Response to EDPB Guidelines 02/2026 on Anonymisation, 30 October 2026

General observations

The three-criteria framework (No Record Isolation, No Linkage, No Inference) provides a practical, structured approach for anonymisation assessments. The contextual and simplified approaches are a sensible addition, giving controllers flexibility without lowering the standard.

Two aspects deserve particular recognition. The acknowledgement that anonymity is perspective-dependent reflects the reality that the same dataset may be personal data for one entity and anonymous for another. The guidelines' contextual approach is consistent with the clarification given by the CJEU in *EDPS v SRB* (C-413/23 P). The CJEU's reasoning in *Breyer* (C-582/14), that data which is not directly identifying may nonetheless constitute personal data where the controller has legal means to obtain additional information, is directly on point and frequently misunderstood in practice. The statement that controllers should not use misleading terminology such as "anonymous" or "de-personalised" when data remains identifiable addresses a problem we see regularly.

These guidelines do not carry forward the technique-specific assessment from the Article 29 Working Party's Opinion 05/2014. Given the techniques and the threat landscape have evolved considerably since 2014, we invite the Board to update this assessment to reflect the state of the art.

Our observations below focus on areas where the guidelines could be strengthened, particularly as they apply to personal data within AI systems.

The No Inference criterion and generative AI

The No Inference criterion is the most consequential addition to the framework, and the worked examples are highly useful. We invite the Board to consider whether this treatment is sufficient for generative AI.

A large language model trained on data assessed as anonymous under the three criteria may nonetheless produce outputs that relate to identifiable individuals, because the model has learned patterns across billions of parameters that, combined with a well-crafted prompt, can reconstruct or approximate personal data never explicitly present in any single training record. Large language models can also memorise and reproduce verbatim sequences from training data, including names and addresses; where a model regurgitates identifying information, the No Inference criterion has failed in a concrete and verifiable way. We invite the Board to require controllers to test for memorisation and regurgitation as part of the No Inference assessment, document the results, and implement mitigations where memorisation is detected.

The guidelines' acknowledgement that AI models may enable inferences representing "correlations between elements of the original information, rather than that information itself" is important: an anonymisation assessment conducted at the point of training may not hold at the point of deployment, or after downstream modification by a third party. We invite the Board to provide explicit guidance that the No Inference criterion must be assessed against the reasonably foreseeable capabilities of systems that will process the data, and revisited when the deployment context changes materially.

Re-identification risk and dataset combination

The guidelines acknowledge that re-identification risk increases over time as techniques become cheaper and more accessible. Two points follow from this.

The emergence of agentic AI systems compounds this. An autonomous agent capable of chaining searches, cross-referencing databases and iterating on partial matches can execute re-identification strategies that would be disproportionately costly for a human analyst. Techniques excluded from the "means reasonably likely" assessment because of their cost may become routine when delegated to an agent operating at marginal cost. The guidelines could usefully acknowledge that agentic AI shifts the "reasonably likely" threshold and require controllers to factor autonomous-agent capabilities into their assessments.

The exclusion of unlawful means from the "means reasonably likely" assessment is sound as a starting position, but we suggest this should be rebuttable. Where specific evidence shows that unlawful means are realistic rather than speculative, for example where the data type has a documented history of targeted attacks, the controller cannot rely on the classification of those means as "unlawful" to treat the data as anonymous.

The risk profile of a dataset changing materially when it is combined with other datasets, is particularly acute in AI supply chains; where a model provider may combine training data from multiple sources, each anonymised independently. The entity performing the combination may have access to a volume and diversity of data that no individual source could have foreseen. We invite the Board to include guidance on situations in which data is anonymised for release into a supply chain or for training a general-purpose AI model, that the anonymisation assessment must account for the reasonably foreseeable combinations of data the receiving entity will have access to. Each entity in the data supply chain should be responsible for conducting its own anonymisation assessment from its own perspective.

Periodic reassessment

The guidelines frame periodic reassessment as "good practice." We respectfully submit that periodic reassessment is a requirement of the Article 5(2) accountability principle and the Article 32 obligation to ensure ongoing security and confidentiality of processing. Where the re-identification landscape shifts and data assessed as anonymous reverts to personal data, detecting that shift is an accountability obligation, not a discretionary exercise. In our experience, controllers who anonymised data years ago may not choose to revisit the assessment absent an explicit regulatory trigger, while the technological landscape shifts beneath them.

We invite the Board to present reassessment as a mandatory component of the accountability framework and to specify triggers: combination with new datasets; advances in re-identification techniques relevant to the data type; incorporation into a new processing context (such as AI model training) not envisaged in the original assessment; and the passage of a defined period since the last assessment.

The inference problem and special categories

We invite the Board to address the risk that inferences drawn from non-special-category data may constitute special-category data under Article 9. A dataset containing browsing history, purchase records or location data may yield inferences about health, religious beliefs, political opinions or sexual orientation. The CJEU's ruling in *OT v Vyriausioji tarnybinė etikos komisija* (C-184/20) held that data which is liable indirectly to reveal sensitive information, following an intellectual operation involving deduction or cross-referencing, falls within Article 9; the earlier ruling in *GC & Others* (C-136/17) supports a broad reading of "revealing" in the same provision.

Where an AI model is trained on data assessed as anonymous but its outputs enable inferences that stray into special categories, the original assessment may have been insufficient: not because the three criteria were poorly applied to the data as given, but because the assessment did not account for such inferences at a time before they were in reasonable contemplation. The No Inference criterion, when applied to data intended for AI model training, should specifically assess whether the model could produce outputs that constitute or reveal special-category data.

Co-ordination with the EU AI Act

The guidelines make no reference to the EU AI Act (Regulation 2024/1689). While we appreciate that the EDPB does not hold jurisdiction over the EU AI Act, we invite the Board to clarify, within its GDPR mandate, which entity bears the obligation to reassess anonymity when data crosses the GDPR-AI Act boundary, and to recommend co-ordination with the AI Office. Without that co-ordination, an upstream entity's assertion that data is anonymous risks functioning as an unchallenged compliance assertion: placing the data outside the GDPR's scope while downstream providers inherit undetected liabilities when re-identification becomes possible.

The Board's parallel Guidelines 03/2026 on web scraping address related controller obligations; cross-references between the two sets of guidance would similarly assist controllers.

Summary of recommendations

Theme	Recommendation
Generative AI and the No Inference criterion	• Require controllers to test for memorisation and regurgitation as part of the No Inference assessment • Assess the criterion against the reasonably foreseeable capabilities of systems that will process the data; revisit when the deployment context changes • Address the risk that inferences from non-special-category data may produce special-category data under Article 9
Re-identification risk	• Acknowledge that agentic AI shifts the "reasonably likely" threshold • Clarify that the exclusion of unlawful means operates as a rebuttable presumption • Require supply-chain anonymisation assessments to account for foreseeable data combinations; confirm each entity must conduct its own assessment
Periodic reassessment	• Present reassessment as mandatory under the accountability framework, not discretionary good practice • Specify triggers: new dataset combinations; relevant advances in re-identification techniques; new processing contexts; passage of a defined period
Technique-specific guidance	• Update to at least the level of WP29 Opinion 05/2014, covering current techniques and the evolved threat landscape
Co-ordination with the EU AI Act	• Clarify which entity bears the obligation to reassess anonymity at the GDPR–AI Act boundary • Recommend co-ordination with the AI Office • Cross-reference the parallel Guidelines 03/2026 on web scraping

We are available to discuss any of these observations with the Board or its working parties.

About the respondent

Spriggun Ltd is a UK-based advisory firm specialising in AI governance, data protection and cyber risk. This response draws on its principals' experience leading AI, privacy and data governance inside global businesses, and advising boards and executive teams on the GDPR, the EU AI Act, NIS2 and related frameworks. Spriggun welcomes the opportunity to contribute to this consultation.