Comment on Guidelines 02/2026 on Anonymisation

Submitted by: Meem Arafat Manab

Affiliation: Marie Skłodowska-Curie Doctoral Fellow, Ontology Engineering Group, Universidad Politécnica de Madrid (HARNESS Doctoral Network, Horizon Europe); Research Associate (pro bono), School of Religion, Theology, and Peace Studies, Trinity College Dublin

Consultation: Guidelines 02/2026 on Anonymisation, Version 1.0, adopted 07 July 2026

Summary

These comments concern a single structural asymmetry in the Guidelines and its consequences.

Paragraph 35 establishes a substantive rule of reversion: where the likelihood of re-identification rises to a level that is no longer insignificant, previously anonymous data "should again be considered personal and the controller will be accountable for any processing of that personal data," citing *Deutsche Wohnen* (C-807/21) on controller liability. The reclassification itself operates automatically, on the state of the world, without regard to whether the controller knows the state of the world has changed, and the substantive obligations that follow attach with it. (*Deutsche Wohnen* also holds that an administrative fine requires intentional or negligent conduct, so the sanction, unlike the reclassification, is fault-based; paragraph 76 of that judgment sets a low threshold, satisfied where the controller "could not have been unaware" of the infringing nature of its conduct.)

The same paragraph then makes the only mechanism by which a controller could detect that reversion, namely reassessment, a "good practice," something "wherever possible and as appropriate," and periodic. Footnote 3 goes further, providing that controllers who assessed a dataset as anonymous under Opinion 05/2014 before publication of these Guidelines "are not expected to conduct a new assessment", leaving only the same good-practice periodic reassessment that applies generally.

The result is a regime in which reclassification and the obligations that follow attach automatically to a state change the controller has no enforceable obligation to notice. Comments 1–3 below address this asymmetry, the specific mechanisms that make it acute rather than theoretical, and the near-total absence of independent verification anywhere in the Guidelines themselves.

Comment 1: Paragraph 35 imposes automatic reversion but only good-practice diligence for detecting it

Location: Paragraph 35; footnote 3 (to paragraph 2); paragraphs 29(e) and 42

Issue

Paragraph 35 correctly recognises that the likelihood of re-identification "typically increases over time," and correctly states the legal consequence: the data reverts to personal data, and the controller becomes accountable for processing it. Recital 26 GDPR supports this, and paragraph 29(e) carries its forward-looking terms into the assessment, listing among the objective factors "the available technology at the time of the processing, as well as reasonably foreseeable technological developments." The qualifier "reasonably foreseeable" is the Guidelines' addition; Recital 26 refers simply to "technological developments."

The difficulty is the mismatch in normative force between the two halves of paragraph 35:

- The reversion rule is stated in mandatory terms ("should again be considered personal"; "the controller will be accountable").
- The detection obligation is stated in waffled discretionary terms ("it is, therefore, good practice ... wherever possible and as appropriate ... to periodically reassess").

A controller who never reassesses is, on the face of these Guidelines, processing personal data unlawfully while having complied with everything the Guidelines actually require of it. Whether a fine follows turns on the *Deutsche Wohnen* fault standard, but that standard is not demanding, and a controller relying on an unreviewed assessment after a publicly documented capability advance is not obviously one that "could not have been unaware." That is not a stable allocation of responsibility, and it is likely to produce exactly the wrong incentive: because reassessment is optional but its findings are liability-triggering, **a rational controller is discouraged from looking**.

Footnote 3 delves much, much farther. Controllers holding datasets assessed under Opinion 05/2014 "are not expected to conduct a new assessment," notwithstanding that paragraph 2 of the Guidelines expressly justifies this update on the basis of "major changes in the legal, privacy, engineering and technological landscapes" since 2014. The grandfather clause thus relieves reassessment of those datasets that had been assessed against a twelve-year-old technical standard, while paragraph 35's reversion rule continues to apply to them in full. But it is these datasets that are at the highest risk of reversal after twelve years have already passed.

Paragraph 42 comes closest to recognising that some circumstances call for reassessment outside the periodic cycle: it provides that a security incident "may lead to a reassessment of anonymity", particularly where the assessment relied on certain information being kept confidential. The logic of paragraph 42 is not specific to breaches. It is the general logic that a discrete, identifiable event can invalidate an anonymity determination, and that when such an event occurs there is reason to act on it rather than wait for the next scheduled review. Paragraph 42 is itself permissive, but absolutely nothing in its reasoning confines that logic to internal security incidents.

Proposed amendment

Align the normative force of the two halves of paragraph 35. Either:

(a) Recast the reassessment obligation in mandatory terms proportionate to the reversion rule, an ongoing obligation to maintain anonymity for as long as the data is retained, rather than a periodic confirmation obtained at some past date; or

(b) At minimum, extend the immediate-reassessment trigger in paragraph 42 beyond security incidents to documented external capability advances, and state that reliance on a prior assessment is not reasonable where such an advance has been publicly documented, and the controller has not reassessed.

Also, please reconsider footnote 3, or qualify it so that the grandfather clause does not extend indefinitely to datasets whose assessment predates the technical developments these Guidelines were written to address.

Comment 2: Three distinct mechanisms erode anonymity on different timescales; periodic review addresses none of them well

Location: Paragraphs 29(c), 29(e), 30, 35, 83, 88, 91, 92, 93

Issue

The Guidelines treat technological changes as a single undifferentiated factor. In practice, at least three mechanisms operate, with materially different time signatures, and a calendar-based review interval is poorly matched to all three.

(a) Capability drift, including agentic AI

The Guidelines already recognise this. Paragraph 92 notes that most currently documented re-identification techniques "are becoming increasingly accessible, require limited resources and can be performed within a reasonable timeframe, even on commodity hardware," and that "developments in AI, and especially agentic AI, will likely further reduce the time and costs necessary to access and deploy these techniques." Paragraph 93 adds that AI may reduce the time associated with gathering and combining information from multiple sources, at least where those sources are publicly available. Paragraph 83 contemplates "special-purpose AI agents" deploying sophisticated probabilistic techniques.

This is a candid (and welcome) acknowledgment. But it is difficult to reconcile with paragraph 35's periodic model. If cost and time barriers are falling continuously and unpredictably, the interval between reviews becomes the period of undetected exposure, and the Guidelines offer no basis for setting it. Paragraph 88's recommendation that contextual analyses incorporate "adequate safety margins ensuring that the assessment remains robust as the entities'

capabilities evolve" is the right instinct, but it is a recommendation, it is confined to the contextual approach, and it gives no guidance on how large a margin is adequate against a capability curve the Guidelines themselves decline to characterise.

(b) Scheduled cryptographic obsolescence

Paragraph 91 is the closest the Guidelines come to this, and it is close: it states that "it is, therefore, important to assess" if entities with access to the data "could still decrypt it with means reasonably likely to be used, even if they do not currently have the decryption key (e.g. through cryptanalysis)." Paragraph 30 identifies domestic and foreign intelligence agencies as relevant entities with "high computational power."

What paragraph 91 does not do is distinguish cryptanalysis-as-gradual-erosion from cryptanalysis-as-scheduled-discontinuity. Anonymisation and pseudonymisation techniques resting on computational hardness assumptions do not degrade smoothly; they fail categorically at a threshold. This is not speculative; it is the premise on which NIST published FIPS 203, 204 and 205 on quantum-resistant algorithms and on which public post-quantum migration timelines are built. It is also retroactive in effect: data retained today under an assessment made against classical capability is already exposed to a future break, since an adversary who obtains the data now can decrypt it later. A periodic review conducted before the threshold returns "anonymous"; a review conducted after returns "personal"; and the data was, throughout, the same data.

(c) Growth of the external data ecosystem

Paragraph 29(c) requires assessment of "any additional information that would allow for the identification of individuals, and how likely it is that such information can be obtained," and paragraph 35 notes "the increased availability of additional information." But the Guidelines assess anonymity dataset-by-dataset and controller-by-controller. The auxiliary data that defeats an anonymisation is frequently published by unrelated third parties over whom the assessing controller has no visibility whatsoever, the mechanism documented in the literature the Guidelines themselves cite at footnote 39 (Narayanan & Shmatikov). No amount of diligent self-review by a controller can detect a risk that materialises entirely outside its field of view. This mechanism does not require any new technology either; it requires only time and an ordinarily growing public data ecosystem. New lawful access routes and data-sharing mandates have the same effect without any technological change: an instrument that compels disclosure of a dataset expands, at a stroke, the additional information available to every entity eligible to receive it.

Proposed amendment

Address these as analytically distinct components of the "reasonably foreseeable technological developments" factor at paragraph 29(e), since each requires a different monitoring posture:

- For (a), guidance on setting safety margins under paragraph 88, extended beyond the contextual approach.
- For (b), a requirement that controllers identify which of their anonymisation or access-restriction measures depend on cryptographic hardness assumptions, and plan

reassessment against published post-quantum migration milestones rather than an internal calendar. Paragraph 91 is the natural home for this.

- For (c), acknowledgment that this mechanism is structurally undetectable by individual controllers acting alone, and consideration of whether it warrants a collective or supervisory-authority-level monitoring function rather than being left entirely to per-controller review.

Comment 3: The Guidelines provide for no independent verification at any point

Location: Paragraphs 22, 34, 41 and 90; footnote 36 (to paragraph 48); footnote 29 (to paragraph 34)

Issue

Paragraph 41 provides that controllers "should ensure adequate documentation" of the anonymisation processing, including the testing of the supposedly anonymous datasets, so that it is possible to "demonstrate both the GDPR-compliance of the anonymisation process, and the effectiveness of the anonymisation itself." Nothing requires that documentation, or the underlying determination, to be reviewed or tested by any party other than the controller that produced it.

This is compounded by the asymmetry in footnote 36. The simplified approach may be used by an anonymising controller to establish that data is anonymous and outside GDPR scope, but "cannot be used by third parties to definitively prove that the information is personal data." The framework thus runs in one direction only: it can establish anonymity for the party asserting it, but cannot be turned around by those affected. Combined with the absence of any audit requirement, a controller's own assessment is effectively unreviewable in practice.

A live illustration

The Commission's parallel specification proceeding concerning Google Search data sharing under the Digital Markets Act (Regulation (EU) 2022/1925, "DMA") is instructive on the limits of point-in-time expert validation, and is offered here on that basis only.

In its preliminary findings of 16 April 2026 (Case DMA.100209), the Commission specified the anonymisation measures Alphabet must apply before sharing search data — technical measures complemented by contractual ones — set out, per the published case summary, "taking into account the Joint Guidelines of the Commission and of the European Data Protection Board on the interplay between the DMA and the [GDPR]". The Commission subsequently described the resulting method as developed "in close collaboration with internal and external privacy experts" and in line with those draft Joint Guidelines (press release IP/26/1634, 16 July 2026).

In early May 2026, Google's Sergei Vassilvitskii stated in written comments to Reuters that Google's internal AI red team "managed to re-identify users in less than two hours". Several caveats are essential. The red team's findings have not been published; the claim is available only through media statements and should be treated with corresponding caution. Subsequent commentary has characterised the exercise as relying on linkage attacks assisted by AI tools, but that characterisation does not come from the statement itself. The claim was made by a party with a direct commercial interest in resisting the data-sharing obligation, and other stakeholders have characterised it as such. The Commission proceeded to adopt its decision on 16 July 2026.

What the episode illustrates, independent of whether the specific claim withstands scrutiny, is something structural. An expert-developed anonymisation method, produced with explicit reference to EDPB-aligned guidance, became the subject of a credible contemporaneous re-identification claim. Under the Guidelines as drafted, an equivalent controller would have satisfied paragraph 41 in full by documenting its process, and nothing in the Guidelines would have surfaced the challenge. It surfaced here only because an unusually well-resourced adversarial party had both the capability and the incentive to test the method and publicise the result.

The Commission has since stated that, "[d]epending on future market developments, including based on independent third-party evaluation", it may amend the decision "in particular as regards the relevant anonymisation measures" (IP/26/1634). The case summary also records that beneficiaries are subject to an independent verification mechanism: an independent audit before access and annual monitoring audits, with the scope of the assurance engagement defined by the Commission. Those audits verify recipients' compliance with the contractual terms rather than testing the anonymisation method itself. But, taken together with the amendment clause, they give that architecture a revisability mechanism tied to external evaluation. The Guidelines contain no equivalent of either.

A related divergence is worth flagging. Paragraph 34 states clearly that a contractual prohibition "should not be treated as a prohibition by law," and paragraph 90 that technical, organisational and contractual measures "are not by themselves sufficient to make data anonymous." The DMA case summary describes that division of labour, and assigns the decisive step to the contractual layer: technical measures reduce the likelihood of re-identification "to a residual level", while contractual measures, such as prohibition on re-identification attempts, or linkage with auxiliary datasets, or retention beyond 13 months and onward sharing, "further reduce the likelihood of re-identification of end users to an insignificant level". Insignificance is the legal threshold, as paragraph 22 of these Guidelines confirms. Paragraphs 34 and 90 therefore appear to be in tension with an architecture in which contractual terms are what carry residual risk across that threshold. Given that footnote 29 cites the Joint DMA/GDPR Guidelines, the EDPB may wish to address that interaction expressly.

Proposed amendment

Please provide that for high-risk anonymisation of large-scale datasets, special category data, or data intended for release to multiple or unspecified recipients, documentation under paragraph

41 should be subject to independent adversarial testing, and that the results of such testing form part of the record.

Please reconsider the one-way character of footnote 36, or explain why an assessment framework should be available to establish anonymity but not to contest it.

Finally, please address expressly whether, and to what extent, contractual measures may contribute to a finding of anonymity, given the tension between paragraphs 34 and 90 and regulatory approaches that rely on them.

Closing note

The Guidelines are a significant and careful improvement on Opinion 05/2014, and paragraph 35's recognition that anonymity can lapse is among their most important contributions. The concern raised here is not that the EDPB has misunderstood the problem, because paragraphs 35, 88, 91 and 92 show clearly that it has not. Our concern is that the operative obligations do not yet match the diagnosis the Guidelines themselves set out. A rule that data reverts to personal data automatically, coupled with a duty to notice that is discretionary and no verification mechanism within the Guidelines themselves, places the entire burden of a continuous problem on a periodic solution. We cannot periodically evaluate a system subject to changes whenever a new technology or a new legal instrument arrives, because they do not arrive on the reviewer's schedule. The evaluation must be continuous.