Multi-Stakeholder Cybersecurity Risk Assessment for Data Protection

Majid Mollaeefar[1,2], Alberto Siena[1] and Silvio Ranise[1]
[1]*Security and Trust, Fondazione Bruno Kessler, Trento, Italy*
[2]*DIBRIS, University of Genova, Genova, Italy*
{mmollaeefar, siena, ranise}@fbk.eu

Keywords: Cybersecurity, GDPR, Multi-Stakeholder Risk Assessment, Security and Privacy, Quantitative Risk Assessment.

Abstract: To ensure the effectiveness of the adopted security measures and minimize the impact of security issues on the rights and freedom of individuals, the General Data Protection Regulation (GDPR) requires to carry out a Data Processing Impact Assessment (DPIA). Such an assessment differs from traditional risk analyses in which the actor carrying out the evaluation is also the one interested in reducing its risk. Conflicts may thus arise between the need of protecting data subjects rights and organizations that shall provide adequate security measures while struggling with various types of constraints (e.g., budget). To alleviate this problem, we introduce the *Multi-Stakeholder Risk Trade-off Analysis Problem*, (MSRToAP) and propose an automated technique to solve their instances. We then show how this can help data controllers make more informed decisions about which security mechanisms allow for a better trade-off between their requirements and those of the data subjects. For concreteness, we illustrate the proposed on a simple yet realistic use case scenario.

1 INTRODUCTION

The General Data Protection Regulation (GDPR) has been introduced to guarantee basic rights of European citizens concerning personal data protection and privacy. One of the main goals of the GDPR is to give individuals control over their personal data. Controllers and processors of personal data must provide for appropriate technical and organizational measures to minimize the risks that personal data can be abused, for instance in the case of a data breach. To ensure the effectiveness of the adopted security measures and minimize the risk, the GDPR requires to carry out a Data Processing Impact Assessment (DPIA) in particular for certain data processing activities that are likely to result in a risk to data subjects rights and freedoms.

The focus of the DPIA is thus to consider risks that are likely to violate fundamental rights related to the privacy of individuals; thereby making a clear separation between the actor (the data controller) performing the risk assessment and the one (the data subject) whose risk should be reduced. This is dramatically different from traditional risk analyses in which the actor carrying out the evaluation is also the one interested in reducing its risk. As a consequence of the separation, conflicts may arise between the protection of data subjects rights and organizations that need to provide adequate security measures while struggling with budget constraints and an increasing skill gap in security and privacy.

To alleviate this problem and help data controllers make more informed decisions about which security mechanisms allow for a better trade-off between their requirements and those of the data subjects, we introduce the *Multi-Stakeholder Risk Trade-off Analysis Problem*, MSRToAP (Section 3) and propose an automated technique to solve instances of such a problem (Section 4). This allows designers to perform a what-if analysis and identify a (small) set of risk management policies that can simultaneously minimize the risks associated with the various actors (Section 5). The approach is illustrated by means of a simple but realistic running example (Section 2). We also discuss related work (Section 6) together with conclusions and future work (Section 7).

2 SCENARIO

ACME is the (anonymized) name of a real Italian startup operating in the healthcare domain. It develops a software application – we call it HCare – which exposes an API service to allow its clients to

Mollaeefar, M., Siena, A. and Ranise, S.
Multi-Stakeholder Cybersecurity Risk Assessment for Data Protection.
In *Proceedings of the 17th International Conference on Security and Cryptography (SECRYPT 2020)*, pages 341-348
ISBN: 978-989-758-446-6

work together, as illustrated in Figure 1. Through the API, HCare connects three actors: the API provider (ACME), the Health Service Provider (HSP), and patients. Notice that a HSP in our case can also be an independent developer, who provides IT-only services without offering actual health care support; for example, providing data visualization tools. Finally, the end-user is typically the patient using the app that can send biometric data or user-initiated requests and receive responses from the HSP; e.g., prescriptions from a doctor, medical alerts, etc. HSPs use the APIs to perform some operations such as create, read, update, and delete (CRUD operations) in a compliant way – i.e., by considering proper roles and permissions and storing and accessing the data accordingly.

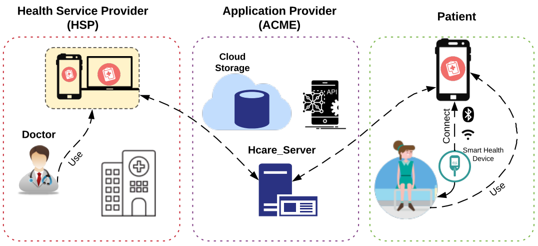

Figure 1: Health service scenario.

The health data is stored in a cloud environment, controlled, and monitored by ACME. Consequently, from a legal perspective, ACME acts as the data processor. However, due to the nature of its offered services, ACME has also to support data controllers to comply suitably. Therefore, it looks at the issue of GDPR compliance from both perspectives, of the data processor and data controllers. This is handled by a service level agreement between ACME and the HSP.

Table 1 shows an example of the impact assessment for the case of ACME. The table has been built through a process modeled on the GDPR data processing impact assessment procedure of ACME. It describes, with a certain level of abstraction, the identified risks with respect to the various principles. For example, the first row reports a risk for the *Confidentiality and integrity* of data (column 1). Data (such as the medical history of the patient) could be lost or corrupted because of a *hardware failure*. The consequences for the patient can be extremely high, because the healthcare data in this scenario is used for providing healthcare services such as medical prescriptions and missing or corrupted data may result in wrong diagnoses or in the impossibility to provide the service (column 2). For this reason, data storage must be reliable, by introducing more frequent backups or data replication (column 3). But these solutions change consequently the risk exposure for the company. In particular, data replication introduces the need for a complex network architecture, with all its associated risks. For example, business risks due to the rising costs, but also process risks (due to the difficulty of decision making and network configuration). Under the same principle, health data leakages could happen because of *unauthorized access* to the sensitive data, which may have bad consequences for the patient, such as *social stigma* and *discrimination*. Because of these consequences, *anonymization* or *pseudonymization* techniques may be required to be applied, but this in turn can introduce additional risks for the data controller, such as degradation of functionalities due to the fact that data can no longer be treated transparently (column 4). This example makes it evident the effects of the law: given that the data subject has some fundamental rights, it is a duty of the data controller to set up the proper technical measures for ensuring that the rights are respected. On the other hand, attempting to minimize the risks for the data subjects gives rise to the possibility that the risk exposure for the data controller increases– not only in terms of data-related risks.

3 PROBLEM DEFINITION

Different stakeholders have different criteria to evaluate the potential impact of threats. As a consequence, risk management policies – intended as the set of technical and organizational measures put in place to deal with risk – have different effects on the risk exposure of the stakeholders. A risk management policy (RMP) should therefore be selected with the purpose of minimizing the risks for all the considered stakeholders, while taking into account, at the same time, additional constraints, such as legal prescriptions or business requirements: failing to comply with applicable laws brings in additional risks and costs due to the expected fines; but even if the law is not formally violated, unbalanced risk levels among stakeholders may point out potential conflicts of interests, which can result in generating additional risks.

From these considerations, we derive the **Multi-Stakeholder Risk Trade-off Analysis Problem** (MSRToAP), which consists of the problem of jointly estimating the potentially conflicting risk levels for a set of actors, in relation to the technical settings and relating threats. We formalize MSRToAP as follows:

- A set of threats $T_1,\ldots,T_t$, each of which has a certain *degree of evidence* of being true or false.
- An RMP, defined as a set of technical and organizational security controls $C_1,\ldots,C_c$, each of them associated to a set of threats.

Table 1: Risk assessment under GDPR.

GDPR principles	Data subject risk	Potential solutions	Data controller risk
Confidentiality and integrity	Patient data losses or data corruption; wrong diagnoses by doctors;	Patient data backup; patient data replication;	Higher inside job possibility; time-consuming; rising costs; recovery procedure;
	Unauthorized access to patient health data; identity theft; loss of reputation;	Anonymization, pseudonymization and obfuscation; access control; encryption;	System slowdown; complexity; Possible implementation faults; functionality degradation;
Purpose limitation	Unintended permission; Unauthorized data disclosure;	Documenting the purposes in a transparent manner; restrict access to users' data;	Loss of public reputation;
Accuracy, Storage limitation, Data minimization	Disclosing undeleted inaccurate or medical history data; incorrect data may drive to discrimination or social pressure for patients;	Ensuring data accuracy; data cleaning algorithms; automated enforcement of deletion policies; regularly checking data collection;	Rising costs; possible implementation faults;

- A set of security goals $G_1,...,G_g$. E.g, confidentiality, integrity,and availability.
- A set of actors (e.g, the data controller (dc) and the data subject (ds) in the GDPR context) and, for each actor a, a set of criteria $P^a_1,...,P^a_p$ that actors use to evaluate the impact of threats. The method can support more actors, it results more complex (also from an algorithmic point of view), but the substance would not change.

We also assume that the risk R^a for actor a depends on its preferences $P^a_1,...,P^a_p$, the threats $T_1,...,T_t$, and a goal G_k for $k = 1,...,g$. An MSRToAP is the problem of finding a subset C of the finite set $\mathcal{C}$ of candidate RMPs that simultaneously minimise the risk R^a for all actors a. We need to consider a set C since, for any situation in which actors have conflicting objectives, there may be no single solution that is simultaneously optimal for each risk. Imagine a scenario in which, according to one RMP the risk of actor 1 is 1 and that of actor 2 is 2 whereas for another RMP the risks for the two actors are swapped; the former is better than the latter with respect to the risk of the first actor, but it is worse with respect to that of the second. An obvious question arises: which solution should be preferred? The answer is to use the notion of Pareto optimality; see, e.g., (Marler and Arora, 2004). In our context, an RMP is Pareto optimal if there does not exist another RMP that improves one risk without detriment to another. Identifying a set C of optimal RMPs is beneficial as it simplifies the work of designers that need to select the right RMP among the hopefully small set C rather than among all candidate RMPs.

To deal with this problem, in this paper we address the following objectives:

- develop a methodological framework, capable to drive the risk analyst in the quantitative estimation of the risks exposure levels for the actors referring to a same accountant; and
- evaluate the capability of the framework to explore alternative risk management policies, by performing a what-if analysis.

4 MULTI-STAKEHOLDER RISK ASSESSMENT

Risk is traditionally defined as R = L × I, where L is the likelihood of the risky events to happen, and I is their impact on the organizational goals, if they happen. We move from the idea that, within the same accountability context (i.e., the scope, for which the same actor is the accountant for the risk management), different stakeholders are exposed to different risks, but typically only one of them is *accountable* for performing the risk analysis and taking the proper measures. This is particularly sharpened in contexts in which laws, such as the recent GDPR, explicitely prescribe that a certain subject (the data controller) is accountable for the risks existing in its organization, but the risk assessment has to be carried out from the point of view and in the interest of the data subject. Given these hypotheses, we redefine risk as $R = \langle\ R_{ds}, R_{dc}\ \rangle = \langle\ (L_{ds} \times I_{ds}), (L_{dc} \times I_{dc})\ \rangle$. From the aforementioned distinction between *accountability* and *interest*, we derive one further property: since the interested parties are many, the perceived impact is different for each party; on the other hand, since the accountable actor is one by definition, the estimation of the likelihood is the same for all the actors. We assume therefore that $L_{ds} = L_{dc}$, and the risk becomes $R = \langle\ (L \times I_{ds}), (L \times I_{dc})\ \rangle$. The generalization to more than 2 actors is obvious. Given that there are plenty of approaches in the literature that cope with the likelihood estimation form risk assessment, and that its actual value is irrelevant for the purpose of comparing the risk exposures, in the present work we focus on the estimation of the relation between

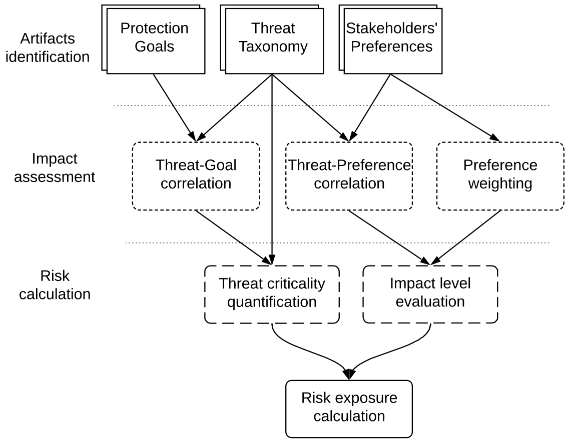

Figure 2: Risk assessment methodology.

the impact levels $I_1,...,I_i$ and the corresponding risk exposures. Given this assumption, the risk exposure depends on two things: (i) the method adopted to estimate the potential impact of threats on each actor a; and (ii) the threats actually associated to the system configuration. In the following we will introduce a methodology to perform a quantitative impact assessment for multiple actors, while in section 5 we will show how to use it to estimate the risk of a specific system configuration, as well as to explore potential risk mitigation solutions.

4.1 Impact Assessment Methodology

The impact assessment methodology consists of three main phases: Artifacts identification, Impact assessment, and Risk calculation, as illustrated in figure 2.

4.1.1 Artifacts Identification

In the initial phase, the following input artifact has to be identified:

Protection Goals. It presents the perspective of the data subject whose rights are at stake. For simplicity, in particular, we consider confidentiality, integrity, availability, unlinkability, intervenability, and transparency (Bieker et al., 2016) as protection goals which may be affected by the threats.

Threat Taxonomy. The potential threats that may exist in the scenario are identified by utilizing threat taxonomies such as ENISA, MITRE, etc. This threat taxonomy is considered as an input artifact. For the sake of simplicity, table 2 shows five possible threats relevant to the ACME scenario. For each threat, one (or more) possible malicious activity is reported, which represents the reason why the threat is considered to be relevant in the scenario.

Stakeholders' Preferences. Different stakeholders have different criteria that define what they consider risky. Data controllers (e.g., companies) typically adopt business impact criteria, such as financial impact or reputation, whereas data subjects (e.g., individuals) evaluate risk on the basis of their impact on their personal sphere. In our scenario, we have considered the *social situation, individual freedom, financial situation,* and *health condition* as the data subject impact criteria. *Reputational situation* and *financial situation* are the impact criteria for ACME, which are linked to indirect and direct pecuniary loss or damage.

4.1.2 Impact Assessment

The second phase is the core of the impact assessment. In risk management, impact assessment is a critical task, due to it is tough to estimate the potential consequences of threats and how they are evaluated by stakeholders. It is mostly a fully subjective evaluation, and this makes it particularly critical when attempting to estimate the impact from the point of view of third parties, such as the data owner. We accept this level of subjectivity as unavoidable, but nevertheless we mitigate it, by conceiving a cross-weighting system, which makes impact assessment more systematic. In a nutshell, in such weighting system the analyst is still required to provide his opinion about the potential impact of threats, but he is asked for an opinion multiple times and on different aspects of the problem, with the twofold objective of (i) inducing the analyst to consider multiple points of view, and (ii) merging these points of view into a more balanced value, which will be done in the third phase. In practice, the impact assessment is articulated into three finer-grained actions: threat-goal correlation, threat-preference correlation, and preference weighting.

Table 2: Possible threats in our scenario.

Threats	Possible malicious activities
T1-Unlimited data storage	Personal data is kept stored more than the time necessary for the purposes by ACME.
T2-Unauthorized access and disclosure	Due to over-privileged or inadequate controls, insiders (i.e., a medical practitioner or an ACME's staff) modify patients' data or disclose by mistake.
T3-Identity theft	Patients and their personal data can re-identify in de-identified data sets by outsiders malicious. (ACME data set)
T4-Denial of service	Attackers can disrupt the communication channel between patients and the healthcare service provider in order to prevent data from being upload to the server.
T5-Threat to intervenability	ACME does not implement a procedure (technical and/or processes) that allows the patients to rectify, erase or block individual data.

Table 3: The affected protection goals by each threat (right side of table), and the assigned impacts to stakeholders' preferences for each threat (left side of table), along with the assigned weights to each preference and the observed ratio for each protection goals in our scenario. The legends are represented as follows: **C**onfidentiality, **I**ntegrity, **A**vailability, **U**nlinkability, **In**tervenability, **T**ransparency, **F**inancial **S**itution, **R**eputational **S**ituation, **H**ealth **C**ondition, **I**ndividual **F**reedom, **S**ocial **S**ituation, and **O**bservation **W**eight. The "✓" symbol means the associated goal is affected by the threat.

Stakeholders' preferences						**Threats**	**Protection goals**						**OW**
Data subject				**Data controller**									
FS	**SS**	**IF**	**HC**	**RS**	**FS**		**C**	**I**	**A**	**U**	**In**	**T**	
0	1	0	0	1	2	**T1**	-	-	-	✓	-	-	1/13
3	2	2	4	2	2	**T2**	✓	✓	✓	✓	-	-	4/13
1	3	4	0	3	3	**T3**	✓	-	✓	✓	-	-	3/13
0	0	3	3	2	3	**T4**	-	✓	✓	-	✓	-	3/13
3	3	3	4	2	2	**T5**	-	-	-	-	✓	✓	2/13
0.1	**0.3**	**0.2**	**0.4**	**0.3**	**0.4**								
Preference weights													

Threat-goal Correlation. Each threat has potentially an impact on one or more security and privacy goals, which depends on the very nature of the threat. For example, every "Denial of service" threat will intuitively have more impact on data availability rather than on integrity, while any "Identity theft" will have more impact on data confidentiality. In table 3, we summarize the correlation between each threat and each goal. The more a threat has impact on multiple goals, the more it is considered pervasive; the more a goal is impacted by multiple threats, the more it is considered scattered. Threat-goal correlation does not directly inform on the threat impact, because a goal impacted severely by a single threat can be more threatened than a goal impacted slightly by many threats. However, it is an information that is used in combination with others to derive the impact.

Threat-preference Correlation. Each threat is perceived as more or less dangerous by each actor in relation of his/her own criteria. E.g., it is very unlikely that excessive storage of patients' health data would damage ACME's reputation. On the other hand, by increasing stored data on its data storage, there is financial damage on ACME cause of cost of storage and management of the IT infrastructure. The reputation of patients is not affected by excessive storage of personal data. However, the increasing amount of stored data increases the risk and impact of data breaches and leaks. To capture this correlation, we assign a value to the level of aversion that each stakeholder is considered to have against each threat, according to his/her own criteria. As shown in table 3, such value is expressed in a scale from 0 to 4 (no impact, low, moderate, critical, catastrophic), has the purpose of capturing the contribution of that specific preference to the overall evaluation of the risk impact.

Preference Weighting. Lastly, each stakeholder has different preferences, which result in different importance given to different criteria. For example, in our scenario for patients the health condition criterion is more momentous than the others. We capture these high level stakeholder preferences by assigning a weight to each stakeholder criterion, which we call it "assigned preference weight" or in short "APW". In table 3, the last row (on the left side) shows the assigned weights to each preference in our scenario.

4.1.3 Risk Caclulation

Finally, in this phase, the risk exposure calculates for each protection goal by following processes.

Threat Criticality Quantification. To quantify the threat criticality value, two values must be considered, (i) the number of goals is affected by the threat, and (ii) the threat evidence value. Basically, the threat criticality shows the level of danger of a threat among all threats. We introduce the "observation weight" for each threat according to the number of goals that are affected by the threat respect to all affected goals in total. For simplicity, we assume all the goals have equal importance. The observation weight for each threat is measured in table 3.

The threat criticality is computed by formula 1 to determine how much a threat has precedence over the other threats. In the formula, OW_i represents the value of the observation weight for threat i, and T_i represents the evidence value of the existence of the threat i, which is a value in [0...1]. In our scenario, in table 3 the observation weights are computed for each threat.

$$\forall_{i=1,..,t}\ TC_i = OW_i \times T_i \tag{1}$$

$$\forall_{i=1,..,t}\ NTC_i = \frac{TC_i}{\sum_{j=1}^{t} TC_j} \tag{2}$$

Indeed, the TC indicates how much a threat can be severe. To calculate risk exposure, we need to normalize the obtained values through formula 2.

Impact Level Evaluation. By having the values which have obtained from the threat-preference correlation process and the weight assigned to each preference, the impact level of each threat is evaluated. Formula 3 evaluates the impact level of each threat for each actor. In this formula, n is the number of preferences a stakeholder has, as well as, $I^a_{(i,j)}$ and $APW^a_{(i,j)}$ are the impact level and the assigned weight for the j-th preference of actor a for the threat i, respectively. For example, according to table 3, the impact level of the *health condition* for the second threat is 4. As can be seen in table 3, we have assigned a weight to each preference. E.g., in case of the data subject, the *health condition* has a higher weight than the *social situation*. As the same performed for the threat criticality, the obtained impact from formula 3 is normalized in order to evaluate the actual impact for each goal.

$$\forall_{i=1,\ldots,t}\, ImpactT^a_i = \sum_{j=1}^{n} I^a_{(i,j)} \times APW^a_{(i,j)} \tag{3}$$

Risk Exposure Calculation. Finally, the risk exposure calculates in terms of how much the protection goals are at risk for both perspectives. By formula 4, we can calculate the risk exposure from different perspectives for each protection goal, which helps us understand the risk exposure gap between stakeholders' interests. In this formula, X_j is the total number of affection of the j-th goal by threats, or other words, how many threats affect the j-th goal. For instance, according to table 3 the *transparency* goal has affected one time (out of five) which means X_5 is 1.

$$\forall_{j=1,\ldots,g}, GoalImpact^a_j = \frac{\sum_{i=1}^{t} NTC_i \times ImpactT^a_i}{X_j} \tag{4}$$

5 RISK POLICY SELECTION

The risk impact assessment methodology described above has the purpose of defining *how* the impact has to be evaluated for each stakeholder, when the existing threats are known. Knowing which threats exist, on the other hand, depends on the specific setting of the system. E.g., in the previous section, we have identified five threats in the scenario, and consequently, it needs to define a set of technical and organizational measures to see how much these threats harm the system. Each threat may have a truth value from 0 to 1, which is estimated by using controls. It is possible to obtain a high level characterization of the system threats by putting in place technical controls for each threat. E.g., a passed control provides evidence that the threat isn't effective, while a failed control provides evidence of the threat's existence.

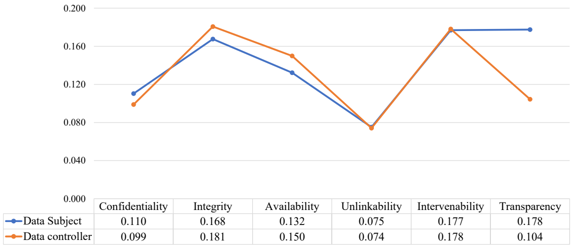

Figure 3: Risk impact exposure for data subject and data controller in the running example.

Out of all the controls for a given threat, the evidence that the threat exists can be simply evaluated as $T_k = 1 - (C_0 \times C_1 \times \ldots \times C_c)$, where $C_k = \{0 \ldots 1\}$. For simplicity, we currently use for C_k the truth values 0 (not implemented (No)), 0.5 (partially implemented) and 1 (fully implemented (Yes)). Table 4 shows the possible controls, along with their implementation status for the five identified threat. The RMP in table 4 is considered as the initial RMP for our scenario, which with that, we performed the risk impact assessment. In our scenario, table 5 shows the obtained result for the *TC* and its normalized value along with the normalized impact for both actors which are calculated for each threat by formula 1-3.

Table 4: Security and privacy controls in our scenario.

Threats	Controls	Implemented?
T1	1) Purpose specification	Yes
	2) Ensuring limited data processing	Yes
	3) Ensuring purpose related processing	Partially
	4) Ensuring data minimization	Partially
	5) Enabling data deletion	No
T2	1) Ensuring data subject authentication	Yes
	2) Ensuring staff authentication	Yes
	3) Ensuring device authentication	Partially
	4) Logging access to personal data	Partially
	5) Performing regular privacy audits	No
	6) Ensuring data anonymization	Partially
	7) Providing confidential communication	Yes
	8) Providing usable access control	Partially
	9) Ensuring secure storage	Yes
	10) Ensuring physical security	Partially
T3	1) Providing confidential communication	Yes
	2) Logging access to personal data	Partially
	3) Ensuring data subject authentication	Yes
	4) Ensuring data anonymization	Partially
T4	1) Enabling offline authentication	No
	2) Network monitoring	Partially
	3) Implementing the preventing denial-of-service attack mechanisms like firewalls, IDS, etc.	No
T5	1) Informing data subjects about data processing	Partially
	2) Handling data subject's change requests	Partially
	3) Providing data export functionality	No

To better understand the risk gap between stakeholders' interests, in figure 3, the risk exposure result has plotted for each protection goal. In this figure, the risk impact of each goal is expressed. By aggregating the impact of all goals, the overall risk impact from the point of the data subject is 0.84, which is the aggregation of goal exposures, and this value for the data controller is 0.78. Notice that these numerical values, per se, do not aim to have an absolute meaning; they are rather intended to serve for comparing the risk exposures of the stakeholders, against each other as well as against different risk management policies, as described in the next section.

5.1 What-if Analysis

The described approach consists of two main phases: a configuration stage, in which the characteristics of the specific organization are captured and transformed into parameters (basically, phases 1 and 2 of the methodology) of the risk evaluation method; and an execution stage, in which the existing threats are actually estimated on the basis of the controls, and their impact level evaluated, which in turn defines the risk exposure for the various stakeholders. What is interesting here is that, once the framework has been configured and the controls defined, it is possible to repeat the risk evaluation several times by changing the input values (i.e., the truth values associated with the controls). In other words, it is possible to perform a what-if analysis of the potential system configuration settings, in order to find the best one. To this purpose, moving from the risk policy currently in force in our scenario, and summarized in table 4, we define 5 additional risk management policies, which are potential alternative configurations of the system to be evaluated. The outcomes of this what-if evaluation are reported in table 6. As mentioned above, table 4 represents the adopted risk policy, while in table 6 policies 1-5 are the potential alternatives. In each policy, we have selected three controls; these controls and their implementation status are mentioned in table 6. For example, in policy 1, the three controls are as follows: (i) enabling data deletion for T1, (ii) ensuring device authentication for T2, and (iii) ensuring data anonymization for T3. Each policy relies on the same stakeholder preferences, protection goals and threat taxonomy, but implies different truth values for the described controls, corresponding to different technical solutions in the system. The truth values for a control can be 1, 0.5, or 0, which represents a control that has been fully or partially implemented, or not considered. We perform different combinations of controls to see how the risk exposure result. We conduct this analysis only for five different combinations under the mentioned condition. In table 6, RMPs have set, and their risk exposure result is calculated for two actors.

By testing all possible RMPs, we can observe how the risk exposure changes for the stakeholders. In a nutshell, we are looking for an optimal setting in which the overall risk exposure becomes minimum for the stakeholders. We see from the table for the DS, policy 1 and 2 are the worst. In contrast, policy 3 and 5 are the best for the DS. From the DC point of view, policies 1 and 2 have similar risk exposures, and policy 3 is the worst. In contrast, policies 4 and 5 are the best for the DC. However, out of this analysis, we can imply the best in both perspectives is policy 5.

6 RELATED WORK

In the scope of information security a wide range of risk assessment approaches have been proposed by standard institutes and organizations like NIST SP 800-30 (NIST, 2012), ISO/IEC 27005 (2011), etc. Regardless of the particular processes each of these security risk assessment approaches has, they all point out to the risk as an unexpected incident that would damage business assets, either tangible (e.g., organization's infrastructures) or intangible (e.g., organization's services). The ultimate goal of an information security program based on risk management is to aug-

Table 5: The obtained threat criticality and the normalized impact of each threat for stakeholders.

Threats	**T1**	**T2**	**T3**	**T4**	**T5**
Threat criticality (TC)	0.03	0.11	0.06	0.19	0.10
Normalized threat criticality (NTC)	0.06	0.22	0.12	0.39	0.21
Data subject nomalized impact	0.075	0.725	0.525	0.450	0.850
Data controller nomalized impact	0.393	0.500	0.750	0.643	0.500

Table 6: Stakeholders' risk exposure under different sets of implementing controls.

Policies	**Implemented Controls**	**DS risk**	**DC risk**
Set 1	T1-5) partially T2-3) fully T3-4) fully	0.85	0.79
Set 2	T1-5) partially T3-4) fully T5-1) fully	0.84	0.79
Set 3	T1-5) partially T2-5) fully T5-3) fully	**0.77**	**0.80**
Set 4	T2-5) partially T4-1) partially T5-3) partially	0.82	0.78
Set 5	T4-2) fully T5-2) fully T5-3) fully	**0.77**	**0.78**

ment the organization's output (product and service) while simultaneously limiting the unexpected adverse outcomes generated by potential risks. These methodologies have several limitations when intending to use them to analyze the risk in multi-stakeholder perspectives. Apart from that, for example, these frameworks are restricted in terms of what are risks related to data subjects and how to evaluate these risks, which is requested by the law. Numerous methodologies and frameworks in the context of privacy impact assessment (PIA) have proposed, such as legal frameworks for data protection authorities in several countries (Act, 2014; OAIC, 2014; CNiL, 2018), as well as academic researchers (Oetzel and Spiekermann, 2014; Clarke, 2009; Wright, 2012), and for specific purposes like PIA for RFID and Smart Grids (Commission, 2014; Oetzel et al., 2011). There a lot of risk assessment approaches which consider multi-criteria to calculate risk exposure, e,g., in (Zulueta et al., 2013) risk analysis is modeled as a Multi-Criteria Decision Making (MCDM) problem in which experts express their preferences for each risk. However, a few approaches that have defined risk impact criteria for different stakeholders. E.g., in the context of cloud computing, in (Albakri et al., 2014) a security risk assessment framework proposed that can enable cloud service providers to assess security risks in the cloud computing environment and allow cloud clients with different risk perspectives to contribute to risk assessment. In analyzing the conflict of interest between stakeholders in (Rajbhandari and Snekkenes, 2012) authors proposed the conflicting incentives risk analysis method in which risks are modelled in terms of conflicting incentives. The goal of it is to provide an approach in which the input parameters can be audited more easily. In (Wright, 2012), the authors have declared that privacy risk shall be assessed from both data subjects and system perspective. Similarly, in (Iwaya et al., 2019) a privacy risk assessment is proposed by considering both perspectives in the case of mobile health data collection system.

7 CONCLUSION

We have formalized the Multi-Stakeholder Risk Trade-off Analysis Problem together with an automated technique to identify a set of risk management policies that simultaneously minimize the risks associated with the data subjects and data controllers. This assists designers with conducting a DPIA, as mandated by the GDPR, by supporting a what-if analysis to explore various alternatives at design time or when there is a need to re-evaluate risks because of evolving requirements.

As future work, we plan to mechanize the proposed approach on top of an automated solver for multi-objective optimization problems. To simplify the practical application of the methodology, we will also identify indicators for threat detection. Finally, we are going to evaluate the integration of the methodology in existing risk assessment approaches.

REFERENCES

Act, D. P. (2014). Conducting privacy impact assessments code of practice. Technical report, Technical Report. Information Commissioners Office.

Albakri, S. H., Shanmugam, B., Samy, G. N., Idris, N. B., and Ahmed, A. (2014). Security risk assessment framework for cloud computing environments. *Security and Communication Networks*, 7(11):2114–2124.

Bieker, F., Friedewald, M., Hansen, M., Obersteller, H., and Rost, M. (2016). A process for data protection impact assessment under the european general data protection regulation. In *Annual Privacy Forum*, pages 21–37. Springer.

Clarke, R. (2009). Privacy impact assessment: Its origins and development. *Computer law & security review*, 25(2):123–135.

CNiL (2018). Privacy risk assessment (pia).

Commission, E. (2014). Data protection impact assessment template for smart grid and smart metering systems.

Iwaya, L. H., Fischer-Hübner, S., Åhlfeldt, R.-M., and Martucci, L. A. (2019). Mobile health systems for community-based primary care: Identifying controls and mitigating privacy threats. *JMIR mHealth and uHealth*, 7(3):e11642.

Marler, R. T. and Arora, J. S. (2004). Survey of multi-objective optimization methods for engineering. *Structural and multidisciplinary optimization*, 26(6):369–395.

OAIC (2014). Guide to undertaking a privacy impact assessment. https://www.oaic.gov.au/.

Oetzel, M. C. and Spiekermann, S. (2014). A systematic methodology for privacy impact assessments: a design science approach. *European Journal of Information Systems*, 23(2):126–150.

Oetzel, M. C., Spiekermann, S., Grüning, I., Kelter, H., and Mull, S. (2011). Privacy impact assessment guideline for rfid applications. *German Federal Office for Information Security (BSI).*

Rajbhandari, L. and Snekkenes, E. (2012). Intended actions: Risk is conflicting incentives. In *International Conference on Information Security*, pages 370–386. Springer.

Wright, D. (2012). The state of the art in privacy impact assessment. *Computer Law & Security Review*, 28(1):54–61.

Zulueta, Y., Martell, V., Martínez, J., and Martínez, L. (2013). A dynamic multi-expert multi-criteria decision making model for risk analysis. In *Mexican International Conference on Artificial Intelligence*, pages 132–143. Springer.