ISSUE BRIEF

U.S. POLICY

PETs Workshop Proceedings

FUTURE OF PRIVACY FORUM

Executive Summary

The Future of Privacy Forum (FPF) and the Mozilla Foundation hosted a PETs Workshop featuring short, expert panels exploring new and emerging Privacy Enhancing Technology (PETs) applications. A group of technical, legal, and policy experts from industry, civil society, and academia participated in the workshop, discussing the implications of the use cases. Technology and policy experts presented several leading PETs use cases, and the group analyzed how PETs work with other privacy protections and discussed how PETs may intersect with data protection rules.

Key points

- The lack of clear regulatory guidance forces organizations to rely on best practices, trust, and adherence to local laws between/across jurisdictions;
- No one PET or technology can address every risk or threat; in many cases, multiple PETs and privacy protections are beneficial and can help reduce risk;
- Synthetic data and regulatory sandboxes were essential for PETs' development and proof of concept;
- Organizations need more incentives to develop and implement PETs, especially given their high cost and regulatory uncertainty.

Use Cases Presented

- Preventing Financial Fraud Across Different Jurisdictions with Fully Homomorphic Encryption - Mastercard.
- Differential Privacy for End-of-Life Data - Oblivious.

PETs Use Case 1:

Preventing Financial Fraud Across Different Jurisdictions with Fully Homomorphic Encryption

Organization: Mastercard
Presented by: Caroline Louveaux, Chief Privacy Officer, Mastercard
Discussant: Joseph Lorenzo Hall, Distinguished Technologist, Internet Society

Summary: Using Fully Homomorphic Encryption (FHE), Mastercard created an anti-fraud detection system that enables lawful cross-border fraud checks while protecting customers' data privacy.

Financial firms are required and incentivized to engage in a range of anti-fraud programs. Firms often seek to determine whether a particular International Bank Account Number (IBAN) involved in a cross-border transaction is at high risk for fraud. Traditionally, this process involves cooperation between financial institutions that can raise compliance questions and challenges under rules related to cross-border data transfers/data localization, data protection/data privacy, and bank secrecy. This use case uses various privacy-protective techniques to facilitate anti-fraud measures while minimizing compliance risks.

Key features include:

- Fully Homomorphic Encryption (FHE) to engage queries against encrypted data without transferring personal or confidential data;
- A hub structure that "fans" queries to keep confidential the identity of the querying institution and the identities of any institutions that flag the IBAN as high-risk;
- A data minimization feature that encrypts the queries and ensures they only persist for 2 minutes.

Privacy Unit:

- In this use case, the key pieces of information being protected are the sources of both query and response. That is, an entity in one country can ask the others if they've detected fraud, and those countries can respond without either the asking or responding institutions knowing from where or to whom their data are going.

Discussion Points:

- PETs might address a wide range of compliance questions. For example, data localization/ bank secrecy, data protection/data privacy, and cross-border transfer rules are often codified and analyzed in different silos and manners.
- One technology cannot address every risk or threat; in many cases, multiple PETs and privacy protections are beneficial and can help reduce risk.
- Different aspects of the privacy-protective framework address different compliance uncertainties.
- Synthetic data and a regulatory sandbox were used for a proof of concept during system development and testing.
- The high cost of building and deploying PETs and similar frameworks means there is a need for additional incentives for PETs adoption and use of PETs.
- Regulators need to be highly engaged with PETs, understand how they can be used, and what risks and benefits are possible.

What is the key privacy advantage of FHE?

- It allows computations and analytics on fully encrypted data—no access to the data itself is required for computation/analysis;
- Encrypted mathematical representations of data are the points that are queried.
- A benefit is that the query is also encrypted and allows for complete confidentiality;
- The encrypted query is sent to a hub in an encrypted form to participating banks, and then a response is returned from the banks with a signal of "yes" or "no" to identify banking numbers that are or are highly likely to be fraudulent or otherwise involved in fraud;
- It can return indications of fraud and would not even reveal/need to reveal the bank that returned the Y/N signal indicating the likelihood of fraud;
- FHE can also help with key legal requirements, such as data localization and privacy requirements, which is a massive benefit.
    - One of the hardest problems is data minimization because you need to collect data to identify fraud. But how do you protect privacy if you are collecting data? FHE presents a potential in-between solution: collect data that can't be read, only computed.

    - Queries persist for a maximum of 2 minutes before dissolving (the key point and bonus feature of this process/tech is that it is not permanent), which aligns with the spirit of data localization requirements.

How do you weigh the benefit of determining fraudulent information here?

- The validity/accuracy of the process, privacy preservation, and the ability to comply with jurisdictional requirements regarding data protection, privacy, data localization, security, etc.
- Ultimately, the goal was to protect the data—how can this tech squarely advance data protection?

Future Goals:

- Get regulators on board with using PETs, namely FHE and multi-computational analysis;
- Incentivize more companies to use PETs.

PETs Use Case 2:

Differential Privacy for End-of-Life Data

Organization: Oblivious
Presented by: Robert Pisarczyk, CEO & Co-Founder of Oblivious
Discussant: Bennett Hillenbrand, Adjunct Professor, Georgetown University

Summary: Meeting regulatory requirements to delete personal data at the end of the data's lifecycle presents a critical challenge for organizations. Under many privacy laws, companies are obligated to erase personal data once its retention period expires or when users exercise their "right to be forgotten." However, this can result in the loss of valuable macro-level insights and historical trends that could help inform strategic decisions. Differential Privacy can help organizations extract and retain aggregate-level information while ensuring that individual-level data is permanently erased and regulatory obligations are met. By anonymizing data before deletion, Differential Privacy allows businesses to generate summaries, trends, and patterns that do not compromise individual privacy. When data reaches its scheduled deletion, the organization processes it into Differentially Private flat files for future analysis. The flat file is saved, and the underlying microdata is deleted. In this use case, an insurance company uses Differential Privacy to maintain more accurate actuarial tables.

Key features include:

- An insurance organization uses Differential Privacy to create anonymized "snapshots" of important statistics about a population at different points in time that it can keep for future use and analysis, even after the data the snapshots came from is deleted.
- The organization agrees upon the statistics it wants from the data, which are influenced by the intended use and sector (health insurance, life insurance, etc.). A script automates statistical calculations without the need for ad hoc analysis by any data analyst or data scientist.
- Differentially private snapshots are taken at regular time intervals.
- Determining the length of those intervals is critical as it influences the choice of epsilon.
- After the organization applies Differential Privacy to the data, the Differentially Private outputs are securely stored, and the original data is permanently deleted.
- The process requires integration with the organization's enterprise cloud environment, including data connectors and a governance framework to manage privacy budgets.

- Better modeling could lead to better premiums, which would benefit customers in the insurance use cases.

Guiding Questions

- How do you define deletion?
  - Is it the case that deletion seeks to advance a normative outcome, that it is an exercise of a right, and an exercise of privacy? Deletion, thus, may not be a goal unto itself but instead a means. Others may disagree on what 'deletion' is.
- If personal data is deleted but aggregate insights are retained, does that meet statutory "deletion requirements"? Given the value to be derived from aggregate insights, this remains an open question worth considering.
- Must systems be built with "unlearning" in mind? Is that possible in strongly privacy-protected data sets? If so, how would we do it, particularly in the context of differential privacy, where removing someone from a dataset reveals more about them than keeping them in?
- How do anti-fraud efforts intersect with deletion requirements, especially when fraudulent actors request the deletion of personal and aggregate information regarding their attempted fraud?
- How do you pledge to establish trust in this process for data subjects and regulators alike?

Discussion centered on:

- How can we create regulatory clarity when it comes to data deletion?
- Seeking to make data more useful without making it more dangerous.
- Resilience is built through record keeping of past events that occurred, but the key to protecting privacy may be doing so without actually keeping *individual-level* records.
- For model training, if a data subject comes in post-training and says they want their data deleted, does the data need to be "untrained" from a model?
  - There can be a comparison between AI training and consent withdrawal in research ethics. Thinking about a public interest research trial, if a person engages in a research trial and then later withdraws, must you withdraw all the data and insights collected up to that point? There is an argument for "no" – that past data is kept and only future data is not collected – and that the same logic could be applied to data uses and data used for training AI models. When a data subject requests data deletion, there is an imperative to no longer use data for training, but no obligation to un-train or unlearn data. But this is a space that requires significantly more theorizing *and* regulatory guidance.

- If you have data, create a graph or ML model using it, and then get a deletion request and need to remove the data from the dataset, do you then need to adjust the graph or model in relation to the removal of the data from the dataset, or can you maintain those insights?
  - Doing so seems counterintuitive and would cause a set of privacy, utility, and practical problems if necessary (e.g., adjusting a graph or other research output/finding to account for withdrawal from continued participation in research/training).
- Are differential privacy, summary statistics, etc., a way to preserve population-level knowledge? However, given that the theme in the first two discussions was that any one PET is not an end-to-end solution for privacy preservation, would that also apply in this context? What combination of PETs here is sufficient for this kind of work?
  - That said, the key question of this case is whether using Differential Privacy is sufficient for data deletion requirements. Depending on the differential privacy results, other PETs/mitigations will need to be determined later based on the intended use of the data and the threat model.
- From a GDPR standpoint (including precursor work like the Article 29 working paper), Differential Privacy is seen as a way to move through aggregate statistical preservation while maintaining legal compliance.

FUTURE OF
PRIVACY
FORUM

1350 EYE STREET NW | SUITE 350 | WASHINGTON, DC 20005

info@fpf.org | FPF.ORG