

On the Integration of Privacy-Enhancing Technologies in the Process of Software Engineering

Alexandra Klymenko^a, Stephen Meisenbacher^b, Luca Favaro and Florian Matthes

Technical University of Munich, TUM School of Computation, Information and Technology,

Department of Computer Science, Garching, Germany

{alexandra.klymenko, stephen.meisenbacher, luca.favaro, matthes}@tum.de

Keywords: Privacy-Enhancing Technologies, Requirements Engineering, Privacy Requirements, Privacy Engineering.

Abstract: The class of technologies known as *Privacy-Enhancing Technologies (PETs)* has been receiving rising attention in the academic sphere. In practice, however, the adoption of such technologies remains low. Beyond the actual implementation of a PET, it is not clear where along the process of software engineering PETs should be considered, and which activities must take place to facilitate their implementation. In this light, we aim to investigate the placement of PETs in the software engineering process, specifically from the perspective of privacy requirements engineering. To do this, we conduct a systematic literature review and interview 10 privacy experts, exploring the integration of PETs into the software engineering process, as well as identifying associated challenges along with their potential solutions. We systematize our findings in a unified process diagram that illustrates the roles and activities involved in the implementation of PETs in software systems. In addition, we map the identified solution concepts to the diagram, highlighting which stages of the software engineering process are vital in tackling the corresponding challenges and supporting the adoption of PETs.

1 INTRODUCTION


The emphasis placed on data privacy and data protection in today's technological landscape stems not only from the strict mandate put forth by modern privacy regulations such as the General Data Protection Regulation (GDPR), but also from the mounting pressure companies face from the threat of data breaches. The consequences of both non-compliance and privacy incidents can be significant, highlighting the utmost importance of proper technical and organizational measures to safeguard privacy.


As a technical response to the issue of data privacy, the class of technologies known as Privacy-Enhancing Technologies (PETs) has arisen from the research sphere, boasting provable privacy guarantees in a variety of settings. Amongst these technologies are notable leaders such as Differential Privacy or Secure Multi-Party Computation. While the promise of such technologies cannot be denied, the question remains of how to implement them in practice, i.e., ensure their transition from research to industry.

This operationalization of *academic* privacy into software engineering practices is seen as an integral

aspect of Privacy Engineering (Kostova et al., 2020), yet doing so comes with many challenges. A main category of these challenges, as pointed out by Kostova et al., relates to PETs, the adoption of which can be hindered by *changing requirements*, particularly in agile software environments. Another identified issue comes with the development of these technologies in the context of engineering practices, such as in the evaluation of privacy guarantees or the consideration of the effect of PETs on the deployment environment.

The challenges to fostering adoption beyond the requirements aspect are numerous (Klymenko et al., 2023). Many of the challenges revolving around the adoption of PETs as safeguards for data privacy pertain to the inherent complexity of such technologies, requiring expertise and resources to implement correctly. In addition, the relation of PETs to regulatory compliance is unclear. Other challenges only exacerbate the issue, such as the lack of defined structures and roles in place to facilitate PET adoption, as well as the question of the role of Privacy Engineering in this process. These challenges relating to PETs are echoed by Martin and Kung, who state that "*Privacy-Enhancing Technologies remain unknown for most engineers, due to the uncoupling between the PETs and the practice of systematic engineering and de-*

^a  <https://orcid.org/0000-0001-7485-2933>

^b  <https://orcid.org/0000-0001-9230-5001>

velopment; which makes engineers unaware or unknowledgeable of the proper applicability of such solutions.” (Martin and Kung, 2018)

Motivated by this statement, we aim to investigate the role of Privacy-Enhancing Technologies in the software engineering process. This is done in the light of privacy requirements engineering, which can be coupled with the stages of software engineering. Guided by these theoretical foundations, we seek practical perspectives on the integration of PETs into the Software Development Life Cycle (SDLC) by conducting 10 interviews with privacy experts in the software development field. From this, we hope to gain insight into the placement of PETs in the above-mentioned process, the challenges therein, and the potential solutions to help facilitate a more widespread adoption of PETs in software development.

Our findings include valuable insights into the diverse roles and activities involved in the process of privacy requirements engineering, particularly with the implementation of PETs. These findings are captured in a process diagram, which we create to give structure to the process of integrating and implementing PETs in the SDLC. Along with these insights, we identify nine unique challenges in the adoption of PETs in software engineering, which are mapped to nine solution concepts.

The contributions of our work are as follows:

1. We investigate PETs from the perspective of Privacy Engineering, with a focus on the roles involved and the activities that take place.
2. We present an artifact in the form of a process diagram that integrates PETs into the software development life cycle.
3. We propose solution concepts for increasing PET adoption and map them to the related activities in the process diagram.
4. We evaluate our artifact, supported by expert interview insights and quantitative survey results.

2 FOUNDATIONS

2.1 Requirements Engineering, Privacy by Design and Privacy Engineering

Requirements Engineering (RE) is the branch of software engineering that focuses on identifying and operationalizing goals, functions, and constraints of software systems (Zave, 1997). It plays an important aspect in software and system development, ensuring that the development process is aligned with real-world problems (Laplante and Kassab, 2022). More

specifically, RE involves eliciting, analyzing, modeling, validating, and managing the requirements of a system (Laplante and Kassab, 2022).

In addition to functional requirements, it is important to consider non-functional requirements such as privacy and security (Cysneiros and do Prado Leite, 2020), and to embed privacy measures directly into the design of a system as mandated by *Privacy by Design (PbD)* (Cavoukian, 2009). PbD prioritizes privacy as an essential requirement that must be addressed throughout the software development life cycle and promotes the proactive consideration of privacy from the early stages of the SDLC, i.e., ensuring that privacy features are incorporated into the system’s design before the implementation begins (Stallings, 2019).

Privacy Engineering (PE) is an emerging field of research and practice (Gürses and Del Alamo, 2016) that operationalizes the principles of PbD, providing “approaches for the inception and application of privacy-oriented solutions throughout systems and software development processes” (Martín García and Álamo Ramiro, 2017). PE involves implementation, deployment, and ongoing management of privacy features in a system, with the main goals of satisfying privacy requirements, protecting Personally Identifiable Information (PII), and mitigating the impact of possible breaches of personal data (Stallings, 2019). While Privacy Engineering is a distinct concept that, in the strict sense, builds upon Privacy by Design, the term is often used more broadly to refer to the range of privacy-related activities throughout the SDLC (Stallings, 2019). In this work, we likewise utilize the term in its wider context.

2.2 Privacy-Enhancing Technologies

With the introduction of modern privacy regulations such as the GDPR, implementing Privacy by Design principles has become a legal requirement for organizations. In particular, Article 25 of GDPR titled “Data protection by design and by default” mandates organizations to implement “appropriate technical and organisational measures” to protect the personal data of data subjects. In line with this, recent years have witnessed a significant increase in the research and development of technological solutions aimed at preserving the privacy of individuals, giving rise to a class of technologies commonly referred to as *Privacy-Enhancing Technologies (PETs)*. More formally, PETs can be defined as “*System of ICT measures protecting informational privacy by eliminating or minimizing personal data thereby preventing unnecessary or unwanted processing of personal data,*

without the loss of the functionality of the information system” (Borking and Raab, 2001). PETs can support PbD by ensuring data minimization, incorporating effective anonymization or pseudonymization solutions, and mitigating risks from personal data breaches by making the data inaccessible to unauthorized individuals (ICO, 2023).

Despite modern PETs such as Zero-Knowledge Proofs or Differential Privacy showing great potential for protecting privacy while allowing to derive value from data, they remain highly academic and are not widely adopted in practice (Klymenko et al., 2022; Klymenko et al., 2023). Among the main reasons for this are the lack of awareness of PETs, their inherent complexity, and little incentive for organizations to go beyond the “bare minimum” required for compliance and invest in more complex technologies (Klymenko et al., 2023). From the organizational side, the absence of defined structures, roles, and processes can likewise hinder the progress of such adoption (Klymenko et al., 2023), making the path to ensuring privacy ambiguous. In this work, we aim to provide structure to the processes behind the integration of PETs in software engineering, while also incorporating strategies to increase their adoption.

2.3 PETs and Privacy Requirements Engineering

While the body of related work is quite limited, several prior studies have addressed the challenge of managing privacy requirements and incorporating PETs into the software engineering process. The approach by Deng et al. involves translating broad privacy objectives and data protection standards into specific, detailed requirements that can be mapped to appropriate PETs based on their privacy protection goals (Deng et al., 2011). The proposed LINDDUN methodology includes developing a Data Flow Diagram based on the high-level system description and mapping privacy threats to this diagram. The identified privacy threats are then documented as misuse cases, essentially compiling potential threat scenarios. Based on these cases, privacy requirements are derived, and appropriate PETs are selected.

The PriS method views privacy requirements as organizational goals, using privacy-process patterns to illustrate how privacy requirements influence business processes, and recommending a set of PETs to meet these requirements (Kalloniatis et al., 2008). Although both abovementioned methods are comprehensive from a technical perspective, their focus may be too specific for them to be easily integrated into the broader landscape of software engineering.

A study by Hoffmann et al. takes a holistic view of privacy requirements, providing insight into how PETs should be integrated by design into the software engineering process (Hoffmann et al., 2008). The authors propose a four-step process for privacy requirements engineering to identify appropriate measures, technologies, and mechanisms that provide a balance between stakeholder interests and user constraints: 1) Description of Ambient Environments, 2) Identification of Stakeholders and Assets, 3) Analysis of Threats and Risks, and 4) Establishment of Privacy-Enhancing Technologies. The main limitation of this work is that the proposed requirements engineering process terminates with the selection of suitable PETs, leaving out subsequent development stages such as implementation, integration, and maintenance. In our work, we extend the reported study by considering the full software development life cycle. Through the comprehensive diagram presented in Figure 2, we illustrate the involved actors and related activities, emphasizing the interaction between technical, legal, and business-oriented teams.

3 METHODOLOGY

To guide our work, we define two research questions:

RQ1: How and to what extent are PETs included in the process of requirements engineering in the context of software engineering?

RQ2: From insights in the industry, how can the integration of PETs in the software engineering process be supported?

Answering these questions was accomplished by a mixed methodology of a Systematic Literature Review, semi-structured interviews, and a survey study.

3.1 Systematic Literature Review

The first step of our research involved a Systematic Literature Review (SLR), in which the goal was to survey existent literature on the topic of Privacy (Requirements) Engineering, with a focus on PETs. To conduct this study, the methodology of Kitchenham et al. (Kitchenham et al., 2015) was followed. The search engines used were IEEE Xplore, ACM Digital Library, ScienceDirect, Scopus and SpringerLink.

For the SLR, the first task was to define search queries from the above-mentioned search engines. As the initial goal was to focus on the role of requirements in the implementation of PETs, we centered the literature search around these aspects. We developed

a two-part query, which necessitated *privacy enhancing technologies* in the keywords of the paper, as well as *requirements engineering* anywhere in the metadata. As will be discussed in Section 4, the findings from the literature will be extended to the other stages of the SDLC, on the basis of the insights from privacy requirements engineering literature.

This search yielded 91 results. The uncovered sources were then screened by *abstract* and *introduction*. Guided by our inclusion criteria, which primarily necessitated a focus on PETs in requirements engineering, the number of relevant papers was filtered to seven. Finally, a forward/backward search was performed to include other relevant literature from the SDLC, yielding two additional sources, for a total of nine. The complete SLR process is illustrated in Figure 1 and summarized in Table 1, which also includes the references to the final selection of papers.

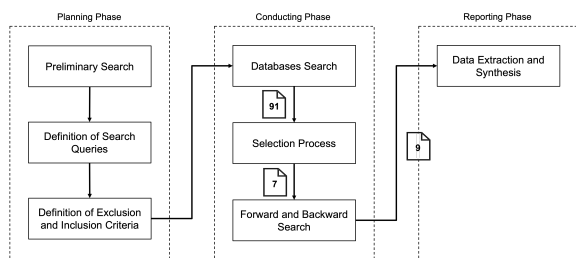


Figure 1: The SLR process.

3.1.1 Analysis

Following the collection of all relevant literature sources, a thematic content analysis was performed on the papers (Braun and Clarke, 2006). In particular, the sources were read and analyzed for recurring themes, which were then aggregated to form a succinct overview of the key topics covered in the nine literature sources. These findings served as the basis for the design of our semi-structured interviews.

3.2 Interviews

With the insights from the SLR, we then proceeded to conduct a round of semi-structured interviews with technical experts working at the intersection of privacy requirements and PETs. The goal of these interviews was to gain a practical perspective on PETs in software engineering, investigating their actual use in practice with a particular focus on privacy requirements. A secondary objective aimed at answering RQ2 by exploring challenges and possible solutions to increase the adoption of PETs.

3.2.1 Design

The interview guide was developed to inquire about the nature of privacy requirements, including their creation as well as the roles involved in this process. In particular, an emphasis was placed on PETs, namely how privacy requirements can be translated into the implementation of a PET.

The interview guide consisted of three main sections: *Background*, *Privacy-Enhancing Technologies and Requirements*, and *Looking Forward*. *Privacy-Enhancing Technologies and Requirements* inquired into the derivation of privacy requirements, the roles involved, the consideration of PETs, and the translation and verification of privacy requirements. *Looking Forward*, on the other hand, focused on investigating challenges for the inclusion of PETs in the software engineering process, as well as observed success factors and potential solutions going forward.

The main target audience of our interview study includes technical experts with a privacy focus. To find candidates, we used LinkedIn, initially employing search strings such as ‘privacy engineer’, ‘privacy champion’, ‘requirements engineer’, and ‘privacy architect’. The profiles of potential candidates were screened for mention of PETs, and candidates with such mention were assigned a higher priority for contacting. Candidates were contacted first informally via direct message. In the case of a positive response, a formal email invitation for an interview was sent. In total, 10 interviews were conducted, and the main demographics of these interviewees can be found in Table 2. All interviews were held via Zoom.

3.2.2 Analysis

After each interview, a full transcription was created using Otter.ai¹, and key points were extracted to answer all questions in the interview guide. Additional insights and themes were extracted by a team of three annotators, in order to mitigate personal or researcher bias. Constant comparison (Glaser, 1965) was employed, in the way that transcripts were annotated after each interview, in order to extract new findings and adapt the conduction of further interviews.

The extraction of themes was performed using the coding techniques proposed by Kitchenham et al. (Kitchenham et al., 2015). First, open coding was done on the raw transcripts to highlight excerpts of importance, which were then combined into themes (or axes) in axial coding. These themes were finally grouped into categories, which are presented in the ensuing Sections 4 and 5.

¹<https://otter.ai/>

Table 1: SLR results and references.

Search Engine	Initial Sample	Final Sample	References
IEEE Xplore	6	3	(Li and Palanisamy, 2019; Diamantopoulou et al., 2018; Anthonyamy et al., 2017)
ACM Digital Library	42	-	-
Science Direct	10	-	-
Scopus	30	4	(Hoffmann et al., 2008; Kalloniatis et al., 2008; Vrakas et al., 2010; Diamantopoulou et al., 2017)
SpringerLink	3	2	(Meis and Heisel, 2017; Deng et al., 2011)

Table 2: Interview study participants. Interviewees marked with an asterisk also participated in the evaluation study.

ID	Date	Role	Industry	Org. Size	Country	Exp.	Dur.
I1	05/23	Senior Privacy and Security Architect	IT Services and Consulting	Medium	Finland	10-20	61'
I2	05/23	Senior Privacy Engineer	Software Development	Medium	Germany	5-10	59'
I3	05/23	Privacy Engineer - Consultant	IT Consulting	-	Germany	5-10	64'
I4*	05/23	Senior Requirements Engineer	IT Services	Small	UK	5-10	47'
I5*	05/23	Staff Site Reliability Engineer	Software Development	Large	Netherlands	5-10	54'
I6	05/23	Senior Privacy Researcher and Developer	IT Services	Medium	Spain	3-5	59'
I7*	05/23	Principal Privacy Engineer	Online Retailing	Medium	Germany	20+	83'
I8*	05/23	Senior Privacy Engineer	Telecommunications	Large	Germany	3-5	61'
I9	06/23	Privacy Director	IT Services and Consulting	Large	Germany	10-20	40'
I10*	06/23	Product Manager - PETs	Software Development	Small	UK	1-3	53'

3.3 Artifact Evaluation

The final phase of our research consisted of an evaluation study in the form of an interview and survey. The goal of these two parts was to evaluate the artifact created after the conclusion of our semi-structured interviews, which will be introduced in Section 4.

3.3.1 Evaluation Interviews

The main goal of the second round of interviews was to refine and improve the artifact produced in the context of RQ1, introduced above. To ensure a fair review, the candidates from the first interview study were contacted, and five agreed to have a second interview to evaluate the artifact. As with the first interviews, an interview guide was prepared beforehand. The candidates participating in the evaluation are marked with an asterisk in Table 2.

The structure of the questions for the evaluation interviews proceeded in the following format:

1. Evaluation of included stakeholders
2. Evaluation of mapping between the SDLC and PE
3. Evaluation of interactions between roles
4. Evaluation of activities
5. Evaluation of solution placement
6. Open-ended suggestions for improvement

3.3.2 Evaluation Survey

Following these interviews, the process diagram was updated according to the provided feedback. Then, the same participants were invited to partake in a follow-up survey that featured the newly revised diagram. In the survey, the respondents were prompted to evaluate the diagram on a series of criteria, namely whether the diagram: (P1) delivers value, (P2) is de-

tailed enough, (P3) is comprehensible, (P4) covers all important stages, (P5) presents all needed activities, (P6) includes all relevant actors, (P7) integrates PETs correctly, (P8) presents solutions at the correct point.

P1-8 were evaluated on the Likert scale: *strongly disagree*, *disagree*, *neutral*, *agree*, *strongly agree*. The evaluation results are presented in Section 6.

4 PETS IN SOFTWARE ENGINEERING

In this section, we aim to answer RQ1, namely *how* and *to what extent* PETs are included in the process of privacy requirements engineering in the context of software engineering. Guided by the SLR findings and insights from the expert interviews, we create a process diagram that integrates the expert insights on the incorporation of PETs into the SDLC. In the creation of a comprehensive process diagram, we hope to guide practitioners in effectively integrating privacy requirements, the PbD approach, and PETs.

As the basis of our process diagram, we utilize the simplified SDLC model, including five main stages: *planning*, *analysis*, *design*, *implementation*, *testing and integration*, and *maintenance* (Akinsola et al., 2020). Parallel to the SDLC stages, we incorporate the corresponding privacy engineering stages (Hoffmann et al., 2008), augmented by the subsequent stages found in the diagram, including the implementation and integration of PETs in software systems.

The final process diagram is found in Figure 2.

4.1 Roles

Through the SLR and the initial round of interviews, we identified three main groups of actors involved in

the privacy requirements engineering process:

1. **Business actors:** mainly external stakeholders who hire another company to develop software or a product, along with the product manager from the developing company responsible for deriving business requirements and overseeing the entire product development process.
2. **Technical actors:** includes architects, software engineers, and privacy engineers, who are responsible for translating requirements into technical solutions and implementing them.
3. **Legal actors:** responsible for deriving legal requirements from applicable regulations and assessing the compliance of a solution before deployment. In Europe, this usually includes a Data Protection Officer (DPO) as mandated by the GDPR, and a team of privacy experts.

In addition, further external actors may be involved if the organization opts for certification or an audit.

In Figure 2, *Roles* illustrate the behavior and responsibilities of the actors involved in the process. While the descriptions for business and legal stakeholders are intentionally kept broad as they are outside the scope of this work, we provide detailed role descriptions for the technical actors below:

- A *Product Manager (PM)* is responsible for supervising the development and management of a product. A PM collaborates with external stakeholders to define goals and scenarios, helping to translate ideas into a concrete project. The PM also establishes the overall strategic direction for the product's privacy features, ensuring alignment with the organization's goals.
- *External Stakeholders (ES)* represent customers who engage an external company to develop a product. Their input is crucial for understanding privacy expectations and market demands, guiding the establishment of privacy requirements.
- *Privacy Engineers (PE)* are responsible for translating privacy requirements into technical specifications, identifying appropriate mitigation strategies, and working alongside Software Engineers to effectively integrate PETs.
- *Software Engineers (SE)* develop the software components and features of a product and collaborate with Privacy Engineers to integrate PETs and other privacy-related functionalities.
- *Requirements Engineers (RE)* capture, analyze, and document privacy requirements. They facilitate the communication between technical and non-technical stakeholders to ensure that privacy concerns are translated into explicit requirements.

- *Legal Teams (LT)* ensure that the product meets legal obligations by analyzing applicable regulations, inferring requirements, and verifying their fulfillment post-implementation. The expertise of the LT helps to bridge the gap between technical implementation and legal compliance.

4.2 Activities

In Figure 2, *Activities* are depicted as boxes either with solid lines if they are mandatory or dashed lines if they are optional. Each activity belongs to one of the three categories: strategic, technical, and legal. As a clear separation between the categories is not always possible, some activities are placed between categories. If activities are strictly related or must sequentially follow one another, they are grouped within an additional box, highlighting the importance of viewing them as a joint activity. Lastly, for each activity, the involved actors are specified. Below, we describe the activities based on the SDLC stages.

4.2.1 Planning and Analysis

The process begins with the decision to develop a new product, initiated either by an external stakeholder or internally by the product manager. In the initial stages, the scope of the project and business goals are defined. A thorough analysis involving the product manager takes place to assess the project's feasibility and potential financial benefits for the company. Requirements engineers must already be involved in this stage of the process in order to translate abstract ideas into specific requirements. On the business side, various usage scenarios and their related stakeholders are identified, subsequently undergoing an initial risk assessment on the technical side. The involvement of privacy engineers in this phase is essential, as some ideas may be rejected if they pose significant privacy risks. Based on the established scenarios and the system's intended use, privacy risks are identified and mapped to protection goals through the definition of adversary models. On the legal side, the requirements from the applicable regulations are derived, which are then merged with the identified protection goals to conclude the Privacy Impact Assessment, resulting in a comprehensive set of structured requirements.

4.2.2 Design and Implementation

The process proceeds with the task of identifying appropriate technical measures to fulfill the defined privacy requirements. At this point, PETs should be considered to provide an appropriate level of privacy and fulfill requirements. A collaborative effort between

software and privacy engineers is essential to determine the best solution, given the available expertise. Once the measures are agreed upon, the system is designed to integrate and operationalize them. The legal team reviews the solution to ensure that it meets compliance standards, and finally, the project is validated by the product manager and the external stakeholders. Several rounds of discussion may occur between the identification of technical measures and project validation, as consensus among all parties is essential before moving on to the implementation phase. Once the project is approved, the choice must be made between developing PETs in-house or purchasing a solution available on the market. The decision is both technical and business-driven: a lack of implementation skills might drive towards purchasing, while budget limitations could push for an in-house approach. The subsequent phases encompass infrastructure setup and system implementation, which are the responsibility of privacy and software engineers. Producing technical documentation is essential, not only due to regulatory requirements but also to highlight any deviations from the original design. During the implementation phase, the system undergoes periodic testing to evaluate its correctness, quality, and fulfillment of the requirements.

4.2.3 Testing & Integration and Maintenance

Once the system is implemented, testing is conducted to verify its performance and assess the impact of PETs, potentially uncovering new problems or requirements not previously considered. Test results and performance metrics are reported to the legal team and customers. The implemented solution is reviewed by requirements and privacy engineers for alignment with defined requirements and then assessed by the legal team for compliance. The results of these evaluations, along with the compiled documentation, are presented to the customer for final approval. Once the system is approved and prepared for deployment, the customer may opt for a privacy compliance certification to enhance its credibility, which can also be obtained after the system's deployment. Finally, periodic checks for new (1) business requirements, (2) attack vectors, and (3) legal requirements must be performed in the maintenance phase.

4.3 Current State of PET Adoption

Through the interviews described in Section 3.2, we investigated the current level of adoption and consideration of PETs in meeting privacy requirements. Our findings indicate a very limited practical use of advanced state-of-the-art PETs. I1 highlighted the fact

that many of such PETs are often overlooked, with the industry mainly relying on encryption and data access controls. I4 further elaborated on this by indicating the lack of knowledge and “skills to actually implement [PETs]” as the main barriers to their wider adoption. I5 indicated that unless PETs are a priority for the company, they are usually not considered. Together with I2, they support the idea that simple and tested solutions are usually preferred. In contrast, more advanced PETs are only considered if there is a strict need or a clear added value. Another interviewee (I8) attributed the limited use of PETs in the industry to the disparity between academic expectations and real-world industrial needs, saying “some requirements between academics and business are clashing” in terms of privacy preservation versus the value of data utility in practice.

The decision to adopt a PET is closely tied to the applicable legislation. Multiple experts (I1, I4, I7) reported that their privacy choices lean towards simple solutions that suffice for regulatory compliance. A prevalent notion appears to be that PETs are mainly oriented at organizations handling large volumes of sensitive data, leaving many to believe they are not relevant to their own business. Even if some organizations might consider PETs during the design phase, they often opt out later due to challenges hindering practical integration. Exceptions include highly regulated sectors, such as finance or healthcare, where advanced PETs are more likely to be used.

While advanced PETs remain underutilized overall, some of our interviewees reported relevant practical experience. One category of PETs that appears to be most recognized and used is data anonymization and modification; for instance, I2 and I5 mentioned their previous use of Differential Privacy and k-anonymity. Interviewees I6 and I8, who serve as researchers and developers of PETs, reported substantial application of these technologies in their field.

5 INCREASING PET ADOPTION

Below we present our findings in response to RQ2, which aims to investigate the ways to increase the adoption of PETs in the software engineering process.

5.1 Challenges

Identifying and understanding the limiting factors is the foundational step in devising viable and effective solutions for improving the practical use of PETs.

According to the interviewed experts, the most prevalent challenge is the *limited knowledge and un-*

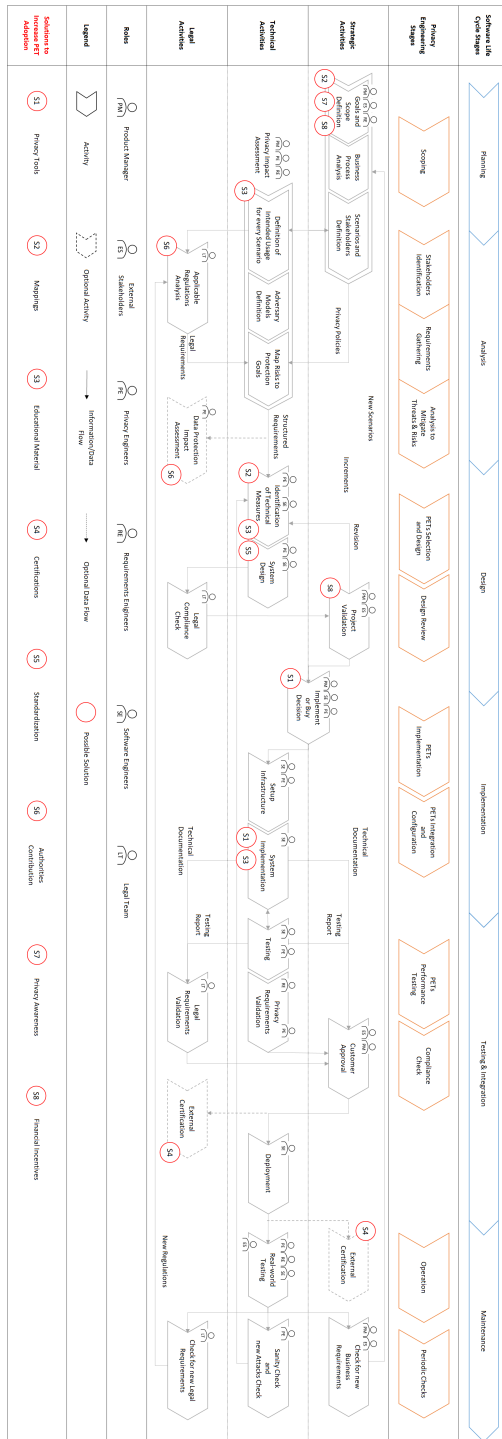


Figure 2: Process Diagram for the Implementation of PETs. In the top lane are the six stages of the Software Life Cycle. Underneath, the corresponding Privacy Requirements Engineering Stages are listed. The next three lanes feature the parallel work streams of (1) *Strategic Activities*, mainly concerned with business functions, (2) *Technical Activities*, which include technical design, implementation, and deployment, and (3) *Legal Activities*, in which legal teams and Privacy Engineers work to ensure compliance and verify requirements. The different role categories are illustrated in the *Roles* lane, and the *Legend* explains the different shapes used in the diagram. The final lane, named *Solutions to Increase PET Adoption*, includes 8 identified solutions (discussed in Section 5), mapped to the most appropriate location in the process diagram.

derstanding of PETs, including their capabilities, possible application contexts, benefits, and limitations. Another common challenge is the inherent *complexity* of these technologies. The challenges of awareness and understanding of PETs are also echoed in recent work (Klymenko et al., 2023; Boteju et al., 2023), the latter of which focuses on the perspective of software developers. The interviewees expressed concerns regarding the feasibility of non-expert developers implementing PETs, maintenance difficulties, and the need for specialized domain knowledge. Other hindrances from a business perspective include the *limited time and resources*, as well as the *cost-effectiveness* of implementing PETs. Research and implementation of PETs requires time, expertise, and financial investment; as long as the business value of PETs is unclear, it is unlikely that their adoption will be on a company's roadmap. Further challenges include *optionality*, meaning regulations do not enforce the adoption of specific technologies, *reluctance* to be an early adopter of a technology due to potential fines, as well as the issue of *legacy systems*. The latter represents a challenge for adopting new PETs due to outdated architectures, incompatibilities with modern solutions, high adaptation costs, and employee resistance to change. Lastly, the *choice of technology* poses another challenge, as identifying the privacy risks and deciding on the right PETs for a given use case and requirements is far from straightforward.

5.2 Solutions

In the following, we briefly describe the insights gained from the interviewee's perspectives on addressing the identified challenges, including observed success factors in the context of adopting PETs.

- S1: *Privacy Tools* can help make PETs more accessible by encapsulating their complex implementation within user-friendly libraries and interfaces. Such tools can provide ready-to-use implementations of various PETs, allowing developers and organizations to integrate strong privacy measures into their systems and leverage the benefits of PETs without the need for a deep understanding of their technical details.
- S2: *Mappings* can serve as a guide for choosing appropriate PETs for particular use cases or requirements by creating a link between the technical properties of PETs and the real-world requirements of specific scenarios. By translating complex technical features into tangible functionalities, mappings help decision-makers, developers, and other stakeholders identify fitting PETs to meet their data privacy objectives.
- S3: *Educational Material* can increase the understanding of PETs among stakeholders, minimizing the perceived difficulty of implementing and integrating PETs into software systems. Such educational content can bridge the gap between theory and practice, providing technical stakeholders with the knowledge needed to implement PETs in their systems, leading to enhanced data privacy.
- S4: *Certifications* are of great significance in the field of data privacy. Possessing a compliance certification (e.g., GDPR) for a product, system, or data processing activity demonstrates commitment to responsible and accountable handling of personal data, increases business credibility, and fosters trust among customers and partners. Certifying technical professionals in accordance with standards like ISO/IEC 17024 ensures that they have the essential knowledge, skills, and comprehension required to implement privacy and risk mitigation practices during the development process.
- S5: *Standardization* can boost the adoption of PETs in software engineering by providing precise and universally accepted guidelines. Standardized procedures can simplify the integration of PETs, reducing complexity and uncertainty for developers who would be able to rely on established best practices and predefined protocols.
- S6: *Contribution from Authorities* can take various forms to promote the adoption of PETs. They can provide guidelines for the effective adoption and integration of PETs or enforce stricter data privacy regulations that explicitly mandate their implementation in specific contexts. Authorities could also reinforce existing regulations such as the GDPR through rigorous enforcement and fines, forcing companies to prioritize privacy practices and adopt PETs to mitigate risks.
- S7: *Privacy Awareness* is essential for business stakeholders to comprehend the potential of PETs. They must recognize the importance of data privacy and the value that PETs can bring to a company, including financial benefits, enhanced reputation, and trust. As such, it is important to make PETs commercially attractive by showing how investing in them can result in significant returns.
- S8: *Financial Incentives* emerge as strategic solutions to overcome the challenge of high costs associated with researching and implementing PETs. Through the provision of financial support (e.g., through government funding pro-

grams), organizations can cover additional expenses, making the integration of PETs more financially viable.

S9: *Tech Leader Adoption* can serve as a driving force for the broader adoption of PETs, as small and medium-sized enterprises (SMEs) often learn about new technologies from larger technology corporations. Furthermore, unlike SMEs with limited budgets, large organizations have the financial capacity to invest and experiment with new technologies, as well as to handle fines in case a technology fails to meet compliance standards or leads to a privacy breach.

Another important aspect that was mentioned by interviewees, although not explicitly as a means to enhance the adoption of PETs, is the incorporation of cross-functional teams. While not a distinct solution concept, it is addressed in Section 4.1, and illustrated in the corresponding interactions in Figure 2.

In Figure 3, we map the identified challenges to potential solution concepts. The mapping is many-to-many: one challenge may potentially be mitigated by several solutions, and one solution may be helpful in mitigating several challenges.

5.3 Solution Integration

To contextualize the proposed solution concepts, we used feedback from the interviewees to map each solution to the activities discussed in Section 4.2. Figure 2 presents the process diagram that incorporates the mapped solutions, which are depicted by red circles with the respective solution code. The final placement of the solutions was determined through iterative refinements performed in the context of evaluation interviews, discussed in Section 3.3.1. The only solution not depicted on the diagram is S9, as the adoption of PETs by big tech companies impacts the overall landscape of PETs and data privacy at large, and as such, cannot be associated with singular activities. Below, we briefly elaborate on the placement of solutions S1-S8 and their impact on the process.

Increasing *privacy awareness* is essential at the early stages of project definition when business stakeholders decide on the privacy investment level; understanding the financial benefits of investing in privacy and PETs can influence these decisions. *Mapping* business cases to appropriate PETs is important from both business and technical perspectives to demonstrate the applicability of the technologies to real-world contexts. The mappings are, therefore, likewise helpful at the beginning of the project, as well as during the identification of technical measures. *Financial incentives* for PETs affect all activities influenced by

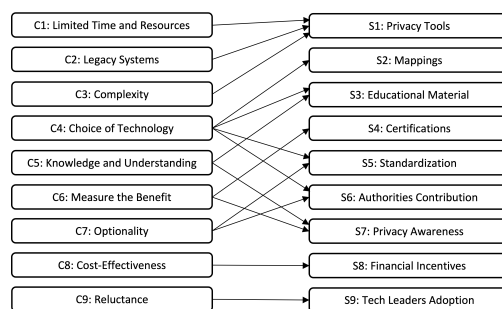


Figure 3: The mapping of challenges and solutions.

the project’s budget, including the initial scoping and the final review before validation. If the designed solution exceeds the budget, the project does not get approved, requiring privacy and software engineers to propose an alternative design using different technical measures. The technical knowledge on PETs obtained through *educational material* is helpful in the initial privacy assessment, as well as for identifying appropriate technical measures, and correctly implementing them in the system. In turn, stricter privacy regulations enforced by *authorities* would impact the general legal requirements and the data protection impact assessment, which may become mandatory for every system. *Standards* can positively influence the system design by ensuring a common understanding of PETs, providing guidance on their use, and building trust in their effectiveness. *Privacy tools* and open-source libraries for PETs can affect the “implement or buy” decision and the system implementation by offering easy-to-use and cost-effective solutions. Finally, requiring *certifications* for new systems would affect the two stages of the development process where the external certification is currently marked as optional. The certification activity would then become mandatory, facilitating compliance.

6 EVALUATION

As introduced in Section 3, the final step in the creation of our process diagram was to validate the artifact in a two-step evaluation study. The qualitative feedback received, as well as the survey results, are presented in the following. Note that Figure 2 shows the final version, after implementing the feedback.

Interview Feedback. Expert feedback in the evaluation interviews took the form of five types. For each feedback category, representative examples are provided, but they do not present the exhaustive list.

1. **Naming:** *Scoping* replaced *Problem Description*, *Configuring* PETs rather than *Tuning*.

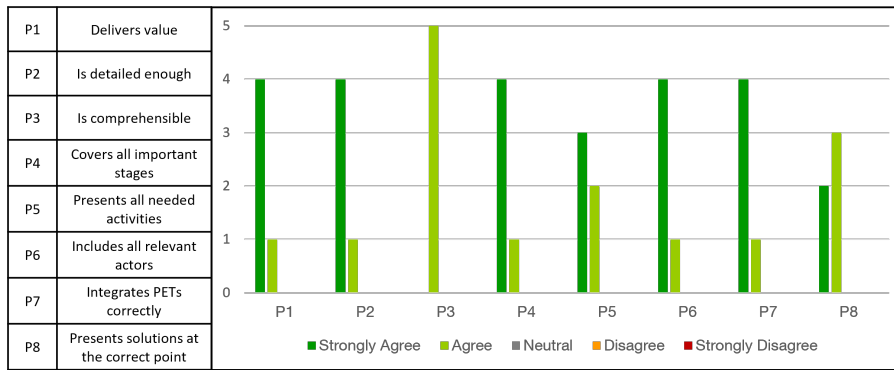


Figure 4: Artifact validation results (n=5). All eight criteria were met with either *strongly agree* or *agree* responses.

- Removing Ambiguities:** *Establishment of PETs* changed to *PETs Selection and Design*, add clarity to *Performance Testing*.
- Missing Stages:** *Operation* in the “Maintenance” stage, add the step *PETs Implementation*, add *Periodic Checks and Compliance Checks*.
- Missing Activities:** *Data Protection Impact Assessment (DPIA)*, *External Certification*, *Identification of Technical Measures*.
- Update Actors:** external stakeholders involved in *Real-world Testing*, continuous interaction between *Testing* and *System Implementation*.

Survey Results. The finalized process diagram with incorporated feedback was presented to the participants in an evaluation survey, introduced in Section 3. The results of this survey are depicted in Figure 4.

The results in Figure 4 show that all respondents agree that the model delivers value and is detailed enough, which implies it can be used to explain how to handle privacy requirements correctly and integrate PETs throughout the entire software life cycle. Regarding comprehensibility, respondents agree that the diagram is understandable and readable. The most crucial point, namely the correct integration of PETs in the process, was also perceived positively overall. Lastly, there is also general agreement on the presence of all needed actors, activities, and stages, as well as on the correct placement of the proposed solutions.

7 CONCLUSION

In this work, we investigate the integration of Privacy-Enhancing Technologies into the software engineering process. Guided by existing literature on the topic, we conducted 10 interviews with practitioners in the field of privacy engineering and PETs. Insights from these interviews served as the basis for the creation of

the *Process Diagram for the Implementation of PETs*, which was iteratively refined and validated.

Our findings show that the road to the implementation of PETs as part of the SDLC involves a large cohort of roles and activities, starting before the implementation itself and extending well after in the maintenance and continuous verification of PETs. Other findings include the perceived challenges hindering a more widespread adoption of PETs, as well as potential solutions. To make concrete where these solutions concepts can start to materialize, they are incorporated into the presented process diagram.

The main limitation of our work comes with the potentially limited generalizability of our findings. While the process diagram, and all of the supporting findings, aim to generalize the insights of a diverse group of experts, there is a possible bias towards larger organizations (>50 employees) with established privacy practices. Naturally, companies without the resources to support such a dynamic process may not be best captured by our artifact. Nevertheless, we believe our work lays the important groundwork for future refinements towards a greater understanding of the nature of PETs in practice.

Therefore, we make concrete paths for future work: (1) further validation of the proposed process diagram, taking into account more organization types and sizes, (2) case studies in companies where PETs have been successfully implemented, where these processes can be recorded to augment our diagram, and (3) feasibility studies on the proposed solutions.

At the core of our work, we seek to address the question of how PETs, which have predominantly remained a topic for research institutions, can begin to be adopted in the industry. Before this can happen in a more widespread manner, though, we argue that a greater understanding of its place in practice should be shared. With this, the true purpose of PETs can be brought to fruition, in the way that meaningful research turns to impactful practice.

REFERENCES

- Akinsola, J. E., Ogunbanwo, A. S., Okesola, O. J., Odun-Ayo, I. J., Ayegbusi, F. D., and Adebisi, A. A. (2020). Comparative analysis of software development life cycle models (sdlc). In *Intelligent Algorithms in Software Engineering: Proceedings of the 9th Computer Science On-line Conference 2020, Volume 1 9*, pages 310–322. Springer.
- Anthonyssamy, P., Rashid, A., and Chitchyan, R. (2017). Privacy requirements: Present & future. In *2017 IEEE/ACM 39th International Conference on Software Engineering: Software Engineering in Society Track (ICSE-SEIS)*, pages 13–22.
- Borking, J. and Raab, C. (2001). Laws, pets and other technologies for privacy protection. *Journal of Information, Law and Technology*, 2001.
- Boteju, M., Ranbaduge, T., Vatsalan, D., and Arachchilage, N. A. G. (2023). SoK: Demystifying privacy enhancing technologies through the lens of software developers. *arXiv preprint arXiv:2401.00879*.
- Braun, V. and Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2):77–101.
- Cavoukian, A. (2009). Privacy by design: The 7 foundational principles. *Information and privacy commissioner of Ontario, Canada*, 5:12.
- Cysneiros, L. M. and do Prado Leite, J. C. S. (2020). Non-functional requirements orienting the development of socially responsible software. In *Enterprise, Business-Process and Information Systems Modeling: 21st International Conference, BPMDS 2020, 25th International Conference, EMMSAD 2020, Held at CAiSE 2020, Grenoble, France, June 8–9, 2020, Proceedings 21*, pages 335–342. Springer.
- Deng, M., Wuyts, K., Scandariato, R., Preneel, B., and Joosen, W. (2011). A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements. *Requirements Engineering*, 16(1):3–32.
- Diamantopoulou, V., Androutsopoulou, A., Gritzalis, S., and Charalabidis, Y. (2018). An assessment of privacy preservation in crowdsourcing approaches: Towards gdpr compliance. In *2018 12th International Conference on Research Challenges in Information Science (RCIS)*, pages 1–9.
- Diamantopoulou, V., Kalloniatis, C., Gritzalis, S., and Mouratidis, H. (2017). Supporting privacy by design using privacy process patterns. In De Capitani di Vimercati, S. and Martinelli, F., editors, *ICT Systems Security and Privacy Protection*, pages 491–505. Cham. Springer International Publishing.
- Glaser, B. G. (1965). The constant comparative method of qualitative analysis. *Social problems*, 12(4):436–445.
- Gürses, S. and Del Alamo, J. M. (2016). Privacy engineering: Shaping an emerging field of research and practice. *IEEE Security & Privacy*, 14(2):40–46.
- Hoffmann, M., Heikkinen, S., Hornung, G., Thuvesson, H., and Schnabel, C. (2008). Privacy-enhanced personalisation in ambient environments. In *2008 IEEE 19th International Symposium on Personal, Indoor and Mobile Radio Communications*, pages 1–6.
- ICO (2023). Privacy enhancing technologies (PETs). *Information Commissioner's Office, United Kingdom*.
- Kalloniatis, C., Kavakli, E., and Gritzalis, S. (2008). Addressing privacy requirements in system design: The pris method. *Requir. Eng.*, 13:241–255.
- Kitchenham, B. A., Budgen, D., and Brereton, P. (2015). *Evidence-Based Software Engineering and Systematic Reviews*. Chapman & Hall/CRC.
- Klymenko, O., Kosenkov, O., Meisenbacher, S., Elahidoost, P., Mendez, D., and Matthes, F. (2022). Understanding the implementation of technical measures in the process of data privacy compliance: A qualitative study. In *Proceedings of the 16th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*, pages 261–271.
- Klymenko, O., Meisenbacher, S., and Matthes, F. (2023). Identifying practical challenges in the implementation of technical measures for data privacy compliance. *AMCIS 2023 Proceedings*. 2.
- Kostova, B., Gürses, S., and Troncoso, C. (2020). Privacy engineering meets software engineering. On the challenges of engineering privacy by design. *arXiv preprint arXiv:2007.08613*.
- Laplante, P. A. and Kassab, M. (2022). *Requirements engineering for software and systems*. Auerbach Publications.
- Li, C. and Palanisamy, B. (2019). Privacy in internet of things: From principles to technologies. *IEEE Internet of Things Journal*, 6(1):488–505.
- Martin, Y.-S. and Kung, A. (2018). Methods and tools for gdpr compliance through privacy and data protection engineering. In *2018 IEEE European symposium on security and privacy workshops (EuroS&PW)*, pages 108–111. IEEE.
- Martín García, Y. S. and Álamo Ramiro, J. M. d. (2017). A metamodel for privacy engineering methods. In *CEUR Workshop Proceedings*.
- Meis, R. and Heisel, M. (2017). Pattern-based representation of privacy enhancing technologies as early aspects. In Lopez, J., Fischer-Hübner, S., and Lambri-noudakis, C., editors, *Trust, Privacy and Security in Digital Business*, pages 49–65. Cham. Springer International Publishing.
- Stallings, W. (2019). *Information privacy engineering and privacy by design: Understanding privacy threats, technology, and regulations based on standards and best practices*. Addison-Wesley Professional.
- Vrakas, N., Kalloniatis, C., Tsohou, A., and Lambri-noudakis, C. (2010). Privacy requirements engineering for trustworthy e-government services. In Acquisti, A., Smith, S. W., and Sadeghi, A.-R., editors, *Trust and Trustworthy Computing*, pages 298–307. Berlin, Heidelberg. Springer Berlin Heidelberg.
- Zave, P. (1997). Classification of research efforts in requirements engineering. *ACM Computing Surveys (CSUR)*, 29(4):315–321.