Privacy enhancing technologies

WHITE PAPER

FEBRUARY 2024

Contents

Executive summary

Today's financial system is subject to an unrelenting stream of increasingly sophisticated attacks from actors seeking to defraud businesses, steal the identities of consumers, conceal the transfer of criminal gains, circumvent sanctions, and fund terrorist activity. As criminals use the expanding availability of AI to mount ever-smarter attacks, and the proliferation of real-time payment rails to automate the rapid exfiltration of their illicit gains, the safety and security of financial networks depends on continuous investment in cutting edge technologies that can stay one step ahead of these attackers. These technologies all have one thing in common - they are fueled by data, and the more data they have the more effective they are. In a world where attackers skillfully exploit limitations of a given organization's siloed data, that means the cornerstone of an effective defense are systems that draw on collaborative datasets that transcend organizational and jurisdictional boundaries. In other words — divided we are vulnerable, while united we are strong.

But alas, the world is never quite so simple. While breaking down data silos is critical to defending financial networks from attack, other important considerations – from privacy and confidentiality to national data localization requirements – impose limitations on the willingness and ability of organizations to share data. Until recently, the need to navigate those limitations have placed frustrating limits on a financial institution's ability to build collaborative defenses – but all that is starting to change thanks to a new family of technologies called privacy enhancing technologies (PETs).

PETs is an umbrella term for a range of tools and techniques, each with a distinct set of analytical capabilities and limitations. These techniques can break the "see-saw" paradigm of needing to reconcile trade-offs between competing goals, allowing expanded data sharing without sacrificing data security and privacy. Collectively, techniques like zero-knowledge proofs, multi-party computation, homomorphic encryption, and differential privacy have the capacity to unlock collaborations that were once infeasible, while potentially improving rather than trading away data security or privacy.

A mix of public and private institutions ranging from Google to the European Commission are currently exploring the use of PETs and Mastercard is no exception. We are eagerly exploring the use of PETs in new products and services, while remaining committed to Mastercard's Data and Tech Responsibility Principles[1] and the discipline of privacy by design. This includes exploring the use of PETs in new multi-party consortiums focused on fraud and financial crime. Most notably, we recently completed a Proof of Concept within the Singapore Info-communications Media Development Authority's PETs Sandbox, demonstrating how a PET technique called "homomorphic encryption" could be used to enable the sharing of vital financial crime intelligence between Singapore, the U.S., India, and the UK.

Successfully unlocking the capacity for PETs to support the safety and security of the financial system will depend on work by a range of stakeholders across the public and private sector.

While traditional efforts to share such data could be stymied by "tipping-off" rules that limit the sharing of reports on suspicious activity between institutions, our proof of concept demonstrated that it is possible to build a system where banks can share sensitive data about the riskiness of a given account without (1) letting the enquiring institution see any of the responding institutions underlying data, (2) without letting the responding institution see who is enquiring or which account is being queried, and (3) without any of the responding institution's underlying data being transferred across borders or outside of the institution.

While PETs bring exciting and ground-breaking capabilities to the table, much work is still required to fully unlock the potential of these technologies, and to assess their limitations. The novel character of PETs means that existing laws and regulations were not designed with these kinds of capabilities in mind, and as a result it may not be clear how the application of PETs should be treated within these frameworks.

Successfully unlocking the capacity for PETs to support the safety and security of the financial system will depend on work by a range of stakeholders across the public and private sector. Privacy and data protection regulators will play a critical role in providing guardrails around the use of these techniques and may wish to develop regulatory "sandboxes" to facilitate private sector experimentation in a safe and controlled environment. Meanwhile, financial regulators may wish to act as a catalyst for public/private experimentation into how PETs could support improved regulatory and supervisory outcomes, while at the same time encouraging global financial standards setters to consider the treatment of PETs under their existing regulatory guidance regarding cybersecurity, operational, systemic, and third-party risk management. Finally, financial institutions, such as banks, fintechs, and payment service providers, may wish to begin their PETs journey with use-case ideation, exploring where new social or commercial value might be unlocked through access to insights from data sharing that — to date — has not been feasible.

Together, we can unlock the potential for PETs to be a meaningful tool in our collective fight against the increasing speed, sophistication, and geographic scope of financial crime. As a leading innovator in the financial sector, Mastercard is excited to play an active role in exploring the application of PETS against a wide range of use cases and is actively seeking opportunities for collaborative experimentation with both the public and private sector.

Our financial system is under siege

Consumers, businesses, and financial institutions around the world are under siege. Sophisticated attackers are leveraging new technologies and techniques in their never-ending effort to defraud businesses, steal the identities of consumers, conceal the transfer of criminal gains, circumvent sanctions, and fund terrorist activity. Recent years have seen an enormous spike in cybercrime, which is estimated to have increased by 600% since the pandemic,[2] and in the United Kingdom alone it's estimated that individuals were tricked into sending £485.2 million to scammers posing as a legitimate payee.[3]

Effectively confronting these bad actors is essential to safeguarding the security of society at large and to preserving confidence in the security of our financial system. Unfortunately, that task is becoming more challenging every day, because every innovation that makes our lives better — from AI to instant payments — is also a tool in the hands of sophisticated criminals. That means:

Financial crime is getting faster
The global proliferation of real-time payment systems has made instant frictionless payments possible any time of day, 365 days a year, which benefits the accessibility of payments, velocity of money, efficiency of billing and overall rate of GDP growth. Unfortunately, criminals have exploited that same access to low- cost and irrevocable payments to support more sophisticated strategies for laundering money, enabling them to move funds through payment systems faster than many legacy anti-money laundering processes can flag suspicious transactions.

Financial crime is getting smarter
Fraudsters and criminal organizations are finding new ways of using cutting edge technology to confuse and deceive their targets. For example, a spate of recent scams in the United States use artificial intelligence to "clone" the voice of a specific individual. Scammers then target the friends and family of that individual with phone calls, using the cloned voice to desperately request funds to address a dire situation, such as a kidnapping or arrest.[4]

Financial crime is getting more global
While the remit of a given jurisdiction's regulatory and law enforcement apparatus ends at "the water's edge," criminal organizations are increasingly trans-national in character. They use this to their advantage, exploiting organizational divisions and compartmentalization of data to conceal illicit cross-border flows and orchestrate globally coordinated cyberheists. For example, sophisticated cash-out attacks on ATM networks often employ precisely timed attacks where teams of individuals operating across multiple countries make simultaneous withdrawal designed to drain compromised accounts before alarms can be raised. The globally distributed nature of these crimes can make it challenging to mount a coordinated defense, and repatriating the stolen funds is, frustratingly, almost impossible.

Together we stand, divided we fall

Of course, criminals using new technologies to ply their trade is nothing new. The advent of postal mail, telegraphs, and checks all saw the proliferation of unique scams, frauds, and other criminal enterprises. A common theme across these threats is the exploitation of divisions – such as a company's internal departments or a country's national borders – to deceive their victims and conceal the flow of ill-gotten funds. The best defense against these tactics has always been to find ways of breaking down barriers, collaborating across national, institutional, and divisional silos in ways that leave the criminals and funds with no place to hide. As the Financial Action Task Force (FATF) - the global watchdog for money laundering and terrorist financing – observes, "information sharing is critical for combating money laundering, terrorist financing and financing of proliferation."[5]

Mastercard has extensive experience fostering collaboration across the financial system in the service of safeguarding trust in the safety and security of payment systems. In doing so, we're helping to address some of the most pressing challenges facing our industry.
For example:

- **We're convening the payments industry to fight real-time money laundering:** In the United Kingdom we've tackled the challenge of increased speed in money laundering flows by convening twelve banks, who collectively represent the overwhelming majority of UK interbank payments volume, around a collaborative system for maintaining a network-level view of money-laundering activity. This system hampers the ability of criminals to hide flows by moving them between banks and is used to train sophisticated machine-learning models that root out "mule" accounts and identify suspicious transactions in real-time. At the same time, in our global card network, our new AML Account Risk service evaluates credit card numbers for the risk that they're engaging in money laundering. We create money laundering scores at the cardholder level using typologies promoted by FATF and our subject matter experts that indicate potential money laundering activities. This will increase our network security by pinpointing cards with higher risk profiles in order to reduce criminal use of our network based products.

A common theme across these threats is the exploitation of divisions – such as a company's internal departments or a country's national borders...

- **We're building smarter tools to fight fraud:** Our Ekata and NuData services help financial institutions outsmart sophisticated synthetic identity and account takeover strategies by aggregating data from a plethora of sources to provide a single point where banks opening new accounts or authorizing tractions can cross-reference the data being provided against multiple sources. Our Consumer Fraud Risk service looks at complex patterns of behavior on account-to-account payments, understanding the risk of the bank accounts on both sides of the transaction, as well as the relationship between them.
- **We're providing a global view of financial crime:** Our Safety Net system is an ecosystem-level monitoring platform that provides a second line of defense for domestic payment networks, central banks, and members of our card network that is independent of fraud safeguards employed at the financial institution level. The global nature of this system enables financial institutions and central banks within a given country to enjoy the benefits of a model trained on patterns of fraud and cybercrime worldwide, rather than needing to learn exclusively from costly experiences within their own jurisdictional silo.

Stopping fast-moving "money mules" in their tracks with Trace

For years, people involved in organized crime have endeavored to set up or take-over a network of "mule" bank accounts that can be used to obscure the source of their funds to anti-money-laundering (AML) authorities. By breaking funds down into small amounts and dispersing them across accounts at multiple financial institutions, sophisticated criminals can take advantage of the limited vision each bank has into the data of (only) its own customers to make the movement of monies associated with financial crime not to appear suspicious.

Mastercard worked with Pay.UK (the U.K.'s Faster Payments Scheme) to create an industry-wide anti-money-laundering solution to detect and respond to criminal activity across the payments network. The team compiled two years' worth of transaction data from 18 participating financial institutions to build a model of the UK's payments network. Together, they connected nearly 87 million accounts detailing over 357 million individual payment relationships, enabling them to produce a statistical understanding of how suspect money flows through the UK banking network, as well as an algorithm to flag those accounts and relationships that demonstrate suspected mule behavior.

Within weeks of the system going live, this collaborative data infrastructure identified multiple large and well-concealed money laundering rings actively moving funds between networks of accounts and institutions. These insights helped to reinforce banks' fraud prevention processes and reduced the time necessary to identify and take action against suspected accounts.

Sharing isn't always easy

We know that working together to break down data silos is the key to stopping new forms of financial crime in its tracks, but like most solutions it's easier said than done. Sharing data to stop fraudsters and criminals needs to be weighed against competing interests and requirements that may place limitations on what we are willing/allowed to share and who we can share it with.
For example:

- **Companies** may be reticent to share sensitive data with their competitors, fearing that it might provide insights into their proprietary strategies, or allow a rival firm to poach a prize customer.
- **Competition regulators** may also be worried about those same companies getting too cozy when it comes to sharing data, perhaps fearing that a system designed to mitigate fraud could serve a dual-purpose of also facilitating some kind of collusive or anti-competitive behavior.
- **Chief information security officers** might well be reticent to have sensitive data transferred to third parties. After all, once the data has left their organization it might be more difficult to monitor and enforce appropriate data handling practices (access control, encryption, etc.). They might also be justifiably worried that bringing together the sum total of their industry's data in shared repository would create an irresistible target for cybercriminals.
- **Customers** stand to benefit from better protections against fraud and cybercrime, but for all of the above reasons they also have good reason to be concerned about data sharing if it puts their sensitive data at greater risk of being stolen, mishandled, or used to facilitate anti-competitive behavior.
- **Privacy and data protection regulators** play a central role in ensuring that personal data is processed fairly, transparently, and securely, in a way that protects individuals' fundamental rights and freedoms with respect to their personal data. How, with whom and for what purpose personal data is shared and how it is protected end-to-end is central to most privacy and data protection frameworks.
- **Financial regulations** may impose additional specific requirements on financial institution's handling of data, for example outlining specific banking secrecy obligations or imposing limitations on the ability to move certain data "off-soil" to data centers located outside of their jurisdiction.
- **Law enforcement agencies** have legitimate concerns that if criminal actors become aware their financial flows are subject of scrutiny, they may be able to alter their behavior or move their funds before legal authorities have the opportunity to act. Consequently, while law enforcement agencies recognize the potential benefits of breaking down data silos, they also impose strict "tipping-off" rules that limit the sharing of reports on suspicious activity across borders, institutions, and even within the internal divisions of a firm.

But all that might be about to change. Rapid progress is being made on an exciting new family of technologies and techniques, collectively called privacy enhancing technologies or PETs, with the potential to break the "see-saw" paradigm of needing to reconcile trade-offs between competing goals, instead allowing expanded data sharing without sacrificing security and privacy.

So where does this leave us? Breaking down the silos that separate data is essential to enabling a robust defense against the increasing sophistication of fraud and other forms of financial crime, but this process must be done in a way that fully aligns with the concerns detailed above. Historically that has meant that responding to new threats has been a study in trade-offs, a kind of "see-saw" where certain insights that could help mount a defense against financial crime can only be employed by putting at risk competitive interests, data security, or privacy – risks that the financial community is often (rightly) unwilling or not permitted to take.

Introducing privacy enhancing technologies (PETs)

In the same way that the term "artificial intelligence" serves as an umbrella term for a wide variety of techniques for extracting insights from data (e.g., machine learning, neural networks, large language models, etc.) the term Privacy Enhancing Technologies captures a range of tools and techniques, each with a distinct set of analytical capabilities (and a distinct set of limitations). Broadly speaking, PETs are a set of emergent technologies and techniques that help to operationalize fundamental data protection principles by minimizing personal data use, transforming data in privacy-preserving ways, and/or maximizing data integrity, confidentiality, and security. When applied appropriately, PETs can help meet data protection requirements while unlocking data utility.

In order to understand how to properly use PETs to achieve these objectives, it is important to understand what these technologies can do. Delving into the details of how exactly these technologies work may seem daunting to those without an advanced degree in mathematics or computer science – but achieving a working knowledge of what they can do is both necessary and within everyone's grasp. Just as important as understanding how these technologies work, is understanding what these technologies can do to help advance public and private sector objectives. At a high level, it helps to think of what PETs can do in three categories:

1. **Some PETs let you share an output similar to the data, rather than the data itself:** Imagine that you had a window into a parallel universe that was eerily similar to our own, but not quite the same. Your mirror-world counterpart might cheer for a different sports team and live in a different neighborhood, while a friend's counterpart might have a different job and favorite food. Those are important differences at the individual level, but you would still expect that the average price of a house in your neighborhood, overall popularity of pizza, and even the profiles of those most likely to fall victim to financial fraud, in the mirror world would look very similar to our own universe. One sub-category seeks to build something a lot like that window into a parallel universe – preserving the utility of the overall dataset while protecting the privacy of individuals within the underlying data. One technique called **differential privacy** introduces noise into the insights that you can extract from the underlying data, designed to limit the ability of an outsider to identify information about specific individuals while sharing useful insights about a group or the data set as a whole. Differential privacy goes a step beyond simple data aggregation by using a sophisticated mathematical framework to protect the privacy any individual within the data set. Another technique called **synthetic data** takes a different approach, using a carefully calibrated computer model to create a simulated data set that approximates the statistical properties of the original data set.

For example: U.S. Census Bureau collects vast quantities of data in the pursuit of its mission to "serve as the nation's leading provider of quality data about its people and economy", but as much of this data is highly sensitive, it is required by law to ensure that it does not release information that could be used to identify an individual within the statistics it publishes. Meeting this obligation is more challenging than it might sound in a world where powerful computers make it much easier to match specific datapoint across massive datasets in ways that make it possible for an individual's census responses to be connected back to them. So easy in fact that a simulated re-identification attack conducted by the US Census Bureau conducted in 2021 using 2010 census data was able to confirm re-identification of the responses of a staggering 179 million respondents (58% of the same population).[6] In response the U.S. Census Bureau has leveraged differential privacy to modernize their practices for introducing statistical noise into statistics released as part of the 2020 Census.[7]

2. **Some PETs let us derive specific pieces of information or insights from a dataset without seeing all of the underlying data:** There are many reasons to be worried about sharing data with a third party, even one that you trust. Another category of PETs can provide ways to let others derive specific pieces of information or insights from your data without letting them see the entirety of the underlying data. Take, for example, a situation where you hold a data set and another party wants you to prove the truth of a statement about the dataset without disclosing the data itself. A technique called **Zero knowledge proofs** allows you to demonstrate this to a degree of mathematical certainty that satisfies the other party. This is a little bit like proving to someone that you know the combination to a safe by locking it and then re-opening it without ever actually sharing the combination to the safe itself. In other circumstances, two different people might have separate data sets and wish to exchange insights based on the combination of those two data sets, without disclosing the entirety of their underlying dataset to the other party or knowing what combined insights the other party asked for.

   Homomorphic encryption does this by encrypting both the underlying data and the query so that analysis can be performed, without the underlying data being visible to the asking party, and without the query itself or the outcome of that query being visible to the other party. In a larger data ecosystem, this could allow a centralized entity, such as a utility, to collect complementary data sets and make it possible for multiple participants to query the combined data without sharing the entirety of an unencrypted dataset with the whole group, and without revealing the queries made by an individual participant to the other participants. This also helps mitigate cybersecurity risks by preventing attackers from accessing data while it is being processed in unencrypted form. As discussed in more detail later in this paper, zero knowledge proofs and homomorphic encryption are computationally intensive and may sometimes be difficult to implement at scale. For more complex data sharing operations, more limited forms of secure information sharing or "data clean rooms" may be a helpful intermediate step.

For example: As part of its Homomorphic Encryption Applications and Technology (HEAT) project the European Commission explored the application of a version of homomorphic encryption to the automated detection of organized crime. The experiment enabled data from multiple databases to be aggregated within the cloud after being homomorphically encrypted, at which point authorized users were able to construct statistical scans for signals of organized criminal behavior across the fused databases.[8] The experiment successfully demonstrated that a French user could perform search queries on an encrypted German database without directly viewing the data, and without the queries themselves being visible to German authorities. While not yet implemented at scale, the experiment was deemed to provide a solid proof of concept for the sharing of data between EU countries to solve crimes, while limiting access to financial crime data (and the queries executed against such data) to those with a legitimate need to know.

3. **Some PETs let us build shared tools without ever sharing the underlying data used to train those tools:** A third category of PETs allows us to ditch the paradigm of bringing data to a central point for analysis, and instead enabling the analytics to go to the data. This can help to do at least two things: first, it can unlock many of the benefits of data sharing without ever requiring the data to leave its organizational or jurisdictional silo (or in some cases, without leaving an individual's device); second, it can add security benefits by eliminating the need to create a centralized data repository that would be an attractive target for hackers. One example of this technique is called **federated learning**, which allows individual organizations to train discrete analytical models locally based on one dataset and then combine that model (but not the data that was used to train it) with models locally trained on other data sets. This technique can be used within a single organization to train a model in a decentralized way across disparately held data sets, or it can be used across different organizations as a way of enhancing shared tools without necessarily sharing the underlying data.

   The result is a stronger model that can be much more effective than an algorithm trained centrally on a local (but limited) data set. Another technique called **secure multi-party computation (MPC)** allows the job of data analysis to be broken up among multiple parties, configuring the task in such a way that calculations spanning multiple datasets can be executed without requiring any sensitive data to be shared between parties.

As various PETs are tested, layered, and combined thoughtfully, these technologies have the capacity to unlock collaborations that were once infeasible and data utility yet unrealized, while potentially improving rather than trading away cyber resilience, privacy, and competition. However, as we will discuss later in the document, the application of PETs must be considered in line with the objective to be achieved and the limitations of the technology available, supplementing the technology as appropriate with administrative controls, contractual limitations, auditing, and other traditional measures for enhancing the security of your data and ensuring accountability over its use.

For example: Google currently uses federated learning to train the speech models that power its "Google Assistant" offering without ever moving audio data to Google's central servers. It does this by saving relevant audio data directly to the device and locally deploying a training algorithm on the data to learn how the speech model could be improved. It then sends a summary of those changes to Google's servers where it is aggregated with changes sent from other user's devices to improve the quality of the overall speech model for all users.[9]

Putting PETs to work at Mastercard

At Mastercard, we're eagerly exploring the application of PETs to numerous aspects of our business, as well as to the creation of new products and services. We are committed to being responsible data stewards, consistent with Mastercard's Data and Tech Responsibility Principles and the discipline of privacy by design.[10]

We also recognize that PETs are new and there is much yet to be learned about their capabilities and limitations. Exploring and testing new use cases — including in regulatory sandboxes, industry groups, and other collaborative settings — will be essential to developing meaningful standards and driving smart adoption of these technologies.

Informed by these principles, we are experimenting with the use of synthetic data for a variety of use-cases across our organization. We are exploring how the use of synthetic data in place of real data can deliver more robust models while also reducing risks to data security and privacy. For example, we have found that synthetic data can be highly effective in augmenting risk analysis in emerging market jurisdictions where existing data quality may be poor. We are also considering how synthetic data can enable technical teams all over the world to collaborate to solve problems specific to a jurisdiction without needing to egress the data from that jurisdiction.

Using synthetic data in this way may also enables insights from markets with localization requirements to be incorporated into global and regional models without ever needing the data to leave its country of origin. Furthermore, we believe that synthetic data has the potential to play an important part in both training AI models and may provide an additional tool to help correct for AI bias.

In addition to our work on synthetic data we are actively exploring the potential for federated learning and homomorphic encryption to support the creation of new multi-party consortiums focused on fraud and financial crime. The remainder of this section will focus on a deep dive into the recently published results of one such proof-of-concept (POC) undertaken by Mastercard in mid-2023 as part of the Singapore Info communications Media Development Authority's PETs Sandbox. This project explores the use of homomorphic encryption - specifically a technique called "Fully Homomorphic" Encryption or FHE using an implementation with software provided by Duality Technologies - to share vital financial crime intelligence between Singapore, the United States of America, India and the U.K.

We believe that synthetic data has the potential to play an important part in both training AI models and may provide an additional tool to help correct for AI bias.

This effort explores a use-case in which one financial institution (the enquiring entity) wishes to determine if a specific international bank account number (IBAN) has been flagged as high-risk by any one of several other financial institutions (the source entities). Historically, financial institutions have been reticent to share such information for fear of contravening "tipping-off" rules, as well as concerns about unintentionally disclosing private or competitively sensitive information. In instances where one or more of the source entities are in different jurisdictions the problem is further complicated by differences in privacy, data protection, and banking secrecy rules. The net result is a lack of data sharing that leaves organizations only able to observe activity or information obtained from their own systems, making organizations blind to criminal actors who exploit those boundaries.

The implementation of a financial intelligence sharing system based on FHE has the potential to address some or all of the impediments to data sharing by establishing a system where source entities can share sensitive data about individuals and businesses (in this case whether their IBAN has been flagged as high risk) without (1) the enquiring entity gaining access to the underlying data or (2) the source entity having visibility into which IBAN is being queried or who the enquiring entity is.

To test the viability of this hypothesis we populated test databases in Singapore, Europe, North America, and India with synthetic IBAN watch lists and established a hub program to facilitate both the distribution of queries to source entities and the aggregation of results. The POC employed FHE software (provided by a third-party supplier) which was deployed on the IT systems of enquiring and source entities.

Figure 1 illustrates a sample query in which:

A Singapore based enquiring entity submits a query to determine if a given IBAN has been flagged as risky by a source entity in either Europe, North America, or India.

- The query is homomorphically encrypted using a public key held by the enquiring entity and is submitted via a hub to the source entities.
- The hub distributes the encrypted query and public key to all of the source entities participating in the system.
- Using the public key, the encrypted query is executed against the data without ever decrypting the query, and an encrypted result is produced.
- The source entities respond with the encrypted result. The source data remains under the control of the source entities at all times and never leaves their data centers.
- The encrypted results are sent back to the hub where they are homomorphically aggregated and sent back the enquiring entity.
- Only the enquiring entity is able to decrypt the result using its private key.

Figure 1: Logical diagram showing how keys, queries, results, and source data flow through the system. Encrypted payloads and operations are outlined in bold and marked with a lock symbol. The query (red) is encrypted with the public key generated by the inquiring entity at request time (teal). The homomorphic operations are executed on the data and the encrypted query at each source data node using the public key, passing encrypted results back to the hub (green). The hub performs the final homomorphic operation to aggregate the results and send them back to the inquiring entity. The inquiring entity uses the private key (orange) to decrypt the results.

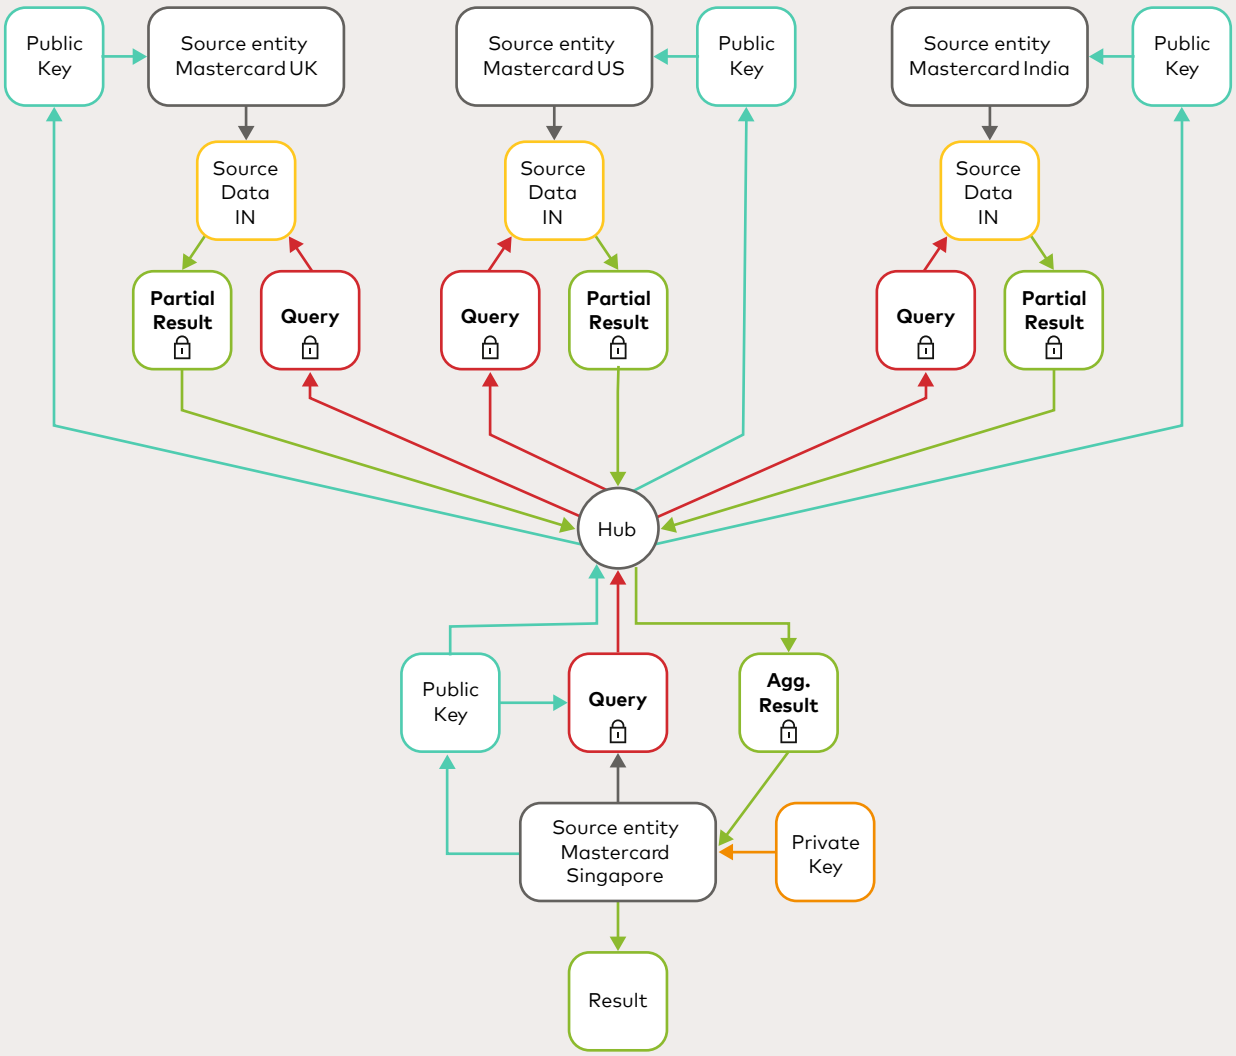

Our successful completion of this POC for a minimum viable product use case demonstrates that it is possible for an inquiring financial institution to gain insights into the riskiness of IBAN without disclosing to the source entities who is asking the question or the subject of the query. In other words, the source entities do not know who is asking the query, nor can they see the specific IBAN. Moreover, the source entities are able to support financial intelligence sharing without disclosing any source data to the enquiring entity or transferring any information within its database across borders.

Critically, the process is bi-directional, meaning that participants in the system can both provide value by serving as source entities and derive value by submitting queries as enquiring entities.

We believe that systems like this one could be the future of financial data sharing and critical to enabling new ways of responding to financial crime. We plan to continue to evolve this POC and to explore the usefulness of PETs in addressing a range of other use cases.

Understanding the limitations of PETs

While PETs bring exciting and ground-breaking capabilities to the table, it's important to understand that they have limitations and can be misused, just like any other technology.

For example, the increased privacy or security provided by many PET techniques comes at a cost of increased computational complexity and intensity. This means that processing data via a PET technique can increase the time required to complete a computation and the cost (in terms of computing cycles) required to complete that computation. This may be particularly true in instances where multiple PET techniques must be combined or layered in order to achieve a desired set of outcomes. While extensive research is underway to identify ways to increase the speed and reduce the costs of deploying PET, current instances of the technology may struggle to accommodate use-cases that must be scaled to accommodate large volumes or low latencies.

The novel character of PETs also means that existing laws and regulations were not designed with these kinds of capabilities in mind, and as a result it may not be clear how the application of PETs should be treated within these frameworks. For example, how do PETs relate to privacy and data protection laws, including limitations on cross-border data transfers or anonymization of personal data? How should competition regulators consider the risks associated with the deployment of new industry collaborations centered around PETs? How should financial regulators consider the services of specialized PET service providers within the context of their third party risk and resiliency standards? Where might existing legislation or regulatory guidance need to be updated to be able to unlock the enormous opportunities of this new technology? Where might regulatory safe harbors be appropriate in order to allow for the testing necessary to fully assess the capabilities of PETs?

Addressing these questions is a significant undertaking that demands extensive public/private collaboration that will likely need to be periodically revisited as the capabilities and potential application of PETs continues to evolve. Collaboration will likely be required across sectors and different regulatory competencies. For example, even if a full-scale financial intelligence sharing system with a PETs-based architecture would be consistent with data protection, data localization, and cross-border data transfer laws of the countries in question, it may not be consistent with the existing banking secrecy requirements in one or more jurisdiction.

Finally, the use of PETs in shared data infrastructure faces an impediment common to all shared infrastructure projects, irrespective of their underlying technology - catalyzing and channeling collective action between organizations with many competing interests and priorities. Shared data infrastructure requires significant investment by founding organizations, often with uncertain gains and the risk that shared governance, liability, and in some cases, dispute resolution mechanisms for the new infrastructure. Companies are also facing an important challenge when considering the integration and selection of the most appropriate PETs for their use cases. While these challenges may sound trivial compared to the technical and legal impediments discussed above, they represent a serious stumbling block that can keep promising ideas from moving beyond the "whiteboard" stage of development.

Recommendations

While the limitations discussed in the previous section mean that full-scale deployment remains rare (particularly for less mature techniques like homomorphic encryption), significant research and experimentation are underway in both the public and private sectors. PETs represent a remarkable opportunity, but their successful deployment is far from a forgone conclusion. Successfully unlocking their potential to support use-cases within financial services, and beyond, will be contingent on effective collaboration between a range of stakeholders across the public and private sector. The work to be done by each group will be complex and multi-faceted, but we offer here a few recommendations of specific areas where key groups might choose to focus their efforts.

1. **Privacy and data protection regulators** face perhaps the most significant challenges and opportunities in responding to the emerging capabilities of PETs. At this early stage, we believe that regulatory approaches that promote additional research and development will help identify where the biggest opportunities exist to strengthen privacy and data protection frameworks using PETs. Regulatory sandboxes allow companies to experiment with PETs in a safe and controlled environment while resulting in learnings that can be used to inform future regulatory standards. Additionally, before PETs become a standard technique in the privacy and data protection toolkit, the workforce at large needs to develop a bigger talent pool of privacy professionals well-versed in PETs. Greater investment in public-private partnerships that promote knowledge transfer between academia, industry, SMEs, and public authorities should be encouraged. Finally, pragmatic guidance and clarity on how PETs can meet various privacy and data protection legal requirements will be needed to drive the multi-year investments required for organizations to apply these technologies effectively.

2. **Financial regulators and law enforcement** can be direct beneficiaries of the application of PETs by financial institutions and financial market infrastructure where the use of these technologies supports their core objectives, such as the safety, security, stability, and compliance of the financial system. Where this is the case, financial regulators may wish to act as a catalyst for public/private experimentation into how PETs might be used to support improved regulatory and supervisory outcomes. At the same time, we would encourage financial regulators and international regulatory standards setters to consider the treatment of PETs under their existing regulatory guidance regarding cybersecurity, operational, systemic, and third-party risk management. Where the application of existing frameworks to the use of PETs may not be clear, we would encourage regulators to take steps to reduce uncertainties that might otherwise limit private sector exploration and investment into these technologies.

3. **Financial institutions**, such as banks, fintechs, and payment service providers, may wish to begin their PETs journey with use-case ideation, exploring where new social or commercial value might be unlocked through access to insights from data sharing that – to date – has not been feasible. Financial institutions may also be uniquely positioned to collaborate with policymakers on identifying potential regulatory barriers to the application of PETs, as well as proposing ways in which those regulation could be modified to encourage the responsible use of PETs, while still preserving the intended regulatory outcomes. Where the application of one or more PET appears to have the potential to deliver meaningful value, financial institutions may wish to conduct preliminary experimentation – similar to the POC described in this paper – to validate the feasibility of the use case in question. Where these use cases involve transaction flows or other forms of payment data, Mastercard would be eager to open a dialogue on the subject.

Together, we can unlock the potential for PETs to be a meaningful tool in our collective fight against the increasing speed, sophistication, and geographic scope of financial crime. As a leading innovator in the financial sector, Mastercard is excited to play an active role in exploring the application of PETS against a wide range of use cases and is actively seeking opportunities for collaborative experimentation with both the public and private sector.

To learn more about how Mastercard is helping businesses, governments, and society use data to better serve their stakeholders check out our recent Signals report on "Digital footprints", which explores how secure data sharing can help to break down data silos and empower consumers to have greater access, more control and benefit from their financial data.[11]

For more information contact
PrivacyEnhancingTechnologies@mastercard.com

Sources

1. Mastercard (2019) 'Mastercard Data Responsibility Principles', available at: https://www.mastercard.us/en-us/vision/corp-responsibility/data-responsibility.html (accessed October 24, 2023).
2. Mastercard (2023) 'Securing the digital economy', available at: https://www.mastercard.us/en-us/business/large-enterprise/safety-security/cybersecurity/cybersecurity-whitepaper.html (accessed October 24, 2023).
3. Payment Systems Regulator (2023) 'APP scams', available at: https://www.psr.org.uk/our-work/app-scams/ (accessed October 24, 2023).
4. Federal Trade Commission (2023) 'Scammers use AI to enhance their family emergency schemes', available at: https://consumer.ftc.gov/consumer-alerts/2023/03/scammers-use-ai-enhance-their-family-emergency-schemes (accessed October 24, 2023).
5. Financial Action Task Force (2017) 'FATF Guidance: Private sector information sharing', available at: https://www.fatf-gafi.org/content/dam/fatf-gafi/reports/Private-Sector-Information-Sharing.pdf (accessed October 24, 2023).
6. US Census Bureau (2021) 'Appendix B – 2010 Reconstruction-abetted re-identification simulated attack', available at: https://www2.census.gov/about/policies/foia/records/disclosure-avoidance/appendix-b-summary-of-simulated-reconstruction-abetted-re-identification-attack.pdf (accessed October 24, 2023).
7. US Census Bureau (2021), 'Differential Privacy Press Kit', available at: https://www.census.gov/newsroom/press-kits/2021/differential-privacy.html (accessed October 24, 2023).
8. European Commission (2015) 'ICT-644209 HEAT - Homomorphic Encryption Applications and Technology, D.1.3, ADOC Use Case', available at: https://ec.europa.eu/research/participants/documents/downloadPublic?documentIds=080166e5a59bfeb8&appId=PPGMS (accessed October 24, 2023).
9. Google (2023) 'Your voice & audio data stays private while Google Assistant improves', available at: https://support.google.com/assistant/answer/10176224?sjid=10933388604693694490-NA#zippy=%2Chow-federated-learning-protects-your-data%2Chow-google-assistant-improves-with-federated-learning (accessed October 24, 2023).
10. Mastercard (2019) 'Mastercard Data Responsibility Principles', available at: https://www.mastercard.us/en-us/vision/corp-responsibility/data-responsibility.html (accessed October 24, 2023).
11. Mastercard (2023) 'Digital Footprints', available at: http://innovationinsights.mastercard.com/digital-footprint-mastercard-signals/p/1 (accessed October 24, 2023).

Designed by Mastercard Creative Studio