

Chapter 9

Risk Evaluation

At this point we have identified the risks and analyzed their likelihood and consequence. From this we can establish the risk level and compare it to the risk evaluation criteria, as explained in Sect. 2.4.4 and Sect. 5.3.5. We also need to consider whether some risks that we have regarded as separate are actually instances of the same risk and therefore should be aggregated and evaluated as one risk. Furthermore, as preparation for the risk treatment, we group risks according to relationships such as shared vulnerabilities or threats. However, as analysis of likelihood and consequence is notoriously difficult, we start by reviewing the results from the previous step in order to check whether any adjustments need to be made.

9.1 Consolidation of Risk Analysis Results

The goal of the consolidation of risk analysis results is to make sure that the correct risk level is assigned to each risk. This is important because the risk levels direct the identification of treatments and provide essential decision support for the management. The central question is not whether each likelihood and consequence estimate is correct, but rather whether the resulting risk level is correct. For example, for risk no. 4 in Table 8.5, we assigned likelihood *Rare* and consequence *Moderate*, which according to the risk evaluation criteria defined by Fig. 6.2 gives risk level *Low*. Even if the likelihood is increased to *Unlikely*, the risk level will remain *Low*. Hence, for this risk, the distinction between these two likelihood levels is not essential for determining the risk level. On the other hand, if we are uncertain whether the consequence for risk no. 15 should remain at *Minor* or perhaps be increased to *Moderate*, then we need to investigate the issue, as this would bring the risk level from *Low* to *Medium*. When consolidating analysis results we direct our attention to the risks where 1) we are uncertain about the likelihood and/or consequence estimate and 2) this uncertainty may affect the risk level or the risk treatment.

We also make sure to check whether there are any risks that are both malicious and non-malicious. This is typically the case if malicious and non-malicious threats

can result in the same incident. In our case, this would mean that the same incident occurs in both Table 8.5 and Table 8.6. In such cases we need to check that the likelihood and consequence estimates are consistent, and that both the malicious and the non-malicious causes have been considered when estimating the likelihood. This can be easy to overlook since we are dealing with the malicious and non-malicious risks separately during much of the risk assessment.

As part of the consolidation we also revisit the risk evaluation criteria defined during the context establishment. Sometimes decision makers will want to adjust the criteria based on any new insights gained through the process so far, or on the results of the analysis.

The results of the consolidation are documented in the same place as the risk analysis results simply by making the necessary corrections and updates, and also adding references if new information sources have been used. For our analysis, this would mean updating the relevant entries in the tables presented in Chap. 8.

9.2 Evaluation of Risk Level

Having consolidated the risk analysis results, we are ready to evaluate the risks. The risk level of each risk is determined by its likelihood and consequence according to the risk matrix. In our case, risk evaluation is performed simply by plotting each risk in the risk matrix defined in Fig. 6.2. The result for malicious risks is shown in Fig. 9.1, where the numbers refer to the risk numbers in Table 8.5. Figure 9.2 shows the result for non-malicious risks from Table 8.6.

		Likelihood				
		Rare	Unlikely	Possible	Likely	Certain
Consequence	Critical		2			
	Major	6				
	Moderate	4,5	8		1	
	Minor				3	
	Insignificant	7				

Fig. 9.1 Risk matrix with malicious risks from Table 8.5

9.3 Risk Aggregation


During the evaluation we need to take into account that some risks may “pull in the same direction” to the degree that they should actually be evaluated as a single risk. There are basically two cases where this may hold.

		Likelihood				
		Rare	Unlikely	Possible	Likely	Certain
Consequence	Critical					
	Major	13				
	Moderate		11,12,14			
	Minor			15		9
	Insignificant				16,17	10

Fig. 9.2 Risk matrix with non-malicious risks from Table 8.6

The first case, which is illustrated by Fig. 9.3, concerns incidents that harm more than one asset of the same party, thereby giving rise to more than one risk for the party in question. Even if the risk of incident *X* harming asset *A* and the risk of incident *X* harming asset *B* are both low, it may be that the combined effect of harm to *A* and *B* warrants a higher risk level for the aggregation of these risks. In this case the likelihood of the aggregated risks remains the same, while the consequence is the joint consequence of the two risks.

Fig. 9.3 Aggregation of risks where one incident harms more than one asset of the same party



The second case is illustrated by Fig. 9.4 and concerns a single asset being harmed by more than one incident. Even if the risk of each individual incident harming the asset in question is low, it may be that the combined effect on the asset yields a higher risk. A typical situation in which we might aggregate is when the incidents are of the same nature, as is the case for Y_1 and Y_2 in Fig. 9.4 a), or when the occurrences of the incidents are triggered by the same threat, as is the case for U and V in Fig. 9.4 b). Notice that this also needs to be taken into account in cases where one of the incidents is malicious and the other is non-malicious.

Whatever the case and whatever the situation, we need not aggregate unless this can bring the aggregated risk to a new risk level. The risk level is, after all, what matters with respect to decision making. For a set of risks that are acceptable only if considered individually, deciding not to aggregate can give a false impression that no treatments are needed. Such decisions should therefore be taken with care.

We now return to our assessment. Going through Table 8.5 and Table 8.6 we find that there are no instances where a single incident harms more than one asset. Hence, the type of aggregation illustrated by Fig. 9.3 is not relevant for us.

However, risk no. 4, *Malware compromises meter data*, and risk no. 11, *Software bug on the metering terminal compromises meter data*, both concern software on the metering nodes and harm the integrity of meter data. They can therefore be viewed as special instances of a more generic incident, which we can call *Software on the*




Fig. 9.4 Aggregation of risk where a) two incidents are special instances of a common, more abstract instance, or b) two incidents are triggered by the same threat

metering node compromises meter data. Hence, they are candidates for aggregation as per Fig. 9.4 a). Looking at their risk levels in Figs. 9.1 and 9.2, we notice that their places in the risk matrix give reason to think that aggregation may yield a higher risk level than is given by either of the individual risks. We therefore decide to perform the aggregation. This is done by aggregating likelihood and consequence values separately, and then combining these to obtain the risk level in the usual way. As a starting point, we list the incidents, likelihoods, and consequences of the original risks, as shown in the upper rows of Table 9.1.

First up are the likelihoods. Here we notice that the incidents of risks nos. 4 and 11 may actually overlap to some degree. For example, malware may compromise meter data that are already compromised by a software bug. Moreover, the likelihoods are given as intervals rather than exact values, which means that adding up likelihoods may yield a new interval that spans more than one step of the likelihood scale defined in Table 6.3. This means that we cannot simply sum up the likelihoods of the contributing incidents, but need to use our judgment. After careful considerations about the nature of the incidents and the degree of overlap, we may for example arrive at likelihood *Possible* for the aggregated risk.

Next up are the consequences. Since the aggregated incident represents a generalization of each of the original incidents, rather than a combined occurrence, it clearly would not make sense to add up their consequences. Unless we are considering instances where simultaneous occurrences of several incidents cause additional harm, the consequence of the aggregated incident should not be greater than the highest of the original consequences. A good rule of thumb is that if all the original incidents have the same consequence, then we use the same value for the aggregated incident. If they do not, we can either use some kind of average value, possibly weighted according to likelihoods, or resolve the issue by consulting representatives of the party of the asset. In our case, we notice that risks nos. 4 and 11 both have consequence *Moderate*, hence this is also the value we use for the aggregated risk. The lowermost row of Table 9.1 shows the result. The plus sign denotes aggregation.

Similarly to the above case, it seems reasonable to aggregate risks nos. 5 and 12, and risks nos. 6 and 13. For the rest we decide to retain the original risks. Fig. 9.5

Table 9.1 Aggregation of risks nos. 4 and 11

No.	Incident	Likelihood	Consequence
4	Malware compromises meter data	Rare	Moderate
11	Software bug on the metering terminal compromises meter data	Unlikely	Moderate
4+11	Software on the metering node compromises meter data	Possible	Moderate

shows the results. All original malicious and non-malicious risks are included, as well as risks aggregated from both kinds.

		Likelihood				
		Rare	Unlikely	Possible	Likely	Certain
Consequence	Critical		2			
	Major	6,13	6+13			
	Moderate	4,5	8,11,12,14	4+11,5+12	1	
	Minor			15	3	9
	Insignificant	7			16,17	10

Fig. 9.5 Risk matrix after aggregation

9.4 Risk Grouping

Overviews like the one provided by Fig. 9.5 give an indication of which risks need treatment. However, as preparation for the risk treatment, we also want to take into consideration the fact that treatments may have an effect on several risks, thereby justifying higher cost than if we only consider individual risks. It can therefore be useful to group risks with this in mind.

The distinction between malicious and non-malicious risks earlier in the assessment has given us two groups. This is already useful, as some treatments will only have an effect on one of these groups. For example, data encryption, firewalls, and intrusion detection systems will usually reduce the likelihood or consequence of (some) malicious risks, without having any effect on non-malicious risks.

In addition to distinguishing between malicious and non-malicious risks, we may typically group risks according to shared vulnerabilities, threats, threat sources, or assets. The purpose of the grouping is to facilitate identification of the treatments that give the best effect for the least cost by placing together risks that may benefit from a common treatment.

In order to find out how to further group risks for our assessment, we systematically go through the results of the risk identification in Sect. 7.2 and Sect. 7.3. Do any of these risks have anything in common that indicates that they will benefit from the same treatment? Here we find, for example, that risk no. 14, *Mistakes during maintenance of the central system disrupt transmission of control data to the choke component*, and risk no. 15, *Mistakes during maintenance of the central system prevent reception of data from metering nodes*, are both related to the threat *Mistakes during update/maintenance of the central system* and to the vulnerability *Poor training and heavy workload*, as illustrated in Table 9.2. As shown in Fig. 9.2,

Table 9.2 Grouping of risks nos. 14 and 15

No.	Incident	Asset	Threat	Vulnerability
14	Mistakes during maintenance of the central system disrupt transmission of control data to the choke component	Provisioning of power to electricity customers	Mistakes during update/maintenance of the central system	Poor training and heavy workload
15	Mistakes during maintenance of the central system prevent reception of data from metering nodes	Availability of meter data	Same as the row above	Same as the row above

risks nos. 14 and 15 are both *Low*, but increasing the likelihood or consequence of either of them by a single step would bring its risk level to *Medium*. Treatments that address both these risks are therefore quite likely to be worth the cost. By grouping such risks we make it easier to take such considerations into account.

Similarly to the above case, we find that risks nos. 4-6 share a common threat and vulnerability, and that the same applies to risks nos. 11-13. Even if each of these risks is part of an aggregated risk with risk level *Medium*, thereby ensuring that they receive attention during the risk treatment, it is still useful to group them together for the purpose of cost-benefit analysis. We therefore create two new groups, one consisting of risks nos. 4-6 and one consisting of risks nos. 11-13.

9.5 Further Reading

For how to deal with uncertainty we refer to Chap. 13, which is dedicated to this particular problem. With respect to risk aggregation and grouping, we are not aware of any standards or similar sources that provide detailed guidelines, although the CORAS method [47] offers some support.



<http://www.springer.com/978-3-319-23569-1>

Cyber-Risk Management

Refsdal, A.; Solhaug, B.; Stolen, K.

2015, XI, 145 p. 32 illus., Softcover

ISBN: 978-3-319-23569-1