Chapter 2
Risk Management

The topic of this chapter is risk management in general. We begin by explaining what risk is and presenting the terminology we need in order to talk about risk. Thereafter we introduce risk management and explain what it involves for an organization to manage risk in a systematic and effective manner. Subsequently we look more into the details of the risk management process and its sub-processes.

2.1 What is Risk?

Basically, risk is the potential that something goes wrong and thereby causes harm or loss. The gravity of a risk depends on its likelihood to occur and its consequence. The consequence is the impact on an asset, and an asset is an object of value that we want to protect.

Definition 2.1 A *risk* is the likelihood of an incident and its consequence for an asset.

In order to convey more precisely what this definition means, we need to explain the concepts it refers to, namely incident, likelihood, consequence, and asset. We start with the notion of incident. When we discuss or assess risk we need to be careful to distinguish between its causes and the potential occurrence of the incident that constitutes the risk. Consider, for example, a burglar who enters a house by breaking in through a window. Understanding how such events unfold and what makes them possible is necessary for understanding how risk arises. But the actual incident itself is only the event that causes the harm or loss. In our example such an event could be the theft of jewelry. The definition of incident makes this precise.

Definition 2.2 An *incident* is an event that harms or reduces the value of an asset.

The definition of the term "incident" makes it clear that risk is about the occurrence of harmful events. But when can we say that an event is harmful and therefore an incident? After all, this depends on who we ask and what our focus is. For example,

A. Refsdal et al., *Cyber-Risk Management*, SpringerBriefs
in Computer Science, DOI 10.1007/978-3-319-23570-7_2

could we not say that a burglar breaking a window lock is a harmful event? And could we not say that a burglar entering a house is at the cost of privacy? The answer to these questions is that it depends on what our assets are. If our concern is the window lock, then the breaking of the lock is an incident. On the other hand, breaking a window lock alone does not harm privacy, and is therefore not an incident of a privacy risk. By including the notion of asset in the definition of incident we are forced to be specific about which events can be understood as incidents.

Definition 2.3 An *asset* is anything of value to a party.

The party is the entity or unit, such as a company or other organization, for which the assets in question have value. In the same way as there is no risk without an asset, there is no asset without a party. Because what is held as assets and how valuable they are depend on the party, we always need to be specific about who the party is when we manage or assess risk.

Definition 2.4 A *party* is an organization, company, person, group, or other body on whose behalf a risk assessment is conducted.

Hence, before we can discuss or assess risk, we must determine the party and the assets of concern to this party. Only then can we speak of incidents and risks in a precise and meaningful manner. In the burglar scenario, for example, the party could be the house owner. But it could also be someone renting the house. Whereas the assets for the renter may include both jewelry and privacy, the house owner would typically worry about damage to the property.

Notice that a party is not the same as a stakeholder. A party may be thought of as a stakeholder, but in a risk assessment situation there are normally many stakeholders that are not parties. A *stakeholder* in this context is basically any person or organization that may affect or be affected by the subject of the assessment. If we conduct a risk assessment on behalf of a company then the company is the party. Within and related to the company there may be many stakeholders (for instance, employees and suppliers) with all kinds of conflicting interests and they are not parties in this risk assessment. When we identify assets on behalf of a party we focus solely on the interests of the party in question. In most risk assessments there is only one party. If, however, there are several parties then the assets of the different parties must be kept apart. The same object, for example, a human life, may be an asset of different values for different parties. For you, the value of your life is perhaps infinite, while for a hospital it may be equal to the amount they have to pay the bereaved in the case of death due to maltreatment by the hospital.

The remaining concepts from the definition of risk are those of likelihood and consequence. Together, these notions characterize the gravity of a risk.

Definition 2.5 A *likelihood* is the chance of something to occur.

The notion of likelihood refers to the chance of something happening, no matter how we measure or represent it. Sometimes we describe it qualitatively, and other times quantitatively. We may describe it in general terms, and we may represent it mathematically as probability or frequency. *Probability* is a measure of the chance of

occurrence expressed as a number between 0 and 1, whereas *frequency* is a measure of the number of occurrences per unit of time.

Definition 2.6 A *consequence* is the impact of an incident on an asset in terms of harm or reduced asset value.

In this book the term "consequence" refers to negative impact only. Some approaches to risk management take a more general view of risk by considering any effect on assets, both positive and negative. This is useful when we conduct risk management with the aim of balancing risk and opportunity; negative outcomes could be accepted given the foreseen gain.

As mentioned above, the notions of likelihood and consequence characterize the gravity of a risk. We measure this gravity in terms of risk level.

Definition 2.7 *Risk level* is the magnitude of a risk as derived from its likelihood and consequence.

Our definitions of risk and risk level are well established and widely used. There are, however, alternative ways of expressing and measuring risk, some of which we discuss in Chap. 11.

The UML class diagram of Fig. 2.1 illustrates how the terms we have defined in this section relate to each other. Risk consists of three ingredients, namely consequence, incident, and likelihood. The relation represented by a line with a black diamond connecting risk and consequence captures that consequence is an ingredient that belongs to risk. The consequence represents the impact of an incident on an asset. Consequence is therefore also connected to the relation between incident and asset, since it is a measure of harm. The diagram also captures that for a given asset, there is a party that values it.

The same incident may give rise to several risks. Risk is therefore connected to incident with a white diamond to express that although incident is an ingredient of risk, it does not necessarily belong uniquely to one risk. Likelihood is a measure of how often the incident occurs. We may therefore see likelihood as an ingredient that belongs to incident. Since incident is an ingredient of risk we also have that likelihood is an ingredient of risk.

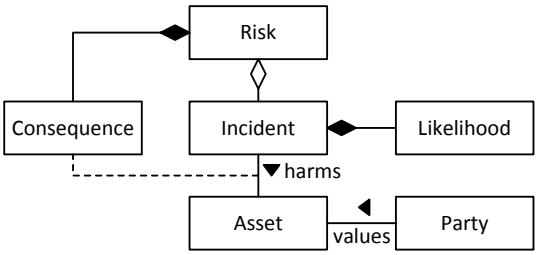

Fig. 2.1 Risk concepts

The terminology we just introduced allows us to explain what risk *is*, but not how risk *arises*. In the rest of this chapter we look more closely into how risks are managed, which includes understanding the sources and causes of risk and how to handle them.

2.2 What is Risk Management?

All organizations are exposed to risk, and most organizations do some kind of risk management. However, if we aim to precisely understand the kinds and nature of the risks, and to manage them in a systematic and effective manner, we need a well-defined process for risk management. We moreover need to understand the underlying principles and framework for the risk management process.

Definition 2.8 *Risk management* comprises coordinated activities to direct and control an organization with regard to risk.

For a risk management process to be adequate, efficient, and effective it should be based on a risk management framework. This framework should in turn comply with the basic principles for risk management. These relationships between the risk management principles, framework, and process are shown in Fig. 2.2. The framework should be subject to continual improvement, partly based on experience, findings, and results from the risk management process. This explains the arrow back from the process to the framework in Fig. 2.2. The purposes of the risk management

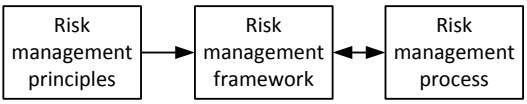

Fig. 2.2 Risk management elements

process must be decided as part of the overall management of the organization. This is why the implementation of the risk management process in the organization should be based on a risk management framework. The framework defines the mandate and commitment of the risk management, the risk management policy and responsibilities, the integration of the risk management into the organizational processes, and the mechanisms for internal and external communication and reporting. The risk management framework should be continuously monitored, reviewed, and improved.

The risk management framework, in turn, must comply with the basic principles for risk management. The principles apply to all kinds of risk management, but organizations need to understand what the principles mean for them and for their own framework for risk management. ISO 31000 lists eleven such principles. Among others, these include the principles that risk management shall create and protect value, that risk management shall be an integral part of all organizational processes,

that risk management shall be part of decision making, and that risk management shall be based on the best available information.

Figure 2.3 presents the risk management process in more detail. *Risk assessment* is a finite process that organizations conduct on a regular basis. The two others, namely *communication and consultation* and *monitoring and review*, are continuous activities. In the following we look more closely into the details of these three components of the risk management process.

Fig. 2.3 Risk management process

2.3 Communication and Consultation

By *communication and consultation* we mean activities aiming to provide, share, or obtain information and to interact with stakeholders regarding the management of risk. A *stakeholder* in a risk management context is a person or organization that may affect or be affected by the organization that is the subject of the risk management.

The interaction and information sharing serve as a basis for decision making. The information of relevance is anything that may determine how the organization should manage risk, including how risks should be communicated to internal and external stakeholders. This includes both external issues such as legislation, market situation, and external sources of risk, and internal issues such as reorganization, business strategies, and risk appetite. Such information can in general relate to the existence, nature, form, likelihood, significance, evaluation, acceptability, and treatment of risk [25].

For the communication and consultation to be efficient and effective it is advisable to establish a dedicated team and to define a plan for the process. This, in turn, helps to ensure endorsement of the risk management process and to communicate risk assessment results as explained in the following.

2.3.1 Establish a Consultative Team

The communication and consultation with internal and external stakeholders may concern any part or activity of the overall risk management process. Efficient and adequate communication and consultation ensures that those responsible for implementing the risk management process understand the basis for decisions and why particular actions are required. As part of the overall risk management it is useful to establish a consultative team with defined responsibilities for the communication and consultation. Such a team typically includes internal stakeholders such as decision makers and risk managers, as well as employees with insight into the organization. The team may also include external stakeholders such as board members, customers, and those with a vested interest. The roles and responsibilities of the team members must be clearly defined and specified. For smaller organizations it may not be an option to establish a team for the communication and consultation. In that case, the organization should still appoint a responsible point of contact and consultation.

2.3.2 Define a Plan for Communication and Consultation

The way risks are judged and perceived varies from person to person, even within the same organization. This may be due to differences in background, position, values, needs, concerns, and so forth. Decision makers need to take such varying perceptions into account when determining how to manage risks. A clear plan and good procedures for communication and consultation aid decision makers in this respect. The consultative team, or those responsible for the communication and consultation, should be involved in defining the plan and procedures. In addition to defining roles and responsibilities, organizations should establish procedures for how to support any of the processes of the overall risk management. This includes, for example, ensuring that different areas of expertise are brought together during risk assessments, that the interests of all relevant stakeholders are considered, that the risk evaluation criteria are appropriate, and that the decision making is informed.

2.3.3 Ensure Endorsement of the Risk Management Process

Communication and consultation support decision making, and aim to give decision makers and other stakeholders a sense of responsibility about the management of risks. Risk communication should furthermore help ensure mutual understanding among decision makers and stakeholders, thereby avoiding that bad decisions are made due to misunderstanding and lack of information. More fundamentally, good procedures for communication and consultation help to ensure endorsement of and support for the risk management process as such.

Effective and efficient management of risk requires decision makers, stakeholders, and any key personnel to pull in the same direction. For this purpose it is important to achieve a common agreement on and mutual understanding of how risk should be managed.

2.3.4 Communicate Risk Assessment Results

The results of the risk assessment are an important part of the information that must be communicated to all relevant stakeholders. This will support decision making, improve the understanding of the sources and nature of risk, strengthen risk awareness, and generally make the organization better positioned for managing risk. The communication of the risk assessment results will help both internal and external stakeholders to understand decisions and prioritizations regarding the management of risk. The risk assessment results may also be important for demonstrating, for example, policy adherence or compliance with directives and regulations. The risk assessment results are also important for justifying treatment plans, including the required resources for risk mitigation. Communicating the results helps those with a vested interest to understand the basis on which decisions are made and why particular actions are required. This, in turn, helps to ensure endorsement of risk treatment plans from key stakeholders.

2.4 Risk Assessment

By risk assessment we mean activities aiming to understand and document the risk picture for specific parts or aspects of a system or an organization. The assessment includes the estimation of the risk level, as well as the identification of options for risk treatment. The results serve as a decision basis for risk management, including the decision of which controls and measures to implement to mitigate risk. The *risk assessment process*[1] is divided into five steps, as illustrated by Fig. 2.4.

2.4.1 Context Establishment

The context establishment is the preparatory step for the subsequent activities and involves the documentation of both the external and the internal context of relevance for the assessment in question. This step defines the goals and objectives of the risk assessment, and therefore requires the participation of decision makers. The *external context* includes the relationships with external stakeholders, as well as the relevant

[1] This is a slight deviation from ISO 31000. See Sect. 2.6 for an explanation.

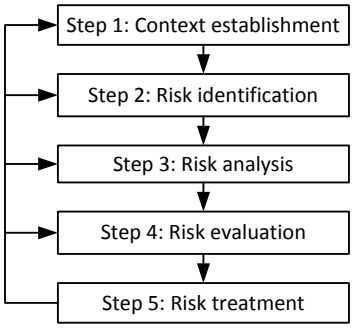

Fig. 2.4 Risk assessment process

societal, legal, regulatory, and financial environment. The *internal context* includes the relevant goals, objectives, policies, and capabilities that may determine how risk should be assessed.

In addition to establishing this general context for the risk assessment, the *context establishment* involves providing all the input that is needed for the following steps of risk assessment. We refer to this as the *context description*, the contents of which are discussed in the following. The *goals and objectives* are what we seek to achieve by the risk assessment. These can be of a high level, such as the achievement of business objectives or the provisioning of business services, but are important in order to understand the target, scope, and focus of the assessment.

Definition 2.9 The *target of assessment* is the parts and aspects of the system that are the subject of the risk assessment.

Definition 2.10 A *system* is a set of related entities that forms an integrated whole and has a boundary to its surroundings.

Notice that our definition of system is very broad. An organization, for instance, may be understood as a system according to this definition. The target of assessment (or target for short) includes the activities, processes, personnel, users, and all other relevant entities constituting the subject of the risk assessment. During the risk assessment we do the risk identification based on the description of the target. It is therefore important that the description is at a level of abstraction that matches the level of detail at which we aim to do the risk identification. It is useful also to decide and explicitly specify the desired scope and focus of the assessment. The *scope of the assessment* is the extent or range of a risk assessment; it defines what is held inside of and what is held outside of the assessment. The *focus of the assessment* is the main issue or central area of attention in a risk assessment; the focus is within the scope of the assessment.

Together with the description of the target of assessment we need to specify our *assumptions* about the target and its environment. An assumption is something we take for granted or accept as true about the system in question, and the risk assessment is valid only given these assumptions. Examples of assumptions could be that

risks are caused by internal personnel only because we have been asked to restrict our attention to company personnel, or that certain service level agreements (SLAs) will be fulfilled by external suppliers. Assumptions are made to focus the risk assessment and avoid duplicating work. The reason for the above SLA assumption could be that the party in question conducts (or has recently conducted) in parallel a separate risk assessment addressing the potential impacts of unfulfilled SLAs. The documentation of all such assumptions is essential because they are needed as input to the risk assessment, and because the results of the assessment are valid only under these assumptions.

A crucial step in the context establishment, and in defining the focus of the assessment, is the identification and documentation of the assets with respect to which the risk assessment is conducted. Before we can do the asset identification, we need to be specific about who the party of the risk assessment is. What is held as assets, how critical, important, or valuable the assets are, and the degree to which they require protection can be determined only by considering the party. A risk assessment is typically conducted with respect to one party, but it is possible to allow for two or more.

Having specified the target and assets we can define the risk scales and the risk evaluation criteria. For defining the risk scales we need scales for consequences and likelihoods. In principle we can use the same consequence values for all kinds of assets, for example in terms of monetary loss. However, this can be challenging in a practical setting where it may be very hard to know the economic implications of risks. Therefore it is often more useful to describe consequences that are specific to the asset in question. For example, for availability of a service the consequences could be given in terms of downtime. For each asset we therefore first consider the nature and kind of consequences that can occur and how they will be measured. Moreover, because the same risk assessment may involve assets of different kinds, we may need to define several consequence scales, one for each kind of asset.

For the documentation of likelihoods we need only one scale, but we need to decide how the likelihoods shall be measured. Sometimes it is suitable to use general terms such as “seldom” or “often,” and other times we use numeric, discrete scales. For some risk assessments the most suitable alternatives are frequencies or probabilities. The consequence and likelihood scales we define can be quantitative or qualitative, and they can be continuous, discrete, or given as intervals. In Chap. 12 we discuss the different alternatives more closely and give advice on which kind of scale to use for which purposes.

Risk levels are given by a function from likelihoods and consequences. This can be a mathematical function, for example by the multiplication of probability and monetary loss. In such a case of quantitative and continuous consequence and likelihood scales, the scale for risk levels is also quantitative and continuous. A more common way of specifying the risk function is by using a risk matrix with the likelihoods on one axis and the consequences on the other. Each cell then corresponds to a specified risk level. What is important when defining the risk function is that it serves as a basis for defining the risk evaluation criteria and that risk assessors and decision makers can distinguish between risk levels when the difference is signif-

icant for the risk evaluation. The risk matrix usually serves this purpose since we can always adjust the granularity by increasing or decreasing the number of cells. In Chap. 11 we present alternative ways of expressing risk and risk levels.

The *risk evaluation criteria* are the terms of reference by which the significance of risk is assessed. Because assets may be of different kinds and significance, we may need to define different evaluation criteria for different assets or different kinds of assets. The context description, which constitutes the collection of the above, serves as the input to and the basis for the risk assessment.

2.4.2 Risk Identification

By *risk identification* we mean activities aiming to identify, describe, and document risks and possible causes of risk. To this end we keep in mind two things. First, according to Def. 2.1, a risk is always associated with an incident. Second, there are three elements without which there can be no risk, namely asset, vulnerability, and threat. Without assets there is nothing to harm, without vulnerabilities there is no way to cause harm, and without threats there are no causes of harm. We therefore conduct the risk identification with respect to the identified assets by identifying threats and understanding how the threats may lead to incidents (and thereby risks) by exploiting vulnerabilities.

Definition 2.11 A *vulnerability* is a weakness, flaw, or deficiency that can be exploited by a threat to cause harm to an asset.

Examples of vulnerabilities are weak window lock and lack of intruder alarm, both of which a burglar can exploit in a break-in. Other examples are broken smoke detectors, insufficient staff training, and lack of back-up copies of critical operator manuals. The criticality of a vulnerability depends on the threats that may exploit them.

Definition 2.12 A *threat* is an action or event that is caused by a threat source and that may lead to an incident.

Threats may lead to incidents, but in order to identify threats and understand how they arise, we need to understand their initial causes, namely the threat sources.

Definition 2.13 A *threat source* is the potential cause of an incident.

A threat source can be human or non-human, and it can be tangible or intangible. Examples of human threat sources are burglars and negligent employees, while natural causes such as lightning or flood are non-human threat sources. Malware is an example of an intangible threat source.

Figure 2.5 illustrates how a threat source causes a threat that can lead to a risk by exploiting vulnerabilities. The arrow pointing backwards illustrates that threats can lead to other threats that eventually cause risks. During the risk identification we seek to understand and document how this can happen with respect to the identified

assets. In practice we often structure the risk identification by starting at one end

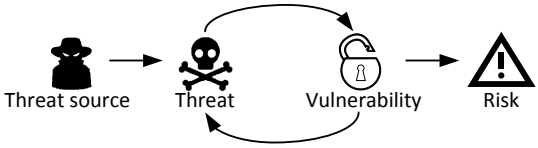

Fig. 2.5 Threat sources cause risks

and working our way to the other end, for example, by first identifying potential incidents and then trying to understand how and why they can arise. As illustrated by Fig. 2.6, we can go back and forth while gradually building the risk picture. For example, a threat that we identify for a given incident can trigger the identification of other incidents. As explained further in Sect. 5.3, where to start and in which order to address the questions depend on the kind of risk we are dealing with. When con-

Fig. 2.6 Risk identification

ducting risk identification in a systematical manner, we need techniques for doing the identification, and we need suitable formats for describing and documenting the results. Techniques for risk identification include brainstorming, interviews, checklists, statistics, and approaches for gathering historical data. Risk assessors may also use modeling techniques such as event trees [24], Bayesian networks [5], attack trees [70], CORAS diagrams [47], or threat modeling [75] to support the description of risks and how they are related to threats and threat sources.

Which technique to use depends on a variety of factors, such as the desired level of detail, the available resources and the expertise and experience of the risk assessors. The same is the case for the choice of documentation format. Apart from plain prose, there are basically two formats for the description and documentation of risks, namely tables and graphical models. Graphical models, such as those mentioned above, are often designed for specific purposes, for example to support brainstorming, to explore causes and/or consequences of incidents and scenarios, or to facilitate more rigorous assessment. Tables are suitable for structuring the information in a systematic way, but are typically used for more high-level risk assessments.

Whatever techniques and level of detail we choose for the risk identification and documentation, we always need to make sure that we describe all the elements of the risk picture that we need for the purpose and objectives of the risk assessment. At the very least the documentation should include threat sources, vulnerabilities, risks, and assets.

2.4.3 Risk Analysis

By *risk analysis* we mean activities aiming to estimate and determine the level of the identified risks. As defined in Sect. 2.1, the risk level is derived from the combination of the likelihood and consequence. The objective of this step, therefore, is to estimate likelihoods and consequences for the identified incidents using the scales defined during the context establishment. An incident represents one risk for each of the assets it harms, and we need to estimate the consequence for each of these assets.

The impact or severity of an incident can be determined only by considering the party in question. The severity of a wrongly addressed postal letter that leads to the exposure of confidential patient information, for example, is likely to be judged differently by the hospital and the patient in question. The consequence estimation should therefore be conducted by a walk-through of all identified incidents and assigning the estimates with the involvement of personnel representing the party or someone who can judge consequences on behalf of the party.

Likelihood estimation is to determine the frequency or probability of incidents to occur using the defined likelihood scale. This requires the use of techniques for gathering empirical data. Such techniques include interviews and brainstorming sessions to gather expert opinions, inspection of logs or other statistical and historical data, and the use of available repositories. Many of the risk-modeling techniques such as Bayesian networks, attack trees, and CORAS diagrams, also come with support for likelihood estimation and documentation. How we choose to model or document the risks during the risk identification may therefore have some implications on which techniques are available for the likelihood estimation.

The desired level of detail of the risk assessment and documentation is another factor. Sometimes we are only interested in the likelihoods of the incidents and make our best estimates directly for these events. Very often, however, we need to understand how risks are most likely to arise, and which threat sources are most important. In that case we should also try to estimate the likelihood that threat sources initiate threats, and the likelihood that such threats may lead to incidents. This information will not only help to understand the most important threat sources and vulnerabilities, but also to determine the likelihood of the resulting incidents.

Once we have estimated the likelihood and consequences for each incident, we calculate the risk level of all identified risks by using the risk function we defined during the context establishment.

2.4.4 Risk Evaluation

By *risk evaluation* we mean activities involving the comparison of the risk analysis results with the risk evaluation criteria to determine which risks should be considered for treatment.

In principle this step is quite straightforward given the risk estimates and evaluation criteria. For example, if we have specified the risk evaluation criteria using the risk matrix, we simply need to plot each risk into the matrix to determine the risk level. However, because the risk evaluation is a decision point in the overall risk assessment process, we take the time to confirm the risk evaluation criteria and consolidate the risk estimates. Decision makers and other personnel that are involved in the risk assessment often gain new insight and knowledge about the risks and their consequences, and we must therefore make sure that the initially defined criteria are still appropriate. For the consolidation of the risk assessment results we focus on the risk estimates that we are uncertain about, and where this uncertainty implies doubt about the actual risk level.

We moreover need to investigate the identified risks to see whether certain sets of risks should be aggregated and evaluated as a single risk. This is to avoid the pitfall of accepting a set of risks that individually are non-critical, yet unacceptable in combination. Even if the likelihood and consequence of the respective risks yield an acceptable risk level for both, it may be that the two are unacceptable taken together. Another situation is when we have a number of separate incidents that harm the same asset, and where the incidents can be understood as special cases of the same, more general, incident, or where the incidents may be caused by the same threat. How to aggregate likelihoods and consequences depends on the kinds of scales we use and any statistical dependencies between incidents.

A final recommendation for the risk evaluation is to group risks that have elements in common. Risks that share threat sources, threats, vulnerabilities, and/or assets may often be treated by the same means. Therefore, in preparation for the risk treatment and to facilitate cost-efficient treatment, we go through the identified risks and group them as we see appropriate.

2.4.5 Risk Treatment

By *risk treatment* we mean activities aiming to identify and select means for risk mitigation and reduction. Sometimes this step is referred to as risk modification to stress the fact that risks can both decrease and increase as a consequence of treatments. In particular, this is the case for approaches to risk management that may involve taking or increasing risk in order to pursue an opportunity. In this book we focus only on the identification of treatments for the purpose of reducing or removing risks. This is reflected by the following definition.

Definition 2.14 A *treatment* is an appropriate measure to reduce risk level.

In principle we should seek to treat all risks that are unacceptable, but in the end this is a question of cost and benefit, no matter the risk level. If a low risk is very cheap to eliminate, we might do so even if the risk in principle is acceptable. And, similarly, if the cost of treating a very high risk is unbearable there may be no other option than to accept it.

The risk treatment activity, therefore, should involve both the identification and the analysis of treatments. The treatment identification can be done similarly to the risk identification, for example via brainstorming or by the use of available lists and repositories. The selection of which treatments to implement should be the result of an analysis of the costs and benefits of the identified treatments. The analysis should take into account that some treatments can create new risks, and that some groups of treatments can reduce the isolated effect of each other. For example, unauthorized access may be mitigated by improved intrusion detection or by stronger access control, but we cannot expect the effect of the two in combination to be the sum of the effects of each of them alone.

There are four main options for risk treatment, namely risk reduction, risk retention, risk avoidance, and risk sharing [32]. We may reduce risk by reducing the likelihood and/or consequence of incidents. To do this we seek options to remove threat sources, remove or reduce the severity of vulnerabilities, or reduce the likelihood of threats by other means. Risk retention is to accept the risk by informed decision. This is typically an option for risks that are acceptable according to the risk criteria, or risks that are too costly to treat given the alternative options. Risk avoidance is simply to avoid the activity that gives rise to the risk in question, which sometimes is the only option for unacceptable risks. Risk sharing is to transfer the risk or parts of it to another party, for example, by insurance or sub-contracting.

2.5 Monitoring and Review

Monitoring and review apply both to the underlying risk management framework and to the risk management process, but specifically also to the identified risks and to the measures that the organization implements in order to treat risks. *Monitoring* is the continual checking, supervising, critically observing, or determining the current status in order to identify deviations from the expected or required status. The *review* activity is to determine the suitability, adequacy, and effectiveness of the risk management process and framework, as well as risks and treatments. The main purposes of the monitoring and review process are as follows [25]:

- Ensure that controls are effective and efficient
- Obtain further information to improve risk assessment
- Analyze and learn lessons from incidents, changes, trends, successes, and failures
- Detect changes
- Identify emerging risks

2.5.1 Monitoring and Review of Risks

Risks are not static and must therefore be monitored and reviewed. This includes all aspects of risks, including assets, threats, and vulnerabilities, as well as likelihoods and consequences. Constant monitoring is necessary for detecting and identifying changes to any of these aspects. Existing risk assessment results and other risk documentation must be reviewed to determine whether they are still valid. The monitoring and review of risk serves as a basis for taking actions, such as modifying the risk picture or conducting new risk assessments. Elements to monitor include the following:

- Assets: The set of assets that are of concern in the overall management of risk must be monitored in order to determine whether there are significant changes in asset value or priority over time. Changes in the internal or external context may moreover introduce new assets and make others obsolete.
- Threats: Internal or external changes could introduce new threats, including changes of assets or asset values. In some cases specific and known threats can be observed directly. In other cases it may be required to conduct risk assessments in order to thoroughly identify new threats.
- Vulnerabilities: Known vulnerabilities can be monitored in order to determine those that potentially could be exposed to new threats. They can also be monitored to detect changes, such as vulnerabilities that become more easy to exploit or more widespread.

Previous risk assessments are an important source of risk factors that should be monitored. In particular, this is the case for residual and acceptable risks that over time could evolve.

2.5.2 Monitoring and Review of Risk Management

Organizations also need to conduct continuous monitoring and review of the risk management framework and process. This is to ensure that the framework and process, as well as all related activities, procedures, roles, and responsibilities, remain relevant, appropriate, and adequate for the organization. Moreover, the review activity is conducted to verify that the risk evaluation criteria are valid over time, and that they are consistent with policies and business objectives.

Generally, any changes in the internal or external context that may affect the adequacy of the risk management need to be monitored and reviewed. This may include the following:

- Legal and environmental context
- Competition context
- Assets and asset values
- Risk evaluation criteria
- Resources required for adhering to the risk management framework

The risk management monitoring and review may result in changes in assets or evaluation criteria. But the required changes may also be more profound, for example, by changing the risk assessment techniques or tools, or by changing risk management procedures and responsibilities.

2.6 Further Reading

The terminology introduced in this chapter is largely based on the risk management vocabulary of ISO Guide 73 [26]. The presentation of the risk management process and how this process relates to the risk management framework and principles is based on the ISO 31000 risk management standard [25]. This standard also makes use of the vocabulary of ISO Guide 73.

Note that ISO 31000 refers to risk assessment as the three activities of risk identification, risk analysis, and risk evaluation. In this book we use the term "risk assessment" in a broader sense. It also include the activities of context establishment and risk treatment. There are two reasons for this: First, ISO 31000 offers no term denoting the process consisting of these five activities. Second, in our view, this better reflects how the term "risk assessment" is used in practice.

In addition to these ISO standards, we refer the interested reader to the ISO/IEC 27005 [32] standard on information security risk management. This standard is much more limited than ISO 31000 as it concerns information security risks, but because it builds closely on the latter, it gives some good insights into many principles of risk management in general.

For a useful and quite comprehensive overview and classification of risk assessment techniques, the reader is referred to IEC 31010 [30]. The overview includes techniques for risk identification, risk analysis, and risk evaluation.