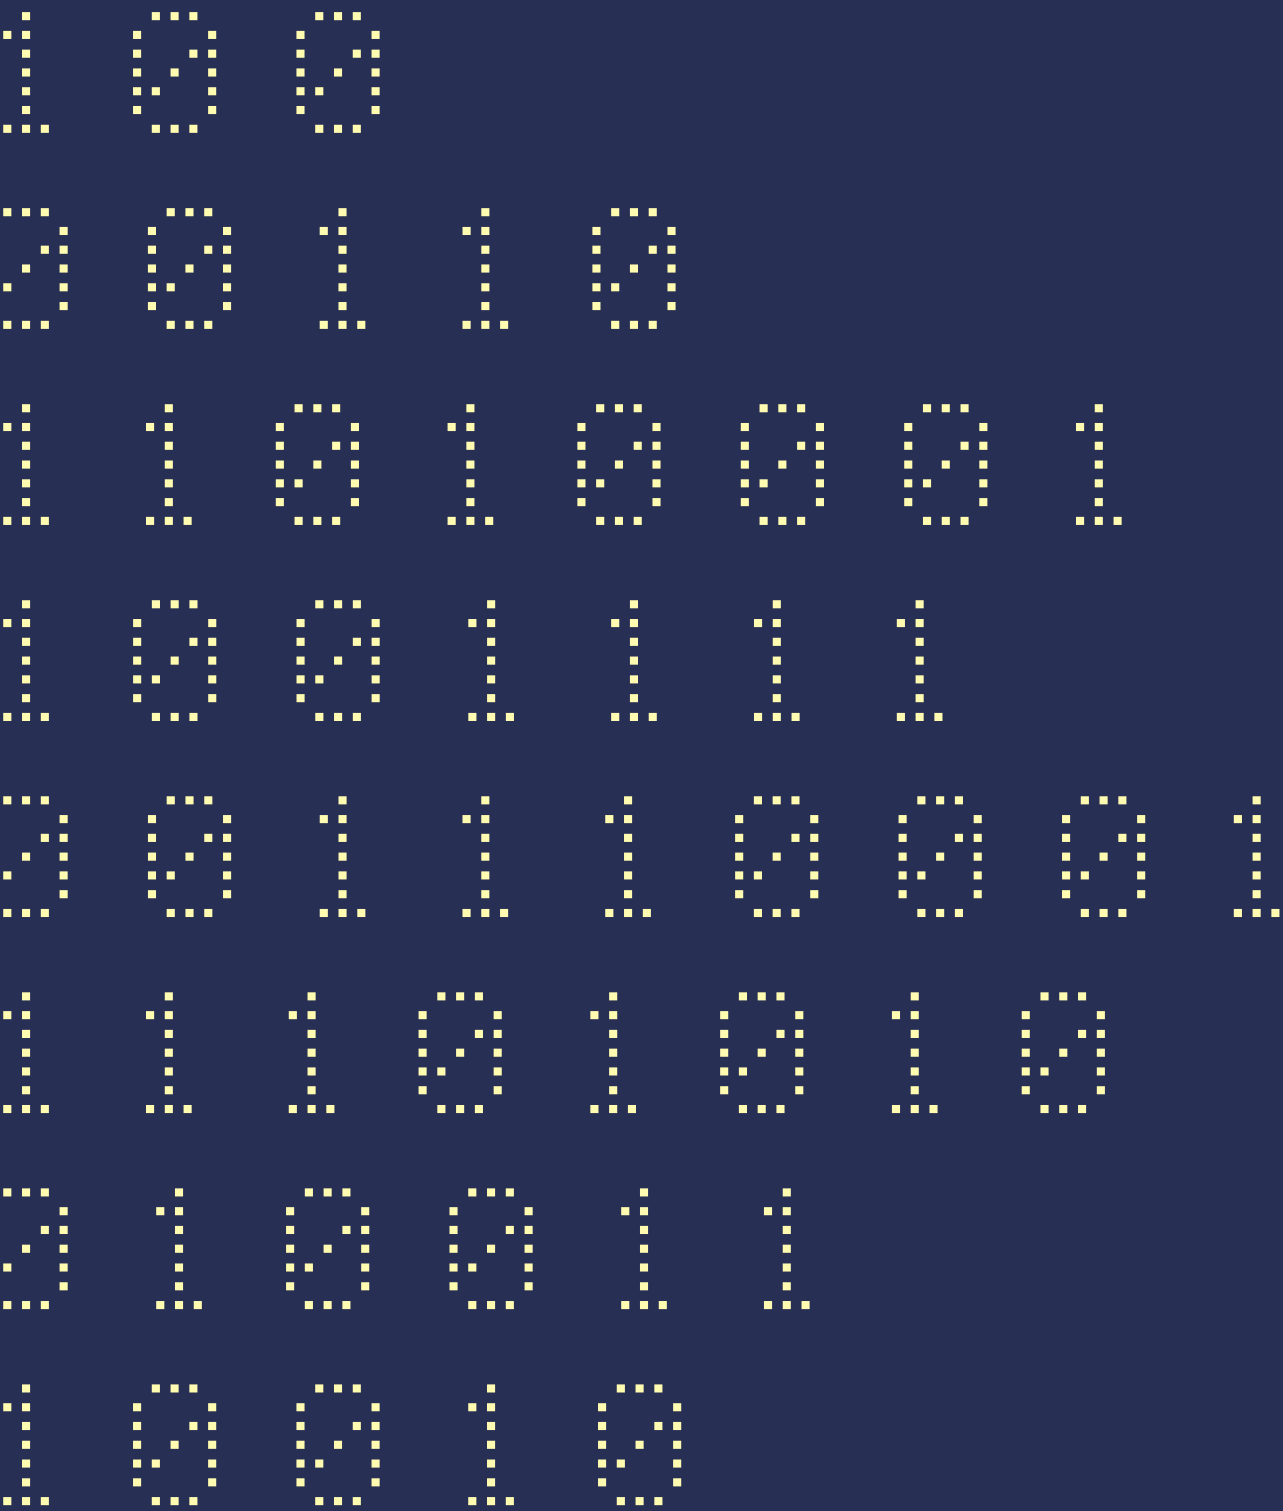
4.1. Cloud service providers

The concept of network-based computing can be traced back to the 1960s but it wasn't until the early 2000s that the idea of "the cloud" became a recognisable concept for describing the process of accessing software, computing resources and data over the web instead of from a local computer. The branding of cloud services increased substantially once they became accessible to small and medium-sized businesses, as well as every-day consumers.

Individuals and businesses have cited "flexibility" and "efficiency of cloud systems" as some of the reasons for this increase in use and today consumers interact with them on a daily basis, whether this is by streaming movies from services such Netflix or sharing photographs through websites such as Photobox.

As a result of these developments, the Cloud Industry Forum cites the overall cloud adoption rate in the UK as now standing at 88%, with 67% of users expecting to increase their adoption of cloud services over the coming year. McAfee further supports these findings, with its Building Trust in a Cloudy Sky survey stating that: "Cloud services are widely used in some form, with 93% of organisations utilising software-as-a service, Infrastructure-as-a-service, or Platform-as-a-Service offerings" *(McAfee, 2017)*.

Research commissioned by Forbes illustrates an estimate of cloud infrastructure services' market share by core providers in this space (*see Figure 2, right*).

Figure 2: Market share for cloud infrastructure services by provider
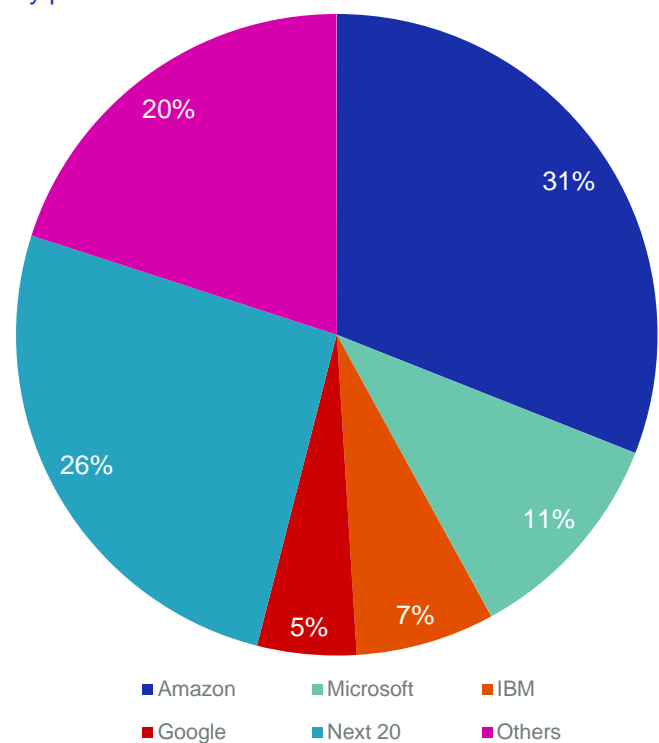
Source: Forbes, 2016

There are three main types of cloud computing services:

– **Infrastructure as a service (IaaS):** physical computing resources such as servers and value-added security services

– **Platform as a service (PaaS):** development environments sold on a subscription basis

– **Software as a service (SaaS):** web-based software and database subscription services

While these services should be considered on their own, they also operate with one another as part of their deployment. PaaS and SaaS stack on top of IaaS as a basic building block to any cloud service. This layering of infrastructure uses technology called a hypervisor to distribute a server's resources to the many virtual machines within it.

Hypervisors

A "virtual machine" is a virtual simulation of a physical computer system, created through software implementation on a single piece of hardware. Virtual machines can be deployed using specialised hardware, software, or a combination of the two, and are sometimes called "hypervisors" or "virtual machine monitors".
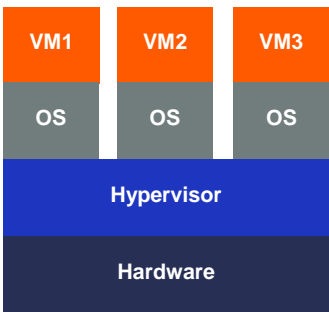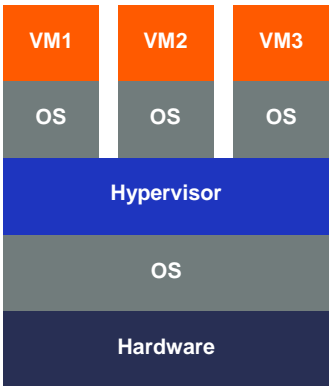
A hypervisor creates a virtual platform on the host computer, on top of which multiple guest operating systems are executed and monitored. This way, multiple operating systems, which are either multiple instances of the same operating system or different operating systems, can share the hardware resources offered by the host *(Oracle, 2014)*.

Hypervisors are one of the fundamental building blocks of the cloud: they maintain the separation and privacy of neighbouring virtual machines, thereby enabling the entire ecosystem to function. Isolation between virtual machines is a key element of security, and privacy between each implementation and therefore any disruptions or breakdown to this function could lead to wide-scale interruption for all virtual machines residing on a server.

Types of hypervisors

There are two types of hypervisors:

Table 2: Hypervisor classifications

Classification	Characteristics and description
Type 1: native or bare metal VM1 VM2 VM3 OS OS OS Hypervisor Hardware	Type 1 hypervisors have directed access to system hardware and are often referred to as a "native","bare metal" or "embedded" hypervisors *(Ruest, 2010)*. With a type 1 hypervisor, there is no operating system to load as the hypervisor that users load is the operating system *(Davis, 2013)*. Examples of current type 1 hypervisors include: Microsoft Windows Server Hyper-V, VMware vSphere ESXi, Xen / Citrix XenServer and Red Hat Enterprise Virtualisation (RHEV) *(Davis, 2013)*.
Type 2: hosted VM1 VM2 VM3 OS OS OS Hypervisor OS Hardware	Type 2 hypervisors run on an ordinary operating system just like other applications do. In this case, guest-operating systems run as a virtual machine on the host and are abstracted from the host's operating system. Examples of current type 2 hypervisors include VMware Workstation, VMware Player, VirtualBox, Parallels Desktop for Mac and QEMU *(Ruest, 2010)*.

Sources of risk

Because of these features of cloud-based infrastructure, there is the potential for systemic risk and interdependency, as companies who are reliant upon common infrastructure will suffer business interruption or outages simultaneously when that infrastructure is compromised or incapacitated. For example, if a major security flaw was found in a commonly used hypervisor, cloud customers of service providers using it to segment their virtual environments could suffer from a breach on all the shared systems connected to that hypervisor.

Attacks on these systems could result in cascading outages within supply chains and the potential for significant losses arising from data breaches and system outages. Insurers offering cyber insurance should therefore consider and explore the potential for substantial losses within portfolios due to possible connections of their insureds sitting within widespread supply chains of core cloud-service providers.

The following section informs the scenario and describes known vulnerabilities that are exploited to cause extreme but plausible losses.

Common sources of risks that have been exploited in known examples include:

– Software vulnerabilities

– Backdoors

– Race conditions

Software vulnerabilities

Hypervisors are software and, as described in Section 1, all software has the potential for vulnerabilities due to its very nature. Figure 3 (*overleaf*) illustrates known vulnerabilities in Xen and VMware "ESXi OS", two of the examples of hypervisors described in Table 2 (*p21*).

Figure 3: Software vulnerabilities by type and year
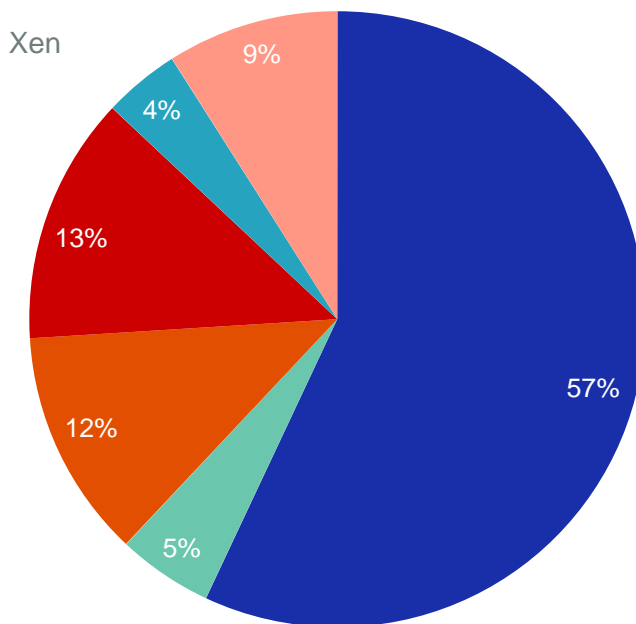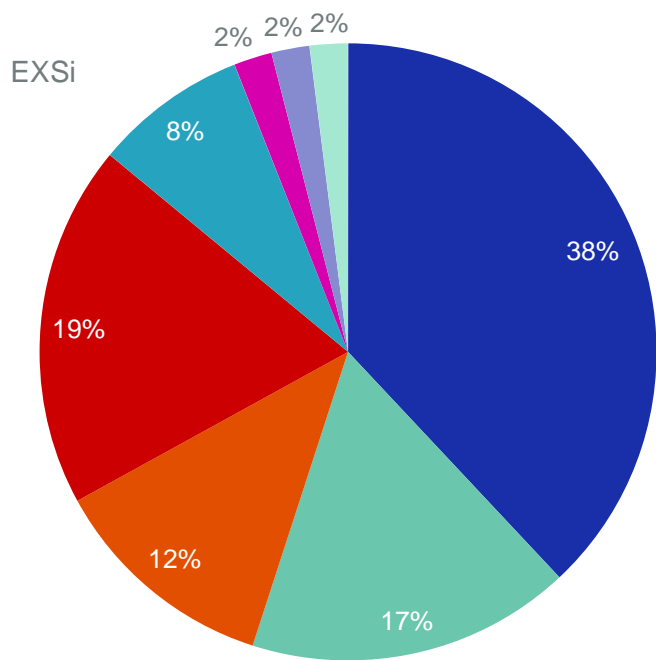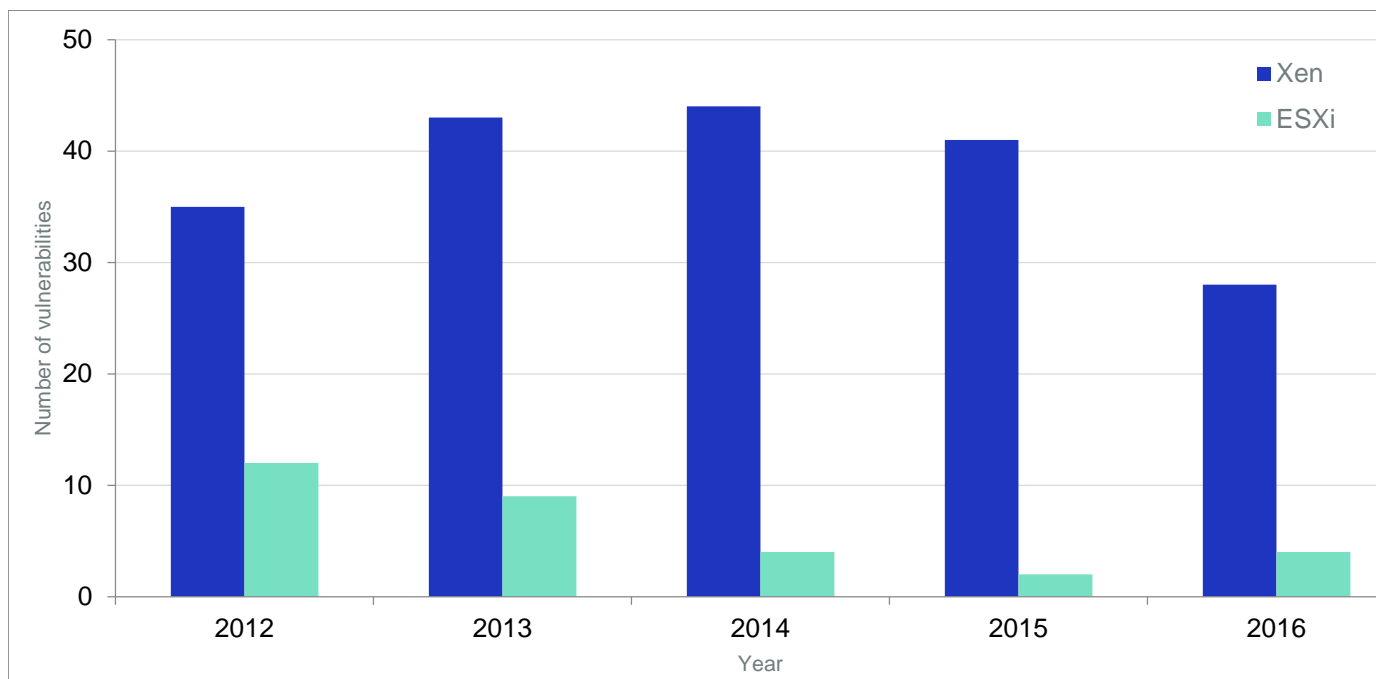
The majority of vulnerabilities allow for Denial of Service.

Source: Özkan, 2017a; b

Open source software has varying levels of review process. When the coding community is small, mistakes can make it past the peer-review process and have real-world implications, as was demonstrated with Heartbleed, a vulnerability in the open source OpenSSL protocol that powers many secure website communications. The vulnerability was introduced by a programmer's added feature. It was not caught by the code review and validation process, and eventually made its way from a development branch into a new release of the software *(Vaughan-Nichols, 2014)*. Similar problems exist in closed sourced software, but this is often retained internally within a company.

Remediating systems with vulnerabilities often requires rebooting the system after the fix has been applied. Cloud providers avoid system reboots at all costs since restarting even a small proportion of their machines becomes an enormous task due to the scale of their operations. Modern applications built to utilise the cloud can be built for high availability[g], incorporating failover capabilities onto other virtual machines to limit the impact of system downtime. However, legacy applications and systems that have been migrated to the cloud typically are less resilient. These older applications and systems can be more difficult to remediate and involve longer downtimes.

Problems with hypervisors can lead to large-scale issues. For example, in 2015, Amazon AWS had to reboot its systems on two occasions due to Xen hypervisor patches that required a full restart of all affected systems. The software flaw, CVE-2015-7835, affecting the Xen hypervisor allowed malicious actors to create guest servers which could access the host computer's memory and take control of the entire system. This security breach went undiscovered for more than seven years.

Examples of past software vulnerabilities

Box 1: Xen

Xen is a market dominant hypervisor for cloud service providers such as Amazon, Rackspace and IBM. It is an open source project managed by the Linux Foundation with more than 50 independent organisations from both private and public sector working on the codebase, in addition to individual contributors. Xen now has more than 1.1 million lines of code *(Corbet, Kroah-Hartman and McPherson, 2012)* and 17 supported integration platforms *(Linux Foundation, 2017)*. This growth has created complications in code management and "The Xen Project" has stated in the past that code review has not caught up *(Kurth, 2016)*. An example of this is that in 2016, 28 public vulnerabilities were discovered in the Xen System *(Özkan, 2017b)*.

Box 2: ESXi

The ESXi hypervisor, created by the company VMware, is another system that has been in production since 2002. VMWare saw great early success in the private cloud space and continues to expand its market share across the globe. From 2012 to 2016, 31 CVE have been assigned to various vulnerabilities found in the ESXi hypervisor platform *(Özkan, 2017a)*. Nine of them have been identified as remote code execution. During that same time, VMWare released 81 bulletins specifically addressing security related issues *(VMWare, 2017)*.

Box 3: KVM

KVM emerged in 2006 as a free open source Linux-based hypervisor. KVM stands for Kernel-based Virtual Machine and was developed by the Open Virtualization Alliance organisation. As of version 2.6.20, KVM has been included in the Linux kernel base and bundled with the Linux operating system. Vulnerabilities affecting KVM have been routinely discovered, such as virtual machine escapes to denial of service *(Steinberg, 2015)*.

---

[g] The introduction of additional equipment to cover against possible failure or malfunction that is not in use for the majority of time, but can support content and traffic if the primary systems fail.

Backdoors

A backdoor is a method of bypassing normal authentication and security measures in a product or computer system. Backdoors are often legitimately included in software and made known by manufacturers as a means of restoring locked user accounts, debugging the product and other use cases.

However, a malicious actor in the role of a developer or open source contributor could surreptitiously include code for backdoors into production code. In this function, backdoors could serve as a means of securing unauthorised remote access to computer systems or as a kill switch for anarchist hackers looking to cause a wide-scale system outage large enough to force people to take notice. It is also possible to use malware as a method of installing a backdoor on affected systems as a means for malicious actors to get back into systems after they have been remediated.

With minor and unintended flaws in code having such serious implications, the problem becomes a cryptographic puzzle with endless combinations. An example of this was in February 2017 when Amazon experienced a four-hour interruption at one of its largest data centres due to a typographic error *(Stevens, 2017)*. See Table 3 for further examples (*overleaf*).

An analyst isn't looking for an obvious error rather the potential for an error, and as the code looks the same and the difference is in assessing the context and interpretation, the process is a never-ending, complex task.

Race conditions

Many software protocols have time dependencies that need to be accurate for the overall system to function. Inaccuracies are known as race errors, which can be described as anomalous behaviour due to unexpected critical dependence on the relative timing of events. For example, if one process writes to a file while another is reading from the same location then the data read may be the old contents, the new contents or mixture blend of the two, depending on the relative timing of the read-and-write operations. This causes a bug in the software *(FOLDOC, 2002)*. If this bug is in a critical portion of the code such as the booting procedures of a system, it could cause persistent fatal errors that are capable of incapacitating it.

Examples of known race conditions

One example of a race-condition error affecting a critical computer system occurred in August 2003 in parts of the north-eastern and mid-western US and the Canadian province of Ontario. A widespread electric grid outage in these areas occurred when a manageable number of powerlines were knocked down by a storm. At the same time, General Electric Energy's Unix-based "XA/21" energy management system had a race-condition error that stalled the grid operators' alarms for more than an hour, which resulted in cascading blackouts that took between several hours and days to resolve. Roughly 500 generating units at 265 power plants shut down during the outage, resulting in a peak 80% decrease in power output *(Wander, 2007)*.

Examples of known backdoors

Table 3: Timeline of notable backdoors

Date	Vendor Product	Description
2016	WordPress	Custom Content Type Manager, a WordPress plugin with more than 10,000 active installations, started stealing admin credentials via a backdoor. The culprit appeared to be an auto-update.php file recently added to the plugin, which was actually a backdoor that could download files from a suspicious wordpresscore.com domain *(SecurityWeek News, 2016)*.
2014	WordPress	A serious vulnerability on the MailPoet WordPress Plugin, a very popular plugin with almost two million downloads, was detected. The vulnerability allowed an attacker to inject anything they wanted on the site, which could be used for malware injections, defacement, spam and many more nefarious acts. Once they succeed in uploading the malicious theme, they access their backdoor inside /wp-content/uploads/wysija/themes/mailp/: and have complete control of the site *(Cid, 2014)*.
2013	Dual_EC_DRBG (encryption software)	As reported in an article in 2013, in 2006, a US federal agency, the National Institute of Standards and Technology, helped build an international encryption system to help countries and industries fend off computer hacking and theft. "Unbeknown to the many users of the system, a different government arm, the National Security Agency, secretly inserted a "back door" into the system that allowed federal spies to crack open any data that was encoded using its technology" *(The New York Times, 2013)*. This backdoor was referred to as the DUAL EC DRBG algorithm.
2008	Juniper Networks	In 2015, Juniper Networks announced that unknown attackers had added unauthorised code to ScreenOS, the operating system for their NetScreen VPN routers. This code created two vulnerabilities: an authentication bypass that enabled remote administrative access and a second vulnerability that allowed passive decryption of VPN traffic. Analysis by researchers at various sources indicates that the backdoor was inserted in 2008 into the versions of firmware ScreenOS from 6.2.0r15 to 6.2.0r18 and from 6.3.0r12 to 6.3.0r20 *(Checkoway et al., 2016)*.
2003	Linux kernel	Software developers detected and thwarted a hacker's scheme to submerge a backdoor in the next version of the Linux kernel. Security experts say the attempt indicates that subtle source-code tampering is more than just paranoid speculation. The backdoor was a two-line addition to a development copy of the Linux kernel's source code, carefully crafted to look like a harmless error-checking feature added to the wait4() system call - a function that's available to any program running on the computer and which, roughly, tells the operating system to pause execution of that program until another program has finished its work *(Poulsen, 2003)*.
2001	Borland Interbase	Borland Interbase versions 4.0 to 6.0 had a backdoor feature, an innocent addition to the code in 1994 that enabled one part of the database software to communicate with another password-protected part put there by the developers. The server code contained a compiled-in-backdoor account (username: politically, password: correct), which could be accessed over a network connection. Once a user logged in with it, they could take full control over all Interbase databases. The backdoor was detected in 2001 and patch was released *(Shankland, 2001)*.
1998	Microsoft Windows	Back Orifice was a backdoor created in 1998 by hackers from the Cult of the Dead Cow, apparently to highlight Microsoft's lack of security. The backdoor allowed the sender to remotely control and monitor a computer running Windows 95 or 98. Once installed, the program did not show up in the user's task manager, giving it the potential to run undetected. Microsoft issued a patch *(Festa, 1998)*.

The examples above demonstrate the prevalence of attack vectors across various software platforms, both proprietary and open source.

4.2. Modelled scenario: Cloud service provider hack

Note: the report does not name particular companies and details should be substituted to explore scenario variants within portfolios.

Background

A sophisticated group of "hacktivists" sets out to disrupt cloud-service providers and their customers. Their aim is to draw attention to what the group considers the environmental impacts of business and the modern economy. The hacktivist group determines that a modification to a hypervisor controlling the cloud infrastructure could trigger many customer servers to fail, causing wide spread service outages and business interruption.

Planning

The opportunity came about when one of the hackers, who is on a team creating an important new feature for ABC hypervisor that was offered publically as open source, introduced modified code into the feature. The code modification was designed to easily blend into the added feature and pass the code review. The code modification is designed to create a race condition that will trigger system crashes on a specific date and time in the future.

Figure 4: Timeline of events for cloud service provider malicious code outage in days

Day 1	Day 185	Days 185 – 250	Day 365
Code inserted in a feature update of a particular version of ABC	Feature widely used and integrated into ABC at update	Other CSP's running slightly differing versions of ABC integrate feature	Timer goes off, Malicious code enters hypervisor and triggers persistent shut down command

Day 366	Day 367	Day 368	Ramp up
Forensics and security researchers identify vulnerability	Tier 1 CSP's start to deploy patch	Tier 2 CSP's receive patch and begin deployment	CSP and server restoration takes place over the following days

Activation

After a year of dormancy, the malicious code has been integrated into most ABC versions running at cloud service providers. When the system clocks reach the specific hard-coded date and time, the malicious code executes with no further action being taken by the group, causing many servers to crash and wide spread system outages to ensue.

All affected servers crash and will not reboot properly. Large numbers of virtual machines are immediately brought down simultaneously.

Identification

The cloud service providers quickly notice their systems are having trouble across multiple locations and that the servers are stuck in a cycle of rebooting and crashing again. Administrators immediately begin evaluating how many systems are impacted and forensics teams are called in to begin their analysis.

Since the malicious code is exploiting a race condition, the actual event is not easily reproducible despite the significant resources and expertise deployed on the issue.

After 24 hours of searching the hard drives, network traffic and logs, these teams identify the backdoor and malicious code within ABC that has caused the server crashes. A massive joint effort begins immediately to create an emergency patch that takes an additional 24 hours to complete.

At that time, thousands of customers are affected by the incident and the affected cloud service providers are moved to prioritise key customers' systems reboots.

Response

Due to the nature of virtualisation it is difficult to prioritise individual customers, and as a result of this the cloud service providers focus on high-priority regions while installing the latest patch and bringing systems back online. This process takes a minimum of six hours to get the high-priority regions back online.

Cloud service providers without sufficient resources lack the capabilities needed to conduct these operations at speed and scale. It takes CSPs in this category an additional 24 hours to receive the patch and 24-48 hours to reboot, install and ramp up their services.

Once a cloud service provider is able to restore service, there is additional reboot time for the end-users to get their systems back online and running. Large online retailers often have robust resiliency and restoration times set out for their critical servers as part of business continuity planning that may take minutes to restore.

The majority of companies impacted will require at least several hours to bring systems back online and it could be anticipated to take longer than a day if the appropriate IT expertise isn't immediately available - small and medium-sized end users with outsourced or inexpert IT functions, for example.

Box 4: Cloud service provider scenario: outage summary and assumptions

- Hypervisor ABC is updated once or twice a year, information publicly available and known to hackers.

- To diversify risk, top-tier providers use multiple versions of the ABC hypervisor; however, the code was obfuscated and made its way into 25-50% of the versions. These were similar enough for the malware to function. There are four major supported versions at the time of writing.

- The malicious code continually crashed the operating systems upon reboot. The system clock could not be changed due to cryptography software dependencies.

- 24 hours to identify backdoor includes the time to bring security experts onsite.

- 24 hours to develop the first patches to remediate the backdoor vulnerability and system crashing

- 24 hours additional time for second tier providers due to fewer resources and the need to hire outside IT and security teams to help investigate, secure, and remediate the problem.

- 6-48 hours ramp-up time (prioritised by CSP region customer impact)

- 1-12 hours additional time for affected companies to bring their systems online after the CSP had restored service.

- Total outage time = 55 hours minimum for sophisticated clients on tier 1 cloud service providers; up to 5 days, 19 hours for less sophisticated organisations on tier 2 cloud service providers.

Impacts

The global economy experiences impacts, with individual companies experiencing severe business interruption and cloud service providers suffering massive reputational damage due to the scale of the event, the repercussions of which last for years.

The primary effects of the cloud service provider outage are experienced by:

– Cloud service providers

– Customers of cloud service providers

– Their customers in the form of service interruptions originating from the impaired servers

– Businesses that lose income and incur additional expenses as a result of the impaired functionality of cloud services

The duration of impairment will be largely dependent upon the ability of the cloud service provider to implement the complete patch and get the systems back up and running, as well as the ability of the end-users to restart their systems and the efficacy of their contingency planning.

Impacts on the economy

To quantify the economic impact that could occur due to incidents of this nature, a stochastic cyber-risk model created by Cyence in 2016 was used.

Primary effects

Cyence estimates the year return period, ground-up losses and 95% confidence intervals (*see Table 4, below*) as well as approximate duration associated with each return period caused by the cloud service power outage described in this section for companies in the US, Canada, the United Kingdom and Europe.

Connection and dependency of each company upon cloud service providers are modelled based on network traffic, redundancy of cloud service and publically available industry information.

Table 4: Return period losses for cloud service provider outage

Sector	% of all businesses analysed (including those that are uninsured)	Return period losses (US Dollars)	
		Large loss	**Extreme loss**
Financial services	10%	$1.29bn	$16.72bn
Software and tech services	4%	$214m	$1.79bn
Hospitality / Retail trade	11%	$332m	$3.08bn
Healthcare	3%	$60m	$853m
Other	72%	$2.70bn	$30.60bn
All industries	**100%**	**$4.60bn**	**$53.05bn**
		95% CI: ($1.60bn-$10.85bn*)*	*95% CI: (*$15.62bn-$121.41bn*)*
Duration		**12-18 hours**	**2.5-3 days**

Losses in Table 4 represent contingent business interruption losses in the form of lost income as well as extra expenses typically covered under a cyber liability insurance policy.

Industries highlighted in the report include financial services, software and tech services, hospitality and retail trade, and healthcare. These have been highlighted as they represent the main sectors purchasing cyber insurance coverage. The overall calculated return period losses have been calculated to represent the full economic costs of the incident to the greater economy and, as such, include all industry sectors.

Confidence intervals have also been included on the "All industry" losses category in order to provide a sense of the variability of the projected loss given the level of data available on the risk.

Losses are expected to vary between sectors based on the nature of their service provision. For example, financial services are expected to see a US$1.29bn loss in a large loss, whereas healthcare organisations will only see a US$60m loss.

This is because companies in industries such as financial services and online retail that are highly reliant upon computer systems to conduct business - more so than other industries – are expected to incur greater losses as their operations and revenue suffers as a result. Healthcare providers may be able to maintain some level of output despite their cloud service provider's outage as the nature of their service provision differs in comparison.

Examples of this are evidenced by past events. In 2012, banks in the US began to see distributed denial of service attacks that interrupted customers' ability to log into their accounts, access money at ATMs and longer sustained outages that could impede their ability to conduct businessin any way, bringing operations to a halt. In comparison, Erie County Medical Center in Buffalo, NY, saw a virus shut down its information technology systems for several days, yet business was able to continue by using paper records.

Organisation size analysis

Table 5: Organisation losses by size for cloud service provider outage

Size (by annual revenue in US Dollars)	% of all Businesses analysed (including those that are uninsured)	Return period losses (US Dollars)	
		Large loss	**Extreme loss**
Small *(Greater than $20M, Less than $100M)*	97.9%	$118m	$2.31bn
Medium *(Greater than $100M, Less than $1B)*	1.8%	$333m	$5.92bn
Large *(Greater than $1B)*	0.3%	$4.15bn	$44.82bn
All industries	**100%**	**$4.60bn** **95% CI: ($1.60bn-$10.85bn)**	**$53.05bn** **95% CI: ($15.62bn-$121.41bn)**

These figures cover the US, Canada, the United Kingdom and Europe

Small organisations make up the majority of businesses in the considered regions by number and represent 97.9% of all businesses included in the analysis. Many of these smaller organisations are reliant on cloud service providers and are likely to lack the formal business continuity planning and testing that can be considered more common place in larger organisations. Although business income losses are smaller in magnitude for these smaller organisations, these organisations may not have the balance sheet liquidity to manage the interruption of cash generation and therefore may see a relatively larger impact.

Company dependency factors

While every industry will generally have a varying degree of reliance on information technology system availability, the true measure of granularity can only be known when considering the environment at an individual company level. Business continuity planning and controls such as redundant systems and alternative procedures aim to reduce the severity of losses by keeping as much of the business operating despite a cloud service provider outage.

For example, an organisation can typically purchase cloud redundancy, or may have a hybrid cloud or separate datacentre available that can be used the moment a cloud service provider goes down. In the case of the Erie County Medical, there was a defined process that could be deployed using paper files. While this process was more time consuming and expensive than the electronic system, it was tested and reliable as a functioning back-up to mitigate the impact to the organisation and patients undergoing care.

Insured losses

For the purposes of assessing the effect of the scenario on the insurance market, it is necessary to model the actual population of organisations that purchase cyber coverage, as well as their coverage limits and retentions.

Insured losses are estimated in Table 6 (*below*):

Table 6: Insured loss

Large loss	Extreme loss
$620m	$8.14bn
95% CI: ($259m-$1.34bn*)*	*95% CI: (*$2.13bn-$18.42bn*)*

Losses in Table 6 (*above*) represent the insured portion of contingent business interruption losses in the form of lost income. Additional potential insured losses could include extra expenses for workarounds, or costs for IT staff overtime and IT consultants to continue business operations.

These values are comprised of the following assumptions:

– **Penetration rates** of cyber insurance have been modelled between 2-15% and policy structures have been modelled with limits 1-3% of annual turnover and retentions of 1% of limits. These estimates are based on public reporting published by Advisen, Marsh, The Council of Insurance Agents and Brokers, as well as industry experts interviewed as part of the report consultation (*see Section 2 to view the methodology behind the project, p12*).

– **Sublimits:** One key element specific to the cloud service provider scenario is that the ensuing business interruption would be considered covered under the contingent business interruption portion of a cyber policy. This portion of a policy is intended to provide business income and extra expenses during a business interruption caused by a third-party supplier. This is not a standardised area of cyber coverage.

Cyence's review of the cyber marketplace uncovered varying coverage offerings – ranging from full limits, to sub-limits, or in some cases no coverage at all.

This report assumes the following sub-limits of coverage:

Table 7: Sublimits by annual revenue

Small (Greater than $20m, Less than $100m)	Medium (Greater than $100m, Less than $1bn)	Large (Greater than $1bn)
20%	30%	50%

Secondary effects

The modelled losses in this scenario only include the direct costs associated with the cloud service provider outage which would be foreseeably covered under cyber liability policies, namely under contingent business interruption coverage which provides affected companies with compensation for additional expenses and lost business operating profits (*see coverage assumptions Section 3 for details, p16*).

There are a number potentially significant secondary effects that sit outside the scope of this report. For the purpose of considering the full range of potential impacts and the development of forward looking exposure management strategies, they can be briefly described as follows:

– Property damage and loss of life

Some life-critical databases or functions may be hosted in the cloud, such as healthcare records, remote surgery or critical SCADA alarm systems. While such systems should have significant redundancy and fall-back planning to avoid such losses, impacts on them caused by cloud service provider outages cannot be ruled out. Outages of such systems could have implications on many insurance policies depending on the nature of the loss (property, Workers Comp, General liability, healthcare professional, medical malpractice, technology E&O, etc.)

– Reputational loss

Affected cloud service providers, and potentially the industry overall, may experience reputational harm and it has become common practice for a company to send apology notes to customers when cloud systems go down *(Amazon, 2011)*. Organisations reliant on cloud service providers may also experience knock-on impacts to their business model. For example, an online news provider experienced impacts when customers didn't get their news, paid advertisers lost their visibility and contracted freelance designers were unable to complete scheduled work *(Nichols, 2017)*. Experiences such as these may slow migrations to the cloud and incentivise businesses to reduce their reliance on cloud services. These shifts could cause stock prices to suffer at cloud service providers and result in securities class action lawsuits.

– Litigation against the cloud service providers affected customers

Companies whose services are heavily reliant upon cloud service providers could potentially be sued by their customers or investors which would result in defence costs and if found negligent, legal damages. There may be a higher risk for such litigation in sectors such as technology services. For example, a company that provides software as a service powered by an affected cloud service provider may see an outage which causes financial harm to its customers.

4.3 Mass vulnerabilities

Security vulnerabilities are pieces of software code that contain an error or weakness that could allow a hacker to compromise the integrity, availability or confidentiality of information accessed by that software.

Once vulnerabilities are identified, malicious actors are able to create exploits that can use the security weakness. Vulnerabilities can occur in many forms with some easier for hackers to exploit than others and with impacts that can range from the inconvenient to the systemic.

As the capabilities of software have grown, code volumes have expanded and development has become more diffuse; vulnerabilities are widely understood to be inevitable.

The more installed code there is, the greater the potential for vulnerability and resultant damages. For example, Microsoft products, which utilise billions of lines of code – some of which is considered outdated by experts – require so many patches the company has instituted what has commonly become known as "patch Tuesday". This is a scheduled monthly update, with further unscheduled patches rolled out on an ad hoc basis as and when issues require quicker fixes. This process is not unique and is seen across a range of large operating systems and applications *(Alhazmi, Malaiya and Ray, 2005)*.

Zero day vulnerabilities

"Zero day" vulnerabilities are a particularly severe sub-set of vulnerabilities that are unknown to a software vendor or the information security community. Zero refers to the amount of time that the security vendor has been aware of the vulnerability to patch it. The moment a zero day vulnerability is discovered by malicious actors, a window of vulnerability begins for attacks exploiting the weakness.

When zero day vulnerabilities are disclosed publicly, malicious actors will act to leverage the vulnerability before companies are able to patch it. Once an effective patch is written and applied, the vulnerability is no longer

called a zero day. If a patch is developed but still leaves a company vulnerable, it can still considered as a zero day as companies and individuals remain vulnerable even if they patch their systems.

It is important to keep in mind that these vulnerabilities and subsequent attacks are rarely discovered right away. It can take anywhere from days to years before a developer is made aware of the vulnerability that led to an attack and software vendors are regularly alerted of potential vulnerabilities from internal or external teams.

Knowledge of a zero day vulnerability creates a delicate balancing act around public disclosure. For example, while public disclosure may help users protect themselves, circulating that knowledge also alerts malicious actors to the weakness. Yet exploitation of the vulnerability may already be happening, and only by disclosing information can remediation and action take place.

Further complicating the issue, much like creating the original code, creating patches can be a complex process. In addition to fixing the vulnerability, patches can introduce new errors into the system. Therefore, software companies need to do extensive testing to ensure robustness of the patch. For example, it is widely known that Google had to issue multiple patches before ultimately rebuilding the core Android Nougat Operating System from the ground up to get at the root of the issue that caused problems with the software *(Brandom, 2016)*.

Responsible disclosure principles

As a result of the challenges and the potentially catastrophic implications of zero day vulnerability exploits, software vendors have begun to develop responsible disclosure principles and policies to alert relevant actors to the need to fix the security vulnerability while minimising the chances of alerting malicious actors to the issue *(Evans et al., 2010)*.

Companies will typically not disclose vulnerabilities publicly before patches are ready unless there is evidence of attacks occurring; however, the standard of

evidence for proof of such attacks varies across companies:

Table 8: Responsible disclosure principles

Organisation	Disclosure deadline
Google	7 - 90 days
Cisco	8 days
CERT	45 days
Yahoo	90 days
Zero Day Initiative	120 days

Companies set tight timelines for the maximum number of days that they will keep vulnerability information confidential before advising the public. Once a patch is created or it becomes apparent that an exploit is being used, disclosure before the timeline will take place.

Sources: Evans, Chris, Hawkes, Ben, Adkins, Heather, Moore, Matt, Zalewski, Michal, Eschelbeck, 2015; CISCO, 2017; Zero Day Initiative, 2017; Rosenblatt, 2016

Many large software providers set standard timelines around public disclosure in order provide clarity around how such circumstances should unfold. This can become contentious when a company with disclosure principles gains knowledge of vulnerability in a third-party and the software provider in question has not yet developed a patch and is determined not to "disclose".

The following section informs the scenario and describes known vulnerabilities that are exploited to cause extreme but plausible losses.

Potential sources of risk

Common sources of risks that have been exploited in known examples include:

− Unscheduled disclosure

− The dark web

− Online communities

Unscheduled disclosure

In addition to software vendors themselves, many parties work to identify zero day vulnerabilities, including security researchers, nation states, security agencies (including their third-party contractors) and organised criminal entities *(Ablon and Bogart, 2017)*. For example, a security researcher named Yang Yu received a US$50k "bounty" from Microsoft in 2016 for identifying the BadTunnel exploit, which was effective against all versions of Windows going back to Windows '95 *(Microsoft, 2017)*. These third-party organisations

generally try to keep a low profile; however, many have become publically known after something went wrong.

One of the most recognisable instances in recent years occurred at Booz Allen Hamilton in 2013, when an employee named Edward Snowden revealed extensive confidential details of the US National Security Agency (NSA) surveillance programs *(Greenberg, 2013)*. While the information disclosed by Snowden did not include specific details on zero day exploits, in 2016 and 2017 there were three incidents that did *(Strohm, 2016)*.

In one case, the Equation Group, an organisation ostensibly working with the NSA Tailored Access Operations, had its data compromised and samples of its vulnerability exploits were posted and offered for auction *(Mimoso, 2016)*.

In another example, a "hacker-for-hire" service out of Italy called the Hacking Team was subject to a large hack revealing three zero day vulnerabilities and more than 400 gigabytes of internal documents and emails were made available online describing the business model of selling them *(Cox, 2017)*.

Just two years after Snowden's disclosure, Booz Allen Hamilton had a second incident when an employee put several terabytes of sensitive information on their personal private home server *(Cameron, 2017)*.

Increasingly, government employees, including ranking government officials, are found to be using personal devices for their work-related activity, potentially putting important sensitive information at risk *(Macri, 2015)*. While it has yet to be disclosed if a breach has occurred in such circumstances, poorly secured networks with highly valuable information are often cited as the sources of public doubt regarding the security of the potentially classified information.

The dark web

The dark web provides a vehicle to communicate confidentially and anonymously. Technically speaking, the dark web is a network of computers that requires special tools to access it. The Tor network[h] is a popular network for accessing dark web data. This network uses a technology called "Onion Routing" – instead of making a direct connection to the destination the connection is made through a proxy network *(Klosowski, 2014)*. The proxy uses a series of intermediate systems to hide the true owner of the message and the message is more resilient to both eavesdropping and traffic analysis.

A simple analogy can be related to conventional mail. A user can address a letter to a final destination and put

[h] See footnote [e], p17 for description of Tor.

this message into an envelope. This envelope will then go into an additional envelope with the destination listed as an intermediate recipient and repeated many times. When the intermediate recipient opens the envelope, they will take the original envelope that was inside and forward it to the next recipient.

The dark web gained public notoriety in October 2013, when an online marketplace for illicit materials called "Silk Road" was taken offline and its creator arrested. At that time, it was estimated to conduct US$22m in sales per year *(Soska and Christin, 2015)*. In November 2014, international law enforcement agencies coordinated the arrest of 17 people to take down over a dozen illicit materials marketplaces that operated in a similar function to "Silk Road" *(Greenberg, 2013)*. Even with these efforts by international law agencies, it is estimated that between US$144- $252m in illegal transactions take place every year (Kruithof et al, 2016).

Communities

In addition to illicit marketplaces, the anonymity of the dark web provides the perfect forums for hacker discussion boards. Forums have been found that are explicitly designed to teach hacking (Vitaris, 2016).

It is challenging to build malware from scratch and malicious actors will often reuse code or services from previous versions for resource efficiency. For example, Stuxnet was a unique form of malware that leveraged four separate zero day vulnerabilities to spread across the internet undetected, enter an "air gapped network" through an infected USB, and render an Iranian nuclear enrichment facility inoperable.

As soon as malware is identified, it can be reverse engineered and taught to others. Sean McGurk, former director of Homeland Security's National Cybersecurity Operations Center, state the identification of Stuxnet malware in the public provided *"a textbook on how to attack" (Hudson, 2012)*.

Examples of known historical vulnerabilities

Box 5: Shellshock

The Shellshock family of vulnerabilities affects Bash, a program that various Unix-based systems use to execute commands given by a remote attacker through the "function export" feature *(InfoSec, 2014)*. The shared scripts are assumed to come from another instance, but the new instance cannot verify it, nor can it verify that the command that it had built is a properly formed script definition. Therefore, an attacker can execute commands on the system (Red Hat, 2016).

Examination of the Bash source code's history shows that Shellshock has been present since version 1.03 of Bash, released in September 1989; however, vulnerability was only disclosed on September 21, 2014 *(InfoSec, 2014)*. Within hours of its initial disclosure, malicious actors started exploiting Shellshock creating botnets of compromised computers used to perform distributed denial-of-service. By 30 September 2014, security firms were tracking approximately 1.5 million attacks and probes per day related to the bug. These botnets have been said to be targeting companies like Akamai Technologies and Yahoo as well as the US Department of Defence *(RedHat, 2014)*. It took an estimated 10 days from the initial discovery to develop the first patches.

Box 6: Joomla

Joomla is a free and open source content management system for publishing web content. It is written in PHP and stores data in several types of SQL databases (MySQL, MS SQL, Postgres). Joomla SQL Injection is associated with more than 200 known common vulnerabilities and exposures *(CVE, 2017)*.

SQL injection is a technique used to attack data-driven applications like Joomla and WordPress. The injection exploits a vulnerability in the software to let the attacker spoof their identity, access or modify existing data or transactions and/or become administrators of the website's backend database server.

In October of 2015, this type of vulnerability was discovered in the core module of Joomla version 3.2, leaving millions of websites used in e-commerce and other sensitive industries vulnerable to SQL injection attacks, data breaches and business interruptions *(Goodin, 2015)*. This issue was catalogued under the name "CVE-2015-7297". It took approximately four days to develop a patch and release the upgrade *(CVE, 2015)*.

Box 7: Heartbleed

The Heartbleed vulnerability affects the OpenSSL cryptography library which is a widely-used implementation of the Transport Layer Security (TLS) protocol. This commonly relied-upon security code was maintained by a small group comprised mainly of volunteers – rather than dedicated security professionals –while the code base continued to grow in use *(McMillan, 2014)*.

The vulnerability results from improper input validation in the implementation of the TLS Heartbeat Extension. The Heartbeat Extension for the TLS protocol provides a way to test and keep secure (encrypted) communication links without needing to renegotiate the connection each time. The vulnerability is classified as a buffer over-read, where more data can be read than should be allowed. The data obtained by a Heartbleed attack can include unencrypted exchanges between TLS parties *(CVE, 2014)*, including form post data in user requests. The data exposed can include session cookies, passwords and other user authentication elements. Attacks can also reveal the private keys of compromised parties, which would enable an attacker to decrypt communications.

The vulnerability was disclosed in April 2014. At that time, approximately 17% of the internet's secure web servers certified by trusted authorities were believed to be vulnerable to the attack *(Mutton, 2014)*. By May 2014, the number of affected websites was down to approximately 1.5% of the top 800,000 websites. It took about seven days for a patch and upgrade to be issued *(Leyden, 2014)*.

Box 8: EternalBlue

On 14 April 2017, ShadowBrokers published a compilation of hacking tools that was allegedly compiled by governmental agency. These tools could give anyone with technical knowledge the ability to exploit certain computers running Windows (*Windows 2000, XP, 7, 8) as well as their server-side variants (server 2000, 2003, 2008, 2008 R2 and 2012), as long as they were connected to the internet *(Khandelwal, 2017)*. While these were categorised as zero days when they were released in mid-April, Microsoft had released patches to many of the identified vulnerabilility exploits in March. Once released, these exploits were rolled into open source exploitation frameworks *(Goodin, 2017)*.

While these are well known recent examples, the observed trend of zero day vulnerabilities being identified is increasing every year. According to RAND's "Zero Days, Thousands of Nights" report, the average zero day exploit lasts 6.9 years in the "wild" *(Ablon and Bogart, 2017)*. Once identified, it is believed to take an average of 22 days to develop an exploit, which contrasts unfavourably with the estimated average of 100- 245 days needed to remediate the vulnerability (Whitehat Security, 2016). Table 9 (*below*), illustrates how this figure varies across sectors:

Table 9: Vulnerability fix time by sector

Sector	Time to fix
IT	248 days
Healthcare	208 days
Retail	205 days
Financial Services	160 days

Source: WhiteHat Security, 2016

There have been countless examples of zero day vulnerabilities that were identified and patches promptly offered; however even if a company is diligent in its patching, the frequency of these events means that the overwhelming likelihood is that malicious actors will make their way into a network if determined to do so. According to Mandiant's M-Trends 2015 report, once a malicious actor has made their way in, the average length of time it takes for a company to realise it has been breached is 205 days *(Mandiant, 2015)*.

4.4 Modelled scenario: mass vulnerability attack

Note: the report does not name particular companies and details should be substituted to explore scenario variants within portfolios.

Background

XYZ Corp is one of a number of third-party contractors that a nation state's national intelligence agency uses for cyber reconnaissance missions. XYZ Corp employs what it markets as "elite hacking teams" globally, with expertise in tailored access and cyber-offensive programs.
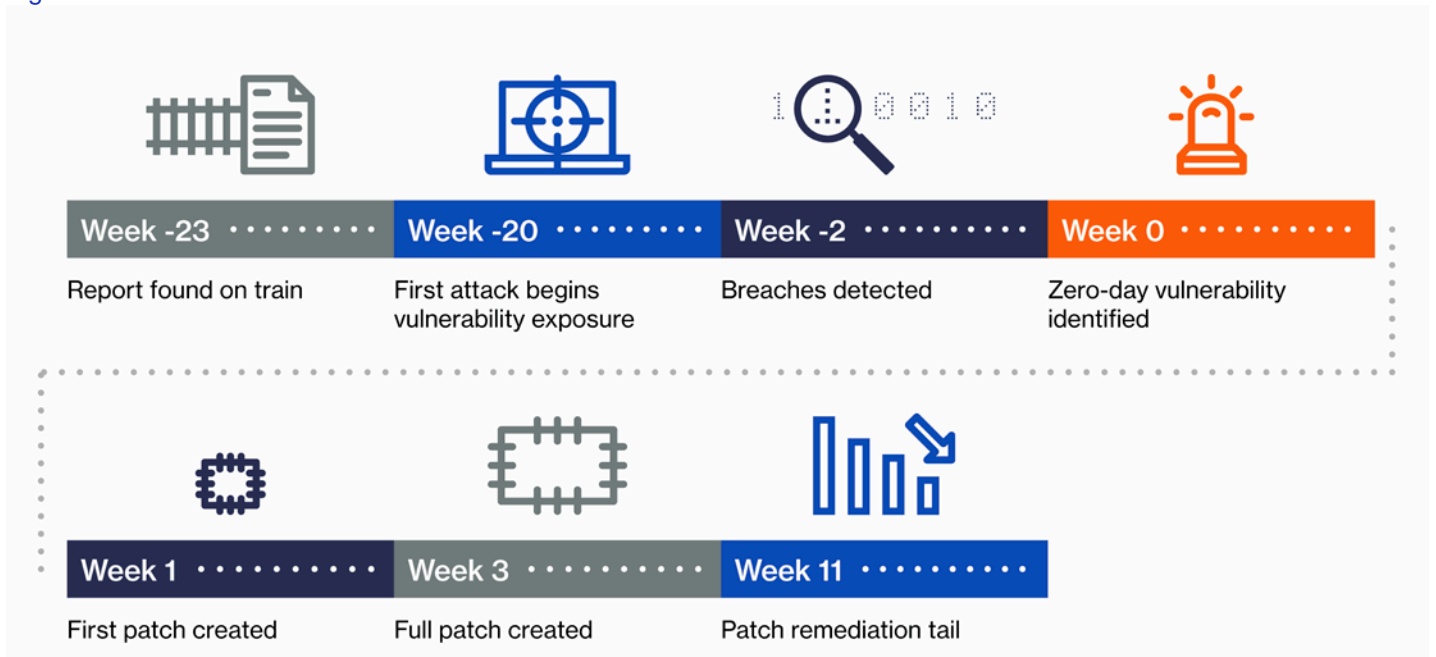
On the way home from work, an analyst accidentally leaves a bag on the train that contains a physical copy of a full report on a recently completed operation.

The vulnerabilities

The report includes details on an identified vulnerability affecting all versions of the DEFG 111 operating system, which is deployed by 45% of the global market. The detailed analysis in the report covers the extensive resources of the organisation, details of the team and individuals who have been working on the reconnaissance mission over the two years, and the specific vulnerability

This information is made available across marketplaces on the dark web at T-23 weeks – approximately six months – on the scenario timeline and is purchased by an undetermined number of unidentified parties. The scenario unfolds as described in Figure 5 (*below*):

Figure 5: Timeline of events

Week -23	Week -20	Week -2	Week 0
Report found on train	First attack begins vulnerability exposure	Breaches detected	Zero-day vulnerability identified

| Week 1	Week 3	Week 11
First patch created	Full patch created	Patch remediation tail

- 1: At t = week -23, the vulnerability report is found on the train and is rapidly uploaded to marketplaces on the dark web.

- 2: At week -20, the first exploit toolkit is developed and the first attacks are launched. This marks the starting point of the window for attacks. At that time, DEFG operating system has approximately 45% market share, which represents the proportion of systems that are vulnerable to this zero day exploit. The proportion of systems on the vulnerable DEFG operating system exploit increases over time as companies transition to it from earlier versions.

- 3: At week -2, organisations begin to detect breaches of their systems.

- 4: Two weeks later, at t = week 0 the vulnerability is identified. At this point some security teams decide to take vulnerable systems offline, leading to a decrease in potential systems to exploit.

- 5: One week later at t = week 1, the first patch is released and many security teams install it. However, the patch still leaves the vulnerability open to some exploits and thus even patched systems remain vulnerable.

- 6: Two weeks later, at t = week 3, a full patch becomes available. The vulnerability is no longer technically a zero day.

- 7: With a full patch available and significant publicity around the incident, security teams implement it quickly to end their vulnerability to the issue for the systems they maintain.

Despite the early patch implementation, there is a long tail towards full remediation. It is important to note that remediation never hits 100% because of poor awareness and patching practices (*see Figure 6, below*). As a result, companies remain vulnerable until the patch is implemented and new systems may come online as companies transition to it from earlier versions of DEFG operating system.

Figure 6: Exposure vulnerability period
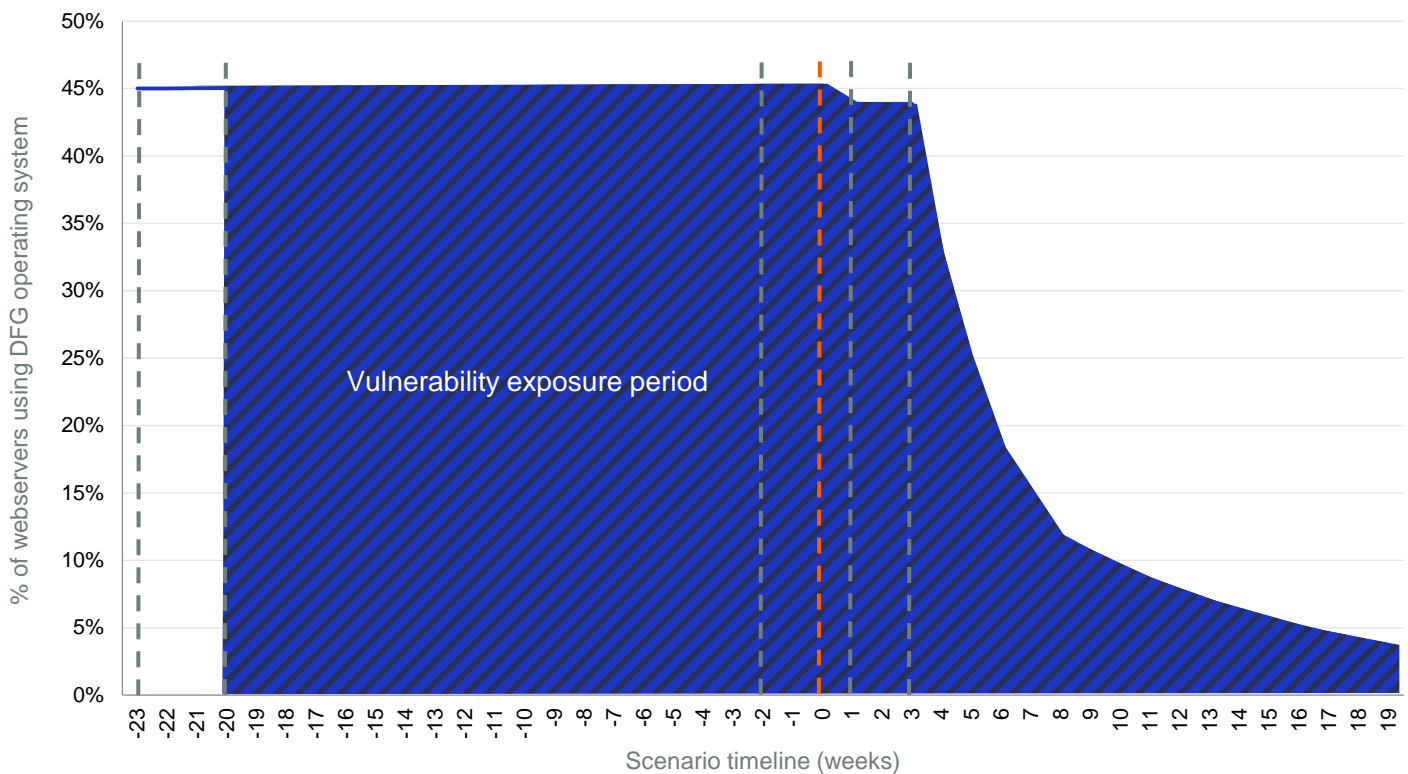
See Figure 5 (p36) for the timeline event details that impact the vulnerability exposure.

Exploitation

Even though the exploit could be scripted targeting specific vulnerabilities, the critical data will be located in different places depending on the network. This means that limited manual effort is still required to parse the data and find the right servers to attack once in a system. By analogy, the exploit allows malicious actors through the back door of a large warehouse, but manual effort is still required to locate the valuable items and extract them.

Using this entry point, various criminal organisations independently develop multiple exploits methodologies and within three weeks begin using them to gain access into corporate networks. The exploits are considered reliable and modular, and are able to be used at scale in custom frameworks by individuals with different motives and resources.

Identification process

Attack frequencies increase against the 45% or so of companies in the market that are running version 111 of DEFG operating system's software as exploits make their way to the dark web and into use by criminal organisations. Due to the nature of early exploits, they are often able to fly under the radar of organisations security teams for a period of time. As malicious actors begin to exfiltrate valuable data, organisations start to identify the breach activity.

Upon identification of the data exfiltration or anomalous network activity, organisations begin to respond by researching the systems with internal employees, or third parties, hiring cyber incident response teams to contain the malware infection and stop the attack. This will depend on the organisations planning and available resources. Because of the nature of the exploits and vulnerability, it required expert security researchers two to three weeks to identify the initial point of entry.

The initial vector is determined to be a vulnerability in DEFG operating system and it is estimated this was introduced by the vendor before the outlined timeline, at least 24 weeks earlier.

Response

This fact is immediately shared through the security community as a requirement of responsible disclosure procedures followed by the teams who make the detection. The software vendor then works with the security community to create patches and some security teams decide to take vulnerable systems offline, leading to a decrease in potential systems to exploit that will either solve the issue or to take vulnerable systems offline before further damages are caused.

− **Patch 1:** Within seven days patches are created and released to high priority organisations, and shortly after this, the patches are posted publicly for wider implementation. However, while monitoring networks using DEFG operating system that have installed the patch, it is identified that 20% of attacks by hackers are still able to leverage the vulnerability in some way to gain system access.

− **Patch 2:** Over the following two weeks, a second, fully tested patch is developed and released. From a company's perspective, remedying the error is a straightforward process once the second patch is ready.

   Depending on reboot times, companies could be back to business as usual within less than an hour. However, not all companies are able to patch their exposed systems immediately.

One reason for a lack of prompt patching is stated as the presence of complex legacy system hierarchies. These circumstances have been seen to increase the difficulty associated with patching because the systems are often built without all the documented dependencies. This means that making changes in one location can have unintended consequences elsewhere.

Another reason for delays in patching is due to a lack of a formal Chief Information Security Officer (CISO) or an equivalent role dedicated to information security within an organisation.

Long-term response

While loss occurrence begins to level off as patches are implemented, the fix is not complete. A capable malicious actor is often able to use a single system penetration to engineer a permanent entry using a range of sophisticated techniques such as immediately installing other backdoors in a network once they have gained entry, or using their initial access to perform reconnaissance that can be used to propagate deeper in to corporate networks.

Due to these techniques, patching a newly detected vulnerability may not remove or prevent future intrusions unless a complete system overhaul is conducted. High severity, customised attacks may continue at a sustained higher rate than those prior to the breach for many years as a result of these techniques.

These attacks may follow in multiple forms, including:

− Pure data exfiltration of sensitive information such as credit cards and health records that have easily transferable economic value

− Theft of sensitive intellectual property such as costly geological estimates of oil reserves, bid negotiation preparations for mergers and acquisitions and financial disclosures that could affect stock prices

− National security information such as the database for security clearance background checks and technical specifications of advanced military technology

Box 9: Mass vulnerability scenario: summary and assumptions

− Vulnerability affects DEFG operating system, allowing for remote execution

− Three weeks to develop the first exploit kit and the first attacks are launched. This begins the window of vulnerability for attacks

− 45% of web servers run on DEFG operating system at the time of the first attack

− 24 weeks before the attacks are identified, during which time the exploits are distributed on the dark web and the operating system slowly gains popularity as users upgrade from an older version

− 2 weeks for white hats to research and identify the exploit vector

− 1 week to develop the first patch, which did not fully fix the vulnerability

− 2 weeks to develop the second patch, which fully fixes the vulnerability

− Remediation path after final patch mirrors that of past zero day patches

Impacts

The cost drivers resulting from this scenario relate to the data breach costs of affected companies.

These include:

− Incident response

− Data forensics

− Notification costs for affected individuals

− Credit monitoring when applicable

These losses begin at week 20 in the scenario timeline and may continue for several years for the following reasons:

1. Slow patching times – more than 3% of companies are still exposed to the vulnerability three months after the second patch is released
2. The exposure period to the vulnerability is 23 weeks – from -20 when the window opens through to -3 the time it takes for the first patch to be developed.

During this time, many companies compromised by the exploits are infected with malware by the malicious actors targeting their system. This may encompass organisations that applied patches promptly, as initial entry into their networks may have provided footing for malicious actors to establish a persistent presence.

This may not have been identified or remediated in the breach response process. The sophistication of these attacks results in many companies experiencing large data breaches anywhere from months to years after the initial vulnerability has been exploited.

Impacts to the economy

To quantify the economic impact that could occur due to incidents of this nature, Cyence used a stochastic cyber risk model it developed in 2016. This enabled an estimation of the ground-up and insured losses for cyber policies based on Monte Carlo simulation.

Primary effects

Cyence estimates the year return period, ground-up losses and 95% confidence intervals associated with each return period caused by the mass vulnerability scenario described in this section for companies in the US, Canada, the United Kingdom and Europe with annual revenue greater than US$20m. Confidence intervals provide a sense of the certainty of the projected loss given the level of data available on the risk.

The dependency of each company upon common operating systems is modelled based on a combination of internet sensor traffic, automated confirmation of software usage and publically available industry information.

Losses in Table 10 (*below*) represent ground-up losses from the mass vulnerability scenario described. Losses include first and third-party data breach costs as well as business interruption that causes lost income.

Table 10: Industry-level losses for mass vulnerability exploitation

Sector	% of population	Return period losses (US Dollars)	
		Large loss	Extreme loss
Financial Services	10%	$2.41bn	$7.37bn
Software and Tech Services	4%	$311m	$784m
Hospitality / Retail Trade	11%	$1.19bn	$2.93bn
Healthcare	3%	$615m	$1.75bn
Other	72%	$5.15bn	$15.89bn
All industries	**100%**	**$9.68bn**	**$28.72bn**
		95% CI: ($4.12bn - $15.63bn)	**95% CI: ($20.50bn - $34.22bn)**

Losses cover the US, CA, the UK and the EU with annual revenue greater than US$20m. Industries highlighted in the report include financial services, software and tech services, hospitality and retail trade, and healthcare. These have been highlighted as they represent the main sectors purchasing cyber insurance coverage. The overall calculated return period losses have been calculated to represent the full economic costs of the incident to the greater economy and as such include all industry sectors.

Confidence intervals have also been included on the "All industry" losses category in order to provide a sense of the variability of the projected loss given the level of data available on the risk.

First-party breach losses
First-party data breach losses include the costs of forensics, notification, ID monitoring, other consultants such as "breach coaches"[i] and the provision of call centres.

Third-party breach losses
Third-party losses include the costs of legal defence for lawsuits (including class actions) and regulatory investigations as well as the resulting fines, liabilities and settlement amounts.

Business interruption losses
Business interruption losses from lost income are included in the loss figures provided. As infected businesses remediate their systems, some may experience interruptions to their operations and suffer lost income.

Additional losses
Additional potential losses include extra expenses for workarounds, or costs for IT staff overtime and IT consultants to continue business operations. However, these are not included in the loss figures described in Table 10 (*see p40*).

Some impacted companies may determine that the costs of the incident response are sufficient to justify litigation against the DEFG operating system vendor for the vulnerability that led to the breach, or to hold liable any third party information security firms that may have had responsibility for patching and network monitoring while the incident occurred.

There are significant protections in the US against such an argument and there are very few circumstances where major damages have been awarded by the courts. However, this action could serve to transfer some portion of the costs of the incident back to the operating system vendor. Either way, this course of action would result in legal costs for DEFG operating system vendor, as well as the affected organisations bringing the lawsuit.

Additional IT resources
In addition, the provider of the DEFG operating system had to invest significant resources into identifying the vulnerability and in creating patches to remediate it.

Breach costs and values
Based on historical events, malicious actors with access to exploits targeting "zero days" can be expected to focus their attacks on organisations where the potential payout for a successful breach may be very high.

[i] A breach coach works with an organisation to isolate the affected data, notify customers, retain necessary forensics professionals and manage crisis communications *(Travelers, 2017)*.

Records
Costs vary widely across industries as a result of their service provision and the value of the data they hold. The black-market value of a record is associated with the potential actions it can be used for. There is more incentive to attack organisations with especially sensitive data and the associated cost of the post-breach remediation is more likely to be expensive. Ponemon estimate the average cost for each lost or stolen record containing sensitive and confidential information at $141 in 2016 *(Ponemon Institute, 2017)*. This ranges from $280 for healthcare records to $101 for public records.

From this, it is reasonable to expect malicious actors to target organisations to maximise their return on investment. All else being equal, healthcare ($280 per record) and financial services ($101 per record) would be more attractive to attackers than retailers ($154 per record).

Target values
A single breached health record was estimated to be worth as much as US$500 in 2015 *(Shahani, 2015)*, whereas a credit card details are thought to be worth US$7-20 per credit card *(SecureWorks, 2016)*. It is important to note that prices for records on the black market are subject to supply and demand constraints, and will vary depending on the number in circulation. For example, in the case of large data breaches flooding the market with new records, credit card record prices have been seen to drop to single dollar figures *(Raj Samani, 2015)*.

Information that would allow a malicious actor to pose as an affected individual for significant financial transactions can have significant real-world implications. This includes collecting tax refunds, taking out loans or receiving health benefits for costly procedures such as surgery. Compare this to the breach of a retailer's credit card system where fraud alerts are quickly triggered based on profiled customer data, and cards are reissued in a timely manner, thereby nullifying the value of the breached data. Hence the value to criminals is much reduced.

Organisational size analysis
The modelled zero day vulnerability provides an entry point into organisations through a virtual door. In the scenario, large organisations experience targeted focus and levels of sophistication. Smaller and middle market-sized companies are also likely to see attacks, as well as using a more generalised form of attack.

This is likely to occur because scripted attacks can be automated to attack the entire population at once. However, given the lower potential value of sensitive records at such organisations, criminal organisations will scale their efforts according to potential financial again and will devote resources accordingly.

Losses are expected to be distributed accordingly:

Table 11: Organisation losses by size for mass vulnerability exploitation

Size (by annual revenue in US Dollars)	% of all businesses analysed (including uninsured)	Return period losses (US Dollars)	
		Large loss	Extreme loss
Small *(Greater than $20m, Less than $100m)*	97.9%	$163m	$683m
Medium *(Greater than $100m, Less than $1bn)*	1.8%	$770m	$3.16bn
Large *(Greater than $1bn)*	0.3%	$8.75bn	$24.88bn
All industries	**100%**	**$9.68bn** **95% CI: ($4.12bn - $15.63bn)**	**$28.72bn** **95% CI: ($20.50bn - $34.22bn)**

Based on companies in the US, Canada, the United Kingdom and the European Union, with annual revenues greater than US$20m

Company security factors

To maintain operational security, organisations must be proactive. The time when a properly configured firewall amounted to a resilient defence has passed and today's network environment is often a porous system that is flexible enough to deal with trends that include:

– Increased business-process outsourcing

– Bring-your-own-device approaches

– Remote employees

– Software as a service

Currently there are no single tools that can prevent an attack, due in part to the active and reactive environment that detects and exploits software developments.

Insured loss

For the purposes of assessing the effect of the scenario on the insurance market, it is necessary to model the actual population of organisations that purchase cyber coverage, as well as their coverage limits and retentions.

Penetration rates for coverage vary by size of business, industry class and the country of domicile for the insured. As such, the insurance industry will only see the portion of the loss that is covered under a valid cyber policy.

Insured losses are estimated in Table 12 (*below*):

Table 12: Insured loss

Large loss (US Dollars)	Extreme loss (US Dollars)
$762mn	$2.07bn
95% CI: ($337m-$1.28bn*)*	*95% CI: (*$1.58bn- $2.44bn*)*

Losses include all modelled primary effects described earlier, which include insurable first and third-party data breach costs such as breach response, liability and lost income from business interruption.

Due to low penetration rates and adequacy of purchased limits, the insurance industry is projected to pay less than 10% of the event's ground-up costs in claims on cyber policies.

Secondary effects

The losses modelled in this scenario include the direct costs associated with the breach such as:

- Breach response
- Extortion payments
- Business interruption
- Legal liability (which may find coverage under typical cyber liability policies)

However, there are many types of secondary losses that organisations may face over time related to the initial breach from the mass vulnerability.

There are a number potentially significant secondary effects that sit outside the scope of this report. For the purpose of considering the full range of potential impacts and the development of forward looking exposure management strategies, they can be briefly described as follows:

Reputational losses
Any firms with noteworthy losses are likely to experience reputational losses as a result of either increased customer "churn" or a decline in trust, both of which may lead customers to avoid doing business with a company going forwards.

Box 10: Reputational loss examples

Target 2013 breach
Target's sales fell by 46% year-on-year in the fourth quarter of 2013 after it disclosed its point-of-sale system data breach during the key holiday sales period *(McGrath, 2014)*.

Reputational losses also have the potential to affect investors. For example, Target's stock price dropped by 10% in the period following its breach. While this can be terminal for SMEs, large established organisations have recovered from these events with minimal impacts visible 12 months after a breach is reported.

Bloomberg Twitter feed hack
Another hacking related "flash crash" was caused in April 2013, when the Syrian Electronic Army hacked Bloomberg's Twitter feed to report that a bomb went off in the White House, which momentarily sent the market plunging 1.5%, representing US$136bn in shareholder value *(Fisher, 2013)*. While the market recovered, stock was not necessarily in the same hands meaning losses for some and gains for others.

Many of the effects of reputational loss can be managed with an effective breach response plan and stand-alone cyber policies may contain features that can help to address these issues in a few key areas: First, breach response services help organisations respond in an effective and organised manner; and second, stand-alone cyber policies often have a coverage sublimit that includes dedicated costs to retain a public relations consultant to manage the messaging of the incident.

Loss of intellectual property
The realisation period for measuring losses due to intellectual property (IP) theft or brand reputation can play out over years. Loss of first-party IP is not currently covered under cyber-liability policies mainly due to concerns around the speculative nature of valuations over time.

A 2014 report by the Center for Responsible Enterprise and Trade and PwC, based on extrapolations from national research and development spending and its associated benefits *(CREATe and PwC, 2014)*, estimated that anywhere from 1-3% of GDP is lost each year in more economically developed countries due to espionage.

Malicious actors engaged in this space can be grouped into the following:

- Nation states
- Malicious insiders
- Competitors
- Transnational organised crime
- Hacktivists
  (Passman, Subramanian and Prokop, 2014)

It is recognised that cyber-enabled hacking and espionage play a significant role in this figure *(Passman, Subramanian and Prokop, 2014)* and an extreme loss mass vulnerability attack may contribute to a further increase in this GDP-loss estimate.

Costs of additional security controls
Many affected companies are likely to respond in the short-term by implementing additional security controls beyond simply patching the specific vulnerability. For example, companies may choose to:

- Hire additional staff to implement more robust patching practices
- Purchase a higher quality security information and event management system to catch data leakage
- Implement application whitelisting[j] or data segmentation

---

[j] Specifying a list of approved software applications that are allowed to be present and active on a system.

Increased identity theft and backlash against digitisation

Increases in zero day vulnerability disclosures and the resulting mass data thefts may leave the public weary of engagement with increasingly digital platforms. These circumstances increase the likelihood of movement towards a more nationalistic view, resulting in increased borders and regulations of cyber space. Some people believe that a series of large hacks could remove trust in the economy, causing governments to impose new regulations and institutions to slow down the pace of technology innovation *(Chinn, Kaplan and Weinberg, 2014)*. The Risk and Responsibility in a Hyperconnected World report *(Chinn, Kaplan and Weinberg, 2014)* estimated that such a trend could forgo as much as US$3 trillion in value creation over the next five to seven years.

Long-tail cyber risk

Cyber is increasingly viewed as a peril that can be covered under many existing insurance policies. "Non-affirmative" cyber is when exposures may exist but the policies don't specifically mention the cyber cover.

There have been many circumstances where coverage was sought, but so far insurance companies have been relatively successful in only providing the coverage that was explicitly included in the policy. However, an extreme loss event may provide circumstances that demonstrate instances where coverage could be found in some policies depending on the outcome of legal rulings. Courts typically adjudicate ambiguity in insurance contracts in favour of the insured.

Potential coverage

The following section provides brief descriptions of types of commercial insurance coverage and a subjective discussion on how cyber "as a peril" losses may find coverage.

- Personal lines identity theft

  Personal lines identity theft insurance can be purchased by individuals as well as through programs offered via group affiliations, such as banks or associations. This coverage may be offered by breached organisations to affected individuals through breach response coverage sections. Due to the frequency of breaches, individuals have monoline identify theft coverage, or have included it as part of a homeowner's policy. The mass vulnerability described in the scenario could lead to many individuals' sensitive records being compromised, and fraud that may trigger such policies.

- Professional liability

  Professional liability (E&O) policies may receive claims in the form of litigation arising from the security breach. These lawsuits may target the operating system provider, technology and security

consultants and breached third-party business providers. If these entities carry professional liability coverage but not cyber liability coverage, the E&O policy may respond to the claim.

It is important to note that many E&O policies have exclusions for cyber losses, especially when the insurance company offers that coverage as a separate purchased endorsement or coverage agreement and are vulnerable to claims arising from the scenario described.

- D&O

  Management liability (D&O) policies may be affected by claims brought against the software vendor and its management team. These claims may be brought as derivative lawsuits against the boards of directors for negligence in management duties or securities class action lawsuits due to a stock-price drop.

- Product liability

  Product liability covering IoT and electronically enabled devices may be impacted by data breaches of a company or a company's corporate network resulting from an initial vulnerability. Software providers have been insulated from such claims in the past as software is not a physical product and proving there was a defect and a reasonable alternative design for a complex piece of software is near impossible. Tort law also pursues contractual remedies, which the software vendors waive. However, these protections are less effective for physical devices with extensive electronic components, which would be the cause of automotive, property and homeowners' losses described above.

- Automotive liability

  The market is already seeing automotive liability coverage adapting to the progression of automation within vehicle systems, as they essentially become computers on wheels. Manufacturers have already had to recall vehicles because of vulnerabilities in the software that can potentially allow the introduction of lethal commands. This was illustrated with the hack of Fiat Chrysler's Jeep in 2015, where an outside party was able to cause both unintended acceleration and the ability to disable the brakes remotely with code.

  Looking forwards, Tesla cars currently receive software updates through the internet on a frequent basis *(Greenberg, 2016)*. This capability creates – or increases – the potential for aggregation of cyber insurance losses. Any accidents caused because of such a loss may trigger an automotive liability policy for resulting bodily injury or property damage (BI/PD).

–   Property

Commercial property policies could be triggered by covered triggers of losses, such as physical damage to buildings and equipment that are caused by cyber-attacks. There are several examples of this, including the German steel mill attack in December 2014 *(BSI, 2014)*.

Box 11: Physical property damage

A report released at the end of 2014 by Germany's Federal Office for Information Security, confirmed physical property damage caused by a suspected cyber attack. The report indicated the attackers gained access to the steel mill through the plant's business network, after which point they were able to pivot their way into production networks to access systems controlling plant equipment.

The malicious actors were then able to manipulate and disrupt control systems to such a degree that a blast furnace could not be properly shut down, finally resulting in what the German Federal Office for Information Security described as "massive" though unspecified physical damage.

Source: Zetter, 2015; BSI, 2014

Conclusion

Emerging Risks Report 2017

Technology
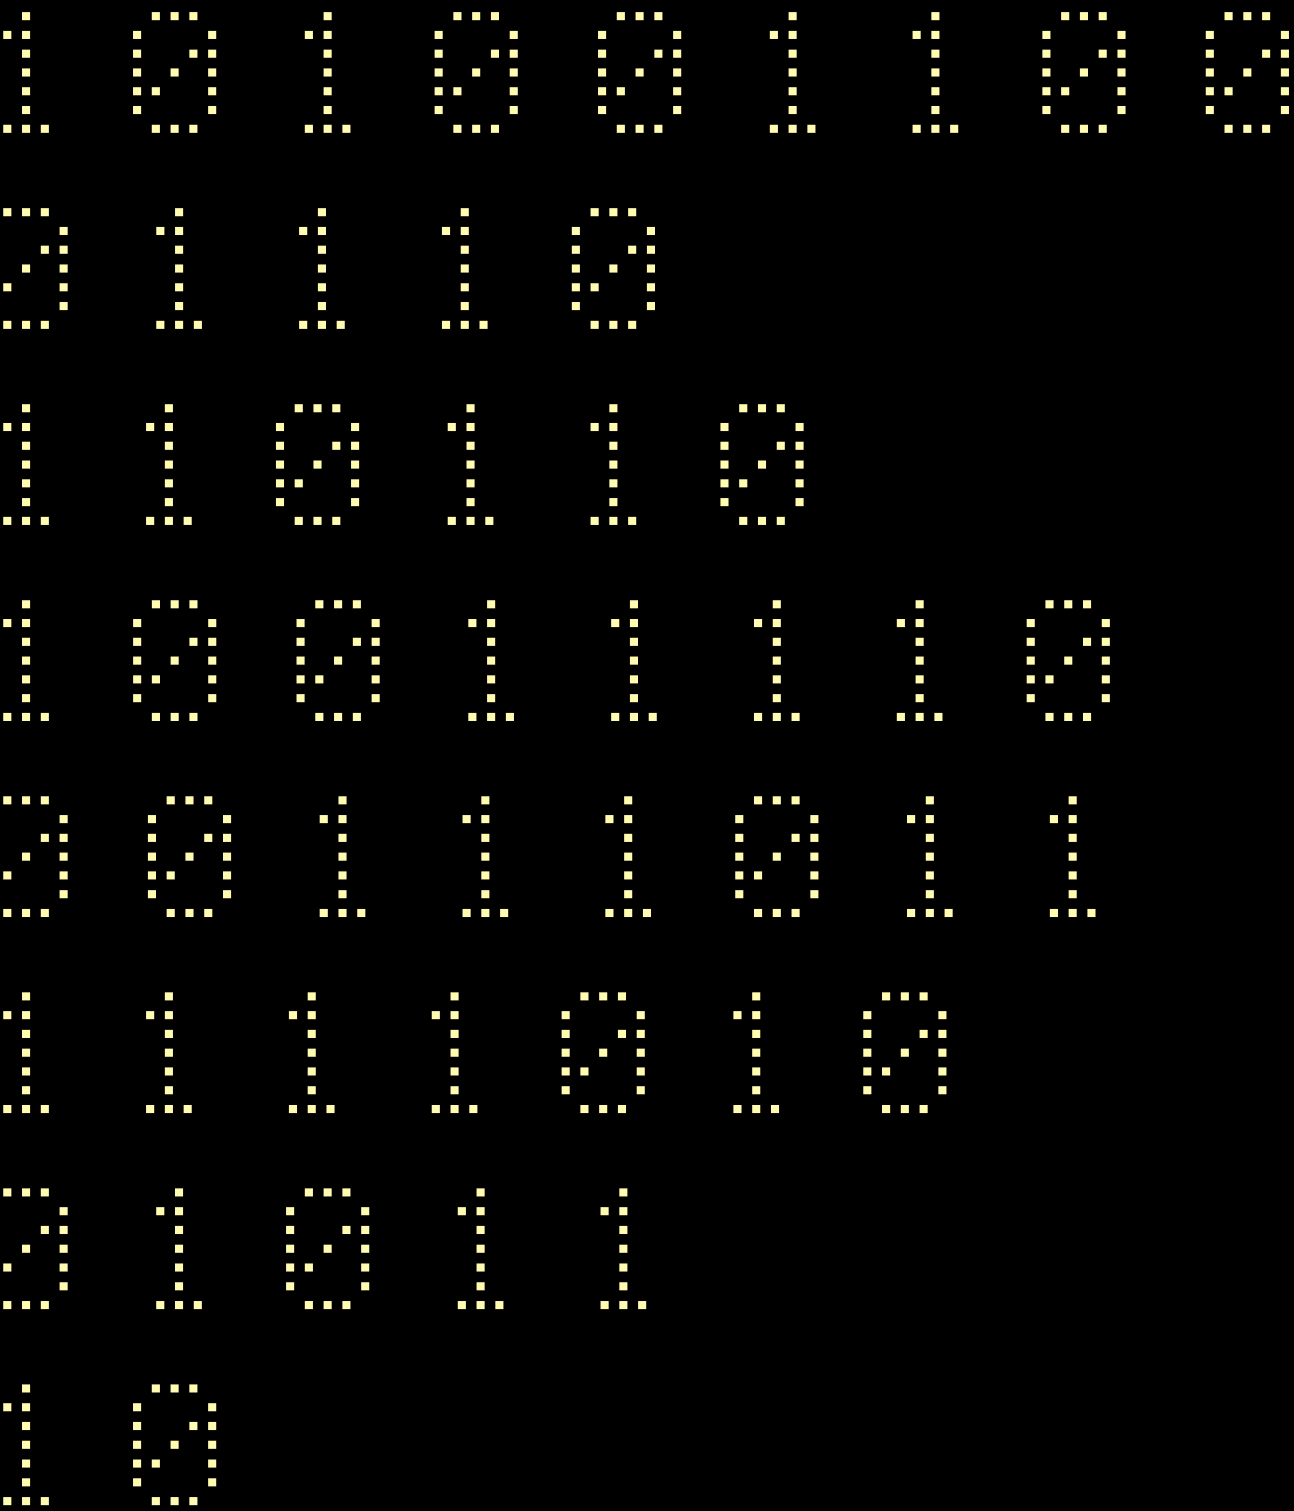
5. Conclusion

This report is designed to increase insurers' and risk managers' understanding of cyber-risk liability and aggregation. It analyses aggregation through the prism of six trends that contribute to digital vulnerability. The two modelled scenarios share many common elements but highlight the wide variety of damages that can occur as a result of cyber-attacks.

Scenario 1: Cloud service provider hack

The modelled scenario demonstrates that there is additional potential for cyber risk aggregation even across cloud service providers who are using common technology – the hypervisor being one such example. This event is the cyber equivalent of a hurricane; it sweeps through many organisations simultaneously and the resulting interruptions cause immediate and potentially severe business impacts.

Scenario 2: Mass vulnerability attack

The mass vulnerability scenario models a zero day vulnerability which makes its way into the hands of malicious actors with the means to exploit it. Due to inconsistencies in patching and a host of operational factors, the remediation of new vulnerabilities tends to vary substantially across organisations. Commonly known and highly publicised vulnerabilities from years ago like Heartbleed and Shellshock are still frequently found on corporate networks despite remediation options. These factors lengthen the tail of events and create uncertainty for insurers managing those liabilities with their risk capital. As attackers gain greater sophistication, these initial entry points will have increasingly significant real world implications.

Key findings

The report makes five important key findings:

– The direct economic impacts of cyber events lead to a wide range of potential economic losses. For the cloud service disruption scenario in the report, these losses range from US$4.6 billion for a large event to US$53 billion for an extreme event; in the mass software vulnerability scenario, the losses range from US$9.7 billion for a large event to US$28.7 billion for an extreme event[k].

– Economic losses could be much lower or higher than the average in the scenarios because of the uncertainty around cyber aggregation. For example, while average losses in the cloud service disruption scenario are US$53.1 billion for an extreme event, they could be as high as US$121.4 billion or as low as US$15.6 billion[l], depending on factors such as the different organisations involved and how long the cloud-service disruption lasts for.

– Cyber-attacks have the potential to trigger billions of dollars of insured losses. For example, in the cloud-services scenario insured losses range from US$620 million for a large loss to US$8.1 billion for an extreme loss. For the mass software vulnerability scenario, the insured losses range from US$762 million (large loss) to US$2.1 billion (extreme loss).

– The scenarios show there is an insurance gap of between US$4 billion (large loss) and $45 billion (extreme loss) in terms of the cloud services scenario – meaning that between 13% and 17% of the losses are covered, respectively. The underinsurance gap is between US$9 billion (large loss) and $26 billion (extreme loss) for the mass vulnerability scenario meaning that just 7% of economic losses are covered.

The "Top 10 world's costliest natural catastrophes by insured losses, 1980-2016" *(MunichRe, 2017)*, saw an average of 30% covered by insurance. Compare this to the projected coverage for the modelled cyber incidents and it is apparent there is opportunity for market development.

Table 13: Estimated coverage for the modelled scenarios (US Dollars)

Event	Overall losses		Insured losses		% loss covered	
	Large loss	Extreme loss	Large loss	Extreme loss	Large loss	Extreme loss
Cyber CSP interruption	$4.60bn	$53.05bn	$620m	$8.14bn	13%	17%
Cyber mass vulnerability	$9.68bn	$28.72bn	$762m	$2.07bn	7%	7%

– When assessing current estimated market premiums against the forecasted cyber scenario insurance loss estimates set out in the report, it is apparent that a single cyber event has the potential to increase industry loss ratios by 19% and 250% for large and extreme loss events, respectively. This illustrates the catastrophe potential of the cyber-risk class.

Table 14: Industry loss ratios

Current market premium (US Dollars) $3,250bn

Scenario	Large loss	Extreme loss
Cyber CSP interruption	$620m	$8.14bn
Loss ratio	19%	250%
Cyber Mass vulnerability	$762m	$2.07bn
Loss ratio	23%	64%

These modelled loss ratios are associated with a single catastrophic event. Note that these are additive to the existing loss ratios that would have occurred if the catastrophe hadn't occurred.

---

[k] These figures represent the mean values of simulated loss year severities for large and extreme loss events, and take into account all expected direct expenses related to the events. Impacts such as property damage, bodily injury, as well as indirect losses such as the loss of customers and reputational damage are not taken in to account.

[l] These are illustrated as 95% confidence ranges – the range of values that act as good estimates to cover known and unknown parameters.

The aggregation potential of the losses from these scenarios shows that cyber risks should be considered as cat exposed classes, and that economic loss events have the potential to be as large as a major hurricane. In property classes that are exposed to aggregating risks it is typical to include catastrophe loading in technical premium calculations and capital models *(Kreps, 1990, 1993)*, and this or similar approaches may be appropriate for cyber business going forward, especially as insurance penetration rates rise.

.

Next steps

The purpose of this report is to assist insurance markets writing cyber coverage to understand potential aggregation paths, and enable better capital management and risk understanding. Cyber liability is still at an early stage compared to other coverage lines and deeper understanding of exposures will help the market move towards more expansive coverage and set adequate limits that meet the insurance needs presented by cyber risk.

This report's findings suggest economic losses from cyber events have the potential to be as large as those caused by major hurricanes. Insurers could benefit from thinking about cyber cover in these terms and make explicit allowance for aggregating cyber-related catastrophes. To achieve this, data collection and quality is important, especially as cyber risks are constantly changing.

Traditional insurance risk modelling relies on authoritative information sources such as national or industry data, but there are no equivalent sources for cyber-risk and the data for modelling accumulations must be collected at scale from the internet. This makes data collection, and the regular update of it, key components of building a better understanding of the evolving risk.

For the insurance industry to capitalise on the growing cyber market, insurers would benefit from a deeper understanding of the potential tail risk implicit in cyber coverage.

Risk managers could use the cyber-attack scenarios to see what impacts cyber-attacks might have on their core business processes, and plan what actions they could take to mitigate these risks.

References

A.M Best. 2016. AM Best Special Report: AM Best Increases Estimate for Net Ultimate US Asbestos Losses to $100 Billion [online]. A.M Best. Available at:
http://www3.ambest.com/ambv/bestnews/presscontent.aspx?refnum=24668&altsrc=23

Ablon, L. and Bogart, A. 2017. Zero Days, Thousands of Nights: The Life and Times of Zero-Day Vulnerabilities and Their Exploits [online]. Available at: https://www.rand.org/pubs/research_reports/RR1751.html

AGC, V. 2017. Virtual Apollo Guidance Computer software [online]. Available at: https://github.com/virtualagc/virtualagc

Alhazmi, O., Malaiya, Y. and Ray, I. 2005. Security vulnerabilities in software systems: A quantitative perspective Lecture Notes in Computer Science, [online] 3654, pp.281–294. Available at:
http://www.scopus.com/inward/record.url?eid=2-s2.0-26444516466&partnerID=40&md5=292087b19973814d509854a20283200c

Amazon Web Services. 2017. Summary of the Amazon S3 Service Disruption in the Northern Virginia (US-EAST-1) Region [online]. Available at: https://aws.amazon.com/message/41926/

Amazon. 2011. Summary of the Amazon EC2 and Amazon RDS Service Disruption in the US East Region [online]. Amazon. Available at: https://aws.amazon.com/message/65648/

Aon Benfield. 2017. Cyber Update: 2016 cyber insurance profits and performance [online]. Available at:
http://thoughtleadership.aonbenfield.com/Documents/20170504-ab-cyber-naic-supplemental-study.pdf

Atwood, J. 2007. The Best Code is No Code At All [online]. Coding horror. Available at:
https://blog.codinghorror.com/the-best-code-is-no-code-at-all/

Brandom, R. 2016. Google rebuilt a core part of Android to kill the Stagefright vulnerability for good [online]. The Verge. Available at: https://www.theverge.com/2016/9/6/12816386/android-nougat-stagefright-security-update-mediaserver

BSI. 2014. BSI: Bericht zur Lage der IT-Sicherheit in Deutschland 2014 [online]. Bundesamt für Sicherheit in der Informationstechnik. Available at: https://www.wired.com/wp-content/uploads/2015/01/Lagebericht2014.pdf

Cameron, D. 2017. Top Defense Contractor Left Sensitive Pentagon Files on Amazon Server With No Password [Updated] [online]. Gizmodo. Available at: http://gizmodo.com/top-defense-contractor-left-sensitive-pentagon-files-on-1795669632

Checkoway, S., Cohney, S., Garman Matthew Green, C., Heninger, N., Maskiewicz, J., Rescorla, E., Shacham, H. and Weinmann San Diego, R.-P.U. 2016. A Systematic Analysis of the Juniper Dual EC Incident. [online]. Available at:
http://dualec.org/DualECJuniper-draft.pdf

Chelf, B. 2009. In search of perfect code [online]. The Chartered Institute for IT (BCS). Available at:
http://www.bcs.org/content/conWebDoc/22549

Chinn, D., Kaplan, J. and Weinberg, A. 2014. Risk and responsibility in a hyperconnected world: Implications for enterprises. [online]. World Economic Forum In collaboration with McKinsey & Company, Available at: http://www.mckinsey.com/~/media/mckinsey/business functions/mckinsey digital/our insights/risk and responsibility in a hyperconnected world implications for enterprises/risk and responsibility in a hyperconnected world.ashx

Chuvakin, A. 2016. Our 'Understanding Insider Threats' Paper Publishes [online]. Gartner. Available at: http://blogs.gartner.com/anton-chuvakin/2016/05/09/our-understanding-insider-threats-paper-publishes/

Cid, D. 2014. MailPoet Vulnerability Exploited in the Wild - Breaking Thousands of WordPress Sites [online]. SucuriBlog. Available at: https://blog.sucuri.net/2014/07/mailpoet-vulnerability-exploited-in-the-wild-breaking-thousands-of-wordpress-sites.html

CISCO. 2017. Security Vulnerability Policy [online]. Cisco. Available at: http://www.cisco.com/c/en/us/about/security-center/security-vulnerability-policy.html#rsvifc

Corbet, J., Kroah-Hartman, G. and McPherson, A. 2012. Linux Kernel Development How Fast it is Going, Who is Doing It, What They are Doing and Who is Sponsoring It [online]. The Linux Foundation. Available at: https://www.linux.com/publications/linux-kernel-development-how-fast-it-going-who-doing-it-what-they-are-doing-and-who-5

Cox, J. 2017. DEA Used Malware Without Laying Out Privacy Risks - Motherboard [online]. Vice Motherboard. Available at: https://motherboard.vice.com/en_us/article/dea-used-malware-without-laying-out-privacy-risks

CREATe. and PwC. 2014. Economic Impact of Trade Secret Theft [online]. The Center for Responsible Enterprise And Trade (CREATe.org) and PricewaterhouseCoopers LLP. Available at: https://create.org/resource/economic-impact-of-trade-secret-theft/

CVE. 2014. CVE-2014-0160 [online]. MITRE Corporation. Available at: https://cve.mitre.org/cgi-bin/cvename.cgi?name=cve-2014-0160

CVE. 2015. CVE-2015-7297 [online]. MITRE Corporation. Available at: https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2015-7297

CVE. 2017. Joomla!: Security Vulnerabilities [online]. MITRE Corporation. Available at: https://www.cvedetails.com/vulnerability-list/vendor_id-3496/product_id-16499/Joomla-Joomla-.html

Davis, D. 2013. What Is a Type 1 Hypervisor? [online]. VirtualizationSoftware. Available at: http://www.virtualizationsoftware.com/type-1-hypervisors/

Delange, J., Nichols, W., Mchale, J., Hudak, J. and Nam, M.-Y. 2015. Evaluating and Mitigating the Impact of Complexity in Software Models Software Engineering Institute, [online] (December). Available at: https://resources.sei.cmu.edu/asset_files/TechnicalReport/2015_005_001_448093.pdf

Department of the Treasury. 2016. Guidance Concerning Stand-Alone Cyber Liability Insurance Policies Under the Terrorism Risk Insurance Program [online]. Available at: https://www.federalregister.gov/documents/2016/12/27/2016-31244/guidance-concerning-stand-alone-cyber-liability-insurance-policies-under-the-terrorism-risk

Evans, C., Grosse, E., Mehta, N., Moore, M., Ormandy, T., Tinnes, J., Zalewski, M. and Team, G.S. 2010. Rebooting responsible disclosure: a focus on protecting end users [online]. Google Security Blog. Available at: https://security.googleblog.com/2010/07/rebooting-responsible-disclosure-focus.html

Evans, Chris, Hawkes, Ben, Adkins, Heather, Moore, Matt, Zalewski, Michal, Eschelbeck, G. 2015. Feedback and data-driven updates to Google's disclosure policy [online]. Google Project Zero. Available at: https://googleprojectzero.blogspot.com/2015/02/feedback-and-data-driven-updates-to.html

Festa, P. 1998. Windows "back door" raises flags [online]. CNET. Available at: https://www.cnet.com/news/windows-back-door-raises-flags/

Fisher, M. 2013. Syrian hackers claim AP hack that tipped stock market by $136 billion Is it terrorism? [online]. The Washington Post. Available at: https://www.washingtonpost.com/news/worldviews/wp/2013/04/23/syrian-hackers-claim-ap-hack-that-tipped-stock-market-by-136-billion-is-it-terrorism/?utm_term=.921471be86e6

FOLDOC. 2002. Race Condition [online]. Free Online Dictionary of Computing. Available at: http://foldoc.org/race condition

Forbes: Great Speculations. 2016. Amazon Continues To Gain Share In Cloud Infrastructure Services Market [online]. Forbes. Available at: https://www.forbes.com/sites/greatspeculations/2016/08/17/amazon-continues-to-gain-share-in-cloud-infrastructure-services-market/#572554ae15b8

Friedman, S. and Thomas, A. 2017. Demystifying cybersecurity insurance [online]. Deloitte. Available at: https://dupress.deloitte.com/dup-us-en/industry/financial-services/demystifying-cybersecurity-insurance.html

Gelles, D., Tabuchi, H. and Dolan, M. 2015. Complex Car Software Becomes the Weak Spot Under the Hood [online]. The New York Times. Available at: https://www.nytimes.com/2015/09/27/business/complex-car-software-becomes-the-weak-spot-under-the-hood.html?_r=0

General Data Protection Regulation. 2014. Official Journal of the European Union.

Goodin, D. 2015. Joomla bug puts millions of websites at risk of remote takeover hacks [online]. Ars Technica. Available at: https://arstechnica.com/security/2015/10/joomla-bug-puts-millions-of-websites-at-risk-of-remote-takeover-hacks/

Goodin, D. 2017. Risk assessment - Fearing Shadow Brokers leak, NSA reported critical flaw to Microsoft [online]. ArsTechnica. Available at: https://arstechnica.com/security/2017/05/fearing-shadow-brokers-leak-nsa-reported-critical-flaw-to-microsoft/

Graham, L. 2017. Cybercrime costs the global economy $450 billion [online]. CNBC Cyber Security. Available at: http://www.cnbc.com/2017/02/07/cybercrime-costs-the-global-economy-450-billion-ceo.html

Greenberg, A. 2013. NSA Contractor Booz Allen Hamilton Rushes To Distance Itself From Staffer Who Leaked Top Secret Docs [online]. Forbes. Available at: https://www.forbes.com/sites/andygreenberg/2013/06/09/nsa-contractor-booz-allen-hamilton-rushes-to-distance-itself-from-staffer-who-leaked-top-secret-docs/#7126c2855799

Greenberg, A. 2016. Tesla Responds to Chinese Hack With a Major Security Upgrade [online]. WIRED. Available at: https://www.wired.com/2016/09/tesla-responds-chinese-hack-major-security-upgrade/

Hudson, J. 2012. Here's How the Stuxnet Virus Could Be Used Against the US [online]. The Atlantic. Available at: https://www.theatlantic.com/technology/archive/2012/06/heres-how-stuxnet-virus-could-be-used-against-us/327389/

InfoSec. 2014. Exploiting and Verifying Shellshock: CVE-2014-6271 [online]. InfoSec Institute. Available at: http://resources.infosecinstitute.com/bash-bug-cve-2014-6271-critical-vulnerability-scaring-internet/#gref

Insurance Services Office. 2016. ISO Cyber Risk Solutions ISO's Cyber Insurance Program. [online]. Available at: www.verisk.com/downloads/iso-cyber-insurance-program.pdf

Johnson, P. 2012. Curiosity about lines of code [online]. IT World. Available at: http://www.itworld.com/article/2725085/big-data/curiosity-about-lines-of-code.html

Jones, C. 2006. The Economics of Software Maintenance in the Twenty First Century [online] Available at: www.compaid.com/caiinternet/ezine/capersjones-maintenance.pdf

Kalinich, K. 2017. US Treasury Makes Standalone Cyber Insurance Policies More Valuable. [online]. Available at: http://www.aon.com/attachments/risk-services/cyber/TRIA-2017Update.pdf

Keizer, G. 2015. XcodeGhost used unprecedented infection strategy against Apple [online]. Computerworld. Available at: http://www.computerworld.com/article/2986768/application-development/xcodeghost-used-unprecedented-infection-strategy-against-apple.html

Khandelwal, S. 2017. Turns Out Microsoft Has Already Patched Exploits Leaked By Shadow Brokers [online]. The Hacker News. Available at: http://thehackernews.com/2017/04/window-zero-day-patch.html

Klosowski, T. 2014. What Is Tor and Should I Use It? [online]. Available at: https://lifehacker.com/what-is-tor-and-should-i-use-it-1527891029

Kreps, R. 1990. Reinsurer risk loads from marginal surplus requirements In: PCAS LXXX. [online] pp.196–203. Available at: https://www.casact.org/pubs/proceed/proceed90/90196.pdf

Kreps, R. 1993. Reinsurer risk loads from marginal surplus requirements Insurance: Mathematics and Economics, [online] 12(1), p.73. Available at: https://doi.org/10.1016/0167-6687(93)91028-S

Kurth, L. 2016. Xen Project Blog, Introducing the Xen Project Code Review Dashboard [online]. Xen Project Community, A Linux Foundation Collaborative Project. Available at: https://blog.xenproject.org/tag/code-review-dashboard/

Laux, J. and Kerman, C. 2017. Cyber Update: 2016 Cyber Insurance Profits and Performance. [online]. Available at: http://thoughtleadership.aonbenfield.com/Documents/20170504-ab-cyber-naic-supplemental-study.pdf

Levine, J. 2012. Mars miracle The Dryden X-Press, [online] 54(9), Aug., p.8. Available at: https://www.nasa.gov/sites/default/files/files/08_03_12.pdf

Leyden, J. 2014. Thousands of websites still spilling their crypto blood on carpets everywhere [online]. The Register. Available at: https://m.theregister.co.uk/2014/05/20/heartbleed_still_prevalent/

Linux Foundation. 2017. Xen Project Source Repositories [online]. Linux Foundation. Available at: https://xenbits.xen.org/

Macri, G. 2015. Hillary's Not the Only Government Employee Using a Personal Device for Work [online]. InsideSources. Available at: http://www.insidesources.com/hillarys-not-the-only-government-employee-using-a-personal-device-for-work/

Mandiant. 2015. M-Trends 2015: A view from the front lines. [online]. Available at: https://www2.fireeye.com/rs/fireeye/images/rpt-m-trends-2015.pdf

Massman, S. 2001. War Risk Exclusion Legal History Outlined [online]. PropertyCasualty360. Available at: http://www.propertycasualty360.com/2001/10/01/war-risk-exclusion-legal-history-outlined?&slreturn=1499559037

McAfee. 2017. Building Trust in a Cloudy Sky. [online]. Available at: https://prod2.secureforms.mcafee.com/verify?docID=9638eaf9ef0e1ac769dbc96af0c38fc2&tag=csr&region=us&eid=17SC_CSGLQ1_WP_WW&elqCampaignId=12696

McConnell, S. 2004. Code Complete: A Practical Handbook of Software Construction. 2nd ed.

McGrath, M. 2014. Target Profit Falls 46% On Credit Card Breach And The Hits Could Keep On Coming [online]. Forbes. Available at: https://www.forbes.com/sites/maggiemcgrath/2014/02/26/target-profit-falls-46-on-credit-card-breach-and-says-the-hits-could-keep-on-coming/#38ea17d97326

McMillan, R. 2014. How Heartbleed Broke the Internet — And Why It Can Happen Again [online]. WIRED. Available at: https://www.wired.com/2014/04/heartbleedslesson/

Microsoft. 2017. Bounty Hunters: The Honor Roll [online]. Microsoft Security Response Center. Available at: https://technet.microsoft.com/en-us/security/dn469163.aspx

Mimoso, M. 2016. Shadowbrokers leak has 'strong connection' to Equation Group [online]. Kapersky. Available at: https://threatpost.com/shadowbrokers-leak-has-strong-connection-to-equation-group/119941/

MunichRe. 2017. NatCatSERVICE Annual Statistics. Available at: https://www.munichre.com/touch/naturalhazards/en/natcatservice/annual-statistics/index.html

Mutton, P. 2014. Half a million widely trusted websites vulnerable to Heartbleed bug | Netcraft [online]. Netcraft. Available at: https://news.netcraft.com/archives/2014/04/08/half-a-million-widely-trusted-websites-vulnerable-to-heartbleed-bug.html

Mutton, P. 2015. Millions still running the risk with Windows Server 2003 [online]. Netcraft. Available at: https://news.netcraft.com/archives/2015/08/12/millions-still-running-the-risk-with-windows-server-2003.html

NCSL. 2017. Security Breach Notification Laws [online]. National Conference of State Legislatures. Available at: http://www.ncsl.org/research/telecommunications-and-information-technology/security-breach-notification-laws.aspx

Nichols, S. 2017. AWS's S3 outage was so bad Amazon couldn't get into its own dashboard to warn the world [online]. The Register. Available at: https://www.theregister.co.uk/2017/03/01/aws_s3_outage/

Noyes, K. 2010. 10 Reasons Open Source Is Good for Business [online]. PCWorld. Available at: http://www.pcworld.com/article/209891/10_reasons_open_source_is_good_for_business.html

OECD. 2012. OECD Internet Economy Outlook 2012. Organisation for Economic Co-operation and Development.

Open Source Initiative. 2017. Community & Collaboration [online]. Open Source Initiative. Available at: https://opensource.org/community

Oracle. 2014. Concepts Guide for Release 33 - 13 What are Hypervisors? [online]. Oracle. Available at: https://docs.oracle.com/cd/E50245_01/E50249/html/vmcon-hypervisor.html

Özkan, S. 2017a. VMware » Esxi: Vulnerability Statistics [online]. CVEdetails. Available at: http://www.cvedetails.com/product/14180/?q=Esxi

Özkan, S. 2017b. XEN : Vulnerability Statistics [online]. CVEdetails. Available at: https://www.cvedetails.com/vendor/6276/XEN.html

Passman, P., Subramanian, S. and Prokop, G. 2014. Economic Impact of Trade Secret Theft: A framework for companies to safeguard trade secrets and mitigate potential threats. [online]. Available at: https://www.pwc.com/us/en/forensic-services/publications/assets/economic-impact.pdf

Ponemon Institute. 2017. 2017 Cost of Data Breach Study: Global Overview [online]. Available at: https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=SEL03130WWEN

Poulsen, K. 2003. Thwarted Linux backdoor hints at smarter hacks [online]. SecurityFocus. Available at: http://www.securityfocus.com/news/7388

PwC. 2015. Insurance 2020 & beyond: Reaping the dividends of cyber resilience [online]. Available at: http://www.pwc.com/gx/en/insurance/publications/assets/reaping-dividends-cyber-resilience.pdf

Raj Samani. 2015. What Is Your Customer Data Worth? [online]. McAfee Blogs. Available at: https://securingtomorrow.mcafee.com/executive-perspectives/customer-data-worth/

Ralph, O. 2017. Cyber insurance market expected to grow after WannaCry attack [online]. Financial Times (online). Available at: https://www.ft.com/content/25bf97e8-3a27-11e7-821a-6027b8a20f23

RedHat. 2014. Frequently Asked Questions about the Shellshock Bash flaws - Red Hat Customer Portal [online]. RedHat. Available at: https://access.redhat.com/blogs/766093/posts/1976393

Rosenblatt, S. 2016. When to disclose a zero-day vulnerability [online]. The Parallax. Available at: https://www.the-parallax.com/2016/12/06/disclose-zero-day-vulnerability/

Rowley, J. 2017. Open Source Software Is Big Business With Big Funding [online]. Crunchbase. Available at: http://about.crunchbase.com/news/open-source-software-big-business-big-funding/

Ruest, D. 2010. Virtualization hypervisor comparison: Type 1 vs Type 2 hypervisors [online]. TechTarget. Available at: http://searchservervirtualization.techtarget.com/tip/Virtualization-hypervisor-comparison-Type-1-vs-Type-2-hypervisors

SecureWorks. 2016. Underground Hacker Markets Annual Report [online]. Available at: https://www.secureworks.co.uk/resources/rp-2016-underground-hacker-marketplace-report

SecurityWeek News. 2016. Backdoor in WordPress Plugin Steals Admin Credentials [online]. SecurityWeek News. Available at: http://www.securityweek.com/backdoor-wordpress-plugin-steals-admin-credentials

Shahani, A. 2015. The Black Market For Stolen Health Care Data [online]. NPR. Available at: http://www.npr.org/sections/alltechconsidered/2015/02/13/385901377/the-black-market-for-stolen-health-care-data

Shankland, S. 2001. Borland InterBase backdoor detected [online]. ZDNet. Available at: http://www.zdnet.com/article/borland-interbase-backdoor-detected/

Soska, K. and Christin, N. 2015. Measuring the Longitudinal Evolution of the Online Anonymous Marketplace Ecosystem 24th USENIX Security Symposium (USENIX Security 15), [online] pp.33–48. Available at: https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/soska

Stanley, C. 2017. Cyber market estimate (interview 26 June, Christian Stanley, Casualty Executive, Class of Business Underwriting Performance, Lloyd's).

Steinberg, J. 2015. Security Vulnerability Discovered In Millions Of Business Computer Systems -- Here's What You Need To Know [online]. Forbes. Available at: https://www.forbes.com/sites/josephsteinberg/2015/05/13/major-vulnerability-discovered-in-millions-of-business-computer-systems-heres-what-you-need-to-do/#bfedd7c74efe

Stevens, L. 2017. Amazon Finds the Cause of Its AWS Outage: A Typo [online]. The Wall Street Journal. Available at: https://www.wsj.com/articles/amazon-finds-the-cause-of-its-aws-outage-a-typo-1488490506

Strohm, C. 2016. Another NSA Breach Hits Booz Allen Will Anything Change? - Bloomberg [online]. Bloomberg Markets. Available at: https://www.bloomberg.com/news/articles/2016-10-07/another-breach-at-nsa-involves-booz-allen-will-anything-change

The New York Times. 2013. Close the NSA's Back Doors [online]. The New York Times. Available at: http://www.nytimes.com/2013/09/22/opinion/sunday/close-the-nsas-back-doors.html

Travelers. 2017. What is a Data Breach Coach and How Do I Get One? [online]. Travelers. Available at: https://www.travelers.com/resources/cyber-security/what-is-a-data-breach-coach.aspx

Tufekci, Z. 2017. The World Is Getting Hacked Why Don't We Do More to Stop It? [online]. The New York Times. Available at: https://www.nytimes.com/2017/05/13/opinion/the-world-is-getting-hacked-why-dont-we-do-more-to-stop-it.html?_r=2

Vaughan-Nichols, S. 2014. Heartbleed: Open source's worst hour [online]. ZDNet. Available at: http://www.zdnet.com/article/heartbleed-open-sources-worst-hour/

VMWare. 2017. VMware Security Advisories (VMSAs) [online]. VMWare. Available at: https://www.vmware.com/security/advisories.html

Wander, S. 2007. Powerless National Aeronautics and Space Administration - System Failure Case Studies, 1(10), pp.1–4.

Weisbart, S.N. 2017. Commentary on 2016 year-end results. [online]. Available at: http://www.iii.org/sites/default/files/docs/pdf/2016_fullyear_commentary_050417-2.pdf

WhiteHat Security. 2016. Web applications security statistics report 2016 [online]. Available at: https://www.whitehatsec.com/info/website-stats-report-2016-wp/

WikiLeaks. 2017. Vault 7: Projects - Marble Framework [online]. WikiLeaks. Available at:
https://wikileaks.org/vault7/?marble#Marble Framework

Zero Day Initiative. 2017. Disclosure Policy [online]. Zero Day Initiative. Available at:
http://www.zerodayinitiative.com/advisories/disclosure_policy/

Zetter, K. 2015. A Cyberattack Has Caused Confirmed Physical Damage for the Second Time Ever [online]. WIRED.
Available at: https://www.wired.com/2015/01/german-steel-mill-hack-destruction/

Zivtech. 2015. The Benefits of Open Source Software [online]. Zivtech. Available at:
https://www.zivtech.com/blog/benefits-open-source-software