CYBER RISK ASSESSMENT AND MANAGEMENT

LIBFEXDLMCSECRAM01_E

LIBF

LEARNING OBJECTIVES

The **Cyber Risk Assessment and Management** course starts with an explanation why there is a business need for risk management and the anatomies of a data exfiltration attack. Basic information of cyber catastrophes and cyber risks will be provided.

To manage cyber risks, threats need to be measured. This course book will describe how threats can be measured, what metrics can be applied, and how organizations can measure threats. Measuring threats can only be done if the threat itself is known and understood. You will learn about three threat modeling methodologies: attack trees, STRIDE, and LINDDUN. These tools and methodologies can then be combined to form the risk assessment process. Furthermore, this course book will give you an introduction to standardizations, such as the NIST Risk Management Framework, the ISO/IEC 27005, and the BSI 100-3.

In the last part of this course book, the cyber-resilient organization and cyber insurance are explained. You will learn how to create a cyber-resilient organization that has a changing approach to risk management, incident response, crisis management, resilience engineering, and security solutions.

UNIT 1

ORGANIZATIONAL IT RISK MANAGEMENT

STUDY GOALS

On completion of this unit, you will be able to …

- understand why risk management is needed.
- identify different cyber catastrophes.
- define a cyber risk.

1. ORGANIZATIONAL IT RISK MANAGEMENT

Introduction

Risks are a part of every person's life, whether it is the risk of losing apartment keys or the risk of being robbed. After assessing a risk, we formulate a strategy to mitigate it: We might buy a keychain or avoid places where we might be pickpocketed. Businesses are also exposed to risks, and these need to be managed in some way. Risks are usually more complex than losing apartment keys. An example of such a risk is the possibility of losing data in a data breach, of which the possible outcome is, in most cases, a damaged reputation or financial impact (e.g., losing customers or fines). In short, this is why businesses need risk management. In this unit, the need businesses have for risk management will be justified through real life examples of cyber catastrophes. Furthermore, the base goals of risk management and how risk management can prevent cyber incidents will be addressed. To conclude, a formal definition of "risk" and "threats" will be provided.

1.1 Business Need of Risk Management

Building and leading a business offers excellent opportunities to achieve something great. Taking these opportunities, however, can carry great risks. The people in charge of the business need to know what those risks are and what the resulting impact of such risks can be. Knowing what the risks are is the first step; the next step is to know the possible mitigation measures. Is it possible to reduce the risk or is it better to transfer it to another party (e.g., an insurance company)? With proper risk management, the decision-makers are made aware of the risks and associated costs of a certain choice. The management can then decide if it is worth implementing such measures or if it can just be accepted. The following is a tabletop exercise of what a risk management process can look like. To conduct this simulation, we need to assume some facts:

- The business is a medium-sized corporation.
- The topic to be discussed is the European General Data Protection Regulation (GDPR).
- The company processes personal data in a web application.
- The website has a design flaw, which might leak data.

With this information, a risk can be identified. An obvious risk would be that the design flaw of the website can be used to extract data from the company. The worst case scenario is that this is made public and a fine has to be paid, as per GDPR guidelines. In this case, the mitigation strategy would be to fix the design flaw on the website. The cost of fixing the design flaw is X euros, and the cost of the fine and reputational damage is Y euros. With this, and other more detailed information, the decision-makers can decide whether they want to accept or fix the risk. In this example, the risk should be fixed if $\mathrm{Y}>\mathrm{X}$.

In the end, risk management is needed in a business to determine whether it is worth investing money in fixing an issue. From a financial point of view, it may not be worth eliminating the risk if its impact or likelihood to happen is low. Resolving issues with a high impact, however, is worth the costs. Risk management is the art of presenting decision-makers with a fact-based list of risks and how they are calculated. This list of risks can also help prioritize and plan corporate resources: A higher risk needs resources and attention more urgently than a lower risk (Siegel & Sweeney, 2020).

1.2 Anatomy of a Data Exfiltration Attack

Data exfiltration is a high risk for corporations. Some big corporations were victims of such attacks (e.g., Facebook in 2019, Capital One in 2019, and Equifax in 2017). In a data exfiltration attack, a malicious actor gains access to internal data. Most of the time, these are customer data or sensitive information. The following examples highlight the anatomy of a data exfiltration attack. The events and the consequences of some cases of data exfiltration will be presented.

Equifax Data Breach in 2017

In 2017 Equifax, a consumer credit reporting agency in the United States (US), was breached, and unauthorized actors gained access to customers data (Equifax, 2017). The data breach affected the data of 143 million US customers (Fruhlinger, 2020). The data included the following details of an individual (Electronic Privacy Information Center, 2020):

- name
- Social Security number
- birthdate
- address
- driver's license number

The main entry point was an unpatched vulnerability in **Apache Struts** (CVE-2017-5638). The hackers exploited this vulnerability to access internal servers of the corporate network. They gathered internal information, including employee credentials, and used those to gain further access into the network. They then scanned and exfiltrated information, undetected, for 76 days. To mask their activities the intruders encrypted the data and only exfiltrated small archives (Mort, 2017).

Apache Struts
The Apache Struts tool is an open-source framework for Java EE web applications.

As a result of this data breach, Equifax paid at least $575 million in a settlement (Federal Trade Commision, 2019). Equifax did not lose only money: The company lost the public's trust and received bad publicity after the breach. Brian Krebs, a well-known cybercrime journalist, called the response to the breach "a dumpster fire" (Krebs, 2017).

An interesting part of this data breach is that none of the exfiltrated data were sold on the darknet. This led to the theory that "normal" cybercriminals were not behind the data breach, but a nation state. This theory was confirmed when the US charged Chinese mili-

tary officers with the data breach (Benner, 2020). The Chinese government denied these accusations (CBS Interactive, 2020). In the end, victims' data are now in possession of an unauthorized third party (Fazzini, 2019), regardless of who was responsible.

Capital One Data Breach in 2019

In 2019, Capital One, a bank holding company in the US, suffered a massive data breach (Capital One Financial Corporation, 2019). An unauthorized person accessed more than 100 million customer user account data. The breach contained the following personal data:

- names
- addresses
- zip codes/postal codes
- phone numbers
- email addresses
- birthdates
- self-reported income
- customer status data, e.g., credit scores, credit limits, balances, payment history, and contact information
- Social Security numbers
- linked bank account numbers

Web Application Firewall
A Web Application Firewall is a Layer 7 firewall used to protect web applications from common attacks.

AWS metadata service
The AWS metadata service is a backend service in the AWS environment that provides credentials to resources.

The initial attack vector for this data exfiltration was a misconfigured **Web Application Firewall (WAF)**. This misconfigured WAF was used to obtain security credentials from the **AWS metadata service**. The credentials obtained from the metadata service for the WAF were able to list and sync data from the S3 buckets of Capital One. With this capability, the attacker downloaded nearly 30 GB of data (Novaes Neto et al., 2020). Capital One was fined $80 million for this data breach. The US Office of the Comptroller of the Currency (OCC) said "that the bank failed to identify and manage risks leading up to the move to cloud storage, and lacked sufficient network security and data loss prevention controls" (as cited in Schroeder, 2020, para. 5).

Facebook Data Breach in 2019

In 2019, a database with over 419 million records was found online (Holmes, 2021). Facebook denied a hacking attempt on their internal systems. The database contained the following data (Whittaker, 2019):

- user name
- Facebook ID
- phone number
- gender
- location by country

This data breach was not caused by a hack of the internal systems, but by an abused "feature." A technique called scraping was used to gather information from Facebook. Web scraping is used to harvest available data from websites. In the case of Facebook, the

search function was abused, and users could find friends on the platform with their phone numbers. Malicious actors used this to extract Facebook user IDs. Most likely, they gathered phone numbers and checked if a Facebook profile matched that phone number. After a previous incident (the Cambridge Analytica scandal), Facebook disabled this feature in April 2018, but the data had already been scraped (O'Sullivan, 2019).

In April 2021, the data from the database found in 2019 was made public in an underground forum which, once again, cast a bad light on the data security practices at Facebook. Initially, the data were obtained by abusing a legitimate function of the system. Disregarding the regulations, the data were given to an unauthorized third party without user consent. This is a violation of the general data protection regulation and can result in a fine of up to four percent of Facebook's total global turnover of the preceding fiscal year (intersoft consulting, 2018).

Risk Management of Data Exfiltration Attack

As seen in the examples above, a data breach or a data exfiltration attack can have a major impact. This impact ranges from reputation loss to hefty fines. Most of the time, it is cheaper to address the risk instead of paying the fines. Therefore, risk management must include the risk of data exfiltration and propose controls to mitigate such a risk. If the organization tackles the risk of data exfiltration, they are less likely to be fined by regulatory authorities.

1.3 Cyber Catastrophes

Catastrophes are part of our life on this planet. There are natural catastrophes (e.g., bush fires in Australia in 2020) and human-made catastrophes (e.g., the Chernobyl nuclear accident in 1986). These are all major catastrophes with severe impact, such as loss of lives and high ecological damage. Such catastrophes can also occur in cyber space. For example, cyber criminals attacked a water treatment plant in the US in 2020 and tried to add chemicals to the fresh water to poison it. Luckily, an employee intervened and reverse the change before the chemicals were added. If this attack had been successful, the poisoned water would have affected up to 15,000 residents. This attack could have been on the same level as the previous mentioned catastrophes (BBC, 2021).

General Catastrophes

Do events always have to be this severe to count as a catastrophe? To answer this question the definition of catastrophe needs to be analyzed. Cambridge Dictionary (n.d.) defines a catastrophe as "a sudden event that causes very great trouble or destruction." Firstly, this means casualties are not a factor to classify an event as a catastrophe. Secondly, "great trouble" is a subjective perception. It can be, of course, death or just a disruption to "normal" life. Lastly, the scale of a catastrophe is important. A small event might not be a catastrophe for the whole world, but it might indeed be one for a smaller group of people or a company. Another definition of a catastrophe is provided by the American Academy of Actuaries. They state that "catastrophes are infrequent events that

cause severe loss, injury, or property damage to a large population of exposures" (American Academy of Actuaries, p. 5). Summing up the definitions from Cambridge and the American Academy of Actuaries, the following factors determine whether an event is a catastrophe:

- low likelihood
- high impact
- large group impact

Unfortunately, these parameters of a catastrophe cannot be defined in an accurate and precise way. The severity of a catastrophe depends on the harm caused to the victim, be it individuals or organizations.

Cyber Catastrophes

Cyber catastrophes do not need to have an impact in the "real" world. A cyber catastrophe is a cyber event for which a high impact causes severe harm to an organization. If an event has the following properties, it can be considered a cyber catastrophe (Bashan & Lo Giudice, 2020):

- blast radius. A large group of users is affected by the event.
- outage. A service provided to the end user is degenerated. The end user is not able to use this service which impacts their work or life.
- uncontrollability. The organization is no longer able to control an event that affects it. The only possible action that the organization can take is to recover from the catastrophe.

The blast radius of an event can be a global radius (i.e., the global internet is affected) or a smaller radius with only a (bigger) group of users are affected. Hence, we can define two kinds of catastrophes: A local catastrophe, which is a catastrophic event for group of people or an organization, and a global catastrophe, which is an event with a global impact. The following table presents examples of recent global cyber catastrophes.

Table 1: Selected List of Recent Global Cyber Catastrophes

Year	Event	Description
2001	Code Red	Code Red was a computer worm that infected servers. The worm caused Denial of Service and defacement on the victims' system (Boyce, n.d.).
2008	Conficker	Conficker was a computer worm that slowed infected systems. The worm infected millions of computers and established a botnet from these systems (Burton, n.d.).
2016	Dyn DDoS attack	Dyn is a DNS server provider. A Distributed Denial of Service (DDoS) attack disrupted the service and brought down a major part of the internet. The DDoS was so widespread that it caused problems at internet service providers (Woolf, 2019).

Year	Event	Description
2017	WannaCry / Not Petya Ransomware	Ransomware attacks became common in 2017, disrupting big corporations and national organizations like hospitals. The attacker encrypts the data on the devices and demands a ransom to unencrypt the data (Hern, 2017).

Source: Created on behalf of IU (2021).

These examples show that a global event is not bound to a single group. Most internet users were affected by these incidents. Local catastrophes, however, are bound to a smaller user group, as seen in the following table.

Table 2: Selected List of Recent Local Cyber Catastrophes

Year	Event	Description
2019	AWS Cloud disruption	A power outage and failure of generators caused an incident in the AWS US-EAST-1 data center. This incident led to EC2 and EBS instances becoming unavailable. After power was restored, storage volumes suffered hardware damage and data were lost. Customers without backups were not able to restore their data and services (Abrams, 2019).
2020	Garmin ransomware attack	A ransomware attack encrypted Garmin systems. It is believed that Garmin was unable to recover the files and paid the $10 million ransom (BBC, 2020).
2021	Azure AAD outage	A new deployment of Azure AAD disrupted the service. Due this disruption, major Microsoft services (including Office, Teams, and the Azure Portal) failed. It took two hours to resolve the disruption (Foley, 2021).

Source: Created on behalf of IU (2021).

1.4 Cyber Risk

What is a risk? More specifically, what is a cyber risk? A risk is the level of impact a threat can have on an organization, combined with the likelihood of that threat occurring. Cyber risks are the result of a cyber threat and how likely it is that such a threat will impact an organization. Cyber risks originate from cyber threats, and cyber threats are limited to cyber space. For example, a data center burning down is not a cyber threat, as it is caused by a non-cyber event. If the fire is caused by a cyberattack, the threat is a cyber threat. Risk management in an organization needs to tackle both of these risks (Refsdal et al., 2015).

What Is a Threat?

A threat is the source of a risk. Generally speaking, a threat is an event with the potential to harm the organization. This resulting harm can affect the organization itself or an **asset** of the organization. The source of a threat (threat source) can be malicious (intentional) or

Asset
An asset is a resource, owned by an organization, that is valuable for the operation or business.

non-malicious (unintentional). Malicious threats are always the result of human action. To continue with the example of the burned down data center, the threat source is malicious if a threat actor purposely sets it on fire. In contrast, a non-malicious threat can be caused by a human or a non-human event. In the example of the burning data center, a human, non-malicious threat would be a fire caused by an electrician's mistake. A non-human, non-malicious threat would be a fire caused by a lightning strike.

Figure 1: Threat Source

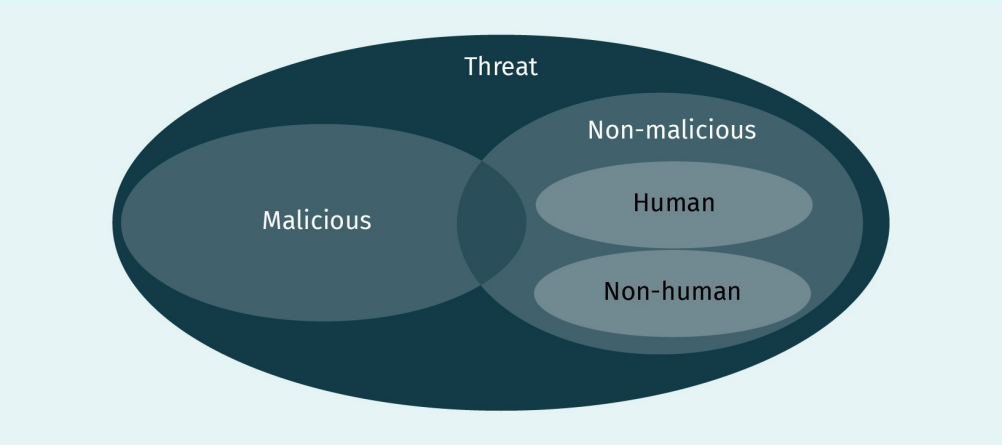

Source: Created on behalf of IU (2021), based on Refsdal et al. (2015).

Link between Risk and Threat

As previously mentioned, a threat does not equal a risk. Not all threats to an organization are also risks, but all risks contain a threat. A threat that stands on its own, independent from an organization, would not cause an issue. For example, an exploit for a specific system is a severe threat. However, if the organization is not using this particular system, it does not constitute a risk as it is impossible that this threat would cause an issue. In essence, a threat becomes a risk for the organization if there is a likelihood that this threat will cause an impact.

Abstraction of Risks

Not all risks are relevant for different management or reporting levels. Sometimes it is helpful to abstract or summarize risks, as shown in the following figure. The highest level is the abstract risk of "reputation loss." This risk is made up of several smaller risks that could harm the organization's reputation. The two examples here are "data exfiltration" and "employee misbehavior" (indecorous actions in public). These risks are, again, made up of even more specific risks. This abstraction can go on as long as there are risks to specify.

Figure 2: Risk Abstraction Diagram

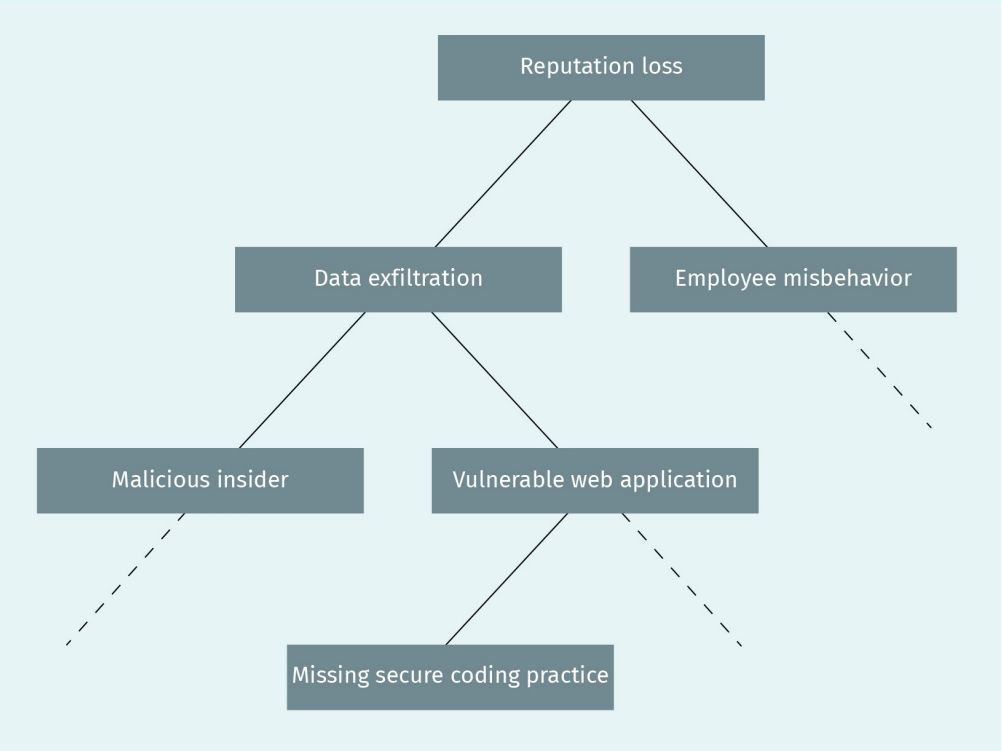

Source: Created on behalf of IU (2021).

SUMMARY

Risk management is a business need. Risk management offers tools to businesses to decide where to invest money to improve the resiliency of the organization. It provides facts-based aid for the decision-makers, used to prioritize and plan the investment to fix vulnerabilities.

A severe risk for many organizations is the risk of a data exfiltration attack. In a successful data exfiltration attack (also known as data breach), an attacker steals sensitive or confidential data from an organization. The data can be customer information or intellectual property of the company. Most of the time, a data breach has a high impact on the business. The impact can be financial (e.g., fines or settlements) or reputational (e.g., bad press and loss of customers).

Cyber catastrophes are events with the maximum impact on an organization. A catastrophe can have a global or local impact. Catastrophes are always severe. They have a blast radius of affected individuals or organizations. When there is an outage, for example, and the affected service

cannot be provided or is highly degenerated, the victim organization has failed to contain the event and can now only try to recover from this catastrophe.

A risk is the level of impact a threat has on an organization and the likelihood that this threat will affect the organization. A cyber risk is the result of a threat that only exists in cyber space (e.g., a malware attack). A risk for a cyber device is not always a cyber risk (e.g., hardware failure of a server). A threat is an incident that might cause harm to an organization or one of its assets. The threat source can be malicious or non-malicious. A malicious threat is always human made, whereas a non-malicious threat can be initiated by a human or a non-human event (natural causes). A threat can be abstracted into different layers to ease the understanding of the containing threats.

UNIT 2

MEASURING THE CYBER THREAT

STUDY GOALS

On completion of this unit, you will be able to …

- understand threat measurement.
- calculate the metrics of a risk.
- identify the differences between a catastrophe and a black swan event.
- explain the likelihood of major cyber events.

2. MEASURING THE CYBER THREAT

Introduction

Most of the time, multiple risks are present in an organization. The challenge is comparing those risks to facts, removed from feeling. This challenge is addressed by assessing the risk with measurable values. In our personal lives, we can measure a risk with simple metrics. For example, if we leave a bicycle somewhere without a lock, there is a risk that the bicycle will be stolen. There are two factors to be considered: the impact of the risk and the likelihood of it happening. The impact of the risk in the example is the amount of money we would lose when the bike is stolen. The likelihood of the bike being stolen could be investigated with questions, such as “Is it on a private property? Is it old so no one wants it?”

Corporations need to measure the threats and risks to their organization. Only with measurable and comparable data do the decision-makers have grounds for their decisions. Risks can be divided into separate values, and a risk level can be calculated. With these values, a manager can easily decide if they want to tackle a risk (e.g., is the cost of fixing higher than the cost of the risk?). The main concerns are to determine how risks can be calculated and how they can be managed. This unit will introduce a method of measuring a risk. Moreover, we will use metrics to calculate the likelihood of factors in real life events.

2.1 Measurement and Management

Risk management cannot be based on gut feelings. To be manageable, risks need to be properly measured.

Measurement

To measure the risk of a threat, a method for classifying risks with measurable parameters must be defined. A risk is the possibility of something bad happening. This means that a risk is the likelihood that a threat will generate a bad impact for an organization. This definition can now be used to sketch a simple formula to calculate the risk of a threat:

$$\text{Risk} = \text{Impact} \cdot \text{Likelihood}$$

This formula describes the risk. To have real value, the likelihood and the impact need to be measurable. Each one of them has to be split into more detailed metrics that are easier to define and measure. The following figure shows one way to do this.

Figure 3: Attack Path

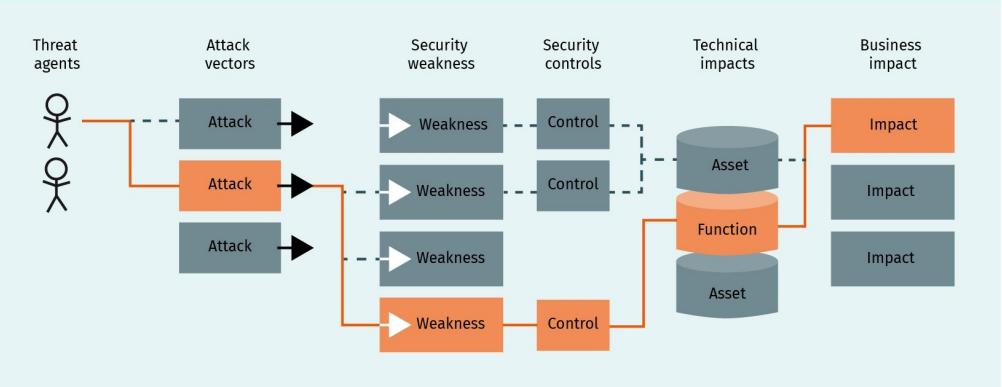

Source: Williams (2020). CC BY-SA 4.0.

The image shows the path of a threat actor (or threat agent) on their way to damage or impact the business. The impact can be described as technical and business impact. The technical impact describes the impact on an IT system, and the business impact describes the impact of the threat to the business. A technical impact must always be present for a business impact. No threat to a cyber-system has a direct impact on a business: First, there has to be an impact on the system itself.

A system can be exploited by a threat actor. Some systems are easy to exploit, and others are more secure, requiring specialized knowledge from the threat actor. Before a threat actor can impact a system, they need to find a way to do so. This can be described as the likelihood: How likely is it for a threat actor to find this way? Firstly, the threat actor needs to have an attack vector that targets a vulnerability in the organization's systems or processes. Such systems or processes are normally protected by security controls. These controls normally prevent a weakness from being used to impact the system. Consequently, the attacker needs to find a weakness that has no security control or does not detect the attack vector, which also attacks the weakness. Summing up this chain of attack, the likelihood and impact of a threat can be visualized as in the following figure (Williams, 2020).

Figure 4: Likelihood and Impact of a Threat

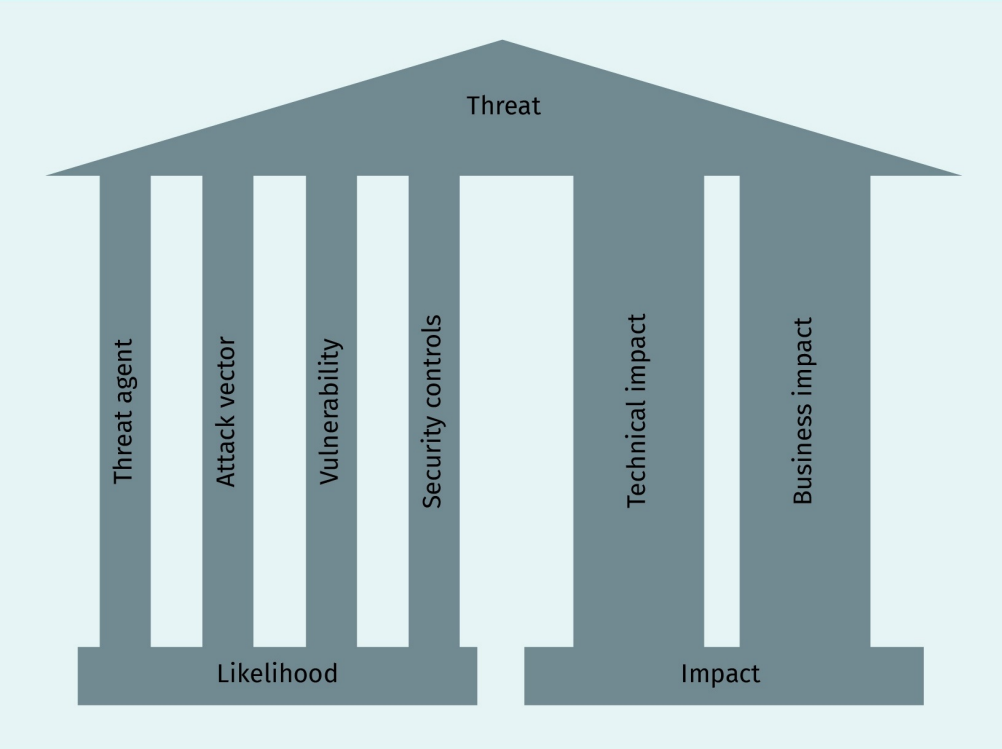

Source: Created on behalf of IU (2021), based on Williams (2020).

Management

After successfully measuring a risk, it needs to be managed. This means that the responsible person needs to decide how to mitigate the risk. Normally, there are four actions that can be taken to tackle the risk (Niedbala, 2021):

1. Avoid the risk
2. Reduce or mitigate risk
3. Transfer the risk
4. Accept the risk

Risk owner
The risk owner is the person or group responsible for handling the risk.

All of these measures have a cost attached to them. In the risk management process, the **risk owner** has to decide which measure is viable and has a lower cost than the actual cost of the risk, should the attack occur.

Avoiding the risk

To avoid the risk, the responsible person needs to define actions to make it impossible for the threat to occur. This can be done by remodeling technical aspects of a system or design processes, which will change the usage of the system. The following scenario illustrates this action.

Assume a system handles confidential information for an organization. This system is accessible from the internet and the front end website of this organization. This front end has a critical vulnerability, which is easily exploitable. This now presents the risk that the confidential data are being stolen by a third-party via the front end. To avoid the risk of the third party accessing the data via the front end, a decision can be made to split the system into two separate systems. The public system will be moved to a third party provider and the access from the internet to the corporate network will be turned off. This action has the clear goal of avoiding the risk that the confidential data are stolen via the public accessible front end. The front end is still vulnerable, but the data are safer.

Reduce or mitigate risk

Reducing the risk is the action of mitigating the vulnerabilities or adding measures to reduce the risk. Risk reduction or mitigation always contains an action that reduces the likelihood or the impact. Risk mitigation has different methods. These methods depend on the threat of the risk and the means of the organization. Some of these mitigation measures are

- fixing vulnerabilities,
- adding security controls,
- mitigating security flaws,
- changing the design, and
- adding prevention and detection mechanisms.

To illustrate this measure, the previous example can be used. In this scenario, the origin of the risk is that the organization's web application is vulnerable. To reduce or mitigate this risk, the people responsible have to reduce the likelihood that this can happen. The obvious action would be to fix the security vulnerability in the application. Another method would be to install an **Intrusion Detection and Prevention System (IDPS).** This system detects and prevents the access to the confidential data. With these actions, the likelihood of the data being compromised by the vulnerability in the frontend is reduced.

Intrusion Detection and Prevention System (IDPS)
An IDPS monitors the network or systems to find suspicious traffic or actions. If the system detects such activities, it can automatically block them.

Transfer the risk

We have the option to transfer the risk to a third party. This third party is then the owner of the risk and has to carry it. The transfer of the risk is often handled with legally binding contracts. These contracts state that the third party is now responsible for the risk. These third parties can be insurance companies or a service providers. In the example in this section, the organization could approach a cyber insurance company to cover the theft of their confidential data. They could also contract a managed service provider to care for the system and state in the contract that the security of the confidential data has to be warranted by this provider.

Accept the risk

On paper, accepting the risk is the easiest action a risk owner can take. If this action is analyzed further, it is no longer that easy. Risk acceptance is not about ignoring the risk. When accepting the risk, the risk owner needs to know what it would cost if this risk came true,

and how often this could be the case. The risk also needs to be monitored and not forgotten. The monitoring of the risk is needed so that the business can adjust to it. This could be through a change in the vulnerability or a change in the business.

Accepting the risk is not the lazy option. Sometimes there is simply no way to avoid, reduce, or transfer the risk. Some risks need to be accepted to either start or improve a business. We must balance taking risks and avoiding or mitigating them. Good risk management processes find this balance.

In the presented example, the risk owner accepts the risk that an attack might be able to steal data through the vulnerable frontend. To comply with the organization's risk management, the risk owner sets up a risk management plan to monitor the risk and to find a better solution later on.

2.2 Cyber Threat Metrics

A metric is the easiest way to quantify a specific topic. In risk management, quantification is key for success. Unfortunately, "likelihood" and "impact" are not easily quantifiable. Nevertheless, to measure the risk for an organization they need be quantifiable. The **Open Web Application Security Project (OWASP)** has published measurable factors for the elements of the impact and the likelihood. To quantify the factors, a value from zero to nine is added to each level of the factor. With these values, proper metrics can be obtained. In what follows, a possible quantification of the likelihood and the impact of a risk will be described. These metrics are general values and may need to be adapted to the particular needs of a business.

Open Web Application Security Project
The OWASP is a nonprofit organization focused on improving software security.

Likelihood

As previously described, the likelihood can be split into different parts. These metrics are the foundation of the quantified likelihood.

Threat actor

The threat actor is the first factor of the likelihood. This factor describes the parameters the threat actor needs to exploit a vulnerability. The worst-case threat actor should be used for the calculation. The following factors can be used for the description (Williams, 2020, Threat actor section).

Table 3: Threat Actor Factors

Skill level	What is the skill of the threat actors?	(1) No technical skill (3) Some technical skills (5) Advanced computer user (6) Network and programming skills (9) Security penetration skills

Motive	How high is the motivation of the threat actors?	(1) Low or no reward (5) Possible reward (9) High reward
Opportunity	Which resources and what access is required?	(0) Full access or expensive resources required (4) Special access or resources required (7) Some access or resources required (9) No access or resources required
Size	How large is the group of threat actors?	(2) Developers (2) System administrators (4) Intranet users (5) Partners (6) Authenticated user (9) Anonymous internet users

Source: Created on behalf of IU (2021), based on Williams (2020). CC BY-SA 4.0.

Vulnerability

After defining the factor for the threat actor, the same is applied to the vulnerability. The following factors can be used to describe the vulnerability (Williams, 2020, Vulnerability factors section).

Table 4: Vulnerability Factors

Ease of discovery	How easily can the threat actor find the vulnerability?	(1) Practically impossible (3) Difficult (7) Easy (9) Automated tools available
Ease of exploit	How easily can the vulnerability be exploited by the threat actor?	(1) Theoretical (3) Difficult (5) Easy (9) Automated tools available
Awareness	How well-known is the vulnerability to the threat actor?	(1) Unknown (4) Hidden (6) Obvious (9) Public knowledge
Intrusion detection	How easily can the exploit be detected?	(1) Active detection in application (3) Logged and reviewed (8) Logged without review (9) Not logged

Source: Created on behalf of IU (2021), based on Williams (2020). CC BY-SA 4.0.

Impact

After defining the factors for the likelihood, we define the impact factors. Normally, the impact can be split into two parts. The technical impact is the direct impact the threat has on the technical system (e.g., the server or database). The business impact is the threat to the organization's business, i.e., the losses that the organization has if the threat is executed, and is the more important of the two.

Technical impact

The technical impact is the factor that describes what happens to the system it is attacked or once the threat is carried out. The technical impact can be described with the confidentiality, integrity, and accountability (CIA) principle. To complete the factors for the technical impact, accountability can be added (Williams, 2020, Technical impact factors section).

Table 5: Technical Impact Factors

Loss of confidentiality	How sensitive are the data and how much can be disclosed?	(1) Minimal non-sensitive data disclosed (6) Minimal critical data disclosed (6) Extensive non-sensitive data disclosed (7) Extensive critical data disclosed (9) All data disclosed
Loss of integrity	How much corruption damage can be done?	(1) Minimal, slightly corrupt data (3) Minimal, seriously corrupt data (5) Extensive, slightly corrupt data (7) Extensive, seriously corrupt data (9) All data totally corrupt
Loss of availability	How much can the service be disrupted?	(1) Minimal secondary services interrupted (5) Minimal primary services interrupted (5) Extensive secondary services interrupted (7) Extensive primary services interrupted (9) All services completely lost
Loss of accountability	Can the threat actor be traced to an individual?	(1) Fully traceable (7) Possibly traceable (9) Completely anonymous

Source: Created on behalf of IU (2021), based on Williams (2020). CC BY-SA 4.0.

Business impact

The business impact is the most important factor for the risk because all effort from risk management is targeted to minimize the risk for the business. This factor justifies the action to fix a risk. Common factors are listed in the table below (Williams, 2020, Business impact factors section).

Table 6: Business Impact Factors

Financial damage	How much financial damage could this threat cost?	(1) Less than the cost to fix the vulnerability (3) Minor effect on annual profit (7) Significant effect on annual profit (9) Bankruptcy
Reputation damage	How much would the reputation suffer?	(1) Minimal damage (4) Loss of major accounts (5) Loss of goodwill (9) Brand damage
Non-compliance	How much does this threat violate compliance?	(2) Minor violation (5) Clear violation (7) High profile violation
Privacy violation	Could the threat result in a privacy violation and how would it be affected?	(3) One individual (5) Hundreds of people (7) Thousands of people (9) Millions of people

Source: Created on behalf of IU (2021), based on Williams (2020). CC BY-SA 4.0.

2.3 Measuring the Threat for an Organization

An organization needs a defined metric to measure a threat. These metrics are used to calculate the overall risk with the formula $\mathrm{Risk} = \mathrm{Impact} \times \mathrm{Likelihood}$. To have a repeatable method, the metrics can be used to calculate both the impact and the likelihood. These factors are then multiplied to obtain the risk. For a visual representation of the risk, a simple traffic light table can be used. For this overview, a color is assigned to each value.

Figure 5: Likelihood and Impact Levels

Value	Level
0 to <3	LOW
3 to <6	MEDIUM
6 to 9	HIGH

Source: Created on behalf of IU (2021), based on Williams (2020). CC BY-SA 4.0.

It is important to have a repeatable process for the variable definition. This helps to trace the calculation of the risk and increases the credibility of the risk management process. Still, a lot of the defined values are based on assumptions. Common assumptions can include that the employees of the organization are always acting beningly (non-malicious) or the underlying infrastructure has no known vulnerabilities. For this process, it is important to justify the selected values. Therefore, the question “Why is this rating at level X?” should be always answerable.

Measuring the Likelihood

Measuring the likelihood is done with the previously defined metrics. Here, the questions to define the factors of the risk’s likelihood (threat actor and vulnerability) need to be answered.

Threat actor	Vulnerability
Skill level: What is the skill of the threat actors?	Ease of discovery: How easily can the by the threat actor?
Motive: How high is the motivation of the threat actors?	Ease of exploit: How easily can the vulnerability be exploited by the threat actor?
Opportunity: Which resources and what access is required?	Awareness: How well-known is the vulnerability to the threat actor?
Size: How large is the group of threat actors?	Intrusion detection: How easily can the exploit be detected?

The answers to these questions are inserted into a table or a formular. The values are then added and divided by the number of values to obtain the average. The following are examples of such tables.

Table 7: Threat Actor Factor Example

Skill level	Motive	Opportunity	Size
6	9	7	6
= 7 (HIGH)			

Source: Created on behalf of IU (2021), based on Williams (2020). CC BY-SA 4.0.

Table 8: Vulnerability Factor Example

Ease of discovery	Ease of exploit	Awareness	Intrusion detection
3	1	4	9
= 4.25 (MEDIUM)			

Source: Created on behalf of IU (2021), based on Williams (2020). CC BY-SA 4.0.

With the two values, the likelihood can now be calculated. This is also done with the sum of both the threat actor factor and the vulnerability factor. These are then divided by two to obtain the likelihood. In this case, it is $\text{Likelihood} = (7 + 4.25)/2 = 5.625$ (medium).

Measuring the Impact

Measuring the impact is done the same way as the likelihood. First, the questions used to define the metric need to be answered.

Technical impact	Business impact
Loss of confidentiality: How sensitive is the data and how much can be disclosed?	Financial damage: How much financial damage could cost this threat?
Loss of integrity: How much corruption damage can be done?	Reputation damage: How much would the reputation suffer?
Loss of availability: How much can the service be disrupted?	Non-compliance: How much does this threat violate compliance?
Loss of accountability: Can the threat actor be traced to an individual?	Privacy violation: Could the threat result in a privacy violation and how any would be affected?

Again, the average is calculated for the values of the answers from these questions. The following tables show an example.

Table 9: Technical Impact Example

Loss of confidentiality	Loss of integrity	Loss of availability	Loss of accountability
9	5	1	7

Loss of confidentiality	Loss of integrity	Loss of availability	Loss of accountability
= 5.5 (MEDIUM)			

Source: Created on behalf of IU (2021), based on Williams (2020). CC BY-SA 4.0.

Table 10: Business Impact Example

Financial damage	Reputation damage	Non-compliance	Privacy violation
7	9	5	3
= 6 (MEDIUM)			

Source: Created on behalf of IU (2021), based on Williams (2020). CC BY-SA 4.0.

After addressing the factors of the impact, the overall impact can be calculated. This is also done by calculating the arithmetic average of the technical and the business impact. In this case, *Impact=5.5+6/2=5.75*(medium).

Since the business impact is usually more critical than the technical impact, it can be a more weighted calculation. This means that the business impact factor is stronger than the technical impact factor. It is also possible to use only the technical impact as a factor for the resulting risk.

Calculating the Overall Risk

After the likelihood and the impact are calculated, the overall risk level can be identified. This can be done with a matrix of the calculated likelihood and impact. The inputs of this matrix are the levels of the two factors (low, medium, or high). The following figure shows the matrix with the resulting risk level.

Figure 6: Overall Risk Level (Risk Severity)

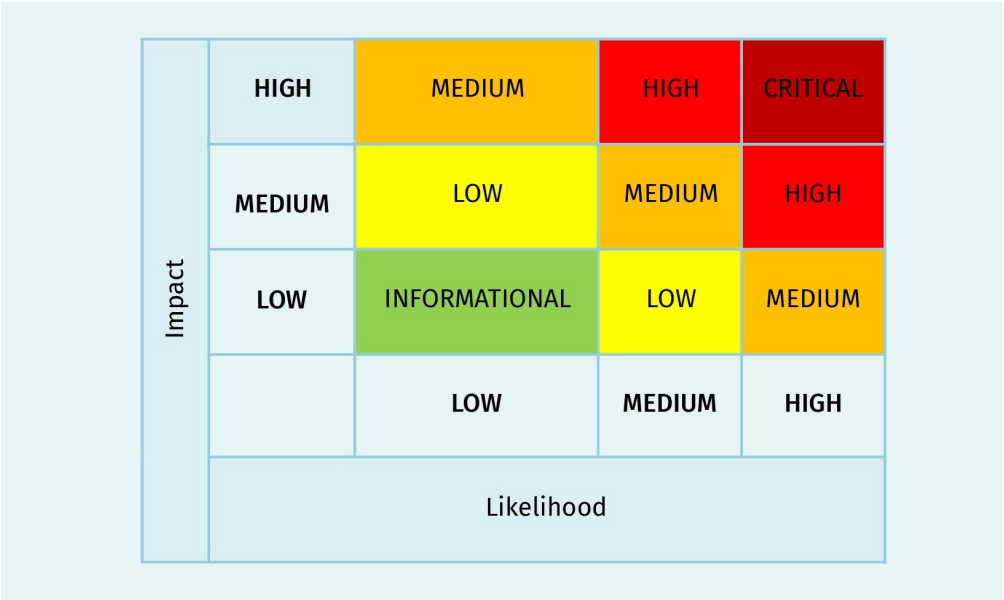

Source: Created on behalf of IU (2021), based on Williams (2020). CC BY-SA 4.0.

In the example, the likelihood was medium, and the impact was medium. The resulting risk level would then be medium.

2.4 The Likelihood of Major Cyber Attacks

The likelihood metrics presented above can also be used to classify major cyber attacks in terms of their likelihood. This is useful to understand how the classification works. To classify a cyber attack, the attack vector needs to be known. Public cyber attacks are often followed up by an analysis of the attack. This analysis can be used to define the likelihood of such an event. The following is an example of how to analyze such an attack.

Equifax Data Breach in 2017

In 2017, Equifax, a consumer credit reporting agency in the United States, was breached and unauthorized actors gained access to customer data. The data breach affected the data of 143 million US customers (Equifax, 2017).

Scenario

The main entry point for the data breach was an unpatched vulnerability in Apache Struts (CVE-2017-5638). The hackers used this vulnerability to gain access to internal servers of the corporate network of Equifax. They gathered internal information including, employee credentials from the system and used those to access the network. They then scanned and exfiltrated information undetected for 76 days. To mask their activities, the intruders encrypted the data and only exfiltrated small archives (Equifax, 2017).

Likelihood analysis

To analyze this cyber attack, the likelihood value of the threat actor is described.

Skill level

What is the skill of the threat actors? As the scenario describes, the threat actor exploited a known vulnerability and scanned the network. The threat actor then also encrypted data to stay undetected. This leads to the conclusion that the threat actor had knowledge about networking and knew how to evade detection mechanisms. This is a threat actor with security penetration skills (9).

Motive

How high is the motivation of the threat actors? Equifax processes personal data from a lot of people. These data can be sold on the darknet. If the attack succeeds, a high reward (9) is expected.

Opportunity

Which resources and what access is required? The website was accessible from the internet which was vulnerable to exploitation. Therefore, no access was required. Since the assets were not stored on the server directly, but rather in the network, the threat actor needed additional resources to access these data. Hence, access and resources were required (7).

Size

How large is the group of threat actors? The data exfiltration was done by an anonymous threat actor. This means the threat actor is an anonymous group of internet users (9).

The selected values for the properties of the threat actor can now be put into a table for calculation. The result is then the overall likelihood of *8.5* for the threat actor.

Table 11: Equifax Threat Actor Factor

Skill level	Motive	Opportunity	Size
9	9	7	9
= 8.5 (HIGH)			

Source: Created on behalf of IU (2021), based on Williams (2020). CC BY-SA 4.0.

After calculating the threat actor factor, the vulnerability factor can be calculated. This is again done by answering the question about the reviewed vulnerability.

Ease of discovery

How easily can the threat actor find the vulnerability? The exploited vulnerability is known, and scanner or untargeted attacks were available. This means that there are automated tools available to detect the vulnerability (9).

Ease of exploit

How easy can the vulnerability be exploited by the threat actor? Exploits for this vulnerability were available on the internet. The exploitation of the internal network, however, required manual work. This means that it is somewhat difficult (3) to exploit this vulnerability to exfiltrate data.

Awareness

How well-known is the vulnerability to the threat actor? This vulnerability had a Common Vulnerabilities and Exposures (CVE) assigned, and patches were available. Consequently, it is easy to say that such vulnerabilities were public knowledge (9).

Intrusion detection

How easily can the exploit be detected? As reported, the security team of Equifax noticed malicious traffic in the network and started investigating. They only noticed it in traffic logs. This means that there were no application logs. The intrusion was logged and reviewed (3).

Given the values, the vulnerability factor result is (6).

Table 12: Equifax Vulnerability Factor

Ease of discovery	Ease of exploit	Awareness	Intrusion detection
9	3	9	3
= 6 (High)			

Source: Created on behalf of IU (2021), based on Williams (2020). CC BY-SA 4.0.

Now that the likelihood factors have been calculated, the overall likelihood can be calculated, resulting in a likelihood of 7.25 (HIGH). This likelihood value should have alerted risk management because patching the system was highly advised. Reasons why this particular system was not patched are unknown.

2.5 Black Swan Events

In the past, black swans were not known to exist in the wild. When a black swan was seen for the first time, it was very surprising. Indeed, this event was devastating for people, as it turned their understanding of nature upside down (Taleb, 2007). This phenomenon is translated to risk management. In risk management, an event with devastating impact and a likelihood close to zero is known as a black swan event. This event is a surprise for the involved parties. In the U.S., the events of September 11, 2001 could be classified as a black swan event. No one thought, at that time, that a terror attack of that magnitude could happen in the US. The impact of this attack was disastrous. Thousands died or were injured in the attack. The US started the "War on Terror" with more casualties and even more financial costs. The attack had a major economic impact, with the stock market temporarily closing and losing value.

In cyber risk management, black swan events are also possible. Some events might have the impact of a black swan event but are not classified as such because the likelihood of the event is higher. This is the major difference between a black swan event and a cyber catastrophe. Just like a black swan event, a cyber catastrophe has a major impact on the affected parties, but its likelihood is not close to zero. For example, the ransomware attacks at WannaCry are considered a cyber catastrophe as they became common in 2017, disrupting big corporations and national entities (Hern, 2017). The attacker encrypts the data on the devices and demands a ransom to unencrypt the data. For the victims, this attack is devastating but not unlikely. WannaCry accesed the operating system to spread and encrypt files. These vulnerabilities were there before, with the presumption that someone would exploit them (Microsoft, 2017).

Classifying cyber events as black swan events is rather complicated as most catastrophes are known to be possible. An example of this is the supply chain attack on SolarWinds, the impact of which can be classified as devastating. The attacker gained access to the update server of SolarWinds and infected the updates with malware. These infected updates were then deployed to organizations and opened backdoors for the attacker (FireEye, 2020). Before this attack happened, there were proof of concepts available showing that supply chain attacks were possible. Furthermore, security experts warned of such attacks. After investigating the incident, it was clear that SolarWinds' risk management failed. No enforced password policy existed for their server and, according to the official statement, an intern was responsible for the weak security (Moore, 2021).

If such a devastating event might not be classified as a black swan event, what is? Such an event can be drawn as a table top simulation. An example target is the financial sector. This guarantees a global impact that likely results in economic disaster. A scenario could be that attackers exploit the software used in the global banking network or the stock markets in the world.

Preparing for a Black Swan Event

The risk management of an organization also needs to plan for such black swan events. The difficulty of such planning is that the events occur unexpectedly, so the likelihood cannot be properly calculated. The best solution for the preparation for a black swan

event is to reduce its possible impact. This requires an in-depth cyber security strategy to lighten the effect of a single event. Besides that, response planning is key to managing such events.

SUMMARY

In risk management, it is important to measure and quantify the risks and threats. To do that, the simple formular $\mathrm{Risk} = \mathrm{Impact} \cdot \mathrm{Likelihood}$ can be used. This formula expresses the risk as the product of the likelihood and the impact.

Both the likelihood and the impact can be measured in terms of their factors. The factors of the likelihood are the threat actor and the vulnerability score of the threat. The factors of the impact are the technical impact and the business impact.

A value can be assigned to each of these factors to calculate the overall likelihood and impact. These values are on a scale from zero to nine. This numeric value can then be compared with other risks. For readability, the numeric values can also be changed into severity categories. These categories are low, medium, and high for the likelihood and impact. For the resulting risk, these metrics are informational, low, medium, high and critical.

These metrics can be used to classify major cyber events. This is basically reverse engineering the event. For this exercise, the known factors of an event can be used to calculate the likelihood of that event. This then can be used to improve the metrics themselves or review the classification of a similar event.

Black swan events are special events. These events are devastating, like a catastrophe, but will occur unexpectedly, meaning that it is difficult to plan for them. The best way to prepare for such events is to reduce their impact and manage them through response planning.

UNIT 3

THREAT MODELING

STUDY GOALS

On completion of this unit, you will be able to …

- draw attack trees.
- describe threats with STRIDE.
- evaluate threats with LINDDUN.

3. THREAT MODELING

Introduction

In order to successfully identify and manage risks, an organization first needs to detect underlying threats. In a sense, threat modeling, the process of identifying and searching for threats, is an art. Similarly to how an artist has an arsenal of tools at their disposal, a threat analyst employs various tools and methodologies to model threats. These tools can vary from simple requirement catalogs and questions for the threat analyst to consider, to full methodologies used to model threats. Selecting the correct tool is crucial to making the process as simple as possible. This unit will introduce various threat modeling methodologies and explain how to use them.

3.1 Attack Tree Methodology

The attack tree methodology is one of the more comprehensive methods used to model threats. Attack trees are graphical representations that hierarchically illustrate threats and determine the probability that attacks might succeed. Bruce Schneier (1999) defined attack trees as "a formal, methodical way of describing the security of systems, based on varying attacks […] represent attacks against a system in a tree structure, with the goal as the root node and different ways of achieving that goal as leaf nodes" (Enter Attack Tree section).

Various Forms of Attack Trees

Attack trees can take a number of forms, the most common of which are lists and graphs. The former involves an attack tree written as a list with indexes. Each index represents an attack on the target system. The root element is the attacker's overall goal, as exemplified in the following list.

ATTACK TREE EXAMPLE

1. Open safe
   - Open safe door
   - Eavesdrop on password combination
   - Pick the safe lock
2. Open with force (&)
   - Access to power-tools
   - Break through safe wall

An attack tree can also take the form of a graph with a root node and an infinite number of child nodes. In this format, the graph defines the root attack and the steps required for the specified attack to succeed (refer to graph below).

Figure 7: Graphical Attack Tree: Opening a Safe

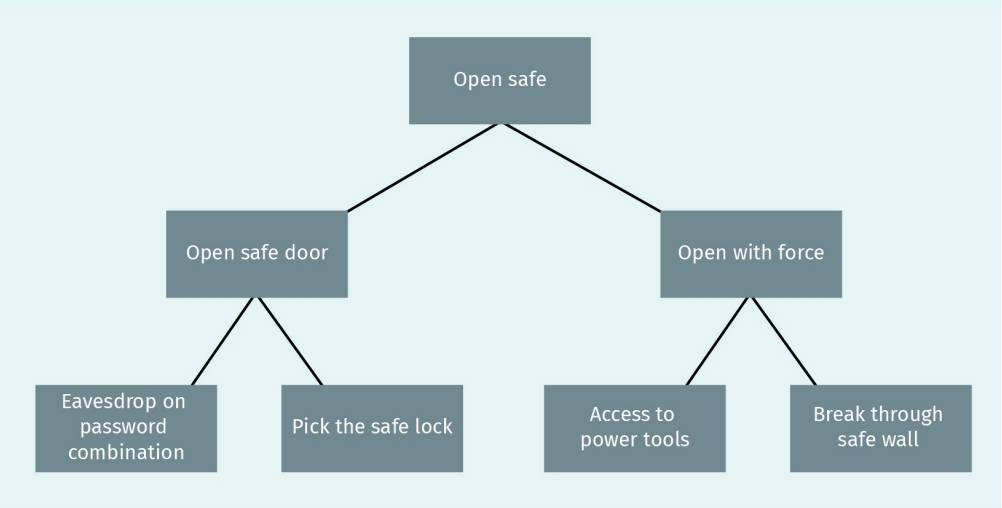

Source: Created on behalf of IU (2021), based on Schneier (1999).

While both versions contain the same structure and information, the graph provides a more comprehensive and legible representation. The list version can be easily converted into the graph version and vice versa.

Definition of an Attack Tree

In order to create an attack tree for a specific threat, we must first understand the basic principle of attack trees. In principle, an attack tree consists of two elements: nodes and edges. Nodes are the attack actions of the tree (e.g., open the safe door). The nodes are connected with edges that represent various conjunctions. A simple tree may feature two different conjunctions: OR or AND. These conjunctions illustrate the relationship between the children of a node. In the following examples, the OR conjunction is represented with a normal edge between two nodes. The AND conjunction is represented by a dot at the beginning of the edge between two nodes in the graphical representation, or as an ampersand (*&*) in the list representation.

Drawing an Attack Tree

When drawing a graph, we first decide on which attack path or threat needs to be modeled for a specific system. This threat provides the root node of the attack tree. All further nodes are then children of the root node. Each attack tree can only have one root node. If different attacker goals need to be modeled, different attack trees will need to be drawn. Developing attack trees is an iterative process for each node. Once a node is drawn, the threat analyst needs to identify its subnodes and then move on to the next one.

Figure 8: Process for Drawing a Node

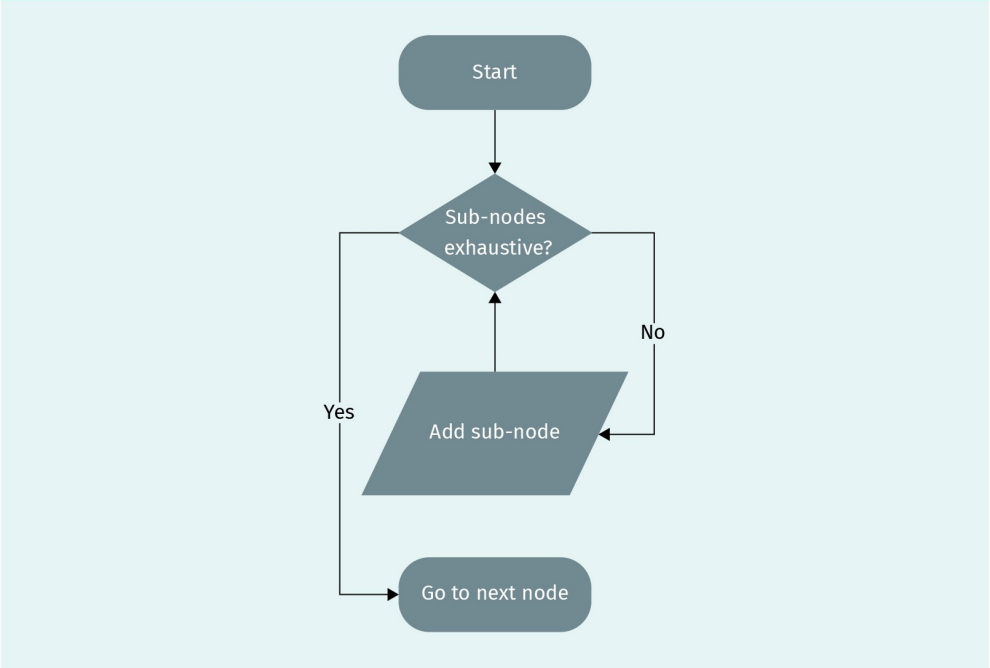

Source: Created on behalf of IU (2021).

It can often be difficult to identify the subnodes for each node. In this case, the threat analyst can contemplate a number of questions to find more subnodes. Some questions include

- What is required to successfully execute the described attack?
- Is this attack subject to certain constraints?
- Can this attack be executed in a number of ways?

After a certain number of subnodes, the attack tree could be considered complete. In reality, however, an attack tree is never truly complete; the threat analyst can always add more detail to a node or think of new ways to achieve the root goal. It is therefore important that the threat analyst determines the sufficient level of completeness of the tree for their particular case. For instance, a high-level attack tree for an entire organization does not need to contain technical details on how a specific system might be compromised. As a rule, an attack tree should always fit on one page to ensure it remains legible and clear.

Once the attack tree is deemed complete, the threat analyst can start pruning the branches, i.e., going through the nodes and checking them for duplicates. Attacks considered impossible in the scope of the tree also need to be marked. They must not be erased to ensure the tree remains complete.

Expansion of Attack Trees

The attack tree explained above covers the root form of a tree. However, an attack tree can be extended in terms of function and used by adding threats. It can also be converted into an attack-defense tree.

Adding threat metrics

Threat metrics can be used to characterize the attack scenario (the node) in quantitative terms. These metrics can be assigned to leaf nodes and, subsequently, further propagated for the remaining nodes. The analyst can decide which threat metrics are used, typically starting with metrics used to determine the probability of being attacked and the potential impact thereof. The probability defines the likelihood that this attack on the specific node might succeed, and the impact describes the effect on the system if the attack is successful.

Attack-defense trees

Attack-defense trees provide an extension to the concept of attack trees with the addition of defenses or countermeasures to the nodes. These defense nodes offer protection against the threat represented by the specific attack node. Once the attack-defense tree is complete, each attack node should have a countermeasure or be mitigated in the subtree. The system will then be protected against all identified attacks. Action is required if a subtree lacks countermeasures or defenses.

3.2 STRIDE

Another methodology used to model threats is STRIDE. Developed by two Microsoft engineers, Loren Kohnfelder and Praerit Garg, STRIDE was initially used to identify threats in Microsoft products (Kohnfelder & Garg, 1999). The mnemonic describes the following threats (Shostack, 2014):

- **Spoofing** is pretending to be someone or something you are not.
- **Tampering** is modifying something you're not supposed to modify.
- **Repudiation** means claiming you didn't do something (regardless of whether you did or not).
- **Information disclosure** is exposing information to people who are not authorized to see it.
- **Denial of service** are attacks designed to prevent a system from providing service, including crashing it, making it unusably slow, or filling all its storage.
- **Elevation of privileges** is when a program or user is technically able to do things that they're not supposed to do.

Each of the STRIDE threats is the direct opposite of a crucial property that a software or a system should have. The pairings are as follows:

- **Spoofing ↔ Authentication:** A system should authenticate the user input and validate the user.
- **Tampering ↔Integrity:** A system should prevent the unauthorized manipulation of data and protect the integrity of its data.
- **Repudiation ↔ Non-Repudiation:** A system should be able to track users and administration action so they can be traced or validated if needed.
- **Information Disclosure ↔Confidentiality:** A system should be able to protect its data from unauthorized access.
- **Denial of Service ↔ Availability:** A system needs to be designed in such a way that malicious user interactions cannot impact the system's ability to serve its purpose.
- **Elevation of Privileges ↔ Authorization:** A system should always check if a user is authorized to execute the requested action. Non-authorized users must be blocked from executing actions.

Using STRIDE to Identify Threats

To identify threats using STRIDE, the threat analyst needs to consider all threats that may compromise the system. The precise way in which they pose a threat to the system can be defined at a later point, e.g., "someone may be able to impersonate an administrator." At first, this might be considered a valid risk, but it remains vague. To further investigate this threat, the threat analyst can consult the system architects or developers. If no one is able to provide a reason to explain why this incident would not occur, the threat analyst has found a verified threat. In certain cases, the threat analyst may be met with the response "no, this could not happen because we always validate user authentication." This means that a mitigation measure is in place for that threat, which can and should be tested to validate the statement.

Threat modeling diagrams

UML
The Unified Modeling Language is a general-purpose modeling language to visualize systems in software engineering.

It can be difficult to identify threats without an overview of the entire system environment, and diagrams often act as helpful threat modeling tools. Unified modeling language (**UML**) diagrams, such as swimlane or state diagrams, can help us gain an understanding of how a system operates (Booch et al., 2005). The most important diagram for threat modeling is the data flow diagram (DFD). This diagram contains all information on entities and their relationship within the system. The DFD helps the threat analyst to identify links and threats to all entities used in the system. A data flow diagram comprises the following elements:

- **process,** which is any executed code controlled by the system.
- **data flow**, which is any data flow between processes, data stores, or external entities.
- **data store**, which is any sub-system that stores data controlled by the system.
- **external entity**, which is any external entity (user or system) that interacts with the system but is not controlled by the system itself.
- **trust boundary**, which is a boundary between two entities that marks a change of trust (e.g., corporate network to the internet)

Figure 9: DFD Symbols

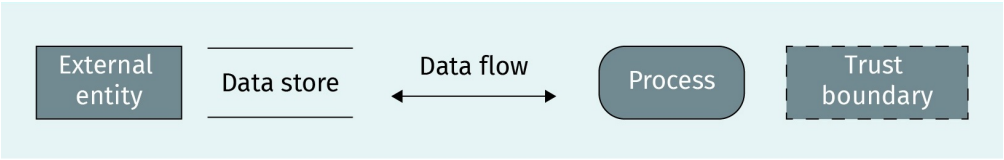

Source: Created on behalf of IU (2021), based on Shostack (2014).

A DFD needs to include everything relevant to the system and everything that has relevant interactions with it. In principle, the DFD defines how the system works, who it interacts with, and how all communications and the data store run. To ensure that a DFD can be used to model threats, the following rules need to be applied when creating the DFD (Howard & LeBlanc, 2009):

- A process needs to connect to at least one data flow.
- A data flow needs to begin or end at a process.
- A data store needs to be connected to a process with a data flow.
- Data stores can only connect to a process.

Figure 10: Example DFD

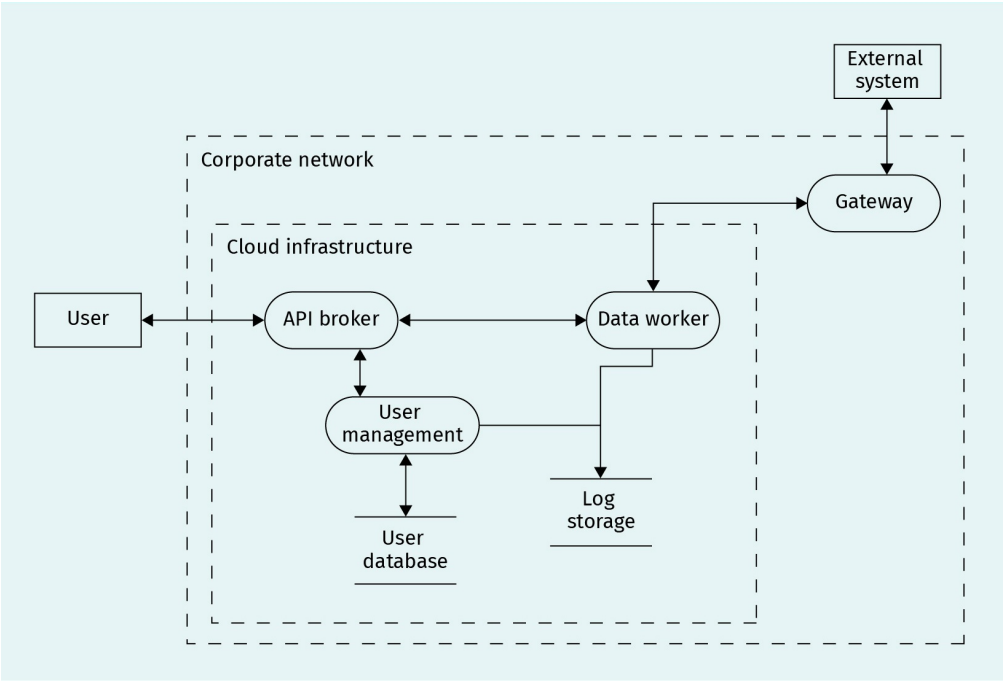

Source: Created on behalf of IU (2021).

Spoofing

Spoofing is the act of pretending to be someone or something else, i.e., “Can I trick the system into thinking I am someone else?” In the example DFD, a threat may be posed by the user spoofing their identity to subvert user management. The external system may also succeed in convincing the gateway that it is another system. Technical vulnerabilities include the following:

- The user ID is part of the request and can be changed.
- An external system is contacted via a subdomain that can be hijacked.

Tampering

Tampering involves changing or modifying data that should not be changed or modified by the user, i.e., “Can I trick the system into modifying specific files, or can I modify them myself?” In the example DFD, tampering threats might include

- modifying the data worker to gain access to different external systems.
- modifying the user database to change values.

Repudiation

Repudiation involves the acceptance or denial of responsibility of specific users or systems for performing certain actions. The goal of the attacker is to prevent the traceability of these actions, i.e., “Can I trick the system into thinking I did or did not do something?” In this example, repudiation threats might include

- User logging is manipulated to prevent the logging of certain actions.
- A user can trigger log events without performing the actual action.
- It is not possible to trace which user accessed which application programming interface (API).

Information disclosure

Information disclosure involves the extraction of information the user is not authorized to access from the system, i.e., “What data can I extract and how?” In the DFD, information disclosure threats might be posed by

- finding a bug in the API to extract the user data.
- receiving error messages that might expose credentials or other sensitive information.

Denial of service

A denial-of-service (DoS) attack is a type of cyber attack in which the attacker uses resources to disrupt the availability or costs of the system, i.e., “How can I overflow the system?” In the example DFD, threats might be posed by

- flooding the API broker with a network DoS.

- generating unnecessary log events to flood the log storage to consume unnecessary cloud resources.

Elevation of privilege

Elevation of privilege is the act of exploiting a system to gain elevated access for a certain user or application. Privileges are typically escalated via broken authentication and access control, or disrupting and corrupting processes, i.e., “What can I manipulate to gain more privileges?” In the example DFD, threats might be posed by

- using vulnerabilities in the user management to elevate privileges.
- breaking out of the data worker process to execute arbitrary code.

STRIDE Variants

STRIDE can be applied to a wide range of threats and provides an important mnemonic for threat modeling. The basic STRIDE variant has no restrictions in terms of identifying threats, which can often prove challenging for the threat analyst and other participants. As a result, a number of STRIDE variants have been developed to facilitate a simpler approach to identify threats. We will now explore two of these variants: STRIDE-per-Element and STRIDE-per-Interaction.

STRIDE-per-Element

STRIDE-per-Element follows every element in the data flow diagram with an emphasis on which threats are prevalent for each element. The following table illustrates the entities that may be subverted by the defined attacks. The question mark in the data store row indicates that data stores may face attacks if they store logs. The table shows that the external entity may become a victim of spoofing, but not of denial-of-service attacks, for example.

Table 13: STRIDE-per-Element

	S	T	R	I	D	E
External entity	X		X			
Process	X	X	X	X	X	X
Data flow		X		X	X	
Data store		X	?	X	X	

Source: Created on behalf of IU (2021).

One weak point of the STRIDE-per-Element approach, as exemplified in the table, is that it does not provide a full representation of the use case with the required scope. This can be attributed to the fact that STRIDE-per-Element was developed by Microsoft and primarily tailored to its needs (Shostack, 2014). The STRIDE-per-Element table may therefore need to be adapted to the organization's needs.

STRIDE-per-Interaction

STRIDE-per-Interaction focuses on the interaction between entities instead of each individual element. Nevertheless, both STRIDE variants will identify the same number of threats. STRIDE-per-Interaction also uses a table to identify threats for each interaction. This table includes the following rows:

- reference number
- entity
- entity interaction
- STRIDE threats for this interaction

The STRIDE-per-Interaction table below features examples from the previously used DFD. X is used to map potential threats to interactions. The following example only uses the gateway, data worker, and log storage entities.

Table 14: STRIDE-per-Interaction

	Element	Interaction	S	T	R	I	D	E
1	Process (gateway)	The process sends outgoing data flows to an external entity.			X			
2		The process receives data from an external entity.	X				X	X
3		The process sends data to another process (data worker).	X		X	X	X	X
4		The process receives data from another process (data worker).	X		X		X	X
5	Process (data worker)	The process sends data to another process (gateway).	X		X	X	X	X
6		The process receives data from another process (gateway).	X		X		X	X
7		The process sends data to another process (API broker).	X		X	X	X	X
8		The process receives data from another process (API broker).	X		X		X	X
9		The process sends data to a data store (log storage).	X			X		

	Element	Interaction	S	T	R	I	D	E
10	Data flow (gate-way/ external)	Crosses environment boundary		X		X	X	
11	Data flow (gate-way/data worker)	Crosses environment boundary		X		X	X	
12	Data store (log stor-age)	The data store has an inbound data flow.		X	X	X	X	
13	External interac-tor (external system)	The external interactor converts inputs to processes.	X		X	X		
14		The external interactor receives inputs from processes.	X					

Source: Created on behalf of IU (2021).

3.3 LINDDUN

LINDDUN is a methodology used to model privacy aspects in a system (DistriNet Research Group, 2020d). LINDDUN focuses on the systematic elicitation and mitigation of privacy threats in a system. Similarly to STRIDE, LINDDUN is a mnemonic for the privacy threat categories it supports (Sion et al., 2018, p. 2).

Linkability	An adversary is able to link two items of interest without knowing the identity of the data subject(s) involved.
Identifiability	An adversary is able to identify a data subject from a set of data subjects through an item of interest.
Non-repudiation	The data subject is unable to deny a claim (e.g., having performed an action or sent a request).
Detectability	An adversary is able to determine whether an item of interest about a data subject exists, regardless of being able to read the contents.
Disclosure of information	An adversary is able to learn the content of an item of interest about a data subject.
Unawareness	The data subject is unaware of the collection, proc-essing, storage, or sharing activities (and corre-sponding purposes) of the data subject's personal data.

Non-compliance	The processing, storage, or handling of personal data does not follow legislation, regulation, and/or policy.

Difference between Privacy and Security Threat Modeling

Security threat models focus on assets and how they can be protected from (external) threats. However, privacy threat models take a different approach. Instead of focusing on the system assets, privacy threat models focus on the data subjects' data in order to protect the **items of interest (IOI)** and not the system. Users' data privacy is at risk from both (external) attacks and the system, if it misbehaves.

Items of interest (IOI)
An item of interest is a piece of data that is protected by privacy law and linked to a person.

Using LINDDUN to Identify Threats

LINDDUN offers more extensive threat modeling than other approaches, such as STRIDE, which does not cover privacy threats. The LINDDUN methodology consists of three main steps:

1. System modeling (modeling the system and drawing a DFD as a basis for threat elicitation)
2. Threat elicitation (using the LINDDUN mnemonic to identify threats)
3. Threat management (finding suitable mitigation measures for the discovered threats)

System modeling

The first step in LINDDUN threat modeling involves gaining an understanding of the system in order to analyze it. This requires the same approach as STRIDE, i.e., drawing a DFD. This DFD then provides the baseline for the scope of the analysis and threat identification.

Threat elicitation

Once the DFD has been produced, the threats need to be identified. Similarly to the STRIDE approach, LINDDUN also provides a table to map potential threats to each DFD element (marked with an X in the following table).

Table 15: LINDDUN Threats for Each Element

	L	I	N	D	D	U	N
External entity	X	X				X	
Process	X	X	X	X	X		X
Data flow	X	X	X	X	X		X
Data store	X	X	X	X	X		X

Source: Created on behalf of IU (2021).

Based on the table, the analyst can now determine whether they pose a threat for the DFD element. LINDDUN also provides a catalog of threat trees structured according to the threat category, which contains the most common attack paths for the particular threat category and the respective link to each DFD element. The following table illustrates the consequences and the **impact actions** defined by the LINDDUN framework for each threat. The information in the following tables can be found on the LINDDUN website.

Impact actions
An action that improves the situation [+] or worsens the situation [-] is an impact action.

Table 16: Linkability

Consequences	• "Can lead to identifiability (see Identifiability trees) when too much linkable information is combined" • "Can lead to inference: when "group data" is linkable, this can lead to societal harm, like discrimination"
Impacted by	**+** "Data minimization: the less info is available, the better" **-** "Identifiability: if the subject's identity is known, all related data can obviously be linked" (DistriNet Research Group, 2020e)

Source: Created on behalf of IU (2021), based on DistriNet Research Group (2020e).

Table 17: Identifiability

Consequences	• "Can lead to severe privacy violations (when subject assumes they are anonymous)"
Impacted by	**+** "Data minimization: the less info is available, the better " **-** "Linkability: the more information is linked, the higher the chance the combined data are identifiable (the more attributes are known, the smaller the anonymity set)" (DistriNet Research Group, 2020c)

Source: Created on behalf of IU (2021), based on DistriNet Research Group (2020c).

Table 18: Non-Repudiation

Consequences	• "Accountability: when a person is not able to repudiate an action or piece of information, they can be held accountable."
Impacted by	**-** "Identifiability: if data are identifiable, it will be hard to repudiate" (DistriNet Research Group, 2020g)

Source: Created on behalf of IU (2021), based on DistriNet Research Group (2020g).

Table 19: Detectability

Consequences	• "Inference: by detecting whether an IOI exists, one can deduce certain information, even without actually having access to that information" (DistriNet Research Group, 2020a)
Impacted by	

Source: Created on behalf of IU (2021), based on DistriNet Research Group (2020a).

Table 20: Disclosure of Information

Consequences	• Can lead to severe privacy violations (when personal data are leaked)
Impacted by	**+** If the system is properly secured it is harder that data might be leaked (DistriNet Research Group, 2020b)

Source: Created on behalf of IU (2021), based on DistriNet Research Group (2020b).

Table 21: Unawareness

Consequences	• "Linkability/identifiability: the more information is available, the easier it can be linked (and identified)" (DistriNet Research Group, 2020h)
Impacted by	

Source: Created on behalf of IU (2021), based on DistriNet Research Group (2020h).

Table 22: Non-Compliance

Consequences	• Can lead to fines (when violating legislation, or not adhering to the communicated corporate policies)" • "Can lead to loss of image, credibility, etc."
Impacted by	**+** "Security officer/legal audit/[…]: a person responsible for the system's compliance" **-** "Tampering with the policy data store: when the policy database is not tamper-proof, attackers can alter the access control policies and user consents of the system" (DistriNet Research Group, 2020f)

Source: Created on behalf of IU (2021), based on DistriNet Research Group (2020f).

Each identified threat needs to be documented in the attack trees provided by the framework. Other documentation, such as databases, can also be used for this purpose.

Threat management

Once all threats have been identified, they need to be rated and prioritized. LINDDUN does not provide a specific risk assessment model for rating threats. The Open Web Application Security Project (OWASP) Risk Rating Methodology can be used to rate the risk of each threat, though it may need to be adapted to privacy threats. Mitigation steps can be planned once the risks have been rated and prioritized. LINDDUN offers a number of mitigation strategies and solutions.

The framework provides two mitigation strategies. The first strategy is the proactive approach, which entails controlling associations between users and their actions and personal information to ensure that the user shares as little information as necessary with the system. This follows the principle that "data that are not shared can't be at risk."

The second mitigation strategy is reactive in nature, whereby the data are protected by keeping the associations of data to individuals to a minimum after disclosure by

- removing unnecessary data,
- replacing data with non-linkable objects,
- generalizing the data, and
- hiding connections between IOIs.

IOIs can be protected, both on a data level and on a technical or system level. This highlights the close link between data protection and system security, namely, data cannot be protected without a secure system.

SUMMARY

Threat modeling is used to find threats to assets in a defined scope. A number of different tools or frameworks can be used to model threats.

One of these is an attack tree: a verbose method used to hierarchically illustrate threats and determine the probability that attacks might succeed. An attack tree features a root node, which defines the attacker's goal. The children of the root node further break down the attacker's goal into subgoals. Different conjunctions can be used to connect children nodes to their parent node. The attack tree can be drawn as a list with indexes or a graphical tree. It can also be expanded with the addition of defense nodes to form an attack-defense tree to include defensive actions that can be taken to mitigate attacks.

Another threat modeling methodology is STRIDE, a mnemonic for the following threat categories: spoofing, tampering, repudiation, information disclosure, denial of service, and elevation of privileges. The STRIDE approach employs data flow diagrams (DFD) to identify threats. DFDs

represent the logical data flow in the scope. Analysts use the DFD to identify threats for each entity and subsequently map them to the STRIDE threat categories.

The LINDDUN framework can be used to model privacy threats. LINDDUN is also a mnemonic for the privacy threat categories: linkability, identifiability, non-repudiation, detectability, disclosure of information, unawareness, and non-compliance. This methodology also employs DFDs to identify threats. Additionally, LINDDUN provides a number of mitigation strategies and solutions for the determined threats.

UNIT 4

STANDARDIZATION AND COMPLIANCE

STUDY GOALS

On completion of this unit, you will be able to …

- describe the NIST Risk Management Framework process.
- understand the content of the ISO/IEC 27005.
- explain the objectives of BSI Standard 100-3.

4. STANDARDIZATION AND COMPLIANCE

Introduction

Decision-makers require a repeatable and standardized framework on which they can rely. As a solution, standardization offers a unified approach for products and processes. We encounter the concept of standardization on a daily basis in the modern world. Standardized items can be found all around us, e.g., electrical outlets and traffic regulations. They help to make our lives easier. The field of IT is also governed by a number of risk management and risk assessment standards that define the necessary processes and measures an organization needs to enforce to ensure compliance with the respective standards. The majority of standards provide a high-level overview and define the overall processes that need to be implemented by organizations in accordance with the standard. Three of the most common standards are ISO/IEC 27005, NIST Special Publication 800-37 and, in Germany, BSI Standard 100-3.

4.1 NIST Risk Management Framework

NIST
The National Institute of Standards and Technology (NIST) is a non-regulatory agency of the United States Department of Commerce. This agency provides public standards and guidelines (e.g., for cyber security topics).

The National Institute of Standards and Technology (**NIST**) Special Publication 800-37 provides a risk management framework for information systems and organizations. The latest version of this framework, Revision 2, was published in 2018 (NIST, 2018). The publication was developed with the aim of providing a practical guide for risk management within organizations. It contains comprehensive information on how to apply a risk management framework (RMF) and explores risk management processes. The following sections provide an overview of NIST Special Publication (SP) 800-37.

Fundamentals

This section describes the basic concepts behind risk management for cyber security and privacy with an emphasis on the following topics (NIST, 2018, Table of contents):

- organization-wide risk management,
- risk management framework steps and structure,
- information security and privacy in the RMF,
- system and system elements,
- authorization boundaries,
- requirements and controls,
- security and privacy posture, and
- supply chain risk management.

Organization-wide risk management

Risk management is an organization-wide task. It starts at the highest level within an organization and filters down to the system engineering level. The risk management process can be split into three levels:

1. **Organization.** The organization level constitutes the highest level. The organization-wide risk management processes are managed at this level.
2. **Mission/business process.** Activities conducted at this level do not address technical risks, but instead focus on the management of product or project risks.
3. **Information system.** This level addresses risk from the perspective of the information system with an emphasis on the technical and organization risks faced by an information system.

All three levels are interconnected; risks are communicated and reported between the levels, and risk management is performed across all levels.

Figure 11: Risk Management Pyramid

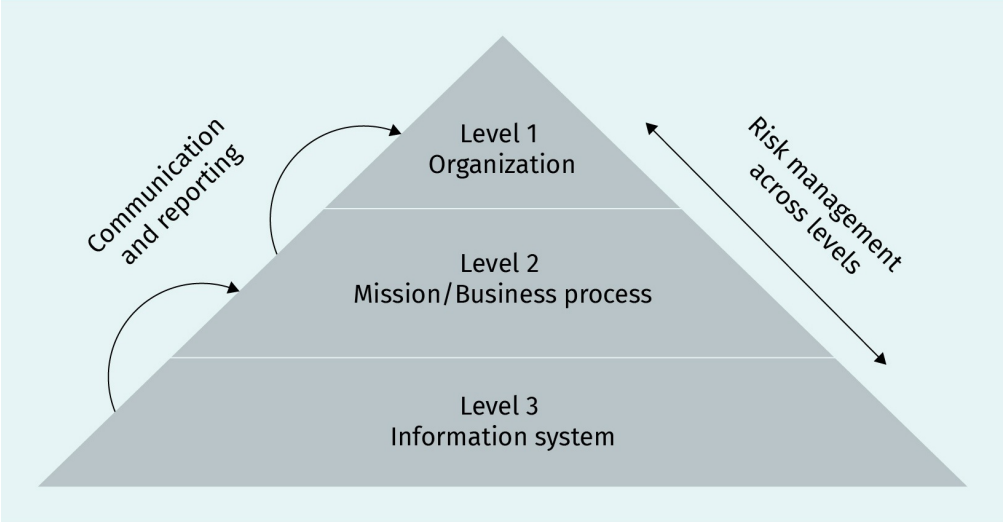

Source: Created on behalf of IU (2021), based on NIST (2018).

The activities at level 3 rely on the preparation activities conducted at levels 1 and 2. Without an organization-wide risk management strategy, information system level risk management will not be very useful for the organization. Risk management is not an isolated activity and requires cooperation at all levels.

NIST SP 800-37 lists activities performed at levels 1 and 2 to prepare the organization to execute the RMF (NIST, 2018, p. 7):

- "assigning roles and responsibilities for organizational risk management processes"
- "establishing a risk management strategy and organizational risk tolerance"
- "identifying the missions, business functions, and mission/business processes the information system is intended to support"

- "identifying key stakeholders (internal and external to the organization) that have an interest in the information system"
- "identifying and prioritizing assets (including information assets)"

Risk management framework steps and structure

A management framework requires a defined process to generate measurable results. NIST SP 800-37 defines seven steps for the risk management process (NIST, 2018, pp. 8—9):

1. **"Prepare** to execute the RMF from an organization- and a system-level perspective by establishing a context and priorities for managing security and privacy risk."
2. **"Categorize** the system and the information processed, stored, and transmitted by the system based on an analysis of the impact of loss."
3. **"Select** an initial set of controls for the system and tailor the controls as needed to reduce risk to an acceptable level based on an assessment of risk."
4. **"Implement** the controls and describe how the controls are employed within the system and its environment of operation."
5. **"Assess** the controls to determine if the controls are implemented correctly, operating as intended, and producing the desired outcomes with respect to satisfying the security and privacy requirements."
6. **"Authorize** the system or common controls based on a determination that the risk to organizational operations and assets, individuals, other organizations, and the Nation is acceptable."
7. **"Monitor** the system and the associated controls on an ongoing basis to include assessing control effectiveness, documenting changes to the system and environment of operation, conducting risk assessments and impact analyses, and reporting the security and privacy posture of the system."

Information security and privacy in the RMF

Privacy and information security measures are often contradictory. For example, data retention periods need to be kept as short as possible to safeguard to privacy of log data, yet information security requires storage for as long as possible. An organization needs to introduce risk management to ensure both aspects satisfy the pertinent requirements.

The risk management framework described in NIST Special Publication 800-37 is applicable to both security and privacy risks. Nevertheless, an organization needs to understand the requirements of both aspects and their relationship.

System and system elements

Information systems can be broken down into systems and system elements. ISO 15288 defines a system as a set of interacting elements that are organized to achieve one or more stated purposes (ISO/IEC/IEEE, 2015). Each of these systems has specific capabilities and specific requirements that are fulfilled by the system elements. The RMF ensures that security and privacy requirements are satisfied by information systems throughout the Software Development Life Cycle (**SDLC**) by identifying the risks for each system.

SDLC
The Software Development Life Cycle defines

processes to manage the life cycle of software development.

Authorization boundaries

Authorization boundaries establish the scope of protection for an information system. Sets of system elements defined as an authorization boundary can be directly managed by one system, as the elements included in the authorization boundary generally involve similar risks. The authorization boundaries are defined by the scope of the systems. According to NIST SP 800-37, system elements within the same authorization boundary should have following characteristics (NIST, 2018, p. 17):

- “support the same mission or business functions”
- “have similar operating characteristics and security and privacy requirements”
- “process, store, and transmit similar types of information””
- “reside in the same environment of operation“

Requirements and controls

Requirements and controls are crucial to risk management in terms of understanding the security and privacy requirements for information systems. Requirements include regulations or laws that need to be adhered to (e.g., GDPR), and guidelines on how to protect assets and systems. Whereas controls refer to safeguards and protection mechanisms that can be implemented to protect the security and privacy objectives of an organization.

Security and privacy posture

The purpose of the RMF is to provide authorizing officials with accurate information on the security and privacy posture of the organization’s information systems to facilitate risk-based decisions. Consequently, the security and privacy posture needs to describe the status of the information system, outline how identified risks are managed, and define how the organization will react to changes within the organization or systems.

Supply chain risk management

With the ever-increasing occurrence of supply chain attacks and or supply chain threats, as exemplified by the 2021 Suez Canal obstruction, management of these risks is becoming increasingly important. As a result, the security postures of suppliers form a key aspect of the risk management process. The supply chain risk management (SCRM) policy must cover all potential risks that could impact the organization.

Process

The NIST RMF tasks can be executed concurrently to SDLC processes due to the similarity between the risk management roles and process steps. Concurrent execution of RMF tasks with SDLC processes helps to optimize the efficiency of an organization in terms of security and privacy risk management. Furthermore, outputs required for the RMF, such as risk assessments, mitigation plans, and privacy concepts, can be obtained from the SDLC process. More information on the complete process can be found in NIST Special Publication 800-37.

4.2 ISO/IEC 27005

ISO/IEC 27005 is a standard for information security risk management published by the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC). The official name of this standard is "Information technology—Security techniques—Information security risk management." It is a core part of the ISO/IEC 27000-series of standards, commonly referred to as ISO27k. The current version of this standard is the third version, which was published in 2018 (ISO/IEC, 2018).

ISO/IEC 27005 lists the following targets that need to be achieved by a proper risk management process (ISO/IEC, 2008, p. 4):

- "risks being identified"
- "risks being assessed in terms of their consequences to the business and the likelihood of their occurrence"
- "the likelihood and consequences of these risks being communicated and understood"
- "priority order for risk treatment being established"
- "priority for actions to reduce risks occurring"
- "stakeholders being involved when risk management decisions are made and kept informed of the risk management status"
- "effectiveness of risk treatment monitoring"
- "risks and the risk management process being monitored and reviewed regularly"
- "information being captured to improve the risk management approach"
- "managers and staff being educated about the risks and the actions taken to mitigate them"

Security Risk Management Activities

In order to achieve the defined targets, ISO/IEC 27005 breaks down security risk management activities into six tasks (ISO/IEC, 2008, p. 3):

1. "Context establishment"
2. "Risk assessment"
3. "Risk treatment"
4. "Risk acceptance"
5. "Risk communication"
6. "Risk monitoring and review"

As risk management needs to be included in an information security management system (**ISMS**), the activities can be aligned as follows to the Plan-Do-Check-Act (PDCA) cycle of an ISMS.

Information security management system
An ISMS is a set of policies and procedures to manage security-relevant assets.

Table 23: Alignment of ISMS and Information Security Risk Management Processes

ISMS process	Information security risk management process
Plan	• Establishing the context • Risk assessment • Developing risk treatment plan • Risk acceptance
Do	• Implementation of risk treatment plan
Check	• Continual monitoring and reviewing of risks
Act	• Maintaining and improving the information security risk management process (ISO/IEC, 2018)

Source: ISO/IEC (2018).

Context establishment

The first objective of the standard is to establish the context of information security risk management. This involves determining the basic criteria for this process, defining the scope, and establishing an organization that operates the risk management process. The basic criteria for this process cover evaluating the risk, defining the impact of specific risks, and determining the level of risk the organization is willing to accept, i.e., the risk evaluation criteria, impact criteria and risk acceptance criteria. Defining the scope refers to determining the scope and boundaries of the information security risk management approach. These aspects need to be defined to ensure that all relevant assets are taken into account in the risk assessment.

Risk assessment

The risk assessment process defines how risks should be assessed. ISO/IEC 27005 specifies the following steps in this regard:

- **Risk identification** involves finding and rating potential risks to the organization in terms of possible resulting losses. This involves the following steps (ISO/IEC, 2008, pp. 10—14):
  - “identification of assets”
  - “identification of threats”
  - “identification of existing controls”
  - “identification of vulnerabilities”
  - “identification of consequences”
- **Risk estimation** is needed to plan and assess actions and mitigation measures for risks faced by an organization. The following steps are used to rate the identified risks (ISO/IEC, 2008, pp. 14—15):
  - “assessment of consequences”
  - “assessment of incident likelihood”
- **Risk evaluation** involves comparing the level of risk to the acceptance criteria defined in the context establishment step.

Risk treatment

Following identification and evaluation, a risk can be managed in a number of ways. The following options are specified by ISO/IEC 27005:

- **risk reduction.** The level of risk should be reduced through the selection of controls so that the residual risk can be reassessed as being acceptable.
- **risk retention.** The decision on retaining the risk without further action should be taken depending on risk evaluation.
- **risk avoidance.** The activity or condition that gives rise to the particular risk should be avoided.
- **risk transfer.** The risk should be transferred to another party that can more effectively manage the particular risk depending on risk evaluation.

Risk acceptance

After the identified risks have been treated with the desired mitigation measure, the formal decision needs to be made about whether the residual risk can be accepted by the organization, and documented.

Risk communication

As risks form a crucial part of decision-making within organizations, it is important that they are communicated between the decision-makers and stakeholders. ISO/IEC 27005 lists a number of reasons behind this communication requirement (ISO/IEC, 2008, p. 22):

- "to provide assurance of the outcome of the organization's risk management"
- "to collect risk information"
- "to share the results from the risk assessment and present the risk treatment plan"
- "to avoid or reduce both occurrence and consequence of information security breaches due to the lack of mutual understanding among decision makers and stakeholders"
- "to support decision-making"
- "to obtain new information security knowledge"
- "to co-ordinate with other parties and plan responses to reduce consequences of any incident"
- "to give decision makers and stakeholders a sense of responsibility about risks"
- "to improve awareness"

Risk monitoring and review

The risk management process is never complete. It is an ongoing process that requires the continuous monitoring and review of systems and risks. ISO/IEC 27005 specifies that a risk review must be conducted following each of the system changes listed below (ISO/IEC, 2008, pp. 22—23):

- "new assets that have been included in the risk management scope"
- "necessary modification of asset values, e.g., due to changed business requirements"

- “new threats that could be active both outside and inside the organization and that have not been assessed”
- “possibility that new or increased vulnerabilities could allow threats to exploit these new or changed vulnerabilities”
- “identified vulnerabilities to determine those becoming exposed to new or re-emerging threats“
- “increased impact or consequences of assessed threats, vulnerabilities and risks in aggregation resulting in an unacceptable level of risk”
- “information security incidents”

4.3 BSI Standard 100-3

BSI Standard 100-3 is a German standard published by the Federal Office for Information Security (BSI). The standard forms part of the IT baseline protection (*IT-Grundschutz*) systematic approach to information security issued by the BSI. IT-Grundschutz standards help private and public organizations to minimize their IT risks. Certain public sector requests for proposal (RfP) require compliance with the IT-Grundschutz as a prerequisite for submission of a bid. BSI Standard 100-3 describes how to conduct a risk analysis based on IT-Grundschutz standards (BSI, 2008). The document is divided into seven sections (BSI, 2008, p. 3):

1. Preliminary work
2. Preparing the threat summary
3. Determination of additional threats
4. Threat assessment
5. Handling risks
6. Consolidation of the security concept
7. Feedback to the security process

Preliminary Work

The first section describes the work that needs to be completed prior to a risk assessment. According to the standard, the steps are as follows (BSI, 2008, p. 7):

- “A systematic information security process must have been initiated.”
- “A scope for the security concept must be defined.”
- “A structure analysis must have been performed for the information domain.”
- “An assessment of protection requirements must have been performed.”
- “A modeling process must have been performed.”
- “Prior to the risk analysis, a basic security check must be performed.”
- “A supplementary security analysis must have been performed.”

Preparing the Threat Summary

The threat summary is the first thing viewed by a decision-maker. Accordingly, it is important that it is informative and provides an overview of the security status of the analyzed system. The standard includes work steps for producing a threat summary.

Determination of Additional Threats

The IT-Grundschutz model features a defined set of threats to systems. However, in certain circumstances, additional isolated threats faced by a system may go beyond the scope of those specified in the IT-Grundschutz model. These threats also need to be analyzed.

BSI Standard 100-3 lists a number of questions that should be considered when determining additional threats (BSI, 2008, pp. 12—13):

- "Which potential *force majeure* events represent particular threats for the information domain?"
- "Which organizational failures must be avoided [...] to guarantee information security?"
- "Which human errors adversely affect the security of information and applications?"
- "Which special security problems could occur to the target object under review due to technical failure?"
- "Which particular threats arise from deliberate attacks by outsiders?"
- "How can insiders affect the proper, secure operation of the target object under review through deliberate actions?"
- "Are there objects not located in the information domain being examined that represent special threats?"

Threat Assessment

BSI Standard 100-3 describes checking whether measures already implemented or planned in the security concept, generally standard security measures taken from the IT-Grundschutz Catalogs, offer adequate protection for each threat, or whether gaps exist. The standard defines the following test criteria (BSI, 2008, p. 15):

- "completeness. Do the standard security measures provide protection for all aspects of each threat?"
- "mechanism strength. Do the protection mechanisms recommended in the standard security measures counteract each threat adequately?"
- "reliability. How difficult is it to circumvent the planned security mechanisms?"

Handling Risks

As a rule, the threat assessment generally identifies a number of risks that are not adequately mitigated by the security measures taken from the IT-Grundschutz Catalogs. The standard provides the following strategies for handling a risk (BSI, 2008):

- risk reduction through additional security measures
- risk avoidance through restructuring

- risk acceptance
- risk transfer

Consolidation of the Security Concept

The security concept will need to be consolidated if additional security measures are added to handle the remaining risks. In addition, all risks need to be documented in the security concept, along with their risk handling strategy. The standard provides the following criteria for the consolidation of the security concept (BSI, 2008, p. 21):

- "Inappropriate security measures should be rejected and after a detailed analysis replaced by effective measures."
- "Contradictions or inconsistencies in the security measures should be resolved and replaced by homogenous mechanisms that are coordinated with each other."
- "Security measures that are not accepted by users have no effect. Practical solutions that restrict or hinder users as little as possible should be found."
- "Security measures that are too difficult or costly should either be re-worked or rejected and replaced by appropriate protective measures. [...] measures that are too weak endanger IT security [...] should also be reworked or replaced."

Feedback to the Security Process

Following successful risk analysis, the outcome of the analysis and the security concept can be used to modify the security process. The analysis forms the basis for the following working steps (BSI, 2008, p. 23):

- implementing the security concept;
- reviewing the information security process on all levels; and
- information flow in the information security process.

If compliance with the IT-Grundschutz is required, ISO 27001 certification can also be obtained. "ISO 27001 Certification in compliance with IT-Grundschutz" can also be attested by an auditor.

SUMMARY

NIST Special Publication 800-37 is a standard for risk management in organizations. It describes how to build a risk management framework. The first part of the publication defines the fundamentals of risk management. The second part describes the risk management process as follows: prepare, categorize, select, implement, assess, authorize, and monitor.

ISO/IEC 27005 is a publication in the ISO 27k family of standards, which focuses on information security risk management. ISO/IEC 27005 breaks down the security risk management process into six activities: context establishment, risk assessment, risk treatment, risk acceptance, risk communication, and risk monitoring and review.

BSI Standard 100-3 published by the Federal Office for Information Security (BSI) provides a framework for risk analysis. It describes how risk management should be executed in compliance with the BSI IT-Grundschutz (baseline protection) Catalogs. The BSI Standard 100-3 process features seven steps: preliminary work, preparing the threat summary, determination of additional threats, threat assessment, handling risks, consolidation of the security concept, and feedback to the security process. The emphasis of this process is compliance with the BSI IT-Grundschutz. If an organization is compliant with the BSI 100 series, ISO 27001 compliance can be attested by an auditor.

UNIT 5

RISK ASSESSMENT

STUDY GOALS

On completion of this unit, you will be able to …

- conduct a risk assessment.
- plan a risk assessment.
- factor black swan events into the risk assessment.
- describe continuous re-evaluation.

5. RISK ASSESSMENT

Introduction

Risk management is the process of managing known or unknown risks for an organization. In life, it is easier to manage an issue if there is a known input for the decision. For example, when driving to a friend or family, we give an approximate time of arrival. There might be an accident or bad traffic, and without an indication of how long or severe a traffic jam will be, we cannot pinpoint our arrival time. Therefore, we consult our preferred navigation app or listen to the traffic updates on the radio to assess the situation. With this new information, we can better specify when we'll arrive.

Risk assessments do the same for the risk management process in organization. A risk assessment combines threat modeling with a surrounding process to gain specific and valuable information about the project's or organization's risks. With this process, the decision-makers can determine how they want to manage risks.

5.1 Methodologies

In a risk assessment, an analyst or a security professional is searching for risks in a defined environment. These risks are then classified and categorized to help the decision-makers with their risk management strategy. A risk assessment should follow a clearly defined methodology so all assessments and resulting risks can be compared. In the end, everything should be documented. Risks are not static, so the risk management process can't be static either. Therefore, the risk assessment needs to be done in an evolving way. A basic framework for a risk assessment process should contain the following steps: assessment preparation, threat modeling, risk assessment, and mitigation proposal.

Assessment Preparation

The preparation phase lays out the foundation of the risk assessment process. In this phase, the scope of the assessment is defined, assumptions are set, and the assets of the scope are listed.

Defining the scope

Defining the scope of the assessment is a crucial part in the analysis. The scope defines what should and should not be analyzed (White, 2014). A proper scope is narrow enough to not include systems that aren't under the control of the project team, but wide enough to include all connections to and from the systems in scope. All systems that are owned or operated by the project team and all connections to and from those systems should be in scope. In this case, a project team is the team that operates or develops the core of the scope systems. With this project definition, it is possible to divide the risk assessment into different projects. Most of the time, this scope definition only contains the technical part

of the project. It is also important to look at the surrounding processes of the system that needs to be analyzed. There are some cases where the technical system might seem properly implemented but the process around it might cause major threats to the organization (e.g., missing patch management or life cycle processes).

A normal three-tier web application running on a common public cloud provider is one example. The frontend and the application services are running on a virtual machine (VM), and a **Platform-as-a-Service (PaaS)** is used for the database, as in the figure below.

Platform-as-a-Service (PaaS)
A Platform-as-a-Service is a service model in cloud computing that provides a fully managed computing service as a platform for the user to build their application.

First, the out-of-scope elements are identified. In this case, the user is out-of-scope as the project team cannot control them. It might only be possible to manage the risk for the user, e.g., through awareness programs. Additionally, the cloud infrastructure is out of scope as the cloud provider manages this infrastructure. The virtual machine is in scope where the frontend and backend service. The database as a data store should also be in scope. The database server itself is out of scope as the cloud provider manages this one. Only the cloud customer facing configuration could be in scope. Normally, all connections to or from in-scope elements should also be in scope. In the example case, these are all drawn connections.

Figure 12: Dataflow Diagram of a Web Application

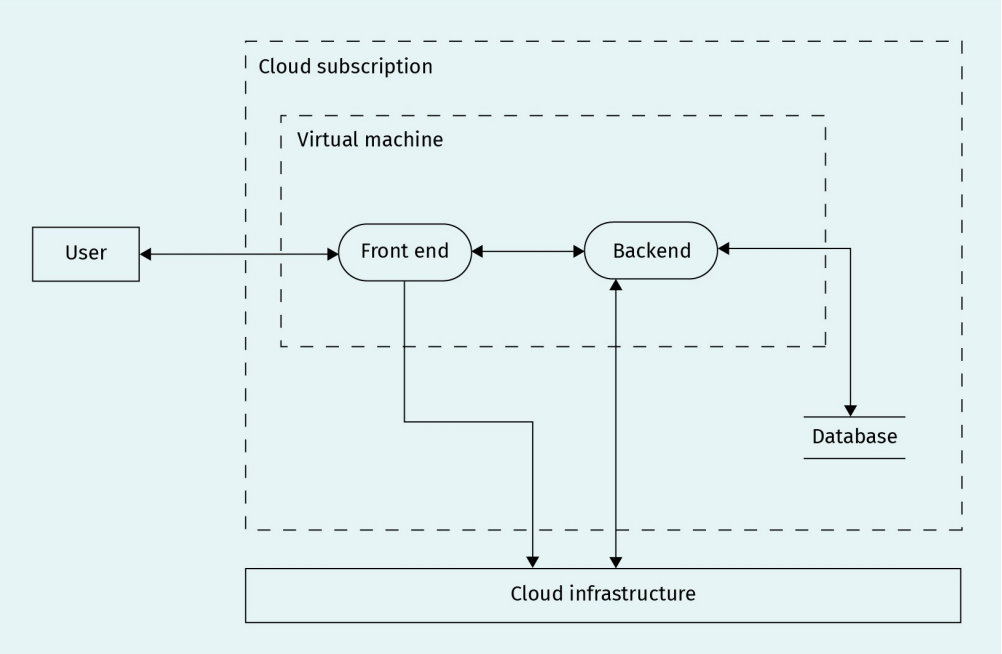

Source: Created on behalf of IU (2021).

Defining the assumptions

After the scope is defined, assumptions about the project are made. These assumptions are the baseline of the risk analysis. Assumptions tighten the scope. Without assumptions, the depth of the analysis can become uncontrollable. Common assumptions include the following:

- “The employees of the organization are trusted.”
- “The cloud provider does not act in a malicious way.”
- “The used hardware has no known vulnerabilities.”
- “The physical security is assured.”
- “The cloud infrastructure acts as described.”

In a nutshell, assumptions help the analyst to focus on the relevant part of the scoped project and not waste time on uninteresting parts. For example, with the assumption that the physical security is assured, the analyst can directly skip all the threats targeting the data center and start at the network threats of a server.

Listing the assets

Once the scope and the assumptions are set, the analyst can start searching for and listing the assets relevant for this analysis. Assets are items that are interesting to attackers and therefore need protection (e.g., intellectual property). These assets can be divided into two categories: primary and secondary or supporting assets (International Organization for Standardization [ISO]/International Electrotechnical Commision [IEC], 2018). Primary assets focus on the business and the attacker. These assets “make the money.” Business activities and information are primary assets. Examples include

- intellectual property,
- business plans,
- customer contacts,
- customer data, and
- personal identifiable information (PII).

The secondary assets support the primary assets and are only relevant when a primary asset is affected by them. A server can, for example, be classified as a supporting asset when customer data are stored on it. If the server hosts no primary asset, then it is not relevant in the scope since there is no risk. Supporting assets include

- servers,
- network,
- cryptographic material,
- secrets,
- processes, and
- employees.

The primary assets are the data stored on the database and the information from the user. Logging data could also be classified as a primary asset. The supporting assets are the virtual machine, the cloud processes, the keys and credentials used for operation, and the source code itself for the frontend and backend server.

In a risk assessment, the process of finding assets should be done together with the project team. With supporting questions from the analyst, they can identify the important assets and where these are stored or processed. With this information, the analyst can then describe the primary and supporting assets.

Threat Modeling

Following the preparation phase of the assessment process threat modeling can start. In the threat modeling phase, the analyst searches for threats in the defined scope (Shostack, 2014). The threats are all indirectly targeting one of the defined assets. Threat modeling should be done with a single methodology to be consistent in the project. Therefore, the analyst needs to decide which methodology should be used. Some methodologies as follows:

- STRIDE
- LINDDUN
- attack trees
- Process for Attack Simulation and Threat Analysis (PASTA)
- Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE)
- Visual, Agile, and Simple Threat modeling (VAST)

These threat model methodologies have different benefits and drawbacks. For example, STRIDE is very mature and easy to use, and attack trees are very expressive and easy repeatable.

Threat modeling should begin with drawing a data flow diagram (DFD) for the scope. This diagram shows the logical data flow and contains all information about entities and their relations of the system. A data flow diagram helps the analyst better understand the data flows and the processes of the scope. A DFD also shows if the scope is complete and nothing was forgotten. The DFD can then be used as a foundation for the threat modeling. Other supporting materials should be used for the threat modeling process. Talking to the architects and responsible project managers is always a good start. This can be done in workshops where everyone sits in the same room and the project can be discussed. Separate interviews are also a possibility. These interviews can then also be used to cross check if the different stakeholders have the same understanding of the scope and the project.

Asking the proper questions might sound tricky in the beginning. With more experience, it will be a lot easier to know what to ask to find threats. To get to this point, documents like requirement catalogs (e.g., **OWASPApplication Security Verification Standard**), best practices, or compliance standards are needed (e.g., Health Insurance Portability and Accountability Act [HIPAA]). Most of the time, it is not possible to find all potential risks since information might be missing or the threat landscape may have changed. Therefore, a completeness for the threat model needs to be defined.

Risk Assessment

Once the threat model is defined as complete, the assessment of these threats can start. The risk assessment takes the threats and calculates a risk for the organization. Threats without a risk to the organization can be ignored. The risk classification can be done with the formula $\mathrm{risk} = \mathrm{likelihood} \times \mathrm{impact}$. This formula describes the risk as a multiple of the likelihood of the threat and its potential. This helps the management to make an

informed decision. To have a repeatable output of the risk assessment, a proper framework should be chosen. Such frameworks (e.g., the OWASP risk rating methodology) describe the components of each factor and how to rate them (Williams, 2020).

Mitigation Proposal

Mitigation measures are directed to the threats and should reduce the risk (Shostack, 2014). Normally, these measures will reduce the likelihood since the impact can't usually be reduced. For example, when the impact of customer data theft is high and can't be changed, only mitigation measures can prevent the theft.

Different types of mitigation measures are possible. These can be related to processes (e.g., regarding to the software development process or change management) or technical measures (e.g., changes in the architecture or specific technical requirements, such as installing an anti-malware solution). Since all mitigation measures have a direct impact on the risk rating, a residual risk should be calculated to have an overall risk rating after the mitigation measures are in place. This helps to prioritize the implementation of the mitigation measures.

Documentation and Reporting

Documentation and reporting the risk assessment is essential to good risk assessment. An easy and formal way is to write a report. The report should include preparation steps, findings, and an explanation of what was done. A sample table of contents is shown below.

1. Management summary
2. Scope and prerequisites
    a) Assumptions
    b) Scope
    c) Assets
3. Threat modeling
    a) Diagrams
        i) Architecture diagram
        ii) Data flow diagram (DFD)
    b) Threats
        i) Threat-01
        ii) Threat-02
4. Risk assessment
5. Appendix

Another way to document the findings and the mitigation measures is directly in a ticket or tracking system of the project. In such a system, the project team can include the necessary work in their development cycle.

5.2 Factoring in Black Swan Events

Black swan events are events with a very high impact but a low likelihood. Thus, it might seem that these events are not that relevant in the risk assessment process. This might be true if only small parts of an organization are analyzed but, in reality, black swan events can put an organization out of business (Taleb, 2007).

From the analyst's point of view, the easiest way to factor in black swan events is to make assumptions that scope out these events. This method might be needed to leave the risk assessment in a tight scope and prevent going into a deep rabbit hole. On the other hand, when the risk for the organization needs to be assessed, black swan events need to be part of the assessment. There are two strategies for finding black swan events and assess the risk of these: 1) assessing events and finding risks to assets and 2) assessing threats to assets and finding events.

Assessing Events

The first assessment strategy for black swan events is to assess events and analyze what would happen to the scope. These events can be previous black swan and disaster events (e.g., the September 11, 2001 attack in the US). Wargame events, i.e., events that are imagined by the analysts, can also be played out and the resulting risks for the assets can be analyzed.

Assessing Threats

The second assessment strategy is basically the first one in reverse. In this strategy, threats for the assets are assessed and events that might lead to risks are found. These risks can be results from well-known threats from threat model strategies (e.g., the threats from STRIDE). For example, customer data might be assets to protect. A threat to that might be data loss (Denial of Service in STRIDE terms). With this threat in mind, we could find corresponding scenarios. Natural events might be a possibility. For example, a flood or a hurricane erases the data center where the live data are hosted and the backup data center is destroyed.

Mitigation Measures for Black Swan Events

As the likelihood is relatively low for black swan events, mitigation measures need to tackle the impact of such an event. Technical measures can be used to reduce the risk. In the example of the natural event, the data might be split in more data centers around the world. Another possibility is to avoid the risk. This is done with measures and controls that change the environment so that the threat is not possible. A third option is to transfer the risk with insurances. An organization can be insured for threats or disaster events. This does not prevent the threat from happening, but the impact is reduced as the insurance reimburses the monetary loss.

5.3 Continuous Re-Evaluation

Because the risk landscape for an organization can and will change over time, risk management and risk assessments are not one-time processes. A continuous re-evaluation is necessary to stay on track with the potential risks. The re-evaluation can be done yearly to re-assess the open risks or find new ones. When the in-scope target changes, a new risk assessment might be needed to assess the deviation between the original implementation and the change.

SEI Risk Management Paradigm

PDCA
The plan–do–check–act management method is used to continuously manage and improve a product or process.

The continuous re-evaluation can be part of the standard plan–do–check–act (**PDCA**) continuous management process (Tague, 2005). The Software Engineering Institute of the Carnegie Mellon University (CMU/SEI) went a step further and defined a risk management paradigm to continuously manage risk. This management process is divided into six steps: identify, analyze, plan, track, control, and communicate (Van Scoy, 1992).

Figure 13: SEI Risk Management Paradigm

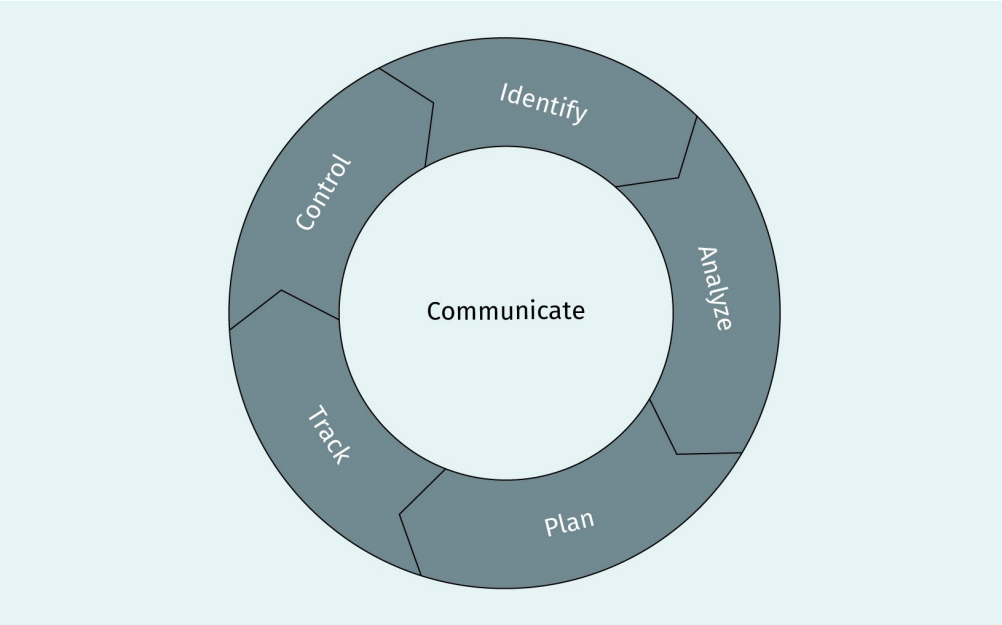

Source: Created on behalf of IU (2021), based on Van Scoy (1992).

Identify	Risks are discovered and identified. This is typically done in the risk assessment.
Analyze	Analyzing the risks includes evaluation, i.e., defining the likelihood and the impact of the risks to the organization. This is important to help the decision-makers focus on the proper risks

Plan	A mitigation plan needs to be established for each risk that contains the measure and timeframe for implementation. The possible measures are the normal risk management measures (e.g., as defined by the ISO/IEC 27005): risk reduction, risk retention, risk avoidance, and risk transfer.
Track	We track the implementation for the measure and determine whether it has had the planned effect. Tracking can be done with triggering events or milestones.
Control	The defined tracking needs to be controlled to be sure that the risk is properly managed. The risk and the measures should be controlled if they meet the desired target. Risk controlling should blend into the normal project management process and not be an unwanted side activity.
Communicate	Communication is key for the risk management process and in at the center of the paradigm. The results of all steps need to be communicated to the appropriate teams and decision-makers.

Tool Support for Continuous Re-Evaluation

Tools can help the re-evaluation process. The tools can be simple lists with all the risks and tracking triggers, plus the implementation plan. Previously used tracking tools (e.g., Atlassian Jira) can be used for the documentation of the evaluation. Some vendors have developed special risk management tools to help organizations carry out risk management (e.g., SimpleRisk). In the end, it is up to the requirements of the organization how they want to support the re-evaluation process with tools.

Figure 14: Risk Assessment and Evaluation List

Date	ID	Risk Name	Plan	Deadline	Control	Done
12.04.2021	R-01	Information disclosure due to missing encryption in transit	Implement encryption in transit	16.04.2021	Test if the connection is encrypted	✓
20.05.2021	R-02	Possible lateral movement in the corporate network	Implement network segregation	01.04.2022	Review the network architecture and check the possible conection between devices	✗
01.06.2021	R-03	SLA cannot be met in case of a cloud region failure	Build a backup environment in another cloud region	01.01.2022	Do failover tests and confirm the proper availability of the second region	✗

Source: Created on behalf of IU (2021).

SUMMARY

Risk assessments are needed in order to have a defined process to establish a specific scope and rate the found risks. The risk assessment process is comprised of assessment preparation, threat modeling, risk assessment, and mitigation proposal. Assessment preparation can be divided into defining the scope, defining the assumptions, and listing the assets. The scope definition describes a concrete area in which the risk assessment should be done. To tighten the scope, the assumptions are statements that the analyst defines as true.

The assets are the things which can be targets of threats. The assets can be split into primary and supporting assets. Primary assets are business relevant, such as intellectual properties or customer data. The supporting assets are systems, which are needed to process or store the primary assets (e.g., servers or software). Once all of these points are defined, the threat model can start with a threat model framework. After that, the threats from the threat model are assessed and risks are identified and rated. Based on these risks, the analyst can propose mitigation measures. Everything is documented for the decision-makers.

In a risk assessment, black swan events must also be considered. As the likelihood is near zero for those events, the impact and how it can be lowered needs to be analyzed. Assets can be analyzed and potential harmful events can be searched, or the other way around, and the impact to assets can be searched with a set of catastrophic events.

As the landscape of risks is changing, continuous re-evaluation is necessary to stay on track with these new risks. The Software Engineering Institute has developed a risk management paradigm to tackle this issue and carry out continuous risk management. The steps in this paradigm are identify, analyze, plan, track, control, and communicate.

UNIT 6

THE CYBER-RESILIENT ORGANIZATION

STUDY GOALS

On completion of this unit, you will be able to …

- describe how risk management needs to change for an Agile environment.
- understand why crisis management and incident response are part of risk management.
- explain why resilience engineering and security tools are needed for a cyber-resilient organization.

6. THE CYBER-RESILIENT ORGANIZATION

Introduction

The complexity of environments and infrastructure is rising, and with it, the threats to an organization. Therefore, organizations need to plan and implement measures to be more resilient against such threats. This preparedness can be exemplified with a home; if someone lives in an area with more crime, they might install better locks or a security camera. They may also lock the door to protect against intruders. The same goes for an organization working against cyber threats. If the threats or the environment changes, organizations adapt and change their risk management strategy, add measures and policies that might help the organization in an incident case, or re-think the way software or systems are built.

6.1 Changing Approaches to Risk Management

As companies change their development processes from a more static development process (e.g., the waterfall process) to an Agile process (e.g., Scrum), the risk management approach needs to change. In Agile projects, the risk management process can't be a static process. Furthermore, risk assessments need to be conducted more often as the scope will change with each software iteration or environment change (Brunton-Spall et al., 2017). For this cultural change, the whole approach to risk management need to change. Areas for these changes are governance, decision making, training, and tools (Deloitte, 2019).

Governance

In an Agile environment, the governance teams are not able to assess all risks and decide on mitigation measures for each of them. Therefore, the governance teams need to decide on clear guidelines for the delivery team about how risks should be managed. Thus, the decision for the risk strategy can be shifted from the governance teams to a lead member of the delivery team.

These policies or guidelines should define the following:

Risk appetite
The risk appetite defines the level of risk an organization is willing to accept without further consequences. Mitigating these risks is considered a waste of resources.

- What is the **risk appetite** of the organization?
- What is the **risk tolerance** of the organization?
- What is the **risk capacity** of the organization?
- At what level can the project decide on risks?
- How should risks be classified (i.e., how is the impact and likelihood rated)?
- How should risks be managed in an Agile team?

Decision-Making

Risk tolerance
The risk tolerance defines the level of risk an organization can accept. Mitigating the risk is balanced with the costs of the risk and the mitigation.

Risk capacity
The risk capacity defines the level of risk an organization can tolerate. Mitigating risks above this level is a must as these risks can harm the organization significantly.

The decision on how to handle risks needs to be in the Agile team (e.g., the project or product owner can be the risk owner). This is necessary in changing environments. Risks will change fast if the design changes in an iteration. In such a scenario, the risks need to be re-evaluated, and new decision need to be made. To stay on track with development, the decision needs to be made quickly and in cooperation with the developers. Mitigation can be planned easily into the Agile framework. For example, Scrum has many opportunities for risk management. In the sprint planning risks can be reviewed and weighted. During the sprint storywriting, the focus can be shifted to the implementation of the mitigation measure. In the design phase, threat modeling can identify risks. During the retrospectives, the implementation of the mitigation measures can be assessed and processes improved.

Training

Most developers are either not trained in risk management or have deeper knowledge in application security. Security experts are also rarer than developers. For this issue, the security community designed the security champion program (SAFECode, 2019). A security champion is part of a developer team, e.g., as a developer, a product manager or a tester. These champions then help implementing and improving security on their team. The duties of the security champion may vary from team to team. Some basic duties are

- risk management execution
  - threat modeling
  - risk assessment
  - risk tracking
- awareness raising
  - motivating the team to follow security controls
  - checking or writing security best practices
  - defining secure coding guidelines for the project team
- implementation check
  - security reviews
  - automated security scanners implementation

Since security champions are developers, and they get this extra task, they need to be volunteers. If the developers are forced to be security champions, then they would see security only as an extra unnecessary task (i.e., they won't be a big help). If they volunteered, they are more motivated doing the tasks. This approach will have a bigger impact and a better result.

Tools

Tools and automation are a big factor when developing in Agile. Tools and automation help the project teams to carry out their desired tasks quickly and efficiently so they can focus more on the development of their product. Security and risk management are no exception. Tools can help the security experts or champions conduct tests or threat models more easily. These can be tools that track and rate risks. Checklists are additional tools

that can be used to find standard issues in the design or the software. One such checklist is the Open Web Application Security Project Application Security Verfication Standard (OWASP ASVS), which has requirements for secure software development (OWASP, 2021).

Automation can be used to generate, for example, dataflow diagrams or other diagrams. This helps the threat modeler to threat model changes rapidly in the design or the system. Additionally, automation can be used to track and test that mitigation measures were correctly implemented.

6.2 Incident Response and Crisis Management

For a cyber-resilient organization, the risk management for development, pro-active measures, and planning for the worst case and incidents are key. Most of the time, it is not a question of "if" an organization will be hit by a cyberattack or a major incident, but "when." For an organization to be prepared, it must apply incident response and crisis management.

Crisis Management

Crisis management is needed to manage unexpected situations that endanger an organization. Crisis management can be divided into three steps (Deloitte, 2016): readiness, response, and recovery. All steps are accompanied by the crisis communication. The communication is important to keep all stakeholders informed about the situation.

Readiness

In this first step, an organization needs to prepare for possible scenarios and general crises. The preparation starts with a crisis strategy and choosing who should be involved in managing the crisis. The preparation can also contain playbooks for specific events or a checklist with all important points of a crisis. The tools for handling a crisis also need to be prepared. This is comprised of setting up communication channels and defining how specific things needs to be communicated and to whom. Training or simulating a crisis is also an important step in this phase. Wargames or simulations help the crisis team to find issues in the crisis management and become adept at handling such situations.

Response

In this phase, a real crisis is handled. The crisis team convenes and uses the policies from the first step to tackle the new crisis. The team also communicates the crisis to the needed parts to resolve this event as soon as possible.

Recovery

Once the crisis is handled, steps need to be taken to return to a normal status of operations. These include assessing the damage and setting up a strategy to fix the damage. The root cause of the event must also be determined and mitigated. As a last step, the crisis process needs to be assessed through a "lessons learned" meeting. In this meeting, the crisis team defines steps to improve the process to better tackle future crises.

Incident Response

The incident response process is similar to the crisis management process. A difference is that the incident response process also handles small events that disturb the normal operation of the organization and the crisis management only handles real crises for the organization (Cichonski et al., 2012). The IT incident response process also handles more technical problems than the crisis management process (e.g., system failures or cyberattacks). A cyber incident response process can be split into four main parts: 1) organizational preparation, 2) technical preparation, 3) handling, and 4) post-processing.

Organizational preparation

An organization needs to plan and prepare the incident response strategy and processes. The requirements for the full incident response process are defined in this step. The following areas should be covered in the organizational preparation (Cichonski et al., 2012):

- **asset management.** The asset management stores all information about all possible assets of an organization. These can be primary or supportive assets.
- **roles and responsibilities.** The roles and responsibilities define who does what in the incident response process. Roles include the on-call team, the incident investigator, the incident manager, and IT operations.
- **access concept.** The access concept describes how the incident response team can access systems or assets. This includes normal read-only access, for example, monitoring and logging or full administrative access in the incident event.
- **communication.** A definition of the communication is needed, as in an incident case, communication needs to be established quickly. This includes internal communication, communication to customers, and communication to law enforcement.
- **process documentation.** All incident response management tasks need to be defined in processes. This gives teams a guideline to prepare and handle an incident.
- **governance and policies.** The incident response management needs to be unified for the organization; therefore, governance and policies should be used to define this process. These policies can contain the definition of cyber incidents, the general reporting of such events, the escalation process, or how to rate and prioritize incidents.
- **training.** The response to incidents must be taught. This should be done to keep the incident response teams on track with new technologies, the environment, and processes of the organization.

Technical preparation

In the technical preparation, the incident response team plans and prepares their technical tools to handle incidents. This also contains the planning and the implementation of the following detection mechanisms:

- **Logging** is a key source of information. Logs can contain traces of the attacker and a list of events, which might lead to the incident.
- **Monitoring** can be used to track the health of the assets, and alarms can be sent to the response teams if anomalies are detected.
- **Access** needs to be implemented so that the incident response team can quickly access systems where the incident happens.
- **Response tool chains** of helpful incident response tools should be collected.

Handling

Handling an incident can be tricky. To properly handle an incident, follow these steps:

1. **Classification.** The input event that led to the incident is analyzed and categorized as real or false-positive.
2. **Triage.** The impact of the incident is analyzed. Here, the affected assets need to be found. With new information, the triage is repeated in the handling process.
3. **Containment.** After the affected assets are found, the impact of the incident needs to be contained. This can be done by, for example, shutting down servers, revoking accesss, or cutting the network connection to systems
4. **Eradication or recovery.** After the incident is contained, the environment needs to be cleaned up and all traces of the incident removed or fixed. Short-term mitigations can also be applied in this step.
5. **Reconstruction or root-cause analysis.** When the incident is contained and under control, the root-cause analysis can start. In this analysis step, evidence is analyzed to determine how the incident could happen.

Post-processing

Post processing of the incident finalizes the whole event. This is done in following steps:

1. **Lessons learned.** The lessons learned meeting is used for improving and learning. In the meeting, the team reviews the event and identifies improvements.
2. **Documentation.** Documentation is key for a proper incident response process. It is particularly helpful if similar events are happening.
3. **Enhancement of the organization.** With the identified root cause or issues, the organization can be improved. This can be done via advisories or new guidelines for the organization.

6.3 Resilience Engineering, Security Solutions, and Finances

Being a cyber-resilient organization implies that the deployed systems and services are also cyber-resilient. To achieve this goal, these systems need to be engineered with resiliency in mind. It is also possible to add security solutions to the environment to detect and prevent cyber incidents.

Resilience Engineering

Resilience engineering is designing and implementing systems with a security focus (Brunton-Spall et al., 2017). The level of security depends on the threat model of the organization and the risk appetite. Therefore, different systems might have different security controls in focus. Systems need to be designed to resist a compromise. This means the attacks that are more severe than the threshold from the risk appetite will fail, and the system is resilient against those. With this in mind, a common assumption is that security and usability have opposite goals. In reality, this should not be the case, because the organization needs to provide its users systems that are usable and secure. For example, an over-engineered system that heavily focuses on security will, in the end, reduce the overall security as users will find ways around the security controls to get their work done. Therefore, engineering secure and resilient systems needs a holistic approach to security. The system needs to be as secure as possible without breaking the usability. To this end, there are common controls to utilize. Such controls need to be defined by the organization through technical or process controls. An implementation should contain the following:

- **Resistive controls** are used to slow down and frustrate the attacker as a first hurdle. Rate-limiting or obfuscation are some of the measures in this control set.
- **Protective controls** are designed for the attack prevention, e.g., firewall access control lists (ACLs).
- **Detective controls** should contain measures to detect an attack or a malicious event. Such controls include logging and auditing controls.
- **Compensation controls** are second line defenses if a first line defense is not possible, e.g., if it is not possible to restrict a service to the internal network, other measures need to be implemented. These measures include enabling multifactor authentication or adding a web application firewall.

Security Solutions and Finances

To improve the cyber-resilience of the organization, it is also possible to buy security solutions. These solutions can be appliances or software to improve the overall security level of the organization. Such solutions need to be carefully planned and designed so they fit the organization's need. Some of these solutions include

- firewalls
- intrusion detection and prevention systems
- spam protection

- endpoint protection and response (EDR)
- security information and event management (SIEM) solutions

Firewalls

Firewalls are devices or software that inspect incoming and outgoing traffic according to a specific set of traffic rules. Modern firewalls (also called Next Generation Firewall [NGFW]) have more built-in capabilities. These are hardware devices or virtual machines that may offer application control, threat intelligence and prevention features, or deep-packet inspection (Brook, 2020).

Intrusion detection and prevention systems

As the name suggests, intrusion detection and prevention systems (IDS/IPS) detect intrusions in the network or a host. These systems can detect intrusions on a pattern basis or with machine-learning and anomaly detection (Barracuda, 2021).

Spam protection

Spam protection filters out unwanted spam or phishing emails. This drastically reduces the likelihood that employees will click on a malicious email (DuoCircle, 2021).

Endpoint protection and response (EDR)

An endpoint protection and response (EDR) solution is used to protect the endpoints (e.g., personal computer or server) of an organization from malware or other threats. EDRs collect data on the endpoint and prevent malware from executing. It also provides response capabilities as audit logging or execution traces from binaries (Malwarebytes, 2021).

Security information and event management (SIEM) solutions

A security information and event management (SIEM) solution is used to manage all security events in a central tool. This includes logging, correlation, and automatic notifications. Such solutions help to quickly react to security incidents and offer a broad overview of what happened (impreva, 2021).

6.4 Cyber Insurance

Cyber insurance is used to shift the identified risk to the insurance. This is used to complete the effort to be a cyber-resilient organization. Cyber insurance policies are broadly available and cover a variety of services and protections. Cyber insurance can cover the following events (AXA, 2021):

- malware attacks
- hacker attacks
- data misuse

- data manipulation
- data disclosure
- data loss
- system outage
- misappropriation of the system

Most cyber insurance policies have a clause explaining that the insured party needs to implement specific measures. If these measures are not implemented, the insurance company can refuse to pay or reduce the payout amount. AXA (2020) provide the following example: “If breach of duties of care, security regulations or other obligations are culpably breached, the indemnity can be reduced commensurate with the scope to which the breach has caused or influenced the damage” (p. 5). This means that the policyholder still needs to implement security controls and can’t rely purely on insurance. Therefore, cyber insurance should only be used to cover potential losses, even if all mitigation measures and security controls are in place.

SUMMARY

To be a cyber-resilient organization, changes to the approach need to be made. First of all, the approach to risk management will change as organizations increasingly prefer an Agile approach in developing software or systems. The classic risk management approach is not suitable to stay on track with the fast-changing Agile environment. Decisions need to be made in the development teams, and the governance department needs to provide guidelines and policies to steer these project teams. The next step in the changing risk management approach is to provide trainings and tools to development teams. This helps them to carry out risk assessments more independently from the central security teams.

The organization needs to engineer their software and systems to be more resilient. This can be done with different sets of controls that define how to securely develop the software. These controls are resistive, protective, detective and compensation. Security solutions can also be implemented to further strengthen the cyber-resiliency of the organization. Reactive measures need to be planned. Here, crisis management and incident response management are important. Cyber-resilience is further safeguarded through cyber insurance that can be used to transfer the leftover risk to an insurance company. These insurance policiess cover a variety of cyber events.