European Digital Rights

Contribution to the public consultation on EDPB Guidelines 02/2026 on Anonymisation

European Digital Rights (EDRi) welcomes the Version 1.0 of the EDPB's Guidelines 02/2026 on Anonymisation[1]. We have waited a long time for updated guidance in this area, and it comes at a rather sensitive moment. The definition of personal data, what it means for someone to be identifiable, and when data can really be said to be anonymous are now being argued about across several EU files and the reality of GDPR enforcement, sometimes with very direct consequences for the scope of the GDPR and, with it, the protection of fundamental rights online. The Guidelines themselves recognise that the legal and technological landscape has moved considerably since Article 29 Working Party's 2014 Opinion 05/2014 (the 2014 Opinion)[2], including because of CJEU case law, EU data spaces and AI and technology in general (para 2).

Our reading of the Guidelines is broadly positive. A lot of the approach from the 2014 Opinion is still there. The assessment remains contextual rather than purely formal, singling out still matters, and the Guidelines look beyond the data held in one particular database at one particular moment. They also take seriously other actors (para 30), changes in technology, security failures and the fact that contractual restrictions are not magic (paras 34 and 90). All of that is useful, especially given the direction of some of the current political debate. **We do, however, think that some of the new qualifications could benefit from more clarification, as detailed hereunder.**

Executive summary and main recommendations

The concept of personal data sits right at the entrance of the GDPR. If information is considered anonymous, the rest of the Regulation does not get a chance to operate. This is why the debate is much bigger than what some are presenting nowadays as 'anonymisation techniques.' The way 'identified or identifiable' is understood affects profiling, data subject rights, accountability, and in some cases the safeguards around automated decision-making too. If the threshold is drawn too narrowly, we can lose the legal protections before we even get to questions such as fairness, purpose limitation or lawfulness. Our Recommendations are the following:

- **Keep the Recital 26 test broad and grounded in the real processing contex**t. The final Guidelines should continue looking at objective factors, additional information, technological capabilities, onward transfers, chains of means and reasonably foreseeable unauthorised access (paras 20–22, 28–35 and 88–93);
- **Keep anonymisation and pseudonymisation clearly separate:** peudonymisation is a safeguard applied to personal data which does not become anonymisation merely because a particular recipient lacks the usual attribution key. The final Guidelines should say this explicitly and cross-reference the forthcoming final Guidelines 01/2025 on Pseudonymisation[3] (paras 37, 50 and 90–91);
- **Explain the relationship between record isolation and singling out more carefully:** a unique record does not always mean that a person is identifiable, as paras 98-100 and Example 24 rightly show. At the same time, the singling-out analysis should be read together with paras 19 and 25, Example 6, para 50 and para 87: persistent recognition and differential treatment can identify a person even where nobody discovers their civil identity;

1 European Data Protection Board (EDPB), Guidelines 02/2026 on Anonymisation Version 1.0 (adopted on 07 July 2026) https://www.edpb.europa.eu/system/files_en?file=2026-09/edpb_guidelines_202602_anonymisation_v1_en_0.pdf

2 Article 29 Data Protection Working Party, WP216 Opinion 05/2014 on Anonymisation Techniques (adopted on 10 April 2014) https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2014/wp216_en.pdf

3 EDPB, Guidelines 01/2025 on Pseudonymisation Version 1.0 https://www.edpb.europa.eu/public-consultations/guidelines-012025-on-pseudonymisation_en, paras 16–23

- **Deal with AI and its implications more explicitly**, rather than leaving the relevant points scattered across the text. The Guidelines already address extraction from supposedly anonymous AI models, membership inference, prompting, AI-assisted re-identification, high-dimensional data and the effect of AI on time and cost (paras 2, 77, 82–83 and 92–93). It would help to bring these points together and spell out what they mean for supposedly anonymised training data and for AI models themselves;
- **Be much clearer about the limits of CJEU's ruling *EDPS v SRB*, which should not be read as creating a general recipient-only test, and certainly not as making pseudonymisation an easy way out of the GDPR.** The Guidelines already contain important limits on an over-broad recipient-only reading of the judgment: para 14 and Example 2, para 15, and paras 20–21 and 28–30 are especially useful. The rather specific facts of *SRB* matter, as does the rest of the Court's case law[4], and the final version should make those safeguards harder to miss when paras 6–8 on entity-specific identifiability are read or quoted on their own;
- **Keep the approach coherent across EU law**: the Guidelines should avoid language that can be lifted out and used to justify the very broad reading of *SRB* found in parts of the so-called 'Digital Omnibus on Data' debate[5]. They could also benefit from explaining how the general GDPR standard relates to the anonymisation arrangements now being developed under Article 6(11) DMA for Google Search data[6];
- **Make anonymity claims properly demonstrable and revisitable:** if an organisation wants to rely on the conclusion that data are anonymous, it should be able to show why with actual analysis, documentation and testing. That assessment must remain open to supervisory review and to reassessment when technology, access conditions or available data change (paras 35, 40–42, 88 and 95–100).

1. Identifiability is a central boundary of EU data protection law

Article 4(1) GDPR defines personal data as information relating to an identified or identifiable natural person. Recital 26 then tells us to look at all the means reasonably likely to be used, including singling out, by the controller or by another person. The Guidelines rightly frame this as a threshold question and make clear that the answer can depend on the relevant entity and context (paras 6–10). Once information is genuinely anonymous, processing of that information as such falls outside the GDPR. So the identifiability test decides, in practice, whether the protections of EU data protection law apply in the first place.

---

4 Court of Justice, Case C-582/14, Patrick Breyer v Bundesrepublik Deutschland, EU:C:2016:779, in particular paras 43–46; Case C-434/16, Peter Nowak v Data Protection Commissioner, EU:C:2017:994, in particular paras 34–35; Case C-319/22, Gesamtverband Autoteile-Handel eV v Scania CV AB, EU:C:2023:837, in particular paras 46–49; Case C-604/22, IAB Europe v Gegevensbeschermingsautoriteit, EU:C:2024:214, in particular paras 43–48; Case C-479/22 P, OC v European Commission, EU:C:2024:215, in particular paras 51–64; and Case C-413/23 P, European Data Protection Supervisor v Single Resolution Board, EU:C:2025:645, in particular paras 74–86.

5 European Commission, Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL amending Regulations (EU) 2016/679, (EU) 2018/1724, (EU) 2018/1725, (EU) 2023/2854 and Directives 2002/58/EC, (EU) 2022/2555 and (EU) 2022/2557 as regards the simplification of the digital legislative framework, and repealing Regulations (EU) 2018/1807, (EU) 2019/1150, (EU) 2022/868, and Directive (EU) 2019/1024 (Digital Omnibus) https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52025PC0837; see also EDRi, Reopening GDPR and ePrivacy through the Digital Omnibus: a risky path for EU digital rights https://edri.org/our-work/reopening-gdpr-and-eprivacy-through-the-digital-omnibus-a-risky-path-for-eu-digital-rights/

6 European Commission, Decision pursuant to Article 8 of Regulation (EU) 2022/1925 specifying measures to be implemented by Alphabet in relation to Article 6(11) DMA, Case DMA.100209 – Alphabet – Article 6(11), adopted 16 July 2026 https://digital-markets-act-cases.ec.europa.eu/cases/DMA.100209 ; see also Sophie Stalla-Bourdillon, Cracking the Anonymisation Test: Early Lessons following EDPS v SRB, March 2026 https://www.scl.org/cracking-the-anonymisation-test-early-lessons-following-edps-v-srb/ ; or Sophie Stalla-Bourdillon, Alfred Rossi, Anonymisation in motion: from data protection to digital markets law, Law, Science, Technology & Society https://researchportal.vub.be/en/publications/anonymisation-in-motion-from-data-protection-to-digital-markets-l/

That is also why identifiability cannot be reduced to the question 'can somebody find the person's name?'[7]. **The Guidelines are actually quite strong on this. Paragraph 19 defines identification as distinguishing a person from others within a given context in a way that makes different treatment possible.** Paragraphs 23–25 add that identifiers can include device identifiers, fingerprints and combinations of attributes. Example 6 then applies that logic directly to targeted advertising: a persistent pseudonym based on a device fingerprint can identify a person because it allows the same individual to be distinguished, profiled and treated differently across sites, even without a name or email address.

This topic is critical for the current political economy, because of the centrality of profiling and automated decision-making. Article 4(4) GDPR defines profiling by reference to automated processing of personal data, while Article 22 provides safeguards for certain decisions based solely on automated processing, including profiling. If a persistent individual profile were treated as non-personal simply because the controller could not attach a civil name to it, those safeguards might never even come into play. That would be particularly difficult to justify where the processing is designed precisely to distinguish between people and change how each one is treated.

We therefore think **the final Guidelines should preserve, and perhaps make more visible, the functional understanding already present in paras 16, 19, 25, 50 and 87: the analysis is not only about discovering who somebody 'really is'.** It also asks whether the actor can distinguish that person in the relevant context and use information relating to them in a way that allows them to be recognised or treated differently. Profiling, ranking, targeting or exclusion are obvious settings where this can matter.

2. The Guidelines largely preserve the WP29 approach, while refining it
We are glad that the Guidelines do not throw away the approach developed in the 2014 Opinion. The 2014 Opinion looked at singling out, linkability and inference as the central risks to test. The new Guidelines refine that architecture into No Record Isolation, No Linkage and No Inference, while retaining a separate legal assessment of identifiability (see paras 2–3, 52 and 95–103). They also keep the basic point that no particular technique comes with an automatic legal guarantee: context and the means reasonably likely to be used still matter.

There is one point worth correcting from some of the discussion around the 2014 Opinion. The Opinion did not say that failing any one of the three criteria automatically made the data personal: it treated satisfaction of all three as a strong indication of robustness, while failure on one criterion required a proper assessment of the remaining identification risk. The new Guidelines make that structure even clearer: para 52 says expressly that violation of a criterion does not necessarily make the data personal, and paras 95–103 set out the further analysis.

A number of the additions are genuinely helpful for understanding this topic. For starters, Para 30 recognises that relevant entities can include receiving entities, rogue employees, investigative journalists, intelligence authorities, unethical companies and cybercriminals. In addition to that, Para 31 is sensibly sceptical of arguments based only on a supposed lack of motivation to identify somebody. Paras 32–34 then distinguish legal prohibitions from contractual restrictions: a binding contract may affect what means are reasonably likely to be used, but it is not the same as a prohibition by law and should complement technical protection rather than replace it. We also welcome the more careful treatment in paras 32–33 of unlawful means. *Breyer* should not be reduced to a rule that illegal means simply do not count. Where a legal prohibition does not meaningfully constrain the actors concerned, or there is concrete evidence that it may be

7 See for instance Frederik J. Zuiderveen Borgesius, 'Singling out people without knowing their names – Behavioural targeting, pseudonymous data, and the new Data Protection Regulation' (2016) 32(2) Computer Law & Security Review 256–271; and Douwe Korff, The Concepts of "Data Subject", "Personal Data", "Pseudonymous Data" and "Anonymous Data" in the EU GDPR, 11 November 2025.

ignored or circumvented, unlawful access and re-identification can still be part of the realistic Recital 26 assessment.

The same goes for time, as Para 35 notes that the likelihood of re-identification will often grow as technology develops and more information becomes available. As pointed out in para 29, future technological developments should be taken into account for the assessment of whether the data can reasonably be regarded as anonymous today. We recommend to include the (expected) period of processing for the anonymised data as an explicit factor to consider when exploring possible future technological developments that could enable re-identification. If the data is made available online, this period will generally be very long indeed. Recital 26 is therefore not only asking whether identification is feasible at the instant of release. Reasonably foreseeable developments during the expected lifetime of the data matter to the conclusion reached today. Paras 88 and 92–93 return to the same point in the technical analysis, including the effect of AI on the time and cost of re-identification. **And where assumptions later change in ways that were not reasonably foreseeable, the anonymity assessment should be revisited rather than treated as permanent.**

3. Entity-specific assessment must not fragment responsibility across processing chains
We support the safeguard in para 15 for processors. If a processor handles information on behalf of a controller for whom those data are personal, the assessment uses the controller's perspective. The processor should not be able to treat the same information as outside the GDPR simply because it does not itself possess the identifying information. Otherwise outsourcing would become a very easy way to hollow out protection (and it already is). It would be better to emphasise that the CJEU took the position in the SRB-Case that the data transferred to a third-party was, from the controller's perspective still personal data and that even if from a service provider's perspective the data does not constitute personal data, this does not affect the data subject's obligations under Article 28 GDPR.

But this problem goes beyond processors. Modern processing chains are messy: adtech, cloud services, data brokers, AI supply chains and data-sharing arrangements often distribute information and capabilities between several actors. No single company may hold every piece needed to identify someone, while the overall system is still organised around persistent recognition, enrichment and recombination. The Guidelines recognise this in several places: para 9 refers to chains of entities, para 13 to relationships between parties, paras 20–21 to indirect identification and transfers, and paras 28–30 to chains of means and a broad range of relevant actors.

The final Guidelines should preserve that real-world approach. **The assessment should include foreseeable transfers, processors and sub-processors, public or commercially available datasets, technical interoperability, contractual access and realistic chains through which data can be combined.** Para 29 already requires the processing and release context, additional information, time, cost and foreseeable technology to be considered; para 90 adds that access restrictions have to be assessed holistically. If a system is in practice organised around singling out, linkage or inference across actors, the analysis should not reduce it to a set of artificial organisational islands.

Recent evidence from the advertising ecosystem makes this concern particularly concrete. Identity-resolution has become a commercial service in its own right: data brokers and advertising intermediaries maintain identity graphs whose purpose is precisely to connect pseudonymous behavioural, transactional and device-level signals across services and contexts. This matters for the Recital 26 assessment, because where the processing environment contains actors whose business model is to perform that linkage, the availability of those capabilities, the

commercial incentives to use them and the relevant data flows cannot be abstracted away from the assessment of the means reasonably likely to be used[8].

Recent measurements of major conversational AI assistants make this point rather concrete. Chat URLs and conversation identifiers, AI-generated titles, prompt content, account or advertising identifiers and hashed email addresses were observed being transmitted to third parties, in some cases alongside persistent tracking identifiers. The relevant identifiability assessment is therefore about the combination of data and the recipient's actual means, rather than whether each field looks identifying when considered on its own[9].

4. Anonymisation and pseudonymisation need a clearer boundary

This is one place where we would like the final text to be more explicit. Anonymisation and pseudonymisation do different legal work: pseudonymisation under Article 4(5) GDPR is a way of processing personal data and a safeguard for reducing risk, and it deliberately preserves the ability to link records relating to the same person and, in authorised circumstances, to attribute those data back to them. The consultation version of Guidelines 01/2025 on Pseudonymisation is clear on this: 'prevention of attribution does not imply anonymity', and paras 21–23 explain that pseudonymised data remain personal where they can be combined with additional information through means reasonably likely to be used. Even after the controller deletes the additional information it retained, the data become anonymous only if the separate conditions for anonymity are actually met.

Guidelines 02/2026 are, in substance, largely consistent with that approach. Para 37 says that data which are anonymous for a recipient can still remain personal for the anonymising controller. Para 50 says that data are plainly not anonymous where they contain pseudonyms that can be reverse engineered or where the pseudonym itself can be used to identify the individual. Para 90 says that technical, organisational and contractual access controls are not, by themselves, enough to make data anonymous. Para 91 makes a similar conceptual point for encryption: a measure that restricts access but is designed to be reversible is different from anonymisation.

Still, **the entity-specific language in paras 6–8 makes the boundary worth spelling out in one place.** The fact that a recipient does not possess the usual 'key', lookup table or direct identifier is not sufficient to establish anonymity, as is also stated by the CJEU in its case law. A separate Article 4(1) and Recital 26 assessment is still needed, including whether the pseudonym or other attributes themselves allow the person to be distinguished and treated differently (paras 19, 25, Example 6 and para 50), whether additional information or methods are available through relevant chains of actors (paras 20–21 and 28–30), and whether access restrictions really make those means unlikely in practice (paras 32–34 and 89–91).

This distinction is especially important in sectors where pseudonymous identifiers are deliberately retained because their commercial value lies in continued linkage. Calling such identifiers 'anonymous' does not change the fact that the same individual can remain recognisable across interactions, nor that specialised actors may exist precisely to connect those identifiers to other information[10].

We would therefore suggest an explicit cross-reference between these Guidelines and the final Guidelines 01/2025 on Pseudonymisation once those are adopted. At the time of this consultation, Guidelines 01/2025 are still published as the version for public consultation. Whichever text is finalised first, the two documents should use compatible language and make

8 Helena Malikova and Johnny Ryan, 'Re-identification-Based Advertising', TechREG Chronicle, July 2026 https://www.pymnts.com/cpi-posts/re-identification-based-advertising/

9 Oliveira et al., 2026, Prompt like a Butterfly, Sting like a Tracker: A Privacy Analysis of Web and Mobile Conversational AI Agents

10 Helena Malikova and Johnny Ryan, 'Re-identification-Based Advertising', TechREG Chronicle, July 2026 https://www.pymnts.com/cpi-posts/re-identification-based-advertising/

clear that pseudonymisation is a safeguard applied to personal data, not a lesser form of anonymisation or a shortcut to treating data as non-personal. A separate Recital 26 assessment is still needed before that conclusion can be reached.

5. No Record Isolation and singling out: a useful distinction that needs a clearer boundary
One of the most delicate conceptual refinements would be the separation between No Record Isolation and singling out. Paras 54–57 treat No Record Isolation as a technical criterion, while Paras 98–100 then say, quite rightly, that failure of that criterion does not automatically mean that the person is identifiable. The controller still has to ask whether the isolated record allows the individual to be singled out. Example 24 explains why: a unique combination of answers in a genuinely secret survey may violate No Record Isolation while still leaving no reasonably likely way to link the answers to the respondent or treat them differently.

We agree with that distinction: uniqueness inside a dataset and legal identifiability are not the same thing. A record can be unique while, in the circumstances, providing no realistic route to distinguish the person in the relevant context. This is also consistent with para 52, which makes clear that violation of one of the three technical criteria triggers further analysis rather than an automatic conclusion that the data are personal.

Our concern is narrower and is mainly about how para 98 may read if it is lifted away from the rest of the Guidelines. The piece of text asks whether an isolated record permits singling out by matching that record with the individual's corresponding attributes, for example through interaction with the person or additional information. That is a sensible route to identification, but it should not be read as the only route. Read too narrowly, the paragraph could sound like singling out requires a later match back to a conventional identity or to an external set of 'real-world' attributes.

The rest of the Guidelines point in a broader direction: Para 19 defines identification as distinguishing a person from others in context so that they can be treated differently, para 25 and Example 6 expressly recognise fingerprinting and persistent pseudonyms as identifiers where they let an actor recognise and profile the same person over time, para 50 again says that a pseudonym can itself identify an individual, and para 87 focuses on whether a technique produces results accurate enough to distinguish and treat the person differently. **We would therefore keep the record-isolation/singling-out distinction, but add a cross-reference from paras 98–100 back to those provisions. Singling out must not silently turn into 'can we recover their name?'.**

6. The Guidelines should address the implications for AI more directly
There is already quite a lot about AI in the Guidelines and we welcome that, but it is spread across different parts of the technical analysis. Para 77 refers to attacks that extract data from supposedly anonymous AI models; para 79 covers membership inference; para 82 addresses prompting AI models with additional information to elicit information about a particular individual which could violate the 'no inference' criterion; paras 83–87 discuss AI-assisted re-identification and the risks of high-dimensional and high-resolution data; and paras 92–93 recognise that AI can reduce the time and cost of finding, combining and exploiting information. We think these pieces should be brought together more visibly, given this is one of the areas where arguments about anonymity are already becoming politically and economically important.

<u>6.1 Supposedly anonymised data used for AI training</u>
Whether a training dataset is considered anonymous can decide whether very large-scale collection and reuse for model development is treated as subject to the GDPR at all. That makes the initial assessment particularly sensitive for rich training data. Search queries, location histories, behavioural sequences, device data, user-generated text and other high-dimensional records can remain highly distinctive even after names and obvious account identifiers have been removed. The Guidelines already recognise the importance of dimensionality, resolution and

additional information when assessing re-identification techniques (paras 85–86), and para 29 requires the properties of the data and the processing context to be taken into account.

The final Guidelines could say more plainly what follows from this for training data. Removing direct identifiers is only the beginning of the analysis. The assessment still has to ask whether combinations of attributes distinguish individuals, whether records can be linked, what can be inferred, and what external data or AI-assisted methods are realistically available to the developer or other relevant actors (paras 23–30, 50–52 and 83–86). The same caution applies to PETs, which can potentially substantially reduce risk and should be encouraged where appropriate, but whose use does not itself answer the legal question whether the resulting data are anonymous.

This is also a matter of purpose and effect: Para 16 says that information can relate to a person because it is used or likely to be used to evaluate, treat or influence them, or because its processing creates a concrete possibility of different treatment. And Para 50 similarly treats a dataset as obviously non-anonymous where it is clearly used to evaluate, treat or influence a specific individual. If a dataset is being assembled for a system that will personalise, rank, recommend, assess, predict or target people at individual level, the identifiability analysis cannot sensibly ignore that processing function. The developer may genuinely have no interest in a person's name and still derive value from persistent individual patterns and outputs about that person.

6.2 Anonymity of the training data and anonymity of the model are separate questions

The Guidelines should also keep two questions separate. Was the training dataset anonymous? And is the resulting model itself anonymous? These are related, but they are not interchangeable. Opinion 28/2024 already says that an AI model can only be treated as anonymous where, for any data subject, the likelihood of extracting personal data about them from the model or obtaining such data through queries is insignificant, taking account of all means reasonably likely to be used.[11]

Guidelines 02/2026 seem to go in the same direction in paras 77, 79 and 82. A model may disclose information from the training data, reveal that somebody was present in a dataset, or make it possible to derive information about a person when prompted with outside data. Those are model-level questions. **Equally, a conclusion that the finished model is anonymous does not answer the separate question whether the personal-data processing used to develop it complied with the GDPR. The final text could make that separation more explicit.**

Although machine unlearning is a different question from anonymisation, recent work on vision-language models is useful for the model-level assessment. Sun et al. find that suppressing particular answers or attributes can leave identity-level knowledge available through alternative modalities, reverse associations, identity binding and relational reconstruction. A model-level anonymity assessment should therefore not stop at whether one training record or one fact can still be directly extracted; it should also consider whether information relating to the same person remains accessible, linkable or reconstructable through other routes[12].

6.3 AI changes what means are reasonably likely to be used

The Recital 26 test also has to move with technology, and we are grateful that the Guidelines are already strong on this. Para 35 notes that re-identification risk changes over time; para 88 recommends safety margins because capabilities evolve; and paras 92–93 expressly say that AI, including agentic AI, is likely to reduce the time and cost of deploying re-identification techniques and gathering information from multiple sources. The practical consequence is

11 EDPB, Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models, adopted 17 December 2024, para 43.

12 Xiongtao Sun, Hui Li, Tiantong Wu, Jiaming Zhang, Fuyao Zhang and Wen Jun Tan, 'What Does It Mean to Forget a Person? Individual-Level Unlearning in Vision-Language Models', arXiv:2609.33481 (2026).

straightforward: something that once required disproportionate time, cost or labour may become a means reasonably likely to be used when widely available tools remove those barriers.

Recent empirical work makes that shift measurable. Lermen et al. show that LLM-assisted pipelines can deanonymise pseudonymous online accounts from unstructured text at scale, reaching up to 68% recall at 90% precision and reducing work that previously required skilled manual investigation to minutes and very low per-profile costs. Wen's study reaches the same problem from a different angle: named-entity masking left substantial residual re-identification risk, and search-enabled agents could retrieve or assemble identifying information even when their final answer did not reveal the person's name. For agentic systems, the threat model therefore has to include retrieval and intermediate combination with public information, not only the final output[13].

That is also why anonymisation cannot be treated as a one-off event. If valuable datasets remain in use for years, or are repeatedly shared, changes in models, auxiliary datasets and attack techniques can change the result. Paras 35, 42 and 88 already support reassessment in different ways. Organisations relying on anonymity should document the assumptions on which that conclusion depends and revisit them when those assumptions stop being true. This is basic accountability, but it is easy to lose once somebody has put an 'anonymous' label on the data.

6.4 Profiling and automated decision-making should be part of the examples

We would add an example where an AI system never learns a person's name but uses a persistent identifier, embedding or distinctive feature vector to score, rank or personalise treatment for that person. As mentioned above, the absence of a civil identity should not end the analysis. If the system can keep distinguishing the same individual and use information relating to them to decide how they are treated, the example should test whether that already amounts to identification in the sense of Recital 26.

This would make the connection with profiling and automated decision-making much less abstract, and would additionally also help prevent the new No Record Isolation analysis being used later to argue that pseudonymous profiling or AI-based decision systems fall outside data protection law simply because the controller has no reason to know the person's name.

7. *EDPS v SRB* should not become a general GDPR exit route

We undestand the Guidelines have to engage with the Court's judgment in the *SRB* case. Overall, we think they handle it more carefully than some of the political debate around the judgment. Paras 6–8 recognise that identifiability can depend on the relevant entity's perspective, but the rest of the legal analysis immediately places limits around that idea. Para 14 and Example 2 show that the applicable perspective can depend on the GDPR obligation at issue; para 15 prevents a processor from escaping the GDPR simply because it cannot identify the data subject on its own; and paras 20–21 make clear that means and information can be spread across different entities, including through foreseeable transfers.

The facts of *SRB* also matter, and their highly case-specific nature sometimes disappear from the current debate. Deloitte received comments that had been filtered, categorised and aggregated; duplicates were merged, and the alphanumeric code attached to a comment did not function like an ordinary stable person-level pseudonym that let Deloitte follow the same participant across a dataset (*EDPS v SRB*, paras 27–28). Only the SRB could use the code to link comments back to identification data. That is a rather particular factual setting. It is a poor basis for broad claims about behavioural profiles, stable device identifiers or rich pseudonymised datasets where the whole point is to keep records about the same person linkable.[14]

13 Simon Lermen, Daniel Paleka, Joshua Swanson, Michael Aerni, Nicholas Carlini and Florian Tramèr, 'Large-scale online deanonymization with LLMs', arXiv:2602.16800 (2026); Jianing Wen and Tianshi Li, 'AgentDoxx: Agentic Re-identification of Anonymized Text with Web Search', arXiv:2610.05586 (2026).

This is critical because the Commission has relied on the recent case law in the Digital Omnibus proposal to justify changes to Article 4(1) GDPR. The Proposal states that information is not personal for a given entity where that entity cannot identify the person using means reasonably likely to be used by that entity, and that the information does not become personal merely because a potential subsequent recipient has such means. Proposed Article 41a would also allow Commission implementing acts to specify means and criteria for determining when data resulting from pseudonymisation no longer constitute personal data for certain entities. The EDPB and EDPS have themselves warned that these proposals go beyond the CJEU case law and could de facto redefine the material scope of EU data protection law.[15]

Our ask to the EDPB is therefore fairly specific. **The safeguards which are already in Guidelines 02/2026 should be made difficult to detach from the entity-specific language in paras 6–8.** Contextual identifiability is not a purely subjective exercise based on what one recipient happens to hold today or says it intends to do. The Guidelines themselves require attention to expected uses, access paths, transfers, other entities, chains of means, available outside information and realistic changes in capability (paras 9, 13, 20–21, 28–35 and 88–93). We think **some additional cross-referencing would help, particularly because SRB is already being relied on for claims that go much further than the judgment.**

Article 11 GDPR is worth mentioning here as well. The Regulation already deals with situations in which a controller does not need to identify a data subject for the purposes of its processing, and adjusts the operation of some data-subject rights where the controller can demonstrate that it is not in a position to identify the person. It does not follow that the information is automatically anonymous or outside the GDPR. Online behavioural advertising is an example where the controller can distinguish between different persons (and processes personal data) without being in a position to identify a specific data subject[16]. The final Guidelines should keep that distinction visible, especially when discussing an entity's inability to attribute data to a conventional identity.

8. The wider EU context makes precision especially important

8.1 Digital Omnibus

The 'Digital Omnibus on Data' makes this discussion more than theoretical. The European Commission proposes to amend the definition of personal data so that the actor-specific dimension is written directly into Article 4(1), including the proposition that information does not become personal for one entity merely because a potential later recipient has means to identify the person. The explanatory material presents this as a clarification drawn from recent CJEU case law, including *SRB*.[17]

As mentioned, we do not believe *SRB* supports that broad legislative move,. The judgment is fact-specific, and the rest of the case law does not justify simply disregarding recipients, anticipated uses or other reasonably available means. More importantly, writing this into the definition risks turning a contextual legal assessment into something much closer to a scope exemption that actors apply to themselves. The consequences for rights and enforcement are obvious.

---

14 Case C-413/23 P, European Data Protection Supervisor v Single Resolution Board, EU:C:2025:645, judgment of 4 September 2025, especially paras 27–28, 74–86 and 102–112; see also F Bieker, 'Schroedinger's data: SRB and the Digital Omnibus' (European Law Blog, 20 January 2026) doi:10.21428/9885764c.25fc449c accessed 28 October 2026

15 EDPB and European Data Protection Supervisor, Joint Opinion 2/2026 on the Proposal for a Regulation as regards the simplification of the digital legislative framework (Digital Omnibus), adopted 11 February 2026, in particular paras 17–25

16 Helena Malikova and Johnny Ryan, 'Re-identification-Based Advertising', TechREG Chronicle, July 2026 https://www.pymnts.com/cpi-posts/re-identification-based-advertising/

17 European Commission, Staff Working Document accompanying the Proposal for a Regulation as regards the simplification of the digital legislative framework ("Digital Omnibus"), 19 November 2025, section 1.2.2.1, 'The definition of personal data'.

The Guidelines do not need to become a commentary on the Omnibus negotiations, but should avoid leaving openings that can easily be quoted as support for this actor-relative reading. That means being firm on the objective elements of Recital 26, on other actors and chains of means, on the unusual facts of *SRB*, and on the difference between genuine anonymity and data that are simply pseudonymised or kept behind access controls.

8.2 Article 6(11) DMA and Google Search data

The Article 6(11) DMA proceedings on Google Search are another useful stress test. On 16 July 2026, the Commission adopted specification measures requiring Google to share anonymised ranking, query, click and view data with eligible third-party search engines, including AI chatbots with search functionality[18]. The measures combine technical transformation with organisational and contractual restrictions, independent audits and periodic review.

According to the Decision, Google has to remove direct identifiers and other identifying attributes, suppress very long queries and queries containing rare terms, generalise metadata and apply a high k-anonymity threshold. The contractual layer then restricts linking with other datasets, onward disclosure, re-identification and retention. The Commission also bars use of the shared data for training general-purpose AI models and for unrelated consumer profiling or advertising. And the framework can be revisited as new attacks or technical capabilities emerge.

This is a good illustration of why the general meaning of anonymisation needs to be clear. Search data are difficult because they are multi-dimensional, often contain free text and are valuable precisely because they are specific. Those same characteristics make linkage and inference easier. The DMA arrangement also shows that contracts and access controls can have a useful supporting role in a controlled environment, but only alongside actual transformation of the data and independent verification.

We would therefore **add a search-data example, or at least explain how the horizontal criteria apply to this type of setting.** Otherwise we risk slowly developing one practical meaning of anonymisation for DMA data sharing and another one under the GDPR. That would be a very strange outcome for concepts that are supposed to remain coherent across EU law.

9. Accountability, burden of demonstration and enforcement

Calling data anonymous has unusually large legal effects: later processing of genuinely anonymous information sits outside the GDPR. The actor relying on that conclusion should therefore be able to demonstrate it. The Guidelines already move in this direction. Para 40 warns against misleading labels such as 'anonymous' or 'de-identified' where people remain identifiable, para 41 requires documentation of both the process and testing, and paras 95–103 require further analysis where the technical criteria are not all satisfied.

We would go a little further and make the expected content of a credible assessment explicit. Depending on the context, it should cover:

- the relevant actors and realistic threat models, and the additional information or means available to them;
- the structure, dimensionality, resolution and diversity of the dataset;
- the tests used for record isolation, linkage and inference, and what happens where a criterion fails;
- technical, organisational and contractual controls, including how reliable and enforceable they are;
- foreseeable security failures, leaks and other unauthorised access paths;
- expected uses of the data and the assumptions being made about present and foreseeable technology

18 European Commission, Decision pursuant to Article 8 DMA, Case DMA.100209 – SP – Alphabet – Article 6(11), 16 July 2026.

- the expected period for which the data will remain available or be processed, and the extent to which dissemination may make that period effectively open-ended.

9. Conclusion
Overall, we believe Guidelines 02/2026 are a good basis for the EDPB's final position. They retain much of the the 2014 Opinion approach and, in several places, make it more realistic by looking at processing context, transfers, chains of means, unauthorised actors, security incidents and technological change. The legal analysis in paras 14–15 and 19–35 in particular contains useful safeguards against an excessively narrow reading of identifiability.

The main issue now is to make sure that some of the new qualifications are not detached from those safeguards when the Guidelines are used in practice. The record-isolation/singling-out distinction is the clearest example. It is useful to avoid saying that every unique record is automatically identifying. But paras 98–100 should be read together with the Guidelines' own broader definition of identification: persistent recognition and differential treatment can matter without a civil name. The same caution applies to entity-specific identifiability. It should remain a contextual, evidence-based assessment, not slide into an actor-by-actor opt-out from the GDPR.

We would also like the final text to draw the line between anonymisation and pseudonymisation more explicitly. The current Guidelines already contain the ingredients for that distinction, and the consultation version of Guidelines 01/2025 is clear that prevention of attribution does not itself imply anonymity. Once the final pseudonymisation Guidelines are adopted, the two texts should cross-reference each other and use compatible language. This is especially important now, when pseudonymisation is being discussed politically as a possible route to data falling outside the GDPR. AI makes all of these questions more urgent as well: training datasets, model extraction, membership inference, automated linkage and individual-level profiling all put pressure on old assumptions about what it means to identify somebody. The final Guidelines should address those implications directly rather than leave the reader to reconstruct them from examples scattered across the document.