CYBERNETICA

Feedback on EDPB Guidelines 02/2026 on Anonymisation

Public consultation response

30.10.2026

Public

©2026 Cybernetica AS

1 Introductory remarks

Cybernetica is an Estonian research and technology company that designs and builds privacy-preserving data processing systems. Our work includes the Sharemind secure computation product line[1] and the JOCONDE project (Joint On-demand COmputation with No Data Exchange)[2] commissioned by Eurostat, where Cybernetica helped design a multi-party secure private computing infrastructure enabling members of the European Statistical System to compute aggregate statistics across confidential micro-data held in different organisations and Member States without those organisations exchanging their data. As part of the JOCONDE project, Cybernetica also carried out the legal and data-protection-impact analysis that must accompany such architectures under EU data protection law.

We welcome the draft Guidelines and provide our feedback from the perspective of an engineering organisation who is expected to translate the EDPB's legal standard into system specifications and whose customers need to justify substantial investment in privacy-enhancing technologies on the basis of that standard.

Our comments are driven by three basic concerns:

1. the current level of data protection in the EU should not be weakened;
2. the current risk-based approach to identifiability/anonymity should be maintained (e.g. a simple replacement of names and ID numbers in a database should not be considered as effective an anonymisation technique as advanced encryption solutions and privacy enhancing technologies like multi-party computation, trusted execution environments, use of synthetic data etc.);
3. the legal interpretation of the term "anonymous information" in the meaning of Recital 26 GDPR should not include anonymisation methods which have been proven to be ineffective or easily circumvented in scientific research.

[1] Cybernetica's Sharemind MPC. – Internet: https://cyber.ee/products/sharemind-mpc/

[2] JOCONDE project (Joint On-demand COmputation with No Data Exchange). – Internet: https://cros.ec.europa.eu/book-page/joconde

2 Feedback on EDPB Guidelines 02/2026

2.1 General remarks

We note that any forthcoming changes to EU policy — including those reflected in the guidelines, in related regulatory initiatives or court practice — should be carefully calibrated so as not to unduly weaken the overall level of data protection afforded to individuals. While we recognise the need to balance innovation, administrative burden, and legal certainty, this should not come at the expense of the fundamental protections that the GDPR and related data protection framework are intended to guarantee.

2.1.1 Contradictions regarding relative assessment

Per Guidelines, the general principle is that whether given data constitutes personal data or anonymous data must be assessed from the perspective of the entity carrying out the processing.

Paragraph 15 of the Guidelines departs from this principle in the specific case of an entity processing data on behalf of a controller for whom that data is personal: in that scenario, whether the data is personal for the processing entity is to be determined by reference to the controlling entity's perspective, with no independent test applied. The practical effect of this approach is that it becomes impossible to process anonymous data on behalf of a controller for whom the underlying data is personal. The entity's role is thereby automatically deemed that of a processor within the meaning of the GDPR, foreclosing the possibility that the entity is in fact processing anonymous data and is therefore not a processor under the GDPR at all.

This approach is, in our view, inconsistent with the general rule and introduces an exception that cannot be reconciled either with the underlying rationale of that rule or with the GDPR itself. We therefore recommend that paragraph 15 of the Guidelines shall be revised and clarified.

2.1.2 Risk management as the central concern

Whether data is anonymous is a legal threshold. The actual protection of individuals, however, depends on how effectively the risk of identification is identified, mitigated and monitored over time. The Guidelines themselves recognise that anonymity may be lost as technology evolves (para 35) and recommend safety margins (para 88). We suggest that the Guidelines state more explicitly that a correctly executed risk management process, with appropriate safeguards, is the core obligation. The anonymous/personal classification should be presented as one outcome of that process, not as a substitute for it.

2.1.3 Input and output privacy

The framework in Section 3 assesses the "given data", which in effect is the output of the anonymisation process. Statistical disclosure control applied at the publication stage (output privacy) is one established route to anonymity, but it is not the only one. Modern privacy-enhancing technologies (hereinafter "PETs") also make it possible to protect the input of the processing, as well as data in use and any intermediate results during computation (input privacy). Depending on the technique, this ensures that no single computing party has access to intelligible personal data (e.g. secure multi-party computation or homomorphic encryption, under defined security

assumptions). Alternatively, it ensures that data processed in clear by a machine remains inaccessible to any human, including the operator of the system (e.g. trusted execution environments).[3]

We recommend that the Guidelines explicitly distinguish between input, throughput and output privacy. The assessment should cover each relevant party at each stage of the process: the input data, the data in use during computation, and the resulting output. The fact that the data under processing is intelligible to an automated system for a limited time, but not to any human at any time during the process, can reduce the identification risk significantly. It would be particularly helpful to have the Guidelines address the issue whether, and under which conditions, such processing must be conducted in order to qualify as anonymisation or "rendering anonymous" in the meaning of Recital 26 GDPR.

We also invite the EDPB to clarify how such input privacy measures interact with para 15. That paragraph treats a processor's perspective as identical to the controller's, even where the processor technically cannot access intelligible data, for example when computing on encrypted or secret-shared data in use. Para 9 recognises that actual access to the data matters, as it does not require an assessment for entities that cannot access or process the data. However, it does not address parties that process data only in a form they cannot read. Further guidance on how such configurations should be assessed, and how they affect the identification risk, would be valuable.

2.1.4 Anonymisation techniques, PETs and encryption

The test of "means reasonably likely to be used" is dynamic. As the Guidelines recognise (para 35), the likelihood of re-identification tends to increase over time as techniques advance and more additional information becomes available. Irreversibility can therefore rarely be guaranteed indefinitely. For this reason, appropriate protection should not rely only on traditional anonymisation techniques, but should also make use of innovative solutions, in particular PETs.

However, the Guidelines do not refer to PETs as such, nor to specific techniques such as secure multi-party computation, homomorphic encryption, trusted execution environments, differential privacy or federated learning. Where related techniques appear, they are treated as potential vulnerabilities rather than safeguards (paras 64, 77 and 82). To encourage wider uptake of PETs, the EDPB's view on when their use may result in anonymous data, or reduce the identification risk to an insignificant level, would be very helpful. At the same time, correct implementation is essential, since misapplied PETs may fail to achieve their purpose and create new risks. We therefore recommend technology-neutral criteria for correct implementation. Such criteria would cover the threat model, security assumptions, parameter choices, key management and resilience to foreseeable advances such as quantum cryptanalysis.

The examples in the Guidelines mainly illustrate simple techniques, such as removing direct identifiers, generalisation and aggregation, and largely show how these can fail. Examples of established disclosure control techniques and PETs applied successfully, and mapped onto the three

[3] See, e.g. UN Committee of Experts on Big Data and Data Science for Official Statistics, *UN Handbook on Privacy-Preserving Computation Techniques*, pp. 14–15. – Internet: https://unstats.un.org/bigdata/task-teams/privacy/UN%20Handbook%20for%20Privacy-Preserving%20Techniques.pdf;
JOCONDE project, Deliverable D1.2 *Usage Scenarios and System Requirements – Final version*, section 1.3;
Ballhausen, H. *et al*. *Privacy-friendly evaluation of patient data with secure multiparty computation in a European pilot study.* – Internet: https://www.nature.com/articles/s41746-024-01293-4;
Bogdanov, D., Kamm, L., Laur, S. and Sokk, V., *Rmind: A Tool for Cryptographically Secure Statistical Analysis* in IEEE Transactions on Dependable and Secure Computing, vol. 15, no. 3, pp. 481-495, 1 May-June 2018, doi: 10.1109/TDSC.2016.2587623. – Internet: https://ieeexplore.ieee.org/document/7505613

criteria, would help controllers assess whether data can be considered anonymous after such measures are implemented.

Encryption also deserves a more differentiated treatment. Para 91 does not distinguish between symmetric and asymmetric encryption, although they differ significantly in their security properties, key management requirements and resilience to attack, including quantum attacks. Nor does it address where keys are stored, who has access to them, and whether and how they are deleted. These factors are central to assessing whether decryption is a means reasonably likely to be used.

Finally, not all pseudonymisation is equal. Replacing national identification numbers with arbitrary codes differs substantially from strong pseudonymisation combined with PETs in the effort required for re-identification. The latter may, in practice, offer stronger protection than anonymisation that still carries a non-negligible residual risk. We encourage the EDPB to acknowledge a graduated assessment based on the complexity of re-identification, consistent with the relative approach and Article 11 GDPR. Such an acknowledgement would encourage controllers to invest in stronger safeguards rather than settling for minimal measures that formally achieve the same legal status.

2.1.5 Techniques known to be easily reversible

Measures such as removing the name column, hashing low-entropy identifiers, or simple generalisation of high-dimensional data have repeatedly been shown in the scientific literature to be insufficient. The Guidelines themselves cite such research (para 64). We suggest stating clearly that a technique for which practical re-identification attacks are publicly documented cannot, by itself, produce anonymous data within the meaning of Recital 26 GDPR. An indicative, regularly updated annex of such techniques would help practitioners.

2.1.6 A methodology for assessing anonymity

Although the three criteria proposed in the Guidelines are useful, controllers would benefit more from guidance on how to measure and document the likelihood of identification, what level of residual risk may be regarded as insignificant, and how to structure the process so that the identifiability of data is minimised at every step. Currently, controllers are often reluctant to conclude that data is anonymous, because there is no recognised basis on which to rely. A clearer acknowledgement that "rendering anonymous" in the meaning of Recital 26 GDPR is a multi-step process, combining technical measures with documented risk analysis, would provide much-needed legal certainty.

2.1.7 Assessing the process rather than every perspective

The contextual approach requires each relevant entity’s capabilities to be assessed individually. The Guidelines acknowledge that this can be complex and prone to error (para 47). For PET-based processing, much of the relevant information is determined by the design of the process itself: what each party receives, what can be computed from it, and which steps are reversible. We suggest that the Guidelines allow the assessment to be anchored in the properties of the anonymisation process. Analysing the process as a whole may often be more reliable and practicable than enumerating perspectives one by one.

3 Conclusion and recommendations

We thank the EDPB for its valuable work on these Guidelines, which provide a solid legal foundation for the assessment of anonymity. We believe that their practical value could be further enhanced by a more process- and risk-oriented approach. In particular, we would welcome further clarification of the approach set out in para 15. We would also welcome more concrete and technically granular guidance to complement the three criteria, which would enable controllers to apply the Guidelines with greater legal and technical certainty.

To this end, we respectfully invite the EDPB to consider:

1. emphasise risk management and appropriate safeguards as the core obligation, with the anonymous/personal distinction as one outcome of that process;
2. distinguish input and output privacy, and clarify the position of parties that cannot access intelligible data (paras 9 and 15);
3. explicitly recognise correctly implemented PETs, including a broader view of cryptographic techniques, as means of mitigating identification risk, together with criteria for correct implementation, in a technology-neutral manner;
4. acknowledge differences in the strength of pseudonymisation according to the complexity of re-identification;
5. state that techniques with publicly documented re-identification attacks cannot by themselves achieve anonymisation;
6. provide a structured methodology, with illustrative PET-based examples, for assessing and documenting residual identification risk; and
7. allow the assessment to be grounded in the properties of the anonymisation process as a whole.