



2024

Cyber Realistic Disaster Scenario Development and Modelling

Triple threat: a new malware model for
systemic cyber insurance industry losses

A Beazley, Gallagher Re, and Munich Re Collaboration



Table of Contents

Executive Summary	4	Results	33
Key Conclusions	5	Impact of portfolio composition	35
Scenarios	6	Sensitivity testing	36
Insurance Portfolio	6	Conclusion	37
Model Methodology	7	Project learnings	37
Next Steps	7	Relevance to the ILS market	37
How to Read this Document	8	Looking to the future	37
Introduction	10	Invitation for comments	38
Motivation for the Project	11	Case Study — Use Case for Model: July 2024 CrowdStrike Faulty Update	39
Systemic cyber-risk: positioning this paper	11	Project Limitations	41
Malware presents a systemic cyber-threat	13	Supplementary Material Scenarios	43
The future of the evolving cyber-threat landscape	14	Autolycus Widespread Software	
Simple and transparent modelling	15	Supply Chain Attack	43
Insurance market	15	Lernaean Hydra Self-Propagating Malware Attack	46
Project Aims	17	Demeter's Curse Targeted Industry Loss Event	49
Modelling strategy	17	The Future of the Evolving Cyber-Threat Landscape: Detail	54
Intended outcomes	17	Additional Detail to the Approach	57
Excluded topics	17	Appendices	61
A New Approach to Modelling Systemic Cyber-Risk	18	Appendix 1: Definitions	61
Investigating the current state of modelling	18	Appendix 2: Summary of Cyber Insurance Coverage	62
Identifying Areas for Improvement	19	Appendix 3: Parameter Details	63
Principles for Scenario Development	20	Regional spread parameter	63
Adoption of an attack path methodology	21	Risk category	63
Counterfactual analysis	21	Fixed additional costs	65
Parameterisation	22	Attack path split for additional costs	67
Review of cost components	22	CBI parameters	67
Efficacy of security controls	23	Gross margin rates	67
Scenario development challenges	23	Percentage of daily revenue lost	68
Description of Selected Scenarios	24	Industry-specific event	68
Model Development	25	Footprint	69
Exposure set	25	Outage days	73
Top-down approach	27	Appendix 4: Impact of Portfolio Composition	75
Model construction	28	Appendix 5: Why the Model is Not Probabilistic	77
Scenario Quantification	32		

Authors

Michael Georgiou

Michael is the Head of Cyber Advisory at Gallagher Re, responsible for supporting clients across non-placement services with gross strategies and developing Views of Risk. He joined Gallagher Re in 2021 and supports clients across the value chain to build their understanding of systemic risk modelling and how this is applied within the business to drive improved capital efficiency across both gross and net strategies. Michael is a Fellow of the Institute and Faculty of Actuaries.

Stephan Brunner

Stephan is Senior Cyber Actuary in Munich Re's cyber actuarial services team. He joined Munich Re more than 5 years ago and is heading the topic of cyber accumulation modelling. Before that, he had different roles at another Reinsurance company. Stephan's background is mathematics and economics (in which he completed his PhD.), and he is also a certified actuary.

Ed Pocock

Ed is the Head of Cyber Security at Gallagher Re. He's responsible for supporting (re)insurers in developing and evaluating cyber catastrophe scenarios, as well as leading Gallagher Re's cyber-risk engineering capability. Ed has experience with Munich Re and in financial services cyber security consultancy with PwC. Ed holds the CISSP (Certified Information Systems Security Professional) qualification and an MSc in Cyber Security.

Tim Marshall

Tim is a Senior Underwriter in Munich Re's Cyber Centre of Excellence. He is responsible for treaty relationships with some of the largest cyber underwriters in the world. Tim has worked in various cyber (re)insurance roles for the past 10 years and is a Fellow of the Institute and Faculty of Actuaries.

Aidan Flynn

Aidan is the Head of Cyber Underwriting Strategy at Beazley. He is responsible for underwriting strategy, appetite, product innovation, and other related areas across the cyber portfolio. More recently, Aidan has focused on Beazley's strategy for management of systemic cyber-risk. He has 20 years experience across a range of senior underwriter roles in the London Market.

Henry Skeoch

Henry is the Exposure Management Lead within the Cyber-risks Underwriting Management function at Beazley with responsibility for leading and coordinating the analysis and communication of systemic cyber-risk within the Cyber-risks division. Henry completed a PhD in the Economics of Information Security at the UCL Centre for Doctoral Training in Cybersecurity, which followed 10 years professional experience as a macroeconomic strategist in the Investment Banking industry.

Alex Jackson

Alex leads the Exposure Management team covering cyber and specialty risks, at Beazley. This role includes leading the development and refinement of Beazley's View of Risk for cyber and liability catastrophe risk. Alex's background is in cyber security and he was previously a consultant at KPMG for 10 years, where his most recent role was a Senior Manager covering cyber-risk and cyber insurance.

Sioned Bentley

Sioned is a Cyber Security Consultant at Gallagher Re; she works closely with the Catastrophe Modelling Team to analyse, improve and develop cyber realistic disaster scenarios. With a background as an Information Security Officer, Sioned brings experience in cyber compliance and assurance, including data governance and ISO 27001 implementation and compliance.

Tim Davy

Tim is a Senior Cyber Security Specialist at Munich Re where he applies his cyber security knowledge within the Cyber Modelling team. He has over 20 years of experience in the IT security space including experience in Systems and Security Architecture where he has built large and complex defence and analysis platforms as well as Security operational roles. He has extensive knowledge in evaluating and understanding both security and technology market trends, as well as understanding the evolving security threats and business transformations across a variety of business verticals including Finance, Retail, Industrial, Government and Telecommunications.

Rishi Shamla, Capital Actuary at Beazley contributed extensively to the model and parameter development.

Executive Summary

This paper presents fresh estimates of the potential systemic losses that the cyber insurance industry could face from a significant malware event. These estimates are the output of a new accumulation model developed during a year-long collaborative partnership between Beazley, Munich Re and Gallagher Re ('The Partnership'). The Partnership brought together experts in actuarial modelling, technical cybersecurity, and underwriting spanning insurance, reinsurance and broking to produce a modelling paper on systemic cyber-risk whose outputs are fully transparent and available to any interested party.

The key aims the Partnership set itself for the model presented in this paper is that it should:

- Reflect the underwriting risk rating methodologies used in the cyber insurance market.
- Contain parameters influenced by actual insured losses and cyber-incidents.
- Be runnable by any interested party without specialised technology infrastructure or coding.
- Be fully transparent.
- Be representative of the composition of the current cyber insurance market.
- Be understandable without a strong technical background in cybersecurity or actuarial science.

The model is constructed by applying three distinct malware scenarios to a synthetic portfolio that is representative of the cyber insurance market aiming to address what the partnership regards as potential limitations to current estimation of systemic-cyber-risk. Many existing systemic cyber-risk scenarios focus on economic losses or contain elements that standalone cyber insurance policies are likely to exclude (for example, a cloud outage due to widespread power failure). Others contain parameters that are either not fully visible or are unintuitive to those without significant technical expertise.

The lack of significant catastrophic loss events impacting the cyber insurance market means that the parameters in a systemic cyber-risk model will inherently contain a degree of subjectivity. However, this does not mean that the inputs to a model cannot be fully justified, explained, and evidenced. Great care has been taken by the Partnership to make the model as simple and transparent as possible and also well-evidenced. External experts were consulted during the paper development and existing academic and industry research is referenced throughout the paper. The claims and incident response costs assumptions used in the paper are based on actual claims experience for realism, but are aggregated and smoothed to preserve client confidentiality. It is hoped that this serves to provide modellers of systemic cyber-risk with a benchmark for calibration of incident costs.

The Partnership hope that this paper, the scenarios and model contained in it serve both to enhance understanding of the systemic threat from malware and to act as an enabler of discussions on the scope of potential systemic insured losses from such events and their modelling.

All references to the 'Partnership' are to be understood as references to a collaborative partnership in the colloquial sense. Beazley, Munich Re and Gallagher Re did not establish any legal partnership, joint venture or similar for the purposes of producing this paper and none of the participants are constituted the agent of another or otherwise authorised to act on another participant's behalf.

Key Conclusions

While acknowledging that systemic cyber-risk is not confined to malware alone, the Partnership decided to focus on malware because of its accepted potential for loss; the perceived difficulty of constraining the insured loss with policy language; the potential volatility and variance of the event; and the cyber-threat landscape. The outputs are intended to be complementary rather than contradictory to existing academic and industry work on systemic cyber-risk. The key conclusions are:

- It is possible to construct a simple and transparent yet insurance-relevant model for systemic cyber-risk based on malware scenarios, though parameterisation is a challenge.
- It is possible to generate realistic systemic losses without needing a large and complex set of parameters or scenarios. Whilst a probability cannot be objectively assigned to these scenarios or estimated from any current data, the Partnership regards the outputs of the model in this paper as a representation of the tail risk that exists in cyber.
- It is important to note that the modelled loss outputs are highly sensitive to the choice of parameter values, as with any model. The values assigned to the model from the developed scenarios are deliberately extreme and designed to represent close to an upper bound to what the Partnership felt was technically possible.
- Despite the extremity of the scenarios, the model suggests that if they occurred then they would not exhaust a significant proportion of the deployed limit. The modelled losses are over twice the premium collected by the market, meaning if the most severe of these events occurred, market performance would be in the region of 300-330% Combined Ratio, including significantly adverse attritional performance. The loss ratios incurred suggest that the market would survive a systemic event but highlight the importance of a strong capital base and diversified portfolio.
- When considering malware attacks, a widespread software supply chain or self-propagating malware pose the greatest risk of systemic loss to the cyber insurance market. A targeted industry loss event, while significant, lacks sufficient footprint to reach a similar magnitude of insured losses although the economic loss could be significant.

It should be noted that the model is intended as guidance in how to frame a systemic event and it is not a substitute for a fully-specified capital model. Beazley, Munich Re, and Gallagher Re individually make use of a variety of models to manage their capital and business with parameters that are carefully documented and extensively reviewed using appropriate expert and prudent risk judgement. To use this model for a probabilistic risk appetite, the parameters would need to be mapped to probabilities via parameter curves. This approach would increase the complexity of the model and it is crucial to note that validation of probabilistic parameter curves for systemic cyber-risk is non-trivial as the dependence of insured companies on technology is complex.

The model is intended to represent a plausible worst case from a specific set of scenarios and to complement, not contradict, existing work by offering a third-party data point contextualising a severe malware event. It is somewhat reassuring that the outputs of this simple, jointly developed, and conservatively parameterised model are broadly similar in magnitude to the tail of commercially available models. The intention of the Partnership is that this model simplifies understanding of systemic cyber-risk and that the model can be used to further validate views of risk or to serve as a starting point for those whose view of cyber-risk is at an early stage of development.



Scenarios

The model is constructed using three distinct malware-driven scenarios developed by combining previously experienced cyber-incidents and expert judgment. The scenario narratives are intended to be sufficiently illustrative to make them readily understandable by a broad audience yet provide key elements for parameterisation: appropriate scale, technology components and unintended consequences.

The scenarios use a condensed version of the MITRE ATT&CK framework, split into four stages: initial access, privilege escalation, lateral movement and impact.

Each of the three scenarios investigates how significant losses could arise via a distinct pathway:

Scenario	Name	Description	Footprint	Per-Insured Loss
1	Autolycus	Widespread Software Supply Chain Attack (WSSC)	High	High
2	Lernaen Hydra	Self-propagating Malware (SPM)	Highest	Moderate
3	Demeter's Curse	Targeted Industry Loss Event (TILE)	Concentrated	Highest

Insurance Portfolio

To support loss quantification, the Partnership created a synthetic portfolio of 20,000 risks intended to broadly reflect the global industry exposure, considering the observed mix of policies by geography, industry, and size (revenue). This portfolio was derived from the Gallagher Re Industry Exposure Database as of mid-2022 consisting of over 1.2m insured policies and \$13b of premium.

Further testing was performed on sub-segments of the portfolio, considering an SME book and a global large corporate risk book to explore the sensitivity of the results to portfolio composition and to consider the differences based on varying portfolio characteristics.

Please note that no individual, specific client data was disclosed to Beazley or Munich Re in creating the synthetic portfolio. Any data exposed was in duly aggregated and anonymised form.

The below tables summarise the portfolio by geography, revenue, and industry, noting this is broadly reflective of the market but is not intended to be an exact replica.

Size	Max Revenue	Premium %	Count %
Nano	1,000,000	9%	57%
Micro	10,000,000	8%	26%
Small	250,000,000	20%	14%
Medium	1,000,000,000	16%	1%
Large	Unlimited	46%	2%
TOTAL		100%	100%

Region	Premium %	Count %
North America	74%	63%
Europe	20%	27%
Asia	1%	1%
Other	5%	9%
TOTAL	100%	100%

Users should consider how their own portfolios may differ to the market, and how the relevant market segments have evolved since.

Model Methodology

The Partnership used the three scenarios to construct a simple and transparent model breaking down the insured losses for each scenario into three components:

Direct losses

Predominately business interruption losses following a successful attack against an insured.

Partial losses

Investigation and recovery costs incurred when an attack against an insured fails to make it to the end of the kill chain but is partly successful.

CBI' losses

Losses to an insured from an attack against an entity on which its operations are reliant (e.g. a supplier), but which is not under its direct control.

The model places emphasis on the relative risk quality of an average risk in different industry and revenue bands. It is evident that not all companies will react the same way to a given scenario. To capture the volatility in the loss per insured seen in an event, calculations use the LogNormal distribution (as is standard in many insurance loss calculations) and a sampling approach to allow an approximation of the expected loss to an insurance layer. This avoids the need for cumbersome simulations, while still ensuring that higher insurance layers are not treated as always loss free.

Next Steps

The Partnership recognise that in cyber modelling there is always more that can be investigated but hope this work makes a useful contribution to ongoing market discussions around systemic risk and market development. Within the model, the areas the Partnership would prioritise for further investigation and refinement are:

- 1 Other families of event, such as non-malicious widespread cloud outage.
- 2 Collation of feedback within Insurance, Government, and Technology.
- 3 Deeper analysis on single meaningful parameters i.e., worm spread, revenue, dependency on IT, efficacy of network segmentation, challenge outage time.
- 4 Use of more current and varied portfolio information.
- 5 Simplistic stochastic model or other method to ascribe return periods.
- 6 Breaking out more loss components from the fixed cost element.

We welcome feedback from any parties that, having read this paper, wish to engage in the discussion of systemic risk. We hope this generates useful discussion and helps advance accumulation modelling excellence, market sustainability and the flow of capital within the cyber insurance industry.

Feedback is welcomed via email: CyberSystemicPartnership@gallagherre.com

We remind respondents not to submit any commercially and/or competitively sensitive information in feedback, in order to ensure competition law compliance.

How to Read this Document

We appreciate that this document is extensive, and, in recognising the diverse range of potential interests and expertise among its readership, the authors have produced this guide to direct each reader to the sections most relevant to them with the goal of helping digest its findings efficiently.

Table 1 below provides a brief synopsis of each of the major sections within this paper. Table 2 then provides guidance on which sections those holding roles within the cyber insurance industry may wish to prioritise. The whole paper is designed to be accessible to those without existing technical expertise but it is recognised that not every reader will have time to read its entire contents.

Table 1: Section Synopsis

Section	Synopsis
Motivation for the Project	This section sets out the Partnership's motivation for initiating this project. It talks to our collective view of the current cyber threat landscape, the limitations in cyber-risk modelling and our reasons for choosing malware as the peril of consideration.
Project Aims	This section outlines the specific goals of the project, detailing what it seeks to achieve in terms of developing and applying a new systemic model.
A New Approach to Modelling Systemic Cyber-risk	In this chapter, the reader is provided with an overview of the development process utilised by the Partnership. This includes analysis of the current state of modelling systemic cyber-risk.
Principles for Scenario Development	This section presents the principles utilised by the Partnership to develop realistic and relevant malware threat scenarios. These include adoption of an attack methodology and use of counterfactual analysis.
Description of Selected Scenarios	A brief segment of the report which provides readers with an overview of the three narratives developed (and subsequently modelled) by the Partnership.
Model Development	This section introduces our modelling methodology. It provides detail into various components of the model, including our exposure set, the top-down approach utilised and insight into model construction.
Results	This chapter of the report details the results from our modelling efforts at both a scenario and thematic level (this includes sensitivity testing).
Conclusion	An extensive review of our findings, project learnings and how the Partnership plans for this work to continue to contribute to the future of systemic cyber event modelling. Project limitations are also included within this section.
Supplementary Material: Scenarios	This provides the reader with a complete description of each of the three narratives developed by the Partnership for those who wish to examine them in greater detail. Additional supplementary material includes further detail to the scenario development approach (including how the Working Group was established and operated throughout the project).
Appendix 1: Definitions	A definitions table for key terms that are used throughout the paper.
Appendix 2: Summary of Cyber Insurance Coverage	Complementary material that enables the reader to have a deeper understanding of existing/common cyber coverage approaches.
Appendix 3: Parameter Details	This Appendix provides details of each parameter within the Partnership's model. This section includes guidance for potential future use and current observations regarding the impact of each parameter. Parameters include regional spread, risk categories, fixed costs, gross margin rates and outage time.
Appendix 4: Impact of Portfolio Composition	Additional outputs from the sensitivity testing conducted by the Partnership.
Appendix 5: Why the Model is Not Probabilistic	A final section dedicated to explaining the logic behind the decision to develop a deterministic model as opposed to a probabilistic one.

Table 2: Section Recommendations by Role

	Insurance Executive	Capital Provider/ Investor	Cyber Executive	Underwriter	Actuary	Cyber Security Specialist	Exposure Manager	Portfolio Manager
Motivation for the Project	✓	✓	✓	✓	✓	✓	✓	✓
Project Aims	✓	✓	✓	✓	✓	✓	✓	✓
A New Approach to Modelling Systemic Cyber-risk					✓	✓	✓	✓
Principles for Scenario Development					✓	✓	✓	
Description of Selected Scenarios	✓	✓	✓	✓	✓	✓	✓	✓
Model Development					✓		✓	✓
Results			✓	✓	✓	✓	✓	✓
Conclusion	✓	✓	✓	✓	✓	✓	✓	✓
Supplementary Material: Scenarios						✓	✓	
Appendix 1: Definitions								
Appendix 2: Summary of Cyber Insurance Coverage	✓	✓						
Appendix 3: Parameter Details				✓	✓	✓		
Appendix 4: Impact of Portfolio Composition				✓	✓	✓		
Appendix 5: Why the Model is Not Probabilistic				✓		✓		

Introduction

Foreword

This paper introduces a fresh perspective on understanding the damage that the cyber insurance industry might experience from systemic cyber-attacks. The concept of systemic cyber-risk is most typically associated with media depictions of societal breakdown resulting from a failure of technology. In the insurance industry, there is a common practice of quantifying possible yet unlikely systemic events using scenarios that describe how such events might generate significant claims on policies issued by insurers. These are commonly known as ‘Realistic Disaster Scenarios’ or RDS. The important element to emphasise here is realism; what makes for riveting viewing is often inherently unrealistic.

The cyber insurance market has seen impressive growth since 2010. What started as a market largely covering technology-related liability losses has evolved to cover a range of potential losses related to the use of technology: everything from business interruption due to IT systems failing, to third-party liability stemming from a malicious data breach. The breadth of coverage is to the benefit of buyers of cyber insurance but presents a complex modelling challenge. Not only is technology sometimes hard to understand and has complex interdependencies, but when considering potential malicious activity, the threat landscape has to be considered. The threat landscape is constantly evolving, as are defences. On top of this, whilst the insurance industry is accumulating sufficient data to price policies for normal expected losses, there has not been a large-scale event that could be used to directly calibrate catastrophe models. The WannaCry and NotPetya malware strains are often referenced, but their footprint was relatively limited with only a few companies seriously affected across their global IT estate and they are over six years old.

In 2023, three of the leading companies in the cyber insurance landscape came together to form a collaborative research effort known as the ‘The Partnership’. The Partnership is formed of Beazley, a leading FTSE-100 listed speciality insurer and writer of standalone cyber insurance policies; Munich Re, a global reinsurer; and Gallagher Re, one of the leading reinsurance brokers (collectively the ‘Parties’).

There are many existing publications, models and industry fora discussing systemic cyber-risk, yet in regular business discussions it became clear that a better understanding of the catastrophic loss potential associated with cyber insurance is needed to support the continued growth of the cyber insurance market. The uncertainty associated with cyber insurance as a relatively new class of business with significant accumulation potential means that most capital providers limit their exposure to a level below that of more well-understood perils. In a growing line, where every new policy that is written requires marginally more allocated capital, there is a constant need to increase the capital allocated to the line of business to prevent a scarcity or ‘capital-crunch’ from stunting the growth and relevance of cyber insurance.

The aim of this paper is to demonstrate that there is a quantifiable ceiling on the potential losses the cyber insurance market could experience from a plausible yet remote significant malware event in the hope that it provides a degree of reassurance to capital providers. Of course, malware is not the only potential source of systemic losses, but as will be extensively discussed in this paper, there are good reasons why it is the family of scenarios the Partnership feels makes most sense to consider. The Partnership hopes that this paper helps convey one insurance perspective on systemic cyber-risk to the broader community and looks forward to engaging broadly on the topic with interested stakeholders including but not limited to academia, insurance market participants, law enforcement, regulators and the cyber-security community.

All references to the ‘Partnership’ are to be understood as references to a collaborative partnership in the colloquial sense. Beazley, Munich Re and Gallagher Re did not establish any legal partnership, joint venture or similar for the purposes of producing this paper and none of the participants are constituted the agent of another or otherwise authorised to act on another participant’s behalf.

Motivation for the Project

Systemic cyber-risk: positioning this paper

The concept of systemic cyber-risk is not new. In 2016, the World Economic Forum (WEF) defined systemic cyber-risk as:

- The risk that a cyber event (attack(s) or other adverse event(s)) at an individual component of a critical infrastructure ecosystem will cause significant delay, denial, breakdown, disruption or loss, such that services are impacted not only in the originating component but consequences also cascade into related (logically and/or geographically) ecosystem components, resulting in significant adverse effects to public health or safety, economic security or national security.^{12,3}

This paper starts from accepting the premise that cyber-attacks can cause systemic risk and that this risk can originate and be categorised in a number of ways. For example, the WEF noted in 2022 that systemic cyber-risk can originate via⁴:

Common cause risks

Risks that originate when multiple organisations utilise the same hardware, software, or communication tools, which create the possibility that multiple failures may arise from a single underlying defect.

Shared service risks

Risks generated by organisations that leverage the same cloud providers or social media platform.

Operational dependency risks

Risks which occur when the disruption in one organisation's operations, such as a shutdown of an electricity grid, disrupts many other organisations' operations, creating a cascading effect across multiple entities.

Shared trust and confidence risks

Risks which stem from activities with over-reliance on — and subsequent loss of — the trust that data and processes are accurate and reliable.

The WEF also notes three ways that systemic cyber-risk can manifest⁵:

Flow risks

Which ‘include the risks that flow from one organisation to another through a multitude of connection and interlinkages. This includes risks that transfer along physical or operational connections between organisations (sometimes described separately as chain risks).’

Simultaneous emergence risks

Which are ‘those risks that appear simultaneously across many different organisations.’

Behaviour risks

Which ‘are the risks propagated by many people or organisations changing their behaviour in a short period of time, such as when the COVID-19 pandemic caused many people to work from home.’

This paper is not intended to debate or challenge the definition, categorisation or theories associated with systemic cyber-risk, rather, it aims to pragmatically try to quantify some plausible yet remote scenarios. Cyber-risk has some uncommon attributes which make it hard to model — as the Geneva Association noted in its November 2023 report:

- ‘Cyber is an anthropogenic peril and the extent of any harm depends on the interplay between the incentives, motives and resources of both victims and attackers, which often involve complex, non-linear relationships among multiple factors.’⁶

This paper focuses on a small part of the discourse on systemic cyber-risk — namely, a structured way of assessing the potential level of insured loss that a systemic cyber-risk event could cause to the insurance market, and some indicative ways that this could occur (i.e. scenarios). Other scenarios and papers have sometimes primarily detailed economic loss, as opposed to insured loss, for example in relation to the impact that a cyber-attack and associated consequences could cause to Gross Domestic Product (GDP), or investment returns.

The October 2023 Lloyd’s Futureset report, in collaboration with the Cambridge Centre for Risk Studies, received widespread attention for a report detailing potential 5-year economic loss of \$3.5trn in relation to a cyber-attack on a major financial services payment system.⁷

By firmly focusing on insured loss, this project can be seen as complementary to past and current work being performed on systemic cyber-risk, by think tanks, regulators, insurance communities and academia.

Malware presents a systemic cyber-threat

There is no explicit and singular definition of a 'malware event', as malware can be a primary causal factor of loss, a distribution mechanism for other attacks, a catch-all term for many different pieces of software, one of a number of tools used in an extensive attack chain of threat actors, and many other things.

For the purposes of this project, we have deemed malware events to be loosely defined, but broadly seen as events whereby the insured losses themselves can be directly and explicitly tied to the deployment and execution of malicious code in a widespread, and potentially self-propagating, way.

The focus on malware was chosen for the following primary reasons:

- **Accepted potential for loss:** it is widely accepted that malware can cause high levels of insured loss. For example, this can be seen in:
 - » Academic and industry publications.
 - » Regulatory publications – for example, in insurance stress tests or realistic disaster scenarios (RDS). For example, the European Supervisory Authority EIOPA (European Insurance and Occupational Pensions Authority) notes that 'no fully systemic ransomware event has yet been observed' but believes that this could occur.⁸
 - » Widely used third-party cyber accumulation models provided by vendors ('vendor models') to the insurance market. Malware scenarios are generally seen as contributing the most of all scenario-types (e.g. versus cloud scenarios or data breach scenarios) to the low frequency, high severity events that make up the 'tail' of the insured loss curve.⁹
- **Relevance to the Partnership:** the Parties collectively agreed that a focus on malware was relevant as we considered it to be one of the most significant scenarios for our respective portfolios. We expect that this is the case for the wider insurance market as well.

- **The perceived difficulty in constraining the insured loss with policy language:** it is challenging to strictly define a malware event and it is therefore more difficult to use levers such as sub-limits or exclusions, or other policy wording, to manage the exposure of (re)insurers to this type of event. This is mostly because there are a variety of different ways that malware events could materialise. There can also be challenges understanding when losses form part of an event, or whether individual attacks are unrelated. By comparison, cloud outages are relatively well confined to the Contingent Business Interruption coverage, and therefore any changes to this cover will impact a cloud outage event.
- **The volatility and variance of the event:** malware events can be particularly volatile and variable, both in concept and in practice. Partially this can be down to the nature of the attack because the term 'malware' covers a large number of possibilities, including ransomware and 'wiperware' or destructive malware. The volatility can also occur because malware events typically involve many individual companies' own IT 'network/s' being compromised, and recovery/minimisation of insured loss relies on the individual ability of companies. By contrast, in a cloud outage event there is a more defined point of failure, and recovery would typically rely on the controls and resilience of the cloud service provider, whose business model is built around resilience and availability, and individual IT networks of customers would not typically be compromised.
- **The cyber threat landscape:** malware is a prevalent threat type, in historic, current and future intelligence on the cyber threat landscape.

By focusing on malware events, we are not focusing on other types of cyber events, but this does not mean that we believe that high levels of insured loss cannot be caused by other types of events. That view is reflected in other publications, not least by regulators, who are concerned by the potential impact of an outage to cloud service providers and other critical third parties, whether that impact is triggered by cyber-attack or other means. In the cyber scenarios contained in the UK Prudential Regulation Authority's (PRA) Insurance Stress Test in 2022 it was the extensive cloud outage scenario that caused the highest level of insured loss, but this scenario was also deemed as being the most remote in likelihood, by respondents. As stated elsewhere in this paper, other types of events could be addressed in future, in a similar project.

The future of the evolving cyber-threat landscape

We tried not to be constrained by the historic or current cyber threat landscape. As such, we aimed to embrace a forward-looking view and consider incorporating this into our scenarios, in order to be more future-proof. As EIOPA notes:

- 'A widespread shock impacting the cyber insurance market would involve a common cause that may not have been clearly identified until now. It may involve new techniques developed by cyber criminals, but also new IT practices or usage happening across one or several industries. Given the pace at which the digital industry is evolving in the late years and given the pace at which new hacking techniques are invented, such an event cannot be considered as unlikely.'¹²

This was also the feedback received from several external third-party experts consulted during the project — historic events might not provide the full scope of what may be possible in the future. As a result, the Partnership performed some horizon scanning through conversations with experts and review of publications on the future cyber threat landscape (c.2023–2030). The analysis aimed to answer questions such as the following:

- » How will technology change over the next few years and how could this impact the nature of malware cyber-attacks and the susceptibility of companies to them?
- » What does the cyber threat landscape look like in the future? For example, how does Artificial Intelligence (AI) impact the ransomware crime ecosystem?
- » What impact could geopolitical developments have on threat actors and cyber-attacks? For example, work sponsored by the UK National Cyber Security Centre (NCSC) has suggested that there are new groups of threat actors that were created during the Russia-Ukraine war, and these groups may turn their attention to performing other cyber-attacks once the conflict ends.¹³

The analysis showed three primary themes that should be considered for our scenarios, per below. Further detail on the analysis and each of the themes can be seen in the supplementary material.

- AI abuse/machine learning (AI/ML).
- Internet of Things (IoT)/Industrial Internet of Things (IIoT).
- Quantum computing/cryptography.

In summary, our key conclusions from the horizon scanning analysis were:

- The threat landscape moves rapidly, with significant change in how malware operates and how organisations deploy security controls since the NotPetya event in 2017.
 - » The tactics, techniques and procedures preferred by threat actors will continue to evolve, but the fundamental way computers operate is not expected to change in the short-to-medium term.
- Many of the trends observed had the potential for significant economic loss or social and political unrest but had limited direct impact to our scenarios or insured losses. Near-term emerging threat trends, such as from generative AI, are unlikely to impact the fundamental nature of our malware scenarios and subsequent insured losses.
- However, the analysis did help to act as input data into the scenarios, as follows:
 - » Inputs for challenging the boundaries of what may be possible in scenario plausibility. For instance, assumptions have been made that there may be a greater ability in the future for a cyber-attack to impact more companies than at present (e.g. see the UK NCSC paper on the impact of AI on cyber-attacks).¹⁴
 - » Inputs for the data parameters in the scenario modelling. For example, based on the example above, consideration of AI has been fed into the decision-making process for the data parameters for each scenario.

The scenarios have also been constructed to enable re-parameterisation, in future, to account for material shifts in threat trends.

Simple and transparent modelling

We had four core focuses when developing this model:

- 1 Create a simple model anyone can recreate and validate.
- 2 Consider very severe events for risk management purposes (not the entire loss curve for pricing).
- 3 Be transparent about all assumptions and parameters.
- 4 Make use of data and other insights from the Parties, as well as third-party external experts.

Thereby, we are giving a starting point for open discussion within the market about the scale of potential cyber losses and whether there are different views on fundamental parameters.

To develop a loss curve would require complex parameters and simulations to attempt to model, instead we have aimed to describe and then parameterise severe but plausible scenarios which push towards the upper bound of a loss we can expect to see. As many loss curves tend to flatten in higher return period ranges, rather than aiming to assign a granular probability, the aim was to ensure that the scenarios are sufficiently extreme to represent an obvious tail position. For risk management purposes such an extreme scenario view is sufficient, however without the full loss curve it is not possible to use such a model for pricing considerations. This is discussed further in Appendix 5.

Insurance market

For common understanding, we have included a short (simplified) explanation of the operation of the cyber insurance and reinsurance markets, covering the risk transfer chain from insured to capital.

Market structure

The global cyber insurance market is estimated to have generated approximately \$14b premium in 2023 and to grow to \$15b in 2024¹⁵, with two-thirds of the premium originating from North American businesses.¹⁵ Of this premium, Gallagher Re estimated that of those insurers purchasing quota shares, an average of 50% of premium was ceded to reinsurers in 2022 through proportional Quota Share¹⁷ placements, and this is estimated to reduce to around 45% in 2024. This means that up to half of all the risk is sitting with the reinsurance industry.

Most policies are sold to businesses through broker intermediaries, but there is some direct distribution at the SME end of the market. There has been an increasing trend, particularly in the mid-market and large risks, for cyber insurance to be sold as a standalone product. Many SME business now also have standalone cyber policies, but a large amount of SME cyber exposure is sold as a policy add-on or as part of a package with other business risks.

For SME policies, limits are typically small, averaging \$780k¹⁸ for insureds with revenue between \$1m and \$10m per year; this is usually sold by one carrier. Many insurers can only offer limits up to \$5m or \$10m. With large companies often purchasing total limits in excess of this, into the low hundreds of millions, the brokers build 'towers' of cover with layers (policies) from different carriers placed on top of one another to reach the desired amount of insurance, or syndicated risks with each carrier taking a proportion of a larger limit (e.g. one insurer providing 10% of a \$100m layer pays just \$1 in every \$10 of loss).

This structure for large risks allows carriers to manage their exposure to any one insured and also build a more diverse portfolio of risk, not just by industry, country, and revenue, but also by the size of loss required before their layer is affected.

Contract wording/coverage

For there to be an insured loss from the scenarios developed in this paper, the cause and consequences of the malware scenarios created need to 1) trigger an insuring agreement; and 2) not trigger one of the exclusions. Whilst there are many policy forms and variations of coverage provided in the market, the majority have similar coverage, albeit with potentially different wording.

However, whilst the language differs the core coverage is often similar, Appendix 2 provides a summary of the most common coverage components in a standalone cyber insurance policy. An example coverage used by Beazley has also been referred to during the work, but we make no representation that this coverage is market standard.

Standard malware would usually be seen to trigger the data restoration, incident response and business interruption insuring agreements as a Security Breach of the Insured's computer systems.

There are two exclusions, often seen in cyber policies, which are particularly relevant when developing widespread malware scenarios:

- **Critical infrastructure:** losses resulting from the failure of a public utility such as power, Internet or telecoms. This means we do not need to consider the secondary effects of a malware disrupting a major public utility, or scenarios which focus predominately on such infrastructure.
- **War:** it is usual to also exclude cyber losses stemming from war or war like activities of a nation-state. Therefore, our scenarios do not need to consider the skills and motivations only held by nation-states.

For the rest of the paper, we will assume that there are no coverage issues related to these scenarios, except where we specifically consider the critical infrastructure exclusions and war exclusions in the construction of the scenarios.

Capital: protecting against extreme events

The working assumption of the Partnership is that a large-scale issue which triggers many cyber policies in a short period of time is possible. This can be described as accumulation/systemic or catastrophe (CAT) risk ('systemic cyber-risk'), all of which we have taken to mean the risk that a portfolio of cyber insurance policies has the potential, in certain situations, to be liable to pay losses far in excess of the premium collected.

Insurers are required by regulation to hold money (capital) on the balance sheet to cover such eventualities, and of course the insurer needs to have access to this capital and deliver sufficient return on it to make the whole proposition viable. Specific regulatory requirements vary by geographical location and line of business, and a discussion of these lies outside the scope of this paper. Where an insurer does not want to hold capital to support their entire portfolio they are able to pass (cede) some of the risk to reinsurers, effectively making use of the reinsurer's balance sheet at a cost. This may also be done to balance different lines of business within an insurer's portfolio or generally to control volatility, even if sufficient capital is available.



Project Aims

Early discussions amongst the Partnership identified that widespread damage from malware scenarios would be a meaningful modelling target for cyber insurance accumulation scenarios, in terms of both loss size and potential for further refinement.

The aim of the modelling is to consider the chains of events that might lead to a systemic loss from malware, which are remote yet plausible. The focus is predominantly on what might happen rather than the precise details of how it happens. Some of the more significant malware events such as Stuxnet and WannaCry used techniques that were previously not widely known. A detailed malware scenario is highly specific to the systems it targets and there is a danger of false precision in ascribing arbitrary technical details. The skill in constructing remote scenarios is to leverage existing technical knowledge yet not be overly constrained by it. As such, the modelling assumptions developed by the Partnership draw on technical expertise from its constituents and also from external third-party experts;¹⁹ this is known in the insurance industry as 'expert judgement' and is standard for modelling the impact of extreme perils on insurance portfolios.

Modelling strategy

- Investigate representative malware scenarios that could generate potentially capital-depleting losses, supported by detailed narrative and plausible technical mechanisms.
- Bring together expertise from across the cyber insurance ecosystem to challenge thinking on systemic cyber, collaborating with a broad range of perspectives, expertise and experience.
- Produce a simple model that takes into account insurance portfolio conditions.
- Propose best estimates for model parameters at capital-relevant levels.
- Produce cyber insurance industry-wide loss estimates derived from these assumptions.

Intended outcomes

- Deliver a research project on systemic cyber-risk that is fully transparent, available to any interested party.
- Address what the Partnership regards as potential limitations to current estimation of systemic-cyber-risk arising from narratives not grounded in technical realism or relevant to the insurance industry.

- Act as a third-party data point to complement existing quantitative views of systemic cyber-risk.
- Communicate the limitations of existing thinking of systemic cyber-risk.
- Stimulate an informed and public discussion of how to best model these catastrophic perils.


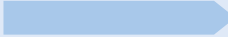
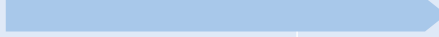
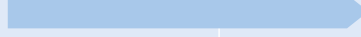
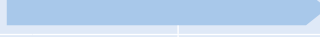
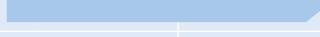
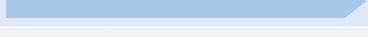
Excluded topics

- Non-malware scenario-types (e.g. cloud or mass hardware failure).
- Defining what constitutes a single event; the scenarios considered are written without consideration of how an event within (re)insurance may be formally defined in a policy wording or reinsurance contract. Instead, we focus on mechanisms which can cause a large amount of loss in a short period of time.
- Insured losses from non-cyber insurance covers that may pay out in the event of a cyber event (e.g. Directors and Officers insurance).
- Cyber events not covered by standalone cyber insurance policies, given the typical exclusions in the insurance market such as public infrastructure.
- Cyber events explicitly deemed as meeting the definition of 'cyber war' under a typical standalone cyber policy – in practice, meaning we discounted capabilities only held by nation states.
- 'Physical damage' impacts associated with cyber events, and the insurability of any resulting losses, under a standalone cyber policy.
- 'Attritional' cyber losses such as those stemming from ransomware campaigns by multiple threat actors requiring manual input for different targets over a long time span.

We hope that by adding to the public discourse on various critical scenarios, and parameters we can provide key stakeholders with increased confidence around the potential downside and associated capital charges.

A New Approach to Modelling Systemic Cyber-risk

The Partnership completed the work detailed in this paper over an approximate 12-month period starting in Q1 2023. The structure and governance around the work can be seen in the supplementary material. The work was completed in the following key stages:

Phases	2023				2024		
	Q1 2023	Q2 2023	Q3 2023	Q4 2023	Q1 2024	Q2 2024	Q3 2024
Investigating the current state of modelling							
Identifying areas for improvement							
Principles for scenario development							
Scenario definition							
Modelling							
Scenario quantification							
Reporting and communication							

Investigating the current state of modelling

This phase involved looking at the current landscape in relation to systemic cyber-risk. The analysis was performed to help identify any potential areas for improvement in the current thinking.

Whilst the analysis was not exhaustive of all sources and thinking, three key bodies of material were covered, which collectively helped to inform the project:

- Widely used third-party cyber accumulation models provided by vendors to the insurance market.
- Published cyber disaster scenarios, which were split into those published by:
 - » Regulators (e.g. as part of supervision activities and market stress tests).²⁶
 - » Other bodies (e.g. hypothetical stress tests published to help insurers with risk management).²⁷

- Publications on systemic cyber-risk by academic centres/think tanks,²⁰ industry bodies and marketplaces (e.g. Geneva Association²¹), non-governmental organisations and supra-national agencies (e.g. World Economic Forum²²). Relevant peer-reviewed academic papers include Eling et al²³, Hillairet & Lopez²⁴ and Baldwin et al²⁵.

In conclusion, the analysis suggested:

- Large-impact malware events do not necessarily have to be multi-sector, global scenarios — sector or system-specific malware can cause large impact.
- Scenarios can generally be split into:
 - » Widespread compromise due to vulnerabilities in commonly used systems/software (distributed attack across many firms e.g. WannaCry attack).
 - » Focused attacks through compromise of a single point of failure — either causing impact on its own, or as a distribution channel to attack users of the provider (e.g. distribution of malware through a software update from a trusted provider, such as the SolarWinds attack).
- The types of malware used in existing scenarios varies but included ransomware (very commonly), wiperware/destructive malware, data scrambling (corruption) malware, data integrity-changing malware, malware on rogue hardware and malware for industrial controls systems.
- The types of vulnerabilities leveraged in existing scenarios included vulnerabilities in operating systems (OS), database software, various industrial control systems and IoT devices, web applications, chip architecture, mail servers and physical hardware.

Further detail on the analysis can be seen in the supplementary materials.

Identifying Areas for Improvement

The Partnership identified a number of current limitations, challenges and opportunities for improvement in relation to existing published scenarios. Several of these areas were prioritised during this project, as the Partnership felt that it would have additional information that could enhance them (e.g. in terms of methodology, expertise or data sources). These efforts were supported with consultation with various stakeholders and industry experts, alongside additional actuarial and catastrophe modelling support.

The areas prioritised were as follows:

- **Geographical nuances:** focus on global events with allowance for regional differences.
- **Underwriting risk quality and claims data:**²⁸ use of real-world data to model parameters.
- **Event response:** improved consideration of event response mechanisms, redundancy planning and dependency factors.
- **Attack propagation:** transparent use and application of attack path modelling and propagation methods.

Further detail on the areas prioritised can be seen in the supplementary material.

Principles for Scenario Development

The primary objective in the scenario development phase was to create scenarios that addressed some of the constraints described above and could be modelled. Our aim was to design narratives that would be both illustrative enough for the scenarios to be feasible, yet also provide key elements of parameterisation, appropriate scale, technology components and the concept of intended and unintended consequences. By using these steps, the Partnership were able to build scenarios that illustrate key characteristics of events that would provide significant losses, whilst remaining generic enough that adaptations can be made in future. In order to achieve this, we determined a threshold for specificity — for example, although generally available software, technology, and common vulnerabilities are implied they are not explicitly mentioned, in keeping with the volatile nature of cyber-risk. To this end, we also looked to model the scenarios such that the efficacy of security controls were detrimentally affected, to take into account the ‘perfect storm’ of vulnerability, exploit, payload and affected technology, that results in global security efforts in more of a reactive mode than may be the norm. This is in line with a stress test approach conducted by regulators, such as the UK Financial Policy Committee, which assumes disruption has occurred and does not examine preventative or detective controls.²⁹


The following themes were key parts of the scenario development process:

- **Adoption of an attack path methodology:** use of an industry-standard MITRE ATT&CK framework (MITRE framework) to structure the scenarios.
- **Counterfactual analysis:** using what-if analysis, in relation to historical events, to develop and stress severity factors for our scenarios.
- **Parameterisation:** development of an achievable list of data parameters for each scenario.
- **Review of cost components:** use of real-world data,²⁸ for example use of claims data to inform the way that losses were parameterised and distributed both per event and per insured.
- **Efficacy of security controls:** consideration of the effectiveness of different controls and the likelihood of control failure and the resultant impact on scenarios.
- **Scenario development challenges:** acknowledgement that scenarios would need to be sufficiently applicable across industries and technologies, whilst also being technically feasible from a cyber security perspective.

Adoption of an attack path methodology

It was important to the Parties that an appropriate methodology was utilised when developing these scenarios. Munich Re have been working for some time with The Adversarial Tactics, Techniques, and Common Knowledge or MITRE ATT&CK (MITRE framework) which is an industry-standard guideline for classifying and describing cyberattacks and intrusions, alongside the relevant mitigations for defenders.³⁰

When adapting the framework for our own use, we chose to condense the MITRE framework 'Tactics' or stages into four: initial access, privilege escalation, lateral movement and impact (see graphic below).



Within each of these we identified the success factor of the stage and the population exposed. Condensing the detailed MITRE framework into a more user-friendly model enhanced accessibility and usability without compromising on the depth of detail inherent in the framework. By simplifying the presentation of information and streamlining the structure, we were able to make it easier for future users to navigate and understand key concepts and tactics. This approach enables stakeholders, regardless of their level of expertise, to effectively leverage the framework.

Counterfactual analysis

A counterfactual analysis is defined by Woo et al. (2017)³¹ as a what if exercise designed to explore hypothetical alternatives to historical events by modifying them in some way. We utilised this exercise early in the narrative development process in order to conceptualise feasible new cyber events based on historical incidents. Our focus here was on understanding and developing the severity parameters of each narrative by considering severity factors of previous cyber-attacks and identifying how these could be made more extreme (without undermining plausibility).

We identified that altering certain severity factors would often result in a decrease in others (e.g. an attack that targets critical infrastructure would increase the scope and severity of the impact experienced but may actually decrease insured loss due to the assumed critical infrastructure exclusions). Therefore, when implementing counterfactual findings, we had to take care to consider how the alteration of different severity parameters would influence the overall event.

Parameterisation

In acknowledging the difficulties in parameterising cyber events, we first identified a longlist of relevant parameters and then prioritised modelling those which we have observed as being most critical. Our final parameterisation output gravitated more heavily towards capturing downtime and thus, business interruption.

We opted to focus on fewer parameters but with more detail to ensure a deeper understanding, practical implementation and more meaningful insights. This approach allowed us to delve into each parameter. Whilst external scanning data or other technical data could potentially enrich the modelling process (e.g. by seeing how many companies might be exposed to certain vulnerabilities), we decided against its implementation in this iteration to maintain focus and clarity.

Review of cost components

One of the key focus areas for the Partnership was to ensure that real-world claims data²⁸ was used to inform the way that losses were parameterised and distributed both per event and per insured. To develop an accurate cost model, we included an in-depth analysis of existing claims data (predominately Beazley intellectual property) and also engaged with subject matter experts.

The cost model categorised key losses into two main buckets:

- Business interruption (BI) costs: encompassing both direct and contingent interruption, capturing the impact of the attack on the company's operations and revenue.
- Additional costs: includes expenses related to extortion, data recovery, notification, and incident response (including forensics).

There are inherent obstacles that can introduce inconsistencies and unknowns into our cost components. One major challenge is the inconsistency in claims data, as different sources may use different reporting methods or categorisations, making it difficult to compare and aggregate the data accurately. Additionally, there is an inherent challenge of transposing attritional claims data into a systemic cyber event and assuming the same principles/values will remain accurate.

Furthermore, limited insight into certain aspects of incident response services can make it challenging to accurately estimate costs. An example of this is both first and third-party legal and regulatory defence costs; not only did we face challenges in sourcing reliable data to accurately estimate these costs, but we ultimately determined that their contribution to the overall losses was minimal and thus, we excluded such legal costs from our analysis. Likewise, we have determined that for single-risk (non-CAT) incidents, ransomware payments may form a large part of standard losses; in an accumulation event the infrastructure to administer payments and many unique decryption keys is likely to fail as there is unlikely to be sufficient resources to process the volume of tens or hundreds of thousands of demands. Whether there was ransom demand or not, we have effectively considered this as malware without decryption. Therefore, we did not include ransom payments as a separate cost type, but rather saw them forming part of the data restoration costs within our 'additional costs' component.

We also decided to exclude those types of physical damage potentially recoverable under affirmative cyber coverage (e.g. computer system replacement (bricking)) from our additional costs, as a separate element. When evaluating the impact of hardware damage, we determined that hardware replacement fell under the category of data recovery costs. This is because only severe data loss would necessitate the complete replacement of physical assets, and usually only if such cost was less than pure data restoration.

Efficacy of security controls

We underwent a detailed process to ensure that we accurately captured and implemented reference to security controls, incident response and data recovery within our own narratives and model. To achieve this, we engaged in discussions to assess the effectiveness of different controls and the likelihood of control failure. These discussions helped us determine the success rates of attackers, represented as percentages, within the model. Additionally, we took into account the unique characteristics of different industries and companies, considering their response capabilities and the controls they were likely to have in place. To capture this variation, we adopted a seven-tier risk categorisation approach (RC1 to RC7), where companies in lower numbered categories were deemed to have a more mature security posture and were therefore less likely to be successfully impacted by attackers.³² This approach allowed us to incorporate industry-specific nuances and provide a more nuanced representation of the likelihood of successful attacks within our model.

Scenario development challenges

Ensuring longevity and relevance in cyber event narratives posed a significant obstacle; by focusing on the underlying principles and common vulnerabilities inherent in cybersecurity incidents, we ensured that our narratives remained applicable across various industries, software platforms, and technological advancements. This not only enabled us to future-proof our narratives but also ensured their adaptability and utility across diverse contexts and scenarios.

Another significant obstacle that we encountered was developing scenarios that could scale to a loss that was significant enough in the insurance industry to warrant consideration. This was difficult, as a number of narratives we considered equated to large economic losses or individual industry or company losses but could not scale to the level of insured loss that would be relevant. For the three selected scenarios, we had to consider what existing security controls would need to fail, and in what order, for the events to reach a significant financial impact. In addition to this, we had to consider how much of the impact would be a result of unintended consequence (i.e. an attack that spread further or had a greater impact than even the threat actor originally intended).



Description of Selected Scenarios

We created three narratives that illustrate potential ways in which substantial insured losses could arise from malware events. When selecting names for the three scenarios, we opted for three Greek myths that effectively conveyed the narratives we were aiming to portray through our scenarios. Greek mythology not only presented us with an opportunity to be creative, but is renowned for its symbolism.

Autolycus

Widespread Software Supply Chain Attack (WSSC):³³ named after a successful robber who had the power to metamorphose or make invisible the things he stole, this concept aligned closely with the attack that inspired the scenario (SolarWinds 2020), in which sophisticated coding was utilised to hide the data exfiltration process. This scenario investigates the hypothesis that the loss could primarily arise from a targeted supply chain attack impacting a subset of companies using a compromised software product. This leads to a large number of affected companies, with losses incurred being relatively high per company.

Lernaean Hydra

Self-Propagating Malware (SPM):³⁴ named after a serpentine water monster with nine heads and thus the ability to attack multiple targets at the same time. This scenario, inspired by the WannaCry incident, examines a vicious and uncontrollable attack, based on the hypothesis that the loss could primarily stem from the widespread distribution of a self-replicating 'wormable' malware, impacting a larger footprint of organisations when compared with our other scenarios, but with relatively lower loss incurred per insured.

Demeter's Curse

Targeted Industry Loss Event:³⁵ according to mythology, Demeter laid a curse on the world that caused plants to wither and die, rendering the land desolate. This represents a direct and targeted attack to a valued critical resource and thus, aligned with our narrative, and that which it was based on – the Colonial Pipeline cyber-attack of 2021. This scenario tests the hypothesis that the loss could primarily result from a small number of significant losses for individual firms in specific sectors. The event, stemming from a targeted malware attack, results in higher loss per insured when compared with the other scenarios.

Detail on the scenarios can be seen in the supplemental material.

Model Development

Exposure set

Through their underwriting and portfolio management approach, carriers can have varying exposure across geographies, industries, and insured revenues (size) along with many other technographic and firmographic factors. The nature of using an RDS approach means that the results from a specific RDS may be material on one portfolio but not on another (e.g. a US-focused event may have reduced impact on a European portfolio).

When considering the results of such an exercise it is important to ensure the scenarios are relevant to the portfolio under assessment. As part of this project, we also wanted to consider the effect of portfolio composition on the resultant losses and explore the nuances of different events of typical market portfolios.

Gallagher Re has developed an Industry Exposure Database ('IED') consisting of ~1.2 million policies and \$13 billion of premium. This is based on data from over 70% of the insurance market and scaled up based on market expectations at the time of development. A representative sample of the Gallagher Re IED was used to construct an industry loss view for the scenarios developed.

To support this analysis, we constructed a sample, synthetic portfolio of 20,000 policies with similar proportions of industry, country, and revenue to the Gallagher Re IED. This portfolio represented exposure seen in the insurance market. This analysis focusses on standalone, affirmative cyber cover.



Summary of portfolio

Premium and Rate on Line ('ROL') by industry:

Industry	Premium	Rate on Line	Count
Agriculture, Forestry and Fishing	2,100,000	1.14%	287
Education	2,200,000	0.86%	422
Entertainment and Recreation	3,300,000	0.37%	677
Finance	28,600,000	1.41%	1,555
Healthcare	12,500,000	0.94%	1,495
Information Technology	15,500,000	0.84%	1,834
Manufacturing	24,200,000	2.24%	938
Mining & Primary Industries	300,000	1.62%	15
Miscellaneous & Unlisted	22,400,000	0.88%	3,251
Oil & Gas	2,500,000	1.26%	292
Professional, Technical and Business Services	23,400,000	0.72%	5,729
Public Administration and Non-Profit	8,500,000	1.21%	1,048
Real Estate, Property and Construction	4,700,000	1.11%	639
Retail & Wholesale Trade	14,900,000	1.21%	1,092
Telecommunications & Media	500,000	0.83%	64
Tourism & Hospitality	100,000	0.25%	37
Transportation & Logistics	8,000,000	1.62%	545
Utilities & Energy	1,300,000	1.64%	80
TOTAL	175,100,000	1.05%	20,000

Premium and ROL by revenue band:

Segment	Premium	Rate on Line	Count
Cat Size 1: 0–20m revenue	38,300,000	0.34%	17,827
Cat Size 2: 20m–100m revenue	18,600,000	1.16%	1,256
Cat Size 3: 100m–1b revenue	37,800,000	2.17%	576
Cat Size 4: 1b–10b revenue	51,300,000	3.68%	249
Cat Size 5: 10b and above	29,000,000	5.01%	92
Total	175,100,000	1.05%	20,000

Premium and ROL by region:

Segment	Premium	Rate on Line	Count
North America	128,800,000	1.05%	12,558
Europe	35,400,000	1.06%	5,447
Asia	2,000,000	0.95%	240
Other	8,900,000	1.00%	1,755
Total	175,100,000	1.05%	20,000

Top-down approach

Estimating bounds and necessary conditions for defined losses

To provide some bounds on the modelling and provide a high-level sense check on the outputs, we performed some simple analysis to investigate key parameters required for certain size loss events. The example below focusses on a 100% loss ratio industry loss event, or roughly \$14b.

For any scenario, the loss amount depends on the ‘footprint’, number of companies impacted, and the severity of loss. The simple top-down approach facilitated discussions about the required size of loss to generate a material impact on the cyber insurance industry.

The Gallagher Re IED was used to provide benchmarking around the policy count and average limits deployed, while expert judgement was used for the reasonableness of the average losses. We note however that at the micro / nano end of the market, the insured limit can sometimes be many multiples of insured revenue and therefore unlikely to be exhausted or suffer a full limit loss.

Note that the revenue bandings are different to the above due to the availability in the data set.

Size	Revenue	Implied market extrapolated from IED		Total Limit Deployed
		Count	Avg limit per policy	
Large	>\$1b	12,000	27,000,000	300,000,000,000
Med	\$250m-\$1b	24,000	4,500,000	100,000,000,000
Small	\$10m-\$250m	150,000	1,700,000	230,000,000,000
Micro	\$1m-\$10m	300,000	800,000	240,000,000,000
Nano	<\$1m	900,000	250,000	220,000,000,000
TOTAL				1,090,000,000,000

To achieve a 100% market loss ratio, the below table shows the required number of assumed ‘average’ losses by revenue band assuming a split of loss across revenue in the same proportion as total limit deployed. We considered this under the three scenarios we defined, described in more detail below.

Size	Average Loss			Implied Number of Avg Losses			Insured Footprint Required		
	Scenario 1	Scenario 2	Scenario 3	Scenario 1	Scenario 2	Scenario 3	Scenario 1	Scenario 2	Scenario 3
Mega	–	–	300,000,000	–	–	47	–	–	N/A
Large	20,000,000	25,000,000	0	196	156	0	2%	1%	0%
Med	3,000,000	5,000,000	0	420	252	0	2%	1%	0%
Small	500,000	750,000	0	5,877	3,918	0	4%	3%	0%
Micro	100,000	150,000	0	30,505	20,337	0	9%	6%	0%
Nano	50,000	75,000	0	56,790	37,860	0	6%	4%	0%
TOTAL				36,998	24,663	47	7%	4%	

For Scenario 3, a targeted industrial attack, we intended to focus on industries with a concentration of a few very large companies with tens to hundreds of billions of revenue and introduce ‘Mega’ here as a subset of ‘Large’ companies (revenue >\$1bn). This was to consider the impact of fewer but very high severity losses.

While this was a very rough analysis it highlights the fairly significant percentage of companies which would need to be impacted to start producing a bad earnings event. This becomes a good anchor to ensure the scenarios under consideration were extreme enough to cause relevant impact to the insurance industry.

Model construction

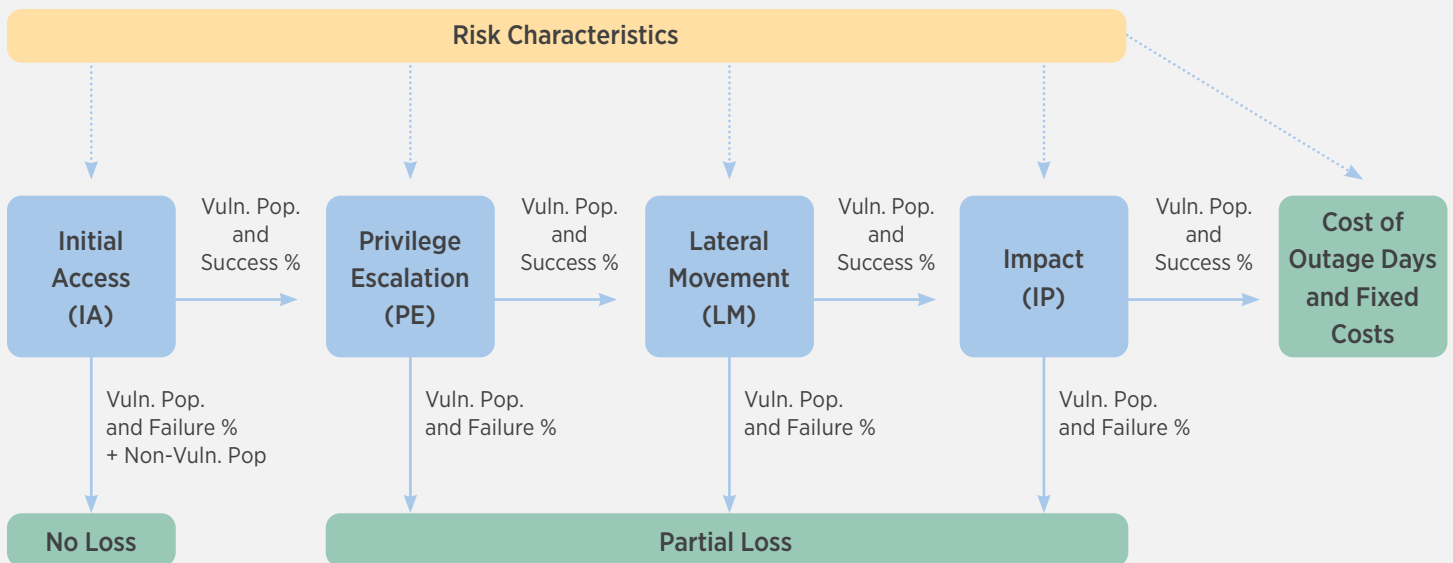
In order to assess the scenarios from an economic and insurance perspective, we have built a model to calculate the (insured) loss for the representative worldwide cyber portfolio discussed above. The model is deterministic in the sense that we assume that the described events have happened and we only calculate the loss amount.

We consider three different contributions to the final loss:

- 1 Full losses:** losses where the attack was successful and there are BI costs, alongside additional costs (e.g. response costs).
- 2 Partial losses:** losses where the attack was not successful but there are still some costs (e.g. response costs).
- 3 CBI losses:** losses from risks not directly affected by the malware but dependent on those which are.

As in all cyber models, we consider two model components to come up with the modelled loss. Firstly, the footprint (who is impacted) and secondly, the severity (the loss amount in case of impact). As both of these components depend on the underlying risk, we calculate both as a function of the risk information (industry, size and country). Specifically, based on historic experience each combination of Size and Industry was assigned to one of seven Risk Categories (see Appendix 3 for more detail).

We start by calculating the probability that a company is impacted by an attack and multiply this by the size of the loss if the company is affected. As we have a large portfolio on average this should lead to the same results as a simulation of random attacks with the same assumptions.



Footprint

Full losses

The footprint is the percentage of the overall population that is impacted; for this we are using a condensed MITRE framework attack path, shown in the figure above. This attack path has four steps and each step has a specific percentage of the population which is exposed to/attacked in that step and a likelihood that the attack is successful for this group. If an attack runs successfully through all four steps of the attack path, we have a BI loss, the likelihood of which is just the multiplication at each step of the percentage which is vulnerable by the percentage for which the step is successful.

Partial losses

The footprint in the calculation of full losses only considers risks where the attack was successful and leads to a BI loss. We are omitting the risks where some steps of the attack path were successful, but the malware did not lead to a BI loss. These risks might also have costs, for example incident response costs. We reflect this in the loss calculation by allocating a cost at each step of the attack path for attacks that fail to make it further through the attack path.

CBI losses

For the CBI losses, we assume that a number of risks are indirectly impacted. Further, we assume that the percentage of the indirectly impacted risks is proportional to the percentage of directly impacted risks. We have used 20% in our model which is more fully discussed in Appendix 3.

Severity

Once we know the percentage of infected risks, we estimate the loss per risk. It consists of two parts; the BI loss and additional costs (e.g. incident response, data restoration). The latter costs are a fixed number depending on the industry and size of the risk but the BI losses depend on several factors that are a function of the revenue and the industry.

Full losses

The loss for a specific company is calculated by firstly deriving the from ground-up ('FGU') loss and then applying the insurance structure, limit, attachment point and deductibles.

The FGU loss consists of two loss components – additional costs and the BI loss.

Full losses = additional costs (based on industry, size) + BI losses.

The additional costs are a function of the industry and the size and are uniform across all three scenarios. The full amount of these costs is applied to all risks where an attack successfully causes BI.


The BI costs are the product of the daily revenue, the percentage of daily revenue lost, the gross margin rates, and the number of outage days.

BI loss = daily lost profit x BI days (based on industry, size).

BI loss = daily revenue x BI revenue lost factor (size) x gross margin rate (industry) x BI days (industry, size).



Taking only the average outage days when calculating the severity would understate potential losses, even though this would be the average BI loss. Since we then later apply the insurance structure on the loss, if we took the average loss, higher layers would show no potential contribution to the loss. This would not be accurate, as across the entire portfolio and with the variation within the BI loss, we expect some losses to be much higher than the average and therefore incur insured losses to even higher layers. This issue is demonstrated on the left of the figure below.



Therefore, we are performing a 'quasi-simulation' to also reflect larger losses which impact a given layer.

To do this, it is necessary to define an expected distribution of BI outage days and, by extension, the severity of loss. We assume a LogNormal distribution of the outage days which is often used in Property and Casualty (P&C) insurance for claims severity. To define a LogNormal distribution two points are required; for example, we can set the average value as well as the value at which only 10% of outages are longer. The distribution of the outage days needs to be defined for each combination of the three events and the seven Risk Categories, so we require 21 LogNormal distributions in total.

By using the inverse Cumulative Distribution function of the LogNormal distribution, we can return the outage day value. This value has a given probability (probability of the outage being smaller than the outage day value). We can use this to calculate height of the loss (yellow bars in figure above) at each probability step by multiplying it by the revenue per day, derived from the annual revenue and the gross profit margin for the industry and BI revenue lost factor.

Finally, taking a spread of probabilities between 0 and 1, (we used 0.03, 0.97 and the values between 0.05 and 0.95 in 0.05 increments), delivers a discrete approximation of the (continuous) distribution to produce the quasi-simulation. These steps were chosen to closely approximate the expected loss of the distribution but this is not a prescriptive requirement and the step choices could be varied according to modelling objectives or even adapted to produce a true random sampling of the full distribution. However, as discussed, a full stochastic simulation is a complexity step beyond the purpose of this project and given the uncertainty associated with the modelling, the increased accuracy of such an approach is arguably not a beneficial trade-off relative to the complexity.

In each case of a Full Loss we assume the additional costs also apply, so these can simply be added to the expected loss of insured before the application of the layer structure.

$$\text{Impact to Layer}(i) = \min \left(\text{Layer Limit}, \max \left(0, F_x^{-1}(i) * \text{revenue per day} * \text{gross profit margin} * \text{BI loss factor} + \text{additional costs} - \text{Attachment Point} \right) \right)$$

where $\ln(X) \sim N(\mu, \sigma^2)$ and $F_x^{-1}(i)$ is the inverse cumulative distribution function of X at point i , parameterised for each policy based on the specific risk characteristics of the insured. The expected loss to the layers may then be computed by taking the average of the layer impacts at the 21 different probability points ($p=0.03, 0.05, 0.10 \dots 0.95, 0.97$).

$$\text{therefore Exp Loss to the Layer} = \frac{1}{21} \left(\left(\sum_{k=1}^{19} \text{Impact to Layer} \left(\frac{k}{20} \right) \right) + \text{Impact to Layer}(0.03) + \text{Impact to Layer}(0.97) \right)$$

By summing up the expected loss of each policy in the portfolio this gives the expected loss of the portfolio for the given event.

CBI losses

The severity of the CBI losses is a function of the average BI losses of the infected risks. We use the same input as for the BI calculations above, but do not perform the quasi-simulation. Instead, we apply a CBI severity factor at the end. We set this factor to 50% since we assume that there are some interruption losses but as the risks are not infected by the malware (so they do not incur costs in the same way as risks that are directly compromised do) these losses are probably smaller and varied depending on how reliant companies are on suppliers.

Severity = CBI factor x average annual daily BI of the same company if it were affected by the malware directly.

After that, we apply the insurance structure and then have the expected insured losses.

Partial losses

For risks which were attacked but not all steps of the attack path were successful, we still assume that there are some losses. Depending on how far the attack was successful, we distribute parts of the additional costs to the corresponding steps of the attack. We then apply the insurance structure on these losses to get the expected insured losses.

Scenario Quantification

In order to run the model, there need to be parameters. These were probably the most challenging elements of the project to determine, as we are trying to model events that have not happened yet. Therefore, the parameters and associated values should be assessed with caution.


We have summarised in the table below all parameters that go into the model and stated where they are used, how we derived them and how sensitive the model is to changes in them. In addition, we have also stressed some of the parameters specifically to understand how this impacts on the model results.

Parameter	Derived from	Sensitivity	Comment
Regional spreading parameter	Conservative estimation	Depending on portfolio and starting region of the event; up to 1-1.	The idea behind this parameter is that geography has an impact on the model. For example, there is an assumed natural time lag between geographies (e.g. if a malware event starts in the US in the morning, it is already afternoon in Europe and night in Asia and so fewer computers may be affected). Geography may impact in other ways, for example regional differences in software used, limiting the scope of potential of compromise (see further detail on geography in the supplementary material).
Risk category	Expert judgement and underwriting data (risk assessment)	The risk category has an impact on the infection probability and the outage days.	A risk specific modifier that is applied on every risk depending on the size and industry.
Fixed additional costs	Loss and pricing data	The additional costs are part of the final costs and therefore have an effect smaller than 1-1.	The additional costs include notification, cyber extortion, data recovery, incident response (including forensics) and other costs but exclude BI and CBI costs.
Attack path split of additional costs	Expert judgement	Small impact since it only influences the partial losses (which are a small portion of the overall losses).	In case only certain attack path steps are successful but don't lead to a BI loss, there will still be costs (e.g. incident response), that are dependent on how many steps of the attack path were successful.
CBI parameters	Expert judgment	Has a 1-1 effect on the CBI losses which are part of the overall losses (not the majority).	We apply a factor on the footprint and the average severity and then calculate the loss per risk, for the risks which are not directly impacted by the malware but are impacted via CBI.
Gross margin rates	Derived from data (external and publicly available)	Has an impact on the revenue at risk which is a factor in the BI/CBI loss calculation.	Since insurance policies usually cover only lost profit and not lost revenue, we apply these rates here.
Percentage of daily revenue lost	Expert judgement	Has a 1-1 impact on the severity, which has a smaller than 1-1 impact on the final loss (due to insurance structures).	It is assumed that for larger corporates, in particular, parts of the company are still productive in case of a malware event – this is reflected here.
Industry-specific event	Expert judgement	Has a strong impact on the results depending on the portfolio as some industries are not affected.	For the Demeter's Curse scenario, which is an industry-specific event, we have chosen specific industries which are hit in this case. These selected industries are those that are perceived to have a greater take-up of operational technology (OT) and IIoT devices.
Footprint	Expert judgment	Has a 1-1 impact on the final loss	The footprint is derived from the four steps of the attack path.
Outage days	Expert judgement supported by loss data	Has an impact smaller than 1-1 on the final losses (due to insurance structures)	Based on some loss data. However, this has to be considered with caution since the actual loss data does not come from systemic events but from single attacks.


Results

Running the portfolio in the model for the three scenarios and scaling it up to a market loss by the premium (multiply the loss for the portfolio by the ratio market premium/portfolio premium where the market premium is assumed to be \$14b) leads to the following results (insured market loss and market loss ratio):

Insured market losses (in USD)



The insured loss in the chart above is the sum of the direct losses (companies where the attack is successful and results in a BI loss), partial losses (companies which are attacked but the attack is not completely successful and therefore leads to no BI loss) and CBI losses (companies which are not attacked but suffer a loss as an indirect consequence of the attack on suppliers). The split of the loss contributors in the three events is depicted in the pie charts below.



The larger two scenarios lead to a market loss around 200% and 250% loss ratio. The biggest loss contributors are direct losses followed by partial losses. We see these events as unlikely but still plausible. Therefore, the loss ratios sit in the tail of the distribution of potential portfolio losses.

It is important to note that each of these events are independent so there is not a need to stack them with one another.

The final scenario (Event 3) is the smallest, since it has a much smaller and more targeted footprint and is restricted to industries heavily reliant on operational technology. This aligns to our finding from our top-down analysis, detailed in the previous section – to produce significant loss at a market level, large losses to individual companies alone is not sufficient; a significant footprint of the event is also required, potentially spanning multiple sectors. The wider spread and fairly industry-agnostic footprint of the larger two scenarios outweighs the more significant outage times assumed in the final scenario. This could also be accentuated by the fact that the layer structure of the insurance industry limits the impact of per risk severity to any one insurer.


For the first scenario, the assumption is that the compromised software has highly privileged access, therefore, a higher percentage of the attacked companies go on to suffer BI losses resulting in relatively lower partial losses. In the second scenario, there is a larger set of vulnerable companies, but the sets of Initial Access, Privilege Escalation, Lateral Movement and Impact are independent per company therefore, many companies are infected but do not experience a BI loss, resulting in a higher percentage of partial losses. The final scenario also highlights the potential for successful Initial Access which fails at a later step of the kill-chain without causing major BI loss, in this case because Lateral Movement into the OT environment is considered more difficult, particularly in an automated way. This results in a high percentage of partial losses (e.g. incident response and system restoration) and relatively low direct BI losses.



Impact of portfolio composition


As with all models, not only are the results sensitive to the parameters chosen but also the portfolio entered. For example, in the sample portfolio we used, companies with revenue under USD 1 billion make up 54% of the premium. However, as shown below, these smaller companies disproportionately contribute to Event 2 (Lernaean Hydra (SPM)). By comparison, in Event 1 (Autolykus (WSSC)) larger firms are over-represented, as it is assumed they will have a faster patching cadence and more security software so are more likely to be caught in this attack, even if it doesn't lead to a BI.

Contribution to Insured Loss by Revenue Bracket



Similar effects can be seen for primary policies vs excess layers, in this case primary policies provide roughly 65% of the premium used but contribute 80% to 90% of the loss amount, as partial losses tend to only impact primary layers. These two effects are likely interlinked as there is a higher chance that within the portfolio the primary layers belong to smaller firms.

Contribution to Loss by Policy Layer



This is particularly relevant when we consider the scaling to a market loss, as any deviation between the sample portfolio and the market composition, including in these two attributes, will distort the veracity of the final market loss figure calculated. As will any change in the rate levels between the sample portfolio and the market, or a misestimation of current market premium. All these uncertainties need to be kept in mind when considering any market loss figure derived from scaling up a sample portfolio.

Sensitivity testing

Since the parameters, especially the most crucial ones like footprint, have a strong impact on the modelled results we performed a stress test to better understand the effect of parameter variation. We applied four different stresses (we were only testing what would happen if the parameter choice underestimates the losses):


The footprint is too small (increase of the footprint by 50%).

The risk quality is too optimistic (all risk categories reduced by 1).

The outage time is too short (increase outage time by 50%).

The additional costs are too small (increase the additional costs by 50%).

In the chart below, there are the results on the overall losses for the different events with the different shocks.



The footprint shock has an impact that is increasing the overall losses to the same extent as the shock, so a 50% increase in the footprint also increases the overall loss by 50%.

An overestimation of the risk category of the specific risks in the portfolio has different effects depending on the scenario. For Event 1, a lower cyber security is reflected also in lower patching cadence and therefore reduces the footprint of this event. For the other events, the effect of overestimating the cyber security has an impact of an increased footprint that then shows in the increased overall losses.

The outage time has a smaller effect than the footprint, since it only affects the CBI losses and the direct losses. Partial losses don't have a BI and therefore do not react on the increased outage time.

A shock on the additional costs increases the losses in all three events and has an impact on the direct and the partial losses. CBI losses are not affected, therefore, the effect is smaller than the shock on the footprint.

For a more details on the sensitivity analyses, please see the Appendices.

Conclusion

A clear finding of this project is that there is still a large amount of uncertainty and complexity in relation to systemic cyber-risk and the insured losses that might accumulate from such an event. We strongly encourage all parties attempting to quantify cyber accumulation risk to use this model to compare and communicate the results relative to their own modelling results and assumptions. Only through robust technical debate and a broad range of voices can we drive the understanding of such risk forward.

Despite the significant uncertainty that remains, the Partnership hopes that the modelled results presented in this paper offer some reassurance that a significant malware event would be survivable by the cyber insurance market. The loss ratios generated by the model are significant commensurate with the intended severity of the modelled scenarios, but the Partnership believe that they are not large enough to be ruinous for the market. It is our sincere hope that this third-party, open source contribution will deepen the commitment of existing providers and encourage potential new providers of cyber (re)insurance capital into the market. The dynamic nature of cyber-risk means that capital is an essential ingredient for enabling growth in the cyber insurance industry, particularly as technology continues to be further integrated into society.

Project learnings

Having successfully completed this year-long, collaborative research project, the paper authors wish to share the following reflections as a challenge to our peers and colleagues engaged in modelling cyber accumulation risk:

- It is hard even for experienced modellers to visualise and consider averages during extreme events. In particular we needed many rounds of parameterisation to reach a consensus on the outage times.
- There are some simple, assumably measurable attributes, such as percentage of revenue lost in an IT outage, which have not been widely measured in a systematic way.
- Risk quantification in insurance relies heavily on return periods, but these present a significant challenge for a dynamic peril such as cyber.
- The universe of events of significant magnitude, when taken to general abstraction, are similar. However, patient and persistent actors may be able to achieve compromise on a surprising scale as witnessed by the XZ Utils backdoor which occurred as we were writing this paper

Relevance to the ILS market

One of the most visible uses of cyber modelling in 2023 was the pricing of insurance-linked securities (ILS), which are also known as catastrophe ('cat') bonds. The publicly tradable cyber-cat bond market totalled some \$415mn as of 1.1.2024. Prior to 2023, the ILS market was confined largely to natural catastrophe perils, and so the growth in this area as a source of capital to the cyber insurance market has been rapid. ILS are priced using a vendor model as a view of risk, sometimes with a second model as an alternative view. The model presented in this paper is not a substitute for a commercially developed model but offers an alternative contribution in helping describe a plausible magnitude for modelled losses from an extreme cyber event. While there have been attempts to produce such a figure in academia³⁶ and via Lloyd's Futureset,³⁷ to the best of the Partnership's knowledge, this project is the most detailed insurance industry collaborative attempt to produce a worst-case whole-market loss that is underpinned by the structure of the insurance market.

Looking to the future

Having completed this year long journey it is natural to ask what is next, as a group the Partnership thinks the following steps would be beneficial to consider in the future:

- In the short term there is an intention to collect and consider any and all feedback from interested parties; the best way to share, respond to or incorporate such feedback is an open question.
- It is also possible that others take this framework and publish amendments better reflecting their view of the world.
- Now there is a simple model framework in place, it is natural to consider spending significant time refining specific parameters with the support of relevant experts
- A clear extension of this collaboration would be to consider another family of models such as cloud-outage or mass-hardware failure
- Focusing on a simple modelling approach has precluded considering a range of return periods, if this becomes important this could be a next step. But given the complexity of building such a model such work may be limited to refining parameters or working with adapting an existing model.

Invitation for comments

The Partnership now invites all stakeholders, clients and other industry professionals to submit comments on the Partnership's first iteration of the malware event set and associated modelling approach.

The goal of this paper is to provide all interested parties with an in-depth and transparent view on how the Partnership developed and subsequently modelled a malware cyber-attack event set. We therefore welcome your feedback and commentary in the hopes that it will highlight room for growth, hence allowing for us to ensure that this work is as successful and useful as intended.

You can submit your feedback here: CyberSystemicPartnership@gallagherre.com. We remind respondents not to submit any commercially and/or competitively sensitive information in feedback, in order to ensure competition law compliance.

Case Study — Use Case for Model: July 2024 CrowdStrike Faulty Update

The model presented in this paper was designed to contemplate systemic cyber events that are remote, and accordingly the framework may appear relatively abstract in relation to the type of cyber events the world has seen up until the point of writing. However, the framework has been carefully designed to be flexible and can be readily customised to model different views of risk and parameters and for new scenarios to be modelled.

A case in point is the 19 July outage of Windows machines running the CrowdStrike Falcon endpoint detection/response sensor. A 'channel update' (similar in concept to antivirus definitions updates) caused the CrowdStrike software to enter an error state owing to a mismatch between the data in the channel file and the input expected by the CrowdStrike Falcon sensor software.³⁸ As is common with security software, CrowdStrike Falcon runs with significant privileges. This means that rather than generating an error message such as a dialog box popup, under certain conditions the operating system on the affected machine crashes totally and displays an error message (commonly called the 'Blue Screen of Death'). Under normal circumstances, these errors are rare and do not repeat on a correctly configured system, however in the case of the CrowdStrike update, this was not universally the case with Microsoft estimating 8.5 million devices were impacted.³⁹ Removing the problematic channel file resolved the issue, but the speed with which this could be achieved across an insured's IT environment has potential to vary significantly depending on multiple factors such as physical access to the affected computers, number of available IT support staff, ability to assist users self-resolve to name but a few

A natural question for readers to ask is whether the model presented in this paper could be used to estimate the range of losses for a specific portfolio from the CrowdStrike outages. We describe below how the model parameters might be adjusted to consider events involving a widespread, simultaneous outage of business-critical IT services. It is important to understand that the results are highly dependent on the extent of outage experienced by individual insureds and The Partnership offers no opinion or indication in regard to this, nor should any of the following analysis be viewed as indicative of any potential losses or claims experienced by any of the Parties. The methodology outlined below is intended to assist users of the model in understanding what parameters in the model would deliver comparable results to either their own or third-party loss estimates. A prerequisite is having a view of the resilience of organisations at least based on firmographic parameters to form these estimates.

Regional lag

The regional lag of an event can be reflected in the model. The faulty CrowdStrike update was pushed to customers during the daytime for the Asia-Pacific region, meaning Asia could be chosen as the start region, reflecting the time lag with a fix available by the time the US entered work hours. The factors which will determine the spreading for the other regions can be adjusted in the parameter sheet. In the three events described in this paper, we gave all 'non-beginning' regions a spread factor of 75% which means that the footprint is 25% reduced compared to the region where the malware is observed first. For CrowdStrike, the user can change these parameters to reflect their own understanding of the spread.

The model framework also supports relatively simple counterfactual analysis to be performed here by instead considering the impact had the event started during US daytime, potentially increasing the Business Interruption.

The four steps of the kill chain (Footprint)

Each of the four steps of the kill chain can be parameterized independently in the model. In the CrowdStrike event, the blue screen appeared after the system was patched. Therefore, we would only see the initial access as a driver of the model and set all other steps to 100% (which is a conservative approach and can be questioned). In the Initial Access, different risk classes were impacted with a different footprint, and this can be reflected in the corresponding parameter selection. CrowdStrike is understood to be more prevalent in larger insureds by revenue and their customer base also varies by country which could be reflected in the exposure parameter choice.

BI time

One key question in determining potential insured losses is the length of outage experienced by the insureds and the proportion of revenue lost to IT service disruption. The variance in BI times is expressed in the model by varying risk category as a proxy for operational resilience. By the nature of how the losses are modelled, the potential for prolonged outages is reflected in the model (assuming a LogNormal distribution for the outage days). The tail volatility can be captured by altering the outage length distribution. Interested parties could conduct sensitivity tests for different outage lengths and ranges to build a loss distribution.

The current model does not consider the impact of time based 'Waiting Periods' which would remove losses which fail to reach a minimum length, instead it only applies the monetary deductible. This is not material when considering the extreme tail events in the rest of the paper, but may become more relevant in situations such as this case study where the outage length is shorter and may fall within the Waiting Period. As it stands the model construction can be viewed as conservative on this point.

Additional costs

Additional costs can easily be varied from the default assumptions used in the model, including on a granular basis by size/industry bucket. This allows for different costs in recovery for different risks to be reflected as desired.

Risk classes

The model also allows for reallocation and redefinition of risk classes. If only differentiation between size of companies and not between industries is desired, risk categories can be defined only for the size (eg RC1 is corporates above 10 bn USD revenue, RC2 is corporates between 1 and 10 bn USD revenue, ...) and then set the parameters (for BI time and the footprint) accordingly. If some industries are more impacted than others, this can also be reflected in the tool by choosing a risk category for only one industry and then setting the parameters accordingly.

Summary

As described above, the parameters of the tool can be modified to calculate the potential range of losses from an event like the CrowdStrike outages. However, estimation of the parameters is subjective and will be reliant on data available to each user and their own expert judgement.

Project Limitations

This project has a number of primary limitations – these are grouped into themes below and shared for transparency:

Scope

The project has not covered those items deemed out of scope, as per the Introduction. For example, because we are focusing on sudden systemic cyber events, we are not considering events that could involve losses associated with a vulnerability occurring over a long period of time.

Scenarios

The use of scenarios is a common way of providing structure when thinking about possible events and losses. The scenarios detailed in this paper are a small number of indicative ways that large levels of insured loss could be generated from malware events. They are not the only ways that this could occur.

Simplification of the attack framework

We have used a condensed version of the MITRE framework to structure our scenarios and model. This simplifies the attack path in our narrative but could result in more of a 'cliff-edge' between stages in the model. Since the footprint for our scenarios is determined by multiplying the impacted population at each stage in the model, it is possible that expanding from four stages into a larger number, could result in different figures for footprint, although whether those different figures would be more or less accurate is unknown.

Parameters covered

Given constraints (e.g. time, complexity, data), only some parameters of a cyber-attack have been considered in the project. In practice, there may be an innumerable number of parameters, with a vast number of possibilities (e.g. a distribution for each parameter at the individual company level). It is not practical to consider or model all of this — simplification and prioritisation have been necessary. The parameters focused on during this project were those that the Partnership believes it is most able to provide a unique viewpoint (e.g. due to expertise and available data from the Parties). Within the parameters modelled, some could be more granular — for example, the 'additional costs' for each scenario are an agglomeration of different costs (e.g. data recovery, incident response) and are given as a uniform figure (dependent on size of firm) across all scenarios, but as an industry we need to collect more granular data to improve the understanding here.

Portfolio used

The portfolio plays a crucial role in the results and the results are sensitive to any changes. We have chosen a portfolio that roughly represents (from our perspective) the worldwide cyber market but this is not guaranteed. The portfolio consists of mostly smaller companies (>80% in count) and is skewed to the US (around two-thirds of the risks are located in the US), which matches our expectation of the market. One major drawback is the data consisting of policies written in the 12 months to the middle of 2022, therefore any structural and pricing changes since then are not considered. This can of course be combatted by the use of more up-to-date data.

Model assumptions

The insured loss results have not been generated using a stochastic model – there are positives and negatives associated with this as detailed previously. For example, it is difficult to estimate or validate the frequency associated with the insured losses generated. There are also a number of other model assumptions made, some of which cannot be back-tested, given the current scarcity of data and the belief of many commentators that we have not yet seen a truly systemic cyber event. In particular it was difficult to determine the length of outage which is likely in a mass event, therefore it was necessary to err on the side of caution in parameterisation.

Modelled losses

The project focuses on insured losses, as that is where the expertise and data from the Partnership can be deployed (e.g. based on claims data). However, some external stakeholders, such as governments, may be more interested in total economic losses than insured losses. Invariably, if the focus was on economic losses, different scenarios might have been used (e.g. focusing on cyber-attacks on critical infrastructure).

Application of Waiting Periods

The model currently only applies the monetary deductible and attachment point. In fact, often cyber policies include a 'Waiting Period' (10-12 hours is typical in Beazley's experience)⁴⁰ which outage periods must exceed before coverage may apply. For the tail events considered in this paper, it is assumed most outages will exceed the Waiting Period and therefore this element does not need to be considered. However, for more granularity or lower severity events it could become more relevant.

Footprint calculation

To get the proper footprint, we would have to run a simulation and draw a random number to decide whether a risk is impacted or not. Since this would make the model more complicated and complex, we decided to use an 'infection rate share approach' which means, that we do not decide for each risk whether it is impacted or not but rather calculate the loss in case it is impacted and then multiply it with the footprint likelihood. Especially for small portfolios, this will lead to too low results since random effects are not reflected but the results are converging for large portfolios. For example, if we have an infection rate of 10% and a portfolio of ten risks, then we would expect one risk being affected. In case we have 2 affected risks (which is not too unlikely), we are twice as high as the number of expected risks. In case we have now a portfolio of 100k risks, we would assume 10k risks being affected. Having 20k risk affected (twice the expected number of risks) is very unlikely. Therefore, the larger the portfolio, the better the calculation approach we are using.

Fixed costs

An assumption was made to consider all non-business interruption losses (termed additional costs) as fixed costs across scenarios. Some of these additional costs were incurred during the attack and some only after a successful attack that progressed through the entire attack path. While there are strong arguments to de-emphasise the third-party costs and ransomware payments, this is a limitation that could warrant more investigation to make the model more precise. This may also reduce the contribution from partial losses which in turn may make the output more constructive.

Extent of consultation

A variety of third-party experts were consulted, alongside many individuals from within the Parties, and there was widespread review and analysis of existing publications. However, no amount of consultation could ever be deemed comprehensive and there will invariably be experts who disagree with elements of the paper.

Falsification

Given the limited history and uncertainty around modelling cyber tail events, there are a number of factors in the risk landscape that could change, which may falsify parts of this study, model, or assumptions. We have considered a non-exhaustive selection of these below:

- Regulatory landscape: Increased, or changing, regulation e.g. compulsory insurance or change in coverage, terms, or minimum standards.
- Cyber attack/defence breakthroughs or maturity: Many of the parameters are based on the understanding of today's cyber security landscape and current trends in attack and defence capabilities. The scenario narratives are derived from knowledge of how attacks may be able to enter or lateral move around systems, and the defences in place to limit this. However, if there is a significant shift in the power balance or capabilities between sides then events may play out in notably different ways.
- Significant change in business mix or rate change: The penetration rate for cyber insurance is relatively low, particularly in SME or non-US segments. The results in the paper are dependent on the overall business mix assumed in the portfolio. Individual insurance portfolios may also differ materially in mix, e.g. SME vs Corporate. Given the assumptions can vary notably between the size of insureds and risk categorisation, changing is mix may change the loss ratios.
- Systemic event occurs: An event occurring would provide data into how events play out, from initial access and propagation, and particularly incident response into potential demand surge or economies of scale with a collective response to secure systems or kill the code. The post event market may also take a different form, with increased penetration and rates, and improved cyber security awareness and posture from the end insureds potentially changing the probability of future events.

Supplementary Material: Scenarios

Autolycus | Widespread Software Supply Chain Attack

Scenario context

A supply chain attack (also referred to as a third-party attack) is when a threat actor infiltrates an organisation's system via a third-party partner/vendor that the organisation uses (e.g. an attacker exploits an IT service provider who has privileged access to their client networks which are managed through an individual portal or system. The attacker uses this to access many of the IT service provider's customers' systems). Whilst not unheard of before, 2017 saw a substantial increase in the number of software supply chain attacks due to Operation Cloud Hopper which was a cyber espionage campaign attributed to a Chinese hacking group known as APT10 (Advanced Persistent Threat 10). The campaign targeted managed service providers (MSPs) and their clients in various countries, aiming to steal sensitive information.

The perceived success of the operation quickly made this a valuable attack technique for organised crime groups (OCGs). The largest supply chain attack to date is the SolarWinds supply chain attack. This was an attack carried out against the Orion software, a widely used IT management tool; threat actors inserted a backdoor into the software updates, allowing them to infiltrate the networks of a number of organisations including government agencies and major corporations. Our own scenario has built upon these historical events and has sought to identify how events such as these will continue to change in the coming years and what this will look like for (re)insurers.

Scenario development

Supply chain attacks continue to rise in popularity as an effective technique for cyber threat actors, as the attacker potentially needs only to exploit a single organisation in order to compromise many more with significantly less friction. Despite this, some insurers are citing limited impact to the insurance industry as a result of these attacks, namely because performing such attacks usually requires a much higher level of skill, attack capability and funding. When developing this scenario, we considered two potential narratives (termed Option One and Option Two):

Option One

A destructive wiperware. This narrative explored the concept that a wiperware payload is released via a malicious update with the goal of impacting endpoints. The attack results in significant business interruption and destroys hardware. This option was far more capable of achieving widespread disruption and loss than Option Two. Critically, we differentiated this narrative from the Lernaean Hydra scenario via considering an alternate attack vector and the prospect of a widespread wiperware that moves without wormable characteristics. Our focus on hardware destruction also helped differentiate the two events.

Option Two

A targeted data exfiltration for espionage purposes. This narrative considered an ongoing, undetected data exfiltration attack conducted via supply chain infiltration. Whilst an interesting attack path, we determined that the manual ('hands-on-keyboard') nature of the attack, paired with the small, targeted footprint resulted in a loss profile lower than that which would be relevant for this project.

Through investigation, we have theorised that of these two primary scenario pathways, a destructive wiperware delivered via malicious update would achieve the largest loss profile. Therefore, we decided on Option One and, the narrative and associated details below follow a supply chain wiperware attack.

Scenario description

A number of permatemp employees within a leading software company's development team are left disheartened after their request to obtain permanent employment status after multiple years at the company is refused again. Employees are individually targeted by a sophisticated Organised Crime Group (with unconfirmed links to a Nation State sponsor) and offered sufficient funds in return for the assistance in conducting their attack. The OCG use these individuals to unwittingly release malicious code within a servicing stack update, with the goal of targeting as many of the supplier's customers as possible. We considered various softwares for this attack, and have parameterised this event around such softwares with a large market share to achieve sufficient footprint and already with enhanced privileges, (e.g. anti-virus software), to increase the potential damage after the successful deployment of a malicious update.

- 1 Initial access:** In the lead up to the attack, the OCG's internal recruits share their knowledge of security controls in their respective environments and any other relevant obstacles. This includes a review of firewalls (e.g., sharing firewall configuration (rules and policies)), intrusion detection systems, company processes and monitoring tools. They adopt relevant evasion techniques, such as using encryption, obfuscation, or anonymisation methods (e.g., running through network proxy servers) to avoid detection and bypass security controls. They also consider the timing for their attack — with the goal to deploy the malware at a time of low staffing/staffing unavailability and with the goal of accessing companies within as many global regions as possible. Utilising their access privileges across development, testing and production environments, the insiders embed malicious code within the update in the regular update channel. The malware acts as a logic bomb, detonating 14 hours after the update is released, to maximise instillation-base, without giving companies sufficient time to identify malicious code before the payload is triggered.
- 2 Privilege escalation and lateral movement:** Software, especially in the security space, often inherently enjoys privileged access within networks. The escalation of privilege and ability to move laterally across the estate is achieved almost by default alongside initial access, with few controls which are able to mitigate it. Once the code is embedded, wormable malware that carries two payloads is released in every device that the update is successfully installed on. The first payload creates a backdoor by changing source file code names (the malware just slightly changes the source file code names to deploy a backdoor). Following the creation of the backdoor, the second payload begins overwriting data on the storage device with random information pre-programmed by the attackers. By overwriting the original content, the data becomes irrecoverable.
- 3 Impact:** The malicious update is successfully installed in a huge number of devices; success of installation is largely due to the update being released by a legitimate channel with a correct hash, therefore, tricking endpoint detection and response (EDR) tools into accepting the changes as a trusted update. The coding is purposefully lean to evade detection. Measures incorporated by the threat actor to limit the malware's spread are ineffective, meaning that there is no specification to ensure that only targeted systems or datasets are hit. The malicious update is installed much faster than initially expected; in the first few hours, the malware works rapidly, impacting emergency services, government laptops and healthcare systems to name a few. For many companies, their hardware is irrecoverable.
- 4 Mitigation and response:** Fortunately, the destructive malware operates in a repeatable manner, which is useful when standardising and coordinating response efforts. Upon identification, it is learned that the malware is quick to reverse engineer. Due to the 'noise' of the malware, the infection is discovered quickly and mitigating controls are marginally more effective than for similar attacks (i.e., SolarWinds). Mitigation controls include ensuring that users are proactively segmenting parts of their networks to limit damage and spread. Upon identification, the software provider immediately releases a statement telling all users to shut down their devices/or refrain from turning them on if they have not already done so. As the initial release took place at 10:00am Eastern Time Zone, the majority of impacted customers are in North America. Australia and New Zealand have a far lesser impact. The software provider releases a statement that identifies indicators of compromise and produces a rapid update that ensures a clean version of the software is being run on customer devices. There is significant downtime and business interruption experienced as users are encouraged not to use the software in the period that it takes to release the patch that overwrites the malicious update. An average of between five and 21 days of business interruption is experienced by companies, depending on their incident response and recovery capabilities. For example, companies in the RC1 – RC3 risk categories are more likely to be utilising services such as Citrix which will drastically improve recovery time. Comparatively, companies who have less IT and cyber security resource (RC4–RC7 risk categories) will be more likely to experience longer downtimes. Those who require mass hardware replacement are most likely to fall into the longer outage length range.

Cost components

Cost component		Applicable(Y/N)	Narrative commentary
Business interruption costs	(Direct) business interruption	Yes	<p>The software company suffers some business interruption in the initial aftermath of the event — this is largely due to an initial shut down and subsequent investigation period in which business as usual operations and functionalities cannot persist.</p> <p>Companies using the software who are impacted by the malicious update also suffer Business Interruption losses as a result of being unable to use their systems and devices until they can be replaced or restored. Companies utilising Citrix, mobile devices or similar services enabling employees to access company data remotely suffer lower business interruption. Larger companies without these redundancies in place are forced to replace upwards of thousands of computers in the aftermath of the attack.</p>
	Contingent business interruption	Yes	In addition, some downstream companies will experience financial impact as a result of disruptions or damages experienced by their suppliers, customers, or other key business partners.
Non-business interruption costs (also referred to as additional costs)	Notification costs	Yes	Notification costs include identifying affected companies, preparing and sending out notifications, and managing public relations and communication efforts.
	Extortion	No	N/A
	Data recovery	Yes	<p>Whilst some customers have protected and available back-up procedures in place, both immediate and downstream data recovery costs are high, with many businesses unable to recover large percentages of business-critical data.</p> <p>The most significant losses are those associated with hardware replacement. Due to the number of companies attempting to replace endpoints, an acute hardware shortage prompts a demand surge resulting in increased business interruption (for those unable to purchase replacements immediately).</p>
	Incident response (including forensics)	Yes	Whilst the software provider had an effective incident response plan in place, and communications with customers was ongoing and relatively transparent, the investigation costs were extremely high, as were the costs taken to fast-track their patching efforts.

Scenario constraints and limitations

A key limitation of the scenario is that it is not designed with a specific software in mind, which makes it difficult to accurately map parameters. In understanding that different software will present varying levels of business functionality (and therefore, varying levels of customer dependency), the scenario may not align perfectly with any particular software. As a result, the mapping of parameters from the scenario to a specific software may not be completely accurate, leading to potential discrepancies in the implementation. Whilst the Partnership was unable to address this limitation in its entirety without detracting from the flexibility of scenario application, our research and expert consultations enabled us to identify a series of different software products that closely matched the requirements and objectives of the scenario. Therefore, whilst not a complete match, we were largely able to tailor the parameterisation approach in a manner that is applicable across a broad range of software providers.

We recommend regular testing and feedback loops to ensure that this scenario does not only remain relevant and applicable, but so that parameterisation can be refined in order to improve accuracy over time.

Lernaean Hydra | Self-Propagating Malware Attack

Scenario context

Since the emergence of WannaCry, along with predecessors such as Conficker and the ILoveYou virus, the mass propagation of malware has consistently posed a significant threat to the digitally interconnected world. The insurance industry has long regarded these incidents as potential catastrophes due to their inherent unpredictability and their capacity to inflict widespread damage, driving cyber losses within a short timeframe until underlying IT vulnerabilities are addressed (through patch releases and applications) or until global response efforts effectively control and mitigate the impact.

A defining characteristic of WannaCry was its utilisation of the EternalBlue exploits, which were state-sponsored capabilities that played a crucial role in its widespread propagation and its impact on the businesses it targeted. Regarding Organized Crime Groups (OCGs), numerous press releases from law enforcement and government agencies have highlighted the increasing sophistication of these groups in terms of their organisational structure. This heightened sophistication enhances their ability to execute ransomware operations with greater efficiency. Modern technology has become a key tool for these groups, with phishing serving as a common initial access vector that has been further augmented by the use of generative AI. This technology enables the creation of more convincing email content with improved language support. While defensive measures strive to keep pace, they often find themselves lagging behind, particularly during periods of increased attacker capability.

In this hypothetical scenario, it was considered what if a perfect storm of events occurred: the leakage and adaptation of a state-sponsored toolkit by OCGs to incorporate into their existing ransomware campaigns, coupled with the utilisation of generative AI-enhanced phishing techniques and additional initial access from the leaked toolkit. The result is the emergence of a highly virulent strain of malware capable of bypassing many existing defences, leading to a surge in ransomware infections that even the initial OCGs cannot adequately address. This surge causes both financial demands and substantial damage within affected organisations. This scenario serves to stress-test incident response capabilities as the industry contends with a surge in demand resulting from mass infections overwhelming its resources for a period of time.

Scenario development

The primary focus of the Partnership's discussion regarding this widespread malware outbreak initially centred around formulating the hypothesis that the individual components of the malware attack path (such as initial access, lateral movement, payload, and impact) were indeed feasible. This involved assessing whether these components met certain criteria justifying the scale of the potential event. Factors under consideration included the automation of the malware upon delivery and its ability to propagate within and across organisations. Unlike the other scenarios developed for this project, only one option was developed.

The Partnership concluded that given the extreme nature of this scenario and drawing upon historical precedents, it was plausible that a combination of these components could be assembled with the intent of orchestrating a mass, untargeted exploitation on a global scale. Such an event would likely exploit previously undisclosed vulnerabilities within operating systems and widely used software, possibly utilising unknown bypasses for standard security measures. Once these exploits became known, swift patching and countermeasures would be anticipated. The ensuing event would trigger a significant surge in demand, overwhelming incident response and IT services within initially infected companies, leading to extended outage periods. Subsequently, a global response effort would kick in, facilitating a more rapid recovery phase through the deployment of patches and the implementation of defences, which could be more readily administered.

Furthermore, the event would catch OCGs off guard with the unexpected effectiveness of the newly deployed capabilities. This would hamper their ability to process ransoms promptly, potentially forcing a temporary shutdown of their operations, particularly if all their cryptocurrency payment addresses were taken offline by law enforcement. Consequently, newly infected victims might find themselves unable to contact the OCGs for ransom payment, exacerbating the impact on businesses akin to a wiperware scenario.

Scenario description

With the Organised Crime Group in possession of more sophisticated tools due to a leak of a Nation State actor, they quickly combine them into their current operations and are able to greatly increase the efficiency of their toolkit.

- 1 Initial access:** With the adaption of the new capabilities, the OCG launches fresh attacks – predominately through phishing. Success is greatly improved with the addition of AI to create more believable emails resulting in more downloads of the now more potent malware.
- 2 Privilege escalation and lateral movement:** Once downloaded and deployed, the malware quickly goes to work undetected and automatically spreads via the newly acquired Nation State capabilities. Once a sufficient footprint has been established the malware beacons back to the OCG so that they can set up the ransomware component.
- 3 Impact:** The characteristics of the impact are largely the same, except the scale and velocity is greatly improved. The malicious update is successfully installed in a huge number of devices. The devices are now encrypted and over time, as the OCG is overwhelmed in response, decryption keys are not registered therefore the malware turns more wiper in nature as no recovery method is available.
- 4 Mitigation and response:** The widespread of the malware, originating in North America, does provide an opportunity for a global response. Initially the spread is unable to be stopped, but as the event progresses the malware and its techniques become more widely known and initial mitigations attempt to stem the tide of more infections. Within days, the initial patches are released which helps to prevent reinfection. Within a week most organisations have at least one form of defence against the threat and have actively mitigated any initial breach they may have suffered. As the OCG actors step away from the attack infrastructure due to its initial overwhelming success, new attacks start to fade within one week, leaving opportunistic copycat attacks from other groups who were slower in merging the original Nation State code with their own. These are less successful due to the defences in place by the global security community.

Cost components

Cost component		Applicable(Y/N)	Narrative commentary
Business interruption costs	(Direct) business interruption	Yes	For every company infected due to the volatile and infectious nature, a short period of interruption will occur as general IT is disrupted. For some companies the infection propagates into more critical IT systems resulting in a longer interruption due to increased clean up and recovery.
	Contingent business interruption	Yes	As companies are affected, given the interconnection and reliance of digital services between companies some larger companies will be affected as key suppliers are disrupted.
Non-business interruption costs (also referred to as additional costs)	Notification costs	Yes	Notification costs include identifying affected companies, preparing and sending out notifications, and managing public relations and communication efforts.
	Extortion	Yes	Initially, some companies affected will be targeted with extortion. As the attacks persist extortion will not be a primary factor as the attack groups are also overwhelmed in their success.
	Data recovery	Yes	With much wide scale disruption to general IT systems customers will rely on restoring from backups where infections reach critical systems. Recovery costs are also exacerbated by the need to restore personal IT equipment at scale and within a short space of time hardware replacement may be seen as a quicker, cheaper option. This may result in hardware shortages and demand surges. Networks may also be under serious load if backups are based on remote or cloud services.
	Incident response (including forensics)	Yes	There will be an initial demand surge of Incident Response services which will overwhelm capacity of many small players in the market and stretch large vendors to capacity. The demand will be lessened over time by greater understanding of the malware used, and as the global threat response progresses more tools and techniques will be available to react and respond to incidents. Typical activities will initially be containment and recovery, as more defence capabilities are available over time the response will move to more proactive defence techniques.

Scenario constraints and limitations

Malware in this scenario is sophisticated but largely automated; this is reflected in the short time span where global remediation begins to deal with the event.

Ransom payments are factored into losses, however, due to the uncertainty of a payment system being sustainable by the OCGs, the primary cost component is the interruption.

OCGs will be selective in the companies that they feel would give the greatest reward, leaving many to fend for themselves.

Overall, the scenario also requires a number of sophisticated components to make up a complete toolkit that would infect at a high rate and evade detection from common security controls. Additionally, the initial access vector would need to be successful within a very short time frame to avoid detection and adaption of defences.

Initial business impact and losses would be expected to be moderate to start with, compounded by the demand surge in incident response, but tail off in the mid-term as the global security community reacts and fixes vulnerabilities/better detects the malware components.

Demeter's Curse | Targeted Industry Loss Event

Scenario context

Previous cyber-attacks have caused some large events for the insurance market, for example, either:

- Large individual losses, both in overall and insured losses. For example, Merck losses are believed to have exceeded \$1b,⁴¹ with claims believed to be made for \$700m insured losses (albeit not covered under a standalone cyber line policy). Some of these very large individual losses have been part of wider linked common-cause events (e.g. NotPetya ransomware attacks).
- Large notification events (i.e. impacting many insureds) but with relatively little insured loss per impacted insured, leading to limited overall insured loss for carriers.⁴²

There have also been events that have been impactful and high-profile cyber-attacks, causing notable levels of economic loss, but with limited insurance impact. For example, the Colonial pipeline attack in 2021 was sufficiently severe as to cause shortages of oil in parts of the US, as well as causing the US President to issue emergency orders.⁴³ The associated insurance loss has been as assumed as relatively small but the period of outage associated with the attack lasting a week, plus ongoing knock-on impacts and the potential of a critical firm/supplier/supply chain being impacted by a cyber-attack and causing large losses remains a concern.

The Demeter's Curse scenario has built upon a Colonial-type scenario and has sought to identify whether a capital-depleting level of loss for the cyber insurance market could be caused by a small selection of very large tower losses, alongside additional contingent business interruption losses, for individual firms in particular sectors. This is in contrast to the other two scenarios in this paper, which focus on all-sector events.

Scenario development

The Partnership looked into the possibility of a Colonial-type attack (e.g. major aggregation point for a particular industrial sector) on a larger scale (e.g. multiple aggregation points or sectors). In particular, the Partnership was interested in looking at industrial sectors, given that interruptions to operational technology (OT) have the potential to cause extended business interruption losses and it would enable an investigation into the potential impact of the following key policy wording concepts:

Pay-outs for voluntary shutdown

Some firms might shut down parts of their operations, even where those operations are not directly impacted by cyber-attack. For example, because other parts of their business have been impacted by the cyber-attack (e.g. shutting down the operational technology environment for fear of the spread of compromise from their IT environment, or due to an inability to safely maintain operational technology systems). Pay outs due to such shutdowns will be dependent on the policy wording.

Infrastructure exclusions

Depending on the exact firms impacted, and policy wordings, there may be exclusions based on critical infrastructure companies being impacted by cyber-attack. However, this may come down to particular firms, as not all firms affiliated with a sector deemed critical infrastructure, will necessarily be excluded.

Contingent business interruption (CBI) cover

Firms are impacted when other firms are hit by the cyber-attack but pay outs for outages due to this will be dependent on policy wordings and what type of upstream providers are deemed as in-scope of the policy. Additionally, CBI cover may not be subject to full policy limits and may be sub-limited.

As a result of discussions within the Partnership and third-party expert consultation, two broad options were considered for the scenario:

Option One

Event impacting a specialist service provider for a several industrial sectors, with ensuring loss to users of the service provider.

Option Two

Event driven by a vulnerability in a system commonly used by many companies in several industrial sectors, with ensuring loss to a number of those companies and additional contingent business interruption to some other companies in the sectors.

Ultimately, the Partnership did not pursue Option One because the data and current market conditions suggested that the level of insured loss generated from such an event would not cause concern to the market as a whole. This was for two primary reasons:

- Contract structure:
 - » Insured loss across a number of firms would largely be generated as CBI only, as the direct business interruption loss would impact on the service provider only. Whilst the direct business interruption loss could be very high for the service provider, a single firm's limits being exhausted would not cause a capital-depleting event for the market. Additionally, the typical contract structure for CBI loss is such that cyber insurance policies tend to only pay out for CBI loss in relation to Tier 1 suppliers (those directly contracting with the firm suffering the cyber-attack). That is to say that CBI is only applicable for losses caused to a firm if they are a direct, Tier 1 supplier of the firm suffering the disruption. As a result, the CBI losses for Option One would be limited by this structure.
 - » Some permutations of this scenario, if they impacted a certain sector or type of firm, could trigger the critical infrastructure exclusion present in most standalone cyber insurance policy. In this instance, the CBI losses may be excluded and would not pay out, although the direct business interruption loss for the service provider would likely pay out.
- Current standalone cyber market exposure to individual sectors: based on the Partnership's understanding of current limits available in the standalone cyber insurance market, there ostensibly does not seem to be any combination of a single sector that might be impacted by this event and a sector that also has a sufficient number of very large cyber insurance towers (e.g. \$100m limits and upwards) to have the potential to cause a capital-depleting level of insured loss.

The analysis therefore suggests that, at the current time, an Option One-type scenario would probably need to focus on a fairly-universal service provider type that is not specific to particular sectors. For example, cloud service providers or similar equivalents (e.g. unlikely to be industrial service providers). This could change over time as the cyber insurance market profile changes.

As a result of the above analysis, Option Two was the option chosen to complete modelling on.

Scenario description

A group, or groups, of threat actors are motivated by one or more political (including terroristic), social or economic drivers decide to target a selection of large companies operating primarily in the United States, including a desire to cause damage and destruction. For example, the motivation may be ideological, a desire to cause significant impact, or due to incidental reasons, such as the availability of known zero-day vulnerabilities for a system, or systems, used widely in the industrial sectors. If the motive is financial, then the threat actor/s are likely to be high-end organised crime, or state-sponsored actors for certain states that use cyber operations to gain funds.

The threat actors are high capability actors who can identify or acquire one of more zero-day vulnerabilities and integrate those vulnerabilities into a practical 'cyber-attack path' at scale. The actors have a high appetite for risk, and significant start-up capital to build this attack. This is particularly relevant as the actors likely require specialist knowledge of the industrial sectors being targeted, including knowledge of particular IT and OT systems that support those sectors, and the associated degree of dependency on those systems to maintain business operations. The actors' tolerance for risk is particularly high given known responses by countries to certain cyber-attacks in the past. For example, the Colonial Pipeline attack was purported to be using the capabilities from an OCG called Darkside; the group then avowed to avoid similar attacks, due to the attention and impact on critical infrastructure, and therefore vet targets before launching ransomware in future. The attack in our hypothetical scenario would be significantly more disruptive and high-profile so would need a high tolerance for risk.

- 1 Initial access:** This phase involves the widespread infiltration of companies by the threat actors, after potentially months of reconnaissance and development of capabilities (e.g. identification of vulnerabilities, creation of custom malware). The nature of this infiltration will depend on the exact mechanisms of initial access employed by the threat actor/s and how they gain a 'beachhead' in a company's IT network. For the purpose of this narrative, it is assumed that the infiltration would rely on vulnerabilities (potentially zero-day vulnerabilities) in externally-accessible systems (specific to the impacted industries), thereby enabling widespread compromise across a number of companies concurrently. In practice, there may be other ways of enabling initial access to companies on a large scale (e.g. watering hole-type attacks, widespread phishing). Once widespread compromise has occurred across many firms, initially in the US but also spreading globally, the threat actors begin to perform discovery to understand the companies' IT networks from the inside, as well as deploy additional tools and malware within the compromised companies. By the end of this phase, there is no widespread knowledge of compromise by companies or the security community, and threat actors have embedded themselves into compromised companies, either through a combination of human operator-driven actions (e.g. through remote access via a command-and-control channel), or through malware enacting its own actions based on conditions pre-set for it during the coding process.
- 2 Privilege escalation and lateral movement:** During this phase, threat actors aim to escalate their privileges to gain a greater level of access within networks. Depending on the initial access vector that was successful for a particular compromised company, the threat actor may already have a degree of privileged access to certain systems. It is assumed that lateral movement across and within networks is required either on a self-spreading basis, or through human involvement via command-and-control channels. The privilege escalation and lateral movement may rely on exploiting additional vulnerabilities, including zero-days, and is focussed on moving towards being able to access business critical systems, which could include IT or OT systems. Given the footprint in this scenario, little to no human involvement by threat actors is assumed during the attacks on most compromised companies. Where human operators are driving the cyber-attack through command-and-control channels it is assumed that they will focus primarily on the targets deemed most valuable, which are likely to be the largest companies.

3 Impact: This phase involves the crystallisation of the business impact across companies, as well as the recovery from the attack. Of the impacted companies, there are varying degrees of impact, for example dependent on where in the lifecycle of the attack that the event was stopped. Those attacks that completed the Privilege Escalation and Lateral Movement phases are assumed to have a greater chance of causing higher level of impact. Overall, the threat actors have been able to perform a number of actions to further their objectives, such as destruction of data, encryption of systems and therefore disruption to business operations, including potentially through causing damage, or shutdown (voluntary or otherwise) of operational technology and critical IT systems in some cases. Pay outs due to shutdowns such as this will be dependent on the policy wording.

4 Mitigation and response: Mitigation and response occurs throughout the phases of the attack and helps determine if, and to what degree, companies suffer losses from the attack. A number of factors are key in the response and recovery phase and will impact companies' ability to recommence full operations. For example:

- Physical recovery or restoration is assumed to be needed in certain cases (e.g. engineers travelling to remote sites).
- Speed and delivery of patch for vulnerable system and the timely application of the patch to address vulnerabilities used in the attack.
- The overall degree of impact and the resultant government and security community involvement in the response.
- How similar the attack paths are within companies (e.g. same initial access vectors, same indicators of compromise), as this may reduce the overall level of forensic analysis and incident response required.
- Risk categories of impacted companies, as it is assumed that those companies with the best security and resilience are more able to stop the attack before the crystallisation of the business impact (e.g. avoiding lateral movement between IT and OT environments), or recover more quickly. This is why the best risk category companies in the model have the lowest overall footprint percentage, as well as the lowest average business interruption outage.

Cost components

Cost component		Applicable(Y/N)	Narrative commentary
Business interruption costs	(Direct) business interruption	Yes	Primary driver of loss in this scenario given that each individual company directly compromised via an intrusion into their network will incur business interruption costs if the attack path completes.
	Contingent business interruption	Yes	Given the interconnection and reliance within the impacted sectors, some companies will be affected as key suppliers are disrupted.
Non-business interruption costs (also referred to as additional costs)	Notification costs	Yes	Limited notification costs, to include identifying affected companies, preparing and sending out notifications, and managing public relations and communication efforts.
	Extortion	Yes	Initially, some companies affected may be targeted with extortion if ransomware is deployed but the insured costs incurred and not expected to be high, given attention on the threat actors and their likely inability to maintain open cryptocurrency wallets, as well as widespread attention that decryption keys may not be working.
	Data recovery	Yes	Whilst some customers have protected and available back-up procedures in place, both immediate and downstream data recovery costs are high, with many businesses unable to recover large percentages of business-critical data. The most significant losses are those associated with hardware replacement. Due to the number of companies attempting to replace endpoints, an acute hardware shortage prompts a demand surge resulting in increased business interruption (for those unable to purchase replacements immediately).
	Incident response (including forensics)	Yes	Substantial driver of loss given that many individual companies are compromised and therefore there is a need for widespread incident response activities. In practice, the volume of incident response hours needed may partially depend on how similar the attack paths are within companies (e.g. initial access vectors, same indicators of compromise), as very similar attacks across companies may reduce the overall level of forensic analysis and incident response required per company.

Scenario constraints and limitations

This scenario relies on a vulnerability in a system commonly used by many companies in the impacted industrial sectors. This paper does not define the exact vulnerability or system/type of system impacted, beyond supposing that it could be a commonly used Industrial Internet of Things (IIoT) related one, as IIoT-related systems can be cross-sector (but likely not all-sector). This could include a wide variety of possibilities, such as platforms, chips, code libraries and many others.

The inability to define the specific system and vulnerability makes it more complicated to write a compelling narrative and model but is necessary in this case, as the Partnership does not have sufficient evidential weight to be sure enough to define realistic examples of aggregation points in relation to IT and OT that impact a number of different sectors, particularly given the complexity and opaqueness of underlying components and the variations between industries. The challenges of the supply chain in this regard, including unknown aggregation, is one of the reasons that, for example, regulators are looking into the topic, such as via the EU Cyber Resilience Act.⁴⁴ Other publications have theorised about IoT-related scenarios that could cause major insured losses.⁴⁵

A significant amount of due diligence was performed prior to the decision not to select a particular system/type of system, and a number of external third-party experts were consulted, but there was no universal agreement on what could conclusively meet the criteria. For example, the diligence included investigating a number of more commonly-known potential points of aggregation (e.g. airline software, baggage handling firms, card payment providers, industrial control systems) as well as some more obscure ones (e.g. battery providers for EV automotives, gas turbine suppliers).

The Future of the Evolving Cyber-Threat Landscape: Detail

To understand the emerging threats with the greatest propensity to impact our scenarios, we looked back at previous cyber-attacks and technologies to identify common attributes and features. For example, the NotPetya cyber-attack in 2017 saw destructive and widespread malware distributed through a malicious software update for a Ukrainian accounting software impact many companies with operations in Ukraine. This can be considered a pivotal moment in the threat landscape, setting the stage for the ransomware-dominated period from 2018 to 2022. Recognising the crucial factors that made NotPetya possible and understanding its significance in the history of cyber threats have been essential considerations for the Partnership. These insights have guided our horizon scanning efforts and the testing of scenarios presented in this paper, helping us understand their resilience against future technological and social shifts.

The following are the key factors identified by the Partnership which, combined, created the conditions for the NotPetya attack to have a significant and far-reaching impact:

- Emergence of sophisticated exploits for widely unpatched vulnerabilities increases the risk of rapid and widespread propagation of malware. For NotPetya, this role was played by EternalBlue, which was leaked by the Shadow Brokers. Whilst EternalBlue has been used in a large number of varying threat actor campaigns, it was its combination with the below attributes that resulted in NotPetya's impact.
- Presence of new threat actor techniques increases the likelihood existing cyber security controls will have a diminished capability. The WannaCry cyber-attack of 2017 used the EternalBlue exploit as well, but primarily spread as a worm that scanned the Internet for vulnerable systems, infecting them directly. NotPetya, on the other hand, focused on lateral movement within networks, using stolen credentials to infect other vulnerable systems within the same organisation on a largely automated basis, greatly increasing its damage. Organisations with limited network segmentation were particularly vulnerable to this novel attack technique and the ability of security controls to flag malicious activity was reduced by NotPetya's use of legitimate Windows administrative tools. The global impact of a regionally targeted attack also highlighted the raised possibility of threat actor 'miscalculation' and 'unintended consequences' when working with novel attack techniques or new technologies. This was considered by the project team when testing scenarios.
- Targeting trusted software and infrastructure provides threat actors with a staging point to launch more disruptive and damaging attacks enjoying wider reach or access privileges.⁴⁶ NotPetya specifically targeted the software update mechanism of M.E.Doc, a trusted Ukrainian accounting software. The malware was able to distribute itself to numerous organisations that relied on this software, facilitating rapid and extensive propagation. The Partnership identified applications with strong market penetration, particularly cyber security and identity management tools, as some of the propagation vectors posing the highest risk for potential significant losses.

We analysed many viewpoints from relevant industry experts and publications to identify common themes in perceived cyber threats. The overall results of this study are summarised in the table below — themes that were present in multiple publications appear higher up in the table:

Theme	Publication									Count of theme
	ENISA	US National Intelligence Council — Global Trends 2040	Microsoft	NordLayer	10Guards	Field Effect	LinkedIn	Manage Engine Blog	Iplocation	
AI Abuse/Machine Learning	X	X	X	X	X	X	X	X	X	9
IoT-enabled Advanced Attacks	X	X	X	X	X	X	X	N/A	X	8
Supply Chain Attacks	X	N/A	X	X	X	X	X	N/A	N/A	6
Advanced Disinformation Campaigns	X	X	X	X	N/A	X	X	N/A	N/A	6
Data Regulations	N/A	X	X	X	N/A	X	X	N/A	X	6
Cyber-physical Human Errors	X	N/A	N/A	X	N/A	X	X	N/A	X	5
Advanced Hybrid Threat	X	X	X	N/A	N/A	X	X	N/A	N/A	5
Skill Shortage in Cyber	X	X	X	N/A	X	X	N/A	N/A	N/A	5
Cyber Physical Systems	X	X	X	X	N/A	N/A	N/A	N/A	N/A	4
Biometric Data	X	X	X	X	X	N/A	N/A	N/A	N/A	4
Zero Trust Security Framework	N/A	N/A	N/A	X	N/A	X	X	N/A	X	4
Digital Surveillance Authoritarianism	X	X	X	N/A	N/A	N/A	N/A	N/A	N/A	3
Industry Disruption and Jobs	X	X	X	N/A	N/A	N/A	N/A	N/A	N/A	3
Telco as a Single Point of Failure	X	N/A	N/A	X	N/A	N/A	N/A	N/A	N/A	2
Quantum Computing	X	N/A	N/A	X	N/A	N/A	N/A	X	N/A	2
Blockchain/Decentralised Systems	N/A	N/A	X	X	N/A	N/A	N/A	N/A	N/A	2
Crime as a Service	N/A	N/A	X	X	N/A	N/A	N/A	X	N/A	2
5G Networks	N/A	N/A	X	N/A	N/A	N/A	N/A	X	N/A	2
Quantum Cryptography	N/A	N/A	X	X	N/A	N/A	N/A	N/A	N/A	1
Open-Source Tools	N/A	N/A	N/A	N/A	N/A	X	N/A	N/A	N/A	1
Secure Access Service Edge	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	X	1
Cyberterrorism	N/A	N/A	N/A	N/A	N/A	N/A	N/A	X	N/A	1

The commentary on the threat trends identified in the studies above diverged in the depth of analysis provided. There was limited consistency in the threats identified as most likely to prompt threat landscape changes. Our focus when evaluating different emerging trends was whether they had the capacity to change the nature of our malware narratives or change their associated insured loss profiles. We continually found that many of the trending event types that we observed had the potential for significant economic loss or social and political unrest (e.g., through ‘widespread disinformation’ campaigns), but had limited direct impact to our scenarios or insured losses. We also observed that cyber events carried out against critical infrastructure was a common phenomenon, something we ignore (given general insurance market exclusions) with our focus on insured loss. Some threat trends could be grouped by the types of threats they presented or the systems and organisations they were theorised to impact; many centred around the inherent risk of aggregation points, whilst others highlighted operational technology threats.

Evaluating these threats against our chosen scenarios and the criteria that made NotPetya so impactful resulted in the Partnership focusing on threats presented by generative AI abuse/machine learning, IoT/IloT and quantum cryptography in scenario construction and testing. These three themes are explored further below.

Artificial Intelligence abuse/Machine Learning (AI/ML)

Generative AI technology, such as LLMs (large language models) like ChatGPT, is highly likely to impact the threat landscape. AI tools enable greater automation and streamlining of cyber-attacks at multiple points within a cyber-attack attack path. Automation is a key driver of widespread malware event scalability and can impact both speed of propagation and footprint. Some of the key areas where AI automation could change the nature of a cyber event could be:

- **Social engineering/phishing:** AI enables material improvements in automated and scalable social engineering attacks. Through natural language processing and generative models, AI can imitate legitimate communication, increasing the difficulty for users and security tools to identify social engineering attempts. AI can also be used to create deepfake videos and audio recordings, potentially enabling impersonation for malicious purposes.
- **Adaptive malware:** AI can generate malware with self-learning capabilities, enabling it to operate more like biological viruses where it can adapt to its environment and decrease the prospect of detection by both signature and behavioural security measures.

Internet of Things/Industrial Internet of Things (IoT/IloT)

The exponential growth of IoT or interconnected devices, from smart homes to industrial control systems, changes the attack surface. Many IoT devices are less secure by design than traditional endpoints, servers or even mobile devices and firmware updates are often intermittent or non-existent. Mitigating risk from these technologies' hinges upon asset management, network segmentation and strong authentication mechanisms.

We can anticipate IoT devices to increasingly play a role in cyber-attacks, either by giving a threat actor more points of data collection, an initial access vector, or a device to be added to a botnet. For the purposes of our malware scenarios, the Partnership believe that IoT devices might be impacted by/part of the propagation of widespread destructive malware.

Quantum computing/Cryptography

Another emerging threat of note is quantum cryptography, which will force an evolution of how users are authenticated and data is encrypted. This could have the potential to impact many of our parameters, but it is too early to identify whether this would alter our scenarios, since the full impact of quantum cryptography on currently used encryption methods is likely to be many years away. Defences are already being prepared and quantum-resistant cryptographic methods being developed. In addition, the Partnership have determined that multiple commonly used cryptographic methods being subjected to active exploitation by multiple threat actors overnight is highly unlikely, given the difficulty of obtaining quantum processing power.



Additional Detail to the Approach

The structure and governance of the work was broadly consistent across the 18-month period of the project and included the following:

- Content Working Group meeting weekly to distribute and track actions, with relevant sub-groups and ad-hoc meetings also deployed.
- Steering Committee meeting approximately monthly to oversee progress of the Content Working Group.
- Frequent collaboration workshops, lasting approximately 1-3 days per workshop. The aim of these workshops was to perform intense periods of collaborative work on the project and identify actions that would be distributed for completion by Working Group members in the intervening period before the next collaboration workshop.
- Internal outreach (to experts within the Partnership) and external outreach (to third-party external experts) to provide inputs into various discrete parts of the project (e.g. opining on the appropriateness of the scenarios, whether any other scenarios had been missed, as well as technical cyber aspects of the scenarios, including response and recovery from a theoretical systemic event). This outreach included experts in a number of different organisations and fields, including cyber underwriting, cyber security and resilience, cyber exposure management, public policy and operational specialists in certain industry sectors.

Investigating the current state of modelling

This sub-section provides further detail on the analysis of the existing academic and insurance publications, as well as third-party accumulation models, performed by the Partnership. Additionally, commentary is provided on the differences between RDS and stochastic models for cyber modelling.

Publications on systemic cyber-risk

The Partnership identified bodies of literature on systemic cyber-risk based on existing knowledge, consultation with third-parties, attendance at events and structured Internet analysis.

The publications reviewed primarily provided background materials to help inform the Partnership's thinking in relation to this project, as opposed to directly informing the scenario development. However, the publications helped with contextualising the position of this project within the wider discourse on systemic cyber-risk.

Alongside the background materials on systemic cyber-risk, the Partnership also completed extensive review and analysis of the current cyber threat landscape through document analysis, consultation and additional mediums, all with the aim of providing additional background and context ahead of scenario development.

Third-party cyber accumulation models

The Partnership identified some of the most widely used vendor models in the cyber insurance market and determined which scenarios within each model would meet the definition of a 'malware' scenario for the purposes of this project.

Each malware scenario was then assessed against Gallagher Re's proprietary Vendor Model Scenario Assessment Framework, by running the scenarios in the model, using a simulated, synthetic portfolio which was deemed representative of the wider cyber market.

This framework provided a repeatable and structured way of assessing each scenario through a technical and non-technical lens. For example, the framework assessed the narrative, frequency, cost components and other assumptions and attributes for each scenario.

As an output, this exercise provided a list and analysis of the malware scenarios that vendor models (and by extension a number of external experts) deemed could cause systemic cyber-risk. Additionally, it provided the Partnership with an informed viewpoint around potential areas for improvement, which was carried through into the approach described in the following phases.

Published cyber disaster scenarios

The Partnership identified a list of malware scenarios, outside of third-party cyber accumulation models. These scenarios could broadly be split into those published by regulators and those published by other bodies (e.g. industry and academia).

Some scenarios were not considered for the analysis — for example, scenarios that were already included in the third-party accumulation models referenced above, scenarios that were not explicitly deemed as malware-driven scenarios, or some regulatory-published scenarios (e.g. either because they are perceived as insufficiently detailed for this exercise, or they were not finalised at the point of analysis, such as those included in the EIOPA paper on cyber stress testing).⁴⁷

For the scenarios considered, a bespoke scenario review matrix was completed identifying common data points that could be compared across scenarios (e.g. type of malware involved, footprint of the scenario, geographical scope, industry/sector scope).

The analysis provided data on the coverage and comparability of existing scenarios and therefore information on how other expert parties perceived that systemic cyber-risk could arise for insurers from malware scenarios. This informed the Partnership's thinking for the project.

RDS vs stochastic

For exposure management of insurance portfolios, it is common to frame risks in terms of return periods – in simple terms, what would be the expected loss from a 1-in-100 year hurricane (e.g. the size of the loss expected to be seen on average every 100 years). This type of modelling is often referred to as probabilistic modelling. For natural perils, which are often either random, driven by probabilistic processes or at least periodic, this makes intuitive sense but is harder for cyber-risk which mainly arises from human-led attacks on systems and has a relatively short history. The solution that has been adopted by the insurance industry is to use RDS around themes such as cloud outage, self-propagating malware, electrical blackouts or similar themes. This is where scenarios are described with a methodology to determine the amount of loss a portfolio will sustain if such an event occurred. Of course, such scenarios are only as useful as their assumptions, but RDS or equivalents are currently considered an important regulatory tool.

One obvious criticism of the RDS approach is that while it provides a sound methodology for estimating the potential severity of catastrophic losses, it does not provide an estimate of the frequency of losses. A variety of possible solutions to bridge this gap have been theorised.

The goal of modelling teams is to move toward developing fully probabilistic and stochastic models which can simulate events of all different severities. However, determining the parameters for such a model is difficult and their construction is complicated, which makes them opaque and counter to the goals of this project, which seeks to offer a simple and transparent model that can be easily assessed. The model presented in this work provides a bridge between a purely deterministic approach and a full, data-driven probabilistic model by deploying assumptions about the distribution shape for losses based on the characteristics of the affected insured. This allows the varying severity of loss to a company to be simulated, even if the overall probability of a scenario occurring is deterministic. These characteristics are necessarily subjective but have been evaluated using expert judgement for technical consistency.

There are ways to try to gain more data, such as using honeypot data to examine the frequency of attacks and then to map these to claims made against the portfolio. This might then be used to calibrate a stochastic model for losses (see, for example, Bessy-Rolland et al⁴⁸). At present, a key challenge with such a stochastic modelling approach is modelling the potential claims stemming from operational disruption related to IT incidents, and so the default is for the industry to rely to an extent on external third-party accumulation model vendors who have a number of prescribed scenarios that then generate a probabilistic-type output based on the composition of the portfolio.



Identifying areas for improvement

This sub-section provides further detail on the areas which were prioritised during this project, given the Partnership felt that it would have additional information that could enhance them (e.g. in terms of methodology, expertise or data sources).

Geographical nuances

For a cyber event to be systemic and impact multiple policyholders within a short timeframe, it is believed the attack needs to have certain characteristics, such as follow a common vulnerability, system, software, or dependency (see the Introduction section on how systemic cyber-risk can manifest). Essentially the ‘victims’ of the attack may all need some technological dependency in common for the attack to materialise and/or be propagated. Furthermore, all the victims need to be vulnerable within the attack timeframe and before a patch (or otherwise) is available and, more importantly, deployed.

We accept there is global connectivity with the possibility of global propagation (e.g. NotPetya), but there can also be disaggregating factors on a regional level which include:

- Differing attacker motivation and risk management (e.g. attackers might not want an attack to spread beyond certain countries).
- Differing software, including varying market share by region.
- The impact of time zones and the associated availability of machines and users.
- Varying network architecture (e.g. between and within companies).
- Level of Internet connectivity in the region.
- Language differences.
- For service providers — regional infrastructure separation with differing protocols and security (e.g. regional cloud infrastructure architecture, or requirements to compartmentalise systems/data based on legal and regulatory requirements).

For instance, the WannaCry event in 2017 saw regional differences in the firms impacted, and this was in part due to the time zones associated with those regions. Computers need to be turned on and running to be vulnerable to attacks from a propagating virus and thus, depending on the time an attack is initially launched, swathes of the global population may be offline and not immediately vulnerable. This provides internal security professionals with the chance to implement initial defences or simply shut down networks before users start their day, hence limiting the immediate impact. Of course, a patch against such an attack would need to be deployed for a long-term solution.

Many of the existing RDS or third-party accumulation model vendor scenarios, particularly in the tail, have been focused on global events. Given the presence of disaggregating factors at regional level we consider that this could be an area for improvement and have sought to reflect some benefits of diversification in a portfolio when constructing the scenarios.

Underwriting risk quality and claims data

Modelling cyber losses poses a significant challenge due to the limited availability of claims data. Unlike traditional insurance lines with well-established historical data and long-standing physical processes, the dynamic and rapidly evolving nature of cyber-risks makes it difficult to amass a comprehensive dataset. Cyber incidents vary widely in terms of scale, sophistication, and impact, and many organisations hesitate to disclose the full extent of their breaches due to concerns about reputation damage and regulatory implications.⁴⁹ As a result, insurers and risk modelers face a scarcity of reliable information to calibrate their models. This limitation hinders the development of robust predictive models, and thus, the industry’s ability to assess and quantify cyber-risks effectively.

In comprehending this limitation, the Parties collaborated in order to establish a reliable dataset comprised of real-world claims data.²⁸ This dataset was used to inform the way that losses were parameterised and distributed both per event and per insured in our model.

Additionally, data from underwriting experience was used as an input into the Risk Categories in the model, with some minor expert adjustment to help reflect accumulation relevant covers.

Event response

In a real event, companies, both impacted and those not, take active measures to prevent or mitigate impact. Unlike in Property CAT insurance where it is impossible to move the building out of the hurricane's path, as a cyber event unfolds, both before and after an attack on an insured, there are many steps that can be taken to either avoid the event before impact (e.g. through taking the company out of the attack 'path', such as via a patch), during the event (e.g. voluntary shutdown), or post event (e.g. isolating affected endpoints). The application of controls at any point (or cumulatively) throughout the life of the event may fail. However, it is less likely that all the controls will fail and more likely that at least one (or more) of the controls will have a partially positive impact to the insured. For example, business interruption cover is typically provided on an indemnity basis and therefore even some reduction in impact, or quicker recovery, would reduce the overall impact of a systemic cyber event. Pay outs will always be dependent on the policy wording.

Modelling the activities of humans is a far less precise science than that of modelling the physical effects of hurricanes, as companies can make a plethora of decisions in response to a wide variety of attackers, attack vectors, and propagation methods. Furthermore, given the limited event history available to cyber event modellers, there is insufficient data to support our understanding of by exactly how much security measures reduce losses. In taking all of this into account, we acknowledge the difficulties in modelling such considerations and thus understand why parameterisation of event response mechanisms has not been fully integrated into current modelling approaches. However, we believe that these challenges in modelling should not be ignored at a portfolio level, and even after allowing for potential control failure, we have considered some partial benefit of action taken.

Attack propagation

Within existing scenarios, there can be limited transparency into the way in which a cyber-attack can disseminate within a computer network or system and how that impacts on the resulting modelling. Cyber-attacks can propagate through various methods and vectors, and their ability to spread can depend on the nature of the attack and the vulnerabilities present in the targeted systems. For example, models may not ostensible show the key distinctions between how worms can spread both between networks and within them. Failure to incorporate these nuanced propagation dynamics into risk models undermines the models' predictive accuracy and limits their capacity to provide a comprehensive understanding of potential losses. As cyber threats continue to evolve in sophistication and scope, an imperative exists for insurance models to evolve concurrently, adopting a more nuanced and inclusive approach that embraces the multifaceted nature of cyber propagation methods. Our research and model have sought to improve this limitation by providing a simple and transparent 'attack path' based model, combined with propagation narratives within each of our scenarios.

Appendix 1:

Definitions

Below are some of the key definitions used in this paper.

CBI (contingent business interruption) losses — Losses in the model from risks not directly affected by the malware but dependent on those which are.

Cost component — Costs that are grouped according to their specific impact as the result of an insurable event.

Multiple cost components can contribute to an overall loss (financial and/or insured). For example, incident response (including forensics) is a cost component in the additional costs type in the model.

Counterfactual analysis — A framework used in the scenario development to assess the parameters and outcomes of historical events with a view of understanding how these elements could have had more severe consequences if altered.

Direct losses (also called full losses) — Predominately business interruption losses following a successful attack against an insured in the model.

Giant companies — Risks with annual revenues of over \$10b.

Industry Exposure Database (IED) — Proprietary Industry Exposure Database from Gallagher Re, consisting of ~1.2 million policies and \$13 billion of premium. This is based on over 70% of the insurance market.

Large companies — Risks with annual revenues of up between \$1b–\$10b.

Medium companies — Risks with annual revenues of up between \$100m–\$1b.

MITRE framework — The Adversarial Tactics, Techniques, and Common Knowledge or MITRE ATT&CK is a guideline for classifying and describing cyberattacks and intrusions. It was created by the Mitre Corporation and released in 2013. For the purposes of the project, the MITRE framework refers to a condensed version of ATT&CK used to structure the modelling and scenarios.

Micro companies — Risks with annual revenues of up to \$20m.

Modelling — The use of analytics and cyber-risk data (including scenarios) to quantify cyber-risks in order to inform appropriate underwriting, portfolio management, and risk transfer decisions.

Parameter — A numerical or other measurable factor forming one of a set that defines a scenario or sets the conditions of its impact.

Partial losses — Losses where the attack was not successful but there are still some costs (e.g. response costs) for the purposes of modelling.

Realistic Disaster Scenario (RDS) — A narrative that presents a plausible and catastrophic event with a low likelihood/frequency attached. An RDS is a means of stress testing a (re)insurer's portfolio.

Risk category — A risk specific modifier that is applied on every risk depending on the size and industry.

The risk classification in the model goes from RC1 (strongest controls) to RC7 (weakest controls). The allocated risk category for an industry and size bucket is derived from several factors, including underwriting experience of Beazley and Munich Re, with some minor expert adjustment to help reflect accumulation relevant covers.

Small companies — Risks with annual revenues of up between \$20m–\$100m.

SPM — Abbreviation for Self-Propagating Malware. This is the attack type that is identified for our Scenario 2 narrative.

Systemic cyber-risk (used synonymously with accumulation/systemic or catastrophe risk) — The risk that a portfolio of cyber insurance policies has the potential, in certain situations, to be liable to pay losses far in excess of the premium collected, due to a large-scale issue which triggers many cyber policies in a short period of time.

TILE — Abbreviation for Targeted Industry Loss Event. This is the attack type that is identified for our Scenario 3 narrative.

Vendor models — Third-party accumulation models responsible for producing analytics platforms and associated cyber RDS for use by (re)insurers.

Vendor Modelling Assessment Framework — A proprietary Gallagher Re framework designed to consider both technical and non-technical elements of vendor model scenarios in order to understand whether the narratives are feasible, appropriately informed and a true representation of both current and predicted cyber trends.

WSSC — Abbreviation for Widespread Software Supply Chain. This is the attack type that is identified for our Scenario 1 narrative.

Appendix 2:

Summary of Cyber Insurance Coverage

The following are commonplace coverage elements in cyber, but policies can offer more and less depending on the situation:

- **Information Security Liability** — coverage of third-party claims and associated claims costs arising from unauthorised access, theft of or destruction of data, denial of service attacks and virus transmission resulting from computer security breaches
- **Privacy Liability** — coverage of third-party claims and associated claims costs arising from theft, loss, or unauthorised use of personally identifiable non-public information in computer and hard copy form; failure to properly follow breach notification laws; and coverage for failure to comply with the insured's privacy policies.
- **Breach Response Costs** — such as forensic and legal expenses, notification and credit monitoring for affected individuals as well as crisis management expenses.
- **Regulatory Defence and Penalties** — coverage to defend the insured against regulatory inquiries and proceedings arising from a covered data breach or security breach and, where legally permitted, the indemnification of fines or penalties against the insured.
- **First party Coverage, includes coverage for:**
 - » Restoration of damaged data.
 - » Business interruption, the loss of profit and extra expenses caused by a covered cyber event.
 - » Cyber extortion costs arising from threats to alter, delete, or corrupt data, prevent access to computer systems or data, perpetrate unauthorised use of computers, or steal, misuse, or publicly disclose personally identifiable non-public information.
 - » System failure coverage for unintentional outage of a computer system may be offered.
- **Contingent Business Interruption** — coverage for an insured's business interruption, loss of profit and extra expenses due to the downtime of an outsourced service provider upon which the Insured relies to conduct their business pursuant to a written contract.
- **eCrime** — coverage for Fraudulent Instruction, Funds Transfer Fraud and Telephone Fraud.

Appendix 3:

Parameter Details

Regional spread parameter

We believe that geography has an impact on propagation of malware. For example, in case of a severe malware event, we assume that the effect will be different in the various regions of the world, caused by the time zones or other geographic factors (as detailed previously in the supplementary material). Evidence on the time zone lag was seen during the 2017 WannaCry attack. Therefore, we define a starting region (Asia, Europe, North America and Other) and give this starting region a weighting of 1 and all other regions a factor of 0.75, which is simply multiplied on the probability of being affected (footprint). We recognise that this singular factor is a simplification, and the value itself is blunt, but we believe the intent of the principle to be necessary and worthy of further research. The Parties have begun further exploring the similarities and differences in software used regionally with additional data sets, but this is not factored into the model at present.

Risk category

Like in all lines of insurance, we try to differentiate risks by some of their characteristics. We are defining a risk category which shall represent a number of factors/attributes such as IT security, cyber resilience and business continuity management (BCM), attractiveness as a target and historic trends. We group all risks by industry and size band and allocate a risk category to all of them. Our risk classification goes from RC1 (strongest controls) to RC7 (weakest controls). The allocated risk category for an industry and size bucket is derived from several factors, including underwriting experience of Beazley and Munich Re, with some minor expert adjustment to help reflect accumulation relevant covers.



Risk category

Risk Category	Size Cat 1	Size Cat 2	Size Cat 3	Size Cat 4	Size Cat 5
1.1 IT-Software	RC5	RC4	RC3	RC2	RC1
1.2 IT-Hardware	RC5	RC4	RC3	RC2	RC1
1.3 IT-Services	RC5	RC4	RC3	RC2	RC1
2 Retail	RC6	RC5	RC4	RC3	RC2
3.1 Finance-Banking	RC5	RC4	RC3	RC2	RC1
3.2 Finance-Insurance	RC6	RC5	RC4	RC3	RC2
3.3 Finance-Investment management	RC5	RC4	RC3	RC2	RC1
4 Healthcare	RC7	RC6	RC5	RC4	RC3
5 Business & professional services	RC6	RC5	RC4	RC3	RC2
6 Energy	RC6	RC5	RC4	RC3	RC2
7 Telecommunications	RC5	RC4	RC3	RC2	RC1
8 Utilities	RC7	RC6	RC5	RC4	RC3
9 Tourism & hospitality	RC6	RC5	RC4	RC3	RC2
10 Manufacturing	RC7	RC6	RC5	RC4	RC3
11 Pharmaceuticals	RC7	RC6	RC5	RC4	RC3
12 Defence/Military contractor	RC5	RC4	RC3	RC2	RC1
13 Entertainment & Media	RC5	RC4	RC3	RC2	RC1
14 Transportation/Aviation/Aerospace	RC6	RC5	RC4	RC3	RC2
15 Public authority; NGOs; Non-profit	RC7	RC6	RC5	RC4	RC3
16 Real estate, property & construction	RC7	RC6	RC5	RC4	RC3
17 Education	RC6	RC5	RC4	RC3	RC2
18 Mining & primary industries	RC6	RC5	RC4	RC3	RC2
19 Food & agriculture	RC7	RC6	RC5	RC4	RC3
20 Other	RC6	RC5	RC4	RC3	RC2

■ RC1 Strongest controls
 ■ RC2
 ■ RC3
 ■ RC4
■ RC5
 ■ RC6
 ■ RC7 Weakest controls

The following revenue bands have been used for the size categories, although users are able to customise the risk categorisation to their own internal views as required.

Category	Revenue
Size Cat 1	0–20m
Size Cat 2	20m–100m
Size Cat 3	100m–1000m
Size Cat 4	1000m–10b
Size Cat 5	10b and above

Fixed additional costs

If a company is infected by a malware, there will be BI losses, as well as losses coming from other sources like incident response and data recovery costs. We group all these non-BI losses under ‘additional costs’. These costs are constant for specific industry and size buckets and are derived from single risk loss models of Beazley and Munich Re. For example, the table below provides a breakdown of typical average fixed costs of claims by revenue to provide a sense check against fixed cost assumptions in the model.⁵⁰

Average (claims >\$10k per category)	Cyber Data Recovery	Cyber Crisis Management	Cyber Forensic	Cyber Notification	Cyber Legal
A. Under \$5M	\$23,000	\$9,000	\$21,000	\$8,000	\$17,000
B. \$5M–\$10M	\$44,000	\$15,000	\$25,000	\$7,000	\$20,000
C. \$10M–\$20M	\$46,000	\$21,000	\$27,000	\$10,000	\$21,000
D. \$20M–\$35M	\$53,000	\$16,000	\$32,000	\$13,000	\$23,000
E. \$35M–\$100M	\$84,000	\$30,000	\$47,000	\$18,000	\$27,000
F. \$100M–\$250M	\$146,000	\$18,000	\$50,000	\$19,000	\$30,000
G. \$250M–\$500M	\$172,000	\$49,000	\$64,000	\$28,000	\$37,000
H. \$500M–\$1B	\$211,000	\$32,000	\$70,000	\$62,000	\$70,000
I. \$1B–\$3B	\$373,000	\$36,000	\$89,000	\$71,000	\$57,000
J. Greater than \$3B	\$273,000	\$89,000	\$82,000	\$90,000	\$52,000

Source: Beazley

Within the model, for simplicity, we do not differentiate by industry or scenario, but this is clearly a further refinement that can be made.

Additional Costs (IR, data restoration etc.) excluding Business Interruption	0-20m	20m-100m	100m-1000m	1000m-10b	10b and above
1.1 IT - Software	10,000	300,000	1,000,000	10,000,000	25,000,000
1.2 IT - Hardware	10,000	300,000	1,000,000	10,000,000	25,000,000
1.3 IT - Services	10,000	300,000	1,000,000	10,000,000	25,000,000
2 Retail	10,000	300,000	1,000,000	10,000,000	25,000,000
3.1 Finance - Banking	10,000	300,000	1,000,000	10,000,000	25,000,000
3.2 Finance - Insurance	10,000	300,000	1,000,000	10,000,000	25,000,000
3.3 Finance - Investment management	10,000	300,000	1,000,000	10,000,000	25,000,000
4 Healthcare	10,000	300,000	1,000,000	10,000,000	25,000,000
5 Business & Professional Services	10,000	300,000	1,000,000	10,000,000	25,000,000
6 Energy	10,000	300,000	1,000,000	10,000,000	25,000,000
7 Telecommunications	10,000	300,000	1,000,000	10,000,000	25,000,000
8 Utilities	10,000	300,000	1,000,000	10,000,000	25,000,000
9 Tourism & Hospitality	10,000	300,000	1,000,000	10,000,000	25,000,000
10 Manufacturing	10,000	300,000	1,000,000	10,000,000	25,000,000
11 Pharmaceuticals	10,000	300,000	1,000,000	10,000,000	25,000,000
12 Defense / Military Contractor	10,000	300,000	1,000,000	10,000,000	25,000,000
13 Entertainment & Media	10,000	300,000	1,000,000	10,000,000	25,000,000
14 Transportation/Aviation/Aerospace	10,000	300,000	1,000,000	10,000,000	25,000,000
15 Public Authority; NGOs; Non-Profit	10,000	300,000	1,000,000	10,000,000	25,000,000
16 Real Estate, Property & Construction	10,000	300,000	1,000,000	10,000,000	25,000,000
17 Education	10,000	300,000	1,000,000	10,000,000	25,000,000
18 Mining & Primary Industries	10,000	300,000	1,000,000	10,000,000	25,000,000
19 Food & Agriculture	10,000	300,000	1,000,000	10,000,000	25,000,000
20 Other	10,000	300,000	1,000,000	10,000,000	25,000,000

Attack path split for additional costs

As mentioned previously in the paper, we also want to reflect the losses that an insured experiences even if the cyber-attack is not successful or more correctly is only partially successful (i.e. attack does not complete the entire attack path). We therefore allocate portions of the additional costs to the four different steps of the attack (Initial Access, Privilege Escalation, Lateral Movement, Impact). This distribution of the losses is subjective since only limited (or even no) data is available. From our perspective, additional costs appear at the point of a successful initial access since once an attacker is successfully in the system, there are likely some incident response costs (e.g. proactive compromise assessment to determine how widespread access was) even if the attack is believed to have been stopped after initial access (e.g. isolation and containment of affected endpoints). Therefore, we are allocating 50% of the additional costs to the Initial Access step of the attack path; this is also a prudent allocation given the uncertainty of how these costs might be split in practice. The Privilege Escalation and the Lateral Movement steps get another 10% each of the allocation and the successful attack (Impact attack path step) then generates the final 30% of the additional costs.

CBI parameters

Companies which are not infected by the malware might still have a loss in case they are dependent on a company that is impacted directly by the malware. The insurability of that loss would be dependent on policy wording; for example, the directly impacted company may need to be a direct, Tier 1 supplier of the company suffering the CBI loss.

We assume that the severity for these indirectly impacted companies is smaller than for the ones which are hit directly. A factor of 50% on the average BI loss of a directly impacted company seems prudent, for example to allow for the fact that most companies are not completely reliant on a third-party for their operations. As for the footprint, we assume a factor of 20% of companies might be impacted indirectly compared to the directly impacted ones. Both the assumptions for severity and footprint are subjective and based on expert judgement.

Gross margin

Based on typical insurance policy wording, the insured BI loss does not cover the revenue but only the profit (policies are an indemnity-based product). Therefore, we apply gross margin rates to the revenue. We use the second highest out of the last six years' gross margin rates from the NYU Stern database⁵¹ of publicly available company performance. At the time of the analysis this covered 2017 to 2022. The second highest year was used for each industry independently, as this was considered to be prudent and also accounted for the fact that some of the years used in the data may have experienced impact from COVID-19. If we used an average across the five-year period, the impact of COVID-19 would impact the values used. However, we note that there was limited difference to the average gross margin of the last five years, with our selected measure typically less than 1% higher per industry. The gross margin rates used in the model are shown below.⁵²

Gross Margin	
1.1 IT — Software	0.66
1.2 IT — Hardware	0.29
1.3 IT — Services	0.21
2 Retail	0.28
3.1 Finance — Banking	1.00
3.2 Finance — Insurance	0.24
3.3 Finance — Investment management	0.56
4 Healthcare	0.40
5 Business & Professional Services	0.32
6 Energy	0.40
7 Telecommunications	0.47
8 Utilities	0.42
9 Tourism & Hospitality	0.40
10 Manufacturing	0.29
11 Pharmaceuticals	0.65
12 Defense / Military Contractor	0.23
13 Entertainment & Media	0.40
14 Transportation/Aviation/Aerospace	0.25
15 Public Authority; NGOs; Non-Profit	0.30
16 Real Estate, Property & Construction	0.33
17 Education	0.46
18 Mining & Primary Industries	0.26
19 Food & Agriculture	0.20
20 Other	0.30

Percentage of daily revenue lost

In case of a malware attack, some companies will be able to maintain at least partial profit generation (e.g. not all operations may be halted during the attack, new or additional channels for generating revenue may be leveraged). We model this by applying a factor on the BI loss. This factor is dependent on the size of the company and decreases with increasing revenue. The idea behind this behaviour is that larger companies, due to their size, will be able to at least maintain part of their operations. There may also be other factors — for example, the largest global companies are often formed of multiple operating units, based on product and/or region, and may be on independent systems which limit the internal spread of malware (e.g. private equity firms that buy other firms may not integrate technology systems with them).

We have taken a prudent approach with this parameter, as we consider the principle to be important but recognise further assessment is needed, as it is based on expert judgement. The factors used in the model are shown below.

Percentage of Daily Revenue due to BI	Percentage of BI
0-20m	90%
20m-100m	83%
100m-1000m	75%
1000m-10b	68%
10b and above	60%

Industry-specific event

For Scenario 3, we are modelling an event that is sector-specific (not all-sector). As detailed in the supplementary material, this paper does not define the exact vulnerability or system/type of system impacted for the scenario. However, we suppose that it could be a commonly used Industrial Internet of Things (IIoT) related one, and therefore the industry sectors chosen are those perceived to have a greater take-up of operational technology (OT) and IIoT devices. The sectors modelled for this scenario are shown in the table below.

Industry	Included in scenario
1.1 IT - Software	N
1.2 IT - Hardware	N
1.3 IT - Services	N
2 Retail	N
3.1 Finance - Banking	N
3.2 Finance - Insurance	N
3.3 Finance - Investment management	N
4 Healthcare	Y
5 Business & Professional Services	N
6 Energy	Y
7 Telecommunications	N
8 Utilities	Y
9 Tourism & Hospitality	N
10 Manufacturing	Y
11 Pharmaceuticals	Y
12 Defence/Military Contractor	N
13 Entertainment & Media	N
14 Transportation/Aviation/Aerospace	Y
15 Public Authority; NGOs; Non-Profit	N
16 Real Estate, Property & Construction	Y
17 Education	N
18 Mining & Primary Industries	Y
19 Food & Agriculture	Y
20 Other	N

Footprint

To define the footprint of a malware event, we split the attack in four steps: Initial Access, Lateral Movement, Privilege Escalation and Impact. Each of these attack path steps have a population which is attacked (as the percentage of the overall population/the population where the step before was successful plus any adjustment) and a probability that this attack is successful. These attack path steps are parameterised separately for each scenario. The parameters were chosen in a way so that the scenario is extreme in the percentage of the population which is impacted. The final footprint is the probability of being attacked multiplied with the probability of success for all four steps (multiplied).

We have chosen a probability of being attacked and a success probability of each of the four attack path steps for all risk categories (RC1-RC7). This is because we want to reflect that different risks have a different probability of being affected by a malware attack.

Explanations for the rationale of the footprint parameters is detailed below, split by scenario.

Scenario 1 — Autolycus

Attack path step	Population exposed	Success factor
Initial Access	The purpose of this scenario was to consider the impact of a supply chain attack in a leading software provider. The assumption underpinning the narrative is that this provider would have some of the highest level of access privileges. Therefore, we opted for the largest market share for software providers that we considered as viable targets for an attack such as this (e.g., leading anti-virus software providers). Therefore, population exposed is set to 30%.	In order to align with the faster patching cadence of prominent software providers and account for potential failures arising from comprehensive deduction systems and sandboxing, we set a success rate of 60% for the initial access phase. However, for lower quality risks, the success factor can decrease to a minimum of 20% in scenarios where patching cadence is slower.
Privilege Escalation	Due to the nature of the event (see success factor explanation), we assumed that population exposed would be 100%.	Due to the inherent nature of a software supply chain event, there is an assumption of near-automatic success; however, for higher quality risks, a failure rate of up to 10% is accounted for to accommodate potential limitations imposed by sandboxing or network segmentation. Thus, the success factor rate ranges from between 90–100%.
Lateral Movement	Due to the nature of the event (see success factor explanation), we assumed that population exposed would be 100%.	As with the privilege escalation parameterisation, there is a similar assumption at play for the lateral movement component of the attack path. There is a significant likelihood of success; however, it is important to acknowledge that the initial release took place at 10:00am Eastern Time Zone, the majority of impacted customers are in North America and as such, the sequential impact of the attack across different regions needs to be taken into account and may result in attack execution being far less likely (i.e., countries in the Southern Hemisphere). Therefore, the success factor at this stage ranges from 50% to 85% (with RC1 category companies presenting at the lowest end of this range and RC7 at the highest).
Impact	Upon achieving complete control of the system, the population exposure is assumed to be at 100%.	The impact experienced by companies is contingent upon their incident response and recovery capabilities, with variations observed across different risk categories. For instance, organisations classified under RC1–RC3 are more likely to leverage services like Citrix, resulting in significantly improved recovery time. Conversely, companies with limited IT and cyber security resources (RC4–RC7 risk categories) are more prone to longer downtimes. As a result of these considerations, success rate factor ranges from 70% (RC1) to 90% (RC7).

Scenario 2 — Lernaean Hydra

Attack path step	Population exposed	Success factor
Initial Access	<p>This scenario considers phishing to be the main initial access vector for the newly developed malware.</p> <p>It was considered that a high percentage of the population was targeted by sophisticated malware (75% for all risk categories).</p>	<p>However, even with the addition of AI generated phishing emails we perceived the probability of success for risks with better controls and awareness to be somewhat limited (RC1 set to 10%).</p> <p>Lesser mature companies however were expected to be much more susceptible to the attack with a much higher success rate (RC 7 set to 60%).</p>
Privilege Escalation	<p>Of the initial vector it was considered a high number of the population to be further exposed to privilege escalation, however some allowance was made for the malware either not acting as intended or being interrupted, therefore the exposure figure was set to 75% for all risk categories.</p>	<p>Given the component parts of the malware solution in this scenario it was considered that even the most defended and able companies would have their defences compromised and privilege escalation achieved by the malware. This gave a variance of between 60% for RC1 and 80% for RC7 based on risk categories of the organisations.</p>
Lateral Movement	<p>Lateral movement was considered to be largely successful, again due to the combined nature of the component parts of the malware solution. Again, some allowance has been made for the execution of the malware not working for reasons unknown, therefore a population exposure of 70% for all risk categories was considered.</p>	<p>Similar to privilege escalation the malware components are considered to be highly capable of moving laterally and executing without disruption or detection. Some more well-prepared companies may be able to identify and contain the threat but the variance for success were still considered high (70% for RC1 and 90% for RC7)</p>
Impact	<p>Of those remaining within the infected population it is expected all will be affected by the impact of the malware.</p>	<p>With the malware acting quickly the scenario does not suggest a long waiting period before the payloads impact the companies infected. Going largely undetected we consider the success factor to remain high (70% for RC1 and 100% for RC7)</p>

Scenario 3 — Demeter's Curse

Attack path step	Population exposed	Success factor
Initial Access	<p>This scenario is focused on a selection of industrial sectors and the population exposed is linked to a vulnerability in a system commonly used by many companies in the impacted industrial sectors.</p> <p>As the system in question is concentrated in particular sectors, we follow a premise that a higher market share, and therefore potential exposed population, is appropriate versus Scenario 1, whereby software is assumed to be more generic and all-sector. Therefore, a prudent value chosen for this scenario is 50%.</p>	<p>The success factor for this scenario is relatively higher than Scenario 2, for example, given the high capability actors envisioned for this scenario, as the actors are assumed to have specialist knowledge of the industrial sectors being targeted, including knowledge of particular IT and OT systems that support those sectors, and the associated degree of dependency on those systems to maintain business operations.</p> <p>Therefore, they are able to build a specific initial access attack step to compromise companies. The success factor is based on an inverse ratio with defence capabilities, as represented by risk category. The success factor ranges from 50% for RC1 to 75% for RC7.</p>
Privilege Escalation	<p>Of the initial vector it was considered a high number of the population to be further exposed to privilege escalation, however some allowance was made for the malware either not acting as intended or being interrupted, therefore the exposure figure was set to 75% for all risk categories.</p>	<p>For this step we assume significant variations in the success factors based on risk category, ranging from 25% for RC1 to 75% for RC7. This is to reflect the multitude of preventative and detective controls that may mitigate progress by a threat actor at this point (e.g. privileged identify and access management capabilities). The significant variation also accounts for the fact that it is this step in the attack path that we assume is likely to trigger some more widespread detection of attacks at an aggregate level and some known indicators of compromise might start to be shared across the security community, even if attacks are linked to each other.</p> <p>Depending on initial access vector that was successful for a particular compromised company, the threat actor may already have a degree of privileged access to certain systems and this is factored into the success factor values.</p>
Lateral Movement	<p>Lateral movement was considered to be largely successful, again due to the combined nature of the component parts of the malware solution. Again, some allowance has been made for the execution of the malware not working for reasons unknown, therefore a population exposure of 70% for all risk categories was considered.</p>	<p>Lateral movement success factors vary significantly – from 10% for RC1 to 50% for RC7. This wide variety reflects the particular difficulties of lateral movement from the initial access location to more business-critical systems (which are revenue-generating or support those critical business services), and the increased risk of detection and response associated with this attack path step. Regardless of whether zero-day vulnerabilities are deployed at this point, companies with better defence capabilities may identify indicators of compromise, either through detective actions or proactive threat hunting. The particularly low success factors for RC1, for example, also reflects a nuance of this scenario in that it is industrial sectors that are targeted by the threat actor and our premise is that the largest outage periods will be caused by actions that can compromise operational technology (OT) systems. Therefore, lateral movement activities may include attempts to cross between OT and IT environments, which may be blocked by effectively deployed security practice (e.g. network segmentation, zero trust environments) by those companies with the best defence capabilities (e.g. RC1). Additionally, given the footprint in this scenario, little to no human involvement by threat actors is assumed during the attacks on most compromised companies, we assume that lateral movement, particularly between OT and IT environments is especially difficult where you do not have human operators driving the cyber-attack 'hands-on-keyboard'.</p>
Impact	<p>Of those remaining within the infected population it is expected all will be affected by the impact of the malware.</p>	<p>If the attack reaches this step in the attack path within a company a high degree of success is assumed (i.e. at least some degree of business interruption occurs), even within RC1 companies (70% success factor) and especially within RC7 companies (90% success factor).</p> <p>Detection and response capabilities of RC1 companies, for example, provide some of the rationale for lower success factors, whilst recovery capabilities will help to manage the length of the outage period.</p>

	Initial Access					Privilege Escalation				
	Technique		Population Exposed	Success Factor	% of risks impacted after this step	Technique		Population Exposed	Success Factor	% of risks impacted after this step
Event 1	Software Supply Chain	RC1	30%	60%	18%	Established with Initial Access	RC1	100%	90%	16%
		RC2	30%	60%	18%		RC2	100%	92%	17%
		RC3	30%	52%	16%		RC3	100%	93%	15%
		RC4	30%	44%	13%		RC4	100%	95%	13%
		RC5	30%	36%	11%		RC5	100%	97%	10%
		RC6	30%	28%	8%		RC6	100%	98%	8%
		RC7	30%	20%	6%		RC7	100%	100%	6%
Event 2	Phishing	RC1	75%	10%	8%	Zero day	RC1	75%	60%	3%
		RC2	75%	18%	14%		RC2	75%	63%	7%
		RC3	75%	27%	20%		RC3	75%	67%	10%
		RC4	75%	35%	26%		RC4	75%	70%	14%
		RC5	75%	43%	33%		RC5	75%	73%	18%
		RC6	75%	52%	39%		RC6	75%	77%	22%
		RC7	75%	60%	45%		RC7	75%	80%	27%
Event 3	Exploit of Public-facing App/External Remote Services	RC1	50%	50%	25%	Zero day	RC1	75%	25%	5%
		RC2	50%	54%	27%		RC2	75%	33%	7%
		RC3	50%	58%	29%		RC3	75%	42%	9%
		RC4	50%	63%	31%		RC4	75%	50%	12%
		RC5	50%	67%	33%		RC5	75%	58%	15%
		RC6	50%	71%	35%		RC6	75%	67%	18%
		RC7	50%	75%	38%		RC7	75%	75%	21%


	Lateral Movement					Impact					Avg BI time (days)		Final Infection Rate
	Technique		Population Exposed	Success Factor	% of risks impacted after this step	Technique		Population Exposed	Success Factor	% of risks impacted after this step			
Event 1	Initial Access already widespread	RC1	100%	50%	8%	Data destruction/wipe	RC1	100%	70%	5.70%	RC1	5	5.67%
		RC2	100%	56%	9%		RC2	100%	73%	6.70%	RC2	5	6.75%
		RC3	100%	62%	9%		RC3	100%	77%	7.00%	RC3	7	6.95%
		RC4	100%	68%	9%		RC4	100%	80%	6.80%	RC4	9	6.82%
		RC5	100%	73%	8%		RC5	100%	83%	6.30%	RC5	15	6.33%
		RC6	100%	79%	7%		RC6	100%	87%	5.70%	RC6	18	5.68%
		RC7	100%	85%	5%		RC7	100%	90%	4.60%	RC7	21	4.59%
		Event 2	Remote service	RC1	70%		70%	2%	Data encrypted	RC1	100%	70%	1.20%
RC2	70%			73%	3%	RC2	100%	75%		2.50%	RC2	5	2.51%
RC3	70%			77%	5%	RC3	100%	80%		4.30%	RC3	7	4.29%
RC4	70%			80%	8%	RC4	100%	85%		6.60%	RC4	8	6.56%
RC5	70%			83%	10%	RC5	100%	90%		9.40%	RC5	10	9.38%
RC6	70%			87%	14%	RC6	100%	95%		12.80%	RC6	12	12.84%
RC7	70%			90%	17%	RC7	100%	100%		17.00%	RC7	14	17.01%
Event 3	Remote service			RC1	70%	10%	0%	Data destruction/wipe		RC1	100%	70%	0.20%
		RC2	70%	17%	1%	RC2	100%		73%	0.60%	RC2	5	0.58%
		RC3	70%	23%	1%	RC3	100%		77%	1.10%	RC3	8	1.15%
		RC4	70%	30%	2%	RC4	100%		80%	2.00%	RC4	12	1.97%
		RC5	70%	37%	4%	RC5	100%		83%	3.10%	RC5	17	3.11%
		RC6	70%	43%	5%	RC6	100%		87%	4.70%	RC6	23	4.67%
		RC7	70%	50%	7%	RC7	100%		90%	6.60%	RC7	30	6.64%

Outage days

For the outage days we are following a similar approach as above for the footprint. Each of the three scenarios leads to a specific outage for each of the risk categories. We are assuming a LogNormal distribution for the outage times and therefore need two points to parameterize the distribution. We have estimated the average outage time for each specific risk category in each specific event based on a combination of expert judgement, prior events and extrapolation of existing claims data from Beazley (represented in the chart below).

This dataset includes attritional claims and so the 50% percentile of 2.5 days from this dataset was judged to be too low set against the extremity of the scenarios considered. To set the shape of the distribution, we assume that twice the average outage time is the 90% percentile of the distribution.

The selection that the 90% percentile is twice the average outage days was considered a reasonable approximation based on the expert judgement selections and the distribution of normal BI losses from two different claims data sets.



Days_Outage calculated as 'loss cost/(revenue per day * gross margin)'



One assessment found that when looking at the total incurred loss as a proportion of insured revenue for attritional claims (for incurred values and also considering the full known towers), the 90% percentile was around 2–2.5 times the mean which supports the assumption made — this has held true to claims for companies with revenue above \$100m and also where we have been able to identify the claims as malware related. We caution however that this data was based on attritional claims for, often targeted, attacks. It is also based on total incurred values rather than just the BI component, but this was the best data available. A parallel in natural catastrophe events however can be made that the distribution of loss from all magnitude earthquakes is broader than the loss distribution from one specific earthquake and thus we do not consider this unreasonable to help validate our distribution spread assumption.

Given the limited historical event data, much of the historical claim analysis has been performed on attritional claims. We considered how these claims may differ to a systemic attack and whilst we debated a number of principles that may apply in a systemic event, the reality is unknown. In terms of principles, we considered a number of components, for example:

- ‘Demand surge’ may play a part in lengthening certain restoration services and hardware supplies.
- Systemic attacks can be more indiscriminate in nature, given the large footprint and fact that direct targeting of critical business services in compromised firms is more challenging, particularly if the attack path is more automated and there is little or no human operator (‘hands-on-keyboard’) involvement. This may mean that the outage period for revenue-generator services is reduced.
- In response to mass events, the cyber security community has often been collaborative and shared indicators of compromise and effective response and recovery options. This may reduce response time and costs.

Explanations for the rationale of the outage days parameters is detailed below, split by scenario.

Scenario 1 — Autolycus

This parameter represents the average business interruption (BI) time for each risk category. The average BI time in this event ranges from 5 (RC1) to 21 days (RC7). It is key to note that this is an average only and that BI could therefore be longer or shorter at each risk category level.

The parameter was informed by expert consultation and an assessment of the mitigation strategies available to different companies (dependent on their cyber maturity), these assumptions represent the current capabilities of impacted parties to recover from a supply chain attack. Longer durations are likely to be attributed to difficulty in replacing hardware.

Scenario 2 — Lernaean Hydra

Despite the high capability of the malware the outage days for business interruption were more determined by the perceived abilities for the companies to respond to the incident. Therefore, it was considered that for RC1 companies the outage time would be relatively low (2 days), compared to RC7 with an outage period of 14 days. This reflects resources, execution of response and restoration from backups performed by the affected companies.

Scenario 3 — Demeter’s Curse


The range of outage days for this scenario is between 2 days (RC1) to 30 days (RC7). This scenario has the widest range of all the scenarios, as well as the highest upper bound for the average. The width in range reflects the diverse attack paths that could occur within this scenario and the degree to which operational technology (OT) systems are impacted. For lower risk categories, there is assumed to be a greater degree of operational impact, combined with less mature recovery capabilities, which equate to an extended outage. In particular, it is assumed that there may need to be physical recovery or restoration in certain cases, given the OT impact (e.g. engineers travelling to remote sites), as well as difficulties in replacing hardware, which could prolong the outage.

Appendix 4: Impact of Portfolio Composition


Below are the results of the sensitivity analyses of the three scenarios with the unstressed parameters as the baseline:



Event 2



Event 3



Appendix 5:

Why the Model is Not Probabilistic

There has been a trend from deterministic to probabilistic modelling and reporting of cyber-risk appetite. For example, in November 2023, Beazley announced that it was moving to the 1:250 OEP⁵³ net loss figure from its internal model as the basis of its internal risk appetite⁵⁴. Likewise, Munich Re has long set its appetite at a 1:1,000 level developing proprietary models to measure this. A key challenge that might be posed to the modelling in this paper is: why is it not probabilistic? The Partnership discussed the merits of such an approach but quickly ruled it out.

It is also worth considering the theoretical implications of deploying a probabilistic model. Rather than a single parameter for each input corresponding to, for example, different risk categories, a curve of inputs would be needed for each risk category, scenario parameter and success factor at a minimum. Additionally, the expected frequency for each scenario would need to be specified as a distribution. The Partnership concluded that such efforts would not deliver any additional benefits in terms of explanatory power of systemic cyber-risk from malware and would simply create an illusion of sophistication.

That is not to say that the Partnership does not believe there is value in probabilistic estimates, but rather that it would not be beneficial for the project. Based on the Partnership's knowledge of existing models, it has concluded that the developed scenarios in the model are sufficiently remote as to lie in the tail of systemic cyber-risks and that this is sufficient information for users of the model to make judgements based on the scenario and model conclusions as they see fit.

To achieve the aim of quantifying extreme events it is not necessary to introduce all the complications that come with the parameters mentioned above. As we have seen with this model the results will be sensitive to each parameter which is selected.



Footnotes

¹Contingent Business Interruption.

²[White_Paper_GAC_Cyber_Resilience_VERSION_2.pdf](#) (weforum.org).

³See Appendix 1 for definitions/a glossary of terms used in this paper.

⁴[WEF_GFC_Cybersecurity_2022.pdf](#) (weforum.org).

⁵[WEF_GFC_Cybersecurity_2022.pdf](#) (weforum.org).

⁶[Cyber-risk Accumulation: Fully tackling the insurability challenge](#) (genevaassociation.org).

⁷[Counting the economic cost: How vulnerable could you be? — Lloyd's](#) (lloyds.com); [Lloyd's finds major hack of a payments system could cost \\$3.5tn](#) (ft.com); [Major cyber attack could cost the world \\$3.5 trillion — Lloyd's of London | Reuters](#).

⁸[METHODOLOGICAL PRINCIPLES OF INSURANCE STRESS TESTING](#) (europa.eu).

⁹For example, see the 2023 report from Guy Carpenter on the cyber market – [Guy_Carpenter_Cyber_\(Re\)insurance_Market_Report_Publish_rev.pdf](#) (guycarp.com).

¹⁰For example, in the UK see the PRA consultation paper from 2023 (CP26/23 – Operational resilience: Critical third parties to the UK financial sector | Bank of England). In the EU, see the Digital Operational Resilience Act on critical third-party service providers ([ESAs publish first set of rules under DORA for ICT and third-party risk management and incident classification | European Banking Authority](#) (europa.eu)).

¹¹[Thematic findings from the 2022 cyber stress test](#) (bankofengland.co.uk).

¹²[METHODOLOGICAL PRINCIPLES OF INSURANCE STRESS TESTING](#) (europa.eu).

¹³New analysis highlights strength of Ukraine's defence... – [NCSC.GOV.UK; ECCRI_REPORT_The-Cyber-Dimensions-of-the-Russia-Ukraine-War-19042023.pdf](#). See also the NATO Cooperative Cyber Defence Centre of Excellence cyber threat landscape report looking ahead to 2030 (written in 2020), which also notes a number of similar sentiments, even though written prior to the current conflict in Ukraine – [https://ccdcoe.org/uploads/2020/12/Cyber-Threats-and-NATO-2030_Horizon-Scanning-and-Analysis.pdf](#).

¹⁴[The near-term impact of AI on the cyber threat](#) – NCSC.GOV.UK.

¹⁵[The vital role of capital in cyber \(re\)insurance](#) (ajg.com), [Cyber Insurance: Risks and Trends 2024](#) | Munich Re.

¹⁶Gallagher Re Industry Exposure Database.

¹⁷[The vital role of capital in cyber \(re\)insurance](#) (ajg.com).

¹⁸Gallagher Re Industry Exposure Database.

¹⁹External third-party experts were used from across a number of different organisations and fields included cyber underwriting, cyber security and resilience, cyber exposure management, public policy and operational specialists in certain industry sectors.

²⁰See, for example, Lloyd's many collaborations with third parties on the topic, such as with the Cambridge Centre for Risk Studies (CRS), part of Cambridge University's Judge Business School. Examples include: [CyRiM Scenario: Bashe Attack — Technology and space — Cambridge Judge Business School](#); [crs-shen-attack-cyber-risk-in-asia-pacific-ports.pdf](#) (cam.ac.uk); [crs-sybil-logic-bomb-cyber-catastrophe-stress-test.pdf](#) (cam.ac.uk).

²¹For example, [Cyber-risk Accumulation: Fully tackling the insurability challenge](#) (genevaassociation.org).

²²For example, [White_Paper_GAC_Cyber_Resilience_VERSION_2.pdf](#) (weforum.org) or [WEF_GFC_Cybersecurity_2022.pdf](#) (weforum.org).

²³[https://www.tandfonline.com/doi/abs/10.1080/10920277.2022.2034507](#).

²⁴[https://www.tandfonline.com/doi/abs/10.1080/03461238.2021.1872694](#).

²⁵[https://link.springer.com/chapter/10.1007/978-3-030-97124-3_23](#).

²⁶See for example, [General Insurance Stress Test 2022 – Scenario Specification Guidelines Instructions Final V2](#) (bankofengland.co.uk); [Realistic Disaster Scenarios – Lloyd's](#) (lloyds.com); [2022cyberscenariotechnicalspecificationsv13.pdf](#) (nbb.be).

²⁷See, for example, the Lloyd's collaborations with CRS mentioned in footnotes above.

²⁸Please note that no raw real-world claims data was disclosed by any partner to another; any claims data submitted by a partner was duly anonymized and aggregated prior to disclosure.

²⁹[Thematic findings from the 2022 cyber stress test](#) (bankofengland.co.uk).

³⁰Alongside being used in the cyber security industry, MITRE ATT&CK is also used by regulators – for example, the UK Financial Authorities use it to structure tactics, technique and procedures used in intelligence-led penetration tests against UK firms, as part of the CBEST scheme.

³¹Woo, Gordon, et al. *Emerging Risk Report 2017 Understanding Risk Reimagining History Counterfactual Risk Analysis*. 2017.

³²EIOPA, for example, stated 'Larger undertakings might have a better cyber hygiene leading to a faster recovery' – [METHODOLOGICAL PRINCIPLES OF INSURANCE STRESS TESTING](#) (europa.eu).

³³Autolycus, WSSC, Scenario 1 and Event 1 are terms used interchangeably.

³⁴Lernaean Hydra, SPM, Scenario 2 and Event 2 are terms used interchangeably.

³⁵Demeter's Curse, TILE, Scenario 3 and Event 3 are terms used interchangeably.

³⁶Eling et al (2023), *The Economic Impacts of Extreme Cyber-risk Scenarios*, [https://www.tandfonline.com/doi/pdf/10.1080/10920277.2022.2034507](#).

³⁷[Illuminating cyber crime](#) (lloyds.com).

³⁸A detailed, technical explanation of the event from the perspective of CrowdStrike is provided here: [https://www.crowdstrike.com/wp-content/uploads/2024/08/Channel-File-291-Incident-Root-Cause-Analysis-08.06.2024.pdf](#)

³⁹[Helping our customers through the CrowdStrike outage — The Official Microsoft Blog](#)

⁴⁰[https://beazley.com/en-US/cyber-and-breach-response-portal/cyber-basics/#:-:text=Waiting%20periods%20of%2010%20to,waiting%20periods%20may%20be%20available.&text=Cyber%20business%20interruption%20coverage%20may,dependent%20or%20contingent%20business%20interruption](#).

⁴¹[Merck Settles With Insurers Over \\$700m NotPetya Claim - Infosecurity Magazine](#) (infosecurity-magazine.com).

⁴²For example, the cyber-attack impacting the software provider Blackbaud.

⁴³[FACT SHEET: The Biden-Harris Administration Has Launched an All-of-Government Effort to Address Colonial Pipeline Incident | The White House](#).

⁴⁴[Cyber resilience act: Council and Parliament strike a deal on security requirements for digital products - Consilium](#) (europa.eu).

⁴⁵For example, see the 2021 report by Lloyd's, CyberCube and Guy Carpenter [Microsoft Word - The Emerging Cyber Threat to Industrial Control Systems Final PDF_12022021](#) (lloyds.com).

⁴⁶See for example [2021 Trends show increased globalised threat of ransomware.pdf](#) (ncsc.gov.uk).

⁴⁷[METHODOLOGICAL PRINCIPLES OF INSURANCE STRESS TESTING](#) (europa.eu).

⁴⁸[https://www.cambridge.org/core/journals/annals-of-actuarial-science/article/multivariate-hawkes-process-for-cyber-insurance/6A007F9E2CCDA46919EDF79AC791A08D](#).

⁴⁹Some requirements exist for reporting cyber incidents (e.g. new Security Exchange Commission rules – see [https://www.sec.gov/news/statement/gerding-cybersecurity-disclosure-20231214](#)) but reported incidents do not always provide useful data for cyber insurance modelling purposes.

⁵⁰As with all other real-world costs and data sets, no individual, specific client data was disclosed by either Munich Re or Beazley for the purposes of preparing these average figures, and all cost estimates were duly anonymized and smoothed prior to disclosure.

⁵¹[https://www.stern.nyu.edu/-adamodar/New_Home_Page/data.html](#).

⁵²The 'Public Authority; NGOs; Non-Profit' rates have been matched to the 'Other' rates as the dataset used did not include equivalent data for the former.

⁵³Occurrence exceedance probability – the probability of a single event of a particular size or greater occurring in a given period.

⁵⁴[https://www.beazley.com/globalassets/ir-documents/presentations/2023/capital-markets-day-presentation-november-2023.pdf](#).

All references to the 'The Partnership' are to be understood as references to a collaborative partnership in the colloquial sense. Beazley, Munich Re and Gallagher Re did not establish any legal partnership, joint venture or similar for the purposes of producing this paper and none of the participants are constituted the agent of another or otherwise authorised to act on another participant's behalf. Without limiting the foregoing, no Author shall have any responsibility for any act or omission of any other Author

© 2024, Beazley Furlonge Limited ("Beazley"), Münchener Rückversicherungs-Gesellschaft Aktiengesellschaft in München ('Munich Re') and Arthur J. Gallagher (UK) Limited ("Gallagher").

All rights reserved.

The contents of this paper (including, without limitation, the text, pictures and graphics) are protected under copyright law and other protective legislation.

The contents herein are provided for informational purposes only and do not constitute and should not be construed as professional advice. Any and all examples used herein are for illustrative purposes only, are purely hypothetical in nature, and offered merely to describe concepts or ideas. They are not offered as solutions for actual issues or to produce specific results and are not to be relied upon. The reader is cautioned to consult independent professional advisors of their choice and formulate independent conclusions and opinions regarding the subject matter discussed herein. Beazley Furlonge Limited ('Beazley'), Münchener Rückversicherungs-Gesellschaft Aktiengesellschaft in München ('Munich Re') and Arthur J. Gallagher & Co. and subsidiaries ('Gallagher'), (each an 'Author'), are not responsible for the accuracy or completeness of the contents herein and expressly disclaims any responsibility or liability based on any legal theory or in any form or amount, based upon, arising from or in connection with for the reader's application of any of the contents herein to any analysis or other matter, nor do the contents herein guarantee, and should not be construed to guarantee any particular result or outcome.

Beazley Furlonge Limited (Company Registration Number: 01893407 and VAT Number: 649 2754 03) is a managing agent for Syndicates at Lloyd's and is authorised by the Prudential Regulation Authority and regulated by the Financial Conduct Authority and the Prudential Regulation Authority (Firm Reference Number: 204896). Beazley Furlonge Limited is registered in England and Wales with its Registered Office at 22 Bishopsgate, London EC2N 4BQ.

Email: info@beazley.com Tel: +44 (0)20 7667 0623 Fax: +44 (0)20 7082 5198

Münchener Rückversicherungs-Gesellschaft Aktiengesellschaft in München (Commercial Register Munich Number: HRB 42039) is supervised by the German Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin). Munich Re is a reinsurance company organized under the laws of Germany with its Registered Office at Koeniginstrasse 107, 80802 Munich. In some countries, including in the United States, Munich Reinsurance Company holds the status of an unauthorized reinsurer.

Email: contact@munichre.com Tel : +49 (89) 38 91 -0

Gallagher Re is a trading name of Arthur J. Gallagher (UK) Limited, which is authorised and regulated by the Financial Conduct Authority. Registered Office: The Walbrook Building, 25 Walbrook, London EC4N 8AW. Registered in England and Wales. Company Number: 1193013. www.ajg.com/uk. FPI179-2024, Expiry 071025

