



Centre for Information Policy Leadership

HUNTON ANDREWS KURTH

Privacy-Enhancing and Privacy-Preserving Technologies: Understanding the Role of PETs and PPTs in the Digital Age

December 2023

Table of Contents

Foreword	8
Executive Summary	9
1. Introduction and Methodology	13
A. Introduction.....	13
B. Methodology	13
2. Review of Policy and Regulatory Action on PETs	14
A. Canada.....	14
B. European Union	15
C. Singapore.....	15
D. Spain.....	15
E. United Kingdom	16
F. United States	16
G. International Organizations.....	17
3. Adoption and Deployment of PETs within Organizations	18
A. Incentives for Adopting PETs	18
B. Proponents for the Adoption of PETs within Organizations	19
C. Process for Deciding Whether to Develop and Implement PETs	19
D. Challenges during Implementation and Development	20
E. Measuring the Success of PETs	20
F. Roadblocks to Adoption of PETs.....	21
4. Support of Data Protection Principles	22

Table of Contents (continued)

A. PETs can facilitate secure data processing	22
B. PETs can support data minimization and purpose limitation principles	22
C. Organizations can demonstrate accountability by investing in and using PETs	23
D. Anonymization.....	23
E. PETs can assist cross-border data transfers	23
5. The Need for Clear Guidance and Pragmatic Application of Laws.	24
6. Overview of Emerging PETs	25
A. Cryptographic Tools	25
i. Homomorphic Encryption	25
ii. Secure Multi-Party Computation	28
iii. Trusted Execution Environments	30
iv. Zero-Knowledge Proofs.....	32
B. Distributed Analytics	34
i. Federated Learning	34
C. Pseudonymization and Anonymization Tools	38
i. Differential Privacy	38
ii. Synthetic Data	41
Appendix 1—Additional Case Studies.....	44
Appendix 2—Glossary	46

List of Tables

Table 1: PETs Summary Table	11
Table 2: Different Types of Homomorphic Encryption	26
Table 3: Limitations and Potential Solutions of Homomorphic Encryption	27
Table 4: Different Types of Secure Multi-Party Computation Methods	28
Table 5: Limitations and Potential Solutions of Secure Multi-Party Computation	30
Table 6: Limitations and Potential Solutions of Trusted Execution Environments	32
Table 7: Different Types of Zero-Knowledge Proofs.....	32
Table 8: Limitations and Potential Solutions of Zero-Knowledge Proofs	33
Table 9: Different Approaches to Federated Learning.....	34
Table 10: Limitations and Potential Solutions of Federated Learning.....	37
Table 11: Different Types of Differential Privacy.....	38
Table 12: Limitations and Potential Solutions of Differential Privacy	41
Table 13: Different Types of Synthetic Data	41
Table 14: Limitations and Potential Solutions of Synthetic Data	43

List of Figures

Figure 1: Homomorphic Encryption	25
Figure 3: Trusted Execution Environment.....	30
Figure 4: Zero-Knowledge Proof.....	33
Figure 5: Federated Learning.....	34
Figure 6: Differential Privacy	38
Figure 7: Synthetic Data.....	42

List of Case Studies

CASE STUDY 1: Homomorphic encryption for national security	26
CASE STUDY 2: Homomorphic encryption for training machine learning models	26
CASE STUDY 3: Homomorphic encryption for verifying election results	27
CASE STUDY 4: Secure multi-party computation for username or password breach warnings.....	29
CASE STUDY 5: Secure multi-party computation for AI model validation	29
CASE STUDY 6: Secure multi-party computation and homomorphic encryption for personalizing health and insurance plans	29
CASE STUDY 7: Trusted execution environments for processing biometric data	31
CASE STUDY 8: Trusted execution environments for computing on joint data	31
CASE STUDY 9: Federated learning and trusted execution environments for training AI models	35
CASE STUDY 10: Federated learning and secure multi-party computation for improving speech recognition models	35
CASE STUDY 11: Federated learning for training edge devices.....	36
CASE STUDY 12: Federated learning and differential privacy for next-word prediction	36
CASE STUDY 13: Differential privacy for learning when locations are busy	39
CASE STUDY 14: Differential privacy for understanding digital inequity within a country	40

List of Case Studies (continued)

CASE STUDY 15: Differential privacy for training speech recognition models	40
CASE STUDY 16: Differential privacy for understanding economic opportunity through the measurement of social capital.....	40
CASE STUDY 17: Synthetic data and differential privacy for learning about human trafficking	42
CASE STUDY 18: Synthetic data for training and testing algorithms on biometric data	42
CASE STUDY 19: Synthetic data for detecting fraud in healthcare insurance claims	43
CASE STUDY 20: Homomorphic encryption for protecting passwords.....	44
CASE STUDY 21: Trusted execution environments for cloud computing	44
CASE STUDY 22: Synthetic data for training chatbots	45
CASE STUDY 23: Federated learning and secure multi-party computation for text selections	45
CASE STUDY 24: Federated learning and secure multi-party computation for suggesting message replies	45



Foreword

Data and transformative technologies, such as AI, are enabling revolutionary progress in our society—in industry and government agencies, in health and scientific research, and in the pursuit of aspirational social goals. As the need for data that fuels technological and societal progress increases, so must our respect for the privacy, security and safety of individuals also increase. In this context, privacy-enhancing technologies (PETs) and privacy-preserving technologies (PPTs) are gaining increased attention and investment for their ability to integrate privacy and security controls into the design and architecture of systems, technologies and products, while enabling rich analysis and beneficial uses and sharing of data. PETs and PPTs play an especially important role in addressing privacy concerns in AI and machine learning models. While PETs and PPTs are not a panacea, they do yield real benefits across a range of use cases.

This White Paper is a culmination of the Centre for Information Policy Leadership’s year-long examination of, research into, and interviews with businesses, experts and regulators on the state of the art of PETs and PPTs. In particular, this paper:

- provides insights into the types of PETs available;
- demonstrates their application through various case studies;
- explores the extent to which PETs support data protection principles and legal compliance, as well as innovation;
- discusses obstacles to the development and adoption of PETs; and
- suggests ways to overcome those obstacles.

Importantly, we call on policy and lawmakers to understand the business drivers calling for the increased use of data and to proactively incentivize and support the creation and broader implementation of PETs and PPTs.

I hope that our findings and recommendations will spark rich debate and advance the widespread adoption of these technologies so that the privacy rights of individuals can continue to be protected as businesses and governments generate new insights and benefits from the use and sharing of data.

Bojana Bellamy

President
Centre for Information Policy Leadership

Executive Summary

Privacy-enhancing technologies (PETs) and privacy-preserving technologies (PPTs) generally refer to innovations that facilitate the processing and use of data in a way that preserves the privacy of individuals whose data is being used. These technologies not only enhance privacy protections, but also maintain the informational value of data to varying degrees.

While there is no unified definition denoting a technology as a PET, there are several technologies generally understood to qualify. For the purposes of this paper, PETs are separated into three categories:

- 1. Cryptographic tools**—such as homomorphic encryption, secure multi-party computation, trusted execution environments and zero-knowledge proofs—that allow certain data elements to remain hidden while in use;
- 2. Distributed analytics tools**—such as federated learning—where data is processed at the source; and
- 3. Tools for pseudonymization and anonymization**, which would include solutions such as differential privacy or the use of synthetic data.

The table at the end of this section summarizes key findings regarding specific types of PETs falling within the general three categories mentioned above.

No single PET constitutes a silver bullet; each has its own strengths and weaknesses. Given that, PETs can be and are often used in combination depending on the use case.

Despite the increased recognition of their potential to facilitate responsible and privacy-protective data use, several hurdles prevent the widespread adoption of PETs:

- **Businesses and enforcement authorities need further education on PETs.** Many privacy professionals lack a clear understanding of the capabilities and limitations of PETs, of their benefits and risks and of their concrete uses in specific circumstances. Both business leaders and regulators require up-to-date knowledge of available technologies in relation to current business drivers in order to foster the realistic adoption and effective use of PETs.
- **PETs can be challenging to develop and costly to implement.** Not all PETs are equally well developed and mature. Some are relatively new, technically complex and resource-intensive, making them cost-prohibitive to some businesses, especially SMEs.
- **The policy and regulatory landscape for PETs is still evolving.** Few jurisdictions have provided regulatory guidance about PETs or are actively encouraging their use. As mentioned above, many regulators are still developing their understanding of the technology. To the extent a jurisdiction has not directly addressed PETs or their use within the context of the applicable legislative/regulatory landscape, legal uncertainty exists about the benefits of their adoption.

- **Lack of industry standards hinders international convergence and trust.** PET developers and other stakeholders lack a common framework or set of technical standards to foster agreement on reliable solutions and best practices. Researchers and engineers need to seek international convergence on these issues in order to build trust, thereby enabling seamless integration of privacy-enabling solutions.

Notwithstanding the above challenges, several developments are fostering wider adoption of PETs. These include the availability of open-source, off-the-shelf and customizable solutions; the promotion of innovation through government-funded competitions and a significant increase in cost-effective computing power. Moreover, some privacy regulators have taken concrete steps to support the development and use of PETs. Notably, the UK Information Commissioner’s Office (ICO) has issued guidance on how PETs relate to data protection, strongly recommending organizations to adopt PETs in specific circumstances. And Singapore’s Infocomm Media Development Authority (IMDA) has launched a PET Sandbox to support businesses facilitating experimentation with PETs.

CIPL’s Recommendations

Building on these promising initiatives, CIPL offers the following recommendations to help overcome the hurdles to widespread adoption:

- **Issue regulatory guidance and incentives regarding PETs.** Organizations respond to legal clarity. Official regulatory guidance addressing PETs in the context of specific legal obligations or concepts (such as anonymization) will certainly drive adoption, as will “safe harbors” from liability for organizations that implement PETs. By supporting such initiatives, regulators and policymakers will also incentivize greater private sector investment in fundamental and application-specific research to advance these technologies.
- **Increase education and awareness about PETs.** To achieve widespread adoption, PET developers and providers need to show tangible evidence of the value of PETs and how such technologies can facilitate responsible data use. Case studies of deployments are especially useful for this purpose. Equally, businesses must understand the limitations of PETs and the conditions that determine which PET or combination of PETs is most suitable for a given use case. Individuals whose data is being processed via PETs also need a better understanding of the technology and the protection measures put in place. This will foster further trust and digital confidence.
- **Develop industry standards for PETs.** The lack of industry standards for many PETs is an obstacle to their wider adoption. While standards do exist for some PETs (such as homomorphic encryption), other PETs (like differential privacy) are at an earlier stage of development. Industry standards would help facilitate interoperability among PETs across jurisdictions. Common frameworks would establish compatibility and consistency, enabling different PETs to communicate and work together. Standards would also help codify best practices, thereby ensuring a level of sophistication and technical reliability to foster trust in the technologies.
- **Recognize PETs as a demonstrable element of accountability.** PETs complement robust data privacy management programs that are grounded in principles of organizational accountability, such as CIPL’s Accountability Framework. By helping to mitigate risk and avoid harm, PETs support compliance efforts and demonstrate effective accountability. Organizations developing, deploying and investing in PETs are able to demonstrate their commitment to protecting privacy, while at the same time enabling beneficial uses of data in a systematic, sustainable and an accountable way.

Table 1: PETs Summary Table

PET	Use	Data Protection Principles Supported	Key Benefits	Limitations	Solutions
Homomorphic encryption	Perform computations on encrypted data without revealing data	Provides secure data processing May have potential to anonymize data, ¹ removing the need to comply with personal data protection requirements and restrictions	Secure data sharing and collaboration Secure outsourcing of computations or processing to third parties Reduces the risk of data leaks as data remains hidden	High computing resources required during encryption and processing for specific types of data and computations	Algorithmic and hardware improvements improve overall performance
				Complexity	Open-source compilers and libraries are making it easier to use and implement, as well as vendors that provide solutions that already incorporate homomorphic encryption
Secure multi-party computation	Allows computation on data held across multiple parties without revealing or transferring any data beyond the results of the computation	Can support the data minimization principle Provides secure data processing Avoids the restrictions on data sharing	Secure data sharing and collaboration	Communications between parties may allow data to be reconstructed	The use of auditing and accountability measures or other PETs, such as homomorphic encryption, can help address the risk of collusion between parties
				High communication costs, leading to scalability issues	Advances in protocol designs and techniques
				Complexity	Easy-to-use compilers and libraries are available
Trusted execution environments	Keep data and code secure during use	Provides secure data processing Supports the purpose limitation principle	Provides high level protection against attackers	Security (side-channel attacks)	Addressing side-channel attacks is difficult, but this technology can act as a great security layer and should be used in combination with other PETs to maximize security
				Trust	Attestation mechanisms
				Speed	Cloud computing can offer greater computing resources
				Lack of standards	Project communities are bringing together different stakeholders to accelerate the adoption of technologies and standards
Zero-knowledge proofs	Verify knowledge without revealing the data necessary to prove it	Supports the data minimization principle	Secure data sharing by proving certain properties of the data	Performance overhead on the parties generating and verifying the proofs	More efficient zero-knowledge proofs are being developed, and computational power is continuously increasing over time
				Security	Some zero-knowledge proofs require a trusted setup process which must be carefully implemented
				Lack of standards	Initiatives are bringing together different stakeholders to develop standards and best practices

¹ Although there are debates as to whether homomorphic encryption meets current interpretations of anonymization under European law, it is important to note the European Data Protection Board (EDPB) is in the process of revising its guidance on anonymization. Also, a recent decision of the EU General Court ([Single Resolution Board v. European Data Protection Supervisor](#) Case T-557/20)—which seemed to emphasize that determining whether data has been anonymized requires a risk-based and contextual assessment of the risk of reidentification—may be interpreted to mean that homomorphic encryption could be used to anonymize data. Discussed *infra*, Section 5.

PET	Use	Data Protection Principles Supported	Key Benefits	Limitations	Solutions
Federated Learning	Enables different parties to train models using distributed data without exposing it to each other or a third party	Eliminates centralized data collection and its associated requirements for data protection and trust in the central collection entity Supports the data minimization principle Supports compliance with international data transfers restrictions	Can help address data localization requirements by removing the need for data transfers Enables secure data sharing and collaboration for model training	High communication costs on the devices/nodes involved	Model compression techniques and communication-efficient algorithms can reduce communication requirements
				Complying with the right to be forgotten	Researchers are exploring efficient methods to remove the influence of parties on the model when required, and differential privacy can help ensure that models do not overfit to a single entity, making it easier for entities to leave the federation
				Does not provide data anonymization	Differential privacy methods can be combined with federated learning to provide a target level of anonymization, addressing membership inference and model inversion attacks
				Bias	The weight given to devices under or over-represented can be adjusted accordingly
				Fairness	Federated analytics can be used to securely learn summary information about the different parties
Differential Privacy	Prevents the identification of any individual's data	Anonymizes data to a given target, removing the need to comply with personal data protection requirements and restrictions	Offers quantifiable privacy guarantees	Lack of standards for target privacy levels	Different stakeholders need to develop guidelines and best practices around determining the right epsilon value that provides privacy protection in the relevant context
				Trade-off in data accuracy	Using adaptive privacy levels to adjust the amount of noise relative to the sensitivity
				Restricted to use cases of aggregated data	Adaptive privacy levels or local differential privacy can improve the number of appropriate use cases
				Sensitivity to outliers	Using data pre-processing techniques to identify and remove or cap outliers
				Complexity	Open-source projects and platforms are easing implementation
Synthetic data	Replace real data with artificial data to safeguard privacy	May help anonymize data, removing the need to comply with personal data protection requirements and restrictions	Secure data sharing and collaboration	Data leakage	Using differential privacy with synthetic data
				Low accuracy	Thorough use case assessment and model evaluation
				Potential information loss	Careful model design and testing
				Dependency on real-world data	Verifying the accuracy and reliability of the original data
				Bias in synthetic data	Rigorous model testing

1. Introduction and Methodology

A. Introduction

Many modern data protection laws integrate the concept of privacy by design. This means that when collecting and processing personal data, data protection legal requirements must be considered and addressed from inception. A range of technological solutions can advance privacy by design, including but not limited to PETs. Generally, such privacy technologies support organizations in adhering to the different requirements of privacy laws. For example, notice delivery tools help organizations inform individuals when they are collecting their personal data, and consent management tools help implement individuals' consent preferences.

We distinguish PETs from other privacy technologies by their focus on facilitating the processing and use of data in an inherently privacy-preserving way. These technologies not only enhance privacy protections but do so while maintaining the informational value of data to varying degrees.

CIPL has worked with leading organizations involved in the development and application of PETs to inform our understanding of the factors driving the adoption of these technologies, common use cases and the obstacles to further adoption. This white paper explores how organizations are approaching PETs, how PETs can advance data protection principles and how specific types of PETs work. We also explore potential challenges to their use, as well as possible solutions to those challenges.

B. Methodology

Our research encompassed interviews and written questionnaires with organizations building and/or deploying PETs, as well as secondary research drawing on the growing body of literature and regulatory guidance focused on PETs. Through the interviews, we sought to understand how participating organizations are approaching PETs, what role they see PETs playing in the data privacy landscape and how they are currently implementing and deploying PETs.

Organizations that participated in our research included CIPL members Amazon, Amazon Web Services, Apple, Cisco, DoorDash, Google, IBM, Mastercard, Meta, Microsoft, Telefónica, TikTok and ZoomInfo. In addition, we received input from other organizations specializing in development and deployment of PET solutions, including Duality, Enveil, Scuba Analytics and Truata.

The case studies in this paper were provided by participating organizations or drawn from publicly available sources.

PETs are often called privacy-preserving technologies (PPTs). While PPTs may be a better description of what these technologies do, both terms are widely used. We use the terms interchangeably in this paper.

2. Review of Policy and Regulatory Action on PETs

While PETs have seen a recent upsurge in attention from technologists, researchers and policymakers, they are not a new concept; the term “PETs” has been in use since at least the early 2000s.²

That said, the adoption of data privacy laws in many jurisdictions and the advent of new data-intensive technologies have expanded the interest in PETs, as organizations embrace data-driven innovation within frameworks of responsible data practices. At the same time, consumers are giving increasing importance to digital trust. They want to know how organizations intend to use their data before buying products or services.³ Accountable organizations use PETs to ensure they are good data stewards, and to earn and maintain the trust of regulators and consumers alike. Furthermore, the use of PETs enables organizations to maximize the value of data and minimize risk, by making potentially sensitive data less vulnerable to unauthorized sharing or breaches.

Consistent with these developments, PETs have also been the focus of a number of policy and regulatory initiatives across the globe in recent years. With the revolution in the development and deployment of AI technologies, many realize the potential PETs can play in delivering responsible and trusted AI systems and applications.

An overview of specific initiatives from several jurisdictions follows.

A. Canada

In November 2017, the Office of the Privacy Commissioner of Canada (OPC) issued a report on PETs highlighting the different ways PETs can protect individual privacy.⁴ The report illustrated how these technologies can help, e.g., by improving the consent process and by giving end-users more control over what personal information is used.

As a follow-up, the OPC published a blog post in April 2021, specifically examining federated learning and differential privacy.⁵ The blog concluded that these two PETs have strong potential to preserve privacy. Although the complexity of these technologies had restricted the number of business use cases at the time of writing, the OPC predicted improvements that would generate a greater uptick in the future.

2 Early PETs included the use of pseudo-identities to hide the identity of data subjects, and encryption to prevent unauthorized parties from interpreting data, available at https://www.andrewpatrick.ca/pisa/handbook/Handbook_Privacy_and_PET_final.pdf.

3 Why digital trust truly matters, McKinsey, September 12, 2022, available at <https://www.mckinsey.com/capabilities/quantumblack/our-insights/why-digital-trust-truly-matters>.

4 Privacy Enhancing Technologies—A Review of Tools and Techniques, The Office of the Privacy Commissioner of Canada (OPC), November 2017, available at https://www.priv.gc.ca/en/opc-actions-and-decisions/research/explore-privacy-research/2017/pet_201711/.

5 Privacy Tech-Know blog: Privacy Enhancing Technologies for Businesses, The Office of the Privacy Commissioner of Canada (OPC), April 12, 2021, available at <https://www.priv.gc.ca/en/blog/20210412/>.

B. European Union

The European Union has addressed the use of PETs in several contexts.

The European Union Agency for Cybersecurity (ENISA) has referenced PETs in two reports. In a report on pseudonymization techniques in January 2021,⁶ where ENISA encourages broader adoption of the described practices, zero-knowledge proofs and secure multi-party computation are included as advanced techniques. Similarly, its January 2022 report on data protection engineering highlights several PETs,⁷ including differential privacy, homomorphic encryption, secure multi-party computation, trusted execution environments, synthetic data and zero-knowledge proofs, in order to support organizations adopting a data protection by design and by default approach.

The European Data Protection Board has issued recommendations in the context of international data transfers wherein it states that secure multi-party computation is an effective “supplementary measure” to ensure compliance with the level of protection required under EU law in a particular third country.⁸

The European Commission’s legislative proposal on “European statistics on population and housing” in January 2023 explicitly supports the use of PETs when sharing data in accordance with EU law, specifically requiring the use of PETs that implement “data minimization by design.”⁹ Additionally, Article 13(3) of the proposal states that the use of PETs is preferred when sharing confidential data, and Recital 30 similarly expresses the preference of using data sharing mechanisms based on PETs rather than directly transferring data.

C. Singapore

The Infocomm Media Development Authority (IMDA) launched a PET Sandbox in July 2022 to support companies interested in exploring the application of PETs.¹⁰ In the invitation to participate, IMDA (with the support of the Personal Data Protection Commission (PDPC)) described PETs as techniques that either allow data to be shared in a different form or enable data analysis without disclosing the data.¹¹ The PET Sandbox matches use case owners with PET solution providers in order for both parties to improve their understanding of how PETs can be used to tackle real-world problems and what the technical and regulatory boundaries might be.

A year after its launch, Google partnered with the PET Sandbox to offer industry guidance to participants and to teach organizations ways to access data without using third party cookies.¹² The IMDA and PDPC intend to use their observations from these pilot projects to create guidelines and tools to support the wider adoption of PETs.

D. Spain

In September 2023, Spain’s data protection authority, the Agencia Española de Protección de Datos (AEPD), published a blog post on the use of PETs, stating that they can be used to implement privacy principles and governance policies to increase trust and data sovereignty.¹³ The AEPD noted that PETs are “dual-use technologies” for their ability to facilitate GDPR compliance, as well as stakeholders’ control and trust over the use of data.

6 Data Pseudonymisation: Advanced Techniques and Use Cases, European Union Agency for Cybersecurity (ENISA), January 28, 2021, available at <https://www.enisa.europa.eu/publications/data-pseudonymisation-advanced-techniques-and-use-cases>.

7 Data Protection Engineering, European Union Agency for Cybersecurity (ENISA), January 27, 2022, available at <https://www.enisa.europa.eu/publications/data-protection-engineering>.

8 Recommendations 01/2020 on measures that supplement transfer tools to ensure compliance with the EU level of protection of personal data, European Data Protection Board, June 18, 2021, available at https://edpb.europa.eu/our-work-tools/our-documents/recommendations/recommendations-012020-measures-supplement-transfer_en.

9 Proposal for a Regulation of the European Parliament and of the Council on European statistics on population and housing, January 20, 2023, available at <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52023PC0031>.

10 Privacy Enhancing Technologies (PETs) Sandbox, Infocomm Media Development Authority (IMDA), July 20, 2022, available at <https://www.imda.gov.sg/how-we-can-help/data-innovation/privacy-enhancing-technologies-sandbox#:~:text=IMDA%20launched%20Singapore%27s%20first%20PET,use%20cases%20and%20pilot%20PETs>.

11 Invitation to Participate in Privacy Enhancing Technology Sandbox, Infocomm Media Development Authority (IMDA), July 20, 2022, available at <https://www.imda.gov.sg/-/media/imda/files/programme/pet-sandbox/pet-sandbox-cfp.pdf>.

12 New Joint Partnership Between IMDA and Google to Help Singapore Businesses Prepare For a Privacy-First Future, 18 July, 2023, available at <https://www.pdpc.gov.sg/news-and-events/announcements/2023/07/new-joint-partnership-between-imda-and-google-to-help-singapore-businesses-prepare-for-a-privacy-first-future>.

13 Data Spaces, sovereignty and privacy by design, Agencia Española de Protección de Datos, September 28, 2023, available at <https://www.aepd.es/en/prensa-y-comunicacion/blog/data-spaces-sovereignty-and-privacy-by-design>.

E. United Kingdom

A number of UK government and expert bodies have addressed the use of PETs.

In July 2021, the Centre for Data Ethics and Innovation (CDEI), a government expert body whose purpose is to promote and support the responsible use of data and AI, published an interactive PETs Adoption Guide.¹⁴ The tool is intended to encourage the adoption of PETs by helping organizations understand how these technologies can be used to unlock new opportunities. The guide includes a repository of real-world use cases and a question-based flowchart to help decision-makers determine which PETs are most appropriate for their projects.

In June 2023, the UK Information Commissioner's Office (ICO) launched a joint project with the CDEI to develop a cost-benefit analysis tool.¹⁵ This tool is intended to help organizations interested in adopting PETs to improve their understanding of the costs and benefits involved. The ultimate goal is to encourage wider adoption of PETs by addressing common challenges in implementation.

Also in June 2023, the ICO published its final guidance on PETs after a draft guidance was released in September 2022 with an opportunity for consultation.¹⁶ The ICO defines PETs as “technologies that embody fundamental data protection principles by minimizing personal data use, maximizing information security or empowering people,”¹⁷ and makes clear that PETs can facilitate data protection by design and by default and can support data protection principles, such as data minimization and data security. The guidance aims to provide practical direction for data protection officers on their journey toward adopting PETs. Overall the ICO looks at PETs in two ways: providing *input privacy*, i.e., reducing access to personal data, and *output privacy*, i.e., reducing the risk of obtaining or inferring personal data.

In addition to describing how PETs operate and how they can mitigate risk, the guidance provides insight into the ICO's view on anonymization. In particular, the ICO argues for a flexible approach to anonymization, which it calls “effective anonymization.” This recognizes that the risk of reidentifying data need not be reduced to zero in order to anonymize the data effectively. Rather, the risk can be reduced to a sufficiently low, negligible level. Further, the ICO explains that a host of factors goes into the determination of risk, including what technical measures, such as PETs, have been applied.

The Royal Society, an independent scientific academy of the UK, has published two reports on PETs. The first report, published in March 2019, provides an overview of five selected PETs and includes an assessment of their respective maturity levels and purposes through case studies.¹⁸ The 2019 report also provides recommendations on how the UK can maximize the potential of PETs. The second report, published in January 2023 in collaboration with the Alan Turing Institute, considers the role of PETs in data governance and the barriers to adoption.¹⁹

F. United States

The United States joined forces with the UK to advance responsible innovation in PETs in the form of a multi-stage competition known as “Prize Challenges.”²⁰ The challenges, which commenced in July 2022, involved a white paper submission, prototype development and a red-teaming phase. Participants could choose to create a federated learning solution for either financial crime prevention or pandemic response capabilities. For extra points, participants could develop a solution that worked

14 Privacy Enhancing Technologies Adoption Guide, Centre for Data Ethics and Innovation (CDEI), July 14, 2021, available at <https://cdeiuk.github.io/pets-adoption-guide/>.

15 Working with the ICO to encourage the adoption of PETs, Centre for Data Ethics and Innovation (CDEI), June 20, 2023, available at <https://cdei.blog.gov.uk/2023/06/20/working-with-the-ico-to-encourage-the-adoption-of-pets/>.

16 Privacy-enhancing technologies (PETs), Information Commissioner's Office (ICO), June 2023, available at <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/data-sharing/privacy-enhancing-technologies/>.

17 Ibid.

18 Protecting privacy in practice, The Royal Society, March 2019, available at <https://royalsociety.org/-/media/policy/projects/privacy-enhancing-technologies/Protecting-privacy-in-practice.pdf?la=en-GB&hash=48A28CDF4FB012663652BE671CFFED08>.

19 From privacy to partnership, The Royal Society, January 2023, available at <https://royalsociety.org/-/media/policy/projects/privacy-enhancing-technologies/From-Privacy-to-Partnership.pdf?la=en-GB&hash=4769FEB5C984089FAB52FE7E22F379D6>.

20 Privacy-Enhancing Technologies Prize Challenges, available at <https://petsprizechallenges.com/>.

for both. The winners were announced in March 2023.²¹ The US and the UK plan to continue their collaboration through the development of tools and guidance to support the adoption of PETs technologies.

In June 2022, the White House Office of Science and Technology Policy (OSTP) invited input on how to advance the wider adoption of PETs in a responsible way.²² Based on the submissions received, the OSTP published a report in March 2023: “National Strategy to Advance Privacy-Preserving Data Sharing and Analytics.”²³ The report cited the challenges preventing widespread adoption—such as limited awareness and understanding of PETs and the lack of standards and legal clarity—but it also identified a plan of action. The priorities included the formation of a steering group to foster greater clarity on the use of such technologies, the promotion of PETs research and deployment, the advancement of training and education and the expansion of international collaboration.

When US President Joseph Biden issued an executive order to support the development of “Safe, Secure, and Trustworthy AI” in October 2023,²⁴ he called for the strengthening of privacy-preserving research and privacy-enhancing technologies. The executive order promised to support the research and development of these technologies with new federal programs and to develop guidelines for federal agencies to help evaluate their effectiveness.

G. International Organizations

With a view to accelerate PET adoption, the UN Committee of Experts on Big Data and Data Science for Official Statistics (UN-CEBD) launched the UN PET Lab in January 2022.²⁵ This international initiative consists of three pillars: experimentation, outreach and training, and support services. The objective is to advance the adoption of PETs within the community of official statistics through pilot projects, knowledge sharing and collaborations.

In February 2023, the UN also released a PETs guide²⁶ which provides an overview of the latest PETs, including their evolution, security considerations and costs. Case studies are used to illustrate the diverse range of contexts where PETs can be used, however, most of these case studies are still in the concept or pilot stage. The guide further looks at the legal and regulatory issues facing PETs, concluding that regulatory guidance is vital to advancing adoption.

The Organisation for Economic Co-operation and Development (OECD) published a report in March 2023 to help policymakers and regulators consider PETs for privacy protection and data governance.²⁷ The report describes the different types of PETs, their associated opportunities and challenges, and regulatory developments. The OECD asserts that the promise of PETs requires a reassessment of the application of regulations on data collection and processing, including a focus on how PETs protect privacy.

21 The UK winners were the University of Cambridge and STARLIT (Privitar, University College London, Cardiff University), Faculty and Featurespace. The US winners were Scarlet Pets (Rutgers University), PPML Huskies (University of Washington Tacoma, Delft University of Technology, University of Brasilia), ILLIDAN Lab (Michigan State University, University of Calgary), puffle (Carnegie Mellon University), MusCAT (Broad Institute, MIT, Harvard Business School, University of Texas Austin, University of Toronto), ZS_RDE_AI (ZS Associates).

22 Advancing a Vision for Privacy-Enhancing Technologies, Office of Science and Technology Policy (OSTP), June 28, 2022, available at <https://www.whitehouse.gov/ostp/news-updates/2022/06/28/advancing-a-vision-for-privacy-enhancing-technologies/>.

23 National Strategy To Advance Privacy-Preserving Data Sharing and Analytics, March 2023, available at <https://www.whitehouse.gov/wp-content/uploads/2023/03/National-Strategy-to-Advance-Privacy-Preserving-Data-Sharing-and-Analytics.pdf>.

24 Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence, Executive Office of the President, October 30, 2023, available at <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>.

25 UN launches first of its kind ‘privacy lab’ to unlock benefits of international data sharing, UN Committee of Experts on Big Data and Data Science for Official Statistics (UN-CEBD), January 25, 2023, available at <https://unstats.un.org/bigdata/events/2022/unsc-un-pet-lab/UN%20PET%20Lab%20-%20Press%20Release%20-%2025%20Jan%202022.pdf>.

26 The PET Guide, UN Committee of Experts on Big Data and Data Science for Official Statistics (UN-CEBD), February 2023, available at <https://unstats.un.org/bigdata/task-teams/privacy/guide/2023-UN%20PET%20Guide.pdf>.

27 Emerging privacy-enhancing technologies: Current regulatory and policy approaches, Organisation for Economic Co-operation and Development (OECD), March 2023, available at https://www.oecd-ilibrary.org/science-and-technology/emerging-privacy-enhancing-technologies_bf121be4-en?sessionid=KTCF9oxJROgkund4Em-NkYBo1KQ3CgugzIVa4335.ip-10-240-5-172.

3. Adoption and Deployment of PETs within Organizations

As part of CIPL's research, we engaged with organizations (both developers and deployers of PETs) to understand how they are addressing PETs, including their motivations for using PETs and perceived barriers to wider adoption. Our findings are set out below.

A. Incentives for Adopting PETs

Many organizations are still exploring whether and how to fully integrate PETs into their data privacy management and compliance programs as they await additional guidance from regulators on how PETs may be used to comply with legal and regulatory obligations. Others, however, especially larger, data-intensive technology companies, are early adopters and have already included PETs as part of their technical organizational measures and privacy and security by design efforts.

A number of common factors drive organizations to adopt PETs across diverse industries and use cases. These include:

- **Supporting data use consistent with data protection principles.** PETs can be instrumental in operationalizing core privacy principles, such as anonymization and data minimization. Organizations are developing and implementing PETs in ways that adhere to these principles without affecting performance. Indeed, in many cases PETs are helping to improve systems. This is particularly critical in the context of AI systems, where PETs can safeguard user data, thereby promoting trust, enabling compliance and supporting the secure deployment of AI.
- **The potential for PETs to aid compliance.** As more jurisdictions adopt data privacy laws and regulators enforce digital privacy rights, the potential for PETs to aid compliance would motivate adoption. At the same time, questions remain concerning the extent to which PETs may help to achieve compliance across different use cases and jurisdictions. Organizations will need to monitor developments for additional guidance from regulators.
- **Demonstrating accountability and responsible data practices.** In addition to supporting data protection principles, organizations have come to recognize that accountable privacy practices can give rise to a broader range of benefits, e.g., deepening customer trust and loyalty, building resiliency and competitiveness, retaining talented staff and attracting investment.²⁸ PETs can serve as an important tool to support these goals by helping to demonstrate accountability and responsible business practices.
- **Meeting customer requirements.** Some organizations are implementing PETs to address specific customer demands, e.g., for keeping customer data private when collaborating across multiple organizations or when sharing customer data to gain business insights. Indeed, where transformative technologies have enabled new uses of customer data, PETs have permitted organizations to participate in this market transformation while continuing to safeguard consumer privacy.

²⁸ CIPL and Cisco, "Business Benefits of Investing in Data Privacy Management Programs," January 2023, https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cisco-cipl_report_on_business_benefits_of_investing_in_data_privacy_management_programs_10_jan_2023_.pdf.

- **Improving business processes and technologies.** PETs can be used to facilitate collaborative analysis and data partnerships that are ethical, legal and responsible. Organizations are able to share data for the benefit of all stakeholders and improve products and services for customers. The revolution in AI technology will drive further developments in PETs, as companies seek to train, build and test algorithms to ensure reliability and safety and avoid bias and discrimination.
- **Facilitating information-sharing and breaking down silos within organizations.** PETs can enable data sharing among different teams within an organization and with third-party service providers. Such sharing can increase efficiency, cultivate new insights and ultimately contribute to an organization's success.
- **Security.** Data security professionals have encouraged the use of PETs to mitigate risks associated with potential security threats. By making data indecipherable to hackers and other adversaries, PETs can help prevent the identification of the individuals associated with the data.

B. Proponents for the Adoption of PETs within Organizations

In most organizations, **a variety of teams are involved in the process of considering the implementation of PETs.** Privacy teams, researchers, engineers, security teams and product teams often collaborate to determine utility and effective use of PETs. That said, at times, decision-making may differ according to the purpose in question. For instance, a product team may suggest a PET solution in order to enhance the performance of a particular product. Similarly, a data strategy team may suggest the use of PETs to share information within the organization. In general, however, privacy teams and data strategy teams often look across the organization to identify use cases that could benefit from the application of PETs.

C. Process for Deciding Whether to Develop and Implement PETs

The decision-making process for the development and implementation of PETs varies from organization to organization, but CIPL found the following commonalities in the decision-making process:

- **Business considerations.** First, organizations identify the problem to be solved, the business practice at issue, the impact on operations, business goals and the intended output.
- **Cost and effort.** Organizations then assess the level of effort required for implementation, including whether the PET can be built internally or externally, the extent of the technical requirements, the cost of developing and managing the tool and whether the costs and effort justify the implementation.
- **Prioritization with competing initiatives.** Organizations must also evaluate implementation in light of other privacy and business initiatives.
- **Privacy and security considerations.** Organizations examine how implementation would mitigate legal and security risks, how it would affect legal obligations and internal policies, and how it compares with best practices and industry standards.
- **Wider benefits.** Organizations address whether implementation can be applied to other/future situations and whether there are additional benefits (such as those supporting the adoption, development and deployment of AI technologies). PETs can enable data training and algorithmic training to help ensure the development of responsible AI processes and controls.

D. Challenges during Implementation and Development

As with any technology, organizations sometimes encounter challenges during the implementation and development of PETs. These include:

- **Integration with new technologies.** With the emergence of large language models, organizations have found it sometimes challenging to adapt PETs to these models, given the huge amounts of data such models demand, which in turn require high computational capabilities.
- **PETs may be use-case specific.** Some PETs are more effective in specific scenarios. For example, synthetic data may be better for building AI models than for extracting granular insights. Other PETs may work better with particular data types or with a certain volume of data. Given that technologies in the PETs family are often complementary, several technologies may be used in conjunction to achieve the desired outcomes of a given use case.
- **High technical requirements and cost.** The high technical requirements and cost of PETs can be a prohibitive factor, especially for SMEs (although resource-intensity and complexity vary widely across technologies and use cases, as discussed further below).
- **Ensuring appropriate security controls.** Bringing a new technology into a production or staging environment takes time (and testing) to ensure that the new technology does not introduce unintended security vulnerabilities.
- **Gaps in understanding.** Mapping PETs solutions to organizational policy goals remains a significant challenge, as engineers and policy professionals often have different perceptions of data privacy objectives. Organizations need to educate all business units in order to facilitate a common understanding of PETs and their potential use cases.

E. Measuring the Success of PETs

To gauge whether the implementation or use of a given PET has proven to be successful, organizations may look at several factors:

- Where a PET has been implemented for a specific product, organizations may consider the **performance metrics of the product**, such as usage, demand or customer feedback.
- Organizations may also consider **whether the PET has functioned as intended**, e.g., whether it has successfully enabled a business practice while preserving privacy and security.
- Organizations can review whether the PET has **made processes easier for the engineering team**. For example, has the use of synthetic data enabled the engineering team to resolve issues that could not be resolved with a limited data set?
- Organisations can look at the **impact of PETs on legal postures and exposures**. In particular, they may consider how PETs have enabled compliance with data privacy laws and regulations, how they have mitigated legal risks and lowered legal exposure and/or costs and whether such benefits represent a return on investment.

F. Roadblocks to Adoption of PETs

In order to foster greater use of PETs, it is important to understand the factors preventing organizations from implementing them. Our research reveals the following as the most significant barriers to adoption:

- **Education and awareness.** There is a general lack of understanding of what PETs offer. This lack of knowledge is prevalent among industry leaders, government officials, regulators, and policymakers. It extends to customers and internal teams as well. Further education is needed if widespread adoption is to be achieved. A greater awareness of PETs use cases can help demonstrate what PETs can (and cannot) do; specific use cases can also show how PETs can promote business, data, and legal objectives.
- **Regulatory guidance.** Despite ample regulatory interest in PETs, regulators have provided little guidance on PETs' potential to satisfy legal and regulatory obligations. As such, there is little incentive to implement them. Consider, for example, how many organizations would be adopting PETs if regulatory guidance had clarified that data sufficiently anonymized with the use of PETs would no longer be considered personal data. Organizations need regulatory guidance on how PETs should be used and what they can help achieve. Regulators must be attuned to the use of PETs in the context of data sharing and collaboration.
- **Complexity.** Some PETs are technically complex to implement. This has caused compatibility issues with legacy systems, which may require modifications to interoperate with the proposed PETs. For example, where legacy systems were not designed to handle encrypted or anonymized data, problems arise when trying to read or process data in that format.
- **Cost.** Some PETs can be relatively expensive to deploy, limiting their use by smaller organizations and start-ups. However, costs vary significantly by technology, and deployments are generally becoming more affordable as the technologies mature.
- **Tension between utility and privacy.** While PETs can enable privacy-protective analytics, some PETs risk obfuscating potentially useful data, therefore lowering the suitability for business implementation. This makes it difficult for organizations to determine the best balance to strike in particular use cases.
- **Reluctance for digital transformation.** Despite the opportunity offered by PETs, organizations may take a conservative approach to adopting cutting-edge technologies. Organizations that continue to use outdated systems and inefficient processes may be especially hesitant to embrace PETs. Organizations more advanced in their digital transformation process trajectory are more likely to explore and consider PETs, and vice versa.

4. Support of Data Protection Principles

The potential to support the use of data consistent with data protection principles is one of the biggest drivers for organizations developing and deploying PETs. PETs can do this in several ways. PETs can be responsive to privacy by design requirements by demonstrating that privacy and data protection were considered from the outset when processing data or when designing new technology. They may also have an impact on whether data is deemed personal and thus subject to protections specific to such data, or anonymized and non-personal. As we discuss further below, additional guidance from regulators to clarify compliance impacts is important. A pragmatic application of the law, grounded in a risk-based approach is needed: where the likelihood of a PET working effectively to protect personal data is high, the law and regulators should no longer treat this as “personal data.” Such interpretations could play a valuable role in incentivizing greater development and deployment of PETs.

Specifically, PETs have the following benefits and impact on organizations’ efforts to uphold data protection principles.

A. PETs can facilitate secure data processing

PETs can be considered the quintessential example of a technical measure that may be appropriate in a particular context to ensure the security of personal data. PETs can reduce the risk of security threats and other incidents where personal data is being accessed, used or shared in an unauthorized or unlawful way. Some PETs can even heighten security, helping organizations comply with legal requirements for appropriate technical safeguards.²⁹ For example:

- **Homomorphic encryption and secure multi-party computation** keep data hidden at all times.
- **Trusted execution environments** offer a secure setting by limiting access only to authorized entities while keeping data encrypted in storage.
- By keeping sensitive data on a device or within silos, **federated learning** reduces the risk of unauthorized access to sensitive information, as data is not transmitted outside its point of origin, nor is it left sitting in potentially vulnerable centralized data stores.

B. PETs can support data minimization and purpose limitation principles

PETs can provide technological alternatives to the use of personal data for a particular processing purpose or facilitate the exchange of necessary insights from data while simultaneously limiting the visibility, access or sharing of underlying data sets. For example:

- **Zero-knowledge proofs** provide parties with the ability to prove the truth of a mathematical statement, without enabling the verifier to prove anything else.

²⁹ EU GDPR, Article 32; HIPAA, 45 CFR Part 160.

- **Secure multi-party computation** enables different parties to collaborate without sharing any information other than the final output.
- In **federated learning**, access to data is limited at all stages and is only processed as far as necessary to produce model updates to inform a global model.

C. Organizations can demonstrate accountability by investing in and using PETs

PETs are an integral part of data privacy management programs and an example of a risk-based approach to data protection and accountability. Organizations apply PETs where this is specifically warranted or needed to address heightened risks of data processing.

Also, PETs can be used as a way for organizations to demonstrate organizational accountability. They can be instrumental in showing the measures that have been taken to protect personal data and data privacy and embed data protection principles. Organizations investing in the building or deployment of PETs go an extra mile to demonstrate their accountability and mitigation of risks. For example, trusted execution environments can illustrate the steps an organization has taken to ensure data is processed for a specified purpose³⁰ and in a secure manner.³¹

D. Anonymization

PETs may also have an impact on whether data is deemed personal—or non-personal, if anonymized—and thus whether it is subject to specific requirements. Data protection laws primarily protect personal data, and it is generally understood that these laws do not apply to anonymous data. For example, the GDPR is not applicable if data is anonymized so that it is not reasonably likely to identify an individual. By using certain PETs and techniques, organizations ensure that personal data is no longer “personal data” within the meaning of the law. Hence, such data is no longer subject to various data protection legal restrictions and requirements when being used, shared or even transferred across borders. This enables organizations to use data more freely, to avoid high compliance costs and sometimes to simply be able to proceed with a project, which would have been impossible due to data protection legal restrictions.

- Depending on the parameters used, **differential privacy** can help anonymize data. By adding sufficient random noise to the analytical outputs from datasets, this can prevent the identification of any individuals’ data involved.
- **Synthetic data** may be regarded as anonymous if individuals cannot be identified from such data. By carefully generating synthetic data, artificial data that resembles real data can be created where it is not reasonably likely to identify an individual.
- It may also be possible for **homomorphic encryption** to have the same practical results of anonymizing data. Under homomorphic encryption, unlike other encryption schemes, data is hidden at all times.³²

E. PETs can assist cross-border data transfers

PETs can offer companies a robust means to enable data transfers to third countries. In navigating the complexities of cross-border data flows, a risk-based approach is critical, allowing organizations to assess and manage potential risks associated with such transfers. By leveraging PETs and embracing a risk-based approach, organizations can strike a balance between the necessity of international data transfers and the imperative to protect individual privacy.

- **Federated learning** can facilitate cross-border data transfers by enabling model training on data without the need to transfer the raw data itself across borders.
- **Secure multi-party computation** can be used by different parties to obtain an output of a computation from their joint data without sharing the raw data.

³⁰ EU GDPR, Article 5(1)(b).

³¹ Ibid Article 5(1)(f).

³² This being said, it should be noted that although the underlying data remains encrypted, the results of the computation may disclose insights about the underlying dataset.

5. The Need for Clear Guidance and Pragmatic Application of Laws

To foster greater adoption of PETs, policymakers and regulators need to carefully consider their interpretations and guidance on the legal impact of PETs. Providing greater clarity and revisiting some past interpretations could incentivize organizations to invest more in PETs' development and deployment. At the same time, regulators should also not mandate the adoption of any particular technology; PETs should be treated as one among a range of avenues to arrive at desired privacy data protection outcomes.

Deployment within the context of the GDPR and similar laws provide a useful example. Before using the original dataset to generate synthetic datasets, the party processing the data needs to ensure it has a legal basis for this processing. In jurisdictions that follow the GDPR and similar laws, applying anonymization techniques to transform personal data into anonymous information may count as processing, therefore requiring a lawful basis and a clearly defined purpose. Legitimate interest may be the most appropriate legal basis in these circumstances; however, for special category or sensitive personal data, legitimate interest alone is not a sufficient basis for processing. Regulators could revisit this interpretation of legitimate interest to expand opportunities for organizations to protect privacy through the use of synthetic data.

Similarly, the lack of clarity around encryption and its ability to effectively anonymize data is another example where more pragmatic application of the law is desirable. Encryption is generally recognized to pseudonymize (but not anonymize) data,³³ meaning that the personal data cannot be assigned to a particular data subject without the use of additional information (such as a decryption key).³⁴

However, in a recent case, the European General Court ruled that in order to determine whether an individual is identifiable, account should be taken of all means reasonably likely to be used, and that this test must be performed from the perspective of the recipient/holder of the data.³⁵ In other words, the determination of whether data is personal or not, and anonymized or not, must be made from the point of view of the organization that has (and is using) the data—a processor or a third party service provider, or another business receiving the data set.

This ruling has been interpreted by some to mean that if the decryption key is inaccessible, then the data could be deemed anonymous.³⁶ If this position is confirmed, this would significantly incentivize organizations to invest in PETs like homomorphic encryption because it would allow for greater access to data, in turn improving business services and other technologies.

Finally, regulators can incentivize the use of PETs to facilitate safe data transfers by issuing clear guidance around circumstances in which data transferred will be deemed not subject to restrictions that apply to cross-border transfers of personal data.

33 Recommendations on shaping technology according to GDPR provisions—An overview on data pseudonymisation, European Union Agency For Network and Information Security (ENISA), November 2018, available at <https://www.enisa.europa.eu/publications/recommendations-on-shaping-technology-according-to-gdpr-provisions>.

34 EU GDPR, Article 4.

35 *Single Resolution Board v. European Data Protection Supervisor* (Case T-557/20), <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:6202T%3A2023%3A219>.

36 The current view of the EDPB is that encryption does not result in anonymization: https://ec.europa.eu/justice/article-29/documentation/opinion-recom2014/wp216_en.pdf.

6. Overview of Emerging PETs

This section sets out different PETs, explains how these different tools work, demonstrates their application through case studies and discusses potential limitations, challenges and solutions associated with the use of each.

A. Cryptographic Tools

i. Homomorphic Encryption

Homomorphic encryption enables encrypted computations to be performed on data without first having to decrypt them. With non-homomorphic encryption schemes, data must be decrypted before any computations can be performed on it. This means that there is a period of time during which the data is vulnerable to interception or other attacks. By removing this opening, homomorphic encryption can help preserve the privacy and security of sensitive data.

How it works:

- Key generation: The first step in homomorphic encryption is to generate a pair of public and private keys, and an evaluation key. The public key is used for encrypting data, while the private key is kept secret and used for decrypting the encrypted results. The evaluation key is used to perform computations on the encrypted data.
- Encryption and computation: Once the keys have been generated, the data is encrypted using the public key. The evaluation key is then used to perform computations on the encrypted data.
- Decryption: Once the computation has been performed, the private key can be used to decrypt the output and obtain the result.

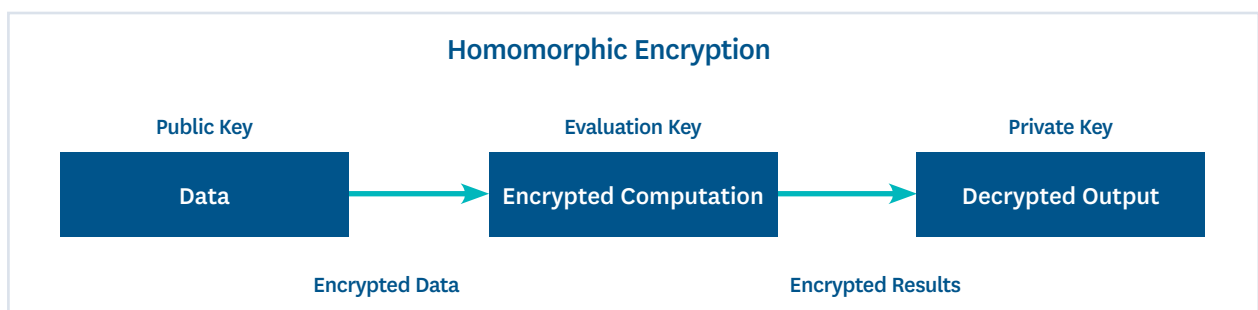


Figure 1: Homomorphic Encryption

Source: CIPL

There are three different types of homomorphic encryption: fully, somewhat and partial.

Table 2: Different Types of Homomorphic Encryption

Type	Description
Fully	The most powerful type of homomorphic encryption, allowing general computations to be performed on encrypted data. It also has no limits on the number of operations. However, its complexity means it can be computationally expensive.
Somewhat	Permits general computation up to a limit determined when the keys are generated.
Partial	Supports only specific operations and not general computation, such as addition or multiplication.

Source: CIPL

This technology supports secure data sharing and outsourcing. It can enable secure data sharing by allowing computations to be performed on encrypted data without exposing the raw data to other parties. In this way, it can also enable secure outsourcing of computations to third-party service providers. Some fully homomorphic encryption systems are also completely resistant to quantum computing attacks, providing the highest level of security.³⁷

CASE STUDY 1: Homomorphic encryption for national security³⁸

A government agency needed to examine the financial dealings of individuals associated with terrorist activities to learn whether these individuals were financing terrorism or behaved well while on parole. In order to ensure all parties complied with data privacy and data protection laws during the investigations, the government agency used homomorphic encryption to collaborate with the relevant financial institutions.

To confidentially check the financial activity of individuals on terror watchlists, the government agency created and encrypted their query, and sent this to the collaborating financial institution. Homomorphic encryption enabled the encrypted query to run on the financial institution's database and have the encrypted results sent back to the government agency where the results could be decrypted.

By leveraging homomorphic encryption, the government agency was able to achieve its national security goals while preserving privacy. Throughout the process, search terms and results were hidden. No external parties beyond the government agency itself knew what the government searched for, who was being investigated or what the results were. This not only protected the privacy of the individuals being investigated and ensured compliance with data protection and privacy laws, but also helped to avoid unnecessary de-risking.³⁹ Financial institutions did not learn which of their customers were on terror watchlists, making it easier for these individuals to integrate back into society. Importantly, as only the results of the query were returned, the financial data of innocent individuals not suspected of terrorism were not exposed, also maintaining their privacy.

CASE STUDY 2: Homomorphic encryption for training machine learning models⁴⁰

Machine learning models are increasingly being trained in the cloud, which is more cost-effective and provides for more expansive computing resources. In order to improve accuracy and safeguard privacy, a platform offers users the ability to leverage homomorphic encryption to train machine-learning models or run inference requests on the cloud using encrypted data.⁴¹

Previously, data scientists would have had to go through various time-consuming anonymization stages that could lead to lower-accuracy models. This is because, by adding random noise, blurring or stripping personally identifiable information, significant details required to create precise models can be removed. As homomorphic encryption ensures that data remain hidden at all times, no data is removed or altered.

Using homomorphic encryption to train a model in the cloud using encrypted data, while keeping the decryption key

³⁷ Mache Creeger, The Rise of Fully Homomorphic Encryption: Often called the Holy Grail of cryptography, commercial FHE is near 20(4) ACM (2022), <https://dl.acm.org/doi/10.1145/3561800>.

³⁸ Participating organization, 24 March, 2023.

³⁹ De-risking refers to financial institutions closing accounts of clients perceived as high risk for illegal financial activity.

⁴⁰ The ultimate tool for data privacy: Fully homomorphic encryption, IBM Research, December 8, 2022, available at <https://research.ibm.com/blog/fhe-cloud-security-h4cloud>.

⁴¹ IBM FHE Cloud Service, available at <https://he4cloud.com/public>.

secure and hidden from the cloud, can create a secure model that only the data owner can manipulate. This preserves privacy as the data is never exposed, and it enhances security as it protects from membership inference and model inversion attacks because only the data owner can decrypt outputs.

CASE STUDY 3: Homomorphic encryption for verifying election results⁴²

Voting is a crucial element of democratic societies, and technology can play a critical role in increasing confidence in the integrity of the voting and wider election process. Homomorphic encryption has made it possible for voters to verify that their votes in an election were counted and not changed, and that the election results are correct, while keeping their vote secure.

Homomorphic encryption allows votes to be counted while keeping it secure during the entire voting process. Individual ballots are encrypted and given a unique identifier, enabling voters to track their vote to check that it is successfully included in the final tally without being altered. The encryption ensures that no one can learn which way voters had voted. Voters can also run ballot checks which empower anyone to check that the technology is working correctly and all votes have been correctly tallied. This involves the creation of a ballot that is not counted but instead unlocked afterward to review the accuracy of the encryption.

Table 3: Limitations and Potential Solutions of Homomorphic Encryption

Challenge	Description	Solution
High computing resources required during encryption and processing for specific types of data and computations	Fully homomorphic encryption is computationally intensive, making some deployments expensive and time-consuming. Large performance overheads remain for unstructured data, such as images, and complex data science models, like large language models.	Due to technological progress, increasing computational power and algorithmic and hardware improvements, overall performance is improving.
Complexity	Homomorphic encryption requires specialized knowledge and expertise to implement. This can make it difficult for organizations to use, especially smaller organizations with limited resources.	Leading organizations are supporting wider use and adoption. ⁴³ For instance, Google has released an open-source general-purpose compiler for fully homomorphic encryption. It enables developers to write code and transform it into a form that can run on encrypted data. ⁴⁴ Duality Technologies has recently integrated Google's compiler into its open-source fully homomorphic encryption library, making the technology more accessible. ⁴⁵ However, the use of open source tools by an organization could reduce the risk of complexity, but it may also insert new security risks. A risk assessment is, therefore, needed on a case-by-case basis. There are also vendors providing solutions that already incorporate homomorphic encryption, avoiding the need to use libraries or write code. This can mitigate some of the complexity.

Source: CIPL

⁴² What is ElectionGuard?, Microsoft, March 27, 2020, available at <https://news.microsoft.com/on-the-issues/2020/03/27/what-is-electionguard/>.

⁴³ For example, IBM provides a fully homomorphic encryption toolkit for Linux, available at https://github.com/IBM/fhe-toolkit-linux?mhsrc=ibmsearch_a&mhq=fully%20encryption%20toolkit.

⁴⁴ Our latest updates on Fully Homomorphic Encryption, Google, June 14, 2021, available at <https://developers.googleblog.com/2021/06/our-latest-updates-on-fully-homomorphic-encryption.html>.

⁴⁵ OpenFHE and the Google Transpiler, Duality Technologies, September 9, 2022, available at <https://dualitytech.com/blog/openfhegoogle-transpilerfully-homomorphic-encryption-practical-reality/>.

ii. Secure Multi-Party Computation

Secure multi-party computation provides a solution to allow multiple parties to compute on their combined data, without either party revealing any information about their input data. The objective is to protect the privacy of the parties while still enabling them to perform useful computations.

Secure multi-party computation is implemented using a technique called “secret sharing.” This technique is used to split each party’s data (the secret) into different shares, which are then distributed among the other parties. The secret can only be reconstructed by combining a minimum number of shares. Each party then computes on their shares and may distribute their results to the other parties to help reach their target answer.

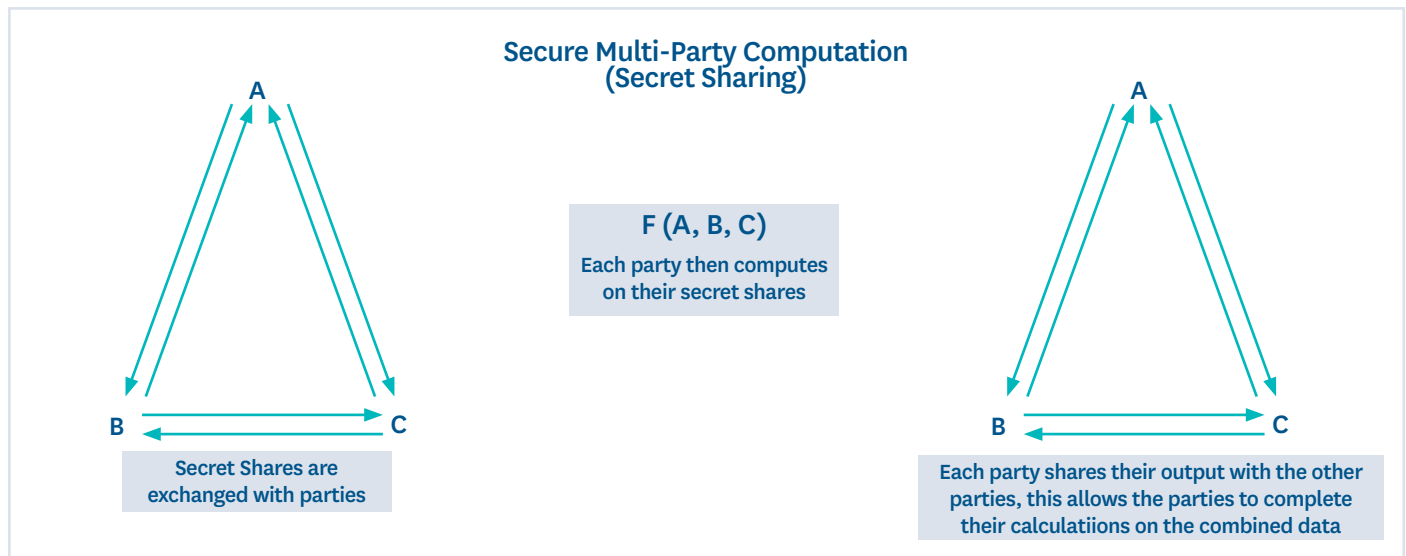


Figure 2: Secure Multi-Party Computation (Secret Sharing)

Source: CIPL

There are several different ways that secure multi-party computation has been applied; the most common are listed in the table below.

Table 4: Different Types of Secure Multi-Party Computation Methods

Method	Description
Secure aggregation protocol	This protocol enables multiple parties to aggregate their data without revealing their individual inputs. This is achieved by encrypting the input data and sharing this with the other parties. Computations are then performed on the encrypted data, using homomorphic encryption, before decrypting the final result.
Private set intersection	This technique allows parties to compare their datasets and find the common elements while keeping the rest of their data private.

Source: CIPL

Homomorphic encryption can also be used as a building block in secure multi-party computation. The parties can use homomorphic encryption to encrypt their private inputs and perform computations. Parties cannot see the initial inputs of the other parties. To obtain the results of the computations, the parties jointly decrypt the final encrypted output. This can be particularly useful in health research settings. For instance, healthcare data can be shared to enable more accurate machine learning models for disease prediction, without exposing sensitive data.⁴⁶

⁴⁶ How to Use Homomorphic Encryption in the Real World, Duality Technologies, August 8, 2022, available at <https://dualitytech.com/blog/how-to-use-homomorphic-encryption-in-the-real-world/>.

CASE STUDY 4: Secure multi-party computation for username or password breach warnings⁴⁷

In order to inform end-users that their username or password was part of a data breach, a secure multi-party computation technique can be used to determine securely whether credentials were compromised, so only the end-user learns the result. Private set intersection ensures that throughout the process, end-users' credentials are never exposed.

First, a hashed and encrypted copy of usernames and passwords are stored when they are revealed by a data breach. When a user signs into a website, the browser sends a hashed and encrypted version of the user's username and password, enabling private set intersection to then compare the user's encrypted username and password with all the encrypted, breached usernames and passwords, without revealing any information about other end-users' usernames and passwords. For additional security, the final check takes place locally on the end-user's device. If there is a match, the user is alerted to change the password.

Using secure multi-party computation, end-users are notified of data breaches without exposing their usernames or passwords at any time.

CASE STUDY 5: Secure multi-party computation for AI model validation⁴⁸

Organizations seeking to use AI developed by an external vendor may want to undertake "model validation" to test models from multiple vendors. This ensures that the best model is selected for the organization's needs. This is usually done using test data, but the two most commonly used methods for providing test data pose a risk of exposure. In the first method, the model owner could share its model data, but that risks exposure of intellectual property and could reveal training data used for the model from model inversion and membership inference attacks. The second method is for the organization to share its own data, but this, of course, may expose sensitive data. A secure multi-party framework has been developed to address these problems.

The secure multi-party computation framework requires users to first send their code, stating the computation and the parties providing the inputs. A compiler then takes this code and converts it into an efficient secure multi-party computation protocol. Through the exchange of random bits, the protocol allows the parties to accurately perform AI model validation without sharing any data. This framework has successfully been tested in the healthcare domain.⁴⁹

CASE STUDY 6: Secure multi-party computation and homomorphic encryption for personalizing health and insurance plans⁵⁰

In order to provide customers with more personalized health and insurance plans, an organization may need to collect customer data from multiple sources. This could include medical data from a hospital and physical activity data from the gym. Using a combination of secure multi-party computation and homomorphic encryption allows data from multiple institutions to be used to build models.

The data from multiple sources such as hospitals, clinics and gyms are first encrypted. Using secure multi-party computation, this data from various institutions is linked while remaining encrypted. Homomorphic encryption then allows the model to be trained while the data is encrypted.

This combination of PETs enabled customers to receive personalized care plans, providing those at greater risk of hospitalization with wellness plans.

47 Better password protections in Chrome—How it works, Google, December 10, 2019, available at <https://security.googleblog.com/2019/12/better-password-protections-in-chrome.html>.

48 EzPC: Increased data security in the AI model validation process, Microsoft, January 12, 2022, available at <https://www.microsoft.com/en-us/research/blog/ezpc-increased-data-security-in-the-ai-model-validation-process/>.

49 Multi-institution encrypted medical imaging AI validation without data sharing, Microsoft, August 2021, available at <https://www.microsoft.com/en-us/research/publication/multi-institution-encrypted-medical-imaging-ai-validation-without-data-sharing/>.

50 Duality Empowers NTT DATA to Revolutionize Privacy Enhanced Secure Data Collaboration, Duality Technologies, April 4, 2022, available at <https://dualitytech.com/blog/duality-empowers-ntt-data-to-revolutionize-privacy-enhanced-secure-data-collaboration/>.

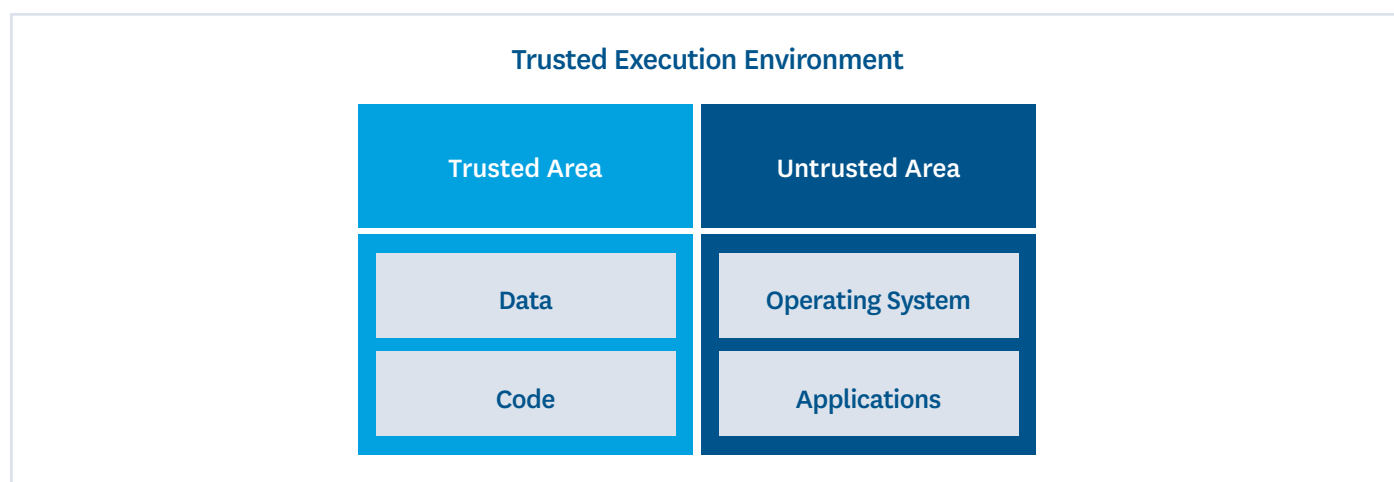
Table 5: Limitations and Potential Solutions of Secure Multi-Party Computation

Challenge	Description	Solution
Communications between parties may allow data to be reconstructed	There is the risk of collusion. For example, secret sharing may make it possible for the input data to be reconstructed if parties secretly communicate.	Auditing and accountability measures can deter collusion attempts. By logging and monitoring actions taken by the parties, malicious or colluding behavior can be detected and investigated. Furthermore, homomorphic encryption can be used to address this challenge by ensuring the data is kept hidden at all times.
High communication costs leading to scalability issues	Depending on the method used, secure multi-party computation can lead to high communication costs and, therefore, scalability issues. For example, in a garbled circuit protocol, the communication scales with computation.	Data reduction techniques can reduce the size of the inputs or intermediate results that need to be communicated. For example, data compression algorithms can be used to compress the data before transmission to reduce communication costs.
Complexity	Implementing and deploying secure multi-party computation protocols correctly can require expertise and employment of skilled professionals.	New compilers are being released to foster adoption, such as CrypTen. ⁵¹ This software framework facilitates secure multi-party computation for machine learning by enabling machine learning researchers to easily experiment with secure multi-party computation techniques.

Source: CIPL

iii. Trusted Execution Environments

A trusted execution environment is a secure and isolated area within a computing system that provides a platform for running code and accessing data in a protected way. Applications running outside the trusted execution environment cannot access data within it, but applications running inside the trusted execution environment can access the data outside it.

**Figure 3: Trusted Execution Environment**

Source: CIPL

A trusted execution environment enables confidential data processing and computing. Users can specify all the hardware and software that can have access to their data and code. A significant aspect of trusted execution environments is that they provide attestation mechanisms to remotely verify that the user's security and privacy requests have been addressed. This can support the purpose limitation principle.⁵² This requires data to be used only for the purpose clearly set out from the start. By restricting the data only to permitted applications and code, a trusted execution environment safeguards against others accessing and processing the data in unauthorized ways.

⁵¹ CrypTen, available at <https://crypten.ai/>.

⁵² EU GDPR, Article 5(1)(b).

Trusted execution environments are commonly used with cloud computing to enhance security and privacy in the cloud. The cloud, in turn, offers greater computing resources for trusted execution environments. Trusted execution environments allow code and data to remain hidden even from the cloud provider.

Trusted execution environments also offer a setting that can be leveraged by other PETs:

- Secure multi-party computation allows for different parties to collaborate in a trusted execution environment while attestation mechanisms ensure that all collaborators must approve before specific steps are performed.
- Homomorphic encryption can add an extra layer of protection to trusted execution environments. Without homomorphic encryption, data is usually encrypted in transit and at rest, but is decrypted for computations. With homomorphic encryption, data can remain encrypted at all times, even when performing calculations on the data.

CASE STUDY 7: Trusted execution environments for processing biometric data⁵³

Today, secure authentication is often performed using biometric data; individuals can unlock devices with either a finger press or a glance. This gives end-users easier and faster access to their devices and can be more secure than ordinary passwords. However, biometric data must be kept secure, as it can put end-users at risk of identity-based attacks. And while end-users can change normal passwords, they cannot change their fingerprints or faces. In order to keep this biometric data secure, a type of trusted execution environment called a secure enclave can be used to store user fingerprints and facial data.

When a user sets up biometric authentication on the device, the biometric sensor captures the biometric image from the user and sends this to the secure enclave where it is encrypted and stored. When a user then wishes to unlock the device, the secure enclave compares the new data coming from the biometric sensor against the stored data to decide whether to unlock the device.

The biometric data stored in the secure enclave is stored without any identification information, and the data never leaves the user's device. The data also cannot be accessed by the device's operating system or applications. The trusted execution environment provides a high level of privacy and security, enabling the use of sensitive data to make authentication easier and more secure.

CASE STUDY 8: Trusted execution environments for computing on joint data⁵⁴

Organizations often wish to enrich their existing data with data belonging to other organizations. However, security and privacy requirements make this difficult. In order to overcome this difficulty, a trusted execution environment can be used to enable one party to compute on the joint data of multiple parties, without revealing any data.

To illustrate, a bidding service wishes to securely process two parties' data. The trusted execution environment uses an attestation document to verify its identity and ensure that only authorized code is running. The parties then encrypt their respective bids and store them in their own accounts. The bidding service then runs the application, retrieving the encrypted bids from each parties' account and decrypts the bids. This allows the bidding service to determine the highest bidder without disclosing their values to any party.

By using a trusted execution environment, access to sensitive data is only allowed from the application running within it, safeguarding privacy and security.

⁵³ Face ID and Touch ID security, Apple, February 18, 2021, available at <https://support.apple.com/en-gb/guide/security/seco67eb0c9e/1/web/1>.

⁵⁴ Use AWS Nitro Enclaves to perform computation of multiple sensitive datasets, Amazon Web Services, June 29, 2022, available at <https://aws.amazon.com/blogs/compute/leveraging-aws-nitro-enclaves-to-perform-computation-of-multiple-sensitive-datasets/>.

Table 6: Limitations and Potential Solutions of Trusted Execution Environments

Challenge	Description	Solution
Security (side-channel attacks)	Trusted execution environments are susceptible to side-channel attacks. For example, information may be revealed from the way the trusted execution environment communicates with other parts of the computer. Timing attacks are the most common and refer to attacks based on measuring how much time various computations take to perform.	Addressing side-channel attacks is difficult and the most robust techniques impose significant performance costs. The technology is not foolproof, but it can act as a valuable security layer and can be used in combination with other PETs to maximize security.
Trust	Trust is placed in the cloud service provider or the specific computer system, rather than a mathematical formula that has guarantees.	Attestation can help prove the security of a trusted execution environment. It confirms that the code is executing inside the secure environment of the PET.
Speed	Application may be slow due to the segmentation and extra steps required, as opposed to a computer system that does not implement a trusted execution environment.	Cloud computing offers greater computing resources and can speed up the use of trusted execution environments while the PET simultaneously enhances the security and privacy of the cloud itself.
Lack of standards	There are no formal standards that describe what a trusted execution environment is, how different trusted execution environments should interact with each other or the best attestation mechanisms.	The Confidential Computing Consortium brings together hardware vendors, cloud providers and software developers to develop and drive adoption of solutions and standards for more secure confidential computing. ⁵⁵

Source: CIPL

iv. Zero-Knowledge Proofs

Zero-knowledge proof is a technique that enables one party (the prover) to prove a claim to another party (the verifier) without revealing anything more than the truth of the claim. Through the use of complex mathematical algorithms, the proof is generated in such a way that it is computationally infeasible for someone who does not know the claim to generate a similar or related proof.

A zero-knowledge proof has three main properties:

- **Completeness:** If the claim being proved is true, then an honest verifier will be convinced of this fact with high probability.
- **Soundness:** If the prover does not know the claim, then he cannot deceive the verifier with high probability.
- **Zero-knowledge:** The verifier does not learn anything other than the validity of the claim.

There are two main types of zero-knowledge proofs: interactive and non-interactive.

Table 7: Different Types of Zero-Knowledge Proofs

Type	Description
Interactive	The prover and verifier engage in a number of interactions where the verifier questions the prover, and the prover provides responses to these queries. These interactions are repeated until the verifier is convinced of the validity of the prover's claim. The disadvantage of this type of zero-knowledge proof is that it can be slow and inefficient when the proof is complex and when communications must be performed with each new verifier. For this reason, interactive methods are more appropriate when the claim requires little computational effort to verify or when there are few verifiers.
Non-interactive	There is no interaction between the prover and verifier, as the proof is self-contained and can, therefore, be independently verified. Although this method does not have multiple iterations and can, therefore, be faster, it does require greater computing power.

Source: CIPL

⁵⁵ Confidential Computing Consortium, available at <https://confidentialcomputing.io/>.

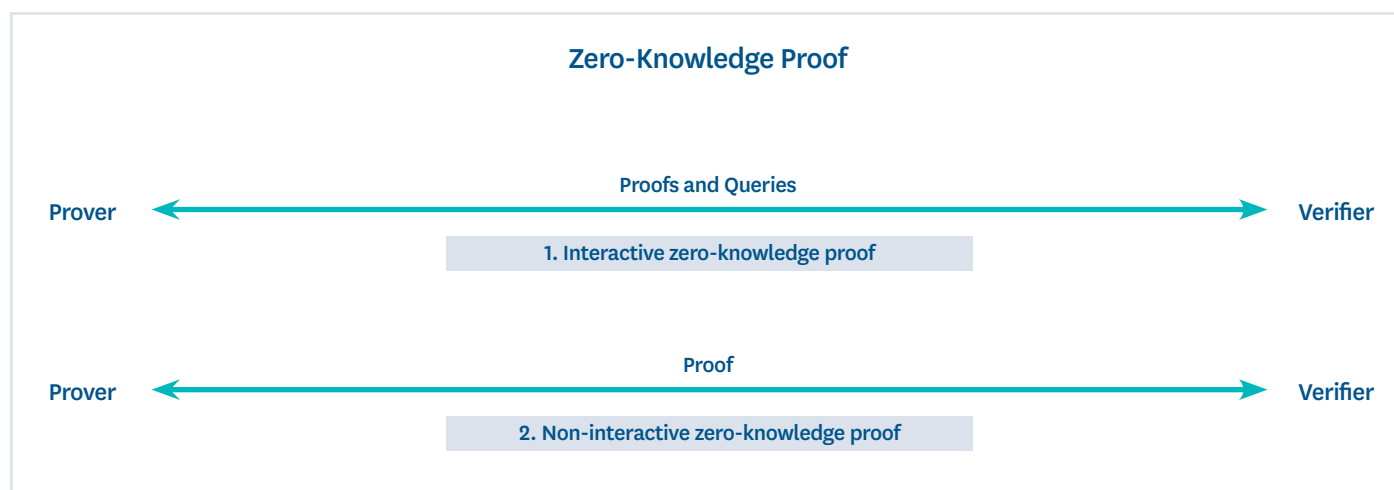


Figure 4: Zero-Knowledge Proof

Source: CIPL

Zero-knowledge proofs can offer a mechanism to prove the correctness and integrity of the performance of other PETs:

- By using zero-knowledge proof with homomorphic encryption, it becomes possible to verify the correctness of computations on encrypted data without revealing the data.
- Zero-knowledge proof can be used with secure multi-party computation to prove specific properties of private inputs without revealing the inputs themselves.

There are fewer deployment examples to date for zero-knowledge proofs compared to other PETs, but sources we consulted with for this paper see potential applications to blockchain, cryptocurrency and decentralized finance. Zero-knowledge proofs can help prove that an account balance is adequate for a transaction without revealing the balance, can enable transactions to be validated without accessing transaction data and can verify the identity of users without revealing information beyond what is necessary to establish trust, thereby helping to prevent identity theft and fraud. Zero-knowledge proofs are earlier in the development-deployment lifecycle and can be challenging to implement.

Table 8: Limitations and Potential Solutions of Zero-Knowledge Proofs

Challenge	Description	Solution
Performance overhead on the parties generating and verifying the proofs	Algorithms need intense computational resources and may slow down transaction processing times. Some require many interactions between verifiers and provers whereas others are very computationally intense.	Researchers are constantly developing more efficient zero-knowledge proofs, and as computational power increases over time, this should not be a significant issue.
Security	Some zero-knowledge proofs require a “trusted setup” where private and public keys are generated to create and verify the proof. If the numbers used to create the keys are not destroyed, they could be used to create false proofs.	Careful implementation and management of the trusted setup process are vital.
Lack of standards	Currently, there are no international standards for the use of zero-knowledge proofs. Without standards, widespread use of the technology will be restricted, as it limits interoperability and reliability.	There are initiatives underway with the aim of standardizing zero-knowledge proofs. For example, ZKProof brings together researchers, practitioners and industry experts to develop standards and best practices for designing and implementing zero-knowledge proofs. ⁵⁶

Source: CIPL

⁵⁶ ZKProof, available at <https://zkproof.org/>.

B. Distributed Analytics

i. Federated Learning

Federated learning is a technique that enables different parties to train a shared machine learning model without sharing their data. Each party trains the model on its own device using its own data, and then sends model updates to inform a single global model. The significance of federated learning is that, in contrast to traditional machine learning model training, data is neither collected nor stored in one location. Federated learning enhances privacy because the raw data is never shared. Moreover, security is strengthened by removing the need to store data in one place.

There are generally two approaches to federated learning: “server-coordinated” and “fully decentralized.” The first model is more conventional, involving a central coordinating server. The second model entails different data owners coordinating directly among themselves.

Table 9: Different Approaches to Federated Learning

	Server-Coordinated Federated Learning	Fully Decentralized Federated Learning
How it works	The server sends out a copy of the machine learning model to each party. The parties train their model using their own data, updating the model parameters in the process. The updated model parameters from each party are sent back to the server, which aggregates them into a new global model. This process repeats itself several times with the server broadcasting the newly updated global model until a satisfactory level of accuracy is reached.	In this model, as no central server is involved, each party exchanges its model updates with each other. A consensus mechanism, such as averaging, is used to combine the updates and create the updated global model. Like the server-coordinated approach, the process is repeated until the desired accuracy is achieved.
Advantages	<p>Ease of management: The single point of control makes the training process easy to manage.</p> <p>Scalability: It can accommodate large amounts of data and participants.</p> <p>Accuracy: The quality of data and participants can be regulated, leading to improved model accuracy.</p>	Resilience: With multiple points of control, there is no single point of failure.

Source: CIPL

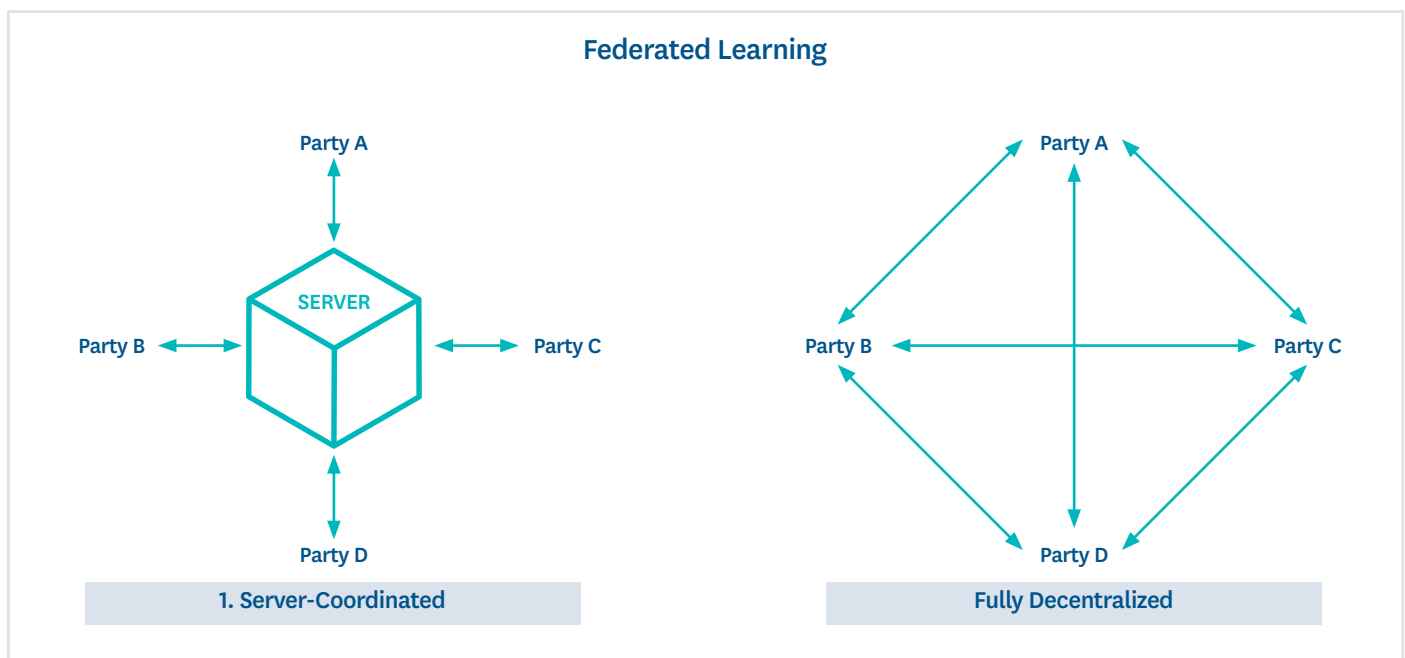


Figure 5: Federated Learning

Source: CIPL

Federated learning can also be combined with other PETs to further enhance privacy.

- Homomorphic encryption can be used with federated learning to encrypt the model updates before they are shared with the central server. This means that neither the server nor malicious actors can get access to this data.
- Differential privacy can be used to prevent reconstruction of private data from model updates and/or from the final aggregates, depending on the stage of data processing where protective noise is applied.
- Secure multi-party computation can be used to hide model updates through various encryption schemes. For example, with secure aggregation, each device computes its own update, uses secret sharing to encrypt it and sends it to the central server.⁵⁷ The central server aggregates the encrypted updates and decrypts the final result to obtain the updated model parameters. This means that no individual party's update can be inspected before aggregation.
- Trusted execution environments can provide additional protections for federated learning through attestation and verification of the computation process for updates (in data silos) and/or updates aggregation (in the centralized coordination server).

CASE STUDY 9: Federated learning and trusted execution environments for training AI models⁵⁸

Trusted execution environments can be leveraged to enhance the security of federated learning processes. Trusted execution environments enable data owners to control who can use their data and for which specific tasks, while keeping the data secure. This reduces the need for data owners to ensure trust in a centralized coordinator server.

The aggregation server can also be run in a trusted execution environment, enabling the proper execution of code with attestation mechanisms and verifying that data is being used for the intended purpose.

Local training can also take place in a trusted execution environment, ensuring that each party is training its model correctly. This can improve transparency and accountability and prevent participants from using biased data or purposely corrupting the training process.

CASE STUDY 10: Federated learning and secure multi-party computation for improving speech recognition models⁵⁹

Many devices, such as phones, laptops and voice assistants, use speech recognition technologies. These technologies use deep neural networks to create faster and more accurate systems, but they require large amounts of data to function and improve. To address this issue, federated learning can be used to process greater amounts of user data in a privacy-preserving way, facilitating improvement of the models.

For federated learning to take place, audio data is first saved onto the user's device. An AI model on the device is then trained on this data. The model changes are sent to the organization's servers, where these changes are aggregated with updates from other participating devices using secure aggregation, a secure multi-party computation technique. By aggregating these updates, the model is improved for the benefit of all end-users. Furthermore, it protects the privacy of end-users by preventing the organization from learning the audio data contents or model changes from individual devices.

⁵⁷ Federated Learning: Collaborative Machine Learning without Centralized Training Data, Google, April 6, 2017, available at <https://ai.googleblog.com/2017/04/federated-learning-collaborative.html>.

⁵⁸ Unlocking the potential of Privacy-Preserving AI with Azure Confidential Computing on NVIDIA H100, Microsoft, March 24, 2023, available at <https://techcommunity.microsoft.com/t5/azure-confidential-computing/unlocking-the-potential-of-privacy-preserving-ai-with-azure/ba-p/3776838>.

⁵⁹ Learn how Google improves speech models, Google, available at <https://support.google.com/assistant/answer/11140942?hl=en#zippy=%2Cfederated-learning>.

CASE STUDY 11: Federated learning for training edge devices⁶⁰

Server-coordinated federated learning is used to train edge devices, such as digital assistants, on data from multiple devices without moving the data from individuals' devices.

A cloud server is used to orchestrate the federated learning process by sending the current global model to the parties, where they each train their local models in order to send back model parameters to the server. The server then aggregates these parameters to update the global model. Adaptive aggregation can be used for improved personalization. Whereas most approaches take into account each party's dataset size when aggregating the model parameters, this adaptive method means that where a local model's training data departs from the global averages, more weight is given to it. This has enabled better tailoring of devices to the types of data they are likely to see, improving user experiences while preserving privacy.

CASE STUDY 12: Federated learning and differential privacy for next-word prediction⁶¹

Many devices use next-word prediction, which can help reduce keystrokes and minimize spelling mistakes. Federated learning can be used with differential privacy to protect user privacy, while improving the on-device accuracy of this feature.

Raw data remains on user devices and is used to train local machine learning models. Model updates are sent from the devices to inform a global model, which empowers the technology behind next-word prediction across all devices.

Differential privacy can further protect the data of each device. By adding random noise during training, differential privacy ensures that next-word prediction would not meaningfully change even if the training data from any specific device were removed. The combination of these two PETs enhances user privacy when using data from multiple devices to train machine learning models.

60 Personalized federated learning for a better customer experience, Amazon, December 5, 2022, available at <https://www.amazon.science/blog/personalized-federated-learning-for-a-better-customer-experience>.

61 Federated Learning with Formal Differential Privacy Guarantees, Google, February 28, 2022, available at <https://ai.googleblog.com/2022/02/federated-learning-with-formal.html>.

Table 10: Limitations and Potential Solutions of Federated Learning

Challenge	Description	Solution
High communication costs on the devices/nodes involved	The frequent communication of model updates between devices and the central server means communication overhead can be significant. This communication can also lead to increased latency and be a burden on network bandwidth.	Model compression techniques can reduce communication requirements ⁶² and the bandwidth necessary to download the current model. ⁶³ Furthermore, communication-efficient algorithms, such as Google's Federated Averaging Algorithm, can help address both bandwidth and latency limitations by computing higher quality updates. Fewer iterations of high-quality updates are required to produce a good model. ⁶⁴
Complying with the right to be forgotten	Removing the influence of a party on the central model when they leave the federation is challenging but necessary to meet requirements of privacy laws that empower end-users with the right to be forgotten. ⁶⁵	Currently, in order to remove a party's training data without compromising the model, the remaining parties are required to retrain the model from the beginning, but researchers are exploring methods to make this necessary only to the point at which the now-erased data were added. ⁶⁶ Furthermore, differential privacy could be used in combination with federated learning to train models not to overfit to a single entity. This can make it easier for entities to leave the federation.
Does not provide data anonymization	Researchers have demonstrated that by comparing the differences between a model before and after updates, information can be revealed about changes in the training data. ⁶⁷ Model parameters could also be intercepted during transmission and can be used to learn information about training data. Researchers have also shown that by querying large language models, private data used to train the model can be extracted. ⁶⁸ These models can be vulnerable to both membership inference and model inversion attacks.	It is recommended that other PETs, such as differential privacy, be used with federated learning to provide a target level of anonymization to help protect against these privacy and security concerns.
Bias	There is also a risk of biases stemming from federated learning. ⁶⁹ For example, if devices are selected for updates when their output can be computed more quickly, ⁷⁰ devices with faster processors, likely correlated with socioeconomic status, could be over-represented.	Recent efforts have begun to tackle bias. For instance, to ensure devices with more data are not over-represented, devices with less data can be given higher relative weight to encourage less variance in the final accuracy distribution. ⁷¹
Fairness	Having the ability to access and analyze sensitive information is often required to ensure individuals are treated fairly, but federated learning obfuscates such distinctions. ⁷²	Federated analytics can be used to securely perform computations on data held on devices, including sensitive personal data, without transferring data. ⁷³ The result from the computation is sent back to the server where secure aggregation encrypts and combines the outputs from all devices in order to safely learn the results. Federated analytics can help learn summary information about the data on devices that could be used to promote fairness.

Source: CIPL

62 Compressed-VFL: Communication-Efficient Learning with Vertically Partitioned Data, IBM, July 17, 2022, available at <https://research.ibm.com/publications/compressed-vfl-communication-efficient-learning-with-vertically-partitioned-data>; Jakub Konečný et al., *Federated Learning: Strategies for Improving Communication Efficiency* (2016), <https://arxiv.org/abs/1610.05492>.

63 Song Han, et al., *Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding* (2015), <https://arxiv.org/abs/1510.00149>.

64 H Brendan McMahan, et al., *Communication-Efficient Learning of Deep Networks from Decentralized Data* 54 PMLR 1273 (2016), <https://proceedings.mlr.press/v54/mcmahan17a.html>.

65 Ibid.

66 Federated Unlearning: How to Efficiently Erase a Client in FL?, IBM, July 17, 2022, available at <https://research.ibm.com/publications/federated-unlearning-how-to-efficiently-erase-a-client-in-fl>.

67 Analyzing Information Leakage of Updates to Natural Language Models, Microsoft, November 2020, available at <https://www.microsoft.com/en-us/research/publication/analyzing-information-leakage-of-updates-to-natural-language-models/>.

68 Nicholas Carlini, et al., *Extracting Training Data from Large Language Model* (2021), <https://arxiv.org/abs/2012.07805>.

69 Peter Kairouz, et al., *Advances and Open Problems in Federated Learning* (2021), <https://arxiv.org/abs/1912.04977>.

70 Takayuki Nishio and Ryo Yonetani, *Client selection for federated learning with heterogeneous resources in mobile edge* (2019), <https://arxiv.org/abs/1804.08333>.

71 Tian Li, et al., *Fair resource allocation in federated learning* (2019), <https://arxiv.org/abs/1905.10497>.

72 Peter Kairouz, et al., *Advances and Open Problems in Federated Learning* (2021), <https://arxiv.org/abs/1912.04977>.

73 Federated Analytics: Collaborative Data Science without Data Collection, Google, May 27, 2020, available at <https://ai.googleblog.com/2020/05/federated-analytics-collaborative-data.html>.

C. Pseudonymization and Anonymization Tools

i. Differential Privacy

Differential privacy is a technical solution that uses a mathematical framework to safeguard privacy. By adding the right amount of random noise to analytical outputs from datasets, individual privacy is preserved while minimizing the trade-off on data accuracy. The purpose of differential privacy is to alter the data in a way that prevents the identification of any individuals' data involved.

In differential privacy, the privacy loss parameter or privacy budget controls the amount of noise to be added. This parameter is measured in epsilon (ϵ) and regulates the trade-off between privacy and accuracy. Smaller values lead to greater privacy but lower accuracy. For example, $\epsilon=0$ completely protects privacy at the cost of no accuracy as only noise is present. In this context, accuracy is defined as the proximity of the output from a differentially private dataset to the real output when analyzing the data.

There are two types of differential privacy: central and local.

Table 11: Different Types of Differential Privacy

Type	How it Works
Central differential privacy	In this type of differential privacy, the different parties send their data to a trusted central server which adds noise after a query is identified in order to create a differentially private output. The disadvantages of this approach are that there is a single point of failure and the parties must trust a central server with their data.
Local differential privacy	Local differential privacy has each party add noise to its own data, making this method more secure. However, because each party adds noise, in order to achieve the same level of privacy as central differential privacy, more noise is required. Therefore, this approach tends to be more appropriate for large datasets for which security and privacy take precedence over utility.

Source: CIPL

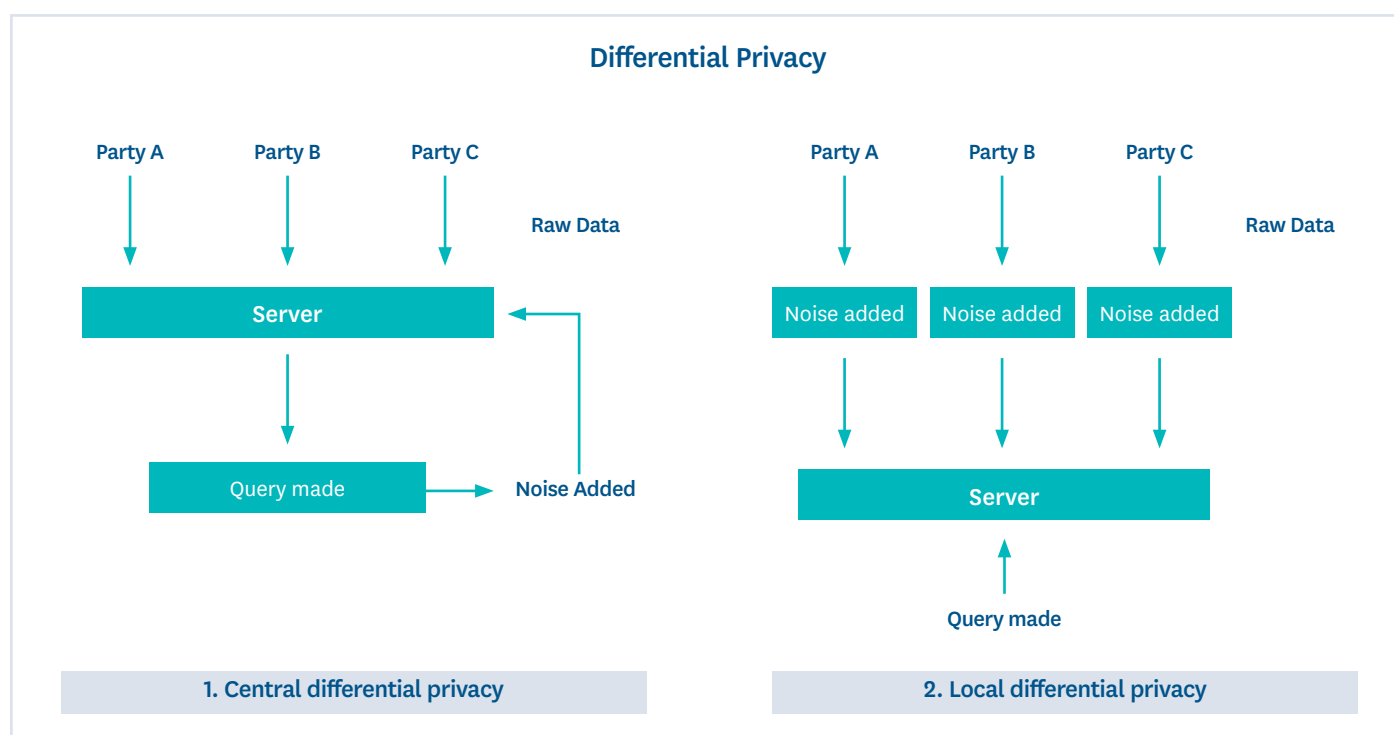


Figure 6: Differential Privacy

Source: CIPL

Differential privacy is most effective for seeking statistical queries from large datasets. For example, organizations leverage it to improve their devices, as it allows them to collect analytics while maintaining privacy.⁷⁴ It has also been used to analyze the impact of social media by measuring content interactions⁷⁵ and to provide organizations with insights into how their employees are getting work done.⁷⁶

Differential privacy is also useful in machine learning processes.⁷⁷ By adding sufficient noise during training, the contributions of individual parties are hidden. Therefore, the final machine learning model learns only general trends in the data. This offers a secure approach to model training.

It should be noted that differential privacy cannot fully remove the risk of data leaks; instead, it quantifies and manages the risk. For instance, one of the main benefits of differential privacy is its robustness under composition. Whereas malicious actors can identify individuals from multiple analyses on an anonymous dataset, differential privacy ensures that all the data extracted remains private. To illustrate, in the case of two differentially private computations, if the first computation uses ϵ_1 and the second uses ϵ_2 , the overall privacy risk is equal to $\epsilon_1 + \epsilon_2$. Accordingly, while differential privacy cannot prevent multiple analyses from inevitably increasing the privacy risk, it is the only known framework that can measure precisely how privacy risk accumulates across multiple analyses. However, privacy can be reinforced by limiting the number of queries from users.

Differential privacy can be used with several other PETs, including secure multi-party computation, homomorphic encryption and federated learning. By applying noise to data before or during these processes, differential privacy can provide further privacy protection.

CASE STUDY 13: Differential privacy for learning when locations are busy⁷⁸

Before visiting or exploring different places, tourists may wish to know in advance how busy certain destinations may be and at what times they may be less crowded. Such information can help people avoid places at peak times and plan to visit them when they are more accessible. In order to provide this service, a mapping platform uses data from individuals who are visiting certain locations and have shared their data. Differential privacy is used to protect the identities of the individuals providing location data.

Differential privacy aggregates the data collected from individuals and ensures that it cannot be used to identify them by adding sufficient random noise after a query is identified. Moreover, the exact location of an individual is never learned. To safeguard privacy, if there is insufficient data to provide accurate results and still maintain privacy, no data is published.

⁷⁴ Privacy—Control, Apple, available at <https://www.apple.com/uk/privacy/control/>.

⁷⁵ New privacy-protected Facebook data for independent research on social media's impact on democracy, Meta, February 12, 2020, available at <https://research.facebook.com/blog/2020/2/new-privacy-protected-facebook-data-for-independent-research-on-social-medias-impact-on-democracy/>.

⁷⁶ Differential Privacy, Microsoft, May 26, 2023, available at <https://learn.microsoft.com/en-us/viva/insights/Privacy/differential-privacy>.

⁷⁷ Privacy Preserving Machine Learning: Maintaining confidentiality and preserving trust, Microsoft, November 9, 2021, available at <https://www.microsoft.com/en-us/research/blog/privacy-preserving-machine-learning-maintaining-confidentiality-and-preserving-trust/>.

⁷⁸ Get information about busy areas from Google Maps, Google, available at <https://support.google.com/maps/answer/11323117?hl=en>.

CASE STUDY 14: Differential privacy for understanding digital inequity within a country⁷⁹

In order to understand broadband-related opportunity gaps across the United States, a dataset was created to show broadband coverage by zip code. Its purpose was to inform researchers and policymakers on ways to improve broadband access. To ensure this dataset was supplied in a privacy-preserving way, differential privacy was used.

A leading technology company collected data from its hardware, software and services on a continuous basis, including connection and download speeds. To determine zip code level, IP addresses were used, and devices with internet connection speeds over 25Mbps were deemed to have broadband.

Differential privacy was used to aggregate and anonymize outputs from the data collected, thus limiting the amount of information the output revealed about individuals. As random noise was added to the output, the amount of noise was dependent on the sensitivity of the query. For instance, if a query sought to learn the broadband coverage of a zip code with relatively few houses, more noise would be added in order to protect the privacy of individual households. These results afforded accurate data on broadband coverage across the US—vital information for policymakers seeking to direct public funds to reduce digital opportunity gaps.

CASE STUDY 15: Differential privacy for training speech recognition models⁸⁰

Speech recognition models must use live, real-world data from individuals using devices in order to perform at a high level. However, privacy protections need to be implemented to keep such data secure. If performed successfully, a privacy attack could learn the identity of the speakers whose data was used for training and, in some cases, even the inputs themselves. Differential privacy can be used to protect the data of individuals who contribute to model training.

By adding sufficient random noise during training, attackers are unable to learn the relationship between inputs and outputs and the data used in training. This protects against model inversion attacks and shields the privacy of individuals.

CASE STUDY 16: Differential privacy for understanding economic opportunity through the measurement of social capital⁸¹

Researchers wanted to learn whether individuals can improve their economic status through their social network. The researchers hypothesized that an individual's social connections play a major role in shaping success in school and the workforce. Using data from friendships on a social media platform, differential privacy was used to provide measures of social capital across zip codes, high schools and colleges in the United States, while protecting privacy.

Social capital was organized into three categories: (1) connectedness (the extent to which people with different traits are friends with each other), (2) cohesiveness (the degree to which friendship networks are separated into groups and/or supported by mutual friends), and (3) civic engagement (indications of trust or participation in civic organizations). Differential privacy was then used to measure these different categories by aggregating measurements and adding noise to protect individual privacy, while maintaining a high level of statistical reliability.

Using the data provided, researchers found that social connections play a significant role in helping individuals achieve economic mobility. Communities that cultivate more connections between low-income and high-income individuals tend to have higher levels of mobility.

79 U.S. Broadband Coverage Data Set: A Differentially Private Data Release, Microsoft, March 2021, available at <https://www.microsoft.com/en-us/research/publication/u-s-broadband-coverage-data-set-a-differentially-private-data-release/>.

80 Better differential privacy for end-to-end speech recognition, Amazon, January 11, 2023, available at <https://www.amazon.science/better-differential-privacy-for-end-to-end-speech-recognition>.

81 Social Capital Atlas, Meta, available at <https://dataforgood.facebook.com/dfg/tools/social-capital-atlas#accessdata>.

Table 12: Limitations and Potential Solutions of Differential Privacy

Challenge	Description	Solution
Lack of standards for target privacy levels	There is currently no best practice regarding the level at which the privacy loss parameter should be set. Researchers focused on preserving privacy generally advocate for setting ϵ at a small value, but many practitioners prefer larger values, as less accuracy is sacrificed.	Setting the privacy loss parameter is a policy and case-by-case decision, as different types of data require different levels of protection and accuracy. Through research, experimentation and discussion, stakeholders need to develop guidelines and best practices.
Trade-off in data accuracy	By design, differential privacy works by sacrificing accuracy, especially in applications with small populations. If highly accurate statistics are desired, this objective may be difficult to achieve.	Researchers are continually developing new methods to improve the accuracy-privacy trade-off. For example, some algorithms use adaptive privacy levels that can adjust the amount of noise added to the data based on the sensitivity of the query. ⁸² For machine learning models, the amount of noise applied to each attribute is dependent on the feature's importance and data type. ⁸³ IBM's differential privacy library allows users to explore the impact of differential privacy on machine learning accuracy, using classification and clustering models. ⁸⁴
Restricted to use cases of aggregated data	Differential privacy is not suitable for low counts. This is because the amount of noise applied will significantly affect the results.	Local differential privacy can be used, so minimal noise is added. Adaptive privacy levels can be used to adjust the amount of noise added to the data based on the sensitivity of the data and the query.
Sensitivity to outliers	Differential privacy can be sensitive to outliers. Outliers can cause more noise to be added than necessary in order to protect their privacy. This can reduce the accuracy of the results for the majority of the data points that are not outliers.	To address this issue, data pre-processing techniques can be used to identify and remove or cap outliers before applying differential privacy.
Complexity	Differential privacy can require expertise to implement correctly. For example, deciding the appropriate value for epsilon is difficult.	Initiatives such as OpenDP are helping address this obstacle. This open-source project develops tools for differential privacy that are ready for use. It also provides methods of analysis for the researchers who study the data. ⁸⁵ Similarly, BigQuery is a platform that enables analysts and data scientists to apply differential privacy to their datasets with ease. ⁸⁶

Source: CIPL

ii. Synthetic Data

Synthetic data refers to artificially generated data that resembles real data. Personal information may be replaced with fake data, or all original data may be removed. By carefully generating synthetic data, an alternative to real data can be provided that protects privacy, without losing the value from data. It can be used for various purposes, including data analysis, machine learning model training, testing and data sharing without the risk of exposing sensitive information.

There are two broad types of synthetic data: fully synthetic and partially synthetic.

Table 13: Different Types of Synthetic Data

Type	Description
Fully synthetic	Describes a dataset that contains no real-world data and is completely artificial. This type offers the strongest privacy protection.
Partially synthetic	Select sensitive values are replaced with synthetic versions. As not all original values are removed, the risk of reidentification is greater than that of a fully synthetic dataset.

Source: CIPL

⁸² Asma Alnemari, *An Adaptive Algorithm for Range Queries in Differential Privacy* (2016), <https://scholarworks.rit.edu/theses/9252/>.

⁸³ Assem Utaliyeva, Jinmyeong Shin and Yoon-Ho Choi, *Task-Specific Adaptive Differential Privacy Method for Structured Data* 23(4) *Sensors* 1980 (2023), <https://www.mdpi.com/1424-8220/23/4/1980>.

⁸⁴ Diffprivlib: The IBM Differential Privacy Library, available at <https://github.com/IBM/differential-privacy-library>.

⁸⁵ OpenDP, available at <https://opendp.org/>.

⁸⁶ Introducing BigQuery differential privacy and partnership with Tumult Labs, Google, May 9, 2023, available at <https://cloud.google.com/blog/products/data-analytics/introducing-bigquery-differential-privacy-with-tumult-labs>.

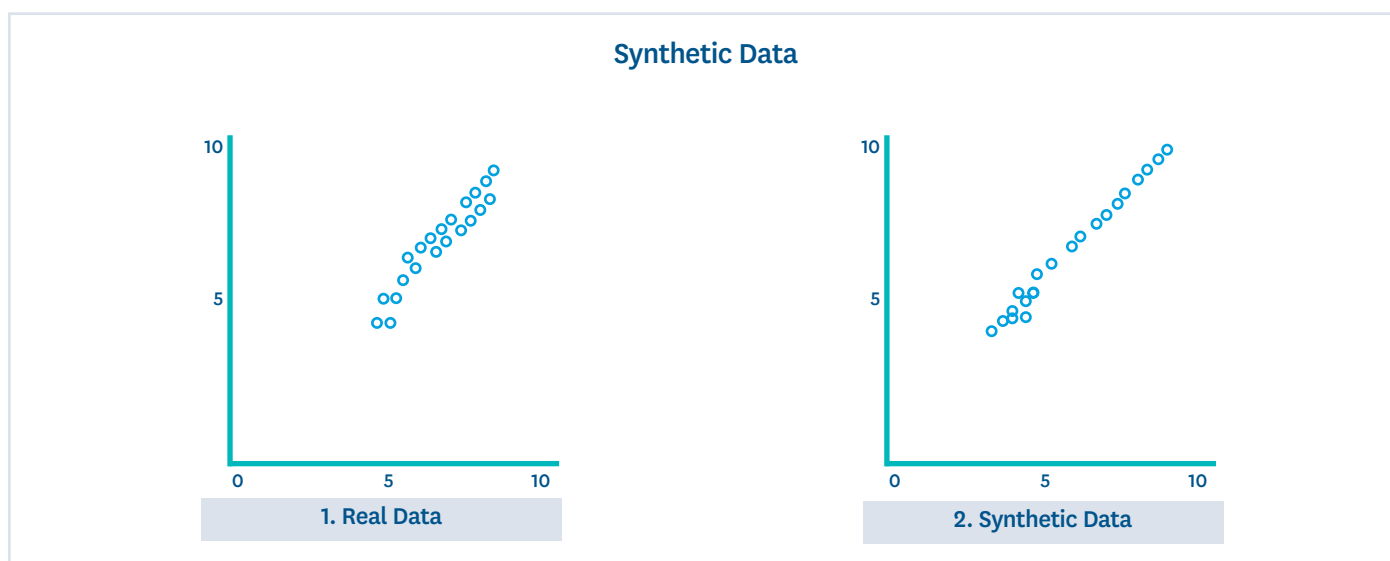


Figure 7: Synthetic Data

Source: CIPL

Synthetic data is often combined with differential privacy to add a second layer of privacy protection. This helps to prevent the disclosure of any original data from the synthetic dataset by also adding random noise.

CASE STUDY 17: Synthetic data and differential privacy for learning about human trafficking⁸⁷

Sensitive data is sometimes unavailable for research purposes. However, through the use of synthetic data and differential privacy, a research institute's sensitive case records on trafficking victims were able to be used to produce insights on the relationships between victims and perpetrators.

Synthetic data was generated from the original sensitive data. By creating this artificial data, the identities of the real victims were hidden, preserving their privacy, while simultaneously maintaining the statistical properties of the original data.

Differential privacy was also leveraged to provide an extra layer of privacy by providing quantifiable privacy guarantees, which facilitated the disclosure of more data. The combination of these PETs allowed the release of data on trafficking victims and their descriptions of the perpetrators.

CASE STUDY 18: Synthetic data for training and testing algorithms on biometric data⁸⁸

Synthetic data can be used to enable research based on biometric data while preserving privacy.

Synthetic data has been used to provide large volumes of diverse data to an organization wishing to create a system empowering individuals to use their palms to make payments. The system requires customers to hover their palm above a device, which then identifies their unique palm signature. The unique palm signature is associated with the customer's payment card to charge them for their purchases. Using synthetic data helped address the shortage of public data on images of palms and enabled the organization to train its algorithms across a diverse range of data.

⁸⁷ The Global Victim-Perpetrator Synthetic Dataset, Counter-Trafficking Data Collaborative, December 1, 2022, available at <https://www.ctdatacollaborative.org/global-victim-synthetic-dataset>.

⁸⁸ Four key physical retail technology takeaways from re:MARS 2022, Amazon, June 22, 2022, available at <https://www.aboutamazon.com/news/retail/in-store-shopping-technology-re-mars-2022>.

CASE STUDY 19: Synthetic data for detecting fraud in healthcare insurance claims⁸⁹

Synthetic data is being used to help health insurance companies detect fraud and abuse in claims. Synthetic datasets based on data about medical histories, healthcare claims and other medical data are being used to train and validate AI algorithms to identify fraudulent claims and irregularities in an individual's healthcare records.

In addition to removing the need for sensitive healthcare data and protecting privacy, by leveraging synthetic data, bias from real-world data can be reduced and the larger volume of data can allow the algorithm to scale better.

Table 14: Limitations and Potential Solutions of Synthetic Data

Challenge	Description	Solution
Data leakage	Data can be leaked from synthetic data. For example, if the real data contains outliers which are then captured by the synthetic model, malicious actors would be able to learn information about the real data.	Using differential privacy with synthetic data can help address data leakage. By adding random noise, outliers will reveal less information.
Low accuracy	Deployers must be aware that not all synthetic data will replicate the content of the original data. They must understand how the synthetic data was generated, how accurate it is and, by extension, the circumstances for which its use would be appropriate. In cases where highly accurate data is important, synthetic data may not be appropriate.	This is an inherent limitation of synthetic data. Each situation must be individually evaluated to decide whether synthetic data is appropriate and if the model being used to generate the synthetic data is suitable.
Potential information loss	In the process of generating synthetic data, some information from the original data may be lost.	Models being used to generate synthetic data must be carefully designed and tested before use.
Dependency on real-world data	Synthetic data relies upon real-world data. If real-world data is inaccurate, incomplete or biased, it can negatively impact the synthetic data being generated. Also, if the original data is dynamic and changes over time, this adds a further complication and requires the synthetic data to be kept up-to-date where necessary.	Checks are necessary to assess the accuracy and bias of the original data before generation. Data imputation can help complete datasets by replacing missing values. There are also automated tools that can be used to automatically update the synthetic data when the original data changes.
Bias in synthetic data	Models used to generate synthetic data must be designed carefully to avoid biases or inaccuracies that can then impact the synthetic datasets generated from them. Bias from the original dataset will also be replicated into the synthetic data.	Thorough model testing and assessments of the synthetic data once it has been generated can help address biases that can creep into the manufactured data.

Source: CIPL

⁸⁹ Anthem Looks to Fuel AI Efforts With Petabytes of Synthetic Data, The Wall Street Journal, May 17, 2022, available at <https://www.wsj.com/articles/anthemfuel-ai-efforts-with-petabytes..>

Appendix 1—Additional Case Studies

CASE STUDY 20: Homomorphic encryption for protecting passwords⁹⁰

Individuals possess many passwords to prevent unauthorized entities from accessing their online accounts, but passwords are sometimes compromised through data breaches by attackers. Organizations, therefore, need to be aware if an individual's password has been breached so that a new password can be created as soon as possible. When checking whether a user's password has been breached, one organization uses homomorphic encryption to ensure the user's password is not exposed to any parties during this process.

First, the user's credentials are hashed and sent from the server to the client. The client uses homomorphic encryption to encrypt the hashed value and send the output to the server. While the data remains encrypted, the server computes a matching function to learn whether the credentials have been part of a data breach. The encrypted result is returned to the user where it is decrypted and the result is learned.

Homomorphic encryption prevents the organization from learning end-users' passwords, while enabling the monitoring of the security status of those passwords. Additionally, third parties are prevented from accessing this information even if they successfully intercept this data while it is in transit between the user and the organization's servers, thus protecting against man-in-the-middle attacks.

CASE STUDY 21: Trusted execution environments for cloud computing⁹¹

Cloud service providers commonly leverage trusted execution environments to enable organizations to collaborate and share data securely. For example, one organization provides a trusted execution environment as part of its cloud service, allowing data to be processed in the cloud with high security and in accordance with the customer's wishes.

By using a trusted execution environment, this organization ensures that the data provided by customers is not exposed to unauthorized applications, codes or parties. For example, a trusted execution environment is used on a healthcare platform to safeguard patient data, while selected shared data is being processed.⁹² An attestation mechanism is used to verify that the trusted execution environment is secure and that the customer-defined restrictions have been applied. The user's data is then encrypted and stored in the trusted execution environment for processing. The data is, therefore, processed securely in the cloud, providing only user-defined hardware and software access to the data and preventing unauthorized data access by cloud providers, administrators and users.

⁹⁰ Password Monitor: Safeguarding passwords in Microsoft Edge, Microsoft, January 21, 2021, available at <https://www.microsoft.com/en-us/research/blog/password-monitor-safeguarding-passwords-in-microsoft-edge/>.

⁹¹ Azure confidential computing, Microsoft, available at <https://azure.microsoft.com/en-us/solutions/confidential-compute/#overview>.

⁹² Confidential computing on a healthcare platform, Microsoft, available at <https://learn.microsoft.com/en-us/azure/architecture/example-scenario/confidential/healthcare-inference#architecture>.

CASE STUDY 22: Synthetic data for training chatbots⁹³

In order to enhance chatbot performance, synthetic data has been used to improve the quality and accuracy of chatbots' outputs and to train these programs without using sensitive data that could reveal features such as ethnicity or user identity.

Individuals have different accents and speak in different styles. Instead of collecting this data, an algorithm produces artificial sentences and then checks them for accuracy. Using synthetic data eliminates the risk that attackers could learn information about individuals whose data was used to train the chatbot.

CASE STUDY 23: Federated learning and secure multi-party computation for text selections⁹⁴

In order to improve a device's ability to predict the word or collection of words a user intends to select, an organization leverages server-coordinated federated learning to preserve user privacy. This enables the neural network model to be trained on real user interactions responsibly.

To begin, the central server initializes the model. Participating devices are selected and sent a copy of the model. When a user selects text and corrects the device's suggestion on one of these devices, the device can use this data to improve its own model, and, in turn, send the updated model weights to the server. The server uses this information to refine the global model before the process restarts until the model is as precise as possible. Privacy is protected because raw data is not transmitted, and the user's selections are only temporarily kept on the device.

Privacy is supported further by using secure aggregation, a form of secure multi-party computation. This means that the model weights of multiple devices are averaged to prevent model weights from individual devices from being read. This results in a more accurate model while preserving privacy.

CASE STUDY 24: Federated learning and secure multi-party computation for suggesting message replies⁹⁵

An organization has used federated learning to improve message suggestions on its devices. Use of this technology has created a better user experience while keeping user data (i.e., the content of conversations) on the device, hidden from the organization at all times.

Here, a machine learning model is sent to individual devices. Summaries of recent model changes (based on the user's actions) are then sent back to the organization's servers, where secure multi-party computation, specifically secure aggregation, is used to combine these model changes with many other devices. These combined changes are used to improve the suggested reply feature for all end-users.

By using federated learning, the organization does not see the content of messages or individual user data, as this remains on the device. Secure aggregation strengthens privacy by preventing the organization from viewing model changes from individual devices.

93 What is synthetic data?, IBM, February 8, 2023, available at <https://research.ibm.com/blog/what-is-synthetic-data>.

94 Predicting Text Selections with Federated Learning, Google, November 22, 2021, available at <https://ai.googleblog.com/2021/11/predicting-text-selections-with.html>.

95 How Messages improves suggestions with federated technology, Google, available at <https://support.google.com/messages/answer/9327902?hl=en#zippy=%2Chow-federated-technology-works>.

Appendix 2—Glossary

Included below are a list of key terms used in this paper.

- **Cloud computing:** the delivery of computing services over the internet, enabling users to access and use shared computing resources instead of relying on local infrastructure.
- **Compiler:** a software tool that translates source code written in a high-level programming language into machine code that can be executed by a computer.
- **Data compression:** the process of using compression algorithms to encode data in a more concise form, resulting in reduced storage space or bandwidth transmission requirements.
- **Data reduction:** the process of reducing the size or volume of data while preserving its essential characteristics.
- **Deep neural network:** a machine learning method that uses multiple layers of interconnected nodes to process large amounts of data.
- **Edge devices:** pieces of equipment that transmit data between the local network and the internet. Examples include smartphones, Internet of Things devices and routers.
- **Hashed:** the process of applying a hash function to data to produce a fixed-size string of characters that uniquely identifies the data.
- **Keys:** a random sequence of characters that is used by an encryption algorithm to transform the original data into an encrypted format known as ciphertext, or vice versa. The key acts as input to the encryption algorithm, determining how the encryption process is performed.
- **Large language models:** a type of machine learning model that is trained on huge amounts of data to understand and generate human-like language.
- **Machine learning:** the ability of computer systems to learn and improve from experience without being programmed.

- **Man-in-the-middle attacks:** cyberattacks where the attacker secretly intercepts and relays communication between two parties without their knowledge. The attacker can eavesdrop on the communication and manipulate the data being transmitted.
- **Membership inference attacks:** when an adversary attempts to determine, based on the output of a machine learning model, whether a specific data point was used in the training data.
- **Model inversion attacks:** a privacy attack that aims to extract sensitive data used to train a machine learning model by exploiting the model's output.
- **Outliers:** data points that are significantly different from the rest of the data.
- **Parallel computing:** a computational approach where many calculations or processes are executed simultaneously.
- **Side-channel attacks:** security attacks that exploit unintended information leakage from a system's physical implementation to obtain sensitive data.

About the Centre for Information Policy Leadership

CIPL is a global data privacy and cybersecurity think tank in the law firm of Hunton Andrews Kurth LLP and is financially supported by the law firm and 90 member companies that are leaders in key sectors of the global economy. CIPL's mission is to engage in thought leadership and develop best practices that ensure both effective privacy protections and the responsible use of personal information in the modern information age. CIPL's work facilitates constructive engagement between business leaders, privacy and security professionals, regulators and policymakers around the world. For more information, please see CIPL's website at

<http://www.informationpolicycentre.com/>.



Centre for Information Policy Leadership

HUNTON ANDREWS KURTH

DC Office

2200 Pennsylvania Avenue
Washington, DC 20037
+1 202 955 1563

London Office

30 St Mary Axe
London EC3A 8EP
+44 20 7220 5700

Brussels Office

Avenue des Arts 47-49
1000 Brussels
+32 2 643 58 00