How to mitigate ransomware risk through data and risk quantification

Received (in revised form): 20th May, 2024

Erik Sørup Andersen
Partner and Chief Executive Officer, Risk Measure, Denmark

Erik Sørup Andersen has worked professionally within the field of cyber security since 1999, when he joined a network security solutions provider as a software developer, designing and developing network security products. Later he managed the virtual private network (VPN) product portfolio and roadmap, working with global Internet Service Providers and system integrators to develop secure network solutions for their customers. For five years Erik led the central government information security risk management programme and was responsible for the adoption and subsequent representation of Denmark in the Common Criteria recognition arrangement. For the past 16 years Erik has provided consultancy to highly targeted and regulated organisations in Northern Europe. In recent years he has assisted organisations in adopting risk quantification techniques into their cyber risk management programme. Today Erik is Chief Executive Officer (CEO) of a risk management consultancy specialising in developing cyber risk quantification capabilities for organisations.

E-mail: esa@riskmeasure.dk

Abstract Ransomware attacks have, over the past years, been the most frequent cyberattack type and a growing community of adversaries continues to innovate methods for extorting organisations into paying ransom. Yet this risk is still, to many organisations, not well understood. Some refer to the averages reported in the media of the size of ransom and cost of ransomware attacks. But these numbers can be very far from the actual risk of a particular organisation. The nature of the risk, comprising many attack techniques and paths through an organisation's IT assets affecting a range of systems, data and the processes they support, makes it complex to describe and analyse. By using a risk analysis technique, where the risk scenario is decomposed to account for the contributions to the risk from different attack techniques, the vulnerabilities they exploit and the different forms of impact the attack inflicts on an organisation, it is possible to describe the risk in a more nuanced way unique to an organisation. Having created a model of the risk scenario that accounts for the factors relevant to the target organisation, it is possible to study mitigation options more consistently and simulate effects of implementing potential controls. Collecting data used to estimate the individual contributions to the total risk reduces the uncertainty of the risk measure and enables calculation of mitigation effects. This paper introduces the concept of quantitative risk assessment by highlighting results from quantitative studies of ransomware risk and providing examples of how data can be collected. Common pitfalls when using high-level data are demonstrated by showing examples of insights gained from collecting data about controls effectiveness. Being more effective in mitigating ransomware risk will both benefit the organisation directly and, by making ransomware attacks less profitable, society.

KEYWORDS: cyber risk quantification, ransomware risk, controls effectiveness, mitigation strategies

DOI: 10.69554/ZTGT3456

INTRODUCTION

According to Statista, 66 per cent of organisations worldwide were hit by ransomware attack in the period March 2022 to March 2023.[1] In Q4 2022, the share of ransomware in overall cyberattacks worldwide was 68.42 per cent, with a total of 154.93 million attacks registered in that quarter.[2]

With the potential to significantly disrupt business operations and cause reputational and financial damage, ransomware attacks are considered among the most persistent cyber threats worldwide. Organisations will therefore benefit from having a good understanding of this category of cyber risk, to support better mitigation decisions. To gain this understanding, a model of a ransomware risk scenario considering factors relevant to the target organisation can be build.

The nature of ransomware attacks varies, depending on the kind of organisation attacked and the intention of the attacker. For example, some threat actors deploy various other attacks alongside ransomware, such as information theft or distributed denial of service (DDoS). The loss incurred by a ransomware attack also varies. The differences in the ways attacks are carried out play a role, and so do organisational characteristics, such as the effectiveness of existing controls against a ransomware attack, the size and 'surface' of the organisation, and how the organisation depends on the processes affected by the attack.

The primary objective of this paper is to illustrate how an organisation, by use of data, can improve ransomware mitigation. The paper introduces a quantitative method for taking the factors affecting likelihood and loss into account when analysing the risk of ransomware. It further demonstrates how to determine effectiveness of mitigations.

It should be noted that cyber risk quantification (CRQ) is a discipline. The paper will briefly describe the elements of CRQ to help a reader unfamiliar with the discipline to understand the key concepts. The first section explains the difference between quantitative and qualitative risk assessment methods and why it is necessary to use a quantitative method to ensure an accurate result.

The next section, 'Two studies of ransomware risk', provides examples of insights gained from quantitative analysis of ransomware risk. Each study gives insights into the differences in how industries are affected by ransomware and shows that control effectiveness depends on the industry of the target organisation. From the first study, selected results from analysing 422 attacks are presented. From the second study, selected results from three model organisations in different industries are presented.

To determine the optimal choice of controls for a particular organisation, the unique characteristics of the organisation must be accounted for in the risk analysis. The following sections, 'Methodology for quantitative risk assessment', 'Data collection' and 'Risk quantification result', describe how the risk analysis is performed, from building a model of the risk to estimating the factors that must be accounted for in the risk analysis.

Having built a model of the ransomware risk scenario, the next step is considering 'Mitigation'. This section continues the data collection description, but now focusing on mitigation effectiveness which in turn helps prioritise mitigation recommendations provided by organisations such as SysAdmin, Audit, Network and Security (SANS), the Cybersecurity and Infrastructure Security Agency (CISA) and the National Institute of Standards and Technology (NIST).

QUANTITATIVE VERSUS QUALITATIVE RISK ASSESSMENT METHODS

The most widely used methods for risk assessments are qualitative. While there are

many flavours, what they have in common is the use of ordinal values for scoring likelihood and impact on a scale from typically 1–5 and then multiplying these values to determine the risk and place the result in a risk matrix — a heatmap. There are several problems with this approach:

- They force a precise assessment of likelihood and impact even though it is rarely true that all risks can be accurately placed in one of the values 1–5. For example, in most cases the loss from a ransomware attack is rather small but can potentially be very severe.
- It is not possible to add risk scenarios to provide a total risk. If an organisation considers a number of ransomware attack scenarios with each probability, the result of a qualitative risk assessment would be a number of red, yellow and green risks with values ranging from 1–25. Simply adding these numbers makes no sense as it does not show the amount of risk from the scenarios combined.
- This representation of risk cannot be used in any practical way to study mitigation effectiveness. For example, it is not clear whether reducing one of the red risks would be better than reducing three of the yellow risks.

Quantitative methods use ranges instead of ordinal values. Likelihood is measured as a frequency of loss events in a period. For example, the likelihood can be expressed as minimum 5 per cent, most likely 10 per cent and maximum 25 per cent, meaning that a loss event will happen once every 4–20 years. Impact is expressed in loss ranges. For example, the loss can be expressed as minimum US$50,000, most likely US$1m and maximum US$100m. Every loss within the range is included in the analysis, using probability theory and mathematically consistent ways to add different scenarios, yielding an expected loss distribution as the result.

To summarise, the quantitative method provides a representation of the risk that enables mathematically consistent studies of mitigation effects, as opposed to the qualitative method which is mathematically inconsistent. The next section demonstrates the usefulness of quantitative methods.

TWO STUDIES OF RANSOMWARE RISK

In this section, to illustrate the topic of this paper in practice — how to mitigate ransomware risk using quantitative methods and data — findings from two studies performed by Cambridge Centre for Risk Studies are presented. The first study shows that some controls are more effective than others in mitigating ransomware risk. The findings can be useful for prioritising among recommended controls. The second study finds that the expected loss from a ransomware attack varies significantly for different organisations. Using quantitative methods, the magnitude of the expected loss and where the loss comes from is calculated. Both provide valuable information when deciding how much to invest in mitigation and where in the organisation the investments should be made. The two studies will be referred to in the remainder of the paper.

Insights from ‘Mitigating Ransomware Risk: Determining Optimal Strategies for Businesses’

A study published in December 2022 by Cambridge Centre for Risk Studies in collaboration with Kivu Consulting[3] was able to see patterns in controls effectiveness on ransomware risk. The objective of this study was to provide an evidence base for controls prioritisation with respect to ransomware risk. Data from 300 organisations which paid ransom between May 2019 and January 2022, and another set of data about controls implemented at 180 organisations that were hit by a ransomware attack between

January 2021 and March 2022, was used for the study. There were 68 different types of ransomware variants and the differences in variants were considered.

Some of the key results are as follows:

- The study was able to correlate effectiveness of CIS 20 controls (V7)[4] with ransomware attack frequency and payment, and to identify groups of controls that are most strongly correlated.
- CIS Control 4 (controlled use of administrative privileges), Control 6 (maintenance, monitoring and analysis of audit logs) and Control 8 (malware defences, which in V7 includes endpoint detection and response [EDR]) were identified as the top three most effective controls at preventing or mitigating the attack. Control 3 (vulnerability management) was the fourth most effective control.
- Control 10 (data recovery capabilities) was not found to be effective at preventing or mitigating the attack. The study found that in many cases attackers were able to destroy data recovery or the control was irrelevant due to data exfiltration being used in the attack.

Industry differences were also seen in the data (see Figure 1). For example, Control 12 (network infrastructure management) is particularly effective for sectors with operational technology (OT) environments such as the energy sector and the consumer staples sector.

The study also investigated the cost-effectiveness of implementing controls properly compared to the amount paid in ransom and found that incident response (Control 18) was the most effective, followed by Controls 3 and 6 (see Figure 2). A possible explanation is that both controls limit the range and duration, and the less affected the victim is, the lower the ransom they will be willing to pay.

It is encouraged to use the data with caution and carefully examine the notes in the report. Nevertheless, the results found in this study can lead a company's efforts to mitigate ransomware risk in directions with highest control cost savings and effectiveness.

It should be noted that the study was inconclusive regarding some of the CIS 20 controls as the data points in those cases were too few.

In the section 'Mitigation' it is shown how a similar analysis can be made for a particular

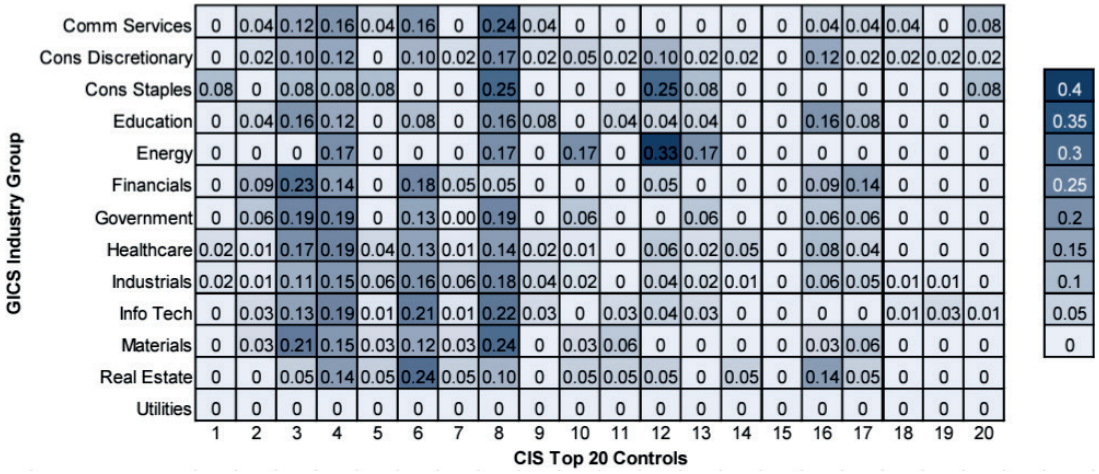

Figure 1: The effectiveness of controls on a ransomware attack, by industry. The darker the colour/bigger the number, the stronger effect the control has on the ransomware risk. Note that some controls have similar effectiveness across industries while other controls have big variations across industries, which is not always evident in general ransomware mitigation guidance

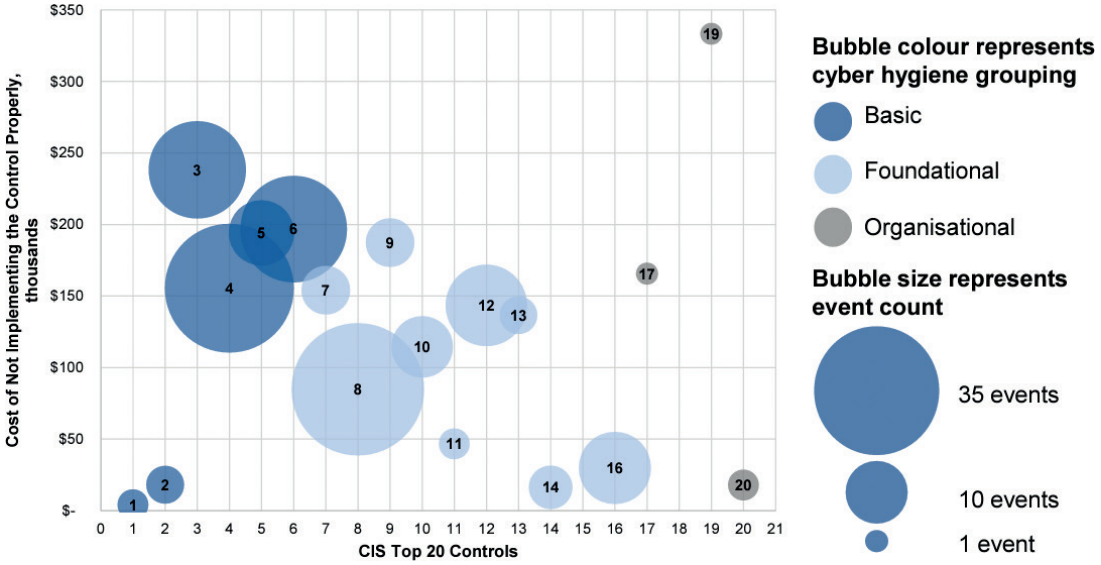

Figure 2: Cost-effectiveness of implementing controls properly. The bubble size shows the sample size for which the cost-effectiveness is derived

organisation, yielding a more accurate conclusion of mitigation effectiveness.

Insights from 'Cyber Security Cost Effectiveness for Business Risk Reduction'

Where the first study focused on mitigation effectiveness, the second study investigates the magnitude and nature of the loss from a ransomware attack. It makes an important point, that numbers like 'average ransom paid' or the extreme cases of ransom reported in the media are not good indicators of how an organisation will be affected. The study demonstrates how quantitative methods can provide insights that enable better mitigation decisions.

This study was performed in collaboration with BitSight. It describes a way to quantify the risk of ransomware for three model companies in the transportation, apparel retail and manufacturing industry respectively, and discusses the results.[5]

BitSight provided security ratings of the companies, using an outside-in telemetry approach[6] — essentially the attack surface and vulnerabilities that can be observed from outside a company. Ratings of a particular company were compared to an industry average and were correlated with ransomware statistics of the three industries, to provide a likelihood of each model company being attacked by ransomware.

In addition to the rating data, organisational characteristics and security controls information was collected for each model company.

Using these datasets, the study performed a quantitative risk analysis to calculate an expected total annual loss from ransomware attacks.

Some of the key results are as follows:

- The total annual loss expectancy (ALE) in percentage of earnings value was estimated at 0.23 per cent for the apparel retail company, 0.68 per cent for the transportation company and 0.74 per cent for the manufacturing company.
- In other words, the risk for the manufacturing company was 3.2 times higher than the risk for the apparel retail company.
- The revenue loss was the largest loss factor and accounted for 77–89 per cent of total loss.

- The ransom itself accounted for 5–10 per cent of total loss.
- Legal settlements accounted for 2–10 per cent of total loss.
- The other loss categories included in this study — labour costs, marketing and PR costs, data software and maintenance, impairment of property, plant and equipment, incident response cost, regulatory investigation and fines and compensation costs — were all insignificant compared to the above loss categories, except for the manufacturing company, where labour costs, data software and maintenance contributed 6 per cent. For example, performing manual overrides and defaulting to manually running a production can be a significant factor.

The key observation is that the loss from disrupting revenue-generating processes is 3.3–8.1 times larger than all other costs combined, making this the most interesting loss for which to seek mitigation.

The loss distribution is shown in Figure 3.

This level of detail of the ransomware risk allows for a more nuanced discussion about how to mitigate the risk.

- Is the overall risk acceptable?
- What processes or cashflows are driving the revenue loss?
- What controls would be most effective in mitigating revenue loss?
- What is the return on mitigation?

The result of the analysis performed in this study also shows the value of performing a quantitative analysis and how this presentation of the risk can be used for better informing mitigation decisions.

To take these concepts to work in practice for a particular organisation, the following section introduces the quantitative risk assessment approach applied to ransomware risk.

METHODOLOGY FOR QUANTITATIVE RISK ASSESSMENT

While the studies introduced in the previous section demonstrated the value

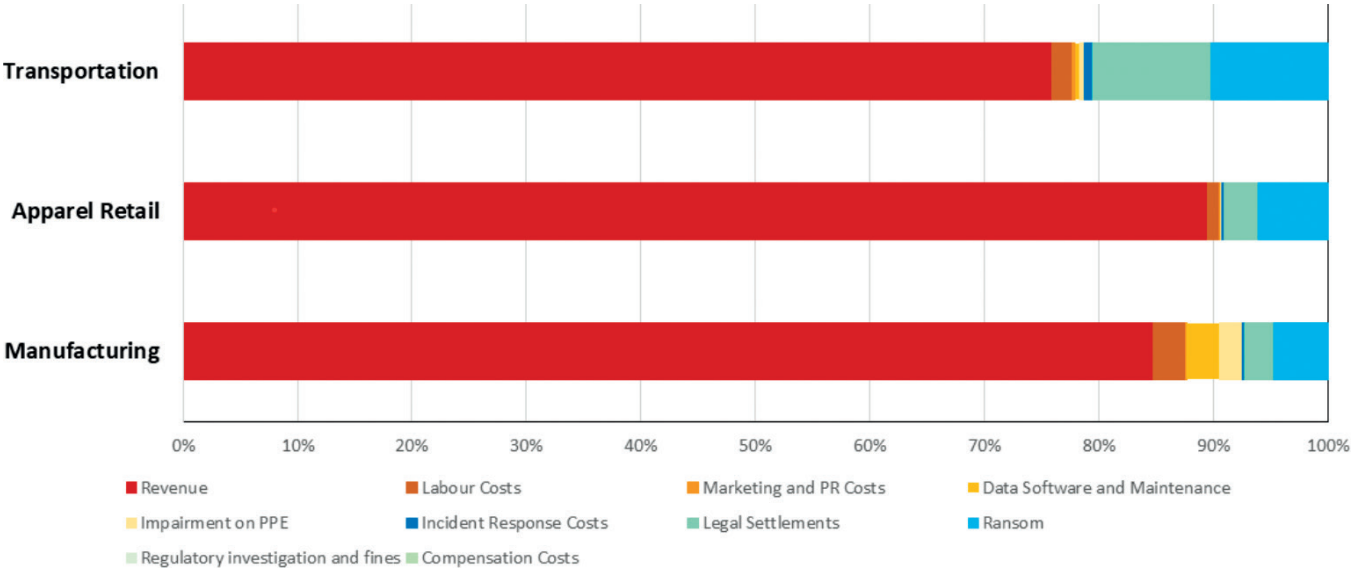

Figure 3: Loss distribution. The relative magnitude of loss forms from a ransomware attack on three model companies. It is seen that in all cases revenue loss is the primary factor of loss. Revenue loss comes from, for example, not being able to manufacture, ship products or take orders from customers. An example of labour cost would be having to perform tasks manually. Marketing and PR costs are costs associated with communicating to press and external stakeholders. Legal settlements may come from, for example, class action lawsuits or contractual clauses

of quantitative methods for risk analysis, the results presented are still not accurate enough for informing mitigation decisions in a particular organisation. It is necessary to perform the analysis using the target organisation's characteristics.

The factor analysis of information risk (FAIR) risk assessment methodology[7] described in this section can be used for an organisation-specific analysis. The FAIR method is the most widely used method and is supported by many tools available in the market. It is a value at risk (VAR) measurement technique.[8]

Performing a quantitative risk assessment involves the elements shown in Figure 4.

Each of the elements is briefly described below.

Scenario building

The first step of the analysis is to describe the risk scenario. This considers whether the whole organisation is affected, or only particular business units, platforms, production sites, regions, etc. The scenario must also describe threats in scope for the analysis. In the case of ransomware, the scope could, for example, be 'any kind of ransomware attack', 'targeted ransomware attacks' or 'ransomware groups that target industry x'.

Often, it is beneficial to start with a simple scenario and then subsequently expand the scenario or scope. For example, the main concern might be a manufacturing facility. Starting with one site or group of sites will help refine the model of the ransomware risk. Adding more sites or organisational units tends to be more straightforward than dealing with the complexity of a large scope initially.

Expert calibration

The estimates of frequency and impact bounds are performed by experts with knowledge about threats, vulnerabilities, the digital platforms in scope and other organisational characteristics. As the research of Daniel Kahneman concluded,[9] however, humans are not by nature good at making objective estimates, for a range of reasons. Humans tend to rely on intuition and expectation more than logic. The mistakes humans make are not random; there is a specific pattern to these mistakes called cognitive bias. Humans are typically overconfident and inconsistent and reach wrong conclusions.

It is, however, possible to train experts to perform accurate estimates of the lower and upper bounds of loss event frequency and loss. Training — or rather, calibrating experts — is a necessity to ensure useful estimates.

Estimation with ranges

Instead of using a point estimate, as the qualitative risk analysis techniques prescribe

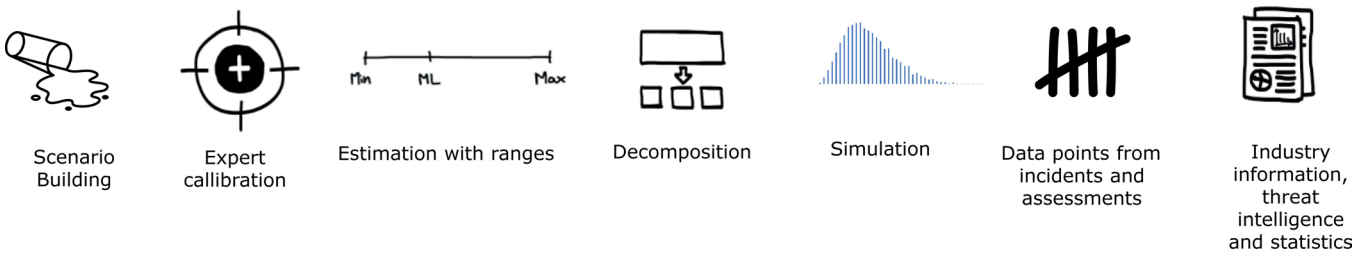

Figure 4: Elements of a quantitative risk assessment. The first step is 'scenario building'. Then the scenario is decomposed ('decomposition') into factors for which the range of values the factor can take is estimated. Correct estimation is achieved by a combination of (internal) datapoints, industry information and proper calibration of the participants delivering the estimates. A 90 per cent confidence interval is created, ie minimum is the 5 per cent fractile and maximum is the 95 per cent fractile of the full range of possible values for the factor

(see for example ISO 27005:2022),[10] quantitative assessments use ranges. A range more accurately describes the different degrees of impact an attack may have instead of being forced to choose a particular score or magnitude. To factor in which values in the range are most likely, a suitable distribution function is chosen. The distribution function then describes the relative probability of impact magnitudes in the range, as in the example shown in Figure 5, where the more frequent low-impact events are weighted higher, using a lognormal distribution.

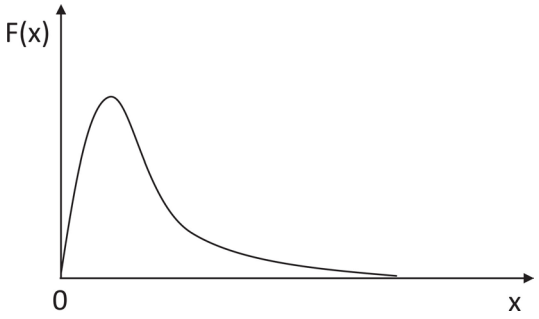

Figure 5: The lognormal distribution weighs the lower part of the range higher than the long tail. History shows that losses from ransomware attacks follow a lognormal distribution. Note that contrary to a normal (Gaussian) distribution, mode, median and mean are not the same. Hence the mean and median are not good predictors of the most likely loss (mode) experienced from an incident

Decomposition and data

Having described the scope, the analysis can be decomposed, as shown in Figure 6.

The decomposition is the basis for identifying which data is of interest. In general, data sources can be divided into internal data, collected from within the organisation, and external sources.

Simulation

When the risk has been decomposed and each factor has been estimated, the risk scenario is composed of several possible events with different frequencies (probability of happening within a year) and, if it occurs, a range of possible losses with a certain distribution of the losses within the range.

Using eg excel, R, Python or specialised software, it is possible to perform a simulation of the scenario. The technique used in the FAIR methodology is called Monte Carlo simulation. Monte Carlo simulation involves taking random sample values from each of the input distributions, performing calculations to derive a result value, and then repeating the process through a series of iterations to build up a distribution of the calculated risk from each iteration in the repetition. From this it is possible to draw a probability of the loss exceeding 'x' (see Figure 8).

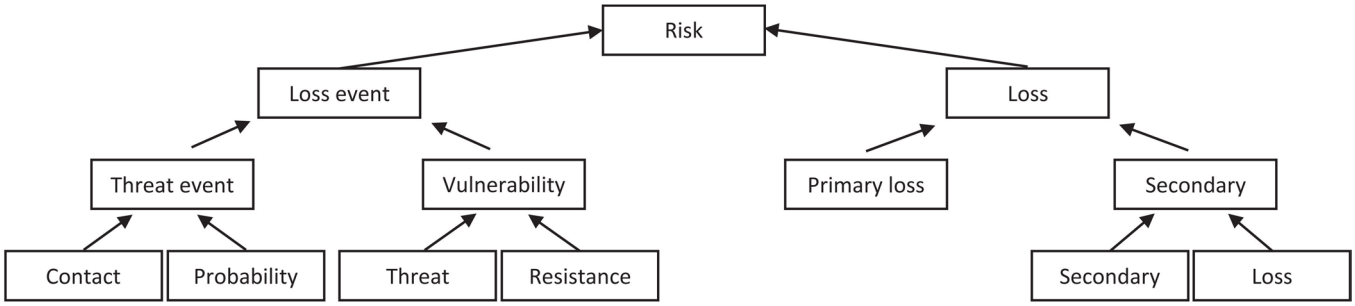

Figure 6: Decomposition of a risk scenario. The left side of the tree (loss event) is a frequency range. The right side is a loss magnitude range. Thus, the risk is calculated as a frequency range multiplied by a loss magnitude range modulated by probability distributions that weigh the values within the ranges according to how likely they are to occur in a given event. Monte Carlo simulation is used to 'replay' the event many times and count the number of times each risk value occurred. The counts divided by number of replays then represent the likelihood of a certain risk value

Data points from incidents and assessments

Reports of incidents are studied and used to provide good range estimates. Also, incidents that are different from the scenario under assessment may provide insights, as the investigations may indicate the strength of existing controls. In a similar manner, assessments and tests provide indications of the control strength.

Industry information, threat intelligence and statistics

External incident statistics are also useful for range estimation. Today, it is possible to acquire high-quality data about types of incidents for a particular industry in each region. Similarly, information from computer emergency response teams (CERTs) will provide current trends in attack methods which, together with the incident statistics, can be used to adjust range estimates.

The use of data is further explained in the following section.

DATA COLLECTION

To ensure accurate estimations, data that is significant, knowable and objective is identified and collected. Collected data is assessed with regard to how it is derived, what are the assumptions and what were the sources, before being used for estimation.

As a case in point for assessing how the data can be used, consider the study of the three model companies described above. The frequency of attack for each model company was estimated using an attack frequency risk profile industry-specific correlation. For some risk scenarios the approach used in the study is both effective and reasonably precise. The decisive points to assess before using an approach like the one in the study is whether this type of data is representative for the scenario in question, and whether the risk profiling has the relevant factors incorporated. Some variants of ransomware have unique characteristics and may not follow the general trend in the dataset. If the scenario is limited to certain ransomware variants, this potential source of error should be considered. The risk profiling data used in the study came from an outside-in telemetry approach. Krebs[11] compared the outside-in approach to 'judging the fire risk of a company from a photograph taken from the other side of the street'. And, as Jan Lemnitzer[12] pointed out, running honeypots on the internal network for threat research purposes could affect the rating negatively, whereas it could be argued that the rating should be affected positively. Both authors speak to the importance of including relevant internal risk parameters to achieve a more accurate risk profile.

In the following, examples of categories of data useful for estimation of ransomware attacks are discussed.

Threat event frequency

Ransomware attacks can be divided into opportunistic attacks and targeted attacks.

Opportunistic attacks using common 'spray and pray' tactics — such as phishing, social engineering and exploit kits — can target many organisations randomly and infect numerous desktops, laptops and servers with little effort. Therefore, the attack surface of an organisation is a good measure of the exposure and hence the frequency. Mapping the attack surface can be done with help from tools for the most part, but the non-technical part of the attack surface requires manual work.

Table 1, from the Federal Financial Institutions Examination Council (FFIEC),[13] illustrates this approach.

For each of the surface categories a set of relevant data points is collected, eg number of personal devices, employees, Internet-facing applications, third parties, sites, branches. Financial institutions may use the FFIEC approach directly to create their inherent risk profile and determine whether they should estimate the frequency to be above or below average. Other

Table 1: Attack surface

Surface category	Least	Minimal	Moderate	Significant	Most
Technology and connection types (14 data points)					
Delivery channels (3 data points)					
Online/mobile products and technology services (14 data points)					
Organisational characteristics (7 data points)					
External threats (1 data point)					

industries may have different attack surface characteristics but can use the structure as inspiration for identifying relevant data points.

Further refinements can be made if a particular ransomware threat actor or group of actors is considered. For example, geopolitical aspects may matter even for opportunistic attacks, of which NotPetya[14] is an example.

The threat event frequency for targeted ransomware attacks depends not only on the attack surface. The attack surface does matter, as the larger it is, the more likely is initial access to be acquired by a broker; however, before putting effort into an attack, the ransomware group will consider attractiveness of the potential targets. Some groups are known to attack certain industries, so industry matters too. In general, the following factors can play a role in the exposure to targeted ransomware attacks:

- *Size of the organisation*: Some groups have specialised in organisations that meet specific size characteristics.
- *Economical or reputational value*: A market leader that puts high value on brand and reputation.
- *Clients or customers*: Serving clients where discretion, confidentiality or privacy is critical.
- *Sector*: Some groups have specialised in healthcare, others in manufacturing/OT.
- *Digital platform*: Companies that have critical processes depending on a digital platform are more attractive.
- *Region*: The US is known as the most targeted region. Geopolitical factors may also play a role.

Vulnerability

Relevant vulnerability data is identified from a description of how an attack is performed. Threat reports and repositories such as MITRE ATT&CK[15] contain this data. As an example, phishing is a commonly used technique and the number and categories of employees vulnerable to phishing is therefore relevant vulnerability data.

Figure 7 shows the common techniques used in the first step of a ransomware attack and the controls prohibiting the execution of these techniques.

With the attack technique and surface elements identified, the vulnerability can be estimated by assessing the existence and maturity of the prohibiting controls.

Loss

The impact of the attack on the organisation is quantified in monetary units, the VAR also referred to as financial loss. In the decomposition shown in Figure 6, two categories are defined: the primary loss and the secondary loss. These two categories are then further decomposed, in order to collect data points for estimation.

Primary loss accounts for the things that happen immediately because of an event. Examples of primary loss include loss due to process disruptions, cost of handling the event and the ransom itself.

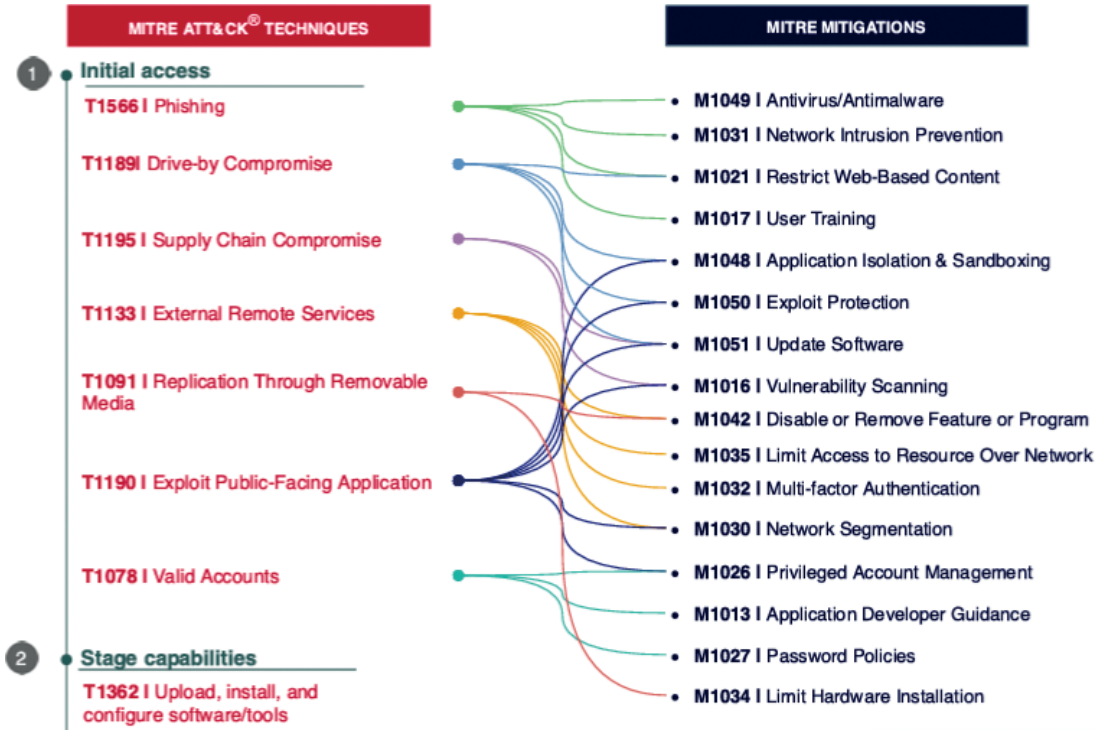

Figure 7: Mitigations to limit the first step in a typical ransomware attack[16]

The secondary loss category accounts for losses that are inflicted when external stakeholders react to the event. Examples of external stakeholders are customers, competitors, authorities, shareholders, media and business partners. Table 2 lists common loss categories and the data to collect for estimation of loss ranges for a ransomware attack.

For each loss form a range is estimated, eg if the loss per day from not being able to take orders is x, then the range is determined by multiplying x by minimum number of days order intake is disrupted for the lower bound of the range, and multiplying x by maximum number of days for the upper bound. Table 2 can serve as a starting point for further decomposition, eg revenue loss can be decomposed into one factor for each revenue-generating process type or even further for each process. In the case of order intake, if a company has several order intake processes that have different characteristics, a more precise estimate can be achieved by estimating each order intake process separately.

To estimate revenue loss, for example, the relevant processes are listed, and data about cost of disruption per unit time is then gathered from within the organisation. This can then be combined, as shown in Table 3.

Performing this decomposition will immediately reveal which process disruption will account for most revenue loss. If for example the revenue loss primarily results from not being able to process orders, it would make sense to investigate low-cost mitigation options immediately and assess more costly mitigation options later.

RISK QUANTIFICATION RESULT

The frequencies and loss ranges are used to calculate the total loss distribution, which then can be represented in the form of a loss exceedance curve (LEC), as shown in Figure 8.

Table 2: Factorisation of losses into loss forms

Loss category	What to look for	Data
Revenue	Production of goods and products Delivery of products and services Order intake	Loss per unit of time
Opportunities	Marketing events, product launch, planned events, mergers and acquisition (M&A) valuation	Loss from postponements or missed opportunity, clauses in contracts
Response	Restoring IT Incident response and forensics Manually reconstructing data Communication to external parties Internal communication Incident management Notification	Hours spent by internal employees External assistance from experts, marketing, cost of communication, eg letters, advertisements.
Replacement	Discarded products stored Damaged manufacturing facilities Emergency operations	Price per product or unit of raw material Cost of rebuilding Cost of relocation overtime costs to employees
Competitive advantage	Intellectual property leaked Market intelligence Performance differentiators	Premium pricing Market share Cost of delivery per unit
Legal	Contracts, licence to operate, fines, shareholder and class action lawsuits	Penalty units in contracts Legal assistance Cost per shareholder/stakeholder
Ransom	Ransom negotiation, ransom, payment assistance	Historic ransom claims Repeated ransom claims Repeated attack if ransom is paid
Reputation	Share price	Communication to market Restoring trust, eg additional security or guarantees

Table 3: Revenue loss

	Loss per unit	Units per day	Minimum days	Maximum days	Minimum loss	Maximum loss
Production of goods and products lost						
Delivery of products and services lost						
Order intake loss						
Total revenue loss					**Sum**	**Sum**

The curve depicts the likelihood of experiencing a loss of at least an amount x. Presenting the risk in this format makes it easier for non-IT executives and stakeholders to participate in a more nuanced discussion about risk tolerance by addressing questions such as:

- Is an 80 per cent likelihood of losing 0.2 per cent of annual revenue on average per year acceptable?
- What is the maximum loss we will tolerate every 50 years?
- What options do we have for bringing the 15 per cent likelihood of 2 per

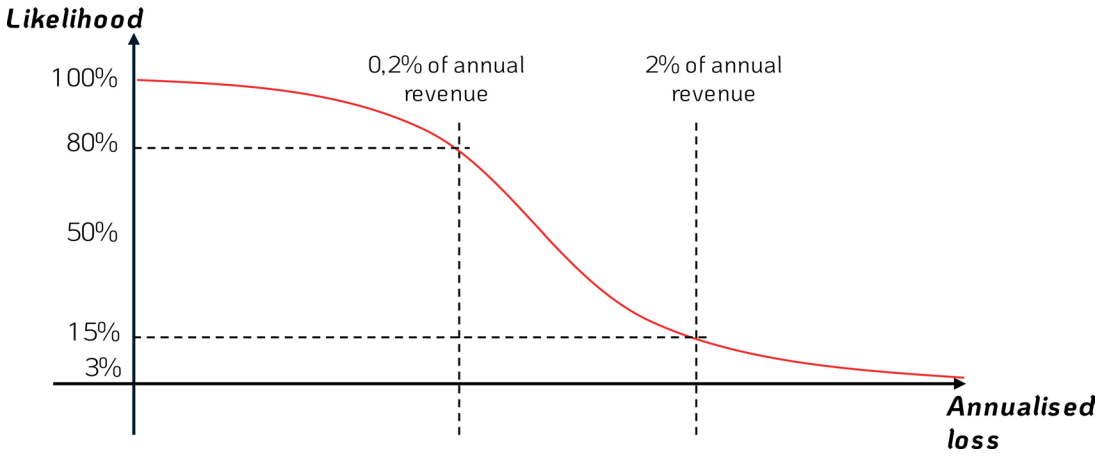

Figure 8: The LEC shows the probability of a loss exceeding a certain amount (the x-axis). The graph shows an 80 per cent likelihood of a loss greater than 0.2 per cent of annual revenue, each year, ie this loss happens four out of five years on average. It also shows a 15 per cent likelihood of a loss greater than 2 per cent of annual revenue, ie this loss happens once every 6.7 years on average

cent revenue loss down to a 5 per cent likelihood of a 2 per cent revenue loss?

With the guidance from stakeholders about risk acceptance and willingness to pay for risk reduction, the risk analysis team can then start to investigate mitigation options.

MITIGATION

Figure 9 illustrates conceptually how different categories of controls affect the risk. Generally, preventive controls reduce frequency and extent of an attack, whereas detective, responsive and recovery controls reduce impact.

Using Figure 9 as a reference, mitigation effects can be investigated.

Several sources provide suggestions for controls to mitigate ransomware risk:

- The studies referenced in the beginning of this paper (using Center for Internet Controls CIS 20 as reference) as they indicate effectiveness of different types of controls.
- Attack vector data, eg Coveware reports,[17] contain data about the current vulnerabilities exploited. Figure 10 shows data from the latest Coveware ransomware report.
- CISA Stop Ransomware guide[18] has controls recommendations for each step in the life cycle of a ransomware incident, some with very high cost that a quantitative analysis may justify and help prioritise.
- SANS has a selection of resources regarding ransomware[19] with controls recommendations, some of which are industry-specific as well as specific to company size.

Using this as a starting point, further data to estimate the effectiveness of mitigation can be collected.

In the following sections, two of the most common vectors, 'phishing' and 'vulnerability management', are discussed.

Phishing

Many organisations perform simulations and other types of training to mitigate phishing attempts. In a study published by Cyentia[20] based on data from 2,000 departments, it was not possible to measure the effect of awareness. The conclusions from the study were:

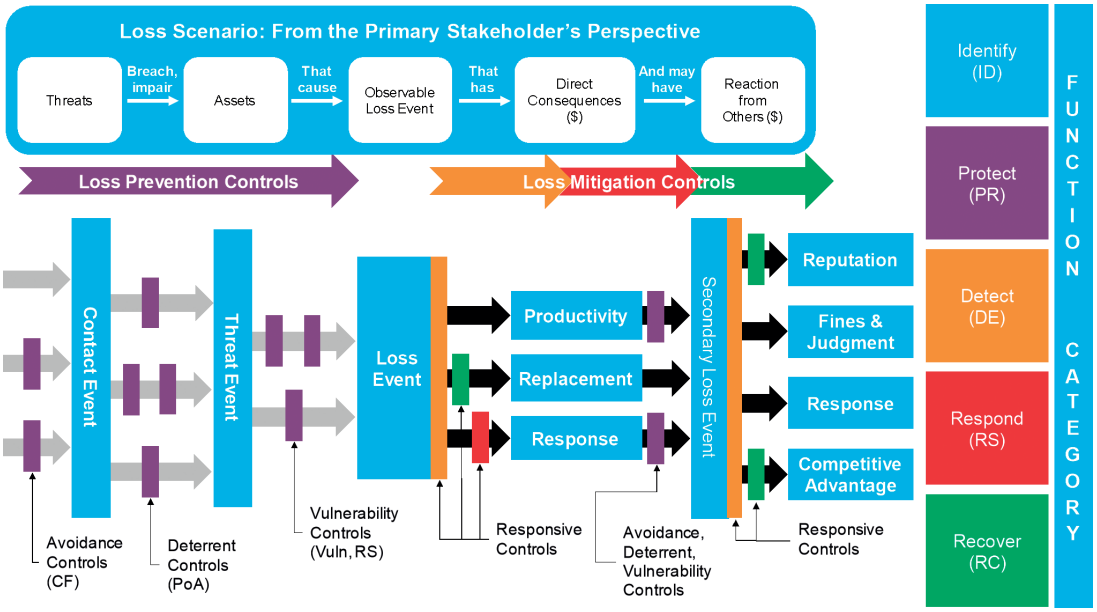

Figure 9: Decomposing an open FAIR loss scenario, including the open FAIR control categories and the NIST cyber security framework (CSF) five function categories: identify, protect, detect, respond and recover. It is seen that controls from the protect domain can reduce both likelihood and loss magnitude and must therefore be considered for both frequency and loss magnitude estimates[21]

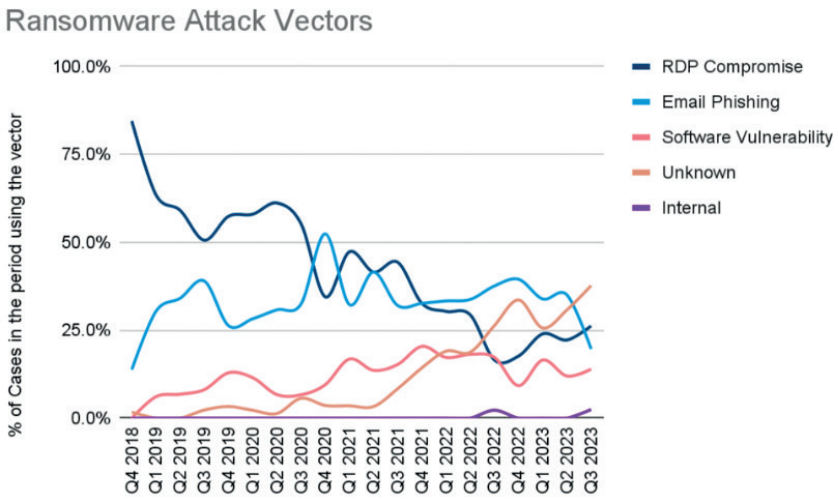

Figure 10: Trends in the attack vectors from ransomware attacks. Data like this helps experts perform better estimates
Source: Coveware ransomware reports

- It was not possible to see an effect of awareness on the frequency of successful attacks. Due to the departmental structure of access and credentials, a click rate of 6 per cent was shown to be enough for a full compromise of the organisations in the study.
- Users in the front line — the leaves in the org tree — are ten times more likely to fall for a phishing attempt.

- The malware incident frequency of users drops from 7.4 per cent to 0.4 per cent, ie a 19 times reduction, when a password manager is used.
- Users who participated in five training sessions were more likely to click (14.2 per cent) than users who only participated in one training session (11.2 per cent).

Using these findings, the return on targeted awareness, introducing password managers or improving access rights management can be estimated.

Vulnerability management

Studies have shown[22] that organisations are capable of patching approximately 10 per cent of all vulnerabilities in their environment in a given month. Identifying the relevant vulnerabilities to patch is thus essential for the effectiveness of vulnerability management. The following data can be used to identify which vulnerabilities to patch:

- Threat intelligence data from sector CERTS or threat intelligence providers can help understand which vulnerabilities are known to be exploited for ransomware attacks common in the sector.
- MITRE ATT&CK can be a useful source of details about the specific tactics, techniques and procedures (TTPs) of the group or groups known to attack the industry or sector in question.
- The description of a ransomware attack in the publication 'The anatomy of a targeted ransomware attack'[23] (referred to in Figure 3) can be helpful to identify vulnerability types used to perform an attack.
- The Exploit Prediction Scoring System (EPSS) project hosted by the Forum of Incident Response and Security Teams (FIRST)[24] shares data about the frequency with which specific vulnerabilities are exploited.

Insurance effectiveness

The effectiveness of an insurance can be measured as the cost of transferring a certain amount of risk to the insurance company and the percentage of the ALE transferred.

Insurance covers specific forms of losses and has an upper limit, that should be compared to the long-tail risk depicted in Figure 8. When estimating the effectiveness of insurance, the following should be considered:

- What loss forms are covered by the insurance, to what limit and how big a portion of the loss do they account for (see Figure 4)?
- What controls does the insurance company require to be in place?, eg malware protection, multi-factor authentication (MFA). If meeting the requirements implies investments in new or better controls, a new simulation should be performed, assuming the required controls are implemented.
- The likelihood of receiving the insurance claim in case of an attack. Negligence or evidence that required controls were not properly implemented may cause the insurance company to reject the claim. If a backup control was not working, it may not be possible to get coverage for the cost of rebuilding data.
- If the insurance conditions were found not to be fulfilled, there is a risk that the insurance company will charge for the cost of providing incident management support.

Data to support the estimation of the effectiveness of a cyber insurance can be gathered from internal controls assessments, threat simulations, incident reports, etc. Percentage of rejected claims with reasons for the rejection may also be requested from the insurance provider.[25]

It should be noted that, in some cases, cyber insurance is a contractual requirement. Therefore, assessing the cost of the insurance compared to the risk

transferred may not be the only factor in deciding on an insurance.

The long tail

The long tail of the LEC, ie the most severe but also more rare events, is often reputational damage, lost opportunities, irrecoverable data or very long recovery times.

Mitigations that will affect the long tail include business continuity, crisis management, ensuring that backups are regularly tested and well protected, regularly testing the company response playbook to a ransomware attack and having a good communication strategy to reduce reputational loss. Having described these mitigation initiatives, the effect is analysed by estimating how they affect the upper bound of the respective loss factors in the risk scenario decomposition.

Reputational damage

Often, reputational damage is perceived to be a major factor in the risk. Several studies have been undertaken to determine the effect on share price[26-28] and the general conclusion is that the correlation is weak. Some explanations proffer that 'good' communication is an effective way to mitigate reputation loss and stress the importance of timing of the communication, eg announcing an attack on a day with high news pressure attenuates the effect. Hence, assessing reputational damage can, in many cases, be performed by estimating the cost of communication and potentially the cost of implementing and running an elevated security level[29] for a given amount of time to reassure stakeholders that security is taken seriously.

Paying ransom

Paying ransom is both a potential loss and a potential mitigation. As indicated in the Cambridge study described above, ransom pay itself contributes about 5–10 per cent to the total loss. A closer examination shows a more nuanced picture. Figure 11 depicts a LEC created from the Cambridge study dataset.

Figure 11 shows that negotiation tends to halve the demand and that high payments are very rare. The data here is not compared to company size, but it is reasonable to assume that the ransom is correlated with company size. Using this assumption, the distribution of company sizes could be used to make a more qualified estimate of the likely ransom size, ie for a top 1 per cent company, a range around the ransom size at the 1 per cent mark could be used as an estimate.

The duration of negotiation and payment has an impact on revenue loss and can be estimated using the revenue loss estimates.

Reportings from ransomware attacks document a practice of using different encryption keys for subsets of data, forcing the victim to repeat payments.

The risk of not being able to recover data despite paying ransom is 25 per cent, according to a recent study by Veeam Software.[30]

If the motivation for paying ransom is to avoid data leakage in cases where the attack features info stealing, it should similarly be factored in that data may be leaked anyway.

The risk of being attacked again is higher if ransom is paid than if payment is denied. According to Arntz,[31] 38 per cent of those who denied payment experienced a repeat attack, compared to 80 per cent who paid ransom.

Assessing mitigation effectiveness

When the effect of the selected mitigation options is understood and estimated, the risk is recalculated incorporating the mitigations of choice. Several tools support what-if analyses where the effect of one or more mitigations can be simulated instantly. The result of a what-if analysis is illustrated in Figure 12.

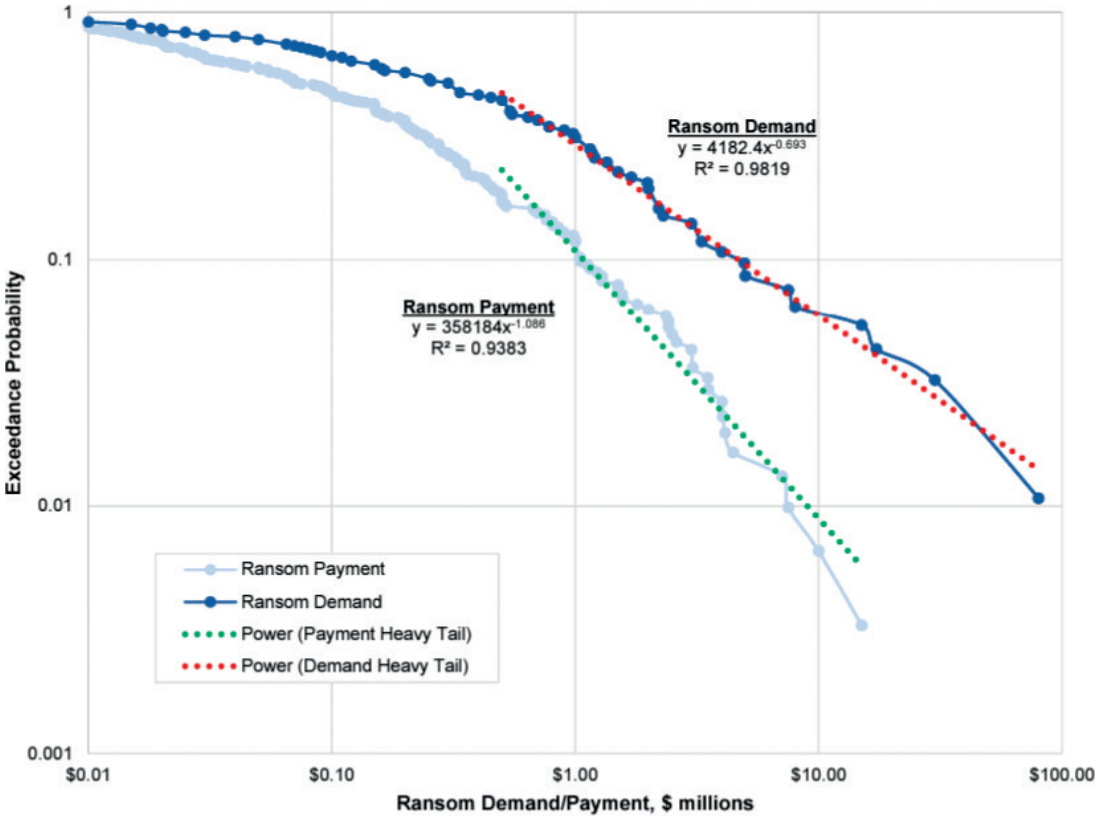

Figure 11: The likelihood of a given ransom size[32]

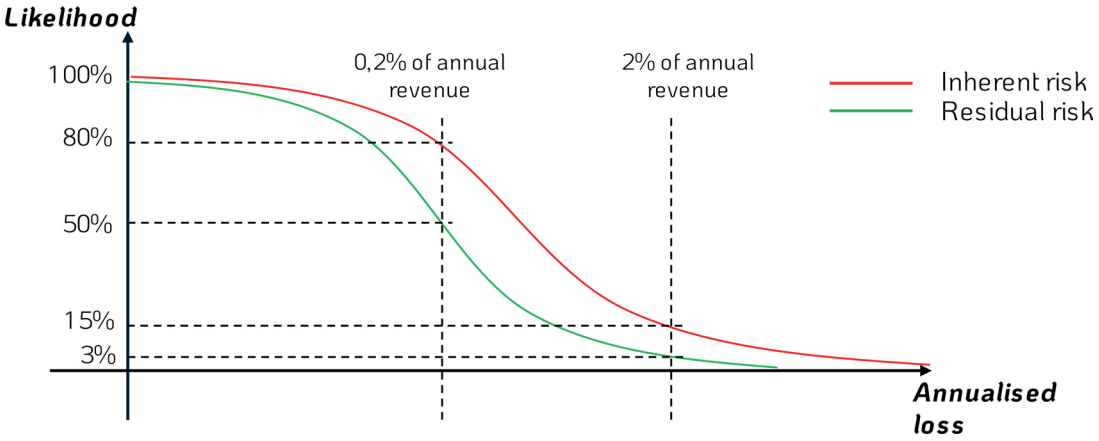

Figure 12: LEC: the red line depicts the risk prior to mitigations; the green line depicts the residual risk, after implementation of mitigations. The green line is calculated by adjusting the estimates of relevant factors in the decomposed scenario and then recalculating the LEC

SUMMARY

The decomposition of risk into factors that contribute to the ransomware risk and using data to estimate the individual risk factors makes it possible, by using a quantitative method like FAIR, to accurately calculate how much of the risk individual mitigation options could

remove and thus enables better mitigation decisions.

In the first section it was shown that quantitative studies were able to determine which controls are most effective at mitigating ransomware in general, for different industries and across the different ransomware variants seen in the dataset used for the study. Using quantitative methods, it was possible to determine most effective combinations of controls and is therefore a useful supplement to general guidance on ransomware mitigation including Coveware,[33] CISA Stop Ransomware[34] and SANS,[35] also referenced in the mitigation section.

The second study referenced showed that quantitative methods can help understand the overall magnitude of the ransomware risk as well as the primary loss drivers for a particular organisation. This insight provides a significantly better foundation for deciding whether to invest in any further controls at all and if so, where the biggest opportunities for risk reduction are seen. Compare this to a qualitative assessment where the ransomware risk is reported being 'red' or '25'.

In the subsequent sections the steps in performing a quantitative assessment were described. The intent was to give the reader an understanding of the work and tasks to perform to assess the ransomware risk for a particular organisation. The importance of assessing general assumptions about controls effectiveness was demonstrated by challenging assumptions for selected controls that are typically recommended for ransomware mitigation.

It was shown how a loss exceedance curve facilitates a more nuanced discussion and assessment of mitigation options that both IT and non-IT stakeholders can participate in.

GETTING STARTED

For organisations that have not yet adopted quantitative risk assessment methods for cyber risk assessment, 'Measuring and managing information risk'[36] and 'How to Measure Anything in Cybersecurity Risk'[37] can serve as introductions to the techniques.

References

1. Statista (2022), 'Volume of ransomware infections worldwide 2022', available at https://www.statista.com/statistics/1351338/global-volume-ransomware-organizations/ (accessed 21st November, 2023).
2. Statista, 'Topic: Ransomware', available at https://www.statista.com/topics/4136/ransomware/ (accessed 21st November, 2023).
3. Cambridge Judge Business School (2022), 'Mitigating ransomware risk: Determining optimal strategies for business', available at https://www.jbs.cam.ac.uk/faculty-research/centres/risk/publications/technology-and-space/mitigating-ransomware-risk-determining-optimal-strategies-for-business/ (accessed 19th November, 2023).
4. Cambridge Judge Business School (2022), 'BitSight cyber security cost effectiveness for business risk reduction', available at https://www.jbs.cam.ac.uk/faculty-research/centres/risk/publications/technology-and-space/bitsight-cyber-security-cost-effectiveness-for-business-risk-reduction/(accessed 19th November, 2023).
5. Center for Internet Security (CIS), 'CIS critical security Controls v7.1', available at https://www.cisecurity.org/controls/v7 (accessed 2nd December, 2023).
6. BitSight, 'How does BitSight work? How to use security ratings', available at https://www.bitsight.com/blog/how-does-bitsight-work (accessed 12th December, 2023).
7. Freund, J. (2014), *Measuring and Managing Information Risk: A FAIR Approach*, Butterworth-Heinemann, Woburn, MA.
8. International Organization for Standardization (ISO) (2019), 'IEC 31010:2019 Risk assessment techniques', available at https://www.iso.org/standard/72140.html (accessed 19th November, 2023).
9. Kahneman, D. (2013), *Thinking Fast and Slow*, Farrar, Straus, and Giroux, New York.
10. International Organization for Standardization (ISO) (2022), 'ISO/IEC 27005:2022 Information security, cybersecurity and privacy protection – Guidance on managing information security risks', available at https://www.iso.org/standard/80585.html (accessed 19th November, 2023).
11. KrebsonSecurity (December 2018), 'Scanning for flaws, scoring for security', available at https://krebsonsecurity.com/2018/12/scanning-for-flaws-scoring-for-security/ (accessed 3rd December, 2023).
12. Lemnitzer, J. M. (2020), 'Do we need an EU cybersecurity rating agency?', blog, Directions, available at https://directionsblog.eu/do-we-need-an-eu-cybersecurity-rating-agency/ (accessed 3rd December, 2023).
13. Federal Financial Institutions Examination Council

(FFIEC), 'Cybersecurity Assessment Tool', available at https://www.ffiec.gov/cyberassessmenttool.htm (accessed 22nd November, 2023).
14. Columbia University (2021), 'NotPetya: A Columbia university case study', available at https://www.sipa.columbia.edu/sites/default/files/2022-11/NotPetya%20Final.pdf (accessed 19th November, 2023).
15. MITRE ATT&CK, 'Groups', available at https://attack.mitre.org/groups/ (accessed 19th November, 2023).
16. Centre for Cybersecurity (November 2020), 'The anatomy of targeted ransomware attacks', available at https://www.cfcs.dk/globalassets/cfcs/dokumenter/rapporter/en/cfcs-the-anatomy-of-targeted-ransomware-attacks.pdf (accessed 12th December, 2023).
17. Coveware, 'Ransomware quarterly reports', available at https://www.coveware.com/ransomware-quarterly-reports (accessed 1st December, 2023).
18. Cybersecurity and Infrastructure Security Agency (CISA) (October 2023), '#StopRansomware Guide', available at https://www.cisa.gov/sites/default/files/2023-10/StopRansomware-Guide-508C-v3_1.pdf (accessed 24th April, 2024).
19. SysAdmin, Audit, Network and Security (SANS), 'Ransomware', available at https://www.sans.org/mlp/ransomware/ (accessed 24th April, 2024).
20. Elevate Security (2021), 'Elevating human attack surface management', available at https://elevatesecurity.com/resource/cyentia-elevating-human-attack-surface-management/ (accessed 19th November, 2023).
21. The Open Group, 'O-Ra 2.0.1', available at https://pubs.opengroup.org/security/o-ra/ (accessed 23rd November, 2023).
22. Baker, W. (2023), 'The Pithy P2P: 5 years of vulnerability remediation & exploitation research', Cyentia Institute, available at https://www.cyentia.com/pithy-p2p/ (accessed 13th December, 2023).
23. Centre for Cybersecurity, ref 16. above.
24. Forum of Incident Response and Security Teams (FIRST), 'Exploit Prediction Scoring System (EPSS)', available at https://www.first.org/epss/ (accessed 13th December, 2023).
25. Rasch, M. D. (2023), 'Coverage challenges in ransomware claims: Cyber insurance policies and trends in denials', Kohrman Jackson Krantz, available at https://kjk.com/2023/07/28/coverage-challenges-in-ransomware-claims-cyber-insurance-policies-and-trends-in-denials/ (accessed 19th November, 2023).
26. Ford, A. (May 2023), 'How do information security announcements affect stock markets?', BCS, available at https://www.bcs.org/articles-opinion-and-research/how-do-information-security-announcements-affect-stock-markets/ (accessed 20th November, 2023).
27. Kannan, K., Rees, J. and Sridhar, S. (December 2014), 'Market Reactions to Information Security Breach Announcements: An Empirical Analysis', *International Journal of Electronic Commerce*, Vol. 12, No. 1, pp. 69–91.
28. Foerderer, J. and Schuetz, S. W. (February 2022), 'Data breach announcements and stock market reactions: A matter of timing?', *Management Science*, Vol. 68, No. 10, pp. 7298–7322.
29. Kannan, *et al.*, ref. 27 above.
30. Veeam Software (2003), '2023 Ransomware Trends Report', available at https://www.veeam.com/ransomware-trends-report-2023 (accessed 1st December, 2023).
31. Arntz, P. (September 2023), 'The main causes for ransomware reinfection', blog, Malwarebytes, available at https://www.malwarebytes.com/blog/news/2023/09/the-main-causes-for-ransomware-reinfection (accessed 19th November, 2023).
32. Cambridge Judge Business School, ref. 3 above.
33. Coveware, ref. 17 above.
34. Cybersecurity and Infrastructure Security Agency, ref. 18 above.
35. SysAdmin, Audit, Network and Security (SANS), ref. 19 above.
36. Freund, ref. 7 above.
37. Hubbard, D. W. and Seiersen, R. (2016), *How to Measure Anything in Cybersecurity Risk*, John Wiley & Sons, Nashville, TN.