Finnish Information Security Cluster
Kyberala

30.10.2026

Status: Preliminary position

Guidelines 02/2026 on Anonymisation – EDPB feedback by 30.10.2026

Subject: **EDPB Guidelines 02/2026 on Anonymisation** (version 1.0, adopted for public consultation on 7 July 2026) – preliminary comments

Finnish Information Security Cluster (FISC) is a branch association of the Technology Industries of Finland (TIF) representing the cybersecurity industry operating in Finland. We present our preliminary comments on the European Data Protection Board's (EDPB) draft guidelines on anonymisation. Our focus is on whether the guidelines are fit for purpose for cybersecurity industry and, more broadly, for technology industries.

Summary
The guidelines are a clear improvement on WP29 Opinion 05/2014. They take recent case law of the Court of Justice of the European Union (CJEU) as their starting point (in particular C-413/23 P EDPS v SRB), recognise that anonymity is entity-specific, i.e. relative, and offer a structured three-criteria test. However, the guidelines are written mainly from the perspective of research, health and consumer data. As they stand, they are unfortunately rather poorly suited to the types of data typical of the cybersecurity sector, such as logs, telemetry, threat intelligence and AI models.

Several positions in practice tilt the test towards zero risk, which does not reflect the CJEU's requirement that the likelihood of identification be insignificant. These include the perspective applied to processors, the "cannot be ruled out" threshold for onward transfers and leaving the required level of security undefined.

The importance of the guidelines is underlined by the state of play in the Digital Omnibus negotiations. It is our understanding that the Council's compromise text has removed the Commission's proposed clarification of the definition of personal data (Article 4(1) GDPR) and has left the matter to EDPB guidance. Even if the final outcome of the Omnibus were different, the guidelines are in any event set to become the de facto key instrument determining the status of pseudonymised and anonymised data in the EU data economy. They must therefore faithfully reflect the case law of the CJEU.

Positive aspects

- The relative approach based on CJEU case law (C-413/23 P EDPS v SRB): anonymity is assessed from the perspective of each entity (paras 6–7).
- The risk need not be zero but insignificant (para 22), and failing a criterion does not automatically make data personal data (para 52).
- Assessments carried out in accordance with WP29 Opinion 05/2014 do not need to be repeated (footnote 3).
- Third parties cannot use the simplified approach to prove that data is personal data (footnote 36).
- Inferences based on general knowledge are excluded from the No Inference criterion (Examples 13 and 14), which is important for analytics and AI.
- The role of security controls is recognised (paras 89–91), and immediate anonymisation is a positive factor in the legitimate interest balancing test (para 39).

Teknologiateollisuus ry
Kyberala ry

Eteläranta 10, P.O. Box 10, FI-00131 Helsinki
Tel. +358 9 192 31
www.teknologiateollisuus.fi
Business ID: 0215289-2

Key recommendations

- Add examples of cybersecurity and technology-sector data (logs, threat intelligence, telemetry, AI models, synthetic data), including examples of successful anonymisation.
- Clarify the position of the processor where a service provider derives anonymous threat intelligence from customer data (para 15).
- Align the onward-transfer threshold with the CJEU's "insignificant likelihood" standard (para 21).
- Establish a presumption that security in line with the controls that are decisive under the current threat landscape, together with recognised encryption, excludes unauthorised actors from the relevant entities.
- Confirm the contextual approach as the main rule and give concrete content to safety margins.
- Refer to ISO/IEC 20889 and ISO/IEC 27559 and develop a lighter assessment path and a documentation template for SMEs.
- Align reassessment following security incidents, and the related notification obligations, with the single-entry-point principle.

1. Fitness for purpose for cybersecurity and technology companies

In the course of protecting their customers, cybersecurity companies process large volumes of technical data: IP addresses, logs, network traffic metadata, endpoint telemetry, malware samples, phishing messages and indicators of compromise (IOCs). These datasets are record-level, high-dimensional and high-resolution. According to the guidelines, this is precisely the type of data that is most difficult to anonymise (para 86). At the same time, other EU legislation encourages the sharing of threat intelligence: Article 29 of the NIS2 Directive provides a framework for voluntary cybersecurity information-sharing, and the revision of the Cybersecurity Act (CSA2) and the Cyber Solidarity Act pursue the same aim. The information shared is often aggregated or anonymised.

The guidelines contain not a single example drawn from cybersecurity, the industrial internet, cloud services or software development. Device identifiers and browser fingerprints are mentioned only as means of identification (para 25 and Example 6). The status of industrial machine and device data (which is known to have personal data dimensions) is left entirely open, even though Technology Industries of Finland, among others, has long highlighted the resulting legal uncertainty.

From the cybersecurity sector's perspective, the guidelines are double-edged. On the positive side, they recognise the importance of security for anonymity: appropriate confidentiality controls limit the entities that have access to the data (paras 89–91). On the other hand, the guidelines make anonymity dependent on the level of security without providing any yardstick for it (Examples 8 and 22). We consider it possible that a security incident could thereby retroactively turn anonymous data into personal data and trigger the notification obligations under Articles 33 and 34 GDPR (para 42).

Recommendation: *The guidelines should include examples from the technology sector: (a) processing Security Operations Center (SOC) or Managed Detection and Response (MDR) service logs into aggregated threat intelligence, (b) sharing IOCs within an information-sharing community under Article 29 NIS2, (c) telemetry from industrial equipment and predictive maintenance, (d) AI model training, and (e) synthetic data. The examples should also show successful anonymisation, not only failures.*

2. The processor's perspective and the service provider's own purposes (para 15)
According to the guidelines, anonymity is assessed for the processor from the controller's perspective, so the data is personal data for the processor as well. The starting point is understandable, as the processor construct prevents the GDPR from being circumvented by outsourcing processing. However, cybersecurity services (MSSP, SOC, EDR, cloud services) typically act as processors and derive aggregated threat intelligence, detection rules and models from customer data in order to protect all of their customers. The guidelines do not make clear whether the service provider may in such cases act as an independent controller for the anonymisation, from whose perspective the result is assessed, and what the legal basis for the processing is. The CJEU did not decide the processor scenario in SRB; the guidelines extend the interpretation on teleological grounds.

Recommendation: *The guidelines should clarify (i) that a processor may anonymise data in accordance with the controller's instructions and that the anonymity of the result is assessed from the perspective of the entities for which the data is intended, and (ii) how the compatibility assessment under Article 6(4) and Recital 49 (network and information security as a legitimate interest) apply when a service provider produces threat intelligence for its own purposes.*

3. Onward transfers and the "cannot be ruled out" threshold (paras 21 and 28)
According to para 21, data is personal data both for the transfer and for the recipient's processing, and indirectly also for the transferring entity, if it is reasonably likely that the data will be transferred to a third party and it cannot be ruled out that the recipient has means reasonably likely to be used for identification. "Cannot be ruled out" is a clearly lower threshold than the insignificant likelihood required by the CJEU (C-413/23 P, para 82), and it conflates two different likelihood assessments.

In threat-intelligence sharing, the group of recipients is open-ended (information-sharing communities, CERTs, authorities, customers). The "cannot be ruled out" threshold would therefore in practice make almost all shared technical information personal data for every party in the chain. Example 7, according to which a recipient can always hire a third party to carry out re-identification, reinforces this impression.

Recommendation: *The threshold in para 21 should be aligned with the CJEU's case law: the recipient's means must be reasonably likely to be used, not merely possible. The guidelines should also recognise the significance of a combination of distribution restrictions (e.g. TLP classification), terms of use and technical controls.*

4. Threat assessment, the role of security and encryption (paras 30–34, 42, 88–93)
We share the EDPB's threat assessment. The cybersecurity sector's own situational synthesis confirms that AI agents have caused the unit cost of data breaches and re-identification to collapse. Reconnaissance, tool development and the processing of stolen data take place at machine speed. A single individual can now combine leaked datasets and build a people-search service out of them – something that previously required an organisation. Capability no longer distinguishes state actors, criminal groups and individual activists from one another. We therefore consider it correct that the assessment cannot be based on a lack of motivation (para 31) and that cybercriminals and foreign intelligence services are included among the relevant entities (para 30). It should be emphasised that threat actors are traditionally assessed on the basis of motivation × capability; when AI agents can provide a capability clearly greater than zero, very few threat actors can be considered insignificant compared with the past.

We do, however, disagree with the conclusion. When attackers' capabilities grow, the answer cannot be to treat almost all data as personal data. That would remove the incentive to anonymise, and companies would process personal data where anonymous data would suffice. The protection of data subjects would weaken rather than strengthen. The sustainable answer is to restrict access to the data and to the additional information that enables identification. The guidelines recognise this (para 89) but leave the adequate level of cybersecurity open. The "out-of-date" security in Example 8 is not a yardstick: it would allow a supervisory authority to deem any system inadequate after the fact, and a company cannot assess the risk in advance.

The yardstick should therefore be tied to the controls that are decisive given the threat landscape. Most breaches start from stolen credentials, not software vulnerabilities. Data is exfiltrated in bulk, and through the supply chain – for example via a SaaS provider's application authorisations – attackers gain access to the data of hundreds of organisations at once. Where anonymity relies on the data or additional information remaining confidential, protection should cover at least phishing-resistant multi-factor authentication (MFA), management of secrets and service accounts, segmentation and detection of anomalous bulk exports, management of supplier access rights, and protection of logs and backups. These are already part of the requirements of Article 21 NIS2 and ISO/IEC 27001, so no new obligation is created.

As regards encryption, the guidelines require an assessment of whether encryption could be broken, for example through cryptanalysis (para 91). This is technically justified, but without a presumption no encryption could ever be considered sufficient. The quantum threat ("harvest now, decrypt later") should be managed within the timeline of the EU's transition to post-quantum cryptography. It should not be used as grounds for deeming current appropriate encryption insufficient.

Because a breach can progress from initial unauthorised access to data theft within hours, anonymity that relies solely on confidentiality is weak. Anonymity that exists only because the key enabling identification is under lock and key is one breach away from harm. Under the EDPB's contextual assessment, data may be anonymous for a recipient even though someone else holds information that enables individuals to be identified. Such information may be, for example, a pseudonymisation key, the original dataset or an internal report. The condition is that this additional information remains confidential and the recipient has no access to it (para 89, Examples 8, 21 and 22). In such a case, anonymity is therefore not based on identification being technically impossible, but on no outsider being able to access the key. If the systems of the organisation holding the key are breached, the attacker may obtain both the anonymised dataset and the key and combine them. According to para 42, the data is then personal data again. Detecting and responding to an intrusion typically takes days or weeks, whereas exfiltrating the data may take only a few hours. One therefore cannot rely on the breach being stopped before the key leaks.

Moreover, the relationship between reassessment following a security incident (para 42) and notification obligations is unclear. The same incident may lead to notifications under, among others, both NIS2 and the GDPR, on different timelines and with different content.

Recommendation: *(i) The guidelines should establish a presumption: where confidentiality protection has been implemented in line with recognised requirements and covers the controls listed above, unauthorised actors are not considered relevant entities unless there is concrete evidence to the contrary (cf. para 33). (ii) For recognised, state-of-the-art cryptographic mechanisms combined with appropriate key management, it should be presumed that breaking the encryption is not a means reasonably likely to be used. (iii) The guidelines should encourage keeping additional information that enables identification, such*

as pseudonymisation keys, separate from the anonymised dataset and destroying it once it is no longer needed. (iv) The notification obligations under para 42 should be aligned with the single-entry-point principle ("report once, use many times").

5. The simplified approach, safety margins and the risk-based approach (paras 36, 47–49, 88, 97)

The simplified approach asks whether re-identification is possible at all (para 49). High-dimensional technical data will in practice never pass this test: the more attributes are recorded for each individual or event, the more likely it is that the combination of attributes is unique (para 56). According to the guidelines, the approach is voluntary and is not an alternative legal standard (para 48). However, since the guidelines describe the contextual assessment as burdensome and error-prone, there is a risk that the simplified approach will become the de facto supervisory standard. Footnote 36, according to which third parties cannot use the simplified approach to prove that data is personal data, is important and welcome – it is fair to ask why this point is made only in a footnote.

In addition, the "adequate safety margins" recommended in para 88 are not defined. According to para 36, the entire dataset is considered personal data if even one record is not anonymous. In the table in para 97, a failure of the No Inference criterion leads directly to the conclusion that the data is not anonymous, unlike for the other two criteria, and the difference is not explained. Taken together, these positions shift the emphasis from a risk-based assessment towards zero risk.

Recommendation: *The contextual approach should be confirmed as the legal main rule and the simplified approach as a voluntary "safe harbour". The principle in footnote 36 should also apply to supervisory authorities, and moving it from the footnote into the main text should be considered. Safety margins should be given a concrete yardstick, and the different treatment of the No Inference criterion should be justified or aligned with the other criteria.*

6. AI models, synthetic data and inference (paras 66–82)

Conceptually, the "No Inference" criterion is the strongest part of the guidelines. Excluding inferences based on general knowledge or the population at large from the criterion (Examples 13 and 14) is essential for analytics and AI. The comparison between datasets D and D' (para 71(b) and footnote 43) reflects the core idea of differential privacy. However, the guidelines do not indicate how large a difference between prior and posterior belief is "significant" (footnote 44). For AI models and synthetic data, membership inference and the extraction of data from models are mentioned (paras 77 and 82), but without thresholds developers cannot know when a model is anonymous.

The development of European AI models, including large language models, requires predictable rules on the status of training data and models. An unclear anonymity standard pushes development outside the EU and weakens technological sovereignty.

Recommendation: *The guidelines should include examples of successful privacy-enhancing techniques (differential privacy with its parameters, synthetic data, federated learning) and clarify their relationship with EDPB Opinion 28/2024 and the requirements of the AI Act.*

7. Standards, feasibility and legal certainty over time

The guidelines make no reference whatsoever to international standards, although an established framework exists: ISO/IEC 20889:2018 (de-identification techniques) and ISO/IEC 27559:2022 (de-identification framework). Nor do they refer to ENISA's work on

pseudonymisation and privacy-enhancing technologies. The technical assessment requires specialist expertise (e.g. robust record-linkage attacks and membership inference) that is not realistically within the reach of SMEs. This is at the same time an opportunity for providers of privacy-enhancing technologies (PETs) and security testing services – but only if the assessment criteria are clear enough for the service to be productised.

According to the guidelines, the likelihood of re-identification increases over time, and the controller is accountable for the processing if the data becomes personal data (para 35). We welcome the fact that assessments carried out in accordance with the WP29 Opinion do not need to be repeated (footnote 3). Liability must, however, be based on what the entity knew or ought to have known. A properly documented assessment should protect the entity against retroactive liability.

Recommendation: *The guidelines should refer to ISO/IEC 20889 and ISO/IEC 27559 and state that following them supports the demonstration of anonymity. In addition, a documentation template (para 41), a lighter assessment path for SMEs, and the use of codes of conduct (Article 40) and certification (Article 42) to demonstrate anonymisation are needed. At national level, the Office of the Data Protection Ombudsman should provide practical guidance before resorting to sanctions.*

8. Comments on the questions raised with the national Data Ombudsman

- **When a full technical assessment is required:** Simple cases (paras 50–51) are described only in the direction of obvious personal data. A corresponding lighter path is needed for obviously anonymous data, such as highly aggregated statistics or purely machine-to-machine process data.
- **Practical application:** Large companies with data protection and data analytics expertise can apply the guidelines in part. SMEs cannot do so without templates, tools and a standards-based procedure.
- **Unclear concepts:** The following are unclear: "relevant entity", "applicable perspective" and "receiving entity" in relation to one another; "cannot be ruled out"; "treated differently" in the context of technical security measures; "adequate safety margins"; "general knowledge / population at large"; and "record" in the context of technical data.
- **Unclear examples:** Example 8: what constitutes "out-of-date" security, and on what basis is it assessed? Example 22: the conclusion that a leak is possible but that persons close to the data subjects would not obtain the data is speculative and in tension with para 30. Example 7: when is hiring a third party not reasonably likely? Example 24 (cinema survey) is a good positive example, and more of these are needed. Editorially, the table in para 97 contains a reference to "paras. 98 to 00", and the references in footnote 44 to footnotes 40 and 41 appear to be intended as footnotes 42 and 43.

Further information:

Peter Sund, CEO, Finnish Information Security Cluster (FISC), +358 50 565 0621, peter.sund@teknologiateollisuus.fi