CAS RESEARCH PAPER

EXPOSURE MEASURES FOR PRICING AND ANALYZING THE RISKS IN CYBER INSURANCE

Michael A. Bean, FCAS, CERA, FCIA, FSA, Ph.D.

Sponsored by

Casualty Actuarial Society
and Society of Actuaries

Exposure Measures for Pricing and Analyzing the Risks in Cyber Insurance

Michael A. Bean, FCAS, CERA, FCIA, FSA, Ph.D.

April 2020

Contents

1. Summary

This report identifies and makes recommendations regarding possible exposure measures for pricing and analyzing the risks in cyber insurance. Cyber insurance is an insurance product that is designed to provide protection against the financial consequences of a failure or compromise of an organization's information system as a result of a cyber event. A cyber event is an event that compromises the availability, integrity or confidentiality of an organization's information system or electronic data in some way. Examples of cyber events are a cyberattack on an information system or the unintentional disclosure of electronic medical records due to human error.

Cyber insurance has been available in various forms since the 1990s but is still a relatively new product and continues to evolve. Most carriers of cyber insurance offer a core set of coverages, either as a package or part of a modular policy, as well as a number of supplementary coverages, which can vary from one carrier to another. Core coverages include privacy liability, network security liability, cyber event response, network interruption, recovery and restoration of digital assets, regulatory actions, fines and penalties, and payment card industry assessments. Supplementary coverages include cyber extortion, cybercrime, media content liability, technology errors and omissions, and bodily injury and property damage that results directly or indirectly from a cyber event.

This report uses a conceptual rather than an empirical approach to identify and evaluate potential exposure measures for cyber insurance. In particular, it considers the losses that can arise with each cyber insurance coverage, identifies potential exposure measures that are related to these losses, and then evaluates these potential exposure measures based on a set of criteria. Criteria used to evaluate potential exposure measures include ease of calculation, ability to audit the calculation, strength of relationship to losses, stability over the period of insurance coverage, and extent to which the candidate measure can legally be determined and shared with insurers or other third parties without violating privacy laws or regulatory requirements. An empirical approach to identifying and evaluating potential exposure measures is more

challenging due to the current scarcity of reliable, representative and publicly available loss experience for cyber insurance.

For coverages such as privacy liability where losses primarily involve exposure to individuals, it is convenient to distinguish between individuals who provide a product or service to an organization and those who receive it. Individuals who provide a product or service are considered employees while those who receive it are considered customers. For such coverages, the recommended exposure measure is the number of employees if this is greater than the number of customers; otherwise, it is the number of customers or accounts, depending on the type of data that is collected from customers. In most organizations, the number of customers is greater than the number of employees but in some, such as auto parts manufacturers, the reverse is true.

For coverages such as network security liability or recovery and restoration of digital assets where losses primarily involve exposure to computer systems, the recommended exposure measure is the number of endpoints (desktops, laptops, mobile devices, etc.) or the number of distinct user IDs, whichever is easier to determine.

For all other coverages, e.g., network interruption, the recommended exposure measure is revenue or sales over a specified period.

Whatever exposure measure is used for a coverage, it should be scaled appropriately based on the context.

2. Organization of Report

The report is organized as follows.

- Section 3 describes what cyber risk is, how it can arise, and the potential loss exposure associated with it, whether or not insurance is available to cover that exposure.

- Section 4 provides an overview of the cyber insurance coverages currently available in the market, with a focus on commercial coverages and policies where coverage of cyber risk is intended.

- Section 5 considers the historical experience on cyber events that is available in the public domain and what that experience tells us about the nature of cyber risk exposure.

- Section 6 describes the criteria used in this report to evaluate potential measures of exposure for cyber insurance.

- Section 7 identifies potential candidates for measuring the exposure associated with each cyber insurance coverage and considers the extent to which these candidates satisfy the criteria described in Section 6.

- Section 8 recommends the exposure measure that should be used for each type of cyber insurance coverage based on the observations made in Section 7.

3. The Nature of Cyber Risk Exposure

Cyber risk is the risk of failure or compromise of an organization's information system as a result of a cyber event. A **cyber event** is an event that compromises the availability, integrity or confidentiality of an organization's information system or electronic data in some way. The most obvious type of cyber event is a cyberattack. However, other types of cyber event are also possible, e.g., the unintentional posting of confidential medical information to a publicly accessible website or the unintentional corruption of a client's computer system following the installation of a new software application.

3.1. Exposure Associated with a Cyberattack

A **cyberattack** is an attack on an electronic device with an embedded computer chip or a network of such devices that is perpetrated by introducing or attempting to introduce erroneous or unauthorized electronic information into the device or network for the purpose of:

- damaging, destroying or disrupting the normal operation of the device or network,
- stealing, corrupting, erasing or precluding access to information that is stored on the device or within the network,
- using the resources of the device or network to
    - damage, destroy or disrupt the normal operation of another such device or network,
    - steal, corrupt, erase or preclude access to information on another device or within another network, or
    - steal the resources of another device or network, and/or
- using the device or network to cause harm in some other way.

The attack could be on a device or network that is local or part of a cloud computing structure.

Note that for an attack on a device or network to be considered a cyberattack, the device must contain a computer chip and the attack must be perpetrated by electronic means. So, a power surge that damages a conventional toaster would not be considered a cyberattack nor

would smashing a smart electricity meter with a hammer; however, using a smartphone to hack into a smart TV would be considered a cyberattack.

There are four basic types of cyberattack:

- Attacks that disrupt the normal functioning of a device or network and/or impair access to it without damaging or compromising the electronic information stored on the device or within the network.
- Attacks that damage, destroy or make unusable electronic information such as files, applications or firmware that is stored on a device or within a network.
- Attacks that result in the theft of nonpublic information.
- Attacks that commandeer the resources of a device or network for a malicious purpose.

We consider each type of cyberattack in turn.

3.1.1. Attacks That Disrupt Normal Functioning

A cyberattack can disrupt the normal functioning of a device or network by initiating frivolous processes or tasks that indiscriminately consume large quantities of resources and starve legitimate processes of the resources they need to perform their assigned tasks. The most common example of such an attack is a denial-of-service attack.

Denial-of-service attacks

A **denial-of-service attack** is an attack in which one or more malicious actors flood a device connected to a network or the network itself with superfluous traffic in an attempt to prevent legitimate traffic from reaching the device or using the network. The network in question is usually a public network such as the internet or a cellular telephone network but could also be a private network or one with restricted access. A denial-of-service attack can be likened to a convoy of cars or trucks that enters a superhighway for the sole purpose of disrupting traffic, overloading the highway's capacity and ultimately making the highway inaccessible and impassable.

The impact of a denial-of-service attack can range from relatively minor to quite significant and depends on the target of the attack, what services are disrupted, how critical the disrupted services are, when the disruption occurs and how long the disruption lasts. For example, if the target of the attack is the website of a retailer with a physical location, the disruption occurs during the middle of the night and it only lasts for a couple of hours then the impact is likely to be relatively minor. However, if the target of the attack is the website of an online bank, i.e., a bank with no physical location, the disruption occurs during normal waking hours and it lasts for several days then the impact is likely to be significant, possibly even fatal for the online bank.

Potential targets of a denial-of-service attack

There are many potential targets of a denial-of-service attack, for example:

- Retailers
    - The website of a bricks-and-mortar retailer that accepts online orders
    - The website of an online retailer, i.e., a retailer with no physical location
    - The logistics and inventory control systems that support the operations of a retailer
- Financial services providers
    - The website of a bank, credit union or similar financial services company that provides banking services to the general public
    - The ATM or payment card network of a bank, credit union or similar financial services company
    - The payment processing network of an electronic payments service such as Visa, MasterCard, PayPal or Square
    - The website of an online investment broker
    - The website of an insurance company
- Healthcare providers and insurers
    - The patient information exchange network of a health maintenance organization, hospital or similar group of healthcare providers

  - The claims submission network of a health insurer
  - Advanced diagnostic devices such as CT (computerized tomography) scanners or MRI (magnetic resonance imaging) machines
  - Robots or similar devices that perform high-precision surgical procedures
  - High-speed communication networks that facilitate surgeries in remote locations, i.e., surgeries where the patient and surgeon are in different geographic locations
- Telecommunications providers
  - The high-speed internet network of a telephone, cable or similar telecommunications company
  - The wireless communications network of a telephone, cable or similar telecommunications company
  - The legacy wireline network of a telephone, cable or similar telecommunications company
  - The servers of an email service provider
  - The website of an email service provider
- IT service providers
  - The servers of an information technology company that provides enterprise or cloud computing services to businesses and organizations that use technology to support their operations
  - The communications networks that connect the servers of IT service providers to their clients
- Utilities
  - The power grid of an electrical utility
  - The power generating stations of an electrical utility
  - The smart meters or sensors on the power grid of an electrical utility
  - The smart meters or sensors on the pipeline network of a gas or water utility
  - The wireless communication systems used to communicate with these smart meters or sensors

  - The website of an electrical utility that enables customers to monitor their electricity usage in real time and remotely control the settings of their heating and air conditioning systems
- Transportation systems
  - The communications network that facilitates remote monitoring and control of the signals, switches and trains on a rail network
  - The communications network that facilitates autonomous operation of the trains in a subway or surface passenger rail transit system
  - The communications network that facilitates remote monitoring of the buses in a metropolitan transit system
  - The communications network that facilitates remote monitoring and control of the traffic lights on a city's streets
  - The reservation system of an airline
  - The air traffic control system of an airport
- Manufacturers
  - The logistics and inventory control systems that ensure the components required for a manufacturing process are available at the time they are needed but not before
  - The robots on the production line of a precision manufacturing process

A denial-of-service attack can have many victims other than the intended target. For example, a denial-of-service attack that targets a particular telecommunications provider's network could inadvertently disable a hospital's MRI machine if that machine happens to be connected to the telecommunications provider's network and uses the network to send results electronically to doctors.

Potential losses associated with a denial-of-service attack

There are many different types of losses that can result from a denial-of-service attack. The type and amount of loss depends to a great extent on the victim and the nature of the attack.

For a retailer that relies on the internet for sales, the loss is likely to take the form of foregone sales; if the disruption occurs during a seasonally important time such as on Black Friday or Cyber Monday, this loss could be substantial. On the other hand, for an information technology company that provides enterprise and cloud computing services to other companies, the loss is likely to take the form of contractual penalties for failing to provide customers agreed-upon services; an IT service provider could also find itself liable for the losses its customers suffer as a result of a disruption, which could include losses that arise from third party lawsuits.

Virtually all victims of a denial-of-service attack will incur investigation and recovery costs. Depending on the victim and the severity of the attack, these costs could be substantial. The first step in responding to a denial-of-service attack is to determine the type of traffic and whether it originates from a single IP address or multiple addresses. If it originates from a single address, then blocking or diverting the traffic is fairly straightforward. However, if it originates from multiple addresses then blocking or diverting traffic may be insufficient and other measures such as removing the targeted device from the network or shutting down the network entirely may be necessary. This is particularly true in cases where attackers camouflage their identities (similar to what overseas telemarketers do when they route calls through a local exchange) or change their IP addresses as an attack progresses. Recovery from a denial-of-service attack can be a painstaking process. Most companies do not have the technical expertise required and will need to retain the services of outside experts.

Companies that provide a service that customers expect to be reliable and available when needed, e.g., telecommunications or financial services, can suffer a significant loss of reputation when a denial-of-service attack occurs. If the company providing the service operates in a competitive marketplace, i.e., one in which customers have alternatives, the result can be lost customers or difficulty acquiring new ones; from a financial perspective, a lost customer represents foregone future revenues and profits.

Companies such as banks, telecommunications providers or utilities that are regulated or belong to an industry association that requires adherence to particular standards can be subject

to fines or other penalties for failing to prevent or contain the damage caused by a denial-of-service attack. The fine or penalty levied usually depends on the severity of the lapse and how systemically important the company is to the network.

The variety of potential losses associated with a denial-of-service attack suggests that it could be difficult to determine a single quantity for measuring all the losses. Some losses appear to be related to sales, revenue or profit, others to the size or complexity of a company's computer system, and still others to a company's systemic importance. This observation has important implications when identifying possible exposure measures for cyber insurance.

Perpetrators of denial-of-service attacks

From the preceding discussion, it is clear that victims of denial-of-service attacks can suffer substantial financial losses, but it is less obvious what the perpetrators of such attacks have to gain. Unless the perpetrator is paid by some third party, there would appear to be little or no financial incentive to launch a denial-of-service attack.

The perpetrators of denial-of-service attacks typically fall into one of the following categories:

- State-affiliated actors engaged in various forms of cyber warfare or cyber terrorism, including testing the resilience of potential adversaries' critical infrastructure.
- Activists attempting to advance a particular political agenda, e.g., by targeting companies that sell products or conduct business in a way they find objectionable.
- Hackers or other "thrill-seekers" aiming to demonstrate their technical prowess.

Accordingly, the most frequent targets of denial-of-service attacks tend to be governments, financial service providers and educational institutions. However, any organization can be the victim of a denial-of-service attack at any time.

Denial-of-service attacks are sometimes used in conjunction with other types of attack to steal confidential information from unsuspecting individuals. For example, an attacker could set up an unsecured Wi-Fi network in a hotel lobby with a name similar to the hotel's legitimate

network and then launch a denial-of-service attack on the hotel's legitimate network. If hotel guests use the bogus network instead of the legitimate one, they will unknowingly reveal confidential information to the attacker. An attack of this type is known as a **man-in-the-middle attack**.

3.1.2. Attacks That Can Damage Electronic Information

A cyberattack can damage, destroy or make unusable electronic information that is stored on a device or within a network by penetrating the defenses of the device or a device on the network and inserting malicious software that destroys or makes inaccessible data that is on the device or within the network or compromises the functioning of applications or the operating system in some way.

Malicious software that is inserted into a program or copied to the memory of a device for the purpose of destroying data, running destructive or intrusive programs, or compromising the availability, integrity or confidentiality of data, applications or a device's operating system is generically referred to as **malware**. Note that the term malware encompasses malicious software that is designed to steal confidential information and computing resources as a well as software that simply damages, destroys or makes unusable the electronic information on a device. Hence, malware can be used to launch more than one type of cyberattack.

Types of malware

There are many different types of malware. The most basic types are viruses, worms and Trojan horses.

A **virus** is a piece of malicious software that propagates by inserting copies of itself into files and other programs and is activated when a host file or program is opened and/or executed. Viruses can cause damage in a number of ways; for example, they can corrupt or overwrite data files, cause programs to malfunction or even prevent them from running. There are two basic types of viruses: those executed by an application, known as **interpreted viruses**, and those executed by the operating system, known as **compiled viruses**. Generally speaking, compiled viruses are more difficult to remove than interpreted viruses.

A **worm** is a standalone piece of malicious software that is similar to a virus but does not require a host file or program to propagate; instead it propagates itself by taking advantage of vulnerabilities in a network or tricking users to open and run programs. Worms can cause damage that is similar to viruses. Worms can also be used to deliver other types of malware.

A **Trojan horse**, as the name suggests, is a standalone piece of software that appears to be innocuous but has a hidden malicious purpose. Users are usually tricked into loading the software onto their systems and running it. In addition to damaging files or programs, Trojan horses can steal data, activate or spread other malware, and create secret entrances known as backdoors that allow attackers access to a system without going through the system's normal authentication procedures.

More recent types of malware include ransomware, web-based malware and logic bombs.

Ransomware is malicious software that blocks access to a user's data and threatens to delete the data within a certain period of time unless a ransom is paid. The software blocks access by encrypting the data; if the ransom is paid, the user receives keys for decrypting the data. In theory, these keys would be sufficient to recover the encrypted data; however, in practice, they may only facilitate a partial recovery or no recovery at all, particularly if the perpetrator of the attack does not have expertise in encryption.

Web-based malware, also known as drive-by download, is malicious software that redirects a web browser to an infected website, which then attempts to install tools such as rootkits, keystroke loggers or web browser plug-ins that attackers can use to steal passwords and launch attacks. A **rootkit** is a collection of files or programs that is designed to hide the existence of malware, including the rootkit itself. A **keystroke logger** is a program that surreptitiously records the keys pressed on a keyboard or keypad. A **web browser plug-in** is software that provides additional functionality to a browser; malicious plug-ins can record everything that a browser does and any data entered.

A **logic bomb** is malicious software that remains inactive until a particular logical condition is met. An example is code that maliciously deletes data or files on a particular date such as April 1 or Friday the 13th or when a particular number of transactions has been processed. Logic bombs can be contained in other types of malware or installed directly on a computer system by a disgruntled employee.

Malware can be difficult to detect and even more difficult to remove. It can also cause serious damage to computer systems and, as a result, harm the people and machines that depend on these systems. For example, malware can alter the diagnostic images of CT scans, which could result in incorrect diagnoses and improper treatment. Likewise, it can cause robots in a computerized manufacturing process to work incorrectly, which could result in the production of unsafe goods.

Sources of malware infection

Malware can infect any computing device, whether the device is connected to a network or not. For example, a standalone computer can be easily infected from an infected USB or external hard drive. However, the most common source of infection is email.

Malware frequently enters a device or network through an email that contains an infected attachment or link to an infected website. Recipients of the email are encouraged to open the attachment or follow the link. However, unbeknownst to them, opening the file or clicking on the link triggers the downloading of malware to the recipient's device. Once installed on that device, the malware then propagates to other devices on the network.

More and more users of networked computer systems are being trained to recognize suspicious emails and delete them immediately without opening attachments or clicking on links. However, it only takes one user to open an infected attachment or click on a contaminated link to trigger the installation of malware onto the network. Even users who are normally very security conscious can be fooled from time to time, e.g., if the email is sent to a mobile device and the small size of the device's screen makes it difficult to recognize the email as suspicious. Emails with malware infection are also becoming more difficult to detect. Whereas such emails

used to have a number of telltale signs, e.g., spelling or grammatical errors, corporate logos that are obviously counterfeit, etc., they are now much more professional and can appear to be legitimate.

Potential losses associated with malware infections that damage electronic information

The losses associated with a malware infection that damages, destroys or makes inaccessible the electronic information on a device but does not result in the theft or disclosure of confidential information or involve a malicious actor taking control of the device are similar to the losses associated with a denial-of-service attack, but generally of much greater size. If one attacker has managed to infect a system with malware, others might have as well, so a thorough forensic investigation needs to be conducted.

If a malware infection is localized, then recovery may be straightforward and involve nothing more than removing an infected or damaged data file and replacing it with a backup version that is deemed to be free of infection. However, if the infection is widespread or has impacted critical software such as the operating system then recovery will be much more complicated and take longer. In extreme cases, it may be necessary to wipe a device clean and reinstall its operating system and all the applications and data files that were on the device; if the device in question is a server, this could take a considerable amount of time.

Today's malware is much more damaging and difficult to locate and remove than the malware of the past. Companies that are victims of malware will need to retain the services of computer security and forensics specialists with up-to-date knowledge and technical expertise in the latest types of malware. The costs of these services can be substantial.

Recovery from a cyberattack involving malware generally requires the quarantine of devices, applications or files that are infected or suspected of being infected. A quarantine reduces the chance of infection spreading during the recovery operation but also reduces the computing resources available to support a company's normal operations. As a result, a company may experience an interruption in operations or decline in service standards; in an extreme case, a company may find itself with no computing resources at all and have to suspend operations

entirely. The losses associated with such interruptions will vary from one company to another. If the company is an online retailer, the loss is likely to take the form of foregone sales. On the other hand, if the company is an IT services provider, it is likely to take the form of contractual penalties and liabilities for third party losses as well as lost revenue.

Other losses that a company may suffer as a result of an attack involving malware include:

- Regulatory fines and penalties if, for example, it is determined that the company failed to adequately protect its systems from malware attacks and
- Reputational damage that results in lost customers and lower future revenue and profits.

One potential loss associated with a malware infection that does not arise with a denial-of-service attack is the cost of ransom payments associated with ransomware. Companies that are victims of ransomware may have little choice but to pay the ransom demanded, particularly if the ransomware encrypts all of the company's files and the company has no backup copies. Ransoms can vary from token amounts to very substantial sums of money and are often payable in bitcoin or some other cryptocurrency to maintain the anonymity of the attacker.

As in the case of denial-of-service attacks, there is no single quantity that can be used to measure losses. Some losses are related to the size or complexity of the company's computer system, others to sales, revenue or profit, and still others to a company's systemic importance.

Motives for perpetrating malware infections that damage electronic information

With the exception of ransomware attacks, which are intended to extort money from their victims, there would appear to be little financial incentive for spreading malware infections that damage, destroy or make inaccessible the electronic information on a device but do not steal confidential information or computing resources. This suggests that the primary motive for such attacks is sabotage.

The same groups that perpetrate denial-of-service attacks, namely state-affiliated actors engaged in cyber warfare or cyber terrorism, activists and thrill-seeking hackers, are often

involved in these types of attacks as well. However, in this case, the objective is to do real harm to the victim, not just create mischief. For example, rather than testing a potential adversary's cyber defenses, a state-affiliated actor may wish to cripple or set back a secret research or weapons development program that the actor's sponsor finds alarming. Alternatively, the actor may wish to send a subtle message regarding its ability to destroy or disable critical infrastructure.

Disgruntled employees or employees who fear their jobs are at risk may also perpetrate these types of attacks.

3.1.3. Attacks That Result In the Theft of Nonpublic Information

Although some cyberattacks that penetrate the defenses of a device or network just sabotage the operation of the device or network, many attacks of this type steal confidential information or are part of a program of cyberespionage.

Stolen information can include login credentials, credit card numbers, personal identifying information such as name, address and social security number, purchase histories, email addresses, medical records, trade secrets and the results of proprietary research. Information of this type can be sold to third parties for immediate financial gain or used by the attackers themselves to commit fraud or engage in other illicit activity.

Obtaining unauthorized access to an organization's information system

There are many different ways of obtaining unauthorized access to an organization's information system.

- ***Using stolen login credentials:*** An attacker could use stolen login credentials that the attacker has purchased from a third party. Note that although the login credentials are valid, using them in an unauthorized way is not and hence considered a cyberattack.

- ***Exploiting human vulnerabilities:*** An attacker could use phishing or pretexting to trick an authorized user into providing login credentials for the user or someone else.

Phishing is the act of sending people emails that appear to be legitimate but in reality are not in order to trick them into revealing confidential information about themselves. For example, an attacker could send a person an email with a link to what appears to be the legitimate login webpage for the person's bank but actually is a counterfeit page designed to steal login credentials. **Pretexting** is the act of using (possibly confidential) information about a particular target to establish legitimacy in the mind of the target and induce the target to divulge (even more) confidential information or do something that the target ordinarily would not do. For example, an attacker could steal personal identifying information on the customers of a credit card issuer and then contact each customer to request additional identifying information on the pretext that the customer's account has been frozen and needs to be reactivated.

- ***Exploiting system vulnerabilities:*** An attacker could use known technical deficiencies in public networks or operating systems to gain access to an organization's information system. This task is made easier if the organization's system administrators do not keep abreast of the latest network or operating system vulnerabilities and/or are lax in addressing them. Once inside the system, an attacker can create bogus login credentials to facilitate continuous, unrestricted access to the system; however, more often an attacker will just make a copy of all the existing access credentials and use ones that are least likely to arouse suspicion.

- ***Installing malware:*** By exploiting human and system vulnerabilities, an attacker can install malware such as keystroke loggers or malicious web browser plug-ins onto an organization's information system. Malware of this type allows the attacker to steal login credentials and gain access to the system.

Once access to an organization's system has been established, an attacker can unobtrusively map the network, look for open ports, locate sensitive information, and surreptitiously make copies of data that the attacker believes to be valuable.

Potential adverse consequences of the theft of nonpublic information

There are many potential adverse consequences associated with the theft of nonpublic information.

- ***Impersonation:*** Criminals could use identifying personal information such as usernames and passwords, payment card numbers or the numbers of insurance policies to impersonate an individual or group of individuals. For example, a criminal could use payment card information to make fraudulent purchases or health insurance information to fraudulently bill an insurer for healthcare services or obtain prescription drugs for resale on the black market. A criminal could also steal an individual's identity outright and use it to set up fraudulent credit accounts, obtain fraudulent passports or commit other crimes in the individual's name.

- ***Surveillance:*** Criminals could use information such as an individual's mobile phone location or electricity usage to determine where the individual is and what the individual is doing at a given time, thereby gaining insights into the individual's vulnerabilities, e.g., when the individual is likely to be away from home.

- ***Profiling:*** Criminals could use information such as the customer lists of home security companies or investment brokerages that cater to high net worth individuals to develop profiles of desirable targets and gain insights into their vulnerabilities. For example, a criminal could contact a customer of a home security company claiming to be a service technician and rewire the customer's system to facilitate a break-in without detection at a later date.

- ***Disclosure of sensitive or embarrassing personal information:*** Sensitive medical information such as treatment for substance abuse or a rare disease that an individual does not want a prospective employer to know about may become public. Likewise, embarrassing personal information such as the purchase of controlled substances or dating services like Ashley Madison could become public.

- ***Extortion:*** A criminal could threaten to disclose sensitive or embarrassing personal information or post the credit card information of a company's customers on the web unless a specified amount is paid by a certain date.

- ***Loss of proprietary information, trade secrets or other information of value:*** Sensitive corporate information such as business plans, proprietary research and trade secrets that would be valuable in the hands of a competitor could be made public.

Potential losses associated with an attack that results in the theft of nonpublic information

The losses associated with a cyberattack that involves the theft of nonpublic information are generally much greater than the losses for an attack that only seeks to do damage to an organization's information system. In addition to the usual investigation and recovery costs, companies can face significant third-party liability costs and regulatory fines related to the actual or potential disclosure of nonpublic information.

The first step in responding to an attack that involves the theft of nonpublic information is to determine the extent of the theft and take immediate action to prevent other nonpublic information from being compromised. This usually requires the technical expertise of cybersecurity specialists. A company may become aware of a theft through routine system monitoring, in much the same way a person becomes aware of a house break-in by noticing that an outside door has been forced open or items inside are not in their usual place. However, more likely, it will become aware of the theft when it receives an extortion letter or the stolen information appears on some publicly accessible website. Either way, the company will need to retain the services of experts.

Once the extent of a theft is determined, the company will need to notify the individuals whose confidential information has been or could have been stolen so that they can take prompt action to protect themselves, e.g., by changing login passwords, etc. This is not just good business practice; it is also required under many privacy and data protection statutes and

regulations such as the European Union's General Data Protection Regulation (GDPR), which came into effect in 2018. If the theft involves identifying information such as birth dates or social security numbers that criminals could use to set up fraudulent accounts, the company may also provide free credit monitoring services to those impacted for a specified period of time both as a goodwill gesture and to limit future liabilities. The cost of notification and providing free credit monitoring will depend on the number of individuals impacted; all else being equal, one would expect these costs to be greater for companies with a greater number of customer, patient or employee files.

Companies that have nonpublic information stolen from them may be subject to regulatory fines or class-action lawsuits. The amount of fines levied or damages awarded will depend on a number of factors including the harm caused, the number of individuals impacted, what the company did to limit the harm, the company's size and importance, the extent to which the company took preventative measures to protect against the theft and disclosure of confidential information, and the additional measures the company put in place to prevent a similar theft from happening again. Regulatory fines are often based on a company's revenues while damages associated with class-action lawsuits take into account the number of members in the class. Regulatory fines can also depend on the regulatory framework, e.g., financial services versus healthcare, the number of individuals impacted and the geographic location of the impacted individuals.

In cases involving identity theft, companies will need to meticulously review their records to ensure that any inaccurate information is removed from the impacted records. For example, if the patient records of a health maintenance organization (HMO) are compromised, the HMO will need to ensure that any information related to treatments or prescriptions obtained by persons other than the patient is removed from the patient's record to prevent that information from being inappropriately used in future diagnoses and treatments. Cleaning up records that have been compromised can be time consuming and costly.

Companies that are victims of an attack that steals confidential information will need to assess the damage done to their computer systems and ensure that any malware is removed. As with any attack involving malware, the costs of recovery can be substantial and may depend on the size and complexity of the company's computer system. There could also be costs associated with business interruption if some or all of the company's computing resources need to be taken offline during the recovery operation.

As with the earlier types of attack, there is no single quantity that can be used to measure losses. Some losses are related to the number of patient, customer or employee records, others to the size or complexity of the company's computer system, and still others to the company's sales, revenue or profit.

Note that many attacks involving the theft of confidential information are not discovered for months or even years, particularly if the attacker is skilled or engaged in cyberespionage where the objective is usually to extract as much confidential information for as long as possible. All else being equal, one would expect losses to be greater the longer the discovery period.

3.1.4. Attacks That Commandeer a System's Resources for Malicious Purposes

The final type of cyberattack is an attack that uses malware to steal the resources of a device or network or take control of it for some malicious purpose. Such attacks are known as **command-and-control attacks** or simply C2 attacks.

Examples of command-and-control attacks

There are many different examples of command-and-control attacks.

An attacker could surreptitiously install software on a device for the purpose of mining cryptocurrency such as bitcoin. Mining cryptocurrency involves finding solutions to particular complex mathematical equations; once a solution is found, an amount of cryptocurrency is awarded. The computing resources needed to mine for cryptocurrency can be very significant, in fact so significant that the cost of the electricity used to power the computers needs to be taken into account when calculating how much profit was made from a particular mining

operation. If a crypto miner is able to steal these resources, then the cryptocurrency ultimately mined will be pure profit.

On the other hand, an attacker could use the resources of a device to launch an attack on another device or network. One way of doing this is to install a malicious bot on the device. A **bot** is just an automated process (bot is short for robot). A **malicious bot** is a bot that is connected to and controlled by another computer for some malicious purpose. A group of malicious bots controlled by the same computer is known as a **botnet**. A botnet could be used to launch any number of different attacks, the most common being a denial-of-service attack. More advanced botnets can commandeer the resources of internet-of-things devices such as Wi-Fi enabled electronics or appliances to magnify the impact of an attack.

An attacker could also use a device to hide its identity when attacking another device or network. Indeed, cybercriminals typically go through dozens of networks before hitting their intended targets in order to cover their tracks and make it difficult for investigators to determine the real perpetrators of a crime.

Finally, an attacker could commandeer the resources of a device or network to intentionally cause harm to the device or network and/or the people that use and rely on it. For example, an attacker could take control of the power grid and cause random blackouts.

From these examples, one sees that there can be a variety of motives for launching a command-and-control attack, from financial gain in the case of malicious crypto miners to camouflaging the identity of the attacker in the case of cybercriminals or sabotaging the target computer system in the case of cyberterrorists. However, theft and disclosure of confidential information is not usually the primary objective.

Potential losses associated with command-and-control attacks

Losses associated with attacks that use the resources of the victim's computer for the financial benefit of the attacker, e.g., crypto mining, are usually limited to the costs of investigation and recovery, e.g., locating and removing crypto mining malware from the

victim's computer. By contrast, losses associated with attacks that use the victim's computer to launch an attack against another computer are more substantial and include third party liability claims as well as first party investigation and recovery costs.

3.2. Exposure Associated with Other Cyber Events

A cyberattack is not the only type of cyber event that can give rise to cyber exposure. The availability, integrity or confidentiality of an organization's information system or electronic data can also be compromised by human error, the misuse of privileges by an authorized user, or the manipulation of human behavior, which in cybersecurity contexts is known as social engineering.

3.2.1. Human Error

There are a number of ways in which human error can give rise to a cyber event.

- ***Erroneous delivery of communication:*** An email containing confidential information could unintentionally be sent to the wrong recipients.

- ***Publishing error:*** Confidential information could unintentionally be posted to a website that is accessible to the public.

- ***Misconfiguration of hardware:*** The settings on a remote access server could unintentionally be configured in a way that allows unwanted guests into the organization's computer system.

- ***Erroneous installation of software:*** Software or data on a client's computer could unintentionally be damaged or a backdoor unintentionally created and left open when a consultant installs a new application or updates an existing one.

- ***Failure to address known vulnerabilities:*** Systems administrators could fail to keep abreast of the latest cybersecurity threats or install patches for known vulnerabilities in a timely fashion.

Losses associated with an unintentional disclosure of confidential information, e.g., a publishing error or an erroneous delivery of communication, will be similar to losses when the disclosure is the result of a cyberattack. An investigation will have to be conducted to determine the source and extent of the disclosure; impacted individuals will have to be notified and provided with credit monitoring services as necessary; regulatory fines and penalties will have to be paid; and third party claims will have to be settled. The only losses that do not arise when the disclosure is unintentional are the costs of removing malware and any associated service interruptions.

Losses associated with a misconfiguration of hardware or an erroneous installation of software depend on whether a vulnerability was created and an attacker was able to exploit this vulnerability. Losses will also depend on whether the computer system that is damaged belongs to the organization or a client with which the organization has a business relationship. If no security vulnerability is created then losses will be limited to fixing the damage done to the hardware or software, unless the damage is done to a client's computer in which case the organization will also be responsible for any client losses such as business interruption costs that can be attributed to the organization's errors. On the other hand, if a security vulnerability is created, losses will be similar to those of a cyberattack, particularly if the vulnerability is exploited by an attacker.

3.2.2. Misuse of Privileges

Authorized users of an organization's information system can misuse their privileges in a number of ways.

- ***Privilege abuse:*** An employee of a hospital, out of personal curiosity, could access the medical records of a well-known public figure being treated at the hospital. Similarly, an employee of the tax department who holds a grudge against a particular individual or organization could leak embarrassing financial information about the individual or organization to the media.

- ***Mishandling of data for financial gain:*** Employees who leave a company to start their own business or work for a competitor could take copies of client lists and other valuable company information with them. This is known as **insider threat**.

- ***Unapproved workarounds:*** Employees could install unapproved software on their desktops, laptops or mobile devices or make unauthorized modifications to the installations of existing software, which could create security risks in a number of ways. For example, a backdoor could inadvertently be created or a known vulnerability in the installed or modified software could go unpatched because the systems administrators were unaware that the software was installed on the system.

Any such misuse of privileges is considered to be a cyber event. Note that misuse of privileges is distinct from human error because the actions of the authorized user in this case are intentional.

2.2.3. Social Engineering

Social engineering is the manipulation of human behavior for a malicious purpose. As noted earlier, phishing or pretexting can be used to trick legitimate users of an organization's computer system into providing their login credentials or the login credentials of others to would-be attackers who then use them to launch a cyberattack against the organization. However, techniques such as phishing or pretexting can also be used to trick users into doing things that they otherwise might not without these actions resulting in a cyberattack.

For example, a third party could send an employee in the organization's accounts payable department an email with an attached invoice that appears to be from a legitimate service provider. If the employee is tricked into sending the requested payment to the third party, the event is a cyber event. Note that this event is considered a cyber event because the compromise involves the organization's computer system but it is not considered a cyberattack because there is no damage, destruction or disruption of the organization's computer system and no theft or erasure of information on the system. If the fraud attempt had occurred through regular postal

mail, it would still be considered social engineering, but it would not be considered a cyber event because the organization's computer system is no longer directly involved in the fraud.

2.3. Silent Cyber Risk

The potential for a cyber event to trigger losses on insurance policies where coverage was not intended is known as **silent cyber risk**. Silent cyber risk is an increasingly important issue for writers of traditional property and casualty insurance policies.

As an example, consider a cyberattack that shuts down a state's power grid for an extended period of time. Such an attack would interrupt or shut down the operations of virtually every business or organization that depends on electricity, from grocery stores, general merchandisers and banks to schools, hospitals and factories. Although some businesses and organizations might have backup power generators, most would not, and even the ones with backup generators would only be able to supply power to their most critical functions given the limited capacity of most backup generators. Indeed, backup generators still require fuel – usually gasoline or diesel – to operate, which during a prolonged power outage would likely be in short supply. Hence, even businesses and organizations with backup generators would eventually find themselves without power. Losses from a prolonged interruption of operations would include lost sales or production and damaged or unsaleable goods, e.g., spoiled food.

If a shutdown of the power grid occurred at the same time as a period of heavy rain, homeowners and businesses in low-lying areas or areas prone to flooding could suffer significant property damage. Without power, basement sump pumps and other flood-control devices would eventually fail, even if they were equipped with a battery backup, and the basements of properties in flood-prone areas would become flooded. Moreover, if the properties happened to be located near a river and the shutdown of the power grid impaired the operation of dams further upstream, the river could overflow its banks, which could lead to widespread evacuation of the area and even greater damage to properties.

If the shutdown of the power grid lasted long enough, it could impair the ability of municipalities to provide their populations with clean water. Although municipalities typically

have several days of supply of treated water stored in water towers, this supply would be quickly depleted – and with it the water pressure needed to keep the system functioning – during a prolonged shutdown of the power grid. The reason is that power is needed to pump treated water from surface-level treatment plans to water reservoirs at the top of the water towers. Without a reliable source of clean water, the population would be at risk of contracting water-borne diseases requiring medical attention. Even if the health of the population were not at risk, a drop in water pressure could still make it difficult to fight a major fire during the power outage, with the result that fire-related property damage could be much greater than normal.

Finally, a prolonged shutdown of the power grid could lead to civil unrest and looting, particularly if people start to run out of food or other essential supplies and believe that authorities have lost control of the situation. In this case, losses could be conceivably be similar to those associated with a hurricane or natural disaster.

Note that for each of the scenarios described, the proximate cause of the loss is something other than the cyberattack that shuts down the power grid. In the case of business interruption, it is the failure of the power system; in the case of water-related property damage, it is the heavy rain and flooding; in the case of water-borne illness, it is the failure of the water treatment system; and in the case of fire-related property damage that is made worse by inadequate water pressure, it is the fire. This has important implications for traditional business interruption, property and health insurance policies. Depending on how such traditional policies are worded, coverage could be triggered even though it was not intended.

The focus of this report is on cyber insurance policies. Such policies are intended to cover cyber risk, so the issue of silent cyber risk does not really arise. However, the observations and conclusions in this report are likely to be relevant for policies where silent cyber risk is an issue.

4. Cyber Insurance Coverage

The preceding discussion on the nature of cyber risk exposure identified a number of types of losses that can occur as a result of a cyber event. However, not all losses that are attributable to a cyber event are covered by cyber insurance. This section of the report provides a summary of the cyber insurance coverages that are currently available in the market based on a review of sales literature and sample policy forms for the major carriers of cyber insurance in the U.S. and Europe. The focus in this report is on policies where coverage of cyber risk is intended. Hence, policies where coverage of cyber risk is not intended, e.g., traditional property and casualty insurance policies with silent cyber risk, are not considered.

Although cyber insurance has been available in various forms since the 1990s, it is still a relatively new product and continues to evolve. As a result, coverages and policy forms vary by insurer as well as geography. Some insurers offer modular policies in which customers can pick and choose from among the available cyber coverages and build a policy to suit their needs while others offer package policies in which preselected cyber coverages are bundled together by the insurer. Most insurers offer a core set of cyber coverages, either as a package or part of a modular policy, as well as a number of supplementary coverages, which can vary from one insurer to another.

Coverage particulars, e.g., what is considered to be a covered event, can vary from one insurer to another. Some insurers consider the unauthorized use of login credentials to be a covered event even if the login credentials are obtained by non-electronic means, e.g., by reading them from a piece of paper taped to an employee's computer, while other insurers limit coverage to situations where the stolen credentials are obtained electronically. To avoid confusion with our definition of the term cyber event, we will use the generic term **cybersecurity event** in the coverage descriptions that follow to refer to a cyber event that is covered under a particular cyber insurance policy.

3.1. Core Coverages

The core coverages offered by most cyber insurance carriers are privacy liability, network security liability, cyber event response, network interruption, recovery and restoration of digital assets, regulatory actions, fines and penalties, and payment card industry assessments. Note that core coverages typically exclude any losses related to bodily injury or property damage.

3.1.1. Privacy Liability

Provides coverage for damages and defense costs associated with third party claims that arise from a failure to protect private or confidential information in an insured's possession. Also provides coverage for any violation of privacy statutes alleged in connection with such claims.

In this context, private or confidential information is generally considered to be:

- Information that can be used to uniquely identify an individual, e.g., name, address, telephone number, social security number, account numbers, passwords, account histories, etc.
- Information that is considered nonpublic personal information or personal health information by statutes governing the protection of such information.
- Information that the insured uses to authenticate its customers or clients.
- Trade secrets or similar information of a third party that is not available to the general public.

For clarity in the discussion that follows, a failure to protect private or confidential information in an insured's possession will be referred to as a **privacy event**.

Note that privacy liability coverage is not restricted to disclosures of nonpublic information that are the result of a cyberattack. Disclosures caused by human error or negligence are also covered. This makes sense because most information today is stored in digital form and so determining whether a particular disclosure of confidential information is the result of a cyberattack can be challenging. Moreover, losses depend on the type and amount of private

information that is disclosed and how widely available that information becomes, not whether the cause of the disclosure is a cyberattack. For the same reason, privacy liability coverage is not restricted to information stored in electronic form; private or confidential information in paper or some other non-electronic form is also covered.

3.1.2. Network Security Liability

Provides coverage for damages and defense costs associated with third party claims that arise from a failure to prevent a cybersecurity event from impacting the computer systems or networks of others.

There are many ways in which a cybersecurity event could impact other computer systems and networks. For example, malware installed on the insured's computer could launch a denial-of-service attack on another computer or be part of a botnet that launches a coordinated attack. Alternatively, malware on the insured's email server could send email infected with worms or Trojan horses to all outside email addresses that reside on the server.

Coverage does not include losses related to bodily injury or property damage. For example, it does not include bodily injury to a passenger in a car that crashes as a result of malware that originates from the insured's computer and infects the car's onboard computer. However, some insurers may offer a special endorsement to extend coverage to third party bodily injury and property damage claims that are attributable to a cybersecurity event. See section 3.2 Supplementary Coverages for details.

3.1.3. Cyber Event Response

Provides coverage for costs incurred by the insured to:

- Conduct an investigation to determine the cause of a cybersecurity or privacy event and provide an initial assessment of the damage.
- Obtain legal advice on the insured's obligations.
- Retain the services of public relations and/or crisis management professionals to develop a plan for maintaining or restoring public confidence in the insured.

- Notify individuals or other third parties whose private or confidential information has or may have been compromised as a result of a cybersecurity or privacy event and provide advice on available remedies, e.g., promptly changing passwords, etc.
- Provide credit monitoring or similar identity theft prevention and recovery services to the individuals notified, as appropriate.

Note that coverage is provided whether the costs arise as a result of a cybersecurity event or a privacy event. Although the costs of investigating a cybersecurity event are likely to be different from the costs of investigating a privacy event that does not also involve a compromise of the organization's information system, the costs of obtaining legal services, retaining the services of public relations or crisis management professionals, and notifying and if necessary providing credit monitoring services to the impacted individuals are likely to be the same in either case.

As noted in the discussion of cyber risk exposure, the cost of investigating a cybersecurity event generally varies with the size and complexity of the system while the cost of notifying and providing credit monitoring services to impacted individuals generally varies with the number of patient, customer or employee records in the insured's possession.

3.1.4. Network Interruption

Provides coverage for losses associated with an interruption of an insured's business operations as a result of a cybersecurity event. Covered losses include net income (net profit or loss before income tax) that the insured would otherwise have earned and the insured's normal operating expenses.

Note that coverage is only provided for interruptions that are the result of a cybersecurity event. Interruptions that are the result of physical damage to the insured's computers are not covered.

3.1.5. Recovery and Restoration of Digital Assets

Provides coverage for the cost of removing malware from the insured's computer or network and repairing, restoring or replacing lost or damaged software or data on the insured's computer or network as a result of a cybersecurity event.

4.1.6 Regulatory Actions, Fines and Penalties

Provides coverage for defense costs associated with regulatory actions and any civil fines or penalties that arise from these actions unless the particular fine or penalty is considered uninsurable in the jurisdiction where it is imposed. For the purposes of defining coverage, regulatory action means a request for information made by a government agency or a civil investigation or proceeding launched by a government agency as a result of a cybersecurity or privacy event; regulatory action does not include criminal investigations or proceedings.

4.1.7. Payment Card Industry Assessments

Provides coverage for contractual fines, penalties or other assessments from a payment card industry association such as Visa, MasterCard or American Express or a bank that processes payment card transactions for failing to comply with generally accepted payment card industry data security standards if that failure resulted in a cybersecurity or privacy event.

3.2. Supplementary Coverages

In addition to the core coverages offered by most cyber insurance carriers, there are a number supplementary coverages offered by some carriers. These coverages, which are typically offered by endorsement or as part of an extended package, include cyber extortion, cybercrime, media content liability, technology errors and omissions, third party bodily injury and property damage, enhanced business interruption, and reputation protection. Note that not all coverages are offered by all insurers and among the insurers offering a specific coverage, coverage particulars can vary.

3.2.1. Cyber Extortion

Provides coverage for the cost of responding to extortion involving the insured's digital assets, computer systems and networks, or confidential information in its possession.

The extortion could take a number of different forms:

- An attacker could install ransomware on the insured's computer, blocking access to the insured's files and threatening to delete them unless a ransom is paid.
- An attacker could steal confidential information from the insured's computer and threaten to release it.
- An attacker could threaten to launch a denial-of-service attack.

Covered costs include:

- The services of a consultant with expertise in cyber extortion to provide advice on the appropriate response and assist in negotiations.
- The services of a cybersecurity expert to conduct a forensic investigation to determine the validity, cause and scope of the particular threat.
- The services of a law firm to provide advice on dealing with law enforcement agencies and regulators.
- The services of a public relations firm to develop a plan for maintaining or restoring public confidence in the insured.
- The cost of ransom payments, and if applicable reward payments, made to resolve the threat, including any assistance procuring bitcoin or another cryptocurrency required for payment.
- Losses associated with an interruption in the insured's business operations that is the result of the particular threat, including any extra expenses incurred to mitigate the impact of such an interruption.

3.2.2. Cybercrime

Provides coverage for losses that result from criminal money transfers involving computer systems or networks.

The criminal money transfer could take a number of different forms:

- ***Computer fraud:*** Third party fraudulently accesses the insured's computer system to steal money, e.g., third party steals a bank customer's login credentials, accesses the customer's account and transfers money to an account that the third party has set up at another financial institution, likely using stolen identification.

- ***Funds transfer fraud:*** Third party tricks the insured's bank into transferring funds from the insured's account, e.g., third party tricks a bank's voice recognition software into thinking the third party is the insured and then transfers money to an account the third party has set up at another financial institution.

- ***Social engineering fraud:*** Third party tricks one of the insured's employees into transferring money from the insured's account to the third party, e.g., third party sends an employee who works in the insured's accounts payable department an email with an attached invoice that appears to be from a legitimate service provider and the employee sends payment to the third party.

Coverage may also be provided for losses that an insured incurs as a result of telephone toll fraud, e.g., phone charges for calls fraudulently routed through the insured's telephone system.

3.2.3. Media Content Liability

Provides coverage for damages and defense costs associated with third party claims that arise from the insured's publication, distribution or broadcast of media content.

Media content generally means information that is published in electronic form, e.g., articles or videos posted to online forums or social media websites, but it may also include information published in printed or other forms.

Coverage varies from one insurer to another but can include protection against some or all of the following claims:

- Infringement of copyright, trademark, domain name or similar protected name or mark.

- Plagiarism, piracy or misappropriation of ideas.
- Invasion, infringement or interference with rights of privacy or publicity.
- Eavesdropping, false arrest or malicious prosecution.
- Defamation, libel or slander.

Strictly speaking, media content liability is not a cyber risk or privacy risk exposure. However, most insurers include it in their cyber insurance offerings, either as a core or additional coverage. Consequently, we have included it in our summary of cyber coverages as well.

4.2.4 Technology Errors and Omissions

Provides coverage for damages and defense costs associated with third party claims that arise from the insured's products and services, e.g., software on a client's computer is unintentionally corrupted when a consultant installs a new application or modifies an existing one.

3.2.5. Third-Party Bodily Injury and Property Damage

Provides coverage for damages and defense costs associated with third party bodily injury or property damage claims that arise from

- A failure of an insured to protect private or confidential information in the insured's possession,
- A failure of an insured to prevent a cybersecurity event from impacting the computer systems or networks of others, or
- A cybersecurity event associated with a computer system that is part of an insured's product.

Note that this expands the core coverages of privacy liability and network security liability discussed earlier to include third party bodily injury and property damage claims. It further expands coverage to include cybersecurity events associated with the computer systems that are part of an insured's products. Hence, bodily injury claims made against a car manufacturer that

arise from a car accident caused by malware in the car's onboard computer system would be covered under this provision.

3.2.6. Enhanced Business Interruption

Provides coverage for losses associated with an interruption of an insured's business operations as a result of physical damage to the insured's property due to a cybersecurity event.

This extends the core network interruption coverage to cover first party property damage.

3.2.7. Reputation Protection

Provides coverage for costs incurred by the insured to

- Protect the insured's reputation against a potential threat or
- Respond to an attack on the insured's reputation that has already occurred.

In this context, a threat to the insured's reputation is an action or event that, if publicly disclosed, would be viewed as a material breach of trust by stakeholders and have an adverse impact on the public perception of the insured, while an attack on the insured's reputation is a third party's publication of such information.

Strictly speaking, reputation protection is not a cyber risk or privacy risk exposure. However, many insurers include it in their cyber insurance offerings.

5. Cyber Event Experience

A cursory glance at any of the cybersecurity threat maps available online[1] suggests that thousands of cyberattacks are attempted every minute. Of course, very few of these attempts ever result in a successful cyberattack, i.e., the penetration of a computer system's defenses or launch of a denial-of-service attack. For obvious reasons, cyber security firms prefer to emphasize the number of attacks that their software thwarted rather than the number that got

[1] For example, https://threatmap.checkpoint.com, https://threatmap.fortiguard.com or https://cybermap.kaspersky.com/stats

through. However, if even a tiny fraction of these attempted attacks are successful, that still represents a lot of cyber event experience. Hence, in theory, there should be a meaningful amount of cyber event data available to use when identifying possible exposure measures for cyber insurance.

Unfortunately, the cyber event data that is available in practice is not very good. There are several reasons for this:

- ***Not all events are reported:*** Unless there is a legal requirement to do so, many companies do not report when they are victims of a cyberattack.

- ***Data on reported events is often incomplete:*** Companies that do disclose the occurrence of a cyber event may only provide a minimal amount of information. In particular, important information such as the cause of the event or the ultimate size of the loss may be missing.

- ***Recording of events is inconsistent:*** Companies may capture and record cyber event data in different ways or may not have a consistent framework for recording cyber events at all, with the result that the same event may be recorded differently by different companies or people.

- ***Not all events are known due to the lag between event occurrence and discovery:*** It can take months or even years for some cyber events to be discovered, so by their very nature, many cyber events go unreported.

In addition, the continuing evolution of the threat landscape, due in part to the success of cybersecurity countermeasures, means that the cyber event data that does exist can quickly become out of date. Consider the case of payment card skimmers. A payment card skimmer is a device that is designed to capture the information on the magnetic strip of a payment card and copy it to a counterfeit card without the cardholder's knowledge; the skimmer could be

attached to the card reader of an automated teller machine or payment card terminal or be a standalone device that a compromised store clerk swipes separately when a customer makes a purchase. At one time, payment card skimmers were one of the most popular means of stealing a customer's payment card credentials. However, the introduction of chip technology put a significant dent into this type of fraud. The reason is that cards with an embedded computer chip generate a unique verification code every time they are used, so knowledge of the information on a card's magnetic strip is insufficient to steal the card's credentials; to fraudulently access a cardholder's account, a criminal needs to be in possession of the actual chip-embedded card.

4.1. The VERIS Framework

To address the problem of inconsistency in the recording of cybersecurity events and facilitate research in cybersecurity, Verizon developed the VERIS framework. VERIS, which is short for the Vocabulary for Event Recording and Incident Sharing, is a framework for describing cybersecurity events in a structured and consistent way.[2]

4.1.1. Description of the Framework[3]

Within the VERIS framework, events are recorded based on four descriptors: actor, action, asset and attribute.

Actor

An **actor** is an agent or entity that causes or contributes to a cyber event. There are three primary categories of actor: external, internal and partner.

- An **external** actor is one from outside the organization or its network of partners. Examples include criminal groups, hackers, former employees and state-affiliated entities. Outsiders have no access privileges with the organization.

[2] When it was first released in 2010, VERIS was branded as the Verizon Enterprise Risk and Incident Sharing framework. However, to encourage wider adoption of the framework, Verizon rebranded it as the Vocabulary for Event Recording and Incident Sharing.

[3] A complete description of the VERIS framework is available at http://veriscommunity.net.

- An **internal** actor is one from inside the organization. Examples include full-time employees, independent contractors and interns. Insiders generally have access privileges with the organization, although the level of privilege can vary from one insider to another.

- A **partner** actor is a third party that has some business relationship with the organization. Examples include suppliers, IT service providers and outsourced IT support. Partners usually have some access privileges with the organization, but the level is typically lower than for insiders.

More than one type of actor can be involved in the same event. For example, an external actor could send an email with an infected attachment to an employee of the organization and the employee (an internal actor) could open the infected attachment.

Action

An **action** describes what an actor did to cause or contribute to the cyber event. There are seven primary categories of action: hacking, misuse of privileges, social engineering, error, malware, physical attack and environmental.

- **Hacking** is a malicious action in which an external actor, i.e., an actor that has not been granted access privileges, intentionally attempts to access or do harm to an information asset by circumventing the security mechanisms in place. The information asset could be computer hardware or software, paper documents or electronic media, a telephone system, or something else.

- **Misuse of privileges** is an action in which an internal or partner actor, i.e., an actor that has been granted some degree of access privileges, uses the privileges granted in a way that is contrary to what was intended. Examples include inappropriately accessing a patient's medical records, using proprietary company information for personal gain, and installing unapproved software on a company computer.

- **Social engineering** is a malicious action in which an external actor attempts to manipulate the behavior of a human (usually internal or partner) actor in a way that is advantageous to the external actor but harmful to the human actor or the organization. Examples include phishing, pretexting and bribery.

- **Error** is an action in which an internal or partner actor unintentionally does something incorrectly or inadvertently, or fails to do something that should have been done. Examples include sending email to the wrong recipients, posting private information to a public website, and misconfiguring a remote access server to allow unwanted guests into the company's computer system.

- **Malware** is a category of action that encompasses anything involving the installation and/or use of malicious software (malware). Examples include the installation and use of command and control (C2) malware, backdoors or ransomware.

- **Physical attack** is a category of action in which a human actor intentionally does physical harm to an asset or makes modifications to the environment in which an asset is located that are harmful to the organization. Examples include sabotage, wiretapping, surveillance and assault.

- **Environmental** is a category of action that includes natural disasters, power failures, pipe leaks and other so-called "acts of God."

Every cyber event has at least one action associated with it, but most events have multiple actions, often across multiple categories. An action can be either causal or contributory.

Asset

An **asset** is a person or thing that is impacted or compromised by an action. There are five primary categories of asset: server, network device, end-user device, media and people. Most

categories are self-explanatory. Note that the category end-user device includes desktops, laptops, mobile phones, telephones and peripheral devices such as printers, the category media includes paper documents and electronic storage media, and the category people includes employees and customers.

Attribute

An **attribute** describes how an asset is impacted or compromised by the cyber security event. There are three primary categories of security attribute: availability, integrity and confidentiality.

- **Availability** describes the extent to which the system's resources are disrupted by the cyber event. Examples of compromises in availability include the loss or destruction of an asset and the interruption or degradation of system performance.

- **Integrity** describes the extent to which there has been an unauthorized change to the system and the nature of that change. Examples of compromises in integrity include creating new user accounts without authorization, modifying privileges, permissions or hardware/software configurations without authorization, making unauthorized changes to software, stored data or content, installing unauthorized software, and initiating fraudulent transactions.

- **Confidentiality** describes the extent to which confidential information is disclosed as a result of the cyber event. Examples of compromises in confidentiality include the disclosure of login credentials, bank account or payment card information, medical records, personal or identifying information, and trade secrets.

Note that a cyber event can have more than one attribute associated with it.

VERIS A^4 Grid

The combination of actor, action, asset and attribute descriptors forms a grid known as the **VERIS A^4 grid**. The cells of this grid represent associations within a cyber event, not direct

connections between descriptors. For example, the cell consisting of the actor external, the action malware, the asset server and the attribute confidentiality represents a cyber event that includes an external actor, a malware action, a server asset and a confidentiality attribute at some point in the event sequence. It does not mean that an external actor installed malware that compromised the confidentiality of the server. Likewise, a cell with the action malware and the asset people does not represent the (nonsensical) cyber event in which malware is installed on person; it simply means that both malware and a person were involved in the cyber event.

4.1.2. VERIS Community Database

To facilitate research within the cybersecurity community and support corporate decision making in the area of cybersecurity, Verizon established the VERIS community database in 2013. This database, which is coded in VERIS format, enables organizations to publicly share information on cybersecurity events that have occurred without compromising the security of their information systems or the confidentiality of individuals.[4]

The VERIS community database is still relatively small – there are fewer than 8,000 reported events for the period 2010-2018 and most of these correspond to the years 2012 and 2013 – and is skewed toward events in the health sector. The reason is that most of the data for the period 2010-2018 is from the U.S. Department of Health and Human Services and the websites of various state Attorneys General that report unauthorized disclosures of confidential information. Hence, care must be exercised when making inferences using this database. As more organizations contribute to the database, its scope and comprehensiveness should improve.

4.2. Verizon Data Breach Investigations Reports

In addition to developing the VERIS framework and launching the VERIS community database, Verizon has published the Verizon Data Breach Investigations Report each year since 2008. This report, which provides an accessible and data-driven view of current trends in cybersecurity, has become an invaluable resource for cybersecurity risk managers and others

[4] A link to the database is available at http://www.veriscommunity.net/vcdb.html.

seeking to protect their information systems against cyberattacks, and is considered the most comprehensive and authoritative report of its kind in the field.[5] As of 2018, the database on which the reports are based had accumulated information on more than 330,000 incidents and 16,000 breaches. This is more than 10 times the amount of experience accumulated in the VERIS community database.

The Verizon reports include information on incident and breach patterns across industries, and hence provide an indication of the types of cyber coverages that are likely to have been triggered in the past. From the perspective of identifying possible exposure measures for cyber insurance, this is useful information because it allows one to focus attention on the cyber coverages more likely to give rise to losses.

4.2.1. Incident versus Breach

The Verizon reports define an **incident** to be "a security event that compromises the integrity, confidentiality or availability of an information asset" and a **breach** to be "an incident that results in a confirmed disclosure – not just a potential exposure – of data to an unauthorized party."

Verizon's definition of incident is similar to the definition of cyber event used in this report but also includes events described earlier as privacy events, i.e., events that represent a failure to protect the private or confidential information of others that is in one's possession. The reason is that in the VERIS framework, "people" is a possible category of asset; hence Verizon's definition of incident includes confidential information in paper as well as electronic form. On the other hand, Verizon's definition of breach is more restrictive than what we have described as a privacy event because it only includes events where there is a confirmed disclosure of data.

The terms incident and breach can mean different things in different contexts. In some insurance policies, a breach is what Verizon defines as an incident, i.e., it is a cyber event or a privacy event, while in other policies, it is just a privacy event. Likewise, in some research

[5] The report for the most recent year, as well as some previous years, is available at https://enterprise.verizon.com/resources/dbir/.

reports, a breach includes events in which confidential information is potentially put at risk whether or not there is a confirmed disclosure, whereas in other reports, the definition of breach is more restricted. Consequently, care must be exercised when comparing the findings of different reports and working with experience data.

4.2.2. Data Sources for the Verizon Reports

The incident and breach data underlying the Verizon reports comes from three main sources:

- Publicly available information on security incidents, e.g., information in the VERIS community database
- Information on security incidents investigated by the Verizon Threat Research Advisory Center (VTRAC)
- Security-incident information provided by one of Verizon's external collaborators.

For the 2019 report, over 70 organizations contributed security-incident information. These organizations included Cisco Security Services, Checkpoint Software Technologies, Kaspersky Lab, Fortinet, McAfee, Palo Alto Networks, the Center for Internet Security, the U.S. Computer Emergency Readiness Team (US-CERT), CERT European Union, and Chubb. Unfortunately, most of this data is considered proprietary. Hence, unlike the VERIS community database, the data underlying the Verizon reports is not publicly available.

4.2.3. Dataset Underlying the 2019 Report

To be included in the dataset used for the 2019 report, an incident must have occurred between November 1, 2017 and October 31, 2018 and must pass a number of quality screens, e.g., a sufficient number of fields must be specified in the incident record in order to provide a meaningful description of the incident. Verizon makes clear that this dataset is likely a biased sample of all the security incidents that occurred across all organizations between November 1, 2017 and October 31, 2018 and notes that the degree of bias cannot be measured since there is no way of knowing what proportion of incidents are represented in the sample.

The dataset underlying the 2019 report contains over 100,000 incident records. Of these, over 50,000 correspond to botnet-related breaches, in particular, breaches where customers are victims of a banking Trojan or other credential-stealing malware; over 80% of these botnet-related breaches are in the finance or insurance industries. Given their large number, Verizon removes these botnet-related breaches from the dataset and considers them separately. Verizon also removes an additional 6,500 or so incident records with limited information. The resulting dataset consists of 41,686 confirmed security incidents, of which 2,013 are confirmed data breaches.

4.2.4. Incident and Breach Patterns in the 2019 Report

To gain a better understanding of emerging trends, Verizon classifies incidents and breaches into nine categories: misuse of privileges, denial-of-service attack, web-application attack, error, cyber espionage, crimeware, lost or stolen assets, payment card skimmers, and point-of-sale intrusions. Verizon has found that over 98% of incidents and 88% of breaches in its database of accumulated experience fall within one of these categories. However, the distributions by category can change from one year to another.

The definitions of these categories are similar to those used in the VERIS framework, but with a few nuances:

- The **misuse of privileges** category consists of all incidents where the primary VERIS action category is "misuse of privileges."
- The **denial-of-service attack** category consists of all incidents that are designed to compromise the availability of a network or system.
- The **web application attack** category consists of all incidents in which a web application is the method of attack.
- The **error** category consists of all incidents in which an unintentional action directly compromises a security attribute of an asset.
- The **cyber espionage** category consists of all incidents involving state-affiliated actors or exhibiting the motive of espionage.

- The **crimeware** category consists of all incidents involving malware that are not included in one of the other categories.
- The **lost or stolen assets** category consists of all incidents where an information asset goes missing.
- The **payment card skimmers** category consists of all incidents in which a skimming device is implanted on an asset that reads magnetic stripe data from a payment card.
- The **point-of-sale intrusions** category consists of all incidents in which there is a remote attack against a terminal or controller used to conduct retail payment card transactions.

Table 1 shows the distribution of incidents by category for a selection of industries based on the dataset underlying the 2019 Verizon report. Table 2 shows the corresponding distributions for breaches.[6] All figures have been rounded to the nearest percentage point.[7] Note that the category labeled "Other" includes incidents in the categories payment card skimmers and point-of-sale intrusions as well as incidents that do not fit into any of the other seven categories.

Table 1: Distribution of Incidents by Verizon Category for Selected Industries (2019)

	Category							
Industry	Misuse of Privileges	Denial of Service Attack	Web Application Attack	Error	Cyber Espionage	Crimeware	Lost & Stolen Assets	Other
Public Admin	56%	4%	0%	6%	1%	20%	12%	1%
Education	5%	59%	8%	10%	2%	8%	2%	6%
Healthcare	23%	1%	15%	22%	1%	16%	13%	9%
Finance	11%	62%	8%	4%	3%	6%	1%	5%
Retail	7%	23%	39%	5%	1%	9%	3%	13%
Manufacturing	10%	46%	11%	4%	5%	16%	1%	7%

[6] These tables were created from the data on page 31 of the 2019 Data Breach Investigations Report.

[7] Hence, an entry of 0% in a particular category does not necessarily mean that there were no incidents or breaches in that category.

Table 2: Distribution of Breaches by Verizon Category for Selected Industries (2019)

	Category							
Industry	Misuse of Privileges	Denial of Service Attack	Web Application Attack	Error	Cyber Espionage	Crimeware	Lost & Stolen Assets	Other
Public Admin	12%	0%	10%	18%	42%	2%	5%	11%
Education	9%	0%	24%	35%	5%	3%	3%	21%
Healthcare	28%	0%	21%	32%	1%	0%	9%	9%
Finance	22%	0%	33%	16%	11%	3%	1%	14%
Retail	10%	0%	62%	8%	1%	2%	2%	15%
Manufacturing	16%	0%	40%	13%	14%	6%	2%	9%

From these tables, it is clear that incident and breach distributions vary widely across industries. Denial-of-service attack incidents are most frequent in the educational services, finance/insurance and manufacturing industries with strong representation in the retail industry as well but are virtually nonexistent in public administration and the healthcare industry. On the other hand, misuse of privileges is by far the most frequent incident in public administration, and error and misuse of privileges combined make up almost 50% of incidents in the healthcare industry. The high frequency of misuse-of-privilege incidents in public administration and the healthcare industry could be due to more stringent reporting requirements in those industries.

Cyber espionage is by far the most frequent breach in public administration – in the words of Verizon's report, "cyber espionage is rampant in the public sector" – and is also material in the finance/insurance and manufacturing industries. In the educational services and healthcare industries, the most frequent type of breach is error with web application attack and/or misuse

of privileges a close second. Web application attack is the most frequent breach for the finance/insurance, retail and manufacturing industries while misuse of privileges and error are also meaningfully represented. Across industries, denial-of-service attacks are very rarely associated with breaches.

The Verizon report provides useful information on incident and breach frequencies but very little information on the resulting losses. Hence, while it provides an indication of the cyber coverages likely to be triggered, it does not tell us which coverages are likely to generate the largest losses.

4.3. Other Sources of Information on Historical Cyber Events

There are a number of other organizations that provide information on historical cyber events. These include the following:

- U.S. Department of Health and Human Services (https://ocrportal.hhs.gov/ocr/breach/breach_report.jsf)
- Privacy Rights Clearinghouse (https://www.privacyrights.org)
- Ponemon Institute (https://www.ponemon.org)
- Advisen (https://www.advisenltd.com/data/cyber-loss-data/)

The U.S. Department of Health and Human Services and the Privacy Rights Clearinghouse provide data on actual or potential compromises of personal information. The data is freely available on their websites. As the focus of these organizations is on public awareness, data on losses tends to be limited.

The Ponemon Institute conducts an annual study on the costs associated with a data breach and publishes its findings in a report entitled *Cost of Data Breach Study*. The study is based on interviews with companies that have experienced a data breach during the past year and uses

an activity-based costing method to estimate losses. These studies provide information on benchmark costs rather than actual historical losses. Past reports can be downloaded for free from the Ponemon Institute's website.

Advisen maintains a proprietary database of cyber loss data that, according to its website, is compiled from multiple publicly available sources, goes through a rigorous quality assurance process and is designed to meet the needs of insurance professionals. We did not have access to Advisen's database during the preparation of this report, so we cannot provide any further details on this potential source of historical loss information.

5. Approach for Identifying Potential Exposure Measures

Given the scarcity of reliable, representative and publicly available loss experience data for cyber insurance, this report uses a conceptual rather than an empirical approach to identify potential exposure measures.

For each cyber insurance coverage, losses that can arise in connection with the coverage are considered and quantities related to these losses identified. These quantities are then analyzed to determine the extent to which they possess the properties of a good exposure measure.

A good exposure measure has the following properties:

- ***Simplicity:*** The exposure measure is easy to calculate.

- ***Auditability:*** The calculation of the exposure measure can be readily audited.

- ***Relationship to losses:*** There is a strong, typically linear, relationship between the exposure measure and losses.

- ***Stability:*** The calculated value of the exposure measure is stable over the period of insurance coverage.

- ***Legally determinable:*** The information required to calculate the exposure measure can be legally obtained and shared with insurers or other third parties without violating any privacy laws or regulations.

Quantities that possess all of these properties are identified as candidate exposure measures.

From these candidate exposure measures, the best exposure measure is selected for each coverage. The recommended exposure measure (or collection of exposure measures) for the cyber insurance policy as a whole is then determined by focusing on the coverages that are expected to generate the largest losses.

6. Candidate Exposure Measures for Cyber Insurance Coverages

This section of the report identifies possible exposure measures for the commercial cyber insurance coverages described earlier using the approach outlined in the preceding section. For each coverage, losses that can arise are highlighted and quantities related to these losses are identified. These quantities are then evaluated against the criteria of simplicity, auditability, strength of relationship to losses, stability and legal determinability. As every policy will have at least one core coverage, we focus on those coverages.

6.1. Candidate Exposure Measures for Core Coverages

Recall that the core coverages offered by most cyber insurance carriers are privacy liability, network security liability, cyber event response, network interruption, recovery and restoration of digital assets, regulatory actions, fines and penalties, and payment card industry assessments. In the discussion that follows, we focus on the first three coverages (privacy liability, network security liability and cyber event response). These coverages are more likely to generate large losses. The exposure measures identified for these coverages also turn out to be reasonable candidates for the others.

6.1.1. Privacy Liability Coverage

Losses that arise under privacy liability coverage consist of:

- Damage and defense costs associated with class-action lawsuits

- Fines levied under applicable privacy statutes.

Damages awarded and/or fines levied depend on a number of factors including the harm caused, the number of individuals impacted, the organization's size and importance, and the extent to which the organization took preventative measures to protect against the unauthorized disclosure of confidential information.

Possible bases for measuring exposure

It is impossible to know the harm caused by an event until after the event has occurred, so basing an exposure measure on harm caused is not feasible. Basing an exposure measure on preventative measures taken is also not feasible but here the difficulty lies with quantifying preventative measures taken and connecting that to losses; an organization could have state-of-the-art protection in place but still suffer large losses as a result of an unauthorized disclosure.

The number of individuals impacted by an event is not known until after the event has occurred but it will be some fraction of the total number of individuals associated with the organization, e.g., as customers, patients, employees, etc., which is something that can be quantified in advance, at least in theory. Hence, basing exposure on class size or a related proxy is a possibility. Likewise, organization size and importance is something that can be quantified in advance, e.g., using annual revenue or sales, so basing exposure on organization size or a related proxy is also a possibility.

Class size versus organization size

Judges determining damage awards in class-action cases are likely to put more weight on the size of the class than the organization to ensure that each member of the class receives a reasonable monetary settlement. This suggests that quantities related to class size are likely to be better measures of exposure for losses related to class-action lawsuits.

On the other hand, authorities responsible for administering fines under privacy statutes are likely to put more weight on the size of the organization to ensure that organizations have appropriate incentives to comply with the privacy protection provisions in these statutes. This

suggests that quantities related to organization size may be better measures of exposure for losses related to fines.

That being said, fines can vary widely from one jurisdiction to another depending on the provisions of the particular statute and range from relatively insignificant amounts to substantial sums. For example, under the U.S. HITECH[8] Act, the maximum fine is $50,000 per violation per day or $1.5 million per year for violations of the same type, whichever is smaller, while under the European Union's GDPR, the maximum fine (for serious violations) is €20 million or 4% of global revenues, whichever is greater. Hence, an organization's size is not necessarily a good proxy for the fines it could face. In fact, in jurisdictions where fines are capped by a fixed monetary amount, organization size may not matter much at all because any fine is likely to be dwarfed by damages awarded under class-action lawsuits, particularly if the class is large. This suggests that for the purpose of identifying possible exposure measures for privacy liability coverage, it may be better to focus on quantities related to class size.

Types of classes

There are many different classes of individuals that could be associated with an organization, e.g., customers, patients, students, employees, contractors, etc. For the purpose of measuring class size, it is convenient to distinguish between two types of classes: those consisting of individuals that provide a product or service to the organization, e.g., employees, contractors, etc. and those consisting of individuals that receive a product or service from the organization, e.g., customers, patients, students, etc. For simplicity, we'll refer to individuals of the first type as employees and individuals of the second type as customers. Note that it is possible for an individual to be both an employee and a customer of an organization; however, for the purpose of determining a simple measure of class size, these overlaps can be ignored, at least for the time being.

[8] HITECH is short for Health Information Technology for Economic and Clinical Health

Classes consisting of employees

In most organizations, employees are assigned a unique employee identification number when they begin their jobs, but even in organizations where this is not the case, e.g., small businesses with only a few staff, employees still have a unique social security or taxpayer identification number, which they are usually required to give to their employer before they can get paid. Hence, determining the number of employees associated with an organization should be a relatively straightforward task; for example, one could count the number of distinct employee files or social security numbers. The only potential challenge is that the calculation must include both active and inactive, retired or terminated employees; the reason is that an organization continues to be liable for the unauthorized disclosure of confidential information it maintains on inactive, retired or terminated employees even after they have left the organization.

Classes consisting of customers

Some organizations such as healthcare providers or educational service providers assign a unique identification number to each of their customers (patients, students, etc.) while others such as retailers or hospitality providers cannot or choose not to do so.

Customers uniquely identified

For organizations that assign a unique identification number to each customer and collate customer information such as patient visits, lab test results, courses taken and course credits earned using this identification number, the number of distinct customer identification numbers is usually a reasonable measure of class size. However, there are circumstances where this is not the case.

Consider, for example, a credit card company that groups accounts together by the social security number of the primary cardholder in order to have a better understanding of total exposure to individual customers. If class size is measured by the number of distinct primary cardholder social security numbers, then customers with more than one credit account will be treated the same as customers with a single account even though their potential losses are greater. A better measure of class size in this case would be the number of distinct accounts.

In any case, the measurement of class size must include both current and former customers. The only exception is former customers whose information has been completely purged from the organization's records.

Customers not uniquely identified

For organizations that cannot or choose not to assign unique identification numbers to their customers, measuring class size can be more challenging. Consider the situation of retailers.

A major risk that retailers face with respect to customer privacy is the theft of payment card information. Retailers capture payment card information when processing customer transactions and store it along with other details of the transaction in a transaction file. If payment card information is not encrypted during the transaction process or when stored, it becomes an easy target for criminals since all that one needs to initiate a fraudulent charge are a card number, expiry date, verification code (a three or four digit number printed on the back or front of the card depending on the issuer) and possibly cardholder name. Criminals that are able to access a retailer's transaction file can make a lot of money in a relatively short period of time by selling stolen payment card information to others; for example, at $10 a card, a portfolio of 10 million cards would net $100 million.

There is no easy way to determine the number of distinct payment card numbers in a transaction file. For most retailers, the transaction file is simply too large and contains too much duplication, i.e., there are too many instances where the same card is used in multiple transactions. Even if a retailer uses a program to extract payment card information from the transaction file, that program would still have to be audited. Hence, retailers need some other way of measuring of class size.

One possible measure of class size is the number of transaction records. In theory, this quantity is straightforward to calculate but in practice, the resulting number is likely to be huge, particularly for national retailers. It is also an imperfect measure of class size because many transactions have no payment card information associated with them – they are cash transactions – and very often the same card is used for multiple transactions.

An alternative measure of class size is revenue or sales over a specified period of time. This quantity is generally easier to calculate than the number of transaction records but still has some of the same shortcomings, e.g., it includes cash transactions and does not adjust for duplication. It may also underestimate the exposure associated with transactions occurring in previous periods, e.g., if sales in the current period are lower than they have been historically. That being said, payment cards are generally reissued every few years, so payment card information in older transaction files can quickly become stale.

Measures of customer class size

From these observations, one sees that measuring class size is generally more complicated for customers than it is for employees. If customers are assigned a unique identification number and customer information is collated using this identification number then class size can be measured as the number of distinct customers or accounts, depending on the context. However, if customers are not assigned a unique identification number then some proxy for class size must be used. The simplest, but by no means an ideal, proxy is revenue or sales over some specified period of time.

Employee versus customer class size

For organizations that provide products or services directly to consumers, customer class size is generally greater than employee class size, often by several orders of magnitude, and employee class size can be ignored when measuring exposure for privacy liability coverage. However, for organizations that provide products or services to other organizations, e.g., auto parts manufacturers, this is not the case. Organizations like these have a small number of customers, e.g., in the case of auto parts manufacturers, auto assembly companies such as General Motors, Toyota and Volkswagen, but a workforce that can number in the tens of thousands. For such organizations, privacy liability exposure is primarily attributable to employees and customer class size can be ignored.

There are organizations that at first glance can appear to have a small number of customers relative to employees but in fact have privacy liability exposure that is many times the number of employees. A good example is appliance manufacturers. Appliance manufacturers sell their

products to appliance dealers who in turn sell them to consumers. Although there are more appliance dealers than auto assembly companies, the number of appliance dealers is still generally less than the number of people employed by appliance manufacturers. This suggests that number of employees may be a reasonable measure of privacy liability exposure. However, appliance manufacturers typically require registration of their products for warranties to be valid. Hence privacy liability exposure is not limited to dealers and employees; it also includes the consumers. Consequently, a quantity such as number of units sold is probably a better measure of exposure.

Candidate exposure measures for privacy liability coverage

Taking all of this into consideration, we make the following observations:

- If the number of employees is greater than the number of customers,[9] then a reasonable measure of exposure is number of employees.
- If the number of customers is greater than the number of employees and the organization assigns a unique identification number to each customer then a reasonable measure of exposure is number of distinct customers or accounts depending on the context.
- If the number of customers is greater than the number of employees but the organization does not assign a unique identification number to each customer and stores information by transaction, then a possible proxy for exposure is revenue or sales over some specified period of time.

Note that each of the identified measures of exposure satisfies the criteria stated earlier – simplicity, auditability, strength of relationship to losses, stability and legal determinability – although the degree to which the criteria are satisfied can vary. Note also that the exposure measure used in actual calculations needs to be scaled appropriately. For example, for an organization with tens of thousands of employees, millions of customers and annual sales in the

[9] The customers of the organization's customers should be included in the calculation of the number of customers if the organization maintains data on them.

hundreds of millions, employees should probably be counted in thousands, customers in thousands or possibly millions, and sales in millions.

6.1.2. Network Security Liability Coverage

Losses that arise under network security liability coverage consist of damage and defense costs related to the claims of third-party computer systems or networks.

Damages awarded depend on the harm caused to the third party's computer system or network and the extent to which the insured organization took preventative measures to address vulnerabilities in its systems and protect against the spread of malware infection. As noted earlier, it is impossible to know the harm caused by an event in advance and difficult to quantify preventative measures taken, so neither of these quantities is feasible for measuring exposure. An alternative approach is to consider quantities related to system access such as the number of endpoints (desktops, laptops, mobile devices, etc.) or the number of user IDs. Increases in these quantities are generally associated with increases in the risk of malware infection and hence expected losses.

Calculating the number of endpoints or user IDs on a computer system should be a relatively straightforward task. If not, it may be indicative of poor system management or security practices, something that should give pause to any prospective underwriter of the risk. Auditing these calculations should also be straightforward. Finally, there are no obvious legal barriers to performing these calculations and sharing the results with insurers or other third parties.

This suggests that either number of user IDs or number of endpoints is a reasonable measure of exposure for network security liability coverage. Of the two quantities, number of distinct users is probably slightly simpler to determine. However, for a well-managed system, neither quantity should be difficult to calculate.

6.1.3. Cyber Event Response Coverage

Losses that arise under cyber event response coverage consist of costs related to:

- Conducting an investigation
- Notifying impacted individuals and providing them credit monitoring services
- Retaining the services of legal and public relations experts.

Investigation costs

The cost of conducting an investigation is the cost of retaining the services of cybersecurity experts and/or forensic investigators to determine the cause of the particular cyber or privacy event and provide an initial assessment of the damage. Cybersecurity experts and forensic investigators are in high demand and generally bill by the hour. Hence, the cost of conducting an investigation depends on the time taken to complete the investigation. The larger and more complex an organization's computer system, the longer it is likely to take to determine how the system was compromised and what the extent of the damage is. This suggests that quantities related to the size and complexity of the organization's computer system are possible measures of exposure. One such quantity is the number of endpoints.

Notification and credit monitoring costs

The cost of notifying individuals whose confidential information has been or may have been compromised and providing them credit monitoring services for a period of time depends on the number of individuals impacted and whether these individuals are employees, i.e., individuals who provide a product or service to the organization, or customers, i.e., individuals who receive a product or service from the organization.

Notifying employees of a compromise can be done through an organization's internal communication system and costs almost nothing. On the other hand, notifying customers of a compromise may require the use of postal mail or a telephone contact center, which is much more expensive than email. Providing credit monitoring services is likely to cost the same whether the individual is a customer or employee.

These observations suggest that quantities related to customer class size are possible measures of exposure in this case. In the earlier discussion of privacy liability coverage, we argued that number of customers and number of accounts are the most appropriate measures of

customer class size, provided that these quantities can be determined. In cases where they cannot be determined, we argued that revenue or sales over a specified period is a reasonable, although not ideal, proxy.

Cost of legal and public relations experts

The cost of retaining legal experts to advise the insured organization of its obligations can be substantial. Depending on the nature of the cybersecurity or privacy event and the number of different legal jurisdictions involved, these costs can far exceed the costs of conducting a forensic investigation of the computer system. Although these costs can be difficult to predict in advance, number of customers or accounts could be considered reasonable proxies.

The cost of retaining public relations experts to develop a plan for maintaining public confidence in the organization are likely to be modest compared to the costs of conducting a forensic investigation, notifying and providing credit monitoring services to individuals impacted by a disclosure of confidential information, or retaining the services of legal experts. Hence, for the purposes of determining an exposure measure, these costs can be ignored.

Note that cyber event response coverage does not include coverage for legal costs associated with defending against liability claims or class-action lawsuits. These costs are covered under the privacy liability and network security liability coverages as well as elsewhere.

Candidate exposure measures for cyber event response coverage

The costs of conducting an investigation can be substantial, but in many cases, these costs will be dwarfed by the costs of notification and monitoring and the costs of retaining the services of legal experts. This suggests that a reasonable measure of exposure for cyber event response coverage is number of customers or accounts, depending on the context. If neither of these quantities can be determined, then a reasonable but by no means perfect proxy is revenue or sales over a specified period.

6.1.4. Other First Party Coverages

The remaining core coverages are all first party coverages. Losses that arise under these coverages are generally much smaller than the other core coverages and exposure can be measured using one of the candidates already identified.

Network interruption coverage

Losses that arise under network interruption coverage consist of:

- Lost income during the interruption
- Normal operating expenses during the interruption.

This coverage is effectively business interruption coverage. Lost income can be measured by revenue or sales over a specified period or by some other quantity related to transaction volume. Operating expenses can be measured using number of employees or average expenses; however, since operating expenses are usually targeted at some percentage of revenue or sales, it makes sense to measure them using revenue or sales as well. It follows that revenue or sales over a specified period is a reasonable measure of exposure for network interruption coverage.

Regulatory actions, fines and penalties coverage

Losses that arise under regulatory actions, fines and penalties coverage consist of:

- Costs associated with responding to a government agency's request for information
- Defense costs associated with a civil investigation or proceeding launched by a government agency
- Any resulting civil fines or penalties.

These costs are similar to fines levied in connection with the violation of a privacy statute. Based on the earlier discussion of privacy liability coverage, a reasonable measure of exposure is revenue or sales over a specified period of time.

Payment card industry assessments coverage

Losses that arise under payment card assessments coverage consist of contractual penalties levied by a payment card industry association for failing to comply with generally accepted

payment card industry data security standards. As in the case of regulatory fines and penalties, a reasonable measure of exposure is revenue or sales over a specified period of time.

6.2. Candidate Exposure Measures for Supplementary Coverages

A number of supplementary coverages are available as additions to the core coverages. Some of these coverages such as cyber extortion and cybercrime provide protection against a cyber risk, but several such as media content liability and reputation protection do not. As our primary interest in this report is exposure measures for cyber insurance, we focus on supplementary coverages where the connection to cyber risk is most clear.

6.2.1. Cyber Extortion Coverage

Losses that arise under cyber extortion coverage consist of costs related to:

- Conducting an investigation
- Retaining the services of experts in cyber extortion, law and public relations
- Ransom and reward payments
- Interruptions to operations
- Data restoration.

These costs are similar to those already encountered in connection with cyber event response and network interruption coverage but with some subtle differences. For example, the investigation required to determine the validity, cause and scope of an extortion threat could be more detailed and time-consuming than the investigation into an obvious cyberattack. Moreover, events related to extortion typically require the involvement of law enforcement, which can result in higher costs. Ransom and reward payments can also be difficult to predict and generally depend on the nature of the extortion threat. The same exposures identified earlier should be adequate for cyber extortion coverage as well.

6.2.2. Cybercrime Coverage

Losses that arise under cybercrime coverage consist of amounts that are fraudulently transferred from an account of the insured organization or the account(s) of its customers to an

account of the criminal by stealing login credentials or impersonating an employee or customer of the organization in some other way.

Losses suffered by an insured organization will depend on the number of fraudulent transactions and the amount of money transferred during each transaction. There is no easy way to estimate the amount of each fraudulent transaction. However, if one assumes that the risk of a fraudulent transaction is related to the number of individuals who could potentially be impersonated then the same exposure measures identified for privacy liability coverage should be reasonable candidates for cybercrime coverage as well.

6.2.3. Other Supplementary Coverages

The remaining supplementary coverages are media content liability, technology errors and omissions, third party bodily injury and property damage, enhanced business interruption, and reputation protection. As noted earlier, not all of these coverages are cyber or privacy coverages.

Since these coverages are only offered in conjunction with core coverages such as privacy liability, network security liability or cyber event response and since losses are likely to be smaller than the core coverages, it is reasonable to measure exposure in the same way as the core coverages.

7. Recommended Exposure Measures

Based on the observations and discussion in the previous section, we make the following recommendations.

For coverages where losses primarily involve exposure to individuals, e.g., privacy liability coverage, cyber event response coverage, etc., and the number of employees and customers can be determined, the recommended exposure measure is

- The number of employees if this is greater than the number of customers or

- The number of customers or accounts (depending on the context) if the number of employees is less than the number of customers.

Note that an employee is an individual who provides a product or service to the organization while a customer is an individual who receives a product or service from the organization.

For coverages where losses primarily involve exposure to computer systems hardware or software, e.g., network security liability coverage, recovery and restoration of digital assets coverage, etc., the recommended exposure measure is number of endpoints (desktops, laptops, mobile devices, etc.) or number of distinct user IDs, whichever is easier to determine.

For all other coverages, e.g., network interruption coverage, regulatory actions, fines and penalties coverage, etc., the recommended exposure measure is revenue or sales over a specified period.

Whatever exposure measure is used for a coverage, it should be scaled appropriately based on the context.

Acknowledgment

The author would like to thank the reviewers for their comments, particularly on issues related to cybersecurity, which undoubtedly helped to improve the final report.