NAI
PRIVACY, ACCOUNTABILITY & TRUST

A Primer on Privacy Enhancing Technologies (PETs)

October 2025

TABLE OF CONTENTS

I. Introduction and Overview

The NAI's Primer on Privacy Enhancing Technologies (PETs) is intended to help privacy professionals in digital advertising understand the available methods to protect and enhance the privacy of Personal Data. Our goal is to educate non-technical practitioners about how these methods work and when to consider adopting them for companies' processing of Personal Data. We hope this primer will also assist companies that are processing Personal Data in evaluating the key benefits and trade-offs of using these methods to:

- Mitigate privacy risks to consumers;
- Minimize over-exposure of valuable or confidential data; and
- Reduce the risk of abuses or unauthorized uses of personal or otherwise confidential data.

This primer is directed primarily to a **non-technical** audience to help them understand the fundamentals of different PETs being used in the market today and to demystify how they work. It provides an overview and analysis of the following methods for processing Personal Data:

- Trusted Execution Environments (TEEs)
- Multiparty Computation
- Differential Privacy
- Zero-Knowledge Proofs

For each technique listed above, a dedicated section provides a simple and accessible explanation of key concepts underpinning the technique, as well as both general use cases and advertising use cases to illustrate how the method works in practice. Although there are important differences among these techniques, they are often employed for the shared purpose of allowing businesses to leverage information derived from Personal Data while minimizing the risks.

PETs often rely on a trusted third party to restrict processing that will be applied to inputted data, to filter or aggregate the output from their system, and/or to validate overlap or possession of shared information. Some vendors that employ these or other PETs are billed as "clean room" providers or "data collaboration platforms";[1] however we will forgo use of either of those terms in this primer and instead focus on the specific PETs listed above that different vendors may employ. Some larger organizations may also employ internally-managed PETs to manage or restrict data access and use across the same organization.

To help provide context for the specific PETs addressed in this primer, this Introduction will first review ideas and concepts that apply generally across different methods as follows:

A. Why Use PETs?
B. Limitations of Using PETs
C. Using PETs in the Digital Marketing Lifecycle; and
D. Safeguarding Privacy through Data Transformation.

[1] *See, e.g.*, IAB Tech Lab, Data Clean Rooms, Guidance and Recommended Practices (July 5, 2023) at 10, https://iabtechlab.com/blog/wp-content/uploads/2023/06/Data-Clean-Room-Guidance_Version_1.054.pdf ("A data clean room is a secure collaboration environment which allows two or more participants to leverage data assets for specific, mutually agreed upon uses, while guaranteeing enforcement of strict data access limitations for e.g., not revealing or exposing the personal data of their customers to other parties.")

While some terms used throughout this primer may differ from statutory definitions, which may also vary by jurisdiction, they are used consistently in this document. Please refer to the Glossary in the Appendix of this document for more information.

A. Why Use PETs?

Broadly speaking, organizations adopt PETs to meet different **data governance** goals. Data governance is an umbrella concept that refers to an organization's ability to set policies for its processing of Personal Data and the procedures for ensuring those policies are followed in practice. Three key organizational goals for good data governance include the **privacy**, **confidentiality**, and **security** of data an organization is processing.[2] The use of PETs—along with appropriate organizational measures—can help organizations improve data governance by mitigating risks associated with the confidentiality of business information, uses of data that can raise privacy concerns, and unauthorized access of data that pose security concerns.

1. *Data Confidentiality*

In some cases, a business's choice to employ PETs relates to a business goal for data **confidentiality** rather than compliance goals related to privacy or security. For example, a business may have certain business-proprietary or trade secret information that it does not want competitors to learn, even though collaborating on insights from that data may be beneficial to both parties. In those cases, PETs may be leveraged to restrict certain information from being shared with or accessed by additional recipients such as business partners or vendors while still meeting business goals. This may be true even if there are no privacy or security compliance barriers to sharing that information.

2. *Privacy*

In other cases, PETs may be employed to help manage **privacy** risks or to help meet privacy obligations. Privacy objectives are generally focused on consumer rights and business obligations related to a business's processing of Personal Data and are usually focused on limiting processing to intended, permissible uses. While PETs do not help facilitate consumer-initiated signals to processing entities (e.g., opt-out signals related to tailored or targeted advertising), they can help organizations abide by their obligations. For example, an organization may need to limit how it is processing a consumer's Personal Data after that consumer opts out of targeted advertising or cause that consumer's Personal Data to be deidentified or deleted when requested. Further, a business may need to ensure that consumer Personal Data is deidentified before allowing other data controllers to access it or to analyze it for purposes beyond those disclosed to consumers in a privacy notice.

[2] While there are other organizational goals good data governance typically promotes – such as management of data quality (reliability and accuracy), and metadata management (enhancing utility) – this primer focuses on privacy and confidentiality as two key topics that are more relevant for the application of PETs and compliance with data protection laws. Security objectives usually relate to preventing unauthorized or unintended access, or monitoring use of an organization's data assets. While PETs do not necessarily enhance the security of data held by an organization, they can limit what is accessed by other recipients.

B. What are some Limitations of Using PETs?

Although there are many potential data governance benefits that may accrue from using PETs, businesses should also keep their limitations in mind. Even when using PETs, it is not possible to avoid all vulnerabilities to confidentiality, privacy, or security. Data protection rules also generally focus on whether the mitigation measures put in place by an organization are reasonable in relation to risk. For example, legal and regulatory requirements may vary given the size of the organization[3] as well as the cost and complexity of the measures instituted balanced against the likelihood and severity of foreseeable risks.[4] In addition, the trust associated with recipient organizations that process data often relies on contractual guarantees rather than technology, even when the recipient organizations are chosen because they will employ PETs.

Organizations should also keep in mind that compliance with specific legislation and regulations requires a context-specific analysis relating not only to what data is being collected and processed for specific purposes but also which organizational measures and data protection safeguards have been put in place. While PETs can help mitigate risk, context-specific analysis is also required for legal compliance and handling data responsibly. When safeguarding data processing with a PET provider, the organization remains accountable for its own compliance with data protection regulations. For example, using a platform that allows two controllers to collaborate with data using privacy controls does not absolve each controller of responsibility for how those privacy controls are configured and used. Businesses should also exercise caution in the claims they make about the privacy benefits of using PETs in marketing materials and privacy policy notices.[5]

Further, while implementing PETs involves technical resources, including specialized technologies and vendors, the decision to apply a particular solution is often a **business-initiated** safeguard, with goals established by marketers' objectives and driven by security and privacy compliance leaders, rather than being introduced and driven by IT/Engineering departments. There is also a fundamental business tradeoff with complexity and cost when applying certain approaches that can adversely impact the accuracy and/or usefulness of the data outputs.

It's also important to understand which common privacy concerns PETs can or cannot address. While consumers may have privacy concerns about how their Personal Data might be used by a business, PETs may not be intended to address all of those concerns. For example, use of a PET may prevent a consumer's identity from being exposed to specific recipients; but may not address how that consumer's information is processed within an organization. Consumers often have privacy rights that they can exercise with businesses under state or federal law to address those

---

[3] *See*, *e.g.*, Cal. Civ. Code § 1798.140(d) (exempting from the requirements of the California Consumer Privacy Act (CCPA) smaller businesses whose annual gross revenues are less than $25,000,000 per annum unless it "annually buys, sells, or shares the personal information of 100,000 or more consumers or, households.")

[4] *See*, *e.g.*, *Id.* § 1798.140(a)(15) ("The factors to be considered in determining when processing may result in significant risk to the security of personal information shall include the size and complexity of the business and the nature and scope of processing activities.")

[5] Staff of Office of Technology, FTC Launches New Office of Technology to Bolster Agency's Work, Federal Trade Commission (Feb. 17, 2023) https://www.ftc.gov/news-events/news/press-releases/2023/02/ftc-launches-new-office-technology-bolster-agencys-work

concerns (such as right to request deletion of Personal Data), but PETs on their own may not prevent or address those concerns.

C. When are PETs Helpful for Digital Advertising Use Cases?

Examples of Employing PETs at Different Stages of the Marketing Lifecycle

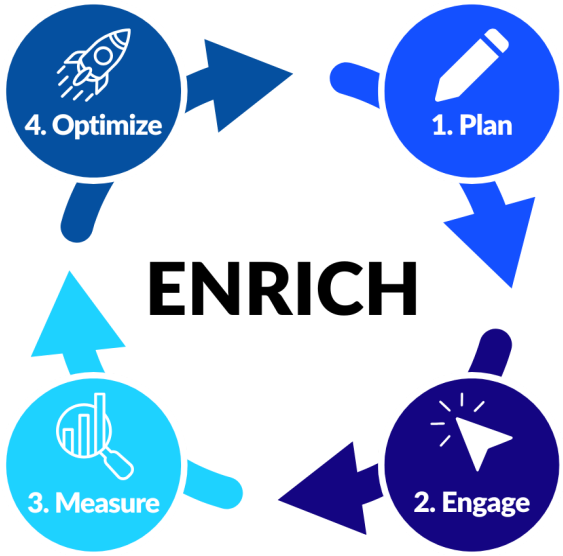

Before explaining how specific PETs work, it is helpful to zoom out and examine how they may generally fit into the use and transfer of data in a typical marketing lifecycle. Campaign marketing is fundamentally an iterative process of planning, execution, measurement, and optimization – leveraging findings from the process to improve future decisions on how to allocate media budgets to achieve effective results. At certain stages of the marketing lifecycle, different PETs may be better suited to achieve the desired outcomes while incorporating the desired elements of data governance described in this primer.

Marketing Stage	Example Activities	Relevant PETs
1. Plan	• Compare overlap of consumer segments available given targeting tactics ("reach")	• Multiparty-compute • Trusted Execution Environments
	• Determine budget allocation to focus media spend on specific tactics ("targeting") and assess the associated addressable audience reach/coverage	• N/A insofar as PETs focus on safeguarding event-level data and planning is conducted using aggregate-level data
2. Engage	• Match paid content to specific audience segments in specific contexts, while limiting frequency of exposures	• N/A insofar as PETs focus on safeguarding event-level data by producing aggregated outputs and engagement may require transfers/use of record-level match keys[6]
3. Measure	• Event-level data collection to determine fraud, billing, campaign pacing, exposure frequency, interaction data (e.g., clicks), and attributable outcome events (sales) • Aggregate analytics of campaign & channel effectiveness (e.g., media mix modeling)	• Multiparty-compute & Federated Learning & Homomorphic Encryption • Trusted Execution Environments
4. Optimize	• Refine campaign based on measurement events and findings that correlate to data enrichment	• N/A insofar as PETs focus on safeguarding event-level data by producing aggregated outputs, while optimization models generally require transfers/use of event-level data for training
Enrich (as needed)	• Add attributes to the match keys associated with the desired targeting dimensions of audience, context, device and geography • Validation queries that two or more processing entities share a common match key or information linked to that match key	• Trusted Execution Environments • Zero-Knowledge Proofs

[6] With respect to PETs, the utility of aggregate data outputs for model training are suspect. In one experiment, over 100 data scientists were unable to train effective models without event-level data. *See* Alexandre Gilotte, Results from the Criteo-AdKDD-2021 Challenge, MEDIUM (Sep. 27, 2021) https://medium.com/criteo-engineering/results-from-the-criteo-adkdd-2021-challenge-50abc9fa3a6

D. Safeguarding Privacy through Data Transformation – How PETs Help with Data Governance for Digital Advertising

The table below illustrates how PETs along with organizational measures can transform different types of input data into Aggregated or Unaggregated output data. Depending on whether the recipient has appropriate organizational measures to protect individuals' identity against reidentification, these techniques may succeed in transforming Personal Data into deidentified or anonymous data in that organization's hands. Appropriate technical and organizational safeguards (represented by the Arrow icon) can transform the input data into output data that poses lower reidentification risk to individuals represented in the input data.

The color coding represents the relative privacy risk associated with the data where Personal Data that remains directly-identifiable is highlighted in orange, pseudonymous data in yellow, and aggregated or non-personal data in green. These are generalizations, and whether privacy-enhancing transformations applied to Personal Data succeed in producing "pseudonymous" or "deidentified" data under applicable laws depends on the specific facts and circumstances of an organization's own data protection measures.

<table>
<tr><th>Input Data</th><th>Type of Input Data</th><th colspan="3">Output Data Status</th></tr>
<tr><td></td><td></td><td colspan="2">unaggregated</td><td>aggregated</td></tr>
<tr><td rowspan="2">Personal Data</td><td>Directly Identified Unaggregated
Personal Data
(e.g., first and last name; unhashed email address)
↓</td><td>When unaggregated output does not have appropriate organizational measures to protect individuals' identity against reidentification it remains Personal Data.</td><td>→</td><td rowspan="2">When aggregated output has appropriate organizational measures to protect individuals' identity against reidentification it may no longer be Personal Data in that organization's hands</td></tr>
<tr><td>Pseudonymous Unaggregated
Personal Data
(e.g., MAID, hashed email address)
↓</td><td>When unaggregated output has appropriate organizational measures to protect individuals' identity against reidentification it may still be Personal Data in the sending organization's hands</td><td>→</td></tr>
<tr><td rowspan="2">Non-personal data provided that all appropriate organizational measures are in place to keep it deidentified</td><td>Non-personal DeIdentified Unaggregated Data
(e.g., log-level auditing data)
↓</td><td rowspan="2">Remains non-personal data</td><td rowspan="2">→</td><td rowspan="2">Remains non-personal data</td></tr>
<tr><td>Non-personal Anonymous Aggregate Data
(e.g., aggregate population size, foot traffic statistics, billing records)</td></tr>
</table>

II. Privacy-Enhancing Technology Assessments and Analyses

A. Trusted Execution Environments

1. Overview

A Trusted Execution Environment (TEE) can be useful for allowing two or more parties to join, analyze, or otherwise compare their datasets. A TEE is a centralized computing environment that enables a data controller to limit the ways a dataset it controls may be processed. This can allow a data controller to reduce the risks of unauthorized manipulation and use of the data, as well as limit secondary uses of data intended only for specific purposes within this environment. It also provides controllers with enhanced audit capabilities to mathematically prove that processing happened as expected. This can help enable data collaboration without unintended additional processing of Personal Data.

A TEE is an area in a computer hardware device that is kept separate from other computer software and hardware components that enables the functional protections described above. The processes running in a TEE can execute and analyze data without exposing that data to the rest of the computer system, improving data security and privacy. A TEE can be located on a personal computer, on a mobile device, or on a cloud-based server. However, where processing occurs is not relevant to what process is occurring or what data is being processed.[7] A TEE limits access to data and code being used within the TEE only to parties that are specifically allowed access. The logical processing limitations implemented at the hardware level can provide enhanced security and privacy protections for the processing of Personal Data, and in particular, sensitive Personal Data. The protections offered by TEEs can also be validated cryptographically in order to confirm that the code running within the TEE has not been altered.

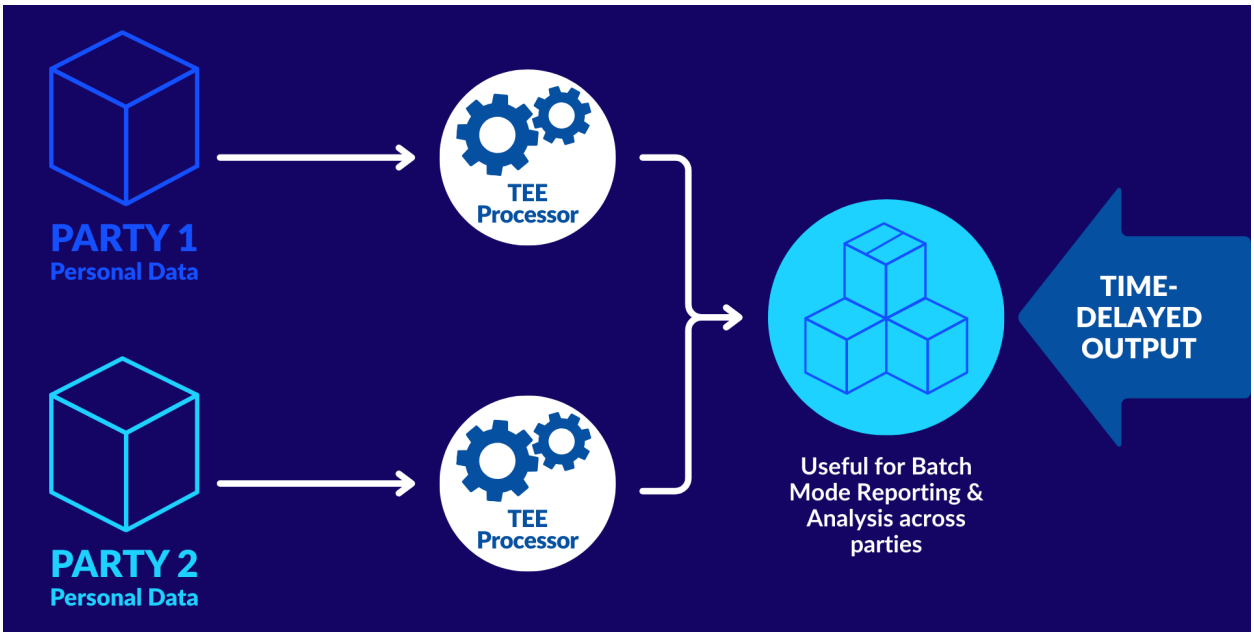

[7] *See, e.g.*, Nissenbaum, et al., "No Cookies For You!: Evaluating The Promises of Big Tech's 'Privacy-Enhancing' Techniques." (Dec. 9, 2023) at 24ff. https://kirstenmartin.net/wp-content/uploads/2023/12/Main-Article-FTC-No-Cookies-For-You-12-09-2023.pdf

Using a TEE can provide a higher level of trust in the validity, isolation, and access control afforded to the information stored and processed within it. TEEs can be used by themselves or in combination with other privacy-enhancing technologies. For example, a data controller can use Differential Privacy to reduce reidentification risk for a given dataset for the output of a TEE; this would limit the processing to be used only for specific types of operations, like attribution reporting (but not targeting). This would reduce risks for both reidentification and secondary uses.

Some of the same functional protections and guarantees offered by TEEs can be achieved through software alone, but would be implemented using logic at the software rather than hardware layer and hence would not share the same implementation characteristics as a TEE (*e.g.*, a separate CPU enclave to carry out the processing). As a result, software-based implementations would rely on software-enabled controls and contracts, instead of hardware-enabled controls and contracts, to govern the processing that will be applied to input data.

2. Key Concepts

TEEs are defined by three key features. The first is **data confidentiality**, meaning unauthorized entities cannot view or access data while it is in use within the TEE. As a TEE is a physically separate part of the CPU, called an "enclave,"[8] one of its core security properties is "isolation," meaning only the CPU can access the data and code inside the enclave. Another related security property is that everything stored and processed in the enclave is encrypted in runtime memory. The second key feature of a TEE is **data integrity**, which means that unauthorized entities cannot add, remove, or alter *data* shared with the TEE. The third key feature of a TEE is **code integrity**, which means unauthorized entities cannot add, remove, or alter *code* while used by the TEE.[9] In this context, unauthorized entities can mean other applications on the host server, the host's operating system, system administrators, service providers, or anyone with physical access to the hardware.

An additional feature of TEEs is the option of an **attestation** report, which validates that the server generating this report is the same server that received and processed the incoming data. In other words, the source of the attestation report being the TEE provider can be independently validated.[10]

The controller of a dataset can use a TEE to grant other parties limited access to facts about the dataset while maintaining contractual and technical safeguards to prevent unauthorized access or use of the underlying data. For example, a TEE can be used to enable a partner to query assets from a data controller, and the strict limitations that are set by the controller are ensured by technical safeguards. These safeguards ensure the efficacy of limits placed on data access and use, and the code can be audited and verified by trusted third parties. Inside this computing environment, the controller of the dataset does not have to trust the other party they are granting access to. The presence of these technical safeguards add additional safeguards beyond any contractual limitations on data use.

---

[8] Felix Schuster, Confidential Computing 101 by Felix Schuster (Edgeless Systems) | OC3 2021, YouTube (Mar. 19, 2021), https://www.youtube.com/watch?v=77U12Ss38Zc

[9] "A Technical Analysis of Confidential Computing," at 6, *The Confidential Computing Consortium*, Nov. 2022 https://confidentialcomputing.io/wp-content/uploads/sites/10/2023/03/CCC-A-Technical-Analysis-of-Confidential-Computing-v1.3_Updated_November_2022.pdf

[10] *Id.*

3. General Examples

i. Streaming Video and Digital Rights Management

Netflix wanted to attest to studio content owners that their videos would not be intercepted or copied during transmission to Netflix customers.[11] In other words, Netflix wanted to demonstrate to content owners that Netflix would access their data assets (videos) for only the limited purpose of real-time streaming to Netflix customers. Netflix relied on TEEs to meet this objective.

Netflix designed a TEE that stored decryption keys for the transmitted video content to ensure the license could not be used outside of Netflix's control. Netflix's code ensures only the authorized recipient (a Netflix customer) receives the encrypted data. Netflix's endpoint processing is the only authorized application to decrypt and render the studios' content. Because Netflix's distribution for this purpose is carried out in a TEE, Netflix can attest that this decryption-and-view use case is the only use case supported with the inputted data. This ensures both data and code integrity. As a result, owners of the data asset (the studios producing the video content) have technical assurance that their videos weren't being intercepted (*e.g.* pirated) while being streamed on Netflix.

ii. Signal's App Messaging Service and Confidential Attributes

Some companies rely on TEEs to append new attributes to existing data, but don't want the existing data to be used for other purposes as a result of the append process. TEEs can be used to enable an authorized party to append a new attribute to a protected dataset while ensuring that the append processing will not alter the protected data or allow it to be used in any way beyond the append, and that the controller of this processing code can attest that this is the only purpose that will be applied to the input data being shared.

For example, Signal is a secure messaging service that offers its users a TEE that identifies any existing Signal users in their address book.[12] A user uploads their address book to Signal's hardware processing environment using end-to-end encryption. This ensures only the authorized recipient (Signal) receives the data. Signal's processing compares the user's contact list with Signal's existing list of customers and appends a new attribute to the user's address book – that a given contact is also a Signal user. Because this processing code is executed in a TEE, Signal can attest that this append use case is the only use case supported with the inputted data and assure its users that their address book data is not being used by Signal for any other purpose.

4. Advertising Uses

TEEs can be used in ad tech to provide owners of an advertising-related dataset (*e.g.* an advertiser's marketing list) an additional layer of security for processing and certainty regarding data access and use controls set for recipients of that dataset. For example:

---

[11] Trusted Execution Environment (TEE) 101: A Primer, *Secure Technology Alliance* (Apr. 2018) at 18-19, https://www.securetechalliance.org/wp-content/uploads/TEE-101-White-Paper-FINAL2-April-2018.pdf

[12] Technology preview: Private contact discovery for Signal (Sep. 26, 2017) https://signal.org/blog/private-contact-discovery

- **Matching:** A TEE is useful for matching disparate datasets to create a targetable audience segment based on overlap between two companies' data.[13] For example, by using a TEE, both companies can be assured that only the overlapping records of the two datasets will be outputted from the TEE. Neither party has access to the non-overlapping data, only the TEE operator does.
- **Attribution Reporting:** A TEE can be helpful in generating an aggregated report, by preventing unauthorized access to raw conversion data. Within a TEE, analytics can be performed on encrypted conversion data. There are some Google Privacy Sandbox proposals that rely on TEEs for this purpose, including: Aggregated Attribution Reporting API (ARA) and Private Aggregation API, which enables the generation of aggregated and noisy reports; and Bidding and Auctions Services API, which aims to implement a secure bidding server.[14]

5. Additional Resources

- *Privacy-enhancing technologies (PETs)*, Information Commissioner's Office https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/data-sharing/privacy-enhancing-technologies
- Fabian Höring, et al., *PETs in Advertising: Scenarios for Trusted Execution Environments*, Criteo Tech Blog (May 25, 2023) https://techblog.criteo.com/pets-in-advertising-scenarios-for-trusted-execution-environments-9d0264c57325
- Joel Timothy, *What is a Trusted Execution Environment (TEE)?*, Duality Tech Blog (Aug. 8, 2022) https://dualitytech.com/blog/what-is-a-trusted-execution-environment-tee
- NextRoll Engineering Team, *TEEs: What They Are and Why They're Critical for Privacy Sandbox Testing*, The NextRoll Blog (May 29, 2024) https://www.nextroll.com/blog/product/trusted-execution-environments
- Carlos Cela et al., *Aggregation Service for the Attribution Reporting API*, Github https://github.com/WICG/attribution-reporting-api/blob/main/AGGREGATION_SERVICE_TEE.md
- Priyanka Chatterjee & Itay Sharfi, *Bidding and Auction Services*, Github https://github.com/privacysandbox/protected-auction-services-docs/blob/main/bidding_auction_services_api.md
- Phillip Lee & Peiwen Hu, *FLEDGE Key/Value service trust model*, Github https://github.com/privacysandbox/protected-auction-services-docs/blob/main/key_value_service_trust_model.md

[13] "Confidential matching", *Google Ads Data Manager Help* https://support.google.com/google-ads-data-manager/answer/14577185?hl=en

[14] Priyanka Chatterjee & Itay Sharfi, Bidding and Auction Services, GITHUB https://github.com/privacysandbox/protected-auction-services-docs/blob/main/bidding_auction_services_api.md; *see also* Carlos Cela et al., Aggregation Service for the Attribution Reporting API, GITHUB https://github.com/WICG/attribution-reporting-api/blob/main/AGGREGATION_SERVICE_TEE.md

B. Multiparty Computation

1. Overview

Multiparty Computation (MPC) refers to a technique for utilizing multiple processing entities to analyze subsets of data without revealing to counterparties the underlying plaintext information being processed. This can allow each participating data controller to reduce the risks of unauthorized access and use of their underlying data by other counterparties while still enabling joint analysis of their respective datasets.

MPC providers accomplish this by transforming and obscuring the subset of underlying information they receive from participating data controllers prior to sending their output to the other processing entities. As a result, each processing entity should remain unaware of the complete set of attribute values of underlying unaggregated information they are processing, as they receive only a subset of this unaggregated data. Given the data transformation involved in this process often relies on encryption, this system is sometimes referred to as "Secure Multiparty Computation."

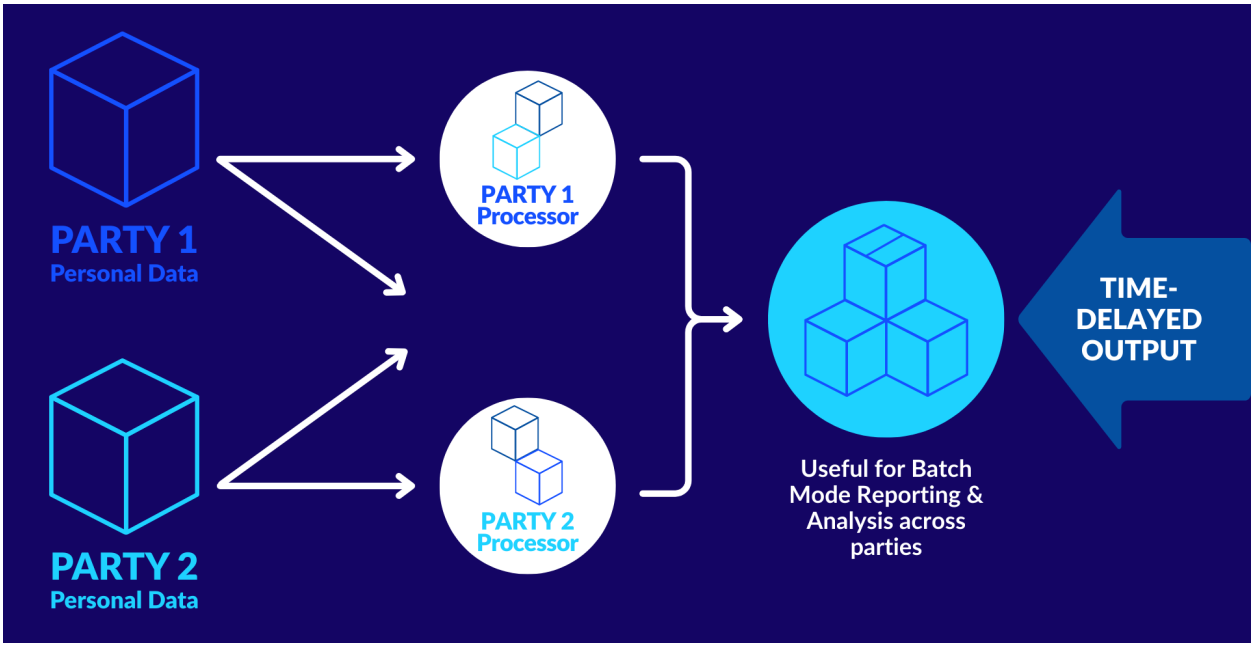

2. Key Concepts

A MPC system relies on three steps:

1. **Data Preparation**: The underlying information is both transformed into a new representation (e.g., encrypted) and split into multiple data sets.
2. **Data Computation**: These smaller encrypted subsets are processed without revealing the underlying information.
3. **Data Aggregation**: The outputs of the sub-processing are recombined to generate a final result.

Because each processing entity sees only a partial set of data, which has been transformed from its raw state, they can learn only a partial amount of information from their processing. The MPC controlling system is the only entity to see the complete set of the raw input and the final output.

3. General Example

a. Imagine three people want to compute their average salaries without sharing their own salary data with one another. If each person splits their own salary into 3 components and shares only these partial amounts with one another they can each get intermediate results that when averaged together yield the correct average without revealing any individual's total salary to any other recipient.

Alice's Salary = $100	Bob's Salary = $200	Carol's Salary = $300
Alice's Split for Alice = $50	Alice's Split for Bob = $15	Alice's Split for Carol = $35
Bob's Split for Alice = $100	Bob's Split for Bob = $40	Bob's Split for Carol = $60
Carol's Split for Alice = $120	Carol's Split for Bob = $150	Carol's Split for Carol = $30
Alice's Sum of splits = $270	Bob's Sum of splits = $205	Carol's Sum of splits = $125
Sum of All Splits = 600		
Average of All Splits = 200		

Note in the above example, each recipient knows a minimum value from the others. If sufficiently large noise is also added prior to processing, then this sharing of even partial information can be further reduced.

4. Advertising Uses

Multiparty Computation can be used in ad tech to provide insights from the combination of multiple datasets. For example, an advertiser may want to analyze the return on ad spend (ROAS) of a campaign based on offline sales. Through a multiparty-compute process, the analysis can be performed without revealing the retailer's or media owners' data to one another.[15]

At a high level, the process would entail:

1. Retailers send sales data for attribution
2. Media owner sends ad exposure data for attribution
3. Trusted MPC vendor uses common match key to compute the attribution

[15] "Privacy Preserving Attribution for Advertising" by Martin Thomson (Feb. 8, 2022) https://blog.mozilla.org/en/mozilla/privacy-preserving-attribution-for-advertising; Martin Thomson, Privacy Preserving Attribution for Advertising (Feb. 8, 2022) https://blog.mozilla.org/en/mozilla/privacy-preserving-attribution-for-advertising

5. Related Techniques

FEDERATED LEARNING: Federated Learning is similar to MPC as both require all computing parties to share the same processing model. Federated Learning differs from MPC by sharing model parameters trained on locally processed subsets of data rather than sharing subsets of unaggregated data among multiple processing entities.

HOMOMORPHIC ENCRYPTION: An approach that transforms raw data prior to computation - is similar to MPC in transforming the input data to be processed, but does not require multiple processing entities nor the addition of noise. The key difference with homomorphic encryption is that the input data is encrypted in a particular way that still allows for mathematical computations and analytics (e.g., finding an average across disparate records) while not exposing the underlying raw values. It is important to note however, that if any recipient can see both the input and output values, then it is possible to reverse the encryption methodology. Additionally, with only the encrypted values, though a recipient will not see raw values, it may be possible for the recipient to understand some information about the data, such as which source is providing a higher input value, and the relative magnitude of the difference.

To make the example above representative of Homomorphic Encryption, the actual salary values would be replaced with encoded data as illustrated in the table below.

Alice = $100	**Bob = $200**	**Carol = $300**
Alice's encrypted bid = A	Bob's encrypted bid = AA	Carol's encrypted bid = AAA
Maximum Bid = AAA		
Average of All Bids = AA		

A data recipient who sees the encrypted outputs would not know the underlying values, but would be able to determine that Carol's value is the largest of the set, and perhaps even determine its magnitude relative to Alice (=3 times more). It is also important to note that if any recipient sees both the raw input and encrypted output values, it is possible to reverse the encryption protocol.

6. Additional Resources

- *Privacy-enhancing technologies (PETs)*, Information Commissioner's Office https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/data-sharing/privacy-enhancing-technologies
- *Highlights of KDD 2022* https://techblog.criteo.com/highlights-of-kdd-2022-6950d10e0248

C. Differential Privacy

1. Overview

Differential Privacy allows the controller of a dataset containing Personal Data to share aggregate information with another party without revealing any specific underlying data elements in the set. Differential Privacy works by introducing "noise" into datasets – essentially, random data points that do not reflect any true features of the individuals in the dataset – to mitigate privacy risks to specific individuals. Added noise reduces the risk that the identity of individuals within the dataset can be discovered while still allowing for statistically useful information to be drawn from the dataset.

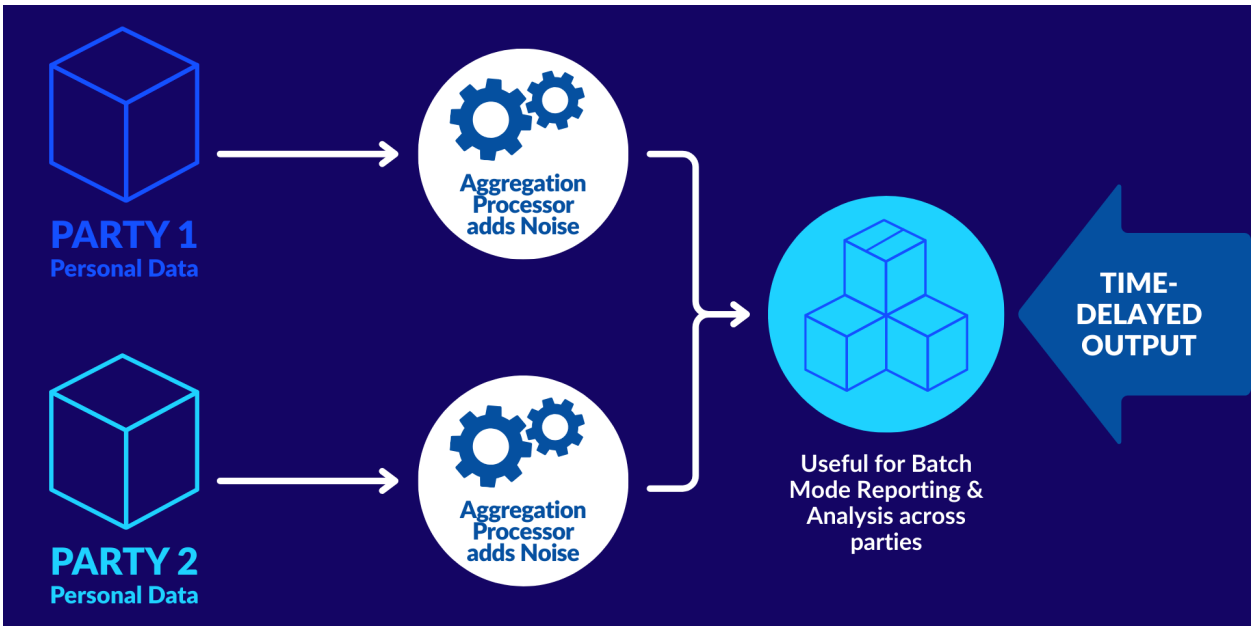

2. Key Concepts

There are two key concepts in Differential Privacy: the **privacy loss budget** (represented by Epsilon, or ε), and the **privacy unit** (*e.g.*, an individual data subject). The privacy unit determines what is being protected, and the privacy loss budget determines how well a privacy unit is protected. The ε represents the maximum tolerance for revealing information through the dataset's outputs. When the controller of a differentially private dataset specifies the value of ε – in other words, establishes the privacy loss budget – it affects the risk that an individual in the dataset could be reidentified, but also affects the utility of the data. Balancing the privacy budget for a dataset with its utility is challenging because lowering the privacy budget also decreases the accuracy of the input information available for processing.

For datasets containing Personal Data, each individual's data is considered a "privacy unit." In other words, the privacy unit refers to the collection of specific data about each individual and it is the identity of each individual that this process aims to protect. In this context, the privacy loss budget (ε) represents the level of protection Differential Privacy provides for each individual's identity within the dataset. The lower the ε value, the greater the level of protection. For example,

a comparatively low ε value for a differentially private dataset means that it would be more difficult to reidentify individuals in that dataset.[16]

However, because Differential Privacy protects information in a dataset by adding noise to it, achieving a higher level of protection involves adding more noise, and hence, decreasing the accuracy of aggregate information drawn from that dataset. In other words, Differential Privacy always involves a trade-off between the risk of reidentification of individuals in the dataset and the accuracy (and hence, utility) of the aggregate information drawn from the dataset.

Delta (δ) is another metric typically set to 1/n, where "n" is the total unique identifying match keys in the input data set. If there are 1,000,000 records in the input data set, and δ=1/100,000, then the disclosure of a single record has only a 1 in 10 chance of being associated with the unaggregated output information.

3. General Examples

a. United States Census Bureau

The United States Census Bureau began to use Differential Privacy in the 2020 United States Census to allow researchers to continue to use census data while protecting the privacy and mitigating the risk of reidentifying any individual census respondents. Adding noise into the census dataset reduces the risk that an outside party can correctly reidentify any individual census respondent, while still preserving the utility of the data for research purposes.[17]

b. Pandemic Community Mobility Reports

At the height of the Covid-19 pandemic, Google looked to analyze how many people went to their workplace or to another specific kind of public location per day, and how long people spent at home. In order to protect the privacy of the individuals whose location data was analyzed, Google employed Differential Privacy techniques by injecting noise into the relevant location datasets in order to compensate for specific instances where the "privacy threshold" (the ε) was not met for certain locations. This enabled Google to analyze useful metrics from the datasets and to maintain statistical reliability while still preserving the privacy of individual consumers on days when fewer individuals visited specific locations. They first generated a set of anonymized metrics from the data of consumer users of Google's tools who opted into Location History. Then they computed percentage changes of these metrics from a baseline based on the historical part of the anonymized metrics. Then they discarded a subset of the data that did not meet Google's bar for statistical reliability, and released the rest publicly in a format that compares the result to the private baseline.[18]

---

[16] For example, when Apple applies Differential Privacy to sensitive health data in some scenarios it sets ε to 2, which represents a relatively high level of protection. Apple, Differential Privacy Overview, § 2 https://www.apple.com/privacy/docs/Differential_Privacy_Overview.pdf

[17] *Understanding Differential Privacy*, U.S. Census Bureau https://www.census.gov/programs-surveys/decennial-census/decade/2020/planning-management/process/disclosure-avoidance/differential-privacy.html

[18] "Community Mobility Reports", Google (2022) https://www.google.com/covid19/mobility

c. Uber Average Trip Data

Uber uses a Differential Privacy method based on elastic sensitivity to protect its drivers and riders. As Uber data scientists query their database to perform analyses, the system limits the amount of Personal Data revealed to ensure every individual's anonymity. This limitation is the ε. Uber can analyze traffic patterns and even calculate revenue from raw data without distinguishing any single consumer's data.

Differential Privacy can provide high accuracy results for the class of queries Uber commonly uses to identify statistical trends. It allows Uber to calculate aggregations (averages, sums, counts, etc.) of elements like groups of consumer users or trips on the platform without exposing information that could be used to infer details about a specific consumer user or trip. Uber adjusts the statistical noise depending on certain factors; larger cities have more trips per day, and removing any one individual trip from the dataset does not change the average trip distance. However, for smaller cities with fewer trips per day, more noise is injected to afford consumers the same degree of privacy.[19] These adjustments prevent individual consumer users from being reidentified, but still allow for Uber to draw conclusions about the aggregate results.

4. Advertising Uses

Within ad tech, Differential Privacy can be useful for measurement and attribution purposes, as well as for modeling certain audiences. Because Differential Privacy is primarily a tool for protecting against reidentification while enabling analysis of aggregate data, its utility is more limited for use cases that require processing or sharing consumer-level data (e.g., auditing specific impressions).[20]

- **ANALYTICS**: Differential Privacy can be applied to datasets containing ad measurement and attribution data to protect the privacy of individual consumers to whom the information in the dataset may relate, all while relaying accurate aggregate information in reports. For example, Differential Privacy can be used to measure conversion rates and opt-out rates, providing valuable insights for advertisers while protecting the privacy of individual consumers. This can be accomplished by adding noise to the dataset containing conversion event data in a way that preserves the accuracy of the conversion rate (and hence, the efficacy of the campaign) but that reduces the risk that an individual in that dataset could be reidentified.[21]

---

[19] Katie Tezapsidis, Uber Releases Open Source Project for Differential Privacy, Medium (Jul. 13, 2017) https://medium.com/uber-security-privacy/differential-privacy-open-source-7892c82c42b6

[20] IAB Tech Lab Privacy Sandbox Task Force, Privacy Sandbox Fit Gap Analysis, IAB TECH LAB (Jun. 2024) at 15, https://iabtechlab.com/wp-content/uploads/2024/06/Privacy-Sandbox-Fit-Gap-Analysis-FINAL.pdf ("Third-party audits are crucial for verifying digital advertising transactions' security, performance, and accuracy today. They objectively assess that an advertising transaction is fraud-free, properly targeted, and meets vital measurement standards.")

[21] In 2024, Zenjob, a job placement platform, wanted to use Differential Privacy to measure the effectiveness of its TikTok campaign. "Anonym matched hashed and encrypted sales data with hashed and encrypted impression data from TikTok. The data was processed using differentially private algorithms for lift and attribution. Differential privacy is a method that adds noise to data sets, making individual data points indistinguishable." https://blog.mozilla.org/en/advertising/anonym-zenjob

- **AUDIENCE AND MODELING**: Differential Privacy can be used for analyzing aggregated consumer activity data to return private query results. Models and lookalike audiences can be built from this analysis, allowing an advertiser to deliver relevant ads while protecting the privacy of individual consumers within an underlying dataset. For example, for sensitive datasets, such as those related to health, an algorithm can analyze the demographics associated with a particular health condition. (This is in line with the guidance the NAI developed for demographic health advertising.[22]) By including noise in the datasets, the downstream algorithms will not be able to identify any specific individuals, but still allow the advertiser to pull out pertinent demographic information to build an audience for an advertisement for a specific treatment or for a clinical trial.[23] By using differentially private techniques to observe certain trends or correlations, ad tech companies can build audiences that match the demographic profiles of these model audiences without actually collecting potentially sensitive information about the audiences they are targeting.[24]

5. Additional Considerations

Setting an appropriate ε can help the controller of a dataset meet legal standards for deidentification if the ε is sufficiently low that the aggregate outputs of differentially private datasets cannot "reasonably" be linked to or used to infer information about any particular individual in the dataset.[25] Data that meets applicable legal standards for deidentification are generally subject to fewer compliance obligations.

6. Additional Resources

- *IAB Tech Lab Differential Privacy Guidance for Digital Advertising*, IAB Tech Lab (2023) https://iabtechlab.com/wp-content/uploads/2023/11/Differential-Privacy-Guidance_PUBLIC-COMMENT_11152023.pdf
- *A Marketer's Guide to Privacy-Enhancing Technologies*, Deloitte Digital/Meta https://www2.deloitte.com/content/dam/Deloitte/us/Documents/us-Meta-PETs-Whitepaper.pdf
- *Differential Privacy Overview*, Apple https://www.apple.com/privacy/docs/Differential_Privacy_Overview.pdf

---

[22] The Network Advertising Initiative, 2023 Demographic Health Advertising Best Practices, *The Network Advertising Initiative* (2023) https://thenai.org/wp-content/uploads/2023/11/NAI-Health-Targeting-Best-Practices-Document-Final.pdf

[23] DeepIntent, Differential Privacy Introduction, DeepIntent (2024) https://www.deepintent.com/differential-privacy-introduction ("DeepIntent uses Differential Privacy to enable advertisers to target protected health segments and extract critical insights and generalized learnings from datasets without linking that information to a specific individual from the dataset.")

[24] Ryan Rogers et al., LinkedIn's Audience Engagements API: A Privacy Preserving Data Analytics System at Scale, 11 J. Privacy & Confidentiality 3 (2021) at 7 https://journalprivacyconfidentiality.org/index.php/jpc/article/view/782/724 ("LinkedIn offers a Differentially Private audience querying solution to enable analysts to understand the attributes associated within a given audience segment. When this solution limits the audience attributes to 3,000 distinct combinations of attributes, they found 93% of analysts' queries would not be impacted by this restriction.")

[25] See, e.g., Cal. Civ. Code 1798.140(m) (defining deidentified data)

- NIST, Guidelines for Evaluating Differential Privacy Guarantees https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-226.ipd.pdf
- *Disclosure Avoidance for the 2020 Census: An Introduction*, United States Census Bureau (Nov. 2021) https://www2.census.gov/library/publications/decennial/2020/2020-census-disclosure-avoidance-handbook.pdf
- *Revealing Wikipedia usage data while protecting privacy*, Tumult Labs https://www.tmlt.io/casestudy/revealing-wikipedia-usage-data-while-protecting-privacy

D. Zero-Knowledge Proof

1. Overview

Zero-Knowledge Proofs enable one party (a Prover) to convey to another party (a Verifier) some truth about a dataset without revealing to the Verifier the actual underlying information in the dataset. By involving a trusted third party in the verification process, Zero-Knowledge Proofs ensure that the truth about the dataset can be confirmed without the Verifier—or any other party—gaining additional access to the underlying information. This can allow the Prover, as a data controller, to reduce the risks of unauthorized disclosure of underlying data to a counterparty while still enabling that counterparty to confirm a truth about the dataset, and without the Verifier gaining any knowledge of the underlying data (hence, the moniker "zero-knowledge" proof).

2. Key Concepts

Zero-Knowledge Proofs are defined as those proofs that convey no additional knowledge other than the validity of the proposition being proved. However, Zero-Knowledge Proofs do not guarantee that the underlying facts are sound or accurate.

A system must satisfy three conditions to be considered a Zero-Knowledge Proof:[26]

1) **Completeness**: If the statement is true, an honest prover can always convince an honest verifier. For example, a digital keypad lock that displays a green light for the correct code and a red light for an incorrect entry would meet this criteria. This is because if an honest prover correctly enters the code without revealing it to the Verifier (e.g., only a green light shows for correct entries), this would ensure the prover can always demonstrate knowledge of this fact to the verifier.
2) **Soundness**: If the statement is false, no dishonest prover can convince an honest verifier, except with very low probability. Keeping with the digital keypad lock example, suppose that the unlock code has enough digits that it is extremely unlikely to guess on the first try. If a dishonest prover enters the wrong code (causing the red light to show for incorrect entries), the Verifier can confirm the Prover does not know the correct code.
3) **Zero-Knowledge**: The verifier learns nothing about the underlying information beyond the fact being proved. In our examples above, nothing about the code itself is shared (the

[26] Zero-knowledge proofs were first described in a 1985 MIT paper from Shafi Goldwasser and Silvio Micali called "The Knowledge Complexity of Interactive Proof-Systems." https://scispace.com/pdf/the-knowledge-complexity-of-interactive-proof-systems-31apre0ecf.pdf

number of digits, the specific digits, or the pattern of digits), but the resulting color feedback demonstrates whether or not the Prover knows the code to unlock the keypad.

3. General Examples

Zero-Knowledge Proofs can prove membership claims without revealing underlying information (e.g., age verification) or knowledge claims without revealing information contents (e.g., password).

The Zero-Knowledge Proof system diagramed below is programmed to evaluate data from the Prover and validate some condition to the Verifier. The condition would be determined by the goals of the system. For example, if a company wants to monetize a segment of its customers based on how much they spend with the company, but they do not want to reveal underlying sales data to a demand-side platform (DSP) activating the segment, the company (the Prover) could deploy a Zero-Knowledge Proof platform to evaluate for the DSP recipient (a Verifier) if a given customer qualifies for the campaign based on whether their sales data is higher than a particular established threshold (e.g., annual spending > $100). The Zero-Knowledge Proof platform can provide the Truth state (Yes or No) to the DSP without revealing the actual value of the total customer sales. It is also a business decision whether the Verifier knows the particular threshold that determines if the result is a True.

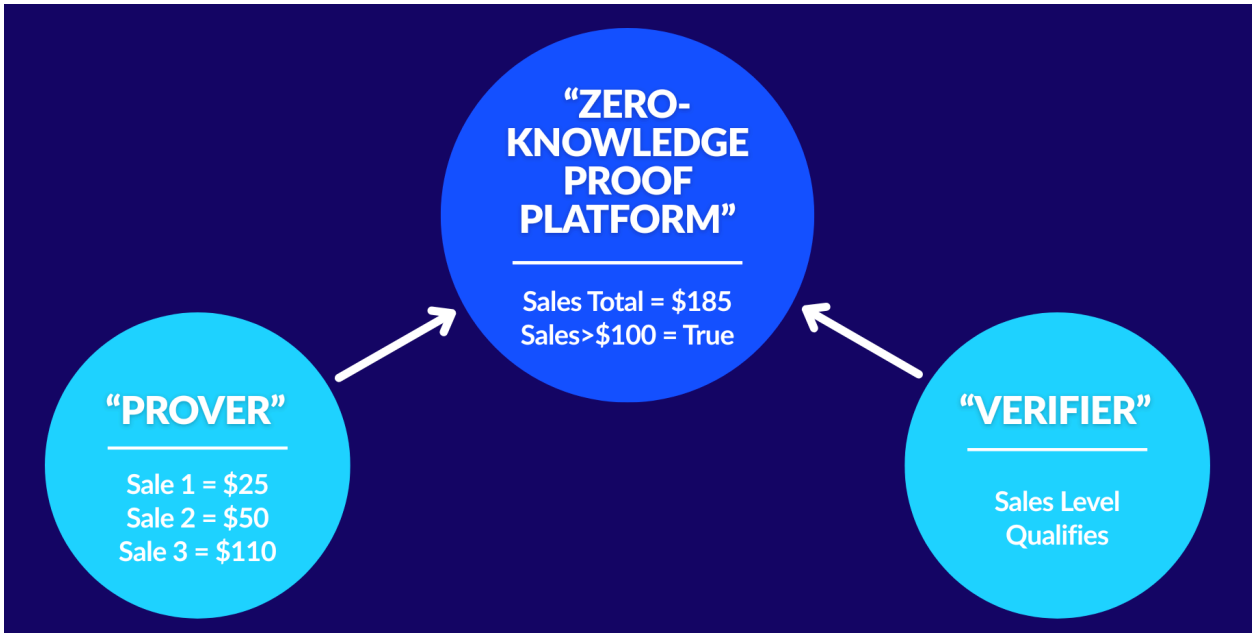

a. Ball Color Test[27]

In this example, suppose Alice is a color-blind verifier possessing a red and green ball, which are otherwise indistinguishable balls. Bob wants to convince Alice he can detect the color of each of these other balls without revealing to Alice which is which.

[27] Oded Goldreich, Silvio Micali & Avi Wigderson, Proofs That Yield Nothing But their Validity or All Languages in NP Have Zero-Knowledge Proofs, 38 J. ACM (1991) at 691- 729 https://dl.acm.org/doi/pdf/10.1145/116825.116852 (The Ball Color Test was first presented by Goldreich, Micali and Wigderson.)

Bob asks a third-party – Charles – to hold both balls behind his back. Charles repeatedly displays them one at a time to Bob, asking him whether he switched the balls each time. Since Charles knows the truth of the color (information), as Bob's replies approach zero error the likelihood Bob can distinguish the balls increases (soundness). If this system is repeated multiple times with the same negligible error, Bob should be able to convince Alice that he knows the truth of the color (completeness). Charles' system can convince Alice that Bob knows this information, but never reveals which ball is which color to Alice (zero knowledge).

b. Age Verification

A visitor registers with a proving authority (e.g., government) providing Personal Data, such as her birthdate. The visitor wants to prove to a digital property that she is at least 18 years old by referring to this trusted authority. By submitting some other authenticating information, the digital property can query the authority (verifier) to receive the answer to whether she is at least 18 years old without learning the actual birthdate. [28]

c. Identity Verification Using Passcodes

An e-commerce property wants to validate whether a visitor to its digital property is the registered account holder but is concerned about another entity from learning or masquerading as the visitor. Imagine the e-commerce website has established over the phone a unique passcode with the registered account holder. The e-commerce property can ask the visitor to enter only the last four characters of the passcode into local software running on the visitor's device which generates an encrypted output using the current time as a salt, the output of which is transmitted back. The visitor's software asks the e-commerce property to use the last four characters of the passcode to encrypt a different set of four characters of the passcode using this same method. The visitor's software then sends the same set of four characters of the passcode selected by the e-commerce property encrypted with the new salt. The e-commerce property can validate if the visitor knows the passcode without asking for the full passcode or revealing it in transit.

4. Advertising Uses

Zero-Knowledge Proof processes can be used in ad tech to provide facts about audiences without revealing the underlying data. For example, the controller of consumer data can provide knowledge of which specific consumers are subject to age restrictions without revealing the age or birthdate of each individual. Another example could be enabling an output recipient to query whether a given consumer has purchased more than a certain amount of a product without the actual sales transactional data or even aggregate purchase amounts being revealed, say for fraud prevention purposes.[29]

---

[28] Google Blog (Apr. 29, 2025) https://blog.google/products/google-pay/google-wallet-age-identity-verifications ("Given many sites and services require age verification, we wanted to develop a system that not only verifies age [without revealing an individual's birthday]. That's why we are integrating Zero Knowledge Proof (ZKP) technology into Google Wallet, further ensuring there is no way to link the age back to your identity.")

[29] What is Zero Trust Zero-Knowledge Proof?, SharedID (Dec. 9, 2024) https://www.shareid.ai/blog/what-is-zero-trust-zero-knowledge-proof ("SharedID relies on Zero Knowledge Proofs to enable the website (verifier) to verify the identity of a user or device without revealing any personal information.")

5. Additional Resources

- Luís Brandão, René Peralta, & Angela Robinson, *NIST comments on the initial ZKProof documentation*, NIST (Apr. 6, 2019) https://csrc.nist.gov/CSRC/media/Projects/pec/documents/20190406-nist-pec-comments-on-zkproof-docs.pdf.
- *Privacy-enhancing technologies (PETs)*, Information Commissioner's Office https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/data-sharing/privacy-enhancing-technologies

III. Appendix

A. Glossary: Terminology required to understand Key Concepts

Understanding the different PETs addressed in this primer requires understanding certain key terms. The definitions of key terms offered here are specific to this document and do not necessarily equate to specific legal or regulatory definitions, which can differ by jurisdiction.

- **Aggregated data**: a summary of record-level data. Aggregated information can be used to show trends or statistical values without identifying individuals or specific record-level transactions within the data. When event-level data from multiple individuals is aggregated and not linked to individuals, it is no longer Personal Data. However, merely aggregating data, while retaining its association to an individual, remains Personal Data (e.g., summing the sales from a single individual).
- **Anonymous Non Personal Information**: information that never was Personal Data or has been processed in such a way that in the hands of a recipient it can no longer be legally attributed to a specific individual. Data protection law does not apply to anonymous information.
- **Attribute Value**: information about a match key.
- **Data**: the combination of match keys and attribute values. Despite the computer software distinction of these two components, many jurisdictions use the terms “data” and “information” synonymously.
- **Data Processing**: This function ingests input data (raw ingredients) and transforms them into Output Data. Sometimes, the output is left with event-level details (unaggregated), and other times it’s combined or grouped together (aggregated) to create a statistic.
- **Deidentified Non Personal Information**: information that may have been Personal Data but has been processed in such a way and has appropriate organizational measures in place such that in the hands of a recipient it can no longer be legally attributed to a specific individual. Data protection law does not apply to non personal information.
- **Directly Identifiable Personal Data**: any data item that, on its own, could uniquely identify a specific individual (or household in certain jurisdictions). Data protection law does apply to all Personal Data. Under privacy laws in the U.S., Personal Data is generally considered to be information that is linked or reasonably linkable to an identified or identifiable individual.
- **Hashing**: a process using a one-way mathematical function that transforms input data into a fixed-length output that does not reveal the original plaintext.
- **Inferences**: the probabilistic attributes that guess or predict details about a match key.
- **Information Value**: The **Value** is the actual information or meaning connected to that Key. For example, if the Key is "age range," the Value could be "25-34." When a Key/Value pair directly identifies a specific person—like "email = john.smith@gmail.com"—then the data is considered **Personal Data**. A Key/Value pair without direct identifiers is not per se “Personal Data” unless the information (e.g., age=25) is also linked to a specific individual, either **directly-identifiable** (e.g. email address; phone number) or a **pseudonymous** identifier (e.g., Hashed Email).
- **Masking:** replacing a direct identifier with a new value while preserving the attributes linked to this identifier. Common examples include replacing names with pseudonyms or masking credit card numbers.
- **Match Key**: a label or direct identifier that, on its own, could uniquely identify a specific object or concept. Match keys are used to connect information across time or systems.

- **Match Key Name:** The **Key Name** is a human-readable label designed to uniquely identify a distinct real-world object (like a person or thing) or describe a concept (like "presence of children" or "age range").
- **Noise**: introducing random numerical data points that do not reflect any true features of the individuals in the dataset.
- **Plaintext**: in cryptography, plaintext refers to information that is not encrypted and is therefore readable.
- **Processing Entity**: A processing entity is a logical group or function within an organization that processes data. Applying PETs may involve a number of processing entities, both within a given organization and across multiple organizations, including vendors that employ PETs. For example, a “clean room” provider is one example of a separate organization that can employ PETs to protect confidentiality among two organizations wishing to share only subsets of their respective pools of information with one another.
- **Pseudonymous Personal Information**: information that remains Personal Data but has been processed in such a way, such as by masking, that in the hands of a recipient it can no longer be attributed, without more information, to a specific individual. Data protection law does apply to all Personal Data.
- **Reidentification**: re-linking the directly identifiable information of a specific individual to deidentified data.
- **Targeting**: a media buyer’s focus on their spend on particular tactics associated with Audience, Context, Device or Geographic enrichment information with an aim to improve the effective return on their investment.
- **Unaggregated data**: record-level information that includes a match key, associated information values, as well as often a timestamp when it is event-level data. An example of unaggregated data could be when a business has ten rows of sales data with each row containing the amount of each sale (e.g., Sale1=$100, Sale2=$150, Sale3=$540, etc.)