PRIPARE: Integrating Privacy Best Practices into a Privacy Engineering Methodology

Nicolás Notario*, Alberto Crespo*, Yod-Samuel Martín‡, Jose M. del Alamo‡, Daniel Le Métayer†, Thibaud Antignac†, Antonio Kung§, Inga Kroener**, David Wright**

*Atos, Madrid, Spain, {nicolas.notario, alberto.crespo} @atos.net
‡Universidad Politécnica de Madrid, Madrid, Spain{samuelm, jmdela} @dit.upm.es
†Inria, Lyon, France, {daniel.le-metayer, thibaud.antignac} @inria.fr
§Trialog, Paris, France, antonio.kung @trialog.com
**Trilateral, London, UK,{inga.kroener, david.wright} @trilateralresearch.com

Abstract—**Data protection authorities worldwide have agreed on the value of considering privacy-by-design principles when developing privacy-friendly systems and software. However, on the technical plane, a profusion of privacy-oriented guidelines and approaches coexists, which provides partial solutions to the overall problem and aids engineers during different stages of the system development lifecycle. As a result, engineers find difficult to understand what they should do to make their systems abide by privacy by design, thus hindering the adoption of privacy engineering practices. This paper reviews existing best practices in the analysis and design stages of the system development lifecycle, introduces a systematic methodology for privacy engineering that merges and integrates them, leveraging their best features whilst addressing their weak points, and describes its alignment with current standardization efforts.**

Keywords— Privacy by Design; Methodology; Privacy Engineering; System Development Lifecycle; Privacy Impact Assessment; Risk management; Requirements Operationalization.

I. INTRODUCTION

The potential benefits of applying privacy-by-design (PbD) principles to system development processes are becoming increasingly acknowledged by policy-makers [1] and data-protection authorities [2]. However, the adoption by system engineers is still severely hindered by the lack of maturity of this discipline in terms of its practical formulation, which confronts engineers with a set of challenges.

First, engineers find hard to translate the abstract, ambiguous privacy requirements coming from the legal realm and based on ethical values and social perceptions into specific technologies and solutions [3][4]. Second, although PbD emphasizes the need to take into account the potential privacy issues from the outset of a project and through its whole lifecycle, very few privacy practices or approaches are specifically addressed at dealing with privacy from a system engineering perspective, language and mindset. Relatively few system engineers have any awareness of these practices [6] – they rather regard privacy as a mere theoretical concept [7]. PbD principles are not operational in their current state, hence the need for an engineering approach that integrates privacy in mainstream engineering methodologies allowing project

activities to move beyond socio-legal principles, and factoring in engineering privacy [5].

Third, a severe disconnection exists among different best practices (e.g., privacy impact assessments and privacy patterns); moreover, existing privacy guidance usually remains domain- or stage-specific. This forces engineers to choose among diverse, sometimes contradicting, approaches, and do their best to integrate them [8]. A standardization effort is required to provide engineers with a recognized methodology to tackle the aforementioned issues.

The EU-funded project PRIPARE (Preparing Industry to Privacy by Design by supporting its Application in Research) is facing these challenges by 1) providing a systematic methodology aimed at the complex ecosystem of all the stakeholders involved in the production of privacy-friendly systems, and which addresses the whole personal data and system development lifecycle (SDLC) of projects and systems, with the total disregard of their size and domain, 2) detailing the engineering processes that allow a move from abstract principles to technical requirements, designs and actual implementation; and 3) merging and connecting existing best practices in the area of PbD into a single methodology, whilst providing different alternatives and criteria to choose the most adequate for each context and at each stage of the SDLC.

This paper describes only some of the results of the PRIPARE project, focusing on the analysis and design stages. Some other aspects such as the integration with mainstream SDLCs (e.g., waterfall, iterative or agile) or processes related to privacy assurance (which are considered complementary to the engineering processes) can be found in the first full version of the methodology [9]. Section II below describes the relation between the two main existing analytical approaches to privacy requirements elicitation, namely goal-oriented and risk-based approaches, which are then independently discussed in sections III and IV. Moving into design, section V introduces three different approaches to designing system architectures for privacy. Section VI describes the current landscape of standardization and what is still missing there, and defines PRIPARE's position and where its contributions may fit. Finally, section VII concludes the paper.

IEEE computer society

II. Overview of Approaches to Privacy Analysis

The goal of a system engineering process is to fulfil the expectations of its stakeholders, by firstly finding the right trade-offs in order to set the balance among sometimes contradictory expectations, then translating them into operational requirements, and finally designing and implementing technical and organizational controls or measures that meet these requirements, in the systems being built.

There are many potential sources for privacy and security requirements: end-users' concerns, self-imposed policies, regulatory framework, prioritized risk scenarios and best practices and standards. However, it is not an easy task to elicit these multi-sourced privacy requirements as "*there is still not a unified view on privacy requirements engineering*" [10].

As a socio-technical issue, engineering privacy-friendly systems requires consideration of complex regulatory goals and user concerns, as well as stakeholders' expectations. Privacy is usually described as a set of high level principles, gathered in a number of sector-specific and generic guidelines and regulations such as the US Federal Trade Commission's (FTC) Fair Information Practice Principles (FIPPs) [11], the EU Data Protection Directive (DPD) [12] and the forthcoming General Data Protection Regulation (GDPR) [1], or the Organisation for Economic Co-operation and Development (OECD) guidelines [13]. In this context, eliciting privacy requirements implies translating these high-level, abstract principles into operational requirements, joining them with requirements derived from end-user concerns and other stakeholders' expectations, and solving the potential conflicts that may arise. These requirements can then lead to the design of technical and organizational measures (privacy controls), which are also suitable for the specific contexts where the system is intended to work. However, privacy principles are abstract concepts expressed in terms often far away from the domain of technical design, and thus they are often difficult to understand by engineers, who require a well-defined systematic methodology to define the appropriate privacy requirements in each case.

Currently, two major approaches coexist to discover and identify operational privacy requirements during a software development process, namely: 1) risk-based and 2) goal-oriented. Both approaches depart from a set of privacy principles, usually established by the pertinent legal framework, with the support of corporate policies. Each approach then develops along different paths (Fig. 1).
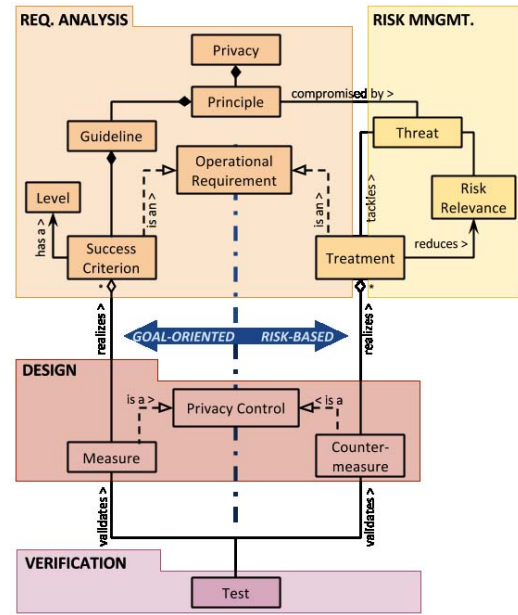

Fig. 1. Goal-oriented vs risk-based privacy requirements elicitation

Risk-based approaches start by identifying the assets to protect in the system under development and the threats that might compromise the accomplishment of the privacy principles on these assets. For instance, given the principle of ensuring the safety of personal data referenced in the EU DPD, a threat might involve an unauthorized party gaining access to some personal data. The threats are analyzed and the relevance of the subsequent risk is assessed, i.e., the combination of their probability and their impact. Then a treatment is proposed to address the risk associated with the threat. This treatment may range from doing nothing (accept the risk) to including requirements that may avoid or reduce the risk, modifying by its impact or its probability. During the design phase, the design team will identify and instantiate countermeasures (privacy controls) within the design to cover these requirements.

In a goal-oriented approach, each principle can be described as a goal that the system must fulfil. For example, data protection authorities' goals in Europe are related to the data protection principles stated in the EU GDPR, such as that of 'accountability' (i.e., ensuring and demonstrating compliance with data protection principles in practice). In turn, a user goal might consist in remaining anonymous while using the system. Each high-level goal (principle) can be deconstructed into a set of lower-level guidelines required to meet the goal, and each guideline can be in turn described as a set of operational requirements. Different requirements may turn out to be more or less critical to meet the privacy goals, which may be reflected by different priority levels attached to each requirement. Consequently, systems may protect privacy at different degrees by showing one of several levels of conformance with privacy goals, depending on the priority of the requirements they implement.

Risk-based and goal-oriented processes are complementary, as both aim to provide an understanding of what the system has to do in order to comply with the privacy principles by means of a set of privacy requirements. However, they differ in the way they tackle this endeavor, as the former focuses on identifying a set of problems that must be fixed, while the latter focuses on identifying a set of features to be built. From another point of view, the two approaches complement each other as the goal-oriented one focuses on preventing privacy risks while the risk-based one considers controls for those risks that could not be prevented.

Both approaches also show some contrasting features. On the one hand, a goal-oriented approach may be easier to follow by practitioners who are educated in systems engineering but have less expertise in privacy engineering. Guidelines detail how to concretize principles based on reusing previous, successful experience and knowledge. Notwithstanding that, these guidelines are not able to grasp a common privacy policy for the diverse systems to which they can be applied, as they do not deal directly enough with the specifics of each of them. In any case, they remove uncertainty by providing a minimum level of privacy at an early stage of the development process, thereby truly following the PbD paradigm. In addition, privacy-protecting decisions do not depend so much on the judgment of the system analyst as on the guidelines used, which are based on community-agreed practices. The decision to apply a specific privacy control (e.g., encrypt all traffic) will ultimately depend on the system developer, but it may rely on informed guidance rather than on their arbitrary judgment. Finally, success criteria provide a very basic but objective privacy metric as a result of the level of conformance selected.

On the other hand, a risk-based approach focuses on the specificities of the system being developed, which allows for a deeper understanding and covers the uncertainties derived from potential privacy threats. This aspect is interesting as there are currently several sources where privacy risks are described, for example [14].

Some frameworks for privacy analysis [15], whilst not explicitly, acknowledge the need to combine a goal-oriented with a risk-based approach in order to develop privacy-friendly systems. These frameworks claim to be risk-based; however, their major activities rather focus on barriers that prevent the achievement of privacy targets, instead of threats to privacy. These targets usually are derived from the corresponding legal framework (e.g., the EU DPD) of high-level privacy principles, which could be extended or adapted according to self-imposed policies or specific user concerns.

In fact, the line that differentiates threats from goals is blurry. While sometimes it can be considered to be just a matter of positive or negative wording of a specific statement, in other cases, each of the approaches are suited for identifying different types of requirements (e.g., when eliciting requirements related to the data subject's rights, it makes more sense to talk of goals than of risks).

The reasons why we consider frameworks such as [15] goal-oriented (instead of purely risk-based) are:

- they follow a systematic approach using catalogues as means to translate the high-level privacy principles into operational privacy requirements;

- they define threats in terms of facts that can be assessed without any intrinsic uncertainty, disregarding the probability of occurrence; and

- the selection of requirements and their specification according to the degree of protection demanded support systems that offer low, medium or high protection of personal data.

PRIPARE's methodology explicitly combines both approaches, by recommending, first, a goal-oriented approach that reduces privacy uncertainty at an early stage of the development process by eliciting a set of detailed requirements by following well-known, community-agreed catalogues (still to be developed); and, second, a system analysis approach for the remaining, system-specific risks, identifying adequate treatment according to several factors such as the risk level. In the end, some residual risks might remain, which must be identified and documented. These risk-based and goal-oriented approaches are respectively detailed in the next two sections.

III. RISK-BASED PRIVACY ANALYSIS AND IMPACT ASSESSMENT

While risk management methodologies have been used since World War II, in the field of security, they only date back to 1974 [16], and it is only in the last few years that they have been re-oriented to address privacy-specific issues by including privacy principles and concepts (identifiability and sensitivity of data, context and consent, impact on data subjects or privacy-specific safeguards).

When an organization projects the development of a new system that deals with personal information, privacy risks may arise. A privacy impact assessment (PIA) is a systematic process conducted by the organization in consultation with stakeholders to identify privacy risks and minimize their impact [17]. The new EU GDPR already foresees mandatory Data Protection Impact Assessments (DPIAs) whenever systems are likely to present risks to the rights and freedoms of data subjects (e.g., monitoring publicly accessible areas or when processing sensitive information), which will mean a major push in the uptake of PIA practice.

Prominent, current best practices on PIA [18], standards [20][21] guidelines [15], frameworks [22] and methodologies [23], all agree that a risk assessment process is key in order to conduct a PIA. Although some of these approaches may be very specific [15] in terms of how to conduct an assessment, others rely on the practitioners' own insight [21]. PRIPARE has analyzed these and other documents in search of differences and common points that may be leveraged during the merger of all those, in particular, to understand how they handle legal compliance, the way they measure risks and their impact, and how they address the privacy issues detected. Following we provide an overview of the key findings, and what elements have been introduced in PRIPARE's methodology.

A. Complying with the Legal Framework

While the whole PIA process can be considered as a means to ensure legal compliance, all of the PIA frameworks analyzed include a point or section that explicitly deals with the legal aspects of the project. Although law compliance could in principle be handled as just another risk category ('legal' or 'compliance risks' [23]), it is more practical that PIA frameworks devote a separate activity for this, whose aim is to identify the relevant legal framework(s) and to ensure that the key defining elements of the project (e.g., business objectives, features or specific implementations) abide by the law. In order to facilitate this, most impact assessment frameworks provide a questionnaire, which attempts to make the practitioners ensure that the project is legally compliant.

We regard these questionnaires, which can be described in layman's terms, as particularly useful for new, inexperienced practitioners as they reflect, in a practical way, specific laws or groups of laws applicable to systems in general or to specific domains. This checklist approach in no way diminishes the need to assess the privacy risks of any system.

B. Measuring Impact

In risk management, the impact of a specific risk is traditionally measured in terms of economic losses. However, in the privacy domain, the impact is split and not only does it affect the organization, which may have financial losses (business impact), but also – and more importantly, in the light of fundamental rights – it falls on the shoulders of the data subject, who may be exposed to impacts in terms of their social standing and reputation, their financial well-being, their personal identity and their personal freedom. A contrasting approach is currently followed by some methodologies [23] that mainly focus on the data subject by measuring features such as the identifiability and sensitivity of all personal data collected, stored or processed.

PRIPARE's methodology recommends following a dual perspective (already present in [15]) considering the impact for both the organization and the data subjects.

C. Measuring Risk

There are several formulae used to calculate the risk index (or relevance) of feared events or threats, which can usually be reduced to a combination of the potential impact and the probability of occurrence. The scales employed to measure these individual factors, and how to react to a given risk index, vary from one framework to another. In [23], for example, a 1 – 4 scale is employed to measure both probability and impact, and it maps the risk index to 4 levels (from negligible to maximum). On the other hand, in [15] a three-level scale (low, medium and high) is associated with privacy targets, and ignores the probability of occurrence of events, which can make the framework considered more similar to a goal-oriented approach than to a risk-based one.

In our opinion, there are no specific inherent benefits in choosing one type of scale over another. Different scales will adjust better to some domains and some regulations may impose some specific granularity of levels. Practitioners will

have to decide on their approach according to the system specificities and their internal and external requirements.

D. Addressing Privacy Issues

Privacy controls are technical and organizational measures that are incorporated into systems and organizations, responsible for, or interacting with, those systems in order to address privacy issues (i.e., threats). After identifying the privacy issues that threaten the system, the risk management process must identify those requirements that may help to treat these risks subject to three different strategic risk management decisions: avoidance, modification/reduction, sharing/transfer. In some cases, mitigation will not be possible and the risks will have to be retained; these remaining risks should be clearly communicated to all the stakeholders involved. Each identified requirement brings with it constraints, costs, limitations and implications that must be correctly balanced to achieve the business objectives in a privacy-friendly way, without compromising other aspects such as performance or usability. Relating privacy requirements to specific threats and/or high level principles guarantees the traceability and accountability of the whole analysis and design process.

The goal-oriented approach, complementary to the risk-based approach, is described in the following section.

IV. GOAL-ORIENTED PRIVACY ANALYSIS

PbD is a process that involves technical and organizational means that embed and implement privacy and data protection principles in systems with distinct functionalities. As these principles are often derived from law, and as they are too often disconnected from engineering practices, it is important to provide a systematic method to make these high-level principles operational. The method described in the operationalization process must be supported with requirements catalogues or guidelines (which may be domain-dependent) that will aid engineers in such a critical task.

A. Requirements Sources

Stakeholder needs and requirements represent the views of those at the business or enterprise operations level—that is, of users, acquirers, customers, and other stakeholders as they relate to the problem (or opportunity), as a set of requirements for a solution that can provide the services needed by the stakeholders in a defined environment [24]. Although users, acquirers, or customers may be familiar with the business domain to which the system belongs, very often the same stakeholders are not aware of or do not have the knowledge required to define the privacy goals of a system. Hence, the need to have a complementary approach for transforming high-level principles—which may stem from regulations, internal policies or a body of knowledge shared by a community of practice—into requirements, minimizing the need to involve the privacy-unaware stakeholders, thus providing a system with a baseline that includes a basic level of privacy.

PRIPARE's goal-oriented approach to privacy analysis is accompanied by a proposed catalogue of requirements [25], which still has to be evolved and harmonized with the privacy practitioners' community. The approach to goal-oriented

requirements provided within PRIPARE will make a set of requirements that are:

1. heuristic, as they are compiled and mapped from the experience reflected in available best practices and domain-specific guidelines [15][26];

2. stakeholder-neutral, as they reflect the variety of perspectives of privacy requirements expressed by the different agents involved;

3. structured and hierarchized, as they are organized into a five-layer, successively refined model, from abstract principles and guidelines, to objective and operable definitions for privacy requirements that can be used to design and embed controls or measures and design test procedures;

4. prioritized, according to the level of protection they provide;

5. and predefined, usable as inputs by system developers, who can apply them straightforwardly to their systems.

This catalogue draws its inspiration from a similar one that was standardized and developed for achieving accessibility within the World Wide Web Consortium's Web Accessibility Initiative (W3C WAI): the Web Content Accessibility Guidelines (WCAG) [27]. A further discussion of the translation of the accessibility requirement concepts into the privacy realm can be found in [28].

B. Operationalization Process

The process that enables the transformation of high-level privacy principles into operational requirements must be systematic, repeatable and easy to follow by engineers who are less privacy-savvy. The process is divided into two main phases: analysis and design.

The analysis phase starts from the set of abstract privacy principles to be applied, and successively refines them to select from the whole set of privacy requirements available in the catalogue to those that are deemed to be implemented as technical or organizational measures. They deal with not only privacy protection ex post facto, but also design strategies that foster users' privacy ex ante. The selection process considers the functional description of the system (dealing with its boundaries, data flows and privacy roles), and the desired level of conformance, in addition to organizational restrictions (e.g., in terms of performance, reliability, and budget). More specifically, it involves three steps:

1. Identify the privacy principles that will guide the selection of a set of privacy-related requirements defined internally or externally, according to the overall definition of the privacy goal assumed by the organization, or established by the regulatory framework.

2. Determine the required level of conformance for the system. It may be self-imposed following internal policies, or imposed by regulations or other stakeholders.

3. Determine the applicability of each privacy requirement, depending on the system specification, the level of conformance desired and/or needed, and other organizational constraints.

A detailed design phase complements the analysis later in the life cycle process, to translate the applicable privacy requirements into technical and organizational measures. This detailed design must match as well the decisions taken at the architectural level, resulting from the processes that will be presented in the next section. The framework also outlines a test suite, which assesses the adherence of a system to the guidelines for a specific conformance level, by stipulating that specific test cases are created to validate conformance with the specified requirements. These processes remain out of the scope of this paper; for details, refer to [25].

V. DESIGNING PRIVACY-COMPLIANT ARCHITECTURES

The result of the requirements analysis phase (whether goal-oriented or risk-based) is a set of requirements for the system. The next phase is the design of the system based on these requirements. PRIPARE follows the approach that PbD should primarily materialize at the architectural level and be associated with suitable methodologies [29]. Many definitions of 'architecture' have been proposed in the literature. In this paper, we will adopt a definition inspired by [30]: The architecture of a system is the set of structures needed to reason about the system, which comprise software and hardware elements, relations among them (including data flows) and properties of both. In practice, it is increasingly necessary to address complex architectures, for example, distributed architectures connecting systems and large system architectures (systems of systems).

Among other benefits, architectural descriptions enable a more systematic exploration of the design space. Architectures are often described in a pictorial way, using different kinds of graphs or semi-formal representations such as Unified Modelling Language (UML) diagrams (class diagrams, use case diagrams, sequence diagrams, communication diagrams, etc.). Even though such pictorial representations can be very useful, thinking about privacy requirements is such a subtle and complex issue that the architecture language used for this purpose must be defined in a formal way. By formal, we mean that some of the properties of the architectures can be defined in a mathematical logic, and reasoning about these properties must be supported by a formal proof or verification system. A source of complexity in the context of privacy is the fact that it often seems to conflict with other requirements, such as functional requirements, integrity requirements, performance, usability, etc. Formal methods make it possible to materialize precisely the concepts at hand (requirements, assumptions, guarantees, etc.) and to help designers explore the design space and reasons about possible choices. These kinds of formal methods have already been applied in order to verify cryptographic protocols [49] and to identify privacy weaknesses in electric vehicle charging protocols [31].

The atomic architecture components are coarse-grained entities, such as modules, components or connectors. In the context of privacy, the components are typically privacy

155

enhancing technologies (PETs), and the purpose of the architecture is their combination to satisfy simultaneously the functional and the privacy requirements of the system. Depending on the initial situation and availability of privacy requirements, code and/or architecture (e.g., initially being or not being privacy-compliant and demonstrating or not such compliance), different strategies can be applied to build a privacy-compliant architecture. In the following, we present three approaches, namely: top-down, bottom-up, and horizontal. Levels in this context refer to degrees of abstraction, with the top level corresponding to requirements (properties) and the low level to code or architecture. In addition, each approach can involve iterative steps.

A. Top-down Approach

The top-down approach is illustrated by the Computer-Assisted Privacy Engineering framework (CAPRIV)[32]. It consists of deriving compliant architectures, starting from the set of requirements (privacy, functional, technical, etc.) resulting from the requirement analysis described in the previous sections. The process can be carried out in either a semi-formal framework or a formal framework (based on specifications of the individual components used in the architecture).

In the top-down approach, the requirements can be expressed in a formal language (for example, a variant of epistemic logic as in [33]), and different choices of architectures can be proposed to the designer based on the trust assumptions between the stakeholders. Different types of trust can be distinguished [32], such as blind trust (assumption that an agent always behaves as expected), verifiable trust (a posteriori verification), or verified trust (a priori verification).

B. Bottom-up Approach

The bottom-up approach is applicable when a first version of the code (or a model of this code) is available. The goal is then to extract properties from this code showing that the desired privacy requirements are satisfied. This approach has been applied in PRIPARE to an electric vehicle charging scenario [34]. A major challenge of these vehicles is their somewhat limited range, requiring the deployment of many charging stations. To effectively deliver electricity to vehicles and guarantee payment, a protocol exists as part of the ISO 15118 [35] standardization effort. A privacy-preserving variant of this protocol, POPCORN [31], has been proposed in recent work, claiming to provide significant privacy for the user, whilst maintaining functionality. We have defined a formal model of the protocol and its expected privacy properties in the applied Pi-Calculus [36] and used ProVerif [49] to check them. This approach has made it possible to identify weaknesses in the protocol and to suggest improvements to address them.

C. Horizontal Approach

The horizontal approach is illustrated by methodologies focusing on privacy-enhancing architectures (PEARs), a term coined by [37]. An example of such methodology could be based on architecture analysis and evaluation methods such as the "Cost–Benefit Analysis Method" [47] (CBAM) and "Architecture Tradeoff Analysis Method" (ATAM) [48],

developed in the Carnegie Mellon Software Engineering Institute. The main objective of this process is to start with an initial architecture and to enhance it in order to achieve the desired business objectives, whilst achieving privacy goals and avoiding privacy risks as well as ensuring security. Changes in the architecture are measured under different scenarios in order to determine if the change provides a positive or negative impact on the overall system.

The resulting PEAR process is based mostly on the use of scenarios, which are "structured means to state attribute requirements". There are six elements to a scenario: the source of a stimulus, the stimulus, the environment of the artifact being stimulated, the artifact itself, the response of the artifact and the measure of such response (Fig. 2).


Fig. 2. PEAR scenario elements

PEAR is an iterative process involving the following steps:

1. Present an initial architecture.

2. Identify and prioritize scenarios and quality attributes.

3. Identify potential architectural security and privacy enhancements.

4. Select and apply privacy and security architectural approaches to the scenario.

The PEAR process relies on the selection and application, at an architectural level, of privacy strategies, patterns and technologies, to address the challenge of resolving conflicts between privacy objectives and other requirements of the system. The process has been endorsed by the European Network and Information Security Agency (ENISA) [6].

VI. Towards Common Privacy-by-Design Practices

The advent of PbD will depend on the availability of an ecosystem through which consensus-based practices and solutions can emerge. In recent years, initiatives have been taken up to promote guidelines or codes of practice, with some of them leading to standard proposals. PIAs are one such example: guidelines are available today, such as in the UK [38], while work on standardization is ongoing (ISO 29134 [20]).

Currently, two major organizations contribute to the standardization work on privacy and privacy by design. The International Organization for Standardisation (ISO), in particular through WG 5 of ISO/IEC JTC 1/SC 27, has focused on standards related to frameworks (e.g., ISO 29100[39] – privacy framework; ISO 29101[40] – privacy architecture framework), and management (e.g., ISO 29134[20] – Privacy impact assessments, ISO 29151[41] – Code of practice for PII protection). OASIS has focused on standards related to the PbD process: OASIS PMRM TC specifies a process to convert

privacy requirements into operational requirements [21], while OASIS PbD-SE TC focuses on software engineering documentation [42]. PRIPARE contributes, firstly, to the provision of a dual viewpoint, by integrating both risk-based and goal-oriented approaches; and secondly, it integrates privacy as an attribute in architectural analysis. Finally, the European Commission has recently issued a mandate for the establishment of European standards on PbD [43]. It is important that such work remains complementary to the aforementioned standards, in order not to create fragmentation.

PRIPARE is currently an organization member of OASIS and actively participates in the PMRM and PbD-SE technical committees. PRIPARE members have already contacted their national standardization bodies in order to participate in the development of the aforesaid EC-mandated PbD standard, and will present the PRIPARE project and methodology to the relevant working group of CEN/CENELEC (European Committee for Standardization and European Committee for Electrotechnical Standardization).

One critical phase in the PbD process is the selection of technology measures (for instance, a security protocol). This requires specialized expertise that is seldom available. Two approaches have been proposed. The first is to create a technology community that would publish and share common solutions, often called privacy patterns. PRIPARE is currently working on the creation of this type of repository [44]. The second is based on the organization of domain-specific consensus (e.g., smart grids, intelligent transport systems, smart cities) on the selection of appropriate technologies. They correspond to the concept of best available techniques (BAT) promoted by the European Data Protection Supervisor [45], also recommended in the case of smart grids [46]. We believe these approaches to be complementary.

VII. Conclusions

PRIPARE has opted for leveraging complementary existing best privacy practices, by integrating them in two dimensions: first, along the different activities of the SDLC; and second, offering alternatives within the same activity, and recommending the situations where each applies best. That way, their synergistic positive impact is maximized and their weak points are cancelled, thereby taking advantage of their touch points.

Following on from the PIA and risk assessment study, PRIPARE has decided to follow a dual approach in its methodology, complementing a goal-oriented approach with a risk-based one, including the PIA approach to directly cover the legal compliance aspects; and end-user and business perspectives when measuring risks, also taking into account sensitivity and identifiability levels of personal data attributes.

This combination provides an objective, systematic, goal-oriented approach, complemented by a more subjective, risk-based approach, which may capture some privacy issues that cannot be addressed by merely applying rigid guidelines.

Including complementary approaches with the same goal in a unified methodology allows its practitioners to adopt the practices that provide the best fit to specific problems,

situations or organizations. PRIPARE has established a unifying reference model that links the common points of each practice, key for the successful merger of these practices. For instance, both goal-oriented and risk-based approaches depart from high-level principles that must be closely matched by identified privacy targets – expressed as guidelines – that finally guide the elicitation of privacy requirements, which help to achieve the privacy targets, and avoid, mitigate or accept the feared events.

The methodology has also identified specific processes to address the new accountability principle (appearing in the forthcoming EU GDPR [1]), and user-empowerment and usability, one of the major current challenges in PbD. This allows positioning PRIPARE's methodology as a tool to achieve truly user-centric PbD systems, whilst demonstrating compliance with the selected privacy principles. In no case does PRIPARE propose what systems need to be built, but it helps to achieve privacy in those systems that are to be created—and it may impact their design decisions.

PRIPARE is moving toward validating its methodology. For that, three research projects have been selected, funded by the European Commission's Research and Technological 7th Framework Programme (FP7), to embed privacy and security, ensuring both that these projects correctly address privacy issues, and that the methodology is validated in practice, in terms of being efficient, practical and aligned with real-world system engineering practices.

References

[1] European Parliament legislative resolution of 12 March 2014 on the proposal for a regulation of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation) (COM(2012)0011 – C7-0025/2012 – 2012/0011(COD)).

[2] 32nd International Conference of Data Protection and Privacy Commissioners, Resolution on Privacy by Design: Jerusalem, 2010. http://www.justice.gov.il/NR/rdonlyres/F8A79347-170C-4EEF-A0AD-155554558A5F/26502/ResolutiononPrivacybyDesign.pdf [Accessed 2-Feb-2015]

[3] S. Gürses, C. Troncoso and C. Diaz, "Engineering privacy by design", Conf. Comput. Priv. Data Protection (CPDP). Brussels, 2011. https://www.cosic.esat.kuleuven.be/publications/article-1542.pdf

[4] S. Spiekermann, "The challenges of privacy by design," Commun. ACM, vol. 55, n. 7, pp. 38–40. New York: ACM, July 2012.

[5] I. Kroener, D. Wright, "A Strategy for Operationalizing Privacy by Design," Inf. Society, v. 30, n. 5, pp. 355-365. London: Routledge, 2014.

[6] G. Danezis, J. Domingo-Ferrer, M. Hansen, J.-H. Hoepman, D. Le Métayer, R. Tirtea and S. Schiffner, Privacy and Data Protection by Design-from policy to engineering. Heraklion, Crete, Greece: European Network and Information Security Agency (ENISA), 2015.

[7] T. Hasson, I. Hadar, O. Ayalon, S. Sherman, E. Toch, and M. Birnhack, "Are Designers Ready for Privacy by Design? Examining Perceptions of Privacy Among Information Systems Designers", 42nd Res. Conf. Commun. Inform. Internet Policy (TPRC). Arlington, VA: 2014.

[8] S. Gürses. "Can you engineer privacy?", Commun. ACM, vol. 57, n. 8, 2014, pp. 20-23.

[9] PRIPARE, D1.2 Privacy and Security-by-design Methodology, December 2014.

[10] K. Beckers, "Comparing Privacy Requirements Engineering Approaches," 2012 Seventh Int. Conf. Availability, Reliab. Secur., pp. 574–581, Aug. 2012.

[11] R. Pitofsky, M. L. Azcuenaga, S. F. Anthony, M. W. Thompson, O. Swindle, M. K. Landesberg, T. Milgrom Levin, C. G. Curtin and O. Lev, Privacy Online: A Report To Congress. Washington, DC: Federal Trade Commission, June 1998.

[12] Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data. Off. J. Eur. Union, vol. 1995, series L, n. 281, pp. 31–50. Luxembourg: Publications Office of the European Union, 23 Nov. 1995.

[13] Organisation for Economic Co-operation and Development (OECD), Recommendation of the OECD Council concerning Guidelines governing the Protection of Privacy and Transborder Flows of Personal Data. Paris: OECD, July 2013.

[14] Open Web Application Security Project Foundation (OWASP), "Top 10 Privacy Risks Project" [Online]. Available: https://www.owasp.org/index.php/OWASP_Top_10_Privacy_Risks_Project. [Accessed: 15-Dec-2014].

[15] M. C. Oetzel, S. Spiekermann, I. Grüning, H. Kelter, S. Mull and J. Cantella (ed.), Privacy Impact Assessment Guideline. Bonn: Bundesamt für Sicherheit in der Informationstechnik (BSI), 2011.

[16] Guidelines for automatic data processing physical security and risk management, Federal Information Processing Standards Publication 31. Washington: U.S. Department of Commerce, National Bureau of Standards, June 1974.

[17] D. Wright and P. De Hert, "Introduction to privacy impact assessment," in *Privacy Impact Assessment*. Dordrecht, Netherlands: Springer. 2012, pp. 3-32.

[18] R. Clarke, "Privacy impact assessment: Its origins and development", Comp. Law Secur. Rev., vol. 25, n. 2, p. 123-135. Elsevier, 2009.

[19] P. De Hert (ed.), D. Kloza (ed.), D. Wright (ed.), K. Wadhwa, G. Hosein and S. Davies, Recommendations for a privacy impact assessment framework for the European Union. Brussels – London: European Commission - Directorate General Justice, November 2012.

[20] International Organization for Standardization (ISO), "ISO/IEC 29134 WD Information technology – Security techniques — Privacy impact assessment – Guidelines," unpublished.

[21] J. Sabo, M. Willett, P. F. Brown, G. Janssen and D. N. Jutla, Privacy Management Reference Model and Methodology (PMRM), Version 1.0. Burlington, MA: Organization for the Advancement of Structured Information Standards (OASIS), July 2013.

[22] Privacy and Data Protection Impact Assessment Framework for RFID Applications, Annex to the Opinion 9/2011 on the revised Industry Proposal. Brussels: Article 29 Data Protection Working Party, January 2011.

[23] Commission nationale de l'informatique et des libertés (CNIL), Methodology For Privacy Risk Management. Paris, June 2012.

[24] Systems Engineering Body of Knowledge (SEBoK), Guide to the Systems Engineering Body of Knowledge. http://sebokwiki.org/wiki/Stakeholder_Needs_and_Requirements [Accessed: 23-Jan-2015]

[25] A. Crespo García, N. Notario McDonnell, C. Troncoso, D. Le Métayer, I. Kroener, D. Wright, J. M. del Álamo and Y. S. Martín, "D1.2: Privacy and Security-by-design Methodology". PRIPARE, 2014.

[26] R. M. Blank, P. D. Gallagher, Joint Task Force Transformation Initiative Interagency Working Group, NIST Special Publication 800-53, Security and Privacy Controls for Federal Information Systems and Organizations, Revision 4. Washington, DC: National Institute of Standards and Technology (NIST), April 2013.

[27] B. Caldwell, M. Cooper, L. G. Reid, and G. Vanderheiden, "Web Content Accessibility Guidelines (WCAG) 2.0," W3C Recommendation 11 December 2008. World Wide Web Consortium (W3C), 2008.

[28] Y.-S. Martin, J. M. del Alamo and J. C. Yelmo. "Engineering privacy requirements valuable lessons from another realm." 1st Workshop Evol. Secur. Priv. Requir. Eng. (ESPRE), pp. 19-24. New York: IEEE, Aug. 2014.

[29] T. Antignac and D. Le Métayer, "Privacy by Design: From Technologies to Architectures" (Position Paper), in Privacy

[30] L. Bass, P. Clements, and R. Kazman. Software architecture in practice (3d edition). SEI Series in Software Engineering, Addison Wesley, 2013

[31] C. Höfer, J. Petit, R. K. Schmidt and F. Kargl, "POPCORN: privacy-preserving charging for eMobility", in 1st Workshop Secur. Priv. Dependability for CyberVehicles (CyCar) at 20th ACM Conf. Comput. Commun. Secur. (ACM CCS 2013). Berlin, 4 Nov. 2013, pp. 37-48.

[32] T. Antignac and D. Le Métayer. Trust Driven Strategies for Privacy by Design, Inria Research Report, hal-01112856, version 1, February 2015.

[33] T. Antignac and D. Le Métayer, "Privacy Architectures: Reasoning about Data Minimisation and Integrity", in Secur. Trust Manage. (STM 2014), Lect. Notes Comput. Sci., vol. 8743, pages 17-32. Cham, Switzerland: Springer International Publishing Switzerland, 2014.

[34] M. Fazouane, H. Kopp, R. W. van der Heijden, D. Le Métayer, F. Kargl, "Formal Verification of Privacy Properties in Electric Vehicle Charging," in Proc. Int. Symp. Eng. Sec. Softw. Syst. (ESSOS 2015), in press.

[35] International Organization for Standardization (ISO), "ISO/IEC 15118:2013 Road vehicles – Communication protocol between electric vehicles and grid". Geneve, 2013.

[36] R. Milner, Communicating and Mobile Systems: The π-calculus. Cambridge, UK: Cambridge University Press, 1999.

[37] A. Kung, "PEARs: Privacy Enhancing Architectures", in Privacy Technologies and Policy, Lect. Notes Comput. Sci., vol. 8450, pp. 18-29. Cham, Switzerland: Springer International Publishing Switzerland, 2014.

[38] Information Commissioner's Office. Privacy by design. https://ico.org.uk/for-organisations/guide-to-data-protection/privacy-by-design/ [Accessed: 15-Dec-2014]

[39] International Organization for Standardization (ISO), "ISO/IEC 29100:2011 Information technology – Security techniques — Privacy framework". Geneve, 2011.

[40] International Organization for Standardization (ISO), "ISO/IEC 29101:2013 Information technology – Security techniques — Privacy architecture framework". Geneve, 2013.

[41] International Organization for Standardization (ISO), "ISO/IEC 29151 WD Code of practice for PII protection," unpublished.

[42] A. Cavoukian, D. Jutla, F. Carter, J. Sabo, F. Dawson, J. Fox, T. Finneran and S. Fieten, Privacy by Design Documentation for Software Engineers Version 1.0. (PbD-SE) Burlington, MA: Organization for the Advancement of Structured Information Standards (OASIS), work in progress.

[43] European Commission. Enterprise And Industry Directorate-General. Draft standardisation request addressed to the European standardisation organisations in support of the implementation of privacy management in the design and development and in the production and service provision processes of security technologies. 12/05/2014. http://ec.europa.eu/ DocsRoom/documents/5290 [Accessed: 15-Dec-2014]

[44] Repository for privacy patterns. http://privacypatterns.eu/ [Accessed: 15-Jan-2015]

[45] EDPS glossary; https://secure.edps.europa.eu/EDPSWEB/edps/lang/en/EDPS/Dataprotection/Glossary/pid/72 [Accessed: 15-Dec-2014]

[46] Commission Recommendation of 10 October 2014 on the Data Protection Impact Assessment Template for Smart Grid and Smart Metering Systems (2014/724/EU). Off. J. Eur. Union, vol. 1995, series L, n. 300, pp. 63–68. Luxembourg: Publications Office of the European Union, 18 Oct. 201

[47] Software Engineering Institute (Carnegie Mellon Institute), Cost Benefit Analysis Method (CBAM)

[48] Software Engineering Institute (Carnegie Mellon Institute), Architecture Tradeofff Analysis Method (ATAM)

[49] Bruno Blanchet, ProVerif: Cryptographic protocol verifier in the formal model