

Towards an Effective Privacy Impact and Risk Assessment Methodology: Risk Analysis

Majed Alshammari and Andrew Simpson

Department of Computer Science, University of Oxford
Wolfson Building, Parks Road, Oxford OX1 3QD, UK
`firstname.secondname@cs.ox.ac.uk`

Abstract. Privacy Impact Assessments (PIAs) play a crucial role in providing privacy protection for data subjects and supporting risk management. From an engineering perspective, the core of a PIA is a risk assessment, which typically follows a step-by-step process of risk identification and risk mitigation. In order for a PIA to be holistic and effective, it needs to be complemented by an appropriate privacy risk model that considers legal, organisational, societal and technical aspects. We propose a data-centric approach for identifying and analysing potential privacy risks in a comprehensive manner.

1 Introduction

It is widely recognised that the potential impacts of data-processing activities need to be proactively assessed in the early stages of the design process [12]. This has led to the emergence of the concept of a *Privacy Impact Assessment* (PIA) — a process that identifies and mitigates the impact of an initiative on privacy with stakeholders’ participation [19]. In order for a PIA to be holistic and effective, it is necessary for it to be complemented by an appropriate privacy risk model that considers legal, organisational, societal and technical aspects.

Privacy is a multifaceted concept that requires multidisciplinary considerations [8]. Privacy engineering, therefore, requires a sufficiently robust privacy risk model to identify potential privacy risks. The identified risks can then be addressed through risk management approaches, which include the selection and application of risk controls. We extend prior work by referring to fundamentals from the broader literature to underpin the main concepts of PIAs along with their meanings and properties. We present a data-centric approach that illustrates the main steps of identifying and analysing potential privacy risks in a meaningful manner. Through a realistic case study, we demonstrated the usefulness and applicability of this approach in a specific context. We argue that this contribution lays the foundation for systematic and rigorous PIA methodologies.

2 Background and Motivation

Ensuring that the processing of personal data is conducted fairly and lawfully is one of the main challenges in the context of data protection. This challenge has

raised concerns over data-processing activities that may lead to privacy violations or harms. *Privacy by Design (PbD)* [7] has been advocated as a response [8].

To realise the concept of PbD in the system development lifecycle (SDLC), potential privacy risks need to be proactively analysed and their potential harms need to be appropriately assessed [15]. In some jurisdictions, ‘legal compliance checks’ [15] or ‘prior checking’ [9] are the most commonly used privacy assessment procedures. These procedures are often not conducted by engineers; rather, auditors, lawyers or data protection authorities utilise a check-list to check compliance with legal frameworks [15]. With the advent of information and communication technologies, holistic and effective impact assessments are considered as complements to, or replacements for, these assessment procedures [15]. This has contributed to the emergence and wide use of the concept of PIAs.

A PIA is an ongoing process that begins at the earliest possible stages [20]. As such, PIAs are considered as a key means to address one of the main concerns of embedding privacy into the early stages of the design process, which is the manifestation of PbD [11]. Existing PIA processes strive to achieve the aim of PbD by applying its foundational principles [15].

The core of a PIA is a risk assessment, which typically follows a step-by-step process of risk identification and risk mitigation [15]. While PIAs are expected to follow the same philosophy, existing PIA processes largely fall short in this regard [15]. These limitations leave a number of open questions: *(1) How can we develop a privacy risk model that defines and/or refines key concepts and assessable risk factors, as well as the relationships among the factors?*; *(2) How can we identify potential privacy risks in a contextual and comprehensive manner to ensure the provision of end-to-end privacy protection?*; and *(3) What is the appropriate level of detail for such a model?*

3 An Analysis of PIA Processes

To identify data-processing activities that may lead to privacy violations or harms, it is essential to represent these activities in a way that is amenable to risk analysis and compliance checking. Rigorous data models need to be adopted to support the management and traceability of the processing and flow of personal data, as well as to help support identifying the planned, actual and potential data flows and processing. Such data models are expected to represent data-processing activities in a comprehensive manner and at an appropriate level of abstraction. This includes: personal data items, data-processing activities, involved actors, and their roles and responsibilities. Such information helps establish the context in which personal data is processed and identify system boundaries.

Some PIA processes, such as the BSI IT-Grundschutz [5], apply security risk analysis to privacy principles, which are typically given at a high level of abstraction, instead of relying upon a set of concrete protection goals. This, in turn, reduces privacy protection to the concepts of anonymity, pseudonymity, unobservability and unlinkability [4, 15]. Thus, targets of evaluation — i.e. personal data and data-processing activities — need to comply with legal frameworks and

standards, and ensure that they will not lead to potential privacy violations and harms. These targets define the scope of PIAs. As privacy principles are semantically different from concrete data-processing activities, it is difficult to use them for assessing these activities and describing design decisions at an architectural level. Accordingly, privacy principles need to be translated into concrete and auditable protection goals to aid engineers in specifying design strategies.

In order to conduct an appropriate privacy risk analysis that goes beyond a traditional security analysis, it is essential to develop a risk model that defines the key risk factors that have an impact on privacy risks, and to establish a conceptual relationship among these factors [12]. Existing PIA guidance documents, however, are not accompanied with proper guidelines or conceptual models that describe key risk factors to sufficiently support privacy risk assessment [15].

PIAs need to be complemented by an appropriate privacy risk model that goes beyond traditional security risk models. Such a model needs to consider not only legal and organisational aspects, but also societal and technical aspects. The model needs to refer to fundamentals from the legal privacy literature to underpin the main concepts, the key risk factors and the conceptual relationship between these factors. This addresses the first question of Section 2 (“How can we develop a privacy risk model that defines and/or refines key concepts and assessable risk factors, as well as the relationships among the factors?”).

Importantly, a privacy risk model needs to adopt a sufficiently robust model that facilitates end-to-end privacy protection and serves as the basis for the identification, analysis and assessment of potential privacy risks in a proactive, comprehensive and concrete manner. Such a robust model needs to sufficiently and contextually represent data-processing activities in a way that is amenable to risk analysis and compliance checking. This addresses the second question of Section 2 (“How can we identify potential privacy risks in a contextual and comprehensive manner to ensure the provision of end-to-end privacy protection?”).

In addition, an appropriate analysis approach needs to be adopted to systematically describe how combinations of risk factors are identified to be analysed. Such an approach needs to consider the appropriateness of the starting points of risk assessment and the level of abstraction in the context of privacy and data protection. This addresses the third question of Section 2 (“What is the appropriate level of detail for such a model?”).

4 A Privacy Risk Model

We review two privacy risk analysis methodologies [12, 10] upon which we build by refining the concepts, risk factors and relationships among these factors. We have chosen these models as they define and distinguish the key notions, risk factors and relationships among these factors in the context of privacy and data protection. To compare, we refer to fundamentals from the legal privacy literature to underpin the key concepts and risk factors along with their meanings, properties and relationships. In particular, we refer to the boundaries of privacy harm [6] to understand the specific characteristics and categories of privacy

harms. In addition, we refer to Solove’s taxonomy [16] to understand the specific characteristics of adverse privacy events and associated categories. Finally, we leverage the concept of contextual integrity [13] to understand the main characteristics of appropriate flow of personal data with reference to context-relative informational norms, from which vulnerabilities can be derived.

We define and/or refine the basic concepts used in conducting risk assessments to be appropriately applied in the context of privacy and data protection.

A *threat* is an event or action with the potential for privacy violation, or which might adversely impact the privacy of data subjects through the processing of personal data via inappropriate collection, retention, access, usage, disclosure or destruction. In our risk model, the threat concept can be decomposed into a threat source and a threat event.

A *threat source* is an entity with the capability to process (lawfully or unlawfully, fairly or unfairly) data belonging to a data subject and whose actions may instantly and/or eventually, accidentally or deliberately manifest threats, which may lead to privacy violations or harms. Each type of a threat source can be characterised by: type (insider or outsider; individual, institution or government; human or non-human), motives (stemming from the value of personal data), resources (including skills and background knowledge that helps re-identify data subjects), role (the way in which a concerned entity participates in processing operations), and responsibility. The specified attributes of a threat source are used to assess the capability of exploiting vulnerabilities. As such, a threat source is more relevant to vulnerability analysis than impact assessment. We use the concept of a threat source to ensure that it can be used appropriately for modelling actors with malicious and benign purposes. Joyee De and Le Métayer [12] use the concept of risk source to refer both to unauthorised entities processing personal data and to entities with legitimate processing capabilities. In [10], risk sources are those who act, accidentally or deliberately, on the supporting assets, on which the primary assets rely. Accordingly, threat sources who act, accidentally or deliberately, on the primary assets are not modelled. As such, we refine these concepts to be used appropriately at an appropriate level of abstraction. With regards to threat sources who act on the supporting assets, we refine the standard definition of threat action. A *threat action* is an intentional act (actively or passively) through which a threat source exploits the vulnerabilities of the supporting assets. It is important to separate the concept of the threat action to engage with the supporting asset and the threat event when a threat source acts against the primary asset.

A *threat event* is a technical event that may happen at specific points in time which has an effect, consequence or impact, especially a negative one, on the privacy of data subjects. Such events involve adverse actions justified by reference to personal data. A threat event is a possible source of privacy violations or harms: it occurs as a result of a successful exploitation of one or more vulnerabilities by one or more sources. Each type of threat event can be characterised by: nature (continuous or discrete; excessive or necessary; anticipated or unanticipated), scope (an individual, a specific group of individuals or whole

society), and category (according to the taxonomy of privacy). Joyee De and Le Métayer [12] and the CNIL methodology [10] use the concept of ‘feared events’. By referring to them as feared events, we may limit those to internal and unpleasant emotions and perceptions caused by the threat. As such, we use the notion of ‘threat events’ to describe harmful or unwanted events that may not be anticipated by data subjects. Since these events not only describe the data subject’s perceptions, we prefer to use threat events to describe unwanted, unwarranted or excessive processing activities that will lead to actual adverse consequences. They refer to a non-exhaustive list of common categories of feared events that an analyst should consider. However, we prefer to consider a well-known classification of such events. For the purpose of this paper, we consider only technical threats that are processing-related. In particular, we focus on data-processing activities, which are composed of adverse actions that are justified by reference to personal data, and events that cause the performance of these actions, which can and do constitute privacy violations or create privacy harms.

A *privacy vulnerability* is a weakness or deficiency in personal data modelling, the specification or implementation of processing operations, or privacy controls, which makes an exploitation of an asset more likely to succeed by one or more threat sources. Successful exploitations lead to threat events that can result in privacy violations or harms. In our context, assets can be classified into *primary assets* and *supporting assets* [10]. The former refers to personal data that is directly concerned with processing operations, as well as processes required by legal frameworks and standards. The latter refers to system components on which the primary assets rely. For the purpose of this paper, we focus on the primary assets and associated vulnerabilities. Each type of vulnerability can be characterised by exploitability and severity. These are used to estimate the seriousness of a vulnerability.

The CNIL methodology [10] uses the concept of vulnerability, which refers to a characteristic of a supporting asset that can be used by risk sources and allowing threats to occur. In contrast, Joyee De and Le Métayer [12] use the concept of ‘privacy weakness’ to refer to a weakness in the data protection mechanisms. By using this concept, they aim to include weaknesses that may not be considered by using the concept of vulnerability, such as inappropriate functionality from which privacy harms may stem. As such, we use the concept of vulnerability with a broader view to not identify them only within data protection mechanisms. Privacy vulnerabilities can be found in the implemented privacy controls and the specified processing operations along with required personal. In addition, we use the classification of assets of [10].

A *privacy violation* is an unfair and/or unlawful action that accidentally or deliberately breaches privacy-related laws, regulations, unilateral policies, contracts, cultural norms or principles. Such actions are triggered by occurrences of threat events that result from the successful exploitation of one or more vulnerabilities. In reality, inappropriate processing of personal data may lead to privacy violations, which may involve a variety of types of activities that may lead to privacy harms [16]. Importantly, the presence of a privacy violation does

not mean that it will necessarily create actual privacy harm. Further, privacy harms can occur without privacy violations [6]. Each type of privacy violation can be characterised by: type (unlawful or unfair), degree (excessive or limited), and scope (an individual, a group of individuals or whole society).

Joyee De and Le Métayer [12] and the CNIL methodology [10] do not distinguish between privacy violations and harms.

A *privacy harm* is the adverse impact (incorporeal, financial or physical) of the processing of personal data on the privacy of a data subject, a specific group of data subjects or the society as a whole, resulting from one or more threat events. A widely held view conceptualises a privacy harm as the negative consequence of a privacy violation [6]. However, privacy harms are related to, but distinct from, privacy violations. This implies that it is not necessary for an actor to commit a privacy violation for a privacy harm to occur and vice versa. Each privacy harm can be characterised by: type (subjective or objective), category (incorporeal, financial or physical), adverse consequences (last for a short time, last for a certain length of time or last for a long time), and affected data subjects (a data subject, a specific group of data subjects, or whole society). Subjective privacy harm represents the perception of inappropriate processing of personal data that results in unwelcome mental states, such as anxiety, embarrassment or fear, whereas objective privacy harm represents the actual adverse consequence, such as identity theft that stems from the potential or actual inappropriate processing of personal data [6].

The CNIL methodology [10] uses the concept of prejudicial effect to assess how much damage would be caused by all the potential impacts. As such, feared events are ranked by estimating their severity based on the level of identification of personal data and the prejudicial effect of these potential impacts. To identify potential impacts of feared events, consequences on the identity and privacy of data subjects and human rights or civil liberties need to be identified. This means that it does not characterise privacy harms to facilitate their identification and analysis. In contrast, Joyee De and Le Métayer [12] use the concept of privacy harms with specific attributes and categories. In our approach, we use the same concept with more details to identify privacy harms at a detailed level of abstraction according to the properties and boundaries identified in [6].

5 An Analysis Approach

Risk analysis approaches differ with respect to the starting points of risk assessments and levels of abstraction. In order for risk assessments to be effective, they need to synthesise multiple analysis approaches to identify the key factors of risk. Potential privacy risks need to be identified, analysed and assessed in a systematic manner. As such, our approach consists of four steps.

Step 1: Context Establishment. Establishing the context in which personal data is processed plays a crucial role in understanding the scope under consideration by identifying all the useful information for privacy risk analysis. This includes the types of personal data to be processed (primary assets that

need to be protected), along with its sources; the purposes for, and the manner in which, this data is processed; involved actors and their assigned roles and responsibilities; relevant legal frameworks and standards; and domain-specific constraints.

As discussed in Section 4, primary assets are classified into *personal data* and *processes*. As such, personal data, associated processes and involved actors need to be represented in a way that is amenable to analysis. While describing systems in multiple views is important [15], we emphasise the importance of data-management models that represent data and associated processing activities at a detailed level of abstraction. We believe that data lifecycles are better at describing processing activities in a detailed level of abstraction.

The Abstract Personal Data Lifecycle (APDL) Model [2] was developed to represent data-processing activities in a way that is amenable to analysis and compliance checking. It represents the personal data lifecycle in terms of lifecycle stages, along with associated activities and involved actors. It can be used to complement a PIA for describing the planned, actual and potential processing of personal data, which, in turn, helps facilitate the management and traceability of the flow of personal data from collection to destruction [2].

Accordingly, we adopt the APDL model to represent the primary assets, along with involved actors. Personal data is represented in the *DataModelling* stage. This stage represents the relevant objects, associated properties, relationships and constraints for the purpose of specifying the minimum amount of required personal data. Processes are abstractly represented in eight stages: *Initiation*, *Collection*, *Retention*, *Access*, *Review*, *Usage*, *Disclosure* and *Destruction*. In each stage, data-processing activities and those required by legal frameworks and standards are concretely represented in *StageActivity*, *StageEvent* and *StageAction*. In addition, involved actors and the way in which they participate in processing activities are represented in *LifecycleRole* and *LifecycleActor*. We use the UML [14] profile for the APDL model proposed in [1] to represent personal data, associated processes and involved actors.

Step 2: Vulnerability Analysis. We assume that identifying and analysing vulnerabilities of the *supporting assets* is part of security risk analysis to ensure availability, integrity or confidentiality of the primary assets. We focus only on vulnerabilities of the *primary assets* to protect the privacy of data subjects and ensure the contextual integrity.

The first step is to define a baseline model of processing that describes the targets of evaluation (primary assets) at an appropriate level of abstraction. To this end, we adopt the concept of contextual integrity [13], which was developed to bring the social layer into view by identifying four main elements: contexts, attributes, actors and transmission principles. These elements constitute *context-relative informational norms*, which govern the flow of information in a particular context to ensure its appropriateness. From a technical perspective, these norms can be adapted by including processing activities as an element to consider both the flow of personal data and the processing of this data (we refer to the adapted

norms as context-relative processing norms). In so doing, contextual integrity is about the appropriate flow and processing of personal data.

In order to comprehensively identify and analyse all possible vulnerabilities of the primary assets, a baseline model, which describes personal data, associated processes and involved actors, needs to be represented in a way that is amenable to analysis. As such, the baseline model of processing can be described in terms of context-relative processing norms. We adopt the APDL model as a source to capture and represent personal data, associated processing activities, involved actors and their assigned roles in each stage of the lifecycle. In addition, processing principles — which can be derived from legal frameworks, standards or domain-specific constraints — are represented as constraints for each data-processing activity in each stage of the data lifecycle. We use the UML profile for the APDL model to describe the context-relative processing norms in a widely-used modelling notation.

Once the context-relative processing norms, which constitute a complete baseline model, are established, vulnerabilities can be derived from how these norms would be breached or disrupted to violate contextual integrity. Crucially, each element of each processing norm (data attributes, data-processing activities, actors and processing principles) need to be considered separately to ensure that: the data attributes are sufficient to fulfil the data-processing activity; the data-processing activity is assigned to authorised actors according to their roles and responsibilities; and the constraints (pre and post-conditions) are modelled in a way that ensures the data-processing activity is specified in conformity with the processing principles. Improper data model and a lack of data minimisation are examples of weaknesses for the elements of *attributes* and *processing principles* respectively. These vulnerabilities may be exploited by a threat source and lead to the identification of a data subject as a threat event. For each vulnerability, its exploitability and severity need to be identified and estimated in relation to its attributes in Section 4.

Step 3: Threat Analysis. *Threat sources.* In order to identify all possible threat sources, it is necessary to establish the context in which personal data is collected and processed (as per Section 5). The context helps support engineers in understanding the scope of analysis, multiple stakeholders, the nature and sensitivity of the processed data. Once the context is established, a list of actors involved in the processing of personal data can be identified, along with assigned roles and responsibilities. In particular, the Initiation stage can be used to concretely identify the types of personal data to be collected and processed, and to abstractly identify involved actors and their roles and responsibilities. In order to identify involved actors at a detailed level of abstraction, we use the basic types of lifecycle roles (data modeller, data subjects, data controllers, data processors and third parties) in each stage of the lifecycle as a source of such details. A lifecycle role is a set of logically related activities that are expected to be conducted together and assigned to different actors as responsibilities according to their capabilities. In addition, a list of entities with interests or concerns in the value of these types of personal data can be identified. All such entities

are potential threat sources. For each threat source, its type, motives, resources, role and responsibilities need to be identified in relation to its attributes.

Threat events. Once the context is established, and vulnerabilities and threat sources are identified, a list of events with the potential to adversely impact the privacy of data subjects can be identified. We adopt the taxonomy of privacy [16] as a means for characterising adverse privacy events. The taxonomy helps facilitate the identification of these events in a comprehensive and concrete manner. It classifies the most common adverse events into four basic groups: information collection, information processing, information dissemination and invasions. Adverse events are arranged around a model that begins with the data subject, from which various entities collect personal data. Data holders process the collected data. They may also disseminate or release the processed data to other entities. The progression from collection through processing to dissemination is indication of the personal data moving further away from the control of the data subject. In the last group of adverse events (invasions) the progression is toward the data subject and does not necessarily involve personal data [16].

The taxonomy was developed to serve as a framework for the future development of the field of privacy law. In our approach, however, we focus only on data-driven adverse events that are more related to primary assets than supporting assets. From a technical perspective, these adverse events need to be arranged around a widely used model in the field of systems engineering for describing the processing of data. The taxonomy classifies the most common adverse events into four basic groups that to a certain extent are arranged around a well-known processing model: the input-process-output (IPO) model. The first three groups (information collection, information processing and information dissemination) represent the input, process and output stages of the model respectively. The fourth group (invasions) is not related to that model as invasions are not only caused by technology and invasive adverse events do not always involve personal data, rather they directly affect data subjects. As such, we consider only some aspects of these events that involve personal data throughout the collection and disclosure stages of the lifecycle. We use the IPO model as a starting point towards describing these events at a detailed level of abstraction. As such, we adopt the APDL as a model around which we arrange these events. We map the basic groups of adverse events onto the stages of the data lifecycle. Additional detail about the conceptual relationship between the categories of the taxonomy of privacy and the stages of the APDL model is illustrated in [3]. Each type of an adverse threat event can be characterised by a set of attributes according to the nature of a processing operation in each stage of the lifecycle that reflects the manner in which personal data is collected, processed and disseminated.

Step 4: Privacy Harm Analysis. *Privacy violations.* Once privacy vulnerabilities, threat sources and threat events are identified, privacy violations can be identified as illegitimate or unanticipated data-processing activities that may result from the occurrence of threat events without negative consequences on data subjects. In particular, for each possible exploitation, privacy violations are activities that can be conducted without adverse actions taken against data

subjects, as well as without their knowledge. For each type of privacy violation, its degree and scope need to be identified in relation to its attributes.

Privacy harms. Once privacy vulnerabilities, threat sources and threat events are identified, privacy harms can be derived from these events as potential adverse consequences on the privacy of data subjects. We use the categories of privacy harms of [12] that have been identified in previous attempts from a legal perspective [16] and [13]. In particular, privacy harms are classified into: physical; economic or financial harms; mental or psychological harms; harms to dignity or reputation; and societal or architectural harms [12]. We arrange these categories of harms around the APDL model according to its lifecycle stages, associated data-processing activities and their corresponding threat events. Additional detail about mapping these categories onto the stages of the APDL model is illustrated in [3]. For each type of privacy harm, its type, adverse consequences and affected data subjects need to be identified in relation to its attributes.

6 A Case Study

6.1 Overview

The European Electronic Toll Service (EETS) aims to support interoperability between electronic road toll systems at a European level to calculate and collect road-usage tolls. The main actors involved in the EETS are service providers, toll chargers and users. EETS providers are legal entities that grant access to EETS to road users [18]. Toll chargers are public or private organisations that are responsible for levying tolls for the circulation of vehicles in an EETS domain [18]. A user is an individual who subscribes to an EETS provider in order to get access to EETS [18]. By signing a contract, a user is required to provide a set of personal data specified by a responsible toll charger, as well as to be informed about the processing of their personal data in relation to applicable law and regulations. Accordingly, the EETS provider provides the user with an On-Board Unit (OBU) to be installed on-board a vehicle to collect, store, and remotely receive and transmit time, distance and location data over time. This data, together with the user's and vehicle's parameters, are specified to declare the toll of circulating a vehicle in a specific toll domain [17].

Due to space limitations, we do not provide an exhaustive list of vulnerabilities, etc. Rather, we give examples to illustrate the usability and applicability of our approach in this particular context.

6.2 Context Establishment

All useful information that helps establish the context has been already captured by the APDL model in [1]. The establishment of the context in which personal data is collected and processed consists of three steps. The first step is to specify the types of personal data along with their attributes (captured by classes stereotyped by «PersonalData») and the main purpose for which this data is

collected and processed (captured by a class stereotyped by «Purpose» along with its lawfulness, fairness and proportionality). With reference to the APDL model, the main purpose is to ‘electronically calculate and collect road-usage tolls’ and the types of personal are:

- Identification and contact data — **EETSUser**: user ID, name, billing address (collected from the EETS user whether the user is the driver, owner, lesser or fleet operator of the vehicle)
- Vehicle classification parameters — **Vehicle**: licence plate, classification code (collected from the EETS user)
- Location data — **LocationData**: time, distance, place (collected by OBUs)

The second step is to specify or model both actual data-processing activities and privacy-related processes required by legal frameworks and standards in each stage of the APDL model. These processes are abstractly captured from classes stereotyped by «Initiation», «Collection», etc. With a focus on location data, we illustrate a data-processing activity in the collection stage of the APDL model: it is abstractly captured from the **CollectingUsageData** class, which is stereotyped by «Collection». The stereotyped class also captures other important details: location data sources (OBUs), available choices (the user is entitled to subscribe to EETS with the EETS providers of their choice among other choices: the national or local manual, automatic or electronic toll services), collection method (OBUs using satellite positioning systems), consent type (implicit by signing a contract) and relevant GPS principles (Collection Limitation). In addition, processes are concretely captured from classes stereotyped by «Stage-Activity», «StageAction» and «StageEvent». Each stage activity contains a set of actions that represent its executable steps and a set of events that cause the execution of these actions. The data-processing activity is concretely captured from the **CollectingLocationData** class, which is stereotyped by «StageActivity». At this level of detail, it aims to collect road-usage data to be used for tolls declaration and calculation. The stereotyped class also captures other important details in terms of constraints: pre-conditions (the privacy notice is communicated to EETS users at or before the collection time in a clear and concise manner; their implicit consent is obtained at or before the collection time in an informed manner. This activity is decomposed into two classes: **CollectLocationData** and **Collect**, which are stereotyped by «StageAction» and «StageEvent» respectively. **CollectLocationData** class captures the time of usage, the covered distance and the place on which the vehicle is circulating on a particular toll domain for tolls declaration and calculation. The **Collect** class captures the occurrence of circulating a vehicle on a particular toll domain to collect location data.

The third step is to specify or model involved actors (captured by classes stereotyped by «LifecycleRole» and «LifecycleActor»). Each lifecycle stage includes a number of lifecycle roles, each of which is played by different actors according to their capabilities and responsibilities. With reference to the APDL model, **CollectionAgent** is a type of the data processor role that consists of logically related activities for collecting road usage data, and **ServiceProvider**

is a type of involved actors who are capable of, and responsible for, performing the activities of the collection agent as a role to which are assigned. Responsibilities are captured from stage activities in which a lifecycle actor participates and to which a lifecycle role is associated.

Establishing the context in which personal data is collected and processed requires specifying or modelling ‘personal data’, ‘data-processing activities’ and ‘involved actors’ along with their roles and responsibilities. The APDL model has served as a preliminary acquisition step to capture all required data that support privacy risk analysis and compliance checking.

6.3 Vulnerability Analysis

In our approach the focus is on vulnerabilities of *primary assets* to protect the privacy of data subjects and ensure contextual integrity. The first step of vulnerability analysis is to develop a baseline model of the processing of personal data. The baseline model captures all appropriate data-processing activities in all stages of the APDL model. In order to develop a baseline model, we need to establish a context-relative processing norm for each data-processing activity. The main elements that constitute these norms are captured from stage activities in the established context. Due to space limitations, we identify only a context-relative processing norm for the *CollectingLocationData* activity:

In the context of EETS, the collection of a certain type of personal data (location data: time, distance, place) about an EETS user (acting as a data subject) by an EETS provider (acting as a data processor on behalf of a toll charger) is governed by processing principles derived from applicable legal frameworks ... and standards ...

In this case, legal framework principles — for example, DIRECTIVE 95/46/EC — are as follows. Personal data must be: processed fairly and lawfully; collected for specified, explicit and legitimate purposes; adequate, relevant and not excessive; and accurate and up to date. In addition, the relevant GPS principle is *Collection Limitation*. Importantly, principles of legal frameworks and standards are modelled as pre- and post-conditions for each stage activity.

Once all context-relevant processing norms are defined in relation to the APDL model, a complete baseline model can be developed to serve as the basis for deriving privacy vulnerabilities. The second step of vulnerability analysis is to derive all possible vulnerabilities of the primary assets from the identified context-relevant processing norms. They can be derived by examining all the main elements that constitute each processing norm — i.e. any possible breach of a processing norm can be derived as a vulnerability. With reference to the above processing norm, a possible vulnerability with regards to *attributes*, as an element, is ‘an improper data model’ (PV.1) that directly or indirectly links location data to users’ IDs. Another possible vulnerability with regards to *processing principles*, as an element, is ‘a lack of data minimisation’ (PV.2) that facilitates inadequate, irrelevant and excessive collection of location data in an

interval basis, which is not necessary for the main purpose. Additional examples of privacy vulnerabilities are listed in [3].

6.4 Threat Analysis

Threat sources. In reference to the established context, EETS providers (TS.1) are involved in the processing of ‘identification and contact data’ and ‘location data’ by playing the role of data processors who grant access to EETS to EETS users. They may act accidentally or deliberately as threat sources while they process personal data lawfully to calculate and communicate personalised fees (road-usage tolls) for each EETS user by the end of the tax period — or unlawfully for further processing with the motivation of profiling EETS users, discriminatory social sorting or providing better services. The utility of ‘location data’ and ‘identification and contact data’ in this context makes such data highly valuable to EETS providers. The value of this data stimulates the motives of EETS providers to exploit vulnerabilities of the primary assets. In particular, it has a market value when it is exploited by EETS providers for administrative and commercial purposes — for example, it gives an EETS provider a competitive advantage with respect to their competitors. According to the attributes of a threat source, EETS providers are insiders and institutions. EETS providers have technical skills and detailed background knowledge about conceptual, logical and physical data models, as well as about the processing operations. It also implies that they have legitimate privileges to collect and process location-related data according to their roles and responsibilities. Based on these, they have access rights to both the ‘fine-grained location data’ and ‘identification and contact data’. They also have reasonable resources (both technical and financial) to get benefit from the values of the collected data by creating comprehensive and identifiable profiles. Additional examples of threat sources are listed in [3].

Threat events. In a straightforward implementation of the EETS architecture, the calculation of road-usage tolls is performed remotely at EETS providers’ back-office systems. The OBU collects, stores, and remotely receives and transmits time, distance and place over time to the EETS provider’s back-office systems. These systems are in charge of processing location data to calculate personalised road-usage tolls and communicate the final premium to EETS user at the end of the tax period. As mentioned, a threat event occurs as a result of a successful exploitation of one or more vulnerabilities by one or more threat sources. With reference to the identified vulnerabilities and threat sources, we identify the most significant threat events with the potential to adversely impact the privacy of EETS users that may happen at specific points in time. The identification of these events needs to be conducted according to the stages of the data lifecycle.

In the collection stage, for example, threat events that may lead to privacy violations or harms are related to the manner in which personal data is collected. By exploiting PV.2, TS.1 may use OBUs to excessively collect irrelevant location data (TE.1) in fine-grained manner about EETS users. With reference to the adapted taxonomy of adverse privacy events, this threat event is a type of

‘surveillance’. It is characterised as continuous, overt and extensive: continuous via the collection of location data over time; overt via informing the EETS user about the manner in which location data will be collected when signing the contract; and extensive via the excessive collection of location data in a fine-grained manner throughout national and international toll domains. Surveillance outside toll domains implicates reasonable expectations of privacy as it may reveal hidden details that would not ordinarily be observed by others. Additional examples of threat events, along with the corresponding threat sources and privacy vulnerabilities, are listed in [3].

6.5 Harm Analysis

Privacy violations. In the collection stage, for example, ‘passive collection of location data outside toll domains’ is a privacy violation that may result from the occurrence of the threat event ‘excessive collection of location data’, which results from the successful exploitation of ‘a lack of data minimisation’ by EETS providers. Its degree is excessive as it collects fine-grained location data outside toll domains, whether they are national and international. Its scope is individuals — i.e. those who are subscribed to EETS. This privacy violation is considered as an illegitimate and unanticipated data-processing activity without adverse consequences. In particular, fine-grained location data is collected in ways EETS users would not reasonably expect, as well as this data is collected passively without the knowledge and consent of EETS users. In addition, the collection of location data outside toll domains does not have legitimate grounds as they are irrelevant and inadequate for the purposes for which location data is collected. Most importantly, this privacy violation is assumed to be without adverse actions taken against EETS users.

Privacy harms. Privacy harm analysis is the most important step of any privacy risk-analysis approach. Harms are derived from the undesirable consequences of threat events as potential adverse actions taken against data subjects. In this paper, we consider only the objective category of privacy harms as the subjective category is mainly about the perception of unwanted observation.

For each stage of the data lifecycle, the potential undesirable consequences of each threat event need to be identified. Then, these consequences need to be analysed to determine whether they can partially contribute to, or completely lead to a negative action that uses personal data against the data subject in an unanticipated or coerced manner. Most broadly, a privacy harm may result from a series of adverse consequences of multiple threat events. In the collection stage, for example, the main undesirable consequence of TE.1 is gathering a large amount of fine-grained location data that has been collected over time as comprehensive driving records (UC.1), which may include complete driving history or driving history for a specific period for EETS users. Additional examples of undesirable consequences are listed in [3].

By analysing the identified undesirable consequences, together with the relevant privacy vulnerabilities and threat sources, we can derive a reasonable set

of privacy harms. For example, the privacy harm ‘increased car insurance premium’ (PH.1) occurs as EETS providers can make excessive inference to derive EETS users’ driving patterns and share anonymised patterns with car insurance providers (TS.7). Insurance providers may make inference to re-identify current and potential customers with the aim of calculating car insurance premium based on the types of vehicle use and health conditions, which are derived from their driving patterns. Additional examples of privacy harms, along with associated threat sources, privacy vulnerabilities, threat events and undesirable consequences of these events, are listed in [3].

7 Conclusion

We have presented an approach that helps support engineers in identifying and analysing potential privacy risks in a comprehensive and contextual manner. It refers to fundamentals from the legal privacy literature to refine key concepts and assessable risk factors, as well as the conceptual relationships among these factors. Such fundamentals help support the distinction between privacy harms and violations and their main sources by providing boundaries and properties of privacy harms. In addition, fundamentals bring the legal and social layers into consideration by defining context-relative processing norms. They also facilitate the identification of adverse events in a systematic manner by providing a taxonomy of harmful activities and their corresponding harms. They also support the taxonomy by providing two main principles: (1) the limiting principle to help protect against reduction of the concept of privacy, and (2) the rule of recognition to support the identification of novel privacy harms as they emerge.

We limit our approach to a risk model and analysis approach that describes how combinations of risk factors are identified and analysed at a consistent level of detail. In order to propose a complete risk-assessment methodology, an assessment approach that associates values with the risk factors needs to be developed to functionally combine the values of those factors and estimate the levels of the identified risks.

Acknowledgments. The authors would like to thank the reviewers for their constructive comments.

References

1. Alshammari, M., Simpson, A.C.: A UML Profile for Privacy-Aware Data Lifecycle Models. In: Proceedings of the 1st International Workshop on SECurity and Privacy Requirements Engineering (SECPRE 2017). pp. 189–209. Springer (2017)
2. Alshammari, M., Simpson, A.C.: Personal Data Management: An Abstract Personal Data Lifecycle Model. In: Proceedings of SPBP’17: Workshop on Security and Privacy-enhanced Business Process Management. pp. 685–697. Springer (2017)
3. Alshammari, M., Simpson, A.C.: Towards an Effective PIA-based Risk Analysis: An Approach for Analysing Potential Privacy Risks. <http://www.cs.ox.ac.uk/publications/publication11663-abstract.html> (2017)

4. BSI (Bundesamt für Sicherheit in der Informationstechnik): Risk analysis on the basis of IT-Grundschutz, BSI Standard 100-3. https://www.bsi.bund.de/EN/TheBSI/thebsi_node.html (2008)
5. BSI (Bundesamt für Sicherheit in der Informationstechnik): IT-Grundschutz-Kataloge. https://www.bsi.bund.de/DE/Themen/ITGrundschutz/ITGrundschutzKataloge/itgrundschutzkataloge_node.html (2011)
6. Calo, R.: The Boundaries of Privacy Harm. *Indiana Law Journal* **86**, 1131–1162 (2011)
7. Cavoukian, A.: Privacy by Design. <https://www.privacybydesign.ca/content/uploads/2009/01/privacybydesign.pdf> (2009)
8. Cavoukian, A., Shapiro, S., Cronk, R.J.: Privacy Engineering: Proactively Embedding Privacy, by Design. <https://www.privacybydesign.ca/content/uploads/2014/01/pbd-priv-engineering.pdf> (2014)
9. Clarke, R.: Privacy Impact Assessment: Its Origins and Development. *Computer Law and Security Review: The International Journal of Technology and Practice* **25**(2), 123–135 (2009)
10. Commission Nationale de l’Informatique et des Libertés (CNIL): Methodology for Privacy Risk Management. <https://www.cnil.fr/sites/default/files/typo/document/CNIL-ManagingPrivacyRisks-Methodology.pdf> (2016)
11. Fineberg, A., Jeselon, P.: A Foundational Framework for a Privacy by Design — Privacy Impact Assessment. <http://privacybydesign.ca/content/uploads/2011/11/PbD-PIA-Foundational-Framework.pdf> (2011)
12. Joyee De, S., Le Métayer, D.: PRIAM: A Privacy Risk Analysis Methodology. In: 11th International Workshop on Data Privacy Management and Security Assurance. pp. 221–229. Springer (2016)
13. Nissenbaum, H.F.: Privacy in Context: Technology, Policy, and the Integrity of Social Life. Stanford University Press (2009)
14. Object Management Group: OMG Unified Modeling Language (OMG UML). <http://www.omg.org/spec/UML/> (2015)
15. Oetzel, M.C., Spiekermann, S.: A Systematic Methodology for Privacy Impact Assessments: A Design Science Approach. *European Journal of Information Systems* **23**(2), 126–150 (2014)
16. Solove, D.J.: A Taxonomy of Privacy. *University of Pennsylvania Law Review* **154**(3), 477–564 (2006)
17. The European Commission: The European Electronic Toll Service (EETS): 2011 Guide for the Application of the Directive on the Interoperability of Electronic Road Toll Systems. http://ec.europa.eu/transport/themes/its/road/application_areas/electronic_pricing_and_payment_en (2011)
18. The European Union: Official Journal of the European Communities: Commission Decision 2009/750/EC of 6 October 2009 on the definition of the European Electronic Toll Service and its technical elements. <http://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32009D0750&from=EN> (2009)
19. Wright, D.: The State of the Art in Privacy Impact Assessment. *Computer Law & Security Review* **28**(1), 54–61 (2012)
20. Wright, D., Wadhwa, K., De Hert, P., Kloza, D.: A Privacy Impact Assessment Framework for data protection and privacy rights. <http://www.piafproject.eu/Deliverables.html> (2011)