Public consultation — EDPB Guidelines 02/2026 on Anonymisation

Comments submitted in response to the consultation (deadline 30 October 2026)

Respondent: Olivier Blazy, Professor of Cybersecurity, École polytechnique (LIX / Surycat team), Institut Polytechnique de Paris. These comments are submitted in a personal, individual capacity as an academic working in cryptography and privacy-enhancing technologies, and do not represent an institutional position.

1. General appreciation

I welcome these Guidelines and support their direction. Anonymisation carries a lot of weight in European data protection law: once data is anonymous the GDPR stops applying, so the whole question turns on the word "anonymous." The Guidelines set a demanding bar, embed the *EDPS v SRB* line of reasoning, and organise the technical analysis around No Record Isolation, No Linkage and No Inference. This is a real improvement on Opinion 05/2014, and it tracks the current re-identification literature instead of a decade-old snapshot of it.

Four choices in particular are right. Treating anonymity as entity-relative and perspective-dependent. Refusing to credit a controller's asserted "lack of motivation." Refusing to equate a contractual prohibition with a legal one. And recognising that re-identification gets more likely over time, and that AI, agentic AI included, keeps lowering the cost of the means. Each of these blocks a common route by which a supposedly anonymous dataset turns out to be personal data.

So the comments below are not an argument against the Guidelines being strict, which I think they should be. They ask the Guidelines to hold themselves to the rigour and precision they demand of controllers. Section 2 is my main point: the entity-relative approach treats the holder as fixed in both identity and capability, when in practice data changes hands through acquisition and asset sale, and a holder's analytic reach jumps with each model upgrade. Either shift can turn "anonymous" data into personal data without any processing event to pin the change on. Sections 3 onward deal with criteria a technical reader cannot apply the same way twice. These Guidelines draw the in-scope/out-of-scope line for the whole Regulation, so vagueness here is not cosmetic; it decides whether a given person is protected. I propose concrete wording throughout.

2. The central objection: entity-relativity does not survive a change in the holder, whether of identity or of capability

This is my main comment, and it goes to the foundation of the framework: anonymity is defined relative to the entity holding the data (paras 20–25, contextual approach at paras 50–64). Grant, for argument, that some relativity is unavoidable, since the same dataset really does carry different risk in different hands. The problem is that the Guidelines take the assessment as a snapshot. They fix both the set of relevant entities *and their means* at the moment of assessment. Neither is stable. The entity can change (acquisition), and the means available to a given entity can change (a model upgrade). Neither touches the data, yet under the Guidelines' logic each can flip its legal nature. And this is not some corner case; it is how the market and the technology normally behave (except if the union is moving away from a market capitalism).

2.1 Change of identity: acquisition and transfer

Consider a corporate acquisition. Company A holds a dataset it has assessed as anonymous, correctly, relative to A's own auxiliary information and means. A is acquired by B, or the dataset is transferred to B as an asset. B holds population registers, or a second dataset, or analytic capabilities that A did not. Relative to B, the records are identifiable. Nothing was done *to the data*: not re-processed, not enriched, not re-consented. The bits are the same the day before the deal and the day after. Under the Guidelines' own logic the data has nonetheless changed nature, and there is no processing event to attach that change to.

The Guidelines never face this. If a line change on an org chart, with nothing done to the data, can turn anonymous data into personal data, then "anonymous" here says nothing about the data. It says something about who holds it today, and holders change all the time, without warning and without the data subject ever knowing. The "reasonably foreseeable developments" test (para 29(e)) does not save this: in a live M&A market the future holder is exactly what you cannot foresee. Periodic reassessment (para 35), which I support, only makes the point sharper. If you have to re-check the classification every time the world moves around the data, the classification was never a property of the data.

I am not, in this comment, asking the Board to abandon the contextual approach. I am asking it to follow that approach to its own logical conclusion, which it currently stops short of:

- **Transfer and acquisition are re-identification-relevant events.** The Guidelines should state explicitly that any transfer of a dataset to a new holder, including transfer by merger, acquisition, insolvency, or asset sale, is an event that *requires* fresh assessment of anonymity relative to the new holder and the auxiliary information reasonably available to it. Reassessment on transfer should be named as a specific trigger, not left implicit inside the general "periodic" duty of para 35.
- **A transfer to a holder for whom the data is identifiable is a re-identification.** Where the receiving entity can identify individuals with means reasonably likely to be used, the data is personal data in its hands from the moment of transfer, and the transfer is the processing event that brings it back into scope. The transferor should not be able to launder the data's status by handing anonymous-to-me data to an entity for whom it is identifiable.
- **The transferor must be accountable for foreseeable re-identifying transfers.** An anonymity assessment that ignores the realistic possibility of the dataset changing hands is incomplete. Where onward transfer to a more-capable holder is a realistic prospect, as it always is for a commercially valuable dataset, the assessing controller should be required to account for it, rather than treat its own present incapacity as a permanent shield. This is the entity-relative counterpart to the Guidelines' own, correct, rejection of "lack of motivation": present incapacity is as unstable a foundation as present disinclination.

This matters for real people. Data that is brokered, sold, or swept up in an acquisition is one of the main ways "anonymised" data actually gets re-identified. A framework that calls data anonymous in the first holder's hands and then says nothing when it moves to the second leaves the person exposed at the worst possible moment, with no event they could see, object to, or bring a right against.

2.2 Change of capability: AI model updates as re-identification events

The same logic applies, and applies more urgently, to a holder whose identity does not change but whose *capability* does. An entity that assessed a dataset as anonymous against the models and tools

it ran at time T may, after a routine upgrade, be able to link or infer at T+1 what it could not before. The data is unchanged and the holder is unchanged; only its reach grew. The Guidelines' own text gives you everything you need to call this a re-identification event, paras 92–93 on agentic AI cutting the cost of assembling auxiliary information, and para 29(e) on foreseeable technological developments. They just never say it. A model upgrade is an acquisition of capability, and it should count the same as acquiring the dataset.

In one respect this is worse than the acquisition case: capability only goes up. An acquisition can in principle move data into weaker hands as easily as stronger ones, but model capability, over the life of a dataset, never falls. So an assessment pinned to the state of the art at time T comes with an expiry date built in, and the Guidelines should say so.

I would ask the Board to add the following, as the capability-side counterpart to the transfer triggers above:

- **The anonymity assessment must state an assumed model-capability envelope.** A controller claiming anonymity should be required to characterise the re-identification reach, for linkage and for inference, of the models and analytic tools it deploys, and of those it could reasonably foreseeably deploy over the data's sensitivity lifetime. This is the threat-model discipline that should govern the "insignificant likelihood" test generally (see Section 3.1), applied with the deployed or deployable model as the adversary. It does not require predicting a specific future model; it requires stating the capability envelope against which the claim of anonymity is made, so that the claim can be audited and revisited when the envelope is exceeded.
- **A capability increase that renders the data identifiable is a re-identification.** Where a deployed model becomes able to single out, link, or infer at the level the Guidelines prohibit, the data has become personal data in that holder's hands, with no processing event other than the capability change itself. A model update should be a named trigger for reassessment, exactly as transfer is.
- **Bounding now; explainability as the goal.** I recognise that full, interpretable explanation of a model's re-identification reach is not achievable for current frontier systems, and I do not ask the Board to require the unachievable. The enforceable requirement today is *bounding,* a documented, conservative envelope of assumed capability, revised upward as the state of the art advances. Full explainability of that reach should nonetheless be named as the direction of travel: as interpretability matures, the assumed envelope should give way to a demonstrated one. Stating the goal now lets the standard tighten over time without having to be re-litigated.

One trap needs closing head-on. A DPIA under Article 35 is only owed once you are processing personal data, so a controller claiming anonymity says no DPIA is due and the model-reach question never comes up. That is circular: the assessment that lets the controller off the hook is the one that has to account for foreseeable model capability. So the envelope belongs in the *anonymity assessment itself:* the analysis that decides whether the data is in scope at all, not pushed off to a DPIA the anonymity claim is being used to dodge. An exemption a controller grants itself against a threat that keeps growing after the fact is not one the data subject can rely on.

Fixing either of these, identity or capability, costs the contextual approach nothing. It only asks the Board to treat the holder as a variable: who it is, and what it can do. An anonymity that has to be re-

checked every time the holder is bought or its models are upgraded is not a property of the data. It is an assessment with a shelf life, and the Guidelines should label it as one.

3. Rigour gaps in the three criteria

3.1 "Insignificant" likelihood is never quantified

The Guidelines repeatedly turn on whether the likelihood of identification is "insignificant in reality" (paras 8, 22, 26, 32). This is the correct legal test, drawn from *EDPS v SRB*, and I do not suggest replacing it with a hard numeric threshold, which would be gamed and would ossify as techniques advance. But the Guidelines give the assessor no method for reasoning about "insignificant" at all. That leaves the word to mean whatever the controller needs it to mean.

Recommendation: add guidance that "insignificant" should be assessed against an explicit, documented threat model, stating the assumed adversary, the auxiliary information assumed available to it, and the success rate the technique would need to achieve to distinguish an individual. The privacy literature already provides principled framings here (e.g. bounding an adversary's posterior advantage over its prior). The Guidelines need not mandate a specific formalism, but they should require that the claim of insignificance be tied to a stated adversary and stated assumptions, not asserted in the abstract. Paragraph 87's notion of "sufficiently high confidence" should be linked to this: confidence relative to what baseline is the operative question.

3.2 No Inference: "specific and meaningful" carries too much weight to be left informal

No Inference (paras 66–82) is the cleverest part of the framework and the most fragile. The line it draws, between an inference that reflects "the population at large" (allowed) and one that "relies on" the individual's own record (not allowed), is the right line. It is essentially the intuition behind differential privacy, and footnotes 43–44 reach for it by comparing outputs on *D* and *D′* (the dataset with the individual removed).

The gap is that the Guidelines describe this leave-one-out reasoning verbally and as one of three fallback options, without acknowledging that it is precisely the counterfactual that differential privacy formalises and bounds. As drafted, an assessor is asked to judge whether an inference "works in the same way" for *D* and *D′* (footnote 43) with no notion of how close is close enough. Two outputs are essentially never bit-identical; the whole question is the tolerance.

- Recommendation (a): state explicitly that where a formal guarantee such as differential privacy has been applied, the privacy parameter (ε, and where relevant δ) is directly relevant evidence for the No Inference assessment, and that a controller relying on such a mechanism should document the parameters and the composition across queries.
- Recommendation (b): make clear that the leave-one-out comparison in footnote 43 requires a stated tolerance "significantly different posterior belief" (footnote 44) needs an operationalisation, however qualitative, or it cannot be tested or audited.
- Recommendation (c): the treatment of composition is currently missing. Repeated queries, or the release of multiple aggregates from the same original data (acknowledged in para 77), degrade any inference guarantee. The Guidelines should say plainly that the No Inference assessment must consider the totality of releases, not each release in isolation.

3.3 No Linkage is stated as a binary but behaves as a probability

Paragraph 65 says the No Linkage criterion "is likely to be met if the information has not been recorded elsewhere." Para 80–81 then correctly acknowledges probabilistic and one-to-many linkage. These two passages sit uneasily together: linkage in practice is rarely a clean 1:1 join, and the Narayanan–Shmatikov result cited in footnote 39 is precisely about robust linkage under noise and sparsity. A criterion phrased as "has it been recorded elsewhere" invites a controller to answer "not to my knowledge" and stop.

Recommendation: reframe No Linkage explicitly in terms of the probability that a record can be matched to auxiliary data by a stated adversary, and state that the criterion is violated where that probability is non-insignificant even absent an exact join: consistent with para 80. The high-dimensional, sparse case (where near-uniqueness makes linkage easy) deserves an explicit worked example, since it is the empirically dominant failure mode for record-level data.

4. Encryption and the state of the art

Paragraph 91 is broadly correct: encryption limits access but is reversible and is not itself anonymisation, and one must ask whether an entity could decrypt "with means reasonably likely to be used, even if they do not currently have the decryption key (e.g. through cryptanalysis)." I suggest two refinements, both in the direction of precision.

- Harvest-now-decrypt-later. The reference to reasonably foreseeable technological developments (para 29(e)) should be connected explicitly to the quantum transition. Ciphertext that is secure today against a classical adversary but is captured and retained may be decrypted once a cryptographically-relevant quantum computer exists. For data with a long sensitivity lifetime, an assessment resting on the confidentiality of classically-encrypted data is already questionable. The Guidelines should name this scenario, consistent with the EU's own post-quantum migration roadmap.
- Distinguish encryption from privacy-preserving computation. The text treats encryption as a single access-control measure. It would be worth distinguishing conventional encryption from techniques whose security properties are qualitatively different, for instance secure multiparty computation or, at the formal end, mechanisms carrying a differential-privacy guarantee, since these bear on the criteria in different ways and should not be collapsed into "encryption is reversible." The PETs issue should not be dismissed in a spray and pray fashion… EUDIW is coming with a SNARK presented as a Zero-Knowledge proof, which will play with this blurry line between true anonymity, and pseudonymity who doesn't say its name.

I strongly support the encouragement in para 84 of "scientifically grounded expertise."[1] I would go further: the Guidelines could commit the EDPB, or national authorities, to maintaining a living reference on the state of the art of re-identification, so that the "state of the art" standard has an authoritative anchor rather than resting on each controller's literature review.

5. Points I endorse and would ask the EDPB to hold firm on

[1] As a cryptographer, this is really icky… Anonymous means anonymous… if information can be retrieved from the data, then it was never anonymous to begin with… As such, the call from being rooted in current scientific expertise means that every ad-hoc definition based on only "relative anonymity" is failing to meet this bar, the law/CJUE cannot redefine scientific fact.

Because parts of the consultation will push in the opposite direction, I want to be explicit about the positions I consider correct and hope the Board will maintain against pressure to soften them.

- Rejection of "lack of motivation" (para 31). Motivation is unobservable, unstable, and re-identification frequently occurs by accident, negligence or compulsion rather than intent. Keeping this out of the analysis is right.
- Contractual prohibitions are not legal prohibitions (para 34). Contracts can be revised, breached, or simply ignored, and they bind only their parties. Treating them as a complement to technical measures rather than a substitute is the correct hierarchy.
- Malicious and out-of-jurisdiction actors are in scope (paras 30, 32–33). The Guidelines are right that the assumption "people follow the law" is rebuttable, and that foreign intelligence services and cybercriminals are relevant entities where the evidence supports it. Example 8 is a good illustration.
- The whole-dataset rule for mixed data (para 36). Requiring insignificant re-identification risk for every individual, and treating a mixed dataset as personal where the parts are not separated, resists the temptation to average away the exposure of outliers: they are precisely the people most at risk!
- Time-dependence and periodic reassessment (para 35). Anonymity is not a permanent property. Building in reassessment is essential and well-judged. Just, be careful, that some occasion may precipitate the need for this reassessment.

6. Minor and drafting points

- Para 22 and the glossary use "insignificant" and "reasonably likely" as the operative thresholds but never place them on a common scale. A short paragraph relating the two would help, they are two sides of the same likelihood question and readers conflate them.
- Example 20 (para 87) is excellent and under-used: the point that a technique which succeeds against exactly one unknown individual can still be effective deserves to be stated as a general principle in the body text, not only in an example.
- The simplified/contextual distinction (paras 45–49) is clear, but the Guidelines should state once, prominently, that the simplified approach can only ever be more protective, never less (footnote 36 makes this point but it is easy to miss and is load-bearing).
- Para 92–93: the observation that agentic AI reduces the cost of gathering and combining public information is important and correct. Consider elevating it from a remark to an explicit instruction that assessments assume the near-term, not merely the current, cost of assembling auxiliary information.

7. Summary

The direction is right and I support it. My main worry is structural: the entity-relative definition treats the holder as fixed in two respects that are in fact moving. Its *identity* changes through acquisition, merger, insolvency and asset sale; the holder's *capability* grows with every AI model upgrade. Either shift can turn anonymous data into personal data with nothing happening to the data itself. None of this asks the Board to drop the contextual approach. It asks four things. Name transfer, acquisition and model upgrade as events that trigger reassessment. Treat a transfer or a capability gain that

makes data identifiable as a re-identification. Put the model-capability envelope in the anonymity assessment itself, bounded now and demonstrated later as interpretability allows. And hold the controller answerable for the re-identifying transfers and capability gains it could have foreseen.

Beyond that, several criteria are stated with less precision than the Guidelines demand of the controllers who apply them. The three criteria are the right ones. But “insignificant” likelihood, “specific and meaningful” inference, and probabilistic linkage each need enough operational content that two competent assessors land in the same place on the same dataset. Tie No Inference to the differential-privacy counterfactual it already implies, ask for a stated threat model and tolerance, deal with composition across releases, and name the harvest-now-decrypt-later problem, and the main gaps close without softening anything. I am happy to help the Board or its technical experts on any of this.

Prof. Olivier Blazy, École Polytechnique (LIX / GRACE), Institut Polytechnique de Paris.