A Principled Approach for Engineering Privacy by Design

Majed Alshammari

Worcester College

University of Oxford

A dissertation submitted for the degree of

Doctor of Philosophy

Hilary Term 2019

Acknowledgements

First and foremost, I would like to express my deepest thanks and sincere gratitude to my supervisor, Dr Andrew Simpson, for his constant encouragement and guidance while carrying out this research.

I would like to give special thanks to my parents, siblings and wife for their patience, support and love. I would also like to extend my thanks and appreciation to my friends, especially Dr Nasser Alanazi, who has provided constant inspiration and confidence.

Last but not the least, I am very grateful to the Saudi Arabian government for the sponsorship of this accomplishment that would not have been possible without their financial support. I am also grateful to the Department of Computer Science and Worcester College for their academic expenses grants that include conference attendance, travel and other costs incurred for research purposes.

Abstract

Privacy is a multi-faceted concept that has several aspects. It is subjective in nature, not least because it is influenced by a variety of factors, including societal demands, which evolve over time, and technological developments. With the advent of information technologies, legal frameworks and standards alone are not sufficient to preserve the privacy of data subjects. As a response, Privacy by Design (PbD) has emerged as a proactive approach for embedding privacy into the early stages of the design process. Challenges involved in engineering PbD include a lack of holistic methodologies that address the plurality and contextuality of privacy and support the translation of its principles into engineering activities. In this dissertation, we investigate various methods for engineering PbD that capture and address privacy issues in the early stages of the design process. We also investigate how to model the key aspects of abstract privacy principles stated in legal frameworks and standards to bridge the semantic gap between technical and normative concepts. This gives rise to the Abstract Personal Data Lifecycle (APDL) model, which serves as an abstract model for personal data lifecycles. We also define a UML profile for the APDL model to represent data-processing activities in a way that is amenable to risk analysis and compliance checking. In addition, we develop a privacy risk model that defines the main factors that have impacts on privacy risks along with their assessable attributes and conceptual relationships. Based on this, we develop analysis and assessment approaches that illustrate how combinations of these factors are analysed and used as inputs to assess the levels of risk. Furthermore, we characterise privacy protection as a quality attribute by means of a general quality attribute scenario to avoid non-operational or overlapping definitions. Based on this, we develop a tactical approach that identifies privacy architectural strategies as collections of tactics, which are described through design patterns, to support the adoption of Privacy-Enhancing Technologies (PETs), and to specify, implement and justify various levels of privacy protection. Together, these contributions give rise to a principled approach for engineering PbD that captures privacy concerns in a comprehensive manner; addresses these concerns at an architectural level; and reasons about the compliance of architectural choices with legal frameworks and standards. It is aided by techniques and tools, which provide procedures with a prescribed language and notation, to accomplish its activities.

Contents

List of Figures

"Privacy seems to be about everything, and therefore it appears to be nothing."

— Daniel J. Solove

1 Introduction

Contents

1.1 Motivation

Privacy is a multi-faceted concept that has several aspects, including legal, social, psychological, economical and political aspects [1], as well as various states, dimensions or types, such as information privacy, bodily privacy, communications privacy and territorial privacy. It is subjective and contextual in nature, not least because it is influenced by a variety of factors, including cultural values and social norms, which evolve according to time, contexts and technological developments [1, 2]. This plurality and contextuality indicates privacy's dynamic nature, which introduces complexity and variability of privacy issues [2].

With a focus on information privacy, it is important to note that the terms 'privacy' and 'data protection' are related but not interchangeable [3, 4]. Privacy

refers to the data subject's perspective whereas data protection refers to the organisation's perspective [4]. In this dissertation, we use the term 'privacy protection' to consider both perspectives, with the aim of ensuring that personal data is collected, processed and disseminated fairly and lawfully.

In recent years, a large amount of personal data has been collected and processed by a wide range of organisations (including governments). It is well recognised that personal data has economic and social value both to data subjects and to those who are interested in, or actually involved in, the processing of personal data. It has, for example, a market value when it is used for administrative or commercial purposes, e.g. targeted advertising [5]. It also has a nuisance value when it is used for unfair or malicious purposes, e.g. spam [5]. In either case, personal data is an 'infrastructural resource' that might be used for unlimited purposes[1]. It can be used to provide a number of growth opportunities, or to bring benefits across society in ways that would not be expected when the data was originally collected, including benefits from scientific or historical research, or the development of new products or services[2].

In contrast, personal data needs to be managed responsibly as the processing of personal data incurs a legal liability on organisations involved in its collection and processing to ensure that such data is processed fairly and lawfully. In practice, inappropriate processing of personal data can lead to privacy violations and harms, which may have adverse consequences on organisations and data subjects, whether these consequences are physical, financial or incorporeal. This raises a number of legitimate concerns over the collection, retention, access, use, destruction and disclosure of personal data to third parties, especially given recent advances in information and communication technologies.

These concerns, in part, have motivated the development of legal frameworks and standards for governing the collection, processing and dissemination of personal data, such as the EU's General Data Protection Regulation (GDPR) [6] and the

[1]The Organisation for Economic Co-operation and Development (OECD): Maximising the economic and social value of data. Retrieved from: `http://www.oecd.org/sti/ieconomy/enhanced-data-access.htm` [Accessed 5 February, 2018]

[2]ibid.

Global Privacy Standard (GPS) [7]. As a result, software engineers are increasingly expected to give appropriate consideration to privacy and data protection issues throughout the system development lifecycle. However, legal frameworks and standards alone are insufficient to preserve the privacy of data subjects [8, 9]. Thus, they need to be accompanied with methodologies, guidelines and tools to aid software engineers in addressing the complexity and variability of privacy issues in the early stages of the design process.

As a response, Privacy by Design (PbD) [10] has emerged as a proactive and creative approach for embedding privacy requirements into the early stages of the design process. It has been advocated by the former Information and Privacy Commissioner of Ontario, amongst others; it aims to achieve adequate levels of privacy protection and ensure compliance with legal frameworks and standards [10]. Its principles are based on the Fair Information Practice Principles (FIPPs) [11], and act as a universal framework for integrating privacy into three main areas of application: information technology, business practices, and physical designs and infrastructures [12].

1.2 Research Problem

Typically, legal frameworks and standards in general, and those related to privacy and data protection in particular, do not rely on rigorous models that specify key concepts along with their properties and relationships [13]; rather, the principles of such frameworks and standards are often given at a high level of abstraction. This leads to practical challenges with respect to translating legal, social and ethical concerns into systems requirements [14] and reasoning about compliance with these abstract principles. From an engineering perspective, it is difficult to use such abstract principles as a means of specifying design options or justifying architectural choices [15]. There is a semantic gap between technical and normative concepts [13]. This semantic gap leads to a disconnect between policy-makers and software engineers in terms of the actual meaning of privacy and the ways in which technical solutions can be developed to comply with the principles of legal

frameworks and standards [14]. As such, abstract privacy principles need to be translated into concrete, auditable and functionally enforceable goals to aid engineers in specifying design decisions that meet privacy compliance requirements [15].

The PbD approach serves as a reference framework and its principles are not operational in their current state [16]. Accordingly, privacy engineering has emerged as a new field of research and practice that aims to apply engineering principles and processes in developing and maintaining software systems in a systematic and repeatable manner, to achieve various levels of privacy protection [2]. To properly align these concepts with each other, PbD explains 'What to do' to achieve adequate levels of privacy protection, and privacy engineering explains 'How to do it' by defining privacy protection as a quality attribute in systems engineering [12]. PbD can be considered as a facilitator for privacy engineering as it provides foundational principles that need to be translated into fundamental engineering activities that are part of the SDLC [2]. In turn, privacy engineering extends PbD [2] by developing and evaluating methods, techniques and tools that capture and address privacy issues in a systematic manner during the system development lifecycle (SDLC) [17]. In this context, however, existing approaches to PbD are usually stage- or domain-specific [16]; they provide partial solutions to the overall problem and aid engineers during specific stages of the SDLC. The deficiency in existing approaches lies in an inability to integrate disconnected methods under a common approach. This indicates that there is a need for a holistic approach that integrates and improves existing approaches to PbD to address the complexity and variability of privacy issues, with a view to understanding how to:

1. identify the need for privacy;
2. find the places where it is needed;
3. determine what privacy issues should be addressed;
4. determine what degree of privacy must be achieved; and
5. illustrate how these issues can be addressed.

Such an approach should capture privacy issues in a comprehensive manner, address these issues in the early stages of the design process, and ensure and demonstrate

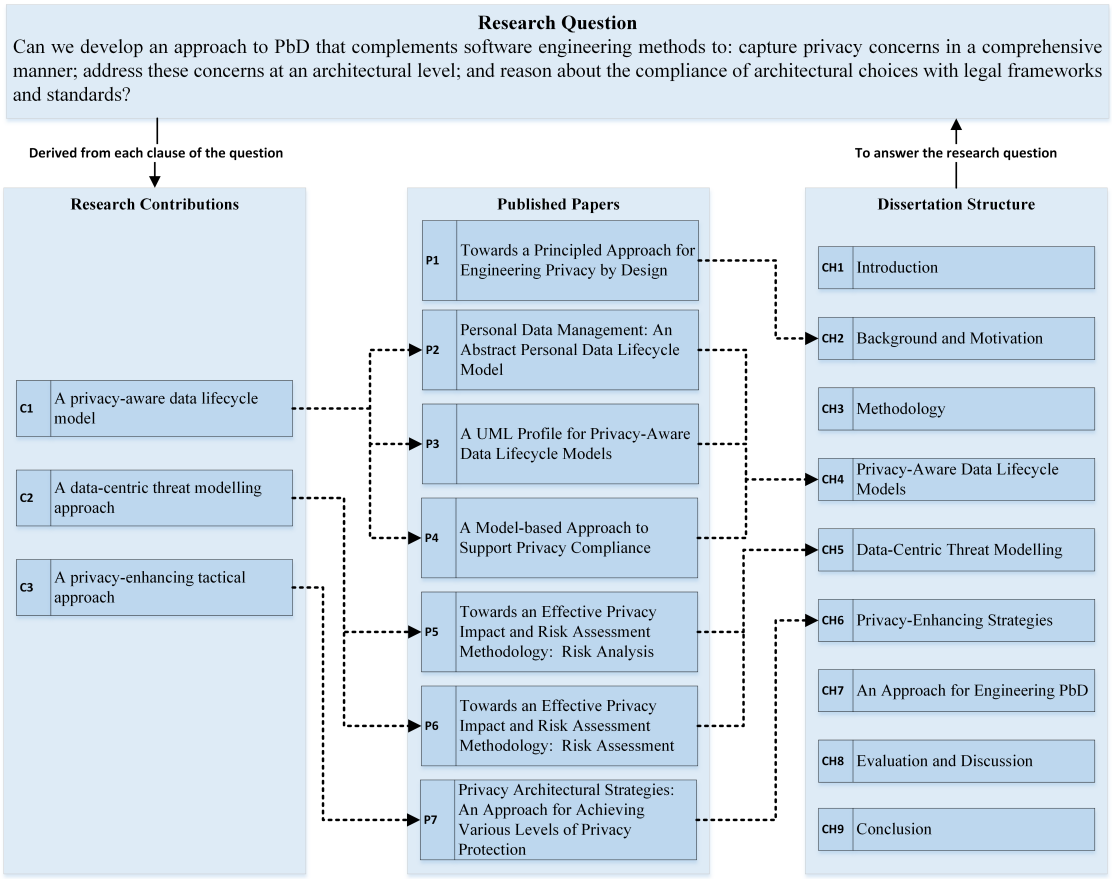

Figure 1.1: The structure of the dissertation

compliance with legal frameworks and standards. It has the potential to aid engineers to strike a balance between the need for personal data, data subjects' concerns, Data Protection Authorities' goals and software architects' strategies when making their architectural choices.

1.3 Contribution

Given these challenges, the research question addressed by this dissertation can be framed as:

> *Can we develop an approach to PbD that complements software engineering methods to: capture privacy concerns in a comprehensive manner; address these concerns at an architectural level; and reason about the compliance of architectural choices with legal frameworks and standards?*

Driven by the research question, three main research contributions are made to the field of study, as shown in Figure 1.1.

C.1 A privacy-aware data lifecycle model. An Abstract Personal Data Lifecycle (APDL) model has been developed to serve as an abstract representation for personal data lifecycles — where a data lifecycle is defined by a set of stages through which personal data moves during its lifetime, associated activities, and involved actors. It distinguishes between the main types of operations that can be performed on personal data during its lifecycle by outlining the various distinct activities for each operation.

A UML profile for the APDL model has been defined to support the modelling of privacy-related concepts along with associated properties and relationships. It also supports the management and traceability of personal data during its lifecycle. Importantly, it represents data-processing activities in a way that is amenable to risk analysis and compliance checking: it facilitates the identification of these activities that may lead to privacy violations and harms, and helps support the demonstration of privacy compliance with legal frameworks and standards. As such, it provides a common language that facilitates a meaningful participation of, and communication between, multiple stakeholders.

C.2 A data-centric threat modelling approach. A privacy risk model that goes beyond traditional security risk models to consider the dynamic and contextual nature of privacy has been developed. It takes into account legal, organisational, societal and technical aspects of privacy. The privacy risk model is built upon fundamentals from the legal privacy literature to underpin the main concepts, the key risk factors, and the conceptual relationship between these factors. It aims to facilitate data-centric system threat modelling, which is focused on protecting particular types of data within systems. In particular, it uses the concept of primary assets to focus on personal data that is directly concerned with processing operations and those processes required by legal frameworks and standards.

In addition, a methodical approach for identifying and analysing potential privacy risks has been developed. The approach is built upon the risk model that defines the main factors that have impacts on privacy risks along with their meanings, properties and relationships. It describes how combinations of risk factors are identified and analysed to ensure adequate coverage of the problem space at a consistent level of detail. Furthermore, a methodical approach for assessing privacy risks has been developed. It is also built upon the risk model. The approach defines the dependencies between the nominal and assessable attributes of key risk factors, and establishes a set of assessment rules that specify the range of values the risk factors can assume.

C.3 A privacy-enhancing tactical approach. An architectural approach for applying privacy tactics has been developed to support the adoption of PETs in the early stages of the design process with the aim of achieving various levels of privacy protection. It adopts a set of privacy protection goals to translate the abstract privacy principles stated in legal frameworks and standards into desired privacy protection goals in a contextual manner. It also defines context-specific architectural strategies as a set of architectural tactics, which are described through design patterns and implemented by PETs, to achieve the desired protection goals. In particular, each architectural tactic specifies fundamental design decisions that contribute to, or achieve, a desired protection goal. The approach considers context-related factors that influence the degree to which privacy is required. It also considers the relationships between these factors and technical measures that might be adopted as privacy controls (in the form of architectural strategies).

The architectural strategies are considered as means for mapping privacy requirements onto suitable software architectures to specify, implement and justify various levels of privacy protection as the default setting. In addition, the approach provides a set of selection criteria to aid software engineers help determine what combination of tactics, design patterns and PETs will

achieve, or contribute to the achievement of, the desired protection goals. These criteria can also aid software engineers to justify and reason critically about their architectural choices.

1.4 Dissertation Structure

The reminder of the dissertation is organised as follows (see Figure 1.1).

- **Chapter 2** gives a detailed background to PbD and illustrates the motivation for the contributions of this dissertation. It also gives a relatively detailed discussion of privacy-related issues together with the main challenges and limitations of adopting the PbD approach in practice from an engineering perspective.

- **Chapter 3** describes the methodology used for carrying out the research described in this dissertation. It also introduces two main case studies: the first one is used as an illustrative case study, whereas the second case study is used for evaluating the contributions of the research.

- **Chapters 4 – 6** give a detailed description of the main contributions of this dissertation, **C.1** – **C.3**.

- **Chapter 7** gives a relatively detailed description of a principled approach for engineering PbD that synthesises the main contributions of this dissertation, **C.1** – **C.3**. It also provides some examples from the illustrative case study (the European Electronic Toll Service (EETS)) to demonstrate the applicability and usefulness of the principled approach in this particular context.

- **Chapter 8** presents an evaluation of the principled approach by discussing the efficacy of each technique in the context of the case study. It also gives a relatively detailed discussion about the significance of the contributions of this dissertation with reference to the criteria of Chapter 2.

- **Chapter 9** summarises the contributions of this dissertation that can be applied to the current practice and highlights their weaknesses and limitations. It also outlines opportunities for future work to extend the contributions of this dissertation.

1.5 Papers Arising from This Work

This research has contributed to the following publications (in the order of their appearance in this dissertation).

P.1 Alshammari, M. and Simpson, A.: Towards a Principled Approach for Engineering Privacy by Design. In Proceedings of the Annual Privacy Forum (APF 2017), Lecture Notes in Computer Science, vol 10518, pp. 161–177. Springer (2017)

P.2 Alshammari, M. and Simpson, A.: Personal Data Management: An Abstract Personal Data Lifecycle Model. In Proceedings of the Business Process Management Workshops (BPM 2017), Lecture Notes in Business Information Processing, vol 308, pp. 685–697. Springer (2017)

P.3 Alshammari, M. and Simpson, A.: A UML Profile for Privacy-Aware Data Lifecycle Models. In Proceedings of the International Workshop on Security and Privacy Requirements Engineering (SECPRE 2017). Lecture Notes in Computer Science, vol 10683, pp. 189–209. Springer (2017)

P.4 Alshammari, M. and Simpson, A.: A Model-based Approach to Support Privacy Compliance. Information & Computer Security 26(4), 437–453 (2018)

P.5 Alshammari, M. and Simpson, A.: Towards an Effective Privacy Impact and Risk Assessment Methodology: Risk Analysis. In Proceedings of the 13th DPM International Workshop on Data Privacy Management (DPM 2018), Lecture Notes in Computer Science, vol 11025, pp. 209–224. Springer (2018)

P.6 Alshammari, M. and Simpson, A.: Towards an Effective Privacy Impact and Risk Assessment Methodology: Risk Assessment. In Proceedings of the 15th International Conference on Trust, Privacy and Security in Digital Business (TrustBus 2018), Lecture Notes in Computer Science, vol 11033, pp. 85–99. Springer (2018)

P.7 Alshammari, M. and Simpson, A.: Privacy Architectural Strategies: An Approach for Achieving Various Levels of Privacy Protection. In Proceedings of the 2018 Workshop on Privacy in the Electronic Society (WPES'18), pp. 143–154. ACM (2018)

Figure 1.1 outlines the aspects of this dissertation that have been published in journals, as well as peer-reviewed workshop and conference proceedings, and shows how these papers relate to the research contributions.

"A right to privacy is neither a right to secrecy nor a right to control but a right to appropriate flow of personal information."

— Helen Nissenbaum

2 Background and Motivation

Contents

2.1 Introduction

This chapter gives a relatively detailed background to data privacy and Privacy by Design (PbD) and illustrates the motivation for the contributions of this dissertation. The methodology undertaken to identify and review relevant studies is described in

Chapter 3. We start the review by analysing a wide array of privacy definitions, states and types, as well as by investigating approaches for conceptualising privacy. This helps understand the plurality and contextuality of privacy issues by giving a brief introduction of the main characteristics of privacy expectations, concerns and harms. Then, we frame these definitions for privacy engineering by proposing an operational definition for data privacy to be used for characterising privacy protection as a quality attribute. Next, we discuss the main challenges of translating the principles of PbD into engineering activities. Based on these challenges, we analyse the underlying goals of these principles together with the principles of the GPS. In particular, we identify the following points of reference when making the review: the identification of privacy concerns in a structured, comprehensive and contextual manner; the determination of various levels of privacy protection; the achievement of these levels at an architectural manner by adopting privacy patterns and PETs; and the reasoning of the compliance of architectural choices with legal frameworks and standards. Then, we analyse existing approaches to PbD with respect to these points of reference to both identify homogeneity of privacy-related concepts, techniques and tools and, importantly, to note discordance and dissonance. Finally, we summarise the key points resulting from the review as a list of challenges associated with engineering PbD, and establish a set of criteria for developing and evaluating holistic engineering approaches to PbD.

This chapter is organised as follows. Section 2.2 gives an overview of privacy dimensions, states and types. It also discusses essential concepts and definitions of privacy to illustrate its complexity and variability. Section 2.3 explains the origins, meaning, philosophy and foundational principles of PbD. It also gives a relatively detailed discussion of privacy-related issues together with the main challenges of translating the principles of PbD into engineering activities. Then, Section 2.4 provides a detailed examination of how existing approaches to PbD address these issues and gives a detailed explanation of the limitations of these approaches. Section 2.5 gives a relatively detailed discussion of how these approaches might be complemented by methods, techniques or tools to overcome their limitations.

Based on these, Section 2.5.1 gives a summary of key points from the review of existing approaches as a list of challenges for engineering PbD (incorporating the PbD approach into software engineering practices). Then, Section 2.5.2 establishes a set of criteria against which engineering approaches to PbD can be evaluated. Finally, Section 2.6 gives a brief summary of the chapter.

2.2 Privacy

In this section, we start by giving an overview of the plurality of privacy in terms of types, states, categories and dimensions. In addition, we explore essential concepts and definitions, and illustrate how these definitions capture specific aspects of privacy. Further, we distinguish between the concepts of privacy and security, and explain the relationship between them.

2.2.1 Essential Concepts and Definitions

Privacy Definitions and Dimensions

Privacy is a multi-faceted concept that has various types [18], states [19], clusters [20], categories [21] and dimensions [22], including information privacy (also known as data privacy, which is the type of privacy this dissertation addresses), bodily privacy, communications privacy and territorial privacy. Due to its dynamic nature, it is difficult to find a global, consistent, satisfying and overarching definition of privacy [23]. In recent decades, there have been several attempts to capture all privacy-related states [1, 19, 20, 21, 24, 25]. However, these definitions are not overarching; rather, they capture specific aspects of privacy. Theories that have attempted to conceptualise privacy have been either very broad or too narrow [23]. In particular, these conceptions of privacy build upon one of the following definitions: the right to be left alone, limited access, secrecy, control of personal information, personhood, or intimacy [23].

The limitation in privacy conceptualisation has led to the importance of elaborating a categorisation that covers all aspects of privacy [26]. In particular, the taxonomy of privacy [26] was developed to classify all possible types of activities

that may lead to privacy violations and harms into categories according to shared characteristics in a comprehensive and concrete manner. According to the taxonomy, privacy is better understood as a family resemblance concept (groups of related but distinct harmful activities) rather than referring to abstract definitions, and addressed in a pluralistic and contextual manner [26]. The rationale behind developing such a taxonomy is that privacy is not only about controlling access over personal data; rather, it goes beyond that to mitigate potential risk arising from inappropriate collection, processing and dissemination of personal data when it is no longer under the data subject's control [27]. However, the taxonomy was developed based upon a set of socially recognised privacy violations addressed in specific sources without defining a set of criteria for inclusion or exclusion of new emerging activities and harms [28]. In addition, the taxonomy focuses on harmful activities without accompanying guidelines about articulating these harms in a contextual manner. As such, an approach [28] was developed to describe the specific boundaries and characteristics of a privacy harm. The approach creates a 'limiting principle' that helps distinguish privacy from other values, such as autonomy, equality, etc. It also creates a 'rule of recognition' to help identify new emerging privacy harms [28]. By defining such principles, the approach can be used to complement the taxonomy of privacy to include and exclude harmful activities into/from its main categories. Furthermore, other frameworks focus on distinguishing the types of privacy rather than identifying and characterising privacy-impacting events and the corresponding harms [22, 18]. Comparing these frameworks with the taxonomy of privacy, it is noticeable that [26] is considered proactive, whereas [22] and [18] are considered protective [18].

In contrast, another theory [29] attempts to conceptualise the right to privacy in terms of contextual integrity: it is neither the right to limit access nor the right to control personal data; rather, it is the right to determine appropriate flow of personal data. The contextual integrity framework [29] was based on social and philosophical theories to bring the social layer into view by understanding reasonable expectations of privacy and their implications, as well as to evaluate the appropriateness of the

flow of personal data in a particular context. In particular, appropriate flow of personal data needs to conform with context-relative informational norms, which consist of four elements: contexts, data types, involved actors and transmission principles. Contextual integrity is violated when these norms are breached or disrupted [29]. However, the contextual integrity framework does not consider privacy concerns that arise as a result of the massive collection of personal data and the nature of processing operations performed on this type of data [17].

In this dissertation, the focus is on data privacy, which other authors sometimes refer to as information privacy. Before we explore the most cited definitions of data privacy, we define the concept of personal data.

Personal Data

Since personal data is at the heart of data privacy definitions, it is critical to understand how personal data is defined and how its definition has evolved over time [2], not least because it is the primary asset protected by data protection legal frameworks, standards, policies, processes and technologies. Typically, it is defined as data that, directly or, in combination with other data, facilitates the identification of a data subject [2]. Inspired by the definition stated in the EU's GDPR, in this dissertation personal data is defined as any data that is sufficiently related to an identified or identifiable individual. A subset of personal data can be considered as 'sensitive data', either culturally or by legislation, or both. The EU's GDPR [6], for example, defines sensitive data as 'special categories of data', which refer to personal data revealing racial or ethnic, political opinions, religious, philosophical beliefs, etc. The collection, processing and dissemination of such categories require careful considerations, which may vary depending on various jurisdiction and regulations.

Data Privacy

Data privacy has been defined by several scholars from different perspectives. Westin [19] defines information privacy as:

> '[t]he claim of individuals, groups or institutions to determine for themselves when, how, and to what extent information about them is communicated to others.'

In a similar manner, Kang [20] defines information privacy as:

> '[a]n individual's control over the processing — i.e., the acquisition, disclosure, and use of personal information.'

Likewise, Fried [30] defines information privacy as:

> 'Privacy is not simply an absence of information about us in the minds of others; rather it is the control we have over information about ourselves.'

Similarly, Froomkin [31] defines information privacy as "[t]he ability to control the acquisition or release of information about oneself." In addition to these definitions, the American Institute of Certified Public Accountants (AICPA) and the Canadian Institute of Chartered Accountants (CICA) [32] give a practical and operational definition of information privacy as:

> 'The rights and obligations of individuals and organisations with respect to the collection, use, retention, disclosure, and disposal of personal information.'

In order to frame data privacy for privacy engineering, we propose an operational definition of data privacy. Data privacy can be defined as the collection, processing and dissemination of personal data in a manner that prevents the occurrence of adverse privacy events and their negative impacts on data subjects.

2.2.2 Privacy and Security

It has been widely recognised that privacy is a 'fuzzy' concept [9]; consequently, it is difficult to protect. Conceptually and methodologically, it is sometimes confused with security [9]. Therefore, it is important to distinguish between security and privacy to ensure protecting the privacy of data subjects in a non-reductive manner (by which we mean privacy is not reduced to the concepts of confidentiality, secrecy, anonymity, pseudonymity, unlinkability, etc.). To avoid such a misconception, appropriate definitions of privacy need to be adopted [17], not least to understand

what to protect and by which means [9]. From an organisational perspective, security aims to protect and control data, whereas privacy aims to make appropriate decisions about the collection, processing and dissemination of personal data, which is governed by laws and regulations, by social norms, by economics, or by policies or contracts [12]. From a privacy and data protection perspective, security is considered as one of several means that ensure privacy [9] by enforcing those decisions, but not by making the decisions [12].

In practice, risk management processes are used in many areas, including security, privacy, safety, etc. Methodologically, both security and privacy risk assessments need to be integral parts of risk management processes. However, security risk assessments address the potential adverse impacts that may result from the loss of confidentiality, integrity and availability of information and information systems. Beyond these protection goals, privacy risk assessments address the potential adverse impacts that may result from the loss of anonymity, pseudonymity, unobservability, undetectability, unlinkability, intervenability and transparency. Importantly, both security and privacy risk assessments need to be fully compliant with international standards for risk-management processes and frameworks to be closely integrated and successfully used.

2.3 Privacy by Design (PbD)

In this section, we give consideration to Privacy by Design (PbD). We start by giving an overview of the origins of PbD and illustrate its philosophy and foundational principles as a proactive and creative approach to privacy. We also illustrate the main challenges of implementing the PbD approach in practice (from an engineering perspective).

2.3.1 The Origins of PbD

Ensuring that the processing of personal data is conducted fairly and lawfully is one of the main challenges in the context of privacy and data protection. This challenge has raised concerns over data-processing activities that may lead to privacy violations

or harms. These concerns have motivated the development of legal frameworks and standards with the aim of governing the processing of personal data.

In 1972, the Fair Information Practice Principles (FIPPs) were developed as core principles of the Code of Fair Information Practices (FIPs), which was recommended as a means of protecting personal data against undesirable consequences [11]. The FIPPs represent the minimum rights of data subjects in terms of the processing of personal data, which, in turn, are considered as legal safeguard requirements [11]. These principles serve as a universal framework for translating privacy and data protection goals into law, regulation, policy and technology [12]. Since then, other guidelines and principles have been developed by a variety of organisations, such as the Organisation for Economic Cooperation and Development (OECD)[3] and the organisation for Asia-Pacific Economic Cooperation (APEC)[4], to codify the FIPPs to protect the privacy of data subjects and ensure that personal data flow across borders is appropriate. Key elements of these guidelines are 'the fairness test' and 'harm prevention' respectively [33]. The former element emphasises the appropriate manner by which personal data is collected, even if this data is collected by lawful means with legitimate purposes, whereas the latter focuses on preventing the negative impacts of actual and potential risks of disclosing personal data [33]. In 2006, Global Privacy Standard (GPS) was accepted as a global instrument at the 28th International Data Protection and Privacy Commissioners Conference [34]. The GPS is considered as a single harmonised set of universal privacy principles that reflects the best of those found in many variants of FIPPs (which all share common fundamentals). It is noteworthy that the GPS, for the first time, identified data minimisation explicitly as a universal privacy principle [34].

In addition to these guidelines and principles, legislation and regulations have been enacted — whether they are generic as is typically the case in, say, Europe or sectoral-based as is typically the case in, say, the US — to regulate the processing of

[3]ibid., p. 2

[4]Asia-Pacific Economic Cooperation (APEC): APEC Privacy Framework. Retrieved from: `http://publications.apec.org/publication-detail.php?pub_id=390` [Accessed 13 March, 2018]

personal data, such as the EU's GDPR [6] and the US Health Insurance Portability and Accountability Act (HIPAA)[5].

Accordingly, software engineers are increasingly expected to develop and maintain privacy-aware systems and services that both comply with such frameworks and standards and meet reasonable expectations of privacy. To ensure and demonstrate privacy compliance, legal and regulatory requirements need to be properly interpreted as operational requirements. However, software-intensive systems have become increasingly large and complex. Thus, the adoption of the FIPPs and/or applicable legal frameworks need to be applied in an appropriate manner to ensure the fairness of data-processing activities and that they will not lead to privacy violations and harms. This implies that the adoption needs to be accompanied with a comprehensive approach and guidelines for the designers through the design process to address privacy issues and corresponding technical measures that have to be taken into account during the development process. In particular, the FIPPs need to be incorporated into the design and operation of information and communications technologies from the early stages of the process [12]. With the increasing complexity and interconnectedness of information technologies, legal frameworks and standards alone are insufficient in protecting personal data and need to be accompanied with effective technical solutions to provide adequate levels of privacy protection [8, 35, 36]. This, in turn, has led to the emergence and growth of Privacy-Enhancing Technologies (PETs) as technical measures to eliminate or mitigate privacy concerns.

Privacy-Enhancing Technologies (PETs)

PETs are defined as technical measures that prevent unnecessary or unlawful collection, use and disclosure of personal data without affecting the functionality of software systems, and provide mechanisms to enable data subjects to exercise control over their personal data [37, 38, 35]. Each technical measure is a building

[5]U.S. Government Publishing Office. Health Insurance Portability and Accountability Act of 1996. Retrieved from: `https://www.gpo.gov/fdsys/pkg/CRPT-104hrpt736/pdf/CRPT-104hrpt736.pdf` [Accessed 13 March, 2018]

block that addresses specific privacy issues and is applicable to specific contexts [39]. PETs may be thought of as complementary to legal frameworks and standards by integrating legal and standard principles into design specifications to provide various levels of privacy protection [40, 41].

Over the past decade, there have been several proposals of PETs [8], which, in turn, emphasise a need to find a suitable approach to adopt a combination of these PETs in practice [17, 42]. Simultaneously, a set of principles that support the classification of PETs [41] and privacy patterns [43] have been proposed to be used as guidelines and tools for engineers to effectively make systematic comparisons of various PETs, and then support the application of appropriate solutions, as well as to encourage the innovation of new solutions for recurring problems. However, existing classifications have mainly considered security-related issues and do not focus on specific aspects of privacy [8]. To avoid such limitations, a universal taxonomy of PETs was proposed in [8] to provide a tool for systematic comparisons of various PETs. It is a dimension-based approach of organising the taxonomy as a tree [8], with the primary dimensions being: scenario, aspect, aim, foundation, data, trusted third party and reversibility. These dimensions describe PETs from different privacy-related perspectives. For each dimension, a sub-tree is defined to describe its main properties at a detailed level — the PETs are assigned to the leafs of the tree. In principle, PETs are used to implement a certain privacy design pattern with concrete technology (as will be further explained in Section 2.4.1).

From an engineering perspective, there is a need for a proactive, integrative and creative approach to effectively adopt combinations of these technologies in the early stages of the design process to achieve various levels of privacy protection. This, in part, has led to the emergence of the concept of PbD [10], which captures the notion of embedding privacy into the design specifications of information technologies and supports the notion of 'privacy by default' by applying the FIPPs and adopting various PETs [35].

2.3.2 The Philosophy of PbD

The philosophy of PbD originated from the concept of PETs [38], which is focused on embedding privacy into the design specifications of various technologies [11]. The intent of PbD was to incorporate the FIPPs into the design, development, operation, and management of information and communication technologies and information systems [10, 35]. Initially, it focused on information technologies as its main application area. Subsequently, its scope of application expanded to cover other areas such as business practices and physical designs and infrastructures to give rise to an effective and comprehensive approach [10, 35]. In this dissertation, we will focus on information technology as the primary area of application of PbD with the aim of embedding privacy at the early stages of design of information technologies, architectures and systems, all of which are beyond the data subject's control [35]. In this context, PbD can be defined as a strategic management and engineering approach that aims to selectively and sustainably minimise potential privacy risks arising from the processing of personal data by applying administrative and technical measures [9, 44].

PbD assumes that appropriate levels of privacy protection cannot be achieved by only complying with legal frameworks and standards; rather, privacy assurance should be a 'default mode of operation' [45]. This requires: engineering innovation; clear commitment at managerial level; data subjects' participation; and social norms and controls — legislation, regulations, code of conducts, and public response [46]. In 2010, PbD was adopted as an international privacy standard at the 32nd International Conference of Data Protection and Privacy Commissioners [36]. Subsequently, PbD has played a role in legislation in various jurisdictions, such as the FTC[6] and the EU's GDPR [6].

PbD is a principled approach that consists of seven foundational principles that affirm the GPS principles. These principles are intended to serve as a comprehensive

[6]Federal Trade Commission (FTC): Privacy By Design and the New Privacy Framework of the U.S. Federal Trade Commission. Retrieved from: `https://www.ftc.gov/sites/default/files/documents/public_statements/privacy-design-and-new-privacy-framework-u.s.federal-trade-commission/120613privacydesign.pdf` [Accessed 5 February, 2018]

framework [34] for integrating privacy requirements effectively and proactively into the early stages of the design process. The principles of PbD are as follows [34, 45].

P.1 **Proactive** not Reactive; **Preventative** not Remedial
P.2 Privacy as the **Default Setting**
P.3 Privacy **Embedded** into Design
P.4 Full Functionality - **Positive-Sum**, not Zero-Sum
P.5 End-to-End Security - **Full Lifecycle Protection**
P.6 **Visibility** and **Transparency** - Keep it **Open**
P.7 **Respect** for User Privacy - Keep it **User-Centric**

Translating PbD's principles into engineering activities (a set of specific, and operationally feasible methods, techniques and tools) involves a number of challenges in practice. Such challenges will be discussed in the following sections.

2.3.3 A Critique of PbD

Incorporating privacy into the early stages of the design process requires an appropriate translation of legal, social and ethical concerns into system requirements [14], which, in turn, presents a difficulty with respect to understanding the actual meaning of privacy and its related issues. PbD has been advocated to address these concerns and achieve the highest standard of privacy protection. However, its implementation has several challenges, including a lack of step-by-step methodologies that: address the complexity and variability of privacy issues; incorporate privacy into the SDLC [9]; and support the translation of its principles into engineering activities [14]. In some ways this is understandable, as PbD was developed to take into account a range of sources and standards. However, a consequence is that its principles are given at a high level of abstraction [47] with vague definitions [14]; consequently, it is not clear how its principles can be appropriately translated into engineering activities — meaning that there is a reliance on software engineers' expertise with regards to translating abstract privacy principles stated in legal frameworks and standards into operational requirements. In addition, some of its principles are defined in a recursive way [14]. For example, the principle "Privacy **Embedded** into Design" is defined thus [34]:

> "**Privacy by Design** is embedded into **the design** and architecture of IT systems and business practices. It is not bolted on as an add-on, after the fact [...]."

According to this recursive definition, it is not clear how PbD is to be embedded into the design and architecture of systems [14]. Such vagueness leads to a disconnect between policy-makers and software engineers in terms of the actual meaning of PbD and the way in which technical solutions can be developed to comply with legal frameworks and standards [14]. As such, the translation of its principles into engineering activities requires a specific type of engineering expertise, contextual analysis, and a balance between security and privacy interests [14].

2.4 A Critique of Existing Approaches to PbD

In this section, we investigate various methods for PbD that capture and address privacy issues in the early stages of the design process. We also investigate how these methods might be aided by techniques or tools, which provide procedures possibly with a prescribed language, notation or means to accomplish privacy-engineering tasks and activities.

In the privacy research literature, it has been observed that there are broadly two distinct areas of research that work towards different goals [27]. The first area of research aims to develop provable cryptographic mechanisms for assuring (or contributing to assuring) privacy protection. Researchers in this area adopt a threat model that assumes sophisticated threat actors who are not deterred by legal frameworks and standards. The second area of research aims to protect personal data from accidental disclosure and misuse, and facilitate notice and choice options. Researchers in this area emphasise the importance of implementing the FIPPs with the assumption that technology enforces legal frameworks and standards, rather than guaranteeing privacy [27].

With reference to these areas of research, existing approaches to PbD can be initially classified into: goal-oriented, risk-based and hybrid approaches. Goal-oriented approaches describe abstract privacy principles as high-level goals that

need to be achieved. Each goal can be decomposed into low-level guidelines, which, in turn, can be described as sets of operational requirements. Risk-based approaches identify the assets to protect and the threats that may compromise the relevant abstract privacy principles. Based on the threat analysis and risk assessment, privacy requirements are elicited from risk treatments. The approaches are complementary: they both aim to support better understanding of how to comply with abstract privacy principles by means of a set of explicit requirements [16]. They differ in the manner by which they address privacy issues: risk-based approaches focus on identifying privacy risks that need to be addressed, whereas goal-oriented approaches focus on identifying privacy goals that need to be achieved [16]. We consider each in turn with the aim of discussing their limitations and establishing a set of criteria for developing and assessing engineering approaches to PbD. These criteria will be collated in Section 2.5.2.

2.4.1 Goal-Oriented Approaches

Privacy Design Strategies

A set of privacy design strategies are proposed in [48] to help engineers facilitate PbD in the early stages of the design process. These strategies — which describe fundamental approaches to achieve specific design goals (such as certain levels of privacy protection) — aim to provide a useful classification of privacy design patterns and their underlying PETs. They also aim to support engineers in fulfilling privacy requirements with practical tools that enable them to make well-founded choices. The strategies were derived from legal frameworks and standards, and build upon the framework of [27] to engineer privacy from two perspectives: 'privacy-by-architecture' and 'privacy-by-policy'.

A critical analysis of these strategies is given in [3], with a view to bridging the gap between the legal and engineering domains by exploring the relationships between the concepts introduced and similar concepts used in software engineering. As a result, privacy design strategies are concretely refined as engineering approaches to PbD that correspond to architectural goals. A design strategy specifies a specific

architectural goal to achieve a certain level of privacy protection [3]. In addition, an additional level of abstraction between strategies and design patterns is proposed: *tactics*. A tactic is an approach to PbD that contributes to the specific goal of an overarching design strategy [3] — similar to software architectural tactics. A set of tactics is proposed in [3] to facilitate privacy protection goals expressed as strategies.

Critique. Although the strategies of [48] and [3] bridge the gap between legal and technical domains, they jump directly from abstract privacy principles into software architectures without defining and/or refining a set of explicit requirements [49] or referring to a set of comprehensive privacy protection goals that help derive requirements and choose techniques (design patterns) and technologies (PETs) that implement these requirements [4]. Semantically, privacy principles are expressed as abstract concepts that are different from technical ones [16]; therefore, they need to be translated into concrete and auditable protection goals [15] to aid in determining appropriate privacy design strategies. Furthermore, these strategies are not enough on their own, as they fall short of relating how they can be applied when building privacy-preserving systems [50]. Although the refined definitions distinguish between strategies, tactics and how they relate to PETs through design patterns, a catalogue of design strategies and tactics is not sufficient to provide insight into the process through which these strategies can be applied in practice [50]. This gives rise to criteria **Cr.12** and **Cr.13** of Section 2.5.2.

Privacy-Enhancing ARchitectures (PEARs)

A methodology that focuses on privacy-enhancing architectures was proposed in [51] to design, analyse and evaluate software architectures through quality attributes, architectural tactics and patterns. Despite the functional requirements of privacy that focus on, for example, purpose specification, the approach focuses only on non-functional requirements, as it considers privacy as a quality attribute. It was based on the software architecture methodology from Carnegie Mellon University [52]. Typically, quality attributes are specified by modelling the most important quality scenarios. The elements of a scenario are: *the stimulus; the source of a stimulus; the*

environment where the system is deployed; the artefact being stimulated; the response as the result of the arrival of the stimulus; and *the response measure* [51]. The proposed architecture design process can be described through three main steps [51]. First, the most important quality scenarios need to be specified. Next, appropriate architectural tactics that control the response element of quality scenarios need to be selected. Finally the impact of these tactics on response measures is verified. In [51], four categories of tactics for PbD are proposed to achieve various levels of privacy protection. To facilitate tactics re-use, tactics are described through architectural patterns and implemented by PETs.

Critique. Although the PEARs methodology uses quality scenarios as structured means to capture quality attribute requirements, it does not refer to a set of comprehensive privacy protection goals around which the proposed architectural tactics can be organised (as previously mentioned above). The environment element of a quality scenario focuses on the technical environment in which a system operates, but other contextual factors are not explicitly considered by the methodology. In addition, the PEARs methodology does not explicitly illustrate how PETs can be mapped onto architectural patterns. This gives rise to criteria **Cr.9**, **Cr.12** and **Cr.13** of Section 2.5.2.

Privacy Design Patterns

Design patterns are considered as an important method for refining the system architecture to achieve certain functional requirements within a given set of constraints [48]. They are a useful means for making design decisions about the architecture of a software system. In the context of privacy and data protection, there has not been a high volume of work on privacy design patterns compared to the work on PETs [53, 43, 48]. Nonetheless, a number of privacy patterns have been proposed at varying levels of abstraction to be applicable to the design of anonymity solutions in various problem domains [53, 43]. Privacy patterns can be classified into three categories [43]: the first category concerns patterns that document general privacy concepts, such as [54, 55]; the second category concerns patterns that ensure

compliance with privacy policies, such as [56, 57]; and the third category concerns patterns that provide sufficient guideline for PET development, such as [43].

In addition, a privacy pattern catalogue has been proposed in [53] to aid the decision-making process for the designers of privacy-preserving systems. Each pattern is described in a manner that illustrates its: *intent, context, problem, forces, solution, design issues, consequences, known uses* and *related patterns*. Each pattern also takes into consideration a specific threat model and a required degree of anonymity, which is scenario-specific. Similarly, a set of privacy patterns are proposed in [56] for web-based activities.

Furthermore, a pattern language has been proposed in [43] for developing PETs. Such a language can be used to narrate a solution that enables anonymity in PETs in a manner that illustrates how a pattern can be applied in various domains. It guides PET developers towards designing a solution for a recurring privacy problem, or innovating a new solution for a new privacy problem. According to this language, PETs are designed with a desired degree of anonymity and a threat model. As such, each pattern takes into account what is being protected (user data or user communication patterns) and who it is being protected from (insider, outsider or partner — based on their relationships with data subjects, whether they work on their own, or may be colluding).

In addition, a decision-based support system that considers contextual factors that influence the degree to which privacy is required for a given context was proposed in [58]. Based on a set of context-specific selection criteria, the system provides a ranked list of candidate patterns that can be used in conjunction to meet privacy requirements [58].

Critique. Although patterns help move from legal and normative requirements to technical designs, they are not enough on their own, as they fall short of refining explicit requirements [49]. Furthermore, patterns are not described uniformly by referring to a standard template, format or classification schema. A consequence is that patterns are presented in a manner that does not support their use [59]. In addition, the decision-based support system of [58] recommends a list of relevant

patterns based on a set of context-specific selection rules. However, each pattern needs to be carefully formulated in a fine-grained manner by adding a criteria field in which a list of parameters relating to context and selection criteria can be defined. Furthermore, the system does not consider how privacy requirements can be mapped onto software architectures through the adoption of PETs as technical measures. This gives rise to criteria **Cr.2**, **Cr.12** and **Cr.13** of Section 2.5.2.

The PriS Method

The PriS method [60] aims to integrate privacy requirements into the early stages of the design process by modelling privacy requirements as organisational goals. It considers eight privacy protection goals: identification, authentication, authorisation, data protection, anonymity, pseudonymity, unlinkability and unobservability. In addition, the method considers stakeholders' expectations and concerns during the elicitation of privacy-related goals in relation to the system's environment. Each of the protection goals has relevant stakeholders who may have different conflicts of interest; therefore, conflict resolution techniques may be utilised [60].

Once privacy-related goals are identified, their potential impact can be analysed. This may lead to the identification of new goals, which, in turn, may lead to new processes or improve existing goals. Next, these processes are modelled using relevant privacy-process patterns. The PriS method adopts goal models to address privacy concerns in each process. In addition, the PriS method supports the mapping of privacy requirements onto appropriate software architectures by providing privacy-process patterns. Each pattern illustrates privacy activities that need to be implemented, which, in turn, aids in deciding where privacy controls (manifested by, for example, PETs) need to be implemented to achieve an acceptable level of privacy protection.

Critique. In spite of the fact that the PriS method emphasises the complexity and variability of privacy as a legal and social concept, it does not refer to specific privacy definitions, principles and comprehensive privacy protection goals. In addition, it does not explicitly illustrate how to identify reasonable expectations

of privacy in a contextual manner. This gives rise to criteria **Cr.2**, **Cr.9** and **Cr.14** of Section 2.5.2.

2.4.2 Risk-Based Approaches

The most frequently used privacy risk models are compliance-based models [12, 61]. Given that the FIPPs establish a basis for most legal frameworks and standards, these principles, and all of their variants, serve as the most common privacy risk models [12]. Such models provide a range of options for developing privacy-compliant systems [61]. Typically, legal frameworks and standards prescribe and proscribe certain aspects related to the structure and behaviour of a system that processes personal data in terms of specifying the purposes for which personal data is processed, the types of collected personal data and the manner in which the data is processed, shared and protected [61]. Often, engineers interpret legal requirements or standard principles and employ these interpretations as a risk model. The identification of privacy risks entails examining the components of the system that implements each specific legal requirement or principle and how they may give rise to additional vulnerabilities [61].

With the advent of ICTs, software-based systems have become increasingly large and complex. This complexity presents challenges for compliance-based models [12]. In the privacy research literature, there are several contributions [26, 28, 29] that propose a basis for developing robust privacy risks models that help facilitate the identification of potential risks that the compliance-based models may fail to detect [12]. However, these attempts do not present complete risk models; rather, they provide 'grounded reference points' around which complete user-centric privacy risk models can be developed [12].

Privacy Impact Assessments (PIAs)

In order to anticipate and prevent the processing operations that may lead to privacy violations or harms, the adverse impacts of these operations need to be proactively assessed in the early stages of the design process [62]. In some jurisdictions, 'legal

compliance checks' [15] or 'prior checking' [63] are the most commonly used privacy assessment procedures. These procedures are often not conducted by engineers; rather, auditors, lawyers or data protection authorities utilise check-lists to check compliance with legal frameworks [15]. With the advent of ICTs, holistic and effective impact assessments are considered as complements to, or replacements for, these assessment procedures. This has contributed to the emergence of the concept of a *Privacy Impact Assessment (PIA)*.

A PIA is defined as a process that identifies and mitigates the impact of an initiative on privacy with stakeholders' consultation [64]. It is more than a tool: it is an ongoing process that begins at the earliest possible stage [65]. The use of PIAs is now common practice in a variety of jurisdictions, not least because they play a crucial role in providing privacy protection for data subjects and in supporting risk management for organisations. As such, PIAs are now mandated by, for example, the EU GDPR [6]. 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 a manifestation of PbD [66].

Many guidance documents have been published to help support organisations in performing PIAs, yet these documents vary noticeably in their comprehensiveness and quality. A systematic methodology for PIAs [15] has been proposed to identify and assess privacy impacts, as well as to achieve the principles of PbD. Existing PIA processes strive to achieve the aim of PbD by applying its foundational principles [15].

Critique. In spite of the fact that PIAs have been mandated in some jurisdictions, there is a lack of standards that illustrate how these PIAs can be conducted systematically [15]. Guidance documents do not precisely illustrate how a risk assessment of a PIA should be performed [62] as a PIA tends to focus more on legal and organisational aspects [62, 39]. 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 [15]. Although PIAs are expected to follow the same philosophy, existing PIA processes largely fall short in this respect [15, 62]. They cannot be applied easily, not least because they are imprecise, lengthy or

improperly structured [15]. Existing PIA guidance documents do not typically support the integration of a PIA into a risk-management process [65]. 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. This gives rise to criteria **Cr.7**, **Cr.8** and **Cr.11** of Section 2.5.2.

Some PIA processes, such as the BSI IT-Grundschutz [67], 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 privacy protection goals. This, in turn, reduces privacy protection to the concepts of anonymity, pseudonymity, unobservability, undetectability and unlinkability [15, 68]. Thus, targets of evaluation — e.g. 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. Abstract privacy principles stated in legal frameworks and standards are different from concrete data-processing activities; therefore, it is difficult to use them for assessing these activities and describing design decisions at an architectural level [15]. Accordingly, abstract privacy principles need to be translated into concrete and auditable protection goals to aid engineers in specifying design strategies. This gives rise to criterion **Cr.12** of Section 2.5.2.

The Methodology for Privacy Risk Management (CNIL)

The CNIL methodology [5] presents an analytical approach for managing the risks that may arise from the collection, processing and dissemination of personal data. It classifies assets into two categories: *primary assets*, which represent personal data, processing operations and legal processes required by legal frameworks, and *supporting assets* on which the primary assets rely, such as hardware, software, people, etc.

Risk model. The CNIL methodology defines a risk model that illustrates four assessable risk factors: feared events, threats, vulnerabilities and risk sources. Importantly, a feared event describes both the adverse event and its potential impacts on data subjects. The CNIL methodology provides five types of feared events that affect the processing operations according to the types of primary

assets. For a feared event to occur, one or more risk sources exploit, accidentally or deliberately, one or more vulnerabilities of supporting assets through different threats. Based on the risk model, a privacy risk is composed of one feared event and all the threats that make it possible.

Assessment approach. The risk level is assessed in terms of severity and likelihood, with levels of risk being based on two key risk factors: feared events are used to assess the severity, which depends on the level of identification of personal data and the prejudicial effect of the potential impacts; threats are used to assess the likelihood, which depends on the level of vulnerabilities of the supporting assets and the level of capabilities of the risk sources. The CNIL methodology uses a semi-quantitative approach that uses a fixed scale of levels (negligible, limited, significant, maximum), along with corresponding numbers. The risk levels are located on a risk map with severity and likelihood on its axes, with the aim of ordering and prioritising these risks. Further, it provides a catalogue of risk-treatment measures that can be adopted to address the identified and assessed risks. The PIA for smart grid and smart metering systems [69] is an example of a PIA that adopts the CNIL methodology to identify, analyse and assess potential privacy risks.

Critique. In spite of the fact that the CNIL methodology presents a risk-management approach that appropriately facilitates the adoption of risk-treatment measures in a commensurate manner, it does not describe the key risk factors in well-defined attributes to help assess their levels. In addition, it does not describe the established assessment rules in a way that reflects the assessable attributes of the key risk factors to facilitate their roles in risk assessments and their translation into qualitative terms for multiple stakeholders. This gives rise to criteria **Cr.7**, **Cr.8** and **Cr.11** of Section 2.5.2.

The Privacy Risk Analysis Methodology (PRIAM)

The Privacy Risk Analysis Methodology (PRIAM) [62] provides step-by-step guidance for conducting privacy risk assessments in a systematic manner. PRIAM defines and characterises the key factors that have impacts on privacy risks and

considers their contribution to the assessment of the overall risks [62]. It consists of two main phases: the information-gathering phase and the risk-assessment phase. The first phase aims to collect all useful and relevant information that helps support the determination of the values of the attributes, whereas the second phase aims to use these values to assess the risk levels. In addition, PRIAM adopts the concept of privacy harm, which has historically been given less attention by engineers during risk assessments. To support the risk analysis, PRIAM presents the concept of the harm tree as an analysis technique, which is used to describe the many-to-many relationships among the key risk factors for each privacy harm.

Based on this methodology, a refinement approach for the reuse of privacy risk analysis results [39] has been developed to improve the cost-effectiveness of PIAs. The approach provides guidance to software designers to select an architecture and justify the choice with respect to the results of a privacy risk analysis.

Risk model. PRIAM concretely defines a risk model that defines key risk factors with well-defined attributes: privacy harms, feared events, privacy weaknesses and risk sources. It also illustrates the relationships among these factors and describes the dependencies between their attributes. A privacy harm results from one or more feared events. Each feared event results from the exploitation of one or more privacy weaknesses by one or more risk sources [62].

Assessment approach. The risk level is assessed in terms of severity and likelihood for each privacy harm. PRIAM estimates the severity of a privacy harm based on its intensity and victims, which are influenced by the 'irreversibility' and 'scale' attributes of associated feared event respectively. The likelihood of a privacy harm is computed from the likelihood of its corresponding feared events derived from the likelihood of successful exploitation of associated privacy weaknesses, which depends on the capabilities of risk sources and the exploitability of privacy weaknesses. PRIAM uses a qualitative assessment approach involving various scales for assessing the severity of privacy harms; it also uses a semi-quantitative assessment approach that adopts a set of rules for assessing their likelihood.

Critique. In spite of the fact that PRIAM presents a risk model that defines the key risk factors and a harm tree that represents the conceptual relationships among these factors, it does not refer to an appropriate classification of feared events that characterises these events according to the nature of processing operations. In addition, it does not consider a model that helps support the identification of reasonable expectations of privacy. This gives rise to criteria **Cr.9** and **Cr.10** of Section 2.5.2.

Although the refined approach provides a systematic way for conducting a generic, architecture-based and context-based risk analysis, privacy is not embedded into the design process of software architectures — i.e. it is considered after the fact by analysing and evaluating a range of pre-defined architectural options from which a suitable architecture is selected. Embedding privacy sometimes necessitates re-inventing existing architectural choices when alternatives are not suitable for preventing the identified adverse privacy events. In particular, privacy requirements need to be considered at the early stages where architectural decisions around the collection, processing and dissemination of personal data can still be made. Further, the refined approach, in its current state, is not accompanied with guidelines that illustrate how to adopt existing privacy patterns and their underlying PETs as technical measures. This gives rise to criterion **Cr.14** of Section 2.5.2. In addition, the refined approach considers contextual factors in terms of risk sources motivation (incentives and disincentives) and resources (background information and technical resources). External and internal contexts need to be considered in a broader way to include the factors influencing design decisions (i.e. legal and regulatory factors, social factors and contractual factors). These factors help capture social and legal aspects, including stakeholders' perceptions and expectations in a contextual manner. This gives rise to criterion **Cr.9** of Section 2.5.2.

LINDDUN

LINDDUN [70] is a methodology to elicit the privacy requirements of systems and select appropriate PETs according to these requirements. It defines a set of privacy

protection goals: unlinkability, anonymity and pseudonymity, undetectability and unobservability, plausible deniability, confidentiality, content awareness, and policy and consent compliance. LINDDUN provides a set of threat types in relation to the elements of Data Flow Diagrams (DFDs). It also provides a catalogue of privacy-specific threat tree patterns to support the identification of potential privacy risks, the documentation of misuse cases, and the elicitation of requirements. The catalogue provides a means to map existing PETs to the main types of privacy requirements.

Critique. In spite of the fact that LINDDUN provides a privacy threat analysis framework that supports the elicitation and fulfilment of privacy requirements, the provision of a catalogue as a means for mapping PETs to the corresponding types of privacy requirements without providing a set of selection criteria is not sufficient in terms of reasoning critically about architectural decisions and choosing among architectural alternatives, especially as new PETs are continually emerging. This gives rise to criterion **Cr.14** of Section 2.5.2.

2.4.3 Hybrid Approaches

Over the past decade, there have been two prominent methodologies that emphasise the importance of engineering systems according to the principle of data minimisation. We discuss both below.

The Framework for Privacy-Friendly Systems Design (PFSD)

The Framework for Privacy-Friendly Systems Design (PFSD) [27] has been developed by Spiekermann and Cranor to provide a holistic view of privacy engineering. It is built upon a three-layer model that translates frequently cited privacy definitions into high-level engineering responsibilities in relation to three technical domains: the user sphere, the recipient sphere and the joint sphere. The identified responsibilities are concerned with ensuring that users can exercise control over their personal data and that engineers are responsible for managing potential privacy risks where personal data is not under their control.

The model also serves as a basis for privacy requirements analysis with the aim of identifying the activities of system operations (transfer, storage and processing) that may raise privacy concerns by compromising data subjects' expectations of privacy. These operations may have impacts on privacy based on several factors: types of personal data, involved actors, the way these operations are performed, and the domain in which personal data is processed. It is also important to understand reasonable expectations of privacy and the extent to which PETs are needed to address privacy concerns and meet legal framework and standards [27]. The adoption of appropriate PETs depends on the threat models around which they are designed. The PFSD framework considers only seven areas of activities that raise privacy concerns identified as a result of the empirical study of [71] in relation to the three system operations and technical domains.

The PFSD framework also provides guidelines for designing privacy-friendly systems according to two main approaches: privacy-by-architecture and privacy-by-policy. The first approach focuses on identifying architectural choices that specify various levels of privacy protection by minimising data collection, and emphasising anonymisation and client-side data storage and processing [27]. The second approach focuses on providing data subjects with essential mechanisms to exercise control over their personal data. In particular, it implements the notice, choice and access principles by providing information about: how and what types of personal data is collected; how they will be processed, shared and protected; and how data subjects will access their personal data and provide and/or withdraw their consent.

Critique. In spite of the fact that the PFSD framework emphasises the importance of understanding privacy concerns and expectations, it is not accompanied by techniques that explicitly illustrate how these can be identified in a contextual and structured manner. It is understandable that the PFSD framework is likely to be domain-specific for e-commerce and ubiquitous computing. However, identifying a set of concerns is not sufficient in considering the variability of privacy, especially as perceptions and expectations of privacy are influenced by legal, social and economic

changes, as well as by technological developments which evolve over time. This gives rise to criteria **Cr.9** and **Cr.10** of Section 2.5.2.

In the context of privacy, adverse events that may lead to privacy violations and harms depend mainly on the manner by which personal data is collected, processed and disseminated. Instead of focusing only on three operations, an abstract model that defines and distinguishes between the main types of operations performed on personal data needs to be considered to represent data-processing activities in a way that is amenable to analysis. This gives rise to criterion **Cr.1** of Section 2.5.2.

Further, the PFSD framework does not explicitly define a privacy risk model that considers the dynamic nature of privacy issues to support the adoption of appropriate PETs and reasoning about architectural choices made. This gives rise to criteria **Cr.7**, **Cr.8** and **Cr.11** of Section 2.5.2.

The privacy-by-architecture approach relies mainly upon making architectural choices that limit the amount of collected personal data to the absolute minimum necessary and enforce anonymous and decentralised storage and processing. However, the PFSD does not explicitly provide guidelines that illustrate how these can be achieved; in particular, ensuring data minimisation is itself a challenge. This gives rise to criterion **Cr.13** of Section 2.5.2.

Although the privacy-by-policy approach provides some degree of control over personal data, it implements only the notice, choice and access principles, which are a subset of the FIPPs tailored to the e-commerce context by the FTC[7]. A set of universal privacy principles that reflects the best of those found in many variants of FIPPs needs to be adopted by the privacy-by-policy approaches to be applied in various contexts, especially with the absence of global privacy and data protection laws and regulations. This gives rise to criterion **Cr.2** of Section 2.5.2. The privacy-by-policy approach considers the implementation of technical mechanisms that can aid in auditing and enforcing compliance with the notice, choice and access principles of the FIPPs. Without providing necessary technical assurances, it is difficult to systematically evaluate whether the system

[7]ibid., p. 20

developed using this approach complies with legal frameworks and standards. This gives rise to criterion **Cr.6** of Section 2.5.2.

A Straw-Man Methodology for Engineering PbD

An incomplete methodology for engineering PbD has been presented by Gürses *et al.* [14] to develop systems according to the principles of PbD. It uses the principle of data minimisation as a necessary and foundational step for privacy engineering activities to mitigate potential privacy risks, avoid function-creep, provide data subjects with control over their personal data, and achieve privacy compliance with legal frameworks and standards. Crucially, the application of data minimisation does not imply anonymity; rather, it can be achieved by leveraging mathematical and computational capabilities of technologies with the aim of minimising different types of personal data [14] in a contextual manner (according to the purposes for which personal data is collected).

Engineering systems according to the principles of PbD requires integrating privacy requirements into the typical systems engineering activities [14]. The methodology consists of six preliminary activities: functional requirements analysis; data minimisation; modelling attackers, threats and risks; multilateral security requirements analysis; and implementation and testing of the design [14].

Critique. In spite of the fact that data minimisation is a critical step for engineering PbD, ensuring data minimisation is itself a challenge. The methodology emphasises the importance of clearly describing the desired functionality before any of the activities are conducted to specify personal data that is absolutely necessary to fulfil the functionality. However, it does not explicitly illustrate how the functionality can be refined into concrete data-processing activities and represented in a way that is amenable for analysis to ensure that the collection and processing of personal data is restricted to the minimum amount necessary for each data-processing activity. This gives rise to criteria **Cr.3**, **Cr.4** and **Cr.5** of Section 2.5.2. In addition, the methodology is not accompanied with techniques or strategies that illustrate how

to identify appropriate types of data minimisation and adopt appropriate PETs to implement these types. This gives rise to criterion **Cr.13** of Section 2.5.2.

Although the methodology emphasises the importance of developing models of potential threat actors and the types of threats those actors may realise, the methodology does not explicitly define a privacy risk model that supports the adoption of appropriate PETs and reasoning about architectural choices made. This gives rise to criteria **Cr.7**, **Cr.8** and **Cr.11** of Section 2.5.2.

Data Minimisation Strategies

The principle of data minimisation conceals a number of design strategies that engineers can apply intuitively when engineering systems according to the principles of PbD [50]. Such strategies can be considered as constraints on the flow of personal data from the user sphere (user domain) to the recipient sphere (service domain). Gürses *et al.* [50] identified six data minimisation strategies that can be used to minimise potential privacy risks by avoiding a single point of failure and minimise the need to trust data holders by retaining data under the data subject's control.

Gürses *et al.* [50] also provided a preliminary description of four activities that engineers may conduct to apply the identified data minimisation strategies. The first activity concerns the classification of the entities of a given system into two domains (the user domain and service domain). The second activity concerns the identification of data necessary at the service domain that sufficiently fulfils the system's functionality. The minimum amount necessary depends on the context in which, and the purposes for which, personal data is collected and processed, as well as technological capabilities. The third activity concerns the distribution of personal data in the system architecture to fulfil its functionality. It aims to map the specified personal data onto the entities in the two domains. The fourth activity concerns the adoption of PETs to implement the strategies.

Critique. In spite of the fact that the main activities are preliminarily described, the specification of personal data necessary for fulfilling the system functionality

is described at a high level of abstraction using three approaches: 'collect-all-data', 'select-before-collect' and 'only-collect-necessary-data'. It is analytically useful to specify the minimum amount necessary for each data-processing activity rather than for the functionality of the system. This gives rise to criteria **Cr.3**, **Cr.4** and **Cr.5** of Section 2.5.2.

The fourth activity provides a non-exhaustive list of fundamental approaches that help engineers to keep personal data under the user control. However, it is not explicitly illustrated how the identified data minimisation strategies can be mapped onto privacy patterns. In addition, it is not accompanied with a technique that helps engineers to prioritise and select appropriate design strategies, along with technological alternatives that implement them, to reason critically about the architectural choices made. This gives rise to criterion **Cr.13** of Section 2.5.2.

2.5 Beyond the Critique

2.5.1 The Challenges of Engineering PbD

In the preceding sections, we reviewed existing approaches to PbD. In this subsection, a summary of the key points from this review are articulated as a list of challenges associated with engineering PbD.

Ch.1 The plurality and contextuality of privacy present complexity and variability of privacy issues. As a consequence, the dynamic nature of privacy challenges engineers to understand and translate abstract privacy principles into operational requirements. It also challenges engineers to understand and consider stakeholders' expectations and concerns, which, in turn, require specific expertise, contextual analysis and resolution of stakeholders' conflicts.

Ch.2 There is a lack of means to ensure and demonstrate privacy compliance. Complying with legal frameworks and standards requires abstract models that represent data-processing activities in a way that is amenable to compliance checking. Such a representation needs to specify: types and sensitivity of personal data; the purposes for, and the manner in which, this data is collected,

processed and disseminated; involved actors; assigned roles and responsibilities; and legal and domain-specific constraints.

Ch.3 There is a lack of systematic methods that help facilitate data-centric threat modelling in a meaningful manner. This challenges engineers to holistically identify and systematically analyse potential privacy risks for deriving privacy requirements. Further, privacy risk assessment needs to go beyond identifying technical risks; however, this requires an understanding of social perceptions and data subjects' expectations that are derived from social norms.

Ch.4 There is a lack of techniques that help determine appropriate degrees to which privacy is required. Importantly, appropriate levels of privacy protection need to be determined in a contextual manner according to several factors: legitimate objectives, appropriate types of data minimisation, stakeholders' expectations, legal frameworks and standards, technological capabilities, and appropriate threat models.

Ch.5 There is a lack of methods that characterise privacy protection as a general quality attribute scenario indicating the range of values its main elements can take to generate concrete quality scenarios. A collection of concrete quality scenarios can be used to derive the quality attribute requirements for a given system. Further, these scenarios can be used to specify, implement and justify appropriate levels of privacy protection.

2.5.2 Criteria for Assessing Approaches to PbD

In the following, we present a set of criteria for developing and evaluating engineering approaches to PbD derived from the analysis of Section 2.4.

Cr.1 An abstract model that distinguishes between the main types of operations performed on personal data (from collection to destruction) needs to be developed.

Cr.2 A single harmonised set of universal privacy principles that reflects the best of those found in many variants of FIPPs needs to be adopted.

Cr.3 The abstract purposes for which personal data is collected, processed and disseminated need to be operationalised by refining these purposes into concrete data-processing activities.

Cr.4 Each concrete data-processing activity needs to be expressed in terms of actions and events that trigger the execution of these actions, and roles that define a set of responsibilities performed by different actors according to their capabilities.

Cr.5 The minimum necessary amount of personal data needs to be derived from the actions of the concrete data-processing activities.

Cr.6 Necessary technical assurances that facilitate reasoning about privacy compliance with legal frameworks and standards need to be considered by any engineering approach to PbD.

Cr.7 A risk model, which defines the key risk factors that have impacts on the privacy of data subjects and establishes a conceptual relationship among these factors, needs to be developed at an appropriate level of detail.

Cr.8 A risk analysis approach, which describes how combinations of risk factors are identified and analysed, needs to be developed to ensure adequate coverage of the problem space at a sufficient level of detail.

Cr.9 The reasonable expectations of privacy need to be identified in a contextual manner in terms of appropriate collection, processing and dissemination of personal data that reflect social norms.

Cr.10 Threats to privacy (adverse privacy events) need to be identified in a structured, concrete and comprehensive manner according to the main types of operations performed on personal data.

Categories	Existing Approaches	Criteria
Goal-Oriented Approaches	Privacy Design Strategies	Cr.12 and Cr.13
	The PEARs Methodology	Cr.9, Cr.12 and Cr.13
	Privacy Design Patterns	Cr.2, Cr.12 and Cr.13
	The PriS Method	Cr.2, Cr.9 and Cr.14
Risk-Based Approaches	PIAs	Cr.7, Cr.8, Cr.11 and Cr.12
	The CNIL Methodology	Cr.7, Cr.8 and Cr.11
	PRIAM	Cr.9, Cr.10 and Cr.14
	LINDDUN	Cr.14
Hybrid Approaches	The PFSD Framework	Cr.1, Cr.2, Cr.6, Cr.7, Cr.8, Cr.9, Cr.10, Cr.11 and Cr.13
	The Straw-Man Methodology	Cr.3, Cr.4, Cr.5, Cr.7, Cr.8, Cr.11 and Cr.13
	Data Minimisation Strategies	Cr.3, Cr.4, Cr.5 and Cr.13

Table 2.1: The mapping of existing approaches onto the established criteria.

Cr.11 A risk assessment approach, which establishes a set of assessment rules that reflect the assessable attributes of the key risk factors and specify the range of values the risk factors can assume, needs to be developed.

Cr.12 A set of comprehensive and agreed-upon privacy protection goals that consider legal, technical and societal aspects of privacy and data protection needs to be adopted.

Cr.13 A set of selection criteria needs to be established to aid engineers to map appropriate design strategies onto architectural tactics, design patterns and their underlying PETs.

Cr.14 A set of selection criteria needs to be established to aid engineers to choose appropriate architectural tactics, design patterns or PETs that fulfil the elicited requirements.

Table 2.1 matches the main categories along with the existing approaches and the established criteria.

2.6 Summary

In this chapter, we have discussed key concepts and definitions of privacy to illustrate its complexity and variability. Then, we gave a detailed background to PbD and illustrated the motivation for the contributions of this dissertation. We have also given a relatively detailed discussion of the main challenges of translating

the principles of PbD into engineering activities. Next, we provided a detailed examination of how existing approaches to PbD address these challenges and gave a detailed explanation of the limitations of these approaches. Based on these, we have given a summary of key points from the review of existing approaches as a list of challenges for engineering PbD with a view to incorporating the PbD approach into software engineering practices. Finally, we have established a set of criteria against which engineering approaches to PbD can be evaluated.

In the next chapter, we will present a research methodology that: is suitable to the research problem; allows the flexibility necessary for investigation in this area; and can be used for assessing the contributions of this dissertation.

Purely technical approaches might prove insufficient for aligning nuanced legal policies with engineering artifacts.

—Seda Gürses and Jose M. del Alamo

3 Methodology

Contents

3.1 Introduction

In this chapter, we investigate how the legal, socio-technical and software engineering research communities address the concepts of data privacy, privacy by design and privacy engineering as research topics. Informed by these various approaches, we present the research methodology used for conducting the research described in this dissertation and assessing the contributions of this dissertation. Section 3.2 gives a relatively brief discussion of the research methodologies that have been

used in these research areas. Based on these, Section 3.3 describes the research approach used in this dissertation. Then, Section 3.4 introduces two case studies that have various privacy issues to support detailed exploratory investigations to understand and explain how to proactively embed privacy 'by design' using primarily qualitative analysis. The first case study will be used for illustrative purposes, whereas the second case study will be used for evaluation purposes. Finally, Section 3.5 gives a brief summary of the chapter.

3.2 Research Approaches

3.2.1 Legal and Socio-Technical Approaches

Legal and socio-technical research communities focus on non-technical considerations of data privacy. In particular, research methodologies in this area tend to concentrate on the identification and articulation of privacy violations, privacy harms (also known as privacy impacts), harmful activities and reasonable expectations of privacy. From a methodological perspective, legal and social-technical research communities tend to focus on two broad methods of reasoning: deductive and inductive research approaches.

The inductive approach moves from specific observations to broader generalisations and theories [72]. Informally, it is a 'bottom-up' approach as it begins with specific observations and measures, begins to detect patterns and regularities, formulates some tentative hypotheses that we can explore, and finally ends up developing some general conclusions or theories [72]. For example, both Solove's taxonomy [26] and Nissenbaum's contextual integrity [29] used the inductive approach to capture harmful activities and reasonable expectations of privacy respectively [12]. Solove analysed various privacy harms that have achieved a significant degree of social recognition from laws, cases, constitutions, guidelines and other sources, and classified harmful activities into 16 categories using a bottom-up approach and within particular contexts rather than in the abstract [26].

In contrast, the deductive approach moves from the more general to the more specific [72]. Informally, it is a 'top-down' approach as researchers might begin

with a theory about their topic of interest. They then narrow that down into more specific hypotheses that they can test. They narrow down even further when they collect observations to address the hypotheses. This ultimately leads them to be able to test the hypotheses with specific data — a confirmation (or not) of their original theories [72]. For example, Calo's objective/subjective dichotomy of privacy harms [28] used the deductive approach [12]. Calo classified privacy harms into two categories — subjective and objective — and identified these categories using the top-down approach by establishing overarching principles for distinguishing privacy violations from privacy harms, as well as for including and excluding privacy harms.

3.2.2 Software Engineering Approaches

Empirical studies in software engineering have been widely used and their contributions to scientific knowledge is continuously growing [73]. Easterbrook *et al.* [74] identified and compared five classes of research methods that are most relevant to software engineering: controlled experiments, case studies, survey research, ethnographies and action research. The choice of suitable methods depends mainly upon several factors, including the theoretical stance of the researchers, access to resources and the research problem. In addition, the validity of the results depends on how well the research design limits the shortcomings of the selected methods [74].

In order to investigate a context-sensitive research problem that is hard to study in isolation, a case study is a suitable research method for such software engineering research [73, 74]. The concept 'case study' is used in abstract with terms like field studies [75] and observational methods [76], each of which focus on a specific aspect of the research methodology [73]. Importantly, a case study is closely related to action research as the former is purely observational, whereas the latter is focused on the change process [77]. Case studies provide in-depth insights of how and why certain issues and their cause-effect relationships occur. In contrast, action research is suitable when researchers attempt to solve a real-world problem, and at the same time study the experience of solving the problem [74, 78]. However, there are

practical considerations that need to be taken into account when selecting action research relating to time, budget and personnel resources, and access to data.

Importantly, the validity of empirical studies depends on several criteria that are acceptable by various philosophical stances [74], such as construct validity, internal and external validity, reliability, etc.

3.2.3 Other Approaches for Qualitative Research

In other research domains (e.g. qualitative health care research), there is a growing interest in meta-synthesis as a method for providing new insights and understanding from qualitative research [79]. In particular, it is an intentional and coherent approach to analysing data across qualitative studies [80]. It is a process that enables researchers to formulate a specific research question and then search for, select, appraise and synthesise existing qualitative studies to address the identified research question [80]. Meta-synthesis has an interpretive, rather than aggregating, intent: it can deepen understanding of the contextual dimensions of research settings [79]. As privacy is contextual in its nature, we believe that this research method is suitable for investigating privacy-related research problems.

3.3 The Research Approach

As discussed in Chapter 1, the aim of the research presented in this dissertation is to investigate the areas of data privacy, privacy engineering and software engineering, propose an approach to PbD that captures and addresses privacy issues in the early stages of the design process of software systems, and evaluate this approach through a typical case study. Working at the intersection of disciplines in this way requires a research methodology that is suitable to the research problem and allows the flexibility necessary for investigation in this area: it needs to take into account technical and non-technical considerations of data privacy.

From a methodological perspective, each research area approaches data privacy by its own research methods. These efforts have produced engineering methods that do not address privacy in a meaningful manner. Each of these research areas

can contribute to pertinent knowledge: software engineering provides models and processes for analysis and design methods; security engineering provides security countermeasures and secure design methods. However, taken in isolation, these approaches lack the necessary coverage for resolving the problem. As a result, it is necessary to adopt a multidisciplinary and flexible research strategy. As such, we can conclude that the used research approach needs to address the following considerations.

- The research needs to investigate the research problem within its context.
- The research methods need to avoid the practical challenges relating to time, budget, personnel resources and access to data.
- The research contributions need to practically improve the current practice of privacy engineering.

3.3.1 Research Methods

In this subsection, a detailed account of the research methods used to conduct this research is provided. We mainly used two research methods: *meta-synthesis* is used to develop models and approaches for engineering PbD; and *case studies* are used to illustrate and validate the final synthesis. We consider each in turn.

Meta-Synthesis Method

The meta-synthesis method helps synthesise findings from related studies [81]. In the context of privacy engineering, there have been several investigations that remained isolated (domain or stage-specific) and esoteric (intended for or likely to be understood by only a small number of people with a specialised knowledge — not by multiple stakeholders) and incapable of influencing or informing either policy, practice or research. Thus, it is crucial to bring together and break down the findings of these investigations, analyse them, identify the main features and in some way combine and refine them to produce integrative findings that are more substantive than those resulting from individual investigations.

To ensure the validity of research and that meta-synthesis is conducted in a way that is unbiased and to a degree repeatable, we went through a research process that consists of well-defined steps, as per Figure 3.1. We started by framing the scope of meta-synthesis by formulating a specific research question (see Chapter 1) to ensure that meta-synthesis considers the right elements that need to be taken into consideration. Then, we identified the relevant literature by iteratively undertaking a robust search on the broader topic area as a method of locating and selecting relevant studies to address the research question. The databases that we searched to locate journal articles, conference proceedings, books and other references during our investigation are: IEEE Xplore Digital Library[8], ACM Digital Library[9], Web of Science[10], Scopus[11] and Google Scholar[12]. We completed an electronic search of the literature using terms such as *privacy, data privacy, information privacy, privacy by design, privacy engineering, privacy impact assessment, privacy patterns, privacy design strategies, privacy enhancing technologies, privacy risk analysis* and *privacy risk assessment*. Additional studies were identified based on a review of publication reference lists. Next, we established a set of inclusion and exclusion criteria for the studies located. These include: a study is not domain-specific; findings of a study can be translated into engineering activities; a study that adopts an abstract model that helps translate normative concepts stated in legal frameworks and standards into technical ones that can be applied in the software engineering domain; a study that defines and distinguishes the key notions, risk factors and relationships among these factors; and a study that identifies the key factors that help support the adoption of appropriate design patterns and PETs. We then established a set of criteria to appraise studies derived from [82]. These include: a study is qualitative; the research question is clearly stated; the approach is appropriate for the research question; the study context is described; analysis is appropriate to the research

[8] https://ieeexplore.ieee.org/

[9] https://dl.acm.org/

[10] http://apps.webofknowledge.com/WOS_GeneralSearch_input.do?product=WOS&search_mode=GeneralSearch&SID=D5EeQjK5ueb4Ewcsmjd&preferencesSaved=

[11] https://www.scopus.com/search/form.uri?display=basic

[12] https://scholar.google.co.uk/

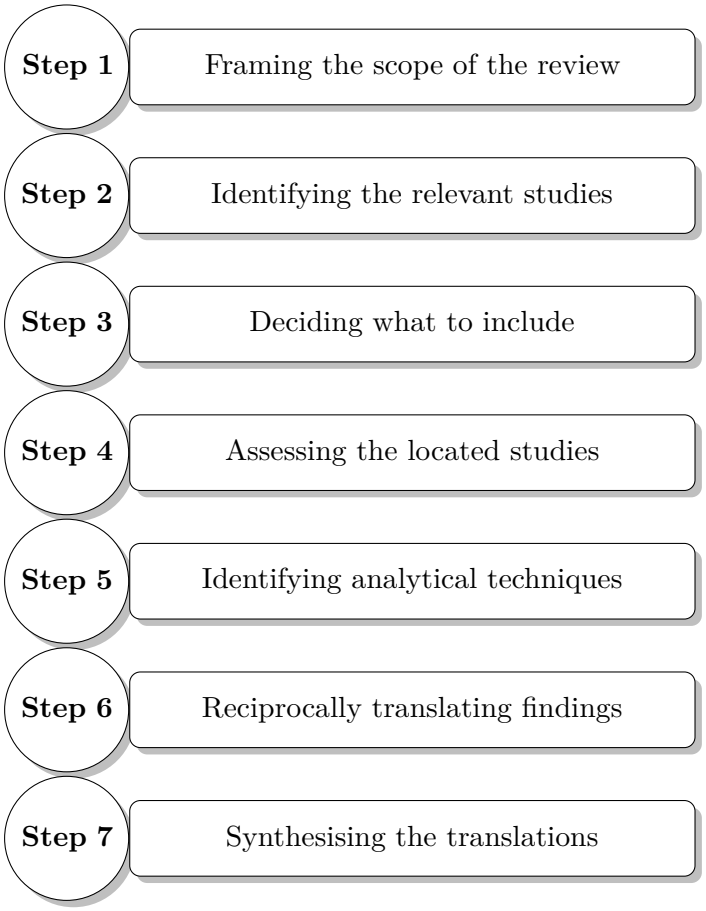

Figure 3.1: The main steps of the meta-synthesis process.

question; and claims are supported by sufficient evidence. Having done that, we defined an analytical technique to determine how studies are related or dissonant. We identified, compared and contrasted key concepts, themes, ideas and relations. Then, we reciprocally translated one study's findings into another using concepts that could be applied to both. Finally, we synthesised the translations to elucidate more refined meanings and new concepts. Ultimately, the final synthesis is the grounds on which the value of meta-synthesis is assessed and it therefore needs to convey how the whole is greater than the sum of the constituent parts.

Case Study Method

The case study research method helps get in-depth understanding of how and why privacy issues occur as a result of inappropriate processing of personal data in a

contextual manner. Further, it can reveal the mechanisms by which cause-effect relationships among these issues occur. The case study research method also provides a flexible design strategy that allows changing key parameters during the course of the study to cope with the complex and dynamic nature of real-world contexts.

To ensure the validity of research and that the case study is conducted in a way that is unbiased and to a degree repeatable, we went through a research process that consists of well-defined steps, as per Figure 3.2. We started by designing and planning the case study by defining the main objective (to investigate how to design a privacy-preserving and data protection-compliant eToll pricing system). This objective is refined into the research question under investigation that needs to be answered through the case study (see Chapter 1). Since the selection of cases is an important step in case study research, we have selected a typical case to gain more insights into common situations. To ensure that the case study considers what is investigated according to the research question, the case and its units of analysis are defined: the case (object of the study) is a software system that contains one unit of analysis (data-processing activities). Then, we defined the study protocol that determines the collection and analysis procedures that are sufficient to fulfil the objective of the case study. By taking practical and ethical considerations into account, the types of data necessary for the study are not collected from the subjects of the study directly; rather, data is collected through *independent* collection methods (i.e. document and text analysis) by analysing multiple data sources, including the relevant literature, legal frameworks, privacy-related standards and documents related to eToll pricing systems. Next, we used general techniques (hypothesis-generating techniques) for data analysis. In particular, data is coded to represent certain themes and concepts that are iteratively used to identify insights. During the iterative process, a set of generalisations can be formulated to result in a formalised body of knowledge, which is the final result of the research attempt. As the selected case study is representative, the aim is to enable analytical generalisation where the results can be extended to cases that have common characteristics and for which the findings are relevant. During the analysis, we maintained a chain of

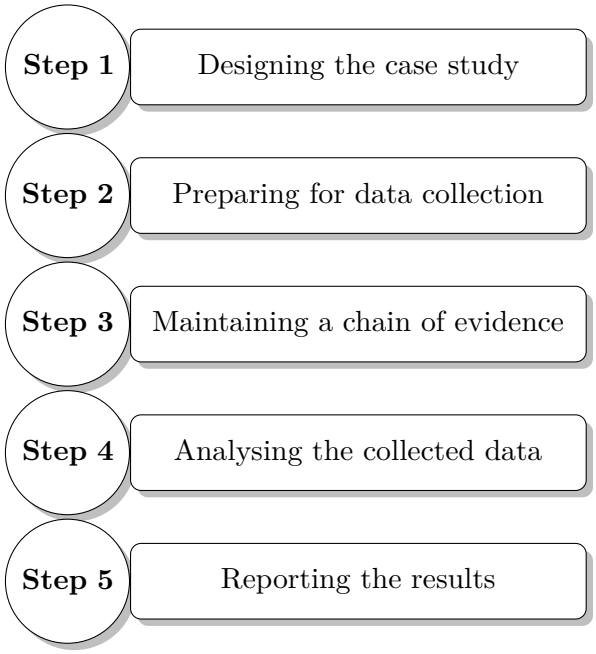

Figure 3.2: The main steps of the case study process.

evidence from the findings to the original data. Finally, we considered the validity of the case study from the beginning by analysing threats to validity and subsequently taking appropriate countermeasures. In particular, triangulation was achieved in different ways: the case study protocol and results were reviewed by peer researchers, and compared to the results that have been critically discussed in the relevant literature. The published papers and presentations in conferences and workshops are referenced as a further means of validating the research approach.

3.3.2 Evaluation

A typical, realistic case study is used to evaluate the applicability and usefulness of our principled approach and its coverage of the problem space. In particular, we use this case study to evaluate the extent to which the abstract personal data lifecycle model: represents data-processing activities and key aspects of privacy principles in a way that is amenable to analysis; captures all necessary and useful information that support the identification of potential privacy risks; and facilitates reasoning about privacy compliance with legal frameworks and standards.

In addition, the case study is used to evaluate the extent to which the data-centric threat modelling approach: defines the key factors that have impacts on privacy risks; establishes relationships among these factors; specifies the range of values these factors can assume during the risk assessment; and describes how combinations of these factors are identified and analysed.

Further, it is also used to evaluate the extent to which the privacy-enhancing tactical approach for achieving various levels of privacy protection supports: the identification of the desired privacy protection goals in a contextual and non-reductive manner; the determination of which combination of tactics, design patterns and PETs will achieve, or contribute to the achievement of, the desired protection goals; and reasoning about architectural choices.

The establishment of an appropriate evaluation framework provides clear evidence of the applicability and usefulness of our principled approach. The application of such a framework reduces the subjectivity during the evaluation, especially when it establishes a set of factors and criteria that can be used as a basis for the evaluation process. As previously mentioned in Chapter 1, the aim of our approach is to complement software engineering methods by capturing and addressing privacy-related issues in the early stages of the design process. Often, the concepts of applicability and usefulness are considered in the context of usability. In software engineering, the concept of usability has been applied to improve the quality of software engineering methods, among others, to measure their applicability and effectiveness/usefulness [83]. However, these concepts have been defined in different ways for different domains. Subsequently, several models and evaluation frameworks have been published [84, 85, 86], some of which cannot be dirrectly applied to engineering methods [83].

For the purpose of this dissertation, we focus only on *applicability* and *usefulness* as two key factors to evaluate our approach in the context of the case study. To do this, we refer to the evaluation framework of [83], not least because it defines applicability and usefulness as key factors, establishes a set of criteria to assess their levels and develops an evaluation model that guides the process

Criteria/(Importance)	Description
Flexibility (Expected)	The ability to change or be easily changed by adopting other techniques or tools to suit a different situation.
Scalability (Mandatory)	The capacity to be changed in size, scale or different application domains.

Table 3.1: Criteria for applicability.

Criteria/(Importance)	Description
Completeness (Essential)	The activities of the approach consider privacy issues throughout the entire data lifecycle.
Capability (Mandatory)	The ability of the approach to solve privacy-related complex problems.

Table 3.2: Criteria for usefulness.

of evaluation. It uses *feature analysis* as an evaluation method, which can be applied in different ways, including case studies. By referring to this framework, we define these factors as follows.

1. ***Applicability:*** the degree to which the approach is convenient for, relevant to, or appropriate for, the context to which it is applied.
2. ***Usefulness:*** the degree to which the approach is successful in producing desired results.

In addition, we select and refine relevant criteria for each factor and assign importance levels for each criterion, as per Table 3.1 and Table 3.2. The mandatory level of importance reflects that the feature is important and the approach must support it, whereas the essential level reflects that the feature is critical and the approach should support it. The expected level reflects that the feature is desirable and the approach is likely to support it. To assess the two factors, we adopt the same fulfilment levels (full support, partial support, little support and no support) to be assigned to the criteria listed in Table 3.1 and Table 3.2 to reflect how applicable and useful the approach is in the context under consideration.

3.4 Case Studies

To help address the research question, we need to choose realistic case studies that meet particular criteria. These case studies need to have a diversity of privacy issues and requirements.

We utilise two case studies — electronic petition systems and electronic toll pricing systems — that have various privacy issues and features. These case studies are derived from the literature to support detailed exploratory investigations to understand and explain how to proactively embed privacy 'by design' using primarily qualitative analysis. We use these case studies to gain insights into the contexts in which personal data is processed and their effects in the design of privacy-preserving solutions.

We classify these case studies into two categories on the basis of the underlying mechanisms that are applied to implement appropriate data minimisation techniques. The first category concerns mechanisms that conceal individuals' identities, while revealing other data. The second category concerns mechanisms that conceal data related to identifiable individuals, while revealing individuals' identities.

3.4.1 The ePetition System

We use the ePetition system, the aim of which is to implement the European Citizens' Initiative (ECI) [87], as an illustrative case study to demonstrate and highlight the usefulness and practical applicability of our contributions in this particular context. In general, a petition can be submitted either in written or in electronic form. In this dissertation, however, we focus on a petition that can be submitted in electronic form.

Overview

The ECI is an online collection system used to support a formal request to an authority for submitting a proposal for a legal act. It enables EU citizens to invite the European Commission to propose a legal act on issues where it has competence to legislate.

Member State	The mandatory fields
France	Full first names, family names, Permanent residence: (street, number, postal code, city, country), Date of birth, Nationality, Personal identification number/ (identification document type and number), Date
Italy	Full first names, family names, Permanent residence: (street, number, postal code, city, country), Date and Place of birth, Nationality, Personal identification number/ (identification document type, number, issuing authority), Date

Table 3.3: The mandatory fields of a statement of support form by France and Italy.

The organisation of an initiative consists of six steps. The first step involves establishing a citizens' committee of at least seven EU citizens. All of the committee's members need to be permanent residents or citizens of the EU Member States and old enough to vote in elections to the European Parliament. This committee acts in its capacity as the official organiser of the initiative and is responsible for preparing and managing the initiative. Second, the organisers need to prepare an initiative and register it with the European Commission. The organisers also need to find a hosting provider when signatures are intended to be collected electronically by an online collection system — either using an instance of the open source software that is provided by the European Commission and hosting it at its site, or by developing their own collection system and using a hosting service provider. For both, organisers need to obtain a certificate from the competent national authority to verify its compliance with minimum technical requirements [88]. Then, the certificate should be posted in the online collection system. Next, individuals, who act as signatories, are able to submit their personal data and their statements of support. To give their support for the initiative, signatories need to provide the specified personal data. In each Member State, the mandatory fields that are required to sign up to an initiative are variable according to relevant national regulations. For illustration purposes, therefore, we illustrate two scenarios that state the mandatory fields of a statement of support form by France and Italy, as per Table 3.3. It is important to ensure that duplicate signatures by the same individual are avoided. Having reached the required number of signatures, organisers should send this personal data to relevant

competent national authorities for verification and certification. Having received all certificates from competent national authorities, organisers should submit the initiative by sending these certificates to the European Commission.

In accordance with the EU's GDPR [6] and the Regulation (EU) No. 211/2011 on the Citizens' Initiative [89], organisers and competent national authorities act as data controllers. In particular, organisers are required to notify the Data Protection Authority in the EU Member State where the personal data will be processed. They are also required to apply appropriate measures to protect personal data in compliance with the EU's GDPR and relevant regulations. This includes that personal data must be "adequate, relevant and not excessive" in relation to the purpose of supporting the initiative and verifying the statements of support. Accordingly, the organisers and the competent national authorities must ensure that collected personal data is not used for purposes other than those specified for supporting the initiative and verifying the statement of support respectively. In addition, the data controllers must destroy all statements of support and any copies one month after submitting the initiative to the Commission or issuing the certificate respectively.

Purposes

The main purpose of the ECI is to verify and certify the number of valid statements of support for a proposed citizens' initiative.

Privacy Issues

In order for signatories to support a specific initiative, they are required to provide 'identifying' personal data (as per Table 3.3), which is typically retained in databases. Since absolute security cannot be achieved in realistic scenarios, it is unrealistic to expect that personal data breaches will not occur — especially as initiatives can be launched and organised by a variety of entities that may have insufficient security expertise and financial means to apply appropriate and reasonable security measures [14].

Code	Type	Threat Actor
TA.1	Insiders	Organisers
TA.2	Outsiders	Intelligence and security services
TA.3	Outsiders	Advertisers
TA.4	Outsiders	Adversaries
TA.4	Outsiders	Employment agencies

Table 3.4: The identified threat actors in the context of the ePetition system.

Code	Vulnerability
VU.1	Improper data models
VU.2	Improper security configurations

Table 3.5: The identified vulnerabilities in the context of the ePetition system.

In addition, complying with the EU's GDPR and the Regulation (EU) No. 211/2011, does not guarantee the elimination of potential privacy risks that may arise as a result of collecting and processing this type of personal data, which may reveal sensitive information related to the individual's political affiliation, religious beliefs and sexual preferences. In particular, the manner in which this data is collected, processed and retained for verification and certification purposes implies that it might be exploited by a wide range of threat sources (insiders or outsiders; individuals, institutions or governments) with capabilities to process this data (lawfully or unlawfully, fairly or unfairly), whether accidentally as a result of weaknesses in privacy controls (technical, organisational or legal), or deliberately for illegitimate and malicious purposes.

We analyse the identified privacy risks of [14, 90] and then we characterise these risks in terms of their factors: threat sources, vulnerabilities and threat events (as per Table 3.4, Table 3.5 and Table 3.6 respectively).

Service Integrity Requirements

In order to achieve the specified purposes in a complete, coherent and accountable manner, there are a set of 'service integrity requirements' [50] that need to be considered. By these requirements, we mean those that allow involved actors to check whether or not others acted responsibly according to their roles and responsibilities.

Code	Source of Data	State of Data	Threat Event
TE.1	C	Data-at-rest	Unauthorised data disclosure
TE.2	C	Data-at-rest	Unauthorised access
TE.3	C, A and/or D	Data-at-rest	Excessive data collection
TE.4	C and/or A	Data-at-rest	Unjustified data integration
TE.5	C, A and/or D	Data-at-rest	Excessive data inference
TE.6	C and/or D	Data-at-rest	Unauthorised secondary use
TE.7	C, A and/or D	Data-at-rest	Unjustified data retention
TE.8	C and/or D	Data-in-motion	Traffic analysis

C: data elements are directly *collected* from the data subject.
A: data elements are *acquired* from external sources.
D: data elements are *derived* from other types of data.

Table 3.6: The identified threat events in the context of the ePetition system.

1. The authenticity of signatures.
2. The eligibility of signatories.
3. The detection and avoidance of duplicate signatures.
4. The correctness of the number of valid signatures.
5. The anonymity of signatories and the assurance of permitting full functionality.
6. The sufficiency of anonymity mechanisms.
7. The possibility of revocation for accountability purposes.

3.4.2 The eToll Pricing System

We now introduce the eToll pricing system, the aim of which is to implement the European Electronic Toll Service (EETS) [91], which we use to evaluate the contributions of this dissertation with reference to the set of criteria established in Chapter 2. We have chosen this case study for the following reasons.

1. The case study is similar to other cases that charge and bill users according to their actual consumption, such as Smart Metering [92] and Pay-As-You-Drive insurance [93]. As such, this case can be considered as a representative or typical case study that informs about common situations.

2. The case study has been critically analysed in the literature [14, 94] with regards to privacy risks and their countermeasures at an architectural level.

3. EETS is regulated by Directive 2004/52/EC of the European Parliament and of the Council of 29 April 2004 on the interoperability of electronic road toll systems in the Community [95] and the related Commission Decision 2009/750/EC of 6 October 2009 on the definition of the European Electronic Toll Service and its technical elements [96]. This gives us the opportunity to reflect on the abstract privacy principles stated in these legal frameworks.

4. The European Commission provides full details about EETS by publishing a guide as a reference manual for all parties concerned by the Directive and the Decision. The guide illustrates references and procedures to help the implementation of electronic road toll systems interoperability and EETS [91].

Overview

EETS aims to support interoperability between national electronic road toll systems at a European level to electronically calculate and collect road-usage tolls. The calculation of road-usage tolls is based on a variety of parameters: the distance, the type of road, traffic density, the time of usage and vehicle classification parameters. EETS is intended to cover all domains and objects that are subject to toll, such as road networks, specific sections of roads (e.g. a bridge, a tunnel or a ferry connection), or specific areas offering services (e.g. a parking lot or access to a protected area in a city). It enables road users to easily pay road-usage tolls throughout the Member States with a single subscription contract with a service provider [91]. EETS is independent of the decisions taken by Member States to levy tolls based on a particular category of vehicles, or a specific road-charging policy; it is only a method of collecting road-usage tolls.

The main actors involved in the EETS are service providers, toll chargers and users. Service providers are legal entities that grant access to EETS to road users [96]. Toll chargers are public or private organisations that are responsible for levying tolls for the circulation of vehicles in an EETS domain [96]. Users are individuals who subscribe to EETS providers in order to get access to EETS, regardless of nationality, country of residence or the Member State in which the

vehicle is registered [96]. By signing a contract, a user is required to provide a set of data — user and vehicle classification parameters — 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 service 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 [91].

Prior to explaining the essential elements of the EETS architecture, we introduce some relevant concepts [96].

- *Toll domain* is a road network or a specific section of a road that is subject to toll, such as a tunnel, a bridge or a ferry.
- *Toll declaration* is a statement sent to a toll charger to confirm the circulation of a vehicle in its toll domain.
- *Toll transaction* is an action or a set of actions in which a toll declaration is sent to a toll charger.
- *Toll context data* is data defined by the responsible toll charger to describe the location of the toll domain, charging policies, and the format of toll declarations.
- *Domain statement* is a statement that is developed and maintained by a toll charger to establish technical specifications and a set of contractual terms and conditions for a service provider for accessing its toll domains.

Figure 3.3 illustrates the essential elements of the EETS architecture, together with interfaces via which data is exchanged between these elements [91]. Service providers and toll chargers are required to implement the relevant interfaces in their OBUs and fixed or mobile Roadside Equipment (RSE) respectively, as well as in their back-office systems.

- *Interface 1* is used for data exchange between the OBU and the service provider's back-office systems. This includes configuring the OBU with contract data and vehicles' classification parameters, updating the OBU with toll context data, and sending toll declarations. The OBU either collects, stores and declares tolls locally, or collects and relays location data to the relevant service provider for further processing by its back-office systems. In both cases, the service provider needs to verify the toll declarations to ensure that they correspond to the actual road-usage of users.

- *Interface 2* is used for data exchange between back-office systems of the service providers and toll chargers. This includes toll declarations, toll context data, and domain statements. In addition, service providers are required to send toll transactions to toll chargers to be used for enforcement purposes.

- *Interface 3* is used for data exchange between toll chargers' back-office systems and RSE, including toll declaration and enforcement data. Using RSE, toll chargers can identify a service provider's contract on the basis of the vehicle's registration number or any attribute of the OBU. This data can be used as a way in which toll charger can manage enforcement by detecting whether a vehicle circulating on associated toll domain is actually equipped with a validated and properly functioning OBU.

- *Interface 4* is used for data exchange between a toll charger's RSE and a service provider's OBU. This includes toll declaration data and vehicles' classification parameters that can be used for compliance checking.

Each Member State is required to keep an up-to-date and accurate national electronic register that contains relevant toll domains, corresponding toll chargers, toll context data, domain statements and service providers along with their area of competence [96]. Importantly, EETS provision entails personal data processing, which must be carried out in compliance with the EU's GDPR [6] and Directive 2002/58/EC [97]. Thus, data controllers are required to apply appropriate security

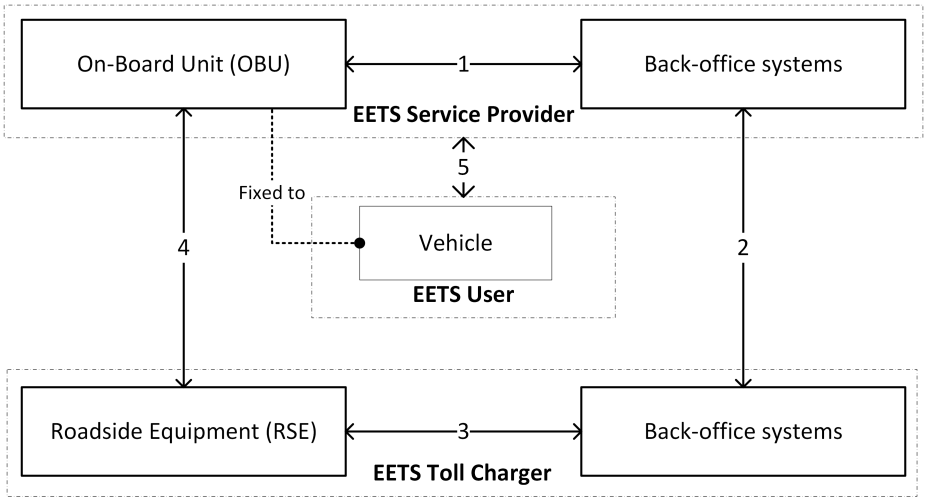

Figure 3.3: The essential elements of the EETS architecture.

measures to protect the collected personal data, ensure it is used only for the specified purposes, and retained only as long as necessary to achieve these purposes or for the stated period of time.

Purposes

The main purpose of the electronic road toll system is to calculate personalised road-usage tolls and communicate the final premium to EETS user at the end of the tax period.

Privacy Issues

In order to calculate and communicate the final fees for EETS users, fine-grained location data needs to be collected over time. This type of data, along with identification and contact data, are typically retained in databases to achieve the specified purposes. However, absolute security cannot be achieved in realistic scenarios. As such, it is unrealistic to expect that personal data breaches will not occur — especially as EETS can be provided by a variety of entities that may have insufficient security expertise and financial means to apply appropriate and reasonable security measures [14].

Complying with the EU's GDPR [6] and Directive 2002/58/EC [97] does not guarantee the elimination of potential privacy risks that may arise as a result of

Code	Type	Threat Actor
TA.1	Insiders	Service providers
TA.2	Outsiders	Department for Transport
TA.3	Outsiders	Police
TA.4	Outsiders	Tax authorities
TA.5	Outsiders	Eavesdroppers
TA.6	Outsiders	Communication providers

Table 3.7: The identified threat actors in the context of the eToll pricing system.

Code	Vulnerability
VU.1	Improper data models
VU.2	Improper security configurations

Table 3.8: The identified vulnerabilities in the context of the eToll pricing system.

collecting and processing these types of personal data, which may reveal sensitive information related to the individual's political opinions, religious beliefs, sexual preferences, philosophical beliefs and so forth. In particular, the manner in which this data is collected, processed and retained for billing purposes implies that it might be exploited by a wide range of threat sources (insiders or outsiders; individuals, institutions or governments) with capabilities to process this data (lawfully or unlawfully, fairly or unfairly), whether accidentally as a result of weaknesses in privacy controls (technical, organisational or legal), or deliberately for illegitimate and malicious purposes.

We analyse the identified privacy risks of [14, 90] and then we characterise these risks in terms of their factors: threat sources, vulnerabilities and threat events (as per Table 3.7, Table 3.8 and Table 3.9 respectively).

Service Integrity Requirements

In order to achieve the specified purposes in a complete, coherent and accountable manner, there are a set of service integrity requirements that need to be considered when making privacy-preserving choices:

1. The correctness of the final fees.
2. The detection and avoidance of fraud.

Code	Source of Data	State of Data	Threat Event
TE.1	C and/or D	Data-in-motion	Traffic analysis
TE.2	C and/or A	Data-at-rest	Unjustified data aggregation
TE.3	C, A and/or D	Data-at-rest	Unjustified identification
TE.4	C and/or D	Data-at-rest	Data inference
TE.5	C and/or D	Data-at-rest	Unauthorised secondary use
TE.6	C and/or D	Data-at-rest	Undesirable disclosure

C: data elements are directly *collected* from the data subject.
A: data elements are *acquired* from external sources.
D: data elements are *derived* from other types of data.

Table 3.9: The identified threat events in the context of the eToll pricing system.

3. The disclosure of minimum amount of location data and the assurance of permitting full functionality.
4. The possibility of revealing the committed values for accountability purposes.

Solutions

A Privacy-Preserving Electronic Toll Pricing (PrETP) has been proposed in [14, 94] to reveal the minimum amount of location data. PrETP has been proposed in terms of a decentralised architecture that meets two main requirements: the service provider needs to know the final fee for each user; and the provider must be reassured that this fee is correctly calculated and that users cannot commit fraud. In accordance with such an architecture, OBUs calculate the final fees locally and transmit them to the service provider at the end of the tax period. To ensure that the final fee is calculated correctly, PrETP uses cryptographic commitments to allow a user to commit to a value without having to disclose it. PrETP relies on the assumption that the toll charger has access to evidence (such as a photograph taken by RSE) proving that a vehicle was at a specific location at a specific time.

3.5 Summary

In this chapter, we have reviewed relevant research approaches from the legal, socio-technical and software engineering research communities. Then, we described the research approach used in this dissertation, and justified its relevance to the

research problem. Finally, we characterised two case studies in a way that is amenable to analysis and investigation.

In the next three chapters, we will illustrate the main activities of our principled approach for engineering PbD: data-processing representation, data-centric threat modelling, and privacy-enhancing strategies.

The task for privacy-aware engineers and systems architects is to translate the PbD conceptual framework into a set of specific, and operationally feasible, tools.

— Ann Cavoukian

4

Privacy-Aware Data Lifecycle Models

Contents

4.1 Introduction

This chapter gives a detailed description of an abstract model for personal data lifecycle models, which contributes to the first activity of our principled approach for engineering PbD. The abstract model provides a straightforward characterisation of the processing operations performed on personal data in a way that is amenable to *risk analysis* and *compliance checking*. In particular, Section 4.2 gives an overview of the Abstract Personal Data Lifecycle (APDL) model as a means for representing the personal data lifecycle in terms of lifecycle stages, along with associated activities and involved actors, as well as for facilitating the management and traceability of personal data. It also illustrates how abstract purposes for which personal data is collected, processed and disseminated can be operationalised

(refined into concrete purposes). In addition, it describes the conceptual model around which the APDL model is developed. Then, Section 4.3 explains the profile that allows the APDL model to be represented in the UML as a meta-model. The UML profile can be used as a means for data-processing representation. Finally, Section 4.4 gives a brief summary of the chapter. The work presented in this chapter previously appeared in [98, 99, 100].

4.2 An Abstract Personal Data Lifecycle Model

In the context of data-centric domains, data undergoes a variety of actions — including creation, use, publication and destruction — by several actors for various purposes. These actions in combination constitute a data lifecycle. It is understandable that each domain is concerned with a specific type of data and each data lifecycle model has its own specific focus. Most importantly, they all consider the same item of interest — data, which is a "living thing" that moves though various stages during its lifecycle and is at the heart of these systems [101]. In the context of privacy and data protection, personal data often moves through various stages that are governed by laws, regulations or standard principles. Accordingly, personal data should be at the heart of methods, techniques and tools that systematically and proactively identify and address privacy-related issues at the early stages of the design process.

The Abstract Data Lifecycle Model (ADLM) [101] was derived from specific instances of data lifecycle models to ensure broad coverage and wide applicability [101]. For each domain, a list of models was analysed in terms of their lifecycle phases, features, roles, actor features and metadata features. By analysing, comparing and contrasting these models, the ADLM was derived as an abstract data lifecycle model for data-centric domains. It establishes five areas of classification: lifecycle phases, features, roles, actor features, and metadata features. The ADLM serves as a generic data lifecycle model for data-centric domains, and can be used as a means to classify, compare and relate other data lifecycle models, as well as to provide the basis for new data lifecycle models [101].

In the context of privacy and data protection, an abstract personal data lifecycle model can similarly play a crucial role in providing privacy-aware data lifecycle models. As such, we consider only the parts of the ADLM that are relevant to representing personal data processing activities in a way that is amenable for analysis: *lifecycle stages* and *roles*. In particular, we analyse the lifecycle stages and roles of the ADLM and specialise these stages and roles for the purpose of this dissertation. Some of its stages will be combined (generalised) according to their characteristics and associated activities, while others will be defined (specialised) to limit associated activities to particular privacy principles. Those stages not relevant to personal data will not be considered.

4.2.1 A Textual Description

The Abstract Personal Data Lifecycle (APDL) model represents personal data processing in terms of *states* (data items), *operations* (processing activities) and *roles* (a set of responsibilities that can be assigned to different actors according to their capabilities). It defines a set of stages through which personal data moves during its lifetime (from collection to destruction) and indicates the *order* and *depth* in which associated activities can occur. As such, the APDL model: serves as an abstract model for personal data lifecycles; and distinguishes between the main types of operations that can be performed on personal data during its lifecycle by outlining the various distinct activities for each operation. We consider each in turn.

Lifecycle Stages

As there are obligations and limitations on the stages of the personal data lifecycle and associated activities, it is essential to adopt a set of universal privacy principles that can be applied in a variety of contexts in various jurisdictions. Due to the dynamic and contextual nature of privacy, legal frameworks vary between jurisdictions, whereas standards do not; to some extent, standards are commonly used practices [13], especially with the absence of global data protection laws and regulations. The GPS principles harmonise various sets of the FIPPs into universal

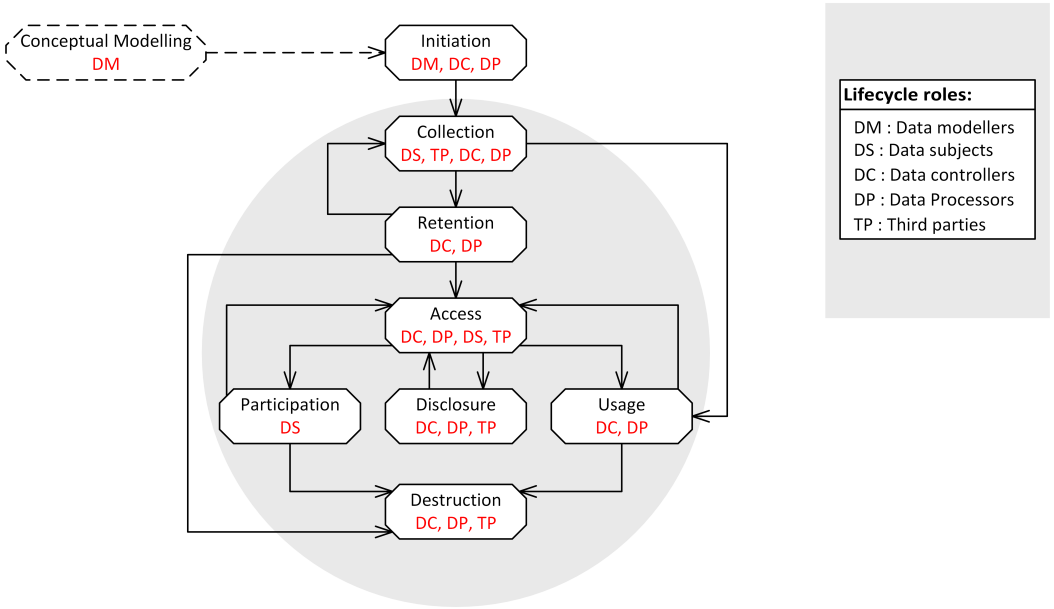

Figure 4.1: The Abstract Personal Data Lifecycle (APDL) Model.

privacy principles. As such, we adopt the GPS principles to impose constraints on the stages of the lifecycle to govern the behaviour of associated activities. For example, 'Purposes' emphasises that the purposes for which personal data is collected, retained, used and disclosed must be clearly specified and communicated to data subjects at or before the time of collection, and 'Collection Limitation' affirms that the collection of personal data must be fair, lawful and limited to the specified purposes. Such constraints can be initially described using activity diagrams for each lifecycle stage. For each stage, we establish a set of assessment criteria (checkpoints or privacy assessment activities) to proactively address privacy-related issues in relation to relevant GPS principles, as per Appendix A.

Figure 4.1 abstractly illustrates the main stages of the APDL model, their logical dependencies on other stages, and relevant lifecycle roles. We describe each stage in terms of associated activities, dependencies on other stages and relevant GPS principles.

0. **Conceptual Modelling:** This is a preliminary stage: it is a prerequisite for any data lifecycle model.

Activities. The Conceptual Modelling stage involves activities to develop a conceptual model that describes the problem and its solution in terms of the domain vocabulary. These include: specifying the purpose of the modelling along with its scope, intended users and uses; and determining the most appropriate technique for deriving useful and potentially usable concepts, together with their properties, meanings, relationships, corresponding actions and associated constraints. The conceptual model provides terms that can be used to facilitate communication with multiple stakeholders.

Dependency. As per Figure 4.1, the Conceptual Modelling stage is represented by a dotted line to distinguish that it is not a core part of the personal data lifecycle. The output of this stage is a conceptual model that represents key and relevant concepts, associated meanings, properties, relationships and constraints that restrict the semantics of the concepts and their conceptual relationships.

Principles. All the GPS principles serve as the foundation for the Conceptual Modelling stage and associated activities.

1. **Initiation:** This stage precedes the collection of personal data and represents the first stage of the lifecycle.

   Activities. The Initiation stage involves activities to specify a 'processing plan' in relation to the context. These include: specifying the items of personal data to be collected; the sources of these items; the choices available to data subjects and the types of consent to be obtained; the methods of collection, retention, retrieval, disclosure and destruction; the purposes for which personal data is collected, processed and/or disseminated; involved actors, assigned roles and responsibilities; relevant legal frameworks and standards; and domain-specific constraints that restrict the processing of personal data.

   It is important to be explicit as to explain whether the purposes for which personal data is collected involve the intention to collect, derive or acquire

additional personal data items, whether from existing items or external sources. In addition, the items of personal data need to be adequate, relevant and not excessive in relation to the purposes for which personal data is collected. The specification of the minimum amount of personal data can be initially achieved by developing data models, i.e. logical and physical data models that represent context-related entities, associated properties and semantic relationships as shared knowledge for a particular application domain.

Dependency. The input is from the Conceptual Modelling stage, as well as relevant privacy policies and procedures. The output is a 'complete processing plan' that determines the purposes for, and the manner in which, personal data is collected, retained, used, disclosed and destroyed. The processing plan is considered as the basis for establishing a privacy notice to be communicated to data subjects.

Principles. All the GPS principles serve as the foundation for the Initiation stage and associated activities.

2. **Collection:** It follows the Initiation stage and precedes the Retention stage.

   Activities. The Collection stage involves activities pertaining to recording, capturing, collecting or preparing personal data values for the storage, whether these values are directly collected from data subjects, or have been acquired from external sources. In both cases, it is important that these values have not existed in the lifecycle before the Collection stage. The most important aspects in this stage are the set of personal data values, associated sources (i.e. whether they are primary or secondary sources), and the methods of collection (i.e. whether they actively or passively collect data values).

   Dependency. The input is a data processing plan from the Initiation stage. The output is a set of recorded, captured or collected personal data values to be used as inputs to the Retention stage.

   Principles. Openness, Purposes, Consent and Collection Limitation.

3. **Retention:** It follows the Collection stage and precedes the Access and Destruction stages.

   Activities. The Retention stage involves activities pertaining to organising, structuring or storing personal data values for a specific period of time in repositories or digital storage media. These include: specifying the retention period of personal data to fulfil the specific purposes or to comply with legal requirements; specifying the storage methods (i.e. whether they are in the user or service domain); organising, structuring or storing personal data values; and making additional copies of the original data for future, compliance or operational purposes, such as archiving and backup. These activities constitute the principal activities of the Retention stage: primary storage, archiving and backup.

   Dependency. The input is a set of personal data values from the Collection stage. The output is a set of personal data values that will be stored and retained for a specific period of time as specified in the processing plan.

   Principles. Use, Retention and Disclosure Limitation.

4. **Access:** It follows the Retention stage. In particular, it follows the activities of primary storage and occasionally follows the activities of archiving for regulatory compliance purposes.

   Activities. The Access stage involves activities pertaining to specifying and retrieving or consulting personal data values stored in repositories or digital storage media: personal data is made accessible for use by involved actors, whether they are internal users, external users or data subjects. Involved actors gain access to the stored personal data and are able to retrieve this data to perform specific actions according to their roles and responsibilities as specified in the processing plan. Data retrieval is not restricted to specific mechanisms, such as using query languages; rather, it can be accomplished by using interfaces or any mechanisms that allow the stored data to be searched, retrieved and appropriately displayed.

Dependency. The input is a set of personal data values from the Retention stage. The output is a set of personal data values to be used as inputs to the Participation, Usage and Disclosure stages.

Principles. Access, Use, Retention and Disclosure Limitation.

5. **Participation:** It follows the Access stage and may precede the Destruction stage.

   Activities. The Participation stage involves activities pertaining to implementing access rights and altering personal data values by data subjects to ensure that their data is accurate, complete and up-to-date. These include: specifying retrieval methods; presenting the previously accessed and retrieved personal data values; altering these values if they are inaccurate; and confirming alterations and inform consequences if consent has been withdrawn or preferences have been changed. These activities constitute the principal activities of the Participation stage: review and alteration. The most important points in this stage are providing data subjects with access to exercise control over personal data and specifying the means by which data subjects can review, update and correct this data.

   Dependency. The input is a set of personal data values from the Access stage. The output is a set of reviewed, updated or corrected personal data values that can be stored and used in the Usage or Disclosure stages.

   Principles. Access and Accuracy.

6. **Usage:** It follows the Access stage and may precede the Destruction stage.

   Activities. The Usage stage involves activities pertaining to manipulating and using personal data values in conformance with the specified purposes. These include classifying, altering, adapting, synthesising, combining, integrating or analysing personal data values. These activities constitute the principal activities of the Usage stage: classification, manipulation, synthesis and analysis.

Dependency. The input is a set of personal data values from the Access stage. The output is a set of altered, adapted, aligned, integrated, derived or used personal data values that can be stored and used again in the Usage or Disclosure stages.

Principles. Use, Retention and Disclosure Limitation.

7. **Disclosure:** It follows the Access stage.

   Activities. The Disclosure stage involves activities for disseminating, making available or transmitting the previously accessed and retrieved personal data for external use by third parties, to perform further activities, including classifying, manipulating, analysing or integrating several data items from various sources. These may include further processing for historical, statistical or scientific purposes. Most importantly, the disseminated personal data values need to be used only for the specified purpose for which data subjects have provided their explicit or implicit consent.

   Dependency. The input is a set of personal data values from the Access stage. The output is a set of disseminated, made available or transmitted personal data values to be used by third parties for further processing.

   Principles. Use, Retention and Disclosure Limitation.

8. **Destruction:** It follows the Retention, Collection, Participation and Usage stages and is the final stage.

   Activities. The Destruction stage involves activities for erasing, destroying, redacting or disposing of personal data in accordance with relevant retention and destruction policies. These include: specifying and removing or redacting specific items of personal data that can serve as identifiers or quasi-identifiers; completely and permanently erasing personal data items and disposing of original, archived and backup copies of these items; and destroying digital storage media in accordance with destruction policies. This indicates that the

most important points in this stage are the permanent destruction, erasure or redaction of personal data, as well as the methods of destruction.

Dependency. The input is a set of personal data values from the Retention, Collection, Participation or Usage stages. The output is a set of erased or redacted personal data values.

Principles. Use, Retention and Disclosure Limitation.

Lifecycle Roles

A lifecycle role is a set of logically-related activities that are expected to be conducted together and assigned to different actors according to their capabilities. They define the ways in which actors participate in the activities of the lifecycle stages. Each actor may be assigned to one or many roles and each role will typically be associated with one or more lifecycle stages.

1. *Data modellers* are involved in the Conceptual Modelling stage and establish the context in which personal data is collected, processed and disseminated. In addition, they are involved in the Initiation stage and are responsible for developing logical and physical models for the application domain.

2. *Data subjects* are involved in the Collection stage with the capability of providing their personal data. Such actors can actively participate in the collection of the personal data values. Data subjects may be involved in the Access and Participation stages with the capability of accessing and rectifying their personal data. Actors with this ability can access, review, update or correct their personal data to ensure that the retained personal data is accurate.

3. *Data controllers* are actors who specify the purposes for, and the manner in which, personal data is to be collected, processed and disseminated. They are involved in the Initiation, Collection, Retention, Access, Usage, Disclosure and Destruction stages with administrative capabilities. Such actors are responsible for handling personal data items without changing its format or meaning.

If the data controller is a data processor, administrators are responsible for archiving, making backup copies, disclosing and destroying personal data items. The administrative capabilities may also include other activities, such as those related to compliance monitoring and audit trails. Data controllers are also involved in the Access and Usage stages with different levels of user capabilities. Such actors manipulate and use the retained personal data items according to the purposes for which this data is collected. They perform data-processing activities, including classification, analysis, manipulation, combination or other actions as per the processing plan.

4. *Data processors* are actors who process the collected personal data on behalf of the data controller. They are involved in the Retention, Access and Usage stages and process personal data items without changing their format or meaning. Such actors are responsible for archiving, making backup copies and destroying personal data items according to the data controller instructions. The role of data processors may also include other responsibilities, such as those related to operations and performance monitoring.

5. *Third parties* are actors other than data subjects, data controllers or data processors. They may be involved in the Collection stage with data-providing capabilities, i.e. they may be secondary sources other than data subjects. Such actors actively participate in the collection of the personal data values. In addition, third parties may be involved in the Disclosure stage of the lifecycle with data-receiving capabilities. Such actors receive and use the disclosed personal data items only for the purposes specified in the processing plan and with the consent or knowledge of data subjects.

4.2.2 A Conceptual Model for the APDL

As previously mentioned, the APDL model is proposed to serve as a common language that supports a meaningful participation of multiple stakeholders and facilitates communication between those stakeholders. It has the potential to

facilitate the understanding of the meaning of privacy-related concepts and the ways in which systems can be developed to comply with legal frameworks and standards, and to meet data subjects' expectations by supporting the traceability and management of the flows and changes in the states of personal data. However, it is informally represented in terms of the lifecycle stages, associated activities and involved actors. In order for it to be integrated into an appropriate software engineering process, a widely-used modelling notation needs to be adopted to help support its main concepts. The Unified Modeling Language (UML) [102] is ideal for this purpose. This, in turn, requires a conceptual model for privacy engineering that helps derive and model the main concepts of the APDL model along with their properties, relationships and the key aspects of abstract privacy principles as constraints.

Modelling Requirements

To derive and model the key aspects of abstract privacy principles, it is important to define a conceptual model that can be used as a common language to express stakeholders' expectations and concerns. It can be used by multiple stakeholders — both those concerned with privacy and data protection, and those responsible for developing and maintaining privacy-preserving systems.

In order to develop such a model, essential requirements for the core parts of the model, which will be used as the basis for our UML profile, need to specified. Crucially, the *purpose* and *scope* of modelling shall be specified in relation to the context of privacy and data protection. In addition, the *most appropriate technique* shall be used for deriving useful and potentially usable concepts, associated properties and relationships.

We partially specify the purpose and scope of the modelling, as well as the appropriate techniques and conceptualisation approach used.

1. The purpose of building a conceptual model is to describe precisely the key privacy-related concepts, associated properties and relationships in the context of privacy and data protection. The model is intended to be used by

multiple stakeholders — both those concerned with data protection and those responsible for developing and maintaining privacy-preserving systems. The model is intended to be used as a common language for privacy engineering to consider protection, manageability and traceability of personal data. Such a language is provided with the ability to express stakeholders' expectations and concerns.

2. The scope of the modelling is identified by a list of concepts (as explained in the following subsections).

3. Informal text analysis is used with the aim of analysing commonly-used concepts that have been already described in privacy standards instead of 'starting from scratch'. We have chosen the GPS principles [7] and the Generally Accepted Privacy Principles (GAPP) [32] to ensure that our modelling of the privacy-related concepts and their meanings are based upon a widely-used set of terms. These standards are based on internationally known FIPPs [11].

4. Concept classification is used with the aim of analysing and classifying relevant terms into concepts and processing activities that can be represented in a fine-grained manner as actions. With regards to concepts, we identify and describe useful and potentially usable concepts, associated properties, meanings and possible values. With regards to actions, we identify and describe useful and potentially usable actions and associated constraints that specify conditions to be satisfied before, or to be guaranteed after, the execution of corresponding actions.

In order to meet these requirements, we analyse the stages, activities and roles of the APDL model, and the key aspects of abstract privacy principles with the aim of classifying the primary terms into *concepts*, along with associated properties, meanings, possible values, and useful and potentially usable *actions*, together with associated constraints. The concepts are to represent the primary terms, whereas the actions are to ensure that the main operations make sense in the context of

privacy and data protection. The constraints specify conditions to be satisfied before, or to be guaranteed after, the execution of corresponding actions. Each aspect of these principles can be modelled by one or more concepts, which can be characterised by a set of properties, or one or more actions, which can be restricted by a set of constraints.

This leads to the importance of starting by operationalising the purposes for which personal data is collected, processed and disseminated.

Operationalisation of Purposes

To represent data-processing activities and the key aspects of abstract privacy principles in a way that facilitates reasoning about privacy compliance, the abstract purposes for which personal data is collected, processed and disseminated need to be operationalised. The abstract purpose is a nonoperational objective to be directly achieved by the collection and processing of personal data. By operationalising it, we wish to move towards a position in which the objective is expressed in terms of *data items*, *data-processing activities* that consist of concrete *actions* and *events* that cause the execution of these actions, and *roles* that define a set of responsibilities performed by different *actors* according to their capabilities. Further, abstract purposes are made operational through constraints (derived from abstract privacy principles) that specify conditions to be satisfied before, or to be guaranteed after, the execution of corresponding operations (activities).

The *abstract purpose* needs to be refined into a set of concrete purposes that can be assigned to actors as responsibilities. This can be achieved by specifying the abstract purpose at a certain level of detail as concrete purposes (*data-processing activities*) to capture how each activity participates in the fulfilment of the abstract purpose. Instead of specifying abstract purposes in terms of either entities or activities, they are better specified as objectives from which entities and activities can be derived. They can be represented in a hierarchical structure of which the lowest level represents concrete purposes as data-processing activities, which can be assigned as the responsibility of actors, as per Figure 4.2. The *activities* are

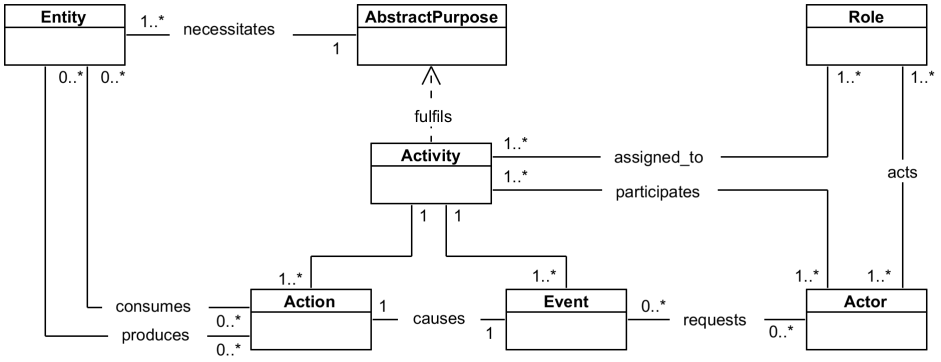

Figure 4.2: The refinement of abstract purposes.

operationalised by *actions* that are triggered by *events*, which are requested by the *actors* responsible for each activity according to the *roles* to which they are assigned in such a way that activities fulfil the abstract purpose. Actors may also process certain entities of *personal data* that are derived when refining abstract purposes. These entities can be consumed or produced by certain actions as inputs or outputs. At a high level of abstraction, the abstract purpose necessitates entities of personal data; at a low level, entities are necessary for an activity if these entities are processed by an action that operationalises the activity.

Example 4.1. In the context of the European Citizens' Initiative (ECI), the *abstract purpose* is to verify and certify the valid number of statements of support for a proposed citizens' initiative. In order for this purpose to be operationalised, it first needs to be refined into a set of concrete purposes. The refinement step is guided by the main operations performed on signatories' personal data in relation to the stages through which this data moves during its lifecycle. The abstract purpose of ECI is operationalised by refining it into five concrete purposes (data-processing activities), as per Figure 4.3.

CollectingStatementsOfSupport is represented as *Activity1*, which collects the specified personal data from at least one million EU citizens who act as signatories for signing-up to a citizens' initiative. ExportingAllStatementsOfSupport is represented as *Activity2*, which exports all statements of support and displays the total distribution of signatures. ExportingSpecificStatementsOfSupport is represented

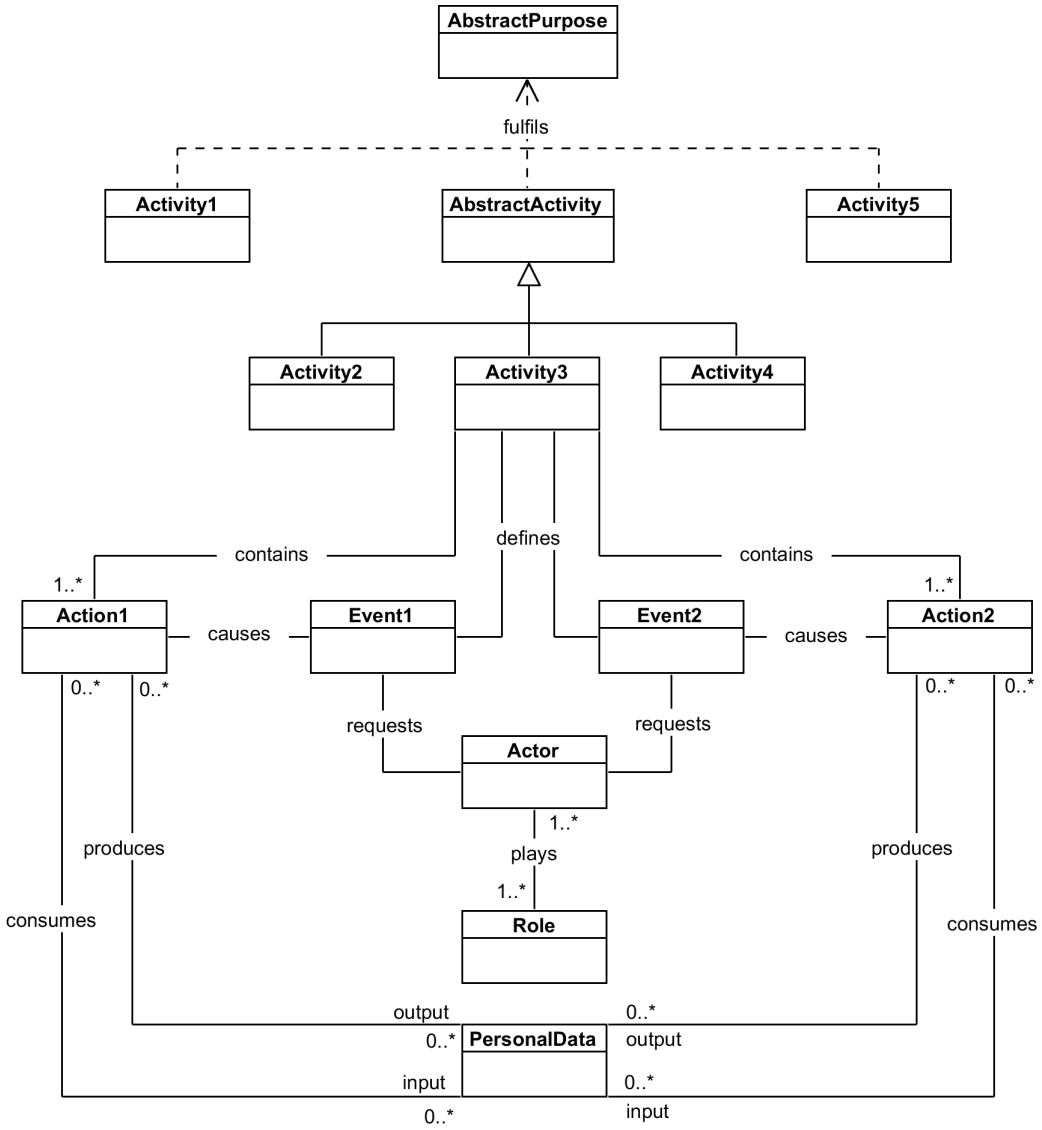

Activity1:	CollectingStatementsOfSupport
Activity2:	ExportingAllStatementsOfSupport
Activity3:	ExportingSpecificStatementsOfSupport
Activity4:	DeletingSpecificStatementOfSupport
AbstractActivity:	MonitorAndExportStatementsOfSupport
Activity5:	ConfirmingValidStatementsOfSupport
Action1:	ExportByMemberState
Event1:	ExportByCountry
Action2:	ExportBySubmissionDate
Event2:	ExportByDate
Actor:	ECIUser
Role:	ECIOrganiser
PersonalData:	StatementsOfSupport

Figure 4.3: The refinement of the ECI's abstract purpose.

as *Activity3*, which exports the selected statements of support according to the Member State and/or the date of submission for reporting purposes. Delet-

ingSpecificStatementOfSupport is represented as *Activity4*, which deletes the selected statement of support according to the signature identifier and optionally the date of submission. *Activity2–Activity4* are generalised by an abstract activity — MonitorAndExportStatementsOfSupport — which is represented as *AbstractActivity*. ConfirmingValidStatementsOfSupport is represented as *Activity5*, which verifies and certifies the number of valid statements of support.

In this example, we illustrate only how to express *Activity3* in terms of actions and events that trigger the execution of these actions. The activity coordinates its execution via two actions and two corresponding events. ExportByMemberState is represented as *Action1*. ExportByCountry is represented as *Event1*. It specifies the occurrence of retrieving specific statements of support according to the Member State to export their corresponding data. ExportBySubmissionDate is represented as *Action2*. ExportByDate is represented as *Event2*. It specifies the occurrence of retrieving specific statements of support according to the submission date to export their corresponding data.

Based on the identified actions, the minimum amount of required personal data items can be initially derived. For *Action1* and *Action2* to accomplish their execution, they consume and/or produce personal data items, including name, address and date of birth. These data elements, the mandatory fields of the statement of support form, are represented as *PersonalData* in Figure 4.3.

Having expressed *Activity3* in terms of actions and events, it needs to be assigned to capable actors who participate in its performance according to their capabilities. ECIUser, represented as *Actor*, is an actor who is capable of, and responsible for, performing the activities of the ECI organiser. ECIOrganiser, represented as *Role*, is a data controller role that involves logically-related activities for monitoring and exporting the collected statements of support.

Abstract Purposes and Personal Data

AbstractPurpose represents goals and reasons for which personal data is collected, processed and/or disseminated. It has the following properties: *informalDescription;*

actualPurpose (a statement in a concrete and explicit manner); *isFair* (indicates whether the processing of personal data has justified adverse effects on the concerned data subjects and is consistent with their reasonable expectations); *isLawful* (indicates whether there are legitimate or legal grounds for collecting and processing personal data); *isProportional* (indicates whether the specified purpose is legally and/or politically proportional); and *relevantPrinciple* (specifies the relevant GPS principles that govern the purpose specification in the sense of placing limitations or constraints).

For an abstract purpose to be fulfilled, a minimum amount of personal data needs to be appropriately specified. As such, ***PersonalData*** represents the minimum necessary amount of data that is sufficiently related to an identified or identifiable individual in support of the specified purpose. It has the following properties: *informalDescription; category* (indicates the category of personal data in terms of its sensitivity and the manner in which it is to be processed, and drawn from {SpecialCategory, Unspecified}[13]); *type* (indicates the type of personal data in relation to the source and manner in which it is created, and drawn from {Collected, Acquired, Derived}); and *linkability* (indicates the linkability of data to personal identifiers, and drawn from {Linked, Linkable with reasonable effort, Not linkable with reasonable effort or Unlinkable}).

In order to initially specify the minimum amount of personal data that fulfils the specified purpose, a data model needs to be constructed; ***DataModel*** represents the relevant objects, associated properties, relationships and constraints for the purpose of specifying the required data. This representation can be used as shared knowledge by multiple stakeholders for a specific application. It has the following properties: *subjectDomain; modellingPurpose;* and *modellingScope* (specifies the set of objects to be represented at an appropriate level of abstraction).

[13]Personal data, by its nature, is considered sensitive data when it is related to special categories, including racial or ethnic origin, etc. [6].

Lifecycle Stages

DataLifecycle represents the main characteristics of the personal data lifecycle in terms of the openness of the processed data and the centrality of its underlying system. It has the following properties: *informalDescription; isOpen;* and *isCentralised.* Each data lifecycle consists of a set of stages. As such, ***LifecycleStage*** represents the concept of a generic lifecycle stage that models all possible stages through which personal data moves during its lifecycle in more repetitive and circular flows. The LifecycleStage is abstractly represented as a general classifier that can be used as a classification of all possible stages of the lifecycle. It is mainly used as a target of generalisations, which can be specialised into eight specific classifiers according to associated activities. We consider each in turn.

Initiation represents a complete processing plan that can be referred to before and during the processing of personal data. It has the following properties: *informalDescription; specifiedPurpose; requiredData* (specifies the minimum amount of personal data as a set of relevant, adequate and not excessive personal data items in support of the specified purpose); *dataSource* (specifies the sources from which personal data items are to be collected, derived or acquired, whether these are internal or external sources); *availableChoice* (describes the choices available to data subjects with regards to the collection, usage and disclosure of their personal data); *consentType* (drawn from {Explicit, Implicit}); *collectionMethod; storageMethod; retentionTime* (specifies the necessary period for which personal data is retained to fulfil the specified purpose or as required by applicable laws and regulations); *retrievalMechanism* (e.g. query languages, command-line, browser, or graphical user interfaces); *disclosureMechanism* (specifies the means and the manner in which personal data is to be disseminated, transmitted or made available); *destructionMethod; applicableRegulations* (indicates the applicable laws and regulations); and *relevantPrinciples* (specifies the relevant GPS principles that govern the stage of lifecycle in the sense of placing limitations or constraints on the associated activities).

Collection represents the act of creating personal data values, whether these are directly recorded, captured or collected from data subjects, or have been acquired

from external sources. It has the properties *informalDescription, createdData, dataSource, collectionMethod, availableChoice, consentType,* and *relevantPrinciples.*

Retention represents the act of organising, structuring, storing and retaining personal data values in repositories or digital storage media for operational, compliance or operational recovery purposes. It has the properties *informalDescription, retainedData, activityType* (drawn from {PrimaryStorage, Archiving, Backup}), *activityPurpose* (drawn from {Operational, RegulatoryCompliance, FutureReference, OperationalRecovery}), *retentionTime, storageMethod,* and *relevantPrinciples.*

Access represents the act of specifying, and retrieving or consulting personal data values that are stored in repositories or digital storage media. It has the following properties: *informalDescription; retrievedData; retrievalMechanism;* and *relevantPrinciples.*

Participation represents the act of implementing access rights and rectifying personal data values by data subjects to ensure that their data is accurate, complete and up-to-date. It has the following properties: *informalDescription; reviewedData; activityType* (drawn from {Retrieval, Alteration, Alignment}); *activityPurpose* (drawn from {Review, Update, Correction}); and *relevantPrinciples.*

Disclosure represents the act of disseminating, making available or transmitting personal data for external use by third parties. It has the following properties: *informalDescription; disclosedData; activityType* (drawn from {InitialProcessing, FurtherProcessing}); *activityPurpose* (drawn from {TheSpecifiedPurposes, HistoricalPurposes, ScientificPurposes, StatisticalPurposes}); *disclosureMechanism;* and *relevantPrinciples.*

Usage represents the act of using, altering, adapting, refining, aligning or combining personal data items. It has the following properties: *informalDescription; usedData; activityType* (drawn from {Alteration, Alignment, Consultation}); *activityPurpose* (drawn from {Use, Derivation}); and *relevantPrinciples.*

Destruction represents the act of erasing, destroying, redacting or disposing of personal data. It has the following properties: *informalDescription; de-*

stroyedData; destructionMethod; competentAuthority (drawn from {InternalUnit, ExternalDataDestructionService}); and *relevantPrinciples.*

Stage Activities, Events and Actions

StageActivity represents data-processing activities (concrete purposes) that constitute the operations performing on personal data in each stage of the lifecycle. It has the following properties: *informalDescription; input; output; preCondition;* and *postCondition.*

StageEvent represents occurrences that may happen at specific points in time that may have consequences for personal data. It has two properties: *informalDescription* and *category* (*implicit* events occur on the change of states or the passage of some interval of time; *explicit* events occur when an operation is directly requested).

StageAction represents single execution steps within an activity. Actions are the fundamental units that describe personal data processing activities in a fine-grained manner. It has the following properties: *informalDescription; inputParameter; outputParameter; localPreCondition;* and *localPostCondition.*

Lifecycle Roles

LifecycleRole represents the way in which a concerned actor participates in a set of related activities of the personal data lifecycle. As such, a role represents a set of responsibilities that are logically-related to each other, either by their objectives or by the actors that may play the role. It has the following properties: *informalDescription; roleType* (drawn from {DataModeller, DataSubject, DataController, DataProcessor, ThirdParty}); and *responsibility.* Lifecycle roles are identified to cover all stages of the data lifecycle. The processing of personal data in various stages identifies actors in these roles.

LifecycleActor represents an external or internal entity that is capable of, and responsible for, performing the activities of the role to which it is assigned. It has the following properties: *informalDescription; actorNature* (drawn from {HumanActor, SoftwareAgent}); and *responsibility.*

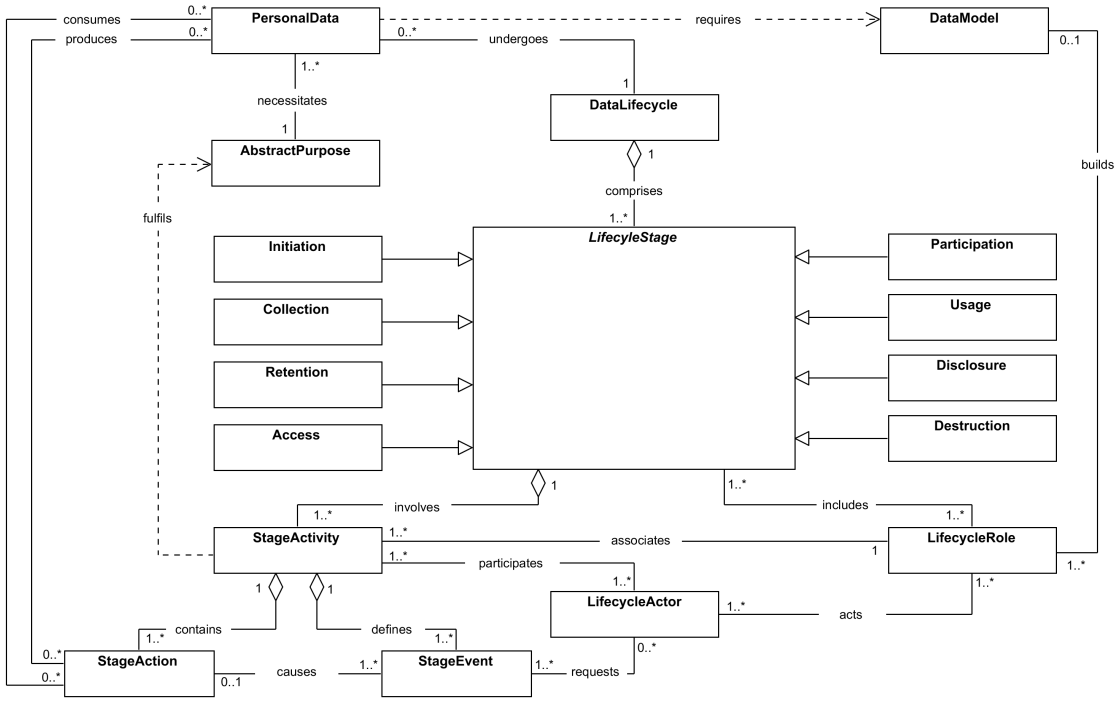

Figure 4.4: The meta-model of the APDL profile.

Figure 4.4 shows the meta-model of the APDL profile with minimal syntax, omitting the attributes of the classes for simplicity and readability.

4.3 A UML Profile for the APDL Model

To model the abstract and concrete purposes, together with the key aspects of abstract privacy principles as a requirements model, it is important to define a UML profile for the APDL model to serve as a means for representation. In order to define such a profile, we map the conceptual model of Section 4.2.2 to the UML profile. The stereotypes of the APDL profile are defined to extend existing metaclasses with the aim of using privacy-related terminology whether in place of, or in addition to, the terminology used for the extended metaclasses. The abstract syntax of the APDL is specified by extending three elements of the UML metamodel — the metaclass *Class*, the metaclass *Association* and the metaclass *Dependency* — with additional properties and constraints. The name of the applied stereotypes are shown within a pair of guillemets. The UML profile defines the concepts needed to model personal data-processing activities using UML 2.5 [102]. The constraints needed to express

Ref.	Stereotype	Base Class	Parent
C.01	«AbstractPurpose»	Class	—
C.02	«PersonalData»	Class	—
C.03	«DataModel»	Class	—
C.04	«DataLifecycle»	Class	—
C.05	«LifecycleStage»	Class	—
C.06	«Initiation»	Class	«LifecycleStage»
C.07	«Collection»	Class	«LifecycleStage»
C.08	«Retention»	Class	«LifecycleStage»
C.09	«Access»	Class	«LifecycleStage»
C.10	«Usage»	Class	«LifecycleStage»
C.11	«Disclosure»	Class	«LifecycleStage»
C.12	«Participation»	Class	«LifecycleStage»
C.13	«Destruction»	Class	«LifecycleStage»
C.14	«StageActivity»	Class	—
C.15	«StageEvent»	Class	—
C.16	«StageAction»	Class	—
C.17	«LifecycleRole»	Class	—
C.18	«LifecycleActor»	Class	—

Table 4.1: The APDL Profile: stereotyped classes.

privacy-related concepts of the APDL model are limited to association multiplicities, pre- and post-conditions of stage activities and actions unless additional constraints are explicitly stated. The stereotyped classes and associations belonging to the APDL profile are listed in Table 4.1 and Table 4.2 respectively.

We distinguish between three levels of abstraction: meta, domain and instance levels. The focus is on meta and domain levels as the instance level is specific to the application domain. At the meta level, we refer to domain-independent abstractions, such as stereotypes, tag definitions, relationships and constraints in relation to the conceptual model. At the domain level (i.e. when a stereotype is applied to a model element), we refer to stereotyped classes, tagged values, domain-specific attributes, relationships and constraints specific to the application domain.

Ref.	Stereotype	Base Class	Relates
A.01	«Necessitate»	Association	C.01 → C.02
A.02	«Require»	Dependency	C.02 → C.03
A.03	«Undergo»	Association	C.02 → C.04
A.04	«Comprise»	Association	C.04 → C.05
A.05	«Involve»	Association	C.05 → C.14
A.06	«Define»	Association	C.14 → C.15
A.07	«Contain»	Association	C.14 → C.16
A.08	«Fulfil»	Abstraction	C.14 → C.01
A.09	«Produce»	Association	C.16 → C.02
A.10	«Consume»	Association	C.16 → C.02
A.11	«Cause»	Association	C.15 → C.16
A.12	«Include»	Association	C.05 → C.17
A.13	«Associate»	Association	C.14 → C.17
A.14	«Build»	Association	C.17 → C.03
A.15	«Act»	Association	C.18 → C.17
A.16	«Participate»	Association	C.18 → C.14
A.17	«Request»	Association	C.18 → C.15

Table 4.2: The APDL Profile: stereotyped associations.

AbstractPurpose, PersonalData and DataModel

Stereotypes. The abstract purpose for which personal data is processed can be specified using the «AbstractPurpose» stereotype, which constrains the semantics of the objects, meaning that only they can be used as abstract purposes. The primary tag definitions of the AbstractPurpose stereotype are *informalDescription*, *isFair*, *isLawful*, *isProportional* and *relevantPrinciple*, as illustrated in Figure 4.5. When the AbstractPurpose stereotype is applied to any class, its primary attributes may include *actualPurpose*. Some aspects are more challenging to model, such as the fairness, lawfulness and proportionality of the specified purpose. These can be represented as *Boolean* tagged values to be specified by competent or authorised actors. The primary tag definitions of the «PersonalData» stereotype are *informalDescription, category, type* and *linkability*, as illustrated in Figure 4.5. When the PersonalData stereotype is applied to any class, its primary attributes

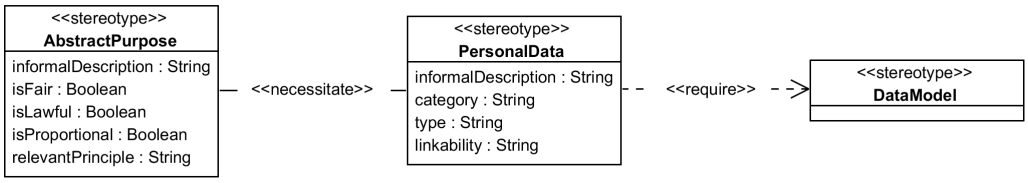

Figure 4.5: AbstractPurpose, PersonalData and DataModel stereotypes at the meta level.

are specific to the application domain.

Relationships. The «necessitate» association denotes a relationship between «AbstractPurpose» and «PersonalData». It is used to specify a minimum amount of data required to fulfil the specified purpose, as illustrated in Figure 4.5. Here, «require» is a usage dependency that denotes a relationship between «PersonalData» and «DataModel». It specifies that the required data that fulfils the specified purpose requires a data model for its full specification, as per Figure 4.5.

Constraints. Each data-processing initiative has an abstract purpose that can be concretely refined to a set of purposes that can be specified in terms of data-processing activities, which, in turn, can be hierarchically structured as concrete actions and events. Consequently, it should not make sense to have multiple instances of the class stereotyped by «AbstractPurpose». As such, it is constrained as a singleton — it is instantiated only once in each particular model.

DataLifecycle, LifecycleStage and its Specialisations

Stereotypes. The stages through which personal data moves during its lifetime can be specified using the «DataLifecycle» stereotype. The primary tag definitions of the DataLifecycle stereotype is: *informalDescription, isOpen* and *isCentralised*, as illustrated in Figure 4.6. When the DataLifecycle stereotype is applied to any class, its primary attributes are specific to the application domain. The DataLifecycle stereotype consists of one or more stages that involve data-processing activities. These stages are represented by the abstract «LifecycleStage» stereotype. The primary tag definitions of its specialisations are illustrated in Figure 4.6.

Relationships. The «undergo» association denotes a relationship between «PersonalData» and «DataLifecycle». It is used to specify that personal data is

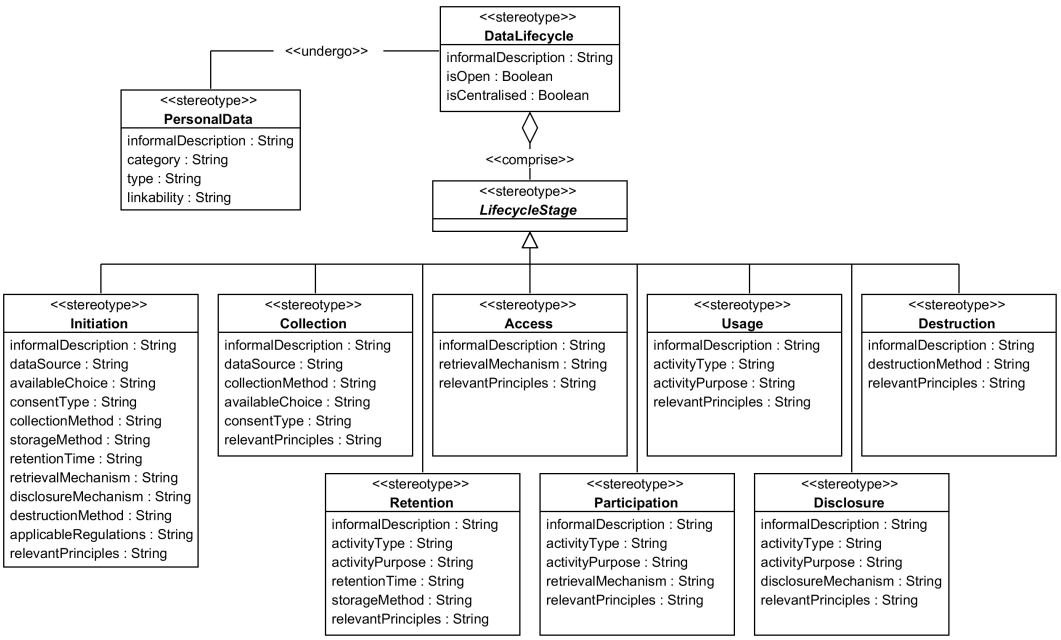

Figure 4.6: DataLifecycle and LifecycleStage stereotypes at the meta-level.

subject to a set of stages, each of which involves a set of processing activities. The «comprise» association denotes a relationship between «DataLifecycle» and «LifecycleStage». It is used to specify that the data lifecycle consists of various stages that involve distinct but related processing activities.

Constraints. «LifecycleStage» represents the concept of a generic lifecycle stage. It is mainly used as a target of generalisation; as such, it is constrained as abstract. The generalisation set, which combines all the special classifiers of the LifecycleStage, has two properties: *complete* and *disjoint*. Each data-processing initiative requires the development of a complete processing plan that may serve as the basis of establishing a privacy notice to be communicated to data subjects. Consequently, it should not make sense to have multiple instances of «Initiation». As such, it is constrained as a singleton.

StageActivity, StageEvent and StageAction

Each lifecycle stage involves a set of activities, which, in turn, consist of a set of concrete actions, as well as a set of events that trigger their execution.

Stereotypes. A data-processing activity is specified using the «StageActivity»

stereotype, which constrains the semantics of the objects, meaning that only they can be used as processing activities. The primary tag definitions of the StageActivity stereotype are *informalDescription, preCondition* and *postCondition*, as illustrated in Figure 4.7. When the StageActivity stereotype is applied to any class, its primary attributes may include *input* and *output*, and the tagged values of preCondition and postCondition can be written as constraints in OCL. Each StageActivity contains a set of actions specified by the «StageAction» stereotype. The primary tag definitions of the StageAction stereotype are *informalDescription, localPreCondition* and *localPostCondition*, as illustrated in Figure 4.7. When the StageAction stereotype is applied to any class, its primary attributes may include *inputParameter* and *outputParameter*, and the tagged values of localPreCondition and localPostCondition can be written as constraints in OCL. Each StageActivity defines a set of events specified using the «StageEvent» stereotype. The primary tag definitions of the StageEvent stereotype are *informalDescription* and *category*, as illustrated in Figure 4.7. When the StageEvent stereotype is applied to any class, its primary attributes are specific to the application domain.

Relationships. Figure 4.7 shows how the LifecycleStage, StageActivity, StageAction and StageEvent stereotypes participate in relationships.

The «involve» association denotes a relationship between «LifecycleStage» and «StageActivity». Each lifecycle stage may involve one or more related but distinct activities. The «contain» association denotes a relationship between «StageActivity» and «StageAction». Each stage activity may consist of one or more concrete actions. The «define» association denotes a relationship between «StageActivity» and «StageEvent». Each stage activity may consist of one or more events that are requested by actors to trigger its execution. The «cause» association denotes a relationship between «StageEvent» and «StageAction». Each action may be triggered by zero or more events that are requested by actors. The «consume» association denotes a relationship between «StageAction» and «PersonalData». The action requires specific personal data items to accomplish its execution. The «produce» association denotes a relationship between «StageAction» and

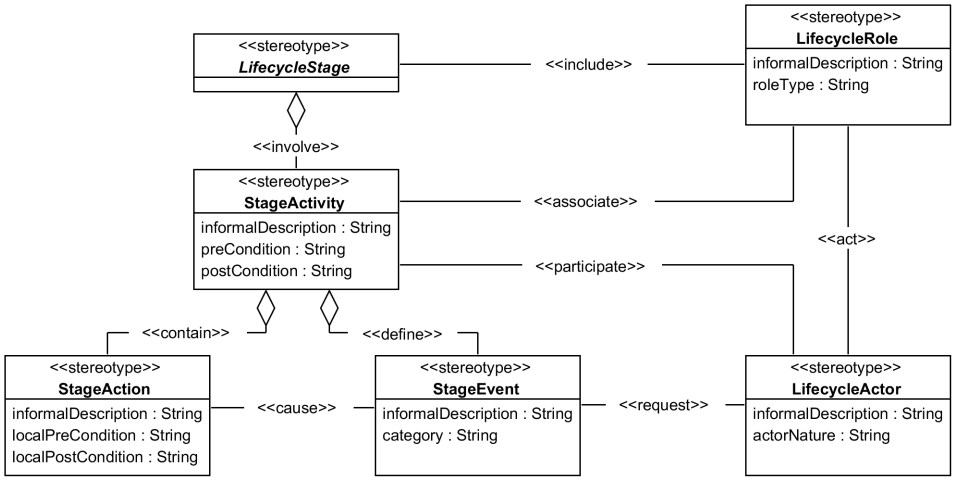

Figure 4.7: StageActivity, StageAction, StageEvent, LifecycleRole and LifecycleActor stereotypes at the meta level.

«PersonalData». The action provides specific personal data items as a result of its execution. The «fulfil» association is an abstraction that donates a refinement relationship between «StageActivity» and «AbstractPurpose». It is a specialisation of the standard abstraction stereotype «Refine». It is used to specify data-processing activities that have already been specified at a certain level of detail as a purpose. It is used to capture how a processing activity participates in the fulfilment of the specified purpose.

Constraints. The constraint of the «consume» stereotype is that values of the association's *inputParameter* property must be attributes of the «PersonalData» stereotyped class. Similarly, the constraint of the «produce» stereotype is that values of the association's *outputParameter* property must be attributes of the «PersonalData» stereotyped class.

LifecycleRole and LifecycleActor

Each lifecycle stage involves a set of roles that may be played by different actors.

Stereotypes. A set of related activities that are expected to be performed together can be represented using the «LifecycleRole» stereotype. The primary tag definitions of the LifecycleRole stereotype are *informalDescription* and *roleType*, as illustrated in Figure 4.7. When the LifecycleRole stereotype is applied to any class, its primary attributes are specific to the application domain. An external

or internal entity that is capable of, and responsible for, performing a set of activities associated with the role to which it is assigned can be specified using the «LifecycleActor» stereotype. The primary tag definitions of the LifecycleActor stereotype are *informalDescription* and *actorNature*, as illustrated in Figure 4.7. When the LifecycleActor stereotype is applied to any class, its primary attributes are specific to the application domain.

Relationships. Figure 4.7 shows how the LifecycleRole and LifecycleActor stereotypes participate in relationships.

The «include» association denotes a relationship between «LifecycleStage» and «LifecycleRole». Each lifecycle stage may include one or more lifecycle roles that participate to accomplish associated activities, and each lifecycle role may participate in one or more lifecycle stages. The «associate» association denotes a relationship between «StageActivity» and «LifecycleRole». Each stage activity may be assigned to exactly one lifecycle role; further, each lifecycle role may involve one or more stage activities. The «act» association denotes a relationship between «LifecycleActor» and «LifecycleRole». Each lifecycle actor may be assigned to one or more lifecycle roles; further, each lifecycle role may involve one or more lifecycle actors. The «request» association denotes a relationship between «LifecycleActor» and «StageEvent». Each lifecycle actor may perform an action by requesting one or more stage events that trigger its execution, and each stage event may be requested by zero or more lifecycle actors. The «participate» association denotes a relationship between «LifecycleActor» and «StageActivity». Each lifecycle actor is capable of performing one or more stage activities that are assigned to one or more lifecycle roles to which the actor is assigned, and each stage activity may involve one or more lifecycle actors. The «build» association denotes a relationship between «LifecycleRole» and «DataModel». Each data modeller may construct zero or one data models, and each data model may be constructed by one or more data modellers.

Example 4.2. Figure 4.8 shows how to represent *Activity3* of Example 4.1 using the UML profile for the APDL model. In this example, we illustrate only *Constraint1* and *Constraint2* that informally specify conditions to be satisfied before, or to

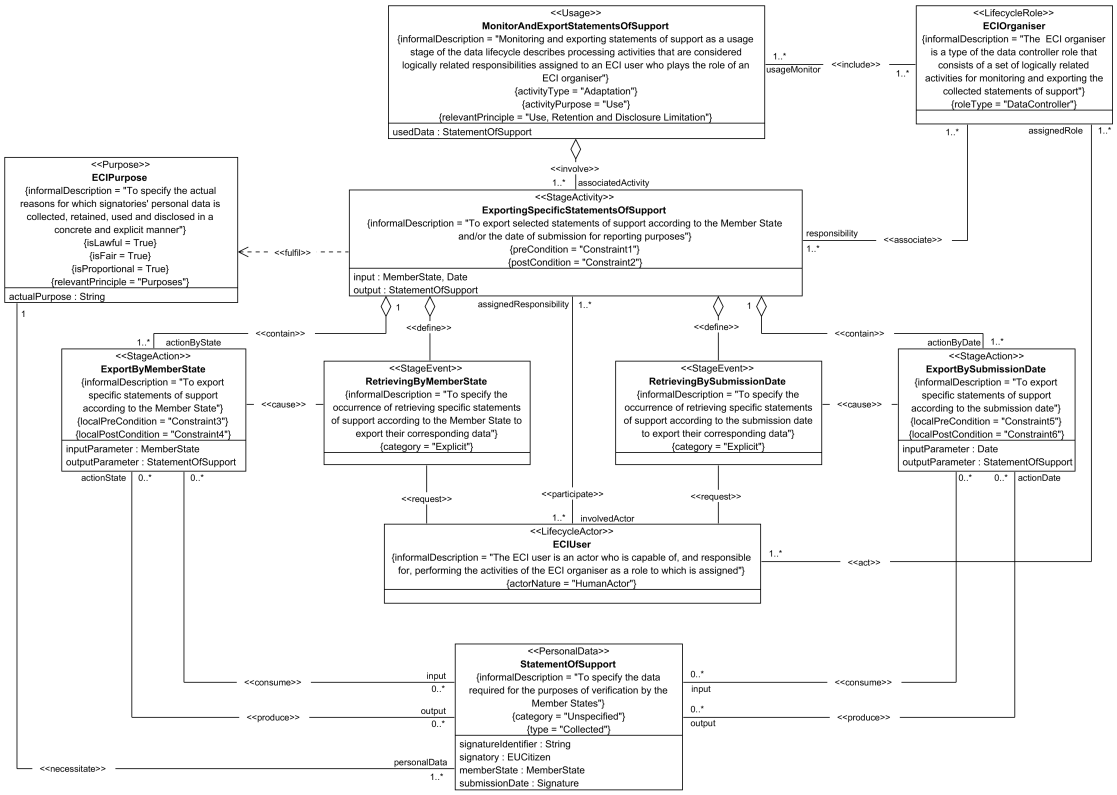

Figure 4.8: The representation of Activity3 of Example 4.1.

be guaranteed after, the performance of *Activity3*. The pre-condition *Constraint1* is to specify that signatories have been informed about the processing of their personal data via an appropriate privacy notice. The post-condition *Constraint2* is to specify that the corresponding data of the selected statements of support has been successfully retrieved with the consent and knowledge of signatories.

By using the UML profile for the APDL model as a means for representation, it is important to mention that all data-processing activities need to be represented in one diagram to illustrate how each activity participates in the fulfilment of the abstract purpose. In this dissertation, however, we represent each activity separately for readability purposes.

4.4 Summary

In this chapter, we have developed the APDL model as an abstract model for personal data lifecycle models. The APDL model represents the data lifecycle as

a sequence of activities performed on personal data by different actors (according to their assigned roles) using various processing methods. These activities are performed in an ordered manner to indicate how and when to move from one stage to another. We have also developed a conceptual model for the APDL, upon which we have defined the UML profile.

The UML profile for the APDL model represents privacy-related concepts using the standard extension mechanisms of the UML meta-model. Stereotypes and tagged values are used to represent data-processing activities and key aspects of abstract privacy principles as requirements, and constraints provide criteria for the evaluation of these aspects to determine whether the representation fulfils these requirements. Crucially, the UML profile represents data-processing at two levels of granularity to indicate whether personal data is completely or partially processed. They are classified as *coarse*, which means that all personal data items are processed in each cycle, or *fine*, which means that only a number are processed in each cycle. The level of granularity helps support the application of the principle of data minimisation as a foundational step for privacy engineering by specifying the processed data items in each single atomic action within an activity. This, in turn, helps analyse and restrict the processing of personal data to the minimum amount necessary according to the purpose of each concrete action. As such, the processing cycle provides a view on how the data flows through the lifecycle stages and is ultimately used in support of the specified purposes.

In the next chapter, we will illustrate the second activity of our principled approach for engineering PbD: data-centric threat modelling. In addition, we will explain how data-processing representation (using the APDL model and its UML profile) serves as a foundation for identifying, analysing and assessing potential privacy risks.

Risk is a key aspect of systems engineering. Privacy engineering, therefore, requires a sufficiently robust approach to risk.

— Ann Cavoukian

5 Data-Centric Threat Modelling

Contents

5.1 Introduction

This chapter gives a detailed description of a data-centric threat modelling approach, which contributes to the second activity of our principled approach for engineering PbD. Data-centric threat modelling is focused on protecting specific types of data, which, in the context of this dissertation, is personal data. Section 5.2 describes

a privacy risk model that goes beyond traditional security risk models to take into consideration the dynamic and contextual nature of privacy. In particular, it considers legal, organisational, societal and technical aspects. Section 5.3 illustrates a method for identifying and analysing potential privacy risks. It also describes how combinations of risk factors are identified and analysed to ensure adequate coverage of the problem space at a consistent level of detail. Then, Section 5.4 illustrates a method for assessing privacy risks. It also describes the dependencies between the nominal and assessable attributes of key risk factors, and establishes a set of assessment rules that specify the range of values the risk factors can assume during the risk assessment. Finally, Section 5.5 gives a brief summary of the chapter. The work presented in this chapter previously appeared in [103, 104].

5.2 A Privacy Risk Model

Typically, a *risk model* defines key terms and assessable risk factors, and the conceptual relationships among these factors. In general, risk factors are distinguishing characteristics that can be used as inputs, with the aim of determining risk levels while conducting risk assessments. In security risk models, typical risk factors are threats, vulnerabilities, impact and likelihood. Often, risk factors can be decomposed into more detailed characteristics; for example, threats can be decomposed into threat sources and threat events. These levels of detail are important as risk assessments rely heavily upon well-defined attributes (nominal and assessable) of these factors to effectively determine risk levels. There is no denying that different models can lead to different levels of detail in characterising key risk factors.

As previously mentioned in Section 2.2.2, security risk assessments consider confidentiality, integrity and availability as security protection goals, whereas privacy risk assessments need to go beyond these goals to consider anonymity, pseudonymity, unobservability, undetectability, unlinkability, intervenability and transparency. Even though security goals need to be considered while assessing privacy risks, decisions about the collection, processing and dissemination of personal data need to be made from a privacy and data protection perspective as security goals may

conflict with privacy goals, which, in turn, require conflict resolution techniques and contextual analysis. As such, privacy risk assessments need to go beyond identifying technical risks of the system being developed; however, this requires a better understanding of social perceptions and expectations that are derived from social norms [14, 105]. To some extent, the nature of privacy impact differs from security impact, as the potential impact of privacy violations might be incorporeal, psychological, or emotional. This means that the negative consequences of privacy violations are not only related to the affected data subjects but also may extend to affect society [26]. In addition, security impact is often measured from a financial perspective, whereas privacy impacts measured from two different perspectives: (1) as a financial impact, whether this impact is tangible, such as legal sanctions, or intangible, such as an entity's reputation; and (2) as personal and societal impacts, such as social standing and an individual's reputation [16].

To carry out an appropriate privacy risk assessment that goes beyond traditional security assessment, we define and/or refine the key risk factors, along with the relationships among these factors, used in conducting security risk assessments to be appropriately applied in the context of privacy and data protection. We refer to fundamentals from the legal privacy literature to underpin the key terms and risk factors along with their meanings, properties and relationships. In particular, we refer to *Calo's scheme* [28] to understand the specific characteristics and categories of privacy harms. In addition, we refer to *Solove's taxonomy* [26] to understand the specific characteristics of adverse privacy events and associated categories. Finally, we refer to *Nissenbaum's contextual integrity heuristic* [29] to understand the main characteristics of appropriate flow of personal data with reference to context-relative informational norms, from which vulnerabilities can be derived. In this dissertation, we consider only risks that arise from the collection, processing and dissemination of personal data that have adverse impacts on the privacy of data subjects.

5.2.1 Privacy Risk Factors

Threats

A *threat* is an event with the potential for a privacy violation, or to adversely impact the privacy of data subjects through inappropriate collection, processing and/or dissemination of personal data. In our risk model, the threat concept is abstractly represented: it can be decomposed into a threat source and a threat event.

A *threat source* is an entity with capability to collect, process and/or disseminate (lawfully or unlawfully, fairly or unfairly) data belonging to data subjects 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 (representing the way in which a concerned entity participates in processing operations, such as normal user, privileged user, service provider, service consumer, etc.), and responsibility.

The specified attributes of a threat source are used to assess the capability of exploiting possible privacy vulnerabilities. As such, a threat source is more relevant to vulnerability analysis than impact assessment, i.e. impact is independent of vulnerability and threat analysis — in practice it is irrelevant whether the threat event flows from an internal or external threat source whose actions are accidental or deliberate. We use the concept of a threat source to ensure that it can be used appropriately for modelling actors with malicious and benign purposes.

Example 5.1. In the context of ECI, *organisers* may act as threat sources as they are able to lawfully collect and process *signatories'* personal data. They are: individuals as they are members of a committee that consists of at least seven EU citizens; and insiders as they are responsible for preparing and managing the initiative. Often, initiatives concern topics that are socially and politically controversial, such as abortion, immigration law, issues related to religious freedom,

etc. The utility of signatories' personal data makes such data highly valuable to organisers. The motives behind collecting, processing and/or disseminated such data (unfairly and/or unlawfully) stem from its value, including profiling, categorising, discriminating or stigmatising signatories based on signatories' ideas. With regards to organisers' resources, they have real and sufficient skills, privileges and technical resources to exploit possible vulnerabilities as they develop conceptual, logical and physical models of the data, and maintain access control models. They act in their capacity as data controllers and they are responsible for collecting, exporting, deleting, monitoring and confirming statements of support.

A *threat event* is a data-driven event that may happen at specific points in time and which leads to a privacy violation or has an effect, consequence or impact, especially a negative one, on the privacy of data subjects. Each threat event involves an adverse action justified by reference to personal data — i.e. *what can go wrong*. 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 privacy vulnerabilities by one or more threat sources. Each type of threat event can be characterised by: category (collection, processing or dissemination), nature (continuous or discrete; excessive or reasonable; anticipated or unanticipated), adverse effects and scope (individuals, a specific group of individuals or whole society).

For the purpose of this dissertation, we consider only data-driven threats that are processing-related, not those caused by natural disasters, power failures, etc. 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.

Example 5.2. In the context of ECI, *unauthorised secondary use* is a threat event that may occur when organisers use signatories' personal data for purposes unrelated to the purposes for which it was initially collected without the knowledge and consent of signatories. With regards to the threat event category, the threat event is related

to processing adverse events. With regards to the scope of this threat event, it affects a specific group of individuals (those who have support a certain initiative). With respect to its nature, the threat event is discrete (as it happens in a series of discrete steps), excessive (as it uses signatories' profiles) and unanticipated (it may be stealthy and invisible to signatories). With regards to its adverse effects, revealing sensitive facts about signatories is one of those undesirable consequences.

Privacy Vulnerabilities

A *privacy vulnerability* is a weakness or deficiency in: data modelling; processing operations specification or implementation; or privacy controls, whether these controls are technical, organisational or legal, that makes an exploitation of an asset more likely to succeed by one or more threat sources. Successful exploitations of privacy vulnerabilities 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*. The former refers to personal data that is directly concerned with processing operations, as well as principles stated in legal frameworks and standards. The latter refers to system components on which the primary assets rely, such as hardware, software, people, etc. As previously mentioned in Section 2.2.2, personal data is governed by laws and regulations, by social norms, by economics, or by policies or contracts. Thus, it needs to be distinguished from other types of data within systems. In addition, adverse actions of threat sources on primary assets and supporting assets happen through different threats with different characteristics. Indeed, there are several types of threat modelling, for example, system threat modelling for operational systems. To consider those exploitations of vulnerabilities that lead to data-driven events that involve adverse actions justified by reference to personal data, we focus on *data-centric system threat modelling* to model relevant aspects of exploits and defence sides of particular logical entities — i.e. primary assets — instead of physical entities hardware, networks, etc. — i.e. supporting assets. As such, in this dissertation, we consider only the primary assets and associated vulnerabilities — i.e. *what we are trying to protect*. Each type of vulnerability

can be characterised by exploitability (how hard is it to exploit the vulnerability?) and severity (what does the vulnerability enable a threat source to do?), which are used to estimate the seriousness of a vulnerability.

We use the concept of privacy vulnerability with a broader view to not associate them only within data protection mechanisms: privacy vulnerabilities can be found in the implemented privacy controls and the specified processing operations along with the required personal data.

Example 5.3. In the context of ECI, *an improper data model* is a privacy vulnerability that might be exploited by *organisers*: it is a weakness in data modelling. With regards to its exploitability, it can be exploited by organisers as they develop conceptual, logical and physical models of the data. With respect to its severity, the successful exploitation of this vulnerability leads to serious impacts as it facilitates inadequate, irrelevant and excessive collection of signatories' personal data, which is not necessary for the main purpose.

Privacy Violations

A *privacy violation* is an unfair and/or unlawful action that accidentally or deliberately breaches privacy-related legal frameworks, principles, unilateral policies, contracts or social norms without an adverse action taken against data subjects, as well as without their knowledge and consent. Such an action is triggered by an occurrence of a threat event that results from the successful exploitation of one or more privacy vulnerabilities. In reality, inappropriate collection, processing and dissemination of personal data may lead to privacy violations, which may involve a variety of types of actions that may not lead to privacy harms. In our risk model, the presence of a privacy violation does not mean that it will necessarily create actual privacy harm. Further, a privacy harm can occur without a privacy violation. Each type of privacy violation can be characterised by: type (unlawful or unfair), degree of breach (excessive or limited) and scope (individuals, a specific group of individuals or whole society).

Example 5.4. In the context of ECI, *passive collection of sensitive data* is a privacy violation that may result from the occurrence of the threat event *excessive data collection*, which results from the successful exploitation of *an improper data model* by *organisers*. Its type is unlawful as the mandatory fields that are required are specified by each Member State according to relevant national regulations. In addition, it is unfair as it is unexpected by signatories to collect such type of data. Its degree is excessive as it collects inadequate and irrelevant signatories' personal data, which is not necessary for the main purpose. Its scope is individuals — i.e. those who are supporting a particular initiative.

Privacy Harms

A *privacy harm* is the adverse impact or consequence (incorporeal, financial or physical) of inappropriate collection, processing and dissemination of personal data on the privacy of data subjects, a specific group of data subjects or 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. However, in our risk model, privacy harms are related to, but distinct from, privacy violations. In the legal privacy literature, privacy harms are classified into two related categories: subjective and objective. To distinguish between these categories, the former represents the perception of undesirable collection, processing and dissemination of personal data that results in unwelcome mental states, such as anxiety, embarrassment or fear, whereas the latter represents the unanticipated or forced collection, processing and dissemination of personal data and the use of this data against the data subject, such as identity theft. In this dissertation, the focus is on the objective category of privacy harms (as they are the potential adverse consequences that result from adverse actions justified by reference to personal data).

Each privacy harm can be characterised by: category (incorporeal, financial or physical), extent of damage and affected data subjects (data subjects, a specific group of data subjects, or whole society).

Example 5.5. In the context of ECI, *denial of a job* is a privacy harm that may result from the occurrence of the threat events *excessive data inference* and *unauthorised data disclosure*, which result from the successful exploitation of *an improper data model* by *organisers*. It is a type of employment discrimination that may occur as organisers can make inferences to create signatories' profiles and share these profiles with employment agencies who may make data inference to re-identify job applicants with the aim of filtering those job candidates according to views or beliefs that are derived from their profiles. As previously mentioned, initiatives are about issues that are socially and politically controversial, such as issues related to immigration law, religious freedom, etc., which may reveal sensitive information about signatories. The adverse effects of associated threat events are: gathering identifiable data about signatories; storing identifiable records for signatories, from which sophisticated profiles can be derived; and sharing profiles for signatories beyond expected boundaries that may be analysed to derive sensitive information that can inhibit certain rational judgements. The duration of these effects may last for specific period of time (in this case, the recruitment period). The affected data subjects are those who support a specific initiative.

5.2.2 Relationships Between the Privacy Risk Factors

Figure 5.1 shows the conceptual relationships between the key risk factors. According to these relationships, a threat source may exploit zero or more privacy vulnerabilities, and a privacy vulnerability may be exploited by zero or more threat sources. The successful exploitation of a privacy vulnerability by a threat source leads to a threat event. Each threat event may lead to one or more privacy harms, and may constitute one or more privacy violations. Each privacy violation is constituted by one or more threat events, whereas a privacy harm is created by zero or more threat events.

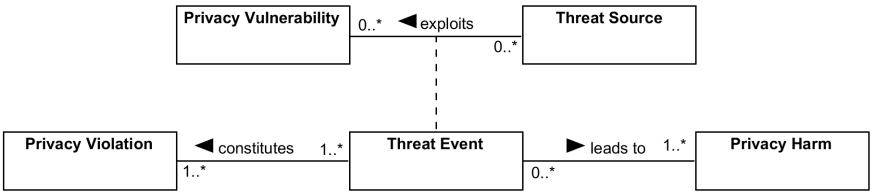

Figure 5.1: The conceptual relationships among the key risk factors.

5.3 A Privacy Risk Analysis

In this section, we describe the main steps of our analysis approach, which aims to identify and analyse potential privacy risks. The approach is built upon the privacy risk model of Section 5.2 that defines the main factors that have impacts on privacy risks along with their meanings, properties and relationships. It describes how combinations of risk factors are identified and analysed to ensure adequate coverage of the problem space at a consistent level of detail.

Typically, analysis approaches differ with respect to the starting points of risk assessments, the level of abstraction and complexity with which threat events are identified. In order for risk assessments to be effective, they need to synthesise multiple analysis approaches (threat-oriented, asset/impact-oriented, or vulnerability-oriented) to identify the key factors of risk. Potential privacy risks need to be identified, analysed and assessed in a systematic manner. As such, our analytical approach consists of four steps, as per Figure 5.2. The first step is to establish the context in which personal data is collected, processed and/or disseminated. The second step is to identify and analyse all possible privacy vulnerabilities in this particular context. The third step is to identify and analyse potential threats (in terms of threat sources and events). The fourth step is to identify and analyse privacy violations and harms. We consider each in turn.

5.3.1 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 relevant information for privacy risk analysis and assessment. Context establishment aims

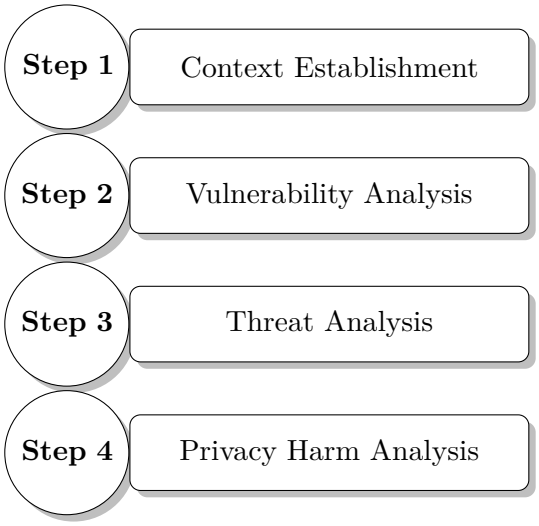

Figure 5.2: The main steps of the analysis approach.

to model primary assets in a contextual manner, representing assets and their relationships, and model the context in which actors participate in the performance of data-processing activities. This includes the types of personal data to be processed, along with their sources; the purposes for, and the manner in which, these types are collected, processed and/or disseminated; involved actors and their assigned roles and responsibilities; relevant legal frameworks and standards; and domain-specific constraints.

As mentioned in Section 5.2.1, primary assets are classified into *personal data*, which relates to identified or identifiable individuals, and *processing operations*, which refer to both data-processing activities and privacy-related principles stated in legal frameworks and standards. As such, personal data, data-processing activities and involved actors need to be represented in a way that is amenable to analysis. While describing systems in multiple views is important, we emphasise the importance of data-management models that represent data and associated processing operations at a detailed level of abstraction and in a comprehensive manner. We believe that data lifecycle models are better at describing data-processing activities at a detailed level of abstraction — they categorise and represent these activities in relation to the main stages of the lifecycle: from collection to destruction.

The APDL model of Chapter 4 was developed to represent data-processing activities in a way that is amenable to risk 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.

To establish the context, we adopt the APDL model as a means of representing the primary assets, along with associated processing operations and involved actors. In order to describe the context in a widely-used modelling notation, we use the UML profile for the APDL model to represent primary assets in terms of the UML. Personal data is represented by developing a *DataModel*, which represents the relevant objects, associated properties, relationships and constraints for the purpose of specifying the minimum amount of required personal data. Processing operations are abstractly represented in eight stages: *Initiation, Collection, Retention, Access, Participation, Usage, Disclosure* and *Destruction* stages. In each stage, data-processing activities are concretely represented in *StageActivity, StageEvent* and *StageAction*, and principles stated in legal frameworks and standards are modelled as tagged values and constraints. In addition, involved actors and the way in which they participate in processing activities are represented as *LifecycleRole* and *LifecycleActor*.

Example 5.6. In the context of ECI, *the statement of support form* specifies the mandatory fields that are required to sign up to an initiative. As per Example 4.1, these data items are subject to five processing activities by different actors according to their roles and responsibilities.

1. Collecting statements of support
2. Exporting all statements of support
3. Exporting specific statements of support
4. Deleting specific statement of support
5. Confirming valid statements of support

Figure 4.8 of Example 4.2 illustrates how 'Exporting specific statements of support', as a data-processing activity, is represented in a way that is amenable to risk analysis. By representing all the five processing activities, the context, in which signatories' personal data is collected and processed, is established.

5.3.2 Vulnerability Analysis

The *first step* in vulnerability analysis is to define a *baseline model* for all data-processing activities to understand reasonable expectations of privacy in each particular context. As such, we adopt the concept of contextual integrity, which was developed from social and philosophical theories 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 data-processing activities as an element to consider all possible operations performed on personal data in a broader view. In so doing, contextual integrity is about the appropriate collection, processing and dissemination of personal data. As such, we define *context-relative processing norms* as a means for describing data-processing activities according to reasonable expectations of privacy in each particular context as follows.

> In a context, the *collection, processing and/or dissemination of personal data* of a certain type (attributes) about *data subjects* (acting in particular capacity/roles) by *actors* (acting in particular capacity/roles) is governed by particular *processing* principles.

In order to comprehensively identify and analyse all possible vulnerabilities of the primary assets, a baseline model, which describes personal data items, associated data-processing activities, involved actors, along with their assigned roles, and privacy principles stated in legal frameworks and standards, needs to be represented in a way that is amenable to analysis. As such, the baseline model of data processing can be described in terms of context-relative processing norms. It is a set of processing norms classified according the stages of the APDL model. Implicitly,

establishing context-relative processing norms reflects the operationalisation of purposes of Section 4.2.2. Thus, we use Step 1 of Section 5.3.1 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 APDL model.

The *second step* is to derive privacy vulnerabilities from each processing norm. In particular, privacy vulnerabilities can be derived from how these norms would be breached or disrupted to violate contextual integrity. This can be achieved by separately analysing each element of the norms (data-processing activities, attributes, actors and processing principles).

- ***Element 1:*** *data-processing activities.* Any data-processing activity that does not explicitly participate in the fulfilment of the abstract purpose is considered as a breach or disruption that violates contextual integrity. From this disruption, an 'improper purpose specification' can be derived as a privacy vulnerability (a weakness in the specification or implementation of a processing operation).

- ***Element 2:*** *attributes.* Any data item that is not necessary for the fulfilment of the concrete purpose of the data-processing activity (i.e. it is not necessary to accomplish the execution of associated actions) is considered as a breach or disruption that violates contextual integrity. From this disruption, an 'improper data model' and 'improper activity specification' can be derived as privacy vulnerabilities (weaknesses in data modelling and in the specification or implementation of a processing operation).

- ***Element 3:*** *actors.* Any actor who is not capable of, or responsible for, the performance of the data-processing activity, which is associated with the role to which the actor is assigned is considered as a breach or disruption that violates contextual integrity. From this disruption, an 'improper role-based

access control model' and 'insufficient logging and monitoring' can be derived as privacy vulnerabilities (weaknesses in a privacy control).

- ***Element 4:*** *processing principles.* Any privacy principle stated in the applicable legal frameworks and standards that is not imposed on the specification of the data-processing activity via its pre- and post-conditions is considered as a breach or disruption that violates contextual integrity. From this disruption, a 'lack of data minimisation', a 'weak of anonymisation technique' and an 'improper preference specification' can be derived as privacy vulnerabilities (weaknesses in a privacy control).

Example 5.7. In the context of ECI, a *baseline model* can be described in terms of five context-relative processing norms that correspond to the data-processing activities of Example 5.6. For illustration purposes, we illustrate only one processing norm pertaining to *Activity3*.

> In the context of ECI, the usage of personal data of a certain type (the mandatory fields of the statement of support form) about signatories (acting as data subjects) by organisers (acting as data controllers) is governed by processing principles derived from applicable legal frameworks (the EU's GDPR and the Regulation (EU) No. 211/2011 on the Citizens' Initiative) and standards (relevant GPS principles).

To derive privacy vulnerabilities from the established processing norm, we separately analyse each element of the norm.

- ***Element 1:*** *data-processing activities.* The above context-relative processing norm is established with reference to Activity3, which is specified as a result of the purpose refinement of Example 4.1. Thus, Activity3 explicitly participates in the fulfilment of the abstract purpose.

- ***Element 2:*** *attributes.* The mandatory fields of 'the statement of support form' are modelled according to relevant national regulations of each Member State. They are necessary for the fulfilment of the concrete purpose of Activity3. This means that they are necessary to accomplish the execution

of associated actions. In relation to the specification of Activity3, it is not explicitly specified in a manner that prevents the linkability of personal data to particular signatories. As such, 'improper data model' can be derived as a privacy vulnerability (a weakness in data modelling).

- ***Element 3:*** *actors.* Organisers (represented by ECIUser) are actors who are capable of, and responsible for, performing Activity3 of the ECIOrganiser as a role to which they are assigned. We assume that a role-based access control model is maintained in a proper way.

- ***Element 4:*** *processing principles.* The relevant processing principle stated in the EU's GDPR, for example, is: "[...] personal data must be adequate, relevant and not excessive in relation to the purposes for which they are collected [...]". The relevant processing principles stated in the GPS are: Use, Retention and Disclosure Limitation. In relation to the specification of the Activity3, these processing principles are specified as pre- and post-conditions. This means that the usage of these types of personal data is in compliance with the processing principles stated in relevant legal frameworks and standards.

5.3.3 Threat Analysis

Threat Sources

The most likely threat sources that can contribute to or cause privacy harms can be identified through two different steps. The *first step* is to identify those who are directly involved in the collection, processing and/or dissemination of personal data. This can be achieved by identifying the actors who participate in the performance of each data-processing activity, along with the roles to which they are assigned. In so doing, a list of actors involved in all stages of the APDL model is identified, together with assigned roles and responsibilities.

The *second step* is to identify those who may be indirectly or unlawfully involved in the collection, processing and/or dissemination of personal data. This can be achieved by identifying third parties who may conduct further processing on personal

data (whether the data is identified, pseudonymous or anonymous) once it has been disclosed by data processors or controllers. In doing so, a list of entities with interests or concerns in the value of these types of personal data is identified. Data Protection Authorities, law enforcement bodies and other governmental agencies are examples of these entities. All such entities may be considered as potential threat sources.

Example 5.8. In the context of ECI, *organisers* (data controllers), *hosting service providers* (data processors) and *competent national authorities* (data controllers) are directly involved in the collection and processing of signatories' personal data. Thus, they might be considered as potential threat sources.

Further, intelligence and security services, employment agencies and advertising companies may indirectly involve in the collection, processing and/or dissemination of such data once it has been disclosed by the organisers, and might be considered as potential threat sources.

Threat Events

In this dissertation, we focus on data-driven events that are related more to primary assets than to supporting assets. As such, we adopt Solove's taxonomy [26] as a means for characterising adverse privacy events (threat events). The taxonomy helps facilitate the identification of these events in a comprehensive and concrete manner. It classifies the main types of adverse events into four basic groups: information collection, information processing, information dissemination and invasions. These events are arranged around a model that begins with the data subject, from which various entities — e.g. data holders, such as individuals or institutions — collect personal data. Data holders process the collected data. They may also disseminate or release the processed data to other entities.

We acknowledge that the taxonomy was developed to serve as a framework for the future development of the field of privacy law. It covers all aspects and dimensions of privacy. In our approach, however, we focus only on those adverse events that have implications on data privacy. From a technical perspective, these events need to be arranged around a widely used model in the field of systems

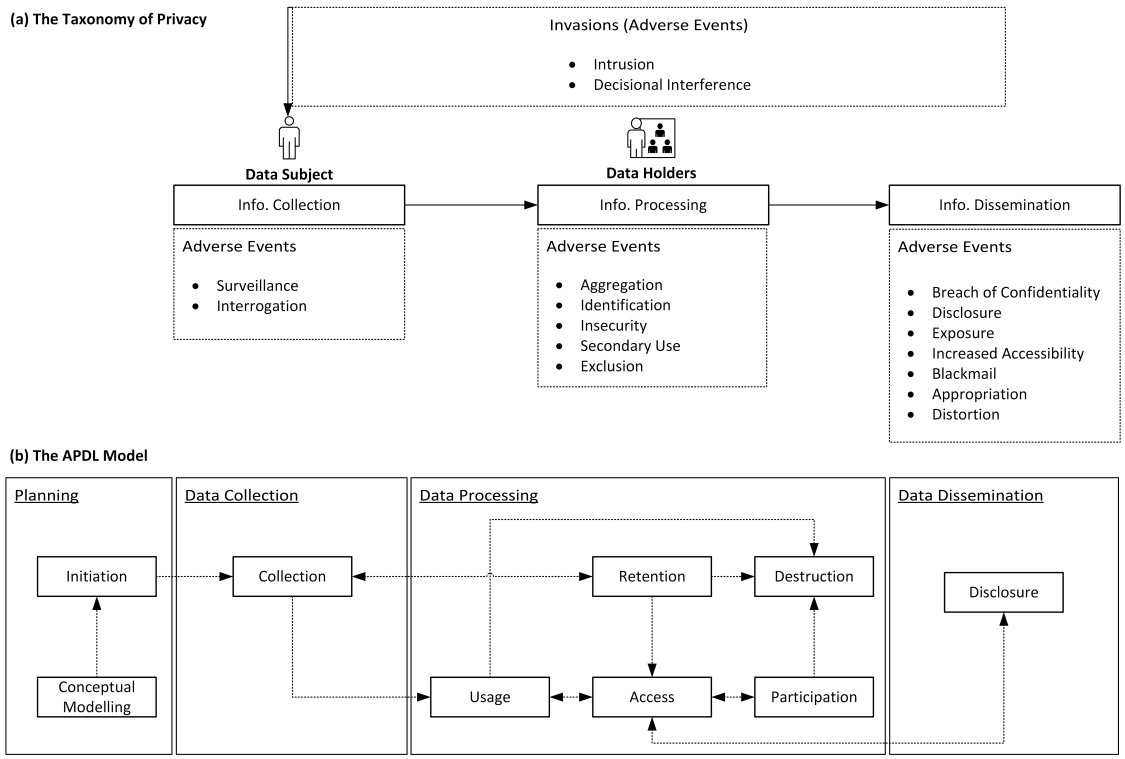

Figure 5.3: The adapted taxonomy of adverse privacy events.

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 define an *adapted taxonomy of adverse privacy events* by mapping the sub-categories of Solove's taxonomy onto the stages of the APDL model. Figure 5.3 shows the conceptual relationship between the taxonomy and the APDL model.

The sub-categories of Solove's taxonomy are considered as the main types of threat events. Each type of an adverse privacy event can be characterised by a set of attributes according to the nature of a data-processing activity in each stage of the APDL model that reflects the manner in which personal data is collected, processed and disseminated. Figure 5.4 shows the attributes of concrete threat events. We consider each in turn.

Collection. In the Collection stage, adverse events are related to the manner in which personal data is collected in terms of available choices and collection methods.

- **Surveillance.** It involves collecting or recording a large amount of personal data about the data subject's activities. Each type of a threat event can be characterised by a set of attributes: manner (continuous or discrete monitoring), type (covert or overt — i.e. passive or active) and nature (extensive or limited).

- **Interrogation.** It involves coercively collecting personal data by asking or probing unwarranted questions. Each type of a threat event can be characterised by a set of attributes: degree of coerciveness, type (direct or indirect) and nature (excessive or limited).

Retention. In the Retention stage, adverse events are related to the manner in which the collected personal data is structured, organised, stored and retained.

- **Aggregation.** It involves structuring, organising, storing or retaining integrated items of personal data about a data subject. Each type of a threat event can be characterised by a set of attributes: manner (anticipated or unanticipated), nature (excessive or limited), and data sources (internal or external).

- **Identification.** It involves structuring, organising, storing or retaining different items of personal data in a manner through which personal data can be linked to particular data subjects logically or physically. Each type of a

threat event can be characterised by a set of attributes: manner (anticipated or unanticipated), identifiability of data (identified, pseudonymous or anonymous) and linkability of data to personal identifiers (linked, linkable with reasonable effort, not linkable with reasonable effort or unlinkable).

- **Insecurity.** It involves improper data protection and handling. Each type of a threat event can be characterised by a set of attributes: nature (data handling or data protection) and type (design flaw, implementation flaw, retention time).

Access. In the Access stage, adverse events are related to the manner in which personal data is retrieved.

- **Insecurity.** In this stage, handling includes retrieval mechanisms. Each type of a threat event can be characterised by a set of attributes: nature (data handling or data protection) and type (design flaw, implementation flaw).

Participation. In the Participation stage, adverse events are related to the manner in which data subjects participate in the processing of personal by exercising their access rights to review or rectify their personal data and ensure that it is accurate, complete and up-to-date.

- **Exclusion.** It involves the failure to provide data subjects with notice and access to their personal data. Each type of a threat event can be characterised by a set of attributes: nature (partial or complete) and source of denial (unjustified, necessary for processing, or required by law or regulation).

Usage. In the Usage stage, adverse events are related to the manner in which personal data is manipulated and used.

- **Aggregation.** It involves altering, adapting, refining, aligning and combining or integrating different items of personal data about a data subject.

- **Identification.** It involves altering, adapting, refining or aligning different items of personal data in manner through which personal data can be linked to particular data subjects.

- **Insecurity.** It involves improper data protection and handling.

- **Secondary Use.** It involves using the collected personal data for purposes unrelated to the purposes for which it was initially collected without the knowledge and consent of the data subject. Each type of a threat event can be characterised by a set of attributes: conformity (with the specified purposes), agreement (with the obtained consent) and compliance (with legal frameworks).

Destruction. In the Destruction stage, adverse events are related to the manner in which personal data is erased, destroyed, redacted or disposed.

- **Insecurity.** It involves improper data protection and handling.

Disclosure. In the Disclosure stage, adverse events are related to the manner in which personal data is disseminated, made available or transmitted to third parties.

- **Breach of confidentiality.** It involves the revelation of confidential personal data about a data subject by violating a trusted relationship. Each type of a threat event can be characterised by a set of attributes: source of breach (individual, group of individuals or institution) and type of relationship (trusted, semi-trusted or untrusted).

- **Disclosure.** It involves the revelation of concealed and true personal data about a data subject to third parties. Each type of a threat event can be characterised by a set of attributes: nature (extensive or limited) and accuracy (accurate or inaccurate).

- **Exposure.** It involves the revelation of concealed personal data that refers to physical or emotional attributes about a data subject. Each type of a threat event can be characterised by a set of attributes: type of data (physical or emotional attributes), source of exposure (individual, group of individuals or institution) and completeness of data (reflects the capability of data for judgement, whether it is complete and can be used for judging a data subject's personality or character).

- **Increased accessibility.** It involves making personal data that is already available to the public easier to access. Each type of a threat event can be characterised by a set of attributes: type of data, nature (excessive or limited) and degree of accessibility.

- **Blackmail.** It involves coercing data subjects by threatening to reveal their concealed personal data for legal or illegal purposes. Each type of a threat event can be characterised by a set of attributes: type of data, degree of coercion and source of coercion.

- **Appropriation.** It involves the use of personal data that shapes a data subject's identity or personality for the purposes and goals of another. Each type of a threat event can be characterised by a set of attributes: value of the name or likeness (reputation, prestige, social or commercial standing) and purpose of another.

- **Distortion.** It involves exposing a data subject to the public inaccurately by revealing false and misleading personal data. Each type of a threat event can be characterised by a set of attributes: type of data and nature of falsehood (untrue, inaccurate or misleading data).

Example 5.9. In reference to Example 5.2, *unauthorised secondary use* is a concrete threat event that relates to 'Secondary Use' threat events that relate to the Usage stage of the APDL model. This type of adverse events is related to the manner in which personal data is used. It can be characterised as: it is not in conformance

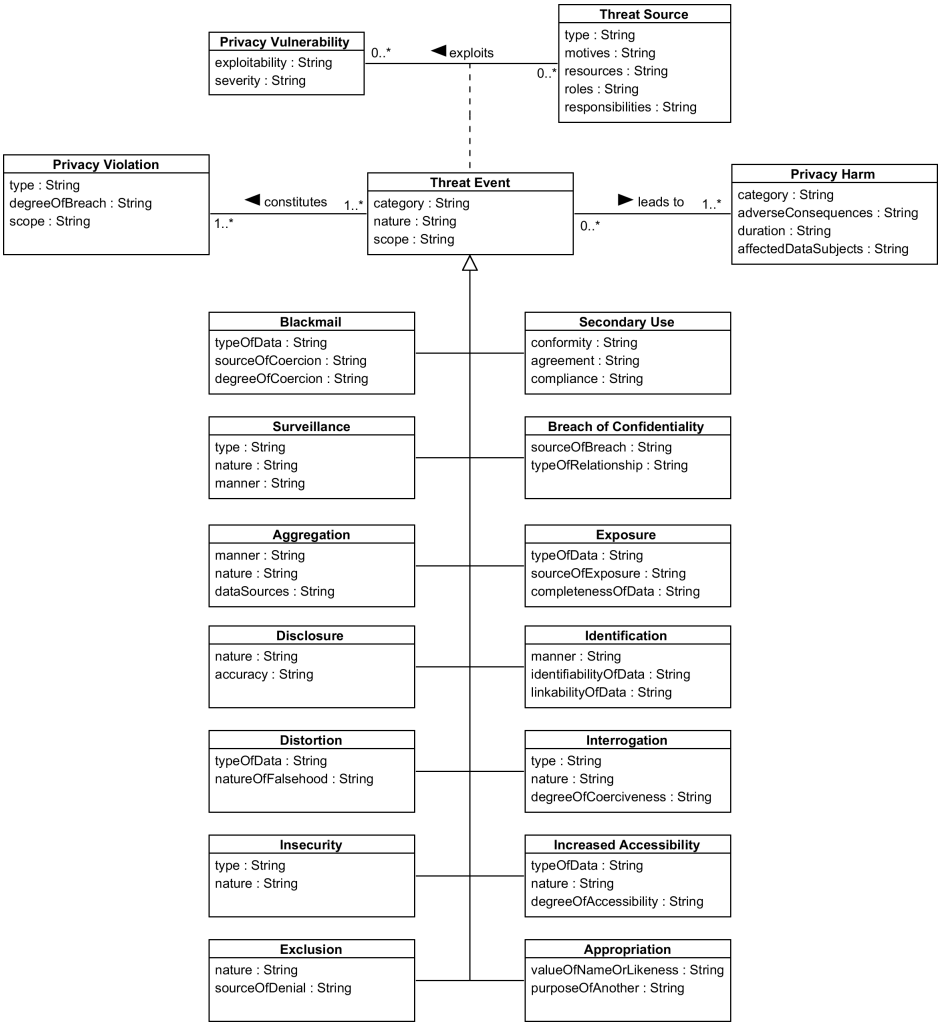

Figure 5.4: The conceptual relationship between the key risk factors along with their attributes.

with the specified purpose (to verify and certify the number of valid statements of support for a particular citizens' initiative); it is not in agreement with the obtained consent; and it is not in compliance with the applicable legal framework (the EU's GDPR and the Regulation (EU) No. 211/2011 on the Citizens' Initiative).

The *first step* of identifying threat events is to analyse the identified threat sources and privacy vulnerabilities for each context-relative processing norm to identify what the vulnerability enables a threat source to do and map this to the main types of threat events identified in the adapted taxonomy. In doing so, threats can be categorised according to the stage of the lifecycle and grouped by context-relative processing norms (which relate to data-processing activities). The

second step is to identify all possible exploitations that lead to each threat event by arranging all threat sources and privacy vulnerabilities that lead to the same threat event in the same stage of the lifecycle. This step is to analyse how a threat event occurs by: multiple threat sources who exploit the same vulnerability; or a threat source who exploits different vulnerabilities.

5.3.4 Privacy Harm Analysis

Privacy Violations

For each stage of the APDL model, the identified threat events need to be analysed to identify harmless and illegitimate or unanticipated actions that may result from the occurrence of each threat event. During the analysis, privacy vulnerabilities need to be considered to determine how they enable threat sources to take such actions. These actions are not able or likely to cause privacy harms to data subjects. In doing so, a list of actions that are categorised according to the stages of the APDL model is identified as a set of privacy violations.

Example 5.10. In the retention stage, *retaining signatories' data beyond the specified retention period* is an illegitimate or unanticipated action that may result from the occurrence of the threat event *unjustified data retention*, which results from the successful exploitation of *improper activity specification* by *organisers*. Its degree is excessive as it retains signatories' data for longer than necessary that exceed the specified retention time without operational or legal justifications. Its scope is individuals — i.e. those who are supported a particular initiative. This privacy violation is considered as illegitimate and unanticipated data-processing activity without adverse effects. In particular, signatories' data is retained in ways signatories would not reasonably expect, beyond the specified retention time, and without their knowledge and consent. In addition, the retention of such data does not have operational and legitimate grounds as it is no longer necessary to fulfil the specified purposes. Most importantly, this privacy violation is assumed to be without adverse actions taken against signatories.

Privacy Harms

For each stage of the APDL model, potential adverse effects (undesirable consequences) that may result from the occurrence of each threat event need to be identified. In doing so, a list of adverse effects is identified and categorised according to the stages of the APDL model. Then, the identified effects need to be analysed (as motives) to determine whether they can partially contribute to, or completely lead to, a harmful (adverse) action that uses personal data against the data subject in an unanticipated or coerced manner (privacy harms). During the analysis, privacy vulnerabilities need to be considered to determine how they enable threat sources to take such adverse actions. Most broadly, a privacy harm may result from a series of adverse effects of multiple threat events caused by multiple threat sources.

In order to classify the identified privacy harms, we use the same categories of privacy harms of [62] that have been identified in previous attempts from a legal perspective [26, 28]. In particular, privacy harms are classified into: physical harms; economic or financial harms; mental or psychological harms; harms to dignity or reputation; and societal or architectural harms [62].

Example 5.11. In the Collection stage, the main adverse effects of 'excessive data collection' is *gathering identifiable data* for signatories. In the Retention stage, the main adverse effects of 'unjustified data retention' is *storing identifiable records* for signatories, from which sophisticated profiles can be created. In the Usage stage, the main adverse effects of 'unjustified data integration' is *creating identifiable profiles* for signatories, from which sensitive data can be derived, such as religious beliefs or political affiliation. In the Disclosure stage, the main adverse effects of 'unauthorised data disclosure' from organisers to 'employment agencies' is *revealing sophisticated profiles* for signatories beyond expected boundaries that may be analysed to derive sensitive information that can inhibit certain rational judgements.

By analysing the identified effects, along with relevant threat sources, 'denial of job' can be derived as a privacy harm (a potential adverse action that may be taken against data subjects).

Summary. The results of the analysis approach can be used to develop and model reasonable *threat scenarios* that describe how the identified threat events attributed to specific threat sources, with capability to exploit privacy vulnerabilities, can contribute to or cause privacy harms.

5.4 A Privacy Risk Assessment

In order for risk assessments to be effective, the conceptual relationships among the key risk factors need to be represented in a way that is analytically useful for assessment by illustrating the dependencies between the nominal and assessable attributes of each risk factor, and the dependencies between the nominal and assessable attributes of all risk factors. Furthermore, the assessment rules that specify the range of values the key risk factors can assume need to reflect the assessable attributes of these factors to facilitate their roles in risk assessments and their translation into qualitative terms for multiple stakeholders. Importantly, potential privacy risks need to be assessed in a structured manner. As such, our assessment approach consists of four steps, as per Figure 5.5. The first step is to represent the conceptual relationships among the key risk factors for each privacy harm from which a reasonable set of threat scenarios can be generated. The second step is to assess the severity of privacy harms. The third step is to assess the likelihood of occurrence. The fourth step is to assess the risk levels of privacy harms in terms of their severity and likelihood.

In general, risk assessments require careful threat and vulnerability analysis to determine the extent to which threat events could adversely impact the privacy of data subjects and the likelihood that such events will occur. Accordingly, we consider the results of the analysis approach of Section 5.3 as relevant information that is necessary for determining the values of the attributes of key risk factors. In addition, we adopt the fixed scale of levels and the corresponding values of [5] (1. Negligible; 2. Limited; 3. Significant; and 4. Maximum) as assessment scales with refined and/or newly defined rules for assessing the key risk factors of the risk model of Section 5.2. These scales can be easily translated for multiple stakeholders

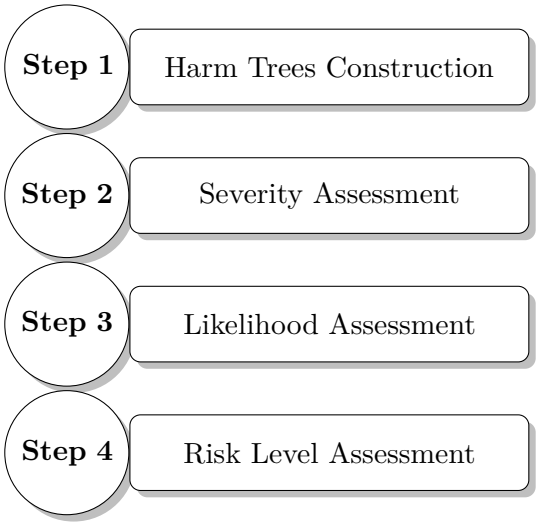

Figure 5.5: The main steps of the assessment approach.

Sum of values	Overall values
< 5	1. Negligible
$= 5$	2. Limited
$= 6$	3. Significant
> 6	4. Maximum

Table 5.1: A reference table for assessing the overall values from combinations of values, inspired by [5].

and allow relative comparisons between values in different scales or even within the same scale. Table 5.1 illustrates a set of rules for assessing overall values from combinations of values that can be applied to the key risk factors.

5.4.1 Harm Trees Construction

In the risk model of Section 5.2, a privacy harm results from one or more threat events, each of which results from the successful exploitation of one or more privacy vulnerabilities by one or more threat sources. Thus, it is useful to generate multiple threat scenarios describing how the threat events caused by the most likely threat sources can contribute to or cause a privacy harm.

In [62], a harm tree describes the relationship between a privacy harm (a root node) and all possible feared events (intermediate nodes) that exploit privacy weaknesses (leaf nodes) by the most likely risk sources, which are both represented

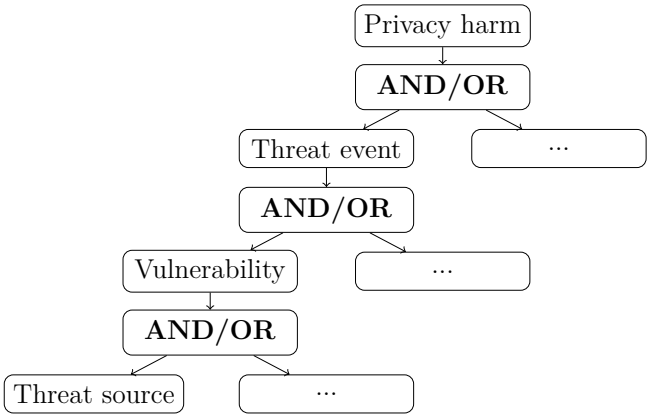

Figure 5.6: The structure of the refined harm tree.

as pairs of the form (privacy weakness, risk source). We slightly refine the harm tree by adding an additional level to present a privacy harm (a root node) and all possible threat events (intermediate nodes) that exploit the vulnerabilities of primary assets (intermediate nodes) by the most likely threat sources (leaf nodes), as illustrated in Figure 5.6. We use 'AND' and 'OR' connectors to combine child nodes and indicate whether all or some child nodes are necessary to enact the parent node. This refinement is to represent all possible exploitations of a vulnerability for each threat event by one or more threat sources. This helps analyse the exploitation of a vulnerability when there is collusion between two or more threat sources (when those threat sources are connected to a privacy vulnerability via 'AND'). In addition, it helps provide focus on the most important vulnerabilities that need to be addressed when a vulnerability is connected to several threat sources or its exploitation may lead to several threat events.

Example 5.12. Figure 5.7 shows the harm tree for the privacy harm PH.1: Denial of job (previously described in Example 5.11). It occurs when organisers (TS.1) makes unjustified data integration to profile signatories based on their ideas (TE.1), and share these profiles with employment agencies (TE.2). An employment agency (TS.2) makes excessive data inference (TE.3) to derive sensitive data, such as religious beliefs or political affiliation. An 'improper data model' (PV.1) and a 'lack of data minimisation' (PV.2) that may be exploited by TS.1 lead to the occurrence of TE.1. An 'improper purpose specification' (PV.3) and a 'lack of logs and audit

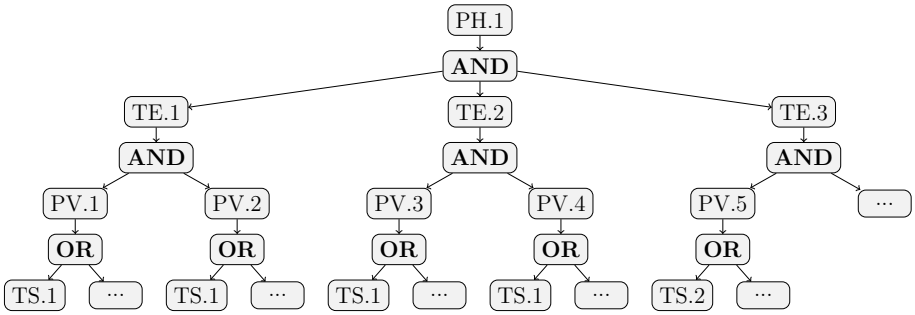

PH.1:	Denial of job	TE.1:	Unjustified data integration
TE.2:	Unauthorised data disclosure	TE.3:	Excessive data inference
PV.1:	Improper data model	PV.2:	Lack of data minimisation
PV.3:	Improper purpose specification	PV.4:	Lack of logs and audit trails
PV.5:	Weak anonymisation technique	TS.1:	Organisers
TS.2:	Employment agencies		

Figure 5.7: The structure of the harm tree for the privacy harm PH.1.

trails' (PV.4) that may be exploited by TS.1 lead to the occurrence of TE.2. A 'weak anonymisation technique' (PV.5) that is exploited by TS.2 leads to the occurrence of TE.3.

5.4.2 Severity Assessment

The severity of a privacy harm essentially depends on the *intensity* of adverse effects of associated threat events and the *range* of these effects suffered by a variety of data subjects.

Intensity

The intensity of a privacy harm represents the level of adverse effects on the affected data subjects. It is based on the 'extent of damage' attribute of the privacy harm, which indicates the level of adverse effects of the corresponding threat event. It is influenced by the 'adverse effects' attribute of the threat event, which, in turn, is influenced by the 'nature', 'category' and 'scope' attributes of the same event that reflect the irreversibility of the effects and the level of difficulty with which these effects might be overcome (duration), as per Figure 5.8. The factors that influence irreversibility vary between threat events according to their classification in the risk model of Section 5.2. As such, the 'nature' attribute abstractly represents other specific attributes of the threat events that are classified according to the stages of

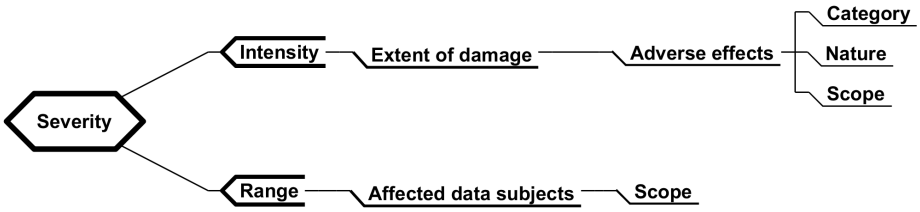

Figure 5.8: The dependencies between the nominal and assessable attributes of privacy harms and threat events.

Values	The affected data subjects encounter ...
1. Negligible	insignificant adverse effects, which can be reversed without difficulty and last for a short time.
2. Limited	slight adverse effects, which can be reversed with some difficulty and do not last for a long time.
3. Significant	serious adverse effects, which can be reversed with great difficulty and last for a certain length of time.
4. Maximum	severe adverse effects, which cannot be reversed at all and last for a long time.

Table 5.2: Intensity of a privacy harm.

the APDL model. These attributes, in turn, help assess the extent of damage caused by the adverse effects and the difficulty with which these effects can be reversed. Table 5.2 illustrates the rules for assessing the intensity value of a privacy harm.

Range

The range of a privacy harm represents the scope of adverse effects of threat events encountered by a variety of data subjects. It is based on the 'affected data subjects' attribute, which is influenced by the 'scope' attribute of the corresponding threat event that reflects the domain of the adverse event, assessed, perhaps, in terms of the number and categories of potential data subjects whose personal data is affected, as per Figure 5.8.

Table 5.3 illustrates the rules for assessing the value of the range of a privacy harm. These are a refinement of the rules of [62, 39] to distinguish between the range of a privacy harm. The severity of a privacy harm is assessed by adding the assessed value of the intensity and the assessed value of the range, then selecting

Values	Description
1. Negligible	Only specific individuals are affected.
2. Limited	Specific individuals and their relatives and/or friends are affected.
3. Significant	Specific categories of individuals are affected.
4. Maximum	The whole of society is affected.

Table 5.3: Range of a privacy harm.

the overall value according to Table 5.1. Trace matrices can be used as a tool to document the severity of each privacy harm according to its risk factors.

Example 5.13. The intensity of PH.1 is based on the extent of damage caused by, the irreversibility of, and the duration of, the adverse effects of associated threat events TE.1, TE.2 and TE.3. TE.1 is categorised as a type of 'identification'. It is characterised as unanticipated by signatories; it is also characterised as extensive. TE.2 is categorised as a type of 'disclosure'. It is characterised as extensive and accurate. TE.3 is categorised as a type of 'identification'. It is characterised as unanticipated by signatories; it is also characterised as extensive. The adverse effects of these threat events are identified as a series of related impacts started by creating identifiable profiles, discovering private facts about signatories, then revealing sensitive information beyond expected boundaries and ending by filtering job candidates according to their ideas. They are assessed as slight because filtering job applications depends on other factors, including qualifications, experience, candidates' qualities and abilities, the requirements of the role, etc. The duration of these effects may last for a certain length of time: the period of a job offer. However, they may last for longer than the period of a job offer when the disclosed data is used by several job offers or several employment agencies to categorise candidates. Once profiles are created, disclosed to employment agencies and sensitive information is inferred, it is technically difficult to reverse these effects. As such, the intensity of PH.1 is assessed as '2. Limited'.

PH.1 may affect specific signatories based on their ideas. They may also affect specific categories of signatories based on specific conditions. Together, these may impact job offers. Thus, the range of this privacy harm is assessed as '3. Significant'.

The severity of PH.1 is assessed as '2. Limited' by adding the assessed value of the intensity and the assessed value of the range, and selecting the overall value according to Table 5.1.

5.4.3 Likelihood Assessment

The likelihood of occurrence of a privacy harm is *the highest value* of the overall likelihood of occurrence of associated threat events. The *overall likelihood* of a threat event is a combination of the likelihood of the threat event occurrence and the likelihood of the threat event resulting in adverse impacts. The likelihood of a threat event occurrence depends on the *capability* of the most likely threat sources and the *exploitability* of the relevant privacy vulnerabilities. The likelihood of a threat event resulting in adverse impacts depends on the *capability* of the most likely threat sources and the *severity* of the relevant privacy vulnerabilities.

Seriousness

The seriousness represents the level of a vulnerability, which is based on its 'exploitability' and 'severity'. The 'exploitability' attribute indicates the level of exploitation of a primary asset's vulnerability, whereas the 'severity' attribute indicates the relative importance of mitigating a primary asset's vulnerability. Both are influenced by the attributes of the relevant element of context-relative processing norms: personal data, data-processing activities, involved actors and processing principles, as per Figure 5.9. These include the attributes of personal data: 'category', 'type' and 'linkability'. All these attributes are used to estimate the degree to which the vulnerability of the primary asset can be exploited, and to enable the threat source to conduct adverse actions that breach these norms and violate contextual integrity.

Table 5.4 illustrates the rules for assessing the value of the exploitability of a vulnerability. Table 5.5 illustrates the rules for assessing the value of the severity of a vulnerability. The seriousness of the vulnerability is estimated by adding the

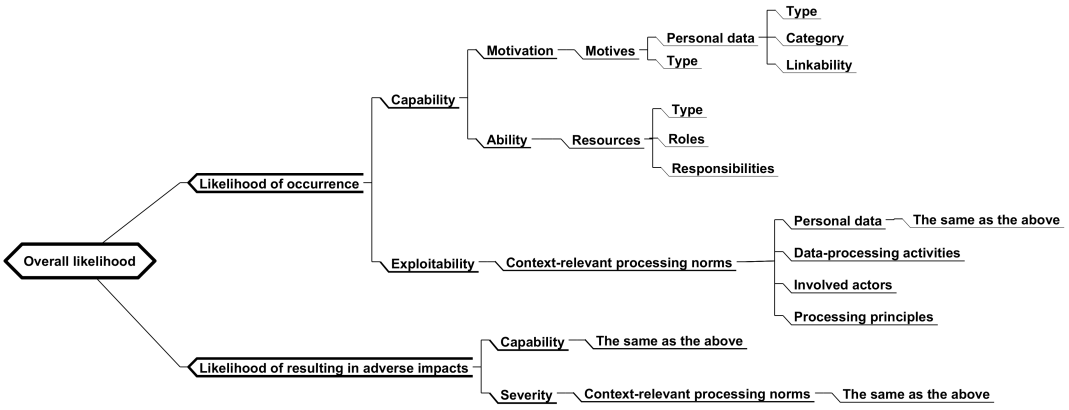

Figure 5.9: The dependencies between the nominal and assessable attributes of privacy vulnerabilities and threat sources.

Values	The exploitation of a vulnerability ...
1. Negligible	does not appear possible.
2. Limited	appears to be difficult.
3. Significant	appears to be possible.
4. Maximum	appears to be extremely easy.

Table 5.4: Exploitability of a vulnerability.

Values	The successful exploitation of a vulnerability ...
1. Negligible	leads to insignificant impacts.
2. Limited	leads to slight impacts.
3. Significant	leads to serious impacts.
4. Maximum	leads to severe impacts.

Table 5.5: Severity of a vulnerability.

estimated value of the exploitability and the estimated value of the severity, then selecting the overall value according to Table 5.1.

Capability

The capability of a threat source represents the motives, skills or resources that make them able to exploit the vulnerabilities of a primary asset. It is mainly estimated based on the 'motivation' and 'ability' of the threat source, which are determined by the values of 'type', 'resources', 'role' and 'responsibility' as attributes of the threat source.

The 'motivation' indicates the value of personal data to threat sources that

Values	The threat source ...
1. Negligible	does not have any specific motives based on the value of personal data.
2. Limited	has indistinct and unreasonable motives based on the value of personal data.
3. Significant	has definite and reasonable motives based on the value of personal data.
4. Maximum	has multiple, definite and strong motives based on the value of personal data.

Table 5.6: Motivation of a threat source.

stimulates their motives to exploit vulnerabilities of the primary assets. In general, personal data has value both to data subjects and to the entities that collect, process and disseminate it. In addition, it has a nuisance value when it is exploited for unfair or malicious purposes. The value of personal data is influenced by the attributes of the personal data as a primary asset: 'category', 'type' and 'linkability'. The motive is influenced by the 'type' attribute of the threat source. Table 5.6 illustrates the rules for assessing the value of the motivation of a threat source.

The 'ability' indicates the level of resources by which a threat source is able to exploit the vulnerabilities of a primary asset. These resources include the skills, background knowledge, privileges, financial and technical resources. These assessable attributes are influenced by the 'type' attribute of the threat source. The 'privileges' attribute is influenced by the roles of the threat source and assigned responsibilities, if any. Further, technical and financial resources are influenced by the 'type' attribute. The 'background knowledge' is influenced by other factors, such as the type of relationship with the data subject. Table 5.7 illustrates the rules for assessing the value of the ability of a threat source.

The capability of the threat source is estimated by adding the assessed values of motivation and ability, then selecting the overall value according to Table 5.1.

The likelihood of occurrence of a threat event is estimated by adding the assessed values of capability and exploitability of the relevant primary assets' vulnerabilities, then selecting the overall value according to Table 5.1. The likelihood of the threat event resulting in adverse impacts is estimated by adding the assessed value of

Values	The threat source ...
1. Negligible	does not appear to have specific skills, background knowledge, privileges, financial and technical resources to exploit the vulnerability.
2. Limited	has insufficient skills, background knowledge, privileges, financial and technical resources to exploit the vulnerability.
3. Significant	has real and significant skills, background knowledge, privileges, financial and technical resources to exploit the vulnerability.
4. Maximum	has definite and unlimited skills, background knowledge, privileges, financial and technical resources to exploit the vulnerability.

Table 5.7: Ability of a threat source.

capability and the assessed value of the severity of the relevant primary assets' vulnerabilities, then selecting the overall value according to Table 5.1. The overall likelihood of occurrence of a threat event causing adverse impacts is estimated by adding the assessed value of the likelihood of occurrence of the threat event and the assessed value of the likelihood of the threat event resulting in adverse impacts, then selecting the overall value according to Table 5.1. The likelihood of occurrence of a privacy harm is the highest value of the overall likelihood of the corresponding threat events. Trace matrices can be used as a tool to document the likelihood of each privacy harm according to its risk factors.

Example 5.14. The likelihood of occurrence of TE.1, TE.2 and TE.3 is assessed on the capability of TS.1 and TS.2, and the exploitability of PV.1–PV.5 for each possible exploit.

The motivations of TS.1 and TS.2 are based on the values of 'initiative data' to those threat sources and their motives according to these values. The utility of ECI-related data makes such data highly valuable to organisers and employment agencies. It also has a nuisance value when it is exploited by employment agencies. Thus, the motivation of TS.1 is assessed as '3. Significant' and the motivation of TS.2 is assessed as '4. Maximum'.

The ability of those sources is based on their skills, background knowledge, privileges, and technical and financial resources. According to the 'type' attribute of TS.1, they are insiders and individuals. This implies that they 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 such data according to their roles and responsibilities. Organisers play the role of data controllers. Based on these, they have access rights to the data. In addition, they have both technical and financial resources to benefit from the values of the collected data by creating comprehensive and identifiable profiles. As such, the abilities of TS.1 are assessed as '4. Maximum'.

According to the 'type' attribute of TS.2, they are outsiders and institutions. Employment agencies are third parties that do not have direct roles with respect to the processing of personal data. Based on this, they do not have access rights to the initiative data; rather, they can legally process this data when it is anonymised. In addition, they have both technical and financial resources to benefit from the values of the disclosed data by making excessive inference. As such, the ability is assessed as '2. Limited'. The capability of TS.1 is assessed as '4. Maximum'; those of TS.2 are assessed as '3. Significant'.

The seriousness of PV.1 – PV.5 is based on the exploitability and severity of each exploit. In this example, we consider only the seriousness of PV.1. The ease of the exploitation of PV.1 is influenced by the relevant element of context-relevant processing norms. The relevant element of the processing norms is 'attributes', which refers to personal data. In this context, personal data is classified as 'identification and contact data'. This type is categorised as 'collected data'. These types of data are not sensitive in themselves; rather, they are valuable and can be used to derive sensitive data, such as political affiliation or religious beliefs. Initiative data can easily be linked with 'identification and contact data' with reasonable effort as they are modelled, collected and processed by a hosting service and accessed by organisers. The vulnerability of 'improper data model' can be easily exploited by organisers. Thus, the exploitability of PV.1 is assessed as '3. Significant'.

The severity of PV.1 is influenced by the relevant element of context-relevant processing norms. The relevant element of the processing norms is 'attributes', which refers to personal data. Thus, PV.1 enables TS.1 to breach the processing

SEVERITY					
4. Maximum					
3. Significant					
2. Limited					
1. Negligible					
	1. Negligible	**2. Limited**	**3. Significant**	**4. Maximum**	LIKELIHOOD

Figure 5.10: Risk Map

norms and violate the contextual integrity by making unjustified data integration (TE.1) with the aim of creating profiles for signatories based on their ideas. This type of aggregation is a threat event that may lead to the privacy harm PH.1. As such, the severity of PV.1 is assessed as '3. Significant'. The seriousness of PV.2, PV.3, PV.4 and PV.5 are similarly assessed.

The threat event TE.1 may result from the exploitation of PV.1 and PV.2: both vulnerabilities are necessary for the occurrence of TE.1. The vulnerability PV.1 may be exploited by TS.1. With reference to the harm tree, those sources are sufficient to exploit the vulnerability PV.1. This leads to one possible exploit (EX.1). In this case, the highest value of the overall likelihood of EX.1 is taken: '4. Maximum'. The overall likelihood of occurrence of all possible exploitations of TE.2 and TE.3 are similarly assessed. As such, the likelihood of occurrence of PH.1 is assessed as '4. Maximum' by taking the highest value of the overall likelihood of occurrence of TE.1, TE.2 and TE.3.

5.4.4 Risk Level Assessment

The risk level of a privacy harm is assessed as a (likelihood, severity) pair. These pairs are used to determine the order in which the identified risks should be managed according to their severity and likelihood. We adopt the 'risk map' of [5] to locate the assessed risks according to their levels. The likelihood of a privacy risk is plotted on the X-axis; its severity is plotted on the Y-axis, as per Figure 5.10.

5.5 Summary

In this chapter, we have described a privacy risk model that considers legal, organisational, societal and technical aspects. The risk model defines the main

factors that have an impact on privacy risks along with their nominal and assessable attributes and conceptual relationships. It was built upon fundamentals from the legal privacy literature to refine key concepts and assessable risk factors, as well as the conceptual relationships among the 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, which can be used to derive privacy vulnerabilities. They also facilitate the identification of threat events in a systematic manner by providing a taxonomy of adverse events 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.

In addition, we have illustrated an analysis approach that describes how combinations of risk factors are identified and analysed at a consistent level of detail. The analysis approach is built upon the privacy risk model. It also adopts the APDL model as a sufficiently robust model that facilitates end-to-end privacy protection to serve as the basis for the identification and analysis of potential privacy risks in a comprehensive manner. Such a robust model sufficiently and contextually represents data-processing activities in a way that is amenable to risk analysis.

Furthermore, we have presented an assessment approach that adopts the risk model that: characterises the risk factors by well-defined attributes (nominal and assessable) to facilitate the identification, analysis and assessment of these factors in a systematic and traceable manner; and describes the dependences between the nominal and assessable attributes of the key risk factors, to refine how each risk factor can be used as an input to estimate the levels of risk. The assessment approach also refines the concept of the harm tree by adding an additional level that separates threat sources from privacy vulnerabilities to represent the conceptual relationships among the key risk factors in a way that is analytically useful for analysis and assessment. In addition, the approach adopts fixed levels of scale that can be easily

translated into qualitative terms for multiple stakeholders and establishes a set of assessment rules that reflect the assessable attributes of the key risk factors.

In the next chapter, we will illustrate the third activity of our principled approach for engineering PbD: privacy-enhancing strategies. In addition, we will explain how to use risk assessment results as the basis for setting objectives from which architectural tactics can be identified with the aim of managing the assessed privacy risks.

Quality attribute scenarios and architectural tactics are some of the tools available for the creation of an architecture.

— Len Bass, Paul Clements and Rick Kazman

6 Privacy-Enhancing Strategies

Contents

6.1 Introduction

This chapter gives a relatively detailed description of a privacy-enhancing tactical approach, which contributes to the third activity of our principled approach for engineering PbD. The tactical approach illustrates how to capture privacy requirements and identify privacy controls to meet these requirements. These requirements are expressed as quality attribute scenarios derived from privacy risk assessments and the corresponding treatment strategies. In particular, Section 6.2

characterises privacy protection as a quality attribute by means of a general quality attribute scenario. Section 6.3 describes the main steps of the tactical approach that uses the concept of architectural strategies to support the adoption of PETs in the early stages of the design process to achieve various levels of privacy protection in a contextual manner. These strategies are collections of privacy architectural tactics, which are described through privacy design patterns and realised by PETs. Finally, Section 6.4 gives a brief summary of the chapter. The work presented in this chapter previously appeared in [107].

6.2 Privacy Protection as a Quality Attribute

In essence, privacy is motivated by several concerns, including compliance, user trust, risk management and ethical concerns. As such, integrating privacy into the early stages of the design process requires capturing both functional and non-functional requirements. The former are typically expressed as use cases, whereas the latter are expressed as a set of constraints or system-specific quality scenarios. In spite of the functional requirements of privacy that focus on, for example, purpose specification, collection, use, disclosure and retention limitation, our approach focuses only on non-functional requirements, as it considers privacy protection as a quality attribute. In the context of software engineering, software architectures can be designed, analysed and evaluated through quality attributes, which are often considered as requirements that have impacts on architectural decisions.

6.2.1 Quality Attribute Scenarios

Typically, general quality attribute scenarios are used as a means of characterising quality attributes to avoid non-operational or overlapping definitions [52]. As per Figure 6.1, the elements of a general scenario are: the stimulus; the source of a stimulus; the environment where the system is; the artefact being stimulated; the response as the result of the arrival of the stimulus; and the response measure [51]. By placing the focus on privacy protection as a quality attribute, a general scenario for the quality attribute of privacy protection needs to be modelled to specify

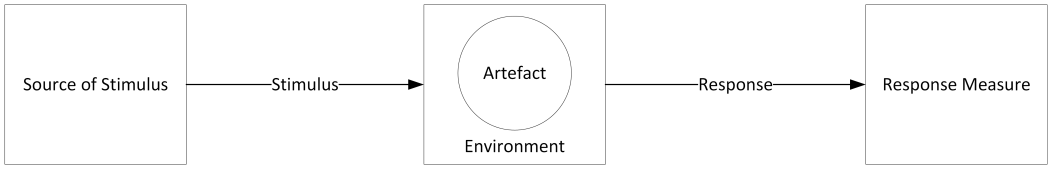

Figure 6.1: The main elements of a quality attribute scenario.

the responses of the system to realise specific goals (privacy protection goals) at an architectural level by indicating the range of values its elements can take. The privacy protection general scenario can then be used for generating the most important concrete (system-specific) quality scenarios, which are considered as quality attribute requirements.

6.2.2 Architectural Tactics

The achievement of quality attribute scenarios relies on fundamental design decisions, namely architectural tactics. In the context of privacy protection, for example, data avoidance, data minimisation, data hiding, generalisation, separation and isolation are all techniques for achieving or supporting the protection goal of unlinkability. These techniques, among others, are considered as fundamental design decisions that control responses to stimuli (adverse privacy events). Further, each is an architectural tactic — or a design option. Importantly, architectural tactics may refine other tactics. For example, data minimisation, as an abstract tactic, can be refined into concrete tactics, such as data avoidance, data reduction, etc. Thus, tactics need to be organised in a hierarchical manner to represent both abstract and concrete tactics. In addition, tactics can be classified into categories according to their underlying goals for privacy protection. Figure 6.2 shows how tactics can be represented in a generic hierarchical structure, categorised according to their goals for privacy protection. Tactics with italics represent abstract tactics that can be refined into concrete ones. The lowest level of a category has only concrete tactics; other levels of the category may have both abstract and concrete tactics.

To classify these tactics according to their underlying goals, we adopt the privacy protection goals of [4] as they provide a common scheme for addressing legal,

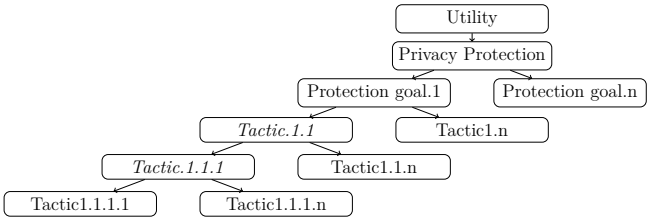

Figure 6.2: The hierarchical structure of privacy architectural tactics.

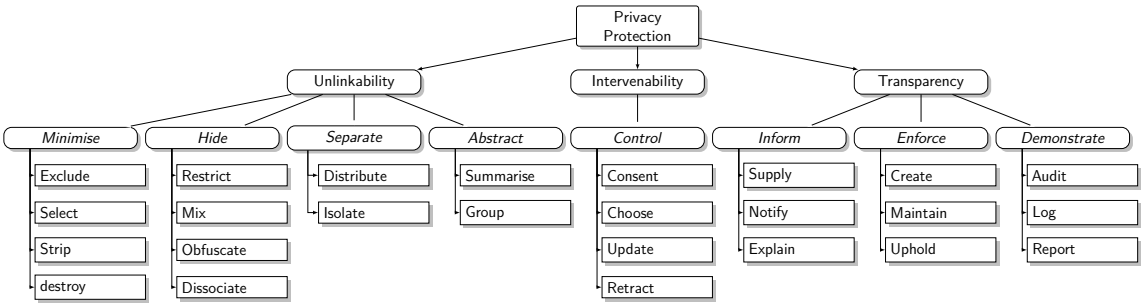

Figure 6.3: The structure of privacy tactics according to the adopted protection goals.

organisational, societal and technical aspects of privacy. These goals distinguish between three protection goals for information security (confidentiality, integrity and availability) and three protection goals for privacy protection (unlinkability, intervenability and transparency). In this dissertation, the focus is on those that pertain to privacy protection. In addition, we consider only the tactics of [3] — we do not propose new tactics. We also refer to the privacy design strategies of [3] as abstract tactics for the goals of privacy protection. These design strategies are not architectural strategies; rather, they are goals of privacy protection tactics — i.e. abstract tactics that can be refined into concrete tactics. We organise these tactics as a hierarchy according to the adopted privacy protection goals. However, we should note that the hierarchy is intended only to demonstrate the tactics of [3] and that any list of tactics is necessarily incomplete. The organisation is intended to provide a path for software engineers to search for appropriate tactics. Figure 6.3 shows how privacy tactics are organised as a hierarchy according to the adopted privacy protection goals. Table 6.1 briefly summarises the underlying goals of these tactics.

Abstract tactic	Underlying goals
Minimise	It aims to limit usage by partly or entirely avoiding the processing, deciding on a case by case basis, removing unnecessary personal data items or completely removing personal data.
Hide	It aims to prevent exposure by preventing unauthorised access to personal data, processing data randomly, preventing understandability of data or removing the correlation between different pieces of personal data.
Separate	It aims to prevent correlation by partitioning personal data or processing parts of personal data independently.
Abstract	It aims to limit detail by extracting commonalities in personal data or inducing less detail from personal data prior to processing.
Control	It aims to provide data subjects with means to give their consent to the processing of personal data, allow for the selection of personal data, keep their personal data accurate and up-to-date or remove any personal data in a timely fashion.
Inform	It aims to provide clarity by supplying extensive resources on the processing of personal data, notifying data subjects of any updates on the processing of personal data in a timely manner or explaining information on personal data processing in a concise and understandable form.
Enforce	It aims to ensure commitment by creating privacy policies, maintaining privacy when designing or modifying features and upholding privacy policies and controls.
Demonstrate	It aims to ensure evidence by auditing all day-to-day activities for any risks to personal data, tracking all data-processing activities without revealing personal data, and analysing collected information on audits and logs periodically to review improvements to data protection.

Table 6.1: The underlying goals of the privacy tactics of [3].

6.2.3 A General Scenario for Privacy Protection

To develop a general scenario for the quality attribute of privacy protection, it is essential to identify all possible values of its elements. To do this, we believe that risk-based and goal-oriented approaches are complementary: they both aim to support better understanding of how to prevent potential privacy risks and comply with abstract privacy principles by means of quality attribute requirements. As such,

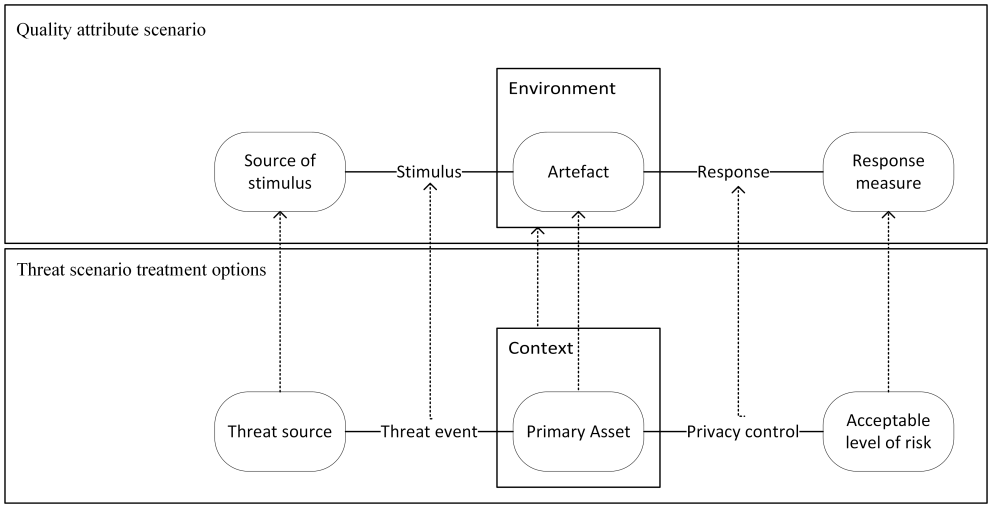

Figure 6.4: Threat scenarios versus quality attribute scenarios.

the elements of a threat scenario are mapped to the elements of a quality attribute scenario, as per Figure 6.4. The threat source is the source of a stimulus and the threat event is the stimulus. The primary asset is the artefact and the context in which personal data is processed is the environment. The privacy control (technical measure) that addresses the relevant privacy vulnerabilities is the response, and the acceptable level of risk is the response measure.

By mapping these scenarios, we identify all possible values of the elements (source of stimulus, stimulus, artefact, environment and response measure) of the general scenario for privacy protection. The values of the response element are identified in relation to the abstract tactics (privacy design strategies) of [3]. Figure 6.5 shows a general scenario for the quality attribute of privacy protection along with possible values of its elements. From this general scenario, concrete quality scenarios can be derived by instantiating each of its elements.

6.3 A Privacy-Enhancing Tactical Approach

Our tactical approach aims to articulate the desired privacy protection goals in a contextual manner. It also aims to aid software engineers in determining combinations of existing tactics, design patterns and PETs that will achieve, or contribute to the achievement of, the desired protection goals. As such, it helps

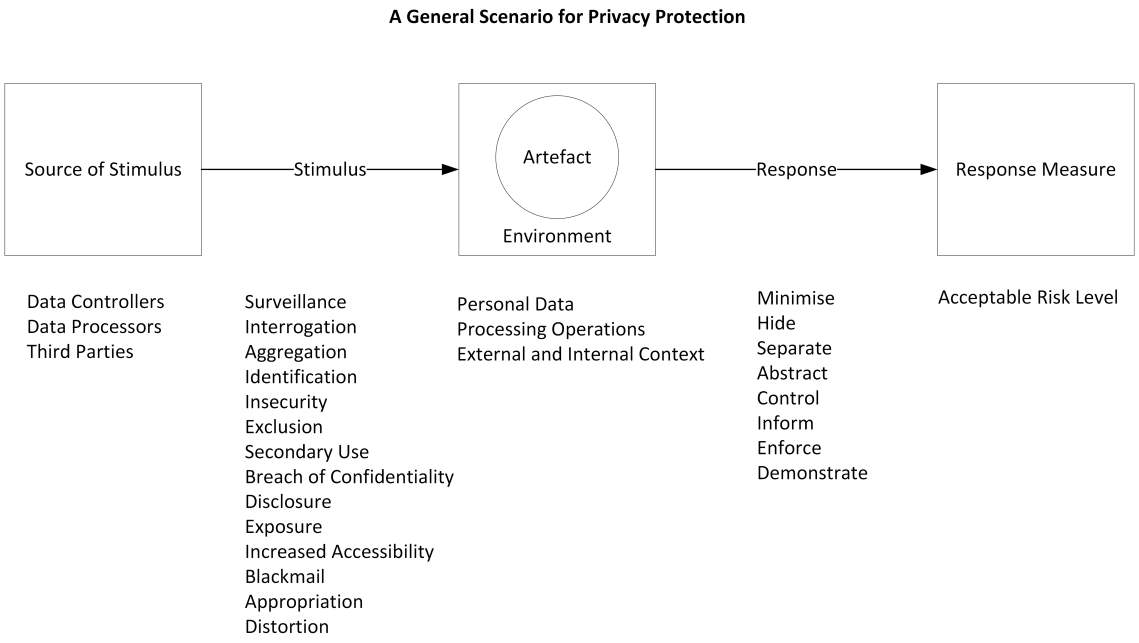

Figure 6.5: A general scenario for the quality attribute of privacy protection.

software engineers to justify and reason about their architectural choices. As per Figure 6.6, it consists of five main steps, which, in combination, establish a link between quality attribute requirements and architectural decisions, with the aim of mapping privacy requirements to suitable software architectures. First, context-relative (desired) privacy protection goals are articulated and expressed as quality attribute scenarios. Second, architectural tactics that control the response element of the generated scenarios are identified to achieve the desired privacy protection goals. Third, the architectural tactics are described through privacy design patterns according to a set of selection criteria. Fourth, design patterns are realised by appropriate PETs according to a set of selection criteria. Finally, architectural strategies are identified as combinations of architectural tactics, along with the corresponding design patterns and PETs, according to the desired protection goals. We consider each in turn.

6.3.1 The Articulation of Privacy Protection Goals

This step aims to articulate desired privacy protection goals by analysing concrete quality scenarios, which are generated by instantiating each element of the general

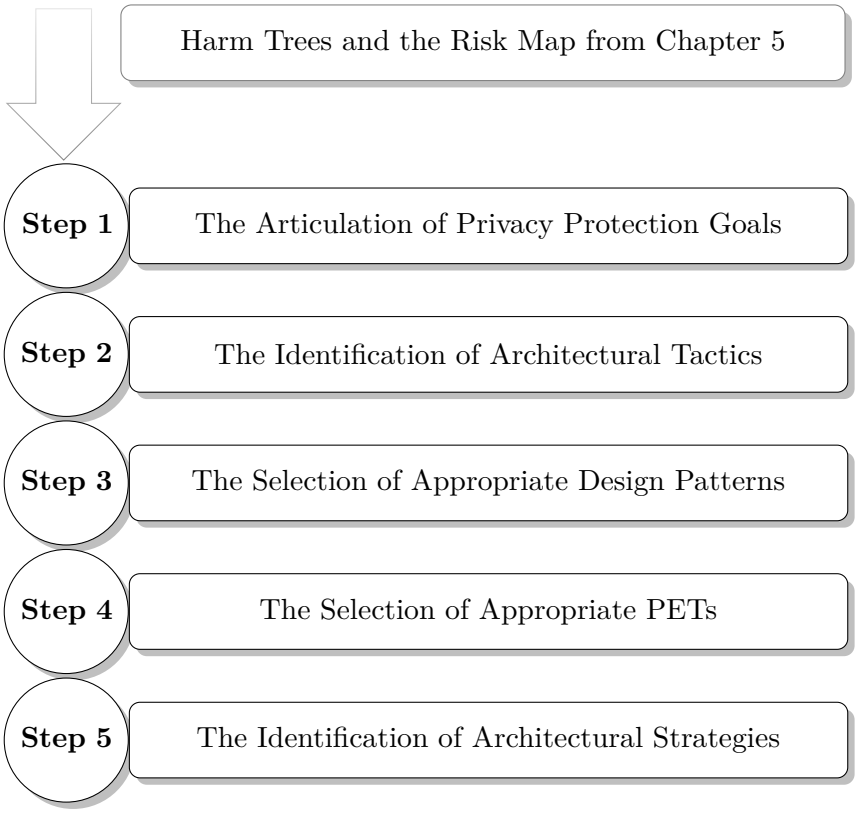

Figure 6.6: The main steps of the tactical approach.

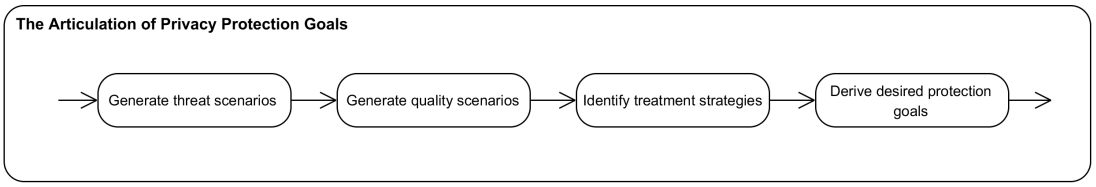

Figure 6.7: The main steps of articulating context-specific protection goals.

scenario in relation to the generated threat scenarios. As previously discussed in Chapter 5, threat scenarios are generated using the concept of the harm tree as a graph-based analysis technique. Each constructed harm tree identifies a set of threat scenarios that lead to a specific privacy harm (each path of the harm tree is considered as a threat scenario). As such, threat scenarios are used as the basis for generating the most important concrete quality scenarios that specify the responses of the system to realise specific privacy protection goals. Figure 6.7 illustrates the main steps of articulating context-relative protection goals.

For each privacy harm, concrete quality scenarios that correspond to its threat

scenarios need to be generated. For each quality scenario, a desired protection goal needs to be articulated as the specification of a reasonable and appropriate response to control the corresponding stimulus. This can be achieved by analysing the levels of the identified risks of privacy harms along with the capability of threat sources, the vulnerabilities that might be exploited by those sources, and the adverse effects of the threat events.

To order and prioritise the identified risks of privacy harms, we identify four distinguished levels. The first level concerns *maximum risks*: risks with significant or maximum severity and likelihood. The second level concerns *significant risks*: risks with significant or maximum severity and negligible or limited likelihood. The third level concerns *limited risks*: risks with negligible or limited severity and significant or maximum likelihood. The fourth level concerns *negligible risks*: risks with negligible or limited severity and likelihood. To treat the located risks in each level, treatment strategies that avoid or reduce either the likelihood or severity of the risk of the privacy harm or both need to be identified according to all possibilities of their severity and likelihood. As such, treatment strategies are considered as a means for determining the degree to which privacy is required. Based on appropriate treatment strategies, desired protection goals can then be derived with reference to the adopted goals (unlinkability, intervenability and transparency), which serve as abstract protection goals as they articulate what is being protected and who from. A desired privacy protection goal is a refinement of the response element — they are special instances of the abstract protection goals.

Example 6.1. In reference to the constructed harm tree of Example 5.12, there are five reasonable threat scenarios for the privacy harm PH.1. For each scenario, a concrete quality scenario needs to be generated. For illustration purposes, we describe only one threat scenario.

Threat scenario #1 PH.1 occurs when organisers exploit the improper data model and the lack of data minimisation vulnerabilities, and make unjustified data

Quality Scenario #1	**Elements**
Source of stimulus	Organisers
Stimulus	Unjustified data integration
Context	ECI
Artefact	Initiative data and signatories' data
Response	Prevent the linkability of initiative data with other types of signatories' personal data
Response measure	Acceptable level of risk without adverse effects on data utility

Table 6.2: An example of privacy protection scenarios.

integration to create identifiable profiles for signatories who give their support to a particular initiative.

Privacy protection quality scenario #1 Organisers may make unjustified data integration to create identifiable profiles for signatories who give their support to a particular initiative. The data integration is conducted from disparate sources by exploiting the improper data model and the lack of data minimisation. The system shall prevent such integration that leads to linking the initiative's data to signatories' data without an adverse effect on its utility. Table 6.2 shows the main elements of privacy protection scenario #1.

Desired privacy protection goal #1 The quality scenario #1 is to specify the response of the system to realise a desired privacy protection goal. In reference to our privacy protection characterisation, 'preventing the integration that can lead to linking initiative data to other types of data across domains' is the desired protection goal, which is related to 'unlinkability' as an abstract privacy protection goal.

6.3.2 The Identification of Architectural Tactics

The specification of concrete quality scenarios enables the refinement of quality requirements and the articulation of desired privacy protection goals, but it does not provide architectural guidance for achieving these goals. By articulating the

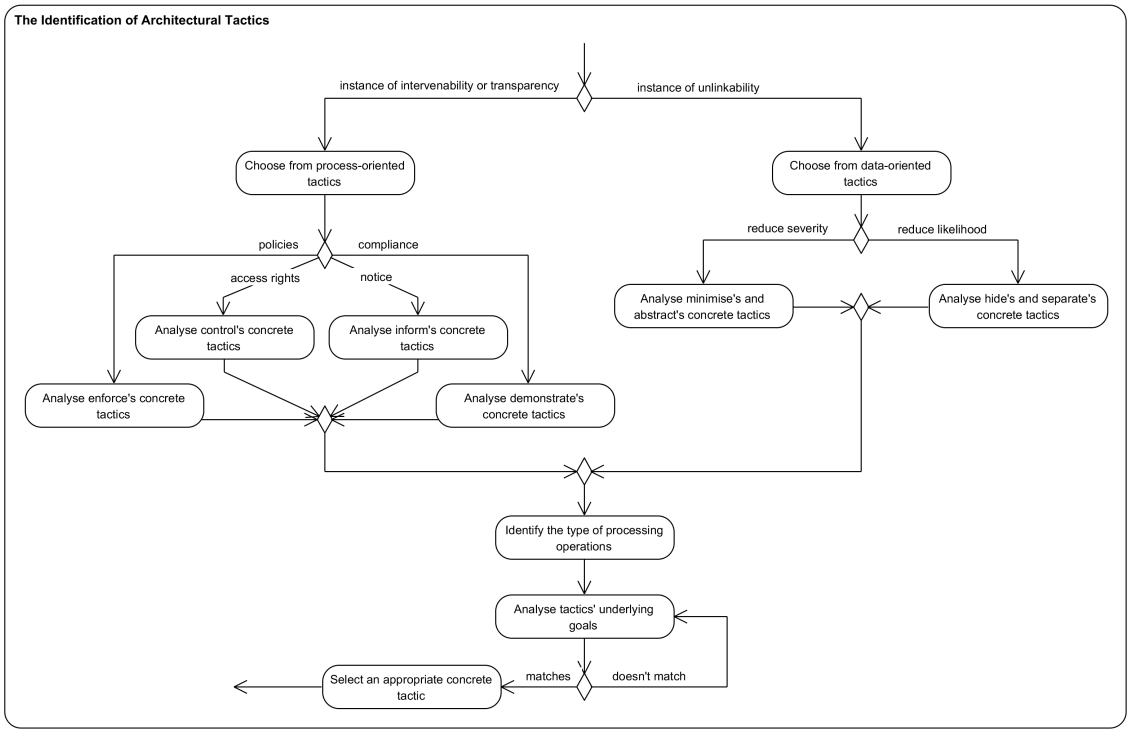

Figure 6.8: The main steps of identifying appropriate architectural tactics.

desired protection goals, appropriate techniques that control the response element of quality scenarios can be identified to achieve these protection goals. In order to choose an appropriate architectural tactic, it is important to determine whether the desired protection goal relates to unlinkability, intervenability or transparency: unlinkability can be achieved by data-oriented tactics, whereas intervenability and transparency can be achieved by process-oriented tactics, as per Figure 6.8.

To choose from data-oriented tactics, it is important to determine whether the treatment strategy is to avoid or reduce either the likelihood or severity of the risk of the privacy harm, or both. Based on the treatment strategy, the intent of the technique can be determined. For example, a risk with maximum severity and likelihood must be avoided, thus data avoidance is an appropriate technique that controls the response to such a risk. In [3], tactics related to separate and hide can be applied to reduce the likelihood of the risk of a privacy harm, whereas tactics related to abstract and minimise can be applied to reduce the severity of the risk of a privacy harm. To choose the most relevant tactic, it is important to consider the type of operations (collection, retention, disclosure, etc.) performed on

personal data. We refer to the APDL model of Chapter 4 to identify the type of operations under consideration. This helps identify appropriate and concrete tactics, as each tactic addresses specific types of operations. Having chosen the appropriate abstract tactic, associated concrete tactics can be analysed with respect to their underlying goals and the constraints imposed on the types of processing operations to determine which concrete tactics can control the response in a way that achieves the desired protection goal. Repeating this step helps identify the most effective tactics that can be used in combination to achieve all the desired protection goals.

Based on the underlying goals of process-oriented tactics, they are more related to privacy violations than privacy harms. Demonstrate tactics are associated with Data Protection Authorities to demonstrate privacy compliance. Enforce tactics are associated with data controllers and processors to enforce polices, whereas inform and control tactics are associated with data subjects to control their personal data through data controllers who inform them about the collection, processing and dissemination of their personal data. To choose from process-oriented tactics, it is important to determine whether the aim is to: inform data subjects; provide them with means to exercise control over their personal data; enforce policies and technical controls; or demonstrate privacy compliance with legal frameworks and standards, as per Figure 6.8.

Example 6.2. The desired privacy protection goal #1 is an instance of unlinkability as an abstract protection goal. In addition, data integration is more concerned with the data itself. According to the structure of privacy protection tactics (shown in Figure 6.3), minimise, hide, separate and abstract are the most relevant abstract tactics to achieve the desired protection goal. By referring to the risk level of this privacy harm in Chapter 5, the likelihood is *maximum*, whereas its severity is *limited*. In this case, the treatment strategy is to reduce the likelihood of the risk. This means that separate and hide can be applied to reduce the likelihood of the risk of the privacy harm. Before analysing their associated tactics, the main type of operation performing in data is *data integration*. Hide aims to prevent exposure of access, association, visibility and understandability of personal data, to reduce the

likelihood of a privacy threat occurrence. Separate aims to prevent the correlation of personal data to reduce the likelihood of a privacy threat occurrence by distributing or isolating processing operations on personal data. As such, hide's tactics are more relevant to achieve the desired protection goal. Restrict is an appropriate tactic as it aims to prevent unauthorised access to personal data. Accordingly, restrict is a concrete tactic that prevents the data inference that can be made by linking initiative data to other types of signatories' data.

6.3.3 The Selection of Appropriate Design Patterns

Once the most relevant architectural tactic has been identified, appropriate design patterns that describe the tactic need to be identified. In the context of software engineering, architectural tactics are often described through design patterns, whereby a design pattern may describe multiple tactics — whether they are desired or not [52]. Some of these tactics may relate to other quality attributes, such as security. In addition, a design pattern may describe one or more abstract tactics. It may also describe more than one concrete tactic. In contrast, a design pattern can be described in terms of its tactics. To identify appropriate design patterns that describe the identified tactic, privacy design patterns need to be analysed with respect to their intent and/or problem, motivation, forces and context, and solution, as per Figure 6.9.

The intent and/or problem of a privacy design pattern are both statements of the problem, which focus typically on what is being protected. As such, the underlying goal of the abstract tactic and the definition of its concrete tactic can be mapped to the intent and/or problem of a privacy design pattern, as per Figure 6.10. By mapping these factors, it is possible to decide whether or not the intent and/or problem of a design pattern abstractly reflects the underlying goal of the architectural tactic and the constraints on operations performed on personal data.

The quality scenarios can be considered as characterisations of quality attribute requirements as they are generated from a reasonable set of threat scenarios. Based on this, a quality scenario can be used as an indicator to the abstract scenario

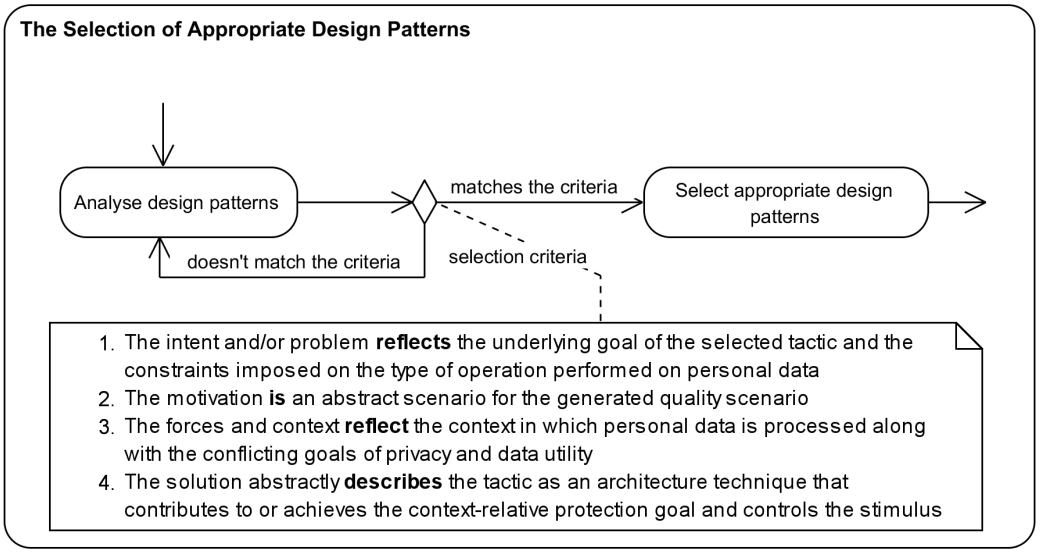

Figure 6.9: The main steps of selecting appropriate design patterns.

of a design pattern. The motivation of a design pattern illustrates the problem by giving an abstract example. As such, the quality scenario can be mapped to the motivation of a design pattern, as per Figure 6.10. By mapping these factors, it is possible to decide whether or not the motivation of a pattern is an abstract scenario for the generated quality scenario.

The quality scenario (especially the environment element) can also be used as an indicator to the context in which personal data is processed, along with non-functional aspects, such as performance, usability, etc. The forces and context of a design pattern illustrate the context in which personal data is processed, along with the conflicting goals. As such, the quality scenario can be mapped to the forces and context of a design pattern, as per Figure 6.10. By mapping these factors, it is possible to decide whether or not the forces and context of a design pattern reflect the conflicting goals of privacy and data utility.

The quality scenario (especially the response element) can be used as an indicator to the solution proposed by a design pattern. The solution of a design pattern is an abstract description that is sufficiently complete to achieve concrete design goals. As such, the quality scenario can be mapped to the solution of a design pattern, as per Figure 6.10. By mapping these factors, it is possible to decide

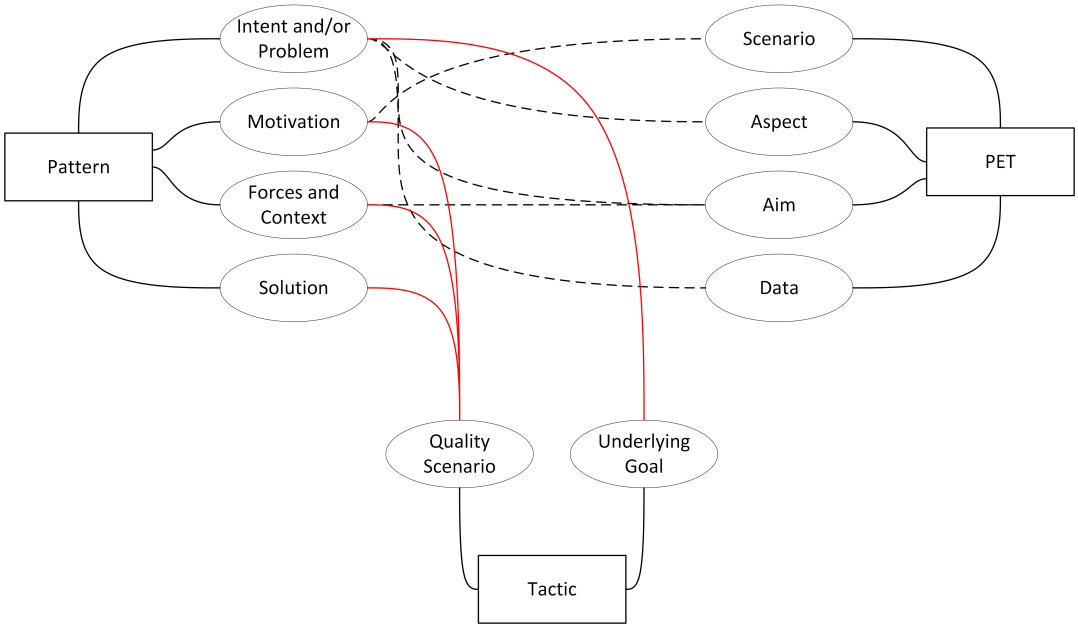

Figure 6.10: The dependencies between tactics, patterns and PETs.

whether or not the solution of a pattern abstractly describes the tactic as an architecture technique that contributes to or achieves the context-relative protection goal and controls the stimulus. Such a technique may abstractly specify a specific approach, such as encrypting data.

Example 6.3. The underlying goal of the restrict tactic is to prevent exposure. The imposed constraint on the usage operation is to prevent unauthorised access to personal data. The Attribute Based Credentials pattern describes the underlying goal of restrict tactic. Its intent is to flexibly and selectively authenticate different attributes about an entity without revealing additional data about the entity. Its motivation gives an abstract scenario from the context of the internet and smart cards that reflects the generated quality scenario (quality scenario #1 of Example 6.1) in an abstract fashion. It can be applied in any context in which the collection and processing of personal data for legitimate purposes may pose threats to the privacy of data subjects. Its solution abstractly describes the restrict tactic by preventing data exposure.

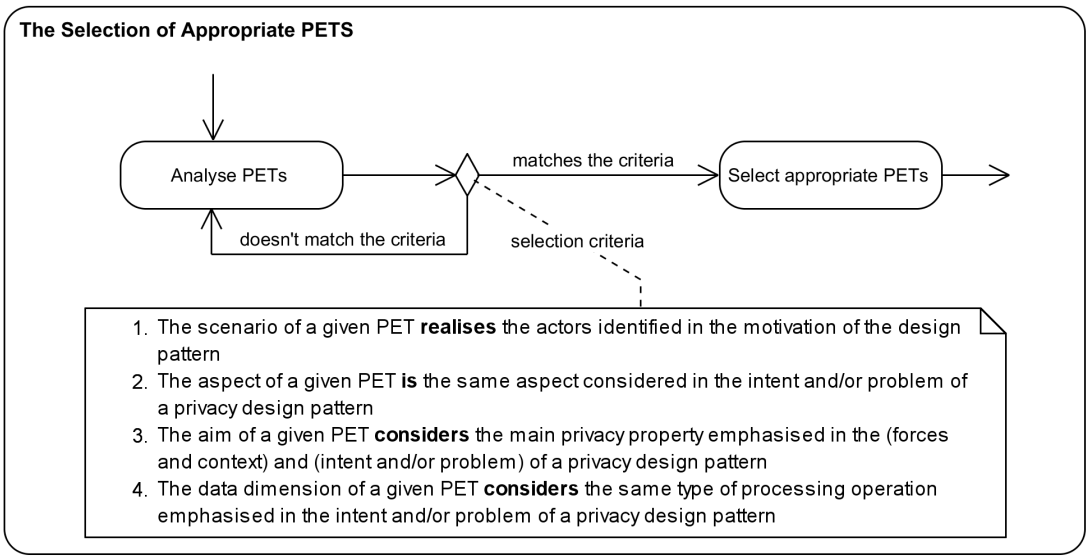

Figure 6.11: The main steps of selecting appropriate PETs.

6.3.4 The Selection of Appropriate PETs

Once the most relevant design pattern has been identified, appropriate PETs that implement this pattern need to be identified. This can be achieved by analysing existing PETs with respect to their scenario, aspect, aim and data dimensions, as per Figure 6.11.

Each privacy design pattern takes into account who personal data is being protected from. These types of threat sources can be determined on the basis of their relationships with the data subject. The motivation of a privacy design pattern is an abstract scenario that illustrates the problem. Based on this, the motivation of the design pattern can be used as an indicator to the threat model around which a PET is developed. The scenario dimension of a PET defines the main actors involved in the processing of personal data and categorises them from a trust perspective with the aim of identifying potential threat sources. As such, the motivation of the pattern can be mapped to the scenario of the given PET, which, in turn, needs to reflect the generated quality scenario, as per Figure 6.10. By mapping these factors, the main actors involved in the quality scenario can be identified and categorised from a trust perspective — whether they are untrusted service providers, untrusted data subjects or both (mutual).

The problem focuses typically on what is being protected; it may also focus on non-functional aspects and normative provisions. Based on this, the intent and/or problem of a privacy design pattern can be used as an indicator of the degree of anonymity required. The aspect dimension of a PET complements the scenario dimension by defining the main target of the PET (identify, content or behaviour). As such, the intent and/or problem of a privacy design pattern can be mapped to the aspect of the given PET, which, in turn, needs to reflect the underlying goal of the corresponding tactic, as per Figure 6.10. By mapping these factors, the main target of the underlying goal can be derived and/or identified.

The forces and context of a design pattern illustrate the context in which personal data is processed, along with the conflicting goals, such as performance, usability, etc. Based on this, the forces and context of a design pattern can be used as an indicator of the non-functional aspects. This factor, along with the intent and/or problem of a privacy design pattern, can be used as indicators of the main aspects that need to be considered to achieve an appropriate level of privacy protection. The aim dimension of a PET defines how privacy is achieved in terms of indistinguishability, unlinkability, deniability and confidentiality. As such, the forces and context and intent and/or problem of a privacy design pattern can be mapped to the aim of the given PET, as per Figure 6.10. By mapping these factors, the main target of the underlying goal can be derived and/or identified and the normative goals can be articulated.

The intent and/or problem of a privacy design pattern can also be used as an indicator of the state of personal data addressed by the PET. The data dimension of a PET describes the states of data that are addressed by a PET (data-at-rest or data-in-motion). As such, the intent and/or problem of a privacy design pattern can be mapped to the data dimension of the given PET, as per Figure 6.10. By mapping these factors, the main types of operations, along with the constraints that need to be imposed on such operations, can be identified.

The PET that matches these factors can then be chosen as a concrete technology that implements the design pattern, which, in turn, describes the identified architectural tactic to achieve the context-relative protection goal.

Example 6.4. By selecting Attribute Based Credentials as a privacy design pattern that describes the restrict tactic by the separation of data, an appropriate PET that realises this pattern needs to be chosen. By analysing existing PETs, IDEMIX is the most relevant PET that implements the selected design pattern.

The scenario dimension of the PET realises the actors identified in the motivation of the design pattern (user, registration authority and service provider). In this case, we choose the 'mutual' scenario to describe the case of signatories, organisers and registration authorities are not trusting each other.

The aspect dimension of the PET is the same aspect considered in the intent of the design pattern (the identity). In this case, the identity of signatories are not required when signing up to a particular initiative. This indicates that anonymity is the best option. The focus is on how to protect the identity. As such, the technical measure should hide the identify of signatories by authenticating different attributes about an entity without revealing additional data about the entity.

The aim dimension of the PET considers the privacy property emphasised in the intent of the design pattern. In this case, the focus is on unlinkability to ensure that 'initiative data' is not linkable with 'identification and contact data'.

The data dimension of the PET considers the same type of operations performed on personal data. In this case, the focus is on the collection and processing of signatories' personal data to ensure that data is processed without revealing the identify of signatories.

6.3.5 The Identification of Architectural Strategies

Once appropriate architectural tactics, design patterns and PETs have been chosen, architectural strategies that combine these artefacts need to be identified to address each quality scenario. This can be achieved by categorising the chosen tactics, along with patterns and PETs, according to the desired protection goals, and

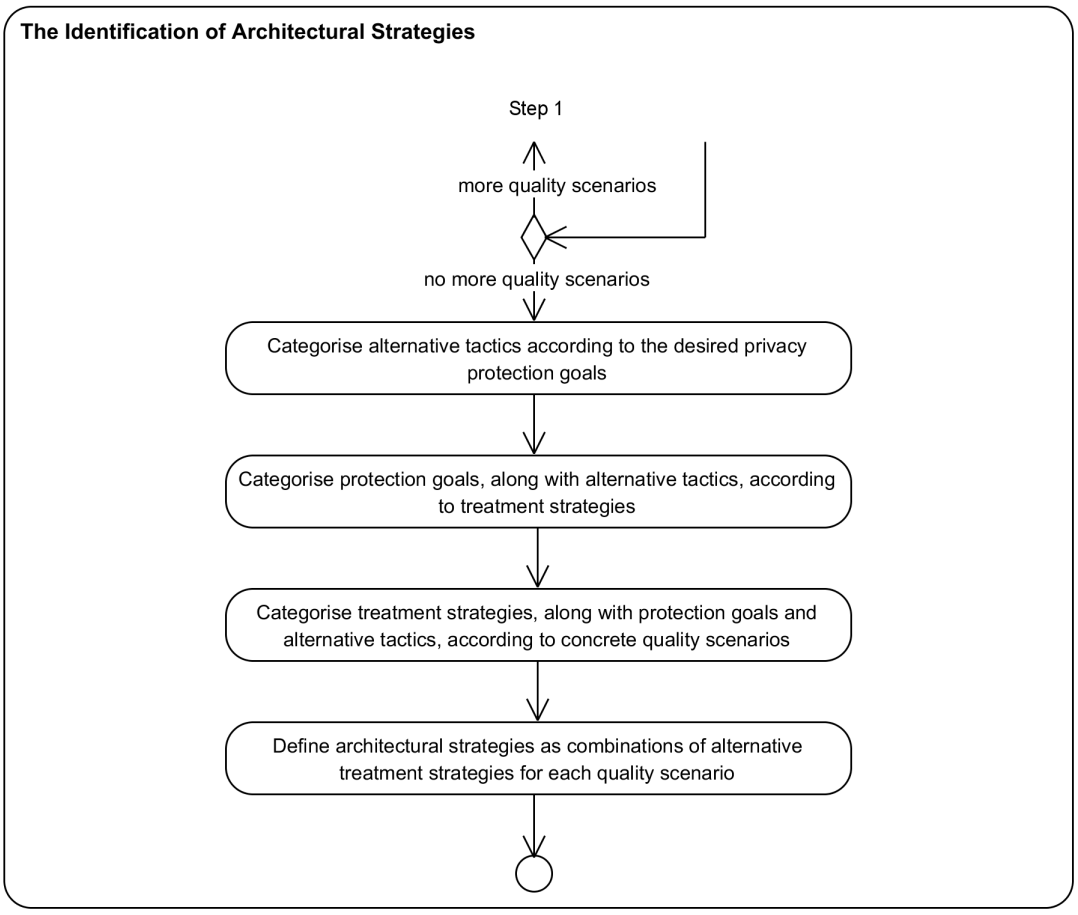

Figure 6.12: The main steps of identifying architectural strategies.

categorising these goals, along with tactics and corresponding patterns and PETs, according to the generated quality scenarios, as per Figure 6.12. In the context of software engineering, architectural strategies are collections of tactics to achieve multiple quality attributes. For the purpose of this dissertation, privacy architectural strategies are collections of architectural tactics that achieve privacy protection goals. Since a design pattern may implement multiple tactics, and a tactic may be described through multiple design patterns, each architectural strategy is a combination of a particular tactic, design pattern and PET. As such, multiple strategies may address the same quality scenario based on the identified treatment strategies, which describe risk-treatment options in a commensurate manner. These options are described through architectural tactics that meet the relevant response measure (the acceptable level of the corresponding risk). As such, architectural strategies are

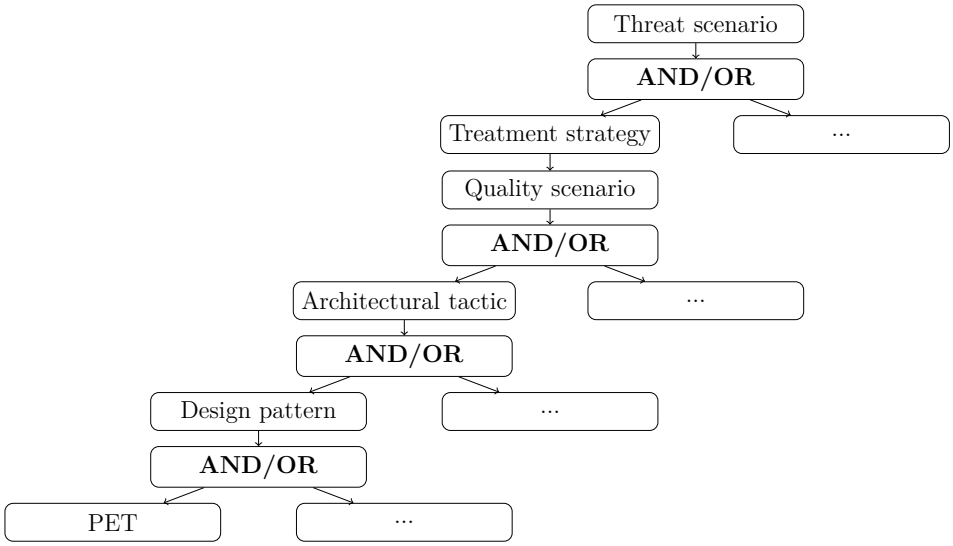

Figure 6.13: The structure of the strategy tree.

considered as architectural options. These options are useful during the analysis and evaluation of the software architecture with multiple stakeholders' participation to make architectural choices that determine the relative levels of privacy protection.

To represent the identified strategies in a way that is analytically useful, we define a concept of a strategy tree. A strategy tree describes the relationship between a threat scenario (a root node) and all possible treatment strategies (intermediate nodes) upon which quality scenarios (intermediate nodes) are generated that can be achieved by architectural tactics (intermediate nodes), which are relaised by design patterns (intermediate nodes) and implemented by PETs (leaf nodes), as illustrated in Figure 6.13. We use 'AND' and 'OR' connectors to combine child nodes and indicate whether all or some child nodes are necessary to enact the parent node.

Example 6.5. Once the architectural tactic is identified, along with its corresponding privacy design pattern and underlying PET, appropriate architectural strategies need to be identified. In this case, we identified only one architectural tactic, design pattern and PET. Accordingly, only one architectural strategy can be identified: the 'anonymous credential strategy', which aims to conceal signatories' identities, while revealing other data.

Architectural strategies are *risk mitigation strategies*. As such, the identified strategy aims to reduce the likelihood of data integration that may contribute to the

occurrence of the risk of the privacy harm. It also aims to address the exploitability of primary assets' vulnerabilities and the adverse effects of data integration that lead to the identification of signatories that may reveal sensitive information, such as religious beliefs, political affiliation, etc.

6.4 Summary

In this chapter, we have presented a tactical approach that provides insight into the process through which existing privacy design strategies, architectural tactics, design patterns and PETs can be applied in practice. It is a hybrid approach that combines risk-based and goal-oriented approaches to address potential privacy risks and meet regulatory compliance requirements in the early stages of the design process (at an architectural level).

The tactical approach uses the concept of quality attributes to characterise privacy protection by means of a general quality scenario to specify the responses of the system to realise desired privacy protection goals. The general quality scenario specifies the range of values its elements can take. As such, it is used as the basis for generating concrete scenarios from which desired protection goals can be derived.

The tactical approach categorises existing tactics based on a set of privacy protection goals that address legal, organisational, societal and technical aspects of privacy, and organises these tactics in a hierarchical manner that facilitates the representation of both abstract and concrete tactics. The adoption of these protection goals facilitates the translation of abstract privacy principles of legal frameworks and standards into context-relative privacy protection goals, which, in turn, facilitate compliance checking.

The approach also adopts the concept of architectural strategies as collections of appropriate privacy architectural tactics, which are described through design patterns and realised by PETs, to address quality scenarios by achieving the corresponding desired protection goals. In particular, each architectural tactic specifies fundamental design decisions that contribute to, or achieve, a desired protection goal. As such, architectural strategies is used as a means for mapping

privacy requirements onto suitable software architectures to specify, implement and justify various levels of privacy protection. In addition, the approach provides a set of selection/mapping criteria to aid software engineers to give consideration to what combination of tactics, design patterns and PETs will achieve, or contribute to the achievement of, the desired protection goals.

In the next chapter, we will illustrate our principled approach for engineering PbD. In particular, we will explain how the privacy-aware data lifecycle models of Chapter 4, the data-centric threat modelling of Chapter 5 and the privacy-enhancing strategies of Chapter 6 can be used in combination to capture privacy concerns in a comprehensive manner, address these concerns at an architectural level, and reason about the compliance of architectural choices with legal frameworks and standards.

Engineering systems with privacy in mind requires integrating privacy requirements into the typical systems engineering activities.

— Seda Gürses, Carmela Troncoso and Claudia Diaz

7 An Approach for Engineering PbD

Contents

7.1 Introduction

This chapter gives a relatively detailed description of a principled approach for engineering PbD, which synthesises the contributions of Chapter 4, Chapter 5 and Chapter 6. The principled approach can complement software engineering methods with the aim of capturing privacy concerns in a comprehensive manner, addressing these concerns at an architectural level and reasoning about the compliance of

architectural choices with legal frameworks and standards. In particular, Section 7.2 describes the main steps of the principled approach that adopt: the UML profile for the APDL model as a means for representing data-processing in a way that is amenable for risk analysis and compliance checking; the data-centric threat modelling approach for identifying and assessing potential privacy risks in a comprehensive and contextual manner; and privacy-enhancing strategies as means for mapping privacy requirements to suitable software architectures to specify, implement and justify various levels of privacy protection. Section 7.3 illustrates how the main activities of our approach are mapped onto the requirements analysis and design phases of the software development lifecycle. Then, Section 7.4 provides some examples from the illustrative case study (the European Electronic Toll Service (EETS)) of Chapter 3 to demonstrate the usefulness and applicability of the principled approach in this particular context. Finally, Section 7.5 gives a brief summary of the chapter.

7.2 An Approach for Engineering PbD

Designing software-based systems according to the principles of PbD requires incorporating privacy requirements into engineering activities that are related to each phase of the software development lifecycle. Our principled approach supports the effective translation of abstract privacy principles, privacy risk models and privacy mechanisms into implementable requirements. It also support the integration of these activities into the system development lifecycle. It is intended to aid software engineers in analysing functional requirements, eliciting privacy requirements, making appropriate design decisions that fulfil these requirements, implementing the design decisions, and reasoning about the compliance of design decisions with legal frameworks and standards in a structured manner. It is principled as it: adopts the GPS principles; supports the translation of abstract privacy principles stated in legal frameworks and standards into context-relative protection goals; and translates the principles of PbD into engineering activities. It consists of three activities, each of which consists of a number of steps, as per Figure 7.1. We consider each in turn.

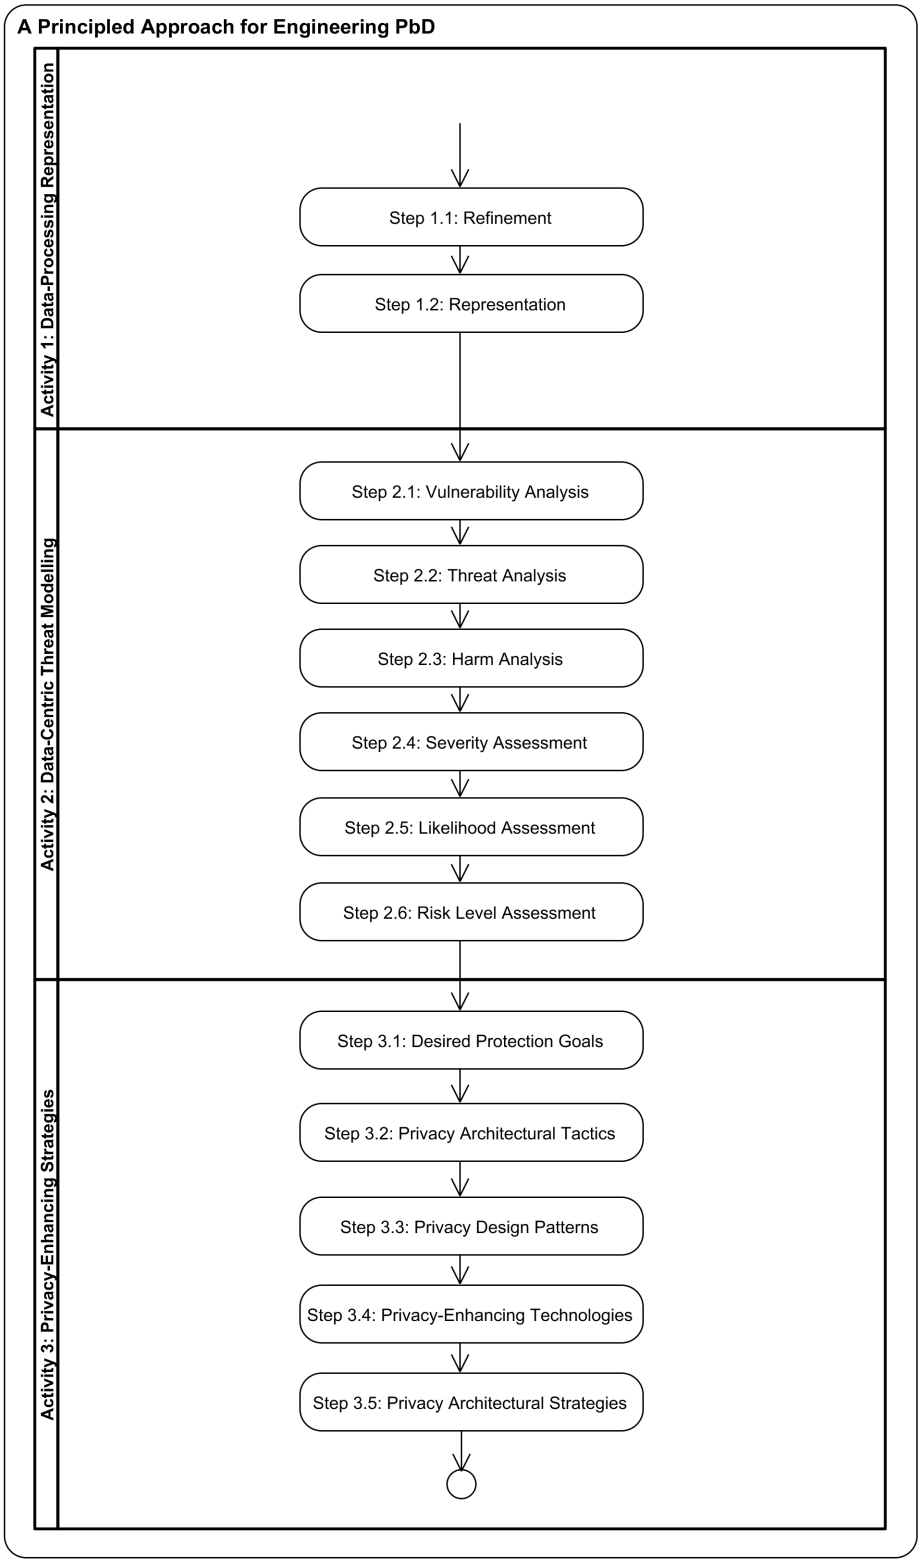

Figure 7.1: The main activities of our approach.

7.2.1 Data-Processing Representation

This activities aims to establish the context in which personal data is collected, processed and disseminated, and represent data-processing activities in a way that is amenable to analysis. It consists of two steps: refinement and representation. As previously discussed in Chapter 4, the refinement step aims to operationalise the abstract purpose by refining it into a set of concrete purposes that can be assigned to actors as responsibilities. The representation step aims to model the abstract and concrete purposes together with the key aspects of abstract privacy principles as a requirements model.

The output of this activity is a *requirements model* that is represented using the UML profile for the APDL model.

7.2.2 Data-Centric Threat Modelling

This activity aims to identify, analyse and assess potential privacy risk in a comprehensive and contextual manner. As previously discussed in Chapter 5, it consists of six steps. The first step is vulnerability analysis, which aims to identify and analyse all possible privacy vulnerabilities in the established context. The second step is threat analysis, which aims to identify and analyse potential threat sources and events. The third step is to identify and analyse potential privacy violations and harms. It also aims to construct harm trees for each particular privacy harm to be used as the basis for generating reasonable sets of threat scenarios. The fourth step is severity assessment, which aims to assess the severity of a privacy harm based on the intensity of adverse effects of a threat event and the range of these effects suffered by a variety of data subjects. The fifth step is likelihood assessment, which aims to assess the likelihood of occurrence of a privacy harm as the highest value of the overall likelihood of occurrence of associated threat events. The sixth step is risk level assessment, which aims to assess the level of the risk of a privacy harm as a pair of (likelihood, severity).

The input of this activity is a *requirements model* from the previous activity. The output of this activity is a set of *harm trees* (used to generate threat scenarios for each privacy harm) and a *risk map*.

7.2.3 Privacy-Enhancing Strategies

This activity aims to provide various levels of privacy protection. As previously discussed in Chapter 6, it consists of five steps. The first step is to articulate context-relative (desired) privacy protection goals. The second step is to select appropriate architectural tactics that contribute to, or achieve, the desired protection goals. The third step is to select appropriate privacy design patterns that describe the identified architectural tactics. The fourth step is to select appropriate PETs that implement the identified design patterns. The fifth step is to identify architectural strategies that serve as architectural choices.

The input of this activity is a set of *threat scenarios* and a *risk map* from the previous activity. The output of this activity is a set of *architectural strategies* that specify architectural choices.

7.3 Software Development Lifecycle: Waterfall Model Stages

In software engineering, a software development process, also known as a software development lifecycle, which is a subset of the system development lifecycle (SDLC), is the process of dividing software development into distinct phases. These phases vary between SDLC models, which include the waterfall model, spiral model and Agile model. For simplicity, we consider the waterfall model. It is a sequential development approach, in which development is seen as flowing steadily downwards through several phases, typically: requirements analysis, design, implementation, testing, integration, deployment and maintenance. In this dissertation, we focus only on the *requirements analysis* and *design* phases of the software development process as represented in the waterfall model with the aim of mapping the main activities of our principled approach onto these phases, as per Figure 7.2.

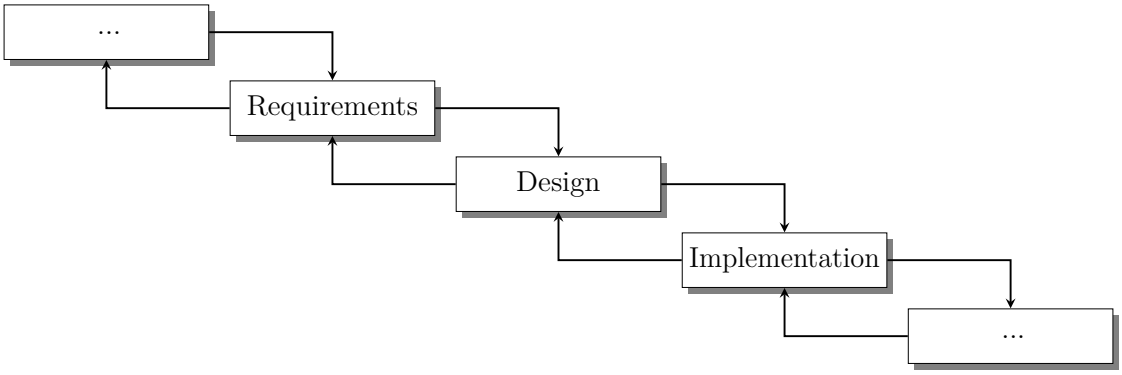

Figure 7.2: The main phases of the software development process represented in the waterfall model.

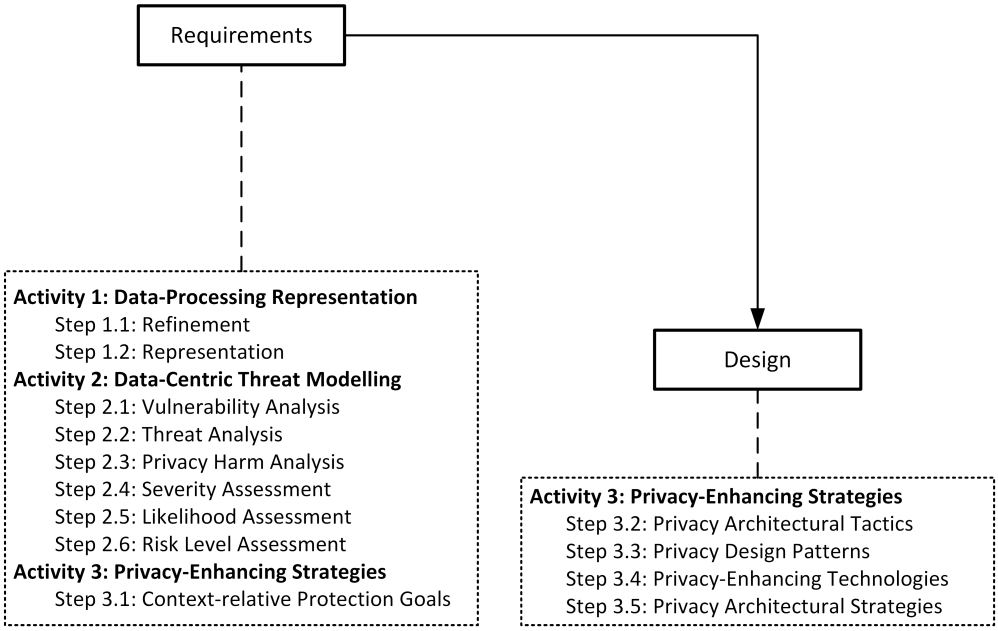

Figure 7.3: Mapping of the activities of our approach to the phases of the SDLC.

In this section, we initially convey the three contributions (which are considered as privacy-related models and approaches) such that they are actionable across the SDLC. Figure 7.3 illustrates how the main activities of our approach are mapped against the requirements analysis and design phases of the SDLC.

7.3.1 Requirements Analysis

The *UML profile for the APDL model* is a technique for analysing, specifying and structuring abstract purposes and data-processing activities (concrete purposes). It defines an abstract purpose as a statement of intent that prescribes the goal for

which personal data is collected, processed and disseminated. It operationalises abstract purposes by processing operations that are expressed in terms of data items, data-processing activities that consist of concrete actions and events that cause the execution of these actions, and roles that define these activities as responsibilities performed by difference actors according to their capabilities. This means that data-processing activities are assigned to actors as responsibilities. Actors are external or internal entities that are capable of, and responsible for, performing the activities of the role to which they are assigned. In this nomenclature, a concrete purpose is defined as a data-processing activity under the responsibility of an actor. Consequently, a requirement is a refinement of a data-processing activity. The first activity of our approach consists of two distinct steps: refinement and representation. The first step aims to operationalise abstract purposes, whereas the second step aims to represent data-processing activities. Thus, this activity can be mapped onto the *requirements analysis phase*, which aims to reach a better understanding of each functional requirement and represent requirements in a well-structured manner that facilitates effective communication with multiple stakeholders, as per Figure 7.3.

The *harm tree* is a graph-based analysis technique that is used for generating a reasonable set of threat scenarios for each privacy harm. In particular, it describes how threat events caused by the most likely threat sources with capability to exploit possible vulnerabilities can contribute to or cause a privacy harm. Further, the *risk map* is a technique that is used to allocate the assessed privacy risks according to their levels (likelihood and severity). It is also used to: determine the order in which the identified risks should be managed based on their levels; and prioritise the ordered risks in each level. In addition, the *general quality scenario* is a technique that is used to characterise privacy protection as a quality attribute by a set of elements that can be mapped onto the factors of threat scenarios. Quality scenarios are used to derive context-relative privacy protection goals. Consequently, a privacy non-functional requirement is a refinement of a quality scenario: it specifies the response of the software to realise a privacy protection goal. The second activity of our approach consists of six distinct steps. These steps aim to generate all

possible threat scenarios for each privacy harm and locate the assessed privacy risks on the risk map. The first step of the third activity aims to generate quality scenarios, from which context-relative protection goals are derived. Thus, the second activity and the first step of the third activity can be mapped onto the *requirements analysis phase*, which aims to reach a better understanding of each non-functional requirement and represent requirements in a well-structured manner that facilitates effective communication with multiple stakeholders, as per Figure 7.3.

7.3.2 Software Design

The software design phase includes both high-level and low-level design. In this dissertation, we focus only on high-level design.

The *architectural tactics* are primary mechanisms used to achieve the desired privacy protection goals and to meet the captured privacy non-functional requirements. In particular, tactics are used to create an architectural design using *design patterns* and *architectural strategies*. Consequently, tactics are fundamental design options that guide the architectural decisions. The third activity of our approach consists of five distinct steps. The first step is related to the requirements analysis phase, whereas other steps aim to identify architectural tactics that are described through design patterns and grouped into architectural strategies. These tactics are chosen to meet privacy non-functional requirements. Thus, the third activity can be mapped onto the *design phase*, which aims to solve privacy-related design problems and describe software solutions at a high-level of abstraction, as per Figure 7.3.

7.4 The Case Study

In this section, we illustrate our approach by means of the case study of Chapter 3 (the European Electronic Toll Service (EETS)). The illustration is conducted through three main activities. The first activity is to represent data-processing activities in a way that is amenable to risk analysis and compliance checking. The second activity is to generate a reasonable set of threat scenarios for each potential privacy harm. It is also to identify, analyse and assess potential privacy risks and locate them on the risk

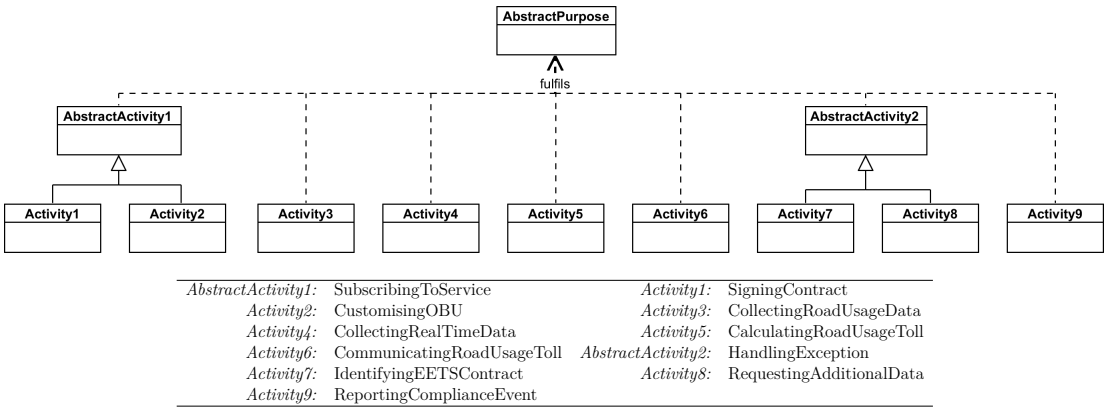

Figure 7.4: The refinement of the abstract purpose of EETS.

map. The third activity is to identify a set of architectural strategies that provide various levels of privacy protection. The case study has several privacy-related aspects; however, in this section, we will focus on those related to the collection and usage of EETS users' personal data. Other aspects are illustrated in Appendix B.

7.4.1 Data-Processing Representation

Refinement

The **abstract purpose** of EETS *is to calculate personalised road-usage tolls and communicate the final premium to EETS users at the end of the tax period.* In order for this purpose to be operationalised, it first needs to be refined into a set of concrete purposes (data-processing activities) according to the stages through which EETS users' personal data moves during its lifecycle. In this case, the abstract purpose is refined into nine concrete purposes, as per Figure 7.4.

SigningContract is represented as *Activity1*, which describes the activity of collecting a user's personal data to subscribe a contract with a EETS service provider. CustomisingOBU is represented as *Activity2*, which describes the activity of collecting a vehicle's characteristics to be used for initialising the OBU with the vehicle's classification parameters and toll context data for tolls calculation. *Activity1* and *Activity2* are generalised by an abstract activity SubscribingToService, which is represented as *AbstractActivity1*. AbstractActivity1 is classified as a collection stage of the data lifecycle to describe processing activities that are considered

logically-related responsibilities that can be assigned to: a new user who plays the role of an unsubscribed user as a type of the data subject role; and a service provider who plays the role of a contract agent as a type of the data processor role.

CollectingRoadUsageData is represented as *Activity3*, which describes the activity of collecting toll declaration data (including location data) to be used for toll calculation. Activity3 is classified as a collection stage of the data lifecycle to describe processing activities that are considered logically-related responsibilities that can be assigned to a service provider who plays the role of a collection agent as a type of the data processor role.

CollectingRealTimeData is represented as *Activity4*, which describes the activity of collecting charging and enforcement data (including location data and vehicle classification parameters) to be used for enforcement and compliance checking. Activity4 is classified as a collection stage of the data lifecycle to describe processing activities that are considered logically-related responsibilities that can be assigned to a toll charger who plays the role of an enforcement authority as a type of the data controller role.

CalculatingRoadUsageToll is represented as *Activity5*, which describes the activity of using toll declaration data for toll calculation according to the relevant toll context data. Activity5 is classified as a usage stage of the data lifecycle to describe processing activities that are considered logically-related responsibilities that can be assigned to a service provider who plays the role of a contract agent as a type of the data processor role.

CommunicatingRoadUsageToll is represented as *Activity6*, which describes the activity of sending claims (by means of invoices) for the final premium at the end of the tax period to the EETS users. Activity6 is classified as a usage stage of the data lifecycle to describe processing activities that are considered logically-related responsibilities that can be assigned to a service provider who plays the role of a contract agent as a type of the data processor role.

IdentifyingEETSContract is represented as *Activity7*, which describes the activity of using a licence plate (registration number) to identify a EETS provider's contract

and its corresponding EETS user for enforcement management. RequestingAdditionalData is represented as *Activity8*, which describes the activity of requesting additional parameters for justifying billing details and payment guarantee for an inferred object. *Activity7* and *Activity8* are generalised by an abstract activity HandlingException, which is represented as *AbstractActivity2*. AbstractActivity2 is classified as an access stage of the data lifecycle to describe processing activities that are considered logically-related responsibilities that can be assigned to a toll charger who plays the role of an enforcement authority as a type of the data controller role.

ReportingComplianceEvent is represented as *Activity9*, which describes the activity of reporting compliance check communication events to EETS providers for monitoring and customer care reasons. *Activity9* is classified as a usage stage of the data lifecycle to describe processing activities that are considered logically-related responsibilities that can be assigned to a toll charger who plays the role of an enforcement authority as a type of the data controller role.

Once the abstract purpose is refined into concrete activities, the next step is to express each activity in terms of actions and events that trigger the execution of these actions. Figure 7.5 shows the operationalisation of *Activity1* and *Activity2* in terms of actions, events, roles and involved actors. The operationalisation of *Activity3* – *Activity9* is illustrated in Section B.1.1 of Appendix B.

Activity1 coordinates its execution via one action and one corresponding event. CollectUserData is represented as *Action1*. Sign is represented as *Event1* to specify the occurrence of signing a contract with the service provider to collect a user's personal data. *Activity2* coordinates its execution via one action and one corresponding event. CollectVehicleData is represented as *Action2*. Configure is represented as *Event2* to specify the occurrence of customising the OBU to collect the vehicle's data.

Having described the activities in terms of actions and events, the next step is to describe these in terms of roles and actors.

Driver, represented as *Actor1*, is an actor who is capable of, and responsible for, performing the activities of the unsubscribed user as a role to which is assigned.

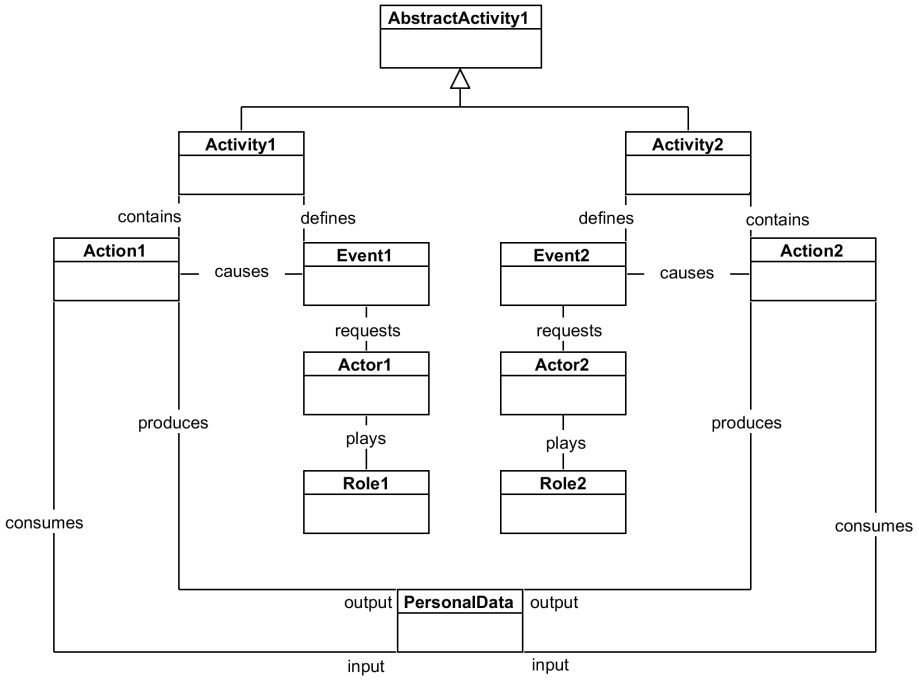

AbstractActivity1:	SubscribingToService	*Activity1:*	SigningContract
Activity2:	CustomisingOBU	*Action1:*	CollectUserData
Event1:	Sign	*Action2:*	CollectVehicleData
Event2:	Configure	*Actor1:*	Driver
Role1:	UnsubscribedUser	*Actor2:*	EETSProvider
Role2:	ContractAgent	*PersonalData:*	EETSUser, Vehicle

Figure 7.5: The refinement of the AbstractActivity1.

UnsubscribedUser, represented as *Role1*, is a specialisation of the data subject role that involves a set of logically-related activities for subscribing to the EETS service.

EETSProvider, represented as *Actor2*, is an actor who is capable of, and responsible for, performing the activities of the contract agent and/or the collection agent as roles to which is assigned. ContractAgent, represented as *Role2*, is a type of the data processor role that involves a set of logically-related activities for customising, installing and verifying OBUs. CollectionAgent, represented as *Role3*, is a type of the data processor role that involves a set of logically-related activities for: collecting toll declaration data using OBUs; and calculating and communicating road-usage tolls.

TollCharger, represented as *Actor3*, is an actor who is capable of, and responsible for, performing the activities of the enforcement authority as a role to which it is assigned. EnforcementAuthority, represented as *Role4*, is a type of the data controller role that involves a set of logically-related activities for managing enforcement

using fixed or mobile RSE.

Representation

Having operationalised the abstract purpose of EETS into a set of concrete purposes (data-processing activities) that are expressed in terms of activities, actions, events, roles and actors, these processing activities can be represented using the UML profile for the APDL model.

AbstractPurpose, PersonalData and DataModel. EETSPurpose is a class stereotyped by «AbstractPurpose» to represent the abstract purpose for which EETS users' personal data is collected and processed. At the domain level, its tagged values informally describe the aim of the stereotyped class, as well as the fairness, lawfulness and proportionality of the purpose, and indicate the relevant GPS principle, as illustrated in Figure 7.6. At the instance level, the value of its actualPurpose attribute is to calculate personalised road-usage tolls and communicate the final premium to EETS users at the end of the tax period.

In order for the specified purpose to be fulfilled, a minimum amount of required data needs to be appropriately specified to support the declaration and calculation of road-usage tolls. Such a specification requires building a data model that represents the relevant objects, their properties and relationships in relation to the context of EETS. Accordingly, we draw a partial data model diagram that represents the following classes: EETSUser, LocationData, UserAccount, Contract, Vehicle, OBU and TollContextData, as illustrated in Figure 7.6. Classes that represent personal data are stereotyped by «PersonalData», whereas those represent generic data are not stereotyped. The relationships in which these classes participate can of course be modelled directly by associations in the UML.

As an example, EETSUser is a class stereotyped by «PersonalData» to represent the user's personal data, whether the user is the driver, owner or fleet operator of the vehicle. At the domain level, its tagged values informally describe the aim of the class, the category and the type of the data, as illustrated in Figure 7.6. Its primary attributes are userId, name and billingAddress.

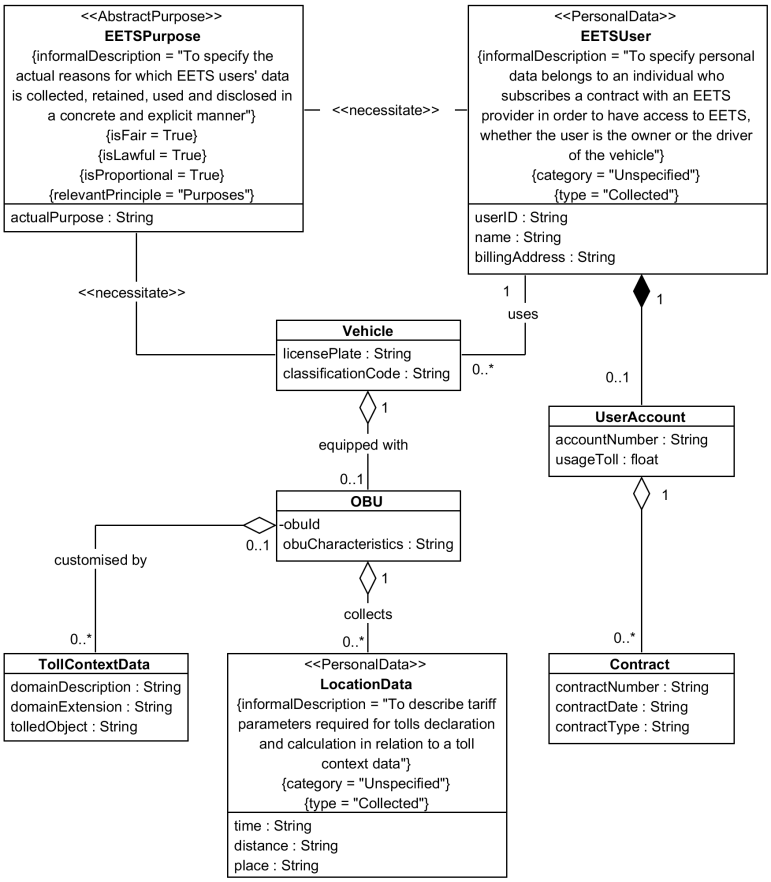

Figure 7.6: EETSPurpose, along with the partial data model diagram for EETS.

The types of personal can be classified as follows:

- Identification and contact data — ***EETSUser:*** user ID, name, billing address (collected from the EETS user, whether the user is the driver, owner 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)

DataLifecycle, LifecycleStage and its Specialisations. In reference to the general architecture of the EETS proposed by the European Commission in Chapter 3, the required personal data needs to be specified by Member States in relation to relevant national regulations. This indicates that the EETS data

lifecycle is closed, i.e. no arbitrary data from external sources will be collected, acquired or derived without initial planning. In accordance with the physical nature of the EETS architecture, the collected EETS users' personal data is not stored and/or retained in a centrally controlled infrastructure. This indicates that the EETS data lifecycle is decentralised.

EETSDataLifecycle is a class stereotyped by «DataLifecycle» to represent the main characteristics of the personal data lifecycle in the context of EETS. At the domain level, its tagged values informally describe the aim of the class, the openness of the processed data and the centrality of its underlying system, as illustrated in Figure 7.7. At the instance level, the values of its lifecycleType and granularityLevel attributes are Evolutionary and FinedGrained respectively.

Each lifecycle stage can be represented by one or more stereotyped classes that represent sets of activities according to: the types of processed data, its sources, the manner in which it is processed, and the assigned roles and responsibilities of the actors. This gives each lifecycle stage the ability to be expressed in terms of cycles, reflecting the repetitive nature of data processing activities. For example, SubscribingToService and CollectingRoadUsageData are two classes of the collection stage, separated according to the nature of collection activities and responsibilities of the involved actors. To subscribe to an EETS service, drivers, as data subjects, provide their personal data by establishing and signing contracts with the service providers. To collect usage data, EETS service providers, as data processors, collect location data of subscribed EETS users via OBUs.

SubscribingToService is a class stereotyped by «Collection» to represent a set of related activities, i.e. signing a contract and customising an OBU, with the aim of collecting users' personal data and vehicles' classification parameters. At the domain level, its tagged values informally describe the aim of the class, data sources, collection methods, available choices, consent type and the relevant GPS principles, as illustrated in Figure 7.7. At the instance level, the values of its createdUserData and createdVehicleData attributes are of the type EETSUser and Vehicle respectively.

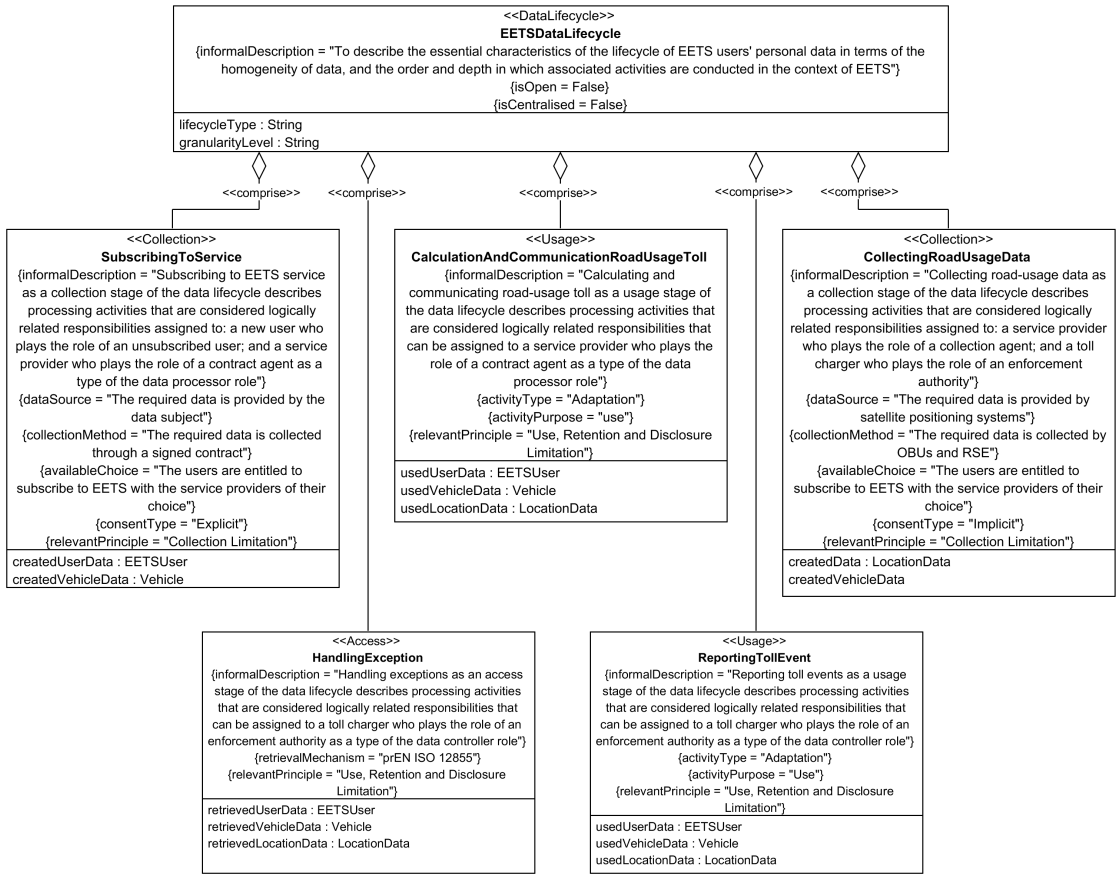

Figure 7.7: EETSDataLifecycle along with its lifecycle stages.

CollectingRoadUsageData is a class stereotyped by «Collection» to represent a set of related activities, i.e. collecting toll declaration data, with the aim of collecting the time of usage, the covered distance and the place on which the vehicle is circulating on a particular toll domain for toll calculation. At the domain level, its tagged values informally describe the aim of the class, data sources, collection methods, available choices, consent type and the relevant GPS principle, as illustrated in Figure 7.7. At the instance level, the value of its createdData attribute is of the type LocationData.

Figure 7.7 partially shows the main lifecycle stages of the EETSDataLifecycle, which reflect the main types of operations performed on EETS users' personal data. Table 7.1 lists the main activities of the EETSDataLifecycle stages.

«LifecycleStage»	«StageActivity»
SubscribingToService (PO.1)	SigningContract (DPA.1)
	CustomisingOBU (DPA.2)
CollectingRoadUsageData (PO.2)	CollectingUsageData (DPA.3)
	CollectingRealTimeData (DPA.4)
StoringContractData (PO.3)	StoringEETSUserData (DPA.5)
	StoringVehicleData (DPA.6)
StoringRoadUsageData (PO.4)	StoringUsageData (DPA.7)
	StoringRealTimeData (DPA.8)
HandlingException (PO.5)	IdentifyingEETSContract (DPA.9)
	RequestingAdditionalData (DPA.10)
CalculatingUsageToll (PO.6)	CalculatingRoadUsageToll (DPA.11)
	CommunicatingRoadUsageToll (DPA.12)
ReportingTollEvent (PO.7)	ReportingComplianceEvent (DPA.13)

Table 7.1: The stages of the EETSDataLifecycle along with associated data-processing activities.

StageActivity, StageEvent and StageAction. Each lifecycle stage involves a set of stage activities, each of which contains a set of actions that represent its executable steps, and a set of events that cause the execution of these actions. SubscribingToService, for example, as a collection stage involves two stage activities: SigningContract and CustomisingOBU.

SigningContract is a class stereotyped by «StageActivity» to represent the activity of collecting EETS user's personal data by signing a contract with a service provider. At the domain level, its tagged values informally describe the aim of the class, the pre- and the post-conditions of the activity, as illustrated in Figure 7.8. At the instance level, the values of its input and output attributes are of the type EETSUser. SigningContract coordinates its execution by containing CollectUserData as an action and defining Sign as an event that causes the execution of this action.

CollectUserData is a class stereotyped by «StageAction» to represent a concrete action of collecting the user's personal data with the aim of subscribing a contract with EETS service and creating a personal account number. At the domain level, its tagged values informally describe the aim of the class, the local pre and post-conditions of the action, as illustrated in Figure 7.8. At the instance level, the values

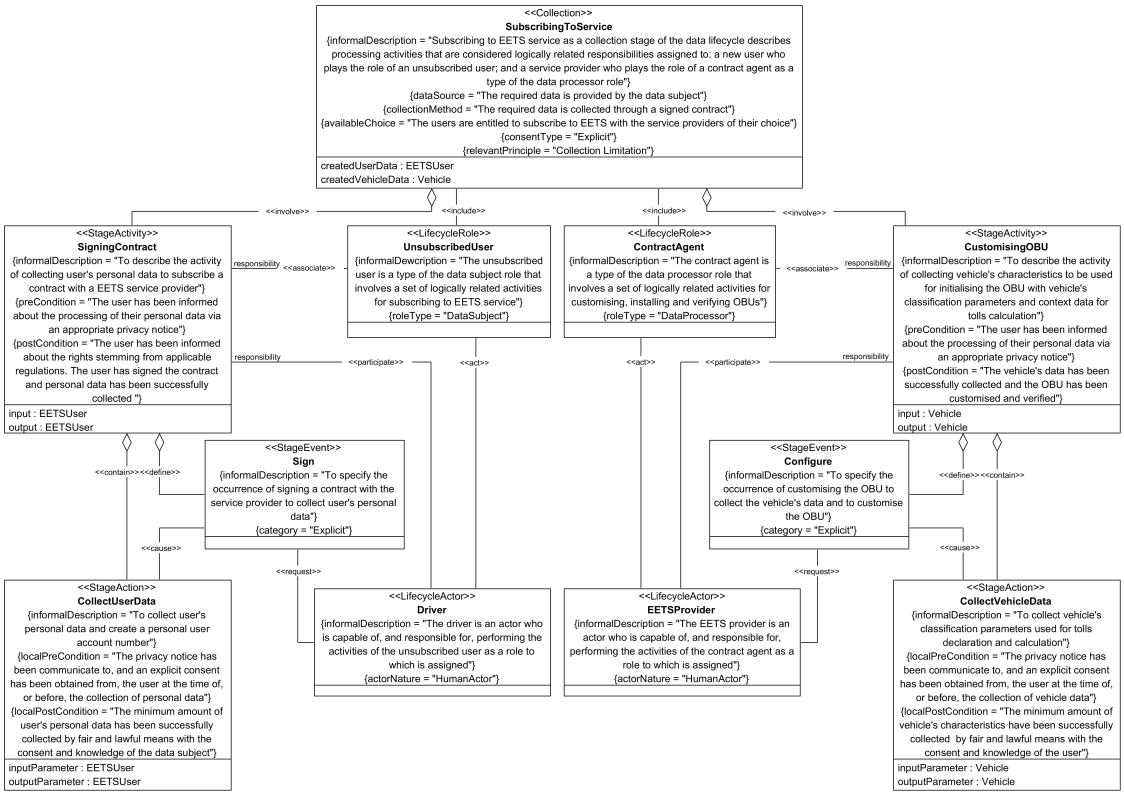

Figure 7.8: The representation of the SubscribingToService as a collection stage.

of its inputParameter and outputParameter attributes are of the type EETSUser. Sign is a class stereotyped by «StageEvent» to represent the occurrence of signing a contract with a service provider, which causes the execution of the CollectUserData as the corresponding action. At the domain level, its tagged values informally describe the aim of the class and the category of the event, as illustrated in Figure 7.8.

CustomisingOBU is a class stereotyped by «StageActivity» to represent the activity of collecting vehicle's classification parameters by customising, initialising and verifying the OBU. At the domain level, its tagged values informally describe the aim of the class, the pre- and post-conditions of the activity, as illustrated in Figure 7.8. At the instance level, the values of its input and output attributes are of the type Vehicle. CustomisingOBU coordinates its execution by containing CollectVehicleData as an action and defining Configure as an event that causes its execution.

CollectVehicleData is a class stereotyped by «StageAction» to represent a concrete action of collecting the vehicle's classification parameters to be used for declaration and calculation of tolls. At the domain level, its tagged values

«StageActivity»	«StageAction»	«StageEvent»
SigningContract	CollectUserData	Sign
CustomisingOBU	CollectVehicleData	Configure
CollectingUsageData	CollectLocationData	Detect
CollectingRealTimeData	CollectRealTimeLocationData	Collect
	CollectVehicleData	Identify
CalculatingRoadUsageToll	CalculateToll	Calculate
CommunicatingRoadUsageToll	CommunicateToll	Communicate
IdentifyingEETSContract	RetrieveEETSContract	Identify
RequestingAdditionalData	RequestPaymentGuarantee	RequestGuarantee
	RequestParametersForBilling	RequestData
ReportingComplianceEvent	ReportCCC	Report

Table 7.2: The representation of EETS stage activities, actions and events.

informally describe the aim of the class, the local pre- and post-conditions of the action, as illustrated in Figure 7.8. At the instance level, the values of its inputParameter and outputParameter attributes are of the type Vehicle.

Configure is a class stereotyped by «StageEvent» to represent the occurrence of customising and acquiring the OBU, which causes the execution of CollectVehicleData as the corresponding action. At the domain level, its tagged values informally describe the aim of the class and the category of the event, as illustrated in Figure 7.8.

The representation of other lifecycle stages, along with their activities, actions and events are illustrated in Appendix B.1.2.

LifecycleRole and LifecycleActor. Each lifecycle stage includes a set of lifecycle roles, each of which is played by different actors according to their capabilities and responsibilities. SubscribingToService, for example, includes UnsubscribedUser and ContractAgent as lifecycle roles, and Driver and EETSProvider as lifecycle actors. Figure 7.8 shows the stage activities, actions and events, as well as lifecycle roles and actors of the UnsubscribedUser as an instance of the Collection stage.

UnsubscribedUser is a class stereotyped by «LifecycleRole» to represent a type of the data subject role as a set of related activities, including SigningContract, that are expected to be performed together for a certain task, i.e. subscribing

«LifecycleRole»	«LifecycleActor»	«StageActivity»
UnsubscribedUser	Driver	SigningContract
CollectionAgent	EETSProvider	CollectingUsageData
ContractAgent	EETSProvider	CustomisingOBU
		CalculatingRoadUsageToll
		CommunicatingRoadUsageToll
EnforcementAuthority	TollCharger	CollectingRealTimeData
		IdentifyingEETSContract
		RequestingAdditionalData
		ReportingComplianceEvent

Table 7.3: The representation of EETS lifecycle roles and actors.

to EETS service. At the domain level, its tagged values informally describe the aim of the class, and the main role type.

ContractAgent is a class stereotyped by «LifecycleRole» to represent a type of the data processor role as a set of related activities, including CustomisingOBU, that are expected to be performed together for a certain task, i.e. toll declaration and calculation. At the domain level, its tagged values informally describe the aim of the class and the main role type.

Driver is a class stereotyped by «LifecycleActor» to represent the driver, owner or fleet operator of the vehicle that is capable of, and responsible for, performing the activity of SigningContract, i.e. associated to UnsubscribedUser, with the aim of subscribing to the EETS service. The performance of this activity can be achieved by requesting the Sign event that causes the execution of the CollectUserData action. At the domain level, its tagged values informally describe the aim of the class and the nature of the actor.

EETSProvider is a class stereotyped by «LifecycleActor» to represent the service provider that is capable of, and responsible for, performing the activity of CustomisingOBU, associated to ContractAgent, with the aim of customising, installing and verifying OBUs. The performance of this activity can be achieved by requesting the Configure event that causes the execution of the CollectVehicleData action. At the domain level, its tagged values informally describe the aim of the class and the nature of the actor.

Table 7.3 shows how lifecycle roles, actors and stages are mapped to each other. The representation of other lifecycle stages, along with their lifecycle roles and actors are illustrated in Appendix B.1.2.

In summary, the UML profile of the APDL model has served as a preliminary acquisition step to capture all required concepts that support requirements analysis — a critical step in the system development lifecycle. It represents data-processing activities in a contextual and fined-grained manner that is amenable to risk analysis and compliance checking.

7.4.2 Data-Centric Threat Modelling

To establish the context in which EETS users' personal data is collected and processed, *primary assets* need to be represented in a way that is amenable to risk analysis. All the relevant information that helps establish the context in which this data is collected and processed has already been captured by the APDL model and represented using the UML profile for the APDL model (see Section 7.4.1). As such, the UML profile for the APDL model serves as a means for establishing the context by modelling the primary assets: personal data, data-processing activities along with abstract principles stated in legal frameworks and standards, and involved actors.

Vulnerability Analysis

To identify and analyse all possible vulnerabilities of primary assets, we describe how to develop a baseline model that helps identify privacy vulnerabilities in two main steps.

The first step is to establish a context-relative processing norm for each data-processing activity. Then, these norms need to be classified according to the stages of the APDL model. The main elements that constitute these norms are captured from data-processing activities specified in Section 7.4.1. In this section, we illustrate how to establish the processing norms of the processing operation: *SubscribingToService (PO.1)*, which consists of two data-processing activities: *SigningContract (DPA.1)* and *CustomisingOBU (DPA.2)*.

- *The context-relative processing norm (PN.1)* for DPA.1 is as follows.

  > In the context of EETS, the collection of personal data of a certain type (EETSUser: user ID, name, billing address) about EETS users (acting as data subjects) by EETS providers (acting as data processors on behalf of toll chargers) is governed by processing principles derived from the EU's GDPR, Directive 2004/52/EC and the related Commission Decision 2009/750/EC.

- *The context-relative processing norm (PN.2)* for DPA.2 is as follows.

  > In the context of EETS, the collection of personal data of a certain type (Vehicle: licence plate, classification code) about EETS users (acting as data subjects) by EETS providers (acting as data processors on behalf of toll chargers) is governed by processing principles derived from the EU's GDPR, Directive 2004/52/EC and the related Commission Decision 2009/750/EC.

PN.1 and PN.2 are categorised as processing norms in the Collection stage of the APDL model because PO.1 is a class stereotyped by «Collection». The processing norms of PO.2 – PO.7 are established in Appendix B.2.1.

By establishing all context-relative processing norms in relation to the stages of the APDL model, a complete baseline model is developed to serve as the basis for deriving privacy vulnerabilities.

The second step is to derive all possible privacy vulnerabilities from the established context-relative processing norms. This can be achieved by analysing all the elements that constitute each processing norm — i.e. any possible breach of a processing norm can be derived as a vulnerability. In this section, we illustrate how to analyse the elements of PN.1 and PN.2 and derive relevant privacy vulnerabilities.

- ***Element 1:*** *data-processing activities.* PN.1 and PN.2 are established with reference to DPA.1 and DPA.2, which are specified as a result of the purpose refinement of Section 7.4.1. Thus, DPA.1 and DPA.2 explicitly participate in the fulfilment of the abstract purpose.

- ***Element 2:*** *attributes.* 'Identification and contact data' and 'vehicle classification parameters' are necessary for the fulfilment of the concrete purposes

of DPA.1 and DPA.2 respectively. This means that they are necessary to accomplish the execution of associated actions: CollectUserData and CollectVehicleData respectively. In relation to the specification of DPA.1, it is not explicitly specified in a manner that prevents the collection of irrelevant data items when signing a contract. As such, 'an improper activity specification' (PV.1) is derived as a privacy vulnerability (a weakness in the specification of DPA.1).

- ***Element 3:*** *actors.* Drivers and EETS providers are the actors who are involved in the performance of the DPA.1 and DPA.2 respectively. The drivers are responsible for performing the activities of the unsubscribed users as a role to which they are assigned with data-providing capabilities. EETS providers are responsible for performing the activities of the contract agent as a role to which they are assigned. Both types of roles are defined in PO.1. We assume that a role-based access control model is maintained in a proper way.

- ***Element 4:*** *processing principles.* The relevant processing principles stated in the EU's GDPR are: "[...] personal data must be":

  – "processed fairly and lawfully";

  – "collected for specified, explicit and legitimate purposes and not further processed in a way incompatible with those purposes"; and

  – "adequate, relevant and not excessive in relation to the purposes for which they are collected and/or further processed".

  The relevant processing principles stated in the GPS are: Collection Limitation and Data Minimisation. In relation to the specification of DPA.1 and DPA.2, these processing principles are specified as pre- and post-conditions, as per Figure 7.8. This means that the collection of these types of personal data is in compliance with the processing principles stated in relevant legal frameworks and standards.

Code	Privacy vulnerability
PV.1	An improper activity specification
PV.2	An improper data model
PV.3	A lack of logs and audit trails
PV.4	An improper purpose specification
PV.5	A weak anonymisation technique

Table 7.4: A list of the most important vulnerabilities in the context of EETS.

Other privacy vulnerabilities that are derived from PO.2 – PO.7 are described in Appendix B.2.1. Table 7.4 shows the most important vulnerabilities in the context of EETS.

Threat Analysis

Threat Sources. The first step is to identify the actors who are involved in the collection, processing and dissemination of EETS users' personal data. By considering DPA.1 – DPA.13, EETS providers (TS.1) and toll chargers (TS.2) are the actors who are directly involved in the collection and processing of EETS users' personal data. Table 7.3 shows those actors and their roles and responsibilities (data-processing activities associate with the roles to which they are assigned).

The second step is to identity third parties who may perform processing operations on EETS users' personal data when it is shared by EETS providers and/or toll chargers. According to the nature and utility of 'location data' specified in the established context, it is obvious that there are external entities who are not directly involved in the processing of this data; rather, they may act as third parties with interests or concerns in the value of these types of personal data with various motives and resources, as per Table 7.5.

- Departments, agencies and public bodies. Fine-grained location data may be of interest to state agencies for several motives.

  - Department for Transport (TS.3): Location data can be used as a source for collecting traffic statistics to improve road mobility by applying congestion charges.
  - Intelligence and security services (TS.4): Location data can be used to facilitate law enforcement investigations by discovering whether individuals are where they claim to have been at any point in time. It can also be used to identify and put individuals under surveillance based on their associations with others or the locations frequented.

- Employment agencies (TS.5): Location data can be used for background checking. Those agencies may make excessive inference with the aim of, for example, deriving health conditions from driving patterns. The agencies may use the derived data for filtering job candidates based on these conditions.

- Health insurance providers (TS.6): Location data can be used to make excessive inference with the aim of deriving health conditions. Those conditions are considered to be one of the main factors for calculating health insurance premiums.

- Car insurance providers (TS.7): Location data can be used to discover whether a driver is where they claim to have been at any point in time. In addition, it can be used to make excessive inference with the aim of deriving driving patterns, which are one of the main factors for calculating car insurance premiums.

- Advertising companies (TS.8): Location data can be used to make excessive inference with the aim of, for example, deriving health conditions or religious beliefs. These companies may use the derived data for sending targeted advertising or unsolicited emails.

Code	Threat source
TS.1	EETS providers
TS.2	Toll chargers
TS.3	Department for Transport
TS.4	Intelligence and security services
TS.5	Employment agencies
TS.6	Health insurance providers
TS.7	Car insurance providers
TS.8	Advertising companies

Table 7.5: The most important threat sources in the context of EETS.

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.

The first step is to analyse the identified privacy vulnerabilities and associated threat sources for each context-relative processing norm grouped by the stages of the APDL model.

- *In the Collection stage,* threat events that may lead to privacy violations or harms are related to the manner in which personal data is collected. In reference to the established processing norms, the Collection stage consists of four processing norms: PN.1 – PN.4. Both PN.1 and PN.3 involve PV.1 as a privacy vulnerability, whereas PN.2 and PN.4 do not involve any privacy vulnerabilities.

  By exploiting PV.1, TS.1 may use contracts (as a method of collection) to actively collect irrelevant personal data (TE.1) about EETS users by requiring unwarranted personal data items as a condition for subscribing to EETS. This threat event is a type of 'interrogation' according to the adapted taxonomy of adverse privacy events.

By exploiting PV.1, TS.1 or TS.2 may use OBUs (as a method of collection) to excessively collect location data (TE.2) about EETS users by recording a large amount of data about their movements. This threat event is a type of 'surveillance' according to the adapted taxonomy of adverse privacy events.

- *In the Retention stage*, threat events that may lead to privacy violations or harms are related to the manner in which personal data is modelled. In reference to the established processing norms, the Retention stage consists of four processing norms: PN.5 – PN.8. Both PN.5 and PN.7 involve PV.2 as a privacy vulnerability, whereas PN.6 and PN.8 do not involve any privacy vulnerabilities.

  By exploiting PV.2, TS.1 may store complete driving records for EETS users (TE.3) by integrating multiple items of personal data from various sources. This threat event is a type of 'aggregation' according to the adapted taxonomy of adverse privacy events.

  By exploiting PV.2, TS.1 or TS.2 may store driving profiles (TE.4) for EETS users by linking driving records to particular EETS users. This threat event is a type of 'identification' according to the adapted taxonomy of adverse privacy events.

- *In the Access stage,* threat events that may lead to privacy violations or harms are related to the manner in which personal data is retrieved. In reference to the established processing norms, the Access stage consists of two processing norms: PN.9 and PN.10. Both PN.9 and PN.10 involve PV.3 as a privacy vulnerability.

  By exploiting PV.3, TS.2 may retrieve identifiable driving profiles (TE.5) for EETS users by specifying specific vehicle classification parameters. This threat event is a type of 'insecurity' according to the adapted taxonomy of adverse privacy events.

- *In the Usage stage,* threat events that may lead to privacy violations or harms are related to the manner in which personal data is manipulated and used. In reference to the established processing norms, the Usage stage consists of three processing norms: PN.11 – PN.13. They involve PV.2 and PV.3 as privacy vulnerabilities.

  By exploiting PV.2 and PV.3, TS.1 may infer excessive sensitive information (TE.6) by analysing the aggregated data — EETS users' profiles — in a particular data analysis technique, for example, data mining to discover useful information, such as driving patterns that may reveal health conditions among others. This threat event is a type of 'identification' according to the adapted taxonomy of adverse privacy events.

  By exploiting PV.2 and PV.3, TS.2 may infer excessive sensitive information (TE.6) by analysing real-time location data that may reveal health conditions among others. This threat event is a type of 'identification' according to the adapted taxonomy of adverse privacy events.

  By exploiting PV.3 and PV.4, TS.1 or TS.2 may use EETS users' profiles for further processing (TE.7). This includes commercial or malicious purposes which are not related to the purposes for which location data was initially collected and for which EETS users have provided implicit or explicit consent. This threat event is a type of 'secondary use' according to the adapted taxonomy of adverse privacy events.

- *In the Disclosure stage,* threat events that may lead to privacy violations or harms are related to the manner in which personal data is disseminated, made available or transmitted to third parties for external use.

  By exploiting PV.3 and PV.4, TS.1 or TS.2 may share driving patterns with third parties (TE.8) that are with interests or concerns in the value of these types of personal data. This threat event is a type of 'disclosure' according to the adapted taxonomy of adverse privacy events.

Code	Threat event
TE.1	Irrelevant data collection
TE.2	Excessive data collection
TE.3	Irrelevant data retention
TE.4	Excessive data retention
TE.5	Unauthorised data retrieval
TE.6	Excessive data inference
TE.7	Unauthorised secondary use
TE.8	Unauthorised data disclosure

Table 7.6: The most important threat events in the context of EETS.

The utility of location data can be of interest to TS.3, TS.4, TS.5, TS.6, TS.7 or TS.8 for several purposes. By exploiting PV.5, TS.3 may excessively make inference (TE.9) with the aim of re-identifying data subjects to facilitate enforcement management for congestion charge payments. TS.4 may excessively make inference (TE.9) with the aim of re-identifying data subjects to facilitate law enforcement investigations. In addition, TS.4 may re-identify data subjects to put them under close surveillance based on their associations with others or the locations frequented. TS.5 may excessively make inference (TE.9) with the aim of re-identifying data subjects for background checking to filter job candidates based on the derived health conditions. TS.6 may excessively make inference (TE.9) with the aim of re-identifying data subjects to use the derived health conditions as a criterion for calculating health insurance premium. TS.7 may excessively make inference (TE.9) with the aim of re-identifying data subjects to use the derived vehicle use and health conditions as criteria for calculating car insurance premium. TS.8 may excessively make inference (TE.9) with the aim of re-identifying data subjects to use the derived religious beliefs or health conditions as references for sending targeted advertising or unsolicited emails. This event is a type of 'identification' according to the adapted taxonomy of adverse privacy events.

Table 7.6 summarises the most important threat events.

Code	Threat event	Threat source	Privacy vulnerability
EX.01	TE.1	TS.1	PV.1
EX.02	TE.2	TS.1	PV.1
EX.03	TE.2	TS.2	PV.1
EX.04	TE.3	TS.1	PV.2
EX.05	TE.4	TS.1	PV.2
EX.06	TE.4	TS.2	PV.2
EX.07	TE.5	TS.2	PV.3
EX.08	TE.6	TS.1	PV.2 and PV.3
EX.09	TE.6	TS.2	PV.2 and PV.3
EX.10	TE.7	TS.1	PV.3 and PV.4
EX.11	TE.7	TS.2	PV.3 and PV.4
EX.12	TE.8	TS.1	PV.3 and PV.4
EX.13	TE.8	TS.2	PV.3 and PV.4
EX.14	TE.9	TS.3	PV.5
EX.15	TE.9	TS.4	PV.5
EX.16	TE.9	TS.5	PV.5
EX.17	TE.9	TS.6	PV.5
EX.18	TE.9	TS.7	PV.5
EX.19	TE.9	TS.8	PV.5

Table 7.7: A set of reasonable exploitations.

The second step is to identify all reasonable exploitations that lead to each threat event. Table 7.7 shows how the identified threat events result from the successful exploitation of the identified privacy vulnerabilities by the most likely threat sources.

Privacy Harm Analysis

Privacy Violations. In this section, we identify potential privacy violations in the Collection stage of the APDL model. Privacy violations that are related to the Retention, Access and Usage stages of the APDL model are described in Appendix B.2.2.

In the Collection stage, TE.1 and TE.2 are identified as threat events.

By the occurrence of TE.1, TS.1 may passively (without the knowledge and consent of EETS users) collect irrelevant data, for example, technical data (IP address, MAC address, etc.) using cookies as a method of collection.

Code	Privacy Violation
VI.1	Passive personal data collection
VI.2	Passive location data collection
VI.3	Passive data acquisition
VI.4	Unjustified data retention
VI.5	Unwarranted data retrieval
VI.6	Unjustified data inference
VI.7	Unjustified data manipulation

Table 7.8: The most important privacy violations in the context of EETS.

By the occurrence of TE.2, TS.1 may passively collect location data outside toll domains. By conducting such an action, fine-grained location data is collected in ways EETS users would not reasonably expect; in addition, 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 effects to EETS users.

Table 7.8 shows the most important violations in the context of EETS.

Privacy Harms. Privacy harm analysis is the most important step of any analysis approach. Harms are derived from the adverse effects of threat events as potential adverse actions taken against data subjects.

The first step to identify privacy harms is to identify adverse effects of threat events and categorise these effects according the stages of the APDL model.

- *In the Collection stage,* TE.1 and TE.2 are identified as threat events. The main adverse effect of TE.1 is gathering a large amount of unwarranted personal data (AE.1).

  The main adverse effect of TE.2 is gathering a large amount of fine-grained location data that has been collected over time as comprehensive driving records (AE.2), which may include complete driving history or driving history for a specific period for EETS users.

- *In the Retention stage,* TE.3 and TE.4 are identified as threat events. The main adverse effect of TE.3 is storing multiple data items about EETS users (AE.3), from which comprehensive driving records can be derived.

  The main adverse effect of TE.4 is storing identifiable driving records for EETS users (AE.4), from which driving profiles can be derived.

- *In the Access stage,* TE.5 is identified as a threat event. The main adverse effect of TE.5 is retrieving identifiable driving records for EETS users (AE.5), from which driving movements can be derived.

- *In the Usage stage,* TE.6 and TE.7 are identified as threat events. The main adverse effect of TE.6 is creating identifiable driving profiles for EETS users (AE.6), from which driving patterns can be derived.

  The main adverse effect of TE.7 is deriving driving patterns from their driving profile (AE.7), from which sensitive information can derived.

- *In the Disclosure stage,* TE.8 and TE.9 are identified as threat events. The main adverse effect of TE.8 is revealing anonymised driving patterns for subscribed EETS users beyond expected boundaries (AE.8), from which sensitive information about EETS users can be derived.

  The main adverse effect of TE.9 is deriving sensitive information about EETS users from their driving patterns (AE.9), based on which adverse actions against the data subjects can be taken by the relevant threat sources.

Table 7.9 shows the most important adverse effects of TE.1 – TE.9 in the context of EETS.

The second step is to analyse the identified adverse effects as motives for adverse actions. By analysing AE.1 – AE.9, together with the relevant threat scenarios, we can derive a set of privacy harms. In this section, we describe *increased car insurance premium (PH.1)* in terms of its threat scenarios. Other privacy harms that can be derived from these adverse effects are described in Appendix B.2.

Code	Adverse effect	Threat
AE.1	A large amount of unwarranted personal data	TE.1
AE.2	A large amount of fine-grained location data	TE.2
AE.3	Comprehensive driving records	TE.3
AE.4	Identifiable driving records	TE.4
AE.5	Identifiable driving records	TE.5
AE.6	Identifiable driving profiles	TE.6
AE.7	Identifiable driving patterns	TE.7
AE.8	Anonymised driving patterns	TE.8
AE.9	Sensitive facts about EETS users	TE.9

Table 7.9: The most important adverse effects of threat events in the context of EETS.

PH.1 occurs when a EETS provider (TS.1) or a toll charger (TS.2) may make 'excessive data inference to derive driving patterns' (TE.6) for EETS users and 'shares these driving patterns with car insurance companies' (TE.8). An insurance provider (TS.7) may make 'excessive data inference to re-identify its current and potential customers' (TE.9) by linking the derived data to particular drivers to find out their health conditions and vehicle use, or to discover whether a policy holder is where they claim to have been at any point in time. An 'improper data model' (PV.2) and a 'lack of logs and audit trails' (PV.3) that may be exploited by TS.1 or TS.2 lead to the occurrence of TE.6. A 'lack of logs and audit trails' (PV.3) and 'improper purpose specification' (PV.4) that may be exploited by TS.1 or TS.2 lead to the occurrence of TE.8. A 'weak anonymisation technique' (PV.5) that is exploited by TS.7 leads to the occurrence of TE.9.

Table 7.10 shows the most significant privacy harms, along with associated threat sources, privacy vulnerabilities, threat events and adverse effects of these events.

To model all the reasonable threat scenarios for each privacy harms, we construct harm trees for each privacy harm. Figure 7.9 shows the harm tree for the privacy harm PH.1. Other privacy harms (PH.2 – PH.5) are described and assessed in Appendix B.2.3.

Code	Threat source	Vulnerability	Threat event	Adverse effects
PH.1	TS.1 or TS.2	PV.2 and PV.3	TE.6	AE.6
	TS.1 or TS.2	PV.3 and PV.4	TE.8	AE.8
	TS.7	PV.5	TE.9	AE.9
PH.2	TS.1 or TS.2	PV.2 and PV.3	TE.6	AE.6
	TS.1 or TS.2	PV.3 and PV.4	TE.8	AE.8
	TS.6	PV.5	TE.9	AE.9
PH.3	TS.1 or TS.2	PV.2 and PV.3	TE.6	AE.6
	TS.1 or TS.2	PV.3 and PV.4	TE.8	AE.8
	TS.5	PV.5	TE.9	AE.9
PH.4	TS.1 or TS.2	PV.2 and PV.3	TE.6	AE.6
	TS.1 or TS.2	PV.3 and PV.4	TE.8	AE.8
	TS.4	PV.5	TE.9	AE.9
PH.5	TS.1 or TS.2	PV.2 and PV.3	TE.6	AE.6
	TS.1 or TS.2	PV.3 and PV.4	TE.8	AE.8
	TS.8	PV.5	TE.9	AE.9

Table 7.10: The most significant privacy harms along with associated threat scenarios.

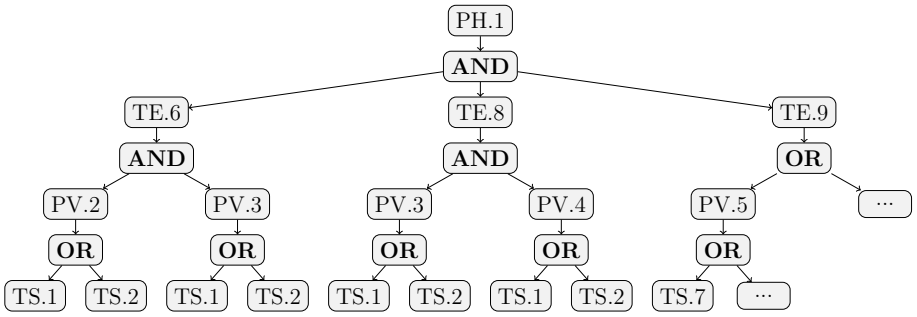

PH.1:	Increased car insurance premium	TE.6:	Excessive data inference
TE.8:	Unauthorised data disclosure	TE.9:	Excessive data inference
PV.2:	Improper data model	PV.3:	Lack of logs and audit trails
PV.4:	Improper purpose specification	PV.5:	Weak anonymisation technique
TS.1:	EETS providers	TS.2:	Toll chargers
TS.7:	Car insurance providers		

Figure 7.9: The structure of the harm tree for the privacy harm PH.1.

Severity Assessment

The intensity of PH.1 is based on the extent of damage caused by, the irreversibility of, and the duration of, the adverse effects of associated threat events TE.6, TE.8 and TE.9. TE.6 is categorised as a type of 'identification'. It is characterised as unanticipated by EETS users; it is also characterised as extensive in terms of the

Privacy harm	Intensity	Range	Severity
PH.1	2. Limited	3. Significant	2. Limited
PH.2	2. Limited	3. Significant	2. Limited
PH.3	2. Limited	2. Limited	2. Limited
PH.4	4. Maximum	1. Negligible	2. Limited
PH.5	1. Negligible	3. Significant	1. Negligible

Table 7.11: The severity of the identified privacy harms.

volume of data. TE.8 is categorised as a type of 'disclosure'. It is characterised as extensive and accurate. TE.9 is categorised as a type of 'identification'. It is characterised as unanticipated by EETS users; it is also characterised as extensive. The adverse effects (AE.6, AE.8 and AE.9) of these threat events are identified as a series of related impacts started by discovering private facts about EETS users, then revealing sensitive information beyond expected boundaries and ended by charging higher rates of insurance premium. They are assessed as slight because calculating insurance premium depends on other factors, including address, occupation, claims history, etc. The duration of these consequences may last for a certain length of time: the period of cover. However, they may last for longer than the period of cover when the disclosed data is used as 'driving history' by insurance providers to calculate car insurance quotes. Once profiles are created, disclosed to insurance providers and sensitive information is inferred, it is technically difficult to reverse these effects. As such, the intensity of PH.1 is assessed as '2. Limited'.

PH.1 may affect specific EETS users based on their driving patterns. They may also affect specific categories of EETS users based on their health conditions. Together, these may impact insurance premiums. Thus, the range of this privacy harm is assessed as '3. Significant'. The severity of PH.1 is assessed as '2. Limited' by adding the assessed value of the intensity and the assessed value of the range, and selecting the overall value according to Table 5.1. The severity of PH.2 – PH.5 are similarly assessed. Table 7.11 shows the severity of the identified privacy harms (PH.1 – PH.5).

Threat source	Motivation	Ability	Capability
TS.1	3. Significant	3. Significant	3. Significant
TS.2	3. Significant	3. Significant	3. Significant
TS.7	4. Maximum	3. Significant	4. Maximum
TS.6	4. Maximum	3. Significant	4. Maximum
TS.5	3. Significant	3. Significant	3. Significant
TS.4	3. Significant	4. Maximum	4. Maximum
TS.8	4. Maximum	3. Significant	4. Maximum

Table 7.12: The capability of the identified threat sources.

Likelihood Assessment

In reference to the harm tree of PH.1 (see Figure 7.9), the likelihood of occurrence of TE.6, TE.8 and TE.9 is assessed based on the capability of TS.1, TS.2 and TS.7 and the exploitability of PV.2 – PV.5 for each possible exploit.

The motivations of TS.1, TS.2 and TS.7 are based on the values of location-related data to those sources and their motives according to these values. The utility of 'identification and contact data' and 'location data' makes such data highly valuable to EETS providers, toll chargers and insurance providers (they have definite motives based on the value of this data). It may also have a nuisance value when it is exploited by insurance providers. Thus, the motivation of TS.1 and TS.2 are assessed as '3. Significant' and the motivation of TS.7 is assessed as '4. Maximum'.

The ability of those sources is based on their skills, background knowledge, privileges, and technical and financial resources. According to the 'type' attribute of TS.1 and TS.2, they are insiders and institutions. This implies that they 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. EETS providers and toll charges play the roles of data processors and data controllers respectively. Based on these, they have access rights to both the 'fine-grained location data' and 'identification and contact data'. In addition, they have both technical and financial resources to benefit from the values of the collected data by creating comprehensive and

Privacy vulnerability	Exploitability	Severity	Seriousness
PV.2	4. Maximum	3. Significant	4. Maximum
PV.3	2. Limited	3. Significant	2. Limited
PV.4	3. Significant	3. Significant	3. Significant
PV.5	2. Limited	4. Maximum	3. Significant

Table 7.13: The seriousness of PV.2, PV.3, PV4 and PV.5.

identifiable profiles. As such, the abilities of TS.1 and TS.2 are assessed as '3. Significant'. According to the 'type' attribute of TS.7, they are outsiders and institutions. Insurance providers are third parties that do not have direct roles with respect to the processing of personal data. Based on this, they do not have access rights to the 'fine-grained location data' and 'identification and contact data'; rather, they can legally process this data when it is anonymised. In addition, they have both technical and financial resources to benefit from the values of the disclosed data by making excessive inference. As such, the ability is assessed as '3. Significant'. As such, the capability of TS.1 and TS.2 are assessed as '3. Significant' and the capability of TS.7 is assessed as '4. Maximum'. The capabilities of TS.6, TS.5, TS.4 and TS.8 are similarly assessed, as per Table 7.12.

The seriousness of each of PV.2 – PV.5 is based on the exploitability and severity of the vulnerability. The ease of the exploitation of PV.2 is influenced by the relevant element of context-relevant processing norms. The relevant element of the processing norms is 'attributes', which refers to personal data. In this context, personal data is classified into two types: 'identification and contact data' and 'location data'. Both types are categorised as 'collected data'. These types of data are not sensitive in themselves; rather, they are valuable and can be used to derive sensitive information, such as driving history or patterns, and health conditions. The fine-grained location data can easily be linked with 'identification and contact data' with reasonable effort as they are modelled, collected and processed by an EETS provider and accessed by a toll charger. The vulnerability of 'improper data model' can be easily exploited by EETS providers and toll chargers. Thus, the exploitability of PV.2 is assessed as '4. Maximum'.

Code	Exploit	T.E.	L.O.	L.R.	O.L.
EX.01	TS.1 - PV.2 and PV.3	TE.6	S	S	S
EX.02	TS.2 - PV.2 and PV.3		S	S	S
EX.03	TS.1 - PV.2 and TS.2 - PV.3		S	S	S
EX.04	TS.2 - PV.2 and TS.1 - PV.3		S	S	S
EX.05	TS.1 - PV.3 and PV.4	TE.8	L	S	L
EX.06	TS.2 - PV.3 and PV.4		L	S	L
EX.07	TS.1 - PV.3 and TS.2 - PV.4		L	S	L
EX.08	TS.2 - PV.3 and TS.1 - PV.4		L	S	L
EX.09	TS.7 - PV.5	TE.9	S	M	M
EX.10	TS.6 - PV.5	TE.9	S	M	M
EX.11	TS.5 - PV.5	TE.9	L	M	S
EX.12	TS.4 - PV.5	TE.9	S	M	M
EX.13	TS.8 - PV.5	TE.9	S	M	M

T.E.: Threat Event. **L.O.:** Likelihood of occurrence. **L.R.:** Likelihood of resulting in adverse impacts. **O.L.:** The overall likelihood.
M: 4. Maximum, **S:** 3. Significant, **L:** 2. Limited, **N:** 1. Negligible.

Table 7.14: The overall likelihood of occurrence of the identified threats.

The severity of PV.2 is influenced by the relevant element of context-relevant processing norms. The relevant element of the processing norms is 'attributes', which refers to personal data. Thus, PV.2 enables TS.1 or TS.2 to breach the processing norms and violate the contextual integrity by making unjustified data inference (TE.6) with the aim of deriving driving patterns for EETS users based on their deriving history. This type of identification is a threat event that may lead to the privacy harm PH.1. As such, the severity of PV.2 is assessed as '3. Significant'. The seriousness of PV.3, PV.4 and PV.5 are similarly assessed based on the corresponding exploitability and severity, as per Table 7.13.

The threat event TE.6 may result from the exploitation of PV.2 and PV.3: both vulnerabilities are necessary for the occurrence of TE.6. The vulnerability PV.2 may be exploited by two different threat sources: TS.1 or TS.2. With reference to the harm tree, either one of those sources is sufficient to exploit the vulnerability PV.2. This leads to four possible exploits: EX.01 – EX.04. Similarly, TE.8 has four possible exploits (EX.05 – EX.08), whereas TE.9 has only one exploit: EX.09. Table 7.14 shows possible exploits of TE.6, TE.8 and TE.9 along with their likelihoods. In

Privacy harm	Severity	Likelihood
PH.1	2. Limited	4. Maximum
PH.2	2. Limited	4. Maximum
PH.3	2. Limited	3. Significant
PH.4	2. Limited	4. Maximum
PH.5	1. Negligible	4. Maximum

Table 7.15: The likelihood of occurrence of the identified privacy harms.

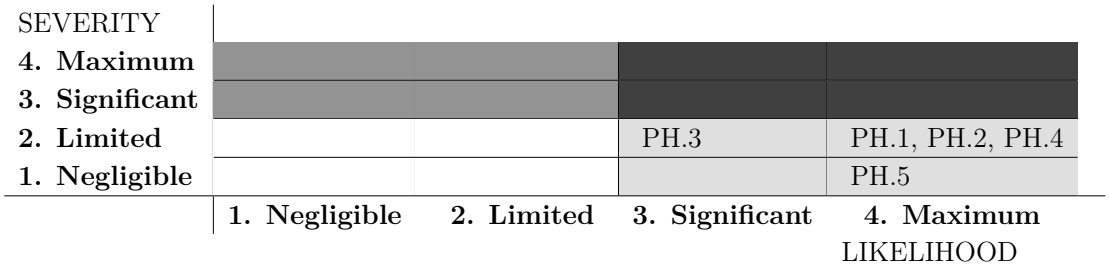

Figure 7.10: The risk levels of the assessed privacy harms.

this case, the highest value of the overall likelihoods of TE.6 based on its exploits is taken: '4. Maximum'. Similarly, the highest value of the overall likelihoods of TE.8 based on its exploits is taken: '2. Limited'. The overall likelihood of the threat event TE.9 is assessed as '4. Maximum'. As such, the likelihood of occurrence of PH.1 is assessed as '4. Maximum' by taking the highest value of the overall likelihood of occurrence of TE.6, TE.8 and TE.9.

EX.01 – EX.08 are common between PH.1 – PH.5, whereas EX.09 – EX.13 relate to PH.1 – PH.5 respectively. The likelihood of occurrence of PH.2 – PH.5 are similarly assessed, as per Table 7.15.

Risk Level Assessment

Figure 7.10 shows the risk levels of the assessed privacy harms PH.1 – PH.5 on the risk map.

7.4.3 Privacy-Enhancing Strategies

Privacy Protection Goals

Figure 7.9 shows the constructed harm tree of PH.1, from which nine threat scenarios for PH.1 can be generated, as per Table 7.16. For each scenario, a quality scenario

Threat scenario	Threat event	Threat source	Privacy vulnerability
TScenario 1	TE.6	TS.1	PV.2 and PV.3
TScenario 2	TE.6	TS.2	PV.2 and PV.3
TScenario 3	TE.6	TS.1/TS.2	PV.2/PV.3
TScenario 4	TE.6	TS.2/TS.1	PV.2/PV.3
TScenario 5	TE.8	TS.1	PV.3 and PV.4
TScenario 6	TE.8	TS.2	PV.3 and PV.4
TScenario 7	TE.8	TS.1/TS.2	PV.3/PV.4
TScenario 8	TE.8	TS.2/TS.1	PV.3/PV.4
TScenario 9	TE.9	TS.7	PV.5

Table 7.16: A set of reasonable threat scenarios pertaining to PH.1.

needs to be generated. Table 7.17, Table 7.18 and Table 7.19 show the quality scenarios for TE.6, TE.8 and TE.9 respectively.

By analysing the response element of each quality scenario, we can derive three context-relative protection goals.

1. By analysing QScenario 1 – QScenario 4, we derive PPGoal 1: The system shall prevent the data inference that can be made by linking location data to identfication and contact data.

2. By analysing QScenario 5 – QScenario 8, we derive PPGoal 2: The system shall prevent the data disclosure that can be made by sharing concealed personal data about EETS users to third parties.

3. By analysing QScenario 9, we derive PPGoal 3: The system shall prevent the identification of data subjects that can be made by linking driving patterns to particular data subjects.

Privacy Architectural Tactics

PATactic 1. To identify an appropriate architectural technique that controls the response element of QScenario 1 – QScenario 4, PPGoal 1 needs to be analysed and mapped onto the most relevant privacy protection goal. PPGoal 1 is an instance of unlinkability as an abstract protection goal. Importantly, data inference is more concerned with the data itself. According to the structure of privacy protection

QScenario 1	**Elements**
Source of stimulus	TS.1
Stimulus	TE.6
Context	EETS
Artefact	PV.2 and PV.3
Response	Prevent the linkability of location data with other types of EETS users' personal data
Response measure	Acceptable level of risk without adverse effects on data utility
QScenario 2	**Elements**
Source of stimulus	TS.2
Stimulus	TE.6
Context	EETS
Artefact	PV.2 and PV.3
Response	Prevent the linkability of location data with other types of EETS users' personal data
Response measure	Acceptable level of risk without adverse effects on data utility
QScenario 3	**Elements**
Source of stimulus	TS.1/TS.2
Stimulus	TE.6
Context	EETS
Artefact	PV.2/PV.3
Response	Prevent the linkability of location data with other types of EETS users' personal data
Response measure	Acceptable level of risk without adverse effects on data utility
QScenario 4	**Elements**
Source of stimulus	TS.2/TS.1
Stimulus	TE.6
Context	EETS
Artefact	PV.2/PV.3
Response	Prevent the linkability of location data with other types of EETS users' personal data
Response measure	Acceptable level of risk without adverse effects on data utility

Table 7.17: The generated quality scenarios for TE.6.

tactics (shown in Figure 6.3), minimise, hide, separate and abstract are the most relevant abstract tactics to achieve PPGoal 1. By referring to the risk level of this

QScenario 5	**Elements**
Source of stimulus	TS.1
Stimulus	TE.8
Context	EETS
Artefact	PV.3 and PV.4
Response	Prevent the exposure of EETS users' driving patterns to third parties without removing personal identifiers
Response measure	Acceptable level of risk without adverse effects on data utility
QScenario 6	**Elements**
Source of stimulus	TS.2
Stimulus	TE.8
Context	EETS
Artefact	PV.3 and PV.4
Response	Prevent the exposure of EETS users' driving patterns to third parties without removing personal identifiers
Response measure	Acceptable level of risk without adverse effects on data utility
QScenario 7	**Elements**
Source of stimulus	TS.1/TS.2
Stimulus	TE.8
Context	EETS
Artefact	PV.3/PV.4
Response	Prevent the exposure of EETS users' driving patterns to third parties without removing personal identifiers
Response measure	Acceptable level of risk without adverse effects on data utility
QScenario 8	**Elements**
Source of stimulus	TS.2/TS.1
Stimulus	TE.8
Context	EETS
Artefact	PV.3/PV.4
Response	Prevent the exposure of EETS users' driving patterns to third parties without removing personal identifiers
Response measure	Acceptable level of risk without adverse effects on data utility

Table 7.18: The generated quality scenarios for TE.8.

privacy harm, the likelihood is *maximum*, whereas its severity is *limited*. In this case, the treatment strategy is to reduce the likelihood of the risk. This means

QScenario 9	Elements
Source of stimulus	TS.7
Stimulus	TE.9
Context	EETS
Artefact	PV.5
Response	Prevent the re-identification by linking driving patterns to other types of personal data
Response measure	Acceptable level of risk without adverse effects on data utility

Table 7.19: The generated quality scenario for TE.9.

that separate and hide can be applied to reduce the likelihood of the risk of the privacy harm. Before analysing their associated tactics, the main type of operation performing in location data is *data analysis*, which is a type related to the Usage stage of the APDL model. Hide is primarily concerned with access control and sharing; the underlying goal of hide is to prevent exposure of access, association, visibility and understandability of personal data, to reduce the likelihood of a privacy threat occurrence. Separate aims to prevent the correlation of personal data to reduce the likelihood of a privacy threat occurrence by distributing or isolating processing operations on personal data. As such, separate's tactics are more relevant to achieve PPGoal 1. By analysing the definitions of concrete tactics of separate, isolate is an appropriate tactic as it aims to process parts of personal data independently, without access or correlation to related parts. Accordingly, isolate is a concrete tactic that limits data inference that can be made by linking location data to other types of information.

PATactic 2. To identify an appropriate architectural technique that controls the response element of QScenario 5 – QScenario 8, PPGoal 2 needs to be analysed and mapped onto the most relevant privacy protection goal. PPGoal 2 is an instance of unlinkability as an abstract protection goal. According to the structure of privacy protection tactics (shown in Figure 6.3), minimise, hide, separate and abstract are the most relevant abstract tactics to achieve PPGoal 2. Similar to the reasoning of PPGoal 1, the treatment strategy is to reduce the likelihood of the risk. This

means that hide and separate can be applied to reduce the likelihood of the risk of the privacy harm. Before analysing their associated tactics, the main type of operation performing in location data is *data transmission*, which is a type related to the Disclosure stage of the APDL model. As previously mentioned above, the underlying goal of hide is to prevent exposure of access, association, visibility and understandability of personal data, to reduce the likelihood of a privacy threat occurrence. As such, hide's tactics are more relevant to achieve Goal 2. By analysing the definitions of concrete tactics of hide, dissociate is an appropriate tactic as it aims to remove the correlation between different pieces of personal data. Accordingly, dissociate is a concrete tactic that prevents data disclosure that can be made by sharing concealed driving patterns for EETS users to third parties.

PATactic 3. To identify an appropriate architectural technique that controls the response element of QScenario 9, PPGoal 3 needs to be analysed and mapped onto the most relevant abstract protection goal. PPGoal 3 is an instance of unlinkability as an abstract protection goal. Importantly, data inference is more concerned with the data itself. According to the structure of privacy protection tactics (shown in Figure 6.3), minimise, hide, separate and abstract are the most relevant abstract tactics to achieve PPGoal 3. Since the treatment strategy of PPGoal 1 is to reduce the likelihood of the risk, the treatment strategy of PPGoal 3 is to reduce the severity of the risk. This means that minimise and abstract can be applied to reduce the severity of the risk of the privacy harm. Before analysing their associated tactics, the main type of operation performing in location data is *data analysis*, which is a type related to the Usage stage of the APDL model. Minimise is primarily concerned with data minimisation; the underlying goal of minimise is to limit usage of personal data to encourage the non-collection of purposeless data with the aim of reducing the severity of a privacy threat occurrence. Abstract aims to limit detail of personal data to reduce the severity of a privacy threat occurrence by summarising or grouping personal data to the coarsest granularity still useful for performing operations. As such, abstract's tactics are more relevant to achieve

PPGoal 3. By analysing the definitions of concrete tactics of abstract, group is an appropriate tactic as it aims to induce less detail from personal data prior to processing, by allocating into common categories. Accordingly, group is a concrete tactic that limits the identification of data subjects that can be made by linking driving patterns to particular data subjects.

Privacy Design Patterns

In this case, we refer to two different catalogues of privacy patterns: the initiative of the UC Berkeley School of Information[14] and the PRIPARE project[15].

To identify appropriate design patterns that describe the idntified tactics, design patterns described in the above catalogues are analysed with respect of their intents and/or problems, motivations, forces and contexts, and solutions.

PDPattern 1. The underlying goal of the isolate tactic is to prevent correlation. The imposed constraint on the usage operation is to process parts of personal data independently, without access or correlation to related parts. The User Data Confinement pattern describes the underlying goal of the isolate tactic. Its intent is to collect and/or process personal data in a domain that is physically controlled by the data subject. Its motivation gives an abstract scenario from the context of smart grid that reflects the generated quality scenario QScenario 1 – QScenario 4 in an abstract fashion. It can be applied in any context in which the collection and processing of personal data for legitimate purposes may pose threats to the privacy of data subjects. Its solution abstractly describes the isolate tactic by avoiding the central collection of personal data and shifting some of the processing of personal data to the user device. However, the application of this pattern has a consequence: EETS providers will require some guarantees from EETS users to reassure that road-usage tolls are correctly calculated and EETS users cannot commit fraud. This may involve the consideration of cryptographic algorithms and

[14] http://privacypatterns.org/
[15] https://privacypatterns.eu/

protocols. Further, this pattern implements the distribute tactic, which aims to partition personal data so that more access is required to process it.

PDPattern 2. The underlying goal of the dissociate tactic is to prevent the exposure of personal data. The imposed constraint on the disclosure operation is to remove the correlation between different items of personal data. The Anonymity Set pattern describes the underlying goal of the dissociate tactic. Its intent is to aggregate multiple entities into a set such that distinguishing between them becomes infeasible — it prevents location tracking and analyses the behaviour of users. Its motivation gives an abstract scenario from the context of healthcare that reflects the generated quality scenarios QScenario 5 – QScenario 8 in an abstract fashion. It can be applied in any context in which the processing of personal data for legitimate purposes may pose threats to the privacy of data subjects. Its solution abstractly describes the dissociate tactic by aggregating different entities into equivalence classes. Further, this pattern implements the hide tactic, which aims to prevent exposure by processing personal data randomly within a large enough group to reduce correlation.

PDPattern 3. The underlying goal of the group tactic is to limit detail. The imposed constraint on the usage operation is to induce less detail from personal data prior to processing, by allocating the data into common categories. The Location Granularity pattern describes the underlying goal of the group tactic. Its intent is to limit detail as much as possible by reducing its accuracy and precision, so that data is provided only in a coarse-grained manner. Its motivation gives an abstract scenario from the context of location-based services that reflects the generated quality scenario QScenario 9 in an abstract fashion. It can be applied in any context in which the processing of personal data for legitimate purposes may pose threats to the privacy of data subjects. Its solution abstractly describes the group tactic by adjusting the granularity of location data according to the amount of individuals that share the same location at the same moment — sharing only the necessary

level of granularity a service may be able to maintain the same functionality without requesting or distributing potentially sensitive data.

Privacy-Enhancing Technologies

In this case, we refer to various sources for PETs, such as the Stanford PET wiki[16] and the proceedings of the annual Privacy Enhancing Technologies Symposium[17].

To identify appropriate PETs (privacy measures) that implement the identified design patterns, PETs described in the above sources are analysed with respect to four dimensions: aspect, scenario, aim and data.

PMeasure 1. By selecting User Data Confinement as a privacy design pattern that describes the isolate tactic by processing personal data locally, an appropriate PET that realises this pattern needs to be chosen. To consider its consequences, a privacy-preserving cryptographic protocol needs to be considered to allow a EETS user to commit to a chosen value while keeping it hidden from others, with the ability to reveal the committed value later by a EETS provider. By analysing existing PETs, Cryptographic Commitments is the most relevant PET that implements the selected design pattern. The scenario dimension of the PET realises the actors identified in the motivation of the design pattern (user and service provider). In this case, we choose the 'mutual' scenario to describe the case of both EETS providers and EETS users not trusting each other. The aspect dimension of the PET is the same aspect considered in the intent of the design pattern (the content). In this case, the identity of EETS users are required when signing the contract for billing purposes. This indicates that anonymity is not an option. The focus is on how to protect the content (fine-grained location data). As such, the main approach to hide location data is to employ encryption and privacy-preserving cryptographic protocols. The aim dimension of the PET considers the privacy property emphasised in the intent of the design pattern. In this case, the focus is on unlinkability to ensure that 'location data' is not linkable with 'identification and contact data'.

[16] http://cyberlaw.stanford.edu/wiki/index.php/PET
[17] http://petsymposium.org/

The data dimension of the PET considers the same type of operations performed on personal data. In this case, the focus is on the collection and processing of location data to ensure that data is processed locally in the user domain.

PMeasure 2. By selecting Anonymity Set as a privacy design pattern that describes the dissociate tactic by removing the correlation between different items of personal data, an appropriate PET that realises this pattern needs to be chosen. By analysing existing PETs, K-anonymity Model is the most relevant privacy-enhancing technique that implements the selected design pattern. The scenario dimension of the PET realises the actors identified in the motivation of the design pattern (user and service provider). In this case, we choose the 'mutual' scenario to describe the case of both EETS providers and EETS users not trusting each other. The aspect dimension of the PET is the same aspect considered in the intent of the design pattern (the identity). To legally share driving patterns with third parties, those patterns need to be anonymous. Thus, the identity of EETS users are not required. The focus is on how to protect the identity of EETS users. The aim dimension of the PET considers the privacy property emphasised in the intent of the design pattern. In this case, the focus is on anonymity to ensure that EETS users cannot be distinguished. The data dimension of the PET considers the same type of operations performed on personal data. In this case, the focus is on the processing of 'location data' and 'identification and contact' to ensure that driving patterns are anonymous.

PMeasure 3. By selecting Location Granularity as a privacy design pattern that describes the group tactic by inducing less detail from personal data prior to processing by allocating into common categories, an appropriate PET that realises this pattern needs to be chosen. By analysing existing PETs, a Location Hierarchy is the most relevant privacy-enhancing technique that implements the selected design pattern. The scenario dimension of the PET realises the actors identified in the motivation of the design pattern (user and service provider). In this case, we choose the 'mutual' scenario to describe the case of EETS providers and insurance providers not trusting each other. The aspect dimension of the PET is the same

aspect considered in the intent of the design pattern (the identity). In this case, the identity of EETS users are not required and their driving patterns need to be anonymous. The focus is on how to protect the identity. The aim dimension of the PET considers the privacy property emphasised in the intent of the design pattern. In this case, the focus is on anonymity to ensure that 'location data' is not linkable with 'identification and contact data'. The data dimension of the PET considers the same type of operations performed on personal data. In this case, the focus is on the collection and processing of location data to ensure that it is shared at different levels of granularity.

Privacy Architectural Strategies

As previously mentioned in Section 6.3.5, architectural strategies are *risk mitigation strategies*. As such, strategies need to be identified with the aim of reducing the likelihood and/or severity of threat events that contribute to the occurrence of the risk of the privacy harm. These strategies need to address the exploitability of primary assets' vulnerabilities and the adverse effects of data inference that lead to the identification of driving patterns that may reveal sensitive information, such as health conditions, religious beliefs, political affiliation. etc.

In this case, we identified only one architectural tactic, design pattern and PET for each privacy protection goal. Accordingly, only one architectural strategy can be identified to achieve the corresponding protection goal, as per Table 7.20. The concept of strategy tree is useful when a threat scenario has more than treatment strategies and the corresponding quality scenario has more that tactics. In this case, we outlines the identified strategies without a strategy tree.

PAStrategy 1. Data Isolation Strategy is identified to achieve PPGoal 1 by processing parts of EETS users' personal data independently, without access or correlation to related parts.

Strategy	Goal	Tactic	Pattern	PET
PAStrategy.1	PPGoal 1	PATactic 1	PDPattern 1	PMeasure 1
PAStrategy.2	PPGoal 2	PATactic 2	PDPattern 2	PMeasure 2
PAStrategy.3	PPGoal 3	PATactic 3	PDPattern 3	PMeasure 3

Table 7.20: The identified privacy-enhancing strategies in the context of EETS.

PAStrategy 2. Data Disassociation Strategy is identified to achieve PPGoal 2 by aggregating multiple entities into a set such that distinguishing between them becomes infeasible.

PAStrategy 3. Data Grouping Strategy is identified to achieve PPGoal 3 by inducing less detail from EETS users' driving patterns prior to processing by allocating into common categories.

A low level of privacy protection can be specified by identifying architectural strategies that apply the tactics of inform, control, enforce and demonstrate, with the aim of protecting 'identification and contact data' and 'location data' from accidental disclosure and misuse.

7.5 Summary

In this chapter, we have illustrated a principled approach for engineering PbD. As previously mentioned in Chapter 2, PbD emphasises that privacy concerns should be addressed early in the requirement analysis phase, as well as when designing the software architecture. As such, our approach provides insight into the process through which privacy requirements are integrating into the early stages of the SDLC. In particular, we illustrated how the main activities of our approach can be mapped onto the requirements analysis and design phases of the SDLC.

Our approach supports the translation of abstract privacy principles, privacy risk models and privacy mechanisms into implementable requirements. It also supports the integration of these activities into the early stages of SDLC where architectural decisions about data processing can be still made. It aids software

engineers in analysing functional requirements, eliciting privacy requirements and making appropriate design decisions that fulfil these requirements.

In the next chapter, we will evaluate our principled approach with reference to the evaluation framework of Chapter 3. In addition, we will discuss the significance of the contributions of this dissertation.

Designing privacy into systems at the beginning of the development process necessitates the effective translation of privacy principles, models, and mechanisms into system requirements.

— Stuart S. Shapiro

8 Evaluation and Discussion

Contents

8.1 Introduction

This chapter presents a process evaluation of our approach by discussing the efficacy of each technique in the context of the case study. It further gives a relatively detailed discussion of the significance of the contributions of the dissertation. In particular, Section 8.2 critically analyses the contributions with reference to the established evaluation framework of Chapter 3 to evaluate how applicable, useful and successful these contributions are in addressing the research problem. Then, Section 8.3 gives a detailed discussion of the significance of the contributions with reference to the

criteria of Chapter 2 to illustrate how these contributions fit with the existing body of knowledge. Finally, Section 8.4 gives a brief summary of the chapter.

8.2 Evaluation

In this section, we critically analyse the activities of our approach and the roles played by the applied techniques and/or tools to evaluate how applicable and useful our approach is in relation to the factors and criteria of Chapter 3 (see Section 3.3.2).

8.2.1 Data-Processing Representation

The activity of data-processing representation consists of two steps. The refinement step is aided by *the APDL model*, which provides all stages through which personal data moves during its lifecycle, associated activities and involved actors, to operationalise the abstract purpose. The refinement of the abstract purpose is straightforward: the APDL model guides the identification of data-processing activities and involved actors by considering all the stages of data lifecycle. The APDL model is an abstract model that can be instantiated to develop context-relative personal data lifecycle models for each particular domain. It is context-dependent: not all application domains have all of the stages. According to these features, flexibility and scalability are supported by the refinement step. The APDL model considers all the stages of the personal data lifecycle (from collection to destruction). It also played a vital role in operationalising the abstract purpose and refining it into concrete purposes. Thus, completeness and capability are supported by the refinement step.

The representation step is aided by *the UML profile for the APDL model* as a graph-based technique, which provides a generic extension mechanism for customising UML models, to represent the abstract and concrete purposes, together with the key aspects of abstract privacy principles in a fine-grained manner. The representation of data-processing activities is straightforward as it is based on the refinement step. The UML profile can be created using any UML CASE tool supporting UML 2.0. It can also be used for different application domains. As such,

flexibility and scalability are supported by the representation step. The UML profile for the APDL model considers all the stages of the personal data lifecycle (from collection to destruction). It also plays a vital role in modelling the abstract purpose and concrete purposes, together with the key aspects of abstract privacy principles. Hence, completeness and capability are supported by the representation step.

8.2.2 Data-Centric Threat Modelling

The activity of data-centric threat modelling consists of six steps. The vulnerability analysis step uses *context-relative processing norms* as an analysis technique to articulate the reasonable expectations of privacy as a baseline model, upon which possible privacy vulnerabilities can be derived. They can be documented using any suitable tool. They can also be established for any data processing initiative in either scale. Thus, flexibility and scalability are supported by the vulnerability analysis step. Processing norms are established in relation to data-processing activities in each stage of the APDL model. This means that completeness is supported. However, the establishment of these processing norms presents challenges. In particular, processing principles cannot be easily derived from applicable legal frameworks. As such, the capability is partially supported by the vulnerability analysis step.

The threat analysis step applied *the adapted taxonomy of adverse privacy events* as a technique to identify potential threat events. The adapted taxonomy is built upon on the stages of the APDL model. It is a generic classification: it can be represented by different ways (e.g. a tree or a tabular form) and used for various domains. Accordingly, flexibility and scalability are supported by the threat analysis step. The adapted taxonomy characterises adverse privacy events by a set of attributes according to the nature of data-processing activities in each stage of the APDL model. As such, completeness is supported. However, the identification of threat events is not straightforward: the main types of adverse privacy events are too fine-grained — many of these events involve basically the same methods at a technical level. This implies that the capability is partially supported by the threat analysis step.

The harm analysis step is aided by *harm trees* as a technique to analyse privacy harms resulting from the identified threat events. The concept of harm trees is generic and can be constructed for various application domains. Harm trees can be represented in a different way (e.g. a tabular form) to be suitable for other domains. Harm trees can be used as a basis for the generation of reasonable threat scenarios for each privacy harm. Flexibility and scalability are supported by the harm analysis step. The construction of a harm tree considers all adverse effects of the threat events identified according to all the stages of the APDL model. This indicates that completeness is supported. Harm trees are ideal for generating all reasonable threat scenarios for each privacy harm. Based on these, capability is supported.

The severity and likelihood assessment steps are aided by *trace matrices* as a tool to illustrate the severity and likelihood of occurrence of privacy harms based on associated risk factors. In particular, trace matrices are used to manage and represent relationships between risk factors for each privacy harm, as well as between risk analysis and assessment, and other artefacts (quality scenarios). They can be easily changed by adopting any technique to suit various application domains. This implies that flexibility and scalability are supported by the severity and likelihood assessment steps. They consider all the identified privacy harms (for each harm tree) that consider all the identified threat events according to all the stages of the APDL model. As such, completeness and capability are supported by these steps.

The risk level assessment step is aided by *the risk map* as a tool to order the assessed risk of privacy harms based on their likelihood and severity. It can be easily changed by using a risk matrix to suit various application domains. Thus, flexibility and scalability are supported by the risk level assessment step. It considers all the identified privacy harms (for each harm tree) that consider all the identified threat events according to all the stages of the APDL model. As such, completeness and capability are supported by the risk level assessment step.

8.2.3 Privacy-Enhancing Strategies

The privacy-enhancing strategies activity consists of five steps. The first step is the articulation of privacy protection goals: it is aided by *quality scenarios* as a technique to derive desired protection goals. The generation of concrete quality scenarios was straightforward. However, the specification of the response element of a quality scenario requires analysing the level of the identified risk of a privacy harm with respect to the capability of threat sources, the vulnerabilities that might be exploited by those sources, and the adverse effects of the threat events. The determination of an appropriate treatment strategy is a decision that cannot be made by architects or software engineers; rather, it must be made with multiple stakeholders' consultation. As such, multiple quality scenarios were generated for each threat scenario. Quality scenarios can be used for, and easily documented by any tool suitable for, various application domains. Flexibility and scalability are supported by this step. Threat scenarios were generated according to all stages of the APDL model. Thus, completeness is supported by this step, whereas capability is partially supported by this step.

The second step is the identification of architectural tactics: it is aided by *the hierarchy of architectural tactics* as a technique, which provides a set of mapping criteria to select appropriate tactics that achieve the derived protection goals. The hierarchy is a generic structure that can be used for, and easily documented by any tool suitable for, various application domains. This implies that flexibility and scalability are supported by this step. The hierarchy is intended only to demonstrate existing tactics in the literature and that any list of tactics is necessarily incomplete. This implies that completeness is partially supported by this step. The identification of appropriate tactics is challenging, not least because existing tactics are described in different ways — capability is partially supported by this step.

The third step is the selection of appropriate design patterns: it is aided by *templates* as a technique, which provides four essential elements upon which a set of mapping criteria were established to select appropriate design patterns. The pattern template can be documented in a generic format that does not require

specifics tied to a particular application domain. Thus, flexibility, scalability and completeness are supported by this step. The analysis of existing design patterns requires specific skills — capability is partially supported by this step.

The fourth step is the selection of appropriate PETs: it is aided by *the dimension-based taxonomy of PETs* as a technique, which provides four essential dimensions upon which a set of mapping criteria were established to select appropriate PETs. It is a generic classification that can be used for, and easily documented by any tool suitable for, various application domains. Accordingly, flexibility, scalability and completeness are supported by this step. The taxonomy provides dimensions that can be used to analyse and compare existing PETs. The analysis of PETs requires specific skills — capability is partially supported by this step.

The final step is the identification of privacy architectural strategies: it is aided by *the strategy tree* as a technique to easily explore and navigate various design options from which architectural choices can be identified. The strategy tree is a generic structure that can be used for, and easily documented by any tool suitable for, various application domains. This means that flexibility and scalability are supported by this step. Each threat scenario is used as the basis for generating multiple quality scenarios according to treatment options. Threat scenarios were generated based on data-processing activities according to all the stages of the APDL model. Hence, completeness is supported by this step. The strategy tree shows a threat scenario and the corresponding quality scenarios, protection goals, tactics, design patterns and PETs hierarchically. It helps software engineers in analysing many alternatives and in reasoning critically about design decisions. These strategies are used as a means of specifying, implementing and justifying appropriate levels of privacy protection. Thus, capability is supported by this step.

In summary, the case study demonstrated that our approach are more convenient for, relevant to, or appropriate for, the context to which it was applied. It also demonstrated that the activities of our approach are successful in producing desired results. With respect to applicability, the approach fully supports the expected and mandatory criteria pertaining to flexibility and scalability respectively. With

			Applicability		Usefulness	
Activity	**Step**	**Technique/Tool**	**Flexibility**	**Scalability**	**Completeness**	**Capability**
Activity 1	Step 1.1	The APDL model	F Support	F Support	F Support	F Support
	Step 1.2	The UML profile for the APDL model	F Support	F Support	F Support	F Support
Activity 2	Step 2.1	Context-relative processing norms	F Support	F Support	F Support	P Support
	Step 2.2	The adapted taxonomy of adverse privacy events	F Support	F Support	F Support	P Support
	Step 2.3	Harm trees	F Support	F Support	F Support	F Support
	Step 2.4	Trace matrices	F Support	F Support	F Support	F Support
	Step 2.5	Trace matrices	F Support	F Support	F Support	F Support
	Step 2.6	The risk map	F Support	F Support	F Support	P Support
Activity 3	Step 3.1	Quality scenarios	F Support	F Support	F Support	P Support
	Step 3.2	The hierarchy of architectural tactics	F Support	F Support	P Support	P Support
	Step 3.3	Templates	F Support	F Support	F Support	P Support
	Step 3.4	The dimension-based taxonomy of PETs	F Support	F Support	F Support	P Support
	Step 3.5	Strategy trees	F Support	F Support	F Support	F Support

F Support: Full support, **P Support:** Partial support.

Table 8.1: The applicability and usefulness of our principled approach.

respect to usefulness, it fully and/or partially supports the mandatory and essential criteria pertaining to completeness and capability respectively, as per Table 8.1.

8.3 Discussion

In this section, we give a relatively detailed discussion of the significance of each contribution with reference to the criteria of Chapter 2. In addition, we show how they impact our understanding of the research problem and contribute to filling the existing gaps.

8.3.1 The UML Profile for the APDL Model

The development of appropriate modelling techniques, which are built upon conceptual models that define privacy-related concepts, is a necessary step to capture privacy-related issues and facilitate the ways in which technical solutions can be developed to comply with the principles of legal frameworks and standards. The UML profile for the APDL model (Chapter 4) serves as a modelling technique that represents data-processing activities in a way that is amenable to compliance checking and risk analysis, with the possibility of expressing how activities are performed, their effects in terms of changes of states, when they take place in terms of lifecycle stages, and where they take place in terms of lifecycle roles.

With respect to compliance checking, the UML profile is built upon a conceptual model that distinguishes between the main types of operations that can be performed on personal data (this meets criterion **Cr.1** of Chapter 2). For each operation, it outlines various distinct activities in relation to the GPS principles, which harmonise various sets of the FIPPs into universal privacy principles that can be applied in a variety of contexts in various jurisdictions, with the aim of governing the behaviour of these operations (this meets criterion **Cr.2** of Chapter 2). The conceptual model also captures the key aspects of abstract privacy principles in terms of concepts, actions and constraints. As such, it provides all necessary concepts for operationalising the abstract purposes for which personal data is collected, processed and/or disseminated (this meets criteria **Cr.3** and **Cr.4** of Chapter 2). The UML profile represents the concrete purposes in a hierarchical structure according to the stages of the personal data lifecycle as data-processing activities that can be assigned to different actors as responsibilities. Further, it represents involved actors, along with their assigned roles and responsibilities, and captures how those actors participate in data-processing activities according to their roles. Specifying abstract purposes in such a way (as objectives) is better than specifying them in terms of entities and activities, specifically to derive the minimum amount necessary of personal data from the actions of concrete purposes (data-processing activities) that fulfils the specified purposes (this meets criterion **Cr.5** of Chapter 2). Thus, the APDL model facilitates the application of the principle of separation of duties that manages the responsibility assignment by determining who is responsible for which lifecycle stage and what is the level of authority with respect to the decisions and activities performed when data is collaboratively collected and processed by multiple stakeholders in multiple domains. In addition, the APDL model facilitates the application of the principle of data minimisation as a foundational step for privacy engineering by specifying the processed data items in each single atomic action within an activity. This, in turn, helps analyse and restrict the processing of personal data to the minimum amount necessary according to the purpose of each concrete action. This can support the second activity of applying data minimisation

strategies of [50] (see Section 2.4.3). Further, it helps derive and specify the constraints through which the abstract purpose can be operationalised. These constraints can be expressed using OCL: they can be used to specify conditions on both concepts and actions identified in the conceptual modelling. The constraints are used to establish a set of suitable rules against which privacy compliance checking can be performed (this meets criteria **Cr.6** and **Cr.4** of Chapter 2).

With respect to risk analysis, the UML profile defines the *scope* of analysis under consideration — i.e. the extent of the problem space that the activities of our approach deal with, or to which it is relevant (the activities of risk analysis and assessment upon which the activities of architectural strategies are identified). By refining each abstract purpose into a set of concrete purposes that are expressed in terms of activities, actions, events, roles and actors, it represents these concepts in relation to the stages of the APDL model in a way that covers all relevant and necessary data-processing activities from collection to destruction. These concepts set the boundaries that limit the processing of personal data. The UML profile also helps establish the *context* in which personal data is collected, processed and disseminated. It considers the external and internal context to describe the environment in which risk treatment decisions are made. The external and internal context helps understanding relevant processing operations, roles and responsibilities, the technical environment and the factors influencing design decisions (i.e. legal and regulatory factors, social factors and contractual factors). It models the external context by capturing social, legal and technical aspects, including stakeholders' perceptions and expectations. It considers social factors by refining abstract purposes according to context-relative processing norms that reflect reasonable expectations of privacy. It also models the key aspects of abstract privacy principles as constrains in terms of invariants, pre- and post-conditions. It models the internal context by capturing organisational aspects by refining abstract purposes into concrete purposes as legitimate objectives. These objectives are expressed in terms of personal data and operations performed on this data. It also considers roles and responsibilities by grouping logically-related concrete purposes as roles to which involved actors are

assigned. It also considers personal data and operations (classified into collection, processing and dissemination operations) performed on this data. The UML profile also helps represent data-processing activities at different levels of *granularity*. This means that processing operations can be represented at two different levels of detail: coarse-grained and fine-grained levels. The former refers to representing data processing in terms of personal data and processing operations (abstract activities) classified according to the stages of the APDL model. The latter refers to representing each processing operation in terms of personal data, activities, actions, events, roles and actors.

As discussed in Chapter 2, some of existing approaches to PbD (whether they are goal-oriented or risk-based) use different modelling techniques, while others use implicit or explicit conceptual models for data processing. For example, Antignac *et al.* [13] developed privacy-aware data flow diagrams (PA-DFDs) that extend standard DFDs by a set of privacy-related concepts to be used as a modelling technique, with the aim of bridging the semantic gap between policy-makers and software engineers. The PA-DFD model is built upon a conceptual model for personal data processing that considers three types of activities a system may perform on data: communication, computation and storage. The performance of these activities is modelled as the occurrence of events: collect, store, retrieve, use, disclose and erase. However, it is mainly intended for compliance checking. The LINDDUN methodology [70] uses standard DFDs as a modelling technique to guide privacy threat analysis by mapping privacy threats to DFD elements. PRIAM [62] uses DFDs as a technique to represent data flows as a part of system characterisation. However, DFDs are not equipped by privacy-related concepts that help capture privacy-related issues. In contrast, other approaches use implicit or explicit conceptual models that represent data-processing activities. The PFSD framework [27] refers to three types of processing activities (transfer, storage and processing) at a high level of abstraction. Privacy design strategies [3] were refined based on an implicit conceptual model for data processing that considers three types of activities a system may perform on data: communication, computation

and storage. The performance of these activities is modelled as a set of seven mutually exclusive actions: operate, store, retain, collect, share, change and breach. Other approaches do not either use modelling techniques or conceptual models, such as the CNIL methodology [5].

As such, the UML profile for the APDL model provides a foundation for capturing privacy concerns in a comprehensive manner (including all the stages through which personal data moves during its lifecycle: from collection to destruction).

8.3.2 Data-Centric Threat Modelling

To discus how the data-centric threat modelling approach captures privacy concerns in a contextual and comprehensive manner, we first explain what we mean by contextuality and comprehensiveness in this context. Contextuality is the extent to which data-processing activities are represented in a way that relates to the internal and external contexts surrounding the processing of personal data — it is a required level to be complied with or reached. Comprehensiveness is the extent to which data-processing activities are represented in a thorough manner over a broad scope that includes the relevant and necessary stages that need to be dealt with.

As previously discussed in Chapter 5, the privacy risk model clearly defines the key risk factors that have impacts on the privacy of data subjects at an appropriate level of detail. It goes beyond traditional security risk models to take into consideration the dynamic and contextual nature of privacy. It characterises these factors by sets of nominal and assessable attributes to be used as inputs to estimate risk levels. It also established conceptual relationships between these factors (this meets criterion **Cr.7** of Chapter 2).

In the context of privacy and data protection, a number of privacy risk-management processes, frameworks and methodologies have been proposed, such as PRIAM [62], the CNIL Methodology [5] and the Privacy Risk Management (PRM) [108], which are based on the ISO 31000 Risk management – Guidelines [109]. They vary with respect to their risk factors.

PRIAM uses the concept of risk sources to refer both to unauthorised entities processing personal data and to entities with legitimate processing capabilities. The CNIL methodology uses the concept of risk sources to refer to those who act, accidentally or deliberately, on the supporting assets, on which the primary assets rely. Based on its risk model, risk sources who act, accidentally or deliberately, on the primary assets are not modelled. As such, we refine these concepts to be used appropriately in the context of privacy at an appropriate level of abstraction. With regards to risk 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.

In addition, PRIAM and the CNIL methodology 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. Both methodologies refer to a non-exhaustive list of common categories of feared events that an analyst should consider. However, in our approach, we defined the adapted taxonomy of adverse privacy events as a classification of threat events.

The CNIL methodology 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, PRIAM use the concept of privacy weakness to refer to a weakness in the data protection mechanisms — whether this weakness is legal, technical or organisational. 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 data. In addition, we use the classification of assets of [5].

PRIAM and the CNIL methodology do not distinguish between privacy violations and harms. These methodologies focus on feared events and their potential impact — i.e. consequences that each feared event may have on the identity and privacy of data subjects and human rights or civil liberties. The CNIL methodology 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, PRIAM uses the concept of privacy harms with specific attributes and categories. In our approach, we use the same concept in more detail to identify privacy harms at a detailed level of abstraction according to the properties and boundaries identified in [28].

By using the UML profile of the APDL model as a means for establishing the context in which personal data is collected, processed and dissiminated (as illustrated in Section 8.3.1), the main steps of the analysis approach clearly illustrate how the key risk factors can be identified in relation to all the stages of the APDL model to ensure adequate coverage of the problem space. Such a representation plays a crucial role in establishing the context and understanding the nature, scope and purposes of data processing under consideration by identifying all the relevant information for privacy risk analysis and assessment. Crucially, complementing a vulnerability-oriented approach with an asset/impact-oriented approach improves the rigour of the analysis. The result of the analysis approach is a set of harm trees that represent the conceptual relationships between the risk factors. In particular, the concept of the harm tree can be used as a graph-based analysis technique to

develop and model multiple (or at least a reasonable set of) threat scenarios for each privacy harm (this meets criterion **Cr.8** of Chapter 2).

The data-centric threat modelling approach uses context-relative processing norms as a means for articulating the reasonable expectations of privacy in each particular context as a baseline model that brings the social layer into view. The establishment of a context-relative processing norm for each data-processing activity in relation to all the stages of the APDL model facilitates the development of a complete baseline model to serve as the basis for deriving privacy vulnerabilities. The processing norms help articulate the contextual factors that influence the processing operations performed on personal data, which, in turn, provide a better understanding of reasonable expectations of privacy in a contextual manner in terms of appropriate collection, processing and dissemination of personal data that reflect social norms (this meets criterion **Cr.9** of Chapter 2). The processing norms help support the identification privacy vulnerabilities in a contextual and comprehensive manner.

The arrangement of the categories and sub-categories of adverse privacy events of Solove's taxonomy around the APDL model facilitates the identification of potential threat events in a structured, concrete and comprehensive manner. Specifically, threat events are identified by analysing the privacy vulnerabilities and associated threat sources for each context-relative processing norm grouped by the stages of the APDL model (this meets criterion **Cr.10** of Chapter 2).

The distinction between privacy harms and violations reflects that the presence of a privacy violation does not mean that it will necessarily create an actual privacy harm. Further, the application of these concepts uncovers that privacy impact can be decomposed into: negative impacts on data subjects and negative impacts on organisations. This helps the identification of treatment options in a commensurate manner (as will be further explained in Section 8.3.3). The characterisation of privacy harms and violations in relation to Calo's scheme [28] facilitates their identification in a structured manner. In particular, they can be identified by

analysing the adverse effects resulting from the threat events categorised according to the stages of the APDL model.

As discussed in Section 5.3, it is useful to generate multiple threat scenarios describing how threat events caused by the most likely threat sources, with the capability to exploit privacy vulnerabilities, can contribute to or cause a privacy harm. The case study showed that using harm trees as an analysis technique helps generate multiple threat scenarios. In particular, harm trees can represent the conceptual relationships among the key risk factors in a way that is analytically useful for assessment. Further, the characterisation of the risk factors by well-defined attributes (nominal and assessable) can facilitate the assessment of these factors in a systematic and traceable manner. In addition, the dependencies between the nominal and assessable attributes of each risk factor, and the dependencies between the nominal and assessable attributes of all risk factors facilitate their use as inputs to determine the levels of risk in risk assessments. In addition, the adoption of the fixed levels of scale, along with the corresponding values of [5] as assessment scales with refined and/or newly defined assessment rules that reflect the assessable attributes of the key risk factor facilitates their roles in risk assessments and their translation into qualitative terms for multiple stakeholders. These scales can be easily translated for multiple stakeholders and allow relative comparisons between values in different scales or even within the same scale (this meets criterion **Cr.11** of Chapter 2).

The result of the assessment approach is a risk map that locates the assessed risks according to their levels (likelihood, severity). The use of the risk map as an analysis technique helps determine the order in which the assessed risks should be managed according to their severity and likelihood. By adopting the risk map, the four distinguished levels help: prioritise the located risks in each level; and identify a set of treatment strategies according to all possibilities of their severity and likelihood. Treatment strategies are considered as a means for determining the degree to which privacy is required (as will be further explained in Section 8.3.3). Thus, the main steps of the assessment approach provide insight

into the identification of appropriate treatment strategies, not least because not all privacy risks warrant the same degree of attention.

8.3.3 Privacy-Enhancing Strategies

The use of the constructed harm trees (see Chapter 6) to generate a reasonable set of threat scenarios facilitates the generation of the quality scenarios by mapping the risk factors of the threat scenarios to the elements of the quality scenarios. By using the risk map (see Chapter 5), appropriate treatment strategies were identified based on the four distinguished levels. By establishing the context, external and internal context factors were considered. Treatment strategies, along with contextual factors, influence the degree to which privacy is required. Thus, the use of the threat scenarios as the basis for generating the quality scenarios facilitates the derivation of the desired privacy protection goals in a contextual manner. By distinguishing general quality attribute scenarios (those that are system independent) from concrete quality attribute scenarios (those that are specific to the particular system under consideration), the quality attribute of privacy protection can be characterised by a general scenario indicating the range of values its elements can take, from which concrete or system-specific scenarios can be derived (see Section 6.2). In order to translate the attribute characterisation into requirements for a particular system, the general scenario needs to be made system-specific by instantiating its elements. The concept of the general quality attribute scenario was used as an analysis technique to generate the most important concrete quality scenarios, from which the desired privacy protection goals can be derived in a contextual and non-reductive manner. It also helps characterise privacy protection to avoid nonoperational definitions and overlapping attribute concerns. By referring to a set of privacy protection goals, privacy protection as a quality attribute can be characterised by three abstract protection goals, from which concrete privacy protection goals can be derived (this meets criterion **Cr.12** of Chapter 2). Importantly, contextual factors that influence the degree to which privacy is required for a given context is considered in the early stages of the design process (when deriving desired protection goals). Further,

the relationship between these factors and the privacy controls that might be adopted was taken into account. Furthermore, the generation of the quality scenario facilitates the refinement of quality requirements to ensure the achievement of the desired privacy protection goals. Importantly, the adoption of abstract protection goals provides an abstraction level between the abstract privacy principles stated in legal frameworks and standards, and abstract architectural tactics.

The case study also showed that the identification of the appropriate architectural tactics was straightforward (in a structured manner based on the generated quality scenarios to ensure that each architectural tactic achieves a desired privacy protection goal). In particular, the structure of privacy protection tactics (shown in Figure 6.3) facilitates the determination of the relevant abstract tactics as these tactics are organised around three abstract protection goals. The risk levels of the privacy harm under consideration help facilitate the determination of the most relevant abstract tactics based on the target of whether to reduce the severity or likelihood of the risks of the privacy harm. In addition, the underlying goal of concrete tactics and the type of operations performed on personal data facilitate the determination of the most appropriate concrete tactic. Further, the adoption of abstract protection goals help categorising tactics in a way that facilitates its selection.

The establishment of a set of selection criteria help determine appropriate design patterns that describe the identified architectural tactics. In particular, these criteria facilitate the analysis and comparison of existing design patterns with respect to their (intent and/or problem, motivation, forces and context, and solution) dimensions. In addition, these criteria are established in a manner that allows matching with the underlying goals and processing operations of various tactics (this meets criteria **Cr.13** and **Cr.14** of Chapter 2).

The establishment of a set of selection criteria helped determine appropriate PETs that realise the identified design patterns. In particular, these criteria facilitate the analysis and comparison of existing PETs with respect to their (scenario, aspect, aim and data) dimensions. In addition, these criteria are established in a manner

that allows matching with the main elements of various design pattern templates (this meets criteria **Cr.13** and **Cr.14** of Chapter 2).

The use of architectural strategies as a means for mapping privacy requirements onto suitable software architectures helps specify, implement and justify the appropriate level of privacy protection as the default setting. As such, they can help software architects to reason critically about architectural decisions. To achieve the aim of PbD, these strategies aim to apply preventive measures rather than protective ones. Each architectural strategy is a treatment option identified based on a specific treatment strategy. Thus, the identified strategies are considered as architectural choices. The case study showed that architectural alternatives can be identified as architectural strategies that are all related to the same quality scenario (classified by treatment strategies: reducing likelihood, reducing severity or both).

8.4 Summary

In this chapter, we have analysed the activities of our approach to evaluate its applicability and usefulness in the context of the case study. We have also discussed the key findings presented in Chapter 7 that combine the contributions of Chapters 4, 5 and 6. We divided our discussion into three main themes: the UML profile for the APDL model, data-centric threat modelling and privacy-enhancing strategies. Based on the discussion, we illustrated the significance of the contributions by arguing how these methods and techniques aided by a prescribed language or notation accomplish privacy-engineering activities in relation to the SDLC. In addition, we explained how these contributions consistently contribute to filling the existing the identified gaps discussed in Chapter 2.

To this end, the principled approach for engineering PbD helps facilitate the translation of abstract privacy principles stated in legal framework and standards into context-relative protection goals that can be achieved by context-specific architectural strategies. This can help to bridge the semantic gap between technical and normative concepts that often leads to a disconnect between policy-makers and software engineers. It can also be used to complement PIA methodologies

to sufficiently support privacy risk assessments and illustrate precisely how the technical part of the PIA can be conducted. It also goes beyond traditional security assessments to consider the nature of the risks arising from the processing of personal data.

Security and privacy tend to be articulated at a level of abstraction that often makes their specific manifestations less than obvious.

— Stuart S. Shapiro

9 Conclusion

Contents

9.1 Introduction

This chapter summarises the contributions of the dissertation, discusses their weaknesses and limitations, and outlines directions for further work. Section 9.2 summarises the contributions of this dissertation in relation to the main challenges of engineering PbD discussed in Chapter 2. Then, Section 9.3 gives a brief discussion of how these contributions answered the research question stated in Chapter 1. Section 9.4 highlights the shortcomings and limitations of these contributions. Finally, Section 9.5 outlines opportunities for future work in this area.

9.2 Contributions

In this dissertation, we have presented a principled approach for engineering PbD that consists of three main activities. The first activity aims to represent the processing of personal data in a way that is amenable to risk analysis and compliance checking. The second activity aims to holistically identify and systematically assess potential privacy risks in a contextual and comprehensive manner. The third activity aims to provide various levels of privacy protection.

In this section, we reflect on the significance of these contributions in relation to: the main challenges of engineering PbD outlined in Section 2.5.1; and the principles of PbD outlined in Section 2.3.2.

9.2.1 The Challenges of Engineering PbD

The UML Profile for the APDL Model

The UML profile for the APDL model was developed as a technique for representing personal data processing in a way that is amenable to compliance checking. It serves as a common language that supports a meaningful participation of, and facilitates communication between, multiple stakeholders. It describes 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. Thus, the UML profile for the APDL model serves as a means for checking and demonstrating privacy compliance with legal frameworks and standards (this addresses challenge **Ch.2** of Chapter 2).

Data-Centric Threat Modelling

The data-centric threat modelling approach defines a privacy risk model that goes beyond traditional security risk models to take into consideration the dynamic and contextual nature of privacy. It helps to: understand and translate various perceptions into operational requirements; and understand and consider multiple stakeholders' expectations and concerns (this addresses challenges **Ch.1** and **Ch.3** of Chapter 2).

Privacy-Enhancing Strategies

The privacy-enhancing tactical approach is built upon the concept of quality attribute scenarios. It adopts quality scenarios as a means for considering the main factors of determining the degree to which privacy required: legitimate objectives, stakeholders' expectations, legal frameworks and standards, appropriate threat models and technological capabilities (this addresses challenge **Ch.4** of Chapter 2).

The privacy-enhancing tactical approach defines the concept of architectural strategies as means for providing various levels of privacy protection. Architectural strategies are considered as a means for addressing privacy concerns at an architectural level, as well as for mapping privacy requirements onto suitable software architectures (this addresses challenge **Ch.5** of Chapter 2).

9.2.2 The Principles of PbD

As previously discussed in Chapter 7, our principled approach aims to translates the principles of PbD into engineering activities. We now analyse its activities with respect to the principles of PbD.

The activities of our principled approach proactively anticipate and commensurately address adverse privacy events before they occur in the early stages of the design process (Activity 1 – Activity 3 address principle **P.1** of Chapter 2).

The third activity of our approach provides various levels of privacy protection. In particular, it adopts the concept of architectural strategies to specify, implement and justify acceptable levels of privacy protection at an architectural level. These levels are built into the system by default (Activity 3 addresses principle **P.2** of Chapter 2).

The activities of our approach integrate privacy into the requirements analysis and design phases of the software development process represented in the waterfall model (Activity 1 – Activity 3 address principle **P.3** of Chapter 2).

The third activity of our approach adopts the concept of quality scenarios as a means for considering the main factors of determining the degree to which privacy required: legitimate objectives, stakeholders' expectations, legal frameworks and standards, appropriate threat models and technological capabilities (Activity

3 addresses principle **P.4** of Chapter 2). The generation of quality scenarios is conducted with reference to the identified threat scenarios for each stage of the APDL model (Activity 3 addresses principle **P.5** of Chapter 2).

The first activity of our approach adopts the UML profile for the APDL model to represent personal data, data-processing activities and involved actors for each stage of the data lifecycle. In particular, the Initiation stage of the APDL model specifies a 'complete processing plan' that determines the purposes for, and the manner in which, personal data is collected, processed and disseminated. The processing plan is considered as the basis for establishing a privacy notice to be communicated to data subjects (Activity 1 addresses principle **P.6** of Chapter 2). The Participation stage of the APDL model represents data-processing activities that provide data subjects with access to exercise control over personal data and specify the means by which data subjects can review, update and correct this data to ensure that it is accurate, complete and up-to-date (Activity 1 partially addresses principle **P.7** of Chapter 2).

The second activity of our approach defines a user-centric risk model, rather than a compliance-based risk model. It considers privacy harms, rather than privacy violations: it focuses on data subjects and societal vulnerabilities and harms rather than purely on violations of regulatory compliance requirements (Activity 2 partially addresses principle **P.7** of Chapter 2).

9.3 Research Question Evaluation

In this section, we critically analyse the contributions to determine how successful they are at answering the research question. The evaluation of the contributions confirms that the research question stated in Chapter 1 has been answered by illustrating how the principled approach for engineering PbD (described in Chapter 7) demonstrates the integration of the three contributions of the dissertation described in Chapters 4, 5 and 6.

In order to evaluate how successful the contributions of the dissertation are at answering the research question, we decompose the research question addressed by this dissertation, around which we center our research, into three clauses.

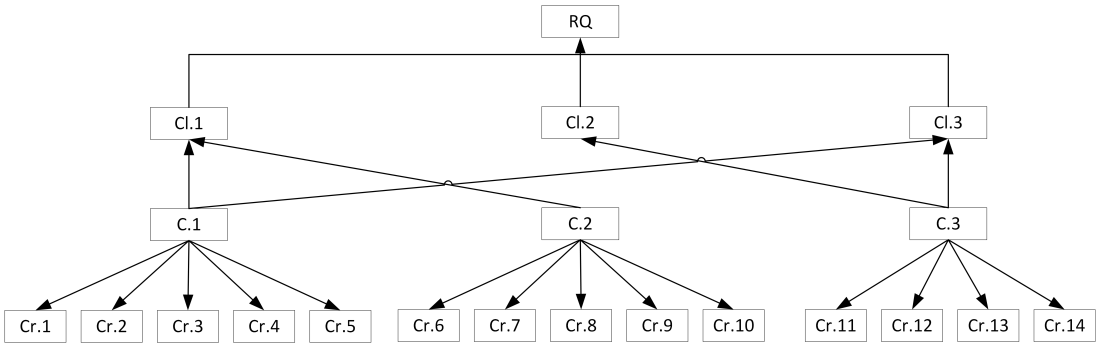

Figure 9.1: The mapping of the contributions onto the clauses of the research question and the established criteria of Chapter 2.

Cl.1 Can we develop an approach to PbD that complements software engineering methods to: capture privacy concerns in a comprehensive manner; [...]?

Cl.2 Can we develop an approach to PbD that complements software engineering methods to: [...] address these concerns at an architectural level; [...]?

Cl.3 Can we develop an approach to PbD that complements software engineering methods to: [...] reason about the compliance of architectural choices with legal frameworks and standards?

The first clause of the research question, ***Cl.1****,* was partially answered by presenting the UML profile for the APDL model (described in Chapter 4), which established the context in which personal data is collected, processed and disseminated, and represented data-processing activities in a way that is amenable to risk analysis and compliance checking. It was also answered by developing the data-centric threat modelling approach (described in Chapter 5), which described the main steps of identifying, analysing and assessing potential privacy risks in a contextual manner.

The second clause of the research question, ***Cl.2****,* was answered by presenting the privacy-enhancing tactical approach (described in Chapter 6), which described the main steps of identifying privacy-enhancing architectural strategies that are considered as a means for addressing privacy concerns at an architectural level.

The third clause of the research question, ***Cl.3****,* was answered by presenting the UML profile for the APDL model (described in Chapter 4), which represented

data-processing activities in a way that is amenable to risk analysis and compliance checking. It was also answered by developing the privacy-enhancing tactical approach (described in Chapter 6), which described the main steps of identifying and justifying privacy-enhancing architectural strategies based on a set of risk-treatment strategies.

Figure 9.1 shows how the three contributions are mapped onto the clauses of the research question and the established criteria of Chapter 2.

9.4 Shortcomings and Limitations

The UML Profile of the APDL model

The conceptual model, upon which the UML profile for the APDL model was built, classifies the main types of operations that can be performed on personal data. However, it does not precisely classify all possible types of each main operation to distinguish processing activities in each stage of the APDL model. The UML profile for the APDL model serves as a preliminary acquisition step to capture privacy-related concepts that support requirements analysis — a critical step in the SDLC. The most interesting possibility resulting from the UML profile is the integration of the APDL model with standard UML models. However, it is not clear how to integrate it with UML's other diagrams, such as use case diagrams, class diagrams, interaction diagrams, etc. The integration facilitates requirements traceability in both requirements analysis and design specification. It also facilitates OCL expressions specification and evaluation.

Data-Centric Threat Modelling

The data-centric threat modelling approach provides insights into the process through which potential privacy risks are holistically identified and systematically assessed. It adopts Solove's taxonomy [26] to characterise threat events. The main types of threat events are mapped to the main types of operations using the APDL model. However, the definitions of the types of threat events are not easily aligned with the types of activities in each stage of the APDL model. In addition, the approach adopts Nissenbaum's contextual integrity heuristic [29] to establish

context-relative processing norms. However, the establishment of these norms has practical challenges. In particular, processing principles need to be derived from legal frameworks and standards in an appropriate manner. The translation of these principles into a set of conditions expressed as constraints is a challenge in itself.

The assessment rules that specify the range of values the key risk factors can assume reflect the assessable attributes of these factors to facilitate their roles in risk assessments. However, two of these rules are established at a high level of abstraction, such as those rules that assess the intensity of a privacy harm and the capability of a threat source. The assessment of the intensity of a privacy harm considers three different factors: the extent of damage; the reversibility of threat events' consequences; and the duration of these consequences. Similarly, the assessment of the ability of a threat source considers five different factors: the skills, background knowledge, privileges, financial and technical resources. In practice, each of these factors needs to be estimated independently.

The assessment approach is semi-quantitative: it associates a numeric score with points on an otherwise descriptive scale (e.g. 5. Maximum). It yields a more structured ranking of risks than purely qualitative approaches. The numeric values are assumed subjectively; they should not be relied upon to compare relative risks with same levels. Further, the approach considers only privacy impact to data subjects (which result from privacy harms), but does not consider privacy impact to organisations (which result from privacy violations).

Privacy-Enhancing Strategies

The privacy-enhancing tactical approach provides insights into the process through which the assessed privacy risks are holistically and systematically treated in a commensurate manner. It adopts a set of abstract privacy protection goals (unlinkability, intervenability and transparency). However, the protection goal of unlinkability is given in a broader view to encompass anonymity, pseudonymity, unobservability, undetectability and unlinkability. It is related to the requirement of data minimisation, which can be achieved by data reduction, generalisation, data

hiding, separation and isolation. This means that the majority of the derived context-relative protection goals and data-oriented tactics are related to unlinkability.

The underlying goals of the privacy design strategies of [3] (we consider them as abstract tactics) cannot be easily linked to the abstract protection goals. Further, the strategies are defined with respect to the main types of actions performed on personal data. However, the privacy-enhancing tactical approach adopts the main types of operations defined in the APDL model along with their constraints. Thus, the mapping of these different types is not straightforward.

Furthermore, existing design patterns and PETs are described in different ways without referring to a standard template. Thus, the use of the established mapping criteria to select appropriate tactics, patterns and PETs is challenging.

9.5 Future Work

Based on the discussion of Chapter 8, and the shortcomings and limitations of Section 9.4, it is possible to extend the work presented in this dissertation in a number of ways:

- The conceptual model, upon which the UML profile for the APDL model was built, can be extended to classify possible activities for each type of processing operation. Such a classification would facilitate the mapping of the main types of adverse events of the Solove's taxonomy to the main types of operations of the APDL model through the corresponding activities.

- Investigating the integration of the UML profile for the APDL model and standard UML models would be worthwhile. It would be interesting to develop a complete UML profile to represent privacy-related concepts using the standard extension mechanisms of the UML meta-model [102]. Stereotypes, along with tag definitions, are used to specify key aspects of privacy as assumptions and requirements, and constraints are used to specify criteria to evaluate whether the requirements are satisfied by the system design.

- The privacy risk model can be refined by characterising the capability of a threat source and the intensity of a privacy harm by a set of concrete and well-defined attributes. Adding such a low level of detail helps facilitate their roles in risk assessments and their translation into qualitative terms for multiple stakeholders.

- An additional area of interest would be to extend the assessment approach to assess the impact of privacy violations on organisations.

- It would be interesting to develop a conceptual model that defines the main protection goals, along with their meanings and properties, for privacy engineering, and establishes the conceptual relationships among these goals. Such a model could support the use of these goals as a universal or agreed-upon set to be used as a means for mapping the underlying goals of abstract and concrete tactics to abstract privacy principles stated in legal frameworks and standards.

- Further research would be useful in investigating how the main types of processing operations of the APDL model, along with the corresponding activities and constraints, may help refine the definitions of abstract and concrete tactics.

- It might be possible to built a decision-based support system (as a tool that supports the third activity of our architectural approach) that considers the main aspects of privacy design patterns and PETs, and deduces a list of relevant patterns and PETs. Such a system would require establishing and modelling precise mapping and/or selection criteria. Further, it would require universal or agreed-upon catalogues for privacy design patterns and PETs that describe these artefacts in a way that is amenable to analysis and comparison.

Bibliography

[1] Y. Onn, M. Geva, Y. Druckman, A. Zyssman, R. Timor, I. Lev, A. Maroun, T. Maron, Y. Nachmani, Y. Simsolo, S. Sicklai, A. Fuches, M. Fishman, S. Packer, and L. Pery, "Privacy in the Digital Environment," *Haifa Center of Law & Technology*, pp. 1–12, 2005.

[2] M. F. Dennedy, J. Fox, and T. Finneran, *The Privacy Engineer's Manifesto: Getting from Policy to Code to QA to Value*. Apress, 2014.

[3] M. Colesky, J. H. Hoepman, and C. Hillen, "A Critical Analysis of Privacy Design Strategies," in *Proceedings of the 2016 IEEE Security and Privacy Workshops (SPW)*, pp. 33–40, IEEE, 2016.

[4] M. Hansen, M. Jensen, and M. Rost, "Protection Goals for Privacy Engineering," in *Security and Privacy Workshops (SPW)*, pp. 159–166, IEEE, 2015.

[5] 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.

[6] Official Journal of the European Union, "General Data Protection Regulation." `http://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32016R0679&from=EN`, 2016.

[7] A. Cavoukian, "Creation of a Global Privacy Standard." `https://www.ipc.on.ca/images/Resources/gps.pdf`, 2006.

[8] J. Heurix, P. Zimmermann, T. Neubauer, and S. Fenz, "A Taxonomy for Privacy Enhancing Technologies," *Computers & Security*, vol. 53, pp. 1–17, 2015.

[9] S. Spiekermann, "The Challenges of Privacy by Design," *Communications of the ACM*, vol. 55, no. 7, pp. 38–40, 2012.

[10] A. Cavoukian, "Privacy by Design." `https://www.privacybydesign.ca/content/uploads/2009/01/privacybydesign.pdf`, 2009.

[11] Secretary's Advisory Committee on Automated Personal Data Systems, *Records, Computers and the Rights of Citizens: Report*. [Cambridge? Mass.]: [MIT Press], 1973.

[12] A. Cavoukian, S. Shapiro, and R. J. Cronk, "Privacy Engineering: Proactively Embedding Privacy, by Design." `https://www.ipc.on.ca/wp-content/uploads/Resources/pbd-priv-engineering.pdf`, 2014.

[13] T. Antignac, R. Scandariato, and G. Schneider, "A Privacy-Aware Conceptual Model for Handling Personal Data," in *Leveraging Applications of Formal Methods, Verification and Validation: Foundational Techniques* (T. Margaria and B. Steffen, eds.), vol. 9952 of *Lecture Notes in Computer Science*, pp. 942–957, Springer, 2016.

[14] S. Gürses, C. Troncoso, and C. Diaz, "Engineering Privacy by Design," *Computers, Privacy & Data Protection*, vol. 14, no. 3, p. 25, 2011.

[15] M. C. Oetzel and S. Spiekermann, "A Systematic Methodology for Privacy Impact Assessments: A Design Science Approach," *European Journal of Information Systems*, vol. 23, no. 2, pp. 126–150, 2014.

[16] N. Notario, A. Crespo, Y. S. Martín, J. M. Del Alamo, D. Le Métayer, T. Antignac, A. Kung, I. Kroener, and D. Wright, "PRIPARE: Integrating Privacy Best Practices into a Privacy Engineering Methodology," in *Security and Privacy Workshops (SPW), 2015 IEEE*, pp. 151–158, IEEE, 2015.

[17] S. Gürses and J. M. del Alamo, "Privacy Engineering: Shaping an Emerging Field of Research and Practice.," *IEEE Security & Privacy*, vol. 14, no. 2, pp. 40–46, 2016.

[18] R. L. Finn, D. Wright, and M. Friedewald, "Seven types of privacy," in *European Data Protection: Coming of Age*, pp. 3–32, Springer, 2013.

[19] A. F. Westin, *Privacy and Freedom*. New York: Atheneum, 1st ed., 1967.

[20] J. Kang, "Information Privacy in Cyberspace Transactions," *Stanford Law Review*, vol. 50, no. 4, pp. 1193–1294, 1998.

[21] S. Davies, *Big brother : Britain's web of surveillance and the new technological order*. London: Pan, 1996.

[22] R. Clarke, "Introduction to dataveillance and information privacy, and definitions of terms." `http://www.rogerclarke.com/DV/Intro.html`, 1997.

[23] D. J. Solove, "Conceptualizing Privacy," *California Law Review*, vol. 90, no. 4, pp. 1087–1155, 2002.

[24] S. D. Warren and L. D. Brandeis, "The Right to Privacy," *Harvard Law Review*, vol. 4, no. 5, pp. 193–220, 1890.

[25] W. L. Prosser, "Privacy," *California Law Review*, vol. 48, no. 3, 1960.

[26] D. J. Solove, "A Taxonomy of Privacy," *University of Pennsylvania Law Review*, vol. 154, no. 3, pp. 477–564, 2006.

[27] S. Spiekermann and L. F. Cranor, "Engineering Privacy," *IEEE Transactions on Software Engineering*, vol. 35, no. 1, pp. 67–82, 2009.

[28] M. R. Calo, "The Boundaries of Privacy Harm," *Indiana Law Journal*, vol. 86, pp. 1131–1162, 2011.

[29] H. F. Nissenbaum, *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press, 2009.

[30] C. Fried, "Privacy," *The Yale Law Journal*, vol. 77, no. 3, pp. 475–493, 1968.

[31] A. M. Froomkin, "The Death of Privacy?," *Stanford Law Review*, vol. 52, no. 5, pp. 1461–1543, 2000.

[32] American Institute of Certified Public Accountants and Canadian Institute of Chartered Accountants (AICPA/CICA), "Generally Accepted Privacy Principles." `https://www.cippguide.org/2010/07/01/generally-accepted-privacy-principles-gapp/`, 2009.

[33] M. Petković and W. Jonker, *Security, Privacy, and Trust in Modern Data Management*. Springer, 2007.

[34] Cavoukian, A., "Privacy by Design: The 7 Foundational Principles - Implementation and Mapping of Fair Information Practices." `https://www.ipc.on.ca/english/Resources/Discussion-Papers/Discussion-Papers-Summary/?id=953`, 2010.

[35] A. Cavoukian, *Privacy by Design ... Take the Challenge*. Office of the Information and Privacy Commissioner of Ontario, 2009.

[36] A. Cavoukian, "Privacy by Design [Leading Edge]," *IEEE Technology and Society Magazine*, vol. 31, no. 4, pp. 18–19, 2012.

[37] Borking, J. and Gardeniers, H. and Rossum, H. V., *Privacy Enhancing Technologies: The Path to Anonymity*. Registratiekamer, 1995.

[38] S. Kenny and J. J. Borking, "The Value of Privacy Engineering," *Journal of Information, Law and Technology (JILT)*, vol. 7, no. 1, pp. 203–219, 2002.

[39] S. Joyee De and D. Le Métayer, "A Refinement Approach for the Reuse of Privacy Risk Analysis Results," in *Proceedings of the 5th Annual Privacy Forum (APF 2017)*, pp. 52–83, Springer, 2017.

[40] J. J. Borking and C. Raab, "Laws, PETs and Other Technologies for Privacy Protection.," *Journal of Information, Law and Technology*, vol. 1, pp. 1–14, 2001.

[41] G. W. v. Blarkom, J. J. Borking, and J. G. E. Olk, "Handbook of Privacy and Privacy-Enhancing Technologies," *Privacy Incorporated Software Agent (PISA) Consortium*, 2003.

[42] T. Antignac and D. Le Métayer, "Privacy by Design: From Technologies to Architectures," in *Privacy Technologies and Policy: Second Annual Privacy Forum (APF 2014)*, pp. 1–17, Springer, 2014.

[43] M. Hafiz, "A Pattern Language for Developing Privacy Enhancing Technologies," *Software: Practice and Experience*, vol. 43, no. 7, pp. 769–787, 2013.

[44] A. Cavoukian, "Operationalizing Privacy by Design: A Guide to Implementing Strong Privacy Practices." `https://www.privacybydesign.ca/content/uploads/2013/01/operationalizing-pbd-guide.pdf`, 2012.

[45] Cavoukian, A., "Privacy by Design: The 7 Foundational Principles." `https://www.ipc.on.ca/wp-content/uploads/Resources/7foundationalprinciples.pdf`, 2009.

[46] G. Duncan, "Privacy by Design," *Science*, vol. 317, no. 5842, pp. 1178–1179, 2007.

[47] C. Bier, P. Birnstill, E. Krempel, H. Vagts, and J. Beyerer, "Enhancing Privacy by Design from a Developer's Perspective," in *Privacy Technologies and Policy: First Annual Privacy Forum (APF 2012)*, pp. 73–85, Springer, 2012.

[48] J. H. Hoepman, "Privacy Design Strategies," *ICT Systems Security and Privacy Protection*, pp. 446–459, 2014.

[49] Y. S. Martín, J. M. Del Alamo, and J. C. Yelmo, "Engineering Privacy Requirements: Valuable Lessons from Another Realm," in *2014 IEEE 1st Workshop on Evolving Security and Privacy Requirements Engineering (ESPRE)*, pp. 19–24, IEEE, 2014.

[50] S. Gürses, C. Troncoso, and C. Diaz, "Engineering Privacy by Design Reloaded.," 2015.

[51] A. Kung, "PEARs: Privacy Enhancing ARchitectures," in *Privacy Technologies and Policy: Second Annual Privacy Forum (APF 2014)*, pp. 18–29, Springer, 2014.

[52] Bass, L. and Clements, P. and Kazman, R., *Software Architecture in Practice*. Addison-Wesley, 2003.

[53] M. Hafiz, "A Collection of Privacy Design Patterns," in *Proceedings of the 2006 Conference on Pattern Languages of Programs (PLoP '06)*, pp. 1–13, ACM, 2006.

[54] M. Schumacher, "Security Patterns and Security Standards," in *Proceedings of the European Conference on Patterns Language of Programming (EuroPLoP'02)*, pp. 289–300, 2002.

[55] M. Schümmer, T., "The Public Privacy–patterns for Filtering Personal Information in Collaborative Systems," in *Proceedings of the Conference on Human Factors in Computing Systems (CHI2004)*, 2004.

[56] S. Romanosky, A. Acquisti, J. Hong, L. F. Cranor, and B. Friedman, "Privacy Patterns for Online Interactions," in *Proceedings of the 2006 Conference on Pattern Languages of Programs (PLoP '06)*, pp. 1–12, ACM, 2006.

[57] M. Sadicoff, M. M. Larrondo-Petrie, and E. B. Fernandez, "Privacy-aware Network Client Pattern.," in *Proceedings of the 2005 Conference on Pattern Languages of Programs (PLoP '05)*, pp. 1–12, ACM, 2005.

[58] S. Pearson and Y. Shen, "Context-Aware Privacy Design Pattern Selection," in *Proceedings of the International Conference on Trust, Privacy and Security in Digital Business*, pp. 69–80, Springer, 2010.

[59] M. Colesky, J. Caiza, J. Del Álamo, J. Hoepman, and Y. Martín, "A System of Privacy Patterns for User Control," in *Proceedings of the Symposium on Applied Computing (SAC 2018)*, ACM, 2018.

[60] C. Kalloniatis, E. Kavakli, and S. Gritzalis, "Addressing Privacy Requirements in System Design: The PriS Method," *Requirements Engineering*, vol. 13, no. 3, pp. 241–255, 2008.

[61] T. Breaux, *Introduction to IT Privacy*. IAPP, 2014.

[62] S. Joyee De and D. Le Métayer, "PRIAM: A Privacy Risk Analysis Methodology," in *11th International Workshop on Data Privacy Management and Security Assurance*, pp. 221–229, Springer, 2016.

[63] R. Clarke, "Privacy Impact Assessment: Its Origins and Development," *Computer Law and Security Review: The International Journal of Technology and Practice*, vol. 25, no. 2, pp. 123–135, 2009.

[64] D. Wright, "The State of the Art in Privacy Impact Assessment," *Computer Law & Security Review*, vol. 28, no. 1, pp. 54–61, 2012.

[65] D. Wright, K. Wadhwa, P. De Hert, and D. Kloza, "A Privacy Impact Assessment Framework for Data Protection and Privacy Rights." `http://www.piafproject.eu/Deliverables.html`, 2011.

[66] A. Fineberg and P. Jeselon, "A Foundational Framework for a Privacy by Design — Privacy Impact Assessment." `http://privacybydesign.ca/content/uploads/2011/11/PbD-PIA-Foundational-Framework.pdf`, 2011.

[67] BSI (Bundesamt für Sicherheit in der Informationstechnik), "IT-Grundschutz-Kataloge." `https://www.bsi.bund.de/DE/Themen/ITGrundschutz/ITGrundschutzKataloge/itgrundschutzkataloge_node.html`, 2011.

[68] 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.

[69] European Commission, "Data Protection Impact Assessment Template for Smart Grid and Smart Metering systems." http://ec.europa.eu/energy/sites/ener/files/documents/2014_dpia_smart_grids_forces.pdf, 2014.

[70] M. Deng, K. Wuyts, R. Scandariato, B. Preneel, and W. Joosen, "A Privacy Threat Analysis Framework: Supporting the Elicitation and Fulfillment of Privacy Requirements," *Requirements Engineering*, vol. 16, no. 1, pp. 3–32, 2011.

[71] H. J. Smith, S. J. Milberg, and S. J. Burke, "Information Privacy: Measuring Individuals' Concerns about Organizational Practices," *MIS Quarterly*, vol. 20, no. 2, pp. 167–196, 1996.

[72] S. M. A. Burney, "Inductive and Deductive Research Approach," 2008.

[73] P. Runeson and M. Höst, "Guidelines for Conducting and Reporting Case Study Research in Software Engineering," *Empirical Software Engineering*, vol. 14, no. 2, pp. 131–164, 2009.

[74] S. Easterbrook, J. Singer, M. A. Storey, and D. Damian, "Selecting Empirical Methods for Software Engineering Research," in *Guide to Advanced Empirical Software Engineering*, pp. 285–311, Springer, 2008.

[75] T. Lethbridge, S. Sim, and J. Singer, "Studying Software Engineers: Data Collection Techniques for Software Field Studies," *Empirical Software Engineering*, vol. 10, no. 3, pp. 311–341, 2005.

[76] M. V. Zelkowitz and D. R. Wallace, "Experimental Models for Validating Technology," *IEEE Computer*, vol. 31, no. 5, pp. 23–31, 1998.

[77] C. Robson, *Real World Research*. Blackwell, 2002.

[78] R. Davison, M. G. Martinsons, and N. Kock, "Principles of Canonical Action Research," *Information Systems Journal*, vol. 14, no. 1, pp. 65–86, 2004.

[79] D. Walsh and S. Downe, "Meta-synthesis Method for Qualitative Research: A Literature Review," *Journal of Advanced Nursing*, vol. 50, no. 2, pp. 204–211, 2005.

[80] E. J. Erwin, M. J. Brotherson, and J. A. Summers, "Understanding Qualitative Metasynthesis: Issues and Opportunities in Early Childhood Intervention Research," *Journal of Early Intervention*, vol. 33, no. 3, pp. 186–200, 2011.

[81] D. L. Finfgeld, "Metasynthesis: The State of the Art — So Far," *Qualitative Health Research*, vol. 13, no. 7, pp. 893–904, 2003.

[82] S. Atkins, S. Lewin, H. Smith, M. Engel, A. Fretheim, and J. Volmink, "Conducting A Meta-Ethnography of Qualitative Literature: Lessons Learnt," *BMC Medical Research Methodology*, vol. 8, no. 1, p. 21, 2008.

[83] Z. Masood, S. Xuequn, and J. Yousaf, "Usability Evaluation Framework for Software Engineering Methodologies," *Lecture Notes on Software Engineering*, vol. 2, no. 3, pp. 225–232, 2014.

[84] B. A. Kitchenham, "Evaluating Software Engineering Methods and Tools Part 1: The Evaluation Context and Evaluation Methods," *ACM SIGSOFT Software Engineering Notes*, vol. 21, no. 1, pp. 11–14, 1996.

[85] B. A. Kitchenham and L. Jones, "Evaluating Software Engineering Methods and Tools Part 6: Identifying and Scoring Features," *ACM SIGSOFT Software Engineering Notes*, vol. 22, no. 2, pp. 16–18, 1996.

[86] S. Hesari, H. Mashayekhi, and R. Ramsin, "Towards a General Framework for Evaluating Software Development Methodologies," in *Proceedings of the 2010 IEEE 34th Annual Computer Software and Applications (COMPSAC) Conference*, pp. 208–217, IEEE, 2010.

[87] European Commission, "The European Citizens' Initiative." http://ec.europa.eu/citizens-initiative/public/welcome, 2012.

[88] European Commission, "Commission Implementing Regulation (EU) No 1179/2011 of 17 November 2011 laying down technical specifications for online collection systems pursuant to Regulation (EU) No 211/2011 of the European Parliament and of the Council on the Citizens' Initiative." http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2011:301:0003:0009:EN:PDF, 2011.

[89] European Commission, "Regulation (EU) No 211/2011 of the European Parliament and of the Council on the Citizens' Initiative." http://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:02011R0211-20131008&from=EN, 2011.

[90] C. Diaz, E. Kosta, H. Dekeyser, M. Kohlweiss, and G. Nigusse, "Privacy Preserving Electronic Petitions," *Identity in the Information Society*, vol. 1, no. 1, pp. 203–219, 2008.

[91] 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.

[92] A. Rial and G. Danezis, "Privacy-Preserving Smart Metering," in *Proceedings of the 10th annual ACM Workshop on Privacy in the Electronic Society (WPES 2011)*, pp. 49–60, ACM, 2011.

[93] Troncoso, C. and Danezis, G. and Kosta, E. and Balasch, J. and Preneel, B., "PriPAYD: Privacy-Friendly Pay-As-You-Drive Insurance," *IEEE Transactions on Dependable and Secure Computing*, vol. 8, no. 5, 2011.

[94] J. Balasch, A. Rial, C. Troncoso, B. Preneel, I. Verbauwhede, and C. Geuens, "PrETP: Privacy-Preserving Electronic Toll Pricing," in *Proceedings of the 19th USENIX Security Symposium*, pp. 63–78, 2010.

[95] The European Union: Official Journal of the European Communities, "Directive 2004/52/EC of the European Parliament and of the Council of 29 April 2004 on the interoperability of electronic road toll systems in the Community." http://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32004L0052R(01)&from=EN, 2004.

[96] 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.

[97] The European Union: Official Journal of the European Communities, "Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector." http://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32002L0058&from=EN, 2002.

[98] M. Alshammari and A. C. Simpson, "Personal Data Management: An Abstract Personal Data Lifecycle Model," in *Proceedings of the Business Process Management Workshops (BPM 2017)*, pp. 685–697, Springer, 2017.

[99] M. Alshammari and A. C. Simpson, "A UML Profile for Privacy-Aware Data Lifecycle Models," in *Proceedings of the International Workshop on Security and Privacy Requirements Engineering (SECPRE 2017)*, pp. 189–209, Springer, 2017.

[100] M. Alshammari and A. C. Simpson, "A Model-based Approach to Support Privacy Compliance.," *Information & Computer Security*, vol. 26, no. 4, pp. 437–453, 2018.

[101] K. Möller, "Lifecycle Models of Data-centric Systems and Domains: The Abstract Data Lifecycle Model," *Semantic Web*, vol. 4, no. 1, pp. 67–88, 2013.

[102] Object Management Group, "OMG Unified Modeling Language (OMG UML)." `http://www.omg.org/spec/UML/`, 2015.

[103] M. Alshammari and A. C. Simpson, "Towards an Effective Privacy Impact and Risk Assessment Methodology: Risk Analysis," in *Proceedings of the 13th DPM International Workshop on Data Privacy Management (DPM 2018)*, pp. 209–224, Springer, 2018.

[104] M. Alshammari and A. C. Simpson, "Towards an Effective Privacy Impact and Risk Assessment Methodology: Risk Assessment," in *Proceedings of the 15th International Conference on Trust, Privacy and Security in Digital Business (TrustBus 2018)*, pp. 85–99, Springer, 2018.

[105] P. Dourish and K. A., "Collective Information Practice: Exploring Privacy and Security as Social and Cultural Phenomena.," *Human–Computer Interaction*, vol. 21, no. 3, pp. 319–342, 2006.

[106] T. Kim, S. Lee, and E. Lee, "Privacy engineering in ubiComp," *Computational Science And Its Applications*, vol. 3482, pp. 1279–1288, 2005.

[107] M. Alshammari and A. C. Simpson, "Privacy Architectural Strategies: An Approach for Achieving Various Levels of Privacy Protection," in *Proceedings of the 2018 Workshop on Privacy in the Electronic Society (WPES'18)*, pp. 143–154, ACM, 2018.

[108] A. Cavoukian, M. Monica, A. Fariba, R. Dan, and K. Jeff, "Privacy Risk Management: Building Privacy Protection into a Risk Management Framework to Ensure that Privacy Risks are Managed, by Default." `https://www.ipc.on.ca/images/Resources/pbd-priv-risk-mgmt.pdf`, 2010.

[109] International Organization for Standardization, "ISO 31000 - Risk Management – Principles and Guidelines." `http://www.iso.org/iso/home/standards/iso31000.htm`, 2009.

Appendices

A

The APDL Model

A.1 Assessment Criteria

In this chapter, we establish a set of assessment criteria (checkpoints) for the principal activities of the APDL model stages. In addition, we describe these activities along with associated criteria via activity diagrams.

A.1.1 The Collection Stage

Based on the relevant GPS principles, we establish four essential assessment criteria in conformity with the identified processing plan: personal data items need to be adequate, relevant and not excessive in relation to the specified purposes; the specified personal data has been collected by lawful and fair methods; privacy notice has been communicated to data subjects at or before the time of collection; and the consent of data subjects have been obtained. If these criteria are satisfied, data values can be created or recorded; otherwise, corrective actions can be carried out. Depending on the unsatisfied criterion, the process will continue, either by specifying the minimum amount required of data items, specifying lawful and fair collection methods, communicating the privacy notice to data subjects at the collection time, or by obtaining their explicit or implicit consent. Figure A.1 shows

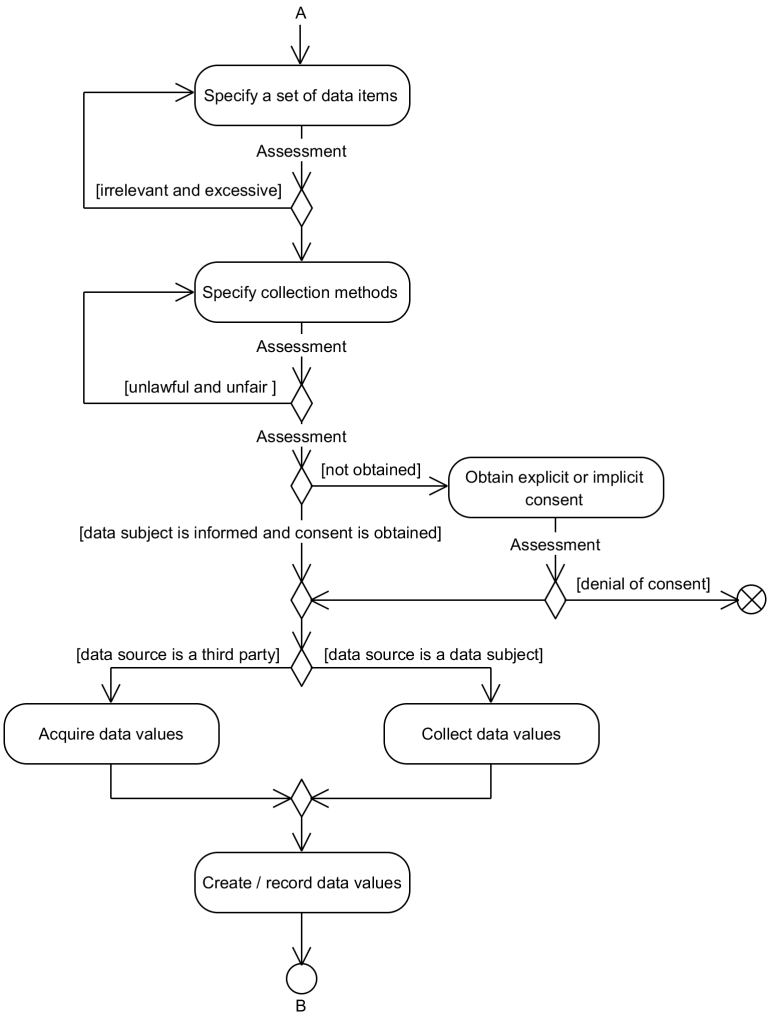

Figure A.1: The principal activities of the Collection stage

an activity diagram that represents the principal activities of the Collection stage, along with the essential assessment criteria.

A.1.2 The Retention Stage

Based on the relevant GPS principles, we establish two essential assessment criteria: personal data will not be kept beyond the specified retention time and the storage methods conform with those identified and reviewed in the processing plan. If these criteria are satisfied, data values can be organised, structured or stored, and archival and backup copies can be made; otherwise, corrective actions can be carried out. Depending on the unsatisfied criterion, the process will continue, either by specifying a retention time according to the retention policy or by specifying

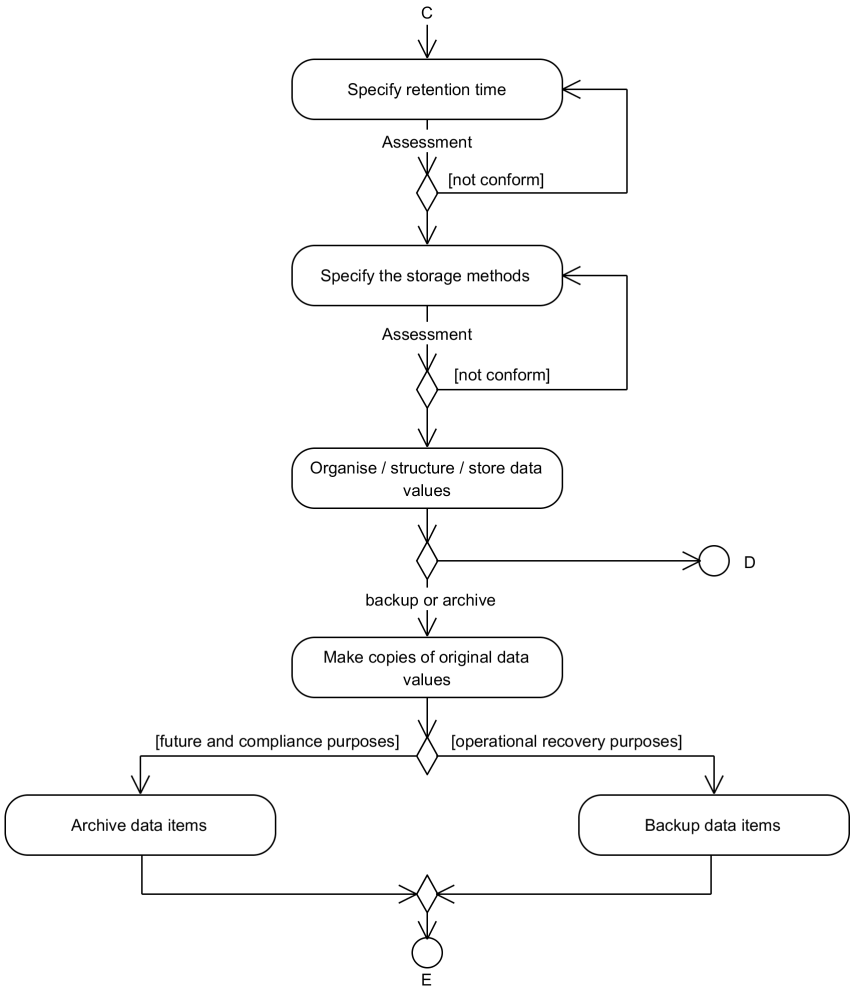

Figure A.2: The principal activities of the Retention stage

storage methods according to the processing plan. Figure A.2 shows an activity diagram that represents the principal activities of the Retention stage, along with the essential assessment criteria.

A.1.3 The Access Stage

Based on the relevant GPS principles, we establish three essential assessment criteria: personal data will not be accessed and retrieved when it is retained beyond the specified retention time for unjustified and unlawful reasons; the specified set of personal data items are limited to the specified purpose for the Participation, Usage and Disclosure stages; and the retrieval mechanisms conform with those identified and reviewed in the processing plan. If these criteria are satisfied, data values

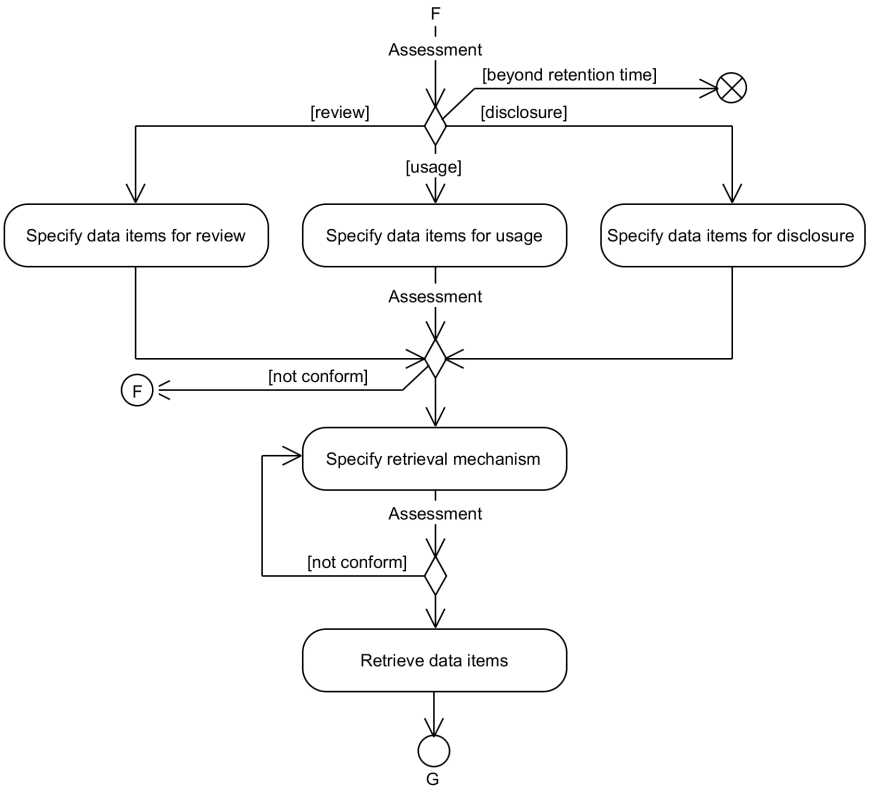

Figure A.3: The principal activities of the Access stage

can be specified, and retrieved or consulted; otherwise, corrective actions can be carried out. Depending on the unsatisfied criterion, the process will continue, either by specifying relevant, adequate and not excessive data items according to the following stage or by specifying retrieval mechanisms according to the processing plan. Figure A.3 shows an activity diagram that represents the principal activities of the Access stage, along with the essential assessment criteria.

A.1.4 The Participation Stage

Based on the relevant GPS principles, we establish three essential assessment criteria: the means by which data subjects can review their personal data are described in a clear and concise manner; personal data values are displayed in an understandable format; and consent and preferences changes are confirmed and data subjects are informed about the consequences of these changes. If these criteria are satisfied, data values can be reviewed and altered; otherwise, corrective actions can be carried

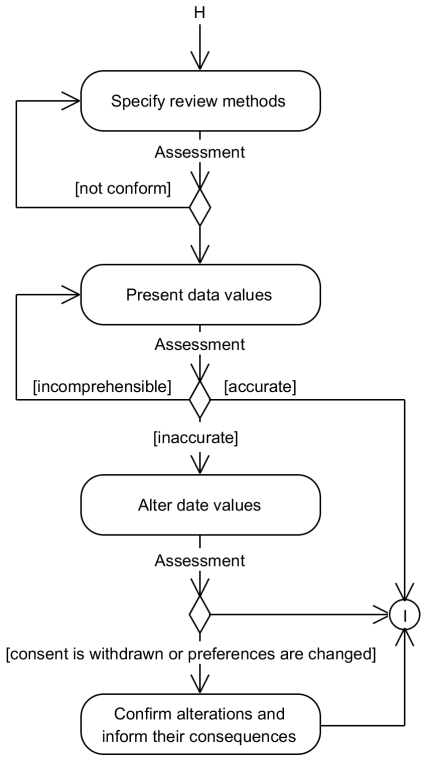

Figure A.4: The principal activities of the Participation stage

out. Depending on the unsatisfied criterion, the process will continue, either by specifying review methods according to the processing plan, by presenting personal data values in a comprehensible manner, or by confirming and informing data subjects about alterations and their corresponding consequences. Figure A.4 shows an activity diagram that represents the principal activities of the Participation stage, along with the essential assessment criteria.

A.1.5 The Usage Stage

Based on the relevant GPS principles, we establish an essential assessment criterion: the retrieved personal data values are used only for the specified purposes for which data subjects have provided their explicit or implicit consent, unless required by applicable law or regulations. This implies deriving personal data values by mining or combining several data values from internal or external sources. If these criteria are satisfied, data values can be classified, manipulated, synthesised or analysed; otherwise, corrective actions can be carried out. The process will continue by

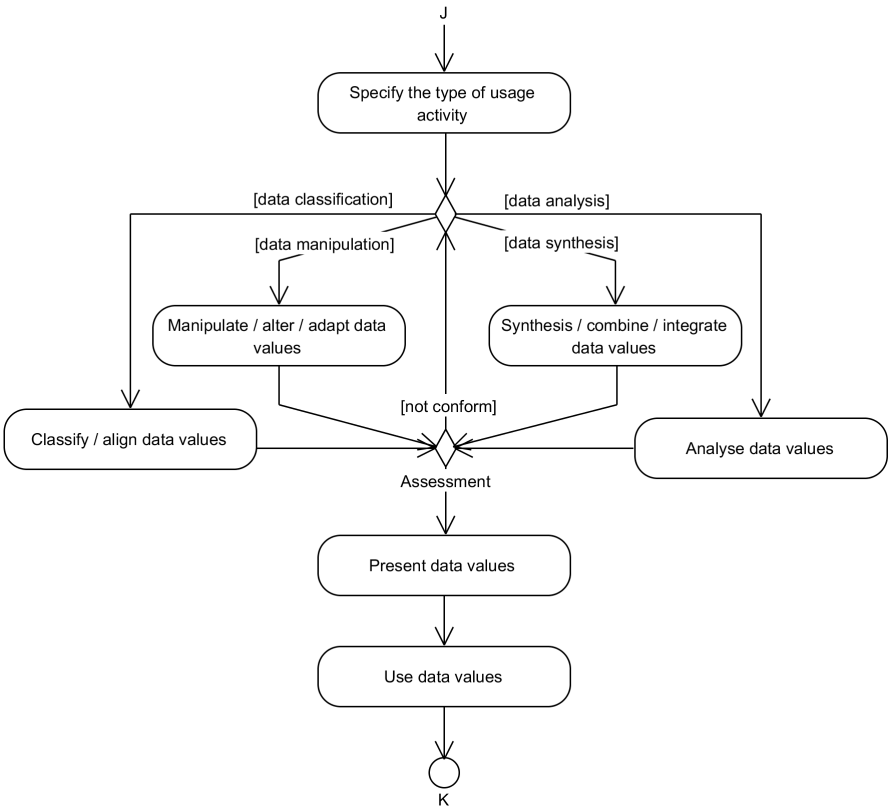

Figure A.5: The principal activities of the Usage stage

ensuring that personal data values conform with the specified purpose, in agreement with the consent obtained from data subjects and in compliance with applicable law and regulations. Figure A.5 shows an activity diagram that represents the principal activities of the Usage stage, along with the essential assessment criterion.

A.1.6 The Disclosure Stage

Based on the relevant GPS principles, we establish two essential assessment criteria: the retrieved personal data values are disclosed only for the specified purposes for which data subjects have provided their explicit or implicit consent, unless required by applicable law or regulations; and the means by which personal data is disclosed are limited to those identified and reviewed in the processing plan. If these criteria are satisfied, data values can be disseminated, made available or transmitted; otherwise, corrective actions can be carried out. Depending on the unsatisfied criterion, the process will continue, either by ensuring that personal

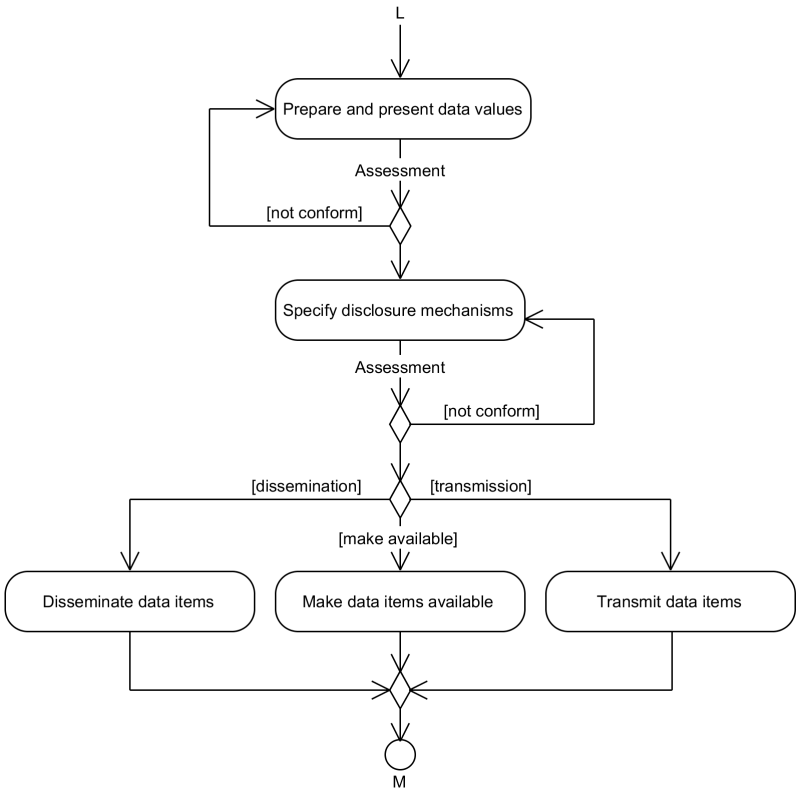

Figure A.6: The principal activities of the Disclosure stage

data values are disclosed in conformity with the specified purposes, in agreement with the consent obtained from data subjects and in compliance with applicable law and regulations, or by specifying disclosure mechanisms according to the processing plan. Figure A.6 shows an activity diagram that represents the principal activities of the Disclosure stage, along with the essential assessment criteria.

A.1.7 The Destruction Stage

Based on the relevant GPS principles, we establish two essential assessment criteria: personal data values are not required by applicable law or regulations; and the means by which personal data is destroyed are limited to those identified and reviewed in the processing plan. If these criteria are satisfied, data values can be erased, redacted or destroyed; otherwise, corrective actions can be carried out. Depending on the unsatisfied criterion, the process will continue, either by ensuring that personal data values are destroyed in conformity with the specified purpose and in compliance with

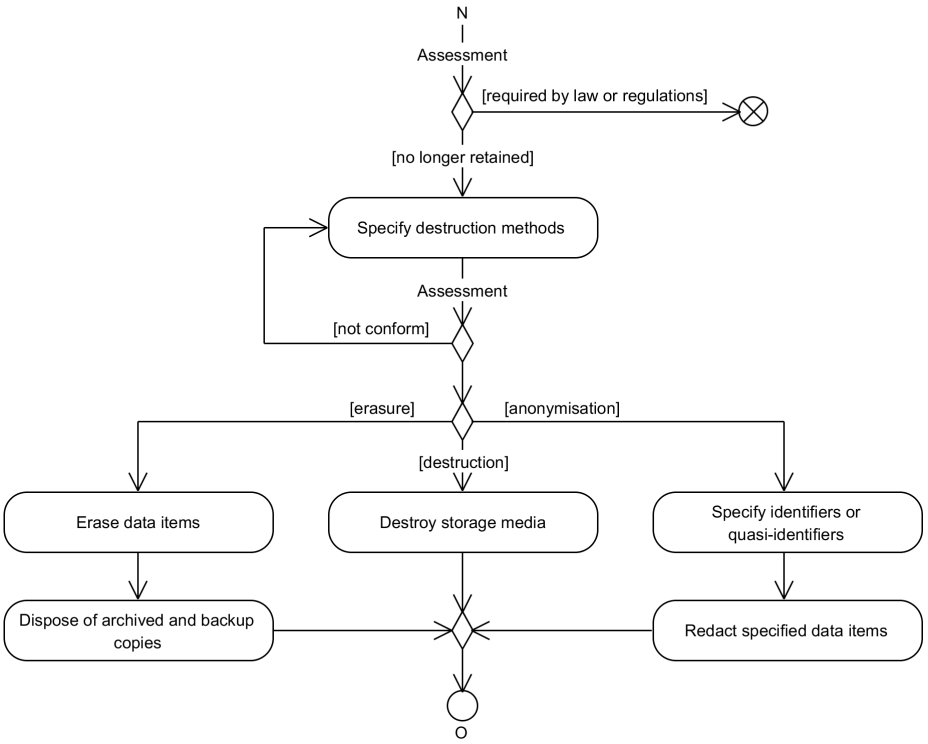

Figure A.7: The principal activities of the Destruction stage

applicable law and regulations, or by specifying destruction methods according to the processing plan. Figure A.7 shows an activity diagram that represents the principal activities of the Destruction stage, along with the essential assessment criteria.

B
The Case Study: EETS

B.1 Data-Processing Representation

B.1.1 Refinement

Figure B.1 – Figure B.6 show the operationalisation of *Activity3, Activity4, Activity5, Activity6, AbstractActivity2* and *Activity9*, respectively, in terms of actions, their corresponding events, roles and involved actors.

Activity3 coordinates its execution via one action and one corresponding event. CollectUsageData is represented as *Action3*. Detect is represented as *Event3* to specify the occurrence of circulating a vehicle on a particular toll domain to collect location data for toll declaration and calculation.

Activity4 coordinates its execution via two actions and two corresponding events. CollectVehicleData is represented as *Action4*. Identify is represented as *Event4* to specify the occurrence of identifying a vehicle circulating on a particular toll domain to collect vehicle data for enforcement management. CollectRealTimeLocationData is represented as *Action5*. Collect is represented as *Event5* to specify the occurrence of identifying a vehicle circulating on a particular toll domain to collect location data for enforcement management.

Activity5 coordinates its execution via one action and one corresponding event. CalculateToll is represented as *Action6*. Calculate is represented as *Event6* to

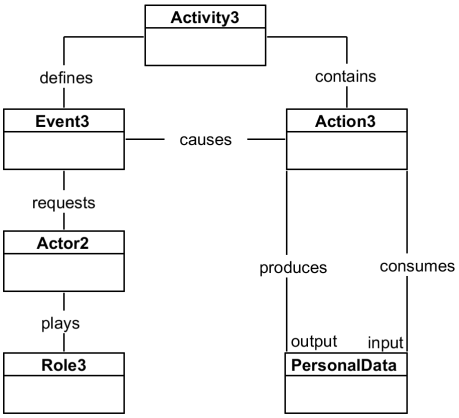

Activity3:	CollectingRoadUsageData	*Action3:*	CollectUsageData
Event3:	Detect	*Actor2:*	EETSProvider,
Role3:	CollectionAgent	*PersonalData:*	EETSUser, Vehicle, LocationData

Figure B.1: The operationalisation of Activity3.

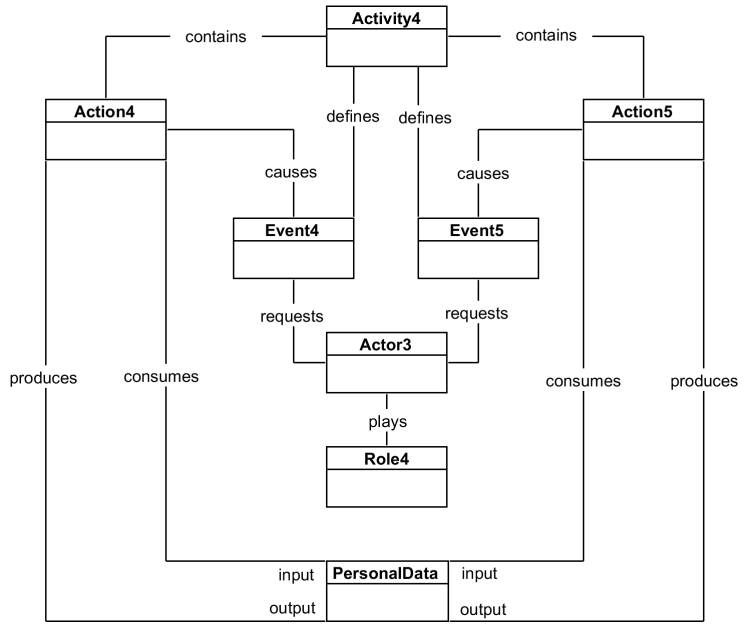

Activity4:	CollectingRealTimeData	*Action4:*	CollectVehicleData
Event4:	Identify	*Action5:*	CollectRealTimeLocationData,
Event5:	Collect	*Actor3:*	TollCharger
Role4:	EnforcementAuthority	*PersonalData:*	EETSUser, Vehicle, LocationData

Figure B.2: The operationalisation of Activity4.

specify the occurrence of using toll declaration data to calculate road-usage toll at the end of the tax period.

Activity6 coordinates its execution via one action and one corresponding event. CommunicateToll is represented as *Action7*. Communicate is represented as *Event7* to specify the occurrence of sending a claim to a EETS user by means of an

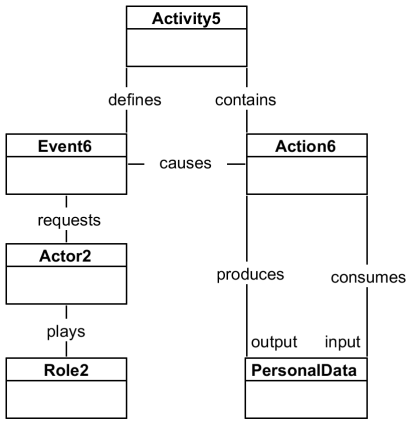

Activity5:	CalculatingRoadUsageToll	*Action6:*	CalculateToll
Event6:	Calculate	*Actor2:*	EETSProvider,
Role2:	ContractAgent	*PersonalData:*	EETSUser, Vehicle, LocationData

Figure B.3: The operationalisation of Activity5.

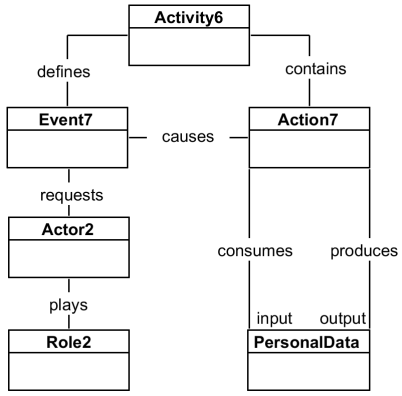

Activity6:	CommunicatingRoadUsageToll	*Action7:*	CommunicateToll
Event7:	Communicate	*Actor2:*	EETSProvider,
Role2:	ContractAgent	*PersonalData:*	EETSUser, Vehicle, LocationData

Figure B.4: The operationalisation of Activity6.

invoice at the end of the tax period.

Activity7 coordinates its execution via one action and one corresponding event. RetrieveEETSContract is represented as *Action8*. Identify is represented as *Event8* to specify the occurrence of identifying a EETS provider's contract using fixed or mobile RSE equipment.

Activity8 coordinates its execution via two actions and two corresponding events. RequestPaymentGuarantee is represented as *Action9*. RequestGuarantee is represented as *Event9* to specify the occurrence of requesting payment guarantee for an inferred object. RequestParametersForBilling is represented as *Action10*.

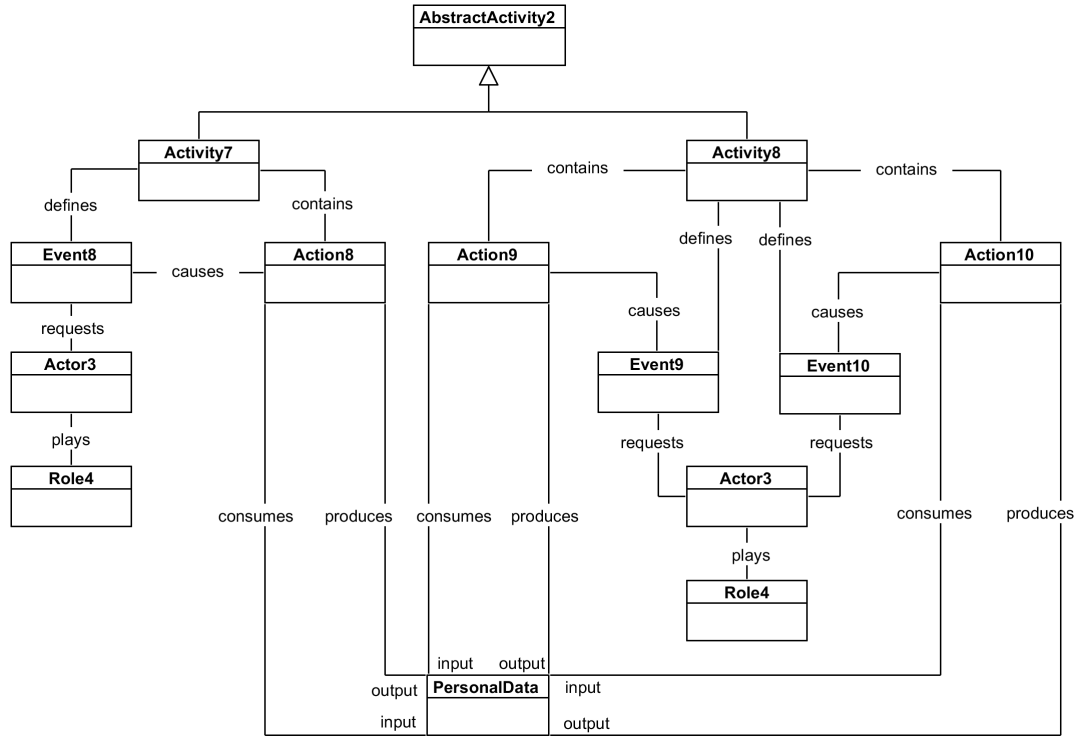

AbstractActivity2:	HandlingException	*Activity7:*	IdentifyingEETSContract
Activity8:	RequestingAdditionalData	*Action8:*	RetrieveEETSContract
Event8:	Identify	*Action9:*	RequestPaymentGuarantee
Event9:	RequestGuarantee	*Action10:*	RequestParametersForBilling
Event10:	RequestData	*Actor3:*	TollCharger
Role4:	EnforcementAuthority	*PersonalData:*	EETSUser, Vehicle, LocationData

Figure B.5: The operationalisation of AbstractActivity2.

RequestData is represented as *Event10* to specify the occurrence of requesting additional parameters for billing details.

Activity9 coordinates its execution via one action and one corresponding event. ReportComplianceCheckCommunication is represented as *Action11*. Report is represented as *Event11* to specify the occurrence of reporting compliance check communication events to EETS providers.

B.1.2 Representation

Figure B.7 – Figure B.10 show the representation of the CollectingRoadUsageData, CalculatingUsageToll, HandlingException and ReportingTollEvent respectively as data-processing activities.

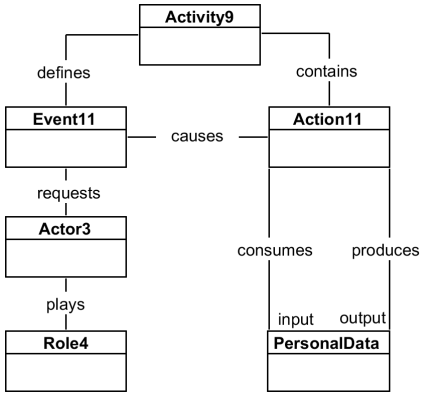

Activity9:	ReportingComplianceEvent	*Action11:*	ReportComplianceCheckCommunication
Event11:	Action11	*Actor3:*	TollCharger,
Role4:	EnforcementAuthority	*PersonalData:*	EETSUser, Vehicle, LocationData

Figure B.6: The operationalisation of Activity9.

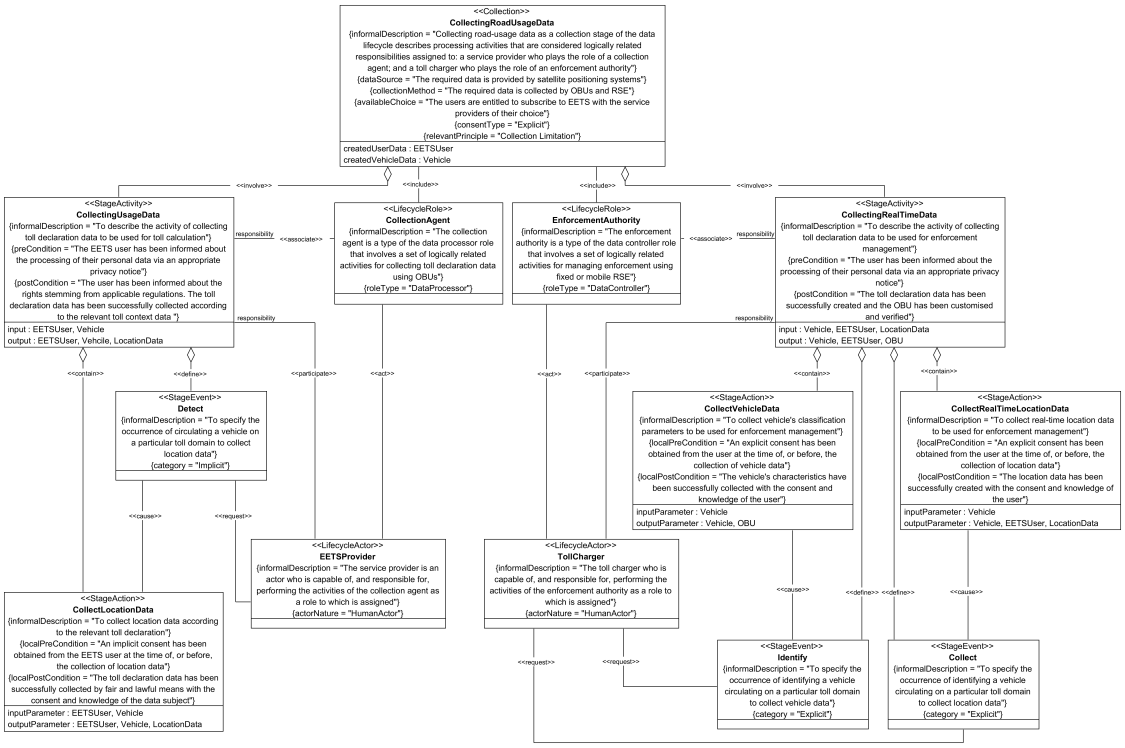

Figure B.7: The representation of CollectingRoadUsageData as a Collection stage.

B.2 Data-Centric Threat Modelling

B.2.1 Vulnerability Analysis

Context-Relative Processing Norms

Collection Stage

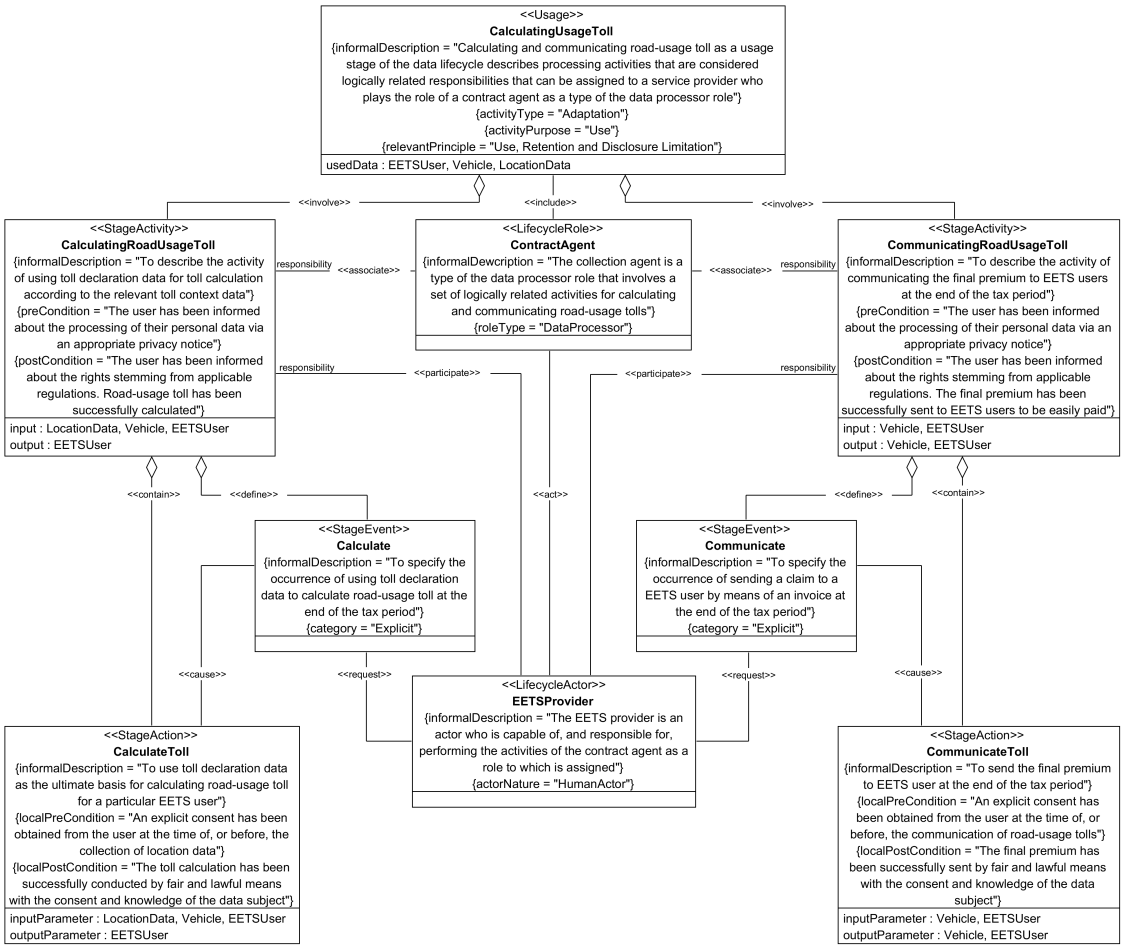

Figure B.8: The representation of CalculatingUsageToll as a Usage stage.

- *The context-relative processing norm (PN.3)* for DPA.3 is as follows.

  In the context of EETS, the collection of personal data of a certain type (LocationData: time, distance, place) about EETS users (acting as data subjects) by EETS providers (acting as data processors on behalf of toll chargers) is governed by processing principles derived from the EU's GDPR, Directive 2004/52/EC and the related Commission Decision 2009/750/EC.

- *The context-relative processing norm (PN.4)* for DPA.4 is as follows.

  In the context of EETS, the collection of personal data of a certain type (LocationData: time, distance, place) about EETS users (acting as data subjects) by toll chargers (acting as data controllers) is governed by processing principles derived from the EU's GDPR, Directive 2004/52/EC and the related Commission Decision 2009/750/EC.

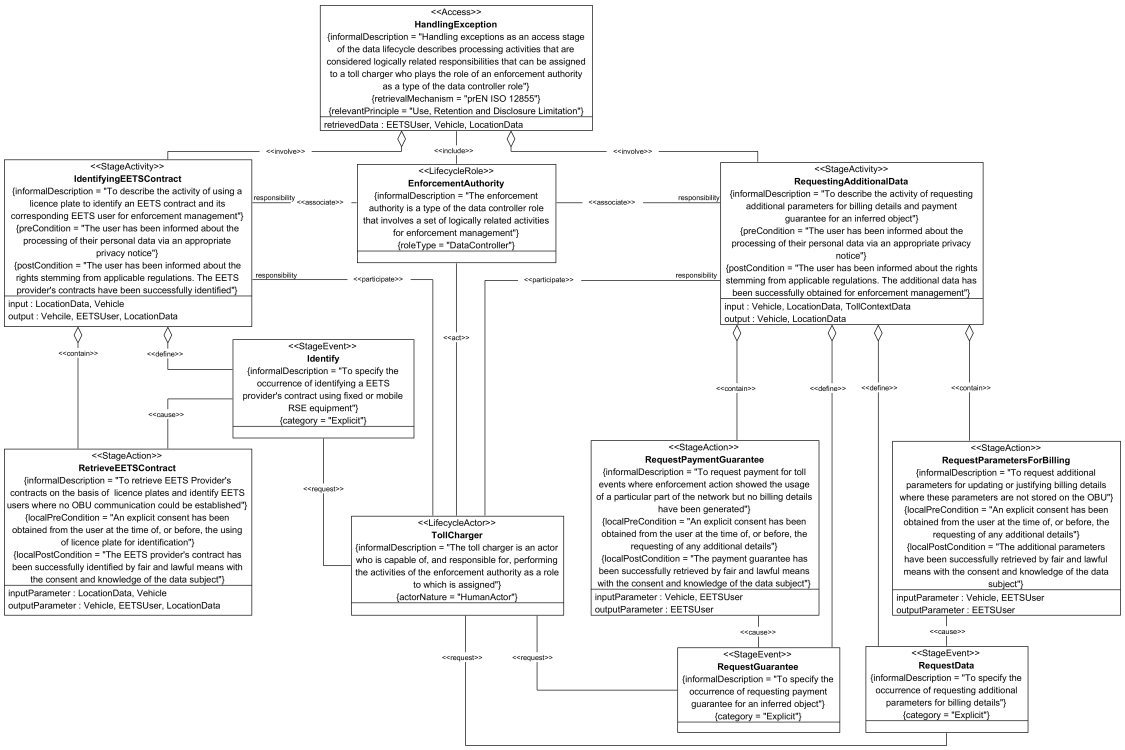

Figure B.9: The representation of HandlingException as an Access stage.

Retention Stage

- *The context-relative processing norm (PN.5)* for DPA.5 is as follows.

  In the context of EETS, the retention of personal data of a certain type (EETSUser: user ID, name, billing address) about EETS users (acting as data subjects) by EETS providers (acting as data processors on behalf of toll chargers) is governed by processing principles derived from the EU's GDPR, Directive 2004/52/EC and the related Commission Decision 2009/750/EC.

- *The context-relative processing norm (PN.6)* for DPA.6 is as follows.

  In the context of EETS, the retention of personal data of a certain type (Vehicle: licence plate, classification code) about EETS users (acting as data subjects) by EETS providers (acting as data processors on behalf of toll chargers) is governed by processing principles derived from the EU's GDPR, Directive 2004/52/EC and the related Commission Decision 2009/750/EC.

- *The context-relative processing norm (PN.7)* for DPA.7 is as follows.

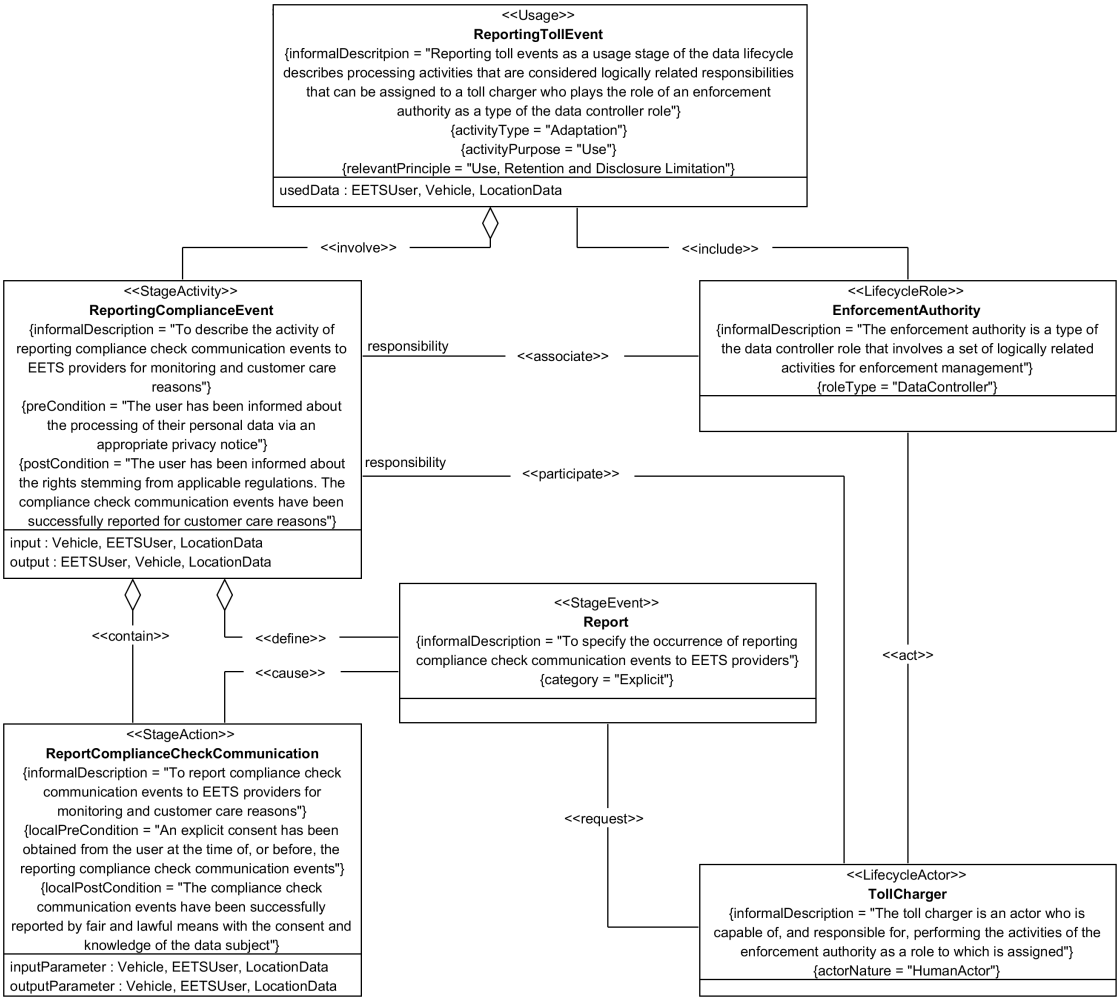

Figure B.10: The representation of ReportingTollEvent as a Usage stage.

In the context of EETS, the retention of personal data of a certain type (LocationData: time, distance, place) about EETS users (acting as data subjects) by EETS providers (acting as data processors on behalf of toll chargers) is governed by processing principles derived from the EU's GDPR, Directive 2004/52/EC and the related Commission Decision 2009/750/EC.

- *The context-relative processing norm (PN.8)* for DPA.8 is as follows.

  In the context of EETS, the retention of personal data of a certain type (LocationData: time, distance, place) about EETS users (acting as data subjects) by toll chargers (acting as data controllers) is governed by processing principles derived from the EU's GDPR, Directive 2004/52/EC and the related Commission Decision 2009/750/EC.

Access Stage

- *The context-relative processing norm (PN.9)* for DPA.9 is as follows.

  In the context of EETS, the retrieval of personal data of a certain type (EETSUser: usageToll) about EETS users (acting as data subjects) by a toll charger (acting as a data controller) is governed by processing principles derived from the EU's GDPR, Directive 2004/52/EC and the related Commission Decision 2009/750/EC.

- *The context-relative processing norm (PN.10)* for DPA.10 is as follows.

  In the context of EETS, the retrieval of personal data of a certain type (EETSUser: user ID, name, billing address) about EETS users (acting as data subjects) by a toll charger (acting as a data controller) is governed by processing principles derived from the EU's GDPR, Directive 2004/52/EC and the related Commission Decision 2009/750/EC.

Usage Stage

- *The context-relative processing norm (PN.11)* for DPA.11 is as follows.

  In the context of EETS, the use of personal data of a certain type (LocationData: time, distance, place; EETSUser: user ID, name, billing address) about EETS users (acting as data subjects) by an EETS provider (acting as a data processor on behalf of a toll charger) is governed by processing principles derived from the EU's GDPR, Directive 2004/52/EC and the related Commission Decision 2009/750/EC.

- *The context-relative processing norm (PN.12)* for DPA.12 is as follows.

  In the context of EETS, the use of personal data of a certain type (EETSUser: user ID, name, billing address) about EETS users (acting as data subjects) by an EETS provider (acting as a data processor on behalf of a toll charger) is governed by processing principles derived from the EU's GDPR, Directive 2004/52/EC and the related Commission Decision 2009/750/EC.

- *The context-relative processing norm (PN.13)* for DPA.13 is as follows.

> In the context of EETS, the use of personal data of a certain type (LocationData: time, distance, place; EETSUser: user ID, name, billing address) about EETS users (acting as data subjects) by a toll charger (acting as a data controller) is governed by processing principles derived from the EU's GDPR, Directive 2004/52/EC and the related Commission Decision 2009/750/EC.

Privacy Vulnerabilities

In this section, we illustrate how to derive privacy vulnerabilities from context-relative processing norms (PN.3 – PN.13). We refer to Table 7.1 that illustrates the main processing operations and their associated data-processing activites.

- The context-relative processing norms (PN.3 and PN.4) for PO.2
  - ***Element 1:*** *data-processing activities.* PN.3 and PN.4 are established with reference to DPA.3 and DPA.4, which are specified as a result of the purpose refinement of 7.4.1. Thus, DPA.3 and DPA.4 explicitly participate in the fulfilment of the abstract purpose.
  - ***Element 2:*** *attributes.* Time, distance and place are necessary for the fulfilment of the concrete purposes of DPA.3 and DPA.4 (toll declaration and enforcement management) respectively. This means that they are necessary to accomplish the execution of associated actions: CollectLocationData and CollectRealTimeData respectively. In relation to the specification of DPA.3, location data is collected over time and the collection is not restricted to relevant national and international toll domains. DPA.3 is not explicitly specified in a manner that prevents the collection of location data outside the boundaries of toll domains. As such, 'an improper activity specification' (PV.1) is derived as a privacy vulnerability (a weakness in the specification of DPA.3).
  - ***Element 3:*** *actors.* EETS providers and toll chargers are the actors who are involved in the performance of DPA.3 and DPA.4. EETS providers are responsible for performing the activities of the collection agent as a role to which they are assigned. Toll chargers are responsible for

performing the activities of the enforcement authority as a role to which they are assigned. Both types of roles are defined in PO.2. We assume that a role-based access control model is maintained in a proper way.

- ***Element 4:*** *processing principles.* The relevant processing principles stated in the EU's GDPR are: "[...] personal data must be":
  - "processed fairly and lawfully";
  - "collected for specified, explicit and legitimate purposes and not further processed in a way incompatible with those purposes"; and
  - "adequate, relevant and not excessive in relation to the purposes for which they are collected and/or further processed".

  The relevant processing principles stated in the GPS are: Collection Limitation and Data Minimisation. In relation to the specification of DPA.3 and DPA.4, these processing principles are specified pre- and post-conditions, as per Figure B.7. This means that the collection of these types of personal data is in compliance with the processing principles stated in relevant legal frameworks and standards.

- The context-relative processing norms (PN.5 and PN.6) for PO.3
  - ***Element 1:*** *data-processing activities.* PN.5 and PN.6 are established with reference to DPA.5 and DPA.6, which are specified as a result of the purpose refinement of 7.4.1. Thus, DPA.5 and DPA.6 explicitly participate in the fulfilment of the abstract purpose.
  - ***Element 2:*** *attributes.* 'Identification and contact data' and 'vehicle classification parameters' are necessary for the fulfilment of the concrete purposes of DPA.5 and DPA.6 respectively. This means that they are necessary to accomplish the execution of associated actions: StoreUserData and StoreVehicleData respectively. In reference to the data model, the required data items are modelled in a manner that directly or indirectly facilitates the integration of additional data items with 'identification

and contact data'. As such, 'an improper data model' (PV.2) is derived as a privacy vulnerability (a weakness in data modelling).

- ***Element 3:*** *actors.* EETS providers are the actors who are involved in the performance of the DPA.5 and DPA.6. They are responsible for performing the activities of the contract agent as a role to which they are assigned. This type of role is defined in PO.2.1. We assume that a role-based access control model is maintained in a proper way.
- ***Element 4:*** *processing principles.* The relevant processing principles stated in the EU's GDPR are: "[...] personal data must be":
  - "processed fairly and lawfully";
  - "collected for specified, explicit and legitimate purposes and not further processed in a way incompatible with those purposes"; and
  - "adequate, relevant and not excessive in relation to the purposes for which they are collected and/or further processed".

  The relevant processing principle stated in the GPS is: Use, Retention and Disclosure Limitation. In relation to the specification of DPA.5 and DPA.6, these processing principles are specified as pre- and post-conditions, as per Figure 7.8. This means that the retention of these types of personal data is in compliance with the processing principles stated in relevant legal frameworks and standards.

- The context-relative processing norms (PN.7 and PN.8) for PO.4
  - ***Element 1:*** *data-processing activities.* PN.7 and PN.8 are established with reference to DPA.7 and DPA.8, which are specified as a result of the purpose refinement of 7.4.1. Thus, DPA.7 and DPA.8 explicitly participate in the fulfilment of the abstract purpose.
  - ***Element 2:*** *attributes.* Time, distance and place are necessary for the fulfilment of the concrete purposes of DPA.7 and DPA.8 (toll declaration and enforcement management) respectively. This means that they are

necessary to accomplish the execution of associated actions: StoreLocationData and StoreRealTimeData respectively. In reference to the data model, the required data items are modelled in a manner that directly or indirectly facilitates the linkability of 'location data' to 'identification and contact data'. As such, 'an improper data model' (PV.2) is derived as a privacy vulnerability (a weakness in data modelling).

- ***Element 3:*** *actors.* EETS providers and toll chargers are the actors who are involved in the performance of DPA.7 and DPA.8. EETS providers are responsible for performing the activities of the collection agent as a role to which they are assigned. Toll chargers are responsible for performing the activities of the enforcement authority as a role to which they are assigned. Both types of roles are defined in PO.2.2. We assume that a role-based access control model is maintained in a proper way.
- ***Element 4:*** *processing principles.* The relevant processing principles stated in the EU's GDPR are: "[...] personal data must be":
  - "processed fairly and lawfully";
  - "collected for specified, explicit and legitimate purposes and not further processed in a way incompatible with those purposes"; and
  - "adequate, relevant and not excessive in relation to the purposes for which they are collected and/or further processed".

  The relevant processing principle stated in the GPS is: Use, Retention and Disclosure Limitation. In relation to the specification of DPA.7 and DPA.8, these processing principles are specified as pre- and post-conditions, as per Figure 7.8. This means that the retention of these types of personal data is in compliance with the processing principles stated in relevant legal frameworks and standards.

- The context-relative processing norm (PN.9 and PN.10) for PO.5

- ***Element 1:*** *data-processing activities.* PN.9 and PN.10 are established with reference to DPA.9 and DPA.10, which are specified as a result of the purpose refinement of 7.4.1. Thus, DPA.9 and DPA.10 explicitly participate in the fulfilment of the abstract purpose.
- ***Element 2:*** *attributes.* 'Identification and contact data' is necessary for the fulfilment of the concrete purposes of DPA.9 and DPA.10 (enforcement management and payment guarantee) respectively. This means that they are necessary to accomplish the execution of associated actions: RetrieveEETSContract and RequestPaymentGuarantee respectively.
- ***Element 3:*** *actors.* Toll chargers are the actors who are involved in the performance of DPA.9 and DPA.10. They are responsible for performing the activities of the enforcement authority as a role to which they are assigned. This type of role is defined in PO.5. We assume that a role-based access control model is maintained in a proper way. However, DPA.9 and DPA.10 require sufficient logging and monitoring — not least because they retrieve identification and contact data, vehicle classification parameters and location data. As such, 'a lack of logs and audit trails' (PV.3) is derived as a privacy vulnerability (a weakness in a privacy control).
- ***Element 4:*** *processing principles.* The relevant processing principles stated in the EU's GDPR are: "[...] personal data must be":
  - "processed fairly and lawfully";
  - "collected for specified, explicit and legitimate purposes and not further processed in a way incompatible with those purposes"; and
  - "adequate, relevant and not excessive in relation to the purposes for which they are collected and/or further processed".

  The relevant processing principles stated in the GPS are: Use, Retention and Disclosure Limitation. In relation to the specification of DPA.9 and DPA.10, these processing principles are specified pre- and post-conditions,

as per Figure B.9. This means that the retrieval of these types of personal data is in compliance with the processing principles stated in relevant legal frameworks and standards.

- The context-relative processing norm (PN.11 and PN.12) for for PO.6
  - ***Element 1:*** *data-processing activities.* PN.11 and PN.12 are established with reference to DPA.11 and DPA.12, which are specified as a result of the purpose refinement of 7.4.1. Thus, DPA.11 and DPA.12 explicitly participate in the fulfilment of the specified purpose.
  - ***Element 2:*** *attributes.* 'Identification and contact data', 'location data' and 'vehicle classification parameters' are necessary for the fulfilment of the concrete purposes of DPA.11 and DPA.12 (toll calculation and toll communication) respectively. This means that they are necessary to accomplish the execution of associated actions: CalculateToll and CommunicateToll respectively. In reference to the data model, the required data items are modelled in a manner that directly or indirectly facilitates the linkability of 'vehicle classification parameters' to 'identification and contact data' and 'location data'. As such, 'an improper data model' (PV.2) is derived as a privacy vulnerability (a weakness in data modelling).
  - ***Element 3:*** *actors.* EETS providers are the actors who are involved in the performance of DPA.11 and DPA.12. EETS providers are responsible for performing the activities of the contract agent as a role to which they are assigned. This type of role is defined in PO.4. We assume that a role-based access control model is maintained in a proper way. However, DPA.11 and DPA.12 require sufficient logging and monitoring — not least because they use identification and contact data, vehicle classification parameters and location data. As such, 'a lack of logs and audit trails' (PV.3) is derived as a privacy vulnerability (a weakness in a privacy control).

- ***Element 4:*** *processing principles.* The relevant processing principles stated in the EU's GDPR are: "[...] personal data must be":
  - "processed fairly and lawfully";
  - "collected for specified, explicit and legitimate purposes and not further processed in a way incompatible with those purposes"; and
  - "adequate, relevant and not excessive in relation to the purposes for which they are collected and/or further processed".

  The relevant processing principles stated in the GPS are: Use, Retention and Disclosure Limitation. In relation to the specification of DPA.11 and DPA.12, these processing principles are specified pre- and post-conditions, as per Figure B.8. This means that the usage of these types of personal data is in compliance with the processing principles stated in relevant legal frameworks and standards.

- The context-relative processing norm (PN.13) for for PO.7
  - ***Element 1:*** *data-processing activities.* PN.13 is established in relation to DPA.13, which is specified as a result of the purpose refinement of 7.4.1. Thus, DPA.13 explicitly participates in the fulfilment of the specified purpose.
  - ***Element 2:*** *attributes.* 'Identification and contact data', 'location data' and 'vehicle classification parameters' are necessary for the fulfilment of the concrete purpose of DPA.13 (monitoring and customer care reasons). This means that they are necessary to accomplish the execution of associated action: ReportComplianceCheckCommunication. In reference to the data model, the required data items are modelled in a manner that directly or indirectly facilitates the linkability of 'vehicle classification parameters' to 'identification and contact data' and 'location data'. As such, 'an improper data model' (PV.2) is derived as a privacy vulnerability (a weakness in data modelling).

- ***Element 3:*** *actors.* Toll chargers are the actors who are involved in the performance of DPA.13. They are responsible for performing the activities of the contract agent as a role to which they are assigned. This type of role is defined in PO.5. We assume that a role-based access control model is maintained in a proper way. However, DPA.13 requires sufficient logging and monitoring — not least because it uses identification and contact data, vehicle classification parameters and location data. As such, 'a lack of logs and audit trails' (PV.3) is derived as a privacy vulnerability (a weakness in a privacy control).
- ***Element 4:*** *processing principles.* The relevant processing principles stated in the EU's GDPR are: "[...] personal data must be":
  - "processed fairly and lawfully";
  - "collected for specified, explicit and legitimate purposes and not further processed in a way incompatible with those purposes"; and
  - "adequate, relevant and not excessive in relation to the purposes for which they are collected and/or further processed".

  The relevant processing principles stated in the GPS are: Use, Retention and Disclosure Limitation. In relation to the specification of DPA.13, these processing principles are specified pre- and post-conditions, as per Figure B.10. This means that the usage of these types of personal data is in compliance with the processing principles stated in relevant legal frameworks and standards.

In the context of EETS, there are other privacy vulnerabilities that can be identified by analysing additional processing norms. For example, 'an improper purpose specification' (PV.4) can be derived as a privacy vulnerability (a weakness in the specification of processing operations) from any data-processing activity that is not specified in relation to the purpose refinement of 7.4.1. In addition, when EETS users' personal data is anonymised, it is lawfully to be shared with third parties for further processing (historical, statistical or scientific purposes).

However, this requires sufficient anonymisation techniques. As such, 'a weak anonymisation technique' (PV.5) can be derived as a privacy vulnerability (a weakness in privacy controls).

B.2.2 Privacy Violations

In this section, we illustrate how to identify privacy violations in the Retention, Access and Usage stages of the APDL model.

- *In the Retention stage,* TE.3 and TE.4 are identified as threat events.

  By the occurrence of TE.3, TS.1 may passively (without the knowledge and consent of EETS users) acquire and integrate multiple data items about EETS users from various sources.

  By the occurrence of TE.4, TS.1 may passively (without the knowledge and consent of EETS users) retain EETS users' driving profiles for longer than necessary to fulfil the specified purposes or for a period specifically required by legal frameworks.

- *In the Access stage,* TE.5 is identified as a threat event. By the occurrence of TE.5, TS.2 may unwarrantably access driving profiles for specific EETS users without their knowledge and consent.

- *In the Usage stage,* TE.6 and TE.7 are identified as threat events.

  By the occurrence of TE.6, TS.1 or TS.2 may derive driving patterns for EETS users without any further adverse actions and without their knowledge and consent.

  By the occurrence of TE.7, TS.1 or TS.2 may unjustifiably classify, alter and align EETS users' driving profiles in a way facilitate future applications without their knowledge and consent.

B.2.3 Privacy harms

In this section, we analyse the identified adverse effects as motives for adverse actions. By analysing AE.1 – AE.9, together with the relevant threat scenarios, we can derive additional privacy harms as follows.

- *Increased health insurance premium (PH.2).* PH.2 occurs when a EETS provider (TS.1) or a toll charger (TS.2) may make 'excessive data inference to derive driving patterns' (TE.6) for EETS users and 'shares these driving patterns with health insurance providers' (TE.8). An insurance provider (TS.6) may make 'excessive data inference to re-identify its current and potential customers' (TE.9) by linking the derived data to particular drivers with the aim of calculating health insurance premium based on health conditions. An 'improper data model' (PV.2) and a 'lack of logs and audit trails' (PV.3) that may be exploited by TS.1 or TS.2 lead to the occurrence of TE.6. A 'lack of logs and audit trails' (PV.3) and 'improper purpose specification' (PV.4) that may be exploited by TS.1 or TS.2 lead to the occurrence of TE.8. A 'weak anonymisation technique' (PV.5) that is exploited by TS.6 leads to the occurrence of TE.9.

- *Denial of a job (PH.3).* PH.3 occurs when a EETS provider (TS.1) or a toll charger (TS.2) may make 'excessive data inference to derive driving patterns' (TE.6) for EETS users and 'shares these driving patterns with employment agencies' (TE.8). An employment agency (TS.5) may make 'excessive data inference to re-identify its job applicants' (TE.9) by linking the derived data to particular drivers with the aim of filtering those job candidates according to their health conditions or religious beliefs. An 'improper data model' (PV.2) and a 'lack of logs and audit trails' (PV.3) that may be exploited by TS.1 or TS.2 lead to the occurrence of TE.6. A 'lack of logs and audit trails' (PV.3) and 'improper purpose specification' (PV.4) that may be exploited by TS.1 or TS.2 lead to the occurrence of TE.8. A 'weak anonymisation technique' (PV.5) that is exploited by TS.5 leads to the occurrence of TE.9.

- *Being under close surveillance (PH.4).* PH.4 occurs when a EETS provider (TS.1) or a toll charger (TS.2) may make 'excessive data inference to derive driving patterns' (TE.6) for EETS users and 'shares these driving patterns with intelligence and security services' (TE.8). An intelligence and security agency (TS.4) may make 'excessive data inference to re-identify a number of EETS users' (TE.9) by linking the derived data to particular drivers with the aim of identifying and putting under surveillance (a special type of intrusion) a number of EETS users based on their associations with others or the locations frequented. An 'improper data model' (PV.2) and a 'lack of logs and audit trails' (PV.3) that may be exploited by TS.1 or TS.2 lead to the occurrence of TE.6. A 'lack of logs and audit trails' (PV.3) and 'improper purpose specification' (PV.4) that may be exploited by TS.1 or TS.2 lead to the occurrence of TE.8. A 'weak anonymisation technique' (PV.5) that is exploited by TS.4 leads to the occurrence of TE.9.

- *Receipt of targeted advertising (PH.5).* PH.5 occurs when a EETS provider (TS.1) or a toll charger (TS.2) may make 'excessive data inference to derive driving patterns' (TE.6) for EETS users and 'shares these driving patterns with advertising companies' (TE.8). An advertising agency (TS.8) may make 'excessive data inference to re-identify specific groups of individuals' (TE.9) by linking the derived data to particular drivers with the aim of sending targeted advertising emails that implicitly make reference to their religion beliefs or health conditions. An 'improper data model' (PV.2) and a 'lack of logs and audit trails' (PV.3) that may be exploited by TS.1 or TS.2 lead to the occurrence of TE.6. A 'lack of logs and audit trails' (PV.3) and 'improper purpose specification' (PV.4) that may be exploited by TS.1 or TS.2 lead to the occurrence of TE.8. A 'weak anonymisation technique' (PV.5) that is exploited by TS.8 leads to the occurrence of TE.9.

Figure B.11 – Figure B.14 show the harm tree for the privacy harms PH.2 – PH.5.

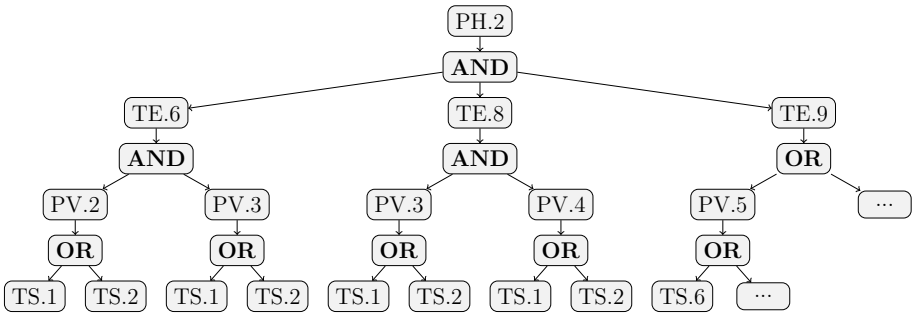

PH.2:	Increased health insurance premium	TE.6:	Excessive data inference
TE.8:	Unauthorised data disclosure	TE.9:	Excessive data inference
PV.2:	Improper data model	PV.3:	Lack of logs and audit trails
PV.4:	Improper purpose specification	PV.5:	Weak anonymisation technique
TS.1:	EETS providers	TS.2:	Toll chargers
TS.6:	Health insurance providers		

Figure B.11: The structure of the harm tree for the privacy harm PH.2.

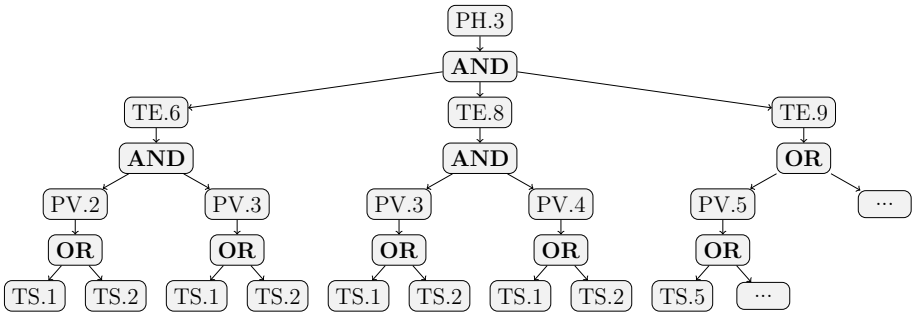

PH.3:	Denial of a job	TE.6:	Excessive data inference
TE.8:	Unauthorised data disclosure	TE.9:	Excessive data inference
PV.2:	Improper data model	PV.3:	Lack of logs and audit trails
PV.4:	Improper purpose specification	PV.5:	Weak anonymisation technique
TS.1:	EETS providers	TS.2:	Toll chargers
TS.5:	Employment agencies		

Figure B.12: The structure of the harm tree for the privacy harm PH.3.

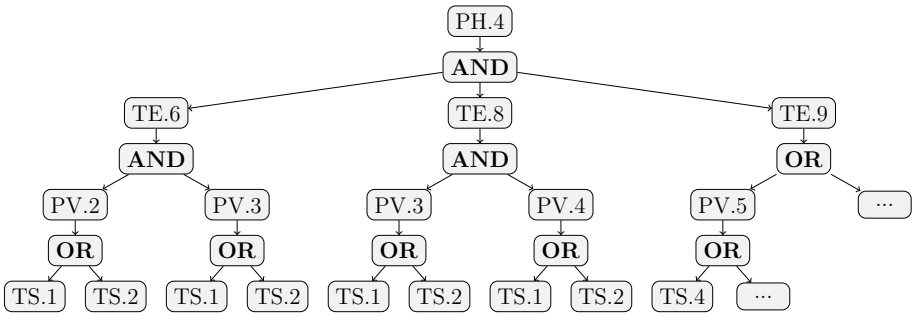

PH.4:	Being under close surveillance	TE.6:	Excessive data inference
TE.8:	Unauthorised data disclosure	TE.9:	Excessive data inference
PV.2:	Improper data model	PV.3:	Lack of logs and audit trails
PV.4:	Improper purpose specification	PV.5:	Weak anonymisation technique
TS.1:	EETS providers	TS.2:	Toll chargers
TS.4:	Intelligence and security services		

Figure B.13: The structure of the harm tree for the privacy harm PH.4.

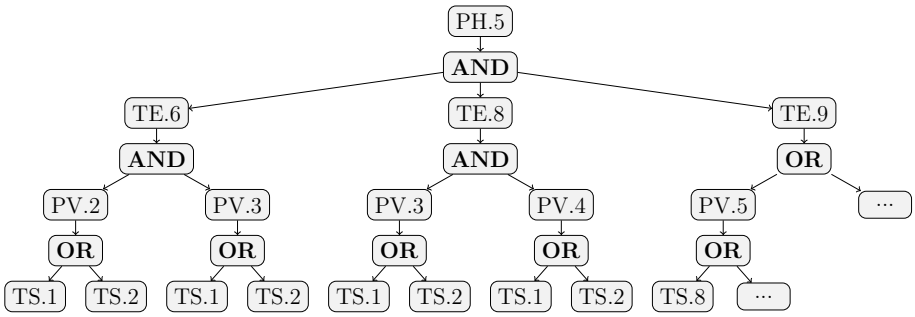

PH.5:	Receipt of targeted advertising	TE.6:	Excessive data inference
TE.8:	Unauthorised data disclosure	TE.9:	Excessive data inference
PV.2:	Improper data model	PV.3:	Lack of logs and audit trails
PV.4:	Improper purpose specification	PV.5:	Weak anonymisation technique
TS.1:	EETS providers	TS.2:	Toll chargers
TS.8:	Advertising companies		

Figure B.14: The structure of the harm tree for the privacy harm PH.5.