# Between Promise and Practice: Challenges and Misperceptions of Applying Privacy Enhancing Technologies in Business Contexts

Johannes Lohmöller, Hajeong Jeon, Jael Hentschel, Klaus Wehrle, Jan Pennekamp
RWTH Aachen University
{lohmoeller,jeon,hentschel,wehrle,pennekamp}@comsys.rwth-aachen.de

## Abstract

*Applying privacy-enhancing technologies (PETs), such as homomorphic encryption or differential privacy, promises to improve organizational cybersecurity strategies. However, in business contexts, significant gaps manifest between their technical capabilities and organizational perceptions, indicating a mismatch between promise and practice. This paper presents the first comprehensive meta-analysis of organizational PET perceptions through a systematic review of 34 empirical studies. Our findings reveal that while regulatory pressures and reputational considerations drive adoption, organizations face substantial practical challenges, including complexity management and insufficient understanding of technological capabilities. Even experienced practitioners show misperceptions about PET functionality, leading to misconfigurations that undermine promised privacy benefits. Thus, misperceptions directly impact cybersecurity effectiveness, as organizations may overestimate deployed protections or underutilize available capabilities. Consequently, our analysis highlights the need for and recommends implementing improved education, regular re-assessments of current beliefs regarding PETs, and transparency mechanisms to translate potential into successful enterprise cybersecurity.*

**Keywords: Privacy-Enhancing Technologies, security awareness, privacy literacy, organizational security**

## 1. Introduction

Organizational cybersecurity increasingly relies on advanced technological mechanisms for protecting business, customer, and personal data across its entire lifecycle. Well-established Privacy-Enhancing Technologies (PETs) like homomorphic encryption, differential privacy, secure multiparty computation, and trusted execution environments, are examples for such mechanisms, each contributing specific capabilities for privacy-preserving data analysis and processing. Indeed, the integration of PETs within corporate environments has become increasingly pertinent in emerging data sovereignty initiatives and data ecosystems such as Gaia-X (Braud et al., 2021), and the prevalence of recent data breaches (Schlackl et al., 2022). Despite these technological advances and their substantial potential, perceptions of these mechanisms within organizations vary considerably, impacting their practical deployment (Agrawal et al., 2021).

In this study, we conceptualize organizational PETs perception as the collective understanding and beliefs that stakeholders within an organization hold regarding their capabilities, limitations, and practical implications. We find organizational perceptions shaped by various factors including prior experience, organizational culture, available expertise, and external influences such as academic communications and regulatory guidance. Importantly, these perceptions may diverge significantly from objective technical realities, creating gaps that influence adoption decisions, implementation strategies, and ultimately, the success of cybersecurity initiatives.

Previous research indeed indicates mismatches between the actual technical capabilities of PETs and organizational stakeholders' perceptions of the privacy guarantees these technologies offer. For instance, organizations reportedly struggle with configuration aspects, e.g., parameter selection, which are integral to the effective deployment of PETs (Dwork et al., 2019). Furthermore, research (Prince et al., 2023) highlights a notable gap in privacy literacy among decision-makers and technical experts, potentially leading to misconceptions regarding the strengths and limitations of these technologies. While prior work has already started documenting both technical and organizational challenges encountered by businesses attempting to adopt PETs (Lohmöller et al., 2024), these efforts lack comprehensive understandings of the motivations underlying PET adoption, encountered challenges, and comparisons across technologies. Besides, questions remain whether stakeholders comprehend the technology's limitations and benefits, and whether the technologies align with their initial expectations. Currently, this knowledge gap hinders realizing the full potential of PETs within business contexts.

Therefore, this paper systematically assesses organizational perceptions of PETs through a structured meta-analysis of prior studies that report on businesses' motivations and experiences. We address two primary research aspects: ① *What motivates businesses to implement PETs, and what specific challenges do they encounter in this process?* ② *Do stakeholders truly understand the privacy gains and technical limitations of PETs, and how does this understanding align with their initial motivations for adoption?* Our findings indicate that while external and internal factors, such as regulatory compliance and reputational considerations, serve as key drivers for PET adoption, organizations face substantial implementation challenges, including usability limitations, complexity management, and insufficient understanding of PET capabilities and constraints. By highlighting these issues, this paper offers—for the first time—a comprehensive meta-analysis of the interplay between organizational decision-making, human perceptions, and technical capabilities of the analyzed building blocks in enterprise cybersecurity contexts and thereby contributes to better aligning these factors.

## 2. A Primer on Privacy-Enhancing Technologies

PETs are a diverse collection of techniques enabling an individual's data security, confidentiality, or anonymity to preserve privacy (Kaaniche et al., 2020). Although privacy primarily centers around an individual's right to control its personally-identifiable information (PII), privacy also affects corporate settings when processing customer's or other individual's data, for instance, by enforcing certain compliance requirements with regulation (Pennekamp et al., 2019). Beyond PII, PETs also foster protection of valuable information including business documents, trade secrets, and other intellectual property.

Both of these use cases boil down to well-known building blocks like *homomorphic encryption* (HE), *differential privacy* (DP), *trusted execution environments* (TEEs), *secure multiparty computation* (SMPC) or *zero-knowledge proofs* (ZKPs), which, if applied correctly, help implement the above-mentioned goals. These building blocks either build upon math (DP), cryptography (HE, SMPC), or hardware support (TEEs), and induce varying levels of complexity, exhibit performance penalties, or require extensive adaptation of algorithms. For instance, configuring DP is non-trivial, while governing the amount of information leakage over time, which either restricts the number of analyses or results in weaker data security than anticipated (Dwork et al., 2019). Thus, the chosen technology and configured parameters are crucial.

Compared to our building-block-centered definition of PETs, others also include VPNs and anonymization networks, encrypted email, or authentication schemes when defining PETs (Kaaniche et al., 2020). These technologies protect *data at rest* and *in transit* or, in the case of authentication, serve orthogonal purposes. While relevant, useful, and widely applied, they cannot protect *data in use*, rendering them insufficient for collaboration across organizations or other use cases requiring special precautions. In this paper, we thus focus on the introduced set of more complex but also more capable PETs.

## 3. Perception of PETS in Related Work

Prior work on the perception of PETs in corporate contexts revealed several influencing factors, often in the context of analyzing technology adoption. For instance, education and regulation are important aspects (Geppert et al., 2022) to increase adoption rates. Likewise, larger organizations tend to be better equipped with necessary infrastructure and human resources for adoption, whereas SMEs suffer from financial constraints (Hasani et al., 2023). From a managerial perspective, moral and ethical considerations were identified as main drivers for PET adoption (Klymenko et al., 2024), whereas employees' opinions reveal that the interplay between technical, organizational, and social factors, such as changed processes or increased transparency are frequent prohibitors of adoption (Bosse et al., 2023; Gan et al., 2019).

Further improving employee education is a frequent recommendation to tackle today's poor adoption rates (Agrawal et al., 2021; Boteju et al., 2023). However, which exact aspects require further training remains unclear. For example, with DP, this inexperience, paired with non-trivial parameter choices, has caused significantly diverging recommendations for privacy budgets (Dwork et al., 2019), risking data leakage despite employing PETs. Such incidents indicate a mismatch between perceived security and actually provided protection.

While the anecdotic focus in prior work provides insightful individual takeaways concerning the perception of PETS in business settings, a systematic and reliable assessment pertaining to motivations and challenges when deploying them in organizations is still missing.

## 4. Survey Methodology

To systematize current perceptions centering around PETs applications in business contexts, we first need to understand challenges and motivations for doing so, that we can then validate for soundness. To this end, we perform a systematic literature review (SLR) (Okoli, 2015). We query Scopus and the ACM DL in August 2025 using the keywords shown in Fig. 1.

TITLE-ABS-KEY(
(“privacy enhancing technology” OR “differential privacy” OR “homomorphic encryption” OR “k-anonymity” OR “secure multiparty computing” OR “trusted execution environment” OR “zero knowledge proof”)
AND
(“interview” OR “perception”))

**Figure 1. Keywords used for queries against SCOPUS and ACM. Querying technology-only keywords yields a significant amount of false positives, largely irrelevant to our study goals. In contrast, including either “interview” or “perception” reduces the number of papers to a manageable size. A manual review of the included papers did not indicate further frequent terms suitable as keywords.**

To extract only relevant paper from the literature corpus, we apply several inclusion and exclusion criteria:

- **IC1** Work analyzes *business context*.
- **IC2** Work discusses *adoption challenges*.
- **IC3** Work discusses *motivations* for adoption.
- **IC4** Work reports *experiences* from implementation.
- **EC1** We were unable to access the paper’s full text.
- **EC2** The study is a duplicate of an other included work.
- **EC3** Matched keywords do not reflect the paper content.
- **EC4** Study reports on technology only.
- **EC5** Use case is unrelated to business context.
- **EC6** The work considers none of our technical building blocks (cf. Section 2).

Matching a single criterion is sufficient for inclusion or exclusion and one domain expert reviewed the 520 papers based on the approach from Okoli (2015). We resolved ambiguities via discussion among the authors[1]. For included papers based on our keyword search, we also performed a backward search, adding 12 papers to our final corpus. Fig. 2 summarizes this workflow detailing numbers for each criterion. Despite our crafted keywords, we find a large fraction of papers solely focusing on the technology but not reporting results on their perception. Research foci unrelated to the keywords and use cases beyond business scenarios are other frequent reasons for exclusion. Overall, this strategy yields a final corpus of 34 papers.

We assess these 34 works in detail by extracting mentioned motivations, challenges, experiences, and other reported findings. More precisely, we employ the Technology-Organization-Environment (TOE) framework (Tornatzky et al., 1990) to systematize these factors influencing the perception of PETs in business contexts. Other well-known technology adoption models like TAM (Marangunić & Granić, 2015) or UTAUT (Williams et al., 2015) are not applicable because we have no control of the original study designs. We thus rely on the TOE framework to provide a structured analysis of the factors influencing the perception of PETs in business contexts.

[1]For reproducibility, we provide the full list of papers besides review decisions as an artifact to this paper online: https://dx.doi.org/10.5281/zenodo.17175132

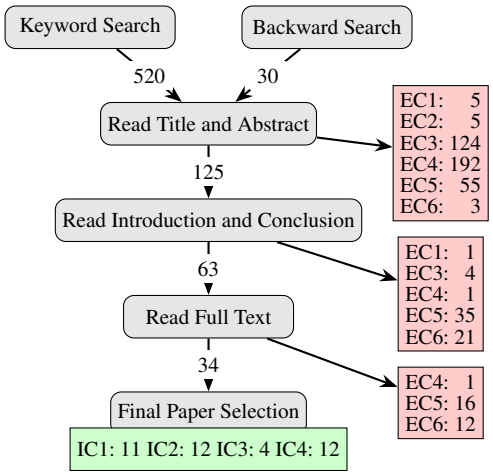


**Figure 2. Included and excluded papers per step.**

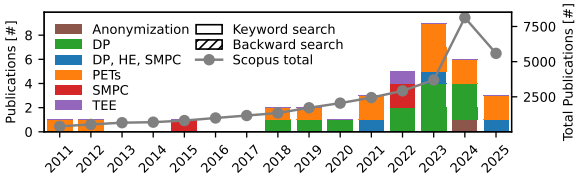


**Figure 3. Included studies tend to be recently published, but altogether cover more than a decade of research.**

**TOE-Framework (Tornatzky et al., 1990)**

TOE analyzes how contexts shape the adoption and implementation of technological innovation by attributing factors to these three categories. The *technological context* involves technologies already in use, under consideration, besides available external technologies that collectively set limits and opportunities for the firm. The *organizational context* covers internal factors, such as structures, communication processes, size, and available resources, influencing the complexity and effectiveness of adoption and implementation processes. Lastly, the *environmental context* captures external influences including industry dynamics, regulatory environments, competitive pressures, and support infrastructures like skilled labor and consultancy services.

## 5. Systematizing PETs Adoption in Business Contexts

In this section, we give a quantitative overview of our paper corpus obtained from our SLR.

Our 34 selected papers have been published between 2012 and 2024, with an increase over time. Fig. 3 shows their distribution over the years. While the volume of papers covering PETs generally increases over the years, the shift from technical analyses to also include organizational and managerial perspectives is a recent trend.

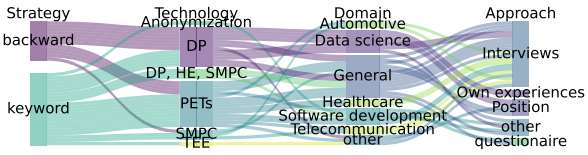


**Figure 4. Distribution of included studies across search strategy, covered technologies, domains, and research approaches.**

Interestingly, the first two papers (2011/12) already criticize slow adoption rates, an effect that persists over the years. Only after 2020, we note an increase in volume and diversity of domains and methods: Out of the 34 works, 23 report on self-conducted interviews with 12 participants on average, five conduct only surveys with an average of 241 respondents, and two papers report own experiences after deploying PETs to business scenarios. With general data science (9), software development (5), automotive (2), healthcare (3), telecommunications (2), and finance (1), we consider these covering a wide diversity of domains in which PETs are relevant.

Fig. 4 visualizes the interdependencies between analyzed technologies, domains, and research methodologies, showing the rich diversity. Interestingly, we find no papers published before 2011, despite concepts like differential privacy and PETs being significantly older.

We identify nine researchers authoring more than one (at most four) publications in our dataset. Most journals and conference only appear once in our dataset, except for the HDSR journal and the ACM CHI conference, with three and two papers, respectively. Overall, two papers were published at an A*-ranked conference and four papers were published in journals with an IF above 5.

Technology-wise, we count 15 papers explicitly discussing differential privacy, whereas SMPC (5), HE (2), and TEEs (1) building blocks receive less frequent consideration. Surprisingly, 14 works refrain from studying a concrete building block, pointing out only general conclusions. Due to the small corpus size, these numbers are unlikely to reflect real-world deployment quantity.

Overall, we find the growing interest in the perception of PETs to be primarily driven by research on DP (44 % of papers). Still, studies also cover other building blocks in various domains while originating from a heterogeneous group of authors.

## 6. Qualitative Analysis

We examine the selected papers qualitatively, focusing on motivators behind PET adoption (Section 6.1), perceived challenges and experiences (Section 6.2), the understanding of privacy gains (Section 6.3), and understood technical limitations (Section 6.4).

**Methodology and Validation.** Based on our corpus, we manually code described factors into motivators and challenges across the TOE framework dimensions, before clustering similar factors. We summarize these findings in Table 1, ordered by frequency of occurrence. To ensure that our coding is reliable and not solely based on our interpretation, we embed each sentence in our corpus using the `all-MiniLM-L6-v2` model, and then calculate cosine similarity between individual sentences and paraphrased variants of the identified factors as a means to validate our labeling. This approach yields an AUC of 0.61, with a precision of 0.29 and recall of 0.76 when labeling all sentences with a similarity above 0.3 with any of the analyzed factors in Table 1. Manually inspecting labeled sentences reveals that this strategy detects relevant factors with high accuracy, but is often unable to distinguish between challenges and motivations, thereby creating a high number of false positives. Still, this method confirms our manual coding process.

### 6.1. Motivations for Adopting PETs

Organizations adopt PETs due to multiple factors, although most of them indicate some external pressure. Table 1 lists those factors mentioned in our corpus along the TOE categories. The most common motivators include compliance with regulations (EM1, 47 % of papers), organizational reputation (EM2, 35 %), improved collaboration on data (EM3, 32 %), and competitive advantages (EM4, 21 %). From these, we find regulatory compliance particularly motivating organizations in highly regulated sectors such as healthcare, finance, and life sciences.

Internally, organizations recognize PETs for reducing legal risks (OM1, 38 %), contributing to their corporate identity (OM2, 32 %), and improving processes (OM4, 24 %). Ethical motivations (OM3, 24 %) also significantly influence adoption decisions, with organizations increasingly viewing privacy protection as part of corporate social responsibility.

Comparatively fewer papers report primarily technology-driven adoption. When technological factors do motivate adoption, innovation (TM1, 26 %) and novelty–for instance, privacy-preserving utilization of cloud applications (Geppert et al., 2022)–are reported to help maintain competitive advantages. Access to previously untapped data (TM3, 21 %) serves a similar motivational purpose, while reduced risks of data breaches (TM2, 24 %) functions as prevention. However, these technological factors appear less frequently than environmental and organizational factors, and tend not to be the primary motivation for introducing PETs.

Factors such as firm size, resource availability, tech-

| | Technological | ∑ 30 | Organizational | ∑ 31 | Environmental | ∑ 29 |
|---|---|---|---|---|---|---|
| Motivations | TM1 Technological innovation | 9 | OM1 Privacy and risk | 13 | EM1 Regulatory compliance | 16 |
| | TM2 Risk reduction | 8 | OM2 Corporate identity | 11 | EM2 Trust and reputation | 12 |
| | TM3 Data utility | 7 | OM3 Ethical commitments | 8 | EM3 Improved collaboration on data | 11 |
| | | | OM4 Developing and improving services | 8 | EM4 Market positioning | 7 |
| | | | OM5 Optimizing business processes | 6 | EM5 Availability of products and support | 6 |
| | | | OM6 Financial incentives | 5 | EM6 Customer and partner influence | 6 |
| | | | OM7 Internal drivers | 5 | | |
| | | | OM8 Building trust | 4 | | |
| Challenges | TC1 Choosing (privacy) parameters | 12 | OC1 Lack of training and education | 14 | EC1 Lack of guidance regarding regulation | 8 |
| | TC2 Privacy-utility trade-off | 11 | OC2 Changing workflows | 11 | EC2 Privacy communication to customers | 6 |
| | TC3 Ease of use | 8 | OC3 Adoption costs | 11 | EC3 Lack of documentation resources & support | 6 |
| | TC4 Complexity | 8 | OC4 Interdisciplinary communication | 11 | EC4 Measuring successful implementation | 5 |
| | TC5 Risk and uncertainty | 8 | OC5 Lack of information on PETs | 10 | EC5 Lack of information | 5 |
| | TC6 Compatibility | 7 | OC6 Complacency | 8 | EC6 Research vs. real-world | 4 |
| | TC7 Computational demands | 4 | OC7 Organizational culture | 8 | EC7 Lack of education | 4 |
| | TC8 Technological limitations | 4 | OC8 Regulatory compliance | 7 | EC8 Lack of knowledge about practical needs | 4 |
| | TC9 Lack of software support | 3 | OC9 Choosing parameters | 7 | EC9 Lack of incentives | 3 |
| | TC10 Availability & Robustness | 2 | OC10 Lack of experts | 7 | EC10 Lack of ready-made solutions | 3 |
| | TC11 Differences in algorithms | 1 | OC11 Identifying best PET for task | 6 | EC11 Low adoption rates | 3 |
| | | | OC12 Complexity of technology | 5 | EC12 Side-effects of increased transparency | 2 |
| | | | OC13 Lack of technological maturity | 4 | | |
| | | | OC14 Invisibility of benefits | 2 | | |

**Table 1. Motivations and challenges mentioned across included papers categorized via the TOE framework.**

nical readiness, managerial support, and employee acceptance further moderate these motivations and can create significant challenges during adoption processes, as discussed in the following (Bada et al., 2023).

## 6.2. Reported Challenges and Experiences

Given the diverse building blocks employed in our dataset (cf. Section 5), we observe a broad range of challenges and experiences in deploying these technologies. Except for zero-knowledge proofs, the included papers examine all analyzed technologies (cf. Section 2). While we observe a focus on DP that leads to certain technology-specific challenges, such as parameter selection (TC1, 35 %), being mentioned particularly frequently, many challenges affect all building blocks. Since papers often address multiple technologies collectively, our study design does not permit sharp distinctions; therefore, we consider a cross-sectional view for all technologies while highlighting DP-specific aspects where relevant.

**6.2.1. Technological Perspective** We identify three challenging themes: design tensions, implementation barriers, and operational requirements.

Design tensions comprise parameter selection (TC1, 35 %) and the privacy-utility trade-off (TC2, 32 %), both particularly important for DP (53 % of DP papers mention either challenge) but also affecting other technologies (16 % of others). Especially the privacy-utility trade-off is perceived as a significant barrier.

Regarding the implementation and integration of PETs, organizations encounter several interconnected challenges: compatibility (TC6, 21 %) with legacy systems and existing architectures, frequently requiring significant workarounds and introducing complications; the inherent technological complexity (TC4, 24 %), encompassing both computational demands (TC7, 12 %) and algorithms (TC11, 3 %). Reportedly low technological readiness levels complicate this situation (Agrawal et al., 2021). Here, the perceived complexity can be so daunting that practitioners avoid learning about certain technologies altogether.

Operationally, usability (TC3, 24 %) compared to transparent and well-understood legacy systems and workflows, availability (TC10, 6 %) due to increased complexity and novel failure modes, and efficiency considerations (TC7, 12 %) represent commonly reported challenges across all technologies, albeit they do not seem specific for PETs, but summarize general challenges of technology adoptions. Additionally, certain technologies like SMPC impose technical limitations, for example, regarding dataset sizes which are not reported for, e.g., DP.

**6.2.2. Organizational Perspective** The organizational dimension encompasses knowledge gaps, structural barriers, and governance issues.

Knowledge and capability gaps, primarily involve a perceived lack of training and education (OC1, 41 %), creating cascading implementation difficulties. Organizations report employees lacking confidence in software that incorporates PETs effectively due to insufficient training, while educational materials remain largely inaccessible to non-academic audiences (cf. EC6).

Participants report existing documentation as primarily targeting academics, while perceiving a significant entry burden for developers. This knowledge gap directly contributes to a reported insufficient information availability (OC5, 29 %) that leads, e.g., to employees holding on to misunderstandings about the technology–such as be-

lieving that analysis on encrypted data is impossible–that impede adoption efforts eventually. From the reported findings, companies rarely perceive their in-house PET expertise as sufficient (OC10, 21 %) and struggle to find external experts.

Organizational adoption of PETs is consistently reported as requiring significant workflow and process changes (OC2, 32 %), although varying across technologies. For instance, data scientists perceive SMPC as challenging data management practices by thwarting raw data analysis and exploratory queries (Agahari et al., 2022). Cultural barriers include low privacy prioritization (OC7, 24 %) and communicative challenges between developers, legal, and management teams (OC4, 32 %), that are perceived as having distinct vocabularies and knowledge bases. Complacency (OC6, 24 %)——assuming existing measures are "good enough"——and perceived PET invisibility (OC14, 6 %) complement cultural barriers.

Lastly, governance and compliance constitute organizational challenges. How should users be involved in parameter selection processes? Who is generally authorized to translate privacy principles into concrete technical parameters? (OC9, 21 %). Related questions create additional liability challenges that are reported as complicating regulatory compliance (OC8, 21 %).

**6.2.3. Environmental Perspective** Environmental challenges encompass regulatory frameworks, knowledge dissemination, support infrastructures, and market dynamics. They reside beyond organizational control yet significantly influence internal adoption decisions.

Regulatory ambiguities (EC1, 24 %) pose barriers through the absence of clear PET implementation guidance. A perceived lack of technology-specific mandates incentivizes minimal compliance solutions rather than comprehensive privacy implementations (Klymenko et al., 2024). Evaluation guideline deficits (EC4, 15 %) reinforce this effect through absent standardized assessment frameworks. As a result, organizations also perceive difficulties in establishing evaluation and audit processes for monitoring the efficiency of measures (Cummings & Sarathy, 2023).

Besides, gaps in translating academic results into practice (EC6, 12 %) create implementation obstacles, with academia addressing simplified scenarios inadequately reflecting real-world complexity (Klymenko et al., 2024). Here, organizations perceive a disconnect between theoretical development and practical application requirements. Academic results are perceived as insufficiently integrated in curricula (EC7, 12 %), contributing to expert shortages and largely confining PET coverage to advanced computer science courses (Garfinkel et al., 2018).

Perceived deficiencies in documentation and support (EC3, 18 %) manifest through absent technical frameworks, limited implementation literature, inadequate development toolkits, and a perceived insufficiency of external consultancy services (Klymenko et al., 2023; Panavas et al., 2024). A lack of off-the-shelf solutions (EC10, 9 %) further constrains adoption through the scarcity of pre-built implementations and inadequate open-source library development. Lastly, incentivization issues (EC9, 9 %) undermine deployment through insufficient rewards for PET implementation, as users are often unwilling to pay extra (Pape & Harborth, 2023). In a similar direction, adoption of PETs also creates communicative challenges (EC2, 18 %) requiring organizations to balance technical accuracy with accessibility when communicating PET benefits to customers and partners and may create suspicion about underlying motivations of introduction. Besides, stringent transparency about privacy measures (EC12, 6 %) has side effects, such as DP producing customer complaints based on misconceptions about data quality (Garfinkel et al., 2018). Collectively, these challenges highlight the complexity inherent in incorporating PETs as cybersecurity measures for organizations.

## 6.3. Understanding of Privacy Gains

While motivations largely originate from external pressures and high-level organizational goals, we analyze reported understanding of privacy gains based on practical experiences and implementations.

Internally, PET adoption produces notable operational improvements, including enhanced data management processes and privacy-by-design workflows (Munilla Garrido et al., 2023). Organizations report substantial gains from better data handling practices, reducing privacy breach likelihood and unauthorized disclosures, that might otherwise result in loosing competitive advantages (Kühtreiber et al., 2023). Notably, legal departments rarely serve as driving force behind implementations of these rules. Instead, for DP, scientific or engineering divisions were pointed out as primary drivers (Dwork et al., 2019).

However, despite implementation experience, differences in technological capabilities remain unclear even to seasoned engineers, sometimes resulting in wrong conclusions. For instance, Kühtreiber et al. (2023) report participants fail to recognize DP's advantages over k-anonymity, incorrectly stating that k-anonymity is favorable for added security. Similarly, confusion between pseudonymization and k-anonymity occurs, with practitioners incorrectly assuming that k-anonymity offers no additional benefits on already pseudonymized data.

In this vein, SMPC and HE are frequently considered "black-boxes", even in expert-only interview studies (Balebako et al., 2014). While stakeholders acknowl-

edge PETs' potential for outsourcing computation and interorganizational collaboration, they remain skeptical about deploying these technologies for business-critical data sharing. In automotive contexts, concerns remain about the potential loss of competitive advantages to collaborators (Kühtreiber et al., 2023). Similarly, SMPC and HE are often perceived as technological "overkill," without adequate consideration of alternative approaches or recognition of inherent limitations in methods such as encryption without runtime protection (Kühtreiber et al., 2023). Conversely, less experienced respondents repeatedly expressed a preference for privacy-as-a-service solutions, envisioning pre-configured modules that could integrate seamlessly into existing processes to provide privacy-preserving equivalents (Agrawal et al., 2021).

Within our corpus, we find consistent conclusions across the investigated technologies for corporate environments. However, beyond the corporate setting, interviews with medical practitioners yield distinctly different perspectives (Alaqra et al., 2021), primarily assessing data sensitivity through legal regulations rather than evaluating actual privacy breach risks. When asked to compare different technologies, they perceive corresponding tools as black-boxes and defer to IT experts for technical specifications while desiring "sufficient knowledge" about data storage and utilization to explain aspects to patients, leading them to favor simpler PET implementations.

These findings demonstrate that organizations generally recognize PET advantages, yet misconceptions lead to reduced protection levels compared to technical possibilities. Similar patterns emerge regarding technological limitations, albeit presenting greater risks when overestimating capabilities.

### 6.4. Understanding of Limitations

PETs have been characterized as "black magic" or "holy grails", reflecting both perceived potential and implementation mystique (Balebako et al., 2014). While acknowledging PET functionality, such characterizations reveal concerning gaps in stakeholder understanding regarding limitations and risks. We find organizations and practitioners frequently operating with incomplete comprehension of constraints, as discussed in the following.

For the parameter selection of DP, even experienced data scientists exhibit challenges in selecting appropriate $\varepsilon$ values, while displaying high confidence in defending their choices, despite significant variability in parameter selection across studies (Ngong et al., 2024). Some DP libraries issue warnings indicating that the mitigation of such errors is feasible, though it remains underutilized (Song et al., 2024). More concerning, certain widely-used DP libraries facilitate misconfigurations rather than preventing them (Ngong et al., 2024).

Beyond parameter selection, practitioners criticize DP for excessive data distortion, with some characterizing noise addition as fundamentally deceptive (Kühtreiber et al., 2023). One participant even remarked that adding noise feels akin to lying (Agrawal et al., 2021). The reviewed studies agree that such effects are particularly pronounced in small datasets but tend to disappear on scale. Repeated calls for improved training and education are called out as a solution; however, required training amounts remain uncertain, as even interviewed experts do not consistently exhibit better parameter choosing practices (Ngong et al., 2024).

Concerningly, discussion on other technology limitations, particularly HE, SMPC, and TEEs, remains limited. Some interviewees highlighted poor HE performance compared to SMPC (Agrawal et al., 2021), though context on the specific use case this critique pertains to is missing, as this is not inherently a technology-induced limitation. Introducing HE was noted to "significantly slow down a project" (Kühtreiber et al., 2023) and require deep mathematical understanding (Agrawal et al., 2021). While recent advancements in general-purpose compilers aim to mitigate this issue (Viand et al., 2021), these developments appear to be underutilized or unknown among practitioners. Notably, inherent security factors related to HE and SMPC, such as key management and non-collusion assumptions, are not addressed in any of our analyzed studies. Both aspects are fundamental for cybersecurity and, if neglected, can undermine privacy.

Independent of specific technologies, a recurring concern is the lack of transparency and accountability. If regulators, customers, or consumers remain oblivious of the implementation, e.g., key distribution in the case of HE, implementation of countermeasures in case of TEEs, or parameter choices for DP, security standards may be perceived higher than actually implemented (Dwork et al., 2019). As a potential solution, multiple interviewees expressed a desire for improved transparency, for instance, through open-source software, public privacy parameters (Dwork et al., 2019), or establishing standards. However, consensus on the coverage these standards should have and how they could be made comprehensible for both practitioners and consumers is missing.

### 6.5. Recommendations for Improvement

The tension between technical complexity and expertise in handling PETs emerges as a recurring barrier. Studies in our corpus already propose a set of recommendations and improvement strategies, often aligned in their direction, that we distill and assess as follows.

Several DP studies suggest *improved tooling sup-*

*port*. Comparisons of competing libraries demonstrate that the potential for error prevention is not equally well exploited (Ngong et al., 2024; Song et al., 2024). For other PETs, interviewees express recurring desires to make existing tools more efficient, practical, and user-friendly—up to the point of creating an all-encompassing solution (Kühtreiber et al., 2023). Future work should therefore investigate common sources of errors in greater detail. Whether the aspiration for a fully automated "black-box" solution is feasible remains to be seen; if feasible, such a solution would then also need to account for the configuration pitfalls and understanding difficulties summarized in our study.

A frequently proposed alternative to a privacy "black-box" is *enhanced education* (Alaqra et al., 2021; Sarathy et al., 2023). Indeed, training outcomes within the reviewed studies, particularly among non-experts, show promising results (Munilla Garrido et al., 2023). Overall, the analyzed studies suggest that even basic awareness and basic training can sufficiently alert stakeholders, which then initiate further steps with intrinsic motivation. External certifications and visibility for customers and consumers could further accelerate these processes while simultaneously uncovering typical vulnerabilities.

Ultimately, many of the described challenges and perceptions are not primarily technical in nature but stem from the *gap between academia and practice*. Interviewees expressed optimism about adopting additional PETs in the future (Boteju et al., 2023); however, exact inhibitors remain unclear. Given the overall imbalance favoring technical solutions, which are rarely complemented by experience reports from business applications, the true extent of the problem remains uncertain. Future work should therefore focus more explicitly on identifying ways to analyze this issue.

These findings suggest that PET adoption requires coordinated intervention across technical, educational, and organizational dimensions. The security research community's focus on technical advancement, while necessary, is insufficient without corresponding attention to implementation barriers and practitioner needs. For DP, the case of the US Census Bureau has created a blueprint here. Something similar would also be desirable for other technologies and represents a promising opportunity to reduce misperceptions toward these technologies. Moving forward, research initiatives should explicitly incorporate practitioner perspectives and real-world deployment constraints to ensure that innovations translate accordingly.

# 7. Discussion and Implications

This discussion synthesizes our findings on the perception of PETS in business contexts. We first examine our study's limitations (Section 7.1), then contextualize findings and derive implications and recommendations for organizational cybersecurity practices (Section 7.2).

## 7.1. Limitations of our Systematic Approach

As a mostly qualitative survey across other primary research, this study does not claim statistical representativeness. Moreover, due to the sensitivity of reported insights and selection bias of interviewees, public interviews might not realistically reflect the actual implementation status. For the same reason, participants might have been biased toward reporting positive results only, an inherent limitation in interview-based studies. Lastly, with "perception" and "interview" as keywords, this survey specifically targeted primary research taking an empirical, often qualitative approach, potentially missing technical implementation experiences reported in other contexts. Nevertheless, by synthesizing and systematizing organizational PET perceptions, this paper provides valuable insights to guide both future research and practical adoption efforts.

## 7.2. Discussion

Answering our first research aspect on motivations and challenges of PETs adoption, we identified organizational, technical, and environmental motivators (cf. Section 6.1) driving PETs adoption in organizations, thereby contributing to an improved cybersecurity. Contrary to (our) expectations, the primary inhibitors challenging adoption (cf. Section 6.2) are not technical limitations but rather soft factors such as regulatory constraints, lack of expertise, and the complexity of the technology and existing processes, which are perceived as significant barriers to the adoption of PETs.

Regarding our second research aspect, i.e., the understanding of privacy gains and limitations and their alignment with initial motivations, we rarely find in-depth responses demonstrating a complete understanding of relevant influencing factors. For instance, few studies (8.8 %, all of them discuss DP) address threat models or security-relevant design decisions, such as collusion assumptions, questioning whether PETs uphold their academic promises when utilized in practice. While participants demonstrate some awareness of security differences, their understanding is inconsistent and sometimes contradicts established principles, such as the relationship between k-anonymity and DP (Kühtreiber et al., 2023).

In the following, we summarize implications from our findings, and derive eight *concrete recommendations* R1 to R8 for practice and research.

**7.2.1. Practical Implications** Interviewees and study participants demonstrate the capability to understand key PET concepts during training, often abandoning prior misconceptions regarding the capabilities and limitations of these cybersecurity measures. Thus, we recommend **R1**: *regular training sessions* to account for evolving capabilities, as even concise sessions were shown to be effective. In addition, we advocate an **R2**: *increased transparency of chosen parameters and, in particular, design decisions* in HE and SMPC, to facilitate externally verifiability through certification measures. Such transparency would enable customers and consumers to choose secure service providers better, thereby creating added value. Eventually, added value would also **R3**: *solve persisting incentivation issues*, including invisibility and the lack of benefits (cf. Section 6.2). On a broader scale, our findings imply that organizations need to consider a multitude of primarily organizational prohibitors and need to **R4**: *reassess or regularly update their perceptions* due to evolving PET capabilities.

**7.2.2. Scholarly Implications** For cybersecurity research, complexity, usability, and implementation costs remain implementation barriers, which are being investigated but, according to the included studies, are still not mature enough. This status quo highlights the importance of **R5**: *facilitating easy-to-understand and straightforward solutions*. However, albeit repeatedly mentioned across studies, blackbox or privacy-as-a-service solutions might overshoot in the sense that we already see exaggerated expectations, which might as well create further misperceptions of their own. A related approach could be to better **R6**: *indicate easy-to-follow pathways for translating cryptographic or technical advancements into practice*, which is increasingly fostered by security communities (Grobler et al., 2021).

We also call for **R7**: *research on implementation failures* and vulnerabilities, which we rarely find, possibly due to topic sensitivity. Future research should incorporate longitudinal analyses to address this gap and focus on long-term user experience and implementation challenges, as current studies provide only point-in-time snapshots, ignoring security perception evolution during (lengthy) implementation processes. Finally, while this study outlined factors for organizational PETs adoption, future work should **R8**: *scale our analysis* to a broader paper corpus, for instance, by scaling our semi-automated validation strategy, or confirming findings through empirical studies, e.g., based on Technology Acceptance Methodology (TAM), that allow statistic conclusions between the different technologies.

# 8. Conclusion

Our systematic literature review of the experiences and perceptions of PETs in business contexts revealed that while external pressures, such as regulatory compliance and reputation, drive adoption, significant organizational and technical challenges prevail. Most notably, these challenges include the complexity of PETs, usability issues, and a lack of training and understanding of their capabilities and limitations. Reported misperceptions, particularly regarding parameter selection and security guarantees, highlight the need for joint efforts in improving tooling, education, and transparency. Despite these challenges, PETs, nonetheless, provide organizations with substantial potential for improving cybersecurity and protecting intellectual property. Thus, future research should focus on better translating academic advancements into practical implementations, fostering (transparency) standards, and addressing long-term usability and adoption barriers. Thereby, PETs can increase cybersecurity and data privacy in increasingly data-driven business environments.

# Acknowledgments

This work has been funded by the German Federal Ministry of Research, Technology and Space (BMFTR) under funding reference number 02J24A030. The responsibility for the content of this publication lies with the authors.

# References

Agahari, W., Ofe, H., & De Reuver, M. (2022). It is not (only) about privacy: How multi-party computation redefines control, trust, and risk in data sharing. *Electron. Mark.*, *32*(3).

Agrawal, N., Binns, R., … Shadbolt, N. (2021). Exploring Design and Governance Challenges in the Development of Privacy-Preserving Computation. *CHI*.

Alaqra, A. S., Kane, B., & Fischer-Hübner, S. (2021). Machine Learning–Based Analysis of Encrypted Medical Data in the Cloud: Qualitative Study of Expert Stakeholders' Perspectives. *JMIR Hum. Factors*, *8*(3).

Bada, M., Furnell, S., … Dymydiuk, J. (2023). Supporting Small and Medium-Sized Enterprises in Using Privacy Enhancing Technologies. *HCI-CPT*.

Balebako, R., Marsh, A., … Faith Cranor, L. (2014). The Privacy and Security Behaviors of Smartphone App Developers. *USEC*.

Bosse, C. K., Feth, D., & Schmitt, H. (2023). Challenges, conflicts, and solution strategies for the intro-

duction of corporate data protection measures. In *Human Factors in Privacy Research*.

Boteju, M., Ranbaduge, T., … Arachchilage, N. A. G. (2023, December 30). *SoK: Demystifying Privacy Enhancing Technologies Through the Lens of Software Developers*. arXiv: 2401.00879.

Braud, A., Fromentoux, G., … Le Grand, O. (2021). The Road to European Digital Sovereignty with Gaia-X and IDSA. *IEEE Netw.*, *35*(2).

Cummings, R., & Sarathy, J. (2023). Centering Policy and Practice: Research Gaps Around Usable Differential Privacy. *IEEE TPS-ISA*.

Dwork, C., Kohli, N., & Mulligan, D. (2019). Differential Privacy in Practice: Expose your Epsilons! *J. Priv. Confidentiality*, *9*(2).

Gan, M. F., Chua, H. N., & Wong, S. F. (2019). Privacy Enhancing Technologies implementation: An investigation of its impact on work processes and employee perception. *Telemat. Inform.*, *38*.

Garfinkel, S. L., Abowd, J. M., & Powazek, S. (2018). Issues Encountered Deploying Differential Privacy. *WPES*.

Geppert, T., Anderegg, J., … Ebert, N. (2022). Overcoming Cloud Concerns with Trusted Execution Environments? Exploring the Organizational Perception of a Novel Security Technology in Regulated Swiss Companies. *HICSS*.

Grobler, M., Gaire, R., & Nepal, S. (2021). User, Usage and Usability: Redefining Human Centric Cyber Security. *Front. Big Data*, *4*.

Hasani, T., Rezania, D., … Mohammadi, M. (2023). Privacy enhancing technology adoption and its impact on SMEs' performance. *Int. J. Eng. Bus. Manag.*, *15*.

Kaaniche, N., Laurent, M., & Belguith, S. (2020). Privacy enhancing technologies for solving the privacy-personalization paradox: Taxonomy and survey. *J. Netw. Comput. Appl.*, *171*.

Klymenko, O., Meisenbacher, S., … Matthes, F. (2024). On the Integration of Privacy-Enhancing Technologies in the Process of Software Engineering: *ICEIS*.

Klymenko, O., Meisenbacher, S., & Matthes, F. (2023). Identifying Practical Challenges in the Implementation of Technical Measures for Data Privacy Compliance. *AMCIS*.

Kühtreiber, P., Pak, V., & Reinhardt, D. (2023). "A method like this would be overkill": Developers' Perceived Issues with Privacy-preserving Computation Methods. *IEEE TrustCom*.

Lohmöller, J., Pennekamp, J., … Wehrle, K. (2024). The unresolved need for dependable guarantees on security, sovereignty, and trust in data ecosystems. *Data & Knowl. Eng.*, *151*.

Marangunić, N., & Granić, A. (2015). Technology acceptance model: A literature review from 1986 to 2013. *Univers. Access Inf. Soc.*, *14*(1).

Munilla Garrido, G., Liu, X., … Song, D. (2023). Lessons Learned: Surveying the Practicality of Differential Privacy in the Industry. *PoPETs*, *2023*(2).

Ngong, I. C., Stenger, B., … Feng, Y. (2024). Evaluating the usability of differential privacy tools with data practitioners. *SOUPS*.

Okoli, C. (2015). A Guide to Conducting a Standalone Systematic Literature Review. *Commun. Assoc. Inf. Syst.*, *37*.

Panavas, L., Sarker, A., … Mahyar, N. (2024). Illuminating the Landscape of Differential Privacy: An Interview Study on the Use of Visualization in Real-World Deployments. *IEEE Trans. Visual. Comput. Graphics*.

Pape, S., & Harborth, D. (2023). Acceptance Factors of Privacy-Enhancing Technologies on the Basis of Tor and JonDonym. In *Human Factors in Privacy Research*.

Pennekamp, J., Henze, M., … Wehrle, K. (2019). Dataflow Challenges in an Internet of Production: A Security & Privacy Perspective. *ACM CPS-SPC*.

Prince, C., Omrani, N., … Kraus, S. (2023). Are We Living in Surveillance Societies and Is Privacy an Illusion? An Empirical Study on Privacy Literacy and Privacy Concerns. *IEEE Trans. Eng. Manage.*, *70*(10).

Sarathy, J., Song, S., … Vadhan, S. (2023). Don't Look at the Data! How Differential Privacy Reconfigures the Practices of Data Science. *CHI*.

Schlackl, F., Link, N., & Hoehle, H. (2022). Antecedents and consequences of data breaches: A systematic review. *Inf. Manag.*, *59*(4).

Song, P., Sarathy, J., … Vadhan, S. (2024). "I inherently just trust that it works": Investigating Mental Models of Open-Source Libraries for Differential Privacy. *Proc. ACM Hum.-Comput. Interact.*, *8*(CSCW2).

Tornatzky, L. G., Mitchell, F., & Chakrabarti, A. (1990). *The processes of technological innovation*. Lexington books.

Viand, A., Jattke, P., & Hithnawi, A. (2021). SoK: Fully Homomorphic Encryption Compilers. *IEEE SP*.

Williams, M. D., Rana, N. P., & Dwivedi, Y. K. (2015). The unified theory of acceptance and use of technology (UTAUT): A literature review. *J. Enterp. Inf. Manag.*, *28*(3).