November 2026



IAB EUROPE'S GUIDE TO PRIVACY ENHANCING TECHNOLOGIES (PETs) IN DIGITAL ADVERTISING

iabeurope.eu

Introduction

Privacy-Enhancing Technologies (PETs) are a broad suite of engineering techniques that safeguard and enhance privacy and security by minimising the collection, use, retention and exposure of data, while still enabling the insights that power products and services. In digital advertising, PETs offer opportunities to further strengthen existing privacy and security safeguards, while preserving data utility and supporting continued innovation.

IAB Europe's *Guide to PETs in Digital Advertising* builds on existing work, notably FEDMA's and the OECD's to provide a practical taxonomy grounded in how PETs function in digital advertising. It is designed to give regulators and industry a clear reference point on what these technologies are, how they are recognised, how they can be audited and how they are already used.

The Guide begins with definitions of the key PET categories and an overview of how PETs are currently recognised by regulators and institutions. It then explores Auditability, both of System Architecture and of Processing, before presenting practical use cases in Measurement and Audience Targeting to illustrate how PETs work in real-life examples. The Guide also addresses the most common misconceptions about PETs, and concludes with recommendations for their greater recognition in policy-making.

The Guide is written with regulators and policymakers in mind, and the final section sets out concrete suggestions for how PETs could be better recognised by regulators when assessing compliance under the current EU regulatory framework, and for how their use could be further incentivised through the Digital Omnibus proposal. It is also relevant to the digital advertising industry, which the Guide invites to converge on a layered approach to Auditability.

Table of Contents

Definitions **4**

Current Regulatory Recognition **7**

Auditability **9**

Practical Use Cases **12**

Measurement 12

Audience Targeting 17

Addressing Common Misconceptions **23**

Recognising PETs in Policy-Making **26**

Definitions

[Based on recent FEDMA[1] and OECD[2] papers]

Privacy-Enhancing Technologies (PETs) is an umbrella term to designate multiple tools, technologies and techniques used to improve security and maintain user privacy through an additional layer of protection, for instance, by minimising the amount of data processed by third parties. [...] PETs are a broad suite of engineering techniques that can safeguard and enhance privacy and security by minimising the collection, use, retention, and exposure of data with technical assurance/verification while enabling insights from data that power products and services.

We can group PET technologies into two broad themes:

- Those that add isolation protection, and
- Those that pseudonymise or anonymise data to make it safer for processing.

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

Overview of Key PET Categories, Including (But Not Limited To):

Differential Privacy[3]

Differential Privacy makes small changes (sometimes referred to as 'adding noise') to the raw data to mask the details of individual inputs while maintaining the explanatory power of the data. The idea is that small changes to individual records can securely de-identify the inputs without having a significant impact on the aggregated results. Noise can be added at the time of data collection (distributed) or at the central location before the data are released (centralised).

Federated Learning[4] (FL)

FL is a privacy-enhancing technology that enables machine learning (ML) models to be trained without the need for centralised data collection. As set out by McMahan et al. in the foundational paper that introduced the approach, rather than uploading raw data to a central server, the training data remains on participating devices (known as clients). Each client uses its local data to compute an update to the current shared model, and only that update is sent to a central server, which aggregates the contributions into an improved global model. This applies the principle of

1. The marketers' perspective on unlocking Privacy Enhancing Technologies (PETs) in personalised marketing and advertising communications, FEDMA, 2026
2. Emerging Privacy Enhancing Technologies, OECD digital economy papers, March 2023 No. 351
3. GDMA, Global Data Privacy, What the Consumer Really Thinks, Foresight Factory, 2022
4. Communication-Efficient Learning of Deep Networks from Decentralized Data, 26 Jan 2023

data minimisation, since the updates are specific to improving the model, are not needed once applied, and contain no more information than the underlying data. Some trust in the coordinating server is still required, and the strength of the privacy benefit depends on what the updates reveal, which is why FL is often combined with other PETs, such as differential privacy and secure multi-party computation.

Homomorphic Encryption (HE)[5]

Standard data processing methods require data to be decrypted before an organisation can use it, because, as Acar et al. explain, conventional encryption cannot be worked on without first being unlocked. HE removes this constraint by allowing a third party, such as a cloud or service provider, to compute on encrypted data without ever seeing the underlying content. The data subject locks the data (with a key only they hold) before passing it on to the data processor. The processor performs the requested calculation directly on the encrypted data, without needing the key, and returns an encrypted result that only the data subject can unlock. The range of calculations depends on the type of scheme: partially homomorphic schemes support one type of operation, somewhat homomorphic schemes a limited number of operations, and fully homomorphic encryption (FHE), first shown to be feasible by Craig Gentry in 2009, an unlimited number of operations, and therefore any computable function. The authors note, however, that the computational cost of FHE remains a barrier to practical use in many settings.

Pseudonymisation[6]

Pseudonymisation involves removing potentially identifiable information from the data to reduce the risk of identification of the data subject, although some residual risk may remain. Pseudonymised data may preserve their potential to be reconstructed when combined with additional information that is kept separately. Most recently, in Europe, the CJEU has reasserted the relative nature of personal data (see EDPS v SRB C-413/23P). The case confirmed that pseudonymised data held by a recipient without reasonable means to re-identify the dataset (i.e., the key remains with the controller) should not be considered personal data and therefore falls outside of the scope of the GDPR. It is expected that further clarification on pseudonymisation and anonymisation will arise from regulatory bodies in the near future. The relative approach to personal data has already been reflected in the Digital Omnibus proposal through the amendment to the definition of personal data, as well as a mandate for an implementing act to specify criteria for determining when data resulting from pseudonymisation is no longer considered personal data.[7]

5. A Survey on Homomorphic Encryption Schemes: Theory and Implementation, July 2018
6. OECD, Emerging Privacy Enhancing Technologies: Current Regulatory and Policy Approaches, OECD Digital Economy Papers No. 351 (March 2023)
7. Proposal for a regulation of the European Parliament and of the Council amending Regulations [...] and Directives 2002/58/EC, (EU) 2022/2555 and (EU) 2022/2557 as regards the simplification of the digital legislative framework, and repealing Regulations (EU) 2018/1807, (EU) 2019/1150, (EU) 2022/868, and Directive (EU) 2019/1024 (Digital Omnibus), Arts. 4 and 41a.

Anonymisation[8]

Anonymisation is the process of removing identifying elements from data to prevent re-identification of the data subject, taking into account all means reasonably likely to be used for identification. Anonymised data can no longer be linked back to an individual even when combined with additional data sets. Anonymisation has been used widely to enable the use of data for analytical or statistical purposes. Once data is truly anonymous and individuals are no longer identifiable, the data will not fall within the scope of the GDPR.

Secure Multi-Party Computation (MPC)[9]

MPC is a set of tools that enables the participating parties to jointly compute a function over their input data while keeping that input data private. Essentially, it removes the need for a trusted third party to view and manage the data. MPC can aggregate sensitive data without requiring any data contributor to disclose their own data. As a result, secret sharing techniques, or Homomorphic encryption (HE), can be used to aggregate and compute data from multiple parties.

Trusted Execution Environments (TEEs)[10]

A TEE is a dedicated area on a computer processor that is separated and secured from the operating system. It can process sensitive, immutable data and can run secure code within its secure confines. TEE assumes the operating system is corruptible and untrustworthy. Consequently, under TEE, the operating system cannot access information in the secure area of the processor or read the stored secrets. TEEs provide an isolated processing environment that adds a hardware-backed layer of security for storing and utilising data.

8. ibidem
9. OECD, Emerging Privacy Enhancing Technologies: Current Regulatory and Policy Approaches, OECD Digital Economy Papers No. 351 (March 2023)
10. ibidem

Current Regulatory Recognition

As the following overview demonstrates, PETs are widely accepted as essential "building blocks" for data protection by design and security in areas such as digital identity and data spaces. Regulators have started to acknowledge PETs in these domains, but advertising-specific recognition has been slower to materialise. This regulatory delay often leaves a vacuum where technical solutions are available but lack the explicit endorsement needed to be fully integrated into the AdTech landscape.

Recognition as Privacy-by-Design

- EDPB, in its guidelines, has recognised that PETs can "be employed as a measure in accordance with the DPbDD [data protection by design and by default] requirements if appropriate in a risk-based approach". However, they also noted that "PETs in themselves do not necessarily cover the obligations of Article 25. Controllers shall assess whether the measure is appropriate and effective in implementing the data protection principles and the rights of data subjects". [11]
- ENISA indicates that PETs are designed to support the implementation of data protection principles, in particular often connected with data protection by design and default. "PETs, as technical solutions, can be perceived as building blocks towards meeting data protection principles and the obligations under GDPR Art. 25 on data protection by design".[12]
- AEPD also perceives PETs as tools supporting privacy by design.[13]
- OECD also recognises PETs as supporting the privacy-by-design and -by-default approach[14]. It describes PETs as solutions that allow information to be collected, processed, analysed, and shared while protecting data confidentiality and privacy.
- ICO guidelines indicate that PETs can support data protection by design and by default but also other data protection principles, such as data minimisation[15].

11. EDPB, Guidelines 4/2019 on Article 25 Data Protection by Design and by Default Version 2.0 Adopted on 20 October 2020.
12. ENISA, DATA PROTECTION ENGINEERING From Theory to Practice, JANUARY 2022
13. AEPD, A Guide to Privacy by Design. A Guide to Privacy by Design
14. EMERGING PRIVACY-ENHANCING TECHNOLOGIES, CURRENT REGULATORY AND POLICY APPROACHES OECD DIGITAL ECONOMY PAPERS, March 2023, No. 351
15. How can PETs help with data protection compliance? (ICO)

Recognition as Robust Security Techniques

- The CNIL has a statutory duty to promote the use of PETs, in particular data encryption technologies, by virtue of the Law No. 2016-1321 of October 7, 2016, for a Digital Republic[16]
- PETs such as Trusted Execution Environments, anonymisation or federated learning are being deployed as a part of the security measures for the new EU initiative in particular for the digital identity wallets[17] and Common European Data Spaces[18].
- EDPB indicated that PETs can be used as a security measure for blockchain technology[19].
- Both researchers[20] and ENISA[21] indicate that some PETs are already mature and safe enough to be effective to use, but they also need to be combined with other privacy measures to provide strong protection. The most widely used and mature PETs identified are Differential Privacy, Trusted Execution Environments, Federated Learning, Secure Multi-Party Computation and Anonymisation.

16. LOI n° 2016-1321 du 7 octobre 2016 pour une République numérique (1), LOI n° 2016-1321 du 7 octobre 2016 pour une République numérique (1) - Légifrance
17. TechDispatch #3/2025 - Digital Identity Wallets | European Data Protection Supervisor
18. Auñón JM, Hurtado-Ramírez D, Porras-Díaz L, Irigoyen-Peña B, Rahmian S, Al-Khazraji Y, Soler-Garrido J, Kotsev A. Evaluation and utilisation of privacy enhancing technologies – A data spaces perspective. Data Brief. 2024 May 31;55:110560. doi: 10.1016/j.dib.2024.110560. PMID: 38948408; PMCID: PMC11214200.
19. Guidelines 02/2025 on processing of personal data through blockchain technologies Version 1.1 Adopted on 08 April 2025.
20. Auñón et al., Evaluation and Utilisation of Privacy Enhancing Technologies – A Data Spaces Perspective, 2024.
21. Readiness Analysis for the Adoption and Evolution of Privacy Enhancing Technologies Methodology, Pilot Assessment, and Continuity Plan APPROVED VERSION 1.0 PUBLIC DECEMBER 2015

Auditability

For PETs to maximise their utility as advanced compliance and accountability tools, the Digital Advertising industry must converge on a layered approach that combines mathematical verifiability, attestation mechanisms, auditable process logs, open-source transparency, and independent third-party certification.

Such a layered approach naturally separates into two questions: is the architecture of a PET system open to inspection, and can any given instance of processing be verified after the fact? The first is a matter of system design; the second is a matter of live operational proof.

Auditability of the System Architecture

Open algorithm: a known, signed method. Rather than relying on a bespoke, undisclosed method, a PET can be built on a publicly documented, peer-reviewed algorithm whose mathematical guarantees have already been scrutinised by the wider community. Cryptographically signing the specific implementation running in production closes the gap between what a company claims to run and what is actually deployed: any output can be traced back to a verifiable, unmodified version of the known algorithm, rather than an undisclosed variant substituted after the fact.

Open TEE: a known, signed execution environment. Trusted Execution Environments (TEEs) isolate a computation from the rest of the system, including the company operating it, and produce verifiable cryptographic proofs that the code executed exactly as specified without tampering.

Open stack: an open-source foundation. Open-sourcing the wider software stack (not just a single algorithm or enclave, but the code that orchestrates and deploys a PET) allows independent researchers, regulators, and civil society to audit the system's design. For example, the IAB Tech Lab's PETs Working Group has explicitly adopted open-source software development as a core pillar, working to define standards and projects that can be integrated into adtech stacks.

These three layers should be considered as independent axes of transparency, not a single pass/fail test, since a PET implementation could not be obliged to be open across all three to be considered auditable. What matters is that an organisation can give a considered answer at each layer when called upon to demonstrate the system's trustworthiness. A company might reasonably keep its stack proprietary while still pointing to a signed algorithm and a TEE backed by

a valid Attestation Certificate, and the strength of the overall case rests on being able to account for the posture adopted at every layer, not on maximising openness across all three at once.

Auditability of the Live Processing

Mathematical proofs as a foundation for trust. The strongest form of auditability is one that requires no trust at all, where operational trust is established with absolute mathematical verification. Multi-Party Computation (MPC) protocols are designed with this property: they provide cryptographic proofs that a joint computation was executed correctly over private inputs without any party learning those inputs. Differential Privacy (DP) goes further by offering quantifiable and formally auditable guarantees through its core parameter, epsilon (ε), which bounds the maximum privacy loss any individual can suffer from a query.
Zero-Knowledge Proof (ZKP) constructions extend this logic further: in a Private Set Intersection (PSI), for example, ZKPs let the parties involved verify a match by cryptographically signing the payload to be intersected, without exposing the underlying PII or revealing any detail of the encrypted payload itself. This same property allows companies more broadly to demonstrate adherence to regulatory standards to auditors without needing to transfer or centralise the underlying customer datasets, since the proof itself is mathematically verifiable.

Auditable logs without data exposure. Beyond proof-of-computation, PETs are capable of generating auditable logs of code execution and data access that preserve privacy at the record level. Rather than exposing individual data points, these logs capture what operations were performed, when, and under what parameters, providing the procedural accountability that compliance frameworks require without reversing the advanced data minimisation benefits that PETs are designed to deliver. Audit logs should also include the initiator of the operations – to guarantee a matching role and prevent identity misuse. The ISACA 2024 white paper on PETs[22] highlights this approach in practice: a PET-based anonymisation and risk scoring process enabled a large automotive enterprise to conduct predictive analytics on sensitive data while generating repeatable, shareable audit trails that could be integrated within DPIA documentation.

22. Exploring Practical Considerations and Applications for Privacy Enhancing Technologies, ISACA, 2024

Proof of code integrity. Trusted Execution Environments (TEEs) address a different but equally important dimension of the auditability problem: the question of whether the code running inside a privacy-preserving system is the code that was promised. TEEs provide “remote attestation”, a cryptographic proof that a specific, verified, and untampered piece of code is executing inside a secure enclave isolated from the host operating system and any operator interference. Attestation tokens can be embedded in communications between MPC nodes, and any node that cannot produce a valid token can be treated as potentially compromised. This hybrid architecture not only strengthens security but also provides an auditable chain of evidence, demonstrating that the privacy-enforcing mechanism was intact. The auditability of a TEE can be verified without running any private data through the TEE and therefore trust can be established in advance.

Practical Use Cases

Measurement

Defining Measurement in Digital Advertising

At its core, **Measurement** is the process of quantifying the impact of an advertisement on a specific business objective. It is the analytical bridge between a consumer seeing a brand and that consumer taking a desired action.

It is generally categorised into three layers:

1. **Reach and Frequency:** Did the right people see the ad, and how often?
2. **Attribution/Conversion Measurement:** Which specific ad, channel, or touchpoint gets "credit" for a final sale or sign-up?
3. **Incrementality (Lift):** Would this sale have happened anyway, or did the ad specifically cause it?

Why Measurement Matters for the Digital Ads Ecosystem

Without robust measurement, digital advertising shifts from a data-driven exercise to a speculative expense. Its importance is defined by three pillars:

1. Accountability and ROAS (Return on Ad Spend)

Measurement allows businesses to treat advertising as a predictable investment. By linking a $1,000 ad spend to $5,000 in sales, companies can justify their budgets to stakeholders. It provides the "Proof of Performance" that a campaign actually worked.

2. Optimisation and Efficiency

Measurement acts as a continuous feedback loop. It tells advertisers which creative versions are failing and which platforms are working.

- **Example:** If measurement shows that video ads on Platform A convert at half the cost of static ads on Platform B, the advertiser can shift their budget in real-time to save money and increase total sales.

3. Strategic Decision Making

Beyond short-term sales, measurement helps brands understand long-term consumer behaviour. It identifies "Customer Acquisition Cost" (CAC) and "Lifetime Value" (LTV), helping a business decide if they can afford to grow or if they need to pivot their product strategy.

How Traditional Measurement Used to Work

For the past two decades, digital measurement has primarily relied on **Deterministic Tracking**, a method that creates a direct link between an ad exposure and a consumer action. This model is still being used today extensively due to its high level of accuracy and ease of implementation for small and large businesses alike.

The Mechanism: This approach typically utilises **Third-Party Cookies** or **Mobile Ad IDs (MAIDs)**. These are unique identifiers that stay with a user as they navigate between different websites or apps.

The Process:

1. **Tagging:** When a user interacts with an ad, a "pixel" or "tag" records their unique ID.
2. **Mapping:** When that same user later visits the advertiser's site to make a purchase, the system recognises the ID.
3. **Attribution:** Because the ID is the same in both places, the system can definitively state: *"The person who saw the ad at 10:00 AM is the same person who bought the product at 2:00 PM."*

The Current Landscape: While highly effective for optimisation, this model is undergoing a period of transition. There is an increasing recognition across the industry that the approach can be technically optimised. The industry is moving toward more data-minimisation strategies to streamline how information is processed.

The Proactive Shift: This transition is not merely a response to regulation; it is driven by a fundamental **evolution in industry thinking**. Some advertisers, platforms, and technologists are now actively collaborating to move toward "proving impact" through privacy-preserving architectures.

Modern Measurement with Privacy-Enhancing Technologies (PETs)

Building on the goals of traditional measurement, the industry is developing new measurement solutions that use **Privacy-Enhancing Technologies (PETs)**. This represents a technical evolution: achieving the same high-quality business insights and appropriate security while enabling data to be processed locally, minimising the need for data transfer.

The Mechanism: Modern solutions are rarely "one size fits all". They often utilise a **modular stack of PETs**, selecting the specific technology – or combination of technologies – that best fits the use case. Here are examples of PETs that can be used for measurement solutions.

- **Secure Multi-Party Computation (SMPC):** Used when two parties want to "co-compute" an answer (like a sales match) without either party ever seeing the other's data.
- **Trusted Execution Environments (TEE):** Used to create a hardware-secure "enclave" where data can be processed in total isolation from the rest of the system.
- **On-Device Processing:** In some models, the measurement logic is moved directly to the user's browser or smartphone, ensuring that personal raw data never leaves the device.
- **Differential Privacy (DP):** Often used as a final "safety layer" across various methods, adding mathematical noise to aggregated reports to ensure no individual can be re-identified from the results.

The Process:

1. **Data Minimisation:** The publisher and advertiser prioritise "local" data. They only share the absolute minimum amount of information required for the specific measurement task.
2. **Encrypted Interoperability:** Whether through a "Data Clean Room" or a decentralised protocol, the matching of "Ad Exposure" to "Sale" happens in a mathematically obscured state.
3. **Aggregated Reporting:** Instead of individual logs, the system outputs a summary of group behaviour. This ensures that the advertiser learns *what* worked (e.g., "This campaign had a 4x ROAS") without knowing *who* specifically converted.

Real-Life Example: Google On-Device Conversion Measurement

A leading example of the shift toward privacy-preserving measurement is Google's **On-Device Measurement (ODM)** for iOS app campaigns. This solution is designed to help advertisers understand if their ads led to an app install or a specific action without ever disclosing the identity of the individual user.

How It Works (Simplified):

In traditional advertising, a company would share a list of customers with a platform to see who "matched". In the ODM model, the matching process is decentralised:

1. **The Ad Interaction:** A user clicks an ad for an app (e.g., a gaming app). This interaction is recorded by the platform.
2. **The Conversion:** The user later installs and opens that app. The app's measurement tool (SDK) recognises this event.
3. **The On-Device Check:** Instead of the two parties swapping lists, the matching process is sent to the **user's own device**. The phone checks if the person who clicked the ad is the same person who opened the app.

How PETs Are Built-In

This solution is powered by a combination of multiple Privacy-Enhancing Technologies to ensure that no identifying data ever leaves the device:

- **On-Device Processing:** The most sensitive part of the measurement—the "matching"—happens entirely on the user's smartphone. In this architecture, the required data processing occurs locally without needing to sync user lists with external servers.
- **Private Set Intersection (PSI):** This is a cryptographic protocol that allows two parties to find common data points (like an email or an Event ID) without either party ever seeing the other's full list. It's like two people comparing locked suitcases to see if they contain the same item without ever opening them.
- **Cryptographic Hashing:** All identifiers are turned into scrambled, one-way codes. In newer versions, these codes are "transient", meaning they change daily and are unique to each app, making it impossible to track a user across different advertisers.
- **Differential Privacy:** To further protect anonymity, "placeholder" or "fake" data is added during the process. This ensures that a single user's actions are hidden within a larger group.

Real-Life Example: Trusted Execution Environments (TEEs) for Data Matching

When two different companies – such as an advertiser (who knows who bought a product) and a publisher (who knows who saw an ad) – want to combine their information to measure success, they use a **Trusted Execution Environment (TEE)**.

Think of a TEE as a "Digital Safe Room" inside a computer server where data can be processed, but even the person who owns the computer cannot see what is happening inside.

How It Works:

1. **Encrypted Delivery:** The Advertiser and the Publisher both encrypt their data (like customer email lists) before sending it to the TEE. Because it is encrypted, no one can read it while it's in transit.
2. **Isolated Processing:** Once the data is inside the TEE, the "Safe Room" locks. The computer then decrypts the data and looks for matches (e.g., *"Did this person who saw the ad also buy the shoes?"*).
3. **Hardware Protection:** The security isn't just in the software; it's built into the computer's physical chips. This ensures that even if a hacker (or a dishonest employee) had total control over the server, they still couldn't "peek" inside the TEE while it was working.
4. **Verified Results:** Once the matching is done, the TEE creates an aggregate report (e.g., *"The campaign resulted in 500 sales"*) and then deletes the raw data.

How PETs Are Built-In

This architecture relies on several "layers" of PETs working together:

- **Confidential Computing (Hardware Isolation):** This is the core PET. It uses specialised hardware to create a memory space that is invisible to the rest of the computer's operating system.
- **Attestation:** Before any data is sent, the TEE provides a cryptographic proof (an "attestation") to the companies. This proves exactly what code the TEE is running, ensuring that no "spyware" has been added to the process.
- **K-Anonymity or Differential Privacy:** To make sure the final report is truly private, the TEE often ensures that data is only shared in groups. If only one person bought a product, the TEE might refuse to report it to prevent that person from being identified.

Audience Targeting

Defining Audience Targeting in Digital Advertising

At its core, **Audience Targeting** is the process of identifying and reaching specific groups of consumers based on their interests, behaviours, or existing relationships with a brand. It ensures that advertisements are relevant to the person seeing them, rather than being shown at random.

A core task is to predict the probability that a user will click on a given advertisement, so that campaigns can select relevant audiences, improve ad delivery, and optimise performance. This normally requires training models on examples that link user characteristics to individual outcomes, but privacy constraints may limit access to such outcome data. The industry is therefore moving toward architectures that preserve the ability to model audience intent and optimise campaigns while reducing the exposure and transfer of raw personal data.

In a modern context, targeting can be broken down into four key layers:

1. **First-Party Targeting:** Reaching your existing customers using data they have shared with you directly, such as email lists or phone numbers.
2. **Retargeting:** Reaching users who have previously interacted with your website or app - for example, by adding an item to a cart - to encourage them to complete a desired action.
3. **Interest-Based Targeting:** Reaching new people based on their likely hobbies, lifestyle, or intent signals. This includes reaching "Fitness Enthusiasts" or users currently looking to purchase a new car based on their activity across the web.
4. **Lookalike / Predictive Modelling:** Using machine learning to find new potential customers who share similar characteristics with your existing customer base.

Why Audience Targeting Matters for the Digital Ads Ecosystem[23]

Effective targeting is what makes the modern internet economy viable. It transforms advertising from a nuisance into a helpful service for both consumers and businesses.

1. Relevance for the Consumer

Targeting increases the likelihood that users see ads for products they might actually need or enjoy, irrespective of whether the advertiser knows the identity of the user. Without it, a consumer might see ads for car insurance when they don't drive or pet food when they don't own a pet.

Targeting is especially important on general audience digital properties (websites, apps) and content (such as specific shows on digital TV or specific pages on a news website). With digital properties or content that have a niche audience, it is easier to determine relevance on the basis of the context of the digital property or content. With general audience ones, on the other hand, relevance is more difficult to establish without targeting.

2. Cost Efficiency for Small Businesses

Small businesses benefit from targeting in a significant manner. Businesses that have an eminently local audience and that cannot easily sell beyond a specific region do not benefit from advertising far beyond that location; businesses that have a global reach but are unknown can use targeting to reach consumers they could never reach through traditional advertising. For instance, targeting allows a local bakery to show ads only to people within a 5-mile radius who like "Pastries," rather than wasting their limited budget showing ads to an entire country. It levels the playing field for smaller spenders.

3. Sustaining the Open Web

Most of the "free" internet - from news sites to social media - is funded by advertising. Effective targeting allows publishers to charge more for their ad space because the ads are more likely to result in a sale, which in turn keeps high-quality content and services free or at lower cost - and to reduce the number of adverts that need to be visible (even for financial viability) to a user at any given time.

23. A closely related subset of Audience Targeting is **Retargeting**. While standard targeting looks for new potential customers based on interests, retargeting serves ads to users who have already interacted with a business, such as visiting a website or leaving an item in a digital shopping basket. This allows businesses to re-engage highly interested consumers, helping them find products they have already considered while significantly optimising ad spend.

How Traditional Audience Targeting Used to Work

For the past two decades, digital advertising has utilised **Deterministic Identifiers** to understand consumer journeys across the web. This model remains a highly precise and effective way for businesses to reach the right audience and is still a primary method used across the industry today.

The Mechanism: This approach relies on a variety of stable identifiers, including **Third-Party Cookies**, **Mobile Ad IDs (MAIDs)**, and other **Deterministic IDs** (such as hashed emails). These identifiers act as a consistent signal as a user moves between different websites or mobile applications.

The Process:

1. **Observation:** As a user interacts with different digital properties, these identifiers allow the system to note specific actions, such as visiting a travel blog or searching for a specific product.
2. **Mapping Intent:** These various signals are associated with the identifier to build a picture of the user's current interests and needs.
3. **Delivery:** When the user arrives at a site with ad space, the system recognises the identifier, checks the associated interest or intent signals, and serves a relevant ad.

The Current Landscape: While these methods provide the high level of accuracy that many businesses rely on for growth, the industry is increasingly focused on developing new architectures that achieve this same precision without relying on long-term logs of individual browsing activity.

The Proactive Shift: This transition is driven by a fundamental **change in industry thinking**. Advertisers, platforms, and technologists are now actively collaborating to move toward models focused on aggregate impact measurement through advanced, privacy-by-design architectures. The goal is to reach a nearly equivalent level of ad relevance and business utility by transitioning to next-generation technical standards that process data at an aggregate level.

Modern Audience Targeting with Privacy-Enhancing Technologies (PETs)

While the deterministic identifiers discussed previously have been effective, the industry is continuously innovating to develop new frameworks that deliver equivalent value through different technical mechanisms. The goal is clear: allow brands to continue delivering relevant ads and reaching their customers through technical frameworks that align with evolving data minimisation principles and technical standards.

The central challenge for advertisers and publishers today is how to continue leveraging their valuable first-party data—such as customer email lists or purchase histories—for effective ad campaign targeting. The industry is developing complementary solutions that allow brands to activate this data locally, generating insights without needing to centralise or transfer the underlying datasets to advertising platforms. This is the critical opportunity that **Privacy-Enhancing Technologies (PETs)** are now being deployed to solve.

When paired with advancements in **Artificial Intelligence (AI)**, PETs can be used to build ad products that reach a new balance between privacy and commercial utility. As PETs obscure individual identities, AI steps in with improved prediction capabilities, allowing systems to model high-performing audiences based on aggregated, privacy-safe signals rather than individual tracking.

The Mechanism: Rather than relying on a single tracking ID, modern targeting uses a **modular stack of PETs** and machine learning to ensure data remains secure. For instance, some of the following techniques can be used:

- **Trusted Execution Environments (TEE):** Very useful when leveraging first-party data. Advertisers can match their customer lists against a platform's audience within a hardware-secured "Digital Safe Room", ensuring raw data is never exposed to the platform or the public.
- **On-Device Processing:** Instead of a central server learning a user's interests, the user's own phone or browser categorises their behaviour locally. The raw browsing history never leaves the device.
- **Secure Multi-Party Computation (SMPC):** Allows multiple companies to analyse overlapping audiences (such as mutual customers) without any single company seeing the other's private, underlying data.
- **Differential Privacy (DP):** Adds mathematical noise to audience segments so that a group cannot be granular enough to identify a single individual within it.

The Process:

1. **First-Party Data Protection:** The advertiser and publisher prioritise keeping their raw data on their own servers. When data must be used for targeting, it is cryptographically hashed or encrypted before leaving their servers.
2. **Privacy-Safe Matching:** Whether through local on-device learning or a secure TEE, the system identifies relevant audience segments (e.g., matching an encrypted email to serve a "Retargeting" ad) in a mathematically obscured state.
3. **AI-Driven Prediction:** Because individual-level tracking is minimised, advanced AI models analyse these privacy-safe, aggregated data points. These improved prediction capabilities allow the system to accurately identify intent and build "Lookalike" audiences by finding patterns in the aggregate data, all without needing to know exactly who the underlying users are.
4. **Actionable Delivery:** The ad system serves highly relevant ads to these mathematically verified cohorts. The advertiser gets an audience that provides high chances of driving sales, while the consumer's identity remains completely anonymous.

Real-Life Example: Publisher Advertiser Identity Reconciliation (PAIR)

A leading industry standard for first-party data collaboration is the **PAIR (Publisher Advertiser Identity Reconciliation)** protocol, which is managed as an open standard by the IAB Tech Lab. PAIR is designed to allow an advertiser and a publisher to securely match their respective first-party audiences (e.g., matching a brand's customer list with a news site's logged-in readers) without relying on traditional identifiers or exposing raw data.

How It Works (Simplified):

1. **Independent Encryption:** Both the advertiser and the publisher independently encrypt their first-party data (such as authenticated user emails) using their own unique private keys.
2. **Clean Room Collaboration:** These independently encrypted lists are sent into a secure Data Clean Room environment.
3. **The Double-Blind Match:** The clean room applies both keys to mathematically compare the lists and find overlapping users. Because of how the encryption works, the system can verify a match without ever needing to decrypt the data back to its raw, readable state.
4. **Privacy-Safe Activation:** For every matched user, the system generates a unique, abstract "PAIR ID." The advertiser's buying platform (DSP) uses this abstract ID to bid on and serve the ad on the publisher's site. Once the campaign is over, the ID cannot be used to reverse-engineer the user's identity.

How PETs Are Built-In:

- **Commutative Cryptography:** This advanced mathematical PET allows data to be encrypted by multiple parties in any order and still matched successfully. It acts like a "multi-key safe"—it can only be utilised when both the advertiser and publisher authorise the match, ensuring that neither party can access or copy the other's underlying data.
- **Data Clean Rooms (Interoperability):** The matching logic is executed within isolated, secure environments. PAIR is designed to be interoperable, meaning the advertiser and publisher can even use different Clean Room vendors, and the PETs ensure the data remains secure and unpooled across the entire supply chain.

Addressing Common Misconceptions

As Privacy-Enhancing Technologies move from theoretical research to large-scale business deployment, they have naturally drawn scrutiny from various stakeholders.

Addressing these concerns requires clear communication about how they operate, how they are audited, and their intended outcomes. In this chapter, we talk about the most common criticisms we heard.

Many of these misconceptions make for intriguing soundbites, but they either ignore the technical reality of PETs or are based on assumptions about the behaviour of certain actors, assumptions that can themselves sometimes be brought into question but that do not affect in any event the technical reality of how PETs actually work.

> ***"PETs Are a Way for Tech Companies to Circumvent Privacy Laws"***
>
> The Misconception
>
> Critics argue that PETs are a form of "technological solutionism". They claim companies use these tools for "privacy-washing" while maintaining the same intrusive data practices.
>
> The Reality
>
> This view creates a false dichotomy. In reality, PETs turn legal requirements into verifiable code. Transitioning to PETs provides objective enhancements in data security and enables strict data-minimisation practices. PETs offer a new technical mechanism to operationalise privacy controls, providing programmatic ways to audit and verify data workflows at scale.

"PETs Make Processing Impossible to Audit"

The Misconception

PETs technologies obscure data processing, making it impossible to audit.

The Reality

As set out in the Auditability section, PETs can be audited in several complementary ways. Trusted Execution Environments provide remote attestation, a cryptographic proof of exactly which code is running inside the enclave, which can be verified before any personal data is shared. Differential Privacy offers quantifiable and auditable parameters: its epsilon value bounds the privacy loss any individual can suffer, and the noise added is calibrated accordingly. Multi-Party Computation protocols provide cryptographic proofs that a joint computation was executed correctly, without any party learning the others' inputs.

"PETs Are Theoretical and Not Practical for Real-World Use Cases"

The Misconception

PETs work in theory but aren't viable for real business operations.

The Reality

The deployment of PETs in digital advertising is a relatively new development, despite the technologies being well-established in other fields. However, major platforms are already operating them in live products. As the Practical Use Cases section shows, Google's On-Device Measurement for iOS app campaigns combines on-device processing, Private Set Intersection, cryptographic hashing and Differential Privacy to measure conversions without disclosing the identity of individual users. Additionally, the PAIR protocol, managed as an open standard by the IAB Tech Lab, already allows advertisers and publishers to match first-party audiences through interoperable clean rooms without exposing raw data.

"Anonymisation Is Not a PET"

The Misconception

Some claim that anonymisation is not a PET but just a data transformation.

The Reality

Anonymisation is recognised as one of the most mature PETs by the OECD[24]. From the perspective of data protection law, the process of anonymisation is, moreover, the ultimate PET, as anonymous data can, by definition, not be linked to an identified or identifiable natural person, rendering data protection laws inapplicable.

"Data Cleanrooms Are Inherently 'High Risk' Brokerage Activities"

The Misconception

Data Cleanrooms are categorised as "High Risk" in the context of data brokerage[25], and are seen as tools for surveillance rather than privacy protection.

The Reality

The technology itself is a PET designed specifically to prevent the risk of re-identification when combining large datasets. Cleanrooms provide a "zero-trust" environment where data is matched via encryption or pseudonymisation without the underlying raw data ever being exposed to the parties involved. Labelling them "High Risk" ignores the fact that they are a highly sophisticated technical solution designed to facilitate data collaboration while keeping datasets completely decentralised and mathematically obscured.

24. OECD, Emerging Privacy Enhancing Technologies: Current Regulatory and Policy Approaches, 2023. [4] or Auñón et al., Evaluation and Utilisation of Privacy Enhancing Technologies – A Data Spaces Perspective, 2024. [7].
25. EDPB, R. d'Hauwers, Data Brokers Market Study

Recognising PETs in Policy-Making

PETs can play a greater role in how the Digital Ads Ecosystem implements the GDPR and how regulators assess compliance. Built correctly, PETs can help organisations reduce the amount of personal data exposed during processing and limit how readily individuals can be re-identified, while still preserving the underlying commercial model, personalised advertising, that funds ad-supported content and services.

That dual function of reducing risks on one side, while preserving data utility and enabling innovation on the other, is what makes PETs a strong candidate for more deliberate recognition, both in the legal bases organisations rely on, in the assessments that justify those choices, and in how regulators enforce the GDPR itself.

Risk Reduction

By further advancing data minimisation during processing and reducing the identifiability of individuals, PETs introduce additional technical safeguards that can be objectively measured. This strengthens organisations' ability to demonstrate that appropriate technical and organisational measures are in place to support compliance with Articles 32 and 5(2) of the GDPR.

Regulators should adopt an enforcement approach that accounts for the advanced technical safeguards and enhanced data minimisation capabilities PETs deliver. During investigations, the use of PETs should count as evidence of compliance with Data Protection by Design and by Default and should be treated as a mitigating factor where appropriate.

Similarly, when setting corrective measures and fines, regulators should take the use of (and investment in) PETs into account as at least a mitigating factor, since failing to do so removes much of the incentive to adopt these tools in the first place.

Finally, a clear, consistent and context-based approach to the definition of personal data is also essential to encourage the development and deployment of PETs, in line with the European Commission's proposal in the Digital Omnibus. Indeed, where PETs effectively reduce the identifiability of individuals, this should be duly reflected in the assessment of whether the resulting information qualifies as personal data, taking into account the means reasonably likely to be used to identify individuals in the specific context. Failing to recognise the legal significance of such safeguards risks creating a disincentive to invest in privacy-enhancing technologies, by exposing organisations to legal uncertainty[26] precisely when they take steps to reduce privacy risks.

The above is especially true for smaller players, for which it is harder to enter the market due to high barriers to entry. Deploying PETs effectively requires specialised cryptographic and engineering expertise, substantial computing resources, and the capacity to integrate these tools into existing advertising and measurement workflows. Large, established companies are far better placed to fund these investments than start-ups, independent publishers or smaller advertisers. Legal uncertainty compounds the problem: where it is unclear how regulators will treat the output of a PET, or whether the resulting data will still be regarded as personal data, smaller organisations are the least able to absorb the cost. Left unaddressed, this risks entrenching the position of the largest players and weakening the competition on which privacy-protective innovation depends. Regulators can help lower these barriers by providing clear and consistent guidance on how PETs are assessed, by recognising their use as evidence of compliance and as a mitigating factor, and by supporting open standards and interoperable, reusable solutions, so that proportionate privacy protection is within reach of organisations of every size.

26. FEDMA Paper, Sections 4.1 and 5.2, The marketers' perspective on unlocking Privacy Enhancing Technologies (PETs) in personalised marketing and advertising communications

Supporting Reliance on Legitimate Interests

In the GDPR, legitimate interest as a legal ground is on equal footing with consent (as well as contract), yet regulators have often sought to dismiss this legal ground or heighten the requirements for a Legitimate Interest Assessment (LIA) to be deemed acceptable in their view. Regulators' concerns regarding an unjustified, poorly thought-out reliance on legitimate interest as a legal ground may be valid, but they should not detract from that legal ground's core principles and permissibility. In practice, this regulatory approach means consent remains the default that companies reach for when processing personal data under the GDPR and when accessing or storing information on a user's device under the ePrivacy framework.

In this context, where PETs further advance data protection by design, for instance through robust pseudonymisation, there are good reasons for regulators to take them into account as a means of reducing the privacy risk of a given processing activity. PETs can thus strengthen an organisation's conclusion that the rights, freedoms and interests of data subjects are taken into account and that the balancing test underpinning a LIA ends up in favour of the legitimate interest invoked. Recognising PETs within the LIA framework in this way gives organisations a genuine incentive to invest in them: a credible route to relying on legitimate interest, rather than defaulting to consent.

Indeed, the deployment of PETs allows organisations to demonstrate effective safeguards in support of a LIA, and greater regulatory recognition of PET-backed safeguards would provide organisations with legal certainty that their assessment is robust and acceptable. This is especially visible in online advertising, where consent remains the default requirement for processing regardless of the privacy-enhancing safeguards deployed. Yet consent to a non-PET approach is then just as valid and likely allows the use of more data (in raw form), without the need to invest in PETs. IAB Europe's Position on the Draft Digital Omnibus[27] includes recommendations to give PETs explicit standing and encourage broader adoption.

Giving PETs greater standing in this way would also help ease consent fatigue: reserving consent requests for processing that is genuinely high-impact would make those requests more meaningful and easier for users to evaluate.

27. IAB Europe's Position on the Draft Digital Omnibus on the Digital Acquis, iabeurope.eu/wp-content/uploads/IAB-Europe_Digital-Omnibus_Position-Paper_Feb-2026.pdf

Extending PET Recognition to AI-Driven Processing

The recognition argument for PETs should not be confined to advertising; it extends naturally to the AI systems increasingly embedded within the advertising stack, from audience modelling to campaign optimisation. Regulators are already treating this as a distinct area of technical scrutiny. The EDPS has discussed PETs at length in its recent Tech Sonar reports (2024-2025[28] and 2025-2026[29]), specifically exploring technologies such as machine unlearning and confidential computing. Furthermore, the EDPB's recent Guidelines[30] on web scraping in the context of generative AI highlight machine unlearning and PETs as potential future tools to improve the privacy protections of AI development. This growing body of regulatory engagement suggests that PETs are being read by supervisory authorities as a powerful mechanism for further embedding data protection by design into AI systems, not merely a nice-to-have.

Building on Emerging Regulator Guidance

The EDPB's Work Programme for 2026–2027[31] includes both finalisation of its Guidelines on Pseudonymisation and new guidance on anonymisation, and its own stakeholder consultation process has surfaced a clear industry ask: stakeholders have requested more detailed guidance on technical measures, including concrete examples of PETs that organisations can rely upon, alongside guidance on appropriate key management practices[32]. That gap is precisely what IAB Europe's guide can help fill by offering regulators a specific taxonomy grounded in how PETs function in digital advertising.

To encourage wider adoption of PETs, further consideration should be given by regulators and policymakers to promote greater standardisation in the development and use of PETs, including through assessment methodologies and criteria for evaluating their effectiveness or regulatory sandboxes. Building on the Commission's proposal in the Digital Omnibus to develop criteria for pseudonymisation, establishing common frameworks for PETs more broadly would help ensure technological neutrality in their development and a fair level playing field in their use.

28. TechSonar Report 2024-2025
29. TechSonar Report 2025-2026
30. Guidelines 03/2026 on web scraping in the context of generative AI (open for feedback until October 2026)
31. EDPB work programme 2026-2027: easing compliance and strengthening cooperation across the evolving digital landscape
32. Report on stakeholder event on anonymisation and pseudonymisation of 12 December 2025

Lucio Gagliardi
Digital Advertising Product Manager
gagliardi@iabeurope.eu

iab europe
Rond-Point Robert
Schumanplein 11
1040 Brussels
Belgium

@iabeurope
/iab-europe

iabeurope.eu