ITI Promoting Innovation Worldwide

ITI Feedback on the EDPB Guidelines 02/2026 on Anonymization

30 October 2026

The Information Technology Industry Council (ITI) is the premier voice, advocate, and thought leader for the global information and communication technology (ICT) industry. Our member companies include the world's leading innovation companies, with headquarters worldwide and value chains distributed around the globe.

I. Introduction

ITI welcomes the EDPB's attention to the essential issue of anonymization. Members of ITI from across the tech stack process and utilize data for a variety of purposes. Anonymization can act as a valuable tool for these activities, to ensure both protection of the user and compliance with a variety of regional legislation on data protection, including the EU's General Data Protection Regulation (GDPR).

As drafted, the Guidelines do not yet deliver the clarity organizations need. Several provisions point in different directions on questions that determine the outcome of an assessment, and the draft does not illustrate what successful anonymization looks like in practice. The result is uncertainty for controllers and processors seeking to apply the standard in good faith.

Furthermore, this consultation takes place while the EU co-legislators are negotiating the Digital Omnibus, including key provisions on the definition of personal data and entity-relative identifiability. The EDPB must ensure its Guidelines do not pre-empt these outcomes and should commit to reviewing the guidance once the Omnibus is finalized.

ITI's recommendations address the legal standard, the legal basis for anonymization, the outcome of an assessment, and its technical application. They also explain how the EDPB can better align its guidance with the principles of the EDPB Helsinki Statement.

II. Applying the legal standard

2.1 Assess the likelihood of identification in context:
The draft states that it rejects an absolute or zero-risk standard for anonymization, but several provisions reinstate one in substance, and it leaves key phrases such as "insignificant in reality" and "means reasonably likely to be used" insufficiently defined. The assessment should consider factors that make identification more likely and those that make it less likely. Without explaining how these should be weighed, the Guidelines provide limited

United States
700 K Street NW, Suite 600
Washington, D.C. 20001, USA
+1 202-737-8888

Europe
Rue Froissart 95, 1040
Brussels, Belgium
+32 (0)2 321 10 90

info@itic.org
www.itic.org
@iti_techtweets

practical help in applying the approach reflected in the CJEU's judgment in Breyer v Bundesrepublik Deutschland (C-582/14)[1].

The final Guidelines should set out a non-exhaustive framework for assessing likelihood. Relevant factors should include the availability and quality of auxiliary data, the cost, time, expertise and computing resources required for identification, conditions of access to the data, applicable safeguards, and likelihood that an identification attempt would succeed. The EDPB should explain how these factors interact, rather than treating the theoretical possibility of identification as decisive. The draft is presently inconsistent on this point: paragraph 26 lists time, cost and labor among the factors capable of taking means outside the assessment, while paragraph 92 discounts them.

The draft recognizes that identifiability depends on an entity's circumstances, but several provisions dilute that principle. Paragraph 21 suggests that data is personal for both a sender and a potential recipient if a transfer is reasonably likely and it "cannot be ruled out" that the recipient has a means of identification. This formulation replaces the clear reasonableness standard in GDPR Recital 26 with an overly strict requirement to exclude every theoretical possibility of identification, directly contradicting the jurisprudence. -Breyer (C582/14), Gesamtverband Autoteile-319/22)[2][3]413/23 -)[4], unambiguously reaffirm that[5] must be grounded in the practical circumstances of the specific transfer and recipient, rather than an unworkable zero-risk standard. Instead of stating that identification "cannot be ruled out", we would -a reference to "whether identification is reasonably likely in the practical circumstances." In Breyer (C582/14), for example, the Court directed the referring court to examine whether the conditions for obtaining additional identifying data were met. The draft Guidelines undermine the Court's judgment in stating that anonymized data should be considered personal data by the recipient.

In paragraph 26, the EDPB rewords the Court's use of "practically impossible" to simply "impossible" which is a significant change: the former can be achieved whereas the latter cannot. The EDPB should replace the "cannot be ruled out" with a reasonableness standard and restore the qualifier "practically" in paragraph 26. Furthermore, Recital 26 of the GDPR anchors the assessment in the technology available "at the time of the processing." Paragraph 88 departs from this by treating a possible future data leak as capable of

[1] Judgment of 19 October 2016, *Patrick Breyer v Bundesrepublik Deutschland*, C-582/14, ECLI:EU:C:2016:779.
[2] Judgment of 9 November 2023, *Gesamtverband Autoteile-Handel eV v Scania CV AB*, C-319/22, ECLI:EU:C:2023:837.
[3] Judgment of 9 November 2023, *Gesamtverband Autoteile-Handel eV v Scania CV AB*, C-319/22, ECLI:EU:C:2023:837.
[4] Judgment of 4 September 2025, *European Data Protection Supervisor v Single Resolution Board*, C-413/23 P, ECLI:EU:C:2025:645.
[5] Judgment of 4 September 2025, *European Data Protection Supervisor v Single Resolution Board*, C-413/23 P, ECLI:EU:C:2025:645.

supplying the information needed for identification, so that a hypothetical breach would alter the status of data without the breach having occurred.

The treatment of unauthorized access, negligence, and malicious conduct requires clarification. Paragraph 30 brings rogue employees, cybercriminals and other potential sources of unauthorized access into the assessment, but their mere existence should not establish that re-identification is reasonably likely as no entity can eliminate every risk of intrusion. The list in paragraph 30 is unhelpful for those seeking guidance on compliance.

The same concern arises with paragraph 30's identification of foreign intelligence agencies as a relevant actor for identifiability. Read with paragraph 29, which provides that legal mechanisms enabling access to additional information, such as a compelled-disclosure provision, may constitute a 'means reasonably likely to be used', these provisions suggest that the existence of a foreign legal authority to compel access is sufficient to prevent data from being regarded as anonymous. This is irrespective of the likelihood that the authority would actually be exercised in the circumstances of the particular processing. The Guidelines should clarify that the existence of a legal power to obtain additional information should not, in itself, be determinative. This would preserve the Guidelines' ability to account for genuine government-access risks while avoiding an effective categorical barrier for organizations subject to foreign legal regimes.

The combined effect of paragraphs 21, 28 and 30 risks being understood as requiring organizations to investigate re-identification possibilities across an indeterminate chain of recipients, onward recipients, public data sources and malicious actors, even where those actors do not have a realistic path to receive the data or re-identify an individual. The GDPR already provides a governance framework for processing chains, including the controller's accountability obligations, Article 28 processor and sub-processor arrangements, technical and organizational measures, and transparency obligations. The Guidelines should build on this framework without introducing an unstandardized expectation of downstream due diligence through the anonymization assessment.

These risks are properly weighed in the anticipatory security assessment under Article 32 of the GDPR, rather than treated as sufficient to establish that information constitutes personal data under Article 4(1). Paragraph 89 correctly states that, where appropriate controls are in place, access is a means reasonably likely to be used by authorized users, but not by unauthorized actors. If the assessment must account for the capabilities of any actor anywhere, no entity's perspective differs from any others, and this deprives the draft's entity-relative test of practical meaning. This should be clearly expressed in Paragraph 30.

The assessment should consider evidence such as actual access pathways, prior incidents, strength of security controls, and the likelihood that access would result in successful re-identification. Similarly, paragraph 31's reference to identification resulting from "an accident or negligence" should concern reasonably foreseeable events, rather than any

conceivable mistake. The treatment of unlawful means raises a related concern: paragraph 32 accepts that means prohibited by law may be disregarded and that people may generally be assumed to follow the law, while paragraph 33 allows that assumption to be rebutted by evidence that a prohibition is not effectively monitored or enforced. Whether a counterparty will observe a prohibition ultimately depends on future conduct that no controller can establish with certainty in advance. Moreover, Breyer (C582/14) excludes means prohibited by law without making that exclusion conditional on the effectiveness of monitoring or enforcement.

The EDPB should provide further guidance on how AI-driven reductions in the time needed to combine data sources affect the assessment of "means reasonably likely to be used." As drafted, the test risks becoming open-ended if it relies on the resources, cost, labor and technical capabilities theoretically available to a large organization. This would mean that large entities could effectively never conclude that data is sufficiently de-identified, undermining any incentive to invest in de-identification. The assessment should focus on means that currently exist and are realistically available at the time of the processing, rather than speculative future capabilities. Practical criteria or examples that distinguish existing from hypothetical means would help controllers apply the test consistently and can be updated as technology develops.

The assessment should reflect how the data is shared. A controlled business-to-business transfer between vetted partners using secure infrastructure presents different risks from broad dissemination or sharing that gives more parties access. Paragraph 48's simplified approach may be a useful optional screen, but assessing all entities together without regard to their capabilities should not displace the recipient-relative legal test. The EDPB should replace the "cannot be ruled out" formulation with a reasonableness standard, restore the qualifier "practically" at paragraph 26, reconcile paragraph 92 with paragraph 26, reconcile paragraphs 30, 31 and 89, clarify that the existence of a legal power to obtain additional information is not determinative, and that threat modelling should reflect actual and reasonably foreseeable risks of the sharing environment. As written, the draft implies that no data shared outside the EU, even with GDPR protections, would ever be anonymous.

III. Giving effect to a recipient relative assessment

3.1 Illustrate anonymity from the recipient perspective

ITI welcomes the recognition that whether data is anonymous depends on the means reasonably available to the entity handling it. Data may remain personal for an organization that retains the means to re-identify, while being anonymous to a separate recipient that lacks them. The EDPB recognizes that retaining underlying data does not necessarily prevent the data shared with external parties from being anonymous, unless applicable law requires deletion. The guidelines should illustrate this distinction with an example in which the

organization producing anonymous data has a legitimate reason to retain the underlying data, such as maintaining customer records to provide an ongoing service.

The same reasoning should apply within a single organization. Data may be anonymous inside a defined environment that an organization has ring-fenced through technical and organizational measures, even where other parts of that organization retain identifiable data. Paragraph 10 already accepts anonymization "whether or not the original data is deleted", but every illustration of anonymization inside an organization is coupled with deletion of the source. That framing withholds recognition from organizations that de-identify data for their own internal use and so removes the incentive to do so, which is at odds with the data minimization principle in Article 5(1)(c). The EDPB should confirm that anonymity may be assessed intra-organization, and that retention of source data elsewhere in the organization does not of itself defeat anonymity where measures make re-identification not reasonably likely from the relevant internal perspective.

Furthermore, this recipient-relative perspective must extend to disclosures to law enforcement and governments. While law enforcement authorities may possess legal powers and resources to require ISP records and re-identify individuals from data such as IP addresses, the mere existence of these specialized government powers does not mean that other private entities handling the same data possess reasonable means of identification.

3.2 Recognize contractual and other safeguards

The draft provides that a contractual prohibition should not be equated with a prohibition by law, and that contractual terms may only complement technical measures since they can be disregarded. The EDPB should adopt a differentiated approach, considering the consequences attached to a breach of terms, including cases involving rogue employees. Contractual measures reinforce technical and organizational measures that can prevent a recipient from identifying individuals. Breaking contractual agreements often has serious consequences for the recipient and should decrease the likelihood that they will attempt to re-identify the data they have received under that contract. Where, for example, a breach results in termination of the contract and carries consequences comparable to those of breaching a legal prohibition, the contractual restriction may operate in practice as the equivalent of a legal prohibition. The draft guidance appears to be very reluctant to consider contractual measures meaningful, in contradiction to the recipient-relative analysis endorsed in EDPS v. SRB.

There is a difficulty of principle. Recital 26 requires account to be taken of all objective factors influencing the likelihood of identification, and Article 4(1) poses a single question. Neither supports a hierarchy in which technical measures are constitutive and legal ones merely complementary. In Breyer (C582/14), the Court excluded means "prohibited by law" from the assessment, and its concern was the existence of a legal impediment rather than the instrument in which the prohibition is located: obligations validly assumed by contract

bind the parties and are enforceable through the courts. The GDPR relies on contract in the same way, making contractual necessity an autonomous legal basis under Article 6(1)(b), the instrument that constrains a processor under Article 28, and an appropriate safeguard capable of rendering third-country transfers lawful under Articles 46(2)(c) and 46(3)(a).

The EDPB should clarify that the likelihood of identification must be assessed considering the contractual restrictions applicable to the particular recipient, alongside the technical and organizational measures in place. Once technical measures (such as encryption, tokenization, or data segregation) have reduced the practical ability to identify individuals, enforceable contractual terms (e.g., clear prohibitions on re-identification, purpose limitations, etc.) carry substantial weight for authorized enterprise counterparties. These are not merely commercial promises; under Article 28, they form part of the GDPR accountability framework.

This assessment should consider evidence of whether the recipient is likely to comply. A history of contractual breaches, for example, may limit the weight given to those restrictions. Where a recipient has a record of respecting its obligations and restrictions are supported by effective safeguards, this should reduce the assessed likelihood of identification. The absence of a prior relationship should not count against a recipient, since that would disadvantage first-time data sharing precisely where contractual assurance matters most. The possibility that a contract could be breached should not, on its own, justify disregarding it. If entities must assume they cannot trust the rule of law, then no data in any circumstance could be considered anonymous.

3.3 Keeping the Simplified Approach Optional

The draft presents the simplified approach as a voluntary way for an anonymizing controller to assess data without distinguishing between entities' access to information or their ability to identify individuals. It states that this is not an alternative legal standard. Paragraph 48 describes it as "a way to voluntarily shift the risk from false positives to false negatives, where the anonymizing controller treats anonymous data as personal because they have (in effect) overestimated the likelihood that certain means will be used". Without clearer safeguards, the approach could become a default in practice, leading organizations to classify data as personal even where it would be anonymous to a particular recipient under a contextual assessment.

A conservative initial screen may be useful, but the Guidelines should make clear that the contextual, entity-relative assessment determines whether data is personal data under the GDPR. The simplified approach should remain a permissible, optional tool, and its use should not create a presumption that the data is legally personal. Controllers should be free to move from a simplified assessment to a documented contextual assessment without any adverse inference. This clarity would help avoid systematic over-classification and preserve incentives to use privacy enhancing techniques and appropriately protected data.

3.4 Recognizing when data may be anonymous for a processor

In several passages, the draft requires a processor to assess relevant perspectives, purposes, and reasonably likely means in a manner that may effectively place the processor in the controller position. For example, Paragraph 15 of the draft appears to say that a processor must always assess the data from the controller's perspective, so information that is personal data for the controller can never be anonymous for the processor. This overlooks the contextual approach articulated in the draft guidelines and appears to go beyond the CJEU jurisprudence in the SRB case.

The final Guidelines should not treat paragraph 15 as resolving a legislative issue that remains under negotiation in the Digital Omnibus. They should explain how the controller–processor role allocation interacts with the entity-relative assessment of identifiability and the actual technical and organizational capabilities of the processing entity.

At minimum, the Guidelines should distinguish between:

- an entity that can access identifying information, combine datasets or use pseudonymous identifiers to link individuals across contexts; and
- an enterprise service provider that receives only pseudonymized or encrypted data, does not hold the key or additional information needed to identify individuals, operates in a controlled environment and cannot use the data to attribute information to, profile or otherwise treat a particular person differently.

Even if in the majority of cases, personal data transferred from a controller to a processor will be identifiable and will trigger the obligations of article 28 GDPR, there may be situations in which a controller may share only a limited dataset with a processor that does not provide them with the means reasonably likely to enable identification of the data subject. A controller may use contractual and technical safeguards that prevent the processor from accessing identifying information or attempting re-identification, particularly through the deployment of modern Privacy-Enhancing Technologies. An absolute attribution rule fails to reflect multi-party computation and collaborative architectures, such as Secure Multi-Party Computation (SMPC) or clean rooms operating within Trusted Execution Environments (TEEs). In these environments, such as cross-institutional fraud detection or shared research on rare diseases, multiple controllers combine encrypted inputs to generate aggregate, anonymous outputs. Because intermediate data is ephemeral, unintelligible, and technically inaccessible to the processor, attributing a contributing controller's perspective creates an artificial hurdle. The Guidelines should clarify that where PETs technically prevent the processor from accessing underlying information, identifiability must be evaluated against the processor's own actual capabilities.

The Guidelines should clarify how the controller-perspective approach in paragraph 15 operates alongside the contextual assessment of identifiability and 'chain of means' analysis

in paragraphs 20–21 and 28–29. They should clarify how it applies where a service provider cannot itself identify individuals but processes information on behalf of a controller that can, and how this should be distinguished from transfers to independent recipients. Paragraph 15, Example 3 should be revised: an example should be added involving pseudonymized data shared with a provider that lacks the key or additional information needed for identification.

IV. Anonymization is not a processing activity requiring its own purpose and legal basis

Paragraph 38 states that anonymization must have a legal basis under Article 6 of the GDPR, and an exemption under Article 9(2) where special category data is involved, with both presumed to coincide with those of the preceding processing where the anonymization forms part of the same processing activity and the controller pursues the same purposes.

Both conditions presuppose that anonymization has a purpose of its own to be matched against the purpose of what came before. Practically, it does not, and these conditions cannot be satisfied. Anonymization is not performed on data collected for a single purpose under a single legal basis, rather it is performed on consolidated datasets assembled from many collections. These collections can serve several purposes under different bases, for example, the basis of consent, the basis of legitimate interest, and the basis of legal obligation. Applied to such a dataset, the operation does not belong to a single activity, and there is no single set of "the same purposes" to match.

The Guidelines do not say what follows in that case. It remains unclear whether the presumption must be applied separately to each legal basis for every record, whether it ceases to apply and requires an independent legal basis, or whether it is intended merely as an illustration. The first would be unworkable on any scale. The second would leave organizations without a route to one of the most protective steps available to them. This would be particularly problematic where the resulting anonymous data are intended for a different downstream purpose, such as analytics or model development.

Purpose and legal basis attach to the processing activity as a whole, rather than to the individual operation. Article 4(2) GDPR defines processing as "any operation or set of operations". Articles 5(1)(b) and 30 locate the purpose at the level of the set: a processing activity is the purpose-bearing whole, which may comprise operations such as collection, recording, storage, alteration, retrieval, transfer and erasure. Article 6(1) links the assessment of lawfulness to the purpose served: where a legal basis requires necessity, that requirement can only be assessed once the relevant purpose has been identified. Anonymization is one such operation: it removes the personal character of the data, rather than using them for a purpose directed at an individual.

If the reasoning in paragraph 38 were followed, it would apply to every processing operation with equal force, requiring a separate legal basis under Article 6 for storing,

retrieving, restricting access to or deleting data. The Guidelines offer no principle to distinguish anonymization from the operations around it. This is illustrated by the examples of erasure and encryption: both are considered processing within Article 4(2) GDPR, and both are performed on personal data, yet neither requires a basis of its own. Encryption is also expressly recognized by Article 32(1)(a) GDPR as a security measure. Article 6(4)(e) GDPR goes further, listing pseudonymization among the safeguards supporting compatible further processing rather than among the operations requiring justification. Furthermore, the EDPB's own guidance describes pseudonymization as a security and data protection measure. Anonymization is the strongest measure in that family and should not be the only one requiring authorization.

Anonymization should be treated by default as compatible further processing, regardless of whether the personal data was originally processed under Article 6 GDPR or collected within the scope of the ePrivacy Directive. Controllers should not be required to conduct a separate compatibility test before taking a step that reduces privacy risk; imposing such a requirement would discourage anonymization and create legal uncertainty around a protective measure.

The practical effect of the draft is also counterproductive: an organization would need to justify reducing identifiability, while one that simply continues to process identifiable data doesn't need to do anything further. The Guidelines should clarify that anonymization, like collection, storage, transfer and erasure, is an operation governed by the purpose and legal basis of the processing activity of which it forms part. It does not require a separate purpose or legal basis. The condition that the controller must pursue "the same purposes" should therefore be removed. They should confirm that anonymization need not be presented to data subjects as a distinct processing purpose under Articles 13 and 14; the information provided about the underlying activity is the appropriate place to explain that data will be de-identified or anonymized.

V. Clarifying the Outcome of an Assessment

5.1 State the effect of meeting all three criteria

The draft is unclear about the conclusion organizations may draw when a dataset meets the No Record Isolation, No Linkage and No Inference criteria. The executive summary says such data "can be safely considered anonymous," while paragraph 52 says it "may be regarded as anonymous." The final Guidelines should use a consistent formulation and make clear that a documented assessment satisfying all three criteria gives rise to a presumption of anonymity, which a supervisory authority asserting the contrary would need to rebut, subject to reassessment if material circumstances change. Otherwise, organizations may have little certainty even after completing the assessment the EDPB prescribes.

5.2 Place the burden of proof on the party asserting that the GDPR applies

The GDPR applies only to personal data. The accountability obligation under Article 5(2) arises only once it has been established that the Regulation applies. The burden of establishing that information is personal data rests on the party asserting applicability. Paragraph 47 points the other way: by describing a conclusion of anonymity as a "false positive" to be guarded against (including where the assessor was unaware of a particular entity's means), it treats data as personal unless the assessor can exclude the means of every entity, including those it does not know of.

An organization cannot reasonably be asked to prove a negative against every conceivable actor and every future technique. Paragraph 41 presents documentation as a means of demonstrating compliance, but its practical value lies in supporting a conclusion that the Regulation does not apply. The final Guidelines should confirm that where an organization has documented a reasoned assessment and implemented safeguards, a supervisory authority asserting that the data is personal should identify a concrete re-identification pathway rather than a theoretical possibility.

5.3 Allow Separate Treatment of Mixed Datasets

The treatment of mixed datasets in paragraph 36 also needs clarification. The draft states that anonymization must be effective for every individual and that, where personal and anonymous parts “are not treated separately,” the dataset is personal data in its entirety. A successful re-identification of one individual, as contemplated in paragraph 87, should prompt an assessment of the affected data and any wider implications; it should not automatically determine the status of every separable record. The *Bundeskartellamt* passage cited in paragraph 36 concerned data collected together where the items could not be separated at collection, and whether the resulting personal data fell under Article 9 GDPR. It does not establish that separable non-personal data becomes personal data solely because it is handled alongside a personal record.

The EDPB should distinguish between data that cannot be separated and data that an organization has simply not treated separately and explain how organizations can identify and manage anonymous portions of a mixed dataset. The operative questions should be whether the likelihood of re-identification is insignificant across the dataset as a whole, and whether identifiable records are capable of being separated, not whether risk is zero for every individual record.

The final Guidelines should also provide greater clarity on how organizations should address outliers, including whether appropriately identified and segregated subsets may be treated separately, what level of evidence is expected to demonstrate that the likelihood of re-identification is insignificant for the relevant population, and what constitutes proportionate testing, suppression and segregation of higher-risk records.

VI. Making the technical criteria operational

6.1 Explain how established methods satisfy the criteria

The draft explains the legal standard for anonymization but gives organizations limited guidance on how to apply it to specific techniques. The draft instead centers on three criteria (No Record Isolation, No Linkage, and No Inference) without explaining how organizations should assess these established methods against them. As a result, organizations using strong safeguards may still be uncertain whether, and under what conditions, their approach meets the EDPB's criteria.

The Guidelines should map common anonymization techniques and privacy enhancing technologies to the three criteria, with practical examples explaining the role and limits of each measure. Such a mapping should illustrate which technique can satisfy the criteria rather than prescribe a set of measures and should be stated to be illustrative rather than a minimum specification. The underlying standard should remain technology-neutral and outcomes-based: it should state the privacy result required and not prescribe the means of achieving it, so techniques not yet in common use are not excluded by implication.

6.2 Address modern privacy enhancing technologies

The Guidelines should give greater attention to privacy enhancing technologies (PETs), including trusted execution environments, secure multi-party computation, and advanced cryptographic protocols. These tools can substantially reduce access to identifying information and the risk of re-identification, but their effect depends on how they are implemented. The EDPB should provide practical guidance on how such measures factor into the assessment of whether identification is reasonably likely, including strengths, limitations and use alongside other safeguards. Clear recognition of effective PETs would give organizations greater confidence to invest in stronger data protection while avoiding the assumption that any particular technology automatically renders data anonymous. The Guidelines should confirm that a properly implemented PET is capable of supporting a conclusion of anonymity. As presently drafted, each technique is paired with a caveat, and none are recognized as sufficient in any circumstances, which gives organizations no outcome to aim for and discourages the investment that Article 25 is intended to promote.

The guidance should distinguish methods that alter the identifiability of data from safeguards that restrict access and explain how each affects the overall legal assessment. This includes synthetic data, trusted execution environments, secure multi-party computation, and advanced cryptographic protocols.

6.3 Distinguish individual inference from general analytics

The No Inference criterion addresses genuine risks, including membership inference, extraction, de-aggregation, and AI model memorization. However, the draft does not distinguish between a specific, meaningful inference about an identifiable individual and population-level analysis, statistical relationships, or ordinary machine-learning generalization. Without that distinction, organizations may be uncertain whether statistics, synthetic data or trained models should be treated as personal data simply because they could reveal some information about a person.

The Guidelines should clarify when an inference makes an individual identifiable, applying the draft standard of whether identification is reasonably likely. The existence of an attack technique should be weighed against that threshold rather than treated as an automatic failure of the No Inference criterion. They should distinguish findings about groups or populations from information that can be attributed to an identifiable individual and explain when a trained model retains or reveals information sufficient for identification. Paragraph 71(b) is helpful in this respect and is a good example in the draft where an organization is told what a sufficient demonstration looks like: it recognizes that a lack of reliance can be shown by comparing outputs produced with and without a given individual's records. That approach merits fuller development. This clarification should be consistent with the EDPB's existing guidance on AI models and should confirm that the question is whether identification is reasonably likely in practice, rather than whether a controller can exclude every possibility of extraction.

VII. Supporting implementation and ongoing review

7.1 The need for positive use cases

While ITI appreciates the EDPB's examples of when attempts at anonymization do not meet the legal standard, as well as its illustration of how the success of anonymization depends on the specific context of the data recipient, there are few examples of how organizations can successfully anonymize data. Positive use cases would give organizations a clearer sense of how to approach anonymization in different contexts and demonstrate successful applications of these techniques. Focusing on negative examples could give the impression that anonymization is more likely to fail than succeed, discouraging companies from investing in these approaches despite their potential benefits for privacy and innovation.

7.2 Provide practical guidance for product teams and AI use cases

The draft provides a detailed explanation of the legal principles of anonymization but gives engineers and product teams limited guidance on how to apply them when designing and assessing systems. It does not explain how to assess identifiability in AI model training, synthetic data generation, or model evaluation, including where privacy enhancing technologies are used.

The final Guidelines should include actionable methodologies, practical examples and, where appropriate, measurable indicators to help teams assess re-identification risk. Examples should address training data, synthetic datasets and model outputs, identify the factors relevant at each stage, and explain how technical safeguards affect the assessment. This would help organizations translate the legal standard into product decisions and apply it consistently when developing privacy preserving AI systems. We are reminded of the EDPB's own commitment in Statement 6/2024 and its 2024-2027 Strategy to produce more practical guidance.

7.3 Make documentation and reassessment proportionate to risk

The draft calls for controllers to document their anonymization assessments and periodically reassess whether data remains anonymous, but provides limited guidance on the extent of documentation, the frequency of review, or when reassessment is not necessary. For large datasets, an open-ended obligation to revisit whenever data evolves could impose substantial burdens without a corresponding privacy benefit.

The final Guidelines should provide a model assessment record organized around the No Record Isolation, No Linkage, and No Inference criteria, with documentation expectations proportionate to the risk of re-identification. Reassessment should be prompted by a material change, such as wider access, new linkage data, a relevant security incident, a change in anonymization technique, or a demonstrated advance in re-identification capability. Reassessment should be proportionate to the nature of the data, access conditions, and residual risk. A well-documented assessment made in good faith should remain valid unless new facts call its conclusions into question.

7.4 Clarify terminology used in transparency information

Paragraph 40 cautions controllers against describing data as “anonymous,” “de-identified” or “de-personalized” where individuals remain identifiable. These terms do not necessarily make the same claim. In the AI context, organizations often use “de-identified” or “de-personalized” to indicate that measures have reduced identifiability, while recognizing that the resulting data may still be personal data. Treating these terms as equivalent to “anonymous” could therefore create confusion about the measures applied and the status of the data.

The EDPB should revise paragraph 40 to caution against misleading claims that data has been anonymized or is anonymous, while allowing controllers to use other terms where they accurately explain the measures taken and make clear that individuals may remain identifiable. These terms may create challenges for multinational organizations, including companies that operate under privacy frameworks using concepts such as ‘de-identified data’ that do not necessarily correspond to the GDPR concept of anonymous information. A dataset may be appropriately described as ‘de-identified’ under an applicable third country legal framework while not meeting the threshold for anonymous information under the GDPR, and such terminology should not be treated as inherently inaccurate.