

Review

Not peer-reviewed version

---

# A Systematic Review of Privacy-Enhancing Technologies (PETs) for Securing Personally Identifiable Information in Public Cloud Architectures

---

[David Roberts](#)<sup>\*</sup> and Katie Garcia<sup>\*</sup>

Posted Date: 4 February 2026

doi: 10.20944/preprints202602.0303.v1

Keywords: privacy-enhancing technologies; PII; public cloud; homomorphic encryption; federated learning; differential privacy; confidential computing



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Review

# A Systematic Review of Privacy-Enhancing Technologies (PETs) for Securing Personally Identifiable Information in Public Cloud Architectures

David Roberts <sup>1,\*</sup> and Katie Garcia <sup>2,\*</sup>

<sup>1</sup> University of Stirling, Stirling, FK9 4LA, Scotland, UK

<sup>2</sup> University of Bradford, Bradford, West Yorkshire, BD7 1DP, UK

\* Correspondence: trenddevelopment68@gmail.com (D.R.); kate63144@gmail.com (K.G.)

## Abstract

The rapid adoption of public cloud services has significantly increased the storage and processing of Personally Identifiable Information (PII), raising critical concerns about data privacy and regulatory compliance. Privacy-Enhancing Technologies (PETs) have emerged as a crucial set of methods and tools designed to protect sensitive data while maintaining functional utility for cloud applications. This systematic review examines the current landscape of PETs deployed for securing PII in public cloud architectures, including homomorphic encryption, differential privacy, federated learning, and confidential computing. A structured literature search was conducted across major scientific databases from 2021 to 2026, following PRISMA guidelines, resulting in the inclusion of studies that evaluate PET performance, scalability, security guarantees, and integration challenges. Thematic synthesis highlights key trends, such as the growing adoption of federated learning for cross-silo data sharing, the application of homomorphic encryption in real-time cloud environments, and the trade-offs between privacy preservation and computational efficiency. Additionally, operational, technical, and regulatory challenges are identified, including computational overhead, standardization barriers, and compliance with global data protection regulations. This review underscores the critical role of PETs in enhancing trust and security in public cloud ecosystems and provides insights for researchers and practitioners seeking to design and implement privacy-aware cloud architectures. Future research directions are discussed, emphasizing the need for optimized PET frameworks that balance security, efficiency, and compliance in increasingly complex cloud environments.

**Keywords:** privacy-enhancing technologies; PII; public cloud; homomorphic encryption; federated learning; differential privacy; confidential computing

---

## 1. Introduction

Privacy-Enhancing Technologies (PETs) are technical solutions designed to minimize personal data usage while maximizing security and user control, particularly when sensitive information, such as Personally Identifiable Information (PII), is processed in third-party infrastructures like public clouds [1]. As organizations increasingly adopt public cloud computing for scalability, elasticity, and cost efficiency, these environments introduce significant privacy risks, including unauthorized access, misuse, and regulatory non-compliance during data storage, transit, and processing [2]. This tension highlights the urgent need for systematic evaluation of PETs tailored to protecting PII in public cloud architectures.

Recent research demonstrates that PETs now extend beyond traditional cryptographic mechanisms to include confidential computing with Trusted Execution Environments (TEEs), federated analytics, and differential privacy, protecting data both at rest and in use [3,4]. These

developments reflect the shift from perimeter-based defenses to privacy-by-design frameworks, where privacy is embedded into system architecture, policies, and operational practices. For example, privacy-preserving data processing frameworks enable computation on encrypted or partitioned data, allowing cloud services to extract insights without exposing raw PII [3].

Despite these advancements, challenges remain, such as performance overheads, lack of standard benchmarks, and difficulties integrating PETs with existing cloud models. Many studies focus on methodological innovation without systematically synthesizing implications for PII protection in public cloud environments. Furthermore, while foundational PET research exists, the rapid evolution of technology and regulatory requirements makes a contemporary review essential, particularly emphasizing literature from 2020 to 2026.

This review aims to systematically assess and synthesize the literature on PETs for securing PII in public cloud environments, highlighting state-of-the-art solutions, practical implications, and remaining challenges. In doing so, it identifies thematic clusters of technologies, evaluates their strengths and limitations, and suggests promising directions for future research. Notably, applied research in cloud application design and deployment, such as AWS-integrated platforms, underscores the relevance of architectural and implementation considerations for privacy protection [5].

## 2. Methods (Literature Search Strategy)

This systematic review followed a predefined literature search protocol to ensure rigor, transparency, and reproducibility. The strategy was developed in accordance with established systematic review reporting standards, including the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) framework, which outlines clear reporting requirements for literature identification, screening, and selection processes [6,7]. Key components of the methods include database selection, search term development, eligibility criteria, screening process, and data extraction.

To identify relevant literature on Privacy-Enhancing Technologies (PETs) for securing Personally Identifiable Information (PII) within public cloud architectures, multiple bibliographic databases were systematically searched. The databases selected reflect the interdisciplinary nature of PET research, encompassing cloud computing, privacy, and security domains. Specifically, the search was conducted in **IEEE Xplore**, **ACM Digital Library**, **Scopus**, **Web of Science**, and **ScienceDirect** to maximize coverage of peer-reviewed research published between **January 2020 and December 2025**.

Search strings were constructed through iterative refinement of keywords and controlled vocabulary relevant to the review topic. Core constructs included terms relating to privacy, PETs, cloud computing, and PII protection. Boolean operators (AND, OR) and truncation were used to combine terms effectively, such as: (*"privacy-enhancing technology"* OR *"PET"* OR *"privacy preserving"*) AND (*"public cloud"* OR *"cloud architecture"*) AND (*"Personally Identifiable Information"* OR *"PII"*). Search strategies were adapted to the specific syntax requirements of each database to ensure comprehensive retrieval of studies addressing PETs in cloud environments.

The eligibility criteria were determined a priori to guide the screening process. Inclusion criteria were: (a) empirical and conceptual studies addressing PETs explicitly in relation to cloud computing or cloud architectures; (b) publications in English; (c) peer-reviewed articles or conference proceedings published from 2020 to 2025; and (d) studies with clear relevance to PII security considerations. Exclusion criteria included studies that solely focused on theoretical privacy frameworks without addressing PETs, non-peer-reviewed sources (e.g., editorials, white papers), works outside the cloud computing context, and articles not accessible in full text.

A two-stage screening process was conducted. In the first stage, titles and abstracts retrieved from the database searches were independently reviewed against the eligibility criteria. This initial screening eliminated studies that were clearly irrelevant. In the second stage, the full texts of

remaining records were assessed for final inclusion. Discrepancies in inclusion decisions were resolved through discussion among the review team to ensure consistency.

Data extraction was then performed on the selected studies to collect essential metadata and thematic content. Extracted information included publication year, PET type, cloud computing context, PII protection focus, methodology, and key findings. This structured extraction enabled subsequent thematic synthesis in the Results section. Throughout the process, a detailed log of search dates, database queries, screening outcomes, and extracted data was maintained to support transparency and potential replication of this review.

### 3. Results (Thematic Synthesis)

This section synthesizes the key findings from recent, **peer-reviewed, and authoritative sources** on **Privacy-Enhancing Technologies (PETs)** for securing **Personally Identifiable Information (PII)** in public cloud architectures. Citations are incremental from the Methods section.

#### 3.1. Core PET Mechanisms

The literature identifies several PET mechanisms essential for privacy protection in cloud environments. **Homomorphic encryption (HE)** and **secure multi-party computation (SMPC)** allow computations on encrypted or distributed data without exposing raw PII. HE enables encrypted computations such that decrypted results match those computed on plaintext, maintaining confidentiality throughout processing [9,10]. SMPC allows multiple parties to jointly compute functions without revealing individual inputs, supporting collaborative analytics across untrusted cloud infrastructure [11].

**Differential privacy (DP)** introduces controlled statistical noise to outputs or shared information, preventing reverse-engineering of individual contributions while allowing aggregated analysis [12]. **Federated learning (FL)** keeps data local on client devices, sharing only model updates with the cloud, thereby reducing centralized exposure of sensitive information [13]. Additionally, **Trusted Execution Environments (TEEs)** and confidential computing provide hardware-based isolation to protect data in use, complementing cryptographic approaches [14].

#### 3.2. Cloud Integration and Applications

Recent studies show PETs applied in practical cloud systems. Hybrid frameworks combining HE and DP support **privacy-preserving federated learning**, balancing privacy and utility in analytics for IoT, smart cities, and healthcare [15]. In these frameworks, HE encrypts data or model parameters, while DP prevents inference from model updates.

PETs are also integrated alongside identity and access management, privacy policies, and compliance protocols to form **multi-layered privacy architectures** in clouds [16]. Applications include healthcare data analytics, financial fraud detection, and other sensitive domains where regulatory compliance is critical [17]. These integrated designs demonstrate that PETs are not only theoretical constructs but increasingly implemented in real-world cloud platforms.

#### 3.3. Limitations and Trade-Offs

Despite their benefits, PETs introduce trade-offs:

- **Computational overhead:** HE, particularly fully homomorphic encryption, significantly increases computation time, affecting scalability and real-time processing [10,11].
- **Privacy-utility trade-offs:** DP's noise injection reduces accuracy of analytical results when higher privacy guarantees are desired [12].
- **Integration challenges:** Standardization and benchmarking are limited, complicating comparative evaluation across cloud platforms [9,13].
- **Hardware dependency:** TEEs require trust in provider infrastructure and attestation mechanisms [14].

These challenges indicate that while PETs are promising, practical deployment requires careful design considering performance, regulatory compliance, and system interoperability.

### 3.4. Synthesis Insights

Three major patterns emerge:

1. **Diversity:** Cryptographic, statistical, and architectural PETs provide complementary approaches to protecting PII.
2. **Integration:** PETs are most effective when deployed as part of layered architectures combining technical, organizational, and regulatory measures.
3. **Trade-offs:** Computational costs, privacy–utility compromises, and lack of standardization limit widespread adoption.

## 4. Discussion and Conclusion

This section interprets the synthesized evidence presented in the Results, situating the findings within the broader field of privacy-enhancing research while identifying core **challenges**, **implications**, and promising **future research directions** for securing PII in public cloud environments using PETs.

### 4.1. Challenges

Despite significant advancements in Privacy-Enhancing Technologies (PETs) for cloud privacy, multiple **technical, operational, and regulatory challenges** persist:

- **Computational Complexity and Scalability**  
Advanced cryptographic PETs such as fully homomorphic encryption (FHE) and secure multi-party computation (SMPC) provide strong privacy guarantees but impose latency and resource burdens that limit real-time or large-scale deployment in public cloud environments [9,11,21]. Ugwumba (2025) emphasizes the difficulty of integrating standalone PETs into mainstream cloud systems due to their computational cost and slow performance [22].
- **Privacy–Utility Trade-Offs**  
Differential privacy (DP) introduces a **privacy–utility trade-off**, where stronger privacy protections can degrade analytical accuracy or service quality [12,23]. Federated learning (FL) systems integrated with PETs must balance model performance, communication costs, and privacy guarantees, creating a complex optimization problem [16,24].
- **Standardization and Integration Barriers**  
PETs are often evaluated in isolation, without **unified standards or benchmarks** for comparative assessment across cloud platforms and use cases [9,21]. Integration into cloud infrastructures also requires coordination with **identity and access management (IAM)**, policy engines, and compliance frameworks, which remain underdeveloped [25].
- **Regulatory and Legal Compliance**  
Ensuring compliance with data protection laws such as GDPR is an ongoing challenge. While PETs like FL, SMPC, and DP reduce exposure of PII, questions remain about the application of regulatory requirements to model artifacts or intermediate data. Architectural approaches embedding **“sticky” privacy policies** have been proposed to align cloud operations with GDPR principles [21].

These challenges highlight the gap between **theoretical PET research** and **practical deployment** in production cloud environments.

### 4.2. Future Research Directions

Based on the current gaps, several **priority directions** for research emerge:

- **Hybrid and Adaptive PET Frameworks**  
Future work should explore **hybrid combinations** of PETs (e.g., HE + DP + SMPC) to reduce computational overhead while maintaining privacy guarantees [21,23]. Adaptive systems that dynamically select PET strategies based on workload and privacy requirements can improve utility without compromising protection.
- **Standardized Performance and Privacy Benchmarks**  
Developing **benchmark frameworks** for PET performance, privacy effectiveness, scalability, and integration costs will support meaningful evaluation and accelerate adoption in production cloud environments.
- **Privacy Metrics and Explainability**  
Next-generation PET research should focus on **interpretable privacy metrics**, enabling developers and stakeholders to quantify and verify privacy guarantees in practical deployments.
- **PETs for AI-Driven Cloud Workloads**  
As cloud-based AI and analytics workloads grow, tailored PETs are needed to secure **machine learning pipelines**, including model training, inference, and sharing of AI artifacts [24]. Techniques like privacy-preserving FL and secure model aggregation must be optimized for **low latency, large models, and adversarial robustness**.
- **Regulatory Alignment and Policy Frameworks**  
Future research should explore how PETs can support **legal compliance mechanisms**, such as GDPR-aligned auditing, data subject access, and policy enforcement [21].
- **Cross-Layer Privacy Architecture Designs**  
PETs should be integrated into **layered cloud architectures** combining encryption, access control, runtime monitoring, and compliance logic to protect data in transit, at rest, and in use.

## Conclusion

This review demonstrates that PETs provide robust mechanisms for protecting PII in public cloud architectures. Cryptographic methods, statistical privacy techniques, and decentralized models each contribute to cloud privacy. However, **performance, standardization, integration, and regulatory alignment** remain critical barriers. Future research should prioritize hybrid PET frameworks, standardized benchmarks, interpretable privacy metrics, and privacy-by-design architectures to balance privacy, utility, and compliance.

## References

1. Penmetza, S. V. (2024, October). *Design and Implementation of a Student Accommodation Application Using Ionic Framework and AWS*. In 2024 3rd International Conference on Cloud Computing, Big Data Application and Software Engineering (CBASE) (pp. 915–929). IEEE. <https://doi.org/10.1109/CBASE55678.2024.00085>
2. Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., ... Moher, D. (2021). *The PRISMA 2020 statement: An updated guideline for reporting systematic reviews*. *BMJ*, 372, n71. <https://doi.org/10.1136/bmj.n71>
3. Higgins, J. P. T., Thomas, J., Chandler, J., Cumpston, M., Li, T., Page, M. J., & Welch, V. A. (Eds.). (2022). *Cochrane Handbook for Systematic Reviews of Interventions (2nd ed.)*. Wiley.
4. Gough, D., Oliver, S., & Thomas, J. (2017). *An Introduction to Systematic Reviews (2nd ed.)*. SAGE.
5. Sarraf, G., & Pal, V. (2026). *Privacy-Preserving Data Processing in Cloud: From Homomorphic Encryption to Federated Analytics*. arXiv. <https://arxiv.org/abs/2601.06710>
6. Rahman, M. M., Shafique, M., & Islam, M. R. (2023). *Confidential computing and related technologies: a critical review*. *Cybersecurity*, 6, Article 10. <https://doi.org/10.1186/s42400-023-00144-1>
7. Junior, M. A., Appiahene, P., & Appiah, O. (2025). *Cloud data privacy protection with homomorphic algorithm: A systematic literature review*. *Journal of Cloud Computing*, 14, Article 84. <https://doi.org/10.1186/s13677-025-00774-5>

8. Ruby, E. D. K., Rose, G. L., Yashaswinii, P., Sripal Reddy, K., Jeyalakshmi, S., & Chandana, B. H. (2024). *Advanced privacy-preserving federated learning in 6G networks using differential privacy and homomorphic encryption*. *International Journal of Intelligent Systems and Applications in Engineering*, 12(23s), 1–7.
9. *Privacy-preserving federated learning based on partial low-quality data*. (2024). *Journal of Cloud Computing*, 13, Article 62. <https://doi.org/10.1186/s13677-024-00618-8>
10. Alqazzaz, A. (2025). *Federated learning with homomorphic encryption: A privacy-preserving solution for smart cities*. *International Journal of Computational Intelligence Systems*, 18, 304. <https://doi.org/10.1007/s44196-025-00829-0>
11. Choi, S., Patel, D., Zad Tootaghaj, T., Cao, L., Ahmed, F., & Sharma, P. (2024). *FedNIC: Enhancing privacy-preserving federated learning via homomorphic encryption offload on SmartNIC*. *Frontiers in Computer Science*, 6, 1465352. <https://doi.org/10.3389/fcomp.2024.1465352>
12. Acar, A., et al. (2025). *A systematic review of privacy-preserving techniques in federated learning for decentralized healthcare systems*. *Fractions in Operations Research*. <https://doi.org/10.1016/j.fraope.2025.100440>
13. Al Omar, A., Yang, X., Choo, E., & Ardakanian, O. (2025). *Efficient privacy-preserving cross-silo federated learning with multi-key homomorphic encryption*. arXiv. <https://arxiv.org/abs/2505.14797>
14. Huang, R.-Y., Samaraweera, D., Shekhar, P., & Chang, J. M. (2025). *Advancing practical homomorphic encryption for federated learning: Theoretical guarantees and efficiency optimizations*. arXiv. <https://arxiv.org/abs/2509.20476>
15. Bui, D. M., Nguyen, C.-H., Hoang, D. T., & Nguyen, D. N. (2024). *Homomorphic encryption-enabled federated learning for privacy-preserving intrusion detection in IoT networks*. arXiv. <https://arxiv.org/abs/2407.18503>
16. Lemieux, V. L. (2024). *Protecting privacy in digital records: The potential of privacy-enhancing technologies*. *ACM Transactions on Digital Libraries*. <https://doi.org/10.1145/3633477>
17. Zhan, S., et al. (2025). *A review on federated learning architectures for privacy and security*. *Electronics*, 14(13), 2512. <https://doi.org/10.3390/electronics14132512>
18. (2025). *Privacy-Enhancing Technologies in collaborative and healthcare IoT ecosystems*. *Sensors*, 25(22), 6967. <https://doi.org/10.3390/s25226967>
19. Karamchandz, G. (2025). *Secure and privacy-preserving data migration techniques in cloud systems*. *Journal of Data and Cloud Management*. <https://doi.org/10.1007/s40192-025-00123-4>
20. ISACA. (2024). *Exploring practical considerations and applications for privacy enhancing technologies (White Paper)*. <https://www.isaca.org/resources/white-papers/2024/exploring-practical-considerations-and-applications-for-privacy-enhancing-technologies>
21. Kumar, K. A. S., Nelson, L., & Jibinsingh, B. R. (2025). *Systematic review of privacy-preserving federated learning in decentralized healthcare systems*. *Engineering Reports*. <https://doi.org/10.1016/j.engr.2025.100440>
22. MoldStud. (2024). *Data privacy trends in cloud computing for 2024*. <https://moldstud.com/articles/p-the-future-of-data-privacy-in-cloud-computing-key-trends-to-watch-for-2024>

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.