

Defining Privacy Engineering as a Profession

Nikita Samarin

University of California, Berkeley
Berkeley, CA, USA
nsamarin@berkeley.edu

Liam Webster

University of California, Berkeley
Berkeley, CA, USA
liamwebster2001@berkeley.edu

Nandita Rao Narla

Future of Privacy Forum
Washington, D.C., USA
nnarla@fpf.com

Daniel Smullen

CableLabs
Louisville, CO, USA
d.smullen@cablelabs.com

Abstract

Rapid technological advancements, evolving legal frameworks, and increasingly heightened public concern over personal data have catalyzed the emergence of privacy engineering as a critical discipline. However, the “privacy engineer” role remains loosely defined, with significant variability in responsibilities, required competencies, and organizational positioning. This paper presents a qualitative investigation into the practices, challenges, and professional profiles of privacy engineers through 27 semi-structured interviews with US-based practitioners from diverse organizational contexts. Our thematic analysis reveals four primary themes: (1) the conceptual ambiguity surrounding privacy engineering roles, (2) a blend of ethical motivation, intellectual curiosity, and the desire for career growth driving professionals into the field, (3) organizational and regulatory challenges, such as misaligned incentives and the difficulty of translating abstract legal requirements into actionable technical solutions, and (4) the critical competencies required, including robust technical skills, effective cross-functional communication, and risk management expertise. Our findings contribute to a deeper scholarly understanding of privacy engineering as a multidisciplinary practice and offer practical guidance for organizations aiming to integrate privacy more effectively into their product development cycles.

Keywords

privacy engineering, professional practice, interviews, qualitative study, challenges, skills

1 Introduction

The field of privacy protection has undergone a profound transformation in recent years, driven by fast-paced technological innovation, evolving legal frameworks, and changing public attitudes toward personal information. Advances in data analytics, cloud computing, and artificial intelligence have enabled companies to collect and derive insights from vast volumes of user data. At the same time, consumers and advocacy groups have become more vocal about the need for clear, effective, and enforceable safeguards

that protect individual rights and prevent misuse of sensitive information.

As a result, today’s data protection environment imposes increasingly stringent requirements on companies collecting and processing personal information, with many laws and regulations in different jurisdictions. In early 2025, the International Association of Privacy Professionals (IAPP) reported that 144 countries had enacted some form of national data privacy legislation, covering approximately 82% of the world’s population, an increase of 3% in less than a year [2]. In the United States, a total of twenty states have passed comprehensive privacy laws [51]. This regulatory complexity is further supported by robust enforcement actions, as exemplified by the record-breaking €1.2 billion fine imposed on Meta in 2023 for violating international data transfer requirements of the EU General Data Protection Regulation (GDPR) [18].

Beyond regulatory pressures, privacy protection has become a critical factor in maintaining consumer trust and competitive advantage. A 2023 Pew Research Center survey found that 81% of American adults expressed concern about how companies use the data collected about them [45]. Scholarship has also noted that consumers are willing to pay a premium for privacy protection, particularly on Internet of Things (IoT) devices [24, 65] or when privacy information is available to them [64].

Consumer privacy preferences have direct implications for business strategy and outcomes. For example, the Cisco 2024 Consumer Privacy Survey reported that 75% of respondents “would not purchase from an organization they do not trust with their data” [17]. The competitive implications of privacy practices were further illustrated by the mass exodus of millions of users from WhatsApp to Signal following a global backlash over WhatsApp’s privacy practices that allowed data sharing with its parent company, Meta [31].

Despite the growing importance of embedding privacy into software products by design, addressing legal requirements and user privacy concerns has proven challenging in practice. This challenge has led researchers and practitioners to develop various methods, techniques, tools, and other solutions that consider privacy throughout the software engineering process. Many of these approaches eventually formed the basis for an emerging and rapidly expanding field of *privacy engineering*, receiving significant attention from industry, government, and academic stakeholders.

Rapid growth and recognition have contributed to an increase in scholarship on privacy engineering goals (the “what”) and methods of achieving them (the “how”). **However, a critical gap remains in understanding *who* is, or should be, responsible for putting**

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.

Proceedings on Privacy Enhancing Technologies 2025(4), 549–565

© 2025 Copyright held by the owner/author(s).

<https://doi.org/10.56553/popets-2025-0144>



these methods into practice and achieving these goals within organizations. Although the need for embedding privacy into the design and implementation of software and systems is clear, little is known empirically about the professionals who perform this function and how they do so effectively in practice.

Historically, societies held engineers responsible for ensuring that their products embody collectively desired values, including safety, quality, and reliability [26]. Beyond basic functional requirements, quality attributes such as performance, security and privacy are values for which engineers are responsible and are considered essential to consumer choices, to the extent that they may even be taken for granted as included in products they purchase [24]. Consequently, today’s consumers reasonably expect engineers to design software products and applications with their privacy and security in mind from the outset.

Researchers have consistently shown that software developers struggle to meet these expectations for various reasons, including misplaced responsibility for privacy protection [58], lack of awareness around privacy principles, frameworks and laws [53], communication gaps with other stakeholders [34], and the inherent complexity and context dependence of privacy itself [46, 47]. Given these difficulties, some have even questioned whether it is realistic to expect system engineers to address the challenge of privacy engineering on their own [5].

This combination of factors—the increasing complexity of privacy demands and the documented struggles of general developers to address privacy protection requirements—has fueled speculation and nascent efforts around the need for a dedicated specialist: the *privacy engineer* [20, 29, 55]. While the demand for professionals in this area continues to surge, and prior work has examined related roles like privacy officers [4] or privacy champions [60], or studied practitioners involved in privacy tasks more broadly [38, 40], a fundamental knowledge gap persists regarding the specific nature of the self-identified “privacy engineer” role. Specifically, there is a lack of empirical understanding around the following questions:

- (1) What are the day-to-day responsibilities and activities undertaken by those who identify as privacy engineers?
- (2) What are the core competencies, skills (both technical and non-technical), and experience required or valued in this role?
- (3) How is this role defined, integrated within organizational structures, and differentiated from adjacent roles like security engineering or legal compliance?
- (4) What are the practical challenges privacy engineers face and what methods do they actually employ, particularly given the acknowledged gap between theoretical frameworks (like PbD) and practical implementation [52, 56]?

To address this knowledge gap, we conducted 27 semi-structured interviews with privacy engineers and other professionals working in privacy engineering. We designed our interview protocol to investigate the roles, responsibilities, skill sets, deliverables, metrics, reporting structures, daily practices, and challenges associated with the privacy engineer role from the perspective of the practitioners themselves.

Our analysis provides a practitioner-driven account that moves beyond broad industry surveys, offering novel, qualitative insights

into the lived experiences, core competencies, and day-to-day realities of *self-identified* privacy engineers. We observe that these insights describe privacy engineering in a way which is substantially different from broadly accepted definitions used by organizations such as IAPP [36] and ISACA [37]. Specifically, our work addresses the research questions by revealing how these professionals use strong technical skills, cross-functional communication, and risk management expertise to effectively translate legal requirements, conduct privacy reviews, and bridge communication between legal and engineering teams. We illuminate the practical challenges they face, including conceptual ambiguity, misaligned incentives, and regulatory complexity. Furthermore, we demonstrate that privacy engineers perceive their roles as highly variable, fundamentally technical yet requiring broad legal awareness, and distinct from adjacent functions like security engineering or legal compliance. Finally, we also provide the de-identified and anonymized transcripts from our interviews, so that future researchers will be able to perform additional analysis (including reproducing our results and performing comparisons over time with new data in the future)¹.

Specifically, we found that the majority of our interviewed participants view the privacy engineering role as not strictly defined or used to mean different roles in different organizations in context. However, we observed several commonalities in the responsibilities, challenges, and motivations of the participants. We found that almost all of the participants regularly communicate and engage with various stakeholders, including legal and product teams, to achieve the incorporation of privacy requirements in product design. This strong emphasis on cross-functional communication corroborates and extends the findings of Kilhoffer et al., who recently identified privacy engineers as a crucial link between legal and engineering teams [40].

We observed other commonalities among our interview participants. For example, the participants stressed the importance of software engineering skills and processes as part of their role, highlighting the *engineering* part of privacy engineering. The majority of the participants also mentioned several other core competencies in their role, including understanding and translating legal principles into privacy requirements, conducting privacy risk assessments and threat analysis, and being able to make privacy visible within their organizations. Finally, almost all of the participants discussed the informal nature of their evaluation of their work and expressed the need for better metrics and evaluation techniques.

Key Contributions. Our analysis of the interviews revealed the following key findings.

- Bridging legal and technical domains: Privacy engineers serve as intermediaries between legal, compliance and engineering teams, translating privacy laws and policies into actionable technical implementations.
- Differentiating from security roles: Although it overlaps with security, privacy engineering focuses on data governance, regulatory compliance, and privacy-enhancing designs rather than preventing unauthorized access.
- Core technical competencies: The role requires a mix of software engineering, knowledge of system architecture, data

¹The dataset is available at <https://github.com/blues-lab/priv-eng-dataset/>.

protection strategies, and the ability to evaluate technical implementations for privacy risks.

- Cross-functional collaboration: Privacy engineers work closely with legal, product and engineering teams, requiring strong communication, negotiation, and influence skills to advocate for privacy priorities.
- Lack of standardization in an evolving profession: Privacy engineering is an emerging field with varying definitions across industries, requiring adaptability, continuous learning, and often a self-directed approach to defining responsibilities.

Our work contributes to a scholarly understanding of privacy engineering as an emerging discipline and offers guidance reflecting the essential skills, knowledge, and considerations that underpin an effective privacy engineering capability in an organizational setting. This guidance holds crucial lessons for academic, industry, and policy stakeholders. By providing a more explicit definition of privacy engineering roles and responsibilities, organizations can better align their privacy strategies, improve recruitment processes, and improve the integration of privacy engineering functions within their existing structures. Individuals interested in pursuing privacy engineering as a career will gain valuable insight into the skills, experience, and knowledge required to succeed in this field.

The findings of this study also create opportunities to address knowledge gaps by incorporating observed practices and skill sets into the academic curricula of engineers-in-training, as well as designing tools, solutions, and other interventions to assist engineers in privacy protection. Defining privacy engineering roles enables privacy engineers to be more effective conduits for regulators to interact with engineers and product developers. Privacy engineers can help give policymakers a clear view of technical privacy practices. Bridging the gap between engineering and policy helps ensure that regulations are technically feasible to implement and achieve their policy aims more precisely without knock-on effects. This clarity also improves compliance assessment and reduces ambiguity, leading to more effective and enforceable privacy laws.

2 Background and Related Work

This section explores the evolving landscape of privacy engineering, the challenges associated with implementing high-level privacy principles, and the role of a privacy engineer as a professional tasked with addressing privacy concerns and protections.

The notion of systematically embedding privacy considerations in technological systems has conceptual roots dating back decades [32], yet “privacy engineering” as a formalized term gained traction relatively recently in academic circles, such as the work of Gürses et al. [28]). Early scholarly discourse on privacy focused on legal frameworks and ethical theories rather than explicit engineering practices. However, as technologies became more data intensive and regulations like the Health Insurance Portability and Accountability Act (HIPAA) in the United States and the General Data Protection Regulation (GDPR) in Europe introduced stringent controls, academics began to investigate methods to ensure privacy was addressed throughout the systems development life cycle (SDLC).

Other recent work has attempted to describe the privacy engineer as a profession using an assumed definition² of privacy engineering [36, 37]. In contrast, our work seeks to define privacy engineering based on the experiences and opinions of self-described privacy engineers, without making such assumptions.

2.1 The Evolution of the Privacy Engineering Concept

Over the past three decades, researchers have tackled the challenge of “engineering privacy” using various approaches. In what follows, we describe several ways that the scholarship has framed “privacy engineering” that are most pertinent to our work.

Early academic efforts viewed “privacy engineering” as an activity within the *requirements engineering* process. Requirements engineering involves *defining, implementing, and validating the requirements* of software systems [10]. Kenny and Borking defined privacy engineering “as a systematic effort to embed privacy-relevant legal primitives into technical and governance design” from the outset of system development [39]. They proposed a framework for generating technical requirements from the articles of the EU Directive 95/46/EC (“Data Protection Directive”) [50]. Similarly, Breaux et al. proposed a methodology for extracting rights and obligations from legal and regulatory documents using formal models [8]. They applied their methodology to derive software requirements from legal rules to comply with privacy [8] and accessibility [9] regulations in the United States. Spiekermann and Cranor analyzed user perceptions and user privacy expectations to derive system requirements to address user privacy concerns and proposed architectural and governance guidelines to build privacy-friendly systems [57].

Subsequent research also framed “privacy engineering activities” under the broader umbrella of “Privacy by Design” (PbD) [29]. The PbD framework consists of high-level principles that, taken together, promote the incorporation of privacy considerations throughout the entire system engineering process [14]. However, the PbD framework does not offer prescriptive guidance on embedding its principles into traditional SDLC activities. In fact, Cavoukian, who developed and popularized the concept of “Privacy by Design” in the 1990s [12], has emphasized that there is “no single way to implement, operationalize, or otherwise roll out a PbD-based system” [13]. Other researchers have also noted the challenge of translating abstract privacy principles into real-world systems engineering [15, 43, 52, 56]. Through this lens, “privacy engineering” conceptually can be seen as a set of methods and activities aimed at integrating abstract privacy principles into systems using appropriate design patterns, models, and mechanisms.

Regardless of the specific problem context, there is a long-standing consensus that “engineering privacy” into systems is a challenging problem. In 2001, Feigenbaum et al. proposed a privacy engineering framework for digital rights management systems based on the Fair Information Principles (FIPs), noting that the research community “could make its largest contribution through the development of a practical methodology for privacy engineering” [25]. During the past two decades, academic, government and industry stakeholders

²Nadita Rao Narla, one of the co-authors of this paper, was a member of the IAPP board that created this definition.

have responded to this call by investing time and effort in developing various approaches to address privacy concerns under the collective heading of “privacy engineering” [11, 21, 35, 48].

2.2 The Role of the Privacy Engineer

Despite recent meta-studies on different privacy engineering methods and activities [1, 33, 44, 66], little is known about the desired knowledge, skills, and qualities of a professional tasked with performing the activities prescribed by these and other privacy engineering methodologies. Instead, scholarship has mainly focused on the barriers that have prevented engineers and developers from effectively considering privacy throughout early system design and into implementation. In addition to the absence of established tools and methodologies [54, 61] or their inapplicability to modern software development practices [42], researchers have identified other factors that impact the ability of system engineers to adequately embed privacy throughout SDLC activities. These factors include developers’ inability to design requirements based on a concept as contextual and ambiguous as privacy [46, 47], their limited knowledge of privacy principles, frameworks and laws [53], misplaced responsibility for privacy protection [58], misaligned incentives around privacy [23], and the lack of organizational culture and support around privacy [30]. These challenges have prompted some researchers to question whether the expectations placed on engineers to take responsibility for privacy were unrealistic, yet engineers were assuming this responsibility nonetheless [5].

We can also acknowledge the inherent complexity of privacy engineering methods and activities without minimizing efforts to develop better education, tools, and approaches for systems engineers to address the most common or problematic privacy issues. Scholarship has speculated that addressing this complexity requires specialized knowledge, distinct career paths, and professional development that might not yet be representative of the average system engineer. In 2010, Shapiro noted that there “doesn’t yet appear to be such a thing as a privacy engineer” when discussing the need for appropriately trained specialists akin to security engineers [55]. Gürses et al. also noted that engineering privacy requires “a specific type of expertise,” which could be achieved by “training future experts who are informed about the state-of-the-art research in security and privacy technologies, legal frameworks and the current privacy and surveillance discourses” [29]. Around the same time, Cranor and Sadeh noted that companies that embrace the PbD approach struggle to fill vacant positions for privacy engineers, which they described as “technologists responsible for ensuring that privacy is an integral part of the design process” [20].

Subsequent scholarship eventually recognized the presence of a “privacy engineer” as a practitioner in the emerging field. However, descriptions of this role remain vague or anecdotal. Several commonalities emerged between these descriptions, such as the prerequisite ability to work in multidisciplinary teams [19, 20], navigate complex legal and policy mandates [67], participate in communities of practice [7] and engineer (primarily technical) solutions for privacy protection [41]. Despite these descriptions, few academic papers have described the role of a privacy engineer in practice based on the accounts of actual practitioners in the field.

2.3 Privacy Engineering in Practice

Most existing studies have focused on the perceptions and behaviors of system engineers and software developers about privacy, rather than specialists who explicitly self-identify as privacy engineers. In what follows, we discuss the prior work most relevant to our study of the privacy engineer’s role.

Bamberger and Mulligan conducted semi-structured qualitative interviews with 53 corporate privacy officers identified as “industry leaders” to understand their role in privacy compliance and protection [4]. They found that organizations had better privacy outcomes when they embedded staff with specialized privacy protection expertise and personal responsibility for privacy into their business units. Tahaei et al. conducted 12 interviews with “privacy champions,” defined as professionals within software development teams who advocated for privacy and could influence privacy decision-making [60]. They found several approaches that privacy champions found effective in promoting better privacy protection, including privacy-focused meetings, informal discussions, managerial support, and cross-functional stakeholder communication (e.g., between legal and product teams). Horstmann et al. interviewed 30 participants involved in communicating and implementing privacy requirements in a company setting, including 10 privacy experts described as “individuals with a high level of knowledge or skill in data protection law and data privacy practices” [34]. They identified significant communication issues between privacy experts and software engineers. These issues were found to have the potential to degenerate into adversarial relationships, as developers may perceive privacy requirements as a hindrance during software development. Iwaya et al. interviewed 30 privacy engineering practitioners, individuals who met the following criteria: first, they are software practitioners, and second, they work with systems that process personal data. Participants, from nine countries, were interviewed to understand their perspectives, organizational aspects, and current privacy-relevant practices [38]. They found that the use of well-known or industry-standard privacy methodologies was not a typical practice among the interviewed practitioners; rather, their organizations’ own privacy culture influenced their behavior.

A recent study by Kilhoffer et al. interviewed 14 privacy engineers who met the following criteria: two or more years of experience in privacy engineering, active in privacy engineering, and a high level of familiarity with privacy and security standards, guidelines, and controls [40]. They found that privacy engineers spent significant time educating others in the organization about the importance of privacy and acted as a crucial link between the legal and engineering teams.

Our study’s findings both corroborate and extend the insights from prior research on privacy professionals. For example, the critical role of privacy engineers in bridging communication between legal, product, and engineering teams, a central theme in our results, resonates strongly with Kilhoffer et al.’s [40] portrayal of privacy engineers as crucial links and Tahaei et al.’s [60] work on the cross-functional communication approaches of privacy champions. Furthermore, our observation that the privacy engineering role often lacks a strict definition and is significantly shaped by organizational context aligns with Iwaya et al.’s [38] findings on the

prevalence of organization-specific privacy cultures over standardized methodologies. However, our research distinctively contributes by focusing explicitly on individuals who *self-identify* as privacy engineers. This allows us to move beyond functionally-defined roles, as seen in studies by Bamberger and Mulligan [4] on privacy officers or Horstmann et al. [34] on broader groups of privacy-involved professionals, to provide a practitioner-driven account of the privacy engineer's specific responsibilities, core competencies (such as the stressed importance of software engineering skills and the ability to conduct privacy risk assessments, as highlighted by our participants), and the practical challenges they encounter.

It is further important to note that the aforementioned studies adopt a sampling frame *functional*, selecting participants only based on the tasks they perform, rather than on *self-identified* participation as privacy engineers. This imposes limits on the ability to draw from skills and tasks that may not have been defined as related to privacy engineering in prior work. In contrast, our study purposely foregrounds self-identification, allowing practitioners themselves to define the boundaries of the emerging privacy engineering profession.

3 Methodology

In this section, we explain our IRB-approved interview study methodology, including data collection, analysis, and validation.

We chose semi-structured interviews because of the open-ended nature of our research questions. Furthermore, semi-structured interviews allowed us to further investigate participants' responses and skip questions as needed while maintaining the structure of our interview guide. We limited the interviews to US-based professionals who self-identified as working as privacy engineers or in a privacy engineering role (hereafter simply referred to as "privacy engineers") irrespective of their official job titles. This geographic restriction followed our IRB protocol, which approved only US-based participants. The interviews were held between late 2023 and late 2024, concurrently with the analysis. We designed our interview protocol to address the following general research questions:

- RQ1:** How do privacy engineers conceptualize their roles?
- RQ2:** What motivates individuals to pursue privacy engineering?
- RQ3:** What are the core competencies associated with privacy engineering?
- RQ4:** What do privacy engineers find challenging in their roles, and what strategies do they find effective in overcoming these challenges?

To support generalizable and rigorous qualitative results, we conducted interviews until new themes stopped emerging and we reached saturation [16]. Our subject pool was larger than the 12-20 participants recommended as best practice in the prior qualitative research methodological literature [27]. Therefore, our work can provide a foundation for future quantitative research and generalizable design recommendations.

3.1 Instruments

In this section, we describe the process of iteratively developing our screening survey and the interview guide.

Screening survey. The screening interview was aimed at identifying participants based in the United States who work full-time as employees or consultants in a privacy engineering role. To gain a better understanding of the context in which the participants worked, we included questions about employment status, job title, years of experience, sector and area of employment, privacy and cybersecurity certifications, and membership in relevant industry associations. It also included basic demographic questions, which were optional. After conducting the pilot interviews (which we describe later), we adjusted the text of the questions and emphasized the optional nature of the demographic questions. We include the final set of survey questions in Appendix B.

Interview guide. We developed our interview guide to highlight the unique characteristics of a privacy engineering role and to enable comparison with other similar but distinct privacy roles. To achieve this goal, we divided the interview into six distinct sections that cover: (1) participants' understanding of privacy engineering, (2) their motivation to pursue privacy engineering as a profession, (3) responsibilities and skills, (4) reporting and deliverables, (5) challenges and strategies, (6) evaluating success. All authors reviewed and provided suggestions for the interview protocol, including two authors with extensive experience as privacy engineering practitioners within their respective organizations. We include our final interview guide in Appendix C.

Pilot interviews. Before recruiting participants for the main study, we conducted three pilot interviews with privacy engineers from our personal networks. We tested and iteratively adapted both the screening survey and the interview guide after each pilot interview. These pilot interviews, although not included in the final analysis, were used to validate our interview script, timing, and overall approach, leading to refinements in our interview script and improved consistency among the different interviewers.

3.2 Recruitment

The goal of our recruitment efforts was to identify privacy engineers. In our recruitment messages, we invited prospective interviewees to participate in our study if they "work in privacy engineering" or were "privacy engineers or a professional in a similar role." Therefore, we relied on participants to self-identify as privacy engineers to avoid imposing our definition of the role of a privacy engineer. This inclusive criterion follows social identity theory, which holds that group membership is principally defined by self-categorization and shared meaning rather than externally imposed labels [3, 63]. Accepting participants on the basis of self-identification therefore allowed us to capture the emergent, practitioner-defined identity of "privacy engineers", even when employers still use disparate titles.

Recruitment messages were posted on LinkedIn to reach a broad segment of the professional privacy community. Furthermore, we used snowball sampling [6] and encouraged participants to recommend other potential participants within their professional networks (both offline by word of mouth and on social media). This method proved particularly effective in accessing privacy engineers who might not have seen our online recruitment messages. The recruitment message included a link to a landing page informing the candidates about the study purpose and providing them with

our study consent form (Appendix A) and our contact details. The participants who consented were then directed to complete our screening survey, which took an average of 4.6 minutes to complete. Furthermore, we asked eligible candidates to provide their email addresses to schedule a 60-minute remote interview through Zoom.

We excluded candidates who were students or worked part-time and invited eligible participants for an interview. We expressed our thanks to survey respondents who did not meet our selection criteria for their interest in our research and asked them to share information about our study with other potential candidates.

3.3 Interviews

We performed semi-structured video interviews via Zoom, each lasting approximately 60 minutes, though we remained flexible to extend this time if participants had more to share. A team of four interviewers, all extensively trained in qualitative research methods and privacy concepts, conducted the interviews. As discussed previously, we conducted three pilot interviews to ensure consistency before starting the main study.

Before starting each interview, we read the key information in the consent form aloud as a reminder, ensuring that the participants fully understood the nature of the study and how their data would be used and protected. We started the audio recording and the interview after receiving the verbal consent of the participants. The interview covered a wide range of topics, including the definition of privacy engineering, motivation and interest, required skills and responsibilities, reporting and deliverables, challenges and strategies for overcoming them, and evaluation metrics. Our complete interview guide is in Appendix C.

3.4 Qualitative Analysis

A qualitative approach was chosen for this study because it prioritizes depth of understanding over numerical measurement, making it well suited for exploratory research aimed at conceptually defining the profession of privacy engineering. Unlike quantitative methods that focus on statistical frequency, qualitative analysis allows for a nuanced exploration of the lived experiences, perspectives, and expertise of privacy engineers, capturing the complexity of their roles in ways that predefined categories or numerical data might overlook. This approach is particularly advantageous when studying emerging or evolving professions, as it enables the identification of subtle distinctions, implicit expectations, and underlying competencies that may not yet be standardized or widely recognized. By emphasizing rich, contextual insights rather than quantifiable trends, this method facilitates a more comprehensive and flexible definition of privacy engineering, accommodating the diversity of perspectives across different organizations, industries, and professional backgrounds.

Our data analysis process was iterative, employing a qualitative open-coding process that allowed the organic emergence of core concepts [59]. First, we used Zoom's audio transcription service³ to obtain interview transcripts, which we manually reviewed for transcription errors, then manually de-identified and anonymized to mitigate the risks of participant identification. All transcripts were reviewed by at least two researchers, and all references to

identifiable concepts (such as names, organizations, roles, etc.) were redacted. Indirect identifiers were also redacted. De-identified and anonymized transcripts were reviewed a final time before loading into MAXQDA⁴ for analysis. All researchers coded the first transcript independently before meeting to create an initial codebook. To reduce attentional fatigue, we divided each subsequent interview script into two parts, each containing three sections (see Appendix C for the questions in each section). Two pairs of researchers then independently expanded on the initial codebook based on a detailed analysis of their assigned section of the interview transcript. These codebooks were then compared by both pairs of researchers after coding 2-3 transcripts per iteration, discussed in depth, and merged. Any disagreements were resolved through careful consideration and consensus building [16].

The same two pairs of researchers independently coded additional transcripts to test and refine this consolidated codebook. This step led to further refinement of the codebook, with codes being added, merged, or clarified as needed to capture the nuances in the data. The final codebook was then applied to the remaining interviews, with both pairs of researchers coding all their respective interview segments to ensure consistency and reliability. Disagreements in coding were resolved through in-depth discussions between the coders.

Saturation and final validation. We determined that theoretical saturation had been reached when all the themes raised in new interviews fit within the existing codebook structure without requiring significant additions or modifications. At this point, we stopped further recruitment and revisited the initially coded interviews, recoding them with the final comprehensive codebook to ensure uniform analysis across all data.

After the interview transcripts had undergone grounded analysis (manually performed by human researchers with the assistance of the MAXQDA tool), a large language model (LLM) was used to validate that saturation had been reached, improving confidence that no further findings could be extracted which had not already been extracted through the manual analysis. OpenAI's GPT-4o LLM was chosen for this purpose, as GPT-type models have previously been shown to be effective in extractive analysis tasks of textual and interview data [22, 49, 62]. In order to comply with our data protection obligations, the OpenAI model execution environment was configured to prohibit sharing, training, and other secondary use of the de-identified interview data. In our approach, this additional validation step was used to provide greater assurances that the manual analysis which had already been performed was sufficiently rigorous, and that saturation had indeed been reached. The themes extracted using the LLM prompt were compared with those extracted through manual analysis to provide further assurance that the coding manual was complete. A complete coding manual meant that all themes had been extracted to the point of saturation and no additional themes would emerge through additional analysis. First, each of the interview transcripts was individually processed using a standardized LLM prompt designed to extract five concise bullet points summarizing the most important findings on the skills, competencies, expectations, and distinguishing characteristics of privacy engineers. The following prompt was used,

³https://support.zoom.com/hc/en/article?id=zm_kb&sysparm_article=KB0064927

⁴<https://www.maxqda.com/>

with the de-identified interview transcript attached to the prompt: *I'm working on a research project which aims to define the profession of "privacy engineer". Attached is a transcript from an interview between a researcher and a self-proclaimed privacy engineer. Extract 5 bullet points with the most important findings from the interview as it pertains to the definition of a privacy engineer; focus on the skills, competencies, and expectations that exist for these professionals, as well as what makes them distinct from other disciplines. Keep the points extremely succinct.* Additional prompts varying the number of bullet points (from 3 up until 10) were also used, but none of the prompts yielded any new themes which had not already been identified in the manual analysis.

Once all transcripts had been analyzed in this way, a second prompt was applied to synthesize the extracted bullet points generated across all interviews. This step identified the five most common themes emerging in the data set, while also isolating ten outlier findings that did not align as closely with the dominant themes. The following prompt was used, with the outputs from the previous prompts appended together and attached to the prompt: *I'm working on a research project which aims to define the profession of "privacy engineer". Attached is a file containing bullet points extracted from a series of interviews between a researcher and a self-proclaimed privacy engineer. Each interview was with a different set of people, and each interview resulted in 5 bullet points. Extract the 5 key bullet points that are common across all of the interviews. If there are any bullet points which seem not to fit with the rest, list them separately, limit 10 in total.*

By integrating this LLM-assisted approach with the initial human-led grounded analysis, both the depth and efficiency of the thematic synthesis were enhanced, providing a structured and scalable method for defining the profession of privacy engineering. If additional findings that had not been revealed during the manual analysis were synthesized, they would serve as evidence that saturation had not been reached. If no additional findings were synthesized, that is, all the findings synthesized by the LLM simply repeated those that had been generated previously through manual analysis. LLM-assisted analysis provides further assurance that saturation had been reached and that there were no additional findings to extract.

Crucially, the LLM-assisted analysis was performed only after the human researchers had completed their grounded analysis of the interview transcripts. This sequential approach ensured that the initial identification of key themes was driven by human expertise, minimizing the risk of overlooking nuanced or context-dependent insights. By first establishing a well-founded understanding of the data, we were able to critically assess the accuracy and relevance of the LLM output, ensuring that no hallucinated responses were introduced into our findings. Rather than relying on the LLM to generate novel insights, the role of LLM-assisted validation was to synthesize patterns that we had ideally already found in the data, highlighting potential gaps in our manual analysis if new themes were identified that had not been documented during manual analysis. This allowed us to cross-check our findings and determine whether any important themes had been underemphasized or missed, ultimately strengthening the rigor and comprehensiveness of our analysis.

3.5 Ethical Considerations

We took great care to ensure the confidentiality of the study participants. An institutional review board (IRB) at the University of California, Berkeley reviewed and approved the study protocol, including the analysis approach, data collection and retention practices, and the consent process. All participants provided their informed consent before collecting any personal information. No compensation was provided to participants to avoid any potential bias. All findings were rigorously de-identified and anonymized to ensure that disclosure of interviewee affiliations or reidentification of participants would be mitigated, thus mitigating potential harm upon public release of the redacted interview transcripts. To ensure anonymity, each participant was assigned a unique identifier (e.g., P1) which served as the only way to identify the interview.

As we previously mentioned in 3.4, we validated the results of our manual analysis by using an LLM on the de-identified and anonymized transcripts. This LLM analysis was conducted exclusively on the de-identified data, after the multi-pass manual redaction process described above. We acknowledge that while participants consented to the indefinite use of their de-identified data for future research by ourselves or others, as stated in the consent form and approved by the IRB, the consent form did not explicitly specify the potential use of AI or LLM-based analysis tools. The application of tools to analyze the de-identified transcripts, as well as share the transcripts, aligns with the consent provided under our IRB-approved protocol. All measures were taken to ensure participant confidentiality and prevent re-identification throughout our study.

We carefully removed any identifying information from our data and results. We replaced specific responses to the demographics survey (e.g., age or years in role) with ranges, each containing an approximately equal number of observations. Additionally, a senior privacy and data ethics researcher from the Future of Privacy Forum reviewed Tables 1 and 2 was consulted to ensure that the presented information is not identifying. Each transcript was manually analyzed and de-identified by multiple researchers in multiple sequential passes. We ensured that no facts or phrases that could serve as a way to identify the participant were mentioned. Identifiable facts and phrases were replaced with placeholders (e.g., [INSTITUTION], [LOCATION], [PROJECT] and so forth) and any responses from interviewees in which identifiable information was self-disclosed were fully redacted.

4 Results

In this section, we present an overview of the demographics and professional experience of the participants. We also remark on the results of the analysis process, including the number of interviews which were analyzed before reaching saturation, as well as the resulting interview recruitment and participation statistics. Finally, we detail the themes we observed after analyzing the interview transcripts.

4.1 Participants

All participants first completed a screening survey consisting of demographic and professional history questions. These questions

Table 1: Summary of participants’ demographics.

ID	Age	Education	Gender	Income
P1	26-32	Bachelor’s	M	\$200,001-300,000
P2	N/A	Ph.D./Equiv.	M	N/A
P3	40+	Bachelor’s	M	\$300,001-500,000
P4	N/A	Master’s	M	N/A
P5	N/A	Master’s	M	\$200,001-300,000
P6	N/A	Master’s	M	\$150,001-200,000
P7	N/A	Bachelor’s	M	\$300,001-500,000
P8	26-32	Ph.D./Equiv.	W	\$500,001+
P9	40+	Associate’s	M	\$100,001-150,000
P10	26-32	Master’s	M	\$200,001-300,000
P11	33-39	Master’s	W	\$75,001-100,000
P12	26-32	Master’s	M	\$200,001-300,000
P13	40+	Ph.D./Equiv.	M	\$500,001+
P14	40+	Ph.D./Equiv.	W	\$200,001-300,000
P15	26-32	Master’s	W	\$150,001-200,000
P16	33-39	Master’s	M	\$300,001-500,000
P17	N/A	Master’s	M	N/A
P18	N/A	Ph.D./Equiv.	M	\$500,001+
P19	33-39	Bachelor’s	M	\$500,001+
P20	33-39	Bachelor’s	M	\$500,001+
P21	26-32	Master’s	M	\$300,001-500,000
P22	N/A	N/A	NB	N/A
P23	33-39	Ph.D./Equiv.	N/A	\$500,001+
P24	N/A	N/A	N/A	N/A
P25	40+	Ph.D./Equiv.	M	N/A
P26	N/A	Master’s	W	N/A
P27	N/A	Master’s	W	\$200,001-300,000

‘ID’ column refers to participant ID. No response or ‘prefer not to disclose’ option is abbreviated as ‘N/A.’ Values in ‘Gender’ column: ‘NB’ = non-binary, genderqueer, or gender nonconforming; ‘W’ = woman; ‘M’ = man.

focused on the professional history and workplace of the respondents (e.g., their industry of employment, years of experience in privacy, size of the organization) and their demographics (e.g., age, education, income). In total, 27 participants completed the interviews before the recruitment was closed.

Demographics. Table 1 presents information on the demographics of the participants. The results of our survey indicated that the median age of the respondents was 39 years and the majority of the participants (64%) self-identified as men. Furthermore, most of the respondents (68%) had advanced degrees, including master’s and doctorate degrees, and 35% reported earning \$300,000 or more as their total annual compensation. We also found that 32% identified as being part of communities known to be historically disadvantaged; we refer the reader to Appendix B for a complete formulation of these questions.

Professional context. We found that all except one of the participants work in the industry full-time, of which the majority (93%) indicated working in the ‘technology and software’ sector. Other participants indicated, in addition to other areas of work, to work in ‘consulting’ (five participants), ‘e-commerce’ (four participants), and ‘banking and finance’ (four participants). Two participants

also indicated working for federal government agencies within the United States. Furthermore, half of the respondents work in large organizations employing at least 100,000 people, and most (56%) of the participants are based in California, where the most prominent technology companies are headquartered. We did not retain information on the specific companies in which the interviewees worked.

In terms of role titles, we found that 78% of the respondents had titles that contain the word “privacy” and 63% have titles that include the phrase “privacy engineer”. 44% of the respondents had roles that explicitly indicate seniority (e.g., “senior”, “lead”, or “leader”). We also found that participants had on average 7.5 years of privacy experience working with teams of an average size of approximately 20 individuals. 60% of the participants belong to the International Association of Privacy Professionals (IAPP), and 48% held IAPP Certified Information Privacy Technologist (CIPT) certifications. None of the participants made specific references to other forms of legal training. Table 2 presents information on the demographics of the participants and their professional context.

4.2 Interviews

Twelve interview transcripts were analyzed before we determined that saturation was reached. Further manual analysis of 15 additional transcripts that did not result in any changes to our codebook increased our confidence in this finding. The interview recruitment was then closed with 27 interviews completed and the corpus of transcripts was frozen. Finally, all interviews were examined using LLM-assisted analysis to attempt to reveal findings that did not match those extracted by manual analysis. The LLM-assisted analysis did not identify any additional themes or findings beyond those already identified through manual analysis. The lack of new findings further provided evidence that saturation had been reached after 12 interviews.

Our interview findings are centered on the common themes that emerged during the analysis of the interview transcripts. Our analysis revealed four main themes characterizing the experiences and perspectives of privacy engineers: (1) conceptualizing privacy engineering, (2) motivations to become privacy engineers, (3) common challenges faced in their roles, and (4) competencies and evaluation. These themes illuminate how privacy engineers understand their profession, what drives them to enter and remain in the profession, the obstacles they encounter, skills, and practices they rely on to fulfill their responsibilities. In the following subsections, we present each theme along with a selection of illustrative examples extracted from interview transcripts.

4.3 Conceptualizing Privacy Engineering

During interviews, participants were asked to provide definitions of privacy engineering and described various related professional roles. We found that two themes emerged among the responses to these portions of the interviews: the variability of the privacy engineering role and the importance of having strong technical and software knowledge.

Not strictly defined and highly variable. The majority of participants consistently emphasized that there is no universally agreed-upon definition of privacy engineering. For example, P4

Table 2: Summary of participants' professional context.

ID	'privacy engineer' job title?	'senior' or 'staff' job title?	'manager' or 'lead' job title?	Sector	Area	Years, privacy	Years, role	Org size	Team size	Associations	Certs.	Loc.
P1	●	○	○	I, N	T,E,B,C	3-4	1-3	1-1,000	2-7	N/A	N/A	WA
P2	●	●	○	I	T	9+	4-5	10,001-50,000	13+	IAPP	CIPT	N/A
P3	●	○	●	I	T	9+	1-3	100,001+	2-7	IAPP	CIPP/T/M, CISSP, CISM	WA
P4	○	○	●	I	T	5-8	6+	1,001-5,000	2-7	N/A	N/A	CA
P5	○	●	○	I	B	5-8	1-3	100,001+	8-12	IAPP	CIPT	NC
P6	●	○	○	I	T	9+	1-3	1-1,000	2-7	IAPP	CIPP/T/M	CA
P7	●	○	○	I	T,E,B,O	3-4	1-3	100,001+	N/A	N/A	OSCP, OSCE	CA
P8	●	○	○	I	T	5-8	4-5	5,001-10,000	8-12	IAPP, USENIX	N/A	CA
P9	○	●	○	I	T, E, B, C, O	5-8	4-5	1,001-5,000	2-7	IAPP, (ISC) ²	CIPP/T, Security+	MI
P10	●	○	○	I	T	3-4	1-3	100,001+	8-12	N/A	N/A	CA
P11	○	○	○	I, G	T, C	3-4	1-3	100,001+	8-12	IAPP	CEH	TX
P12	●	○	○	I	T	3-4	1-3	1-1,000	2-7	IAPP	CIPT	CA
P13	○	○	○	I, A	T	9+	6+	100,001+	N/A	IAPP, ISACA, (ISC) ²	CIPP, CISM	MD
P14	●	○	○	I	T	5-8	6+	100,001+	13+	N/A	N/A	CA
P15	●	●	○	I	T, O	3-4	1-3	5,001-10,000	2-7	IAPP	CIPT	NY
P16	○	○	○	I	T	5-8	6+	1,001-5,000	N/A	N/A	N/A	CA
P17	●	●	○	I, A	T, E, O	3-4	1-3	100,001+	13+	IAPP, USENIX	CIPP/T/M	WA
P18	●	●	○	I	T	5-8	1-3	100,001+	8-12	IAPP, ISACA, USENIX	CIPT/M	CA
P19	●	●	○	I	T	9+	6+	100,001+	13+	N/A	CIPT	CA
P20	○	○	○	I	T	5-8	6+	50,001-100,000	8-12	IEEE, ACM	N/A	CA
P21	●	○	○	I	T	3-4	4-5	100,001+	13+	IAPP	CIPT	CA
P22	●	○	○	I, N	T	9+	1-3	100,001+	13+	IAPP	CIPT	WA
P23	●	●	○	I	T, O	9+	1-3	100,001+	8-12	N/A	N/A	CA
P24	○	●	○	G	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
P25	○	○	○	I	T, C	3-4	6+	1-1,000	N/A	IAPP	N/A	WA
P26	○	○	●	I	T, C	9+	6+	1-1,000	N/A	N/A	N/A	CA
P27	●	○	○	I	T	5-8	4-5	100,001+	13+	IAPP	CIPP/T	CA

'ID' refers to participant ID; 'job title?' refers to phrases appearing (●) or not appearing (○) in a participant's job title; 'Sector' and 'Area' refer to participant's employment. No response or 'prefer not to disclose' option is abbreviated as 'N/A.' Values in 'Sector' column: 'I' = Industry, 'G' = Government, 'A' = Academia, 'N' = NGO. Values in 'Area' column: 'T' = Technology and Software; 'B' = Banking and Finance; 'C' = Consulting; 'E' = E-commerce; 'O' = Other (multiple choices allowed). Values in 'Loc.' column refer to U.S. states. For 'Associations' and 'Certs.' columns refer to Appendix B.

stated that “the role of a privacy engineer [...] is also very nebulous and has not been strictly defined.” Other participants reported that job titles and responsibilities vary widely between organizations, even when these roles share the label of “privacy engineer”. For example, some respondents described their work as deeply technical: writing code, designing systems, or conducting architectural reviews. Others considered their remit as more advisory, focused on communicating policy requirements and providing internal guidance. Several participants noted that this lack of standardization led to confusion both internally and externally, with colleagues sometimes misunderstanding the scope of their role or conflating it with more familiar functions like security engineering. Participants who mentioned this idea framed the ambiguity as a natural outcome of a relatively new field that continues to evolve, making role boundaries fluid and subject to ongoing negotiation.

Requires technical expertise. Across the interviews, almost all participants stressed that effective privacy engineering is based on a strong technical foundation: “I think it's definitely a technical role. I think it's someone who helps solve technical challenges related to privacy which can be wide-ranging. But like it's an engineer for a reason” [P9]. They highlighted the need to understand complex systems, data flows, and software architectures to identify privacy risks and integrate privacy-enhancing measures. Many participants recounted instances where their ability to navigate code, evaluate cryptographic tools, or design novel data minimization strategies allowed them to propose practical and implementable solutions. However, participants noted that while technical skill is critical, it must often be combined with broader knowledge of privacy, including familiarity with legal frameworks and organizational policies.

This combination, they suggested, is what distinguishes privacy engineers from purely legal or purely technical professionals.

4.4 Interest and Motivations

The participants described a variety of reasons for entering the field of privacy engineering. These motivations often reflected personal interests and values, as well as a desire for career growth and intellectual stimulation. Three themes stood out: viewing the field as novel and exciting, enjoying continuous learning and opportunities to educate others, and acting upon moral and ethical convictions related to user privacy.

Novel and exciting field. A prominent theme among almost all participants was the attraction of privacy engineering as a dynamic and emerging discipline. Participants described the role as constantly evolving, with frequent changes in technologies, regulations, and user expectations. This state of flux was seen as intellectually stimulating and drove individuals to stay up-to-date with emerging standards and tools. The participants remarked that privacy engineering never feels boring, as each new project presents unique technical puzzles: *“It’s always an interesting problem, because [...] it’s always changing. But then how each company deals with it is kind of unique.”* [P8]. The perceived newness of the field also gave the participants a sense of pioneering work, compared to working in security: *“I feel like cybersecurity in a lot of cases has been figured out. [...] Privacy is very much still being developed and understood.”* [P4].

Moral and ethical values. Several participants cited moral and ethical considerations as major drivers for pursuing careers in privacy engineering. They viewed privacy as a fundamental human right and expressed a sense of responsibility to protect user data. This framing with respect to moral virtues inspired participants to navigate organizational pressures that could deprioritize privacy considerations. Participants sometimes linked their commitment to privacy with broader ethical principles, noting that their desire to protect personal information was part of a larger personal ethos. Such motivations underscored a sense of purpose, setting their work apart from more purely commercial or technical undertakings.

4.5 Skills, Responsibilities, and Competencies

Finally, participants also discussed various competencies they believed were essential for privacy engineers. They also reflected on how their performance was assessed, both formally and informally. Four key themes emerged: the common reliance on informal evaluation methods, the need for cross-functional stakeholder management, the importance of risk and threat management expertise, and the centrality of conducting privacy reviews and offering informed advice.

Cross-functional stakeholder collaboration. Almost all participants emphasized the importance of working effectively with a wide range of stakeholders, including product teams, legal counsel, marketing departments, user researchers and external clients. They detailed how successful privacy engineering requires not only technical acumen, but also communication, negotiation, and diplomacy. These interactions required them to explain complex technical concepts to non-technical audiences, manage conflicting priorities,

and, at times, negotiate compromises that balance privacy with other organizational needs. Ensuring broad alignment and shared understanding was considered a critical competency that directly influenced their impact and the success of privacy initiatives. P1 said it plainly: *“A lot of my time is spent on stakeholder management. So it’s both understanding kind of what my team is doing, removing roadblocks, helping them. [...] So it’s, you know, a range of things. But I would say a lot of what I do is really influencing, understanding, and influencing.”*

Risk and threat management. Handling privacy threats and assessing potential risks emerged as central responsibilities. With only a few exceptions, all participants described their work as identifying vulnerabilities, evaluating data handling practices, and recommending mitigations before issues escalated: *“I focus on looking at perspective features and products and working with the engineering and product teams to understand what are the privacy risks. How can we mitigate them? And then verifying that they actually did mitigate the risks”* [P2]. Some used privacy threat modeling or “red teaming” exercises to anticipate potential problems. Their competence in risk management was closely related to their ability to translate abstract threats into concrete technical recommendations. Some participants recounted how demonstrating foresight and providing actionable solutions contributed significantly to their credibility within the organization.

Perform reviews and offer advice. With only a few exceptions, all participants mentioned that a significant part of their workload involved conducting privacy reviews, performing systems or products assessments, and offering customized advice. These activities often took the form of providing input on design documents, reviewing feature proposals, or helping product teams understand privacy implications. Participants reported that offering timely, accurate, and actionable guidance was one of their most visible contributions, serving as a tangible sign of their value as internal privacy experts. They emphasized the importance of clarity, accessibility, and practicality in their advice, ensuring that stakeholders could implement the recommended changes with minimal confusion.

Learning and teaching. Many respondents found personal fulfillment in the opportunity to learn and teach others about privacy. Some highlighted the satisfaction they derive from gaining new skills or deepening their understanding of privacy techniques, privacy law, and human-centric design considerations. Others described their role as educators within their organizations, helping development teams understand privacy requirements and coaching them in implementing controls effectively. Participants noted that as they gained expertise, they became key knowledge brokers, sharing insights and fostering awareness that went beyond the day-to-day engineering tasks.

4.6 Challenges and Strategies

Although participants expressed enthusiasm for their work, they also highlighted persistent challenges. These ranged from broader organizational and market forces to difficulties in translating complex regulations into actionable technical requirements. Three main themes emerged: misaligned incentives that undermine privacy

efforts, the complexity of translating laws and regulations into technical standards, and structural obstacles posed by reporting and organizational hierarchies.

Informal evaluation. When asked how their performance was evaluated, with only a few exceptions, participants described informal ad hoc evaluation processes. Rather than a standardized set of metrics, many said that their work was gauged through conversations, internal feedback, or their ability to respond effectively to emerging privacy risks. For example, P4 stated that *“So how a manager would look at [this work deliverable] is like: is this something that I completely agree with? Or how can we get it to a state that we can both agree on what this means?”* Some participants noted that their success was measured by the absence of privacy incidents, an outcome that can be difficult to attribute directly to their interventions. Others mentioned that peer recognition, project successes, and positive stakeholder feedback were the leading indicators of meeting organizational expectations: *“One of the things is impact. But also how peers perceive you”* [P6].

Misaligned incentives. A recurring challenge that the participants mentioned was the misalignment of incentives within their organizations. Many felt that privacy engineers are tasked with championing user data protection even when it may conflict with short-term business goals, such as quickly releasing new features or collecting more user data to enhance services: *“And how do you like deal with like the business struggles of trying to collect more versus protect? [...] It’s almost like a political struggle”* [P9]. Participants shared accounts of struggling to convince product managers or other stakeholders to invest time and resources in proactive privacy initiatives. Some participants described a reality where privacy gains visibility only after an incident or regulatory inquiry, thus incentivizing reactive rather than preventive measures. For example, P11 believed that other functions, such as security, are taken more seriously: *“The repercussions are questionable compared to if you have a security incident. It’s much more severe. The CSO can land up in prison. In that sense, I feel security is taken still more seriously than privacy.”*

Translating law and regulation. The majority of participants consistently indicated that navigating the legal and regulatory environment was a complex aspect of their work: *“I think one of [challenges] is like just the sheer fragmentation of all the different privacy laws and what’s involved. [...] Are you healthcare? Are you for profit? How many customers do you have? I need to know all of this”* [P8]. They often felt responsible for transforming high-level rules, such as those in GDPR or CCPA, into actionable technical requirements. This translation process was not straightforward. Many described lengthy interpretations, negotiations with legal teams, or trial-and-error attempts to incorporate policies into system designs. The shifting legal landscape exacerbated this challenge. Participants noted how new or evolving regulations required them to continuously adjust their interpretations and solutions, contributing to a sense of uncertainty.

Reporting structures. Interviewees highlighted how organizational positioning affected their ability to implement privacy measures. Some privacy engineers were embedded within security teams, while others reported to legal or product units. The

participants observed that the department they were housed in influenced the kind of support they received, as well as their perceived authority. For example, privacy engineers embedded in legal teams found it easier to leverage regulatory mandates but struggled when proposing technical changes. In contrast, those in engineering departments found it easier to influence system design, but more challenging to secure the approval of legal experts. Ultimately, these structural arrangements shaped their day-to-day interactions and determined how readily they could advocate for privacy initiatives.

Privacy championing. Interviewees consistently emphasized the importance of championing privacy within their organizations to address complex engineering challenges. Many privacy engineers reported that proactively increasing the visibility of privacy issues enabled them to secure critical support from senior leadership and cross-functional teams. For example, several participants recounted experiences in which they escalated privacy concerns directly to executive levels, effectively influencing policy revisions and garnering additional resources for privacy initiatives. This approach, which often involves regular communication of the risks and benefits associated with improved privacy measures, proved instrumental in changing organizational mindsets. By positioning themselves as trusted advocates for privacy, these professionals not only elevated the discussion around data protection, but also facilitated more agile responses to emerging regulatory and technical challenges. P2 saw this as especially important: *“Many different subcultures and niche groups have different views on [privacy]. [...] So, being aware of that and being able to advocate on their behalf. I think, as an expectation. That’s not directly in the job role, but it’s a core fundamental part of it.”*

5 Discussion

In this section, we reflect on our findings and situate them within the broader context of the evolving field of privacy engineering. We revisit the key themes identified in our analysis—conceptual ambiguity, motivational drivers, organizational and regulatory challenges, and essential competencies—and discuss their implications for practitioners, organizations, researchers, and educators. We conclude by describing practical recommendations, acknowledging limitations, and suggesting directions for future research.

5.1 Evolving Landscape of Privacy Engineering

Our study revealed that privacy engineering remains a nascent and fluid discipline, one whose boundaries and practices continue to evolve. Participants consistently noted that their roles are not strictly defined, reflecting a broader industry reality in which privacy engineers must adapt to varied organizational structures, team compositions, and product domains. This ambiguity underscores a critical need for more formalized frameworks, job descriptions, and skill set benchmarks that allow practitioners, recruiters, and organizational leaders to better understand what privacy engineers do, how they add value, and where they fit within existing hierarchies.

Compared to earlier conceptualizations of privacy engineering that focus on technical solutions or legal compliance alone, the experiences of our participants highlight a more holistic view. Privacy engineering emerges not as a narrow specialization, but rather as a

multidisciplinary pursuit. Privacy engineers blend technical expertise with strong communication and cross-functional coordination capabilities. As privacy concerns increase in scope and complexity, the ability to align technical solutions with regulatory mandates, stakeholder values, and organizational priorities becomes a key differentiator. This multifaceted identity suggests that privacy engineering is both a technical craft and a sociotechnical mediation role, one that will continue to gain prominence as data-driven business models evolve.

5.2 Motivations for Privacy Engineering

Our findings suggest that privacy engineers are often driven by more than professional ambition or technological interest. Many participants reported personal and ethical motivations to participate in privacy work, seeing their role as safeguarding user rights and promoting responsible data stewardship. This moral dimension distinguishes privacy engineers from more traditional engineering roles focused primarily on functionality, performance, or security. The emphasis of engineers on ethical considerations aligns with increasing public scrutiny and heightened expectations of responsible data use, trustworthiness, and fairness.

This ethical orientation can have several positive impacts. It may help organizations anticipate reputational risks and improve brand trust by embedding privacy values early in the development process. At the same time, this dimension calls for careful organizational support. Privacy engineers who view their work as ethically significant may become disheartened if faced with persistent misaligned incentives. Organizations must therefore create an environment that recognizes and leverages the moral commitments of privacy engineers through policies, reward structures, and leadership support to ensure that these motivations translate into sustainable privacy outcomes.

5.3 Organizational and Regulatory Challenges

Participants highlighted persistent organizational barriers, including the difficulty in securing the support of product teams or executives focusing primarily on revenue and rapid deployment. Misaligned incentives remain a core obstacle. While regulators and public sentiment increasingly demand proactive privacy measures, internal metrics often do not recognize or reward preventive privacy work. Addressing these challenges may require organizations to adopt new performance indicators, reconfigure reporting lines, or invest in privacy education to ensure that privacy engineering is not perceived as a simple compliance cost.

Regulatory complexity emerged as another key challenge, with privacy engineers struggling to translate evolving high-level legal requirements into concrete technical practices. This issue points to the potential for greater collaboration between legal and engineering teams and the development of clear guidance and standardized industry frameworks to operationalize legal mandates. Over time, the field could benefit from consolidated guidelines that help privacy engineers bridge the gap between legal text and system design, potentially reducing the uncertainty and resource expenditures involved in continuous regulatory interpretation.

5.4 Competencies and Informal Evaluations

Our analysis showed that privacy engineers rely on a mix of technical skills, cross-functional communication abilities, and risk management competencies. Although technical proficiency is a cornerstone, especially in identifying and integrating privacy-enhancing technologies, the softer, but equally crucial skill of stakeholder management emerged as vital. Privacy engineers must negotiate among multiple parties; developers, product owners, legal experts, and marketing teams to achieve privacy goals. This finding strengthens the case for conceptualizing privacy engineering as a role that thrives at the intersection of technology and organizational culture.

However, participants also noted that their performance is often not measured by formal metrics, instead relying on informal feedback, trust, and the absence of incidents. For organizations to foster more systematic improvement, clearer evaluation frameworks could be established. These could include measures related to the integration of privacy controls in product roadmaps, user satisfaction with data handling practices, or the number and severity of privacy-related incidents avoided. More robust assessment methods could provide privacy engineers with recognition and career advancement pathways, reinforcing their strategic value to the organization.

5.5 Implications

The insights of this study offer several practical recommendations for different stakeholder groups:

Industry and Practice. Organizations may benefit from developing standardized role descriptions and career ladders for privacy engineers, clarifying expectations and needed skills. Doing so can facilitate recruitment and help new hires integrate more smoothly. Second, incorporating privacy-focused key performance indicators into product development cycles can align incentives, ensuring that proactive privacy measures are recognized and rewarded. Third, regular training sessions for both engineers and nontechnical stakeholders can enhance organizational privacy literacy, improving communication and collaboration.

Another implication is the potential role of privacy engineers as internal advocates and educators. Organizations may formalize this function, encouraging privacy engineers to hold periodic workshops or office hours for developers, product managers, and designers. Such activities could help embed a culture of privacy by design and reinforce compliance as a shared responsibility, rather than a niche concern.

Aspiring Privacy Engineers. For aspiring privacy engineers, our findings illuminate a path that requires a versatile skill set and a proactive mindset. Given the field's evolving nature and the lack of strict role definitions, individuals should focus on developing strong technical fundamentals and interpersonal skills to bridge the gap between technical teams and legal or product counterparts. Aspiring professionals should also have an understanding of legal principles and the ability to translate them into technical requirements.

Educators. Educators designing privacy engineering curricula should take note of the diverse competencies our study identifies as critical for practicing privacy engineers. The findings underscore

the need for programs that extend beyond purely technical instruction. The significant emphasis on cross-functional collaboration and communication implies that educational programs must incorporate training in stakeholder management, negotiation, and ethical reasoning to prepare students for the sociotechnical realities of the role. Moreover, given the evolving nature of the profession and the lack of standardization, curricula should encourage adaptability, continuous learning, and the ability to critically assess and apply emerging privacy-enhancing technologies and privacy engineering methodologies. Finally, certification organizations such as IAPP and ISACA should consider our findings when performing future research on the professional community. Notably, we observe that the definition that has emerged from the consensus of our participants differs substantially from the definitions used by these organizations [36, 37]. We argue that the definition of privacy engineering should include the core competencies, skill set, methods employed, daily practices, and responsibilities we have uncovered in our research.

6 Limitations and Future Work

While our study provides valuable insights into the nascent field of privacy engineering, we acknowledge several limitations. It is important to reiterate that this study employs a qualitative approach that does not produce statistical findings or quantitative measurements of the frequency of any theme. Instead, our goal is to explore and conceptually define the profession of privacy engineering through an in-depth analysis of expert perspectives. As such, our findings should not be interpreted as statistically representative of the entire privacy engineering field, but rather as a rich, exploratory synthesis of the insights shared by our participants. Although we identify recurring themes and commonalities, we do not claim that these findings can be fully generalized to all privacy engineers. Instead, this study provides a foundation for understanding the evolving nature of the profession, offering insights that may inform further research and discussion within the field. Our study is both qualitative and exploratory, and while our participant pool was larger than that recommended by some qualitative guidelines, it remains subject to self-selection biases. Those who chose to participate may be particularly motivated or established privacy engineers. Furthermore, because recruitment relied on convenience and snowball sampling, participation may be skewed toward individuals linked to the authors' networks and to well-connected practitioners in large US technology firms, potentially under-representing privacy engineers in smaller organizations or non-tech sectors. In addition, because our IRB approval covered only U.S. participants, the study is U.S.-centric and most of the participants were based in California, a major technology hub. Privacy engineering roles and perceptions may differ in other cultural or regulatory environments. Furthermore, the dynamic nature of privacy engineering means that our findings offer a snapshot in time. Roles, definitions, and practices can change as the field matures, as regulatory regimes change, or as new technologies emerge.

Our study lays the foundation for further research. Although we identified common themes and challenges, future research could quantify their prevalence across different industry verticals, organizational sizes, and cultural contexts. Surveys and larger-scale

quantitative studies could complement our qualitative findings, enabling more generalizable conclusions. Longitudinal studies might also observe how the role of privacy engineers evolves as regulations, technologies, and societal norms continue to shift. In addition, future work could apply the same self-identification-based methodology in a multi-national context to assess whether themes identified in this study hold across different cultural, regulatory, and organizational environments.

Another avenue of research involves exploring the interplay between privacy engineering and emerging areas such as machine learning fairness, explainability, and data governance. As organizations increasingly rely on data-intensive processes, privacy engineers might collaborate with professionals who address algorithmic transparency or responsible AI development. Studying these intersections could inform a more holistic approach to data ethics, trustworthiness, and accountability.

Finally, more nuanced regulatory and policy-oriented research could examine whether and how policymakers could rely on insights from privacy engineers to refine legal frameworks and provide clearer and more practical guidance. Such a feedback loop could help harmonize legal requirements with engineering realities, ultimately producing more effective privacy protections for users.

7 Conclusion

This research provides a novel in-depth look at the realities of privacy engineering roles, uncovering the complexities and nuances that shape this emerging field. Far from being a narrowly defined technical job, privacy engineering involves blending robust technical foundations with a flexible understanding of legal and ethical considerations, organizational cultures, and collaborative communication skills. Privacy engineers act as linchpins, translating high-level privacy principles into concrete, system-level implementations, and ensuring that user rights and trust are respected in an increasingly data-driven world.

As privacy considerations move to the forefront of regulatory and consumer attention, the importance of well-defined, and well-resourced privacy engineering functions will only grow. By illuminating how privacy engineers understand their roles, what motivates them, and what challenges they face, this study aims to help organizations, policymakers, and educators recognize the critical value of privacy engineering. With greater clarity, structured support, and meaningful incentives, privacy engineers can more effectively guide the development of responsible, compliant, and user-centric digital products, ultimately shaping a more trustworthy and privacy-aware technological landscape.

Acknowledgments

This work was supported by the U.S. National Science Foundation under grant CCF-2217771. We especially thank Serge Egelman for his support and advice, R. Jason Cronk, Mohammad Tahaei, and Mira Olson for their feedback on the survey and interview protocol, Amie Stepanovich and Shea Swauger for their review of the paper and our anonymization technique. The views expressed in this paper are solely those of the authors and do not necessarily reflect the official policy or position of their affiliated institutions.

References

- [1] Yaqoob Al-Slais. 2020. Privacy engineering methodologies: A survey. In *2020 International Conference on Innovation and Intelligence for Informatics, Computing and Technologies (3ICT)*. IEEE, 1–6.
- [2] A Apacible-Bernardo and K Bushey. 2025. Data protection and privacy laws now in effect in 144 countries. IAPP, 28 January 2025. Available at: <https://iapp.org/news/a/data-protection-and-privacy-laws-now-in-effect-in-144-countries/>. Accessed: 4 February 2025.
- [3] Blake E Ashforth and Fred Mael. 1989. Social Identity Theory and the Organization. *The Academy of Management Review* 14, 1 (1989), 20–39. <http://www.jstor.org/stable/258189>
- [4] Kenneth A Bamberger and Deirdre K Mulligan. 2010. Privacy on the Books and on the Ground. *Stanford Law Review* 63 (2010), 247.
- [5] Kathrin Bednar, Sarah Spiekermann, and Marc Langheinrich. 2019. Engineering Privacy by Design: Are engineers ready to live up to the challenge? *The Information Society* 35, 3 (2019), 122–142.
- [6] Patrick Biernacki and Dan Waldorf. 1981. Snowball sampling: Problems and techniques of chain referral sampling. *Sociological Methods & Research* 10, 2 (1981), 141–163.
- [7] Travis D Breaux. 2020. *An Introduction to Privacy for Technology Professionals*. International Association of Privacy Professionals.
- [8] Travis D Breaux and Annie Antón. 2008. Analyzing regulatory rules for privacy and security requirements. *IEEE Transactions on Software Engineering* 34, 1 (2008), 5–20.
- [9] Travis D Breaux, Annie I Antón, Kent Boucher, and Merlin Dorfman. 2008. Legal requirements, compliance and practice: an industry case study in accessibility. In *16th IEEE International Requirements Engineering Conference*. IEEE, 43–52.
- [10] Travis D Breaux, Matthew W Vail, and Annie I Anton. 2006. Towards regulatory compliance: Extracting rights and obligations to align requirements with regulations. In *14th IEEE International Requirements Engineering Conference*. IEEE, 49–58.
- [11] Sean Brooks, Michael Garcia, Naomi Lefkowitz, Suzanne Lightman, and Ellen Nadeau. 2017. NISTIR 8062: An Introduction to Privacy Engineering and Risk Management in Federal Systems. U.S. Department of Commerce, National Institute of Standards and Technology, Information Technology Laboratory. Available at: <https://doi.org/10.6028/NIST.IR.8062>.
- [12] Ann Cavoukian. 2010. Privacy by design: the definitive workshop. A foreword by Ann Cavoukian, Ph. D. *Identity in the Information Society* 3, 2 (2010), 247–251.
- [13] Ann Cavoukian. 2012. Operationalizing privacy by design: a guide to implementing. *Commun. ACM* 55, 9 (2012), 7.
- [14] Ann Cavoukian et al. 2009. Privacy by design: The 7 foundational principles. *Information and privacy commissioner of Ontario, Canada* 5, 2009 (2009), 12.
- [15] Aaron Cross and Andrew Simpson. 2018. Rethinking the proposition of privacy engineering. In *Proceedings of the New Security Paradigms Workshop*. 89–102.
- [16] Kathy Charmaz. 2006. *Constructing grounded theory: A practical guide through qualitative analysis*. Sage.
- [17] Cisco. 2024. Cisco 2024 Consumer Privacy Survey. Cisco. Available at: <https://www.cisco.com/c/en/us/about/trust-center/consumer-privacy-survey.html>.
- [18] Data Protection Commission. 2023. *Decision in the matter of Meta Platforms Ireland Limited (previously known as Facebook Ireland Limited)*. Technical Report DPC Inquiry Reference: IN-20-8-1. Data Protection Commission, 21 Fitzwilliam Square South, Dublin 2, Ireland. https://www.edpb.europa.eu/system/files/2023-05/final_for_issue_ov_transfers_decision_12-05-23.pdf Decision made by Commissioner Helen Dixon on 12 May 2023 pursuant to Section 111 of the Data Protection Act, 2018 and Articles 60 and 65 of the General Data Protection Regulation.
- [19] Lorrie Faith Cranor and Norman Sadeh. 2013. Privacy engineering emerges as a hot new career. *IEEE Potentials* 32, 6 (2013), 7–9.
- [20] Lorrie Faith Cranor and Norman Sadeh. 2013. A shortage of privacy engineers. *IEEE Security & Privacy* 11, 2 (2013), 77–79.
- [21] George Danezis, Josep Domingo-Ferrer, Marit Hansen, Jaap-Henk Hoepman, Daniel Le Métayer, Rodica Tirtza, and Stefan Schiffner. 2015. Privacy and Data Protection by Design – from policy to engineering. ENISA Reports, ENISA (European Union Agency for Cybersecurity), January 12, 2015. Available at: <https://www.enisa.europa.eu/publications/privacy-and-data-protection-by-design>.
- [22] Andreas Dengel, Rupert Gehrlein, David Fernes, Sebastian Görlich, Jonas Maurer, Hai Hoang Pham, Gabriel Großmann, and Niklas Dietrich genannt Eisermann. 2023. Qualitative Research Methods for Large Language Models: Conducting Semi-Structured Interviews with ChatGPT and BARD on Computer Science Education. *Informatics* 10, 4 (2023). <https://doi.org/10.3390/informatics10040078>
- [23] Anirudh Ekambaranathan, Jun Zhao, and Max Van Kleek. 2021. “Money makes the world go around”: Identifying Barriers to Better Privacy in Children’s Apps From Developers’ Perspectives. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [24] Pardis Emami-Naeini, Janarth Dheenadhayalan, Yuvraj Agarwal, and Lorrie Faith Cranor. 2023. Are Consumers Willing to Pay for Security and Privacy of IoT Devices?. In *32nd USENIX Security Symposium*. 1505–1522.
- [25] Joan Feigenbaum, Michael J Freedman, Tomas Sander, and Adam Shostack. 2001. Privacy engineering for digital rights management systems. In *ACM Workshop on Digital Rights Management*. Springer, 76–105.
- [26] Don Gotterbarn, Keith Miller, and Simon Rogerson. 1997. Software engineering code of ethics. *Commun. ACM* 40, 11 (Nov. 1997), 110–118. <https://doi.org/10.1145/265684.265699>
- [27] Greg Guest, Arwen Bunce, and Laura Johnson. 2006. How many interviews are enough? An experiment with data saturation and variability. *Field Methods* 18, 1 (2006), 59–82.
- [28] Seda Gürses and Jose M Del Alamo. 2016. Privacy engineering: Shaping an emerging field of research and practice. *IEEE Security & Privacy* 14, 2 (2016), 40–46.
- [29] Seda Gürses, Carmela Troncoso, and Claudia Diaz. 2011. Engineering privacy by design. *Computers, Privacy & Data Protection* 14, 3 (2011), 25.
- [30] Irit Hadar, Tomer Hasson, Oshrat Ayalon, Eran Toch, Michael Birnhack, Sofia Sherman, and Arod Balissa. 2018. Privacy by designers: software developers’ privacy mindset. *Empirical Software Engineering* 23 (2018), 259–289.
- [31] Isobel Asher Hamilton and Grace Dean. 2021. Signal downloads skyrocketed 4,200% after WhatsApp announced it would force users to share personal data with Facebook. It’s top of both Google and Apple’s app stores. Business Insider, Jan 13, 2021. Available at: <https://www.businessinsider.com/whatsapp-facebook-data-signal-download-telegram-encrypted-messaging-2021-1>.
- [32] Jay John Hellman. 1970. *Privacy and information systems: An argument and an implementation*. RAND Corporation.
- [33] Guntur Budi Herwanto, Fajar J Ekaputra, Gerald Quirchmayr, and A Min Tjoa. 2024. Towards a Holistic Privacy Requirements Engineering Process: Insights from a Systematic Literature Review. *IEEE Access* (2024).
- [34] Stefan Albert Horstmann, Samuel Domiks, Marco Gutfleisch, Mindy Tran, Yasemin Acar, Veelasha Moonsamy, and Alena Naiakshina. 2024. “Those things are written by lawyers, and programmers are reading that.” Mapping the Communication Gap Between Software Developers and Privacy Experts. *Proceedings on Privacy Enhancing Technologies* (2024).
- [35] Dominik Huth and Florian Matthes. 2019. “Appropriate Technical and Organizational Measures”: Identifying Privacy Engineering Approaches to Meet GDPR Requirements. In *AMCIS 2019 Proceedings*. 5. https://aiselaisnet.org/amcis2019/info_security_privacy/info_security_privacy/5
- [36] International Association of Privacy Professionals (IAPP). 2023. Defining Privacy Engineering. https://iapp.org/media/pdf/resource_center/defining_privacy_engineering_infographic.pdf. Accessed: 2024-05-26.
- [37] ISACA. 2025. Privacy Engineering. <https://www.isaca.org/career-center/career-journey/privacy-engineering>. Accessed: 2024-05-26.
- [38] Leonardo Horn Iwaya, Muhammad Ali Babar, and Awais Rashid. 2023. Privacy engineering in the wild: Understanding the practitioners’ mindset, organizational aspects, and current practices. *IEEE Transactions on Software Engineering* 49, 9 (2023), 4324–4348.
- [39] Steve Kenny and John Borking. 2002. The Value of Privacy Engineering. *The Journal of Information, Law and Technology (JILT)* 1 (2002), 02–1.
- [40] Zachary Kilhoffer, Devyn Wilder, and Masooda Bashir. 2024. Compliance as Baseline, or Striving for More? How Privacy Engineers Work and Use Privacy Standards. In *IEEE European Symposium on Security and Privacy Workshops*. IEEE, 9–18.
- [41] Lea Kissner and Lorrie Faith Cranor. 2021. Privacy engineering superheroes. *Commun. ACM* 64, 11 (2021), 23–25.
- [42] Blagovesta Kostova, Seda Gürses, and Carmela Troncoso. 2020. Privacy Engineering Meets Software Engineering: On the Challenges of Engineering Privacy by Design. (2020). arXiv:2007.08613 [cs.SE] <https://arxiv.org/abs/2007.08613>
- [43] Inga Kroener and David Wright. 2014. A strategy for operationalizing privacy by design. *The Information Society* 30, 5 (2014), 355–365.
- [44] Yod Samuel Martin Garcia and José María del Álamo Ramiro. 2017. A metamodel for privacy engineering methods. In *CEUR Workshop Proceedings*.
- [45] Colleen McClain, Michelle Faverio, Monica Anderson, and Eugenie Park. 2023. How Americans view data privacy. *Pew Research Center* 18 (2023).
- [46] Deirdre K Mulligan, Colin Koopman, and Nick Doty. 2016. Privacy is an essentially contested concept: a multi-dimensional analytic for mapping privacy. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 374, 2083 (2016), 20160118.
- [47] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Washington Law Review* 79 (2004), 119.
- [48] Nicolas Notario, Alberto Crespo, Yod-Samuel Martin, Jose M Del Alamo, Daniel Le Metayer, Thibaud Antignac, Antonio Kung, Inga Kroener, and David Wright. 2015. PRIPARE: Integrating Privacy Best Practices into a Privacy Engineering Methodology. In *IEEE Security & Privacy Workshops*. 151–158. <https://doi.org/10.1109/SPW.2015.22>
- [49] Stefano De Paoli. 2024. Performing an Inductive Thematic Analysis of Semi-Structured Interviews With a Large Language Model: An Exploration and Provocation on the Limits of the Approach. *Social Science Computer Review* 42, 4 (2024), 997–1019. <https://doi.org/10.1177/08944393231220483>

- [50] European Parliament and Council. 1995. 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. Official Journal of the European Communities, L281, 31 October, pp. 1–15. Available at: <https://eur-lex.europa.eu/eli/dir/1995/46/oj/eng>. Accessed: February 3, 2025.
- [51] F Paul Pittman. 2023. US Data Privacy Guide. *White & Case LLP* 15 (2023).
- [52] Ira S Rubinstein and Nathaniel Good. 2013. Privacy by design: A counterfactual analysis of Google and Facebook privacy incidents. *Berkeley Technology Law Journal* 28 (2013), 1333.
- [53] Awanthika Senarath and Nalin A.G. Arachchilage. 2018. Why developers cannot embed privacy into software systems? An empirical investigation. In *Proceedings of the 22nd International Conference on Evaluation and Assessment in Software Engineering*. 211–216.
- [54] Awanthika Senarath, Marthie Grobler, and Nalin Asanka Gamagedara Arachchilage. 2019. Will they use it or not? Investigating software developers' intention to follow privacy engineering methodologies. *ACM Transactions on Privacy and Security* 22, 4 (2019), 1–30.
- [55] Stuart S Shapiro. 2010. Privacy by design: moving from art to practice. *Commun. ACM* 53, 6 (2010), 27–29.
- [56] Sarah Spiekermann. 2012. The challenges of privacy by design. *Commun. ACM* 55, 7 (2012), 38–40.
- [57] Sarah Spiekermann and Lorrie Faith Cranor. 2008. Engineering privacy. *IEEE Transactions on Software Engineering* 35, 1 (2008), 67–82.
- [58] Sarah Spiekermann, Jana Korunovska, and Marc Langheinrich. 2018. Inside the organization: Why privacy and security engineering is a challenge for engineers. *Proc. IEEE* 107, 3 (2018), 600–615.
- [59] Anselm Strauss and Juliet Corbin. 2008. *Basics of Qualitative Research* (3 ed.). SAGE Publications Inc.
- [60] Mohammad Tahaei, Alisa Frik, and Kami Vaniea. 2021. Privacy champions in software teams: Understanding their motivations, strategies, and challenges. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [61] Mohammad Tahaei, Kami Vaniea, and Awais Rashid. 2022. Embedding privacy into design through software developers: Challenges and solutions. *IEEE Security & Privacy* 21, 1 (2022), 49–57.
- [62] Robert H Tai, Lillian R Bentley, Xin Xia, Jason M Sitt, Sarah C Fankhauser, Ana M Chicas-Mosier, and Barnas G. Monteith. 2024. An Examination of the Use of Large Language Models to Aid Analysis of Textual Data. *International Journal of Qualitative Methods* 23 (2024). <https://doi.org/10.1177/16094069241231168>
- [63] Henri Tajfel and John C Turner. 2001. An integrative theory of intergroup conflict. In *Intergroup Relations: Essential Readings*, Michael A Hogg and Dominic Abrams (Eds.). Psychology Press, New York, 94–109.
- [64] Janice Y Tsai, Serge Egelman, Lorrie Faith Cranor, and Alessandro Acquisti. 2011. The effect of online privacy information on purchasing behavior: An experimental study. *Information Systems Research* 22, 2 (2011), 254–268.
- [65] Swaathi Vetrivel, Brennen Bouwmeester, Michel van Eeten, and Carlos H Gañán. 2024. IoT Market Dynamics: An Analysis of Device Sales, Security and Privacy Signals, and their Interactions. In *33rd USENIX Security Symposium*. 7031–7048.
- [66] Samuel Wairimu, Leonardo Horn Iwaya, Lothar Fritsch, and Stefan Lindskog. 2024. On the evaluation of privacy impact assessment and privacy risk assessment methodologies: A systematic literature review. *IEEE Access* 12 (2024), 19625–19650.
- [67] Joseph Williams and Lisa Nee. 2022. Privacy engineering. *Computer* 55, 10 (2022), 113–118.

A Consent Form

[Each participant viewed and consented to the following form before starting the screening survey and beginning the interview.]

Purpose. The study aims to delve into the roles and experiences of individuals working in a privacy engineering role. This research will provide insights into their responsibilities, skills, experience prerequisites, organizational hierarchies, and the challenges they encounter.

Procedure. If you agree to participate in this interview, you will be asked to answer a series of questions related to the purpose of the study. The entire interview will take no more than 90 minutes.

Eligibility. Participation will be restricted to respondents in the U.S., 18 years of age or older, and speakers of English.

Data Confidentiality. As part of this interview, we will record audio for transcription. However, we will delete raw audio data as soon as we finish transcribing it, no later than the anticipated study completion date. Furthermore, data will be kept confidential to the degree permitted by the technology being used, although the confidentiality of data transmitted over the Internet cannot be completely guaranteed. De-identified data will be retained indefinitely for possible use in future research done by ourselves or others.

Benefits and Risks to the Participant. Participants will obtain the experience of participating in a research study and will help contribute to the body of knowledge about online security, digital safety, and privacy. There are minimal risks associated with participating in this study. The results from the research will be shared with the participants at the conclusion of the research.

Voluntary Nature of the Study. Your participation is voluntary. You may stop participating at any time or withdraw from the study by letting us know. Partial data will not be analyzed or retained.

B Survey Questionnaire

[After the participant read the participant information sheet and consent form and agreed to participate in the study.]

- (1) Please choose all the sectors that closely align with your current employment area (select all that apply).
 - Industry/Business • Academia/Education • Government
 - Non-profit • Other: (free text)
- (2) What is your current employment status? Select all that apply.
 - Full-time employee (or contractor equivalent) • Part-time employee (or contractor equivalent) • Freelance/Consultant
 - Furloughed (temporarily laid off) or on leave • Unemployed
 - Student • Retired
- (3) If employed, what is your job title? (Free text)
- (4) Do you identify as a person who works in privacy engineering irrespective of job title?
 - Yes • No
- (5) If you'd like to participate in the Zoom-hosted interview, what email address should we use to contact you? (Free text)
- (6) What area do you currently work in? Select all that apply. (Optional)
 - Technology and Software • E-commerce • Healthcare

- Banking and Finance • Telecommunications • Government and Public Sector • Education • Retail • Media and Entertainment • Automotive • Real Estate • Utilities • Travel and Hospitality • Consulting • Other (free text)
- (7) How many years of experience do you have in your current role? (Free text, optional)
- (8) How many years of experience do you have working in privacy? (Free text, optional)
- (9) How many individuals do you work with directly (e.g. on your team)? (Free text, optional)
- (10) How many employees work in your organization? (Optional)
 - 1-1,000 • 1,001-5,000 • 5,001-10,000 • 10,001-50,000 • 50,001-100,000 • 100,001 and above
- (11) Where are you currently based in the US? Please indicate the city and state. (Free text, optional)
- (12) What is your age?
 - Age: (free text) • Prefer not to disclose
- (13) What is the highest educational degree you have obtained by now?
 - High school or less • Some college/university study without earning a degree • Associate degree (A.A., A.S., etc.) • Bachelor's degree (B.A., B.S., B.Eng., etc.) • Master's degree (M.A., M.S., M.Eng., MBA, etc.) • Professional degree (JD, MD, Ph.D, Ed.D, etc.) • Other (free text) • Prefer not to disclose
- (14) What is your gender?
 - Man • Woman • Non-binary, genderqueer, or gender non-conforming • Prefer to self-describe: (free text) • Prefer not to disclose
- (15) What is your individual gross annual income (total compensation)?
 - <\$50,000 • \$50,001-75,000 • \$75,001-100,000 • \$100,001-150,000 • \$150,001-200,000 • \$200,001-300,000 • \$300,001-500,000 • >\$500,001 • Prefer not to disclose
- (16) What is your marital status?
 - Single • In a relationship not recognized by law • In a domestic partnership • Married or in a civil union • Separated • Divorced • Widowed • Prefer not to disclose
- (17) Do you identify with any of the following identities or communities: racial or ethnic minority, LGBTQ+, person with a disability, indigenous or native peoples, immigrant or refugee, religious minority, low socio-economic status or background?
 - Yes • No • Prefer not to disclose
- (18) Do you belong to any of the following industry associations? Select all that apply. (Optional)
 - IAPP (International Association of Privacy Professionals)
 - IEEE (Institute of Electrical and Electronics Engineers)
 - ISACA (Information Systems Audit and Control Association) • (ISC)² (International Information System Security Certification Consortium) • ACM (Association for Computing Machinery) • CIPL (Centre for Information Policy Leadership) • EPIC (Electronic Privacy Information Center)
 - USENIX (The Advanced Computing Systems Association)
- (19) Do you have any privacy- or cybersecurity-related certifications? Select all that apply. (Optional)
 - CIPP (Certified Information Privacy Professional) - IAPP
 - CIPT (Certified Information Privacy Technologist) - IAPP

- CIPM (Certified Information Privacy Manager) - IAPP
- CISSP (Certified Information Systems Security Professional) - (ISC)² • CISM (Certified Information Security Manager) - ISACA • CRISC (Certified in Risk and Information Systems Control) - ISACA • CDPSE (Certified Data Privacy Solutions Engineer) - ISACA • CEH (Certified Ethical Hacker) - EC-Council • Others (free text)

C Interview Guide

[After the interviewer has introduced themselves and obtained verbal consent.]

Topic 1 - Introduction

- (1) Can you tell me briefly about what you do in your job?
- (2) Could you also define the term “privacy” as you normally use it in your work context?
- (3) How would you describe the roles in the industry related to privacy engineering?
- (4) How would you define a “privacy engineer”?

Topic 2 - Motivation

- (1) How did you become interested in privacy engineering as a career (or a function of your career)?
- (2) Could you share your career journey and how you arrived at your current position?
- (3) What motivates you to continue pursuing privacy engineering as part of your profession?
 - What are some personal goals you have for this work?
 - What value do you get from it?
 - What do you enjoy about it?
- (4) Now I am going to ask you a question about the future. A year from now, do you see yourself in the same position? More specifically, doing what it is that you currently do in your position.
 - Can you tell me more about why you answered this way?

Topic 3 - Responsibilities and Skills

- (1) Could you give me an idea of what a typical day at work looks like for you?
- (2) What responsibilities does your employer expect you to take on at work?
- (3) Why do you think there is such a [difference/similarity] between the expectation and the reality?
- (4) Are there any additional responsibilities you feel you are expected to take on in your role, such as to society, others in the organization, or even yourself? For instance, serving your broader community or other privacy professionals, mentoring others, volunteering your time, and so on.
- (5) What skills were demanded of you when you started your current role?
- (6) What are the skills you currently use in your job?
- (7) Is there a difference or not between the skills you were expected to demonstrate during the interviewing process and those required of you in your role?

Topic 4 - Reporting and Deliverables

- (1) Who do you report to?
- (2) Does anyone report to you?

- (3) What are the typical reporting structures that you see in your profession?
 - What are the teams and their composition of reporting?
 - What methods do you use to report to others (e.g., meetings, emails, project management platforms)?
 - What is the actual organizational structure (e.g., flat vs hierarchical)?
- (4) What deliverables are required from you in your role? For example, do you write code, research reports, Privacy-by-Design (PbD) advice, etc?
 - Can you tell me more about why these deliverables are important in your role?
 - Do you think these deliverables are typical or not typical for someone in your profession?
- (5) How are those deliverables evaluated by your manager?

Topic 5 - Challenges and Strategies

- (1) Are there any tools, techniques, or standards that create challenges for you?
 - What are the most common challenges that you encounter?
 - Do you think these challenges are typical or not typical for your profession?
- (2) Are there any challenges related to your organizational or reporting structures that you face?
 - What are the most common challenges that you encounter?
 - Do you think these challenges are typical or not typical for your profession?
- (3) Can you tell me more about the strategies that you use to overcome the challenges you mentioned?
 - Which ones do you find the most effective? Why? How do you know it's effective?
 - Which ones do you find the least effective? Why? How do you know it's ineffective?

Topic 6 - Success Metrics

- (1) How would you define 'success' in the work that you do?
- (2) What do you think the overarching goal is?
- (3) How do you think others evaluate the impact of your work?
 - Do you think there are any metrics associated with these evaluation criteria?