Google

May 2025

Responsible Development And Use of Privacy Enhancing Technologies: How Governments Can Help

This is a working paper, reflecting views as of May 2025. We will update the principles and recommendations outlined in this document as the technology continues to develop.

01 Executive Summary

02 Overview of PETs and Google's Approach

03 Challenges To Development and Use

04 Opportunities for Governments to Advance PETs

Championing their use and leading by example by implementing PETs in government data collection and use

Investing in research and skillbuilding to continue to advance the field

Fostering openness and cross-sector, cross-border collaboration

Implementing smart policies and a risk-based regulatory framework, including regulatory clarity and incentives to encourage responsible use.

05 Conclusion

Executive Summary

Privacy Enhancing Technologies (PETs[1]) are generally understood to be a broad suite of techniques and technologies that enable responsible use of data to power valuable insights, products, and services, with privacy and security protections.

PETs have been around for years, and people have benefited from their use. At Google, we use PETs across a wide range of products, including Google Trends, Busyness in Maps, Gboard, and more. Our ongoing commitment includes significant investments in research and development to yield technological advancements. By sharing technological tools and research findings, we aim to foster a more responsible, private, and secure data ecosystem for all stakeholders.

Generative AI (GAI) is accelerating innovation which can spur economic growth and solve complex societal problems. Policymakers around the world are considering how to best guide this innovation so that its full potential can be unlocked for society responsibly. Top of mind is how generative AI model training and applications influence data collection and processing. As GenAI adoption expands, along with it, we see new innovations in the space of PETs.

PETs are an AI opportunity: when used responsibly, they play an increasingly important role in protecting people and their data and preventing suboptimal outcomes like costly harms that can come from data misuse, such as ID theft and fraud. Additionally, with PETs, businesses can develop and grow new lines of revenue, access new markets, and continue to fortify trust in digital services and data flows that are critical to the global economy.

Despite the promise of PETs and high levels of interest from industry, governments, academia, and civil society,[2] adoption of PETs is not widespread, particularly outside the tech sector. Technical complexity, high implementation costs, a lack of necessary expertise and awareness of PET options and how to implement them responsibly, and regulatory uncertainty can pose high barriers to entry. **These challenges are shared across the private and public sectors but governments are uniquely suited to spur investment at scale and expand access to PETs.**

This paper builds on our previous work[3] on data protection frameworks and white papers to provide recommendations for governments around the world to consider. It begins with an overview of PETs, including what they are, how they can be used, and their opportunities for improving data protection, before outlining Google's approach to developing and using PETs in the products we build for our users and customers. The paper then discusses challenges to the further development and use of PETs and concludes with concrete recommendations in four core areas for government intervention:

- Championing their use and leading by example by implementing PETs in government data collection and use;
- Investing in research and skillbuilding to continue to advance the field;
- Fostering openness and cross-sector, cross-border collaboration; and
- Implementing smart policies and a risk-based regulatory framework, including regulatory clarity and incentives to encourage responsible use.

01 Overview of PETs and Google's Approach

When used responsibly, PETs can provide meaningful technical privacy and data protections in a broad range of applications.

We can group PET technologies into two broad themes:

- Those that add isolation protection (e.g. trusted execution environments or TEEs, federated learning and analytics, secure multi-party computation), and
- Those that anonymize data to make it safer for processing and exploration (e.g. leveraging techniques including differential privacy or DP, personally identifiable information filtering, synthetic data generation, clustering and more).

As the UK's Royal Society notes, PETs are particularly helpful in terms of collaboration across organizations, where data sharing risks cause friction or may be prohibited, thereby hindering the pace of promising scientific breakthroughs such as healthcare research across public and private sector institutions, which benefit society at large.

Like building blocks, PETs can be used alone or in combination. In the case above, we see technologies from both themes combined to ensure an end-to-end privacy-protective environment for collaborative data sharing.

Using multiple PETs together is common, and yields comprehensive protection. In machine learning contexts, application of personal data filtering and differential privacy (DP) or k-anonymization can be used to reduce memorization risk of training data. Filtering model outputs at inference time can further eliminate or reduce the risk of inferred (or hallucinated) data.

PETs can also be used at multiple points of the AI lifecycle–data collection, model training, fine tuning, and model use–to provide comprehensive protection. Retaining flexibility to combine and develop new technologies is critical to allowing privacy protections to keep pace with technological and product innovation writ large.

Organizations need to constantly find new ways to scalably match privacy and security protections with new products and surfaces, new modalities, new business needs, and with all of the above, new kinds of vulnerabilities to data privacy and security.

This is especially important given the accelerating pace of new products and services that generative AI will make possible in the coming months and years.

Google's Approach

Google continues to make significant, long-term investments in researching, developing, and deploying the industry's most effective, intuitive, and innovative privacy technologies. We leverage a range of efforts to deploy privacy protections across many Google products and services, enabling us to offer more secure, private, and personalized experiences to our users. We have even used PETs to estimate carbon emissions[4] and discover new common words as languages evolve, including identifying 3,600 previously missing Indonesian words from Gboard, a popular Android keyboard app, in just two days.[5] Recent key advances include:

- PARFAIT, or Private Aggregation and Retrieval, Federated, Analytics, Inference, and Training, an open source project that showcases technologies for private AI for transparency, data minimization, data anonymization, and external verifiability.[6]
- Protecting users through use of differentially-private synthetic training data, our approach to detect unsafe content in device settings.[7]

We want to make it easier for people to experiment with projects and code we have open-sourced. For example:

- Organizing the first Machine Unlearning Challenge[8], to drive collaboration on a promising area of research to address the technical challenges in deleting data from deep learning models by enabling certain data to be forgotten to protect privacy or and other rights.
- Publishing a Practical Guide to Machine Learning with Differential Privacy to discuss the current state of research, common techniques, engineering challenges, mitigation techniques and current open questions.[9]

Examples of PETs at Google today

Search Trends provides public access to a random sample of actual search requests made to Google. The data is anonymized (no one is personally identified), categorized (determining the topic for a search query) and aggregated (grouped together) to prevent reidentification; searches made by very few people, duplicate searches, and special characters are excluded.[10] Trends provides insights to journalists on emerging topics of interest, provides content creators information about what people care about, and the public to see what searches are trending near them, and enables analysis of interest over time and by location.

Busyness feature in Maps: Google uses a combination of PETs to aggregate and anonymize data, from users who have opted in to Google Timeline, to show popular times, visit duration, and wait times in a specific location. With this feature, small businesses can gain insights into their customer flow and optimize operations, such as staffing and inventory management. For customers, it can help to know the optimal time to visit a local restaurant or service provider to minimize wait time or, as we saw during the pandemic, to make informed decisions about social distancing implications.[11]

Confidential matching in ads: Google uses TEEs in our confidential matching solution. For advertisers, using a TEE means added protections for their customers' information when used in marketing. It provides technical assurances including transparency into confidential matching's code and the ability to receive proof, known as an "attestation," that data is processed as intended. Its architecture is available to the public, in addition to a few other open-source examples to help others build confidential solutions.[12]

Zero-knowledge Proof technology for age information: Google is actively integrating Zero-knowledge Proof (ZKP) technology into Google Wallet where stored digital ID passes will be used for age verification purposes to support access to age-gated content. The ZKP technology helps mitigate risks that ecosystem players could work together to reveal the users' underlying identity. Previously, zero-knowledge proofs have been so computing-intensive as to be impractical. Google has developed an approach to ZKP that makes it more practical for the use case of age verification and intends to open source its ZKP technology in the near future.[13]

Android's Private Compute Core (PCC) is a secure, isolated data processing environment inside of the Android operating system that gives users control of the data inside, such as deciding if, how, and when it is shared with others. This way, PCC can enable features like Live Translate without sharing continuous sensing data (e.g., audio, etc.) with service providers, including Google.[14] PCC is a part of the larger Protected Computing toolkit[15] which includes myriad techniques such as cloud enclaves, edge processing, and end-to-end encryption to ensure sensitive data remains in exclusive control of the user. The publicly-verifiable architectures in PCC demonstrate how we strive to deliver confidentiality and control, and do it in a way that is verifiable and visible to users. We provide this transparency through blogs,[16] public documentation,[17] and open-source code.[18]

Workspace's Client-side Encryption (CSE) feature supports Cloud customers who want to ensure the privacy and security of their email and work documents. CSE allows users to add another layer of encryption to their communications and collaboration data and gives them greater control over access to that data with end-to-end encryption that third parties or Google can't decrypt because only the customer has and controls the keys.[19]

Our research and engineering teams regularly publish their research advances and open-source tooling so that anyone, anywhere can benefit from our work and contribute to progress in the field of PETs.[20] We built and operate the world's largest open-source library of differential privacy algorithms, helping everyone—from cancer researchers to census analysts—apply privacy-preserving technology to their work.[21] We have continued to improve our open-source offering with the goal of democratizing access, including to differentially private algorithms.

As mentioned above, we also work to expand the range of advanced privacy features available in Google Cloud's products and services so that customers are able to take advantage of these advancements without having to make significant independent investments in research and tooling and tackle pressing challenges.[22]

Looking to the future, Google Cloud is working with Swift, the global provider of secure financial messaging services, to develop anti-fraud technologies that use advanced AI and federated learning to better combat fraud in cross-border payments. The lack of visibility across the payment lifecycle creates vulnerabilities that can be exploited by criminals. Incorporating federated learning provided a solution for collaborative AI model training without compromising privacy and confidentiality. Instead of requiring financial institutions to pool their sensitive data, the model training occurs within financial institutions on decentralized data. The overarching solution incorporates additional PETs and has the potential to redefine collective efforts to combat fraud in a fragmented context and contribute to a more trustworthy system.[23]

02 Challenges to Development and Use

Though PETs have been available for some time, adoption is still relatively nascent and not evenly distributed across private and public sectors. This section discusses the organizational, resource, technical, and regulatory challenges to further development and use of PETs.

Despite the encouraging interest amongst industry, government, and academia, in PETs and the significant promise they afford, understanding of PETs remains limited. Many in these settings are unclear about how PETs work, about different PET capabilities and limitations, potential use cases, and their impact on existing organization-critical applications. Further, we observe a lack of understanding of how to appropriately balance privacy and other factors, which we discuss further below. As with any technology innovation, even more so for those particularly complex, it can take time for new applications to emerge. Multiparty Computation was developed in the 1980s but is only now experiencing a promising wave of adoption in new applications, including in organizations outside of the technology sector.

Organizations may not immediately understand how to get started: for example, how best to implement PETs in their workflows, and may not have the necessary expertise to implement them correctly, or to accurately assess which solutions (whether open-source, commercial, or a combination of both) are appropriate for their near and long-term needs. Additionally, organizations implementing PETs may encounter challenges in helping external partners, customers, and subcontractors understand and trust their usage. Some organizations directly integrating PETs find that, as with any new technology investment, PETs do have costs associated with implementation (e.g., ramp-up, training, headcount integration, and maintenance costs) and risk assessments.

Organizations, particularly businesses, may have difficulty calculating the opportunity cost of integrating PETs into product and business planning, further complicated by the absence of greater regulatory certainty. Return-on-investment calculations may not be persuasive, especially when benefits may be longer-term and harder to quantify.

Just as there is no single "correct" model or technique for all tasks, there is no single approach to PETs. Like any technology, PETs have constraints and tradeoffs, and decisions about deployment should be made in the context of specific product surfaces, privacy threats such as leakage, memorization, or re-identification; fairness to different entities; utility and accuracy; fraud and abuse detection, and regulatory obligations. There are settings in which tradeoffs are pronounced and priorities are competing; implementing organizations (and all stakeholders) may have to contend with ambiguity.

A number of factors influence which PETs are appropriate for a product or scenario and the legal obligations they may implicate or address. Analyses should take into account threat models (e.g., the risks an organization is trying to protect against); how the data will be used, including the nature, scope, context, and purposes of processing; the data protection risks of the processing under applicable law; whether and how PETs will address the recognized data protection risks; the technical complexity of implementation; the impact on functionality to end-users or others; the costs and the resources available to the implementing organizations.

PETs are not a magical solution nor should they grant a blank check for the use of data or technology that should otherwise not be pursued on ethical, legal, or other grounds. Like any technology or tool, PETs must be implemented correctly to achieve the intended outcomes.

Organizations should consider frameworks to select, evaluate, and test PETs to realize effective use and guard against 'privacy washing' or the false promise of protection, which could not only harm users but also stymie collective efforts across government, industry, academia, and civil society to advance PETs. Resources such as those from data protection authorities and the Future of Privacy Forum's proposed risk-utility framework for PETs in advertising[24] are instructive. Google will publish a paper in the coming weeks proposing a technical framework to help practitioners evaluate PETs. It challenges the notion that privacy is a zero-sum tradeoff with utility and cost and provides some illustrative examples and practical guidance.

As part of a complex assessment, organizations must consider whether the protections provided by a particular PET or PETs adequately address regulatory, legal, and service-level requirements. As a result, a lack of clarity or certainty on legal risk is a significant impediment to use, potentially causing some organizations to hesitate about investing in the necessary people, infrastructure, and tooling to integrate PETs into their workflows and product and business roadmaps.

03 Opportunities for Governments to Advance PETs

Governments have a particularly critical role to play in advancing the development and adoption of PETs. An effective policy and regulatory strategy should be multifaceted, responsive to the aforementioned challenges, and be flexible to account for ongoing technological advancement and the myriad use cases and organizations seeking to implement them.

In this section we discuss areas where government intervention is particularly needed: championing use of PETs and leading by example, making investments in research and development, fostering openness and collaboration, and perhaps most urgently, implementing smart policies and a risk-based regulatory framework that provide much-needed clarity and incentives for industry. **As we've seen with AI, countries that expressly encourage the development and use of PETs will gain a competitive advantage. They will be better positioned to benefit from the innovation, societal progress, and economic growth that come from investments in technologies that support a digital ecosystem respectful of privacy.**

Champion Use of PETs and Lead by Example

By necessity, governments collect and use a lot of data. This provides an excellent **opportunity for governments to use PETs in their applications and services, and expand PETs use more broadly through incentives in procurement and contracts.** Doing so can bolster privacy protections and lower risk for citizens, improve national security, build much-needed capacity across government entities, and bring important insights that can inform future decisions. It can also promote public understanding and confidence in PETs, to further their development and use.

Outside of direct usage, **governments can use their visible platforms to provide fact-based information across sectors, disciplines, and perspectives to demystify and boost visibility of PETs,** collect and share use cases, and disseminate sound practices for stakeholders. **Governments can send important signals through public commentary** that encourages businesses to invest in and apply PETs in their work. Specific activities to consider could include:

- publishing national and/or regional strategies on PETs with calls for public comment;
- including PETs in government budgets and funding requests;
- commissioning an expert taskforce to review how to advance use of PETs in local, state, national, and regional governments; and
- hosting or participating in workshops and conferences to discuss the opportunities and challenges presented by PETs and AI.

Invest in Research and Skill-Building

Several factors slow down the wider development and use of PETs: an incomplete understanding of the technologies, uncertainty over their applicability to myriad use cases, limited awareness of the existence of available tools, and the costs and expertise required to implement them. **Increasing access to resources, including open-source tools and code libraries, relevant data, expertise, and computing power, will assist in lowering those barriers to implementation,** particularly for smaller organizations. It's also a smart way to leverage finite resources, identify common problems, and build a shared understanding.

Many challenges in advancing PETs are similar for everyone—working together, we can identify effective and scalable solutions.

Many governments already invest in education and training. Inclusion of PETs in that work will encourage a larger workforce to support new techniques, vulnerabilities, and solutions, whether across government, academia, civil society, or business, and will pay dividends for civil society and industry.

It is important that research, particularly fundamental research, does not solely come from the private sector. High-risk, high-reward research, such as development of new techniques, may only be possible with public funding and some applications may be best suited for government or government-funded entities to drive. Direct support to educational and research institutions through appropriations, grants, scholarships, innovation sprints, and internships, fellowships, and faculty awards can support fundamental and application-specific research. Separately, governments can explore means to ensure entities have the computing power and infrastructure necessary to continue their efforts. Governments may want to consider extending this support to civil society organizations and small businesses.

Additional activities governments could implement include designating funding for applications of national importance like healthcare and national security; directing government labs and agencies to undertake and/or fund research; allocating head count for priority agencies to boost capacity; establishing centers of excellence for research and applications; and running studies or surveys to assess the industry's awareness of AI, interest in using it, and barriers preventing its uptake.

Promote Openness and Cross-Sector, Cross-Border Collaboration

Collaboration is an engine for progress and good ideas can come from anywhere. Governments should consider requiring or incentivizing government-supported research and datasets to be made publicly available, and when relevant, **support contributions to open-source ecosystems** like code libraries, which can ease the deployment costs of PETs and broaden their accessibility. Promoting openness and collaboration can be as simple as **creating or amplifying opportunities for researchers and developers to share their work** through conferences and citing research in government publications.

Coordination within and across governments and international stakeholder groups can also help identify shared challenges and leverage finite resources to develop effective solutions. **Governments should consider how to work on PETs through multilateral fora**[25] to share sound practices and learn from experiences. Governments can work together to foster public and private sector research and collaboration, and identify where regulations may stymy use of PETs.

Encourage Development of Technical Standards and Implementation Guidelines

Industry standards can go a long way in reducing uncertainty and encouraging investment. National and international consensus standards and frameworks are essential to the health of the global technology ecosystem. They promote cross-border and cross-application interoperability, forge pathways for collaboration across industries, provide opportunities for improved transparency and consumer confidence, and minimize barriers to trade and innovation. **Governments can drive or support benchmarking that fortifies trust in the privacy benefits of PETs.**

Standards may also be more effective than defining specific terms or thresholds in law because they are better able to keep pace with the technical state of the art. Where possible, governments should align policies and regulations with, or otherwise encourage the development of, national

and international standards and recommended technical practices, whether through international standards bodies or government standards-setting entities, that reflect consensus among a variety of stakeholders. In turn, standard-setting bodies and technical groups can provide technical guidance to policymakers.[26] **Best practices and codes of conduct can be particularly useful in sector - or scenario-specific instances,** or where the domain or pace of development may be unique, such as advertising or biometrics in ID verification; they could also be helpful in rapidly maturing techniques like TEEs. Governments should establish pathways to approve or formalize industry or civil society codes and practices.

Implement a Risk-Based Regulatory Framework that Motivates Privacy by Innovation

When carefully crafted, a risk-based, context-aware, flexible regulatory framework can clearly show the benefits of using PETs and create strong incentives for their adoption. We offer the suggestions below for guidance to enhance regulatory clarity and promote the deployment of PETs, with robust safeguards, to meet data protection compliance obligations.

Creating common frameworks for PETs use and spaces that enable experimentation—with oversight—can be a good starting point. They allow stakeholders to test out new techniques or use cases, address liability concerns, and mutually improve understanding about technologies, their opportunities, constraints, and effective mitigation strategies. **Regulatory sandboxes**, such as those operated by the UK ICO and Singapore's IMDA, are good examples of collaboration between regulators and organizations.

The strongest incentive for businesses to consider use of PETs is in the context of legal certainty, and there are three ways policymakers can provide this:

- Take a Tiered Approach to Identifiability of Data: Defining and treating de-identified, anonymous, pseudonymous, personal, and sensitive personal data differently reflects the different risks of harm the processing of these categories of data can pose: the more personally identifying the data, the higher the potential risk of harm that can arise from its processing, and the most obligations placed on that processing.[27] A risk-based approach acknowledges the trade-offs between costs and protections given by different styles of PETs and compliance benefits from data minimization, including through the use of PETs. Anonymous data doesn't pose privacy risks to users so it isn't regulated by privacy law. But many laws treat data that is not fully anonymous in one bucket. **Treating data that has low-risk of re-identification without the same high standards for personally identifying data encourages the use of PETs.** Governments considering new or updates to existing laws should define terms like anonymization clearly and at an appropriately high level while also ensuring technical flexibility. However, these definitions should not be so vague or so prescriptive that they set an unreachable bar. **For example, defining anonymization in law as "irreversible and permanent" is a remote and theoretical standard, whereas setting the bar where data is "no longer identifiable" can be effective.** Governments should also provide clear and concrete guidance so that those using these technologies have confidence in whether they meet the threshold and confidence of the compliance benefits of doing so.

Legal privacy frameworks should allow for data processing that provides value to collaborating organizations in ways that do not pose the risks that sharing personal information might. These techniques may involve the processing of data that are not necessarily anonymous, but otherwise provides significant risk-reducing safeguards through protected environments, such as limitation of access and highly protected sharing. **We encourage implementation of clear compliance benefits for this type of processing that, through use of PETs and other protections such as transparency, fairness, and privacy principles, enables safer sharing and collaboration.** There's a robust field of techniques that are capable for this purpose. TEEs, depending on the context, are one. TEEs are physically and logically separated from the rest of a server or other device. Only the people who collected the information directly, and with permission, know who the data subjects are—no one else.

Policymakers could also explore where the responsible use of PETs can justify the reduction or removal of data transfer, processor agreement, or other regulatory obligations. Multi-stakeholder processes like taskforces or requests for information that invite a broad spectrum of perspectives and commentary can assist in identifying these areas and solutions.

- Provide Guidance that Legally Recognizes the Value of PETs in Legitimate Interests Balancing Tests and Data Subject Rights Requirements: Regulations should provide organizations a presumption of sufficiency in meeting certain regulatory obligations for processing if they meet specified criteria when using PETs. Many legal frameworks place a responsibility on organizations to consider individuals' interests when processing personal data and to protect against potential risks, permitting processing that balances what is often described as the "legitimate interests" of an organization against the impact on the rights and interests of the individual. Where processing of personal data satisfies this balancing test and the organization incorporates transparency, fairness, and privacy principles, legitimate interests can be used as a lawful basis for processing.[28] **Regulatory guidance that expressly indicates that use of PETs is a compelling factor presumptively supporting an organization's legitimate interests balancing test, as documented in their risk or data protection impact assessments, would be beneficial to organizations and users alike, providing protections, flexibility, and accountability.**

  Similarly, clear guidance that says data processed with PETs that meet specified criteria and transparency, fairness, and privacy principles, can be exempt from certain data subject rights requests and can presumptively be retained for a longer period of time would be a helpful encouragement for adoption. The aforementioned criteria should be robust and outcomes-based, focusing on the effective level of user protection being provided. They should leverage privacy by design, risk assessments, data protection impact assessments, and other governance and accountability provisions as appropriate for organizations to document their selection, evaluation, and testing of PETs for particular purposes or risk mitigations. The documentation would enable regulators to check the organizations' work while providing organizations with the necessary flexibility to match the appropriate PETs to the context, taking into account the purpose, product or surface, privacy risk, and technical modality (e.g., classification, analytics, generative).

To be effective, guidance should aim to provide a holistic vision for adopting PETs and avoid overly prescriptive, technology-specific approaches that can slow down innovation, and discourage organizations from taking advantage of the most advanced PETs. Instead, policy or regulatory guidance should **help identify specific risks or threats to privacy (e.g., data leakage, re-identification), set expectations about recommended thresholds for risks, and provide guidance on the applications of appropriate protections (rather than prescribing a specific technique).** Outcomes-based guidance that takes contextual factors into account can assist organizations in adapting their privacy risk taxonomy and develop their own evaluation criteria. **Specifically, governments could consider providing guidance and/or creating or tasking expert bodies to provide guidance for researchers and practitioners.**

- Take Use of PETs into Account in Enforcement: Use of PETs can signal a good-faith effort and policymakers could consider if and how an organization used PETs when contemplating enforcement actions. **Policymakers should consider making use of PETs as a mitigating factor in assessment of sanctions and/or levels of fines.** Doing so would mark PETs as a best practice and incentive, and send a very strong signal to organizations that PETs are an important consideration in their work and worthy of the investment needed to use them responsibly.

PETs offer a promising way to protect privacy in the digital age. PETs should be viewed as a suite of technologies which can be combined in new ways to mitigate pressing privacy concerns and enable continued technological innovation while preserving user trust and safety.

The opportunities and challenges in further development and deployment are shared across the public and private sectors. By championing use of PETs, investing in research and development, fostering collaboration, and implementing smart policies and regulations, **governments can play a crucial role in advancing these technologies and ensuring that individuals' privacy is protected, while still fostering innovation and spurring PET development and adoption.**

As policymakers around the world grapple with the accelerating pace of generative AI and what it means for society, we encourage the consideration of PETs among available options to policymakers, as a means to encourage responsible AI development and use.

1 PETs are also known as Privacy Preserving Technologies (PPTs) and Google uses the terms interchangeably.

2 PETs are a robust area of academic research. Several organizations and government entities are actively engaged in this field, providing essential information and hosting convenings. A non-exhaustive list of these includes the Centre for Information Policy Leadership (CIPL), the Future of Privacy Forum (FPF), the Organisation for Economic Cooperation and Development (OECD), and data protection authorities in Canada, France, Singapore, Spain, the United Kingdom, and the European Data Protection Board.

3 See A Framework for Responsible Data Protection Regulation, Google (2018), https://services.google.com/fh/files/blogs/google_framework_responsible_data_protection_regulation.pdf, Responsible Data Practices, Google (2022), https://services.google.com/fh/files/blogs/google_framework_responsible_data_protection_regulation.pdf, and Generative AI and Privacy Policy Recommendations Working Paper (2024), https://storage.googleapis.com/gweb-uniblog-publish-prod/documents/Google_Generative_AI_and_Privacy_-_Policy_Recommendations_Working_Paper_-_June_2024.pdf.

4 Christopher Bian, et al., Mayfly: Private Aggregate Insights from Ephemeral Streams of On-Device User Data, https://arxiv.org/pdf/2412.07962.

5 Timon Van Overveldt and Daniel Ramage, Discovering new words with confidential federated analytics, Google Research (March 4, 2025), https://research.google/blog/discovering-new-words-with-confidential-federated-analytics/

6 Chloé Kiddon and Prem Eruvbetine, Parfait: Enabling private AI with research tools, Google Research (January 22, 2025), https://research.google/blog/parfait-enabling-private-ai-with-research-tools/.

7 Alexey Kurakin and Natalia Ponomareva, Protecting users with differentially private synthetic training data, Google Research (May 16, 2024), https://research.google/blog/protecting-users-with-differentially-private-synthetic-training-data/. See also Alexey Kurakin et al., Harnessing Large-language Models to Generate Private Synthetic Text (January 11, 2024), https://arxiv.org/abs/2306.01684.

8 Fabian Pedregosa and Eleni Triantafillou, Google Research (June 29, 2023), https://research.google/blog/announcing-the-first-machine-unlearning-challenge/. Fully erasing the influence of the data requested to be deleted is technically challenging. Machine unlearning is an emergent subfield of machine learning that aims to remove the influence of a specific subset of training examples — the "forget set" — from a trained model to positively impact user privacy.

9 Natalia Ponomareva and Alexey Kurakin, Making ML models differentially private: Best practices and open challenges, Google Research Blog (May 19, 2023), https://research.google/blog/making-ml-models-differentially-private-best-practices-and-open-challenges/.

10 FAQ about Google Trends data, Trends Help Center, Google, https://support.google.com/trends/answer/4365533?hl=en#:~:text=Google%20Trends%20provides%20access%20to,and%20aggregated%20(grouped%20together) (last visited May 5, 2025).

11 Maps Help Center, Google, https://support.google.com/maps/answer/11323117?hl=en (last visited May 5, 2025).

12 Kamal Janardhan, Simpler data privacy for advertisers with confidential matching, The Keyword (September 12, 2024). https://blog.google/products/ads-commerce/google-confidential-matching-data-privacy/.

13 Alan Stapelberg, It's now easier to prove age and identity with Google Wallet, The Keyword (April 29, 2025), https://blog.google/products/google-pay/google-wallet-age-identity-verifications/. See also Matteo Frigo et al., Anonymous Credentials from ECDSA, https://eprint.iacr.org/2024/2010.pdf.

14 PCC is designed to enable innovative features while keeping the data needed for them confidential from other subsystems. We do this by using techniques such as limiting Interprocess Communications (IPC) binds and using isolated processes. These are included as part of the Android Open Source Project and controlled by publicly available surfaces, such as Android framework APIs. For features that run inside PCC, continuous sensing data is processed safely and seamlessly while keeping it confidential. See Suzanne Frey, Introducing Android's private compute services, Google Research Blog (September 9, 2021), https://security.googleblog.com/2021/09/introducing-androids-private-compute.html, and Pixel Phone Help Center, Google, https://support.google.com/pixelphone/answer/11209263 (last visited May 5, 2025).

15 Jen Fitzpatrick, How we make every day safer with Google, The Keyword (May 11, 2022), https://blog.google/technology/safety-security/how-we-make-every-day-safer-with-google/.

16 Dave Kleidermacher, Dianne Hackborn, and Eugenio Marchiori, Trust in transparency: Private Compute Core, Google Security Blog, (December 8, 2022), https://security.googleblog.com/2022/12/trust-in-transparency-private-compute.html

17 Eugenio Marchiori, Sarah de Haas, Sergey Volnov, Ronnie Falcon, Roxanne Pinto, Marco Zamarato, Android Private Compute Core Architecture, (September 21, 2022), https://arxiv.org/abs/2209.10317

18 https://github.com/google/private-compute-services

19 How Google Workspace Uses Encryption to Protect your Data White Paper, Google Workspace (June 2021), https://services.google.com/fh/files/misc/google-workspace-encryption-wp.pdf.

20 Google regularly releases publications through the Google Research and GoogleDeepMind Research blogs, conference proceedings, as well as posting on external sites such as arXiv.org and https://github.com/google.

21 Miguel Guevara, Enabling developers and organizations to use differential privacy, Google for Developers (September 5, 2019), https://developers.googleblog.com/en/enabling-developers-and-organizations-to-use-differential-privacy/.

22 Google Cloud was an early adopter and investor in Confidential Computing products and services. Confidential Computing involves fortifying the very foundation of our users' compute infrastructure – the Compute Engine virtual machines (VMs). Confidential VMs protect the confidentiality of data in the cloud by encrypting data-in-use while it's being processed. BigQuery data clean rooms provides another example. In this instance, Google Cloud gives customers the ability to analyze their data without exposing individual records. While standard data exchanges provide a way to share data across organizational boundaries at scale, data clean rooms help our customers address sensitive and protected data-sharing use cases. Data clean rooms provide additional security controls to help protect the underlying data and enforce analysis rules that the data owner defines. For customers who want to ensure sensitive data does not leave their boundaries, Cloud DLP which is now part of Sensitive Data Protection, a family of services designed to help you discover, classify, and protect your most sensitive data. With BigQuery differential privacy SQL building blocks, analysts and data scientists can anonymize their data, especially useful to prevent revealing underlying sensitive datasets while still being able to share inferences safely.

23 Vineet Dave and Arun Santhanagopalan, Google Cloud and Swift pioneer advanced AI and federated learning tech to help combat payments fraud, Google Cloud Blog (December 10, 2024), https://cloud.google.com/blog/products/identity-security/google-cloud-and-swift-pioneer-advanced-ai-and-federated-learning-tech?e=48754805.

24 Advertising in the Age of Data Protection: Background for a Proposed Risk-Utility Framework for Novel Advertising Solutions Discussion Draft (v 1.0) (April 2024), https://fpf.org/wp-content/uploads/2024/04/FPF-Proposed-Advertising-Risk-Utility-Framework-April-2024-v1.0.pdf.

25 E.g., the OECD.

26 See e.g., IEEE and ISO. The US National Institute for Standards and Technology recently issued NIST SP 800-226 Guidelines for Evaluating Differential Privacy Guarantees (March 2025), available at: https://csrc.nist.gov/pubs/sp/800/226/final.

27 The GDPR implements this view by limiting the definition of "personal data" to data related to an "identified or identifiable" person. It also expressly acknowledges this issue in Recital 26, which states that the "principles of data protection should…not apply to anonymous information," which is data that cannot reasonably be linked to an identified or identifiable person. See GDPR, Recital 26 ("The principles of data protection should therefore not apply to anonymous information, namely information which does not relate to an identified or identifiable natural person or to personal data rendered anonymous in such a manner that the data subject is not or no longer identifiable. This Regulation does not therefore concern the processing of such anonymous information, including for statistical or research purposes.").

28 Google supports data protection regulatory frameworks that enable organizations to balance individual interests and collective benefits to other users, as well the interests of an ecosystem of service providers, merchants, content creators, rivals, and more. See A Framework for Responsible Data Protection Regulation, Google (2018), https://services.google.com/fh/files/blogs/google_framework_responsible_data_protection_regulation.pdf and Responsible Data Practices, Google (2022), https://services.google.com/fh/files/blogs/google_framework_responsible_data_protection_regulation.pdf.