

Cloud Down
Impacts on the
US economy

Lloyd's of London disclaimer

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

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

© Lloyd's 2018
All rights reserved

AIR disclaimer

This report has been co-produced by Lloyd's and AIR Worldwide (AIR) for general information purposes only. While care has been taken in gathering the data and preparing the report AIR makes no representations or warranties as to its accuracy or completeness and expressly excludes to the maximum extent permitted by law all those that might otherwise be implied. AIR accepts no responsibility or liability for any loss or damage of any nature occasioned to any person as a result of acting or refraining from acting as a result of, or in reliance on, any statement, fact, figure or expression of opinion or belief contained in this report. This report does not constitute advice of any kind.

About Lloyd's

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

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

About AIR Worldwide

AIR Worldwide (AIR) provides risk modeling solutions that make individuals, businesses, and society more resilient to extreme events. In 1987, AIR Worldwide founded the catastrophe modeling industry and today models the risk from natural catastrophes, terrorism, pandemics, casualty catastrophes, and cyber attacks, globally. Insurance, reinsurance, financial, corporate, and government clients rely on AIR's advanced science, software, and consulting services for catastrophe risk management, insurance-linked securities, site-specific engineering analyses, and agricultural risk management. AIR Worldwide, a Verisk (NASDAQ:VRSK) business, is headquartered in Boston with additional offices in North America, Europe, and Asia. For more information, please visit www.air-worldwide.com.

Key contacts

Trevor Maynard

Head of Innovation

trevor.maynard@lloyds.com

Scott Stransky

Assistant Vice President & Principal Scientist

sstransky@air-worldwide.com

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

Acknowledgements

AIR project team and area of expertise

- Dr Carol Aplin, Senior Scientist, Cyber
- Mark Banks, Business Development Executive
- Gian Calvesbert, Senior Product Marketing Manager
- Dr Eric Dallal, Scientist, Cyber
- Vineeta Gabriel, Scientist, Cyber
- Dr Tomas Girnius, Manager & Principal Scientist, Cyber
- Dr Jayanta Guin, Executive Vice President & Chief Research Officer
- Catherine Jeannette, Senior Technical Writer
- Nan Ma, Marketing Strategist
- David Pigott, Senior Risk Consultant
- Patricia Stevens, Risk Consultant
- Scott Stransky, Assistant Vice President & Principal Scientist, Cyber

AIR external partners

- Frank Cilluffo, Associate Vice President & Director, Center for Cyber and Homeland Security, George Washington University
- Nate Lesser, Senior Fellow, George Washington University

Lloyd's project team

- Dr Trevor Maynard, Head of Innovation
- Dr Keith Smith, Innovation team
- Anna Bordon, Innovation team
- Linda Miller, Marketing and Communications
- Flemmich Webb, Speech and Studies
- Nathan Hambrook-Skinner, Marketing and Communications
- Lizzie Lowe, Marketing and Communications

Further thanks go to the following for their expertise, feedback and assistance with the study:

Verisk cyber team

- Prashant Pai, Vice President, Cyber Strategy
- Caitlin Plunkett, Cyber Lead, Commercial Lines Coverage Products
- Stephen Whelan, Director, Product Management Insurance Coverages

Lloyd's Market Association

- Mel Goddard, Market Liaison Director
- Tony Ellwood, Senior Technical Executive – Underwriting

The following people took part in workshops or roundtables, or commented on earlier drafts of the report; we would like to thank them all for their contributions:

Insurance industry workshops and consultation

- Tom Allen, Channel 2015
- Scott Bailey, Markel
- David Baxter, Barbican
- Marcus Breese, Hiscox
- Stephanie Bristow, Hiscox
- Robert Brown, Neon
- Wesley Butcher, Atrium
- Danny Clack, Pembroke
- Jason Clark, Faraday
- Nils Diekmann, MunichRe
- Daniel Fletcher, QBE
- Matt Harrison, Hiscox
- Matthew Hogg, Liberty
- Adam Holdgate, AM Trust
- Jerry Hyne, Aegis
- Laila Khudairi, Tokio Marine Kiln
- Nick Leighton, Aegis
- Alessandro Lezzi, Beazley
- Ben Maidment, Brit
- Kelly Malynn, Beazley
- Phil Mayes, Talbot
- Alastair Nappin, MunichRe
- Raheila Nazir, Aspen
- Matt Northedge, AM Trust
- Andrew Pearson, Barbican
- Scott Sayce, AXA
- David Singh, MS Amlin
- Dan Trueman, Novae
- Stephen Wares, MS Amlin

Contents

Executive summary.....	5
1. Cloud computing trends.....	9
2. Modelling approach.....	13
3. Industry exposures.....	17
4. Scenario classes description	21
5. Analysis.....	27
6. Implication of cloud failure on (re)insurance	43
7. Applicability of modelling methodology for other scenario analyses	46
8. Conclusion	49
Appendix A. Historical cloud events	51
Appendix B. Cloud resilience	53
Appendix C. E-business factors.....	56
References.....	57

Executive summary

The use and adoption of cloud computing services is proliferating throughout society and it is no coincidence that cyber risk is increasing as well. Not only there are more companies relying on “the cloud” to operate their businesses but economies of scale have created a select few cloud service providers that dominate the market. This reliance on a relatively small number of companies has resulted in systemic risk for businesses using their services. In the event of sustained downtime of a top cloud service provider, simultaneous damage for all its clients and dependents could lead to catastrophic financial losses. According to McKinsey & Company (*Elumalai, Starikova, and Tandon, 2016*), as of 2015, 77% of global companies used traditionally built IT infrastructure (i.e. with computers and servers set up on premises) as the primary environment for at least one workload (i.e. a computing task); this is forecast to drop to 43% in 2018. While only about 25% of companies in 2015 used public infrastructure as a service as the primary environment for at least one workload, that percentage is expected to rise to 37% in 2018.

Following Lloyd’s previous study, *Counting the cost: Cyber risk decoded*, this study analyses cloud service provider failure risk and specifically highlights the expected financial impact of such an event on 12.4 million businesses in the US, the most established cyber insurance market for this emerging line of business.

The insurance industry has been asked by Lloyd’s, regulators, and its own senior management to understand its exposure to this type of cyber risk. To address this, AIR has developed a comprehensive database of industry exposures that provides the information insurers need for accurate modelling and has used it to form the basis of the alternative modelling approach described in this report.

The results of this cloud downtime scenarios analysis could help insurance managers gain insights into how to grow their cyber business in a controlled and prudent manner.

Methodology and approach

Unlike natural disaster risk, which can be aggregated using easily verifiable information such as geographic location, cyber risk aggregates around sources of risk such as third-party IT providers or software vulnerabilities present in the organisation’s systems. This information is hard to capture at the point of underwriting and may not be transferred to the portfolio management level. Technologies are now available that can use external data sets to evaluate a company’s exposure to cyber risk.

This paper provides estimates of e-business interruption costs to the full set of United States companies and the subset of Fortune 1000 companies that arise from the sustained loss of access to a business service, namely a cloud service provider. These e-business interruption costs are modelled using data from the US Census bureau and include costs from e-commerce sales/turnovers, e-shipments, m-commerce sales/shipments and electronic order management systems. The US Census Bureau states that: “E-commerce sales/turnovers are sales of goods and services where the buyer places an order, or the price and terms of the sale are negotiated over the Internet, mobile device (m-commerce), Extranet, Electronic Data Interchange (EDI) network, electronic mail, or other comparable online system. E-commerce shipments (e-shipments) are online orders accepted for manufactured products from customers, including shipments to other domestic plants of the same company for further manufacture, assembly, or fabrication where price and terms of sale are negotiated over the Internet, Extranet, EDI network, electronic mail, or other online system. Payment may or may not be made online.” (*US Census Bureau, 2016*.) The term “e-business”, as used in this paper, is defined in Appendix C.

The results published in this report are based on the top 15 cloud providers in the US, which account for a 70% market share.

This report describes another approach to modelling cyber aggregation risk that uses company specific risk attributes. Detailed accumulation approaches differ from market share approaches because the underlying database of exposures means the modelled loss reflects the true risk insurers are exposed to. By identifying which insureds companies would be impacted by the scenario and omitting those that would not, detailed accumulation approaches are distinct from other approaches that only use broad assumptions such as a provider's industry market share. These market share statistics give no indication as to which organisations are at risk, meaning that only generic scenarios can be created. By contrast, the scenario classes presented in this report considers the impact of disruption to several key cloud service providers for different periods of time. Losses can be split accordingly, providing a deeper understanding of the actual risk. One benefit of this approach is that it provides a framework for measuring the systemic risk associated with any vulnerability that may be common across a group of organisations; it is *not* limited to the analysis of service provider business interruption.

A multiple scenario approach

This report examines multiple scenarios that completely disrupt a cloud service provider in the US, leaving all their clients with no access to the information technology services their businesses rely on. The previous report looked in detail at one scenario leading to cloud provider service failure. This report provides four threat sources and more than 30 additional vectors that could lead to a cloud service provider failure. There are multiple ways to bring down a cloud service provider, and some attacks can be combined assaults (e.g. DDoS attack plus malware plus theft). However, this analysis is agnostic to the causes of downtime – it may arise from environmental, adversarial, accidental, or structural vectors. This report provides some commentary on the causes of such events but its focus is on what happens in the aftermath. The detailed accumulation approach provides estimates for a number of scenario variants, each representing a specific cloud service provider and different outage durations.

Key findings

This report draws the following key conclusions about the various approaches to estimating systemic risk and the losses that can be expected from the sustained downtime of a major cloud service provider:

- The cyber insurance market is still developing, and can be characterised by relatively low take-up rates and coverage limits. As a result, there is a significant difference between the ground-up losses and industry insured losses. This means there is an opportunity for the insurance industry to help society to prepare for and recover from extreme scenarios of cyber risk aggregation.
- The business interruption losses associated with the disruption of a cloud service provider are varied and depend on how many businesses use its services in the US market and the duration of the downtime event. Given the state of the cyber insurance industry today, a cyber incident that takes a top three cloud provider offline in the US for 3-6 days would result in ground-up loss central estimates between \$6.9 and \$14.7 billion and between \$1.5 and \$2.8 billion in industry insured losses. A cyber incident that takes offline a cloud provider that has between the 10th and 15th highest market share in the US for 3-6 days, would result in ground-up loss central estimates between \$1.1 billion to \$2.1 billion and between \$220 million and \$450 million in industry insured losses. These insured loss values are based on affirmative cyber policies only.
- Fortune 1000 companies will carry 37% of the ground-up losses and 43% of the insured losses arising from a 3-6 days downtime event. Smaller companies might be more likely to use the cloud in order to avoid building the business infrastructure in-house, but insurance take-up is low compared to the Fortune1000 companies. As the cyber insurance market grows rapidly, the distribution of risk will have to be monitored carefully.
- A cyber incident that takes down a top three cloud service provider for 3-6 days would result in \$4.2-\$8.6 billion of ground up losses for the manufacturing industry, followed by \$1.4-\$3.6 billion for the wholesale and retail trade industry. These two industries will be the most affected, which holds true for the Fortune 1000 companies as well.

- A comparison of the loss estimates obtained by using the detailed accumulation and market share approaches indicates there is agreement between them at the broad industry level. However, as the scope of organisations evaluated becomes narrower, differences become apparent. Identifying these is critical for risk management purposes. For a hypothetical cyber portfolio and a 3-6 day downtime event, ground-up losses could be in the range of \$640 million - \$1 billion (mean: \$850 million) for the detailed accumulation and \$870 million - \$1.4 billion (mean: \$1.1 billion) for the market share approaches, respectively, a 32% difference. For other specific insurance portfolios or sectors, the difference could be even larger.

Conclusion

Cyber insurance is an emerging market that is outperforming most existing lines of business but this growth track can only be sustained if society's understanding of the nature of risk continues to grow as well.


This report's findings suggest that disruption of a cloud service provider in the US market will significantly impact the manufacturing and retail trade industry, due to their heavy reliance on cloud services, and that Fortune 1000 companies would carry almost half of the insured losses.

The analysis methodology outlined in this report can be used by the insurance sector to standardise and improve risk selection and portfolio management processes in order to inform decisions such as setting underwriting guidelines, deploying capital, and identifying risk transfer needs.

The detailed accumulation report highlights the importance and value of collecting high quality exposure data and having it at hand at the point of exposure management and portfolio analysis. Although there are use cases where market share approaches are viable, (re)insurers who are looking to differentiate their view of the risk should strive to invest in processes for collecting and incorporating detailed exposure risk data.

Lloyd's and AIR hope this report and the discussion it will provoke are a further step towards creating an insurance sector that is more resilient to systemic cyber risk.

Cloud computing trends



1. Cloud computing trends

Introduction

In this paper, the term “cloud” refers to the technologies that allow people to access computing resources from anywhere through the internet. The National Institute of Standards and Technology (NIST) describes the cloud as a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction” (*Mell and Grance, 2011*).

The growth of cloud computing has been driven by improvements in internet availability, reliability and performance, the commoditisation of the required hardware components, and the adoption of design practices such as hardware virtualisation and service-oriented software architectures.

These technology trends have given rise to a booming industry of companies offering cloud computing services in exchange for usage fees or for permissions or rights over the data being uploaded.

Cloud computing service models

There are several options for providing cloud computing resources through a service model that does not require the acquisition of the necessary hardware and software. Each model is unique in that it provides varying degrees of access and control over the underlying software and hardware. The three main cloud service models are:

Software as a Service (SaaS)	Platform as a Service (PaaS)	Infrastructure as a Service (IaaS)
<p>Clients are given access to software applications using a thin user interface, such as a web-browser that are supported in the backend by the service provider’s cloud infrastructure.</p> <p>The service provider is responsible for managing or controlling the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.</p>	<p>The cloud infrastructure is provided as a foundation for the client to deploy applications that were acquired or created using programming languages, libraries, services, and tools supported by the provider.</p> <p>Control of the underlying cloud infrastructure of networks, servers, operating systems, and storage is the responsibility of the service provider, with the user having control over the deployed applications and possibly configuration settings for the application-hosting environment.</p>	<p>The client is provided with computing resources, such as processing, storage and network capabilities, where they have the choice of deploying their desired software applications and operating systems.</p> <p>Management or control of the underlying cloud infrastructure is the service provider’s responsibility with the user having control over operating systems, storage, and deployed applications; and possibly limited control of certain networking components (e.g. host firewalls).</p>

Underlying every cloud solution is a network of physical servers located within a single site or across multiple locations. Who owns this infrastructure or can access the data within it is determined by the chosen cloud deployment model.

In order to meet the different requirements an organisation may have, several types of cloud deployment models have taken shape:

Private cloud	Community cloud	Public cloud	Hybrid cloud
The cloud infrastructure is used exclusively by one organisation. That same organisation may own and manage the cloud infrastructure and/or may outsource that responsibility to a third party. The infrastructure itself may exist on or off the client organisation's premises.	The cloud infrastructure is used exclusively by a specific group of organisations who may own and manage the cloud infrastructure and/or may outsource that responsibility to a third party. The infrastructure itself may exist on or off the client community's premises	The cloud infrastructure is made available to the general public with ownership and management responsibility of the cloud infrastructure falling on one organisation, typically for commercial interests. The infrastructure resides on the provider's premises	The cloud infrastructure is deployed using a combination of private or community and public resources. The resources are separate but are accessed by the user organisation using technology or other internal processes.

A cloud solution is typically architected with multiple regions, where a region is a geographical location where users can run their resources, and is typically made up of multiple zones. All major cloud providers have multiple regions, located across the globe, with Rackspace having the fewest at 7 regions and Microsoft Azure having the most at 36. Additionally, all major cloud providers have multiple regions within the United States.

A typical user sets up resources and data within one zone of a region, with failover capabilities to another zone within the same region. For added resiliency, a user may also set up resources to failover to another region, should their region experience a downtime event. Most cloud providers architect their cloud such that regions are independent of each other and zones within a region are isolated from each other.


Cloud computing adoption trends

Enterprises are shifting from a "build" to a "consume" paradigm for their information technology (IT) needs and that is driving the increasing adoption of cloud computing services. According to the McKinsey & Company (*Elumalai, Starikova, and Tandon, 2016*), as of 2015, 77% of companies used traditionally built IT infrastructure (i.e., with computers and servers set up on premises) as the primary environment for at least one workload, a percentage expected to drop to 43% in 2018.

While only about 25% of companies in 2015 used public infrastructure as a service as the primary environment for at least one workload, that percentage is expected to rise to 37% in 2018.


These business decisions are fuelling the growth of the public cloud service industry and it is projected that this industry's turnover will grow at a compound annual growth rate of 36% between 2014 and 2026. (Note that "turnover" and "revenue" are synonymous for the purposes of this paper.) See Figure 1 for turnover projections for the three cloud service model types, and Figure 2 for the top clouds as measured by adoption rate.

Figure 1: Public cloud vendor turnover projections




Source: Wikibon's Public Cloud Market Forecast 2015-2026 report (Finos, 2015)

Figure 2: Adoption of the top four enterprise public clouds



Source: 2017 State of the Cloud report (RightScale, 2017).

Modelling approach



2. Modelling approach

Cyber risks accumulate around sources of risk such as cloud providers. These sources of risk are challenging to identify because most insurers do not know which cloud vendors their insured customers use or to what extent. In a scenario where a cloud provider is disabled, a traditional market share approach provides a broad, relatively uncertain view of the risk. It assumes that if the cloud provider has 30% market share, 30% of the insurer's portfolio is affected. This might be true, or the portfolio might have more or fewer insured customers who use that cloud provider. Unless the insurer has painstakingly gathered this data, there is no way of knowing which companies would be affected by the outage or how much.


Detailed accumulation

In Figure 3, the circles represent hypothetical companies in an insurance portfolio. Circle colours indicate different industry types and circle sizes indicate total insured values. The market share approach illustrated on the left shows the segment of the portfolio that might experience business interruption if a cloud vendor goes down, based on the market share of that cloud vendor. The segment includes companies of all industry types and all sizes because the market share approach does not consider actual relationships of these companies to the disabled cloud vendor.

The detailed accumulation approach illustrated on the right provides a more accurate view of the risk by using data in the industry exposure database to identify relationships between specific vendors and insured companies. With this data, the same portfolio of seemingly unassociated exposures can be aggregated around four separate cloud providers.

Figure 3: The detailed accumulation approach identifies aggregation points

Market Share Approach



Detailed Accumulation Approach





Figure 4 shows the process for this aggregation analysis. First, the model uses the Industry Exposures to determine which common vendors are used by the insured companies. (The example shows that there are four common vendors.) Next, the model runs various periods of down time against each vendor to determine

accumulated ground up losses for each scenario. This paper presents the impact of downtime on financial losses to the economy and the insurance industry. Finally, the model applies insurance terms to the ground-up losses to determine the gross loss.

Figure 4: Detailed accumulation approach process




This process is important because most portfolios do not behave like the market as a whole. As an example, consider the Dyn DNS (Domain Name System) outage of October 2016. This was not a cloud downtime example, but it is nonetheless useful to illustrate the advantage of using a detailed accumulation approach. Dyn has an approximately 4% market share among DNS providers, represented by the blue bar in Figure 5.


From a large set of companies with DNS provider information, researchers repeatedly took random samples of 100 companies (a number chosen to be representative of a small cyber portfolio), and counted the number of those companies in each sample which used Dyn as their DNS provider.

While using a market share analysis would indicate that exactly four companies in every 100 company sample use Dyn, the analysis found that this was the case in only about 20% of the samples. In fact, about 20% (11% + 5% + 3% + 1%) of the samples (see dotted area in Figure 5) included six or more companies that used Dyn, making an underestimation of 50% or more just as likely as getting the correct result in a market share analysis. Figure 5 illustrates the study. The chart was formed by tabulating the number of companies that used Dyn for each random sample of 100 companies. For each number from 0 to 100, we determined the proportion of the random samples that contained exactly that number of companies using Dyn. The probabilities in the chart are simply these proportions.

Figure 5: Hypothetical distribution of companies impacted by Dyn outage



Industry exposures



3. Industry exposures

To demonstrate the detailed accumulation approach, a rich set of nearly 12.4 million cyber-specific exposures were used. This dataset includes information on US businesses such as industry, turnover, employee count, and location details, along with details on service provider usage and insurance policy terms like limits and waiting periods. The industry exposures were developed using data from the following companies:

Risk Based Security™ (RBS) provides historical privacy breach and incident data on more than 26,000 breaches, including industry-specific details on threat vectors and vulnerabilities and data breach information on business, industries, and geographies.

BitSight Technologies™ provides BitSight Security Ratings and supply chain or network connectivity data. The company's Security Rating Platform continuously analyses vast amounts of external data on security behaviours to provide objective, evidence-based security ratings on thousands of companies by industry, company size, and company headquarters location. BitSight Discover is a key element, as described below.

Nielsen provides turnover, headcount, and other data on businesses in the U.S.

U.S. Census freely makes available data on technology spending by industry and data on percent of turnover from e-business by industry (*US Census Bureau, 2015*).

Verisk Analytics collects and analyses billions of records, drawing on unique data assets and deep domain expertise. Verisk offers predictive analytics and decision support solutions to customers in rating, underwriting, claims, catastrophe and weather risk, global risk analytics, natural resources intelligence, economic forecasting, and many other fields. Verisk has information on cyber policies that was used in this report.

Yahoo Finance offers diverse financial information for many companies throughout the world. Yahoo Finance was used for turnover estimates for large companies.

AIR clients in the London and US markets have provided valuable exposure and claims data that helped AIR better understand insurance policy conditions.

This exposure set represents a comprehensive view of affirmative-only cyber insurance in the United States. Non-affirmative covers are not accounted for in this paper, though the proposed framework could equally be applied to them.

BitSight Discover

Data from BitSight Discover is a key component of the industry exposures used in this study. BitSight pinpoints connections between an organisation, its vendors, and their vendors' service providers, mapping their connections to domains and companies associated with each domain to reveal the level of reliance on a common set of service providers among all insureds within a portfolio.

Data collected by BitSight is sourced from extensive analysis of a company's externally observable webpage content, source IP addresses, public filings, DNS (Domain Name Server) records, and mentions of technologies using NLP (Natural Language Processing) in a company's job postings.

BitSight Discover observes DNS records to identify third and fourth party connections. It is common for organisations to setup CNAME (canonical name) records to point to third party solutions for support, hiring, management of legal documentation, content delivery networks, and more. If an organisation has a CNAME record pointing to .madgexjb.com, for example, it indicates that the company is using Madgex Job Board Technology as a vendor.

As time goes on, more granular data on how the clouds are being used by organisations will become available, for example, whether they are solely hosting their website on the cloud or are running numerous critical business functions on it.

BitSight continuously monitors service provider connections from outside each organisation, identifying new connections daily. BitSight provides visibility into more than 70 different types of service providers, including web hosting, analytics, content delivery networks, mapping providers, domain name servers, payment processors, shopping (storefront providers), SSL certificate authorities, security services, and others.

Using industry averages to derive vendor information

If the industry exposures do not include data for a particular organisation, the model can derive an expected value using industry averages based on market share for different providers in different industry turnover bands. Separate turnover bands are used for each 2-digit NAICS code (North American Industry Classification System). In this way, even if vendors are not known for all organisations, the aggregated data shows the most likely

and most accurate results. Table 1 shows the turnover bands and associated turnover ranges used to determine industry averages.


Table 1: Turnover bands and associated ranges

Turnover band	Range (\$)
A	>= 2 Billion
B	50 Million – 2 Billion
C	10 Million – 50 Million
D	< 10 Million

Exposure statistics for the Fortune 1000 companies and all US industry exposures

The occurrence limits and Business Interruption (BI) sublimits for the US Exposures were estimated using the limit derivation procedure presented in Section 5. The breakdown of limits by industry for both the Fortune 1000 and the US Exposures are shown below in Figure 6.


Figure 6: Cyber business interruption (BI) limits by industry



NAICS classifies certain industry sectors using multiple two-digit codes. For instance, the manufacturing industry is defined by NAICS codes 31, 32 and 33 as seen in Figure 6. The limits shown in this figure are broken down by subsectors when there is enough data available. The significant disparity for industries such as manufacturing and finance is a direct consequence of the difference in

composition by industry between the Fortune 1000 and the full US Exposures. The Fortune 1000 companies account for around 50% of the total turnover in five industries: Mining, Quarrying, Oil and Gas Extraction, Utilities, Transport and Warehousing, Manufacturing, and Information. See Figure 7.

Figure 7: Ratio of turnover, Fortune 1000 to US industry exposures



Scenario classes description



4. Scenario classes description

In this paper, three variants of a cloud outage scenario classes affecting a leading cloud services provider by market share are considered. The variants are:

- 0.5 - 1 day
- 3 - 6 days
- 5.5 - 11 days

The actual duration of losses for companies lasts beyond the initial point of service recovery. Although a 1-day downtime is commensurate with recent experiences, it is AIR's considered professional judgement that any continuous downtime longer than a week is extreme, but not impossible. Cyber aggregation events have been increasing in both frequency and severity.

Clouds can fail or be brought down in many ways. Likely causes of interrupted cloud service include malicious cyber-attacks by external agents, errors by internal workers, as well as hardware and software failures.

Some of these ways, or "vectors", are described in Table 2, mapped to the NIST Taxonomy of Threat Sources. A combination of multiple vectors is also possible. For example, a DDoS attack may be used as a "smokescreen" for an attack involving a zero-day exploit.

A zero-day vulnerability is a security hole in software that is unknown to software creators or antivirus vendors, while a zero-day exploit is the code used by attackers to take advantage of a zero-day vulnerability. "Zero-day" refers to the number of days that a software vendor has known about the vulnerability, meaning no patch is yet available to fix it. Zero-day vulnerabilities and associated exploit codes are extremely valuable, not only to criminal hackers, but to nation-state spies (*Wired, 2014*).

Indeed, it was found that 50% of companies targeted in a DDoS attack were also victims of some form of theft during the event, and that 36% of companies were also infected by malware during the event (*Neustar, 2015*). While not every vector listed could bring down a cloud provider in its entirety, some vectors have led to an all-region downtime event. For example, in February 2013, Microsoft Azure experienced an all-region event when updated HTTPS certificates were not pushed out prior to the expiration of the existing certificates. And in October 2013, Microsoft Azure again experienced an all-region event when an update pushed out to all their data centres exposed an underlying bug.

Table 2: Vectors that could lead to cloud downtime

Threat Source	Vector
Environmental	<p>Lightning strike on data centre</p> <p>Flooding of data centre</p> <p>Solar flare damages electronics</p> <p>Earthquake near data centre</p> <p>Bombing of data centres by terrorists or nation state actors</p> <p>Nearby hazardous materials facility explodes, damaging data centre</p> <p>Accidental cutting of buried power lines, leading to power outages that tax the back-up power systems, which eventually fail</p> <p>Accidental cutting of fibre line of ISP</p> <p>Destruction of data centre via kinetic attack (i.e., crashing a truck or flying a plane into the data centre)</p> <p>Localised or widespread use of EMP</p> <p>Intentional destruction of power grids on which a data centre depends</p> <p>Intentional destruction of nearby dam, flooding a data centre</p> <p>Destruction of network cables by vandals stealing copper, leading to loss of network access</p> <p>Intentional destruction of power grid, leading to widespread power outages that tax the back-up power systems, which eventually fail</p>
Adversarial	<p>Distributed denial of service attack on a cloud service provider</p> <p>Intentional deletion of a large number of virtual machines by a malicious insider</p> <p>Intentional stoppage of a core cloud service by a malicious insider, such as storage</p> <p>Use of zero-day exploit by hackers to compromise the hypervisor</p>
Accidental	<p>Accidental deletion of a large number of virtual machines</p> <p>Accidental stoppage of a core cloud service, such as storage*</p> <p>Accidental simultaneous rebooting of all servers within an availability zone*</p> <p>Use of incorrect configuration settings during routine upgrades leads to loss of availability of front end servers*</p> <p>Insufficient capacity of backup servers during routine maintenance*</p> <p>Errors introduced during routine maintenance leads to a cascading failure and a flood of internal traffic, resulting in a self-caused denial of service type incident (see Box 1)*</p> <p>Human errors introduced during routine maintenance or upgrades interact with the underlying complex system to lead to large-scale downtime events</p> <p>Expiration of HTTPS certificates when renewed certificates were not released as part of routine maintenance*</p> <p>Improperly tested updates expose underlying bugs*</p>
Structural	<p>Failure of environmental management systems*</p> <p>Loss of primary, secondary, and back-up power systems*</p> <p>Short-circuit of power distribution panel</p> <p>Failure of networking devices</p> <p>Failure of file servers</p> <p>Data server with an undetected severe capacity constraint continually crashes and reboots, until automated failure detection systems take the server offline. The loss of server capacity puts increased pressure on remaining data servers until they degrade and fail, leading to cascading failures*</p> <p>Undiagnosed errors masked by automated failure detection systems lead to catastrophic failure of core systems and large-scale downtime events</p>


*Previously encountered vectors that have led to any cloud downtime

Cloud downtime analysis

To understand the likelihood of an extreme cloud downtime event, it is useful to consider both the way a cloud offering is architected and events that have led to cloud outages in the past. Appendix A gives an overview of a number of events that lead to cloud downtimes, and Appendix B discusses cloud architecture and resiliency. Additionally, Gunawi, et al (2016), conducted an in-depth review of past cloud outages. Their cloud outage study looked at events from a subset of cloud providers (e.g. AWS, Microsoft Azure) and cloud based services (e.g. Netflix, Blackberry) from 2009 through 2015.

Taking a further subset of their data to include only cloud providers and events for which the downtime is documented results in 74 distinct downtime events over the seven year period. Four cloud service providers are represented in the dataset: AWS, Google, Microsoft Azure, and Rackspace. The number of downtime events is plotted against the length of the event in Figure 8 below.

Figure 8: Cloud downtime events, 2009-2015



The data set consists of twelve events that lasted for half a day or longer, five events that lasted 1.5 days or longer, and no events that were greater than five days. Of the events in the data set, the longest was a 3.5 day outage affecting AWS in 2011, in which their utility provider suffered a failure of a transformer. Additionally, two all-region events are included in the data set: a nearly half-day downtime event experienced by Microsoft Azure in February 2013, and an 8-hour downtime event experienced by Microsoft Azure in October 2013.


Cloud service providers implement new technologies, controls, and automations in an effort to reduce both the number and severity of any future cloud downtime events, which could potentially limit the utility of historical data as a predictor of future cloud performance. However, industry experts acknowledge that while technology continues to improve, legacy cloud deployment may not keep pace with this change, and that in order to minimize risk in a rapidly evolving technological space continuous improvement of architectural processes is also required. In short, improving technology or increasing the number of data centres cannot prevent cloud downtime events without additional advances in cloud architecture and deployment. This is illustrated in the October 2013 Microsoft Azure downtime event, where an update to a module called Red Dog Front End (RDFE) caused a worldwide outage of the Azure compute service. Because of the way that Azure is architected, only a single version of RDFE can run on the entirety of the Azure cloud, meaning that Azure engineers were unable to deploy the updated RDFE to a subset of Azure data centres to fully test the update. Instead, the update had to be deployed across all Azure regions, at which point an underlying bug was exposed, leading to the outage (*Availability Digest, 2013*).

Recovery schedule

The scenarios assume that the service provider in question goes down in its entirety, i.e., in all regions. The final component of the scenario description is a recovery schedule describing the percentage of the provider's services that are recovered as a function of the elapsed time since the beginning of the outage. Lacking sufficient past examples of cloud downtime (and especially of lengthy downtime) to derive such a curve from data, we have instead produced the recovery curve shown in Figure 9 based on of two assumptions:

1. It will take time to diagnose the cause of downtime, come up with a plan to mitigate it, and begin executing it.
2. Given the distributed nature of the cloud and the (mostly) independent nature of the servers that constitute it, service recovery will progress incrementally.

Figure 9: The percentage of companies with service recovered as a function of time since the beginning of the scenario, for a 3-6 days scenario



Curves for other scenario durations follow the same form, with a linear recovery beginning at half the time it takes for all companies to have recovered service. In particular, the AWS outage of February 2017 satisfied both of the above assumptions, with recovery constituting a multi-step process in which services were restored at distinct times.

Likely causes of interrupted cloud service include malicious cyber-attacks by external agents, errors by internal workers, as well as hardware and software failures.

Box 1: April 2011 Amazon Web Services outage

On April 21st 2011, a widely used storage service within AWS (the Elastic Block Store, or EBS) went down, leading to widespread service disruptions in the Amazon US East Region. The outage affected many popular websites, such as Reddit, Quora, and Foursquare (*Bright, 2011*).


The outage began during a routine configuration change that AWS was making to an Availability Zone within the US East Region as part of an upgrade to the primary network capacity. At 12:47 AM PDT, the EBS cluster in the Availability Zone was moved from the primary router (*Amazon, 2011*). While the cluster was supposed to be moved to a redundant router, it was instead moved to a secondary lower-capacity router. This lower-capacity router could not handle the volume of traffic between the EBS nodes. With the primary network offline and the secondary network overburdened, the result was that each EBS node was effectively isolated from all other EBS nodes (*Bright, 2011*).

EBS nodes in the AWS cloud all have partner nodes, with each node pair storing exact replicas of their data. Should an EBS node lose communication with its partner node, the EBS node immediately begins searching for a new node to on which to replicate its data, a process managed by the EBS control plane. In normal operations, this node behaviour is beneficial, as it ensures that data is not lost. However, during the outage this behaviour led to further issues within the Availability Zone. Once AWS brought the primary network back online, all the EBS nodes immediately began replicating their data onto new partner nodes. This replication behaviour meant that all available storage capacity was quickly filled, leaving many nodes “stuck,” searching for free storage space that was not available (*Amazon, 2011*).

Up to this point, the failures were restricted to a single Availability Zone within the US East Region. However, the “stuck” EBS nodes continued to make requests for partner nodes to the EBS control plane, which operates across all Availability Zones within a region. The control plane, which in addition to handling EBS partner node requests, also processes requests for new volumes, became backed up with requests that it could not fulfil. By 5:30 AM PDT, the control plane began failing all requests, not just within the originally affected Availability Zone, but in all zones within the region. Eventually, Amazon engineers had to disable communications between the affected EBS clusters and the EBS control plane to minimize impact on the other Availability Zones within the region. By 12:04 PM PDT, the outage was confined to only the originally affected Availability Zone (*Amazon, 2011*).

To address the original problem of “stuck” EBS nodes, Amazon had to install additional capacity in the affected Availability Zone, which meant physically moving and then installing additional servers in the Availability Zone. Amazon then gradually restored communications between the affected EBS nodes and the EBS control plane. The process of restoring communications took the entire day of April 23rd. By then, the majority of affected EBS nodes were functional, a minority of which required a manual recovery process. The manual recovery process was completed by 3:00 PM PDT on April 24th, at which point all but 0.07% of affected nodes were fully restored (*Amazon, 2011*).

Analysis



5. Analysis

This section presents a detailed analysis of both ground-up and gross losses for the modelled scenarios, specifically, three combinations of ground up, gross insurable, and gross insured losses using the detailed accumulation approach described in the preceding section. Because of the very low take-up rates of cyber policies outside the US, this analysis focuses on the nearly 12.4 million companies in the US. The same analysis is also applied to the Fortune 1000 companies. The results will also be split by a number of variables, allowing for a comparison of losses by:

- Industry
- Company size (Fortune 1000 vs. industrywide)

Additionally, this report provides a comparison of loss numbers using the detailed accumulation approach to those loss numbers that would be obtained using a market share approach.

Ground-up loss computation

The preceding section fully describes the downtime of the modelled cloud service provider. However, additional parameters are required to model the impact of the above scenario on U.S. businesses:

- Set of affected companies
- Business interruption and contingent business interruption losses due to inability to access the cloud
- Backup plans and recovery process

This section describes and provides values for these parameters, and briefly discusses how they were obtained.

Set of affected companies

The first step is to determine how many companies are affected by the cloud's downtime.


This can be done through either a *market share* approach, or a *detailed accumulation* approach. In the market share approach, the set of impacted companies is determined by randomly sampling a population of companies. A particular company is deemed to be impacted by a cloud's downtime with probability p , where p is equal to the market share of the cloud service provider. To obtain a more accurate estimate of losses, the population may be sampled multiple times, with losses averaged over the various samples to create a sensible estimate, or to infer the degree of uncertainty in loss estimates.

In the detailed accumulation approach, the set of impacted companies is known precisely using data from an external source or the insured themselves. Because current market practice for insurers is generally not to ask potential insureds for a detailed list of service providers, many insurers do not have such data available. Additionally, small companies may not be tracked by external sources. Therefore, the market share approach is also used to supplement the detailed accumulation. The analysis in this paper compares both approaches. In the future, systematic collection of this type of data will enable a better understanding of the technology at risk from cyber-attacks.

Business interruption losses from lost cloud access

Once the set of affected companies has been determined, it is necessary to determine what business interruption losses are incurred as a consequence of a failure to access the cloud. Clearly, different companies use the cloud for different purposes, and a single company may use the cloud for different purposes that are not equally important to its operations. In this paper, business interruption losses are computed using an industry dependent set of factors equal to the percentage of a company's turnover from e-business. These factors are found in Figure 10.

Figure 10: E-Business factors by industry



Source: US Census Bureau, 2015

Business interruption estimation factors

The numeric factors used to calculate daily business interruption losses are computed as percentages of turnover that would be affected by downtime for a given industry. Figure 10 considers both business to consumer (B2C) and business to business (B2B) e-business transactions. Business to consumer e-business is primarily found in the retail sector, while business to business e-business is found in manufacturing and wholesale industries. In fact, the data show that B2B e-business far exceeds B2C commerce, with B2B e-business projected to achieve \$12 trillion in sales worldwide by 2020, for a compound annual growth rate of 8.11% (Frost & Sullivan, 2017). 2014 U.S. Census data shows that as a percentage of total shipments, sales, and turnover, manufacturing e-business leads the way at 60.9%, followed by merchant wholesale trade at 27.7%, and trailed by total retail trade at 6.4%. The US Census Bureau indicates that “E-commerce sales/turnovers are sales of goods and services where the buyer places an order, or the price and terms of the sale are negotiated over the Internet, mobile device (m-commerce), Extranet, Electronic Data Interchange (EDI) network, electronic mail, or other comparable online system.

E-commerce shipments (e-shippments) are online orders accepted for manufactured products from customers, including shipments to other domestic plants of the same company for further manufacture, assembly, or fabrication where price and terms of sale are negotiated over the Internet, Extranet, EDI network, electronic mail, or other online system. Payment may or may not be made online” (US Census Bureau, 2016). The data in this study includes the components mentioned above. The term e-business, as used in this paper, is defined in more detail in Appendix C.

The business interruption estimation factors in this paper are calculated by using the US e-commerce data source (US census, e-commerce) that includes shipments, sales and revenues from various sectors of the economy: manufacturing, wholesale, services and retail. The services sector includes industries like utilities, finance and insurance, educational services, information, healthcare, etc.

AIR has calculated the business interruption estimation factors for the other industries including agriculture, mining and quarrying, by assuming the e-business factor to be the minimum of the set of all available e-business factors.

Deriving business interruption cost per day

The model uses AIR-computed factors to derive a cost. These factors are based on industry, turnover, e-business percentage of the industry type, and the gross profit ratio of the industry type.

Expenses will also diminish during a business interruption. Specifically, all fixed expenses not tied to the actual production of goods/services (e.g., wages, rent, etc.) continue to be incurred, whereas the total cost of producing goods/services is assumed to be reduced by the same factor as the quantity of goods/services produced. The gross profit ratio, equal to one minus the cost of turnover ratio, will be multiplied by each company's turnover to determine the quantity of potential BI and CBI losses. The gross profit ratio was obtained from public data (*Butler Consultants, 2017*) and the various industries in the data were mapped to NAICS codes.

For manufacturing firms specifically, a "waiting period" of 12 hours is included during which time there is no loss. This is because these firms traditionally stockpile inventory and therefore might not experience an instant impact on operations. Note that this 12 hour waiting period is applied to the ground up losses, indicating that no loss to the business is incurred during that time. *Any waiting period included in financial terms is additional, since losses for manufacturers are assumed to begin accruing only after the 12 hour waiting period.*

Backup plans and recovery speed

The final component in determining ground-up losses is a model of company behaviour during and after service downtime. During downtime, companies may implement some form of back-up plan to mitigate (but not entirely eliminate) losses. This back-up may be to a different cloud provider, or it may be a different method of back-up, such as switching to paper. Table 3 and Table 4, describe the prevalence of and time to implementation of back-up plans for low and high turnover companies. Table 5 describes the percentage reduction in losses for companies that implement such a back-up plan. Implementation of a backup plan may involve switching to a method that increases expenses, but any such increase in expenses can be "absorbed" into the aforementioned loss reduction factor. After downtime ends, companies do not recover instantaneously; rather they do so progressively. This is modelled through a table describing how quickly a company moves from maximal business interruption (or mitigated business interruption, if it implemented a back-up plan) to nominal operation. These tables are presented in "Company-specific parameters and data sources" on pg.30.

Gross loss computation

To compute gross losses from ground up losses, the model needs information about take-up rates (when calculating insured gross losses) and insurance terms. Take-up rates are determined as a function of industry and turnover. Because this is a (contingent) business interruption scenario, the relevant insurance terms include waiting periods and limits.

Company-specific parameters and data sources

This section begins by detailing the company-specific parameters described above, complete with a description of how the relevant parameters were obtained from data.

Cloud usage

Where provider data was available, precise information about provider use was used to determine losses. Where this information was not available, a market share approach was applied. The data for this approach was constructed from those companies where there was data. For each combination of industry, provider, and 4 revenue bins, the market share of that provider was determined. Additionally, a table summarising the average number of cloud providers used by companies, split by industry and revenue bin was developed. These values were used to determine the fraction of losses that a company experiences when a single provider of many goes down. See “Distribution of losses” on pg.33 for more information.

Company back-up plans

Table 3 and Table 4 describe the distribution of time that companies take to successfully implement back-up plans, for companies whose annual turnovers are under \$1 billion and over \$1 billion, respectively. The distribution adds up to 0.8 for large companies, but only to 0.5 for small companies. This is because large companies generally have sufficient infrastructure to avoid dependence on the cloud. Smaller companies might be more likely to use the cloud to avoid the expense of building the business infrastructure in-house. This is especially true for those companies using IaaS, in which case there may be no backup plan possible. We also note that back-up plan implementation may fail, a factor that is incorporated in Tables 3 and 4 as one reason companies may fail to implement a back-up plan within 7 days.

The percentage reduction in losses from the implementation of a back-up plan is also modelled probabilistically according to the distribution in Table 5. Thus, it is assumed that (C)BI losses are reduced by an average of 40% as soon as a back-up plan is implemented. That is, the model presented here assumes a discrete transition from peak (C)BI losses to 50%, 60%, or 70% of those losses at the point in time when the back-up plan is implemented. These tables were derived from sources like the 2017 Hiscox Cyber Readiness report, Disaster Recovery Preparedness Benchmark Survey (*Hiscox, 2017*).

Table 3: Probability of implementing a back-up plan as a function of the number of days since the start of the outage, for companies with turnovers of under \$1 billion

Day	1	2	3	4	5	6	7	None/fail
Probability	0.27	0.08	0.05	0.04	0.03	0.02	0.01	0.50

Table 4: Probability of implementing a back-up plan as a function of the number of days since the start of the outage, for companies with turnovers of over \$1 billion

Day	1	2	3	4	5	6	7	None/fail
Probability	0.44	0.13	0.08	0.06	0.04	0.03	0.02	0.20

Table 5: Assumed probability distribution governing reduction of (C)BI losses upon successful implementation of a back-up plan

Loss Reduction	30%	40%	50%
Probability	0.3	0.4	0.3

Company recovery

Table 6 describes the distribution of time that companies take to recover from the level of (C)BI losses experienced at the time of service restoration to nominal productivity. This table was derived from sources such as the 2017 Hiscox Cyber Readiness report, Disaster Recovery Preparedness Benchmark Survey (*Hiscox, 2017*), as well as through consultations with insurers.

Contrary to the model of back-up plans, AIR assumes a progressive recovery of turnover rather than a discrete change. Specifically, losses will diminish linearly over the recovery time. This is not a restrictive assumption, as the same losses would be obtained by a discrete end of losses occurring at half the company recovery duration.


Table 6: Probability distribution governing the time to reach zero (C)BI losses once service to the cloud has been recovered

Day	1	2	3	4	5	6	7
Probability	0.39	0.19	0.13	0.1	0.08	0.06	0.05

Box 2: Recovery example

The following simple example and sample calculation illustrate how service recovery, implementation of back-up plans, and company recovery determine ground-up losses. Suppose that the downtime period under consideration has a duration of 3 - 6 days. Then companies will recover their service at some point from 3 to 6 days, with all times in this range being equiprobable. Suppose then that company X regains its service after 5 days. Suppose further that it implements a back-up plan after 2 days, that this reduced their losses by 40%, and that their recovery time was only 1 day. Finally, assume that their (C)BI losses are \$1/day at the beginning of the scenario. As stated in the "Company back-up plans" section on pg.30, the implementation of a back-up plan reduces these losses instantaneously, bringing them to \$0.60/day. As described in the "Company recovery" section on pg.31, the (C)BI loss then diminishes linearly from \$0.60/day to \$0/day over the recovery period of one day. This yields the plot of Figure 11 for (C)BI losses as a function of time. The total ground-up loss is \$4.10, obtained by integrating the area under the curve.

Figure 11: (C)BI cost experienced as a function of time



Take-up rates

Take-up rates of cyber insurance by companies are presented in Table 7, as a function of industry and turnover. These rates were obtained by combining published data from Advisen (*Advisen, 2014*) and Marsh (*Marsh, 2017*). Marsh provides data on take-up rates by industry and Advisen provides take-up rates by turnover

bands. The median of the original Marsh take-up rates across all the industries was set as a baseline and the variation around this baseline was determined for every take-up rate. These values were then rescaled using the original Advisen take-up rates by turnover for each turnover bin to get the combined data.

Table 7: Cyber insurance take-up rates by industry and turnover

	Original Take-Up Rates (Marsh)	< \$2.5M	\$2.5M to \$5M	\$5M to \$15M	\$15M to \$100M	\$100M to \$300M	\$300M to \$1B	\$1B to \$5B	> \$5B
Original Take-Up Rates (Advisen)		5.20%	6.50%	9.10%	13.00%	23.40%	27.30%	28.60%	33.80%
Manufacturing	14.00%	1.90%	3.20%	5.80%	6.00%	16.40%	20.30%	21.60%	26.80%
Financial Institutions	26.00%	4.30%	5.60%	8.20%	9.40%	19.80%	23.70%	25.00%	30.20%
Retail/Wholesale	28.00%	4.70%	6.00%	8.60%	11.00%	21.40%	25.30%	26.60%	31.80%
Services	30.00%	5.10%	6.40%	9.00%	12.60%	23.00%	26.90%	28.20%	33.40%
Power and Utilities	31.00%	5.30%	6.60%	9.20%	13.40%	23.80%	27.70%	29.00%	34.20%
Hospitality and Gaming	36.00%	6.30%	7.60%	10.20%	17.40%	27.80%	31.70%	33.00%	38.20%
Education	47.00%	8.50%	9.80%	12.40%	26.20%	36.60%	40.50%	41.80%	47.00%
Communications, Media and Technology	50.00%	9.10%	10.40%	13.00%	28.60%	39.00%	42.90%	44.20%	49.40%
Health Care	57.00%	10.50%	11.80%	14.40%	34.20%	44.60%	48.50%	49.80%	55.00%
All Other Industries	26.00%	4.30%	5.60%	8.20%	9.40%	19.80%	23.70%	25.00%	30.20%

Total occurrence limits

The total limits for the organizations are derived based on industry and turnover, and rounded to the nearest predefined limit represented in Table 8. For organisations with higher turnovers, there is no variation in the limits across different industries as seen in various client books. These models were constructed by aggregating and analysing the cyber exposures of over 20 insurers.

Table 8: Predefined limits

Limits
1. 25,000
2. 50,000
3. 100,000
4. 250,000
5. 500,000
6. 750,000
7. 1,000,000
8. 2,500,000
9. 5,000,000
10. 7,500,000
11. 10,000,000
12. 15,000,000
13. 25,000,000
14. 50,000,000
15. 75,000,000
16. 100,000,000
17. 200,000,000 – Select large companies
18. 500,000,000 – Select large companies

Business interruption sublimits

For limits under \$1 million, no separate business interruption sublimit was assumed. For limits of \$1 million or greater, a business interruption sublimit of 56% of the overall cyber limit was assumed. This percentage was determined using actual exposure data from several cyber insurers.

Waiting periods

The paper presents results for 8, 12, and 24-hour waiting periods.

Scenario losses

Prior to presenting estimates of loss, we detail all additional assumptions made and then present the method by which the results were computed.

Calculating business interruption losses: procedure and assumptions

In the business interruption scenarios, loss is calculated with the following procedure, in which the key bit of exposure data is business interruption cost per day (cost per day of no access to service). The model derives it using industry and turnover data in conjunction with AIR-developed factors. For more information, see the “Deriving business interruption cost per day” section on pg.29.

1. Estimate ground up losses as in the recovery example of Box 2 (pg.26)
2. Apply policy terms including waiting periods and limits.

Factors based on e-business percentage

The model calculates the business interruption cost per day with the organisation's daily turnover and a business interruption estimation factor derived from industry-wide e-business data based on the two- or three-digit NAICS code. AIR treats the business interruption estimation factor as an indicator of the extent to which the company uses the internet. (Note that the model uses a 3-digit NAICS code when available and the 2-digit code when it is not.) Using this factor, the model calculates the business interruption cost with this equation:

$$\begin{aligned} &\text{Business interruption cost per day} \\ &= [\text{Daily company turnover}] \\ &\times [\text{Business interruption estimation factor}] \\ &\times [\text{Gross profit ratio}] \end{aligned}$$

Distribution of losses

There are several ways to distribute losses, depending on the number of organisation domains affected by the disabled provider, the numbers of providers the domains use, and the type of provider that goes down.

In the relatively straightforward cases, losses are either incurred entirely by a single domain or are divided equally among domains. If a domain uses, for example multiple content delivery networks or domain name services, then no losses are incurred because the domain's traffic can route through an alternate network in a backup scenario.

Although many organisations have single or multiple domains that can be served by single or multiple providers, cloud scenarios are not necessarily backup scenarios. The model does not assume zero losses in the cloud provider simply because a backup is theoretically possible. Since there is no data available on the importance of any one cloud to a particular company, we make the assumption that, given an organisation which uses X clouds, the failure of one cloud results in greater than $1/X$ of the loss.

This paper does not examine multiple simultaneous cloud failures but it is instructive to consider such a scenario to understand why the factor is greater than $1/X$.

Specifically, consider as an example the problem of estimating the loss of productivity that occurs from a lack of access to the following two resources: Microsoft Office and intranet connectivity. If intranet connectivity is lost, a person loses the ability to access documents located on other computers, but can still access documents on one's own computer. If Microsoft Office fails, then one can no longer edit Office documents, irrespective of where they are located. In particular, any task consisting of editing an Office document that is located on another computer can only be done if both Office and the intranet are working. Thus, *the business interruption loss from the failure of both resources is smaller than the sum of the losses*, since the sum "double counts" some tasks.

In general, the sum of losses from individual components (e.g., clouds) failing will always be greater than the loss that occurs if they all fail simultaneously. It follows that the average loss from a single failure will be greater than that provider's fractional contribution. In particular, the *incremental* losses from each additional failure will *decrease* with the number of failures. This means a curve describing losses as a function of the number of failures should have a *concave* shape.

Loss computation

A number of scenario parameters define probabilities rather than fixed values. To deal with these parameters, losses are computed by sampling and averaging. In particular, this approach is used to determine: the time at which companies recover service, the time to implement a back-up plan, the reduction in losses from doing so, and the time to return to normal operation after service recovery. A total of 5,000 samples were taken overall.

Results

Tables 9, 10, 11, and 12 show results for the three combinations of ground up, gross insurable, and gross insured losses for 8, 12 and 24 hour waiting periods for a sampling of the top 15 providers. Specifically, we provide losses for providers 1, 3, 10, and 15 in the US market, without specifying what the providers are.

The tables provide both central estimates and 95% confidence interval values to give a sense of the uncertainty in these estimates. Tables 12 and beyond assume 8 hour waiting periods. Table 12 shows analogous results to Table 9, but limited to the companies making up the Fortune 1000.

Tables 13 and 14 provide central estimates of losses for a 3-6 day cloud downtime, broken down by industry. Manufacturing and wholesale and retail Trade are the industries that would be mostly affected by the failure of a cloud provider in the US.

Finally, Table 15 provides a comparison of gross insured losses for the detailed accumulation and market share approaches.

Table 9: Central estimates of loss (with 95% Confidence interval) for different ranges of downtime and various cloud providers (8 hour waiting period)

		0.5-1 Day			3-6 Days			5.5-11 Days		
Size of the event		Large			Extreme			Very extreme		
Provider	Ground Up Loss (Billions)	Gross Insurable Loss (Billions)	Gross Insured Loss (Billions)	Ground Up Loss (Billions)	Gross Insurable Loss (Billions)	Gross Insured Loss (Billions)	Ground Up Loss (Billions)	Gross Insurable Loss (Billions)	Gross Insured Loss (Billions)	
Provider 1	5.89	4.83	1.08	14.74	12.70	2.75	23.80	19.91	4.23	
CI 95%	(3.73-10.48)	(2.79-8.97)	(0.62-2.00)	(11.33-19.02)	(9.73-16.17)	(2.13-3.46)	(18.56-29.65)	(15.93-24.18)	(3.43-5.07)	
Provider 3	2.80	2.25	0.58	6.92	5.78	1.45	11.16	8.93	2.19	
CI 95%	(1.79-4.81)	(1.34-4.18)	(0.35-1.06)	(5.33-8.45)	(4.50-7.04)	(1.13-1.75)	(8.60-13.28)	(6.99-10.53)	(1.74-2.58)	
Provider 10	0.85	0.72	0.17	2.12	1.96	0.45	3.43	3.20	0.73	
CI 95%	(0.54-1.55)	(0.41-1.43)	(0.10-0.33)	(1.63-2.62)	(1.47-2.44)	(0.34-0.56)	(2.61-4.18)	(2.45-3.88)	(0.56-0.89)	
Provider 15	0.43	0.36	0.08	1.07	0.99	0.22	1.73	1.64	0.36	
CI 95%	(0.26-0.78)	(0.20-0.71)	(0.04-0.16)	(0.82-1.34)	(0.74-1.25)	(0.16-0.28)	(1.31-2.10)	(1.23-2.00)	(0.27-0.44)	

Table 10: Central estimates of loss (with 95% Confidence interval) for different ranges of downtime and various cloud providers (12 hour waiting period)

Provider	0.5-1 Day			3-6 Days			5.5-11 Days		
	Ground Up Loss (Billions)	Gross Insurable Loss (Billions)	Gross Insured Loss (Billions)	Ground Up Loss (Billions)	Gross Insurable Loss (Billions)	Gross Insured Loss (Billions)	Ground Up Loss (Billions)	Gross Insurable Loss (Billions)	Gross Insured Loss (Billions)
Provider 1	5.89	4.37	0.98	14.74	12.23	2.65	23.80	19.49	4.14
CI 95%	(3.73-10.48)	(2.34-8.51)	(0.52-1.90)	(11.33-19.02)	(9.24-15.72)	(2.02-3.37)	(18.56-29.65)	(15.50-23.79)	(3.34-4.99)
Provider 3	2.80	2.04	0.53	6.92	5.57	1.39	11.16	8.74	2.15
CI 95%	(1.79-4.81)	(1.11-4.00)	(0.29-1.02)	(5.33-8.45)	(4.28-6.85)	(1.08-1.70)	(8.60-13.28)	(6.79-10.36)	(1.69-2.54)
Provider 10	0.85	0.65	0.15	2.12	1.88	0.43	3.43	3.13	0.71
CI 95%	(0.54-1.55)	(0.34-1.36)	(0.08-0.32)	(1.63-2.62)	(1.30-2.37)	(0.32-0.54)	(2.61-4.18)	(2.37-3.81)	(0.55-0.87)
Provider 15	0.42	0.32	0.07	1.07	0.95	0.21	1.73	1.60	0.35
CI 95%	(0.26-0.78)	(0.16-0.68)	(0.04-0.15)	(0.82-1.34)	(0.70-1.22)	(0.15-0.27)	(1.31-2.10)	(1.19-1.96)	(0.26-0.43)

Table 11: Central estimates of loss (with 95% Confidence interval) for different ranges of downtime and various cloud providers (24 hour waiting period)

	0.5-1 Day			3-6 Days			5.5-11 Days		
Size of the event	Large			Extreme			Very extreme		
Provider #	Ground Up Loss (Billions)	Gross Insurable Loss (Billions)	Gross Insured Loss (Billions)	Ground Up Loss (Billions)	Gross Insurable Loss (Billions)	Gross Insured Loss (Billions)	Ground Up Loss (Billions)	Gross Insurable Loss (Billions)	Gross Insured Loss (Billions)
Provider 1	5.89	3.02	0.68	14.74	10.82	2.34	23.80	18.23	3.88
CI 95%	(3.73-10.48)	(1.10-7.09)	(0.25-1.59)	(11.33-19.02)	(7.75-14.36)	(1.70-3.08)	(18.56-29.65)	(14.17-22.58)	(3.06-4.74)
Provider 3	2.80	1.41	0.37	6.92	4.92	1.23	11.16	8.19	2.01
CI 95%	(1.79-4.81)	(0.51-3.45)	(0.13-0.85)	(5.33-8.45)	(3.60-6.26)	(0.91-1.56)	(8.60-13.28)	(6.17-9.83)	(1.54-2.41)
Provider 10	0.85	0.45	0.10	2.12	1.64	0.38	3.43	2.90	0.66
CI 95%	(0.54-1.55)	(0.16-1.16)	(0.04-0.27)	(1.63-2.62)	(1.15-2.14)	(0.26-0.49)	(2.61-4.18)	(2.14-3.60)	(0.49-0.82)
Provider 15	0.43	0.22	0.05	1.07	0.83	0.18	1.73	1.48	0.33
CI 95%	(0.26-0.78)	(0.07-0.58)	(0.015-0.13)	(0.82-1.34)	(0.58-1.10)	(0.13-0.24)	(1.31-2.10)	(1.07-1.85)	(0.24-0.41)

For the rest of this paper, the loss results will be for 8 hour waiting periods.

Table 12: Central estimates of loss (with 95% Confidence interval) for the Fortune 1000 for different ranges of downtime and various cloud providers (8 hour waiting period)

		0.5-1 Day			3-6 Days			5.5-11 Days		
Size of the event		Large			Extreme			Very extreme		
Provider #	Ground Up Loss (Billions)	Gross Insurable Loss (Billions)	Gross Insured Loss (Billions)	Ground Up Loss (Billions)	Gross Insurable Loss (Billions)	Gross Insured Loss (Billions)	Ground Up Loss (Billions)	Gross Insurable Loss (Billions)	Gross Insured Loss (Billions)	
Provider 1	2.20	1.73	0.50	5.43	4.12	1.18	8.77	5.86	1.68	
CI 95%	(1.37-3.88)	(0.97-3.16)	(0.28-0.90)	(4.10-6.78)	(3.37-4.92)	(0.97-1.40)	(6.67-10.70)	(4.77-6.75)	(1.37-1.94)	
Provider 3	1.07	0.87	0.25	2.62	2.26	0.64	4.23	3.54	1.01	
CI 95%	(0.65-1.93)	(0.48-1.65)	(0.14-0.47)	(1.99-3.22)	(1.71-2.80)	(0.49-0.80)	(3.24-5.23)	(2.78-4.23)	(0.79-1.20)	
Provider 10	0.36	0.30	0.08	0.88	0.80	0.23	1.42	1.29	0.37	
CI 95%	(0.21-0.67)	(0.15-0.61)	(0.04-0.17)	(0.65-1.10)	(0.58-0.98)	(0.17-0.28)	(1.06-1.78)	(0.97-1.59)	(0.28-0.45)	
Provider 15	0.17	0.14	0.044	0.42	0.38	0.11	0.67	0.64	0.18	
CI 95%	(0.09-0.31)	(0.07-0.28)	(0.02-0.08)	(0.31-0.52)	(0.27-0.49)	(0.08-0.13)	(0.50-0.83)	(0.47-0.79)	(0.13-0.22)	

As discussed in Section 5, losses for each company were determined by a combination of deterministic (e.g., revenue, NAICS code, etc.) and probabilistic factors (e.g., time at which service is recovered, time at which backup is implemented, etc.). There were four distinct probabilistic factors and these four numbers were generated for each company. However, the values corresponding to the time of service recovery, the percent loss reduction from implementing a backup plan, and the time to recover fully to normal after service has been recovered were all correlated between companies. For the time of service recovery, this correlation models the fact that, while distinct companies may recover service at different times, a realistic scenario is likely to have only a few distinct time points at which a large number of companies simultaneously recover service. For the percent loss reduction from implementing a backup plan and the time to recover fully to normal after service has been recovered, the correlation models the fact that these values are not only a function of the company itself, but also of the event to which they are responding. Thus, these two correlations can be

understood to model uncertainty in the parameters used. On the other hand, the time to implement a backup plan was deemed to be dependent solely on the company itself, and hence was not correlated from one company to the next. Thus, the losses are fully defined by four length N vectors (where N is the number of companies). These length N vectors were generated by using Algorithm 2 of (*Dukic and Marić, 2013*) with a correlation of 0.49 for the three vectors whose components were correlated.

This entire procedure was performed repeatedly to obtain the 95% confidence intervals, which are shown in Tables 9, 10, 11 and 12. Along with the aforementioned parameter uncertainty, the repeated sampling also models sampling error.

Tables 13 and 14 provide ground up and gross insured loss estimates for a 3-6 day cloud downtime, split by industry. Manufacturing and Wholesale and Retail Trade are the industries that would be most affected by the failure of a cloud provider in the US.

Table 13: Central estimates of ground up losses by industry for 3-6 day downtime and various cloud providers (8 hour waiting period) (\$ millions)

Industry	Provider 1	Provider 3	Provider 10	Provider 15
Accommodation and Food Services	251.57	111.03	40.82	16.82
Administrative and Support and Waste Management and Remediation Services	70.40	18.00	11.41	4.81
Arts, Entertainment, and Recreation	55.77	18.68	7.07	3.07
Educational Services	44.10	10.32	4.07	0.68
Finance and Insurance	447.31	182.08	77.29	20.96
Information	846.68	351.12	117.58	33.22
Manufacturing	8,555.62	4,179.71	1,187.52	697.98
Other Services (except Public Administration)	87.19	34.44	10.80	3.88
Professional, Scientific, and Technical Services	136.02	51.40	17.53	7.37
Real Estate and Rental and Leasing	212.98	133.39	33.34	10.48
Transportation and Warehousing	438.97	161.18	67.33	22.33
Wholesale and Retail Trade	3,558.02	1,371.61	582.64	247.17
Others (Agriculture, Construction, Management of Companies and Enterprises, Healthcare, Oil and Gas, Public Administration, Utilities, Unclassified)	31.51	9.92	4.97	1.21

Table 14: Central estimates of gross insured losses by industry for 3-6 day downtime and various cloud providers (8 hour waiting period) (\$ millions)

Industry	Provider 1	Provider 3	Provider 10	Provider 15
Accommodation and Food Services	45.91	20.67	7.40	2.79
Administrative and Support and Waste Management and Remediation Services	11.47	3.09	2.12	0.98
Arts, Entertainment, and Recreation	9.36	2.91	1.10	0.52
Educational Services	9.45	2.74	1.05	0.06
Finance and Insurance	106.73	44.61	19.13	4.90
Information	307.25	130.31	46.55	12.11
Manufacturing	1,477.13	765.34	245.23	144.88
Other Services (except Public Administration)	6.43	3.96	0.89	0.32
Professional, Scientific, and Technical Services	25.29	10.15	3.86	1.35
Real Estate and Rental and Leasing	28.50	17.99	6.05	1.48
Transportation and Warehousing	96.20	38.19	15.93	5.06
Wholesale and Retail Trade	623.98	292.03	112.02	42.79
Others (Agriculture, Construction, Management of Companies and Enterprises, Healthcare, Oil and Gas, Public Administration, Utilities, Unclassified)	4.74	1.81	0.88	0.17

Table 15 below compares gross insured losses for the detailed accumulation and market share approaches. The data in this table was obtained by computing the percentage difference between the two approaches for a number of providers and then summarizing this information by providing the mean of the *absolute value* of the percentage differences. AIR also repeated the same analysis for a hypothetical book of business, which illustrates that larger differences between the two approaches can be expected for an insured's book of business than for the industry as a whole.

Comparing the percentage differences between the detailed accumulation and market share approaches for the various considered providers showed a mix of all the possible cases: the case where the detailed accumulation approach yielded higher losses than the market share approach for both the industry as a whole and the sample book, the case where the market share approach yielded higher losses for both, and the case where one of the approaches yielded higher losses for the industry as a whole but lower losses for the sample book.

Table 15: Comparison of detailed accumulation and market share gross insured losses for different ranges of downtime


0.5-1 Day Gross Insured Loss (billions)		3-6 Days Gross Insured Loss (billions)		5.5-11 Days Gross Insured Loss (billions)	
Overall Mean Absolute Diff (%)	Portfolio Mean Absolute Diff (%)	Overall Mean Absolute Diff (%)	Portfolio Mean Absolute Diff (%)	Overall Mean Absolute Diff (%)	Portfolio Mean Absolute Diff (%)
7.46	24.06	9.34	23.46	10.16	23.20

As an example, gross insured losses for the hypothetical portfolio in the 0.5-1 day scenario for a particular provider were determined to be in the range of \$104 - \$313 million (mean: \$177 million) for the market share approach versus \$74 - \$220 million (mean: \$124 million) for the detailed accumulation approach, a 42% difference, even for a scenario where total US losses for the provider differ by only 9.3%. For the hypothetical cyber portfolio and a 3-6 day downtime event to the same provider, ground up losses could be in the range of \$643 million - \$1.06 billion (mean: \$847 million) for the detailed accumulation and \$868 million - \$1.43 billion (mean: \$1.12 billion) for the market share approaches, respectively, a 32% difference. For other specific insurance portfolios or sectors, the difference could be even larger. The gross insurable losses can be considered a substantial portion of the ground up loss to

United States companies due to the downtime event. Note that the loss does not include financial impacts on the cloud providers themselves, which one source (*Lynley, 2013*) lists as \$1,104 per second, in the case of AWS for example. Nor does the loss include third party liability or any data breach losses.

The hypothetical portfolio used in this analysis is one of many obtained from sample client books or extracted from AIR's full industry exposure database. The data in Table 15 was constructed from the hypothetical portfolio that showed the largest difference between detailed accumulation and market share losses among those that AIR ran, though others might have bigger or smaller differences. AIR thinks this is representative of a realistic difference for a syndicate's portfolio compared to a market share.

Implication of cloud failure on (re)insurance



6. Implication of cloud failure on (re)insurance

Insurance

Cyber insurance can cover many aspects of a loss from a cyber incident, including breach response costs such as forensics, legal fees, notification and credit monitoring, and third-party liability. However, these are more applicable to security breach events, whereas the scenarios contemplated here are more relevant to business interruption coverage discussed below. This report can help insurers and reinsurers understand their potential exposure to business interruption aggregation events. In addition, cyber incidents can cause loss to “non-affirmative” coverage, such as directors and officers, errors and omissions, and even property policies.

Business interruption and contingent business interruption

The first known business interruption (BI) cover was sold in the London Market in 1868 (*LMI Group, 2017*). One hundred fifty years later, BI insurance is a common part of a company’s insurance protection, providing cover for lost or interrupted business income sustained from a number of perils.

However, in the present day, businesses are increasingly reliant on cyber technology to control and optimise production, digitally store and access data, and transact and market their operations. This leaves businesses at risk from cyber attacks on these digital systems that, if successful, could cause significant business interruption and turnover loss.

Furthermore, in a digitally connected world, companies are also increasingly reliant on data from other businesses to conduct their operations. This could also result in significant business interruption if a third-party business providing critical products or services is attacked. This is known as contingent business interruption (CBI).

As such, this supply chain of digital interdependencies is now widely recognised as a significant source of risk aggregation by insurers. If a cyber attack occurs on a critical node of the cyber supply chain, such as a major cloud vendor, the attack could cause systemic business interruption to all associated businesses that rely on the vendor’s services and systems to operate.

Liability

Cloud use will increase as businesses continue to realise the cost and scalability benefits of using cloud providers to store, manage, and process data. However, as a cloud vendor and a customer enter into business with each other, the issue of liability often becomes contentious. Customers want cloud providers to assume unlimited liability for outages and any resultant business interruption, while vendors want to restrict and cap their liability.

For example, Salesforce’s Master Services Agreement describes the company’s commitment to provide services 24 hours a day, except for planned downtime and a number of specific circumstances out of the company’s control, such as service outage from events such as denial-of-service attacks. Another practice that is becoming increasingly common in service agreement contracts is a \$0 valuation of the data being held or processed, further limiting liability for the cloud provider (*Gilbert, 2011*).

The mode of engagement that a customer may have with a cloud service provider also adds complexity regarding liability. For example, if a service provider such as YouTube experiences downtime, many digital businesses that rely on YouTube videos as a turnover stream via direct advertising income or via indirect turnover through social media marketing would be likely to incur business interruption losses. However, as there is no fee paid to YouTube to use the service, users do not receive any compensation for the loss of availability. The business rationale is that if the service is provided for no fee, there

is no financial loss for the user, regardless of any business interruption.

If liability is accepted by the cloud provider in the event of service outage, monetary compensation is rarely issued to customers. Instead, major cloud providers such as Amazon Web Services, Google, Salesforce, or Microsoft all issue credits that entitle the customer to a certain amount of free usage of the cloud provider's services. Credits are calculated as a percentage of the fees paid for a cloud provider's services that were adversely affected by a failure within the current monthly billing period and are applied at the end of the billing cycle (*Cohen, 2013*).

Another issue is jurisdiction. The cloud may be thought of conceptually as an entity not defined by geography. However, servers, customers utilising them, and attackers trying to compromise them all have a physical location. As such, there are considerable complexities involved with which jurisdiction's laws apply during a particular downtime event. For example, if a cloud provider's customer is European, the cloud storage is based in Asia, and a cloud provider is incorporated in the United States, there is no clear answer as to whose laws would govern (*Sullivan, 2015*).

The legislative landscape of data protection is also changing rapidly with, for example, the General Data Protection Regulation coming into effect in 2018 throughout the EU, which will add further uncertainty to cloud data breaches of an international nature (*Ungerleider, 2014*).

Reinsurance

With the growing realisation of both the business opportunity as well as the scale of cyber risk and its mechanisms of accumulation, reinsurance of cyber risk is growing rapidly. However, cyber attacks are a relatively new phenomenon and the reinsurance market is still developing robust solutions for managing and mitigating these risks.

Where exclusions do exist, these are increasingly being removed. This is due to competitive pressures created by soft market conditions in which some clients and brokers are not willing to accept them. One example is the widely used London Market Cyber Attack Exclusion (CL 380), which covers a broad range of marine, energy and industrial property policies. Another example is the Electronic Data Exclusion (NMA2914), which excludes data breaches following a computer virus (*Z/Yen Group, 2015*).

However, by deleting exclusions, there are concerns as to whether appropriate coverage is being provided within this very technical class. Even when exclusions are included, within the rapidly evolving nature of the risk, an exclusion crafted today could become obsolete within six months. In this case, it seems likely that the cyber market may develop in a fashion similar to the terrorism market after September 11. That is, as more tailored cyber coverages are developed and the market matures, reinsurers will be able to incorporate suitable cyber exclusions into the coverages for traditional classes (*Cook, 2016*).

There have also been discussions of a public-private cyber reinsurance scheme (or to extend an existing scheme, e.g., Pool Re), whereby the government helps the insurance industry fund the extreme losses of cyber risk. For example, the government could take responsibility for business interruption risks above a point, say £100 million. Below that point, normal insurers assume the risk from the cyber policies they write, while educating customers about security and best practices (*Z/Yen Group, 2017*).

Applicability of modelling methodology for other scenario analyses



7. Applicability of modelling methodology for other scenario analyses

The detailed accumulation approach is not limited to scenario analysis of cloud downtime and its business interruption impact. This approach describes a framework for measuring the systemic risk associated with any vulnerability that may be common across a group of companies. The key to applying this framework for multiple other scenarios is collecting the same detailed information about the companies being evaluated and storing that in a homogeneous format such as the Lloyd's Common Core Cyber Exposure Data Standard, which was established through collaboration with AIR, Lloyd's, the London Market Association, the Cambridge Centre for Risk Studies and RMS. With such a dataset, risk managers can then identify the most common points of aggregation within the group and develop the appropriate scenarios to estimate the severity of their systemic risk. Examples are provided in the following sections.

Domain name system (DNS) provider

Domain Name Systems (DNS) providers are Internet "traffic directors" that map domain names to internet protocol (IP) addresses so Internet traffic can get to its destination. The typical Internet user relies on Domain Name Systems for accessing websites using manageable domain names (such as lloyds.com) and not the more challenging to remember IP addresses (such as 125.114.175.220). A meaningful scenario would be a DDoS attack that targets one DNS provider, causing an interruption of its service. All the companies who have outsourced their DNS to that provider would find that their websites may no longer be accessible and turnover may be lost or business operations may be halted.

Payment acquirer/processor

Payment acquirers and processors are companies appointed by merchants to handle the transfer of funds from an issuing bank when a customer makes a purchase using a credit or debit card. Acquirer/processors represent the greatest sources of risk in the payment processing chain since a small number of providers dominate the market. A cyber attack that shuts down an acquirer's/payment processor's service would leave merchants without the ability to collect turnover from clients.

Content delivery network (CDN) provider

A content delivery network—also known as a content distribution network or CDN—is a large, geographically distributed network of servers that accelerate the delivery of web content and rich media to internet-connected devices. Content providers contract CDN providers to deliver their internet content to their users. If a CDN service were to be compromised by a cyber attack, a large number of company websites may become inaccessible or sensitive user data such as emails and passwords may be exposed.

SSL certificates provider

Secure Socket Layer (SSL) certificates are technology that enables encrypted communication between a web browser and web server. If a company's SSL certificates expire, its website becomes inaccessible and a business interruption loss could occur. Certificate authorities (CAs) issue these certificates. If a cyber attack leaves a CA incapable of providing this service and the makers of browsers or operating systems decide to revoke trust in the CA's certificates, companies that rely on that provider for security certificates are no longer protected. No e-business can be done from their websites until the certificate is re-authenticated and certified, perhaps with a different CA. Such a revocation of trust occurred in the past to the Dutch company Diginotar, following which it filed for bankruptcy (*Zetter, 2011*), and a similar

revocation of trust in Symantec certificates by Google's Chrome browser is in the process of being implemented (O'Brien, 2017).

Ad network provider

Ad network companies are contracted to publish advertisements on the web, desktop or mobile devices. They have algorithms that target specific audiences, making the ads more effective. Some companies may rely on these ads to drive turnover. If the contracted ad network goes down as a result of a cyber attack, its ads are no longer visible, its clients' website traffic may go down, and turnover may be lost.

Vulnerable or unsupported software


Hackers are constantly looking for holes within popular software products to exploit. These vulnerabilities do not get addressed by vendors when a product is no longer supported. Despite this risk, many people continue to rely on unsupported products. For example, about 52% of businesses are still running Windows XP on some

devices (Tsai, 2017), which is no longer supported by Microsoft. If a new flaw is discovered by hackers, it will not be patched, potentially leading to a coordinated cyber attack on all users of this software or a cyber infiltration of many of companies.

Data aggregators

Suppliers who choose to deliver their services using a cloud model can have customers across many industries, meaning they may be continuously collecting sensitive data such as healthcare records, employee or payroll data, or financial transactions. When these data aggregators capture significant market share, they can become attractive targets for hackers who seek to obtain a large prize. Cloud companies like this typically operate on a shared security model. This means that the cloud provider is responsible for the security of the infrastructure, but the customer is ultimately responsible for the security of its data. If a cloud provider suffers a data breach, it's possible that multiple customers may be liable for significant expenses related to regulatory fines and lawsuits.

Conclusion



8. Conclusion

Cyber events have the potential to cause catastrophic losses, with far-reaching impacts that extend beyond the breached businesses and their insurers. Cyber risk is a high severity risk like the natural perils that traditional catastrophe models were created to address. And, as is the case with perils like hurricanes and earthquakes, the historical record alone for cyber events is not sufficient to estimate the full spectrum of impacts from future attacks. Many attacks go undetected, and companies are often hesitant to publicise that they have been breached. Consequently, for cyber events – such as cloud downtimes – whose impact is distributed over many victim companies, unless the methodology defined by this paper to reliably identify likely aggregations of victim companies is applied, one may not reasonably expect to account for the resulting damage comprehensively. Furthermore, the constantly evolving nature of cyber risk makes it all the more challenging to use past events to project future losses, and the human element in cyber attacks adds to the uncertainty. Cyber criminals are becoming increasingly sophisticated, attacks are happening on a larger scale and are harder to stop, and the ever-expanding internet of things is broadening the range of possible targets.

Among insurers, there is widespread recognition of the potential for extreme accumulated losses from a cyber event, be it from an attack on a cloud provider or payment processor, a power grid attack, massive data theft aggregation event, exploiting a weakness in a commonly used software application, or any one of a number of other nightmare scenarios. Incidents with widespread impact like the Dyn attack in October 2016 and the WannaCry ransomware attack in May 2017 will only become more frequent, and a truly catastrophic cyber event has yet to occur. Although the scope of this paper is limited to downtimes for cloud service providers, the methodology described herein can be extended and adapted to handle DNS service provider failures, ransomware attacks, and other aggregation events.

While awareness of the risk is rising, the penetration of cyber insurance is estimated at less than 30% in the United States (where 90% of premiums are currently being underwritten), and much lower in the UK and other countries. Many insurers set relatively low limits, waiting periods exempt from claims, and a multitude of exclusions to try to control their losses. Methods

commonly used for managing accumulations, based on estimated market share, can offer a crude approximation of the risk, but miss the mark more often than not because they cannot reveal the hidden correlations and sources of risk aggregation in cyber portfolios. Detailed accumulation approaches are a seminal advance in risk aggregation modelling. Furthermore, most cyber policies today cover the direct cost of breaches and third-party liability, while business interruption and contingent business interruption coverage are less common. It is clear that there is much room for growth for cyber insurance, and cyber risk management is still a nascent discipline. Cyber models allow re/insurers to objectively quantify the risk so that they can more confidently provide coverage.

Improving insurance take-up will help businesses and communities become more resilient to the potentially catastrophic impact of cyber-related losses. While larger companies may have reserves and loss response measures to withstand an attack, small and medium enterprises can easily fail without the essential protection that insurance affords. In fact, many underwriters offer incidence response services that help companies in the aftermath of an attack to resume business operations quickly. Such services include eradicating viruses, notifying stakeholders, and improving customer and public relations. Insurance can also incentivize companies to improve their cyber security measures by way of risk-based premiums reductions, much like mitigation discounts work for property owners.

Of course, it is not only corporations that are at risk. Hospitals, utilities, transportation, and other critical infrastructure can be targeted, with worrisome impact to communities. Recognizing the potentially catastrophic impact of such attacks, the EU has recently signed new legislation that requires the providers of essential services to report incidents and demonstrate sound cybersecurity measures. Notably, the classification of “essential services” includes not only infrastructure, healthcare, and financial institutions, but also cloud and search engine providers, a testament to today’s significant reliance on digital services. Similarly, in the United States, the Federal Government could recognize cloud service providers as “critical infrastructure” and manage risks to them in similar ways as they would manage risks to hospitals or power plants.


This white paper describes a novel approach to modelling aggregate losses resulting from cloud downtimes. The methodology for modelling cloud downtimes is of interest in and of itself. The modelled ground up lossestimates derived using that methodology serve as a proof of concept that this peril scales as a viable insurable risk. That is to say, even extremely severe scenarios – those with exceedance probabilities indicative of return periods in the hundreds of years – can be planned for and covered by policies with reasonable premiums, deductibles, and limits. So, for example, the leading cloud service provider in the US suffering a downtime of 3-6 days is estimated to cause insured industry losses of about \$3 billion, given today's market conditions. The loss to a typical book of business would be less, because even the largest cyber book today has approximately a 20% market share. It is important to

note, however, that whereas a crude estimate of how much loss a particular book of business would sustain in such a scenario may be obtained by down-scaling the industry estimates, far more accurate estimates can be obtained by utilizing the detailed aggregation approach that was described in this paper.

Cyber risk has become a top of mind concern for risk managers across all public and private sectors. Organisations large and small are investing in risk and loss mitigation, including preventative security and post-event recovery measures. The continued expansion of the cyber insurance market is both necessary and inevitable. Taking proactive measures now to build a risk-based cyber insurance ecosystem, ahead of the next truly catastrophic event, is essential to establishing more resilient communities and businesses.

Appendix A. Historical cloud events

Figure 17: Major cloud downtime



A cloud downtime event can occur as a result of natural causes (such as lightning storms) as well as human error or malicious intent. Cloud vendors may not publically report attacks in an effort to protect their reputation, or an attack such as a data exfiltration may go unnoticed by the cloud vendor. Over the last decade, as cloud services have become more widely adopted, there have been many cloud downtime events due to human error or natural causes. A few major events are described in Figure 17.

Microsoft Azure "Leap Year Bug"

February 2012. An outage was caused by a software bug for Microsoft Azure security certificates on the 29th of February "Leap Day". It was a complicated disruption which impacted many Microsoft Azure components (Laing, 2012). Customers were unable to access Azure hosted services during the outage, some for up to 8 hours. The Azure team gradually deployed an update to the software to fix the issue. Microsoft Azure reviewed the incident and its response procedures in detail to ensure this "Leap Year Bug" would be prevented in the future. Service credits were provided to customers who were directly affected by the outage.

Amazon Web Services US-East-1 Outage

June 2012. Amazon Web Services US-East-1 region suffered downtime due to multiple generator failures that left the region without power and drained the emergency

uninterruptible power supply (UPS) units. This event occurred only two weeks after a similar outage also caused by generator and electrical equipment problems. During the event, customers should have been able to share their workload across other data centres via the Elastic Load Balancing (ELB) system. However, a bug in the ELB led to a backlog of traffic shift requests for affected customers that could not be met in a timely fashion (AWS Team, 2012).

Microsoft Azure Global Outage (SSL Certificate Expiration)

February 2013. Microsoft's Azure cloud platform experienced a worldwide outage in its storage services because of an expired SSL certificate (Ribeiro, 2013). The global outage lasted almost 12 hours. Azure proactively stated that it would provide credits to affected customers who were running any of five services in the form of a 25% credit for charges for these services for the impacted billing period, as well as a 25% credit on any data transfer usage (Neil, 2013).

Google Outage by Lightning Strike

August 2015. In Saint-Ghislain, Belgium, four consecutive lightning strikes to the local power grid caused failures in the Google Compute Engine of Google's cloud data centre. Although storage systems had battery backup, some recently written data was lost because some storage systems had experienced extended or repeated power drain. In total, Google claims

that it lost only 0.000001% of its data in the affected location (Google, 2015). The original estimate was higher, but some of the disks considered lost became accessible, and customers that had previously backed up data could recover the lost data (Hinks, 2017).

Verizon Communications Inc Outage

January 2016. A power outage was experienced at a Verizon Communications Inc. data centre caused by a maintenance operation (Gates, 2016). This outage lasted three hours (Nichols, 2016) with immediate effects experienced by JetBlue airways. Flights were grounded for several hours as there was no access to essential business IT systems. Verizon did not comment on the specific cause of the outage.

Microsoft Office 365 Email Outage

January 2016. A faulty update prevented Microsoft Office 365 users from downloading emails from the Exchange Online database via internet message access protocol (IMAP) (Curtis, 2016). The problem was caused by a code issue. A fix was implemented, but users continued to experience access issues, due to a configuration that delayed deployment of this fix. The outage lasted up to five days for some users (Tsidulko, 2016).

Salesforce Circuit Fault to NA14

May 2016. Salesforce's North American 14 (NA14) site experienced an outage for one day, causing major

disruption for users along the West Coast of the United States (Davis, 2016). Salesforce posted an extensive review of the cause and preventative action was taken. The cause was determined to be a failed circuit breaker, which caused a power failure to a Salesforce data centre in Washington, D.C. (Salesforce, 2016). A backup secondary data centre in Chicago was brought online to restore service but NA14 continued to experience poor performance. A database failure occurred due to an increase in traffic. Full functionality was reported six days later. Salesforce audited all its data centres to isolate those with the same potential defect as the Washington D.C. centre and replaced the faulty circuit breakers to prevent future outage scenarios from occurring.

Amazon Web Services S3 Outage

February, 2017. The Amazon Web Services AWS Simple Storage Service (S3), which provides hosting for images, entire websites, and app back ends, experienced a severe, four hour disruption in the US-EAST-1 region that affected some websites for up to 11 hours (Eide and O'Shea, 2017). As AWS explained in a posted message, a debugging team executed a command to remove a few servers in a subsystem used by S3. However, due to a typo, a larger set of servers was removed than intended, including servers that supported two additional subsystems. Each of these systems required a full restart, during which time S3 could not service requests (Amazon, 2017). The outage affected major websites and service like Quora, Coursera, Expedia, GitHub, Trello, and many more, as well as devices in the Internet of Things (IoT) such as Nest thermostats and cell phone apps (Novet, 2017).

Appendix B. Cloud resilience

Cloud computing cybersecurity overview

When it comes to security in the cloud, meeting requirements is much different than for on premises IT operations. With two parties involved, the service provider and the customer, the questions of ownership and control over cloud deployment are at the heart of cloud security. There are two key differences between cyber risk on premise and cyber risk on the public cloud, and they reflect the fact that just because a business has chosen to move to cloud, it doesn't mean that they have offloaded all their risk to the cloud service provider.

Shared Responsibility

Shared responsibility in the public cloud means that both cloud providers and cloud customers are responsible for the security of data stored on the cloud. Who is responsible for what, in terms of security, depends on the cloud service model being used (IaaS, PaaS, or SaaS).

On the IaaS end of the spectrum, the cloud service provider is responsible for core infrastructure security, which includes storage, networking, and computing, at least at the physical level. As cloud customers choose to move from IaaS, to PaaS, and then to SaaS, they'll find that they're responsible for less and the cloud service provider is responsible for more. Notably, the shared responsibility model leaves the cloud customer fully accountable for the data that is being stored outside the business, which in the event of a breach makes them most liable for any third-party damages or responsible for regulatory action. Figure 18 below is a representation of Microsoft's approach to the shared responsibility model. Note that On-Prem stands for "on premises," i.e., computing resources not hosted on the public cloud.

Figure 18: Microsoft's approach to the shared responsibility model

Responsibility	On-Prem	IaaS	PaaS	SaaS
Data classification & accountability	Cloud Customer	Cloud Customer	Cloud Customer	Cloud Customer
Client & end-point protection	Cloud Customer	Cloud Customer	Cloud Customer	Cloud Customer / Cloud Provider
Identity & access management	Cloud Customer	Cloud Customer	Cloud Customer / Cloud Provider	Cloud Customer / Cloud Provider
Application level controls	Cloud Customer	Cloud Customer	Cloud Customer / Cloud Provider	Cloud Provider
Network controls	Cloud Customer	Cloud Customer / Cloud Provider	Cloud Provider	Cloud Provider
Host infrastructure	Cloud Customer	Cloud Customer / Cloud Provider	Cloud Provider	Cloud Provider
Physical security	Cloud Customer	Cloud Provider	Cloud Provider	Cloud Provider

Legend: ■ Cloud Customer, ■ Cloud Provider

Source: Shinder, 2016

Most other cloud service providers follow a similar model. For example, Amazon's shared responsibility model (Amazon, 2017(2)) makes the distinction between "security of the cloud", which Amazon is responsible for, and "security in the cloud", which the cloud customer is responsible for (See Appendix A for more information).

Multi-Tenancy

One of the characteristics of a cloud service is multi-tenancy, which allows the provider to group together the same IT resources to serve multiple customers. By doing so, cloud service providers can optimise the use of their assets and lower the costs for the customer.


Similar to the shared responsibility model, multi-tenancy can be described in degrees depending on the cloud deployment model. IaaS and PaaS models tend to be multi-tenant with clients consuming from the same pool of computing resources. On the other hand, SaaS can be multi-tenant or not. Some SaaS offerings leverage a common database schema or business logic, resulting in data from multiple customers being stored in the same location, while others offer customised and dedicated software applications.

The primary security concern of multi-tenancy is how to ensure that proper security and isolation protect consumers from the risks they pose to one another. Some strategies for keeping data separate in a multi-tenant environment include:

- **Database and Virtual machine segmentation.** Firewalls that separate databases and virtual machines owned by different clients that may have different security settings but are hosted within the same infrastructure prevent data from being accessed or spilling over unintendedly.
- **Virtual machine introspection.** Virtual machine introspection lets the cloud customer gather information about virtual machines, virtual network security, and virtual environment settings. With this data customers can tailor the security policies to address to reality of operating in a multi-tenant environment.
- **Virtual private cloud subscription.** Virtual private clouds are similar to public clouds with the exception that exclusively provisioned hardware, network, and storage configurations are provided, usually at a higher cost.

What influences cloud service downtime

Figure 19: Microsoft's approach to the shared responsibility model



The resilience of a cloud service can also be assessed by examining the risk mitigation approaches used to avoid any significant service downtime. These approaches can be broken down as activities to ensure the availability of the essential functions of a cloud service provider's facilities, hardware, software, and networking components (Figure 19).

Facilities

A cloud service provider's data centres are dependent upon the successful and integrated operation of electrical, mechanical, and building systems. Organisations such as the Uptime Institute have created frameworks that evaluate various data centre facilities in terms of potential site infrastructure performance, or uptime. Below are the descriptions of the Uptime Institute's tiered approach (Uptime Institute, 2014) to assessing the functionality, capacity, and expected availability of data centres.

- **Tier I: Basic capacity.** At the most basic level, all data centres should have dedicated space for the IT systems, and uninterruptible power supply (UPS), dedicated cooling equipment, and an engine power generator.
- **Tier II: Redundant capacity components.** Redundant power and cooling components such as UPS modules, chillers or pumps, and engine generators are added to increase the resilience to outages.
- **Tier III: Concurrently maintainable.** Tier III data centres must be able to continue operating even when equipment is scheduled to be replaced or maintained upon.

- **Tier IV: Fault tolerance.** A Tier IV site infrastructure builds on Tier III, adding the ability to withstand unexpected equipment failures without an interruption of service.

Lacking from this standard is any consideration to external factors that may affect a data centre's facilities from operating uninterrupted. These factors include meeting local building and occupancy code regulations; resilience to seismic shocks, extreme weather; flooding; or the impacts of adjacent property uses, union or other labour organisation force activities; and physical security.

Hardware

The resilience of cloud-based services and applications is also dependent on the functional availability of the hardware component of the provider's technology stack. Hardware typically found in data centres includes file servers, database servers and storage areas. The following are examples of risk mitigation strategies for managing the failure hardware components without compromising the cloud service.

- **Link aggregation.** Link aggregation combines multiple network connections to increase the network's capacity for sudden increases in demand, thereby improving the network's reliability.
- **Hot-swappable interfaces.** Hot swapping is a way to add or replace hardware components on the fly without interrupting service.
- **Use of high MTTF hardware.** Mean Time to Failure (MTTF) is a measurement of a hardware component's expected life. Cloud providers can use this metric to design data centres with appropriate lifespans and to schedule maintenance and repair accordingly.

Software

Similar to the hardware components, there are layers of software that cloud providers need in order to provide their service. This includes operating systems, virtualization, middleware, data management and other functionality. The following are examples of risk mitigation strategies for software.

- **Use of application pools.** An application pool is a group of separate applications that share the same resources distributed throughout the network. This creates a level of isolation between each application so that errors in one application pool will not affect the applications running in other application pools.
- **Risk mitigation for critical applications.** Some resources may not be able to be replicated. Cloud providers should identify these and have risk mitigation plans in place in the case of failure.
- **Virtual machine migration.** In the event of a hardware failure, virtual machines can be replicated and migrated to other servers or data centres without interrupting service.
- **In-service software upgrades (ISSU).** ISSU updates software without taking the resource offline. This allows cloud providers to resolve any software issues without interrupting service.

Networking

The core value proposition of the cloud—the ability to access the service from anywhere with an internet connection—relies on the availability of networking capabilities. The networking component is not only required from a client delivery perspective, but also for the successful execution of internal processes. One type of risk mitigation for cloud networks is redundancy, in which alternate network paths can be used if the primary path becomes unavailable.

Appendix C. E-business factors

This paper accounts for e-business costs to the set of United States companies and the subset of Fortune 1000 companies that arise from the sustained loss of access to a business service, namely a cloud service provider. Although e-business can be broader in scope, this paper accounts for costs only from e-commerce sales and turnovers, e-shipments, m-commerce sales and shipments, and electronic order management systems. These costs include both business to consumer (B2C) and business to business (B2B) transactions.

Although there are other forms of e-business, the kinds of transactions listed above define comprehensively those e-business transactions that are in the scope of this paper. By contrast, a few examples of other e-business processes and costs that are *not* accounted for in this paper are outlined below.

- Losses arising from the impact of a cloud provider failure on certain internal business processes though they may be hosted on the cloud. These processes can be industry specific and a few cases are listed below:
 - Healthcare industry - loss of access to patient records hosted on the cloud
 - Banking – loss of access to client information hosted online and not being able to conduct business as usual
 - Education – loss of access to student records and not being able to collect fees etc.
 - Insurance - Using excel online for underwriting and not be able to conduct business as usual
 - Any industry – loss of access to human resource management systems dealing with payroll activities, benefits administration and other human resource purviews.
 - Any industry - Business processes aimed at cost savings, improvements in efficiency, quality assurance and control
- Losses arising from lawsuits due to the cloud downtime
- Fines, reputation damage and public relations costs
- Productivity costs that can be measured in terms of salaries, wages and benefits of employees made idle by the cloud downtime.

References

- Advisen. 2014. *Cyber insurance underwriting: A high-tech, evolving discipline*. Available from: <http://www.advisenltd.com/wp-content/uploads/cyber-insurance-underwriting-high-tech-evolving-discipline-white-paper-2014-11-06.pdf>
- Amazon, 2011. *Summary of the Amazon EC2 and Amazon RDS Service Disruption in the US East Region*. Available at: <http://aws.amazon.com/message/65648>.
- Amazon. 2017. *Summary of the Amazon S3 Service Disruption in the Northern Virginia (US-EAST-1) Region*. Available at: <https://aws.amazon.com/message/41926/>
- Amazon. 2017(2). *Shared Responsibility Model*. Available at: <https://aws.amazon.com/compliance/shared-responsibility-model/>
- Availability Digest, 2013. *Windows Azure Downed by Single Point of Failure*. Available at: http://www.availabilitydigest.com/public_articles/0811/azure.pdf
- AWS Team. 2012. *Summary of the AWS Service Event in the US East Region*. AWS Website. Available at: <https://aws.amazon.com/message/67457/>
- Bright, P. 2011. *Amazon's lengthy cloud outage shows the danger of complexity*. Ars Technica. Available at: <https://arstechnica.com/business/2011/04/amazons-lengthy-cloud-outage-shows-the-danger-of-complexity/>
- Butler Consultants. 2017. *Free Industry Statistics: Gross Margin*. Available at: <http://research.financial-projections.com/IndustryStats-GrossMargin.shtml>.
- Cohen, R. 2013. *New Cloud Computing Insurance Attempts to Solve Cloud Liability Concerns for Service Providers*. Forbes. Available at: <https://www.forbes.com/sites/reuvencohen/2013/04/24/new-cloud-computing-insurance-tries-to-solve-cloud-liability-concerns-for-service-providers/#7daad0721970>
- Cook, S. 2016. *Cyber Risks and Reinsurance*. XL Catlin. Available at: <http://xlcatlin.com/fast-fast-forward/articles/cyber-risks-and-reinsurance>
- Curtis, J. 2016. *Faulty Microsoft customers still can't access Office 365 emails*. CloudPro. Available at: <http://www.cloudpro.co.uk/collaboration/productivity/5770/microsoft-customers-still-can-t-access-office-365-emails>
- Davis, J. 2016. *Salesforce Outage: Can Customers Trust The Cloud?* InformationWeek. Available at: <https://arxiv.org/pdf/1010.6118.pdf>
- Dukic, V. M., Marić, N. 2013. *On Minimum Correlation in Construction of Multivariate Distributions*. arXiv. Available at: <http://www.informationweek.com/cloud/platform-as-a-service/salesforce-outage-can-customers-trust-the-cloud/d/d-id/1325499>
- Eide, N., O'Shea, D. 2017. *Amazon S3 disruption dramatically slowed top retail websites*. CIODIVE. Available at: <http://www.ciodive.com/news/amazon-s3-disruption-dramatically-slowed-top-retail-websites/437270>
- Elumalai A., Starikova I., and Tandon S. 2016. *IT as a service: From build to consume*. Available at: <http://www.mckinsey.com/industries/high-tech/our-insights/IT-as-a-service-From-build-to-consume>
- Finos R. 2015. *Public Cloud Market Forecast 2015-2026*. Wikibon. Available at: <https://wikibon.com/public-cloud-market-forecast-2015-2026/>.
- Frost & Sullivan. 2017. *Future of B2B Online Retailing*. Available at: <https://www.researchandmarkets.com/reports/4377289/future-of-b2b-online-retailing#pos-2>

-
- Gates, R. 2016. *JetBlue, Verizon data center downtime raises DR, UPS questions*. TechTarget. Available at: <http://searchdatacenter.techtarget.com/news/4500271203/JetBlue-Verizon-data-center-downtime-raises-DR-UPS-questions>
- Gilbert, F. 2011. *Cloud computing contracts and cloud outages*. TechTarget. Available at: <http://searchcloudsecurity.techtarget.com/tip/Cloud-computing-contracts-and-cloud-outages>
- Google. 2015. *Google Compute Engine Incident #15056*. Available at: <https://status.cloud.google.com/incident/compute/15056#5719570367119360>
- Gunawi, H., et al. 2016. *Why Does the Cloud Stop Computing? Lessons from Hundreds of Service Outages. Proceedings of the Seventh ACM Symposium on Cloud Computing (SoCC)*, pp. 1-16, Available at: <http://ucare.cs.uchicago.edu/pdf/socc16-cos.pdf>
- Hinks, J. 2017. *Google suffers data loss following lightning strikes*." TechRadar. Available at: <http://www.techradar.com/news/networking/routers-storage/google-suffers-data-loss-following-lightning-strikes-1302271>
- Hiscox. 2017. *The Hiscox Cyber Readiness Report 2017*. Available from: <https://www.hiscox.co.uk/cyber-readiness-report/docs/cyber-readiness-report-2017.pdf>
- Marsh. 2017. *Addressing Cyber Risk*. Available from: https://www.treasury.gov/initiatives/fio/Documents/1-Cyber_Insurance_Market_MarshLLC.pdf
- Laing, B. 2012. *Summary of Windows Azure Service Disruption on Feb 29th, 2012*. Microsoft Azure Identity & Access Management Blog. Available at: <https://azure.microsoft.com/en-us/blog/summary-of-windows-azure-service-disruption-on-feb-29th-2012/>
- LMI Group, 2017. *History of Business Interruption Insurance*. Available at: <http://cms.lmigroup.com/bi-explained/au/history-of-business-interruption-insurance/>
- Lloyd's and Cyence, 2017 *Counting the cost*. Available at: <https://www.lloyds.com/~media/files/news-and-insight/risk-insight/2017/cyence/emerging-risk-report-2017---counting-the-cost.pdf>
- Lynley, M. 2013. *The High Cost Of An Amazon Outage*. BuzzFeed. Available at: <https://www.buzzfeed.com/mattlynley/the-high-cost-of-an-amazon-outage>
- Mell, P. and Grance, T. 2011. *The NIST Definition of Cloud Computing*. National Institute of Standards and Technology. Available at: <http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf>
- Nichols, S. 2016. *JetBlue blames Verizon after data center outage cripples flights*. The Register. Available at: https://www.theregister.co.uk/2016/01/14/jetblue_blames_verizon_data_center_crash/
- Neil, M. 2013. *Details of the February 22nd 2013 Windows Azure Storage Disruption*. Microsoft Azure. Available at: <https://azure.microsoft.com/en-us/blog/details-of-the-february-22nd-2013-windows-azure-storage-disruption/>
- Neustar, 2015. *Neustar Global DDoS Report Reveals Seven Key Trends, Including Attacks Evolving from Disruptive to Continuous Threat*. Neustar. Available at: <https://www.neustar.biz/about-us/news-room/press-releases/2015/ddos-report-pr>
- Novet, J. 2017. *AWS is investigating S3 issues, affecting Quora, Slack, Trello (updated)*. Venture Beat. Available at: <https://venturebeat.com/2017/02/28/aws-is-investigating-s3-issues-affecting-quora-slack-trello/>
- O'Brien, D., Sleevi, R., Whalley, A. 2017. *Chrome's Plan to Distrust Symantec Certificates*. Google Security Blog. Available at: <https://security.googleblog.com/2017/09/chromes-plan-to-distrust-symantec.html>
- Ribeiro, J. 2013. *Microsoft Azure service restored after being downed by expired SSL certificate*. PCWorld. Available at: <http://www.pcworld.com/article/2029188/microsofts-azure-service-falls-to-expired-ssl-certificate.html>
- RightScale. 2017. *2017 State of the Cloud Report*. Rightscale. Available at: <https://www.rightscale.com/lp/2017-state-of-the-cloud-report>
- Salesforce. 2016. *2RCM for NA14 Disruptions of Service - May 2016*. Available at: https://help.salesforce.com/articleView?id=Root-Cause-Message-for-Disruption-of-Service-on-NA14-May-2016&language=en_US&type=1
- Shinder, T. 2016. *What Does Shared Responsibility in the Cloud Mean? – April 2016*. Available at: <https://blogs.msdn.microsoft.com/azuresecurity/2016/04/18/what-does-shared-responsibility-in-the-cloud-mean/>
- Sullivan, C. 2015. *Who Is Liable When the Cloud Is Hacked?* FindLaw. Available at: <http://blogs.findlaw.com/technologist/2015/07/who-is-liable-when-the-cloud-is-hacked.html>

-
- Tsai, P. 2017. *Windows 10 adoption surges, yet businesses still hang on to Windows XP and Vista*. Spiceworks. Available at: <https://community.spiceworks.com/networking/articles/2628-windows-10-adoption-surges-yet-businesses-still-hang-on-to-windows-xp-and-vista>
- Tsidulko, J. 2016. *The 10 Biggest Cloud Outages Of 2016 (So Far)*. CRN. Available at: <http://www.crn.com/slideshows/cloud/300081477/the-10-biggest-cloud-outages-of-2016-so-far.htm/pgno/0/3>
- Ungerleider, N. 2014. *Who Is Liable When Cloud Services Are Hacked?* Fast Company. Available at: <https://www.fastcompany.com/3035104/who-is-liable-when-cloud-services-are-hacked>
- United States Census Bureau. 2016. *E-Stats 2014: Measuring the Electronic Economy*. U.S. Census Bureau. Available at: <https://www.census.gov/content/dam/Census/library/publications/2016/econ/e14-estats.pdf>
- United States Census Bureau. 2015. *2015 E-commerce Multi-sector Data Tables*. U.S. Census Bureau. Available at: <https://www.census.gov/data/tables/2015/econ/e-stats/2015-e-stats.html>
- Uptime Institute. 2014. *Data Center Site Infrastructure Tier Standard: Topology*. Available at: <https://uptimeinstitute.com/publications/asset/tier-standard-topology>
- Zetter, K. 2014. *Hacker Lexicon: What is a Zero Day?* Wired.com. Available at: <https://www.wired.com/2014/11/what-is-a-zero-day/>
- Zetter, K. 2011. *Diginotar Files for Bankruptcy in Wake of Devastating Hack*. Wired.com. Available at: <https://www.wired.com/2011/09/diginotar-bankruptcy/>
- Z/Yen Group. 2015. *Promoting UK Cyber Prosperity: Public-Private Cyber-Catastrophe Reinsurance*. Long Finance. Available at: http://www.longfinance.net/images/Promoting_UK_Cyber_Prosperty_28July2015.pdf
- Z/Yen Group. 2017. *Cyber-Catastrophe Reinsurance*. Long Finance. Available at: <http://www.longfinance.net/cyber-articles/780-cyber-reinsurance.html>