TERLOUW.LEGAL / PUBLIC CONSULTATION

GUIDELINES 02/2026

Anonymisation

Technical comments on evidencing context, capability and re-identification risk

17 September 2026 | Netherlands

CORE CONTRIBUTION

Anonymous should not be a label attached to a file. It should be a claim with a defined perspective, an evidence base and a review trigger.

terlouw.legal is a Dutch legal services and legal technology organisation. Our contribution focuses on the point where legal qualification meets operational evidence. The draft gives a detailed framework for assessing anonymity. We address how an organisation can preserve enough evidence to show which context, entities, means, tests and uncertainties supported an anonymity conclusion at a particular time.

SCOPE OF THIS SUBMISSION

This contribution focuses on operational accountability, evidence quality and decision traceability. It does not propose a new technical standard, statistical threshold or sector specific rule. The aim is to test whether the draft can be implemented in a way that remains auditable after the relevant event.

Organisation: terlouw.legal | Amsterdam, Netherlands | https://terlouw.legal

Contact details are provided through the EDPB submission form and are not repeated in this public PDF.

Executive summary

The draft moves away from treating anonymity as an intrinsic property of a dataset. It requires attention to perspective, context, reasonably likely means, the availability of additional information, evolving technology and the three criteria of No Record Isolation, No Linkage and No Inference. That legal and technical structure is strong. The practical challenge is proving that the structure was actually applied.

PROPOSED IMPLEMENTATION LENS

Treat an anonymity conclusion as a decision record with a scope. The record should identify the data version, intended context, relevant entities, available means, external information considered, test results, uncertainty, safety margin and the events that require reassessment.

Six points for clarification

- Anonymity conclusions should identify the perspective and release context to which they apply.
- The relevant entity analysis should be preserved as a map of who can access the data, who can provide additional information and which chains of means are realistic.
- Reassessment should be linked to concrete triggers such as wider release, new auxiliary data, capability changes, incidents and material advances in re-identification tools.
- Security, contracts and encryption can support the context but should be recorded separately from the technical anonymity conclusion.
- AI and synthetic data require repeatable evidence of inference and membership testing, including tool and model versions.
- The conclusion that identification is insignificant in reality should state its factual basis, uncertainty and safety margin without pretending that one universal numeric threshold exists.

01 Record the perspective and time of the anonymity claim

Draft reference: paras. 11-13, 22 and 35

Observation. The draft states that information may be anonymous from one entity's perspective but personal from another, and that re-identification likelihood can change over time.

Implementation issue. Organisations often mark a dataset simply as anonymous in a data catalogue. That label can travel into new uses without the original recipient assumptions, access restrictions or technical state being visible.

Clarification requested. Please clarify that a defensible anonymity decision should state the perspective or perspectives for which the conclusion was reached, the intended use or release context, the date of assessment and the data version. Where the conclusion depends on restricted access or a defined recipient group, that dependency should travel with the dataset or release decision.

> **PRACTICE EXAMPLE**
>
> A dataset assessed for disclosure to one research partner may not have the same status if later published openly. The underlying bytes may be identical while the relevant entities and available means are materially different.

02 Turn the relevant entity analysis into a visible evidence map

Draft reference: paras. 28-30

Observation. The draft gives a broad meaning to means and recognises chains of means across entities, including recipients, rogue employees, journalists, authorities, unethical companies and cybercriminals depending on context.

Implementation issue. A narrative conclusion such as re-identification is unlikely can hide which entities were actually considered and which auxiliary information or services they could obtain. This is particularly difficult to reconstruct after a security incident or a later disclosure.

Clarification requested. Please clarify that the assessment can be evidenced through a relevant entity map. For each material entity or category, the record can identify access to the given data, additional information reasonably obtainable, technical or organisational capability, likely chains to other entities, and the objective factors supporting inclusion or exclusion from the assessment.

03 Use event-driven reassessment in addition to periodic review

Draft reference: paras. 35 and 88-93

Observation. The draft notes that re-identification likelihood increases over time, encourages periodic reassessment and recognises that data leaks, lower costs and developments in AI can change the means reasonably likely to be used.

Implementation issue. A fixed annual review can miss the event that actually changes the legal conclusion. Conversely, constant reassessment is not realistic for every dataset.

Clarification requested. Please clarify that periodic review can be complemented by defined reassessment triggers. Examples include a new recipient, public release, a material change in access controls, a security incident, a new auxiliary dataset, a significant reduction in attack cost, a new model or agent capability, or a change in the intended use. The trigger should reopen only the parts of the analysis affected by the changed fact.

> **PRACTICE EXAMPLE**
>
> This approach turns the draft's temporal reasoning into a manageable control process and makes it possible to explain why a conclusion remained valid until a particular change occurred.

04 Keep supporting controls distinct from the anonymity conclusion

Draft reference: paras. 34 and 89-91

Observation. The draft states that contractual, organisational and technical measures can restrict access but are not by themselves sufficient to make data anonymous, and that encryption is reversible and differs from anonymisation.

Implementation issue. Operational documentation often merges these questions. A dataset may be described as anonymised because it is encrypted, contractually restricted or held in a secure environment. That can obscure whether the three anonymity criteria are satisfied and from whose perspective.

Clarification requested. Please clarify that the assessment record should separate: the technical and contextual anonymity analysis, the access and security controls that affect available means, and the contractual restrictions that affect behaviour. This separation preserves the legal distinction while still allowing those controls to influence the contextual assessment where appropriate.

05 Make AI inference testing reproducible enough to audit

Draft reference: paras. 66-82 and 83-87

Observation. The draft recognises that inference can make information personal, applies this to AI models and synthetic data, and refers to techniques ranging from simple searches to special-purpose AI agents.

Implementation issue. A statement that an AI system was tested for membership inference or re-identification is difficult to evaluate if the model version, prompts, attack method, auxiliary information and success criterion are not preserved. Results may also change with tool updates.

Clarification requested. Please clarify that organisations may evidence AI-related anonymity testing through a reproducible test record. At minimum, this can identify the data or model version, attack or inference class, tool and version, auxiliary information used, test population, success criterion, observed result and known limitations. This is evidence of the assessment, not a universal certification of anonymity.

> **PRACTICE EXAMPLE**
>
> Where direct tracing is impossible because the processing is opaque, the draft itself discusses comparative approaches. A controlled record of those tests would make the reasoning reviewable later.

06 Document the basis for insignificant likelihood without a false universal threshold

Draft reference: paras. 22, 29 and 87-93

Observation. The draft frames anonymity around whether identification is insignificant in reality and recognises that successful techniques need not have perfect accuracy or certainty.

Implementation issue. Organisations may respond by inventing a fixed numeric threshold and applying it across contexts. That can create false precision because the assessment also depends on the effect of identification, available means, external data, access context and the relevant entity.

Clarification requested. Please clarify that the decision record should state the factual basis for the conclusion, the confidence and uncertainty of relevant tests, any safety margin applied and the assumptions that would invalidate the conclusion. The final Guidelines could usefully distinguish evidence of a reasoned threshold in context from an expectation of one universal percentage.

ANNEX / ANONYMITY DECISION RECORD

The following is not proposed as a mandatory form. It illustrates how the draft framework can be preserved as auditable evidence without reducing anonymity to a permanent label.

FIELD	EVIDENCE TO RETAIN
Data asset and version	Dataset, output, model or synthetic data version; hash or immutable identifier; date of assessment.
Intended context	Purpose, release model, access conditions, expected recipients, onward transfer assumptions and retention context.
Applicable perspective	Entity or entities for whom the data is intended to be anonymous and why those perspectives are legally relevant.
Relevant entity map	Direct and indirect recipients, possible unauthorised actors, chains of means, access to given data and access to auxiliary information.
Reasonably likely means	Search, linkage, inference, external services, computational resources, time, cost, legal access routes and foreseeable capability changes.
No Record Isolation	Technique tested, data structure, result, confidence, outliers and limitations.
No Linkage	Auxiliary datasets and attributes tested, match method, result, confidence and any remaining linkage paths.
No Inference	Inference or membership tests, AI or statistical tools used, prompts or attack class where relevant, result and limitations.
Supporting controls	Security, access, contracts, organisational measures and encryption, recorded separately from the anonymity test itself.
Decision and safety margin	Conclusion by perspective, uncertainty, assumptions, approver and reasons supporting insignificant likelihood in reality.
Reassessment triggers	Public release, new recipient, incident, new auxiliary data, material attack capability change, lower cost, or changed purpose.

CONTROL PRINCIPLE

An anonymity conclusion is strongest when the organisation can show not only the result, but also the perspective, means, evidence, uncertainty and trigger that defined the result.

Closing observations

The draft provides a sophisticated legal and technical framework. Its practical value will depend on whether organisations can translate that framework into decisions that remain explainable when a recipient changes, a dataset leaks, a new linkage source appears or AI makes a previously expensive attack routine.

terlouw.legal therefore invites clarification on the evidence that should accompany an anonymity decision. The objective is not to turn anonymisation into a paperwork exercise. The objective is the opposite: to ensure that the conclusion can be tested against the facts that actually made it valid.

THREE QUESTIONS FOR THE FINAL GUIDELINES

1. Should an anonymity conclusion explicitly identify the perspective, release context and data version to which it applies?
2. Which events should normally trigger reassessment when capabilities or available information change?
3. What minimum testing record would allow No Record Isolation, No Linkage and No Inference conclusions to remain reviewable for AI models and synthetic data?

Source basis

This submission comments on the public consultation version of the EDPB draft. Paragraph references in the body refer to that version.

- **EDPB Guidelines 02/2026 on Anonymisation, version 1.0, adopted for public consultation on 7 July 2026:** https://www.edpb.europa.eu/system/files/2026-07/edpb_guidelines_202602_anonymisation_v1_en_0.pdf
- **EDPB public consultation page:** https://www.edpb.europa.eu/public-consultations/guidelines-022026-on-anonymisation_en
- **GDPR, Regulation (EU) 2016/679, including Recital 26 and Article 4(1):** https://eur-lex.europa.eu/eli/reg/2016/679/oj
- **CJEU case law referenced by the draft, including C-413/23 P EDPS v SRB:** https://curia.europa.eu/juris/liste.jsf?num=C-413/23%20P

PUBLICATION NOTE

The EDPB consultation page states that submitted comments are published after screening for unauthorised submissions. The PDF therefore contains no personal signature and no personal contact details. The organisation name is intentionally visible.

FACT. SOURCE. RULE. DECISION. ACTION.